

Network Underlay Characterization and HCX Performance Outcomes



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

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

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

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

Contents

About This Document	4
1 Network Underlays and HCX	5
2 Characterizing a Network Underlay for HCX	9
Network Underlay Bandwidth	10
Network Underlay Loss and Latency	13
Network Underlay Maximum Transmission Unit (MTU)	14
3 Network Underlay Characterization Reference	21
Network Underlay Reference for 50 Mbps Bandwidth	21
Network Underlay Reference for 100 Mbps Bandwidth	23
Network Underlay Reference for 150 Mbps Bandwidth	26
Network Underlay Reference for 200 Mbps Bandwidth	28
Network Underlay Reference for 500 Mbps Bandwidth	31
Network Underlay Reference for 1 Gbps (1000Mbps) Bandwidth	34
Network Underlay Reference for 2 Gbps (2000Mbps) Bandwidth	37
Network Underlay Reference Beyond 2 Gbps Bandwidth	40
4 Performance Factors Beyond the Network Underlay	42
5 About the Author	44

About This Document

The *HCX Network Underlay Characterization and Performance Outcomes* technical paper provides information to help HCX users understand the relationships between the network underlay and VMware® HCX.

With HCX performance, various dimensions of environmental and load data need to be considered. One of the dimensions is the network underlay and the HCX performance derived from the underlay capabilities. In this regard, **Characterizing an Underlay for HCX** means understanding whether the underlay meets the minimum HCX requirements for providing successful virtual machine migrations and network extension services, and understanding baseline performance outcomes for given underlay conditions.

This document attempts to put these considerations in perspective and also tries to provide some guidance on how to verify whether the performance is optimal for the given environment and parameters.

Intended Audience

This information is intended for users who want to better understand HCX performance. To provide considerations and guidance for the elements influence performance outcomes. The information is written for systems administrators who are familiar with virtual machine technology and datacenter operations.

VMware Technical Publications Glossary

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

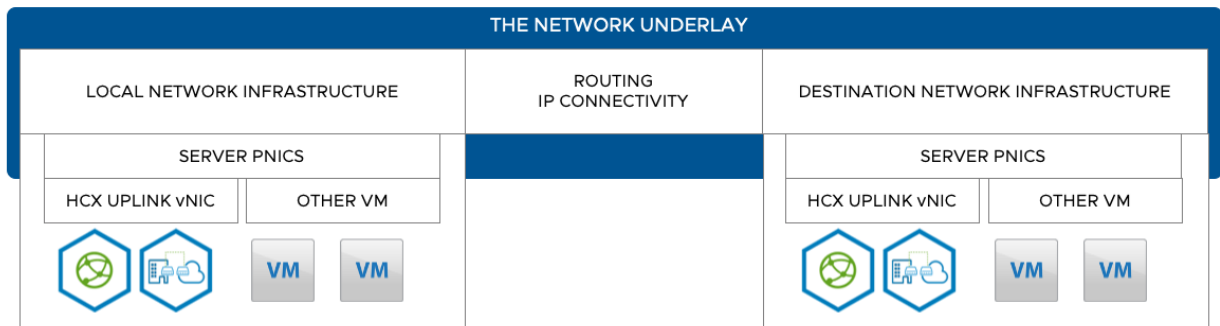
Network Underlays and HCX

1

The HCX data plane uses virtual appliances to connect private and public vSphere environments using tunneling technology over existing network infrastructure. The relationship between the HCX Service Mesh and the network underlays are described in this document.

What is the Network Underlay

An underlay is a network that provides transport of traffic between HCX appliances at the source and destination sites. This includes the physical network infrastructure in the customer data center, any physical network services connecting customer locations(i.e. SD-WAN, IPsec VPN), and (if applicable) the cloud provider virtual infrastructure . The underlay network does not generally have awareness that it is carrying HCX packets.

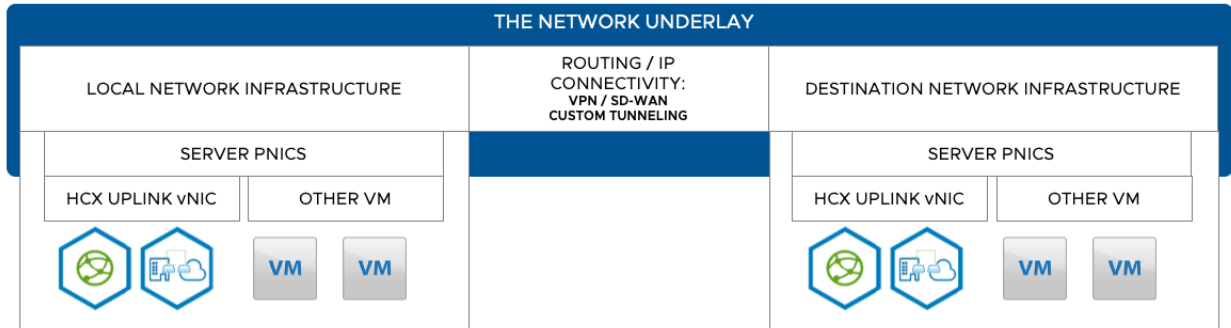


A network underlay can vary from purpose built, lossless, high bandwidth, low latency networks connecting server racks in a data center, to lower bandwidth, variable latency, Internet-based connectivity. In this document, the term network underlay encompasses all elements that affect the performance characteristics of an underlay, including the servers, network devices and connections between the vSphere/HCX environments. **The network underlay requirements must be satisfied across all elements of the underlay.**

VPN based Network Underlays

IPSec Virtual Private Networks (VPNs) are frequently used for creating secure data center connections to private and public VMware clouds over the Internet. SD-WAN and custom tunneling solutions are used over the internet to provide redundancy and improve throughput. SD-WAN, IPSec VPN, and other tunneling solutions are collectively referred to as VPN in this document.


The term network underlay includes connections with VPN configurations. The network underlay requirements must be satisfied when considering all elements of the underlay, including any performance overhead associated with VPN configurations in the path.



HCX Services on the Underlay

An HCX Service Mesh is the instantiation of HCX services across HCX enabled sites. When a service is enabled, appliances are deployed in both environments simultaneously, and service tunnels are automatically established (see the figure below). The migration (IX) appliance and network extension (NE) data plane appliances (and specifically their site to site transport tunnels) are in-scope when discussing HCX on a network underlay throughout this document.

HCX appliances have at least one uplink, and support up to three uplinks. Each uplink can be connected to a unique Port Group, and may utilize a distinct network underlay. Every network underlay that provides transport for HCX traffic should be evaluated when planning for an HCX deployment.



Network Underlay Minimums & HCX Performance Maximums

For a fully supported deployment, all of the minimum network underlay requirements must be satisfied by every network underlay, irrespective of their expected performance outcomes. The minimum requirements are detailed in the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

The HCX maximums below refer to the known limitations of the HCX services. For additional HCX limits, visit configmax.vmware.com.

Caution The HCX performance maximum values below are known HCX product limitations based on controlled testing. They are not indicators of what performance should be expected in all HCX deployments and with all network underlays. For example:

- 1.5 - 2.0 Gbps of migration throughput is not possible on a 500 Mbps network underlay, regardless of any other environmental factors.
- 1.5 Gbps of migration throughput may not be possible on a 2 Gbps network underlay depending on the underlay's latency or packet loss characteristics or if there are other environmental constraints (like poor storage performance) creating a performance choke point.

Migration Maximums
<ul style="list-style-type: none"> ■ 1.5-2.0 Gbps IX appliance maximum. <ul style="list-style-type: none"> ■ Total migration traffic across all concurrent vMotion and Replication based migrations and protections on one IX appliance. ■ Multiple Service Meshes are required to utilize bandwidth beyond the limits of a single appliance. ■ A service mesh can be deployed for each source/destination cluster pair.
<ul style="list-style-type: none"> ■ 400 Mbps (180GB/hr) maximum per VMDK per Replication. <ul style="list-style-type: none"> ■ Applies to HCX Bulk, DR Protection and RAV (initial/delta sync phase). ■ A single VMDK virtual machine replication is subject to this limit. ■ 1 Gbps (450GB/hr) maximum per migration operation.
<ul style="list-style-type: none"> ■ 900Mbps - 1.25Gbps per Relocation <ul style="list-style-type: none"> ■ Applies to Cold migration, HCX vMotion and the RAV switchover. ■ Only one relocation operation per IX appliance. ■ Additional IX appliances must run on different ESXi hosts to increase the maximum relocation concurrency. Additional IX appliances running on the same host will not be able to start a relocation.
<ul style="list-style-type: none"> ■ 1 Gbps WAN Optimization appliance maximum. <ul style="list-style-type: none"> ■ The HCX WAN Optimization service does not increase the IX maximums ■ A WAN Opt increases efficiency by providing data reduction when deployed with network underlays with bandwidth below 1 Gbps.

Network Extension Maximums

- 4-6 Gbps Network Extension maximum throughput per appliance.
 - By default the NE appliance shares a network underlay with the IX.
 - Multiple IX appliances can be deployed per Service Mesh.
- 850Mbps - 1Gbps Network Extension maximum throughput per Network Flow.

Characterizing a Network Underlay for HCX

2

"Characterizing an underlay" refers to measuring the network underlays and understanding what the performance is possible from the observed underlay parameters. A network underlay characterization effort is necessary for understanding what is supported and what performance should be expected with HCX on that underlay.

What is Characterization?

Characterizing an underlay means measuring the network underlays and ultimately understanding what performance is possible from the given underlay parameters. A characterization effort improves the success of HCX migration projects by ensuring requirements are properly identified and satisfied. This effort also helps identify risks that could prevent users from successfully using HCX for migrations.

Network Underlay Unsupported Characteristics

The table below describes the unsupported network underlay characteristics. In the process of measuring and characterizing the various aspects of the underlay, look out for measurements that cross the thresholds in this table. For details about the minimum requirements, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Network Underlay Parameters	HCX vMotion	Replication Assisted vMotion	HCX Bulk Migration	OS Assisted Migration
Bandwidth < 50 Mbps	NO SUPPORT	NO SUPPORT	NO SUPPORT	NO SUPPORT
Packet Loss 0.1 - 1.0 %	NO SUPPORT (< 0.1% OK)	NO SUPPORT (< 0.1% OK)	OK	OK
Packet Loss > 1.0%	NO SUPPORT	NO SUPPORT	NO SUPPORT	NO SUPPORT

Network Underlay Parameters	HCX vMotion	Replication Assisted vMotion	HCX Bulk Migration	OS Assisted Migration
Latency > 150ms	NO SUPPORT	NO SUPPORT	NO SUPPORT	NO SUPPORT
MTU < 1150	NO SUPPORT	NO SUPPORT	NO SUPPORT	NO SUPPORT

Caution In HCX deployments where the underlay includes a VPN tunnel, or the MTU is below 1350bytes, HCX version 4.2 or later is required. For more information, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

This chapter includes the following topics:

- [Network Underlay Bandwidth](#)
- [Network Underlay Loss and Latency](#)
- [Network Underlay Maximum Transmission Unit \(MTU\)](#)

Network Underlay Bandwidth

The bandwidth of a network is usually defined as the maximum quantifiable capacity of the network to transfer data in a given amount of time. This section provides instructions for measuring bandwidth on a network underlay and implications of this characteristic in for HCX installations by measuring the maximum throughput possible on that underlay.

About Network Underlay Bandwidth

This parameter refers to the network bandwidth available to across the migration and network extension appliances across the service mesh. The following considerations apply:


- The method used to measure the network bandwidth on the underlay is a measurement of networking infrastructure and should never be correlated to represent maximum migration throughput.
- When calculating the bandwidth requirements for parallel migrations, the disk churn for all virtual machines that will be migrated in parallel should be considered.
 - Disk churn refers to disk space changes, or disk usage for a virtual machine.
 - The reference tables at the end of this document provide "Maximum Delta Transfer Supported". This refers to the disk churn in GB.
- It's typical for the HCX Network Extension service to use the same network underlay bandwidth as the migration data traffic. This can be a very significant factor to how much bandwidth is available for migrations.
- Bandwidth consumed by non-HCX related workloads on the same underlay should be understood when evaluating the bandwidth available for HCX operations.

Understanding Network Bandwidth Contention

Network contention occurs when two or more hosts on a network compete for bandwidth. Consider a typical home network as an example. A teenager at home playing an online game is competing for network bandwidth with a sibling that is streaming video. If the internet connection is not fast enough to transmit both gaming and video streaming traffic, one of the two application's performance will suffer. The same principle applies to enterprise network at a greater scale. This is an important concept to understand, with the goal of eliminating network contention as much as possible for HCX traffic. Eliminating or reducing contention can be achieved by one or more of these approaches:

- Increase link speeds to provide more bandwidth (e.g. upgrade from 1 Gbps links to 10 Gbps links).
- Implement dedicated network equipment and circuits to a specific application or set of applications.
- Implement Quality of Service (QoS) controls to prioritize specific traffic. Properly implementing QoS for HCX traffic is beyond the scope of this document.

In general the bandwidth distribution with HCX on the network underlay is represented, by the following figure:



Measuring the Network Underlay's Bandwidth

Bandwidth on an underlay can be measured using "perftest uplink" from the HCX Central CLI, or by using equivalent tools on virtual machines connected similarly to the HCX Service Mesh appliances. In this document, the examples are provided using CCLI tools on the HCX appliances.

In the perftest example, the source HCX IX appliance (the initiator) has an uplink IP of 198.51.100.3 and destination IX appliance (the receiver) has an uplink IP of 198.51.200.3.

Use a terminal or an SSH client to connect to the HCX Manager using SSH, and invoke the CCLI:

```
[admin@hcxmgr ~]$ ccli
Welcome to HCX Central CLI
```

```

[admin@hcxmgr] list
|-----|
| Id | Node | Address | State | Selected |
|-----|
| 0 | Interconnect-ConfigIXTest-IX-I1 | 198.51.100.3:9443 | Connected |
|-----|
[admin@hcxmgr] go 0
Switched to node 0.
[admin@hcxmgr:Interconnect-ConfigIXTest-IX-I1] perfctest uplink
Testing uplink reachability...
Uplink-0 round trip time:
PING 198.51.200.3 (198.51.200.3) from 198.51.100.3 : 56(84) bytes of data.
64 bytes from 198.51.200.3: icmp_seq=1 ttl=64 time=0.502 ms
64 bytes from 198.51.200.3: icmp_seq=2 ttl=64 time=0.293 ms
64 bytes from 198.51.200.3: icmp_seq=3 ttl=64 time=0.684 ms
64 bytes from 198.51.200.3: icmp_seq=4 ttl=64 time=0.618 ms
64 bytes from 198.51.200.3: icmp_seq=5 ttl=64 time=0.766 ms

--- 198.51.200.3 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4000msrtt min/avg/max/mdev =
0.293/0.572/0.766/0.166 ms

Uplink native throughput test is initiated from LOCAL site.
+++++++ StartTest ++++++
----- Uplink-0 [te_0, 198.51.100.3 >>> 198.51.200.3] -----
Duration Transfer Bandwidth Retransmit
server workload started
[ 4] 0.00-60.00 sec 37.4 GBytes 5.35 Gbits/sec 6469 sender
[ 4] 0.00-60.00 sec 37.4 GBytes 5.35 Gbits/sec receiver
Done
----- Uplink-0 [te_0, 198.51.100.3 <<< 198.51.200.3] -----Duration Transfer
Bandwidth Retransmit
server workload started
[ 4] 0.00-60.00 sec 40.4 GBytes 5.78 Gbits/sec 666 sender
[ 4] 0.00-60.00 sec 40.4 GBytes 5.78 Gbits/sec receiver
Done
+++++++ StopTest ++++++

```

Caution The **CCLI Perfctest Uplink** command is used to provide a measurement of the network underlay from HCX Uplink to HCX Uplink, without factoring overhead from the HCX IPsec tunnels. The perfctest uplink result may provide a good indication of what the network underlay is able to handle.

The **CCLI Perfctest IPsec** command is used to provide a measurement for a single flow through per IPsec tunnel. The migration appliance (IX) and network extension (NE) appliances use multiple IPsec tunnels to forward HCX traffic.

The **CCLI Perfctest Site** command performs a service level test, for multiple flows across all ipsec tunnels, including the WAN Optimization service, if deployed.

The **CCLI Perfctest All** command runs all performance commands..

The example above uses **CCLI Perfctest Uplink**. The output shows a network underlay capable of handling 5.35 Gbps - 5.78 Gbps of bandwidth. The following interpretations can be made:

- Based on bandwidth alone, the network underlay is able to support the known IX appliance maximums.
- Based on bandwidth alone, the network underlay is able to support the 4 - 6Gbps HCX NE appliance forwarding maximum for a single appliance.
- The network underlay may not be able to support both appliances forwarding traffic at the appliance limit.

Additional considerations regarding this example:

- Bandwidth measurements with CCLI Perfctest or equivalent external tools provide a value that is a point in time measurement. It is an estimation of the available bandwidth at the time the test was done. It is always possible that some other traffic can start moving through the same underlay when migrations are started. This can skew expectations based on the perfctest results.
- Available bandwidth for migration on an underlay also changes as the migration wave progresses. This occurs as more and more VMs are distributed across the two sites. Any East-West traffic that was previously confined to a single site boundary may now traverse the connection between the sites. If the underlay network does not reserve bandwidth for migration traffic, any other inter-VM traffic will compete for underlay bandwidth with the migrations.

Network Underlay Loss and Latency

Latency and loss are metrics that can be gathered for an underlay. This section describes how to measure them, and their implications on HCX services.

Network Latency

Here's a way to understand what is meant by latency:

- Consider the time at which a packet is sent by a host. The precise time is recorded as timestamp A.
- Consider the precise time at which the packet is received at its destination as timestamp B.
- The latency is timestamp B minus timestamp A as per the relevant definition from RFC 1242.

Network Underlay Packet Loss

Packet loss refers to the percentage of transmitted packets that are not received at their destination. This happens for a variety of reasons in networks, like when network device is overloaded and cannot accept additional packets at a given moment. In general, packet loss may or may not be disruptive to the recipient of the data, depending on the protocol used and the severity of the loss. Migration performance is always impacted negatively by packet loss.

The VMware vMotion, Host Based and OS Assisted Replications can tolerate packet loss , depending on the specific migration type. Refer to the previous table, or the [Network Underlay Minimum Requirements](#) page for details. Packet loss is directly correlated to migration performance, even when loss is within the tolerable packet loss levels.

Effects of High Network Latency on HCX Traffic

In general, high latency affects HCX throughput.

- HCX depends on technologies like vMotion and vSphere Replication that perform differently at high latencies:
- Replication buffer limits on ESX hosts are affected at higher latencies, greatly reducing the peak performance.
- High latency can cause HCX WAN Optimization service to read data off from the source more aggressively than can be sent to the destination. This can cause vMotion to detect a stall in progress, and in some cases connections can fail. HCX has mechanisms to reduce these cases but high latency can introduce risk of failed vMotion connections.

Measuring Network Underlay Packet Loss and Latency

A common way to check for loss and latency is to measure several pings sent from the HCX appliance. Performing the ping test with the Uplink IPs will test only the underlay.

```
# ping -I <Local_uplink_ip> -i 0.001-c 10000 <Remote_uplink_ip>
```

This causes a ping at 1000 PPS for 10seconds. A report will be printed at the end of the 10 seconds that looks something like this:

```
-----
--- <Remote_uplink_ip> ping statistics ---
10000 packets transmitted, 10000 received, 0% packet loss, time 10066ms rtt min/avg/max/mdev
= 0.068/0.086/1.563/0.037 ms
-----
```

The time it takes for a ping to be sent and the response to be received is referred to as the round-trip time (RTT). Latency calculations should consider the mean deviation (mdev), which is also referred to as jitter. Latency is determined as follows:

$$\text{RTT} = 0.086 + 0.037 = 0.12 \text{ ms}$$

$$\text{One-way latency} = \text{RTT}/2 = 0.06 \text{ ms}$$

Network Underlay Maximum Transmission Unit (MTU)

Path MTU is equal to the minimum of the MTUs of each hop in a network underlay. This section describes the network underlay path MTU, how to measure and configure it in HCX.

MTU & Path Maximum Transmission Unit (MTU)

The term Maximum Transmission Unit (MTU) refers to the maximum size of a layer 3 packet that can be sent over a network. Path MTU (PMTU), and it is equal to the minimum of the MTUs of each hop in an underlay. Ethernet standards originally specified an MTU of 1500 bytes, which also applies to packets traversing the public Internet. 1500 bytes remains a common MTU for many networks, although modern network equipment can typically support an MTU of 9000 bytes.

Maximum Transmission Unit (MTU) is the largest packet that can be transmitted without fragmenting it. Fragmentation occurs when a packet is received that is larger than the MTU, resulting in that packet being split into two or more packets to accommodate the MTU. Fragmented packets are reassembled when they reach their destination. Fragmentation adds overhead and high fragmentation rates have a negative effect on overall throughput.

When one IP host has a large amount of data to send to another host, the data is transmitted as a series of IP packets. It is usually preferable that these packets utilize the largest MTU possible that does not require fragmentation anywhere along the path. Network devices make a forwarding decision for each packet. 1500 byte packets simply transfer more data than smaller packets, resulting in fewer packets sent. This reduces the number of forwarding decisions each network device must perform, reducing overhead and increasing throughput.

If the end-to-end path has multiple hops, implying multiple path segments, then the minimum of all the segment MTUs is called the Path MTU (PMTU) and it is equal to the lowest MTU in the path.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <<https://www.rfc-editor.org/info/rfc1191>>.

Choosing the correct MTU for HCX network profiles requires knowledge of the underlay and the path that packets will traverse. In environments where administrators have control of the path end-to-end, determining the MTU is a matter of finding the smallest MTU in the path and using that value. Otherwise, testing should be performed to determine the correct PMTU. In many cases the default MTU of 1500 will work, but if there are IPsec tunnels or legacy WAN links in the path, a smaller MTU must be used. This will ensure that no HCX migration traffic is fragmented. Failing to set the correct MTU will result in sub-par migration throughput.

Effects of Setting Incorrect MTU

One of the main concerns in HCX deployments is the potential for MTU mismatch across the different segments in the end to end network path. MTU mismatches can manifest in different ways, and proper functioning of HCX depends on how the intermediate devices handle MTU mismatches.

An intermediate device on encountering an MTU mismatch can:

- 1 Fragment the packet to its required MTU size.

Net effect of fragmentation is it increases number of packets to be processed for the same amount of data while reducing overall performance and bandwidth. Fragmentation at higher network speeds can cause excessive packet drops due to limited fragmentation handling capabilities in most intermediate devices like routers.

2 Drop the packet.

If the “Don’t Fragment” (DF) bit is set in the IP header for a packet, it will be dropped instead of fragmented. In some cases (e.g. specific firewall configurations), packets requiring fragmenting may be dropped even if the DF bit is not set.

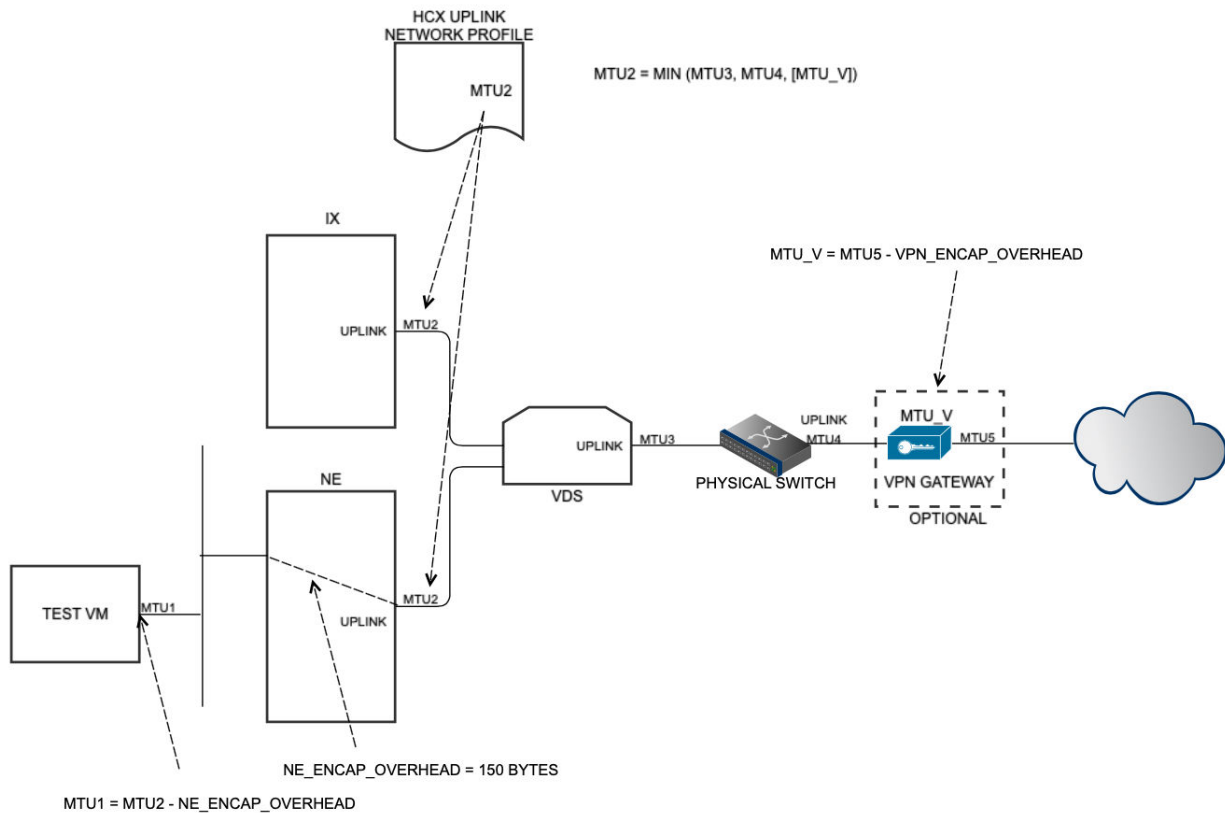
Intermediate devices particularly VPN Gateways can also display a non-conforming behavior on how the DF bit is handled.

- 1 Devices can drop the packet if there is a MTU mismatch such that the packet needs to be fragmented to send to the next hop and the DF bit is set. Conforming devices will send an ICMP message back upstream to signal the reason for the drop.
- 2 Devices can fragment the packet even if DF bit is set (Ignore DF). This is a configurable setting in many of the commercial VPN devices.
- 3 Devices can provide a choice of either ignoring or pulling the DF bit onto the outer IP header that is encapsulating the inner original packet. If DF bit is ignored, Path MTU Discovery (PMTUD) can fail in the remaining path after the VPN terminates.

End to End MTU Example A

Consider the following figure that shows a typical HCX setup with one Uplink Network Profile and two HCX appliances (IX and NE). The figure also shows an optional VPN or tunneling gateway in the underlay that can affect MTU:

Figure 2-1. Example A Diagram



In Example A, each of the interfaces of each of the path devices have an MTU defined at that interface:

- MTU 1 is the MTU used by TEST VM on an HCX extended network.
- MTU2 is the MTU that is provided as a user input in the HCX Uplink Network Profile and is the MTU at the Uplink interface of the HCX IX and NE appliances.
- MTU3 is the Distributed Switch (VDS) Uplink MTU setting
- MTU4 is the physical switch Uplink MTU
- MTU5 is the MTU of the physical interface for the VPN gateway or the Router Endpoint that is providing the VPN.
- MTU_V is the derived MTU of the VPN or the tunnel itself. This is the physical interface MTU on which the VPN is plumbed minus the VPN encapsulation overhead in bytes.

Configuring MTU Based on Path MTU Investigation (Option 1 - The Manual Approach)

Every MTU is discovered through manual verification on known devices in the path.

Following Example A depicted above:

- The vSphere Distributed Switch MTU, Physical Switch/Router, and VPN Gateway MTUs are first learned through a manual discovery of their interface configurations.
- Uplink MTU or MTU2 is then added in the HCX Uplink Network Profile.
 - In Example A, this value should be calculated as the minimum of MTU3, MTU4 and MTU_V. This ensures that packets leaving HCX appliances on the uplink are not fragmented by the VPN Gateway or any other known intermediate device.
- For optimal NE operation, the workload MTU (MTU1) should be set to MTU2 minus the NE's Encapsulation overhead (150 bytes). This ensures virtual machines sending data on stretched networks don't see fragmentation. Managing MTU1 addresses fragmentation cases for both TCP and UDP flows, but requires modification on workload VMs, which is not practical in most cases:
 - The TCP Flow conditioning feature in HCX allows the virtual machine endpoints are free from fragmentation without requiring the MTU to be managed on the workloads.
 - TCP Flow conditioning should always be enabled in deployments with VPN underlays so that TCP data is free of fragmentation without modifying the workloads.

Caution It is still possible to get the MTU wrong if an intermediate device with smaller MTU is not considered. It may not be possible to find out all the intermediate MTUs. This then essentially becomes a best-effort rather than a guarantee. Path MTU Discovery is recommended in the underlay.

Configuring MTU Based on Path MTU Discovery (Option 2 - Rely on PMTUD)

This approach discards the manual verification of MTUs in the path and instead relies on Path MTU Discovery (PMTUD) to set the correct MTU on the underlay. This approach may be used on the basis that it is sometimes difficult to know MTU settings of every segment in a path, so PMTUD is used to automate the discovery of the correct Path MTU.

There are two possible, and separate approaches to PMTUD for the HCX underlay exercise:

- 1 PMTUD initiated by a workload VM on a stretched network. In this case HCX NE is the intermediate device that needs to respond with the correct RFC 4821 behavior by adjusting the encapsulation overheads the appliance introduces before the next hop. HCX does not support this as of the time this document was written.
- 2 PMTUD initiated by HCX on the uplink to determine the smallest MTU on the WAN path.

PMTUD is enabled by default on the HCX appliances. (There is no option to turn it off through configuration). This means the uplink MTU is adjusted accordingly if intermediate routers send ICMP messages back to adjust MTU.

As an example, the below snippet shows that a new MTU of 1438 was discovered while talking to the peer IP 172.17.12.11 over the uplink. The packetization will always consider this MTU (1438) as long as the route cache entry is valid (in the example it says it is valid for another 585 seconds)

```
[root@SM_625f48de_77accd45-IX-I1 ~]# ping 172.17.12.11 -s 1450 -M do -c 1
PING 172.17.12.11 (172.17.12.11) 1450(1478) bytes of data.
From 172.16.12.254 icmp_seq=1 Frag needed and DF set (mtu = 1438)

--- 172.17.12.11 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

[root@SM_625f48de_77accd45-IX-I1 ~]# ip route get to 172.17.12.11
172.17.12.11 via 172.16.64.253 dev vNic_0 src 172.16.64.201
cache expires 585sec mtu 1438
```

Caution Pitfalls of relying purely on Path MTU Discovery:

- 1 PMTUD can break with router/firewall misconfigurations, or when ICMP traffic is blocked. These result in silent failures that can become difficult to debug. The PMTUD behavior is overridable by setting the MTU manually.
 - 2 Hardened routers or firewall policies may prevent these devices from sending ICMP packets back to the source or drop received ICMP packets from downstream routers causing the PMTUD to fail.
 - 3 PMTUD does not solve MTU mismatches within the L2 domain. As an example, if the HCX appliance Uplink MTU and the VDS MTU are different then PMTUD will not catch this difference and packets may be dropped as a result.
 - 4 PMTUD cannot catch end-to-end MTU mismatches. For example, if one appliance has an uplink MTU of 1500 and the peer appliance has an uplink MTU of 9000. Assuming the entire path between these two appliances supports an MTU of 9000, any packet coming from the peer side larger than 1500 would still be dropped because PMTUD could not figure out that one side has an appliance with MTU 1500.
-

Configuring MTU Using a Hybrid Approach (Preferred)

The hybrid approach:

- 1 Manually set the HCX Uplink MTUs by first investigating the MTUs of path devices in the local environment.
- 2 Use the "ip route get to <peer>" command to understand if a new PMTU is discovered.
- 3 If a lower MTU is discovered, change the local uplink MTU value to the discovered MTU value. (Note: It is possible that the two directions can see two different PMTUs.)

- 4 Compare the Uplink MTUs from both the sides and select the minimum of the two. Now change the MTU at the end that had the larger MTU and make it equal to the smaller MTU. This addresses the complete end to end MTU issues.

As an example consider the following path:

Figure 2-2. Example B Diagram



Manually investigating the local environments at the two sites would give the MTU to set on the uplink as 1500.

One site receives PMTU of 1450 while the other site in the other direction receives PMTU of 1435.

With the hybrid approach, the two uplink profiles will be set with 1450 and 1435 respectively at this point.

Now the final step is to take the minimum of these two values i.e. 1435 and set this in the uplink profile on both sides.

MTU with Multiple HCX Uplinks

In HCX deployments with multiple uplinks, the uplink with the smallest MTU is used for all packets. In other words, two uplinks with one 9000 MTU and another with 1500 MTU would end up using the 1500 MTU for all packets.

Non-Uplink MTUs in HCX

HCX appliances connect to other networks for migration and management. The IX appliance connects to networks for vMotion and replication workflows as well as management networks. This document does not provide guidance on setting MTUs for these networks.

Network Underlay Characterization Reference

3

Reference information for HCX performance and supportability outcomes based on sample characterized network underlays.

This section is intended as reference information for what may be expected and what is supported based given a specific network underlay. This information should only be used once the network characterization effort as described in the [Chapter 2 Characterizing a Network Underlay for HCX](#) chapter of this document has been completed.

This chapter includes the following topics:

- [Network Underlay Reference for 50 Mbps Bandwidth](#)
- [Network Underlay Reference for 100 Mbps Bandwidth](#)
- [Network Underlay Reference for 150 Mbps Bandwidth](#)
- [Network Underlay Reference for 200 Mbps Bandwidth](#)
- [Network Underlay Reference for 500 Mbps Bandwidth](#)
- [Network Underlay Reference for 1 Gbps \(1000Mbps\) Bandwidth](#)
- [Network Underlay Reference for 2 Gbps \(2000Mbps\) Bandwidth](#)
- [Network Underlay Reference Beyond 2 Gbps Bandwidth](#)

Network Underlay Reference for 50 Mbps Bandwidth

The tables in this section provide referenceable information based on a network underlay with 100 Mbps of Bandwidth.

HCX vMotion at 50 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX vMotion.

Caution HCX vMotion is not supported with 50 Mbit/sec Network Underlays. For details about the minimum requirements, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

HCX Bulk and OSAM Migration at 50 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX Bulk Migration and OS Assisted Migration.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section)

- Testing is based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Maximum Parallel Migrations Supported and Maximum Delta Transfer values are provided for a given network underlay and presuppose other performance contributors (e.g. hypervisor or datastore performance is not a limiting factor).
- Maximum Delta Transfer for each two hour RPO cycle (the migration RPO is not a configurable setting).

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	Bulk Migration Supported?	Maximum Parallel Migrations Supported on this underlay	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 50 Mbps ■ > 1.0 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 50 Mbps ■ < 100 ms ■ < 0.1 % loss 	Yes	5	30 GB [0.1 % loss and 100 ms latency]
Threshold Packet Loss <ul style="list-style-type: none"> ■ 50 Mbps ■ < 100ms ■ 0.1 - 1.0% loss 	Yes	5	15 GB [1.0 % loss and 100 ms latency]
Threshold Packet Loss & Threshold Latency <ul style="list-style-type: none"> ■ 50 Mbps ■ 100 - 150ms ■ 0.1 - 1.0% loss 	Yes	5	6 GB [1.0 % loss and 150 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 50 Mbps ■ > 150ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Replication Assisted vMotion at 50 Mbps

Caution Replication Assisted vMotion is not supported with 50 Mbit/sec Network Underlays. For details about the minimum requirements, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Reference Data Transfer Times at 50 Mbps

This table provides a theoretical minimum transfer duration, assuming the following conditions:

- Factors like latency, packet loss, MTU fragmentation, and any additional performance constraints have already been accounted for in the transfer rate, and a sustained transfer rate.
- In actual deployments there are other factors like the operational overhead in migration processes completing, and configuring new migrations, and general maintenance activities. All of these factors will reduce achievable data transfer rates.

Caution Data transfer estimations are calculated using public data transfer formulas, for convenience. Values are rounded for readability.

Data Transfer Rate	100 GB VM Disk Data	250 GB VM Disk Data	1 TB VM Disk Data
47 Mbps	4hr 44 min	11hr 49min	2d

Network Underlay Reference for 100 Mbps Bandwidth

The tables in this section provide referenceable information based on a network underlay with 100 Mbps of Bandwidth.

HCX vMotion at 100 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX vMotion.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section).

- Reference values are based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Example Throughput Maximum column is intended to highlight the expected effects of packet loss and latency on data transfers on the given underlay, assuming no other environmental constraints observed.

Network Underlay Characterization Results (Bandwidth, Latency and Packet Loss %)	HCX vMotion Supported with this Underlay?	Parallel Migrations Supported Per IX (Per Service Mesh)	Example Throughput Maximum (Theoretical Best Case Estimate)
High Packet Loss Range <ul style="list-style-type: none"> ■ 100 Mbit/sec ■ > 0.1 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 100 Mbit/sec ■ < 100ms latency ■ < 0.0001 % loss 	Yes	1	~ 94 Mbps (.0001% loss & 10ms latency)
Threshold Underlay Conditions <ul style="list-style-type: none"> ■ 100 Mbit/sec ■ 100-150ms latency ■ 0.0001 - 0.1 % loss 	Yes	1	~ 77 Mbps (.0001% loss & 150ms latency)
High Latency Range <ul style="list-style-type: none"> ■ 100 Mbit/sec ■ > 150 ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

HCX Bulk and OSAM Migration at 100 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX Bulk Migration and OS Assisted Migration.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section)

- Testing is based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Maximum Parallel Migrations Supported and Maximum Delta Transfer values are provided for a given network underlay and presuppose other performance contributors (e.g. hypervisor or datastore performance is not a limiting factor).
- Maximum Delta Transfer for each two hour RPO cycle (the migration RPO is not a configurable setting).

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	Bulk Migration Supported?	Maximum Parallel Migrations Supported on this underlay	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 100 Mbps ■ > 1.0 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 100 Mbps ■ < 100 ms ■ < 0.1 % loss 	Yes	10	60 GB [0.1 % loss and 100 ms latency]
Threshold Packet Loss <ul style="list-style-type: none"> ■ 100 Mbps ■ < 100ms ■ 0.1 - 1.0% loss 	Yes	10	30 GB [1.0 % loss and 100 ms latency]
Threshold Packet Loss & Threshold Latency <ul style="list-style-type: none"> ■ 100 Mbps ■ 100 - 150ms ■ 0.1 - 1.0% loss 	Yes	10	12 GB [1.0 % loss and 150 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 100 Mbps ■ > 150ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Replication Assisted vMotion

Caution Replication Assisted vMotion is not supported with 100 Mbit/sec Network Underlays. For details about the minimum requirements, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Reference Data Transfer Times at 100 Mbit/sec

This table provides a theoretical minimum transfer duration, assuming the following conditions:

- Factors like latency, packet loss, MTU fragmentation, and any additional performance constraints have already been accounted for in the transfer rate, and a sustained transfer rate.
- In actual deployments there are other factors like the operational overhead in migration processes completing, and configuring new migrations, and general maintenance activities. All of these factors will reduce achievable data transfer rates.

Caution Data transfer estimations are calculated using public data transfer formulas, for convenience. Values are rounded for readability.

Data Transfer Rate	100 GB VM Disk Data	250 GB VM Disk Data	1 TB VM Disk Data
94 Mbps	2hr 25min	6hr 3min	1d 1hr
77 Mbps	2hr 54min	7hr 13min	1d 5hr

Network Underlay Reference for 150 Mbps Bandwidth

The tables in this section provide referenceable information based on a network underlay with 100 Mbps of Bandwidth.

HCX vMotion at 150 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX vMotion.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section).

- Reference values are based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Example Throughput Maximum column is intended to highlight the expected effects of packet loss and latency on data transfers on the given underlay, assuming no other environmental constraints observed.

Network Underlay Characterization Results (Bandwidth, Latency and Packet Loss %)	HCX vMotion Supported with this Underlay?	Parallel Migrations Supported Per IX (Per Service Mesh)	Example Throughput Maximum (Theoretical Best Case Estimate)
High Packet Loss Range <ul style="list-style-type: none"> 150 Mbit/sec > 0.1 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> 150 Mbit/sec < 100ms latency < 0.0001 % loss 	Yes	1	~ 142 Mbps (0.0001% loss & 10ms latency)
Threshold Underlay Conditions <ul style="list-style-type: none"> 150 Mbit/sec 100-150ms latency 0.0001 - 0.1 % loss 	Yes	1	~ 116 Mbps (0.0001% loss & 100ms latency)
High Latency Range <ul style="list-style-type: none"> 150 Mbit/sec > 150 ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

HCX Bulk and OSAM Migration at 150 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX Bulk Migration and OS Assisted Migration.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section)

- Testing is based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Maximum Parallel Migrations Supported and Maximum Delta Transfer values are provided for a given network underlay and presuppose other performance contributors (e.g. hypervisor or datastore performance is not a limiting factor).
- Maximum Delta Transfer for each two hour RPO cycle (the migration RPO is not a configurable setting).

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	Bulk Migration Supported?	Maximum Parallel Migrations Supported on this underlay	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 150 Mbps ■ > 1.0 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 150 Mbps ■ < 100 ms ■ < 0.1 % loss 	Yes	15	90 GB [0.1 % loss and 100 ms latency]
Threshold Packet Loss <ul style="list-style-type: none"> ■ 150 Mbps ■ < 100ms ■ 0.1 - 1.0% loss 	Yes	15	45 GB [1.0 % loss and 100 ms latency]
Threshold Packet Loss & Threshold Latency <ul style="list-style-type: none"> ■ 150 Mbps ■ 100 - 150ms ■ 0.1 - 1.0% loss 	Yes	15	18 GB [1.0 % loss and 150 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 150 Mbps ■ > 150ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Replication Assisted vMotion at 150 Mbps

This table primarily describes supportability for a given loss and latency conditions when using Replication Assisted vMotion.

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	RAV Migration Supported?	Maximum Parallel Migrations Supported	Maximum Concurrent Switchovers	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> 150 Mbps > 0.1 % loss * 	No	N/A	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> 150 Mbit/sec < 100ms latency < 0.0001 % loss 	Yes	15	1	90 GB
High Latency Range <ul style="list-style-type: none"> 150 Mbps > 150ms latency * 	No	N/A	1	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Reference Data Transfer Times at 150 Mbps

This table provides a theoretical minimum transfer duration, assuming the following conditions:

- Factors like latency, packet loss, MTU fragmentation, and any additional performance constraints have already been accounted for in the transfer rate, and a sustained transfer rate.
- In actual deployments there are other factors like the operational overhead in migration processes completing, and configuring new migrations, and general maintenance activities. All of these factors will reduce achievable data transfer rates.

Caution Data transfer estimations are calculated using public data transfer formulas, for convenience. Values are rounded for readability.

Data Transfer Rate	100 GB VM Disk Data	250 GB VM Disk Data	1 TB VM Disk Data
142 Mbps	1hr 34min	3hr 55min	15hr 39min
116 Mbps	1hr 55min	4hr 48min	19hr 10min

Network Underlay Reference for 200 Mbps Bandwidth

The tables in this section provide referenceable information based on a network underlay with 100 Mbps of Bandwidth.

HCX vMotion at 200 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX vMotion.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section).

- Reference values are based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Example Throughput Maximum column is intended to highlight the expected effects of packet loss and latency on data transfers on the given underlay, assuming no other environmental constraints observed.

Network Underlay Characterization Results (Bandwidth, Latency and Packet Loss %)	HCX vMotion Supported with this Underlay?	Parallel Migrations Supported Per IX (Per Service Mesh)	Example Throughput Maximum (Theoretical Best Case Estimate)
High Packet Loss Range <ul style="list-style-type: none"> ■ 200 Mbit/sec ■ > 0.1 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 200 Mbit/sec ■ < 90ms latency ■ < 0.0001 % loss 	Yes	1	~ 189 Mbps (0.0001% loss & 10ms latency)
Threshold Underlay Conditions <ul style="list-style-type: none"> ■ 200 Mbit/sec ■ 90-150ms latency ■ 0.0001 - 0.1 % loss 	Yes	1	~ 129 Mbps (0.0001% loss & 90 ms latency)
High Latency Range <ul style="list-style-type: none"> ■ 200 Mbit/sec ■ > 150 ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

HCX Bulk and OSAM Migration at 200 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX Bulk Migration and OS Assisted Migration.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section)

- Testing is based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.

- The Maximum Parallel Migrations Supported and Maximum Delta Transfer values are provided for a given network underlay and presuppose other performance contributors (e.g. hypervisor or datastore performance is not a limiting factor).
- Maximum Delta Transfer for each two hour RPO cycle (the migration RPO is not a configurable setting).

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	Bulk Migration Supported?	Maximum Parallel Migrations Supported on this underlay	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 200 Mbps ■ > 1.0 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 200 Mbps ■ < 100 ms ■ < 0.1 % loss 	Yes	20	120 GB [0.1 % loss and 100 ms latency]
Threshold Packet Loss <ul style="list-style-type: none"> ■ 200 Mbps ■ < 100ms ■ 0.1 - 1.0% loss 	Yes	20	60 GB [1.0 % loss and 100 ms latency]
Threshold Packet Loss & Threshold Latency <ul style="list-style-type: none"> ■ 200 Mbps ■ 100 - 150ms ■ 0.1 - 1.0% loss 	Yes	20	24 GB [1.0 % loss and 150 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 200 Mbps ■ > 150ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Replication Assisted vMotion at 200 Mbps

This table primarily describes supportability for a given loss and latency conditions when using Replication Assisted vMotion.

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	RAV Migration Supported?	Maximum Parallel Migrations Supported	Maximum Concurrent Switchovers	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 200 Mbps ■ > 0.1 % loss * 	No	N/A	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 200 Mbit/sec ■ < 90ms latency ■ < 0.0001 % loss 	Yes	20	1	120 GB [0.1 % loss and 100 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 200 Mbps ■ > 150ms latency * 	No	N/A	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Reference Data Transfer Times at 200 Mbps

This table provides a theoretical minimum transfer duration, assuming the following conditions:

- Factors like latency, packet loss, MTU fragmentation, and any additional performance constraints have already been accounted for in the transfer rate, and a sustained transfer rate.
- In actual deployments there are other factors like the operational overhead in migration processes completing, and configuring new migrations, and general maintenance activities. All of these factors will reduce achievable data transfer rates.

Caution Data transfer estimations are calculated using public data transfer formulas, for convenience. Values are rounded for readability.

Data Transfer Rate	100 GB VM Disk Data	250 GB VM Disk Data	1 TB VM Disk Data	100 TB VM Disk Data
189 Mbps	1hr 11min	2hr 57min	11hr 46min	49 days
129 Mbps	1hr 44min	4hr 19min	17hr 14min	71 days 19hr

Network Underlay Reference for 500 Mbps Bandwidth

The tables in this section provide referenceable information based on a network underlay with 100 Mbps of Bandwidth.

HCX vMotion at 500 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX vMotion.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section).

- Reference values are based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Example Throughput Maximum column is intended to highlight the expected effects of packet loss and latency on data transfers on the given underlay, assuming no other environmental constraints observed.

Network Underlay Characterization Results (Bandwidth, Latency and Packet Loss %)	HCX vMotion Supported with this Underlay?	Parallel Migrations Supported Per IX (Per Service Mesh)	Example Throughput Maximum (Theoretical Best Case Estimate)
High Packet Loss Range <ul style="list-style-type: none"> ■ 500 Mbit/sec ■ > 0.1 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 500 Mbit/sec ■ < 50ms latency ■ < 0.00005 % loss 	Yes	1	~ 474 Mbps (0.00005% loss & 10ms latency)
Threshold Underlay Conditions <ul style="list-style-type: none"> ■ 500 Mbit/sec ■ 50-150ms latency ■ 0.00005 - 0.1 % loss 	Yes	1	~ 330 Mbps (0.00005% loss & 50 ms latency)
High Latency Range <ul style="list-style-type: none"> ■ 500 Mbit/sec ■ > 150 ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

HCX Bulk and OSAM Migration at 500 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX Bulk Migration and OS Assisted Migration.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section)

- Testing is based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.

- The Maximum Parallel Migrations Supported and Maximum Delta Transfer values are provided for a given network underlay and presuppose other performance contributors (e.g. hypervisor or datastore performance is not a limiting factor).
- Maximum Delta Transfer for each two hour RPO cycle (the migration RPO is not a configurable setting).

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	Bulk Migration Supported?	Maximum Parallel Migrations Supported on this underlay	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 500 Mbps ■ > 1.0 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 500 Mbps ■ < 100 ms ■ < 0.1 % loss 	Yes	50	300 GB [0.1 % loss and 100 ms latency]
Threshold Packet Loss <ul style="list-style-type: none"> ■ 500 Mbps ■ < 100ms ■ 0.1 - 1.0% loss 	Yes	50	150 GB [1.0 % loss and 100 ms latency]
Threshold Packet Loss & Threshold Latency <ul style="list-style-type: none"> ■ 500 Mbps ■ 100 - 150ms ■ 0.1 - 1.0% loss 	Yes	50	60 GB [1.0 % loss and 150 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 500 Mbps ■ > 150ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Replication Assisted vMotion at 500 Mbps

This table primarily describes supportability for a given loss and latency conditions when using Replication Assisted vMotion.

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	RAV Migration Supported?	Maximum Parallel Migrations Supported	Maximum Concurrent Switchovers	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 500 Mbps ■ > 0.1 % loss * 	No	N/A	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 500 Mbit/sec ■ < 50ms latency ■ < 0.00005 % loss 	Yes	50	1	300 GB [0.1 % loss and 100 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 500 Mbps ■ > 150ms latency * 	No	N/A	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Reference Data Transfer Times at 500Mbit/sec

This table provides a theoretical minimum transfer duration, assuming the following conditions:

- Factors like latency, packet loss, MTU fragmentation, and any additional performance constraints have already been accounted for in the transfer rate, and a sustained transfer rate.
- In actual deployments there are other factors like the operational overhead in migration processes completing, and configuring new migrations, and general maintenance activities. All of these factors will reduce achievable data transfer rates.

Caution Data transfer estimations are calculated using public data transfer formulas, for convenience. Values are rounded for readability.

Data Transfer Rate	100 GB VM Disk Data	250 GB VM Disk Data	1 TB VM Disk Data	100 TB VM Disk Data
474 Mbps	28min	1hr 11min	4hr 42min	19days 13hr
330 Mbps	41min	1hr 41min	6hr 45min	28days 2hr

Network Underlay Reference for 1 Gbps (1000Mbps) Bandwidth

The tables in this section provide referenceable information based on a network underlay with 100 Mbps of Bandwidth.

HCX vMotion at 1000 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX vMotion.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section).

- Reference values are based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Example Throughput Maximum column is intended to highlight the expected effects of packet loss and latency on data transfers on the given underlay, assuming no other environmental constraints observed.

Network Underlay Characterization Results (Bandwidth, Latency and Packet Loss %)	HCX vMotion Supported with this Underlay?	Parallel Migrations Supported Per IX (Per Service Mesh)	Example Throughput Maximum (Theoretical Best Case Estimate)
High Packet Loss Range <ul style="list-style-type: none"> ■ 1000 Mbit/sec ■ > 0.1 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 1000 Mbit/sec ■ < 50ms latency ■ < 0.00001 % loss 	Yes	1	~ 949 Mbps (0.00001% loss 10ms latency)
Threshold Underlay Conditions <ul style="list-style-type: none"> ■ 1000 Mbit/sec ■ 50-150ms latency ■ 0.00001 - 0.1 % loss 	Yes	1	~ 738 Mbps (0.00001% loss & 50ms latency)
High Latency Range <ul style="list-style-type: none"> ■ 1000 Mbit/sec ■ > 150 ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

HCX Bulk and OSAM Migration at 1000 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX Bulk Migration and OS Assisted Migration.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section)

- Testing is based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.

- The Maximum Parallel Migrations Supported and Maximum Delta Transfer values are provided for a given network underlay and presuppose other performance contributors (e.g. hypervisor or datastore performance is not a limiting factor).
- Maximum Delta Transfer for each two hour RPO cycle (the migration RPO is not a configurable setting).

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	Bulk Migration Supported?	Maximum Parallel Migrations Supported on this underlay	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 1000 Mbps ■ > 1.0 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 1000 Mbps ■ < 100 ms ■ < 0.1 % loss 	Yes	100	600 GB [0.1 % loss and 100 ms latency]
Threshold Packet Loss <ul style="list-style-type: none"> ■ 1000 Mbps ■ < 100ms ■ 0.1 - 1.0% loss 	Yes	100	300 GB [1.0 % loss and 100 ms latency]
Threshold Packet Loss & Threshold Latency <ul style="list-style-type: none"> ■ 1000 Mbps ■ 100 - 150ms ■ 0.1 - 1.0% loss 	Yes	100	120 GB [1.0 % loss and 150 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 1000 Mbps ■ > 150ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Replication Assisted vMotion at 1000 Mbps

This table primarily describes supportability for a given loss and latency conditions when using Replication Assisted vMotion.

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	RAV Migration Supported?	Maximum Parallel Migrations Supported	Maximum Concurrent Switchovers	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 1000 Mbps ■ > 1.0 % loss * 	No	N/A	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 1000 Mbit/sec ■ < 50ms latency ■ < 0.00001 % loss 	Yes	100	1	600 GB [0.1 % loss and 100 ms latency]
High Latency Range <ul style="list-style-type: none"> ■ 1000 Mbps ■ > 150ms latency * 	No	N/A	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Reference Data Transfer Times at 1000Mbit/sec

This table provides a theoretical minimum transfer duration, assuming the following conditions:

- Factors like latency, packet loss, MTU fragmentation, and any additional performance constraints have already been accounted for in the transfer rate, and a sustained transfer rate.
- In actual deployments there are other factors like the operational overhead in migration processes completing, and configuring new migrations, and general maintenance activities. All of these factors will reduce achievable data transfer rates.

Caution Data transfer estimations are calculated using public data transfer formulas, for convenience. Values are rounded for readability.

Data Transfer Rate	250 GB VM Disk Data	1 TB VM Disk Data	100 TB VM Disk Data	1 PB VM Disk Data
949 Mbps	35 min	2 hr 21 min	9 days 19 hr	97 days 14 hr
738 Mbps	46 min	3 hr 1 min	12 days 13 hr	125 days 12 hr

Network Underlay Reference for 2 Gbps (2000Mbps) Bandwidth

The tables in this section provide referenceable information based on a network underlay with 100 Mbps of Bandwidth.

HCX vMotion at 2000 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX vMotion.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section).

- Reference values are based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.
- The Example Throughput Maximum column is intended to highlight the expected effects of packet loss and latency on data transfers on the given underlay, assuming no other environmental constraints observed.

Network Underlay Characterization Results (Bandwidth, Latency and Packet Loss %)	HCX vMotion Supported with this Underlay?	Parallel Migrations Supported Per IX (Per Service Mesh)	Example Throughput Maximum (Theoretical Best Case)
High Packet Loss Range <ul style="list-style-type: none"> ■ 2000 Mbit/sec ■ > 0.1 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 2000 Mbit/sec ■ < 30ms latency ■ < 0.00001 % loss 	Yes	1	~ 1898 Mbps (0.00001% loss & 10ms latency)
Threshold Underlay Conditions <ul style="list-style-type: none"> ■ 2000 Mbit/sec ■ 30-150ms latency ■ 0.00001 - 0.1 % loss 	Yes	1	~ 1231 Mbps (0.00001% loss & 30ms latency)
High Latency Range <ul style="list-style-type: none"> ■ 2000 Mbit/sec ■ > 150 ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

HCX Bulk and OSAM Migration at 2000 Mbps

This table primarily describes supportability for a given loss and latency conditions when using HCX Bulk Migration and OS Assisted Migration.

The reader should identify the single row that matches their underlay's packet loss and latency characteristics (measured during the underlay characterization effort in the previous section)

- Testing is based on 1500 end to end Path Network Underlay MTU with correctly configured HCX Network Profile Uplink MTU.

- The Maximum Parallel Migrations Supported and Maximum Delta Transfer values are provided for a given network underlay and presuppose other performance contributors (e.g. hypervisor or datastore performance is not a limiting factor).
- Maximum Delta Transfer for each two hour RPO cycle (the migration RPO is not a configurable setting).

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	Bulk Migration Supported?	Maximum Parallel Migrations Supported on this underlay	Maximum Delta Transfer Supported
High Packet Loss Range <ul style="list-style-type: none"> ■ 2000 Mbps ■ > 1.0 % loss * 	No	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 2000 Mbps ■ < 100 ms ■ < 0.1 % loss 	Yes	200	1200 GB (0.1 % loss and 100 ms latency)
Threshold Packet Loss <ul style="list-style-type: none"> ■ 2000 Mbps ■ < 100ms ■ 0.1 - 1.0% loss 	Yes	200	600 GB (1.0 % loss and 100 ms latency)
Threshold Packet Loss & Threshold Latency <ul style="list-style-type: none"> ■ 2000 Mbps ■ 100 - 150ms ■ 0.1 - 1.0% loss 	Yes	200	240 GB (1.0 % loss and 150 ms latency)
High Latency Range <ul style="list-style-type: none"> ■ 2000 Mbps ■ > 150ms latency * 	No	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Replication Assisted vMotion at 2000 Mbps

This table primarily describes supportability for a given loss and latency conditions when using Replication Assisted vMotion.

Network Underlay Characterization Results (Bandwidth, Latency and packet loss %)	RAV Migration Supported?	Maximum Parallel Migrations Supported	Maximum Concurrent Switchovers	Maximum Delta Transfer Supported (During two hour cycle)
High Packet Loss Range <ul style="list-style-type: none"> ■ 2000 Mbps ■ > 0.1 % loss * 	No	N/A	N/A	N/A
Preferred Underlay Conditions <ul style="list-style-type: none"> ■ 2000 Mbit/sec ■ < 30ms latency ■ < 0.00001 % loss 	Yes	200	N/A	1200 GB (0.1 % loss and 100 ms latency)
High Latency Range <ul style="list-style-type: none"> ■ 2000 Mbps ■ > 150ms latency * 	No	N/A	N/A	N/A

Caution * This underlay parameter does not meet the minimum requirements. For details, see the [Network Underlay Minimum Requirements](#) page, in the HCX User Guide.

Reference Data Transfer Times at 2000Mbit/sec

This table provides a theoretical minimum transfer durations, assuming the following conditions:

- Factors like latency, packet loss, MTU fragmentation, and any additional performance constraints have already been accounted for in the transfer rate, and a sustained transfer rate.
- In actual deployments there are other factors like the operational overhead in migration processes completing, and configuring new migrations, and general maintenance activities. All of these factors will reduce achievable data transfer rates.

Caution Data transfer estimations are calculated using public data transfer formulas, for convenience. Values are rounded for readability.

Data Transfer Rate	250 GB VM Disk Data	1 TB VM Disk Data	100 TB VM Disk Data	1 PB VM Disk Data
1898 Mbps	18 min	1 hr 11 min	4 days 22 hr	48 days 19 hr
1231 Mbps	27 min	1 hr 48 min	7 days 13 hr	75 days 6 hr

Network Underlay Reference Beyond 2 Gbps Bandwidth

The tables in this section provide referenceable information based on a network underlay with higher multi-gigabit connections.

HCX Network Underlay Maximum

The maximum underlay bandwidth supported by a single HCX-IX migration appliance is 2Gbps. At maximum concurrency (when not limited by other environmental factors) bandwidth beyond 2 Gbps cannot be utilized by a single IX appliance.

An appliance scale out approach can be used to increase concurrency and total utilization of the high bandwidth network underlays.

HCX has a 200 migration concurrency maximum. Any portion of this limit can be distributed using parallel IX appliances. This assumes that the IX VMs are distributed to different ESX hosts, and the storage itself is capable of supporting a high transfer rate.

HCX IX appliances can be increased by adding Service Meshes to an HCX installation. One IX is deployed per Service Mesh.

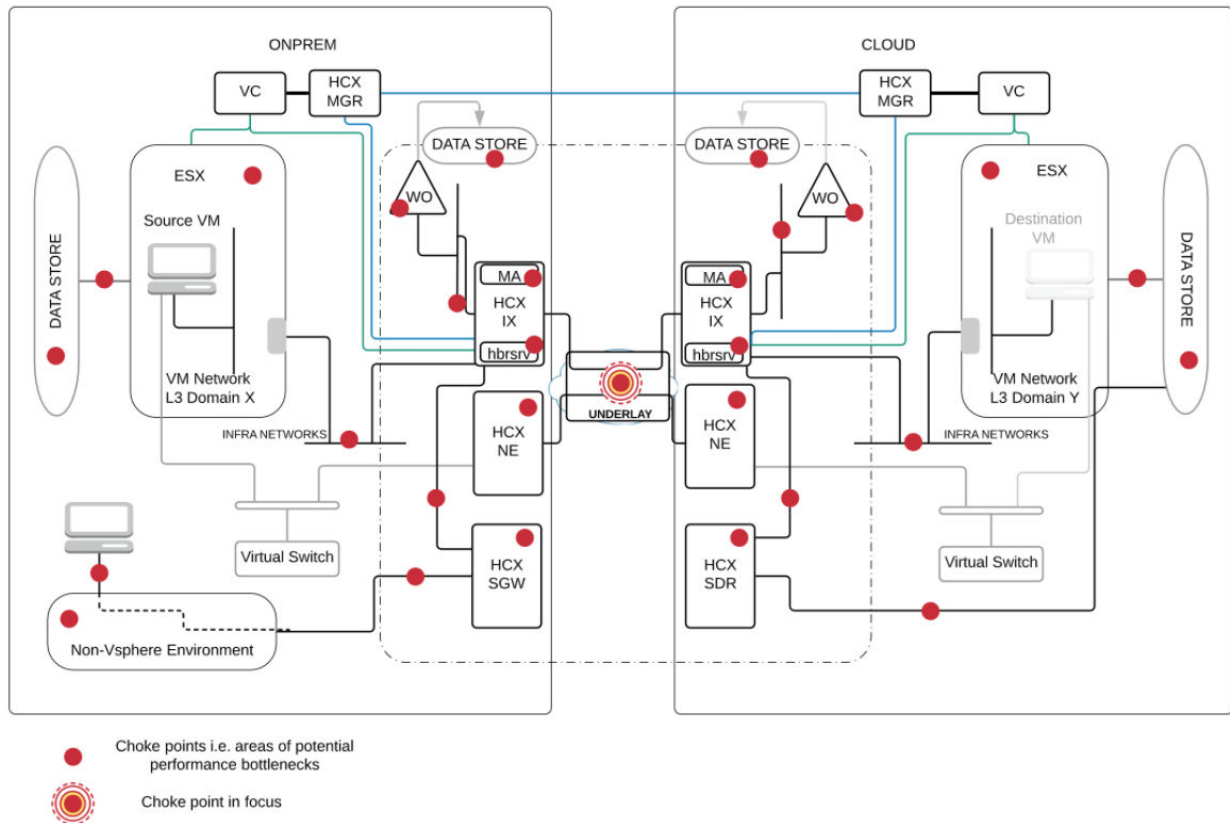
Performance Factors Beyond the Network Underlay

4

The performance of HCX migrations is affected by multiple elements. Many of these elements are beyond the scope of this document, but it is important for users to understand any possible performance bottleneck.

In general, the achievable performance of a migration takes many elements into consideration, like CPU, Memory, Storage and Network performance across all components. Existing load and other factors like variable availability of those resources can also affect performance. Any single element can become a choke point to the overall achievable performance. The purpose of this section is to provide awareness of potential performance limiting factors in HCX deployments.

Figure 4-1. Performance Chokepoints in an HCX deployment.



Other Common Performance Factors

- Compute (CPU/Memory) resource allocations for the HCX appliances.
- Compute load on ESXi servers hosting HCX appliances.
- Network path for ESXi management traffic.
- Storage allocations for the IX,NE,SDR, and SGW appliances.
- Storage (5000 IOPS) per WO appliance.
- Storage load on HCX deployment datastores.
- Storage load on virtual machine datastores.
- Number of virtual machines to be migrated.
- Number of VMDKs per virtual machine affects total possible concurrency.
- Number of VMDKs per virtual machine affects checksum operations.
- Virtual machine activity (churn/rate of change).
- Virtual machine north/south traffic can affect how much traffic is transmitted over the Network Extension tunnels, and can reduce the amount of network resources available for migration.
- ESXi virtual machine load affects replication or vMotion performance.
- ESXi host count available for replication affects the overall performance.
- ESXi load - host/cluster operations like DRS rebalancing and manual cluster vMotions affect HCX migration performance.
- ESXi maintenance can result in reconfiguration of replications and impact performance.
- The cluster count affects how many HCX Service Meshes are configurable.
- Non-vSphere environment performance effects.
- In multiple service mesh deployments:
 - Appliance placement (IX and WO appliance affinity)
 - Appliance placement (IX to IX appliance anti-affinity)

About the Author

5

Authors

Gabe Rosas is a Staff Technical Product Manager for all aspects of VMware HCX in the Networking and Security Business Unit at VMware. He is experienced in designing and operating traditional and software-defined datacenter and cloud infrastructure.

Siddharth Ekbote is a Senior Staff Engineer in the VMware HCX R&D organization. He has deep experience in designing, implementing and troubleshooting networking appliances, including WAN Optimizers, Load Balancers and cloud networking Infrastructure.

Contributors and Reviewers

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

Jeffrey Moore is a Senior Staff Cloud Architect at VMware and is the Lead Worldwide for Azure VMware Solution (AVS). His primary focus is to provide solutions for AVS architectures as well as collaborate with the Microsoft and VMware Product Teams to drive Roadmap discussion based on Customer/Partner discussions.

Bilal Ahmed is Customer Success Cloud Architect at VMware. He loves all things vSphere and has spent the last few years helping customers migrate their VMware workloads into private and public clouds. Bilal is VCDX #251