Meraki: A study in a one-arm MX Auto VPN concentrator deployment

1. Preface

This case study describes the one-arm MX Auto VPN deployment on Cisco Meraki MX Firewalls. The documentation on the Meraki website leaves a lot to be desired in my opinion, I just did not find any good descriptions of this deployment.

Hence, I decided to make my own lab and see how it works. And it works... but there are a number of remarkable “features” (pitfalls?) to this deployment.
I’ll try to describe them here…

2. The topology

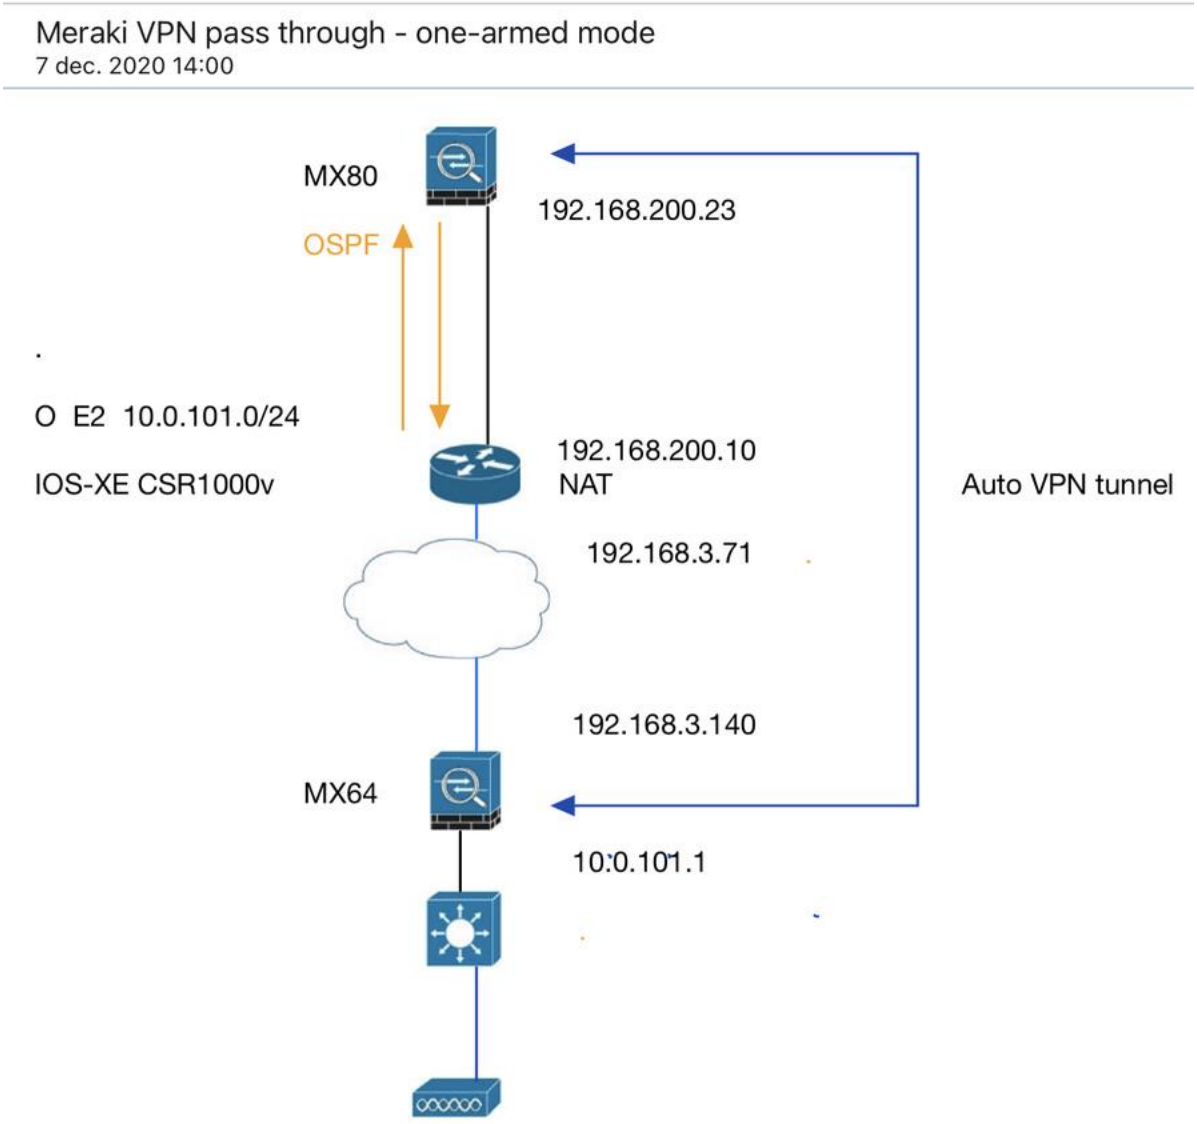

3. Explanation and background of the Topology and the deployment

Anything above the cloud symbolizes a remote Datacenter "DC".

The recommended Datacenter deployment for Meraki in the DC is that you do NOT use the Routed / NAT mode for the Meraki Firewall. That means you are going to LOSE OUT on all those goodies like L7 filtering, advanced malware protection etc. etc. Strange..
Instead, you should use the "one-arm" deployment.

The 2nd remarkable feat is that you are advised NOT to deploy the MX in this deployment directly on the Internet, as there are some vulnerabilities known in one-arm deployments. (???!!!)

What is so "one-arm" about this? (sounds like a Harison Ford movie to me..)
The fact that you can ONLY USE 1 WAN INTERFACE AND THAT IS WAN1.
But.. No failovers then?! IN THE DATACENTER?!!!!!
Well.. you can configure a "warm spare" and cluster the MX.
Hmm.. Can I at least use a portchannel / aggregation on the MX?.....
Nope, NO aggregations are NOT available on the MX platform. ☹

Then.. if you are NOT doing all the Routed stuff.. Who is going to perform the NAT / routing to get from the remote DC to our site?
Well.. for that you use a DIFFERENT VENDOR,- Like Cisco. ☺

In the topology above, I have deployed a CSR1000v to perform the NAT translation between the 192.168.200.0/24 network and the "internet" on 192.168.3.0/24. (which of course is impossible because of RFC1918)

Do I need to setup port-forwarding for this on the CSR?..... Here is the cool thing;
In most cases you DO NOT. The MX platform registers it's UDP port-numbers with the VPN registry in the cloud. Then it will UDP hole-punch an opening to the remote side on the UDP port that you learned that it has just opened up for you.

The remote side will come back on your high UDP port number that it just learned from the cloud VPN registry as well.

If you UDP session does not timeout… (!!!!) then you can have a cool Auto VPN connection without needing the extra complexity of port-forwarding. Neat!

If however you do need to port-forward… then things get messy as Auto VPN uses a whole range of UDP ports. (dynamic in nature… Yak ☹ ☹)

4. OSPF..

When you configure the top MX Firewall in VPN concentrator mode, you can set the area, router-id and the cost for the interfaces. (no auto-reference bandwidth…)

Meraki: A study in a one-arm MX Auto VPN concentrator deployment

The MX80 and the CSR1000v are happily becoming neighbors which is cool.
BUT… just how functional is that OSPF implementation on the Meraki side?

The weird thing is:
The Auto VPN mechanism will share the local MX80 network 192.168.200.0/24 with the MX64 automatically provided you select it for VPN usage, via OSPF.

The MX80 in the DC will learn the 10.0.101.0/24 network via OSPF and Auto VPN automatically (again, if selected) and share it with the DC CSR1000v router.
So far so good..

BUT… here's the thing.. **the CSR1000v router CAN NOT PUBLISH ANY ROUTES TO THE DC MX80 via OSPF**! (by design..!)

→ So.. If you want to make ANY other Networks within the DC available across the VPN, then you have to add a static route on the MX80 which in turn will share the route via OSPF with the MX64.
Weird stuff..

Here is an example where I created an IP address on the loopback adapter of the CSR1000v. 10.10.10.10/32 needs to be added to OSPF on the MX80 in order to see it on the MX64..

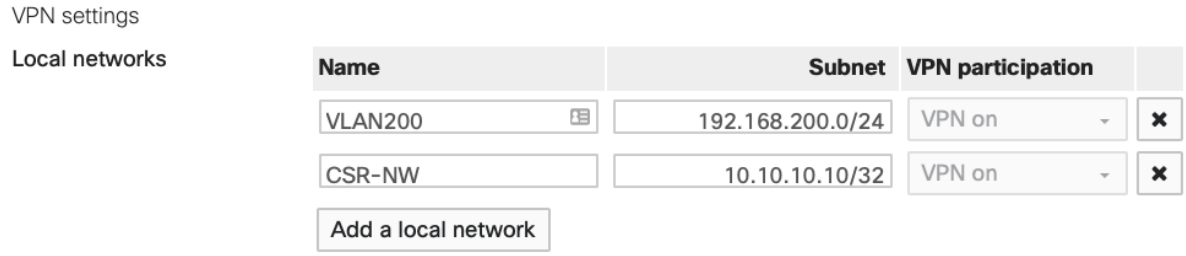

The result in the routing table of the MX64 is:

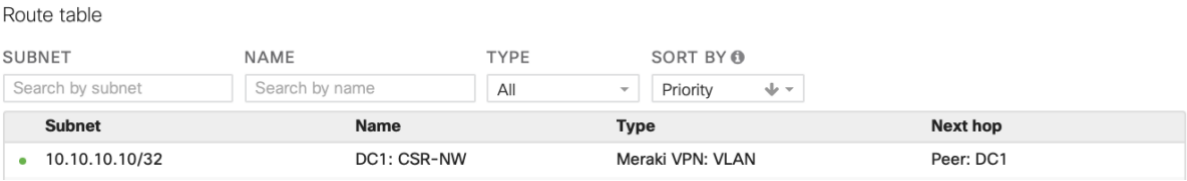

And on the CSR1000v the return – route towards the MX64 has been learned via OSPF:

```
O E2     10.0.101.0/24 [110/1] via 192.168.200.23, GigabitEthernet3
```

The way it works is that the route to the loopback on the CSR is now shared from the MX80 with the MX64.

But how will the MX80 reach the loopback network? Since you did not enter a next-hop?!
It uses the Default-Gateway as next hop for the loopback network.

5. Caveats..

-No L2 redundancy on the MX platform in the form of aggregations / bonds / Portchannels.
-Only WAN1 can be used, even if WAN2 is available.
-Only half working OSPF implementation on the MX in this mode.
-You can combine this deployment with BGP however.. see:
https://documentation.meraki.com/MX/Networks_and_Routing/BGP

-Due to the concentrator mode, you lose out on all the cool L7 fw rules and features. (all the advanced stuff basically)
-You need **EXTRA KIT** on the edge of the DC because the Meraki is not supposed to be safe to deploy it on the Internet edge in this mode. Pfffff…
-Having only 1 interface for both the WAN and LAN TRAFFIC (routing on a stick) means that this solution might prove not to scale so well..
-You can NOT use ANY OF THE LAN INTERFACES ON THE MX!!! It's ONLY THE 1 WAN INTERFACE! Such a waste really…
-You are using a fancy "routing on a stick" implementation, so EVERY FRAME will both ENTER and LEAVE the same 1 interface and there is NO way of getting around this.

The manageability of the dashboard in the above setup is awesome, but I think you'll agree that there is room for improvement…
Oh did I mention that IPv6 is still not supported?...