Meraki MX Firewall High Availability study

1. Topology

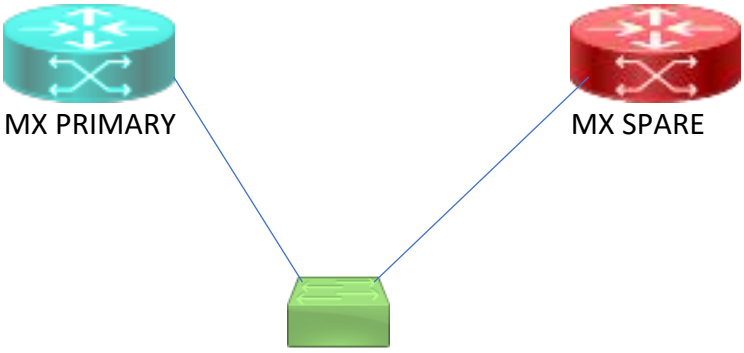

2. Prerequisites for Meraki MX Firewall HA

-Both MXs must be of the same MX model and firmware.
-Only 1 license is needed, as there is only 1 MX.
-The 2nd MX should not be assigned to any Network prior to warm-spare deployment.
-The MXs do NOT need a separate "heartbeat interconnect", just a L2 interconnect will do.
-The Simple VRRP v.3 (protocol 112) is used for the failover configuration.

Make sure STP is enabled on the downstream switching infrastructure, as a properly-configured HA topology will introduce a loop on the network. See the notes under 4.

1. Preparation for the Warm Spare configuration

Here we'll create the Warm Spare HA configuration with – MX 64 Firewalls:

Under ORGANIZATION → Inventory → remove the 2nd MX from the current NW.

2. Configuration of the Warm Spare setup

On the first, active MX :
SECURITY & SD-WAN → MONITOR → Appliance Status → Configure Warm Spare:

Meraki MX Firewall High Availability study

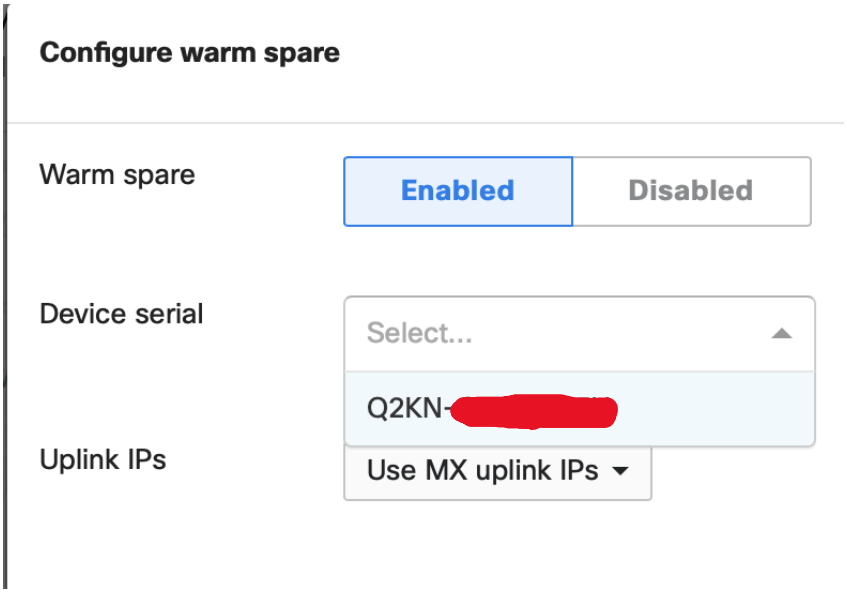

You can either use the standard Internet uplink IP address, or configure VRRP on the WAN side as well.

Click “Update” and you should be in business within a few minutes.

Note: this WILL take a few minutes so be prepared to wait.

After a few minutes:

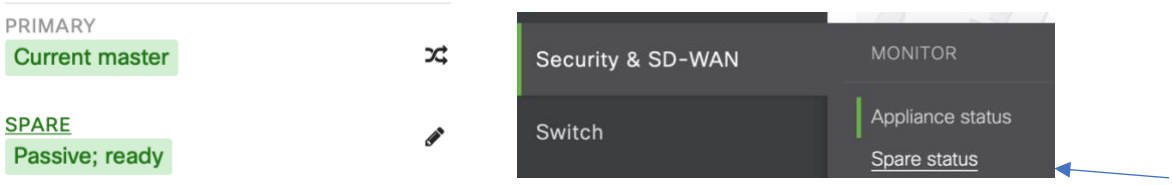

The strange thing is.. the SPARE can currently NOT ping to the internet, despite that you can administer it from the Internet cloud hosted Dashboard; it seems like the whole Routing engine is down as long as you are not the Master. (Juniper SRXs show a similar behavior in a Chassis Cluster)

The Addressing & VLANs page STILL shows the SAME IPs for the PRIMARY MX.
So contrary to eg Cisco VRRP, there are NO IP addresses configured on the Interfaces of the WARM spare! It “just listens for VRRP messages”. The interface IP address is used as VIP address.

What do the VRRP v.3 messages that PRIMARY Master MX1 sends, look like?

Meraki MX Firewall High Availability study

```
    4 0.085819    10.0.11.1         224.0.0.18        VRRP          74 Announcement (v3)
> Frame 4: 74 bytes on wire (592 bits), 74 bytes captured (592 bits)
v Ethernet II, Src: CiscoMer_91:50:00 (cc:03:d9:91:50:00), Dst: IPv4mcast_12 (01:00:5e:00:00:12)
  > Destination: IPv4mcast_12 (01:00:5e:00:00:12)
  > Source: CiscoMer_91:50:00 (cc:03:d9:91:50:00)
    Type: 802.1Q Virtual LAN (0x8100)
v 802.1Q Virtual LAN, PRI: 0, DEI: 0, ID: 10
    000. .... .... .... = Priority: Best Effort (default) (0)
    ...0 .... .... .... = DEI: Ineligible
    .... 0000 0000 1010 = ID: 10
    Type: IPv4 (0x0800)
> Internet Protocol Version 4, Src: 10.0.11.1, Dst: 224.0.0.18
v Virtual Router Redundancy Protocol
  > Version 3, Packet type 1 (Advertisement)
    Virtual Rtr ID: 13
    Priority: 255 (This VRRP router owns the virtual router's IP address(es))
    Addr Count: 7
    0000 .... = Reserved: 0
    .... 0000 0110 0100 = Adver Int: 100
    Checksum: 0x9783 [correct]
    [Checksum Status: Good]
    IP Address: 10.0.11.1
    IP Address: 192.168.102.1
    IP Address: 10.0.31.1
    IP Address: 10.0.51.1
    IP Address: 10.0.2.1
```

Notice how the VRRP messages list ALL the SVI IP addresses that it is willing to serve.

Now we'll disable the link switch → active MX from the MX.

Once the SPARE has had time to converge, its status is as follows:

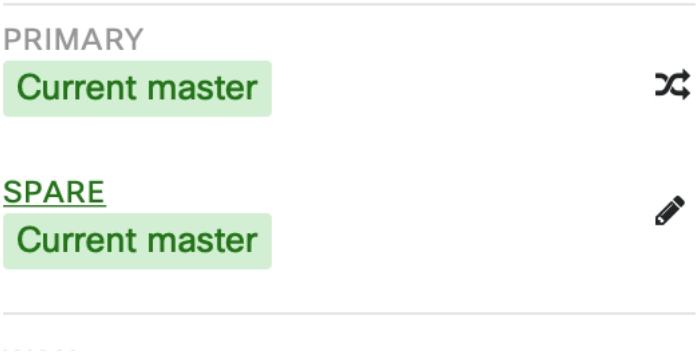

So this results in a Split brain, which makes sense as there is no dedicated heartbeat interface.

But… does the failover work?....

→ a test shows that the failover worked and the clients have regained Internet access.

After the failover, a Wireshark capture of the VRRP traffic from the SPARE MX looks as follows:

Meraki MX Firewall High Availability study

```
    4 0.542288    10.0.11.1         224.0.0.18        VRRP          74 Announcement (v3)

> Frame 4: 74 bytes on wire (592 bits), 74 bytes captured (592 bits)
> Ethernet II, Src: CiscoMer_91:50:00 (cc:03:d9:91:50:00), Dst: IPv4mcast_12 (01:00:5e:00:00:12)
> 802.1Q Virtual LAN, PRI: 0, DEI: 0, ID: 10
v Internet Protocol Version 4, Src: 10.0.11.1, Dst: 224.0.0.18
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0xc0 (DSCP: CS6, ECN: Not-ECT)
    Total Length: 56
    Identification: 0x23dd (9181)
  > Flags: 0x00
    Fragment Offset: 0
    Time to Live: 255
    Protocol: VRRP (112)
    Header Checksum: 0xa1a5 [validation disabled]
    [Header checksum status: Unverified]
    Source Address: 10.0.11.1
    Destination Address: 224.0.0.18
v Virtual Router Redundancy Protocol
  > Version 3, Packet type 1 (Advertisement)
    Virtual Rtr ID: 13
    Priority: 235 (Non-default backup priority)
    Addr Count: 7
    0000 .... = Reserved: 0
    .... 0000 0110 0100 = Adver Int: 100
    Checksum: 0xab83 [correct]
    [Checksum Status: Good]
    IP Address: 10.0.11.1
    IP Address: 192.168.102.1
```

→ even the MAC addresses look the same, as is to be expected. Only the Priority number betrays the switchover.

3. Observations

-ALL your IP addresses failover and there is *no separate VIP address*.
-**ALL** your VLAN interfaces / SVI will Multicast to 224.0.0.19 with VRRP v3 / protocol 112.
-Can you load balance amongst the 2 MXs? No way; the whole thing fails over, that's it.

Once the link to the PRIMARY MX is enabled again, will it Fallback?!
→YES

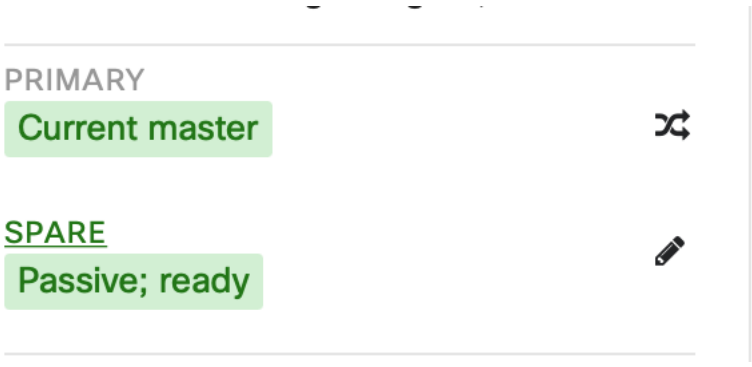

Can you Influence this pre-emption? NO
Can you set the Priority? NO, the Master is leading and uses Priority 255, whereas the SPARE uses Priority 235, which seems hardcoded.

But.. how DO you configure eg a Trunk interface on the Warm spare?
YOU DO NOT. YOU ONLY configure the MASTER and ALL your settings will be replicated to the SPARE MX. So make sure you use the same ports on both MXs for the same functionality.

SO..
If you connect the downstream switch to eg Port 1 on BOTH MXs, and you disable this port on the Primary switch, (there is no menu for controlling ports on the SPARE) do you loose ALL connections to the downstream switch as this is replicated to the SPARE MX?
YES!

How long does It take in case of eg a power failure on the PRIMARY MX, before the MX will failover and Internet connectivity is regained?
I tested this: about 25 seconds for the MX 64 Firewalls.

The failover time during a link failure to the primary MX was about 9 seconds.

Can you tune this? Nope.

4. MX Failover and Spanning Tree

In order for Spanning Tree NOT to block the port to the PRIMARY MX, I had to switch the cables on the downstream switch so it would favor the PRIMARY MX port. (the lowest port-ID in spanning tree)

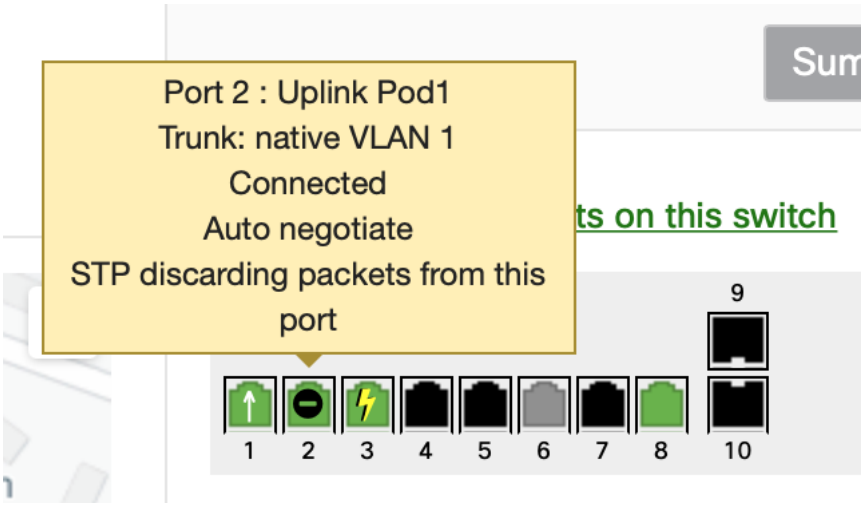

The **MXs do NOT run spanning tree** so do not produce PBDU's. Since the MXs used in the Lab are connected to a Cisco switch, they probably forward the same BPDU from the Cisco switch here. Hence the choice of the Port will now determine which port is Root – Port and which port will be Discarding.

How do you remove the 2nd MX from the HA deployment?

Meraki MX Firewall High Availability study

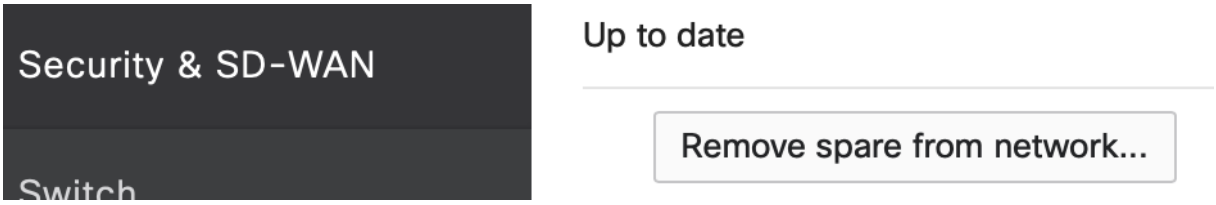

Well.. now that was not so hard.. was it?!

I would prefer to have more control over the HA behavior (priorities, pre-emption, convergence delays etc.) but that's not in the Meraki cards.
Pity the Meraki documentation doesn't do a better job at how the mechanism actually works and I had to find this all out in a Lab.