Cumulus-Nutanix Solution Installation Guide

CONTENTS

THE CUMULUS-NUTANIX PARTNERSHIP

Nutanix provides an industry-leading hyperconverged infrastructure (HCI) solution combining storage and virtualization into a single, disaggregated platform. By leveraging disaggregated data center server hardware, Nutanix collapses multiple independent services into a single converged platform, optimizing both VM and storage infrastructures.

Cumulus Networks, with their flagship Cumulus Linux, has been a leader in disaggregated data center networking supporting traditional layer 2 and layer 3 features as well as advanced overlay and automation technologies such as VxLAN and Ansible. Combined with Nutanix technology, Cumulus Linux and other Cumulus products deliver an unparalleled HCI experience featuring lower operating costs, improved performance, greater scalability, streamlined management and simplified procurement.

CHALLENGES IN HCI DEPLOYMENTS

IT groups in a wide range of industries are deploying Nutanix solutions to take advantage of the benefits of HCI. Unfortunately, many are falling behind schedule, hampered by connectivity and configuration issues as well as organizational challenges.

One aspect of the problem lies in the sheer number of steps that must be executed perfectly to establish network connectivity, ensure redundancy and manage changes in the HCI environment. For starters, newly installed top-of-rack (TOR) switches require the latest Cumulus Linux release, along with license credential and a baseline configuration. Many deployments require redundant uplinks based on the multi-chassis link aggregation protocol (MLAG), which further complicates configuration.

Once the system is properly installed and configured, there are ongoing operational challenges. When a new virtual machine (VM) is deployed, Nutanix automatically assigns a VLAN within the hypervisor, AHV, to allow the VM to send traffic to other network components. However, the VLAN must also be provisioned on the TOR switches, bringing the networking team into the process. Sometimes the solution is for network engineers to provision VLANs as needed, which places a burden on busy networking teams. Another approach is to configure all VLANs on all server-facing switch ports, which creates spanning-tree challenges and impairs switch scaling.

When a VM is removed or moved, Nutanix updates the associated VLAN within the appropriate AHV host. However, the change must also be made on the TOR switch—another manual task for the network team. In addition, network operators have historically had no way to monitor the creation and movement of VMs. This lack of visibility makes it difficult to correlate network issues to changes in the Nutanix environment. These challenges are universal to HCI deployments but the Cumulus-Nutanix integration eases these operational burdens.

THE CUMULUS-NUTANIX SOLUTION

Cumulus engineers worked closely with their Nutanix counterparts to create a solution to these problems. The heart of this solution is the **Cumulus Hyperconverged Service (HCS)**, a new feature in Cumulus Linux version 3.7.3. Cumulus HCS is tightly integrated with the Nutanix AHV hypervisor. Cumulus HCS discovers Nutanix nodes, provisions network resources, and programs new VLANs when notified by Nutanix Acropolis Service (AOS)—all without human intervention.

Another pillar of the solution is **Cumulus Zero Touch Provisioning**—also known as Cumulus on a stick—which allows you to automatically install Cumulus Linux and configure Cumulus switches for bare-metal deployments. With Cumulus, there is no need for additional equipment or servers to manage or provision the network.

Two other Cumulus products enhance Nutanix deployments. **Cumulus RMP** brings all the flexibility and efficiency of Cumulus Linux on an out-of-band management switch such as Cumulus Express. **Cumulus NetQ**, a fabric telemetry solution, provides unprecedented real-time network visibility for dynamic environments as well as historical analysis.

EXAMPLE NUTANIX DEPLOYMENT

For instructional purposes, this guide assumes a typical Nutanix deployment with the following network characteristics (see Figure 1):

- Leaf switches with 48 10Gb or 25Gb ports
- Four or more 40Gb or 100Gb uplinks
- Nutanix servers attached to the leaf switches on 10Gb or 25Gb ports
- MLAG peer link on the first two uplink ports, swp49 and swp50
- Other infrastructure connections on ports swp51 and above
- eth0 management interface configured for management VRF via DHCP
- Gateway IP addresses for all VMs, including the CVM, do not exist on the Cumulus Linux switches

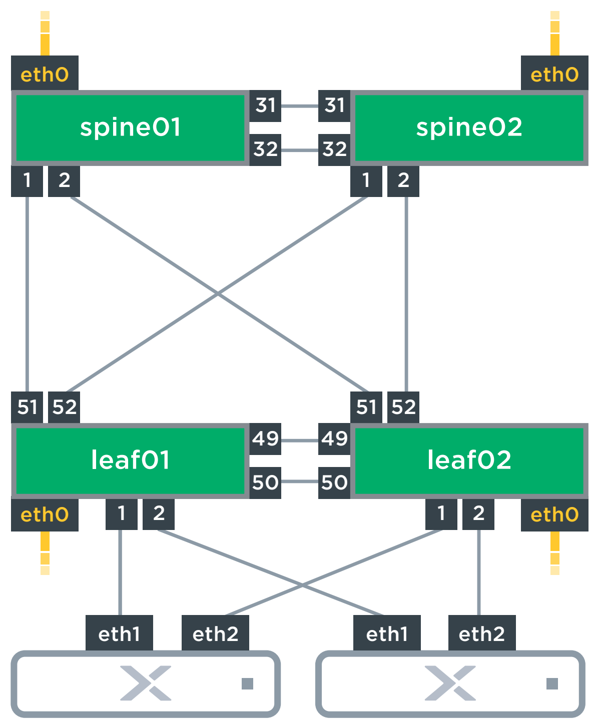

FIGURE 1. EXAMPLE NUTANIX DEPLOYMENT

The following chapters offer step-by-step installation and configuration instructions for the Cumulus-Nutanix joint solution. Because Cumulus Linux runs on more than 90 hardware platforms, you may encounter minor differences in switch port layouts and configurations from the examples in this guide. Note that these procedures only deal with leaf switches—spine switch configuration is beyond the scope of this document.

> Have design questions? The Cumulus Community Slack channel is a great place to ask Cumulus customers and employees for expert opinions.

CONFIGURING BARE-METAL SWITCHES (ZERO TOUCH)

The following procedure explains how to install Cumulus Linux and fully configure Cumulus HCS in new deployments of bare-metal switches. If Cumulus Linux is already installed, see the following chapter.

Requirements

- Cumulus on a Stick disk image
- USB stick

Step 1. Copy the Cumulus disk image onto the USB stick.

Step 2. Set your Nutanix username and password and server IP address by editing the configuration file on the USB stick.

```
### /etc/default/cumulus-hyperconverged config file
# username for Prism (required)
USERNAME=admin
# password for Prism (required)
PASSWORD=1CumulusLinux!
# CVM address used by the service (required)
SERVER=10.1.1.11
# Hook server address (optional)
#HOOK_SERVER=10.0.0.0
# Hook port (optional)
#HOOK_PORT=9440
# Socket timeout (optional)
#SOCKET_TIMEOUT=10.0.0.0
# single/multi rack configuration (optional)
#VXLAN_CONFIG=False
# loglevel: verbose/debug (optional)
LOGLEVEL=debug
# periodic sync timeout (optional)
#PERIODIC_SYNC_TIMEOUT=60
```

Step 3. **Place the USB stick into the switch and cycle the switch power.**

The Cumulus Linux image will be automatically installed, along with your Cumulus Linux license and a baseline configuration. The switch will reboot multiple times during this process. Depending on your specific hardware platform, this process can take up to 20 minutes.

After the installation is complete, you will see the LEDs corresponding to the ports connected to Nutanix nodes green.

Step 4. **When the installation is complete, remove the USB stick.**

CONFIGURING EXISTING CUMULUS SWITCHES

This procedure is for deployments in which Cumulus Linux is already installed on the switches. For bare metal installations, see the previous chapter.

Requirements

- Cumulus Linux-supported switches running Cumulus Linux 3.7.3 or later
- Nutanix AOS 5.8 or later
- LLDP enabled on Nutanix (the default value in AOS 5.8 and later)
- LACP enabled on Nutanix (the default value in AOS 5.8 and later)
- IP connectivity between the switches and Nutanix controller VMs (CVMs)

Step 1. **Enable Cumulus MLAG.**

```
cumulus@leaf01:~$ net add interface swp49,swp50 mtu 9216
cumulus@leaf01:~$ net add clag peer sys-mac 44:38:39:FF:40:00 interface swp49,swp50
primary
cumulus@leaf01:~$ net commit

cumulus@leaf02:~$ net add interface swp49,swp50 mtu 9216
cumulus@leaf02:~$ net add clag peer sys-mac 44:38:39:FF:40:00 interface swp49,swp50
secondary
cumulus@leaf02:~$ net commit
```

Note: sys-mac is a MAC address from the reserved Cumulus Networks MAC address space and must be the same on both MLAG peers as shown in the example above. If more than one MLAG pair is being deployed, sys-mac must be unique across MLAG sets.

Step 2. Configure the Default Layer 2 Bridge.

```
cumulus@leaf01:~$ net add bridge bridge ports peerlink
cumulus@leaf01:~$ net add bridge bridge pvid 1
cumulus@leaf01:~$ net commit

cumulus@leaf02:~$ net add bridge bridge ports peerlink
cumulus@leaf02:~$ net add bridge bridge pvid 1
cumulus@leaf02:~$ net commit
```

Note: pvid is the native VLAN address. If unknown, set the value to 1 as shown above.

Step 3. Set Nutanix username and password and server IP address.

```
cumulus@leaf02:~$ cat /etc/default/cumulus-hyperconverged
### /etc/default/cumulus-hyperconverged config file
# username for Prism (required)
USERNAME=admin
# password for Prism (required)
PASSWORD=1CumulusLinux!
# CVM address used by the service (required)
SERVER=10.1.1.11
# Hook server address (optional)
#HOOK_SERVER=10.0.0.0
# Hook port (optional)
#HOOK_PORT=9440
# Socket timeout (optional)
#SOCKET_TIMEOUT=10.0.0.0
# single/multi rack configuration (optional)
VXLAN_CONFIG=False
# loglevel: verbose/debug (optional)
LOGLEVEL=verbose
# periodic sync timeout (optional)
#PERIODIC_SYNC_TIMEOUT=60
```

Note: The server IP address may be a specific Nutanix CVM address or may be the cluster IP.

Step 4. Enable and start Cumulus HCS.

```
cumulus@leaf01:~$ sudo systemctl enable cumulus-hyperconverged.service
cumulus@leaf01:~$ sudo systemctl start cumulus-hyperconverged.service

cumulus@leaf02:~$ sudo systemctl enable cumulus-hyperconverged.service
cumulus@leaf02:~$ sudo systemctl start cumulus-hyperconverged.service
```

Step 5. Verify that the service is running.

```
cumulus@leaf02:~$ sudo systemctl status cumulus-hyperconverged.service
cumulus-hyperconverged.service - Cumulus Linux Hyperconverged Daemon
Loaded: loaded (/lib/systemd/system/cumulus-hyperconverged.service; enabled)
Active: active (running) since Mon 2019-01-07 03:36:26 UTC; 56min ago
Main PID: 4207 (cumulus-hyperco)
CGroup: /system.slice/cumulus-hyperconverged.service
    ├─4207 /usr/bin/python /usr/bin/cumulus-hyperconverged.service
    └─4300 /usr/sbin/lldpcli -f json watch
```

If the service fails to start more information may be available in the service logs.
View the log with `sudo journalctl -u cumulus-hyperconverged`

Step 6 (optional). Enable server-facing ports and configure jumbo MTU. LLDP is automatically enabled.

```
cumulus@leaf02:~$ net add interface swp1-48 mtu 9216
cumulus@leaf02:~$ net commit
```

At this point, the service is fully configured. It may take up to 60 seconds for LLDP frames to be received to trigger Cumulus HCS.

VERIFICATION PROCEDURES

Use the following procedures to verify proper operation of Cumulus HCS and provide necessary troubleshooting information to Cumulus support.

Dynamic Bond Creation

Check to see that bonds are being dynamically created as shown below. This example shows how the interfaces `bond_swp1`, `bond_swp2` and `bond_swp3` are automatically created. If Cumulus HCS is operating correctly, you should see a similar table.

```
cumulus@leaf01:~$ net show interface bonds
    Name       Speed   MTU  Mode     Summary
--  ---------  -----  ----  -------  ---------------------------------
UP  bond_swp1  1G     1500  802.3ad  Bond Members: swp1(UP)
UP  bond_swp2  1G     1500  802.3ad  Bond Members: swp2(UP)
UP  bond_swp3  1G     1500  802.3ad  Bond Members: swp3(UP)
UP  peerlink   2G     1500  802.3ad  Bond Members: swp49(UP), swp50(UP)
```

LLDP Messages Received

If bonds are not being created, the LLDP messages may not be getting through. You can check this possibility using *net show lldp*:

```
cumulus@leaf01:~$ net show lldp
LocalPort  Speed  Mode           RemoteHost       RemotePort
---------  -----  -------------  ---------------  ----------
swp1       1G     BondMember     NTNX-e08c61ec-A  ens3
swp2       1G     BondMember     NTNX-d618a06d-A  ens3
swp3       1G     BondMember     NTNX-4e6eac27-A  ens3
swp49      1G     BondMember     leaf02           swp49
swp50      1G     BondMember     leaf02           swp50
swp52      1G     NotConfigured  spine01          swp1
swp52      1G     NotConfigured  spine02          swp1
```

Cumulus HCS Configuration

Use this procedure to view the complete Cumulus HCS configuration.

```
cumulus@leaf01:~$ net show lldp swp1
-----------------------------------------------------------------------------
LLDP neighbors:
-----------------------------------------------------------------------------
Interface:    swp1, via: LLDP, RID: 37, Time: 0 day, 01:42:23
  Chassis:
    ChassisID:    mac 2c:c2:60:50:f6:8a
    SysName:      NTNX-e08c61ec-A
    SysDescr:     CentOS Linux 7 (Core) Linux 4.4.77-1.el7.nutanix.20180425.199.x86_64
#1 SMP Thu Apr 26 01:01:53 UTC 2018 x86_64
    MgmtIP:       10.1.1.10
    MgmtIP:       fe80::2ec2:60ff:fe50:f68a
    Capability:   Bridge, on
```

```
    Capability:   Router, off
    Capability:   Wlan, off
    Capability:   Station, off
  Port:
    PortID:       mac 2c:c2:60:50:f6:8a
    PortDescr:    ens3
    TTL:          120
    PMD autoneg:  supported: yes, enabled: yes
      Adv:          10Base-T, HD: yes, FD: yes
      Adv:          100Base-TX, HD: yes, FD: yes
      Adv:          1000Base-T, HD: no, FD: yes
      MAU oper type: 1000BaseTFD - Four-pair Category 5 UTP, full duplex mode
-----------------------------------------------------------------------------
```

CONFIGURING UPLINKS

The process of configuring uplinks depends on the installation procedure, Zero Touch or manual configuration.

Zero Touch Provisioning

When zero-touch provisioning method is used, you can define the VLANs assigned to VMs by editing the ZTP settings file. See Appendix B for a sample ZTP file.

Manual Provisioning

For manual installations on existing Cumulus switches, enable the uplinks and define the associated VLANs as shown below.

```
cumulus@leaf01:~$ net add interface swp51-52 mtu 9216
cumulus@leaf01:~$ net add interface swp51-52 bridge vids 1-2999,4000-4094
cumulus@leaf01:~$ net add interface swp51-52 bridge pvid 1
cumulus@leaf01:~$ net commit
```

In this example, all VLANs are allowed on the uplink ports. Configuring any set of VLANs is allowed. Be aware that VLANs 3000-3999 are reserved on Cumulus Linux. This example assumes the untagged or native VLAN is VLAN ID 1. Change the VLAN ID as needed.

ADDING LOCAL DEFAULT GATEWAYS

Regardless of the deployment method (Zero Touch or manual), you can add one or more local default gateways on the switch pair to provide a redundant solution as shown in this example:

```
Leaf01:
cumulus@leaf01:~$ net add vlan 1 ip address 10.1.1.11/24
cumulus@leaf01:~$ net add vlan 1 ip address-virtual 00:00:5e:00:01:01 10.1.1.1/24
cumulus@leaf01:~$ net commit
```

```
Leaf02:
cumulus@leaf02:~$ net add vlan 1 ip address 10.1.1.12/24
cumulus@leaf02:~$ net add vlan 1 ip address-virtual 00:00:5e:00:01:01 10.1.1.1/24
cumulus@leaf02:~$ net commit
```

The first configuration line defines the unique IP address assigned to each switch. This is required and must be unique. On leaf01, this is represented by the IP address `10.1.1.1/24`.

The second line defines the virtual IP address that will be used as the default gateway address for any hosts in this VLAN. On both leaf01 and leaf02 this is represented by the IP address `10.1.1.1/24`.

The `address-virtual` MAC address is assigned from a reserved pool of Cumulus Networks MAC addresses. The address must start with `00:00:05:00:01:` and end with any value between `00` and `ff`. Both switches leaf01 and leaf02 must have the same MAC address. Outside of this switch pair, this MAC address must be unique and only be assigned to a single switch pair in your network.

Cumulus Linux relies on the VRR feature to make TOR switches act as gateways for the dual-attached Nutanix servers.

CONFIGURING OOB MANAGEMENT SWITCHES

Cumulus supports out-of-band (OOB) management in two ways. The recommended approach is to deploy Cumulus RMP. You can also install Cumulus Linux on generic switches as described below.

Cumulus RMP

Cumulus RMP is a ready-to-deploy solution that enables out-of-band management for web-scale networks. With Cumulus RMP, you can directly manage and support Nutanix systems in the rack without relying on the rest of the network.

To install Cumulus RMP, connect the Nutanix 1Gb IPMI, 1Gb Shared IPMI and 1Gb ports and the TOR eth0 ports to the Cumulus OOB switch. No additional configuration is required.

Cumulus RMP does not support MLAG or active/active connections across RMP switches. Connections across more than one RMP switch rely on traditional Spanning Tree Protocol for redundancy.

Generic Switch

For non-RMP 1Gb switches that support Cumulus Linux, the switch will need the Cumulus Linux software and license to be manually installed, along with the baseline configuration. The Cumulus on a Stick image has this information. Follow the standard Cumulus on a Stick instructions to apply a license and Cumulus Linux image.

Once the software is installed the following configuration can be used to configure all ports for a single, untagged, management VLAN, including any uplinks.

```
cumulus@oob-switch:~$ net add interface swp1-52 bridge access 1
```

A management IP address can also be assigned to this same untagged bridge interface. Use an appropriate IP address for your infrastructure.

```
cumulus@oob-switch:~$ net add vlan 1 ip address 192.0.2.1/24
```

Finally apply the configuration

```
cumulus@oob-switch:~$ net commit
```

These examples assume that the untagged VLAN ID is 1. To use a different VLAN, simply change the value in both configuration examples.

APPENDIX A. TECHNICAL DESCRIPTIONS

This appendix contains technical information on selected topics to help advanced users have a better understanding of how the Cumulus-Nutanix solution works.

Cumulus HCS Functionality

Running in the background of Cumulus Linux, Cumulus Hyperconverged Service has two major components:

Nutanix LLDP Switch Agent. When enabled, the agent listens for directly connected Nutanix servers via LLDP and enables MLAG bonding on the relevant ports. The Nutanix default of active-backup LACP load balancing is supported as well as optional active-active LACP load balancing.

Nutanix Webhook VLAN Provisioner. Cumulus switches register with the Nutanix CVM and wait to receive Nutanix webhooks. When a new VM is deployed on a server in the cluster, the CVM sends a message to the Cumulus switch with the physical server name and relevant VLANs. The Cumulus switch then dynamically provisions the configuration on the ports of the specific physical server.

Dynamic VLAN Provisioning

Cumulus HCS periodically polls Nutanix Prism for information about VMs in the cluster. When a new VM is discovered, Cumulus HCS automatically identifies the physical Nutanix server hosting the VM and discovers any VLANs that are required for a given VLAN. Cumulus HCS then automatically adds these VLANs to the default Cumulus VLAN-aware bridge, the MLAG peer link and the automatically created bond to the Nutanix node. When a VM is powered off, removed or moved and the associated VLAN has no other VMs, the VLAN is automatically removed from the bridge, peerlink, and dynamic bond.

Virtual Router Redundancy (VRR)

To provide redundant gateways, Cumulus Linux relies on the VRR feature. Virtual Router Redundancy (VRR) enables hosts to communicate with any redundant router without reconfiguration, running dynamic router protocols, or running router redundancy protocols. This means that redundant routers will respond to Address Resolution Protocol (ARP) requests from hosts. Routers are configured to respond in an identical manner, but if one fails, the other redundant routers will continue to respond, leaving the hosts with the impression that nothing has changed.

APPENDIX B: CONFIGURATION REFERENCE

This appendix gives the complete Cumulus configurations for the sample Nutanix deployment used in this guide.

Leaf01 Switch

/etc/network/interfaces

```
auto eth0
iface eth0 inet dhcp
    vrf mgmt

auto swp1
iface swp1
    alias Nutanix Server
    mtu 9216

auto swp2
iface swp2
    alias Nutanix Server
    mtu 9216

auto swp49
iface swp49
    alias MLAG Peerlink
    mtu 9216
```

```
auto swp50
iface swp50
    alias MLAG Peerlink
    mtu 9216
    bridge-pvid 1
    bridge-vids 1-4094

auto swp51
iface swp51
    alias Uplink
    bridge-pvid 1
    bridge-vids 1-4096

auto swp52
iface swp52
    alias Uplink
    bridge-pvid 1
    bridge-vids 1-4096

auto bridge
iface bridge
    bridge-ports peerlink
    bridge-vids 1-4096
    bridge-vlan-aware yes

auto peerlink
iface peerlink
    bond-slaves swp49 swp50

auto peerlink.4094
iface peerlink.4094
    clagd-peer-ip linklocal
    clagd-priority 1000
    clagd-sys-mac 44:38:39:FF:40:00

auto vlan1
iface vlan1
    address 10.1.1.200/24
    address-virtual 00:00:5e:00:01:01 10.1.1.1/24
    vlan-id 1
    vlan-raw-device bridge
```

```
/etc/default/cumulus-hyperconverged
### /etc/default/cumulus-hyperconverged config file
# username for Prism (required)
USERNAME=admin
# password for Prism (required)
PASSWORD=nutanix/4u
# CVM address used by the service (required)
SERVER=10.1.1.123
# Hook server address (optional)
#HOOK_SERVER=10.0.0.0
# Hook port (optional)
#HOOK_PORT=8888
# Socket timeout (optional)
#SOCKET_TIMEOUT=30
# single/multi rack configuration (optional)
VXLAN_CONFIG=False
# configure the vxlan local tunnel ip (optional)
#VXLAN LOCAL IP=10.0.0.0
# loglevel: verbose/debug (optional)
LOGLEVEL=verbose
# periodic sync timeout (optional)
#PERIODIC_SYNC_TIMEOUT=60
```

Leaf02 Switch

```
/etc/network/interfaces
auto eth0
iface eth0 inet dhcp
    vrf mgmt

auto swp1
iface swp1
    alias Nutanix Server
    mtu 9216

auto swp2
iface swp2
    alias Nutanix Server
    mtu 9216

auto swp49
iface swp49
    alias MLAG Peerlink
    mtu 9216
```

```
auto swp50
iface swp50
    alias MLAG Peerlink
    mtu 9216
    bridge-pvid 1
    bridge-vids 1-4096

auto swp51
iface swp51
    alias Uplink
    bridge-pvid 1
    bridge-vids 1-4096

auto swp52
iface swp52
    alias Uplink
    bridge-pvid 1
    bridge-vids 1-4096

auto bridge
iface bridge
    bridge-ports peerlink
    bridge-vids 1-4096
    bridge-vlan-aware yes

auto peerlink
iface peerlink
  bond-slaves swp49 swp50

auto peerlink.4094
iface peerlink.4094
    clagd-peer-ip linklocal
    clagd-priority 2000
    clagd-sys-mac 44:38:39:FF:40:00

auto vlan1
iface vlan1
    address 10.1.1.201/24
    address-virtual 00:00:5e:00:01:01 10.1.1.1/24
    vlan-id 1
    vlan-raw-device bridge
```

```
/etc/default/cumulus-hyperconverged
### /etc/default/cumulus-hyperconverged config file
# username for Prism (required)
USERNAME=admin
# password for Prism (required)
PASSWORD=Nutanix/4u
# CVM address used by the service (required)
SERVER=10.1.1.123
# Hook server address (optional)
#HOOK_SERVER=10.0.0.0
# Hook port (optional)
#HOOK_PORT=8888
# Socket timeout (optional)
#SOCKET_TIMEOUT=30
# single/multi rack configuration (optional)
VXLAN_CONFIG=False
# configure the vxlan local tunnel ip (optional)
#VXLAN LOCAL IP=10.0.0.0
# loglevel: verbose/debug (optional)
LOGLEVEL=verbose
# periodic sync timeout (optional)
#PERIODIC_SYNC_TIMEOUT=60
```

OOB Switches

```
/etc/network/interfaces
auto eth0
iface eth0

auto swp1
iface swp1

auto swp2
iface swp2

auto swp3
iface swp3

auto swp4
iface swp4

auto swp5
iface swp5
```

```
auto swp6
iface swp6

auto swp7
iface swp7

auto swp8
iface swp8

auto swp9
iface swp9

auto swp10
iface swp10

auto swp11
iface swp11

auto swp12
iface swp12

auto swp13
iface swp13

auto swp14
iface swp14

auto swp15
iface swp15

auto swp16
iface swp16

auto swp17
iface swp17

auto swp18
iface swp18

auto swp19
iface swp19

auto swp20
iface swp20
```

```
auto swp21
iface swp21

auto swp22
iface swp22

auto swp23
iface swp23

auto swp24
iface swp24

auto swp25
iface swp25

auto swp26
iface swp26

auto swp27
iface swp27

auto swp28
iface swp28

auto swp29
iface swp29

auto swp30
iface swp30

auto swp31
iface swp31

auto swp32
iface swp32

auto swp33
iface swp33

auto swp34
iface swp34

auto swp35
iface swp35
```

```
auto swp36
iface swp36

auto swp37
iface swp37

auto swp38
iface swp38

auto swp39
iface swp39

auto swp40
iface swp40

auto swp41
iface swp41

auto swp42
iface swp42

auto swp43
iface swp43

auto swp44
iface swp44

auto swp45
iface swp45

auto swp46
iface swp46

auto swp47
iface swp47

auto swp48
iface swp48

auto swp49
iface swp49

auto swp50
iface swp50
```

```
auto swp51
iface swp51

auto swp52
iface swp52

auto bridge
iface bridge
   bridge-ports swp1 swp2 swp3 swp4 swp5 swp6 swp7 swp8 swp9 swp10 swp11 swp12 swp13
swp14 swp15 swp16 swp17 swp18 swp19 swp20 swp21 swp22 swp23 swp24 swp25 swp26 swp27
swp28 swp29 swp30 swp31 swp32 swp33 swp34 swp35 swp36 swp37 swp38 swp39 swp40 swp41
swp42 swp43 swp44 swp45 swp46 swp47 swp48 swp49 swp50 swp51 swp52
   bridge-vlan-aware yes

auto vlan1
iface vlan1
   address 192.168.1.100/24
   vlan-id 1
   vlan-raw-device bridge
```

Zero-Touch Provisioning (ZTP) Script

The source code and details about the latest ZTP script can be found on the Cumulus Networks Github page, https://github.com/CumulusNetworks/nutanix-hcs.