cisco Meraki

Campus Deployment Guide

JUNE 2016

This guide provides information and guidance to help network administrators deploy Meraki Access and Distribution Swithing in a Campus environment.

Table of Contents

Purpose

Cisco Meraki switches combine powerful enterprise features with intuitive centralized management via the cloud. The Meraki cloud provides a seamless management experience for networks of all sizes, coupled with deep network visibility and control. Meraki switches can be set up for deployment to a complete site in a matter of minutes without touching the hardware, and managed for the life of the deployment, all via an intuitive browser-based user interface called Dashboard.

This guide provides information and guidance to help network administrators deploy the Meraki Switch (MS) line in a Campus environment.

Introduction

Campus networks typically adopt a tiered design, scaled according to the specific needs of the individual campus. These larger networks generally comprise WAN access, a core, an aggregation/ distribution layer and an access/edge. This blueprint is used over and over again as it's proven to be scalable and fit the majority of use cases. An example of this template/blueprint can be found below.

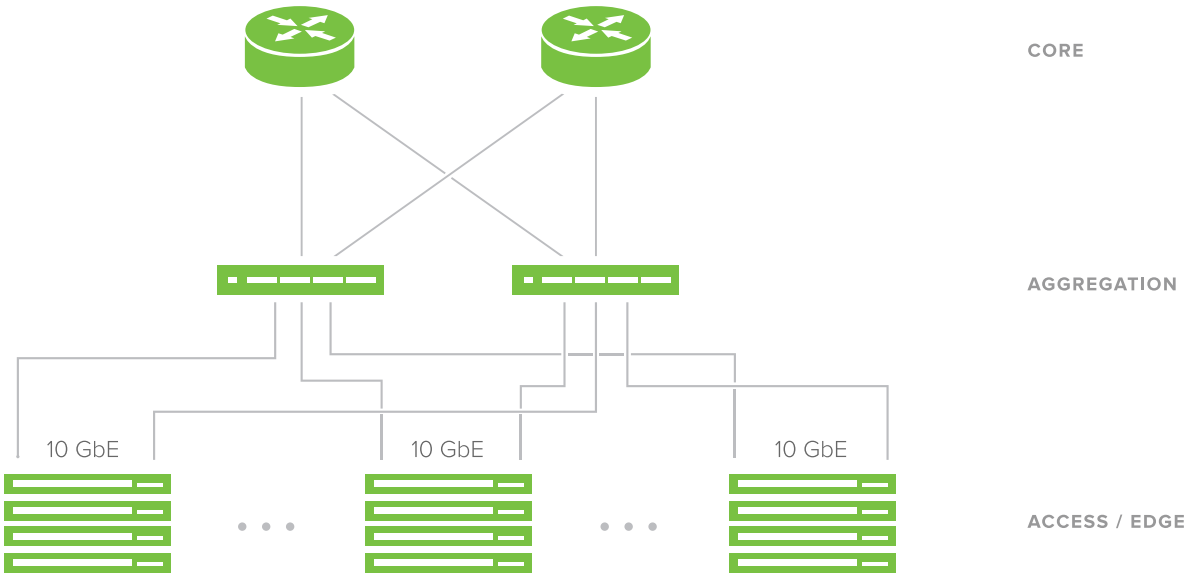

While the underlying blueprint remains the same, the devices used ultimately dictate the ease of implementation and insight into the network, characteristics which are cornerstones of the Meraki platform.

Cisco Systems, Inc. | 500 Terry A. Francois Blvd, San Francisco, CA 94158 | (415) 432-1000 | sales@meraki.com

The Meraki Life

Let's take a moment to briefly discuss all of the Cisco Meraki services and benefits before continuing with the deployment guide. Meraki hardware operates via a cloud-hosted configuration and monitoring software suite aptly referred to as 'Dashboard'. Since Dashboard is cloud-hosted, all it requires is that Meraki devices be able to reach the internet – and thus the cloud – for configuration and data reporting. With this model of service, new features are deployed using firmware and are all included in a single license, one per device. This provides an ever evolving feature set to better serve networks as features are developed. The other benefit of a cloud managed solution is that client tracking and traffic analytics are included in the management tool and the full stack of Cisco Meraki products (switches, wireless, security and MDM) can be managed via a single pane of glass.

This last point is something that will save overhead and time for network engineers trying to deploy, maintain, and troubleshoot a network. Anyone who's ever managed a full network stack knows that different vendors have different configuration syntaxes and methods.

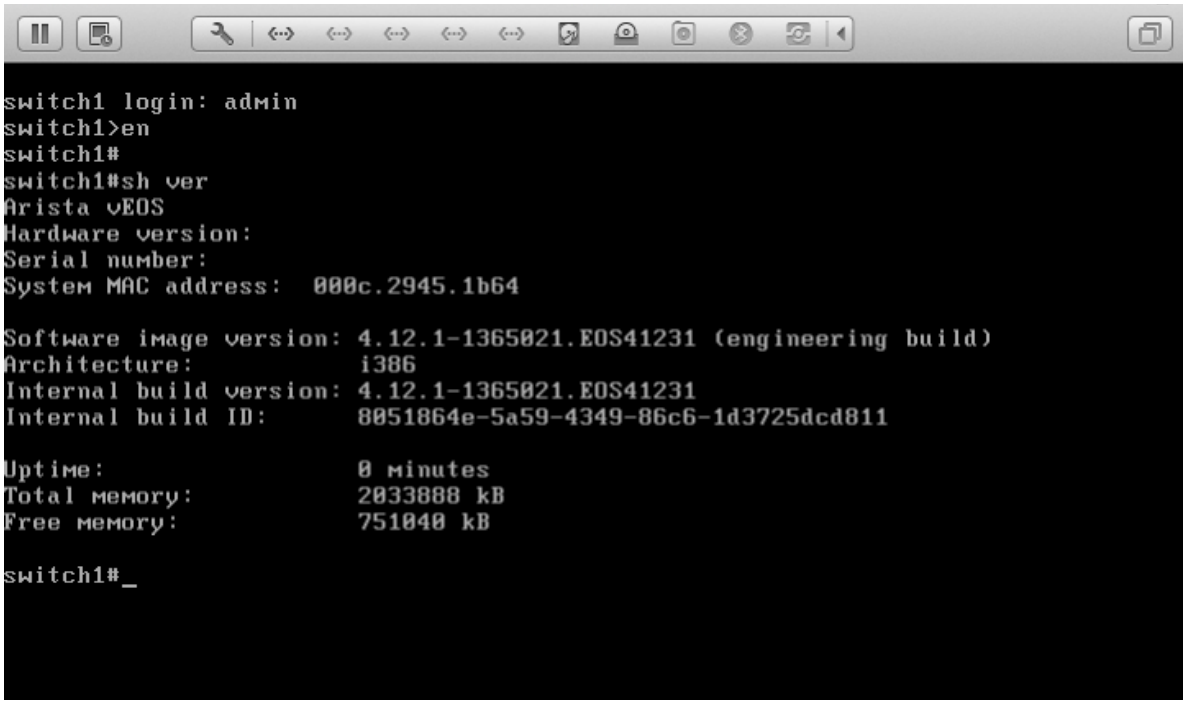

These methods often even differ between device types; whereas the Meraki user interface or Dashboard doesn't require special syntax – everything is as intuitive as a modern website. This makes configuration far simpler, requiring less expertise on syntax than on actual network design.

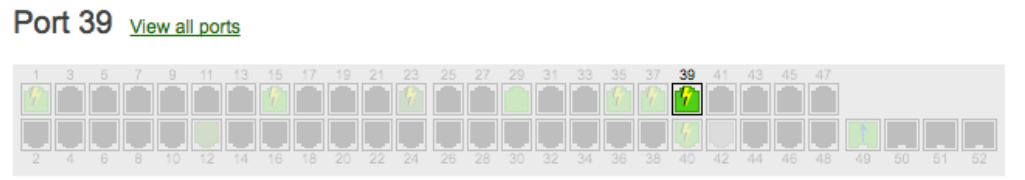

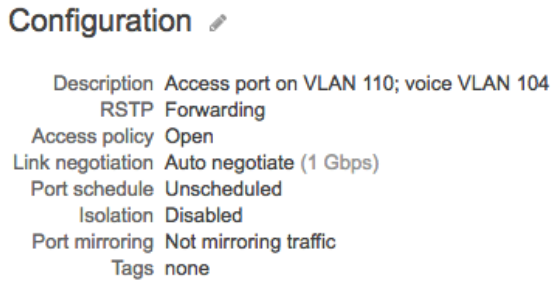

Not only is configuration straightforward without requiring unique syntax, but lots troubleshooting tools are also built directly into Dashboard. Packet captures can be run remotely, cable tests on switch ports, counters, connected clients, ping tools, and various other troubleshooting tools are readily availble in Dashboard and can be run with just the click of a button. This cuts down on time spent figuring out the best command to run to find information pertaining to an issue.

Many enterprise switches require direct or separate management access to be setup. This isn't the case with the cloud based solution - all that's needed is an internet connection. While this helps in all cases that the internet is reachable, the natural question becomes what happens when there's no internet? How can a new device be brought online initially or some information obtained into what might be happening? This can all be done via the local status page available on all Meraki devices. This page is hosted on each individual device and contains basic functionality to help bring equipment online and see a status for internet and cloud connectivity.

This covers the tip of the iceberg that is the Cisco Meraki solution. For more reading on the Meraki Dashboard and cloud architecture, please consult our documentation.

Meraki Switch Benefits

As the purpose of this document is to provide insight into larger deployments, it makes sense to provide information on what value the Cisco Meraki switch can bring to an enterprise network. To begin with, switching is the core of any deployment, so it is extremely valuable to be able to provide insight and visibility into this critical part of the network.

The Meraki switch line does this via an impressive lineup of visibility options and tools in Dashboard. The first and most basic of these is link state on the switch status page.

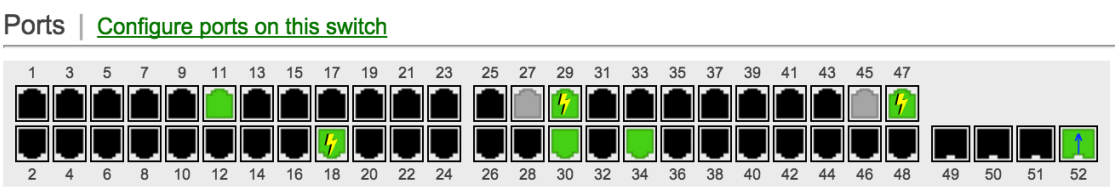

This is a great way to see port utilization on a switch quickly and efficiently. At a glance it is easy to identify ports providing PoE (⚡), the port status (green up, black down, grey disabled), and negotiated link speed (brighter green 1Gb/10Gb/s full duplex, darker green 10/100Mb/s). The switch view even indicates the link used as the uplink, denoted by the blue 'up' arrow (↑). In addition to this quick overview we can get further information on an individual port simply by clicking on it.

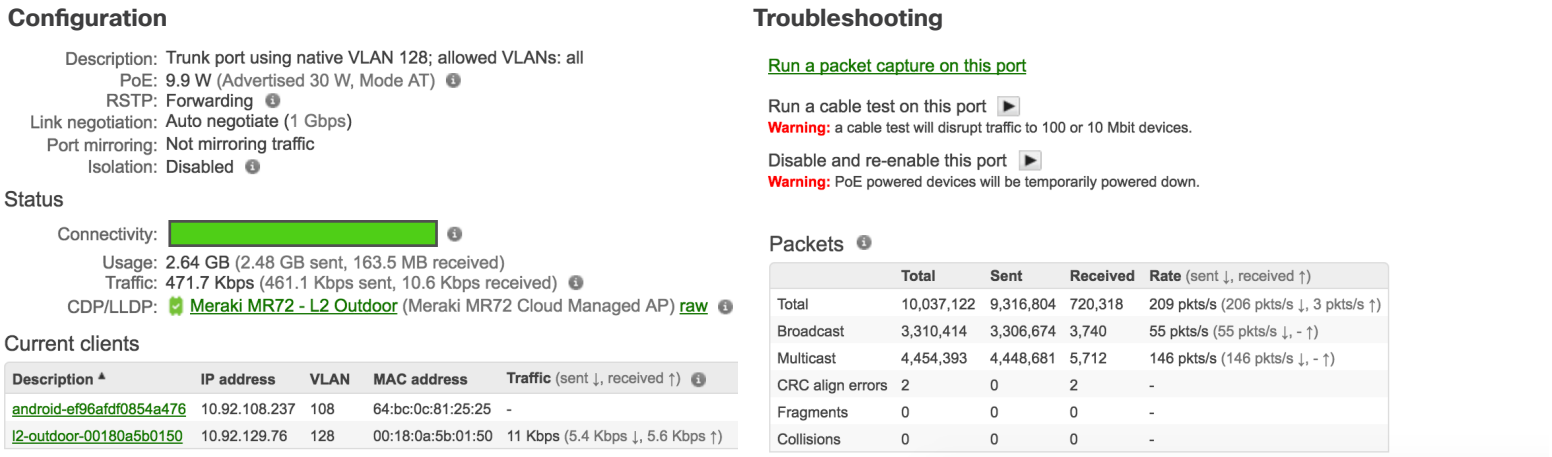

From the port view we can see additional details about the connected devices and the amount of traffic, as well as run various troubleshooting and debugging tools on the selected interface. This leads into one of the most beneficial aspects of the Cisco Meraki product offering - client monitoring and traffic analytics. While the information seen on the port page is useful, we can get much more detailed information on a specific client, or on the entire network by navigating to Network-wide > clients. The client view provides a very unique way of exploring how the network is being utilized, all the way to the application layer. All of these features and more are available to customers without requiring additional licenses.

In addition to these unique Dashboard features and platform benefits it is also possible to integrate Meraki switches with monitoring systems such as Cisco Prime, or via industry standard protocols such as SNMP and syslog.

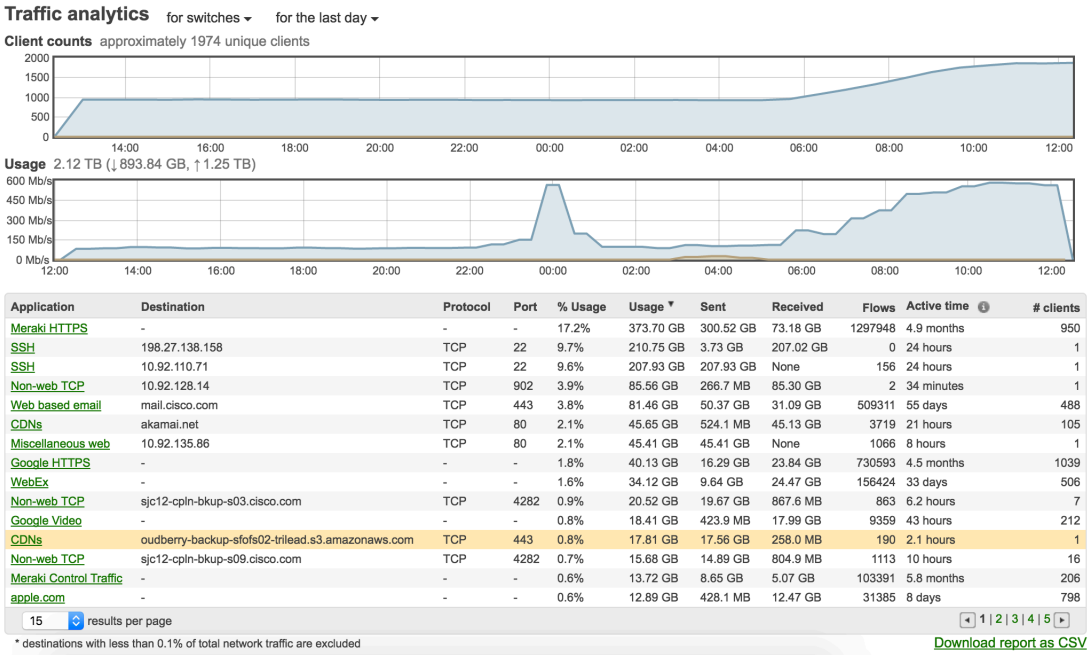

Application	Destination	Protocol	Port	% Usage	Usage ▼	Sent	Received	Flows	Active time	# clients
Meraki HTTPS	-	-	-	17.2%	373.70 GB	300.52 GB	73.18 GB	1297948	4.9 months	950
SSH	198.27.138.158	TCP	22	9.7%	210.75 GB	3.73 GB	207.02 GB	0	24 hours	1
SSH	10.92.110.71	TCP	22	9.6%	207.93 GB	207.93 GB	None	156	24 hours	1
Non-web TCP	10.92.128.14	TCP	902	3.9%	85.56 GB	266.7 MB	85.30 GB	2	34 minutes	1
Web based email	mail.cisco.com	TCP	443	3.8%	81.46 GB	50.37 GB	31.09 GB	509311	55 days	488
CDNs	akamai.net	TCP	80	2.1%	45.65 GB	524.1 MB	45.13 GB	3719	21 hours	105
Miscellaneous web	10.92.135.86	TCP	80	2.1%	45.41 GB	45.41 GB	None	1066	8 hours	1
Google HTTPS	-	-	-	1.8%	40.13 GB	16.29 GB	23.84 GB	730593	4.5 months	1039
WebEx	-	-	-	1.6%	34.12 GB	9.64 GB	24.47 GB	156424	33 days	506
Non-web TCP	sjc12-cpln-bkup-s03.cisco.com	TCP	4282	0.9%	20.52 GB	19.67 GB	867.6 MB	863	6.2 hours	7
Google Video	-	-	-	0.8%	18.41 GB	423.9 MB	17.99 GB	9359	43 hours	212
CDNs	oudberry-backup-sfofs02-trilead.s3.amazonaws.com	TCP	443	0.8%	17.81 GB	17.56 GB	258.0 MB	190	2.1 hours	1
Non-web TCP	sjc12-cpln-bkup-s09.cisco.com	TCP	4282	0.7%	15.68 GB	14.89 GB	804.9 MB	1113	10 hours	16
Meraki Control Traffic	-	-	-	0.6%	13.72 GB	8.65 GB	5.07 GB	103391	5.8 months	206
apple.com	-	-	-	0.6%	12.89 GB	428.1 MB	12.47 GB	31385	8 days	798

15 results per page

◂ 1 | 2 | 3 | 4 | 5 ▸

* destinations with less than 0.1% of total network traffic are excluded

Download report as CSV

Cisco Systems, Inc. | 500 Terry A. Francois Blvd, San Francisco, CA 94158 | (415) 432-1000 | sales@meraki.com

Campus Design

The Core

Most designs start with the network core. We'll be exploring two common design options - One for very large networks with chassis switches at the core and the other for small to medium environments that do not require chassis switching.

Let's begin with a large campus example, one that consists of multiple buildings and floors. Sometimes, these even traverse multiple geographic locations. In these scenarios it is most often a requirement to be able to aggregate many links and process large amounts of high-bandwidth routed traffic. In this scenario we will explore a hybrid network architecture, utilizing Cisco Nexus 9000 series at the core and Meraki switches at both the aggregation and access layer. Of course, Catalyst 4000 or 6000 series switches may be your core switch as well and similar design considerations will apply

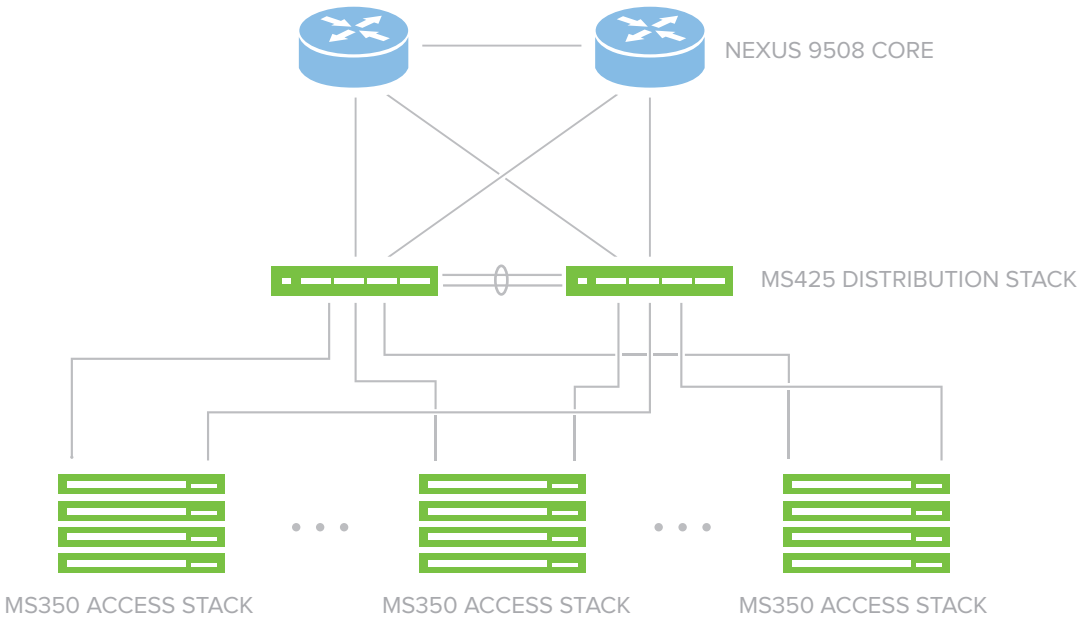

Cisco Systems, Inc. | 500 Terry A. Francois Blvd, San Francisco, CA 94158 | (415) 432-1000 | sales@meraki.com

The very first thing we'll want to do is configure core redundancy. For this example we will use two Nexus 9K switches configured in a vPC (Virtual Port Channel) pair. This lets us utilize both switches and introduce high failure resiliance. To achieve this we'll configure vPC using the following commands:

```
configure terminal
feature vpc
vpc domain <domain number>
peer-gateway
```

After this basic configuration we're ready to go ahead and configure the peer-link for our two cores to sync up and operate as one. To do this we'll setup a port-channel for this link and enable it using the following commands:

```
configure terminal
interface port-channel <port-channel number>
vpc peer-link
```

With our core set up and connected, we'll want to now dual-home downstream devices using the virtual port channels in order to add failure redundancy. This is relatively easy as we'll set up the port-channels and configure them to be part of the vPC using the following:

```
configure terminal
Interface port-channel <port-channel number>
vpc <number>
```

The last consideration is to configure (per-vlan spanning tree) PVST+ or (multiple spanning tree) MSTP to run for loop prevention. Either will work and is interoperable with Meraki switches so long as the native VLAN is present and passing untagged bridge port data units (BPDUs). This is because Meraki switches run rapid spanning tree (RSTP) and BPDUs are transmitted on the native VLAN (i.e. untagged).

At this stage we can add routed interfaces and enable any WAN or LAN routing services. These further configuration options can typically be found on cisco.com and tend to vary depending on how each network is designed.

Let's now discuss smaller network deployments using a collapsed-core architecture. As with larger deployments, to add redundancy we will want to deploy a solution that is both powerful and fully redundant, so we'll leverage a stackable distribution switch capable of dynamic rounting, first-hop redundancy and additional layer-3 features - the Meraki MS425. This product supports 40 Gigabit connections to interconnect the two core switches for physical redundancy as well as add protocol failover and gateway redundancy. In this example we will use two MS425s but feel free to adapt and expand on this example as it best suits your environment's needs.

We will begin by connecting the MS425s to their gateway upstream, along with providing both switches with an active internet connection. This will initiate a download of any available firmware updates directly from the cloud, in addition to fetching any configuration changes. To verify that both switches are online and connected to the Meraki Cloud, you can always check the status LED on either switch which should be solid white in color. A flashing LED indicates a software upgrade is in progress. You can always refer to our online documentation: MS425 Series Installation Guide for troubleshooting. Additionally, if you are logged into Dashboard, the switches should show a green status indicating they are connected and ready to go.

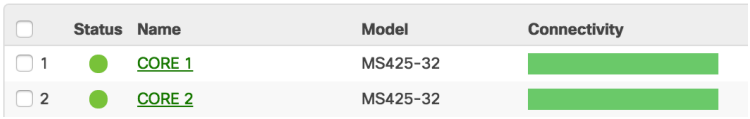

	Status	Name	Model	Connectivity
1	●	CORE 1	MS425-32	
2	●	CORE 2	MS425-32	

With the switches online, we'll need to configure the QSFP+ interfaces to act as stacking interfaces. You can use any of the ethernet interfaces for this purpose. This is very straightforward and can be done via Switch > Switch Ports or by selecting each individual switch and clicking on the ports. This feature is similar to Cisco VSS (Virtual Switching System) in that you are modifying standard ethernet interfaces to act as dedicated stacking interfaces.

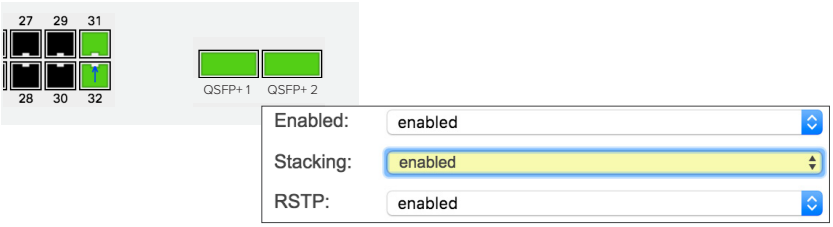

Once we have opened the configuration dialog, simply choose "stacking: enabled" as seen above, and save your changes. This will push the change to the switches and the ports will initiate and begin running Meraki stacking protocol. You can now provision the newly-created stack in Dashboard. This can be done via Switch > Switch Stacks in Dashboard. Simply select the switches using the "new stack" option or if you've already completed the previous step and the switches have downloaded their configuration, you can select them from the detected stacks list as shown below:

Choose a new name for the switch stack and save your changes. At this stage the switches are setup and ready to be provisioned with routed interfaces, static routes and any routing protocols that are desired. Routed VLANs can be configured on individual switches or on a switch stack via Switch > Routing and DHCP. Simply choose 'Add a static route' or 'Add an interface' and fill out the appropriate information, making sure to select the switch or switch stack that was defined earlier.

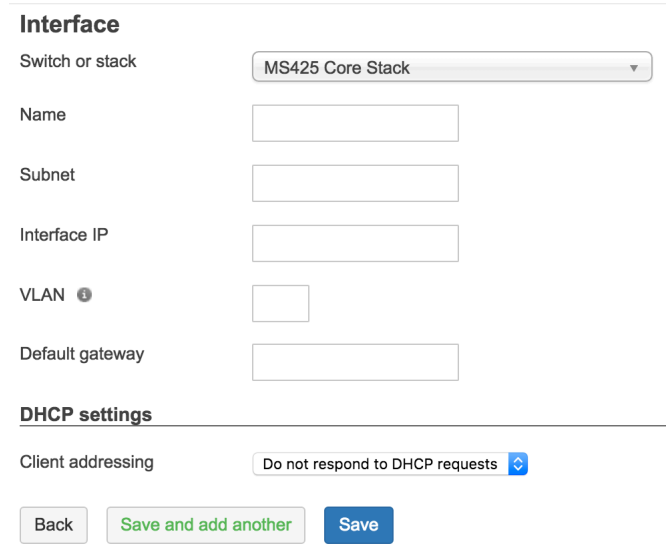

Once we have finished this configuration we can also enable OSPF depending on the size of the campus so that we can dynamically route between buildings. This is done via Switch > OSPF and will provide us with dynamic failover capabilities for redundant WAN paths as well as connections between buildings when applicable. Enable OSPF on the desired routed interface(s) and be sure to match configuration timers to their desired values.

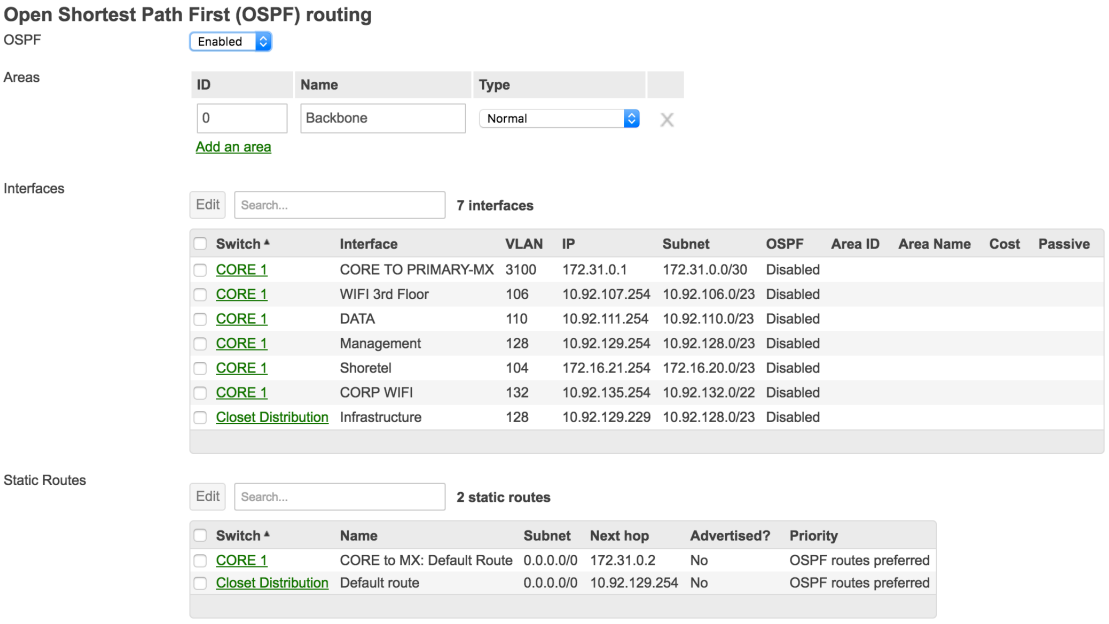

Open Shortest Path First (OSPF) routing

OSPF: Enabled

Areas

ID	Name	Type
0	Backbone	Normal

Add an area

Interfaces

Edit | Search... | 7 interfaces

Switch	Interface	VLAN	IP	Subnet	OSPF	Area ID	Area Name	Cost	Passive
CORE 1	CORE TO PRIMARY-MX	3100	172.31.0.1	172.31.0.0/30	Disabled				
CORE 1	WIFI 3rd Floor	106	10.92.107.254	10.92.106.0/23	Disabled				
CORE 1	DATA	110	10.92.111.254	10.92.110.0/23	Disabled				
CORE 1	Management	128	10.92.129.254	10.92.128.0/23	Disabled				
CORE 1	Shoretel	104	172.16.21.254	172.16.20.0/23	Disabled				
CORE 1	CORP WIFI	132	10.92.135.254	10.92.132.0/22	Disabled				
Closet Distribution	Infrastructure	128	10.92.129.229	10.92.128.0/23	Disabled				

Static Routes

Edit | Search... | 2 static routes

Switch	Name	Subnet	Next hop	Advertised?	Priority
CORE 1	CORE to MX: Default Route	0.0.0.0/0	172.31.0.2	No	OSPF routes preferred
Closet Distribution	Default route	0.0.0.0/0	10.92.129.254	No	OSPF routes preferred

This configuration will provide a solid foundation for the network core, allowing us to move to the access layer of the network.

Stacking at the Access Layer

With the network core fully provisioned, we are now ready to focus on configuring the access layer of the network. For most campus environments, end-to-end redundancy is key so it makes sense to introduce physical stacking capabilities into the edge of the network. This will ensure a resilient access layer, minimizing the impact of any physical or logical failures that may occur while giving us the bandwidth needed for today's demanding applications. The Meraki MS350 has been engineered for this purpose, with fully redundant power supplies and fans, along with support for stacking of a maximum of eight switches, providing up to 384 edge ports in a single logical stack. This is typically enough to accommodate a floor or a building wing. Coupled with 160Gbps stacking bandwidth, we can utilize multiple uplinks with cross-stack link aggregation (MLAG) for non-blocking throughput to the aggregation or core layers of the network. MS350 stacks seamlessly integrate with the stacked core of MS425s that were configured previously to provide end-to-end redundancy and minimize network downtime.

To set up our MS350 switch stacks, we will begin by connecting our stacking links. A Meraki stacking cable is included with each MS350 switch, and a ring topology is recommended. To create a full ring, start by connecting switch 1 / stack port 1 to switch 2 / stack port 2 and so forth, with the bottom switch connecting to the top switch to complete the ring.

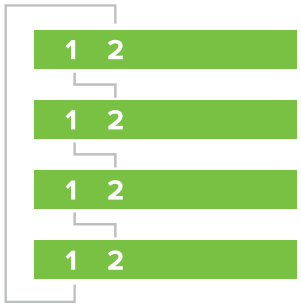

Once your switch stacks are properly cabled, we can power each switch and provide an uplink so we can connect each stack to the Meraki cloud. Dashboard will automatically detect if a stack not already been provisioned and will prompt you to both provision and name the stack via Switch > Switch Stacks:

Detected potential stacks

Stack Members	Actions
Switch 3 Switch 1 Switch 5 Switch 6 Switch 2 Switch 7 Switch 4 Switch 8	Provision this stack

10 results per page < 1 >

Cisco Systems, Inc. | 500 Terry A. Francois Blvd, San Francisco, CA 94158 | (415) 432-1000 | sales@meraki.com

Once each stack has been named, you can proceed to configure the individual ports in each stack using Meraki's Virtual Stacking technology (Configure > Switch Ports). To learn more on Meraki's virtual stacking technology, please see our Meraki Stacking whitepaper.

With the access layer deployed and provisioned, you are now ready to configure redundant, active/active uplinks from each stack to the network core. Using Link aggregation (LACP), you are able to bundle up to eight links per switch or switch stack. You will want to be considerate of bandwidth demands, and it is recommended you configure at least two links for any stack, with each link spanning multiple stack members in order to limit outages caused by downed links or cabling problems. Link aggregation is simple to configure using Meraki's Virtual Stacking, simply navigate to Configure > Switch Ports, select the interfaces to be configured in a link aggregate and click the aggregate button towards the top, as demonstrated in the figure below:

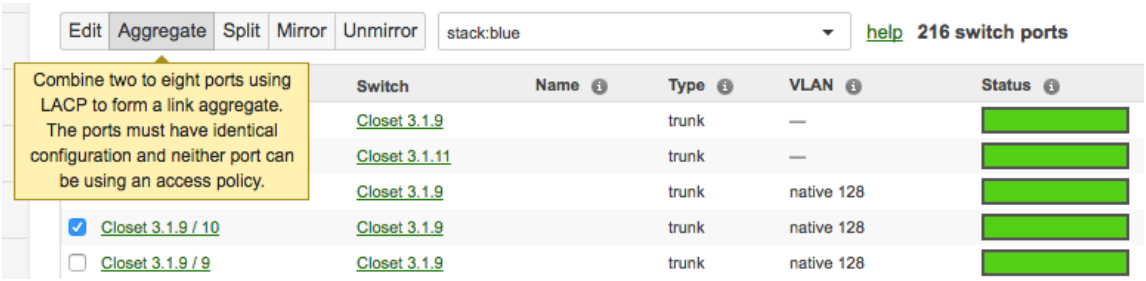

The MS350 switch family includes a multigigabit ethernet model that pairs with our MR53 access point to provide speeds greater than a gigabit over a single run of twisted-pair ethernet. This is a perfect solution for any campus looking to deploy 802.11ac Wave 2 without costly cable upgrades. For more information on Multigigabit and 802.11ac Wave 2, please see our multigigabit technology brief.

QoS Considerations in the Campus

In any campus deployment, traffic prioritization is key to keeping critical network applications running, even under heavy load. This is done through the use of Quality of Service (QoS) configuration. The simplest explanation of QoS is the prioritization of traffic, ensuring that important or latency-sensitive traffic will get bandwidth before less demanding traffic. To read more on this topic please see our guide, QoS on Meraki.

The table below lists the commonly used DSCP values as described by RFC 2475. To keep things standardized it's recommended to utilize these values, unless the deployment already uses a different set of values.

DSCP VALUE	DECIMAL	VALUE MEANING	DROP PROBABILITY	EQUIVILANT IP PRECDENCE VALUE
101 110	46	High Priority Epedited Forwarding (EF)	N/A	101 - Critical
000 000	0	Best Effort	N/A	000 - Routine
001 010	10	AF11	Low	001 - Priority
001 100	12	AF12	Medium	001 - Priority
001 110	14	AF13	High	001 - Priority
010 010	18	AF21	Low	001 - Immediate
010 100	20	AF22	Medium	001 - Immediate
010 110	22	AF23	High	001 - Immediate
011 010	26	AF31	Low	011 - Flash
011 100	28	AF32	Medium	011 - Flash
011 110	30	AF33	High	011 - Flash
100 010	34	AF41	Low	100 - Flash Override
100 100	36	AF42	Medium	100 - Flash Override
100 110	38	AF43	High	100 - Flash Override

DSCP VALUE	DECIMAL	VALUE MEANING	DROP PROBABILITY	EQUIVILANT IP PRECDENCE VALUE
001 000	8	CS1	1	-
010 000	16	CS2	2	-
011 000	24	CS3	3	-
100 000	32	CS4	4	-
101 000	40	CS5	5	-
110 000	48	CS6	6	-
111 000	56	CS7	7	-
000 000	0	Default	-	-
101 110	46	EF	-	-

To help with the reasoning behind the above chart, here's the IP precedence priority from lowest to highest as per RFC 791.

VALUE	DESCRIPTION
000 (0)	Routine or Best Effort
001 (1)	Priority
010 (2)	Immediate
011 (3)	Flash - Mainly used for Voice Signaling for for Video
100 (4)	Flash Override
101 (5)	Critical - mainly used for Voice RTP
110 (6)	Internet
111 (7)	Network

We'll want to mirror the rest of the network deployment on the access switches so that QoS is the same across the network, so match default Cisco settings we'll be setting the following:

CoS Value	0	1	2	3	4	5	6	7
DSCP Value	0	8	16	24	32	40	48	56
DSCP Values	0	8, 10	16, 18	24, 26	32, 34	40, 46	48	56
CoS Values	0	1	2	3	4	5	6	7

You'll configure these under Switch > Switch Settings and define the QoS simply through the user interface:

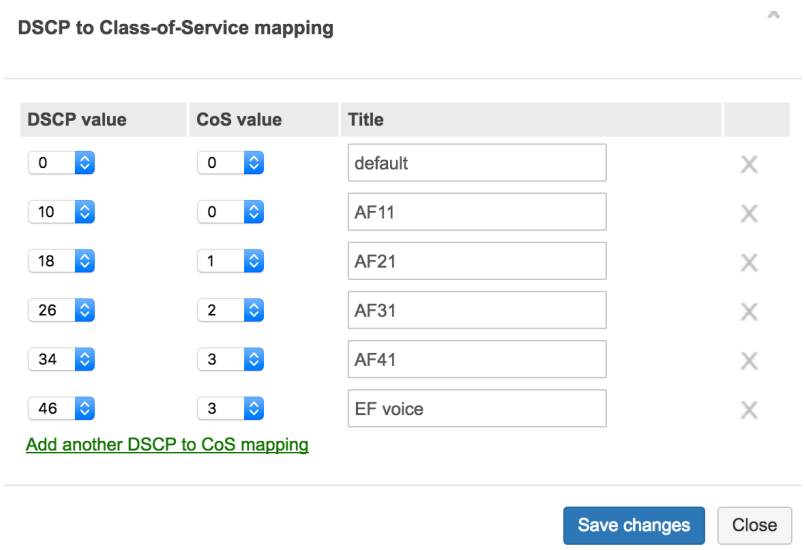

Security Settings

With QoS properly configured and traffic being prioritized for your network, any administrator who wants to make sure they have the most control of their network will want to implement security settings. To achieve this we can use the DHCP server detection mechanism and set it to automatically block new DHCP servers on the network. This will help protect agains malicious attacks and accidents that can cause significant downtime. You will want to whitelist the DHCP servers that should be permitted to keep the network up and operational.

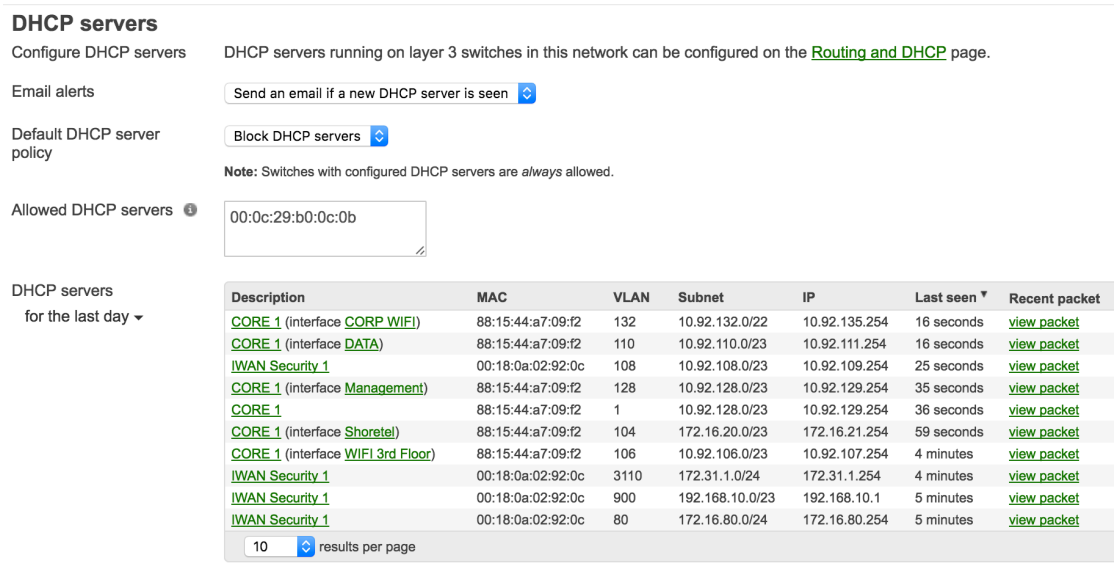

DHCP servers

Configure DHCP servers — DHCP servers running on layer 3 switches in this network can be configured on the Routing and DHCP page.

Email alerts — Send an email if a new DHCP server is seen

Default DHCP server policy — Block DHCP servers

Note: Switches with configured DHCP servers are *always* allowed.

Allowed DHCP servers — 00:0c:29:b0:0c:0b

DHCP servers for the last day

Description	MAC	VLAN	Subnet	IP	Last seen	Recent packet
CORE 1 (interface CORP WIFI)	88:15:44:a7:09:f2	132	10.92.132.0/22	10.92.135.254	16 seconds	view packet
CORE 1 (interface DATA)	88:15:44:a7:09:f2	110	10.92.110.0/23	10.92.111.254	16 seconds	view packet
IWAN Security 1	00:18:0a:02:92:0c	108	10.92.108.0/23	10.92.109.254	25 seconds	view packet
CORE 1 (interface Management)	88:15:44:a7:09:f2	128	10.92.128.0/23	10.92.129.254	35 seconds	view packet
CORE 1	88:15:44:a7:09:f2	1	10.92.128.0/23	10.92.129.254	36 seconds	view packet
CORE 1 (interface Shoretel)	88:15:44:a7:09:f2	104	172.16.20.0/23	172.16.21.254	59 seconds	view packet
CORE 1 (interface WIFI 3rd Floor)	88:15:44:a7:09:f2	106	10.92.106.0/23	10.92.107.254	4 minutes	view packet
IWAN Security 1	00:18:0a:02:92:0c	3110	172.31.1.0/24	172.31.1.254	4 minutes	view packet
IWAN Security 1	00:18:0a:02:92:0c	900	192.168.10.0/23	192.168.10.1	5 minutes	view packet
IWAN Security 1	00:18:0a:02:92:0c	80	172.16.80.0/24	172.16.80.254	5 minutes	view packet

10 results per page

Another important area of network security is authentication. Most security-minded deployments will have a RADIUS server configured to authenticate wired clients on the network. Meraki switches can integrate seamlessly to provide authentication via 802.1x, MAC Authentication Bypass (MAB) or a hybrid of these on any switch port. This provides fine-tuned control on who connects and which resources are accessible to each client. With hybrid authentication configured, a switchport will first require 802.1x and fall back to MAB before choosing restricting client access or assigning a guest / remediation VLAN. It is highly recommended to set up a remediation VLAN to quartantine unauthorized, guest or non-compliant devices.

In addition to the authentication methods mentioned above, Meraki also includes RADIUS server monitoring to enable use of Hybrid Auth, whereby if the RADIUS server is offline, client sessions will be reinitiated once a RADIUS server is available after an outage. This allows for seamless recovery of client sessions and access to network resources when authentication is available, without manual intervention. Use of this particular is extremely useful if the network utilizes a data center or cloud hosted radius server that may become unreachable.

As with any of the above security considerations, no network is complete without the use of Access Control Lists (ACLs) to be able to keep traffic restricted between VLANs. While this is achievable on Meraki switches, the actual implementation of ACLs is typically done on the upstream core, where the routing is done in this deployment scenario. For more information on configuring ACLs on the core switches please visit the following: Configuring IP Access Lists

Multiple VLANs

Another important design consideration is segmenting the network into VLANs.. Segmentation is important for three key reasons. First, it creates logically distinct 'pieces' of the network that can be organized by function or security importance, so that policies can be more easily applied to the correct pieces. Second, it limits the reach of broadcast packets, lessening the amount of traffic that the switch has to process and reducing the potential impact of broadcast storms. Third, segmentation can be valuable in troubleshooting, since it can make it easier to locate a user or device based on their IP address.

For example, the network design may utilize VLAN 3 for Data and VLAN 4 for Voice on the first floor, then apply VLAN 5 for Data and VLAN 6 for Voice on the second floor. When configuring the distribution ports that provide connectivity to the first floor, only VLANs 5 and 6 would be required. Such a design can easily be scaled up to a multi-building campus. Configuring which VLANs are allowed on a switchport is very straightforward in dashboard and can be accomplished by using a comma separated list. As an example, to allow VLANs 1,2,3,4, 5, 6, 10, and 20, simply type '1-6,10,20 into the port's 'allowed VLANs' field.

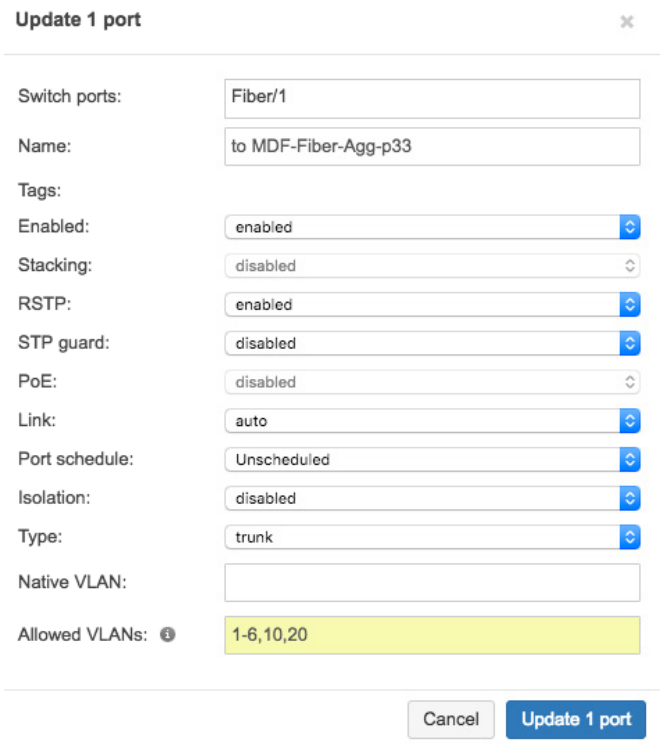

It is important to note that Meraki switches do not require VLANs to be manually created or added in order for them to be accepted on interfaces. Therefore, protocols such as Cisco VTP are not necessary for VLAN configuration, though they are in most cases compatible with MS switches.

For any deployment using Voice over IP (VOIP) it's a good idea to ensure the voice traffic gets assigned to the correct VLAN. A lot of times this voice VLAN is defined in addition to the normal traffic VLAN and modern phones will often have a PC connected through them. With Meraki switches we can easily assign a voice VLAN that the phone will get passed so that it can tag its traffic into the appropriate VLAN. This transaction is done via CDP or LLDP, depending on the phone model being used. The configuration of this VLAN is straightforward in dashboard.Simply configure the port as an access port, defining the normal VLAN and then additionally the Voice VLAN.

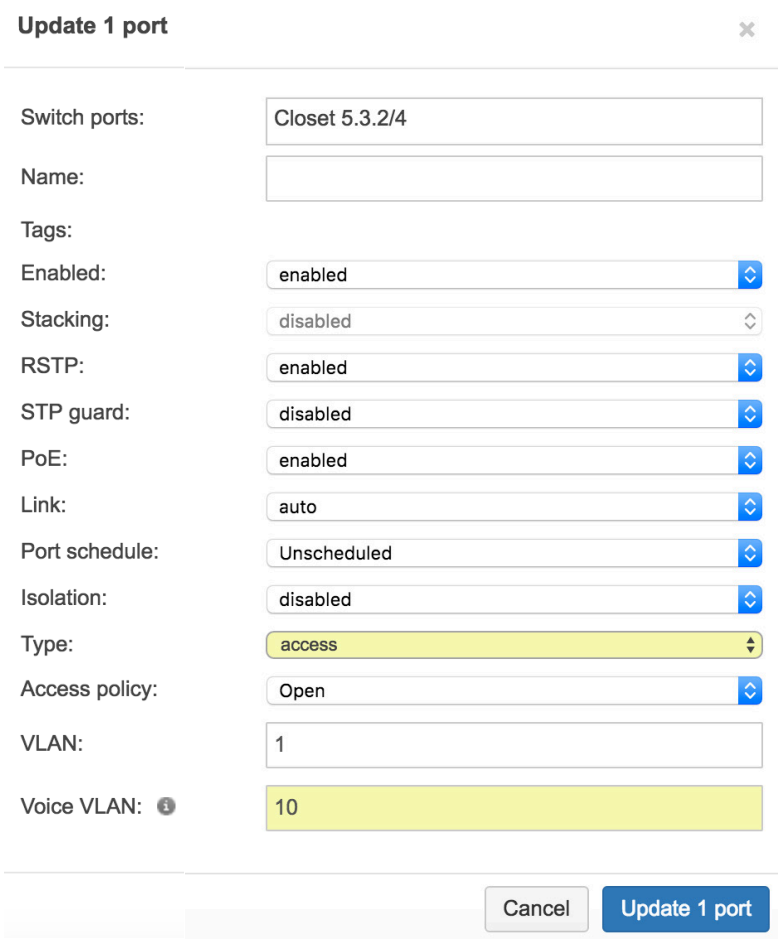

Administrative Access Control

Multiple levels of administrative visibility and access can be configured within the Meraki Dashboard. These privileges can be assigned at the network, device, or port level. Administrators can also be given privileges to certain groups of ports so that different functions in the network can be easily divided up amongst the appropriate IT stakeholders.

For example, you may want your helpdesk staff to have visibility and control over end-user access ports, your networking team to have control over infrastructure, server, or distribution ports,, and non-network operations stakeholders to have visibility, but no ability to modify configuration. This is accomplished with the use of tags. Tags can be applied to switches or switchports, and administrators can be granted access to only resources that bear the appropriate tag.

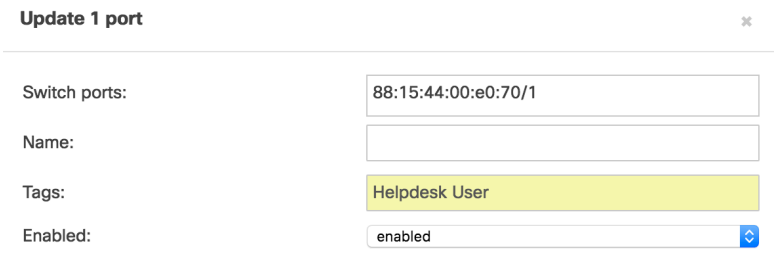

Cisco Systems, Inc. | 500 Terry A. Francois Blvd, San Francisco, CA 94158 | (415) 432-1000 | sales@meraki.com

Visibility

At this point the campus network setup is complete, utilizing the cloud. The next step would be ensuring users transition smoothly to the brand new network. The Dashboard can be used to provide excellent reporting and knowledge of what's happening in the new campus deployment. Dashboard will provide an overview of what clients are doing with Meraki's built-in, powerful traffic analytics.

Meraki's layer 7 application visibility is enhanced to dynamically detect applications running on the network, and provide hostname and IP address visibility. This information can be used to understand user behavior on the network and make policy decisions, such as creating custom traffic shaping rules, or applying group policies to specific users.

ENABLING TRAFFIC ANALYTICS

This is an opt-in feature, and can be enabled under the Configure > Network-wide settings page by selecting the 'Enable Hostname Visibility' functionality:

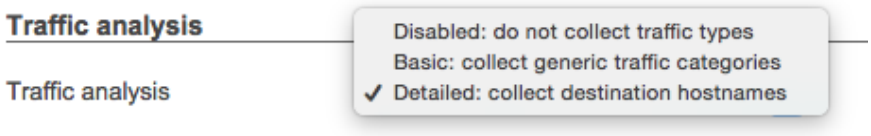

TRAFFIC ANALYTICS

The enhanced Traffic analytics page will be visible under the Monitor menu whenever hostname visibility is enabled.

This page will provide a total unique client count across an entire network over time. The view can be customized for different time periods (last 2 hours, week, day, and month), and on a per-SSID basis. This page will show the following information on a per-network or per-SSID basis:

- Application
- Specific destination for broad application categories such as 'Miscellaneous secure web'
- Protocol information
- Port information
- Usage breakdowns by %, data sent and received, number of flows
- Total active time on application across all clients
- Number of clients

SIGNATURE OR APPLICATION-LEVEL ANALYTICS

By clicking on an application signature (e.g. 'Dropbox' or 'Non-web TCP'), it's possible to see a complete breakdown of hostnames and IP addresses comprising this application category. Use this information to understand the communication patterns of certain types of traffic.

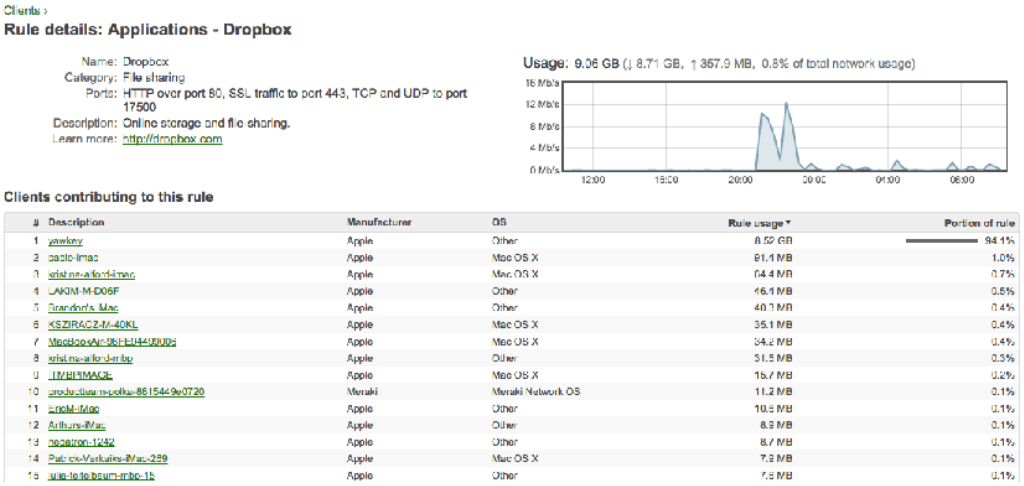

This page will show you the following information on a per-application basis:

- Application name, category, ports, description
- Usage over time
- List of users
- Destination list - hostnames and IP addresses contributing to this application
- Total number of clients per destination
- Time spent per client

USER LEVEL ANALYTICS

By clicking on a specific user, it's possible to see a complete breakdown of hostnames and IP addresses this user has visited, including the time spent on each destination. Use this information to understand individual user behavior and apply policies on a per-user basis.

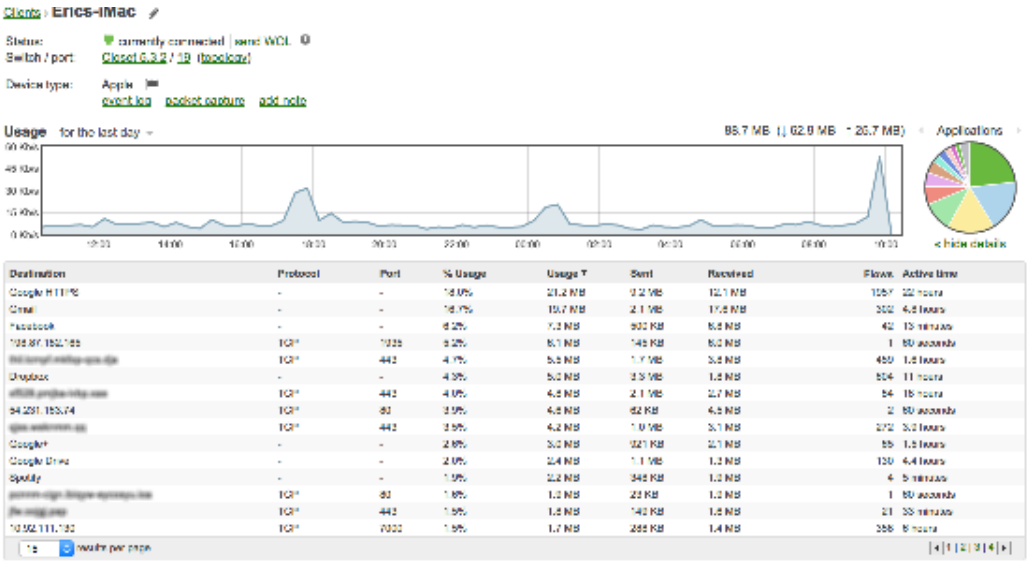

Summary reports can be mailed directly to the network administrator, relieving them of at least one task when their plate is full. These reports provide updates on the network and can be shared directly to key stakeholders with little to no effort.

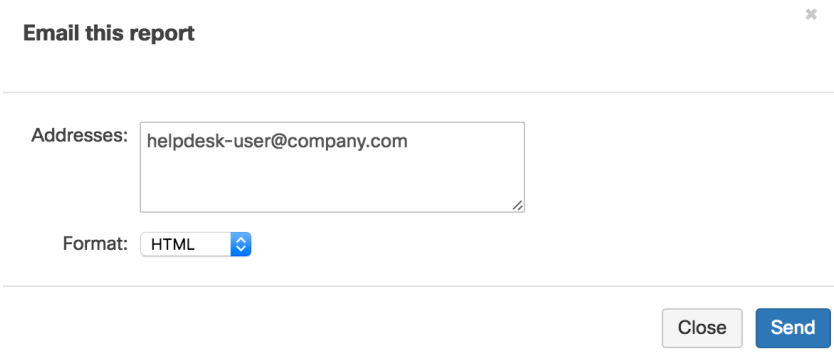

Cisco Systems, Inc. | 500 Terry A. Francois Blvd, San Francisco, CA 94158 | (415) 432-1000 | sales@meraki.com

Troubleshooting

Traffic analytics and summary reports emailed directly to the inbox help ease the burden of the engineer. That is until the first trouble ticket comes in from someone unable to print to their local printer or unable to access a resource on the file server. While this used to be a headache, the task can be done quickly utilizing the visibility of dashboard. Simply open up the clients page and type in the name associated with the user or their PC. This will filter the list directly to the machine in question.

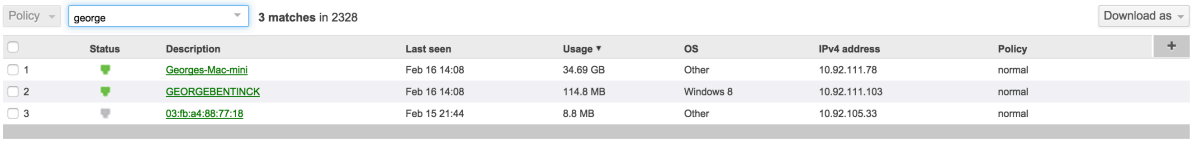

Once the machine in question has been located there's more information that can be provided. By clicking on the device we're interested in we can get details such as what switch/port or even AP it's connected to as well as IP address, MAC address and even firewall information.

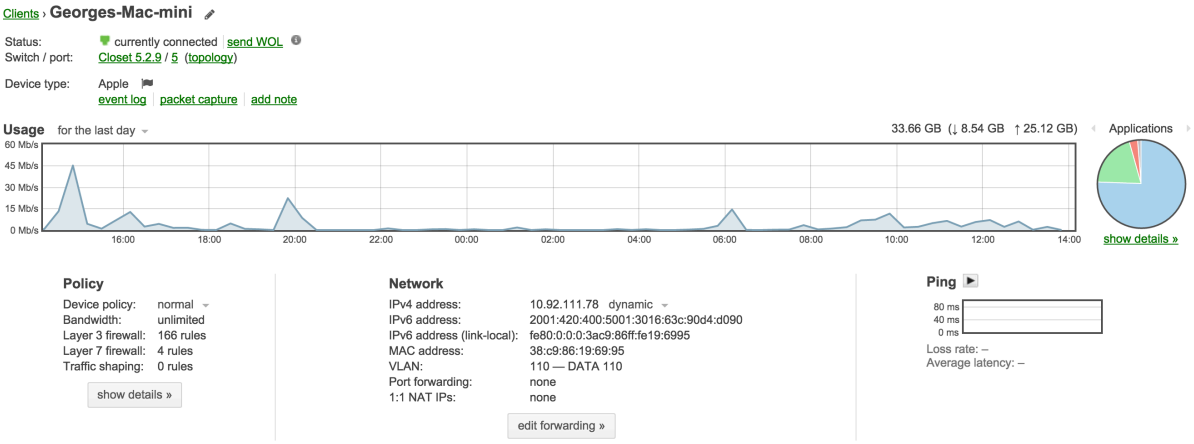

This screen allows us to quickly get an overview of the client and even try to ping it. We can also simply take a packet capture of the traffic while having the end user attempt the action that was failing. By being able to do this remotely and quickly via dashboard we save time and money by not having to trace cables and figure out the specific information about the end user device. There's even additional information if we're tracking clients using Systems Manager (SM), enabling us to see what software is installed and making sure it's in sync with any updates or policies that might be applied via SM.

Returning to troubleshooting, if we want to rule out a physical layer issue this can be done straight from dashboard using the cable test utility available on the switches.

Cable test

Warning: this test will disrupt traffic to 100 or 10 Mbit devices.

Ports (eg. 1 or 1,2,3 or 1 - 3): 17

Port	Link speed	Length (± 3m)	Status	Pair 1	Pair 2	Pair 3	Pair 4
17	1Gfdx	98.25 m	OK	ok	ok	ok	ok

This allows us to rule out the physical layer simply and easily without trying to find a cable tester or sending someone to a remote location with cables and a tester to verify, saving time and money. With the information provided by the cable test, extra cables or a new cable run can be implemented if the test comes back as failed. Another excellent dashboard tool is the topology view, which provides insight into the network and how it's connected, even showing a redundant link that's not plugged in. In the case of a hybrid deployment, information can be pulled directly from directly connected devices to see that things are wired properly in the infrastructure.

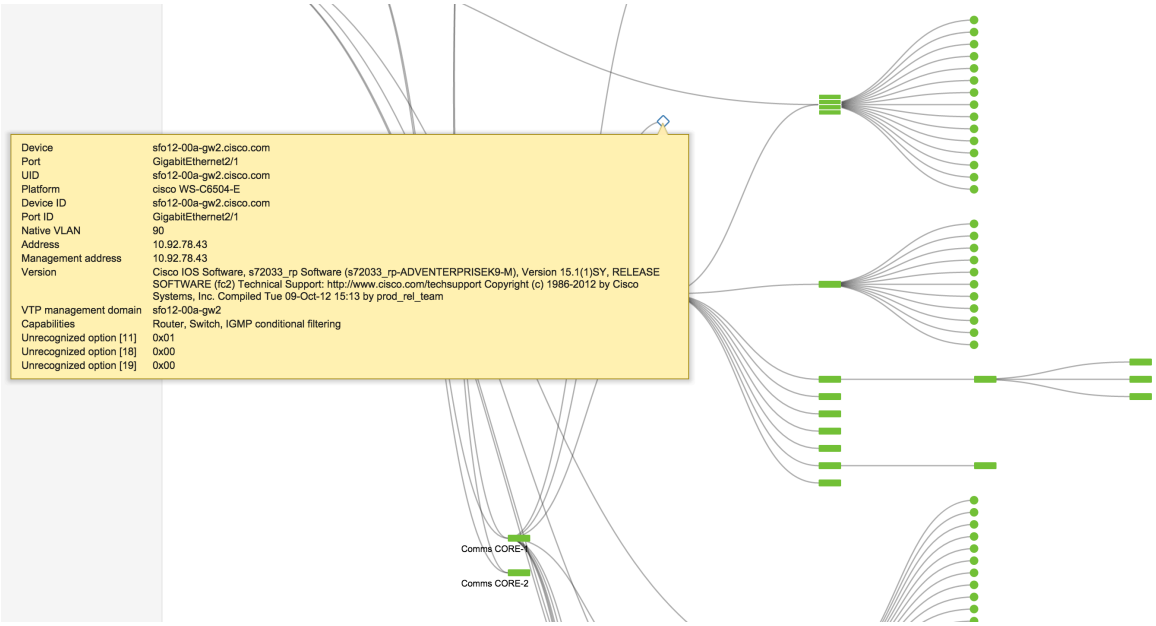

Cisco Systems, Inc. | 500 Terry A. Francois Blvd, San Francisco, CA 94158 | (415) 432-1000 | sales@meraki.com

Conclusion

Building large scale networks for today's dynamic world of wired and wireless devices is no simple task. Businesses are required to be far more agile in both IT application delivery and geography, while pivoting quickly to meet the demands of their customers and markets. When many of today's network operating systems were developed, the demands placed on the network were more straightforward and predictable. Today's networks needs are far more complex and rarely constant, yet administrators are still required to support performance-sensitive and often mission-critical apps like voice and video.

This document has illustrated the way in which a campus networking environment can benefit from Meraki technology. True zero–touch provisioning changes forever the way networks are staged and deployed. Powerful application monitoring provides at–a–glance visibility into how network resources are being used. Convenient remote troubleshooting tools transform the ability of an IT team to respond to low level problems.

All of this makes a significant difference to the cost of running an IT operation. With all aspects of support covered, with feature upgrades covered, with lower operational costs, customers can expect a Meraki network to deliver significant real world savings over the course of its lifetime.

Let us know how we can improve even further. Look for the Make–a–wish box on the lower right of every dashboard page.