NSX-T Troubleshooting Guide

Modified on 21 DEC 2017
VMware NSX-T Data Center 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

Contents

NSX-T Troubleshooting Guide

The *NSX-T Troubleshooting Guide* provides information on how to troubleshoot issues that might occur in an NSX-T environment.

Intended Audience

This guide is for system administrators of NSX-T. A familiarity with virtualization, networking, and datacenter operations 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.

Troubleshooting Layer 2 Connectivity

1

If there is a communication failure between two virtual interfaces (VIFs) that are connected to the same logical switch, for example, you cannot ping one VM from another, you can follow the steps in this section to troubleshoot the failure.

Before you start, make sure that there is no firewall rule blocking traffic between the two logical ports. It is recommended that you follow the order of the topics in this section to troubleshoot the connectivity issue.

This chapter includes the following topics:

Check the NSX Manager and NSX Controller Cluster Status

Verify that the status of NSX Manager and the NSX Controller cluster is normal, and the controllers are connected to the NSX Manager.

Procedure

1 Run the following CLI command on the NSX Manager to make sure the status is stable.

```
NSX-Manager> get management-cluster status
Number of nodes in management cluster: 1
- 192.168.110.47 (UUID 45A8869B-BB90-495D-8A01-69B5FCC56086) Online

Management cluster status: STABLE

Number of nodes in control cluster: 3
- 192.168.110.201 (UUID 45A8869B-BB90-495D-8A01-69B5FCC56086)
- 192.168.110.202 (UUID 45A8869B-BB90-495D-8A01-69B5FCC56086)
- 192.168.110.203 (UUID 45A8869B-BB90-495D-8A01-69B5FCC56086)
```

2 Run the following CLI command on an NSX Controller to make sure the status is active.

```
NSX-Controller1> get control-cluster status
uuid: db4aa77a-4397-4d65-ad33-9fde79ac3c5c
is master: true
in majority: true
uuid                                   address              status
0cfe232e-6c28-4fea-8aa4-b3518baef00d 192.168.110.201       active
bd257108-b94e-4e6d-8b19-7fa6c012961d 192.168.110.202       active
538be554-1240-40e4-8e94-1497e963a2aa 192.168.110.203       active
```

3 Run the following CLI command on an NSX Controller to make sure it is connected to the NSX Manager.

```
NSX-Controller1> get managers
- 192.168.110.47   Connected
```

Check the Logical Ports

Check that the logical ports are configured on the same logical switch and their status is up.

Procedure

1 From the NSX Manager GUI, get the logical ports UUIDs.

2 Make the following API call for each logical port to make sure the logical ports are on the same logical switch.

```
GET https://<nsx-mgr>/api/v1/logical-ports/<logical-port-uuid>
```

3 Make the following API call for each logical port to make sure the status is up.

```
GET https://<nsx-mgr>/api/v1/logical-ports/<logical-port-uuid>/status
```

Check the Transport Node Status

Check the status of the transport node.

Procedure

- Make the following API call to get the state of the transport node.

  ```
  GET https://<nsx-mgr>/api/v1/transport-nodes/<transport-node-ID>/state
  ```

  If the call returns the error `RPC timeout`, perform the following troubleshooting steps:

  - Run `/etc/init.d/nsxa status` to see if `nsxa` is running.
  - Run `/etc/init.d/nsx-mpa status` to see if `nsx-mpa` is running.
  - To see if `nsx-mpa` is connected to the NSX Manager, check the `nsx-mpa` heartbeat logs.
  - To see if `nsxa` is connected to the NSX Manager, check the `nsxa` log. You will see the following message if `nsxa` is connected to the NSX Manager.

    ```
    NSXA_LOG(LVL_INFO, "[%s] Connected to mpa, cookie:[%d]\n", __func__ ,
    _mpaCookieId );https://opengrok.eng.vmware.com/source/xref/nsx.git/mpa/clients/nsxa/src/core/mp
    aClient.cpp#419
    ```

  - To see if `nsxa` is stuck processing `HostConfigMsg`, check the `nsxa` log. If so, you will see an RMQ request message but the reply is not sent or sent after a long delay.
  - Check to see if `nsxa` crashed while executing `HostConfigMsg`.
  - To see if the RMQ messages are taking a long time to be delivered to the host, compare the timestamps of log messages on the NSX Manager and the host.

  If the call returns the error `partial_success`, there are many possible causes. Start by looking at the `nsxa` logs. On the ESXi host, check `hostd.log` and `vmkernel.log`. On KVM, `syslog` holds all the logs.

Check the Logical Switch Status

Check the status of the logical switch.

Procedure

- Make the following API call to get the state of the logical switch.

  ```
  GET https://<nsx-mgr>/api/v1/logical-switches/<logical-switch-ID>/state
  ```

If the call returns the error `partial_success`, the reply will contain a list of transport nodes where the NSX Manager failed to push the logical switch configuration or did not get a reply. The troubleshooting steps are similar to those for the transport node. Check the following:

- All required components are installed and running.
- `nsx-mpa` is connected to the NSX Manager.
- `nsxa` is connected to the switching vertical.
- Grep the logical switch ID in `nsxa.log` and `nsxaVim.log` to see if the logical switch configuration was received by the transport node.
- Check the `nsxa` and `nsx-mpa` uptime. Find out when `nsxa` was started and stopped by grepping `nsxa` log messages in the syslog file.
- Find out `nsxa`'s connection time to the switching vertical. If the logical switch configuration is sent to the host when `nsxa` is not connected to the switching vertical, the configuration might not be delivered to the host.

On KVM, no logical switch configuration is pushed to the host. Therefore, most of the logical switch issues are likely to be in the management plane.

On ESXi, an opaque network is mapped to the logical switch. To use the logical switch, users connect VMs to the opaque network using vCenter Server or vSphere API.

Check the CCP for the Logical Switch

Verify that the logical switch is in the central control plane (CCP).

Procedure

- Run the following CLI command on an NSX Controller to make sure that the logical switch is present.

```
NSX-Controller1> get logical switches
VNI   UUID                                 Name
52104 feab22ec-94b2-46f4-88f8-f9d44a416272 ls1
```

Note This CLI command does not list VLAN-backed logical switches.

Check the Local Control Plane Status

For an overlay logical switch, check that the `netcpa` on the host is connected to the central control plane.

Prerequisites

Find the controller that the logical switch is on. See Check the CCP for the Logical Switch.

Procedure

1. SSH to the controller that the logical switch is on.

2 Run the following command and verify that the controller shows the hypervisors that are connected to this VNI.

```
get logical-switch 5000 connection-table
```

3 On the hypervisors, run the command `/bin/nsxcli` to start NSX CLI.

4 Run the following command to get the CCP sessions.

```
host1> get ccp-session
Session Index State Controller
Config  0     UP    10.33.74.163
L2      5000  UP    10.33.74.163
```

You should see a Config session on one of the CCP nodes in the CCP cluster. For every overlay logical switch, you should see an L2 session to one of the CCP nodes in the CCP cluster. For VLAN logical switches, there are no CCP connections.

Troubleshoot Config Session Issues

If the CCP config session is not up, check the status of MPA and `netcpa`.

Procedure

1 Make the following API call to see if MPA is connected to the NSX Manager.

```
GET https://<nsx-mgr>/api/v1/logical-ports/<logical-port-uuid>
```

2 On the hypervisor, run the command `/bin/nsxcli` to start NSX CLI.

3 Run the following command to get the node-uuid.

```
host1> get node-uuid
0c123dd4-8199-11e5-95e2-73cc1cd9b614
```

4 Run the following command to see if the NSX Manager pushed the CCP information to the host.

```
cat /etc/vmware/nsx/config-by-vsm.xml
```

5 If config-by-vsm.xml has CCP information, check if a transport node is configured on the hypervisor.

The NSX Manager sends the host certificate for the hypervisor in the transport node creation step. The CCP must have the host certificate before it accepts connections from the host.

6 Check the validity of the host certificate in `/etc/vmware/nsx/host-cert.pem`.

The certificate must be the same as the one that the NSX Manager has for the host.

7 Run the following command to check if the status of `netcpa`.

On ESXi:

```
/etc/init.d/netcpad status
```

On KVM:

```
/etc/init.d/nsx-agent status
```

8 Start or restart `netcpa`.

On ESXi, start `netcpa` if it is not running, or restart it if it is running.

```
/etc/init.d/netcpad start
```

```
/etc/init.d/netcpad restart
```

On KVM, start `netcpa` if it is not running, or restart it if it is running.

```
/etc/init.d/nsx-agent start
```

```
/etc/init.d/nsx-agent restart
```

9 If the config session is still not up, collect the technical support bundles and contact VMware support.

Troubleshoot L2 Session Issues

This applies to overlay logical switches only.

Procedure

1 On the hypervisor, run the command `/bin/nsxcli` to start NSX CLI.

2 Run the following command to see if the logical switch is present on the host.

```
host1> get logical-switches
```

3 Check that the state of the port is not `admin down`.

On ESXi, run `net-dvs` and look at the response. For example,

```
port 63eadf53-ff92-4a0e-9496-4200e99709ff:
com.vmware.port.extraConfig.opaqueNetwork.id = … <- this should match the logical switch UUID
com.vmware.port.opaque.network.id = …. <- this should match the logical switch UUID
com.vmware.port.opaque.network.type = nsx.LogicalSwitch , propType = RUNTIME
com.vmware.common.port.block = false, ... <- Make sure the value is false.
com.vmware.vswitch.port.vxlan = …
com.vmware.common.port.volatile.status = inUse ... <- make sure the value is inUse.
```

If the logical port ends up in the blocked state, collect the technical support bundles and contact VMware support. In the meantime, run the following command to get the DVS name:

```
[root@host1:~] net-dvs | grep nsx-switch
com.vmware.common.alias = nsx-switch , propType = CONFIG
```

Run the following command to unblock the port:

```
[root@host1:~] net-dvs -s com.vmware.common.port.block=false <DVS-NAME> -p <logical-port-ID>
```

On KVM, run `ovs-vsctl list interface` and verify that the interface with the corresponding VIF UUID is present and `admin_state` is up. You can see the VIF UUID in OVSDB in `external-ids:iface-id`.

Troubleshoot Dataplane Issues for an Overlay logical Switch

The steps in this section are for troubleshooting connectivity issues between VMs on different hypervisors through the overlay switch when the config and runtime states are normal.

If the VMs are on the same hypervisor, go to Troubleshoot ARP Issues for an Overlay Logical Switch.

Procedure

1. Run the following command on the controller that has the logical switch to see if CCP has the correct VTEP list:

   ```
   controller1> get logical-switch 5000 vtep
   ```

2. On each hypervisor, run the following NSX CLI command to see if it has the correct VTEP list:

   On ESXi:

   ```
   host1> get logical-switch <logical-switch-UUID> tep-table
   ```

   Alternatively, you can run the following shell command for the VTEP information:

   ```
   [root@host1:~] net-vdl2 -M vtep -s vds -n VNI
   ```

   On KVM:

   ```
   host1> get logical-switch <logical-switch-UUID or VNI> tep-table
   ```

3. Check to see if the VTEPs on the hypervisors can ping each other.

   At the ESXi shell prompt:

   ```
   host1> ping ++netstack=vxlan <remote-VTEP-IP>
   ```

At the KVM shell prompt:

```
host1> ping <remote-VTEP-IP>
```

If the VTEPs cannot ping each other,

a Make sure the transport VLAN specified when creating the transport node matches what the underlay expects. If you are using access ports in the underlay, the transport VLAN should be set to 0. If you are specifying a transport VLAN, the underlay switch ports that the hypervisors connect to should be configured to accept this VLAN in trunk mode.

b Check underlay connectivity.

4 Check if the BFD sessions between the VTEPs are up.

On ESXi, run `net-vdl2 -M bfd` and look at the response. For example,

```
BFD count: 1
==========================
Local IP: 192.168.48.35, Remote IP: 192.168.197.243, Local State: up, Remote State: up, Local
Diag: No Diagnostic, Remote Diag: No Diagnostic, minRx: 1000000, isDisabled: 0
```

On KVM, find the GENEVE interface to the remote IP.

```
ovs-vsctl list interface <GENEVE-interface-name>
```

If you don't know the interface name, run `ovs-vsctl find Interface type=geneve` to return all tunnel interfaces. Look for BFD information.

If you cannot find an GENEVEinterface to remote VTEP, check if nsx-agent is running and OVS integration bridge is connected to nsx-agent.

```
[root@host1 ~]# ovs-vsctl show
96c9e543-fc68-448a-9882-6e161c313a5b
  Manager "tcp:127.0.0.1:6632"
    is_connected: true
  Bridge nsx-managed
    Controller "tcp:127.0.0.1:6633"
      is_connected: true
    Controller "unix:ovs-l3d.mgmt"
      is_connected: true
    fail_mode: secure
```

Troubleshoot Dataplane Issues for a VLAN logical Switch

The steps in this section are for troubleshooting connectivity issues between VMs on different hypervisors through the configured VLAN on the underlay when the config and runtime states are normal.

If the VMs are on the same hypervisor and all the configuration and runtime states are normal, go to Troubleshoot ARP Issues for an Overlay Logical Switch.

Procedure

- Check that the underlay is configured for the VLAN for the logical switch in trunk mode.

  On ESXi, verify VLAN is configured on the logical port by running net-dvs and looking for the logical port. For example:

  ```
  port 63eadf53-ff92-4a0e-9496-4200e99709ff:
    com.vmware.common.port.volatile.vlan = VLAN 1000 propType = RUNTIME VOLATILE
  ```

  On KVM, the VLAN logical switch is configured as an openflow rule on integration bridge. In other words, for traffic received from the VIF, tag it with VLAN X and forward it on the patch port to the PIF bridge. Run `ovs-vsctl list interface` and verify the presence of the patch port between the NSX-managed bridge and the NSX-switch bridge.

Troubleshoot ARP Issues for an Overlay Logical Switch

The steps in this section are for troubleshooting where packets are being lost for an overlay switch.

For a VLAN-backed logical switch, go to Troubleshoot Packet Loss for a VLAN logical Switch or When ARP Is Resolved.

Before performing the following troubleshooting steps, run the command `arp -n` on each VM. If ARP is successfully resolved on both VMs, you do not need to perform the steps in this section. Instead, go to the next section Troubleshoot Packet Loss for a VLAN logical Switch or When ARP Is Resolved.

Procedure

- If both endpoints are ESXi and ARP proxy is enabled on the logical switch (only supported for overlay logical switches), check the ARP table on the CCP and the hypervisor.

  On the CCP:

  ```
  controller1> get logical-switch 5000 arp-table
  ```

  On the hypervisor, start NSX CLI and run the following command:

  ```
  host1> get logical-switch <logical-switch-UUID> arp-table
  ```

  Fetching the ARP table only tells us whether we have the correct ARP proxy state. If the ARP response is not received via proxy, or if the host is KVM and does not support ARP proxy, the datapath should broadcast the ARP request. There might be a problem with BUM traffic forwarding. Try the following steps:

  - If the replication mode for the logical switch is `MTEP`, change the replication mode to `SOURCE` for the logical switch from the NSX Manager GUI. This might fix the issue and ping will start working.
  - Add static ARP entries and see if the rest of the datapath works.

Troubleshoot Packet Loss for a VLAN logical Switch or When ARP Is Resolved

You can use the automated traceflow tool or manually trace the packets to troubleshoot packet loss.

To run the traceflow tool, from the NSX Manager GUI, navigate to **Tools > Traceflow**. For more information, see the *NSX-T Administration Guide*.

Procedure

- To manually trace the packets,

  On ESXi, run `net-stats -l` to get the switchport ID of the VIFs. If the source and destination VIFs are on the same hypervisor, run the following commands:

  ```
  pktcap-uw --switchport <src-switch-port-ID> --dir=0
  pktcap-uw --switchport <dst-switch-port-ID> --dir=1
  ```

  If the source and destination VIFs are on different hypervisors, on the hypervisor hosting the source VIF, run the following commands:

  ```
  pktcap-uw --switchport <src-switch-port-ID> --dir=0
  pktcap-uw --uplink <uplink-name> --dir=1
  ```

  On the hypervisor hosting the destination VIF, run the following commands:

  ```
  pktcap-uw --uplink <uplink-name> --dir=0
  pktcap-uw --switchport <dest-switch-port-ID> --dir=1
  ```

  On KVM, if the source and destination VIFs are on the same hypervisor, run the following command:

  ```
  ovs-dpctl dump-flows
  ```

Other Troubleshooting Scenarios 2

This section describes how to troubleshoot various error scenarios.

This chapter includes the following topics:

Failure to Add or Delete a Transport Node

You cannot delete or add a transport node.

Problem

The error occurs in the following scenario:

1. An ESXi host is a fabric node and a transport node.
2. The host is removed as a transport node. However, transport node deletion fails. The state of the transport node is `Orphaned`.
3. The host is removed as a fabric node immediately.
4. The host is added as a fabric node again.
5. The host is added as a transport node with a new transport zone and switch. This step results in the error `Failed/Partial Success`.

Cause

In step 2, if you wait for a few minutes, the transport node deletion will succeed because NSX Manager will retry the deletion. When you delete the fabric node immediately, NSX Manager cannot retry because the host is removed from NSX-T. This results in incomplete cleanup of the host, with the switch configuration still present, which causes step 5 to fail.

Solution

1. Delete all vmknics from vCenter Server on the host that are connected to the NSX-T switch.
2. Get the switch name using the `esxcfg-vswitch -l` CLI command. For example:

```
esxcfg-vswitch -l
Switch Name      Num Ports   Used Ports  Configured Ports  MTU     Uplinks
vSwitch0         1536        4           128               1500    vmnic0

  PortGroup Name        VLAN ID  Used Ports  Uplinks
  VM Network            0        0           vmnic0
  Management Network    0        1           vmnic0

Switch Name      Num Ports   Used Ports  Uplinks
nsxvswitch       1536        4
```

3. Delete the switch name using the `esxcfg-vswitch -d <switch-name> --dvswitch` CLI command. For example:

```
esxcfg-vswitch -d nsxvswitch --dvswitch
```

NSX Manager VM Is Degraded

NSX Manager that is deployed on a KVM host returns an error when running CLI commands such as `get service` and `get interface`.

Problem

The CLI command get service returns an error. For example,

```
nsx-manager-1> get service
% An error occurred while processing the service command
```

Other CLI commands might also return an error. The `get support-bundle` command indicates that the `/tmp` directory has become read-only. For example,

```
nsx-manager-1> get support-bundle file failed-to-get-service.tgz
% An error occurred while retrieving the support bundle: [Errno 30] Read-only file system:
'/tmp/tmpHzXF1u'
```

The `/var/log/messages-<timestamp>` log has the a message such as the following:

```
Nov 17 07:26:48 no kernel: NMI watchdog: BUG: soft lockup - CPU#5 stuck for 23s! [qemu-kvm:4386]
```

Cause

One or more file systems on the NSX Manager appliance were corrupted. Some possible causes are documented in https://access.redhat.com/solutions/22621.

To resolve the issue, you can repair the corrupt file systems or perform a restore from a backup.

Solution

1 Option 1: Repair the corrupt file systems. The following steps are specifically for NSX Manager running on a KVM host.

   a Run the `virsh destroy` command to stop the NSX Manager VM.

   b Run the `virt-rescue` command in write mode on the qcow2 image. For example,

   ```
   virt-rescue --rw -a nsx-unified-appliance-2.0.0.0.0.6522097.phadniss-p0-DK-to-DGo-on-rhel-prod_nsx_manager_1.qcow2
   ```

   c In the `virt-rescue` command prompt run the `e2fsck` command to fix the `tmp` file system. For example,

   ```
   <rescue> e2fsck /dev/nsx/tmp
   ```

   d If necessary, run the `e2fsck /dev/nsx/tmp` again until there are no more errors.

   e Restart NSX Manager with the `virsh start`.

2 Option 2: Perform a restore from a backup.

   For instructions, see the *NSX-T Administration Guide*.

NSX Agent Times Out Communicating with NSX Manager

In a large-scale environment with many transport nodes and VMs on ESXi hosts, NSX agents, which run on ESXi hosts, might time out when communicating with NSX Manager.

Problem

Some operations, such as when a VM vnic tries to attach to a logical switch, fail.
The `/var/run/log/nsxa.log` has messages such as:

```
level="ERROR" errorCode="MPA41542"] [MP_AddVnicAttachment] RPC call [0e316296-13-14] to NSX management plane timout
2017-05-15T05:32:13Z nsxa: [nsx@6876 comp="nsx-esx" subcomp="NSXA[VifHandlerThread:-2282640]" tid="1000017079" level="ERROR" errorCode="MPA42003"] [DoMpVifAttachRpc] MP_AddVnicAttachment() failed: RPC call to NSX management plane timout
```

Cause

In a large-scale environment, some operations might take longer than usual and fail because the default timeout values are exceeded.

Solution

1 Increase the NSX agent timeout value.

a On the ESXi host, stop the NSX agent with the following command:

```
/etc/init.d/nsxa stop
```

b Edit the file `/etc/vmware/nsxa/nsxa.json` and change the `vifOperationTimeout` value from 25 to, for example, 55.

```
"mp" : {
    /* timeout for VIF operation */
    "vifOperationTimeout" : 25,
```

Note This timeout value must be less than the hostd timeout value that you set in step 2.

c Start the NSX agent with the following command:

/etc/init.d/nsxa start

2 Increase the hostd timeout value.

a On the ESXi host, stop the hostd agent with the following command:

```
/etc/init.d/hostd stop
```

b Edit the file `/etc/vmware/hostd/config.xml`. Under `<opaqueNetwork>`, uncomment the entry for `<taskTimeout>` and change the value from 30 to, for example, 60.

```
<opaqueNetwork>
    <!-- maximum message size allowed in opaque network manager IPC, in bytes. -->
    <!-- <maxMsgSize> 65536 </maxMsgSize> -->
    <!-- maximum wait time for opaque network response -->
    <!-- <taskTimeout> 30 </taskTimeout> -->
```

c Start the hostd agent with the following command:

/etc/init.d/hostd start

Failure to Add an ESXi Host

You are not able to add an ESXi host to the NSX-T fabric.

Problem

From the NSX Manager GUI, adding an ESXi hosts fails with the error `File path of ... is claimed by multiple non-overlay VIBs`". The log file shows messages such as the following:

```
Failed to install software on host. Failed to install software on host. 10.172.120.60 :
java.rmi.RemoteException: [DependencyError] File path of '/usr/lib/vmware/vmkmod/nsx-vsip' is claimed
by multiple non-overlay VIBs
```

Cause

Some VIBs from a previous install are still on the host, probably because a clean uninstall did not occur.

Solution

1. From the error message, get the names of VIBs that are causing the failure.
2. Use ESXi commands to uninstall the VIBs.

Incorrect NSX Controller Status

Some controllers in an NSX Controller cluster report incorrect status for one of the controllers.

Problem

After a controller is powered off and on a number of times, the other controllers report that it is inactive when it is up and running.

Cause

An internal error involving the ZooKeeper module sometimes occurs when a controller is powered off and on and causes a communication failure between this controller and the other controllers in the cluster.

Solution

- Remove the controller node that is reported to be inactive from the cluster, remove the cluster configuration from the node and rejoin the node to the cluster. For more information, see the section "Replace a Member of the NSX Controller Cluster" in the *NSX-T Administration Guide*.

Management IPs on KVM VMs Not Reachable with IPFIX Enabled

When IPFIX is enabled on multiple VMs on a KVM host and the sampling rate is 100%, the management IPs on some of the VMs might intermittently be unreachable.

Problem

When you enable IPFIX for multiple VMs on the same host and you set the sampling rate to be 100%, there can be a large amount of IPFIX traffic. This can impact management traffic, causing the management IPs to be intermittently unreachable, even if the production traffic and management traffic go through different OVSes.

Cause

The workload is too stressful for the host and the VMs.

Solution

- Reduce the load on the host by reducing the number of VMs with IPFIX enabled or reducing the sampling rate.

KVM Transport Node Unreachable

KVM transport node's bridged interface loses connectivity to the physical interface.

Problem

Intermittently, a KVM transport node with active VMs and BFD tunnel loses management connectivity while idle. Rebooting the host might or might not resolve the problem.

Solution

- Run the following command:

```
ovs-ofctl add-flow nsx-switch.0 actions=NORMAL
```