
Fuel NSX-T plugin documentation

Release 1.0-1.0.0-1

Mirantis Inc.

Nov 22, 2016

CONTENTS

1	Documentation contents	2
1.1	Installation	2
1.2	OpenStack environment notes	2
1.3	Configuration	4
1.4	Limitations	7
1.5	Usage	8
1.6	Troubleshooting	9
1.7	Release notes	10
1.8	How to build the plugin from source	10

The Fuel NSX-T plugin allows end user to use VMware NSX Transformers SDN as network backend for Neutron.

The plugin supports VMware NSX Transformers version 1.0.0 and 1.0.1.

The pre-built package of the plugin is in [Fuel Plugin Catalog](#).

DOCUMENTATION CONTENTS

Installation

1. Download the plugin .rpm package from the [Fuel plugin catalog](#).
2. Upload the package to the Fuel master node.
3. Install the plugin with the `fuel` command-line tool:

```
[root@nailgun ~] fuel plugins --install nsx-t-1.0-1.0.0-1.noarch.rpm
```

4. Verify that the plugin installation is successful:

```
[root@nailgun ~] fuel plugins list
id | name | version | package_version | releases
---+-----+-----+-----+-----
6  | nsx-t | 1.0.0   | 4.0.0           | ubuntu (mitaka-9.0)
```

After the installation, the plugin can be used on new OpenStack clusters; you cannot enable the plugin on the deployed clusters.

Uninstallation

Before uninstalling the plugin, ensure there no environments left that use the plugin, otherwise the uninstallation is not possible.

To uninstall the plugin, run following:

```
[root@nailgun ~] fuel plugins --remove nsx-t==1.0.0
```

OpenStack environment notes

Environment configuration

The Fuel NSX-T plugin cannot deploy NSX Transformers.

Before you start OpenStack deployment, verify that your VMware NSX Transformers is configured and functions properly:

- VLAN transport zone must present
- Overlay transport zone must present

- tier0 router must be created
- uplink profile for OpenStack nodes must be created
- NSX edge cluster must be formed
- IP address pool for OpenStack controllers and compute nodes must exist

To use the NSX-T plugin, create a new OpenStack environment using the Fuel web UI by doing the following:

1. On the *Networking setup* configuration step, select *Neutron with NSX-T plugin* radio button

Create a new OpenStack environment ×

<p>Name and Release</p> <p>Compute</p> <p>Networking Setup</p> <p>Storage Backends</p> <p>Additional Services</p> <p>Finish</p>	<p><input checked="" type="radio"/> Neutron with NSX-T plugin ✓ NSX Transformers uses STT tunneling protocol. NSX must be up and running before OpenStack deployment!</p> <p><input type="radio"/> Neutron with ML2 plugin ✓ Framework that enables simultaneous utilization of the layer 2 networking technologies through drivers.</p> <p><input type="checkbox"/> Neutron with VLAN segmentation Your network hardware must be configured for VLAN segmentation. This option supports up to 4095 networks.</p> <p><input type="checkbox"/> Neutron with tunneling segmentation By default VXLAN tunnels will be used. This option supports millions of tenant data networks.</p>
--	---

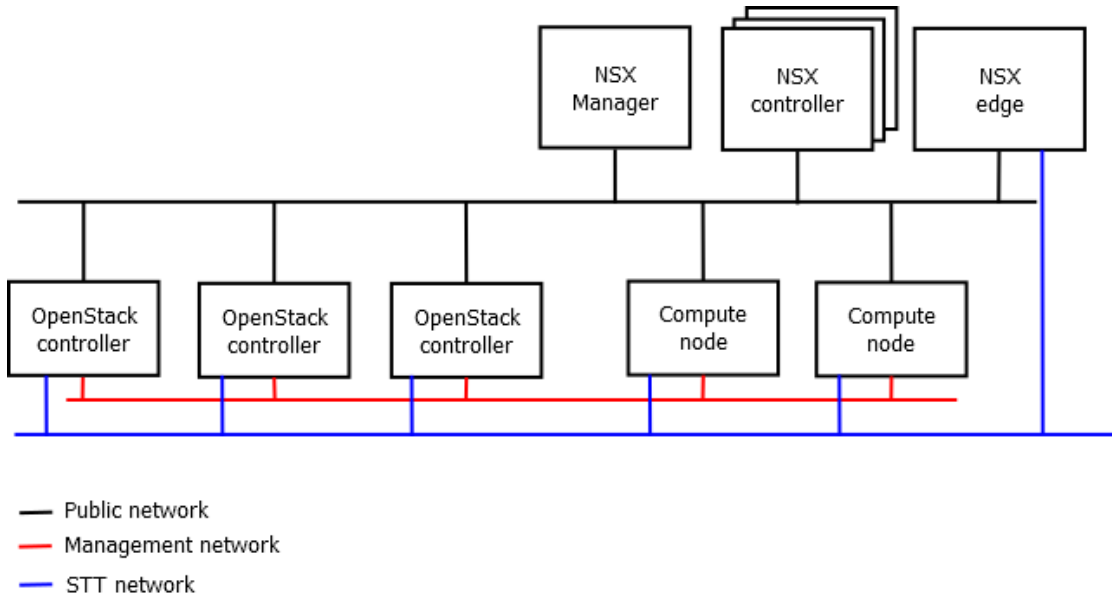
Cancel
← Prev
Next →

ESXi hosts that participate in vCenter cluster that is used on the VMware tab must be manually added as transport nodes in NSX Manager. Hosts must be added prior OpenStack deployment.

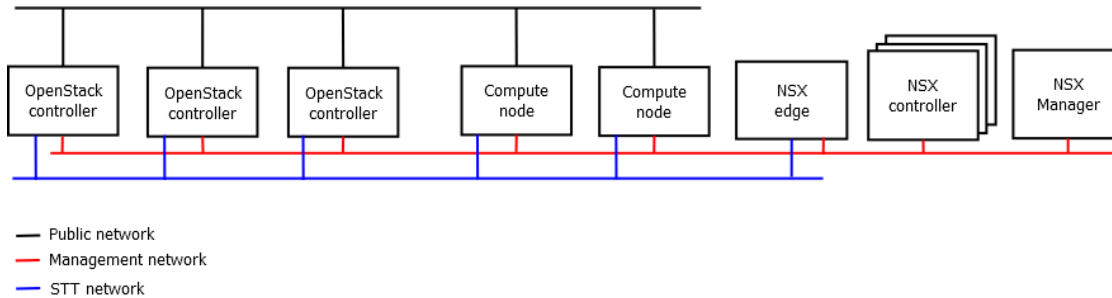
Network setup

Pay attention to on which interface you assign the *Public* network. The OpenStack controllers must have connectivity with the NSX Manager host through the *Public* network since the *Public* network is the default route for packets.

If NSX Manager and NSX Controllers are going to communicate with OpenStack controllers and computes through Public network then setting *Assign public network to all nodes* (Networks tab -> Settings/Other label) must be enabled. Otherwise compute node will be communicating with NSX Manager through controller that perform NAT which will hide compute node IP addresses and will prevent them to register in NSX management plane.



Another way is to locate NSX nodes in OpenStack management network. In this setup there is no need to assign public network to all nodes, because OpenStack and NSX nodes has L2 connectivity and no NAT is performed. OpenStack controllers and computes will still use Public network as default route.

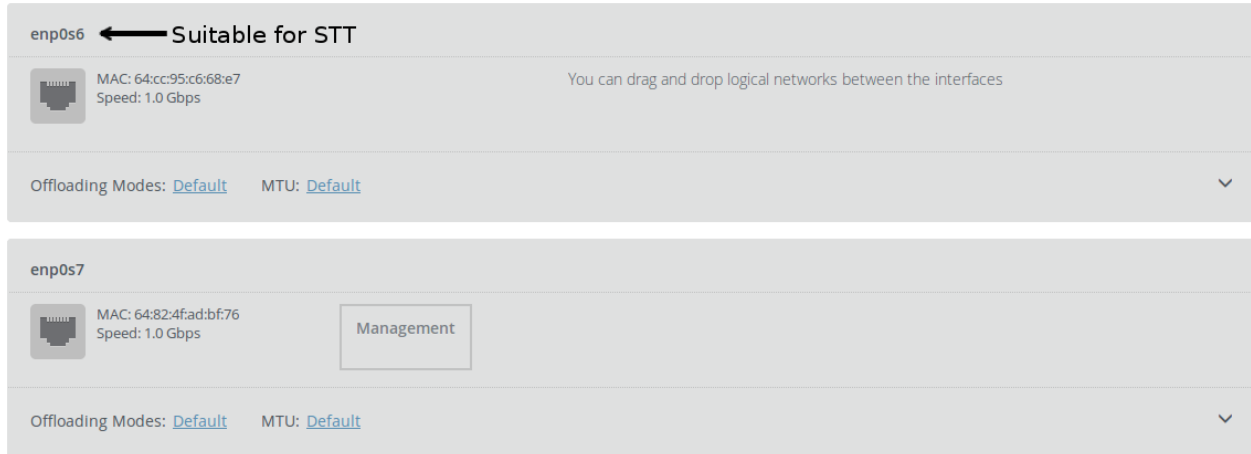


During the deployment, the plugin creates a simple network topology for the admin tenant. The plugin creates a provider network which connects the tenants with the transport (physical) network: one internal network and a router that is connected to both networks.

Configuration

Node interfaces for overlay traffic

NSX Transformers uses STT protocol to carry virtual machines traffic. Plugin requires that interfaces which are going to be used for STT traffic must not carry any other traffic (PXE, storage, openstack management).



Switch to the *Networks* tab of the Fuel web UI and click the *Settings/Other* label. The plugin checkbox is enabled by default. The screenshot below shows only the settings in focus:

NSX Transformers plugin

You have just activated NSX Transformers plugin. The developer of this environment may stop functioning properly if this plugin is applied in developer on further appropriate course of actions.

Versions 1.0.0 Bypass NSX Manager certificate verification

NSX Manager	<input type="text" value="172.16.0.249"/>
User	<input type="text" value="admin"/>
Password	<input type="password" value="••••••••"/> <input type="checkbox"/>
Overlay transport zone ID	<input type="text" value="0eeb1b85-c826-403d-8762-6a9c23a4f1"/>
VLAN transport zone ID	<input type="text" value="8efe20d2-e71a-4d6e-acdd-f78a2ec2e9"/>
Tier-0 router ID	<input type="text" value="ed9d5f1d-8524-4ded-9cdc-59eb2a008"/>
Edge cluster	<input type="text" value="c53d602a-4010-47cc-a8b1-4ef11d0a3e"/>
Uplink profile ID	<input type="text" value="99864272-b34f-46a5-89c8-5657fa7042"/>
IP pool ID for controller VTEPs	<input type="text" value="2e06fcb2-7c5b-4515-a7a9-98809c7b8e"/>
Colon separated pnics pairs for controller nodes	<input type="text" value="enp0s6:uplink"/>
IP pool ID for compute VTEPs	<input type="text" value="2e06fcb2-7c5b-4515-a7a9-98809c7b8e"/>

The plugin contains the following settings:

1. Bypass NSX Manager certificate verification – if enabled, the HTTPS connection to NSX Manager is not verified. Otherwise, the two following options are available:
 - The setting “CA certificate file” appears below making it possible to upload a CA certificate that issued the NSX Manager certificate.
 - With no CA certificate provided, the NSX Manager certificate is verified against the CA certificate bundle that comes with Ubuntu 14.04.
2. NSX Manager – IP address or hostname, multiple values can be separated by comma. If you are going to use hostname in this textbox, ensure that your OpenStack controller can resolve the hostname. Add necessary DNS servers in the *Host OS DNS Servers* section.

OpenStack Controller must have L3 connectivity with NSX Manager through the Public network which is used as default route.

3. Overlay transport zone ID – UUID of overlay (STT) transport zone which must be pre-created in NSX Manager.
4. VLAN transport zone ID — UUID of transport zone which represents network that connects virtual machines with physical infrastructure.
5. Tier-0 router ID – UUID of tier0 router that must exist in NSX Transformers.
6. Edge cluster – UUID of NSX edge nodes cluster that must be installed and configured.
7. Uplink profile ID – UUID of uplink profile which specifies STT interfaces configuration (e.g. MTU value).
8. IP pool ID for controller VTEPs – UUID of IP pool that will be assigned to controllers STT interfaces.
9. Colon separated pnics pairs for controller nodes – this field sets correspondence between physical NIC and uplink name that is used in “Uplink profile ID” on controller nodes, e.g. `enp0s1:uplink`. Each pair must be one separate line.

Warning: Uplink name must exactly match value of “Active uplink” or “Standby uplink” in uplink switch profile that was specified above.

10. IP pool ID for compute VTEPs – UUID of IP pool that will be assigned to STT interfaces of compute nodes.
11. Colon separated pnics pairs for compute nodes – this fields sets correspondence between physical NIC and uplink name that is used in “Uplink profile ID” on compute nodes, e.g. `“enp0s1:uplink”`. Each pair must be one separate line.
12. Floating IP ranges – dash-separated IP addresses allocation pool for external network, e.g. `“172.16.212.2-172.16.212.40”`.
13. External network CIDR – network in CIDR notation that includes floating IP ranges.
14. Gateway – default gateway for the external network; if not defined, the first IP address of the network is used.
15. Internal network CIDR – network in CIDR notation for use as internal.
16. DNS for internal network – comma-separated IP addresses of DNS server for internal network.

Limitations

Fuel NSX-T plugin does not support SSL verification

Plugin does not allow user to enable SSL certificate verification during deployment. As workaround it is possible to store certificate at post deployment stage and enable it in Neutron NSX configuration file (`/etc/neutron/plugins/vmware/nsx.ini`) (Neutron NSX plugin supports SSL certificate).

NSX-T plugin cannot be used simultaneously with NSXv plugin

Since both plugins provide the same network component called `network:neutron:core:nsx` it is not possible to use both plugins together.

The plugin is not hotpluggable

It is not possible to enable plugin on already existing OpenStack.

Ubuntu cloud archive distribution is not supported

Fuel 9.0 provides two options for OpenStack release. One is plain Ubuntu distribution, another one includes Ubuntu cloud archive. The plugin does not support Ubuntu cloud archive packages.

Ironic service is not supported

Ironic service requires control of top of rack switches and can not be used with NSX Transformers.

OpenStack environment reset/deletion

The Fuel NSX-T plugin does not provide a cleanup mechanism when an OpenStack environment is reset or deleted. All registered transport nodes, logical switches and routers remain intact, it is up to the operator to delete them and free resources.

Usage

The easiest way to check that the plugin works as expected is to create a network or router using the `neutron` command-line tool:

```
[root@nailgun ~]# ssh node-4      # node-4 is a controller node
root@node-4:~# . openrc
root@node-4:~# neutron router-create r1
```

You can monitor the plugin actions in `/var/log/neutron/server.log` and see how edges appear in the list of the Networking & Security -> NSX Edges pane in vSphere Web Client. If you see error messages, check the *Troubleshooting* section.

STT MTU considerations

NSX Transformers uses STT protocol to encapsulate VM traffic. The protocol adds additional data to the packet. Consider increasing MTU on the network equipment connected to hosts that will emit STT traffic.

Consider the following calculation:

Outer IPv4 header == 20 bytes

Outer TCP header == 24 bytes

STT header == 18 bytes

Inner Ethernet frame == 1518 (14 bytes header, 4 bytes 802.1q header, 1500 Payload)

Summarizing all of these we get 1580 bytes. Consider increasing MTU on the network hardware up to 1600 bytes.

Troubleshooting

Neutron NSX plugin issues

The Neutron NSX-T plugin does not have a separate log file, its messages are logged by the neutron server. The default log file on OpenStack controllers for neutron server is `/var/log/neutron/server.log`

Inability to resolve NSX Manager hostname

If you see following message:

```
2016-10-18 ... INFO vmware_nsx.plugins.nsx_v3.plugin [-] Starting NsxV3Plugin
2016-10-18 ... INFO vmware_nsx.nsxlib.v3.cluster [-] Endpoint 'https://nsxmanager.
↳mydom.org'
    changing from state 'INITIALIZED' to 'DOWN'
2016-10-18 ... WARNING vmware_nsx.nsxlib.v3.cluster [-] Failed to validate API_
↳cluster endpoint
    '[DOWN] https://nsxmanager.mydom.org' due to: HTTPSConnectionPool(host='nsxmanager.
↳mydom.org',
    port=443): Max retries exceeded with url: /a..nes (Caused by NewConnectionError(
    '<requests.packages.urllib3.connection.VerifiedHTTPSConnection object at_
↳0x7ff69b3c4b90>:
    Failed to establish a new connection: [Errno -2] Name or service not known',))
2016-10-18 ... ERROR neutron.service [-] Unrecoverable error: please check log for_
↳details.
2016-10-18 ... ERROR neutron.service Traceback (most recent call last):
...
2016-10-18 ... ERROR neutron.service ServiceClusterUnavailable: Service cluster:
    'https://nsxmanager.mydom.org' is unavailable. Please, check NSX setup and/or_
↳configuration
```

It means that the controller cannot resolve the NSX Manager hostname (`nsxmanager.mydom.org` in this example) that is specified in the configuration file. Check that the DNS server IP addresses that you specified in the *Host OS DNS Servers* section of the Fuel web UI are correct and reachable by all controllers; pay attention to that the default route for controllers is *Public* network. Also, verify that the hostname that you entered is correct by trying to resolve it via the `host` or `dig` programs.

SSL/TLS certificate problems

```
2016-10-28 12:32:26.086 2832 INFO vmware_nsx.nsxlib.v3.cluster [-] Endpoint
    'https://172.16.0.249' changing from state 'INITIALIZED' to 'DOWN'
2016-10-28 12:32:26.087 2832 WARNING vmware_nsx.nsxlib.v3.cluster [-] Failed to
    validate API cluster endpoint '[DOWN] https://172.16.0.249' due to: [Errno 1]
    _ssl.c:510: error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate_
↳verify failed
```

This error indicates that you enabled the SSL/TLS certificate verification, but the certificate verification failed during connection to NSX Manager. The possible causes are:

1. NSX Manager certificate expired. Log into NSX Manager web GUI and check certificate validation dates.
2. Check if the certification authority (CA) certificate is still valid. The CA certificate is specified by `ca_file` directive in `nsx.ini` (usually `/etc/neutron/plugins/vmware/nsx.ini` on controller node).

User access problems

```
2016-10-28 12:28:20.060 18259 INFO vmware_nsx.plugins.nsx_v3.plugin [-] Starting
↳NsxV3Plugin
2016-10-28 12:28:20.201 18259 WARNING vmware_nsx.nsxlib.v3.client [-] The HTTP
↳request returned error code 403,
  whereas 200 response codes were expected. Response body {u'module_name': u'common-
↳service',
  u'error_message': u'The username/password combination is incorrect or the account
↳specified has been locked.', u'error_code': u'98'}
2016-10-28 12:28:20.202 18259 INFO vmware_nsx.nsxlib.v3.cluster [-] Endpoint 'https://
↳172.16.0.249' changing
  from state 'INITIALIZED' to 'DOWN'
2016-10-28 12:28:20.203 18259 WARNING vmware_nsx.nsxlib.v3.cluster [-] Failed to
↳validate API cluster endpoint
  '[DOWN] https://172.16.0.249' due to: Unexpected error from backend manager (['172.
↳16.0.249']) for GET https://172.16.0.249/api/
  vl/transport-zones : The username/password combination is incorrect or the account
↳specified has been locked.
```

Verify that username and password that are entered on the plugins pane are correct.

Release notes

Release notes for Fuel NSX-T plugin 1.0.0:

- Plugin is compatible with Fuel 9.0 and 9.1.

How to build the plugin from source

To build the plugin, you first need to install `fuel-plugin-builder 4.1.0`

```
$ pip install fuel-plugin-builder==4.1.0
```

Build the plugin:

```
$ git clone https://git.openstack.org/openstack/fuel-plugin-nsx-t
$ cd fuel-plugin-nsx-t/
```

The `librarian-puppet` ruby package is required to be installed. It is used to fetch upstream `fuel-library` puppet modules that the plugin uses. It can be installed via the `gem` package manager:

```
$ gem install librarian-puppet
```

or if you are using ubuntu linux, you can install it from the repository:

```
$ apt-get install librarian-puppet
```

and build the plugin:

```
$ fpb --build .
```

fuel-plugin-builder will produce an .rpm package of the plugin which you need to upload to the Fuel master node:

```
$ ls nsx*.rpm  
nsx-t-1.0-1.0.0-1.noarch.rpm
```