



Cisco Compute Hyperconverged with Nutanix

Disjoint Layer 2 Networking Configurations
and Distributed Virtual Switch Migrations

Document Information

Revision History

Version	Date	Foundation VM version	Foundation Central version	AOS LTS version	AOS STS version	LCM Version	Notes
1.0	Feb 2024	5.5	N/A	6.5.4.5	6.7.1	2.7	Initial Release for FI based deployments with M6 generation servers.



Contents

- [Cisco UCS Network Configuration](#)
- [Nutanix Installation](#)
- [Modify UCS Configurations](#)
- [Guest VM Networking](#)
- [Distributed Virtual Switch Migration](#)

Common use cases for Disjoint Layer 2 VLAN configs

Use case 1

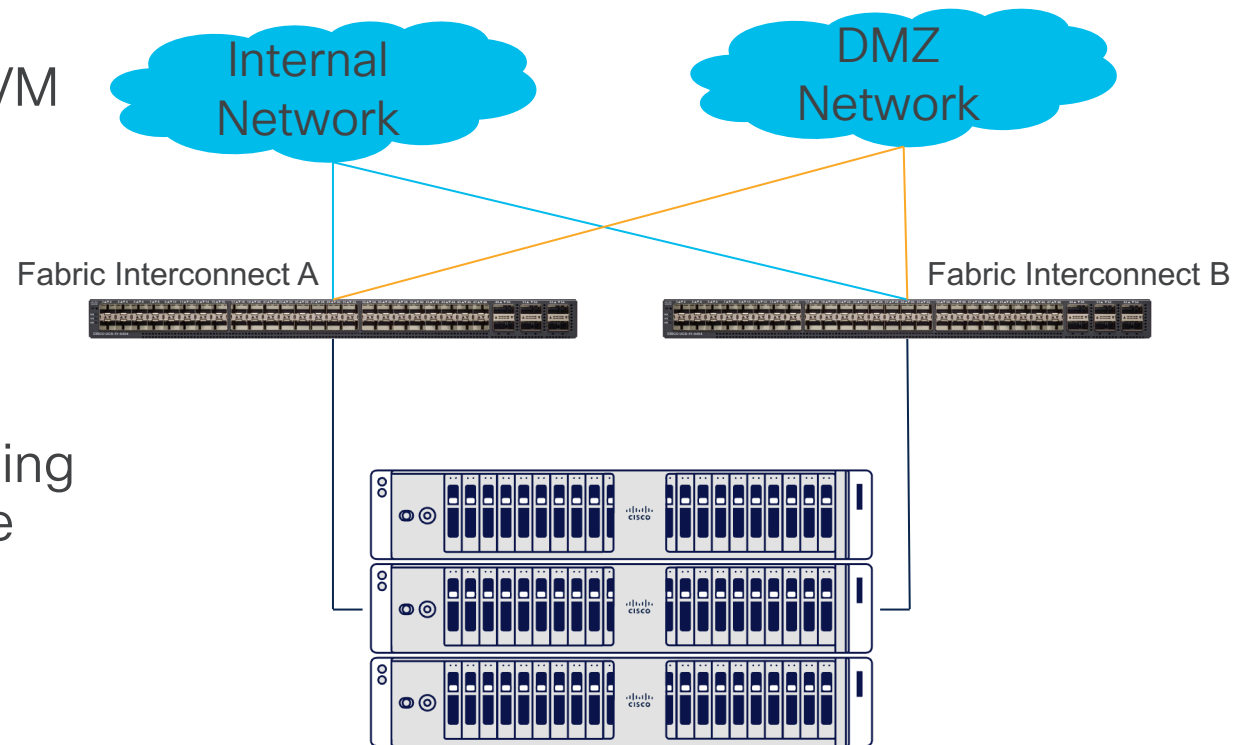
Management (OOB and Storage) and DMZ (VM traffic) are connected to different upstream networks – a common DMZ deployment for UCS for Banks

Use case 2

VMs deployed are part of different groups using common storage infrastructure. Networks are isolated and disjoint

Use case 3

Management (OOB and Storage) and Backup (via a backup proxy VM) are connected to different upstream networks



Summary of supported configurations

- For ESXi deployments – This document covers the steps to add additional vNICs to the hosts, the disjoint L2 configuration in UCSM and the virtual switch configuration in VMware vCenter.
- For AHV deployments – Support for adding additional vNICs is not supported. However, a disjoint L2 configuration is still possible in a dual-VIC hardware configuration, i.e. a Cisco VIC mLOM + Cisco VIC PCIe card, because the default configuration built by Foundation will have 4 vNICs.
- Guest VM Networking – In FI managed mode, we support the addition of 3rd party NIC cards that are connected directly to ToR switches. This is for guest VM traffic only and the physical links of the card should be used by their own dedicated vSwitches. These NICs should be disconnected during install. Note: Cisco VIC cards can't be used for this purpose

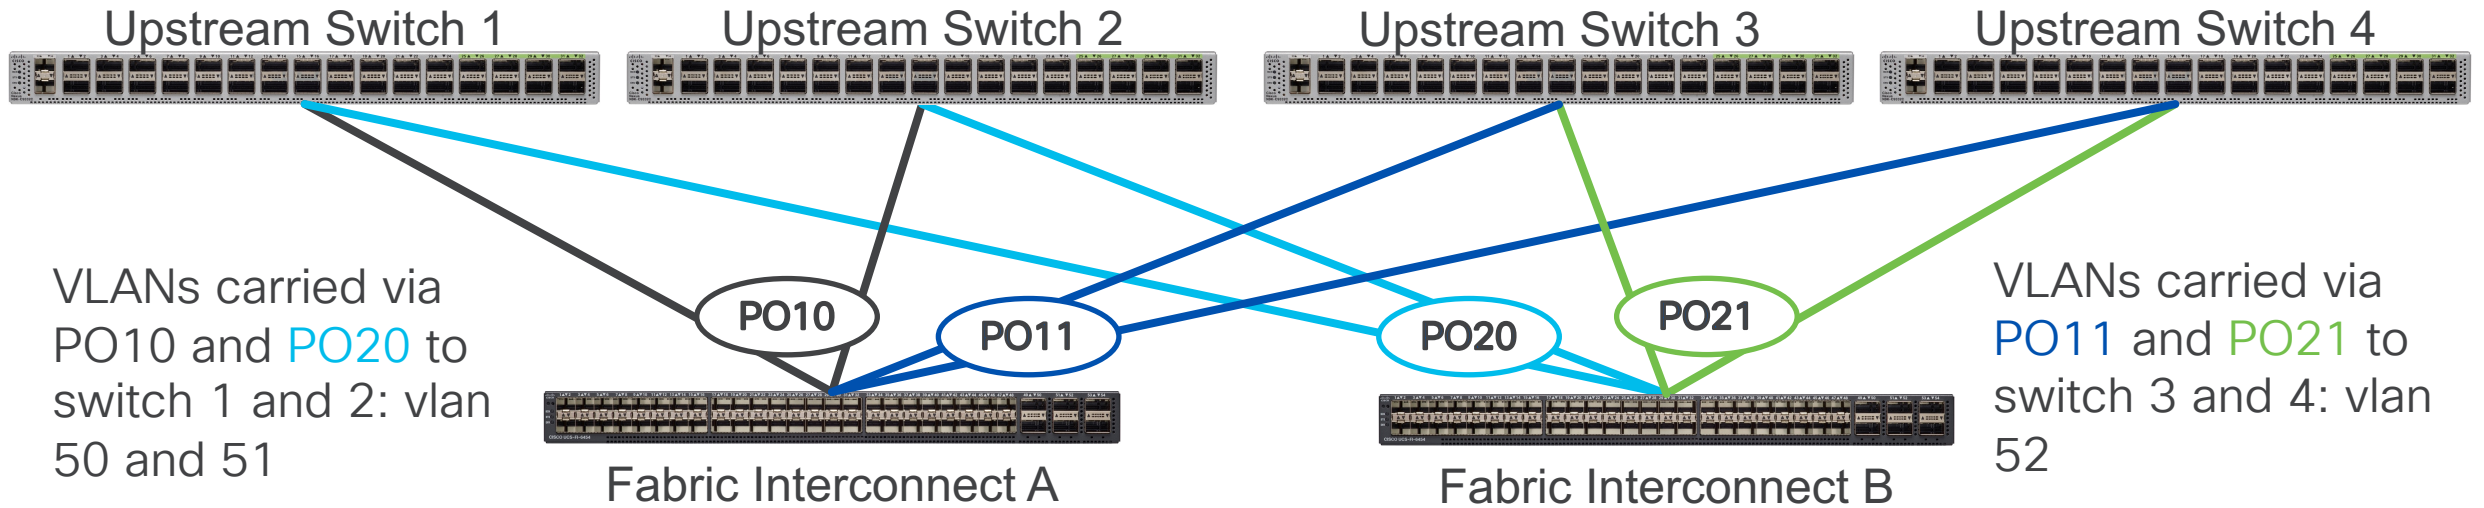
Cisco UCS Network Configuration



Assumptions and Prerequisites

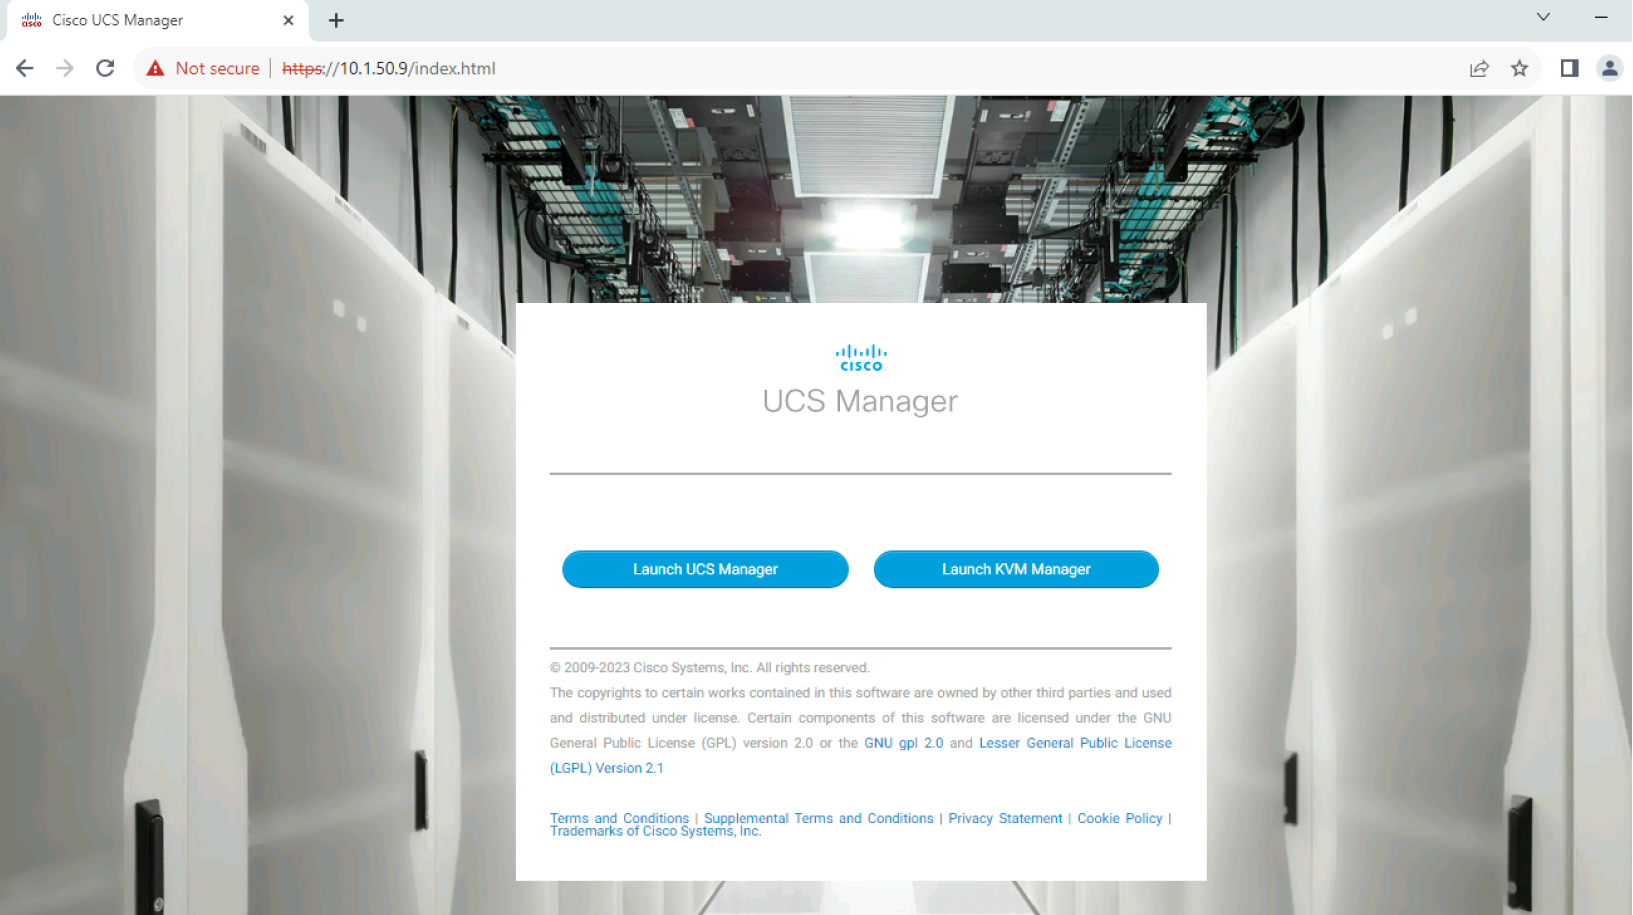
- This document assumes all northbound network connections have been established between the Fabric Interconnects and their northbound switches, and the configuration of the allowed VLANs has already been completed on the northbound switches.
- This document assumes a basic disjoint layer 2 configuration as shown in the diagram following this slide, where two northbound paths exist from each Fabric Interconnect, each carrying a distinct set of VLAN IDs.
- This document does not cover the initial installation of Nutanix but does cover post-setup configuration tasks necessary to enable disjoint L2 functionality for both ESXi and AHV hypervisors.
- Disjoint L2 with AHV is only possible with dual Cisco VIC cards per server which has 4 vNICs defined by Foundation during the installation.

Disjoint Layer 2 Config



Note: In this example all northbound connections are vPC port channels, these uplinks could be single connections, for example a single link from FI-A to switch 3 and a single link from FI-B to switch 4, each configured to carry only VLAN 52. This matches the examples shown in the document as built in our lab.

Log in to Cisco UCS Manager



Connect to the roaming cluster IP address, not an individual FI's IP address

Enable Uplink Ports

The screenshot shows the Cisco UCS Manager interface. The left sidebar contains navigation menus for Equipment, Fabric Interconnects, and Fixed Modules. The main area displays a table of Ethernet Ports. A context menu is open over a row, with 'Configure as Uplink Port' highlighted.

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	37	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	38	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	39	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	40	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	41	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	42	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	43	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	44	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	45	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	46	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	47	00:08:31:2A:F5...	Unconfigured	Physical	Admin Down	Disabled	
1	0	48	00:08:31:2A:F5...	Unconfigured	Physical	Admin Down	Disabled	
1	0	49	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	50	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	51	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	52	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	53	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	54	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	55	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	56	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	57	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	58	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	59	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	
1	0	60	00:08:31:2A:F5...	Unconfigured	Physical	Sfp Not Pre...	Disabled	

If not already done, configure the uplink ports for both FI-A and FI-B for all interfaces that will be used as uplinks

Create Uplink Port Channels (optional)

The first screenshot shows the UCS Manager interface with the navigation path: LAN > LAN Cloud > Fabric A > Port Channels > Create Port Channel. The second screenshot shows the 'Create Port Channel' dialog with 'Add Ports' selected, displaying two ports in a table:

Ports			
Slot ID	Aggr. Po...	Port	MAC
No data available			

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
1	0	47	00:08:3...
1	0	48	00:08:3...

The third screenshot shows the UCS Manager interface with the navigation path: LAN > LAN Cloud > Fabric B > Port Channels > Create Port Channel.

If using port-channels from the FIs to the uplink switches, create the port channels for FI-A and FI-B, adding the uplink ports that were enabled earlier.

Create UCS VLANs

The screenshot shows the UCS Manager interface. On the left, the navigation menu is visible with 'VLANs' highlighted in blue and a green checkmark. The main content area shows the 'LAN / LAN Cloud / VLANs' path. A table lists the existing VLANs:

Name	ID	Type	Transport	Native
VLAN default (1)	1	Lan	Ether	Yes

Below the table, there are 'Add', 'Delete', and 'Info' buttons. The 'Add' button is highlighted with a red box. At the top right of the interface, there are four status icons with counts: a red 'X' with '0', an orange triangle with '2', a yellow triangle with '0', and a green circle with '2'.

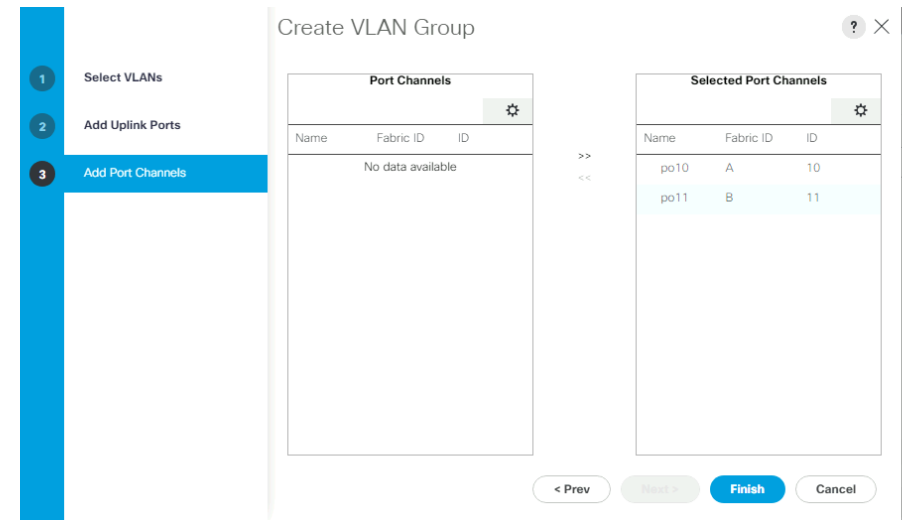
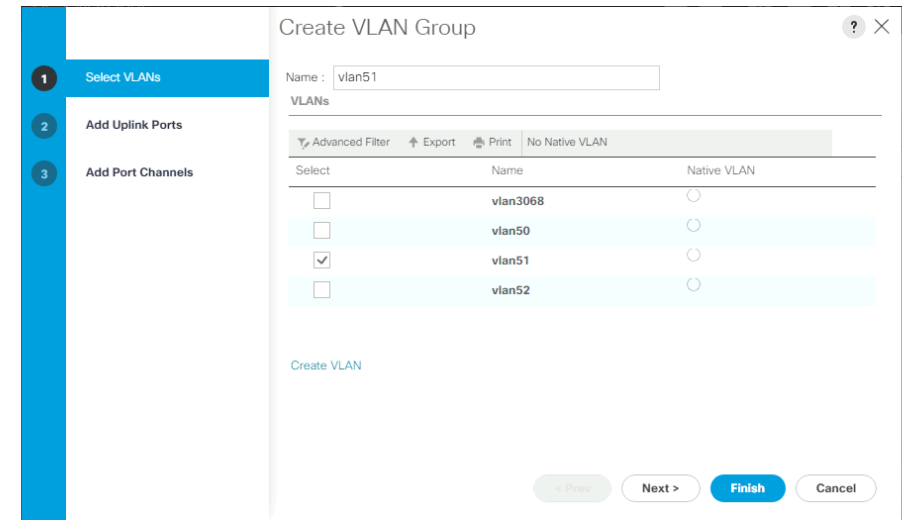
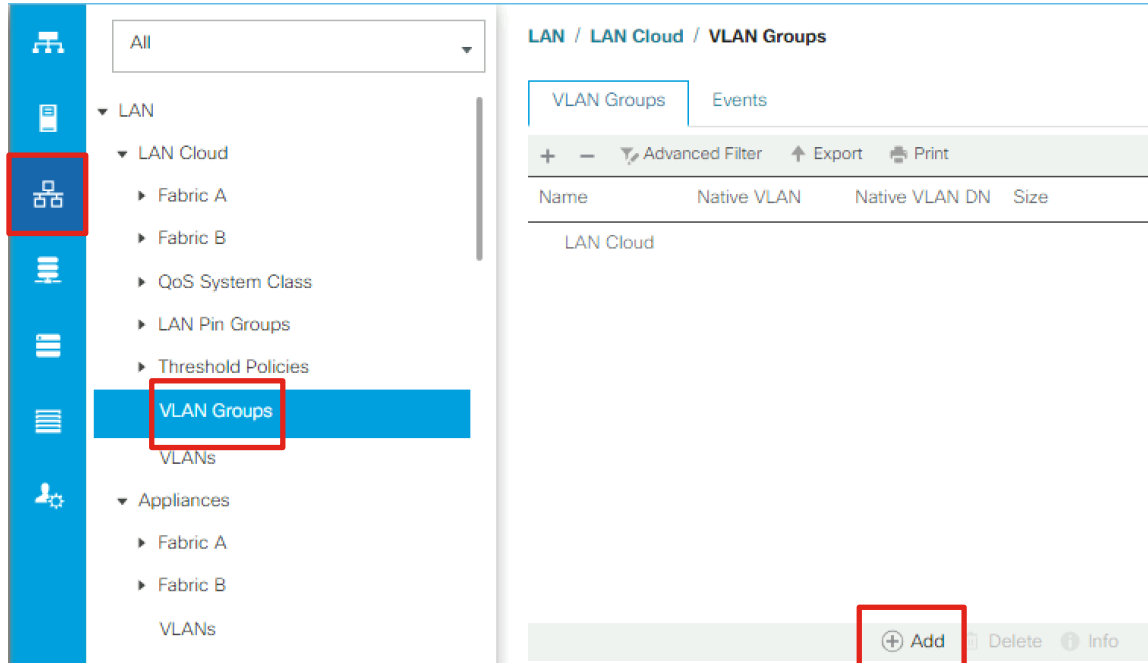
The 'Create VLANs' dialog box is shown. It contains the following fields and options:

- VLAN Name/Prefix :
- Multicast Policy Name : [Create Multicast Policy](#)
- Sharing Type : Common/Global Fabric A Fabric B Both Fabrics Configured Differently
- VLAN IDs :

Below the dialog, there are three buttons: 'Check Overlap', 'OK', and 'Cancel'.

Create all of the VLANs needed in UCS Manager, for example the VLAN for Nutanix cluster CVMs and ESXi/AHV host management, and multiple VLANs for the guest VMs which would be divided across the uplinks in the disjoint L2 configuration.

Create UCS VLAN Groups



Create a VLAN Group for each set of uplinks, adding to it the VLAN IDs carried by those uplinks, and adding the port channels or the individual uplink ports to the group as well. For our example we will create two VLAN groups, one with VLAN 51, and one with VLAN 52.

UCS VLAN Groups Continued

The screenshot shows the UCS management interface. On the left is a navigation menu with a tree structure. The 'VLAN Groups' item is highlighted with a red box. The main content area is titled 'LAN / LAN Cloud / VLAN Groups'. It has two tabs: 'VLAN Groups' (active) and 'Events'. Below the tabs is a toolbar with '+', '-', 'Advanced Filter', 'Export', and 'Print' icons. A table with the following columns is visible: 'Name', 'Native VLAN', 'Native VLAN DN', and 'Size'. The table contains one entry: 'LAN Cloud'. At the bottom right of the main area, there is a toolbar with '+ Add', 'Delete', and 'Info' icons. The '+ Add' icon is highlighted with a red box.

The first screenshot shows the 'Create VLAN Group' wizard. It has a progress bar on the left with three steps: '1 Select VLANs', '2 Add Uplink Ports', and '3 Add Port Channels'. The '1 Select VLANs' step is active. The 'Name' field contains 'vlan52'. Below is a table of available VLANs:

Select	Name	Native VLAN
<input type="checkbox"/>	vlan3068	<input type="radio"/>
<input type="checkbox"/>	vlan50	<input type="radio"/>
<input type="checkbox"/>	vlan51	<input type="radio"/>
<input checked="" type="checkbox"/>	vlan52	<input type="radio"/>

At the bottom right, there are buttons for '< Prev', 'Next >', 'Finish', and 'Cancel'. The 'Finish' button is highlighted.

The second screenshot shows the 'Create VLAN Group' wizard at the '2 Add Uplink Ports' step. It has two tables side-by-side. The left table is titled 'Uplink Ports' and contains the text 'No data available'. The right table is titled 'Selected Uplink Ports' and contains the following data:

Fabric ID	Slot ID	Aggre...	Port ID
A	1	0	46
B	1	0	46

At the bottom right, there are buttons for '< Prev', 'Next >', 'Finish', and 'Cancel'. The 'Finish' button is highlighted.

Create a VLAN Group for each set of uplinks, adding to it the VLAN IDs carried by those uplinks, and adding the port channels or the individual uplink ports to the group as well. For our example we will create two VLAN groups, one with VLAN 51, and one with VLAN 52.

Nutanix Installation



Install a new Nutanix Cluster

Perform a standard installation of Cisco Compute Hyperconverged with Nutanix, following the latest version of this document at Cisco Communities: <https://community.cisco.com/t5/unified-computing-system-knowledge-base/cisco-compute-hyperconverged-with-nutanix-field-guide/ta-p/4982563>

Modify UCS Configurations

- [Modify Configuration for ESXi](#)
- [Modify Configuration for AHV](#)

Modify UCS Config for ESXi

Add vNICs to the Nutanix Service Profiles

The screenshot displays the Nutanix Prism Central interface. On the left, the navigation pane shows the hierarchy: Servers > Service Profiles > root > fdtnWMP27210026. The 'Network' tab is selected in the top navigation bar. The main content area shows the configuration for the service profile 'fdtnWMP27210026'. The 'vNICs' section contains a table with two entries:

Name	MAC Address	Desired Order	Actual Order	Fabric ID	Desired Place...	Actual Placem...
vNIC 1-fabric-A	00:25:B5:A0:F...	1	1	A	Any	1
vNIC 1-fabric-B	00:25:B5:A0:F...	2	2	B	Any	1

At the bottom of the interface, there are buttons for 'Delete', 'Add', and 'Modify'. The 'Add' button is highlighted with a red box.

Add two new vNICs to each of the Nutanix Service Profiles for all the servers in the cluster. In our example, these two vNICs will carry only VLAN 52, as the original two vNICs were configured to carry VLAN 51 by Foundation during the installation. The first vNIC will be configured on Fabric A, the second on Fabric B.

Add the First New vNIC

Create vNIC

Name : 2-fabric-A

MAC Address

MAC Address Assignment : Nutanix(52/64)

Create MAC Pool

The MAC address will be automatically assigned from the selected pool.
The MAC address assignment change will be effective only after server reboot.

Use vNIC Template :

Fabric ID : Fabric A Fabric B Enable Failover

Enable QinQ :

Note: The QinQ VLAN selection is considered only when the Enable QinQ check box is checked.

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash.

VLANs	VLAN Groups			
<input type="checkbox"/>	default	<input type="radio"/>	1	<input type="radio"/>
<input type="checkbox"/>	vlan3068	<input type="radio"/>	3068	<input type="radio"/>
<input type="checkbox"/>	vlan50	<input type="radio"/>	50	<input type="radio"/>
<input type="checkbox"/>	vlan51	<input type="radio"/>	51	<input type="radio"/>
<input checked="" type="checkbox"/>	vlan52	<input type="radio"/>	52	<input type="radio"/>

Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

Pin Group : Create LAN Pin Group

OK Cancel

Adapter Policy : <not set> Create Ethernet Adapter Policy

QoS Policy : <not set> Create QoS Policy

Network Control Policy : fdtnWMP27210026 Create Network Control Policy

- Give the new vNIC a descriptive name to differentiate it from the original vNICs.
- Select a MAC address pool with enough available addresses for all the new interfaces.
- Select Fabric A, and **DO NOT** enable failover.
- Select only the vlan ID or IDs which will be carried on the secondary uplinks, in our example only VLAN ID 52. **DO NOT** select any of them as native.
- **NOTE:** you cannot mix VLANs from the different uplinks or VLAN groups on a single vNIC, this will cause a config error.
- You may select a VLAN Group instead of VLANs one-by-one.
- Modify the MTU to 9000.
- Select the network control policy which matches the name of the server's service profile.

Add the Second New vNIC

Create vNIC

Name:

MAC Address

MAC Address Assignment:

[Create MAC Pool](#)
The MAC address will be automatically assigned from the selected pool.
The MAC address assignment change will be effective only after server reboot.

Use vNIC Template:

Fabric ID: Fabric A Fabric B Enable Failover

Enable QinQ:

Note: The QinQ VLAN selection is considered only when the Enable QinQ check box is checked.

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash.

VLANs | VLAN Groups

<input type="checkbox"/>	default	<input type="radio"/>	1	<input type="radio"/>
<input type="checkbox"/>	vlan3068	<input type="radio"/>	3068	<input type="radio"/>
<input type="checkbox"/>	vlan50	<input type="radio"/>	50	<input type="radio"/>
<input type="checkbox"/>	vlan51	<input type="radio"/>	51	<input type="radio"/>
<input checked="" type="checkbox"/>	vlan52	<input type="radio"/>	52	<input type="radio"/>

[Create VLAN](#)
CDN Source: vNIC Name User Defined

MTU:

Pin Group: [Create LAN Pin Group](#)

Adapter Policy: [Create Ethernet Adapter Policy](#)

QoS Policy: [Create QoS Policy](#)

Network Control Policy: [Create Network Control Policy](#)

- Give the new vNIC a descriptive name to differentiate it from the original vNICs.
- Select a MAC address pool with enough available addresses for all the new interfaces.
- Select Fabric B, and **DO NOT** enable failover.
- Select only the vlan ID or IDs which will be carried on the secondary uplinks, in our example only VLAN ID 52. **DO NOT** select any of them as native.
- **NOTE:** you cannot mix VLANs from the different uplinks or VLAN groups on a single vNIC, this will cause a config error.
- You may select a VLAN Group instead of VLANs one-by-one.
- Modify the MTU to 9000.
- Select the network control policy which matches the name of the server's service profile.

Verify the Desired Order and Save Changes

vNICs

Name	MAC Address	Desired Ord...	Actual Order	Fabric ID	Desired Place...	Actual Placem...	Admin Host Port	Actual Host Port
vNIC 1-fabric-A	00:25:B5:A0:F...	1	1	A	Any	1	ANY	NONE
vNIC 1-fabric-B	00:25:B5:A0:F...	2	2	B	Any	1	ANY	NONE
vNIC 2-fabric-A	Derived	3	unspecified	A	Any	Any	ANY	NONE
vNIC 2-fabric-B	Derived	4	unspecified	B	Any	Any	ANY	NONE

Advanced Filter Export Print

Delete Add Modify

Save Changes

Reset Values

Verify that the new vNICs are listed with their desired order as 3 and 4, then save the changes in the service profile. The changes will result in pending changes which will automatically take effect on the next reboot of the node.

Save Changes

Your changes:

- Create: **vNIC 2-fabric-A** (*org-root/ls-fdtnWMP27210026/ether-2-fabric-A*)
- Create: **Network vlan52** (*org-root/ls-fdtnWMP27210026/ether-2-fabric-A/if-vlan52*)
- Create: **vNIC 2-fabric-B** (*org-root/ls-fdtnWMP27210026/ether-2-fabric-B*)
- Create: **Network vlan52** (*org-root/ls-fdtnWMP27210026/ether-2-fabric-B/if-vlan52*)

Will require User Acknowledgement before the Reboot of:
Service Profile fdtnWMP27210026 (*org-root/ls-fdtnWMP27210026*) [Server: **sys/rack-unit-4**]

Or pending changes will be applied during the next reboot.

Are you sure you want to apply the changes?

Yes No

Pending Activities

There are activities in the system requiring reboot.
Click the **Pending Activities** button on the toolbar to acknowledge.

Do not show again

Repeat for all the Nutanix Service Profiles

Servers / Service Profiles / root / Service Profile fdtnWMP27...

General Storage **Network** iSCSI vNICs vMedia Policy Boot Order Virtual Machines FC Zones Policies

Actions

Modify vNIC/vHBA Placement

Dynamic vNIC Connection Policy

Nothing Selected

vNIC/vHBA Placement Policy

Nothing Selected

LAN Connectivity Policy

LAN Connectivity Policy : <not set>

LAN Connectivity Policy Instance :

Create LAN Connectivity Policy

vNICs

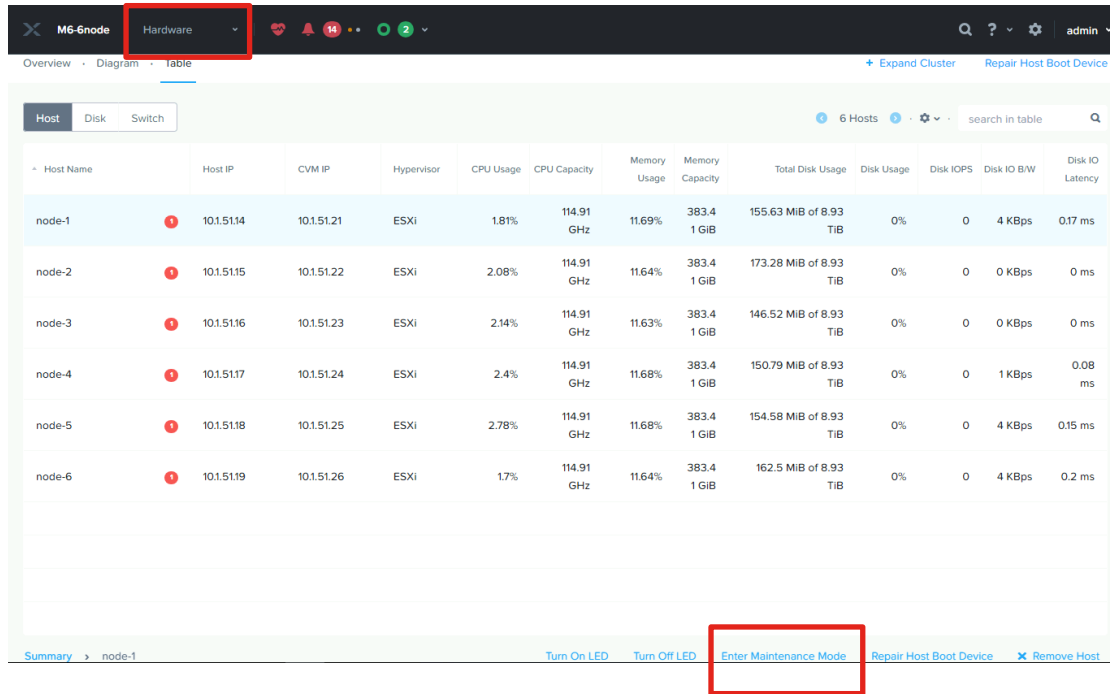
Advanced Filter Export Print

Name	MAC Address	Desired Order	Actual Order	Fabric ID	Desired Place...	Actual Placem...
vNIC 1-fabric-A	00:25:B5:A0:F...	1	1	A	Any	1
vNIC 1-fabric-B	00:25:B5:A0:F...	2	2	B	Any	1
vNIC 2-fabric-A	00:25:B5:A0:F...	3	3	A	Any	1
vNIC 2-fabric-B	00:25:B5:A0:F...	4	4	B	Any	1

Delete Add Modify

Repeat the process until all the service profiles in the cluster have 4 vNICs set with the appropriate VLAN IDs, then proceed with rolling reboots of the nodes in the cluster to apply the new configuration.

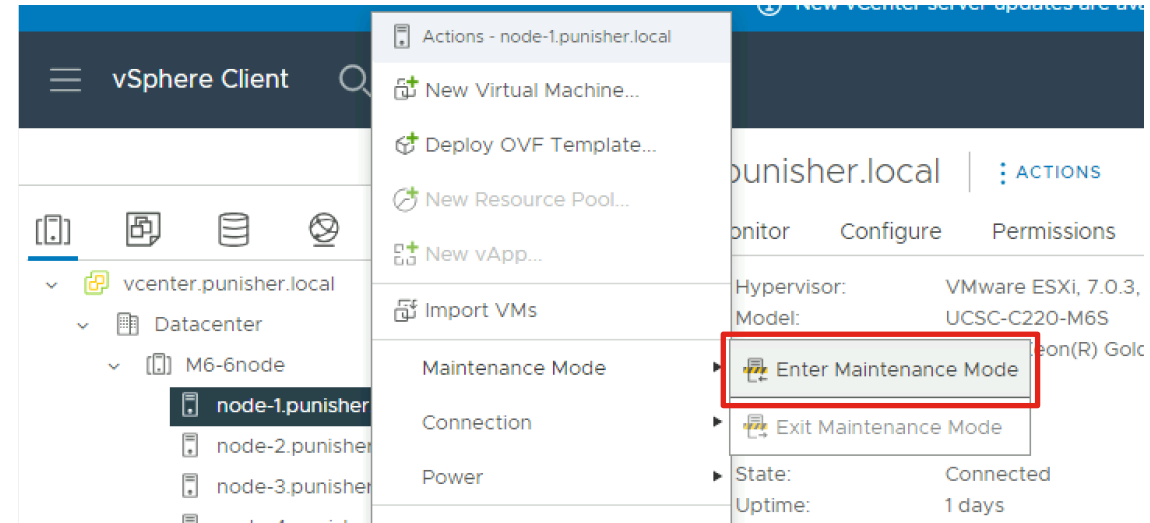
Place First Node in Maintenance Mode and Reboot



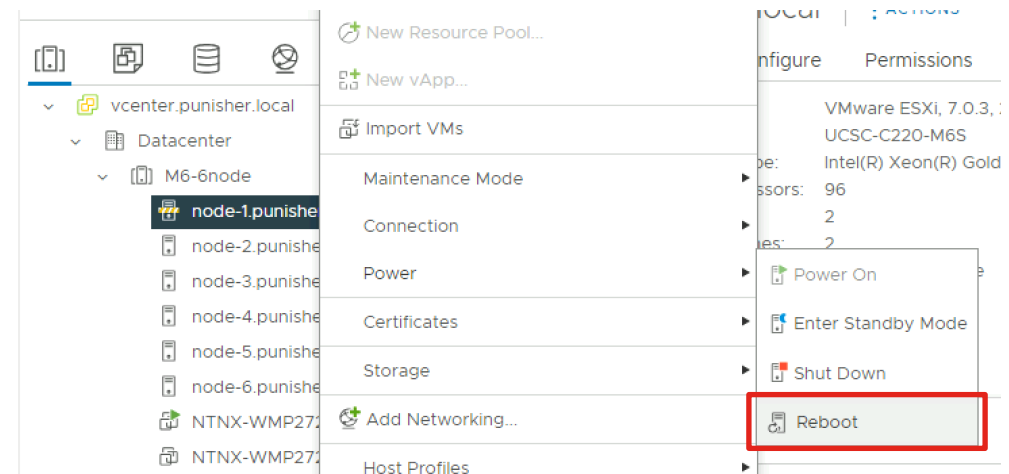
The screenshot shows the vCenter Hardware tab for a cluster named 'M6-6node'. A table lists six hosts (node-1 to node-6) with their respective IP addresses, CVM IP addresses, hypervisors, and resource usage. The 'Enter Maintenance Mode' button is highlighted in red at the bottom of the table.

Host Name	Host IP	CVM IP	Hypervisor	CPU Usage	CPU Capacity	Memory Usage	Memory Capacity	Total Disk Usage	Disk Usage	Disk IOPS	Disk IO B/W	Disk IO Latency
node-1	10.1.51.14	10.1.51.21	ESXi	1.81%	114.91 GHz	11.69%	383.4 1 GiB	155.63 MiB of 8.93 TiB	0%	0	4 KBps	0.17 ms
node-2	10.1.51.15	10.1.51.22	ESXi	2.08%	114.91 GHz	11.64%	383.4 1 GiB	173.28 MiB of 8.93 TiB	0%	0	0 KBps	0 ms
node-3	10.1.51.16	10.1.51.23	ESXi	2.14%	114.91 GHz	11.63%	383.4 1 GiB	146.52 MiB of 8.93 TiB	0%	0	0 KBps	0 ms
node-4	10.1.51.17	10.1.51.24	ESXi	2.4%	114.91 GHz	11.68%	383.4 1 GiB	150.79 MiB of 8.93 TiB	0%	0	1 KBps	0.08 ms
node-5	10.1.51.18	10.1.51.25	ESXi	2.78%	114.91 GHz	11.68%	383.4 1 GiB	154.58 MiB of 8.93 TiB	0%	0	4 KBps	0.15 ms
node-6	10.1.51.19	10.1.51.26	ESXi	1.7%	114.91 GHz	11.64%	383.4 1 GiB	162.5 MiB of 8.93 TiB	0%	0	4 KBps	0.2 ms

Place the first node into maintenance mode via vCenter, Prism Element or Prism Central. If using vCenter you will likely have to manually shut down the controller VM (CVM) on the node for the node to enter maintenance mode. Then reboot the node to apply the new vNIC changes.



The screenshot shows the vSphere Client interface. The context menu for 'node-1.punisher.local' is open, and the 'Enter Maintenance Mode' option is highlighted in red. Other options include 'New Virtual Machine...', 'Deploy OVF Template...', 'New Resource Pool...', 'New vApp...', 'Import VMs', 'Connection', and 'Power'.



The screenshot shows the vSphere Client interface. The context menu for 'node-1.punisher.local' is open, and the 'Reboot' option is highlighted in red. Other options include 'New Resource Pool...', 'New vApp...', 'Import VMs', 'Maintenance Mode', 'Connection', 'Power', 'Certificates', 'Storage', 'Add Networking...', and 'Host Profiles'.

Verify Pending Change Applied and Cluster State

Pending Activities

User Acknowledged Activities | Scheduled Activities

Service Profiles | Fabric Interconnects | Servers | Chassis Profiles

Advanced Filter | Export | Print | Show Current User's Activities | Acknowledge All

Name	Overall Status	Server	Acknowledgment St...	Config. Trigger State	Reboot Now
Service Profile fd...	Pending Reboot	sys/rack-unit-2	Waiting For User	Waiting For Next Boot	<input type="checkbox"/>
Service Profile fd...	Pending Reboot	sys/rack-unit-3	Waiting For User	Waiting For Next Boot	<input type="checkbox"/>
Service Profile fd...	Pending Reboot	sys/rack-unit-4	Waiting For User	Waiting For Next Boot	<input type="checkbox"/>
Service Profile fd...	Pending Reboot	sys/rack-unit-5	Waiting For User	Waiting For Next Boot	<input type="checkbox"/>
Service Profile fd...	Pending Reboot	sys/rack-unit-6	Waiting For User	Waiting For Next Boot	<input type="checkbox"/>

Data Resiliency Status

Warning

Data resiliency is reduced
System is self-healing

Rebuild Progress 66%

Failure Domain **Node**

Fault Tolerance **1**

Verify the pending activity in UCSM for the first server is automatically applied during the reboot. Observe that the Nutanix cluster Data Resiliency Status will be in either a **Warning** or **Error** status depending on the configured redundancy factor setting. After the node reboots, exit from maintenance mode and ensure the controller VM boots. Do not place any additional nodes into maintenance mode or reboot them until the first server has booted and the Resiliency Status returns to **OK**. Repeat the reboots, rolling through all servers one-by-one.

Data Resiliency Status

OK

Data resilient as per configuration

Failure Domain **Node**

Fault Tolerance **2**

Modify UCS Config for AHV

Requirements and Caveats

- Extra vNICs cannot be added to the UCSM configuration when running AHV
- Disjoint L2 configuration is therefore only possible with dual VIC configurations which are built by Foundation with 4 vNICs by default
- Foundation will build the system with all 4 vNICs as uplinks on a single OVS virtual switch carrying one native VLAN
- The process is to remove one pair of uplinks from the default OVS virtual switch, and create a new virtual switch using those uplinks
- Adding or removing uplinks from an existing OVS virtual switch, and creating the new virtual switch will each result in rolling reboots of the hypervisors

Determine in UCSM which uplinks to remove from each node

Servers / Service Profiles / root / Service Profile fdtmWMP27...

General Storage **Network** iSCSI vNICs vMedia Policy Boot Order Virtual Machines FC Zones Policies Server Details CIMC Sessio >

Name	MAC Address	Desired Order	Actual Order	Fabric ID	Desired Plac...	Actual Place...	Admin Host P...	Actual Host P...
vNIC 1-fabric-A	00:25:B5:A0:F9:02	unspecified	2	A	1	1	ANY	NONE
vNIC 1-fabric-B	00:25:B5:A0:F9:00	2	1	B	1	1	ANY	NONE
vNIC 2-fabric-A	00:25:B5:A0:F9:03	1	1	A	2	2	ANY	NONE
vNIC 2-fabric-B	00:25:B5:A0:F9:01	unspecified	2	B	2	2	ANY	NONE

LAN Connectivity Policy

LAN Connectivity Policy : <not set>

LAN Connectivity Policy Instance :

Create LAN Connectivity Policy

vNICs

Advanced Filter Export Print

Delete Add Modify

Removing these two vNICs would be an example of a Split Card configuration

Removing these two vNICs would be an example of a Spread vNIC configuration

Two strategies:

1. Split cards: Use both A and B side vNICs from one card as the uplinks for the new virtual switch.
2. Spread vNICs: Use the A side from one card and the B side from the other card as the uplinks

In UCSM, the placement designates which Cisco VIC the vNIC is placed on. Note the MAC addresses to verify the correct interfaces.

Add VLANs in UCS Manager

The screenshot shows the UCS Manager interface. On the left, a navigation menu has 'VLANs' highlighted with a red box and a green checkmark. In the main area, the 'LAN / LAN Cloud / VLANs' page is displayed. A table lists existing VLANs: 'VLAN default (1)' with ID 1 and 'VLAN vlan50 (...)' with ID 50. A red box highlights the '+ Add' button at the bottom right of the table.

Name	ID	Type	Transport	Native
VLAN default (1)	1	Lan	Ether	Yes
VLAN vlan50 (...)	50	Lan	Ether	No

The 'Create VLANs' dialog box is shown. It contains the following fields and options:

- VLAN Name/Prefix :
- Multicast Policy Name : [Create Multicast Policy](#)
- Sharing Type : Common/Global Fabric A Fabric B Both Fabrics Configured Differently
- VLAN IDs :

Buttons at the bottom: Check Overlap, **OK**, Cancel.

Create additional VLANs in UCS Manager for guest VMs, if not already created earlier

Modify the VLAN(s) of the Nutanix Host vNICs

Servers / Service Profiles / root / Service Profile fdtnWMP27...

General Storage Network ISCSI vNICs vMedia Policy Boot Order Virtual Machines FC Zones Policies

Name	MAC Address	Desired Order	Actual Order	Fabric ID	Desired Place...	Actual Placem...
vNIC 1-fabric-A	00:25:B5:A0:F...	unspecified	2	A	1	1
vNIC 1-fabric-B	00:25:B5:A0:F...	2	1	B	1	1
vNIC 2-fabric-A	00:25:B5:A0:F...	1	1	A	2	2
vNIC 2-fabric-B	00:25:B5:A0:F...	unspecified	2	B	2	2

LAN Connectivity Policy

LAN Connectivity Policy : <not set>

LAN Connectivity Policy Instance :

Create LAN Connectivity Policy

vNICs

Advanced Filter Export Print

Modify vNIC

Name : 2-fabric-A

MAC Address

MAC Address Assignment: Nutanix(52/64)

Create MAC Pool

MAC Address : 00:25:B5:A0:F9:03

The MAC address will be automatically assigned from the selected pool.

The MAC address assignment change will be effective only after server reboot.

Use vNIC Template :

Create vNIC Template

Fabric ID : Fabric A Fabric B Enable Failover

Enable QinQ :


Note: The QinQ VLAN selection is considered only when the Enable QinQ check box is checked.

VLANs VLAN Groups

VLAN	Group	Order	Selected
default		1	<input type="radio"/>
vlan3068		3068	<input type="radio"/>
vlan50		50	<input type="radio"/>
vlan51		51	<input checked="" type="radio"/>
vlan52		52	<input type="radio"/>

OK Cancel

Modify the VLANs to the vNICs of the Nutanix service profiles. For example, change the two vNICs on the second VIC to carry only the disjoint L2 VLAN ID, leaving the other two vNICs carrying the other VLAN(s). These VLANs must be non-native (i.e. tagged), while the original VLAN used during installation is native. Repeat for all the hosts in the cluster.



Guest VM Networking

- [Guest VM Networking for ESXi](#)
- [Guest VM Networking for AHV](#)

Configure Guest VM Networking for ESXi

Verify New Physical Adapters

The screenshot shows the vSphere Client interface for node-1.punisher.local. The left sidebar shows a tree view with 'node-1.punisher.local' selected. The main pane is in the 'Configure' tab, showing the 'Physical adapters' section under 'Networking'. A table lists four physical adapters: vmnic0, vmnic1, vmnic2, and vmnic3. The vmnic2 and vmnic3 rows are highlighted with red boxes. Below the table, it says 'No items selected'.

Device	Actual Speed	Configured Speed	Switch	MAC Address
vmnic0	25 Gbit/s	25 Gbit/s	vSwitch0	00:25:b5:a0:f9:0a
vmnic1	25 Gbit/s	25 Gbit/s	vSwitch0	00:25:b5:a0:f9:07
vmnic2	25 Gbit/s	25 Gbit/s	--	00:25:b5:a0:f9:0e
vmnic3	25 Gbit/s	25 Gbit/s	--	00:25:b5:a0:f9:0f

Add New Virtual Switch

The screenshot shows the vSphere configuration page for a host named 'node-1.punisher.local'. The 'Configure' tab is active, and the 'Networking' section is expanded to show 'Virtual switches'. The 'Standard Switch: vSwitch0' is selected, and the 'ADD NETWORKING...' button is highlighted with a red box. The interface displays a central diagram of the vSwitch and its connections to various network components.

node-1.punisher.local | ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

Storage (expanded)
Storage Adapters
Storage Devices
Host Cache Configuration
Protocol Endpoints
I/O Filters

Networking (expanded)
Virtual switches
VMkernel adapters
Physical adapters
TCP/IP configuration

Virtual switches

ADD NETWORKING... REFRESH

Standard Switch: vSwitch0 | ADD NETWORKING EDIT MANAGE PHYSICAL ADAPTERS ...

- Backplane Network** ...
VLAN ID: --
Virtual Machines (1)
- Management Network** ...
VLAN ID: --
VMkernel Ports (1)
vmk0 : 10.1.51.14 ...
- Physical Adapters** (expanded)
 - vmnic0 25000 Full ...
 - vmnic1 25000 Full ...

Add New Virtual Switch Continued

node-1.punisher.local - Add Networking

1 Select connection type

Select connection type

2 Select target device

Select a connection type to create.

3 Connection settings

4 Ready to complete

VMkernel Network Adapter

The VMkernel TCP/IP stack handles traffic for ESXi services such as vSphere vMotion, iSCSI, NFS, FCoE, Fault Tolerance, vSAN, host management and etc.

Virtual Machine Port Group for a Standard Switch

A port group handles the virtual machine traffic on standard switch.

Physical Network Adapter

A physical network adapter handles the network traffic to other hosts on the network.

CANCEL BACK NEXT

node-1.punisher.local - Add Networking

1 Select connection type

Select target device

2 Select target device

Select a target device for the new connection.

3 Create a Standard Switch

4 Connection settings

5 Ready to complete

Select an existing standard switch

BROWSE ...

New standard switch

MTU (Bytes) 1500

CANCEL BACK NEXT

node-1.punisher.local - Add Networking

1 Select connection type

2 Select target device

3 Create a Standard Switch

4 Connection settings

5 Ready to complete

Create a Standard Switch

Assign free physical network adapters to the new switch.

Assigned adapters

+ | ✖ | ↑ | ↓

Active adapters

- (New) vmnic2

Standby adapters

- (New) vmnic3

Unused adapters

Adapter Name	Location	Driver
Cisco Systems In vmnic2	PCI 0000:35:00.0	enic

Status

Status Connected

Actual speed, Duplex 25 Gbit/s, Full Du

Configured speed, Duplex 25 Gbit/s, Full Du

Networks No networks

Network I/O Control

Status Allowed

SR-IOV

Status Not supported

Cisco Discovery Protocol

ⓘ Cisco Discovery Protocol is not available on t

CANCEL BACK NEXT

node-1.punisher.local - Add Networking

1 Select connection type

2 Select target device

3 Create a Standard Switch

4 Connection settings

5 Ready to complete

Connection settings

Use network labels to identify migration-compatible connections common to two or more hosts.

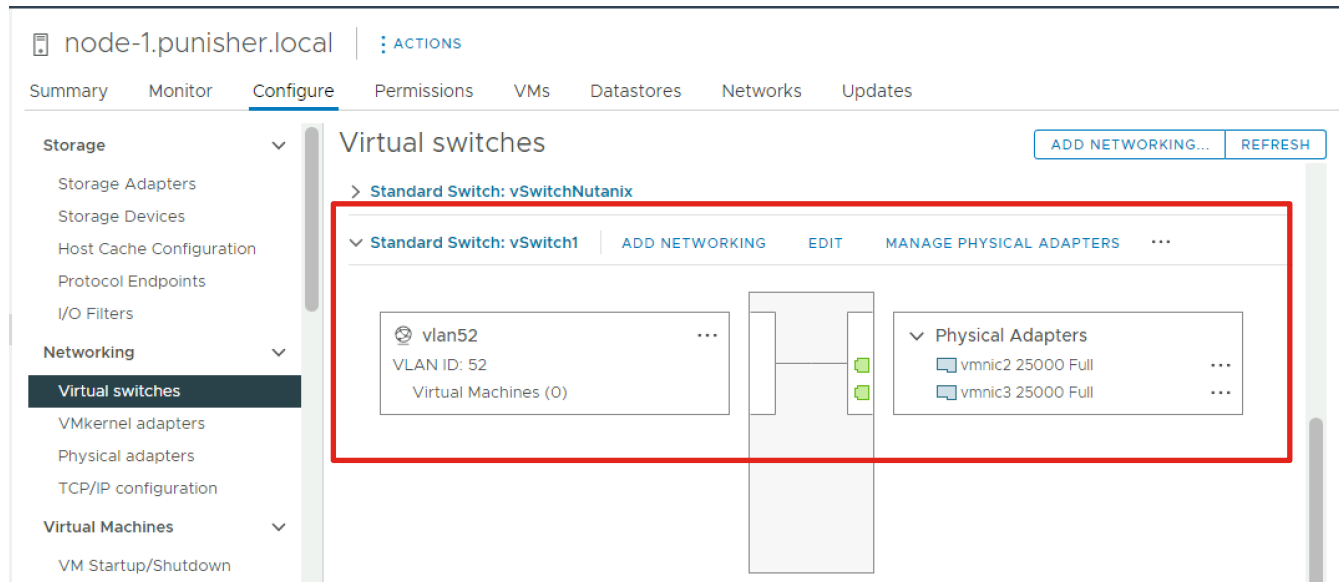
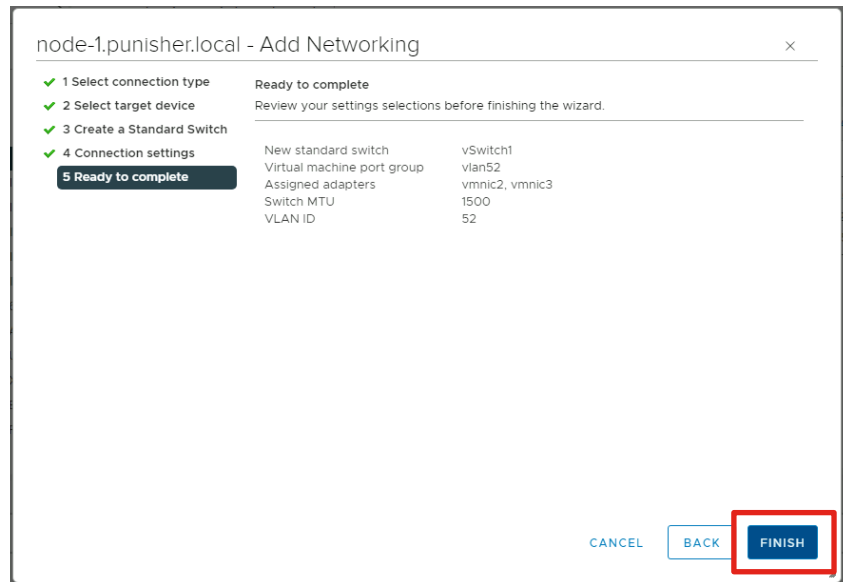
Network label vlan52

VLAN ID 52

CANCEL BACK NEXT



Add New Virtual Switch Completed



Add a new standard virtual switch, using the two new vNICs as the uplinks, with vmnic2 as active and vmnic3 as standby, plus one or more port groups carrying the disjoint L2 VLAN IDs. Repeat the steps for all the servers in the cluster. Afterwards, new or existing VMs can be created or migrated to use the additional disjoint layer 2 VLANs.

Configure Guest VM Networking for AHV

Verify in Prism Element the uplinks to remove

The screenshot shows the Prism Element interface for a host named 'node-4'. The 'Hardware' menu is selected in the top navigation bar. The 'Host NICs' tab is highlighted in the 'Host Details' section. The table below lists the host's network interfaces.

Host Nic	Speed (in KBps)	MAC Address	Rx Pkts	Tx Pkts	Dropped Rx Pkts	Dropped Tx Pkts	Rx Pkt Errors	Tx Pkt Errors
eth0	25000000	00:25:b5:a0:f9:03	2	2	0	0	0	0
eth1	25000000	00:25:b5:a0:f9:01	2	2	0	0	0	0
eth2	25000000	00:25:b5:a0:f9:00	12,887	13,647	0	0	0	0
eth3	25000000	00:25:b5:a0:f9:02	2	2	0	0	0	0

Verify in Prism Element the Host NICs as seen by AHV and match their MAC addresses to determine which uplinks to remove from the default virtual switch. In our example we will use Split cards, so we will remove eth0 and eth1, as they are the A and B pair from the second Cisco VIC as seen in UCSM.

Remove the uplinks from the default OVS virtual switch

Overview · Table + Create VM **Network Config**


VM Include Controller VMs · 3 VMs · search in table

VM Name	Host	IP Addresses	Cores	Memory Capacity	Storage	CPU Usage	Memory Usage	Controller Read IOPS	Controller Write IOPS	Controller IO Bandwidth	Controller Avg IO Latency	Backup...	Flash Mode
NTNX-node-4-CVM	node-4/AHV	10.151...	12	32 GiB	- / 0 GiB	5.24 %	21.94 %	-	-	-	-	No ?	No

Network Configuration

Subnets Internal Interfaces **Virtual Switch**

+ Create VS

Name	Bridge	MTU (bytes)	Bond Type	
vs0	br0	1500	Active-Backup	

Edit Virtual Switch

1 General 2 **Uplink Configuration**

Virtual Switch Name: vs0

Description: Default Virtual Switch

Physical NIC MTU (bytes): 1500

MTU has to be a value in the range 1280 to 9216 inclusive.

Select Configuration Method

Standard (Recommended)

This method will ensure minimal disruption to the workloads by placing the hosts in maintenance mode during the configuration. The process can take a significant amount of time to complete depending on the number of workloads.

Quick

This method can briefly interrupt the workloads on the cluster. The hosts will not be placed in maintenance mode in this method. Use this method only if the cluster is not running production workloads at this time.

Cancel **Next**

Select uplinks to remove

Edit Virtual Switch

Select Hosts

All Hosts x

Select Uplink Ports

Show

Connected and Unconnected Uplink Ports

Only Connected Uplink Ports

On Switches (with LLDP)

All switches

Uplink Port Speeds

All Speeds

Select All Clear All

Hosts	10G NICs
node-4	4 ports
	<input type="checkbox"/> eth0 ?
	<input type="checkbox"/> eth1 ?
	<input checked="" type="checkbox"/> eth2 ?
	<input checked="" type="checkbox"/> eth3 ?

No simultaneous configuration changes to the virtual switches of this cluster are allowed until this operation is complete.

Back Cancel Save

Uncheck the box next to the uplinks to be removed from this virtual switch. Expand the section for each host, unselecting the uplinks to be removed from all the hosts, then click Save.



Monitor rolling reboots until completed

M6-3node-dualVIC | Tasks | 7 | 1

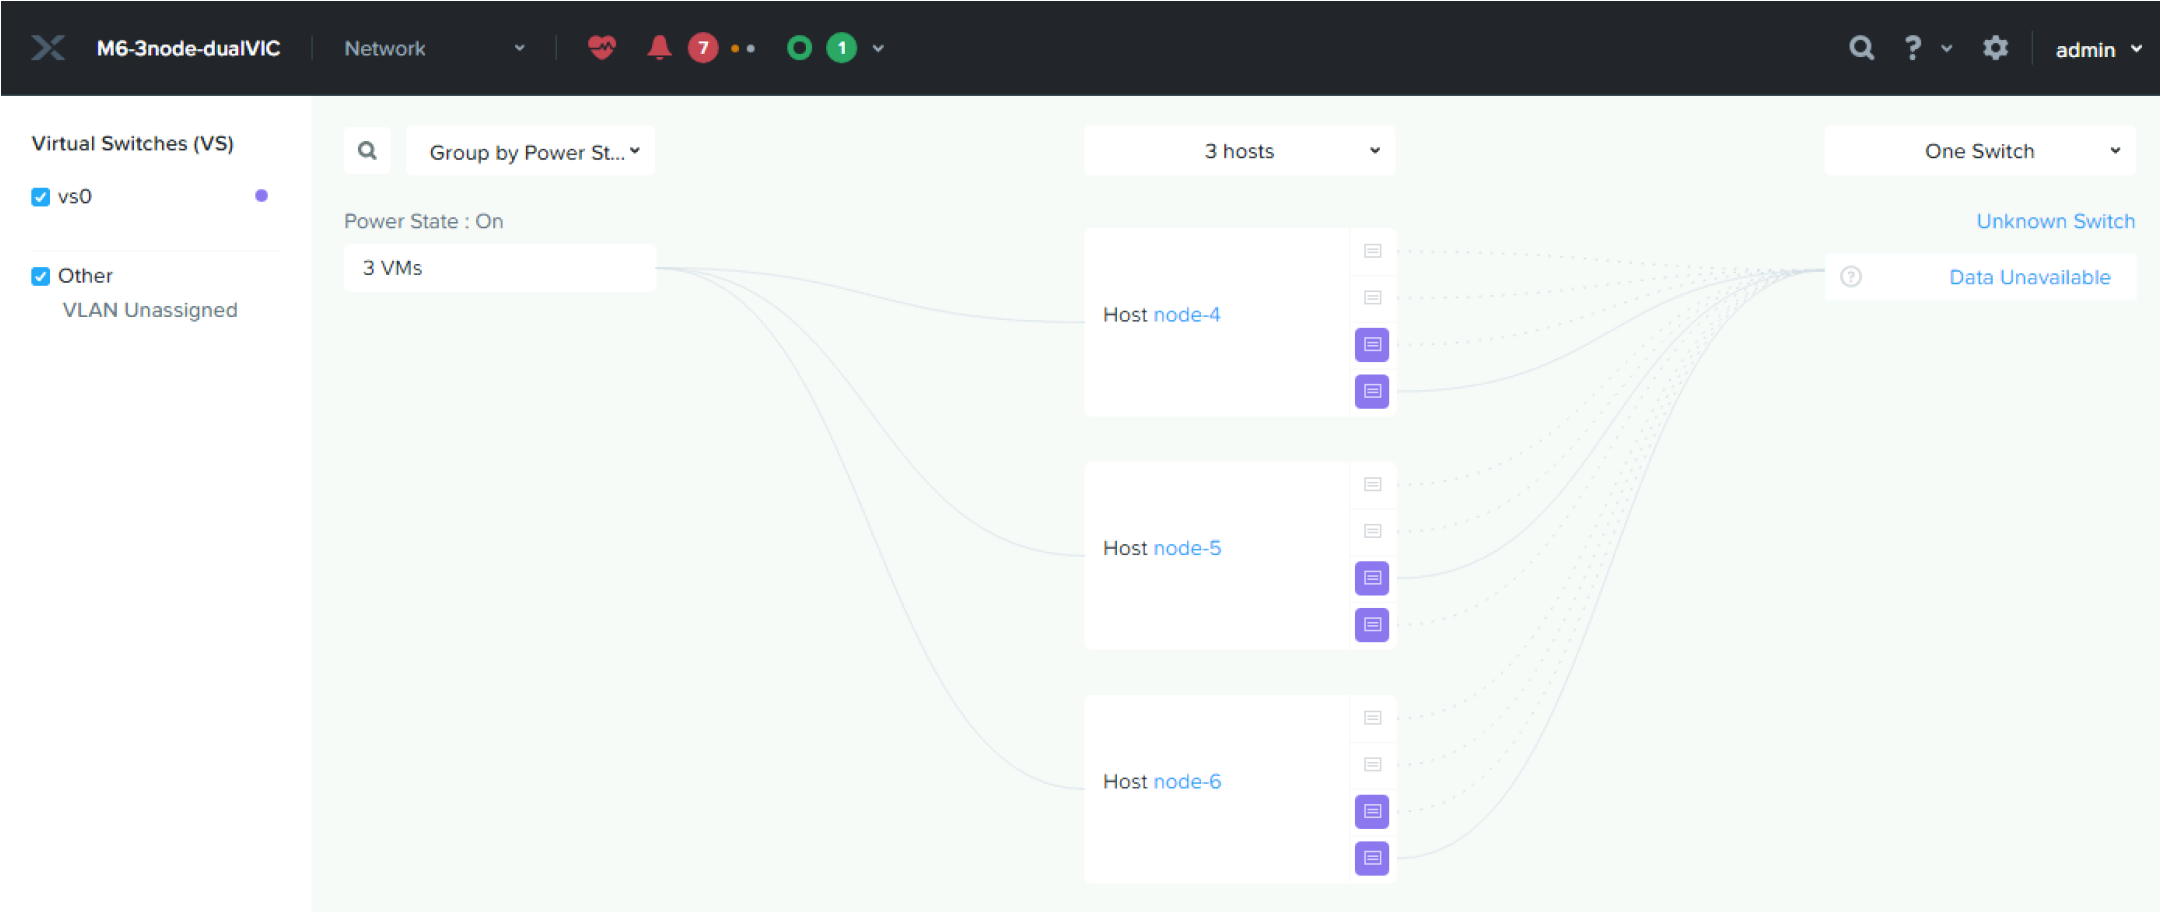
Overview

Type text to filter by

Viewing all 16 Tasks

Task	Entity Affected	Progress	Status
Hypervisor rolling restart	Cluster Details	<div style="width: 28%;"><div style="width: 28%;"></div></div> 28%	Running

Verify the nodes now only use 2 uplinks



Create a new virtual switch

The screenshot shows the VMware vSphere interface. At the top, the navigation bar includes 'M6-3node-dualVIC' and a 'VM' dropdown menu. Below the navigation bar, there are tabs for 'Overview' and 'Table'. A '+ Create VM' button is visible, along with a 'Network Config' button highlighted with a red box. Below the navigation, there is a 'VM' tab and a search bar. A table lists VMs with columns for VM Name, Host, IP Addresses, Cores, Memory Capacity, Storage, CPU Usage, Memory Usage, Controller Read IOPS, Controller Write IOPS, Controller IO Bandwidth, Controller Avg IO Latency, Backup..., and Flash Mode. One VM is listed: 'NTNX-node-4-CVM' on host 'node-4/AHV' with IP '10.1.51...' and 12 cores.

The screenshot shows the 'Network Configuration' dialog box with the 'Virtual Switch' tab selected. A '+ Create VS' button is highlighted with a red box. Below the button is a table with columns for Name, Bridge, MTU (bytes), and Bond Type. One entry is shown: 'vs0' with bridge 'br0', MTU '1500', and bond type 'Active-Backup'.

The screenshot shows the 'Create Virtual Switch' configuration page. It has two tabs: 'General' (selected) and 'Uplink Configuration'. The 'Virtual Switch Name' field contains 'vs2'. The 'Description' field contains 'disjoint L2 vlans'. The 'Physical NIC MTU (bytes)' field contains '1500'. Below this, there is a note: 'MTU has to be a value in the range 1280 to 9216 inclusive.' Under 'Select Configuration Method', the 'Quick' option is selected. Below the 'Quick' option, there is a description: 'This method can briefly interrupt the workloads on the cluster. The hosts will not be placed in maintenance mode in this method. Use this method only if the cluster is not running production workloads at this time.' At the bottom right, there are 'Cancel' and 'Next' buttons, with 'Next' highlighted by a red box.

Choosing the standard or quick option appears to result in rolling reboots of the hypervisors.

Create a new virtual switch continued

Create Virtual Switch

Bond Type
Active-Backup

One adapter in the bond is active. Additional adapters act as backup until the active adapter fails. **Link Aggregation must not be configured on the connected switch.**

Select Hosts
All Hosts x

Select Uplink Ports
Show

Connected and Unconnected Uplink Ports
 Only Connected Uplink Ports

On Switches (with LLDP)
All switches

Uplink Port Speeds
All Speeds

Select All Clear All

Hosts	10G NICs
node-4	4 ports
	<input checked="" type="checkbox"/> eth0
	<input checked="" type="checkbox"/> eth1

< Back Cancel Create

Leave the bond type as Active-Backup.

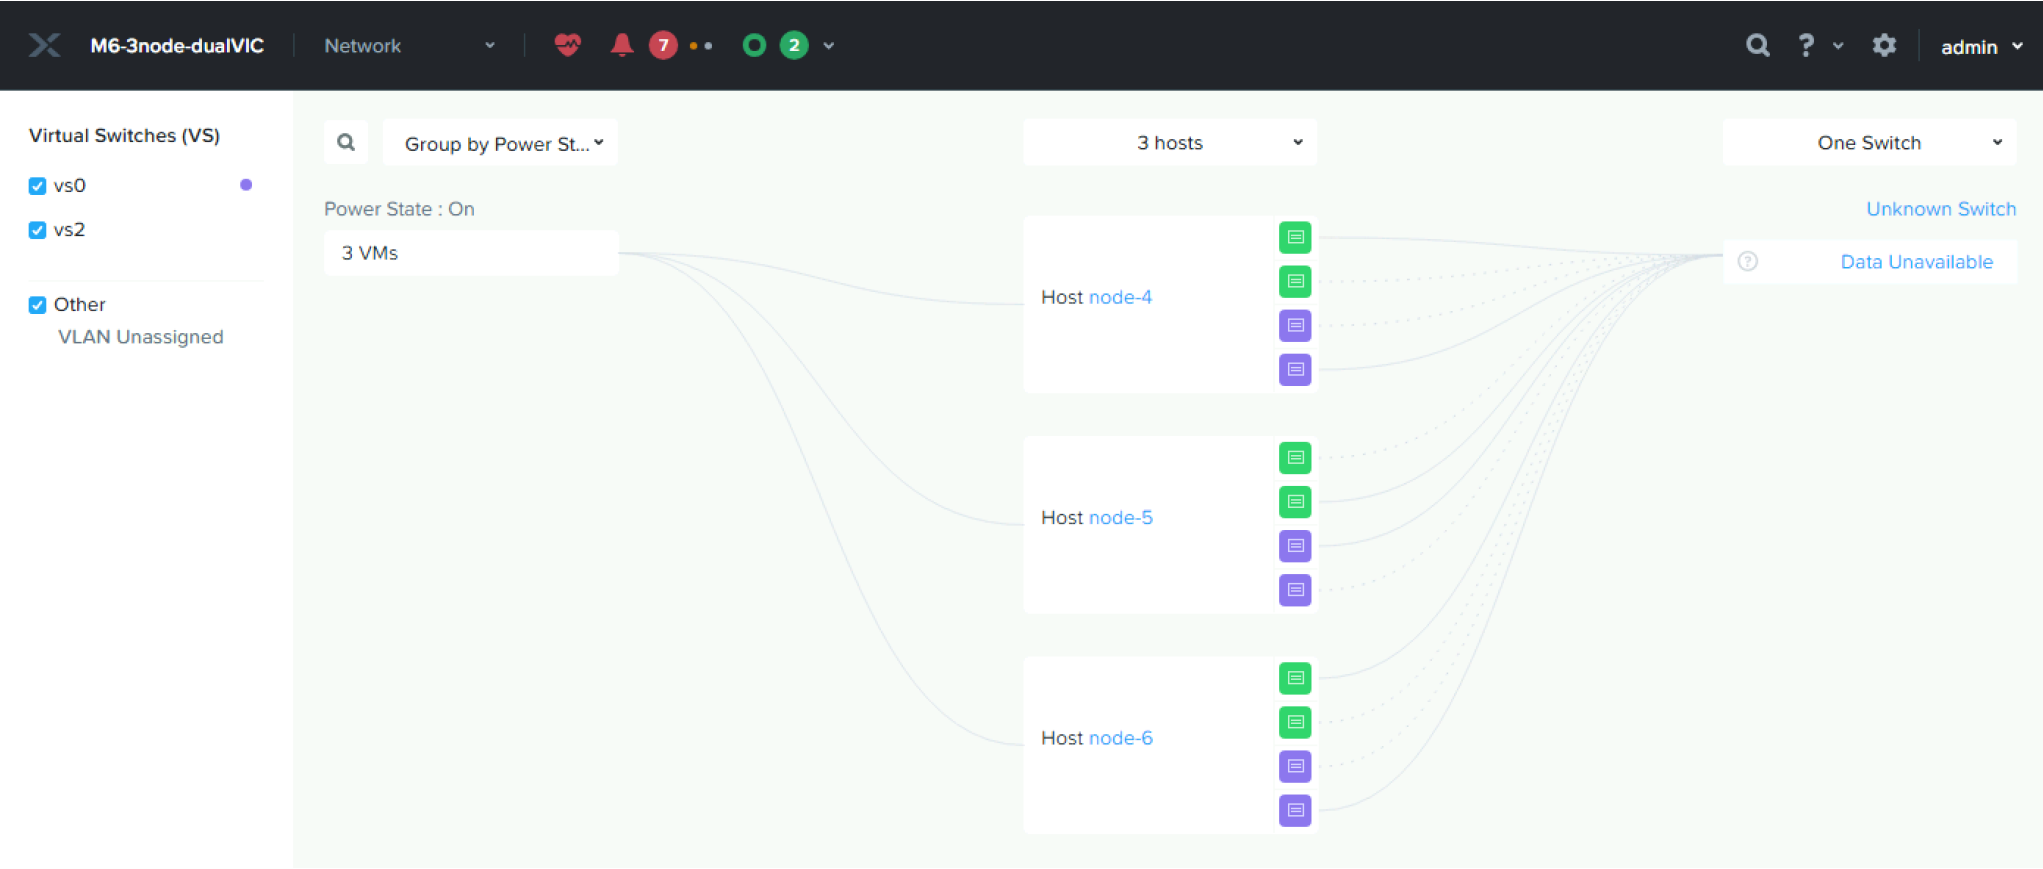
Expand each node and select the two uplinks which were removed from the default virtual switch. Click Create.

Monitor rolling reboots until completed

The screenshot shows a management interface for a device named 'M6-3node-dualVIC'. The 'Tasks' section is active, displaying a list of tasks. A search bar is present with the placeholder text 'Type text to filter by'. Below the search bar, it indicates 'Viewing all 16 Tasks'. A table lists the tasks, with the first row showing a task in progress.

Task	Entity Affected	Progress	Status
Hypervisor rolling restart	Cluster Details	<div style="width: 28%;"><div style="width: 28%;"></div></div> 28%	Running

Verify the configuration



Create VM Subnet(s)

Overview · Table + Create VM **Network Config**

VM Include Controller VMs · 3 VMs · search in table

VM Name	Host	IP Addresses	Cores	Memory Capacity	Storage	CPU Usage	Memory Usage	Controller Read IOPS	Controller Write IOPS	Controller IO Bandwidth	Controller Avg IO Latency	Backup...	Flash Mode
NTNX-node-4-CVM	node-4/AHV	10.151...	12	32 GIB	- / 0 GIB	5.24 %	21.94 %	-	-	-	-	No ?	No

Network Configuration

Subnets Internal Interfaces Virtual Switch

No subnets have been configured.

Create Subnet

Create Subnet

Subnet Name: vlan51

Virtual Switch: vs0

VLAN ID: 0

Enable IP address management

This gives AHV control of IP address assignments within the network.

Cancel Save

Create Subnet

Subnet Name: vlan52

Virtual Switch: vs2

VLAN ID: 52

Enable IP address management

This gives AHV control of IP address assignments within the network.

Cancel Save

Create the needed subnets with the necessary VLAN IDs and associating them with the appropriate virtual switches.

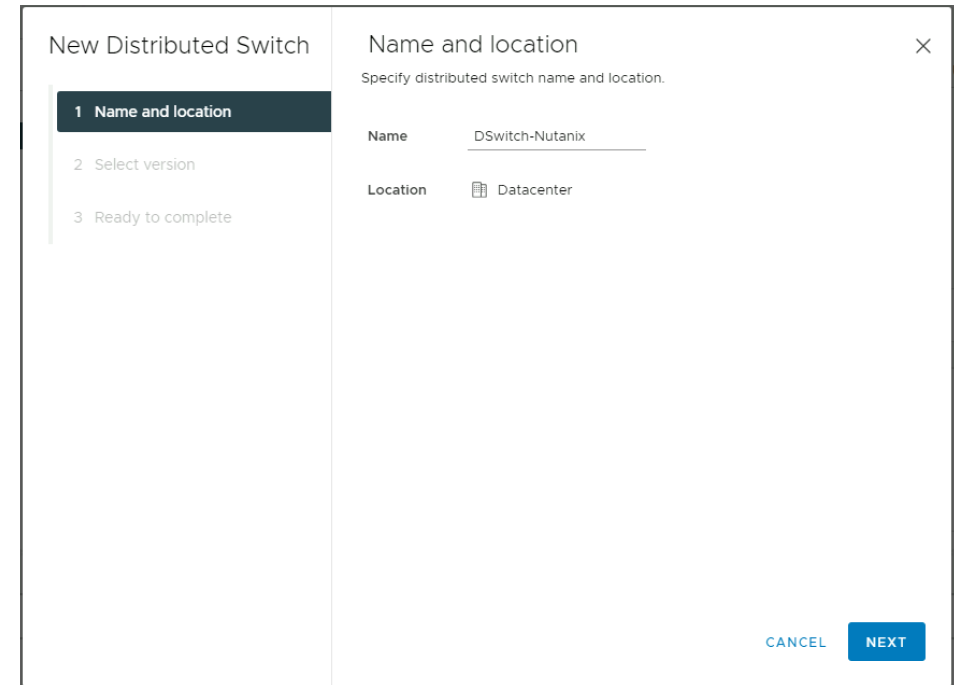
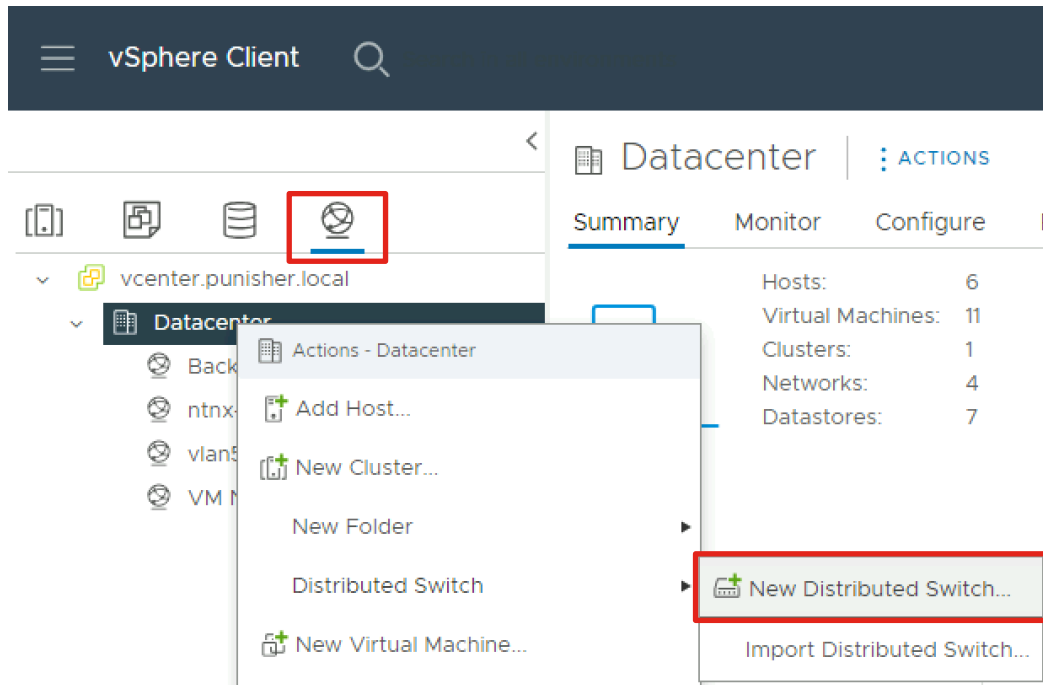
Network Configuration

Subnets Internal Interfaces Virtual Switch + Create Subnet

Subnet Name	Virtual Switch	VLAN ID	Used IP Addresses	Free IPs in Subnets	Free IPs in Pool	Actions
vlan51	vs0	0	N/A	N/A	N/A	Edit · Delete
vlan52	vs2	52	N/A	N/A	N/A	Edit · Delete

Distributed Virtual Switch Migration

Create New Distributed Virtual Switch(es)



Refer to the official documentation for migration to distributed virtual switches available here:
https://portal.nutanix.com/page/documents/details?targetId=vSphere-Admin6-AOS-v6_7:vsp-cluster-migrate-std-dist-without-lacp-vmware-c.html


Create New Distributed Virtual Switch(es) Continued

The screenshot shows the 'New Distributed Switch' wizard at the 'Select version' step. The left sidebar has four steps: '1 Name and location', '2 Select version' (highlighted), '3 Configure settings', and '4 Ready to complete'. The main area is titled 'Select version' and contains the instruction 'Specify a distributed switch version.' There are five radio button options: '7.0.3 - ESXi 7.0.3 and later' (selected), '7.0.2 - ESXi 7.0.2 and later', '7.0.0 - ESXi 7.0 and later', '6.6.0 - ESXi 6.7 and later', and '6.5.0 - ESXi 6.5 and later'. A light blue information box states: 'The multicast filtering mode on the switch will be set to IGMP/MLD snooping if you continue with the selected version.' Below this is a link for 'Features per version'. At the bottom are 'CANCEL', 'BACK', and 'NEXT' buttons.

The screenshot shows the 'New Distributed Switch' wizard at the 'Configure settings' step. The left sidebar has four steps: '1 Name and location', '2 Select version', '3 Configure settings' (highlighted), and '4 Ready to complete'. The main area is titled 'Configure settings' and contains the instruction 'Specify number of uplink ports, resource allocation and default port group.' A red box highlights the configuration fields: 'Number of uplinks' (set to 2), 'Network I/O Control' (set to Disabled), 'Default port group' (checked 'Create a default port group'), and 'Port group name' (set to vlan51-pg). At the bottom are 'CANCEL', 'BACK', and 'NEXT' buttons.

Create the new vDS for the cluster with 2 uplinks and a default port group to match the existing port group already in place. If you are not building a disjoint L2 setup and have dual Cisco VIC cards, you can configure with 4 uplinks. For a disjoint L2 setup, create a second vDS with 2 uplinks, also with port groups for the disjoint L2 VLANs, set with the correct VLAN IDs.

Create Additional Distributed Port Groups



New Distributed Port Group

Name and location
Specify distributed port group name and location.

- Name and location**
- Configure settings
- Ready to complete

Name	management-pg
Location	DSwitch-Nutanix

New Distributed Port Group

Configure settings
Set general properties of the new port group.

- Name and location
- Configure settings**
- Ready to complete

Port binding	Static binding
Port allocation	Elastic ⓘ
Number of ports	8
Network resource pool	(default)
VLAN	
VLAN type	None
Advanced	
<input type="checkbox"/> Customize default policies configuration	

CANCEL BACK NEXT

Create distributed port groups for management, vmotion and the Nutanix storage backplane traffic. For a disjoint L2 configuration, ensure that the port groups are created on the appropriate vDS, and that their VLAN type and ID are properly set.

- DSwitch-Nutanix
 - backplane-pg
 - DSwitch-Nutanix-DVUplinks-...
 - management-pg
 - vlan51-pg
 - vmotion-pg

Edit Distributed Port Group Settings



backplane-pg | ACTIONS

Summary Monitor **Configure** Permissions Ports Hosts VMs

Settings

- Properties
- Policies**
- Traffic filtering and marking
- Alarm Definitions

Policies

Security		
Promiscuous mode		Reject
MAC address changes		Reject
Forged transmits		Reject
Ingress traffic shaping		
Status		Disabled
Average bandwidth		--
Peak bandwidth		--
Burst size		--
Egress traffic shaping		
Status		Disabled
Average bandwidth		--
Peak bandwidth		--
Burst size		--
VLAN		
Type		None



Distributed Port Group - Edit Settings

backplane-pg

General

Advanced

VLAN

Security

Traffic shaping

Teaming and failover

Monitoring

Miscellaneous

Load balancing Route based on originating virtual port

Network failure detection Link status only

Notify switches Yes

Failback Yes

Failover order

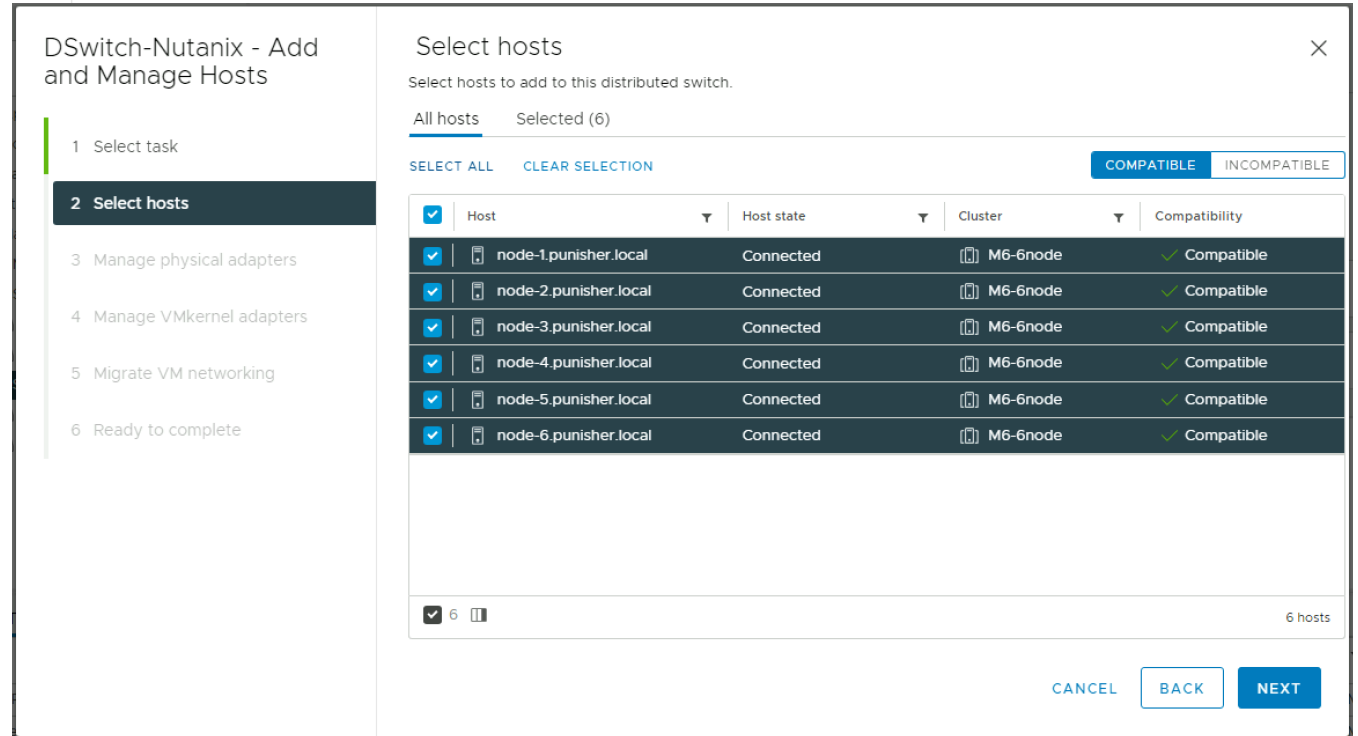
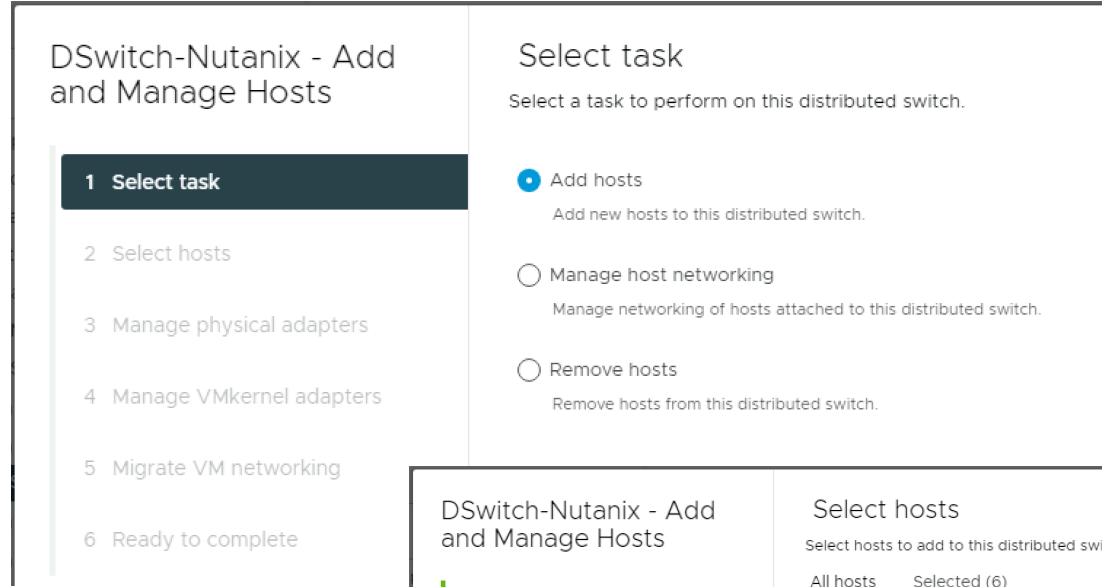
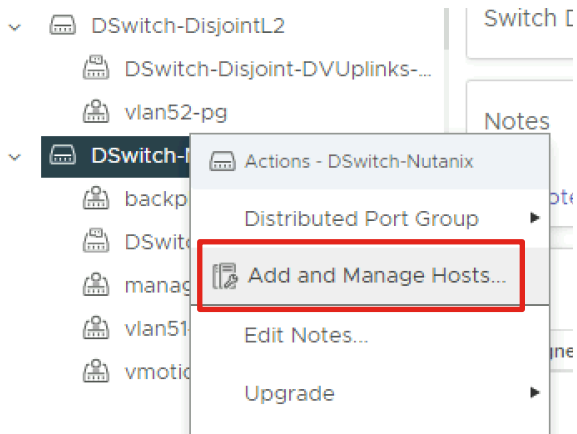
MOVE UP MOVE DOWN

- Active uplinks
- Uplink 1
- Standby uplinks
- Uplink 2
- Unused uplinks

CANCEL OK

Edit the uplink failover order for all the distributed port groups to use uplink 1 as active and uplink 2 as standby.

Add Hosts to the Distributed Switch



Add Hosts to the Distributed Switch Continued

DSwitch-Nutanix - Add and Manage Hosts

- Select task
- Select hosts
- Manage physical adapters**
- Manage VMkernel adapters
- Migrate VM networking
- Ready to complete

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

Adapters on all hosts Adapters per host

To associate a physical network adapter with an uplink, use "Assign uplink". This assignment would be applied to all the hosts that have the same physical network adapter available.

Physical network adapters	In use by switch	Assign uplink
>> vmnic0	This switch	Uplink 1
>> vmnic1	This switch	Uplink 2
>> vmnic2	6 hosts / 6 switches	None
>> vmnic3	6 hosts / 6 switches	None

4 physical network adapters

CANCEL BACK NEXT

Select vmnic0 and vmnic1 as uplinks 1 and 2 in a standard cluster with only two vNICs, or in a cluster with four vNICs in a disjoint L2 config. For a cluster with four vNICs not using disjoint L2, you can select all four vmnics here.

Add Hosts to the Distributed Switch Continued

DSwitch-Nutanix - Add and Manage Hosts

- Select task
- Select hosts
- Manage physical adapters
- Manage VMkernel adapters**
- Migrate VM networking
- Ready to complete

Manage VMkernel adapters

Manage and assign VMkernel network adapters to the distributed switch.

Adapters on all hosts Adapters per host

To assign vmkernel network adapter to port group, click on the arrow or "Assign port group" button. This assignment would be applied to all the hosts that have the same vmkernel network adapter available.

Name	In use by switch	Destination port group
>> vmk0	This switch	management-pg
>> vmk1	6 hosts / 6 switches	ASSIGN PORT GROUP

2 Vmkernel network adapters

CANCEL BACK NEXT

Assign vmkernel port vmk0 to the management distributed port group. **DO NOT** modify any setting related to vmk1.

Add Hosts to the Distributed Switch Continued

DSwitch-Nutanix - Add and Manage Hosts

- Select task
- Select hosts
- Manage physical adapters
- Manage VMkernel adapters
- Migrate VM networking**
- Ready to complete

Migrate VM networking

Select virtual machines or network adapters to migrate to the distributed switch.

Migrate virtual machine networking

Configure per network adapter Configure per virtual machine

Select an individual virtual machine to migrate its network to different source network

Virtual machine	Network adapter	Source port group	Destination port group
>> NTNX-WMP2721002A-A	Network adapter 1	VM Network	 management-pg ×
>> NTNX-WMP2721002A-A	Network adapter 2	ntnx-internal-pg	ASSIGN PORT GROUP
>> NTNX-WMP2721002A-A	Network adapter 3	Backplane Network	 backplane-pg ×
>> vlan51-vm	Network adapter 1	VM Network	 vlan51-pg ×
>> NTNX-WMP2721004X-A	Network adapter 1	VM Network	 management-pg ×
>> NTNX-WMP2721004X-A	Network adapter 2	ntnx-internal-pg	ASSIGN PORT GROUP
>> NTNX-WMP2721004X-A	Network adapter 3	Backplane Network	 backplane-pg ×

20 Virtual machines

CANCEL BACK NEXT

For the Nutanix controller VMs, migrate network adapter 1 to the management distributed port group, and network adapter 3 to the backplane distributed port group. **DO NOT** modify any setting related to network adapter 2 on the controller VMs. Also, migrate the adapters of any guest VMs already running on the system.

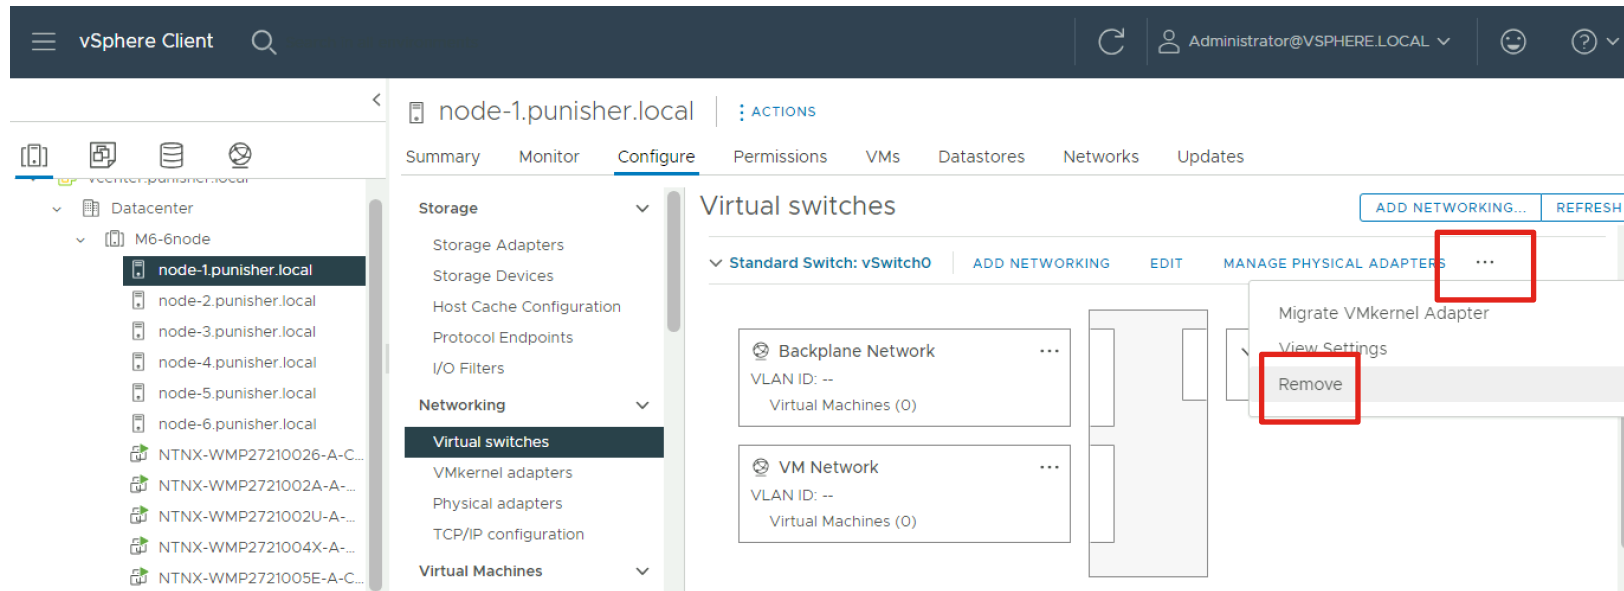
Add Hosts to the Distributed Switch Complete

The screenshot shows a wizard window titled "DSwitch-Nutanix - Add and Manage Hosts". On the left, a vertical progress bar indicates six steps: 1. Select task, 2. Select hosts, 3. Manage physical adapters, 4. Manage VMkernel adapters, 5. Migrate VM networking, and 6. Ready to complete. The "Ready to complete" step is highlighted in dark blue. The main content area is titled "Ready to complete" and includes a close button (X) in the top right corner. Below the title, it says "Review your settings selections before finishing the wizard." There are two expandable sections: "Number of managed hosts" showing "Hosts to add" as 6, and "Number of network adapters for update" showing "Physical adapters" as 12, "Reassigned VMkernel adapters" as 6, and "Virtual machine adapters" as 13. At the bottom right, there are three buttons: "CANCEL" (light blue), "BACK" (blue outline), and "FINISH" (green).

Section	Item	Count
Number of managed hosts	Hosts to add	6
	Number of network adapters for update	
Number of network adapters for update	Physical adapters	12
	Reassigned VMkernel adapters	6
	Virtual machine adapters	13

For a disjoint L2 configuration, you must also add the hosts to the second vDS after adding them to the first. For the second vDS, there are no vmkernel ports to migrate, only guest VM adapters. For the second vDS, you choose vmnic2 as uplink 1 and vmnic3 as uplink 2.

Remove Standard Virtual Switches



The standard vSwitch0 can be removed, along with any other standard virtual switches you may have manually created. **DO NOT** modify or remove the standard virtual switch named vSwitchNutanix.



The bridge to possible