



VMware NSX-T 2.2.1 Release Notes

VMware NSX Container Plug-in 2.2.1 | 19 JULY 2018

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 2.2.1 is a maintenance release specifically for the NSX Container Plug-in (NCP) feature of NSX-T 2.2. This release resolves a number of issues found in previous releases and has the following new features:

- Support for limiting SNAT IP pool to specific Kubernetes or OpenShift namespaces or PCF orgs.
- Pivotal Application Service (PAS) NSX-T tile improvements.
 - Open vSwitch is now a separate package.
 - Improvements in Open vSwitch kernel module compilation.

Compatibility Requirements

Product	Version
NCP / NSX-T Tile for PAS	2.2.1
NSX-T	2.1, 2.2
Kubernetes	1.9, 1.10, 1.11
OpenShift	3.7, 3.9
Kubernetes/OpenShift host VM OS	Ubuntu 16.04, RHEL 7.4, RHEL 7.5
	OpsManager 2.1.x + PAS 2.1.x (except PAS

Resolved Issues

- **Issue 2127607: After removing the TLS specification from Kubernetes Ingress, the default pool is not moved to the HTTP virtual server**

A Kubernetes Ingress is configured with TLS and a default backend. After removing the TLS specification, the rules are transferred from the HTTPS virtual server to the HTTP virtual server but the default pool for the backend service remains in the HTTPS virtual server.

Workaround: Delete the Kubernetes Ingress, restart NCP, and re-create the Kubernetes Ingress.

- **Issue 2164378: After deleting the NCP pod, the newly created NCP pod crashes.**
If NCP is restarted because the NCP pod is deleted, new virtual servers in the existing NSX-T load balancer for the cluster are created, and the old virtual servers are not cleaned up properly. This can cause NCP to fail to restart.
- **Issue 2164363: In a PKS environment, when NCP restarts, new load balancer virtual servers for Ingress are created and the old ones are not cleaned up.**
If NCP is restarted because the NCP pod is deleted or the node where NCP runs is recreated, extra L7 virtual servers with port 80 and port 443 will be created, even though the required L7 virtual servers with port 80 and port 443 already exist.
- **Issue 2165096: After deleting a load balancer service with multiple ports, not all the virtual servers of the NSX-T load balancer are deleted.**
After deleting a load balancer service with multiple ports, there might be virtual servers of the NSX-T load balancer that are not deleted. If a service of type LoadBalancer is recreated with the same name and in the same namespace, NCP does not handle it properly.

Known Issues

- **PAS 2.1.0 CNI change**
Due to the CNI plugin change in PAS 2.1.0, none of the NSX-T Tile 2.1.x and 2.2.x will work with PAS 2.1.0. This is fixed in PAS 2.1.1.
- **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. Than re-create the Kubernetes Ingress.

- **Issue 2160806: Updating the TLS spec of an active Ingress when NCP is not running is not supported**

If NCP has assigned an external IP to an Ingress resource, and you update the TLS spec of the Ingress when NCP is not running, for example, by removing or changing the parameter secretName, NCP will not be aware of the changes. When NCP runs again, the certificate corresponding to the old secret still exists and will not be deleted.

Workaround: Perform the following steps:

- Stop NCP.
- Delete the Ingress.
- Restart NCP. The NSX resources corresponding to the old Ingress will be deleted.
- Recreate the Ingress.

- **Issue 2193714: Error in documentation regarding the rewrite annotation**

In the NSX-T Container Plug-in for Kubernetes and Cloud Foundry - Installation and Administration Guide, an example illustrating how to specify a rewrite annotation (https://docs.vmware.com/en/VMware-NSX-T/2.2/com.vmware.nsx.t.ncp_kubernetes.doc/GUID-FF86F9C2-E70D-418B-AFDB-B3CF8DE106C2.html) contains the following line:

```
ncp/rewrite_target: ""
```

This line is incorrect. The correct annotation should be:

```
ingress.kubernetes.io/rewrite-target: /
```

- **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

Copyright © 2021 VMware, Inc. All rights reserved.