

HCX Proximity Routing for Layer 3 Aware VM Mobility

HCX Proximity Routing – v.2

Table of Contents

1. Introduction	3
About the HCX Service	3
HCX Network Extension (without Proximity Routing)	3
HCX Network Extension with Proximity Routing	4
2. Use Cases for HCX Proximity Routing	4
2.1 Local Egress for Internet Access	4
2.2 Optimal Routing to Target environment's Native Networks	4
2.3 Gradual Network Swing-Over	4
2.4 Application Level Failover with HCX Disaster Recovery	4
3. Understanding Virtual Machine Traffic Flows in Proximity Routed Networks	5
3.1 Virtual Machine Traffic within the Proximity Routed Extended Network	5
3.2 Virtual Machine Routed Traffic (External to the Proximity Routed Network)	5
4. How Does HCX Proximity Routing Work?	6
4.1 Proximity Routing and Bulk Migration	6
4.2 Proximity Routing and HCX vMotion	6
4.3 Proximity Routing and HCX Disaster Recovery	Error! Bookmark not defined.
5. Supported NSX Architectures for HCX Proximity Routing	7
5.1 Proximity Routing to NSX Edge Services Gateway (ESG)	7
5.2 Proximity Routing to NSX DLR (with No NSX ESG, or 3 rd party Virtual Router)	Error! Bookmark not defined.
5.3 Proximity Routing to NSX DLR connected to NSX ESG	8
5.4 Proximity Routing to NSX DLR with Multiple NSX ESG with ECMP Enabled	8
6. Proximity Requirements and Restrictions	9
7. Enabling Proximity Routing in the HCX Interface	10
About the Authors	11
Gabe Rosas	11
Jesse Schachter	11

1. Introduction

WHAT IS HCX ?

The HCX service provides application mobility, infrastructure hybridity and compatibility without retrofit.

NETWORK EXTENSION

HCX Network Extension provides secure L2 adjacency at the target environment, for various types of virtual machine networks at the source.

About the HCX Service

VMware HCX delivers secure and seamless application mobility and infrastructure hybridity across environments running vSphere 5.0 and above, both on-premises and in the cloud. HCX abstracts on-premises and cloud resources and presents them as one continuous hybrid environment. VMware HCX mobility services run over high-performance, secure and optimized multi-site interconnects.

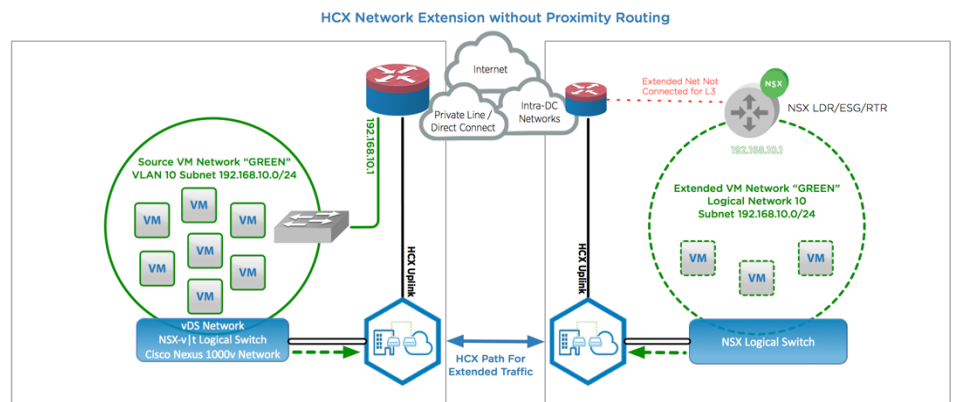
HCX Network Extension (without Proximity Routing)

The HCX Network Extension service securely stretches VLAN and VXLAN backed Virtual Machine networks from a vSphere or 3rd party Distributed Switch to a remote environment as a seamless operation within the vSphere Client.

Virtual Machines on the extended network can be moved with zero-downtime and retain their MAC and IP address identity. Network extension is useful in the following ways:

- Migrated virtual machines on the extended network utilize existing routing and security policies on-premises, expediting consumption of the target site's resources.
- Easy cloud bursting without the complexities of architecting the remote site networks.
- Accelerates migration timelines by removing the need to refactor the workloads, which tends to expose unaccounted for application dependencies.

The Network Extension creates a Logical Network at the target environment and connects it with the source network without needing to make configuration changes to the physical network layer.

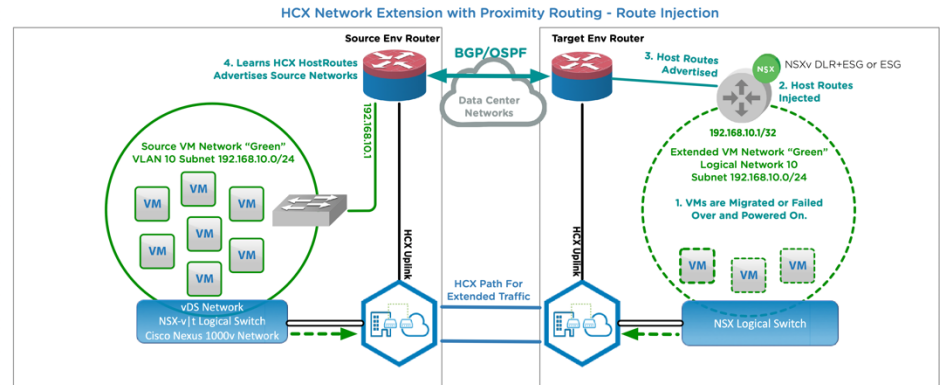


PROXIMITY ROUTING

Proximity Routing integrates HCX Network Extension with NSX Dynamic Routing, enabling virtual machines to route with a local software gateway.

HCX Network Extension with Proximity Routing

Proximity Routing is an advanced setting for the Network Extension that can be enabled (per network) when networks are extended with HCX. Proximity Routing integrates HCX with the NSX dynamic routing configuration. Virtual machines that have been migrated to the target environment, or failed over with disaster recovery, will use proximity routing to exit the network without tromboning via the target environment's router.



PROXIMITY ROUTING USE CASES

- Local egress for Internet Access.
- Local routing to reach the target environment's networks and virtual machines.
- Gradual network swing-over.
- Application-level failover with HCX Disaster Recovery.

2. Use Cases for HCX Proximity Routing

VMware HCX Proximity Routing is useful in the following scenarios:

2.1 Local Egress for Internet Access

Virtual machines migrated to the proximity routed stretched network will access the Internet through the local NSX gateway at the target environment rather than using a trombone path via the original default gateway at the source site.

2.2 Optimal Routing to Target Environment's Native Networks

Virtual machines migrated to the proximity routed stretched network will communicate with other target networks through the NSX gateway at the target environment rather than using a trombone path via the original default gateway at the source site.

2.3 Gradual Network Swing-Over

Proximity Routing allows the source site network to be "moved" to the target environment gradually. Once all virtual machines on the extended network are migrated to the target environment, the unextend operation can be executed to enable the NSX gateway for the migrated network, and the source gateway can be disabled.

2.4 Application Level Failover with HCX Disaster Recovery

With Proximity Routing, recovered virtual machines can continue to operate using the target environment's NSX gateway in the event of a network outage in the source environment.

PROXIMITY ROUTING TRAFFIC FLOWS

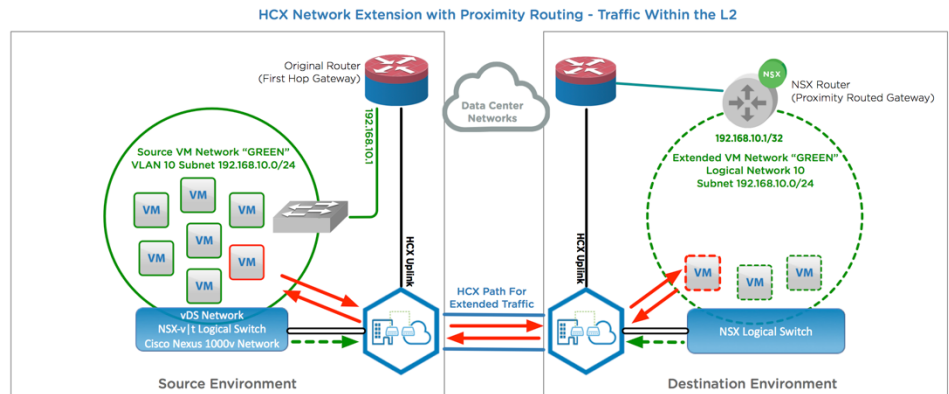
Proximity Routed Extended Network Behavior:

- Remote/migrated virtual machines on the same extended network will be reached through the HCX Network Extension path.
- Source site virtual machines on the proximity routed network will use the original source site's gateway to reach other networks.
- Target environment virtual machines on the proximity routed network will use the target site NSX gateway to reach other target site networks, including external networks or the Internet.

3. Understanding Virtual Machine Traffic Flows in Proximity Routed Networks

3.1 Virtual Machine Traffic within the Proximity Routed Extended Network

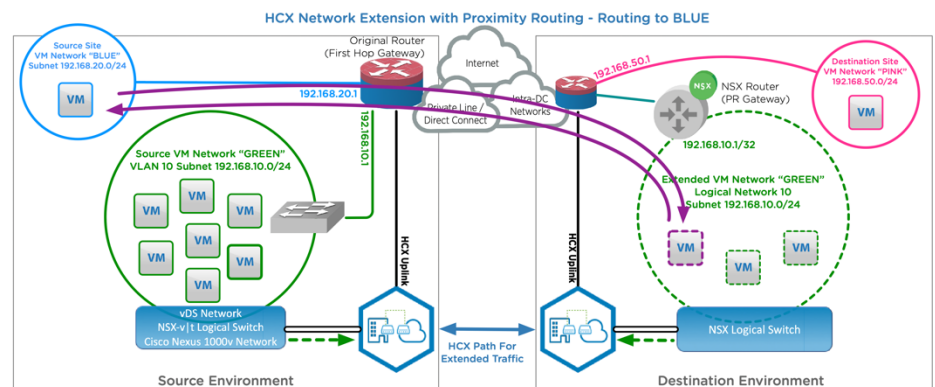
Virtual machine traffic within the same extended network uses the Network Extension path. This is represented with red arrows in the illustration below.



3.2 Virtual Machine Routed Traffic (External to the Proximity Routed Network)

Traffic from virtual machines on the proximity routed network to destinations outside of the extended network will use the gateway in the datacenter where the virtual machine resides.

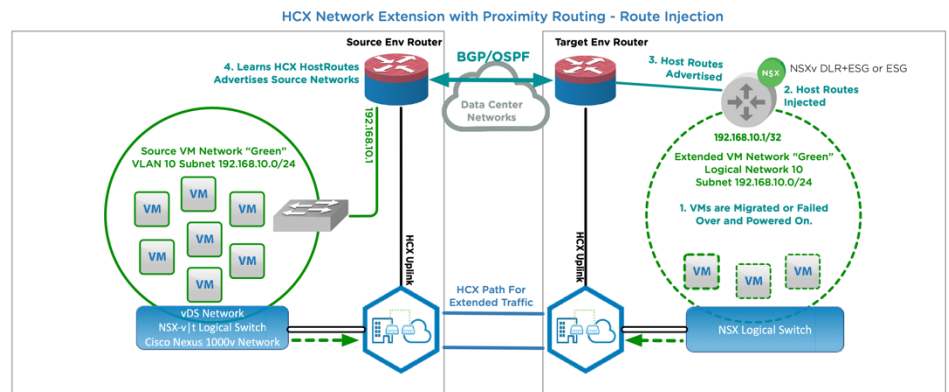
Inbound traffic to virtual machines on proximity routed networks will use the same path as the egress traffic. Proximity Routing prevents asymmetrical traffic flows. This is represented with red arrows in the illustration below.



4. How Does HCX Proximity Routing Work?

Proximity Routing allows an extended network to have an active gateway on both the source and target at the same time by using ARP filtering and route injection on the NSX router at the target environment.

- The injected routes are present so the target environment routers (via dynamic routing) have awareness of the virtual machines accessible locally to them.
- For proximity routing to work as intended, the injected routes must be learned in the source environment's routers. In the illustration below, the source router (on the left side) should learn the /32 host routes for the three depicted virtual machines - directing traffic over the data center networks.



4.1 Proximity Routing with Bulk Migration and Disaster Recovery

Virtual machines moved to the target environment with HCX Bulk Migration and Disaster Recovery will use the target side NSX gateway for communicating with other networks.

4.2 Proximity Routing with HCX vMotion

Virtual machines moved to the target environment with HCX vMotion will temporarily use the source site router until the virtual machine is power cycled. This is done with vMotion migrations to prevent disruption of live traffic flows. Once the migrated virtual machine is power cycled, it will switch to using the target side NSX gateway for communicating with other networks.

“Proximity Routing is like being able to take your cellphone number to any provider and knowing calls will work.”

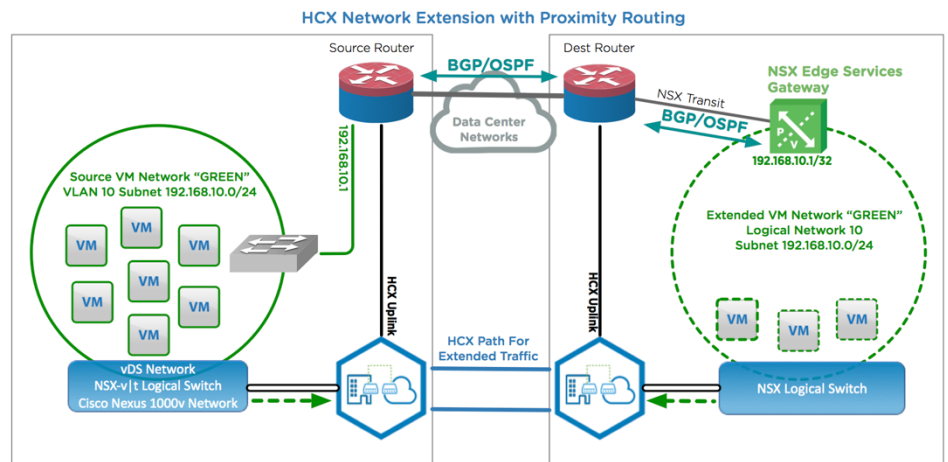
SERGE MASKALIK
SR DIRECTOR OF ENGINEERING, HCX
VMWARE

5. Supported NSX Architectures for HCX Proximity Routing

HCX Proximity Routing works with the following NSX Topologies in the target environment:

5.1 Proximity Routing to NSX Edge Services Gateway (ESG)

- In this topology, the target environment NSX ESG requires dynamic routing to be enabled.
 - The NSX ESG redistributes HCX-injected /32 static routes into the routing protocol.
 - The source router learns the injected /32 routes.

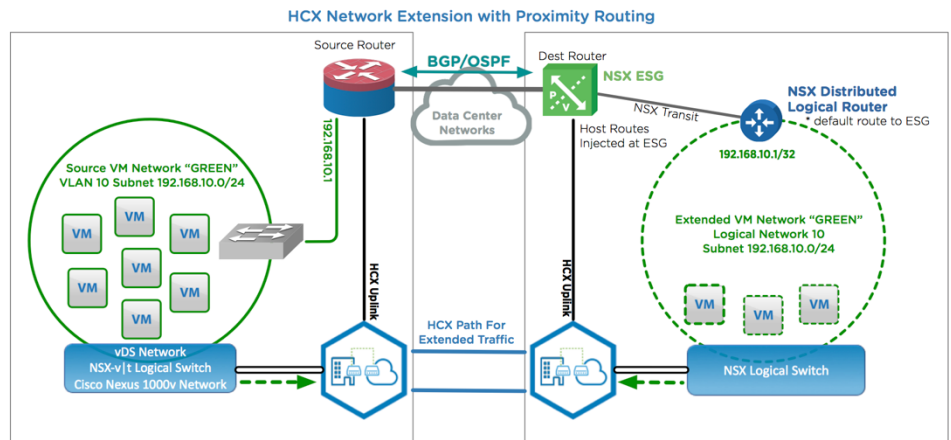


PROXIMITY ROUTING SUPPORTED NSX TOPOLOGIES

- PR-Extension to an NSX ESG.
- PR-Extension to an NSX DLR+ESG.
- PR-Extension to an NSX DLR + ESG with ECMP Scale OUT.

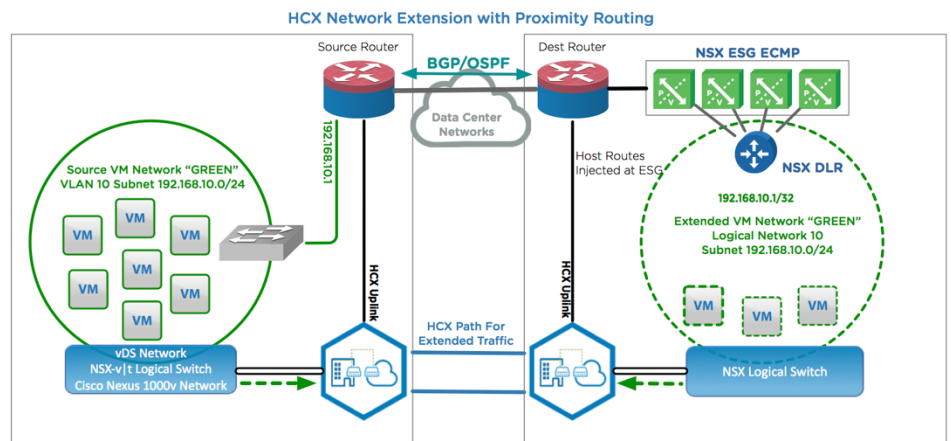
5.2 Proximity Routing to NSX DLR connected to NSX ESG

- HCX supports topologies using NSX DLR at the target connected to an NSX ESG. During the proximity routed network extension operation, the NSX DLR functions as the first hop router for migrated virtual machines.
- The target environment requires dynamic routing to be enabled.
 - The NSX DLR has a default route with the NSX ESG as the next hop.
 - The /32 host routes will be injected in the NSX ESG and will point to the NSX DLR as the next hop.



5.3 Proximity Routing to NSX DLR with Multiple NSX ESG with ECMP Enabled

- HCX supports topologies using an NSX DLR at the target connected to NSX ESGs in ECMP scale-out.



6. Proximity Requirements and Restrictions

The following requirements should be met for a successful implementation of HCX Network Extension with Proximity Routing.

PROXIMITY ROUTING REQUIREMENTS	
Virtual Machine Requirements	
Virtual Machine VMware Tools	Migrated virtual machines must have VMware Tools installed.
Detectable IP Address	The virtual machine IP address must be detectable by VMware Tools.
Environmental Requirements & Restrictions	
Dynamic Routing	<p>The proximity routing feature integrates HCX Network Extension and Migrations with NSX Dynamic Routing at the target environment. For routing to function correctly, the following conditions must be met:</p> <ul style="list-style-type: none"> • The first-hop router at the source site must be participating in dynamic routing. • The first-hop router at the source must learn HCX-injected /32 host routes from the target environment. • All participating routers at the source and target environments must be capable of learning the /32 host routes that HCX injects into NSX. <p>Important: Proximity Routing can always be enabled, but it will not work and is not supported in environments without Dynamic Routing.</p>
Firewall Policies at the Target Environment	Because the first hop gateway function “follows” the virtual machine with proximity routing, it is necessary to pre-configure security policies to ensure the virtual machine's routed traffic does not become blocked after it is migrated to the target.
NSX at the Target	NSX for vSphere 6.4 or above is required.
NSX at the Source	NSX is not required at the source environment for Proximity Routing.

7. Enabling Proximity Routing

Open the Network Extension interface in the HCX Plugin.

Proximity Routing is enabled by toggling the PR switch during the Network Extension operation.

Figure 7-1 Extension Operation Without Proximity Routing





Figure 7-2 Extension Operation With Proximity Routing Enabled



About the Authors

Gabe Rosas

Gabe Rosas is a Technical Product Manager with the VMware HCX team. He is a rounded technologist with comfortable designing and operating traditional and software defined data center solutions, and HCX powered multi-site multi-cloud architectures.

Jesse Schachter 🙋

Jesse Schachter is a Staff Engineer in the HCX team. He has worked on products such as vCloud Director and NSX Edge. He has expertise in Internet network security, virtualized routing and switching.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.
Copyright © 2019 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: vmw-wp-temp-word 2/19