



VMware NSX-T 1.1 Release Notes

Updated on: 15 FEB 2017

VMware NSX-T | 02 FEB 2017 | Build 4789008

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 and System Requirements](#)
- [Known Issues](#)
- [API Reference Information](#)

What's New

NSX-T 1.1 introduces a new platform architecture that meets customer needs for flexible, scalable, and agile network and security infrastructure.

The following new features and enhancements are now available as part of the NSX-T 1.1 release.

New NSX-T Features

DHCP Server

DHCP server feature supports dynamic IP addresses as well as static IP to MAC address binding.

Metadata Proxy Server

Metadata proxy server feature allows VMs to quickly retrieve instance-specific metadata from an Openstack Nova server.

Geneve

The Generic Network Virtualization Encapsulation (Geneve) protocol is used to establish tunnels across the transport nodes to carry overlay traffic. Geneve replaces STT protocol used in the earlier release.

MAC Learning

MAC Learning capability on a logical switch provides network connectivity to deployments where multiple MAC addresses are configured behind one logical switch port, for example, in a nested

hypervisor deployment.

IPFIX

Granular IPFIX configuration support. Enable IPFIX at the granularity of logical switch or logical port.

Port Mirroring

Support for capturing traffic on a transport node for troubleshooting purposes. User can deploy a sniffer tool in VM running on a transport node, and then configure port mirroring to send virtual machine or uplink traffic to that sniffer tool for troubleshooting.

Backup and Restore

Support for automated backup. Allows you to define regular backup plan to send the backup files to a remote server.

Tech Support Bundle

Central place to collect log bundles from the NSX Manager, NSX Controllers, and transport nodes for troubleshooting purposes.

API Spec/Schema

Open API/Swagger specification support helps third parties to generate language bindings as well as other tools such as, PowerShell components and Postman collection. Users can build automation around NSX-T feature with their choice of programming languages.

NSX-T Enhancements

Grouping, Tagging and Search

Enhancements to the grouping, tagging, and search features make it easy to organize and search for objects in the NSX-T environment.

Routing

Enhancements include the support for BFD on tier-0 uplinks to detect node or path failures quickly. Route Map support for setting BGP path attributes such as weight, AS Path Prepend, and MED.

Port Connectivity

Enhancement to the port connectivity user interface that represents transport nodes connectivity to the Physical Top of Rack Switch. This connectivity information is collected through Link Level Discovery Protocol (LLDP).

Traceflow

Enhancements to the trace flow user interface where you can select VMs as source or destination from drop-down instead of logical port IDs.

Logging Improvements and Log Insight Content Pack

Consistent logging framework with error codes helps identify issues quickly in the platform. Log Insight Content pack utilizes these error codes and provides monitoring and troubleshooting dashboard for the NSX-T environment.

Compatibility and System Requirements

For compatibility and system requirement information, see the [NSX-T Installation Guide](#).

API Reference Information

The latest API reference is located in the [NSX-T Information Center](#). Use the latest API reference instead of the version available from the NSX Manager user interface.

Deprecated API Calls and Properties

The following API calls and properties are deprecated. They are marked as deprecated in the API reference. You can continue to use them at your discretion, but be aware that they will be removed from NSX-T in a future release.

Deprecated API Calls

Available API Call	Deprecated API Call
POST /api/v1/licenses	PUT /api/v1/license
GET /api/v1/licenses	GET /api/v1/license
POST /api/v1/dhcp/relay-profiles with schema DhcpRelayProfile	POST /api/v1/service-profiles
GET /api/v1/dhcp/relay-profiles with schema DhcpRelayProfileListResult	GET /api/v1/service-profiles
PUT /api/v1/dhcp/relay-profiles/ <service-profile-id> with schema DhcpRelayProfile</service-profile-id>	PUT /api/v1/service-profiles/ <service-profile-id></service-profile-id>
DELETE /api/v1/dhcp/relay-profiles/	DELETE /api/v1/service-profiles/
GET /api/v1/dhcp/relay-profiles/ <service-profile-id> with schema DhcpRelayProfile<service-profile-id>	GET /api/v1/service-profiles/
POST /api/v1/dhcp/relays with schema DhcpRelayService	POST /api/v1/services
GET /api/v1/dhcp/relays with schema DhcpRelayServiceListResult	GET /api/v1/services
DELETE /api/v1/dhcp/relays/	DELETE /api/v1/services/
PUT /api/v1/dhcp/relays/ <service-id> with schema DhcpRelayService<service-id>	PUT /api/v1/services/
GET /api/v1/dhcp/relays/ <service-id> with schema DhcpRelayService<service-id>	GET /api/v1/services/

Deprecated API Properties

--	--

Deprecated API Property	Description
Deprecated Property in BgpConfig (type)	<ul style="list-style-type: none"> • as_number (Use as_num instead.)
Deprecated Properties in BgpNeighbor (type)	<ul style="list-style-type: none"> • filter_in_ipprefixlist_id (Use address_family instead.) • filter_in_routemap_id (Use address_family instead.) • filter_out_ipprefixlist_id (Use address_family instead.) • filter_out_routemap_id (Use address_family instead.) • remote_as (Use remote_as_num instead.) • source_address (Use source_addresses instead.)
Deprecated Property in HostSwitch (type)	<ul style="list-style-type: none"> • static_ip_pool_id (Use ip_assignment_spec instead.)
Deprecated Property in TransportNode (type)	<ul style="list-style-type: none"> • host_switches (Use host_switch_spec instead.)
Deprecated Property in PortConnectionHypervisor (type)	<ul style="list-style-type: none"> • pnics
Deprecated Property in AddControllerNodeSpec (type)	<ul style="list-style-type: none"> • control_plane_server_certificate

Known Issues

The known issues are grouped as follows.

- [General Issues](#)
- [Installation Issues](#)
- [NSX Manager Known Issues](#)
- [NSX Edge Known Issues](#)
- [Logical Networking Known Issues](#)
- [Security Services Known Issues](#)
- [Operations and Monitoring Services Known Issues](#)
- [KVM Networking Known Issues](#)
- [Solution Interoperability Known Issues](#)
- [API Known Issues](#)
- [Documentation Errata and Additions](#)

General Issues

- **Management cluster supports only a single node.**
- **Issue 1769925: Occasionally, support bundle collection from the user interface might fail on NSX Edge nodes**
Collecting support bundles from the NSX Manager user interface might fail with an error message that the NSX Edge nodes are unreachable.

Workaround: Log in to NSX Manager and restart the support bundle application.

- **Issue 1747453: Status of a node in the NSX Controller might become inactive because of a known Zookeeper problem**
After powering on and powering off a node in the NSX Controller several times, this node might become inactive. See <https://issues.apache.org/jira/browse/ZOOKEEPER-1549>

Workaround: Complete the following steps.

1. Remove the node with this problem out of the cluster.
`detach control-cluster <controller_IP>`
2. Deactivate the clustering configuration on the node.
`deactivate control-cluster`
3. Join the node to the control cluster.
`join control-cluster thumbprint <thumbprint>`
4. Activate the control cluster on the node.
`activate control-cluster`

- **Issue 1764415: vMotion of a VM connected to a logical switch fails because of a network not accessible error**

vMotion of a VM can fail in the following scenarios:

1. NSX-T vSwitch status is down.
2. NSX-T vSwitch status changed from down to up but the change was not reflected.

Workaround: In the first case, failure is expected.

In the second case, invoke the Refresh NetworkSystem API via the Virtual Center NGC user interface on the ESXi host where the VM is deployed.

- **Issue 1774379: Issues in the open-source swagger-codegen project affecting NSX-T OpenAPI**

Due to open bugs in the open-source swagger-codegen project

<https://github.com/swagger-api/swagger-codegen>, the NSX-T OpenAPI spec inlines all type definitions except for leaf-node type definitions in the inheritance hierarchy. The NSX-T OpenAPI does not include the supertype via an allOf JSON Schema directive, but instead copies all of the superclass properties into the subclass for any classes that are supertypes of another class. When this issue is fixed in swagger-codegen, an updated NSX-T OpenAPI specification is posted to the VMware downloads Web site.

- **Issue 1691716: Disabling BFD configuration from the external router black holes the traffic**

When BFD configuration is deleted or disabled from the external router, the external physical router sends an ADMIN_DOWN message. NSX Edge does not remove the static route whereas the external router removes the static route, which black holes the traffic at the NSX Edge.

This problem only occurs in case of the ADMIN_DOWN message.

Workaround: Manually delete the static route from the NSX Edge before removing the BFD configuration from the external router.

Installation Issues

- **Issue 1747450: After successfully installing a fabric node from the NSX Manager user interface, the status might show "Installation Failed" for a noticeable period of time before it changes to "Installation Successful"**

Workaround: None required. Wait until the status changes to **Installation Successful**.

- **Issue 1630494: When htmlClient rule is enabled on ESXi host, host registration/deregistration to NSX Manager fails**

The register-node and deregister-node commands on ESXi use port TCP 443 in the NSX Manager cluster. Hence, a rule is added to ESXi firewall with allowedip=all to enable outbound traffic to TCP port 443. However, there is another rule in ESXi firewall named htmlClient - this by default has allowedip=all and is not enabled. However, if the user is to enable this rule and specify a list of allowed IP addresses, the ESXi firewall matches all traffic destined to TCP port 443 to this rule. It can thus drop traffic destined to the NSX Manager cluster, causing NSX Manager node registration/deregistration to fail.

Workaround: Disable htmlClient firewall rule and any rules that applies filters on TCP port 443.

- **Issue 1564210: NSX Manager and NSX Controller appliances must use static IP addresses**

NSX Manager and NSX Controller appliances must use static IP addresses. Changing the IP address of an NSX Manager or NSX Controller is not supported.

Workaround: If you install an NSX Manager or NSX Controller appliance without a static IP configuration and it gets an IP address from DHCP, you must install a new appliance. Do not change the IP address of the appliance.

- **Issue 1576541: The detach management-plane command error message states Fetch the latest copy of the object**

The detach management-plane command might fail and return the following error message:

```
% Node deregistration failed: 'The object unknown was modified by somebody else. Fetch the latest copy of the object and retry operation'.
```

Workaround: You do not need to fetch the latest copy of the object. Run the detach management-plane command again to detach the node from the management plane.

- **Issue 1618327: NSX Edge nodes are missing deployment_type**

NSX Edge nodes are missing `deployment_type` in the response of this API request: `https://<NSX_Manager>/api/v1/fabric/nodes/<node-id>`. Configuration of the NSX Edge fails because the system does not meet the [system requirements](#).

Workaround: Install the NSX Edge on hardware with Westmere-based or Sandy Bridge-based CPU. See the system requirements for details.

- **Issue 1538016: vSphere Client does not support configuring OVF extra configuration properties, such as IP address, when directly connected to ESXi**

If you are using ESXi without vCenter, and deploy an NSX-T appliance with the vSphere Client, you cannot edit the OVF extra configuration properties.

Workaround: If you need to install NSX Edge, NSX Manager, and NSX Controller appliances on ESXi without vCenter, use the `ovftool` command. For NSX Edge only, you can install with the vSphere Client, and then log in to the CLI and configure the network interface manually.

- **Issue 1576112: KVM hypervisors require manual configuration of gateway if they reside in different Layer 2 segments**

If you configure an IP pool on NSX Manager and use that IP pool for creating transport nodes, Ubuntu KVM boxes do not show a route for the gateway that was configured in the IP Pool configuration. As a result, the overlay traffic between VMs that reside on hypervisors that are in different L2 segment fail because the underlying fabric host does not know how to reach the fabric nodes in remote segments.

Workaround: Add a route for the gateway so that it can route traffic to other hypervisors that reside in different segments. If this configuration is not done manually, then the overlay traffic would fail since the fabric node does not know how to reach the remote fabric nodes.

- **Issue 1617459: Host configuration for Ubuntu does not support sourcing of interface configuration files**

If the `pnic` interface is not in the `/etc/network/interfaces` file, then MTU is not configured correctly in network configuration file. Because of this, MTU configuration on transport bridge is lost after every reboot.

Workaround: Move PNIC interface configuration to `/etc/network/interfaces`

NSX Manager Known Issues

- **Issue 1762527: NSX Manager might trigger a cache-poisoning vulnerability warning during a security scan**

NSX Manager might trigger a cache-poisoning vulnerability warning during a security scan because the NSX Manager uses the host HTTP request header to construct the redirect response.

For example:

```
GET /nsxapi/ping.json?_dc=1474040351698 HTTP/1.1
```

```
Host: 10.32.41.238
```

```
User-Agent: Mozilla/5.0 (X11; Ubuntu; Linux x86_64; rv:32.0) Gecko/20100101 Firefox/32.0
```

Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8

Accept-Language: en-US,en;q=0.5

Accept-Encoding: gzip, deflate

Cookie: JSESSIONID=DE5258CE9FAD8C160B2BC94E2A63EC0C

Connection: keep-alive

The risk of cache-poisoning is mitigated by disallowing HTTP responses from being persisted at intermediate caches.

- **Issue 1742510: NSX Manager database can be corrupted by a power-off event**

Occasionally, NSX Manager cannot start up after a power-off event. The NSX Manager status does not report **STABLE** even after 5 minutes.

The following log message appears in the nsxapi.log file: [nsx comp="nsx-manager" errorCode="MP4113" subcomp="manager"] GemFire is in illegal state, restating Proton process com.vmware.nsx.management.container.dao.gemfire.GemFireInitializationException: java.lang.NullPointerException

Workaround: Restore the latest NSX Manager backup.

- **Issue 1710152: NSX Manager GUI does not work on Internet Explorer 11 in compatibility mode**

Workaround: Go to **Tools > Compatibility View Settings** and verify that Internet Explorer does not display the NSX Manager GUI in compatibility mode.

NSX Edge Known Issues

- **Issue 1762064: Configuring the NSX Edge VTEP IP-pool and uplink profile immediately after rebooting the NSX Edge causes the VTEP BFD session to become unreachable**

After rebooting the NSX Edge, the broker requires some time to reset the NSX Edge connections.

Workaround: Wait about five minutes after rebooting the NSX Edge to allow the broker to reset the NSX Edge connections.

- **Issue 1747919: Static route created with Nexthop as VIP IP not pushed to the NSX Edge node**

Static route created with Nexthop as VIP IP not pushed to the NSX Edge node when the VIP subnet is different than the uplink interface subnet.

Workaround: For static route with VIP, always use the VIP IP address from one of the existing uplink subnet IP.

- **Issue 1765087: Kernel interfaces that NSX Edge creates to transfer packets from the datapath to Linux kernel only supports MTU up to 1600.**

Kernel interfaces between datapath and kernel does not support the jumbo frame. BGP packets size that exceed 1600 are truncated and dropped by the BGP daemon. SPAN packets size that exceed 1600 are truncated and the packet capture utility displays a warning. The payload is not truncated and remains valid.

Workaround: None.

- **Issue 1738960: If a DHCP server profile NSX Edge node is replaced with an NSX Edge node from another cluster, then IP addresses given to VMs by the DHCP server change**

This issue is caused by a lack of coordination between the node that is replaced and the new node.

Workaround: None.

- **Issue 1629542: Setting a forwarding delay on single NSX Edge node causes an incorrect routing status to be displayed**

When running an NSX Edge as a single NSX Edge node (not in an HA pair), configuring a forwarding delay might result in an incorrect reporting of the routing status. After the forwarding delay is configured, the routing status incorrectly appears as **DOWN** until the forwarding timer expires. If router convergence is complete but the forwarding delay timer has not yet expired, the datapath from south to north continues to flow as expected, even if the routing status is reported as **DOWN**. You can safely ignore this warning.

- **Issue 1601425: Cannot clone NSX Edge VM that is already registered with the NSX Manager cluster**

Cloning of an NSX Edge VM once it is registered with the NSX Manager cluster is not supported. Instead, a fresh image should be deployed.

Workaround: None.

- **Issue 1580586: Redistribution rules do not support LE or GE configurations in the PrefixList**

Redistribution rules do not support LE or GE configurations in the prefix-list, the NSX Manager is not validating this configuration, and the NSX Edge is not supporting this. As a result, the user might see that the configuration is not taking effect.

Workaround: Do not use le or ge configurations in the IP-prefix-list.

- **Issue 1604923: Removing or changing NSX Edge cluster member indexes referred by the Tier-1 logical router link port disrupts the north-south traffic**
- **Issue 1585575: Cannot edit NSX Edge cluster details on Tier-1 router attached to a Tier-0 router**

If you have enabled NAT on a Tier-1 logical router, you must specify an NSX Edge node or NSX Edge cluster before connecting the Tier-1 router to a Tier-0 router. NSX does not support editing the NSX Edge cluster details on a Tier-1 router that is already attached to a Tier-0 router.

Workaround: To edit the NSX Edge cluster details on a Tier-1 router that is already attached to a Tier-0 router, disconnect the Tier-1 router from the Tier-0 router, make the changes, and reconnect again.

Logical Networking Known Issues

- **Issue 1769922: NSX Controller cluster plane might show internal IP address 172.17.0.1 on vSphere Client rather than actual IP address**

On vSphere Client, the IP address for NSX Controllers is incorrectly shown as 172.17.0.1 rather than the actual IP address. For NSX Manager, the IP address is shown correctly.

Workaround: None needed. This cosmetic issue does not affect any functionality.

- **Issue 1771626: Changing the IP address of the NSX Controller node is not supported**

Workaround: Redeploy the NSX Controller cluster.

- **Issue 1753468: Enabling Spanning Tree Protocol (STP) on bridged VLAN causes the bridge cluster status to display as down**

When STP is enabled on VLANs that are used for bridging with LACP teaming, the physical switch port-channel is blocked resulting in the bridge cluster on the ESX host to display as down.

Workaround: Disable STP or enable the BPDU filter and BPDU guard.

- **Issue 1753468: Tier-0 logical router does not aggregate the routes, instead the logical router redistributes them individually**

Tier-0 logical router does not perform route aggregation for a prefix which does not cover all the sub-prefixes connected to it and instead the logical router distributes the routes separately

Workaround: None.

- **Issue 1763570: Deleting old IP Pool after the transport node is updated with new IP Pool fails**

Workaround: Delete all the logical router's configuration and apply the new IP Pool on all transport nodes.

- **Issue 1536251: Copying VMs from an ESX host to another ESX host which is attached to same logical switch is not supported**

Layer 2 network fails when a VM is copied from one ESX host and the same VM is registered on another ESX host

Workaround: Use VM Cloning if the ESX host is part of Virtual Center.

If you do copy a VM between ESX hosts, the external ID must be unique in the VM .vmx file for the layer 2 network to work.

- **Issue 1747485: Removing any uplink from the LAG interface brings all of the BFD protocol down and flaps BGP routes**

When any interface is deleted from the configured LAG interface, it brings all of the BFD protocol down and flaps BGP routes, which impacts traffic.

Workaround: None.

- **Issue 1773703: BGP stops advertising all PERMIT ip-prefixes, when a route map is**

added in the OUT direction with an ip-prefix-list in a single sequence without the GE and LE options provided

BGP stops advertising all PERMIT ip-prefixes, when an ip-prefix is set as PERMIT and another ip-prefix is set as DENY, then the ip-prefix-list single sequence action in the route map is set as PERMIT, and that route map in the BGP neighbor filter is in the OUT direction.

Workaround: You can perform one of the following tasks:

- Instead of a route map you can use the ip-prefix-list in the BGP neighbor filter.
- Use the GE and LE options while adding one of the ip-prefixes in the ip-prefix-list in the route map.
- Create a separate ip-prefix-list for each ip-prefix and add them in the route map in separate sequences.

- **Issue 1769491: Deleting route map added in the BGP neighbor filter OUT direction causes an Assertion error and flaps BGP routes**

Workaround: No workaround needed because BGP connection gets reestablished in a few seconds.

- **Issue 1736536: Maximum number of supported logical switches with MDProxy service is 1024**

Configuring more than 1024 logical switches with MDProxy service does not allow the back end MDProxy service to start the nginx Web server.

Workaround: You can add more than 1024 logical switches on all the NSX Edges that support the MDProxy service.

1. Navigate to /etc/init.d/nsx-edge-mdproxy with root privilege.
2. Delete the line ulimit -n 1024.
3. Navigate to etc/init/nsx-edge-nxa.conf with root privilege.
4. Add the line limit nofile 20000 20000.
5. Restart the MDProxy service in the admin console.
restart service local-controller

- **Issue 1721716: Logical router port can be deleted even if there are existing static routes configured on that port**

When you delete a logical router port with static routes configured the static routes remain in the system.

Workaround: You can manually delete the static routes or leave these routes in the system which do not cause any harm.

- **Issue 1763576: Hypervisors are allowed to be removed as transport nodes even when they have VMs on the NSX-T network**

NSX-T does not prevent you from deleting a transport node even when there are VMs on the node that are part of the NSX-T network. The VMs lose connectivity after the transport node is deleted.

Workaround: For both ESXi and KVM hosts, recreate the transport node again with the same host switch name.

- **Issue 1780798: In a large-scale environment, some hosts might get into a failed state**
In a large-scale environment with 200 or more host nodes after running for some time, some hosts might lose connectivity with NSX Manager and the log contains error messages such as:

```
2016-12-09T00:57:58Z mpa: [nsx@6876 comp="nsx-esx" subcomp="mpa" level="WARN"] Unknown routing key: com.vmware.nsx.tz.*
```

Workaround: Restart the MPA process on the failed hosts.

- **Issue 1741929: In a KVM environment, when port mirroring is configured and truncation is enabled, jumbo packets from the source are sent in fragments but are re-assembled at the mirror destination**

Workaround: No workaround needed because the re-assembly is performed by the destination VM vNIC driver.

- **Issue 1754187: In a multi-transport zone environment, after removing one of the transport zones from a transport node, VMs in the removed transport zones can still participate in the NSX-T network**

Workaround: Before removing a transport zone from a transport node, disconnect VMs from logical switches that are in the transport zone.

- **Issue 1770041: Between BGP peers when a route-map is configured to prepend ASN, the standby fails to prepend it immediately**

Between 2-byte BGP peers, when a route-map is configured to prepend 4-byte ASN, it takes the standby 15 seconds to prepend the 4 byte ASN properly.

Between 4-byte BGP peers, when a route-map is configured to prepend 2-byte ASN, it takes the standby 15 seconds to prepend the 2 byte ASN properly.

Workaround: None.

- **Issue 1585874: IP address binding needs to be configured with port SpoofGuard on a logical switch profile**

If port SpoofGuard is enabled on a logical switch profile, IP address binding also needs to be configured on the VM ports belonging to the logical switch. This is especially important for ports connecting to vCenter VMs, because without the binding configured, traffic on VM ports could be black holed due to an empty whitelist configuration with SpoofGuard.

Workaround: None.

- **Issue 1619838: Changing a transport zone connection of a logical router to a different set of logical switches fails with a mismatch error**

Logical router only supports a single overlay transport zone for downlink ports. Therefore, without deleting the existing downlink or routerlink ports you cannot change a transport zone connection to a different set of logical switches.

Workaround: Complete the following steps.

1. Delete all of the existing downlink or routerlink ports.
2. Wait for some time for the system to update.
3. Retry changing the transport zone connection to a different set of logical switches.

- **Issue 1620144: NSX-T CLI, get logical-switches lists logical switches with status UP, even after the transport node is deleted**

The NSX-T CLI might mislead the user that there is a functional logical switch. Even when logical switches are seen, they are not functional. The opaque switch is disabled when the transport node is deleted, thus no traffic gets through.

Workaround: None.

- **Issue 1625360: After creating a logical switch, the NSX Controller might not show the newly created logical switch information**

Workaround: Wait 60 seconds after creating logical switch to check the logical switch information on the NSX Controller.

- **Issue 1581649: After logical switch creation and deletion, VNI pool range cannot be shrunk**

Range shrink fails because VNIs are not released immediately after a logical switch is deleted. VNIs are released after 6 hours. This is to prevent reuse of VNIs when another logical switch is created. Due to this you cannot shrink or modify ranges until 6 hours after the logical switch deletion.

Workaround: To modify the range from which VNIs had been allocated for logical switches, wait for 6 hours after the deletion of logical switches. Alternatively, use other ranges from the VNI Pool, or reuse the same range without shrinking or deleting the range.

- **Issue 1516253: Intel 82599 NICs have a hardware limitation on the Queue Bytes Received Counter (QBRC) causing an overflow after total received bytes exceeds 0xFFFFFFFF**

Because of the hardware limitation, the CLI output of `get dataplane physical-port stats` does not match the actual number if overflow occurs.

Workaround: Run the CLI once such that the counters is reset and run again in shorter durations.

Security Services Known Issues

- **Issue 1759369: Weak TLS configuration on the NSX Controller allows attackers to read or modify the TLS traffic between the NSX Controller and the hypervisor and vice-versa**

TLS listener on port 1234 is configured to support TLS 1.0, which is vulnerable to external attacks.

Workaround: Physically secure the network between NSX Controller and hypervisors.

- **Issue 1643872: After you apply a filter to a firewall rule, you cannot disable the rule**
In NSX Manager user interface, if you apply a filter to a rule, the user interface does not allow you to disable it.

Workaround: Remove the filter and then disable the rule.

- **Issue 1765476: NS Groups synchronization with the NSX Controller might take some time after reboot**

After rebooting a node in the management plane, synchronizing the NS Groups with the NSX Controller might take about 30 minutes or more depending on the number of NS Groups.

Workaround: None.

- **Issue 1680128: DHCP communication between client and server is not encrypted**

Workaround: Use IPSEC to make the communication more secure.

- **Issue 1711221: IPFIX data is sent over the network in plaintext**

By default, the option to collect IPFIX flows is turned off.

Workaround: None.

- **Issue 1721519: Connection between the metadata proxy server and the Nova server is not encrypted or authorized**

Workaround: None.

- **Issue 1726081: Geneve tunnel traffic (UDP) is rejected in KVM**

Workaround: Perform the following steps to allow Geneve tunnel traffic

If KVM is using firewalld, create a hole in the firewall with the following command:

```
# firewall-cmd --zone=public --permanent --add-port=6081/udp
```

If KVM is using IPtables directly, create a hole with the following command:

```
# iptables -A INPUT -p udp --dport 6081 -j ACCEPT
```

If KVM is using UFW, create a hole with the following command:

```
# ufw allow 6081/udp
```

- **Issue 1520694: In RHEL 7.1 kernel 3.10.0-229 and earlier, FTP ALG fails to open negotiated port on data channel**

For an FTP session, where both client and server reside in VMs on the same hypervisor, the FTP application level gateway (ALG) does not open up the negotiated port for the data channel. This issue is specific to Red Hat and is present in RHEL 7.1 kernel 3.10.0-229. Later RHEL kernels are not affected.

Workaround: None.

- **Issue 1590888: Warning needed that logical ports selected in Ethernet section apply only within same L2 network**

For the NSX-T distributed firewall, in the Ethernet section, when any logical port or MAC

address is entered in the source/destination section, a warning should be displayed that MAC addresses or logical ports should belong to VM ports in same L2 network (attached to same Logical switch). Currently, there is no warning message.

Workaround: None.

Operations and Monitoring Services Known Issues

- **Issue 1764175: The Port Connection and Traceflow tools take a long time to discover VMs**

In a scenario where an NSX-T deployment starts with a small number of hosts (10 or fewer) and grows to a large number of hosts, and NSX Manager or the proton service has not been restarted since the initial small number of hosts, if NSX Manager loses connectivity with a large number of hosts for any reason without the proton service being restarted, it takes a long time for NSX Manager to synchronize with those hosts and discover the VMs that run on them. The proton log file has messages such as the following: Processing sync_init requests, batch size <calculated batch size>

Workaround: Restart the proton service. You do not need to reboot NSX Manager.

- **Issue 1743476: Moving a VM from a non-NSX-T managed switch to a NSX-T managed switch might disable the firewall on the VM.**

Moving a VM deployed through the NSX-T workflows from a non-NSX-T managed switch to a NSX-T managed switch might disable the firewall on the VM.

Workaround: Power off the VM and power the VM on.

- **Issue 1749078: After deleting a tenant VM on an ESXi host and the corresponding host transport node, deleting the ESXi host fails**

Deleting a host node involves reconfiguring various objects and can take several minutes or more.

Workaround: Wait several minutes and retry the delete operation. Repeat if necessary.

- **Issue 1761955: Unable to connect a VM's vNIC to an NSX-T logical switch after registering the VM**

If an existing vmx file is used to register a VM on an ESXi host, the register operation ignores following vNIC-specific errors:

- vNICs that are configured with invalid network backing.
- VIF attachment failures for vNICs that are connected to an NSX-T logical-switch.

Workaround: Complete the following steps.

1. Create a temporary port group on a standard vSwitch.
2. Attach the vNICs that are in the disconnected state to the new port group and mark them as connected.
3. Attach the vNICs to a valid NSX-T logical switch.

- **Issue 1774858: On rare occasions, the NSX Controller cluster becomes inactive after running for multiple days**

When the NSX Controller cluster becomes inactive, all transport and NSX Edge nodes lose connectivity to the NSX Controllers and changes to the configuration cannot be made. However, data traffic is unaffected.

Workaround: Complete the following steps.

- Fix disk latency issues if they exist.
- Restart the cluster-mgmt service on all NSX Controllers.

- **Issue 1576304: Dropped-byte count is not included as part of the Port Status and Statistics report**

When using `/api/v1/logical-ports/<lport-id>/statistics` or NSX Manager to view logical port status and statistics, there is dropped-packet count with a value of 0. This value is not accurate. Regardless of the number of dropped packets, the number displayed here is always blank.

Workaround: None.

KVM Networking Known Issues

- **Issue 1717061: In a KVM environment, network performance is poor for tunneled traffic with older Linux kernels**

NSX-T 1.1 uses Geneve instead of STT that was used in previous releases for encapsulation. Older Linux kernels were not optimized for Geneve.

Workaround: Run the supported Ubuntu or RHEL release that has the most recent version of the Linux kernel.

- **Issue 1775916: The resolver API POST `/api/v1/error-resolver?action=resolve_error` does not resolve errors after a RHEL KVM host fails to be added to the fabric**

After a RHEL KVM host fails to be added to the fabric and the NSX Manager user interface shows the installation status as failed, the resolver API POST `/api/v1/error-resolver?action=resolve_error` is run to resolve errors. However, adding the host to the fabric again results in the following error messages:

```
Failed to install software on host. Un-handled deployment plug-in perform-action.  
Install command failed.
```

Workaround: Complete the following steps.

1. Manually remove the following packages.

```
rpm -e glog-0.3.1-1nn5.x86_64  
rpm -e json_spirit-v4.06-1.el6.x86_64  
rpm -e kmod-openvswitch-2.6.0.4557686-1.el7.x86_64  
rpm -e nicira-ovs-hypervisor-node-2.6.0.4557686-1.x86_64  
rpm -e nsx-agent-1.1.0.0.4690847-1.el7.x86_64  
rpm -e nsx-aggservice-1.1.0.0.4690845-1.el7.x86_64  
rpm -e nsx-cli-1.1.0.0.4690892-1.el6.x86_64  
rpm -e nsx-da-1.1.0.0.4690845-1.el7.x86_64  
rpm -e nsx-host-1.1.0.0.4690932-1.x86_64 rpm -e nsx-host_node_status_reporter-1.1.0.0.4690845-1.el7.x86_64  
rpm -e nsx-lldp-1.1.0.0.4690845-1.el7.x86_64
```

```
rpm -e nsx-logical_exporter-1.1.0.0.4690845-1.el7.x86_64
rpm -e nsx-mpa-1.1.0.0.4690845-1.el7.x86_64
rpm -e nsx-netcpa-1.1.0.0.4690924-1.el7.x86_64 rpm -e nsx-sfhc-1.1.0.0.4690845-1.el7.x86_64
rpm -e nsx-support-bundle-client-1.1.0.0.4690845-1.el7.x86_64
rpm -e nsx-transport_node_status-1.1.0.0.4690845-1.el7.x86_64
rpm -e nsxa-1.1.0.0.4690845-1.el7.x86_64
rpm -e openvswitch-2.6.0.4557686-1.x86_64
rpm -e openvswitch-selinux-policy-2.6.0.4557686-1.noarch
rpm -e python-simplejson-3.3.3-1.el7.x86_64
```

If there are any error while running the rpm -e command, include the --noscripts flag to the command.

2. Run the resolver API POST `/api/v1/error-resolver?action=resolve_error`.
3. Add the KVM host to the fabric again.

- **Issue 1602470: Load balance teaming is not supported on KVM**

- **Issue 1611154: VMs in one KVM transport node cannot reach VMs located in another transport node**

When multiple IP pools are used for VTEPs that belong to different networks, the VM on the KVM host might not reach the VM deployed on other hosts that have VTEP IP addresses from a different IP pool.

Workaround: Add routes so that the KVM transport node can reach all of the networks used for VTEP on other transport nodes.

For example, if you have two networks 25.10.10.0/24 and 35.10.10.0/24 and the local VTEP has the IP address 25.10.10.20 with gateway 25.10.10.1, you can use the following command to add the route for another network:

```
ip route add dev nsx-vtep0.0 35.10.10.0/24 via 25.10.10.1
```

- **Issue 1587627: Ops/BFD APIs might report unusual information**

Ops information is collected from all platforms BFD/tunnel implementation on the ESXi, NSX Edge, or KVM platforms is not consistent.

- NSX Edge and KVM keep the BFD/tunnel session alive forever.
- ESXi destroys these sessions when not needed and creates them on demand.

Because of this there might be unusual BFD state where one end shows tunnel is down and another end does not have that tunnel. Additionally, BFD does not aid in debugging because when a BFD session is deleted, the last cause of failure is lost.

Workaround: None

- **Issue 1654999: Connection tracking of underlay traffic reduces available memory**

When establishing a large number of connections between virtual machines, you might experience the following symptoms.

In the `/var/log/syslog` or `/var/log/messages` file, you see entries similar to:

```
Apr 26 11:45:44 prmh-nsx-perf-server149 kernel: [1625289.950872] net_ratelimit: 239 callbacks suppressed
Apr 26 11:45:44 prmh-nsx-perf-server149 kernel: [1625289.950875] nf_contrack: table full, dropping packet
Apr 26 11:45:44 prmh-nsx-perf-server149 kernel: [1625289.958436] nf_contrack: table full, dropping packet
```

The issue seems to manifest itself when default firewall rules have been configured. The

issue does not manifest itself if firewall rules are not configured (For example: Logical switches are put in the firewall exclusion list).

Note: The preceding log excerpts are only examples. Date, time, and environmental variables might vary depending on your environment.

Workaround: Add a firewall rule to disable connection tracking for UDP on port 6081 on underlay devices.

Here is an example command:

```
# iptables -A PREROUTING -t raw -p udp --dport 6081 -j CT --notrack
```

This should be configured to run during boot. If the platform also has a firewall manager enabled (Ubuntu: UFW; RHEL: firewalld), the equivalent rule should be configured through the firewall manager. See related [KB 2145463](#).

Solution Interoperability Known Issues

- **Issue 1588682: Putting ESXi hosts in lockdown mode disables the user nsx-user**
When an ESXi host is put into lockdown mode, the user vpxuser is the only user who can authenticate with the host or run any commands. NSX-T relies on another user, nsx-user, to perform all NSX-T related tasks on the host.

Workaround: Do not use Lockdown mode. See [Lockdown Mode](#) in the vSphere documentation.

API Known Issues

- **Issue 1781225: The API GET `https://<NSX-Manager>/api/v1/fabric/nodes/<node-id>/modules` does not work for Ubuntu**
The API GET `https://<NSX-Manager>/api/v1/fabric/nodes/<node-id>/modules` works for ESXi and RHEL but not for Ubuntu.

Workaround: None

- **Issue 1781233: The API GET `https://<NSX-Manager>/api/v1/fabric/nodes/status` returns an error**
The API GET `https://<NSX-Manager>/api/v1/fabric/nodes/status` , which gets the status of multiple nodes, might return an error.

Workaround: Use GET `https://<NSX-Manager>/api/v1/fabric/nodes/<node-id>/status` to get the status of individual nodes.

- **Issue 1770207: When making an API call to change authentication policies, the changes do not take effect**
When you make the API call PUT `https://<NSX-Manager>/api/v1/node/aaa/auth-policy` to change any of the following policies, the changes do not take effect.

- `api_failed_auth_lockout_period`
- `api_failed_auth_reset_period`
- `api_max_auth_failures`
- `minimum_password_length`

Workaround: Restart the proxy service.

- **Issue 1605461: NSX-T API logs in syslog show system-internal API calls. NSX-T logs**

both user-invoked API calls as well as system-invoked API calls to syslog

The logging of an API call event in syslog is not evidence of a user directly calling the NSX-T API. You see NSX Controllers and NSX Edge API calls in the logs, even though these NSX-T appliances do not have a publicly exposed API service. These private API services are used by other NSX-T services such as, the NSX-T CLI.

Workaround: None.

- **Issue 1619450: Test vertical is returned by polling frequency configuration APIGET /api/v1/hpm/features**

GET /api/v1/hpm/features returns the list of all features for which polling frequency can be configured. This API returns some internal, test-only features. There is no functional impact on the user other than extra noise.

Workaround: Ignore the extraneous API response.

- **Issue 1641035: Rest call to POST/hpm/features/<feature-stack-name? action=reset_collection_frequency> does not restore the collection_frequency for overwrite statistics**

If you attempt to reset the collection frequency to its default by using this REST call, it won't be reset.

Workaround: Use PUT /hpm/features/<feature-stack-name> and set collection_frequency to the new value.

- **Issue 1648571: On-demand status and statistics requests can intermittently fail. HTTP failure code is inconsistent**

In certain situations, on-demand requests fail. Sometimes these requests fail with an HTTP 500 error instead of an HTTP 503 error, even though the API call succeeds on retry.

For statistics APIs, the timeout condition might result in spurious message-routing error logs. These occur because the response returns after the timeout period has expired.

For example, errors such as the following might occur: java.lang.IllegalArgumentException:

Unknown message handler for type

com.vmware.nsx.management.agg.messaging.AggService\$OnDemandStatsResponseMsg.

For status APIs, the timeout condition, a response returns after timeout, could cause the cache to be updated prematurely.

Workaround: Retry API request.

Documentation Errata and Additions

- **Issue 1372211: Two interfaces on same subnet**

Tunnel traffic can leak out to the management interface if the tunnel endpoint is on the same subnet as the management interface. This happens because tunneled packets can go through the management interface. Make sure that the management interfaces are on a separate subnet from the tunnel endpoint interfaces.

- **Issue: API request results in 403 Forbidden, Bad XSRF token error**

While you are logged in to the NSX Manager Web user interface, the NSX Manager cookie is only usable within the Web user interface. If you send API requests via an application that uses the same cookie source (for example, a browser extension), you get a 403

Forbidden / Bad XSRF token? response.

```
{  
"module_name": "common-service",  
"error_message": "Bad XSRF token",  
"error_code": "98"  
}
```

Workaround: You must either log out of the NSX Manager Web user interface, or use a REST client that uses a different source of cookies.

For example, use one browser for web user interface access, and an extension on a different browser to access the API. You can also get session-cookies that don't require the XSRF header. See the Session-based Authentication section of the NSX-T API Guide for more information.

- **Issue 1622719: Disconnect associated VMs from logical switches before making certain transport node and transport zone changes**

Before performing the following procedures, disconnect the associated VMs from the logical switch.

- Before disconnecting a logical switch from a transport node.
- Before moving a transport node from one transport zone to another transport zone.
- Before deleting a transport zone.

- **Issue 1590888: Logical ports selected in Ethernet section apply only within same Layer 2 network**

If you create a firewall rule with a logical port or a MAC address for the source or destination and you apply the distributed firewall rule in the Ethernet section, the MAC addresses or logical ports must belong to VM ports that are in the same Layer 2 network. In other words, they must be attached to same Logical switch.

- **Issue 1537399: IPFIX on ESXi and KVM sample tunnel packets in different ways**

On ESXi the tunnel packet is sampled as two records:

Outer packet record with some inner packet information.

- SrcAddr, DstAddr, SrcPort, DstPort, Protocol refers to the outer packet.
- Contains some enterprise entries to describe the inner packet.

Inner packet record.

- SrcAddr, DstAddr, SrcPort, DstPort, Protocol refers to the inner packet.

On KVM the tunnel packet is sample as one record:

The inner packet record with some outer tunnel information.

- SrcAddr, DstAddr, SrcPort, DstPort, Protocol refers to the inner packet
- Contains some enterprise entries to describe the outer packet.

- **Issue 1520687: Possible performance degradation when using IPFIX**

Certain IPFIX configurations can cause performance degradation on ESXi and KVM hypervisors. For example, setting low values for idle timeout, flow timeout, or max flows, combined with high values for sampling probability can cause performance degradation. When setting the IPFIX configuration, monitor the impact on hypervisor performance.

- **Issue 1622362: IP address range 100.64.0.0/10 used for external transit network**

Do not configure the IP address range 100.64.0.0/10 on the uplink interface of a tier-0 logical router. This IP address range is used for external transit network on the NSX Edge. See RFC 6598 for more information.

- **Issue 1444337: Support bundle contains auditing information including usernames**
You can generate a support bundle which contains information from NSX-T appliances that can be used by VMware Support to diagnose customer reported issues. This support bundle contains auditing information, including usernames of valid users on the appliance.

- **Issue 1608972: Changing NSX Manager HTTP connection timeout or session timeout requires a restart**

Changing NSX Manager's HTTP connection timeout or session timeout requires a restart. This is not covered in the NSX documentation. You can restart the HTTP service through the API or the CLI.

API- `POST /api/v1/node/services/http?action=restart`

CLI- `restart service http`

Copyright © 2021 VMware, Inc. All rights reserved.