HCX User Guide

VMware HCX 4.8

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

https://docs.vmware.com/

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

Contents

About the VMware HCX User Guide

The VMware® HCX™ *User Guide* describes how to plan for, install, and operate VMware HCX services in a vSphere data center. The information includes step-by-step configuration instructions and operational procedures.

Intended Audience

This information is for anyone who wants to install, upgrade, or use VMware HCX. The information is for Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms used in the VMware technical documentation, go to http://www.vmware.com/support/pubs.

Overview

1

VMware HCX is an workload mobility platform that is designed for simplifying workload migration, workload rebalancing, and business continuity across data centers and clouds.

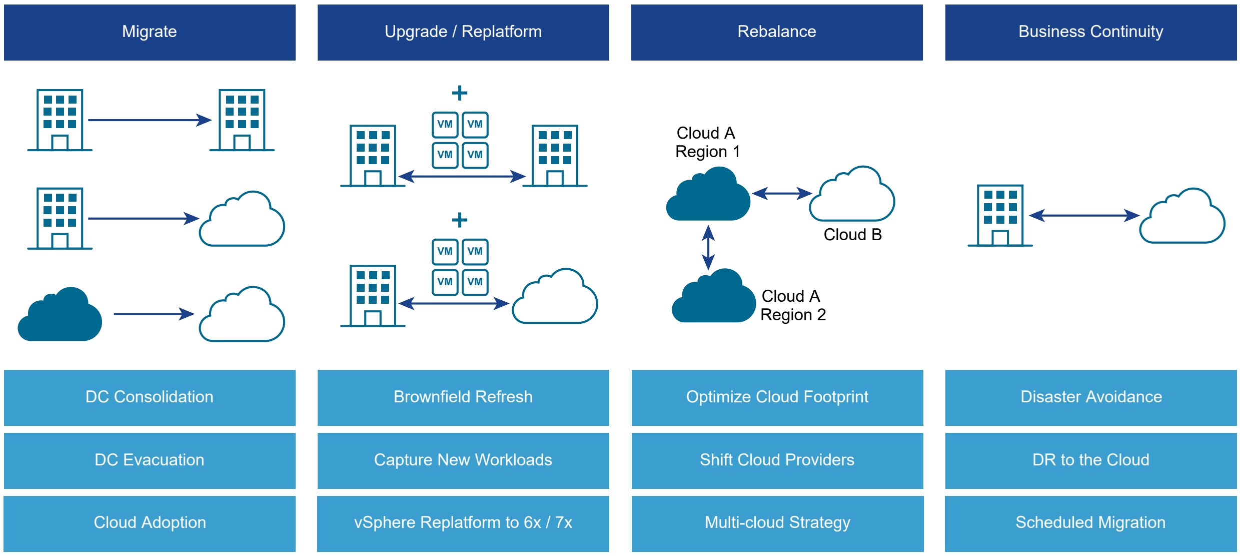

Migrate	Upgrade / Replatform	Rebalance	Business Continuity
DC Consolidation	Brownfield Refresh	Optimize Cloud Footprint	Disaster Avoidance
DC Evacuation	Capture New Workloads	Shift Cloud Providers	DR to the Cloud
Cloud Adoption	vSphere Replatform to 6x / 7x	Multi-cloud Strategy	Scheduled Migration

VMware HCX use cases:

- Application migration

  You can schedule and migrate thousands of vSphere virtual machines within and across data centers without requiring a reboot.

- Change platforms or upgrade vSphere versions

  With HCX, you can migrate workloads from vSphere and from non-vSphere (KVM and Hyper-V) environments within and across data centers or clouds to current vSphere versions without requiring an upgrade.

- Workload rebalancing

  HCX provides a multi-cloud mobility platform across cloud regions and cloud providers to allow customers to move applications and workloads at any time to meet scale, cost management, compliance, and vendor neutrality goals.

- Business continuity and protection

Using HCX capabilities, administrators can protect workloads by replicating them to other HCX sites. Workload migration is available on-demand, or it can be scheduled for business or for maintenance planning.

System Services

2

To build and maintain a secure optimized transport fabric between paired sites, and to provide feature rich functionality, HCX managed sites run various connectivity and mobility services.

HCX offers various services based on the type of license installed with the system.

HCX is available with an Advanced and an Enterprise license. HCX Advanced delivers basic connectivity and mobility services to enable hybrid interconnect and migration services. HCX Enterprise offers add-on functionality for scalability and performance when transforming large data centers or moving large quantities of virtual machines to cloud infrastructures.

HCX Advanced is a requirement for HCX Enterprise in all deployments except for VMware HCX for Cloud on AWS, which includes support for all HCX Advanced services and select HCX Enterprise features and services with no additional license requirement and at no additional cost.

For more information, see VMware HCX Licensing and Packaging Overview.

Advanced Services	Description
Interconnect	This service creates and secures connections between HCX installations, supporting management, migration, replication, and disaster recovery operations. This service is deployed as a virtual appliance.
WAN Optimization	The WAN Optimization service works with the HCX Interconnect service to improve the network performance through a combination of deduplication, compression, and line conditioning techniques. This service is deployed as a virtual appliance.
Network Extension	This service extends the Virtual Machine networks from an HCX source site to an HCX remote site. Virtual Machines that are migrated or created on the extended segment at the remote site are Layer 2 adjacent to virtual machines placed on the origin network. This service is deployed as a virtual appliance.
Bulk Migration	This service uses VMware vSphere Replication protocol to move virtual machines in parallel between HCX sites.

Advanced Services	Description
vMotion Migration	This migration method uses the VMware vMotion protocol to move a single virtual machine between HCX sites with no service interruption.
Disaster Recovery	The HCX Disaster Recover service replicates and protects virtual machines to a remote data center.

Enterprise Services	Description
Mobility Groups	This service supports assembling one or more virtual machines into logical sets for migration and monitoring as a group. Group migration provides the flexibility to manage migrations by application, network, or other aspects of your environment.
Mobility Optimized Networking (MON)	MON is an enterprise capability of the VMware HCX Network Extension (HCX-NE) feature. MON enables optimized workload mobility for virtual machine application groups that span multiple segmented networks or for virtual machines with inter-VLAN dependencies, as well as for hybrid applications, throughout the migration cycle. Migrated virtual machines can be configured to access the internet and cloud provider services optimally, without experiencing the network tromboning effect.
Network Extension High Availability	Network Extension High Availability protects extended networks from disruptions associated with Network Extension appliance downtime. Network Extension High Availability creates a redundant pair, or group, of connections for Network Extension appliances. In the event that a Network Extension appliance fails or is taken offline, the active connection fails over to the standby connection, and the extended network operations continue without interruption.
OS Assisted Migration	This migration service moves Linux- or Windows-based non-vSphere guest virtual machines from their host environment to a VMware vSphere data center. This service comprises two appliances. The HCX Sentinel Gateway appliance is deployed the source site, and the HCX Sentinel Data Receiver appliance at the destination site. This service also requires the installation of HCX Sentinel software on each guest machine.
Replication Assisted vMotion (RAV)	This service uses both VMware Replication and vMotion technologies for large-scale, parallel migrations with no service interruption.

Enterprise Services	Description
Workload Migration for NSX V2T	This service combines the wave orchestration and operational capabilities of HCX with the performance and concurrency characteristics of HCX Assisted vMotion to expedite virtual machine migration from an NSX for vSphere environment to an NSX environment.
Traffic Engineering ■ Application Path Resiliency ■ TCP Flow Conditioning	VMware HCX provides settings for optimizing network traffic for HCX Interconnect and Network Extension services. ■ The Application Path Resiliency service creates multiple tunnel flows, for both Interconnect and Network Extension traffic, those may follow multiple paths across the network infrastructure from the source to the destination data centers. The service then intelligently forwards traffic through the tunnel over the optimal path and dynamically switches between tunnels depending on traffic conditions. Application Path Resiliency forwards traffic over one tunnel at a time and does not load balance across multiple paths. ■ The TCP Flow Conditioning service adjusts the segment size during the TCP connection handshake between end points across the Network Extension. This optimizes the average packet size to reduce fragmentation and lower the overall packet rate.

System Components

3

VMware HCX comprises a virtual management component at both the source and destination sites, and up to five types of VMware HCX Interconnect service appliances depending on the HCX license. VMware HCX services are configured and activated at the source site, and then deployed as virtual appliances at the source site, with a peer appliance at the destination site.

HCX Connector and HCX Cloud Installations

In the HCX site-to-site architecture, there is a notion of a source and a destination environment. Depending on the environment, there is a specific HCX Manager (HCX) installer: HCX Connector or HCX Cloud. HCX Connector is always deployed as the source. HCX Cloud is typically deployed as the destination, but it can be used as the source in cloud-to-cloud deployments. In HCX-enabled public clouds, the cloud provider deploys HCX Cloud. The public cloud tenant deploys HCX Connector on-premises.

The source and destination sites are paired together for HCX operations.

Note An HCX Connector cannot be paired with another HCX Connector.

In both the source and the destination environments, HCX Manager is deployed to the management zone, next to each site's vCenter Server, which provides a single plane for administering VMware HCX. HCX Manager provides a framework for deploying HCX service virtual machines across both the source and destination sites. HCX administrators are authenticated with the Cloud Services Portal, and each task authorized through the existing vSphere SSO identity sources.

HCX mobility, extension, protection actions can be initiated from the HCX User Interface or from within the vCenter Server Navigator screen's context menus.

With HCX, you initiate mobility and connectivity actions from the HCX Console.

In the NSX Data Center Enterprise Plus (HCX for Private to Private deployments), the tenant deploys both the source and the destination HCX Managers.

HCX-IX Interconnect Appliance

The Interconnect service (HCX-IX) appliance provides replication and vMotion-based migration capabilities over the Internet and private lines to the destination site while providing strong encryption, traffic engineering, and virtual machine mobility.

The Interconnect appliance includes deployment of the Mobility Agent service that appears as a host object in the vCenter Server. Mobility Agent is the mechanism that HCX uses to perform vMotion, Cold, and Replication Assisted vMotion (RAV) migrations to a destination site.

WAN Optimization Appliance

The WAN Optimization (WAN-OPT) appliance improves performance characteristics of the private lines or Internet paths by applying WAN optimization techniques like the data de-duplication and line conditioning.

Network Extension Appliance

The Network Extension (IX-NE) appliance provides layer 2 connectivity between sites. Network Extension provides the ability to keep the same IP and MAC addresses during virtual machine migrations. When the Network Extension service is enabled, a pair of virtual appliances is deployed: one in the source and one in the destination site.

Sentinel Gateway Appliance

Using OS Assisted Migration (OSAM) service, you can migrate guest (non-vSphere) virtual machines from on-premise data centers to the cloud. The OSAM service has several components: the HCX Sentinel software that is installed on each virtual machine to be migrated, a Sentinel Gateway (SGW) appliance for connecting and forwarding guest workloads in the source environment, and a Sentinel Data Receiver (SDR) in the destination environment.

HCX Sentinel Data Receiver Appliance

The Sentinel Data Receiver (SDR) appliance works with the HCX Sentinel Gateway appliance to receive, manage, and monitor data replication operations at the destination environment.

Preparing for Installation

4

This section describes the system requirements, network ports, and protocols that must be allowed and various other requirements, like software versions and feature interoperability requirements.

Read the following topics next:

System Requirements

Before deploying HCX, consider the required resources for both the source and the destination environment site managers.

Virtual Hardware Requirements for HCX Manager Appliances

Appliance	vCPU	Memory	Disk Space/IOPS
HCX Manager	4	12 GB	60 GB
HCX Interconnect (HCX-IX)	8	6 GB	2 GB
HCX Network Extension (HCX-NE)	8	3 GB	2 GB
HCX WAN Optimization (HCX-WAN-OPT)	8	14 GB	100 GB / 5000 IOPS

Appliance	vCPU	Memory	Disk Space/IOPS
HCX Sentinel Gateway (HCX-SGW [source only])	8	8 GB	■ HCX: 21 GB (disks: 2 GB, 6 GB, 4 GB, and 9 GB)
HCX Sentinel Data Receiver (HCX-SDR [destination only])	8	8 GB	■ HCX: 21 GB (disks: 2 GB, 6 GB, 4 GB, and 9 GB)

Note The storage requirement per appliance is doubled during the upgrade and redeploy operations, as a second appliance is created for operation.

Migration Host and Datastore Requirements

During a virtual machine migration operation, the destination compute and storage resources are selected. These resources must meet the following general requirements:

Requirement ID	Requirement	Description
Host-Req-1	Host CPU and Memory Capacity	The destination host must support the original vCPU and Memory configurations of the virtual machine.
Host-Req-2	Host CPU and Memory Usage	The destination host capacity must satisfy admission control requirements.
Host-Req-3	Host Compatibility with VMTools	The destination hosts must support the VMTools version of the original virtual machine. For more information see the VMware Certified Compatibility Guides.
Host-Req-4	Host Compatibility with VM Hardware	The destination hosts must support the VM Hardware version of the original virtual machine. For more information see VMware KB 2007240.

Requirement ID	Requirement	Description
Datastore-Req-1	Datastore Type	The destination datastore type must be VMFS5, VMFS6, NFS3, NFS4.1, or vSAN.
Datastore-Req-2	Datastore Capacity	The destination datastore capacity must be able to store the following: ■ Complete original virtual machine data (everything that makes up the virtual machine) ■ Additional replica instance vmdk file on target datastore (applicable for Bulk or RAV Migration during the delta sync) ■ Redo logs (virtual machine data changes over a one hour interval) ■ Additional data for multiple point in time (MPIT) recovery points ■ Full copy of the virtual machine when using test recovery ■ With VSAN, additional data required by the Primary level of failures to tolerate (PFTT) configuration

System Requirements for Scaling Out Deployments

Consider the following requirements when scaling out HCX deployments:

Note For information regarding the scaling of HCX resources, see VMware KB article 93605.

- When scaling HCX deployments, the hardware requirements apply for each additional appliance. For appliance limit information, refer to the VMware Configurations Maximums.
- In environments with multiple vCenter Servers, a maximum of one site manager can be deployed per vCenter Server.
- By default, the migration service encompasses all source and destination clusters in a service mesh with a single IX appliances. The migration service can be scaled out when there are additional clusters in either the source or destination environments (you can create a service mesh for each unique source/destination cluster pair).

  For example, if there is source site with clusters A and B paired to a destination cloud with clusters X, Y and Z, you can scale out by creating service meshes AX, AY, AZ, BX, BY, BZ.

  Note In this scale out example, cluster A is hosting three IX appliances, and you can configure host anti-affinity to enable parallel switchovers with HCX vMotion and RAV migrations.
- Multiple Network Extension appliances can be deployed per distributed virtual switch (DVS) or NSX Transport Zone.

- When using Network Extension High Availability (HA), a maximum of two Network Extension appliance can be deployed for each VLAN or NSX segment. In the HA group, one appliance is Active and one is Standby.

  Note HA Standby appliances must be considered when calculating the overall hardware resource requirements.

- When not using Network Extension HA, a maximum of one Network Extension appliance can be deployed for each VLAN or NSX segment.
- The Network Extension appliance count in the Service Mesh configuration must equal the sum of all planned standalone, HA Active, and HA Standby appliances.

Software Version Requirements

HCX Manager appliances must be running the supported software versions.

Software Version Requirements for HCX Installations

Note In cloud-to-cloud deployments, the HCX Cloud environment requirements apply to both sides of a site pair.

Note For granular compatibility information, select VMware HCX in the VMware Product Interoperability Matrix.

For information about General Availability (GA), End of Support (EoS), End of Technical Guidance (EoTG) for VMware software, see VMware Product Lifecycle Matrix.

Component Type	HCX ConnectorEnvironment Requirements	HCX Cloud
HCX	For private cloud installations, use latest HCX Connector version. For new public cloud installations, a compatible version of the HCX Connector is provided as a download link in the System Updates page of the HCX Cloud Manager interface (:443).	■ For private cloud installations, use latest HCX Cloud Manager version. ■ For new public cloud installations, the latest HCX Cloud Manager version is provided for the hypervisor environment by VMware.
vSphere version	vSphere versions within Technical Guidance or newer. For details about the evacuation of versions that are no longer generally available, see VMware KB article 82702.	Generally available vSphere versions listed in the VMware Product Interoperability Matrix.
NSX-T	NSX-T versions within Technical Guidance or newer. See VMware KB article 82702.	Generally available NSX versions listed in the VMware Product Interoperability Matrix.

Component Type	HCX ConnectorEnvironment Requirements	HCX Cloud
NSX for vSphere	Provided for evacuation only. Refer to VMware KB article 82702.	Not supported.
Cloud Director	Not supported at the source.	Generally available Cloud Director versions listed in the VMware Product Interoperability Matrix.

Additional Software Version Considerations

This section highlights requirements currently in the footnotes of the *VMware Product Interoperability Matrix* for clarity.

HCX Component	Version Requirements
HCX Connector with End of Support software	vSphere and NSX versions in the End of Support phase are supported only with HCX Connector installations, for evacuation purposes, until the End of Technical Guidance. Refer to VMware KB article 82702.
HCX Cloud Manager with vSphere 7.x	NSX-T 3.0.1 or later.
HCX Cloud Manager with VMware Cloud Director	NSX-T 3.1.1 or later.
Mobility Optimized Networking (MON)	NSX-T 3.0 or later.
Replication Assisted vMotion (RAV)	vSphere 6.5 U3F+ or vSphere 6.7 U3+.
Coexistence with vSphere Replication	vSphere Replication 8.1 or later.

Network Port and Protocol Requirements

HCX deployments require allowing various ports for communication between services on the HCX Manager appliance itself and between HCX pairs at the source and destination sites.

You must allow the following connections in HCX Manager deployments:

- The perimeter firewalls must allow HCX Manager connections to connect.hcx.vmware.com.
- The perimeter firewalls must allow HCX Manager connections to hybridity-depot.vmware.com.
- The source site firewalls must be configured to allow outbound connections to the destination HCX Manager systems.
- The destination site firewalls must be configured to allow inbound connections from the source HCX Manager system.
- All local connections (within a single HCX Manager) either at the source or destination environment. These connections never traverse from source to destination or from destination to source.
- A proxy server can be configured for HTTPS connections. Refer to Configure a Proxy Server.

- Connections made when the HCX Manager is added as a solution in a vRealize Operations installation.

For a complete list of network port and protocol requirements, see VMware Ports and Protocols.

Network Underlay Minimum Requirements

The infrastructure providing network connectivity to an HCX deployment (the Underlay) must meet the minimum requirements. The underlay includes any intermediate system that is customer managed, cloud provider managed, or part of an Internet service provider network.

What Is the Network Underlay

A network underlay provides the physical or logical connectivity on which HCX transport packets are tunneled, where an HCX transport packet contains an overlay header. The underlay network does not need to know that it is carrying HCX transport packets. This includes the physical routing infrastructure on the customer data center and (if applicable) the cloud provider infrastructure, and any physical network services joining the connected locations.

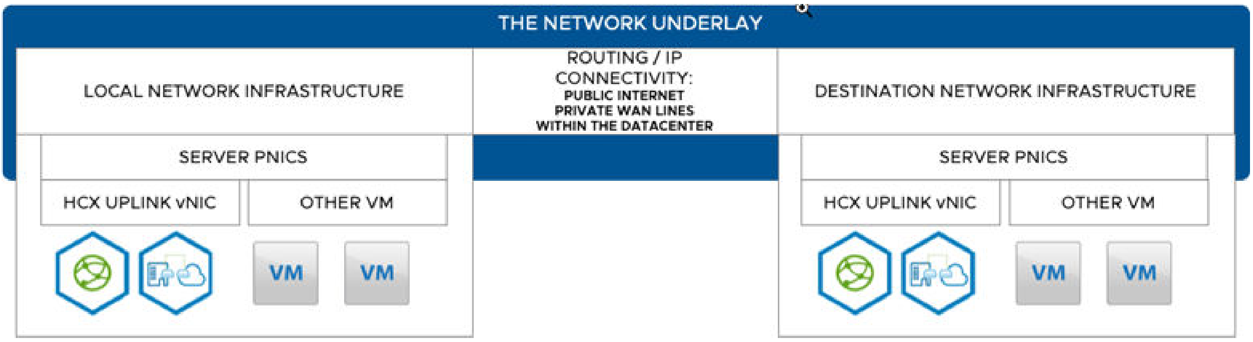

A network underlay can vary from high-bandwidth low-latency private paths between server racks in a data center, to lower-bandwidth higher-latency Internet based connectivity. In this document, the term "network underlay" encompasses all elements that affect the performance characteristics of an underlay, including the servers and network devices and connections between the vSphere environments. The network underlay requirements must be satisfied when considering all elements of the underlay.

VPN Based Network Underlays

Virtual Private Networks are frequently used for creating secure network connection to private and public vSphere clouds over the Internet. SDWAN and custom tunneling solutions are used over the internet to improve data traffic transmissions. The SDWAN, VPN, and other tunneling solutions are collectively referred to as VPN in this document.

The network underlay includes connections with VPN configurations. The network underlay requirements must be satisfied when considering all elements of the underlay.

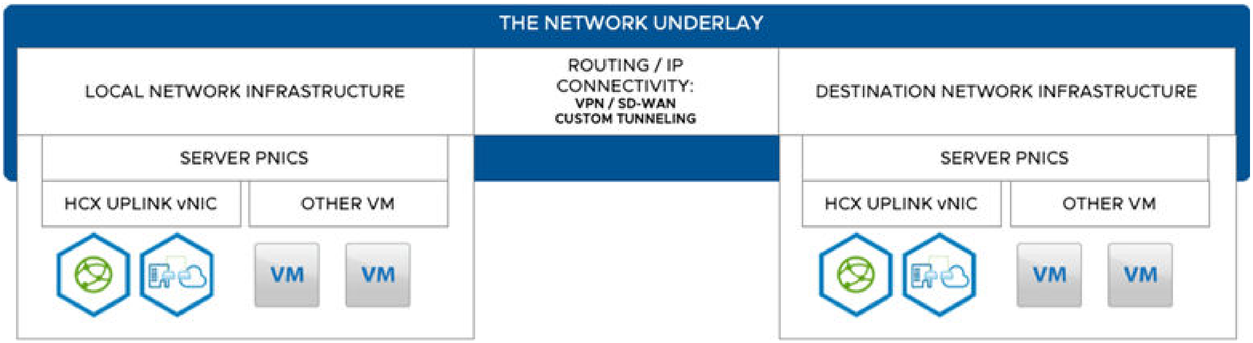

General Network Underlay Requirements

HCX supports multiple uplinks and each uplink can be connected to a different network underlay. Examples of different network underlays include private line, public Internet, and multi-homed connectivity.

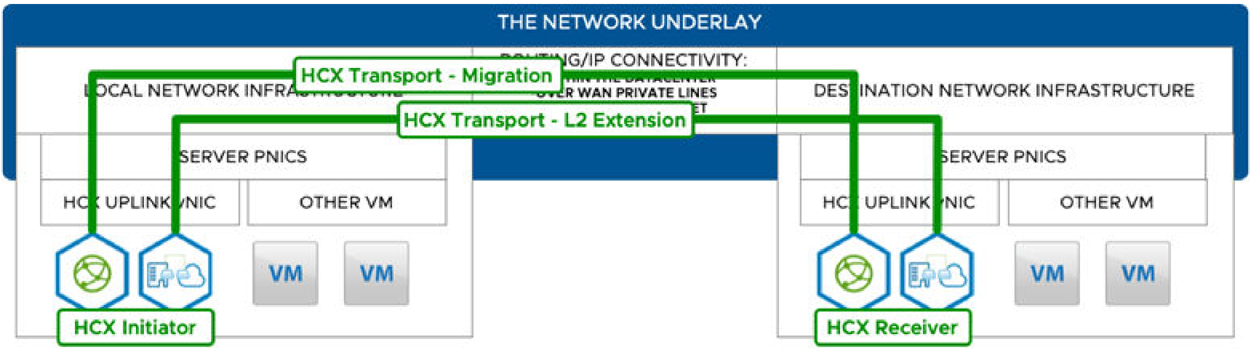

The following table summarizes the requirements from the Network Underlay to use the HCX migration and extension services:

- This table applies to HCX Migration and Extension Overlays (HCXService Mesh appliances).
- This table does not apply to HCX Connector or HCX Cloud Manager or the management connections.

HCX Requirement ID	Requirement Summary	Requirement Details
hcx-overlay-req-1	IP Addressing & IP Reachability	Requires a valid IP address and IP connectivity for end to end communication between the HCX Uplink IP.
hcx-overlay-req-2	Bandwidth, Loss and Latency, MTU	All underlays must comply with minimum parameters requirements for services to be supported at the minimum performance level. The minimum requirement applies to all network underlays and is provided in the next table. MTU configuration must be applied to the HCX Site Resource Profile prior to deploying the IX/WO appliances. If the MTU is changed on existing appliances, the appliances must be redeployed. ■ MTU 1150 – 9000 is valid for IX (no WANOPT). ■ MTU 1150 – 1500 is valid for IX (with WANOPT).
hcx-overlay-req-3	Source Network Address Translation (SNAT)	SNAT is not required, but can be used to translate HCX Uplink private IP packets to public IP addresses for connections over the Internet. SNAT can only be applied to the HCX Initiator (the HCX source appliances).
hcx-overlay-req-4	Destination Network Address Translation (DNAT) Load Balancing Reverse Proxy	Inbound DNAT, load balancing, or reverse proxy configurations in the underlay are not supported for the HCX Migration and Extension Transport tunnels.
hcx-overlay-req-5	VPN	Any VPN configuration in the network path is treated agnostically as an underlay, and is supported while the measured underlay parameters meet the documented requirements. Any additional encapsulation and performance degradation, overhead, or cost in addition to the characteristics of the underlay they ride on must be considered when measuring underlay outcomes. HCX does not support VPN configurations where the NSX Tier-0 router provides the VPN termination AND connectivity to the HCX Manager uplinks through NSX Service Insertion.

Minimum Network Underlay Requirements for HCX

HCX has network underlay minimums for HCX migration and disaster recovery operations. HCX operations with lesser performance than the minimum values are not supported.

The table below lists the minimum requirements for individual operations at minimum performance. For vMotion and Replication Assisted vMotion, the bandwidth requirement varies based on whether the WAN-Optimization service is running in the Service Mesh for the site pair. WAN Optimization can improve the network performance through a combination of deduplication, compression, and line conditioning techniques..

In addtion to activating the WAN Optimization service in the Service Mesh, review these additional considerations related to network bandwidth performance:

- Minimizing latency, loss, jitter can result in improved migration performance outcomes.
- Parallel HCX operations (migration and extension) can result in increased bandwidth requirements.

Network Parameter	HCX vMotion	Replication Assisted vMotion	Bulk Migration & DR (Protection)	OS Assisted Migration
Min Bandwidth (Mbps)	150 with WAN Optimization 250 without WAN Optimization	150 with WAN Optimization 250 without WAN Optimization	50	50
Min MTU (bytes) (1350 if version < HCX 4.2)	1150	1150	1150	1150
Max Packet Loss (%)	0.1	0.1	1.0	1.0
Max Latency (ms)	150	150	150	150

Bandwidth distribution with HCX Manager on a network underlay can be visualized as a set of nested pipes, where the underlay network is the main channel. HCX Manager and non-HCX Manager traffic is carried through main channel. Migration and Network Extension traffic can be thought of as separate pipelines through the site manager channel. The Network Extension pipe provides the throughput for all the extended network traffic. The migration pipe handles the vMotion, Bulk, Protection, and OS Assisted Migration service traffic.

Note The number of parallel migrations allowed depends on the bandwidth of the migration pipe.

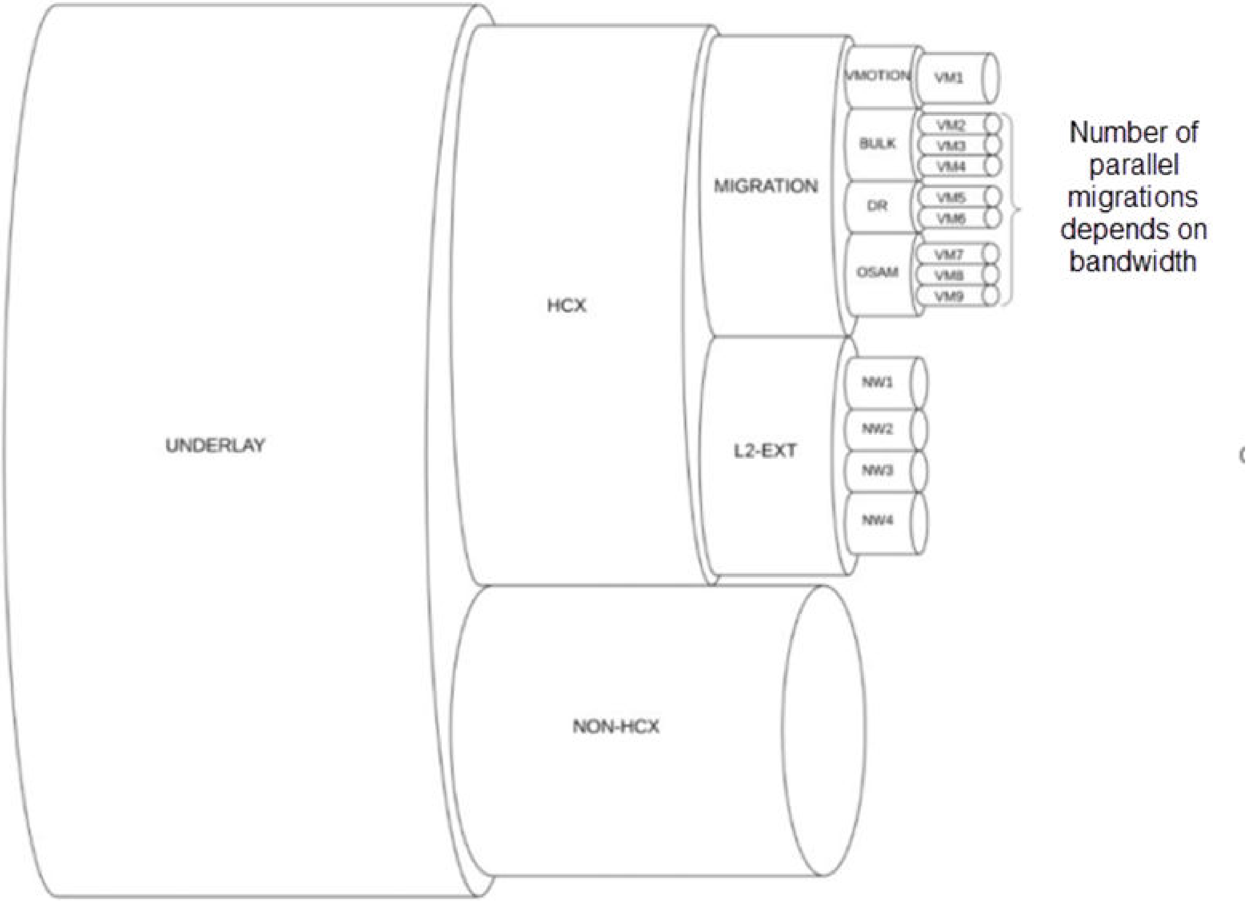

HCX Manager User Account and Role Requirements

HCX Manager (HCX) configuration and operation requires an understanding of the various accounts and roles involved in deploying, managing, and operating the system.

User Accounts

HCX has the following account requirements for site manager deployments:

Account	Requirements	Additional Information
admin	- The admin password must be set. - The root password must be set.	- Created during the site manager OVA deployment. - Used in the Appliance Management interface (https://*hcx-ip-or-fqdn*:9443) - Used for CLI/terminal shell access.
Account for vCenter Server Registration	The account must belong to the vSphere administrators group and have the administrator role assigned.	- The administrator@vsphere.local account is suggested by default, but not required. - Alternate vSphere SSO local users that meet the requirements can be used. - Active Directory service accounts that meet the requirements can be used.
Account for NSX Registration	This account must have the Enterprise Admin role assigned. If NSXv, this account must have the Enterprise Administrator role assigned.	- The NSX admin account is suggested by default, but not required. - Alternate NSX local accounts that meet the requirements can be used. - Active Directory service accounts that meet the requirements can be used. - Prior to NSX-T Data Center 3.0, it is mandatory to use the NSX admin account. **Note** This account is not required for HCX Connector installations. It is required only when extending NSX Segments, or migrating NSX tags.
Account for vCloud Director Registration	The account must have the System Administrator role assigned.	- The VMware Cloud Director sysadmin account is suggested by default, but not required. - An alternate local account that meets the requirements can be used. - LDAP service accounts that meet the requirements can be used. **Note** This account is only required for provider installations of VMware HCX with vCloud Director. A tenant does not require this account.

Account	Requirements	Additional Information
Accounts for HCX Role Mapping (This account refers to SSO User accounts that map to an HCX role.)	The user's group must be included in the HCX Role Mapping configuration.	■ HCX supports two user roles: Administrator and Tenant. ■ The Administrator role is for users who configure and operate HCX (create and manage the Compute Profiles, Site Pairings, Service Meshes, Network Extensions, Migrations, and DR operations). ■ The Tenant role is for Service Provider installations only. This role does not support adding or deleting Network Profiles. ■ The vsphere.local\Administrators vSphere SSO Group is added by default to HCX Administrator. However, it is not mandatory to use this SSO group. For the HCX Tenant role, no default group is provided. ■ A common practice is to create an hcx-administrators vSphere SSO Group. SSO and Active Directory users are populated into the hcx-administrators vSphere SSO group. The default vsphere.local\Administrators HCX Administrator user group entry in the Role Mapping configuration is replaced with the new hcx-administrators vSphere SSO group.
Site Pairing Accounts	The user's group must be included in the HCX Role Mapping configuration (on the remote HCX Cloud system being paired). The user's group can be in either the Administrators group or the Tenant group.	The site pairing user is entered along with the HCX Cloud URL in the site pairing configuration on the source HCX system. The following are typical scenarios: ■ In a private data center HCX deployment, the site pairing user is traditionally the administrative user for the destination vSphere environment. ■ In a dedicated public cloud HCX deployment, the site pairing user is traditionally the SDDC administrator account provided to the tenant. ■ In a VMware Cloud Director deployment, the site pairing user is the Organization Administrator account.

Note The vCenter Server and the NSX Manager registration accounts (“service accounts”) must have global object access.

VMware Cloud Director registration accounts must also have global object access.

Role Mapping

Access to HCX services and features depends on the assigned user role. User roles are assigned in the HCX appliance management interface during the initial HCX activation and configuration.

Administrator

SSO groups assigned to the Administrator role have unrestricted access to perform all HCX configurations and operations.

Tenant

This role is intended for use by Service Providers. SSO groups assigned to the Tenant role cannot add or delete HCX Network Profiles.

Note The Tenant role is not available in HCX Connector deployments.

vSphere Privileges for Migration Operations

User groups assigned to the Administrator or the Tenant role must have these vSphere vCenter Server privileges to perform migrations.

vCenter Resource Type	User Privilege	Description
ComputeResource	■ VirtualMachine.Inventory.Create ■ VirtualMachine.Interact.PowerOn ■ VirtualMachine.State.CreateSnapshot ■ VirtualMachine.State.RemoveSnapshot	Privileges required on the destination compute resource object when performing a migration operation.
HostSystem	■ VirtualMachine.Inventory.Create ■ VirtualMachine.Interact.PowerOn ■ VirtualMachine.State.CreateSnapshot ■ VirtualMachine.State.RemoveSnapshot	Privileges required on the destination HostSystem object when performing a migration operation.
ClusterComputeResource	■ VirtualMachine.Inventory.Create ■ VirtualMachine.Interact.PowerOn ■ VirtualMachine.State.CreateSnapshot ■ VirtualMachine.State.RemoveSnapshot	Privileges required on the destination ClusterComputeResource object when performing a migration operation.

vCenter Resource Type	User Privilege	Description
ResourcePool	■ VirtualMachine.Inventory.Create ■ VirtualMachine.Interact.PowerOn ■ VirtualMachine.State.CreateSnapshot ■ VirtualMachine.State.RemoveSnapshot	Privileges required on the destination ResourcePool object when performing a migration operation.
Folder	■ VirtualMachine.Inventory.Create ■ VirtualMachine.Interact.PowerOn ■ VirtualMachine.State.CreateSnapshot ■ VirtualMachine.State.RemoveSnapshot	Privileges required on the destination **Folder** object when performing a migration operation.
Datacenter	■ VirtualMachine.Inventory.Create ■ VirtualMachine.Interact.PowerOn ■ VirtualMachine.Interact.PowerOff ■ Resource.HotMigrate ■ Resource.ColdMigrate ■ Folder.Create ■ Folder.Delete	Privileges required on the destination data center objects when performing a migration operation.
Datastore	■ Datastore.UpdateVirtualMachineMetadata ■ Datastore.DeleteFile	Privileges required on the destination datastore objects when performing a migration operation.
DistributedVirtualPortgroup/Network	Network.Assign	Privileges required on the destination network objects when performing a migration operation.
VirtualMachine	■ VirtualMachine.Interact.PowerOn ■ VirtualMachine.Interact.PowerOff ■ Resource.HotMigrate ■ Resource.ColdMigrate ■ VirtualMachine.State.CreateSnapshot ■ VirtualMachine.State.RemoveSnapshot ■ VirtualMachine.Hbr.ConfigureReplication ■ VirtualMachine.Hbr.MonitorReplication	Privileges required on the source virtual machines when performing a migration operation.

Configuration and Service Limits

When you are configuring, deploying, and operating VMware HCX, you must stay within the supported limits.

When using HCX, it's important to be aware of the system and operational limits in the following categories:

- Sites and Service Components
- Migrations
- Disaster Recovery
- Network Extension
- Migration-centric Virtual Machine Limits

For a detailed list of the system limits, refer to the VMware Configurations Maximum tool.

Using the Interface

You can access HCX services and system-level operations through one of several interfaces.

VMware HCX has the following user interfaces:

Interface	Description
vSphere Client	You can perform all operations related to HCX services from the HCX Plug-in in the vSphere Client. This interface is not available in VMware Cloud Director clouds.
HCX Manager UI	At both the source or destination site, you perform all operations related to HCX services by logging in to the HCX Manager at <https://hcx-ip-or-fqdn:443>.
HCX Manager Appliance Management UI	At both the source or destination site, you perform system-level management, licensing, and upgrade operations by logging in to the HCX Manager appliance system interface at https://*hcx-ip-or-fqdn*:9443.
Central CLI	You access the CCLI for debugging or troubleshooting HCX issues with VMware representatives. For access to the CCLI, see Logging in to the HCX Manager Shell.

HCX Activation and Licensing

The HCX service features are available based on the installed license.

HCX licenses are available in two types: Advanced and Enterprise. The Advanced license is packaged with NSX Data Center Enterprise Plus, VMware Cloud on AWS, VCF Enterprise and from VMware Cloud Provider Partners. The HCX Enterprise license is available for purchase to NSX Enterprise Plus customers. For a list of services available with each license type, see Chapter 2 System Services.

Activating or Licensing New HCX Systems

Activation requirements refer to any information required to activate a newly deployed HCX system.

HCX Activation Requirements

The HCX system must be activated before it can trigger services like migration and extension.

During the initial configuration of the HCX Manager, the wizard displays an activation screen.

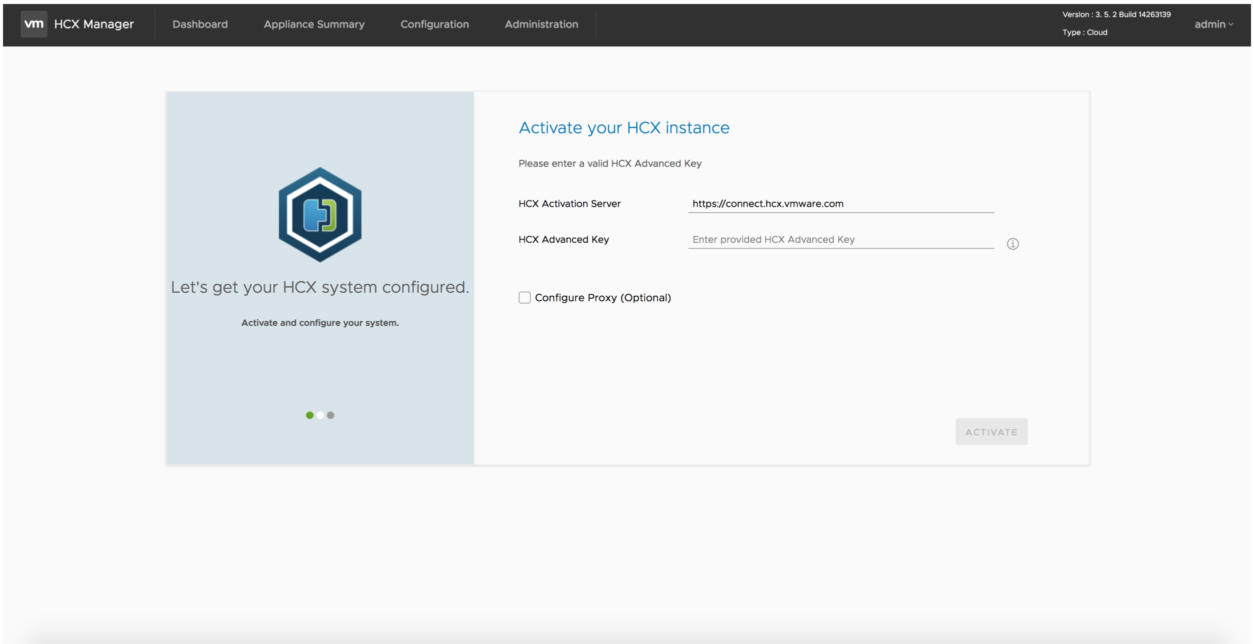

Version : 3. 5. 2 Build 14263139

Type : Cloud

Activate your HCX instance

Please enter a valid HCX Advanced Key

HCX Activation Server: https://connect.hcx.vmware.com

HCX Advanced Key: Enter provided HCX Advanced Key

□ Configure Proxy (Optional)

Requirement	Details
Activating the system requires network access from the HCX Manager system to https://connect.hcx.vmware.com and a valid activation key or license key.	To test connectivity from the HCX Manager, use SSH to connect to the HCX Manager shell. `curl -k -v https://connect.hcx.vmware.com`
Network access from the HCX Manager system to `https://connect.hcx.vmware.com` when there is a proxy for outbound HTTPS connections.	If there is a proxy server in the environment, the proxy server must be configured on the HCX Manager. The proxy settings can be configured in the Administration interface. To resume the Initial Configuration Wizard, click the dashboard tab. Caution By default, when you configure a proxy server, the system uses that server for all HTTPS connections, including the local vCenter Server, ESXi, NSX, and HCX-IX. For a successful deployment, define all related proxy exceptions when you configure a proxy server.

Requirement	Details
Activating an HCX Connector in a private cloud HCX installation.	An HCX Connector does not require unique activation keys. It uses the same HCX Advanced and HCX Enterprise licenses used on the destination HCX Cloud. Enter NSX Datacenter Enterprise Plus license when prompted for the **HCX Advanced Key**. The **HCX Enterprise Key** can be added to the HCX Connector after providing an HCX Advanced license key.
Activating an HCX Cloud system in a private cloud when using private vSphere with NSX as the destination environment. Activating HCX Cloud systems when using Cloud Director as the destination environment.	HCX Advanced is included with NSX Data Center Enterprise Plus. Use NSX Data Center Enterprise Plus licenses from my.vmware.com. Enter this license when prompted for the HCX Advanced Key. NSX Data Center Enterprise Plus evaluation licenses can be used, but they must be updated to full keys for operations exceeding the trial limits.
Activating HCX Connector with VMware Cloud on AWS. Activating HCX Connector with VMware Cloud on AWS GovCloud.	Obtain the activation keys for the HCX system following the cloud service provider procedures. . To obtain the HCX Connector activation key for VMware Cloud services, follow this procedure: 1 Log in to the VMware Cloud console: ■ VMware Cloud on AWS ■ VMware Cloud on AWS GovCloud 2 For your SDDC, Navigate to **View Details > AddOns**, and click **Open HCX**. 3 Select **Activation Keys > Create Activation Key**. 4 For System Type, select **HCX Connector**. 5 Click **Confirm**.
Activate Later	This option allows HCX activation to be temporarily skipped. To complete the installation while waiting for proxy or firewall allow additions, choose this option. The activation keys can be entered in the Appliance Management Configuration interface.
Grace Period	A small grace period allows the installation of HCX components. After the grace period expires, the system stops all associated services and operations. After the installation, the HCX systems must maintain an outbound connection to the central service URL, `connect.hcx.vmware.com`.

Updating an HCX Evaluation License Key

You can update VMware HCX installations using evaluation or trial activation keys to use a standard HCX Advanced License key.

Note This procedure is applicable to both the source and destination HCX systems activated with NSX Data Center Enterprise Plus trial licenses (or expiring licenses).

This procedure, however, does not apply for HCX systems connecting with a VMware HCX-activated public cloud.

Prerequisites

- Administrative access to the HCX system.
- NSX Data Center Enterprise Plus purchased license.

Procedure

1. Navigate to the HCX Appliance Management interface: `https://hcx-ip-or-fqdn:9443`.
2. Navigate to the **Configuration** tab.
3. Select **License** on the side menu and click **Edit**.
4. Enter the new HCX Advanced license (NSX Enterprise Plus key), and click **UPDATE**.

Removing or Adding the HCX Enterprise Upgrade Key

You can update evaluation or trial activation keys to use a premium HCX Enterprise license, or if no license exists, you can add the HCX Enterprise license.

This procedure is applicable to both the source and destination HCX Manager systems.

Prerequisites

- Administrator access to the HCX Manager system.
- HCX Enterprise purchased license.

Procedure

1. Navigate to the Appliance Management Interface `https://hcx-ip-or-fqdn:9443`.
2. Navigate to the **Configuration** tab.
3. Select **License** on the side menu and click **Edit**.
4. Remove or add the HCX Enterprise license key:
   - Remove an HCX Enterprise license key.

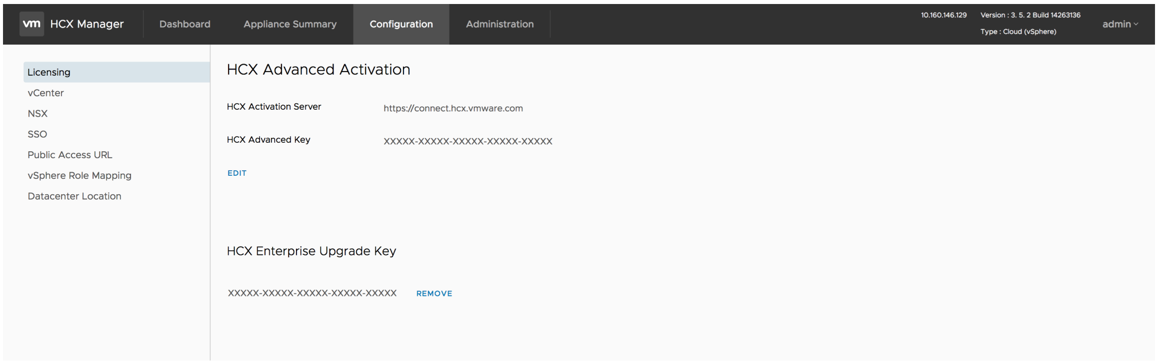

10.160.146.129 Version : 3. 5. 2 Build 14263136
Type : Cloud (vSphere)

HCX Advanced Activation

HCX Activation Server https://connect.hcx.vmware.com

HCX Advanced Key XXXXX-XXXXX-XXXXX-XXXXX-XXXXX

HCX Enterprise Upgrade Key

XXXXX-XXXXX-XXXXX-XXXXX-XXXXX

a Click **REMOVE** to remove the existing license key.

◆ Add an HCX Enterprise license key.

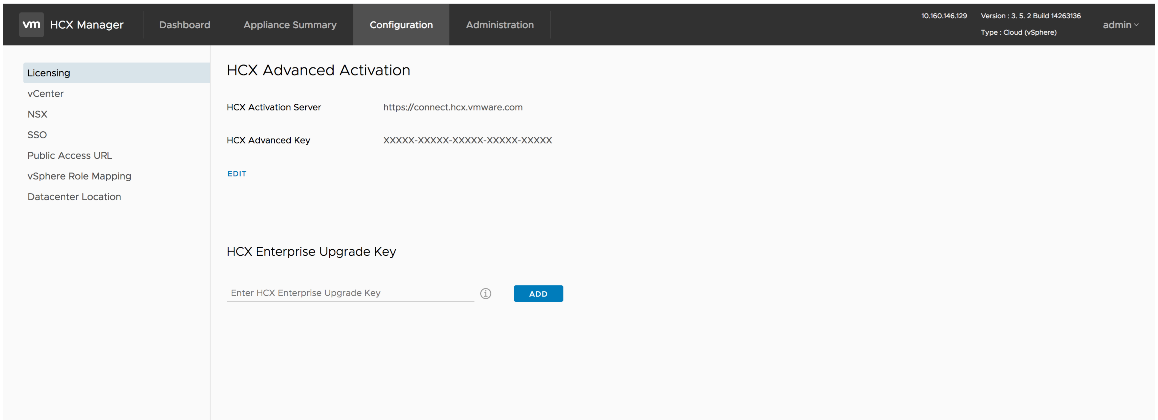

10.160.146.129 Version : 3. 5. 2 Build 14263136
Type : Cloud (vSphere)

HCX Advanced Activation

HCX Activation Server https://connect.hcx.vmware.com

HCX Advanced Key XXXXX-XXXXX-XXXXX-XXXXX-XXXXX

HCX Enterprise Upgrade Key

a Enter the HCX Enterprise license, and click **ADD**.

NSX Requirements for HCX Manager Deployments

In VMware HCX installations connecting private environments, NSX must be installed and configured before deploying HCX Manager. This section details the requirements.

NSX Requirements for HCX Deployments

For environments requiring NSX virtual networking, you must install and configure NSX, including integration with the vCenter Server, before deploying HCX Manager.

- In the destination environment, the NSX Manager must be installed and integrated with the vCenter Server. For a list of supported NSX versions, see the VMware Product Interoperability Matrix.
- An NSX Data Center Enterprise Plus license is required to activate HCX systems and provide access to HCX Advanced features.

Note HCX Advanced is a requirement for HCX Enterprise in all deployments except for VMware HCX for VMware Cloud on AWS, which includes support for all HCX Advanced services and select HCX Enterprise features and services.

- The NSX Manager must be registered during the HCX Manager install with the admin user.
  - If the NSX Manager IP or FQDN uses self-signed certificates, it might be necessary to trust the NSX system manually using the Import Cert by URL interface in the HCX Appliance Management interface.
- HCX requires an NSX configured with an Overlay Transport Zone.
- When NSX-T is registered, both Overlay and VLAN segments can be used during the Network Profile creation.
- When NSX-T is registered, both Overlay and VLAN segments can be used during the Site Resources Profile creation.
- In cross-vCenter NSX environments, the HCX is connected to each vCenter Server and the Primary or Secondary NSX Manager, respectively. The HCX system connected to the Secondary NSX Manager must also connect to the primary NSX Manager in the provided field. This configuration is required for Universal Logical Switch extension.
- In NSX-T deployments, the HCX supports integration with networking objects created with the NSX Simplified UI/API only.
- Deployment of the HCX in environments with multiple Data Centers prepared for NSX and a single vCenter (sharing transport zones and datastores) is not supported.

Note The Network Extension service has additional NSX requirements. See Requirements for Network Extension.

NSX Requirements for the HCX Connector Installation

NSX is not required for HCX Connector installation, but requirements apply if you extend NSX overlay networks with HCX.

For more information, see NSX Requirements for HCX Deployments.

Installing HCX Manager

5

This section describes how to install and activate the HCX Manager component services: HCX Connector and HCX Cloud .

These components work together to provide the HCX services. In cloud-to-cloud environments, you deploy HCX Cloud at both the source and the destination sites. In legacy vSphere-to-cloud (private or public) deployments, you install HCX Connector at your on-premises or legacy site and HCX Cloud at the destination cloud site.

Read the following topics next:

HCX Manager Installation Workflow

This section provides an overview of the HCX Manager installation workflow for supported installation scenarios.

Deployment Types and General Workflow

The HCX deployment type varies depending on the environments being connected. Regardless of the deployment type, the overall installation workflow is the same.

HCX Deployment Types

HCX deployments use the following terminology:

Software Defined Data Center (SDDC)

Software Defined Data Center (SDDC) refers to an environment using current VMware software. An SDDC can refer to a private (on-premises) cloud or a public cloud that meets the requirements of the HCX destination. The SDDC is typically the destination for HCX migrations and network extension. See Software Version Requirements.

Legacy vSphere

A legacy environment uses vSphere Version 6.0 or higher and optionally uses NSX. These environments typically contain the workloads to migrate and the networks to extend.

Public Cloud

An SDDC that is offered as a service by HCX activated public cloud providers. VMware Cloud on AWS is a public cloud. A public cloud is typically the destination for migrations and network extension.

Note VMware Cloud on AWS uses the term "SDDC" to describe a compute instance in the cloud.

HCX is available for the following hyperscaler and VMware partner environments:

- Azure VMware Solution
- Google Cloud VMware Engine
- Oracle Cloud VMware Solution
- VMware Cloud on AWS
- VMware Cloud on Dell EMC
- VMware Cloud Provider Program

HCX deployments fall into several types:

Deployment	Description
Legacy vSphere to SDDC	In this deployment type, the HCX Connector at the Legacy site initiates Site Pairing, and the Service Mesh appliances initiate the Interconnect tunnels. The HCX Cloud manager and the Service Mesh appliances at the SDDC site are the receivers.
Legacy vSphere to Public Cloud	In this deployment type, the HCX Connector at the Legacy site initiates Site Pairing, and the Service Mesh appliances initiate the Interconnect tunnels. The HCX Cloud Manager and the Service Mesh appliances at the Public Cloud are the receivers.
Cloud-to-Cloud (Public Cloud to Public Cloud, SDDC to SDDC, or SDDC to Public Cloud)	In this deployment type, the HCX Manager at the SDDC or the Public Cloud can initiate or receive Site Pairing requests and act as the initiator or receiver during the Interconnect tunnel creation. **Note** Cloud-to-cloud deployment is not available for VMware Cloud Director.

HCX Installation Workflow

In any of these deployment types, HCX operation is functionally the same, and the same general workflow applies:

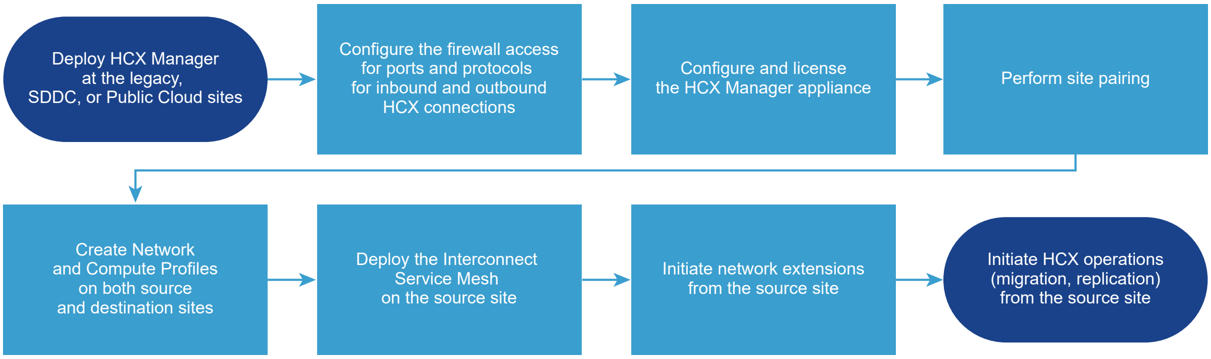
Deploy HCX Manager at the legacy, SDDC, or Public Cloud sites
Configure the firewall access for ports and protocols for inbound and outbound HCX connections
Configure and license the HCX Manager appliance
Perform site pairing
Create Network and Compute Profiles on both source and destination sites
Deploy the Interconnect Service Mesh on the source site
Initiate network extensions from the source site
Initiate HCX operations (migration, replication) from the source site

HCX Installation Workflow for HCX Public Clouds

A sample public installation workflow using HCX on the VMware Cloud on AWS.

Important Follow the HCX installation procedures provided by your public cloud service.

This section provides an example procedure demonstrating how to use HCX with the VMware Cloud on AWS. Not all these steps must be repeated for each source and destination site pair:

- Steps 1-3, 8, 9 must be performed for each SDDC.
- Steps 4-7 are only required once for each source site.

An HCX Connector installation can pair with many VMware Cloud on AWS SDDCs when the Network Profiles are configured to support them.

1. Prepare the deployment configurations using Checklist B in Getting Started with VMware HCX.
2. Activate HCX in the **Add Ons** tab of your VMware Cloud on AWS SDDC. See Deploying HCX from the VMware Cloud on AWS Console.
3. In the VMware Cloud on AWS SDDC Console go to the Network and Security tab to perform the following actions:

   a. Configure the Management Gateway to allow the HCX Cloud Manager (use the pre-defined HCX group as the destination) to receive inbound TCP-443 connections.

   b. If configuring HCX to use AWS Direct Connect with a Private Virtual Interface, see Configuring VMware HCX for Direct Connect Private Virtual Interfaces.

   For more details on the network port configuration, see Network Port and Protocol Requirements.
4. Download the HCX Connector for the Source on-premises installation and site pairing.

   a. On the **Add Ons** tab of your SDDC, click **Open HCX** on the HCX card.

   b Navigate to the SDDC tab and click **Open HCX**.

   - For VMware Cloud on AWS, the browser redirects to `hcx.sddc-*.vmwarevmc.com`.
   - For VMware Cloud on AWS GovCloud, the browser redirects to `https://hcx.sddc-*.vmwarevmcgov.com`.

   c Enter the cloudadmin@vmc.local user and password and click **Log In**.

   d Under the Administration tab, select **System Updates** and click **Request Download Link**.

   An option is provided to download the HCX Connector OVA locally or copy the download link.

5 Deploy HCX Manager in the source environment using the HCX Connector OVA. See Deploy the HCX Manager OVA in the vSphere Client.

Note For VMware Cloud on AWS GovCloud environments, be sure sure to select **Enable GovCloud Mode** when deploying the HCX Connector OVA in the vSphere Client.

6 Configure the source site (on-premises) firewall to allow inbound and outbound connectivity based on the HCX services or features used. See https://ports.vmware.com/home/VMware-HCX.

7 Create and activate keys for the source site HCX Connector that will be paired with the HCX Cloud Manager in VMware Cloud on AWS.

   a Log in to console.cloud.vmware.com.

   b On the **Add Ons** tab of your SDDC, click **Open HCX** on the HCX card.

   c Navigate to the **Activation Keys** tab.

   d Create an Activation Key for the source HCX Connector.

   e Enter the created activation key in the source HCX Connector and click **Activate**.

8 Pair HCX Connector with HCX Cloud. See Adding a Site Pair.

9 Activate the HCX services to deploy the HCX Interconnect.

   See Configuring and Managing the HCX Service Mesh.

HCX Installation Workflow for vSphere Private Clouds

This topic summarizes a fully private HCX installation, where both the destination/modernized environment and the source/legacy environments must be considered.

1 Prepare the deployment configurations using Checklist A in Getting Started with VMware HCX.

2 Download the HCX Installer for the destination site first. See Downloading the Installer OVA.

3 Deploy HCX Manager in the destination environment using the HCX Cloud OVA. See Deploy the HCX Manager OVA in the vSphere Client.

4 Configure firewall rules at source and destination sites to allow inbound and outbound connectivity based on the HCX services or features used. See https://ports.vmware.com/home/VMware-HCX.

5 Activate and Configure the HCX Cloud system. See Activating and Configuring HCX Cloud Manager

6 Configure the Compute Profile on the HCX Cloud system, see: Configuring and Managing the HCX Service Mesh.

Compute Profiles are defined in both HCX Connector and HCX Cloud systems. Later in the workflow, the **Multi-Site Service Mesh** wizard is used to deploy HCX Interconnect services.

7 Deploy the HCX Manager in the source environment using the HCX Connector OVA. See Deploy the HCX Manager OVA in the vSphere Client.

8 Activate and Configure the HCX Connector system. See Activate and Configure HCX Connector.

9 Pair HCX Connector with HCX Cloud. See Adding a Site Pair.

10 To deploy the HCX Interconnect, activate HCX services at the source site. See Configuring and Managing the HCX Service Mesh.

HCX Installation Workflow for VMware Cloud Director Private Clouds

This topic summarizes a fully private HCX installation, where both the destination/modernized environment and the source/legacy environments must be considered. In this workflow, the destination system is integrated with Cloud Director.

1 Prepare the deployment configurations using Checklist A in Getting Started with VMware HCX.

2 Deploy HCX Manager in the destination environment using the HCX Cloud OVA. See VMware Knowledge Base article 94012.

3 Configure firewall rules at source and destination sites to allow inbound and outbound connectivity based on the HCX services or features used. See https://ports.vmware.com/home/VMware-HCX.

4 Activate and Configure the HCX Cloud system. See Activating and Configuring HCX Cloud Manager.

During this step, select Cloud Director as the installation type and select the additional Cloud Director-specific details (for example, Public Access URL, AMQP).

5 Prepare the destination site's HCX Cloud system for Interconnect Deployments using the Multi-Site Service Mesh, see: Configuring and Managing the HCX Service Mesh. Define Compute and Network Profiles.

6 Deploy HCX Connector in the source environment using the HCX Connector OVA. See VMware Knowledge Base article 94012.

7 Activate and Configure the HCX Connector system. See Activating and Configuring HCX Cloud Manager.

8 Pair HCX Connector with HCX Cloud. See Adding a Site Pair. Note the Cloud Director-specific information when connecting a Cloud Director-based target site.

9 To deploy the HCX Interconnect, activate HCX services at the source site. See Configuring and Managing the HCX Service Mesh.

Downloading the HCX Manager OVAs

You use separate OVA files to deploy HCX Connector and HCX Cloud site managers.

The installer OVA provides the image necessary to deploy HCX Cloud Manager. You obtain the installer OVA from the VMware downloads site. You obtain the HCX Connector OVA file by accessing the HCX Cloud Manager service interface after you have fully deployed and activated HCX Cloud Manager.

Note For information regarding HCX for VMware Cloud Director deployments, see VMware Knowledge Base article 94012.

Downloading the Installer OVA

The installer OVA is used for deploying HCX Cloud Manager in a vSphere cloud environment.

In VMware Cloud on AWS environments, the cloud service provider deploys and configures HCX Cloud Manager.

Use this procedure to download the installer OVA.

Procedure

1 Navigate to `https://downloads.vmware.com`.

2 Search for **HCX**.

3 Select **VMware HCX**.

4 Click **Download Now**.

Results

This installer updates itself to the most current service updates.

What to do next

Deploy the downloaded installer OVA in the vCenter Server. See Deploy the HCX Manager OVA in the vSphere Client.

Deploy the downloaded installer OVA in the vCenter Server. See

Downloading the HCX Connector OVA

The HCX Connector OVA is used when deploying VMware HCX at the legacy site in legacy-to-vSphere cloud environments.

You obtain the HCX Connector OVA from the **System Updates** selection in the HCX Cloud Manager service UI.

Prerequisites

Before you can download the HCX Connector OVA, you must install, activate, and configure HCX Cloud Manager.

Procedure

1 Navigate to the HCX Cloud Manager service interface: **https://*hcxcloudmgr-ip-or-fqdn***.

2 Log in using HCX administrator credentials.

> **Note**
>
> - For HCX Cloud Manager registered with VMware Cloud Director, log in using HCX system administrator credentials.

3 Navigate to the **Administration** tab.

4 Navigate to **System Updates** using the left-side menu.

5 Select **Current Version**.

6 Click **Check for Updates**.

7 Click **Request Download Link**.

What to do next

Deploy the downloaded HCX Connector OVA in the vCenter Server. See Deploy the HCX Manager OVA in the vSphere Client.

Deploy the HCX Manager OVA in the vSphere Client

Use this procedure to deploy the downloaded HCX Manager OVA using a standard OVA template installation through the vSphere Client.

Use the HCX installer OVA to install the latest version of HCX Cloud Manager in the destination vSphere Client site for private cloud deployments.

To deploy the HCX Manager OVA with VMware Cloud on AWS, see Deploying HCX from the VMware Cloud on AWS Console.

Procedure

1 Log in to the vSphere Client.

2 Right-click any inventory object that is a valid parent object of a virtual machine (such as a data center, folder, cluster, resource pool, or host) and select **Deploy OVF Template**.

3 On the **Select an OVF template** page, enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive or a network share for a local file, and click **Next**.

4 On the **Select a name and folder** page, enter an unique virtual machine name and the inventory location. Click **Next**.

5 On the **Select a compute resource** page, select a compute resource location, and click **Next**.

6 On the **Review details** page, verify the OVA template details, and click **Next**.

7 On the **License agreements** page, read and accept the VMware End User License Agreement, and click **Next**.

8 On the **Select storage** page, select the virtual disk format, storage policy, storage name, and then click **Next**.

9 On the **Select networks** page, select the destination network, and click **Next**.

10 On the **Customize Template** page, set the appropriate deployment properties:

- **Passwords** - Configure the CLI admin password and the root user password.
- **Network Properties** - Enter the network properties for the default gateway.
- **Static Routes** - Add a static route for a destination subnet or host.
- **DNS**
  - **DNS Server List** - Enter the list of DNS servers for this virtual machine.
  - **Domain Search List** - Domains that you enter are searched in the order you list them, and the search stops when a valid name is found.

    Note For HCX 4.4.0 and later, add the .local subdomain in environments without multicast DNS correctly configured as in RFC-6762.
- **Services Configuration**
  - **NTP Server List** - Enter the list of NTP servers and ensure that the NTP server can be reached from the virtual machine. If the NTP time is out of sync, services fail to start.
- **Deployment** (Connector OVA installation only)
  - **Enable GovCloud Mode** - If you are deploying the HCX Connector OVA for VMware Cloud on AWS GovCloud installations, you must check this box.

Note During OVA deployment, use “space” or “comma” separation for multiple NTP/DNS/Domain Name servers.

11 Click **Next**.

12 Review the deployment settings and click **Finish**.

What to do next

Allow up to 5 minutes for initialization, then browse to the appliance management interface for the initial activation using `https://hcx-ip-or-fqdn:9443`.

Activating and Configuring HCX Cloud Manager

After you have deployed the installer OVA file, activate the HCX Cloud Manager and perform the initial configuration immediately when you next open the appliance management interface.

Note This procedure is for deploying HCX Cloud Manager in private clouds. For public cloud deployments, the HCX Cloud Manager deployment and activation is automated.

Prerequisites

- The HCX Installer OVA is deployed in the vSphere Client. Allow up to five minutes after the installer OVA deployment for the services to initialize.
- Configure firewall rules at source and destination sites to allow inbound and outbound connectivity based on the HCX services or features used. See https://ports.vmware.com/home/VMware-HCX.
- Obtain the activation key. See Activating or Licensing New HCX Systems.

Procedure

1 Browse to the appliance management interface and log in using the admin user credentials.

Browse to `https://hcx-ip-or-fqdn:9443`.

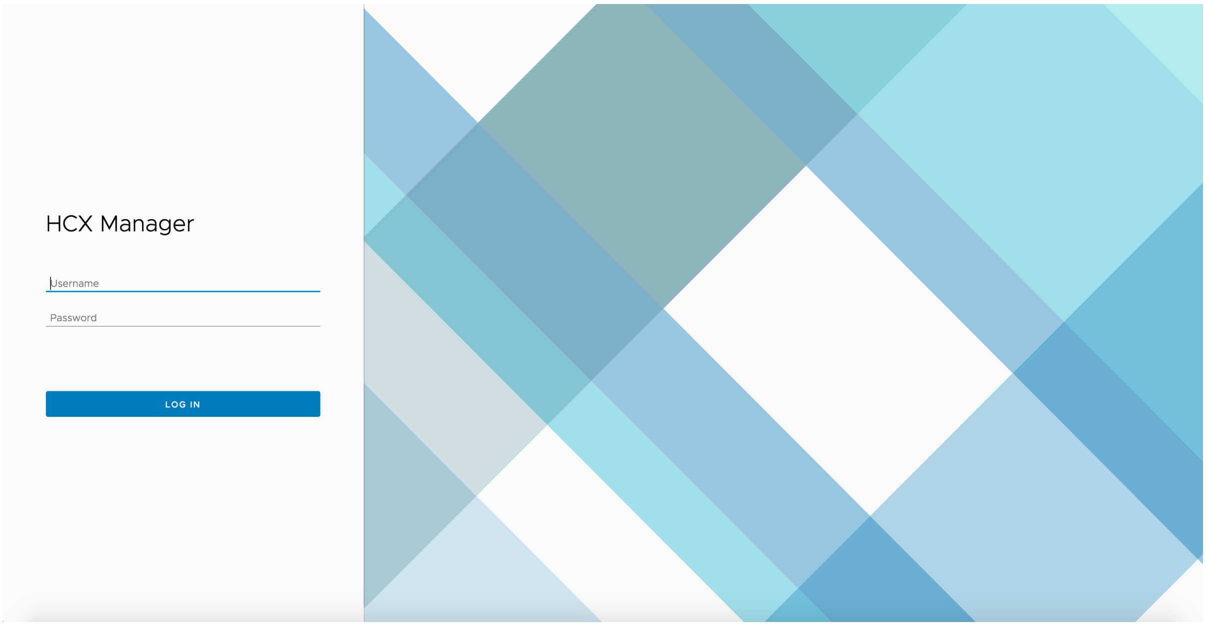

The **Configure your HCX System** screen appears, which includes a brief explanation of using the installer OVA to activate, configure, and set up the system.

2 Click **Continue**.

The activation screen appears.

3 Enter the **HCX Activation Server URL** and the **HCX License Key**.

The HCX Activation Server URL is preset to connect with VMware, but you must enter the HCX License Key. For information about obtaining a license key, see Activating or Licensing New HCX Systems.

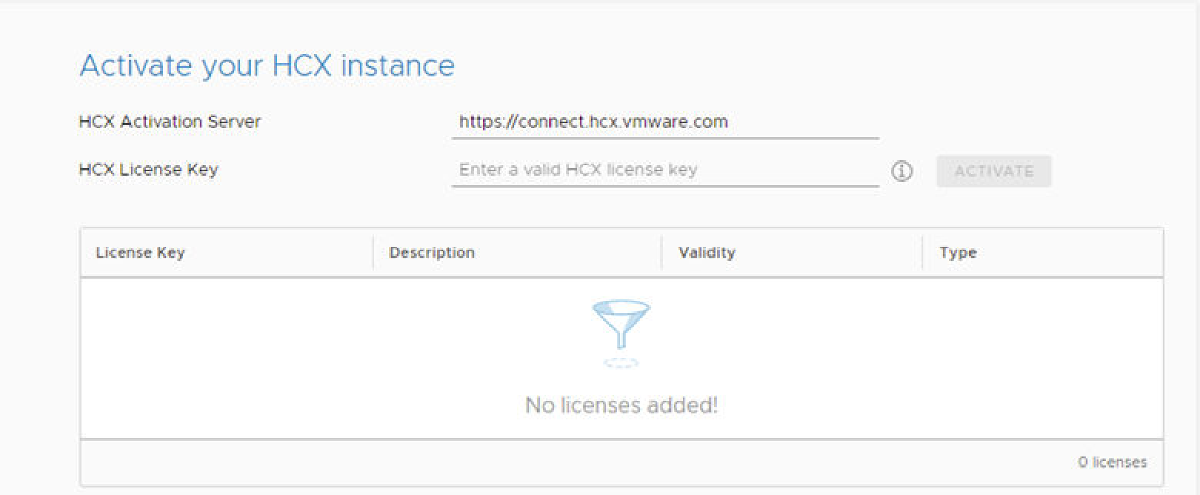

Activate your HCX instance

HCX Activation Server https://connect.hcx.vmware.com

HCX License Key Enter a valid HCX license key

License Key	Description	Validity	Type

No licenses added!

0 licenses

4 (Optional) If there is a proxy server in the environment in the path for outbound HTTPS connections, check **Configure Proxy**.

If a proxy server is entered, add the local vCenter, ESXi, NSX, SSO, and HCX-IX systems as exceptions not to be sent to the proxy server.

5 Click **Activate**.

The system prompts you to confirm the deployment type. The system detects the deployment type based on the license key and displays a graphic illustrating the installation component.

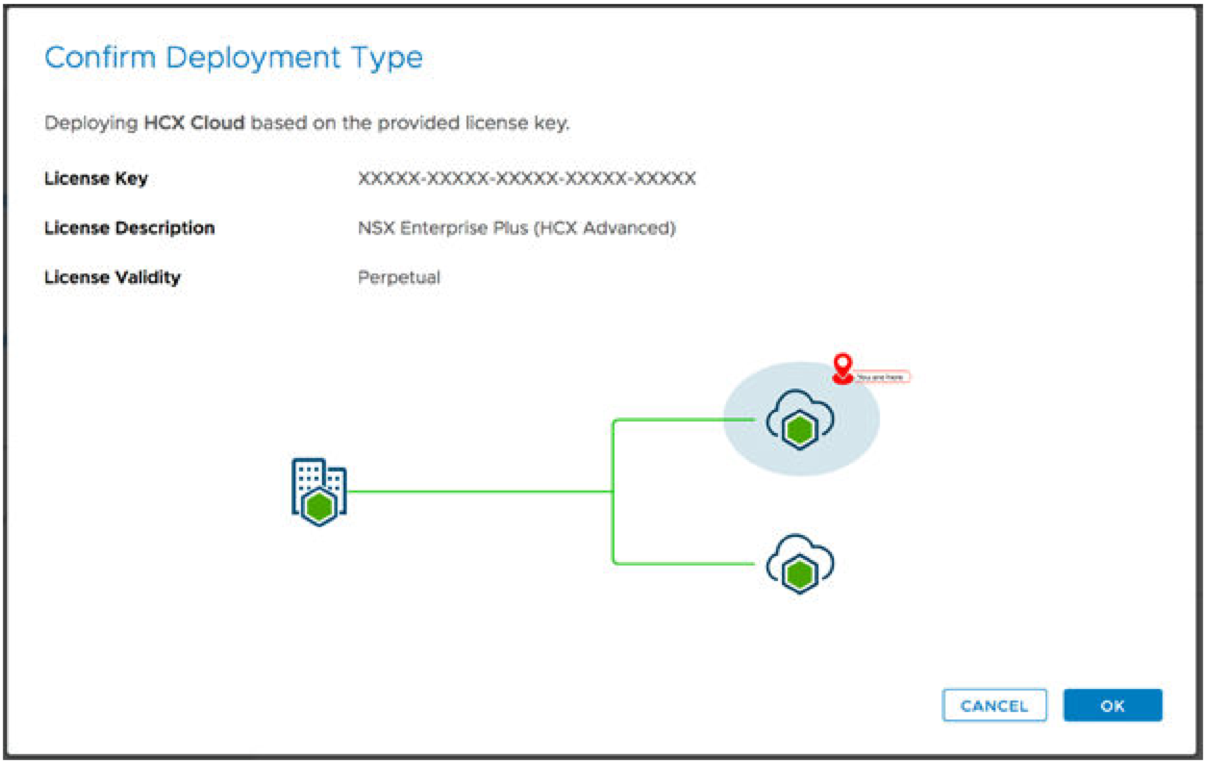

Confirm Deployment Type

Deploying **HCX Cloud** based on the provided license key.

License Key	XXXXX-XXXXX-XXXXX-XXXXX-XXXXX
License Description	NSX Enterprise Plus (HCX Advanced)
License Validity	Perpetual

6 Click **OK**.

The system begins downloading and upgrading your HCX instance.

The Manage License Keys screen appears.

7 (Optional) If you have an HCX Enterprise License (upgrade) key, enter it in the HCX License Key field, and click **Add**.

The upgrade license key is added to the license key table with the activation key. The table includes information about each license key and its duration.

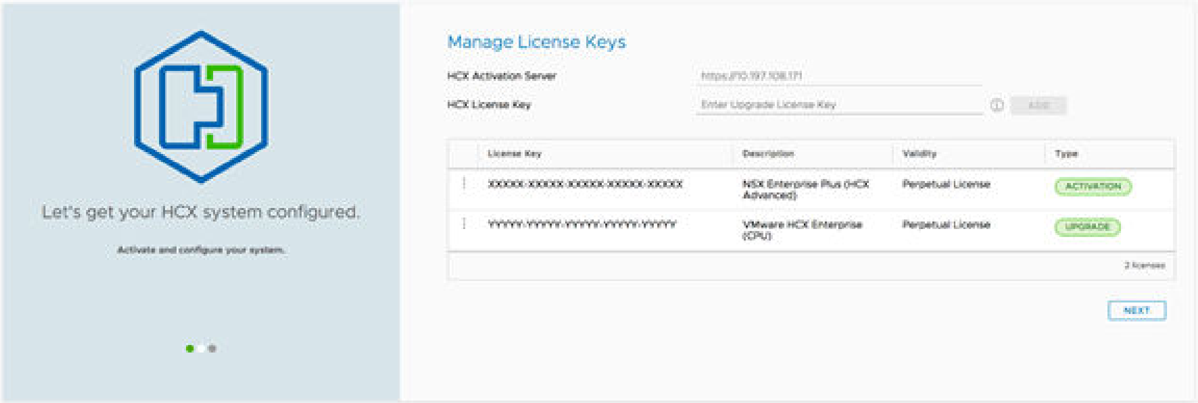

Let's get your HCX system configured.
Activate and configure your system.

Manage License Keys

HCX Activation Server: https://10.197.108.171

HCX License Key: Enter Upgrade License Key

License Key	Description	Validity	Type
XXXXX-XXXXX-XXXXX-XXXXX-XXXXX	NSX Enterprise Plus (HCX Advanced)	Perpetual License	ACTIVATION
YYYYY-YYYYY-YYYYY-YYYYY-YYYYY	VMware HCX Enterprise (CPU)	Perpetual License	UPGRADE

2 licenses

8 Click **Next**.

9 Observe the system download information.

After you enter the license information and confirm the deployment type, the system begins downloading the image file that is specific to the deployment type. If upgrades are available, they are applied before the download. The download process can take several minutes depending on your environment. A display screen provides the download status.

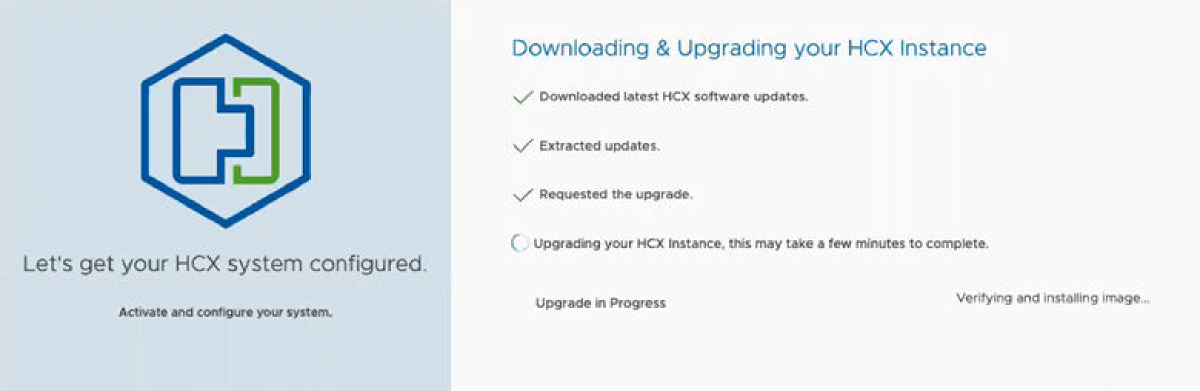

Let's get your HCX system configured.

Activate and configure your system.

Downloading & Upgrading your HCX Instance

✓ Downloaded latest HCX software updates.

✓ Extracted updates.

✓ Requested the upgrade.

○ Upgrading your HCX Instance, this may take a few minutes to complete.

Upgrade in Progress

Verifying and installing image...

When the download is complete, the system reloads, and the login screen appears.

10 To start the configuration wizard, log in to the system using the admin user credentials.

The system location screen appears.

11 Enter the location where you are deploying the system.

Select the nearest major city to where the HCX system is geographically located. HCX sites are represented visually in the Dashboard.

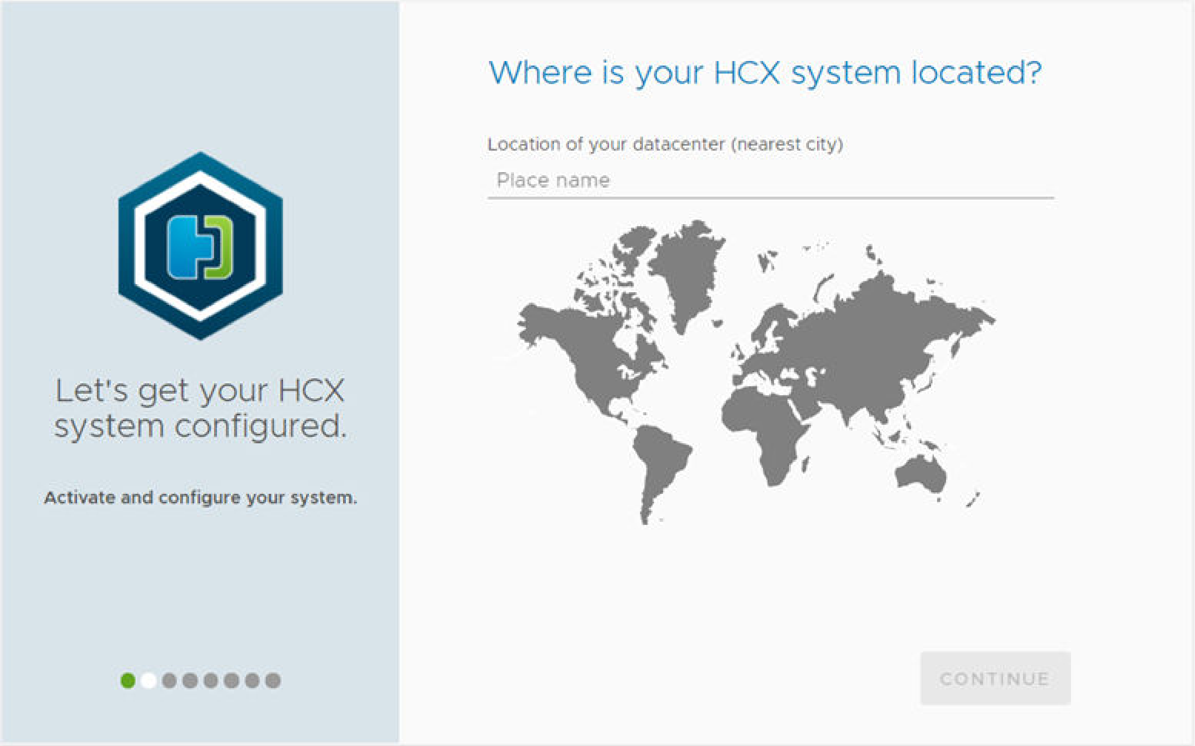

Let's get your HCX system configured.

Activate and configure your system.

Where is your HCX system located?

Location of your datacenter (nearest city)

Place name

12 Click **Continue**.

A screen appears prompting you for a system name.

13 Enter the system name, and click **Continue**.

A screen appears prompting you to select the cloud instance type.

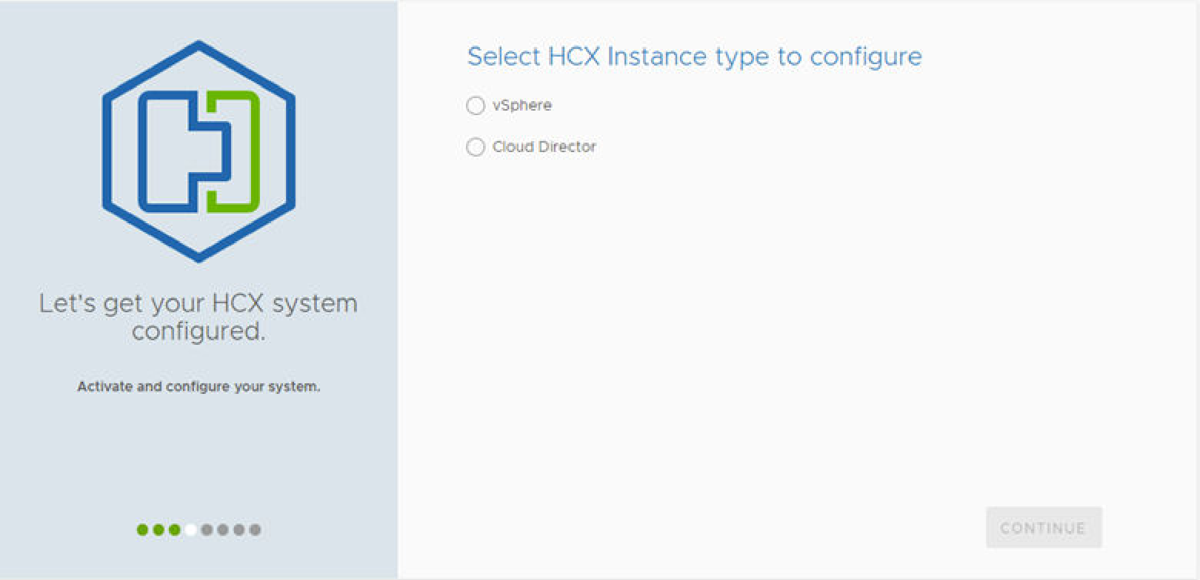

Let's get your HCX system configured.

Activate and configure your system.

Select HCX Instance type to configure

○ vSphere

○ Cloud Director

14 Select the cloud instance to which VMware HCX will be connected.

The HCX can connect to only one cloud instance per deployment.

Note Kubernetes is not available for VMware HCX.

15 Click **Continue**.

A series of screens appears, prompting you for the selection details.

16 Enter the configuration details for the selected cloud instance.

After entering the information, click **Continue** to proceed to the next screen

Cloud Instance	Configuration Parameters
vSphere	a vCenter Server and NSX details 1 vCenter Server ■ vCenter URL ■ User name ■ Password 2 NSX ■ NSX URL ■ User name ■ Password b SSO details ■ vCenter Server or Platform Services Controller URL c Public Access URL details ■ URL through which the HCX Manager is accessed. **Note** This is typically the HCX Manager services UI: https//*<hcx-mgr-fqdn-or-ip>*.
Cloud Director	a Cloud Director details ■ Cloud Director URL ■ System Administrator user name ■ System Administrator password b vCenter Server and NSX details **Note** The HCX Manager automatically fetches the vCenter Server and NSX URLs. 1 vCenter Server ■ User name ■ Password 2 NSX ■ User name ■ Password c AMQP details **Note** The HCX Manager automatically fetches the AMQP parameters. Edit the parameters as appropriate. ■ AMQP Host name ■ Port ■ vHost ■ User name ■ Password ■ Use SSL

The system verifies the configuration and then displays a configuration summary screen. The summary information lists the Location, System Name, vCenter Server, NSX Manager, SSO information, and Public Access URL. The summary includes instructions to restart the HCX Application Service and Web service for the changes to take effect, and to configure vSphere Roles after restarting the services.

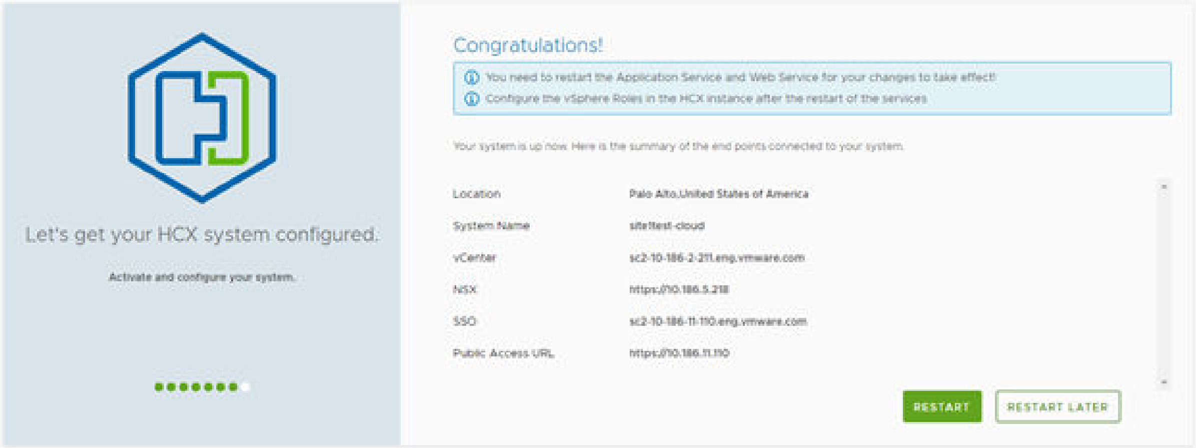
Let's get your HCX system configured.
Activate and configure your system.

Congratulations!

ⓘ You need to restart the Application Service and Web Service for your changes to take effect!

ⓘ Configure the vSphere Roles in the HCX instance after the restart of the services

Your system is up now. Here is the summary of the end points connected to your system.

Location	Palo Alto,United States of America
System Name	site1test-cloud
vCenter	sc2-10-186-2-211.eng.vmware.com
NSX	https://10.186.5.218
SSO	sc2-10-186-11-110.eng.vmware.com
Public Access URL	https://10.186.11.110

17 To reload the system, click **Restart**.

It can take several minutes to reinitialize the system completely. During this process, the appliance management interface is not available.

After the system reloads, it displays the Appliance Management Dashboard. For more information about the dashboard, see Understanding the Appliance Management Dashboard.

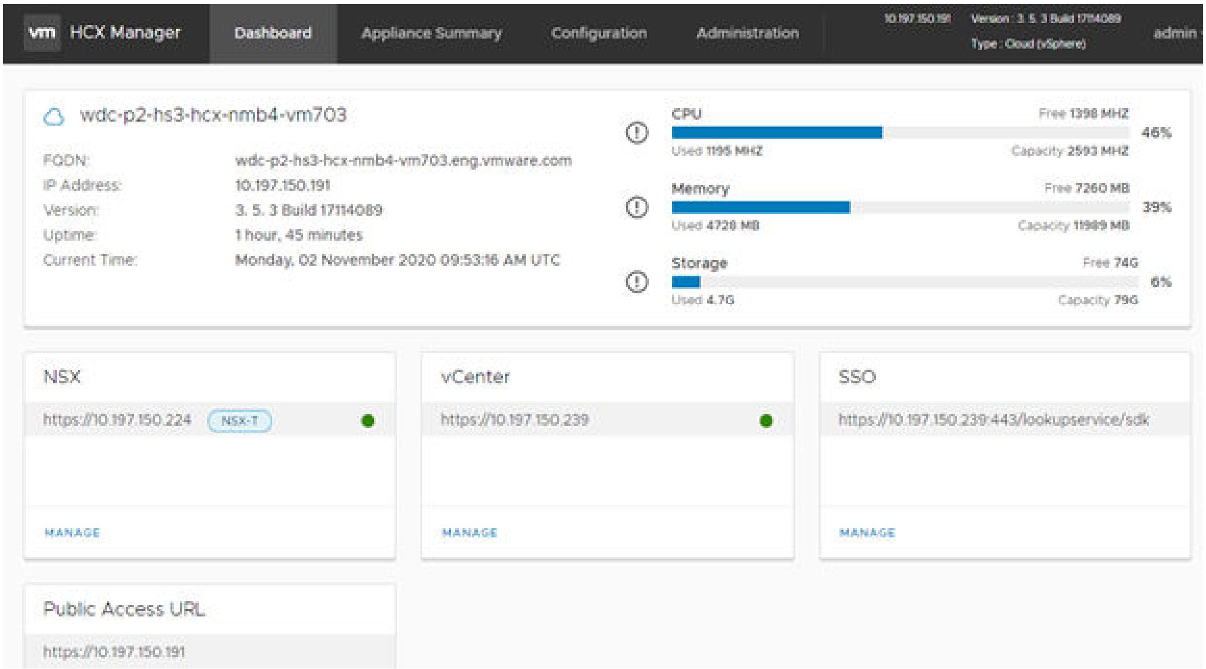

wdc-p2-hs3-hcx-nmb4-vm703

FQDN:	wdc-p2-hs3-hcx-nmb4-vm703.eng.vmware.com
IP Address:	10.197.150.191
Version:	3. 5. 3 Build 17114089
Uptime:	1 hour, 45 minutes
Current Time:	Monday, 02 November 2020 09:53:16 AM UTC

	Free	Used	Capacity	%
CPU	Free 1398 MHZ	Used 1195 MHZ	Capacity 2593 MHZ	46%
Memory	Free 7260 MB	Used 4728 MB	Capacity 11989 MB	39%
Storage	Free 74G	Used 4.7G	Capacity 79G	6%

NSX
https://10.197.150.224 NSX-T ●
MANAGE

vCenter
https://10.197.150.239 ●
MANAGE

SSO
https://10.197.150.239:443/lookupservice/sdk
MANAGE

Public Access URL
https://10.197.150.191

18 Configure HCX roles:

   a In the appliance management dashboard, navigate to **Configuration > HCX Role Mapping**.

   b Assign the HCX Roles to the vCenter User Groups that are allowed to perform HCX operations.

     By default, the HCX Administrator role is mapped to the local vSphere administrator group. For HCX Cloud Manager deployments, the system displays the optional HCX Tenant role, which is intended for use by HCX Service Providers.

   c Click **Save**.

Results

The HCX Cloud Manager system configuration is complete.

What to do next

Deploy any additional HCX systems, and then go to "Activating and Configuring HCX Connector."

Activate and Configure HCX Connector

This section describes how to activate and configure HCX Connector for private cloud (on-premises) sites.

Prerequisites

- HCX Cloud Manager is activated and configured. See Activating and Configuring HCX Cloud Manager.
- You downloaded HCX Connector OVA from the Cloud Manager Interface and deployed it in the on-premises vSphere Client. See "Downloading the HCX Connector OVA."
- Obtain the HCX Connector activation key. See Activating or Licensing New HCX Systems.

Procedure

1 Browse to the appliance management interface and log in using the admin user credentials.

  Browse to *https://hcx-ip-or-fqdn*:9443.

  The activation screen appears.

2 If you are activating HCX for VMware Cloud on AWS GovCloud, accept the monitoring notice by clicking **I Agree**.

  The monitoring notice appears each time you log in to a VMware secure environment.

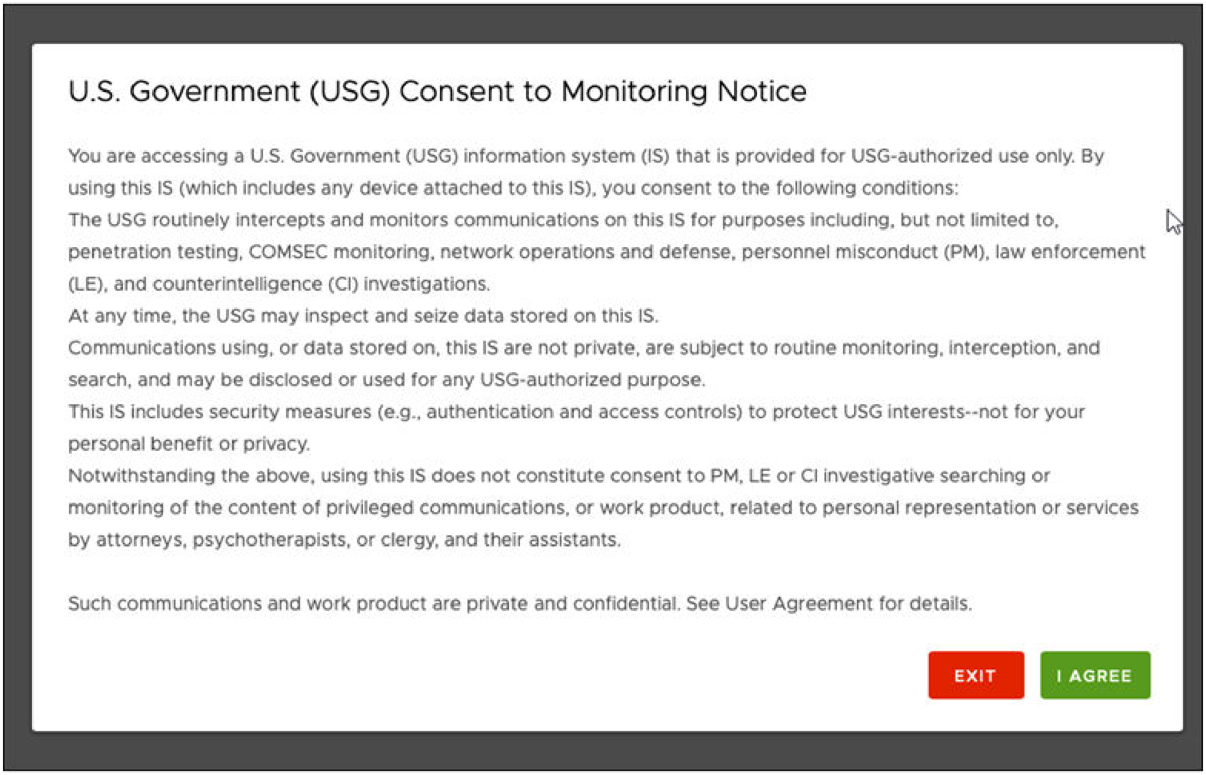

U.S. Government (USG) Consent to Monitoring Notice

You are accessing a U.S. Government (USG) information system (IS) that is provided for USG-authorized use only. By using this IS (which includes any device attached to this IS), you consent to the following conditions:
The USG routinely intercepts and monitors communications on this IS for purposes including, but not limited to, penetration testing, COMSEC monitoring, network operations and defense, personnel misconduct (PM), law enforcement (LE), and counterintelligence (CI) investigations.
At any time, the USG may inspect and seize data stored on this IS.
Communications using, or data stored on, this IS are not private, are subject to routine monitoring, interception, and search, and may be disclosed or used for any USG-authorized purpose.
This IS includes security measures (e.g., authentication and access controls) to protect USG interests--not for your personal benefit or privacy.
Notwithstanding the above, using this IS does not constitute consent to PM, LE or CI investigative searching or monitoring of the content of privileged communications, or work product, related to personal representation or services by attorneys, psychotherapists, or clergy, and their assistants.

Such communications and work product are private and confidential. See User Agreement for details.

3 Enter the **HCX Activation Server URL** and the **HCX License Key**.

The HCX Activation Server URL is automatically set to connect with VMware, but you must enter the HCX License Key.

Note If you are activating HCX for VMware Cloud on AWS GovCloud enter the **Secure Cloud Activation Proxy URL** and **HCX License Key**. The Secure Cloud Activation Proxy URL is the HCX Manager URL in the GovCloud SDDC: `https//`*`hcx.sddc-*.vmwarevmcgov.com`*.

4 Click **Activate**.

The system location screen appears.

5 Enter the location where you are deploying the system.

Select the nearest major city to where the system is geographically located. HCX sites are represented visually in the Dashboard.

6 Click **Continue**.

A screen appears prompting you for a system name.

7 Enter the system name, and click **Continue**.

A screen appears indicating that you have successfully activated HCX with a prompt to continue setting up HCX Manager.

8 Click **Yes, Continue.**

A screen appears prompting you to enter the information for connecting HCX Manager with the vCenter Server.

9 Enter the vCenter Server information:

- vCenter Server IP address or FQDN: *https://vc_ip/fqdn*
- Username
- Password

10 (Optional) If your on-premises site uses NSX, at the bottom of the screen check the box **Connect your NSX Manager**.

A panel appears for entering the NSX Manager information.

11 Enter the NSX Manager information:

- NSX Manager IP address or FQDN: https://*nsx_ip/fqdn*
- Username
- Password

12 Click **Continue**.

A screen appears prompting you to configure the SSO/PSC information.

13 Enter the Identify Sources information: *https://sso_ip/fqdn*.

The system verifies the configuration and displays a system configuration summary screen. The summary information lists the Location, System Name, vCenter Server, NSX Manager, and SSO information. The summary includes instructions to restart the Application Service and Web Service for the changes to take effect, and to configure vSphere roles after restarting the services.

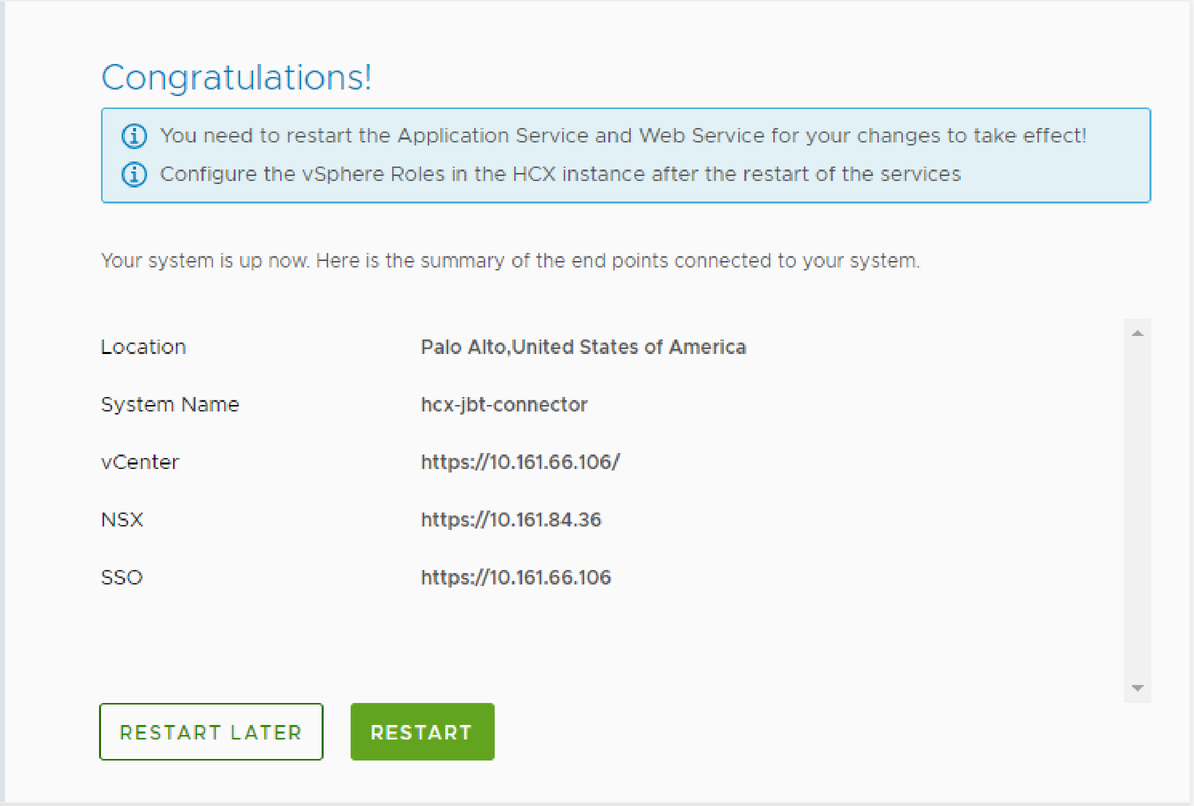

Congratulations!

ⓘ You need to restart the Application Service and Web Service for your changes to take effect!

ⓘ Configure the vSphere Roles in the HCX instance after the restart of the services

Your system is up now. Here is the summary of the end points connected to your system.

Location	Palo Alto,United States of America
System Name	hcx-jbt-connector
vCenter	https://10.161.66.106/
NSX	https://10.161.84.36
SSO	https://10.161.66.106

14 To reload the system, click **Restart** at the bottom of the screen.

Reinitializing the system can take several minutes. During this process, the appliance management interface is not available.

After the system reloads, it displays the appliance management dashboard. For more information about the dashboard, see Understanding the Appliance Management Dashboard.

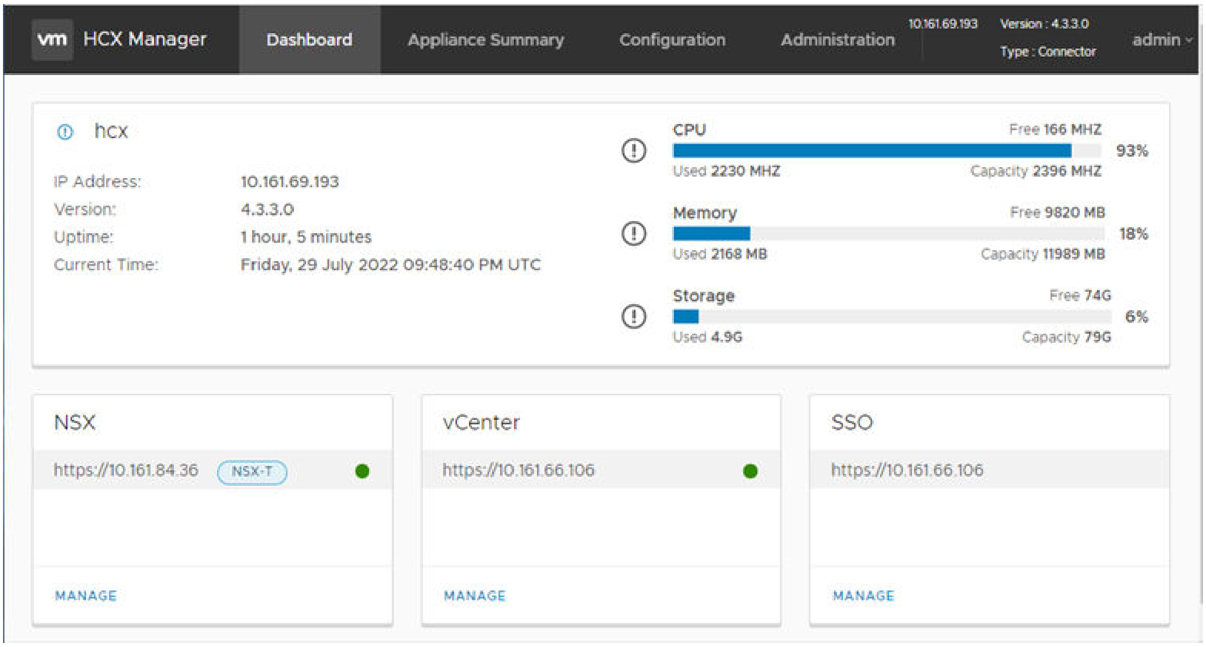

hcx

IP Address: 10.161.69.193
Version: 4.3.3.0
Uptime: 1 hour, 5 minutes
Current Time: Friday, 29 July 2022 09:48:40 PM UTC

CPU Free 166 MHZ 93% Used 2230 MHZ Capacity 2396 MHZ

Memory Free 9820 MB 18% Used 2168 MB Capacity 11989 MB

Storage Free 74G 6% Used 4.9G Capacity 79G

NSX https://10.161.84.36 NSX-T ●

vCenter https://10.161.66.106 ●

SSO https://10.161.66.106

15 Configure roles:

   a In the Appliance Management Dashboard, navigate to **Configuration > HCX Role Mapping**.

   b Assign roles to the vCenter User Groups that are allowed to perform HCX operations.

     By default, HCX Manager maps the Administrator role to the local vSphere administrator group. For HCX Cloud deployments, the system displays the optional Tenant role, which is intended for use by Service Providers.

   c Click **Save**.

Results

The system-level activation and configuration is complete.

What to do next

Deploy any additional sites, and then go to "Configuring and Managing the HCX Interconnect."

Configuring and Managing the HCX Interconnect

6

The VMware HCX Interconnect provides a secure pipeline for migration, extension, and Virtual Machine protection between two connected VMware HCX sites.

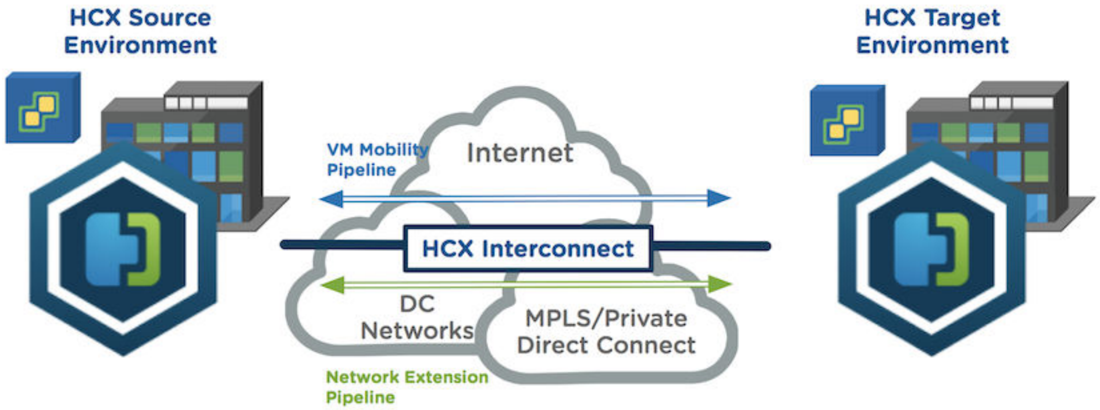

Read the following topics next:

- Overview of the HCX Interconnect Services Deployment with Multi-Site Service Mesh
- Configuring and Managing the HCX Service Mesh
- Sentinel Management

Overview of the HCX Interconnect Services Deployment with Multi-Site Service Mesh

HCX services are deployed and managed using the Multi-Site Service Mesh.

The following steps use the Interconnect Multi-Site Service Mesh interface. Configuration preparation steps are symmetrical.

1 Site Pairing: Register the destination HCX system at the source.

2 Create a Compute Profile in the source and destination HCX environments.

3 Add the Service Mesh at the source:

   a Select a source and destination Compute Profile.

   b Activate a Multi-Site Service Mesh.

Configuring and Managing the HCX Service Mesh

The Multi-Site Service mesh is used to create a secure optimized transport fabric between any two sites managed by HCX.

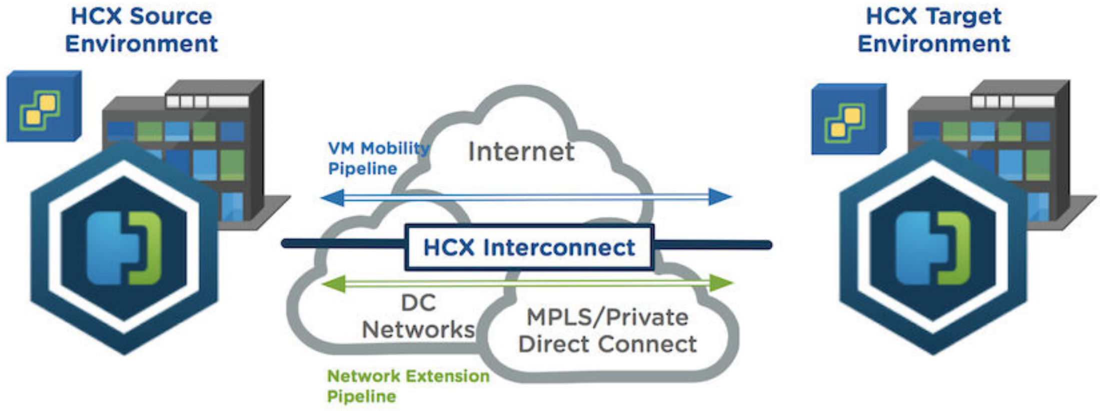

About the HCX Multi-Site Service Mesh

When HCX Migration, Disaster recovery, Network Extension, and WAN Optimization services are activated, HCX deploys Virtual Appliances in the source site and corresponding "peer" virtual appliances on the destination site. The Multi-Site Service Mesh activates the configuration, deployment, and serviceability of these Interconnect virtual appliance pairs.

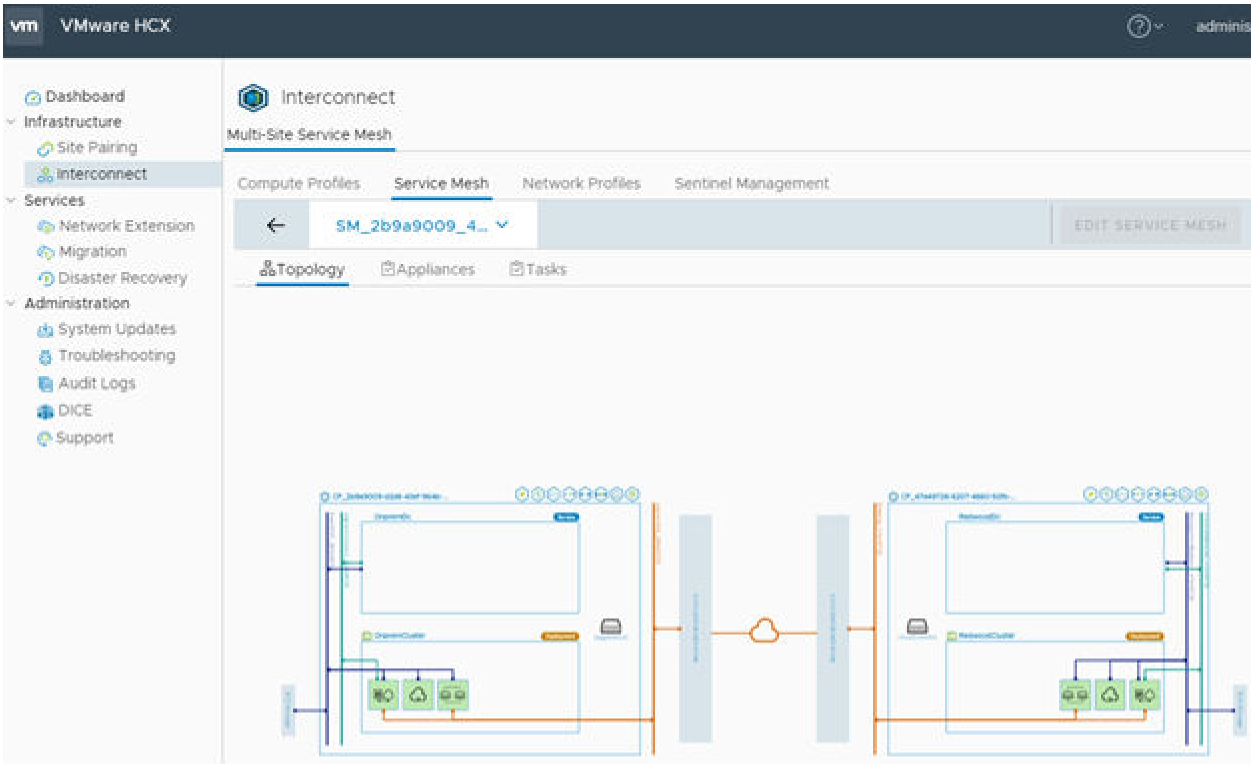

Multi-Site Service Mesh Benefits

New Configuration Options:

- Uniformity: the same configuration patterns at the source and remote sites.
- Reusability: Once a compute profile is created, it can be used to connect to multiple HCX sites.
- Multi-site ready: Compute Profiles and Network Profiles can be shared across multiple sites.
- Ease of reconfiguration: New capability to pool datastores or modify them after deploying an Interconnect network structure.
- Scale-out deployment: The HCX-IX can be deployed per cluster or a single HCX-IX can be shared across multiple clusters.

Performance Enhancements:

- Parallel execution ensures faster Interconnect deployments (in under 5 minutes).
- The new lockless model ensures parallel configuration of network stretches.

  Note When extending or unextending multiple networks that are part of the same Network Extension High Availability (HA) group, only one network extension operation runs at a time. Other network extension requests are queued until the current operation finishes.

Usability Enhancements:

- Improved interfaces display a clear deployment diagram.
- New task-tracking features provide incremental details for each step of the progress of operations.
- Preview of required firewall rules to avoid configuration difficulties.

Multi-Site Service Mesh Site Pairs

You register the destination HCX system in the Site Pairing Interface at the source site. Pairing the source and destination sites is a requirement for creating a Service Mesh.

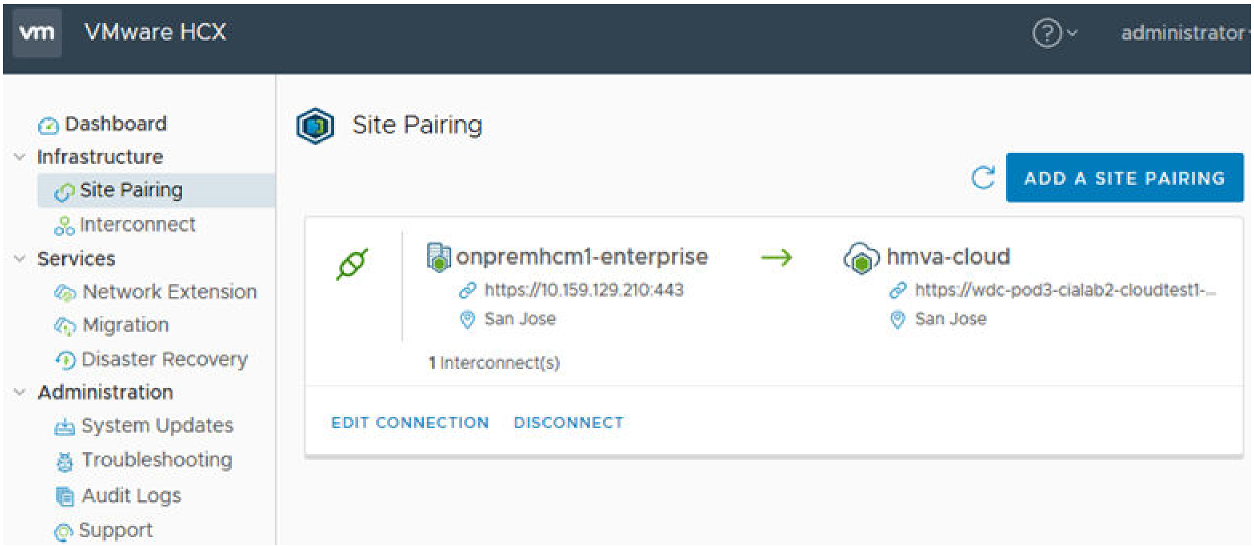

onpremhcm1-enterprise → hmva-cloud

https://10.159.129.210:443 — https://wdc-pod3-cialab2-cloudtest1-...

San Jose — San Jose

1 Interconnect(s)

Compute and Network Profiles

The compute profile defines the structure and operational details for the virtual appliances used in a Multi-Site Service Mesh deployment architecture. The compute profile:

- Provisions the infrastructure at the source and the destination site.
- Provides the placement details (Resource Pool, Datastore) where the system places the virtual appliances.
- Defines the networks to which the virtual appliances connect.

The following conditions apply when deploying a service mesh network:

- The integrated compute profile creation wizard can be used to create the compute and network profiles (or Network Profiles can be pre-created).
- HCX Interconnect service appliances are not deployed until a service mesh is created.

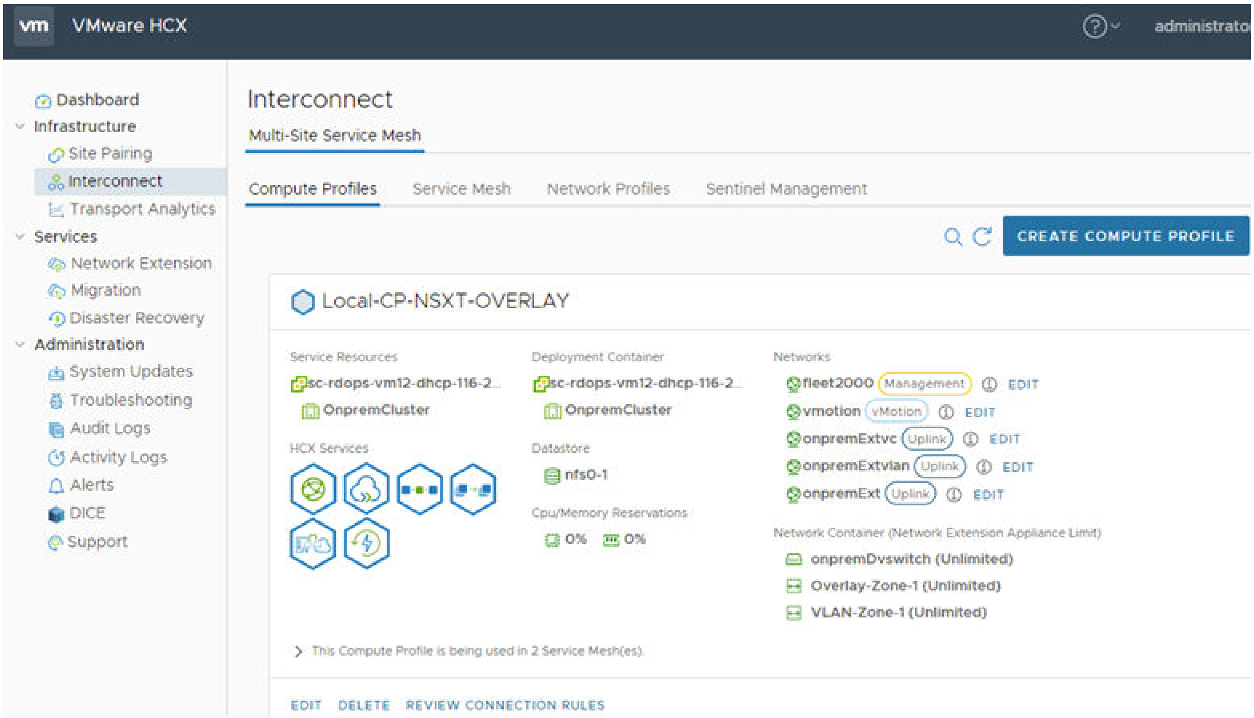

Local-CP-NSXT-OVERLAY

Service Resources
sc-rdops-vm12-dhcp-116-2...
OnpremCluster

Deployment Container
sc-rdops-vm12-dhcp-116-2...
OnpremCluster

Networks
fleet2000 Management EDIT
vmotion vMotion EDIT
onpremExtvc Uplink EDIT
onpremExtvlan Uplink EDIT
onpremExt Uplink EDIT

HCX Services

Datastore
nfs0-1

Cpu/Memory Reservations
0% 0%

Network Container (Network Extension Appliance Limit)
onpremDvswitch (Unlimited)
Overlay-Zone-1 (Unlimited)
VLAN-Zone-1 (Unlimited)

This Compute Profile is being used in 2 Service Mesh(es).

Service Mesh

A **Service Mesh** specifies a local and remote Compute Profile pair. When a **Service Mesh** is created, the HCX Service appliances are deployed on both the source and destination sites and automatically configured by HCX to create the secure optimized transport fabric.

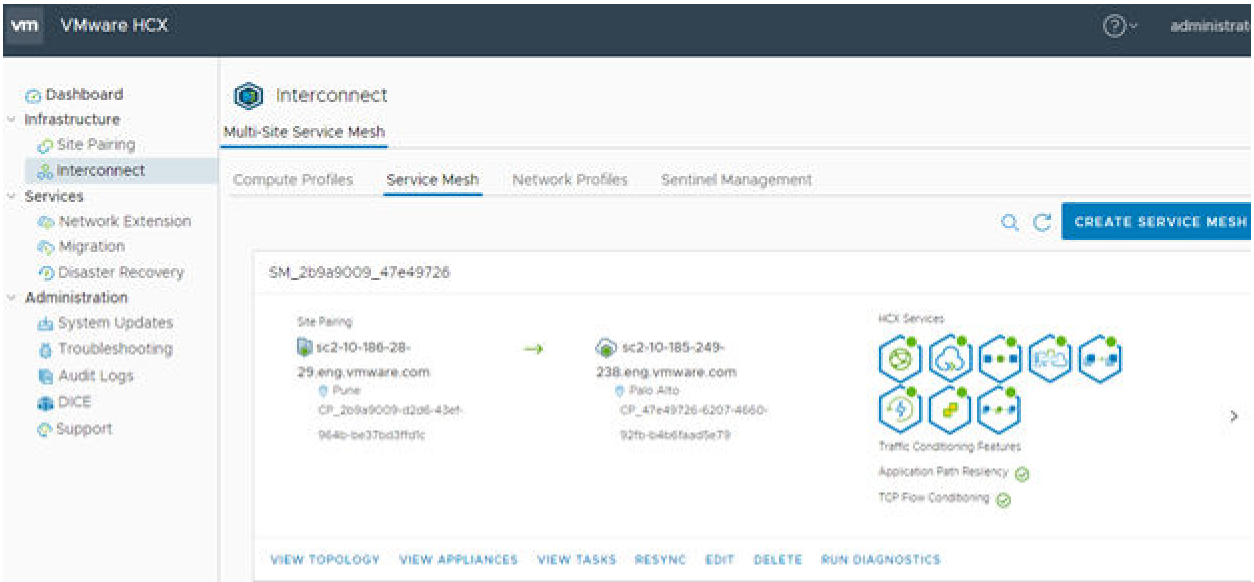

SM_2b9a9009_47e49726

Site Pairing
sc2-10-186-28-29.eng.vmware.com → sc2-10-185-249-238.eng.vmware.com
Pune
CP_2b9a9009-d2d6-43ef-964b-be37bd3ffd1c
Palo Alto
CP_47e49726-6207-4660-92fb-b4b6faad5e79

HCX Services

Traffic Conditioning Features
Application Path Resiliency ✓
TCP Flow Conditioning ✓

Sentinel Management

You must install HCX Sentinel on all guest virtual machines requiring migration using HCX OS Assisted Migration. Sentinel gathers the system configuration from the guest VM and assists with the data replication.

The source system information is used by various HCX OS Assisted Migration service processes. In part, the information is used to create an inventory of guest VM systems for migration and to help replication processes prepare the disks on the replica VM for replication and migration.

Sentinel also helps with the data replication by reading data written to the source disks and passing that data to the SDR appliance at the destination site.

Adding a Site Pair

A Site Pair establishes the connection needed for management, authentication, and orchestration of HCX services across a source and destination environment.

In HCX Connector to HCX Cloud deployments, the HCX Connector is deployed at the legacy or source vSphere environment. The HCX Connector creates a unidirectional site pairing to an HCX Cloud system. In this type of site pairing, all HCX Service Mesh connections, Migration and Network Extension operations, including reverse migrations, are always initiated from the HCX Connector at the source.

In HCX cloud-to-cloud deployments, site pairing can be unidirectional or bidirectional:

- In unidirectional site pairing, the HCX Cloud containing the virtual machine inventory and networks (similar to HCX Connectors) will site pair to the destination HCX Cloud. In this type of site pairing, all HCX Service Mesh connections, Migration and Network Extension operations, including reverse migrations, are always initiated from the source HCX Cloud system. In this case, an administrator might see the message URL not available when viewing site pairing from the destination site. This is expected behavior because HCX Connector to HCX Cloud site pairing is unidirectional.
- In bidirectional site pairing, the HCX Cloud systems are site paired with each other, share a common Service Mesh, and can initiate Migration and Network Extension operations from either HCX Cloud system.

In the case of unidirectional site pairing, an administrator may see the message `URL not available` when viewing site pairing from the destination site. This is expected behavior because HCX Connector to HCX Cloud site pairing is unidirectional. In the case of bidirectional site pairing, the URLs for the paired sites are visible from either the source or destination.

An HCX Connector cannot be the target for a site pairing.

Prerequisites

- HCX Manager installed and configured in the source and destination environments.
- The Site URL and User:
  - When the destination is a private vSphere-based private cloud, the Site URL refers to the HCX Cloud Manager at the target site:

    ```
    https://hcx-cloud-ip-or-fqdn
    ```

    Provide a user from the destination site's SSO configuration. The user must be included in the HCX Role-Mapping Group configuration.

The `administrator@vsphere.local` user is included by default.

  - When the destination system is a Public Cloud, the Site URL can use a trusted domain name pre-created by the cloud provider or an IP address.

    Use the credentials for a user that holds a cloud admin role in your private cloud.

  - When registering a Cloud Director Organization as the HCX destination endpoint, the Site URL refers to the HCX Cloud system with a suffix referencing the Org:

    ```
    https://hcx-cloud-ip-or-fqdn/cloud/org/<orgname>
    ```

    Provide a Local or LDAP Organization User with the Organization Administrator role. Use the format `username@orgname`.

- The destination Site URL must use a CA signed trusted certificate or be manually trusted on the source HCX system. See Chapter 11 Managing CA and Self-Signed Certificates.

Procedure

1. From the HCX dashboard, go to **Infrastructure** > **Site Pairs**.
2. Click **Add a Site Pairing**.
3. Enter the Remote HCX URL and credentials, then click **Connect**.
4. (Optional) To achieve bidirectional site pairing in cloud-to-cloud deployments, repeat this procedure at both cloud sites.

Results

If all validations succeed, the system displays the remote site in the list as a connected site. With bidirectional site pairing, both sites show up in the list.

Example: Connected Site

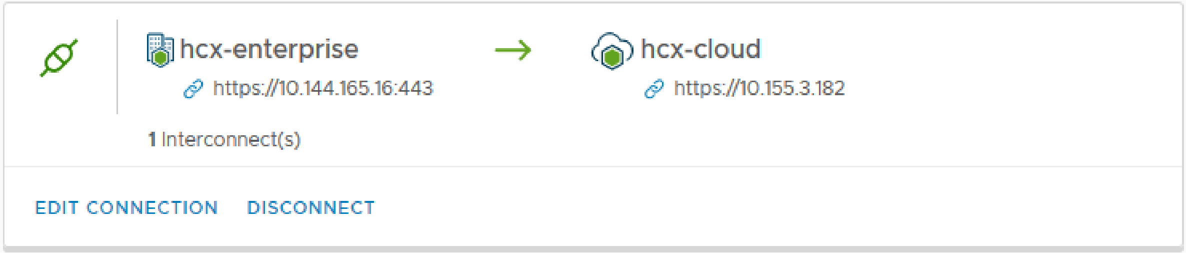

hcx-enterprise → hcx-cloud

https://10.144.165.16:443

https://10.155.3.182

1 Interconnect(s)

What to do next

Create the Network and Compute Profiles, followed by the Multi-Site Service Mesh.

Creating a Network Profile

The Network Profile is an abstraction of a Distributed Port group, Standard Port group, or NSX Logical Switch, and the Layer 3 properties of that network. A **Network Profile** is a sub-component of a complete **Compute Profile**.

Create a **Network Profile** for each network you intend to use with the HCX services. The extension selects these network profiles when creating a **Compute Profile** and assigned one or more of Network Profile functions.

Note Although a Network Profile can be assigned any of the functions during the Compute Profile configuration, consider creating a separate profile for each function as a best practice.

Note In Federated NSX environments, the NSX Global Manager populates local NSX managers with all global network segments known to the Global Manager. These networks are flagged in HCX as Global Networks in the HCX inventory, and become available for use for Bulk, RAV, and vMotion migrations. These global segments, however, are not supported for HCX Interconnect configuration, meaning Network Profile and Compute Profile creation with Global Segments or Global transport Zones.

- Management Network:

  The HCX Interconnect appliances use this network to communicate with management systems like the HCX Manager, vCenter Server, ESXi Management, NSX Manager, DNS, NTP.

- Uplink Network:

  The HCX Interconnect appliances use this network for WAN communications, like TX/RX of transport packets.

- vMotion Network:

  The HCX Interconnect appliances use this network for the traffic exclusive to vMotion protocol operations.

  Important The HCX Interconnect uses a Network Profile configuration dedicated to vMotion traffic. This configuration does not include the vMotion NFC traffic. HCX always uses its Management interface for vMotion NFC traffic.

- vSphere Replication Network:

  The HCX Interconnect appliances use this network for the traffic exclusive to vSphere Replication.

  Important In deployments where ESXi servers use a dedicated VMkernel configuration for vSphere Replication services, the HCX Interconnect uses a Network Profile configuration dedicated to the vSphere Replication traffic. This configuration does not include the vSphere Replication NFC traffic. HCX always uses its Management interface for vSphere Replication NFC traffic.

- Guest Network for OS Assisted Migration

The Sentinel Gateway appliances use this vSphere network to connect with non-vSphere virtual machines.

Important When creating a separate Network Profile for vMotion or vSphere Replication services, although the option is available to configure a GW as a standard Network Profile, traffic for those services will only use the default GW in the Management Network Profile to attempt to access resources in a different subnet. If ESXi resources are not L2 adjacent to the IX appliance on those networks, there is a requirement to configure “Static Routes” as part of the “Advance Configurations” option in the Compute Profile to ensure traffic is directed to the default GW on those networks.

Prerequisites

- The HCX Manager must be installed and configured.
- Use the planned network configurations prepared using the checklist described in Getting Started with VMware HCX.

Procedure

1 Navigate to the **Network Profiles** interface:

   a In the vSphere Client, navigate To **HCX** > **Interconnect** > **Multi-Site Service Mesh** > **Network Profiles**.

   b At the destination site, navigate to https://hcx-cloud-ip-or-fqdn > **Multi-Site Service Mesh** > **Network Profiles**.

2 Click **Create Network Profile**.

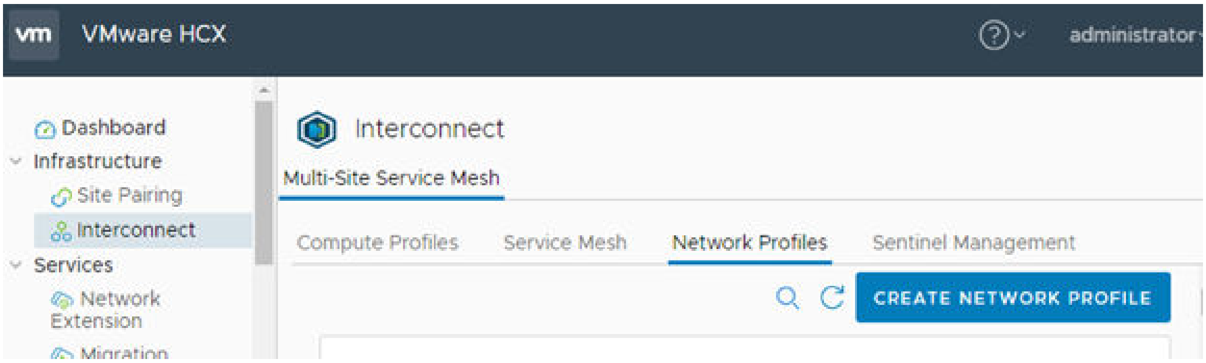

3 Select a vCenter Server and existing Network.

   a Select a vCenter Server from the drop-down menu.

   b Select Distributed Port Group, Standard Switch Port Group, NSX Logical Switch, or External Network to filter the available networks by type.

   > **Note** Additional options to create a Network Profile backed by a VMware Cloud Director External Network display only in VMware Cloud Director deployments.

   c Select one of the available networks.

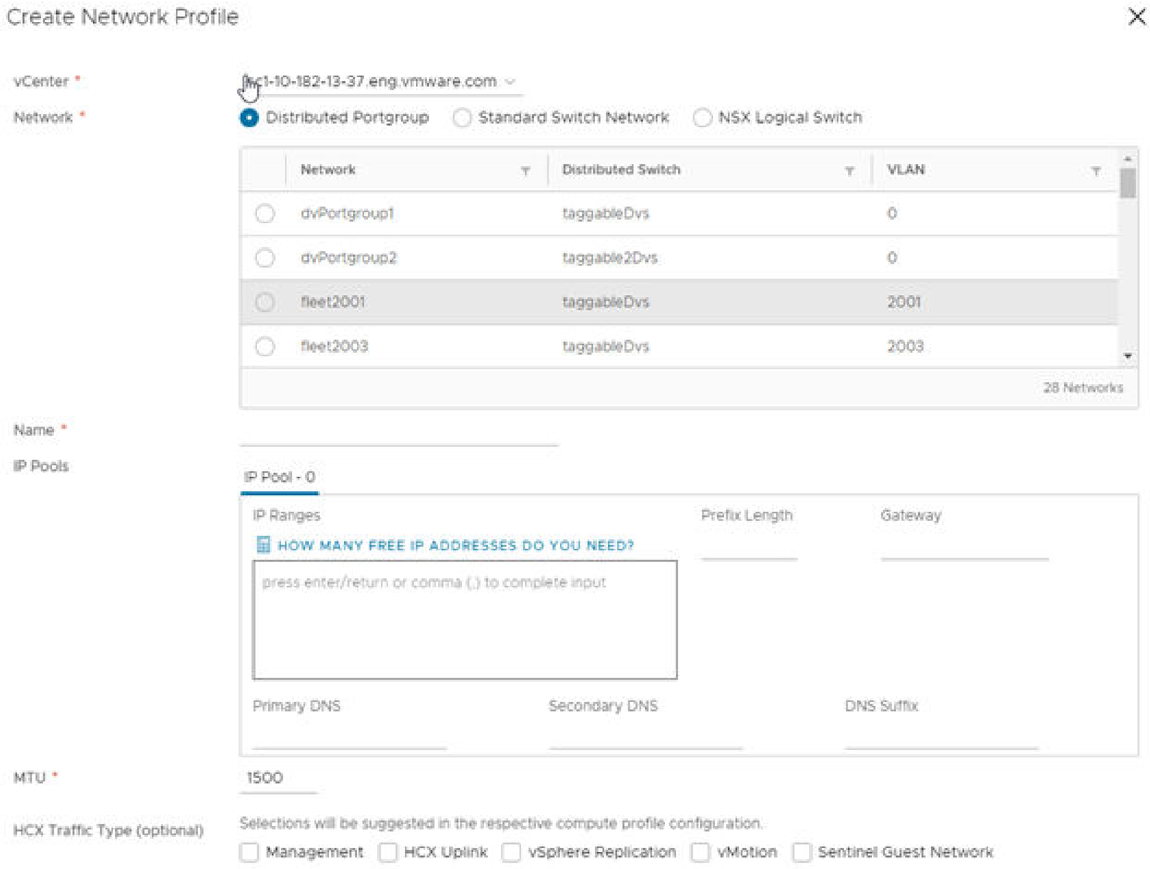

Create Network Profile

vCenter * c1-10-182-13-37.eng.vmware.com

Network * ● Distributed Portgroup ○ Standard Switch Network ○ NSX Logical Switch

	Network	Distributed Switch	VLAN
○	dvPortgroup1	taggableDvs	0
○	dvPortgroup2	taggable2Dvs	0
○	fleet2001	taggableDvs	2001
○	fleet2003	taggableDvs	2003

28 Networks

Name *

IP Pools: IP Pool - 0

IP Ranges

HOW MANY FREE IP ADDRESSES DO YOU NEED?

press enter/return or comma (,) to complete input

Prefix Length

Gateway

Primary DNS

Secondary DNS

DNS Suffix

MTU * 1500

HCX Traffic Type (optional)

Selections will be suggested in the respective compute profile configuration.

□ Management □ HCX Uplink □ vSphere Replication □ vMotion □ Sentinel Guest Network

4 Name the Network Profile.

5 Provide the IP address pool details for the network profile.

   a Provide an IP address range available for the HCX appliances. Use a comma to separate multiple discontiguous ranges within the same subnet.

   b Enter the **Prefix Length** for the network containing the IP ranges provided.

   c Enter the **Default Gateway Address** for the network.

   d Specify the DNS server information.

6 Enter the MTU value.

7 (Optional) Using the check boxes, associate one or more suggested traffic types with the network selection: Management, HCX Uplink, vSphere Replication, vMotion, Sentinel Guest Network.

   The traffic type selection appears as a suggestion of which networks to use when creating the Compute Profile. It does not prevent the network from being used for other types of network traffic.

8 To complete the **Network Profile** configuration, click **Create**.

What to do next

Created **Network Profiles** are designated to one or more specific functions during a **Compute Profile** configuration, or when to override default Network Profiles when creating a **Service Mesh**.

Note To edit an existing Network Profile, navigate to the specific Network Profile and click **Edit**.

Creating a Compute Profile

A **Compute Profile** contains the compute, storage, and network settings that HCX uses on this site to deploy the Interconnect-dedicated virtual appliances when a Service Mesh is added.

Create a Compute Profile in the Multi-Site Service Mesh interface in both the source and the destination HCX environments using the planned configuration options for each site, respectively.

Prerequisites

- Install and configure HCX Manager.
- To obtain the optimum system usage, assign resource configurations based on HCX deployment considerations.
- Use the planned configurations collected using the checklist described in Getting Started with VMware HCX.

Procedure

1 Navigate to the **Compute Profiles** interface:

   a At the source site, open vSphere Client and navigate to the **HCX** plug-in > **Interconnect** > **Multi-Site Service Mesh** > **Compute Profiles**.

   b At the destination site, navigate to https://hcx-cloud-ip-or-fqdn > **Multi-Site Service Mesh** > **Compute Profiles**.

   The system displays all the defined **Compute Profiles**. If no profiles have been configured, the system highlights the **Create Compute Profile** option.

2 Click **Create Compute Profile**.

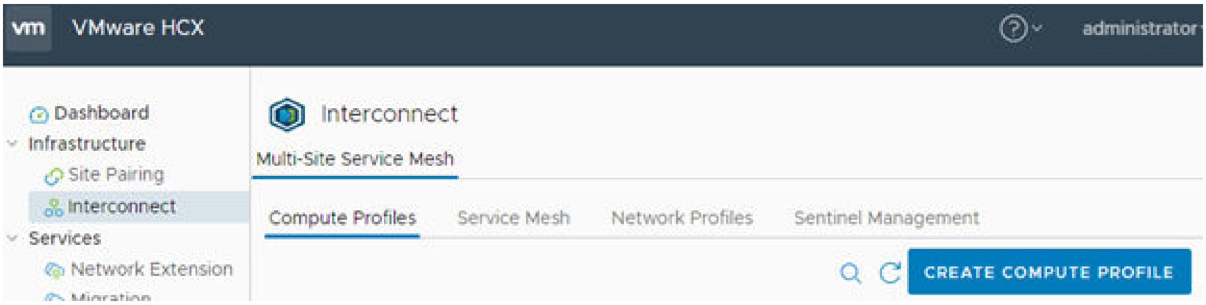

3 Name the Compute Profile:

a Enter a name for the Compute Profile.

b Click **Continue**.

4 Select the HCX services to be activated. Click **Continue**.

Note Premium services require the HCX Enterprise license.

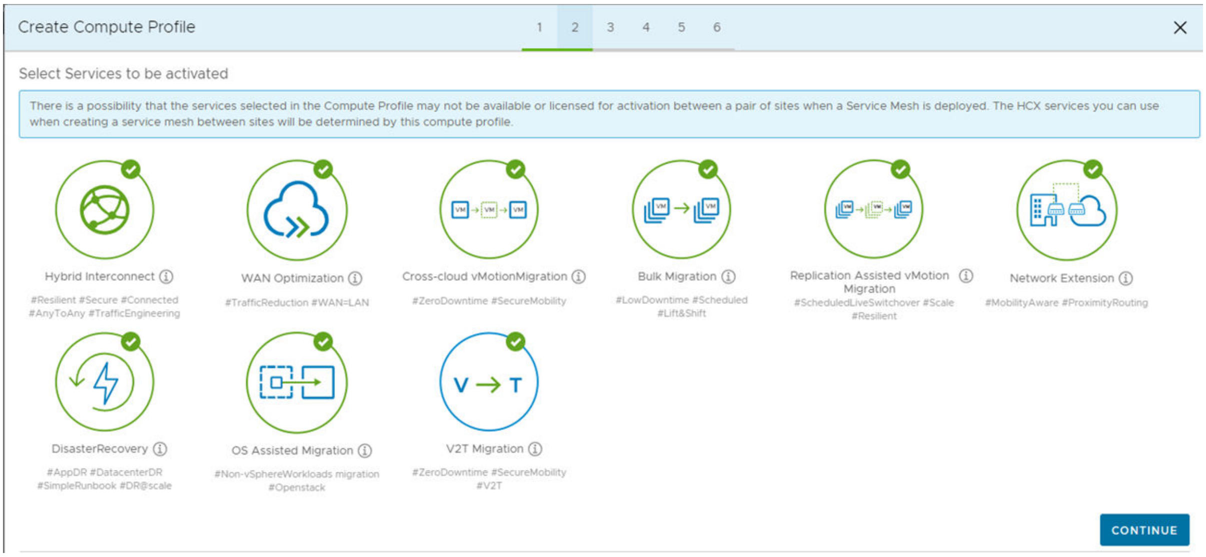
Create Compute Profile
1
2
3
4
5
6
×

Select Services to be activated

There is a possibility that the services selected in the Compute Profile may not be available or licensed for activation between a pair of sites when a Service Mesh is deployed. The HCX services you can use when creating a service mesh between sites will be determined by this compute profile.

Hybrid Interconnect ⓘ
#Resilient #Secure #Connected #AnyToAny #TrafficEngineering
WAN Optimization ⓘ
#TrafficReduction #WAN=LAN
Cross-cloud vMotionMigration ⓘ
#ZeroDowntime #SecureMobility
Bulk Migration ⓘ
#LowDowntime #Scheduled #Lift&Shift
Replication Assisted vMotion Migration ⓘ
#ScheduledLiveSwitchover #Scale #Resilient
Network Extension ⓘ
#MobilityAware #ProximityRouting
DisasterRecovery ⓘ
#AppDR #DatacenterDR #SimpleRunbook #DR@scale
OS Assisted Migration ⓘ
#Non-vSphereWorkloads migration #Openstack
V2T Migration ⓘ
#ZeroDowntime #SecureMobility #V2T
CONTINUE

5 Select the Service Resources:

a Click the **Select Resources** drop-down menu.

b Select each cluster to be activated for HCX services.

If there is only one cluster, it is selected automatically.

c Click **Continue**.

The Select Deployment Resources and Reservations screen appears.

6 Make your resource, and resource reservation selections.

   a From the **Select Resource** drop-down menu, and select each cluster or resource pool to be used when deploying HCX Interconnect appliances.

   Note For HCX WAN-OPT appliance deployment in NSX-T enabled environments, NSX-T Overlay Transport Zone must be associated to all hosts that are part of the deployment cluster. This should be done in addition to the NSX-T Overlay or NSX-T VLAN network extension configuration part of Network Containers section.

   b From the **Select Datastore** drop-down menu, and select the datastore to be used when deploying HCX Interconnect appliances.

   When multiple compute resources or datastores are selected, HCX uses the first selection until its capacity is exhausted.

   c (Optional) From the **Select Folder** drop-down menu, and specify a folder in which to deploy the HCX appliances.

   d Using the slide bar, select the amount of CPU and memory to reserve for HCX operations.

   As a best practice, set the CPU and memory reservation to 100 percent.

   For example, setting **Memory Reservation** to 100 percent ensures that all of the memory allocated for HCX appliances is always available for HCX operations.

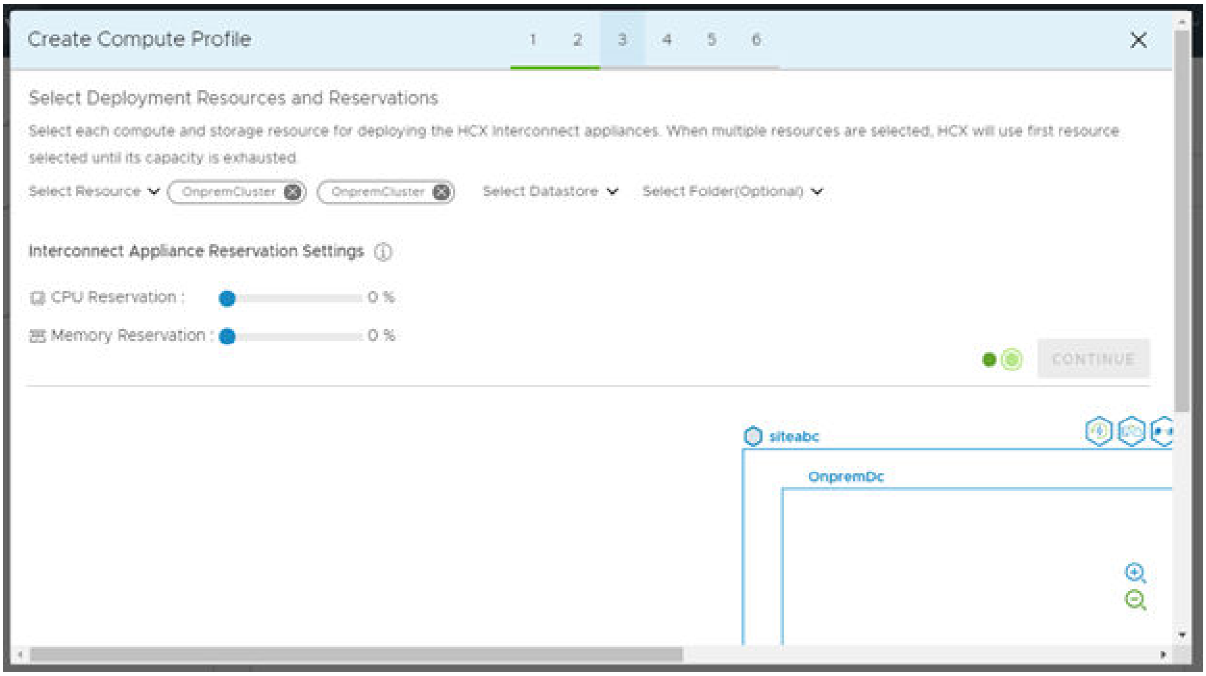

Create Compute Profile

Select Deployment Resources and Reservations

Select each compute and storage resource for deploying the HCX Interconnect appliances. When multiple resources are selected, HCX will use first resource selected until its capacity is exhausted

Select Resource: OnpremCluster, OnpremCluster

Select Datastore

Select Folder(Optional)

Interconnect Appliance Reservation Settings

CPU Reservation : 0 %

Memory Reservation : 0 %

7 Select the Management Network Profile:

   a Click the **Select Management Network Profile** drop-down menu.

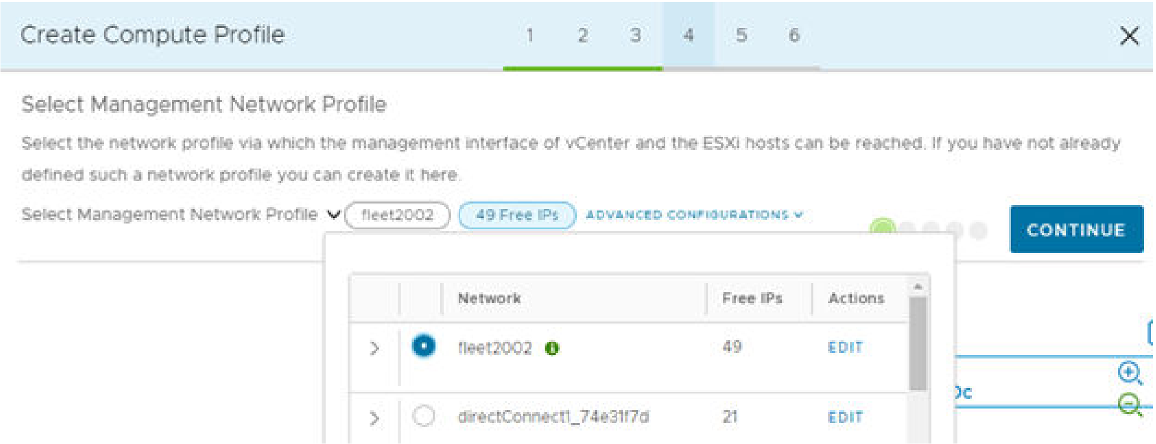

Select Management Network Profile

Select the network profile via which the management interface of vCenter and the ESXi hosts can be reached. If you have not already defined such a network profile you can create it here.

Select Management Network Profile: fleet2002 — 49 Free IPs

		Network	Free IPs	Actions
>	●	fleet2002	49	EDIT
>	○	directConnect1_74e31f7d	21	EDIT

   b Select an existing Network Profile or click **Create Network Profile** to create it.

   Reference the Creating a Network Profile topic for more details.

   Note Networks identified with an information icon have been suggested for use with this type of network in the Network Profile. This is only a suggestion, and you can select other networks for this traffic type.

   c Expand the selected Management Network Profile to view its details and free IP Addresses. Click **Close** when done reviewing.

   d Click **Continue**.

8 Select the Uplink Network Profile:

- The **Network Profile** previously selected for another function, like Management can also be assigned as the **Uplink Network Profile**.
- Multiple Network Profiles can be selected.

   a Click the **Select Uplink Network Profile** drop-down menu.

   b Select one or more existing **Network Profile**, or click **Create Network Profile** to create it.

   Reference the Creating a Network Profile topic for more details.

   c Expand the selected **Uplink Network Profile** to view its details and free IP Addresses. Click **Close** when done reviewing.

   d Click **Continue**.

HCX Manager now updates the topology view to depict the configured Network Profile. As shown in the diagram, the Compute Profile configuration tool displays a symbolic map of the network links between the Interconnect appliance virtual machines to be deployed for the selected Uplink network.

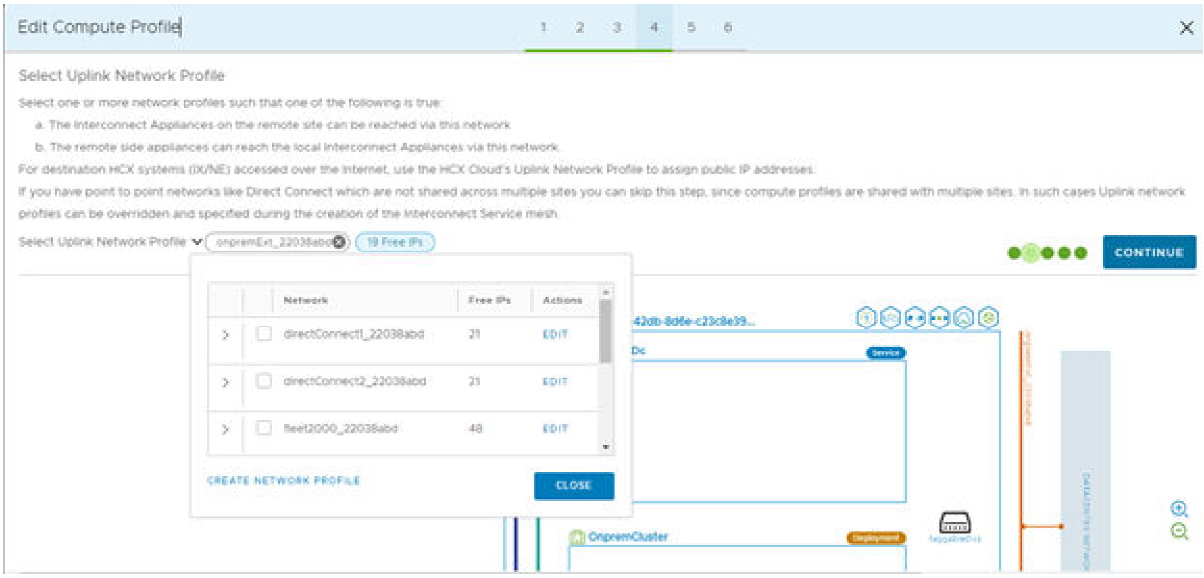

Edit Compute Profile

1 2 3 4 5 6

Select Uplink Network Profile

Select one or more network profiles such that one of the following is true:

a. The Interconnect Appliances on the remote site can be reached via this network

b. The remote side appliances can reach the local Interconnect Appliances via this network

For destination HCX systems (IX/NE) accessed over the Internet, use the HCX Cloud's Uplink Network Profile to assign public IP addresses.

If you have point to point networks like Direct Connect which are not shared across multiple sites you can skip this step, since compute profiles are shared with multiple sites. In such cases Uplink network profiles can be overridden and specified during the creation of the Interconnect Service mesh.

Select Uplink Network Profile: onpremExt_22038abd | 19 Free IPs

		Network	Free IPs	Actions
>	□	directConnect1_22038abd	21	EDIT
>	□	directConnect2_22038abd	21	EDIT
>	□	fleet2000_22038abd	48	EDIT

CREATE NETWORK PROFILE | CLOSE | CONTINUE

9 Select the vMotion Network Profile:

a Click the **Select vMotion Network Profile** drop-down menu.

b Select an existing Network Profile or click **Create Network Profile** to create it.

Reference the Creating a Network Profile topic for more details.

c Expand the selected vMotion Network Profile to view its details and free IP Addresses. Click **Close** when done reviewing.

d Click **Continue**.

The Network Profile tool now displays a topology view that shows how the selected vMotion Network connects the HCX Interconnect appliances assigned to the profile.

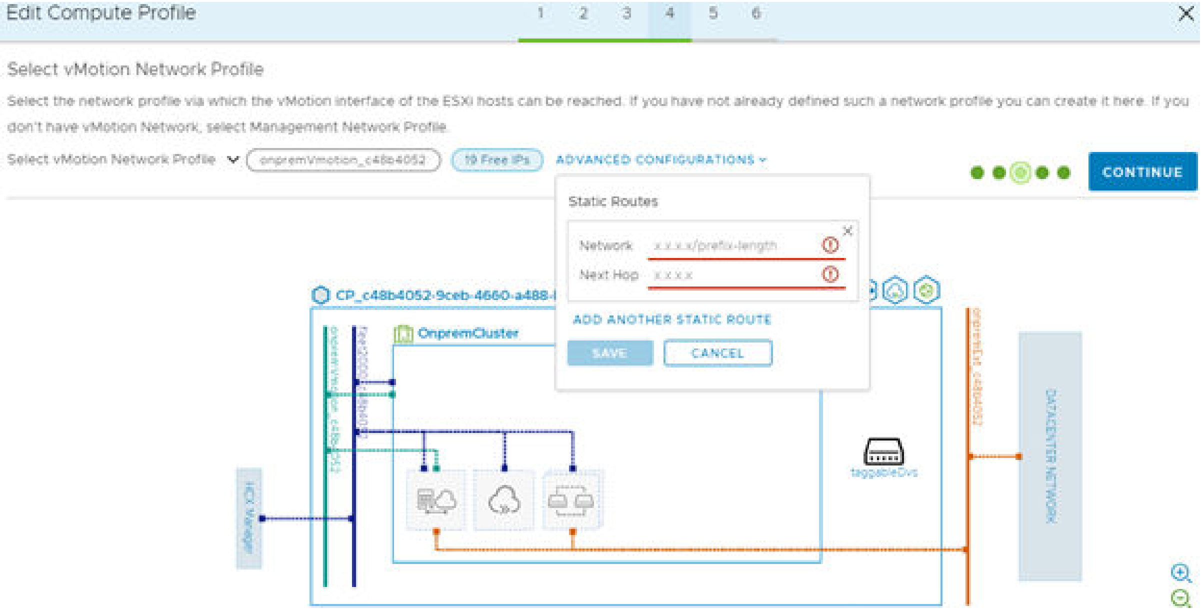

Edit Compute Profile

1 2 3 4 5 6

Select vMotion Network Profile

Select the network profile via which the vMotion interface of the ESXi hosts can be reached. If you have not already defined such a network profile you can create it here. If you don't have vMotion Network, select Management Network Profile.

Select vMotion Network Profile: onpremVmotion_c48b4052 | 19 Free IPs | ADVANCED CONFIGURATIONS

Static Routes

Network: x.x.x.x/prefix-length

Next Hop: x.x.x.x

ADD ANOTHER STATIC ROUTE | SAVE | CANCEL | CONTINUE

10 Select the vSphere Replication Network Profile:

Assigning a vSphere Replication Network Profile is useful when there is a VMkernel interface for the network traffic that is exclusive to vSphere Replication operations. If the Management Network Profile is used for Replication operations, click **Continue** to skip this step.

a Click the **Select vSphere Replication Network Profile** drop-down menu.

b Select an existing Network Profile or click **Create Network Profile** to create it.

Reference the Creating a Network Profile topic for more details.

c Expand the selected vSphere Replication Network Profile to view its details and free IP Addresses. Click **Close** when done reviewing.

d Click **Continue**.

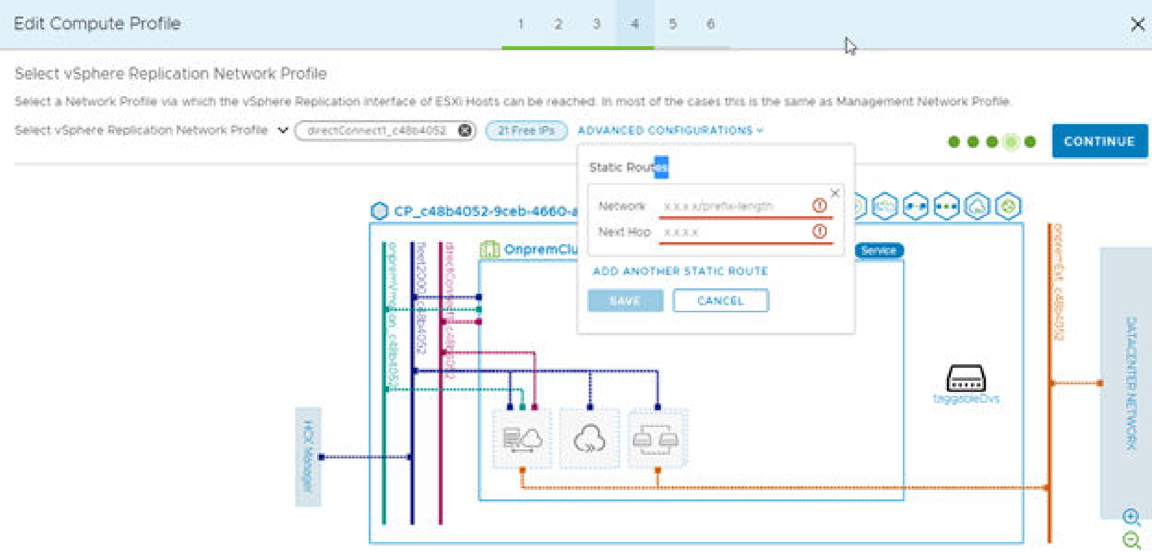

Edit Compute Profile

Select vSphere Replication Network Profile

Select a Network Profile via which the vSphere Replication interface of ESXi Hosts can be reached. In most of the cases this is the same as Management Network Profile.

Select vSphere Replication Network Profile: directConnect1_c48b4052

11 If HCX OS Assisted Migration is activated, select the Guest Network Profile:

- This is the network on which guest virtual machines communicate with the HCX SGW for OS Assisted Migration.
- The **Network Profile** previously selected for another function, like Management can also be assigned as the **Guest Network Profile**.

Note This step appears on only the source Compute Profile interface, and only if you have selected OS Assisted Migration from the list of available services.

a Click the **Select Guest Network Profile** drop-down menu.

b Select a network, or click **Create Network Profile** to create it.

For additional information, see Creating a Network Profile.

c Expand the selected **Guest Network Profile** to view its details and free IP Addresses, and click **Close**.

d Click **Continue**.

The Guest Network is depicted in the topology.

12 Select Network Containers Eligible for Network Extension:

Note NSX-T Overlay and NSX-T VLAN networks can be extended. These extensions are always created as NSX Overlay networks at the destination. HCX network extension is always to NSX Overlay networks. Pre-created segments in the vSphere Distributed Switch, CVDS, or VLAN transport zone, cannot be used as a destination.

a Click the **Select Network Containers** drop-down menu.

The system displays the list of network containers found in the service resources selected in a previous step.

b Select entries with virtual machine networks. Click **Close**.

You can select vDS, N-VDS, Transport Zones, or some combination for Network Extension. If only one network container is found in the previously selected service resources, it is pre-selected.

Note Only the virtual machine networks on the selected switches or Transport Zones can be used during network extension operations.

c Click **Advanced Configurations** and set the number of Network Extension appliances that can be deployed for this container.

If this number is not known, select Unlimited. This number impacts the number of Network Extension appliances that can be deployed through the Service Mesh. If the number is too low, it can prevent the Service Mesh from deploying needed appliances. If the number is too high, it is possible to deploy too many Network Extension appliances than the planned resources can support.

d Click **Continue**.

The topology view is dynamically updated, depicting your selections.

13 Pre-Deployment Validation:

HCX checks whether the selected configurations are valid for interconnect deployments.

a Address any errors reported by the validation.

b To see which firewall rules might be required if the service mesh uses this service profile, click **Review Connection Rules**.

c To export the rules, click **COPY ALL** to copy them to the clipboard in the JSON format.

d Click **Next**.

14 Ready to Complete:

   a Review the configuration. The topology diagram depicts the selected configurations.

   b To create the Compute Profile, click **Finish**.

Results

A Compute Profile is created, and can be used when creating a service mesh.

What to do next

Once there are valid Compute Profiles in the source and destination environments, use the HCX Manager UI at the source site to create the Interconnect Service Mesh.

Note To edit an existing Compute Profile, navigate to the specific Compute profile, and click **Edit**.

Creating a Service Mesh

An HCX Service Mesh is the effective HCX services configuration for a source and a destination site. You can add a Service Mesh to a connected Site Pair that has a valid Compute Profile created on both of the sites.

Adding a Service Mesh initiates the deployment of HCX Interconnect virtual appliances on both of the sites. An interconnect Service Mesh is always created at the source site.

Prerequisites

Creating a Service Mesh requires:

- A connected Site Pair.
- A valid compute profile at the HCX Source site.
- A valid compute profile at the HCX destination site.
- For each switch that is present in the Compute Profile at both the source and the destination sites, the switch must span all hosts in at least one of the compute clusters. If the switch does not span all hosts in the compute cluster, then it is possible that the Network Extension appliance is deployed on a different host in a compute cluster and spans across a different switch. In this case, the Service Mesh deployment can fail.

Procedure

1 Navigate to the **Service Mesh** interface:

   a In the vSphere Client, navigate to **HCX** > **Interconnect** > **Multi-Site Service Mesh** > **Service Mesh** tab.

   Created **Service Mesh** configurations are listed.

2 Click **Create Service Mesh**:

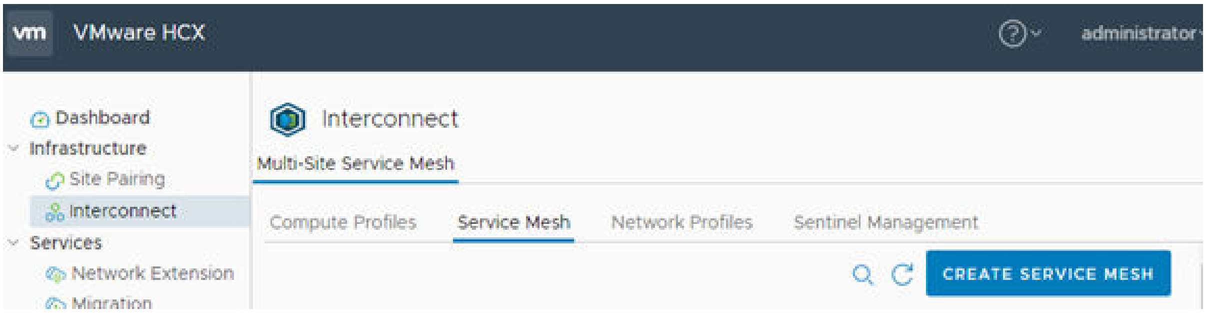

3 Select Sites:

a Click each drop-down and select a source and a destination site. Only connected **Site Pairs** are displayed.

b Click **Continue**.

4 Select Compute Profiles:

a Click the **Select Source Compute Profile** drop-down and select a **Compute Profile**.

b Click the **Select Remote Compute Profile** drop-down and select a **Compute Profile**.

c Click **Continue**.

5 Select the HCX services to be activated, and click **Continue**:

Note Premium services require an additional HCX Enterprise license.

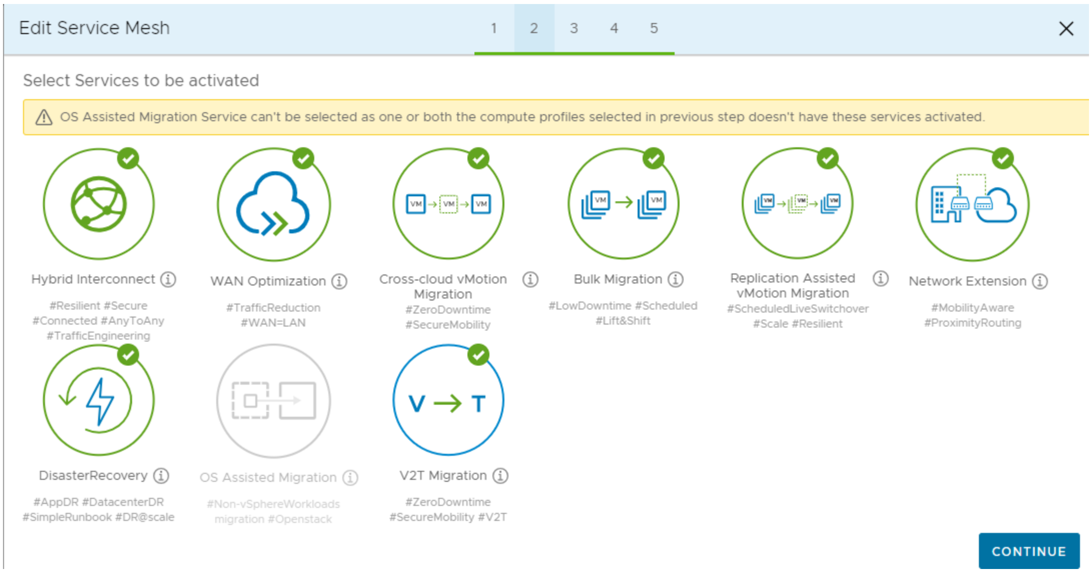

Edit Service Mesh

Select Services to be activated

⚠ OS Assisted Migration Service can't be selected as one or both the compute profiles selected in previous step doesn't have these services activated.

1 2 3 4 5
Hybrid Interconnect ⓘ
#Resilient #Secure #Connected #AnyToAny #TrafficEngineering
WAN Optimization ⓘ
#TrafficReduction #WAN=LAN
Cross-cloud vMotion Migration ⓘ
#ZeroDowntime #SecureMobility
Bulk Migration ⓘ
#LowDowntime #Scheduled #Lift&Shift
Replication Assisted vMotion Migration ⓘ
#ScheduledLiveSwitchover #Scale #Resilient
Network Extension ⓘ
#MobilityAware #ProximityRouting
DisasterRecovery ⓘ
#AppDR #DatacenterDR #SimpleRunbook #DR@scale
OS Assisted Migration ⓘ
#Non-vSphereWorkloads migration #Openstack
V2T Migration ⓘ
#ZeroDowntime #SecureMobility #V2T
CONTINUE

6 (Optional) Override the default Uplink Network Profile:

By default, the HCX interconnect uses Uplink Network Profiles defined in the Compute Profile for the source and the destination sites. You can override the default.

As an example, an override can be useful in Cloud Director-based deployments where an uplink network that deviates from a common configuration is created for an Organization to consume during the **Service Mesh** creation.

a Click the **Select Source Uplink Network Profile** drop-down.

b Select one or more networks. Click **Close**.

The HCX Service Mesh can use up to three HCX Uplinks, adding network path failover and improving overall resiliency for HCX services. Multiple HCX Uplinks are not aggregated for increased throughput capacity. The following specific behaviors apply:

- HCX attempts to load balance traffic on the Network Extension (HCX-NE) appliance based on characteristics of the flow and the performance of the uplinks.
- HCX does not load balance migration traffic on the Interconnect (HCX-IX) appliance. Additional uplinks might or might not be used.

c Click **Continue**.

d Optionally, repeat these steps for the destination site.

7 Configure the Network Extension appliances deployed per switch or Transport Zone:

As an example, this advanced configuration can be useful when deploying Network Extension appliances to extend high volume source networks.

a In **Advanced Configuration - Network Extension Appliance Scale Out**, review the default Extension appliances per Network Container.

b For each entry, set the number of Network Extension appliances that HCX deploys when it activates the Service Mesh configuration.

Extended network service can be carried by a single (standalone) Network Extension appliance at each site, or an HA group that consists of two Network Extension appliances at each site. For example, to create two standalone Network Extension appliances and one HA group for a container entry, set the scale-out number to 4 (2 + 1 x 2 = 4) in the Service Mesh.

The default setting is 1. This setting restricts the Service Mesh to deploying one Network Extension appliance.

Note You must configure the Network Extension Appliance limit in the Compute Profile at both the source and remote sites to equal or exceed the number of scale-out appliances set in the Service Mesh.

For the system resource considerations, see System Requirements.

c Click **Continue**.

8 (Optional) Configure HCX Traffic Engineering features:

The Application Path Resiliency and TCP Flow Conditioning features are available with the HCX Enterprise license.

a To create multiple transport tunnels for directing the HCX traffic to a destination site, check **Application Path Resiliency**.

Enabling Application Path Resiliency (APR) creates up to eight transport tunnels between each Interconnet and Network Extension appliance uplink interface IP address pair between sites. If a few tunnels fail, there is no impact in the data traffic as only one transport tunnel out of eight is used always to provide secure data transfer across the Wide Area Network (WAN) or Internet connection.

Application Path Resiliency forwards traffic over one tunnel at a time and does not load balance across multiple paths.

Note To view the available tunnels after completing the Service Mesh configuration, navigate to **Interconnect > Multi-Site Service Mesh > Service Mesh > View Appliances** and expand the HCX-WAN-IX appliance.

Important For additional dynamic tunnel requirement, the source Interconnect (IX) and Network Extension (NE) appliances uses a random source UDP port in the 4500 – 4628 range and target UDP port as 4500 to create a different flow for each subsequent tunnel. The reverse tunnel originated by target IX/NE appliances have source port as UDP 4500 and destination ports from same random ports used by source appliances for the forward direction in the range 4500 – 4628.

Ensure the firewall settings on either side allow for that connectivity.

b To dynamically manage the TCP segment size and optimize the transport performance for the HCX Network Extension service traffic, check **TCP Flow Conditioning**.

This option is available only after activating the HCX Network Extension service.

c To manage the bandwidth consumed for migrations across all uplink networks, use the up and down arrows to change the bandwidth setting.

This option is available only after activating the HCX WAN Optimization service.

Note It is a best practice to retain the default setting of 10000 Mb/S.

9 Review Topology Preview:

a Review the selected clusters and the resources.

b Click **Continue**.

10 Ready to Complete:

   a To view a summary of the **Service Mesh** selections, click the **here** link.

   b Name the **Service Mesh**.

   The Service Mesh name has a limit of 50 characters.

   c To create the service mesh, click **Finish**.

   After the Service Mesh configuration is complete, verify the underlay network performance for each Uplink Network. The underlay network performance must meet the minimum requirements for HCX services. See Understanding HCX Transport Analytics.

What to do next

If it is necessary to make any direct changes to an existing Service Mesh, such as activating or deactivating services and overriding uplinks, select **Interconnect > Service Mesh > Edit**. The editing workflow includes a preview screen, listing the changes and describing the impact of those changes on related services prior to finishing the procedure. You can select to complete or cancel the update.

Updating and Synchronizing the Service Mesh

The Service Mesh is the effective HCX service configuration between a source and a destination site, and it must be kept updated between the pair following changes in the Compute Profile, Network Profile, or Service Mesh itself.

When there is an update to the Compute Profile, Network Profile, or Service Mesh, it can have an impact on the system configuration, including possible downtime. The HCX Manager UI provides two operations for addressing Service Mesh changes: **Edit** and **Resync**. Which operation is required depends on the change. The Edit operation edits the Service Mesh. The Resync operation takes changes to the Compute Profile and applies them to the Service Mesh.

If the Service Mesh configuration includes Network Extension High Availability (HA) groups, a Resync or Edit operation redeploys all Network Extension appliances in the HA group.

To perform an Edit operation, go to **Interconnect > Service Mesh > Edit**.

To perform a Resync operation, go to **Interconnect > Service Mesh > Resync**.

Note Resync is intended for use only in healthy Service Mesh environments. Resync is not for troubleshooting a Service Mesh in an error state.

Editing and synchronizing the Service Mesh is available from the source site HCX Manager UI where the Service Mesh was created. These update operations are not available at the destination site.

Procedure

1 Log in to the HCX Manager at the source site: `<https://hcxmgr-ip-or-fqdn>:443`.

2 Go to **Interconnect > Service Mesh**.

3 Click **Resync**.

Verify that the changes appear in the Service Mesh configuration.

Renaming Service Mesh Appliances

You can update the HCX Service Mesh appliance names for either HCX Connector or HCX Cloud Manager.

The HCX Service Mesh, which can comprise the HCX-WAN-IX, HCX-NE, and HCX-WAN-OPT appliances, provides the interconnect services between paired source and destination sites. You can change the appliance names in your environment.

Note Do not use the same name for appliances in a site pair.

Procedure

1 In the HCX Manager UI, navigate to **Interconnnect > Service Mesh.**

2 Click **View Appliances**.

A list of appliances appears.

3 Select an appliance from the list, and click the **Rename Appliance** tab.

4 Enter the new appliance name, and click **Rename**.

HCX initiates the name change operation and syncs the changes with the peer site. To monitor the status of the rename operation, click the **Task** tab.

What to do next

Repeat the procedure for any further name changes.

Understanding HCX Transport Analytics

You can test and review the underlay network performance for each uplink network using HCX Transport Analytics.

Transport Analytics help to characterize the health and performance of the underlay network. Run bandwidth tests on each uplink network in each Service Mesh configuration to determine whether the underlay network supporting that uplink meets the minimum requirements for HCX services. Understanding the underlay network performance can help in planning migration wave size and in analyzing data transfer and performance issues.

Note HCX Transport Analytics test the end-to-end underlay network supporting each uplink, and not the uplink network itself. The test checks the underlay network between the source and the destination site. Throughout this section, underlay network results are organized and reported against each uplink network in the Service Mesh.

To support seamless system operation between sites, Network Extension and migration operations have minimum requirements for underlay network bandwidth, latency, and loss. These services all use underlay networks for transporting data from the source site to the destination site. If the underlay network does not meet the minimum requirements, it can result in degraded service performance or failure. For a complete description of network underlays and the minimum requirements for HCX services, see Network Underlay Minimum Requirements.

Verifying Underlay Network Performance for Service Mesh Uplinks

HCX Transport Analytics help with monitoring the connectivity health between sites for a Service Mesh and with characterizing the underlay network requirements for HCX services.

The Transport Analytics page lists all Service Mesh configurations for a site pair with a separate card for each Service Mesh. Each Service Mesh card has the same information:

- Connected Sites.
- Performance Insights
- Services Overview.
- Uplinks Overview.

You can use the search option at the top right of the page to find a specific Service Mesh.

You access Transport Analytics from the HCX Manager interface after logging in to the site manager: https://*your-site-manager-ip*:443. After logging in, navigate to **Infrastructure > Transport Analytics**.

Prerequisites

- Service Mesh appliances are at version 4.4.0 or later.
- If you have changed the uplink networks in the Compute Profile, navigate to **Interconnect > Service Mesh,**, and click **Resync** to apply those changes to the Service Mesh.

Procedure

1 In the HCX Manager UI, click **Transport Analytics**.

The transport analytics page displays the Service Mesh inventory.

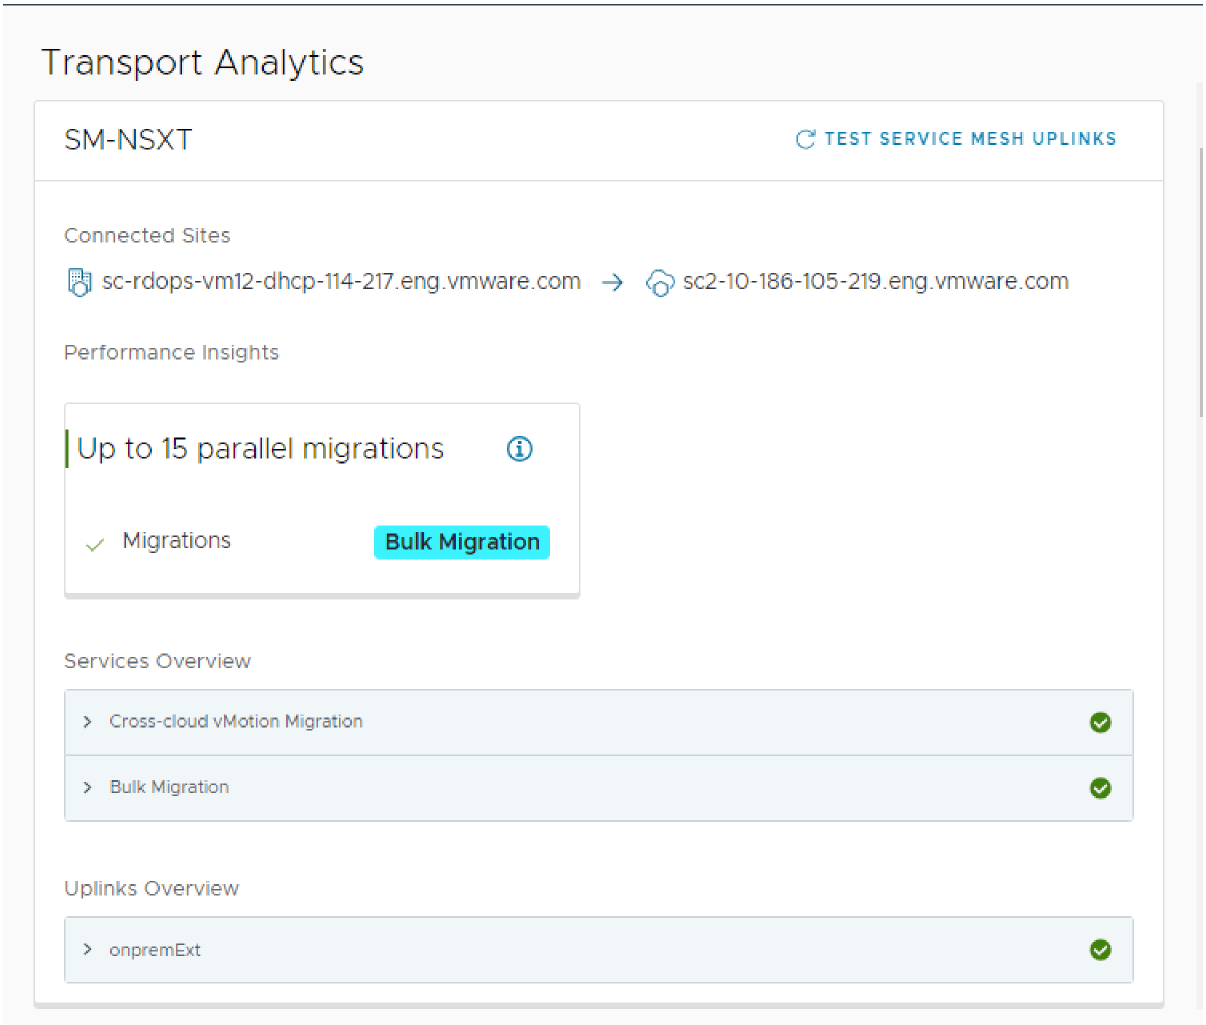

Transport Analytics

SM-NSXT

Connected Sites

sc-rdops-vm12-dhcp-114-217.eng.vmware.com → sc2-10-186-105-219.eng.vmware.com

Performance Insights

Up to 15 parallel migrations

✓ Migrations Bulk Migration

Services Overview

Cross-cloud vMotion Migration	✓
Bulk Migration	✓

Uplinks Overview

onpremExt	✓

2 Click **Test Service Mesh Uplinks**.

Note It is a best practice to test the Service Mesh uplinks independent of other tests. Running tests like CCLI perftest in parallel can skew test results.

3 Review each section of the Service Mesh card for health details:

Section	Description
Connected sites	Identifies the source and the destination sites for the site pair.
Performance insights (Currently available for Bulk Migration only)	Based on the bandwidth reported by the uplink test and the bandwidth requirement of the HCX service, HCX calculates the underlay network performance associated with the uplink. From that calculation, HCX recommends a number of simultaneous migrations that the underlay network can support at the point when the test was run.

Section	Description
Services Overview	Lists the HCX services enabled in the Service Mesh. Expand each service entry to view the transport thresholds: Minimum Bandwidth, Maximum Latency, and Maximum Loss. The transport thresholds reflect the underlay network requirements for each service. The health of a service relative to the underlay network is indicated by either a checkmark or an exclamation mark. A checkmark indicates that the underlay network performance meets the minimum requirements for the service, and the uplink network is in a healthy state for the specified HCX service. An exclamation mark indicates that the uplink network is in a degraded state relative to the underlay network performance, and that uplink is not recommended to use with the specified HCX service.
Uplinks Overview	Provides the transport analysis and the transport monitoring of the underlay network for each uplink in the Service Mesh. Transport analysis lists the tested transport thresholds of the underlay network. Expand each uplink network entry to view Available Bandwidth, Latency, and Loss metrics. Select **Re-run Test** to start a new test for the selected uplink network. The health of a service relative to the underlay network is indicated by either a checkmark or an exclamation mark. A checkmark indicates that the service thresholds are met for all the services configured in the Service Mesh. An exclamation mark is shown if any one of the service thresholds are not met for the services configured in the Service Mesh. If a new network uplink is added to the Service Mesh, you must run an initial test on that uplink to generate the transport thresholds. Following the initial test, re-run the test to evaluate the current state of the underlay network. **Note** Available bandwidth is what is currently available at the time of the test. Transport threshold metrics are static and are based on the most recent test. To update the metrics for an uplink network, you must re-run the test. When an uplink network is used in multiple Service Mesh configurations, each Service Mesh using that uplink is updated with new metrics. Transport monitoring provides detailed, time-series bandwidth, latency, and loss information for each uplink network. For more information, see Monitoring Underlay Network Performance Using Transport Monitor.

4 (Optional) To view the detailed metrics, select **Transport Monitor** for each uplink network.

5 For each Service Mesh card listed under Transport Analytics, review the health details to verify that the underlay network performance meets the requirements for the HCX services enabled for that Service Mesh.

What to do next

If HCX Transport Analytics indicate that the uplink networks are healthy, proceed to using HCX services. If the uplink network status is Degraded, use Transport Monitor to help identify and resolve underlay network performance issues before using HCX services.

Monitoring Underlay Network Performance Using Transport Monitor

HCX Transport Monitor provides detailed, time-series throughput, latency, and loss information for the uplink networks in a Service Mesh.

Transport Monitor provides a graphical representation of underlay network performance for each uplink network. Transport Monitor collects the metrics every minute. For each selected uplink, the page provides the same information:

- Metrics selections: Throughput, Latency, Loss
- Duration selections: Real Time, Last Day, Last Week, Custom
- Graphic displays based on selected metrics and duration

In some scenarios, metrics data is not available for an uplink and underlay analytics information is not displayed:

- The Service Mesh appliance tunnels are down.
- The source environment is using Network Address Translation (NAT).
- The firewall security policies disallow the ICMP traffic.

Procedure

1 From the Transport Analytics page in the HCX Manager UI, select **Transport Monitor**.

HCX displays the Transport Monitor page.

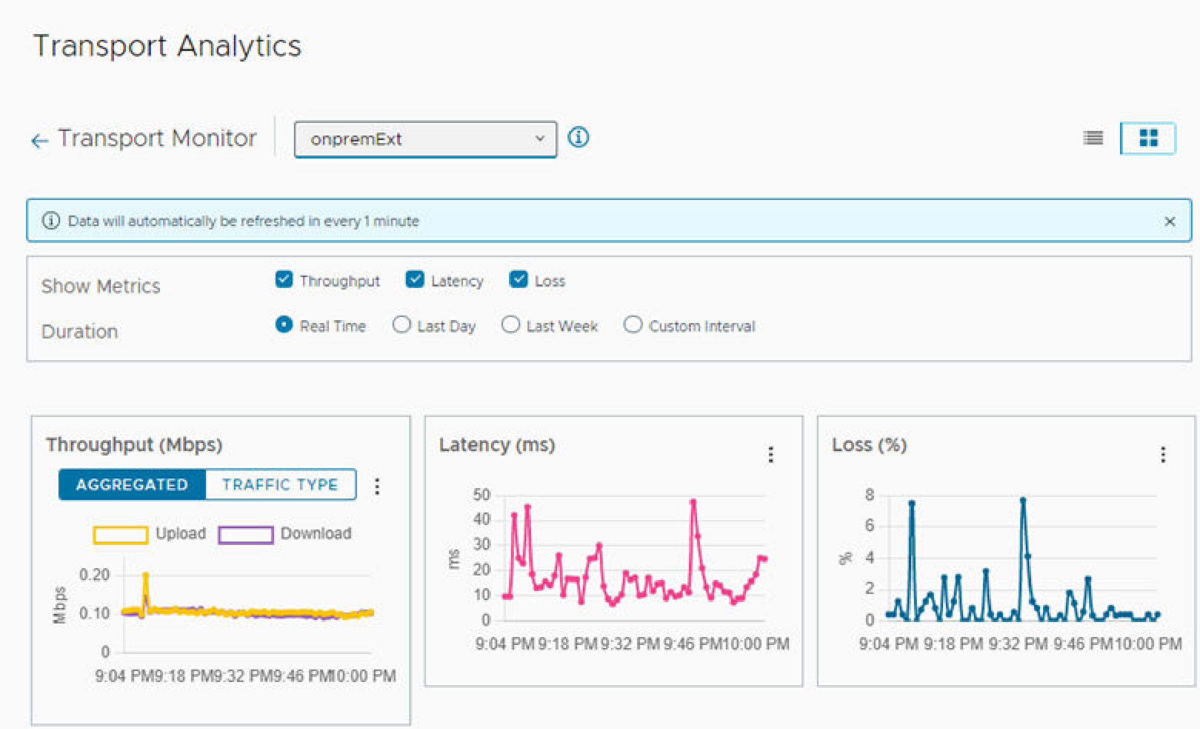

Transport Analytics

← Transport Monitor | onpremExt ⓘ

ⓘ Data will automatically be refreshed in every 1 minute

Show Metrics ☑ Throughput ☑ Latency ☑ Loss

Duration ● Real Time ○ Last Day ○ Last Week ○ Custom Interval

Throughput (Mbps)
AGGREGATED
TRAFFIC TYPE
Upload
Download
Mbps
0.20
0.10
0
9:04 PM9:18 PM9:32 PM9:46 PM10:00 PM
Latency (ms)
ms
50
40
30
20
10
0
9:04 PM 9:18 PM 9:32 PM 9:46 PM10:00 PM
Loss (%)
%
8
6
4
2
0
9:04 PM 9:18 PM 9:32 PM 9:46 PM10:00 PM

2 At the top of the page, use the pull-down menu to select an uplink network.

The information icon to the right of the pull-down menu specifies which Service Mesh configurations use the selected uplink network.

3 At the top right of the page, click one of the icons to toggle the orientation for graphs between the inline and the portrait mode.

4 For **Show Metrics**, click the check box in front of each option that you want to display: **Throughput**, **Latency**, **Loss**.

5 For **Duration**, select the time frame to display in the graph:

Duration	Description
Real Time	The system automatically updates the information in the graph every 1 minute.
Last Day	Displays information from the last 24 hours.
Last Week	Displays information from the last 7 days
Custom Interval	Displays information for the selected date range.

6 Review the graphs for each metric type:

Graph	Description
Throughput	Displays the amount data traffic using the uplink network from all HCX services. At the upper-right of the graph, you can select from the Aggregated or the Traffic Type data. Aggregated data combines the traffic throughput for all HCX migration and Network Extension services using the uplink. Traffic Type displays the data separated by the migration or the Network Extension traffic. In either case, the graph displays both upload and download throughput. To display the minimum threshold for each migration type, use the **Show Thresholds** pull-down menu in the upper-right corner of the graph.
Latency	Displays the latency of underlay network supporting the uplink. To display the minimum threshold for each migration type, use the Show Thresholds pull-down menu in the upper-right corner of the graph.
Loss	Displays the percentage of data loss for the underlay network supporting the uplink. To display the minimum threshold for each migration type, use the Show Thresholds pull-down menu in the upper-right corner of the graph.

Sentinel Management

You must download and install HCX Sentinel on all guest virtual machines requiring migration using HCX OS Assisted Migration. Sentinel gathers the system configuration from the guest VM and assists with the data replication.

The guest VM information is used by various HCX OS Assisted Migration service processes. In part, the information is used to create an inventory of guest VM systems for migration and to help replication processes prepare the disks on the replica VM for replication and migration.

Sentinel also helps with the data replication by reading data written to the source disks and passing that data to the SDR appliance at the destination site.

The Sentinel Management tab, which provides access to downloading the Sentinel software, appears in the HCX Interconnect interface when an HCX Enterprise license is activated, and you have a deployed a service mesh with an SGW/SDR pair deployed. For more information about OS Assisted Migration, see Understanding VMware HCX OS Assisted Migration.

Downloading and Installing HCX Sentinel Agent Software

When performing migrations from non-vSphere virtual machines, you must install the HCX Sentinel on all guest virtual machines requiring migration using HCX OS Assisted Migration. The sentinel agent gathers the system configuration from the guest virtual machine and assists with the data replication.

Prerequisites

HCX Enterprise license is activated.

HCX OS Assisted Migration is activated in Service Mesh.

Procedure

1 In the vCenter Server for the HCX Connector, navigate to **Interconnect** > **Multi-Site Service Mesh** > **Sentinel Management**.

2 Download the software bundle appropriate for the environment that you are migrating.

   The Sentinel software bundle is downloaded to the local machine with the name `<SGW-name>-linux-sentinel-installer.sh` or `<SGW-name>-windows-sentinel-bundle.zip`.

3 Install the Linux or Windows software on all guest VMs that require migration.

   ◆ HCX Sentinel installation for Linux

   a Connect to your guest system using SSH.

   b Copy the linux-sentinel-installer.sh file to the guest system.

   c At the terminal, enter the command `bash linux-sentinel-installer.sh`.

      The software prompts you for permission to start the installation.

   d Enter **yes**, and press **Enter**.

   ◆ HCX Sentinel installation for Windows

   a Log in to the guest system.

   b Copy the `windows-sentinel-bundle.zip` file to the guest system.

   c Unzip the bundle.

   d To run the installer, double-click **install-sentinel.exe**.

   e Click **Next** to continue.

   f Accept the license agreement and click **Next** to continue.

   g Choose the location where you to install the software and click **Next**.

   h Click **Finish**.

Uninstalling HCX Sentinel Agent Software

The HCX OS Assisted Migration (OSAM) service automatically uninstalls the Sentinel software from the guest system after a successful migration. Alternatively, you can manually remove the software using the Sentinel Management interface.

Following a successful migration, the OSAM service automatically sends instructions to the guest virtual machine to power off and uninstall the Sentinel agent software upon reboot. The OSAM service then removes the VM from the inventory of non-vSphere virtual machines on the HCX.

You can manually uninstall the software from a source VM using the **Uninstall** button in the Sentinel Management interface. The action taken by the OSAM service to uninstall the Sentinel software depends on whether the service has access to the source system:

- OSAM service has a connection to the source VM—OSAM uninstalls the Sentinel software from the source VM. Also, OSAM removes the source VM from the inventory of non-vSphere virtual machines on the HCX.
- OSAM service has no connection to the source VM—OSAM removes the source VM from the inventory of non-vSphere virtual machines on the HCX, but the Sentinel software remains installed on the source VM. In this case, if a connection to the source VM is reestablished with the OSAM service, the source VM reappears in the inventory of non-vSphere virtual machines on the HCX. To remove the Sentinel agent software and delete the source VM from the inventory, repeat the uninstall the procedure.

Note The OSAM service prevents you from uninstalling the HCX Sentinel software during the source VM migration.

To uninstall the Sentinel agent software manually, use the following procedure .

Procedure

1. Go to **Interconnect > Multi-Site Service Mesh > Sentinel Management**.

   The system displays the list of source VMs installed with the Sentinel agent software.

2. Select the source systems.
3. Click **Uninstall**.

   The system prompts you to verify the action.

4. Click **Yes**.

   The OSAM service begins the process of uninstalling the software from the source VM.

5. In the Sentinel Management interface, verify that the entry is removed from the inventory of non-vSphere virtual machines.

Upgrading HCX Sentinel Agent Software

To maintain compatibility with OS Assisted Migration (OSAM) service appliances, update the HCX Sentinel agent software on guest virtual machines.

The OS Assisted Migration (OSAM) upgrade bundle includes the HCX Sentinel agent, Sentinel Gateway (SGW) appliance, and Sentinel Data Receiver (SDR) appliance software. This software is downloaded to the HCX only after you upgrade the HCX Manager. This means you must upgrade to the latest HCX software to get the latest OSAM updates.

The OSAM upgrade bundle has two versions depending on the HCX deployment: On-premise or Cloud. The Sentinel agent software is only downloaded with on-premise HCX deployment upgrades.

Prerequisites

The Sentinel Gateway and Sentinel Data Receiver appliances are updated to the latest version as described in Upgrade the HCX Service Mesh Appliances.

Procedure

1 Navigate to the HCX Dashboard, and select **Interconnect > Multi-Site Service Mesh > Sentinel Management**.

   The system displays the inventory of guest virtual machines. Each entry lists the current Sentinel software version and the available software version installed after upgrading the HCX Manager.

2 Select the guest VMs to update:

   It is a best practice to update all guest virtual machines to the same version at the same time.

   - To update all guest VMs, check the box at the top of the **Hostname** column.
   - To update individual VMs, check the box next to each VM.

3 Click **Upgrade**.

   The upgrade begins for the selected guest VMs. For the upgrade status, review **Task Details**.

   Note Sentinel upgrade is allowed only when the migration is not in the switchover phase for that guest virtual machine. An attempt to upgrade Sentinel on a guest virtual machine that has a switchover in progress results in an upgrade request failure.

Extending Networks 7

Using the VMware HCX Network Extension service, you can create Layer 2 networks at the destination HCX site and bridge the remote network to the source network over a multi-gigabit-capable link. The new stretched network is automatically bridged with the network at the source HCX data center.

Read the following topics next:

- About Network Extension
- Extending Networks Using VMware HCX
- Understanding Network Extension with Mobility Optimized Networking
- Viewing Network Extension Details
- In-Service Upgrade for Network Extension Appliances
- Removing a Network Extension
- Understanding Network Extension High Availability

About Network Extension

You can bridge networks between HCX Manager activated data centers with HCX Network Extension.

With VMware HCX Network Extension (HCX-NE), you can extend the virtual machine (VM) networks to an HCX remote site. Virtual machines that are migrated or created on the extended segment at the remote site behave as if on the same L2 segment as virtual machines in the source environment. With Network Extension, the default gateway for the extended network is only connected at the source site. Traffic from virtual machines in remote sites that must be routed to a different L3 network flow through the source site gateway.

Using Network Extension with VMware HCX Migration you can:

- Retain the IP and the MAC addresses of the VM, and honor existing network policies.
- Extend VLAN-tagged networks from a VMware vSphere Distributed Switch.
- Extend NSX segments.

Requirements for Network Extension

HCX supports extending networks from VMware vSphere Distributed Switches, NSX segments, and NSXv logical switches.

The following information and requirements apply when extending networks:

- General requirements:
  - Never extend the networks used to create the network profiles.
  - Never use HCX to extend the vSphere Management network or other VMkernel networks (for example: vMotion, vSAN, replication) to the remote site.
- Networks with NSX-T Data Center at the source:

  Note Registering NSX-T in an HCX Connector is optional, except when extending NSX networks or migrating NSX Security Tags. If NSX-T is registered, the following requirements apply:

  - The NSX-T Manager must be registered during the HCX Manager deployment.
  - The NSX-T Manager must be Version 2.4 or higher.
  - NSX-T Overlay or VLAN Transport Zones must be configured in the vCenter Server where the network originates.
    - The ESX hosts where the NSX-T segment originates must be configured as NSX-T transport nodes.
    - NSX-T Overlay and NSX-T VLAN networks can be extended.
- Networks with NSX-T Data Center at the destination:
  - HCX connects to Tier-1 Gateways and Segments created in the Networking tab (policy UI). NSX configurations created in the manager tab cannot be used with HCX Network Extension and migration operations.
  - Additional network extension appliances are required when extending more than 8 networks.
  - NSX-T Overlay must be configured in the destination vCenter Server or SDDC.
  - NSX-T Overlay and NSX-T VLAN networks are always created as NSX Overlay networks at the destination. HCX Network Extension is always to NSX Overlay networks. Pre-created segments in the vSphere Distributed Switch, CVDS, or VLAN transport zone, cannot be used as a destination.

- Networks with NSX-T Data Center at the source and the destination:
  - HCX extensions must be connected to Tier-1 Gateways and NSX-T segments that are created with the NSX-T Policy UI/API.

  Note In addition to these requirements for both sites, individual requirements for networks with NSX-T at the source and the individual requirements for networks with NSX-T at the destination also apply.

- HCX does not support networks with NSX for vSphere at the destination.
- NSX for vSphere Logical Switches at the source.
  - The NSX-V Manager must be registered during the HCX Manager deployment.
  - The NSX-V Manager must be Version 6.4.0 or higher.

Restrictions and Limitations for Network Extension

HCX Network Extension is allowed or prevented under certain conditions.

Detected and Restricted Source Network Types

The HCX Network Extension service detects and prevents several non-supported Network Extension scenarios (items are dimmed in the Network Extension UI):

- vSphere infrastructure networks (ESXi VMkernel networks).
- HCX Network Profile networks (Distributed Port Groups or Segments selected in a Network Profile).
- Untagged Distributed Port Groups (Distributed Port Groups with VLAN type None, ID 0 or NULL).
- Private VLAN (PVLAN) networks.
- vCenter Server backed Port Groups configured with ephemeral binding cannot be extended.
- NSX-T logical switches.

Unsupported Source Configurations

HCX Network Extension does not support the following source configurations:

- vSphere Standard Switch (vSS) networks.
- Cisco Nexus 1000v or other third-party switches.
- Cisco Application Centric Infrastructure (ACI) with VMware Virtual Machine Monitor (VMM).
- Virtual machine networks must only be extended with a single solution. HCX does not support Network Extension for networks already extended to the same NSX router by an external solution. For example, either HCX Network Extension or NSX L2 VPN can be used to provide connectivity, but both must not be used simultaneously. Using multiple bridging solutions simultaneously can result in a network outage.

- Virtual machine networks with shared or overlapping VLAN configurations must not be extended to the same destination router. This can result in a network outage.
- Secondary subnets in a single layer-2 network.
- NSX-T Global Federation configurations.

  HCX does not integrate with the NSX Global Manager for extending networks (only the NSX Local Manager).

Unsupported Destination Configurations

HCX Network Extension does not support the following destination configurations:

- NSX-T Global Federation configurations.

  HCX does not integrate with the NSX Global Manager for extending networks (only the NSX Local Manager).
- NSX-V at the cloud is unsupported for HCX Cloud Manager.
- VMware Cloud Director enabled with data center group networking backed by NSX.

Additional Considerations

- HCX supports extending the same network to a maximum of 3 distinct destinations or routers.
- One Network Extension configuration cannot be extended multiple times to the same destination router.
- Daisy-chain "L" network extension (extending extensions) is only supported to one additional environment in the same data center, public cloud provider, and region.
  - Daisy-chain extension is not supported with source networks based on NSX distributed routing.
  - Daisy-chain extension can lower end-to-end network performance due to the combined latency and additional layer of packet and encryption processing.
- One Network Extension appliance can only connect to one Distributed Virtual Switch or NSX Transport Zone.
- Networks can only be extended between one appliance pair (source and destination appliances) per site, and multiple network appliances cannot be used to increase throughput.
- Network Extension does not detect or mitigate loops.
- Virtual machine networks that span more than one vCenter Server must not be extended from more than one vCenter to the same destination router. This can result in a network outage.
- Network Extension does not detect or mitigate IP conflicts.
- Network Extension does not detect or mitigate MAC conflicts.

- For a cloud/site pair, a given network can be extended through only one appliance and is subject to the resource and the performance limitations of that appliance.
- Network Extension connects to an existing segment on the target site if it has the same gateway IP and Prefix configured for the extension, and it disconnects the NSX router interface from the network. If the NSX tier-1 router interface was previously connected and in service, all communication to the gateway on that cloud network is disrupted.
- NSX-T Overlay and NSX-T VLAN networks can be extended. These extensions are always created as NSX Overlay networks at the destination. HCX Network Extension is always to NSX Overlay networks. Pre-created segments in the vSphere Distributed Switch, CVDS, or VLAN transport zone, cannot be used as a destination.
- VMware NSX Traceflow does not work with extended networks.
- During disaster recovery (DR) events when the onpremises site is not available or powered off, the existing extended networks are required to be unextended from target/Cloud Network Extension (NE) wizard to ensure Cloud NSX DLR is connected for all extended segments where migrated cloud VMs are hosted.

  Follow the below practices to avoid downtime as much as possible:

  - Once the onpremise HCX Manager/NE Appliances are powered off or no longer accessible during the DR event, refer to the destination HCX Cloud Manager dashboard and verify if the site pair status is showing Disconnected (connection from the remote site may be down).
  - Go to destination HCX NE wizard, and perform “Force Unextend Network” for segments one at a time.

    Note Do not trigger the standard unextension workflow, which is designed to be functional when both sites are Up and running.

    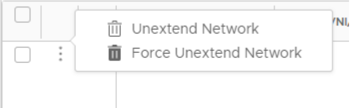

  Important During DR events, the impact will be low if existing extended segments already have the HCX MON (Mobility Optimized Networking) feature enabled, which makes Cloud NSX DLR to be CONNECTED for extended segments and helps provide E-W routing between two cloud VMs (depending upon user configured policy-routes). Refer to Knowledge Base Article 83375 for more info.

  Note For VMware Cloud on AWS specific deployments, MON may not help to optimize traffic between extended segments that are not directly connected to the same Tier-1 router.

Network Extension to Destinations with Universal Distributed Logical Routers

When working with destination environments with Cross-vCenter NSX configurations, HCX supports extending source networks to destination environments using a Universal Distributed Logical Router (UDLR). When a UDLR is selected during the network extension operation, HCX creates a Universal Logical Switch on the destination across multiple vCenter Servers.

The following information and requirements apply for Network Extension when specifying a UDLR as the gateway.

- The HCX Cloud Manager configuration includes all secondary NSX systems that exist in **Configuration** > **NSX Manager** at time of installation.
- Secondary NSX systems must have the administrative credentials of its associated vCenter Server.
- HCX Network Extension does not support the local egress feature of UDLRs.

Extending Networks Using VMware HCX

VMware HCX Network Extension is a layer-two bridging function initiated at the source site.

If you are using the HCX Manager UI (standalone or vSphere Client plug-in), you can extend networks by selecting one or more Distributed Port Groups or NSX segments. When you extend a network, a corresponding NSX segment is created at the destination site.

If you are using the vSphere Client Networking interface, you can select a Distributed Port Group and extend it.

Note For a list of restrictions regarding Network Extension, see Restrictions and Limitations for Network Extension.

For the operational limits supported with HCX Network Extension, see Configuration and Service Limits.

Procedure

1 If you are using the HCX Manager UI, follow these steps to select a network for extension:

   a In the HCX Services menu, select **Network Extension**.

   A summary screen appears displaying all configured site pairs. Expand a site pair to see the associated Service Mesh information. Expand a Service Mesh to see the associated Network Extensions.

   b At the top of the page, select **Extend Networks**.

   A screen appears prompting you for the target site network selections.

   c Select a Service Mesh.

   Note If you have only one Service Mesh, it is selected by default.

d Select one or more Distributed Port Groups or NSX Logical Switches.

You can use the available filters to hide networks that are ineligible for extension, hide networks that do not have virtual machines associated with them, or hide networks without extension.

e Click **Next**.

2 If you are using the vSphere Client Networking interface, follow these steps to select a network for extension:

a From the vSphere menu, select **Networking**.

b Right-click a Distributed Port Group.

c Locate HCX Actions near the bottom of the list, and select **Extend Network to HCX Target Site**.

A screen appears prompting you for the target site network selections.

d Expand Remote Site Connection and select a site.

Note If you have only one site pairing, it is selected by default.

3 Use the drop-down menu to select the Destination First Hop Router.

4 Provide the Gateway IP address and Prefix Length for the network being extended in the format <gateway IP/Prefix Length>. For example: 192.168.10.1/24.

For a VMware Cloud Director target cloud, click the extended option drop-down menu and optionally specify the DNS configuration.

5 Select the Extension appliance.

Select a Network Extension appliance or a Network Extension High Availability group.

6 (Optional) For each source network, expand **Settings - optional** and select the appropriate options:

Allow Overlapping VLAN

The HCX Manager prevents you from extending networks that have the same VLAN ID and Gateway IP address. Select this option to override system and allow duplicate VLAN IDs.

DNS entries

For a VMware Cloud Director target cloud, optionally specify the DNS configuration: **Primary DNS**, **Secondary DNS**, and **DNS Suffix**.

7 (Optional) Depending on the NSX version running in your data center, select Mobility Optimized Networking for all workloads that require routing through the local gateway at the target site.

For more information about MON and additional configuration settings, see Understanding Network Extension with Mobility Optimized Networking.

8 To finish, click **Submit**.

To view the task status, navigate to the HCX Dashboard and scroll down to the Activity Logs display.

Note If any conflicting operations on the same appliances are running, then network extension requests are queued. Examples of conflicting operations include Service Mesh edits or resyncs, appliance operations such as redeploy or force resync, and network extension using the same high availability (HA) group.

Results

Return to the Network Extension page for a summary of information that includes the number of extended networks, extension appliances, mobility optimized networks, and network extension served by an HA group. The HA group name is used in place of appliance name with an HA tag to indicate the HA activation.

Understanding Network Extension with Mobility Optimized Networking

Mobility Optimized Networking (MON) is feature of the HCX Network Extension Service. 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 through 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.

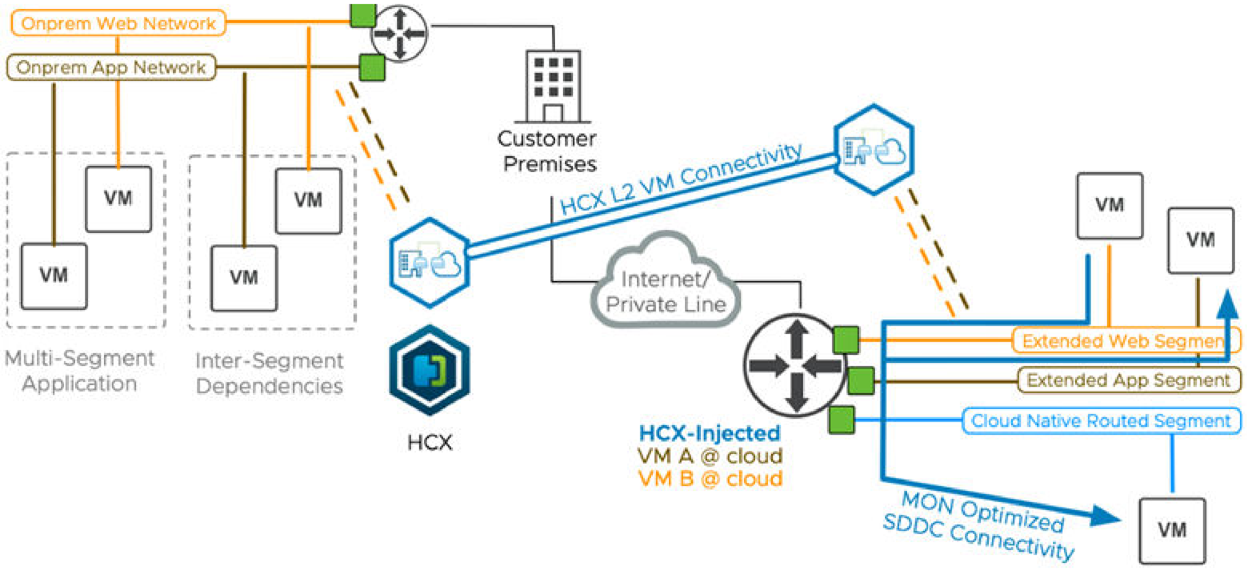

About Mobility Optimized Networking

This section provides an overview of workload traffic flows using HCX Network Extension with and without Mobility Optimized Networking.

Use Cases for Mobility Optimized Networking

MON improves network performance and reduces latency for virtual machines that have been migrated to the cloud on an extended L2 segment. MON provides these improvements by allowing more granular control of routing to and from those virtual machines in the cloud.

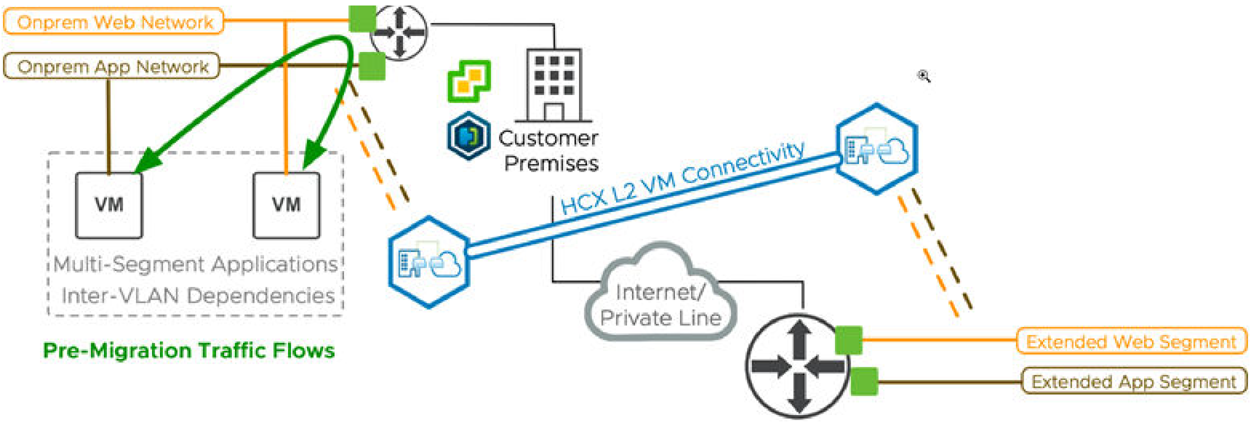

Without MON, HCX Network Extension expands the on-premises layer-2 network to the cloud SDDC while the default gateway remains at the source. The network tromboning effect is observed when virtual machines in the destination connected to different extended segments communicate.

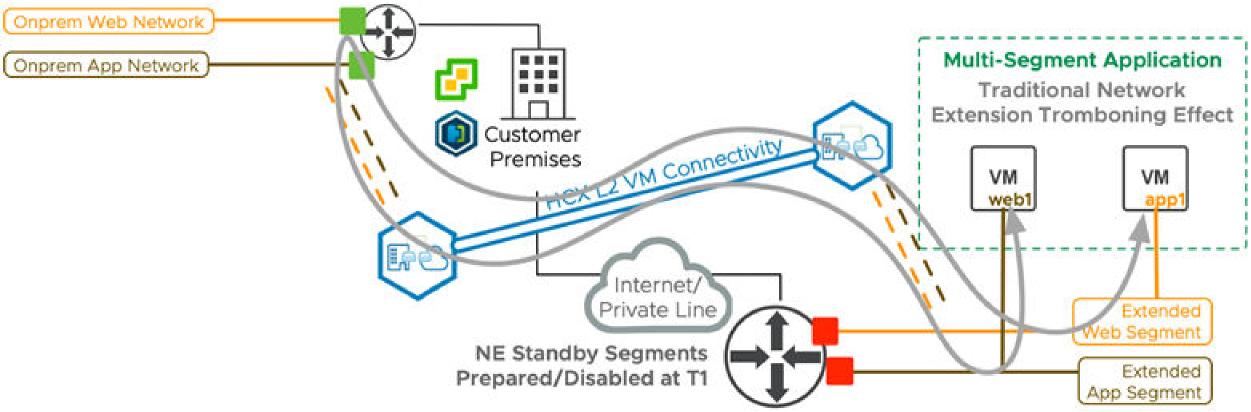

MON enables migrated virtual machines to reach segments within the SDDC without sending packets back to the source environment router.

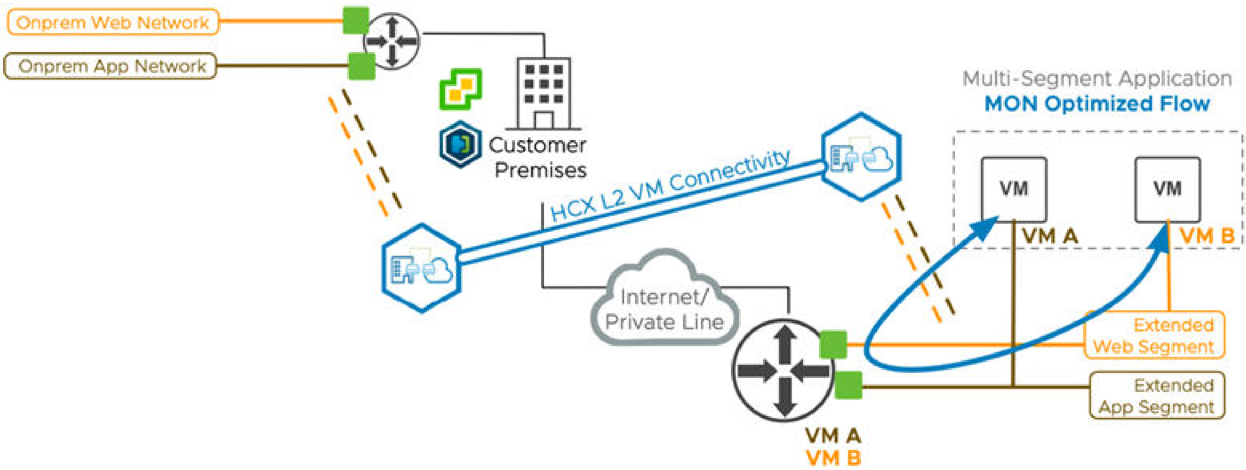

MON can be configured to allow migrated virtual machines to reach services hosted within a public cloud.

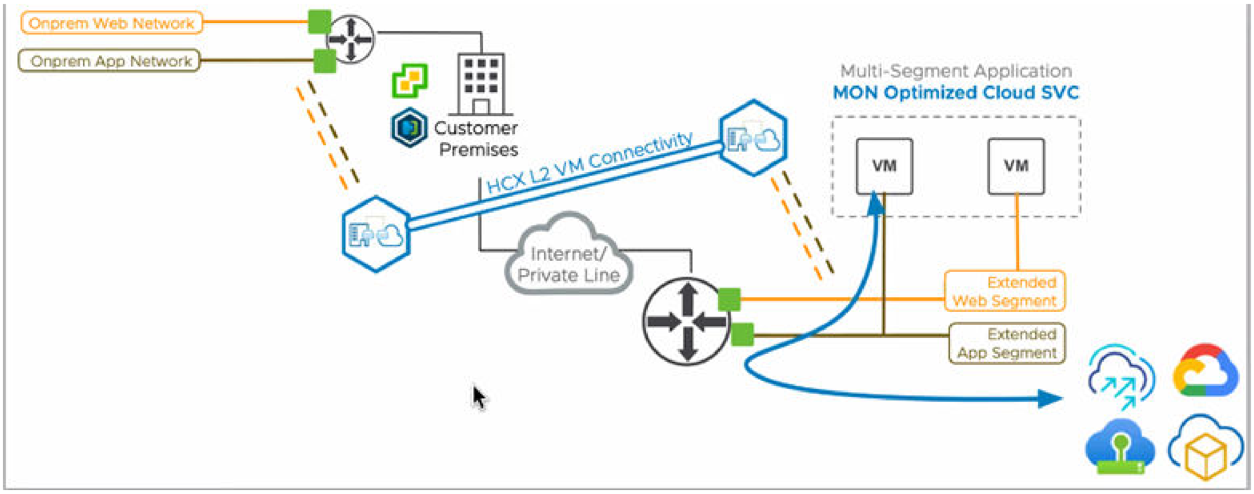

MON enables migrated virtual machines to use the SDDC Internet interface (with SNAT).

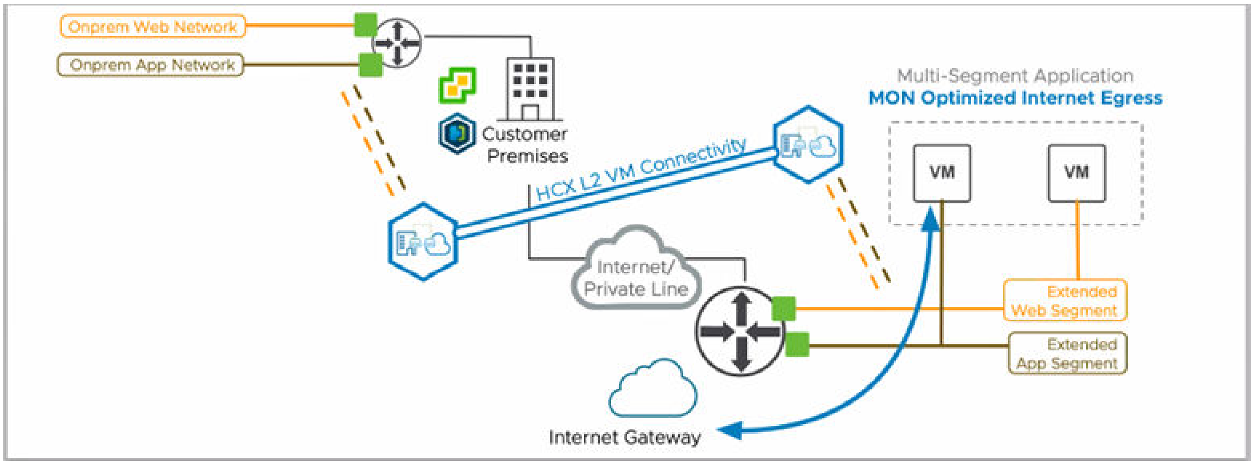

Mobility Optimized Networking Outcomes by Migration Type

- HCX Bulk-migrated virtual machines are automatically MON-enabled in the SDDC.
- HCX vMotion- and RAV-migrated virtual machines use the on-premises gateway until they are specifically configured to use the cloud gateway in the HCX UI/API.
- Virtual machines attached to the segment prior to enabling MON use the on-premises gateway until they are specifically configured to use the cloud gateway in the MON interface.

Mobility Optimized Networking Operation

Network Extension with Mobility Optimized Networking provides the following functionality:

- Enable or deactivate MON at the time of stretching a network.
- Enable or deactivate MON for already extended networks.
- Enable or deactivate MON on an individual VM basis for VMs residing on extended networks in the SDDC.
- Display which VMs are using MON.
- When using HCX to vMotion a VM, preserve existing network connections while providing the option to activate Mobility Optimized Networking on that VM after migration.
- Configure MON Route Policy to define on-premises (non-SDDC) subnets or exception/deny subnets for local egress.

The following process explains what happens during the various phases of Mobility Optimized Networking.

1. Mobility Optimized Networking is enabled for an HCX extended segment.

   HCX enables the network ID (gateway IP) in the SDDC Compute Gateway. It is enabled with a limited /32 255.255.255.255 network mask.

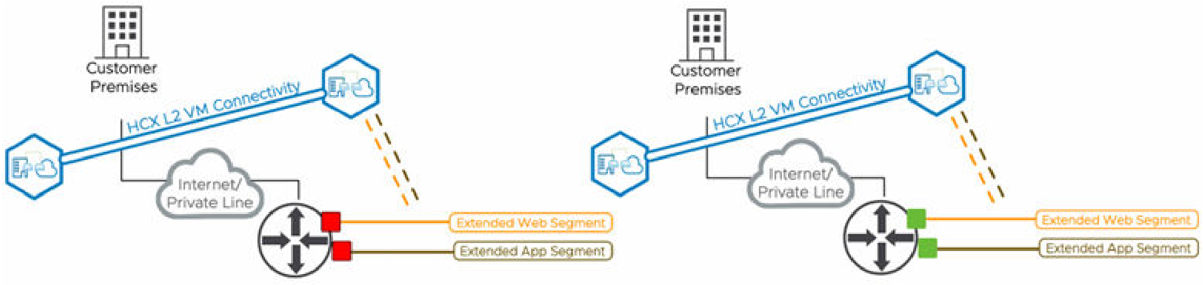

2. Static routes are added in the SDDC Compute Gateway for migrated virtual machines on HCX extended network.

   HCX adds reachability information for the migrated virtual machine (in the form of a virtual machine specific static route) to the SDDC Compute Gateway, allowing reachability within the SDDC. This VM static route is not advertised to the on-premises environment. The HCX L2 path is used to reach subnets not in the SDDC.

3 Using SDDC forwarding technology, the virtual machine uses the SDDC Compute Gateway to reach the SDDC networks.

For reachability outside of the SDDC tier-1, the MON policy configuration is evaluated according to the MON policy configuration. Matching subnets are sent to the original premises router. Nonmatched subnets are sent to the SDDC tier-0 router. For more information on MON policy routes, see Mobility Optimized Networking Policy Routes.

Mobility Optimized Networking Policy Routes

When the destination network for a traffic flow is not within the SDDC, the Mobility Optimized Networking policy is evaluated.

MON policy routes define which traffic is routed through the source gateway versus traffic that is routed through the cloud gateway. The Advanced tab in the HCX Network Extension interface provides an option for configuring policy routes.

When the destination network for a traffic flow is not within the SDDC tier-1 router, the MON policy is evaluated:

- If the destination IP is matched and configured as allow in the MON policy configuration, the packet is forwarded to the premises gateway using the HCX Network Extension appliance.
- If the destination IP is not matched, or configured to deny in the MON policy, the packet is forwarded to the SDDC Tier-0 to be routed.

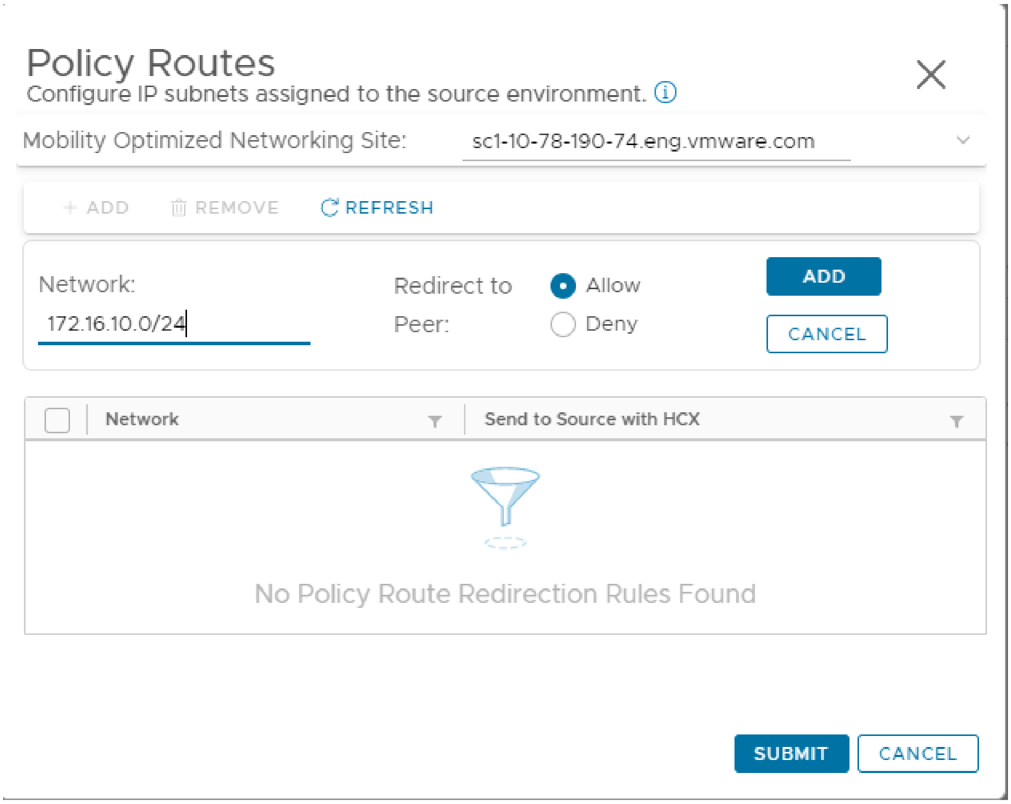

Policy Routes

Configure IP subnets assigned to the source environment.

Mobility Optimized Networking Site: sc1-10-78-190-74.eng.vmware.com

Network: 172.16.10.0/24

Redirect to Peer: ● Allow ○ Deny

□	Network	Send to Source with HCX

No Policy Route Redirection Rules Found

Example Policy Route Configurations

Mobility Optimized Networking policy route configuration settings can vary depending on the HCX deployment.

Important The examples in this section are generalized approaches that might not be suitable for all deployments. The policy route configuration defines how routed traffic is forwarded for MON enabled virtual machines. The configurations should be well understood in the context of the site to site routing design. Incorrect configurations can result in disrupted traffic for the MON enabled virtual machines.

Default MON Policy Configuration

The default MON policy includes all RFC-1918 networks. This policy configuration forwards private subnet traffic (not destined to segments within the SDDC) to the on-premises router and sends internet egress traffic to the SDDC tier-0 router.

Policy Configuration for Internet Egress On-premises

For MON deployments where security policies require internet access on-premises, replace the default MON Policy Configuration:

- Remove the default RFC-1918 entries from the Policy Routes interface.
- Add a single Allow entry for network 0.0.0.0/0.

  This policy configuration forwards private subnet traffic (not destined to segments within the SDDC) to the on-premises router and internet egress traffic, while maintaining routing symmetry.

Policy Configuration for Cloud Services

MON policy routing can be revised to achieve cloud service reachability.

- Configure the IP address ranges for the cloud based service as Deny entries (exclusions) to the MON Policy.
- Deny entries are sent to the SDDC tier-1 router.

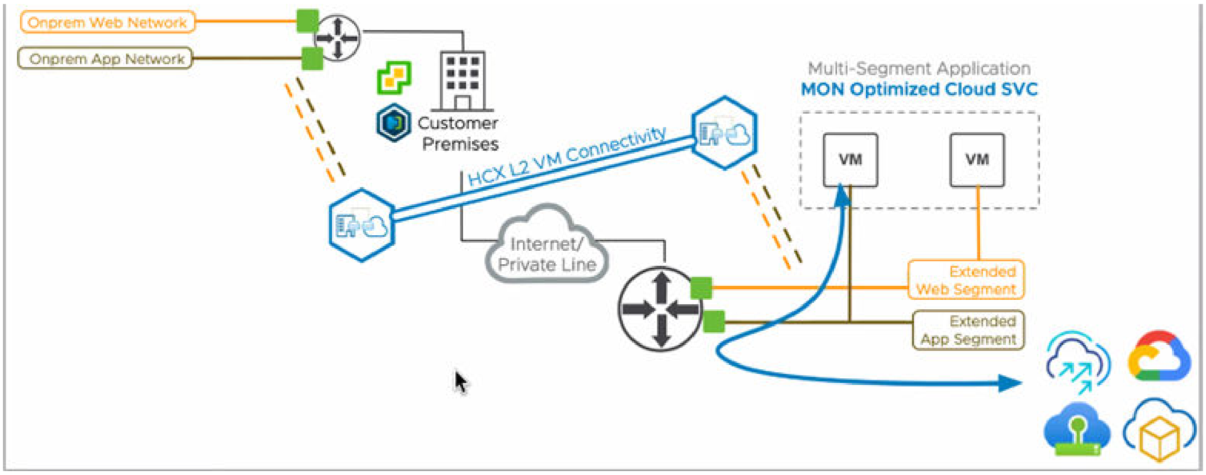

Requirements and Limitations for Mobility Optimized Networking Topologies

The HCX Mobility Optimized Networking (MON) feature routes network traffic based on locality of the source and destination virtual machines. MON operation requires the specific configuration of the Network Extension parameters and the network environment between the source and the destination sites.

General Requirements for MON

To get started, MON requires an HCX Enterprise license or activation.

Requirements for MON-Enabled Virtual Machines

- VMware Tools must be installed.

Environmental Requirements for MON

- Extended networks meet all NSX requirements. See Requirements for Network Extension
- An NSX Tier-1 router connected to a Tier-0 router must exist in the cloud environment prior to enabling this feature.

  Note Layer-2 network extensions can only terminate on a Tier-1 router.
- The default Gateway IP and the DHCP Server IP cannot be the same. The Gateway for the extended segment can provide DHCP services but must have a different IP address used for the DHCP server.

Limitations of MON

- MON does not provide traffic optimization from MON-enabled virtual machines to virtual machines on other extended networks without MON.
- MON provides optimization in the Cloud. MON does not provide optimization in the source.
- For additional information related to running MON on VMware Cloud on AWS, see HCX Mobility Optimized Networking for VMware Cloud on AWS.
- When a virtual machine with multiple NICs is connected to multiple MON enabled extended networks, and Target Router Location is changed, that change applies to all the interfaces.

Configuring Mobility Optimized Networking

For data centers using NSX-T, configure Mobility Optimized Networking for workloads that require routing through the local gateway.

Configuring MON for an extended network is available as one of the Advanced options when creating the extended network.

In some cases, you might not configure MON for an extended network. In other cases, the virtual machines that are members of the extended network are migrated using vMotion or Replication Assisted vMotion, which requires setting MON manually for those VMs.

Use the following procedure to configure MON for an extended network and for virtual machine workloads.

Prerequisites

The destination environment meets the requirements described in Requirements and Limitations for Mobility Optimized Networking Topologies.

Procedure

1 In the HCX Manager UI, navigate to **Services > Network Extension**.

2 To see the extended networks in the Network Extension screen, expand a site pair.

The system highlights Network Extensions activated for MON with an icon.

3 To display network details, expand each extension.

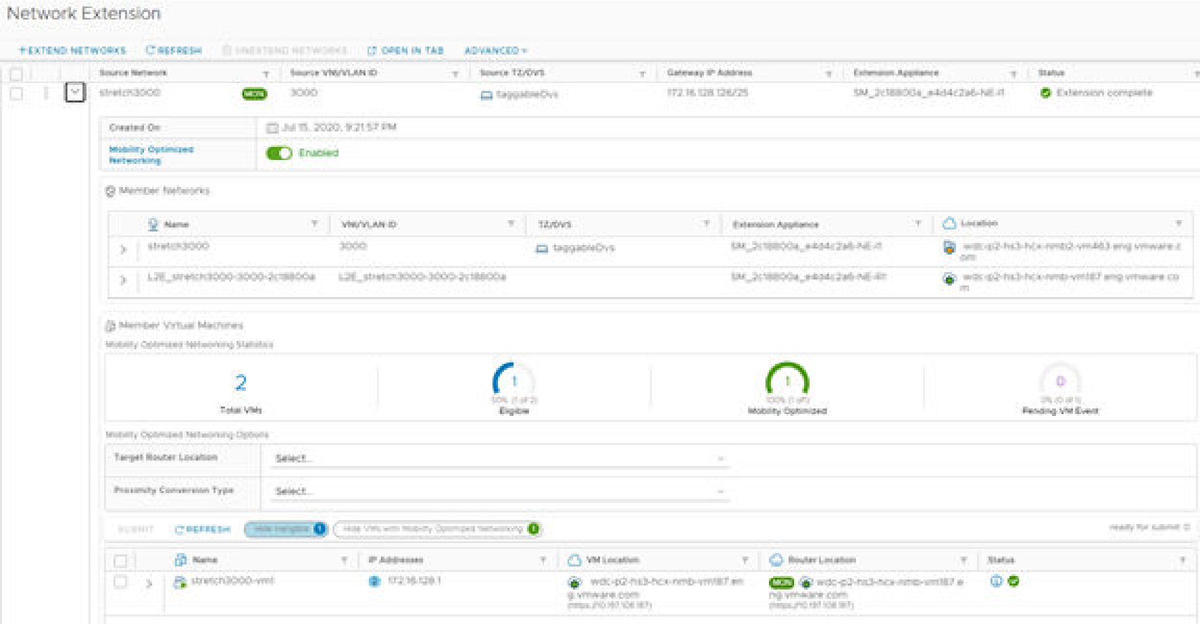

Network Extension

+EXTEND NETWORKS | REFRESH | UNEXTEND NETWORKS | OPEN IN TAB | ADVANCED

Source Network	Source VNI/VLAN ID	Source TZ/DVS	Gateway IP Address	Extension Appliance	Status
stretch3000 MON	3000	taggableDvs	172.16.128.126/25	SM_2c18800a_e4d4c2a6-NE-I1	Extension complete

Created On	Jul 15, 2020, 9:21:57 PM
Mobility Optimized Networking	Enabled

Member Networks

Name	VNI/VLAN ID	TZ/DVS	Extension Appliance	Location
stretch3000	3000	taggableDvs	SM_2c18800a_e4d4c2a6-NE-I1	wdc-p2-hs3-hcx-nmb2-vm463.eng.vmware.com
L2E_stretch3000-3000-2c18800a	L2E_stretch3000-3000-2c18800a		SM_2c18800a_e4d4c2a6-NE-R1	wdc-p2-hs3-hcx-nmb-vm187.eng.vmware.com

Member Virtual Machines

Mobility Optimized Networking Statistics

Total VMs	Eligible	Mobility Optimized	Pending VM Event
2	1 50% (1 of 2)	1 100% (1 of 1)	0 0% (0 of 1)

Mobility Optimized Networking Options

Target Router Location	Select...
Proximity Conversion Type	Select...

SUBMIT | REFRESH | Hide Ineligible | Hide VMs with Mobility Optimized Networking | Ready for submit: 0

Name	IP Addresses	VM Location	Router Location	Status
stretch3000-vm1	172.16.128.1	wdc-p2-hs3-hcx-nmb-vm187.eng.vmware.com (https://10.197.108.187)	MON wdc-p2-hs3-hcx-nmb-vm187.eng.vmware.com (https://10.197.108.187)	✓

4 If MON is not enabled on the Network Extension, set the slider for Mobility Optimized Networking, and click **Enable**.

   HCX updates the Network Extension screen, displaying the member virtual machines and the **Router Location** selection drop-down menu. The **Router Location** identifies the target Tier-1 cloud router for the member virtual machines.

   Note The Bulk migration and the DR recovery operations automatically set the **Target Router Location** as the cloud Tier-1 router. Virtual machines migrated using HCX vMotion and Replication Assisted vMotion (RAV), or VMs attached to the segment prior to enabling MON, continue to use the on-premises router until the **Target Router Location** is set manually in the next step.

5 To enable MON optimization on VMs that have been migrated using vMotion or RAV, or VMs attached to the segment prior to enabling MON, set the **Target Router Location** for each VM.

   a Select a VM and expand the row.

      You can select multiple VMs using the check box next to each workload.

   b Set the **Target Router Location** for each VM by selecting the destination (**Cloud**) option from the drop-down menu.

   c Click **Submit**.

   HCX configures all selected VM workloads for Mobility Optimized Networking.

   Note To deactivate MON optimization for a VM, select **Target Router Location** and select the source (**On-premises**) option from the drop-down menu.

What to do next

If required, configure the policy routes. See Configuring Policy Routes for Mobility Optimized Networking.

Configuring Policy Routes for Mobility Optimized Networking

With Mobility Optimized Networking, you have the option to control which traffic is routed locally using the cloud gateway versus traffic that goes out through the source gateway. Policy routes define which traffic is routed through the source gateway. All other traffic is routed through the cloud gateway.

Procedure

1 In the HCX Manager UI, navigate to **Network Extension**.

2 In the Network Extension screen, click the **Advanced** tab.

3 Click **Policy Routes**.

4 Using the pull-down menu, select a destination site.

5 In the Network field, for which you want traffic routed through the source gate, click **Add**.

6 Complete the entries for **Network IP Address** and **Prefix Length**.

7 Set **Redirect to Peer** to **Allow** or **Deny**.

8 Click **Add**.

The policy is applied to the MON extended network.

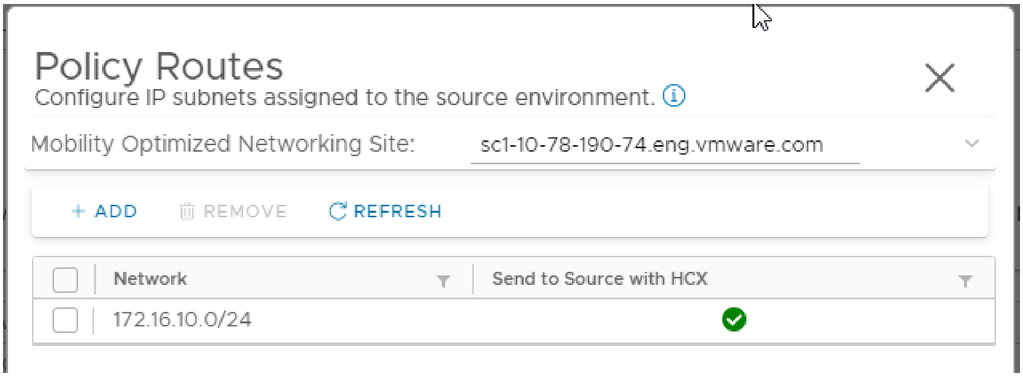

Policy Routes

Configure IP subnets assigned to the source environment.

Mobility Optimized Networking Site: sc1-10-78-190-74.eng.vmware.com

□	Network	Send to Source with HCX
□	172.16.10.0/24	✓

Viewing Network Extension Details

HCX provides detailed tunnel state information for the Network Extension appliance, and connection information for each extended network associated with that appliance.

HCX maintains Up or Down state information regarding the tunnel used for Network Extension. The information includes the tunnel ID, local IP address and port number, remote IP address and port number, and tunnel status.

Connection statistics for extended networks includes the bit rate, bytes transferred and received, packet rate, and packets transferred and received. The statistics are updated every 1 minute and stored in the HCX database. The bytes and packets transferred and received information reflects the total number since the Network Extension appliance was powered on.

Network Extension statistics also detect and display the MAC address of each virtual machine NIC on the extended network, which can be helpful in determining the status of a particular virtual machine on that network. A search option is provided to filter the list of addresses.

Procedure

1 From the HCX Dashboard, go to **Infrastructure > Interconnect > Service Mesh**.

2 Click **View Appliances**.

HCX displays a list appliances that have been activated in the Service Mesh.

3 Expand the Network Extension appliance.

The system displays options for selecting **Tunnel Details** or **Network Extension Details**. **Tunnel Details** is selected by default.

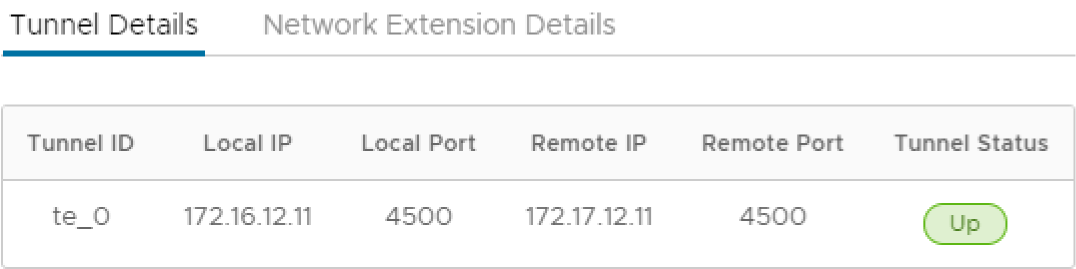

Tunnel ID	Local IP	Local Port	Remote IP	Remote Port	Tunnel Status
te_0	172.16.12.11	4500	172.17.12.11	4500	Up

4 To view details regarding extended networks, click **Network Extension Details**.

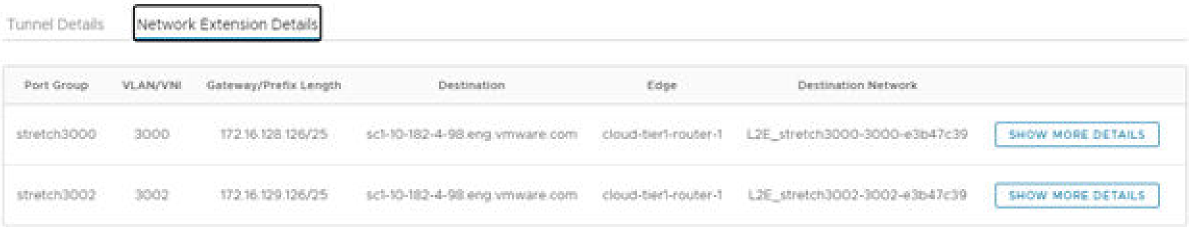

Port Group	VLAN/VNI	Gateway/Prefix Length	Destination	Edge	Destination Network	
stretch3000	3000	172.16.128.126/25	sc1-10-182-4-98.eng.vmware.com	cloud-tier1-router-1	L2E_stretch3000-3000-e3b47c39	SHOW MORE DETAILS
stretch3002	3002	172.16.129.126/25	sc1-10-182-4-98.eng.vmware.com	cloud-tier1-router-1	L2E_stretch3002-3002-e3b47c39	SHOW MORE DETAILS

5 To view the connection statics information for a specific network, click **Show More Details**.

HCX displays the connection statistics information. To see updated information, click the refresh icon.

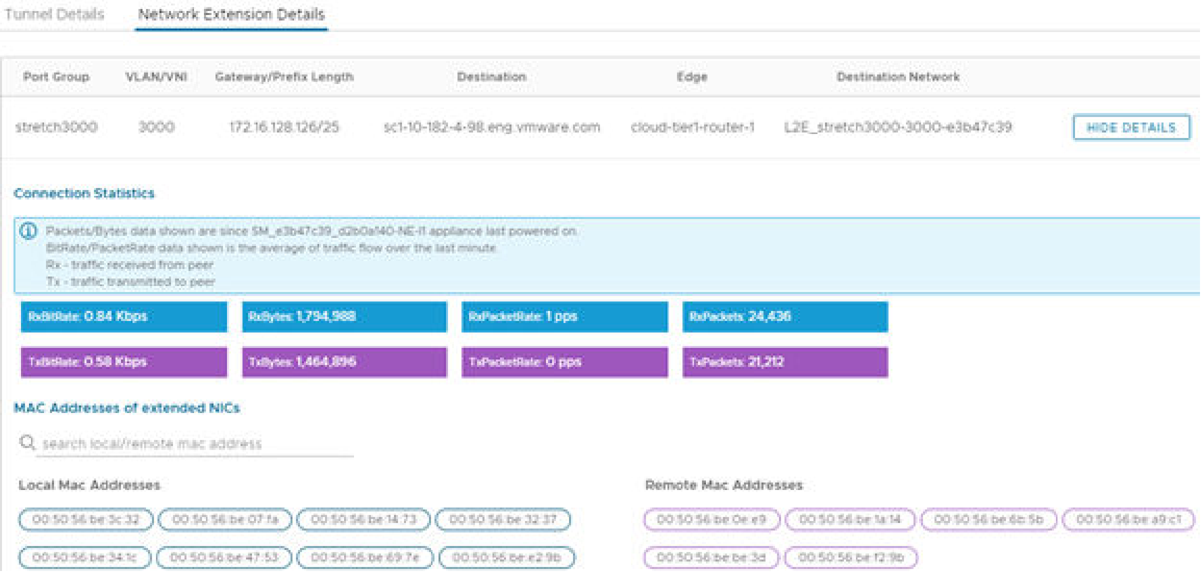

Port Group	VLAN/VNI	Gateway/Prefix Length	Destination	Edge	Destination Network	
stretch3000	3000	172.16.128.126/25	sc1-10-182-4-98.eng.vmware.com	cloud-tier1-router-1	L2E_stretch3000-3000-e3b47c39	HIDE DETAILS

Connection Statistics

ⓘ Packets/Bytes data shown are since SM_e3b47c39_d2b0a140-NE-I1 appliance last powered on.
BitRate/PacketRate data shown is the average of traffic flow over the last minute.
Rx - traffic received from peer
Tx - traffic transmitted to peer

RxBitRate: 0.84 Kbps
RxBytes: 1,794,988
RxPacketRate: 1 pps
RxPackets: 24,436
TxBitRate: 0.58 Kbps
TxBytes: 1,464,896
TxPacketRate: 0 pps
TxPackets: 21,212

MAC Addresses of extended NICs

search local/remote mac address

Local Mac Addresses	Remote Mac Addresses
00:50:56:be:3c:32	00:50:56:be:0e:e9
00:50:56:be:07:fa	00:50:56:be:1a:14
00:50:56:be:14:73	00:50:56:be:6b:5b
00:50:56:be:32:37	00:50:56:be:a9:c1
00:50:56:be:34:1c	00:50:56:be:be:3d
00:50:56:be:47:53	00:50:56:be:f2:9b
00:50:56:be:69:7e	
00:50:56:be:e2:9b	

6 To close the display, click **Hide Details**.

What to do next

Log in to the peer HCX site to view connection statistics information from that site.

In-Service Upgrade for Network Extension Appliances

HCX provides options for Network Extension upgrade or redeployment that help to minimize service downtime and disruptions to on-going L2 traffic.

The Network Extension appliance is a critical component of many HCX deployments, not only during migration but post migration as extended networks continue to be used after migration in a hybrid environment. HCX operations using extended networks can be impacted during Network Extension appliance upgrades or when a change to the HCX Compute Profile or Service Mesh requires redeploying the Network Extension appliance.

Network Extension appliances are available for **In-Service** or **Standard** (default) upgrade.

Note In-Service upgrade is not available for Network Extension High Availability (HA) groups. HA groups use the failover process to complete the upgrade. In this case, the Standby pair is upgraded first. After the Standby upgrade finishes, a switchover occurs and the Standby pair takes on the Active role. At that point, the previously Active pair is upgraded and takes on the Standby role.

Note HCX Network Extension In-Service/Standard Upgrade or Redeploy operation will fail, if one of the stretched networks has Port Bindings set as “Ephemeral – no binding” for a DVPG.

With an In-Service upgrade or redeployment, the following high-level workflow applies:

- A new appliance is provisioned at the source and destination site.
- New Uplink and Management IP addresses are reserved for each new Network Extension appliance.
- NICs on the new appliances are connected, including NICs for extended networks, except the NIC connection state is flagged as Down.
- Secure tunnel connections are established between sites.
- Old appliance Bridge NICs are disconnected. New Appliances Bridge NICs are connected.
- The old appliance is deleted. IP addresses used for the old appliance are released and made available.

As a result of this workflow, switchover from the old appliance to the new appliance requires only minimal action and can happen within a few seconds or less. The actual time it takes to return to forwarding traffic depends on the overall environment.

With a Standard upgrade or redeployment of the Network Extension appliance, the new appliances use the same Uplink and Management IP addresses as the existing appliance. Using the same IP addresses means that HCX must disconnect the existing appliance so that the IP addresses are available for the new appliance. In this case, tunnel connections are established only after switchover happens, requiring 30 seconds or more to re-establish data traffic across extended networks.

Network Extension upgrade or redeployment operations display a pop-up window with the option for In-Service or Standard mode deployment.

For more information about upgrading HCX appliances, see Chapter 14 Updating VMware HCX.

Prerequisites

The following prerequisites apply for In-Service upgrade or deployment.

Note In the event that the prerequisites for In-Service mode are not met, use Standard mode to complete Network Appliance upgrade or redeployment operations.

- Existing HCX-NET-EXT appliances must be running HCX 4.0 or later. This feature cannot be used while upgrading from HCX 3.5.x to HCX 4.0.
- Appliance tunnels must be in Up state as shown by navigating to **Interconnect > Service Mesh > Appliances**. Appliances with Down or Degraded tunnels are not supported.
- HCX Manager must be able to reach HCX-NET-EXT appliances via the management network.
- For each Network Appliance, free IP addresses are available in both the HCX Management and Uplink Network profile address pools during the Upgrade or Redeploy process. For example, upgrading three Network Extension appliances at the same time requires three available Management IP addresses and three available Uplink IP addresses. Once the process completes, the IP addresses used by the previous appliance are released.

Procedure

1 Select the Network Extension appliances for update or deployment.

2 Click **Update** or **Redeploy**.

The selections that appears depends on the operation being performed. The option for selecting Standard or In-Service mode can appear when updating or redeploying the Network Extension appliances from either the Service Mesh View Appliances or View Topology window, or from the Service Mesh Edit or Resync window. The following are examples of the update and deployment windows.

In the following Update Appliances screen example, NE-I1 is a standalone appliance that supports In-Service upgrade:

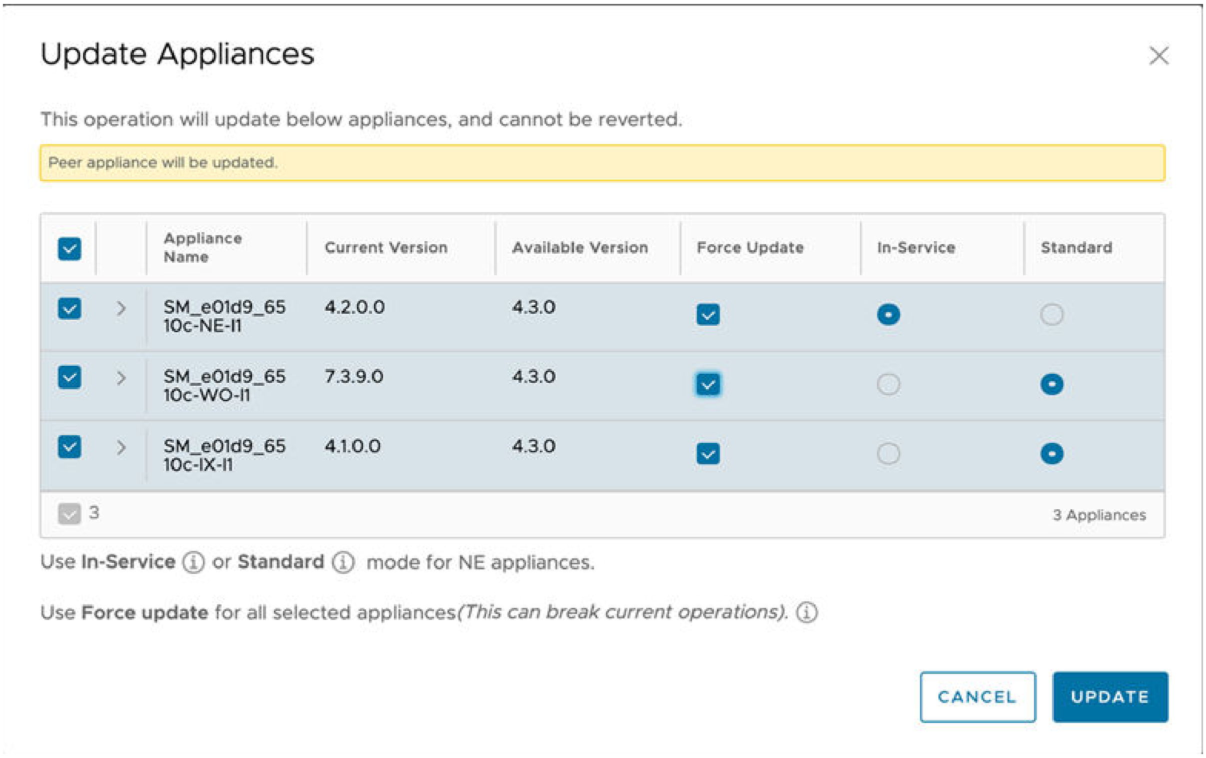

Update Appliances

This operation will update below appliances, and cannot be reverted.

Peer appliance will be updated.

☑		Appliance Name	Current Version	Available Version	Force Update	In-Service	Standard
☑	>	SM_e01d9_65 10c-NE-I1	4.2.0.0	4.3.0	☑	●	○
☑	>	SM_e01d9_65 10c-WO-I1	7.3.9.0	4.3.0	☑	○	●
☑	>	SM_e01d9_65 10c-IX-I1	4.1.0.0	4.3.0	☑	○	●

☑ 3 — 3 Appliances

Use **In-Service** ⓘ or **Standard** ⓘ mode for NE appliances.

Use **Force update** for all selected appliances*(This can break current operations).* ⓘ

In the following Redeployment Appliance screen example, NE-I1 is a standalone appliance that supports In-Service upgrade, while NE-I5 is part of an HA group with no support for In-Service upgrade.

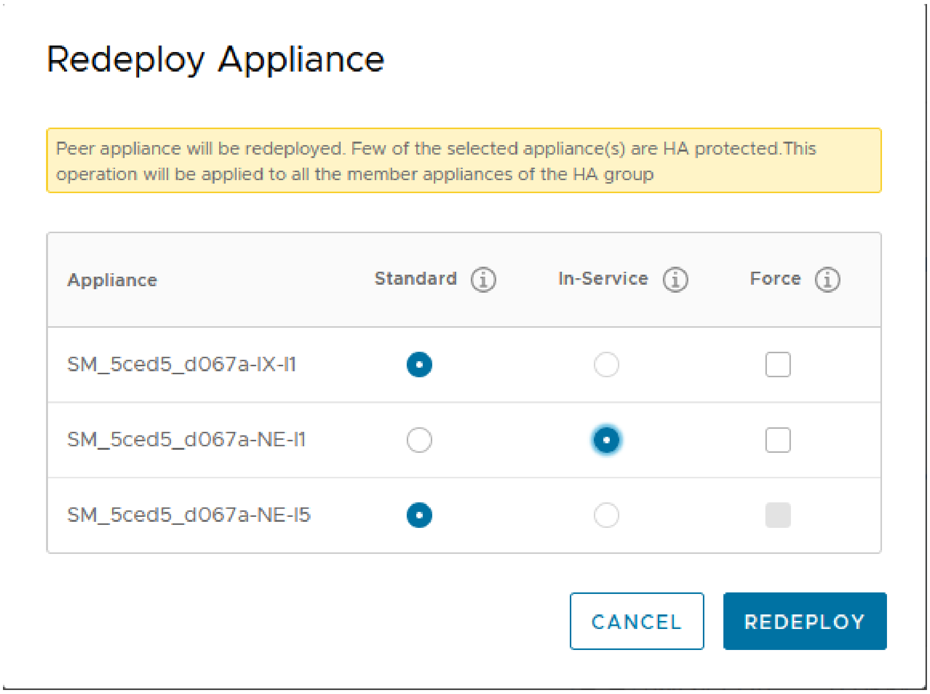

Redeploy Appliance

Peer appliance will be redeployed. Few of the selected appliance(s) are HA protected.This operation will be applied to all the member appliances of the HA group

Appliance	Standard ⓘ	In-Service ⓘ	Force ⓘ
SM_5ced5_d067a-IX-I1	●	○	□
SM_5ced5_d067a-NE-I1	○	●	□
SM_5ced5_d067a-NE-I5	●	○	□

3 Depending on the operation, click **Update** or **Redeploy**.

What to do next

Verify the Update or deployment task by navigating to **Interconnect** > **Service Mesh** > **Tasks**. After the task completes, check that the tunnel status is Up.

Removing a Network Extension

Removing a network extension prevents further cross-site communications between virtual machines residing on that network. This operation is typical when the source side network is vacated.

You can remove a network extension at any time, but be aware that removing an extended network can impact the network infrastructure. Migrated virtual machines that use a source environment DHCP server or that use statically assigned network services like DNS or NTP can lose those services after unextending a network.

Procedure

1 In the HCX Manager UI, select **Services** > **Network Extension**.

   The system displays a list of extended networks.

2 Select the network or networks that must be unextended, and click the ellipsis menu to see a list of actions.

   Note To remove multiple network extensions simultaneously, select the networks and click the **Unextend Networks** tab.

   Available actions for unextending networks are **Unextend Networks** and **Force Unextend Networks**. In most circumstances, use **Unextend Networks**.

   In cases where **Unextend Networks** fails due to the state of the Network Extension appliance, use **Force Unextend** to remove the extension and clear internal operations and processes from the source and the destination sites.

   In cases where the source components are no longer available, **Force Unextend** allows the Unextend operation, removing Network Extension components from the source and the destination sites.

   In cases when the on-premises site is not available or powered off, the existing extended networks are required to be unextended from target/Cloud Network Extension (NE) interface to ensure the Cloud NSX distributed logical router (DLR) is connected for all extended segments where migrated cloud VMs are hosted. Refer to Restrictions and Limitations for Network Extension for more details.

   After selecting an action, a dialog box appears to confirm the action.

3 (Optional) For each network that will be unextended, expand the network entry and select **Connect cloud network to cloud edge gateway after unextending** to connect the remote side gateway.

Dynamic routing can be activated on the Cloud Edge Gateway as part of the OSPF or the BGP configuration. By default, the cloud segment is left disconnected from the Edge Gateway after removing the network extension. This is done to prevent an Edge Gateway from advertising a route to the cloud segment and causing a potential routing conflict with the network in the on-premises data center. Selecting this option connects the segment to the Cloud Edge Gateway after removing the network extension. If dynamic routing is activated, the network is advertised from the Cloud Edge Gateway. Refer to VMware Cloud on AWS Networking and Security guide to ensure proper routing configuration.

Note Unextending a network removes the HCX L2 bridged path without removing the NSX Segment or vSphere Port Group, or NSX interface. The NSX router interface remains disconnected when the option **Connect cloud network** is not used.

4 To confirm the operation, click **Unextend**.

Results

HCX removes the network extension.

Understanding Network Extension High Availability

Network Extension High Availability protects extended networks from a Network Extension failure at either the source or remote site.

Overview

The Network Extension High Availability (HA) setup requires four Network Extension appliances, with two appliances at the source site and two at the remote site. Together, these two pairs form the HA Group, which is the mechanism for managing Network Extension High Availability. Appliances on the same site require the similar configuration and must have access to the same set of resources.

Similar to a standalone Network Extension deployment, where one Network Extension appliance at the source site pairs with another standalone Network Extension appliance at the remote site, each Network Extension appliance of an HA group pairs with another Network Extension appliance of the same HA group at remote site. This pairing relationship between the two appliances does not change.

Through a process of role negotiation, appliances of a pair are either Active or both Standby. During a failover event, the Standby pair takes over the Network Extension service from the Active pair.

The four Network Extension appliances of an HA group negotiate their roles of Active or Standby automatically after the HA group is formed. Following role negotiation, appliances of a pair are either both Active or both Standby. A heartbeat signal between the Active pair and the Standby pair at each site synchronizes the two appliances. A loss of heartbeats between the Active pair or the Standby pair at either the source or the remote site triggers a failover.

HCX uses vSphere DRS host anti-affinity to place the Active and the Standby appliances on separate hosts.

When an HA Group is operating normally, the HA state is Healthy. If a problem is detected with the Active appliance, the peer Standby appliance pair starts failover actions and sets its role to Active. At this point, the HA state for the group switches to Degraded. Following failover, if the failed Network Extension appliance has recovered from its failure, this appliance and its peer appliance in the pair renegotiate their roles to be Standby, and the HA group returns to Healthy state. An HA-enabled Network Extension Appliance can enter Failed state when it encounters unexpected conditions during HA related operations. The HA group also enters the Failed state as one of the appliances is in Failed state. The **Recover** selection in the HA Management tab can recover an HA group from Failed state by redeploying the appliance that is in Failed state.

For a summary of Network Extension High Availability operational states and roles, see Monitoring Network Extension High Availability.

Considerations for Network Extension High Availability

- Network Extension HA requires the HCX Enterprise license.
- Network Extension HA provides only appliance level resilience. Appliance Uplink resiliency is achieved using the Application Path Resiliency feature in the Service Mesh or multiple HCX uplinks.
- Each Active and Standby pair is managed as an HA group, which includes upgrading and redeploying appliances. The process for redeploying and updating HA groups is the same as with standalone appliances, except that the operation is applied to both Active and Standby appliances at both the source or remote site. For a list of HA group management operations, see Managing Network Extension High Availability.
- Network Extension High Availability protects against one Network Extension appliance failure in an HA group. More than one appliance failure in the same HA setup at the same time disrupts the Network Extension service.
- Network Extension HA operates in Active/Standby mode.
- Network Extension HA operates without pre-emption, with no automatic failback of an appliance pair to the Active role.
- Network Extension HA Standby appliances are assigned IP addresses from the Network Profile IP pool.

Limitations for Network Extension High Availability

- Following a failover event, policy-routed Mobility Optimized Networking (MON) traffic takes longer to recover than non-policy routed traffic. This is due to the time needed for the MON service to re-discover the next hop for the Policy Route traffic after the failover.
- Network Extension HA does not support Storage DRS anti-affinity.

Activating Network Extension High Availability

Activate Network Extension High Availability to provide failover protection for extended networks.

Network Extension High Availability activation is initiated only from the source site.

Prerequisites

- In the HCX Compute Profile at all source and destination sites, the Network Extension Appliance Limit is configured for at least two times the number of Network Extension appliances that use Network Extension HA, plus the number of appliances required to support existing non-HA extensions. See Creating a Compute Profile.

  For example, you plan to create two new HA extensions in addition to having two existing non-HA extensions from a single source to a single destination site. In this case, the Compute Profile configuration at all source and destination sites must support at least six appliances each, which is two for each HA extension, and two for non-HA extensions.
- At all source and destination sites, all Network Extension appliances required for the HA pairs are deployed prior to creating Network Extension HA groups. These appliances must not have any existing network extensions assigned to them.

  For example, you plan to create two new HA extensions to take over for two existing non-HA network extensions from a single source to a single destination site. This example requires that you create four new Network Extension appliances at the source site and four new appliances at the destination site before creating the two Network Extension HA groups.
- In the HCX Service Mesh, the Network Extension Appliance Scale Out Appliance Count is set to provide enough appliances to support network extension objectives, including any Network Extension HA groups. See Creating a Service Mesh.
- Only Network Extension appliances upgraded to HCX 4.3.0 or later can be added to HA Groups.
- In environments with vCenter in Linked Mode, all HCX systems have HCX 4.3.0 or later for HA objects to display correctly.

Procedure

1 At the source site, navigate to **Interconnect > Service Mesh > Appliances**.

2 Select an available Network Extension appliance and click **Activate High Availability**.

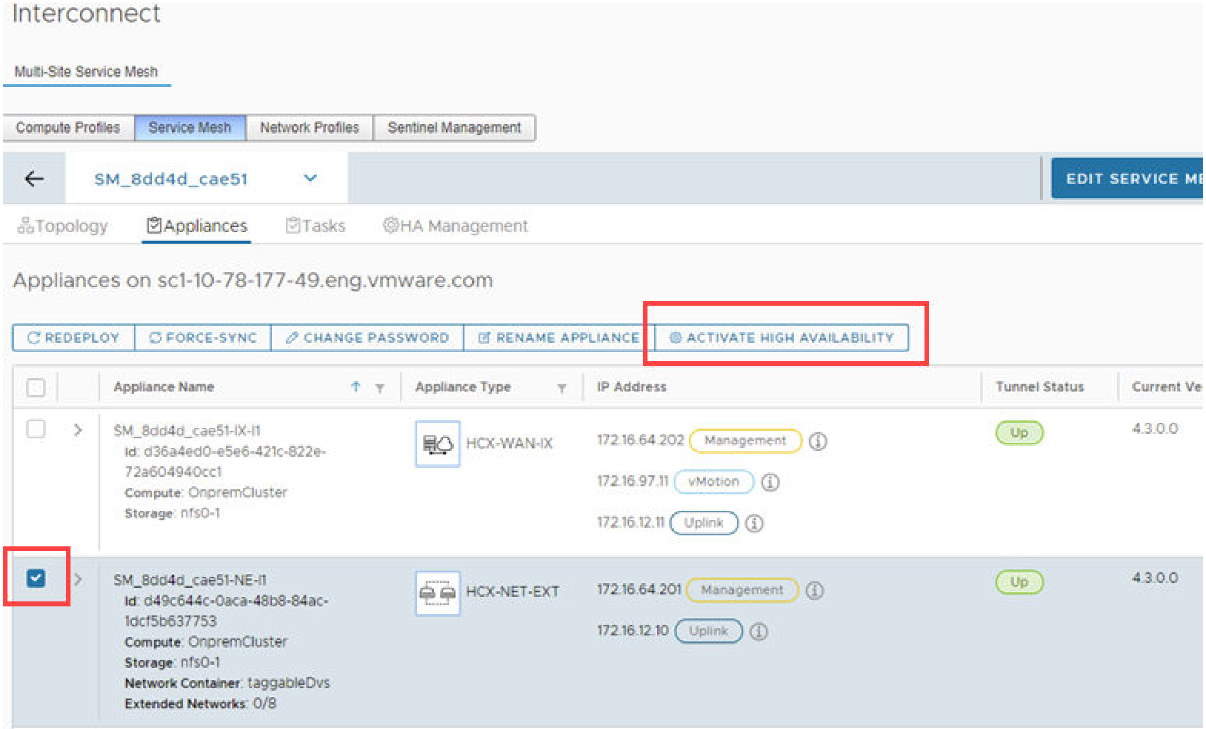

Interconnect

Multi-Site Service Mesh

Compute Profiles | Service Mesh | Network Profiles | Sentinel Management

SM_8dd4d_cae51 — EDIT SERVICE ME

Topology | Appliances | Tasks | HA Management

Appliances on sc1-10-78-177-49.eng.vmware.com

REDEPLOY | FORCE-SYNC | CHANGE PASSWORD | RENAME APPLIANCE | ACTIVATE HIGH AVAILABILITY

	Appliance Name	Appliance Type	IP Address	Tunnel Status	Current Ve
□	SM_8dd4d_cae51-IX-I1 Id: d36a4ed0-e5e6-421c-822e-72a604940cc1 Compute: OnpremCluster Storage: nfs0-1	HCX-WAN-IX	172.16.64.202 Management 172.16.97.11 vMotion 172.16.12.11 Uplink	Up	4.3.0.0
☑	SM_8dd4d_cae51-NE-I1 Id: d49c644c-0aca-48b8-84ac-1dcf5b637753 Compute: OnpremCluster Storage: nfs0-1 Network Container: taggableDvs Extended Networks: 0/8	HCX-NET-EXT	172.16.64.201 Management 172.16.12.10 Uplink	Up	4.3.0.0

A pop-up window prompts you to confirm activation.

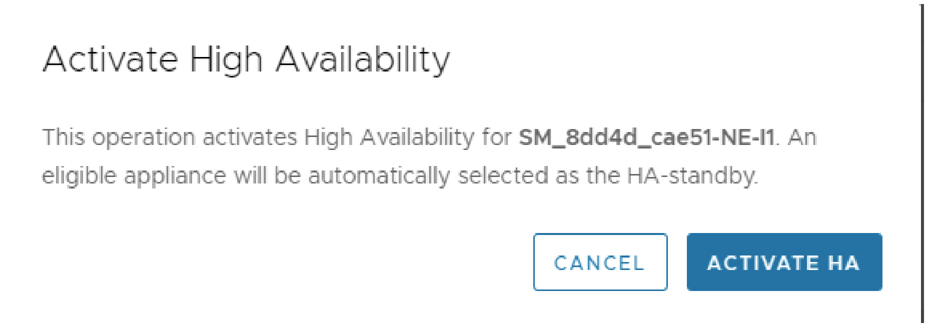

Activate High Availability

This operation activates High Availability for **SM_8dd4d_cae51-NE-I1**. An eligible appliance will be automatically selected as the HA-standby.

CANCEL | ACTIVATE HA

If the Network Extension appliance is already extending a network, or not enough appliances are available to create an HA group, the pop-up window provides information about the issue.

3 Click **Activate HA**.

The HCX system starts the process of creating the HA group. The HCX Manager automatically selects the second pair of Network Extension appliances to form the HA group. To view the activation progress, select the **Tasks** tab.

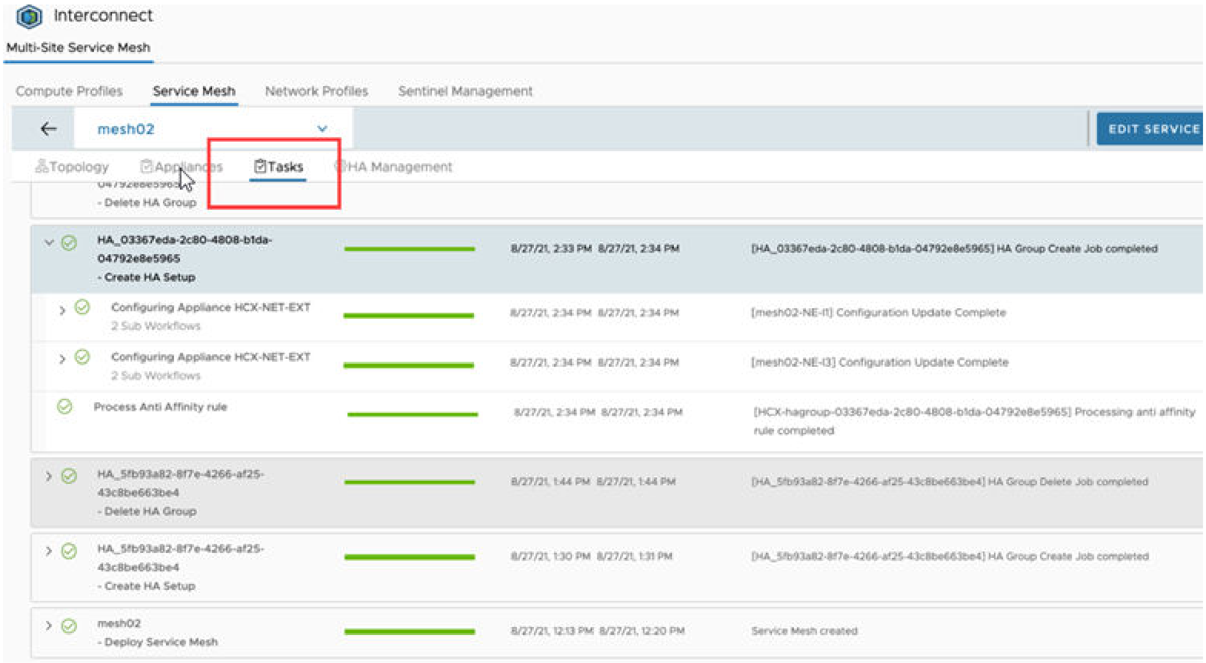

- Delete HA Group

HA_03367eda-2c80-4808-b1da-04792e8e5965 - Create HA Setup | 8/27/21, 2:33 PM 8/27/21, 2:34 PM | [HA_03367eda-2c80-4808-b1da-04792e8e5965] HA Group Create Job completed

Configuring Appliance HCX-NET-EXT 2 Sub Workflows | 8/27/21, 2:34 PM 8/27/21, 2:34 PM | [mesh02-NE-I1] Configuration Update Complete

Configuring Appliance HCX-NET-EXT 2 Sub Workflows | 8/27/21, 2:34 PM 8/27/21, 2:34 PM | [mesh02-NE-I3] Configuration Update Complete

Process Anti Affinity rule | 8/27/21, 2:34 PM 8/27/21, 2:34 PM | [HCX-hagroup-03367eda-2c80-4808-b1da-04792e8e5965] Processing anti affinity rule completed

HA_5fb93a82-8f7e-4266-af25-43c8be663be4 - Delete HA Group | 8/27/21, 1:44 PM 8/27/21, 1:44 PM | [HA_5fb93a82-8f7e-4266-af25-43c8be663be4] HA Group Delete Job completed

HA_5fb93a82-8f7e-4266-af25-43c8be663be4 - Create HA Setup | 8/27/21, 1:30 PM 8/27/21, 1:31 PM | [HA_5fb93a82-8f7e-4266-af25-43c8be663be4] HA Group Create Job completed

mesh02 - Deploy Service Mesh | 8/27/21, 12:13 PM 8/27/21, 12:20 PM | Service Mesh created

4 Review the list of appliances to verify the Active and Standby assignments.

The appliance information shows both the HA group role and the partner appliance in the group.

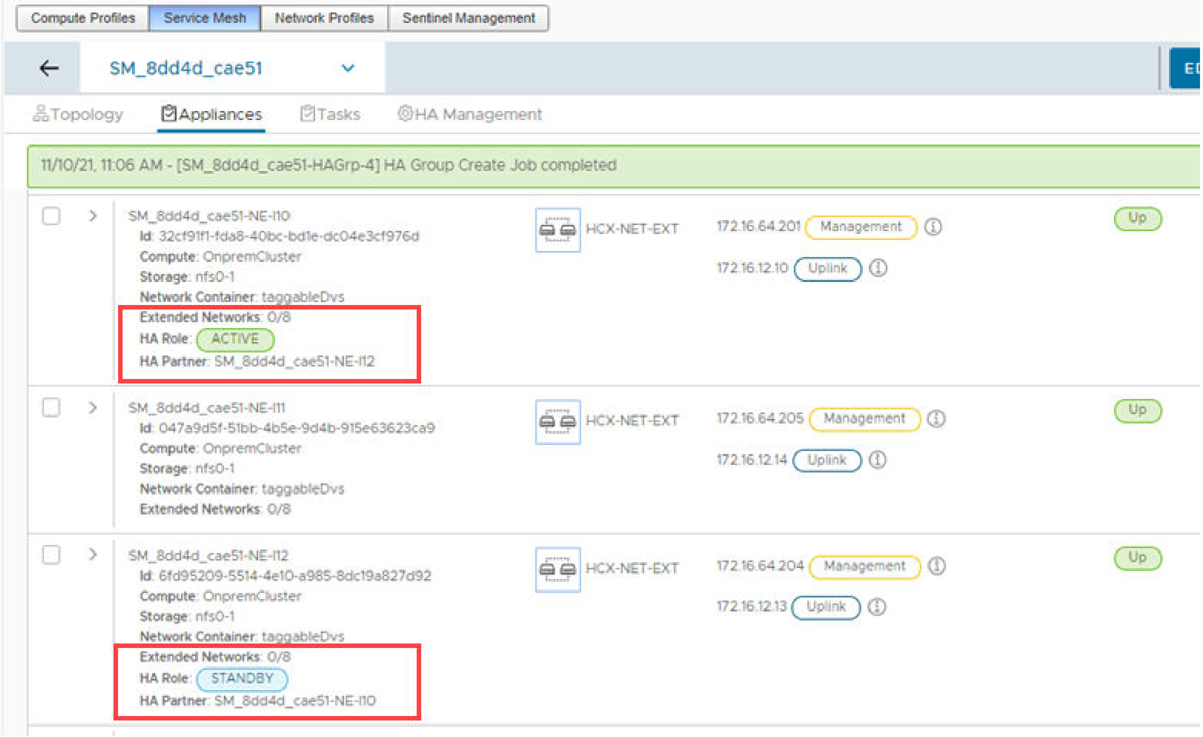

11/10/21, 11:06 AM - [SM_8dd4d_cae51-HAGrp-4] HA Group Create Job completed

SM_8dd4d_cae51-NE-I10
Id: 32cf91f1-fda8-40bc-bd1e-dc04e3cf976d
Compute: OnpremCluster
Storage: nfs0-1
Network Container: taggableDvs
Extended Networks: 0/8
HA Role: ACTIVE
HA Partner: SM_8dd4d_cae51-NE-I12
HCX-NET-EXT | 172.16.64.201 Management | 172.16.12.10 Uplink | Up

SM_8dd4d_cae51-NE-I11
Id: 047a9d5f-51bb-4b5e-9d4b-915e63623ca9
Compute: OnpremCluster
Storage: nfs0-1
Network Container: taggableDvs
Extended Networks: 0/8
HCX-NET-EXT | 172.16.64.205 Management | 172.16.12.14 Uplink | Up

SM_8dd4d_cae51-NE-I12
Id: 6fd95209-5514-4e10-a985-8dc19a827d92
Compute: OnpremCluster
Storage: nfs0-1
Network Container: taggableDvs
Extended Networks: 0/8
HA Role: STANDBY
HA Partner: SM_8dd4d_cae51-NE-I10
HCX-NET-EXT | 172.16.64.204 Management | 172.16.12.13 Uplink | Up

5 To view detailed state and role information about HA groups, click the **HA Management** tab.

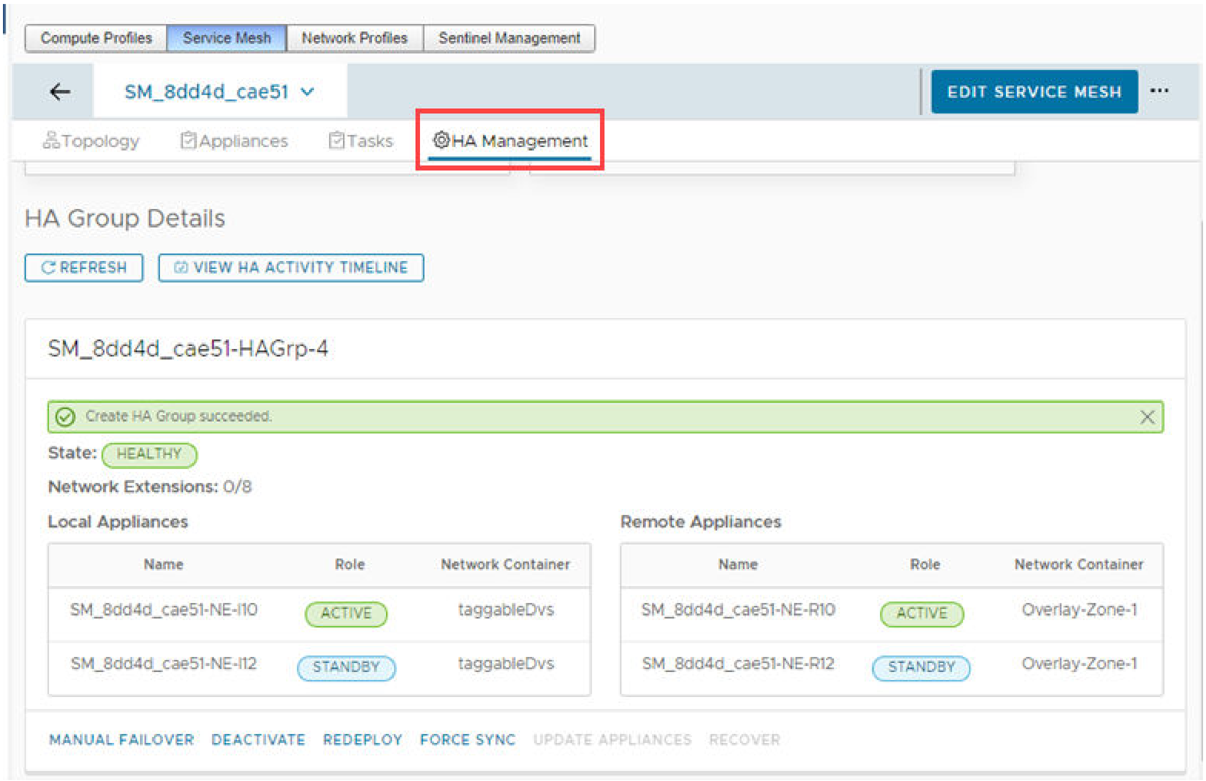

Create HA Group succeeded.

State: HEALTHY

Network Extensions: 0/8

Local Appliances

Name	Role	Network Container
SM_8dd4d_cae51-NE-I10	ACTIVE	taggableDvs
SM_8dd4d_cae51-NE-I12	STANDBY	taggableDvs

Remote Appliances

Name	Role	Network Container
SM_8dd4d_cae51-NE-R10	ACTIVE	Overlay-Zone-1
SM_8dd4d_cae51-NE-R12	STANDBY	Overlay-Zone-1

Results

After the group is created, the Active and Standby roles for the local and remote appliances display on the HA Management page.

What to do next

Select a Network Extension HA group when extending networks. See Extending Networks Using VMware HCX.

Managing Network Extension High Availability

Each Active and Standby pair, at either the source or remote site, is managed as an HA group.

The process for managing Network Extension HA groups is the same as with standalone Network Extension appliances, except that the operation is applied to both Active and Standby appliances at the respective source or remote site.

Perform Network Extension High Availability management operations in the HA Management interface.

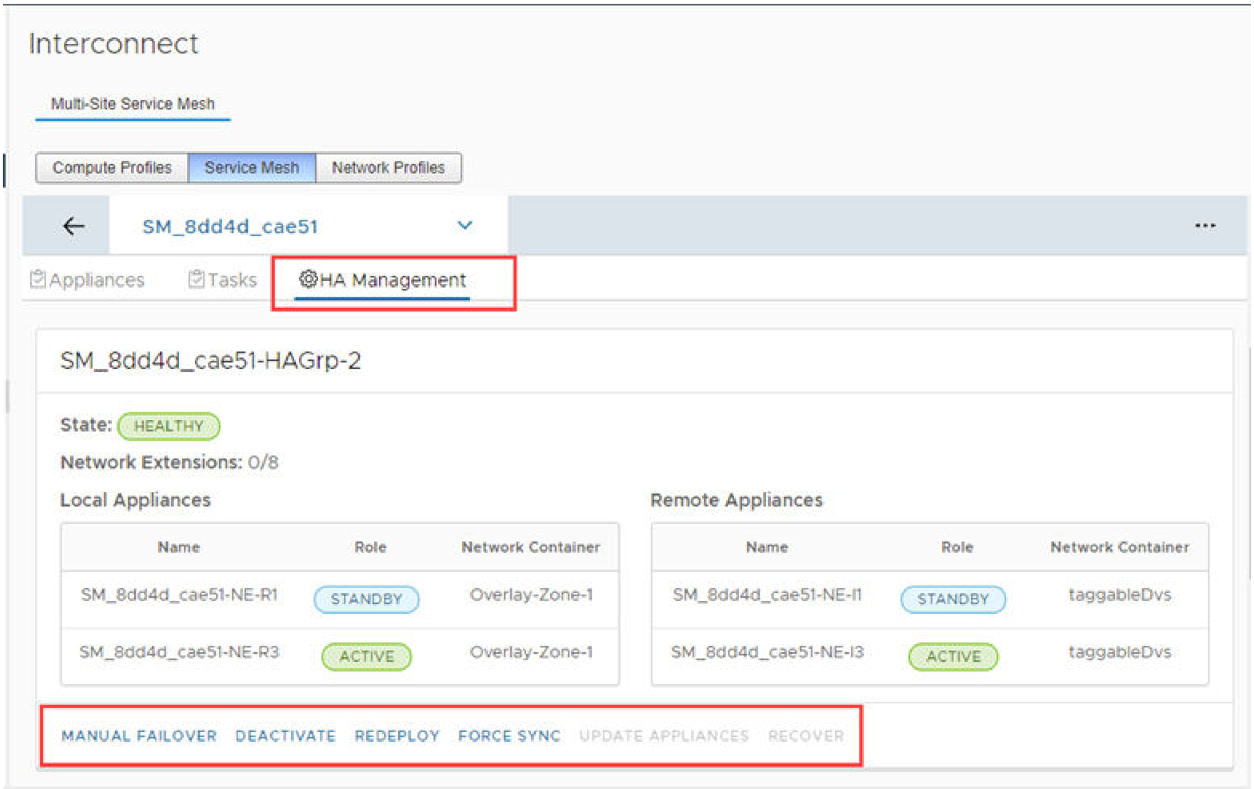

SM_8dd4d_cae51-HAGrp-2

State: HEALTHY

Network Extensions: 0/8

Local Appliances

Name	Role	Network Container
SM_8dd4d_cae51-NE-R1	STANDBY	Overlay-Zone-1
SM_8dd4d_cae51-NE-R3	ACTIVE	Overlay-Zone-1

Remote Appliances

Name	Role	Network Container
SM_8dd4d_cae51-NE-I1	STANDBY	taggableDvs
SM_8dd4d_cae51-NE-I3	ACTIVE	taggableDvs

Operation	Description
Manual Failover	The HA Active role transitions to the Standby Network Extension appliance in the HA Group. Network extensions are unprotected during the operation.
Deactivate	This operation deactivates the HA group. Any association between all the participant appliances is dissolved and the high availability capability for all the associated network extensions is lost. All HCX-NE appliances return to be standalone HCX-NE appliances for use in extending networks. When an HA Group is deactivated, any configured network extensions continue to be served by the standalone appliance pair that was in the Active role. This operation includes a **Force Deactivate** option that attempts a best effort to deactivate HA Group even if any error occurred during the process. Using **Force Deactivate** can result in stranded network extensions on the Standby NE appliances.
Redeploy	This operation redeploys the HA group appliances. The operation also redeploys peer appliances at the remote site.
Force Sync	This operation synchronizes the configuration on HA group appliances, including peer appliances, at the remote site. This operation is applied to all the member appliances of the HA group.

Operation	Description
Update Appliances	Use this option to update Network Extension appliances that are part of an HA group.
Recover	This operation attempts to return an HA group to a Healthy state. This can require redeploying some appliances. To track the state of a recover operation, navigate to **Interconnect > Service Mesh > View Appliances**, and select the **Tasks** tab.

Monitoring Network Extension High Availability

Information for monitoring Network Extension High Availability activity includes the HA group health, group roles, and appliance status.

Information related to Network Extension High Availability operation and status appears on multiple HCX interface pages: HA Management, Site Pairs (Mobility Mesh), and Network Extension.

Information related to Network Extension High Availability operation and status appears on multiple HCX interface pages: HA Management, Service Mesh Appliance, and Network Extension.

Note Although HA operations take effect within a matter of seconds, it can take a couple of minutes for those changes to propagate to the UI display.

HA Management Information

The HA Management page provides the overall HA group health along with detailed information regarding the condition of individual HA groups at both the source and the destination sites.

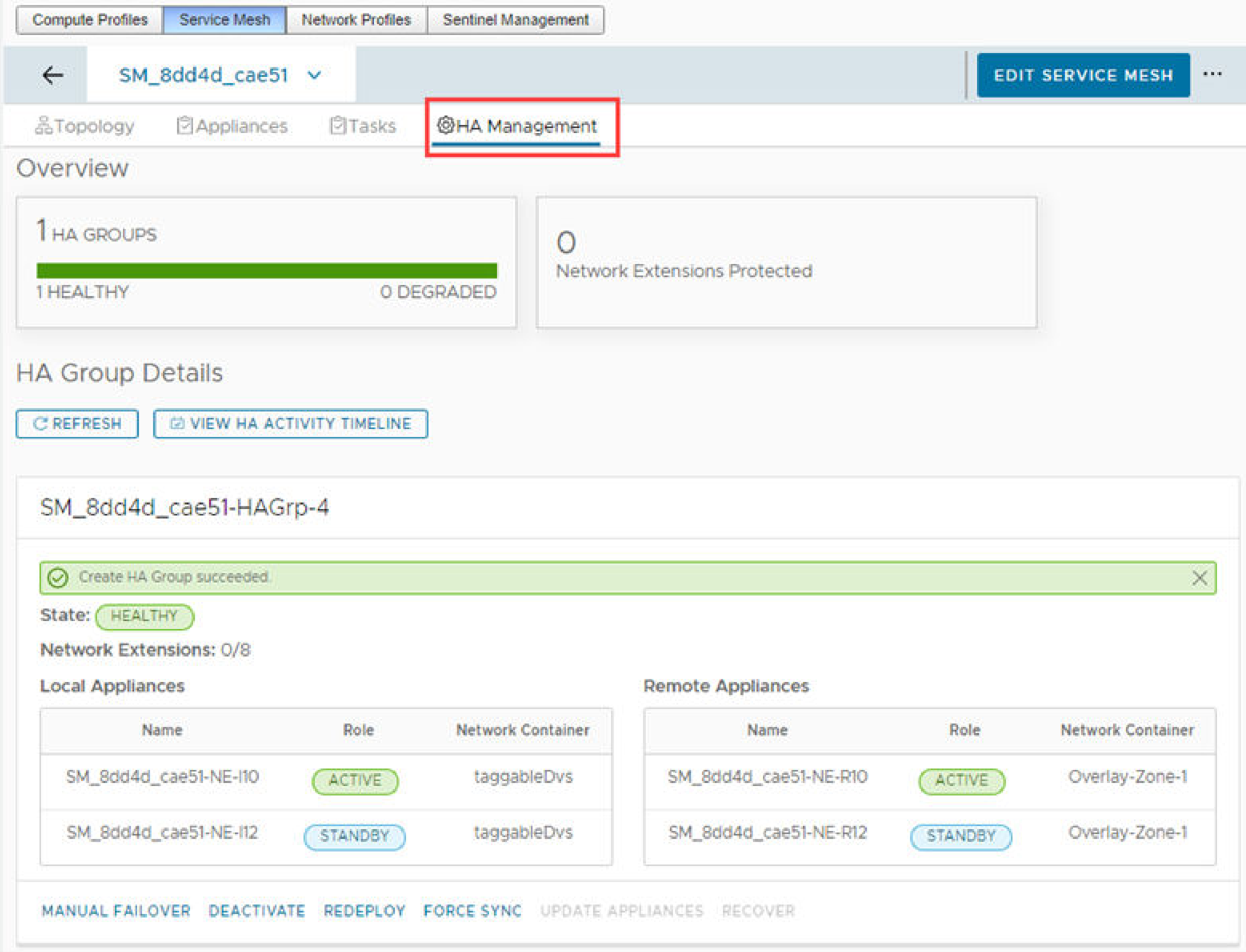

Overview

1 HA GROUPS

1 HEALTHY 0 DEGRADED

0
Network Extensions Protected

HA Group Details

SM_8dd4d_cae51-HAGrp-4

Create HA Group succeeded.

State: HEALTHY

Network Extensions: 0/8

Local Appliances

Name	Role	Network Container
SM_8dd4d_cae51-NE-I10	ACTIVE	taggableDvs
SM_8dd4d_cae51-NE-I12	STANDBY	taggableDvs

Remote Appliances

Name	Role	Network Container
SM_8dd4d_cae51-NE-R10	ACTIVE	Overlay-Zone-1
SM_8dd4d_cae51-NE-R12	STANDBY	Overlay-Zone-1

The Overview information indicates the total number of HA groups and the number of those groups that are in a Healthy or a Degraded state. This information also specifies the total number of extended networks protected by an HA group.

The HA Group Details information provides the group State and Role information. HA Groups have one of the several states:

Group State	Description
HEALTHY	The HA group is functioning as expected.
DEGRADED	One or more of the Network Extension appliances in the HA group are not in a state to perform HA operations. This might not indicate a problem. For example, the HA group state is Degraded when the group is recovering from a failover operation. Administrator intervention might be required to determine the issue. **Note** The Degraded state only refers to the HA group. The Active appliance continues to serve the Network Extension service.
FAILINGOVER	A failover operation is in progress for the HA group.

Group State	Description
FAILED	The HA group cannot be automatically recovered from a failover. Try to recover by clicking the **Recovery** button redeploys the failed appliance pair.
MAINTENANCE	The HA group is in the process of redeploying or synchronizing as initiated by User.

Network Extension appliances in an HA group can have one of the several roles:

Group Role	Description
UNINIT	The Network Appliance is pending completion of the HA specific configuration.
UNDECIDED	The Network Extension appliance is pending role negotiation with other appliances in the HA group, and its role is not set.
ACTIVE	Identifies the Network Extension appliance that is currently carrying out the network extension service by handling the network traffic.
STANDBY	Identifies the Network Extension appliance that is currently monitoring the ACTIVE appliance but does not provide network extension services and handles no network traffic.

The HA Group Details information includes selections to refresh the page data and to view a timeline of HA activity.

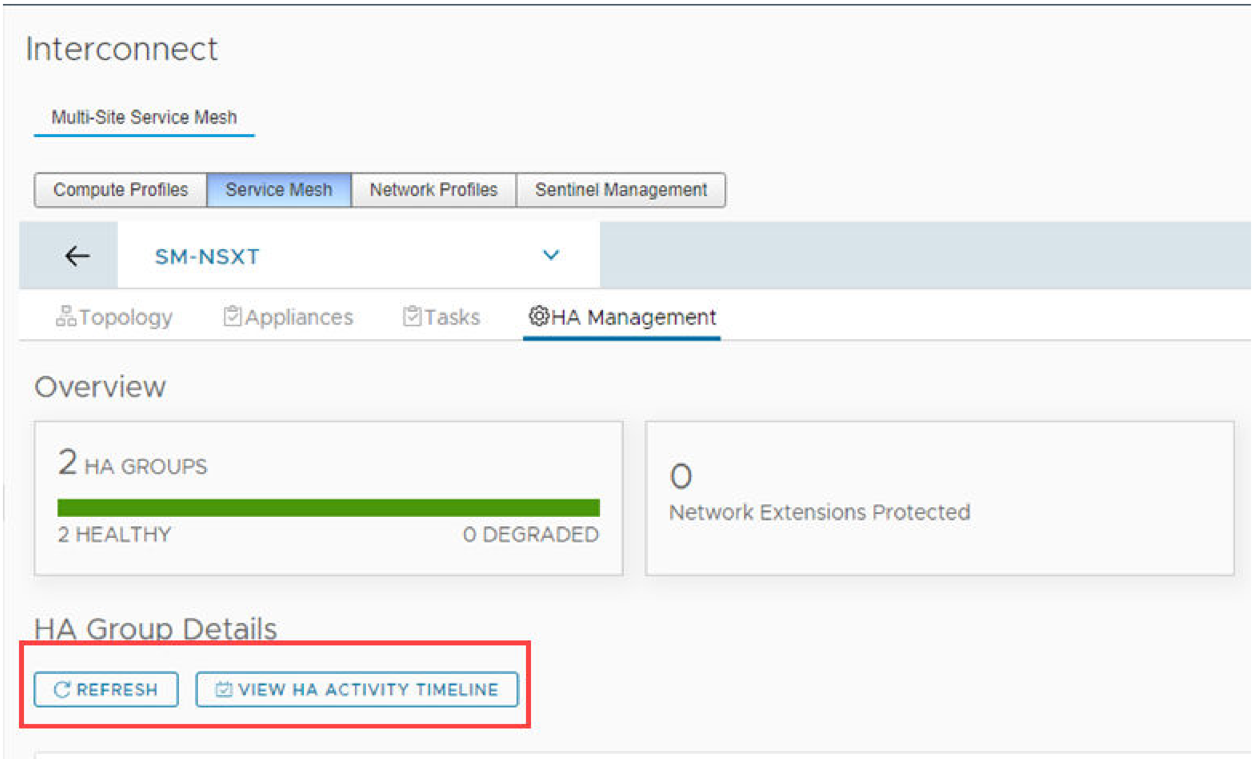

The HA activity timeline provides options selecting the date range and display view, along with an option to filter activity by HA group.

Service Mesh Appliances Information

The Service Mesh Appliances page provides two entries for monitoring Network Extension HA status: HA Role and HA Partner. The HA Partner entry identifies the other appliance on the same site in the HA group, that they exchange heartbeat signals. The HA Role entry identifies which appliance is operating as the Active or Standby partner in the HA group at the moment.

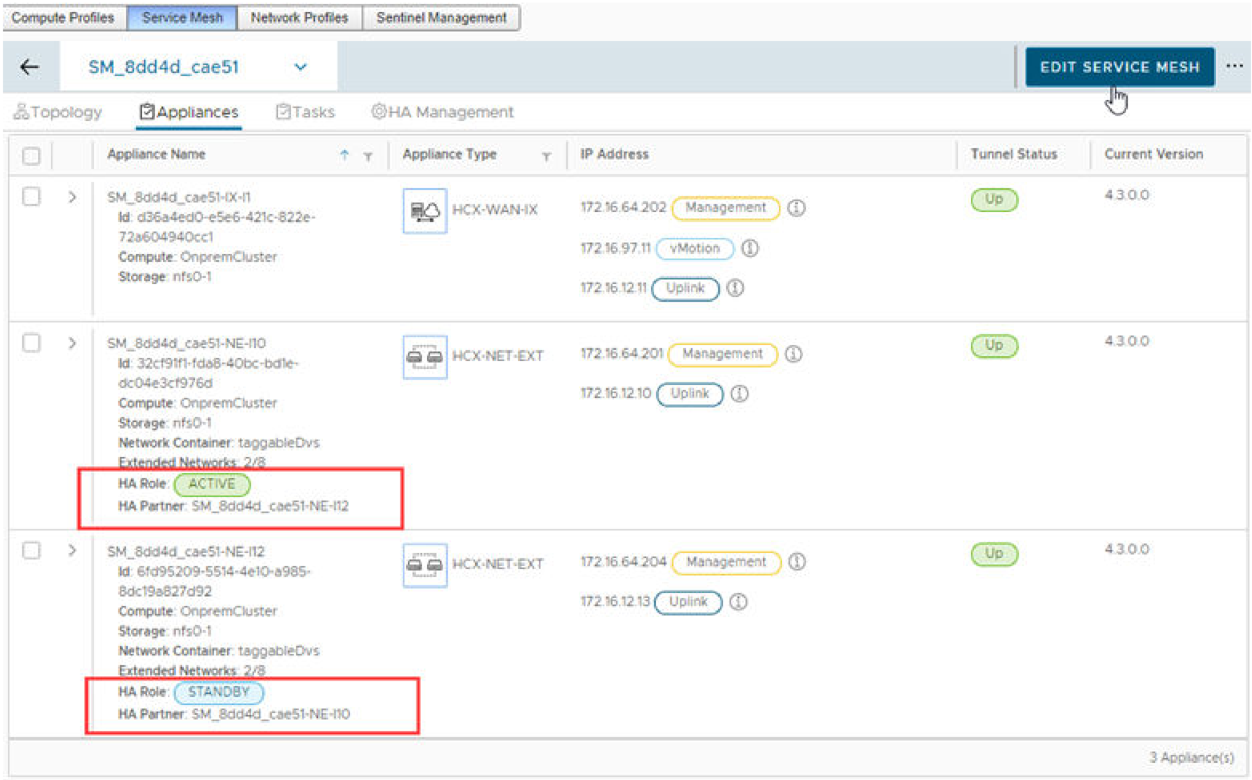

□		Appliance Name	Appliance Type	IP Address	Tunnel Status	Current Version
□	>	SM_8dd4d_cae51-IX-I1 Id: d36a4ed0-e5e6-421c-822e-72a604940cc1 Compute: OnpremCluster Storage: nfs0-1	HCX-WAN-IX	172.16.64.202 Management 172.16.97.11 vMotion 172.16.12.11 Uplink	Up	4.3.0.0
□	>	SM_8dd4d_cae51-NE-I10 Id: 32cf91f1-fda8-40bc-bd1e-dc04e3cf976d Compute: OnpremCluster Storage: nfs0-1 Network Container: taggableDvs Extended Networks: 2/8 HA Role: ACTIVE HA Partner: SM_8dd4d_cae51-NE-I12	HCX-NET-EXT	172.16.64.201 Management 172.16.12.10 Uplink	Up	4.3.0.0
□	>	SM_8dd4d_cae51-NE-I12 Id: 6fd95209-5514-4e10-a985-8dc19a827d92 Compute: OnpremCluster Storage: nfs0-1 Network Container: taggableDvs Extended Networks: 2/8 HA Role: STANDBY HA Partner: SM_8dd4d_cae51-NE-I10	HCX-NET-EXT	172.16.64.204 Management 172.16.12.13 Uplink	Up	4.3.0.0

3 Appliance(s)

Network Extension Monitoring Information

The Network Extension page provides an "HA" tag in the Extension Appliance column to identify networks that are protected using Network Extension High Availability.

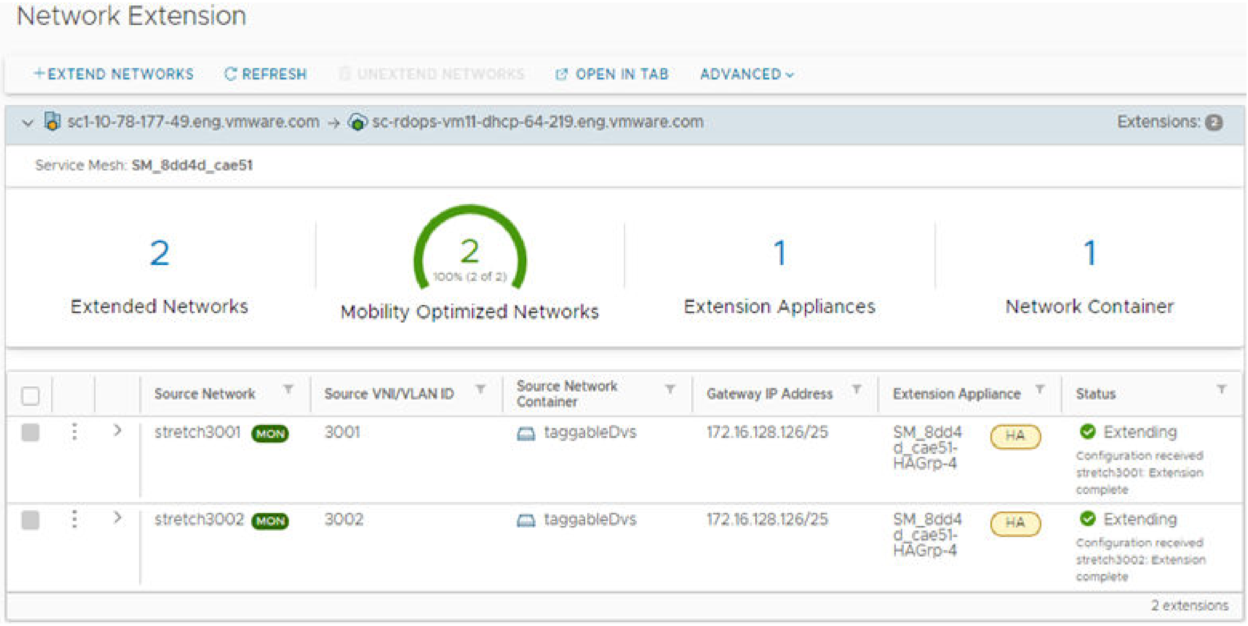

sc1-10-78-177-49.eng.vmware.com → sc-rdops-vm11-dhcp-64-219.eng.vmware.com Extensions: 2

Service Mesh: SM_8dd4d_cae51

2	2 100% (2 of 2)	1	1
Extended Networks	Mobility Optimized Networks	Extension Appliances	Network Container

□			Source Network	Source VNI/VLAN ID	Source Network Container	Gateway IP Address	Extension Appliance	Status
□	⋮	›	stretch3001 MON	3001	taggableDvs	172.16.128.126/25	SM_8dd4d_cae51-HAGrp-4 HA	✓ Extending Configuration received stretch3001: Extension complete
□	⋮	›	stretch3002 MON	3002	taggableDvs	172.16.128.126/25	SM_8dd4d_cae51-HAGrp-4 HA	✓ Extending Configuration received stretch3002: Extension complete

2 extensions

Migrating Virtual Machines

8

Workloads can be migrated bi-directionally between data centers using various VMware HCX migration technologies.

Organizations migrate application workloads for many reasons. From data center consolidation and evacuation to modernization and maintenance, migrating workloads requires analysis and planning. Administrators identify individual workloads for migration, or waves of workloads based on, for example, cluster, network, or application landscape. HCX provides an array of migration types for moving these workloads including cold, warm, and live migration.

HCX also provides procedures for migrating groups, or waves, of virtual machines. And through integration with VMware vRealize Network Insight, Application Group information can be exported to HCX for migration as Mobility Groups. HCX Mobility Groups and vRealize Network Insight integration with HCX are available with the HCX Enterprise license.

Read the following topics next:

- HCX Integration with vRealize Network Insight

VMware HCX Migration Types

Virtual Machines can be moved to and from HCX-activated data centers using multiple migration technologies.

Note For information regarding the number of concurrent migrations supported for a specific migration type, see VMware Configuration Maximums for HCX.

Bulk Migration

This migration method uses the VMware vSphere Replication protocols to move the virtual machines to a destination site.

- The Bulk migration option is designed for moving virtual machines in parallel.
- This migration type can set to complete on a pre-defined schedule.
- The virtual machine runs at the source site until the failover begins. The service interruption with the bulk migration is equivalent to a reboot.

For more information, refer to Understanding VMware HCX Bulk Migration.

vMotion Migration

This migration method uses the VMware vMotion protocol to move a virtual machine to a remote site.

- The vMotion migration option is designed for moving single virtual machine at a time.
- The virtual machine state is migrated. There is no service interruption during an HCX vMotion migration.

For more information, refer to Understanding HCX vMotion and Cold Migration.

Cold Migration

This migration method uses the VMware NFC protocol. It is automatically selected when the source virtual machine is powered off.

For more information, refer to Understanding HCX vMotion and Cold Migration.

Replication Assisted vMotion

Replication Assisted vMotion (RAV) combines advantages from Bulk Migration (parallel operations, resiliency, and scheduling) with vMotion Migration (zero downtime virtual machine state migration).

For more information, refer to Understanding HCX Replication Assisted vMotion.

VMware HCX OS Assisted Migration

This migration method provides for the bulk migration of guest (non-vSphere) virtual machines using OS Assisted Migration to VMware vSphere on-premises or cloud-based data centers. Activating this service requires additional HCX licensing.

For more information, refer to Understanding VMware HCX OS Assisted Migration.

Workload Migrations for NSX V2T

Workload Migrations for NSX V2Toperates within a single vCenter, across vCenters, or to Federated NSX environments for the lift-and-shift of workloads from an NSX for vSphere environment to an NSX environment.

For more information, refer to Understanding Workload Migrations for NSX V2T.

Mobility Agent vSphere Host for HCX Migrations

HCX uses the Mobility Agent service to perform vMotion, Cold, and Replication Assisted vMotion (RAV) migrations.

The HCX-IX appliance adds the Mobility Agent service as a host object in the vCenter Server. The processor, memory, storage and networking resources displayed on this object does not represent actual consumption on the physical hypervisor hosting the IX appliance.

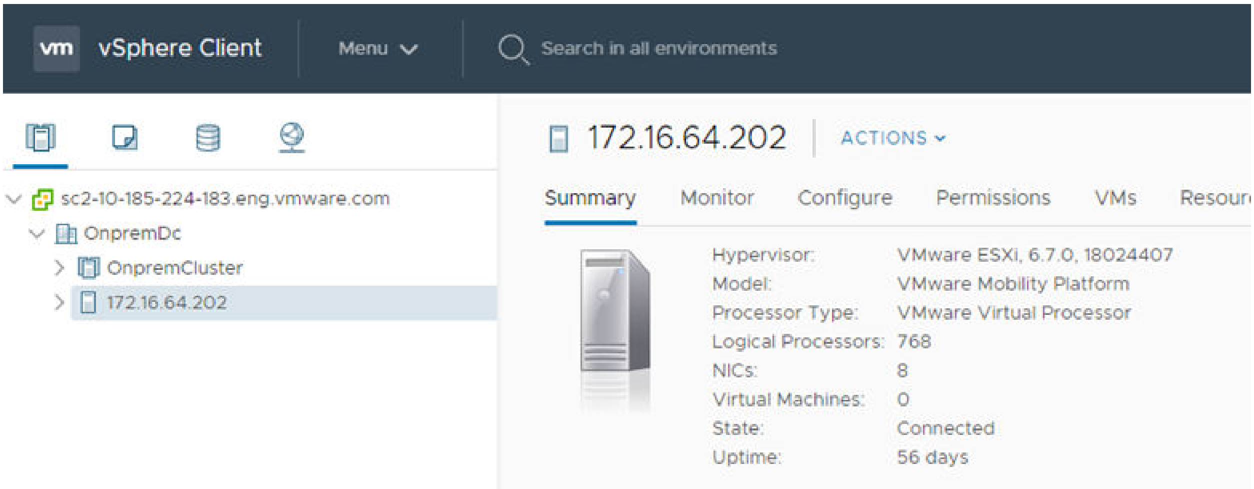

Hypervisor: VMware ESXi, 6.7.0, 18024407
Model: VMware Mobility Platform
Processor Type: VMware Virtual Processor
Logical Processors: 768
NICs: 8
Virtual Machines: 0
State: Connected
Uptime: 56 days

Caution The Mobility Agent host is required. Deleting this host can disrupt HCX Cold, vMotion, and Replication Assisted vMotion (RAV) migrations.

Understanding VMware HCX Bulk Migration

Bulk migration uses the host-based replication to relocate a virtual machine between HCX data centers.

To reduce the downtime, the source VM remains online during the replication and is bootstrapped on the destination ESX host after replication completes.

A Bulk Migration request triggers the following actions:

1 Replication begins an initial full synchronization transfer to the remote site. The time it takes to replicate is a function of the size of the VM, the data change rate on the virtual machine disk files (VMDKs), and available bandwidth.

2 Replication bandwidth consumption varies depending on the number of disks for each VM and the number of VMs being migrated concurrently.

3 The switchover can start immediately after the initial synchronization is completed, or it can be delayed until a specific time using the scheduled migration option. By using the scheduled migration option, the switchover can occur during a maintenance window.

4 A delta synchronization with two-hour recovery point objective (RPO) occurs while waiting for the scheduled switchover, after the initial synchronization is completed.

5 Depending upon data churn on source disk, additional snapshots are being created during the RPO cycle. After each RPO cycle, disk consolidation takes place and creates a “hbrdisk.RDID vmdk” called as replica instance vmdk file on target datastore. Refer VMware KB 87028.

6 During switchover, the source VM is powered-off to perform a final off-line synchronization, data consolidation, and VM instantiation at the target data center.

7 Following the switchover, the migrated VM replica is powered-on and HCX Manager renames the original VM using a POSIX timestamp suffix to avoid a naming conflict with the migrated VM. If the **Retain MAC** option was not selected, the migrated VM obtains a new MAC address.

8 The migration completes and the original VM is copied to the Migrated VMs folder.

Important There must be sufficient resources to power on the VM. If for some reason the new VM cannot power on, it remains powered off, and the original is powered on.

VMware HCX copies the original VM to the Migrated VMs folder. This copy can be accessed in the vSphere Virtual Machines and Templates view, and can be used to recover the migrated VM.

Single vCenter Operation

You can use Bulk migration within a single vSphere vCenter (VC) Server. In this case, both HCX Connector and HCX Cloud Manager are deployed to different clusters within the same VC. Other than deployment within the same VC, the process for configuring HCX is the same, which means configuring the Compute Profile and the Service Mesh to enable Bulk migration.

Note Bulk migration from one cluster to another in a single VC environment does not retain the MAC address of the VM even if the option is explicitly selected. See VMware KB article 219.

Requirements for Bulk Migration

- The Hybrid Interconnect Service and the Bulk Migration Service must be activated and in a healthy state in the relevant service mesh.
- The resources to create, power on and use the virtual machine must be available in the destination environment.
- Virtual machines must be running Hardware Version 7 or later.
- Virtual machines must have VMware Tools installed.
- Virtual machines must reside in a Service Cluster (defined in the Compute Profile).
- Network Extension is required for low downtime migration operations.
- Personalization Scripts and System Identity changes (Hostname, IP, SID) require the system to be rebooted one additional time during the switchover phase.
- Bulk Migration potential throughput can vary depending on bandwidth available for migrations, latency, available CPU/MEM/IOPS, and disk read speed. For a successful switchover phase, the bandwidth and the network conditions must be sufficient to satisfy the operation considering the dataset and virtual machine data change rate. For more information about how to determine bandwidth requirements, see Bandwidth Requirements for vSphere Replication.

For information regarding the migration limits, see VMware Configurations Maximums.

Restrictions for Bulk Migration

- Virtual machines with Raw Device Mappings (RDM) in Physical Compatibility mode cannot be migrated. Virtual machines with RDM in Virtual Compatibility mode can be migrated. They are converted to VMDKs at the destination.
- Virtual machines with mounted ISO images cannot be migrated. The HCX Bulk migration operation can be used for force unmount ISO images.
- Virtual machine snapshots are not migrated. The HCX Bulk migration operation has an option to remove the snapshots.
- Virtual machines with DirectPath I/O configurations cannot be migrated without first removing the DirectPath device.
- Virtual machines cannot be migrated while they are using a shared SCSI bus, flagged for multi-writer, enabled for fault tolerance, or configured for shared VMDK disk sharing.
- Virtual machines that cannot be gracefully powered off cannot be migrated. HCX can override with the Force Power-off VM option.
- Virtual machines using virtual NVMe (vNVME) Controllers cannot be migrated.
- Migration to or from vVOL datastores is not supported.
- The HCX Migration interface does not display port groups that are VLAN trunks.

- VMware software appliances (including, but not limited to vCenter Server and NSX Manager) cannot be migrated with HCX Bulk migration.
- vSphere VM Encryption is not supported with HCX migrations.
- Virtual Based Security (VBS) functionality is not supported with VMware HCX migrations.

Guest OS Customization with Bulk Migration

One characteristic of Bulk Migration is the ability to customize several aspects of the Guest OS.

In general, it is best practice when migrating virtual machines on an extended network to keep a virtual machine's IP address, MAC address, and overall identity. In some scenarios, it can be beneficial to modify a VM’s characteristics. For example, it might be necessary to migrate non-production workloads to free up private network prefixes, and making these changes during the migration can save the effort of manually updating the VM settings after the migration. The following guest customizations are available:

- Guest OS Hostname
- IP Address
- Gateway
- Netmask (Subnet Mask)
- Primary DNS
- Secondary DNS
- Security Identifier (Windows SID)
- Run Pre- or Post-Guest Customization scripts

Guest customization is available only for specific guest OS types. See Guest OS Types for Guest Customization.

You select guest customizations using the **Edit Extended Options** in the Migration interface. Certain customization options only appear when you select the **Edit Extended Options** on a virtual machine level.

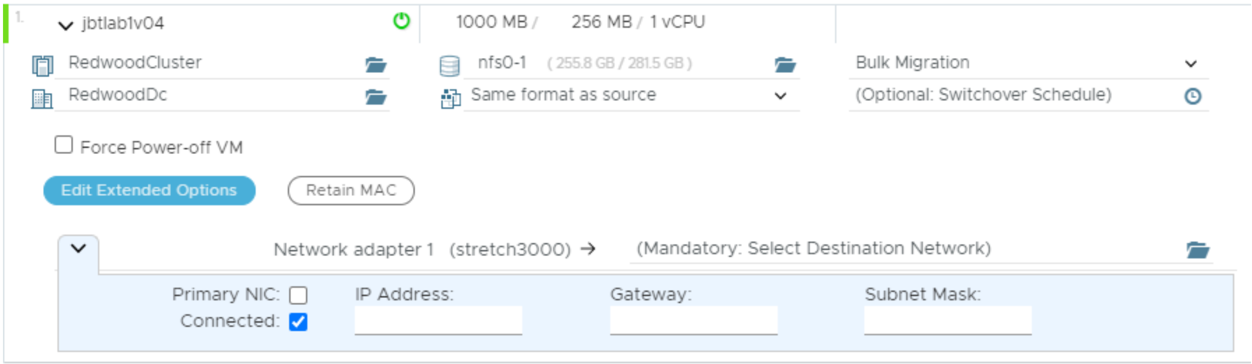

HCX applies the guest customization options during the switchover phase when the virtual machine is powered on.

This functionality is also supported for reverse migrations.

Caution When changing Guest OS information, HCX does not store the original settings.

All values must be specified in the wizard, even those that must remain unchanged.

Values will be cleared on the migrated virtual machine for fields left empty if IP Customization is configured.

For details about the Extended Options available for Bulk migration, see Additional Migration Settings.

Note Changing the Security Identifier of a Windows machine that is already the member of a Windows domain breaks the domain relationship and requires the machine to be re-joined. On a domain controller, this operation can impact the domain. By default, the Generate New Security Identifier (SID) option is not selected.

Guest OS Types for Guest Customization

HCX supports Guest OS customization for specific Windows and Linux operating systems.

For Bulk migrations, HCX supports customizing various aspects of the guest OS on the destination virtual machine. For more information about guest OS customization, see Guest OS Customization with Bulk Migration.

The virtual machine Guest OS type and guestID are reflected as virtual machine **config** parameters in the vCenter Server Managed Object Browser (`https://`*`vcenterfqdn`*`/mob`). For IP Customization to work, the virtual machine **guestId** entry must match the supported **Guest OS Types**.

Windows Operating System Types Supported for Customization

vCenter GuestOS ID (guestID)	Guest OS Type	Minimum HCX Version
win31Guest	Windows 3.1	4.0.0
win95Guest	Windows 95	4.0.0
win98Guest	Windows 98	4.0.0
winntGuest	Windows NT	4.0.0
win2000ProGuest	Windows 2000 Professional	4.0.0
win2000ServGuest	Windows 2000 Server	4.0.0
win2000AdvServGuest	Windows 2000 Advanced Server	4.0.0
winXPProGuest	Windows XP (32-bit)	4.0.0
winXPPro64Guest	Windows XP (64-bit)	4.0.0
winNetEnterpriseGuest	Windows Server 2003 Enterprise (32-bit)	4.0.0
winNetDatacenterGuest	Windows Server 2003 Data Center (32-bit)	4.0.0
winNetStandardGuest	Windows Server 2003 Standard (32-bit)	4.0.0

vCenter GuestOS ID (guestID)	Guest OS Type	Minimum HCX Version
winNetWebGuest	Windows Server 2003 Web (32-bit)	4.0.0
winNetBusinessGuest	Windows Server 2003 Business (32-bit)	4.0.0
winNetEnterprise64Guest	Windows Server 2003 Enterprise (64-bit)	4.0.0
winNetDatacenter64Guest	Windows Server 2003 Enterprise_DC_(64-bit)	4.0.0
winNetStandard64Guest	Windows Server 2003 Enterprise_SE_(64-bit)	4.0.0
winVistaGuest	Windows Vista (32-bit)	4.0.0
winVista64Guest	Windows Vista (64-bit)	4.0.0
winLonghornGuest	Windows Server 2008 (32-bit)	4.0.0
winLonghorn64Guest	Windows Server 2008 (64-bit)	4.0.0
windows7Guest	Windows 7 (32-bit)	4.0.0
windows7_64Guest	Windows 7 (64-bit)	4.0.0
indows7Server64Guest	Windows 7 Server (64-bit)	4.0.0
windows8Guest	Windows 8 (32-bit)	4.0.0
windows8_64Guest	Windows 8 (64-bit)	4.0.0
windows8Server64Guest	Windows 8 Server (64-bit)	4.0.0
windows9Guest	Windows 10	4.0.0
windows9_64Guest	Windows 10 (64-bit)	4.0.0
windows9Server64Guest	Windows 10 Server (64-bit)	4.0.0
windows2019srv_64Guest	Windows 2019 Server	4.2.2

Linux Operating System Types Supported for Customization

vCenter GuestOS ID (guestID)	Guest OS Type	Minimum HCX Version
asianux3Guest	Asianux Server 3 (32-bit)	4.0.0
asianux3_64Guest	Asianux Server 3 (64-bit)	4.0.0
asianux4Guest	Asianux Server 4 (32-bit)	4.0.0
asianux4_64Guest	Asianux Server 4 (64-bit)	4.0.0
asianux5_64Guest	Asianux Server 5 (64-bit)	4.0.0
centosGuest	CentOS 4/5 (32-bit)	4.0.0
centos64Guest	CentOS 4/5 (64-bit)	4.0.0
coreos64Guest	CoreOS (64-bit)	4.0.0
debian4Guest	Debian GNU/Linux 4 (32 bit)	4.0.0
debian4_64Guest	Debian GNU/Linux 4 (64-bit)	4.0.0
debian5Guest	Debian GNU/Linux 4 (32-bit)	4.0.0
debian5_64Guest	Debian GNU/Linux 5 (64-bit)	4.0.0
debian6Guest	Debian GNU/Linux 6 (64-bit)	4.0.0

vCenter GuestOS ID (guestID)	Guest OS Type	Minimum HCX Version
debian6_64Guest	Debian GNU/Linux 6 (64-bit)	4.0.0
debian7Guest	Debian GNU/Linux 7 (32-bit)	4.0.0
debian7_64Guest	Debian GNU/Linux 7 (64-bit)	4.0.0
debian8Guest	Debian GNU/Linux 8 (32-bit)	4.0.0
debian8_64Guest	Debian GNU/Linux 8 (64-bit)	4.0.0
oracleLinuxGuest	Oracle Linux 4/5 (32-bit)	4.0.0
oracleLinux64Guest	Oracle Linux 4/5 (64-bit)	4.0.0
rhel7Guest	Red Hat Enterprise Linux 7 (32-bit)	4.0.0
rhel7_64Guest	Red Hat Enterprise Linux 7 (64-bit)	4.0.0
rhel6Guest	Red Hat Enterprise Linux 6 (32-bit)	4.0.0
rhel6_64Guest	Red Hat Enterprise Linux 6 (64-bit)	4.0.0
rhel5Guest	Red Hat Enterprise Linux 5 (32-bit)	4.0.0
rhel5_64Guest	Red Hat Enterprise Linux 5 (64-bit)	4.0.0
fedoraGuest	Red Hat Fedora Linux (32-bit)	4.0.0
fedora64Guest	Red Hat Fedora Linux (64-bit)	4.0.0
sles12Guest	Suse Linux Enterprise Server 12 (32-bit)	4.0.0
sles12_64Guest	Suse Linux Enterprise Server 12 (64-bit)	4.0.0
sles11Guest	Suse Linux Enterprise Server 11 (32-bit)	4.0.0
sles11_64Guest	Suse Linux Enterprise Server 11 (64-bit)	4.0.0
sles10Guest	Suse Linux Enterprise Server 10 (32-bit)	4.0.0
sles10_64Guest	Suse Linux Enterprise Server 10 (64-bit)	4.0.0
opensuseGuest	OpenSUSE Linux (32-bit)	4.0.0
opensuse64Guest	OpenSUSE Linux (64-bit)	4.0.0
ubuntuGuest	Ubuntu Linux (32-bit)	4.0.0
ubuntu64Guest	Ubuntu Linux (64-bit)	4.0.0
otherlinuxguest	Linux 2.2x Kernel (32-bit)	4.0.0
otherlinux64guest	Linux (64-bit) (experimental)	4.0.0

Understanding HCX vMotion and Cold Migration

HCX integrates with ESXi to perform vMotion migrations for powered on virtual machines, and with Cold Migration for powered off virtual machines.

HCX vMotion

VMware HCX vMotion can transfer a live Virtual Machine from a HCX activated vCenter Server to a HCX activated destination site (or from the destination site towards the local site). The vMotion transfer captures the virtual machine active memory, its execution state, its IP address, and its MAC address. The migration duration depends on the connectivity, including both the bandwidth available and the latency between the two sites.

HCX Cold Migration

Cold migration uses the same network path as HCX vMotion to transfer a powered-off virtual machine. During a cold migration, the Virtual Machine IP address and the MAC address are preserved. Cold migrations must satisfy the vMotion requirements.

Requirements and Limitations for HCX vMotion and Cold Migration

- HCX Interconnect tunnels must be up/active.
- HCX vMotion requires 150 Mbps or higher throughput capability.
- The virtual machine hardware version must be at least version 9 or higher.
- The underlying architecture, regardless of OS, must be x86.
- VMs with Raw Disk Mapping in compatibility mode (RDM-V) can only be migrated using Bulk Migration.
- VMware NFC is used as the primary protocol during Cold migration and as a secondary protocol during HCX vMotion.

  Note HCX Interconnect uses a Network Profile configuration dedicated to the vMotion traffic. This configuration does not include the Cold and vMotion NFC traffic. HCX always uses its Management interface for Cold and vMotion NFC traffic. In deployments where ESXi servers use a dedicated Provisioning vmkernel for NFC traffic, the HCX continues to route Cold and vMotion NFC traffic through the Management interface.

- Migration to or from vVOL datastores is not supported.
- The HCX Migration interface does not display port groups that are VLAN trunks.

For information regarding migration limits, search for VMware HCX in the VMware Configurations Maximums tool.

Virtual Machine Restrictions for HCX vMotion

Virtual machines with the following attributes are not supported for migration.

- Virtual machines cannot be migrated while they are using a shared SCSI bus, flagged for multi-writer, enabled for fault tolerance, or configured for shared VMDK disk sharing.
- Attached virtual media or ISOs.
- Virtual Machine Hardware Version 8 or earlier.

- Although concurrent VMware HCX vMotion migrations can be initiated up to the vSphere limits, VMware only supports serial HCX vMotion migrations between a source and the destination site. For simultaneous migrations in parallel, select Bulk Migration.
- HCX vMotion defaults to **Opportunistic** mode for per-VM vMotion Encryption if it is set to **Required**. During the migration operation - the mode is changed to Opportunistic during the migration initialization, and then set back to **Required** after the migration is completed.
- Virtual Machines with Change Block Tracking (CBT) can be migrated, but HCX deactivates CBT.
- VMware software appliances (including, but not limited to vCenter Server and NSX Manager) cannot be migrated with HCX vMotion migration.
- vSphere VM Encryption is not supported with HCX migrations.
- Virtual machine workloads with an attached serial port device are not supported. For these workloads, detach such devices prior to starting the migration.

  Note You can use Bulk Migration to migrate workloads with an attached serial port device.

- Virtual Based Security (VBS) functionality is not supported with VMware HCX migrations.

Understanding HCX Replication Assisted vMotion

HCX Replication Assisted vMotion (RAV) uses the HCX along with replication and vMotion technologies to provide large scale, parallel migrations with zero downtime.

HCX RAV provides the following benefits:

- Large-scale live mobility: Administrators can submit large sets of VMs for a live migration.
- Switchover window: With RAV, administrators can specify a switchover window.
- Continuous replication: Once a set of VMs is selected for migration, RAV does the initial syncing, and continues to replicate the delta changes until the switchover window is reached.
- Concurrency: With RAV, multiple VMs are replicated simultaneously. When the replication phase reaches the switchover window, a delta vMotion cycle is initiated to do a quick, live switchover. Live switchover happens serially.
- Resiliency: RAV migrations are resilient to latency and varied network and service conditions during the initial sync and continuous replication sync.
- Switchover larger sets of VMs with a smaller maintenance window: Large chunks of data synchronization by way of replication allow for smaller delta vMotion cycles, paving way for large numbers of VMs switching over in a maintenance window.

HCX RAV migration triggers the following events:

1. Replication begins with a full synchronization (replication) of the virtual machine disks to the destination site.
2. Migrated VMs enter a continuous synchronization cycle until a switchover is triggered.

3 Depending upon data churn on source disk, additional snapshots are being created during the RPO cycle. After each RPO cycle, disk consolidation takes place and creates a “hbrdisk.RDID vmdk” called as replica instance vmdk file on target datastore. Refer VMware KB 87028.

4 You can have the switchover process start immediately following the initial sync or delay the switchover until a specific time using the scheduled migration option. If the switchover is scheduled, the synchronization cycle continues until the switchover begins.

5 The final delta synchronization begins when the switchover phase starts. During this phase, vMotion is engaged for migrating the disk delta data and virtual machine state.

6 As the final step in the switchover, the source VM is removed, and the migrated VM is connected to the network powered on.

   Replication Assisted vMotion creates two folders at the destination site. One folder contains the virtual machine infrastructure definition, and the other contains the virtual machine disk information. This is normal behavior for RAV migrations and has no impact on the functionality of the virtual machine at the destination site.

Note In some cases, having two folders might impact other applications, such as backup tools, that require access to the virtual machines folders. If necessary, you can consolidate the contents of these two folders using VMware Storage vMotion.

Requirements for Replication Assisted vMotion

- HCX Interconnect tunnels must be up/active.
- HCX RAV requires 150 Mbps or higher throughput capability. See Network Underlay Minimum Requirements.
- The virtual machine hardware version must be Version 9 or higher.
- The underlying architecture, regardless of OS, must be x86.
- Hybrid Interconnect, Bulk Migration, vMotion, and Replication Assisted vMotion services must be activated and in a healthy state in the relevant Service Mesh.
- The resources to create, power on and use the virtual machine must be available in the destination environment.
- Virtual machines must reside in a Service Cluster (defined in the Compute Profile).
- RAV uses vSphere Replication whose potential throughput can vary depending on bandwidth available for migrations, latency, available CPU/MEM/IOPS, and disk read speed. For more information about how to determine bandwidth requirements, see Bandwidth Requirements for vSphere Replication.

- VMware NFC is used as a secondary protocol during HCX Replication Assisted vMotion migration.

  Note HCX always uses its Management interface NFC traffic. In deployments where ESXi servers use a dedicated Provisioning vmkernel for NFC traffic, HCX continues to route NFC traffic through the Management interface.

For information regarding the migration limits, search for VMware HCX in the VMware Configurations Maximums tool.

Restrictions for Replication Assisted vMotion

- Virtual machines with the following attributes are not supported for migration.
  - Virtual machines cannot be migrated while they are using a shared SCSI bus, flagged for multi-writer, enabled for fault tolerance, or configured for shared VMDK disk sharing.
  - Attached virtual media or ISOs.
  - Virtual Machine Hardware Version 8 or earlier.
- Live switchover of concurrent RAV migrations is run serially.
- HCX vMotion defaults to **Opportunistic** mode for per-VM vMotion Encryption if it is set to **Required**. During the migration operation - the mode is changed to Opportunistic on the migration initialization, and then set back to **Required** after the migration is completed.
- HCX Replication Assisted vMotion does not support the migration of workloads with Independent persistent and Independent non-persistent disks, as taking a snapshot of such virtual machine does not produce delta disks, which are required for the underlying RAV migration technology.
- Virtual machines with Raw Device Mappings (RDM) in Physical Compatibility mode cannot be migrated.
- Virtual machines with Raw Device Mappings in compatibility mode (RDM-V) cannot be migrated using RAV and vMotion.
- Virtual machines with DirectPath I/O configurations cannot be migrated without first removing the DirectPath device.
- To migrate FT-enabled VMs, temporarily turn off Fault Tolerance, and perform RAV. When this operation is complete, turn Fault Tolerance back on.
- Migration to or from vVOL datastores is not supported.
- The HCX Migration interface does not display port groups that are VLAN trunks.
- Virtual machines that cannot be gracefully powered off cannot be migrated. HCX can override with the Force Power-off VM option.
- Virtual Machines with Change Block Tracking (CBT) can be migrated, but HCX deactivates CBT.

- RAV migration to VMFS6 target datastores requires the following minimum vSphere version at the target site: vSphere 6.5U3f or vSphere 6.7U3.
- Virtual machines using virtual NVMe (vNVME) Controllers cannot be migrated.
- Virtual machine Snapshots cannot be migrated.
- VMware software appliances (including, but not limited to vCenter Server and NSX Manager) cannot be migrated with HCX RAV migration.
- vSphere VM Encryption is not supported with HCX migrations.
- Virtual machine workloads with an attached serial port device are not supported. For these workloads, detach such devices prior to starting the migration.

  Note You can use Bulk Migration to migrate workloads with an attached serial port device.
- Virtual Based Security (VBS) functionality is not supported with VMware HCX migrations.

Understanding VMware HCX OS Assisted Migration

The HCX OS Assisted Migration service uses the Sentinel software that is installed on Linux- or Windows-based guest virtual machines to assist with communication and replication from their environment to a VMware vSphere SDDC.

You must install HCX Sentinel on all guest virtual machines requiring migration using HCX OS Assisted Migration. Sentinel gathers the system configuration from the guest virtual machine and assists with the data replication. The source system information is used by various HCX OS Assisted Migration service processes. In part, the information is used to create an inventory of guest virtual machine systems for migration and to help replication processes prepare the disks on the replica virtual machine for replication and migration.

Sentinel also helps with the data replication by reading data that is written to the source disks and passing that data to the SDR appliance at the destination site.

Guest virtual machines connect and register with an HCX Sentinel Gateway (SGW) appliance at the source site. The SGW then establishes a forwarding connection with an HCX Sentinel Data Receiver (SDR) appliance at the destination vSphere site. You specify the network connections between the guest virtual machines and SGW in the compute profile.

You must install the HCX Sentinel software on each guest virtual machine requiring migration to initate the guest virtual machine discovery and data replication. After Sentinel is installed, a secure connection is established between the guest virtual machine and the HCX SGW. HCX builds an inventory of candidates for migration as the Sentinel software is installed on the guest virtual machines.

Using the established connection between the SGW and SDR, replication connections are made between the Sentinel software on the guest virtual machines and the SDR, with one connection each for control operations and data replication.

An OS Assisted Migration request triggers the following events:

1 Replication begins a full synchronization transfer to the destination site. The guest virtual machine remains online during replication until the final delta synchronization.

2 Before the final delta synchronization, the OS Assisted Migration service quiesces the guest virtual machine.

   The OS Assisted Migration service quiesces the guest virtual machine on a best-effort basis. For example, it is possible for a Linux service running on the guest virtual machine to start immediately after OS Assisted Migration has quiesced the services and stopped all known processes. If some process starts after quiescing the system, it can potentially lead to final synchronization not completing and appear as though the switchover process is stuck.

   Note As part of both continuous and final synchronization, the system checks for changed blocks or files and replicates them.

   - On a Windows system, the OS Assisted Migration service reads the entire disk to determine the changed blocks to replicate. The process of reading the entire disk can be time consuming.
   - On a Linux system, the OS Assisted Migration service walks the entire file system to determine the changed files to replicate.

3 When the delta synchronization finishes, a switchover is triggered. You can have the switchover process start immediately following the initial sync or delay the switchover until a specific time using the scheduled migration option. By using the scheduled migration option, the switchover can occur during a maintenance window. The final delta synchronization begins when the switchover phase starts.

   During scheduled migrations, HCX Sentinel performs continuous synchronization by transferring only the deltas since the previous sync cycle. For Windows HCX Sentinel, this synchronization is achieved by identifying the changed file system blocks, whereas for Linux HCX Sentinel this synchronization is achieved by monitoring the changed files. To improve time it takes to reach that final consistency point for Linux systems, a pre-determined set of files and directories listed in `/opt/vmware/hcx/osam/excluded_paths` is excluded from the continuous synchronization. If you have additional files that do not require monitoring, you can exclude them from continuous synchronization by editing the file. Excluding files requires a restart of the Sentinel service named `vmware-hcx-osam-sentinel` using service or systemctl commands.

   Note Excluded files are always synchronized to the target virtual machine during the initial and final synchronization phases.

4 HCX performs a hardware mapping of the replicated volumes to ensure proper operation, including updates of the software stack on the replica. This fix-up process includes adding drivers and modifying the OS configuration files at the destination. The migrated virtual machine reboots during this process.

Note When migrating Windows systems, HCX OS Assisted Migration software creates a temporary local user on the migrated Windows system during the switchover phase. This user gets deleted after the fix-up process is completed.

Note When migrating Linux systems, HCX OS Assisted Migration software uses an independent software stack residing on a separate disk, called the fix-up disk, for the fix-up process. To fix-up 64-bit workload VMs, HCX uses the Photon 3.0 64-bit version of the fix-up disk. The fix-up boot disk is detached and deleted at the end of the switchover process.

5 As the final step in the switchover, the source is powered-off, the migrated replica is connected to the network, and the switchover completes.

The vSphere target virtual machine reboots twice during the switchover phase.

If the synchronization process fails for any reason, such as a broken network connection, by default the synchronization is retried for eight hours. To improve the time it takes to reach a final consistency point for Switchover to begin, you can shorten the retry period to as little as one hour by editing the file `/opt/vmware/hcx/osam/etc/sync.params` and setting the **max_retry_interval** from one to eight hours. After setting the interval, restart the Sentinel service named `vmware-hcx-osam-sentinel` using service or systemctl commands.

6 HCX Manager names the replica virtual machine with the host name of the source virtual machine.

7 VMware Tools is installed on the migrated virtual machine and migration completes.

8 If the source does not power off, an attempt is made to power off the replica virtual machine.

If the replica virtual machine successfully powers off, it remains connected to the NICs. In this case, you can manually power off the source and power on the replica. If the replica does not power off, both the guest virtual machine and the replica remain on, but the replica is not connected to the network. In this case, you activate the NICs manually that are attached to the replica virtual machine using vCenter, power-off source virtual machine (if not already), and power-on Migrated virtual machine.

Considerations for OSAM Deployment

The OS Assisted Migration (OSAM) service includes several components that work together for connecting and forwarding guest workloads in the source environment.

Refer to the following considerations when deploying and operating OS Assisted Migration in your environment.

- The OSAM service converts non-vSphere guests to vSphere virtual machines. This conversion process involves halting OS services to quiesce the guest virtual machine. The downtime for this conversion process can vary from minutes to hours depending on a virtual machine size and activity.
- HCX deployments for OSAM migrations assume that there is (at minimum) an HA-enabled vSphere cluster.
- The HCX Sentinel Agent encrypts all connections to the Sentinel Gateway. The encryption cannot be deactivated.
- The HCX Sentinel Gateway must be deployed in a vSphere environment, and not within KVM or Hyper-V.
- To use HCX Network Extension with OSAM deployments, VLANs in the non-vSphere environment must also exist on the source vSphere environment network switches and in vCenter Server as Distributed Port Groups.
- When the KVM or Hyper-V environment is collocated in the same data center as the destination environment, it is an option to deploy the HCX Connector and source Service Mesh components at the destination vCenter Server.
- Each Service Mesh deploys one Sentinel Gateway (SGW) and its peer Sentinel Data Receiver (SDR), and supports up to 50 active replica disks.
- HCX OSAM deployments support 200 concurrent VM disk migrations across a four Service Mesh scale out deployment. In this Service Mesh scale out model for OSAM, the HCX Sentinel download operation is presented per Service Mesh.
- Guest virtual machines can only be migrated to a datastore that is accessible by the SDR.
- Redeployment of SGW and SDR appliances is not allowed when any migration is in-progress.
- Only "thin" and "thick" disk provisioning types are supported as the disk provision type for the migrated system. The "Same as Source" option is not supported.
- OSAM Migration using PowerCLI:
  - PowerCLI 11.5 is not supported.
  - Migrations with VMware Cloud Director as target are not supported using PowerCLI.
  - Mobility Group migration is not supported through PowerCLI for VMware Cloud Director (all services).
- The OSAM migration service applies the default storage policy to the migrated VMs and their disks. Currently, the OSAM service does not support a user-selected storage policy.
- Changes to source Guest virtual machine configurations while a migration is in progress might not take effect in migrated virtual machines and sometimes might lead to migration failure.
- VMs with Linux GuestOS can have up to 15 hard disks attached.

- VMs with Windows GuestOS can have up to 64 volumes or file-systems.

For information regarding migration limits, search for VMware HCX in the VMware Configurations Maximums tool.

Supported Guest Operating Systems

The OS Assisted Migration service supports migration of virtual machines running non-vSphere guest operating systems in Linux or Windows environments.

Linux Environments

HCX supports various Linux-based 64-bit guest operating systems on KVM or Hyper-V hypervisors.

Supported Linux OS versions on KVM Hypervisor (BIOS and EFI)	Supported Linux OS versions on Hyper-V Hypervisor (BIOS and EFI)
CentOS 6.1 - CentOS 6.10	RHEL 7.1 - RHEL 7.9 (BIOS/GEN-1 & UEFI/GEN-2)
RHEL 6.1 - RHEL 6.10	RHEL 6.4 - RHEL 6.10 (BIOS/GEN-1 Only)
CentOS 7.1 - CentOS 7.9	CentOS 7.0 - CentOS 7.9 (BIOS/GEN-1 & UEFI/GEN-2)
RHEL 7.1 - RHEL 7.9	CentOS 6.4 - RHEL 6.10 (BIOS/GEN-1 Only)
RHEL 8.0 - RHEL 8.8	RHEL 8.0 - RHEL 8.8 (BIOS/GEN-1 & UEFI/GEN-2)
CentOS 8.0 - CentOS 8.4	CentOS 8.0 - CentOS 8.4 (BIOS/GEN-1 & UEFI/GEN-2)
Ubuntu 14.04 LTS	Ubuntu 14.04 LTS (BIOS/GEN-1 & UEFI/GEN-2)
Ubuntu 16.04 LTS	Ubuntu 16.04 LTS (BIOS/GEN-1 & UEFI/GEN-2)
Ubuntu 18.04 LTS	Ubuntu 18.04 LTS (BIOS/GEN-1 & UEFI/GEN-2)

Windows Environments

HCX supports various 64-bit Windows guest operating systems on KVM or Hyper-V supervisors.

Supported OS versions on KVM Hypervisor	Supported OS versions on Hyper-V Hypervisor
Windows Server 2019	Windows Server 2019
Windows Server 2016	Windows Server 2016
Windows Server 2012	Windows Server 2012
Windows Server 2012 R2	Windows Server 2012 R2
Windows Server 2008 R2 (64-bit)	Windows Server 2008 R2 (64-bit)
Windows Server 2008 SP2 (64-bit)	Windows Server 2008 SP2 (64-bit)

Guest Operating System Considerations

The OS Assisted Migration service supports a variety of hypervisors and guest operating systems in both Linux and Windows environments with limitations and requirements that are both general and specific to these environments.

General Operating System Considerations

Some service limitations are common to Linux and Windows environments where OS Assisted Migration is deployed.

- The HCX Sentinel Agents are installed in the guest operating system and automatically make connections to the HCX Sentinel Gateway.
- Guest virtual machines with the locale and the UI language other than English US are not supported.
- UEFI-based source systems using legacy BIOS boot mode are not supported.
- Anti-virus software, or any OS application that is actively accessing file systems, can significantly delay the OSAM switchover phase. It is a best practice to deactivate these types of applications prior to configuring the OSAM migration.
- Network file shares are not supported with OSAM Migrations. This includes NFS, SMB, and CIFS, with varying outcomes:
  - Mounted files shares like CIFS may result in migration failures.
  - NFS shares are ignored during the fix-up phase and do not result in migration failure.

Linux Specific Considerations

- Block devices (partitions) with unrecognized content will not be migrated to the destination.
  - Unsupported file systems (supported file systems: ext2, ext3, ext4, XFS).
  - Unmounted file systems (Linux specific).
  - Unknown content with a partition or a block device.
  - Encrypted file systems.
  - md devices (software RAID).
- Statics routes are not supported.
- VLAN interfaces are not supported.
- On RHEL/CentOS 7.0, 7.1, and 7.2, the XFS file system UUID is not restored to the original UUID for the file system where /boot resides because mkfs.xfs does not support the functionality (-m option). A new random UUID is generated. Modifying the UUID after the file system is created triggers RHEL bug 1579390.
- /etc/fstab entries for removable media (floppy, CD) are not supported; such entries must be commented-out before migration.
- When migrating Linux systems, the HCX OSAM software uses an independent software stack residing on a separate disk for the fix-up process. This fix-up boot disk is detached and deleted at the end of the switchover process.

- The configuration files of Linux system services, such as dhcpd, that reference network interface names are not modified. You must manually modify these files on the migrated system.
- Device names in bonded/teamed interface configuration files are not replaced with the new device names, and this can lead to the bonded interface not showing as UP. You must make sure that the bonded/teamed interface device names match the new name after migration. You might encounter this issue if you use Network Manager tools like nmcli to configure bonded/teamed interfaces.

Windows Specific Considerations

- No support for syncing logical volumes.
- Basic MBR and GPT disk partitions are supported. Dynamic GPT and Dynamic MBR disks partitions are not supported.
  - If the boot disk is dynamic, migration is not supported.
  - If the data disks are dynamic, the data on the disks is not migrated, and the disks appear as raw disks on the migrated system.
- Only NTFS formatted volumes are supported. ESP (EFI system partition) with the FAT32 file system is an exception.
- No support for non-Windows service applications from the system quiescing perspective during the final sync.
- Systems with more than 64 volumes are not supported since VSS allows a maximum of 64 snapshots on a system.
- Any VSS snapshots present on the source Windows system before migration are not usable on the migrated system.
- VLAN interfaces are not supported.
- For Windows systems, in general the pre-requisites for the VMware Tools installation have to be satisfied on the source system. The pre-requisites for the VMware Tools installation can vary based on the target VC and ESX version, and Windows OS version on the source system. For example, if the target ESXi version is 6.5.0 (or higher), VMware Tools version is 10.3.x. To view the list of prerequisites based on different Windows OS versions see VMware KB 55798.
- When migrating Windows systems, the HCX OS Assisted Migration software creates a temporary local user on the migrated Windows system during the switchover phase. This user gets deleted after the fix-up process is completed.
- Do not run Windows updates on the source system during a migration. If Windows updates are in progress, the migration can fail.
- A Windows source system configured as a failover cluster node is not supported.

- To migrate a Windows Server 2008 SP2 system using OSAM, you must pre-install a version of VMware Tools older than 11.1 on the source system before trying out the migration. VMware tools 11.1 or newer will not install on Windows Server 2008 SP2 as per the KB article 75163.

Understanding Workload Migrations for NSX V2T

Working together with NSX Migration Coordinator for V2T Migration, the Workload Migrations for NSX V2T service supports the lift-and-shift of virtual machine workloads from NSX for vSphere environments to more flexible and feature rich NSX environments.

Workload Migrations for NSX V2T combines the wave orchestration and operational capabilities of HCX with the performance and the concurrency characteristics of HCX Assisted vMotion to expedite virtual machine migration. Workload Migrations for NSX V2T supports the lift and shift of NSX for vSphere workloads.

HCX Integration with NSX Migration Coordinator for V2T Migration

Workload Migrations for NSX V2T integrates with NSX Migration Coordinator to migrate seamlessly virtual machines and correctly maintain their security state. This integration is applicable when running User Defined Topology Configuration and Edge Migration mode from an NSX for vSphere cluster to an NSX cluster.

Note In clustered NSX environments, you must register HCX Connector with the node where NSX Migration Coordinator has been initiated.

Workload Migrations for NSX V2T leverages the bridging capability created by NSX Migration Coordinator for virtual machine connectivity on the logical segment during the transition from NSX for vSphere.

vCenter Server Topologies

The Workload Migrations for NSX V2T service is designed for specific migration topologies within the data center:

- Single vCenter Server cluster to cluster migrations (NSXv-cluster to NSXT-cluster).
- Cross-vCenter Server migration (NSXv-VC to NSXT-VC).
- Federated NSX architectures.

In the single vCenter Server migration topology, the HCX Connector Compute Profile configuration must contain the NSX-V prepared source cluster while the HCX Cloud Manager Compute Profile configuration contains the NSX-T prepared cluster.

For all topologies, HCX registration with the vCenter Server and local NSX Manager remains the same for both the HCX Connector and HCX Cloud Manager. This registration is done during HCX Manager installation.

For Federated NSX architectures, apart from registering the HCX Manager with the local NSX-T manager at installation, you explicitly configure HCX Cloud with the Global Manager that is running the V2T Migration Coordinator in addition to the Local Manager. See Configure NSX V2T Migration for Federated NSX Architectures.

How it works

The Workload Migrations for NSX V2T service workflow has three general phases: pre-migration workload configuration, HCX migration, and post migration clean-up. NSX Migration Coordinator for V2T Migration handles the pre-migration and the post-migration operations, which are outside the scope of this document. For information regarding the NSX Migration Coordinator configuration, refer to the section "End-to-end Workflow of Configuration and Edge Migration," in Configuration and Edge Migration Workflow.

The migration phase for Workload Migrations for NSX V2T is similar to other HCX migration types except that no target network mapping in the HCX UI is necessary for the workloads. One exception is standard port groups, which require the network mapping to be done manually in the HCX Migration wizard. Workload Migrations for NSX V2T obtains the network mapping information through API calls to NSX Migration Coordinator.

Note All API calls to Migration Coordinator are at the group level. For example, a request for mapping information gets the information for all VMs in the group. Any changes to individual VMs in the group after this API request are not allowed.

You initiate the Workload Migrations for NSX V2T service from the HCX Connector migration interface by selecting **V2T-Migration** at the group level.

The Workload Migrations for NSX V2T service has the following high-level workflow:

1. Complete the pre-migration configuration steps using NSX Migration Coordinator for V2T Migration.
2. In the HCX Connector migration interface, create Mobility Groups for the selected workloads.
3. Select the **V2T-Migration** for the Mobility Group and configure available Transfer, placement, and extended options.
4. Validate the configuration. At this point, Workload Migrations for NSX V2T makes API calls to Migration Coordinator to identify the source and target network mapping and identify the incoming groups of VMs.
5. Start the migration.
6. HCX begins consuming the pre-configuration migration settings for the VMs and begins migrating VMs.

7 HCX makes post migration API calls to the NSX Migration Coordinator to perform the post-migration clean-up.

Note If at any point the NSX Migration Coordinator fails or restarts, you must re-enter the NSX for vSphere configuration details on the NSX Migration Coordinator before continuing with Workload Migrations for NSX V2T migrations.

If the NSX Migration Coordinator fails or restarts prior to completing post migration operations, then you must manually run the post-migrate NSX Migration Coordinator API call using the group ID from the migration dashboard. See VMware KB article.

Requirements and Limitations

- NSX for vSphere and NSX environments must be explicitly joined using NSX bridging before migrating workloads. HCX Network Extensions cannot be used between NSX for vSphere and NSX environments. Bridge set up is through the NSX Migration Coordinator.
- NSX for vSphere and NSX environments must be accessible to HCX Manager.
- A dedicated vSphere vMotion data path must be set up between the source and the target ESXi hosts.
- Any changes to the source environment after the NSX V2T migration can affect Workload Migrations for NSX V2T.
- The Distributed Virtual Switch version at the source and the destination sites must be the same.
- Workload Migrations for NSX V2T works in one direction. Reverse migration is not possible using Workload Migrations for NSX V2T.

  Note After Workload Migrations for NSX V2T has completed, you can use Bulk migration to move the VMs back to the source site. For reverse migration using Bulk in single vCenter environments, you must specify a different folder name as the **Destination Folder**.
- Standard switches at the source and target clusters must have matching security policies and the teaming and failover policies.
- For Federated NSX environments, the HCX Cloud Manager must be configured with the Global Manager that is running the V2T Migration Coordinator, in addition to the Local Manager.

Restrictions

Workload Migrations for NSX V2T does not support:

- HCX WAN Optimization.
- HCX Network Extension.
- Adding VMs to an HCX Mobility Group after initiation of Workload Migrations for NSX V2T.

- Adding, deleting, or editing a network from a VM after the initiation of Workload Migrations for NSX V2T.
- Retry of failed or canceled Workload Migrations for NSX V2T migration using the Edit Mobility Group wizard is not supported. A new mobility group must be created for these VMs.
- Migration of suspended VMs.
- Migration from HCX vCenter plug-in. Use the HCX standalone UI to trigger Workload Migrations for NSX V2T.
- Encrypted VMs.

Migrating Virtual Machines with HCX

Through the HCX Migration interface, you can configure multiple virtual machine migrations, including reverse migrations.

This section describes migration operations using HCX Advanced License functionality. For information about migration operations using HCX Mobility Groups, which is available as an HCX Enterprise License feature, see Migrating Virtual Machines with Mobility Groups.

Migrations are always configured using the HCX Connector or Cloud system that initiated site pairing. In Cloud-to-Cloud deployments with bi-directional site pairing, HCX in both paired sites can initiate migrations. For more information, see Adding a Site Pair.

The HCX system automatically detects virtual machine disk additions or removals and reconfigures running migrations to accommodate these changes. These disk changes are honored only if the changes occur before the migration switchover phase. If disk changes occur during the switchover phase, the changes are not recognized, which can affect the success of the migration operation. Support for adding or removing disks is available only with Bulk and Replication Assisted vMotion migrations.

Taking snapshots of a VM during migration, either manually or through a third-party backup solution, it can disrupt the migration process. To prevent any impact, it is required to stop those services that may create or remove snapshots during migration. Refer to KB79220 for more information.

For Bulk, vMotion, and RAV migration types, you can select a Service Mesh to use for migration operations. For each HCX site pair, you can have one or more Service Mesh configurations. Each Service Mesh configuration has a specific resource and network configuration. In some cases, you might choose a specific Service Mesh to use for migrations based on the Service Mesh configuration. If no Service Mesh is selected for a migration, HCX automatically selects the Service Mesh to use. To add a Service Mesh for a site pair, see Creating a Service Mesh.

Note The selectable Service Mesh option requires that both the source and destination sites run HCX 4.8 or later.

For the operational limits supported with HCX migrations, see Configuration and Service Limits.

Prerequisites

- The migration service is activated in both the source and destination site Compute Profile.
- The migration service is activated in the HCX Service Mesh.
- For RAV or OSAM migrations, the HCX Enterprise license is activated.

  Note No additional license or activation is required in HCX for VMware Cloud on AWS deployments.

- Sentinel software is installed on all guest virtual machines requiring OSAM migration. See Sentinel Management.

Procedure

1. Navigate to the HCX dashboard.
2. Select the **Services > Migration**.

   The Migrate Tracking window displays a summary of virtual machine migrations.

3. Select **Migrate Virtual Machines**.
4. Select the **Remote Site Connection**.

   The list of virtual machines available for migration appears in the display.

   Note For OSAM, select **Non vSphere Inventory** > **Remote connections** to display the list of guest virtual machines on which you installed HCX Sentinel.

5. (Optional) To display the list remote site virtual machines available for the reverse migration, click the **Reverse Migration** check box.

   Note Reverse migration is not applicable to OS Assisted Migration.

6. Select the virtual machines you want to migrate.

   Note Click **hide unselected** to keep only selected virtual machines on the screen.

7 Set the Transfer and Placement, Switchover, Interconnect Service Mesh, and Extended options.

- To apply default settings for all selected virtual machines, use the green area of the interface at the top of the window.
- To set machine-specific Transfer, Placement, Interconnect Service Mesh, and Switchover options, select a specific virtual machine and expand the entry.

Note

- For Bulk, RAV, and OSAM migrations, you can schedule the migration date and time as part of the Switchover settings. Scheduling vMotion migrations is not available.
- If the VM is powered off, Cold Migration is set by default.
- Extended Options provide additional settings based on the selected migration type.
- For additional information, see Additional Migration Settings.

8 Select the destination network for each virtual machine to be migrated.

In most cases, the stretched network between the source and destination sites is automatically selected. You can change this selection as needed.

a Expand each virtual machine selection.

b Next to each guest virtual machine NIC name, click the folder for a list of available target networks.

c Click the check box next to the network that you want the guest virtual machine to map to, and then click **Select**.

d (Optional) To specify a new guest OS IP address for the virtual machine at the target network, expand the NIC entry and enter the new IP address, gateway, and subnet mask.

9 Click **Finish**.

The HCX Manager validates your selections and starts the migrations. If a warning is generated, click **Finish** again to proceed.

Monitoring Migration Progress with HCX

The HCX Migration Tracking page displays a summary of migrations, reporting the status and progress of individual virtual machine migrations.

Procedure

1 In the HCX dashboard, select **Services** > **Migration**.

The Migration Tracking page provides a list of all ongoing or recent migrations.

2 To determine the migration status, review the Progress information.

While the migration is underway, the Progress column displays a progress bar with the percentage of replication completed for a specific virtual machine.

For Bulk and Replication Assisted vMotion (RAV) migrations, the Progress column includes a best-effort, real-time estimate of the amount of time remaining for the transfer phase of a specific virtual machine. For vMotion migrations, this includes real-time estimates regarding the relocation phase. This estimate is based on sampling the underlying metrics of the environment, such as bytes transferred, rate of transfer, network throughput, and number of disks. Changes in the underlying metrics can impact the estimate. The interval between estimates varies with the size of the migrated virtual machine:

Virtual Machine Size	Estimate Interval
Less than 50 GB	Every 1 1/2 minutes.
Greater than or equal to 50 GB but less than 1 TB	Every 5 minutes.
Greater than or equal to 1 TB	Every 15 minutes.

Note The estimate interval begins after the system has gathered the underlying metrics and completed calculations, meaning there may appear to be a delay in presenting the initial estimate. For relatively small transfers, the transfer may complete before providing an estimate.

3 To sort the information in the list, use the filter option provided in each column of the display.

You can use the search option at the top-right corner of the display to narrow down the list of migrations. You can search by virtual machine name, state message, migration type, or other attributes.

Migrating Virtual Machines with Mobility Groups

Mobility Groups is an HCX Enterprise License feature that supports assembling one or more virtual machines into logical sets, for execution and monitoring of migrations as a group.

With Mobility Groups, you have the flexibility to manage migrations for sets of virtual machines by application, network, pod, or other aspects of your environment.

Migrations are always configured using the HCX Connector or Cloud system that initiated site pairing. In Cloud-to-Cloud deployments with bidirectional site pairing, HCX in both paired sites can initiate migrations. For more information, see Adding a Site Pair.

The HCX system automatically detects virtual machine disk additions or removals and reconfigures running migrations to accommodate these changes. These disk changes are honored only if the changes occur before the migration switchover phase. If disk changes occur during the switchover phase, the changes are not recognized, which can affect the success of the migration operation. Support for adding or removing disks is available only with Bulk and Replication Assisted vMotion migrations.

For Workload Migration for NSX V2T, you can add additional VMs to the group and change configuration settings on a VM only while all VMs in the group are in **Draft** state. If any one VM in the group starts migration, you can make no further changes to any of the VMs that are members of that group.

Taking snapshots of a VM during migration, either manually or using a third-party backup solution, can disrupt the migration process. To prevent any impact, it is required to stop those services that may create or remove snapshots during migration. Refer to KB79220 for more information.

For Bulk, vMotion, and RAV migration types, you can select a Service Mesh to use for migration operations. For each HCX site pair, you can have one or more Service Meshes. Each Service Mesh configuration has a specific resource and network configuration. In some cases, you might choose a specific Service Mesh to use for migrations based on the mesh configuration. If no Service Mesh is selected for a migration, HCX automatically selects the Service Mesh to use. To add a Service Mesh for a site pair, see Creating a Service Mesh.

Note The selectable Service Mesh option requires that both the source and destination sites run HCX 4.8 or later.

For the operational limits supported with HCX migrations, see Configuration and Service Limits.

Prerequisites

- General:
  - The HCX Enterprise license is activated.
  - The migration service is activated in both the source and destination site Compute Profile.
  - The migration service is activated in the HCX Service Mesh.
  - Information regarding each migration type is reviewed and understood. For more information, refer to VMware HCX Migration Types.
- Workload Migration for NSX V2T:
  - NSX Migration Coordinator pre-migration configuration is done prior to configuring Mobility Groups for Workload Migration for NSX V2T. For more information, see "End-to-end Workflow of Configuration and Edge Migration," in Configuration and Edge Migration Workflow.
  - For NSX V2T migrations in a Federated NSX environment, prior to performing migrations you must configure (register) HCX with the Global Manager running the Migration Coordinator in addition to the Local Manager. See Configure NSX V2T Migration for Federated NSX Architectures.

    Note By configuring HCX with the NSX Global and Local managers, you can migrate to or from an NSX Federated environment using all migration types: Bulk, Replication Assisted vMotion (RAV), vMotion.

Procedure

1 In the HCX Manager UI, navigate to **Services > Migration**.

  The Migration Management interface displays a summary of groups. For detailed group information, you can expand each group.

2 Click **Migrate** and select **Remote Site Connection**.

The Workload Mobility interface displays a list of virtual machines (workloads) that are available for migration and that can be added to a group. You can select the **Networks** or **Hosts and Clusters** icon to update the list of virtual machines. Alternatively, you can use a regular expression search to filter the list of virtual machines by name.

Note If you have only one site pair, it is selected by default. For OS Assisted Migrations, select **Non vSphere Inventory > Remote connections** > to populate the list of guest virtual machines on which you installed HCX Sentinel.

3 (Optional) To display a list of remote site virtual machines available for the reverse migration, click the **Reverse Migration** check box.

Reverse migration refers to the migration of virtual machines from an HCX destination site to a source site.

Note Workload Migration for NSX V2T does not support reverse migration.

4 Specify a **Group Name**.

If no group name is provided, the system automatically assigns a five character identifier as the group name. You can change this name later by editing the group information. See Managing Migrations with Mobility Groups.

5 For Workload Migration for NSX V2T, check the box labeled **V2T-Migration**.

6 Select the set of virtual machines to include in the group and click **ADD**.

You can add additional virtual machines to the group at any time.

Note For Workload Migration for NSX V2T, you can add additional VMs to the group and change configuration settings on a VM only while all VMs in the group are in **Draft** state. If any one VM in the group starts migration, you can make no further changes to any of the VMs that are members of that group.

7 Select the group Transfer and Placement, Switchover, Interconnect Service Mesh, and Extended options.

The settings you provide are applied to all members of the group by default. To override the default settings for specific virtual machines in the group, select and expand the virtual machine entry, and set different options. For additional information regarding these options, see Additional Migration Settings.

Note

- For Workload Migration for NSX V2T, the Migration Profile is not selectable and this field is empty.
- You can set the switchover schedule for Bulk, Replication Assisted vMotion (RAV), or OS Assisted Migration (OSAM) migrations. Switchover scheduling is not available for vMotion or Workload Migration for NSX V2T.
- If the VM is powered off, Cold Migration is set by default.

  Note For Workload Migration for NSX V2, powered off VMs are also migrated as part of the group migration.
- Extended Options provide additional settings.

8 Select the destination network for each virtual machine to be migrated.

In most cases, the stretched network between the source and destination sites is automatically selected. You can change this selection as needed.

Important For Workload Migration for NSX V2T, the destination network is pre-configured by NSX Migration Coordinator and is ignored in the HCX Migration wizard and cannot be changed. VMs connected to standard port groups must be explicitly mapped to the target networks.

Note If you configure HCX with an NSX Global Manager for use with V2T migrations in a federated NSX topology, the Global Manager populates local NSX managers with all global network segments known to the Global Manager. These networks then become visible as Global Networks for Bulk, RAV, and vMotion migrations.

a Expand each virtual machine selection.

b Next to each guest virtual machine NIC name, click the folder for a list of available target networks.

c Click the check box next to the network that you want the guest virtual machine to map to, and then click **Select**.

d (Optional) To specify a new guest OS IP address for the virtual machine at the target network, expand the NIC entry and enter the new IP address, gateway, and subnet mask.

9 To complete the Mobility Group migration operation, select **Go**, **Validate**, **Save**, or **Close** to complete the Mobility Group migration operation:

Migration operation	Result
Go	Validates your virtual machine migration selections, saves the group, and then starts the migration.
Validate	Validates readiness of selected virtual machine for migration without starting the migration. Validation can be done at any time on selected virtual machines or a group.
Save	Saves migration selections as drafts for future editing or scheduling without starting the migration.
Close	Cancels your selections without creating a group or starting a migration.

What to do next

HCX provides multiple options for monitoring and managing migrations. To review the activity, progress, and history of your migrations, refer to these topics:

- Monitoring and Estimating Migration with Mobility Groups
- Managing Migrations with Mobility Groups
- Viewing HCX Migration Event Details

Monitoring and Estimating Migration with Mobility Groups

The HCX Migration Management interface provides a summary of group migration progress, the progress of individual virtual machines in the group, migration phase transfer estimates for Bulk and Replication Assisted vMotion (RAV) migration types, and relocation phase transfer estimates for vMotion migrations.

For all migration types, the Migration Management page displays the status and progress of virtual machine migrations. While migrations are underway, the Progress column displays a progress bar with the percentage of migration completed for the group, along with the number of migrations completed. Expanding a group entry displays information about each virtual machine in the group.

For Bulk and Replication Assisted vMotion (RAV) migrations, information includes real-time estimates regarding the migration transfer phase. For vMotion migrations, this includes real-time estimates regarding the relocation phase. Real-time estimation gathers information from the current migration environment. This estimate is based on sampling the underlying metrics of the environment, such as bytes transferred, rate of transfer, network throughput, and number of disks. Changes in the underlying metrics can impact the estimate. The interval between estimates varies with the size of the migrated virtual machine:

Virtual Machine Size	Estimate Interval
Less than 50 GB	Every 1.5 minutes
Greater than or equal to 50 GB but less than 1 TB	Every 5 minutes
Greater than or equal to 1 TB	Every 15 minutes

Note The estimate interval begins after the system has gathered the underlying metrics and completed calculations, meaning there may appear to be a delay in presenting the initial estimate. For relatively small transfers, the transfer may complete before providing an estimate.

For Bulk and RAV migrations, HCX provides predictive estimates for Mobility Groups in Draft state. Predictive estimation uses historical data from completed migrations in combination with machine learning (ML) to estimate the transfer phase.

To view group or individual migrations, review these steps:

Prerequisites

Predictive estimation has the following requirements:

- At least 50 successfully completed migrations. Predictive estimates for both forward and reverse migrations requires at least 50 successful migrations in each direction.
- The Mobility Group is in Draft state.
- The Mobility Group has been validated successfully.
- Migration prediction is available only for datastores and IX appliances that have participated in the previous migrations to train the ML model.

Procedure

1 Navigate to **Services > Migration > Management** .

   The Migration Management window displays Mobility Group information for each site pair. The window displays both forward and reverse migration information.

2 Review the progress column to view the status or see estimates related to migration.

   For detailed information, expand each virtual machine migration entry in the group.

   - For Bulk and Replication Assisted vMotion migrations, the Progress column provides a best-effort, real-time estimation of the time required to complete the transfer phase for a specific virtual machine. For vMotion migrations, the Progress column provides a real-time estimate to complete the relocation phase.

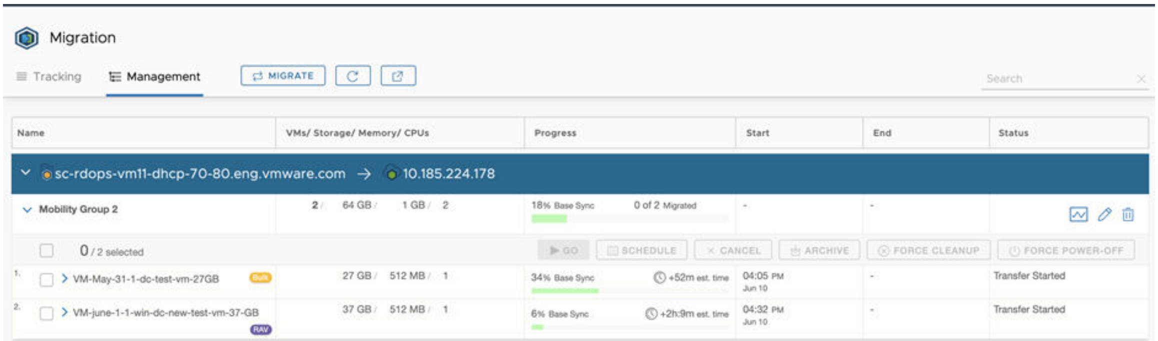

Name	VMs/ Storage/ Memory/ CPUs	Progress	Start	End	Status
sc-rdops-vm11-dhcp-70-80.eng.vmware.com → 10.185.224.178					
Mobility Group 2	2 / 64 GB / 1 GB / 2	18% Base Sync 0 of 2 Migrated	-	-	
□ 0 / 2 selected					
1. □ VM-May-31-1-dc-test-vm-27GB Bulk	27 GB / 512 MB / 1	34% Base Sync +52m est. time	04:05 PM Jun 10	-	Transfer Started
2. □ VM-june-1-1-win-dc-new-test-vm-37-GB RAV	37 GB / 512 MB / 1	6% Base Sync +2h:9m est. time	04:32 PM Jun 10	-	Transfer Started

- For groups of Bulk or RAV migration entries that are in Draft state, the Progress column optionally provides a predictive estimate of the time required to complete the transfer phase for a specific virtual machine. To initiate predictive estimates on the draft entries in the group, click the Get Predictions icon from the group management selections after expanding the group.

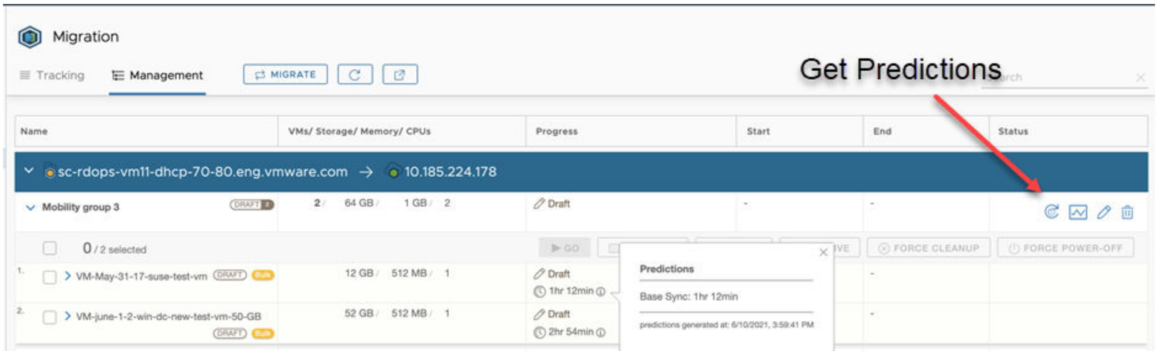

Get Predictions

Name	VMs/ Storage/ Memory/ CPUs	Progress	Start	End	Status
sc-rdops-vm11-dhcp-70-80.eng.vmware.com → 10.185.224.178					
Mobility group 3 DRAFT 2	2 / 64 GB / 1 GB / 2	Draft	-	-	
□ 0 / 2 selected					
1. □ VM-May-31-17-suse-test-vm DRAFT Bulk	12 GB / 512 MB / 1	Draft 1hr 12min		-	
2. □ VM-june-1-2-win-dc-new-test-vm-50-GB DRAFT Bulk	52 GB / 512 MB / 1	Draft 2hr 54min		-	

Predictions

Base Sync: 1hr 12min

predictions generated at: 6/10/2021, 3:59:41 PM

For each estimate in the Progress column, an information icon provides the prediction details and a timestamp for the most recent prediction.

3 To display a list of all virtual machine migrations and migration progress, click the **Tracking** tab.

The Migration Tracking page provides a list of all ongoing or recent migrations regardless of group.

You can sort the tracking information using the filter option provided in each column heading. You can search by virtual machine name, state message, migration type, or other attributes.

4 To return to the Migration Management window, click the **Management** tab.

Managing Migrations with Mobility Groups

From the HCX Migration interface, you can edit any group, delete groups, initiate and stop migrations, and schedule migrations.

Procedure

1 Navigate to **Services > Migration**.

The Migration Management window displays a summary of Mobility Group information for each site pair. The window displays both forward and reverse migration information.

Note You can switch between Migration Management (group migration) and Migration Tracking (individual migration) displays at any time using the menu button.

2 From the Migration Management window, you can edit or delete any group.

Mobility Group Operation	Description
Edit Group	To display the Workload Mobility window, click this option. From this window, you have several options: ■ Add additional virtual machines to the group. **Note** For Workload Migration for NSX V2T, you can add additional VMs to a group and change configuration settings on a VM only while all VMs in the group are in **Draft** state. If any one VM in the group begins the migration, you can make no further changes to any of the VMs that are members of that group. ■ Change the default migration profile for the group. Not supported for Workload Migration for NSX V2T. ■ Change the migration profile of individual virtual machines. Not supported for Workload Migration for NSX V2T. ■ Delete a specific virtual machine from the group. For Workload Migration for NSX V2T only when all workloads are in Draft state. ■ Restart failed or canceled migrations. **Note** For Workload Migration for NSX V2T, you can retry failed or canceled migrations by adding them to a new group.
Delete Group	To delete a group entry, click this option. You can delete a group only when all entries in the group are in the Draft state.

3 To show information about all members of the group, expand the group.

The system displays a list of virtual machines in a group with their migration status.

4 From the expanded group, you can also start, cancel, schedule, or archive one or more **selected** migrations.

Mobility Group Operation	Description
Go	Validates your virtual machine migration selections and then prompts you to start the migration. After the migration starts, the migration progress changes with each phase of the migration.
Schedule	Provides an option to reschedule a switchover for the migration. You can schedule Bulk, Replication Assisted vMotion (RAV), or OS Assisted Migration (OSAM) migrations. You cannot schedule vMotion or Workload Migration for NSX V2T.

Mobility Group Operation	Description
Cancel	Cancels a migration that is in progress. For information about the effects of canceling an OSAM migration, see Canceling a Migration. For information about cleaning up failed or canceled migrations, see Managing Failed or Canceled Migrations.
Archive	Clears the migration entry from the display. Use the **Archive** option to clear failed, canceled, and completed migration activity. Clearing the migration history updates the HCX Dashboard migration counters but does not remove the migration-related details from the HCX log files.
Force Cleanup	Following a failed or canceled migration, you can use Force Cleanup to clear internal operations and processes from source and destination sites. For information, see Managing Failed or Canceled Migrations.
Force Power-off	For Bulk migration only, this option allows for powering off a source virtual machine if the guest shutdown fails during switchover. Not supported for Workload Migration for NSX V2T.

Configure NSX V2T Migration for Federated NSX Architectures

To support Workload Migrations for NSX V2T in a Federated NSX environment, you must explicitly configure HCX with the Global Manager running the NSX Migration Coordinator in addition to the Local Manager.

In Federated NSX architectures, a single NSX Global Manager manages multiple local NSX deployments. For V2T migrations in this federated environment, the Migration Coordinator must be running on the Global Manager. In this case, you must configure HCX Cloud with the Global Manager that is running the Migration Coordinator.

Note The NSX Global Manager populates local NSX managers with all global network segments known to the Global Manager. These networks are flagged in HCX as Global Networks in the HCX inventory, and become available for use for Bulk, RAV, and vMotion migrations. These global segments, however, are not supported for HCX Interconnect configuration, meaning Network Profile and Compute Profile creation with Global Segments or Global transport Zones.

Procedure

1 Log in to the HCX Cloud appliance and log in using the admin user credentials.

```
https://hcx-ip-or-fqdn:9443
```

2 Navigate to the **Configuration** tab.

3 Select **NSX Migration Coordinator** and enter the Global Manager URL.

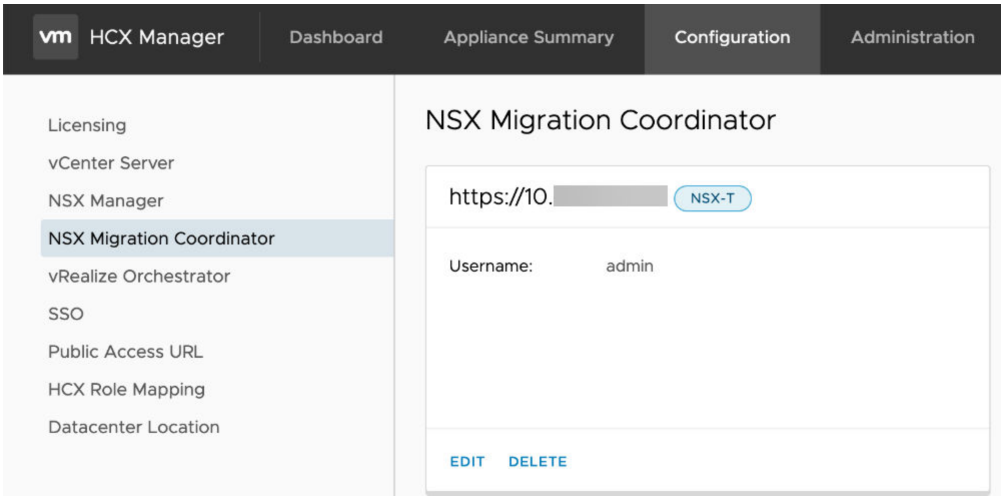

NSX Migration Coordinator

https://10.[redacted] NSX-T

Username: admin

4 Repeat this procedure for each HCX deployment that performs V2T migrations.

5 Navigate to the **Dashboard** tab and verify that the NSX Migration Coordinator appears in the display.

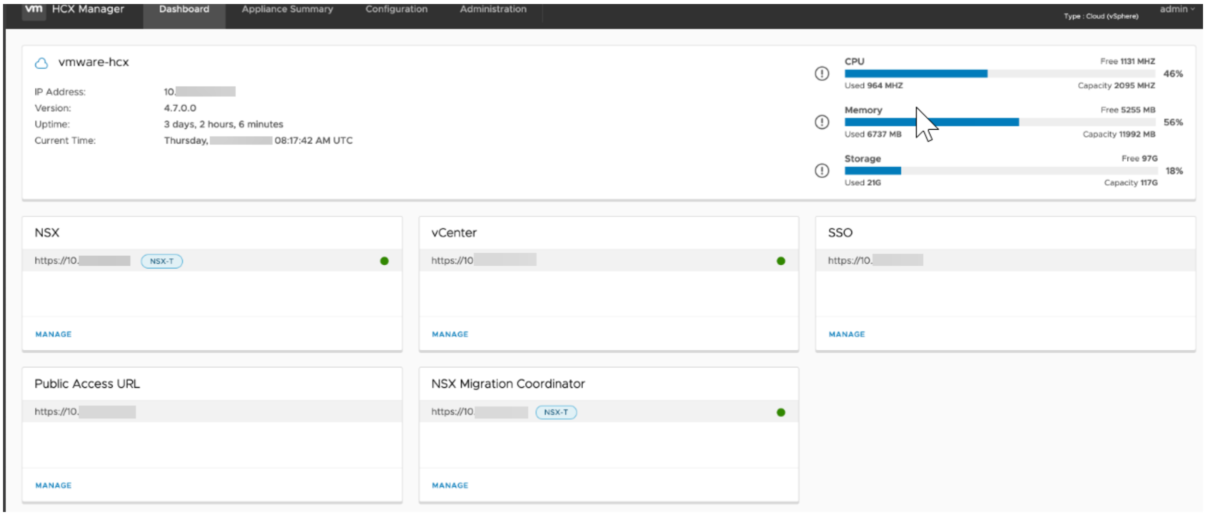

vmware-hcx

IP Address: 10.[redacted]
Version: 4.7.0.0
Uptime: 3 days, 2 hours, 6 minutes
Current Time: Thursday, [redacted] 08:17:42 AM UTC

CPU Free 1131 MHZ 46% Used 964 MHZ Capacity 2095 MHZ
Memory Free 5255 MB 56% Used 6737 MB Capacity 11992 MB
Storage Free 97G 18% Used 21G Capacity 117G

NSX https://10.[redacted] NSX-T

vCenter https://10[redacted]

SSO https://10.[redacted]

Public Access URL https://10.[redacted]

NSX Migration Coordinator https://10[redacted] NSX-T

What to do next

Return to the HCX Console to perform migration operations.

Additional Migration Settings

The HCX migration interface provides a set of options that can be used to tailor the behaviors and conditions of the virtual machine before or after the migration operation.

HCX has several types of optional settings for use when migrating virtual machines: Switchover, Virtual Machine, and Extended.

Note The available options depend on the selected migration type.

Switchover Options

Switchover options are applied during the switchover phase of the migration. Except for the Force Power-Off VM setting, which can be set on an individual virtual machine (VM) migration, these options are set globally for a group of selected VMs.

Force Power-Off VM

By default, HCX attempts to shut down the virtual machine guest gracefully during the HCX Bulk migration operation. If the OS interrupts the termination process, the migration operation fails. Checking this option causes HCX to force the power-off.

Remove Snapshots

Causes HCX to consolidate snapshot files during migration of the virtual machine. For more information, see Understanding Virtual Machine Snapshots During HCX Migrations.

Force Unmount ISO Images

Causes HCX to remove mounted ISO images before migrating the virtual machine.

Virtual Machine Options

The virtual machine (VM) options are available only for individual VM migration.

Enable Seed Checkpoint

In the event that a migration is unsuccessful or canceled, Seed Checkpoint retains the target disks created at the target site. Without Seed Checkpoint, the HCX roll back process cleans up the target disks created during the migration and all transferred data is lost. Seed Checkpoint is available with HCX Bulk Migration and Replication Assisted vMotion.

Note To resume from the migration checkpoint, select **Enable Seed Checkpoint** when you restart the migration. If the Seed Checkpoint option is not selected, the migration creates new disks at the target site and ignores the existing checkpoint.

To use this option, select a specific datastore for migration and not a datastore cluster.

Restarting a Seed Checkpoint migration to a new datastore creates additional target disks on that datastore, and saves new checkpoint information related to that datastore with the virtual machine data in HCX.

Extended Options

Extended options customize the characteristics of the migrated virtual machine with out having to manually update the settings after migration. Extended options can be set globally for a group of virtual machines migrations or for an individual virtual machine.

Retain MAC

Causes a virtual machine to keep its current MAC address during HCX Bulk migration operation, allowing communications to resume gracefully, and honors MAC-based security policies. This option is selected by default for vMotion and Replication Assisted vMotion migration types and cannot be changed.

Upgrade Virtual Hardware

Allows HCX to upgrade virtual machine hardware to the latest supported version as part of the migration operation, making current virtual machine hardware features immediately available to the migrated virtual machine.

Upgrade VMware Tools

Allows HCX to upgrade VMware Tools to the latest supported version as part of the migration operation, making current VMware Tools features immediately available to bulk migrated virtual machine.

Deactivate Per Virtual Machine EVC

vSphere Enhanced vMotion Compatibility (EVC) ensures that workloads can be live migrated, using vMotion, between ESXi hosts in a cluster that is running different CPU generations. EVC is a cluster level setting that supports virtual machine mobility within a cluster. For virtual machines implementing per-VM EVC, the EVC mode becomes an attribute of the virtual machine rather than the specific processor generation it happens to be booted on in the cluster. If this option is selected then this virtual machine might not be able to vMotion back to the source site.

Host Name

Sets the host name of the migrated virtual machines at the destination site.

Domain Name

Sets the domain name for the virtual machine at the destination site. This option is available only with Linux virtual machines.

Personalization Script

Uploads a customization script to a migrated virtual machine. The script runs before and after guest customization. The customization script cannot exceed 1500 characters.

For Linux VMs, Guest Customization using the Personalization Script option requires VMware Tools version 10.1.0 or later.

Note In the VMware Tools configuration, the **enable-custom-scripts** option is deactivated by default for security and must be enabled manually to run the custom script. For information about enabling this option, see VMware KB 74880.

DNS Customization

Sets the Primary and the Secondary DNS servers for the migrated virtual machine.

Generate a new Security Identifier (SID)

Generates a new Security Identifier (SID) for the migrated virtual machine. This option is only for Windows virtual machines. Make sure the virtual machine has an active "local administrator" present. If the virtual machine is connected to the domain, it is moved to workgroup, and no DNS suffix is set.

Resize CPU

Changes the number of vCPUs for the migrated virtual machine. For more information about VMware virtual machine vCPU limitations, see the VMware vSphere product documentation.

Resize Memory

Changes the memory size for the migrated virtual machine. For more information about VMware virtual machine memory limitations, see the VMware vSphere product documentation.

Replicate Security Tag

Replicates VMware NSX for vSphere and NSX security tags associated with a virtual machine undergoing migration. The security tags are replicated from the source site to the destination site.

Migrate Custom Attributes

Migrates VMware Custom Attributes associated with a virtual machine undergoing migration. The Custom Attributes are replicated from the source site to the destination site.

The following table summarizes the options that apply to each migration type.

Option	Cold Migration	vMotion Migration	Bulk Migration	Replication Assisted vMotion Migration	OS Assisted Migration	Workload Migration for NSX V2T
Force Power-off VM	No	No	Yes	No	N/A	No
Remove Snapshots	No	Yes	Yes	Yes	N/A	No
Force Unmount ISO Images	No	Yes	Yes	Yes	N/A	No
Seed Checkpoint	No	No	Yes	Yes	No	No
Retain MAC	No	Yes (default setting; not selectable)	Yes (not supported for single vCenter operation)	Yes (default setting; not selectable)	Yes (default setting; not selectable)	Yes (see note)
Upgrade Virtual Hardware	Yes	Yes (upgrades on reboot)	Yes	Yes	N/A	Yes

Option	Cold Migration	vMotion Migration	Bulk Migration	Replication Assisted vMotion Migration	OS Assisted Migration	Workload Migration for NSX V2T
Upgrade VMware Tools	No	Yes	Yes	Yes	N/A	Yes
Deactivate Per-VM EVC	No	Yes	No	Yes	N/A	No
Hostname	No	No	Yes	No	No	No
Domain Name	No	No	Yes	No	No	No
Personalization Script	No	No	Yes	No	No	No
DNS Customization	No	No	Yes	No	No	No
Generate a new Security Identifier (SID)	No	No	Yes	No	No	No
Resize CPU	No	No	No	No	Yes	No
Resize Memory	No	No	No	No	Yes	No
Replicate Security Tags	Yes	Yes	Yes	Yes	N/A	Yes
Migrate Custom Attributes	Yes	Yes	Yes	Yes	N/A	Yes

Note Workload Migration for NSX V2T does not retain the MAC address if the same MAC address is already present at the target site.

Understanding Virtual Machine Snapshots During HCX Migrations

Virtual Machine snapshot handling varies slightly by migration type. The migration interface includes an option to explicitly remove virtual machine snapshots during the migration operation.

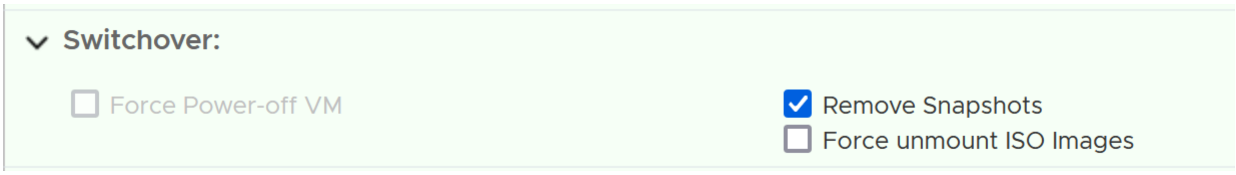

Switchover:

□ Force Power-off VM

☑ Remove Snapshots

□ Force unmount ISO Images

HCX vMotion, RAV, or Cold migration

As part of the migration validation, HCX detects and warns if snapshots are present for the selected VM.

If snapshots are not present, the migration proceeds.

If snapshots are present, check the **Remove Snapshots** option to proceed. HCX will delete snapshots and consolidate changes before proceeding with the migration's relocation.

Not selecting the **Remove Snapshots** option can have the following effects on migration:

- Interruption of HCX migration streams where any of the VMs being migrated have snapshots.
- Failure of pre-migration validation checks or of the migration attempt.

HCX Bulk migration

For Bulk migrations, HCX always removes virtual machine snapshots prior to migration to the remote site. As part of the migration validation, HCX detects whether snapshots are present for a VM. If present, HCX consolidates the LWD streams, and the replicated VMs will have no snapshots.

Note For Bulk migrations, the **Remove Snapshots** option has no effect. Bulk migration always removes snapshots, and migrations continue without interruption due to snapshots.

Viewing HCX Migration Event Details

When migrating content between peer sites, HCX performs a detailed set of actions that are visible in the system as migration events.

Viewing these events provides diagnostic information about the migration. This information provides the detailed migration workflow, the state of the migration, how long the migration remains in a certain state, and whether the migration has succeeded or failed. Understanding what is happening at any point in the migration can provide insight into what infrastructure or configuration changes might be necessary to address any migration issues that might occur.

HCX displays event information for all migration types.

Procedure

1 In the HCX dashboard, navigate to **Services** > **Migration** > **Tracking** or, if you have Mobility Groups activated, **Services** > **Migration** > **Management**.

The system displays a list of completed and ongoing migrations.

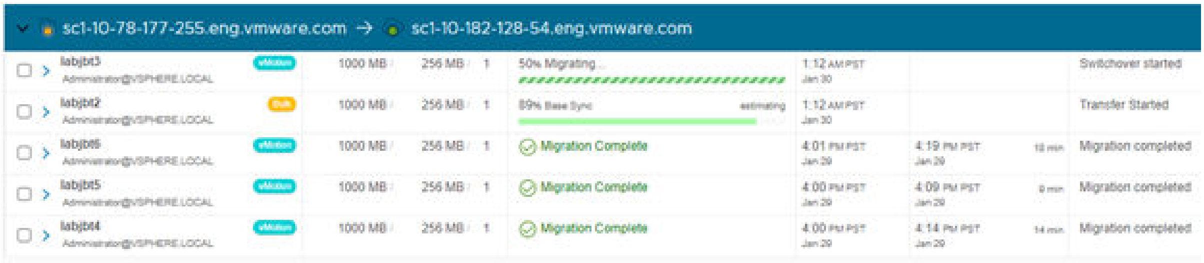

sc1-10-78-177-255.eng.vmware.com → sc1-10-182-128-54.eng.vmware.com

□	labjbt3 Administrator@VSPHERE.LOCAL	vMotion	1000 MB /	256 MB /	1	50% Migrating ...	1:12 AM PST Jan 30			Switchover started
□	labjbt2 Administrator@VSPHERE.LOCAL	Bulk	1000 MB /	256 MB /	1	89% Base Sync estimating	1:12 AM PST Jan 30			Transfer Started
□	labjbt6 Administrator@VSPHERE.LOCAL	vMotion	1000 MB /	256 MB /	1	✓ Migration Complete	4:01 PM PST Jan 29	4:19 PM PST Jan 29	18 min	Migration completed
□	labjbt5 Administrator@VSPHERE.LOCAL	vMotion	1000 MB /	256 MB /	1	✓ Migration Complete	4:00 PM PST Jan 29	4:09 PM PST Jan 29	9 min	Migration completed
□	labjbt4 Administrator@VSPHERE.LOCAL	vMotion	1000 MB /	256 MB /	1	✓ Migration Complete	4:00 PM PST Jan 29	4:14 PM PST Jan 29	14 min	Migration completed

2 For a specific migration in the list, expand the entry.

The system displays information specific to that migration.

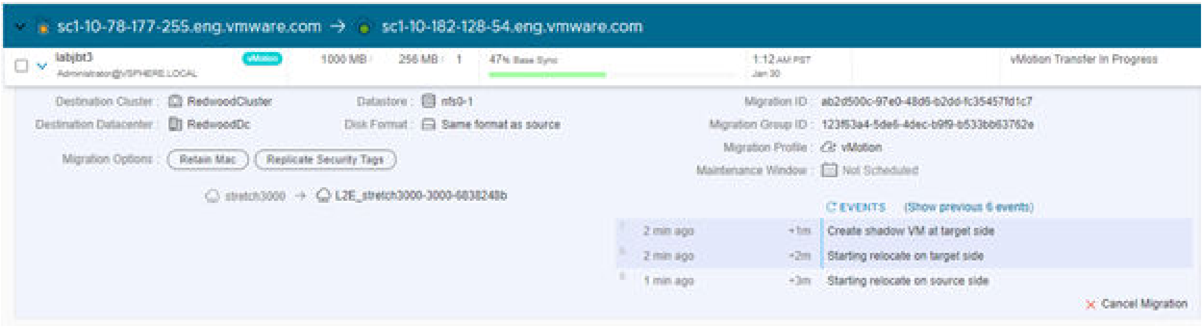

sc1-10-78-177-255.eng.vmware.com → sc1-10-182-128-54.eng.vmware.com

labjbt3 Administrator@VSPHERE.LOCAL vMotion 1000 MB 256 MB 1 47% Base Sync 1:12 AM PST Jan 30 vMotion Transfer In Progress

Destination Cluster : RedwoodCluster

Destination Datacenter : RedwoodDc

Migration Options : Retain Mac Replicate Security Tags

Datastore : nfs0-1

Disk Format : Same format as source

stretch3000 → L2E_stretch3000-3000-6838248b

Migration ID : ab2d500c-97e0-48d6-b2dd-fc35457fd1c7

Migration Group ID : 123f63a4-5de6-4dec-b9f9-b533bb63762e

Migration Profile : vMotion

Maintenance Window : Not Scheduled

EVENTS (Show previous 6 events)

2 min ago +1m Create shadow VM at target side

2 min ago +2m Starting relocate on target side

1 min ago +3m Starting relocate on source side

✕ Cancel Migration

3 Review the Events portion of the screen display for details about the migration.

By default, the latest three events are displayed. Use the refresh button to update the list of events.

Note While the migration is in progress, the systems provides options for managing the migration, including scheduling the migration switchover, canceling the migration, or forcing a power off.

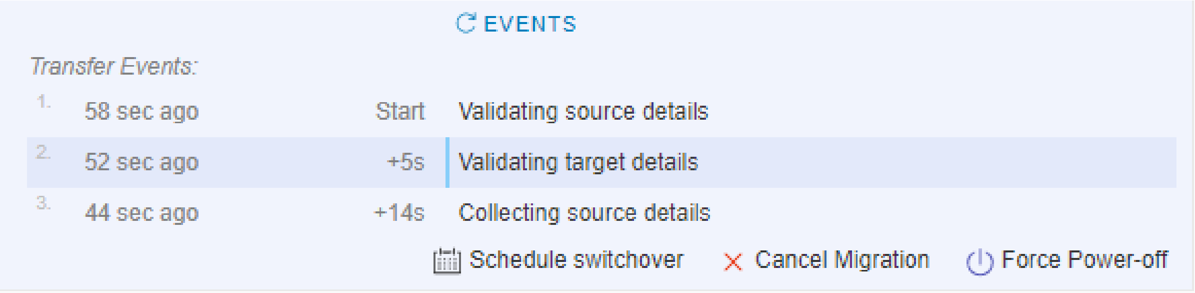

EVENTS

Transfer Events:

1. 58 sec ago Start Validating source details

2. 52 sec ago +5s Validating target details

3. 44 sec ago +14s Collecting source details

Schedule switchover ✕ Cancel Migration Force Power-off

4 To display additional details related to the migration, click **Show previous events**.

The migration events are color-coded for source and destination events. Destination events are shaded, while source events are not. Each event provides an offset time from the "Start" event.

Note Shading is relative to the peer. When viewing the migrations from the peer site, shading appears but is reversed.

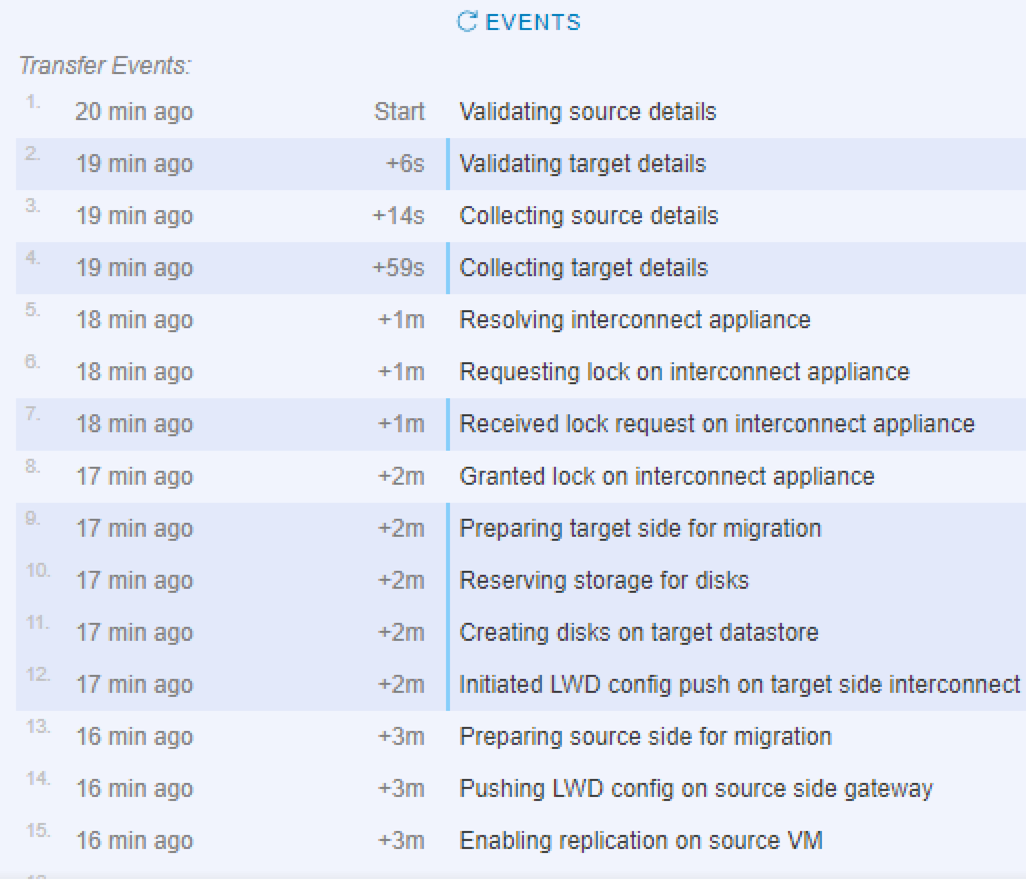

EVENTS

Transfer Events:

1.	20 min ago	Start	Validating source details
2.	19 min ago	+6s	Validating target details
3.	19 min ago	+14s	Collecting source details
4.	19 min ago	+59s	Collecting target details
5.	18 min ago	+1m	Resolving interconnect appliance
6.	18 min ago	+1m	Requesting lock on interconnect appliance
7.	18 min ago	+1m	Received lock request on interconnect appliance
8.	17 min ago	+2m	Granted lock on interconnect appliance
9.	17 min ago	+2m	Preparing target side for migration
10.	17 min ago	+2m	Reserving storage for disks
11.	17 min ago	+2m	Creating disks on target datastore
12.	17 min ago	+2m	Initiated LWD config push on target side interconnect
13.	16 min ago	+3m	Preparing source side for migration
14.	16 min ago	+3m	Pushing LWD config on source side gateway
15.	16 min ago	+3m	Enabling replication on source VM

Note Events that are more than 24 hours old include a date stamp.

When the migration has finished, event messages indicating that the source and destination sides of the migration have been cleaned up and the migration is complete.

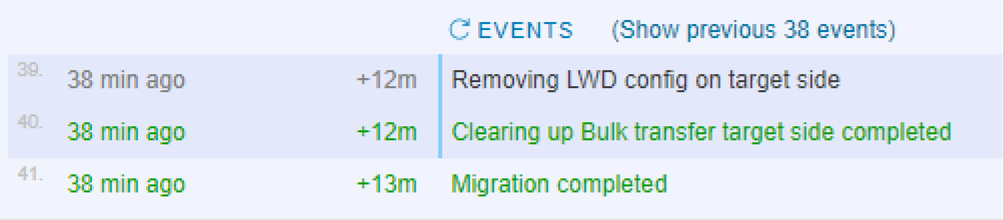

EVENTS (Show previous 38 events)

39.	38 min ago	+12m	Removing LWD config on target side
40.	38 min ago	+12m	Clearing up Bulk transfer target side completed
41.	38 min ago	+13m	Migration completed

Results

The system displays a list of migrated virtual machines and the migration status.

Canceling a Migration

The Migration interface includes an option for canceling in-progress migrations.

For OSAM, the effect of canceling a migration depends on the state of the migration when selecting the **Cancel** option:

- Canceling a migration while an HCX appliance is replicating data to the destination site deletes the associated resources created at the destination site with no effect on the source VM.
- Canceling a migration when the source system is in the final sync phase reboots the source system and deletes the associated resources.
- Canceling a migration after the target VM has been created deletes the destination virtual machine and the associated resources.
- Canceling a migration after the source virtual machine is powered down requires you to restart the virtual machine at the source site. HCX deletes the associated resources at the destination site.

For RAV, vMotion and Cold migration, you can cancel a migration at any point and the source VMs retain their original state.

For Bulk migration, the replication is canceled on the source site VM, and replicated data is deleted from the destination site. Canceling a migration after the target virtual machine is powered on prevents the source virtual machine from returning to a powered-on state. Ensuring only one instance of a virtual machine is powered on prevents duplicate active instances of a virtual machine on both the source and the destination sites.

If Seed Checkpoint is enabled for RAV or Bulk migration, HCX retains any transfered data at the destination site.

Procedure

1. In the HCX dashboard, select **Services > Migration**.

   The Migration Management interface displays a summary of migration information.

2. Identify the virtual machine on which to cancel migration, and expand the entry.
3. In the **Status** column, select **Cancel Migration**.

   This operation can take several minutes. When finished, the UI displays the message **Migration cancelled**.

Managing Failed or Canceled Migrations

For a failed or canceled migration, HCX offers you several options for what to do next.

Depending on the state of the migration, responding to a failed or canceled migration can include the following options:

- Retry a migration.
- Use the Force Cleanup option to clear failed or canceled migration processes.
- Archive the migration to clear the migration history.

Retrying a Migration

For a migration in a failed or canceled state, you can use the **Retry** option to retry the original migration operation. A migration retry operation uses the same parameters as in the original migration operation but assigns a new migration ID to the migration operation. The retry operation retains the same Group ID as the original migration.

If there are multiple migration failures or cancelations, you can choose to retry individual workloads or select the Group that contains the failed or canceled migrations. By selecting the migration group, only those entries in the failed or canceled state are retried.

By default, the **Retry** option enables the **Enable Seed Checkpoint** feature for Bulk and RAV migration types. If the **Enable Seed Checkpoint** option was selected during the initial migration attempt, the retry operation re-establishes the migration effort using the already migrated data. Enabling seed checkpointing prevents having to restart a migration from scratch due to a failure or cancelation.

Note If a cancelation or failed migration rollback does not complete, use the **Force Cleanup** option before the **Retry** option.

To retry a migration, refer to Retry a Failed or Canceled Migration.

Using Force Cleanup to Clear a Failed or Canceled Migration

Following a failed or canceled migration, the system attempts to clean up migration-related processes that started but did not complete. These processes occur on both the source and the destination sites. If the migration clean-up does not succeed, new migration or migration retry operations can fail. To view detailed information regarding the state of a migration failure or cancelation, expand the specific migration entry.

To clear any incomplete processes related to the migration, use the **Force Cleanup** option from either the **Migration Management** or the **Migration Tracking** screens. If **Enable Seed Checkpoint** was selected for the migration, **Force Cleanup** displays an option to remove the checkpoint data.

To clean up processes for a failed or canceled migration, refer to Force Cleanup for a Failed or Canceled Migration .

Archiving a Failed or Canceled Migration

To clear failed or canceled migration activity that requires no further action, use the **Archive** option. Archiving migrations removes them from the **Migration Management** and the **Migration Tracking** screens but does not remove the migration-related details from the HCX log files. Clearing the migration history using the **Archive** option updates the HCX Dashboard migration counters.

To archive migration activity, refer to Clearing the Migration History

Retry a Failed or Canceled Migration

For a migration in a failed or canceled state, you can use the **Retry** option to retry the original migration operation.

Prerequisites

Migrations must be in a failed or a canceled state prior to retrying a migration. In the case where the cancelation or failed migration rollback is not complete, as a best practice, perform **Force Cleanup** before using the **Retry** option.

Procedure

1 Navigate to either the **Migration Management** or the **Migration Tracking** screen.

2 Review the **Progress** and the **Status** columns for failed or canceled migrations, and select the migrations to retry.

   From the **Migration Management** or the **Migration Tracking** screens, you can select individual migrations to retry. If you have multiple failures or cancelations, and those migrations are part of a migration group, selecting the group retries all failed or canceled migrations in that group.

3 Click **Retry**.

   A pop-up window appears prompting you to confirm the retry operation.

4 Click **Retry**.

   The system initiates the migration. To view detailed event information, expand the migration entry.

Force Cleanup for a Failed or Canceled Migration

The **Force Cleanup** option provides the method for manually cleaning up failed or canceled migration processes that were not automatically cleaned up by the system.

Migration processes for a failed or canceled migration must be cleaned up prior to retrying a migration.

Procedure

1 Navigate to either the **Migration Management** or the **Migration Tracking** screen.

2 Review the **Progress** and the **Status** columns for failure or cancelation messages: `Canceling Migration`, `Migration Failed`, or `Migration Canceled`.

3 Review system and event messages for indications that processes related to the migration failure or cancelation were not cleaned up.

4 For any migrations that did not clean up properly, click **Force Cleanup**.

   If there are multiple migrations requiring clean up, you can select all that apply.

A pop-up window appears prompting you to confirm the clean-up operation and offering the option to remove seed checkpoint data.

Caution Removing the seed checkpoint data discards any previous migration progress, and that progress cannot be recovered.

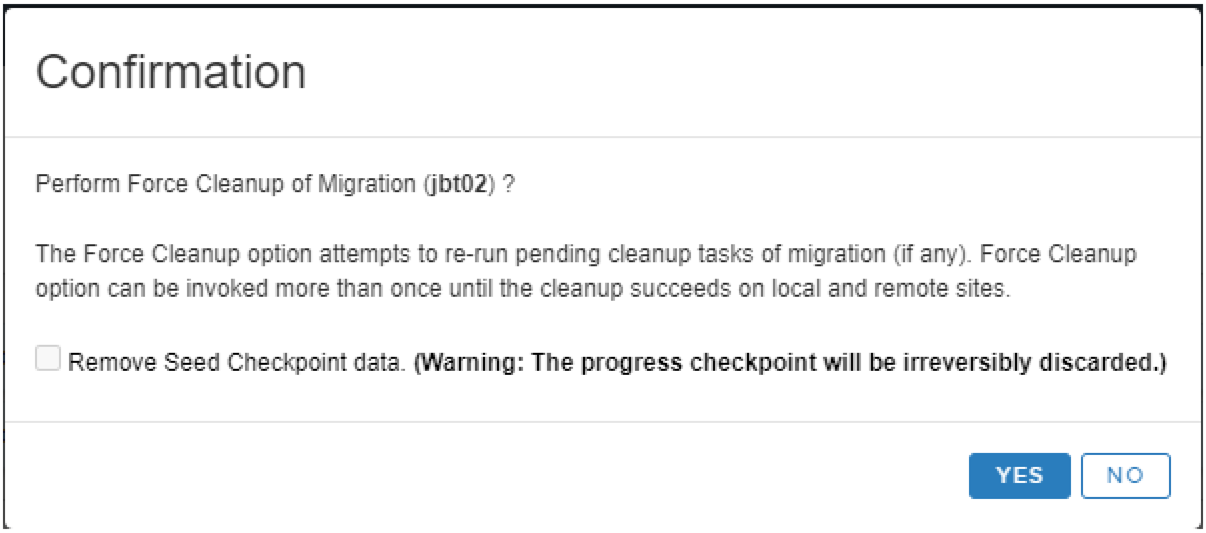

Confirmation

Perform Force Cleanup of Migration (**jbt02**) ?

The Force Cleanup option attempts to re-run pending cleanup tasks of migration (if any). Force Cleanup option can be invoked more than once until the cleanup succeeds on local and remote sites.

☐ Remove Seed Checkpoint data. **(Warning: The progress checkpoint will be irreversibly discarded.)**

5 (Optional) Select **Remove Seed Checkpoint data**.

You might select this option if the seed data at the destination site is not required for future use.

6 To proceed with the clean-up operation, click **Yes**.

Observe the specific migration event messages related to the migration to verify that the clean up was successful.

What to do next

If the clean-up operation did not succeed, repeat the **Force Cleanup** operation.

Clearing the Migration History

You can clear the migration activity for a site using the **Archive** option.

To clear failed, canceled, or completed migration activity that requires no further action, use the **Archive** option. Clearing the migration history updates the HCX Dashboard migration counters but does not remove the migration-related details from the HCX log files. Also, archiving migrations removes them from the **Management** and the **Tracking** screens in the UI.

Procedure

1 Navigate to the HCX Dashboard.

2 Select the **Services** > **Migration** > **Tracking**.

The Tracking window displays a summary of virtual machine migrations for a site pair.

3 Select the migration entries that you want to clear from the display.

   Note You cannot clear migrations that are in progress.

4 Click **Archive**.

   A pop-up screen appears prompting you to acknowledge the request to archive the migration entries.

5 Click **Archive**.

   The system clears the selected entries from the migration history.

HCX Integration with vRealize Network Insight

You can export waves of VMware vRealize Network Insight discovered applications to HCX for migration as Mobility Groups. HCX integration with vRealize Network Insight is available through API calls.

In many cases, the relationships, dependencies, and boundaries among application workloads is complex, and knowing what application to migrate and in which order can be challenging. vRealize Network Insight uses Application Discovery and Dependency Analytics to identify migration waves. From this information, vRealize Network Insight defines Application Groups that are then exported using public APIs to HCX as established Mobility Groups.

After HCX creates the Mobility Groups, you prepare for migration using the HCX Mobility Group configuration procedures.

Note All limitations and requirements for HCX migrations and migration types apply to Mobility Groups created from vRealize Network Insight.

Prerequisites

Public APIs are available for exporting vRealize Network Insight Application Groups to HCX as Mobility Groups. To view the HCX API for creating Mobility Groups, log in to access the HCX API documentation: https://*hcx_ip_or_fqdn*/hybridity/docs. Navigate to **Mobility > Migration Group APIs** in the documentation.

Procedure

1 Import the vRealize Network Insight discovered application groups into HCX using API calls.

2 Navigate to **Services > Migration** and verify that HCX created the Mobility Groups.

   Note Mobility groups created by vRNI have a vRNI label and a timestamp to differentiate them from other Mobility Groups created by HCX admins.

3 Configure the Mobility Groups for migration.

All Mobility Group operations are available for configuration, including setting the migration type, scheduling the migration, and editing the group information. See Migrating Virtual Machines with Mobility Groups.

4 Complete the migration.

Results

The workloads included in the API are created and migrated as Mobility Groups in HCX.

Protecting Virtual Machines with VMware HCX

9

VMware HCX provides services for protecting virtual machines based on the type of license installed.

The HCX Disaster Recovery service, standard with HCX, replicates and protects virtual machines to a remote data center.

Read the following topics next:

- VMware HCX Disaster Recovery

VMware HCX Disaster Recovery

VMware HCX Disaster Recovery (DR) is a service intended to protect virtual workloads managed by VMware vSphere that are either deployed in a private or a public cloud.

Limitations

VMware HCX Disaster Recovery has the following limitations:

- The number of concurrent VM protections is based on each HCX Manager. For HCX appliance limit information, see VMware Configurations Maximums.
- HCX DR does not support using datastore clusters for VM protection operations.
- HCX DR does not support protecting encrypted VMs.
- HCX DR does not support intra-cluster scaling.
- Virtual machines using virtual NVMe (vNVME) Controllers cannot be recovered during Test Recovery or Recovery VM operations.
- Guest customization is not available for HCX DR protection or recovery operations.
- Taking snapshots of a VM while protected, either manually or using a third-party backup solution, can disrupt the replication process. To prevent any impact, it is required to stop those services that create or remove snapshots during replication. Refer to KB79220 for more information.

Benefits

- Secure enterprise-to-cloud and cloud-to-cloud asynchronous replication and recovery of virtual machines.
- Recovery point objective (RPO) and recovery time objective (RTO) policy compliance.

  Note RPO is the interval of time that might pass during a disruption before the quantity of data lost during that period exceeds your Business Continuity/Disaster Recovery maximum allowable threshold. An RTO is the duration of time, and a service level within which data must be restored after a disaster to avoid unacceptable consequences associated with a break in continuity.

- Reverse failover of workflows to the source site.
- Self-service RPO settings from 5 minutes to 24 hours per virtual machine.

  Note RPO policy compliance depends on the available bandwidth from the source site to the destination site.

- Multiple point-in-time recovery snapshots that allows up to 24 previous replication points.
- Optimized replication throughput by using WAN Optimization.
- Route replication traffic through a customer-preferred direct connect network.
- On-premises monitoring and management with the fully integrated vSphere Client.
- Access to production-level support from VMware.
- While workloads are protected by HCX Disaster Recovery, the system automatically detects virtual machine disk additions or removals and reconfigures running protections to accommodate these changes. HCX continues monitoring disk changes until workloads are recovered, or protection is removed on them.

Planning for HCX Protection

The HCX Disaster Recovery service requires planning for the amount of storage consumed at the target location.

HCX Protection Workflow

Replication based operations such as HCX Bulk Migration, Replication Assisted vMotion and HCX Disaster Recovery use the vSphere Replication technologies to transfer virtual machine disk data. When a virtual machine protection operation is first run, the replication engine performs a full synchronization of all the data that makes up the virtual machine to the target location datastore. Following that baseline synchronization, the system performs a delta synchronization, meaning that only changed data blocks are replicated.

Delta synchronization occurs based on the recovery point objective (RPO) interval configured for the virtual machine, creating a replication instance. The selectable RPO ranges from 5 minutes to 24 hours. For example, setting the Recovery Point Objective (RPO) to 2 hours means that the maximum data loss that your organization can tolerate is 2 hours.

Setting an RPO does not mean that replication occurs on a specific interval. A replication instance reflects the state of a virtual machine at the time the synchronization starts. The system schedules replications so that the RPO is not violated. For example, assuming a 15 minute RPO, if the synchronization starts at 12:00 and it takes five minutes to transfer to the target site, the instance becomes available on the target site at 12:05. That instance reflects the state of the virtual machine at 12:00. The next synchronization can start no later than 12:10 so that instance is available no later than 12:15.

Note To determine the replication transfer time, the replication scheduler uses the duration of the last few instances to estimate the next one.

Following a full synchronization, the HCX DR service prompts you to run a test recovery operation to verify the replication.

Using Snapshots with HCX DR Protection

HCX allows for multiple recovery points, or replica instances, which are converted to snapshots when you recover a virtual machine. You set a retention policy for these instances by configuring a snapshot interval along with the number of snapshots to retain for each protected virtual machine. Snapshot intervals range from 1 hour to 7 days. The maximum number of snapshots taken during that interval can range from 1 through 24. For example, setting the number of snapshots to 4 and the snapshot interval to 1 day, means you can restore that virtual machine to any of 4 recovery points over the past 24 hours. In another example, setting the number of snapshots to 24 and the snapshot interval to 3-hours results in 8 snapshots per day for 4 days.

Note The RPO interval and the snapshot interval might not be the same. Snapshots are taken from the latest replication instance based on the RPO. The RPO must be set low enough to create the number of configured snapshots. For example, setting a retention policy of 6 snapshots per day means the RPO period must not exceed 4 hours to create at least 6 replication instances in 24 hours.

With snapshots, delta synchronizations are written to a new (replica) disk created for the snapshot in the same datastore as the baseline. Each new snapshot becomes the child of the previous version. For example, the first snapshot (replica 1) becomes the child and the baseline becomes the parent and all delta synchronization are written to replica 1. When a second snapshot (replica 2) is created, replica 2 becomes the child and replica 1 becomes the parent, and all delta synchronizations are written to replica 2.

Best Practices for HCX Protection Planning

Storage and bandwidth planning for replication at the target site depends on several factors:

Data set size

Consider the data set for replication and the capacity of the virtual disks (VMDK files) that make up the target site virtual machine. Consider whether the target site virtual disks are thick- or thin-provisioned. For example, a 100 GB virtual disk that is thick-provisioned always consumes 100 GB. A 100 GB disk that is thin provisioned consumes only the actual amount of data stored on the disk up to 100 GB. While a thin provisioned disk might initially use only a fraction of the provisioned storage, it can grow to the fill the total storage space.

Data change rate

Consider the amount of data replicated to the target location based on the rate of change in source virtual machine data. For example, a source virtual machine disk with 50 GB of data has an estimated daily change rate of 5 percent, meaning 2.5 GB of data is replicated each day.

Also, consider the maximum amount of data transferred for any one replication instance. Network bandwidth must be capable of meeting the RPO interval for the amount of data transferred.

Recovery Point Objective interval

Assuming consistent rate of change on the source virtual machine, a lower RPO means smaller delta synchronizations but higher bandwidth consumption to meet the lower RPO. Setting the RPO interval to the largest interval that your organization can tolerate can help to reduce network issues.

Network bandwidth

The replication network bandwidth must be sufficient to meet the RPO interval for the amount of data transferred. For example, if the RPO interval is 15 minutes, and the rate of change during that period is 1 GB, the network must capable of transferring that amount of data during the 15 minute interval. Set the number of recovery points as low as possible while still meeting business requirements.

Retention policy

Having multiple recover points means having a copy of the point in time changes for each snapshot, which increases storage requirements by the amount of change over the RPO interval times the amount of snapshots configured.

Protection concurrency with migration operations

Ongoing HCX migrations use the HCX Interconnect (HCX-IX) appliance for virtual machine disk replications. Resources used during a Bulk or a RAV transfer affect the total resources available for HCX Disaster Recovery, and conversely when the same service mesh appliances are used for both services.

Recovery and recovery testing

During recovery operations, or when testing a recovery plan with HCX Disaster Recovery, each recovered virtual machine consumes space. Normally, redo logs are consolidated into the replica base disk or into other redo logs if multiple recover points is activated. During a test recovery, some or all of the redo logs might be in use until the test recovery is cleaned up (completed). If redo logs are in use, HCX cannot consolidate the redo logs. Replication continues during a test recovery, which generates additional redo logs. The actual amount of storage capacity consumed depends on factors such as data change rates, replication frequencies, and how long the test recovery lasts.

DR Protection for a Virtual Machine

The VMware HCX virtual machine (VM) protection operation is used to configure the disaster recovery settings for a virtual machine, with specific remote site resources and recovery point objectives.

Note For the number of concurrent virtual machine protections supported with HCX, see Configuration and Service Limits.

Procedure

1 In the **vSphere Web Client**, navigate to VMware HCX.

2 Navigate to the **Disaster Recovery** tab and click **Protect VMs**.

  The HCX protection screen appears.

3 (Optional) Review the Replication Destination Site setting and check whether to Protect VMs to Primary Site or accept the default.

  By default, replication is to the destination site from the source site, and the system loads the virtual machine inventory for the source site.

  Checking **Protect VMs to Primary Site** causes replication to happen from the destination site to the source site, and the system loads the virtual machine inventory from the destination site.

4 Set the replication options.

  To set the options for all selected virtual machines, use the **Global replication options** area of the interface at the top of the window. To set machine-specific options, select a specific virtual machine and expand the entry.

Virtual Machine Replication Options	Description
Activate Compression	Helps during the seeding process of the VM. Helps if there is a low throughput for LAN/WAN connectivity.
Activate Quiescence	Pauses the virtual machine to ensure that the most consistent copy of the virtual machine is protected on destination site.

Virtual Machine Replication Options	Description
Seed Virtual Machine	This option allows you to select an existing off-line copy of the VM at the destination site. This can reduce the amount of VM data transmitted for DR protection.
Destination Container (optional)	Specifies the Cluster or the Resource Pool where the protected copy of the virtual machine lives.
Storage	Specifies the datastore on which the protected copy of the virtual machine resides. The Storage Policy drop-down menu lists all compatible datastores. If there is no selection, the system uses the Default Storage Policy and the corresponding datastore.
RPO	Indicates the Recovery Point Objective for the VM. With VMware HCX, settings range from 5 minutes to 24 hours. **Note** The 5 minute RPO requires the source host to be ESXi 6.0 or later for vSAN, and ESXi 6.5 for other supported datastores.
Snapshots Interval	Defines the interval between Snapshots. In the event a corrupted change was synchronized to the protected (destination) site, this options provides a way to recover from an earlier point in time. Snapshots provide a Multiple Point in the Time Recovery plan for the protected VM.
No. of Snapshots	Indicates the total number of snapshots within the established snapshot interval.
Destination Network	Corresponding network that the protected virtual machine uses. If the network for the source virtual machine is extended, this field is automatically populated.

Note Always verify the Storage Policy and associated datastore selection, and evaluate the expected storage usage at the DR site. The settings cannot be changed once the protection is in place. Storage Policy selection is NOT available during Recovery or Test Recovery operations.

5 Click **Finish**.

The DR Dashboard is displayed. You can monitor the progress of virtual machine protection. When protection operation is complete, a green shield appears in the Status column.

6 Review the dashboard information.

- Local VMs – Reflects the number of virtual machines on the source site that are protected.
- Remote VMs – Reflects the total number of virtual machines on the remote site that are being protected.
- Activity – Monitors DR related operations.
- Green Shield – Indicates DR protection is active.

  - Yellow triangle – Indicates DR protection has not been tested.
  - Local Site/Remote Site – Arrows indicate the direction of protection.

What to do next

Verify protection using the procedure Performing a Virtual Machine Test Recovery.

Performing a Virtual Machine Test Recovery

A VMware HCX Disaster Recovery protection configuration can be tested by bringing the virtual machine online with a test recovery operation, which does not disrupt the ongoing replication.

Prerequisites

- An initial full synchronization of the protected virtual machine is required. The interface dims the virtual machine test recovery option while the initial synchronization is in progress, to indicate that the option is not available until the initial synchronization procedure completes.
- When working with protected virtual machines on extended networks:
  - Do not connect a test-recovered virtual machine to the extended network. Doing so may impact the original protected virtual machine due to the duplicate IP address.
  - To test the recovery, create or use a test network at the Disaster Recovery site.

Procedure

1. Log in to the **vSphere Web Client** and access the VMware HCX plugin.
2. Go to the **Disaster Recovery** tab.
3. Select the virtual machine and under **Actions**, click **Test Recover VMs**.

   A test recovery screen appears prompting you for the recovery parameters.
4. Make the parameter selections and click **Test**.

   Note If the protected virtual machine is on a stretched network, the option to use **None** is available for the target network.

   After the test completes, a certificate icon appears indicating a successful test. A solid yellow triangle appears indicating that a test cleanup is needed.
5. Perform the clean-up operation,

   a. Select the virtual machine.

   b. Click **Actions**, and then click **Test Recover Cleanup**.

   c. Click **Cleanup**.

Results

The test recovery operation is complete.

Performing a Virtual Machine Recovery

Using the VMware HCX Disaster Recovery operation, you can activate the Virtual Machine replica at the HCX destination site.

Prerequisites

This procedure applies when a protected virtual machine has become unavailable due to a disaster event. This unavailable state is indicated with a red lightning bolt status in the HCX Disaster Recovery interface.

Procedure

1. Open the VMware HCX Cloud interface at the destination site.
2. Navigate to **Services** > **Disaster Recovery**.
3. Click **Actions** and click **Recover VMs**.

   A test recovery screen appears prompting you for the recovery parameters.

   The recovery process starts. After the recovery completes, the Virtual Machine is visible in the VMware HCX Disaster Recovery destination site's vSphere Inventory.
4. Make the parameter selections, and click **Recover**.

Results

After the recovery completes, the virtual machine is available in the VMware HCX Disaster Recovery destination site's vSphere Inventory.

VMware HCX Disaster Recovery - Protect Operations for VMs

In addition to Recovery and Test Recovery operations, HCX provides various operations that provide more control and granularity in replication policies.

The additional operations are available from the Action menu in the Disaster Recovery interface.

Available Operations include:

1. Reverse – After a disaster has occurred. Reverse helps make Site B the source site where the protected VM now lives.
2. Pause – Pause the current replication policy associated with the virtual machine selected.
3. Resume - Pause the current replication policy associated with the virtual machine selected.
4. Remove - Remove the current replication policy associated with the virtual machine selected.
5. Sync Now – Performs out of band sync on the source virtual machine to the protected virtual machine.

Managing System Settings

10

Use the HCX Manager Appliance Management interface for viewing, configuring, and managing system-level functions.

The Appliance Management interface is reached by navigating to the management port: https://*hcx-ip-or-fqdn*:9443. This interface uses the system administration credentials set up during the OVA deployment.

The Appliance Management interface provides the access to the system Dashboard, Appliance Summary, Configuration, and Administration information.

Note Appliance management operations might be done by your cloud service provider .

Read the following topics next:

- Understanding the Appliance Management Dashboard
- Updating the Time Settings
- Updating the System Name
- Configure a Proxy Server

Understanding the Appliance Management Dashboard

The system Dashboard provides the access to status and services, configuration settings, and system-level administration tasks.

The Dashboard is the first screen that appears after you log in to the system Appliance Management interface port (:9443).

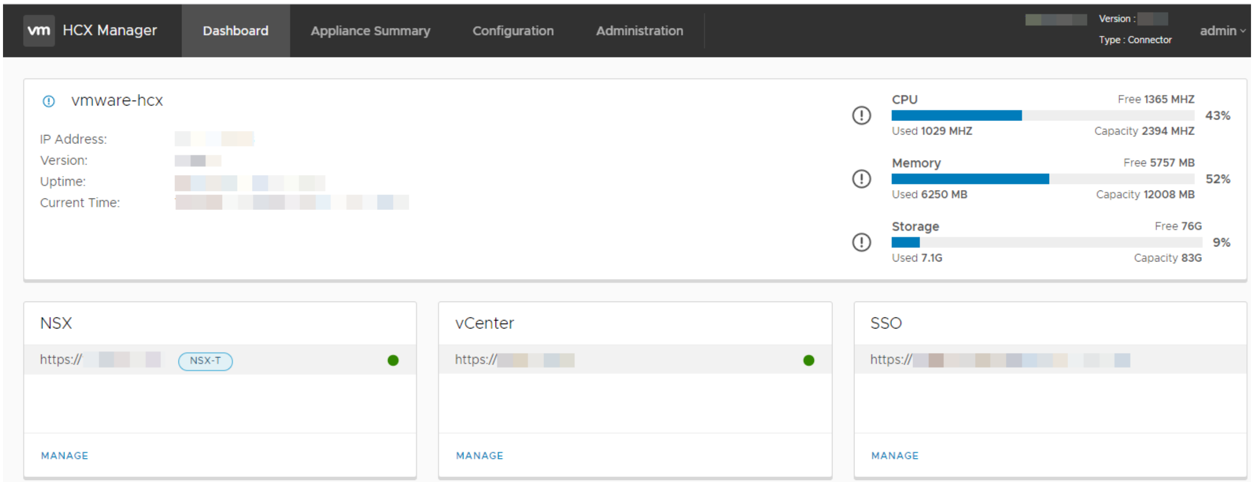

[redacted] Version : [redacted]
Type : Connector

vmware-hcx

IP Address: [redacted]
Version: [redacted]
Uptime: [redacted]
Current Time: [redacted]

CPU Free 1365 MHZ 43%
Used 1029 MHZ Capacity 2394 MHZ

Memory Free 5757 MB 52%
Used 6250 MB Capacity 12008 MB

Storage Free 76G 9%
Used 7.1G Capacity 83G

NSX
https:// [redacted] NSX-T ●

vCenter
https://[redacted] ●

SSO
https://[redacted]

The Dashboard provides the access to various system management settings through a set of tabs at the top of the display.

Note For installations where the vCenter Servers are in linked-mode, the Dashboard includes information from all vCenter Servers registered to a system.

Tab Entry	Description
Dashboard	Displays the appliance status as a set of summary panels: ■ System information and resource usage ■ NSX status ■ vCenter status ■ SSO status ■ Public Access URL status The panels that are visible in the display depend on the installation type. To change the configuration settings for a panel, click **Manage**. The system redirects you to the Configuration tab, where you can update the settings.
Appliance Summary	Displays the status of services running on the system: ■ Hybridity Services ■ Common Services ■ System Level Services Options are provided to stop and restart services. The list of services in the display varies based on the installation type.

Tab Entry	Description
Configuration	Displays the list of service configuration settings. ■ Licensing ■ vCenter ■ SSO ■ Public Access URL ■ vSphere Role Mapping ■ Data Center location To display the current settings, click an item in the list. To modify the current settings, click **Edit**.
Administration	Displays the list of system-level configuration settings. ■ General Settings ■ Time Settings ■ Syslog Server ■ System Name ■ Network Settings ■ General Network ■ DNS Servers ■ Proxy ■ Static Routes ■ Troubleshooting ■ Technical Support ■ Logs ■ Upgrade ■ Backup & Restore ■ Certificate ■ Trusted CA Certificate ■ Server Certificate To display or edit the settings, click an item.

Updating the Time Settings

The system provides initial NTP Server settings during the OVA deployment in the vCenter Server. These settings can be updated in the HCX Manager Appliance Management interface.

Caution Editing NTP Settings requires restarting the Appliance Management Service. You can restart this service from within the **Appliance Summary** tab.

Editing and Removing the NTP Server Configuration

NTP Settings can be modified in the appliance management interface.

HCX requires a valid NTP server synchronized time for integrated systems operations.

1 Navigate to the appliance management interface: `https://hcx-ip-or-fqdn:9443`.

2 Navigate to the **Administration** tab.

3 Select **Time Settings** on the side menu, click **Edit** (or **Unconfigure NTP Servers**).

4 Enter the NTP server.

   Multiple servers can be specified using a separated comma-separated list.

5 Navigate to the Appliance Summary tab in the dashboard, locate the Appliance Management Service, and click **Restart**.

Updating the System Name

The initial Hostname is provided during the OVA deployment. The system name can be updated in the HCX Manager Appliance Management interface.

Important When renaming an HCX Connector or an HCX Cloud , do not use same name for both appliances in a site pair.

Editing the System Name

1 Navigate to the Appliance Management interface: https://*hcx-ip-or-fqdn*:9443.

2 Navigate to the **Administration** tab.

3 Select **System Name** on the side menu, then click **Edit**.

4 Enter the System Name. Click **Save**.

Configure a Proxy Server

A proxy server is an intermediary “message-forwarding agent” selected by a client through its local configuration for outbound HTTPS requests for security or shared caching. Add a proxy server configuration in HCX Manager to send HTTPS requests to a proxy server in the environment.

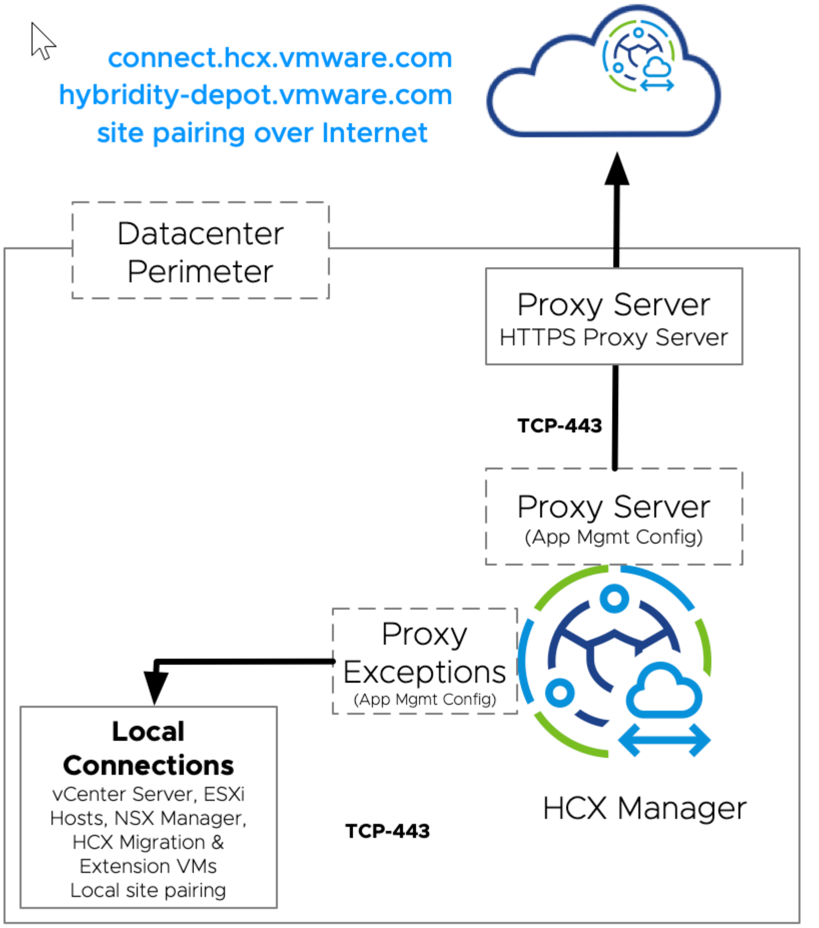

When configuring a proxy server, refer to the following considerations and best practices:

- HCX Manager systems make various HTTPS requests during normal operation:

  Outbound Connections:

  - HCX Manager to `connect.hcx.vmware.com` (for activation /entitlement)
  - HCX Manager to `hybridity-depot.vmware.com` (for updates/downloads)
  - HCX Manager to `console.cloud.vmware.com`
  - HCX Manager to Remote HCX Manager (for site pairing)

  Local Connections:

  - HCX Manager to Registered vCenter Server
  - HCX Manager to Registered vCenter Server's ESXi Hosts
  - HCX Manager to Registered NSX Manager system
  - HCX Manager to Migration and Network Extension Service Mesh appliances deployed by this HCX Manager

- A proxy server is usually intended to handle internet-bound connections from internal systems (to endpoints that resolve to public IP addresses).
- Use the **Proxy Server** field to enable proxy operation.
- For HCX to function correctly when a proxy server is configured, local connections must be explicitly excluded from proxy operation. Use the **Proxy Exclusions** field.
- The destination HCX Manager for site pairing must be configured with the Local Connection when the IP address is internally reachable without traversing the proxy.
- Use the **Proxy Exclusions** field for broad or granular configurations.

  A simple way to restrict Local Connections is to enter one large subnet that includes all internal IP address space for the data center in the **Proxy Exclusions** field. Alternatively, restrict Local Connections by specifying granular subnets in the **Proxy Exclusions** field.
- Kerberos and Windows NTLM Proxy Servers are not supported.

Important Configuring a proxy server without the local exclusions typically results in migration failures and errors during HCX operation. See VMware KB 89180.

Note HCX Service Mesh does not support proxy server configuration.

Procedure

1. Log in to the management interface: https://*hcx-ip-or-fqdn*:9443.
2. Navigate to the **Administration** tab, and select **Proxy**.
3. Enter or edit the proxy server settings:
   a. Proxy Server IP address or FQDN.
   b. Proxy Server Port.
   c. Proxy Server User.
   d. Proxy Server Password.
   e. Proxy Exclusions.

      Using a comma separated list to define all related proxy server exclusions, enter any IP, subnet, host, and/or domain names. Use * for wildcard values and do not include complete URLs (no https://).
4. To verify the configuration, click **Test Connection** and enter the test URL.
5. Click **Save**.
6. Restart the HCX Manager services.

   Restarting HCX Manager services is required for the proxy exclusions to take effect. For more information, see Monitoring HCX Services from the Appliance Management Interface.

Managing CA and Self-Signed Certificates

11

The appliance management interface can be used to add or remove certificates from the system certificate store.

HCX uses self-signed certificates for the HCX Manager Appliance Management interface (port 9443) and HCX Manager UI (port 443) appliances.

When upgrading to HCX 4.4.0 or later, HCX detects and rotates self-signed certificates nearing expiry.

If the HCX Manager certificate is set to expire in less than one year, a new certificate is generated that is shared between the Appliance Management interface and HCX Service UI appliances. Additionally, the common name (CN) for the self-signed certificate is changed from the current FQDN name to hcx.local.

HCX deployments never rotate the customer imported certificates.

If your environment uses a certificate monitoring system, that system must accept the new self-signed certificate.

Read the following topics next:

Importing Trusted Certificates from a Remote Site

This procedure allows you to import manually and trust certificates from remote systems on the HCX Manager appliance.

Procedure

1. Navigate to the appliance management interface: https://*hcx-ip-or-fqdn*:9443.
2. Navigate to the **Administration** tab.
3. Select **Certificate > Trusted CA Certificate** on the side menu.
4. Select the certificate import option: **File**, **URL**, or **Content**.

5 Enter the information for the selected option.

For example, when selecting the URL option, enter the IP or FQDN that the source HCX Manager uses to reach the HCX Cloud Manager.

6 Click **Apply**.

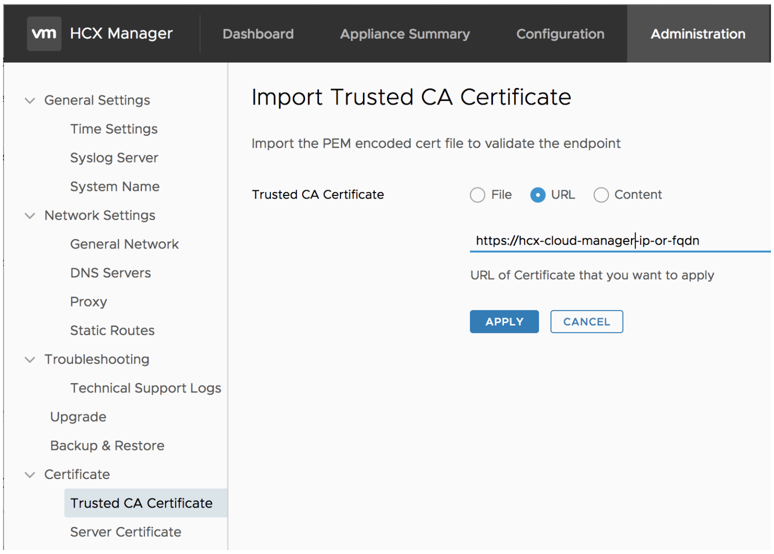

Import Trusted CA Certificate

Import the PEM encoded cert file to validate the endpoint

Trusted CA Certificate ○ File ● URL ○ Content

https://hcx-cloud-manager-ip-or-fqdn

URL of Certificate that you want to apply

Updating the Local Server Certificate on an HCX Manager

This procedure allows you to update manually the certificate on the HCX Manager appliance.

Procedure

1 Navigate to the appliance management interface: https://*hcx-ip-or-fqdn*:9443.

2 Navigate to the **Administration** tab.

3 Select **Certificate > Server Certificate** on the side menu.

4 Enter the **Certificate** and the **Key** information, and click **Apply**.

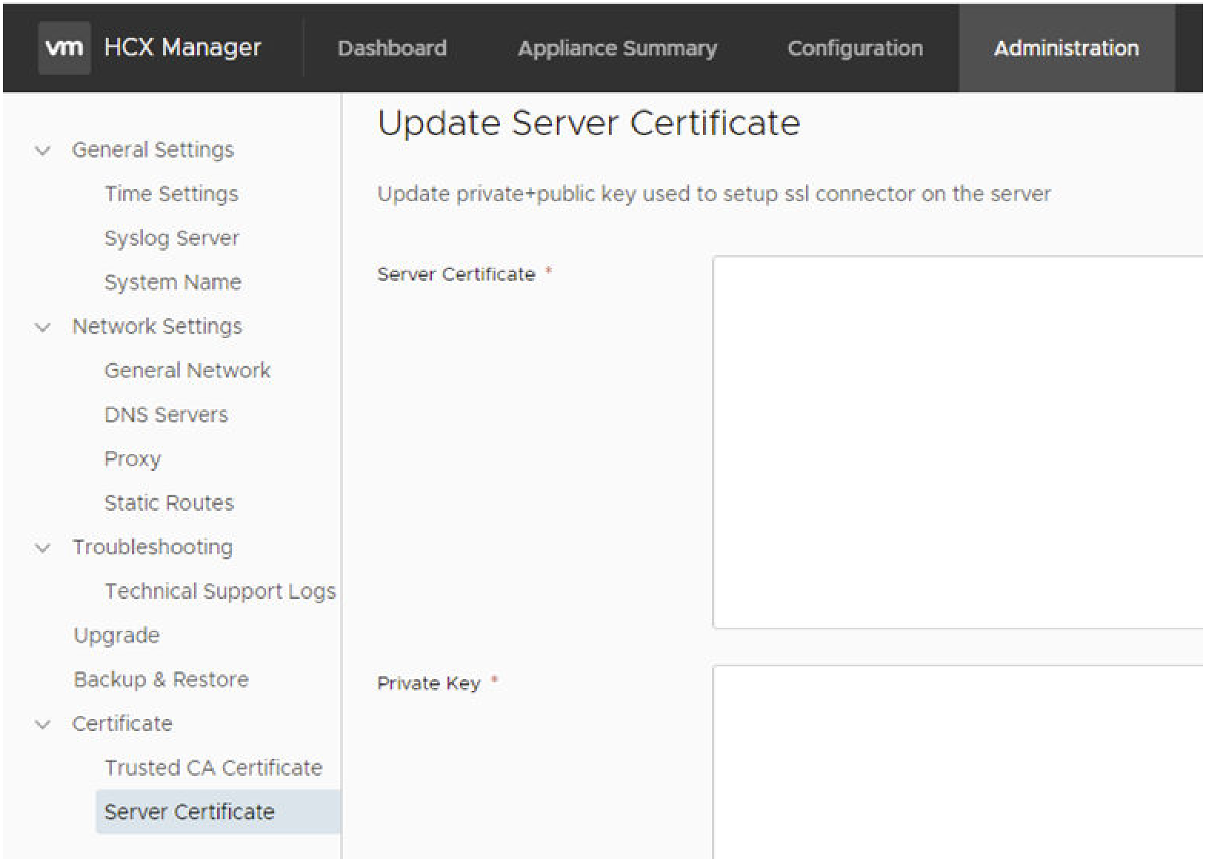

Update Server Certificate

Update private+public key used to setup ssl connector on the server

Server Certificate *

Private Key *

Removing Certificates

You can view the certificate information and remove certificates through the HCX Manager interface.

Caution Deleting non-user imported certificates can impact signing chains or otherwise impact HCX appliances.

Procedure

1 Log in to the HCX Manager interface using the admin user credentials: `https://`*`hcx-ip-or-fqdn`*`:9443`.

2 Navigate to **Administration** > **Certificate** > **Trusted CA Certificate**.

A list of certificates appears. To view certificate information, expand an entry.

3 To delete an entry, click the vertical dots menu and then click **Delete**.

Results

The HCX Manager removes the certificate from the database.

Backing Up and Restoring the System

12

You can back up and restore the appliance from the appliance management interface.

Backup and restore operations are available in the appliance management interface except when restricted by a cloud service provider. You first use the appliance management interface to generate a configuration file and then use that file to restore to a healthy system.

The HCX service appliances, which include HCX-IX and HCX-NE, do not require individual backups. A restored HCX Manager reconnects to existing service appliances that were created within the backup time frame. If the service appliances are no longer functional, the HCX Manager deploys new appliance virtual machines based on the backed-up configuration.

Read the following topics next:

- Backing Up HCX Manager
- Restoring the System

Backing Up HCX Manager

You use the appliance management interface to create a backup file.

This operation backs up the following information:

- Inventory data
- Configuration files
- Certificates
- System UUID

The backup file is saved in tar.gz format.

Procedure

1 Log in to the appliance management interface: <https://hcx-ip-or-fqdn:9443>.

2 Navigate to **Administration** > **Troubleshooting** > **Backup & Restore**.

3 Set up an SFTP server for uploading the backup file:

   a Click the **SFTP server setting** tab.

   b Click **Add**.

   > **Note** The best practice to use a Linux-based OpenSSH host for file transfer operations.

   c Enter the SFTP server information and click **Save**.

4 (Optional) Configure a backup schedule:

   > **Note** The best practice is to schedule **Daily** backups. Restoring from backup files that are more than two days old is not supported due to potential inventory changes from the backup time to present.

   a Click the **Scheduling** tab.

   b Click **Add**.

   The scheduling window appears.

   c Select the Backup Frequency.

   d Enter the hour and minute of the backup.

   e Click **Save**.

5 Click the **Backup and Restore** tab.

6 Click **Generate**.

   If a backup schedule is configured, the system creates the backup file at the scheduled time.

7 For manual backups, save the backup file:

   > **Note** If you have scheduled backups, the system automatically generates the backup file at the scheduled time and saves the file to the SFTP server.

   - To save the generated file to an SFTP server, check the box **Upload to server**.
   - To download the generated file to the client browsing system, click **Download**.

Restoring the System

You use the appliance management interface to restore the system from a backup file. The restore operation is used in cases where the system has become corrupt or unusable due to resource or system failures.

This operation restores the appliance to the state it was in at the time of the backup. The contents of the backup file supersede configuration changes made before restoring the appliance.

Note A restored HCX Manager cannot connect to HCX service appliances that were created during a time after the backup file was generated.

Prerequisites

You have deployed a replacement system that is clean of prior configuration settings. The replacement system has the same software version and IP address as the original system.

Note A clean system deployment requires only the minimum configuration to be manageable and that the system is network reachable from the operator or client system.

Procedure

1. Log in to the appliance management interface: <https://hcx-ip-or-fqdn:9443>.
2. Navigate to **Administration > Troubleshooting > Backup & Restore**.
3. Within the **Restore** section, browse to the backup file and open it.

   Note Restoring from backup files that are more than two days old is not supported due to potential inventory changes from the backup time to present.

4. Click **Continue**.

   The system verifies the uploaded file.

5. Click **Restore**.

   The restoration begins. This process can take several minutes to complete.

6. Verify that the system is operating properly:

   a. Navigate to the Dashboard tab and confirm that the NSX and vCenter Server status is green.

   b. Navigate to the Appliance Summary tab and verify that the Hybridity Services, Common Services, and System Level Services are running.

Manage Alerts

13

HCX generates alerts with different severity levels to flag events such as migration, network stretch failure, or reachability issues connecting to an endpoint or paired site.

Based on the type of alert, you can perform various actions to acknowledge, reset, or suppress the message.

Alert Type	Action	Action Description
Critical	Acknowledge	The alert has been noted and the corrective actions will be taken. HCX records the user who acknowledged the alert, including a time stamp for the entry.
	Reset to Green	Action has been taken to correct the alert. HCX records the user who resets the alert, including a time stamp for the entry. The alert is removed from the display.
Warning	Acknowledge	The alert has been noted and the corrective actions will be taken. HCX records the user who acknowledged the alert, including a time stamp for the entry.
	Reset to Green	Action has been taken to correct the alert. HCX records the user who resets the alert, including a time stamp for the entry. The alert is removed from the display.
Info	Suppress	Signifies that the alert has been reviewed and can be ignored. HCX removes the alert from the display.

Note To take action on an alert, you must log in to the HCX Manager UI.

Procedure

1 Log in to the HCX Manager UI: https://hcx-ip-or-fqdn:443.

The system displays the site manager Dashboard.

2 Navigate to the Alerts interface in the HCX Manager Dashboard, and review the list of alerts for action.

   a (Optional) Expand an entry to see more information regarding an alert.

   b (Optional) To sort the list by status, filter the list by Serverity and select **Info**, **Warning**, or **Critical**.

   c (Optional) To sort the list by **Entity Name** or **Creation Date**, select the respective filter and enter the search information.

3 To take action on an alert, click the ellipsis next to the alert entry, and select the action.

Results

The actions are updated in the Alerts interface in the HCX Dashboard.

Updating VMware HCX

14

The information includes step-by-step instructions for updating HCX components.

Important For details on the support and upgrade requirements for HCX releases, review the HCX Release Notes specific to the update.

Read the following topics next:

- About HCX Service Updates
- Planning for HCX Updates
- HCX Service Update Procedures

About HCX Service Updates

HCX service updates may include new features, software fixes and security patches.

HCX service updates are published periodically as a set for HCX Connector and HCX Cloud types.

Overview of HCX Component Updates

- HCX service updates can be summarized in the following steps:
  - During a new HCX implementation, the latest updates are applied automatically.
  - When VMware releases a service update, metadata for the release is published to the HCX client systems. The HCX Manager displays a notification banner noting the update.
  - The HCX admin identifies site paired HCX client systems, and applies the new service updates to the paired HCX Manager systems. You can update HCX Connector and HCX Cloud systems during separate maintenance windows, but for optimal compatibility update both systems together.
  - Apply service updates during a maintenance window where no new HCX operations are queued up.
    - The HCX Manager and Service Mesh can be upgraded independently, during separate maintenance windows.
    - The upgrade window accounts for a brief disruption to the Network Extension service, while the appliances are redeployed with the updated code.

    - During the window, the Interconnect service components are updated to the new release.

  - Component updates are triggered for each Interconnect or Service Mesh using the source side HCX plugin, but are run symmetrically at the source and destination site.

Planning for HCX Updates

As part of HCX update planning, and to ensure that HCX components are updated successfully, review the service update considerations and requirements.

Service Update Requirements

- HCX Manager systems periodically connect to **connect.hcx.vmware.com** and query the server for published service updates. A continuous connection is required. The VMware HCX UI displays a banner when an updated HCX release is available.
  - VMware HCX client systems must be able to reach **connect.hcx.vmware.com** using HTTPS throughout the entire lifecycle of the system. When this connection is not available, the VMware HCX client system cannot display updates available to other VMware HCX systems.
  - If the connection is not maintained, the client system can miss a published update.
  - A client system without a maintained connection to **connect.hcx.vmware.com** is placed out of support if the connection is not restored. Also, the system displays a banner stating that the system will be deactivated.
- If the HCX service update is not reflected on all site paired HCX systems, contact VMware Support. Partial updates are not supported.
- VMware HCX client systems must be able to reach **hybridity-depot.vmware.com** using HTTPS for the download of update files, without connectivity to the depot, the Update Download fails.
- HCX Site Pairing must reflect healthy connections before applying the service update.
- Unless directed by VMware Support to upgrade to resolve a known issue, HCX components reporting degraded state must be restored to a healthy state before the update.

Service Update Considerations

- The HCX service update file can be downloaded to the HCX Manager systems before the upgrade to reduce the time of the maintenance windows.
  - If Site A is paired with Site B, and Site A is also paired with Site C, plan the updates for Site A, B and C for the maximum compatibility across all environments. The environments can be updated in separate windows.

- Applying a service update causes the HCX Manager system to be rebooted:
  - Existing Network Extensions continue to work during the HCX Manager reboot. New Network Extensions cannot be configured while the HCX Manager is rebooting.
  - Existing VM Protections continue to work during the HCX Manager reboot. New replications cannot be configured while the manager is rebooting.
  - Because upgrading the HCX Managers does not disrupt the Interconnect Service Mesh, the HCX team encourages installing updated releases when they become available to ensure that systems have the most recent fixes and security patches.
- HCX Interconnect (Migration, WAN Optimization, and Network Extension) service component upgrades are performed independently to the manager upgrades:
  - Upgrade the Service Mesh appliances only after all Site Paired HCX Managers are upgraded.
  - Updating the Interconnect service components disrupts those services while the updates are being applied.
    - Ensure that migrations are not running or new migrations or replications are scheduled when updating the IX/CGW or WAN-OPT appliances.
    - Updating the HCX-NE (L2C) appliances disrupts connectivity that crosses the Network Extension path. The tunnel state re-converges in less than one minute after triggering the update.
      - Update the Network Extension components during a maintenance window.
- HCX client systems to be running within the latest three releases to be eligible for support.

Service Update Sequence

1. When a published update is available:
   - Identify the environments connected through HCX Site Pairing. The paired systems are displayed in the two tables in the Administration tab.
     - Connect to all paired HCX Managers and ensure that the update is available.
     - Download the update on all the paired HCX Managers.
     - Ensure that no new migrations, protections, or network extensions are configured during the update.
     - Ensure that all ongoing migrations have finished.
     - Ongoing synchronizations for Disaster Recovery are supported.
     - Ensure that there are no failovers scheduled during the upgrade.
2. Initiate the Upgrade task on all paired HCX Connector and HCX Cloud systems:
   - The HCX Manager system reboots during the upgrade procedure.
   - Allow the system several minutes to complete the initialization process.

- Use the System Updates view to verify that the current version is updated.

3 The HCX Service Mesh can be upgraded once all paired HCX Manager systems are updated and all services have returned to a fully converged state.

  - HCX Interconnect service components can be upgraded from the source HCX system. Use the Service Mesh interface to redeploy or upgrade the VMware HCX Interconnect service appliances:
    - Upgrade or redeploy the HCX-IX (CGW) and HCX-WAN-OPT together.
      - Verify that the required tunnels are functional before resuming services or proceeding to the next component.
    - Upgrade or redeploy the HCX-NE (L2C) appliance.
      - Verify that the required tunnels are functional before resuming services or proceeding to the next component.
    - If the HCX topology has multiple source sites paired to a destination environment, the components upgrade has to be triggered at each source site.

HCX Service Update Procedures

Updating a VMware HCX system installs the latest features, problem fixes, and security patches.

Upgrading the HCX Manager

The HCX update is applied to the HCX Manager systems first.

It is a best practice to create a backup prior to upgrading HCX Manager. See Chapter 12 Backing Up and Restoring the System. This back up option may not be available in some Public Clouds where HCX is managed by the cloud service provider.

In addition to backing up HCX Manager, optionally use the vSphere snapshot feature to take a snapshot of the HCX Manager at the source and the destination sites. If necessary, you can use snapshots to roll back the HCX Manager version. See Rolling Back an Upgrade Using Snapshots.

You can update the HCX manager from the either the HCX Manager UI (https://hcx-ip-or-fqdn:443) or the HCX Manager appliance management UI (https://hcx-ip-or-fqdn:9443). The HCX Manager UI displays system update notifications online through the System Updates interface. Manual updates are done offline, through the management interface. Review the following procedure for either an online or offline update.

Prerequisites

- Verify the HCX Manager system reports healthy connections to:
  - vCenter Server.
  - NSX Manager (if applicable).
  - Cloud Director/RMQ (if applicable).

- Verify that the HCX Manager reports that there are healthy connections to the HCX Interconnect service components.
- Verify that Site Pair configurations are healthy.

Procedure

1 To perform an online update through the HCX Manager UI, follow these steps:

a Open the HCX Manager Service UI.

Note You can update site-paired HCX Managers simultaneously.

b Navigate to the **Administration** tab.

c Navigate to the **System Updates** section.

d In the Local HCX section, under **Available Service Update Versions**, click **Check for Updates**.

In normal operation, the HCX automatically receives the latest service update. But if the HCX is offline or unable to access the Internet when a service update is pushed out, the HCX can miss the update. This selection checks for the latest version and adds it to your available service updates.

e Right-click the available version link and select one of the operations from the drop-down menu.

If Service Updates have not been installed for more than one release, older updates are displayed. The newest updates are on the top.

Option	Description
Download	The upgrade file is downloaded, but not installed.
Upgrade	The file previously downloaded is used during the upgrade. If there is no file available, the option is dimmed.
Download & Upgrade	The upgrade file is downloaded. The upgrade begins immediately after the download completes.
Release Notes	View the Release Notes.

f To begin the selected process, click **OK**.

The system reports that the upgrade is underway. After the upgrade file is downloaded and installed. The HCX system reboots. Allow a few minutes for the system to reinitialize.

g Open the HCX appliance management interface in a browser tab.

This option might not be available in HCX activated Public Clouds.

`https://hcx-ip-or-fqdn:9443`.

h Navigate to the dashboard and verify the registered systems display a healthy connected state.

i Open the **System Updates** interface and confirm that **Current Version** is updated.

If it is ever necessary to gather the logs related to upgrade operations, click the **Download Upgrade Logs** button.

2 To perform an offline (manual) update through the HCX Manager appliance management UI, follow these steps:

a Download the latest HCX Connector and HCX Cloud Manager Upgrade bundles from support.broadcom.com.

b Log in to the HCX Manager appliance management interface: `https://hcx-ip-or-fqdn:9443`.

c Navigate to **Administration > Upgrade**.

The appliance management screen appears.

d Click **Upgrade**.

A screen appears prompting you to upload the new bundle.

e Click **Choose File** and select the appropriate bundle for the site: HCX Connector or HCX Cloud Manager.

f Click **Continue** and follow the prompts to install the bundle.

During this process, HCX reboots and returns to service.

g Verify the upgrade:

- Log in to the HCX Manager UI: `https://hcx-ip-or-fqdn:443`.
- Navigate to **Administration > System Updates** and review the software version running on the site.

Results

With the HCX Managers upgraded, the HCX Service Mesh reflects that an update is available.

The HCX Managers apply the updates, reboot, and become operational in less than five minutes after rebooting. If the HCX Manager does not return to service within that time frame, contact VMware Support.

Upgrade the HCX Service Mesh Appliances

The Service Mesh appliances are upgraded independently of the HCX Managers: HCX Connector and HCX Cloud Manager. Service Mesh appliances are flagged for available updates anytime the HCX Manager has newer software available.

HCX has both system manager and Service Mesh appliances. You deploy HCX Connector and HCX Cloud manager in the vCenter. When creating the Service Mesh, the HCX manager deploys the Service Mesh appliances based on the site entitlements and Compute Profile configuration. The Service Mesh initiates a secure pipeline and activates the mobility and Network Extension operations between sites.

When upgrading Service Mesh appliances, consider the following best practices:

- Upgrade the HCX managers prior to upgrading the Service Mesh appliances.
- Service Mesh appliances (IX and NE) should be upgraded to the same version as HCX Managers.
- If mobility operations are in progress using IX appliances, or networks are extended using NE appliances, you can delay upgrading Service Mesh appliances until the operations have been completed, and plan the appliance upgrade during a separate maintenance window.

Note If Service Mesh appliances are running a lower version than the HCX manager, review the release notes for each HCX release to identify resolved or known issues associated with mobility or network extension operations and update the appliances accordingly. Also, review any Knowledge Base (KB) articles posted by VMware for HCX. For more information about Network Extension appliance versions, refer to KB 90115, 90117, 91086, 93726, and 96352.

Note Starting with HCX 4.4, HCX Connector and HCX Cloud Manager run VMware PhotonOS. If the HCX version of the Service Mesh appliances (IX, NE) is less than 4.4, you must upgrade the Service Mesh appliances to ensure they are also running PhotonOS.

Prerequisites

- You have upgraded site-paired HCX Managers to the same version on both the source and the destination site in a site pair.

Procedure

1. Log in to the HCX Manager at the source site.

   You initiate Service Mesh appliance updates only from the HCX source site.

2. Navigate to **Interconnect** > **Service Mesh** > **View Appliances**.

   Interconnect appliances show a green flag in the Available Versions column if there is an update available.

   Note If you are running OSAM in your environment, the upgrade notification does not appear for SGW and SDR appliances. Continue to follow the steps in this procedure to upgrade SGW and SDR appliances.

3. Select each HCX Service Mesh appliance that must be upgraded.

4 Click **Update Appliance**.

   The **Update Appliance** option is not displayed when there are no available updates.

   Note Network Extension appliances are available for in-service upgrades. For more information, see In-Service Upgrade for Network Extension Appliances.

5 Verify the Current and Available versions are valid.

6 To confirm the operation, click **Update**.

   The selected component and its peer component at the destination site are upgraded at the same time. Click the **Tasks** tab to view the upgrade progress details.

7 (Optional) If you are running OSAM in your environment, redeploy the SGW and SDR appliances:

   a Navigate to **Service Mesh > View Appliances**.

   b Select the SGW appliance.

   c Click **Redeploy**.

      Redeploying the SGW appliance automatically updates both the SGW and the SDR appliances at the source and destination sites.

Results

When the Service Mesh appliances converge to a Tunnel Up state, the upgrade is complete.

The Service Mesh appliances apply the updates, reboot, and become operational in less than two minutes after rebooting.

Rolling Back an Upgrade Using Snapshots

You can roll back an HCX Connector or an HCX Cloud Manager upgrade using VMware snapshots, which preserve the state and data of the HCX virtual machine at a specific point in time.

If a rollback is necessary due to an unexpected issue in the new version that cannot be resolved or workaround in a timely manner, follow this procedure.

Prerequisites

- Take and apply a vSphere snapshot for all site-paired HCX Managers. For the rollback procedure, a snapshot of the source and destination HCX Manager systems must exist. For more information, see Upgrading the HCX Manager.

Procedure

1 Roll back the HCX systems:

Note Use snapshots to roll back only the HCX Connector and the HCX Cloud Manager versions. After rolling back the HCX Manager versions, restore the HCX appliances through the Service Mesh configuration in HCX Interconnect UI. Do not attempt to restore snapshots of any fleet appliances.

Rollback Scenario	Process
The HCX Connector and HCX Cloud Manager version have been updated but the Service Mesh appliances have not been updated.	1 Roll back the HCX Connector and the Cloud Manager versions using the snapshots. 2 To validate the Service Mesh for the roll back version, navigate to **Interconnect > Service Mesh** and click **Resync**.
The HCX Connector and HCX Cloud Manager version have been updated, and the Service Mesh appliances also have been updated.	1 Roll back the HCX Connector and the HCX Cloud Manager versions using the snapshots. 2 For both HCX Connector and the Cloud Manger, navigate to **Interconnect > Service Mesh > View Appliances**. 3 Select all the appliances. 4 To restore the Service Mesh appliances for the rollback version, click **Redeploy**. 5 For each appliance, confirm that the tunnel status is **Up** and the appliance versions are reverted.

2 To verify the rollback, navigate to **Administration > System Updates** and confirm that the HCX Manager has the rollback version.

Removing VMware HCX

15

You can remove the HCX service from your environment by uninstalling the site-paired HCX Managers.

HCX supports on-going operations, including application migration and workload rebalancing while servicing the deployment. If HCX must be removed, stop all ongoing operations.

The procedure for uninstalling HCX can vary based on your environment and privileges. For deployments involving public clouds, uninstalling HCX can require actions from your cloud service provider.

Uninstalling HCX from VMware Cloud on AWS (VMC) deployments requires removing the service from both source and destination site. It is a self-serviceable process and there are no actions required from VMC.

The process for removing HCX has the following general workflow:

1 Stop all migrations and DR protections.

2 Unstretch the network, which removes Network Extension appliances.

3 Delete the Service Mesh, which removes the Interconnect and WAN Optimization appliances.

4 Remove the HCX Manager.

5 Remove the HCX plug-in from the vCenter Server.

For uninstallation details, see Uninstalling VMware HCX.

Read the following topics next:

Uninstalling VMware HCX

Uninstalling HCX requires removing the service from both source and destination site.

Note This procedure applies to non-VMC deployments. To uninstall HCX in VMC environments, see Uninstalling HCX in VMware Cloud on AWS Deployments.

A graceful uninstall of HCX appliances is always initiated from the source side. The process requires that HCX is fully functional, including site pairings and communication between source and destination site appliances.

Prerequisites

All migrations and replications, including DR operations, are finalized.

Procedure

1 Navigate to the HCX Manager Service UI.

2 Verify that no migration or protection operations are running.

3 To remove all network extensions from source-site data centers, complete the following substeps:

   a Go to **Services > Network Extension**.

   b Review each stretched network and decide whether you want the network to be connected on the destination site after uninstalling HCX.

   c Expand each extended network and click **Unextend**.

   The system displays information about the Unextend Network.

   d Under **Cloud Network**, expand the network entry.

   Note By default, the cloud network is disconnected from the cloud Edge Gateway after the network is unextended. This disconnection is done to prevent an edge gateway with dynamic routing activated from advertising the route of the network and causing a potential routing conflict with the network in the source site.

   e (Optional) Use the check boxes to keep the cloud network connected or force unextend the network.

   f Click **Unextend**.

4 In the HCX system containing the Service Mesh configuration, complete the following substeps to delete all Service Mesh instances:

   a Go to **Interconnect > Multi-Site Service Mesh > Service Mesh**.

   b For each Service Mesh, click **Delete**.

   Note Removing the Multi-Service Mesh from the source site also deletes it from the destination site.

5 To disconnect all HCX site pairings, complete the following substeps:

   a From the HCX dashboard, navigate to **Site Pairing**.

   b For each site pair, click **Disconnect**.

6 To remove the HCX Manager, complete the following substeps:

Note For public cloud deployments, contact your cloud service provider to remove HCX.

a At the destination site, navigate to the vCenter **Hosts and Clusters** tab.

b Expand the cluster where the HCX Manager is deployed and locate the virtual machine.

c Right-click on the HCX entry and power off the selection.

d Right-click on the HCX entry and select **Delete from Disk**.

e Repeat this procedure at the source site.

7 Unregister the HCX Plug-in from the vCenter Server using the instructions on how to remove or deactivate unwanted plug-ins using the KB article, https://kb.vmware.com/s/article/1025360.

Note Remove all HCX extensions that include com.vmware.hybridity in the path. Also, remove the following extensions:

- com.vmware.hcsp.alarm
- com.vmware.vca.marketing.ngc.ui

VMware HCX Troubleshooting

16

The following sections contain common HCX Manager troubleshooting scenarios, troubleshooting methodology, general information collection, and how to use built in diagnostic tools like the HCX Manager Central CLI.

Read the following topics next:

Starting SSH on the HCX Manager

This section describes how to start the **SSH Service** on the HCX Manager to access to the command-line interface.

To access to the HCX Manager shell, use a VMware Remote Console session in the vSphere Client or establish an SSH session. If the **SSH Service** was not activated during the initial HCX Manager installation, you must first activate it:

Procedure

1 Log in to the HCX Manager Appliance Management interface: https://*hcx-ip-or-fqdn*:9443.

2 Go to **Appliance Summary**.

3 Under System Level Services, locate the **SSH Service**.

4 Click **Start**.

Logging in to the HCX Manager Shell

This topic contains information on how to connect to the HCX Manager shell.

Prerequisites

You can log in to the HCX Manager shell using VMRC or an SSH session. The first-level access uses the admin account created during the initial installation of the HCX Manager. If requested to do so by support, you can switch the User to root once you log in with the admin account.

Procedure

1 Connect to the HCX Manager using VMRC or SSH.

2 When prompted for credentials, enter *admin* as the user name and then enter the password.

3 Switch to root by typing `su -` and providing the **`root password`**.

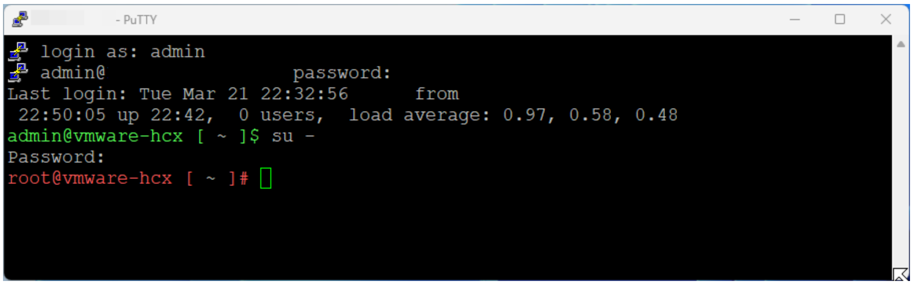

```
login as: admin
admin@                password:
Last login: Tue Mar 21 22:32:56      from
 22:50:05 up 22:42,  0 users,  load average: 0.97, 0.58, 0.48
admin@vmware-hcx [ ~ ]$ su -
Password:
root@vmware-hcx [ ~ ]#
```

Locating the VMware HCX System IDs Using VMware HCX Manager Shell

When working with support, you might have to provide the VMware HCX System IDs. You can get the IDs from the VMware HCX plug-in and from the HCX Manager shell.

Procedure

1 Connect to the HCX Manager shell using VMRC or SSH.

2 Switch user to root: `su`

3 Type `cat /common/location`

4 Note the System ID.

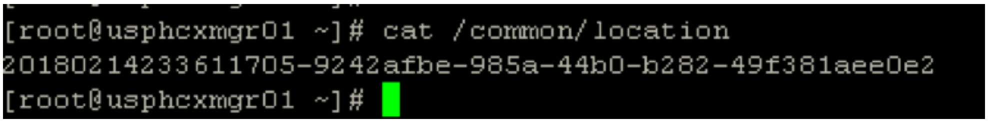

```
[root@usphcxmgr01 ~]# cat /common/location
20180214233611705-9242afbe-985a-44b0-b282-49f381aee0e2
[root@usphcxmgr01 ~]#
```

Locating the VMware HCX System IDs Using VMware HCX Plug-In

When working with support, you might have to provide the VMware HCX System IDs. The IDs can be obtained from the VMware HCX plug-in and from the HCX Manager shell.

Procedure

1 In the vSphere Web Client, navigate to the VMware HCX plug-in > **Administration** > **System Updates**.

2 Under Local HCX, in the Info column, click the ⓘ (information) icon. Doing so copies the System ID to your clipboard. Do the same to obtain the Remote HCX System ID.

3 Note the IDs and provide them to VMware when requested.

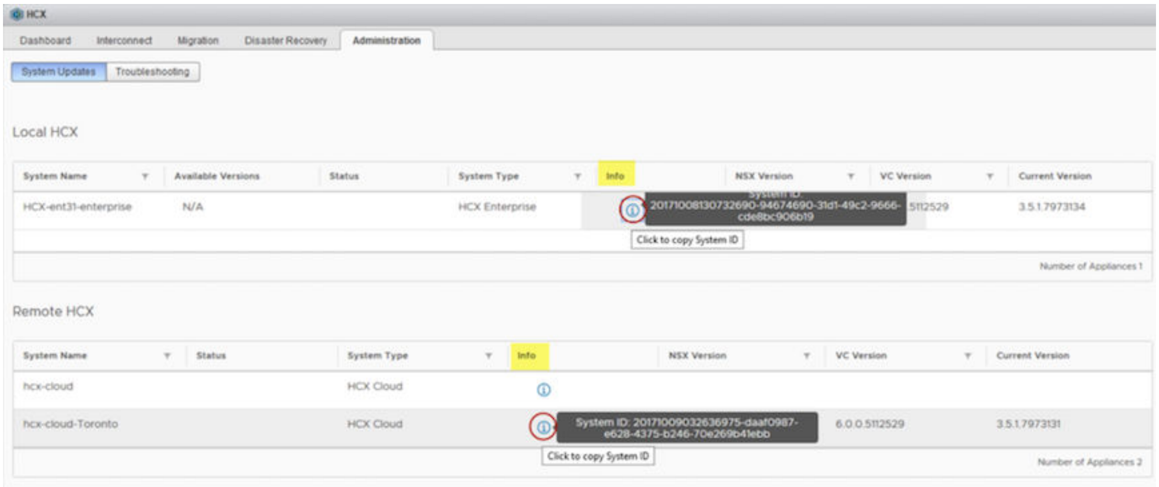

Local HCX

System Name	Available Versions	Status	System Type	Info	NSX Version	VC Version	Current Version
HCX-ent31-enterprise	N/A		HCX Enterprise	ⓘ System ID: 20171008130732690-94674690-31d1-49c2-9666-cde8bc906b19 Click to copy System ID	5112529		3.5.1.7973134

Number of Appliances 1

Remote HCX

System Name	Status	System Type	Info	NSX Version	VC Version	Current Version
hcx-cloud		HCX Cloud	ⓘ			
hcx-cloud-Toronto		HCX Cloud	ⓘ System ID: 20171009032636975-daaf0987-e628-4375-b246-70e269b41ebb Click to copy System ID		6.0.0.5112529	3.5.1.7973131

Number of Appliances 2

Using Central CLI to Connect to HCX Services

From the HCX Manager Central CLI, you can connect to the various VMware HCXHCX services for troubleshooting or gathering information.

Procedure

1 Access the CCLI on the HCX Manager. See Using the HCX Manager Central CLI.

2 Type `list` to view a list of HCX Manager nodes.

3 Identify the Node ID for the HCX Manager service to which you want to connect.

4 Type `go #` where # is the node **ID**.

5 Type `ssh`.

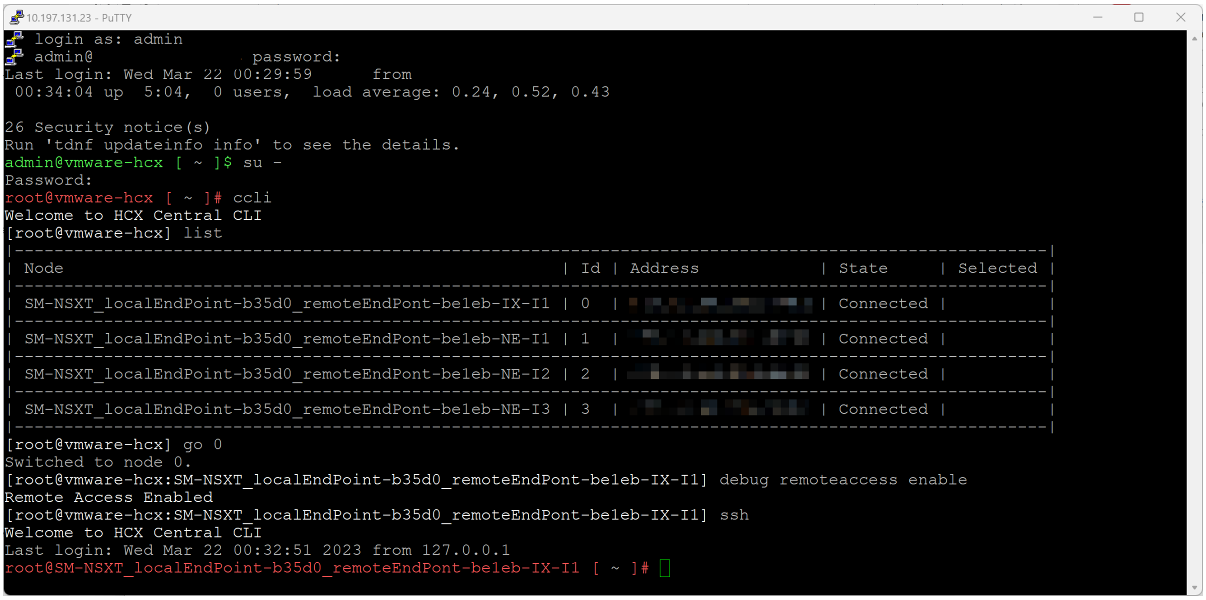

10.197.131.23 - PuTTY

```
login as: admin
admin@                    password:
Last login: Wed Mar 22 00:29:59       from
 00:34:04 up  5:04,  0 users,  load average: 0.24, 0.52, 0.43

26 Security notice(s)
Run 'tdnf updateinfo info' to see the details.
admin@vmware-hcx [ ~ ]$ su -
Password:
root@vmware-hcx [ ~ ]# ccli
Welcome to HCX Central CLI
[root@vmware-hcx] list
|-----------------------------------------------------------------------------------------------------------|
| Node                                                    | Id | Address                  | State      | Selected |
|-----------------------------------------------------------------------------------------------------------|
| SM-NSXT_localEndPoint-b35d0_remoteEndPont-be1eb-IX-I1 | 0  | [redacted] | Connected |          |
|-----------------------------------------------------------------------------------------------------------|
| SM-NSXT_localEndPoint-b35d0_remoteEndPont-be1eb-NE-I1 | 1  | [redacted] | Connected |          |
|-----------------------------------------------------------------------------------------------------------|
| SM-NSXT_localEndPoint-b35d0_remoteEndPont-be1eb-NE-I2 | 2  | [redacted] | Connected |          |
|-----------------------------------------------------------------------------------------------------------|
| SM-NSXT_localEndPoint-b35d0_remoteEndPont-be1eb-NE-I3 | 3  | [redacted] | Connected |          |
|-----------------------------------------------------------------------------------------------------------|
[root@vmware-hcx] go 0
Switched to node 0.
[root@vmware-hcx:SM-NSXT_localEndPoint-b35d0_remoteEndPont-be1eb-IX-I1] debug remoteaccess enable
Remote Access Enabled
[root@vmware-hcx:SM-NSXT_localEndPoint-b35d0_remoteEndPont-be1eb-IX-I1] ssh
Welcome to HCX Central CLI
Last login: Wed Mar 22 00:32:51 2023 from 127.0.0.1
root@SM-NSXT_localEndPoint-b35d0_remoteEndPont-be1eb-IX-I1 [ ~ ]# 
```

6 Use the `help` command to display available commands.

Gather Technical Support Logs

Locating the HCX Manager logs for review and knowing how to generate them is an important part of the troubleshooting process. It is helpful to gather the log files when experiencing an issue and when contacting support.

You can collect and download the HCX Manager Technical Support logs from either the HCX Manager UI (port 443) or the HCX Manager Appliance Management UI (port 9443).

Procedure

1 Access the log files:

   a From the HCX Manager UI, navigate to the log options at > **Administration** > **Troubleshooting** > **Download Log Bundles**.

   b From the HCX Manager System Management interface, navigate to **Administration > Troubleshooting > Technical Support Logs**.

2 Select the box next to one or more logs that you want to generate.

3 Depending on which UI you are in, click **Request** or **Generate** to prepare the log file bundles.

4 After the bundle is prepared, you are prompted to download it.

5 (Optional) To allow the HCX Admin to download logs from your site, click **Auto Approve** at the bottom of the page.

Note This setting is available only in the HCX Manager UI.

Gathering System Technical Support Logs

Locating the HCX Command and Control system logs for review and knowing how to gather them is an important part of the troubleshooting process. It is helpful to include at least the HCX Manager Technical Support log when experiencing an issue and reach for support.

Procedure

1 Log in to the HCX Appliance Management interface: `https://hcx-ip-or-fqdn:9443`.

2 Navigate to **Administration** > **Troubleshooting** > **Technical Support Logs**.

3 Select the box next to one or more logs that you want to generate.

4 Click **Generate**.

5 After the bundle is prepared, you are prompted to download it.

Gathering Technical Support Logs: Script Access

This section describes the process for gathering logs for VMware Technical Support using a standalone script.

In some cases, gathering system logs through the HCX UI or API calls might be unavailable. For example, if the HCX Manager appliance application or web engine stops or fails to start, you cannot generate the log bundles. You can generate and download the DB Dump and Core HCX log bundle using the standalone script included in this procedure.

Note The HCX Fleet Appliances IX/NE and OSAM Sentinel logs export option is not available

Prerequisites

Root user privilege is required for this procedure.

Procedure

1 Using SSH, log in to the HCX Manager as “admin” user and change to “root” user.

2 Change to the /opt/vmware/tools/ directory.

```
# cd /opt/vmware/tools/
```

3 Run the utility to review the script usage and options:

```
# ./export_tech_support_bundle.sh
```

```
Utility script to export Core HCX-Manager Logs and/or the Database Dump
Syntax: sh export_tech_support_bundle.sh [-l|d|h]
```

```
options:
 l Export Core HCX Manager logs
 d Export Database Dump
 h Usage help
```

4 Run the script with Core HCX Manager and Database Dump options:

```
# ./export_tech_support_bundle.sh -l -d
```

The system checks for sufficient disk space, and then exports the bundle under the `/tmp/techsupport` directory. The bundle filename and location display on the screen.

Note Use any file transfer protocols like SFTP/SCP or WinSCP to download the bundle directly from HCX Manager to an external server.

5 To avoid unnecessary utilization of the disk, delete the bundle from the /tmp/techsupport directory once support bundle has been successfully downloaded to an external server.

```
# cd /tmp/techsupport
# rm <file name>
```

What to do next

For assistance, contact VMware Technical Support.

Viewing Logs in the HCX Manager Shell

VMware HCX service logs are useful when troubleshooting failures.

Prerequisites

There are two key logs in the HCX Manager that can be reviewed and used when troubleshooting problems or to monitor system activities. Both are located in `/common/logs/admin` and they are the Application log (`app.log`) and the Web log (`web.log`). The Application log logs all activities for the app-engine service. The Web log logs all activities for the web-engine service. The process requires a good understanding of the HCX Manager system so it is best to review with a VMware support engineer.

Procedure

1 Use VMRC or SSH to connect to the HCX Manager shell.

See Logging in to the HCX Manager Shell.

2 Switch user to root: `su -`.

3 Change directory to `/common/logs/admin`.

4 From within this directory, you can open the relevant logs using standard Linux text commands.

When troubleshooting failures, search using keywords such as Fail, ERROR, exception, or migration.

Monitoring HCX Services from the Appliance Management Interface

You can monitor or restart the HCX services from the HCX Manager Appliance Management interface.

Prerequisites

There are several HCX Manager services critical for VMware HCX operations. Two key services to observe are the Application Service and the Web Service.

When working with a support team at VMware, you might have to confirm that these services are running or might have to restart them. You can view and restart the HCX Manager services in several places.

Important Do not restart services unless directed to do so by VMware Global Support Services.

Procedure

1 Log in to the HCX Manager Appliance Management interface: https://*hcx-ip-or-fqdn*:9443.

2 Navigate to **Appliance Summary**.

3 You can find all services and can monitor or restart them. The only service that is optional is SSH service. All others must always be running.

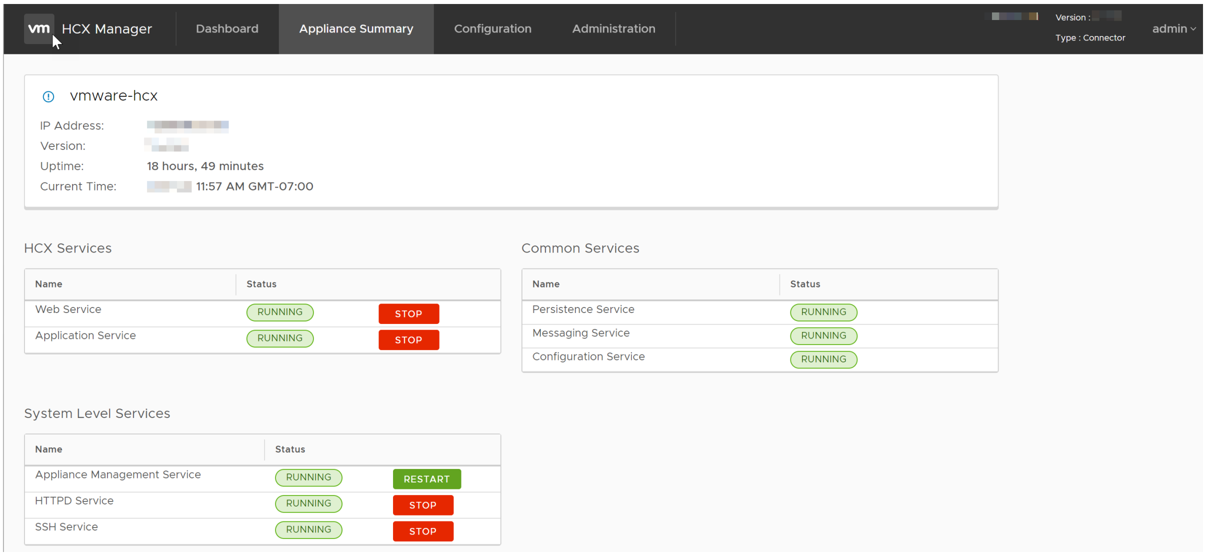

[redacted] Version : [redacted]
Type : Connector

vmware-hcx

IP Address: [redacted]
Version: [redacted]
Uptime: 18 hours, 49 minutes
Current Time: [redacted] 11:57 AM GMT-07:00

HCX Services

Name	Status	
Web Service	RUNNING	STOP
Application Service	RUNNING	STOP

Common Services

Name	Status
Persistence Service	RUNNING
Messaging Service	RUNNING
Configuration Service	RUNNING

System Level Services

Name	Status	
Appliance Management Service	RUNNING	RESTART
HTTPD Service	RUNNING	STOP
SSH Service	RUNNING	STOP

VMware HCX Manager Services from the CLI

You can manage the HCX Manager service using the HCX Manager CLI.

Procedure

1 Use VMRC or SSH to connect to the HCX Manager shell.

  See Logging in to the HCX Manager Shell.

2 Switch user to root: `su -`.

3 Type **systemctl** *action service_name*.

  - You can choose from one of the following actions: Status, Stop, Start, Restart.
  - The service name can the web-engine or the app-engine.

```
#systemctl status web-engine

#systemctl status app-engine

#systemctl stop web-engine

#systemctl restart web-engine
```

Viewing VMware HCX System State

You can view the HCX system state from the appliance management dashboard.

Prerequisites

For HCX to run properly, it is important that it has sufficient available resources. You can view the key system resources such as CPU, memory, and storage from the Dashboard section in the HCX Appliance Management interface. The dashboard section also provides other useful information such as the version that the HCX Manager is running, the uptime, its IP address, and current time. All useful information when reviewing logs or required by support.

Procedure

1 Log in to the HCX Appliance Management interface: https://<hcx-ip-or-fqdn>:9443

2 Navigate to **Dashboard**.

3 Review the CPU, Memory, Storage, Uptime, and Version.

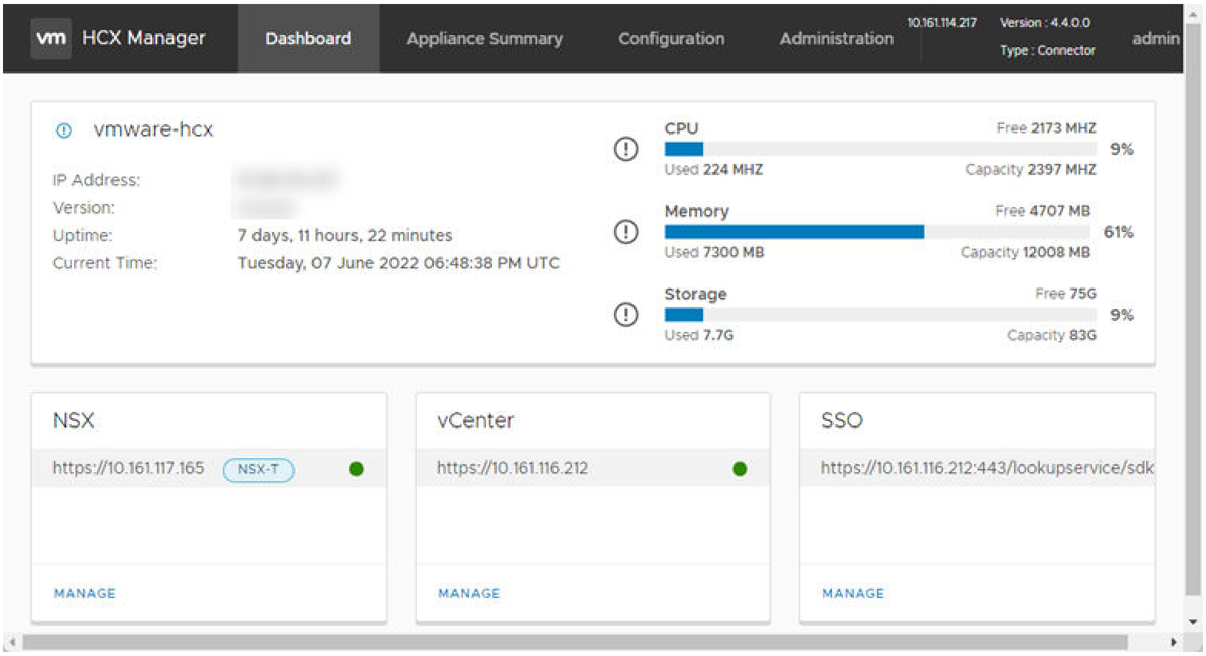

10.161.114.217 Version : 4.4.0.0
Type : Connector

vmware-hcx

IP Address: [redacted]
Version: [redacted]
Uptime: 7 days, 11 hours, 22 minutes
Current Time: Tuesday, 07 June 2022 06:48:38 PM UTC

CPU: Free 2173 MHZ, Used 224 MHZ, Capacity 2397 MHZ, 9%
Memory: Free 4707 MB, Used 7300 MB, Capacity 12008 MB, 61%
Storage: Free 75G, Used 7.7G, Capacity 83G, 9%

NSX: https://10.161.117.165 NSX-T ●
vCenter: https://10.161.116.212 ●
SSO: https://10.161.116.212:443/lookupservice/sdk

Viewing VMware HCX-Related Entries in the vSphere Task Console

Most VMware HCX Operations such as the initial appliance deployment, extending a network, or a migration can be monitored from the vSphere Web Client Task Console.

Procedure

1 Open the vSphere Web Client and navigate to **Home**.

2 Navigate to **Tasks**.

3 In the Task Console, filter the results by using **`HCX`** in the search filter.

4 Look for any failures or errors. If you see an error, you can review the logs to find additional details.

Using the HCX Manager Central CLI

The HCX Manager Central CLI is used for diagnostic information collection and secure connections to the Service Mesh.

The Central CLI (CCLI) on HCX Manager allows you to run commands available centrally on the site manager to view the run time state for HCX Manager services. The CCLI reduces troubleshooting time by providing centralized diagnostics and improves the security posture of the Service Mesh appliances by eliminating the need to run the SSH service on them. To use it, first you must activate the CCLI on the HCX Manager.

Procedure

1. Use VMRC or SSH to connect to the HCX Manager shell.
2. Type `ccli`.

   The CCLI is now available.
3. To begin exploring the CCLI, use the `help` command.

HCX Manager Password Recovery

You can reset the admin or the root password on Photon OS based HCX Connector or HCX Cloud Manager appliances.

HCX 4.4.0 and later releases support VMware Photon OS. If an account password is lost of forgotten, a Photon OS password recovery procedure is available to reset it.

A password reset can impact ongoing migration and configuration workflows. Allow those operations to complete before proceeding to recover the password.

The recovery procedure, which requires a reboot of the HCX Connector or HCX Cloud Manager, does not impact the following scenarios:

- Network Extensions actively forwarding.
- HCX protection operations for virtual machines in continuous synchronization.

To reset the root or admin password, see the VMware KB article 89212.

Updating System Passwords

Update the system admin or the root passwords from the command line interface.

Note For information about recovering a password, see HCX Manager Password Recovery.

Procedure

1. Login to the HCX Manager as the admin user: `ssh admin@`*IP address*.
2. Enter the admin password.

   The system prompt appears.

   ```
   [admin@hcxmgr ~] $
   ```
3. Enter the `passwd` command.

   ```
   [admin@hcxmgr ~]$ passwd
   ```

   Follow the prompts to complete the password change.

4 (Optional) Change the root user password:

   a At the system prompt, elevate to root user using the `su root` command:

   The root prompt appears.

   ```
   root@hcxmgr /home/admin]#
   ```

   b At the root prompt, enter the `passwd` command.

   Follow the prompts to complete the password change.

What to do next

Log out and log back in to confirm the change.

Monitoring VMware HCX Systems 17

VMware HCX native tools and views can be used to collect the current state of the system and general system health. Also, VMware HCX can be integrated with vRealize Log Insight and vRealize Operations using Management Pack.

Read the following topics next:

- Understanding the HCX Manager Dashboard
- vRealize Operations Management Pack for HCX
- DICE Integration for HCX
- VMware vCenter HCX Alarms

Understanding the HCX Manager Dashboard

The Dashboard provides a summary of HCX operations, data center locations, resource usage, status, and activity.

The Dashboard is the first screen that appears when you open the HCX Service UI.

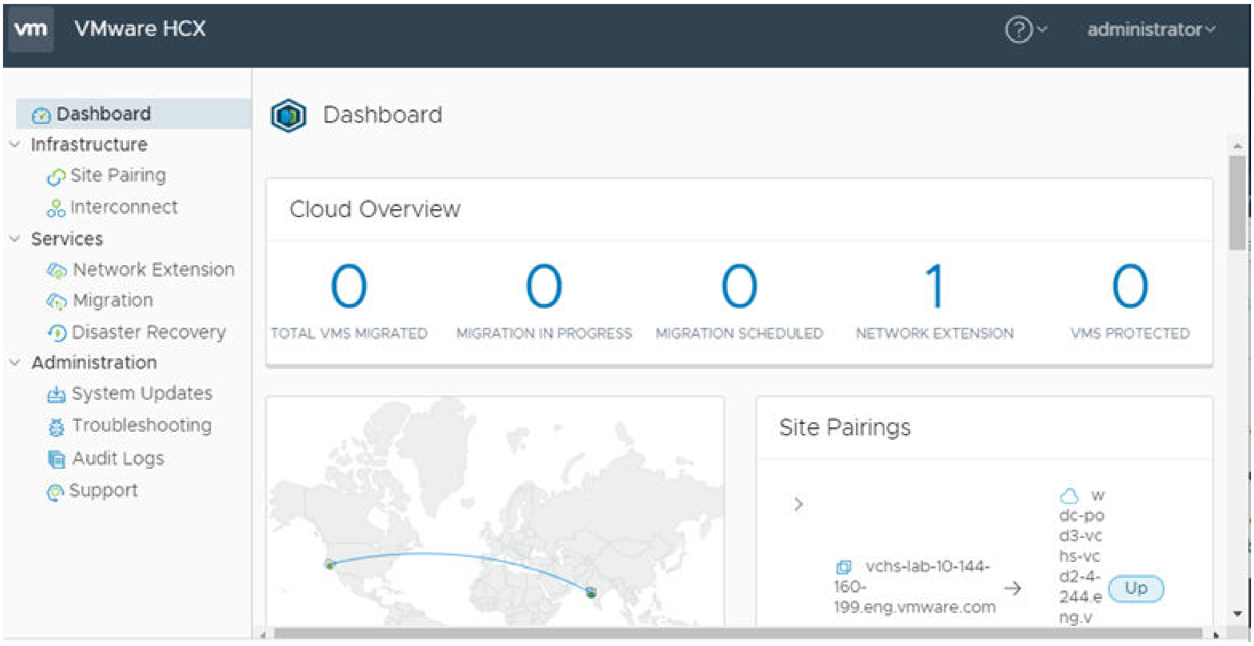

The Dashboard highlights various HCX functions in a set of panels. You can change settings related to those panels.

Note For HCX installations where the vCenter Servers are in linked-mode, the Dashboard includes information from all vCenter Servers registered to an HCX system.

Panel	Description
Cloud Overview	Lists HCX operations: ■ Number of virtual machines migrated ■ Number of migrations in progress ■ Number of scheduled migrations not started ■ Number of extended networks ■ Number of protected virtual machines for business continuity
Site Pairs	Displays connected site pairs and lists the pair status, Up or Down. To create a site pair, click New Site Pairing. To view detailed instructions for adding a site pair, see Adding a Site Pair.
Active Migrations	Displays ongoing migrations for the selected HCX system. Use the pull-down menu to change the source site.
Migrations Overview	Summarizes completed migrations for the selected HCX system. Use the date pull-down menu to display migrations for a specified period: ■ Last 6 Months ■ Last 3 Months ■ This Month
Cloud Resource Usage	For the selected source system, provides a summary view of resource usage for the site pair. Use the pull-down menu to change the source site.

Panel	Description
Alerts	Provides a comprehensive the list of logged Alert messages: Critical, Warning, Info. For a description of Alert messages and available actions, see Chapter 13 Manage Alerts.
Activity Logs	For the selected source system, displays a historical log of system tasks: ■ Job Type ■ Entity Name ■ Percentage of task completed. ■ Task status ■ Task Start Time ■ Task Completion Time To change the source site or to display tasks by status (All, Running, Failed), select from the pull-down menus . Use the search field to identify specific tasks or groups of tasks.

vRealize Operations Management Pack for HCX

The Management Pack (MP) for HCX adds monitoring capabilities with integrated dashboards and reports. It triggers problem alerts for the HCX services.

The Management Pack for HCX extends the Operations Management capabilities of vRealize Operations for HCX Hybrid Mobility, Interconnect Management and Data Center, and Cloud Migrations. For more information, see Management Pack for HCX.

DICE Integration for HCX

The Data Integrated Customer Engagement (DICE) tool uses customer utilization data to model the business benefits of VMware software-defined data center (SDDC) products. Through integration with DICE, you can upload the host and virtual machine inventory of the vCenter Server registered with an HCX.

Contact your account team for help with configuring and using this feature.

Prerequisites

Firewall rules allow access to the DICE portal through port 443.

Procedure

1 Navigate to the HCX Dashboard and select **Administration > DICE**.

The system displays the DICE configuration page.

2 Enter the DICE configuration parameters:

DICE Parameter	Description
API Key	Provides API key information for the REST authentication with the DICE portal. Obtain this key from the DICE website under Account Settings in your profile.
API Secret	Provides API secret key information for the REST authentication with the DICE portal. Obtain this secret from the DICE website under Account Settings in your profile. **Note** If you must change any of the DICE parameters in the future, you must reenter the secret key.
Customer ID	Obtain this ID from your account team.
Model ID	(Optional) Assigned by DICE after the first time you upload the inventory. The Model ID is unique to each HCX. **Note** If the Model ID is deleted from the DICE inventory, edit the configuration to remove the Model ID, and upload the inventory again.

3 Click **Save**.

4 Click **Upload VC Inventory to DICE**.

The HCX uploads the vCenter virtual machine and host inventory. In the DICE portal, a new model is created in the Library, and this Model ID is displayed in the HCX screen.

Note The time it takes to complete the upload depends on the size of the vCenter inventory. To refresh the inventory for the Model ID in the future, click **Upload VC Inventory to DICE** again.

What to do next

Conduct periodic updates once migrations and project milestones are completed to show the overall transformation progress. To compare results before and after, navigate to **Value Realization > Infrastructure Tracking**. Work with your account team for performing analysis on the Model in the DICE portal.

VMware vCenter HCX Alarms

The HCX service generates default vCenter Alarms that are reported to vCenter Server.

You can use these alarms to trigger additional actions or notifications.

HCX Event Alarm in vCenter (Event Code)	Description
HCX RAV Migration Error Encountered (65005)	The HCX RAV Migration did not succeed. See the HCX Migration Tracker for details.
HCX Bulk Migration Error Encountered (65006)	The HCX Bulk Migration did not succeed. See the HCX Migration Tracker for details.
HCX vMotion Migration Error Encountered (65007)	The HCX vMotion Migration did not succeed. See the HCX Migration Tracker for details.
HCX Cold Migration Error Encountered (65008)	The HCX Cold Migration did not succeed. See the HCX Migration Tracker for details.
HCX RAV Migration Cancelled (66001)	The HCX Replication Assisted vMotion migration is cancelled.
vSphere Bulk Migration Cancelled (66002)	The HCX Bulk migration is cancelled.
HCX vMotion Migration Cancelled (66003)	The HCX vMotion migration is cancelled.
HCX Cold Migration Cancelled (66004)	The HCX Cold migration is cancelled.
HCX Cloud Database Upgrade Failed (com.vmware.hcx.cloud.database.upgrade)	The HCX Cloud database upgrade has failed. Contact VMware Support.
HCX Connector Database Upgrade Failed (com.vmware.hcx.enterprise.database.upgrade)	The HCX Connector database upgrade has failed. Contact VMware Support.
HCX Interconnect Service Mesh Tunnel State Change (com.vmware.hcx.interconnect.TunnelStatusDownEvent, com.vmware.hcx.interconnect.TunnelStatusDegradedEvent , com.vmware.hcx.interconnect.TunnelStatusUpEvent, com.vmware.hcx.interconnect.TunnelStatusUnknownEvent)	The HCX Interconnect Tunnel status has changed. See HCX Service Mesh Diagnostics for details. To access Service Mesh Diagnostics, navigate to **Interconnect > Service Mesh > Run Diagnostics**. To view the current topology without running diagnostics, see **Service Mesh > View Topology.**
HCX Service connection is degraded (com.vmware.hcx.communication.HcxCommunicationCritical, com.vmware.hcx.communication.HcxCommunicationWarning,com.vmware.hcx.communication.HcxCommunicationUp)	The HCX Manager system is unable to reach `connnect.hcx.vmware.com`. After a grace period of 7 days, the HCX system becomes deactivated.
HA group State Change (com.vmware.hcx.interconnect.HAGroupFailedEvent, com.vmware.hcx.interconnect.HAGroupDegradedEvent, com.vmware.hcx.interconnect.HAGroupHealthyEvent, com.vmware.hcx.interconnect.HAGroupLicenseExpiredEvent,com.vmware.hcx.interconnect.HAGroupLicenseRenewedEvent, com.vmware.hcx.interconnect.HAGroupRecoveredEvent,)	A Network Extension High Availability appliance within an HA Group has changed state. For additional state information, see Monitoring Network Extension High Availability.

VMware HCX in the VMware Cloud on AWS

18

VMware HCX enables cloud onboarding without retrofitting the source infrastructure, supporting migration from vSphere 6.0+ to VMware Cloud on AWS without introducing the application risk and complex migration assessments.

Read the following topics next:

HCX Services for VMware Cloud on AWS

HCX for VMware Cloud on AWS includes support for HCX Advanced services and HCX Enterprise features and services with no additional license requirement or additional cost.

In addition to supporting HCX Advanced services, installing HCX for VMware Cloud on AWS provides support for these HCX Enterprise class services:

- Mobility Optimized Networking
- Network Extension High Availability
- OS Assisted Migration
- Replication Assisted vMotion
- Traffic Engineering features:
  - Application Path Resiliency
  - TCP Flow Conditioning

- Mobility Groups

> **Note** HCX Mobility Groups support integration with vRealize Network Insight, available as a separate license. This integration allows the creation of mobility groups from VMware vRealize Network Insight discovered applications to HCX for wave migration.

The HCX Connector (source) site inherits the available services from the HCX for VMware Cloud on AWS license, and no additional license is required at the source site.

For a detailed description of HCX services, see Chapter 2 System Services.

Requirements

- Site paring with HCX Cloud Manager is established through Internet network (INET), AWS Direct Connect, or VPN connections.
- HCX Interconnect and HCX Network Extension tunnels are established through INET and AWS Direct Connect only. Connectivity through a VPN tunnel terminated on the NSX Edge for the SDDC is not supported.

HCX Mobility Optimized Networking for VMware Cloud on AWS

Mobility Optimized Networking (MON) improves traffic flows for migrated virtual machines by enabling selective cloud routing (within the SDDC), avoiding a long round trip network path via the on-premises gateway. This feature is available in all VMware Cloud on AWS deployments.

- You can enable Mobility Optimized Networking when extending a network.
- You can enable or disable Mobility Optimized Networking on an existing network extension.

This section provides information specific to running MON in a VMware Cloud on AWS environment. For general configuration steps, and for additional requirements and limitations related to using MON, see Understanding Network Extension with Mobility Optimized Networking.

Limitations for MON on VMware Cloud on AWS

The folllowing items are not supported for MON on VMware Cloud on AWS:

- Modifying AWS "Networking & Security" tab properties .
- Advertising MON-enabled virtual machine static routes over Direct Connect or Transit Connect.
- Route-based VPN connections to native AWS VPCs. MON routes cannot be filtered and reaching the 100-route limit with VMware Cloud on AWS transitions the VPN to a down state.
- Optimization of traffic between MON-enabled migrated virtual machines and the SDDC management networks.
- Traffic between MON-enabled migrated virtual machines and Connected VPC Private IP addresses.

- Traffic between MON-enabled migrated virtual machines and virtual machines in other SDDCs (traffic over private Transit Connect).
- Traffic between MON-enabled migrated virtual machines across Multi-Tenancy Cloud Director Service boundaries.
- Traffic optimization between the extended segment and segments that are not directly connected to the same Tier-1 router.

  Note Reachability outside the Tier-1 can be configured, however, the traffic flows through the Tier-0 router are dependent on the network environment, and the design implementation might not be supported.

HCX Network Extension High Availability for VMware Cloud on AWS

Network Extension High Availability (HA) protects extended networks from disruptions associated with Network Extension appliance downtime.

Network Extension HA uses additional Network Extension appliances to create HA Groups and provide data path and appliance redundancy in an Active/Standby mode. This service functions the same in VMware Cloud on AWS SDDCs as in on-premises or private cloud environments. For more information and additional resources required to activate this feature, see Understanding Network Extension High Availability.

HCX OS Assisted Migration for VMware Cloud on AWS

HCX OS Assisted Migration works the same in VMware Cloud on AWS SDDCs as it does in any other HCX deployment.

Migrating virtual machines using OS Assisted Migration has the following requirements:

- The migration service is activated in both the source and the destination site Compute Profile.
- The migration service is activated in the HCX Service Mesh.
- Sentinel software is installed on all guest virtual machines requiring OSAM migration.

For a complete review of HCX OS Assisted Migration, see Understanding VMware HCX OS Assisted Migration.

To migrate virtual machines using HCX OS Assisted Migration, see Migrating Virtual Machines with HCX.

HCX Replication Assisted vMotion for VMware Cloud on AWS

HCX Replication Assisted vMotion works the same in VMware Cloud on AWS SDDCs as it does in on-premises or private cloud environments.

- The service must be selected in the Compute Profile of both the source and destination sites.
- The service must be activated in the Service Mesh deployed for the respective source and destination Compute Profiles.

To migrate virtual machines using Replication Assisted vMotion, see Chapter 8 Migrating Virtual Machines.

HCX Traffic Engineering for VMware Cloud on AWS

The Application Path Resiliency and TCP Flow Conditioning features define the HCX Traffic Engineering services. These services function the same in VMware Cloud on AWS SDDCs as they do in on-premises or private cloud environments.

- For new installations, optionally select these services when creating the Service Mesh.
- For existing installations, edit the Service Mesh to select these features.
- For existing installations, updating the Service Mesh for Application Path Resiliency has the following operational impact:
  - Redeployment of both the Interconnect and Network Extension appliances.
  - Disruption of Bulk and vMotion migrations. Quiese migration operations prior to finishing the Service Mesh update.
  - Brief disruption of traffic over extended networks.

To configure the HCX Service Mesh for these features, see Creating a Service Mesh.

HCX Mobility Groups for VMware Cloud on AWS

HCX Mobility Groups function the same in VMware Cloud on AWS SDDCs as it does in on-premises or private cloud environments. Support for Mobility Groups includes integration with vRealize Network Insight.

- Mobility Groups support assembling one or more virtual machines into logical sets for execution and monitoring of migrations as a group.
- Migration management functionality allows you to edit and delete groups, initiate and stop migrations, and schedule migrations.
- Through integration with VMware vRealize Network Insight, you can export waves of discovered applications to HCX for migration as Mobility Groups.

For more information about Mobility Groups, see Migrating Virtual Machines with Mobility Groups.

For more information about Mobility Group integration with vRealize Network Insight, see HCX Integration with vRealize Network Insight.

Topology Overview of VMware HCX on VMware Cloud on AWS

This section describes the behavior and features of VMware HCX services on the SDDCs operating with a set of network connectivity features provided by NSX-T.

Summary of Changes to VMware HCX for NSX-T Operations in Support of VMC SDDCs

- Updated component architecture uses the NSX Service Insertion Framework.
- The AWS Direct Connect with Private Virtual Interface is now supported. User-defined Private IP Subnets can be used during the VMware HCX Interconnect configuration.
- Network Extension L2 bridging is done with MAC Address learning on the Network Extension L2 switch port.

VMware HCX Architecture of SDDCs Supported by NSX-T

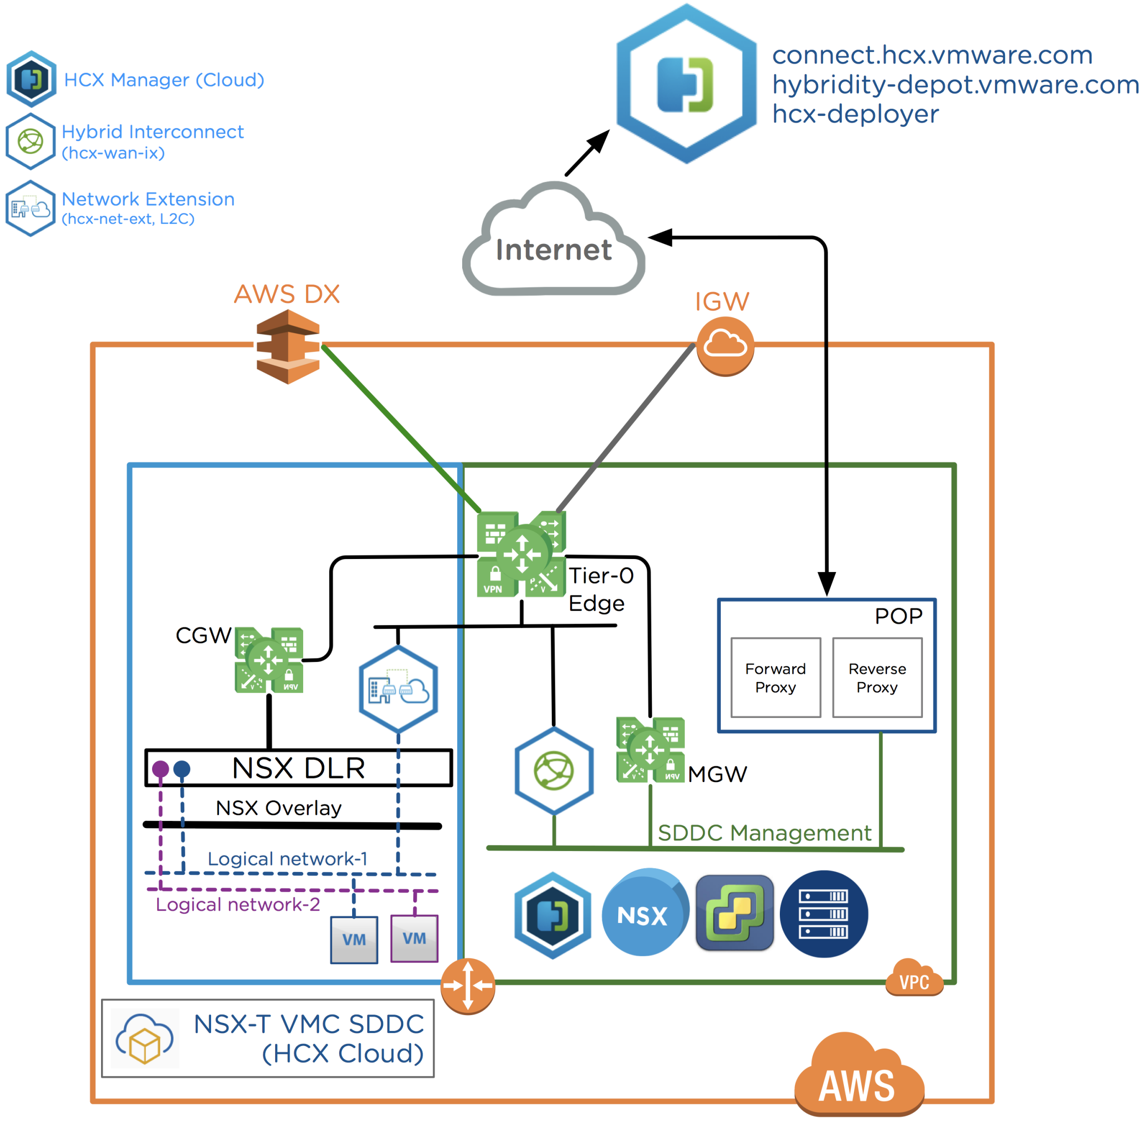

VMware HCX Features of SDDCs Supported by NSX-T

Feature	Details
VMware HCX Virtual Machine Migrations	■ VMware HCX vMotion for serial migrations. ■ VMware HCX Bulk Migration for scheduled, replication-based, parallel migrations. ■ VMware HCX Cold Migrations for powered-off virtual machines.
VMware HCX WAN Optimization	■ Deduplication, compression, and line conditioning of VMware HCX migration and protection network flows.
VMware HCX Network Extension	■ A maximum of eight networks can be extended to the SDDC per VMware HCX Network Extension appliance. ■ After a Network Extension operation, there is a five minute delay until the network is available for a migration operation. ■ Network Extension with Mobility Optimized Networking is available with VMware Cloud on AWS NSX-T SDDCs.
VMware HCX over AWS Direct Connect	■ VMware HCX supports connections over AWS Direct Connect with a Private Virtual Interface.

Deploying HCX from the VMware Cloud on AWS Console

VMware HCX is an add-on to the VMware Cloud on AWS SDDC. After activating the add-on from the VMware Cloud on AWS console, the HCX Cloud components are deployed and the HCX plug-in is available in the vSphere Client.

Prerequisites

- The user performing this procedure must have access to the VMware Cloud on AWS console.

Procedure

1 Log in to the VMware Cloud on AWS console at vmc.vmware.com.

Note For VMware Cloud on AWS GovCloud use https://www.vmc-us-gov.vmware.com.

2 Click **View Details**.

The SDDC interface opens.

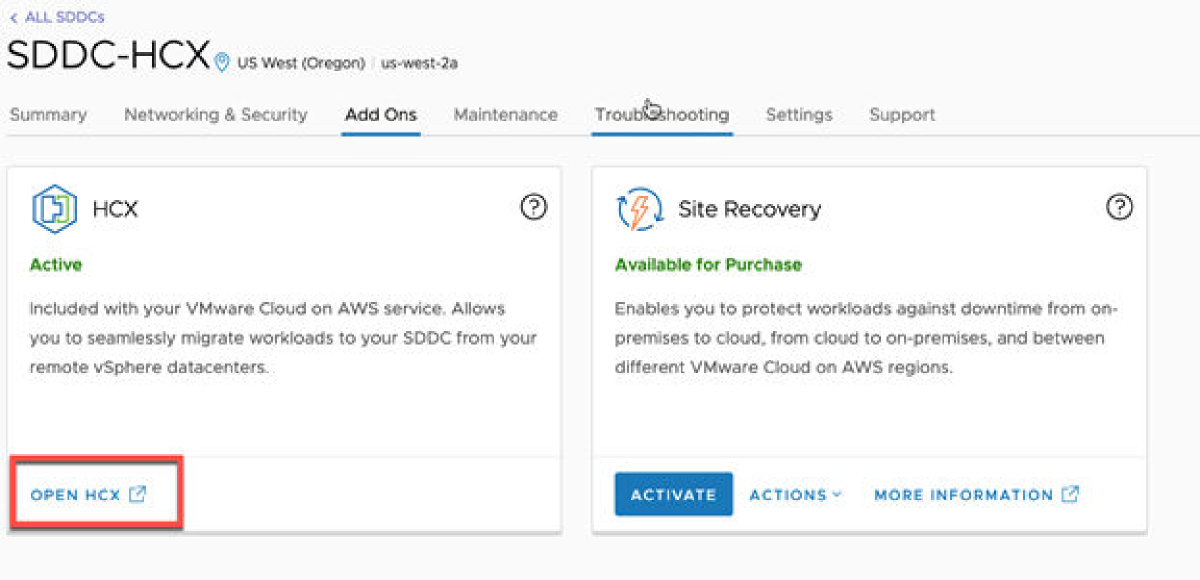

Included with your VMware Cloud on AWS service. Allows you to seamlessly migrate workloads to your SDDC from your remote vSphere datacenters.

Enables you to protect workloads against downtime from on-premises to cloud, from cloud to on-premises, and between different VMware Cloud on AWS regions.

3 On the **Add Ons** tab of your SDDC, click **Open HCX** on the **HCX** card.

The VMware HCX interface opens.

4 Navigate to the SDDC tab and identify the SDDC where you want to deploy HCX.

5 In the SDDC panel where you want to deploy HCX, click **Deploy HCX** and then click **Confirm** to initiate the deployment.

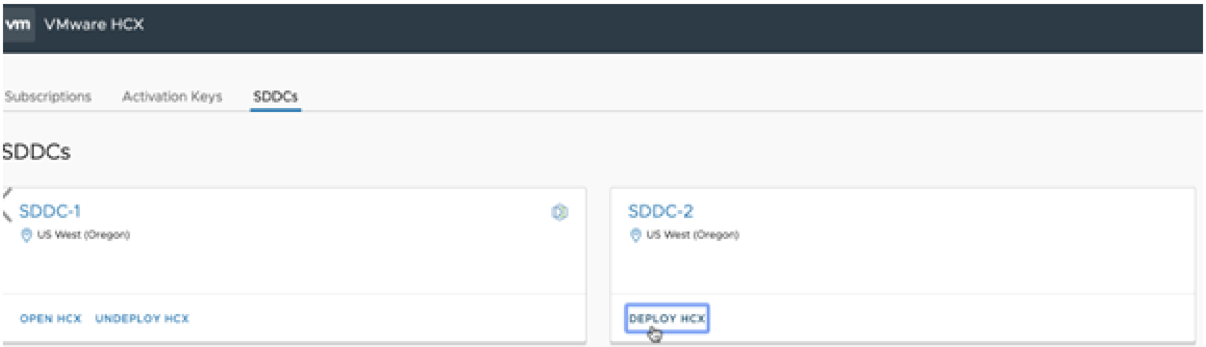

The VMware Cloud on AWS activation is created and displayed and the deployment begins. This step takes several minutes to complete.

6 While the HCX deployment is in progress, click **View Details** for a list of HCX deployment stages.

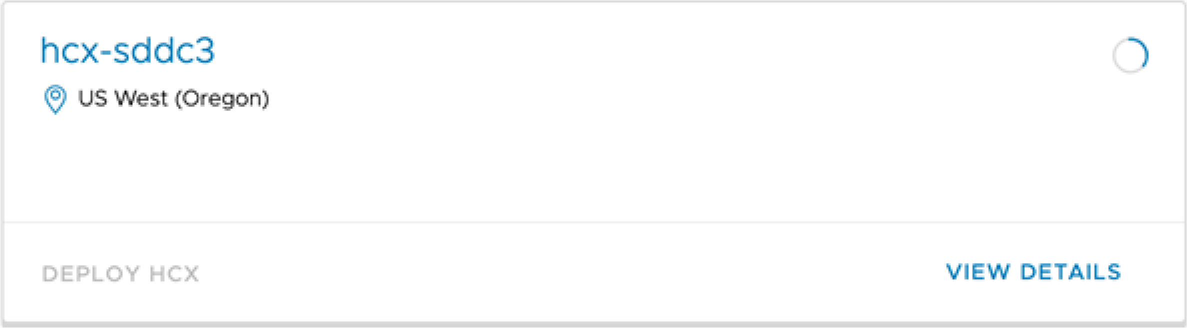

Each stage is a consolidation of the operations performed by VMware services for deploying HCX along with the status of each operation.

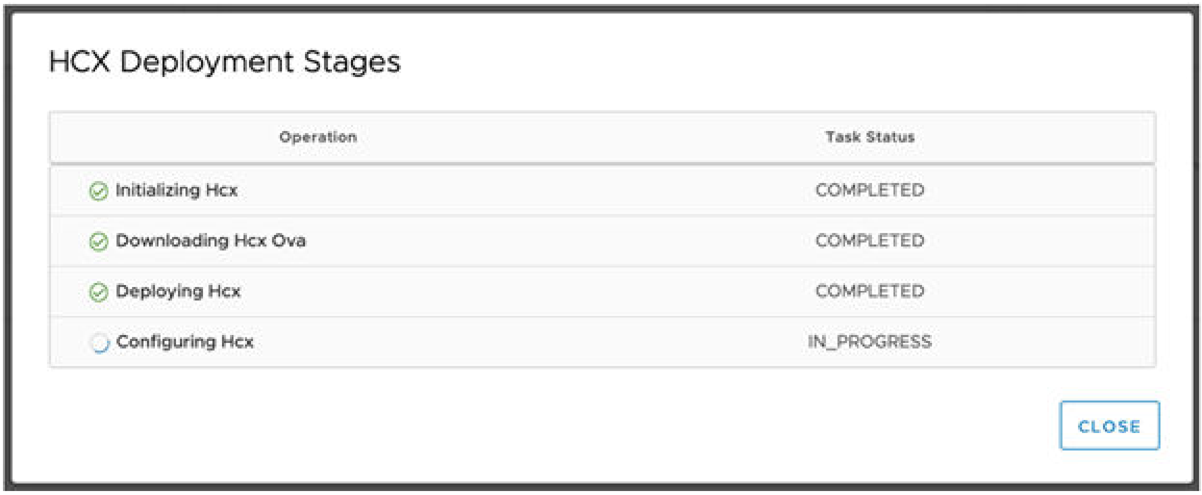

HCX Deployment Stages

Operation	Task Status
✓ Initializing Hcx	COMPLETED
✓ Downloading Hcx Ova	COMPLETED
✓ Deploying Hcx	COMPLETED
○ Configuring Hcx	IN_PROGRESS

After the deployment is complete, hcx_cloud_manager appears in the vCenter console.

7 Create a firewall rule to open the necessary ports to access the HCX Cloud Manager.

a From the VMware Cloud on AWS console, select **Networking & Security**.

b Under **Security**, go to **Gateway Firewall** and select **Management Gateway**.

c Click **Add Rule** and create a new inbound firewall rule with these parameters:

- Source: Where the connection to the HCX manager is coming from.
- Destination: **HCX**
- Services: **HTTPS (TCP 443)**

Note In cloud-to-cloud deployments (where two VMware Cloud on AWS SDDCs are paired) it is required to create an outbound HCX firewall rule. On this rule, set the Source as **HCX**, the Destination as **any**, and Service as **any**.

d To save the new rule, click **Publish**.

8 On the **Add Ons** tab of your SDDC, click **Open HCX** on the **HCX** card.

A new browser tab opens.

9 Navigate to the SDDC tab, and click **Open HCX**.

The VMware HCX Cloud service interface opens.

10 Enter cloudadmin@vmc.local and the password, and click **Log In**.

Note Use the vCenter password.

Results

The HCX Cloud Manager UI is available for HCX operations.

What to do next

Navigate to **Administration > System Updates** and download the HCX Connector OVA, which is needed for the on-premises HCX installation. Downloading the HCX Connector OVA is detailed in Downloading the HCX Connector OVA. For a complete installation workflow, see HCX Installation Workflow for HCX Public Clouds.

Setting DNS Resolution from Public to Private

Use this procedure to route HCX management communications over Direct Connect networks.

HCX for VMware Cloud on AWS can connect the source and the destination site using public or private networks. When switching between public or private networks, set the DNS resolution in VMware Cloud on AWS to use a public or private IP address. For example, changing the Management Network in the HCX Compute Profile to use a Direct Connect network type, means changing the DNS resolution in the VMware Cloud on AWS console to use a private IP address.

Caution Changing the DNS resolution can disrupt site-pairing connectivity while the system updates the local cache for the DNS entry. The time it takes to update the DNS server can vary depending on the TTL value. For HCX, the TTL value is 300 seconds.

Procedure

1 Log in to the VMC Console at vmc.vmware.com.

2 Select the organization .

3 Select the VMware Cloud on AWS service.

4 Click **SDDCs**.

5 Locate the SDDC and click **View Details**.

6 Click the **Settings** tab.

7 In the HCX Information section, expand the selection and click **Edit**.

8 In the Resolution Address field, use the drop-down menu to select the Private IP address, and click **Save**.

Note Use this same field to set a Private IP address to Public.

Configuring VMware HCX for Direct Connect Private Virtual Interfaces

The private virtual interface allows VMware HCX migration and network extension traffic to flow over the Direct Connect connection between your on-premises or cloud source environment and your destination SDDC.

Caution Ensure the IP Address Range configured does not overlap with the VMware Cloud on AWS management subnet CIDR block or any other IP range already in use for services in VMware Cloud on AWS. Overlap can cause routing and network reachability issues for those other components.

Prerequisites

- The AWS Direct Connect with Private Virtual Interface is supported on VMware Cloud on AWS SDDC backed by NSX-T networking.
- The SDDC must be configured to use the Direct Connect Private Virtual Interface.

  See Configure AWS Direct Connect Between Your SDDC and On-Premises Data Center.
- A private subnet that can be reached from on-premises over the Direct Connect with Private VIF or with Transit VIF, ideally reserved for VMware HCX component deployments.

Procedure

1 Log in to the VMware Cloud on AWS console at vmc.vmware.com.

2 Select your organization and data center (SDDC).

3 Select **Add Ons**.

4 Click **OPEN HCX** on the **HCX** card.

5 Navigate to the SDDC tab and click **OPEN HCX**.

6 Enter the cloudadmin@vmc.local user and credentials and click **LOG IN**.

7 Navigate to **Infrastructure > Interconnect**.

8 Click the **Network Profiles** tab.

9 In the Direct Connect network profile template, click **Edit**.

10 Enter the private IP address ranges reserved for VMware HCX.

11 Enter the Prefix Length and the Gateway IP address.

> **Note** A prefix length between 24–30 is recommended. HCX does not support prefix length 0, 31, or 32.

12 Click **Update**.

> **Important** Either **directConnectNetwork1** or **externalNetwork** must be configured as the Uplink Network Profile in the Compute Profile. The **Mgmt-app-network** profile cannot be used and can result in a Service Mesh deployment failure.

Results

When the Service Mesh is deployed, it uses the Uplink Network Profile, private IP addresses assigned by the user. The assigned IP addresses are reachable over the AWS Direct Connect.

Figure 18-1. VMware HCX over Direct Connect Private Virtual Interface

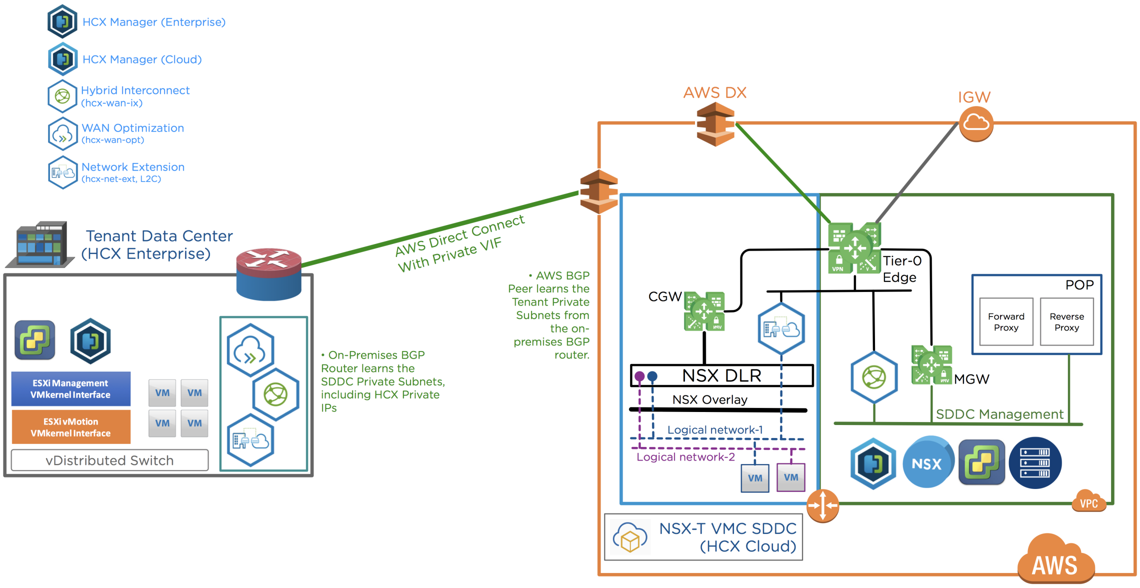
HCX Manager (Enterprise)
HCX Manager (Cloud)
Hybrid Interconnect (hcx-wan-ix)
WAN Optimization (hcx-wan-opt)
Network Extension (hcx-net-ext, L2C)
Tenant Data Center (HCX Enterprise)
ESXi Management VMkernel Interface
ESXi vMotion VMkernel Interface
VM
VM
VM
VM
vDistributed Switch
AWS Direct Connect With Private VIF
• On-Premises BGP Router learns the SDDC Private Subnets, including HCX Private IPs
• AWS BGP Peer learns the Tenant Private Subnets from the on-premises BGP router.
AWS DX
IGW
Tier-0 Edge
CGW
POP
Forward Proxy
Reverse Proxy
NSX DLR
NSX Overlay
Logical network-1
Logical network-2
VM
VM
MGW
SDDC Management
NSX
VPC
NSX-T VMC SDDC (HCX Cloud)
AWS

Configuring HCX for VMware Transit Connect

You can configure VMware HCX to use VMware Transit Connect™ for migration and network extension traffic.

VMware Transit Connect connects your SDDCs and VPCs to provide high-bandwidth, low-latency connections between SDDCs in the group and to other VPCs in the same region. The Direct Connect network profile is used for Transit Connect. For instructions on how to configure Direct Connect, refer to the section Configuring VMware HCX for Direct Connect Private Virtual Interfaces.

Scaling Out HCX Deployments in a Multi-Edge SDDC

An HCX Service Mesh configured with Multi-Edge traffic groups uses dedicated high-bandwidth network paths for HCX Network Extension and Migration operations.

About HCX with Multi-Edge SDDC

In the default configuration, an SDDC network has a single edge (T0) gateway through which all North-South traffic flows. You configure additional bandwidth for North-South traffic flows for direct-connect networks by creating one or more network traffic groups, each of which creates an additional T0 edge router in AWS. Traffic groups are created for each SDDC using the VMC Console.

When traffic groups are created in VMware Cloud on AWS, the HCX Cloud Manager at the destination site automatically detects those groups and adds them to the list of networks in the HCX Network Profile. HCX assigns a logical name to the traffic group, and the HCX Manager communicates any changes to the SDDC. HCX Manager can take advantage of the bandwidth provided by T0 routers to enhance HCX operations.

Requirements for HCX Multi-Edge Deployments

The HCX Multi-Edge solution requires AWS SDDCs configured with VMware Transit Connect.

Best Practices and Limitations

Review these best practices and systems limitations when scaling-out HCX deployments.

- It is a best practice to create a dedicated traffic group for HCX virtual machine migration.
  - The HCX-IX migration services cannot load balance across multiple traffic groups.
- It is a best practice to create a traffic group for HCX Network Extension.
  - If using a single traffic group for Network Extension, assign a /25 prefix to accommodate the maximum Network Extension appliance scale.
  - Create multiple traffic groups to distribute network extension traffic, which can be done using smaller prefixes.
  - The HCX SDDC supports a maximum of 100 HCX Network Extension appliances.

Configuring SDDC Traffic Groups in HCX

You can configure HCX to associate a traffic group with an Uplink Network in the Multi-site Service Mesh for improved migration or workload bandwidth.

By default, VMware HCX uses the management network for all uplink traffic. By overriding the default Uplink Network for the destination site with a specific traffic group, you isolate traffic for the HCX service. By isolating network traffic in this way, you gain any performance advantage by separating the traffic and utilizing the available bandwidth of the T0 router. For example, you can isolate migration and workload traffic by creating one Service Mesh for Bulk Migration and one Service Mesh for Network Extension. Within each specific Service Mesh, you then configure a unique traffic group for the Uplink Network.

Prerequisites

- Each SDDC is configured with traffic groups.
- A separate Service Mesh exists for each HCX service that is using traffic groups. For more information about creating and modifying a Service Mesh, see Creating a Service Mesh.

Procedure

1. In the HCX Cloud Manger UI, Navigate to **Interconnect > Network Profile**.

   The system displays the available SDDC traffic groups.

2. Select a traffic group Network Profile, and click **Edit**.

   They system displays the profile information for the traffic group.

3. Enter a logical name for the Network Profile.
4. Under **IP Pools**, add a range of IP addresses and enter the network prefix.

   The HCX Manager communicates with the SDDC, updating the traffic group with an association map that includes the HCX network profile name and prefix list. At the same time, the traffic group becomes available in the Service Mesh for overriding the uplink networks at the destination site.

5. Click **Update**.
6. In the HCX Manager UI at the source site (HCX Connect), navigate to **Interconnect > Service Mesh**.
7. Select the Service Mesh in which to add the traffic group, and click **Edit**.
8. Step through the Service Mesh dialog until you come to **Override Uplink Network profiles**.
9. Expand the list of **Destination Site Uplink Network Profiles**.
10. Select the named network to associate with the traffic group for the Service Mesh, and click **Continue**.
11. Step through the rest of the Service Mesh dialog, and click **Finish**.

Results

The HCX services activated in the Service Mesh are configured to use the T0 gateway created by the traffic group.

Note Before you can delete a traffic group from an SDDC, you must either delete the HCX Service Mesh that is using the traffic group or override the Uplink Network in the Service Mesh with a different network.

What to do next

Repeat this procedure to isolate traffic and enhance bandwidth for other HCX services.

Uninstalling HCX in VMware Cloud on AWS Deployments

Uninstalling HCX from VMware Cloud on AWS deployments requires removing the service from both the source and destination site.

A graceful uninstall of HCX appliances is always initiated from the source side. The process requires that HCX is fully functional, including site pairings and communication between source and destination site appliances.

Prerequisites

All migrations and replications, including disaster recovery (DR) operations, are finalized.

Procedure

1. Navigate to the HCX Manager Service UI.
2. Verify that no migration or protection operations are running.
3. To remove all network extensions from source-site data centers, complete the following substeps:

    a. Go to **Services > Network Extension**.

    b. Review each stretched network and decide whether you want the network to be connected on the cloud side gateway after uninstalling HCX.

    c. Expand each extended network and click **Unextend**.

    The system displays information about the Unextend Network.

    d. Under **Cloud Network**, expand the network entry.

    Note By default, the cloud network is disconnected from the cloud Edge Gateway after the network is unextended. This disconnection is done to prevent an edge gateway with dynamic routing activated from advertising the route of the network and causing a potential routing conflict with the network in the source site.

e (Optional) Use the check boxes to keep the cloud network connected or to force the network to unextend.

f Click **Unextend**.

4 For any of the unextended networks that are unused, complete the following substeps to remove them from the destination site:

Note Unextending networks does not remove them from the destination.

a Access the VMware Cloud on AWS management interface: `https://console.cloud.vmware.com`

b Select your organization and data center (SDDC).

c Select **Network & Security > Network > Segments**.

d Select the unextended network from the list and click **Delete**.

5 In the HCX Manager containing the Service Mesh configuration, complete the following substeps to delete all Service Mesh instances:

Note Removing a Multi-Service Mesh from the source site also deletes it from the destination site.

a Go to **Interconnect > Multi-Site Service Mesh**.

b For each Service Mesh, click **Delete**.

c Before proceeding to the next step, check that the Service Mesh no longer appears in the HCX Manager Service UI.

6 To disconnect all HCX site pairings, complete the following substeps:

a From the HCX dashboard, navigate to **Site Pairing**.

b For each site pair, click **Disconnect**.

7 (DX only) To remove direct connect private interfaces from the destination (VMware Cloud on AWS) site, complete the following substeps:

a Access the VMware Cloud on AWS management interface: `https://console.cloud.vmware.com`

b Select your organization and data center (SDDC).

c Select **Add Ons**.

d Navigate to the SDDC tab and click **Open HCX**.

e Enter the cloudadmin@vmc.local user and credentials and click **Log In**.

f Navigate to the **Infrastructure > Interconnect**.

g Click the **Network Profiles** tab.

h Select the direct connect network profile and click **Edit**.

   i Clear the IP ranges, Prefix length, and Gateway address.

   j Click **Update**.

8 To remove HCX Manager from the destination (VMware Cloud on AWS) site, complete the following substeps:

   Note For deployments between HCX clouds on VMware Cloud on AWS (cloud-to-cloud), repeat this procedure at the source site.

   a Access the VMware Cloud on AWS management interface: `https://console.cloud.vmware.com`

   b Select your organization and data center (SDDC).

   c Navigate to **Networking & Security > Gateway Firewall > Management Gateway** and delete any rules related to HCX.

   d Click **Add Ons**.

   The system displays all SDDCs with HCX deployed.

   e Click **Undeploy HCX**.

   While the HCX undeploy operation is in progress, click **View Details** for a list of HCX undeployment stages. Each stage is a consolidation of the operations performed by VMware services for undeploying HCX along with the status of each operation. When the undeploy operations have successfully completed, VMware Cloud on AWS automation cleans up SDDC HCX Manager services and removes the HCX Cloud Manager.

9 To remove HCX Connector on-premises, complete the following substeps:

   Note For cloud-to-cloud deployments using VMware Cloud on AWS, skip this step. This step applies only in on-premises to VMware Cloud on AWS deployments.

   a Navigate to the vCenter **Hosts and Clusters** tab.

   b Expand the cluster where the HCX Manager is deployed and locate the virtual machine.

   c Right-click on the HCX Manager virtual machine and power off the selection.

   d Right-click on the HCX Manager virtual machine and select **Delete from Disk**.

10 Unregister the HCX Plug-in from the vCenter Server using the instructions on how to remove or deactivate unwanted plug-ins using the KB article, https://kb.vmware.com/s/article/1025360.

   Remove all HCX extensions that include com.vmware.hybridity in the path. Also, remove entries with the following extensions:

   - com.vmware.hcsp.alarm
   - com.vmware.vca.marketing.ngc.ui

HCX for VMware Cloud on AWS GovCloud

19

VMware HCX is available as an add-on service in VMware Cloud on AWS GovCloud environments.

Note For VMware Cloud on AWS GovCloud, HCX complies with the FedRAMP High baseline requirements. Federal Risk and Authorization Management Program (FedRAMP) is a government-wide program that promotes the adoption of secure cloud services across the federal government by providing a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services.

HCX for VMware Cloud on AWS GovCloud supports similar services and features as HCX for VMware Cloud on AWS. For a list of supported services, see HCX Services for HCX on VMware Cloud on AWS GovCloud.

HCX for VMware Cloud on AWS GovGloud has similarities and differences with commercial versions of HCX from the VMware Cloud on AWS. For a list of these operational differences see, HCX Installation Workflow for VMware Cloud on AWS GovCloud.

You can find the process for deploying HCX for VMware Cloud on AWS GovCloud in Deploying HCX from the VMware Cloud on AWS Console. For a complete workflow describing the deployment process, see HCX Installation Workflow for HCX Public Clouds.

Read the following topics next:

- HCX Services for HCX on VMware Cloud on AWS GovCloud
- HCX Installation Workflow for VMware Cloud on AWS GovCloud

HCX Services for HCX on VMware Cloud on AWS GovCloud

HCX for VMware Cloud on AWS GovCloud includes all HCX Advanced features and select HCX Enterprise features.

HCX for VMware Cloud on AWS GovCloud includes the following Advanced features:

- Interconnect
- Network Extension
- Bulk Migration
- vMotion Migration

- Disaster Recovery

HCX for VMware Cloud on AWS GovCloud includes the following Enterprise features:

- Mobility Optimized Networking
- Network Extension High Availability
- Replication Assisted vMotion
- Workload Migration for NSX V2T
- Traffic Engineering features:
  - Application Path Resiliency
  - TCP Flow Conditioning
- Mobility Groups

Note HCX for VMware Cloud on AWS GovCloud does not support the HCX WAN Optimization service and OS Assisted Migration.

The HCX Connector (source) site activation keys are generated from the HCX for VMware Cloud on AWS GovCloud portal.

For a detailed description of HCX features, see Chapter 2 System Services.

HCX Installation Workflow for VMware Cloud on AWS GovCloud

The HCX installation workflow for HCX for VMware Cloud on AWS GovCloud deployments has similarities and differences from commercial versions of HCX for VMware Cloud on AWS.

The workflow for installing and activating HCX for VMware Cloud on AWS GovCloud is similar to VMware Cloud on AWS public clouds.

The procedure for deploying HCX for VMware Cloud on AWS GovCloud is provided in Deploying HCX from the VMware Cloud on AWS Console.

For a summary workflow describing the deployment process, see HCX Installation Workflow for HCX Public Clouds.

Requirements

- HCX Cloud Manager establishes site pairing through the Internet, AWS Direct Connect, or VPN connections.
- HCX for VMware Cloud on AWS GovCloud cannot be site-paired to non-GovCloud systems.
- HCX Interconnect and HCX Network Extension tunnels are established through the Internet and AWS Direct Connect only. Connectivity through a VPN tunnel terminated on the NSX Edge for the SDDC is not supported.

HCX for VMware Cloud Director with NSX

20

VMware HCX for VMware Cloud Director with NSX supports migration and network extension services between on-premises sites and VMware Cloud Director deployments.

Note For VMware Cloud Director compatibility with VMware HCX, refer to the Product Interoperability Matrix.

Read the following topics next:

HCX Services for VMware Cloud Director with NSX

HCX for VMware Cloud with NSX includes support for HCX Advanced services and for select HCX Enterprise features and services.

Activating HCX for VMware Cloud Director with NSX using the standard license key supports these HCX Advanced Services:

- Interconnect
- WAN Optimization
- Network Extension
- Bulk Migration
- vMotion Migration

In addition to Advanced services, adding the HCX Enterprise license provides access to these HCX Enterprise class services:

- Mobility Groups
- Replication Assisted vMotion Migration
- Traffic Engineering
  - Application Path Resiliency
  - TCP Flow Conditioning

For more information, see HCX Limitations for VMware Cloud Director with NSX-T.

For more informaton about HCX features and services, see Chapter 2 System Services.

HCX System Workflow with VMware Cloud Director

The HCX system installation and configuration workflow for VMware Cloud Director deployments has some differences from HCX deployments in other vSphere vCenter private or public clouds.

The general workflow for installing and activating VMware Cloud Director private clouds is similar to vSphere vCenter deployments. This workflow includes installation of the HCX Cloud Manager OVA at the destination site, installation of the HCX Connector OVA at the source site, and HCX activation at both sites. See HCX Installation Workflow for VMware Cloud Director Private Clouds.

As with other cloud deployments, HCX for VMware Cloud Director with NSX-T has the following Interconnect configuration workflow:

- Add a site pair
- Create a Network Profile
- Create a Compute Profile
- Create a Service Mesh

This section highlights the differences in these workflows for VMware Cloud Director with NSX-T and describes differences in Network Extension and Migration services configuration.

Administrator Roles

Configuring VMware Cloud Director includes creating an Organization Administrator role for each Organization, or Tenant, that is being managed by VMware Cloud Director. The HCX System Administrator has privileges to perform all HCX operations, while the Organization Administrator has access to perform specific HCX Cloud Manager functions related to that organization.

Site Pairing

The following differences apply to site pairing the source and the destination for HCX with VMware Cloud Director:

- Site Pairing can only be done with a Cloud Director managed Organization.
- The Remote HCX URL must be of the format - https://*hcx-cloud-ip-or-fqdn*/cloud/org/*org-name*.
- Credentials provided can be for the Organization Administrator.

Network Profile

Both the HCX System Administrator and Organization Administrator have privileges to update the Network Profile.

For HCX System Administrators, the Create Network Profile window includes an additional network type: External Networks. This type provides an option to create a Network Profile backed by a VCD External Network to use as the Uplink Network in the Compute Profile.

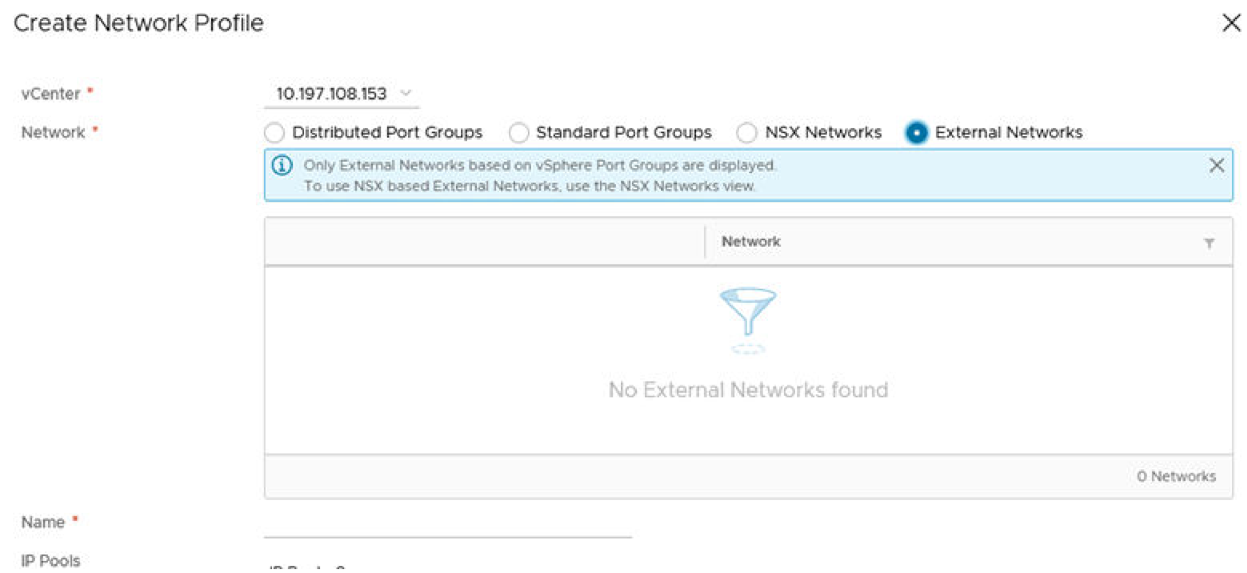

Create Network Profile

vCenter * 10.197.108.153

Network * ○ Distributed Port Groups ○ Standard Port Groups ○ NSX Networks ● External Networks

Only External Networks based on vSphere Port Groups are displayed.
To use NSX based External Networks, use the NSX Networks view.

Network
No External Networks found

0 Networks

Name *

IP Pools

For Organization Administrators, the Create Network Profile window displays only Org-VDC networks.

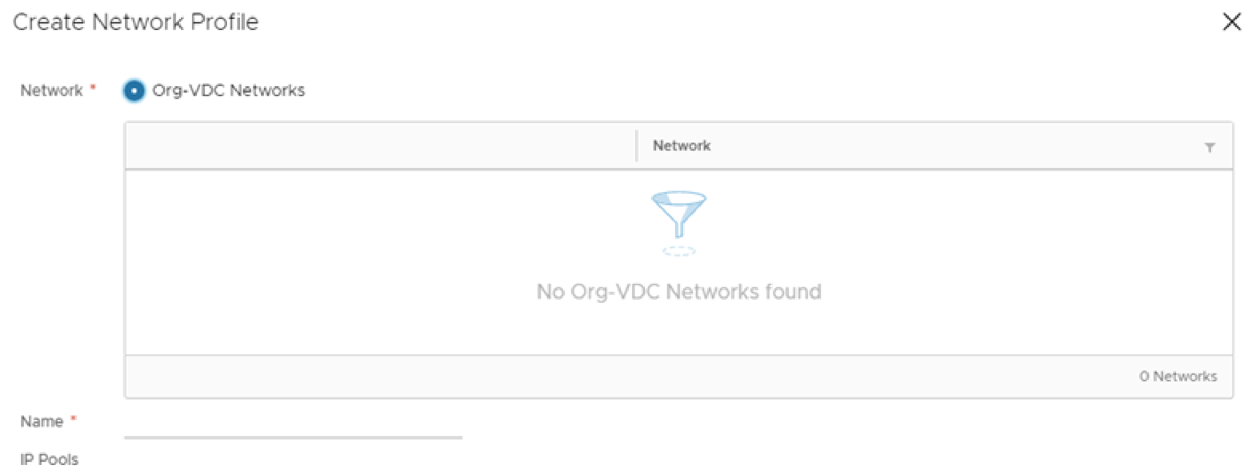

Create Network Profile

Network * ● Org-VDC Networks

Network
No Org-VDC Networks found

0 Networks

Name *

IP Pools

Organization Administrators can create Network Profiles, but the backing for these Network Profiles is limited to these Org VDC Network types: Imported and Routed.

Networks added by the Organization Administrator can only be used to override the Uplink Network Profile in the Service Mesh.

Compute Profile

Compute Profiles for the Cloud Director based environments can be defined only by the HCX System Administrator.

Organization Administrators who are logged in to the HCX Cloud Manager can view the Interconnect interface, but Compute Profile does not display.

Service Mesh

In the Service Mesh, the destination side Uplink Network Profile can be overridden with the Network Profiles created by the Organization Administrator.

Network Extension

Extending a network in HCX for VMware Cloud Director with NSX-T requires the following selections:

- Select the Destination Org VDC in which the extended network needs to be created.
- Select the Destination First Next Hop Router from the list of Edge Gateways associated with the selected Org VDC.

Extended networks appear as the Routed Org Network in the selected Destination Org VDC.

Virtual Machine Migration

The HCX Migration wizard parameters for VMware Cloud Director with NSX-T have the following differences:

- Storage profile selection is not available.
- The Destination Network is the Destination Org VDC network, which can have the type Routed or Imported.

In the VCD cloud environment, the source virtual machine (VM) gets migrated to the destination vCenter and then imported to the VCD as a vApp. Importing has some effect on how the system handles Bulk, vMotion, Replication Assisted vMotion, and Cold migrations:

- For Bulk migrations:
  - During the switchover phase, the VM is instantiated on the target VC. This powered-off VM is then imported to VCD, is powered-on, and has networks configured in the VCD layer.
  - Import is part of the general switchover workflow. If import is unsuccessful, the migration is considered unsuccessful and the system performs a rollback operation.
- For vMotion, Replication Assisted vMotion, and Cold migrations:
  - Import to VCD is part of the post migration actions. If import is unsuccessful, a warning appears, but migration is considered successful because the VM has already been moved to the destination VC.

HCX Limitations for VMware Cloud Director with NSX-T

Some HCX services, features, and functions are not supported in deployments using VMware Cloud Director with NSX-T.

The following restrictions and limitations apply:

- Unsupported HCX Services:
  - Network Extension High Availability
  - Mobility Optimized Networking
  - OS Assisted Migration
  - Transport Analytics
  - HCX Disaster Recovery
  - PowerCLI Module for HCX in VMware Cloud Director with NSX-T deployments
- HCX installation limitations:
  - In Cloud Director environments backed by multiple vCenters, select only a single vCenter Server in the HCX Manager appliance configuration.
  - Creation of HCX Network Profile in HCX Cloud with a vSphere Standard Switch port group.
  - VMware Cloud Director installation in source-side deployments.
  - Selecting the system virtual data center (VDC) resource pool as a deployment container in Compute Profile.

    The resource pool gets created on the vSphere vCenter when the provider VDC is created in VMware Cloud Director. Do not select Org or system VDC resource pools in the Compute Profile setting.
  - When the vCenter Server is configured in HCX, it must be consistent with the Cloud Director vCenter Server configuration using either the FQDN or the IP address in both environments.
- Unsupported VMware Cloud Director configurations:
  - Auto Discovery feature of VMware Cloud Director for migrated VMs.
  - VMware Cloud Director External Networks backed by a Tier-0 router.
  - Shared organization networks spanning multiple Orgs for Network Profile network backings or for Network Extension.
  - VMware Cloud Director enabled with data center group networking backed by NSX.