

# VMware NSX Container Plug-in 2.3.2 Release Notes

Updated on 08/15/2019

VMware NSX Container Plug-in 2.3.2 | 17 January, 2019

Check regularly for additions and updates to these release notes.

## What's in the Release Notes

The release notes cover the following topics:

- [What's New](#)
- [Compatibility Requirements](#)
- [Resolved Issues](#)
- [Known Issues](#)

## What's New

### What's New

NSX Container Plug-in 2.3.2 is a maintenance release specifically for the NSX Container Plug-in (NCP) feature of NSX-T 2.3.x. This release resolves a number of issues found in previous releases and has the following new features:

- Support for HTTP Ingress annotation `kubernetes.io/ingress.allow-http`.
- Support for specifying an external IP pool for Kubernetes LoadBalancer services, native and third-party Ingresses. After a restart, NCP will re-allocate IPs if a different external pool has been specified.
- Support for specifying an SNAT IP pool for a Kubernetes namespace.
- Support for changing the external IP pool for SNAT. After restarting NCP, the SNAT IPs of projects (Kubernetes namespaces or PCF orgs) will be re-allocated from the new external IP pools.
- Support for PAS apps with multiple process. With PCF V3 API, an app can have multiple process types, each with different commands and scale. NCP will create a logical switch port for each instance of all processes.

## Compatibility Requirements

Product	Version
NCP / NSX-T Tile for PAS	2.3.2
NSX-T	2.2, 2.3, 2.3.1
Kubernetes	1.12, 1.13
OpenShift	3.10, 3.11
Kubernetes Host VM OS	Ubuntu 16.04, RHEL 7.4, 7.5, CentOS 7.4, 7.5
OpenShift Host VM OS	RHEL 7.4, RHEL 7.5
PAS (PCF)	OpsManager 2.3.x + PAS 2.3.x OpsManager 2.4.0 + PAS 2.4.0

## Resolved Issues

- Issue 2194845: The PAS Cloud Foundry V3 API feature "multiple processes per app" is not supported

When using the PAS Cloud Foundry V3 API `v3-push` to push an app with multiple processes, NCP does not create logical switch ports for the processes except the default one. This issue exists in NCP 2.3.1 and earlier releases.

## Known Issues

- Issue 2118515: In a large-scale setup, NCP takes a long time to create firewalls on NSX-T

In a large-scale setup (for example, 250 Kubernetes nodes, 5000 pods, 2500 network policies), it can take NCP a few minutes to create the firewall sections and rules in NSX-T.

Workaround: None. After the firewall sections and rules are created, performance should be back to normal.

- Issue 2125755: A StatefulSet could lose network connectivity when performing canary updates and phased rolling updates

If a StatefulSet was created before NCP was upgraded to the current release, the StatefulSet could lose network connectivity when performing canary updates and phased rolling updates.

Workaround: Create the StatefulSet after NCP is upgraded to the current release.

- Issue 2131494: NGINX Kubernetes Ingress still works after changing the Ingress class from `nginx` to `nsx`

When you create an NGINX Kubernetes Ingress, NGINX create traffic forwarding rules. If you change the Ingress class to any other value, NGINX does not delete the rules and continues to apply them, even if you delete the Kubernetes Ingress after changing the class. This is a limitation of NGINX.

Workaround: To delete the rules created by NGINX, delete the Kubernetes Ingress when the class value is `nginx`. Then re-create the Kubernetes Ingress.

- For a Kubernetes service of type ClusterIP, Client-IP based session affinity is not supported

NCP does not support Client-IP based session affinity for a Kubernetes service of type ClusterIP.

Workaround: None

- For a Kubernetes service of type ClusterIP, the `hairpin-mode` flag is not supported

NCP does not support the hairpin-mode flag for a Kubernetes service of type ClusterIP.

Workaround: None

- Issue 2193901: Multiple PodSelectors or multiple NsSelectors for a single Kubernetes network policy rule is not supported

Applying multiple selectors allows only incoming traffic from specific pods.

Workaround: Use matchLabels with matchExpressions in a single PodSelector or NsSelector instead.

- Issue 2194646: Updating network policies when NCP is down is not supported

If you update a network policy when NCP is down, the destination IPset for the network policy will be incorrect when NCP comes back up.

Workaround: Recreate the network policy when NCP is up.

- Issue 2192489: After disabling 'BOSH DNS server' in PAS director config, the Bosh DNS server (169.254.0.2) still appears in the container's resolve.conf file.

In a PAS environment running PAS 2.2, after you disable 'BOSH DNS server' in PAS director config, the Bosh DNS server (169.254.0.2) still appears in the container's resolve.conf file. This causes a ping command with a fully qualified domain name to take a long time. This issue does not exist with PAS 2.1.

Workaround: None. This is a PAS issue.

- Issue 2199504: The display name of NSX-T resources created by NCP is limited to 80 characters

When NCP creates an NSX-T resource for a resource in the container environment, it generates the display name of the NSX-T resource by combining the cluster name, namespace or project name, and the name of the resource in the container environment. If the display name is longer than 80 characters, it is truncated to 80 characters.

Workaround: None

- Issue 2199778: With NSX-T 2.2, Ingress, Service and Secrets with names longer than 65 characters are not supported

With NSX-T 2.2, when `use_native_loadbalancer` is set to `True`, the names of Ingresses, Secrets and Services referenced by the Ingress, and Services of type LoadBalancer, must be 65 characters or less. Otherwise, the Ingress or Service will not work properly.

Workaround: When configuring an Ingress, Secret, or Service, specify a name that is 65 characters or less.

- Issue 2065750: Installing the NSX-T CNI package fails with a file conflict

In a RHEL environment with kubernetes installed, if you install the NSX-T CNI Package using `yum localinstall` or `rpm -i`, you get an error indicating a conflict with a file from the kubernetes-cni package.

Workaround: Install the NSX-T CNI package with the command `rpm -i --replacefiles nsx-cni-2.3.0.xxxxxxxx-1.x86_64.rpm`.

- **Issue 2224218: After a service or app is deleted, it takes 2 minutes to release the SNAT IP back to the IP pool**

If you delete a service or app and recreate it within 2 minutes, it will get a new SNAT IP from the IP pool.

Workaround: After deleting a service or app, wait 2 minutes before recreating it if you want to reuse the same IP.

- **Issue 2218008: Configuring different Kubernetes clusters to use the same IP block causes connectivity problems**

If you configure different Kubernetes clusters to use the same IP block, some pods will not be able to communicate with other pods or external networks.

Workaround: Do not configure different Kubernetes clusters to use the same IP block.

- **Issue 2263536: Kubernetes service of type NodePort fails to forward traffic**

With a service of type NodePort, a Kubernetes node acts like a router that forwards traffic from outside the cluster to the pods. When setting up such node, sometimes the rules in iptables are not configured correctly to allow traffic to pass through.

Workaround: Run the following command to add a rule to iptables manually:

```
iptables -I FORWARD -m conntrack --ctstate RELATED,ESTABLISHED -j ACCEPT
```

Note that this works only for a NodePort service with "externalTrafficPolicy: Cluster". It does not work for "externalTrafficPolicy: Local".