

HyperFlex e a política de controle de rede

Contents

[Introduction](#)

[HyperFlex e a política de controle de rede](#)

Introduction

Este artigo explicará o que é a política de controle de rede no UCS e como ela se relaciona à operação do cluster do HyperFlex em vários cenários.

HyperFlex e a política de controle de rede

Qual é a política de controle de rede? A política de controle de rede (NCP) define os seguintes recursos e ações:

Cisco Discovery Protocol (CDP): Ativado ou desativado

Modo de registro MAC: Somente VLAN nativa ou todas as VLANs de host

Ação de falha no uplink: Link desativado ou aviso

Segurança MAC - Forge: Permitir ou Negar

LLDP - Transmitir/Receber: Desabilitado ou habilitado

O instalador HX criará os dois NCPs a seguir em **LAN / Políticas / raiz / sub-organização / <nome do cluster HX> / Políticas de controle de rede /**

HyperFlex-infra

General	Events
---------	--------

Actions	Properties
Delete	Name : HyperFlex-infra
Show Policy Usage	Description : Network Control policy for infrastructure vNICs Hype
Use Global	Owner : Local
	CDP : <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
	MAC Register Mode : <input checked="" type="radio"/> Only Native Vlan <input type="radio"/> All Host Vlans
	Action on Uplink Fail : <input checked="" type="radio"/> Link Down <input type="radio"/> Warning
	MAC Security
	Forge : <input checked="" type="radio"/> Allow <input type="radio"/> Deny
	LLDP
	Transmit : <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
	Receive : <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled

HyperFlex-vm

General	Events
---------	--------

Actions	Properties
Delete	Name : HyperFlex-vm
Show Policy Usage	Description : Network Control policy for VM vNICs on HyperFlex s
Use Global	Owner : Local
	CDP : <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
	MAC Register Mode : <input checked="" type="radio"/> Only Native Vlan <input type="radio"/> All Host Vlans
	Action on Uplink Fail : <input checked="" type="radio"/> Link Down <input type="radio"/> Warning
	MAC Security
	Forge : <input checked="" type="radio"/> Allow <input type="radio"/> Deny
	LLDP
	Transmit : <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
	Receive : <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled

A política de controle de rede definida acima é usada pelos modelos vNIC criados pelo HyperFlex Installer. Os Modelos vNIC estão localizados em **LAN / Políticas / raiz / sub-organização / <nome do cluster HX> / Modelos vNIC /**

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 :

Os seguintes modelos vNIC usam o NCP **HyperFlex-infra**:

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

Os seguintes modelos vNIC usam o NCP **HyperFlex-vm**:

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

Vamos detalhar os nomes de políticas do NCP HyperFlex-infra e a Ação sobre Falha de Uplink. Por padrão, a Ação em Falha de Uplink é definida como Link Down. Isso significa que a vNIC será instruída a entrar em um estado inativo quando seu uplink correspondente (lógico ou físico) for desativado. Se formos até a guia VIF de um servidor em **Equipment / Rack-Mounts / Servers / Server #**, podemos ver que uplink nossas vNICs estão utilizando:

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	

As vNICs que vão para a interconexão de estrutura A são vinculadas ao canal de porta 1. As vNICs que vão para a interconexão de estrutura B são ligadas ao canal de porta 2. Se o Port-Channel 1 falhar, as vNICs que vão para o Fabric Interconnect A serão instruídas a desligar. Se fizermos login no vCenter, veremos as VMNICs correspondentes desativadas.

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	

Como ainda temos Port-Channel 2 no Fabric Interconnect B, o cluster HyperFlex permanecerá ativo e em execução. Então, o que acontece se também perdermos o Port-Channel 2 no 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...

Como você esperaria, todos os vNICs estão em um estado de link inativo e os VMNICS correspondentes também estão inativos.

```
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:~]
```

Como todas as VMNICS estão inoperantes, a conectividade com o gerenciamento do ESXi é perdida e o cluster do HyperFlex ficará offline, pois as VMs do controlador de armazenamento não podem mais se comunicar entre si.

O uso de canais de porta virtuais, vPC, fornecerá a melhor redundância para o HyperFlex. Atualmente, não oferecemos suporte ao uso de aviso em vez de link inativo. Há uma possibilidade de que o tráfego possa ficar bloqueado e afetar a redundância de rede do HyperFlex.