

# NSX-T Troubleshooting Guide

Modified on 5 JUN 2018  
VMware NSX-T 2.2



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](mailto:docfeedback@vmware.com)

**VMware, Inc.**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

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

# Contents

NSX-T Troubleshooting Guide	4
<b>1 Troubleshooting Layer 2 Connectivity</b>	<b>5</b>
Check the NSX Manager and NSX Controller Cluster Status	5
Check the Logical Ports	6
Check the Transport Node Status	7
Check the Logical Switch Status	7
Check the CCP for the Logical Switch	8
Check the Local Control Plane Status	8
Troubleshoot Config Session Issues	9
Troubleshoot L2 Session Issues	10
Troubleshoot Dataplane Issues for an Overlay logical Switch	11
Troubleshoot Dataplane Issues for a VLAN logical Switch	12
Troubleshoot ARP Issues for an Overlay Logical Switch	13
Troubleshoot Packet Loss for a VLAN logical Switch or When ARP Is Resolved	13
<b>2 Troubleshooting Installation</b>	<b>15</b>
<b>3 Trouleshooting Routing</b>	<b>19</b>
<b>4 Troubleshooting Firewall</b>	<b>21</b>
Determining Firewall Rules that Apply on an ESXi Host	21
Determining Firewall Rules that Apply on a KVM Host	24
Firewall Packet Logs	25
<b>5 Logs and Services</b>	<b>27</b>
Log Files and Log Messages	27
Checking Services	29
<b>6 Other Troubleshooting Scenarios</b>	<b>32</b>
Failure to Add or Delete a Transport Node	32
Transport Node Takes About 5 Minutes to Connect to Another Controller	33
NSX Manager VM Is Degraded	33
NSX Agent Times Out Communicating with NSX Manager	35
Failure to Add an ESXi Host	36
Incorrect NSX Controller Status	36
Management IPs on KVM VMs Not Reachable with IPFIX Enabled	37

# 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](#)
- [Check the Logical Ports](#)
- [Check the Transport Node Status](#)
- [Check the Logical Switch Status](#)
- [Check the CCP for the Logical Switch](#)
- [Check the Local Control Plane Status](#)
- [Troubleshoot Config Session Issues](#)
- [Troubleshoot L2 Session Issues](#)
- [Troubleshoot Dataplane Issues for an Overlay logical Switch](#)
- [Troubleshoot Dataplane Issues for a VLAN logical Switch](#)
- [Troubleshoot ARP Issues for an Overlay Logical Switch](#)
- [Troubleshoot Packet Loss for a VLAN logical Switch or When ARP Is Resolved](#)

## 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/nsx-opsAgent status` to see if `opsAgent` 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 `opsAgent` is connected to the NSX Manager, check the `nsx-opsAgent` log. You will see the following message if `opsAgent` is connected to the NSX Manager.

```
Connected to mpa, cookie: ...
```

- To see if `opsAgent` is stuck processing `HostConfigMsg`, check the `nsx-opsAgent` 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 `opsAgent` 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 `nsx-opsAgent` 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-vd12 -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-vd12 -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 GENEVE interface 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
```

# Troubleshooting Installation

This section provides information about troubleshooting installation issues.

## Basic Infrastructure Services

The following services must be running on the appliances and hypervisors, also on vCenter Server if it is used as a compute manager.

- NTP
- DNS

Make sure that firewall is not blocking traffic between NSX-T components and hypervisors. Make sure that the required ports are open between the components.

To flush the DNS cache on the NSX Manager, SSH as root to the manager and run the following command:

```
root@nsx-mgr-01:~# /etc/init.d/resolvconf restart
[ ok ] Restarting resolvconf (via systemctl): resolvconf.service.
```

You can then check the DNS configuration file.

```
root@nsx-mgr-01:~# cat /etc/resolv.conf
# Dynamic resolv.conf(5) file for glibc resolver(3) generated by resolvconf(8)
#     DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
nameserver 192.168.253.1
search mgt.sg.lab
```

## Checking Communication from Host to Controller and Manage

On an ESXi host using NSX-T CLI commands:

```
esxi-01.corp.local> get managers
- 192.168.110.19 Connected

esxi-01.corp.local> get controllers
Controller IP  Port  SSL      Status      Is Physical Master  Session State  Controller FQDN
192.168.110.16 1234 enabled  connected   true           up             NA
```

On a KVM host using NSX-T CLI commands:

```
kvm-01> get managers
- 192.168.110.19 Connected

kvm-01> get controllers
Controller IP  Port  SSL      Status      Is Physical Master  Session State  Controller FQDN
192.168.110.16 1234 enabled  connected   true           up             NA
```

On an ESXi host using host CLI commands:

```
[root@esxi-01:~] esxcli network ip connection list | grep 1234
tcp          0      0 192.168.110.53:42271      192.168.110.16:1234
ESTABLISHED 67702 newreno netcpa
[root@esxi-01:~]
[root@esxi-01:~] esxcli network ip connection list | grep 5671
tcp          0      0 192.168.110.253:11721    192.168.110.19:5671  ESTABLISHED 2103688
newreno mpa
tcp          0      0 192.168.110.253:30977    192.168.110.19:5671  ESTABLISHED 2103688
newreno mpa
```

On a KVM host using host CLI commands:

```
root@kvm-01:/home/vmware# netstat -nap | grep 1234
tcp          0      0 192.168.110.55:53686     192.168.110.16:1234  ESTABLISHED 2554/netcpa
root@kvm-01:/home/vmware#
root@kvm-01:/home/vmware#
root@kvm-01:/home/vmware# netstat -nap | grep 5671
tcp          0      0 192.168.110.55:50108     192.168.110.19:5671  ESTABLISHED 2870/mpa
tcp          0      0 192.168.110.55:50110     192.168.110.19:5671  ESTABLISHED 2870/mpa

root@kvm-01:/home/vmware# tcpdump -i ens32 port 1234 | grep kvm-01
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on ens32, link-type EN10MB (Ethernet), capture size 262144 bytes
<truncated output>
03:46:27.040461 IP nsxcontroller01.corp.local.1234 > kvm-01.corp.local.38754: Flags [P.], seq
3315301231:3315301275, ack 2671171555, win 323, length 44
03:46:27.040509 IP kvm-01.corp.local.38754 > nsxcontroller01.corp.local.1234: Flags [.), ack 44, win
1002, length 0
```



```

^C
<truncated output>
root@kvm-01:/home/vmware#

root@kvm-01:/home/vmware# tcpdump -i ens32 port 5671 | grep kvm-01
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on ens32, link-type EN10MB (Ethernet), capture size 262144 bytes
03:51:16.802934 IP kvm-01.corp.local.58954 > nsxmgr01.corp.local.amqps: Flags [P.], seq 1153:1222, ack
1790, win 259, length 69
03:51:16.823328 IP nsxmgr01.corp.local.amqps > kvm-01.corp.local.58954: Flags [P.], seq 1790:1891, ack
1222, win 254, length 101
^C
<truncated output>

```

## Host Registration Failure

If NSX-T uses the wrong IP address, host registration will fail. This can happen when a host has multiple IP addresses. Trying to delete the transport node leaves it in the Orphaned state. To resolve the issue:

- Go to **Fabric > Nodes > Hosts**, edit the host and remove all IP addresses except the management one.
- Click on the errors and select **Resolve**.

## KVM Host Issues

KVM host issues are sometimes caused by not enough disk space. The `/boot` directory can fill up quickly and cause errors such as:

- Failed to install software on host
- No space left on device

You can run the command `df -h` to check available storage. If the `/boot` directory is at 100%, you can do the following:

- Run `sudo dpkg --get-selections | grep ^ii` to see all the kernels installed.
- Run `uname -r` to see your currently running kernel. Do not remove this kernel (linux-image).
- Use `apt-get purge` to remove images you don't need anymore. For example, run `sudo apt-get purge linux-image-3.13.0-32-generic linux-image-3.13.0-33-generic`.
- Reboot the host.
- In NSX Manager, check the errors and select **Resolve**.
- Make sure the VMs are powered on.

## Configuration Error when Deploying an Edge VM

After deploying an Edge VM, NSX Manager shows the VM's status as **configuration error**. The manager log has a message similar to the following:

```
nsx-manager NSX - FABRIC [nsx@6876 comp="nsx-manager" errorCode="MP16027" subcomp="manager"] Edge  
758ad396-0754-11e8-877e-005056abf715 is not ready for configuration error occurred, error detail is  
NSX Edge configuration has failed. The host does not support required cpu features: ['aes'].
```

Restarting the edge datapath service and then the VM should resolve the issue.

## Force Removing a Transport Node

You can remove a transport node that is stuck in the Orphaned state by making the following API call:

```
DELETE https://<NSX Manager>/api/v1/transport-nodes/<TN ID>?force=true
```

NSX Manager will not do any validations as to whether you have any active VMs running on the host. You are responsible for deleting the N-VDS and VIBs. If you have the node added through Compute Manager, delete the Compute Manager first and then delete the node. The transport node will be deleted as well.

# Trouleshooting Routing

NSX-T has built-in tools to troubleshoot routing issues.

## Traceflow

You can use Traceflow to inspect the flow of packets. You can see delivered, dropped, received, and forwarded packets. If a packet is dropped, a reason is displayed. For example, a packet can be dropped because of a firewall rule.

## Checking Routing Tables

To see the routing table on a service router, run the following commands:

```
edge01> get logical-router
Logical Route
UUID                               VRF   LR-ID  Name                               Type
Ports
736a80e3-23f6-5a2d-81d6-bbefb2786666  0     0      SR-t0-router                       TUNNEL              3
c9393d0c-1fcf-4c34-889d-2da1eeee25b8  1     10     SR-t0-router                       SERVICE_ROUTER_TIER0 5
9333c94e-5938-46b4-8c7d-5e6ac2c8b7b5  2     8      DR-t1-router01                     DISTRIBUTED_ROUTER_TIER1 6
c91eb7c5-0297-4fed-9c22-b96df1c9b80f  3     9      DR-t0-router                       DISTRIBUTED_ROUTER_TIER0 4

edge01> vrf 1
edge01(tier0_sr)> get route

Flags: c - connected, s - static, b - BGP, ns - nsx_static
nc - nsx_connected, rl - router_link, t0n: Tier0-NAT, t1n: Tier1-NAT
t1l: Tier1-LB VIP, t1s: Tier1-LB SNAT

Total number of routes: 25

b   10.10.20.0/24      [20/0]      via 192.168.140.1
b   10.10.30.0/24      [20/0]      via 192.168.140.1
b   10.20.20.0/24      [20/0]      via 192.168.140.1
b   10.20.30.0/24      [20/0]      via 192.168.140.1
b   30.0.0.0/8         [20/0]      via 192.168.140.1
rl  100.64.80.0/31     [0/0]       via 169.254.0.1
rl  100.64.80.2/31     [0/0]       via 169.254.0.1
rl  100.64.80.4/31     [0/0]       via 169.254.0.1
<TRUNCATED OUTPUT>
```

```

b    192.168.200.0/24    [20/0]    via 192.168.140.1
b    192.168.210.0/24    [20/0]    via 192.168.140.1
b    192.168.220.0/24    [20/0]    via 192.168.140.1
b    192.168.230.0/24    [20/0]    via 192.168.140.1
b    192.168.240.0/24    [20/0]    via 192.168.140.1

```

To get the IP address of interfaces, run the following command:

```

edge01(tier0_sr)> get interfaces
Logical Router
UUID                               VRF  LR-ID  Name                Type
c9393d0c-1fcf-4c34-889d-2da1eeee25b8  1    10     SR-t0-router        SERVICE_ROUTER_TIER0
interfaces
  interface    : 977ac2eb-8ab7-40e9-8abe-782a438c749a
  ifuid        : 285
  name         : uplink01
  mode         : lif
  IP/Mask      : 192.168.140.3/24
  MAC          : 00:50:56:b5:d5:64
  LS port      : 14391f86-efef-4e3d-98c3-f291c17d13f8
  urpf-mode    : STRICT_MODE
  admin        : up
  MTU          : 1600

  interface    : 6af81d72-4d32-5f66-b7ae-403e617290e5
  ifuid        : 270
  mode         : blackhole

  interface    : 015e709d-6079-5c19-9556-8be2e956f775
  ifuid        : 269
  mode         : cpu

  interface    : 3f40f838-eb8a-4f35-854c-ea8bb872dc47
  ifuid        : 272
  name         : bp-sr0-port
  mode         : lif
  IP/Mask      : 169.254.0.2/28
  MAC          : 02:50:56:56:53:00
  VNI          : 25489
  LS port      : 770a208d-27fa-4f8d-afad-a9c41ca6295b
  urpf-mode    : NONE
  admin        : up
  MTU          : 1500

  interface    : 00003300-0000-0000-0000-00000000000a
  ifuid        : 263
  mode         : loopback
  IP/Mask      : 127.0.0.1/8

```

## Advertising T1 Routes

You must advertise T1 routes so that they are visible on T0 router and upwards. There are different types of routes that you can advertise: NSX Connected, NAT, Static, LB VIP, and LB SNAT.

# 4

## Troubleshooting Firewall

This section provides information about troubleshooting firewall issues.

This chapter includes the following topics:

- [Determining Firewall Rules that Apply on an ESXi Host](#)
- [Determining Firewall Rules that Apply on a KVM Host](#)
- [Firewall Packet Logs](#)

### Determining Firewall Rules that Apply on an ESXi Host

To troubleshoot firewall issues with an ESXi host, you can look at the firewall rules that apply on the host.

Get the list of dvfilters on the ESXi host:

```
[root@esxi-01:~] summarize-dvfilter
<TRUNCATED OUTPUT>
world 70181 vmm0:app-01a vcUuid:'50 35 9c 70 18 8e 99 1d-3c f9 8e cc 6b 27 4c 6f'
  port 50331655 app-01a.eth0
  vNic slot 2
  name: nic-70181-eth0-vmware-sfw.2
  agentName: vmware-sfw
  state: IOChain Attached
  vmState: Detached
  failurePolicy: failClosed
  slowPathID: none
  filter source: Dynamic Filter Creation
world 70179 vmm0:web-02a vcUuid:'50 35 2b f3 4a 4b 10 83-54 72 50 f7 25 10 d8 64'
  port 50331656 web-02a.eth0
  vNic slot 2
  name: nic-70179-eth0-vmware-sfw.2
  agentName: vmware-sfw
  state: IOChain Attached
  vmState: Detached
  failurePolicy: failClosed
  slowPathID: none
  filter source: Dynamic Filter Creation
```

Find a dvfilter for a specific VM:

```
[root@esxi-01:~] summarize-dvfilter | less -p web

world 70179 vmm0:web-02a vcUuid:'50 35 2b f3 4a 4b 10 83-54 72 50 f7 25 10 d8 64'
port 50331656 web-02a.eth0
vNic slot 2
name: nic-70179-eth0-vmware-sfw.2
agentName: vmware-sfw
state: IOChain Attached
vmState: Detached
failurePolicy: failClosed
slowPathID: none
filter source: Dynamic Filter Creation
.
.
.
```

Determine firewall rules that apply to a specific dvfilter (in this example, `nic-70227-eth0-vmware-sfw.2` is the dvfilter name):

```
[root@esxi-02:~] vsiopiectl getrules -f nic-70227-eth0-vmware-sfw.2
ruleset mainrs {
rule 3072 at 1 inout protocol tcp from any to addrset 48822ec3-2670-497b-82f9-524618c16877 port 443
accept with log;
rule 3072 at 2 inout protocol tcp from any to addrset 48822ec3-2670-497b-82f9-524618c16877 port 80
accept with log;
rule 3074 at 3 inout protocol tcp from addrset 48822ec3-2670-497b-82f9-524618c16877 to addrset
8b9e75e7-bc62-4d7f-9a58-a872f393448e port 8443 accept with log;
rule 3074 at 4 inout protocol tcp from addrset 48822ec3-2670-497b-82f9-524618c16877 to addrset
8b9e75e7-bc62-4d7f-9a58-a872f393448e port 22 accept with log;
rule 3075 at 5 inout protocol tcp from addrset 8b9e75e7-bc62-4d7f-9a58-a872f393448e to addrset
b695c8df-9894-4068-a5e7-5504fe48d459 port 3306 accept with log;
rule 3076 at 6 inout protocol tcp from ip 192.168.110.10 to addrset rdst3076 port 443 accept with log;
rule 3076 at 7 inout protocol icmp typecode 8:0 from ip 192.168.110.10 to addrset rdst3076 accept with
log;
rule 3076 at 8 inout protocol tcp from ip 192.168.110.10 to addrset rdst3076 port 22 accept with log;
rule 3076 at 9 inout protocol tcp from ip 192.168.110.10 to addrset rdst3076 port 80 accept with log;
rule 2 at 10 inout protocol any from any to any accept with log;
}

ruleset mainrs_L2 {
rule 1 at 1 inout ethertype any stateless from any to any accept;
}
}
```

Get the list of address sets used in a specific dvfilter:

```
[root@esxi-02:~] vsiopiectl getaddrsets -f nic-70227-eth0-vmware-sfw.2
addrset 48822ec3-2670-497b-82f9-524618c16877 {
ip 172.16.10.13,
mac 52:54:00:42:4d:38,
}
addrset 8b9e75e7-bc62-4d7f-9a58-a872f393448e {
```

```

}
addrset b695c8df-9894-4068-a5e7-5504fe48d459 {
ip 172.16.30.11,
mac 52:54:00:64:0e:4f,
}
addrset rdst3076 {
ip 172.16.10.13,
ip 172.16.30.11,
mac 52:54:00:42:4d:38,
mac 52:54:00:64:0e:4f,
}

```

Check the flows through a specific dvfilter:

```

[root@esxi-02:~] vsipioctl getflows -f nic-75360-eth0-vmware-sfw.2
Count retrieved from kernel active(L3,L4)=20, active(L2)+inactive(L3,L4)=0, drop(L2,L3,L4)=0
a5d914f7a5b85fe5 Active tcp 0800 IN 3076 0 0 192.168.110.10:Unknown(51281) -> 172.16.10.11:ssh(22)
513 FINWAIT2:FINWAIT2 4304 5177 34 33
a5d914f7a5b86001 Active tcp 0800 OUT 2 0 0 172.16.10.11:http(80) -> 100.64.80.1:Unknown(60006) 457
SYNSENT:CLOSED 56 819 1 1
a5d914f7a5b86006 Active igmp 0800 IN 2 0 0 0.0.0.0 -> 224.0.0.1 36 0 1 0
a5d914f7a5b86011 Active tcp 0800 IN 3072 0 0 100.64.80.1:Unknown(60098) -> 172.16.10.11:http(80) 320
FINWAIT2:FINWAIT2 413 5411 9 6
a5d914f7a5b86012 Active tcp 0800 OUT 3074 0 0 172.16.10.11:Unknown(46001) ->
172.16.20.11:Unknown(8443) 815 FINWAIT2:FINWAIT2 7418 1230 10 9
a5d914f7a5b86013 Active udp 0800 OUT 2 0 0 172.16.10.11:Unknown(40080) -> 192.168.110.10:domain(53)
268 140 2 2
a5d914f7a5b86014 Active udp 0800 OUT 2 0 0 172.16.10.11:Unknown(59251) -> 192.168.110.10:domain(53)
268 140 2 2
a5d914f7a5b86015 Active ipv6-icmp 86dd OUT 2 0 0 fe80::250:56ff:feb5:a60e -> ff02::1:ff62:5ed4 135 0
0 72 0 1
a5d914f7a5b86016 Active ipv6-icmp 86dd OUT 2 0 0 fe80::250:56ff:feb5:a60e -> ff02::1:ff62:5ed4 135 0
0 72 0 1
a5d914f7a5b86017 Active tcp 0800 IN 3072 0 0 100.64.80.1:Unknown(60104) -> 172.16.10.11:http(80) 320
FINWAIT2:FINWAIT2 413 5451 9 7
a5d914f7a5b86018 Active tcp 0800 OUT 3074 0 0 172.16.10.11:Unknown(46002) ->
172.16.20.11:Unknown(8443) 815 TIMEWAIT:TIMEWAIT 7314 1230 8 9
a5d914f7a5b86019 Active tcp 0800 IN 3072 0 0 100.64.80.1:Unknown(60110) -> 172.16.10.11:http(80) 320
FINWAIT2:FINWAIT2 373 5451 8 7
a5d914f7a5b8601a Active tcp 0800 OUT 3074 0 0 172.16.10.11:Unknown(46003) ->
172.16.20.11:Unknown(8443) 815 FINWAIT2:FINWAIT2 7418 1230 10 9
a5d914f7a5b8601b Active tcp 0800 IN 3072 0 0 100.64.80.1:Unknown(60114) -> 172.16.10.11:http(80) 328
TIMEWAIT:TIMEWAIT 413 5451 9 7
a5d914f7a5b8601c Active tcp 0800 OUT 3074 0 0 172.16.10.11:Unknown(46004) ->
172.16.20.11:Unknown(8443) 815 TIMEWAIT:TIMEWAIT 7262 1218 7 9
a5d914f7a5b8601d Active tcp 0800 OUT 2 0 0 172.16.10.11:http(80) -> 100.64.80.1:Unknown(60060) 457
SYNSENT:CLOSED 56 819 1 1
a5d914f7a5b8601e Active tcp 0800 IN 3072 0 0 100.64.80.1:Unknown(60120) -> 172.16.10.11:http(80) 320
TIMEWAIT:TIMEWAIT 373 5411 8 6
a5d914f7a5b8601f Active tcp 0800 OUT 3074 0 0 172.16.10.11:Unknown(46005) ->
172.16.20.11:Unknown(8443) 815 FINWAIT2:FINWAIT2 7418 1230 10 9

```

```
a5d914f7a5b86020 Active tcp 0800 IN 3072 0 0 100.64.80.1:Unknown(60126) -> 172.16.10.11:http(80) 229
EST:EST 173 5371 3 5
a5d914f7a5b86021 Active tcp 0800 OUT 3074 0 0 172.16.10.11:Unknown(46006) ->
172.16.20.11:Unknown(8443) 815 FINWAIT2:FINWAIT2 7418 1230 10 9
```

## Determining Firewall Rules that Apply on a KVM Host

To troubleshoot firewall issues with a KVM host, you can look at the firewall rules that apply on the host.

Get the list of VIFs that are subject to firewall rules on the KVM host:

```
# ovs-appctl -t /var/run/openvswitch/nsxa-ctl dfw/vif
Vif ID      : da95fc1e-65fd-461f-814d-d92970029bf0
Port name   : db-01a-eth0
Port number : 2
```

If the output is empty, look for connectivity issues between the node and the controllers.

Get the list of rules applied to a specific VIF (in this example, da95fc1e-65fd-461f-814d-d92970029bf0 is the VIF ID):

```
# ovs-appctl -t /var/run/vmware/nsx-agent/nsxa-ctl dfw/rules da95fc1e-65fd-461f-814d-d92970029bf0
Distributed firewall status: enabled

Vif ID      : da95fc1e-65fd-461f-814d-d92970029bf0
ruleset d035308b-cb0d-4e7e-aae5-a428b461db46 {
  rule 3072 inout protocol tcp from any to addrset 48822ec3-2670-497b-82f9-524618c16877 port 443 accept
with log;
  rule 3072 inout protocol tcp from any to addrset 48822ec3-2670-497b-82f9-524618c16877 port 80 accept
with log;
  rule 3074 inout protocol tcp from addrset 48822ec3-2670-497b-82f9-524618c16877 to addrset 8b9e75e7-
bc62-4d7f-9a58-a872f393448e port 8443 accept with log;
  rule 3074 inout protocol tcp from addrset 48822ec3-2670-497b-82f9-524618c16877 to addrset 8b9e75e7-
bc62-4d7f-9a58-a872f393448e port 22 accept with log;
  rule 3075 inout protocol tcp from addrset 8b9e75e7-bc62-4d7f-9a58-a872f393448e to addrset
b695c8df-9894-4068-a5e7-5504fe48d459 port 3306 accept with log;
}

ruleset 3027fed3-60b1-483e-aa17-c28719275704 {
  rule 3076 inout protocol tcp from 192.168.110.10 to addrset b695c8df-9894-4068-a5e7-5504fe48d459 port
443 accept with log;
  rule 3076 inout protocol icmp type 8 code 0 from 192.168.110.10 to addrset b695c8df-9894-4068-
a5e7-5504fe48d459 accept with log;
  rule 3076 inout protocol tcp from 192.168.110.10 to addrset b695c8df-9894-4068-a5e7-5504fe48d459 port
22 accept with log;
  rule 3076 inout protocol tcp from 192.168.110.10 to addrset b695c8df-9894-4068-a5e7-5504fe48d459 port
80 accept with log;
  rule 3076 inout protocol tcp from 192.168.110.10 to addrset 8b9e75e7-bc62-4d7f-9a58-a872f393448e port
443 accept with log;
  rule 3076 inout protocol icmp type 8 code 0 from 192.168.110.10 to addrset 8b9e75e7-bc62-4d7f-9a58-
a872f393448e accept with log;
  rule 3076 inout protocol tcp from 192.168.110.10 to addrset 8b9e75e7-bc62-4d7f-9a58-a872f393448e port
22 accept with log;
```



```

rule 3076 inout protocol tcp from 192.168.110.10 to addrset 8b9e75e7-bc62-4d7f-9a58-a872f393448e port
80 accept with log;
rule 3076 inout protocol tcp from 192.168.110.10 to addrset 48822ec3-2670-497b-82f9-524618c16877 port
443 accept with log;
rule 3076 inout protocol icmp type 8 code 0 from 192.168.110.10 to addrset
48822ec3-2670-497b-82f9-524618c16877 accept with log;
rule 3076 inout protocol tcp from 192.168.110.10 to addrset 48822ec3-2670-497b-82f9-524618c16877 port
22 accept with log;
rule 3076 inout protocol tcp from 192.168.110.10 to addrset 48822ec3-2670-497b-82f9-524618c16877 port
80 accept with log;
}

ruleset 5e9bdcb3-adba-4f67-a680-5e6ed5b8f40a {
rule 2 inout protocol any from any to any accept with log;
}

ruleset ddf93011-4078-4006-b8f8-73f979d7a717 {
rule 1 inout ethertype any stateless from any to any accept;
}

```

Get the list of address sets used in a specific VIF:

```

# ovs-appctl -t /var/run/vmware/nsx-agent/nsxa-ctl dfw/addrsets da95fc1e-65fd-461f-814d-d92970029bf0
48822ec3-2670-497b-82f9-524618c16877 {
mac 52:54:00:42:4d:38,
ip 172.16.10.13,
}

8b9e75e7-bc62-4d7f-9a58-a872f393448e {
}

b695c8df-9894-4068-a5e7-5504fe48d459 {
mac 52:54:00:64:0e:4f,
ip 172.16.30.11,
}

```

Check connections through the Linux Conntrack module. In this example, we look for flows between two specific IP addresses.

```

# ovs-appctl -t ovs-l3d conntrack/show | grep 192.168.110.10 | grep 172.16.10.13
ACTIVE
icmp,orig=(src=192.168.110.10,dst=172.16.10.13,id=1,type=8,code=0),reply=(src=172.16.10.13,dst=192.168.
110.10,id=1,type=0,code=0),start=2018-03-26T04:43:28.325,id=3122159040,zone=23119,status=SEEN_REPLY|
CONFIRMED,timeout=29,mark=3076,labels=0x1f

```

## Firewall Packet Logs

If logging is enabled for firewall rules, you can look at the firewall packet logs to troubleshoot issues.

The log file is `/var/log/dfwpktlogs.log` for both ESXi and KVM hosts.

```
# tail -f /var/log/dfwpktlogs.log
2018-03-27T10:23:35.196Z INET TERM 3072 IN TCP FIN 100.64.80.1/60688->172.16.10.11/80 8/7 373/5451
2018-03-27T10:23:35.196Z INET TERM 3074 OUT TCP FIN 172.16.10.11/46108->172.16.20.11/8443 8/9 1178/7366
2018-03-27T10:23:35.196Z INET TERM 3072 IN TCP RST 100.64.80.1/60692->172.16.10.11/80 9/6 413/5411
2018-03-27T10:23:35.196Z INET TERM 3074 OUT TCP RST 172.16.10.11/46109->172.16.20.11/8443 9/7 1218/7262
2018-03-27T10:23:37.442Z 71d32787 INET match PASS 3074 IN 60 TCP 172.16.10.12/35770->172.16.20.11/8443
S
2018-03-27T10:23:38.492Z INET match PASS 2 OUT 1500 TCP 172.16.10.11/80->100.64.80.1/60660 A
2018-03-27T10:23:39.934Z INET match PASS 3072 IN 52 TCP 100.64.80.1/60720->172.16.10.11/80 S
2018-03-27T10:23:39.944Z INET match PASS 3074 OUT 60 TCP 172.16.10.11/46114->172.16.20.11/8443 S
2018-03-27T10:23:39.944Z 71d32787 INET match PASS 3074 IN 60 TCP 172.16.10.11/46114->172.16.20.11/8443
S
2018-03-27T10:23:42.449Z 71d32787 INET match PASS 3074 IN 60 TCP 172.16.10.12/35771->172.16.20.11/8443
S
2018-03-27T10:23:44.712Z INET TERM 3074 IN TCP RST 172.16.10.11/46109->172.16.20.11/8443 9/7 1218/7262
2018-03-27T10:23:44.712Z INET TERM 3074 IN TCP FIN 172.16.10.12/35766->172.16.20.11/8443 9/10 1233/7418
2018-03-27T10:23:44.712Z INET TERM 3074 IN TCP FIN 172.16.10.11/46110->172.16.20.11/8443 9/9 1230/7366
2018-03-27T10:23:44.712Z INET TERM 3074 IN TCP FIN 172.16.10.12/35767->172.16.20.11/8443 9/10 1233/7418
2018-03-27T10:23:44.939Z INET match PASS 3072 IN 52 TCP 100.64.80.1/60726->172.16.10.11/80 S
2018-03-27T10:23:44.957Z INET match PASS 3074 OUT 60 TCP 172.16.10.11/46115->172.16.20.11/8443 S
2018-03-27T10:23:44.957Z 71d32787 INET match PASS 3074 IN 60 TCP 172.16.10.11/46115->172.16.20.11/8443
S
2018-03-27T10:23:45.480Z INET TERM 2 OUT TCP TIMEOUT 172.16.10.11/80->100.64.80.1/60528 1/1 1500/56
```

# Logs and Services

Logs can be helpful in many troubleshooting scenarios. Checking the status of services is also important.

This chapter includes the following topics:

- [Log Files and Log Messages](#)
- [Checking Services](#)

## Log Files and Log Messages

Log files are available in every NSX-T component.

### NSX-T Appliance Logs

The appliance logs are available through `nsxcli` or host CLI.

Log in as admin and run the following `nsxcli` command:

```
get log-file <auth.log | http.log | kern.log | manager.log | node-mgmt.log | syslog> [follow]
```

Log in as root and run the following host commands:

```
# cd /var/log
# ls -l
-rw-r----- 1 syslog adm          903319 Mar 23 15:24 nsx-audit.log
-rw-r----- 1 syslog adm        1053425 Mar 23 12:17 nsx-audit.log.1
-rw-r----- 1 root   adm        1048850 Mar 22 20:22 nsx-audit.log.2
-rw-r----- 1 root   adm        1054444 Mar 21 16:50 nsx-audit.log.3
-rw-r----- 1 root   adm        1059856 Mar 21 14:01 nsx-audit.log.4
-rw-r--r--  1 root   root         15539 Mar 20 16:21 boot.log
-rw-r----- 1 syslog adm         462193 Mar 23 15:24 iptables.log
-rw-r----- 1 syslog adm       10490922 Mar 23 14:48 iptables.log.1
-rw-r----- 1 syslog adm       12664378 Mar 23 15:24 syslog
-rw-r----- 1 syslog adm       33558646 Mar 23 05:12 syslog.1
```

You can use various commands, such as `tac`, `tail`, `grep`, and more, to display or search the files.

### NSX Manager Syslog

Syslog provides the main source of information when troubleshooting.

Log in as admin and run the following `nsxcli` command:

```
nsxmgr01> get log-file syslog follow
<182>1 2018-03-27T14:28:38.723227Z nsxmgr01 NSX 2645 - [nsx@6876 comp="nsx-manager" subcomp="mpa"
tid="7073" level="INFO"] Processing message for known client HostNodeStatusVertical.
HostNodeStatusVertical d270a9d2-1f9a-4dfd-9218-c27d00f10c11:4600
<182>1 2018-03-27T14:28:38.723532Z nsxmgr01 NSX 834 - [nsx@6876 comp="nsx-manager" subcomp="agg-
service" tid="989" level="INFO"] [SendMsg] Sent message to MPA for app (HostNodeStatusVertical)
<182>1 2018-03-27T14:28:38.725318Z nsxmgr01 NSX 2645 - [nsx@6876 comp="nsx-manager" subcomp="mpa"
tid="7074" level="INFO"] [SendAppMessage] Finishing sending message id: d270a9d2-1f9a-4dfd-9218-
c27d00f10c11:4600 for client HostNodeStatusVertical
<182>1 2018-03-27T14:28:38.725384Z nsxmgr01 NSX 2645 - [nsx@6876 comp="nsx-manager" subcomp="mpa"
tid="7074" level="INFO"] Successfully sent app data to broker
<182>1 2018-03-27T14:28:39.550Z nsxmgr01 NSX - - [nsx@6876 comp="nsx-manager" subcomp="manager"]
Current delta sync queue has #4 transaction(s)
<180>1 2018-03-27T14:28:39.550Z nsxmgr01 NSX - - [nsx@6876 comp="nsx-manager" subcomp="manager"]
Calculated delta batch is null
```

Log in as root and run the following host commands. In this example, we are interested in a specific host node.

```
root@nsxmgr01:~# tac /var/log/syslog | grep 63805f13-3ca9-4100-8107-2cc89a45b10d | more
<182>1 2018-03-27T07:41:30.042Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Deleting Host virtual machines container for host HostModel/63805f13-3ca9-4100-8107-2cc89a45b10d
<182>1 2018-03-27T07:41:30.041Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Deleting hostModel and its relationships in ES: 63805f13-3ca9-4100-8107-2cc89a45b10d .
<182>1 2018-03-27T07:41:30.020Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Deleting virtual machines for host 63805f13-3ca9-4100-8107-2cc89a45b10d .
<182>1 2018-03-27T07:41:02.824Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
TransportNode of the FabricNode [Node/63805f13-3ca9-4100-8107-2cc89a45b10d] has already removed, no
cleaning required.
<182>1 2018-03-27T07:41:02.822Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
FabricNode [Node/63805f13-3ca9-4100-8107-2cc89a45b10d] has been removed, try to clean up TransportNode.
<182>1 2018-03-27T07:41:02.800Z nsxmgr01 NSX - - [nsx@6876 audit="true" comp="nsx-manager"
subcomp="manager"] Username:'admin', ModuleName:'Batch', Operation:'RegisterBatchRequest', Operation
status:'success', New value: [{"atomic":false} {"requests":
[{"method":"DELETE","uri":"/v1/fabric/nodes/63805f13-3ca9-4100-8107-2cc89a45b10d?
unprepare_host=false"}], "continue_on_error":true}]
<182>1 2018-03-27T07:41:02.799Z nsxmgr01 NSX - FABRIC [nsx@6876 audit="true" comp="nsx-manager"
entId="63805f13-3ca9-4100-8107-2cc89a45b10d" reqId="e61e6381-0c02-4a3f-a8b6-364d8
951994e" subcomp="manager"] Username:'admin', ModuleName:'FabricNode', Operation:'DeleteNode',
Operation status:'success', New value: [{"63805f13-3ca9-4100-8107-2cc89a45b10d" {"u
nprepare_host":false}]
<182>1 2018-03-27T07:41:02.785Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Deleted HostNode 63805f13-3ca9-4100-8107-2cc89a45b10d
<182>1 2018-03-27T07:41:02.777Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Deleting HostNode 63805f13-3ca9-4100-8107-2cc89a45b10d .
<182>1 2018-03-27T07:41:02.777Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Deleted all capabilities for node HostModel/63805f13-3ca9-4100-8107-2cc89a45b10d
<182>1 2018-03-27T07:41:02.772Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Delete capabilities of the node 63805f13-3ca9-4100-8107-2cc89a45b10d
<182>1 2018-03-27T07:41:02.764Z nsxmgr01 NSX - FABRIC [nsx@6876 comp="nsx-manager" subcomp="manager"]
Unregister HostNode 63805f13-3ca9-4100-8107-2cc89a45b10d from messaging framework.
```

## Hypervisor Logs

Hypervisor logs are not available through `nsxccli`. You can access them as root in the directory `/var/log` for both ESXi and KVM hosts. Looking at the file `syslog` on KVM or `syslog.log` on ESXi is always a good starting point.

## Syslog Facilities and Message IDs

NSX-T produces the following logs:

- regular logs (facility `local6`, which has a numerical value of 22)
- audit logs (facility `local7`, which has a numerical value of 23)

All API calls trigger an audit log. When configuring syslog and specifying facilities, be sure to include `local6` and `local7`

The `messageid` parameter can be used in the `set logging-server` command to filter which log messages are sent to a specific logging server. Configuring `messageid` incorrectly might cause certain logs not to be received on the syslog server.

## Syslog Configuration Examples

The first command below sends all info-level logs to the syslog server. This is adequate logging for most deployments. The second command send info-level logs with the `messageid` `FIREWALL`. The third command send info-level logs with the `messageid` `SYSTEM` and `FABRIC`.

```
# set logging-server 10.111.222.33 proto udp level info
# set logging-server 10.111.222.33 proto udp level info messageid FIREWALL
# set logging-server 10.111.222.33 proto udp level info facility syslog messageid SYSTEM,FABRIC
```

It is important that syslog is configured on all the nodes.

On NSX-T appliances, run the following command to see if syslog is configured.

```
nsx> get logging-servers
10.2.60.204 proto udp level info
vrli.corp.local proto udp level info
```

On hypervisors, firewall rule logs are kept locally (`/var/log/dfwpktlogs.log`). If you do not configure syslog, you will not see firewall rule logs on the syslog server.

## Checking Services

Services that stop running or fail to start can cause problems. It is important to make sure that all services are running normally.

To check the status of NSX Manager service:

```
nsxmgr> get services
Service name:      cm-inventory
Service state:     stopped

Service name:      http
Service state:     stopped
Session timeout:   1800
Connection timeout: 30
Redirect host:     (not configured)

Service name:      install-upgrade
Service state:     stopped
Enabled:           True

Service name:      liagent
Service state:     stopped

Service name:      manager
Service state:     stopped
Logging level:     info

Service name:      mgmt-plane-bus
Service state:     running

Service name:      node-mgmt
Service state:     running

Service name:      nsx-message-bus
Service state:     running

Service name:      nsx-upgrade-agent
Service state:     running

Service name:      ntp
Service state:     running

Service name:      search
Service state:     stopped

Service name:      snmp
Service state:     stopped

Start on boot:     False
Service name:      ssh

Service state:     running
Start on boot:     True

Service name:      syslog
Service state:     running
```

In the example above, the http service is stopped. You can start the http service with the following command:

```
nsxmgr> start service http
```

## SSH Service

If the SSH service was not enabled when deploying the appliance, you can log in to the appliance as admin and enable it with the following command:

```
start service ssh
```

You can configure SSH to start when the host starts with the following command:

```
set service ssh start-on-boot
```

To enable SSH root login, you can log in to the appliance as root, edit the file `/etc/ssh/sshd_config` and replace the line

```
PermitRootLogin prohibit-password
```

Alternatively, you can enable the SSH service and enable SSH root access by powering off the appliance and modifying its vApp properties.

with

```
PermitRootLogin yes
```

and restart the sshd server with the following command:

```
/etc/init.d/ssh restart
```

# Other Troubleshooting Scenarios

# 6

This section describes how to troubleshoot various error scenarios.

This chapter includes the following topics:

- [Failure to Add or Delete a Transport Node](#)
- [Transport Node Takes About 5 Minutes to Connect to Another Controller](#)
- [NSX Manager VM Is Degraded](#)
- [NSX Agent Times Out Communicating with NSX Manager](#)
- [Failure to Add an ESXi Host](#)
- [Incorrect NSX Controller Status](#)
- [Management IPs on KVM VMs Not Reachable with IPFIX Enabled](#)

## 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
```

## Transport Node Takes About 5 Minutes to Connect to Another Controller

When an ESXi transport node's connected controller goes down, it takes about 5 minutes for the transport node to connect to another controller.

**Problem**

An ESXi transport node is normally connected to a specific controller in a controller cluster. You can find the connected controller with the CLI command `get controllers`. If the connected controller goes down, it takes about 5 minutes for the transport node to be connected to another controller.

**Cause**

The transport node attempts to re-connect to the controller that is down for a certain amount of time before giving up and connecting to another controller. The whole process takes about 5 minutes. This is expected behavior.

## 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-2.0.0.0.6522097.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/nsx-opsagent.log` has messages such as:

```
level="ERROR" errorCode="MPA41542"] [MP_AddVnicAttachment] RPC call [0e316296-13-14] to NSX management
plane timeout
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 timeout
```

### 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 opsAgent with the following command:

```
/etc/init.d/nsx-opsagent stop
```

- b Edit the file `/etc/vmware/nsx-opsagent/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 opsAgent with the following command:

```
/etc/init.d/nsx-opsagent 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/vmkmob/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.