meraki

White Paper

VPN Concentrator Redundancy

JANUARY 2013

This solution guide describes how to deploy a redundant pair of VPN concentrators in “one-armed” concentrator mode.

Table of Contents

Meraki, Inc. | 660 Alabama St, San Francisco, CA 94110 | (415) 432-1000 | sales@meraki.com

1. Introduction

Meraki Auto VPN provides significant operational savings for distributed networks. This paper outlines how to implement high availability (HA) using a primary / warm spare MX pair based on the VRRP protocol, to minimize the downtime in case of a hardware failure. Initially, the HA pairing will be limited to “one-armed” VPN concentrator mode. In future releases, this limitation will be removed.

2. High Availability Overview

2.1 Normal Operation

During normal operation, there are two MX devices both deployed as “one-armed” VPN concentrators in the datacenter. The active VPN concentrator is called the “primary” concentrator and the backup concentrator is called the “warm spare” concentrator. Each concentrator has its own IP address to exchange management traffic with Meraki cloud-based centralized management. However, the concentrators also share a “virtual IP address.”

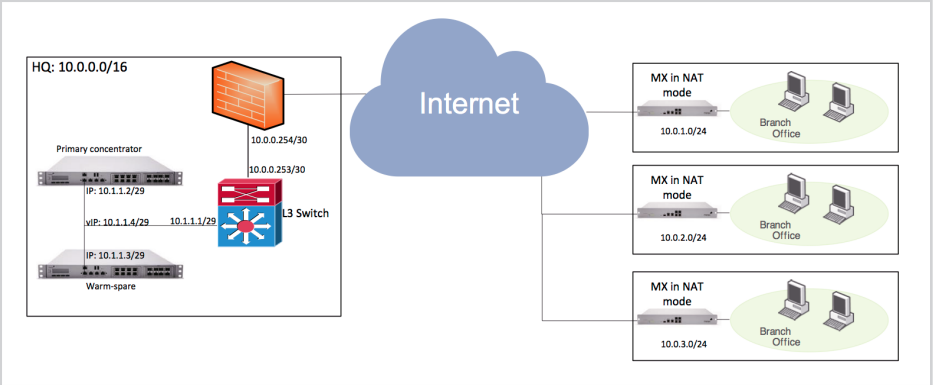

Figure 1: Redundant VPN concentrator pairs in the datacenter.

2.1.1 VIRTUAL IP

The virtual IP address (vIP) is an IP address shared by both the primary and warm spare VPN concentrators. VPN traffic is sent to the vIP rather than the physical IP address of the individual concentrators. The primary and warm spare concentrators use the VRRP protocol to synchronize and select the active concentrator for VPN traffic.

2.2 Failure Detection

The primary / warm spare concentrators share health status information via the LAN they are connected to using the VRRP protocol. In other words, failure detection does not depend on connectivity to the Internet / Meraki dashboard. Upon failure detection, the warm spare concentrator will assume the primary role until the original primary is back online.

2.3 Failback to the Original Primary Concentrator

Once the original primary VPN concentrator is back online and starts advertising its health status via the VRRP protocol, the warm spare concentrator will relinquish VPN concentrator function back to the original primary concentrator.

2.4 Failback Latency

The total time for failure detection, failover to the warm spare concentrator, and ability to start processing VPN packets is typically less than 30 seconds.

Meraki, Inc. | 660 Alabama St, San Francisco, CA 94110 | (415) 432-1000 | sales@meraki.com

3. Supported Configuration

The redundant VPN concentrator feature requires configuring MX security appliances in “one-armed” VPN concentrator mode at headquarters or datacenter.

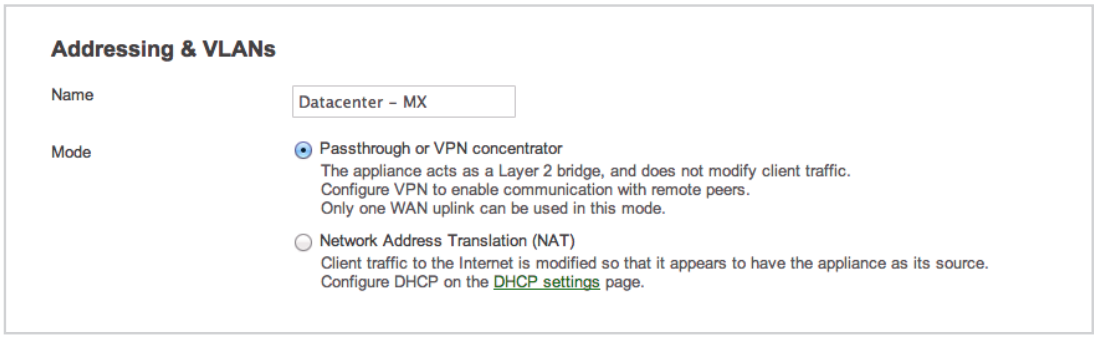

Figure 2: Passthrough or one-armed VPN concentrator mode selector in the Meraki dashboard.

3.1 Prerequisites

- Each MX device requires a license (see FAQ section for details).
- Both MX devices must be in the same Layer 2 broadcast domain.
- Both MX devices must be able to communicate with the Meraki Cloud Management service (i.e., have access to the Internet).
- Both MX devices must be connected as “one-armed” VPN concentrators.

3.2 Connecting the MX in “One-Armed” VPN Concentrator Mode

In one-armed VPN concentrator mode, the MX pair is connected only via their respective Internet ports. Only VPN traffic is routed to the MX, and both ingress and egress packets are sent through the same interface.

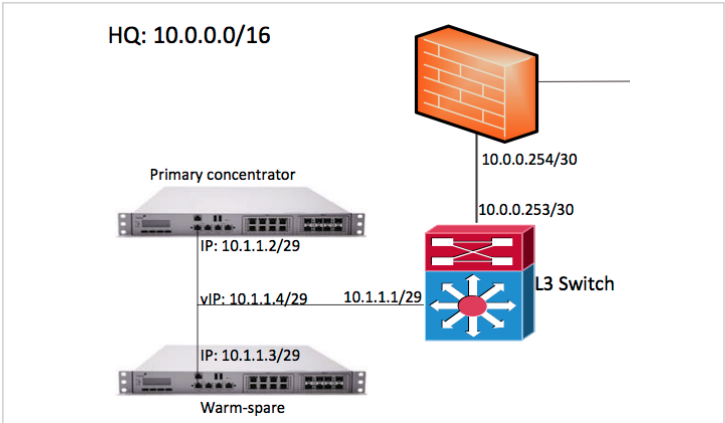

Figure 3: “One-armed” VPN concentrator pairs

3.3 Route Configuration

To send traffic over the VPN tunnel, a new route must be added on the L3 switch. Here is a sample of the Cisco IOS commands for the network diagram above, assuming `10.0.1.0/24`, `10.0.2.0/24`, and `10.0.3.0/24` are branch network routes:

```
ip route 10.0.1.0 255.255.255.0 10.1.1.4
ip route 10.0.2.0 255.255.255.0 10.1.1.4
ip route 10.0.3.0 255.255.255.0 10.1.1.4
```

Note that `10.1.1.4` is the virtual IP (vIP) for the primary / warm spare MX pair. Refer to section 4.2 for setting the virtual IP for the primary / warm spare pair.

Meraki, Inc. | 660 Alabama St, San Francisco, CA 94110 | (415) 432-1000 | sales@meraki.com

4. Configuring the Concentrator Pair

4.1 Setting Up the Primary (active) Concentrator

Create a new network and enter the serial number of the MX. Make sure the appliance is set to `Passthrough or VPN concentrator` mode.

4.2 Setting Up the Warm Spare (passive) Concentrator

1. Click `Configure > Warm spare` in the Dashboard.
2. Enter the serial number or the order number for the warm spare MX appliance
3. Assign a virtual IP for the primary / warm spare MX pair.

4.2.1 VIRTUAL IP REQUIREMENTS:

- The virtual IP must be in the same subnet / VLAN scope.
- The virtual IP must be unique. In particular, it cannot be the same as either the primary or warm spare's physical IP address.

EXAMPLE

In the example above (see Figure 1), if the primary concentrator IP address is `10.1.1.2/29`, and the warm spare concentrator IP address is `10.1.1.3/29`, `10.1.1.4/29` is a valid virtual IP address since it is not used by any other device.

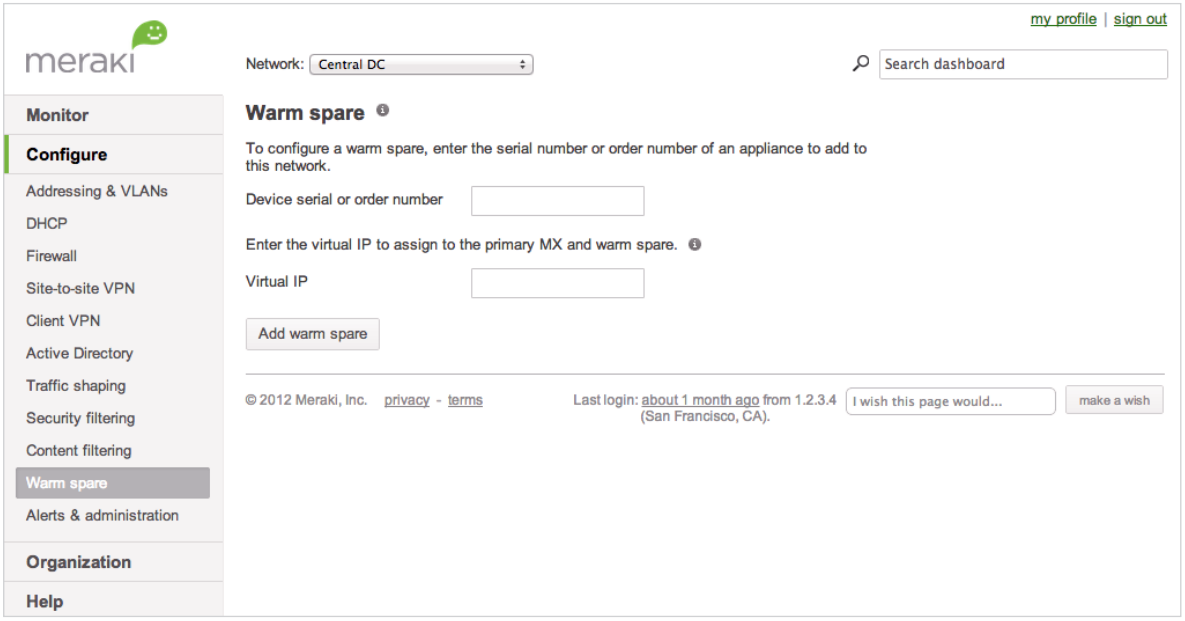

Meraki, Inc. | 660 Alabama St, San Francisco, CA 94110 | (415) 432-1000 | sales@meraki.com

5. FAQ

Q How do I set up HA if the MX is deployed in NAT mode?

A See the NAT Warm Spare section of the Warm Spare documentation for more about configuring HA in NAT mode.

Q Do both MXs have to be the same model?

A While this is not enforced, it is encouraged. For budgetary reasons, some customers may opt to use a lower performance/cost MX as a secondary concentrator. However, it is important to ensure that the secondary MX has sufficient networking power to handle VPN / WAN optimization traffic during failover.

Q Do I need a license for the warm spare unit?

A No, only one license is required to operate a warm spare pair.

Q How long does it take for the system to detect failure of the primary unit and failover to the warm spare unit?

A Failover is typically less than 30 seconds.

Q What happens if the primary unit comes back online?

A It immediately assumes the master VPN concentrator functionality.

Q Will the existing connections (e.g., VoIP calls) get disrupted?

A Yes, there will be a brief disruption, typically less than 30 seconds, during failover.

Q Can we put the warm spare unit in our secondary datacenter for disaster recovery (DR)?

A As long as both datacenters are connected via a Layer 2 connection that allows spanning a single subnet/broadcast domain across both datacenters, primary and warm spare units can be placed in geographically separate datacenters.

6. Conclusion

The MX product line now offers an easy-to-configure and fully automated redundancy for large-scale VPN / WAN optimization deployments using the industry-standard VRRP protocol. Additional documentation and installation instructions can be found at http://docs.meraki.com/mx

Meraki, Inc. | 660 Alabama St, San Francisco, CA 94110 | (415) 432-1000 | sales@meraki.com