AWS vMX-XL Deployment Guide

Overview

More and more customers are moving workloads to the cloud, requiring higher throughputs and greater scale. The new vMX-XL provides;

- A highly scalable new MX control and data plane architecture
- Starting with vMX for new next-gen firmware
- Increased performance to 10Gbps+ and up to 10,000 tunnels on a single instance*
- Highly cost effective by using C6in.2xlarge instance type for compute

This document is a walkthrough for setting up a virtual MX (vMX) appliance in the Amazon Web Services (AWS) Marketplace using the latest next-gen firmware. After completing the steps outlined in this document, you will have a virtual MX appliance running in the AWS Cloud that serves as an Auto VPN termination point for your physical MX devices.

Key Concepts

Before deploying a vMX, it is important to understand several key concepts:

Concentrator Mode

All MXs can be configured in either NAT or VPN concentrator mode. There are important considerations for both modes. For more detailed information, see the article on concentrator modes.

One-Armed Concentrator

In this mode, the MX is configured with a single Ethernet connection to the upstream network. All traffic will be sent and received on this interface. This is the only supported configuration for MX appliances serving as VPN termination points into the AWS Cloud.

NAT Mode Concentrator

In this mode, any traffic coming over auto-VPN or client VPN to the vMX will be NATed to the vMX's IP as it egresses the vMX. Other capabilities of the NAT mode including DHCP, HA or multiple ports (LAN and WAN) are not supported. In each mode the vMX is still a one-armed appliance with one network interface.

> ⓘ Currently NAT mode is not support on the vMX-XL

VPN Topology

There are several options available for the structure of the VPN deployment.

Split Tunnel

In this configuration, branches will only send traffic across the VPN if it is destined for a specific subnet that is being advertised by another MX in the same dashboard organization. The remaining traffic will be checked against other available routes, such as static LAN and third-party VPN routes, and if not matched will be NATed and sent out the branch MX unencrypted.

Full Tunnel

In full-tunnel mode, all traffic that the branch or remote office does not have another route to is sent to a VPN hub.

AWS Cloud Terminology

This document will make reference to several key AWS-specific terms and concepts.

AWS

Amazon Web Services is a hosting platform that allows organizations to host infrastructures and services in the cloud.

VPC

A virtual private cloud is a virtual and private network within the AWS infrastructure. A VPC has a block of associated IP addresses, which can be subdivided into multiple smaller subnets.

EC2

Elastic compute cloud is Amazon's virtual and cloud-based computing infrastructure. EC2 allows you to run virtual machines within your VPC. A virtual machine running on the EC2 platform is commonly referred to as an EC2 instance.

Additional Information

Please refer to the AWS glossary for a dictionary of cloud terminology on the AWS Cloud platform.

Meraki Dashboard Configuration

Begin by creating a new Security Appliance network in your organization. If needed, please refer to the guide on creating a new network in the Meraki dashboard.

1. Add license(s) to the Meraki dashboard

To complete the vMX Meraki dashboard configuration, a vMX license **must** be available for use in your organization.

If your organization has already reached its vMX license limit, you will be unable to create new vMX networks until a vMX network is deleted or additional vMX licensing added.

If you do not have access to a vMX license or require additional vMX licenses, please reach out to your Meraki reseller or sales representative.

2. Create a "Security appliance" network type:

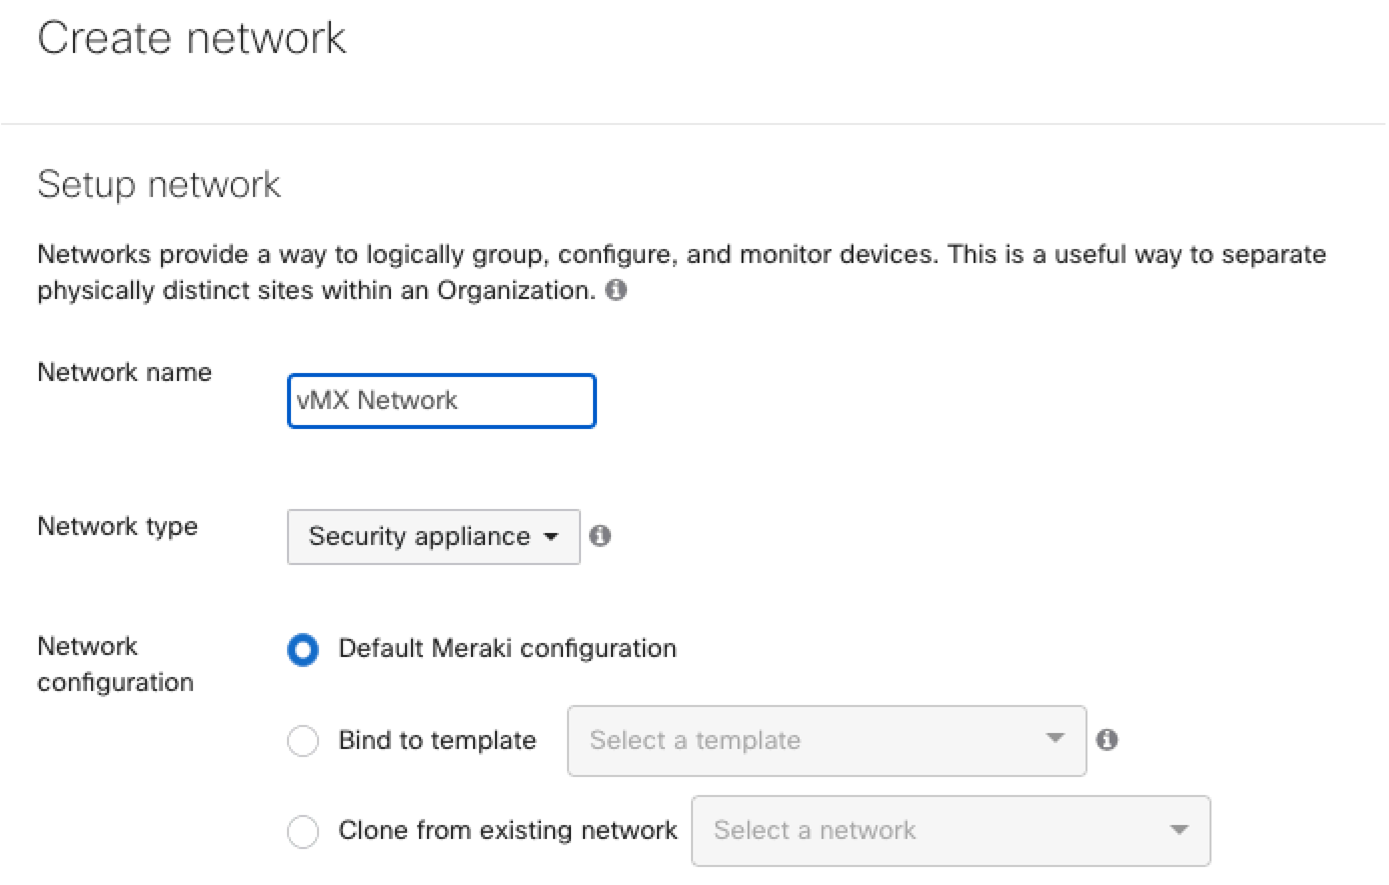

Create network

Setup network

Networks provide a way to logically group, configure, and monitor devices. This is a useful way to separate physically distinct sites within an Organization.

Network name: vMX Network

Network type: Security appliance

Network configuration:
● Default Meraki configuration
○ Bind to template — Select a template
○ Clone from existing network — Select a network

3. Assign vMX type to network

Once you have created the "Security appliance" network and added the appropriate license you will be able to deploy a new vMX to your network by clicking on the 'Add vMX-XL' button.

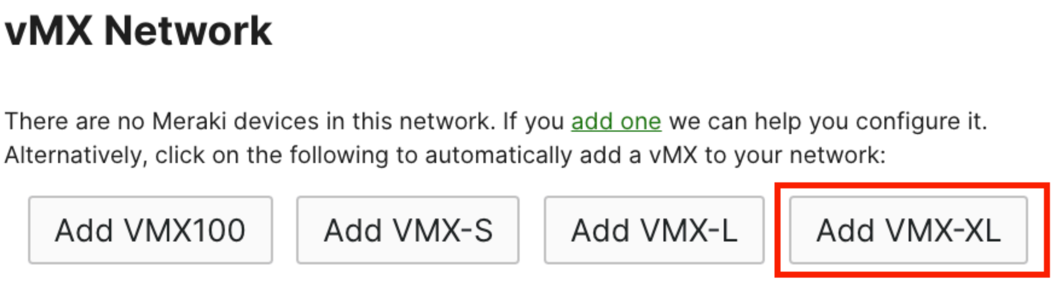

vMX Network

There are no Meraki devices in this network. If you add one we can help you configure it. Alternatively, click on the following to automatically add a vMX to your network:

4. Generate the authentication token

Before generating the token, please verify the firmware is configured for TBD. If the vMX network firmware is set to anything below that, the upgrade will not occur.

After you add the new vMX to your network, navigate to **Security & SD-WAN > Monitor > Appliance status** and select “Generate authentication token” to generate the token for the AWS user-data field.

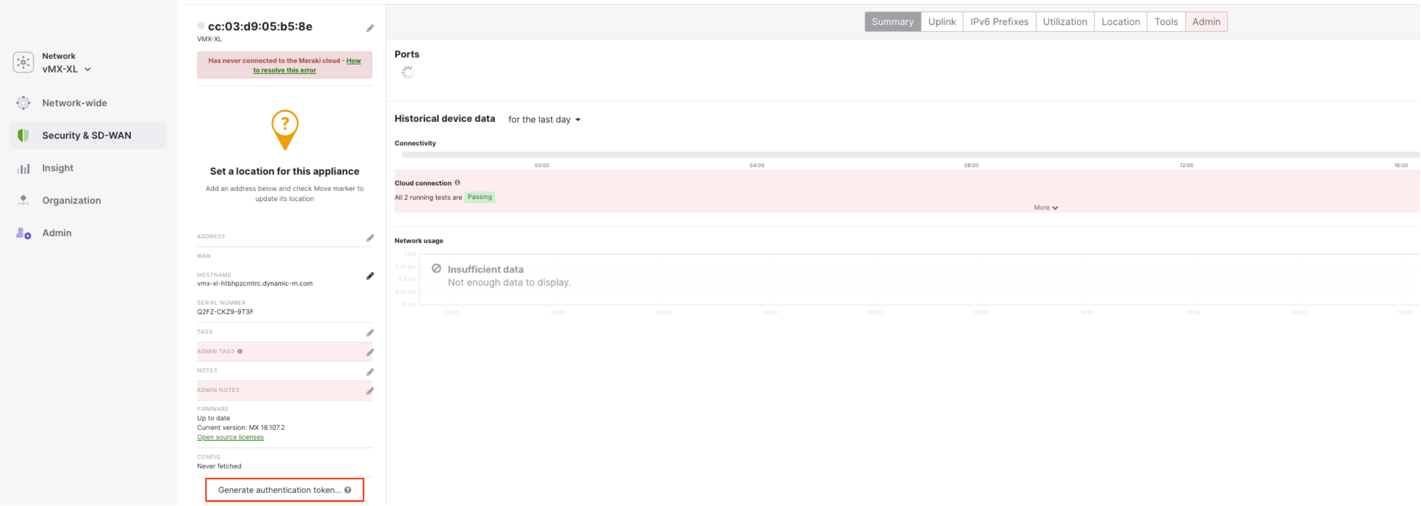

cc:03:d9:05:b5:8e
VMX-XL

Has never connected to the Meraki cloud - How to resolve this error

Set a location for this appliance

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

ADDRESS

WAN

HOSTNAME
vmx-xl-htbhpzcmtrc.dynamic-m.com

SERIAL NUMBER
Q2FZ-CKZ9-9T3F

TAGS

ADMIN TAGS

NOTES

ADMIN NOTES

FIRMWARE
Up to date
Current version: MX 18.107.2
Open source licenses

CONFIG
Never fetched

Cloud connection
All 2 running tests are Passing

Insufficient data
Not enough data to display.

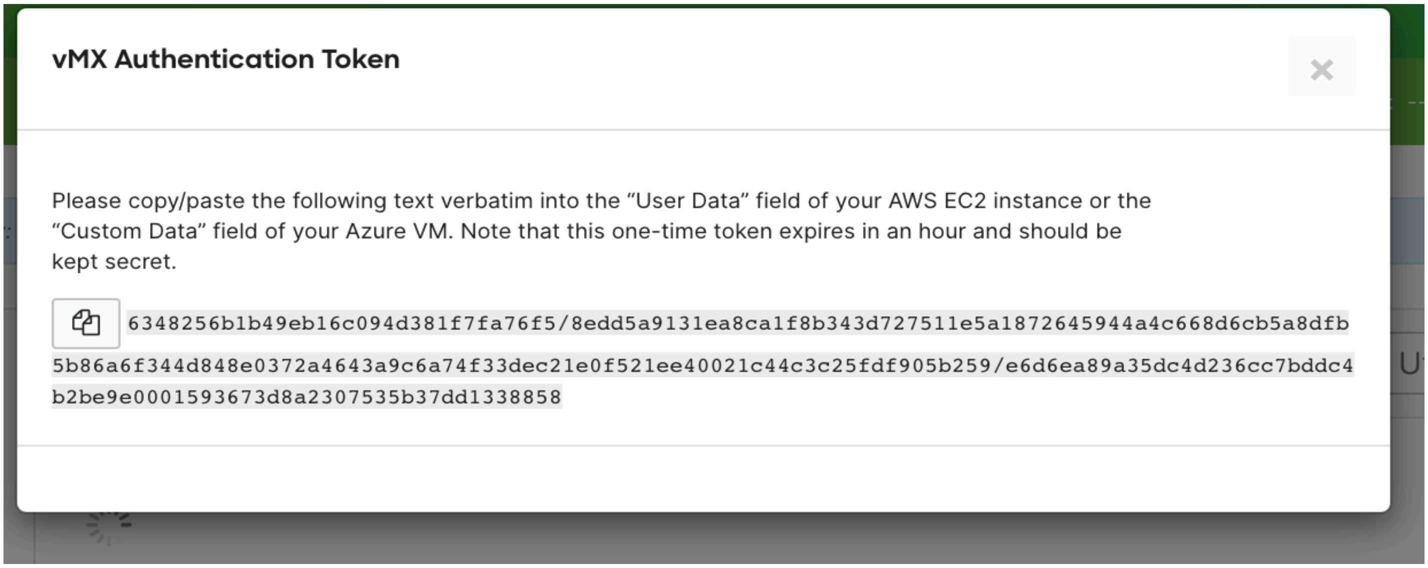

vMX Authentication Token

Please copy/paste the following text verbatim into the “User Data” field of your AWS EC2 instance or the “Custom Data” field of your Azure VM. Note that this one-time token expires in an hour and should be kept secret.

```
6348256b1b49eb16c094d381f7fa76f5/8edd5a9131ea8ca1f8b343d727511e5a1872645944a4c668d6cb5a8dfb
5b86a6f344d848e0372a4643a9c6a74f33dec21e0f521ee40021c44c3c25fdf905b259/e6d6ea89a35dc4d236cc7bddc4
b2be9e0001593673d8a2307535b37dd1338858
```

5. Copy the newly generated token and save it.

The newly generated token will be used in the "System Configurations" section when creating a new AWS Cloud instance.

The authentication token **must** be entered into the AWS Cloud instance within one hour of generating it, otherwise a new token must be generated.

AWS Setup

This section walks you through configuring the necessary requirements within AWS and adding a vMX instance to your virtual private cloud (VPC). For more details on setting up a VPC and other components, refer to Amazon's AWS Documentation.

Accessing the AMI

Click to access the AMI. A screenshot of the AWS Marketplace listing is included below, please note that the vMX-XL offer is different from the marketplace offer used to deploy vMX-S/M/L.

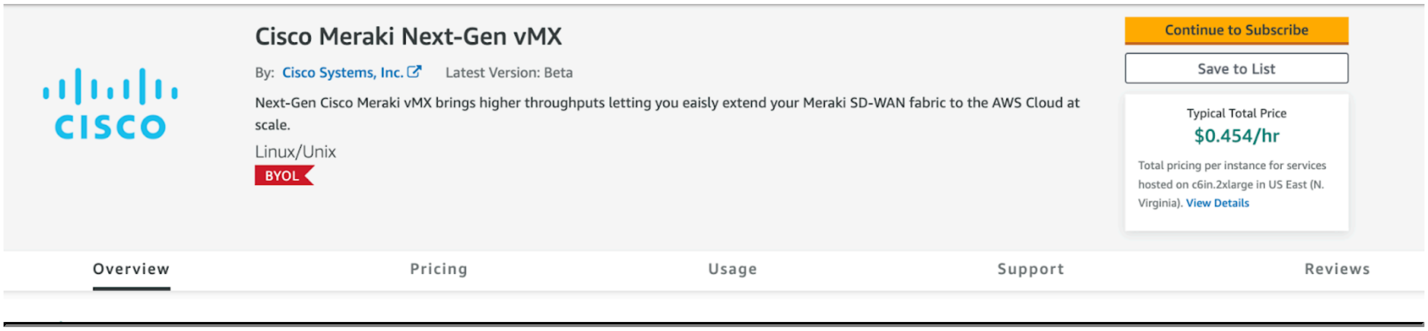

Cisco Meraki Next-Gen vMX

By: Cisco Systems, Inc. Latest Version: Beta

Next-Gen Cisco Meraki vMX brings higher throughputs letting you easily extend your Meraki SD-WAN fabric to the AWS Cloud at scale.

Linux/Unix

BYOL

Typical Total Price
$0.454/hr

Total pricing per instance for services hosted on c6in.2xlarge in US East (N. Virginia). View Details

Select the EC2 instance type and the region to launch the EC2 instance in. This should match the availability zone your VPC resides in.

The recommended instance type for best performance is `c6in.2xlarge`

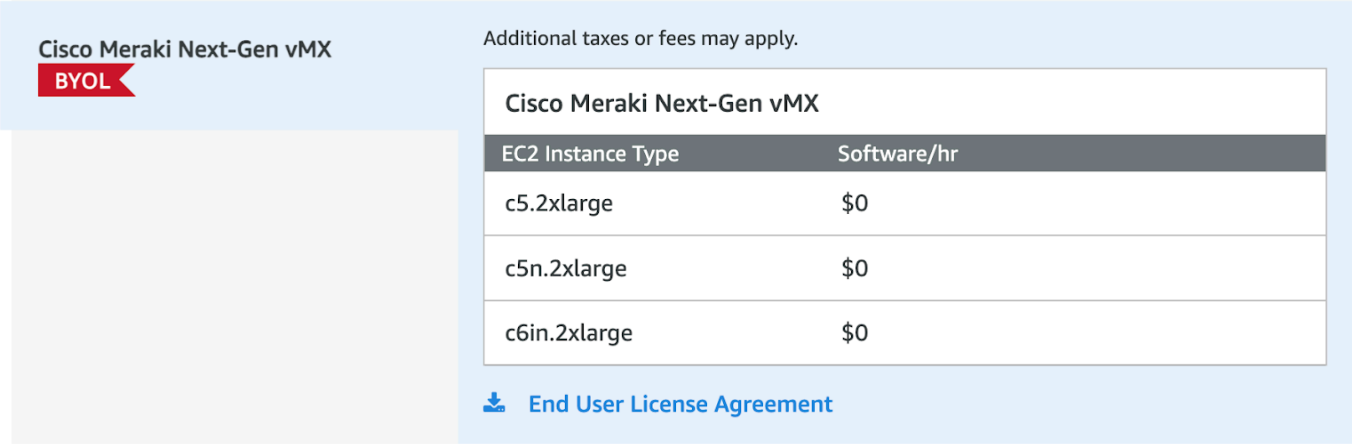

Cisco Meraki Next-Gen vMX

BYOL

Additional taxes or fees may apply.

Cisco Meraki Next-Gen vMX

EC2 Instance Type	Software/hr
c5.2xlarge	$0
c5n.2xlarge	$0
c6in.2xlarge	$0

Click the "Continue to Subscribe" button to configure and finalize the vMX deployment.

Configuring the EC2 Image

After continuing, you will be prompted to configure the EC2 instance settings.

For "Choose action" option select the "Launch through EC2" and click "Launch."

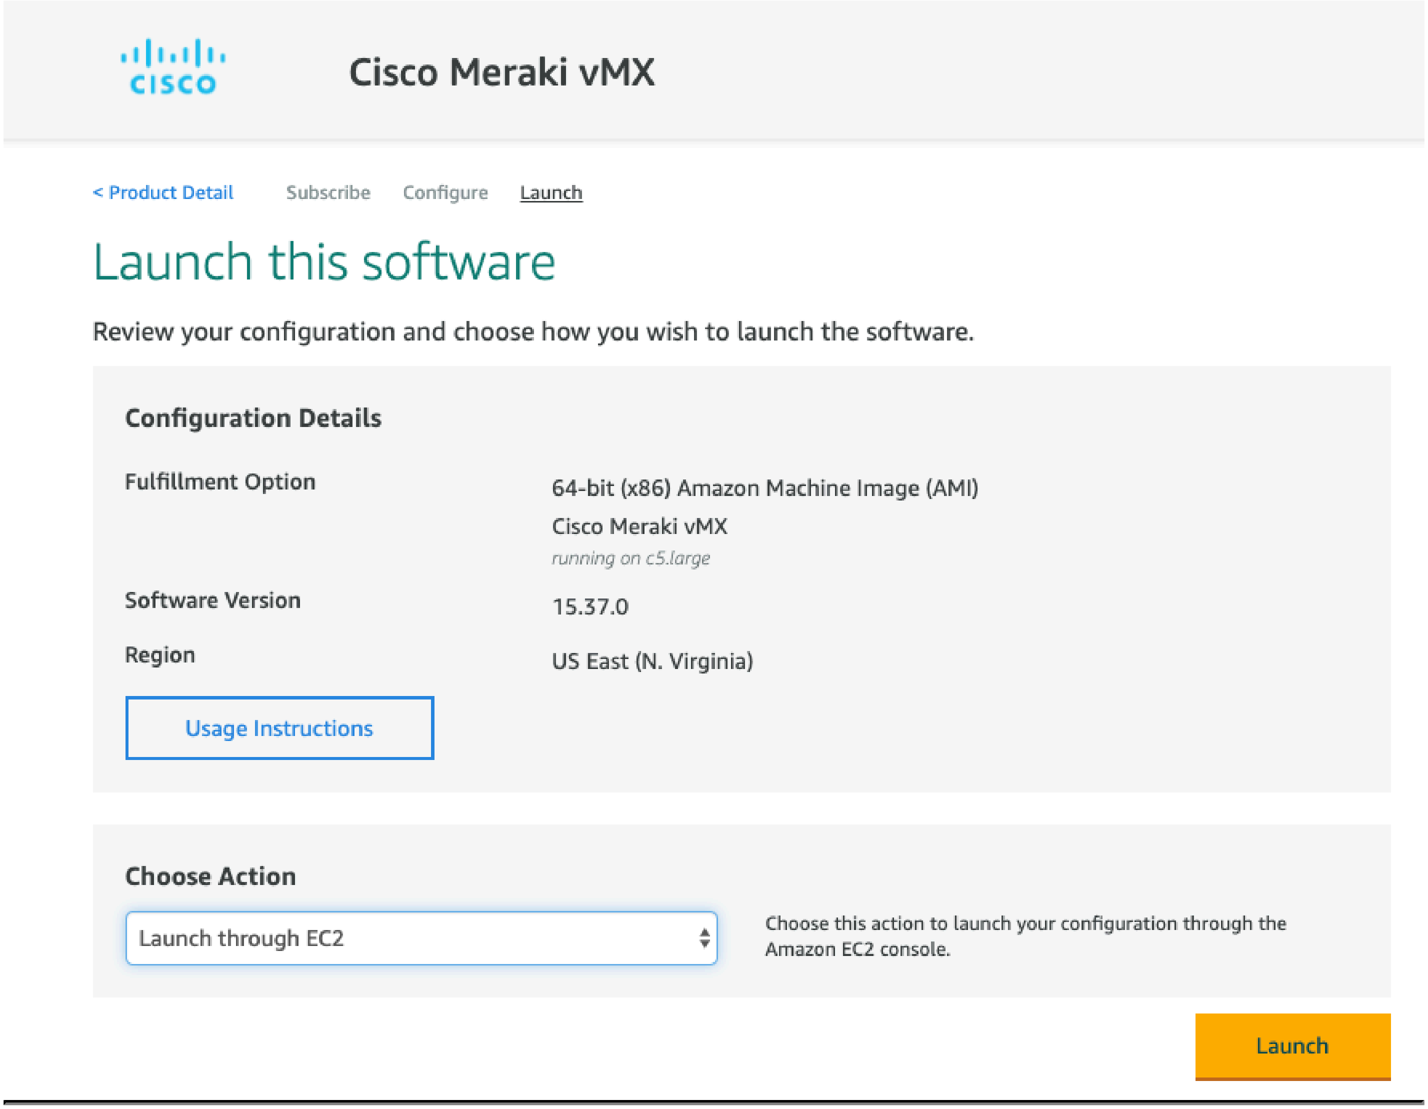

Launch this software

Review your configuration and choose how you wish to launch the software.

Configuration Details

Fulfillment Option: 64-bit (x86) Amazon Machine Image (AMI)
Cisco Meraki vMX
running on c5.large

Software Version: 15.37.0

Region: US East (N. Virginia)

Choose Action

Launch through EC2

Choose this action to launch your configuration through the Amazon EC2 console.

Choose the EC2 instance type

ⓘ C6in.2xlarge is the recomended instance type for optimal performance.

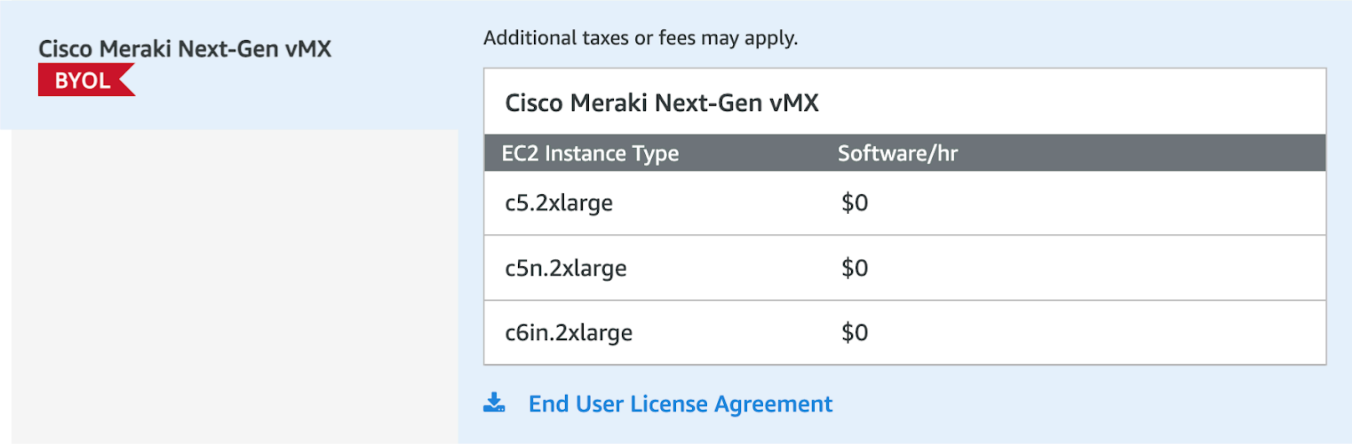

Cisco Meraki Next-Gen vMX

BYOL

Additional taxes or fees may apply.

Cisco Meraki Next-Gen vMX

EC2 Instance Type	Software/hr
c5.2xlarge	$0
c5n.2xlarge	$0
c6in.2xlarge	$0

End User License Agreement

Scroll down to the network settings.

Select the VPC and the subnet the instance will be a part of and make sure the "auto-assign public IP" is enabled.

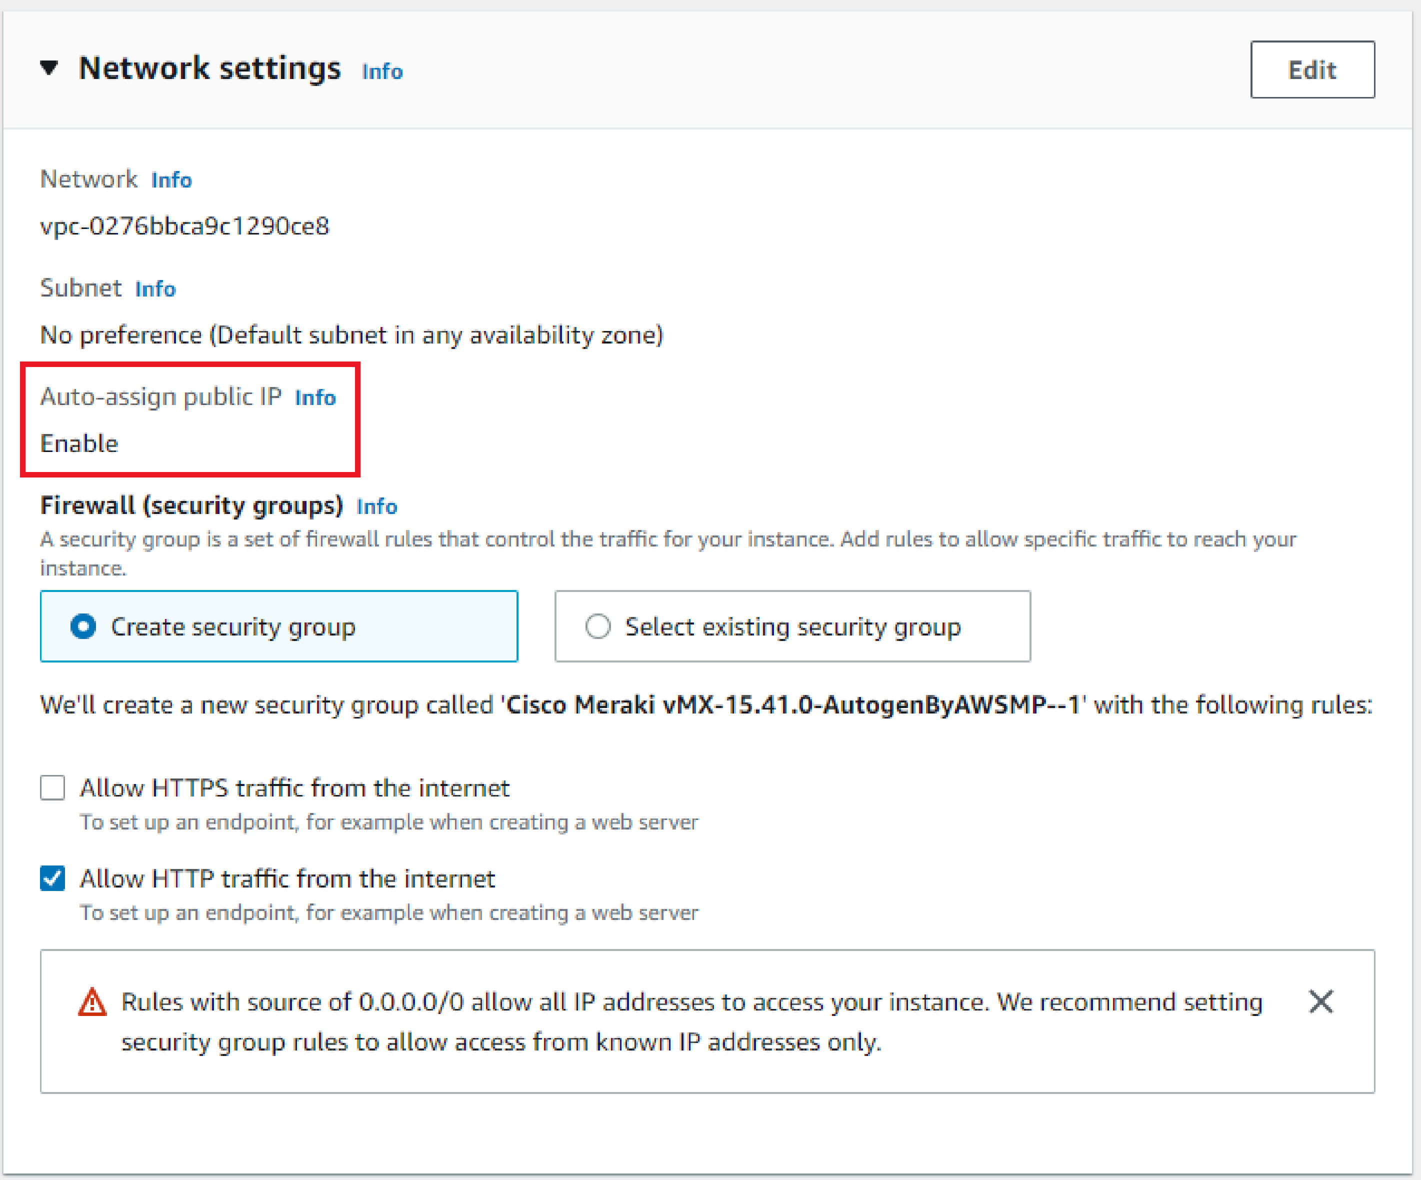

▼ Network settings Info

Edit

Network Info

vpc-0276bbca9c1290ce8

Subnet Info

No preference (Default subnet in any availability zone)

Auto-assign public IP Info

Enable

Firewall (security groups) Info

A security group is a set of firewall rules that control the traffic for your instance. Add rules to allow specific traffic to reach your instance.

● Create security group ○ Select existing security group

We'll create a new security group called '**Cisco Meraki vMX-15.41.0-AutogenByAWSMP--1**' with the following rules:

□ Allow HTTPS traffic from the internet
To set up an endpoint, for example when creating a web server

☑ Allow HTTP traffic from the internet
To set up an endpoint, for example when creating a web server

⚠ Rules with source of 0.0.0.0/0 allow all IP addresses to access your instance. We recommend setting security group rules to allow access from known IP addresses only. ✕

Scroll down to "Advanced Details" enter the vMX authentication token from the dashboard in the user data field.

> The authentication token **must** be entered into the AWS Cloud instance within **one hour** of generating it, otherwise a new token must be generated.
>
> Steps that need to be taken to generate the vMX authentication can be found in the Dashboard Configuration section

> Metadata version 2 is not supported at this time. Using Metadata V2 can result in 'Invalid Token' error message.

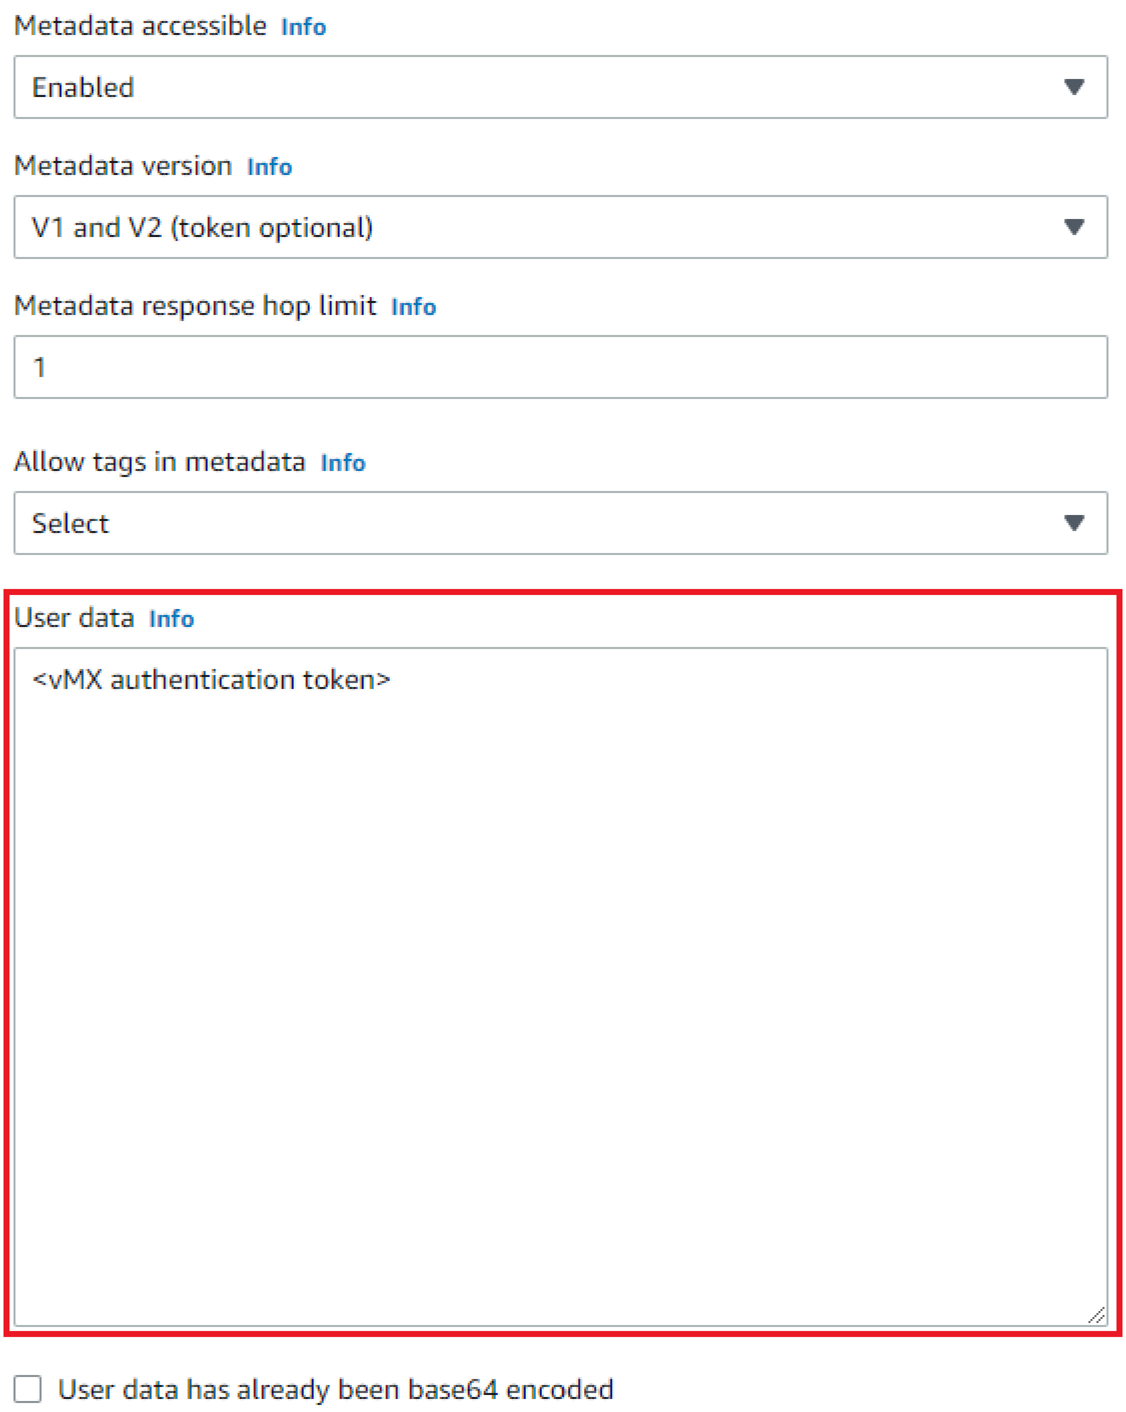

Metadata accessible

Enabled

Metadata version

V1 and V2 (token optional)

Metadata response hop limit

1

Allow tags in metadata

Select

User data

```
<vMX authentication token>
```

□ User data has already been base64 encoded

ⓘ **Note:** the carrot "<>" brackets are to be omitted from your vMX auth token

Click "Review and Launch" to finish creating the instance.

ⓘ It may take several minutes before the software subscription completes and the instance launches.

Additional VPC Configuration

The virtual MX appliance will allow for site-to-site VPN connectivity using Auto VPN between AWS and other remote MXs. In order to have proper bidirectional communication between remote subnets that are terminating into AWS via the vMX and hosts within AWS, the VPC routing table must be updated for the remote Auto VPN-connected subnets.

To do this, navigate to the "VPC" dashboard in AWS, and follow the steps below:

1. Click on "Route Tables" on the left-hand side.
2. Find the route table associated with the VPC or subnet the vMX is hosted in (you can quickly filter the list by typing in your VPC ID or name in the search box).

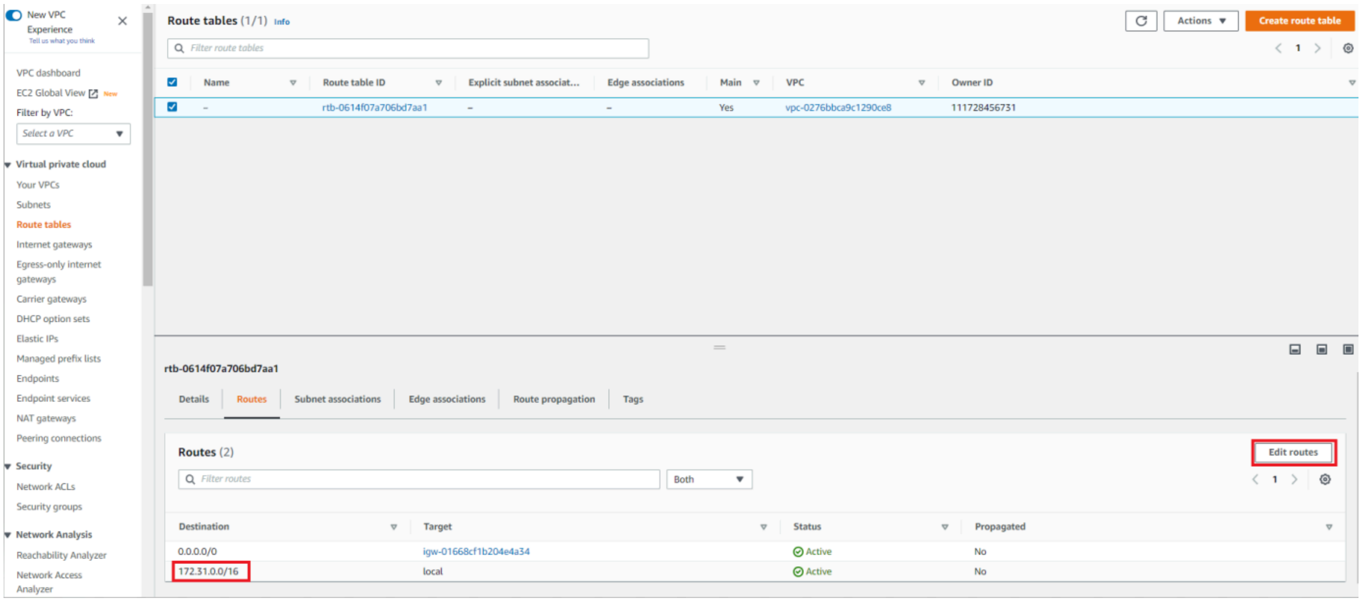

☑	Name	Route table ID	Explicit subnet associat...	Edge associations	Main	VPC	Owner ID
☑	–	rtb-0614f07a706bd7aa1	–	–	Yes	vpc-0276bbca9c1290ce8	111728456731

rtb-0614f07a706bd7aa1

Routes (2)

Destination	Target	Status	Propagated
0.0.0.0/0	igw-01668cf1b204e4a34	Active	No
172.31.0.0/16	local	Active	No

3. Click on the route table in the list to select it, then click on the "Routes" tab in the details pane below the list of routing tables.
4. Click the "Edit" button.
5. Click the "Add another route" button at the bottom of the route table.
6. In the "Destination" column, add the routes available via Auto VPN.
7. In the "Target" column, select the vMX instance or interface ID.
8. Repeat steps 5-7 for each network available via Auto VPN and Client VPN if applicable. You may want to use a summary address. Below is an example of one network available via Auto VPN.

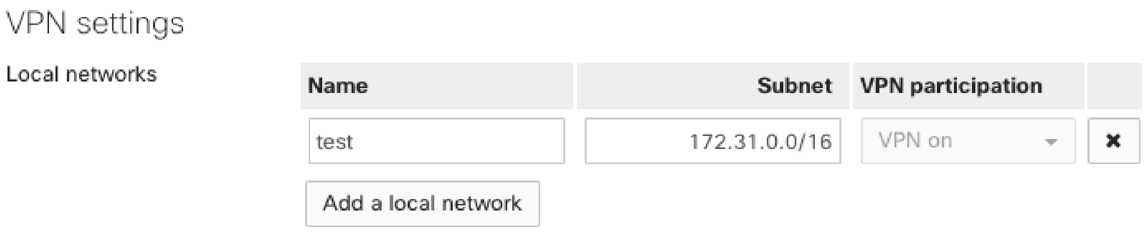

VPN settings

Local networks

Name	Subnet	VPN participation	
test	172.31.0.0/16	VPN on	✖

AWS vMX Setup Video

Troubleshooting

The most common problem people face when deploying a vMX is getting it provisioned and online in their Meraki dashboard in the first place. New troubleshooting/diagnosis messages have been added to the vMX console screen now so you can identify what went wrong during the provisioning process.

When the vMX boots it will execute the following steps during its initial provisioning process:

1. Obtain user-data (vMX auth token)
2. Authenticate with the dashboard (using auth token)
3. Connect to dashboard

Obtain User Data and Authenticate to Dashboard

Once a vMX has successfully connected to a network, it will then attempt to obtain its user data (vMX auth token). There are different user-data mechanisms in each platform that the vMX currently runs on to provide the token to the vMX. In AWS, Azure, GCP, and Alicloud there are user-data fields in the VM configuration where this can be provided.

Token Expired

vMX auth tokens have a lifetime of only one hour for security purposes. If you see the following message on your vMX console it means the token you provided is no longer valid.

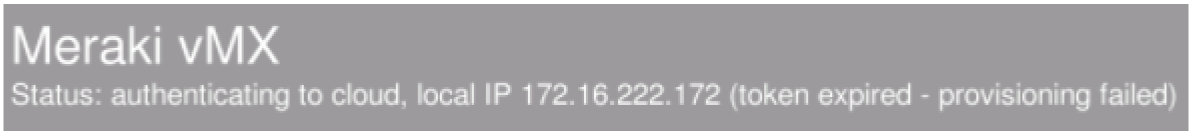

```
Meraki vMX
Status: authenticating to cloud, local IP 172.16.222.172 (token expired - provisioning failed)
```

Invalid Token

If the token provided is incorrect in any way the "*invalid token*" message is displayed on the console.

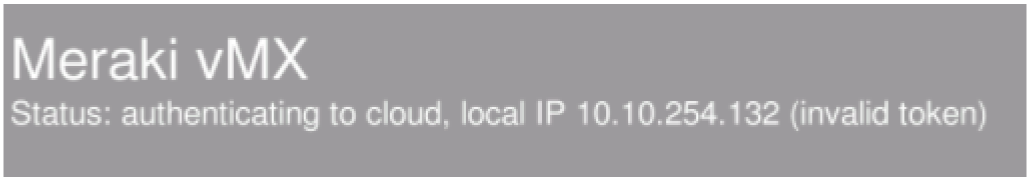

```
Meraki vMX
Status: authenticating to cloud, local IP 10.10.254.132 (invalid token)
```

Unable to Reach Meraki Dashboard

If the vMX is unable to reach the dashboard on TCP port 7734 then the initial provisioning phase will fail and an "Unable to reach Meraki Dashboard" message will be displayed on the console (check the firewall information page for a list of all the firewall rules needed for the Cisco Meraki cloud communication). Please refer to the document on correct ports/IPs that need to be opened for Meraki dashboard communication.
Note: Please ensure that the VPC route table is able to get out to the internet.

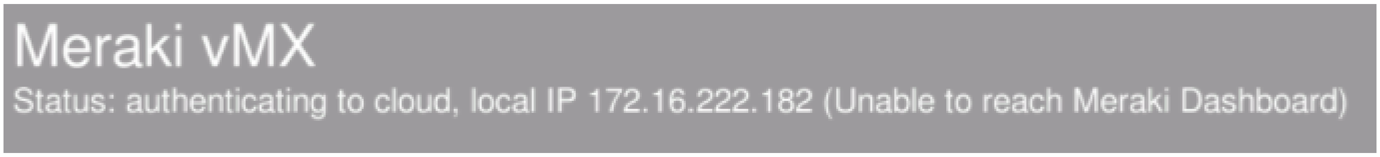

```
Meraki vMX
Status: authenticating to cloud, local IP 172.16.222.182 (Unable to reach Meraki Dashboard)
```

ⓘ **Note:** For any of the following above (Token Expired, Invalid Token, Unable to reach Meraki Dashboard), you may also look to view the instance screenshot, right-click on EC2 instance > Monitor and troubleshoot > Get instance screenshot to view the status. Other resources from AWS:

Getting a screenshot of an unreachable instance

EC2 Instance Console Screenshot

No "Add vMX" Button

When navigating to **Security & SD-WAN > Monitor > Appliance Status**, if there is no "Add vMX" button, please ensure the following two conditions are met:

1. You have available vMX licenses in your license pool.
2. You have created a "Security appliance" network type.

ⓘ Please note that Meraki support does not troubleshoot AWS Cloud-specific firewall rules and deployments.