

Network Adapter Virtualization Design (Adapter-FEX) with Cisco Nexus 5500 Switches and Cisco Nexus 2232 Fabric Extenders

Design Guide

November, 2011

For further information, questions and comments please contact ccbu-pricing@cisco.com

<u>Introduction</u>	3
<u>Software and Hardware Requirements</u>	4
<u>Terminology</u>	4
<u>Static Interface/Static vNIC or vHBA</u>	5
<u>Dynamic Interface/Dynamic vNIC</u>	6
<u>Fixed veth Interface</u>	6
<u>Floating veth Interface</u>	6
<u>Cisco UCS P81E Capabilities</u>	6
<u>Provisioning Model for vethernet Interfaces</u>	10
<u>Communication Between Switchport and Network Adapter</u>	10
<u>Static and Dynamic Provisioning</u>	10
<u>Port-Profiles</u>	11
<u>vNIC Configuration Example</u>	12
<u>Dynamically Provisioned veth</u>	12
<u>Statically Provisioned veth</u>	13
<u>Monitoring vethernet Interfaces</u>	14
<u>Role of vPC in A-FEX</u>	15
<u>Traffic Forwarding in A-FEX</u>	15
<u>Provisioning Model for vFC Interfaces</u>	17
<u>Topology Choice</u>	20
<u>A-FEX Connectivity without vPC</u>	21
<u>The Benefit of Using vPC</u>	22
<u>vPC with FEX Straight-Through</u>	23
<u>vPC Orphan Port Considerations with Layer 2 Uplinks</u>	24
<u>SAN Connectivity with vPC and FEX Straight-Through Mode</u>	25
<u>vPC with FEX Active/Active</u>	26
<u>SAN Connectivity with FEX Active/Active</u>	28
<u>Routing Considerations</u>	30
<u>Sample Configuration Steps</u>	31
<u>Verify Licensing Requirements</u>	31
<u>Configure vPC</u>	31
<u>Configure Fabric Extenders (if needed)</u>	33
<u>Enable the A-FEX Feature</u>	33
<u>Configure All Switchports Connected to VIC-Capable Adapters</u>	33
<u>Configure VIC Adapters to Operate in A-FEX Mode</u>	34
<u>Configure Port-Profiles on Both vPC Peers</u>	34
<u>Configure the FC Connectivity to the SAN with Unified Ports</u>	35
<u>Configure the FCoE connectivity from the A-FEX Adapter</u>	36
<u>Configure Zoning</u>	37
<u>Sample Configurations</u>	39
<u>vPC and FEX Straight-Through with Routed Access</u>	39
<u>Cisco Nexus 5500 Switch 1</u>	39
<u>Cisco Nexus 5500 Switch 2</u>	44
<u>vPC and FEX Active/Active with Routed Access</u>	50
<u>Cisco Nexus 5500 Switch 1</u>	50
<u>Cisco Nexus 5500 Switch 2</u>	57

Introduction

This guide describes how to best design networks with virtualized adapters such as the Cisco Unified Computing System™ (Cisco UCS™) P81E Virtual Interface Card:

http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/data_sheet_c78-558230.html.

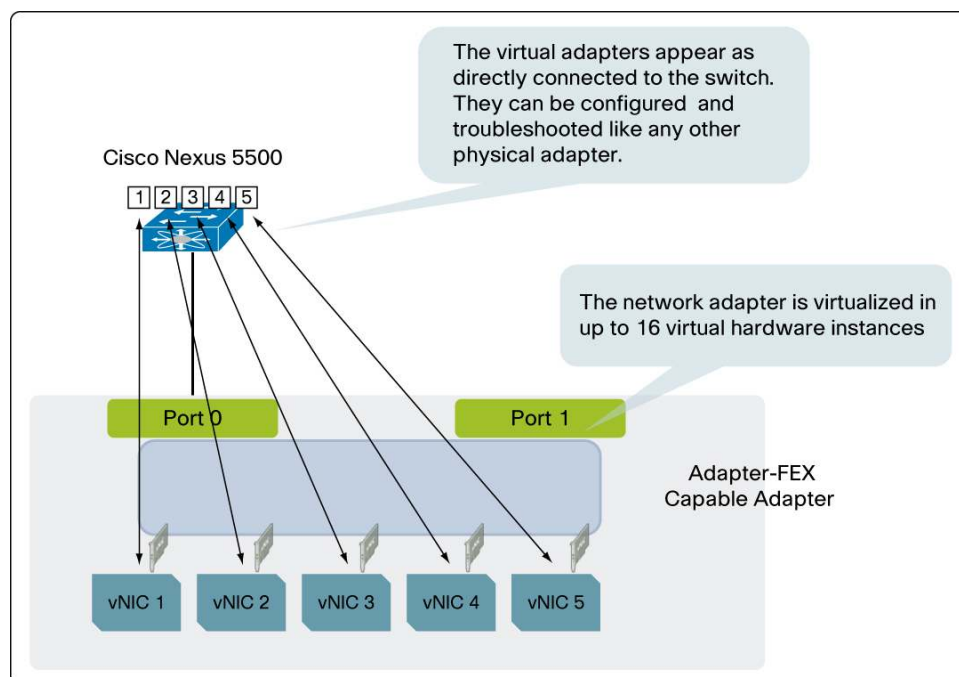
Virtualized adapters that implement the prestandard VN-Tag technology can be connected to a Cisco Nexus® 5500 Switch with Cisco® NX-OS Software Release 5.1(3)N1(1) or later and be remotely operated and configured from the switch itself (which is referred to as an adapter-fabric extender or A-FEX) as described in the following document: http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/data_sheet_c78-657397.html.

As Figure 1 illustrates, a server with a virtualized adapter (called vNICs) can offer the operating system a number of virtual adapters, and with A-FEX technology, they are presented to the Cisco Nexus 5500 platform as directly connected interfaces. All the switching between vNICs occurs on the upstream Cisco Nexus 5500, just as if they were interfaces of a remote linecard or fabric extenders. In addition to this, all features from access control lists (ACLs) to private VLANs, quality of service (QoS), and so on, are available on the remote interfaces.

The redundancy or teaming configuration is not required on the operating system anymore since it is implemented in hardware and controlled by the Cisco Nexus 5500 platform.

The provisioning model allows the network administrator to define profiles with specific network definitions (mode access or trunk, VLAN and so on). The server administrator has the choice of how many vNICs to define and which profile to put them on.

Figure 1. Network Adapter vNICs Appear As Physical Ports on the Cisco Nexus 5500



This guide includes design recommendations for the use of the Cisco UCS P81E network adapter cards in Cisco UCS C-Series Rack-Mount Servers in conjunction with Cisco Nexus 5500 Switches and Cisco Nexus 2232PP Fabric Extenders (FEX).

At the time of this writing, the Cisco Nexus 5000 Series family includes the following products that support A-FEX technology:

- Cisco Nexus 5548P Switch: One-rack-unit (1RU) 10 Gigabit Ethernet, and Fibre Channel over Ethernet (FCoE) switch offering up to 960-Gbps throughput. It has up to 48 ports: 32 fixed 1/10-Gbps Small Form-Factor Pluggable Plus (SFP+) Ethernet and FCoE ports and one expansion slot.
- Cisco Nexus 5548UP Switch: 1RU, 10 Gigabit Ethernet, and FCoE switch offering up to 960-Gbps throughput. It has up to 48 ports: 32 fixed “unified” ports and one expansion slot. The unified ports can be configured as 1/10-Gbps SFP+Ethernet and FCoE ports or as Fibre Channel ports (these two configurations are mutually exclusive).
- Cisco Nexus 5596UP Switch: The Cisco Nexus 5596UP switch is a top-of-rack, 10 Gigabit Ethernet, and FCoE switch offering up to 1920-Gbps throughput and up to 96 ports. The switch has 48 1 and 10 Gigabit Ethernet and FCoE ports and three expansion slots. These 48 ports are “unified ports,” which means that they can be configured as either 1 or 10 Gigabit Ethernet (and FCoE) or as 1-, 2-, 4-, and 8-Gbps native Fibre Channel ports. The use of 1 and 10 Gigabit Ethernet or 1-, 2-, 4-, or 8-Gbps Fibre Channel ports is mutually exclusive but configurable.

At the time of this writing, the Cisco Nexus 2000 Series includes the following product which supports A-FEX:

- Cisco Nexus 2232PP 10 Gigabit Ethernet Fabric Extender: This product has 1 and 10-Gbps Small Form-Factor Pluggable (SFP) and SFP+ Ethernet ports and eight 10 Gigabit Ethernet SFP+ uplink ports. The Cisco Nexus 2232PP is also suitable for carrying Fibre Channel over Ethernet (FCoE) traffic. Servers can attach to the Cisco Nexus 2232PP with twinax cables or optical connectivity in the same way as to a Cisco Nexus 5000 Series Switch.

Software and Hardware Requirements

Adapter-FEX requires the use of the Cisco Nexus 5500 Switch with or without FEX 2232PP running NX-OS Software Release 5.1(3)N1(1) or later.

On the server side, a virtual network tag (VN-Tag)-capable network adapter is required, such as the Cisco UCS P81E Virtual Interface Card (VIC).

If using the Cisco UCS C-Series servers with the Cisco P81E, the following versions of software are required:

- For Cisco Integrated Management Controller firmware, you need a minimum of Version 1.4(1).
- For the C-Series BIOS, you need a minimum of Version 1.4(1).
- For the P81E firmware, you need a minimum of Version 1.6(1).

Terminology

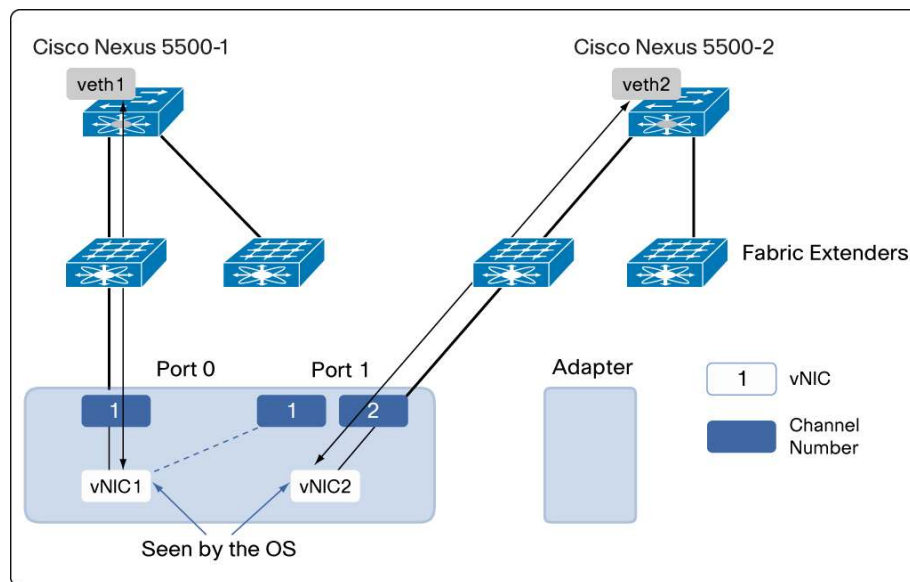
In the context of this document, the following terminology applies:

- vNICs: The hardware instantiations of virtual adapter within a given network adapter.
- vethernet (veth for brief): This term refers to a virtual network adapter “port” (vNIC) as seen by the controlling bridge (that is, by the Cisco Nexus 5500 Switch).

Figure 2 illustrates the association of these two elements. As you can see, the blue box representing the physical network adapter is virtualized in two different instances (vNICs) represented by the grey box. These instances are then using a VN-Tag/channel, which is specific to each wire connecting to the upstream fabric extender (or the Cisco Nexus 5500). For instance, vNIC1 may be using channel 1 on Port 0 as the primary port, and vNIC2 may be configured to use Port 1 with channel 2.

A single vNIC may be able to use both physical ports (Port 0 and Port 1) with a VN-Tag on either port. For instance, in Figure 2, vNIC1 is configured to use both Port 0 and Port 1.

Figure 2. Relationship Between vNICs and veths in a Redundant Network Configuration



The veth box on the Cisco Nexus 5500 Switch represents the instantiation of the vNIC on the Cisco Nexus 5500. All the operations performed by the network administrator, such as looking at counters, shut/no shut, and so on) are performed on the vethernet (veth) interface.

In addition to this terminology, it's also important to distinguish between static and dynamic veth and between fixed and floating veth, as described next.

Static Interface/Static vNIC or vHBA

A static interface is an interface with parameters configured manually by the administrator. A static virtual adapter can be a virtual NIC or a virtual host adapter bus (HBA). A static interface can be a veth or a virtual Fibre Channel (vFC) interface.

For a statically created fixed and floating veth, it is possible for a network administrator to associate configuration to the veth before it is brought up. When creating a static veth, the network administrator specifies which "channel" (for simplicity, you can consider this equivalent to a VN-Tag number) a given veth uses. The server administrator must be sure to define a vNIC on the adapter with the same channel number.

Dynamic Interface/Dynamic vNIC

A dynamic interface is a veth interface that gets configured automatically as a result of adapter and switch communications. The provisioning model of a dynamic interface consists in the configuration of a port-profile on the Cisco Nexus 5500, which appears on the network adapter, followed by the association of the port-profile with the vNIC performed by the server administrator.

A dynamic virtual adapter can be virtual NIC but not virtual HBA. A dynamic virtual interface can be veth interface. Dynamic interfaces have support for hardware-based failover in the Cisco UCS P81E VIC.

Fixed veth Interface

A fixed veth interface is a virtual interface that does not support migration across physical interfaces. When talking about adapter-FEX, the scope is always on fixed veth because adapter-FEX refers to the use of network virtualization by a single (that is, nonvirtualized) operating system.

For fixed veth (static or dynamic), administrators can change configurations, including shut/no shut or create/delete, anytime. The veth-number-to-channel-number binding is persistent unless the administrator changes it.

Floating veth Interface

When the Cisco UCS P81E network adapter is used in a hypervisor environment, each vNIC on the network adapter is associated with one virtual machine (VM). VMs can migrate from one server to another physical server. A virtual interface that “migrates” along with a VM and virtual network link is called **floating virtual interface**.

For a floating (static or dynamic) veth, a configuration change, including shut/no shut, is allowed anytime, regardless of the attached state, except for a binding configuration. Changing the binding configuration is not allowed while a veth is attached. Binding configurations can only be changed if there is no veth associated.

Cisco UCS P81E Capabilities

The Cisco UCS P81E VIC is a PCI Express (PCIe) 2.0 x 8 10-Gbps adapter designed for use with Cisco UCS C-Series Rack-Mount Servers. The virtual interface card is a dual-port 10 Gigabit Ethernet PCIe adapter that can support up to 128 PCIe standards-compliant virtual interfaces, which can be dynamically configured so that both their interface type (network interface card [NIC] or host bus adapter [HBA]) and identity (MAC address and worldwide name [WWN]) are established using just-in-time provisioning. In addition, the Cisco UCS P81E can support network interface virtualization and Cisco VM-FEX technology.

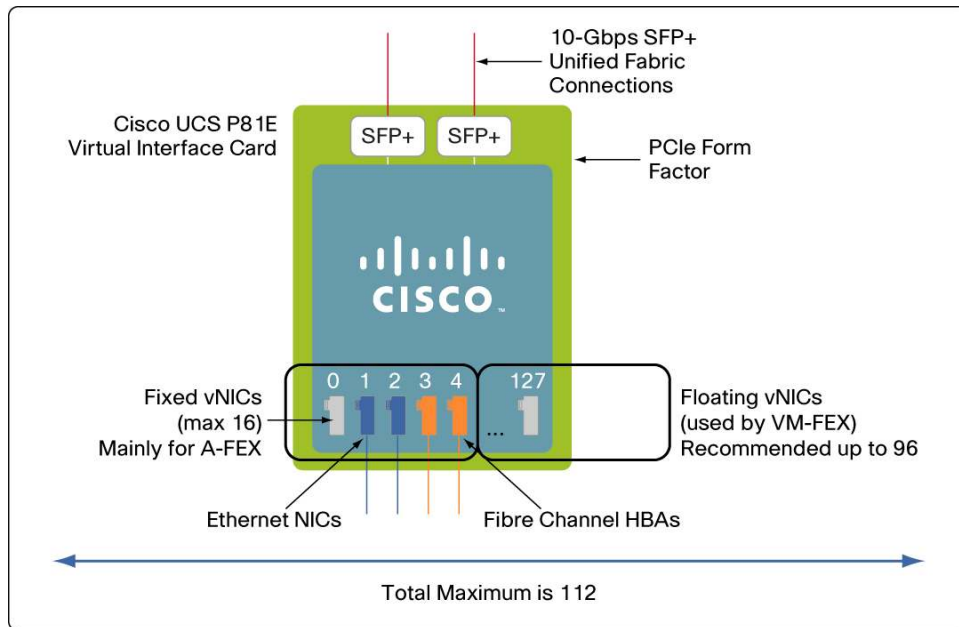
The Cisco UCS P81E can also be configured for virtual HBAs.

The total number of virtual adapters that can be provisioned on the P81E card is 112. As Figure 3 illustrates, this maximum is shared between fixed vNICs and floating vNICs as follows:

- Up to a maximum of 16 fixed vNICs (used for A-FEX purposes) (which leaves space for 96 floating vNICs).
- Up to a maximum of 96 floating vNICs are recommended (used for VM-FEX purposes) (with, for instance two fixed vNICs and two vHBAs).

Note: You may be able to configure a few more vNICs up to, for instance, 115 vNICs, but you will encounter an error message. In addition, the total number of vNICs that you can use is dependent on how many “VN-Tags” can be allocated on a given wire. This concept is explained later in this document.

Figure 3. Up to 112 vNICs Can Be Configured on the Cisco UCS P81E VIC Card

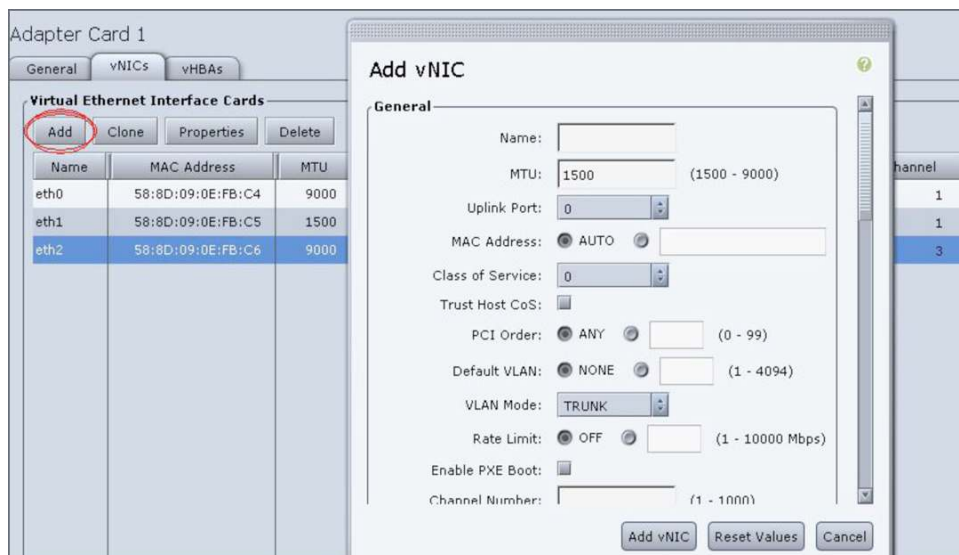


The configuration of the vNICs is performed via the Cisco Integrated Management Controller (CIMC) interface on the UCS C-Series servers. The CIMC is a GUI that provides remote KVM console capabilities and the ability to power up and down the server. In addition, you use the CIMC to configure the network adapter that is installed in the server.

Please refer to the following document to access, configure, and manage the server using the CIMC:
http://www.cisco.com/en/US/products/ps10739/products_installation_and_configuration_guides_list.html.

Figure 4 shows the vNIC tab configuration with the list of vNIC adapters (up to 16). The vNIC tab refers to the adapter-FEX configuration.

Figure 4. The CIMC GUI Allows Adding vNICs and Configuring Their Properties



Each vNIC created gets its own MAC address as shown in Figure 5.

Figure 5. Each A-FEX vNIC Gets a Unique MAC Address

Name	MAC Address
eth0	58:8D:09:0E:F8:5A
eth1	58:8D:09:0E:F8:5B
test1	58:8D:09:0E:F8:5E
test2	58:8D:09:0E:F8:5F
test3	58:8D:09:0E:F8:60
test4	58:8D:09:0E:F8:61

As Figure 6 shows, each vNIC can be configured to use a specific “channel” (that is, VN-Tag) on one of the two physical ports (uplink ports - that is, the adapter ports), and it can be configured for “adapter failover,” called “Uplink Failover” in the CIMC interface.

Figure 6. Configuration Parameters for vNICs in the UCS P81E

vNIC Properties

Uplink Port:

0

MAC Address:

AUTO

58:8D:09:0E:ED:36

Class of Service:

N/A

Trust Host CoS:

PCI Order:

ANY

(0 - 17)

Default VLAN:

(1 - 4094) N/A

VLAN Mode:

N/A

Rate Limit:

(1 - 10000 Mbps) N/A

Enable PXE Boot:

☒

Channel Number:

1

(1 - 1000)

Port Profile:

Enable Uplink Failover:

☐

Failback Timeout:

(0 - 600)

Save Changes

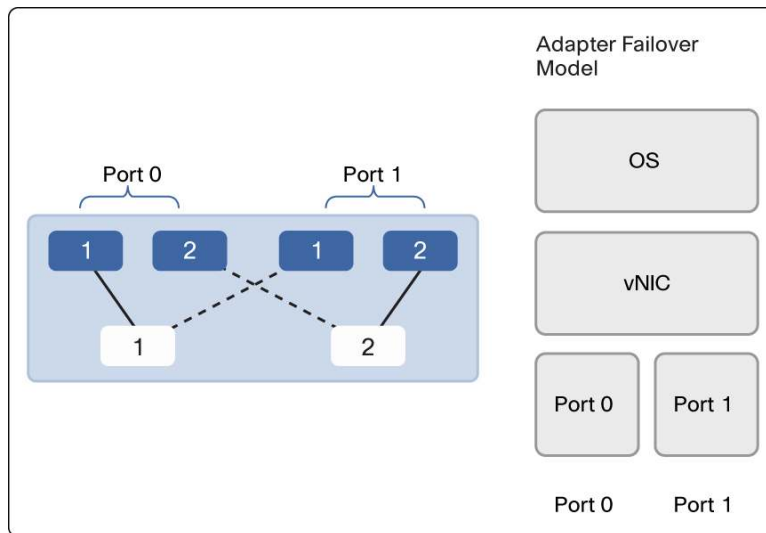
Reset Values

Cancel

Figure 7 illustrates the concept of adapter failover. If adapter failover is enabled, each vNIC can be associated with both adapters (0 and 1), and on each adapter, it is going to use a VN-Tag/channel (which is referred to as “A” on both ports). Hence, when adapter failover is enabled, the channel number of a given vNIC is automatically reserved on both ports: port 0 and port 1.

From an operating system perspective, adapter failover is different from regular NIC teaming in that the operating system is not controlling the failover of the adapter; instead, the network adapter card itself is controlling the failover.

Figure 7. Adapter Failover Model



The P81 card offers a number of features, including TCP offload and so on. The P81 card is also a converged network adapter, so as Figure 8 shows, there are virtual HBAs.

Figure 8. virtual HBAs on the Cisco UCS P81E Card

CPUs Memory Power Supplies Network Adapters Storage PCI Adapters					
Adapter Cards					
PCI Slot	Product Name	Serial Number	Product ID	Vendor	CIMC Management Enab
1	UCS VIC P81E	QC11441A73N	N2XX-ACPCI01	Cisco Systems Inc	no
Adapter Card 1					
General vNICs VM FEXs vHBAs					
Host Fibre Channel Interfaces					
Properties Boot Table Pesistent Bindings					
Name	WWPN	WWNN	Uplink	Boot	Channel
fc0	20:00:58:8D:09:0E:ED:38	10:00:58:8D:09:0E:ED:38	0	disabled	3
fc1	20:00:58:8D:09:0E:ED:39	10:00:58:8D:09:0E:ED:39	1	disabled	4

Provisioning Model for vethernet Interfaces

The provisioning model of A-FEX is based on the concept of port-profiles (of type **vethernet**). Port-profiles are configured on the Cisco Nexus 5500 switch and communicated to the network adapter. From the network adapter management tool (CIMC) the server administrator can associate port-profiles and “vNICs,” which in its turn triggers the creation of the veth on the Cisco Nexus 5500.

Communication Between Switchport and Network Adapter

Before veths can be created, the physical interface connected to the host must be configured with VN-Tag mode and Data Center Bridging Exchange (DCBX) Protocol must run between switchport and host.

In order to enable VN-Tag communication between the switchport and the adapter, you configure the port to operate in VN-Tag mode.

Assume, for instance, that the server P81E card connects to a FEX 2232PP. The configuration would look like this:

```
nexus5500-1(config)# int eth105/1/2
nexus5500-1(config-if)# switchport mode vntag
```

If the server P81E card connects directly to a Cisco Nexus 5500 Switch, the configuration would look like this:

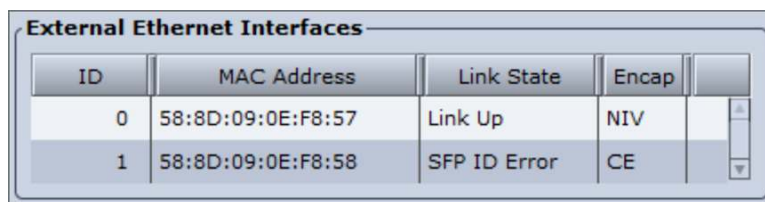
```
nexus5500-1(config)# int eth1/1
nexus5500-1(config-if)# switchport mode vntag
```

The network adapter must be configured for network interface virtualization (NIV) or adapter-FEX.

At this point the host starts to include the NIV capability type-length-value (TLV) fields in its DCBX advertisement. Until then, the network adapter operates in Classical Ethernet (CE) mode. While it operates in CE mode, the host has full connectivity to the network via the physical interface connected to the switch.

You can and should verify the network adapter ports mode of operation from CIMC. The “Encap” field should show “NIV” (Figure 9). If not, verify the configuration of the adapter, and if you are changing from mode CE to mode NIV, be sure to reboot the server (hence the adapter within the server).

Figure 9. UCS P81E Physical Ports Can Operate in NIV Mode or CE Mode



ID	MAC Address	Link State	Encap
0	58:8D:09:0E:F8:57	Link Up	NIV
1	58:8D:09:0E:F8:58	SFP ID Error	CE

Static and Dynamic Provisioning

veth can be either statically configured or dynamically created by the association of a port-profile with a vNIC. The preferred configuration method is the dynamic one.

In this provisioning model, there is no configuration on the switch side that represents the veth. This model depends on the port-profile information that is provided to the switch via the VIC protocol.

For ease of configuration, the switch advertises via the VIC control channel (VIF_SET of 0) the list of port-profile names configured on the switch. The NIV adapter provides a way for server administrator to select from a list of port-profiles to be attached to each vNIC.

These are the protocol steps used by the adapter and the switch to bring up a fixed veth:

1. The NIV adapter sends a VIF_CREATE message with channel number and optional port-profile name.
2. The switch first matches the channel number against its list of static fixed veth configuration. If there is a matching channel number, that veth number is brought up.
3. If there is no static fixed veth with a matching channel number, a dynamic fixed veth is created.
4. If a port-profile name is in the VIF_CREATE message and the fixed veth is not already configured with a port-profile in the switch configuration, the port-profile parameters within the VIF_CREATE message are configured for the veth.

Port-Profiles

Port-profiles are not a functionality that is specific to A-FEX. However, A-FEX uses a particular type of port-profile to interface with the server admin provisioning tool (CIMC in the case of the P81E adapter). Port-profiles provide a template configuration that can then be applied identically to multiple interfaces. The Ethernet port-profile by default is Layer 2 on the Cisco Nexus 5000 Series.

An example of configuration illustrates the configuration template provided by port-profiles:

```
port-profile type ethernet accessvlan
  switchport access vlan 50
  spanning-tree port type edge
  state enabled
```

Whenever you configure **shut** or no **shut** from the port-profile, this gets propagated to the port. Also **state enabled** must be configured for the port-profile configuration to take effect.

Note: Read here for more information about [port-profiles](#).

In the context of A-FEX, the port-profiles are of type **vethernet** (instead of type **Ethernet**).

An example of port-profile that is used for A-FEX is as follows:

```
port-profile type vethernet NIC-VLAN50
  switchport access vlan 50
  switchport mode access
  state enabled
```

When the user creates such a port-profile, it appears on network adapter management tool as one of the available options for vNICs.

Note: When using a virtual PortChannel (vPC), make sure that the same port-profile is configured on both vPC devices. Because the server is dual-homed, it can use either Cisco Nexus 5500 Switch to bring up the veth associated with the port-profile.

vNIC Configuration Example

Figure 10 illustrates the configuration choice for the P81E as a result of the creation of the above port-profile.

By default, the P81E comes with two predefined vNICs: **eth0** and **eth1**. You can see in Figure 10 that you can select which physical port you want the eth0 interface to use by default (in this example it's interface 0). The port-profile NIC-VLAN5 appears in the list. You can also see that the vNIC will try to negotiate using channel 1 for the veth-to-vNIC mapping.

You can also see the option "Enable Uplink Failover." If this box is checked, the server will use port0, for instance, as the primary one, and port1 as the backup (or vice versa if you select port 1 as the "Uplink Port"). In most deployments, you would want to enable uplink failover.

Figure 10. The Properties Window Allows Selecting the Physical Port to Be Used by the vNIC and the Port-Profile

vNIC Properties

Uplink Port: 0

MAC Address: ☐ AUTO ☒ 58:8D:09:0E:F8:5A

Class of Service: N/A

Trust Host CoS: ☐

PCI Order: ☒ ANY ☐ (0 - 17)

Default VLAN: (1 - 4094) N/A

VLAN Mode: N/A

Rate Limit: (1 - 10000 Mbps) N/A

Enable PXE Boot: ☒

Channel Number: 1 (1 - 1000)

Port Profile: NIC-VLAN50

Enable Uplink Failover: ☒

Failback Timeout: (0 - 600)

After you add a vNIC, change the channel, or choose the port-profile, the server might have to be rebooted for the configuration to take effect.

When configuring multiple vNICs in a redundant topology, you may want to make sure that the vNICs alternate in using Uplink Port 0 and Uplink Port 1 in order to maximize traffic distribution across both devices that they are connected to.

Dynamically Provisioned veth

The configuration of the adapter results in the creation of a veth on the Cisco Nexus 5500. If you are using the 5500 in redundancy mode (that is, with vPC), the veth appears on both vPC peers, but the MAC address table configuration appears only on one of the two vPC peers.

For instance, the above configuration results in the following automatic configuration:

```
interface Vethernet32769
  inherit port-profile NIC-VLAN50
  bind interface Ethernet105/1/2 channel 1
```

In addition, the second vNIC has been configured to use VLAN 60 with a different port-profile:

```
interface Vethernet32770
  inherit port-profile NIC-VLAN60
  bind interface Ethernet105/1/2 channel 2
```

The veth numbering for dynamic fixed veth is always above 32768. This is to allow the users to configure static fixed veth in the range below without having to worry about some veth numbers having been taken by dynamic veth.

When you are using a redundant configuration (that is, when using vPC), the veth is going to appear on both vPC peer devices and it will be up on both.

The MAC address will appear only on one of the two vPC peers, as follows:

```
* 50      588d.090e.f85a      static      0      F      F      Veth32769
```

If you save the configuration ("copy run startup") the above veth configuration is saved, and upon reboot of the Cisco Nexus 5500, the association of the veth to the interface and the channel will be already present.

Note: If the veth doesn't get instantiated, it is possible that either the port-profile is not configured on both vPC peers or the server may have to be reloaded. If the veth gets instantiated only on one vPC peer, it's likely that the adapter was not configured for adapter failover.

Statically Provisioned veth

In the static fixed model, the administrator of the switch needs to carefully map the veth being configured on the switch to the channel number of the vNIC that it represents. Since this requires the switch administrator to get the channel information from the server/adapter administrator and this process is prone to errors, the static configuration is considered to be less desirable.

For each static fixed vNIC provisioned in an adapter, a corresponding veth must be created and bound to a **{Ethernet interface, channel-number}** where **Ethernet interface** is the physical interface which the adapter is connected to and **channel-number** is the vNIC instance number provisioned in the adapter.

For instance, if we were to configure a veth for the previously configured vNIC and we were not using the port-profile, a valid configuration would like as follows:

```
int veth10
  bind interface ethernet 105/1/2 channel 1
```

A unique channel number given to each physical interface is required.

Also, if the server administrator changes the channel number on the vNIC interface, the server has to be rebooted (for the adapter to be rebooted) in order for the new channel number to be applied to the configuration of the server.

Monitoring vethernet Interfaces

This section provides some useful commands to monitor the vethernet interfaces that have been provisioned on the Cisco Nexus 5500.

```
nexus5500-1# show vethernet summary
```

Veth Interface	Bound Interface	Channel/ DV-Port	Port Profile	Mac Address	Connectee
Veth11	Eth105/1/2	4			
Veth112	Eth105/1/8	4			
Veth32769	Eth105/1/2	1	NIC-VLAN50		
Veth32769	Eth106/1/2	1	NIC-VLAN50		
Veth32770	Eth105/1/2	2	NIC-VLAN60		
Veth32770	Eth106/1/2	2	NIC-VLAN60		
Veth32771	Eth105/1/8	1	NIC-VLAN50		
Veth32771	Eth106/1/8	1	NIC-VLAN50		
Veth32772	Eth105/1/8	2	NIC-VLAN60		
Veth32772	Eth106/1/8	2	NIC-VLAN60		
Veth32773	Eth105/1/8	100	BACKUP-VLAN		
Veth32773	Eth106/1/8	100	BACKUP-VLAN		

```
nexus5500-1# show vethernet status
```

Interface	VIF-index	Bound If	Chan	Vlan	Status	Mode	Vntag
Veth11	VIF-58	Eth105/1/2	4	1	Up	Active	5
Veth112	VIF-57	Eth105/1/8	4	1	Up	Active	5
Veth32769	VIF-60	Eth105/1/2	1	50	Up	Standby	2
Veth32769	VIF-48	Eth106/1/2	1	50	Up	Active	2
Veth32770	VIF-61	Eth105/1/2	2	60	Up	Active	3
Veth32770	VIF-49	Eth106/1/2	2	60	Up	Standby	3
Veth32771	VIF-59	Eth105/1/8	1	50	Up	Standby	2
Veth32771	VIF-44	Eth106/1/8	1	50	Up	Active	2
Veth32772	VIF-53	Eth105/1/8	2	60	Up	Active	3
Veth32772	VIF-62	Eth106/1/8	2	60	Up	Standby	3
Veth32773	VIF-54	Eth105/1/8	100	100	Up	Standby	6
Veth32773	VIF-65	Eth106/1/8	100	100	Up	Active	6

```
nexus5500-1# show vethernet summary bound interface ethernet 105/1/2
```

Veth Interface	Bound Interface	Channel/ DV-Port	Port Profile	Mac Address	Connectee
Veth32769	Eth105/1/2	1	NIC-VLAN50		
Veth32770	Eth105/1/2	2	NIC-VLAN60		
Veth11	Eth105/1/2	4			
Total 3 Veth interfaces					

Role of vPC in A-FEX

Virtual PortChannels (vPCs) play an important role in A-FEX deployments because the infrastructure provided by vPC ensures uniqueness of veth numbering and synchronization between the vPC peers.

It is fundamental to understand that using vPC with A-FEX doesn't mean that you will be constructing virtual port-channels out of A-FEX veth interfaces. vPC provides the infrastructure to synchronize the database where the information about veth interfaces is stored.

This is useful because a vNIC configured with adapter failover can "attach" to either Cisco Nexus 5500 Switch, depending on the preferred physical port or which of the physical links is up or down. Because of this, the user must configure the same port-profiles on both switches.

The vPC infrastructure makes sure that the veth numbering for the same vNIC is identical on both 5500 switches.

In addition to this, if you configure static fixed veth and you are using a number that is not used on the local Cisco Nexus 5500, vPC will verify whether this number is present in the database. If the number is used on one of the two 5500 switches as part of the same vPC domain, the number won't be allowed.


The vPC primary is the device that decides which veth can be brought up. The vPC secondary, upon receiving **VIF-create** message, verifies with the vPC primary which veth number it can use.

Traffic Forwarding in A-FEX

This section describes how traffic forwarding works in a topology that uses A-FEX technology.

With A-FEX, even if vNICs end up sharing the same physical link, they are really equivalent to two separate physical interfaces, as Figure 11 illustrates. For instance, each one of them can be shut down independently from the upstream Cisco Nexus 5500 Switch and be configured for different security policies.

Figure 11. Multiple vNICs Sharing the Same Physical Link



To the operating system, each vNIC appears as a separate network adapter. The operating system doesn't see whether adapter failover is enabled or not. What's more, there is no requirement to run teaming software on the operating system because adapter failover is implemented in hardware.

In the example shown in Figure 12, the physical adapter is virtualized into two vNICs, each using both ports (uplinks):

- vNIC eth0: uses adapter 0 as the primary one (and adapter 1 as the standby) and is attached to port-profile NIC-VLAN50.
- vNIC eth1: uses adapter 1 as the primary one (and adapter 0 as the standby) and is attached to port-profile NIC-VLAN60.


Figure 12. Example Configuration with vNICs Using Both Physical Ports

Adapter Card 1										
General vNICs VM FEXs vHBAs										
Host Ethernet Interfaces										
Add Clone Properties Delete										
Name	MAC Address	MTU	Uplink Port	CoS	VLAN	VLAN Mode	PXE Boot	Channel	Port Profile	
eth0	58:8D:09:0E:F8:5A	1500	0	N/A	N/A	N/A	enabled	1	NIC-VLAN50	
eth1	58:8D:09:0E:F8:5B	1500	1	N/A	N/A	N/A	enabled	2	NIC-VLAN60	

Figure 13 illustrates the overall connectivity from the operating system to the upstream switches.


The operating system sees two network adapters (eth0 and eth1), which are in reality vNICs, and each one of them is using one specific physical port to connect to the upstream network infrastructure (FEX and the Cisco Nexus 5500). Each virtual network adapter also has a backup path via the other physical port, in case the primary path fails.

Figure 13. Logical Communication Path from vNICs to the LAN




As Figure 14 illustrates, all traffic that is sent by a server over a vNIC is switched at the Cisco Nexus 5500 Switch layer.

Figure 14. vNIC-to-vNIC Switching



Moreover, if one of the physical links fails, adapter failover will move the traffic to the remaining link without any involvement from the operating system, as depicted in Figure 15.

Figure 15. Adapter Failover Provides Uninterrupted Connectivity



Provisioning Model for vFC Interfaces

With A-FEX technology, the benefits of using a network adapter such as the P81E include the possibility to carry Fibre Channel traffic onto the same interface as the LAN traffic. This is because the Cisco UCS P81E card is at the same time a network interface virtualization (NIV) adapter and a converged network adapter.

The Fibre Channel interfaces on the P81E card are called virtual HBAs. FCoE connectivity between the vHBAs and the upstream Cisco Nexus 5500 is configured by defining a virtual Fibre Channel interface. In contrast to veth interfaces, vFC interfaces are configured with static binding; there's no dynamic binding configuration possible with vFCs.

The configuration is achieved by manually associating a veth on the Cisco Nexus 5500 Switch with the vHBA channel and then binding the vFC to the veth as described in the code sample and Figure 16. From the CIMC, you can check the channel number used by the vHBAs (Figure 16). You should also ensure that each HBA uses a different "uplink" - that is, a different network adapter.

Figure 16. vHBA Association with Channel Number

Adapter Card 1						
General vNICs VM FEXs vHBAs						
Host Fibre Channel Interfaces						
Properties Boot Table Persistent Bindings						
Name	WWPN	WWNN	Uplink	Boot	Channel	
fc0	20:00:58:8D:09:0E:F8:5C	10:00:58:8D:09:0E:F8:5C	0	disabled	3	
fc1	20:00:58:8D:09:0E:F8:5D	10:00:58:8D:09:0E:F8:5D	1	disabled	4	

If you don't know which one of the adapters is port0 and port1 and to which FEX or Cisco Nexus 5500 Switch it connects to, you may want to define the veth interface first with one channel and then with the second one. If the veth doesn't go up in the first few seconds, you know that you may have to use the other channel number.


```
interface veth 11
    switchport mode trunk
    switchport trunk allowed vlan 1,11
    bind interface eth 105/1/2 channel 4
    no shut
```

Notice that the veth used for FCoE must be a trunk because of the way FCoE works. Initially in the FCoE Initialization Protocol (FIP) VLAN discovery, FCoE uses the native VLAN, but subsequently FCoE uses the VLAN associated with the VSAN. Because of this, you want to configure the veth to be a trunk and to carry specifically the VLAN used for VSAN purposes as well as the native VLAN.

```
interface vfc11
    bind interface Vethernet11
    no shutdown
nexus5500-1(config)# vsan database
nexus5500-1(config-vsan-db)# vsan 11 interface vfc11
```

Figure 17 illustrates how FCoE connectivity works in an adapter-FEX environment. Each adapter has two virtual HBAs: fc0 and fc1, which are mapped to one of the two physical ports (port 0 and 1). The veth binds to the FEX port and to the channel that is used by the adapter, and the vFC binds to the veth. From the operating system perspective, there are two Fibre Channel storage adapters.


Figure 17. Fibrechannel Dual Fabric Design with vHBAs



The vFC interfaces on the Cisco Nexus 5500 are equivalent to regular Fibre Channel interfaces. So they can be shut down and this will shut down the vHBA on the server without affecting the virtual NICs. Also all properties of zoning - and more properties in general of Fibre Channel fabric configurations - apply to vFCs just as they apply to physical interfaces.

Figure 18 (taken from Cisco Fabric Manager) illustrates how the Cisco Nexus 5500 and the vHBAs appear from a SAN management perspective:

Figure 18. vHBAs As Part of the Fibrechannel Fabric



As you can see the vHBA (Node World Wide Name 10:00:58:8d:09:0e:f8:5d) appears as if it is directly connected to the Cisco Nexus 5500, even if it is just a virtual interface in the P81E card.

In Figure 18, the zone defined for the storage access is highlighted in yellow. The World Wide Name (WWN) of the vHBA can easily be found using the normal Fibre Channel commands: **show flogi database vsan 11** where **vfc11** behaves exactly as any regular Fibre Channel port:

```
nexus5500-1# show flogi database vsan 11
```

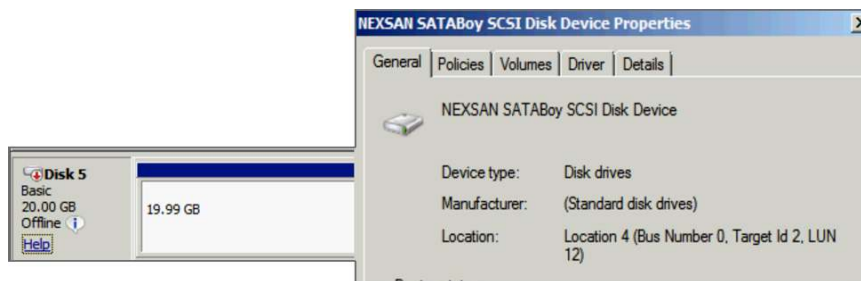
```
-----
INTERFACE      VSAN    FCID          PORT NAME          NODE NAME
-----
vfc11          11     0xa00000     20:00:58:8d:09:0e:f8:5d 10:00:58:8d:09:0e:f8:5d
```

And the zone associated with it is configured as it would be with normal zoning with regular HBAs:

```
nexus5500-1# show zone name tc-esx02
zone name tc-esx02 vsan 11
    pwwn 20:00:58:8d:09:0e:f8:5d
    pwwn 50:00:40:20:01:f4:23:18 [tc-sataboy01-port1]
```

From an operating system perspective, the logical unit number (LUN) targets appear as if they were connected via regular HBAs. For instance, in the case of a Windows 2008 server, the disk management utility would see the screens shown in Figure 19:


Figure 19. The Operating System Can See the Remote LUN via FCoE Over the vHBA



The LUN is visible via two different paths, and as a result you may see a warning message from the operating system indicating that this LUN is reachable via multiple paths (which is why the disk shows as offline in this screen capture).

A look at the disk array (Figure 20) confirms that this is the correct LUN (check the LUN number and compare with the properties from the disk manager):

Figure 20. Storage LUN Assignment from the Disk Array

 4: 'tc-esx02-dedicated' Array: 'Array #1' Capacity: 21.4 GB (20.0 GiB)	LUN12	LUN12
---	-------	-------

Multipath I/O software would operate exactly in the same way as with physical HBAs.

Topology Choice

Servers equipped with an NIV-capable adapter can connect to a Cisco Nexus 5500 system in different ways:

- Directly to the Cisco Nexus 5500 Switch, which is logically equivalent to connecting to a Fabric Extender 2232 in Straight-Through mode (hence this case is not covered).

- To FEX 2232 in FEX Straight-Through mode without vPC. In this topology, each fabric extender module is single-attached to a Cisco Nexus 5500 Switch.
- To FEX 2232 in FEX Straight-Through mode with vPC. In this topology, each fabric extender module is single-attached to a Cisco Nexus 5500 Switch, and the Cisco Nexus 5500 is configured for vPC.
- To FEX 2232 in FEX Active/Active mode with vPC. In this topology, each fabric extender is dual-connected to Cisco Nexus 5500 Switches.

All the topologies - Straight-Through as well as Active/Active - are able to implement a dual fabric topology for the purpose of using Fibre Channel starting from Cisco NX-OS Software Release 5.1(3)N1(1).

Topologies vary also based on the Cisco Nexus 5500 Switch's connectivity to the aggregation layer. This connectivity can be Layer 2 or Layer 3, and in the Layer 2 case it can be based on Spanning Tree Protocol, vPC, or Cisco FabricPath.

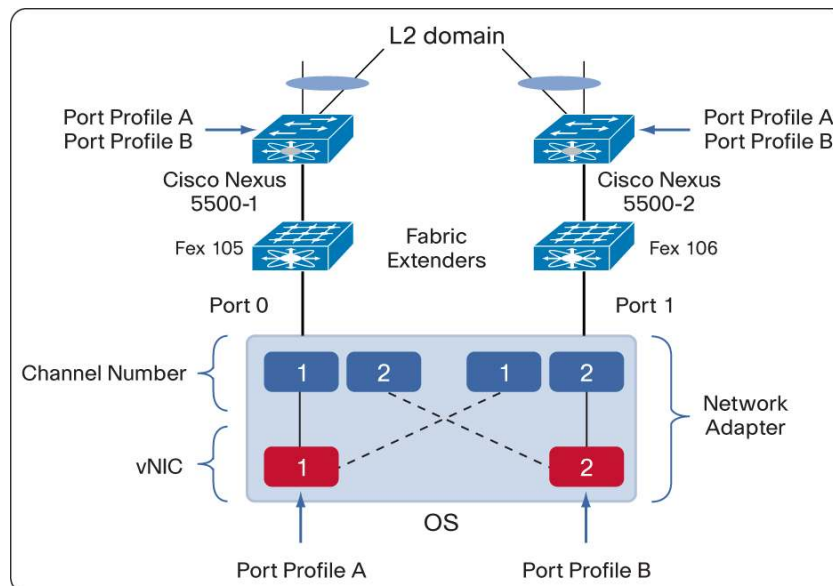
This document covers Layer 3 connectivity to the aggregation layer and vPC connectivity to the aggregation layer.

Note: For simplicity the following examples use the numbering **veth1**, **veth2**, and so on even if for dynamic veth the numbering starts from **veth32768**.

A-FEX Connectivity without vPC

The fabric extender can connect to the Cisco Nexus 5500 Series Switch with a single-homed topology, often referred to as a straight-through connection without vPC, as shown in Figure 21. With the single-homed topology, each fabric extender is attached to a single Cisco Nexus 5500 Series Switch.

Figure 21. Cisco Nexus 2000 Series Straight-Through Topology without vPC



With this topology, the network adapter issues a VIF create message from each vNIC (1 or 2) to the upstream Cisco Nexus 5500. The VIF message creates a veth on the Cisco Nexus 5500 Switches along both the active and the standby path. The port-profile that is used by vNIC1 in this example is port-profile A, and the one used by vNIC2 is port-profile B. Because of redundancy, you need to make sure that both port-profiles are present on both Cisco Nexus 5500 Switches.

With this model, there is no synchronization of the veth namespace, and the failover of MAC addresses would not be automatically triggered but would instead require waiting for the server to send traffic.

For instance, the vNIC1 may be normally active on the left Cisco Nexus 5500. As a result, the MAC table on the left switch would look like this:

```


pods-5548-1# show mac address-table vlan 123
      VLAN      MAC Address      Type      age      Secure NTFY      Ports
-----+-----+-----+-----+-----+-----+-----
* 123      588d.090f.0b3e      static      0          F          F          Veth1

```

No MAC address entry would appear on the other Cisco Nexus 5500 for this vNIC.

In case of failure, the vNIC would start using the alternate path, as depicted in the example shown in Figure 22:

Figure 22. vNIC Failover



When this happens, there's no automatic reprogramming of the MAC address table. For veth5, traffic flow triggers the appearance of the vNIC MAC in the MAC table. Notice that the upstream Layer 2 domain must provide Layer 2 adjacency between the vNICs ports - that is, between port 0 and port 1.

The Benefit of Using vPC

When using vPC in the context of A-FEX, there is no requirement to use the vPC feature to create port-channels. In the context of A-FEX, vPC is used to synchronize the veth database and to optimize the failover behavior of vNICs.

The vPC configuration allows syncing the veth numbers to make sure that the same vNIC uses the same veth numbers on both the switches. For instance, in the previous example, vNIC1 was creating two veths on the Cisco Nexus 5500, veth1 and veth5. With vPC, the vNIC would appear to both switches with the same veth number.

Using vPC also allows syncing the MAC addresses on both the switches, thereby making sure that a failover will cause the switches to still populate the MAC on the new active rather than wait for the new active to send traffic to learn those MACs.

vPC topologies that involve FEX are categorized as:

- FEX Straight-Through
- FEX Active/Active

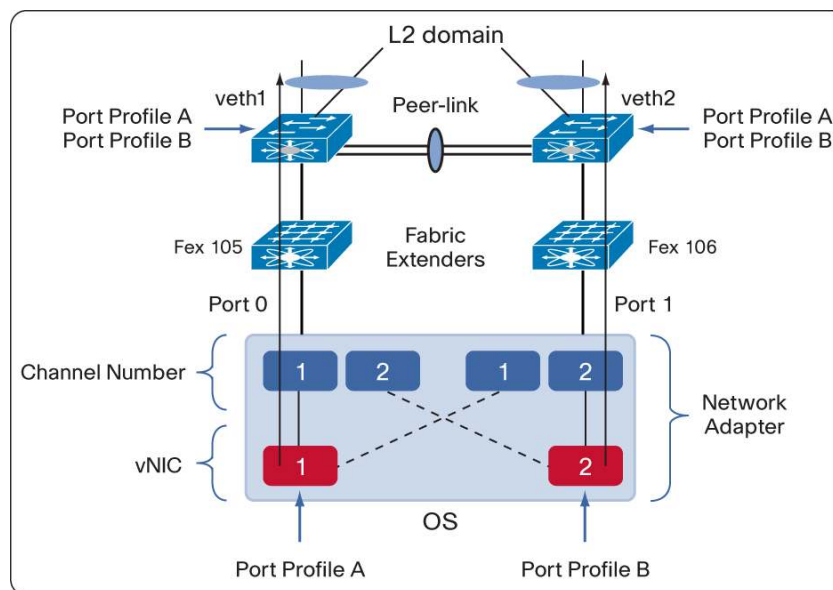
vPC with FEX Straight-Through

Figure 23 illustrates the characteristics of an adapter-FEX topology with FEX Straight-Through mode. The main difference from the non-vPC topology is the configuration of the peer-link and the presence of the vPC domain configuration.

With a vPC design, you would still need to configure the port-profiles type vethernet on both Cisco Nexus 5500 Switches, but the veth that is instantiated by a given vNIC would have the same number on both vPC peers.

The vNIC would still operate in Active/Standby mode with adapter failover. The use of vPC does not enable the vNICs of an A-FEX adapter to create a “port-channel.” For load distribution the port 0 and port 1 uplinks must be utilized by vNICs in a round-robin fashion: that is, half of the vNICs would use port 0 as the primary path, and the other half of the vNICs would use port 1 as their primary path.

Figure 23. Traffic Distribution from vNIC to Upstream Switches



The configuration on the left Cisco Nexus 5500 for the above topology would look like this:

```
interface Vethernet32769
  inherit port-profile NIC-VLAN50
  bind interface Ethernet105/1/2 channel 1
```

The configuration on the right Cisco Nexus 5500 for the above topology would look like this:

```
interface Vethernet32769
  inherit port-profile NIC-VLAN50
  bind interface Ethernet106/1/2 channel 1
```

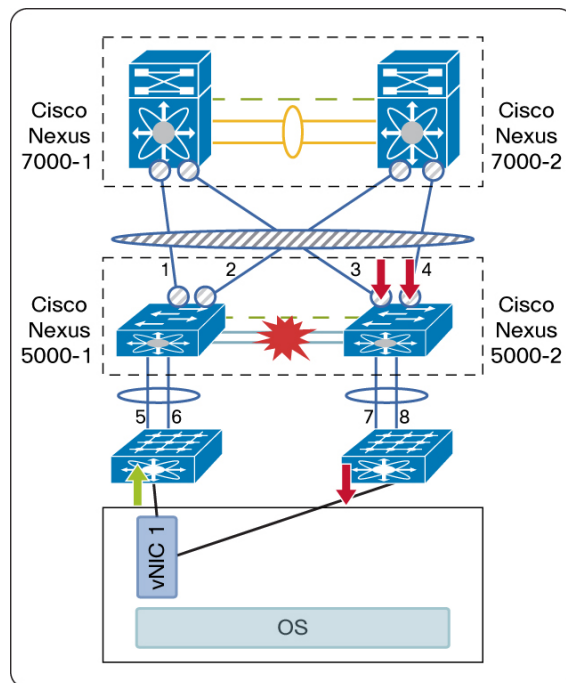
The MAC address table will show the MAC address of the vNIC and the veth interface associated with it only on the 5500 where the vNIC has the active path. Should the active path fail, the MAC address would then be programmed on the other vPC peer.

vPC Orphan Port Considerations with Layer 2 Uplinks

In case that you have orphan ports connected to this topology, you need to consider the use of the command **vpc orphan-ports suspend**. This command applies to non-veth ports only. For veth ports the “vpc orphan-ports” feature is transparently enabled.

As Figure 24 illustrates, when the peer-link is lost, the veth will failover automatically from the vPC secondary to the vPC primary.

Figure 24. Orphan Ports Failover When the vPC Peer-Link Fails



Servers that are single-homed to a FEX port or to the Cisco Nexus 5500 that is vPC secondary should instead be configured with the following option:

```
interface Ethernet1/1
  description to_vcenter_server
  spanning-tree port type edge
  vpc orphan-port suspend
```


The **vpc orphan-port suspend** option should be configured on all ports that are considered vPC orphan ports (minus the veth interfaces) on both the vPC primary and vPC secondary. This feature ensures that when the peer-link goes down, the vPC secondary shuts down these orphan ports, thus forcing the servers to use the path via the vPC primary.

SAN Connectivity with vPC and FEX Straight-Through Mode

The configuration to support FCoE in FEX Straight-Through mode doesn't require any particular explanation. Each Cisco Nexus 5500 Switch would be configured with a veth binding to the appropriate vHBA channel and a vFC binding to the veth.

From a SAN connectivity perspective, with FEX Straight-Through you need to define a static fixed veth on each Cisco Nexus 5500 and bind a vFC to the static fixed veth as shown in Figure 25.

Figure 25. Fibrechannel Connectivity with vPC and FEX Straight-Through



The relevant configuration for Cisco Nexus 5500-1 would look like this:

```
fex 105
  pinning max-links 1
  description "FEX0105"
  type N2232P
vlan 11
  fcoe vsan 11
vsan database
  vsan 11
  vsan 11 interface vfc11
interface vfc11
  bind interface Vethernet11
  no shutdown
vsan database
  vsan 11 interface vfc11
interface Vethernet11
  switchport mode trunk
  switchport trunk allowed vlan 1,11
  bind interface Ethernet105/1/2 channel 4
```

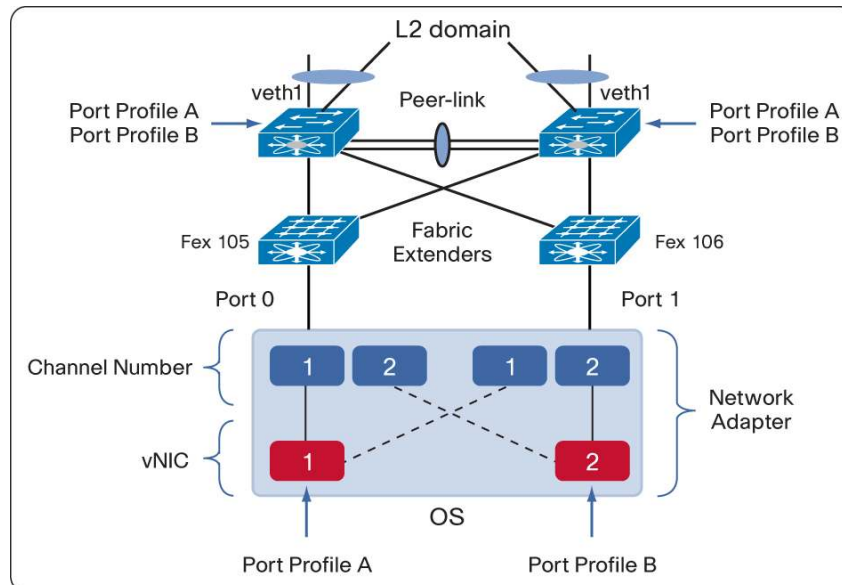
The relevant configuration for Cisco Nexus 5500-2 would look like this:

```
fex 106
  pinning max-links 1
  description "FEX0106"
  type N2232P
vlan 12
  fcoe vsan 12
vsan database
  vsan 12
  vsan 12 interface vfc12
interface vfc12
  bind interface Vethernet12
  no shutdown
vsan database
  vsan 12 interface vfc12
interface Vethernet12
  switchport mode trunk
  switchport trunk allowed vlan 1,12
  bind interface Ethernet106/1/2 channel 3
```

vPC with FEX Active/Active

Figure 26 illustrates the topology for adapter-FEX in a FEX Active/Active topology:

Figure 26. Topology with vPC and FEX Active/Active



The port-profiles type vethernet would have to be programmed on both vPC peers.

The FEX interfaces would have to be configured on both Cisco Nexus 5500 Switches to be in VN-Tag mode. For instance, interface Eth105/1/2 would be configured on both Cisco Nexus 5500-1 and Nexus 5500-2.

The associated configuration for this topology would look like this and would appear on both Cisco Nexus 5500 Switches:

```
interface Vethernet32769
  inherit port-profile NIC-VLAN50
  bind interface Ethernet105/1/2 channel 1
  bind interface Ethernet106/1/2 channel 1
```

Given that the vNIC operates in Active/Standby mode, only one of the two bindings would be active at any given time.

In order to understand which interface is binding where, you can use the command **show interface vethernet <number> detail** and pay attention to the status "active" or "standby":

```
nexus5500-2# show int vethernet 32769 detail
vif_index: 77
-----
veth is bound to interface Ethernet106/1/2 (0x1f690040)
priority: 0
vntag: 2
status: active
channel id: 1
registered mac info:
  vlan 0 - mac 01:00:5e:00:00:01
  vlan 0 - mac 01:00:5e:00:00:fc
  vlan 0 - mac 33:33:00:00:00:01
  vlan 0 - mac 33:33:00:01:00:03
  vlan 0 - mac 33:33:ff:a4:39:0a
  vlan 0 - mac 58:8d:09:0e:f8:5a
  vlan 0 - mac ff:ff:ff:ff:ff:ff

vif_index: 82
-----
veth is bound to interface Ethernet105/1/2 (0x1f680040)
priority: 0
vntag: 2
status: standby
channel id: 1
registered mac info:
  vlan 0 - mac 01:00:5e:00:00:01
  vlan 0 - mac 01:00:5e:00:00:fc
  vlan 0 - mac 33:33:00:00:00:01
  vlan 0 - mac 33:33:00:01:00:03
  vlan 0 - mac 33:33:ff:a4:39:0a
  vlan 0 - mac 58:8d:09:0e:f8:5a
  vlan 0 - mac ff:ff:ff:ff:ff:ff
```

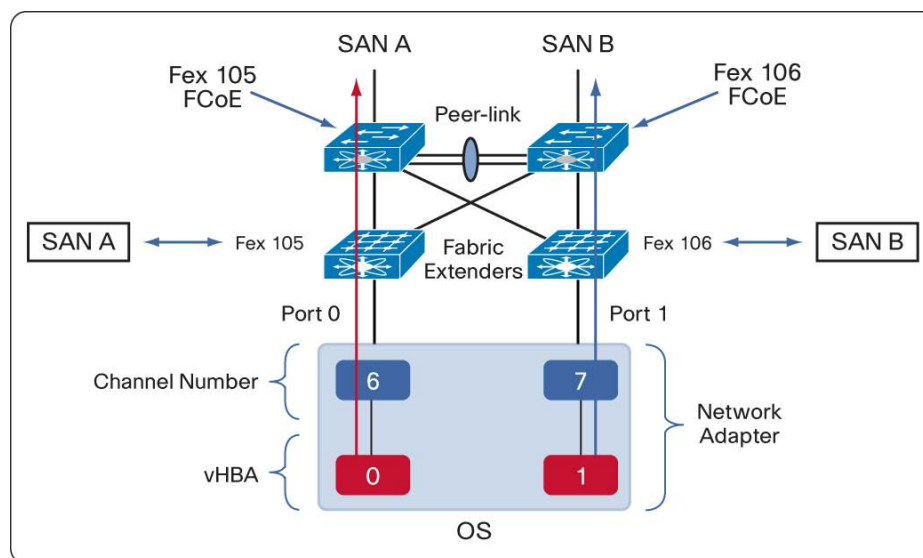
Differently from the FEX Straight-Through topology, the MAC address of the vNIC would be present in both Cisco Nexus 5500 MAC address tables and it will be associated with the same vethernet interface.

SAN Connectivity with FEX Active/Active

The FCoE configuration with FEX Active/Active requires some additional explanation. Each FEX in this topology is attached to both Cisco Nexus 5500 Switches, and so it is potentially connected to both SAN fabrics.

In order to ensure separation of SAN fabrics, each FEX belongs to just one of the two fabrics. This is achieved by typing **fcoe** under the FEX configuration on one of the two Cisco Nexus 5500 Switches (Figure 27).

Figure 27. Fibrechannel Connectivity with FEX A/A Topologies



For instance, in the example shown in Figure 27:

- You can assign FEX 105 to Fabric 1 via Nexus 5500-1 by typing **fcoe** under fex 105 in the configuration for Cisco Nexus 5500-1.
- You can assign FEX 106 to Fabric 2 via Nexus 5500-2 by typing **fcoe** under fex 106 in the configuration for Cisco Nexus 5500-2.

Note: In addition to this, the VLAN that is used for FCoE purposes needs to be defined on both Cisco Nexus 5500 Switches. For instance, if Nexus 5500-1 uses VLAN 11 for FCoE, this VLAN must be created on Nexus 5500-2. Similarly, if Nexus 5500-2 uses VLAN 12 for FCoE, this VLAN must also be created on Nexus 5500-1.

As an example the relevant configuration on Cisco Nexus 5500-1 is as follows:

```
fex 105
  pinning max-links 1
  description "FEX0105"
  type N2232P
  fcoe
fex 106
  pinning max-links 1
  description "FEX0106"
  type N2232P
vlan 11
  fcoe vsan 11
```

```
! vlan 12 is used on the Nexus 5500-2 but must be created
vlan 12
vsan database
    vsan 11
    vsan 11 interface vfc11
interface vfc11
    bind interface Vethernet11
    no shutdown
vsan database
    vsan 11 interface vfc11
interface Vethernet11
    switchport mode trunk
    switchport trunk allowed vlan 1,11
    bind interface Ethernet105/1/2 channel 4
```

The relevant configuration on Cisco Nexus 5500-2 would look like this:

```
fex 105
    pinning max-links 1
    description "FEX0105"
    type N2232P
fex 106
    pinning max-links 1
    description "FEX0106"
    type N2232P
    fcoe
vlan 12
    fcoe vsan 12
! vlan 11 is used on the Nexus 5500-1 but must be created
vlan 11
vsan database
    vsan 12
    vsan 12 interface vfc12
interface vfc12
    bind interface Vethernet12
    no shutdown
vsan database
    vsan 12 interface vfc12
interface Vethernet12
    switchport mode trunk
    switchport trunk allowed vlan 1,12
    bind interface Ethernet106/1/2 channel 3
```

Routing Considerations

If you are using the Cisco Nexus 5500 system for routing, the default gateway for the server is going to be an SVI on the Cisco Nexus 5500. Usual best practices for vPC and SVIs apply. You will configure HSRP gateway as follows:

```
interface Vlan50
  no shutdown
  description server-vlan
  ip address 10.50.1.2/24
  ip ospf passive-interface
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode
  hsrp 1
    preempt
    priority 110
    ip 10.50.1.1
```

Do not forget to install the BASE license; otherwise routing won't work:

```
Licenses: LAN_BASE_SERVICES INCREMENT
```

In addition, you have to remember to configure the **peer-gateway** in the vPC: configuration as follows:

```
vpc domain 2
  role priority 90
  peer-keepalive destination 10.51.35.17 source 10.51.35.18 interval 500 timeout
3
  delay restore 180
  peer-gateway
  auto-recovery
```

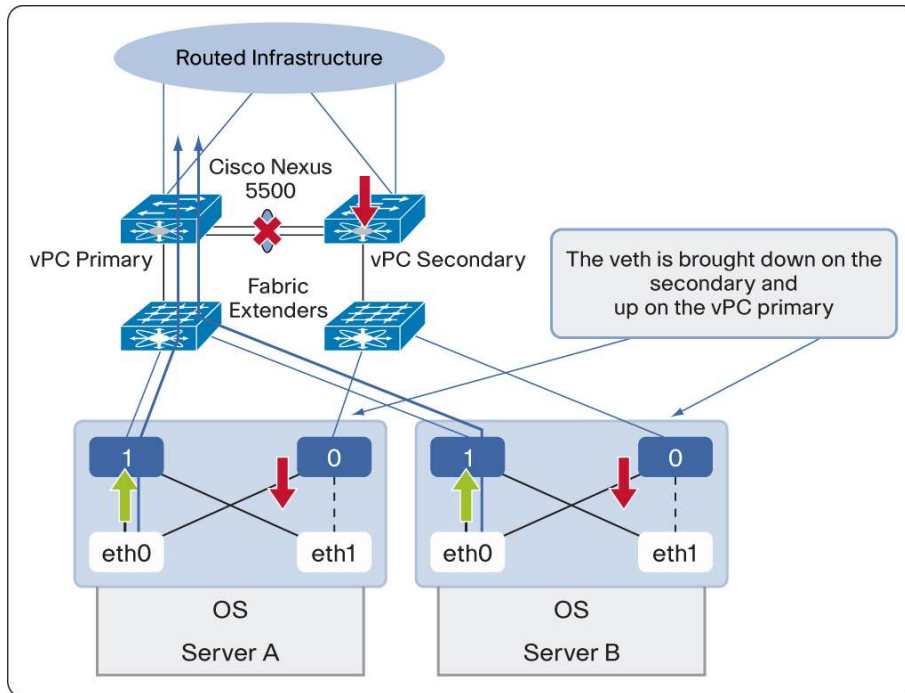
Finally, for all orphan ports (on the Cisco Nexus 5500 or on the FEX) that are not veth, you need to configure the following:

```
interface Ethernet1/1
  description to_vcenter_server
  spanning-tree port type edge
  vpc orphan-port suspend
```

This configuration applies to both the vPC primary and the vPC secondary.

Figure 28 illustrates the fact that when the peer-link goes down, the vPC secondary brings down the SVI. This assumes that all of the above configurations are in place (license, peer-gateway, and **vpc orphan-port suspend** on non-veth orphan ports). As a result, all vNICs will failover to the port that leads to the vPC primary device.

Figure 28. Failover Scenario with a Routed Access Topology



Sample Configuration Steps

The following configurations illustrate how to configure Adapter-FEX on the Cisco Nexus 5500 Switch.

Verify Licensing Requirements

No A-FEX license is required for A-FEX to function.

If using FC or FCoE, make sure the appropriate license is installed.

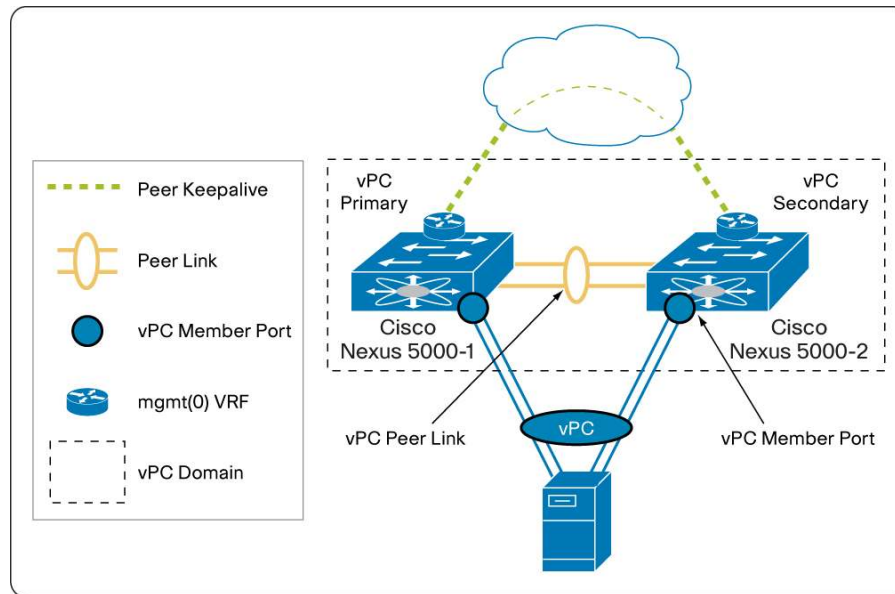
If you are using the Layer 3 card, make sure to install the BASE license.

Configure vPC

Remember that in A-FEX configurations you need to use vPC for the purpose of synchronizing the veth database between the Cisco Nexus 5500 Switches. You don't have to create vPC port-channels, but the vPC infrastructure still is needed for A-FEX.

Figure 29 shows the components of a Cisco Nexus 5000 Series vPC deployment.

Figure 29. vPC Components and Naming Convention



The following list provides a summary of vPC configuration best practices:

- The peer-keepalive is an out-of-band monitoring mechanism that is used for vPC peers to arbitrate roles and to resolve peer-link failures. You should configure the peer-keepalive connectivity either through the mgmt0 interface or a switch virtual interface (SVI) and a separate port. For example, each Cisco Nexus 5000 Series mgmt0 interface may already be connected to the management network, in which case the vPC configuration can simply use the existing connectivity for peer-keepalive purposes.
- The peer-keepalive traffic should never be carried in a VLAN over the peer link; such a configuration would make the peer keepalive useless.
- Make sure to use the **peer-gateway** command.
- Configure **vpc orphan-port suspend** on host orphan ports on both the vPC primary and secondary except for the veth interfaces.

```
feature vpc
vpc domain 2
    role priority 100
    peer-keepalive destination 10.51.35.17 source 10.51.35.18 interval 500 timeout
3
    delay restore 180
    peer-gateway
    auto-recovery

interface Ethernet1/25
    switchport mode trunk
    channel-group 10 mode active

interface Ethernet1/26
```



```
switchport mode trunk
channel-group 10 mode active

interface port-channel10
description 5k1-to-5k2
switchport mode trunk
spanning-tree port type network
vpc peer-link

interface <orphan ports>
vpc orphan-port suspend
```

Configure Fabric Extenders (if needed)

When configuring fabric extenders with the Cisco Nexus 5500, you should use FEX pre-provisioning. FEX preprovisioning is a feature that was introduced in Cisco NX-OS Software Release 5.0(2)N1(1).

```
slot 105
provision model N2K-C2232P
interface Ethernet <a/b>
switchport mode fex-fabric
channel-group 105
interface port-channel105
switchport mode fex-fabric
fex associate 105
```

In FEX Active/Active mode, for SAN connectivity purposes you need to decide which FEX belongs to which SAN fabric. So if FEX 105 should send SAN traffic to the Cisco Nexus 5500 Switch that you are configuring, you need to type **fcoe** under the FEX configuration:

```
fex 105
pinning max-links 1
description "FEX0105"
type N2232P
fcoe
```

Enable the A-FEX Feature

```
install feature-set virtualization
feature-set virtualization
vethernet auto-create
```

Configure All Switchports Connected to VIC-Capable Adapters

```
nexus5500-1(config)# int eth105/1/2
nexus5500-1(config-if)# switchport mode vntag
nexus5500-1(config-if)# int eth105/1/8
nexus5500-1(config-if)# switchport mode vntag
```

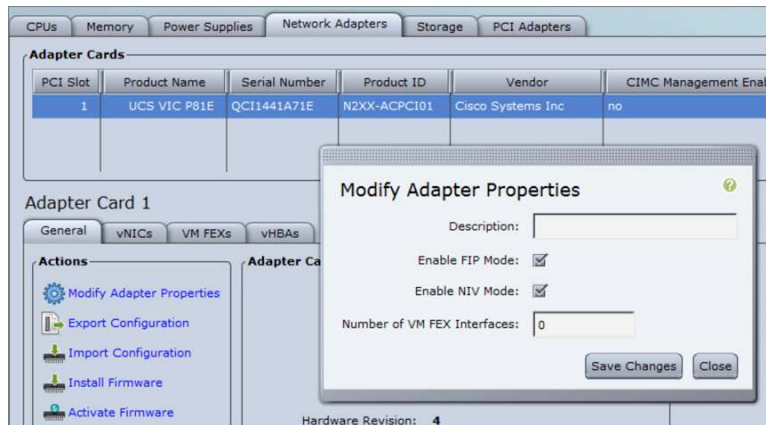
If using FEX Active/Active mode, remember that the above configurations need to be repeated on both Cisco Nexus 5500 Switches.

Configure VIC Adapters to Operate in A-FEX Mode

From CIMC, go to Inventory>Network Adapters>Modify Adapter Properties.

Select Enable NIV Mode as shown in Figure 30:

Figure 30. Enabling A-FEX by Selecting the NIV Option



After configuring NIV mode, reload the server (hence the adapter too) for this change to take effect.

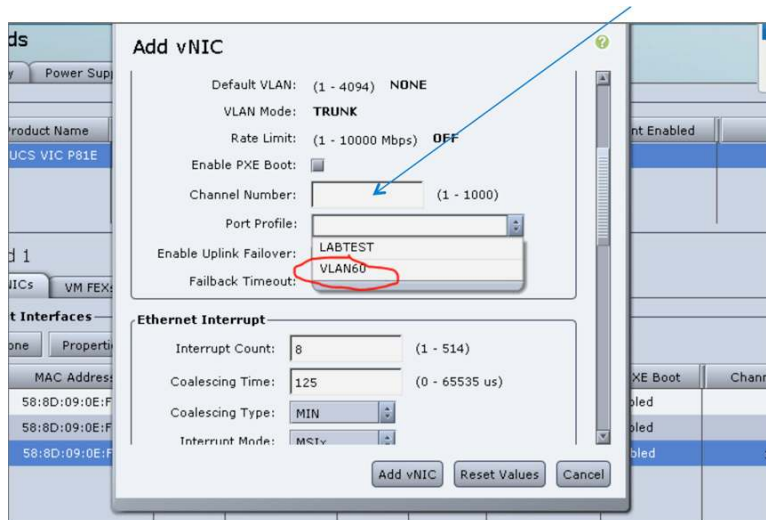
Configure Port-Profiles on Both vPC Peers

For redundancy purposes, you need to configure the same port-profile name on both vPC peers:

```
port-profile type vethernet NIC-VLAN50
switchport access vlan 50
switchport mode access
state enabled
```

From the server, you should associate the vNIC with the port-profile, as shown in Figure 31:

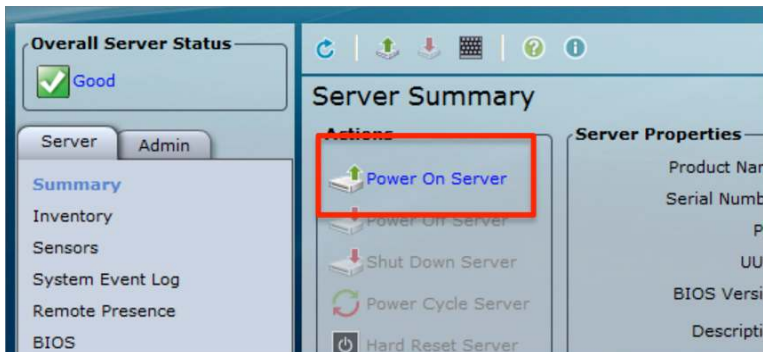
Figure 31. Selection of the Port-Profile for the vNIC



After completing the configuration, you need to reload the server in order for the operating system to recognize the new PCI device and for the network adapter to send a VIF create message to the Cisco Nexus 5500.

So shut down the server then power it up (Figure 32):

Figure 32. CIMC Options to Power Up or Down the Server



Now, on the Cisco Nexus 5500, this configuration will automatically appear:

```
interface Vethernet32769
  inherit port-profile NIC-VLAN50
  bind interface Ethernet105/1/2 channel 1
```

Configure the FC Connectivity to the SAN with Unified Ports

Enable Fibre Channel features:

```
nexus5500-1(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
Warning: Ensure class-fcoe is included in qos policy-maps of all types
```

If you use Unified Ports you need to define the range of ports that operate as Fibre Channel. For instance in the setup that was used to validate this design guide:

```
nexus5500-1# show mod
```

Mod	Ports	Module-Type	Model	Status
1	32	02 32X10GE/Modular Supervisor	N5K-C5548P-SUP	active *
2	16	02 16 port flexible GEM	N55-M16UP	ok

```
slot 2
port 1-16 type fc
```

After changing the ports from Ethernet to Fibre Channel, module 2 must be reloaded:

```
poweroff module 2
no poweroff module 2
```

VSAN and Port VSAN configuration:

```
vsan database
  vsan 11
vsan database
  vsan 11 interface fc2/11
vsan database
```

```
vsan 11
! define alias for WWNs
device-alias database
    device-alias name tc-sataboy01-port1 pwwn 50:00:40:20:01:f4:23:18
    device-alias name tc-sataboy01-port2 pwwn 50:00:40:21:01:f4:23:18

device-alias commit
```

Connect the Fibre Channel interface to the Fibre Channel core:

```
interface fc2/11
    no shutdown
```

Configure the FCoE connectivity from the A-FEX Adapter

Fibre Channel connectivity from the host vHBAs to the SAN is defined via FCoE. For this reason, you should configure FCoE policy-maps to implement priority flow control and proper bandwidth allocation for the Fibre Channel traffic type.

The following policy-maps already exist, but they may not be mapped:

```
policy-map type qos fcoe-default-in-policy
policy-map type queuing fcoe-default-in-policy
policy-map type queuing fcoe-default-out-policy
policy-map type network-qos fcoe-default-nq-policy
```

So in order to enable FCoE, you can simply map them to the **system-qos** (unless they already are):

```
system qos
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type qos input fcoe-default-in-policy
    service-policy type network-qos fcoe-default-nq-policy
```

Define VLAN-to-VSAN mapping:

```
nexus5500-1(config)# vlan 11
nexus5500-1(config-vlan)# fcoe vsan 11
nexus5500-2(config)# vlan 12
nexus5500-2(config-vlan)# fcoe vsan 12
```

Bind veth to the vHBA channel and vFC to veth:

```
interface veth 11
    switchport mode trunk
    switchport trunk allowed vlan 1,11
    bind interface eth 105/1/2 channel 4
    no shut
interface vfc11
    bind interface Vethernet11
    no shutdown
nexus5500-1(config)# vsan database
nexus5500-1(config-vsan-db)# vsan 11 interface vfc11
```

Notice that FIP VLAN discovery is not supported by Linux or ESX servers, so you need to configure the FCoE VLAN information on the vHBA from CIMC itself.

If you are using FEX Active/Active mode, remember that you need to assign each FEX to either fabric by using the **fcoe** keyword under the **fex <number>** configuration within the Cisco Nexus 5500 that connects to a given fabric.

Configure Zoning

The following steps in defining zoning are no different from regular Fibre Channel deployments:

1. Add a zone that includes the Port WWN defined above.
2. Activate the zoneset.
3. Optionally, you can copy the zoneset to the local configuration file: **zone copy active-zoneset full-zoneset vsan <number>**

You can configure zoning and LUN masking on the disk arrays just as if the vHBA were a physical HBA.

First locate the WWN:

```
Nexus5500-1# show flogi database vsan 11
-----
INTERFACE          VSAN    FCID          PORT NAME          NODE NAME
-----
vfc11               11      0xa00000      20:00:58:8d:09:0e:f8:5d 10:00:58:8d:09:0e:f8:5d
```

Normally, the default policy for a zone is a **deny**; if it isn't, you can change it from **permit** to **deny** (do not forget to commit after the change when running enhanced zoning):

```
nexus5500-1(config)# no zone default-zone permit vsan 11
Enhanced zone session has been created. Please 'commit' the changes when done.
nexus5500-1(config)# zone commit vsan 11
nexus5500-2(config)# no zone default-zone permit vsan 12
```

Verify which zones are active:

```
nexus5500-1# show zone status
VSAN: 11 default-zone: deny distribute: full Interop: default
mode: enhanced merge-control: allow
session: none
hard-zoning: enabled broadcast: enabled
[...]
```

```
nexus5500-1# show zone active
```

The zoning can be defined in any of the Fibre Channel switches, and it will be propagated. Initially, you may want to get the existing zoneset database if there's already an existing one:

```
nexus5500-1# zone copy active-zoneset full-zoneset
```

You may then want to check the **fcns** database to find the WWN of the target:

```
nexus5500-1# show fcns database vsan 11
```

```
VSAN 11:
```

```
-----  
FCID          TYPE  PWWN                                (VENDOR)          FC4-TYPE:FEATURE  
-----  
[...]  
0x730004      N      50:00:40:20:01:f4:23:18              scsi-fcp:target  
[tc-sataboy01-port1]  
0x730008      N      50:00:40:20:03:f4:22:64              scsi-fcp:target  
[...]  
0xa00000      N      20:00:58:8d:09:0e:f8:5d              scsi-fcp:init fc-gs
```

You would then add a zone to allow the vHBA to see the remote target:

```
zone name tc-esx02 vsan 11  
  member pwwn 20:00:58:8d:09:0e:f8:5d  
  member pwwn 50:00:40:20:01:f4:23:18
```

After adding this zone to the zoneset, you need to activate the zoneset:

```
nexus5500-1(config)# zoneset activate name tc-esx
```

After doing this, you need to check the zone status:

```
nexus5500-2# show zone status  
[...]  
Active Zoning Database :  
  DB size: 732 bytes  
  Name: tc-esx Zonesets:1 Zones:15  
  Status: Activation completed at 07:06:03 UTC Sep 27 2011
```

Make sure that after making changes, you commit the configuration; otherwise, configuration sessions will be locked on other switches in the fabric. The **session** field tells you which device is changing the zone configuration. Finally, the **Active Zoning Database** tells you which zone is active.

Sample Configurations

vPC and FEX Straight-Through with Routed Access

Cisco Nexus 5500 Switch 1

```
version 5.1(3)N1(1)
feature fcoe
install feature-set virtualization
feature-set virtualization
hostname nexus5500-1
feature telnet
cfs eth distribute
feature ospf
feature pim
feature private-vlan
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp
feature fex

no ip domain-lookup
service unsupported-transceiver
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
system qos
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type qos input fcoe-default-in-policy
    service-policy type network-qos fcoe-default-nq-policy
fex 105
    pinning max-links 1
    description "FEX0105"
    type N2232P
    fcoe
```

```
slot 2
    port 1-16 type fc
slot 105
    provision model N2K-C2232P
vrf context management
    ip route 0.0.0.0/0 10.51.35.1
vlan configuration 1,3,10,50,60
vlan 1
vlan 3
    name l3-vlan
vlan 10
vlan 11
    fcoe vsan 11
vlan 50,60
vlan 100
    name backup-vlan
spanning-tree mode mst
spanning-tree pathcost method long
spanning-tree mst configuration
    name dc1
    revision 3
    instance 1 vlan 1-4093
vpc domain 2
    role priority 100
    peer-keepalive destination 10.51.35.17 source 10.51.35.18 interval 500 timeout
3
    delay restore 180
    peer-gateway
    auto-recovery
port-profile type vethernet NIC-VLAN50
    switchport access vlan 50
    state enabled
port-profile type vethernet NIC-VLAN60
    switchport access vlan 60
    state enabled
port-profile type vethernet BACKUP-VLAN
    switchport access vlan 100
    state enabled
vsan database
    vsan 11
interface Vlan3
    no shutdown
    description l3vlan
    ip address 10.50.3.10/31
    ip ospf cost 50
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode
```



```
hsrp 1

interface Vlan50
  no shutdown
  description server-vlan
  no ip redirects
  ip address 10.50.1.2/24
  ip ospf passive-interface
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode
  hsrp 1
    preempt
    priority 110
    ip 10.50.1.1

interface Vlan60
  no shutdown
  ip address 10.60.1.2/24
  ip ospf passive-interface
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode
  hsrp 1
    preempt
    priority 110
    ip 10.60.1.1

interface port-channel10
  description 5k1-to-5k2
  switchport mode trunk
  spanning-tree port type network
  vpc peer-link

interface port-channel21
  description to-7k1
  no switchport
  lacp suspend-individual
  ip address 10.50.3.1/31
  ip ospf network point-to-point
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode

interface port-channel23
  description to-7k2
  no switchport
  ip address 10.50.3.5/31
  ip ospf network point-to-point
  ip router ospf 1 area 1.1.1.1
```

```
ip pim sparse-mode

interface port-channel105
    switchport mode fex-fabric
    fex associate 105

interface vfc11
    bind interface Vethernet11
    no shutdown
interface vfc112
    bind interface Vethernet112
    no shutdown
vsan database
    vsan 11 interface vfc11
    vsan 11 interface vfc112
interface fc2/11
    switchport trunk allowed vsan 1
    switchport trunk allowed vsan add 11
    no shutdown
interface Ethernet1/5
    description to_fex_105
    switchport mode fex-fabric
    fex associate 105
    channel-group 105
interface Ethernet1/19
    no switchport
    channel-group 21 mode active
interface Ethernet1/20
    no switchport
    channel-group 21 mode active
interface Ethernet1/21
    no switchport
    channel-group 23 mode active
interface Ethernet1/22
    no switchport
    channel-group 23 mode active
interface Ethernet1/25
    switchport mode trunk
    channel-group 10 mode active
interface Ethernet1/26
    switchport mode trunk
    channel-group 10 mode active
interface mgmt0
    ip address 10.51.35.18/27
interface loopback0
    ip address 128.0.0.4/32
    ip router ospf 1 area 1.1.1.1
```

```
ip pim sparse-mode
interface Ethernet105/1/2
    description to_esx02
    switchport mode vntag
interface Ethernet105/1/8
    description to_esx08
    switchport mode vntag
interface Vethernet11
    switchport mode trunk
    switchport trunk allowed vlan 1,11
    bind interface Ethernet105/1/2 channel 4

interface Vethernet112
    switchport mode trunk
    switchport trunk allowed vlan 11
    bind interface Ethernet105/1/8 channel 4

interface Vethernet32769
    inherit port-profile NIC-VLAN50
    bind interface Ethernet105/1/2 channel 1

interface Vethernet32770
    inherit port-profile NIC-VLAN60
    bind interface Ethernet105/1/2 channel 2

interface Vethernet32771
    inherit port-profile NIC-VLAN50
    bind interface Ethernet105/1/8 channel 1

interface Vethernet32772
    inherit port-profile NIC-VLAN60
    bind interface Ethernet105/1/8 channel 2

interface Vethernet32773
    inherit port-profile BACKUP-VLAN
    bind interface Ethernet105/1/8 channel 100
router ospf 1
    area 1.1.1.1 stub no-summary
    auto-cost reference-bandwidth 1000000
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen
ip pim pre-build-spt
no ip igmp snooping mrouter vpc-peer-link
vethernet auto-create
vpc bind-vrf default vlan 1003
zone mode enhanced vsan 11
zoneset activate name tc-esx vsan 11
```

```
zone name tc-esx02 vsan 11
    member pwnn 20:00:58:8d:09:0e:f8:5d
    member pwnn 50:00:40:20:01:f4:23:18
!
[tc-sataboy01-port1]
zone name tc-esx08 vsan 11
    member pwnn 20:00:58:8d:09:0e:ed:39
    member pwnn 50:00:40:20:01:f4:23:18
!
[tc-sataboy01-port1]
zoneset name tc-esx vsan 11
    member tc-esx02
    member tc-esx08
zone commit vsan 11
```

Cisco Nexus 5500 Switch 2

```
version 5.1(3)N1(1)
feature fcoe
install feature-set virtualization
feature-set virtualization
hostname nexus5500-2
feature telnet
cfs eth distribute
feature ospf
feature pim
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp
feature fex

class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
system qos
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type qos input fcoe-default-in-policy
```

```
service-policy type network-qos fcoe-default-nq-policy
fex 106
    pinning max-links 1
    description "FEX0106"
    type N2232P
slot 106
    provision model N2K-C2232P
vrf context management
    ip route 0.0.0.0/0 10.51.35.1
vlan 3
    name l3-vlan
vlan 12
    fcoe vsan 12
vlan 50
    name 10.50.1.x
vlan 60
    name 10.50.2.x
vlan 100
    name backup-vlan
spanning-tree mode mst
spanning-tree pathcost method long
spanning-tree mst configuration
    name dc1
    revision 3
    instance 1 vlan 1-4093
vpc domain 2
    role priority 110
    peer-keepalive destination 10.51.35.18 source 10.51.35.17 interval 500 timeout
3
    delay restore 180
    peer-gateway
    auto-recovery
port-profile type vethernet NIC-VLAN50
    switchport access vlan 50
    state enabled
port-profile type vethernet NIC-VLAN60
    switchport access vlan 60
    state enabled
port-profile type vethernet BACKUP-VLAN
    switchport access vlan 100
    state enabled
vsan database
    vsan 12

interface Vlan3
    no shutdown
    description l3-vlan
```

```
ip address 10.50.3.11/31
ip ospf cost 50
ip router ospf 1 area 1.1.1.1
ip pim sparse-mode

interface Vlan50
  no shutdown
  description server-vlan
  no ip redirects
  ip address 10.50.1.3/24
  ip ospf passive-interface
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode
  hsrp 1
    preempt
    ip 10.50.1.1

interface Vlan60
  no shutdown
  ip address 10.60.1.3/24
  ip ospf passive-interface
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode
  hsrp 1
    preempt
    ip 10.60.1.1

interface port-channel10
  switchport mode trunk
  spanning-tree port type network
  vpc peer-link

interface port-channel22
  description to-7k1
  no switchport
  ip address 10.50.3.3/31
  ip ospf network point-to-point
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode

interface port-channel24
  description to-7k2
  no switchport
  ip address 10.50.3.7/31
  ip ospf network point-to-point
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode
```

```
interface port-channel106
    switchport mode fex-fabric
    fex associate 106

interface vfc12
    bind interface Vethernet12
    no shutdown

interface vfc122
    bind interface Vethernet122
    no shutdown
vsan database
    vsan 12 interface vfc12
    vsan 12 interface vfc122

interface fc2/11
    switchport trunk allowed vsan 12
    no shutdown

interface Ethernet1/6
    description to_fex_106
    switchport mode fex-fabric
    fex associate 106
    channel-group 106

interface Ethernet1/19
    no switchport
    channel-group 22 mode active

interface Ethernet1/20
    no switchport
    channel-group 22 mode active

interface Ethernet1/21
    no switchport
    channel-group 24 mode active

interface Ethernet1/22
    no switchport
    channel-group 24 mode active

interface Ethernet1/25
    switchport mode trunk
    channel-group 10 mode active

interface Ethernet1/26
```

```
switchport mode trunk
channel-group 10 mode active

interface mgmt0
  ip address 10.51.35.17/27

interface loopback0
  ip address 128.0.0.5/32
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode

interface Ethernet106/1/2
  description to_tc-esx02
  switchport mode vntag

interface Ethernet106/1/8
  description to_tc-esx08
  switchport mode vntag

interface Vethernet12
  switchport mode trunk
  switchport trunk allowed vlan 1,12
  bind interface Ethernet106/1/2 channel 3

interface Vethernet122
  switchport mode trunk
  switchport trunk allowed vlan 12
  bind interface Ethernet106/1/8 channel 3

interface Vethernet32769
  inherit port-profile NIC-VLAN50
  bind interface Ethernet106/1/2 channel 1

interface Vethernet32770
  inherit port-profile NIC-VLAN60
  bind interface Ethernet106/1/2 channel 2

interface Vethernet32771
  inherit port-profile NIC-VLAN50
  bind interface Ethernet106/1/8 channel 1

interface Vethernet32772
  inherit port-profile NIC-VLAN60
  bind interface Ethernet106/1/8 channel 2

interface Vethernet32773
  inherit port-profile BACKUP-VLAN
```



```
bind interface Ethernet106/1/8 channel 100

router ospf 1
  area 1.1.1.1 stub no-summary
  auto-cost reference-bandwidth 1000000
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen
ip pim pre-build-spt
no ip igmp snooping mrouter vpc-peer-link
vethernet auto-create
vpc bind-vrf default vlan 1003

!Full Zone Database Section for vsan 12
zone name tc-esx02 vsan 12
  member pwnn 20:00:58:8d:09:0e:f8:5c
  member pwnn 50:00:40:21:01:f4:23:18

zone name tc-esx08 vsan 12
  member pwnn 20:00:58:8d:09:0e:ed:38
  member pwnn 50:00:40:21:01:f4:23:18

zoneset name tc-esx vsan 12
  member tc-esx02
  member tc-esx08

zoneset activate name tc-esx vsan 12
```

vPC and FEX Active/Active with Routed Access

Cisco Nexus 5500 Switch 1

```
version 5.1(3)N1(1)
feature fcoe
install feature-set virtualization
feature-set virtualization
hostname nexus5500-1
feature telnet
cfs eth distribute
feature ospf
feature pim
feature private-vlan
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp
feature fex
```

```
service unsupported-transceiver
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
system qos
    service-policy type queuing input fcoe-default-in-policy
    service-policy type queuing output fcoe-default-out-policy
    service-policy type qos input fcoe-default-in-policy
    service-policy type network-qos fcoe-default-nq-policy
fex 105
    pinning max-links 1
    description "FEX0105"
    type N2232P
    fcoe
fex 106
    pinning max-links 1
    description "FEX0106"
    type N2232P
slot 2
    port 1-16 type fc
slot 105
    provision model N2K-C2232P
slot 106
    provision model N2K-C2232P
vrf context management
    ip route 0.0.0.0/0 10.51.35.1
vlan configuration 1,3,10,50,60
vlan 1
vlan 3
    name l3-vlan
vlan 10
vlan 11
    fcoe vsan 11
vlan 12,50,60
vlan 100
    name backup-vlan
```

```
spanning-tree mode mst
spanning-tree pathcost method long
spanning-tree mst configuration
    name dc1
    revision 3
    instance 1 vlan 1-4093
vpc domain 2
    role priority 90
    peer-keepalive destination 10.51.35.17 source 10.51.35.18 interval 500
timeout 3
    delay restore 180
    peer-gateway
    auto-recovery
port-profile type vethernet NIC-VLAN50
    switchport access vlan 50
    state enabled
port-profile type vethernet NIC-VLAN60
    switchport access vlan 60
    state enabled
port-profile type vethernet BACKUP-VLAN
    switchport access vlan 100
    state enabled
vsan database
    vsan 11

interface Vlan1

interface Vlan3
    no shutdown
    description l3vlan
    no ip redirects
    ip address 10.50.3.10/31
    ip ospf cost 50
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode
    hsrp 1

interface Vlan50
    no shutdown
    description server-vlan
    no ip redirects
    ip address 10.50.1.2/24
    ip ospf passive-interface
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode
    hsrp 1
    preempt
```

```
        priority 110
        ip 10.50.1.1

interface Vlan60
    no shutdown
    no ip redirects
    ip address 10.60.1.2/24
    ip ospf passive-interface
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode
    hsrp 1
        preempt
        priority 110
        ip 10.60.1.1

interface port-channel10
    description 5k1-to-5k2
    switchport mode trunk
    spanning-tree port type network
    vpc peer-link

interface port-channel21
    description to-7k1
    no switchport
    lacp suspend-individual
    ip address 10.50.3.1/31
    ip ospf network point-to-point
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode

interface port-channel23
    description to-7k2
    shutdown
    no switchport
    ip address 10.50.3.5/31
    ip ospf network point-to-point
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode

interface port-channel103
    switchport mode fex-fabric
    fex associate 103

interface port-channel105
    switchport mode fex-fabric
    fex associate 105
    vpc 105
```

```
interface port-channel106
    switchport mode fex-fabric
    fex associate 106
    vpc 106

interface vfc11
    bind interface Vethernet11
    no shutdown

interface vfc112
    bind interface Vethernet112
    no shutdown
vsan database
    vsan 11 interface vfc11
    vsan 11 interface vfc112

interface fc2/11
    switchport trunk allowed vsan 1
    switchport trunk allowed vsan add 11
    no shutdown

interface Ethernet1/5
    description to_fex_105
    switchport mode fex-fabric
    fex associate 105
    channel-group 105

interface Ethernet1/6
    description to_fex_106
    switchport mode fex-fabric
    fex associate 106
    channel-group 106

interface Ethernet1/19
    no switchport
    channel-group 21 mode active

interface Ethernet1/20
    no switchport
    channel-group 21 mode active

interface Ethernet1/21
    no switchport
    channel-group 23 mode active
```

```
interface Ethernet1/22
  no switchport
  channel-group 23 mode active

interface Ethernet1/25
  switchport mode trunk
  channel-group 10 mode active

interface Ethernet1/26
  switchport mode trunk
  channel-group 10 mode active

interface mgmt0
  ip address 10.51.35.18/27

interface loopback0
  ip address 128.0.0.4/32
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode

interface Ethernet105/1/2
  description to_esx02
  switchport mode vntag

interface Ethernet105/1/8
  description to_esx08
  switchport mode vntag

interface Ethernet106/1/2
  description to_tc-esx02
  switchport mode vntag

interface Ethernet106/1/8
  description to_tc-esx08
  switchport mode vntag

interface Vethernet11
  switchport mode trunk
  switchport trunk allowed vlan 1,11
  bind interface Ethernet105/1/2 channel 4

interface Vethernet112
  switchport mode trunk
  switchport trunk allowed vlan 11
  bind interface Ethernet105/1/8 channel 4
```

```

interface Vethernet32769
    inherit port-profile NIC-VLAN50
    bind interface Ethernet105/1/2 channel 1
    bind interface Ethernet106/1/2 channel 1

interface Vethernet32770
    inherit port-profile NIC-VLAN60
    bind interface Ethernet105/1/2 channel 2
    bind interface Ethernet106/1/2 channel 2

interface Vethernet32771
    inherit port-profile NIC-VLAN50
    bind interface Ethernet105/1/8 channel 1
    bind interface Ethernet106/1/8 channel 1

interface Vethernet32772
    inherit port-profile NIC-VLAN60
    bind interface Ethernet105/1/8 channel 2
    bind interface Ethernet106/1/8 channel 2

interface Vethernet32773
    inherit port-profile BACKUP-VLAN
    bind interface Ethernet105/1/8 channel 100
    bind interface Ethernet106/1/8 channel 100
    shutdown

interface Vethernet32774
    inherit port-profile BACKUP-VLAN
    bind interface Ethernet105/1/2 channel 12
    bind interface Ethernet106/1/2 channel 12
router ospf 1
    area 1.1.1.1 stub no-summary
    auto-cost reference-bandwidth 1000000
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen
ip pim pre-build-spt
no ip igmp snooping mrouter vpc-peer-link
vethernet auto-create
vpc bind-vrf default vlan 1003
zoneset name tc-esx vsan 11
    member tc-esx02
    member tc-esx07
    member tc-esx08

zoneset activate name tc-esx vsan 11
zone name tc-esx02 vsan 11
    member pwnn 20:00:58:8d:09:0e:f8:5d

```

```
        member pwnn 50:00:40:20:01:f4:23:18
!          [tc-sataboy01-port1]
```

```
zone name tc-esx08 vsan 11
        member pwnn 20:00:58:8d:09:0e:ed:39
        member pwnn 50:00:40:20:01:f4:23:18
!          [tc-sataboy01-port1]
```

```
zoneset name tc-esx vsan 11
        member tc-esx02
        member tc-esx08
```

```
zoneset name permitany vsan 11
```

```
zone commit vsan 11
```

Cisco Nexus 5500 Switch 2

```
version 5.1(3)N1(1)
feature fcoe
install feature-set virtualization
feature-set virtualization
hostname nexus5500-2
feature telnet
cfs eth distribute
feature ospf
feature pim
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp
feature fex
class-map type qos class-fcoe
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-all-flood
    match qos-group 2
class-map type queuing class-ip-multicast
    match qos-group 2
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
system qos
```



```
service-policy type queuing input fcoe-default-in-policy
service-policy type queuing output fcoe-default-out-policy
service-policy type qos input fcoe-default-in-policy
service-policy type network-qos fcoe-default-nq-policy
fex 105
  pinning max-links 1
  description "FEX0105"
  type N2232P
fex 106
  pinning max-links 1
  description "FEX0106"
  type N2232P
  fcoe
slot 105
  provision model N2K-C2232P
slot 106
  provision model N2K-C2232P
vrf context management
  ip route 0.0.0.0/0 10.51.35.1
vlan 1
vlan 3
  name l3-vlan
vlan 11
vlan 12
  fcoe vsan 12
vlan 50
  name 10.50.1.x
vlan 60
  name 10.50.2.x
vlan 100
  name backup-vlan
spanning-tree mode mst
spanning-tree pathcost method long
spanning-tree mst configuration
  name dc1
  revision 3
  instance 1 vlan 1-4093
vpc domain 2
  role priority 126
  peer-keepalive destination 10.51.35.18 source 10.51.35.17 interval 500 timeout
3
  peer-config-check-bypass
  delay restore 180
  peer-gateway
  auto-recovery
port-profile type vethernet NIC-VLAN50
  switchport access vlan 50
```

```
state enabled
port-profile type vethernet NIC-VLAN60
    switchport access vlan 60
    state enabled
port-profile type vethernet BACKUP-VLAN
    switchport access vlan 100
    state enabled
vsan database
    vsan 12
```

```
interface Vlan3
    no shutdown
    description l3-vlan
    no ip redirects
    ip address 10.50.3.11/31
    ip ospf cost 50
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode
```

```
interface Vlan50
    no shutdown
    description server-vlan
    no ip redirects
    ip address 10.50.1.3/24
    ip ospf passive-interface
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode
    hsrp 1
        preempt
        ip 10.50.1.1
```

```
interface Vlan60
    no shutdown
    no ip redirects
    ip address 10.60.1.3/24
    ip ospf passive-interface
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode
    hsrp 1
        preempt
        ip 10.60.1.1
```

```
interface port-channel10
    switchport mode trunk
    spanning-tree port type network
    vpc peer-link
```

```
interface port-channel22
  description to-7k1
  no switchport
  ip address 10.50.3.3/31
  ip ospf network point-to-point
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode

interface port-channel24
  description to-7k2
  no switchport
  ip address 10.50.3.7/31
  ip ospf network point-to-point
  ip router ospf 1 area 1.1.1.1
  ip pim sparse-mode

interface port-channel105
  switchport mode fex-fabric
  fex associate 105
  vpc 105

interface port-channel106
  switchport mode fex-fabric
  fex associate 106
  vpc 106

interface vfc12
  bind interface Vethernet12
  no shutdown

interface vfc122
  bind interface Vethernet122
  no shutdown
vsan database
  vsan 12 interface vfc12
  vsan 12 interface vfc122

interface fc2/11
  switchport trunk allowed vsan 12
  no shutdown

interface Ethernet1/5
  description to_fex_105
  switchport mode fex-fabric
  fex associate 105
  channel-group 105
```

```
interface Ethernet1/6
    description to_fex_106
    switchport mode fex-fabric
    fex associate 106
    channel-group 106

interface Ethernet1/19
    no switchport
    channel-group 22 mode active

interface Ethernet1/20
    no switchport
    channel-group 22 mode active

interface Ethernet1/21
    no switchport
    channel-group 24 mode active

interface Ethernet1/22
    no switchport
    channel-group 24 mode active

interface Ethernet1/25
    switchport mode trunk
    channel-group 10 mode active

interface Ethernet1/26
    switchport mode trunk
    channel-group 10 mode active

interface mgmt0
    ip address 10.51.35.17/27

interface loopback0
    ip address 128.0.0.5/32
    ip router ospf 1 area 1.1.1.1
    ip pim sparse-mode

interface Ethernet105/1/2
    description to_esx02
    switchport mode vntag

interface Ethernet105/1/8
    description to_esx08
    switchport mode vntag
```

```
interface Ethernet106/1/2
  description to_tc-esx02
  switchport mode vntag

interface Ethernet106/1/8
  description to_tc-esx08
  switchport mode vntag

interface Vethernet12
  switchport mode trunk
  switchport trunk allowed vlan 1,12
  bind interface Ethernet106/1/2 channel 3

interface Vethernet122
  switchport mode trunk
  switchport trunk allowed vlan 12
  bind interface Ethernet106/1/8 channel 3

interface Vethernet32769
  inherit port-profile NIC-VLAN50
  bind interface Ethernet105/1/2 channel 1
  bind interface Ethernet106/1/2 channel 1

interface Vethernet32770
  inherit port-profile NIC-VLAN60
  bind interface Ethernet105/1/2 channel 2
  bind interface Ethernet106/1/2 channel 2

interface Vethernet32771
  inherit port-profile NIC-VLAN50
  bind interface Ethernet105/1/8 channel 1
  bind interface Ethernet106/1/8 channel 1

interface Vethernet32772
  inherit port-profile NIC-VLAN60
  bind interface Ethernet105/1/8 channel 2
  bind interface Ethernet106/1/8 channel 2

interface Vethernet32773
  inherit port-profile BACKUP-VLAN
  bind interface Ethernet105/1/8 channel 100
  bind interface Ethernet106/1/8 channel 100
  shutdown

interface Vethernet32774
  inherit port-profile BACKUP-VLAN
  bind interface Ethernet105/1/2 channel 12
```

```
bind interface Ethernet106/1/2 channel 12

router ospf 1
  area 1.1.1.1 stub no-summary
  auto-cost reference-bandwidth 1000000
ip pim ssm range 232.0.0.0/8
ip pim auto-rp forward listen
ip pim pre-build-spt
no ip igmp snooping mrouter vpc-peer-link
vethernet auto-create
vpc bind-vrf default vlan 1003
zone name tc-esx02 vsan 12
  member pwn 20:00:58:8d:09:0e:f8:5c
  member pwn 50:00:40:21:01:f4:23:18

zone name tc-esx08 vsan 12
  member pwn 20:00:58:8d:09:0e:ed:38
  member pwn 50:00:40:21:01:f4:23:18

zoneset name tc-esx vsan 12
  member tc-esx02
  member tc-esx08

zoneset activate name tc-esx vsan 12
```



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)