


Deploy Layer 2 Disjoint Networks Upstream in End-Host Mode

What You Will Learn

As more customers deploy and expand their Cisco Unified Computing System™ (Cisco UCS™) installations, disjoint Layer 2 networks upstream of Cisco UCS are becoming more prominent. Figure 1 shows a common topology. The production, public, and backup networks are separate Layer 2 domains and do not have adjacencies with each other; however, Cisco UCS blades may need to access all three networks. Furthermore, in the case of a virtualized host, the host may include virtual machines that need access to all three networks.

Figure 1. Disjoint Layer 2 Networks Upstream



Note: The links depicted could be PortChannels or single links. This scenario assumes that there are no overlapping VLAN numbers in the public, production, and backup networks.

This document explains how to deploy disjoint Layer 2 networks upstream of Cisco UCS in Ethernet end-host mode. This feature is being introduced as part of Cisco UCS Manager Release 2.0 and is hardware independent: that is, this feature is supported on both the Cisco UCS 6100 and 6200 Series Fabric Interconnects.

End-Host Mode Overview

In end-host mode, Cisco UCS presents an end host to an external Ethernet network; the external LAN sees the Cisco UCS Fabric Interconnect as an end host with multiple adapters. Fabric interconnects running Cisco UCS Manager Release 1.4 or earlier operating in end-host mode follow certain forwarding rules for handling unicast and multicast and broadcast traffic.

End-host mode offers these main features:

- Spanning Tree Protocol is not run on both the uplink ports
- MAC address learning occurs only on the server ports and appliance ports
- MAC address aging is not supported; MAC address changes are fully supported
- Active-active links are used regardless of the number of uplink switches
- The solution is highly scalable because the control plane is not occupied
- All uplink ports connect to the same Layer 2 cloud

A single Ethernet uplink port (or PortChannel) on each fabric interconnect is chosen to be the broadcast and multicast traffic receiver for all VLANs, and incoming and broadcast traffic is dropped on the other uplinks. This port is called the G-pinned. G-pin port is the designated interface to receive multicast and broadcast in EHM and is selected by the system randomly and is not configurable. If the G-pinned port goes down, another uplink is designated automatically.

To view the elected G-pinned port, use the following commands:

```
UCS-A# connect nxos
UCS-A(nxos)# show platform software enm internal info global | grep -A 6 'Global
Params'

Other Global Params:
    broadcast-if 0x88c1f04(Ethernet1/16)
    multicast-if 0x88c1f04(Ethernet1/16)
    ip_multicast-if 0x88c1f04(Ethernet1/16)
    end-host-mode: Enabled
```

As a result of this behavior, the only network that will operate properly is the network to which the G-pinned port is connected. For example, if the G-pinned port is on the production network, any blade with a virtual network interface card (vNIC) in backup or public VLANs will experience problems. vNICs in those VLANs will not receive any Address Resolution Protocol (ARP) broadcast or multicast traffic from upstream.

For such a network scenario, the recommendation in Cisco UCS Manager Release 1.4 and earlier is to change to switch mode. In switch mode, spanning tree (Per-VLAN Spanning Tree Plus [PVRST+]) is run on uplinks, and broadcast and multicast traffic is handled and forwarded accordingly. Note that in switch mode, fabric interconnects operate like traditional Layer 2 switches, with spanning-tree and MAC address learning enabled on the uplinks.

Support for Disjoint Layer 2 Upstream in End-Host Mode

In Cisco UCS Manager Release 2.0, Cisco UCS provides the flexibility to deploy nondisjoint and disjoint Layer 2 upstream networks in end-host mode. Topological simplification is achieved with end-host mode without the need to turn on switch mode in Layer 2 disjoint network deployments.

Cisco UCS Manager Release 2.0 enables the following functions in end-host mode:

- Capability to selectively assign VLANs to uplinks
- Pinning decision based on uplink and vNIC VLAN membership
- Allocation of a designated broadcast and multicast traffic receiver for each VLAN rather than on a global basis

Layer 2 Disjoint Upstream Packet Forwarding in End-Host Mode

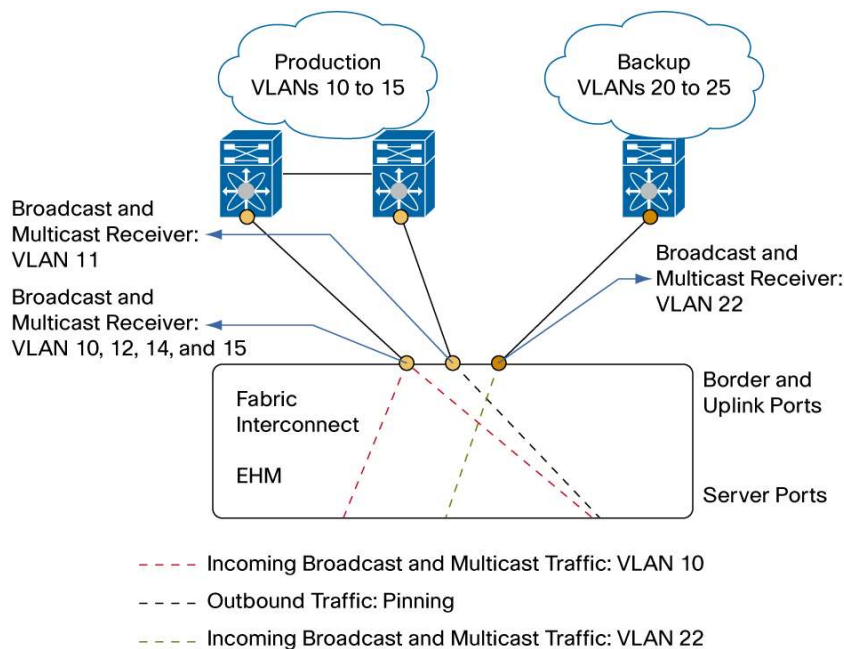
Server links (vNICs on the blades) are associated with a single uplink port (which may also be a PortChannel). This process is called pinning, and the selected external interface is called a pinned uplink port. The process of pinning can be statically configured (when the vNIC is defined), or dynamically configured by the system.

In Cisco UCS Manager Release 2.0, VLAN membership on uplinks is taken into account during the dynamic pinning process. VLANs assigned to a vNIC are used to find a matching uplink. **If all associated VLANs of a vNIC are not found on an uplink, pinning failure will occur.**

The traffic forwarding behavior differs from that in Cisco UCS Manager Release 1.4 and earlier in the way that the incoming broadcast and multicast traffic is handled. A designated receiver is chosen for each VLAN, rather than globally as in Cisco UCS Manager Release 1.4 and earlier.

Note: **The use of overlapping VLAN numbers in the upstream disjoint networks is not supported.**

Figure 2. Packet Flow through Fabric Interconnect in End-Host Mode with Disjoint Layer 2 Networks Upstream



- Each vNIC is pinned exactly to an uplink port based on uplink VLAN membership
- Server-to-server traffic is locally switched
- Server-to-network traffic is transmitted on a pinned uplink port


- Network-to-server unicast traffic is forwarded to the server only if it arrives on a pinned uplink. This feature is called Reverse-Path Forwarding (RPF) check
- Server traffic received on any uplink port, except a pinned uplink port, is dropped. The server MAC address must be learned before traffic can be forwarded to it
- Broadcast and multicast traffic is pinned on per-VLAN basis to an uplink port; it is dropped when received on the other uplink ports

The number of Layer 2 networks for which a blade can send and receive traffic depends on the number of vNICs supported by the mezzanine card because a vNIC can be pinned to only one uplink port (or PortChannel). On a system level, a maximum of 31 disjoint networks are supported.

The Cisco UCS Virtual Interface Card (VIC) or Cisco UCS VIC 1280 is required if a virtualized environment requires individual virtual machines to talk to different upstream Layer 2 domains. Multiple vNICs would need to be created on the service profile in Cisco UCS, and the vNICs need to be assigned to different VMware virtual switches (vSwitches) or have a different uplink port profile assigned to them (Cisco Nexus® 1000V Switch).

For example in Figure 3, if the vNIC definition for vmnic0 has VLANs 10 to 20 trunked to it, pinning will fail because no uplink port matches the VLAN membership. Instead, vmnic0 should have VLANs 10 to 15 (or a subset such as VLANs 10 to 12) trunked to it.

Figure 3. vNIC pinning using End-Host Mode with Disjoint Layer 2 Networks Upstream

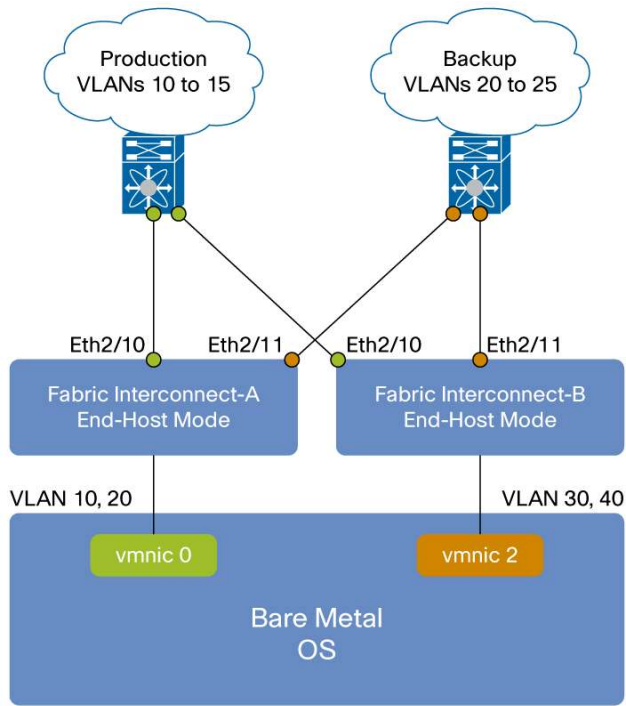


Configuration

The current default behavior of trunking all VLANs to all available uplinks is still present. This default behavior has been retained to maintain backward compatibility. To allow only certain VLANs on certain uplinks, you need to change the configuration.

The configuration steps use the topology shown in Figure 4. These steps assume that links to upstream switches are configured as uplink ports.

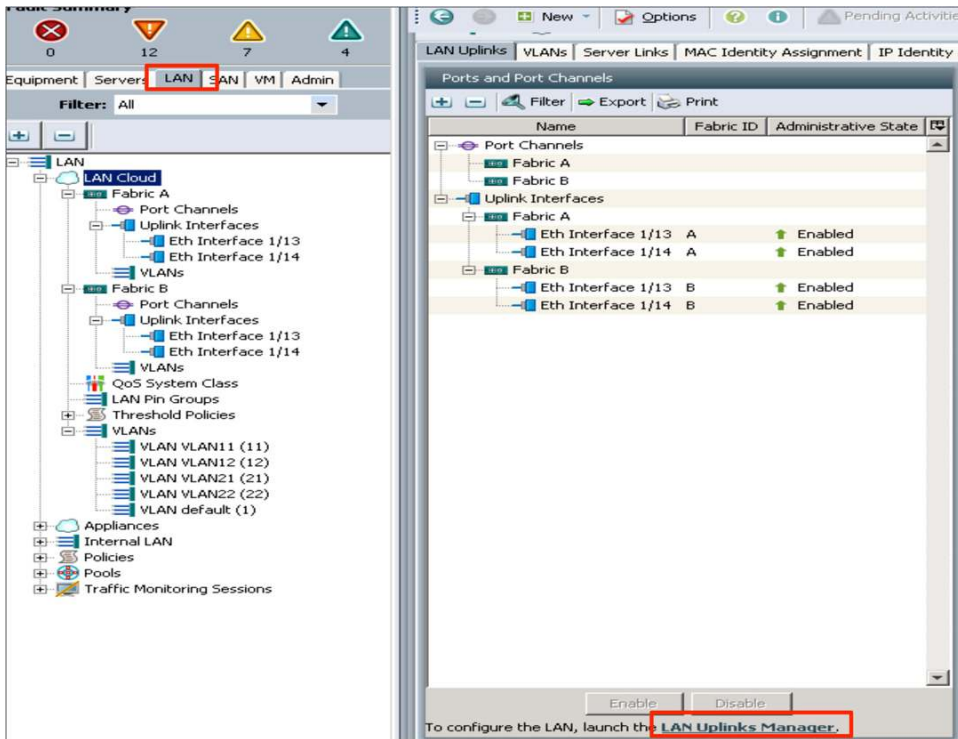
Figure 4. Configuration Topology



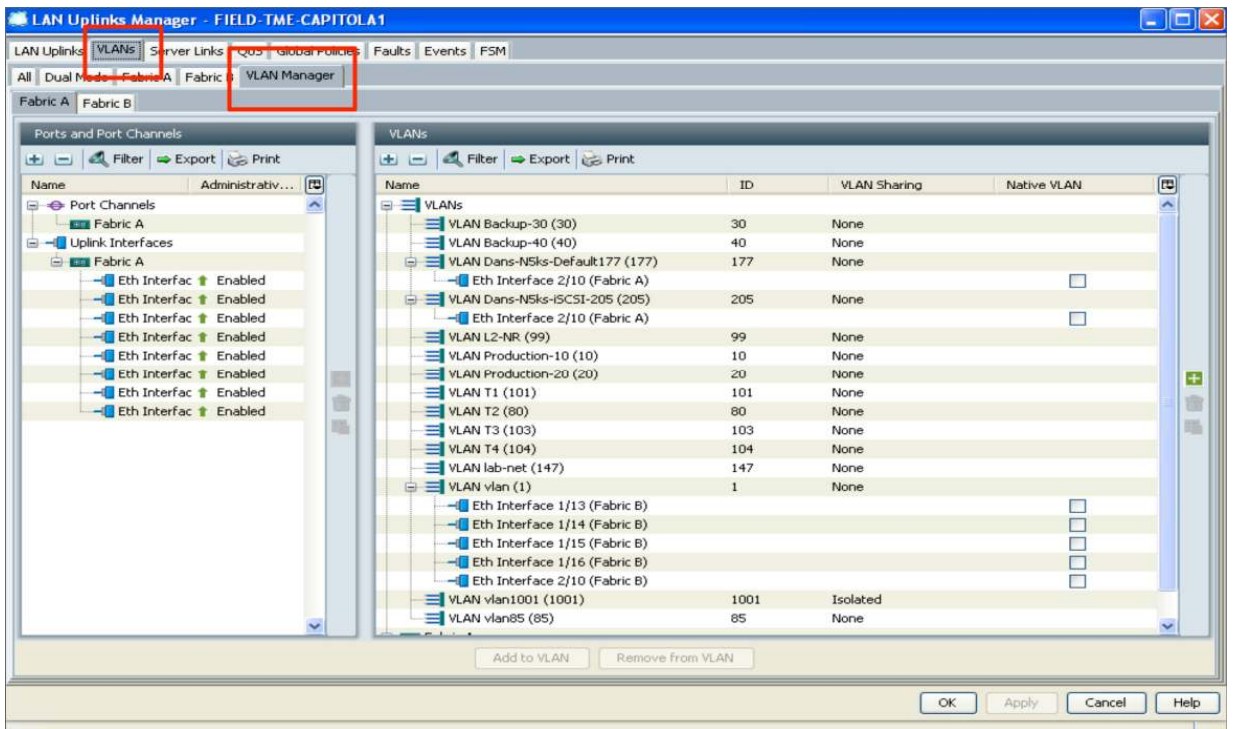
Configuration Steps

1. Define VLANs globally.
2. Define uplinks on Fabrics A and B.
3. Assign uplinks to VLANs on Fabrics A and B.
4. Verify the configuration.

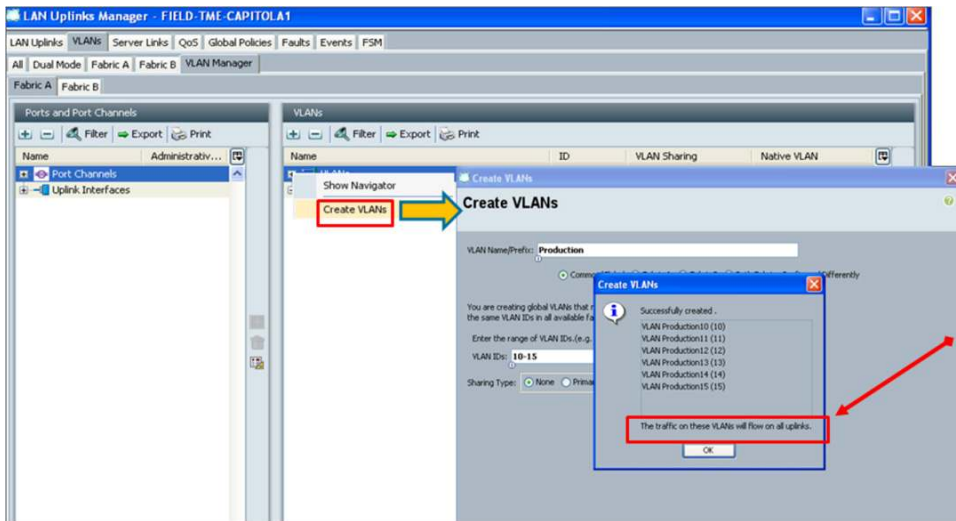
To begin, click the LAN tab and then click LAN Uplinks Manager.



In the LAN Uplinks Manager, click VLANs, then VLAN Manager, and then Fabric A.

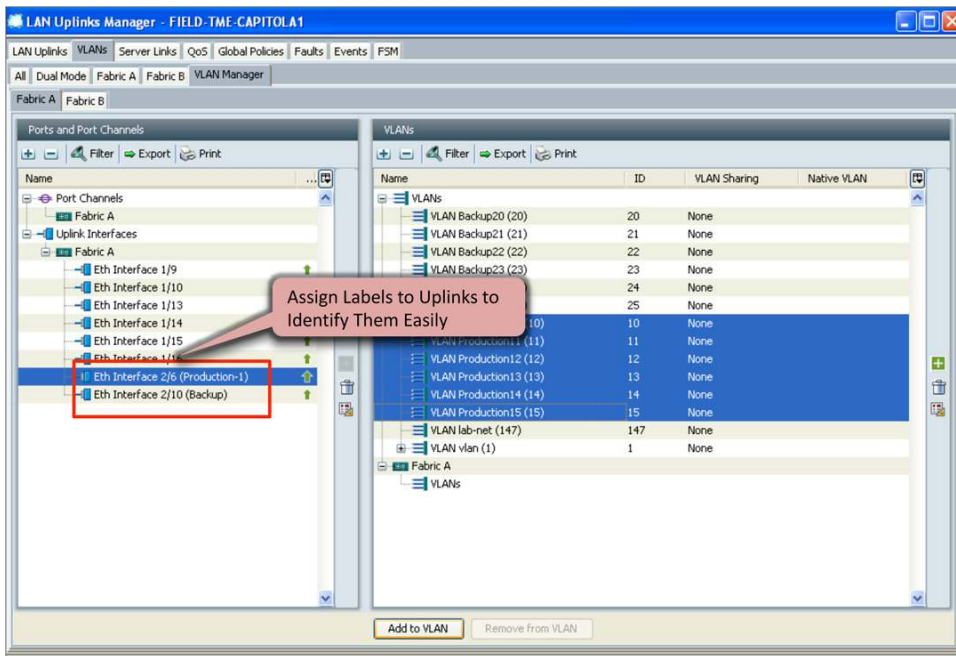


Step 1. Define VLANs globally or reuse the existing VLANs.



When VLANs are newly created and do not have specific uplinks associated, these VLANs will flow on all uplinks.

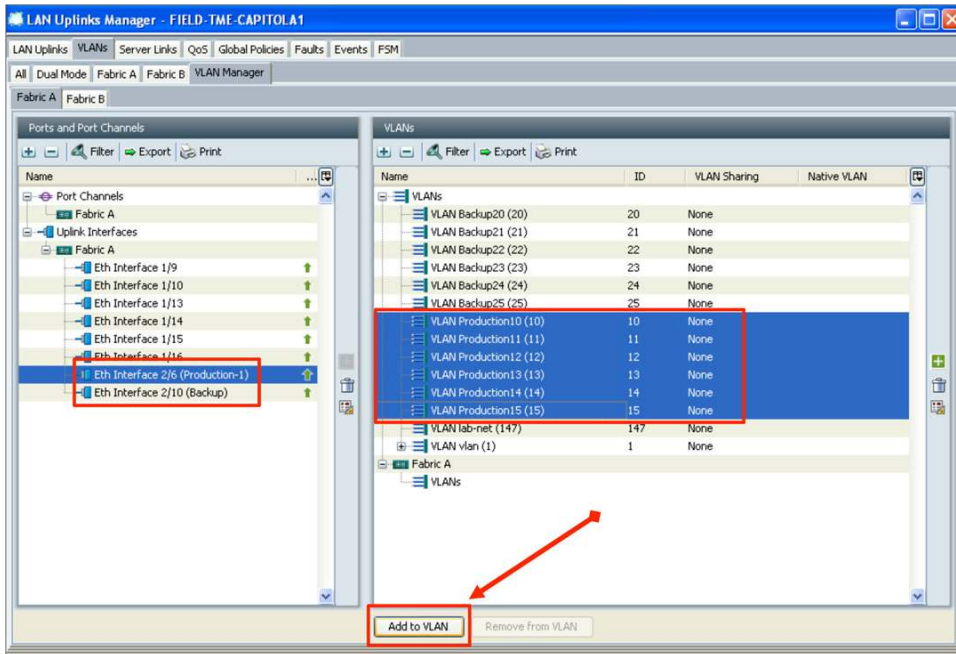
Step 2. Define uplinks on Fabrics A and B.



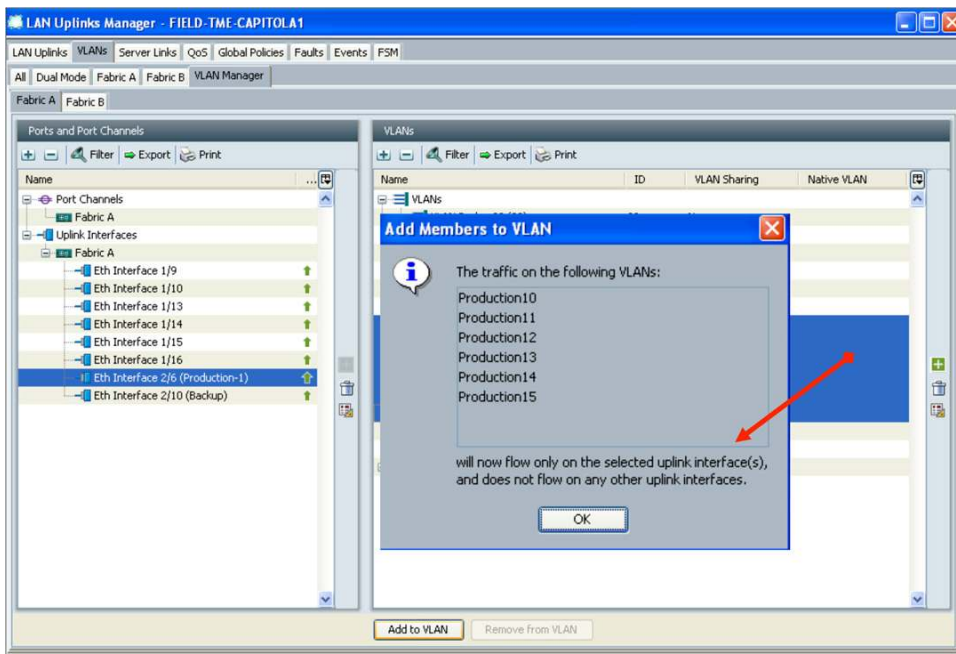
Step 3. Assign uplinks to VLANs.

Select the uplink interface in the left pane and then the VLANs in the right pane.

Click Add to VLAN to add the selected uplinks to the VLANs.



Note: VLANs flow only on selected uplinks, not on the rest of the uplinks.




Follow the same procedure for Fabric B.

Click Save to save the configuration and make it effective.

Removing VLAN Membership

When the last uplink is removed from a VLAN, since no specific uplink is assigned to this VLAN, it inherits the default behavior of the VLAN everywhere and flows on all uplinks. The recommended approach is to remove all unused VLANs from the system.



Note: VLAN 1 still is present on all uplinks and needs to be removed from the upstream switches.

The port designated for multicast and broadcast traffic on a VLAN basis can be queried through the command-line interface (CLI) in the Cisco® NX-OS Software:

```
UCS (nxos)# show platform software enm internal info vlandb id 11
vlan_id 11
-----
Designated receiver: Eth1/13
Membership:
Eth1/13
```


Exceptions and Error Handling

The user must be sure to correctly configure the VLAN on uplinks and vNICs because if the configuration is not correct, vNIC pinning failure may occur. The actions described here are taken in the event of errors.

Single VLAN on a vNIC (No Trunk)

If an invalid pin group is specified (that is, if the VLAN does not exist on the uplink specified in the pin group), the static pin group configuration is applied, the vNIC stays up, and traffic is discarded. A configuration error is flagged in the Status pane as shown in Figure 5.

Figure 5. Configuration Error Flagged in the Status Pane

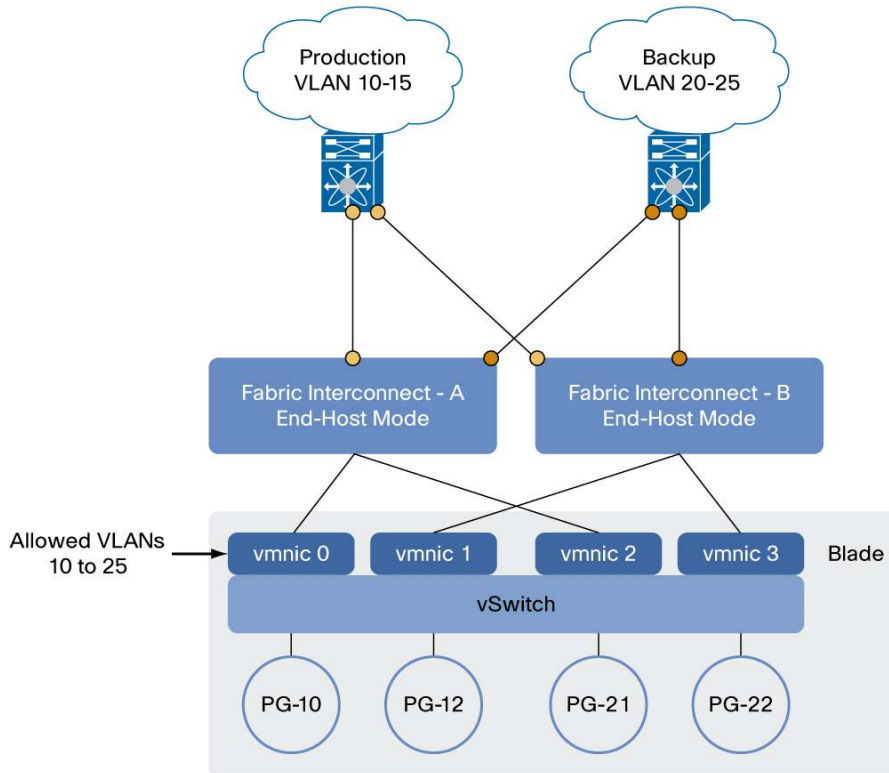


Multiple VLANs on a vNIC (Soft Switch Scenario)

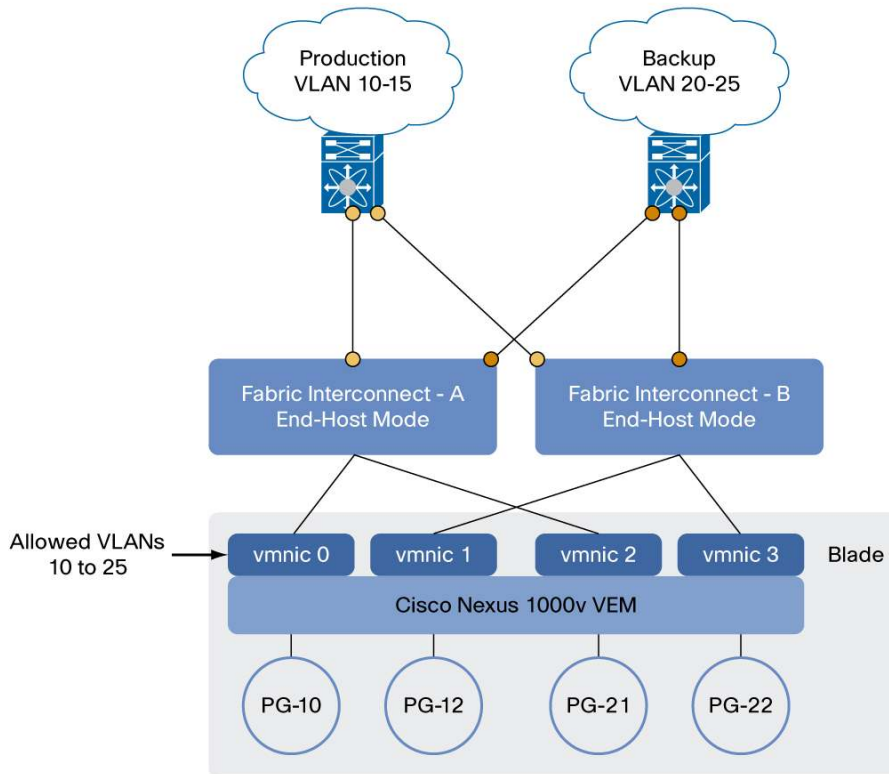
If a vNIC is defined to carry VLANs belonging to two separate disjoint Layer 2 upstream networks, pinning will fail, and a fault will be raised.

For example, in the topology in Figure 6, if a vNIC is defined to carry VLANs 10 through 25 in the service profile, pinning will fail because no uplink matches the pinning criteria.

Figure 6. VLAN mis-configuration on a vNIC



Severity	Code	ID	Affected object	Cause	Last Tran...	Description
▼	F0207	2739535	sys/chassis-1/...	link-down	2011-09-13T...	Adapter host interface 1/1/1/1 link state: down
▼	F0479	2739458	sys/chassis-1/...	vif-down	2011-09-13T...	Virtual interface 924 link state is down
▼	F0479	2739456	sys/chassis-1/...	vif-down	2011-09-13T...	Virtual interface 923 link state is down
▼	F0283	2739455	sys/chassis-1/...	link-down	2011-09-13T...	ether VIF 1 / 1 A-923 down, reason: ENM source pinning failed
▼	F0283	2739454	sys/chassis-1/...	link-down	2011-09-13T...	ether VIF 1 / 1 B-924 down, reason: ENM source pinning failed



Conclusion

Cisco UCS Manager Release 2.0 brings the capability to filter VLANs on fabric interconnect uplinks and enables support for Layer 2 disjoint upstream in end-host mode.



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)