



Technical Report

NetApp HCI Networking Quick Planning Guide for an On-Premises Hybrid Cloud Infrastructure

Mike Geroche, Technical Marketing Engineer, NetApp
March 2020 | TR-4820

TABLE OF CONTENTS

| | | |
|----------|---|----------|
| 1 | Purpose | 3 |
| 2 | NetApp HCI Network Topology | 3 |
| 3 | General Network Requirements for NetApp HCI | 3 |
| 3.1 | Architecture | 3 |
| 3.2 | Connectivity | 4 |
| 3.3 | Configuration and Protocol | 4 |
| 3.4 | Management and Services | 4 |
| 3.5 | Cluster Uplink | 4 |
| 4 | Overview of Networking in NetApp HCI | 4 |
| 4.1 | Management Network | 4 |
| 4.2 | Storage Network | 4 |
| 4.3 | Virtual Machine Network | 4 |
| 5 | NetApp Deployment Engine Network Prerequisites | 5 |
| 5.1 | NDE Networking Information and General Requirements | 5 |
| 5.2 | Compute Node Networking | 5 |
| 5.3 | Storage Node Networking | 7 |
| 6 | Acknowledgments | 7 |
| 7 | Where to Find Additional Information | 7 |

LIST OF FIGURES

| | | |
|-----------|--|---|
| Figure 1) | Simplified network topology: NetApp HCI with NetApp H-Series 1RU compute and storage nodes. | 3 |
| Figure 2) | NetApp H615C 1RU compute node networking, two cables (option A); all storage, VM, migration, and cluster management traffic is segmented with VLANs. | 5 |
| Figure 3) | NetApp H410C 2RU 4-node chassis, compute node networking, two cables (option A); all storage, VM, migration, and cluster management traffic is segmented with VLANs. | 6 |
| Figure 4) | H410C 2RU 4-node chassis, compute node networking with six cables (option B and option C); six-cable option is available only on 2RU chassis compute nodes. | 6 |
| Figure 5) | H610S 1RU storage node networking connectivity; you can optionally assign the physical access port for OOB node management access to a management network port. | 7 |
| Figure 6) | H410S 2RU 4-node chassis, storage networking connectivity; you can optionally assign the physical access port for OOB node management access to a management network port. | 7 |

1 Purpose

NetApp® HCI is an enterprise-scale hybrid cloud infrastructure design that combines storage, compute, and networking. It brings public cloud services to your on-premises private cloud built on a hybrid cloud infrastructure. NetApp HCI delivers an agile, scalable, easy-to-manage solution with guaranteed application performance, mature integrated replication, non-disruptive operations, data protection, and high availability. With a hybrid cloud infrastructure, NetApp HCI delivers a user experience that transcends location: A private cloud essentially becomes another resource, just like the public clouds.

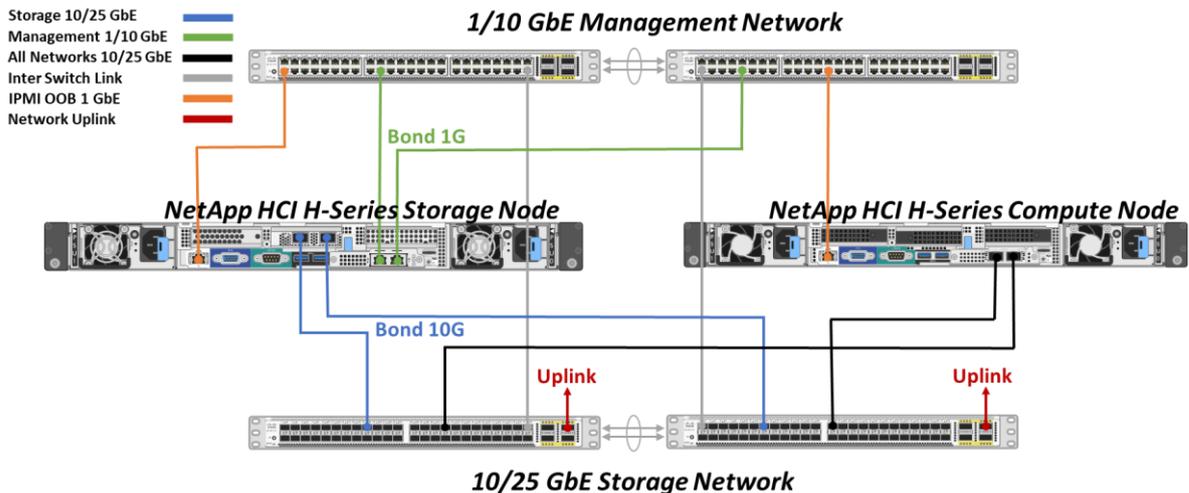
As with the network considerations for any enterprise IT infrastructure, the network considerations for NetApp HCI focus on availability, performance, and extensibility. This guide provides the networking concepts and guidelines to prepare your network for successful deployment of NetApp HCI. NetApp recommends that you implement these recommendations before installing NetApp HCI. Adherence to these guidelines will aid in the installation and overall experience.

2 NetApp HCI Network Topology

NetApp HCI is deployed in clusters that consist of specialized compute and storage resources, redundantly interconnected to top-of-rack network switches. By physically disaggregating the compute and storage nodes, the hybrid cloud infrastructure can independently scale out to meet the needs of the workload. Figure 1 shows a simplified network topology for a single storage node and single compute node.

Note: The minimum cluster size for NetApp HCI is four storage and two compute nodes.

Figure 1) Simplified network topology: NetApp HCI with NetApp H-Series 1RU compute and storage nodes.



3 General Network Requirements for NetApp HCI

The following are the prerequisites of the network infrastructure for NetApp HCI. You must address these requirements before deploying NetApp HCI.

3.1 Architecture

- A fully redundant, fault-tolerant data center–class switch network deployed in redundant top-of-rack (ToR) pairs or in a highly resilient switch chassis
- Nonblocking switch architecture capable of line-rate bidirectional switching on all ports simultaneously
- Redundant network connectivity for all storage, virtual machine (VM) migration, and cluster management traffic

3.2 Connectivity

- High-speed 10GbE or 25GbE interconnects for storage, VMs, and VM migration.
- 1/10GbE connectivity for management traffic.
- Jumbo-frame support for storage, VMs, and VM migration traffic.
- Adequate availability of IP addresses to accommodate NetApp HCI cluster size. The minimum cluster size is 4 storage and 2 compute nodes. NetApp HCI consumes 3 IP addresses per storage node and 5 IP addresses per compute node, and 6 additional IP addresses. Therefore, the minimum configuration (4 storage, 2 compute) consumes 28 IP addresses.

3.3 Configuration and Protocol

- Ability to logically separate management, VM, and storage traffic types by using virtual LANs (VLANs) on different subnets
- Use of Spanning Tree Protocol on all host ports, configured to immediately enter a forwarding state
- Multichassis link aggregation support to ease enablement of Link Aggregation Control Protocol (LACP) on the storage network

3.4 Management and Services

- DNS and Network Time Protocol (NTP) services accessible from the management VLAN
- Cluster management node gateway access to the internet through firewall port or proxy server; for port details, see [NetApp HCI documentation](#)
- Intelligent Platform Management Interface (IPMI) out-of-band (OOB) management of the cluster nodes at 1GbE

3.5 Cluster Uplink

- A minimum of one pair of uplinks (one from each switch) configured as trunks
- Uplinks configured in a virtual port channel or multichassis link aggregation group
- High-speed 10/25/100GbE connectivity to a network spine switch
- Jumbo-frame support preferred

4 Overview of Networking in NetApp HCI

The network segments in NetApp HCI can be categorized as follows.

4.1 Management Network

- VMware vSphere Management
- Existing infrastructure services (NTP, DNS, existing VMware vCenter)
- IPMI for OOB management of the compute and storage system

4.2 Storage Network

- iSCSI traffic
- Storage cluster traffic

4.3 Virtual Machine Network

- VMware vMotion for ESXi hosts
- Guest VM port groups; these VMs can be added post-deployment as the system administrator sees fit

5 NetApp Deployment Engine Network Prerequisites

The NetApp Deployment Engine (NDE) is a management tool that orchestrates the initial deployment and configuration of the NetApp HCI components. The NDE is used to form independent compute and storage nodes into a NetApp HCI system. It is responsible for collecting all user input that is required for completely installing and configuring the system. The NDE also greatly simplifies configuration and deployment of the VMware vSphere and NetApp HCI storage components.

5.1 NDE Networking Information and General Requirements

- Bond1G is a logical interface that combines 1GbE network ports on storage and compute nodes. This network is used for NDE API traffic. All nodes must be able to communicate over the management interface in the same L2 network.
- Bond10G is a logical interface that combines 10/25GbE ports and is used by the NDE for beaconing and inventory. All nodes must be able to communicate over the Bond10G interface with non-fragmented jumbo frames.
- The NDE requires a minimum of one manually assigned IP address on the Bond1G interface on one storage node. The NDE will be run from this node.
- All nodes will have temporary IP addresses assigned by NDE discovery, which is accomplished by Automatic Private IP Addressing (APIPA).
- During the NDE process, all nodes will be assigned permanent IP addresses, and any APIPA-assigned temporary IP addresses will be released.
- Before running the NDE, you must preconfigure separate VLAN networks for management, iSCSI, and vMotion on the switch network.

5.2 Compute Node Networking

Option A: Two Cables for Compute Nodes Using VLAN Tagging (All Compute Node Models)

This mode uses two network cables on the compute nodes for connectivity to all NetApp HCI networks.

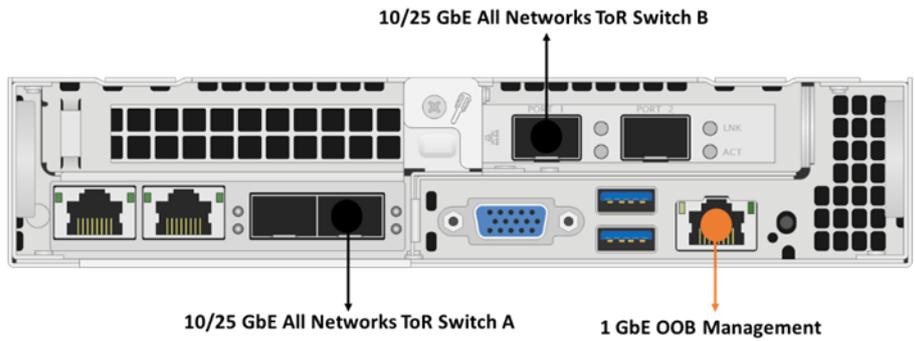
- It consolidates all traffic onto two physical, redundant ports, reducing the cabling and streamlining network configuration.
- This configuration requires that the storage, vMotion, and VM network segments (if any) use VLAN tagging.
- The management network segment can use native or tagged VLAN. However, native VLAN is the preferred mode so that the NDE can assign network resources in an automated manner (using zero-configuration networking, or zeroconf).

Note: This mode requires use of vSphere Distributed Switch, which requires VMware vSphere Enterprise Plus licensing.

Figure 2) NetApp H615C 1RU compute node networking, two cables (option A); all storage, VM, migration, and cluster management traffic is segmented with VLANs.



Figure 3) NetApp H410C 2RU 4-node chassis, compute node networking, two cables (option A); all storage, VM, migration, and cluster management traffic is segmented with VLANs.



Option B: Six Cables for Compute Nodes Using Tagged VLANs (H410C, 2RU 4-Node Compute Node)

This mode uses six network cables for connectivity to all NetApp HCI networks.

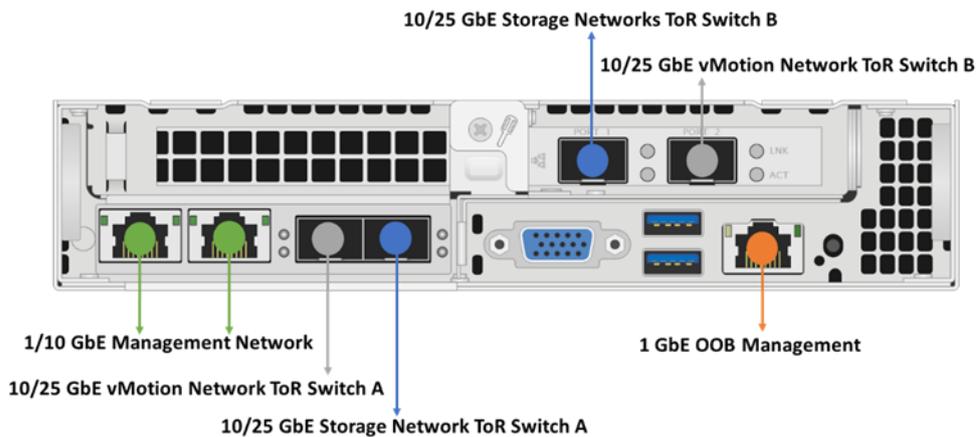
- This option is for use only on 2RU compute nodes; it is not supported for 1RU compute node models.
- This configuration requires that the storage, vMotion, and VM networks use VLAN tagging.
- It provides dedicated bandwidth for data storage traffic, vMotion traffic, and VM network traffic.
- This configuration uses vSphere Standard Switch as the default. Optional use of vSphere Distributed Switch requires VMware Enterprise Plus licensing.

Option C: Six Cables for Compute Nodes Using Native and Tagged VLANs (H410C, 2RU 4-Node Compute Node)

This mode uses six network cables for connectivity to all NetApp HCI networks.

- This configuration is for use only on 2RU compute nodes. It is not supported for 1RU compute node models.
- This configuration uses native VLANs for storage, vMotion, and management. The switch configuration separates network segments.
- VM traffic can use native or tagged VLANs
- This configuration uses dedicated bandwidth for data storage, vMotion, and VM networks.
- It uses vSphere Standard Switch as the default. Optional use of vSphere Distributed Switch requires VMware Enterprise Plus licensing.

Figure 4) H410C 2RU 4-node chassis, compute node networking with six cables (option B and option C); six-cable option is available only on 2RU chassis compute nodes.



5.3 Storage Node Networking

Each storage node must be connected to the 10/25GbE data network through the SFP+/SFP28 interfaces and to the management network through the 1/10GbE RJ-45 interfaces. Additionally, NetApp strongly recommends a connection through 1GbE RJ-45 to the IPMI for OOB management of the storage node.

- The iSCSI storage network interface is named Bond10G within the storage node.
- Configure LACP for the Bond10G interfaces on storage nodes. On storage nodes, you do this at the terminal user interface (TUI).

Note: Also, make sure that the switch network is preconfigured with LACP on the storage iSCSI ports.

- The management network is named Bond1G, with no LACP configured for this port bond.

Figure 5 shows a NetApp H610S 1RU storage node, and Figure 6 shows a NetApp H410S storage node in a 2RU configuration.

Figure 5) H610S 1RU storage node networking connectivity; you can optionally assign the physical access port for OOB node management access to a management network port.

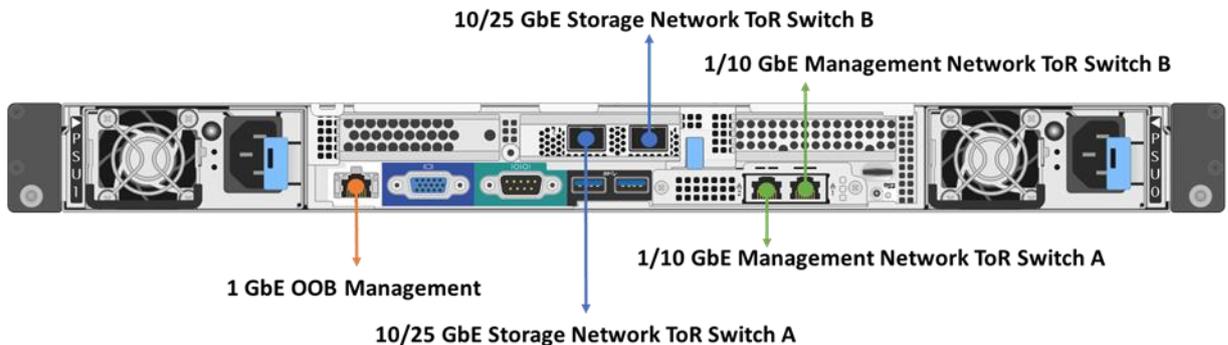
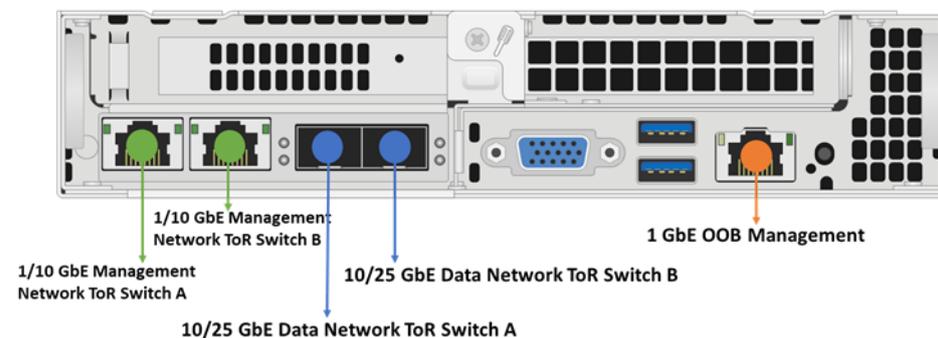


Figure 6) H410S 2RU 4-node chassis, storage networking connectivity; you can optionally assign the physical access port for OOB node management access to a management network port.



6 Acknowledgments

Special thanks to Stephen Carl, Kevin Papreck, Travis Stone, and Michael Sabieri for providing content used in this publication.

7 Where to Find Additional Information

This document provides overview information about setting up the networking for NetApp HCI deployment, including general network infrastructure and setup for the NDE.

For more details about setting up the network for NetApp HCI before you run the NDE, and to learn more about the information that is described in this document, see the following:

- [Setting Up and Using NetApp HCI 1.7P1](#)
- [NetApp Theory of Operations WP-7261](#)
- [NetApp HCI Documentation](#)
- [NetApp HCI Resources](#)

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2020 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4820-0320