

VMWARE HCX DEPLOYMENT CONSIDERATIONS AND BEST PRACTICES

Table of Contents

Introduction	3
About HCX Terminology	3
HCX Services Overview	3
HCX Network Extension.....	3
HCX Virtual Machine Mobility	3
HCX WAN Optimization	4
HCX Disaster Recovery	4
Component Summary.....	4
HCX Installation Workflows.....	5
Considerations for Deployments to HCX Enabled Public Clouds	6
Component Considerations for an HCX Source Site	7
Considerations for Multi-Site HCX Deployments	7
Considerations for Deploying the HCX Manager OVA	9
Considerations for Deploying the HCX Manager in Environments with HTTPS Proxy Servers	10
Considerations for HCX Deployments with vCenter Servers Linked Mode	12
Considerations for Integrating HCX with an Active Directory Domain	13
Considerations for Creating HCX Compute Profiles	15
Considerations for Selecting the Service Clusters in a Compute Profile	16
Considerations for Selecting the Deployment Resources in a Compute Profile	18
Considerations for Creating HCX Network Profiles	19
Considerations for Creating an HCX Service Mesh	20
Considerations for HCX WAN Optimization Service Deployments	21
Considerations for Deploying the HCX Network Extension Service	22
Considerations for Maximum Transmission Unit (MTU)	23
Considerations for Deploying HCX for OS-Assisted Migrations (OSAM)	24
About the Author	26

Introduction

This considerations paper describes VMware HCX best practices and implementation considerations. It has been prepared from iterating through hundreds of HCX implementations on a variety of production architectures and deployment scenarios.

Although there was a considerable effort in collating the best practices information, some deployment scenarios may not be covered. This paper is not intended as a comprehensive guide for implementing VMware HCX in every scenario. Please send any questions or feedback regarding this document to your VMware account team.

For more information see the HCX User Manual at docs.vmware.com/en/VMware-HCX.

About HCX Terminology

The VMware HCX service is operated in the context of multiple environments (specifically a source and destination environment).

We may use Source, On-Premises, Legacy or HCX Enterprise interchangeably to refer to the source vSphere installation in an HCX deployment.

Similarly, we may use Destination, Target Cloud, or HCX Cloud interchangeably to refer to the destination vSphere or vCloud Director(VCD) based installation in an HCX deployment.

In less common single vCenter Server deployments, Source and Destination refers to the clusters.

HCX Services Overview

HCX provides infrastructure abstraction, high performance network extension with advanced services like proximity routing, virtual machine mobility and disaster recovery services with data reduction and WAN line conditioning built in.

HCX Network Extension

HCX Network Extension connects vDS, NSX or Nexus 1000v networks at the source site to an NSX Logical Switch at the destination site. This service can expedite the use of the destination environment's resources by allowing virtual machines to be migrated into the networks without re-IP or complicated VM transformations, by leveraging the routing and security policies at the source site.

HCX Network Extension with Proximity Routing leverages the strong integration with NSX-v Dynamic Routing to achieve local ingress/egress for virtual machines as they are migrated. The HCX Proximity Routing for Layer 3 Aware VM Mobility whitepaper explores this feature with more details.

HCX Virtual Machine Mobility

Virtual machines can be moved to and from HCX-enabled vSphere private and public cloud environments using multiple HCX migration technology. HCX provides version compatibility across legacy and modernized sites.

- HCX Bulk Migration uses the vSphere Replication protocol to transfer up to multiple virtual machines in parallel. virtual machines are “rebooted” into the target site and can be transformed to the latest VM Hardware/VM Tools available. With the Bulk migration option, virtual machines can have their vNIC IP addresses updated as part of the migration operation.
- HCX vMotion uses the VMware vMotion protocol to transfer individual virtual machines. Used with HCX Network Extension for zero-downtime migrations of applications that are sensitive to downtime.
- HCX Cold Migration uses the VMware NFC protocol. This migration type is automatically selected when transferring powered-off virtual machines.
- HCX vMotion with vSphere Replication combines Bulk and vMotion to deliver zero downtime failover for Virtual Machines prepared in parallel (in preview with VMware Cloud on AWS).
- HCX OS-Assisted Migration is a migration technology (currently in beta) which uses an agent in the Windows/Linux Virtual Machines to enable migrations to vSphere from non-vSphere (KVM) environments.

HCX WAN Optimization

The HCX WAN Optimization service improves the performance characteristics of private lines or internet path by applying WAN optimization techniques like data reduction and WAN path conditioning.


HCX Disaster Recovery

VMware HCX uses advanced WAN Optimization to protect on-premises applications by replicating the data to an HCX enabled provider or private cloud deployment. In the event of a disaster, VMware HCX recovers the networking layers. Traffic routes are maintained as before the disaster resulting in high-speed disaster recovery with low downtime. There is no reconfiguration of IPs, which removes complexity and enables either partial or full site recovery.

Component Summary

The HCX technologies are delivered as three distinct service-level virtual appliances, deployed as peers at the source and destination environments.

The HCX enabled sites are paired and then service components are deployed simultaneously at the source and destination sites whenever services are enabled for the selected site pair.



- **HCX Manager**

The HCX manager component is deployed from an OVA, is integrated with the vSphere environment, and enables it to be connected with other HCX-enabled environments to deliver HCX services. HCX Manager is typically deployed one-to-one with each vCenter Server.

HCX Manager is deployed at the source site. Currently this manager is labeled “Enterprise” but this label will be deprecated in the future.

HCX Manager is also deployed at the target site. The HCX manager at the destination automates the deployment of peer appliances when a service mesh is created at the source site.

An HCX Site Pair is created when the source HCX Manager is connected to a destination site’s HCX Manager.

- The HCX Interconnect (HCX-WAN-IX) service appliance is deployed for migration and DR services. This component automatically tunnels to its peer at the remote site and provides an encrypted service path for migration services.
- The HCX WAN Optimization (HCX-WAN-OPT) service appliance is deployed to provide data deduplication and WAN path conditioning (TCP optimization and traffic fairness).
- The HCX Network Extension appliance (HCX-NET-EXT) is deployed when Network Extension services are enabled. This component automatically tunnels to its peer at the remote site and provides an encrypted service path for migration services.

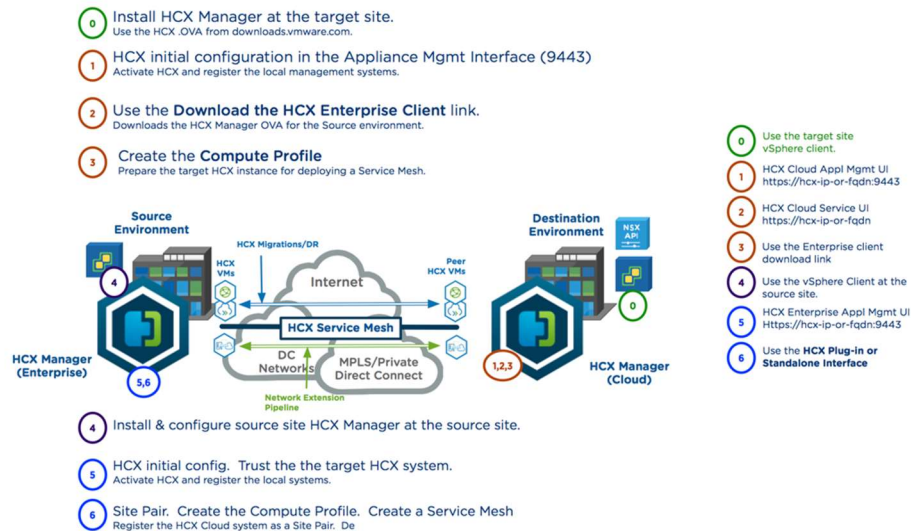
HCX Installation Workflows

There are installation workflows for deploying HCX, depending on the type of destination environment. Use the following links and understand the deployment flow.

- [HCX Installation Workflow for HCX Public Clouds.](#)
Summary of steps when deploying HCX for VMware Cloud on AWS.
- [HCX Installation Workflow for vSphere Private Clouds.](#)
Summary of steps when deploying HCX with a vCenter Server based target environment.

- **HCX Installation Workflow for vCloud Director Private Clouds.**
Summary of steps when deploying HCX with a vCloud Director based target.

Deployment Flow for Private Cloud Installations



Complete the HCX Installation Checklist after reviewing and understanding the deployment and considerations and best practices presented in this document.

NOTE: The remainder of this document is prepared with the assumption of a basic understanding of HCX services, and the source to destination nature of the technologies.

Considerations for Deployments to HCX Enabled Public Clouds

For example, if designing for HCX services between an on-premises environment with a public Cloud (e.g. VMware Cloud on AWS).

- The HCX enabled public cloud environment is considered the destination environment.
- The public cloud provider automation manages the deployment of the HCX Cloud Manager component. This is not true with private targets.
- The public cloud infrastructure will be running modernized SDDC software (i.e. vSphere 6 and above, NSX 6.2 and above.).
- There will be a Public Access URL that an HCX source site can use for the site pairing operation.

<https://hcx-cloud-public-ip-or-fqdn>

- The public access URL for HCX must be resolvable and reachable over TCP-443(HTTPS) from the source site's HCX Enterprise Manager.
- The cloud provider's management layer firewall may need to be updated to make HCX Manager reachable from the manager at the source.
- In a private to public deployment, the activation keys for the HCX Enterprise Manager system on-premises will come from the public cloud provider.
- For various reasons, VMware HCX features may be available on some public cloud providers, but not on others. Consider what is available in the target HCX enabled public cloud, and what is required for a successful HCX deployment.

This paper avoids listing public cloud provider-specific considerations. Public cloud providers publish information for enabling the HCX service in their public cloud.

The public cloud provider published instructions should supersede any conflicting information that may have been included erroneously in this paper.

VMware Software Component Considerations for an HCX Source Site

- Review the **HCX Hardware, Interoperability and Activation Requirements** in the HCX Documentation.
- Using the HCX User Interface (the vCenter Web Client HCX Plugin) for any of the operations above requires vCenter Server 5.5U2 in the on-premises vSphere environment. For older environments, the HCX Standalone UI can be used (by launching HCX HTTPS in a browser).
- HCX can coexist with vSphere Replication 8.1 or later. Older versions of vSphere Replication will disrupt HCX Bulk Migration or any other replication based HCX operation.

Considerations for Multi-Site HCX Deployments

HCX multi-site deployments are supported. In this document, multi-site refers to any installation with more than a single source VC/HCX connected to one or more destination VC/HCX.

- HCX Manager for Enterprise is deployed 1 to 1 with vCenter Server at the Source site(s).
- HCX Manager for Cloud is generally deployed 1 to 1 with vCenter Servers at the Destination site.

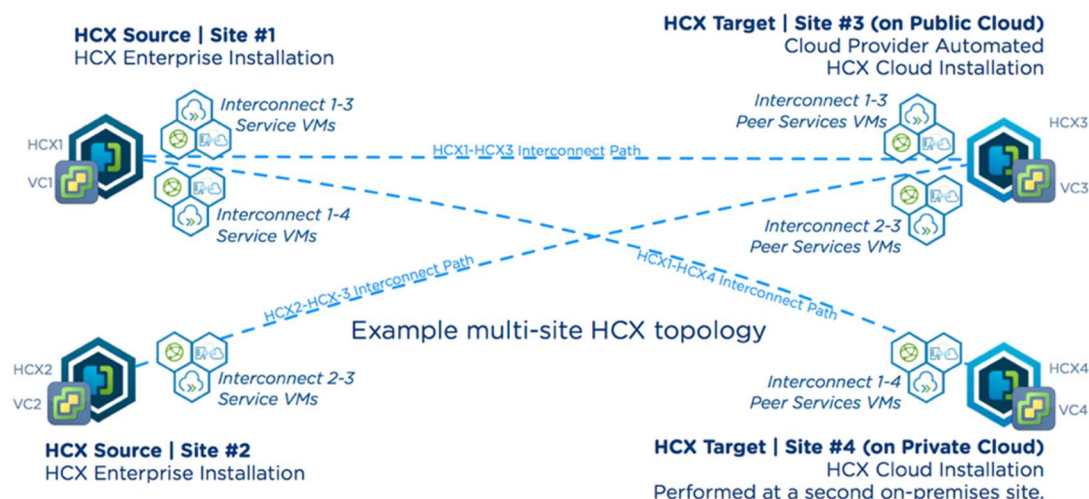
* Under specific conditions, a single HCX Manager may be able to register with multiple VC/NSX sets (Secondary NSX Managers are not supported).

- A source site HCX Enterprise site can connect to many destination sites with HCX Cloud.

Multiple source HCX Enterprise sites can connect to a single destination HCX Cloud

HCX Service appliances will be instantiated for each site pair. In the illustration below, Source Site #1 is paired with destination Sites #3 and #4.

A single HCX/VC pair establishes both of the connections, and service appliances are deployed for each pair.



- Activation Keys for multi-site deployments is based on the deployment type for the destination sites. In the illustration above, destination site #3 is HCX on a public cloud, and site #4 is a private vSphere installation.

In this scenario, activation keys for site #3 come from the HCX Public Cloud. Site #4 activation will come from NSX Data Center + licensing. The source site #1 can be activated with keys from site #3 or site #4.

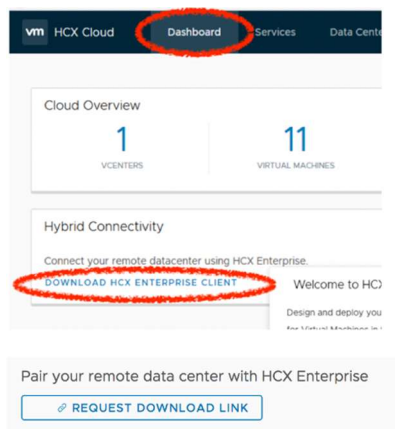
- Multi-Site architectures have been tested for up to 10 site pairs per source or destination HCX Manager system.

Considerations for Deploying the HCX Manager OVA

- The HCX Enterprise and Cloud Manager performs management and control functions for the HCX service. Mobility, Replication and Extension data flows are not transmitted/received or traverse the HCX Managers.
- Deploy the OVA using distributed port group, datastores and compute resources designed for Management virtual machine.
 - For private to private deployments, the HCX Cloud Manager should be deployed first at the destination site. Use the HCX Installer from downloads.vmware.com. This installer will install HCX Cloud and automatically download the latest updates.

Once the HCX Cloud Manager installed Use the Download HCX Enterprise Client link to download the HCX Enterprise Manager OVA for the source site installation.

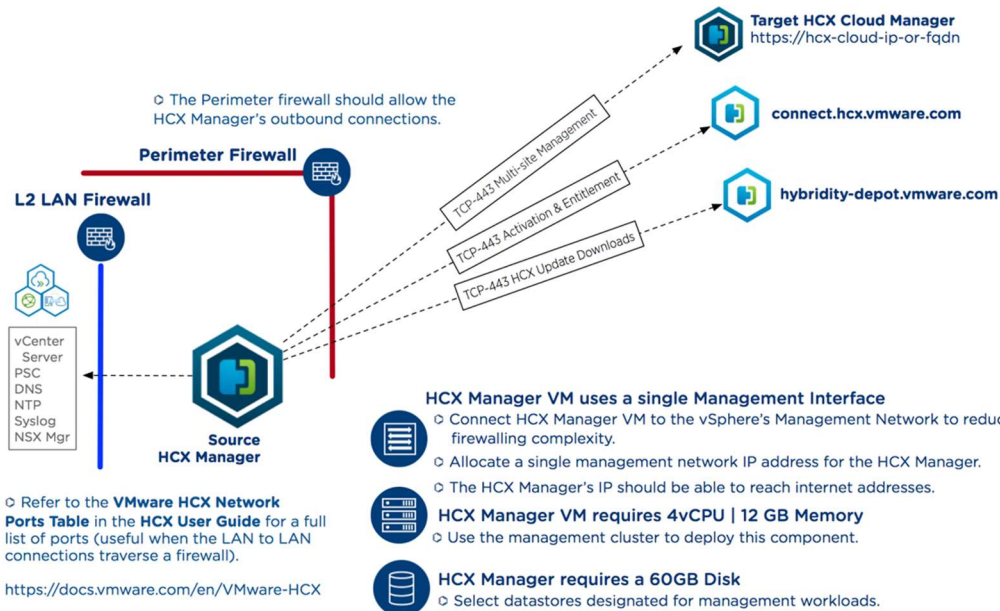
- For private to public deployments, HCX Cloud is already installed by the cloud provider. Access the HCX dashboard, and use the Download HCX Enterprise Client link to download the HCX Enterprise Manager OVA for the source site installation.



The HCX Cloud will generate a download link so a compatible HCX Enterprise Manager OVA can be downloaded.

- During the OVA deployment, provide a functional DNS server that can provide resolution for both external targets like connect.hcx.vmware.com and internal targets, like the vCenter Server and ESXi hosts FQDNs.
- During the OVA deployment, provide live NTP servers that are reachable. Confirm the HCX Manager system is using synchronized time.

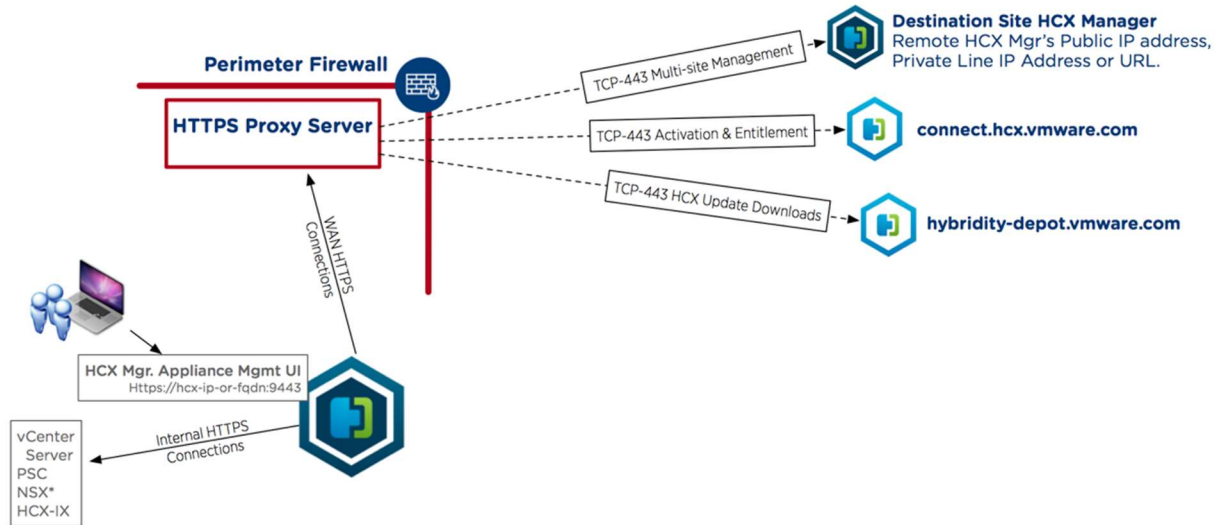
- The HCX Manager IP address should be able to route to both internal and external targets.



- At the source site, the perimeter firewall should allow outbound HTTPS/TCP-443 connections from the HCX Enterprise Manager to:
 - <https://connect.hcx.vmware.com>
 - <https://hybridity-depot.vmware.com>
 - <https://hcx-cloud-mgr-ip-or-fqdn>
- At the destination site, the perimeter firewall should allow outbound HTTPS/TCP-443 connections from the HCX Cloud Manager to:
 - <https://connect.hcx.vmware.com>
 - <https://hybridity-depot.vmware.com>
 Also, inbound connections from the HCX Enterprise Manager.

Considerations for Deploying the HCX Manager in Environments with HTTPS Proxy Servers

HCX supports service operations with an HTTPS proxy server in the path.



- If the environment uses a proxy server for outbound HTTPS connections, it should be defined in the HCX Manager's appliance management interface.
 - Once a proxy server is defined, all HTTPS connections are sent to the proxy server. HCX Manager makes internal HTTPS connections to vCenter Server and the HCX Interconnect appliance.

For communications to work properly, internal subnets should be added as exceptions to the proxy configuration


Note: Kerberos proxy servers are not supported.

Considerations for HCX Deployments with vCenter Servers Linked Mode

VMware HCX supports linked mode for single-pane operation at the source site.

- The HCX Enterprise Manager is deployed and paired with each vCenter Server and registered to a common platform services controller.

HCX with vCenter Servers in Linked Mode

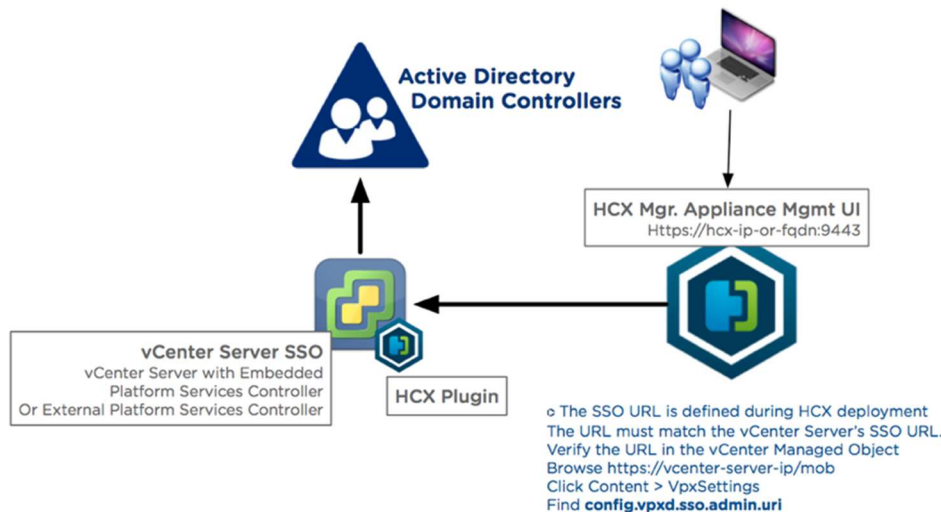


- The resulting behavior is that HCX operations for the combined inventory can be initiated from any of the vCenter Servers that are in Linked Mode.

Considerations for Integrating HCX with an Active Directory Domain

VMware HCX Supports Active Directory logins through integration with vCenter Single Sign-On.

- HCX Enterprise at the source site and HCX Cloud at the destination site can be integrated with distinct Single Sign-On Domains, there is no requirement for the sites to share identity domains.



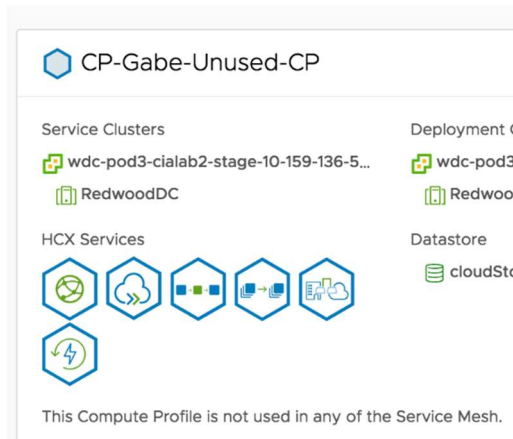
- The HCX appliance management interface should contain a valid SSO URL, matching the vCenter Server's SSO URL.
- The HCX appliance management interface has a role mapping configuration screen where one can define the SSO Active Directory or vSphere.local domain groups that can perform HCX operations.
- An Example best practices configuration for Active Directory based authentication:
 - o Create a new hcx-admins local group in the vCenter Server SSO Users and Groups screen.
 - o Add the Active Directory Groups that will operate HCX services to the hcx-admins@vsphere.local group.
 - o Right click the vCenter and Add Permission. Assign the Administrator Role to the new hcx-admins@vsphere.local group.

Open the HCX appliance management Role Mapping interface. Add the hcx-admins@vsphere.local group to the roles.

Considerations for Creating HCX Compute Profiles

A Compute Profile defines the HCX services that will be allowed to run in the configuration, the compute/network boundaries, as well as the compute, storage and network settings that HCX will use to deploy the Multi-Site Service Mesh virtual appliances.

- Compute Profile creation is identical at the source and destination HCX systems.
- Creating a Service Mesh requires at minimum one Compute Profile in HCX at the source, and one Compute Profile in HCX at the destination.
- A single Compute Profile can be used for all clusters as Service Clusters when the clusters shares a common vMotion/Replication/vSphere-Mgmt networks.
- Multiple Compute Profiles are required when the clusters that will be designated as Service Clusters have different vMotion/Replication/vSphere-Mgmt networks.
- A single Compute Profile can be used with multiple Service Mesh configurations.
- A Compute Profile can be created in advance, without being immediately applied to a Service Mesh configuration.

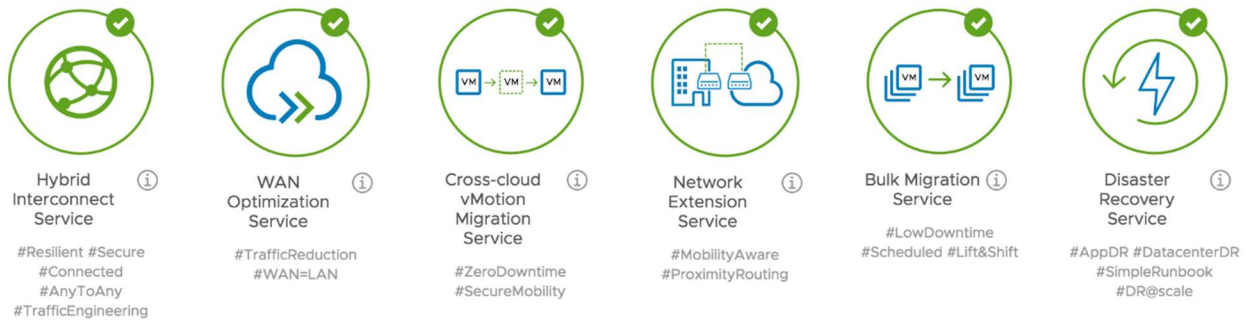


- A Compute Profile contains the following elements:
 - Compute Profile Name
The CP name should be meaningful, especially in environments that will use multiple compute profiles.
 - Services to be enabled (HCX Services)
 - Service Clusters (Clusters in Source Compute Profile)
 - Service Resources (Datacenters or Clusters in the Destination Compute Profile)

- Deployment Resources (Clusters|Resource Pools / Datastores)
- Management Network Profile
- Uplink Network Profile
- vMotion Network Profile
- vSphere Replication Network Profile
- Uplink Network Profile
- Distributed Switches for Network Extension

Considerations for Selecting Services to be Enabled in a Compute Profile

Select Services to be enabled



- This step can be used to explicitly prevent the unselected HCX service from being selected during a service mesh deployment.
- In most deployments all services can be left selected when a Compute Profile is created.
- When a Service Mesh is being configured, if a service was unselected in the source or destination Compute Profile (or both), it will be grayed out in the Service Mesh interface.
- When a Service Mesh is being configured, steps related to the unselected service will be excluded.

For example: If Network Extension is unselected, the Service Mesh creation will exclude the Network Extension Appliance Scale Out step.



Considerations for Selecting the Service Clusters in a Compute Profile

Select the Service Clusters for which HCX services should be enabled.

- The Service Cluster(s) must exist in the single vCenter Server that was registered to the HCX Enterprise or Cloud Manager.
- Virtual Machines in clusters that are designated as Service Clusters in the Compute Profile will be valid objects for HCX migrations and DR operations.
- VM networks in DVS/clusters that are designated as Service Clusters in the Compute Profile will valid objects during HCX Network Extension operations.
- The Service Clusters in a Compute Profile must have common cluster networks.
- When the Service Clusters have different cluster networks (for example, when each cluster has a dedicated vMotion subnet), they should be separated into different Compute Profiles
- As an example, Service Clusters can be used to provide an HCX services boundary for Non-Production and Production, and Test/Dev infrastructure.

Assuming Prod and Non-Prod workloads are separated by clusters.
Assuming all Prod clusters share cluster networking.
Assuming all Non-Prod clusters share common networking


Following the example:

- An initial Non-Prod CP is created to initially allow Non-Production HCX operations, and all non-production clusters are selected as the Service Clusters.
- A Non-Prod Service Mesh is created, the Non-Prod Compute Profile is selected. HCX Migration/Extension appliances will be deployed for Non-Production HCX service traffic.

Once the service mesh appliances are deployed and HCX operations become

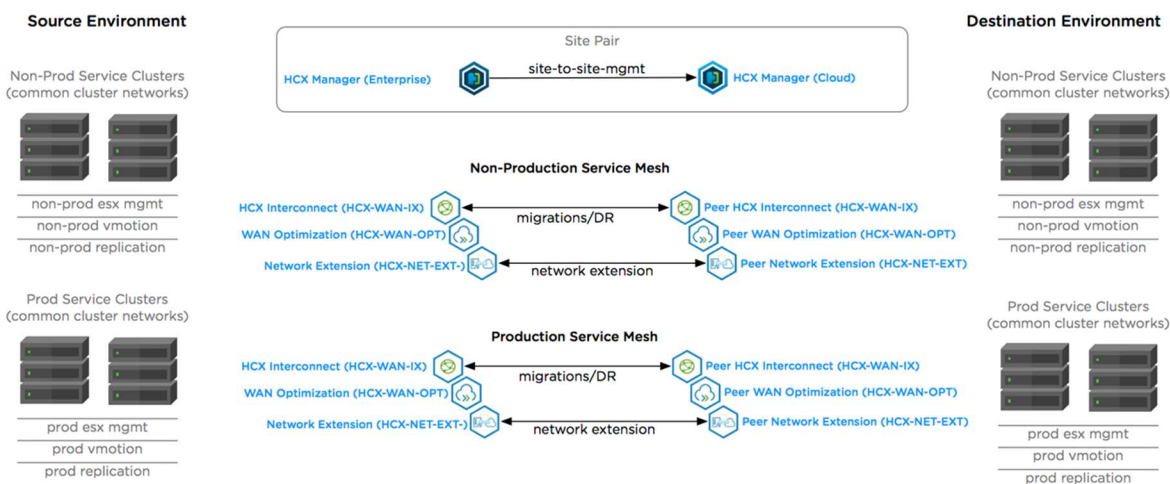
VMWARE HCX DEPLOYMENT CONSIDERATIONS AND BEST PRACTICES

available, HCX Migration/Extension operations on objects outside of the selected Non-Production will fail.

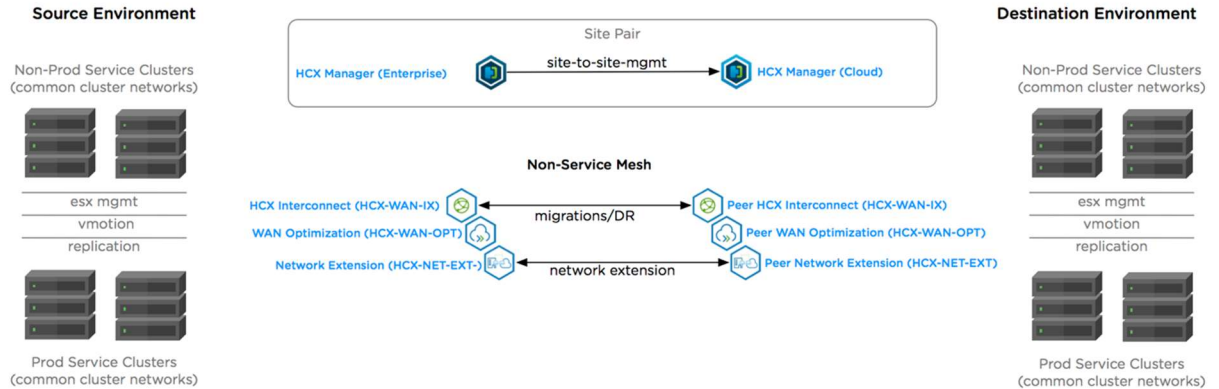


- Later, a Prod-Compute-Profile is created for Production HCX operations. The remaining clusters (which are all production clusters) are selected as the Service Clusters.

Additional HCX Migration/Extension appliances will be deployed for Production traffic when the Prod Service Mesh is created:



- Alternatively, if in this example all Prod and Non-Prod have common cluster networks, the production and non-production Service Clusters can be added later to a single compute profile:



Considerations for Selecting the Deployment Resources in a Compute Profile

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

- The Deployment Clusters, Resource Pools and datastores foremost define where the HCX Migration, Optimization and Extension virtual appliances will be created when the Service Mesh is created.
- The Deployment Cluster must exist in the registered vCenter Server containing the service cluster.

This means a Deployment Cluster can be different than the Service Clusters, but it cannot be in a different vCenter Server.
- The Deployment Cluster selected for the Compute Profile will be used for all appliance types in the service mesh.
- The Deployment Cluster should be able to satisfy the system requirements for the HCX Service virtual appliances. Refer to the User Guide.
- A Resource Pool can be selected as the compute container for deployments (instead of a cluster).
- The Deployment Cluster selected will determine which Deployment Datastores can be selected.

Considerations for Creating HCX Network Profiles

An HCX Network Profile (NP) is a subcomponent of a full Compute Profile.

- A Network Profile configuration is a reference to an existing network, and a range of IP addresses that are allocated for use by HCX service appliances. HCX will not create the networks referenced during this step.
- When configured using the Network Profiles interface, the network profile configuration does specify the HCX function assigned to the network profile. It is purely a network resource allocated for HCX.

A complete network profile configuration includes the friendly name, a portgroup/logical switch name, MTU, IP ranges, Gateway and Prefix Length. Optionally a DNS configuration.

Interconnect
 HCX Components Multi-Site Service Mesh (New)
 Compute Profiles Site Pairing Service Mesh Network Profiles
 CREATE NETWORK PROFILE
 fleet2000 (wdc-pod3-cialab2-stage-10-159-136-48.eng.vmware.com)

Network	MTU	IP Pools			
		IP Ranges	IP Usage(Used/Total)	Prefix Length	Gateway
fleet2000	1500	192.168.4.201 192.168.4.202	1/1 1/1	24	192.168.4.253

EDIT DELETE

- During Compute Profile creation, you will configure how HCX service appliances will perform the following HCX functions:
 - HCX Management (Connections to vCenter/ESX/HCX Mgr/DNS)
 - HCX vMotion (Connections to vMotion vmkernel network)
 - HCX Bulk/Replication (Connections to replication vmkernel)
 - HCX Uplink (Connections to/from the peer IX/ENT appliance)
- The functions above are assigned to a Network Profile during the CP creation:
 - Multiple functions can be assigned to a Network Profile, or dedicated Network Profiles can be created.
 - Best Practice approach is to assign dedicate network profile per HCX function.
 - When multiple HCX functions are assigned using single network profile (e.g. Uplink and Management using the same Network Profile), HCX IX/EXT will use only one vNic with one 1 IP address.

- Jumbo MTU can be configured when the underlying network infrastructure supports it end to end.
- The IP ranges in the Network Profiles should include only IP addresses reserved for HCX (the Manager IP address should not be included in the ranges).
- During HCX service deployments, a single IP in each Network Profile is consumed. It is a good practice to create equal size Network Profiles, as a service deployment will fail when any one of the Network Profiles is fully exhausted.
- Network Profiles can be utilized with multiple Compute Clusters in multiple service mesh configurations they should be sized to accommodate the potential scale.
- The Network Profiles can be expanded while in-use. Service Mesh may need to be resynchronized.
- The DNS setting is only required on the Network Profile that will be used for Management function.
- The Gateway IP is not required in the Network Profile created for non-routed vMotion or Replication networks.
- The Uplink Network Profile is shared by the migration and extension services.

Considerations for Creating an HCX Service Mesh

- Creating a **Service Mesh** configuration requires the source and destination HCX sites to have valid **Compute Profiles**.
- A **Service Mesh** configuration uses only one CP at the source HCX site and one CP at the destination HCX site.
- When the **Service Mesh** is created, HCX interconnect, WAN Optimization and Network Extension appliances are deployed (if selected).
- Service Mesh deployments are flexible:
 - A single **Service Mesh** deployment can be configured to enable every cluster in the source and destination (all clusters added to a single CP).
 - Alternatively, A CP can be created per cluster, or for a subset of clusters and a **Service Mesh** can be created per src/dst CP pair.
- An existing **Service Mesh** between two HCX sites cannot be used with a newly paired site pair.
 - A **Service Mesh** for Source Site A and Destination Site B cannot be used with Destination Site C. A new A-to-C **Service Mesh** needs to be created.
- HCX Services are selected during the **Service Mesh** configuration.
 - Only HCX services that are enabled in the selected CPs will be selectable. HCX services that are not selected in the underlying CP will be grayed out.

Considerations for HCX WAN Optimization Service Deployments

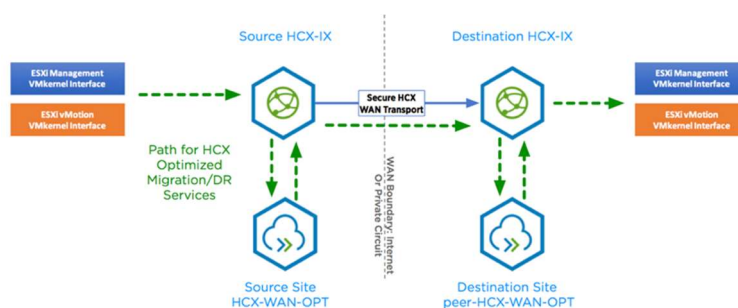
- Enabling the WAN Optimization service is recommended on deployments that are not able to dedicate 10Gbit paths for migrations. Using the HCX Optimization service improves the characteristics of the transport path that HCX services will operate across.
- In 10Gbit, low latency deployments, using WAN Optimization may not yield improved migration performance.
- The HCX-WAN-OPT is always deployed in conjunction with the HCX-WAN-IX. It cannot be deployed to perform stand-alone operations.
- The WAN Optimization service is always deployed in conjunction with the HCX-WAN-IX. It cannot be deployed to service the HCX Network Extension appliances.
- This component uses a simplified deployed model that leverages the compute, storage and network options selected for the HCX-WAN-IX component.
- WAN Optimization Storage

The HCX-WAN-IX appliance is deployed in a fixed form factor that requires 8vCPU, 14GB Memory and two disks: a 30GB OS disk and a 70GB Deduplication backing disk. These resource values are not configurable.


When deploying WAN Optimization service, select SSD Datastores capable of 2500 IOPS when configuring the HCX-WAN-IX.

- WAN Optimization Networking

The WAN Optimization component performs its function in-line (in the path for HCX-WAN-IX flows). Once enabled the HCX-WAN-IX will use internal policy routing to send HCX service packets to its adjacent WAN-OPT component before egressing to the peer site.



- The HCX-WAN-OPT automatically uses the Management Network segment configured during the HCX-WAN-IX deployment. It is auto-assigned special use HCX internal RFC-5735 addresses for communications with the HCX-WAN-IX.



- WAN Optimization Appliance Resource Requirements

CPU	Memory (GB)	Storage IOPS requirements	Storage Throughput Requirements (Mbps)	Guaranteed WAN Throughput Capability (Mbps)
8 vCPU	14	1000-2500	250	1000

Considerations for Deploying the HCX Network Extension Service

- The Network Extension appliance is deployed in pairs. HCX will always deploy an Initiator and a Receiver.
- The HCX-NET-EXT appliance should be deployed per vSphere Distributed Switch containing VM networks that need to be extended.
- Multiple HCX-NET-EXT appliances can be deployed per vSphere Distributed Switch. The deployment can be scaled out to one appliance per VLAN.
- A single HCX-NET-EXT appliance can be used to extend up to 200 networks.
- Additional HCX-NET-EXT appliances can be deployed upfront, or on demand.
- The ESXi hypervisor hosting the HCX-NET-EXT appliance must be connected to the distributed switches of the Virtual Machine networks being extended.
- The HCX-NET-EXT should be configured to use the default 1500 MTU in most deployments. The default MTU setting should not be changed when the internet is traversed to reach the peer HCX-NET-EXT appliance.

- Jumbo MTU can be configured to improve performance when the Jumbo MTU is supported end to end between the source and destination HCX sites.
- The HCX-NET-EXT initiator at the source site initiates UDP-500/UDP-4500 connections to the remote HCX-NET-EXT appliance. Outbound connections should be allowed at the source site's perimeter firewall.

The HCX-NET-EXT receiver at the destination site receives UDP-500/UDP-4500 connections from the source HCX-NET-EXT appliance. Inbound connections should be allowed at the destination site perimeter firewall.

- Use HCX to extend virtual machine networks (VLAN Port Groups and VXLANs/GENEVE Logical Switches).
 - Never extend same VLAN to the same target more than once.
 - Never use HCX to extend the vSphere Management network or other vmkernel/cluster networks (e.g. vmotion/vsan/replication) to the remote site.
 - Never extend the networks used for HCX appliance interfaces.

Considerations for Maximum Transmission Unit (MTU) with HCX

- Jumbo MTU settings can be applied per HCX interface using the Network Profile Configuration.
- Performance may be improved when fragmentation is reduced by the correct use of Jumbo MTU.
- Assign jumbo MTU when the increased MTU size is valid end to end (from the HCX initiator to the HCX receiver appliances).
- **MTU and Network Extension**

The HCX Network extension pipeline adds 150 byte header, when all interfaces in the path support jumbo MTU you can configure MTU to prevent fragmentation during L2 forwarding. Valid Jumbo MTU scenario for HCX:

- VMs in Stretched Network with MTU of 1500.
- Source Site L2 Network Extension appliance Uplink MTU is 1650 or higher.
- Source Site DVS has MTU of 1650 or higher.
- Site to Site interlink has MTU of 1650 or higher
- Destination Site DVS typically has MTU of 9000 (must be 1650 or higher).
- Destination Site L2 Network Extension appliance Uplink MTU of 1650 or higher.

- Target Site VMs use MTU of 1500.
- **MTU and HCX Migrations**

If the vMotion/Replication network is configured for jumbo MTU, it will be beneficial to ensure the HCX vMotion/Replication interface is also configured with the same MTU.

To achieve the best performance result – allocate a vMotion/Replication IP addresses for HCX directly in the cluster vMotion/Replication network using a dedicated vMotion or Replication Network Profile, and configure the MTU to match the cluster setting.

 - In simplified one arm HCX deployments where the HCX Management/Uplink and vMotion are configured using a single network, it may not be possible to achieve the best performance. The best practice is to use dedicated networks mapped to the existing networks.
 - The ESXi cluster may not have a dedicated Replication vmkernel network. It is a good practice but less common than having a dedicated vMotion network. If there is no dedicated Replication network, the management network is used for Replication traffic. In this case it may or may not be possible to use jumbo MTU.

Confirm all environment configurations end to end.

Considerations for Deploying HCX for OS-Assisted Migrations (OSAM)

- This is a new HCX migration option with the HCX service that will enable the migration of non-vSphere virtual machines to an HCX enabled vSphere environment.
 - With the initial release, the service will be certified for KVM to vSphere migrations.
- This service will only be available through a premium SKU called “HCX Enterprise”. Replication Assisted vMotion (RAV), OSAM, and SRM Integration will also be available through the premium SKU.
- **About OSAM:** Conceptually OSAM is similar to Bulk Migration, the source virtual machine remains online during replication.

The source VM is quiesced for a final sync before the migration. HCX performs a software stack adaptation (fixup).

The source VM is powered off, and then the migrated VM is powered on at the target site, for a low downtime switchover.

VMware tools are installed on the migrated VM.

- OSAM Migrations are only supported for certified Operating Systems. In the initial release, the following operating systems are supported for OSAM:

About the Author

Gabe Rosas is a Technical Product Manager with the VMware HCX team. He has expertise in designing and operating traditional and virtualized network Infrastructure, data center virtualization, public cloud environments, and has a passion for driving successful production designs.

