



VMware Cloud: Introduction to HCX

VMware Cloud Migration

Table of contents

VMware Cloud: Introduction to HCX	3
What is HCX?	3
When to Use HCX	4
Data Center Extension	4
Data Center Replacement	4
Disaster Recovery	4
The Technology	5
Components	5
Multi-Site Service Mesh	5
An Overview of HCX Migration	6
Migration Types	8

VMware Cloud: Introduction to HCX

What is HCX?

Although VMware Hybrid Cloud Extension (HCX) provides a number of technical features, its core function lies within its ability to migrate workloads between vSphere environments transparently. While the concept of “workload migration” as a core function seems rather simplistic, the services provided by HCX enable a number of business cases that are otherwise highly impractical.

Specifically, HCX addresses the following technical problems:

- Migration from older versions of vSphere to modern versions of vSphere
- Migrations from non-VMware environments to vSphere
- Migration without the need for IP address changes
- Migration which utilize WAN optimization and data de-duplication
- Data security in the form of [NSA Suite B](#) cryptography for migration traffic

The following demo covers a complete end-to-end demo of workload migration with VMware HCX.

When to Use HCX

The decision of when and if to use HCX should be driven by some “trigger” event in which a specific business need may best be delivered by the functionality which HCX is designed to provide. These triggers are typically broken into 3 categories.

Data Center Extension

Data Center Extension business cases are driven by the desire to quickly expand the capacity of an existing data center; either on temporary or permanent basis. Triggers may include:

- Infrastructure update - Driven by the need to perform a refresh of the hardware or software in the existing production environment. Infrastructure updates are commonly driven by organizational changes which prompt “modernization” projects.
- Capacity expansion - Driven by the need to permanently expand the capacity of a production data center. In this case, HCX provides a means of migrating workloads to and from the SDDC as needed.
- Seasonal bursting - Driven by seasonal expansion. In this case, an SDDC is used to temporarily expand the capacity of a production data center in a transparent manner.

Data Center Replacement

Data Center Replacement business cases are driven by the need to evacuate an existing production data center. In these cases, workloads are permanently migrated to one or more SDDCs. Triggers may include:

- Contract renewal - Driven by an expiring contract on a data center.
- Consolidation - Driven by the desire to consolidate 1 or more data centers into a single SDDC (or multiple SDDCs spread across availability zones or regions). Mergers and acquisitions are a common driver for consolidations.
- Vendor replacement - Driven by the desire to migrate a production data center away from an existing service provider.

Disaster Recovery

Disaster Recovery business cases are driven by the need to use one or more SDDCs as a disaster recovery site. Triggers may include:

- Compliance - Driven by the need to address disaster recovery in order to meet compliance requirements.
- Availability - Driven by the desire to provide increased availability to an existing data center or another SDDC.
- Replacement - Similar to the Data Center Replacement case, but specific to replacing a DR site.

The Technology

For detailed product and documentation, go to the following links:

- [VMware HCX](#)
- [Product Documentation](#)

Components

HCX requires that several appliances be installed within the SDDC and the enterprise environment. These appliances are always deployed in pairs, with one component on the enterprise side and a twin within the SDDC. Installation is driven from the enterprise environment and results in appliances being deployed within the enterprise and the cloud.

HCX uses the following appliances:

- **Manager** - This component provides management functionality to HCX. Within the SDDC, this component is installed automatically as soon as HCX is activated. A download link will be provided for the enterprise HCX manager appliance from within the cloud-side manager. This appliance will be manually installed and used to deploy all other components and drive the operations of HCX.
- **WAN Interconnect Appliance (IX)** - This component facilitates workload replication and migration. The appliance will establish its dedicated IPsec tunnels for communication with its peer within the SDDC.
- **WAN Optimization Appliance (WAN-Opt)** - This component provides WAN optimization and data de-duplication for the WAN Interconnect Appliance. It communicates exclusively with the WAN Interconnect Appliance via private IP (uses addresses from the IPv4 range reserved for carrier-grade NAT).
- **Network Extension (L2C)** - This component provides [Layer-2](#) network extension for “extending” networks between vSphere environments. It will establish its dedicated IPsec tunnels for communication with its peer within the SDDC.
- **Proxy Host** - This is a 'fake' ESXi host deployed silently by the WAN Interconnect appliance. This host is used as the target for vMotion/migrations and is used to “trick” vCenter into thinking the migration target host is local. This host will be visible in the inventory of vCenter.

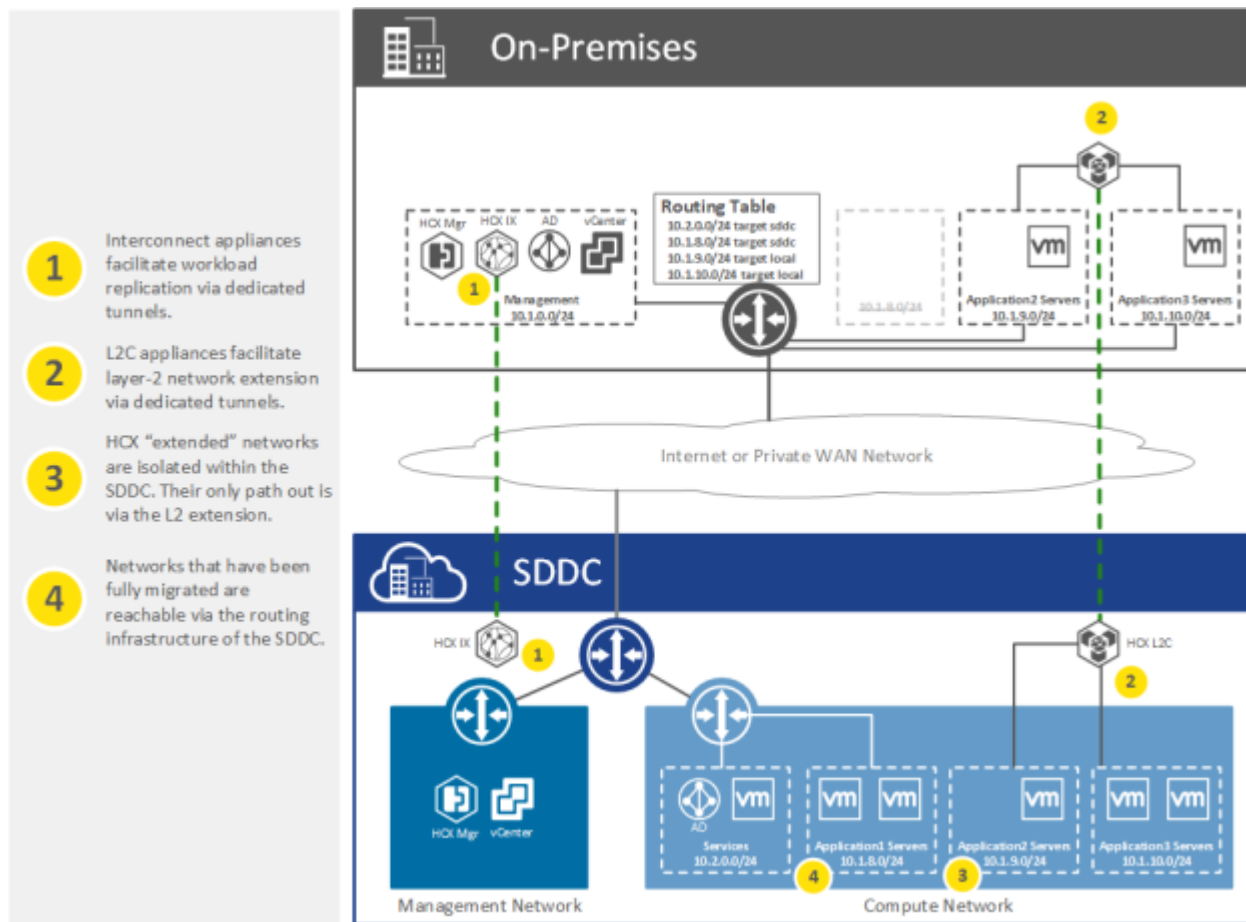
Multi-Site Service Mesh

The Multi-Site Service Mesh feature significantly changes how HCX is deployed and managed. In the previous model, HCX used an appliance-based view of the world where appliances were deployed based on the needed HCX function. In the new model, HCX takes a service-based view where the user picks the needed services, and HCX deploys the appliances required to deliver those services.

All planned new functionality with HCX will be based on Multi-Site Service Mesh, the recommended deployment mechanism for all new installations. If you are a current user of HCX and wish to utilize the service mesh, then you may upgrade your deployment to use the Multi-Site Service Mesh. Note that you must not have any migrations in progress or networks extended to upgrade to Multi-Site Service Mesh.

An Overview of HCX Migration

A typical example of a workload migration project is illustrated below. It assumes that proper wave planning and testing have already been performed.



In this example, HCX is used to migrate application workloads and their associated networks to the SDDC. The WAN Interconnect appliance (IX) is responsible for data replication between the enterprise environment and the SDDC. This replication traffic is carried over a dedicated IPsec tunnel initiated by the enterprise appliance and optimized using the WAN Optimization appliance (not shown). The Network Extension appliance (L2C) is responsible for creating a forwarding path for networks extended to the SDDC. Again, this traffic is carried over a dedicated IPsec tunnel initiated by the enterprise appliance.

In this example, the SDDC contains a "Services" network with critical services such as DNS and Active Directory designed to serve the local environment. In general, it is a good idea to keep such services as close to the workloads as possible, as doing so will not only cut down on the amount of north-south traffic but will also reduce dependencies between sites.

We can see from the diagram that a migration is in progress. In this scenario, we are migrating all the application servers in the networks 10.1.8.0/24, 10.1.9.0/24, and 10.1.10.0/24. The network 10.1.8.0/24 has already been completely migrated, and, as a result, the L2 extension has been disconnected, and the network has been shut down in the on-premises network. This network is now accessible directly through the routing infrastructure of the SDDC. This process of making the SDDC the authority for a migrated network is often referred to as a "network cutover".

Workloads from within the networks 10.1.9.0/24 and 10.1.10.0/24 are still being migrated. As indicated, the L2C is attached to these networks to provide a layer-2 network extension to the SDDC. Due to this network extension, the workloads that have already been migrated have retained their network addresses and, from a routing perspective, appear to reside within the enterprise environment. These "extended" networks are not tied into the routing infrastructure of the SDDC.

Extended networks present an interesting routing scenario. Because they are not tied to the routing infrastructure of the SDDC, the only "way out" for the workloads is via the layer-2 network extension to the enterprise default gateway. This means that all "non-local" traffic to the extended network must pass through the L2C and be routed through the enterprise gateway router. This includes not only communications to resources within the enterprise environment but also communications between other extended networks, as well as to networks that are native to the SDDC or resources within the cross-linked VPC.

This process of forwarding traffic from SDDC -> enterprise -> SDDC between extended networks is referred to as “tromboning” and can result in unexpected WAN utilization and added latency. Since migrations tend to be a temporary activity, this tromboning effect is not typically a significant concern. However, it is essential to keep this in mind when planning migrations such that it may be reduced as much as possible.

One additional note on the diagram above; although the tunnels for the IX and L2C appear to be bypassing the SDDC edge, in reality, this edge is in-path and acts as the gateway device for these appliances. This is an important fact to consider when planning an HCX migration, as bandwidth from replication and the layer-2 extension will contribute to the overall load of the SDDC edge router.

Migration Types

Before planning migrations, it is essential to keep in mind the migration technique that is the most optimal for a given workload or group of workloads. **HCX Cold Migration**

This technique is used for VMs which are in a powered-off state. Workloads that may typically be cold migrated are testing and development VMs, or templates. A very small percentage of workloads tend to use this migration technique.

HCX Bulk Migration

The vast majority of workloads will use this technique. Bulk migration is also commonly referred to as “reboot to the cloud” since workloads are first replicated and at a pre-defined time are “hot swapped”, meaning that the original workload is powered off and archived simultaneously to its replica in the cloud is powered on.

Bulk migration utilizes a scheduler to execute the actual migration. The process for scheduling a migration is as follows:

1. Identify a migration wave and schedule the migration.
2. HCX begins the initial data replication between sites via the IX appliance and utilizes the WAN-Opt appliance for WAN optimization and data de-duplication.
3. Once the initial data replication is completed, HCX will continue replicating changes periodically.
4. When the scheduled migration time arrives, HCX will begin the cutover process.

The cutover process for bulk migration is as follows:

1. Power down the source workload. 1 - 10 minutes, depending on VM shutdown time.
2. Perform final data sync. 1 - 10 minutes.
3. Optional pre-boot fixups (e.g., hardware version upgrade). 1 - 2 minutes, depending on the vCenter load.
4. Boot on target. 10 - 20 minutes, depending on the vCenter load.
5. Optional post-boot fixups (e.g., vmttools upgrade). 1 - 10 minutes.
6. User validation of workload. 0 - n minutes depending on user requirements for validation.

HCX vMotion

This technique is reserved for workloads that cannot tolerate a reboot. Given that additional state must be replicated at the time of the migration, it is more resource intensive than a bulk migration. Typically, only a tiny percentage of workloads will require vMotion.

There are two options for vMotion: standard vMotion and Replication Assisted vMotion (RAV). With standard vMotion, you are limited to a single migration at a time. Initially, this was the only option for vMotion. With the introduction of RAV, vMotion was given a few of the features of bulk migration; specifically, the ability to migrate multiple workloads in parallel, the ability to utilize the WAN-Opt appliance for optimization and de-duplication, and the ability to utilize the scheduler.

The following restrictions for vMotion have been summarized from the HCX user guide. Please refer to that document for details.

- VMs must be running hardware version vmx-09 or higher.
- The underlying architecture, regardless of OS, must be x86 on Intel CPUs.
- VM disk size cannot exceed 2 TB.
- VMs cannot have shared VMDK files.
- VMs cannot have any attached virtual media or ISOs.
- VMs cannot have change block tracking enabled.

Note: Only 1 vMotion operation at a time can be running, and it is not recommended to perform migrations using vMotion while bulk migration replication is also in progress. RAV does not have these restrictions.

