



INSTALLATION RUNBOOK FOR Pluribus Networks + ML2 Plugin

Product Name: **[Pluribus Series Switches]**

Product Version: **[2.5.0]**

MOS Version: **[8.0]**

OpenStack Version: **[Liberty]**

Product Type: **[Network Switching]**

Contents

DOCUMENT HISTORY	2
1. INTRODUCTION	4
1.1 TARGET AUDIENCE	4
2. PRODUCT OVERVIEW	4
3. JOINT REFERENCE ARCHITECTURE.....	5
4. PHYSICAL & LOGICAL NETWORK TOPOLOGY	6
5. INSTALLATION AND CONFIGURATION	7
5.1 ENVIRONMENT PREPARATION.....	7
5.2 MOS INSTALLATION	7
5.3 PLURIBUS NETVISOR ML2 PLUGIN INSTALLATION STEPS	12
5.4 LIMITATIONS.....	14
5.5 TESTING.....	15
5.5.1 TEST CASE AND TEST RESULTS	

Document History

Version	Revision Date	Description
0.1	12-09-2016	Initial Version
0.2	15-09-2016	Rally Test, Comments, CLI Details and additionals updates
0.3	20-09-2016	Updated final comments DHCP comments and iPERF results

1. Introduction

This document is to serve as a detailed Deployment Guide for Netvisor ML2 Plugin. Pluribus Networks offers ML2 Plugin and Switching solutions. This document describes the reference architecture, installation steps for validated MOS+Pluribus Netvisor ML2 Plugin limitations and testing procedures.

1.1 Target Audience

Openstack Administrators, Network Administrator or Architects who are familiar with Openstack and Mirantis Openstack offerings and Pluribus Netvisor Solutions.

2. Product Overview

Pluribus Netvisor® is designed to run your Layer 2 and Layer 3 networking both more efficiently and at a lower cost than traditional switching infrastructure. It provides unparalleled agility, insight and security for your datacenter networking using a distributed controller fabric that eliminates the need for tap or broker infrastructure.

Feature Capabilities include:

Fabric Automation – Agility and simplification via an extensible, open, distributed controller fabric

Fabric Visibility – tap-free, flow-level monitoring and analysis with a network DVR “time machine”

Fabric Virtualization – Virtual network provisioning and Segmentation; support for multi-tenancy; vRouters

Extensibility – Programmability via Ansible, RESTful, Java and C APIs

L2/L3 – Compatible with all existing Layer-2/3 infrastructures CLI

Management – manage network switch hardware, resource pools and network configurations; high availability environment

Monitoring – including SNMP, sFlow, Traceroute and logging

Security – Flow-level (drop/log/mirror), MAC Security, Control Plane Traffic Protection, full L2/L3 ACL

QoS – CoS, DSCP, Strict Priority Queuing, Tos mapping

Audit – system, event and audit logging

Fabric Visibility Features

- **Telemetry** – inspects every individual TCP connection and client-server aggregated connection fabric statistics.

- **vFlow** - filtering fabric-wide data center switching traffic on a granular flow level and applying security/QoS actions or forwarding decisions on each defined flow.
- **vPort** - tracking endpoints/VMs on a global, fabric-wide endpoint table.

Pluribus Netvisor Open Stack ML2 Driver enables neutron to provision VLAN on Pluribus Netvisor OS Switches. This driver communicates to Netvisor OS devices using XML RPC as a transport protocol. This allows the Plugin to work with different Netvisor versions. Once a VLAN is provisioned on a specific TOR Netvisor automatically provisions the VLAN Across the Fabric and enables forwarding on all Switch to Switch Uplink Ports. This works in parallel with the vSwitch (OVS) driver plugin that resides on the compute hosts and provides integration between Computer and Networking

Although Pluribus Netvisor is able integrate with Neutron L3 capabilities, it is out of scope of this document.

3. Joint reference architecture

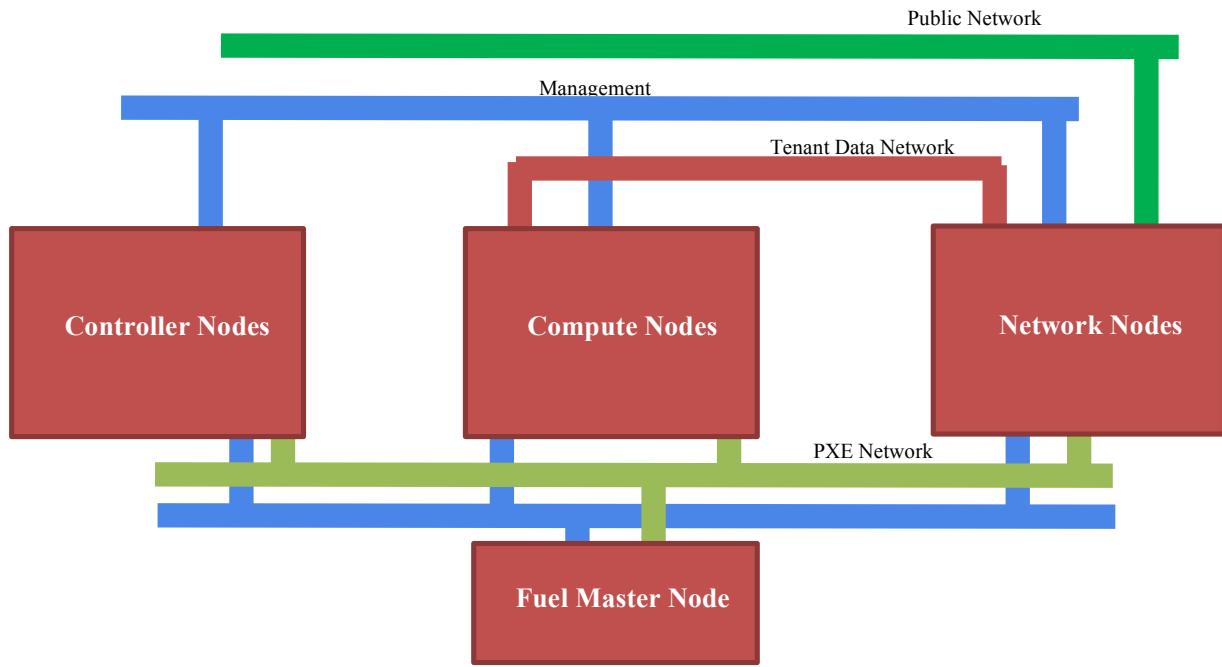
Pluribus Netvisor switches provide Layer 2 functionality + DHCP Services for the VMs for the scope of this document. Netvisors VNET feature creates a unique zone for each VLAN and thus completely segmenting each network and it's L3 Domain

vNET (A Zone is created per vNET for complete L2 Isolation, the Routing/L3 Instances for this L2 domain also resides in this zone)

It is recommended to use unique dedicated VLANs for all control plane communication (Management, PXE and Storage Network). All of these VLANs are present in a single OpenStack VNET.

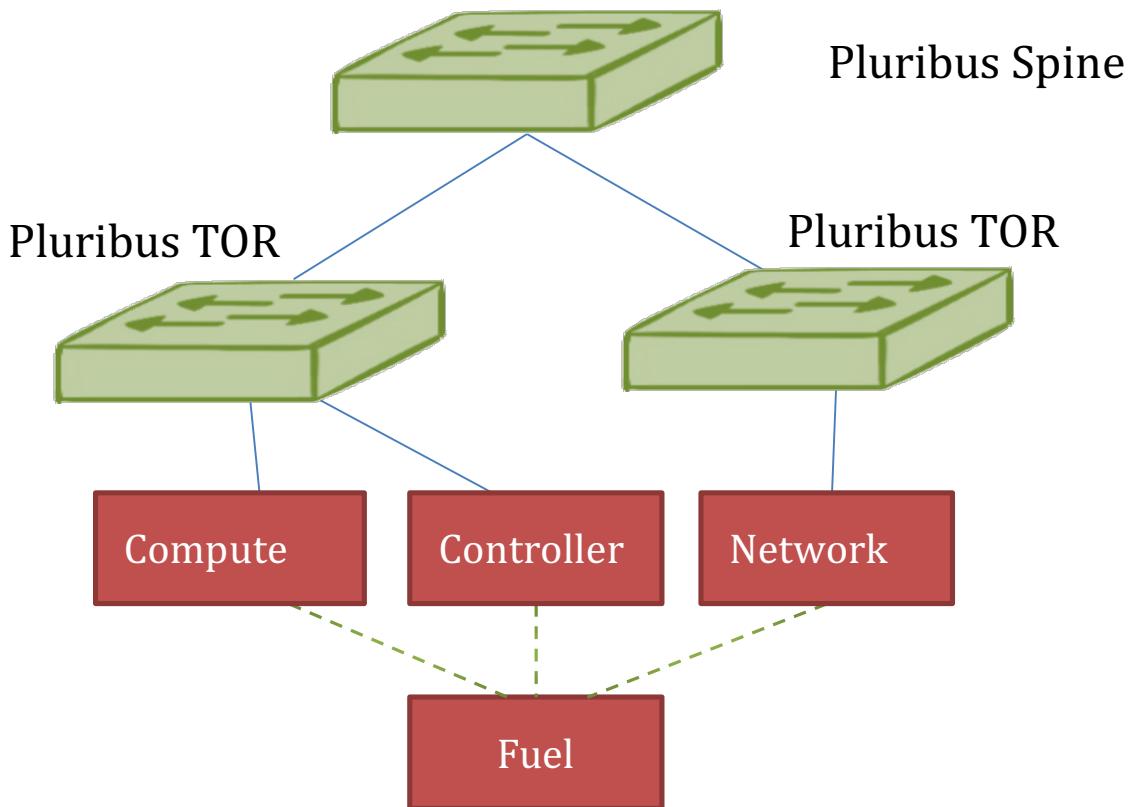
To isolate Tenant Data and Public Network, unique VLANs from the range of 1000-1030 were used in this document and each VLAN + vNET has its own unique VNET present which provides complete isolation of the tenant Network.

Logical Flow of the Network is described here, there are multiple networks for communication between the different services and VLANs are used to provision this



4. Physical & Logical network topology

This explains the physical setup



5. Installation and Configuration

5.1 Environment preparation

In our Installation we had 2 NICs on each node, NIC1 (eth0) was used for Public Network and NIC2 (en0s0p1) had VLAN Tagging enabled and had the Management, Storage, Tenant Data and PXE Networks.

5.2 MOS Installation

Liberty on Ubuntu 14 is the version used from MOS 8.0 Node

Create a new OpenStack environment

Name and Release

Name pluribus

Compute

OpenStack Release

Liberty on Ubuntu 14.04

Networking Setup

Storage Backends

Additional Services

Finish

By default, Fuel uploads the software packages for the Fuel Slave nodes from the external repositories. Please verify the Fuel Master node has the Internet connection. If the Fuel Master node does not have access to the Internet, you must create a local mirror with all required software packages and configure Fuel to use the mirror before you deploy an OpenStack environment.

This option will install the OpenStack Liberty packages using Ubuntu as a base operating system. With high availability features built in, you are getting a robust, enterprise-grade OpenStack deployment.

Cancel

← Prev

Next →

KVM Based Compute

Create a new OpenStack environment

Name and Release

QEMU-KVM

Select this option if you want to use QEMU as a hypervisor with capability of KVM acceleration.

Compute

vCenter ⚠

Select this option if you run OpenStack on VMware vCenter.

Networking Setup

Plugin for DVS/NSX is required to create an environment with vCenter and Neutron.
Please visit Fuel plugins page for details.

Storage Backends

Additional Services

Finish

Cancel

← Prev

Next →

VLAN Based Segmentation

Create a new OpenStack environment

Name and Release

Compute

Networking Setup

Storage Backends

Additional Services

Finish

Neutron with ML2 plugin

Framework that enables simultaneous utilization of the layer 2 networking technologies through drivers.

Neutron with VLAN segmentation

Your network hardware must be configured for VLAN segmentation. This option supports up to 4095 networks.

Neutron with tunneling segmentation

By default VXLAN tunnels will be used. This option supports millions of tenant data networks.

[Cancel](#)

[◀ Prev](#)

[Next ➔](#)

Ceph is used for all Storage

Create a new OpenStack environment

Name and Release

Compute

Networking Setup

Storage Backends

Additional Services

Finish

Block Storage:

LVM

Use default storage providers

Ceph

Use Ceph as backend for Cinder volumes

Object Storage:

Ceph

Use Ceph as backend for Swift objects

Image Storage:

Ceph

Use Ceph as backend for Glance images

Ephemeral Storage:

Ceph

Use Ceph as backend for Nova

[Cancel](#)

[◀ Prev](#)

[Next ➔](#)

L2 Configuration

The screenshot shows the L2 Configuration page for the 'pluribus' environment, which contains 6 nodes. The top navigation bar includes links for Home, Environments, and Networks, along with tabs for Dashboard, Nodes, Networks, Settings, Logs, and Health Check.

The main section is titled 'Network Settings (Neutron with VLAN segmentation)'. It features a 'Node Network Groups' table with a single entry for 'default'. Under 'Neutron L2 Configuration', there is a note about Neutron supporting VLAN, GRE, VXLAN segmentation, and specific parameters like VLAN ID ranges and Base MAC address. The 'VLAN ID range' is set from 1000 to 1030. The 'Base MAC address' is listed as fa:16:3e:00:00:00. A button labeled 'Add New Node Network Group' is visible.

On the left, there are sections for 'Settings' (selected), 'Neutron L2' (highlighted in blue), 'Neutron L3' (disabled), and 'Other'. Below these are 'Network Verification' and 'Connectivity Check' buttons.

At the bottom right are 'Cancel Changes' and 'Save Settings' buttons.

3 Redundant controller nodes, 1 Compute and 2 Ceph Storage node was used

Controller (3)

Select All

<input checked="" type="checkbox"/>	Untitled (cc:3b) CONTROLLER		PENDING ADDITION	CPU: 2 (32) HDD: 5.8 TB RAM: 64.0 GB	
<input checked="" type="checkbox"/>	Untitled (cc:11) CONTROLLER		PENDING ADDITION	CPU: 2 (32) HDD: 5.8 TB RAM: 64.0 GB	
<input checked="" type="checkbox"/>	Untitled (cc:0f) CONTROLLER		PENDING ADDITION	CPU: 2 (32) HDD: 5.8 TB RAM: 64.0 GB	

Compute (1)

Select All

<input checked="" type="checkbox"/>	SUPER MICRO Untitled (fc:47) COMPUTE		PENDING ADDITION	CPU: 2 (32) HDD: 8.7 TB RAM: 32.0 GB	
-------------------------------------	---	--	------------------	--------------------------------------	--

Storage - Ceph OSD (2)

Select All

<input checked="" type="checkbox"/>	SUPER MICRO Untitled (fc:4b) CEPH-OSD		PENDING ADDITION	CPU: 2 (32) HDD: 8.7 TB RAM: 32.0 GB	
<input checked="" type="checkbox"/>	SUPER MICRO Untitled (fc:49) CEPH-OSD		PENDING ADDITION	CPU: 2 (32) HDD: 8.7 TB RAM: 32.0 GB	

5.3 Pluribus Netvisor ML2 Plugin installation steps

The first Step after Mirantis OS setup is complete, is to enable SSH/SFTP Access to the Controller nodes from the Pluribus Switches. Login to the controller nodes and edit the /etc/ssh/ssd_config file and edit the following parameters:

```
PasswordAuthentication yes  
PermitRootLogin yes
```

Perform the following perquisite steps and then follow steps 1-4 as described below to download and install the plugin from the switch:

Prerequisites:

VNET, VLANs and other required configs have already been created on the switch. This includes the following steps:

- a. vNET-create
 - b. vlan-create (of the VLAN assigned to the VNET)

- c. openstack-config-create
- d. openstack-interface-add (Add a management interface to the openstack instance)

Example configuration is shown below:

```
CLI (network-admin@switch) > vNET-create name test-vNET scope fabric vlans 20 vNET-mgr-storage-pool pool-disk2
Creating test-vNET-mgr zone, please wait...
VNET created. Vlans assigned: 20
CLI (network-admin@switch) > vlan-create id 20 scope fabric
Vlan 20 created
CLI (network-admin@switch) > openstack-config-create name test-openstack vNET test-vNET port 9090 vlans 200-300
CLI (network-admin@switch) > openstack-interface-add openstack-name test-openstack ip 10.100.100.227/16 vlan 0 if mgmt
Added interface eth0
```

The pluribus ML2 Driver resides on the switch.

Perform the following steps to download the package to the Openstack controller:

1. Enable SFTP on the switch

Run the following command on the switch CLI and provide a password for SFTP transfer:
CLI> admin-sftp-modify enable

Example:

```
CLI (network-admin@switch) > admin-sftp-modify enable
sftp password:
confirm sftp password:
```

2. Download the package from the switch to the controller:

The pn-ml2-package will be available in the "/sftp/export/pkg/debian/" directory on the switch.
Use sftp to login to the switch from the controller and get the pn_openstack_plugin.deb file.

Example:

```
root@controller:/# sftp sftp@switch.pluribusnetworks.com
Password:
Connected to switch.pluribusnetworks.com.
sftp> ls
disk-lib export import iso-lib nvOS    usr
sftp> cd export/pkg/debian
sftp> ls
pn_openstack_plugin.deb
sftp> get pn_openstack_plugin.deb
Fetching /export/pkg/debian/pn_openstack_plugin.deb to pn_openstack_plugin.deb
/export/pkg/debian/pn_openstack_plugin.deb                                100% 10KB 10.3KB/s
00:00
```

3. Install the package on the controller:

Use command "dpkg -i pn_openstack_plugin.deb" to install the package

Example:

```
root@controller:~# dpkg -i pn_openstack_plugin.deb
Selecting previously unselected package pn-openstack-plugin.
(Reading database ... 113621 files and directories currently installed.)
Preparing to unpack pn_openstack_plugin.deb ...
Unpacking pn-ml2-plugin (2.4.1) ...
Setting up pn-ml2-plugin (2.4.1) ...
```

You can use command "dpkg -L pn-ml2-plugin" to get a list of files installed by the package.

4. Run the installer script to configure the plugin:

Prerequisite:

You will need a "plugin.conf" file as input which contains the details of the openstack configuration.

An example of plugin.conf for the above configuration is shown below:

```
root@controller:# cat plugin.conf
[PLURIBUS_PLUGINS]
vswitch_plugin = neutron.plugins.ml2.plugin.Ml2Plugin
pn_port = 9090
pn_vlans = 200-300
extensions = enable
pn_switch = 10.100.100.227
```

A sample plugin.conf is provided in the package at /usr/lib/python2.7/dist-packages/pluribus/install/py/plugin.conf

Please refer to the sample and fill in the required fields.

a. Run the installer script

The installer script is available in "/usr/bin/python /usr/lib/python2.7/dist-packages/pluribus/install/py/pn_config_pkg.py".

Run this script as follows:

```
/usr/bin/python /usr/lib/python2.7/dist-packages/pluribus/install/py/pn_config_pkg.py -f <path-to-plugin.conf> -e <enable/disable extensions>
```

Example:

```
root@controller:# /usr/bin/python /usr/lib/python2.7/dist-
packages/pluribus/install/py/pn_config_pkg.py -f plugin.conf -e disable
```

b. Restart the neutron-server service:

Use command "service neutron-server restart" , If no errors are seen, the installation went ok.

5.4 Limitations

Currently Netvisor creates a zone per VLAN and this might restrict the number of VLANs a user can deploy.

Redundancy/HA and Layer 3 Neutron plugins are outside the scope of this document.

5.5 Testing

5.5.1 Test cases and Test Results

Testing done with the debian openstack plugin package installation

1. Neutron net-create and vNET-show

```
root@node-21:~# neutron net-create net1
Created a new network:
+-----+-----+
| Field | Value |
+-----+-----+
| admin_state_up | True |
| id | 06adb6ab-8051-4b2a-8e75-3abe5f34c36a |
| mtu | 0 |
| name | net1 |
| port_security_enabled | True |
| provider:network_type | vlan |
| provider:physical_network | physnet2 |
| provider:segmentation_id | 1004 |
| router:external | False |
| shared | False |
| status | BUILD |
| subnets |
| tenant_id | 60828356c8624758a9a690281c0b5b10 |
+-----+
```

```
CLI (network-admin@Spine) > vNET-show
switch name      scope  vlans managed-ports admin      vNET-mgr-name
----- -----
Spine mirantis    fabric 5     none      mirantis-admin  mirantis-mgr
Spine net1-1004   fabric 1004   none      net1-1004-admin net1-1004-mgr
Spine openstack-global fabric none   none
```

2. Neutron subnet-create and ip-pool-show, dhcp-show

```
root@node-21:~# neutron subnet-create net1 192.168.10.0/24 --name subnet1
Created a new subnet:
```

```
+-----+-----+
| Field | Value |
+-----+-----+
| allocation_pools | [{"start": "192.168.10.2", "end": "192.168.10.254"}] |
| cidr | 192.168.10.0/24 |
| dns_nameservers |
| enable_dhcp | True |
| gateway_ip | 192.168.10.1 |
| host_routes |
| id | e82118e8-c551-4ece-909b-f22fee7db105 |
| ip_version | 4 |
| ipv6_address_mode |
```

ipv6_ra_mode	
name	subnet1
network_id	06adb6ab-8051-4b2a-8e75-3abe5f34c36a
subnetpool_id	
tenant_id	60828356c8624758a9a690281c0b5b10

```
CLI (network-admin@Spine) > ip-pool-show
name      vNET      scope  vlan start-ip      end-ip      network
-----
subnet1-0 net1-1004 fabric 1004 192.168.10.3 192.168.10.254 192.168.10.0/24
CLI (network-admin@Spine) > dhcp-show
name          type scope vNET      is-global vNET-service state   gateway pxe-boot
-----
net1-1004-dhcp dhcp fabric net1-1004 false     shared      enabled ::      disabled
```

3. Nova boot and dhcp-host-show

```
root@node-21:~# nova boot --flavor m1.micro --image TestVM --nic net-id=06adb6ab-8051-4b2a-8e75-3abe5f34c36a vm-1
```

Property	Value
OS-DCF:diskConfig	MANUAL
OS-EXT-AZ:availability_zone	-
OS-EXT-SRV-ATTR:host	-
OS-EXT-SRV-ATTR:hypervisor_hostname	-
OS-EXT-SRV-ATTR:instance_name	instance-00000009
OS-EXT-STS:power_state	0
OS-EXT-STS:task_state	-
OS-EXT-STS:vm_state	building
OS-SRV-USG:launched_at	-
OS-SRV-USG:terminated_at	-
accessIPv4	
accessIPv6	
adminPass	49ECMxeNq2bD
config_drive	
created	2016-08-26T06:09:35Z
flavor	m1.micro (1cdb9c44-a617-457f-982f-81f1cffcff5)
hostId	
id	949b006b-424c-4665-a587-a39b6dc95718
image	TestVM (cf5bf88c-8784-493f-93f7-b50799ec818f)
key_name	-
metadata	{}
name	vm-1
os-extended-volumes:volumes_attached	[]
progress	0
security_groups	default
status	BUILD
tenant_id	60828356c8624758a9a690281c0b5b10
updated	2016-08-26T06:09:35Z
user_id	d214c7787bf44976a0109a557926bee4

```
CLI (network-admin@Spine) > dhcp-host-show
dhcp-name      hostname      mac      pxe-boot gateway-ip      fixed-ip
```

```
-----  
net1-1004-dhcp fa:16:3e:48:de:3b fa:16:3e:48:de:3b no 192.168.10.1 192.168.10.5  
-----
```

Further testing:

Scenario 1: Traffic between two VMs in the same network.

Boot up another VM in the same subnet as vm-1 (192.168.10.0/24)

```
root@node-21:~# nova boot --flavor m1.micro --image TestVM --nic net-id=06adb6ab-8051-4b2a-8e75-3abe5f34c36a vm-same
```

Property	Value
OS-DCF:diskConfig	MANUAL
OS-EXT-AZ:availability_zone	-
OS-EXT-SRV-ATTR:host	-
OS-EXT-SRV-ATTR:hypervisor_hostname	-
OS-EXT-SRV-ATTR:instance_name	instance-00000015
OS-EXT-STS:power_state	0
OS-EXT-STS:task_state	scheduling
OS-EXT-STS:vm_state	building
OS-SRV-USG:launched_at	-
OS-SRV-USG:terminated_at	-
accessIPv4	
accessIPv6	
adminPass	DUrdczAuk5mH
config_drive	
created	2016-08-26T09:36:23Z
flavor	m1.micro (1cdb9c44-a617-457f-982f-81f1cffcff5)
hostId	
id	2c298ea7-e5f3-4777-97d5-ded33678401c
image	TestVM (cf5bf88c-8784-493f-93f7-b50799ec818f)
key_name	-
metadata	{}
name	vm-same
os-extended-volumes:volumes_attached	[]
progress	0
security_groups	default
status	BUILD
tenant_id	60828356c8624758a9a690281c0b5b10
updated	2016-08-26T09:36:22Z
user_id	d214c7787bf44976a0109a557926bee4

```

Connected (unencrypted) to: QEMU (instance-00000015)
Send CtrlAltDel
$ ifconfig eth0
eth0      Link encap:Ethernet HWaddr FA:16:3E:39:A2:9B
          inet addr:192.168.10.6 Bcast:192.168.10.255 Mask:255.255.255.0
          inet6 addr: fe80::f816:3eff:fe39:a29b/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
          RX packets:103 errors:0 dropped:0 overruns:0 frame:0
          TX packets:106 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:8763 (8.5 KiB) TX bytes:11612 (11.3 KiB)

$ ping 192.168.10.5
PING 192.168.10.5 (192.168.10.5): 56 data bytes
64 bytes from 192.168.10.5: seq=0 ttl=64 time=5.413 ms
64 bytes from 192.168.10.5: seq=1 ttl=64 time=1.325 ms
64 bytes from 192.168.10.5: seq=2 ttl=64 time=1.383 ms
64 bytes from 192.168.10.5: seq=3 ttl=64 time=0.994 ms
64 bytes from 192.168.10.5: seq=4 ttl=64 time=1.116 ms
64 bytes from 192.168.10.5: seq=5 ttl=64 time=0.952 ms
--- 192.168.10.5 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max = 0.952/1.863/5.413 ms
$ -

```

Ping vm-1 in subnet 192.168.10.0/24 from vm-same in 192.168.10.0/24

Scenario 2: Traffic between two VMs in different networks.

1. Boot up vm-2 in another subnet.
2. Create a router and add interfaces to the subnets: subnet1 and subnet2:

```

root@node-20:~# neutron router-show router1
+-----+-----+
| Field           | Value        |
+-----+-----+
| admin_state_up  | True         |
| distributed     | False        |
| external_gateway_info |
| ha              | False        |
| id              | 835b7d1a-a4de-4dd0-b517-9a9b359ae58d |
| name            | router1      |
| routes          |             |
| status          | ACTIVE       |
| tenant_id       | 60828356c8624758a9a690281c0b5b10 |
+-----+-----+
root@node-20:~# neutron router-port-list router1
+-----+-----+-----+
| id           | name | mac_address | fixed_ips |
+-----+-----+-----+
| 1546741d-156d-4765-bccf-4e1210203326 | fa:16:3e:56:09:cf | {"subnet_id": "d0033218-1972-4ed2-9c1f-60bab6b802a1", "ip_address": "192.168.60.1"} |
+-----+-----+-----+

```

```

| 2a706f19-24ba-4a8c-a268-9a0b779c4578 |      | fa:16:3e:04:09:ec | {"subnet_id": "4c22241e-057e-4ed1-86b0-aefae68b8ab7", "ip_address": "192.168.30.1"} |
| 8949301d-6a75-4d25-be19-78451df2b3f1 |      | fa:16:3e:0d:64:c6 | {"subnet_id": "e8dbe092-a3f4-4e57-9efe-ea6aa3b76409", "ip_address": "192.168.20.1"} |
| 9438a7fd-35be-4edc-853b-ba0a8f06c5d5 |      | fa:16:3e:fc:86:03 | {"subnet_id": "e82118e8-c551-4ece-909b-f22fee7db105", "ip_address": "192.168.10.1"} |
+-----+-----+-----+

```

```

eth0      Link encap:Ethernet HWaddr FA:16:3E:48:DE:3B
          inet addr:192.168.10.5  Bcast:192.168.10.255  Mask:255.255.255.0
          inet6 addr: fe80::f816:3eff:fe48:de3b/64 Scope:Link
             UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
             RX packets:184 errors:0 dropped:0 overruns:0 frame:0
             TX packets:155 errors:0 dropped:0 overruns:0 carrier:0
             collisions:0 txqueuelen:1000
             RX bytes:16022 (15.6 KiB)  TX bytes:16506 (16.1 KiB)

$ ping 192.168.20.5
PING 192.168.20.5 (192.168.20.5): 56 data bytes
64 bytes from 192.168.20.5: seq=0 ttl=63 time=7.608 ms
64 bytes from 192.168.20.5: seq=1 ttl=63 time=1.121 ms
64 bytes from 192.168.20.5: seq=2 ttl=63 time=1.520 ms
64 bytes from 192.168.20.5: seq=3 ttl=63 time=1.160 ms
64 bytes from 192.168.20.5: seq=4 ttl=63 time=1.466 ms
64 bytes from 192.168.20.5: seq=5 ttl=63 time=1.304 ms

--- 192.168.20.5 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max = 1.121/2.363/7.608 ms

```

Ping vm-2 in subnet 192.168.20.0/24 from vm-1 in 192.168.10.0/24 via Neutron router

Scenario 3: Provisioning a VM through Horizon

Hostname	Type	VCPUs (used)	VCPUs (total)	RAM (used)	RAM (total)	Local Storage (used)	Local Storage (total)	Instances
node-23.domain.tld	QEMU	6	32	1.3GB	31.3GB	1GB	13.1TB	6

Looking at the VMs from Horizon

Project	Host	Name	Image Name	IP Address	Size	Status	Task	Power State	Time since created	Actions
admin	node-23.domain.tld	new-vm	TestVM	192.168.10.7	m1.micro	Active	None	Running	5 days, 21 hours	<button>Edit Instance</button>
admin	-	amitsi	TestVM		m1.tiny	Error	None	No State	5 days, 23 hours	<button>Edit Instance</button>
admin	node-23.domain.tld	vm-same	TestVM	192.168.10.6	m1.micro	Active	None	Running	2 weeks, 3 days	<button>Edit Instance</button>
admin	node-23.domain.tld	vm-4	TestVM	192.168.60.5	m1.micro	Active	None	Running	2 weeks, 3 days	<button>Edit Instance</button>
admin	node-23.domain.tld	vm3	TestVM	192.168.30.5	m1.tiny	Active	None	Running	2 weeks, 3 days	<button>Edit Instance</button>
admin	node-23.domain.tld	vm-2	TestVM	192.168.20.5	m1.micro	Active	None	Running	2 weeks, 3 days	<button>Edit Instance</button>
admin	node-23.domain.tld	vm-1	TestVM	192.168.10.5	m1.micro	Active	None	Running	2 weeks, 3 days	<button>Edit Instance</button>

Displaying 7 items

Scenario 4:

A test case was to test stability and reconfiguration. We deleted the existing network configuration from above and added the networks to Public Network (google) connectivity with the default configured Public Router

```
root@node-21:~# neutron subnet-list
+-----+-----+
| id | name |
| cidr | allocation_pools |
+-----+-----+
| 4c22241e-057e-4ed1-86b0-aefae68b8ab7 | subnet3 |
| 192.168.30.0/24 | {"start": "192.168.30.2", "end": "192.168.30.254"} |
| d0033218-1972-4ed2-9c1f-60bab6b802a1 | subnet4 |
| 192.168.60.0/24 | {"start": "192.168.60.2", "end": "192.168.60.254"} |
| e82118e8-c551-4ece-909b-f22fee7db105 | subnet1 |
| 192.168.10.0/24 | {"start": "192.168.10.2", "end": "192.168.10.254"} |
| 28d91cba-e048-4b49-9d49-5c77ea142bd5 | admin_floating_net_subnet |
| 10.9.0.0/16 | {"start": "10.9.101.130", "end": "10.9.101.254"} |
| e8dbe092-a3f4-4e57-9efe-ea6aa3b76409 | subnet2 |
```

```

192.168.20.0/24 | {"start": "192.168.20.2", "end": "192.168.20.254"} |
| 5aff970f-939a-4a02-80f9-bc1b2422b3bc | admin_internal_net_subnet |
192.168.111.0/24 | {"start": "192.168.111.2", "end": "192.168.111.254"} |
+-----+
--+
root@node-21:~# neutron router-list
+-----+
-----+
-----+
----+
| id | name | external_gateway_info |
-----+
| distributed | ha |
+-----+
-----+
----+
| 3a282fa7-7aee-4430-a5f1-85155a32b20c | router04 | {"network_id": "17fab5f6-1dfc-4642-91dc-4027c66177ca", "enable_snat": true, "external_fixed_ips": [{"subnet_id": "28d91cba-e048-4b49-9d49-5c77ea142bd5", "ip_address": "10.9.101.130"}]} |
False | False |
| 835b7d1a-a4de-4dd0-b517-9a9b359ae58d | router1 |
null

| False | False |
+-----+
-----+
----+
root@node-21:~# neutron router-interface-delete router1 4c22241e-057e-4ed1-86b0-aefae68b8ab7
Removed interface from router router1.
root@node-21:~# neutron router-interface-delete router1 d0033218-1972-4ed2-9c1f-60bab6b802a1
Removed interface from router router1.
root@node-21:~# neutron router-interface-delete router1 e82118e8-c551-4ece-909b-f22fee7db105
Removed interface from router router1.
root@node-21:~# neutron router-interface-delete router1 e8dbe092-a3f4-4e57-9efe-ea6aa3b76409
Removed interface from router router1.

root@node-21:~# neutron router-delete router1
Deleted router: router1

root@node-21:~# neutron router-interface-add router04 4c22241e-057e-4ed1-86b0-aefae68b8ab7
Added interface 7899dee2-4ee1-4c72-924d-c129e7e46d82 to router router04.
root@node-21:~# neutron router-interface-add router04 d0033218-1972-4ed2-9c1f-60bab6b802a1
Added interface 83a5721e-562a-4247-87b2-b700be709e59 to router router04.
root@node-21:~# neutron router-interface-add router04 e82118e8-c551-4ece-909b-

```

```
f22fee7db105
Added interface 0b882f4e-5330-46cd-bcd0-56350f88f154 to router router04.
root@node-21:~# neutron router-interface-add router04 e8dbe092-a3f4-4e57-9efe-
ea6aa3b76409
Added interface fd5669a6-d815-470e-8d7a-dbbeaf154e12 to router router04.
```

Configured Interface on VM

```
Connected (unencrypted) to: QEMU (instance-0000001b)
64 bytes from 8.8.8.8: seq=0 ttl=55 time=5.896 ms

--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max = 5.896/5.896/5.896 ms
$ ifconfig
eth0      Link encap:Ethernet HWaddr FA:16:3E:C5:38:41
          inet addr:192.168.10.7 Bcast:192.168.10.255 Mask:255.255.255.0
          inet6 addr: fe80::f816:3eff:fe:3841/64 Scope:Link
            UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
            RX packets:2980 errors:0 dropped:0 overruns:0 frame:0
            TX packets:743 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:1000
            RX bytes:253457 (247.5 KiB) TX bytes:124066 (121.1 KiB)

lo       Link encap:Local Loopback
          inet addr:127.0.0.1 Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
            UP LOOPBACK RUNNING MTU:16436 Metric:1
            RX packets:12 errors:0 dropped:0 overruns:0 frame:0
            TX packets:12 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:1020 (1020.0 B) TX bytes:1020 (1020.0 B)

$ _
```

Testing Public Networking Connectivity

```
Connected (unencrypted) to: QEMU (instance-0000001b)
$  
$  
$ ping 8.8.8.8  
PING 8.8.8.8 (8.8.8.8): 56 data bytes  
z[1]+ Stopped ping 8.8.8.8  
$ ping 8.8.8.8  
PING 8.8.8.8 (8.8.8.8): 56 data bytes  
64 bytes from 8.8.8.8: seq=0 ttl=55 time=6.757 ms  
64 bytes from 8.8.8.8: seq=1 ttl=55 time=5.037 ms  
64 bytes from 8.8.8.8: seq=2 ttl=55 time=4.891 ms  
64 bytes from 8.8.8.8: seq=3 ttl=55 time=5.220 ms  
64 bytes from 8.8.8.8: seq=4 ttl=55 time=3.908 ms  
64 bytes from 8.8.8.8: seq=5 ttl=55 time=4.848 ms  
[2]+ Stopped ping 8.8.8.8  
$  
$  
$ ping 8.8.8.8  
PING 8.8.8.8 (8.8.8.8): 56 data bytes  
64 bytes from 8.8.8.8: seq=0 ttl=55 time=12.378 ms  
64 bytes from 8.8.8.8: seq=1 ttl=55 time=5.129 ms  
--- 8.8.8.8 ping statistics ---  
2 packets transmitted, 2 packets received, 0% packet loss  
round-trip min/avg/max = 5.129/8.753/12.378 ms  
$ ping 8.8.8.8
```

Scenario 5:

Check Health and verify all relevant tests through Fuel Health Check page pass

The screenshot shows the Mirantis OpenStack Fuel web interface. At the top, there's a navigation bar with links like 'Home', 'ENVIRONMENTS', 'EQUIPMENT', 'RELEASES', 'PLUGINS', and 'SUPPORT'. Below the navigation is a user profile icon and a notification bell. The main content area is titled 'pluribus (6 nodes)' and shows a 'Health Check' section. A table lists 12 items under the heading 'OpenStack Health Check'. Each item includes a checkbox, a description, an expected duration, an actual duration, and a status column with a green checkmark. There are buttons for 'Select All', 'Provide credentials', and 'Stop Tests'.

		Expected Duration	Actual Duration	Status
<input checked="" type="checkbox"/>	Sanity tests. Duration 30 sec - 2 min			✓
<input checked="" type="checkbox"/>	Request flavor list	20 s.	0.8	✓
<input checked="" type="checkbox"/>	Request image list using Nova	20 s.	0.7	✓
<input checked="" type="checkbox"/>	Request instance list	20 s.	0.8	✓
<input checked="" type="checkbox"/>	Request absolute limits list	20 s.	0.0	✓
<input checked="" type="checkbox"/>	Request snapshot list	20 s.	0.7	✓
<input checked="" type="checkbox"/>	Request volume list	20 s.	0.6	✓
<input checked="" type="checkbox"/>	Request image list using Glance v1	10 s.	0.0	✓
<input checked="" type="checkbox"/>	Request image list using Glance v2	10 s.	0.0	✓
<input checked="" type="checkbox"/>	Request stack list	20 s.	0.0	✓
<input checked="" type="checkbox"/>	Request active services list	20 s.	0.4	✓

This screenshot shows the same 'Health Check' interface as the previous one, but with a different set of tests. It displays 11 items under the heading 'HA tests. Duration 30 sec - 8 min'. Each item has a checkbox, a description, an expected duration, an actual duration, and a status column with a green checkmark. The items include checks for haproxy backends, MySQL replication, Galera environment state, Pacemaker status, RabbitMQ availability, and RabbitMQ replication.

		Expected Duration	Actual Duration	Status
<input checked="" type="checkbox"/>	Create user and authenticate with it.	80 s.	8.9	✓
<input checked="" type="checkbox"/>	HA tests. Duration 30 sec - 8 min			
<input checked="" type="checkbox"/>	Check state of haproxy backends on controllers	10 s.	1.2	✓
<input checked="" type="checkbox"/>	Check data replication over mysql	10 s.	3.7	✓
<input checked="" type="checkbox"/>	Check if amount of tables in databases is the same on each node	10 s.	4.0	✓
<input checked="" type="checkbox"/>	Check galera environment state	10 s.	1.8	✓
<input checked="" type="checkbox"/>	Check pacemaker status	10 s.	1.5	✓
<input checked="" type="checkbox"/>	RabbitMQ availability	100 s.	12.4	✓
<input checked="" type="checkbox"/>	RabbitMQ replication	100 s.	29.7	✓

Scenario 6:

Test Using Rally

```
root@node-21:~/rally# rally task start
samples/tasks/scenarios/neutron/create-and-delete-networks.yaml
```

```
-----  
---  
Preparing input task  
-----  
---  
  
Input task is:  
---  
NeutronNetworks.create_and_delete_networks:  
-  
  args:  
    network_create_args: {}  
  runner:  
    type: "constant"  
    times: 10  
    concurrency: 1  
  context:  
    users:  
      tenants: 1  
      users_per_tenant: 1  
    quotas:  
      neutron:  
        network: -1  
  
Task syntax is correct :)  
2016-09-16 05:45:03.079 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Starting: Task validation.  
2016-09-16 05:45:03.085 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Starting: Task validation of scenarios names.  
2016-09-16 05:45:03.088 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Completed: Task validation of scenarios names.  
2016-09-16 05:45:03.088 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Starting: Task validation of syntax.  
2016-09-16 05:45:03.107 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Completed: Task validation of syntax.  
2016-09-16 05:45:03.107 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Starting: Task validation of semantic.  
2016-09-16 05:45:03.107 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Starting: Task validation check cloud.  
2016-09-16 05:45:03.348 2331 INFO rally.task.engine [-] Task 04794c49-3996-  
4cb0-92c2-78e830359904 | Completed: Task validation check cloud.  
2016-09-16 05:45:03.358 2331 INFO  
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-  
92c2-78e830359904 | Starting: Enter context: `users`  
2016-09-16 05:45:04.133 2331 INFO  
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-  
92c2-78e830359904 | Completed: Enter context: `users`  
2016-09-16 05:45:04.257 2331 INFO  
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-  
92c2-78e830359904 | Starting: Exit context: `users`
```

```
2016-09-16 05:45:06.961 2331 INFO
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-
92c2-78e830359904 | Completed: Exit context: `users`
2016-09-16 05:45:06.962 2331 INFO rally.task.engine [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | Completed: Task validation of semantic.
2016-09-16 05:45:06.962 2331 INFO rally.task.engine [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | Completed: Task validation.
Task config is valid :)
```

```
-----
---  
Task 04794c49-3996-4cb0-92c2-78e830359904: started
```

```
-----  
---  
Benchmarking... This can take a while...
```

To track task status use:

```
rally task status  
or  
rally task detailed
```

```
Using task: 04794c49-3996-4cb0-92c2-78e830359904
2016-09-16 05:45:06.975 2331 INFO rally.api [-] Benchmark Task 04794c49-3996-
4cb0-92c2-78e830359904 on Deployment 61888ec5-8df6-49b9-aecc-e89f0342351d
2016-09-16 05:45:06.978 2331 INFO rally.task.engine [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | Starting: Benchmarking.
2016-09-16 05:45:06.985 2331 INFO rally.task.engine [-] Running benchmark
with key:
{
  "kw": {
    "runner": {
      "type": "constant",
      "concurrency": 1,
      "times": 10
    },
    "args": {
      "network_create_args": {}
    },
    "context": {
      "users": {
        "users_per_tenant": 1,
        "tenants": 1
      },
      "quotas": {
        "neutron": {
          "network": -1
        }
      }
    }
  }
}
```

```
},
  "name": "NeutronNetworks.create_and_delete_networks",
  "pos": 0
}
2016-09-16 05:45:06.992 2331 INFO
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-
92c2-78e830359904 | Starting: Enter context: `users`
2016-09-16 05:45:07.675 2331 INFO
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-
92c2-78e830359904 | Completed: Enter context: `users`
2016-09-16 05:45:07.675 2331 INFO
rally.plugins.openstack.context.quotas.quotas [-] Task 04794c49-3996-4cb0-
92c2-78e830359904 | Starting: Enter context: `quotas`
2016-09-16 05:45:07.948 2331 INFO
rally.plugins.openstack.context.quotas.quotas [-] Task 04794c49-3996-4cb0-
92c2-78e830359904 | Completed: Enter context: `quotas`
2016-09-16 05:45:07.968 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 1 START
2016-09-16 05:45:09.690 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 1 END: OK
2016-09-16 05:45:09.695 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 2 START
2016-09-16 05:46:04.514 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 2 END: OK
2016-09-16 05:46:04.523 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 3 START
2016-09-16 05:46:58.585 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 3 END: OK
2016-09-16 05:46:58.591 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 4 START
2016-09-16 05:47:53.959 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 4 END: OK
2016-09-16 05:47:53.964 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 5 START
2016-09-16 05:47:55.748 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 5 END: OK
2016-09-16 05:47:55.756 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 6 START
2016-09-16 05:47:57.450 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 6 END: OK
2016-09-16 05:47:57.457 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 7 START
2016-09-16 05:47:59.267 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 7 END: OK
2016-09-16 05:47:59.272 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 8 START
2016-09-16 05:48:52.818 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 8 END: OK
2016-09-16 05:48:52.823 3430 INFO rally.task.runner [-] Task 04794c49-3996-
4cb0-92c2-78e830359904 | ITER: 9 START
```

```
2016-09-16 05:49:46.450 3430 INFO rally.task.runner [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | ITER: 9 END: OK
2016-09-16 05:49:46.456 3430 INFO rally.task.runner [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | ITER: 10 START
2016-09-16 05:50:40.096 3430 INFO rally.task.runner [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | ITER: 10 END: OK
2016-09-16 05:50:40.106 2331 INFO
rally.plugins.openstack.context.cleanup.user [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | Starting: user resources cleanup
2016-09-16 05:50:42.152 2331 INFO
rally.plugins.openstack.context.cleanup.user [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | Completed: user resources cleanup
2016-09-16 05:50:42.152 2331 INFO
rally.plugins.openstack.context.quotas.quotas [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | Starting: Exit context: `quotas`
2016-09-16 05:50:42.310 2331 INFO
rally.plugins.openstack.context.quotas.quotas [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | Completed: Exit context: `quotas`
2016-09-16 05:50:42.310 2331 INFO
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | Starting: Exit context: `users`
2016-09-16 05:50:44.925 2331 INFO
rally.plugins.openstack.context.keystone.users [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | Completed: Exit context: `users`
2016-09-16 05:50:46.122 2331 INFO rally.task.engine [-] Load duration is: 332.126073
2016-09-16 05:50:46.123 2331 INFO rally.task.engine [-] Full runner duration is: 332.152256
2016-09-16 05:50:46.123 2331 INFO rally.task.engine [-] Full duration is 337.936048
2016-09-16 05:50:46.141 2331 INFO rally.task.engine [-] Task 04794c49-3996-4cb0-92c2-78e830359904 | Completed: Benchmarking.
```

```
-----  
---  
Task 04794c49-3996-4cb0-92c2-78e830359904: finished  
-----  
---
```

```
test scenario NeutronNetworks.create_and_delete_networks
args position 0
args values:
{
  "runner": {
    "type": "constant",
    "concurrency": 1,
    "times": 10
  },
  "args": {
    "network_create_args": {}
  }
}
```

```
        },
      "context": {
        "users": {
          "users_per_tenant": 1,
          "project_domain": "default",
          "user_choice_method": "random",
          "user_domain": "default",
          "tenants": 1,
          "resource_management_workers": 20
        },
        "quotas": {
          "neutron": {
            "network": -1
          }
        }
      }
    }
```

Task 04794c49-3996-4cb0-92c2-78e830359904 has 0 error(s)

Response Times (sec)					
Action	Min (sec)	Median (sec)	90%ile (sec)	95%ile (sec)	
(sec)	Max (sec)	Avg (sec)	Success	Count	
neutron.create_network	0.392	0.431	0.447	0.487	
0.528	0.437	100.0%	10		
neutron.delete_network	1.264	53.109	54.476	54.702	
54.929	32.769	100.0%	10		
total	1.694	53.585	54.873	55.12	
55.367	33.206	100.0%	10		

Load duration: 332.126073

Full duration: 337.936048

HINTS:

* To plot HTML graphics with this data, run:

```
rally task report 04794c49-3996-4cb0-92c2-78e830359904 --out  
output.html
```

```

* To generate a JUnit report, run:
    rally task report 04794c49-3996-4cb0-92c2-78e830359904 --junit --out
output.xml

* To get raw JSON output of task results, run:
    rally task results 04794c49-3996-4cb0-92c2-78e830359904

root@node-21:~/rally# rally task report 04794c49-3996-4cb0-92c2-78e830359904
--out output.html
root@node-21:~/rally#

```

```

CLI (network-admin@Spine) > vNET-show
switch name                      scope  vlans managed-ports admin
vNET-mgr-name

-----
-----
Spine mirantis                  fabric 5      none      mirantis-
admin                         mirantis-mgr
Spine net1-1004                 fabric 1004   none      net1-1004-
admin                         net1-1004-mgr
Spine net2-1030                 fabric 1030   none      net2-1030-
admin                         net2-1030-mgr
Spine net3-1001                 fabric 1001   none      net3-1001-
admin                         net3-1001-mgr
Spine net4-1025                 fabric 1025   none      net4-1025-
admin                         net4-1025-mgr
Spine newnet-1005                fabric 1005   none      newnet-1005-
admin                         newnet-1005-mgr
Spine openstack-global           fabric none   none
Spine s_rally_03b1ba3e_3bxyqeyW-1003 fabric 1003   none
s_rally_03b1ba3e_3bxyqeyW-1003-admin s_rally_03b1ba3e_3bxyqeyW-1003-mgr
Spine s_rally_03b1ba3e_5lMPmuCA-1008 fabric 1008   none
s_rally_03b1ba3e_5lMPmuCA-1008-admin s_rally_03b1ba3e_5lMPmuCA-1008-mgr
Spine s_rally_03b1ba3e_D8q4WVih-1014 fabric 1014   none
s_rally_03b1ba3e_D8q4WVih-1014-admin s_rally_03b1ba3e_D8q4WVih-1014-mgr
Spine s_rally_03b1ba3e_Gab4XcBL-1024 fabric 1024   none
s_rally_03b1ba3e_Gab4XcBL-1024-admin s_rally_03b1ba3e_Gab4XcBL-1024-mgr
Spine s_rally_03b1ba3e_MUHzdmLD-1017 fabric 1017   none
s_rally_03b1ba3e_MUHzdmLD-1017-admin s_rally_03b1ba3e_MUHzdmLD-1017-mgr
Spine s_rally_03b1ba3e_OHOBvTMz-1010 fabric 1010   none
s_rally_03b1ba3e_OHOBvTMz-1010-admin s_rally_03b1ba3e_OHOBvTMz-1010-mgr
Spine s_rally_03b1ba3e_PhD6ldKr-1007 fabric 1007   none
s_rally_03b1ba3e_PhD6ldKr-1007-admin s_rally_03b1ba3e_PhD6ldKr-1007-mgr
Spine s_rally_03b1ba3e_Vaqn6IEP-1027 fabric 1027   none
s_rally_03b1ba3e_Vaqn6IEP-1027-admin s_rally_03b1ba3e_Vaqn6IEP-1027-mgr
Spine s_rally_03b1ba3e_ZvrTlWGh-1018 fabric 1018   none
s_rally_03b1ba3e_ZvrTlWGh-1018-admin s_rally_03b1ba3e_ZvrTlWGh-1018-mgr

```

```
Spine s_rally_03b1ba3e_tXWtMSfG-1002 fabric 1002 none
s_rally_03b1ba3e_tXWtMSfG-1002-admin s_rally_03b1ba3e_tXWtMSfG-1002-mgr
```

Scenario 6:

Iperf Traffic throughput testing between VMs

VM Specs

2 Centos 7.0 VMs were used to test performance using Iperf

VM1: Centos70-2

VM2: Centos70

Instances

	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/>	Centos70-2	Centos	192.168.10.30	m1.xlarge	Centos7	Active	nova	None	Running	22 hours, 39 minutes	<button>Create Snapshot</button>
<input type="checkbox"/>	Centos70	Centos	192.168.10.29	m1.large	Centos7	Active	nova	None	Running	22 hours, 51 minutes	<button>Create Snapshot</button>
<input type="checkbox"/>	vm-same	TestVM	192.168.10.6	m1.micro	-	Active	nova	None	Running	1 month	<button>Create Snapshot</button>
<input type="checkbox"/>	vm-4	TestVM	192.168.60.5	m1.micro	-	Active	nova	None	Running	1 month	<button>Create Snapshot</button>
<input type="checkbox"/>	vm3	TestVM	192.168.30.5	m1.tiny	-	Active	nova	None	Running	1 month	<button>Create Snapshot</button>
<input type="checkbox"/>	vm-2	TestVM	192.168.20.5	m1.micro	-	Active	nova	None	Running	1 month	<button>Create Snapshot</button>
<input type="checkbox"/>	vm-1	TestVM	192.168.10.5	m1.micro	-	Active	nova	None	Running	1 month	<button>Create Snapshot</button>

Displaying 7 items

```
CLI (network-admin@spine) > dhcp-host-show
dhcp-name      hostname          mac                  pxe-boot gateway-ip   fixed-ip
-----+-----+-----+-----+-----+-----+-----+
net1-1004-dhcp fa-16-3e-48-de-3b fa:16:3e:48:de:3b no    192.168.10.1 192.168.10.5
net1-1004-dhcp fa-16-3e-39-a2-9b fa:16:3e:39:a2:9b no    192.168.10.1 192.168.10.6
net1-1004-dhcp fa-16-3e-ea-92-54 fa:16:3e:ea:92:54 no    192.168.10.1 192.168.10.29
net1-1004-dhcp fa-16-3e-0d-5b-2f fa:16:3e:0d:5b:2f no    192.168.10.1 192.168.10.30
net2-1030-dhcp fa-16-3e-11-11-ef fa:16:3e:11:11:ef no    192.168.20.1 192.168.20.5
net3-1001-dhcp fa-16-3e-0a-c7-f2 fa:16:3e:0a:c7:f2 no    192.168.30.1 192.168.30.5
net4-1025-dhcp fa-16-3e-7f-90-f9 fa:16:3e:7f:90:f9 no    192.168.60.1 192.168.60.5
CLI (network-admin@spine) >
```

TCP Test

Client Logs:

```
[root@centos70-2 centos]# iperf3 -s
-----
[ ID] Interval      Transfer     Bandwidth
[  5] local 192.168.10.30 port 5201 connected to 192.168.10.29 port 38224
[  5]  0.00-1.00  sec  74.9 MBytes   626 Mbits/sec
[  5]  1.00-2.00  sec  82.1 MBytes   689 Mbits/sec
[  5]  2.00-3.00  sec  90.4 MBytes   757 Mbits/sec
[  5]  3.00-4.00  sec  85.8 MBytes   728 Mbits/sec
[  5]  4.00-5.00  sec  88.9 MBytes   678 Mbits/sec
[  5]  5.00-6.00  sec  81.2 MBytes   682 Mbits/sec
[  5]  6.00-7.00  sec  89.4 MBytes   749 Mbits/sec
[  5]  7.00-8.00  sec  92.0 MBytes   773 Mbits/sec
[  5]  8.00-9.00  sec  78.4 MBytes   656 Mbits/sec
[  5]  9.00-10.00 sec  84.9 MBytes   712 Mbits/sec
[  5] 10.00-11.01 sec  87.9 MBytes   736 Mbits/sec
[  5] 11.01-12.00 sec  88.8 MBytes   746 Mbits/sec
[  5] 12.00-13.00 sec  92.2 MBytes   773 Mbits/sec
[  5] 13.00-14.00 sec  88.0 MBytes   739 Mbits/sec
[  5] 14.00-15.00 sec  86.1 MBytes   722 Mbits/sec
[  5] 15.00-16.01 sec  77.6 MBytes   650 Mbits/sec
[  5] 16.01-17.00 sec  85.9 MBytes   722 Mbits/sec
[  5] 17.00-18.01 sec  79.8 MBytes   662 Mbits/sec
[  5] 18.01-19.00 sec  87.5 MBytes   742 Mbits/sec
[  5] 19.00-20.00 sec  86.4 MBytes   725 Mbits/sec
[  5] 20.00-21.00 sec  81.3 MBytes   683 Mbits/sec
[  5] 21.00-22.00 sec  78.9 MBytes   594 Mbits/sec
[  5] 22.00-23.00 sec  88.2 MBytes   740 Mbits/sec
[  5] 23.00-24.00 sec  87.1 MBytes   731 Mbits/sec
[  5] 24.00-25.01 sec  84.2 MBytes   706 Mbits/sec
[  5] 25.01-26.00 sec  89.6 MBytes   753 Mbits/sec
[  5] 26.00-27.01 sec  81.4 MBytes   681 Mbits/sec
[  5] 27.01-28.01 sec  69.9 MBytes   586 Mbits/sec
[  5] 28.01-29.00 sec  69.5 MBytes   584 Mbits/sec
[  5] 29.00-30.00 sec  85.5 MBytes   716 Mbits/sec
[  5] 30.00-30.02 sec  0.00 Bytes  0.00 bits/sec
-----
[ ID] Interval      Transfer     Bandwidth
[  5]  0.00-30.02 sec  0.00 Bytes  0.00 bits/sec
[  5]  0.00-30.02 sec  2.45 GBytes  701 Mbits/sec
                                         sender
                                         receiver
-----
Server listening on 5201
-----
```

Server Logs:

```
Connecting to host 192.168.10.30, port 5201
[  4] local 192.168.10.29 port 38224 connected to 192.168.10.30 port 5201
[ ID] Interval      Transfer     Bandwidth      Retr  Cwnd
[  4]  0.00-1.00  sec  76.4 MBytes   640 Mbits/sec   0  332 KBytes
[  4]  1.00-2.00  sec  82.7 MBytes   693 Mbits/sec   0  342 KBytes
[  4]  2.00-3.00  sec  90.2 MBytes   757 Mbits/sec   0  342 KBytes
[  4]  3.00-4.00  sec  85.9 MBytes   721 Mbits/sec   0  351 KBytes
[  4]  4.00-5.00  sec  80.8 MBytes   678 Mbits/sec   0  356 KBytes
[  4]  5.00-6.00  sec  81.2 MBytes   682 Mbits/sec   0  362 KBytes
[  4]  6.00-7.00  sec  89.4 MBytes   749 Mbits/sec   0  362 KBytes
[  4]  7.00-8.00  sec  92.0 MBytes   772 Mbits/sec   0  369 KBytes
[  4]  8.00-9.00  sec  78.4 MBytes   658 Mbits/sec   0  369 KBytes
[  4]  9.00-10.00 sec  84.8 MBytes   712 Mbits/sec   0  375 KBytes
[  4] 10.00-11.00 sec  87.9 MBytes   737 Mbits/sec   0  375 KBytes
```

[4]	11.00-12.00	sec	88.7 MBytes	745 Mbits/sec	0	376 KBytes
[4]	12.00-13.00	sec	92.2 MBytes	774 Mbits/sec	0	378 KBytes
[4]	13.00-14.00	sec	88.1 MBytes	739 Mbits/sec	0	380 KBytes
[4]	14.00-15.00	sec	86.1 MBytes	722 Mbits/sec	0	382 KBytes
[4]	15.00-16.00	sec	77.8 MBytes	652 Mbits/sec	0	383 KBytes
[4]	16.00-17.00	sec	85.9 MBytes	721 Mbits/sec	0	385 KBytes
[4]	17.00-18.00	sec	79.7 MBytes	668 Mbits/sec	0	385 KBytes
[4]	18.00-19.00	sec	87.5 MBytes	735 Mbits/sec	0	385 KBytes
[4]	19.00-20.00	sec	86.4 MBytes	724 Mbits/sec	0	385 KBytes
[4]	20.00-21.03	sec	79.4 MBytes	648 Mbits/sec	0	386 KBytes
[4]	21.03-22.00	sec	72.8 MBytes	628 Mbits/sec	0	397 KBytes
[4]	22.00-23.00	sec	88.3 MBytes	741 Mbits/sec	0	397 KBytes
[4]	23.00-24.00	sec	87.1 MBytes	731 Mbits/sec	0	397 KBytes
[4]	24.00-25.00	sec	84.2 MBytes	707 Mbits/sec	0	397 KBytes
[4]	25.00-26.00	sec	89.7 MBytes	752 Mbits/sec	0	397 KBytes
[4]	26.00-27.00	sec	81.3 MBytes	683 Mbits/sec	0	397 KBytes
[4]	27.00-28.00	sec	70.2 MBytes	588 Mbits/sec	0	397 KBytes
[4]	28.00-29.00	sec	69.1 MBytes	581 Mbits/sec	0	397 KBytes
[4]	29.00-30.01	sec	85.5 MBytes	714 Mbits/sec	0	397 KBytes
<hr/>						
[ID]	Interval		Transfer	Bandwidth	Retr	
[4]	0.00-30.01	sec	2.45 GBytes	702 Mbits/sec	0	sender
[4]	0.00-30.01	sec	2.45 GBytes	701 Mbits/sec		

receiver

iperf Done.

UDP Test

Client Logs:

```

root@centos70-2 centos]# iperf3 -s
-----
Server listening on 5281
-----
Accepted connection from 192.168.10.29, port 38238
Warning: Unable to set socket pacing, using application pacing instead
[ 5] local 192.168.10.30 port 5201 connected to 192.168.10.29 port 43591
[ ID] Interval      Transfer     Bandwidth     Jitter      Lost/Total Datagrams
[ 5]  0.00-1.00    sec   1.09 MBytes   9.96 Mbits/sec  0.104 ms  0/139 (0%)
[ 5]  1.00-2.00    sec   1.19 MBytes   9.96 Mbits/sec  0.104 ms  0/152 (0%)
[ 5]  2.00-3.00    sec   1.20 MBytes   10.0 Mbits/sec  0.099 ms  0/153 (0%)
[ 5]  3.00-4.00    sec   1.19 MBytes   9.96 Mbits/sec  0.100 ms  0/152 (0%)
[ 5]  4.00-5.00    sec   1.20 MBytes   10.0 Mbits/sec  0.143 ms  0/153 (0%)
[ 5]  5.00-6.00    sec   1.19 MBytes   9.96 Mbits/sec  0.116 ms  0/152 (0%)
[ 5]  6.00-7.00    sec   1.20 MBytes   10.0 Mbits/sec  0.124 ms  0/153 (0%)
[ 5]  7.00-8.00    sec   1.20 MBytes   10.0 Mbits/sec  0.182 ms  0/153 (0%)
[ 5]  8.00-9.00    sec   1.19 MBytes   9.96 Mbits/sec  0.183 ms  0/152 (0%)
[ 5]  9.00-10.00   sec   1.20 MBytes   10.0 Mbits/sec  0.148 ms  0/153 (0%)
[ 5] 10.00-11.00   sec   1.20 MBytes   10.0 Mbits/sec  0.140 ms  0/153 (0%)
[ 5] 11.00-12.00   sec   1.19 MBytes   9.96 Mbits/sec  0.146 ms  0/152 (0%)
[ 5] 12.00-13.00   sec   1.20 MBytes   10.0 Mbits/sec  0.121 ms  0/153 (0%)
[ 5] 13.00-14.00   sec   1.19 MBytes   9.97 Mbits/sec  0.195 ms  0/152 (0%)
[ 5] 14.00-15.00   sec   1.20 MBytes   10.0 Mbits/sec  0.121 ms  0/153 (0%)
[ 5] 15.00-16.00   sec   1.19 MBytes   9.96 Mbits/sec  0.126 ms  0/152 (0%)
[ 5] 16.00-17.00   sec   1.20 MBytes   10.0 Mbits/sec  0.149 ms  0/153 (0%)
[ 5] 17.00-18.00   sec   1.20 MBytes   10.0 Mbits/sec  0.118 ms  0/153 (0%)
[ 5] 18.00-19.00   sec   1.19 MBytes   9.96 Mbits/sec  0.118 ms  0/152 (0%)
[ 5] 19.00-20.00   sec   1.20 MBytes   10.0 Mbits/sec  0.147 ms  0/153 (0%)
[ 5] 20.00-21.00   sec   1.20 MBytes   10.0 Mbits/sec  0.168 ms  0/153 (0%)
[ 5] 21.00-22.00   sec   1.19 MBytes   9.96 Mbits/sec  0.098 ms  0/152 (0%)
[ 5] 22.00-23.00   sec   1.20 MBytes   10.0 Mbits/sec  0.115 ms  0/153 (0%)
[ 5] 23.00-24.00   sec   1.19 MBytes   9.96 Mbits/sec  0.099 ms  0/152 (0%)
[ 5] 24.00-25.00   sec   1.20 MBytes   10.0 Mbits/sec  0.102 ms  0/153 (0%)
[ 5] 25.00-26.00   sec   1.19 MBytes   9.96 Mbits/sec  0.082 ms  0/152 (0%)
[ 5] 26.00-27.00   sec   1.20 MBytes   10.0 Mbits/sec  0.200 ms  0/153 (0%)
[ 5] 27.00-28.00   sec   1.19 MBytes   9.96 Mbits/sec  0.127 ms  0/152 (0%)
[ 5] 28.00-29.00   sec   1.20 MBytes   10.0 Mbits/sec  0.098 ms  0/153 (0%)
[ 5] 29.00-30.00   sec   1.20 MBytes   10.0 Mbits/sec  0.117 ms  0/153 (0%)
[ 5] 30.00-30.01   sec   0.00 Bytes   0.00 bits/sec  0.117 ms  0/0 (0%)
-----
[ ID] Interval      Transfer     Bandwidth     Total Datagrams
[ 5]  0.00-30.01    sec   0.00 Bytes   0.00 bits/sec  0/4564 (0%)
-----
Server listening on 5281
-----
```

Server Logs:

Connecting to host 192.168.10.30, port 5201

```

[ 4] local 192.168.10.29 port 43591 connected to 192.168.10.30 port 5201
[ ID] Interval      Transfer     Bandwidth     Total Datagrams
[ 4]  0.00-1.00    sec   1.09 MBytes   9.10 Mbits/sec  139
[ 4]  1.00-2.00    sec   1.19 MBytes   9.96 Mbits/sec  152
[ 4]  2.00-3.00    sec   1.20 MBytes   10.0 Mbits/sec  153
[ 4]  3.00-4.00    sec   1.19 MBytes   9.96 Mbits/sec  152
[ 4]  4.00-5.00    sec   1.20 MBytes   10.0 Mbits/sec  153
[ 4]  5.00-6.00    sec   1.19 MBytes   9.96 Mbits/sec  152
[ 4]  6.00-7.00    sec   1.20 MBytes   10.0 Mbits/sec  153
[ 4]  7.00-8.00    sec   1.20 MBytes   10.0 Mbits/sec  153
[ 4]  8.00-9.00    sec   1.19 MBytes   9.96 Mbits/sec  152
[ 4]  9.00-10.00   sec   1.20 MBytes   10.0 Mbits/sec  153
[ 4] 10.00-11.00   sec   1.20 MBytes   10.0 Mbits/sec  153
[ 4] 11.00-12.00   sec   1.19 MBytes   9.96 Mbits/sec  152
[ 4] 12.00-13.00   sec   1.20 MBytes   10.0 Mbits/sec  153
[ 4] 13.00-14.00   sec   1.19 MBytes   9.96 Mbits/sec  152
```

```
[ 4] 14.00-15.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 15.00-16.00  sec  1.19 MBytes  9.96 Mbits/sec  152
[ 4] 16.00-17.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 17.00-18.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 18.00-19.00  sec  1.19 MBytes  9.96 Mbits/sec  152
[ 4] 19.00-20.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 20.00-21.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 21.00-22.00  sec  1.19 MBytes  9.96 Mbits/sec  152
[ 4] 22.00-23.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 23.00-24.00  sec  1.19 MBytes  9.96 Mbits/sec  152
[ 4] 24.00-25.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 25.00-26.00  sec  1.19 MBytes  9.96 Mbits/sec  152
[ 4] 26.00-27.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 27.00-28.00  sec  1.19 MBytes  9.96 Mbits/sec  152
[ 4] 28.00-29.00  sec  1.20 MBytes  10.0 Mbits/sec  153
[ 4] 29.00-30.00  sec  1.20 MBytes  10.0 Mbits/sec  153
- - - - -
[ ID] Interval          Transfer       Bandwidth      Jitter       Lost/Total
Datagrams
[ 4] 0.00-30.00  sec  35.7 MBytes  9.97 Mbits/sec  0.117 ms  0/4564 (0%)
[ 4] Sent 4564 datagrams
```

iperf Done.