Meraki Insight

Meraki Insight Overview

The **Meraki Insight** product is designed to give Meraki customers an easy way to monitor the performance of **Web Applications** on their networks and easily identify if any issues are likely being caused by the Network or the Application. This information is presented in a series of easy to understand graphs and charts that can clearly show if performance issues are being introduced within the local network or if performance issues are the result of something at the Application or WAN level.

> Meraki Insight requires the MX to be running a minimum firmware build of MX 14.20 or greater (14.21 or higher is recommended). Please refer to our Managing Firmware Upgrades article for more information about upgrading firmware.

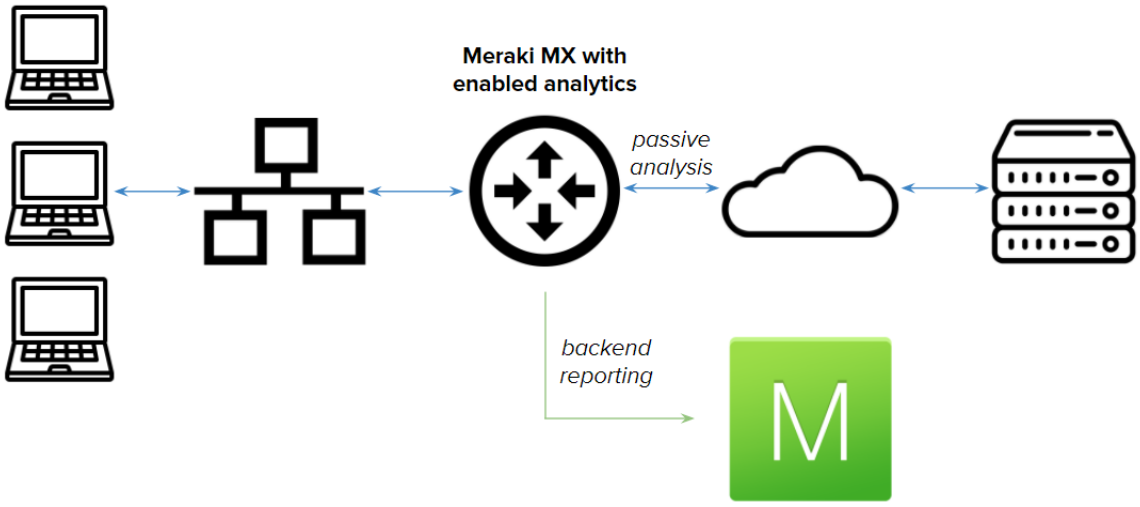

Web Application Performance Overview

Tracked Web Applications

Utilizing **Meraki Insight**, an MX can be configured to monitor and track all traffic associated with specific **Web Applications**. This data is tracked on a per-flow basis at the MX, then the relevant flows are aggregated into categorical groups based on their associated application and sent over an encrypted connection back to the Meraki Cloud Controller for further analysis before populating that flow data into the **Insight** feature on Dashboard. The data gathering process utilizes deep packet inspection, similar to our **Advanced Malware Protection** and **Intrusion Protection** systems, to allow Meraki to gather information at both the Network layer and the Application layer so we can help to identify if performance issues are based around local Network performance or Application layer issues.

Performance Indicator

The **Performance Indicator** is a symbol that is displayed for each application, and is intended to provide a quick reference to the quality of the user experience relative to a specific **Web Application**, based on the thresholds defined by a Meraki Administrator. A **green** performance indicator symbol means that both the Application level and Network level are performing within the defined thresholds and user experience should be optimal. A **red** performance indicator symbol indicates that there may be some noticeable issues with the performance of a specific application for some users.

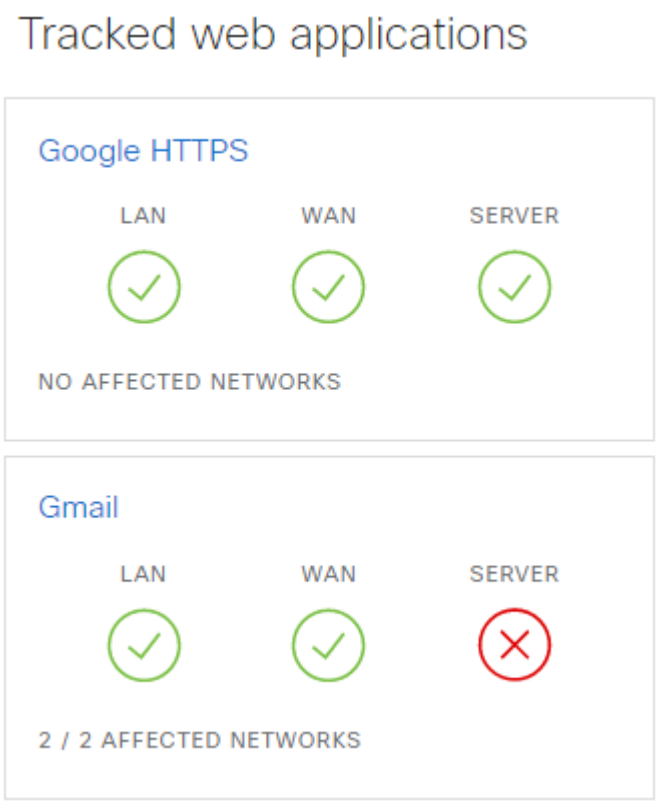

Affected Networks

If an Application has a **Performance Score** of <80% any Networks that are reporting a score below 80% will have a **red** performance indicator symbol and will be listed under the **Affected Networks** list. Clicking on the red performance

indicator symbol for each respective application source (LAN, WAN, Server) will display a list of Affected Networks. Clicking one of these **Affected Networks** bring you to an overview of all tracked **Web Applications** on that Network and their respective **Performance Indicators** for the selected network. For more information about the per-Network view of **Tracked Web Applications** please reference the 'Tracked Web Applications - Per Network View' section of this article.

Configure Web Applications

To begin tracking the performance of **Web Applications** there first has to be at least one category/application selected to monitor. To begin configuring **Web Applications** to monitor, select **Configure Web Applications** from the **Insight > Monitor > Web application performance** page, near the bottom of the applications list. From the popup window you can browse through several categories of applications to monitor, then once a category has been selected, browse until the desired application is listed. Mark the checkbox for that application and select **Save**.

> MI is intended to track your most critical apps. As best practice, we recommend tracking around 10 applications. Although you can track up to 50 apps, high numbers of tracked apps can lead to degraded dashboard performance.

Alternatively, **Web Applications** can be searched for directly by typing the application name into the **Search** bar ar the top of the popup window. Once the selected applications have been saved any MX that are configured to have **Insight** enabled will begin monitoring and reporting information on traffic that matches the applications selected.

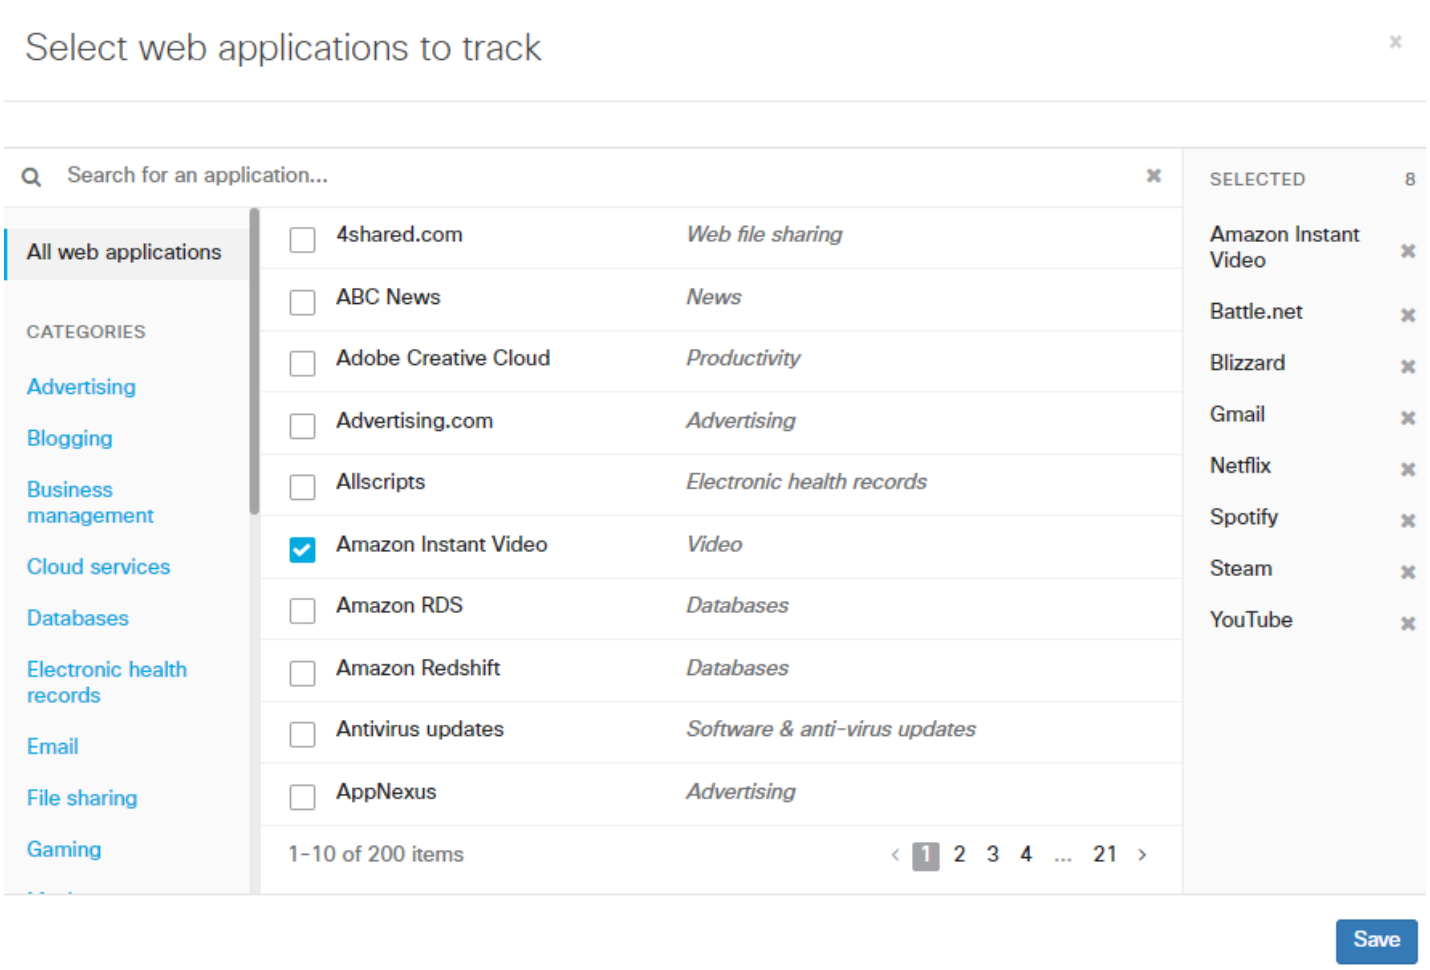

Thresholds

The **Application Performance Score** is calculated based on several defined **Thresholds**, primarily **Per-flow Goodput** and **Application Response Time**. The **Per-flow Goodput** is defined as the "predicted maximum amount of data that could be transmitted per flow, based on network latency and loss," while **Application Response Time** is based on the HTTP response time for requests initiated by the application, excluding any **Network Latency**. Both the desired **Per-flow Goodput** and the desired **Application Response Time** can be configured and customized on a per-Application basis by an Organization Administrator.

The defined minimum acceptable **Per-flow Goodput** can be set as low as 10Kbps or up to 10Mbps, with the default value being 160Kbps. Similarly, the maximum tolerable **Application Response Time** can be configured as low as 100ms and up to 100s with a default value of 3s. **Application Response Times** can also be configured to be ignored when calculating statistics for a specific application if the application is known to use long-polling or WebSockets, as these are expected to have long application response times that do not indicate an issue. This is done by selecting '**Ignore**' instead of '**Choose**' on the dropdown when configuring the **Application Response Time**.

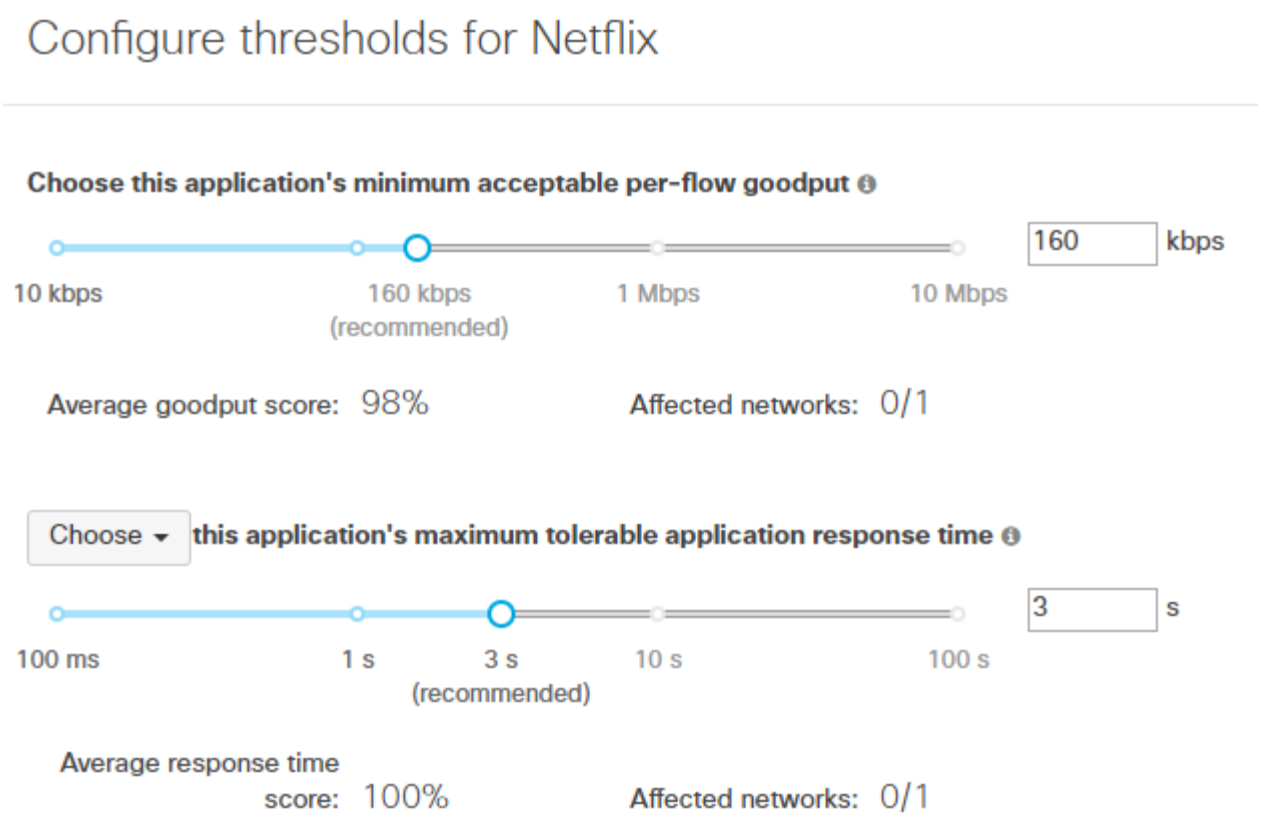

Adding Custom Applications

In addition to being able to add traffic from pre-defined categories, custom applications can be defined, based on hostname. To add or edit custom categories, select **Configure web applications**, and then select **Custom** from the categories.

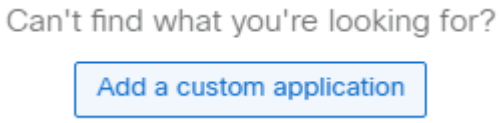

In the **Define a custom application** window, you will be able to fill out the Application name and hostname for the application you wish you track. For example, if you wanted to track traffic to/from google.com, you could enter "Google" as an application name, and "google.com" as the expression.

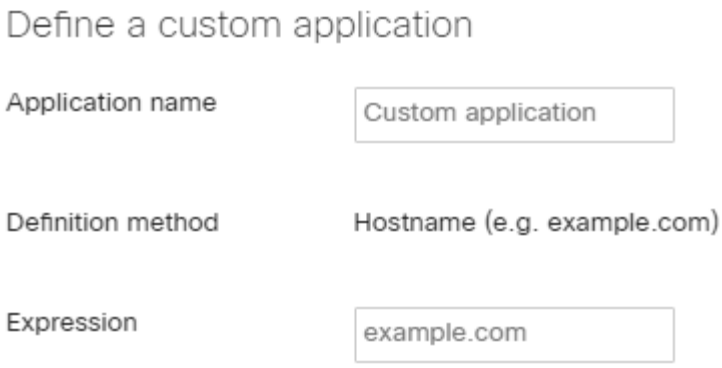

Note that MI will only monitor top level domain names, and will ignore subdomains. So if you add mail.google.com, only google.com will be tracked.

Adding Applications by IP Address

Traffic can also be tracked based on IP address. This can be helpful for tracking traffic to/from a local server or application, or a web application that does not have a hostname. To add an application based on IP address, the definition must first be added to Organization Traffic analysis. To do this, navigate to **Organization > Settings**, and then to the **Traffic analysis** section.

On this page, you can enter a name for the application/traffic you will be tracking, and the IP address for the application. Be sure to select **Save** to commit your changes.

Note that MI will only monitor HTTP/HTTPS traffic (TCP ports 80/443), and that ports cannot be specified.

After the definition has been added, you must set the application to be tracked in MI. To do this, navigate to **Insight > Web application performance**, select **Configure web applications**, and then select **Custom**. The definition you added in **Traffic analysis** should be listed in your custom applications.

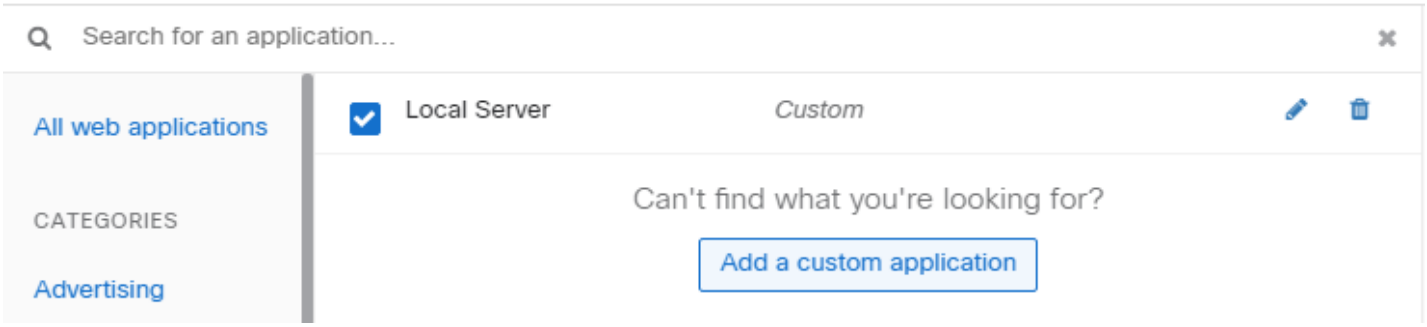

Check the box next to the application you added, then select Save, and you'll find the application in your Tracked web applications list. Note that it may take some before this new application generates some traffic.

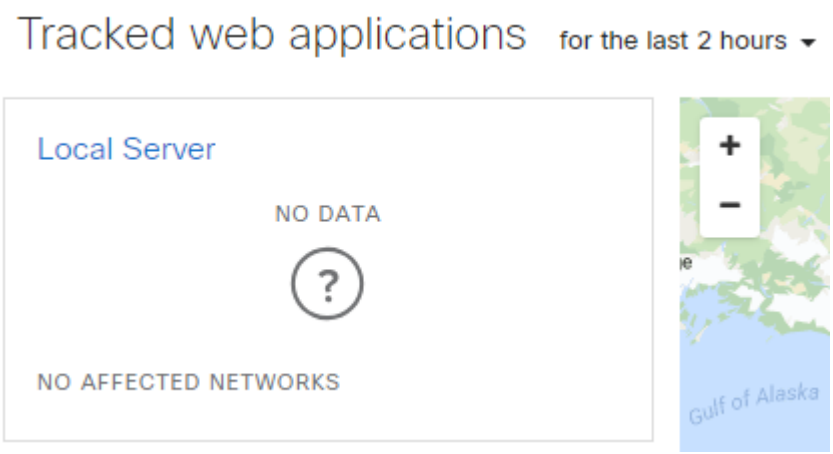

Insight Limitations

Meraki Insight is currently only supported on MX series devices. This does NOT include the Z-Series, MS, or MR series of devices. Additionally, MXs that are acting as Auto VPN Hubs will not be able to analyze traffic arriving over the VPN from Spoke sites. To gather data about traffic from Auto VPN Spoke sites, **Insight** must be enabled on the spoke MX. Finally, traffic originating from the WAN, such as that destined for an onsite web server will not be tracked.

Tracked Web Applications - Per Network View

Selecting an **Affected Network** from the **Tracked Web Applications** page will bring up all **Tracked Web Application** statistics for just the selected network. By clicking on the respective application indicator symbols, this page displays the **Performance Score** for each application for clients in the selected network for both the **Network-Layer** and the **Application-Layer** individually. Both the **Network-Layer** and **Application-Layer** scores are shown in a way that clearly shows how many clients are in the acceptable performance range and how many clients are having experiences that fall outside the configured acceptable range of performance, in addition to how far out of range the experience is.

For example, the image below shows the **Performance Score** for Meraki HTTPS traffic for the last week, which is currently below the respective threshold for two clients. If we click **View Trends**, we can check on the individual clients by clicking on the **Clients** tab. Looking at both the **Application-Layer** score chart we can see that while some clients

might be experiencing noticeable issues with the applications, not all clients are experiencing issues and those that are experiencing issues are experiencing application layer issues, not network layer issues.

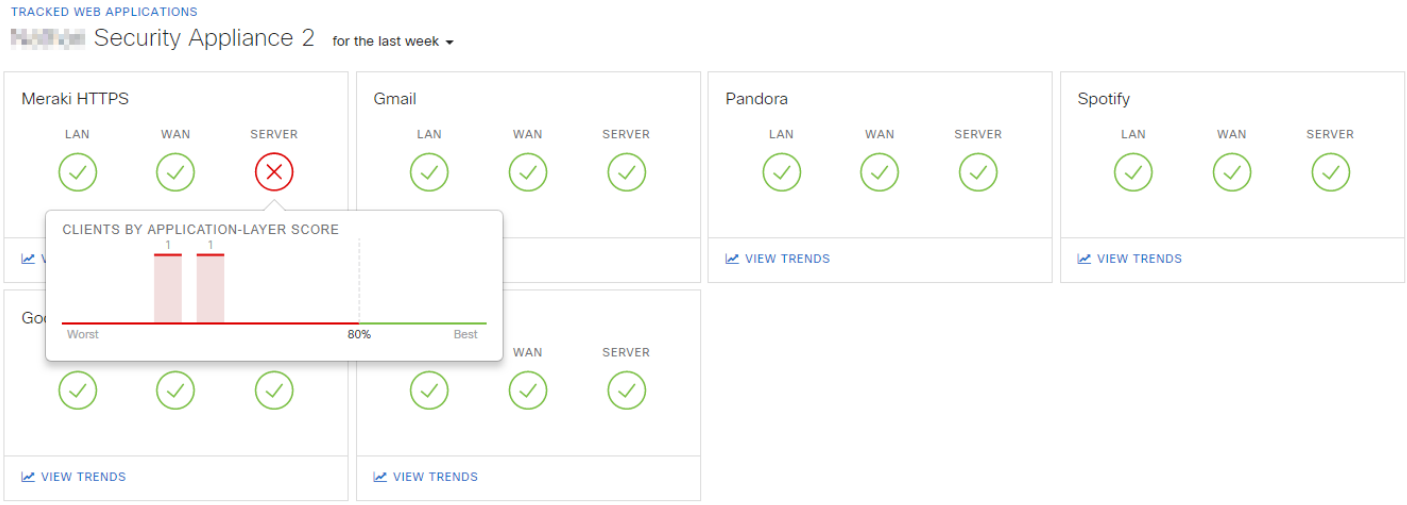

To view more detailed information about a specific **Web Application** click on the **View Trends** link for that Application.

Tracked Web Applications - Client Details View

Selecting a client from the **Network-wide > Monitor > Clients** page will open the **Client Details** page for that specific client. If the network has **Insight** enabled then this page will also list any **Tracked Web Applications** and the respective **Performance Score** for this client for each application. Clicking on the application name from this view will bring up the **Application Trends** page for the chosen application in the current network. For more information about **Application Trends**, please refer to the Viewing Application Trends section of this article.

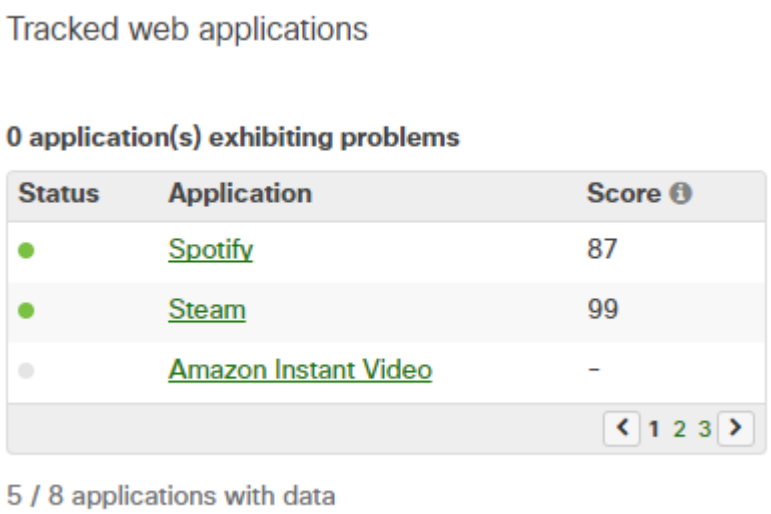

Viewing Application Trends

Overview

To get a better idea of how exactly a specific **Web Application** is performing we can view the trends for that specific application. The **View Trends** page is broken down into several different sub-pages, **Network**, **Application**, **WAN**, **LAN**, **Clients**, **Servers**, and **Domains**, with each one explained in detail below.

Network

The **Network** tab displays information related to the **Network-Layer** performance of data flows that match the definition of the selected **Web Application**.

Events

The **Events Timeline** displays a marker for each Network related event that is logged to the **Event Log**. Events such as VPN route changes and WAN failovers will be displayed here. To view the details of an event simply click on the marker and then select '**Details**' from the right panel under where the number of events is displayed.

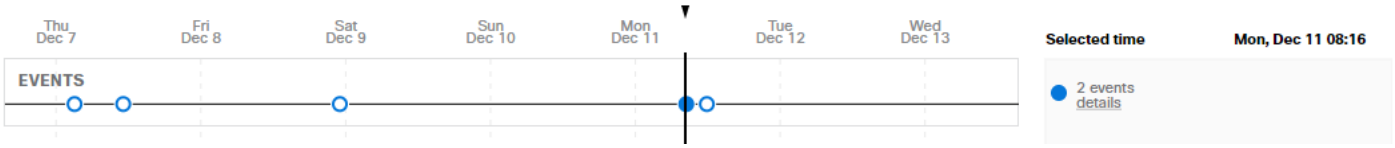

Performance Score

The **Performance Score** chart displays a historical timeline of the **Network Performance Score** for the selected application. The **Network Performance Score** is calculated by comparing the configured **Threshold** values for **Network Latency** and the actual recorded latency of the application.

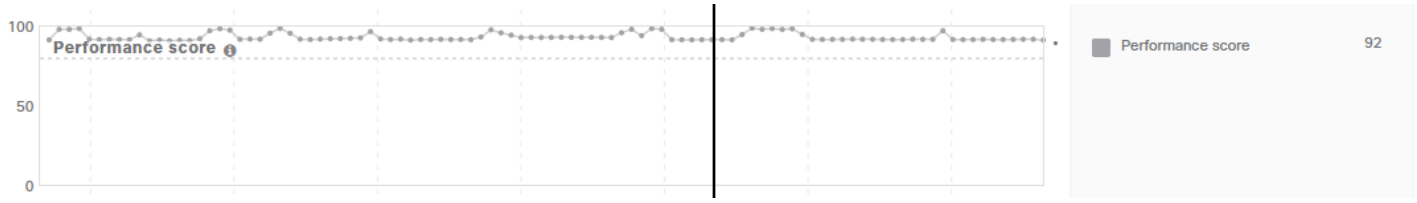

Total Network Usage

The **Total Network Usage** graph displays a historical view of the total bandwidth usage on the WAN of the MX for all traffic types.

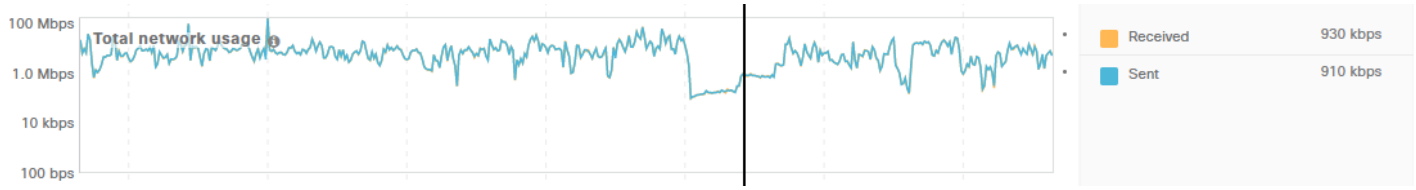

Latency

The **Latency** graph displays a historical view of the recorded TCP round trip time of connections made by the selected application. The network latency for an application is calculated based on the TCP SYN, SYN/ACK response time for connections from that specific **Web Application**.

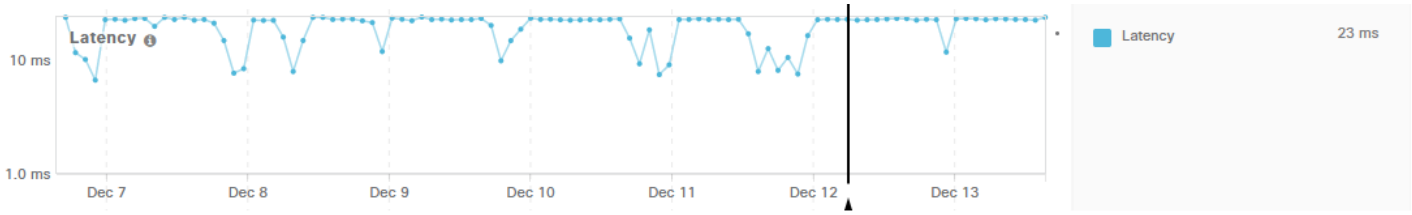

Application

The **Application** tab displays information about the **Application-Layer** performance that has been gathered from traffic flows matching the selected **Web Application**.

HTTP Response Time

The **HTTP Response Time** graph displays the historical average time between the last HTTP Request packet and the first HTTP Response packet for an application flow.

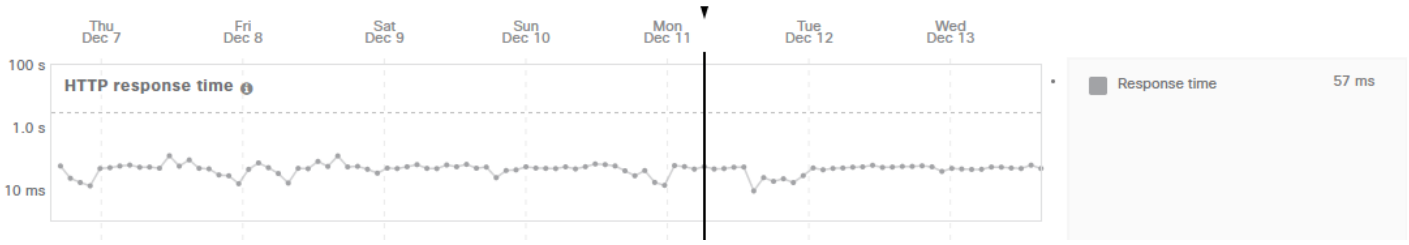

App Usage

The **App Usage** graph displays the total network usage of the selected application as recorded on the WAN of the MX, including the Sent and Received data.

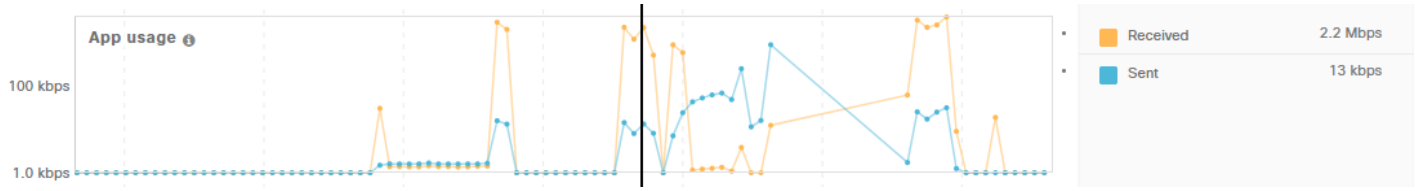

HTTP Request Rate

The **HTTP Request Rate** graph displays the historical average of HTTP Requests per-minute generated by the selected application.

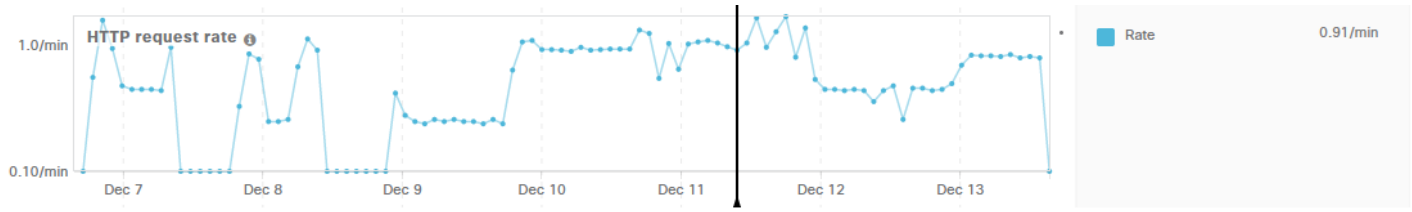

WAN

The **WAN** tab displays information about application performance specific to the WAN side of the MX.

Events

Exactly like the **Events Timeline** on the **Network** tab, the **Events Timeline** on the **WAN** tab displays a marker for each network related event that is logged to the **Event Log**. Events such as VPN route changes and WAN failovers will be displayed here. To view the details of an event simply click on the marker and then select '**Details**' from the panel on the right.

Available Goodput (WAN Limited)

The **Available Goodput** graph displays historical information about the recorded potential goodput on the WAN side of the MX. The **Available Goodput** is defined as the predicted maximum amount of data that could be transmitted per flow, based on network latency and loss on the WAN. The predicted **Available Goodput** is limited to 100Mbps, if the predicted goodput is higher than 100Mbps then the graph will still only display a maximum of 100Mbps.

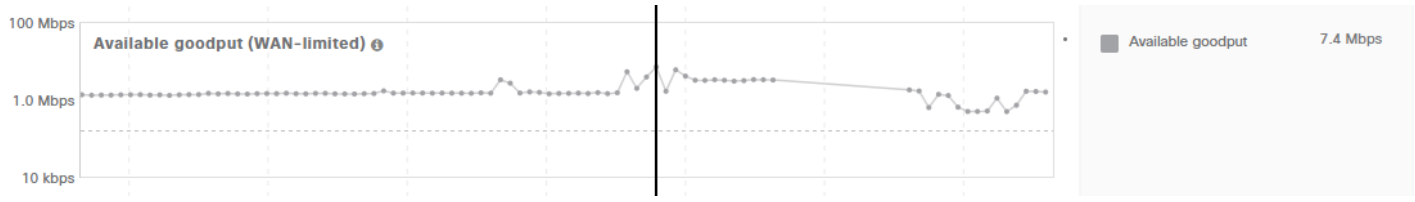

HTTP Response Time

The **HTTP Response Time** graph displays the historical average time between the last HTTP Request packet and the first HTTP Response packet as seen on the WAN.

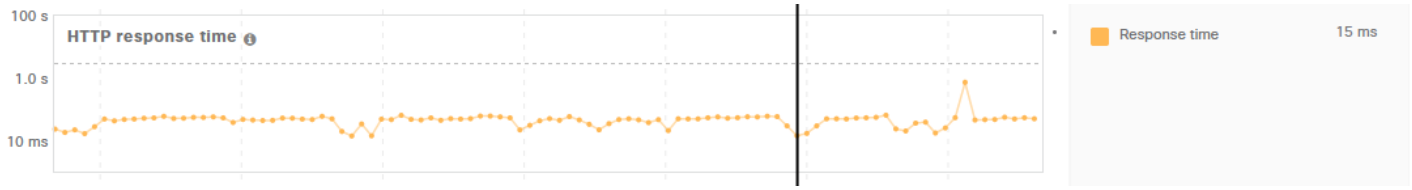

WAN Loss

The **WAN Loss** graph displays the amount of packet loss that has been detected on the WAN side for flows matching this application.

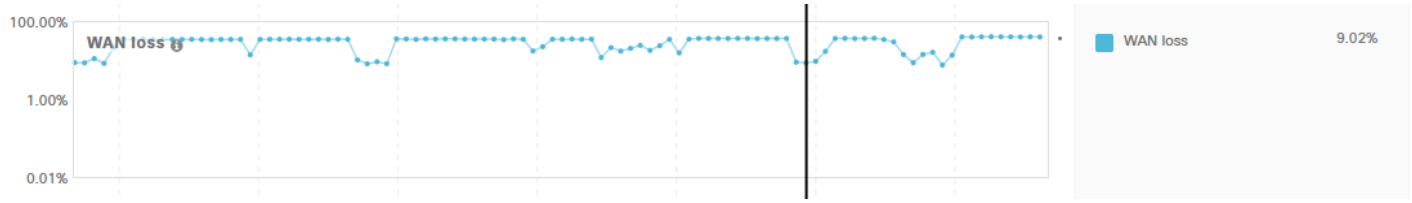

Latency

Exactly like the **Latency** graph on the **Network** tab, the **Latency** graph on the **WAN** tab displays a historical view of the TCP round trip time of connections made by the selected application as seen on the WAN interface. The reported latency is calculated based on the TCP SYN, SYN/ACK response time for the initial connections made by the specified **Web Application**.

Total Network Usage

Exactly like the **Total Network Usage** graph on the **Network** tab, the **Total Network Usage** displays a historical view of the total bandwidth usage on the WAN of the MX for all traffic types.

LAN

The **LAN** tab displays information about the network performance on the LAN side of the MX.

Events

Exactly like the **Network** and **WAN** tabs, the **Events** timeline on the **LAN** tab displays a marker for each network related event that is logged to the **Event Log**. Events such as Active Directory communication failures will be displayed here. To view the details of an event simply click on the marker and then select '**Details**' from the panel on the right.

Available Goodput (LAN Limited)

The **Available Goodput** (LAN Limited) graph displays historical information about the recorded potential goodput on the LAN side of the MX. The **Available Goodput** (LAN Limited) is defined as the predicted maximum amount of data that could be transmitted per flow, based on network latency and loss on the LAN. Like the **Available Goodput** graph on the **WAN** tab, the predicted **Available Goodput** is limited to 100Mbps. If the predicted goodput is higher than 100Mbps then the graph will still only display a maximum of 100Mbps.

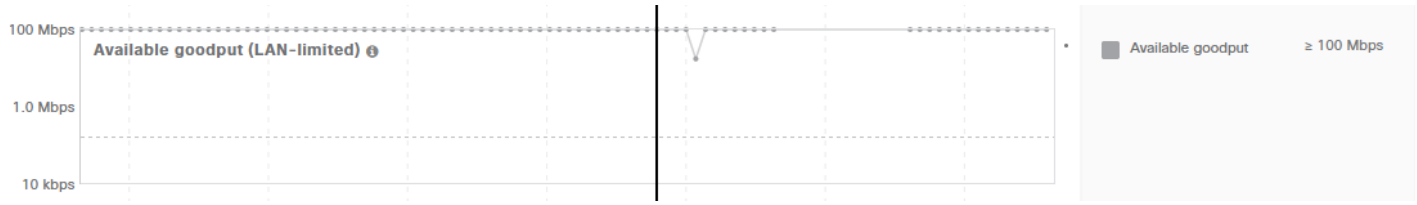

LAN Loss

The **LAN Loss** graph displays the measured packet loss on the LAN for flows matching the specified application.

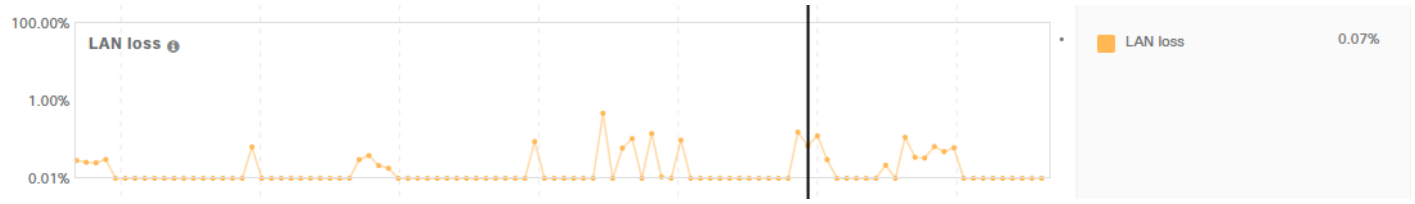

Total Network Usage

Exactly like the **Total Network Usage** graph on the **Network** and **LAN** tabs, the **Total Network Usage** displays a historical view of the total bandwidth usage on the WAN of the MX for all traffic types.

Clients

The **Clients** tab displays information about each client that has used the specified **Web Application** during the selected time period. The information includes the average **Performance Score** for a given client and the current application, the number of requests the client has made, and the average **HTTP Response Time**.

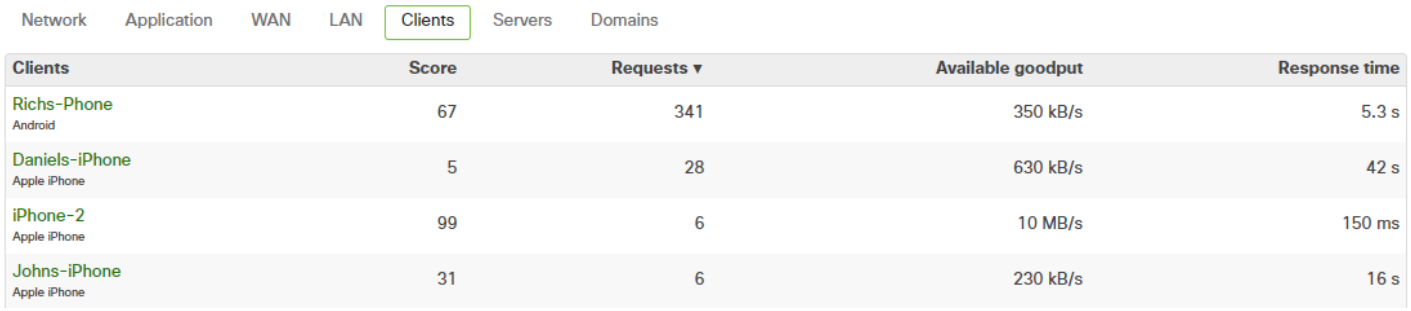

Clients

The **Clients** column lists the **Description** and detected **Operating System** of clients who have generated traffic flows that match the definition of the specified **Web Application**. Clicking on the client **Description** will open the **Client Details** page for that client.

Score

The **Score** column displays the calculated **Performance Score** of a given client for the specified **Web Application** over the selected time period.

Requests

The **Requests** column displays the number of **HTTP Requests** that were made by the **Web Application** from a given client for the specified time frame.

Available Goodput

The **Available Goodput** column displays the average predicted **Available Goodput** to the **Web Application** for a given client for the specified time frame. The **Available Goodput** is defined as the predicted maximum amount of data that could be transmitted per flow, based on network latency and loss.

Response Time

The **Response Time** column displays the average **HTTP Response Time** for the application for a given client for the specified time frame. HTTP Response Time is calculated as the time differential between last HTTP Request packet and the first HTTP Response packet of a flow.

Servers

The **Servers** tab displays information about the remote **Web Servers** that have been utilized by the **Tracked Web Application** during the chosen time period. This can be useful to help identify if there is a specific **Web Server** that could be contributing to application issues.

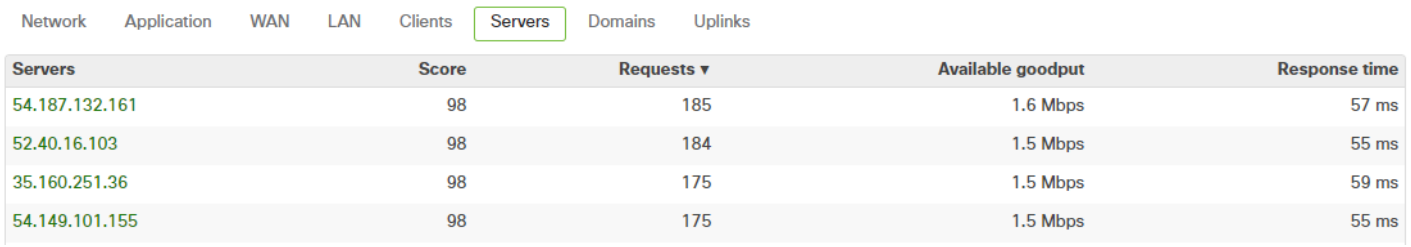

Servers	Score	Requests ▾	Available goodput	Response time
54.187.132.161	98	185	1.6 Mbps	57 ms
52.40.16.103	98	184	1.5 Mbps	55 ms
35.160.251.36	98	175	1.5 Mbps	59 ms
54.149.101.155	98	175	1.5 Mbps	55 ms

Servers

The **Servers** column lists identifying information about the remote server that have been contacted. This could include just the servers IP address or the full address of the server.

Score

The **Score** column displays the average calculated **Performance Score** for flows that are communicating to each specific server over the specified time frame.

Requests

The **Requests** column displays the number of **HTTP Requests** that have been sent to a given server over the chosen time period.

Available Goodput

The **Available Goodput** column displays the predicted maximum amount of data that could be transmitted per flow, based on network latency and loss to a specified server over the chosen time period.

Response Time

The **Response Time** column displays the average **HTTP Response Time**, minus **Network Latency**, for flows destined for a given server over the chosen time period.

Domains

The **Domains** tab displays information about different **Web Domains** that have been contacted by the selected **Web Application**. Similar to the **Servers** tab, this can be useful to determine if there is a specific domain that could be contributing to application performance issues.

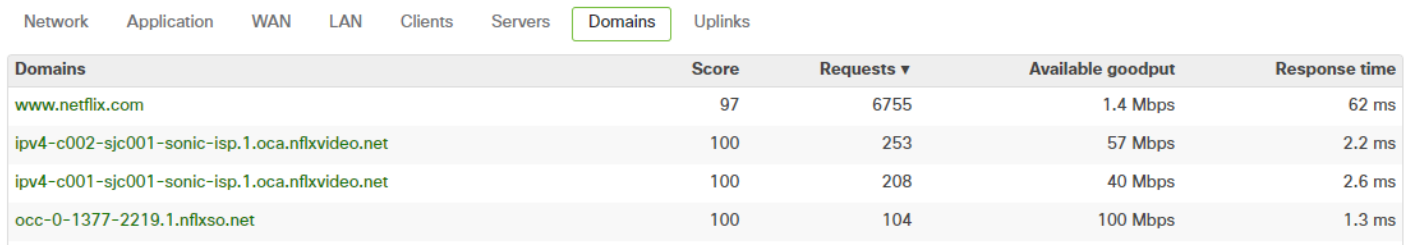

Domains	Score	Requests ▾	Available goodput	Response time
www.netflix.com	97	6755	1.4 Mbps	62 ms
ipv4-c002-sjc001-sonic-isp.1.oca.nflxvideo.net	100	253	57 Mbps	2.2 ms
ipv4-c001-sjc001-sonic-isp.1.oca.nflxvideo.net	100	208	40 Mbps	2.6 ms
occ-0-1377-2219.1.nflxso.net	100	104	100 Mbps	1.3 ms

Domains

The **Domains** column lists domains that have been contacted by the selected **Web Application**.

Score

The **Score** column shows the average calculated P**erformance Score** for flows that are communicating to each specific domain over the specified time frame.

Requests

The **Requests** column displays the total number of **HTTP Requests** sent to each domain by all clients in the network over the specified time period.

Available Goodput

The **Available Goodput** column displays the predicted maximum amount of data that could be transmitted per flow, based on network latency and loss to a specific domain over the chosen time period.

Response Time

The **Response Time** column displays the average **HTTP Response Time** for a given domain over the specified time period.

Licensing

Meraki Insight requires additional licensing that is separate from the standard Dashboard licensing for devices. Like standard licensing, Insight licenses are available in 1yr, 3yr, and 5yr license lengths. Insight licensing is separated by hardware model and expected throughput to be monitored. For example, both the MX84 and MX100 are capable of 250-750Mbps of throughput, so at minimum a 'Medium' Insight license is required to enable Meraki Insight features on a network containing an MX84 or MX100. A network containing an MX84 would not be able to use a 'Small' Insight license since that is capped at a throughput maximum of 250Mbps, making it only applicable to networks containing MX60, MX64, and MX65 devices. The following chart shows the different licensing tiers, throughput ranges for each license type, and their applicable hardware models.

License Type	Supported Throughput	Applicable Hardware
Meraki Insight (Small)	Up to 250 Mbps	MX60(W), MX64(W), MX65(W)
Meraki Insight (Medium)	Up to 750 Mbps	MX84, MX100
Meraki Insight (Large)	Up to 5 Gbps	MX250, MX400, MX600

Meraki Insight (XLarge)	Up to 10 Gbps	MX450

NOTE: Higher tiered licenses can be applied to lower models of MX but lower tiered licenses cannot be applied to higher models of MX. For example, an MX64 can use a 'Large' Insight license but an MX250 cannot use a 'Small' or 'Medium' Insight license.

NOTE: Like standard Dashboard licensing, networks running an MX pair in an HA configuration will only require a single Insight license.

NOTE: Networks bound to a configuration template in Dashboard will still require individual licensing for MI features.

Managing Licensing

Meraki Insight licensing is applied on a per-Network basis, so each network that is implementing Meraki Insight will require the appropriate license to be applied to that Network. After claiming the licenses like normal Dashboard licenses they can be managed by going to **Insight > Configure > Licensing**. From this page we can see how many of each type of Insight license is available and how many are currently applied to existing Networks. This works similarly to the Device Count and License Limit counts on the regular License Info page.

Enabling and Disabling Meraki Insight

To Enable Meraki Insight on a network:

1. Ensure the necessary Licensing is available
2. Select the checkbox next to the Network where Insight should be enabled
3. Click '**Add network(s) to Insight**' at the top left of the table to Enable Meraki Insight on the selected Network(s)

To Disable Meraki Insight on a network:

1. Select the checkbox next to the Network with Insight currently enabled
2. Click '**Remove network(s) from Insight**'

FAQ

Is Meraki Insight supported on all Meraki products?

Currently it is only supported with MX series devices. End-of-Life MX devices and Z-Series devices are currently not supported

What is the firmware version required on MX series devices to add support for Insight?

MX devices need to be running firmware version MX 14.20 at a minimum in order to support Meraki Insight.

What are the protocols and types of applications that this product can track performance for?

Currently web-based applications (HTTP and HTTPS) only.

Does Insight send my network traffic to Meraki?

Meraki Insight aggregates network statistics on the MX and only sends those aggregates back to the Meraki Cloud, so local network traffic is still segregated on your network. These aggregated statistics are sent as TLS-SYSLOG traffic (TCP 6514).

Can Insight track performance for VoIP traffic?

Meraki Insight currently supports tracking of web-based applications only.

Can Insight track performance for custom applications?

Custom Applications can be tracked as long as the application traffic passes through the MX and the application has been defined within **MI**. Custom applications can be configured from the Traffic analysis section in the **Organization > Settings** page by using a hostname or IP address. Once configured the custom application will be available from the list of applications.

What happens if I remove Insight licensing from a Network?

Insight licenses can be moved between networks with just a few clicks, but disabling Insight on a Network will mean any Insight related data for that Network will be lost after several days. If Insight licensing is reapplied prior to the data being lost it will become viewable again, however no data will be gathered and existing data will not be viewable during the time the network is not licensed.

Is Insight included with Advanced Security Edition licensing?

Meraki Insight requires additional licensing as it is a separate product. Please contact your Cisco Meraki Sales representative to inquire.

Why is my application not being tracked?

Only traffic that passes through an MX can be tracked by Meraki Insight. For example, if the MX is handling all routing for a location then both WAN bound and Inter-VLAN bound application traffic will be tracked. However if Inter-VLAN

routing is happening on a downstream device and only WAN bound traffic passes through the MX then only WAN bound application traffic will be tracked.

What back-end technology is MI built on?

MI was built from the ground up by Meraki to fit the Meraki ecosystem. There is nothing to install (even on the back end).

How does MI impact the performance of the MX it is running on?

The MX appliance is used as a 'collector' to gather the data, and the cloud does the heavy lifting from there and provides the analyses of the LAN, WAN, ISP, client/server stats with retrospection. The data is based on end-user HTTP/S data, so there's no need for synthetic probes, and MI leverages the MX’s deep packet inspection that's already happening, so MI has no significant performance impact on the MX itself.