



HPE SIMPLIVITY FOR VSPHERE NETWORKING BEST PRACTICES



CONTENTS

Executive summary.....	3
HPE SimpliVity network overview.....	3
HPE SimpliVity host.....	4
Virtual Controller.....	5
Management Virtual Appliance network interface.....	5
HPE Integrated Lights-Out management interface.....	6
Standard ESXi hosts (compute nodes).....	6
Standard ESXi host (compute node) networking requirements.....	6
Network infrastructure requirements.....	7
Layer 2 networking features.....	8
Jumbo packets/MTU.....	8
VLANs.....	8
Port settings—LAN switch and VMware.....	9
NIC teaming.....	9
Physical network connectivity.....	10
Direct connect clusters.....	10
Network topologies.....	11
Full mesh.....	11
Remote office/Remote branch.....	11
HPE SimpliVity support for NSX-T.....	11
Stretch clusters.....	14
Remote backups.....	15
Bandwidth throttling for HPE SimpliVity backup and restore operations.....	15
Remote vCenter management.....	16
HPE SimpliVity Arbiter.....	16
Aruba switches.....	17
HPE Composable Cloud for HPE ProLiant DL.....	18
Switch requirements.....	18
HPE SimpliVity network checklist.....	18



EXECUTIVE SUMMARY

This document describes best practices for networking in HPE SimpliVity environments. The recommendations are a combination of best practices from VMware®, the networking industry, and successful practices developed by the HPE SimpliVity support, field engineers, and solutions architects. The intended audience is systems engineers and administrators with a variety of backgrounds and skill levels that design, deploy, and manage HPE SimpliVity systems.

The purpose of this document is to provide a basic understanding of the HPE SimpliVity networking environment and common deployment scenarios for an HPE SimpliVity Federation. You should have basic knowledge of VMware vSphere® networking to understand this document. It is not intended to be used as a reference architecture for production deployment.

HPE SIMPLIVITY NETWORK OVERVIEW

HPE SimpliVity hosts run VMware ESXi™ software so the network configuration is based on VMware vCenter Server® networking technology.

An HPE SimpliVity host requires three separate virtual network interfaces: management, storage, and federation. You can configure HPE SimpliVity hosts during deployment through deployment manager to assign either separate NICs or a single NIC to handle HPE SimpliVity network traffic.

Figure 1 shows the configuration of an HPE SimpliVity host with two vSwitches and two physical NICs.

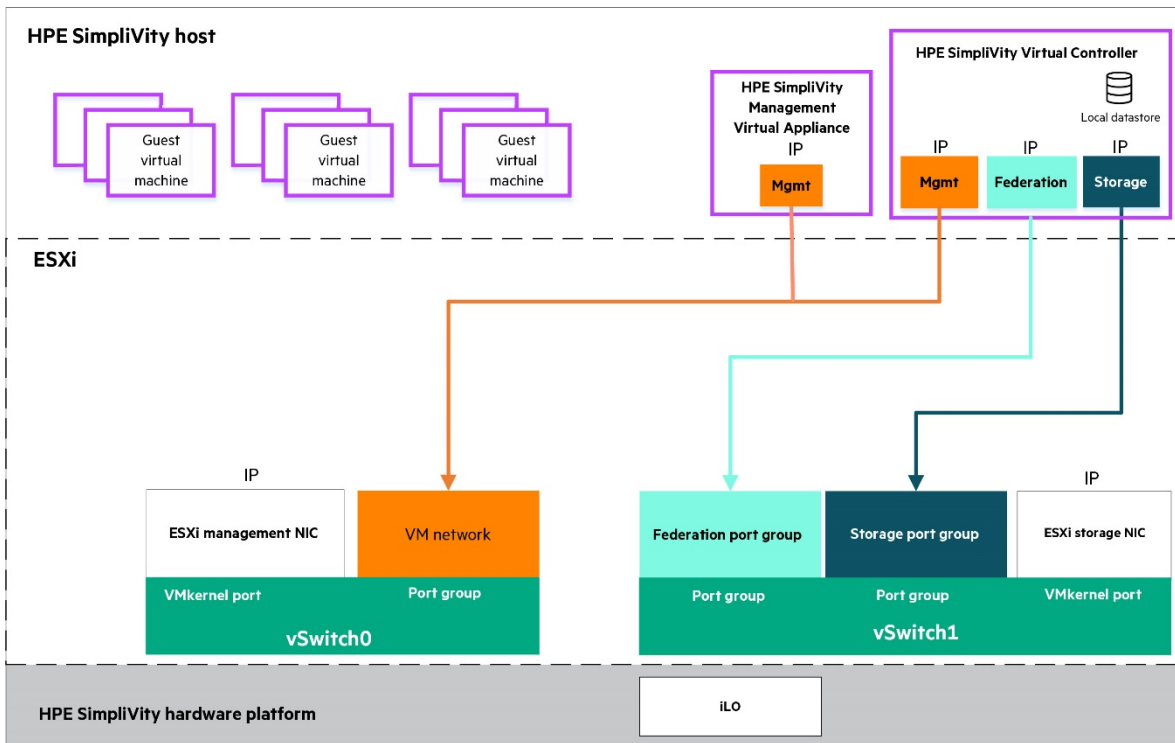


FIGURE 1. HPE SimpliVity host with two vSwitches and two physical NICs

Figure 2 shows the configuration of an HPE SimpliVity hosts with one vSwitch and one physical NIC.



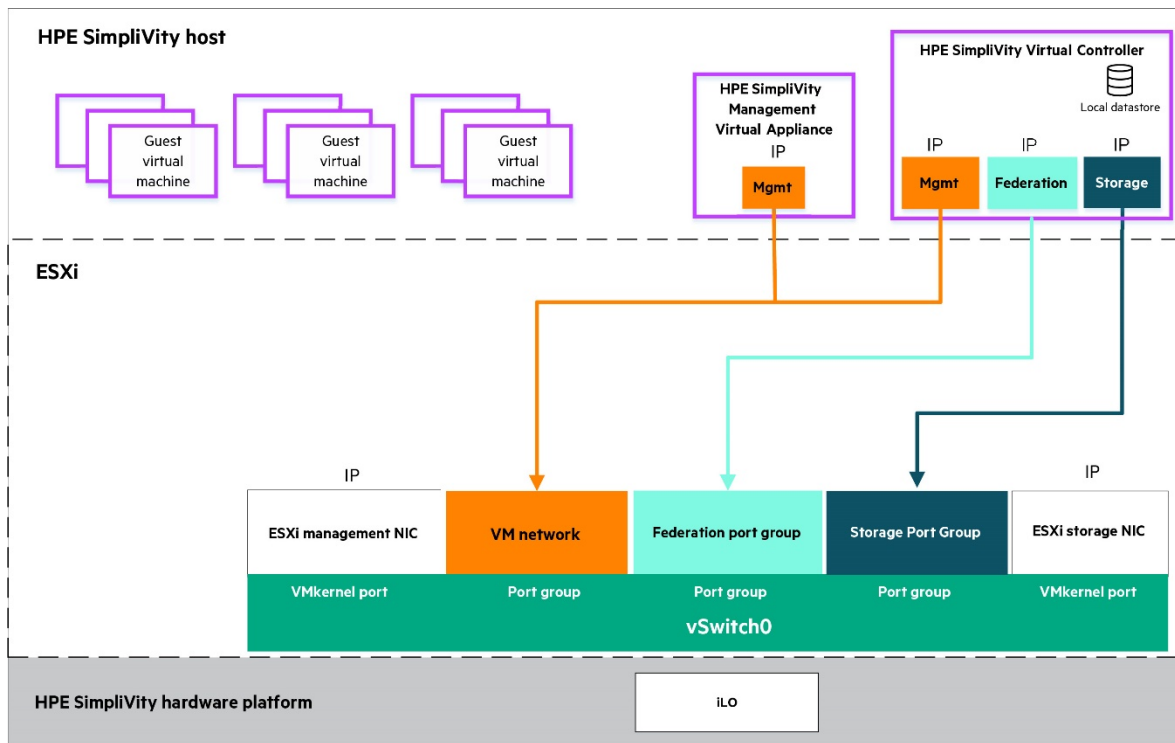


FIGURE 2. HPE SimpliVity host with one vSwitch and one physical NIC

Regardless of the configuration that you choose, connect the management network to vSwitch0. The management port groups associated with vSwitch0 can be connected to any speed network, but the federation and storage port groups should be connected to the fastest available physical NIC.

The VMware ESXi hypervisor on the HPE SimpliVity host must be configured with two VMkernel ports, each with its own static IP address. These IP addresses must be configured on separate networks (subnets). The HPE SimpliVity host management and host storage network interfaces use the two VMkernel ports with static IP addresses.

The networking components for an HPE SimpliVity host are described in the following section.

HPE SimpliVity host

An HPE SimpliVity host runs HPE SimpliVity software and VMware ESXi. The host requires six IP addresses when running in a peer-based management federation and seven IP addresses when running in a centrally managed federation. You assign the IP addresses during deployment. Note that one IP address is for HPE iLO which is not described in this paper.

Two IP addresses are assigned to the HPE SimpliVity host.

TABLE 1. HPE SimpliVity host network interfaces

HPE SimpliVity host network interface	Description
HPE SimpliVity host management network interface	<p>The default port group name for the host is <code>Management Network</code>, and the default kernel adapter is <code>vmk0</code>.</p> <ul style="list-style-type: none"> Used by the HPE SimpliVity host to connect vCenter Server Used by guest virtual machines (by default) for all ingress and egress Ethernet traffic, but customers can also create their own virtual machine networks This is the IP gateway for the ESXi hypervisor; it is also the IP default gateway It can be either 1, 10, or 25 Gbps It can be a single interface or teamed interfaces The maximum latency (WAN) is 300 ms
HPE SimpliVity host storage network interface	<p>The default port group name for the host storage network is <code>SVT_Storage</code>, and the default kernel adapter is <code>vmk1</code>.</p> <ul style="list-style-type: none"> Used by the HPE SimpliVity host to connect to the HPE SimpliVity storage using the NFS protocol via HPE SimpliVity Virtual Appliance (also called HPE OmniStack Virtual Controller) VMkernel port, in the same VLAN as Storage Port Group It can be either 10 or 25 Gbps It can be a single interface or teamed interfaces



Virtual Controller

The HPE SimpliVity Virtual Appliance/HPE OmniStack Virtual Controller is a virtual machine that serves storage to the ESXi host over the NFS protocol. The storage is then delivered to the guest virtual machines as local storage from the ESXi host. The HPE SimpliVity Virtual Appliance requires three static IP addresses that are configured on three separate networks/subnets. Its network interfaces are described in Table 2.

NOTE

The virtual controller is called as the HPE SimpliVity Virtual Appliance also known as the HPE OmniStack Virtual Controller.

TABLE 2. Virtual Controller network interfaces

Virtual Controller network interfaces	Description
Virtual Controller management network interface	<p>The default port group name for the Virtual Controller is <code>VM_Network</code>, and the default network interface is <code>eth0</code>. Used for:</p> <ul style="list-style-type: none"> • North-south communication • Communication with non-local clusters and for communicating with the vCenter Server and the HPE SimpliVity Arbiter service • Management access to the HPE SimpliVity Virtual Appliance operating system • The IP network gateway for the HPE SimpliVity Virtual Appliance • Virtual machine guest data traffic within the HPE SimpliVity host and externally (by default) but guest traffic can also be on a separate dedicated network • L2 traffic, typically using trunked VLANs • L3 routable traffic
Virtual Controller storage network interface	<p>The default port group name for the Virtual Controller storage network is <code>SVT_StoragePortGroup</code>, and the default network interface is <code>eth1</code>. Used by:</p> <ul style="list-style-type: none"> • The Virtual Controller for NFS traffic between HPE SimpliVity hosts and standard ESXi hosts (compute nodes) • The compute VMware vSphere® Storage vMotion® interface • L2 traffic only, not-routable <p>Storage is host local. The network must support at least 10 Gbps. It must use Jumbo frames, typically, the MTU size is 9000, but this size may vary in your environment.</p>
Virtual Controller federation network interface	<p>The default port group name for the Virtual Controller federation network is <code>SVT_FedPortGroup</code>, and the default network interface is <code>eth2</code>. It is the preferred interface for virtual controller-to-virtual controller network communications within the cluster. It must support at least 10 Gbps and be configured on the federation port group. The federation network is local, not routed. Used for:</p> <ul style="list-style-type: none"> • Data replication and backup traffic between HPE SimpliVity Virtual Appliances • Intra-federation communication between Virtual Controllers (TCP and UDP 22122) • L2 traffic only, non-routable • If communication is to a remote cluster, it will be routed out of the management interface through the default gateway, to the management interface of the remote HPE SimpliVity Virtual Appliance. Unless, you have two clusters in the same physical space that are also sharing same VLAN and do not need to route between the two clusters; in that case, the two clusters can share the same subnet for federation

Management Virtual Appliance network interface

The Management Virtual Appliance is a virtual machine dedicated to the federation. It manages the services within the hosts in the federation. It requires one static IP address with one MVA per vCenter. By default, it is configured to use the Virtual Controller management network.

HPE SimpliVity Virtual Appliance networking best practices

The management, storage, and federation network interfaces must be the same on all of the HPE SimpliVity hosts in the cluster. Each HPE SimpliVity host must be configured with matching values for:

- VLAN IDs
- IP subnet/masks
- Port group names

For ease of management, use an IP addressing scheme that makes it easy to identify the host. This is helpful when reviewing logs, planning for new hosts, creating documentation, and more.



For IP addresses configured on the host, for each separate subnet, use the same even number in the fourth octet. For example:

- **HPE iLO:** 10.1.1.8/24
- **Management:** 10.1.2.8/24
- **Storage:** 10.1.3.8/24

For the HPE SimpliVity Virtual Appliance, use the same odd number in the fourth octet

- **Management:** 10.1.2.9/24
- **Storage:** 10.1.3.9/24
- **Federation:** 10.1.4.9/24

HPE Integrated Lights-Out management interface

Enable remote HPE iLO management by configuring an HPE SimpliVity host with a unique iLO address for remote hardware diagnostics and SSH access to ESXi.

HPE iLO best practices

- Configure the HPE iLO management interface with a static IP address on a separate network dedicated to iLO devices
- Configure DNS forward and reverse lookup records for the IP address of the iLO devices

Standard ESXi hosts (compute nodes)

You can add standard ESXi hosts (compute nodes) to an HPE SimpliVity cluster when you need more compute and memory resources. This is especially useful for workloads that are not constrained by storage resources. HPE SimpliVity recommends the following compute nodes configurations:

TABLE 3. Compute node configuration

Use case	Compute node	HPE SimpliVity host
VDI	1	1
General server	2	1

The compute nodes must:

- Run the same version of ESXi software as the HPE SimpliVity hosts
- Reside in the same cluster as the HPE SimpliVity hosts that it is mapped to
- Share an HPE SimpliVity datastore

In a 1:1 mapping, each compute node is mapped to an individual HPE SimpliVity Virtual Appliance storage IP. In a 2:1 mapping, two compute nodes are mapped to one Virtual Controller storage IP.

To improve performance, you can optionally install the VAAI-NAS plug-in on each compute node.

Standard ESXi host (compute node) networking requirements

Compute nodes must be configured with two VMkernel ports. One of the VMkernel ports must be on the ESXi management network interface and the other one on the ESXi storage network interface. The compute nodes use the ESXi management network interface as follows:

- Connects the ESXi hypervisor to vCenter Server
- Used by guest virtual machines for all ingress and egress Ethernet traffic (by default), but guest traffic can also be on a separate dedicated network
- This is the IP Gateway used by the standard ESXi hypervisor, and it is also the IP default gateway

The standard ESXi hosts uses the ESXi storage network interface as follows:

- Provides the standard ESXi host connection to the HPE SimpliVity storage using the NFS protocol via the HPE SimpliVity Virtual Appliance
- Enables VMware vSphere Storage vMotion traffic between standard ESXi hosts and HPE SimpliVity hosts



To allow a standard ESXi host access to an HPE SimpliVity datastore, you must accomplish the tasks in Table 4.

TABLE 4. Steps to enable compute node access to an HPE SimpliVity datastore

Step #	Task	Description
1	Determine the IP address for the storage traffic	<p>You must know the network IP address of the HPE SimpliVity host you plan to use to share the datastore. The HPE SimpliVity host provides two potential paths:</p> <ul style="list-style-type: none"> • HPE SimpliVity Virtual Appliance storage network IP address (recommended for failover) • HPE SimpliVity Virtual Appliance management network IP address Your physical network connectivity also impacts the IP address you use. <p>With the switch-connected method for the 10GbE storage network, use the storage network IP address of the HPE SimpliVity Virtual Appliance. This network provides higher bandwidth and failover capability.</p> <p>With the direct connect method for the 10GbE storage network, specify the management network IP address of the HPE SimpliVity Virtual Appliance. In this scenario, there is no failover capability.</p>
2	Configure advanced settings on the standard ESXi host (compute node)	<p>You must set the following values on the compute node and then reboot it.</p> <ul style="list-style-type: none"> • <code>Net.TcpipHeapMax = 1536</code> • <code>Net.TcpipHeapSize = 32</code> • <code>NFS.MaxVolumes = 256</code> • <code>SunRPC.MaxConnPerIP = 128</code> • <code>NFS.MaxQueueDepth = 128</code>
3	Modify the <code>/etc/hosts</code> file on the standard ESXi host	<p>Edit the <code>/etc/hosts</code> file to include a line similar to the following:</p> <pre>nnn.nnn.nnn.nnn omni.cube.io</pre> <p>Where <code>nnn.nnn.nnn.nnn</code> is the storage IP address you obtained in Step 1. You can optionally use the management IP, but it can cause data unavailability, so it is not a recommended configuration.</p>
4	Enable HPE SimpliVity datastore sharing with an HPE SimpliVity host	<p>In the VMware vSphere® Web Client, navigate to the All HPE SimpliVity Actions → Manage Standard ESXi Hosts, then allow the hosts to share data with an HPE SimpliVity datastore.</p> <p>Or use the <code>svt-datastore-share</code> CLI command.</p>

Network infrastructure requirements

HPE SimpliVity requires that certain network infrastructure services be available on the customer network where it is deployed.

IPv4

The HPE SimpliVity environment requires an IPv4 network. Your network can use both modes (both IPv4 and IPv6). However, if this is the case, you must supply IPv4 addresses for any ESXi host that communicates with the federation across a network.

DNS

HPE SimpliVity requires DNS for the following:

- A default DNS suffix to search for a domain. The primary suffix is used in DNS name registration and name resolution.
- A DNS suffix name for each HPE SimpliVity host.
- Forward and reverse lookup to support SSL certificate validation when logging in to an HPE SimpliVity.

DNS server best practices

- A reliable DNS infrastructure is required for VMware vCenter® to host connections and SSL certificates. Consider using a minimum of two DNS resolvers.
- Hewlett Packard Enterprise strongly recommends using the same DNS resolver (server) for the HPE SimpliVity Virtual Appliance, HPE SimpliVity host, vCenter Server, vSphere Web Client workstation, and the computer running Deployment Manager.
- Each host IP address should have both a forward lookup **A record** and a reverse lookup **PTR record** for name resolution. This should be in place before deployment to allow hosts to be added to vCenter by using hostnames rather than IP addresses.

NTP

HPE SimpliVity Deployment Manager using the U.S. internet timeservers configures the Network Time Protocol (NTP). The default servers are `0.us.pool.ntp.org` through `3.us.pool.ntp.org`. It is essential that all systems in the federation, as well as systems that communicate with the federation are synchronized with NTP.

NTP best practices

- Specify at least three NTP servers
- You might need to modify your firewall to allow NTP to work; refer to ntp.pool.org/en/
- If you decide to use an internal Domain Controller as an NTP server, see: [Synchronizing ESXi/VMware ESX® time with a Microsoft Domain Controller \(1035833\)](#)



LAYER 2 NETWORKING FEATURES

Jumbo packets/MTU

To optimize HPE SimpliVity network performance, consider the following guidelines.

- Maximum transmission unit settings
 - Storage and Federation networks must be configured with a maximum transmission unit (MTU) of 9000 bytes to enable jumbo frames. A frame (or network packet) 9000 bytes or larger is called a jumbo frame and it increases performance in storage networking environments. The compute node storage network MTU should also be set to 9000.
 - MTU setting must be configured end-to-end on all devices in the path (virtual and physical) or it will degrade performance instead of increasing it. For example, if there are multiple switch hops between HPE SimpliVity hosts in a cluster, each switch port in the path must be configured with an MTU of at least 9000.
 - The management network will likely be using a default MTU size of 1500 bytes because of the relative difficulty in ensuring jumbo frame support end-to-end for this network.
- Every physical network adapter connected to the same vSphere standard switch or VMware vSphere® Distributed Switch™ should also be connected to the same physical network.
- Configure the fewest possible number of hardware segments between the HPE SimpliVity hosts in a cluster. The goal is to limit single points of failure and reduce latency.

VLANS

VMware VLANs are not required (but are recommended). It is a common VMware best practice to assign the management network its own VLAN to separate the management traffic from the virtual machine data traffic. It is also common to create separate networks for vMotion as well as other networks, such as datastore heartbeat. Consider following these VMware recommendations in your HPE SimpliVity environment.

For security purposes, configure isolated VLANs for the HPE SimpliVity Federation before deploying the HPE SimpliVity hosts. This protects federation interfaces and data from external discovery or malicious attacks.

Figure 3 illustrates an HPE SimpliVity VLAN configuration.

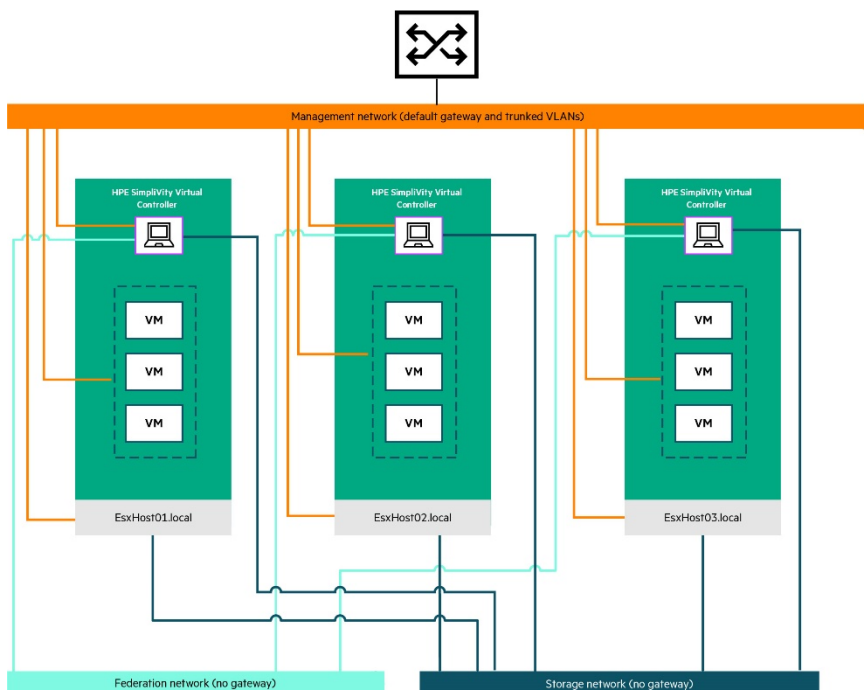


FIGURE 3. HPE SimpliVity VLAN configuration



In the diagram, the HPE SimpliVity Virtual Appliance participates in the storage, federation, and management networks. The ESXi host participates in management and storage networks. The guest virtual machines participate in the management network. The management network can include multiple VLANs and subnets for virtual machine guests.

VLAN best practices

- Configure the ports on the physical switches that an HPE SimpliVity host connects to as trunk ports. This configuration allows the hypervisor to manage the VLAN configuration.
 - You can add VLAN IDs to an allow list. This limits the VLANs allowed to communicate on the trunk ports. In this scenario, the allow list contains the VLAN IDs that you plan to use for the storage and federation networks.
 - In Deployment Manager, when specifying the network settings for the storage and federation networks, enter the VLAN IDs from the allow list in the appropriate fields. For example, you enter the Host VLAN ID and HPE SimpliVity Virtual Appliance VLAN ID for storage network settings and the HPE SimpliVity Virtual Appliance VLAN ID field for the federation network.
 - Ensure that there is network isolation for the Virtual Controller storage and federation network interfaces between clusters when more than one cluster is deployed within a federation.
-

Port settings—LAN switch and VMware

- General (all modes)
 - For best performance, use VMXNET3 virtual NICs (This is at the VM guest level others are at the vSwitch level)
 - Turn off port security
- Non-trunk mode (a single VLAN per physical interface or port channel)
 - STP mode: Edge or PortFast
 - Switch port mode: Access
 - Access VLAN number: configure to match VMware networking
- Trunk mode (multiple VLANs per physical interface or port channel)
 - STP mode: Edge Trunk or PortFast Trunk
 - Switch port mode: Trunk
 - Allowed VLAN number list: configure to match VMware networking, include VLANs of each defined port group and VMkernel port in the associated vSwitch

NIC teaming

After you deploy HPE SimpliVity hosts to a cluster, you can configure NIC teaming to increase network capacity for the virtual switch hosting the team and to provide passive failover if the hardware fails or it loses power. To use NIC teaming, you must uplink two or more adapters to a virtual switch.

For more details on configuring NIC teaming for a standard vSwitch, see the [HPE OmniStack for vSphere Administration Guide](#). You can also find details on the VMware Knowledge Base site (kb.vmware.com) and search on NIC teaming in ESXi.

For a direct connect cluster, use an active/passive configuration.

For an active/standby NIC team—the recommended settings are:

- VMware
 - Default load balancing: Route based on originating port ID
 - Failback: Yes
 - Ensure you maintain that each host has the same active NIC to switch configuration
- LAN switch settings
 - Do not configure a port-channel/link aggregation
 - Connect active and standby NICs to separate physical switches, if possible



For an active/active NIC team—the recommended settings are:

- VMware
 - If the network uses LACP/802.3ad, or EtherChannel, configure Load Balancing (in the vSwitch properties) as: Route based on IP hash.
 - Use Port ID for all other cases.
- LAN switch settings
 - If the LAN supports only standard link aggregation technology such as LACP/802.3ad, EtherChannel, and such connect all active NICs to a single switch.
 - If the LAN is using a switch stacking technology, connect each NIC to a separate switch in the stack.
 - If the LAN supports multi-chassis link aggregation (MLAG), connect each NIC to separate switches in the group.

PHYSICAL NETWORK CONNECTIVITY

HPE SimpliVity uses a shorthand notation to describe a federation's topology, for example a "2+1" federation. A 2+1 is a federation with two HPE SimpliVity hosts in the primary cluster and one HPE SimpliVity host in a secondary cluster (such as a disaster recovery site or ROBO). HPE SimpliVity hosts can be physically connected as 10GbE direct connect and 10GbE switch connect.

Direct connect clusters

Direct connect clusters can be used in a one- or two-node configurations.

Single-node clusters

In a single-node cluster, the storage and federation port groups, as well as the IP addresses for the HPE SimpliVity Virtual Appliance on the federation and storage networks are not actually in use, but should be created for the following reasons:

- The deployment process requires them
- It simplifies deploying additional HPE SimpliVity hosts to the cluster

Two-node, direct connect clusters

Best practices for a two-node cluster where two HPE SimpliVity hosts are connected back-to-back (direct connect), such as a 2+0 or 2+2:

- Configure the 1GbE ports as active/active or active/standby, these must be used for the management and virtual machine data traffic.
- Configure NIC teaming as active/standby on the 10GbE ports, one 10GbE vmnic must be an Active Adapter, while the other 10GbE vmnic is a standby adapter. The 10GbE ports are used exclusively for storage and federation traffic. Data unavailability will occur on failover in active/active direct connect environments.

Configure the following NIC teaming policy exceptions:

- Load balancing: Route based on the originating virtual Port ID
- Network failover detection: Enable and select Link status only
- Notify switches: Enable and select Yes
- Failback: Enable and select Yes

To extend this configuration beyond two HPE SimpliVity hosts requires that you reconnect the hosts through a network switch.

Switch connected clusters

Switch connected clusters are clusters with more than one HPE SimpliVity host that are connected by a 10GbE switch/LAN.

Switch connected cluster best practices

- If using all 10GbE, you may configure all NICs as a NIC team, but you must at a minimum define separate VLANs for the management, storage, and federation networks.
- You may configure vmnics in active/active or active/standby mode.
- You may configure all three HPE SimpliVity VLANs on the same 10GbE trunk port—management, storage, federation, or any additional VLANs used by VMware or other virtual machines. This must be configured after deployment is complete.



NETWORK TOPOLOGIES

HPE SimpliVity supports both full mesh and remote office/branch office (ROBO) network topologies.

Full mesh

With a full mesh topology, every HPE SimpliVity host in every cluster in the federation can communicate directly with every other host in the federation. This configuration enables backup and restore operations to be performed between each cluster in the federation.

Remote office/Remote branch

A ROBO topology supports direct communication between one or more clusters in a hub (for example, clusters with HPE SimpliVity hosts at the company headquarters) and remote clusters (for example, data centers with other HPE SimpliVity hosts) connected via a wide area network (WAN) or a local area network (LAN). In this configuration, all the clusters can communicate directly with the main, centralized clusters, but the remote spoke clusters do not have direct communication with each other. This configuration is also known as a hub-and-spoke topology.

Data movement operations such as moving a virtual machine, restoring a virtual machine or file from a backup, or copying a backup to another cluster are restricted to directly connected clusters. So, in a hub-and-spoke topology, these operations cannot be performed between spoke clusters that are not directly connected.

TABLE 5. Network traffic in a ROBO environment

Type of traffic	Uses	HPE SimpliVity network interface
HA traffic	User data exchange between hosts within a cluster as part of preserving data availability	Federation network interface (10 Gbps)
Backup traffic (WAN)	User data for backups within or between clusters	Preferred: Federation network interface Secondary: Management network interface
Metadata management protocols (WAN)	Metadata sync between all HPE SimpliVity hosts in the federation	Preferred: Federation network interface Secondary: Management network interface

ROBO best practice

In a hub-and-spoke configuration, set up firewall rules to prevent the spoke clusters from communicating with one another.

HPE SIMPLIVITY SUPPORT FOR NSX-T

VMware NSX® provides network by enabling the creation of virtual networks on virtualization the physical network layer. This complements VMware Software-Defined Data Center by providing network virtualization in the network layer beyond the data center infrastructure. Similar to how virtual machines (VMs) are created by abstracting the physical hardware layer, NSX provides the creation of a virtual network by abstracting the physical networking layer.

NSX-T is the next level of network virtualization offered by VMware for applications running on VMs, containers, cloud, bare-metal servers, or any hypervisors. NSX-T provides software-defined network infrastructure that is decoupled from the VMware vSphere environment thus providing support for non-VMware platforms also.

Networking changes with NSX-T

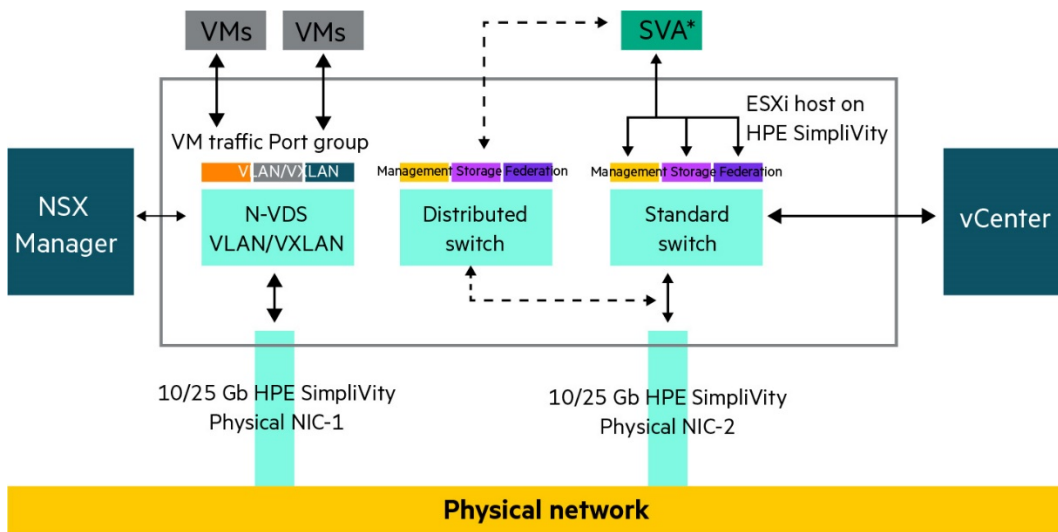
N-VDS like Virtual Distributed Switch requires a minimum of two or more physical NICs in the production environment. N-VDS can only be configured and managed within NSX-T Manager and not from the vCenter as in the case of vDS or vSS. This decoupling from vCenter helps the architects to design the use of N-VDS outside of the vSphere environment.

In NSX-T, transport zone can span endpoints across different platforms or multiple vSphere vCenter hence N-VDS is correlated to a transport zone and created when the transport zone is created. NSX Switches (N-VDS) of the different transport nodes are independent but can be grouped by assigning the same name for centralized management.

VMware NSX-T and HPE SimpliVity

HPE SimpliVity by default communicates with vCenter and other HPE SimpliVity nodes through virtual Distributed vSwitch (vDS) or virtual Standard vSwitch (vSS).





*SVA = HPE SimpliVity Virtual Appliance also known as OVC

FIGURE 4. Logical overview of vSphere switching components, where NSX-T fits and how the HPE SimpliVity platform interacts with them

HPE SimpliVity and NSX-T supported features

- Support for ESXi hypervisors only
- Support for VM traffic only
- The minimum supported version of NSX-T is 3.0.1
- Both VXLAN and VLAN are supported for N-VDS
- NSX-T is supported on HPE SimpliVity version 4.0.1U1 and above
- Support for vSphere 6.7 and above

NOTE

HPE SimpliVity Virtual Appliance traffic is not supported on NSX-T, SVA still has to run on Standard switches or Distributed switches managed by the vCenter.

Network configuration used for qualification

- Management, Federation, Storage networks (SVA traffic) are configured on vSS port groups
- VM traffic is configured on N-VDS (NSX managed Virtual Distributed Switch) port group

Please follow the best practices from VMware for placement of NSX-T manager.



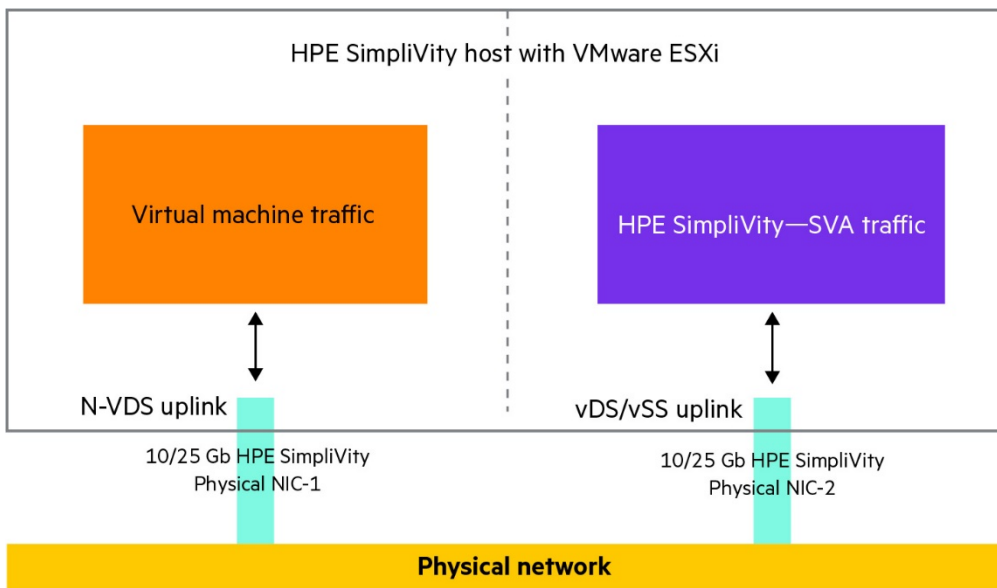


FIGURE 5. SVA and VM traffic logical representation

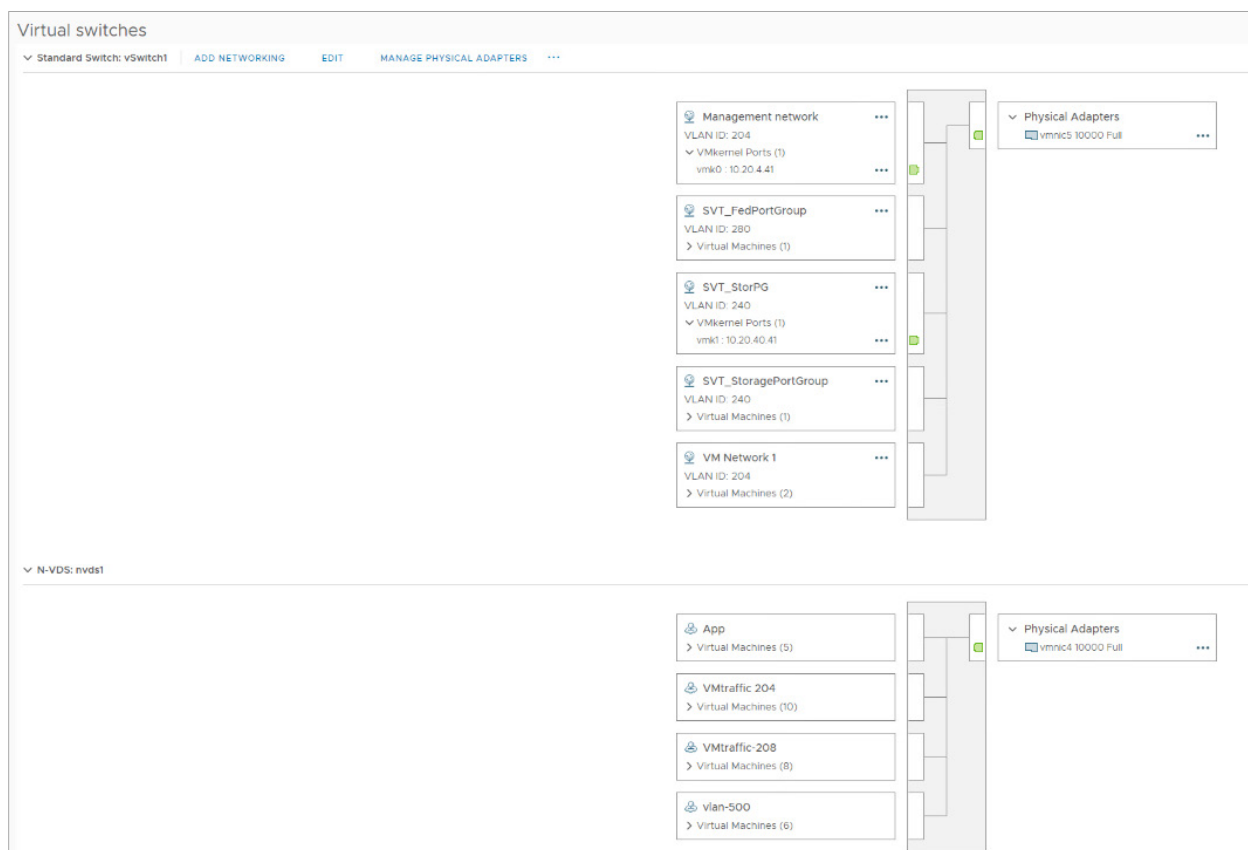


FIGURE 6. Sample configuration with NSX-T, N-VDS, and Standard switch split between two physical NICs in HPE SimpliVity host



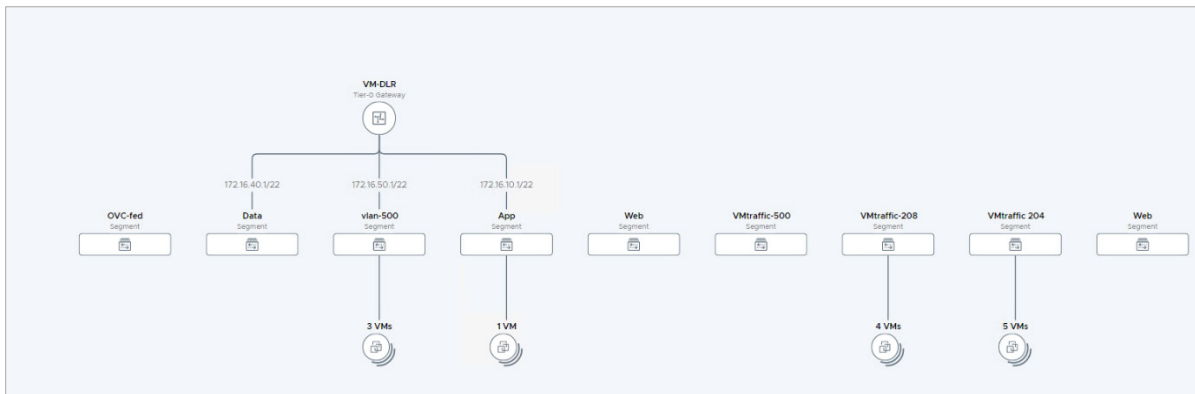


FIGURE 7. Sample topology where VXLAN based and VLAN backed network co-exist on N-VDS Switch

NOTE

1. VMtraffic 204, VMtraffic-208, are VLAN backed port groups
2. Data, vlan-500, and App are VXLAN backed port groups

HPE SimpliVity and NSX-T Caveats

After a restored VM is powered on (HPE SimpliVity backup and restore), all the network adapters of the VM (which are connected to the NSX-T port groups) are in a “Disconnected” state.

Description:

After VMs assigned with NSX-T switches are restored from backup and powered on, the network adapters are in a “Disconnected” state.

Trying to set the network adapters to a “Connected” state through vCenter UI and PowerCLI command also fails.

This behavior is explicit to HPE SimpliVity “Rapid Cloned” VMs whereas VMs cloned through vCenter operation remain in the “Connected” state.

Workaround:

1. After VMs are restored, delete all the network adapters, recreate the network adapters and assign them to the appropriate NSX-T switches.
2. Another workaround is to update the .vmx file of the restored VM with the UUID in the external ID field for each adapter entry.

STRETCH CLUSTERS

The stretch cluster feature ensures service continues in the presence of fault conditions that can impact multiple HPE SimpliVity hosts that are co-located in a rack, lab, or site. This is provided with the use of availability zones, which are collections of hosts that share a common fault-mode. As an administrator, you define these zones to tell the HPE SimpliVity software the nodes that are likely to fail together due to external forces. By configuring an equal number of hosts within two availability zones, and leveraging a third site arbiter, you will not lose access to data following a failure that impacts one of these zones.

Stretch clusters ensure fully committed synchronous writes between two physical sites by only acknowledging writes back to the VM after nodes at both sites have safely persisted the write. This allows the use of VMware HA functionality to automate the recovery of virtual machines after the failure of an entire site, providing RPOs of zero and RTOs of seconds. To ensure this functionality and avoid split-brain scenarios, HPE recommends the HPE SimpliVity Arbiter be deployed at a third site to avoid losing it along with one-half of the HPE SimpliVity cluster. In a stretch cluster environment, you should deploy the Arbiter in a location that is outside the failure domain of both data centers.

Stretch cluster best practices

- Place the HPE SimpliVity Arbiter in a third physical site.
- Equally distribute the HPE SimpliVity hosts in the data center into the two availability zones.
- HPE SimpliVity clusters typically exist within a single physical site but can be extended across two physical sites when the two sites are connected by a high bandwidth (10 Gbps), low-latency interconnect, generally <1 ms. Higher latencies can be deployed depending on application performance requirements.
- Optimal latency for storage networks (stretch cluster) RTT is <2 ms.



REMOTE BACKUPS

Backup replication to a remote cluster requires a number of steps and a varying amount of bandwidth depending on how much of the backup data is already located at the target. WAN link congestion can adversely affect the number of backups that might be replicated to a remote cluster.

If replication fails to complete within the backup window or retention period, check if there is sufficient network throughput. Insufficient throughput may be resolved with a variety of approaches:

- Reduce the frequency, volume, or timing of backups to the relevant cluster
- Provide additional bandwidth
- Reduce latency by identifying and eliminating the source of latency
 - If latency is caused by congestion (WAN or LAN), implement QoS on saturated links in the path or increase bandwidth
 - If latency is caused by distance or link media (such as satellite, wireless), implement TCP acceleration to increase throughput potential

BANDWIDTH THROTTLING FOR HPE SIMPLIVITY BACKUP AND RESTORE OPERATIONS

Backup operations, in general, can consume a significant portion of the WAN bandwidth. This can impact the performance of applications that are sharing the WAN bandwidth.

HPE SimpliVity 4.1.0 now provides a solution that allows customers to control the amount of bandwidth that is consumed by HPE SimpliVity backup and restore operations. This allows other traffic on the network to traverse with much lesser latency.

This feature will benefit customers who have HPE SimpliVity at remote locations, edge sites or any site that has limited bandwidth. Customers can proactively control the backup bandwidth allocated for HPE SimpliVity without impacting other applications.

NOTE

It is recommended to contact support for enabling bandwidth throttling for HPE SimpliVity backup and restore operations.

The feature needs to be enabled on each of the nodes participating in backup and restore operations within a cluster. These changes are at the operation level. They are not at the HPE SimpliVity Virtual Appliance (SVA) network level, i.e., the changes are not applied to management, storage, or federation vNICs of SVA. The bandwidth is set in Kbps. Once the bandwidth throttling is set up it persists across reboots. If bandwidth throttling is not required, it has to be exclusively disabled.

NOTE

Know your environment well before throttling the bandwidth, it shouldn't be done just because it can be done.

The following values and operations could be impacted by throttling the bandwidth for HPE SimpliVity backup and restore operations:

1. Backup and restore SLAs
2. RPOs and RTOs
3. HPE SimpliVity RapidDR performance and DR operations



REMOTE VCENTER MANAGEMENT

High latency may cause certain vCenter tasks to take longer than usual to show in the tasks pane or appear to time out but actually complete.

Remote vCenter management best practice

Network latency between HPE SimpliVity hosts and vCenter must be under 300 ms.

HPE SIMPLIVITY ARBITER

The HPE SimpliVity Arbiter is a Windows service that facilitates communication between HPE SimpliVity hosts in a federation and enables failover and recovery operations.

The Arbiter ensures data availability and integrity by serving as an additional quorum member in two-node and stretch cluster configurations. This arbitration ensures the resiliency of the federation. The Arbiter must be installed and running on a computer that the HPE SimpliVity hosts can access.

The Arbiter can manage up to 4,000 virtual machines. To ensure best performance, install an Arbiter for every 4,000 virtual machines and associate it with one or more clusters in a federation to distribute the workload. In stretch clusters, it is important that the arbiter is installed outside of the failure domain of the hosts that it is arbitrating. For example, suppose that you have deployed HPE SimpliVity hosts in two clusters (cluster 1 and cluster 2); you can deploy the Arbiter on a Windows host that has network access to both clusters. Figure 8 is an example:

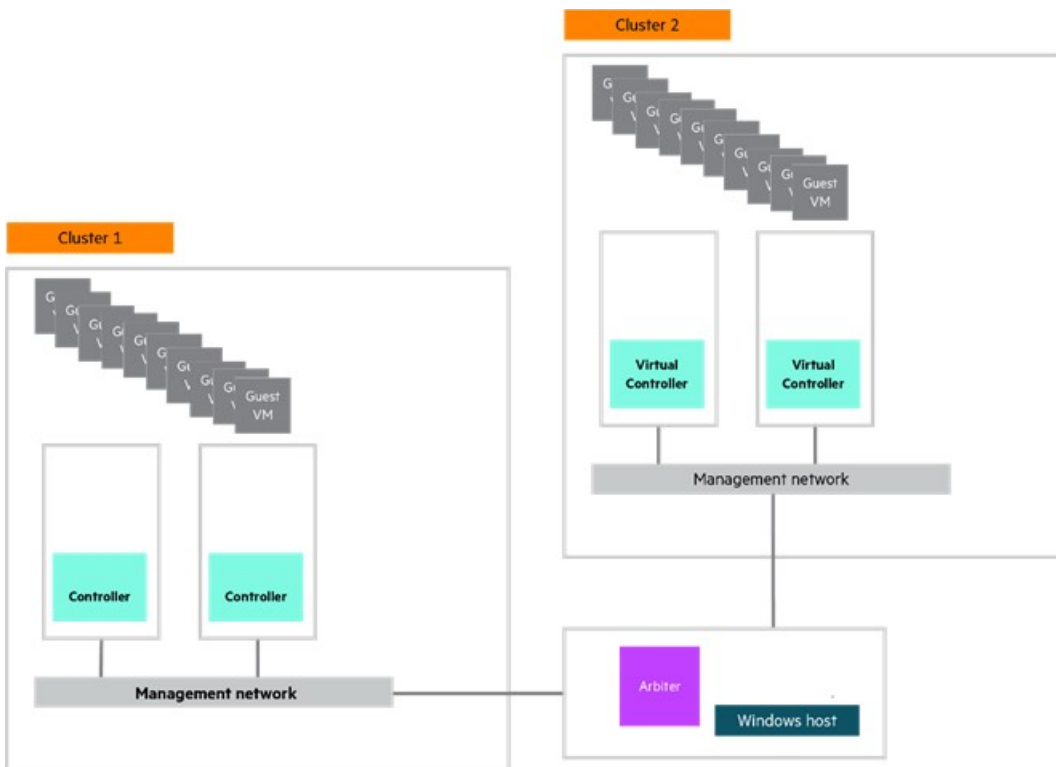


FIGURE 8. HPE SimpliVity Arbiter on a Windows host that has network access to both clusters

Alternatively, you can deploy an Arbiter within each cluster that manages the other cluster. In that configuration, you would not need a separate Windows host. In a two-node cluster scenario, if the Arbiter were to be deployed as a virtual machine within the same cluster that it is supporting, a complete data unavailability would occur if its ESXi server were to go down. This failure would prohibit the surviving node from establishing quorum. For this reason, the Arbiter should be installed on a separate host outside of the cluster that it supports.

In a multi-data center environment, you can install the Arbiter on a Windows virtual machine in one data center that arbitrates the hosts in the other data center and vice versa. Figure 9 is an example.



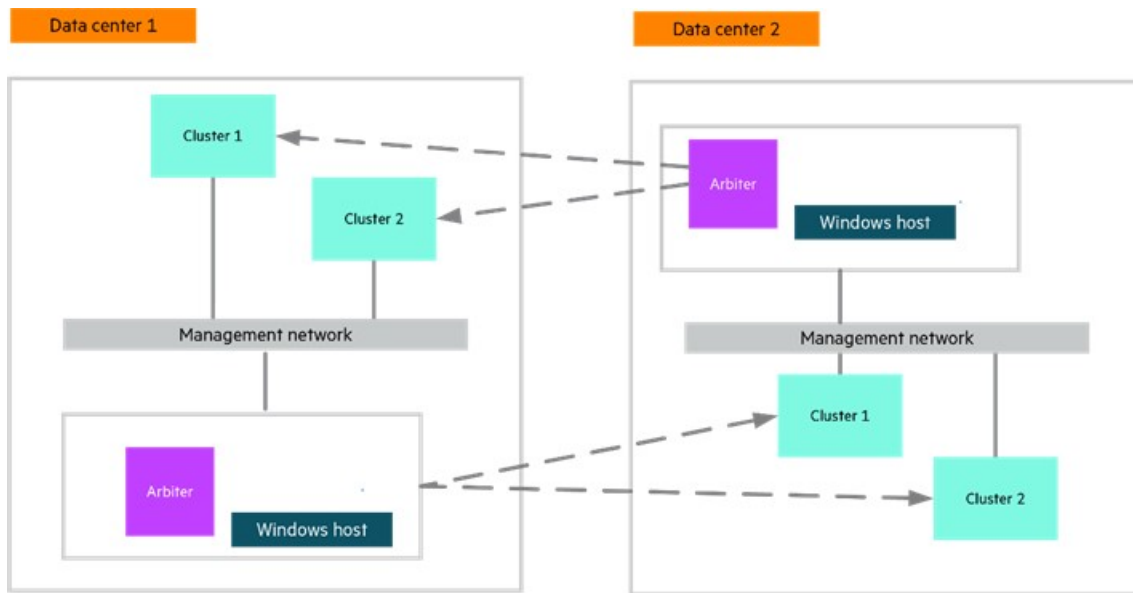


FIGURE 9. HPE SimpliVity Arbitrator in different data centers

Arbitrator best practices

- For HPE SimpliVity 4.0 and later:
 - 2+0 Arbitrator RTT is 300 ms
 - 4+0 non-stretch clusters RTT is 300 ms
 - For stretch clusters max RTT is 50 ms
 - Non-stretch clusters with more than four nodes RTT is 50 ms
 - The latency between HPE SimpliVity hosts remains 2 ms
- The network connection to the Arbitrator should be of high quality with low-packet loss as communication to this component is critical to preserving data availability in the presence of failures.
- Do not recover the Arbitrator storage device to a previous point in time with a snapshot or backup and restore process or you could corrupt the virtual machine data.
- Do not install Arbitrator on a virtual machine managed within the same cluster it serves. If you do install the Arbitrator in this configuration, data unavailability can occur.
- Never restart the Arbitrator for any reason other than resolving problems. Your federation cannot communicate properly when the Arbitrator is not running.
- You can upgrade an arbitrator or move it to another location only when hosts in the federation are healthy. If the federation is not healthy (quorum is lost) do not upgrade or move the arbitrator.

ARUBA SWITCHES

Aruba switch best practices

Minimum requirements:

- Use two supported Aruba switches when deploying hosts. You cannot use a single switch.
- You can only deploy hosts using Aruba switches to a cluster that does not contain previously deployed HPE SimpliVity hosts.

Before deployment:

- Enable all switch interfaces facing HPE SimpliVity hosts. For example, allow multicast, SLP service, and UDP service. (These interfaces come disabled by default.)
- Enable the switch management interface and configure it as needed.
- Ensure that LLDP is enabled.
- Ensure that HTTPS is enabled on a management VRF (virtual routing tables and forwarding).
- Check that HTTPS REST access mode is set to read-write.
- Create a user account with administrator privileges to use during deployment (when prompted by deployment manager).
- HPE SimpliVity uses the default network of 172.16.0.1/24 for the storage and federation network. It is recommended that you not change this range.



HPE COMPOSABLE CLOUD FOR HPE PROLIANT DL

HPE Composable Cloud for HPE ProLiant DL enables you to deploy your hosts to the hypervisor management system with deployment manager and monitor the server infrastructure through HPE OneView. You can then manage the hosts in vCenter Server and use HPE SimpliVity Plug-in features on HPE SimpliVity objects. You can also use HPE OneView to monitor the status of the HPE SimpliVity hardware and networking configuration. This streamlines the ability to monitor the platform, adapter layout, network settings, BIOS settings, firmware, and networks in your data center.

HPE Composable Cloud for HPE ProLiant DL best practices

- All HPE SimpliVity hosts must be booted into deploy installer.
- All HPE SimpliVity hosts must have an initial server profile created through HPE OneView to provide initial access to the management network for discovery.
- Ensure the vCenter Server cluster name where HPE SimpliVity is deployed is unique across all data centers and vCenter Server instances.

SWITCH REQUIREMENTS

HPE SimpliVity does not maintain a specific list of supported LAN switch platforms, but the 10GbE switches must have these minimum capabilities:

- Wire rate 10GbE ports
- Jumbo frames capability
- VLAN capability

HPE SIMPLIVITY NETWORK CHECKLIST

- Reliable DNS infrastructure is available for the HPE SimpliVity federation
- Dependable NTP infrastructure is available for the HPE SimpliVity federation
- Unique subnets are defined for each cluster's federation and storage networks
- Non-routable subnets used for federation and storage networks
- Trunked VLANs should be used on the management network interfaces
- Isolated VLAN IDs defined for each cluster's management, federation, and storage networks
- Consistent MTU sizes defined and configured for each cluster's management, federation, and storage networks
- Use standard switches for deployment and then convert to distributed switches after the HPE SimpliVity host has been deployed

LEARN MORE AT

hpe.com/us/en/integrated-systems/simplivity.html

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

© Copyright 2020–2021 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware ESX, VMware ESXi, VMware vCenter Server, VMware vCenter, VMware vSphere, VMware vSphere Distributed Switch, VMware vSphere Storage vMotion, VMware vSphere Web Client, VMware, and VMware NSX are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.