# ThinkAgile Network Orchestrator

Last Update: December 2018

**Introduces PRISM and vSphere integration and their capabilities**

**Explains configuration of hypervisor integration**

**Shows how manual configuration is automated**

**Explains network topologies for use with PRISM and vSphere**

**B. Scott Lorditch**
**Dillibabu Kodamala**
**Aksh Mittal**
**Chidambaram Bhagavathiperumal**
**Alfred Yang**
**Gangadhar Hariharan**

LENOVO PRESS

# Abstract

Lenovo® ThinkAgile™ Network Orchestrator is a feature of Lenovo RackSwitch® CNOS (Cloud Network OS) network switch firmware that provides the ability for the infrastructure or server administrator to modify the network as needed to support typical tasks, such as the creation, starting, and shutdown of new machines, as well as manipulation of guest virtual machines on VLANs. These tasks are performed through the Nutanix PRISM dashboard and Acropolis hypervisor (AHV), or the VMware vSphere dashboard with the ESXi hypervisor, with the switches configured to detect the changes and act upon them.

The purpose of this software is to enable the infrastructure or server administrator to see a single pane of glass and perform his or her tasks without needing to know details of network switch management. The switches learn of changes in the environment from PRISM or vSphere and dynamically modify their configurations as needed. This feature also enables the networking team to manage security and traffic prioritization of individual virtual machines if desired.

This paper covers the functionality and configuration of ThinkAgile Network Orchestrator. The paper assumes some level of familiarity with CNOS as well as with the Nutanix and VMware hypervisors and their respective management consoles.

**Note**: The release of Nutanix AOS 5.1.2 inadvertently removed the network APIs so it cannot be used with ThinkAgile Network Orchestrator. The APIs will be restored in AOS 5.1.3. ThinkAgile Network Orchestrator is supported with Nutanix AOS v5.1.1.3 or 5.1.3.

**Do you have the latest version?** We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF.

# Contents

# Introduction

In traditional IT environments, an infrastructure administrator relies on network administrators to provision and operate the network, and troubleshoot network related issues. The boundary of the network typically starts upwards from a server NIC. However, the dynamics are different in a virtualized solution.

Cluster architectures for virtualized solutions typically include at least two pairs of top-of-rack (ToR) network switches. As a result, the network infrastructure administrator typically takes on the role of managing the entire solution. It is therefore useful and appropriate to reduce the complexity associated with the tasks of network configuration for a cluster. Provisioning a new host or guest in the cluster will still change the network, but network configuration changes should be automated where possible. The network should integrate seamlessly with the workload/cloud management system.

ThinkAgile Network Orchestrator enables such integration by enabling actions performed on the Nutanix PRISM management console and/or the vSphere management console, when the ESXi hypervisor is chosen, to make appropriate changes in the physical and logical network to which the host nodes are attached. It is possible to use the ESXi hypervisor with or without the use of Nutanix, which can provide hyperconverged storage functionality when using either of those hypervisors.

The architecture of the ThinkAgile Network Orchestrator is shown in Figure 1.

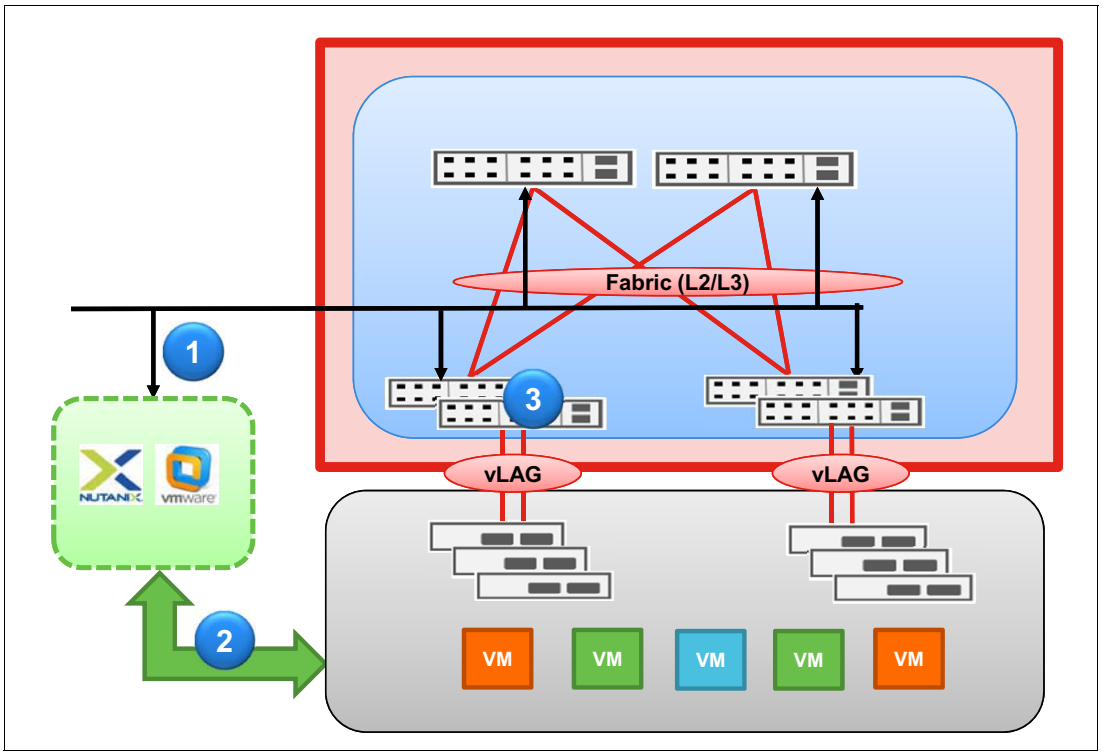


*Figure 1 ThinkAgile Network Orchestrator - Architectural flows*

The steps in this figure are:

1. Integration with the Acropolis or ESXi management portals for notifications and topology discovery.

   **Note:** At present, only the Acropolis and ESXi hypervisors are supported; Microsoft Hyper-V is not currently supported.

   The following notifications occur:

   – VM creation, update, and deletion and power on/off

   – Virtual network creation, update, and deletion

   – Virtual NIC attach and detach from virtual networks and IP address assignment

2. The Virtual Machine Manager (VMM) controls and manages the virtual network infrastructure management and VM workload and overlay provisioning.

3. Workload Visibility and Dynamic Network Provisioning

   – Workload visibility includes VM information and overlay network information. Topology discovery associates VMs with the physical switch interface.

   – Dynamic network provisioning occurs when:

     • The physical network (L2 VLAN) is provisioned based on virtual network and VM associations.

     • CNOS firmware dynamically handles VM migrations and moves L2 VLAN configurations.

   – Configured prioritization, access, and queueing filters (ACLs) will be activated on the appropriate port whenever the guest VM for which they are configured becomes active or migrates to that port. They will be deactivated when the VM is no longer active on the port.

The Cloud Network OS (CNOS) switch firmware integration agent gives increased visibility of the virtual infrastructure, including workload (VM) information and virtual network information. It also provides auto-discovery of any changes to the virtual network topology.

The CNOS firmware integration agent also provides automatic VM-aware provisioning. The physical network is automatically configured based on the auto-discovered topology. The agent provides ongoing, dynamic updates to the physical network configuration in response to new VMs, updated VMs, and deleted VMs, eliminating errors with manual configuration.

The CNOS integration is shown in more detail in Table 1.

*Table 1 Advantages of using the CNOS integration agent*

| Action | Without integration agent | With integration agent |
|---|---|---|
| Initial Network Deployment | Network Administrator engaged. Switch provisioned manually or via automation tools (preplanning required). | Lenovo hyperconverged systems may be delivered to customer preconfigured. |
| Network settings update based on VM and vNW configuration | Pending changes are communicated to the Network Administrator. The Administrator manually reconfigures each physical switch connected to a new VM. This process repeats each time changes in virtual network occurs. | CNOS integration agent communicates with VMM to automatically learn when VM changes occur and update the switch configurations accordingly. No scheduled maintenance window or user intervention is required. |

| Action | Without integration agent | With integration agent |
|---|---|---|
| Live Migration | The Network Administrator must locate switches and update the port settings for the migrated VM. | CNOS integration agent gets notified of new VM settings and automatically syncs the switch configuration. |
| Virtual Network Visibility | Switch has no visibility into the virtual network. | From the switch, details of the entire virtual network infrastructure are available via CLI. |

# Implementation using Nutanix Acropolis

ThinkAgile Nutanix Integration uses a communication channel between PRISM and the Acropolis hypervisor and the Lenovo switches that are part of the cluster where the host nodes are attached.

The communications channel provides the following functions:

- Visibility of the cluster
  - The virtual infrastructure (VLANs and virtual machines) is known to the switches.
  - The switches perform discovery of the vNetwork topology that is known to the hypervisors
- Automatic VM-aware provisioning
  - The physical network is configured based on the discovered topology as it is learned and as changes to it are learned by the switches
  - Ongoing dynamic updates to the physical network configuration are known to the switches, including new, updated, and deleted VMs and new, updated, and emptied or deleted VLANs.

The switches receive information on the following events which occur within the cluster:

- VM power on or off
- VM NIC plugged in to or removed from a vNetwork
- vNetwork creation or deletion
- VM migration to a different host
- Modifications made to a vNetwork – addition, change, or deletion of the vNetwork
- Modifications to a VM such as a change in its name

## Setting up the vNetwork and VM

Illustrations of these functions, and the resulting messages and displays from a switch, are shown below. The screens and messages result from creating a vNetwork, a new VM, assigning that VM to the vNetwork, and powering the VM on.

The following steps show the creation of a new VLAN and the attachment of a new guest VM to that VLAN when PRISM is used in concert with the Acropolis hypervisor.

1. Create vNetwork 9 as shown in Figure 2.

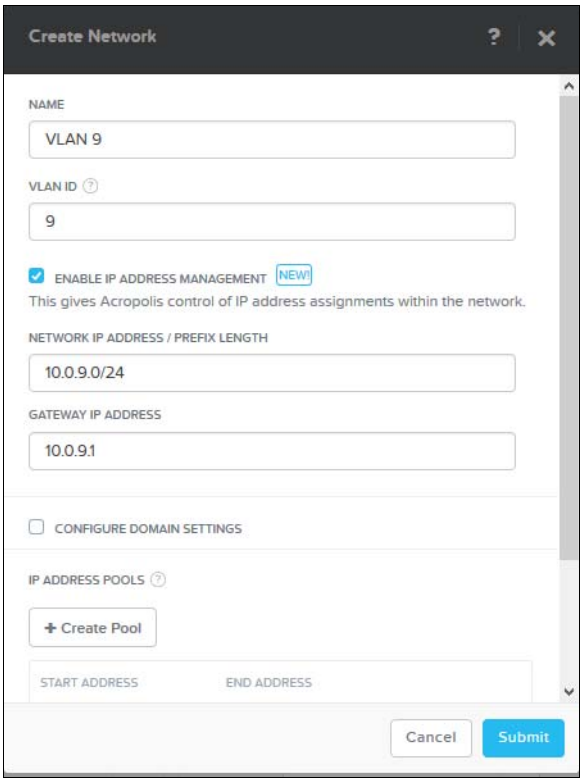


*Figure 2 Create vNetwork 9*

2. Verify vNetwork 9 successfully created, as shown in Figure 3.

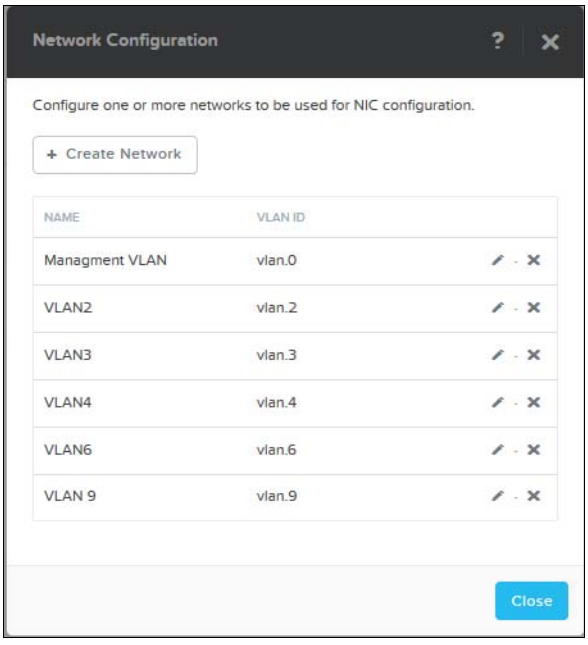


*Figure 3 Verify vNetwork 9 successfully created*

3. Create a new guest, “New VM”, Figure 4.

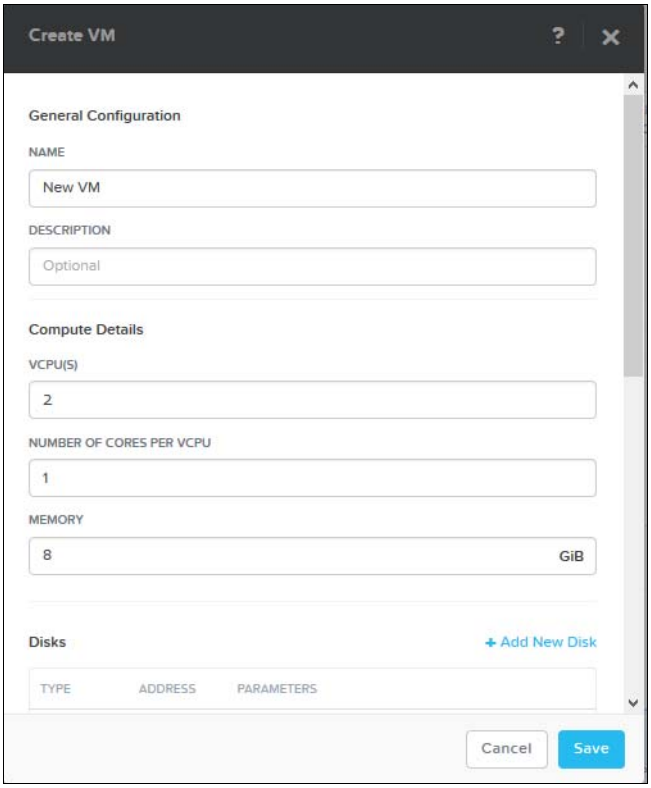


*Figure 4 Step 3: Create a new guest*

4. Define a network attachment (NIC) for the new VM, Figure 5

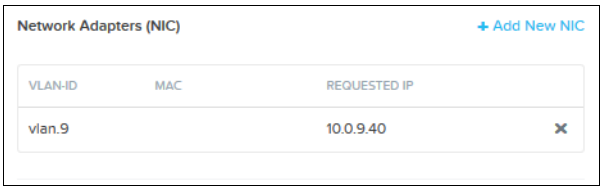


*Figure 5 Step 4: Define Network attachment*

5. Verify successful VM creation and then power it on, Figure 6.

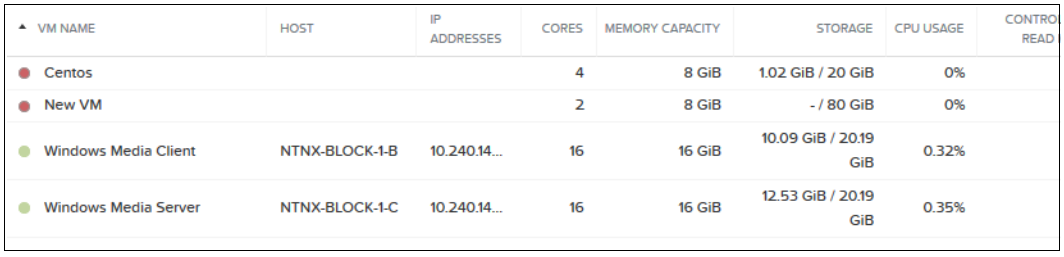


*Figure 6 Step 5: Verify successful VM creation*

6. Review switch log messages resulting from above steps, Figure 7

```
Switch-2#show log last 10
86915 2017-04-04T21:07:40+00:00 Switch-2(data) %NPA-5-VNETWORK_CREATE: Vnetwork 9
with name VLAN 9 with UUID 560e08a1-8c61-47f7-99a5-947090794b0c is created
86916 2017-04-04T21:16:14+00:00 Switch-2(data) %NPA-5-VM_CREATE: VM New VM with UUID
c1dcaea9-1745-4227-a0e8-1f99d4e50ed5 is created
86917 2017-04-04T21:17:03+00:00 Switch-2(data) %NPA-5-VM_NIC_PLUG: VM New VM with
UUID c1dcaea9-1745-4227-a0e8-1f99d4e50ed5 received NIC Plug event
86918 2017-04-04T21:19:01+00:00 Switch-2(data) %NPA-5-VM_ON: VM New VM with UUID
c1dcaea9-1745-4227-a0e8-1f99d4e50ed5 is on
86919 2017-04-04T21:19:06+00:00 Switch-2(data) %VLAN-6-VLAN_CREATE: VLAN 9 created
86920 2017-04-04T21:19:06+00:00 Switch-2(data) %NPA-5-DYNAMIC_VLAN_CREATE: VLAN 9
created
```

*Figure 7 Switch log messages*

Notes on switch messages in Figure 7:

- The first “network creation” message is in response to the creation of VLAN 9 on the PRISM console.
- The two VLAN creation messages at the bottom of the display result from when the new VM is connected to VLAN 9 and powered on – this is the first point at which traffic would need to flow on VLAN 9.

## Switch Configuration and Commands with Acropolis

The network switches that are part of the cluster require configuration to enable integration with PRISM. A sample configuration is shown and discussed in this section.

The supported versions of Lenovo CNOS and Nutanix AOS (Acropolis) are shown in Table 2.

**Note:** Microsoft Hyper-V is currently not supported.

*Table 2 Nutanix/CNOS support matrix*

| | Lenovo CNOS firmware | Nutanix AOS versions | VMware ESXi version | VMware vCenter version |
|---|---|---|---|---|
| R2.5 MR | 10.3.2.0/R2.5 | 5.0.2 | Not applicable | Not applicable |
| R3 SQAE (5/19/2017) | 10.4.1.0/R3 | 5.0.2 | Not applicable | Not applicable |
| R2.5 MS MR SQAE (5/31/2017) | 10.3.3.0/R2.5 MR | 5.0.2<br>5.1 | Not applicable | Not applicable |
| R3 MR (6/30/2017) | 10.4.2.0/R3 MR | 5.0.2<br>5.1<br>5.1.1<br>5.1.3<br>5.1.3.2 | Not applicable | Not applicable |
| R3 MR (10/09/2017) | 10.4.5.0 | 5.0.2<br>5.1<br>5.1.1<br>5.1.3<br>5.1.3.2 | Not applicable | Not applicable |

| | Lenovo CNOS firmware | Nutanix AOS versions | VMware ESXi version | VMware vCenter version |
|---|---|---|---|---|
| R5 (01/26/2018) | 10.6.1.0/10.6.2.0 | 5.0.2<br>5.1<br>5.1.1<br>5.1.3<br>5.1.3.2<br>5.5.0.4[a] | Not applicable | Not applicable |
| R5.5 (3/30/2018) | 10.7.1.0, 10.7.2.0 | 5.0.2<br>5.1<br>5.1.1<br>5.1.3<br>5.1.3.2<br>5.5.0.4[a]<br>5.5.0.5[a]<br>5.6 | Not applicable | Not applicable |
| R5.5 MR vSwitch | 10.7.1.0 | 5.1.3.2<br>5.1.4 | ESXi 6.5.0b 5146846 | 6.5.0e Build 5705665 |
| R5.5 MR vSwitch | 10.7.1.0/10.7.2.0 | 5.1.3.2<br>5.1.4<br>5.5.0.5<br>5.6 | ESXi 6.5 U1 6765664 | 6.5 U1 5973321 |
| R6 vSwitch | 10.8.1.0 | 5.5.0.5<br>5.6<br>5.6.1 | ESXi 6.5 U1 6765664 | 6.5 U1 5973321 |
| R6 dvSwitch | 10.8.1.0 | 5.5.0.5<br>5.6<br>5.6.1 | ESXi 6.5 U1 6765664 | 6.5 U1 5973321 |
| R7/R7MR | 10.9.1.0/10.9.3.0 | Up to 5.8.2<br>5.9.2[b]<br>5.10 | ESXi 6.5 U1 6765664 | 6.5 U1 5973321 |

a. See Support tip below.
b. AOS 5.9.1 not supported

**Support tip:** VM Description Field is Required

**Description of Problem**: Without the VM description on PRISM for user VMs the ThinkAgile Network Orchestrator refresh mechanism would fail for these VMs, this would not update the VM database and not create VLANs attached to these VMs.

**Root Cause**: The POST response for Nutanix API https://<cluster-ip>:9440/api/nutanix/v3/vms/list has “null” value in the VM description field. As a result, VM without description would not be updated when ThinkAgile Network Orchestrator is configured as the refresh mechanism for these user VMs would fail.

**Workaround**: The VM description field on the PRISM should have valid string (name or description) associated to the user VMs.

The topology for the switch configuration is shown in Figure 8.

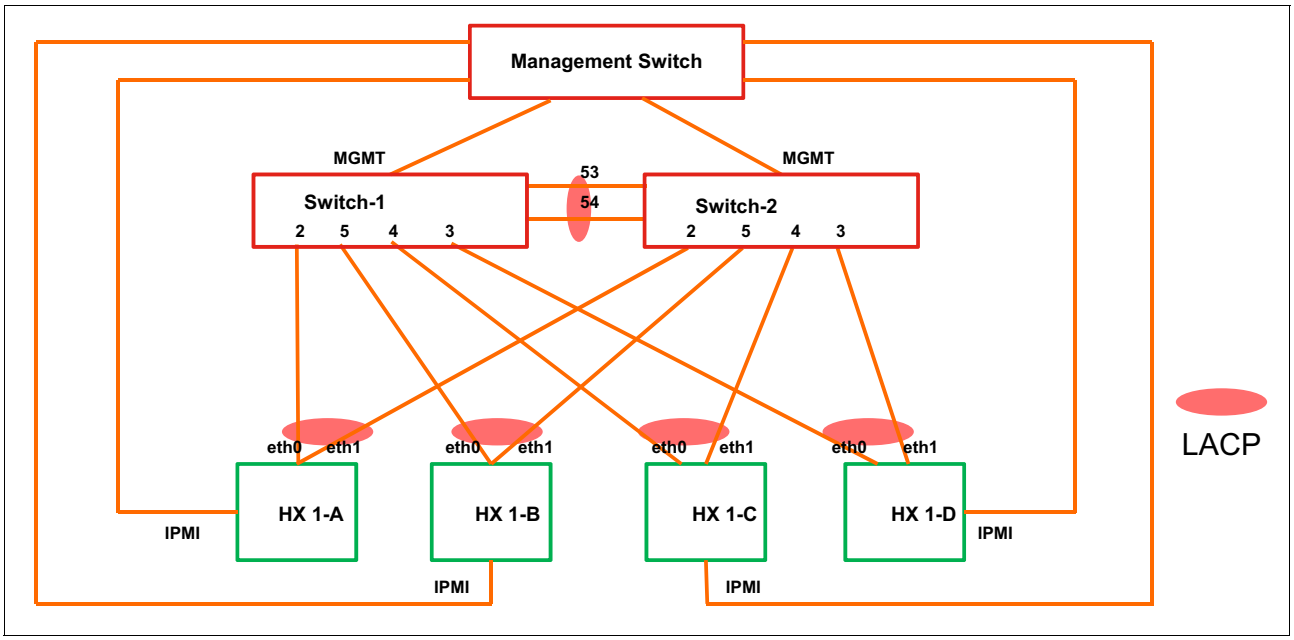


*Figure 8 Test network topology*

The sample configuration is shown in Figure 9.

```
Switch-2#show run vdm
!
vdm nutanix
ip address 10.240.149.233 vrf management
username admin password encrypted 0xb5e222b27807402e
add interface po 2
add interface po 3
add interface po 4
add interface po 5
refresh-vms-url https://10.240.149.233:9440/api/nutanix/v3/vms/list
refresh-vnet-url https://10.240.149.233:9440/api/nutanix/v3/networks/list
topo-discovery-url https://10.240.149.233:9440/PrismGateway/services/rest/v1/vms
subscribe server-url https://10.240.149.233:9440/api/nutanix/v3/webhooks client
-url https://10.240.149.246:443/nos/api/cfg/nutanix/network vnetwork
subscribe server-url https://10.240.149.233:9440/api/nutanix/v3/webhooks client
-url https://10.240.149.246:443/nos/api/cfg/nutanix/vm vm
Switch-2#show run vlag
!
vlag tier-id 10
vlag isl port-channel 100
vlag hlthchk peer-ip 10.240.149.247 vrf management
vlag enable
vlag instance 2 port-channel 2
vlag instance 2 enable
vlag instance 3 port-channel 3
vlag instance 3 enable
vlag instance 4 port-channel 4
vlag instance 4 enable
vlag instance 5 port-channel 5
vlag instance 5 enable
!
… port-channel interface details …
!
interface port-channel2
switchport mode trunk
auto-policy enable
!
interface port-channel3
switchport mode trunk
auto-policy enable
!
interface port-channel4
switchport mode trunk
auto-policy enable
!
interface port-channel5
switchport mode trunk
auto-policy enable
!
```

*Figure 9 Switch configuration example*

Notes on the configuration example:

- This configuration will be used on both of the switches shown in the topology diagram, except as discussed below.

- The **ip address** command specifies the address of the PRISM console and that in this case it is reached via the switch management port.
- The **username** command specifies administrator credentials which are valid on the PRISM console.
- The **add** commands specify interfaces that are used to attach to Nutanix compute nodes (hosts). These can be individual Ethernet interfaces or Port-channel interfaces. Only the interfaces specified will be reviewed for the presence of a host. In this test bed, port 2 on each switch is part of port-channel number 2 (po 2) and vLAG is in use. Ports 2-5 (and the corresponding port-channels) have hosts attached as shown in the diagram; ports 1 and 6 are available and preconfigured for additional hosts if desired.
- The **refresh** and **topo-discovery** commands specify URLs used to retrieve network and VM information from the PRISM console. Both commands are required and are in the form shown, using the IP address of the console, and port 9440.
- The **subscribe** commands specify URLs which are on the switch and are used by PRISM to send events to the switch; there is also an **unsubscribe** command with essentially identical syntax.
- These commands will be different on the other switch. The **client-url** operands will reflect the IP address of each switch where they are configured.
- The **peer-ip** address in the VLAG portion of the configuration in each switch specifies the IP address of the other switch.
- Interfaces where servers are attached need to be configured with the **auto-policy enable** operand. LLDP also needs to be enabled on these interfaces – recommended transmit interval is 5 seconds.

Figure 10 shows the information learned from PRISM about the new VM created in the steps shown in “Implementation using Nutanix Acropolis” on page 5.

```
Switch-2#show virtual-machine info interface po 3
Interface po3
uuid: c1dcaea9-1745-4227-a0e8-1f99d4e50ed5
    name: New VM
        host_reference:
            kind: host
            uuid: 5c834bc8-b382-4ab9-84a3-a57591f840e2
        num_cores_per_vcpu: 1
        hypervisor_type: AHV
        memory_size_mb: 8192
        num_vcpus: 2
        power_state: POWERED_ON
        nic_list:
            kind: network
            uuid: 560e08a1-8c61-47f7-99a5-947090794b0c
            mac_address: 50:6B:8D:76:3F:79
```

*Figure 10 VM display*

Figure 11 shows the information learned from PRISM about VLAN 9. Note that if the VLAN should become empty, or if it is deleted from the PRISM console, it would also be deleted from the switch.

```
Switch-2#show vnetworks

uuid: 560e08a1-8c61-47f7-99a5-947090794b0c
    name: VLAN 9
    vlan_id: 9
    default_gateway: 10.0.9.1
    network_address: 10.0.9.0
    prefix_length: 24
    dhcp_server_address: 10.0.9.254
    pool of IP ranges:
        range: 10.0.9.10 10.0.9.99
```

*Figure 11 Virtual Network Display*

# Implementation using VMware ESXi

When ThinkAgile Network Orchestrator is used in a cluster where the hosts are using the VMware ESXi hypervisor, there are differences in its use and in the switch configuration required. The key difference in the user interface is that many hypervisor functions are configured with vSphere rather than PRISM.

These functions include:

- Creation, migration, power on/off, modification and deletion of guest VMs.
- Creation of virtual networks, which are known as port groups in vSphere. The creation of virtual switches, including distributed virtual switches, is also done in vSphere.

PRISM can still be used for the management of a hyperconverged storage environment if desired, but this is not required.

## Setting up the VLAN and VM

This section describes the steps to create a VLAN (port group) and attach a new guest VM.

ThinkAgile Network Orchestrator functions much the same with ESXi as with Acropolis; the key difference is that the communications channel from the switches connects to vCenter rather than to PRISM.

In both cases, the switches discover and update the topology, and receive notifications of VMs being powered on and off, being migrated, and the creation and deletion of VLANs, and NICs from VMs being attached to or detached from those VLANs. Dynamic VLANs are created in essentially the same way for both hypervisors.

Use the following steps to create a new VLAN and attach a newly created guest VM to it using ESXi as the hypervisor.

1. Create new VLAN 200 as shown in Figure 12.

**Note:** The VLAN must be manually created on every host when the standard VMware switch is used. A future version of Orchestrator will add support for the VMware Distributed Switch, and it will then be possible to create the VLAN on all hosts connected to that switch in a single operation.

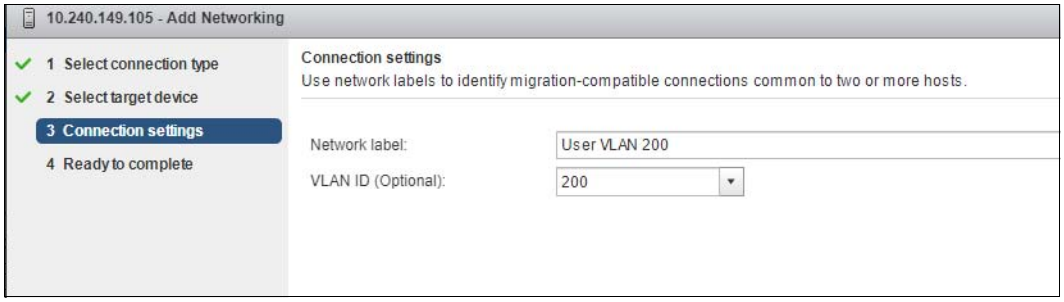


*Figure 12 Create VLAN 200*

2. Verify VLAN creation as shown in Figure 13.

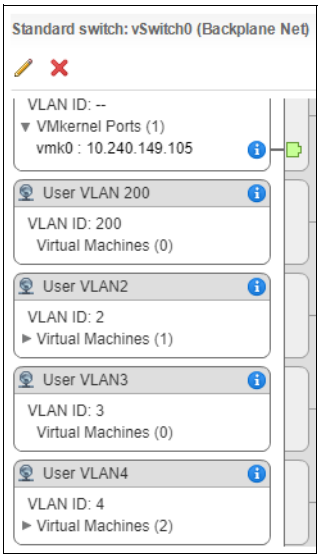


*Figure 13 Verify the VLAN*

3. Create a new guest VM by cloning an existing one as shown in Figure 14.

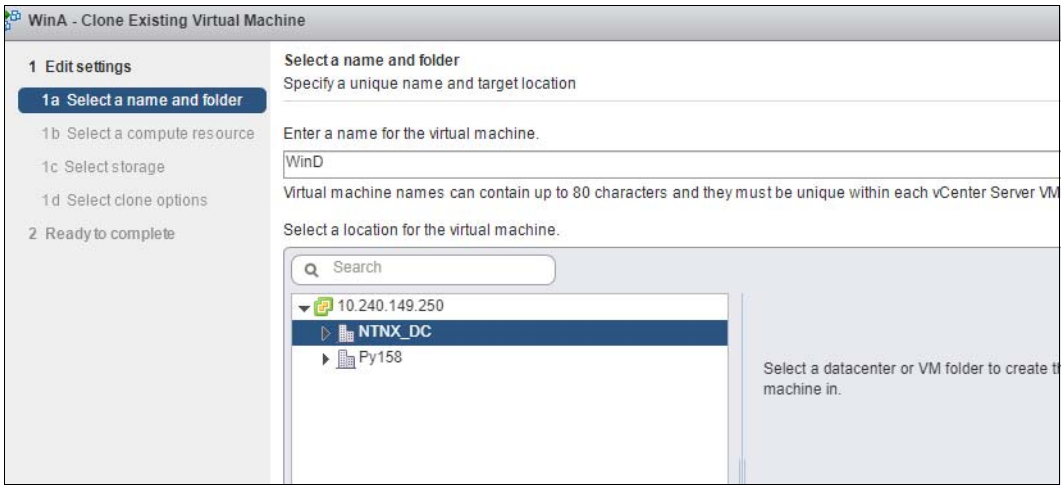


*Figure 14 Create a new VM*

4. Assign the new VM to the new VLAN 200 as shown in Figure 15.

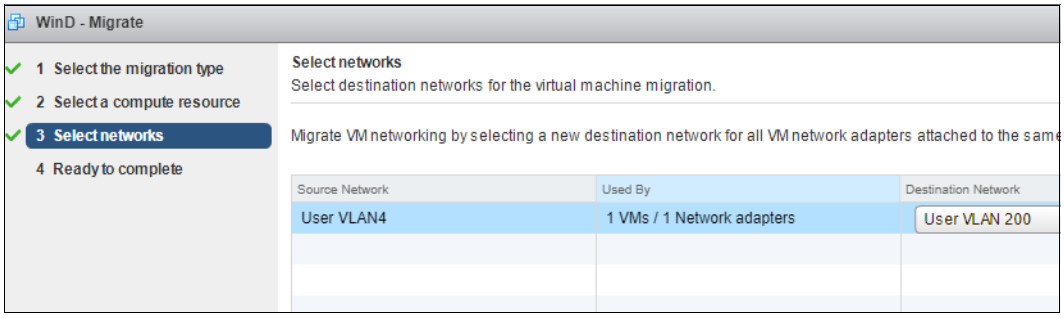


*Figure 15 Assign VLAN 200 to the new VM*

5. Verify the newly created VM – by checking on the PRISM console (which is optional), Figure 16 on page 16. A similar check can be performed through the vSphere console, whether or not PRISM is being used.

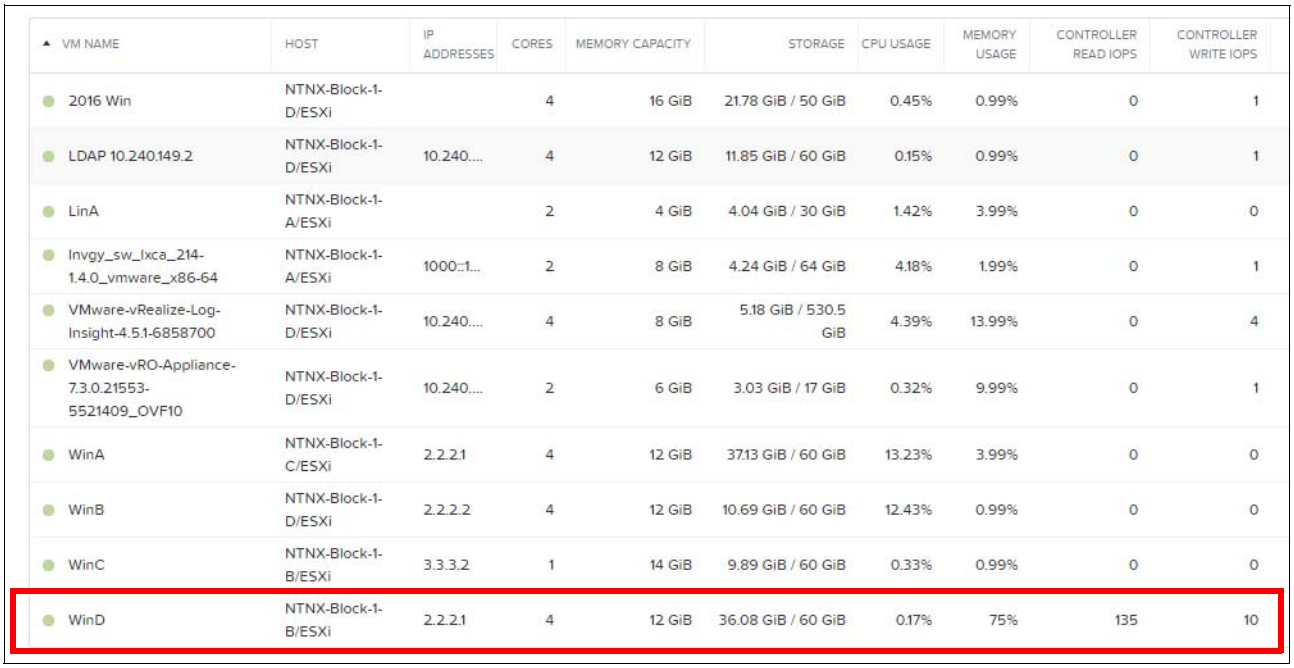

| VM NAME | HOST | IP ADDRESSES | CORES | MEMORY CAPACITY | STORAGE | CPU USAGE | MEMORY USAGE | CONTROLLER READ IOPS | CONTROLLER WRITE IOPS |
|---|---|---|---|---|---|---|---|---|---|
| 2016 Win | NTNX-Block-1-D/ESXi | | 4 | 16 GiB | 21.78 GiB / 50 GiB | 0.45% | 0.99% | 0 | 1 |
| LDAP 10.240.149.2 | NTNX-Block-1-D/ESXi | 10.240.... | 4 | 12 GiB | 11.85 GiB / 60 GiB | 0.15% | 0.99% | 0 | 1 |
| LinA | NTNX-Block-1-A/ESXi | | 2 | 4 GiB | 4.04 GiB / 30 GiB | 1.42% | 3.99% | 0 | 0 |
| Invgy_sw_lxca_214-1.4.0_vmware_x86-64 | NTNX-Block-1-A/ESXi | 1000::1... | 2 | 8 GiB | 4.24 GiB / 64 GiB | 4.18% | 1.99% | 0 | 1 |
| VMware-vRealize-Log-Insight-4.5.1-6858700 | NTNX-Block-1-D/ESXi | 10.240.... | 4 | 8 GiB | 5.18 GiB / 530.5 GiB | 4.39% | 13.99% | 0 | 4 |
| VMware-vRO-Appliance-7.3.0.21553-5521409_OVF10 | NTNX-Block-1-D/ESXi | 10.240.... | 2 | 6 GiB | 3.03 GiB / 17 GiB | 0.32% | 9.99% | 0 | 1 |
| WinA | NTNX-Block-1-C/ESXi | 2.2.2.1 | 4 | 12 GiB | 37.13 GiB / 60 GiB | 13.23% | 3.99% | 0 | 0 |
| WinB | NTNX-Block-1-D/ESXi | 2.2.2.2 | 4 | 12 GiB | 10.69 GiB / 60 GiB | 12.43% | 0.99% | 0 | 0 |
| WinC | NTNX-Block-1-B/ESXi | 3.3.3.2 | 1 | 14 GiB | 9.89 GiB / 60 GiB | 0.33% | 0.99% | 0 | 0 |
| WinD | NTNX-Block-1-B/ESXi | 2.2.2.1 | 4 | 12 GiB | 36.08 GiB / 60 GiB | 0.17% | 75% | 135 | 10 |

*Figure 16 Verify the VM in PRISM*

6. Review switch log messages that show the events in the previous steps, Figure 17.

```
%NPA-5-VNETWORK_CREATE: Vnetwork 200 with name User VLAN 200 is created on host
5c9cc05a-ae82-11e7-ab09-7ed30ae19097
%NPA-5-VM_CREATE: VM WinD with UUID 00000000-0000-0000-0000-000000000000 is created
%NPA-5-VM_NIC_PLUG: VM WinD with UUID 00000000-0000-0000-0000-000000000000 received NIC
Plug event
%NPA-5-VM_MODIFY: VM WinD with UUID 00000000-0000-0000-0000-000000000000 received modify
event for UUID 421b6b33-5254-46c4-bd13-39149cc13b38
%NPA-5-VM_NIC_CONNECTED: VM WinD with UUID 421b6b33-5254-46c4-bd13-39149cc13b38 received
NIC Connected event
%NPA-5-VM_ON: VM WinD with UUID 421b6b33-5254-46c4-bd13-39149cc13b38 is on
%NPA-5-VM_NIC_MODIFY: VM  with UUID 421b6b33-5254-46c4-bd13-39149cc13b38 received vNIC
modify event for VM Power
%NPA-5-DYNAMIC_VLAN_CREATE: VLAN 200 created
%VLAN-6-VLAN_CREATE: VLAN 200 created
%NPA-5-VM_NIC_MODIFY: VM  with UUID 421b6b33-5254-46c4-bd13-39149cc13b38 received vNIC
modify event for IP Address
```

*Figure 17 Switch log messages*

## Switch Configuration and Commands with ESXi

The supported releases of ESXi are shown in Table 3. If ESXi is being used with Nutanix, support is not limited to the releases shown in Table 2 on page 8, because the switch is communicating with vSphere rather than with PRISM.

*Table 3 ESXi / Nutanix / CNOS support matrix*

| | **R5.5 MR vSwitch** |
|---|---|
| Lenovo CNOS versions | 10.7.1.0<br>10.7.2.0 |
| Nutanix AOS versions | 5.1.3.2<br>5.1.4<br>5.5.0.5<br>5.6 |
| VMware ESXi versions | ESXi 6.5b<br>ESXi 6.5 U1 |
| VMware vSphere versions | vSphere 6.5.0e<br>vSphere 6.5 U1 |

The use of Nutanix with the ESXi hypervisor is supported but it is not required. If Network Orchestrator is used in an environment with ESXi, the commands below are used, whether or not Nutanix is included in the configuration. The functionality provided by Network Orchestrator is the same whether or not Nutanix is used for hyperconverged storage when the hypervisor in use is ESXi.

In most cases, the commands use syntax of the form:

```
show vmware <command> <operands>
```

For example, the **`show virtual-networks`** command which would be used with Acropolis is **`show vmware virtual-networks`** when ESXi is used. However, the **`show vdm`** command, which shows which management environment is connected to Orchestrator, works for either hypervisor.

A sample switch configuration for ESXi is shown in Figure 18. Note that the subscription URLs are not used with ESXi, The specification of the interfaces that are attached to the Nutanix appliances is done with the **`add interface`** command just as under Acropolis.

```
Switch-1#sho run vdm
!
vdm vmware
ip address 10.240.149.250 vrf management
username administrator@vsphere.local password encrypted 0xed3f47af2ef08ce1
add interface port-channel 2
add interface port-channel 3
add interface port-channel 4
add interface port-channel 5
vnic-stats enable
vm uuid 420e27e3-1df4-5341-58a7-663bf5ba333f vnic vlan 2 attach queueing-policy pmap-que
!
Switch-1#show run vlag
!
vlag tier-id 10
vlag isl port-channel 100
vlag hlthchk peer-ip 10.240.149.247 vrf management
vlag enable
vlag instance 2 port-channel 2
vlag instance 2 enable
vlag instance 3 port-channel 3
vlag instance 3 enable
vlag instance 4 port-channel 4
vlag instance 4 enable
vlag instance 5 port-channel 5
vlag instance 5 enable
!
... port-channel interface details ...

interface port-channel2
switchport mode trunk
auto-policy enable
auto-policy host-discovery
!
interface port-channel3
switchport mode trunk
auto-policy enable
auto-policy host-discovery
!
interface port-channel4
switchport mode trunk
auto-policy enable
auto-policy host-discovery
!
interface port-channel5
switchport mode trunk
auto-policy enable
auto-policy host-discovery
```

*Figure 18 ESXi VDM configuration*

# Access Control and QoS for Guest VMs

Switches attached to servers in a Nutanix cluster can also be configured with policies which are automatically applied to specifically identified guest VMs when these VMs are activated.

These policies can include the following:

- Access Control filters, which can limit access to a guest VM by IP address or MAC address
- Quality of Service (QoS) filters, which can prioritize certain types of traffic coming in to a guest VM
- Queueing filters, which can choose which hardware queue will be used for traffic

Guest VMs are identified by their name or UUID, and the name or address of the virtual NIC used to attach the guest to a virtual network. No policy will be successfully attached to a VM unless it is explicitly configured on all of the switches (typically two) to which the hosts are attached.

The base VDM configuration for Nutanix is required for the commands shown below to function successfully. If a guest VM is moved from one host to another and both of the hosts are part of the same cluster whose PRISM management is configured on the switch, then the policies configured will be moved by the switch to follow the guest VM to its new host.

The command syntax for this is as shown below. All of the commands shown require that the appropriate policy – ACL, QoS policy map, or Queueing policy map – exist and be configured properly. These policy maps would typically be formulated by the networking team and would be part of a set of policies that were in use across the entire network.

The commands for configuring these policies are shown below and in the current Application Guide and Command Reference documentation.

For an access control (security) filter, using the following commands:

```
(config)# vdm Nutanix|VMware
(config-vdm)#vm {<vm-name> | <vm-uuid>} vnic {vlan <vid> | ip <ip> | mac <mac>} attach
security-policy <ip-acl-name> in
(config-vdm)#vm {<vm-name> | <vm-uuid>} vnic {vlan <vid> | ip <ip> | mac <mac>} attach
security-policy <mac-acl-name> in
```

To see the VM and its security ACLs using the following command:

```
show [vmware] virtual-machine security-policy [vm {<vm-name> | <vm-uuid>}]
```

For a QoS (prioritization) filter:

```
(config)# vdm Nutanix|VMware
(config-vdm)# vm {<vm-name> | <vm-uuid>} vnic {vlan <vid> | ip <ip> | mac <mac>} attach
qos-policy <qos-policy-map-name>
```

Status of the filter can be shown with this command:

```
show [vmware] virtual-machine qos-policy [vm {<vm-name> | <vm-uuid>}]
```

For a queueing filter:

```
(config)# vdm Nutanix|VMware
(config-vdm)# vm {<vm-name> | <vm-uuid>} vnic {vlan <vid> | ip <ip> | mac <mac>} attach
queueing-policy <queueing-policy-map-name>
```

The status of the filter can be checked with the following command:

```
show [vmware] virtual-machine queueing-policy [vm {<vm-name> | <vm-uuid>}]
```

To see the status of security filters:

```
show access-lists summary
```

To see the status of QoS and/or queueing filters:

**show policy-map interface brief**

A sample access-list configuration and related displays are shown in Figure 19:

```
Switch-2#show virtual-machine security-policy info

VM Name : VM2
   VNIC Qualifier   : VLAN 10
   Security Policy  : abc

Switch-2#show run vdm | include "VM2"
 vm name "VM2" vnic vlan 10 attach security-policy abc in

Switch-2#show access-list abc
IP access list  abc
        10 permit any host 1.1.1.1 any
        20 permit any any any
```

*Figure 19 VM access list displays*

# Virtual NIC statistics collection

Statistics on the activity of virtual NICs – those connected from a guest VM to a vNetwork – can be collected periodically by a switch. These statistics include byte counts and other similar measures. This collection can be activated and deactivated by the commands below:

| | |
|---|---|
| **[no]vnic-stats enable** | Turns the collection process on or off |
| **vnic-stats interval <30-3000>** | Controls how frequently the statistics are refreshed – in seconds |

An example of the statistics collected is shown in Figure 20 below. Note that this is for an idle VM.

```
Switch-2#show virtual-machine vnic statistics interface port-aggregation 2

Interface po2

VM uuid  : 17330c49-8150-4b43-a3a8-73109559b20f
VM name  : VM1_1
VNIC uuid: 0f60f792-0e79-4b39-8e21-5a73693fa39f

  Time of capture   :2017/12/28 17:19:01
  RX Bytes          :0                  TX Bytes          :0
  RX Packets        :0                  TX Packets        :0
  RX Mcast Packets  :-1                 TX Mcast Packets  :-1
  RX Bcast Packets  :-1                 TX Bcast Packets  :-1
  RX Dropped Packets:0                  TX Dropped Packets:0
  RX Error Packets  :0                  TX Error Packets  :0
  RX Rate Kbps      :0                  TX Rate Kbps      :0
  Usage Rate Kbps   : -1
```

*Figure 20 VM vnic statistics*

# Layer 2 and Layer 3 upstream connectivity

A cluster can be connected to a customer's existing upstream network in two ways:

- Layer 2 connectivity extends the VLANs that are used by the hosts and guests within the cluster into switches within the existing network. These VLANs will flow over the uplinks between the top-of-rack switches in the cluster and the adjacent upstream switches. Routing between a subnet that originate within the cluster and one that is outside it will take place on upstream switches/routers.
- Layer 3 connectivity routes traffic originating on the subnets and VLANs within the cluster to subnets within the core network. Uplinks between the top-of-rack switches and adjacent upstream switches will either be a subnet of their own or part of a core network subnet.

Figure 21 on page 21 shows the Layer 2 connectivity option.

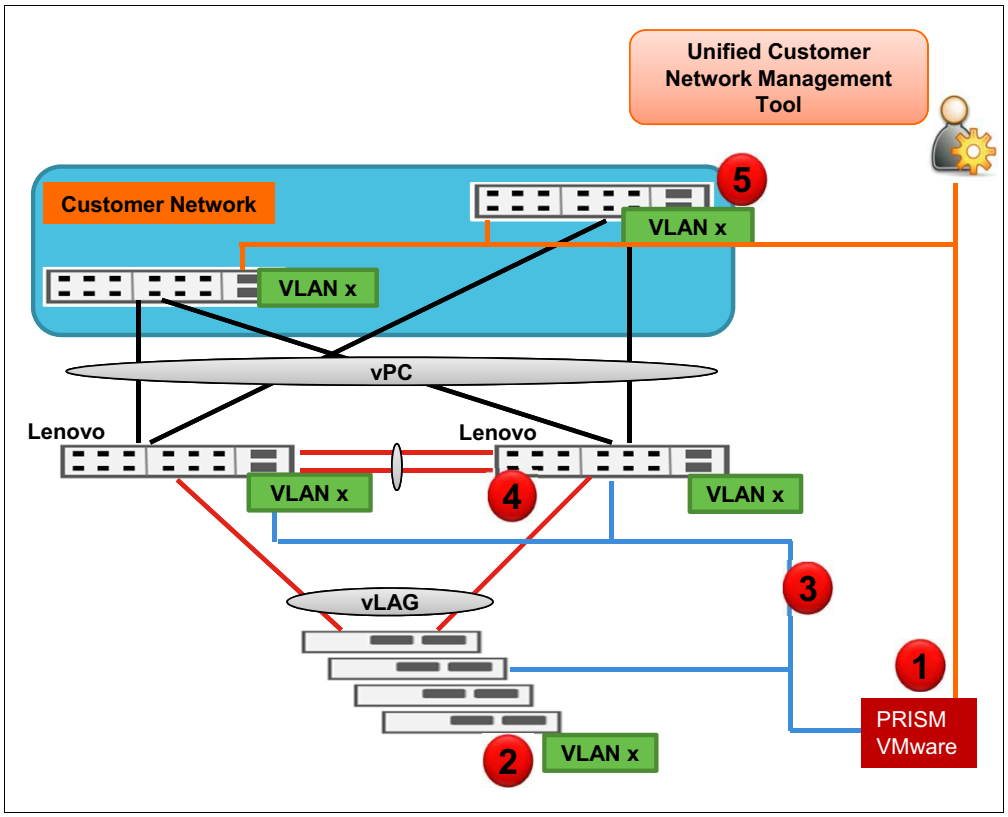


*Figure 21 Layer 2 Uplink - CNOS+PRISM integration*

**Note**: Newly created VLANs are automatically created as a result of a message from the cluster to the Lenovo cluster switches, but not on the upstream switches. They must be created on the upstream switches either manually or by use of another network management tool. This is shown as step 5 on page 7.

Figure 22 on page 22 shows the Layer 3 connectivity option.

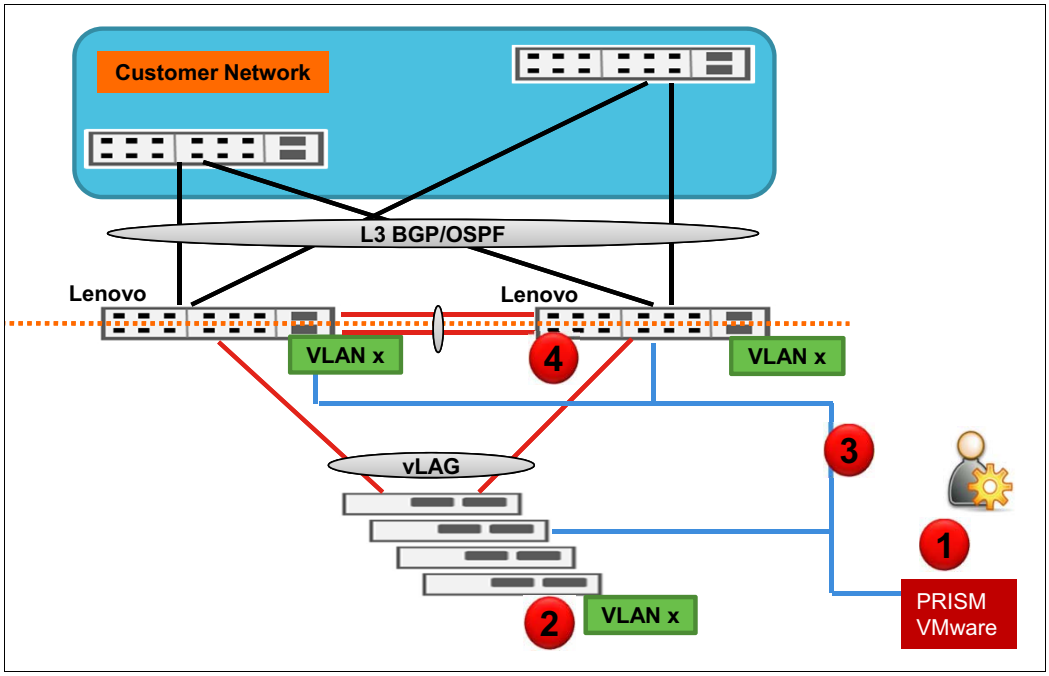


*Figure 22 Express cluster Layer 3 Uplink - CNOS+PRISM integration*

**Note**: A dynamic routing protocol is used so that traffic can be routed between networks in the cluster and those in the customer's existing network.

# Appendix: Complete switch configuration

Example 1 lists the complete configuration of the Lenovo switch. Some of the configuration text in red text is different depending on whether Acropolis or ESXi is used.

Specifically:

- The subscription URLs at the very end of the configuration are only applicable when Acropolis is used as the hypervisor. The configuration option for the port-channel interfaces for **auto-policy host-discovery** is for ESXi systems. These differences apply even when Nutanix is used with ESXi as the hypervisor.
- The port-channels for use with ESXi are static rather than using LACP. This is because our support does not currently include the VMware Distributed Switch, which is required for support of LACP. This support will be added in an upcoming release. The port channels for ESXi therefore have mode **on** rather than **active**.

*Example 1 Switch configuration*

```
hostname Switch-2
!
!
clock timezone PST -8 0
!
vrf context management
ip route 0.0.0.0/0 10.240.149.254
!
microburst-detection interval 100
!
no ip icmp-broadcast control-plane
service-policy input copp-system-policy feature telnet
feature ssh
no feature tacacs+
snmp-server enable snmp snmp-server version v1v2v3
snmp-server view all .1 included
snmp-server community public group network-admin
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
ntp server 10.240.149.253 minpoll 4 maxpoll 6
ntp server 10.240.148.20 minpoll 4 maxpoll 6 username admin role network-admin password
encrypted
$6$bJoWqYo.$QHfpuxc/15jr7SMnhMdHrEftvcIsv5Q9h/5h7aqVe8LU8izyfQ7KtHrlliQjcc7O8fx3ZjkihQ.s6u7
rty9qA1
feature restApi http
vlag tier-id 100
vlag isl port-channel 100
vlag hlthchk peer-ip 10.240.149.247 vrf management vlag startup-delay 0
no vlag vrrp active
vlag peer-gateway
vlag enable
vlag instance 1 port-channel 1
vlag instance 1 enable
vlag instance 2 port-channel 2
vlag instance 2 enable
vlag instance 3 port-channel 3
vlag instance 3 enable
vlag instance 4 port-channel 4
vlag instance 4 enable
vlag instance 5 port-channel 5
vlag instance 5 enable
```

```
vlag instance 6 port-channel 6
vlag instance 6 enable
vlag instance 7 port-channel 7
vlag instance 7 enable
vlag instance 8 port-channel 8
vlag instance 8 enable
vlag instance 10 port-channel 101
vlag instance 10 enable
vlag instance 15 port-channel 15
vlag instance 15 enable
vlag instance 20 port-channel 202
vlag instance 20 enable
spanning-tree mode disable
!! default class-map configuration redacted
!!
vlan 1
!
vlan 100
!
!
interface Ethernet1/1
switchport mode trunk
switchport trunk native vlan 100
channel-group 1 mode active|on
!
interface Ethernet1/2
switchport mode trunk
switchport trunk native vlan 100
channel-group 2 mode active|on
!
interface Ethernet1/3
switchport mode trunk
switchport trunk native vlan 100
channel-group 3 mode active|on
!
interface Ethernet1/4
switchport mode trunk
switchport trunk native vlan 100
channel-group 4 mode active|on
!
interface Ethernet1/5
switchport mode trunk
switchport trunk native vlan 100
channel-group 5 mode active|on
!
interface Ethernet1/6
switchport mode trunk
switchport trunk native vlan 100
channel-group 6 mode active|on
!
interface Ethernet1/7
switchport mode trunk
switchport trunk native vlan 100
channel-group 7 mode active|on
!
interface Ethernet1/8
switchport mode trunk
switchport trunk native vlan 100
channel-group 8 mode active|on
!
```

```
... unused interfaces omitted
!
interface Ethernet1/50
switchport mode trunk
switchport trunk native vlan 100
channel-group 202 mode active
!
interface Ethernet1/51
switchport mode trunk
switchport trunk native vlan 100
channel-group 202 mode active
!
interface Ethernet1/52
!
interface Ethernet1/53
switchport mode trunk
switchport trunk native vlan 100
channel-group 100 mode active
!
interface Ethernet1/54 switchport mode trunk
switchport trunk native vlan 100
channel-group 100 mode active
!
interface loopback0
no switchport
!
interface mgmt0
no switchport
vrf member management
no ip address dhcp
ip address 10.240.149.246/24
!
interface Vlan1
no switchport
!
interface port-channel1
switchport mode trunk
switchport trunk native vlan 100
auto-policy enable
auto-policy host-discovery
!
interface port-channel2 switchport mode trunk
switchport trunk native vlan 100
auto-policy enable
auto-policy host-discovery
!
interface port-channel3
switchport mode trunk
switchport trunk native vlan 100
auto-policy enable
auto-policy host-discovery
!
interface port-channel4
switchport mode trunk
switchport trunk native vlan 100
auto-policy enable
auto-policy host-discovery
!
interface port-channel5
switchport mode trunk
```

```
switchport trunk native vlan 100
auto-policy enable
auto-policy host-discovery
!
interface port-channel6
switchport mode trunk
switchport trunk native vlan 100
auto-policy enable
auto-policy host-discovery
!
interface port-channel7
switchport mode trunk
switchport trunk native vlan 100
auto-policy enable
auto-policy host-discovery
!
interface port-channel8
switchport mode trunk
switchport trunk native vlan 100
auto-policy enable
!
interface port-channel100
lacp suspend-individual
switchport mode trunk
switchport trunk native vlan 100
!
interface port-channel202
switchport mode trunk
switchport trunk native vlan 100
!
line con 0
exec-timeout 0 0
line vty 0 39
!
vdm nutanix|vmware
ip address 10.240.149.10 vrf management
username admin password encrypted 0xb5e222b27807402e
!
! use appropriate credentials for the portal which is being used
!
vnic-stats enable
vnic-stats interval 60
add interface po 1
add interface po 2
add interface po 3
add interface po 4
add interface po 5
add interface po 6
add interface po 7
add interface po 8
refresh-vms-url https://10.240.149.10:9440/api/nutanix/v3/vms/list
refresh-vnet-url https://10.240.149.10:9440/api/nutanix/v3/networks/list
topo-discovery-url https://10.240.149.10:9440/PrismGateway/services/rest/v1/vms
subscribe server-url https://10.240.149.10:9440/api/nutanix/v3/webhooks
client-url http://10.240.149.246:8090/nos/api/cfg/nutanix/network vnetwork
subscribe server-url https://10.240.149.10:9440/api/nutanix/v3/webhooks
client-url http://10.240.149.246:8090/nos/api/cfg/nutanix/vm vm
!
end
```

# Change history

December 17, 2018

- Updated Table 2 on page 8

May 15, 2018:

- ThinkAgile Network Orchestrator now supports VMware vSphere and the ESXi hypervisor
- New section “Implementation using VMware ESXi” on page 13

January 31, 2018:

- Changes to show added function for security, queueing, and QoS filers
- Changes to show added function for vnic statistics
- Added Nutanix support matrix, Table 2 on page 8

September 25, 2017:

- Nutanix AOS 5.1 is now supported, however AOS 5.12 is not supported.
- Support for AOS 5.1.3 is planned.
- IP address correction in various lab examples
- Correction to the lab topology diagram, Figure 8 on page 10

September 8, 2017:

- Grammar and style corrections

May 26, 2017:

- At present, only the Acropolis hypervisor is supported; VMware and Microsoft Hyper-V are not currently supported.

# Authors

This paper was produced by the following team of specialists:

**Scott Lorditch** is a Consulting System Engineer for Lenovo. He performs network architecture assessments and develops designs and proposals for solutions that involve Lenovo Networking products. He also developed several training and lab sessions for technical and sales personnel. Scott joined IBM as part of the acquisition of Blade Network Technologies and joined Lenovo as part of the System x acquisition from IBM. Scott spent almost 20 years working on networking in various industries, as a senior network architect, a product manager for managed hosting services, and manager of electronic securities transfer projects. Scott holds a BS degree in Operations Research with a specialization in computer science from Cornell University.

**Dillibabu Kodamala** is a Software Developer for Lenovo. He has more than 15 years of experience in design, development, and implementation of application/system software in Layer 2-7 switching and data center switching. Dilli joined IBM for Openflow and SDN software development, then joined Lenovo as part of the System x acquisition from IBM. Prior to IBM he is involved in software development and sustenance activities for Nortel/Radware Alteon Switches.

**Aksh Mittal** is an Advisory Engineer for Lenovo. He performs Proof of Concept (PoC) testing and executes conceptual design tests and ensure that designs match with the requirements and expectations. He also performs sustenance and feature testing on various L2 and L3 protocols for Customer field defects on Lenovo Networking products. Aksh joined IBM as a

Technical Co-op after completing his MS degree in Electrical Engineering from California State University, Fullerton and joined Lenovo as part of  the System x acquisition from IBM.

**Chidambaram Bhagavathiperumal** is a Senior Software Engineer for Lenovo. He performs network architecture design and implementation for hyperconverged infrastructure, network telemetry and software-defined networking (SDN) solutions that involve Lenovo products. Chidambaram joined IBM as part of the acquisition of Blade Network Technologies and joined Lenovo as part of the System x acquisition from IBM. Chidambaram spent almost 14 years working on networking products involving L2 and L3 protocols across various hardware platforms. Chidambaram holds a holds a BS degree in Computer Science and Engineering from MS university India.

# Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on December 17, 2018.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp0604

# Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

| | | |
|---|---|---|
| Blade Network Technologies® | Lenovo(logo)® | vNIC™ |
| Lenovo® | System x® | |
| RackSwitch™ | ThinkAgile™ | |

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.