



# Designlet: VMware HCX Mobility Optimized Networking for Azure VMware Solution

VMware Cloud Migration

## Table of contents

Designlet: VMware HCX Mobility Optimized Networking for Azure VMware Solution .....	3
Introduction .....	3
Scope of the Document .....	3
Summary and Considerations .....	4
Use Cases and Scaling .....	5
Uses Cases for Mobility Optimized Networking .....	5
Scaling limits .....	6
Planning and Implementation .....	8
Planning .....	8
Usage Guidance .....	8
Implementation .....	8
Authors and Contributors .....	10

## Designlet: VMware HCX Mobility Optimized Networking for Azure VMware Solution

### Introduction

Mobility Optimized Networking (MON) is an Enterprise-tier capability of the HCX Network Extension feature. MON enabled network extensions improve traffic flows for migrated virtual machines by enabling selective cloud routing (within the destination environment), avoiding a long round trip network path via the source gateway.

The behavior of extended networks is such that all routed traffic for migrated workloads is directed back to the source-site gateway. MON allows you to configure the optimal path for migrated workload traffic to other extended network segments, cloud-native network segments, and Internet egress.

### Scope of the Document

VMware HCX's Mobility Optimized Networking (MON) is used to provide optimized traffic patterns when workloads are spread across on-premises and VMware Cloud SDDC on a L2 extended segment. This document will cover use cases, consideration, and limitations of the feature.

## Summary and Considerations

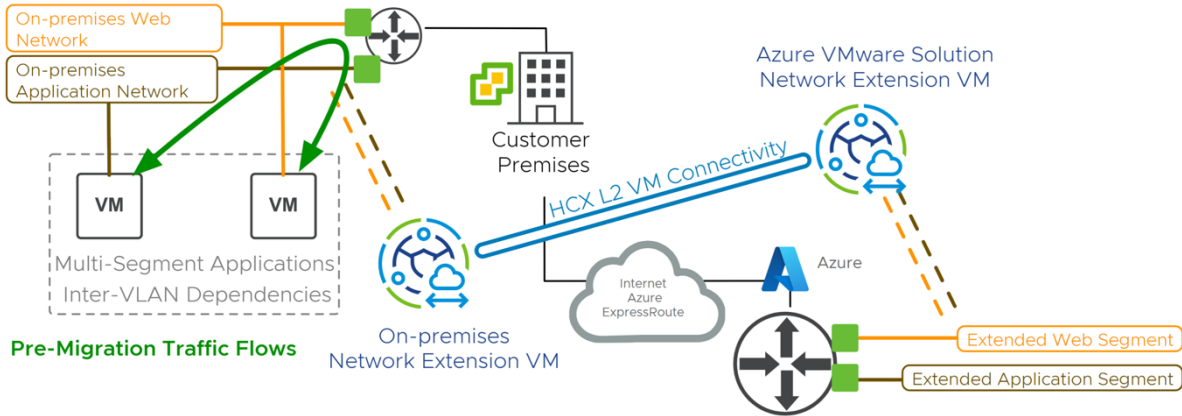
<b>Use Case</b>	
<b>Pre-requisites</b>	
<b>General Considerations/Recommendations</b>	
<b>Performance Considerations</b>	Enabling MON will provide lower latency and higher bandwidth between VMs on extended networks in the same SDDC, and lower latency to the internet due to local egress. There is no other performance impact from using MON beyond the considerations when using Network Extension.
<b>Cost Implications</b>	Traffic egressing the Cloud SDDC to Azure Services or internet may incur network charges based on Azure pricing.
<b>Document Reference</b>	<a href="#">HCX User Guide</a> <a href="#">VMware HCX Configuration Maximums</a> <a href="#">Azure network pricing</a>
<b>Last Updated</b>	August 2022

## Use Cases and Scaling

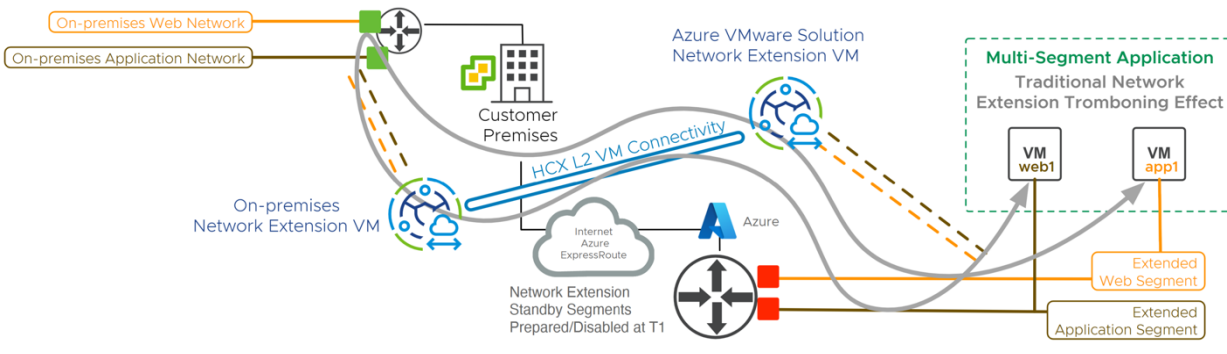
### Uses Cases for Mobility Optimized Networking

This section provides an overview of workload traffic flows using HCX Network Extension with and without Mobility Optimized Networking. This section covers supported use cases. Other use cases may be possible but are not specifically supported.

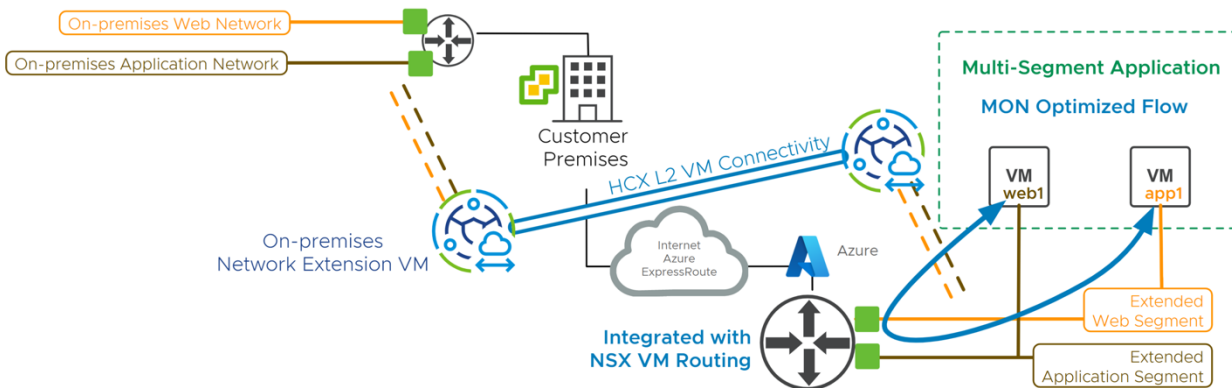
MON improves network performance and reduces latency, in certain traffic patterns, for VMs migrated to the cloud on an extended L2 segment. These traffic patterns are outlined below.



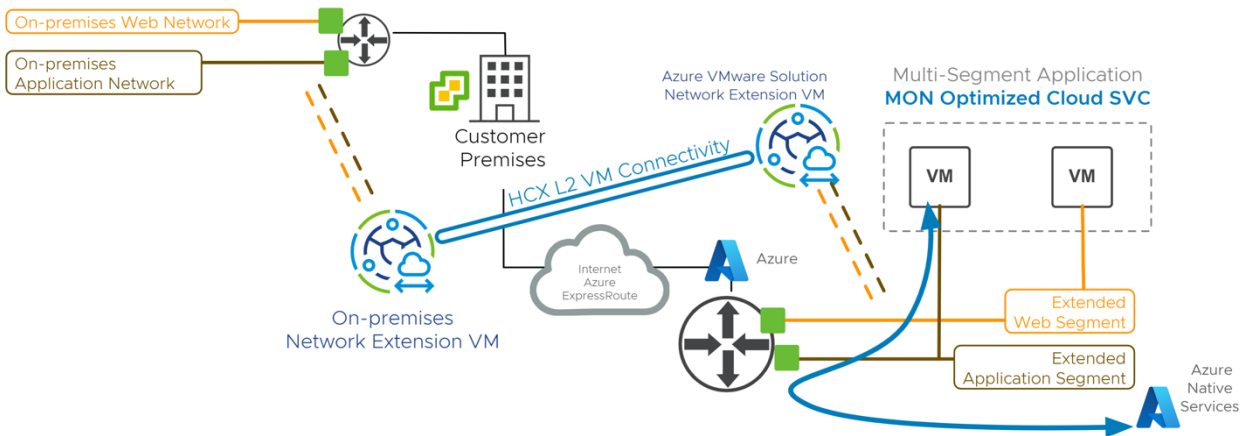
Without MON, HCX Network Extension expands the on-premises layer-2 network to the cloud SDDC while the default gateway remains at the source. The network tromboning effect is observed when virtual machines at the destination, on different extended segments communicate.



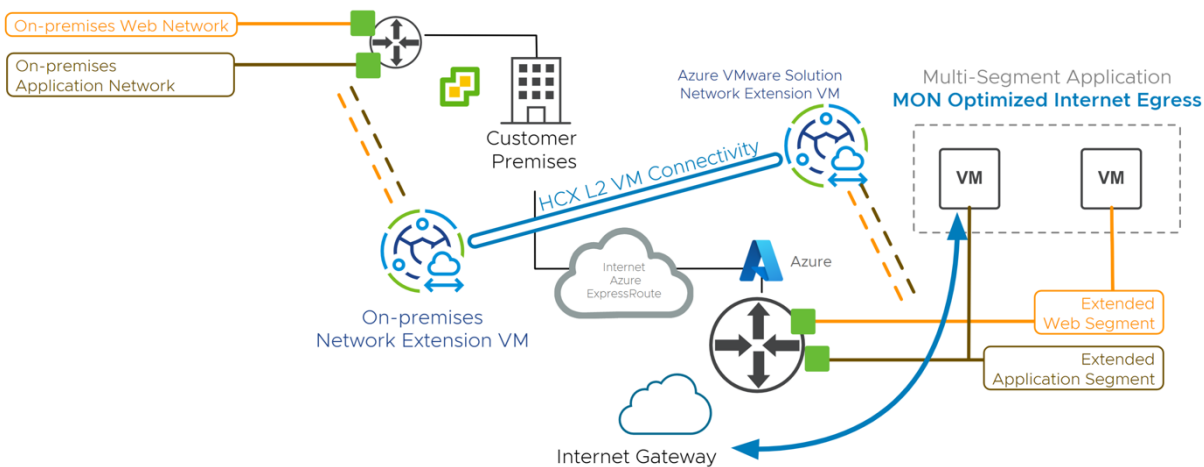
MON enables migrated virtual machines reach segments within the SDDC without sending packets back to the source environment router.



MON can be configured to allow migrated virtual machines to reach services hosted within a public cloud.



MON enables migrated virtual machines to use the SDDC Internet interface (with SNAT).



## Scaling limits

### Introduction

Now that we have reviewed the use cases, we will discuss the MON scale limits, and provide recommendations so you can get the most from the feature while still ensuring HCX performance.

VMware publishes MON scale limits in the configuration maximum's page. The discussion below is current at time of publishing, but please always consult the config max page for up-to-date info. - [VMware HCX Configuration Maximums](#)

Calculating the scale limits for MON required considerations of outside events in an HCX environments. These events consume resources on the HCX manager and were factored into the published limits.

- VM resources on the HCX Cloud Manger
- Number of extended networks with MON enabled
- Number of VMs with MON enabled
- Number of VMs for which MON is being enabled/disabled at the same time
- Number of VMs with network events that will require MON configuration changes at the same time (i.e. during a host failure event, vSphere Distributed Resource Scheduler will relocate a certain number of VMs)

### HCX manager Sizing

VMware supports 2 sizes of the HCX Manager appliance which determine the support maximums for MON. In this section we will look at the 2 sizes and how the scale numbers are impacted.

#### HCX manager - Default Size

HCX manager is always deployed with the "default" size but can be scaled up as a day 2 operation.

**The HCX Manager default resource allocation is: 4 vCPU and 12 GB of Memory.**

Based on previous factors and a default appliance size, MON limits per HCX Cloud Manager

- **250 VMs with MON enabled** - (total number of VMs at the destination environment with MON enabled. New migration waves should not exceed this maximum)
- **100 Network Extensions with MON enabled** - (distributed across all Network Extension appliances)
- **100 concurrent Migrations to MON enabled networks** - (This limit is supported while not exceeding the 250 total VMs with MON enabled)

### HCX Manager - Custom Size

If your scale requirements exceed the default size, we can scale the HCX Cloud Manager to reach higher limits and performance. This is achieved by increasing the vCPU and Memory of the HCX Cloud Manager. Please work with your Cloud Provider to initiate the scale up process.

### The up-scaled HCX Manager resource allocation: 8 vCPU and 24 GB of Memory

With an up-scaled HCX Manager, the scale limits are increased

- **900 VMs with MON enabled** - (total number of VMs at the destination environment with MON enabled. New migration waves should not exceed this maximum)
- **100 Network Extensions with MON enabled** - (distributed across all Network Extension appliances)
- **100 concurrent Migrations to MON enabled networks** - (This max is supported while not exceeding the 250 total VMs with MON enabled)

### Scale examples

To help further explain the scale values, we will look at 2 examples and how they impact your environment. In both examples we will consider an "HCX Manager - Default Size" deployment

**Example 1** - In your destination environment you have 200 MON enabled VMs. You plan to migrate additional VMs to the destination with MON enabled. In this example you can migrate up to 50 VMs with MON enabled through RAV or Bulk Migration. This will bring you to a total of 250 MON enabled VMs at the destination

**Example 2** - In your destination environment you have 100 MON enabled VMs. You plan to migrate additional VMs to the destination with MON enabled. In this example you can migrate up to 100 VMs at one time with MON enabled through RAV or Bulk Migration. This will bring you to 200 MON enabled VMs but stay within the support limits of 100 concurrent migrations.

## Planning and Implementation

### Planning

VMware HCX MON gives a user the ability to optimize traffic flows and reduce latency on extended L2 segments. When planning for your HCX Network extension design, you should consider the impacts of enabling this feature. Changing a traffic pattern can have serious and sometime unintended consequences on applications in your environment.

Before discussing the MON planning questions, let's review the HCX network extension and the default behavior. When a network segment has been extended with the HCX (NE) service, you are expanding the boundary of an L2 domain, while keeping the default gateway for that L2 network at your on prem datacenter. The default behavior has all traffic leaving that L2 segment returning on prem. This is where enabling the MON features on the extended network could improve your overall design. Once enabled, the user has the choice to continuing egressing on-premises or changing this behavior to egress from the Cloud site.

- HCX MON is controlled through the HCX Connector UI (On-premises site)

### Usage Guidance

Now that we have reviewed the MON feature, let's look at optimizing your deployment. We find many customers take a "MON-enabled by default" approach and begin pushing the limits. In our experience, taking a closer look at your VM and traffic patterns will allow you to enable MON where necessary and stay well below the supported maximums.

Considerations to optimize MON usage

- MON optimizes traffic flows and lowers latency but typically isn't required for communication. In other words, you may have test/dev applications that require network communication but can withstand additional latency. In this case, you can consider turning MON off and allowing traffic to trombone.
- Examine traffic flows - We see many customers that make assumptions on VM to VM communication and therefore enable MON. Upon closer inspection, these VMs either don't talk or communicate so infrequently, that MON is not justified.
- Cutover networks - as your migration progress, ensure you are cutting over extended segments as soon as possible. Once a network is cutover, MON is no longer enabled for the VMs in that segment. Cutting over a single segment can drop your MON consumption significantly.

In addition to these considerations, we want to re-enforce the true purpose of MON and clarify misconception.

1. MON **is** a mechanism to provide temporary relief for applications during a migration. These applications are sensitive to latency or have exceptionally large east/west throughput. These applications require the traffic to stay local and not trombone back on-premises.
2. MON **is not** a permanent routing solution. It is intended to be used during a migration exercise.
3. MON **should not** be the central focus of your migration design. It should be used to augment your broader migration strategy for select applications.
4. The focus of your migration should always be to cutover the network as soon as possible, thus staying well within the MON scale limits.

### Implementation

HCX MON is enabled per VM, or for an entire extended network. Optimized routing will apply to migrated VMs after following the steps below to enable MON. VMs that have not been migrated will continue to communicate with their local gateway.

Enable MON

1. In the HCX Manager UI, navigate to Services > Network Extension.
2. In the Network Extension screen, expand a site pair to see the extended networks. Network Extensions enabled for MON are highlighted with an icon.
3. Expand each extension to display network details.
4. Select a Network Extension and enable the slider for Mobility Optimized Networking. Enabling MON applies to all subsequent events, such as VM migrations and new VMs connected to the network. VMs in the source environment and VMs not having VM Tools are ineligible for MON.
5. For any existing migrated VMs requiring MON, follow the steps below:
6. Select a VM and expand the row. You can select multiple VMs using the check box next to each workload.
7. Select Target Router Location and choose the cloud option from the drop-down menu.
8. Select Proximity Conversion Type: "Immediate Switchover" or "Switchover on VM Event". "Immediate Switchover" transfers



the router location immediately. If a workload VM has ongoing flows to the source router, they will be impacted. These impacts can include a few second traffic outage. "Switchover on VM Event" transfers the router location upon VM events like NIC disconnect and connect operations, and VM power cycle operations.

9. Click Submit. All selected VM workloads are configured for MON, which is indicated by a MON icon being displayed.
10. You can verify MON is enabled by measuring the latency to the gateway from a VM with a ping before and after MON is enabled. Viewing the NSX-T routing table in HCX will also show that host routes have been installed in the routing table for MON-enabled VMs.

## Authors and Contributors

A list of authors and contributors should be listed here in the following format:

- Ian Allie, Product Solutions Architect, VMware
- Emad Younis, Director, VMware Cloud Solutions, VMware
- Michael Kolos, Technical Product Manager, VMware
- Gabe Rosas, Staff Technical Product Manager, VMware
- Siddharth Ekbote, Senior Staff Engineer, VMware
- Snehansu Chatterjee, Senior Product Manager, VMware

