



Technical Report

NetApp HCI with Mellanox SN2010 Switch Quick Cabling Guide

Stephen Carl, HCI Tech Solutions, NetApp
December 2018 | TR-4735-1218

TABLE OF CONTENTS

1	Introduction	4
2	NetApp HCI Hardware	4
2.1	Node and Chassis Layout.....	4
2.2	Node Types.....	5
3	Information About NetApp HCI Physical Ports	6
3.1	Flex cabling option for Compute Nodes (Two-Cable)	6
3.2	Connect and Deploy with Six Cables (Optional)	7
3.3	Connect Six Cables and Deploy with Two Cables	8
3.4	Networking.....	8
4	Switch Information	9
4.1	Deploying the Mellanox SN2010 Switch	10
5	Mellanox SN2010 Switch Kit Components and Cabling Diagrams	10
5.1	Sample Two-Cable Compute Node Configuration	14
6	Additional Information	19
7	Acknowledgements	19

LIST OF TABLES

Table 1)	Networking specifications for compute nodes.	5
Table 2)	Networking specifications for storage nodes	5
Table 3)	Node interface requirements.	8
Table 4)	VLAN example used in this document.....	14

LIST OF FIGURES

Figure 1)	NetApp HCI chassis.....	4
Figure 2)	Compute node physical ports.....	5
Figure 3)	Storage nodes physical ports.....	6
Figure 4)	Two-cable compute option.	7
Figure 5)	Six-cable compute option.....	7
Figure 6)	Connect six-cables use two-cable compute option.	8
Figure 7)	Mellanox SN2010 port types overview.	9
Figure 8)	Mellanox SN2010 port numbering.....	10
Figure 9)	Kit cables.	11
Figure 10)	Two-cable setup: two compute nodes and four storage nodes.	12
Figure 11)	Two-cable setup: three compute nodes and four storage nodes.....	13
Figure 12)	Two-cable setup: four compute nodes and four storage nodes.....	14
Figure 13)	Example CLI commands for switch A.....	15

Figure 14) Example CLI commands for switch B..... 17

Figure 15) Port information for example configuration CLI commands..... 19

1 Introduction

This document describes the steps that you must follow for the Mellanox SN2010 Ethernet switches before you configure networking for the NetApp® HCI system. This information includes network cabling, network switch configuration, and other network resources for the Mellanox SN2010 switch. This document assumes that:

- You have completed the NetApp HCI prerequisites checklist.
- You have racked all NetApp HCI chassis and installed all NetApp HCI nodes.
- You have not yet deployed the NetApp HCI system software by using the NetApp Deployment Engine (NDE).
- You have racked the Mellanox SN2010 switches and have the NetApp HCI Mellanox networking kit cables.

For information about system prerequisites or about your NetApp system in general, see [NetApp Documentation: Product Library NetApp Documentation: Product A-Z](#).

For more information about Mellanox Ethernet switches, see [Mellanox Scale-Out SN2000 Ethernet Switch Series](#). This webpage includes links to the SN2010 information and documentation.

2 NetApp HCI Hardware

The NetApp HCI system comes in a 2U chassis that contains up to four independent nodes. In a single chassis, you can install the compute nodes and the storage nodes in any order.

The smallest configuration includes two chassis that contain the following components:

- Two compute nodes
- Four storage nodes

The two-compute-node and four-storage-node NetApp HCI system is the primary focus of this document. Other configuration cabling diagrams are referenced in this document as. See [Setting Up and Using NetApp HCI 1.4](#) for more information.

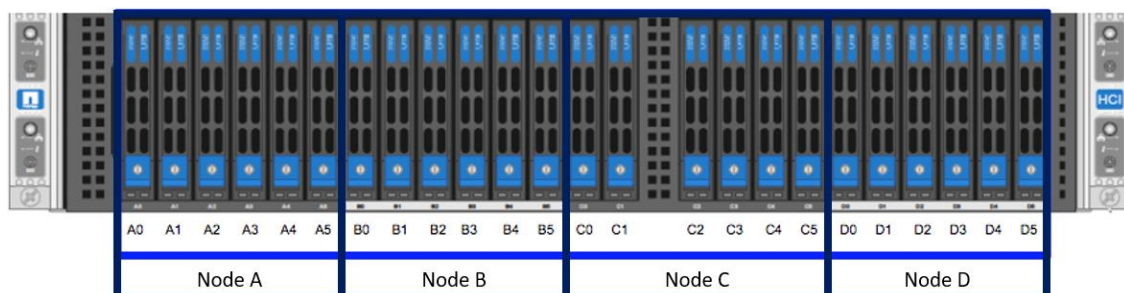
2.1 Node and Chassis Layout

You can load compute nodes into any available slot in the NetApp HCI chassis. Storage nodes require access to the hot-swappable 2.5-inch solid-state drives (SSDs) in the front of the chassis. Because of this restriction, you must place a storage node in a chassis slot that corresponds to its available drives.

The SSDs for the storage nodes are arranged in groups of six across the front of the enclosure. Each storage node has a fixed set of six SSDs connected to it. For example, in Figure 1, the drives for node A are labeled A0 to A5, the drives for node B are labeled B0 to B5, and so on.

The mapping of nodes to drives is fixed. If you move a node from one slot to another, you must also move the corresponding drives in the front of the chassis.

Figure 1) NetApp HCI chassis.



2.2 Node Types

NetApp HCI is available with a range of configuration options for both compute and storage. A minimum starting configuration must have four storage nodes and two compute nodes.

Compute and Storage Nodes

Compute and Storage nodes are available in multiple configurations. You can mix and match compute node sizes as you need to in the NetApp HCI system. Table 1 and Table 2 describes the networking specifications for example configurations (additional configurations for other node models are similar).

Table 1) Networking specifications for compute nodes.

Compute Nodes	H300E (Small)	H500E (Medium)	H700E (Large)
Base networking	4 x 10/25GbE (SFP28) 2 x 1GbE (RJ-45)	4 x 10/25GbE (SFP28) 2 x 1GbE (RJ-45)	4 x 10/25GbE (SFP28) 2 x 1GbE (RJ-45)

Table 2) Networking specifications for storage nodes

Storage Nodes	H300S (Small)	H500S (Medium)	H700S (Large)
Base networking	4 x 10/25GbE (SFP28) 2 x 1GbE (RJ-45)	4 x 10/25GbE (SFP28) 2 x 1GbE (RJ-45)	4 x 10/25GbE (SFP28) 2 x 1GbE (RJ-45)

Detailed Physical NIC Layout

Figure 2 shows the physical ports for the compute nodes.

Figure 2) Compute node physical ports





Figure 3 shows the physical ports for the storage nodes.

Figure 3) Storage nodes physical ports



3 Information About NetApp HCI Physical Ports

It is important to plan for the proper deployment of ports on the Mellanox SN2010 network switches. Two high-availability (HA) cabling configurations are available: one with six cables and one with two cables, known as the *flex cabling option*, for easier installation and configuration.

The six-cable solution requires four ports per storage node (two 1GbE and two 10/25GbE) and six ports per compute node (two 1GbE and four 10/25GbE).

The two-cable flex cabling option requires four ports per storage node (two 1GbE and two 10/25GbE) and two ports per compute node (two 10/25GbE).

NetApp highly recommends that you connect the included out-of-band management (IPMI) ports, but this step is not required.

For the purposes of this document, the two-cable option is shown for the NetApp Mellanox switch kit.

For example, two compute nodes and four storage nodes require the following number of ports:

- By node type:
 - Two compute nodes require four 10/25GbE ports.
 - Four storage nodes require eight 1GbE and eight 10/25GbE ports.
- By interface type:
 - 1GbE: 8 ports total (plus 6 ports for IPMI)
 - 10/25GbE: 12 ports total

Figure 2 and Figure 3 shows the ports for the compute and storage nodes.

3.1 Flex cabling option for Compute Nodes (Two-Cable)


NetApp HCI 1.4 and later versions support the flex cabling option, which gives you the option of configuring the NetApp HCI compute nodes with two cables. When you deploy your system with two cables, only ports D and E are configured by NDE. See Figure 4. All traffic types are on these two ports (management, vMotion, iSCSI, and VM Network traffic). In this case, VLANs are used to separate traffic types.

The Two-Cable configuration leads to three possible deployment scenarios when you run NDE:

- Connect and deploy with two cables (ports D and E).
- Connect and deploy with six cables (ports A–F).
- Connect six cables and deploy with two cables (ports D and E).

Storage nodes require four cables to be connected and configured during deployment.

Figure 4) Two-cable compute option.



3.2 Connect and Deploy with Six Cables (Optional)


Choose this option when you want the best possible performance from your NetApp HCI deployment with the maximum number of options for separating different traffic types.

Note: The six-cable option requires more cables than come with the NetApp HCI Mellanox network kit cables. You can order additional cables with the “Build My Own Networking Kit” option on the order form.

Note: Some networking schemes other than the two-cable compute option requiring additional cable types and lengths are also orderable with the “Build my own Networking kit” option in the NetApp HCI Mellanox switch kit order form.

This option requires all network ports to be cabled. The six-cable option gives you dedicated ports for management (two 1/10GbE), for VMs (two 10/25GbE), and for storage (two 10/25GbE). See Figure 5. Deployment with VMware vSphere Standard Switches (VSS) is the default virtual switch type, and the vSphere Distributed Switch (VDS) is optional. To take advantage of the VDS option, you need a VMware Enterprise Plus license.


Figure 5) Six-cable compute option.



3.3 Connect Six Cables and Deploy with Two Cables

Choose this option when you want to leave ports A, B, C, and F unconfigured by NDE. This option consolidates the NDE managed port groups and traffic types to ports D and E. See Figure 6. This approach is beneficial for installations in which some network resources must be physically separated from the management and data plane of the core infrastructure. Some possible use cases are a network DMZ or VMware NSX deployments. This option requires a VMware Enterprise Plus license. Figure 6 only shows the two cables used for deployment.

Figure 6) Connect six-cables use two-cable compute option.



For more information about the NetApp HCI setup and the different cabling options, see [Setting Up and Using NetApp HCI 1.4](#), and see the “Requirements for NetApp HCI Deployment” section for network cable requirements.

3.4 Networking

During deployment of a NetApp HCI system, the NDE installs and configures multiple components. It therefore requires a properly configured and well-functioning network. In general, the network needs to support the following configuration:

- A maximum transmission unit (MTU) of 9000 on all 10/25GbE ports.
- Separate virtual LANs (VLANs) for management and iSCSI (single VLAN deployments are not supported).
- All host ports should have a spanning tree configured to enter a forwarding state immediately (PortFast).

Table 4 describes the basic requirements for the network ports of each node type.

Table 3) Node interface requirements.

Node Type	Interface	Requirements
Storage node	Bond1G—Management Ports A, B	<ul style="list-style-type: none">• Ping all other Bond1G devices• Dedicated management VLAN• Spanning tree edge port
	Bond10G—iSCSI Ports C, D	<ul style="list-style-type: none">• Ping all other Bond10G devices• MTU set to 9000 (no fragmentation)• Dedicated iSCSI VLAN• Spanning tree edge port• Link Aggregation Control Protocol ([LACP] recommended but optional)

Node Type	Interface	Requirements
Compute node	Bond1G—Management Ports A, B	<ul style="list-style-type: none"> • Ping all other Bond1G devices • Dedicated management VLAN • No LACP • Spanning tree edge port
	Bond10G—iSCSI Ports D, E	<ul style="list-style-type: none"> • Ping all other Bond10G devices • MTU set to 9000 (no fragmentation) • Dedicated iSCSI VLAN • No LACP • Spanning tree edge port
	VM_Network and vMotion Ports C, F	<ul style="list-style-type: none"> • MTU set to 9000 (no fragmentation) • vMotion and VM_Network VLANs • No LACP • Spanning tree edge port

4 Switch Information

Note:

- For the purposes of this document, we cover only the cabling and the basic setup of Mellanox SN2010 switches to help you prepare the NetApp HCI system environment for the Network Deployment Engine (NDE) process. During deployment of a NetApp HCI system, the NDE installs and configures multiple components, including integration with VMware. Therefore, it requires a properly configured and well-functioning network. For more detailed information about network setup for the NetApp HCI system, [Setting Up and Using NetApp HCI 1.4](#).
- To correctly deploy your NetApp HCI system, you need the following network infrastructure:
 - A 1GbE infrastructure for management traffic
 - A 10/25GbE infrastructure for vMotion, VM, and storage traffic

Figure 7 shows the port types that are available on the Mellanox SN2010 switch, and Figure 8 shows the port numbering.

Figure 7) Mellanox SN2010 port types overview.




Figure 8) Mellanox SN2010 port numbering.



4.1 Deploying the Mellanox SN2010 Switch

NetApp HCI servers and nodes must be available on the network so that they can be discovered by NDE during the deployment process.

- NDE 1.3 and higher includes the ZeroConf feature that autoconfigures a private non-routable network on Bond10G interfaces. These private IPs are used to set up a temporary network for NDE to use for initial configuration. This eliminates the need to assign temporary routable IP addresses before starting the NDE. See the document [NetApp Theory of Operations WP-7261](#) for details on ZeroConf.

ZeroConf requires a single routable IP address to be assigned to one of the storage nodes. This IP address is used to drive the NDE. During deployment, the permanent IP addresses are applied to all nodes in the system.

“NetApp H-Series Installation and Setup Instructions” (NetApp H-Series ISI) outlines the steps that you must take to manually assign IP addresses to each interface. This document is provided in the product box and is also available from [NetApp Documentation: Product Library A–Z](#).

The next step is to deploy the NetApp HCI system by using NDE. See the “NetApp HCI Deployment Guide,” also available from [NetApp Documentation: Product Library A–Z](#).

5 Mellanox SN2010 Switch Kit Components and Cabling Diagrams

Table 5 lists the cables in the Mellanox switch kit.

Table 5) Mellanox SN2010 kit and cables.

Marketing Part No.	Mfg. Part No.	Description	Mfg Part Number	Qty	Vendor
MSN2010-C2BF	111-04471	Switch, SN20210	MSN2010-CB2F_C21	2	Mellanox
	111-04473	Holocron Cable pack		1	
X6584-R6	112-00290	1Gb cable, RJ45, 1m	797380093	8	Molex
MCP2M00-A001E30N	112-00620	25Gb Cable, SFP28, 1m	MCP2M00-A001E30N	16	Mellanox
MCP1600-C00AE30N	112-00622	100Gb cable, QSFP28, 0.5m	MCP1600-C00AE30N	2	Mellanox
MC3208411-T	332-00484	SFP/RJ45 Transceiver	MC3208411-T	8	Mellanox

The Mellanox kit also has additional cables that you can order, as listed in Table 6. These cables give you additional capabilities to configure the NetApp HCI system and Mellanox SN2010 switches in your environment.

Table 6) Additional cables that you can order.

Marketing Part No.	Mfg. Part No.	Description	Mfg Part Number	Qty	Vendor
X6585-R6	112-00291	1Gb cable, RJ45, 3m	797380094	1	Molex
MCP2M00-A003E30L	112-00621	25Gb Cable, SFP28, 3m	MCP2M00-A003E30L	1	Mellanox
MCP1600-C001E30N	112-00623	100Gb cable, QSFP28, 1m	MCP1600-C001E30N	1	Mellanox
MCP1600-C003E30L	112-00624	100Gb cable, QSFP28, 3m	MCP1600-C003E30L	1	Mellanox
MCP7F00-A01AR30N	112-00625	100Gb to 4x25Gb cable, 1.5m	MCP7F00-A01AR30N	1	Mellanox
MCP7F00-A003R30L	112-00626	100Gb to 4x25Gb cable, 3m	MCP7F00-A003R30L	1	Mellanox

The switch topology that is outlined in Figures 10 through provides a robust configuration and resiliency for the NetApp HCI system. These figures show the cabling diagrams for a compute node and storage node configuration that uses two chassis. The different colors for the ports and for the cables match for a typical cabling scheme.

In Figure 9 the cables that come with the NetApp HCI Mellanox switch kit are enough for a two-compute node and four-storage node configuration, and up to a four-compute node and four-storage node configuration.

Figure 9) Kit cables.




Figure 10) Two-cable setup: two compute nodes and four storage nodes.




Figure 11) Two-cable setup: three compute nodes and four storage nodes




Figure 12) Two-cable setup: four compute nodes and four storage nodes.




Table 7 is an example of a typical VLAN implementation. All VLANs should be tagged, except for the management VLAN. For more information about a two-cable compute node setup, see [Setting Up and Using NetApp HCI 1.4](#), Configuration option A: Two cables for compute nodes.

Table 4) VLAN example used in this document.

Desired networks, example

Network	VLAN	Switch Port Configuration
HCI Management	100	Native
HCI iSCSI Storage	101	Tagged
HCI vMotion	102	Tagged
HCI VM_Network	103	Tagged

5.1 Sample Two-Cable Compute Node Configuration

To set up this Mellanox switch configuration:

Note: You must modify both switches.

1. Access the Mellanox switches. For instructions on how to connect and to configure the switches for the first time, see the [Mellanox MLNX-OS User Manual for Ethernet](#).

2. Set up an Ethernet connection between the switch and a local network machine by using a standard RJ-45 connector.

When you have input general configuration information and have named, saved, and enabled the configuration, the switchports are ready to be configured for the cabling diagram for NetApp HCI.

Note: Important: Do **not** modify both switches at the same time. Wait for switch A to come up completely before you modify switch B.

The console output in Figure 13 is a sample of the configuration CLI commands to run and to deploy on the Mellanox SN2010 switches for a (two-cable compute) NetApp HCI cabling configuration.

Note: Shown in this example are hybrid ports with all VLANs accessible on all ports. Using access ports for management and normal trunk ports for the 10GbE interfaces with VLAN allowed lists are also supported.

Note: In this example, switchport 18 on both switch A and switch B is used for an uplink to another switch.

Figure 13) Example CLI commands for switch A.

```
#Example 2 cabling for Mellanox SN2010 for NetApp HCI 2x4 config for switch A
#Configuration steps are in order they should be entered

#Enable Services
  protocol mlag
  lacp
  lldp
  ip routing vrf default
  dcb priority-flow-control enable force

#Configure VLANs
  vlan 100-103
  exit
  vlan 100 name "HCI Management"
  vlan 101 name "HCI iSCSI Storage"
  vlan 102 name "HCI vMotion"
  vlan 103 name "HCI VM_Network"

#Configure IPL
  interface port-channel 100
  description IPL
  exit
  interface ethernet 1/20 channel-group 100 mode active
  interface ethernet 1/20 description ISL-SWB_01
  interface ethernet 1/22 channel-group 100 mode active
  interface ethernet 1/22 description ISL-SWB_02
  vlan 4000
  exit
  interface port-channel 100 ipl 1
  interface port-channel 100 dcb priority-flow-control mode on force
  interface vlan 4000
  #IPL Ip should *not* be in the management network
  ip address 10.0.0.1 255.255.255.0
  ipl 1 peer-address 10.0.0.2
  exit

#Configure MLAG VIP
#Both MLAG name and ip need to be unique
#MLAG IP *should* be in the management network
  mlag-vip MLAG-VIP-DOM ip 10.xxx.xxx.174 /24 force
  #mlag system mac can be any unicast MAC
  mlag system-mac 00:00:5E:00:AA:01
  no mlag shutdown

#Configure Uplink
  interface mlag-port-channel 101
  description Uplink UPLINK-SWITCH port 39
  exit
  interface ethernet 1/18 description Uplink to UPLINK-SWITCH E1/39
  interface ethernet 1/18 speed 10000 force
  interface mlag-port-channel 101 mtu 9216 force
  interface ethernet 1/18 mlag-channel-group 101 mode active
  interface mlag-port-channel 101 switchport mode hybrid
```

```

interface mlag-port-channel 101 switchport hybrid allowed-vlan all
interface mlag-port-channel 101 no shutdown

#Configure Management Ports (only for storage nodes)
interface ethernet 1/1-1/4 spanning-tree bpdufilter enable
interface ethernet 1/1-1/4 spanning-tree port type edge
interface ethernet 1/1-1/4 spanning-tree bpduguard enable
interface ethernet 1/1-1/4 speed 1000 force
interface ethernet 1/1-1/4 switchport mode hybrid
interface ethernet 1/1 description HCI-STG-01 eth2:portA
interface ethernet 1/2 description HCI-STG-02 eth2:portA
interface ethernet 1/3 description HCI-STG-03 eth2:portA
interface ethernet 1/4 description HCI-STG-04 eth2:portA
interface ethernet 1/1 switchport hybrid allowed-vlan all
interface ethernet 1/2 switchport hybrid allowed-vlan all
interface ethernet 1/3 switchport hybrid allowed-vlan all
interface ethernet 1/4 switchport hybrid allowed-vlan all
interface ethernet 1/1 switchport access vlan 100
interface ethernet 1/2 switchport access vlan 100
interface ethernet 1/3 switchport access vlan 100
interface ethernet 1/4 switchport access vlan 100

#Configure iSCSI Ports (storage nodes)
interface ethernet 1/7-1/10 mtu 9216 force
interface ethernet 1/7-1/10 spanning-tree bpdufilter enable
interface ethernet 1/7-1/10 spanning-tree port type edge
interface ethernet 1/7-1/10 spanning-tree bpduguard enable
interface ethernet 1/7-1/10 switchport mode hybrid
interface ethernet 1/1 description HCI-STG-01 eth0:portC
interface ethernet 1/2 description HCI-STG-02 eth0:portC
interface ethernet 1/3 description HCI-STG-03 eth0:portC
interface ethernet 1/4 description HCI-STG-04 eth0:portC
interface mlag-port-channel 201-204
exit
interface mlag-port-channel 201-204 no shutdown
interface mlag-port-channel 201-204 mtu 9216 force
interface mlag-port-channel 201-204 lacp-individual enable force
interface ethernet 1/7-1/10 lacp port-priority 10
interface ethernet 1/7-1/10 lacp rate fast
interface ethernet 1/7 mlag-channel-group 201 mode active
interface ethernet 1/8 mlag-channel-group 202 mode active
interface ethernet 1/9 mlag-channel-group 203 mode active
interface ethernet 1/10 mlag-channel-group 204 mode active
interface mlag-port-channel 201-204 switchport mode hybrid
interface mlag-port-channel 201 switchport hybrid allowed-vlan all
interface mlag-port-channel 202 switchport hybrid allowed-vlan all
interface mlag-port-channel 203 switchport hybrid allowed-vlan all
interface mlag-port-channel 204 switchport hybrid allowed-vlan all
interface mlag-port-channel 201 switchport access vlan 100
interface mlag-port-channel 202 switchport access vlan 100
interface mlag-port-channel 203 switchport access vlan 100
interface mlag-port-channel 204 switchport access vlan 100

#Configure data ports (compute nodes)
interface ethernet 1/11-1/12 mtu 9216 force
interface ethernet 1/11-1/12 spanning-tree bpdufilter enable
interface ethernet 1/11-1/12 spanning-tree port type edge
interface ethernet 1/11-1/12 spanning-tree bpduguard enable
interface ethernet 1/11-1/12 switchport mode hybrid
interface ethernet 1/11 description HCI-CMP-01 eth1:portD
interface ethernet 1/12 description HCI-CMP-02 eth1:portD
interface ethernet 1/11 switchport hybrid allowed-vlan all
interface ethernet 1/12 switchport hybrid allowed-vlan all
interface ethernet 1/11 switchport access vlan 100
interface ethernet 1/12 switchport access vlan 100

#Management Interface and DNS
interface mgmt1 ip arp timeout 60
ip name-server 10.xxx.x.250
ip name-server 10.xxx.x.251
hostname Project1_SWA

# DHCP relay configuration
ip dhcp relay instance 1 vrf default
ip dhcp relay instance 1 vrf-auto-helper
ip dhcp relay instance 1 address 10.xxx.xxx.9

```



```
ip dhcp relay instance 1 always-on
```

Both the A-side and the B-side switch in the cabling diagram need the same configuration to be deployed and enabled.

Figure 14) Example CLI commands for switch B.

```
#Example 2 cabling for Mellanox SN2010 for NetApp HCI 2x4 config for switch B
#Configuration steps are in order they should be entered

#Enable Services
protocol mlag
lacp
lldp
ip routing vrf default
dcb priority-flow-control enable force

#Configure VLANs
vlan 100-103
exit
vlan 100 name "HCI Management"
vlan 101 name "HCI iSCSI Storage"
vlan 102 name "HCI vMotion"
vlan 103 name "HCI VM_Network"

#Configure IPL
interface port-channel 100
description IPL
exit
interface ethernet 1/20 channel-group 100 mode active
interface ethernet 1/20 description ISL-SWA_01
interface ethernet 1/22 channel-group 100 mode active
interface ethernet 1/22 description ISL-SWA_02
vlan 4000
exit
interface port-channel 100 ipl 1
interface port-channel 100 dcb priority-flow-control mode on force
interface vlan 4000
#IPL Ip should *not* be in the management network
ip address 10.0.0.2 255.255.255.0
ipl 1 peer-address 10.0.0.1
exit

#Configure MLAG VIP
#Both MLAG name and ip need to be unique
#MLAG IP *should* be in the management network
mlag-vip MLAG-VIP-DOM ip 10.xxx.xxx.174 /24 force
#mlag system mac can be any unicast MAC
mlag system-mac 00:00:5E:00:AA:01
no mlag shutdown

#Configure Uplink
interface mlag-port-channel 101
description Uplink UPLINK-SWITCH port 39
exit
interface ethernet 1/18 description Uplink to UPLINK-SWITCH E1/39
interface ethernet 1/18 speed 10000 force
interface mlag-port-channel 101 mtu 9216 force
interface ethernet 1/18 mlag-channel-group 101 mode active
interface mlag-port-channel 101 switchport mode hybrid
interface mlag-port-channel 101 switchport hybrid allowed-vlan all
interface mlag-port-channel 101 no shutdown

#Configure Management Ports (only for storage nodes)
interface ethernet 1/1-1/4 spanning-tree bpduguard enable
interface ethernet 1/1-1/4 spanning-tree port type edge
interface ethernet 1/1-1/4 spanning-tree bpduguard enable
interface ethernet 1/1-1/4 speed 1000 force
interface ethernet 1/1-1/4 switchport mode hybrid
interface ethernet 1/1 description HCI-STG-01 eth3:portB
interface ethernet 1/2 description HCI-STG-02 eth3:portB
interface ethernet 1/3 description HCI-STG-03 eth3:portB
interface ethernet 1/4 description HCI-STG-04 eth3:portB
interface ethernet 1/1 switchport hybrid allowed-vlan all
interface ethernet 1/2 switchport hybrid allowed-vlan all
interface ethernet 1/3 switchport hybrid allowed-vlan all
```

```

interface ethernet 1/4 switchport hybrid allowed-vlan all
interface ethernet 1/1 switchport access vlan 100
interface ethernet 1/2 switchport access vlan 100
interface ethernet 1/3 switchport access vlan 100
interface ethernet 1/4 switchport access vlan 100

#Configure iSCSI Ports (storage nodes)
interface ethernet 1/7-1/10 mtu 9216 force
interface ethernet 1/7-1/10 spanning-tree bpdufilter enable
interface ethernet 1/7-1/10 spanning-tree port type edge
interface ethernet 1/7-1/10 spanning-tree bpduguard enable
interface ethernet 1/7-1/10 switchport mode hybrid
interface ethernet 1/1 description HCI-STG-01 eth1:portD
interface ethernet 1/2 description HCI-STG-02 eth1:portD
interface ethernet 1/3 description HCI-STG-03 eth1:portD
interface ethernet 1/4 description HCI-STG-04 eth1:portD
interface mlag-port-channel 201-204
exit
interface mlag-port-channel 201-204 no shutdown
interface mlag-port-channel 201-204 mtu 9216 force
interface mlag-port-channel 201-204 lacp-individual enable force
interface ethernet 1/7-1/10 lacp port-priority 10
interface ethernet 1/7-1/10 lacp rate fast
interface ethernet 1/7 mlag-channel-group 201 mode active
interface ethernet 1/8 mlag-channel-group 202 mode active
interface ethernet 1/9 mlag-channel-group 203 mode active
interface ethernet 1/10 mlag-channel-group 204 mode active
interface mlag-port-channel 201-204 switchport mode hybrid
interface mlag-port-channel 201 switchport hybrid allowed-vlan all
interface mlag-port-channel 202 switchport hybrid allowed-vlan all
interface mlag-port-channel 203 switchport hybrid allowed-vlan all
interface mlag-port-channel 204 switchport hybrid allowed-vlan all
interface mlag-port-channel 201 switchport access vlan 100
interface mlag-port-channel 202 switchport access vlan 100
interface mlag-port-channel 203 switchport access vlan 100
interface mlag-port-channel 204 switchport access vlan 100

#Configure data ports (compute nodes)
interface ethernet 1/11-1/12 mtu 9216 force
interface ethernet 1/11-1/12 spanning-tree bpdufilter enable
interface ethernet 1/11-1/12 spanning-tree port type edge
interface ethernet 1/11-1/12 spanning-tree bpduguard enable
interface ethernet 1/11-1/12 switchport mode hybrid
interface ethernet 1/11 description HCI-CMP-01 eth4:portE
interface ethernet 1/12 description HCI-CMP-02 eth4:portE
interface ethernet 1/11 switchport hybrid allowed-vlan all
interface ethernet 1/12 switchport hybrid allowed-vlan all
interface ethernet 1/11 switchport access vlan 100
interface ethernet 1/12 switchport access vlan 100

#Management Interface and DNS
interface mgmt1 ip arp timeout 60
ip name-server 10.xxx.x.250
ip name-server 10.xxx.x.251
hostname Project1_SWB

# DHCP relay configuration
ip dhcp relay instance 1 vrf default
ip dhcp relay instance 1 vrf-auto-helper
ip dhcp relay instance 1 address 10.xxx.xxx.9
ip dhcp relay instance 1 always-on

```

Figure 15 shows the cabling of Mellanox switch A and switch B for the preceding Mellanox configuration CLI commands.

Figure 15) Port information for example configuration CLI commands.

SN2010-SWA

1	4	7	10	13	16	19	21
stg-01-eth2	stg-04-eth2	stg-01-eth0	stg-04-eth0				
stg-02-eth2		stg-02-eth0	esx-01-eth1				
stg-03-eth2		stg-03-eth0	esx-02-eth1		Cust Uplink	ISL	ISL
3	6	9	12	15	18	20	22

SN2010-SWB

1	4	7	10	13	16	19	21
stg-01-eth3	stg-04-eth3	stg-01-eth1	stg-04-eth1				
stg-02-eth3		stg-02-eth1	esx-01-eth4				
stg-03-eth3		stg-03-eth1	esx-02-eth4		Cust Uplink	ISL	ISL
3	6	9	12	15	18	20	22

Legend

STG-iSCSI Ports
ESX-VM Ports
Customer Uplinks
Inter Switch Links
Stg Mgmt (1G)

6 Additional Information

This document provides general and overview information about setting up NetApp HCI with Mellanox SN2010 Ethernet switches. It also provides the VLAN information that you need before you run the NetApp Deployment Engine (NDE) for a completed system setup.

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

- [Setting Up and Using NetApp HCI 1.4](#)
- [NetApp Theory of Operations WP-7261](#)
- [NetApp HCI Documentation](#)
- [NetApp HCI Resources](#)

7 Acknowledgements

- The author would like to thank Aaron Patten for contributing to this technical report.

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 © 1994–2018 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.