Cisco Meraki Voice-over-IP

Deployment Guide

Contents

Preface

Purpose

The purpose of this design guide is to provide guidance and best practices for deploying voice-over-IP (VoIP) services in a branch-office environment using Cisco® Meraki® MS switches, Meraki MR access points, and Meraki MX security appliances in conjunction with Cisco Integrated Services Routers (ISRs) and Cisco Unified Communications Manager.

This guide enhances the **Cisco Meraki Branch-Office Deployment Guide** and should be used along with the other resources mentioned in the reference section.

Audience

This guide is for networking professionals who are responsible for designing, implementing, or administering a network that includes Cisco Meraki MS switches, Meraki MX security appliances, and Meraki MR access points; Cisco ISRs; and Cisco aggregation services routers (ASRs). Readers of this guide are expected to have prior experience working with the Cisco IOS® Software and Cisco Meraki Dashboard, Cisco Unified Communications Manager (part number CUCM), and Cisco Unified Communications Manager Express (part number CME), in addition to being familiar with the concepts and terminology of local-area networking, switching, and VoIP. The "References" section provides additional resources for these specific topics.

Conventions

The publication uses these conventions to convey instructions and information:

- Command names are in **boldface** text.
- System displays are in **italicized** font.

Introduction

This guide describes guidelines and practices around deploying a VoIP solution within a branch-office or satellite location using a Cisco Meraki network and integrating it with existing Cisco solutions. It covers the following independent scenarios:

- A Cisco Meraki branch-office network (MX, MS, or MR) that connects to a Cisco campus network using Cisco Unified Communications Manager for telephony services.
- A Cisco Meraki branch-office network (MX, MS, or MR) that uses a Cisco ISR as the edge router and Cisco Unified Communications Manager Express for telephony services and the Cisco Meraki platform for wireless, LAN, and Universal Threat Management (UTM) services.

Figure 1 illustrates the topology that is referenced throughout this document.

Figure 1. Topology

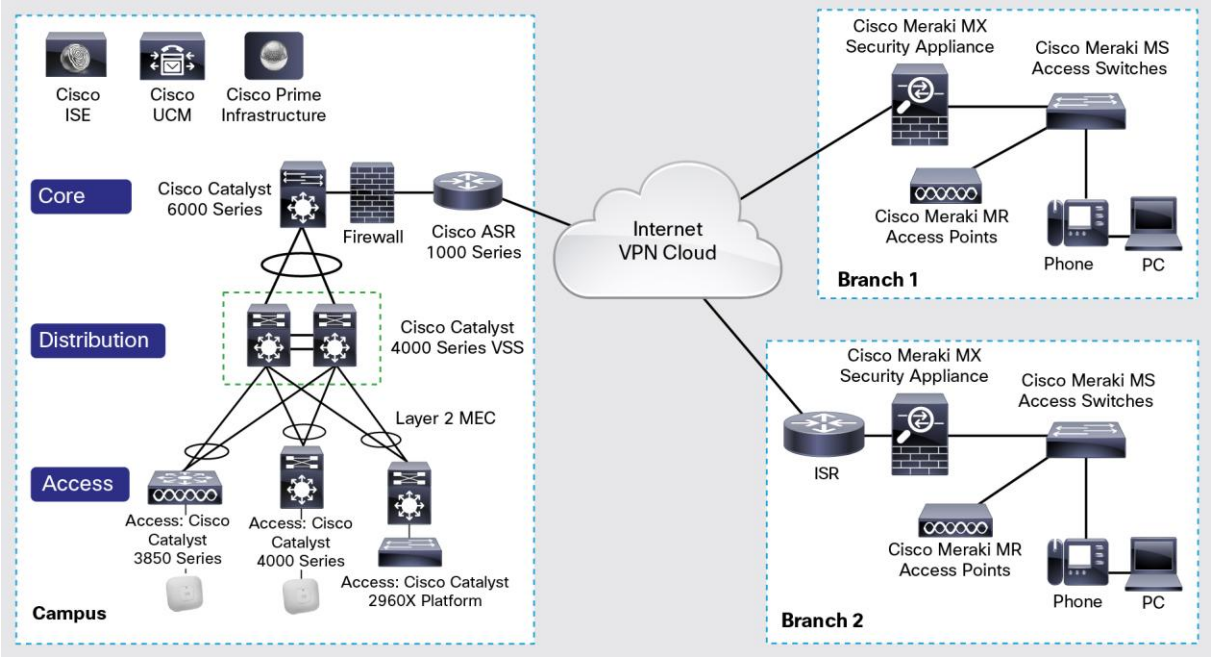

Evolution of the Branch Office

Today's branch-office networks have similar requirements and expectations placed upon them as head-office or campus networks, typically without the onsite expertise to support and maintain those networks.

The new standards for a modern branch office typically include:

- High-performance wireless LAN (WLAN) networks for delivery of rich media and applications
- Network segmentation to separate and secure networks carrying financial and business-critical data from other applications
- Nonstop availability because branch-office environments are usually 24-hour operations
- Integrated security
- Ease of management
- Rapid deployment; branch office in a box
- Network intelligence and analytics
- Centralized call management

Traditionally branch-office networks have been deployed as a mini-representation of the campus network, but the evolution of cloud-managed networks has presented an opportunity to combine powerful remote management tools with the centralized IT services of the headquarters or campus.

Network designers have the option of using cloud-managed infrastructure for elements of the branch-office network that best suits their requirements, including the ability to deploy functionally rich, secure, highly available wired and wireless infrastructure using the Cisco Meraki cloud-managed offerings, alongside rich collaboration and WAN services inherent in the Cisco ISRs.

This flexibility addresses the ease of management, rapid deployment, and network analytics requirements of retail environments, for instance, while substantially reducing the time investment and cost to deliver this capability.

Branch-Office Network Setup

This guide covers two branch-office deployment cases, each deployable independent of the other.

Branch office 1 is a Cisco Meraki cloud-managed branch-office network composed of Cisco Meraki devices (MR access points, MS switches, and an MX security appliance for connectivity to the WAN). The MX security appliance is configured for a site-to-site VPN tunnel to the main campus. The main-campus VPN headend is a Cisco ASR 1000 Aggregation Services Router, although you also can use an MX security appliance. In this branch office, the Cisco Unified Communications Manager at the main campus provides call-processing and telephony services.

Branch office 2 is a Cisco Meraki network (MX, MS, and MR) with a Cisco ISR acting as a gateway device. The Cisco ISR is configured for a site-to-site VPN tunnel to the main campus. The Cisco Meraki MX appliance provides firewall, Dynamic Host Configuration Protocol (DHCP), and intrusion detection and prevention systems (IDSs and IPSs, respectively), as well as utilizing the dashboard interface. The main campus VPN headend is a Cisco ASR 1000. In this branch office, the Cisco ISR running Cisco Unified Communications Manager Express provides call-processing and telephony services locally.

The following network elements are used throughout these deployment scenarios:

- Cisco Meraki MR access points
- Cisco Meraki MS switches
- Cisco Meraki MX security appliance
- Cisco 2900 Integrated Services Router
- Cisco ASR 1000 Aggregation Services Router

LAN

IP Addressing and VLANs

Branch Office 1

In branch office 1, the MX appliance is used to define the subnets and VLANs in the network to help ensure that traffic is segregated in accordance with best practice. In this deployment, at least three VLANs will be created: VLAN 115 for voice traffic, VLAN 120 for data traffic, and VLAN 200 for management traffic. For deployment, follow these steps:

Step 1. Navigate to the **Security Appliance > Addressing & VLANs** page. Ensure that VLANs are enabled.

Step 2. Click **Add a Local** VLAN to create each needed subnet.

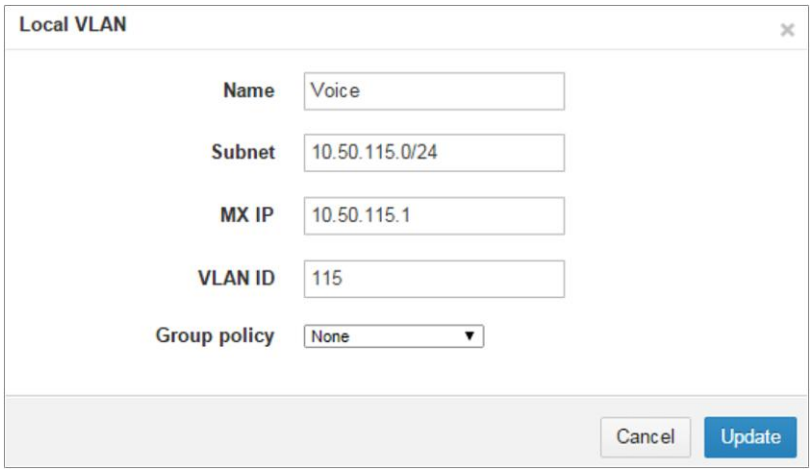

Note:

- **Name** is a descriptive field used only to refer to the created subnet.
- **Subnet** is the subnet to be created, notated in Classless Inter-Domain Routing (CIDR) format.
- **MX IP** is the IP address of the MX appliance within the created subnet. This address will function as the gateway address of devices within the subnet.
- **VLAD ID** is the assignment of the VLAN tag; it must be a number between 1 and 4096.
- **Group Policy** allows you to assign an existing group policy to traffic within the subnet. More information about group policies is available at: https://documentation.meraki.com/MR/Group_Policies_and_Blacklisting/Creating_and_Applying_Group_Policies.

Step 3. Click **Update** to validate the VLAN creation. Repeat this process for each required subnet.

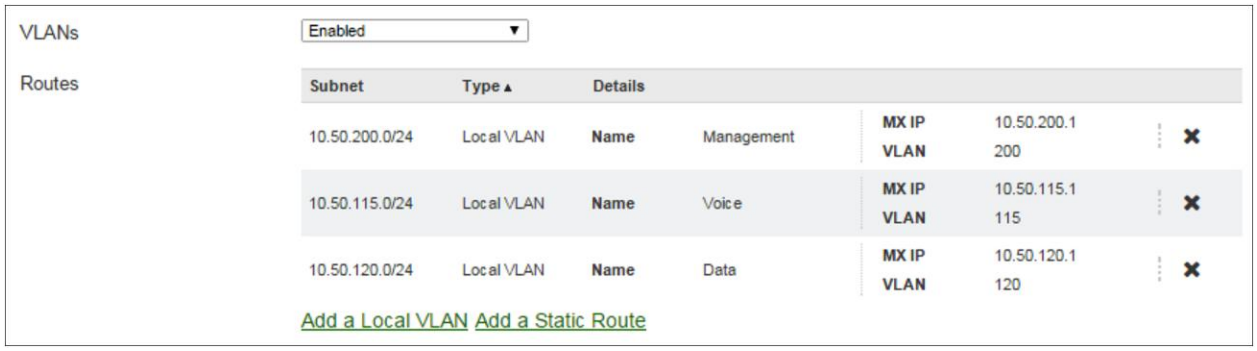

Branch Office 2

In branch office 2, the Cisco ISR provides the Layer 3 termination and routing between VLANs and upstream networks, and provides Network Address Translation (NAT) between the local networks and the Internet. The Meraki MX appliance is configured to operate in passthrough mode as a Layer 2 bridge, and provides services such as firewall, traffic shaping, and security and content filtering.

The specific ISR configuration deployed in this example is included as follows as a quick reference. The VLAN and IP address scheme followed is the same as that used in branch office 1.

```
interface GigabitEthernet0/0
 description UpstreamConnection
 ip address 10.60.0.145 255.255.0.0
 duplex auto
 speed auto
interface GigabitEthernet0/1
 ip address 10.50.10.5 255.255.255.0
 duplex auto
 speed auto
interface GigabitEthernet0/1.115

 description dataVLAN
 encapsulation dot1Q 115
 ip address 10.50.115.1 255.255.255.0
interface GigabitEthernet0/1.120
 description voiceVLAN
 encapsulation dot1Q 120
 ip address 10.50.120.1 255.255.255.0
interface GigabitEthernet0/1.200
 description managementVLAN
 encapsulation dot1Q 200
 ip address 10.50.200.1 255.255.255.0
```

To configure the Meraki MX appliance in passthrough mode, do the following:

Step 1. Navigate to the **Security Appliance > Addressing and VLANs** page.

Step 2. Select **Passthrough Mode**. More information about passthrough mode operation is available at: https://documentation.meraki.com/MX-Z/Networks_and_Routing/Passthrough_Mode_on_the_MX_Security_Appliance_and_Z1_Teleworker_Gateway.

DHCP

ISR

Enabling a DHCP server on a Cisco ISR is covered in detail in the following configuration guide: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/12-4t/dhcp-12-4t-book/config-dhcp-server.html.

The specific DHCP configuration deployed in this example is included as follows as a quick reference. Of special note is the DHCP option 150 in the voice VLAN, which defines the address of a TFTP server that hosts firmware and configuration files for VoIP phones.

```
ip dhcp excluded-address 10.10.10.1
ip dhcp pool ccp-pool
 import all
 network 10.10.10.0 255.255.255.248
```

```
 default-router 10.10.10.1
 lease 0 2
ip dhcp pool VoipPool
 network 10.50.120.0 255.255.255.0
 option 150 ip 10.50.120.1
 default-router 10.50.120.1
 dns-server 8.8.8.8 8.8.4.4
ip dhcp pool DataPool
 network 10.50.115.0 255.255.255.0
 default-router 10.50.115.1
 dns-server 8.8.8.8 8.8.4.4
ip dhcp pool ManagementPool
 network 10.50.200.0 255.255.255.0
 default-router 10.50.200.1
 dns-server 8.8.8.8 8.8.4.4
```

MX

To configure DHCP in the Cisco Meraki MX appliance, do the following:

Step 1. Navigate to the **Security Appliance > DHCP** page. Each created subnet or static route will have an entry. Ensure the **Run a DHCP server** drop-down menu is selected under the appropriate heading for each VLAN.

Step 2. In order to allow IP phones to retrieve a hosted phone firmware, add a custom DHCP option entry with the following information for the voice VLAN:

- For **Option,** choose custom to define an unlisted DHCP option.
- For **Code,** choose 150, indicating the option for a TFTP server address.
- For **Type,** choose IP because the value is an IP address.
- For **Value,** choose the address of your TFTP server.

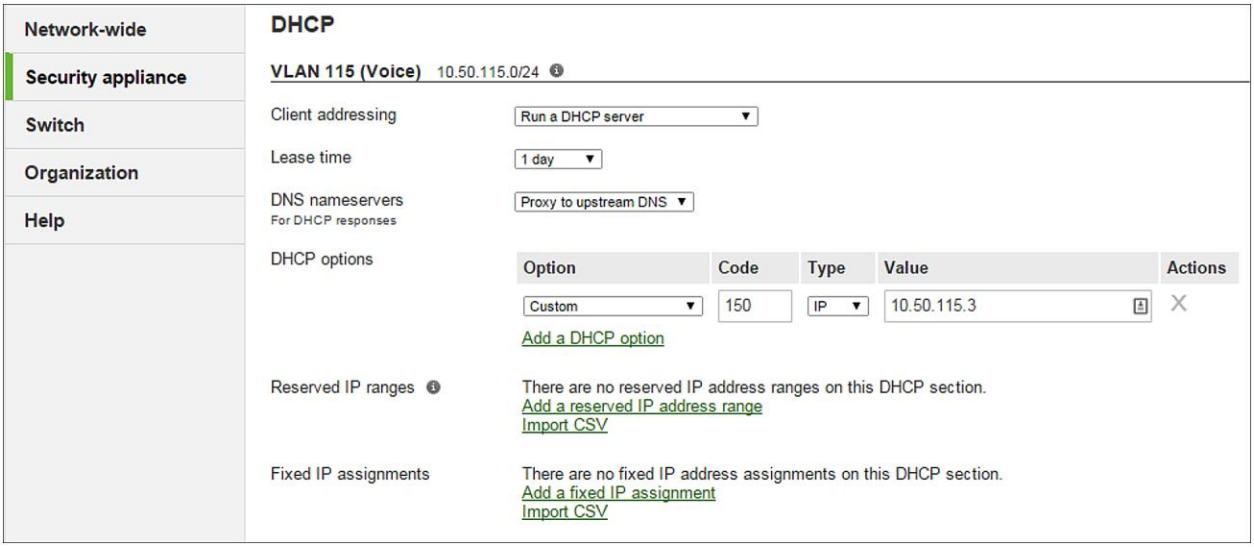

More information about configuring DHCP on the Meraki MX appliance is available at: https://documentation.meraki.com/MX-Z/DHCP/Configuring_DHCP_Services_on_the_MX_and_MS.

Switch-Port Configuration

The MS switch is used to determine the VLANs for both voice and data traffic. To configure the switch port, do the following:

Step 1. Navigate to the **Switch > Switches** page and select the device to configure.

Step 2. VLAN tags are assigned to traffic based on the configuration of the connected port. Navigate to the **Switch > Switch ports** page.

Step 3. To configure a port that will be connected to a Cisco IP Phone, select the desired port(s) and click the **Edit** button.

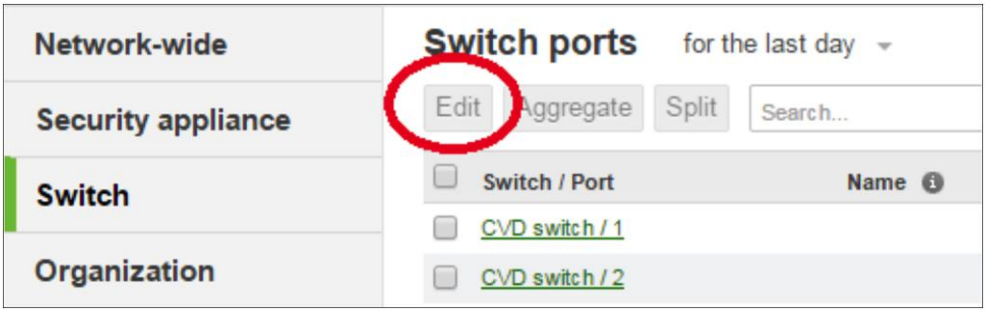

All Cisco Meraki switches support virtual stacking, which allows you to configure up to 10,000 switch ports simultaneously. Virtually stacked switches do not require a physical connection, they can be in different physical locations, and they can be of different switch models, thereby simplifying large-scale, distributed deployments. The search bar on this page supports numerous keywords to allow for the mass selection of ports to configure. The most common keyword is **port:value,** which makes selections based on the port number, but other keywords include **is:uplink** to modify uplink ports as a group, and **lldp:value** to select ports with IP phones attached. A full list of supported terms is available at: https://docs.meraki.com/display/MS/Switch+Ports#SwitchPorts-Searchingforports.

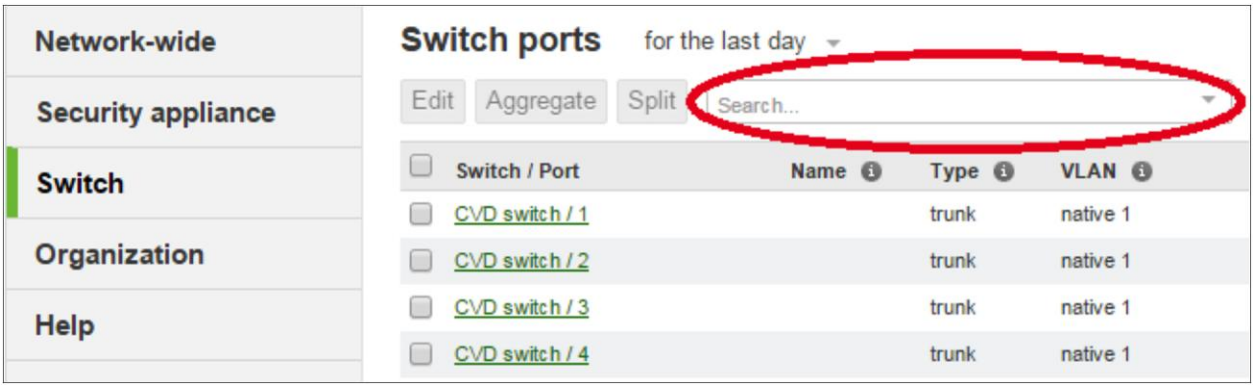

Step 4. Ensure that Power over Ethernet (PoE) is enabled, the Type is set to **access**, and the Voice and Data VLAN IDs match those created previously. The MS switch is compatible with IP phones that can receive Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED) or Cisco Discovery Protocol, and will automatically distinguish voice traffic from computer data traffic behind the IP phone into their respective VLANs.

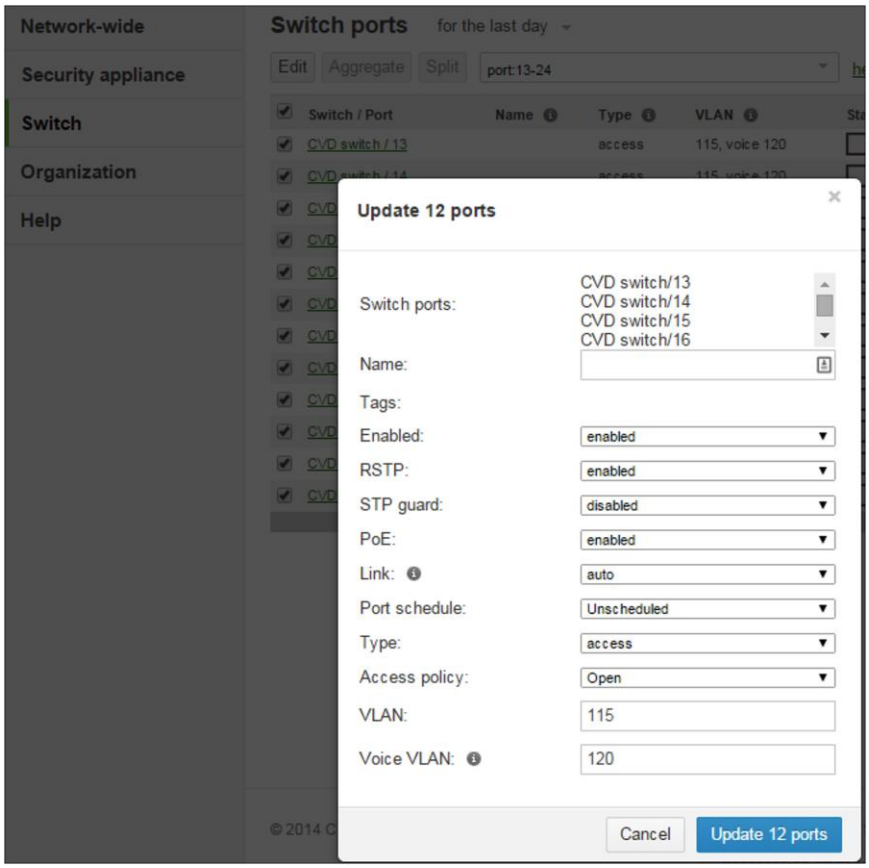

WAN and VPN

VPN to a Cisco ISR

Branch Office 1

To enable site-to-site VPN on the Cisco Meraki MX appliance, do the following:

Step 1. Navigate to the **Security Appliance > Site-to-site VPN** page.

Step 2. Enable the feature with either split-tunnel or full-tunnel mode as needed.

Step 3. Under **Topology,** determine if the appliance will connect to just the headquarters location or if a VPN mesh is desired. If a mesh is desired, select **Connect to all VPN** peers. If only a single tunnel to the headquarters is required, select **Connect directly to only one VPN peer (hub-and-spoke mode).** More information about hub-and-spoke mode is available at: https://documentation.meraki.com/MX-Z/Site-to-site_VPN/Configuring_hub-and-spoke_VPN_connections_on_the_MX_Security_Appliance.

Step 4. Each subnet and static route created on the **Security Appliance > Addressing and VLANs** page will be listed. Select **Yes** to advertise that subnet over the tunnel. If the subnet is not advertised, traffic destined toward that subnet will not be routed through the tunnel.

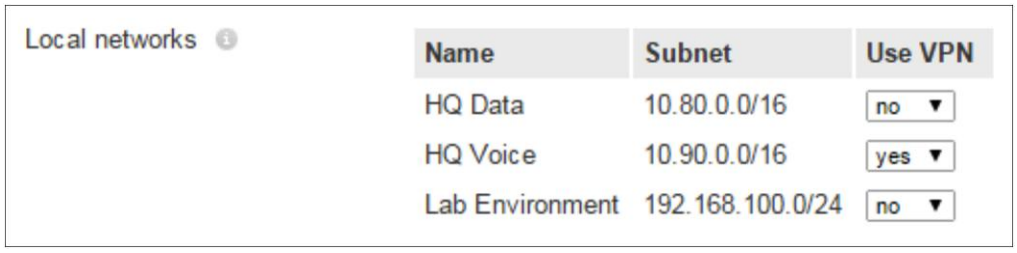

Step 5. To connect to an ASR or other device outside of the Meraki Dashboard organization, use the **Organization-wide settings** section. Add an entry for the headquarters location using the **Add a peer** button. Create the VPN tunnel using the following settings:

- **Name** is a descriptive name for the tunnel.
- **Public IP** is the publicly accessible IP address for the remote device.
- **Private Subnets** is a comma-separated list of subnets advertised by the remote device.
- **IPsec Policies** allows for customization of Internet Key Exchange (IKE) phase 1 and phase 2 options.
- **Preshared secret** is the password created for the VPN tunnel; it must match the secret set on the remote device.

Branch Office 2

Creating an IP Security (IPsec) VPN tunnel on a Cisco Integrated Services Router is covered in detail in the following document: http://www.cisco.com/c/en/us/td/docs/routers/access/1900/software/configuration/guide/Software_Configuration/Secconf1.html#pgfId-1055626.

After you have configured the router, do the following:

Step 1. Ensure the securityk9 Technology Package License is activated on the device, because you must create an Industry Standard Architecture Key Management Protocol (ISAKMP) policy. Details about activation of licensing and features is available at: http://www.cisco.com/c/en/us/td/docs/ios/csa/configuration/guide/csa_commands.html.

The configuration used in the example deployment branch office is included as follows as a quick reference:

```
configure terminal
crypto isakmp policy 10
encryption 3des
hash sha
authentication pre-share
group 2
lifetime 28800

crypto isakmp client configuration group Branch2
key CiscoMeraki
dns 8.8.8.8
domain local.cisco.com
exit
```

```
ip local pool VPNpool 10.50.130.10 10.50.120.50 mask 255.255.255.0

crypto map VPNmap isakmp authorization list Branch2
crypto map VPNmap client configuration address respond

aaa new-model
aaa authentication login Branch 2
aaa authorization network Branch2 local
username user1 password 0 pass1

crypto ipsec profile profile1
crypto ipsec transform-set transform1 esp-3des esp-sha-hmac
crypto ipsec security-association lifetime seconds 28800

interface Gi0/1.200
crypto map VPNmap
exit
exit
```

VPN to a Meraki MX Headend

Branch Office 1

With a Cisco Meraki MX appliance at either end of the connection, create the tunnel using the Cisco Meraki Auto VPN:

Step 1. Follow the steps in the section “VPN to Cisco Headquarters” to select the **Mode**, **Topology**, and advertised **Local Networks**. The Cisco Meraki Auto VPN feature will automatically establish the tunnel between the devices without further configuration.

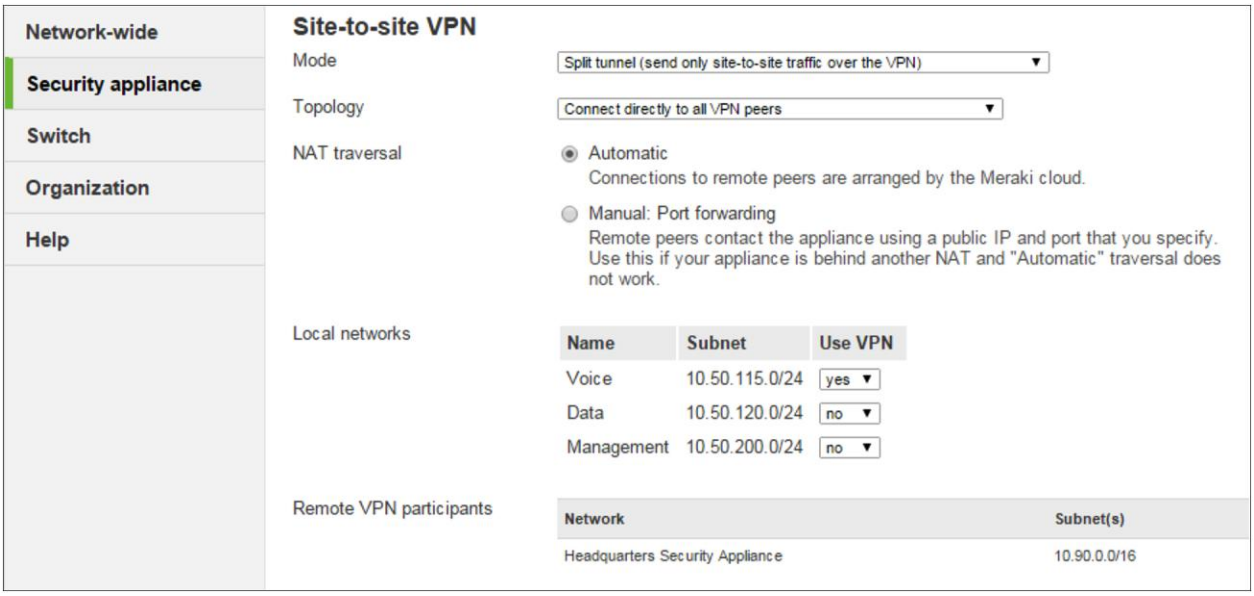

Branch Office 2

The configuration of a VPN connection in this branch office is identical to the tunnel configuration in branch office 1 in the previous section.

Wireless Considerations

Cisco Meraki MR access points provide a smoothly integrated method to expand network coverage away from the wired infrastructure, enabling connectivity and mobility throughout the physical environment. The ability to move throughout an office, warehouse, or lobby while maintaining a smooth voice call and Internet connection adds a tangible value to any business, as well as adding convenience and mobility.

Because voice and video traffic is more sensitive to interruptions and delay than many other protocols, certain concerns and considerations of the network should be stressed. Minimizing the interference of a wireless deployment is tantamount to ensuring the low-latency connection required by a live call. At the same time, dead zones and gaps in wireless coverage must be avoided to prevent unexpected disconnections. As with any wireless deployment, a thorough site survey will allow for planning of the optimal deployment pattern to maximize coverage.

A more thorough guide to wireless deployment of the Cisco Meraki MR access points can be found in the Cisco Meraki Branch-Office Deployment Guide, as well as in resources in the Cisco Meraki Knowledge Base.

Cisco Unified Communications Manager

Cisco Unified Communications Manager is a centralized call-management solution that cancels the requirement for a local telephony service at each branch-office location. The needs of a Cisco Unified Communications Manager deployment hinge upon whether a local Cisco Unified Communications Manager Express service is implemented at a given branch office. In branch office 2, it is assumed that Cisco Unified Communications Manager is the implemented call-management solution and it is not coupled with a local Cisco Unified Communications Manager Express service.

Although it is possible to build a multiple-site meshed VoIP network, this process is more complicated, and out of scope of this document. Local public-switched-telephone-network (PSTN) services, when available, provide for a simpler deployment.

There exist a large number of customization options to help ensure that Cisco Unified Communications Manager serves the needs of any particular deployment. Details about these options, as well as a walkthrough for a basic configuration, are available in the Cisco Unified Communications Manager Administration Guide at: http://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/10_0_1/ccmcfg/CUCM_BK_C95ABA82_00_admin-guide-100.html.

Cisco Unified Communications Manager Express

Cisco Unified Communications Manager Express, as previously stated, has very similar network quality-of-service (QoS) requirements to a centralized call-control deployment model, such as Cisco Unified Communications Manager. These requirements depend on the business model used for IP communications and the level of coupling between branch-office locations from an IP telephony perspective.

In the situation where the business model necessitates the use of the WAN for transporting VoIP traffic between branch offices, an appropriate QoS model needs to be implemented.

The additional consideration in a Cisco Unified Communications Manager Express environment is Call Admission Control (CAC). Although in a centralized call-control deployment the number of concurrent VoIP calls into and out of a branch office can be accurately controlled, the same is not quite true for a Cisco Unified Communications Manager Express environment.

Cisco Unified Cisco Unified Communications Manager Express allows for outbound CAC, but a call-control (Cisco Unified Communications Manager Express) entity at one branch-office location cannot on its own regulate the number of concurrent calls from other branch offices. This level of control requires the assistance of a Cisco Unified Border Element control function, which may not be deployed in many Cisco Unified Communications Manager Express environments.

You have a large number of customization options to help ensure that Cisco Unified Communications Manager Express serves the needs of any particular deployment. Details about these options, as well as a walkthrough for a basic configuration, can be found in the Cisco Unified Communications Manager Express Administration Guide.

QOS

Introduction

Every organization relies on its communications network to function, and disruption or congestion of these communications can slow productivity to a crawl. QoS features account for the importance and priority of various traffic and help ensure that the most urgent and sensitive traffic is given the priority it requires. Voice and video traffic requires low latency to present a proper user experience, and QoS helps ensure that this latency is achieved even while the network connection is under a heavy load.

IP networks must provide a baseline level of service and functions to reliably support voice in the form of IP telephony and VoIP. This baseline level includes security, reliability and predictability, instrumentation, and some level of delivery guarantee.

Network administrators and architects achieve this service level by managing delay, delay variation (jitter), bandwidth provisioning, and packet-loss parameters with QoS techniques.

Elements of QoS

QoS features essentially involve the use of one or more of three core concepts: identification and classification, queuing and scheduling, and traffic shaping and policing. Traffic identification and classification is the act of determining how a particular traffic type or flow should be handled. Rather than performing this determination at every hop in the traffic path, it makes sense to mark the packets of a specific flow with a classification tag at the furthest point from the core of the network, so that the traffic can be treated appropriately by subsequent elements in the network, based on this classification tag. This process offloads the required processing to edge switches, thereby distributing resources more efficiently and helping ensure that traffic is handled on a consistent QoS path from end to end.

This marking is typically done in one of two ways: by using either the 3-bit class-of-service (CoS) portion of the 802.1Q tag of a Layer 2 frame or the 6-bit differentiated-services-code-point (DSCP) tag in the IP address header of a packet. The Layer 2 CoS tag supports up to eight classes of traffic, whereas the Layer 3 DSCP tag supports up to 64 classes.

Queuing and scheduling refers to the mechanisms a given device uses to manage a congested link. These mechanisms determine which traffic is sent first, and which traffic is queued and delayed. Examples of queuing are the priority buckets used with the Cisco Meraki MX appliance. Queuing typically occurs at the gateway or primary router of a given network, although it can be applied at aggregation switches in larger deployments.

Traffic shaping and policing is a broad area that covers the management of traffic flows to limit the amount of bandwidth consumed. Traffic shaping uses queuing and buffering capabilities of the switch or router to help ensure that a particular traffic type or flow conforms to the configured bandwidth limit. Unlike policing, where excess traffic is discarded when a configured bandwidth threshold is reached, shaping buffers the excess traffic for a period of time, with the intent of sending it when the link is less congested, thus using the available bandwidth more cost-effectively.

Elements of QoS at the Switch Level

In a standard deployment, access switches are the devices connected directly to phones and workstations. Their role is to identify incoming traffic as sourced from an IP phone or the machine behind it, mark that traffic with appropriate VLAN tags, and handle that traffic in accordance with its level of QoS.

Cisco Meraki MS switches support LLDP-MED and Cisco Discovery Protocol for the identification of voice and data traffic sourced from a compatible IP phone. With either protocol, the phone sends out an initial identification packet, and the MS switch responds with the voice and data VLAN information configured for the connected port. Using these protocols, traffic sourced from the PC and traffic sourced from the IP phone itself is placed into the appropriate separate VLANs.

When traffic is properly identified, it can be marked with the VLAN ID in the packet header. Most VoIP deployments connect IP phones to access ports, so traffic is handled on the associated VLAN but the packet does not actually carry a tag. More information about this port behavior is available at: https://documentation.meraki.com/zGeneral_Administration/Tools_and_Troubleshooting/Fundamentals_of_802.1Q_VLAN_tagging.

The switch must be configured to handle QoS for the incoming traffic on the **Switch > Switch settings** page under the heading **Quality of service**. Here, traffic can be tagged with DSCP bits based on VLAN and protocol, or configured to trust incoming DSCP tags. Best-practice call-manager deployments with both Cisco Unified Communications Manager and Cisco Unified Communications Manager Express provide phones with detailed configuration files, applying DSCP tags on egress of the phone. For the voice VLAN configuration, this process results in simply trusting the incoming tags.

If desired, the mapping from DSCP tag to CoS queue can be modified in the Meraki Dashboard:

Step 1. Navigate again to the **Switch > Switch settings** page and click **Edit DSCP to CoS map**. You can modify these values to fit the needs of your specific QoS deployment.

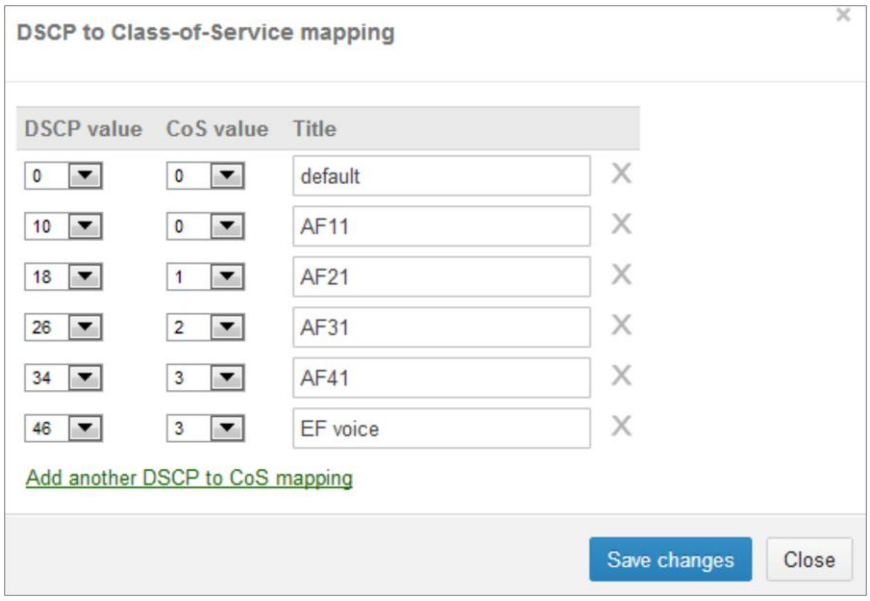

QoS on the MX Security Appliance

In both branch-office deployments explored previously, the Cisco Meraki MX appliance provides traffic-shaping functions to further expand QoS capabilities. All of these settings can be found on the **Security Appliance > Traffic Shaping** page, or alternatively can be configured as part of policies applied to a subset of users. Policy application is covered in further detail later in this document.

When creating a traffic-shaping rule, select the traffic definitions that you want to shape. They can be individual traffic types, such Session Initiation Protocol (SIP), Skinny Client Control Protocol (SCCP), or Skype, or they can be an entire category such as VoIP and video conferencing. You also can create rules to apply to traffic sourced or destined for a particular subnet.

With the traffic definition selected, click the **Bandwidth Limit** drop-down menu and determine if this traffic will follow the network traffic limit, follow a custom traffic limit, or ignore any limits and use bandwidth restriction-free. For video conferencing it is a good practice to ignore bandwidth limits, but voice traffic is generally low bandwidth enough to make it unnecessary.

The **Priority** field allows you to guarantee a fraction of the uplink for each priority level, providing a method to avoid congestion for the most important traffic when the uplink is fully used. Twice as many packets are sent from the normal queue than the low queue at any given time, and the high queue sends twice as many as the normal. For details about the implementation of this feature, please refer to: https://documentation.meraki.com/MX-Z/Firewall_and_Traffic_Shaping/Using_Packet_Prioritization_on_a_Traffic_Shaping_Rule.

The MX appliance can tag traffic with DSCP tags. Traffic in this deployment is tagged before it reaches the MX appliance, so the default **DSCP Tagging** option of Do not set DSCP tag is desired.

The **Global Bandwidth Limit** section allows for the configuration of a maximum throughput on a per-device basis. These limits apply to outbound and inter-VLAN traffic. The **SpeedBurst** feature allows a client to exceed the maximum bandwidth by four times for the first 5 seconds of a given flow. This feature allows users a lower perceived latency because of the rapid downloading of small files.

The MX appliance allows for the use of two Internet uplinks, allowing for an aggregation of lines for an increase in throughput, redundancy in case of failure, or physical segregation of traffic. In the case of MX devices without a dedicated secondary WAN port, the first LAN port can be toggled between LAN and Internet usage through the local status page. Details about this local configuration page are available at: https://docs.meraki.com/display/MX/MX+Local+Status+and+Configuration.

By default, a secondary uplink is used only as a failover. With link aggregation enabled, traffic is shared proportionally across the two connections. The **Internet 1** and **Internet 2** sliders control the limit for both uplinks and the ratio between the two determines the proportions in which traffic is shared. Note that a cellular connection is used only as a failover connection, and does not aggregate with wired uplinks. Also note that Auto-VPN connections and management traffic to the Meraki cloud always use the **primary uplink** regardless of uplink preferences.

Creating an **uplink preference** forces designated traffic over a particular uplink as long as that uplink retains connectivity. A common configuration in deployments with limited bandwidth is to dedicate a stable connection such as a T1 entirely to VoIP traffic, and send all other traffic through a different uplink. To configure this process, disable link aggregation and create a preference for traffic in the voice VLAN to use the secondary uplink.

Note:

- **Protocol** defines whether this configuration applies to TCP, User Datagram Protocol (UDP), or both.
- **Local IP range** is the source address of traffic to be affected.
- **Local port** is the source port of traffic to be affected. Note that many protocols use an ephemeral source port, and the “Any” keyword can be used.
- **WAN IP Range** is the destination address of traffic to be affected.
- **WAN port** is the destination port of traffic to be affected.
- **Preferred uplink** defines which uplink traffic matching these properties uses.

QoS on the ISR

The richness of the Cisco QoS toolset inevitably increases its deployment complexity. To address customer demand for simplification of QoS deployment, Cisco has developed the Automatic QoS (AutoQoS) features.

AutoQoS is an intelligent macro that allows an administrator to enter one or two simple AutoQoS commands to enable all the appropriate features for the recommended QoS settings for an application on a specific interface.

AutoQoS VoIP, the first release of AutoQoS, provides best-practice QoS designs for VoIP on Cisco Catalyst® switches and Cisco IOS® Software routers.

By entering one global and/or one interface command, depending on the platform, the AutoQoS VoIP macro expands these commands into the recommended VoIP QoS configurations (complete with all the calculated parameters and settings) for the platform and interface on which the AutoQoS is applied.

For Cisco IOS Software routers, AutoQoS automatically performs the following tasks:

- Classifies and marks VoIP bearer traffic (to DSCP Expedited Forwarding [DSCP EF]) and call-signaling traffic (to DSCP Class Selector [CS]3). It applies scheduling:
  - Low Latency Queuing (LLQ) for voice
  - Class-Based Weighted Fair Queuing (CBWFQ) for call signaling
  - Fair Queuing (FQ) for all other traffic
- Enables Frame Relay Traffic Shaping (FRTS) with optimal parameters, if required
- Enables link fragmentation and interleaving (LFI), either multilink-point (MLP) LFI or FRF.12, on slow (768-kbps) links, if required
- Enables IP Real-Time Transport Protocol (RTP) header compression (cRTP), if required
- Provides Remote Monitoring (RMON) alerts of dropped VoIP packets

AutoQoS VoIP became available on Cisco IOS Software router platforms in Cisco IOS Software Release 12.2(15)T.

In its second release, for Cisco IOS Software routers only, **AutoQoS Enterprise** detects and provisions for up to 10 classes of traffic, including the following:

- Voice
- Interactive video
- Streaming video
- Call signaling
- Transactional data
- Bulk data
- Routing
- Network management
- Best effort
- Scavenger

The AutoQoS Enterprise feature consists of two configuration phases, completed in the following order:

- Auto Discovery (data collection): This phase uses Network-Based Application Recognition (NBAR)-based protocol discovery to detect the applications on the network and performs statistical analysis on the network traffic.
- AutoQoS template generation and installation: This phase generates templates from the data collected during the Auto Discovery phase and installs the templates on the interface. These templates are then used as the basis for creating the class maps and policy maps for your network. After the class maps and policy maps are created, they are then installed on the interface.

AutoQoS Enterprise became available on Cisco routers in Cisco IOS Software Release 12.3(7)T.

Although the two AutoQoS macro commands offer a very straightforward way of implementing a relatively complex QoS framework and are well suited to the branch-office deployments discussed in this document, a wealth of additional information is available if you want to gain a better understanding of this topic.

NBAR is a classification engine that allows the Cisco ISR to recognize a variety of web-based applications. When classified, the device can guarantee bandwidth for critical services, drop extraneous packets, or tag traffic for handling elsewhere in the network. This option can help ensure that critical and sensitive services such as database queries or communication with networking equipment are given the priority needed to perform.

For a complete list of supported protocols and methods to handle traffic, please visit: http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/network-based-application-recognition-nbar/prod_case_study09186a00800ad0ca.html.

Call Admission Control

Call Admission Control (CAC) is a concept that prevents degradation of voice calls due to oversubscription of the connection used to carry calls. CAC preserves the quality of current calls in case of congestion by disallowing additional calls from completing new connections. For example, if a given link can support only two concurrent calls and a third call is initiated, CAC denies the third call with a reorder tone instead of allowing the third call to degrade the experience of all three lines.

Numerous methods for implementing CAC on a Cisco ISR are available. More information about these methods is available at:

- http://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/voip_solutions/CAC.html.
- http://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/collab09/clb09/cac.html.

Security

Introduction and Overview

Security is a major concern for any organization, and a deployment across multiple branch offices or physical locations only adds complexity. Traditional concerns such as isolating sensitive traffic or preventing access to malicious resources is made more difficult when communications over the Internet are an integral part of daily business functions. Simply creating a tunnel back to a headquarters location greatly increases the attack surface of a given network, and the proliferation of cloud-based storage and collaboration means that more traffic than ever relies on a safe connection to the web.

Goals of Security Features

Although a wide variety of approaches and implementations to security is available, any security-oriented feature basically means controlling one of two things: access or content. Access control determines which users are allowed to use which resources, and the methods of preventing users from bypassing these segregations. Access control includes features such as VLANs or port security. Content control determines what resources should be allowed onto a network, including category filtering or intrusion detection and prevention.

Elements of Access Control

MS

As an access switch, the first and most basic line of defense in a secure deployment is provisioning ports with the correct VLAN information. The purpose of creating VLANs is to segregate traffic, and incorrectly tagged traffic can be considered VLAN hopping, a method used to access traffic in a VLAN that would not normally be accessible. More information about VLAN hopping is available at: https://documentation.meraki.com/zGeneral_Administration/Tools_and_Troubleshooting/Fundamentals_of_802.1Q_VLAN_tagging.

To prevent VLAN hopping, proper port configuration is important. Ports that will have devices such as a PC or IP phone connected to them are best configured as access ports, helping ensure that all incoming traffic from that device is handled in the appropriate VLAN. Aggregation switches, or switches that have other switches connected as opposed to end devices, are typically configured with trunk ports. A trunk port limits what VLANs are able to pass traffic through using a list of **Allowed VLANs**. To configure this feature, do the following:

Step 1. Navigate to the **Switch > Switch ports** page.

Step 2. Select the ports to be modified, and click **Edit**.

Step 3. In the **Allowed VLANs** section, list the VLAN IDs that should be permitted to pass through the selected ports.

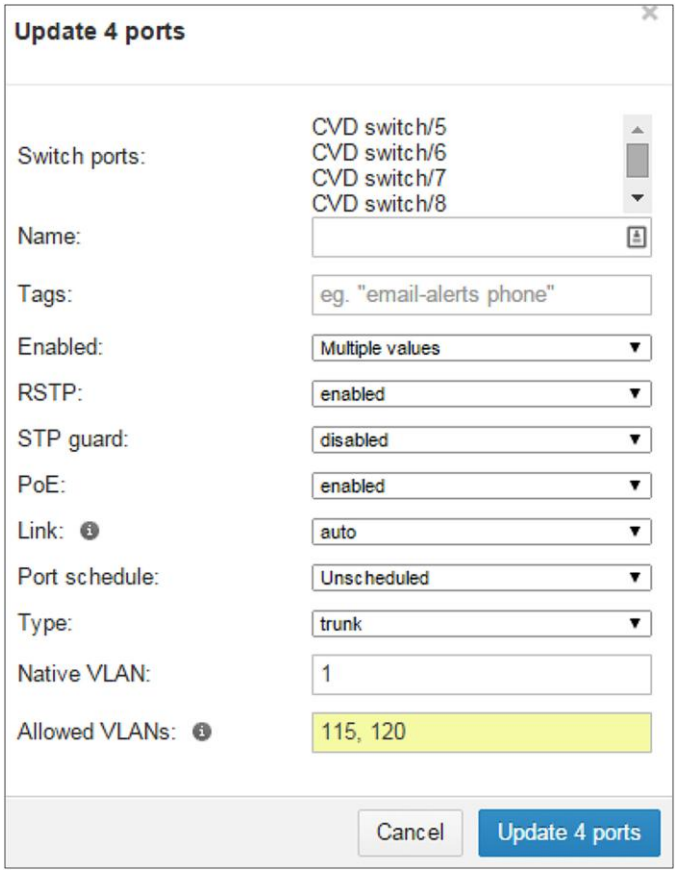

To take this idea of port control a step further, Cisco Meraki MS switches have an 802.1x-based access policy feature for integration with RADIUS servers. The MS switches support not only RADIUS authentication but also accounting, allowing the tracking of additional information about the client connection.

To create an access policy, do the following:

Step 4. Navigate to the **Switch > Access Policies** page.

Step 5. Click **Add an access policy**. Define the following fields:

- **Name** is a descriptive name of the policy.
- **RADIUS servers** is a list of servers to attempt.
- **Host** is the IP address of the server.

- **Port** is the port of the server. Note: The RADIUS default is port 1812.
- **Secret** is the secret key created on the server.
- **RADIUS testing** determines if the Meraki devices proactively check to determine whether the RADIUS server is up.
- **Guest VLAN** determines which VLAN unauthenticated users are placed in. This field can be left blank to prevent guest access.
- **Voice VLAN clients** determines if traffic on the voice VLAN must authenticate as well. If set to bypass, only traffic on the data VLAN will require authentication. This applies to devices connected behind an IP phone using LLDP-MED to assign traffic to VLANs.

Step 6. An access policy is not effective until it is applied to a port. On the **Switch > Switch ports** page, select the desired ports.

Step 7. Click **Edit**, and select the access policy by name.

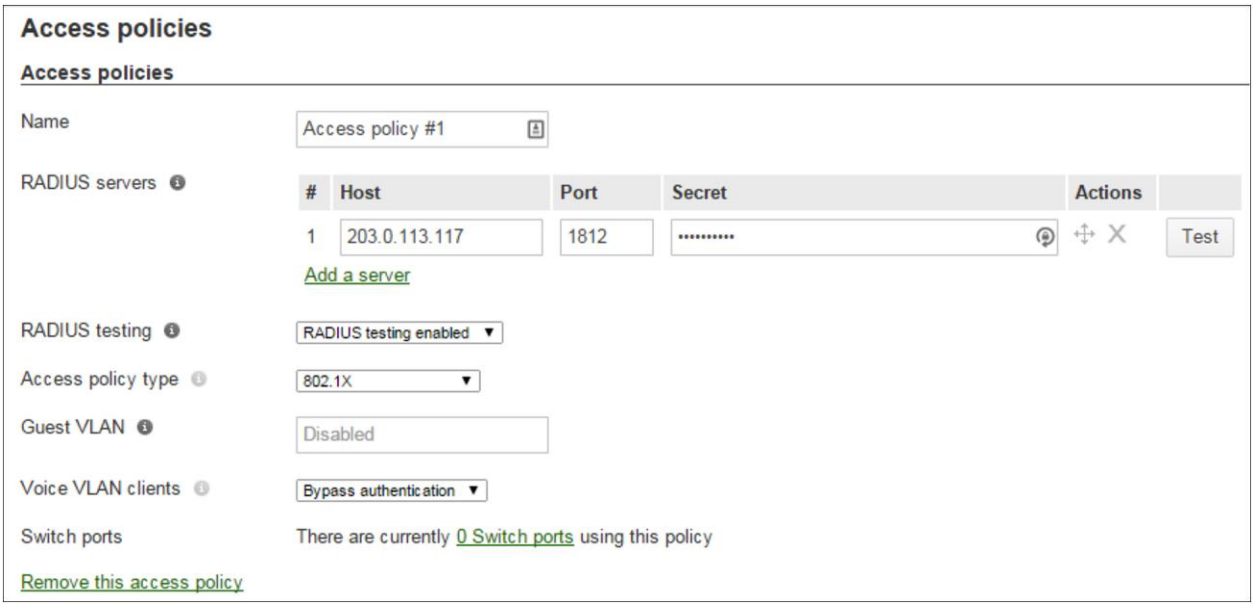

In addition to access policies, a port can be restricted to pass only traffic sourced from specifically whitelisted MAC addresses. This restriction can be configured in the same port configuration menu on the **Switch > Switch ports** page, under **Access Policies**. MAC whitelisting is an effective method of control for smaller deployments, but as the size of a deployment grows, the effort to maintain a whitelist will grow as well. Sticky MAC is an access policy that allows for dynamic whitelisting. By configuring a value for the number of devices to track, the switch learns the MAC address of devices, up to the limit of the configured value, allowing for easy whitelisting of devices that should not move or change often, such as IP phones.

For a more restrictive implementation, port isolation is an available feature that allows an administrator to prevent traffic being sent between specific ports. This feature can be used in conjunction with other VLAN configuration settings, allowing restriction of traffic even within the same subnet. Additional information about this feature is available at: https://documentation.meraki.com/MS/Port_and_VLAN_Configuration/Restricting_Traffic_with_Isolated_Switch_Ports.

Port scheduling enables and disables configured ports based on the local time. You can use this feature to prevent devices from connecting outside of business hours; disable conference room ports unless a need is scheduled or power down PoE phones, cameras, or access points. To configure this feature do the following:

Step 1. Navigate to the **Switch > Port schedules** page.

Step 2. Click **Add a new port schedule** to expand the scheduling section, and define a schedule using the premade template links or configuring each slider to the desired start and end time.

Step 3. Verify the local time zone is set correctly using the **Network-Wide > General** page, particularly when deploying multiple networks.

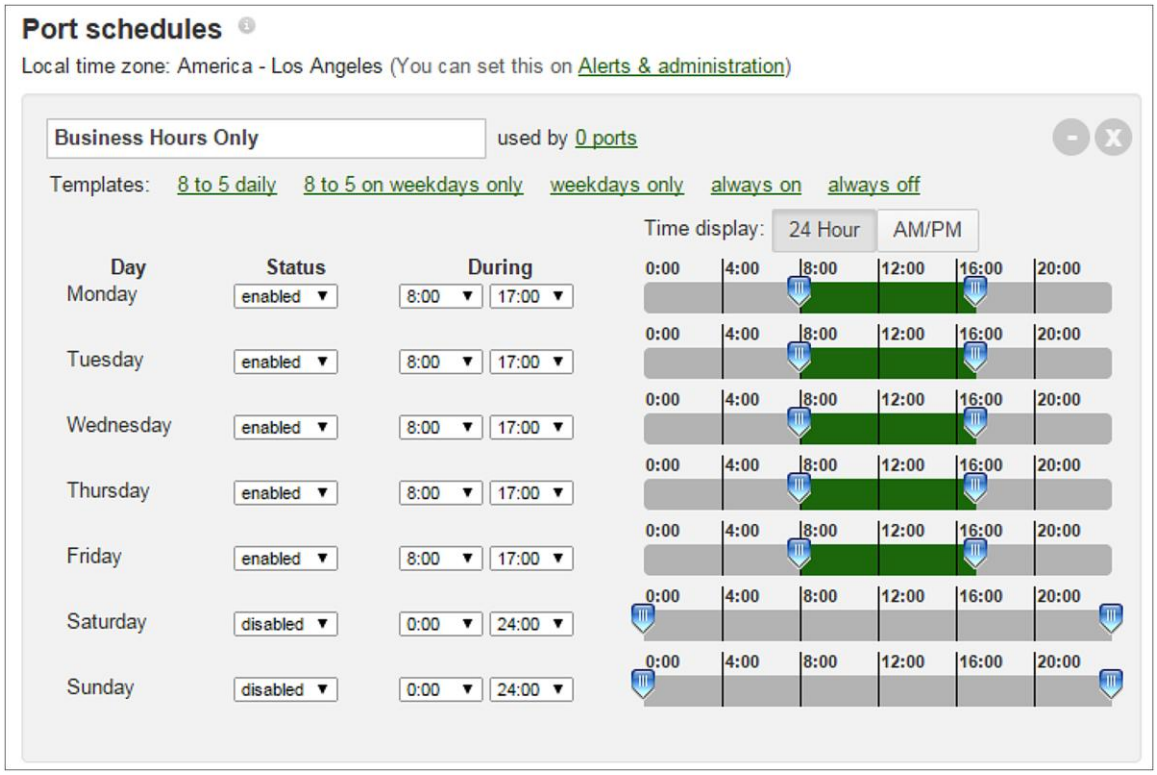

Although DHCP is an extremely useful and common protocol to deploy in nearly any topology, the risk of a rogue DHCP server exists and should be mitigated. Similar to VLAN hopping, a device being assigned an address in an incorrect VLAN can break configurations, or in a worst-case scenario, direct sensitive traffic outside protected networks. The Cisco Meraki MS switch monitors DHCP servers on a network, and gives administrators the ability to block the IP assignment of some or all detected servers, or whitelist only known servers.

To mitigate rogue DHCP servers, follow the steps below:

Step 1. Navigate to the **Switch > DHCP Servers** page. Listed is every DHCP server that has been recently detected, and whether it is allowed or blocked.

Step 2. Under the **Default DHCP server policy** section, select whether servers will be allowed or denied by default. Most branch-office deployments have a single DHCP server, or a primary and secondary server, so the safer option is to block DHCP servers by default, and allow only known servers by MAC address.

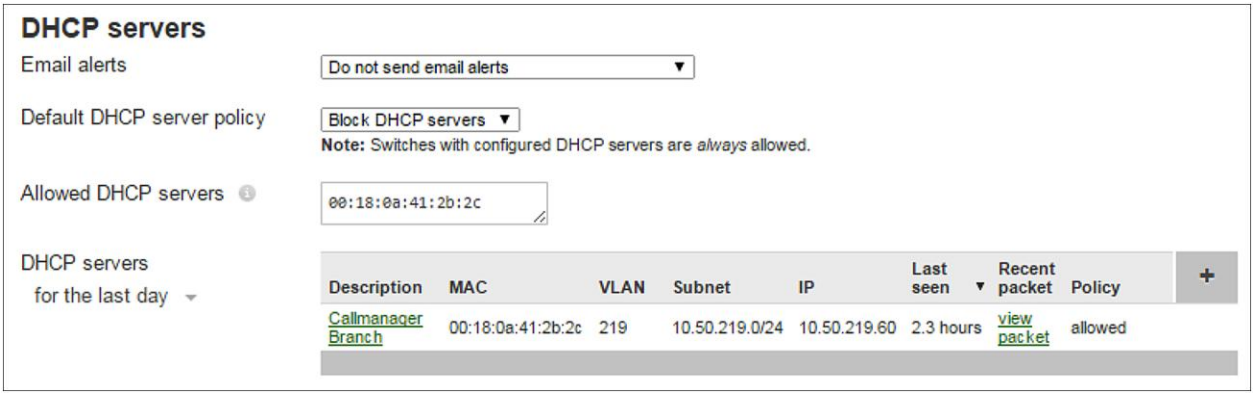

MX

The Cisco Meraki MX appliance can act as a Layer 3 firewall in either NAT or passthrough mode. The firewall is one of the most important components of a secure deployment, and as such, should be tailored to the needs of a given deployment. Refer to the following webpage for more information about implementing firewall rules on the Cisco Meraki MX appliance: https://docs.meraki.com/display/MX/Firewall+settings.

For best practices regarding the configuration of firewall rules, please visit: http://www.cisco.com/web/about/security/intelligence/firewall-best-practices.html.

ISR

The Cisco Integrated Services Router offers many features to control access to the network securely. As the gateway device, access control lists (ACLs) are the most fundamental stage of network control. ACLs define types, direction, source, and destination of traffic, and allow or deny access based on these factors. This type of control relies heavily on the content and location of resources on the network, and should be implemented with considerations to the individual deployment.

For more information about the types of ACLs and methods of implementation, please visit: http://www.cisco.com/c/en/us/support/docs/security/ios-firewall/23602-confaccesslists.html.

Elements of Content Control

MX

In addition to acting as a Layer 3 firewall, the Cisco Meraki MX appliance can also act as a Layer 7 firewall, classifying traffic into numerous categories and restricting access based on those categories. This firewall activity can include blocking entire categories, such as sports or gaming, or blocking traffic from certain identifiable applications such as Skype, iTunes, or Dropbox. In conjunction with bandwidth limits, the Layer 7 firewall is a powerful tool for bandwidth-restricted deployment, and is more effective at blocking web resources where IP addresses can change. This tool is also effective against encrypted traffic, in which the MX examines the domain in the Secure HTTP (HTTPS) certificate to determine blocking status.

The Cisco Meraki MX appliance contains a powerful set of security features, broken into malware detection and IDS or IPS. You can configure these features on the **Security Appliance > Security Filtering** page. Please note, security filtering features require an Advanced Security license.

The malware-detection feature, based on Kaspersky Safestream, screens incoming and outgoing HTTP traffic against the URL or by matching signatures. If a match is made, the suspicious file or traffic to that URL is blocked, and the incident is logged to the event log. Sometimes this feature detects benign traffic, or an administrator may wish to allow access to a resource that triggers the detection in spite of the classification. Whitelists for both URLs and signature IDs can be created through the Meraki Dashboard. Because updating and matching signatures is a dynamic process, if you suspect that a resource is being flagged incorrectly, please contact Meraki Support.

Intrusion detection and prevention, powered by the Sourcefire® Snort® Engine, feeds all traffic flowing between VLANs or between the LAN and WAN interfaces through that engine. The Sourcefire Snort engine compares the traffic against a set of malicious traffic signatures, and blocks that traffic if it detects a match. These signatures represent exploits classified by the Common Vulnerability Scoring System (CVSS). The **Ruleset** field determines how restrictive the security filtering is. Details as to which categories are included in each ruleset are available at: https://docs.meraki.com/display/MX/Security+filtering.

The security report is a compilation of all events detected or blocked by the IDS/IPS feature; it is located on the **Security Appliance > Security Reporting** page. This page graphically represents the quantity and severity of events, and allows for the filtering of certain events, hosts, protocols, and other details. The bar graph on the page details the severity and quantity of events within the selected time period, and the pie chart illustrates more detailed information from that time period, such as source address or affected client. The “Event Log and Syslog” section later in this document details individual events, and the mechanism for sorting through events in a more granular fashion. Signatures are listed by name when available, or by the associated rule ID. These rule IDs are used to further research the nature of a given vulnerability.

The content-filtering feature of a Cisco Meraki MX appliance uses the Brightcloud Web Classification Service to allow administrators to restrict access to entire classes of websites. This feature makes it simple to restrict access to categories such as adult and pornography, academic cheating, sports, or sites known to contain spyware. Content filtering can help ensure that sites likely to cause harm are blocked, and can help keep employees or students on task. To configure this feature, do the following:

Step 1. Navigate to the **Security Appliance > Content Filtering** page.

Step 2. The **URL category list size** option dictates the behavior of the MX appliance when observing HTTP traffic.

If you select “Top sites only”, the MX appliance will download a list of sites for the selected categories and compare visited URLs with that list locally. Because this comparison happens local to the appliance, this option yields a better initial performance compared against the “Full List” option.

If you select “Full list”, any URL requested is checked against a cloud-hosted database for comparison. The MX appliance then stores that result locally, to prevent checking the same URL twice. Because this comparison happens on an external database, a small amount of latency is introduced before the initial HTTP request is completed. As URLs are categorized and the local cache is filled, this effect is minimized. Because the Cisco Meraki MX60 and MX64 appliances do not have a hard drive, only the “Full list” option is available on those platforms.

Web-search filtering is an option that automatically applies adult-content filtering to searches made in supported search engines such as Google, Yahoo, and Bing. This feature takes advantage of the filtering within the search engine itself to help prevent adult content from being accessible. Because this feature applies only to unencrypted searches over the HTTP protocol, the option exists to block HTTPS searches. Enabling these two features in tandem helps ensure that users’ search terms are filtered using the filter that the search engine provides.

URL blocking is a feature that simply allows administrators to allow or deny access to specific resources. These rules take priority over rules configured elsewhere in Meraki Dashboard, such that a URL allowed here but blocked by a content-filtering category and a Layer 7 firewall rule is allowed through. A more detailed explanation of how these URL rules are processed is available on the **Security Appliance > Content Filtering** page.

ISR

Cisco IOS Content Filtering is a category-based URL classification engine that uses Trend Micro's TrendLabs threat database. By querying the threat database when HTTP traffic is detected, the device can determine in which category the destination resource is classified, and allow or disallow the request as needed. This process helps prevent users from accessing sites that are considered dangerous, illegal, or inappropriate.

Group Policies

Group policies allow administrators to define a set of rules, exceptions, and settings and apply this collection of configurations to individuals or groups of users. These policies allow a more granular level of control. Settings applied by group policies override settings applied by that feature elsewhere. Therefore, a popular policy strategy is to create a restrictive set of default rules and apply less-restrictive group policies to authenticated users. To enable group policies, do the following:

Step 1. Navigate to the **Network-wide > Group Policies** page and select an existing policy to edit, or click **Add a group** to create a new one.

Step 2. The configuration page for a given policy shows a large selection of features that you can use to build a profile with the desired behavior. Many options on this page allow for the policy to supplant the network default settings, but you can configure others, such as content-filtering categories, to append to the network settings.

More details about configuring this feature are available at: https://documentation.meraki.com/MR/Group_Policies_and_Blacklisting/Creating_and_Applying_Group_Policies.

Troubleshooting

This section addresses troubleshooting specifically for our example deployment, and troubleshooting that references that topology.

Detailed information about troubleshooting problems that commonly occur on Cisco Meraki devices is available at the Cisco Meraki knowledge base at: https://documentation.meraki.com.

Troubleshooting guides for Cisco products are available at: http://www.cisco.com/c/en/us/support.

One of the most common problems in a VoIP deployment is for one or more phones to fail to register with the telephony service. With the majority of IP phones working, proving that the overall configuration of the network is correct, you have numerous options to troubleshoot the problem.

If an IP phone relies on PoE for connectivity but does not power on, do the following:

Step 1. Navigate to the connected port on the Cisco Meraki MS switch using the **Switch > Switch ports** page.

Step 2. Under the Configuration section, verify that the port is configured correctly, as an access port, and with the correct VLAN and voice VLAN IDs.

Step 3. In this same section, ensure PoE is not disabled. If power negotiation has occurred properly, the PoE entry will have a current power draw as well as an advertised power draw, as shown in the screen shot. If an advertised power is not shown, the IP phone in question may be damaged, or the cable between the phone and the switch may be damaged.

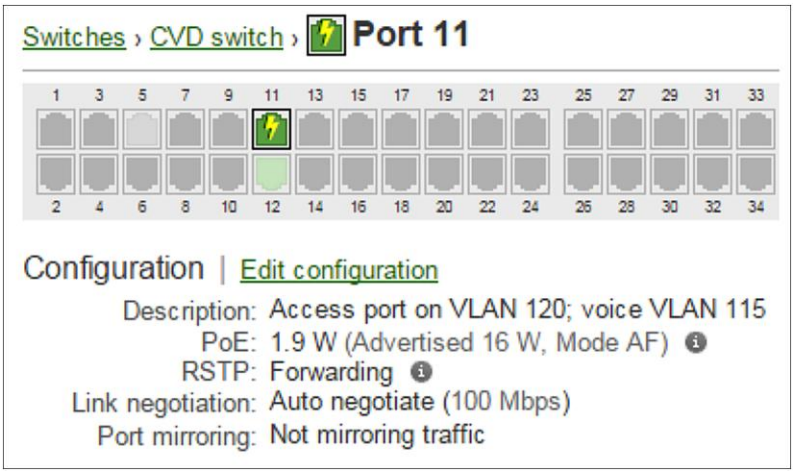

Step 4. If cable damage is possible, the Cisco Meraki MS switches feature a cable test tool that you can run from the dashboard. This test can help to confirm suspicions of cable failure. For details about the cable test tool, please visit: https://documentation.meraki.com/MS/Monitoring_and_Reporting/Using_the_Cable_test_Live_tool.

Step 5. On the port details page of a given port, a connected IP phone should show as a client, complete with name, IP address, and MAC address. Ensure this information is accurate. The VLAN should match the voice VLAN configured previously, and the IP address should be a valid address in the range configured on the DHCP server.

Step 6. To view active DHCP leases on a Cisco Meraki MX appliance, navigate to the **Security Appliance > Appliance status page**, and use the DHCP leases Live Tool.

Step 7. To view active leases on a Cisco ISR, use the command **`show ip dhcp binding`**. Additionally, if a device has an IP address, it will appear as a client on the **Network-wide > Clients** page.

Because the TFTP option is configured on the DHCP server, an IP phone that has successfully obtained an IP address will attempt to reach out to the TFTP server and download the matching firmware files. In branch office 2, the TFTP server is hosted on the same device as the DHCP server, so connectivity is established inherently. In branch office 1, the TFTP server is located at the headquarters location. Ensure the VPN tunnel is established by navigating to the **Security Appliance > VPN Status** page. Across the top of the page should be four status indicators indicating the health of the MX and VPN connectivity.

Because of the UDP hole-punching method Cisco Meraki devices use to establish a VPN tunnel, certain NAT configurations can interfere with VPN connections. It is possible that these configurations can exist on an upstream Internet service provider's network as well. For details about troubleshooting the Automatic NAT Transversal, please visit: https://documentation.meraki.com/MX-Z/Site-to-site_VPN/Troubleshooting_Automatic_NAT_Traversal_for_Meraki_Auto-VPN.

The following section lists the Meraki Dashboard networks this device is configured to establish tunnels to, as well as the advertised subnets. This page does not currently indicate the health of VPN connections with networks without Cisco Meraki devices. For details about troubleshooting VPN connections with a Cisco or other third party device, please visit: https://documentation.meraki.com/MX-Z/Site-to-site_VPN/Troubleshooting_3rd_Party_Site-to-site_VPN.

Step 8. When connectivity to the TFTP server is established, the server sends the default configuration file and the specific configuration file of the phone. Ensure these files exist and are shared on the TFTP server. Finally, the telephony service will complete registration of the phone.

Step 9. If the phone is able to download all these files from the TFTP server, it should reboot and load the correct firmware file. Verify that the firmware loaded on the IP phone is the same firmware that should be running, and that the contents of the firmware file match.

Step 10. Ensure that the telephony service is configured for autoregistration, and that sufficient extensions remain in the pool. Run the **`auto-assign`** command and ensure the phone type is specified. If autoregistration is not enabled, verify the MAC address of the phone and ensure the SIP<MAC>.xml file has been created by running the **`show voice register tftp-bind`** command. If the correct file does not appear in the output of this command, the profile was likely not created. Create the profile using the **`create profile`** command from the voice register global prompt, or recreate the configuration files using the **`create-cnf command`**.

If the IP phone does not register with the telephony service at this point, please contact technical support for further assistance.

Event Log and Syslog

All Cisco Meraki devices use an event log, reachable by navigating to the **Network-Wide > Event Log** page. This page tracks and collects the events, errors, and notifications for MS switches, MX security appliances, and MR access points. Checkboxes across the top of the page allow for the enabling and disabling of various filters, to help focus on the topic or feature in question.

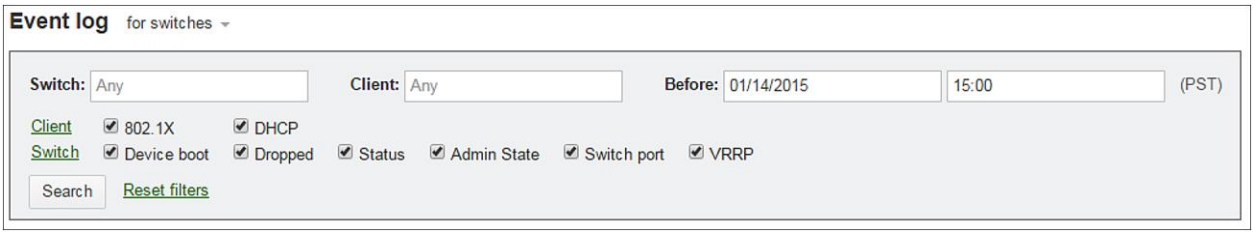

A brief description of the events that can be found in the event log and their meanings is available at: https://documentation.meraki.com/MX-Z/Monitoring_and_Reporting/Navigating_the_Event_Log.

For more detailed reporting, you can configure Cisco Meraki MX security appliances and MR access points to log messages to a syslog server. This feature provides an additional external location for log storage, as well as information that is difficult or impossible to find in the event log, such as traffic flows. Both the MX security appliance and the MR access points support three categories of messages: event log, URL access, and flows. The MX security appliance also supports IDS alerts, whereas the MR access points support air marshal alerts. To add a syslog server to the dashboard, do the following:

Step 1. Navigate to the **Network-Wide > General** page and the **Reporting** section.

Step 2. Click **Add a syslog server** and fill in the IP address and port of the server.

More details about the operation of this feature and configuration of the syslog server are available at: https://documentation.meraki.com/MX-Z/Monitoring_and_Reporting/Syslog_Server_Overview_and_Configuration.

Summary

Cisco Unified Communications Manager and Cisco Unified Communications Manager Express offer robust telephony services with the ability to grow alongside an expanding business. Cisco Meraki MX security appliances, MS switches, and MR access points integrate with these telephony services to provide a powerful and scalable network suitable for mass deployment. Cisco Meraki Dashboard provides an easy-to-manage interface to monitor and administrate users and devices across entire organizations. As VoIP continues to improve over traditional alternatives, the need for top-to-bottom integration of voice services continues to grow in importance.

References

- Cisco Meraki Knowledge Base: https://documentation.meraki.com
- Cisco Meraki Branch-Office Deployment Guide: http://www.cisco.com/c/dam/en/us/solutions/meraki-branch.pdf
- Cisco Collaboration System 10.0 Solution Reference Network Designs: http://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/collab10/collab10.html
- Cisco Unified Communications Manager Administration Guide: http://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/10_0_1/ccmcfg/CUCM_BK_C95ABA82_00_admin-guide-100.html
- Cisco Unified Communications Manager Express System Administrator Guide: http://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucme/admin/configuration/guide/cmeadm.pdf
- VoIP Call Admission Control: http://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/voip_solutions/CAC.html
- Cisco ASA and Cisco PIX® Firewall: Security Appliance to a Cisco IOS Software Router LAN-to-LAN IPsec tunnel configuration example: http://www.cisco.com/c/en/us/support/docs/security/asa-5500-x-series-next-generation-firewalls/63883-ipsec-rtr-2-pix-asa.html
- Medianet WAN Aggregation QoS Design 4.0: http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoSWAN_40.pdf

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C07-734285-00 04/15