



VMware NSX-T 2.1.0.1 and 2.1 Release Notes

VMware NSX-T 2.1.0.1 | 08 FEB 2018 | Build 7725122

VMware NSX-T 2.1 | 21 DEC 2017 | Build 7395507

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](#)
- [API Reference Information](#)
- [Resolved Issues](#)
- [Known Issues](#)

What's New

What's New

NSX-T 2.1 is the incremental upgrade release that enhances the new multi-hypervisor platform delivered for cloud and containers.

NSX-T 2.1.0.1 is a maintenance release specifically for the NSX Container Plug-in (NCP) feature. It has the following improvements:

- Support for wildcard character when configuring Ingress URI paths for load balancing.
- Support for long labels in a Kubernetes environment. NCP normalizes long labels to fit them in NSX tags.
- Support for long names in a Kubernetes environment. NCP can handle resource names that are longer than the maximum NSX tag name (40 characters).
- In a Pivotal Cloud Foundry environment, NCP can handle an app instance that is created in multiple diego cells.

The following new features and enhancements are available in the NSX-T 2.1 release.

New NSX-T Features

Load Balancer

Inline and one-arm load balancer topology support for workloads. NSX-T load balancer support in OpenStack deployments and Ingress in Kubernetes container deployments.

Pivotal Application Service Integration (PAS/PCF)

NSX-T 2.1 integration with Pivotal Application Service 2.0 (CNI integration).

Pivotal Container Service

Networking and security feature support for Pivotal Container Service.

Policy Manager

Policy manager support to allow intent-based firewall policy.

NSX-

T Enhancements

Distributed Firewall

Distributed Firewall now supports total rule hit counts.

Usability Workflows

Enhanced NSX-T dashboard, getting started workflow, and robust search and filter capabilities.

NSX Edge Node Deployment

Updated NSX Edge node deployment.

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 Product Information](#). 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

- `GetLogicalRouterRouteTableInCsvFormat (/logical-routers/<logical-router-id>/routing/route-table?format=csv)` use `/logical-routers/<logical-router-id>/routing/routing-table` for RIB and `/logical-routers/<logical-router-id>/routing/forwarding-table` for FIB. Returns the route table in CSV format for the logical router on a node of the given transport-node-id. Query parameter "transport_node_id=<transport-node-id>" is required. Query parameter "source=cached" is not supported.
- `UpdateServiceProfile (/service-profiles/<service-profile-id>)` modifies the specified service profile. PUT request must include the resource_type parameters. Modifiable parameters include description and display_name. Other parameters might be modifiable, depending on the specified service type.
- `DeleteService (/services/<service-id>)` deletes the specified logical router service.

- ListServices (/services) returns information about all configured logical router services that can be applied to one or more logical router ports. You must create a service-profile before you can create a service. Currently, only DhcpRelayService is supported.
- DeleteServiceProfile (/service-profiles/<service-profile-id>) deletes the specified service profile.
- ListServiceProfiles (/service-profiles) returns information about all service profiles. A service profile is a configuration that you can use to create a service, which is then applied to one or more logical router ports. Currently, only the DhcpRelayProfile is supported.
- CreateServiceProfile (/service-profiles) creates a service profile, which can then be used to create a service. Services are then applied to one or more logical router ports.
- DeleteLicense (/licenses/<license-key>) use POST /licenses?action=delete API instead.
- GetLicense (/license) use the GET /licenses API instead.
- ReadServiceProfile (/service-profiles/<service-profile-id>) returns information about the specified service profile.
- UpdateLicense (/license) use the POST /licenses API instead
- CreateService (/services) creates a service that can be applied to one or more logical router ports. For some service types, you must create a service-profile before you can create a service.
- UpdateService (/services/<service-id>) modifies the specified logical router service. The resource_type parameter is required. The modifiable parameters depend on the service type.
- ReadService (/services/<service-id>) returns information about the specified service.
- GetLicenseByKey (/licenses/<license-key>) use GET /licenses API instead.
- GetLogicalRouterRouteTable (/logical-routers/<logical-router-id>/routing/route-table) use /logical-routers/<logical-router-id>/routing/routing-table for RIB and /logical-routers/<logical-router-id>/routing/forwarding-table for FIB. Returns the route table for the logical router on a node of the given transport-node-id. Query parameter "transport_node_id=<transport-node-id>" is required. Query parameter "source=cached" is not supported.
- PerformNodeAction (/fabric/nodes/<node-id>) supported fabric node actions are enter_maintenance_mode, exit_maintenance_mode for EdgeNode. Use TransportNode maintenance mode API to update maintenance mode, refer to "Update transport node maintenance mode".

Deprecated Type Definitions

- LogicalService
- LogicalServiceResourceTypes resource types of logical services.
- NodeActionParameters fabric node action parameters.
- ServiceProfileResourceTypes resource types of service profiles.

Deprecated API Property Definitions

- BgpNeighbor.filter_in_routemap_id use 'address_family' instead.
- BgpNeighbor.remote_as use 'remote_as_num' instead.
- BgpNeighbor.filter_out_ipprefixlist_id use 'address_family' instead.

- BgpNeighbor.filter_out_routemap_id use 'address_family' instead.
- BgpNeighbor.source_address do not provide a value for this field. Use source_addresses instead.
- BgpNeighbor.filter_in_ipprefixlist_id use 'address_family' instead.
- HostSwitch.static_ip_pool_id the ID of configured Static IP Pool. If specified allocate IP for Endpoints from the Pool. Else assume IP is assigned for Endpoints from DHCP. This field is deprecated, use ip_assignment_spec field instead.
- AddControllerNodeSpec.control_plane_server_certificate do not supply a value for this property.
- BgpConfig.as_number use 'as_num' instead.
- TransportNode.host_switches use 'host_switch_spec' instead. Property 'host_switches' can only be used for NSX-T managed transport nodes. 'host_switch_spec' can be used for both NSX-T managed or manually preconfigured host switches.

Resolved Issues

The resolved issues are grouped as follows.

- [General Resolved Issues](#)
- [Installation Resolved Issues](#)
- [NSX Manager Resolved Issues](#)
- [NSX Edge Resolved Issues](#)
- [Logical Networking Resolved Issues](#)
- [Security Services Resolved Issues](#)
- [Operations and Monitoring Services Resolved Issues](#)
- [KVM Networking Resolved Issues](#)
- [API Resolved Issues](#)
- [Documentation Errata Resolved Issues](#)
- [Upgrade Resolved Issues](#)

General Resolved Issues

- **Issue 1901714: IPv6 traffic is dropped when SpoofGuard is enabled**

If Spoofguard is turned ON for logical-ports, IPv6 packets might be dropped even though the global-scope IP is added under "Manual Address-Binding" whitelist.

Workaround: Add the global scope IPv6 addresses and link-local IPv6 addresses as Manual Address-Binding whitelist if Spoofguard is enabled.

- **Issue 1948580: Rebooting RHEL transport node loses virtual tunnel endpoint DHCP IP address**

If transport node configuration uses DHCP IP for virtual tunnel endpoint, rebooting the transport node does not get DHCP IP address.

Workaround: After the transport node is booted completely, restart dhclient to virtual tunnel endpoint, interface to provide the DHCP IP address to the interface.

- **Issue 1958277: Iptable rule for external syslog server does not persist after appliance upgrade**

After the Management Plane, NSX Controller, and NSX Edge node upgrade from NSX-T

1.1 to NSX-T 2.0, all custom iptable rules automatically added by NSX-T CLI or API while setting external syslog servers are not added. As a result, the configured syslog exporters are unable to forward logs to the external server. In addition, in NSX Manager, dropped packets are logged in the /var/log/iptables.log, which leads to the saturation of the /var/log/iptables.log file after upgrading if the external syslog server was configured in NSX-T 1.1.

Workaround: Delete the external syslog servers after the appliance upgrade and add again these external syslog servers.

- **Issue 1958295: Key Manager registration might fail if hypervisor and Key Managers are registered together**

In cases when both hypervisor and Key Managers are registered together or in parallel, the Key Manager registration might fail.

Workaround: Register hypervisor and Key Manager separately.

- **Issue 1958302: On ESX 6.5, NSX-Exporter fails due to out-of-memory when under heavy flow traffic**

When there are more than 200 connection-per-second on a vNIC, the ESX vsip kernel drops the flow records due to a full queue. This causes the NSX-Exporter to miss the flow updates and those flows are accumulated in the process until restart. This leads to an out-of-memory condition.

Workaround: None.

After the fail, the watchdog restarts the process. Rule stats are off under such conditions.

- **Issue 1905370: Increased number of IP addresses that can be discovered via IP discovery and bigger SpoofGuard port whitelist size**

The number of IP addresses that can be discovered and the size of SpoofGuard port whitelist entries that can be configured to be populated automatically has increased. The supported number of maximum SpoofGuard port whitelist entries is 128. These entries can be manually or dynamically learned using VMware Tools, DHCP, or ARP.

You can exceed 128 port whitelist entries in the following scenarios.

Let A: corresponds to Manual PortWhitelist entries.

Let B: corresponds to all Discovered entries (DHCP, ARP, and VMTOOLS).

Let C: corresponds to SwitchWhitelist through API.

One of the following conditions must be true:

A <= 128, where B = C = 0 (or)

B <= 128, where A = C = 0 (or)

C <= 128, where A = B = 0 (or)

A + C <= 128 (manual bindings override discovered bindings, so number of IPs in B is irrelevant) (or)

B + C <= 128, where A = 0

Workaround: None.

- **Issue 1951653: Logs no longer written to the /var/log directory when the /var/log partition is 80% full**

When the /var/log partition is 80% full, a bug in the Cron job modifies the group ownership of the /var/log directory, preventing any further logs from being written to the /var/log directory.

Workaround: Perform the following steps on all the NSX Manager, NSX Controllers, and NSX Edges right after installation.

1. View the varlog_disk_space_monitor.py file.
`cat /opt/vmware/bin/varlog_disk_space_monitor.py`
2. Locate `_MONITOR_LOG` and check if it is `/var/log/disk_space_monitor.log`.
3. Run this command, `awk 'if ($1=="_MONITOR_LOG") {print "_MONITOR_LOG = \"/var/log/nvpapi/disk_space_monitor.log\"";} else {print $0}}'`
`/opt/vmware/bin/varlog_disk_space_monitor.py > /tmp/varlog_disk_space_monitor.py ; mv /tmp/varlog_disk_space_monitor.py /opt/vmware/bin/`
4. View the varlog_disk_space_monitor.py file.
`cat /opt/vmware/bin/varlog_disk_space_monitor.py`
5. Locate `_MONITOR_LOG` and verify that the log shows `/var/log/nvpapi/disk_space_monitor.log`.
6. Check the group ownership of the /var/log directory.
7. If the group ownership is `www-data`, change the ownership to `syslog`.
`chown root:syslog /var/log`

- **Issue 1917672: Unsuccessful authentication attempts on the vIDM login page is not logged in the NSX-T syslog**

Unsuccessful attempts to login as a remote user on the vIDM login page and account lockout is not sent to the NSX-T syslog as a failure event entry.

Workaround: View the failed authentication attempts log for vIDM users at `./opt/vmware/horizon/workspace/logs/horizon.log`.

Installation Resolved Issues

- **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, log in to the CLI, and configure the network interface manually.

NSX Manager Resolved 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 1955778: High memory consumption on NSX Manager**

Default memory setting of the unified appliance is medium, 4vCPU and 16GB RAM. If it is a scale deployment of for example, 250 hypervisors, Management Plane runs at 90% of RAM after scaling up configuration.

Workaround: The general recommendation for scale environment is to assign a larger setting for unified appliance configured as Management Plane, 8vCPU and 32GB RAM.

- **Issue 1958317: User tab in user interface does not display the local users, admin, and audit**

The User tab in the user interface does not display the local users, admin, and audit.

Workaround: None.

- **Issue 1959887: vIDM configuration detail not saved when edited using Internet Explorer**

In the NSX-T UI, when you navigate to the **System > Users > Configuration** view and edit the vIDM configuration using Internet Explorer, the edited data might not be saved correctly.

Workaround: Use another Web browser.

NSX Edge Resolved Issues

- **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 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 1941888: NSX Edge upgrade fails when the NSX Edge uptime is longer than five days**

NSX Edge upgrade fails with the following error message:

```
[Edge UCP] Upgrade Agent on Edge node <Edge Node ID> is unreachable. Restart the Upgrade agent service and check network connectivity.}, ]
```

Workaround: Manually restart the NSX Edge upgrade service to proceed with the upgrade. On the NSX Edge, type `restart service nsx-upgrade-agent`.

- **Issue 1923325: Traffic disruption might occur during NSX Edge node split brain healing in some configurations**

NSX-T 2.0 introduced a new non-preemptive failover mode for Tier-1 logical routers on NSX Edge nodes. When deployed in non-preemptive mode with firewall or NAT services configured, traffic disruption might occur after a split-brain event occurs and then heals. Split brain occurs when there is a network partition and both logical routers become active.

Workaround: Configuring the Tier-1 logical router in preemptive mode avoids traffic disruption.

- **Issue 1955002: Host switch name length limitation**

When creating a transport zone, if the provided host switch name length exceeds 26 bytes, the operation fails internally. As a result, the UI is unable to display the list of transport nodes for a given logical router.

Workaround: Limit the length of the host switch name to 26 bytes or less. Note that non-ASCII characters take between two to four bytes of space per character in the UTF-8 encoding. Using a host switch name comprised of non-ASCII characters further limits the number of characters that can be accommodated in the 26 bytes space.

Logical Networking Resolved Issues

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

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

- **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-nsxa.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 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 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 1935535: Container logical-ports are not listed under "Effective Logical Port" member of an NSGroup if the respective Containerhost-VM is added as a member of the NSGroup**

When a containerhost-VM is added as a member of a NSGroup via any membership criteria, the logical-ports on the containerhost-VM become the effective members of the same NSgroup, but the container-logical-ports from the containers on the same containerhost-VM do not become effective members of the NSGroup.

Workaround: Add the container-logical-ports to a NSGroup using available entities other than "Virtual Machines".

- **Issue 1955845: Edit transport node to change the transport zone fails, when the same physical NIC is retained**

After editing a transport node with a new transport zone and retaining the same physical NIC, the edit operation shows the status of **Partial Success** with the error message, Host configuration: Physical nics are in use: [vmnic1]. Physical nics available to host-switch [00 84 9b 21 55 c5 49 f0-a3 d7 76 dd a0 41 66 76]: vmnic2

Workaround: Delete the transport node and re-create it with the new transport zone NIC.

Security Services Resolved 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 disable the firewall rule.

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

Workaround: None.

- **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.

Operations and Monitoring Services Resolved Issues

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

When 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 has not been restarted since the initial small number of hosts, if NSX Manager loses connectivity with a large number of hosts, it takes a long time for NSX Manager to synchronize with those hosts and discover the VMs that run on them. The log file has messages such as the following: Processing sync_init requests, batch size <calculated batch size>

Workaround: Reboot the 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 1956097: Router name is truncated when issuing get logical-router command**
Running the `get logical-router` command on NSX Edge NSX CLI might show the truncated router name to fit within the 16-byte limit, which is enforced internally. This does not impact the actual name, as that remains unaffected. However, if the router names share a common prefix, it might be harder to distinguish among them.

Workaround: Rely on other attributes like UUID, VRF, LR-ID, Type etc., to uniquely identify a router. There is no impact to the actual name, you can find the full name in the UI and other locations where the name is exposed.

KVM Networking Resolved Issues

- **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

API Resolved Issues

- **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.

Documentation Errata Resolved Issues

- **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 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 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

Upgrade Resolved Issues

- **Issue 1953721: NSX Edge upgrade fails intermittently**
 During the NSX Edge upgrade, the upgrade process might pause due to failure in the NSX Edge upgrade on first attempt.

 Workaround: After the NSX Edge upgrade fails on the first attempt, re-attempt the upgrade from the UI or API.

Known Issues

The known issues are grouped as follows.

- [General Known Issues](#)
- [Installation Known 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](#)
- [Upgrade Issues](#)
- [Load Balancer Known Issues](#)
- [NSX Container Plug-in \(NCP\) Known Issues](#)

General Known Issues

- **Issue 1842511: Multihop-BFD not supported for static routes**

In NSX-T 2.0, BFD (Bi-Directional Forwarding Detection) can be enabled for a (MH-BGP) multihop BGP neighbor. The ability to back a multihop static route with BFD is not configurable in NSX-T 2.0, only BGP. Note that if you have configured a BFD backed multihop BGP neighbor and configure a corresponding multihop static route with the same nexthop as the BGP neighbor, the BFD session status affects both the BGP session as well as the static route.

Workaround: None.

- **Issue 1931707: Auto-TN feature requires all hosts in the cluster to have the same pnic setup**

When the auto-TN feature is enabled for a cluster, a transport node template is created to apply to all hosts in this cluster. All pnic in the template must be free on all hosts for TN configuration or the TN configuration might fail on those hosts whose pnic were missing or occupied.

Workaround: If the TN configuration failed, reconfigure the individual transport node for the correction.

- **Issue 1909703: NSX admin is allowed to create new static routes, NAT rules and ports in a router created by OpenStack directly from backend**

As part of RBAC feature in NSX-T 2.0, resources like Switches, routers, Security Groups created by the OpenStack plugin cannot be deleted or modified directly by NSX admin from the NSX UI/API. These resources can only be modified/deleted by the APIs sent through the OpenStack plugin. There is a limitation in this feature. Currently NSX admin is only stopped from deleting/modifying the resources created by OpenStack, although admin is allowed to create new resources like static routes, NAT rules inside the existing resources created by OpenStack.

Workaround: None.

- **Issue 1957072: Uplink profile for bridge node should always use LAG for more than one uplink**

When using multiple uplinks that are not formed to a LAG, the traffic is not load balanced and might not work well.

Workaround: Use LAG for multiple uplinks on bridge nodes.

- **Issue 1959647: Using a database server alias name to create a DSN might cause the installation of vCenter Server to fail**

When you use a database server alias name to create a DSN, the installation of vCenter Server with an external Microsoft SQL database fails. The following error appears during the installation of the inventory service: An error occurred while starting invsvc.

Use the IP address or the host name of the database server to create a DSN.

- **Issue 1775315: CSRF attack occurs when the Postman client is opened from Web browser**

For API calls made using Postman, CURL, or other REST clients, you must explicitly provide the XSRF-TOKEN header and its value. The first API call using remote authN or call to /api/session/create(local authN) carries the XSRF-Token in the response object. Subsequent API calls carry the token value in XSRF-TOKEN header as part of the request.

Workaround: Add X-XSRF-TOKEN header.

- **Issue 1970750: Transport node N-VDS profile using LACP with fast timers are not applied to vSphere ESXi hosts**

When an LACP uplink profile with fast rates is configured and applied to a vSphere ESXi transport node on NSX Manager, the NSX Manager shows that the profile is applied successfully, but the vSphere ESXi host is using the default LACP slow timer.

On the vSphere hypervisor, you cannot see the effect of lacp-timeout value (SLOW/FAST) when the LACP NSX managed distributed switch (N-VDS) profile is used on the transport node from the NSX manager.

Workaround: None.

- **Issue 1989407: vIDM users with the Enterprise Admin role cannot override object protection**

vIDM user with the Enterprise Admin role cannot override object protection and cannot create or delete Principal Identities.

Workaround: Log in with the Admin privileges.

- **Issue 1989412: Domain deletion when NSX Manager is not reachable is not reflected when connectivity is restored**

If a domain is deleted from Policy when the NSX manager is not reachable, after the connection is restored to the NSX Manager, the firewall and corresponding rules to the deleted domain still exist.

Workaround: Do not delete a domain from Policy when NSX Manager is not reachable.

- **Issue 1998217: HyperBus interface vmk50 might be missing on vSphere ESXi causing container creation failure**

Container is not created because the HyperBus interface vmk50 might be missing on vSphere ESXi.

Workaround: Complete the following steps.

1. Retrieve the vmk50 port ID using CLI on vSphere ESXi
`net-dvs | grep vmk50 -C 10`
2. Create the vmk50 interface on vSphere ESXi.
`esxcli network ip interface add -P <port-id from step-1> -s DvsPortset-0 -i vmk50 -N hyperbus`
3. Assign an IP address to the vmk50 interface.
`esxcli vmknic -i 169.254.1.1 -n 255.255.0.0 -s DvsPortset-0 -v <port-id from step-1> -N hyperbus`

- **Issue 2018478: Attempting to remove a widget from the dashboard causes a crash with stack trace error**

Custom dashboard user interface changes such as, removing a widget from multiple widgets causes the user interface to crash with a stack trace error.

Workaround: Complete the following steps.

1. Create a widget.
2. Locate the multiple widgets you want to modify.
3. Add a reference to the newly created widget in the multiple widget.

Installation Known Issues

- **Issue 1944678: NSX-T Unified appliance requires valid role type**

When the NSX-T Unified appliance is deployed in KVM without any specified role or an invalid role type, it is deployed in an unsupported configuration with all the roles enabled.

Workaround: A valid role type, nsx-manager, is required as a deployment parameter.

- **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`

- **Issue 1906410: Attempting to delete the host from the UI without first deleting the transport node, causes the host go into an inconsistent state**

Attempting to delete the host from the UI without first deleting the transport node, causes the host to go into an inconsistent state. If you attempt to delete the transport node while the host is in the inconsistent state, the UI does not allow you to delete this host.

Workaround:

1. Before deleting the transport node, power-off all the tenant VMs deployed on this transport node.
2. Remove the transport zone from the transport node.
3. Delete the transport node.
4. If the transport node is deleted successfully then delete the respective Host.

If the transport node deletion fails, complete the steps in the

KB <https://kb.vmware.com/s/article/52068>.

- **Issue 1957059: Host unprep fails if host with existing vibns added to the cluster when trying to unprep**

If vibns are not removed completely before adding the hosts to the cluster, the host unprep operations fails.

Workaround: Make sure that vibns on the hosts are removed completely and restart the host.

- **Issue 1958308: Host preparation or transport node creation fails when host is in lockdown mode**

Host preparation or transport node creation fails when host is in lockdown mode. The following error message appears: Permission to perform this operation was denied.

Workaround: Disable the lockdown mode on the host and retry host preparation.

- **Issue 1944669: Deploying NSX-T appliances on KVM**

When deploying NSX-T appliances on ESX, you can deploy small, medium, and large sizes with different RAM configurations. However, when deploying NSX-T appliances on KVM, the RAM allocation must be explicitly configured.

Workaround: Deploy a VM using KVM on Ubuntu.

```
sudo virt-install --vnc --import --name <VM_NAME> --ram 2048 --vcpus 2 --network=bridge:virbr0,model=e1000 -  
-disk path=/path/to/<IMAGE>.qcow2,format=qcow
```

The --ram command-line option must be in MB.

NSX-T Unified appliance

Small - 2 CPU, 8GB memory virt-install ... --ram 8192 --vcpus 2

Medium - 4 CPU, 16 GB memory virt-install ... --ram 16384 --vcpus 4

Large - 8 CPU, 32 GB memory virt-install ... --ram 32768 --vcpus 8

NSX Controller

4 CPU, 16GB memory virt-install ... --ram 2048 --vcpus 4

NSX Edge

Small - 2 CPU, 4G memory virt-install ... --ram 4096 --vcpus 2

Medium - 4 CPU, 8G memory virt-install ... --ram 8192 --vcpus 4

Large - 8CPU, 16G memory virt-install ... --ram 16384 --vcpus 8

- **Issue 1739120: After restarting Management Plane or Proton service in the Management Plane the Fabric node the deployment status becomes unresponsive**
When you add a new supported host on the Fabric page with host credentials, the status changes to **Install In Progress**. After restarting the Management Plane or the Proton service in the Management Plane, the deployment status of the host shows **Install In Progress** or **Uninstall In Progress** indefinitely.

Workaround: Delete the Fabric node with the unresponsive deployment status and add the host with credentials again.

NSX Manager Known Issues

- **Issue 1978104: Some pages in the NSX Manager user interface are not accessible on Internet Explorer 11**

The Dashboard, Getting Started workflows, and load balancer pages in the NSX Manager user interface are not accessible on the Windows machine that is running Internet Explorer 11.

Workaround: Use the Microsoft Edge, Google Chrome, or Mozilla Firefox browsers on your Windows machine to view the NSX Manager pages.

- **Issue 1950583: NSX Manager scheduled backup might fail after system upgrade to NSX-T 2.0.0**

Some NSX-T environments would fail to execute scheduled backup after upgrading from previous version of NSX-T to 2.0.0. This issue is due to a change in SSH fingerprint format from the previous releases.

Workaround: Reconfigure scheduled backup.

- **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 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.

- **Issue 1928376: Controller cluster member node degraded status after restoring NSX Manager**

Controller cluster member node might become unstable and report degraded health status if the NSX Manager is restored to a backup image that was taken before this member node was detached from the cluster.

Workaround: If cluster membership changes, make sure a new NSX Manager backup is taken.

- **Issue 1954293: vMotion of VMs connected to logical switches fails during Management Plane upgrade**

While Management Plane is being upgraded, if you attempt to vMotion a VM connected to a logical switch, the vMotion fails.

Workaround: Wait for Management Plane upgrade to finish and retry the vMotion.

- **Issue 1954297: If NSX Manager's restore is done, and any new non-VC managed ESX host is registered with NSX Manager and its VMs are connected to existing Logical Switches, then on ESX hosts' MOB, MAC address for VM becomes blank**

If NSX Manager restore is done, and any new non-VC managed ESX host is registered with Management Plane and its VMs are connected to existing Logical Switches, then on ESX hosts' MOB, MAC address for VM becomes blank.

This does not have any effect on VM's Inventory with respect to MAC on NSX Manager.

Workaround: None.

- **Issue 1954986: The license key is shown in the logs when the key is deleted from the UI**

The NSX license key is shown in /var/log/syslog as follows:

```
<182>1 2017-03-24T05:03:47.008Z bb-mgr-221 NSX - SYSTEM [nsx@6876 audit="true" comp="nsx-manager" reqId="3d146f2b-fa34-460f-8ac3-56e3c7326015" subcomp="manager"] UserName:'admin', ModuleName:'License', Operation:'DeleteLicense, Operation status:'success', New value: ["<license_key>"]  
<182>1 2017-03-24T05:03:47.009Z bb-mgr-221 NSX - - [nsx@6876 audit="true" comp="nsx-manager" subcomp="manager"] UserName:'admin', ModuleName:'Batch', Operation:'RegisterBatchRequest, Operation status:'success', New value: [{"atomic":false} {"request": [{"method":"DELETE","uri":"/v1/licenses/<license_key>"}]}
```

If the appliance is configured to send logs to an external log collector, then the key value is visible to any authorized user on the external log collector as well.

Workaround: None.

- **Issue 1956055: Local admin user cannot access tech support bundle from UI when the Management Plane datastore is down**

Local admin user cannot access Tech Support bundle from UI when the Management Plane datastore is down.

Workaround: If the NSX-T UI is not working, use the CLI or API to generate a support bundle.

- **Issue 1956088: Change to Firewall UI view while the rule set in the view has filtering applied might be lost before Saving to Manager if the filters are cancelled**

Change to Firewall UI view while the rule set in the view has filtering applied might be lost

before Saving to Manager if the filters are cancelled

Workaround: None.

- **Issue 1957165: Loading the last page in a search result set that includes 10,040 or more records yields a Search exception**

In a large environment that could return 10,040 or more possible objects for a search query, you might see an exception when trying to load the last few records in the result set from the UI listing.

Workaround: Narrow the search criteria.

- **Issue 1928447: Hypervisors with duplicate virtual tunnel endpoint IP addresses are not logged in the Management Plane node syslog**

Hypervisors with duplicate virtual tunnel endpoint IP addresses are not logged in the Management Plane node syslog. Make sure that unique IP addresses are assigned to the virtual tunnel endpoints of hypervisors and the uplink interfaces of the NSX Edge nodes.

Workaround: None.

- **Issue 1932987: After restoring the Management Plane, the connection between Management Plane and Key Manager Server fails**

When you detach the Key Manager Server from the Management Plane, restore the Management Plane, and attempt to reattach the Key Manager Server, the connection fails.

Workaround: None.

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 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 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.

- **Issue 1955830: Upgrade from NSX-T 1.1 to NSX-T 2.0 fails when the NSX Edge cluster name contains high or non-ASCII characters**

When an NSX Edge cluster is named using high or non-ASCII characters in the NSX-T 1.1 setup, upgrading from NSX-T 1.1 to NSX-T 2.0 fails with an infinite loop error.

Workaround: Rename the NSX Edge clusters to remove high or non-ASCII characters on the NSX-T 1.1 setup instance before upgrading.

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 1940046: When the same static route is added and advertised in multiple Tier-1 logical routers, east-west traffic fails**

If the same static route is added and advertised in multiple Tier-1 logical routers, east-west traffic fails.

Workaround: Static routes should be advertised only from the originating Tier-1 logical router if the prefix resides behind a connected network of the Tier-1 distributed router.

- **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 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 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 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.

- **Issue 1954997: Transport Node deletion fails if VMs on the transport node are connected to Logical Switch at the time of deletion**

1. Fabric Node and Transport Node are created.
2. Attach VIFs to logical switch.
3. Delete transport node without removing VIF attachments to Logical Switch fails.

Workaround: Delete all VIF attachments of the corresponding VMs on the transport node, which need to be removed from NSX, then delete the transport node.

- **Issue 1958041: BUM traffic might not work for Layer 3 flow across physical Layer 2 segments when ESX hypervisor has multiple uplinks**

If all of the following conditions are met, it is possible that BUM traffic from source hypervisor across logical router does not reach the destination hypervisor.

- ESX has multiple uplinks
- Source and destination VMs are connected via logical router
- Source and destination hypervisor are on different physical segments
- Destination logical network is using MTEP replication

This occurs because the BFD module might not have created the session, which means MTEP selection for destination logical network might not have occurred.

Workaround:

1. Start a VM in destination Logical Network on Destination hypervisor or any another hypervisor in same destination Layer 2 physical segment.
2. Change the replication mode of destination logical network to source replication.
3. Disable BFD in the Transport Zone.

- **Issue 1966641: If you add a host and configure it as a transport node, the node status appears as Down if it is not part of a logical switch**

After adding a new host and configuring it as a transport node or when configuring an upgrade plan to NSX-T 2.1, the transport node status appears as Down in the user interface if it is not part of a logical switch.

Workaround: Create a logical switch for the transport node so that tunnels establish connectivity with other transport nodes in that logical switch.

- **Issue 2015445: Firewall state on the active service router might not be duplicated on**

the newly active service router

Tenant logical router (TLR) might have multiple failovers from NSX Edge1 to NSX Edge2 and from NSX Edge2 to NSX Edge1. Firewall or NAT flow states are synchronized between active/standby TLR service routers. When the TLR is configured in a non-preemptive failover mode, the synchronization occurs before the first failover, but does not occur between first and the subsequent failover. As a result, at the second failover, the TCP traffic can time out. This problem does not occur with TLR configured in preemptive mode.

Workaround: Change the logical router configuration to the preemptive mode.

- **Issue 2016629: RSPAN_SRC mirror session fails after migration**

When a VM connected to a port assigned for RSPAN_SRC mirror session is migrated to another hypervisor, and there is no required pNic on the destination network of the destination hypervisor, then the RSPAN_SRC mirror session fails to configure on the port. This failure causes the port connection failure but the vMotion migration process succeeds.

Workaround: To restore port connection failure, complete either one of the tasks.

- Remove the failed port and add a new port.
- Disable the port and enable it.

The mirror session fails to configure, but the port connection is restored.

Security Services Known Issues

- **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 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.

- **DHCP release and renew packets not reaching the DHCP Server when the client is on a different network and routing service is provided by a guest VM**

NSX-T cannot distinguish if a VM is acting as a router, so it is possible that unicast DHCP packets getting routed using a router VM get dropped as the CHADDR field in the packet does not match the source MAC. The CHADDR has MAC of the DHCP client VM, whereas the Source MAC is that of the router interface.

Workaround: If a VM behaves like a router, disable **DHCP Server Block** in the switch security profiles applied to all the VIFs of the router VM.

- **Issue 2008882: For Application Discovery to work properly, do not create a security group that spans multiple hosts**

If one security group has VMs that span across multiple hosts, the Application Discovery session might fail.

Workaround: Create a security group of VMs on one host only. You can create multiple security groups for several hosts and run Application Discovery on them separately.

Operations and Monitoring Services Known Issues

- **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/<|port-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.

- **Issue 1955822: License Usage reporting csv file should also include CPU and VM entitlement along with actual usage**

When querying for licensing usage report (through API/UI), the data contains current usage only.

Workaround: Query for usage limits allowed by current license(s) through the UI or REST API:

Method: GET; URI: `/api/v1/licenses`

- **Issue 1957092: Failed to initialize NSX Controller cluster as error occurs in loading docker image**

The initialize control-cluster command fails with an error message, Control cluster activation timed out. Please try again. There is also the following log information in the syslog:

```
<30>1 2017-08-03T22:52:41.258925+00:00 localhost load-zookeeper-image 1183 - - grpc: the connection is unavailable.
```

Workaround: Run the initialize control-cluster command again.

KVM Networking Known Issues

- **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-1.1.0.0.0.4690847-1.el7.x86_64
rpm -e nsx-aggservice-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-cli-1.1.0.0.0.4690892-1.el6.x86_64
rpm -e nsx-da-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-host-1.1.0.0.0.4690932-1.x86_64 rpm -e nsx-host_node_status_reporter-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-lldp-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-logical_exporter-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-mpa-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-netcpa-1.1.0.0.0.4690924-1.el7.x86_64 rpm -e nsx-sfhc-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-support-bundle-client-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsx-transport_node_status-1.1.0.0.0.4690845-1.el7.x86_64
rpm -e nsxa-1.1.0.0.0.4690845-1.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 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_conntrack: table full, dropping packet
Apr 26 11:45:44 prmh-nsx-perf-server149 kernel: [1625289.958436] nf_conntrack: 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](#).

- **Issue 2002353: Using Linux Network Manager to manage a KVM host uplinks is not supported**

NSX-T manages all the NICs on KVM hosts that are used for N-VDS. Configuration error occurs when the Network Manager is also enabled for these uplinks.

Workaround: For Ubuntu hosts, exclude the NICs to be used for NSX-T from the Network Manager.

Prior to enabling NSX-T on a Red Hat host, modify the NIC configuration script in `/etc/sysconfig/network-scripts` as `NM_CONTROLLED="no"`. If NSX-T has already been enabled for the host, make the same script modification, and restart networking for the host.

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.

- **Issue: 2025624: Splunk dashboards stuck while loading or graphs on the dashboards are blank**
Splunk is fetching the old version of `nsx_splunk_app` because the HTML template is incorrectly pointing to the previous path of the query script. So the dashboards are executing old queries which contain fields such as `vmw_nsxt_comp`, `vmw_nsxt_subcomp`, and `vmw_nsxt_errorcode`, and these fields are named differently in the newer version of the query script. As a result, the queries will return empty results and the dashboards will be blank.

Workaround: Rename the file *nsx_splunk_app.spl* to *logger.spl*, upload the renamed file to Splunk Enterprise Server, and restart the server. This will install the NSX Splunk App version 1.0, and its dashboards will work correctly with the old queries.

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 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 API GET `/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 does not 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.

- **Issue 1954990: Realization API inaccurate status return**

If you use a Realization API to check the realization status for all APIs executed before a barrier, the return status by the Realization API can be misleading relative to the actual status. Because of the complexity of the execution of the DFW inside the Management Plane, DFW API can slip after the barrier they are supposed to follow which leads to this inaccuracy.

Workaround: Do not rely on the Realization API to assess actual realization.

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 do not require the XSRF header.

See the Session-based Authentication section of the NSX-T API Guide for more information.

Upgrade Issues

- **Issue 1930705: vMotion of VMs connected to the logical switches fails during the Management Plane upgrade**

During the Management Plane upgrade, attempting to vMotion VMs connected to a logical switch fails.

Workaround: Wait until the Management Plane upgrade completes and retry the vMotion process.

- **Issue 1944731: DHCP leases might have conflicting records if numerous requests**

are served by the first upgraded NSX Edge during the upgrade of the second NSX Edge

If numerous requests are served by first upgraded NSX Edge during the upgrade of the second NSX Edge, then the DHCP leases might have conflict records.

Workaround: Do not use the DHCP service during the upgrade or manually release the DHCP offer retrieved during upgrade.

- **Issue 1847884: Do not make NSX-T related changes until the upgrade process for the Management Plane has completed**

Performing any changes such as, creating, updating, or deleting a transport zone, transport node, or logical switches during the Management Plane upgrade might corrupt the Management Plane, leading to NSX Edge, host, and data path connectivity failures.

Workaround: Wait until the upgrade completes. Delete the changes made during the upgrade and reconfigure the changes you made earlier.

- **Issue 2005423: KVM nodes upgraded from a previous NSX-T version are not automatically changed to use balance-tcp**

NSX-T does not automatically modify the bond mode of an upgraded KVM host uplink from active-backup to balance-tcp.

Workaround: Edit the transport node, even if there are no configuration changes, to correct the mode setting.

- **Issue 2005709: Upgrade coordinator page becomes inaccessible when you use the NSX Manager FQDN**

When you use the NSX Manager FQDN to open the NSX Manager user interface, the following error message appears in the Upgrade Coordinator page, This page is only available on the NSX Manager where Upgrade Coordinator is running. To enable the service, run the command "set service install-upgrade enabled" on the NSX Manager. If the install-upgrade service is already enabled, try disabling it using "clear service install-upgrade enabled" and then enable it again."

Workaround: Use the NSX Manager IP address to access the user interface.

- **Issue 2022609: Managed hosts are treated as unmanaged host in the upgrade coordinator**

If an environment has more than 128 managed hosts, during the upgrade process the hosts that were part of a cluster appear in the unmanaged ESXi group.

Workaround:

- Manually upgrade the unmanaged ESXi hosts above 128.
- Navigate to the /opt/vmware/upgrade-coordinator-tc-server/webapps/upgrade-coordinator/WEB_INF/classes/config.properties file, change the value of upgrade.host.service.hostNodeListPageSize from 128 to 512, and restart the upgrade coordinator /etc/init.d/upgrade-coordinator restart.

Load Balancer Known Issues

- **Issue 195228: Weighted round-robin and weighted least connection algorithms**

might not distribute traffic properly after a configuration is changed and reloaded

Servers lose connection when a weighted round-robin or weighted least connection configuration is changed and reloaded. After the connectivity loss, the historical traffic distribution information is not preserved which leads to traffic being distributed improperly.

Workaround: None.

- **Issue 2010428: Load balancer rule creation and application limitations**

In the user interface, you can create a load balancer rule from the virtual server only. Load balancer rules created using REST API cannot be attached to the virtual server in the user interface.

Workaround: If you created a load balancer rule using REST API, attach that load balancer rule to the virtual server using REST API. The rules created using REST API now appear in the virtual server from the user interface.

- **Issue 2016489: LCP fails to configure the default certificate when the server name indication is selected**

Default certificate ID should be set first in the certificate list when multiple certificate IDs are used in server name indication (SNI) to avoid LCP ignoring the default certificate.

Workaround: The default certificate should be first in the SNI certificate list.

- **Issue 2018629: Health check table not showing the updated monitor type for the NS group pool**

When you create static and dynamic NS group pools with the same members with a monitor type and change that monitor type on dynamic pool, the dynamic pool health check does not appear in the health check table.

Workaround: Create a dynamic group pool with a monitor and then create a static pool with the same members and a monitor for the health check table to show both of the pool monitoring.

- **Issue 2020372: Passive health check does not consider the pool member down after the maximum fall count is reached**

Passive health check requires additional fall count value than configured to consider the pool member down.

NSX Container Plug-in (NCP) Known Issues

- **Issue 2051265: NCP supports only equality-based label selectors for network policies**

Network policies select pods/namespaces using label selectors. Currently, NCP only supports equality-based label selectors. Inequality-based and set-based label selectors are not supported.

Workaround: Use equality-based label selectors to create network policies.

- **Issue 2011712: NCP does not handle network policy modify events properly**

When NCP is running and a network policy modify event occurs, NCP might not process the event properly. For example, the isolation rule might not be correct, or the source IPset

rule might be missing.

Workaround: Delete and re-create the network policy.

Copyright © 2021 VMware, Inc. All rights reserved.