



The bridge to possible

Meraki MX Inside and Out:

A Support Look at Design, Best Practices,
and Troubleshooting

Serhii Kucherenko


Customer Escalations Engineer

BRKTRS-2007

CISCO *Live!*


#CiscoLive

What can you tell about this place?



What if we put it into perspective –

The ability to see the bigger picture matters during design maintenance and troubleshooting.




A real story ...


Once upon a time, during network planning, I ordered Full POE+ switches that should be powered up from Rack PDUs. PDU power cords were ordered separately ...



Find the way
to put this



into the hole
made for this



by using
nothing but this

...



A word about the speaker – less formal

- My dog smiles more than me.
- He also woofs with less accent than I speak.
- I try to spend as much time as possible with my son.
- He is secretly being prepared to join Cisco one day.



Cisco Webex App

Questions?

Use Cisco Webex App to chat with the speaker after the session

How

- 1 Find this session in the Cisco Live Mobile App
- 2 Click “Join the Discussion”
- 3 Install the Webex App or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 7, 2024.

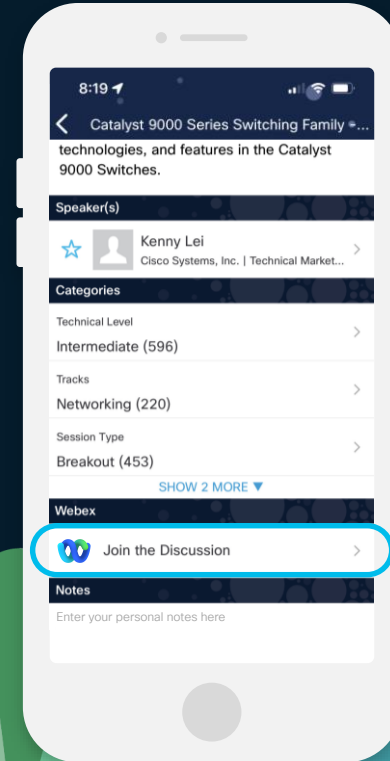
Chakshu Piplani
Senior escalation
engineer (MX)



Abhishek Suresh
Senior escalation
engineer (MX)



<https://ciscolive.ciscoevents.com/ciscolivebot/#BRKTRS-2007>



Additional content

The complete version of the slide deck, along with demos and documents mentioned throughout the session, are available for download –

[Here](#)



In case of any problems, please let me know:

skuchere@cisco.com

Content warning

- The challenge with design and troubleshooting sessions is that all environments are unique.
- Examples in this session may not cover your specific pain points. Yet, they should give you direction.
- If the session content doesn't address your specific problem, see me after the session or use Webex to post questions on your scenarios.





Agenda


- A warm-up
- MX performance - Best practices and caveats
- Unobvious Auto VPN
- Gamified troubleshooting case study



Agenda

- A warm-up
- MX performance - Best practices and caveats
- Unobvious Auto VPN
- Gamified troubleshooting case study

slido



What is your relationship with Meraki MX?

① Start presenting to display the poll results on this slide.

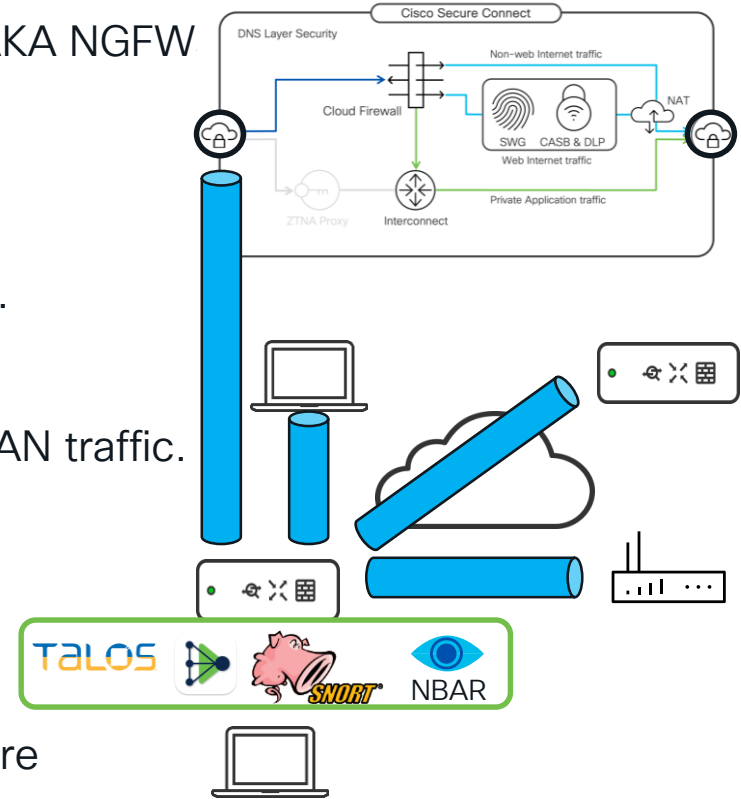
MX – all you need to know in one slide




For your
reference

MX is an SD-WAN Security appliance (AKA UTM, AKA NGFW) with the following main features:

- Stateful Firewall/NAT.
- Member of SD-WAN fabric based on Auto VPN.
- L7 firewall policies with user identity.
- Intelligent path selection for Internet and SD-WAN traffic.
- Snort-based IDS/IPS.
- TALOS-based content filtering.
- AMP/Umbrella/ for end-user protection.
- Third-party VPNs, RA VPNs, Part of Cisco Secure Connect (AKA Meraki SASE)



slido



For what problems do you typically open MX-related cases?

① Start presenting to display the poll results on this slide.

slido



What type of MX problems
caused you the most pain?

ⓘ Start presenting to display the poll results on this slide.



Agenda

- A warm-up
- MX performance – Best practices and caveats
- Unobvious Auto VPN
- Gamified troubleshooting case study

MX performance – where do we start?


A sizing guide is our single source of truth during the design or troubleshooting.

- At the time of design activity, we wish to understand which MX model is the right pick for the given task/location.
- And during troubleshooting, we try to understand if the affected MX is within the recommended numbers.



MX Sizing Guide facts


- The [current](#) sizing guide version is based on MX 18.2.
- The guide was migrated to the Meraki documentation website.
- Older versions of the guide are no longer published.
- Sizing guides for MX18.1, MX17, and MX16 are available in additional materials for this session.



Numbers of the most interest

The most critical MX performance metrics can be placed in the following groups:

- Throughput numbers (Next Generation Firewall, Advanced Security, VPN, etc.).
- VPN numbers (Maximum tunnel count, recommended tunnel count, maximum numbers for remote access VPN).
- Maximum Device Count.
- Maximum flow count.



Sizing guide 18.107 vs 18.2xx

A brief comparison immediately shows a significant increase in the Throughput numbers for the high-end MX platform:

18.107

	MX67/68	MX75	MX85
Max stateful (Layer 3) firewall throughput in NAT mode with large payload file transfers	600 Mbps	1 Gbps	1 Gbps

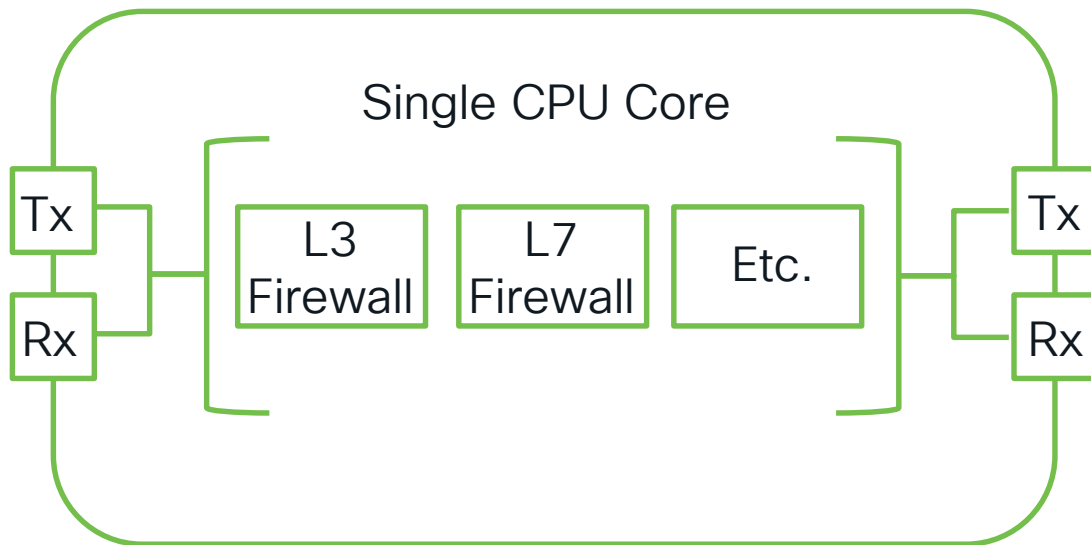
18.2xx

	MX67/68	MX75	MX85
NGFW Throughput RFC2544 - 1518 Byte	600 Mbps	1 Gbps	1 Gbps

MX95	MX105	MX250	MX450
2 Gbps	3 Gbps	4 Gbps	5 Gbps
MX95	MX105	MX250	MX450
2.5 Gbps	5 Gbps	7.5 Gbps	10 Gbps

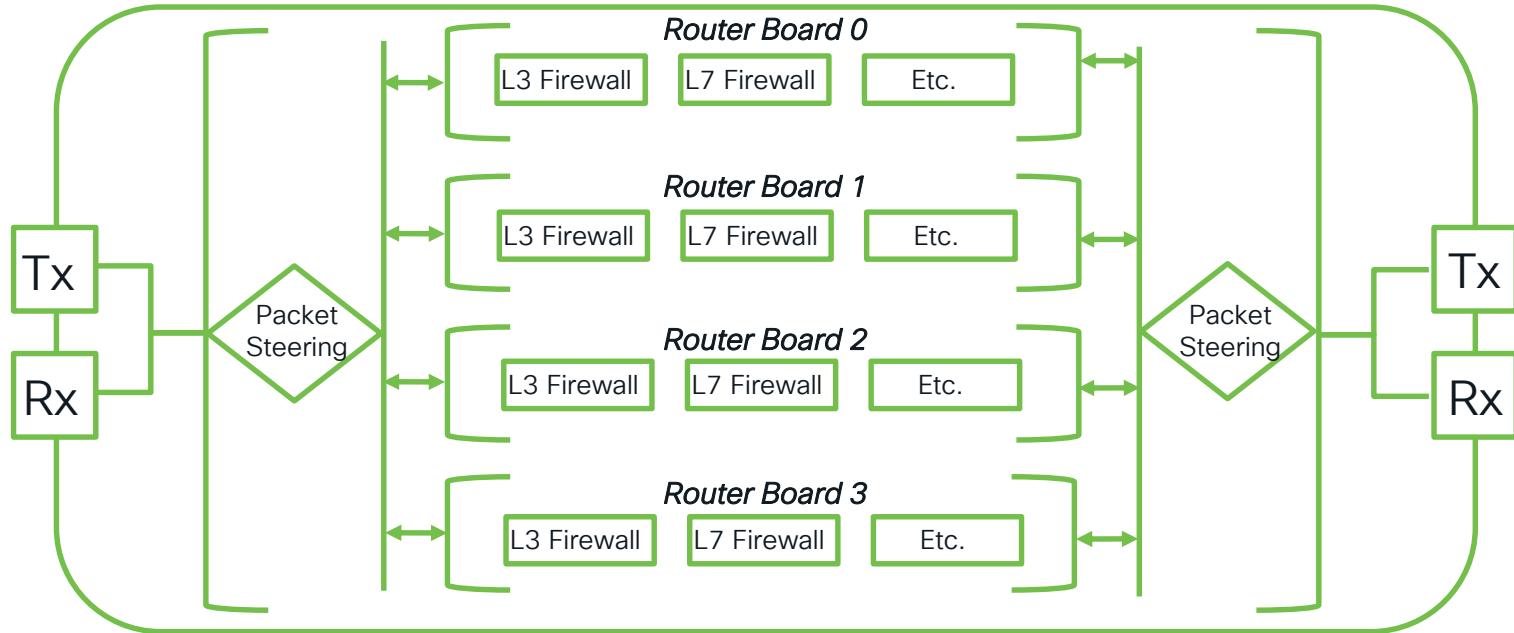
The engine of performance increase

Historically, MX used the below architecture for the packet processing:



The engine of performance increase

Starting from MX 18.2, a new architecture was adopted in which traffic processing was distributed across multiple cores. Each processing engine is called a *Router Board*.



Where is it applicable?

The current state of different MX models:

- One router-board appliances: Z3, Z4, MX67, MX68
- Two router-board appliances: MX75, MX85, MX95, MX105, MX250
- Four router-board appliances: MX450
- Models that don't support 18.2: MX64, MX65, MX84, vMX100



MX scaling parameters that are easily overlooked

- Scaling guide throughput VS real network traffic.
- Recommended Maximum Device Count on MX Hubs.
- Maximum Concurrent Sessions (Flow count)
- Recommended number of tunnels on hubs.




Scaling guide throughput VS real network traffic

- The top number in the guide for each platform says:


*NGFW
Throughput
RFC2544 - 1518 Byte*

- This test is the most favorable for the device.
- The question that arises is - How it's applicable to the real network?



Test VS real life

- The combined volume of 'boxes' in both examples is the same.
- At the same time, due to the difference in the number of 'boxes', the MX has to perform more operations in the second example.
- 1 Gbps of traffic consisting of 1518-byte packets is not equal to 1Gbps of traffic composed of packets with random sizes.



Forecasting MX throughput – Disclaimer

- The forecast method presented further represents an analytical approach.
- The Meraki Escalation team uses it to establish a minimum expected throughput in the given network.
- During the design stage, it is always recommended that you consult with Meraki Sales Engineers when selecting devices for critical network points (data centers, VPN hubs, etc.).




Converting maximum throughput into the PPS

- To calculate the *minimum forecasted throughput*, we use the *maximum platform PPS* for the specific test and the *average packet size* in the given environment.
- The formula to calculate the Maximum platform's PPS:

$$\text{NGFW Throughput RFC2544}^* \text{ (in bps)} / 8 / 1518$$

- An example for MX450 (18.2)

$$10 \times 10^9 / 8 / 1518 = 823452 \text{ PPS}$$




* – RFC2544 tests are used since packet sizes are known for them, which allows throughput to be converted to PPS.

How to get an average packet size

Multiple tools can be used to help with estimating real-life values:

- NetFlow Analysers (MX can be a NetFlow exporter)
- Statistics > Packet Lengths in Wireshark
- Built-in network equipment tools

Packet sizes



■ 400-700 bytes ■ 100-400 bytes
■ 700-1000 bytes ■ 1000-1300 bytes

Topic / Item	Count	Average	Min Val	Max Val	Rate (ms)	Percent	Burst Rate	Burst Start
▼ Packet Lengths	100000	412.86	42	1514	20.6553	100%	57.2300	3.187
0-19	0	-	-	-	0.0000	0.00%	-	-
20-39	0	-	-	-	0.0000	0.00%	-	-
40-79	21995	59.42	42	79	4.5431	22.00%	12.6000	3.187
80-159	3534	117.50	80	159	0.7300	3.53%	2.6300	3.187
160-319	52067	214.55	160	319	10.7546	52.07%	30.5500	4.720
320-639	2573	467.70	320	639	0.5315	2.57%	1.5500	4.020
640-1279	2510	968.73	640	1279	0.5184	2.51%	1.5300	4.037
1280-2559	17321	1429.36	1280	1514	3.5777	17.32%	10.2900	3.187

```
HQ-Distribution-1#show interfaces Fa1/0/2 controller
```

Calculating the minimum forecasted throughput

- After we calculated the Maximum platform's PPS, we can use the below formula for the minimum forecasted throughput:

$$\text{Maximum platform's PPS} * \text{Average Packet Size} * 8$$


- For the network with an Average packet size of 500 bytes:

$$823452 \text{ PPS} * 500 * 8 = 3,29^* \text{ Gbps}$$

- For the network with an Average packet size of 1000 bytes:

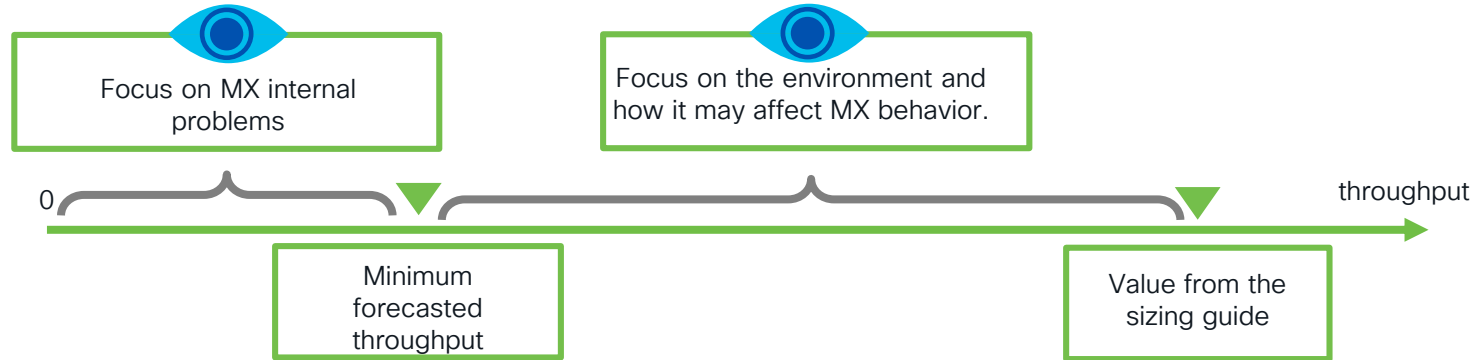
$$823452 \text{ PPS} * 1000 * 8 = 6,59^* \text{ Gbps}$$

*—Analytical results are typically lower than accurate live testing would show, but they are useful as a reference point.



How the minimum forecasted throughput used

- At the time of investigation of MX performance issues, a one-axis performance graph can be established:



- During the design phase, the minimum forecasted throughput value can be used for conservative capacity planning.



Recommended Maximum Device Count

- In standard MX terminology, 'Device' means an active IP or MAC discovered on the appliance's LAN side.
- In the case of 'branch' MXs, getting a number of active devices is straightforward.

1 Network-wide

2 Monitor

Clients

Packet capture

Event log

Map & floor plans

Configure

General

Alerts

Group policies

Users

Add devices

Search for clients


Status, type, OS

Status	Description
<input type="checkbox"/> Wired	10.104.0.10
<input type="checkbox"/> Wired	10.104.4.10

Status

☐ Offline

☐ Online



- Exceeding the number of recommended clients is a rare situation for the 'branch' MXs

slido



How many clients will DC-MX report when tracking by MAC is enabled?

① Start presenting to display the poll results on this slide.

slido



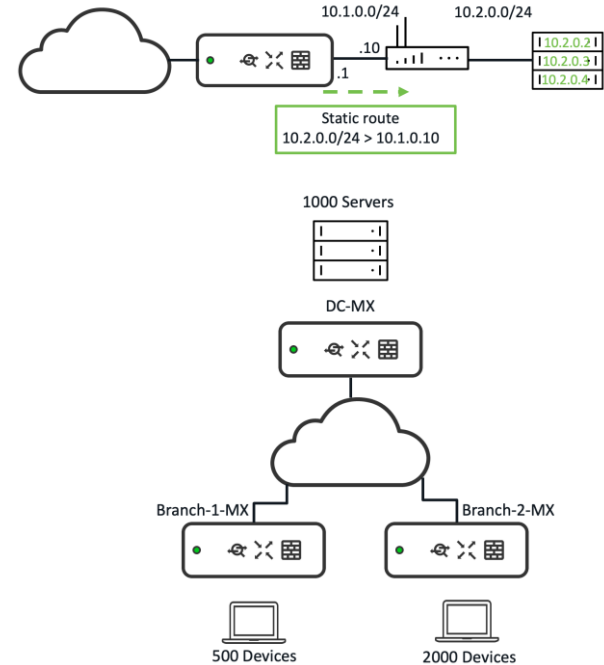
During the design phase, what number of clients should be used as a target for DC-MX if all remote devices use DC applications?

① Start presenting to display the poll results on this slide.

Recommended Maximum Device Count (Hub)

The situation is less straightforward in the case of Routed mode Hub or One-Armed Concentrator.

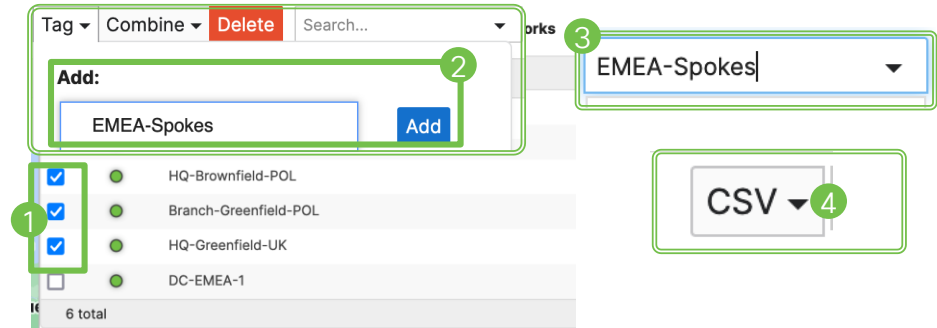
- Clients in MX networks are tracked by MAC (default). In the Routed DC, this does not help much in counting LAN-side clients correctly.
- Routed hubs can be configured to track by IP, but this is not recommended for One-Armed Concentrators (as they start considering external IPs as clients).
- Clients located on remote sites that use resources placed behind the hub should be considered as hub clients at the time of network design and evaluation.



Recommended Device Count – getting a number

Unfortunately, there is no straightforward way to get the number of clients that cross the hub, but some best practices can be applied:

- Add a specific tag to all spokes that use a hub so that they can be filtered in **Organization > Overview** to extract results as CSV.
- Meraki APIs can be used to obtain more accurate data.



The screenshot shows a table of device statistics. The table has columns: Name, Network type, Devices, Offline devices, % offline, Clients, and Usage. The data is as follows:

	Name	Network type	Devices	Offline devices	% offline	Clients	Usage
2	Branch-Greenfield-UK	Security & SD-WAN	1	0	0.0%	1	161 KB
3	HQ-Brownfield-POL	Combined	7	1	14.3%	8	2.36 GE
4	HQ-Greenfield-UK	Security & SD-WAN	1	0	0.0%	157	2.61 TB

Below the table, there are summary statistics: SUM 166, AVERAGE 55,333333333, MIN 1, MAX 157.

```
GET {{baseUrl}}/organizations/{{MyOrgID}}/appliance/vpn/stats?networkIds[]={EMEA-DC-1-Network}}
```

```
GET {{baseUrl}}/networks/{{EMEA-Spoke-1-Network}}/clients?statuses[]=Online
```

MX limits – flow count

- In the latest version of the sizing guide, information on ‘Maximum Concurrent Sessions’ was added for all platforms.

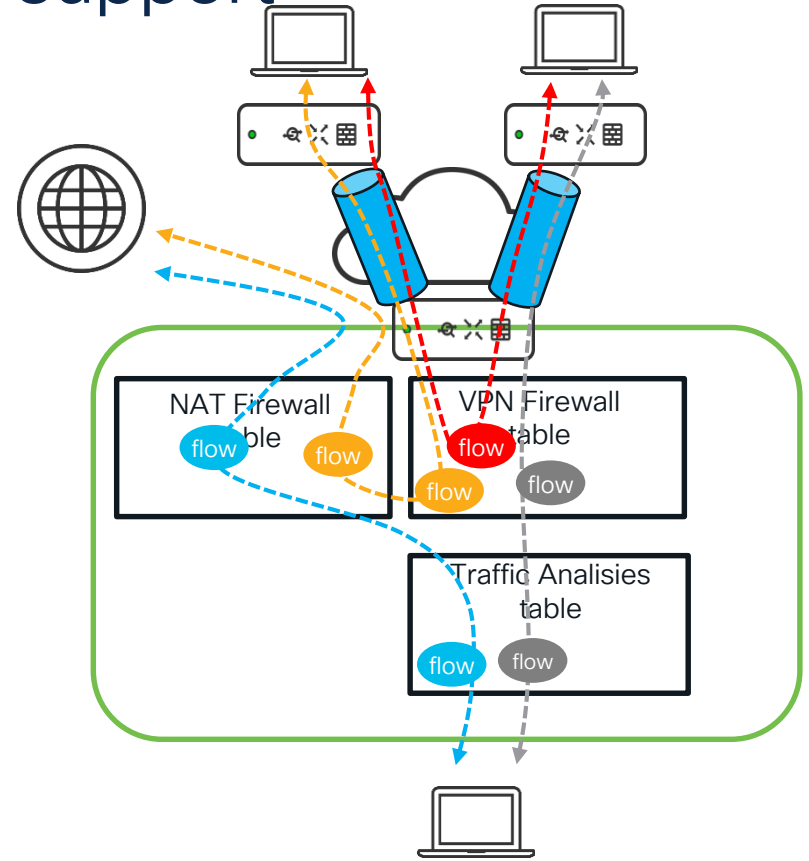
	MX95	MX105	MX250	MX450
Maximum Concurrent Sessions	200,000	250,000	500,000	1,000,000

- The numbers provided show absolute platform capacity and can't be considered as a target for planning.
- As of today, the only way to get to know the number of flows that MX is processing is to contact support.



MX flow limits – talking with support

- MX has multiple flow tables.
- The three most common are *Traffic Analysis*, *NAT Firewall*, and *VPN Firewall*.
- Exceeding flow capacity in any of the tables may result in MX instability.
- Depending on traffic source/destinations, flow can be created in one or more tables.



MX unobvious recommendation for the flow count

- In the latest version of the sizing guide, an implicit number for recommended flows was introduced.

Use Case Recommendations

A use case recommendation is based off of the device throughput; available feature set, and maximum considered to consume up to 50 flows.

A use case recommendation is based off of the device throughput; available feature set, and maximum flow table capacity. In this calculation, each client is considered to consume up to 50 flows.

MX-Series

	MX67	MX68	MX75	MX85				
Recommended Maximum Device Count	50	50	200	250	500	750	2,000	10,000

- This allows us to introduce the formula for the 'maximum recommended' number of flows.

Recommended Maximum Device Count * 50 = 'maximum recommended flow count'

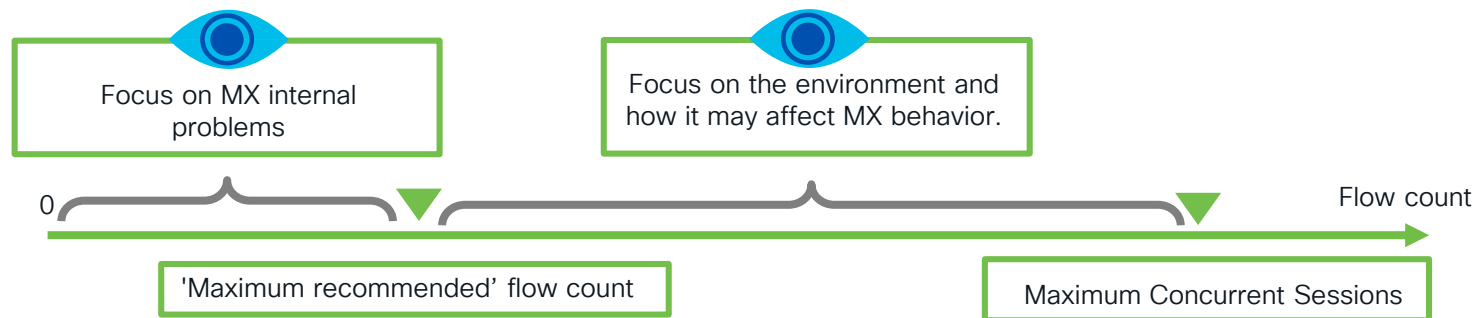
Closure on MX flow count



- After performing a simple calculation, we can get a table with ' the 'maximum recommended' flow count:

	MX67	MX68	MX75	MX85	MX95	MX105	MX250	MX450
'Maximum recommended' flow count	2500	2500	10000	12500	25000	37500	100000	500000

- Now a one-axis graph for the flow count can be created:




MX Recommended number of tunnels

- When it comes to the tunnel counts, we are given two numbers.

	MX67	MX68	MX75	MX85	MX95	MX105	MX250	MX450
Maximum Site to Site VPN Tunnel Count	50	50	75	200	500	1,000	3,000	5,000
Recommended Maximum Site to Site VPN Tunnel Count	50	50	75	100	250	500	1,000	1,500

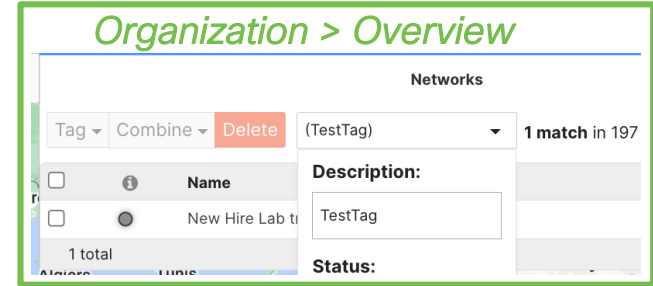
- ‘Maximum Site-to-Site VPN Tunnel Count’ represents the maximum number of AutoVPN and non-Meraki VPN (NMVPN) tunnels that the given model can establish.
- The ‘Recommended Maximum Site-to-Site VPN Tunnel Count’ represents the number of tunnels that MX can handle without a severe increase in utilization.



The math behind tunnel count – NMVPN

To get the number of NMVPN peers for the given MX:

- Locate the MX and check the Tags assigned to it.
- Filter NMVPN peers by the Tags assigned to the MX.
- If no tags are configured, all MXs establish tunnels to all NMVPN peers.



The math behind tunnel count – Auto VPN

- Typically, Hubs are appliances where it's essential to control the number of tunnels.
- By default, hubs establish tunnels to all other hubs and all attached spokes.
- The configuration of 'Active-Active AutoVPN' affects the actual tunnel count.

Security & SDWAN > SD-WAN & traffic shaping

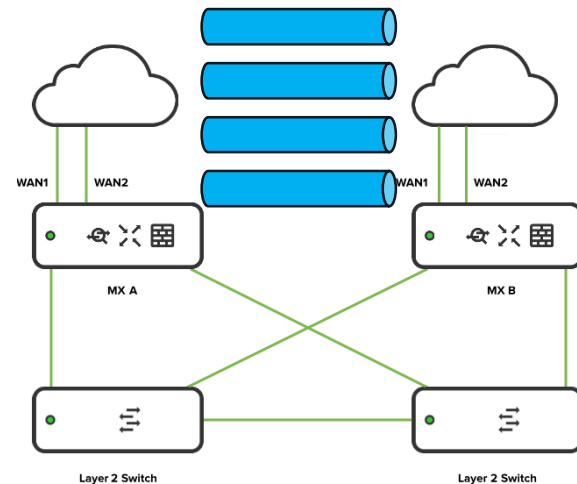
Active-Active AutoVPN

☒ Enabled

Create VPN tunnels over all of the available uplinks (primary and secondary).

☐ Disabled

Do not create VPN tunnels over the secondary uplink unless the primary uplink fails.



The math behind tunnel count Auto VPN


- The [formula](#) below can be used for the calculation. 'S' is the number of spokes, 'H' is the number of Hubs, and 'L' is the number of links on each device.



Formula: Hub Tunnels = $[S + (H-1)] * 2 * L$


- For our example:

$$[100 + (2-1)] * 2 * 2 = 404$$



Performance problems – from theory to practice





Gmail Images

1-BRKTRS-2007-CPU-User.mp4

Google

Search Google or type a URL



It works



Cisco



It works



IIS Windows ...



Web Store



Add shortcut

Customize Chrome

**Welcome to your new starter screen**

We'd love your feedback! Just use the tag on the right side of every new screen or for your full list of networks, check out the "Networks" table below.

Organization Summary New

Devices

2-BRKTRS-2007-CPU-Investigation-1.mp4

Uplinks 8 total**2**Offline ❌**WAN Appliances** 6 total**All**Online ✅**Switches** 8 total**All**Online ✅

Networks

Usage and clients over the last week

Status ▾

Network Type ▾

Tags ▾

6 networks

<input type="checkbox"/>	🕒	Name	Usage	Clients	Tags	WAN Appliances
<input type="checkbox"/>	✅	Branch-Greenfield-POL	507.95 GB	10	Greenfield	✅ 1
<input type="checkbox"/>	✅	Branch-Greenfield-UK	161 KB	1	EMEA-Spokes	✅ 1
<input type="checkbox"/>	✅	EMEA-1	14.00 TB	2542	—	✅ 1
<input type="checkbox"/>	✅	DC-EMEA-2	1.1 MB	1	—	✅ 1
<input type="checkbox"/>	✅	HQ-Brownfield-POL	2.21 GB	4	Brownfield EMEA-Spokes	✅ 1
<input type="checkbox"/>	✅	HQ-Greenfield-UK	14.50 TB	2908	EMEA-Spokes	✅ 1



Network
DC-EMEA-1 ▾



Network-wide



Security & SD-WAN



Insight



Organization

recently-added

NOTES

FIRMWARE

Up to date

Current version: MX 18.107.2

[Open source licenses](#)

CONFIG

Up to date

Remove appliance from network...

Uplink traffic

800 Kb/s
600 Kb/s
400 Kb/s
200 Kb/s
0 Kb/s



Historical device data for the last week ▾

Connectivity to 8.8.8.8 ▾ ⓘ

Latency

40 ms
30 ms
20 ms
10 ms
0 ms

May 18 00:00 May 18 12:00 May 19 00:00 May 19 12:00 May 20 00:00 May 20 12:00 May 21 00:00

Loss

100 %
25 %
5 %
1 %
0 %

May 18 00:00 May 18 12:00 May 19 00:00 May 19 12:00 May 20 00:00 May 20 12:00 May 21 00:00

2-BRKTRS-2007-CPU-
Investigation-1.mp4

Last login: 1 day ago from 2001:420:4880:1250:656c:28e1:591c:f569 Morrisville, NC

Current session started: about 10 hours ago

Summary Report from the last week

3-BRKTRS-2007-CPU-Investigation-2.mp4

NETWORK(S)
DC-EMEA-1

DEVICE TAG
All devices

SSID
All SSIDs

SHOW TOP RESULTS
10

Export to Excel

Customize report

Clients with high usage

2 clients used more than 100.00 GB
"10.103.0.15" and "10.103.0.105"


Usage stats

TOTAL DATA TRANSFERRED
15.24 TB

TOTAL DATA DOWNLOADED
14.88 TB

TOTAL DATA UPLOADED
366.74 GB

Usage over time



Top devices

Name	Model	# Clients	Usage	% Usage
DC-EMEA-1-MX-1	MX250	2573	15.24 TB	100.00%

Top device models by usage

Model	# Devices	Usage	Average Usage per Device
MX250	1	15.24 TB	15.24 TB

Client stats

TOTAL UNIQUE CLIENTS
2573

AVERAGE # OF CLIENTS PER DAY
1471

AVERAGE USAGE PER CLIENT
6.06 GB

Clients per day



Splash page

NUMBER OF CLIENTS TO REQUEST PAGE
0

NUMBER OF CLIENTS GRANTED ACCESS
0

Top clients by usage

Description	Usage	% Usage
10.103.0.15	15.09 TB	99.06%
10.103.0.105	146.93 GB	0.94%
10.102.0.15	17 KB	< 0.01%
10.102.0.16	12 KB	< 0.01%
10.102.0.14	12 KB	< 0.01%

MX performance best practices

- Monitor the average packet sizes/number of routes/number of tunnels/tunnel usage.
- Monitor device utilization and throughput numbers.
- Do not route LAN traffic through the MX device when L3 distribution exists.
- Monitor the number of clients traffic from which passes specific hubs.
- Bypass the IPS for trusted applications.
- Use Direct Internet Access (DIA) when possible.

Topic / Item	Count	Average	Min Val	Max Val	Rate (ms)	Percent
Packet Lengths	18249	872.86	46	1518	0.1303	100%
0-19	0	-	-	-	0.0000	0.00%
20-39	0	-	-	-	0.0000	0.00%
40-79	7323	69.91	46	79	0.0523	40.13%
80-159	624	124.53	82	158	0.0045	3.42%
160-319	155	258.54	166	314	0.0011	0.85%
320-639	82	404.85	328	618	0.0006	0.45%
640-1279	24	1016.67	646	1278	0.0002	0.13%
1280-2559	10041	1517.92	1290	1518	0.0717	55.02%



Formula: Hub Tunnels = $[S + (H-1)] * 2 * L$

Trusted Traffic Exclusions

To increase network performance, select traffic categories and IP addresses or subnets.

Trusted Applications ⓘ

New in 18.2



Streaming & entertainment

Amazon Video, Google Services, Hulu, Netflix, Pandora Radio, Playst... [View all](#)




Online storage

Box, Dropbox, Google Workspace, Microsoft OneDrive, iCloud [View all](#)


MX oversubscription – Horizontal scaling

In cases when current MX can't cap with the network load, horizontal scaling can be used as a solution. There are two main approaches to horizontal scaling:

- **Distribute load by feature** – for example, by adding a separate MX in the HQ to deal with internet traffic.
- **Distribute load by service** – for example, by advertising /32 routes for the highly loaded services (like Proxy servers) through a dedicated Hub.



Name	VPN mode	Subnet
HQ-Subnet	Enabled	10.252.0.0/16



Name	VPN mode	Subnet
HQ-Proxy	Enabled	10.252.1.100/32

Policy-based routing (PBR) to send traffic sourced from Proxy and destined to 10.1.1.0/24 through the Hub2.

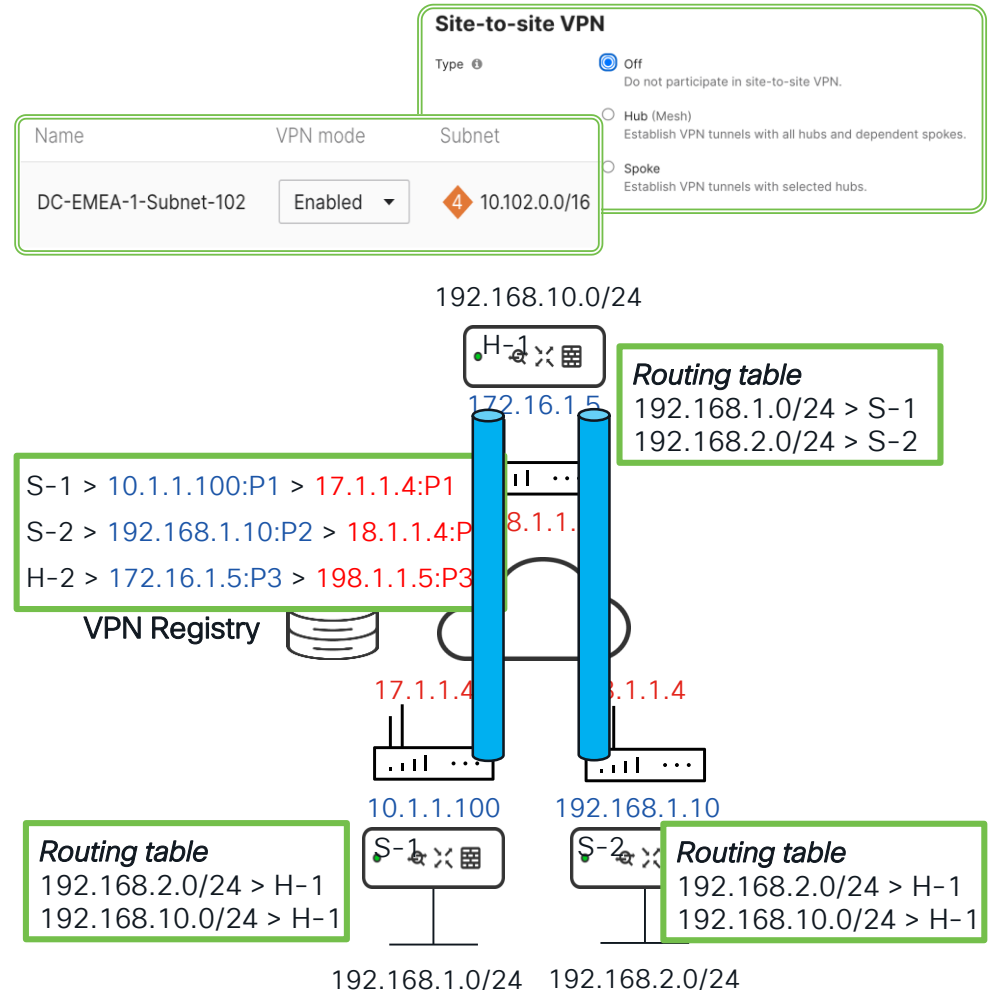


Agenda

- MX performance - Best practices and caveats
- Unobvious Auto VPN
- Gamified troubleshooting case study

Auto VPN is simple ...

- Once S2S VPN is enabled and networks are added, the dashboard pushes routes to MXs.
- MXs contact the VPN registry to report their own public-to-private IP/Port mappings and get mappings for their peers.
- Control plane traffic (Hello messages) starts flowing to activate [NAT hole-punching](#).



Auto VPN is simple until it's not

Auto VPN complexity increases with:

- Growing number of spokes and or networks behind them.
- Growing number of hubs and or networks behind them.
- Introduction of advanced routing (BGP/OSPF)
- Introduction of Data Center Redundancy (DC-DC failover)




Auto VPN – Routing ground rules



For your
reference

- When routes to the *same network* exist through *multiple peers* in *one routing source* (ex: Auto VPN), the *metric (hub priority)* is used to pick the best route.
- When routes to the *same network* exist in *different sources* (ex: Auto VPN and BGP), a *Route Priority is used* to pick the best route (Auto VPN wins over BGP).
- Regular *Auto VPN routes are tracked* (routing is moved to the path with a higher metric, AKA lower priority hub when the path with a lower metric becomes unavailable) *when more than one path is available*.
- BGP routes are tracked using the protocol itself.




Auto VPN – Routing ground rules



For your
reference

- *Hubs* always **send traffic directly** to the **Attached spokes**. Hub to Hub tunnels are used only to route traffic toward networks behind Non-attached spokes and Non-adjacent hubs.
- Due to the [iBGP split horizon](#), when BGP is enabled, spokes always use Auto VPN routes to Adjacent and Non-adjacent spokes.
- OSPF on MX can only advertise Auto VPN routes to neighbors. MX does not learn any routes from OSPF neighbors.



Auto VPN – God from the machine

Auto VPN has multiple features that can enhance the scaling of the SD-WAN fabric. Some are enabled by default, while others can be enabled by Support*. None of those features are customer-visible as of today.

- ***Summarization*** (enabled by default)– Contiguous subnets are summarized to protect small devices. This minimizes the number of routes propagated to spokes by the dashboard.
- ***Track all hub-originated routes*** – This allows tracking to be added on the spoke side for the networks defined on hubs, irrespective of the number of paths.



*—It is recommended that you consult with the Meraki sales engineer before asking for the features to be enabled.

Auto VPN – God from the machine


- ***No spoke routes*** – This feature instructs the dashboard not to install routes available behind other spokes into the spoke config. Enabling this feature requires careful planning, as summary routes will need to be propagated from hubs.
- ***No hub-to-hub tunnels*** – This feature removes tunnels and routes between the hubs. Such change needs to be planned as this will break any connectivity through Auto VPN fabric between Non-adjacent hubs and Non-adjacent spokes.



Auto VPN– from theory to practice (track Hub routes)



Name	VPN mode	Subnet
DC-A-Summary-Subnet	Enabled ▾	10.253.0.0/16
DC-A-HA-Subnet	Enabled ▾	10.254.0.0/16



Name	VPN mode	Subnet
DC-A-Specific-Subnet	Enabled ▾	10.253.0.0/16
DC-A-Standalone-Subn	Enabled ▾	10.254.1.0/24

AutoVPN-
Training-DC-A-HA

AutoVPN-Training-
DC-A-Standalone

5-BRKTRS-2007-Avpn-
Track-noH2H-1.mp4

AutoVPN-
Training-SITE-3

AutoVPN-
Training-SITE-4

10.3.0.0/24
10.3.2.0/24
10.3.3.0/24
10.3.4.0/24
10.3.8.0/24
10.3.255.0/24
DC-SITE-3-CSR

10.4.0.0/24
10.4.2.0/24
10.4.3.0/24
10.4.4.0/24
10.4.8.0/24
10.4.255.0/24
DC-SITE-4-CSR

Routes updated as of Today at 5:43 PM.

IP VERSION

All

SUBNET/PREFIX

Search by subnet/prefix

NAME

Search by name

VLAN

Search by VLAN ID

NEXT HOP

Search by network

DESTINATION

Search by destination

TYPE

All

REPORTED

Current

Stat	Version	Subnet	Name	VLAN	Next hop	Destination
<div></div>	<div></div>	<div><div></div>0.0.0.0/0</div>	Default	—	—	WAN uplink
<div></div>	<div></div>	<div><div></div>10.3.0.0/24</div>	AutoVPN-Training-SITE-3: Single LAN Settings	0	AutoVPN-Training-DC-A-Standalone - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.0.0/24</div>	AutoVPN-Training-SITE-3: Single LAN Settings	0	AutoVPN-Training-DC-A-HA - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.2.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-2	—	AutoVPN-Training-DC-A-Standalone - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.2.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-2	—	AutoVPN-Training-DC-A-HA - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.3.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-3	—	AutoVPN-Training-DC-A-Standalone - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.3.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-3	—	AutoVPN-Training-DC-A-HA - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.4.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-4	—	AutoVPN-Training-DC-A-HA - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.4.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-4	—	AutoVPN-Training-DC-A-Standalone - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.8.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-8	—	AutoVPN-Training-DC-A-HA - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.8.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-8	—	AutoVPN-Training-DC-A-Standalone - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.255.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-255	—	AutoVPN-Training-DC-A-HA - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.3.255.0/24</div>	AutoVPN-Training-SITE-3: SITE-3-Net-255	—	AutoVPN-Training-DC-A-Standalone - appliance	AutoVPN-Training-SITE-3
<div></div>	<div></div>	<div><div></div>10.4.0.0/24</div>	Single LAN Settings	0	10.4.0.1	10.4.0.1
<div></div>	<div></div>	<div><div></div>10.4.2.0/24</div>	SITE-4-Net-2	—	10.4.0.10	—
<div></div>	<div></div>	<div><div></div>10.4.2.0/24</div>	SITE-4-Net-2	—	10.4.0.10	—

Auto VPN– from theory to practice (summarization)





Global Overview



Organization

Lab UK ▾



Network

AutoVPN-Training-SITE-4 ▾



Network-wide



Security & SD-WAN



Insight



Organization

Route table

Rebuild



7-BRKTRS-2007-Avpn-Summary-1.mp4

Routes updated as of Today at 9:01 PM.

IP VERSION

All ▾

SUBNET/PREFIX

Search by subnet/prefix

NAME

Search by name

VLAN

Search by VLAN ID

Stat

Version

Subnet

Name

VL

—

4

0.0.0.0/0

Default

—

—

4

10.3.0.0/24

AutoVPN-Training-SITE-3: Single LAN Settings

0

—

4

10.3.0.0/24

AutoVPN-Training-SITE-3: Single LAN Settings

0

—

4

10.3.2.0/24

AutoVPN-Training-SITE-3: SITE-3-Net-2

—

—

4

10.3.2.0/24

AutoVPN-Training-SITE-3: SITE-3-Net-2

—

—

4

10.3.3.0/24

AutoVPN-Training-SITE-3: SITE-3-Net-3

—

—

4

10.3.3.0/24

AutoVPN-Training-SITE-3: SITE-3-Net-3

—

Global Overview

Organization
Lab UK

Network
AutoVPN-Training-
SITE-4

Network-wide

Security & SD-WAN

Insight

Organization

SITE-4-MX-S

MX64 0c:8d:db:91:c6:c8

PRIMARY



Set a location for this appliance

Add an address below and check Move marker to update its location

ADDRESS

PRIMARY

Current master

SPARE

Unreachable

WAN 1

82.163.125.242

Active

WAN 2

82.163.125.242

Ready

HOSTNAME

autovpn-training-site-4-

8-BRKTRS-2007-Avpn-Summary-2.mp4

Summary

Ports




Historical device data for the last day

Connectivity



Network usage



Auto VPN– from theory to practice (no spoke-to- spoke routes)



Global Overview

Organization
Lab UK

Network
AutoVPN-Training-
SITE-4

Network-wide

Security & SD-WAN

Insight

Organization

SITE-4-MX-S

MX64 0c:8d:db:91:c6:c8

PRIMARY



Set a location for this appliance

Add an address below and check Move marker to
update its location

ADDRESS

PRIMARY

Current master

SPARE

Unreachable

WAN 1

82.163.125.242

Active

WAN 2

9-BRKTRS-2007-Avpn-No-S2S.mp4

Ports




Historical device data for the last day

Connectivity



Network usage





Agenda

- MX performance - Best practices and caveats
- Unobvious Auto VPN
- Gamified troubleshooting case study

And it's time to play a game 😊

- This is like an Escape room, just with the difference that you will not be closed.
- The goal is to find and open a treasure chest.
- The demos in this section have some hints that should help you to progress.



It all started with
an alert



Organization
Ikarem MARS Demo

Network
HQ-Brownfield-POL

Network-wide

Security & SD-WAN

Switching

Wireless

Insight

Organization

10-BRKTRS-2007-Attack-1.mp4

Health
UPLINKS

 1/2 healthy

WAN APPLIANCES

 1/1 healthy

SWITCHES


 2/4 healthy

ACCESS POINTS

 0/1 healthy

Clients





all for the last day



Policy Forget

Search...

6 clients

<input type="checkbox"/>	Status	Description	Last seen	Usage	Client type, OS
<input type="checkbox"/>		00:42:68:72:c8:68	May 28 18:09	None	Other
<input type="checkbox"/>		68:49:92:10:0b:40	May 28 18:09	18.1 MB	Other
<input type="checkbox"/>		6c:de:a9:64:a8:00	May 28 18:09	17.6 MB	Other
<input type="checkbox"/>		6c:de:a9:64:9e:00	May 28 18:09	17.5 MB	Other

Search events

Filter ▾

5 matching events

MX Summary MX Events

Time	Type	Source	Destination	Disposition	Action	Details
May 28 18:07:15	IDS Alert	192.168.0.22	10.9.8.7:1024		Allowed	Alert Check Row 2 Seat 3
May 28 17:56:54	IDS Alert	10.103.0.15:49754	<u>20:cf:ae:1d:33:d1</u>	Blocked	SERVER-WEBAPP	Cisco IOS XE Web UI authentication bypass attempt
May 28 17:56:54	IDS Alert	10.103.0.15:49754	20:cf:ae:1d:33:d1	Blocked	SERVER-WEBAPP	Cisco IOS XE Web UI authentication bypass attempt
May 28 17:56:54	IDS Alert	10.103.0.15:49738	20:cf:ae:1d:33:d1	Blocked	SERVER-WEBAPP	Cisco IOS XE Web UI authentication bypass attempt
May 28 17:56:54	IDS Alert	10.103.0.15:49738	20:cf:ae:1d:33:d1	Blocked	SERVER-WEBAPP	Cisco IOS XE Web UI authentication bypass attempt

< >

1 page

10 ▾

results per page

Filter ▾

5 matching events

MX Summary

MX Events

Time	Type	Source	Destination	Disposition	Action	Details
May 28 18:07:15	IDS Alert	192.168.0.22	10.9.8.7:1024		Allowed	Alert Check Row 2 Seat 3
May 28 17:56:54	IDS Alert	10.103.0.15:49754	20:cf:ae:1d:33:d1		Blocked	SERVER-WEBAPP Cisco IOS XE Web UI authentication bypass attempt
May 28 17:56:54	IDS Alert	10.103.0.15:49754	20:cf:ae:1d:33:d1		Blocked	SERVER-WEBAPP Cisco IOS XE Web UI authentication bypass attempt
May 28 17:56:54	IDS Alert	10.103.0.15:49738	20:cf:ae:1d:33:d1		Blocked	SERVER-WEBAPP Cisco IOS XE Web UI authentication bypass attempt
May 28 17:56:54	IDS Alert	10.103.0.15:49738	20:cf:ae:1d:33:d1		Blocked	SERVER-WEBAPP Cisco IOS XE Web UI authentication bypass attempt



1 page

10 ▾

results per page

12-BRKTRS-2007-Attack-3.mp4

Packet capture

Security appliance: DC-EMEA-1-MX-1

Interface: Internet

Output: Download .pcap file (for Wireshark)

Duration (secs): 60

Filter expression: ether host 00:50:56:b7:ef:47

File name: DC-EMEA-1_MX-DC-EMEA-1-MX-1_IF-Interne

[Start capture](#)

Sample filter expressions

host 10.1.27.253

packets to and from ip address 10.1.27.253

host 10.1.27.253 and port 53

packets to and from ip address 10.1.27.253 and TCP or UDP port 53 (DNS)

icmp[icmptype] != icmp-echo and icmp[icmptype] != icmp-echo-reply

all ICMP packets that are not echo requests/replies (i.e., not ping packets):

ether host 11:22:33:44:55:66

packets to and from ethernet host 11:22:33:44:55:66

pppoe and ip

IP packets encapsulated in PPPoE (Point-to-Point Protocol over Ethernet)

See more [examples](#).

The maximum packet capture duration is 3600 seconds.

This capture will stop after 60 seconds, or when 100,000 packets have been captured.

[Packet capture logs](#)

13-BRKTRS-2007-Attack-4.mp4


Get your gift after the session

- Pins with unique designs are available for pickup at the stage.
- They made a long trip to get here.
- Please don't make them travel back to Poland.



Key takeaway

The ability to see the bigger picture matters during design maintenance and troubleshooting.




Please share your Thoughts

- Don't forget to feel your session survey.
- Let me know in the comments what you like/didn't like
- Share ideas for the subsequent iterations.
- Live your email in the comments, and I'll happily follow up.



Continue your education (Meraki)

- Visit the Meraki booth in TAC clinics to get your support-related questions answered.
- LABMER-1002 - Defending Against Today's and Tomorrow's Threats with Meraki-Talos-Umbrella Integration.
- LABMER-1101 - Demystifying Auto VPN with Cisco Meraki.



Complete Your Session Evaluations



Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to **win 1 of 5 full conference passes** to Cisco Live 2025.



Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.



Level up and earn **exclusive prizes!**



Complete your surveys in the **Cisco Live mobile app**.

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

Contact me at: skuchere@cisco.com



The bridge to possible

Thank you

CISCO *Live!*

#CiscoLive