

NSX-T Container Plug-in for OpenShift - Installation and Administration Guide

Modified on 10 APR 2018

VMware NSX-T 2.1.0.1, 2.1.2

VMware NSX-T 2.1



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2017, 2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

NSX-T Container Plug-in for OpenShift - Installation and Administration Guide 4

- 1 Overview of NSX-T Container Plug-in 5**
 - Compatibility Requirements 5
 - Installation Overview 6

- 2 Setting Up NSX-T Resources 7**
 - Configuring NSX-T Resources 7
 - Create and Configure a Tier-0 Logical Router 10

- 3 Setting Up NSX-T Container Plug-in and OpenShift 12**
 - Deploy OpenShift VMs 12
 - Prepare the Ansible Hosts File 12
 - Install CNI Plug-in and OVS 14
 - Install OpenShift Container Platform 16
 - Run NCP and NSX Node Agent 17
 - Setup Notes 18

- 4 Administering NSX-T Container Plug-in 21**
 - Manage IP Blocks from the NSX Manager GUI 21
 - Manage IP Block Subnets from the NSX Manager GUI 22
 - CIF-Attached Logical Ports 22
 - CLI Commands 23

NSX-T Container Plug-in for OpenShift - Installation and Administration Guide

This guide describes how to install and administer NSX-T Container Plug-in (NCP) to provide integration between NSX-T and OpenShift.

Intended Audience

This guide is intended for system and network administrators. A familiarity with the installation and administration of NSX-T and OpenShift is assumed.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Overview of NSX-T Container Plug-in



NSX-T Container Plug-in (NCP) provides integration between NSX-T and container orchestrators such as Kubernetes, as well as integration between NSX-T and container-based PaaS (platform as a service) software products such as OpenShift. This guide describes setting up NCP with OpenShift.

The main component of NCP runs in a container and communicates with NSX Manager and with the OpenShift control plane. NCP monitors changes to containers and other resources and manages networking resources such as logical ports, switches, routers, and security groups for the containers by calling the NSX API.

The NSX CNI plug-in runs on each OpenShift node. It monitors container life cycle events, connects a container interface to the guest vSwitch, and programs the guest vSwitch to tag and forward container traffic between the container interfaces and the VNIC.

In this release, NCP supports a single OpenShift cluster.

This chapter includes the following topics:

- [Compatibility Requirements](#)
- [Installation Overview](#)

Compatibility Requirements

NSX-T Container Plug-in has the following compatibility requirements.

Software Product	Version
Hypervisor for Container Host VMs	<ul style="list-style-type: none">▪ ESXi 6.5▪ ESXi 6.5 Update 1▪ RHEL KVM 7.3, 7.4
Container Host Operating System	RHEL 7.3, 7.4
Platform as a Service	OpenShift 3.7 (NCP 2.1, 2.1.0.1) OpenShift 3.7, 3.9 (NCP 2.1.2)
Guest vSwitch	OVS 2.6, 2.7, 2.8

Installation Overview

Installing and configuring NCP involves the following steps. To perform the steps successfully, you must be familiar with NSX-T and OpenShift installation and administration.

- 1 Install NSX-T.
- 2 Create an overlay transport zone.
- 3 Create an overlay logical switch and connect the nodes to the switch.
- 4 Create a tier-0 logical router.
- 5 Create IP blocks for the pods.
- 6 Create IP blocks or IP pools for SNAT (source network address translation).
- 7 Deploy OpenShift VMs.
- 8 Prepare the Ansible hosts file.
- 9 Install CNI plug-in and OVS (Open vSwitch) on each node.
- 10 Install OpenShift Container Platform.
- 11 Run NCP and NSX node agent.

Setting Up NSX-T Resources

NSX-T resources must be created to provide networking to OpenShift nodes. You can configure these resources manually using the NSX Manager GUI, or automate the process using an Ansible playbook.

This section describes creating the NSX-T resources manually. To automate the process, see [Install CNI Plug-in and OVS](#).

This chapter includes the following topics:

- [Configuring NSX-T Resources](#)
- [Create and Configure a Tier-0 Logical Router](#)

Configuring NSX-T Resources

NSX-T resources that you need to configure include an overlay transport zone, a tier-0 logical router, a logical switch to connect the node VMs, IP blocks for Kubernetes nodes, and an IP pool for SNAT.

In NCP 2.1 and 2.1.0.1, you must use tags to configure NSX-T resources. Starting with NCP 2.1.2, you must configure NSX-T resources using UUIDs or names in the configuration file `ncp.ini`.

Overlay Transport Zone

Log in to NSX Manager and navigate to **Fabric > Transport Zones**. Find the overlay transport zone that is used for container networking or create a new one.

(NCP 2.1 and 2.1.0.1) Tag the transport zone with one or more `{'ncp/cluster': '<cluster_name>'}` tags to make it specific to one or more clusters, or tag it with `{'ncp/shared_resource': 'true'}` to make it shared by clusters. If the tag `ncp/cluster` exists, the cluster that is named in the tag will only use the resource with the `ncp/cluster` tag and will not use any resource with the `ncp/shared_resource` tag. The cluster name must match the value of the `cluster` option in the `[coe]` section in `ncp.ini`.

(NCP 2.1.2) Specify an overlay transport zone for a cluster by setting the `overlay_tz` option in the `[nsx_v3]` section of `ncp.ini`.

Tier-0 Logical Routing

Log in to NSX Manager and navigate to **Routing > Routers**. Find the router that is used for container networking or create a new one.

(NCP 2.1 and 2.1.0.1) Tag the router with one or more `{'ncp/cluster': '<cluster_name>'}` tags to make it specific to one or more clusters, or tag it with `{'ncp/shared_resource': 'true'}` to make it shared by clusters. If the tag `ncp/cluster` exists, the cluster that is named in the tag will only use the resource with the `ncp/cluster` tag and will not use any resource with the `ncp/shared_resource` tag. The cluster name must match the value of the `cluster` option in the `[coe]` section in `ncp.ini`.

(NCP 2.1.2) Specify a tier-0 logical router for a cluster by setting the `tier0_router` option in the `[nsx_v3]` section of `ncp.ini`.

Note The router must be created in active-standby mode.

Logical Switch

The vNICs used by the node for data traffic must be connected to an overlay logical switch. It is not mandatory for the node's management interface to be connected to NSX-T, although doing so will make setting up easier. You can create a logical switch by logging in to NSX Manager and navigating to **Switching > Switches**. On the switch, create logical ports and attach the node vNICs to them. The logical ports must be tagged with `{'ncp/cluster': '<cluster_name>'}` and `{'ncp/node_name': '<node_name>'}`. The `<cluster_name>` value must match the value of the `cluster` option in the `[coe]` section in `ncp.ini`.

IP Blocks for Kubernetes Pods

Log in to NSX Manager and navigate to **DDI > IPAM** to create one or more IP blocks. Specify the IP block in CIDR format.

(NCP 2.1 and 2.1.0.1) Tag the IP block with one or more `{'ncp/cluster': '<cluster_name>'}` tags to make it specific to one or more clusters, or tag it with `{'ncp/shared_resource': 'true'}` to make it shared by clusters. If the tag `ncp/cluster` exists, the cluster that is named in the tag will only use the resource with the `ncp/cluster` tag and will not use any resource with the `ncp/shared_resource` tag. The cluster name must match the value of the `cluster` option in the `[coe]` section in `ncp.ini`.

(NCP 2.1.2) Specify IP blocks for Kubernetes pods by setting the `container_ip_blocks` option in the `[nsx_v3]` section of `ncp.ini`.

You can also create IP blocks specifically for no-SNAT namespaces.

(NCP 2.1 and 2.1.0.1) Tag the IP blocks with `{'ncp/no_snat': 'true'}` in addition to the `ncp/cluster` tag. If you are upgrading from a previous release and have the tag `{'ncp/no_snat': '<cluster_name>'}`, you must manually change it to `{'ncp/no_snat': 'true'}`.

(NCP 2.1.2) Specify no-SNAT IP blocks by setting the `no_snat_ip_blocks` option in the `[nsx_v3]` section of `ncp.ini`.

If you create no-SNAT IP blocks while NCP is running, you must restart NCP. Otherwise, NCP will keep using the shared IP blocks until they are exhausted.

Note When you create an IP block, the prefix must not be larger than the value of the parameter `subnet_prefix` in NCP's configuration file `ncp.ini`.

IP Pool for SNAT

The IP pool is used for allocating IP addresses which will be used for translating pod IPs via SNAT rules, and for exposing ingress controllers via SNAT/DNAT rules, just like Openstack floating IPs. These IP addresses are also referred to as external IPs.

Multiple Kubernetes clusters use the same external IP pool. Each NCP instance uses a subset of this pool for the Kubernetes cluster that it manages. By default, the same subnet prefix for pod subnets will be used. To use a different subnet size, update the `external_subnet_prefix` option in the `[nsx_v3]` section in `ncp.ini`.

Log in to NSX Manager and navigate to **Inventory > Groups > IP POOL** to create a pool or find an existing pool.

(NCP 2.1 and 2.1.0.1) Tag the pool with `{'ncp/external': 'true'}`. In addition, tag the pool with one or more `{'ncp/cluster': '<cluster_name>'}` tags to make it specific to one or more clusters, or tag it with `{'ncp/shared_resource': 'true'}` to make it shared by clusters. If the tag `ncp/cluster` exists, the cluster that is named in the tag will only use the resource with the `ncp/cluster` tag and will not use any resource with the `ncp/shared_resource` tag. The cluster name must match the value of the `cluster` option in the `[coe]` section in `ncp.ini`.

(NCP 2.1.2) Specify IP pools for SNAT by setting the `external_ip_pools` option in the `[nsx_v3]` section of `ncp.ini`.

You can also configure SNAT for a specific service by adding an annotation to the service. For example,

```
apiVersion: v1
kind: Service
metadata:
  name: svc-example
  annotations:
    ncp/snat_pool: <external IP pool ID or name>
  selector:
    app: example
...
```

NCP will configure the SNAT rule for this service. The rule's source IP is the set of backend pods. The destination IP is the SNAT IP allocated from the specified external IP pool. Note the following:

- The pool specified by `ncp/snat_pool` should already exist in NSX-T before the service is configured.
- In NSX-T, the priority of the SNAT rule for the service is higher than that for the project.
- If a pod is configured with multiple SNAT rules, only one will work.

(Optional) Firewall Marker Sections

To allow the administrator to create firewall rules and not have them interfere with NCP-created firewall sections based on network policies, log in to NSX Manager, navigate to **Firewall > General** and create two firewall sections.

(NCP 2.1 and 2.1.0.1) Tag one section with `{'ncp/fw_sect_marker': 'bottom'}` and the other tagged with `{'ncp/fw_sect_marker': 'top'}`.

(NCP 2.1.2) Specify marker firewall sections by setting the `bottom_firewall_section_marker` and `top_firewall_section_marker` options in the `[nsx_v3]` section of `ncp.ini`.

The bottom firewall section must be below the top firewall section. With these firewall sections created, all firewall sections created by NCP for isolation will be created above the bottom firewall section, and all firewall sections created by NCP for policy will be created below the top firewall section. If these marker sections are not created, all isolation rules will be created at the bottom, and all policy sections will be created at the top. Multiple marker firewall sections with the same value per cluster are not supported and will cause an error.

Create and Configure a Tier-0 Logical Router

The tier-0 logical router connects the Kubernetes nodes to external networks.

Procedure

- 1 From a browser, log in to NSX Manager at `https://nsx-manager-ip-address`.
- 2 Navigate to **Routing > Routers** and click **Add > Tier-0 Router**.
- 3 Enter a name and optionally a description.
- 4 Select an existing edge cluster from the drop-down menu to back this tier-0 logical router.
- 5 Select a high-availability mode.
Select active-standby.
- 6 Click **Save**.
The new logical router appears as a link.
- 7 Click the logical router link.
- 8 Click **Routing > Route Redistribution**.
- 9 Click **Add** to add a new redistribution criterion.
For sources, in a routed (non-NAT) topology, select **NSX Static**. In a NAT topology, select **Tier-0 NAT**.
- 10 Click **Save**.
- 11 Click the newly created router.
- 12 Click **Configuration > Router Ports**

- 13 Click **Add** to add an uplink port.
- 14 Select a transport node.
- 15 Select the logical switch that was previously created.
- 16 Specify an IP address in your external network.
- 17 Click **Save**.

The new logical router appears as a link.

Setting Up NSX-T Container Plug-in and OpenShift

3

This chapter describes installing and configuring NSX-T Container Plug-in (NCP) and OpenShift.

This chapter includes the following topics:

- [Deploy OpenShift VMs](#)
- [Prepare the Ansible Hosts File](#)
- [Install CNI Plug-in and OVS](#)
- [Install OpenShift Container Platform](#)
- [Run NCP and NSX Node Agent](#)
- [Setup Notes](#)

Deploy OpenShift VMs

Before installing NSX-T Container Plug-in, OpenShift must be installed. You must deploy at least one master.

For more information, see <https://docs.openshift.com>.

What to do next

Prepare the Ansible hosts file. See [Prepare the Ansible Hosts File](#).

Prepare the Ansible Hosts File

The Ansible hosts file defines the nodes in the OpenShift cluster.

Procedure

- 1 Clone the NCP GitHub repository at <https://github.com/vmware/nsx-integration-for-openshift>. The hosts file is in the `openshift-ansible-nsx` directory.
- 2 In the `[masters]` and `[nodes]` sections, specify the host names and IP addresses of the OpenShift VMs. For example,

```
[masters]
admin.rhel.osmaster ansible_ssh_host=101.101.101.4

[single_master]
```

```

admin.rhel.osmaster ansible_ssh_host=101.101.101.4

[nodes]
admin.rhel.osmaster ansible_ssh_host=101.101.101.4 openshift_ip=101.101.101.4
openshift_schedulable=true openshift_hostname=admin.rhel.osmaster
admin.rhel.osnode ansible_ssh_host=101.101.101.5 openshift_ip=101.101.101.5
openshift_hostname=admin.rhel.osnode

[etcd]

[OSEv3:children]
masters
nodes
etcd

```

Note that `openshift_ip` identifies the cluster internal IP and needs to be set if the interface to be used is not the default one. The `single_master` variable is used by `ncp`-related roles from a master node to perform certain tasks only once, e.g. NSX-T management plane resource configuration.

- 3 Set up SSH access so that all the nodes can be accessed without password from the node where the Ansible role is run (typically it is the master node):

```

ssh-keygen
ssh-copy-id -i ~/.ssh/id_rsa.pub root@admin.rhel.osnode

```

- 4 Update the `[OSEv3:vars]` section. Details about all the parameters can be found in the OpenShift Container Platform Documentation for the Advanced Installation (search for "advanced installation" in <https://docs.openshift.com>). For example,

```

# Set the default route fqdn
openshift_master_default_subdomain=apps.corp.local

os_sdn_network_plugin_name=cni
openshift_use_openshift_sdn=false
openshift_node_sdn_mtu=1500

# If ansible_ssh_user is not root, ansible_become must be set to true
ansible_become=true

openshift_master_default_subdomain
  This is the default subdomain used in the OpenShift routes for External LB

os_sdn_network_plugin_name
  Set to 'cni' for the NSX Integration

openshift_use_openshift_sdn
  Set to false to disable the built-in OpenShift SDN solution

openshift_hosted_manage_router
  Set to false to disable creation of router during installation. The router has to be
  manually started after NCP and nsx-node-agent are running.

openshift_hosted_manage_registry

```

Set to false to disable creation of registry during installation. The registry has to be manually started after NCP and nsx-node-agent are running.

```

deployment_type
  Set to openshift-enterprise

openshift_hosted_manage_registry
  Set to false to disable auto creation of registry

openshift_hosted_manage_router
  Set to false to disable auto creation of router

openshift_enable_service_catalog
  Set to false to disable service_catalog

(For OpenShift 3.9 only) skip_sanity_checks
  Set to true

(For OpenShift 3.9 only) openshift_web_console_install
  Set to false

```

5 Check that you have connectivity to all hosts:

```
ansible OSEv3 -i /PATH/TO/HOSTS/hosts -m ping
```

The results should look like the following. If not, resolve the connectivity problem.

```

openshift-node1 | SUCCESS => {
  "changed": false,
  "ping": "pong"
}
openshift-master | SUCCESS => {
  "changed": false,
  "ping": "pong"
}

```

What to do next

Install CNI plug-in and OVS. See [Install CNI Plug-in and OVS](#).

Install CNI Plug-in and OVS

The container network interface (CNI) plug-in and Open vSwitch (OVS) must be installed on the OpenShift nodes. The installation is performed by running an Ansible playbook.

The playbook contains instructions to configure NSX-T resources for the nodes. You can also configure the NSX-T resources manually as described in [Chapter 2 Setting Up NSX-T Resources](#). The parameter `perform_nsx_config` indicates whether or not to configure the resources when the playbook is run.

Procedure

- 1 Update the parameter values in `roles/ncp_prep/default/main.yaml` and `roles/nsx_config/default/main.yaml`, including the URLs where CNI plugin RPM, OVS and its corresponding kernel module RPM can be downloaded. In addition, `uplink_port` is the name of the uplink port VNIC on the node VM. The remaining variables pertain to the NSX-T management plane configuration, including:
 - `perform_nsx_config`: whether to perform the resource configuration. Set it to false if the configuration will be done manually, and `nsx_config` script will not be run.
 - `nsx_config_script_path`: absolute path of the `nsx_config.py` script
 - `nsx_cert_file_path`: absolute path of NSX-T client certificate file
 - `nsx_manager_ip`: IP of NSX Manager
 - `nsx_edge_cluster_name`: name of the Edge Cluster to be used by the tier-0 router
 - `nsx_transport_zone_name`: name of the Overlay Transport Zone
 - `nsx_t0_router_name`: name of tier-0 Logical Router for the cluster
 - `pod_ipblock_name`: name of IP block for pods.
 - `pod_ipblock_cidr`: CIDR address for this IP block
 - `snat_ippool_name`: name of the IP block for SNAT
 - `snat_ippool_cidr`: CIDR address for this IP block
 - `start_range`: the start IP address of CIDR specified for this IP pool
 - `end_range`: the end IP address of CIDR specified for this IP pool
 - `os_cluster_name`: name of the OpenShift cluster
 - `os_node_name_list`: comma-separated list of node names
For example, `node1,node2,node3`
 - `nsx_node_ls_name`: name of Logical Switch connected to the nodes
 - `subnet_cidr`: CIDR address for IP administrator will assign to br-int on the node
 - `vc_host`: IP address of vCenter Server
 - `vc_user`: user name of vCenter Server administrator
 - `vc_password`: password of vCenter Server administrator
 - `vms`: comma-separated list of VM names. The order must match `os_node_name_list`.

The names must match the NSX-T resources you created. Otherwise the resources will be created with the names specified. The playbook is idempotent, and ensures resources with specified names exist and are tagged accordingly.

- 2 Change to the `openshift-ansible-nsx` directory and run the `ncp_prep` role.

```
ansible-playbook -i /PATH/TO/HOSTS/hosts ncp_prep.yaml
```

The playbook contains instructions to perform the following actions:

- Download the CNI plug-in installation file.
The filename is `nsx-cni-1.0.0.0.xxxxxxx-1.x86_64.rpm`, where `xxxxxxx` is the build number.
- Install the CNI plug-in installation file.
The plug-in is installed in `/opt/cni/bin`. The CNI configuration file `10.net.conf` is copied to `/etc/cni/net.d`. The rpm will also install the configuration file `/etc/cni/net.d/99-loopback.conf` for the loopback plug-in.
- Download and install the OVS installation files.
The files are `openvswitch-2.7.0.xxxxxxx-1.x86_64.rpm` and `openvswitch-kmod-2.7.0.xxxxxxx-1.el7.x86_64.rpm`, where `xxxxxxx` is the build number.
- Make sure that OVS is running.

```
# service openvswitch-switch status
```

- Create the `br-int` instance if it is not already created.

```
# ovs-vsctl add-br br-int
```

- Add the network interface (`node-if`) that is attached to the node logical switch to `br-int`.
- Make sure that the `br-int` and `node-if link` status is up.

```
# ip link set br-int up
# ip link set <node-if> up
```

- Update the network configuration file to ensure that the network interface is up after a reboot.
- Download the NCP tar file and load the Docker image from the tar file.

What to do next

Install OpenShift Container Platform. See [Install OpenShift Container Platform](#).

Install OpenShift Container Platform

OpenShift Container Platform (OCP) is a platform as a service (PaaS) product that brings together Docker and Kubernetes.

For information about installing OCP, see <https://docs.openshift.com>.

What to do next

Run NCP and NSX node agent. See [Run NCP and NSX Node Agent](#).

Run NCP and NSX Node Agent

Set up and run NCP and NSX node agent.

Procedure

- 1 Edit `roles/ncp/defaults/main.yaml` and specify the OpenShift API server IP, NSX manager IP, and URL's for downloading NCP ReplicationController yaml and nsx-node-agent DaemonSet yaml.
- 2 From the `openshift-ansible-nsx` directory, run the `ncp` role:

```
ansible-playbook -i /PATH/TO/HOSTS/hosts ncp.yaml
```

The `ncp` role performs the following steps:

- Check if `nsx-system` project exists, and create one if it does not.

```
oc new-project nsx-system
```

- Download the `ncp-rbac` yaml file and change the `apiVersion` to **v1**.
- Create the service account for the NCP pod, create a cluster role that specifies resources that NCP can access and bind the cluster role with the NCP service account.
- Create the service account for the `nsx-node-agent` pod, create a cluster role that specifies the resources that the node agent can access and bind the cluster role with the node agent service account.

```
oc apply -f /tmp/ncp-rbac.yaml
```

- Obtain the token associated with the above service accounts, and store it under `/etc/nsx-ujo/<service_account>_token`:

```
secret=`kubectl get serviceaccount ncp-svc-account -o yaml | grep -A1 secrets | tail -n1 | awk {'print $3'}`
kubectl get secret $secret -o yaml | grep 'token:' | awk {'print $2'} | base64 -d > /etc/nsx-ujo/ncp_token
secret=`kubectl get serviceaccount nsx-node-agent-svc-account -o yaml | grep -A1 secrets | tail -n1 | awk {'print $3'}`
kubectl get secret $secret -o yaml | grep 'token:' | awk {'print $2'} | base64 -d > /etc/nsx-ujo/node_agent_token
```

- Download the SecurityContextConstraint (SCC) yaml file for NCP and create the SCC based on the yaml.
- Add the created SCC to the above service accounts:

```
oadm policy add-scc-to-user ncp-scc -z ncp-svc-account
oadm policy add-scc-to-user ncp-scc -z nsx-node-agent-svc-account
```

- Download the yaml files for NCP ReplicationController (RC) and nsx-node-agent DaemonSet (DS) and update the ConfigMap.
- Download and load the NCP image (nsx-node-agent uses the same image).
- Configure the service account and set up the required SecurityContextConstraint for NCP and nsx_node_agent.
- Create the NCP ReplicationController and nsx-node-agent DaemonSet.

Note NCP opens persistent HTTP connections to the Kubernetes API server to watch for life cycle events of Kubernetes resources. If an API server failure or a network failure causes NCP's TCP connections to become stale, you must restart NCP so that it can re-establish connections to the API server. Otherwise, NCP will miss the new events.

Setup Notes

Before setting up OpenShift and NCP, take note of the following information.

- A pod must have no more than 11 labels and a namespace must have no more than 12 labels.
- Labels added for OpenShift internal usage, for example, a label with prefix openshift.io in its key, will be disregarded by NCP and thus user won't see the corresponding tags created on the related NSX resources. Here is a list of label prefixes used by OpenShift, and you should avoid using a label key starting with any of the following:

```
openshift.io
pod-template
```

- The nodes will need access to the pods, for example, for Kubelet health-checks. Make sure the host management interface is able to access the pod network.
- When configuring load balancing for services, you should use the OpenShift router. The NSX-T load balancer is not supported.
- Linux capabilities NET_ADMIN and NET_RAW can be exploited by attackers to compromise the pod network. You should disable these two capabilities of untrusted containers. By default, with restricted and anyuid SCC, NET_ADMIN is not granted. Be wary of any SCC that enables NET_ADMIN explicitly, or enables the pod to run in privileged mode. In addition, for untrusted containers, create a separate SCC based on, for example, anyuid SCC, with NET_RAW capability removed. This can be done by adding NET_RAW to `requiredDropCapabilities` list in the SCC definition.
- There is a defect in the Ansible playbook for upgrading OpenShift 3.6 to 3.7. It ignores the openshift_use_openshift_sdn=false setting in the inventory file and tries to install the OpenShift default SDN during the upgrade. To work around this issue, remove the CNI plug-in before the upgrade and re-install it after the upgrade.

- Allow root access in PODs/Containers (only for testing). Commands below will require root access in all PODs of the oc project you are currently logged in to.

```
oc new-project test-project
oc project test-project
oadm policy add-scc-to-user anyuid -z default
```

- Configure (add) the OpenShift Registry.

```
oc login -u system:admin -n default
oadm registry --service-account=registry --config=/etc/origin/master/admin.kubeconfig
```

- Delete the OpenShift Registry

```
oc login -u system:admin -n default
oc delete svc/docker-registry dc/docker-registry
```

- There is a missing IPtables firewall rule to allow DNS requests from the Docker default bridge containers to the dnsmasq process on the Node. It needs to be opened manually. Edit `/etc/sysconfig/iptables` and add the following Rules at the bottom of the file before COMMIT:

```
-A OS_FIREWALL_ALLOW -p tcp -m state --state NEW -m tcp --dport 53 -j ACCEPT
-A OS_FIREWALL_ALLOW -p udp -m state --state NEW -m udp --dport 53 -j ACCEPT
COMMIT
```

- Restart iptables, docker and origin-node (restarts kube-proxy and kubelet).

```
systemctl restart iptables
systemctl restart docker
systemctl restart origin-node
```

- The internal docker registry of OpenShift needs to be allowed to use non-TLS for OpenShift to work. Normally this should be added automatically by the OpenShift Ansible installer, but it seems that this is currently not working. Edit `/etc/sysconfig/docker` and add:

```
INSECURE_REGISTRY='--insecure-registry 172.30.0.0/16'
```

- Restart Docker.

```
systemctl restart docker
```

- Configure (add) the OpenShift routers (HA-Proxy N/S LBs).

```
oc login -u system:admin -n default
oadm router router --replicas=2 --service-account=router
```

- Delete the created routers.

```
oc login -u system:admin -n default
oc delete svc/router dc/router
```

- Create a sample Ruby-based 2-tier apps.

```
oc login -u system:admin -n default
oc oc new-project nsx
oc process -n openshift mysql-ephemeral -v DATABASE_SERVICE_NAME=database | oc create -f -
oc new-app centos/ruby-22-centos7-https://github.com/openshift/ruby-hello-world.git
oc expose service ruby-hello-world
oc env dc database --list | oc env dc ruby-hello-world -e -
```

- NCP's support for network policies is the same as the support provided by Kubernetes and depends on the Kubernetes version used by OpenShift.
 - OpenShift 3.7, 3.9 - The rule clauses in the network policy may contain at most one selector from namespaceSelector, podSelector and ipBlock.
 - OpenShift 3.7 - Support for egress rules and ipBlock CIDR selector is not available. They are available as beta fields in OpenShift 3.9.
- Certain versions of Kubernetes has a subPath-related issue (see <https://github.com/kubernetes/kubernetes/issues/61076>). If the OpenShift version does not contain a fix for this issue, the creation of the NCP pod will fail with the error CreateContainerConfigError: failed to prepare subPath for volumeMount. You can work around this problem by removing the use of subPath from the NCP yaml. Specifically, remove the line containing subPath: ncp.ini and replace the configuration for volumes with the following:

```
volumes:
  - name: config-volume
    # ConfigMap nsx-ncp-config is expected to supply ncp.ini
    configMap:
      name: nsx-ncp-config
      items:
        - key: ncp.ini
          path: ncp.ini
```

A side effect of this change is that the entire /etc/nsx-ujo directory becomes read-only. As a result, connecting with NSX-T using certificate and private key will not work because NCP will not be able to create a temporary file under /etc/nsx-ujo to move both certificate and private key into a single file.

Administering NSX-T Container Plug-in

4

You can administer NSX-T Container Plug-in from the NSX Manager GUI or from the command-line interface (CLI).

Note If a container host VM is running on ESXi 6.5 and the VM is migrated through vMotion to another ESXi 6.5 host, containers running on the container host will lose connectivity to containers running on other container hosts. You can resolve the problem by disconnecting and connecting the vNIC of the container host. This issue does not occur with ESXi 6.5 Update 1 or later.

Hyperbus reserves VLAN IDs 4093 and 4094 on the hypervisor for PVLAN configuration and the IDs cannot be changed. To avoid any VLAN conflict, do not configure VLAN logical switches or VTEP vmknics with the same VLAN IDs..

This chapter includes the following topics:

- [Manage IP Blocks from the NSX Manager GUI](#)
- [Manage IP Block Subnets from the NSX Manager GUI](#)
- [CIF-Attached Logical Ports](#)
- [CLI Commands](#)

Manage IP Blocks from the NSX Manager GUI

You can add, delete, edit, view details of, and manage the tags for an IP block from the NSX Manager GUI.

Procedure

- 1 From a browser, log in to the NSX Manager at `https://<nsx-manager-IP-address-or-domain-name>`.
- 2 Select **DDI**.

A list of the existing IP blocks is displayed.

3 Perform any of the following actions.

Option	Action
Add an IP block	Click ADD .
Delete one or more IP blocks	Select one or more IP blocks and click DELETE .
Edit an IP block	Select an IP block and click EDIT .
View details about an IP block	Click the IP block name. Click the Overview tab to see general information. Click the Subnets tab to see this IP block's subnets.
Manage tags for an IP block	Select an IP block and click ACTIONS > Manage Tags .

You cannot delete an IP block that has subnets allocated.

Manage IP Block Subnets from the NSX Manager GUI

You can add and delete subnets for an IP block from the NSX Manager GUI.

Procedure

- 1 From a browser, log in to the NSX Manager at `https://<nsx-manager-IP-address-or-domain-name>`.
- 2 Select **DDI**.
A list of the existing IP blocks is displayed.
- 3 Click an IP block name
- 4 Click the **Subnets** tab.
- 5 Perform any of the following actions..

Option	Action
Add an IP block subnet	Click ADD .
Delete one or more IP block subnets	Select one or more subnets and click DELETE .

CIF-Attached Logical Ports

CIFs (container interfaces) are network interfaces on containers that are connected to logical ports on a switch. These ports are called CIF-attached logical ports.

You can manage CIF-attached logical ports from the NSX Manager GUI.

Managing CIF-Attached Logical Ports

Navigate to **Switching > PORTS** to see all logical ports, including CIF-attached logical ports. Click the attachment link of a CIF-attached logical port to see the attachment information. Click the logical port link to open a window pane with four tabs: Overview, Monitor, Manage, and Related. Clicking **Related > Logical Ports** shows the related logical port on an uplink switch. For more information about switch ports, see the *NSX-T Administration Guide*.

Network Monitoring Tools

The following tools support CIF-attached logical ports. For more information about these tools, see the *NSX-T Administration Guide*.

- Traceflow
- Port Connection
- IPFIX
- Remote port mirroring using GRE encapsulation of a logical switch port that connects to a container is supported. For more information, see "Understanding Port Mirroring Switching Profile" in the *NSX-T Administration Guide*. However, port mirroring of the CIF to VIF port is not supported.

Distributed network encryption is not supported in this release.

CLI Commands

To run CLI commands, log in to the NSX-T Container Plug-in container, open a terminal and run the `nsxcli` command.

You can also get the CLI prompt by running the following command on a node:

```
kubectl exec -it <pod name> nsxcli
```

Table 4-1. CLI Commands for the NCP Container

Type	Command
Status	get ncp-master status
Status	get ncp-nsx status
Status	get ncp-watcher <watcher-name>
Status	get ncp-watchers
Status	get ncp-k8s-api-server status
Status	check projects
Status	check project <project-name>
Cache	get project-cache <project-name>
Cache	get project-caches
Cache	get namespace-cache <namespace-name>
Cache	get namespace-caches
Cache	get pod-cache <pod-name>
Cache	get pod-caches
Cache	get ingress-caches
Cache	get ingress-cache <ingress-name>
Cache	get ingress-controllers

Table 4-1. CLI Commands for the NCP Container (Continued)

Type	Command
Cache	get ingress-controller <ingress-controller-name>
Cache	get network-policy-caches
Cache	get network-policy-cache <pod-name>
Support	get ncp-log file <filename>
Support	get ncp-log-level
Support	set ncp-log-level <log-level>
Support	get support-bundle file <filename>
Support	get node-agent-log file <filename>
Support	get node-agent-log file <filename> <node-name>

Table 4-2. CLI Commands for the NSX Node Agent Container

Type	Command
Status	get node-agent-hyperbus status
Cache	(NCP 2.1, 2.1.0.1) get app-cache <app-name> (NCP 2.1.2) get container-cache <container-name>
Cache	(NCP 2.1, 2.1.0.1) get app-caches (NCP 2.1.2) get container-caches

Table 4-3. CLI Commands for the NSX Kube Proxy Container

Type	Command
Status	get ncp-k8s-api-server status
Status	get kube-proxy-watcher <watcher-name>
Status	get kube-proxy-watchers
Status	dump ovs-flows

Status Commands for the NCP Container

- Show the status of the NCP master

```
get ncp-master status
```

Example:

```
kubenode> get ncp-master status
This instance is not the NCP master
Current NCP Master id is a4h83eh1-b8dd-4e74-c71c-cbb7cc9c4c1c
Last master update at Wed Oct 25 22:46:40 2017
```

- Show the connection status between NCP and NSX Manager

```
get ncp-nsx status
```

Example:

```
kubenode> get ncp-nsx status
NSX Manager status: Healthy
```

- Show the watcher status for ingress, namespace, pod, and service

```
get ncp-watcher <watcher-name>
get ncp-watchers
```

Example 1:

```
kubenode> get ncp-watcher pod
Average event processing time: 1174 msec (in past 3600-sec window)
Current watcher started time: Mar 02 2017 10:47:35 PST
Number of events processed: 1 (in past 3600-sec window)
Total events processed by current watcher: 1
Total events processed since watcher thread created: 1
Total watcher recycle count: 0
Watcher thread created time: Mar 02 2017 10:47:35 PST
Watcher thread status: Up
```

Example 2:

```
kubenode> get ncp-watchers
pod:
Average event processing time: 1145 msec (in past 3600-sec window)
Current watcher started time: Mar 02 2017 10:51:37 PST
Number of events processed: 1 (in past 3600-sec window)
Total events processed by current watcher: 1
Total events processed since watcher thread created: 1
Total watcher recycle count: 0
Watcher thread created time: Mar 02 2017 10:51:37 PST
Watcher thread status: Up

namespace:
Average event processing time: 68 msec (in past 3600-sec window)
Current watcher started time: Mar 02 2017 10:51:37 PST
Number of events processed: 2 (in past 3600-sec window)
Total events processed by current watcher: 2
Total events processed since watcher thread created: 2
Total watcher recycle count: 0
Watcher thread created time: Mar 02 2017 10:51:37 PST
Watcher thread status: Up

ingress:
Average event processing time: 0 msec (in past 3600-sec window)
Current watcher started time: Mar 02 2017 10:51:37 PST
```

```

Number of events processed: 0 (in past 3600-sec window)
Total events processed by current watcher: 0
Total events processed since watcher thread created: 0
Total watcher recycle count: 0
Watcher thread created time: Mar 02 2017 10:51:37 PST
Watcher thread status: Up

```

service:

```

Average event processing time: 3 msec (in past 3600-sec window)
Current watcher started time: Mar 02 2017 10:51:37 PST
Number of events processed: 1 (in past 3600-sec window)
Total events processed by current watcher: 1
Total events processed since watcher thread created: 1
Total watcher recycle count: 0
Watcher thread created time: Mar 02 2017 10:51:37 PST
Watcher thread status: Up

```

- Show the connection status between NCP and Kubernetes API server

```
get ncp-k8s-api-server status
```

Example:

```
kubecall> get ncp-k8s-api-server status
Kubernetes ApiServer status: Healthy
```

- Check all projects or a specific one

```
check projects
check project <project-name>
```

Example:

```
kubecall> check projects
default:
  Tier-1 link port for router 1b90a61f-0f2c-4768-9eb6-ea8954b4f327 is missing
  Switch 40a6829d-c3aa-4e17-ae8a-7f7910fdf2c6 is missing

ns1:
  Router 8accc9cd-9883-45f6-81b3-0d1fb2583180 is missing

kubecall> check project default
Tier-1 link port for router 1b90a61f-0f2c-4768-9eb6-ea8954b4f327 is missing
Switch 40a6829d-c3aa-4e17-ae8a-7f7910fdf2c6 is missing
```

Cache Commands for the NCP Container

- Get the internal cache for projects or namespaces

```
get project-cache <project-name>
get project-caches
get namespace-cache <namespace-name>
get namespace-caches
```

Example:

```
kubecall> get project-caches
  default:
    logical-router: 8accc9cd-9883-45f6-81b3-0d1fb2583180
    logical-switch:
      id: 9d7da647-27b6-47cf-9cdb-6e4f4d5a356d
      ip_pool_id: 519ff57f-061f-4009-8d92-3e6526e7c17e
      subnet: 10.0.0.0/24
      subnet_id: f75fd64c-c7b0-4b42-9681-fc656ae5e435

  kube-system:
    logical-router: 5032b299-acad-448e-a521-19d272a08c46
    logical-switch:
      id: 85233651-602d-445d-ab10-1c84096cc22a
      ip_pool_id: ab1c5b09-7004-4206-ac56-85d9d94bffa2
      subnet: 10.0.1.0/24
      subnet_id: 73e450af-b4b8-4a61-a6e3-c7ddd15ce751

  testns:
    ext_pool_id: 346a0f36-7b5a-4ecc-ad32-338dcb92316f
    labels:
      ns: myns
      project: myproject
    logical-router: 4dc8f8a9-69b4-4ff7-8fb7-d2625dc77efa
    logical-switch:
      id: 6111a99a-6e06-4faa-a131-649f10f7c815
      ip_pool_id: 51ca058d-c3dc-41fd-8f2d-e69006ab1b3d
      subnet: 50.0.2.0/24
      subnet_id: 34f79811-bd29-4048-a67d-67ceac97eb98
    project_nsgroup: 9606afee-6348-4780-9dbe-91abfd23e475
    snat_ip: 4.4.0.3

kubecall> get project-cache default
  logical-router: 8accc9cd-9883-45f6-81b3-0d1fb2583180
  logical-switch:
    id: 9d7da647-27b6-47cf-9cdb-6e4f4d5a356d
    ip_pool_id: 519ff57f-061f-4009-8d92-3e6526e7c17e
    subnet: 10.0.0.0/24
    subnet_id: f75fd64c-c7b0-4b42-9681-fc656ae5e435

kubecall> get namespace-caches
  default:
    logical-router: 8accc9cd-9883-45f6-81b3-0d1fb2583180
```

```

logical-switch:
  id: 9d7da647-27b6-47cf-9cdb-6e4f4d5a356d
  ip_pool_id: 519ff57f-061f-4009-8d92-3e6526e7c17e
  subnet: 10.0.0.0/24
  subnet_id: f75fd64c-c7b0-4b42-9681-fc656ae5e435

kube-system:
  logical-router: 5032b299-acad-448e-a521-19d272a08c46
  logical-switch:
    id: 85233651-602d-445d-ab10-1c84096cc22a
    ip_pool_id: ab1c5b09-7004-4206-ac56-85d9d94bfffa2
    subnet: 10.0.1.0/24
    subnet_id: 73e450af-b4b8-4a61-a6e3-c7ddd15ce751

testns:
  ext_pool_id: 346a0f36-7b5a-4ecc-ad32-338dcb92316f
  labels:
    ns: myns
    project: myproject
  logical-router: 4dc8f8a9-69b4-4ff7-8fb7-d2625dc77efa
  logical-switch:
    id: 6111a99a-6e06-4faa-a131-649f10f7c815
    ip_pool_id: 51ca058d-c3dc-41fd-8f2d-e69006ab1b3d
    subnet: 50.0.2.0/24
    subnet_id: 34f79811-bd29-4048-a67d-67ceac97eb98
  project_nsgroup: 9606afee-6348-4780-9dbe-91abfd23e475
  snat_ip: 4.4.0.3

kubenode> get namespace-cache default
logical-router: 8accc9cd-9883-45f6-81b3-0d1fb2583180
logical-switch:
  id: 9d7da647-27b6-47cf-9cdb-6e4f4d5a356d
  ip_pool_id: 519ff57f-061f-4009-8d92-3e6526e7c17e
  subnet: 10.0.0.0/24
  subnet_id: f75fd64c-c7b0-4b42-9681-fc656ae5e435

```

- Get the internal cache for pods

```

get pod-cache <pod-name>
get pod-caches

```

Example:

```

kubenode> get pod-caches
nsx.default.nginx-rc-uq2lv:
  cif_id: 2af9f734-37b1-4072-ba88-abbf935bf3d4
  gateway_ip: 10.0.0.1
  host_vif: d6210773-5c07-4817-98db-451bd1f01937
  id: 1c8b5c52-3795-11e8-ab42-005056b198fb
  ingress_controller: False
  ip: 10.0.0.2/24
  labels:
    app: nginx
  mac: 02:50:56:00:08:00

```

```

port_id: d52c833a-f531-4bdf-bfa2-e8a084a8d41b
vlan: 1

nsx.testns.web-pod-1:
  cif_id: ce134f21-6be5-43fe-afbf-aaca8c06b5cf
  gateway_ip: 50.0.2.1
  host_vif: d6210773-5c07-4817-98db-451bd1f01937
  id: 3180b521-270e-11e8-ab42-005056b198fb
  ingress_controller: False
  ip: 50.0.2.3/24
  labels:
    app: nginx-new
    role: db
    tier: cache
  mac: 02:50:56:00:20:02
  port_id: 81bc2b8e-d902-4cad-9fc1-aabdc32ecaf8
  vlan: 3

kubensode> get pod-cache nsx.default.nginx-rc-uj2lv
cif_id: 2af9f734-37b1-4072-ba88-abbf935bf3d4
gateway_ip: 10.0.0.1
host_vif: d6210773-5c07-4817-98db-451bd1f01937
id: 1c8b5c52-3795-11e8-ab42-005056b198fb
ingress_controller: False
ip: 10.0.0.2/24
labels:
  app: nginx
mac: 02:50:56:00:08:00
port_id: d52c833a-f531-4bdf-bfa2-e8a084a8d41b
vlan: 1

```

- Get network policy caches or a specific one

```

get network-policy caches
get network-policy-cache <network-policy-name>

```

Example:

```

kubensode> get network-policy-caches
nsx.testns.allow-tcp-80:
  dest_labels: None
  dest_pods:
    50.0.2.3
  match_expressions:
    key: tier
    operator: In
    values:
      cache
  name: allow-tcp-80
  np_dest_ip_set_ids:
    22f82d76-004f-4d12-9504-ce1cb9c8aa00
  np_except_ip_set_ids:
  np_ip_set_ids:
    14f7f825-f1a0-408f-bbd9-bb2f75d44666

```

```

np_isol_section_id: c8d93597-9066-42e3-991c-c550c46b2270
np_section_id: 04693136-7925-44f2-8616-d809d02cd2a9
ns_name: testns
src_egress_rules: None
src_egress_rules_hash: 97d170e1550eee4afc0af065b78cda302a97674c
src_pods:
  50.0.2.0/24
src_rules:
  from:
    namespaceSelector:
      matchExpressions:
        key: tier
        operator: DoesNotExist
      matchLabels:
        ns: myns
    ports:
      port: 80
      protocol: TCP
src_rules_hash: e4ea7b8d91c1e722670a59f971f8fcc1a5ac51f1

```

```

kubecall> get network-policy-cache nsx.testns.allow-tcp-80
dest_labels: None
dest_pods:
  50.0.2.3
match_expressions:
  key: tier
  operator: In
  values:
    cache
name: allow-tcp-80
np_dest_ip_set_ids:
  22f82d76-004f-4d12-9504-ce1cb9c8aa00
np_except_ip_set_ids:
np_ip_set_ids:
  14f7f825-f1a0-408f-bbd9-bb2f75d44666
np_isol_section_id: c8d93597-9066-42e3-991c-c550c46b2270
np_section_id: 04693136-7925-44f2-8616-d809d02cd2a9
ns_name: testns
src_egress_rules: None
src_egress_rules_hash: 97d170e1550eee4afc0af065b78cda302a97674c
src_pods:
  50.0.2.0/24
src_rules:
  from:
    namespaceSelector:
      matchExpressions:
        key: tier
        operator: DoesNotExist
      matchLabels:
        ns: myns

```

```
ports:
  port: 80
  protocol: TCP
src_rules_hash: e4ea7b8d91c1e722670a59f971f8fcc1a5ac51f1
```

Support Commands for the NCP Container

- Save the NCP support bundle in the filestore

The support bundle consists of the log files for all the containers in pods with the label **tier: nsx-networking**. The bundle file is in the tgz format and saved in the CLI default filestore directory `/var/vmware/nsx/file-store`. You can use the CLI file-store command to copy the bundle file to a remote site.

```
get support-bundle file <filename>
```

Example:

```
kubnode>get support-bundle file foo
Bundle file foo created in tgz format
kubnode>copy file foo url scp://nicira@10.0.0.1:/tmp
```

- Save the NCP logs in the filestore

The log file is saved in the tgz format in the CLI default filestore directory `/var/vmware/nsx/file-store`. You can use the CLI file-store command to copy the bundle file to a remote site.

```
get ncp-log file <filename>
```

Example:

```
kubnode>get ncp-log file foo
Log file foo created in tgz format
```

- Save the node agent logs in the filestore

Save the node agent logs from one node or all the nodes. The logs are saved in the tgz format in the CLI default filestore directory `/var/vmware/nsx/file-store`. You can use the CLI file-store command to copy the bundle file to a remote site.

```
get node-agent-log file <filename>
get node-agent-log file <filename> <node-name>
```

Example:

```
kubnode>get node-agent-log file foo
Log file foo created in tgz format
```

- Get and set the log level

The available log levels are NOTSET, DEBUG, INFO, WARNING, ERROR, and CRITICAL.

```
get ncp-log-level
set ncp-log-level <log level>
```

Example:

```
kubenode>get ncp-log-level
NCP log level is INFO

kubenode>set ncp-log-level DEBUG
NCP log level is changed to DEBUG
```

Status Commands for the NSX Node Agent Container

- Show the connection status between the node agent and HyperBus on this node.

```
get node-agent-hyperbus status
```

Example:

```
kubenode> get node-agent-hyperbus status
HyperBus status: Healthy
```

Cache Commands for the NSX Node Agent Container

- Get the internal cache for NSX node agent containers.

```
(NCP 2.1, 2.1.0.1) get app-cache <app-name>
(NCP 2.1, 2.1.0.1) get app-caches
(NCP 2.1.2) get container-cache <container-name>
(NCP 2.1.2) get container-caches
```

Example 1:

```
kubenode> get container-cache cif104
ip: 192.168.0.14/32
mac: 50:01:01:01:01:14
gateway_ip: 169.254.1.254/16
vlan_id: 104
```

Example 2:

```
kubenode> get container-caches
cif104:
ip: 192.168.0.14/32
mac: 50:01:01:01:01:14
gateway_ip: 169.254.1.254/16
vlan_id: 104
```

Status Commands for the NSX Kube-Proxy Container

- Show the connection status between Kube Proxy and Kubernetes API Server

```
get ncp-k8s-api-server status
```

Example:

```
kubecall> get kube-proxy-k8s-api-server status
Kubernetes ApiServer status: Healthy
```

- Show the Kube Proxy watcher status

```
get kube-proxy-watcher <watcher-name>
get kube-proxy-watchers
```

Example 1:

```
kubecall> get kube-proxy-watcher endpoint
Average event processing time: 15 msec (in past 3600-sec window)
Current watcher started time: May 01 2017 15:06:24 PDT
Number of events processed: 90 (in past 3600-sec window)
Total events processed by current watcher: 90
Total events processed since watcher thread created: 90
Total watcher recycle count: 0
Watcher thread created time: May 01 2017 15:06:24 PDT
Watcher thread status: Up
```

Example 2:

```
kubecall> get kube-proxy-watchers
endpoint:
Average event processing time: 15 msec (in past 3600-sec window)
Current watcher started time: May 01 2017 15:06:24 PDT
Number of events processed: 90 (in past 3600-sec window)
Total events processed by current watcher: 90
Total events processed since watcher thread created: 90
Total watcher recycle count: 0
Watcher thread created time: May 01 2017 15:06:24 PDT
Watcher thread status: Up

service:
Average event processing time: 8 msec (in past 3600-sec window)
Current watcher started time: May 01 2017 15:06:24 PDT
Number of events processed: 2 (in past 3600-sec window)
Total events processed by current watcher: 2
Total events processed since watcher thread created: 2
Total watcher recycle count: 0
Watcher thread created time: May 01 2017 15:06:24 PDT
Watcher thread status: Up
```

- Dump OVS flows on a node

```
dump ovs-flows
```

Example:

```
kubenode> dump ovs-flows
NXST_FLOW reply (xid=0x4):
  cookie=0x0, duration=8.876s, table=0, n_packets=0, n_bytes=0, idle_age=8, priority=100,ip
actions=ct(table=1)
  cookie=0x0, duration=8.898s, table=0, n_packets=0, n_bytes=0, idle_age=8, priority=0
actions=NORMAL
  cookie=0x0, duration=8.759s, table=1, n_packets=0, n_bytes=0, idle_age=8,
priority=100,tcp,nw_dst=10.96.0.1,tp_dst=443 actions=mod_tp_dst:443
  cookie=0x0, duration=8.719s, table=1, n_packets=0, n_bytes=0, idle_age=8,
priority=100,ip,nw_dst=10.96.0.10 actions=drop
  cookie=0x0, duration=8.819s, table=1, n_packets=0, n_bytes=0, idle_age=8,
priority=90,ip,in_port=1 actions=ct(table=2,nat)
  cookie=0x0, duration=8.799s, table=1, n_packets=0, n_bytes=0, idle_age=8, priority=80,ip
actions=NORMAL
  cookie=0x0, duration=8.856s, table=2, n_packets=0, n_bytes=0, idle_age=8, actions=NORMAL
```