

VMware HCX Availability Guide

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2022 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About This Document	4
1 About VMware HCX	5
Updated Information	6
2 HCX Availability Configurations and Best Practices	7
vSphere Best Practices for HCX Availability	8
NSX Best Practices for HCX Availability	15
HCX Best Practices for Availability	17
Application Path Resiliency	18
Multi-Uplink Service Mesh Resiliency	20
Network Extension High Availability	22
Service Mesh Availability During Upgrades	24
3 About the Author	27

About This Document

The *VMware HCX Availability Guide* provides information to help users understand known configurations that affect the availability of migrated virtual machines, extended networks and VMware® HCX systems. This document provides best practices for improved business continuity outcomes while using HCX.

Intended Audience

This information is for migration and cloud architects, systems administrators and any reader with interest in the implementation of highly available HCX deployments. It is assumed that readers have familiarity with VMware HCX, vSphere and NSX, and have basic knowledge of the systems underpinning HCX services.

VMware Technical Publications Glossary


VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <https://vmware.com/support/pubs>.

About VMware HCX

1

VMware HCX streamlines application migration, workload rebalancing and business continuity across data centers and clouds.

VMware HCX delivers secure and seamless application mobility and infrastructure hybridity across vSphere environments both on-premises and in the cloud. HCX abstracts on-premises and cloud resources and presents them as one continuous hybrid environment, enabling users to connect infrastructure and adopt a hybrid cloud vision, or a full migration to cloud as a consistent experience.



Updated Information

This *VMware HCX Availability Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *VMware HCX Availability Guide*.

Revision	Description
28 Jan 2021	HCX Availability Guide initial release. HCX Service Availability & Resiliency is now deprecated.


HCX Availability Configurations and Best Practices

2

This section includes known best practices for vSphere, NSX and HCX to improve the availability of HCX services and migrated workloads.

VMware HCX provides workload connectivity and mobility functionality in a proxy model, on behalf of distinct environments tethered by HCX.

This document provides HCX related availability guidance following the presupposition that the underlying network, server and storage infrastructure that follows business continuity best practices, without exploring those topics directly. Throughout this document we discuss configurations in vSphere, NSX and HCX that directly contribute to the availability condition of HCX services (workload migrations in-flight, and any virtual machines relying on HCX Network Extension based connectivity



This chapter includes the following topics:

- vSphere Best Practices for HCX Availability
- NSX Best Practices for HCX Availability
- HCX Best Practices for Availability

vSphere Best Practices for HCX Availability

vSphere related considerations and best practices when designing vSphere environments to be used with highly available HCX deployments.

Note The recommendations provided in the vSphere Best Practices sections of the HCX Availability Guide are considered in-scope in on-premises deployments (where the customer manages the source and destination vSphere infrastructure configurations). In public cloud or hybrid deployments, the cloud side infrastructure configurations (DRS settings, HA, vmkernel interfaces, reservations) may be the responsibility of the cloud provider.

In public cloud or hybrid deployments, it may not be possible to implement the recommendations presented in this document.

vSphere Distributed Resource Scheduler

DRS spreads the virtual machine workloads across vSphere hosts inside a cluster and monitors available resources. Fully automated DRS is a powerful vSphere feature that allows a cluster to automatically re-balance virtual machines across the cluster hosts.

Best Practices:

- Use **Partially Automated DRS** mode in the workload clusters being migrated with HCX. In this mode, DRS automatically places powered up virtual machines on a host that is guaranteed to provide the required resources, and make recommendations for rebalancing the cluster resources.
- Avoid the use of **Fully Automated DRS** mode. In this mode, DRS automatically applies rebalancing recommendations on powered on VMs. This mode can cause the following effects on HCX services:
 - Aggressive rebalancing "thrashing" of HCX migration and network extension appliances due to their smaller disk footprint.
 - Contention (between DRS rebalancing and the HCX migration) over vMotion resources.
 - Replication reconfigurations halt migration progress, disrupt checksum operations. A full synchronization may be triggered, delaying migration progress.
 - Excessive network path changes result in the Network Extension flooding RARP for the VM traffic path adjustments.

Note If the Fully Automated DRS configuration must be used, the Service Mesh appliances should be excluded (to maximize service stability). This exclusion is not persistent to upgrade and redeploy operations.

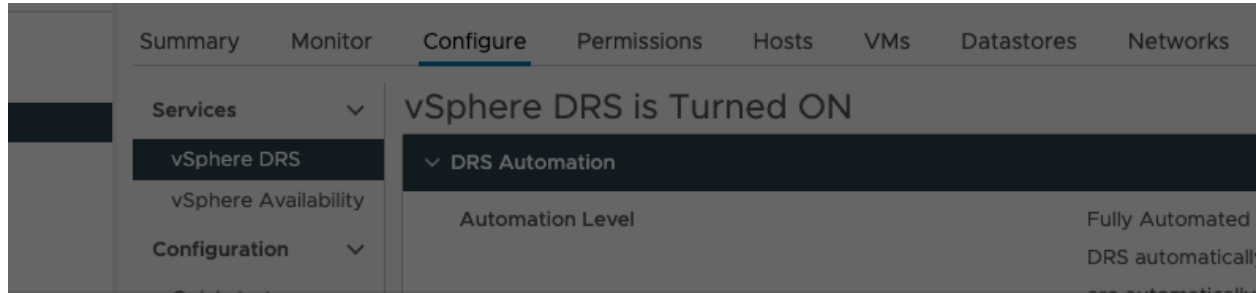
■ **DRS Virtual Machine and Host Affinity Rules** for HCX:

- The HCX Network Extension HA feature (which is distinct from vSphere High Availability and discussed in detail later in the document) configures anti-affinity rules automatically for NE appliances in an HA Group to keep the HA Active and HA Standby on different hosts.
 - HCX NE host anti-affinity uses "should" rules, if there is one host, both appliances in a group will be placed on the same host.
- The HCX migration appliances (IX) may benefit from anti-affinity rules to place IX appliances on different hosts. This allows vMotion/RAV operations to be executed in parallel instead of queuing for serial execution.
- The HCX WAN Optimization appliance may benefit from an affinity rule that places the HCX-WO and HCX-IX on the same host. This rule simplifies and optimizes the service chained traffic data path. When WAN Optimization is enabled, IX sends migration traffic to the WO appliance for the deduplication and compression operations.

Note DRS Affinity Rules for the Network Extension HA service are added automatically by HCX and will persist upgrade and redeploy operations.

DRS Affinity Rules manually added for the IX and WO appliances are not persistent to upgrade and redeploy operations.

Figure 2-1. Configuring DRS Automation in the vCenter Server



Edit Cluster Settings | OnpremCluster

vSphere DRS

Automation | Additional Options | Power Management | Advanced Options

Automation Level

Partially Automated ▼

DRS automatically places virtual machines onto hosts at VM power-on.
Migration recommendations need to be manually applied or ignored.

Migration Threshold ⓘ

Conservative (Less Frequent vMotions) Aggressive (More Frequent vMotions)

DRS provides recommendations when workloads are moderately imbalanced.
This threshold is suggested for environments with stable workloads. (Default)

Predictive DRS ⓘ

Enable Predictive DRS

Virtual Machine Automation ⓘ

Enable Virtual Machine Automation

ESXi Maintenance Mode

You place a host in maintenance mode when you need to service it, for example, to install more memory. A host enters or leaves maintenance mode only as the result of a user request.

Best Practices:

- Perform required ESXi host maintenance after current replication-based migrations have completed.
 - Replication reconfigurations halt migration progress, disrupting checksum operations. A full synchronization may be triggered, delaying the migration progress.
 - A host cannot relocate a virtual machine during HCX Cold Migration, vMotion or the scheduled RAV switchover window.
- While the HCX supports the use maintenance mode, it is a best practice to perform cluster relocation of the HCX Network Extension appliances during a business maintenance window.

vSphere Cluster High Availability

vSphere Cluster HA provides high availability for virtual machines by pooling the virtual machines and the hosts they reside on into a cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts.

Best practice:

- Use **vSphere HA** enabled clusters to deploy HCX components. See the [vSphere Availability](#) guide for more information.
- Configure HCX Network Extension appliances with **High HA Restart Priority**. This ensures the NE appliances are started before virtual machines that will rely on it for connectivity.

vSphere Cluster VM Monitoring

The vSphere VM monitoring feature resets individual virtual machines if their VMware tools heartbeats are not received within a set time.


When the VM Tools process in the HCX Service Mesh appliances is not able to send the heartbeats (and I/O activity is not detected), this is an indicator that the HCX operating system has failed. In this case VM Monitoring determines that the HCX appliance has failed and will reboot it in the same ESXi host.

Best Practice:

- Enable VM Monitoring on the Service Mesh deployment cluster when it is not possible to enable Network Extension High Availability.
- VM Monitoring is not required in deployments with Network Extension High Availability.

Note The VM heartbeat monitoring setting is applicable to all virtual machines in the cluster. The VM heartbeat monitoring setting is not persistent to upgrade and redeploy operations.

Figure 2-2. Enabling Cluster VM Heartbeat Monitoring



Fault Tolerance

FT provides continuous availability for a virtual machine by creating and maintaining another VM that is identical and continuously available to replace it in the event of a failover situation.

Best Practices:

- Do not use vSphere Fault Tolerance with HCX appliances. HCX cannot manage FT nodes, and it is not supported.

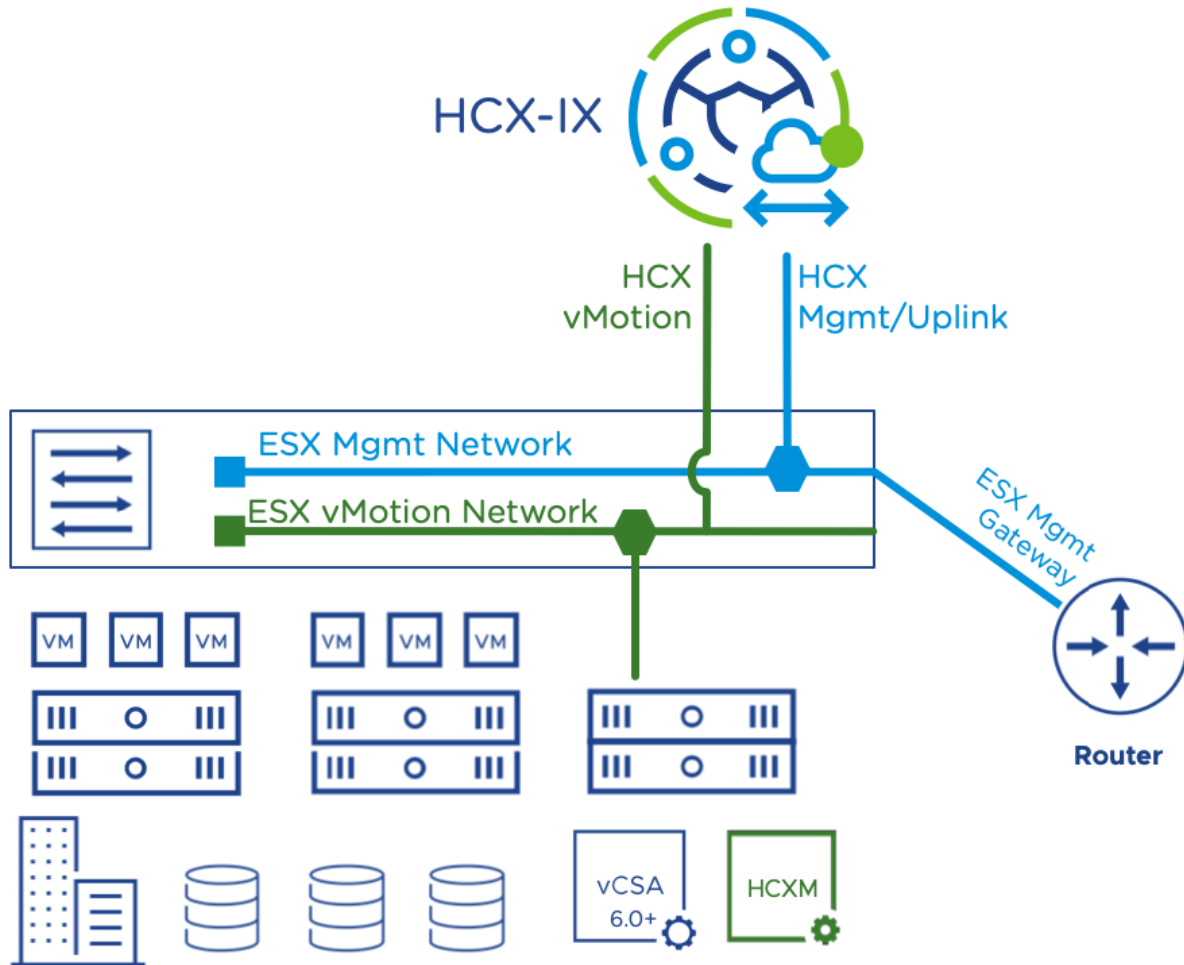
ESXi Host VMkernel vMotion Network

During HCX vMotion and Replication Assisted vMotion, there is vMotion protocol traffic transmitted between the HCX migration appliance (IX) and the ESxi host vMotion network.

Best Practices:

- Allocate IP addresses for the HCX migration appliance (HCX-IX) from existing vMotion networks to optimize the data path and simplify troubleshooting. Configure the HCX Network Profile to use the same portgroup as the vMotion VMkernel interfaces.
- When working with virtual machines connected to a VMware Standard Switch, verify that a vMotion Standard Port Group is created consistently on all hosts in the cluster. It is commonly misunderstood that a Standard Port Group for vMotion is created by default.
- HCX operates at maximum availability when the underlying cluster vMotion configuration is implemented according to best practices. Please see [networking best practices for vMotion](#) for more information.

Figure 2-3. HCX vMotion vNIC connected to the ESXi vMotion Network



ESXi Host VMkernel Replication Network

During Bulk Migration and Replication Assisted vMotion, there is Host-Based Replication protocol traffic transmitted between the HCX migration appliance (IX) and the ESXi host management or replication vmkernel network. It is common for the ESXi management vmkernel interface to be used for replication traffic, but a dedicated VMkernel interface can be configured.

Best Practices:

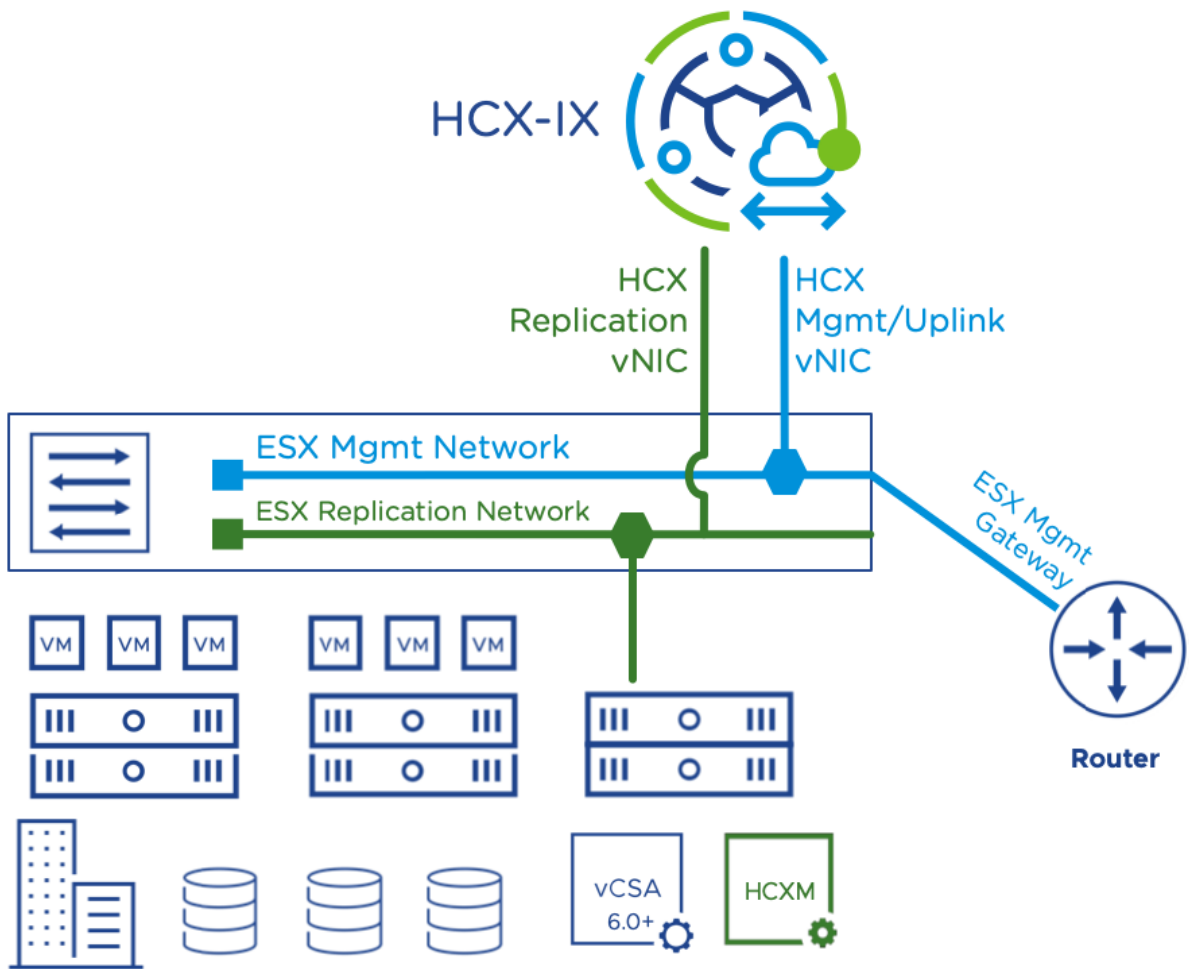
- Allocate IP addresses for the HCX migration appliance (HCX-IX) from existing Management or Replication networks to optimize the data path and simplify troubleshooting.

- When designing cluster networking, use a dedicated cluster replication network . A workload replication vmkernel network can be designed using the same principles as the [networking best practices for vMotion](#). This allows data transfer traffic to be separated from management workflows.

Note When a dedicated Replication VMkernel network is not present, the cluster uses the management network for replication traffic.

In the ESXi vmkernel configurations, vSphere Replication NFC traffic must be configured on the Management VMkernel interface. HCX IX does not support configurations vSphere Replication NFC connections to the ESXi Replication VMkernel interface.

Figure 2-4. HCX Replication vNIC connected to the ESXi Replication Network



Dedicated HCX Mobility Clusters

This refers to the practice of implementing HCX with dedicated cluster hosts for the migration and network extension appliances. Mobility clusters do not contend with the workload virtual machines that will be migrated. This configuration maximizes overall resiliency and recovery by eliminating contention with workloads and optimizes traffic distribution by separating migration and network extension traffic flows from non-migrated workload traffic.

Best Practices:

- Clusters designed specifically for HCX mobility should meet requirements and best practices for vSphere High Availability. For more information See the [vSphere Availability Guide](#).
- Deployment requirements for mobility clusters:
 - A mobility cluster and the workload clusters must be managed by the same vCenter Server.
 - The mobility cluster is selected as the deployment cluster the Compute Profile configuration.
 - The workload clusters are selected as the service clusters in the Compute Profile configuration.
- Requirements for HCX Network Extension with mobility clusters:
 - The mobility cluster hosts must have access to the workload VLANs by association to the existing workload clusters either joined to existing workload Distributed Switches, or by ensuring the new Distributed Switch and Distributed Portgroups can connect to the same underlying VLANs.
- Requirements for HCX Bulk migration:
 - The migration appliance on the mobility cluster hosts must have IP reachability to the workload cluster vmkernel interfaces for Replication/NFC.
- For Cold Migration, HCX vMotion and Replication-Assisted vMotion:
 - The migration appliance on the mobility cluster hosts must have routed reachability to the workload cluster management and vMotion vmkernel interfaces.
 - The mobility cluster host must belong to the workload Distributed Switch.
- In scaleout deployments where multiple service meshes share one mobility cluster:
 - Each workload cluster can only have one migration appliance, multiple workload clusters can share a mobility cluster, but the migration appliances must run on different hosts to provide increased relocation (HCX vMotion/RAV) concurrency.

NSX Best Practices for HCX Availability

VMware NSX-T related considerations and best practices when designing for highly available HCX deployments.

Working with NSX Distributed Firewall

Distributed firewall (DFW) monitors all the East-West traffic on your virtual machines.

Best Practices:

- Use permissive firewall policies to decouple the workload migration from the firewall policy migration.
- In migrations with strict zero-trust microsegmentation, or other restrictive firewall objectives, apply the security policies to the DFW prior to migrating the workloads.
- Use Security Tags for dynamic Security Group membership. In NSX to NSX migrations, HCX migrates the NSX Security Tag.
- Ensure the DFW does not block optimized migration traffic (IX <> WO) packets on the hcx-fleet segment.

NSX Edge Transport Nodes

The NSX Edge provides routing services and connectivity to network NSX Edges that are external to the NSX-T Data Center deployment.

Best Practices:

- HCX migrations and network extensions can add multi-Gbps traffic over a WAN. Separate HCX traffic from application traffic whenever possible to reduce network contention scenarios.

Segment Policies

When migrating virtual machines that rely on DHCP, ensure that NSX Segment Policies are not blocking the virtual machine DHCP requests.

Best Practices:

- When migrating virtual machines that depend on DHCP (without Mobility Optimized Networking), use NSX segment configurations that allow DHCP requests to be forwarded to the source gateway.
- When migrating virtual machines that depend on DHCP (with Mobility Optimized Networking), use NSX segment configurations that provide DHCP services. See [Configure DHCP](#) on a segment.

Route Advertisements

A tier-1 gateway is typically connected to a tier-0 gateway in the northbound direction and to segments in the southbound direction. Route advertisements can optionally be configured on the tier-1 gateway.

Best Practices:

- When using Mobility Optimized Networking, verify that the **All Static Routes** in the Tier-1's **Route Advertisement** is configured to advertise or not advertise according to the design (MON migrations will always add virtual machine static routes). For more information, see [HCX Network Extension with Mobility Optimized Networking for NSX-T](#) in the HCX User Guide.

HCX Best Practices for Availability

VMware HCX configurations and best practices for highly available HCX deployments.


Compute Profile Resource Reservations

A **Compute Profile** contains the compute, storage, and network settings that HCX uses on each site to deploy the Interconnect-dedicated virtual appliances when a Service Mesh is added. A Compute Profile can be used to define CPU and Memory.

Best Practices:

- Apply 100% CPU and Memory resource reservations when Network Extension appliances are sharing resources with workload virtual machines.
- Use the HCX compute profile to configure CPU and Memory reservations. Resource reservations configured directly in vCenter Server are not persistent to HCX lifecycle operations.

Figure 2-5. Configuring CPU and Memory Reservations in the HCX Compute Profile



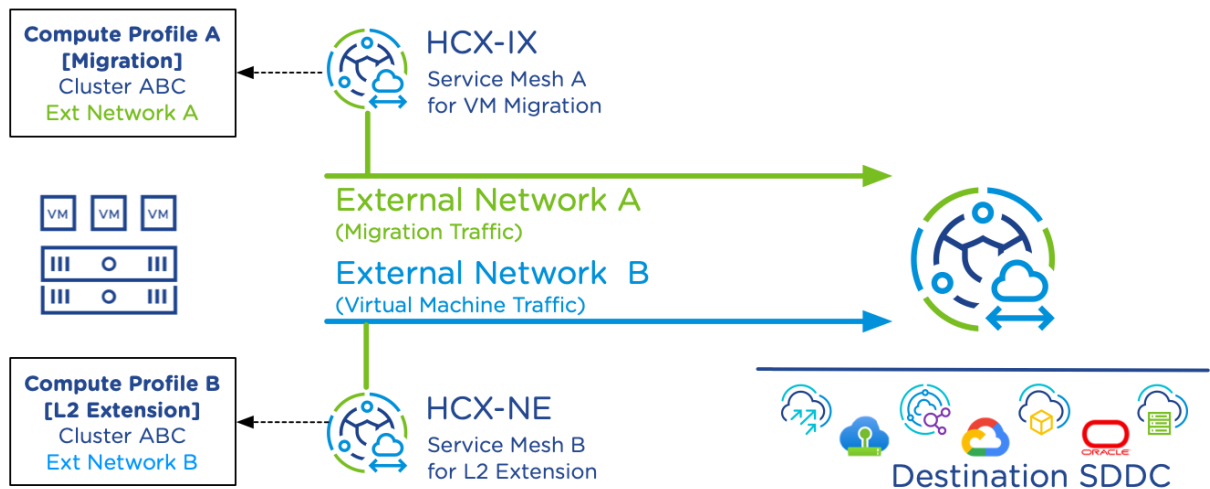
Configuring Dedicated Service Mesh

Typically a single **Service Mesh** is configured for both migration and network extension using a single Compute Profile configuration for both services. HCX allows Network Extension services to be separated into a second Service Mesh, with different uplinks and resource assignments.

Best Practices:

- Design HCX Uplinks where Migration Data traffic can be separated from Virtual Machine traffic (Network Extension traffic):
 - Configuring specific vSphere Distributed Port Groups for HCX Migration Uplink and HCX Extension Uplink enables more flexible policies configurations where network extension traffic can be prioritized, and where migration traffic and resource is isolated and appropriate network policies can be implemented.
 - Create the first compute profile and service mesh to enable migration services only.
 - Create a second compute profile and service mesh to enable Network Extension services.
 - This configuration allows the deployed appliances to use network and compute resources and network policies to be tailored to the service mesh function.

Figure 2-6. Dedicated Service Mesh for HCX Extension and HCX Migration



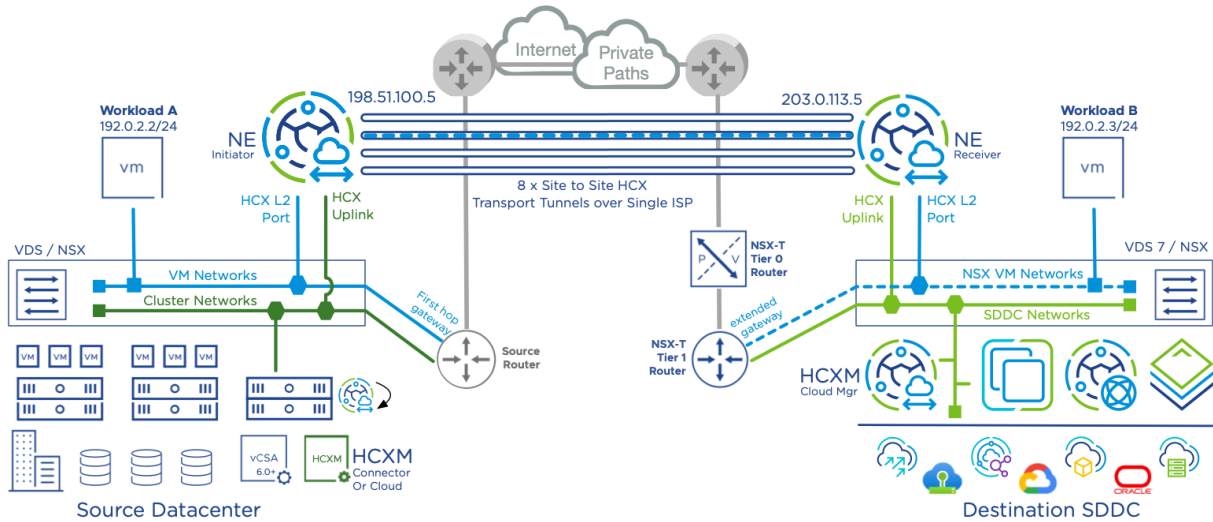
Application Path Resiliency

You can create network path resiliency for single vNIC HCX Service Mesh deployments.

Application Path Resiliency (APR) is an HCX Enterprise capability that creates additional site-to-site tunnels per uplink. HCX tracks the condition of each tunnel and avoids sending traffic over degraded or inoperational paths. When APR is enabled, eight tunnels are used in every migration appliance (HCX-IX) and network extension (HCX-NE) appliance in the service mesh.

APR enables environments that do not have access to multiple network providers or paths an option to achieve additional uplink path resiliency. It accomplishes this by using network flows that have unique 5-tuples, the most common characteristic used by network devices when determining how to distribute flows in Equal Cost Multi-Path (ECMP) routed paths or in switched paths that rely on load balancing hash algorithms .

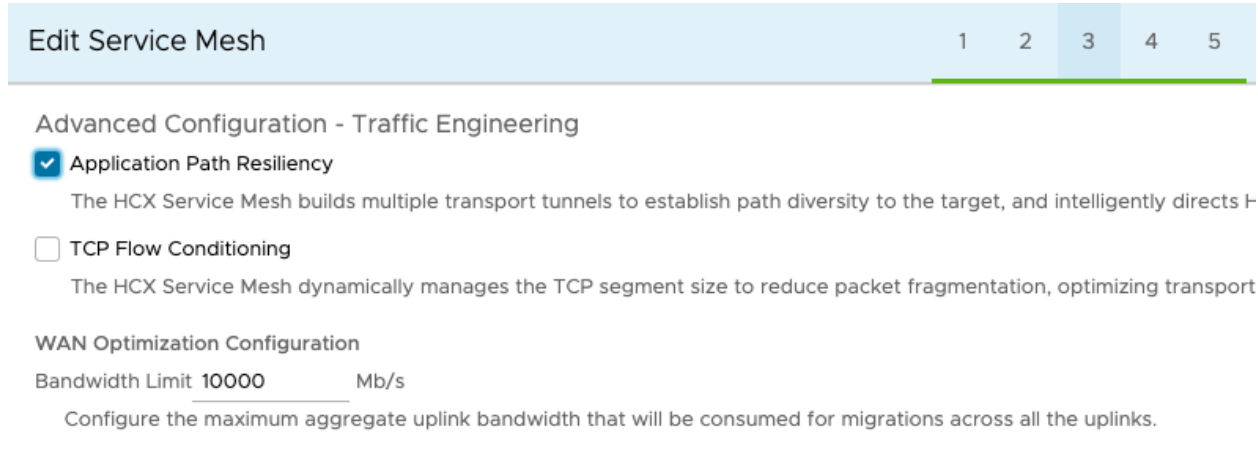
Figure 2-7. Single Uplink Multi-Tunneling for Network Path Resiliency



Best Practices:

- Enable APR in single uplink HCX Service Mesh configurations.
- Enable APR in single ISP configurations.
- Enable APR in single Direct Connect/Private Path configurations.
- Enable APR for network path resiliency. APR can detect and avoid underlying lossy links in a link aggregation bundle. APR does not provide link aggregation or load balancing functionality.
- Using APR may not be optimal, or may be excessive in deployments where HCX is connected to diverse network paths for resiliency using multiple uplinks.
 - As an example, in deployments with 3 uplinks (e.g. Direct Connect 1, Direct Connect 2 & Internet), using APR will result in maintaining 24 service mesh tunnels.
- Firewalls typically do not need to be adjusted for APR. The following characteristics apply:
 - Eight tunnels created per HCX Uplink.
 - Same destination IP address and destination port for the eight APR tunnels.
 - Same source IP address and variable port range (UDP 4500-4600) for the eight APR tunnels.

Figure 2-8. Configuring Application Path Resiliency in the Service Mesh

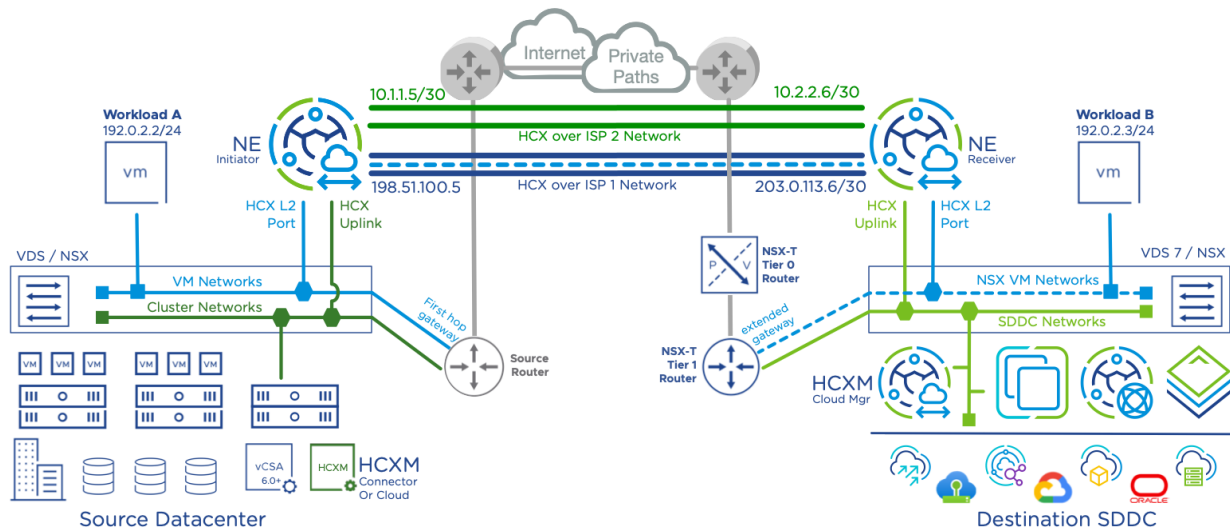


Multi-Uplink Service Mesh Resiliency

You can create network path resiliency with multi vNIC HCX Service Mesh deployments.

A multi-uplink service mesh deployment refers to a configuration that uses up to three uplink vNICs for site-to-site traffic for network path resiliency. Additional uplinks are automatically added to every migration appliance (HCX-IX) and network extension appliance (HCX-NE) in the service mesh.

Figure 2-9. Multi-Uplink Service Mesh for Network Path Resiliency



Best Practices:

- Use a multi-uplink configuration for environments with access to different network underlays between the environments (up to 3).

- Use a multi-uplink configuration to configuring connections on private network underlays with fallback to the public internet.
- In multi-uplink configurations, every HCX Uplink network must connect to a unique VLAN/VNI and a IP subnet.
- Multiple uplinks can be assigned in either the source or destination environment with the same resiliency outcome. For example:
 - Assigning two uplinks at the source and one at the destination will yield two service mesh tunnels.
 - Assigning one uplink at the source and two at the destination will also yield two service mesh tunnels.
- A tunnel will be created for every unique combination of uplinks. As an example, assigning two uplinks at the source and two uplinks at the destination will yield four tunnels:
 - Source Uplink 1 <--> Destination Uplink 1
 - Source Uplink 1 <--> Destination Uplink 2
 - Source Uplink 2 <--> Destination Uplink 1
 - Source Uplink 2 <--> Destination Uplink 2
- Using Application Path Resiliency with multi-uplink configurations means 8 tunnels will be created per unique uplink. In the previous example with 4 uplink combinations, enabling APR will result in 32 independent transport tunnels which is not efficient when resilience is already achieved with the distinct uplinks.

Figure 2-10. Configuring a Service Mesh with Multiple Uplinks

Edit Service Mesh
1 2 3 4 5
✕

Advanced Configuration - Override Uplink Network profiles (Optional)

Uplink network profiles are used to connect to the network via which the remote site's interconnect appliances can be reached. For destination HCX systems (IX/NE) accessed over the Internet, use the HCX Cloud's Uplink Network Profile to assign public IP addresses. By default the uplink network profiles for source and destination sites are used from the compute profiles selected at the source and the destination sites. Here you can optionally override the uplink network profiles. Some of that cases where this is required are point to point links like Direct Connect between two sites.

Please select all the uplink network profiles that you want to connect to or skip this step if you want to continue using the ones specified in the compute profiles.

Select Source Site Uplink Network Profile(s) ▼

Select Destination Site Uplink Network Profile(s) ▼

cloudDirect1 ✕


21 Free IPs

cloudDirect2 ✕

21 Free IPs

● ● ●

CONTINUE



Network Extension High Availability

Network Extension High Availability provides appliance failure tolerance to the HCX Network Extension service.

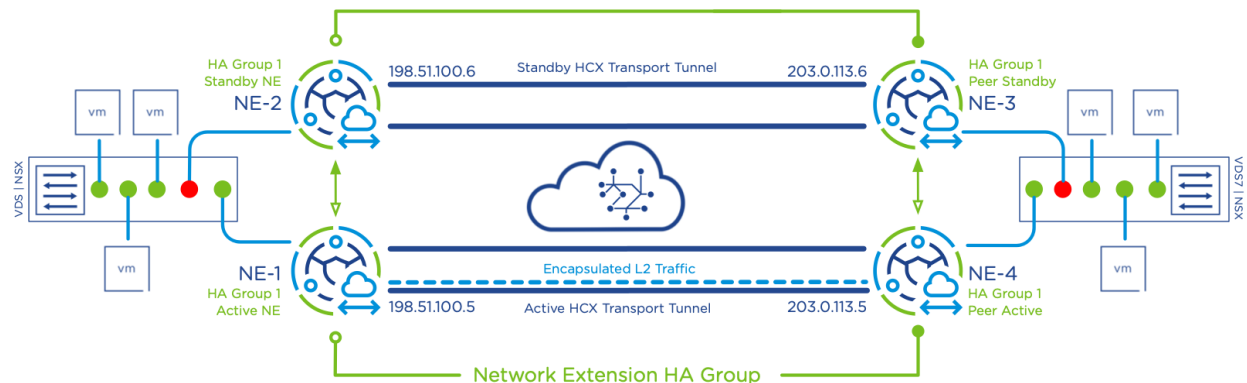
Network Extension High Availability is an HCX Enterprise capability that groups Network Extension appliances into resilient pairs (called an HA Group). When Network Extension High Availability is enabled for a selected appliance, HCX will pair it with an eligible appliance and enable an Active Standby resiliency configuration. This enables highly available configurations that can remain in-service in the event of an unplanned appliance level failure.

When either of the HA Actives fail, both standby appliances take over. The Network Extension High Availability is designed to recover within a few seconds after a single appliance has failed.

Note Network Extension High Availability feature is available as an **Early Adoption** feature with HCX 4.3.0.

Early Adoption (EA) denotes an initial release phase where a feature has limited field exposure and it has strong dependencies on the deployment environment for its functionality. While the feature has completed the entire development process and it is fully supported, it is expected to reach maximum stability and performance through subsequent maintenance releases.

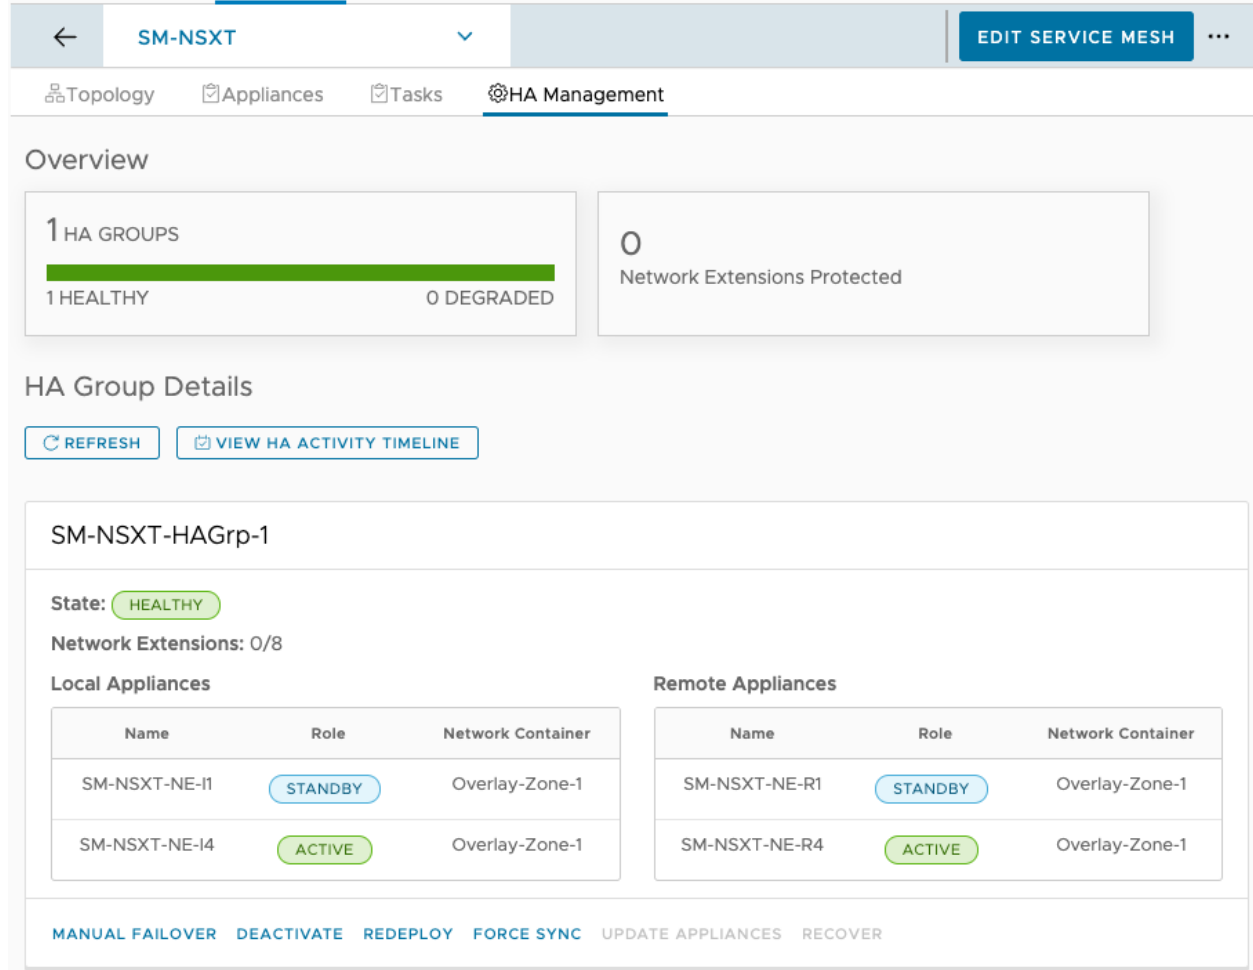
Figure 2-11. Network Extension in a High Available Active Standby Configuration



Best Practices:

- Plan for additional provisioning requirements:
 - Plan for the increased appliance count (additional CPU and Memory is needed to run HA Groups) . Each site will have a local HA Active and Standby. Four appliances per group.
 - Plan IP address consumption in the context of HA Groups. Both the HA Active and HA Standby appliances will establish service mesh transport tunnels with their remote peers.
- Plan for additional ESXi Host requirements:
 - Network Extension HA creates host anti-affinity rules.
 - The selected ESXi cluster should be licensed for DRS capabilities.
 - DRS must be enabled.
 - The ESXi host running the HA Standby appliance should have enough available capacity reserved to take over the forwarding load of the HA Active appliance.
- Always use an uplink redundancy strategy for complete redundancy (e.g. Application Path Resiliency or Multi-Uplink):
 - A failed uplink will not result in declaring an HA node failure if the heartbeats are received on one of the other interfaces. HA heartbeats are delivered using all existing networks to reduce false positives.

Figure 2-12. Active Network Extension HA Configuration



Service Mesh Availability During Upgrades

HCX provides two options for service mesh upgrades. The **Standard** option performs an appliance replacement using existing IP addresses. The **In-Service** option leverages Network Pool to reduce downtime by pre-establishing and validating the Transport Tunnels on the upgraded appliances.

In-Service Upgrade (ISSU)

The Network Extension appliance is a critical component of many HCX deployments, providing workload connectivity during a migration and in multi-cloud deployments.

The (ISSU) upgrade option enables the Network Extension appliances to receive critical patches and capability upgrades while minimizing impact to connected virtual machines. Network Extension ISSU is designed to recover Network Extension at maximum within a few seconds (sub-second recovery is possible when there are no resource constraints).

When In-Service upgrade or redeployment is selected, the following applies:

- A new NE appliance is provisioned at the source and destination site.

- New Uplink and Management IP addresses are assigned for each new Network Extension appliance.
- The NICs on the new appliances are connected, including bridge NICs for extended networks (flagged down)
- Secure tunnel connections are established and verified between the sites.
- The old appliance Bridge NICs are disconnected. And new Appliances Bridge NICs are connected.
- The old appliance is deleted. The IP addresses used for the old appliance are released back into the IP Pool.

Best Practices:

- Always use the ISSU option for Network Extension upgrade and redeploy operations to minimize downtime.
- Plan for additional IP addresses in the Network Profile IP Pools to accommodate the In-Service operations.
- Verify that HCX related firewall rules are configured for the full IP pool ranges.

Figure 2-13. ISSU New Appliances are Connected and Verified


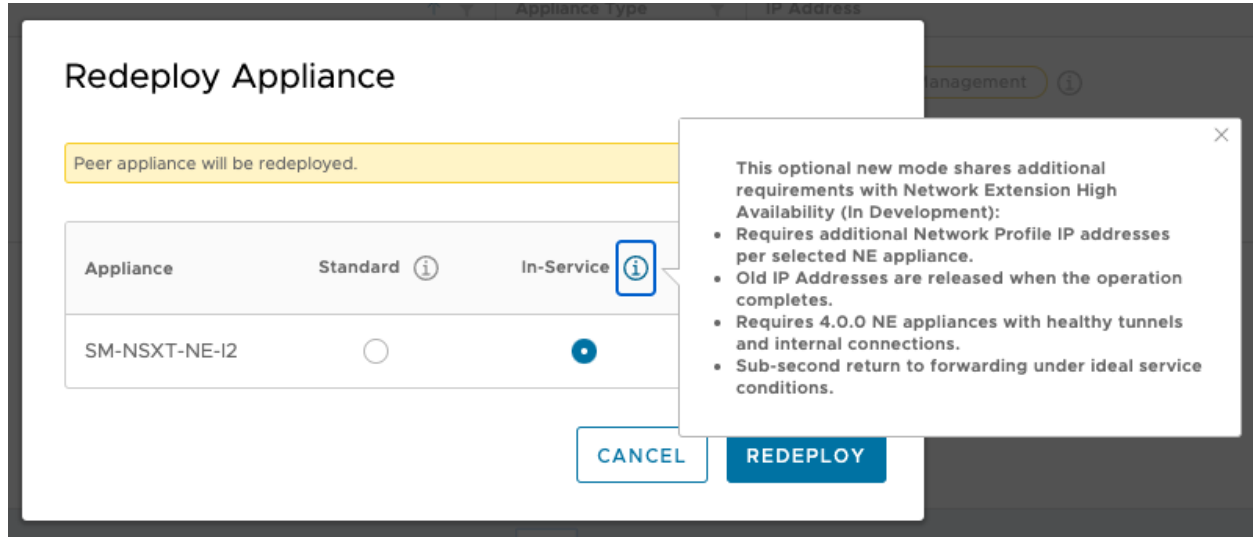


Figure 2-14. Selecting the In-Service Operation



About Standard Upgrades

The **Standard** option creates new Service Mesh appliances using the current IP addresses. It requires the new Service Mesh appliances to be fully disconnected until the final steps and requires tunnel negotiation to delay until the new appliances are connected.

When this option is used for Network Extension upgrades, virtual machine network connections will time out.

Best Practices:

- Use ISSU for Network Extension upgrades.
- Plan for a few minutes VM downtime when Network Extension appliances will be upgraded using the Standard option.

About the Author

3

Authors

Gabe Rosas is a Staff Technical Product Manager for VMware HCX in the Networking and Advanced Security Business Group at VMware. He is experienced in design and operation of network and server virtualization, in datacenter migration and multicloud technologies. Gabe is an active vExpert. You can connect with him in the VMware Communities (VMTN), on Twitter @gabe_rosas or follow his blog hcx.design.

Contributors

Michael Kolos is a Product Solution Architect in the VMware Cloud Product Management team. He specializes in networking and datacenter infrastructure, with over 20 years of experience in technology and IT roles with solution providers, ISPs, software, and financial services organizations. He has presented at VMworld, VMUGs and internal conferences on VMware cloud on AWS and is passionate about sharing knowledge with customers and colleagues.

Tom Zukowski is Senior Staff VCN Solutions Engineer supporting Enterprise accounts in the Great Lakes territory for VMware since 2015. Previously, a Systems Engineer for IBM, Brocade, Foundry Networks and 3Com helping customers implement a variety of networking technologies like Ethernet, WAN, Token Ring and Fiber Channel.

Bilal Ahmed is Customer Success Cloud Architect at VMware. He loves all things vSphere and has spent the last few years helping customers migrate their VMware workloads into private and public clouds. Bilal is VCDX #251 and is Office of the CTO, Ambassador, he is also a big Nandos and Batman fan.

Nathan Thaler is a Customer[0] in the VMware Cloud Product Management team. He serves as voice of the customer to help build better products and services by working closely with Product Management and Engineering. He has over 20 years experience various IT roles within startups and higher education and joined VMware in 2021. He was an early customer and design partner of HCM (now HCX) in 2015 and utilized the service for network extension and frictionless, high velocity migrations across multiple cloud providers.

Matt Elliot is a Staff Multi-Cloud Solution Architect at VMware with IT experience in several sectors, including healthcare, manufacturing, big law and VAR/MSP. He is passionate about all things networking, virtualization, automation, cloud, and monitoring. Matt is CCIE #56011.

Brian Taylor is a Senior Technical Writer with the VMware Information Experience team. Brian has experience documenting network, security, and storage technologies for Enterprise environments. He enjoys working across VMware organizations to deliver complete and concise customer content.

Christopher Dooks is a Senior Consultant in the VMware Professional Services Organisation. He has a background in backup, storage, and networking. His primary role is to help customers with their cloud journey, architecting multi-cloud deployments, their connectivity, and workload migrations.