

HyperFlex and the Network Control Policy

Contents

[Introduction](#)

[HyperFlex and the Network Control Policy](#)

Introduction

This article will explain what the Network Control Policy is within UCS and how it relates to the operation of your HyperFlex cluster under various scenarios.

HyperFlex and the Network Control Policy

What is the Network Control Policy? The Network Control Policy (NCP) defines the following features and actions:

Cisco Discovery Protocol (CDP): Enabled or Disabled

MAC Register Mode: Only Native VLAN or All Host VLANs

Action on Uplink Fail: Link Down or Warning

MAC Security - Forge: Allow or Deny

LLDP - Transmit/Receive: Disabled or Enabled

The HX Installer will create the following two NCPs under **LAN / Policies / root / Sub-Organization / <HX Cluster name> / Network Control Policies /**

HyperFlex-infra

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **HyperFlex-infra**

Description : Network Control policy for infrastructure vNICs Hype

Owner : **Local**

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

Transmit : Disabled Enabled

Receive : Disabled Enabled

HyperFlex-vm

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **HyperFlex-vm**

Description : Network Control policy for VM vNICs on HyperFlex s

Owner : **Local**

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

Transmit : Disabled Enabled

Receive : Disabled Enabled

The Network Control Policy defined above are used by the vNIC templates created by the HyperFlex Installer. The vNIC Templates are located unser **LAN / Policies / root / Sub-Organization / <HX Cluster name> / vNIC Templates /**

LAN / Policies / root / Sub-Organizations / hx-1-sjs / vNIC Templates / vNIC Template hv-m...

General VLANs VLAN Groups Faults Events

Actions

- Modify VLANs
- Modify VLAN Groups
- Delete
- Show Policy Usage
- Use Global

Properties

Name : **hv-mgmt-a**

Description :

Owner : **Local**

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter VM

Template Type : Initial Template Updating Template

CDN Source : vNIC Name User Defined

MTU :

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

Policies

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

The following vNIC templates use the NCP **HyperFlex-infra**:

- hv-mgmt-a
- hv-mgmt-b
- hv-vmotion-a
- hv-vmotion-b
- storage-data-a
- storage-data-b

The following vNIC templates use the NCP **HyperFlex-vm**:

- vm-network-a
- vm-network-b

Let's drill down on NCP policy names HyperFlex-infra and the Action on Uplink Fail. By default, the Action on Uplink Fail is set to Link Down. This means that the vNIC will be instructed to go into a down state when its corresponding Uplink (logical or physical) goes down. If we go to the VIF tab of a server under **Equipment / Rack-Mounts / Servers / Server #**, we can see what uplink our vNICs are utilizing:

Equipment / Rack-Mounts / Servers / Server 4

Inventory Virtual Machines Hybrid Display Installed Firmware SEL Logs CIMC Sessions **VIF Paths** Power Control Monitor Health Diagnostics Faults Events FSM Statistics T >

| Name | Adapter Port | FEX Host Port | FEX Network Port | FI Server Port | vNIC | FI Uplink | Link State | State Qual |
|----------------------|--------------|---------------|------------------|----------------|----------------|-----------|------------|------------|
| ▼ Path A/1 1/2 A/1/8 | | | | | | | | |
| Virtual Circuit 1556 | | | | | hv-mgmt-a | A/PC- 1 | Up | |
| Virtual Circuit 1557 | | | | | storage-data-a | A/PC- 1 | Up | |
| Virtual Circuit 1558 | | | | | vm-network-a | A/PC- 1 | Up | |
| Virtual Circuit 1559 | | | | | hv-vmotion-a | A/PC- 1 | Up | |
| ▼ Path B/1 1/1 B/1/8 | | | | | | | | |
| Virtual Circuit 1560 | | | | | hv-mgmt-b | B/PC- 2 | Up | |
| Virtual Circuit 1561 | | | | | storage-data-b | B/PC- 2 | Up | |
| Virtual Circuit 1562 | | | | | vm-network-b | B/PC- 2 | Up | |
| Virtual Circuit 1563 | | | | | hv-vmotion-b | B/PC- 2 | Up | |

The vNICs going to **Fabric Interconnect A** are pinned to **Port-Channel 1**. The vNICs going to **Fabric Interconnect B** are pinned to **Port-Channel 2**. If **Port-Channel 1** does down, the vNICs that go to **Fabric Interconnect A** will be instructed to go down. If we log into vCenter, we will see the corresponding VMNICs as down.

Equipment / Rack-Mounts / Servers / Server 4

Inventory Virtual Machines Hybrid Display Installed Firmware SEL Logs CIMC Sessions **VIF Paths** Power Control Monitor Health Diagnostics Faults Events FSM Statistics T >

| Name | Adapter Port | FEX Host Port | FEX Network Port | FI Server Port | vNIC | FI Uplink | Link State | State Qual |
|-----------------------|--------------|---------------|------------------|----------------|----------------|-----------|------------|---------------------------|
| ▼ Path A/1 1/2 A/1/8 | | | | | | | | |
| Virtual Circuit 15... | | | | | hv-mgmt-a | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | storage-data-a | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | vm-network-a | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | hv-vmotion-a | unpinned | Down | ENM source pinning fai... |
| ▼ Path B/1 1/1 B/1/8 | | | | | | | | |
| Virtual Circuit 15... | | | | | hv-mgmt-b | B/PC- 2 | Up | |
| Virtual Circuit 15... | | | | | storage-data-b | B/PC- 2 | Up | |
| Virtual Circuit 15... | | | | | vm-network-b | B/PC- 2 | Up | |
| Virtual Circuit 15... | | | | | hv-vmotion-b | B/PC- 2 | Up | |

hx-1-esxi-04.sjs.local | ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks

Storage Adapters Storage Devices Host Cache Configur... Protocol Endpoints I/O Filters

▼ Networking Virtual switches VMkernel adapters Physical adapters TCP/IP configuration

Virtual Machines VM Startup/Shutdo... Agent VM Settings Default VM Connati

Physical adapters

Add Networking... Refresh Edit...

| Device | Actual Speed | Configured Speed | Switch | MAC Address | Observed IP Ranges | Wake on LAN Sup... | SR-IOV Status | S |
|--------|--------------|------------------|--------------------|-------------------|-------------------------|--------------------|---------------|----|
| vmnic0 | Down | Auto negotiate | vswitch-hx-inba... | 00:25:b5:99:a1:02 | 172.16.671-172.16.67... | No | Not supported | -- |
| vmnic1 | Down | Auto negotiate | -- | 00:25:b5:99:a3:02 | No networks | No | Not supported | -- |
| vmnic2 | Down | Auto negotiate | -- | 00:25:b5:99:a5:02 | 0.0.01-255.255.255... | No | Not supported | -- |
| vmnic3 | Down | Auto negotiate | -- | 00:25:b5:99:a7:02 | No networks | No | Not supported | -- |
| vmnic4 | 10000 Mb | 10000 Mb | vswitch-hx-inba... | 00:25:b5:99:b2:02 | No networks | No | Not supported | -- |
| vmnic5 | 10000 Mb | 10000 Mb | -- | 00:25:b5:99:b4:02 | No networks | No | Not supported | -- |
| vmnic6 | 10000 Mb | 10000 Mb | -- | 00:25:b5:99:b6:02 | No networks | No | Not supported | -- |
| vmnic7 | 10000 Mb | 10000 Mb | -- | 00:25:b5:99:b8:02 | No networks | No | Not supported | -- |

Since we still have **Port-Channel 2** on **Fabric Interconnect B**, the HyperFlex cluster will remain up and running. So what happens if we also lose **Port-Channel 2** on **Fabric Interconnect B**.

Equipment / Rack-Mounts / Servers / Server 4

General Inventory Virtual Machines Hybrid Display Installed Firmware SEL Logs CIMC Sessions VIF Paths Power Control Monitor Health Diagnostics Faults Events FSM S>

+ - Advanced Filter Export Print

| Name | Adapter Port | FEX Host Port | FEX Network Port | FI Server Port | vNIC | FI Uplink | Link State | State Qual |
|-----------------------|--------------|---------------|------------------|----------------|----------------|-----------|------------|---------------------------|
| ▼ Path A/1 | 1/2 | | | A/1/8 | | | | |
| Virtual Circuit 15... | | | | | hv-mgmt-a | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | storage-data-a | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | vm-network-a | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | hv-vmotion-a | unpinned | Down | ENM source pinning fai... |
| ▼ Path B/1 | 1/1 | | | B/1/8 | | | | |
| Virtual Circuit 15... | | | | | hv-mgmt-b | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | storage-data-b | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | vm-network-b | unpinned | Down | ENM source pinning fai... |
| Virtual Circuit 15... | | | | | hv-vmotion-b | unpinned | Down | ENM source pinning fai... |

As you would expect, all vNICs are in a Down Link State and the corresponding VMNICS are also Down.

```
The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
[root@hx-1-esxi-04:~] esxcli network nic list
Name      PCI Device  Driver  Admin Status  Link Status  Speed  Duplex  MAC Address  MTU  Description
-----
vmnic0    0000:05:00.0  nenic  Up            Down         0      Half    00:25:b5:99:a1:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic1    0000:06:00.0  nenic  Up            Down         0      Half    00:25:b5:99:a3:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic2    0000:07:00.0  nenic  Up            Down         0      Half    00:25:b5:99:a5:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic3    0000:08:00.0  nenic  Up            Down         0      Half    00:25:b5:99:a7:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic4    0000:09:00.0  nenic  Up            Down         0      Half    00:25:b5:99:b2:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic5    0000:0a:00.0  nenic  Up            Down         0      Half    00:25:b5:99:b4:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic6    0000:0b:00.0  nenic  Up            Down         0      Half    00:25:b5:99:b6:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic7    0000:0c:00.0  nenic  Up            Down         0      Half    00:25:b5:99:b8:02  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
[root@hx-1-esxi-04:~]
```

Since all VMNICS are down, connectivity to the ESXi management is lost and the **HyperFlex cluster will go offline** as the storage controller VMs can no longer communicate to each other.

The use of virtual port-channels, vPC, will provide the best redundancy for HyperFlex. Currently, we do not support using warning instead of link down. There is a possibility that traffic could become blackholed and affect the network redundancy of HyperFlex.