VMware NSX-T Data Center 2.3 Release Notes

VMware NSX-T Data Center 2.3 | 18 SEP 2018 | Build 10085361

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

NSX-T Data Center 2.3 is the incremental upgrade release that enhances the new multi-hypervisor platform delivered for cloud and containers.

The following new features and feature enhancements are available in the NSX-T Data Center 2.3 release.

Introducing NSX-T Data Center Support for Bare-Metal Hosts

Bare-metal support includes Linux-based workloads running on bare-metal servers and containers running on bare-metal servers without a hypervisor. NSX-T Data Center leverages the Open vSwitch, to enable any Linux host to be an NSX-T Data Center transport node.

- **Bare-Metal Server Support:** includes native compute workloads running RHEL 7.4, CentOS 7.4, and Ubuntu 16.0.4 operating systems to allow users to network bare-metal compute workloads over VLAN, overlay backed connections, and to enforce micro-segmentation policies (stateful Layer 4 enforcement) for Virtual-to-Physical and Physical-to-Physical communication flows.
- **Bare-Metal Linux Containers Support:** runs Docker Containers using RedHat OpenShift Container Platform on bare-metal Linux hosts with RHEL 7.4 or RHEL 7.5.

NSX Cloud Enhancements

- **Support for AWS Deployments:** NSX Cloud support for AWS workloads.
- **Automatic NSX Agents Provisioning in Azure VNETs**
- **VPN Support Between On-Premise to Public Cloud:** includes built-in VPN capabilities within the NSX Cloud Public Cloud Gateway using APIs. You can use the VPN capabilities to create IPSEC links between the following:
    - Managed compute Amazon VPCs/Azure VNets and third-party service VMs in transit Amazon VPCs/Azure VNets
    - Managed Amazon VPC/Azure VNET and an on-premise VPN device
- **Expanded OS Support for NSX Cloud Agent:** NSX Cloud supports RHEL 7.5 operating systems in the public cloud.

Security Services Support

Introducing Service Insertion at the Routing Tiers

- **Service Insertion Support on Tier-0 and Tier-1 Routers:** includes the ability to onboard third-party security solutions, deploy a High Availability third-party security solution at Tier-0 or Tier-1 or both and insert the third-party security solution via redirect policy.
  Check the VMware Compatibility Guide – Network and Security for the latest certification status of third-party solutions on NSX-T Data Center.

Distributed Firewall Enhancements

- **Multiple Section Support in NSX Edge Firewall**: adds multiple sections in the NSX Edge Firewall for ease of manageability
- **Firewall Rule Hit Count and Rule Popularity Index**: monitors rule usage and quick identification of unused rules for clean-up
- **Firewall Section Locking:** enables multiple security administrators to work concurrently on the firewall
- **Grouping Objects:** supports an object to be added to a group if it matches all five specified tags, which was previously two tags
- **Tag Length**: increases tag length value from 65 to 256 and tag scope from 20 to 128
- **Application Discovery:** discovers and categorizes (allowing for custom categorization by users too) applications installed inside guest VMs. Applications will contain details about executables, hash, publisher information, and install date.

Network and NSX Edge Services Support

- **Overlay Support for Enhanced Data Path Mode in N-VDS:** in conjunction with vSphere 6.7, the Enhanced data path mode in N-VDS for NSX-T Data Center 2.3 supports NFV style workloads requiring high-performance data path.
- **Support for Stateful NAT and Firewall Services on the Centralized Service Port**
- **API Support to Clear All DNS Entries on DNS Forwarder:** provides the ability to clear all the DNS cache entries in a single API call on a given DNS forwarder. This command is useful when a DNS server is giving wrong answers and to avoid waiting for the DNS entry timeout after the DNS server is fixed.
- **Load Balancer Enhancements**
    - **Support for Pre-Defined Cipher List:** Pre-defined SSL profiles for HTTPS VIP for higher security or performance.

- **Load Balancer Rule Enhancement:** new Load Balancer rules, *delete header action*, *SSL match condition*, and *Assign variable on match condition*.
- **Load Balancer Support on Stand-Alone Service Router:** provides the ability to deploy a load balancing service on a service router that does not have a router port.

User Interface Enhancements

- **New Language Support**: user interface now available in English, German, French, Japanese, Simplified Chinese, Korean, Traditional Chinese, and Spanish.
- **Enhanced Navigation and Home Page**: new home page highlights search and at-a-glance summary of the system.
- **Enhanced Search:** search includes type-ahead suggestions, which are accessible from the home page.
- **Network Topology Visualization**: NSX Policy Manager provides the ability to monitor communications from group-to-group, VM-to-VM, and process-to-process. You can visualize relationships between network objects such as, logical switches, ports, routers, and NSX Edges.

Operations and Troubleshooting Support

- **Install and Upgrade Enhancements**
  - **NSX-T Data Center in a Stateless vSphere Environment:** enables additional deployment options by providing support for stateless ESXi hosts that use vSphere Auto Deploy and Host Profiles. The feature support requires vSphere 6.7 U1 or higher.
  - **Support for NSX Edge VM and Bare-Metal to Co-Exist in the Same NSX Edge Cluster:** NSX Edge nodes VM and bare-metal can now exist in the same NSX Edge cluster to simplify the scaling of services hosted on the NSX Edge node, such as load balancer.
  - **Modular NSX-T Data Center Upgrade:** includes support for modular upgrade in the Upgrade Coordinator. You can upgrade only the NSX-T Data Center components that have changed in the new release version. This added functionality reduces the operational overhead of patching an NSX-T Data Center version.
- **Monitoring and Troubleshooting**
  - **ERSPAN for KVM Hypervisor:** includes support for port mirroring on KVM – ERSPAN Type II and III.
  - **Use Traceflow to and from Tier-0 Logical Router Uplinks:** provides the ability to generate traceflow traffic from the Tier-0 logical router uplinks and report the receiving of traceflow packets on Tier-0 logical router uplinks to simplify the troubleshooting operations to include the northbound interfaces of the NSX Edge nodes in traceflow reporting.
  - **CLI Support to Shut Down DPDK Ports on Bare-Metal Edge Node:** provides the ability to shut down a port claimed by DPDK on the bare-metal NSX Edge node to simplify port isolation during installation and troubleshooting operations.

OpenStack Neutron plugin Support
These features are supported from OpenStack Upstream Queens release onwards.

- **Ability for the Neutron Plugin to Provision Overlay Logical Switch Backed by Enhanced Datapath:** NSX Neutron plugin offers the ability to leverage Enhanced Data

Path mode for overlay, which used to be VLAN only. With this supports you can take advantage of the Enhanced datapath performance in addition to the OpenStack environment for instance, for the NFV related workload.

- **Support for Co-existence of NSX Products with OpenStack:** NSX Neutron Plugin now supports managing both NSX Data Center for vSphere and NSX-T Data Center simultaneously for an OpenStack implementation.
- **Ability to Consume VPN as a Service Feature in OpenStack:** support for OpenStack VPNaaS in the Neutron extension in OpenStack that introduces VPN feature set.

NSX Container Plug-in (NCP) Support

- **Concourse Pipeline to install NSX-T Data Center**
- **Annotation for Load Balancer SNAT IP**: SNAT IP for a load balancer is annotated in a Kubernetes service of type LoadBalancer, ncp/internal_ip_for_policy: <SNAT IP>, and added to the service's status, status.loadbalancer.ingress.ip: [<SNAT IP>, <Virtual IP>]. This IP can be used to create network policy which allows this IP CIDR.
- **Kubernetes Network Policy Enhancement:** provides the ability to select pods from different namespaces with Kubernetes network policy rules.
- **Kubernetes Load Balancer/SNAT Annotation Improvement**
  - If NCP fails to configure a load balancer for a service, the service will be annotated with ncp/error.loadbalancer.
  - If NCP fails to configure an SNAT IP for a service, the service will be annotated with ncp/error.snat.
- **Session Persistence of NSX-T Date Center Load Balancer for Kubernetes Ingress and OpenShift Routes**
- **Cleanup Script Enhancement**

Compatibility and System Requirements

For compatibility and system requirement information, see the NSX-T Data Center Installation Guide.

NSX-T Data Center in a Stateless vSphere Environment - for stateless ESXi hosts that use vSphere Auto Deploy and Host Profiles, the requirement is vSphere 6.7 U1 or higher.

NCP Compatibility Requirements:

Product	Version
NCP / NSX-T Data Center Tile for PAS	2.3.0
NSX-T Data Center	2.2, 2.3
Kubernetes	1.10, 1.11
OpenShift	3.9, 3.10
Kubernetes host VM OS	Ubuntu 16.04, RHEL 7.4, 7.5
OpenShift host VM OS	RHEL 7.4, RHEL 7.5
	OpsManager 2.1.x + PAS 2.1.x (except PAS 2.1.0)

PAS (PCF)	OpsManager 2.2.0 + PAS 2.2.0

General Behavior Changes

Default HA Mode for Tier-1 Logical Routers Changes from Preemptive to Non-Preemptive

When creating a Tier-1 logical router, the default HA mode was preemptive, leading to a traffic slowdown when the preferred NSX Edge node was going back online. With the new default HA mode to non-preemptive, the newly created Tier-1 logical routers do not experience this traffic slowdown. The existing Tier-1 logical routers are not be affected by this change.

Communication Change from Transport Node to NSX Controller

Due to changes in the communication from the transport node to NSX Controller, you must now open the TCP port 1235 for NSX-T 2.2 and higher. See the NSX-T Installation Guide.

When upgrading from NSX-T 2.1 to higher versions, both the TCP ports 1234 and 1235 must be open. After the upgrade is complete, the TCP port 1235 is in use.

API Reference Information

See NSX-T Data Center and NSX Policy deprecated API calls and properties.

The latest API reference is located in the NSX-T Data Canter Product Information.

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
- Load Balancer Resolved Issues
- Solution Interoperability Resolved Issues
- Operations and Monitoring Services Resolved Issues
- Upgrade Resolved Issues
- API Resolved Issues
- NSX Container Plug-in (NCP) Resolved Issues

General Resolved Issues

- **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.

- **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.

- **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.

- **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.

Installation Resolved Issues

- **Issue 1739120: After restarting the Management Plane or the Proton service in the Management Plane the fabric node deployment status becomes stuck**
  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.

- **Issue 1944669: Deploying NSX-T Data Center appliances on KVM requires specifying the exact memory size**
  When deploying NSX-T Data Center appliances on ESX, you can deploy small, medium, and large sizes with different RAM configurations. However, when deploying NSX-T Data Center appliances on KVM, the RAM allocation must be explicitly configured.

- **Issue 1944678: Deploying an NSX-T unified appliance requires a valid role type**
  When an NSX-T unified appliance is deployed in KVM without any specified role or with an invalid role type, it is deployed in an unsupported configuration with all the roles enabled.

- **Issue 1958308: Host preparation or transport node creation fails when host is in lockdown mode**
  Host preparation or transport node creation fails when the host is in lockdown mode. The following error message appears: Permission to perform this operation was denied.

NSX Manager Resolved Issues

- **Issue 1954923: 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.

- **Issue 1954927: After NSX Manager's restore is done, and a new non-VC managed ESX host is registered with NSX Manager and its VMs are connected to existing logical switches, then on the ESX host's MOB, the MAC address of the VMs are blank**
  After NSX Manager restore is done, and a new non-VC managed ESX host is registered with NSX Manager and its VMs are connected to existing logical switches, then on the ESX host's MOB, the MAC address of the VMs are blank.

- **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 when using Internet Explorer on a Windows machine.

- **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.

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

- **Issue 1957165: Loading the last page in a search result set that has 10,040 or more records causes an error**
  In a large environment that can return 10,040 or more objects for a search query, you might see an error when trying to load the last few records in the result set.

NSX Edge Resolved 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.

Logical Networking Resolved Issues

- **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.

- **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.

- **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.

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

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

- **Issue 1763576: Hypervisors are allowed to be removed as transport nodes even when they have VMs on the NSX-T Data Center network**
  NSX-T Data Center does not prevent you from deleting a transport node even when there are VMs on the node that are part of the network. The VMs lose connectivity after the transport node is deleted.

- **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.*

- **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.

- **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.

Security Services Resolved Issues

- **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.

- **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.

Load Balancer Resolved Issues

- **Issue 1995228: 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.

- **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.

- **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.

Solution Interoperability Resolved Issues

- **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. Therefore, 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.

Operations and Monitoring Services Resolved Issues

- **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.

Upgrade Resolved Issues

- **Issue 1847884: Do not make NSX-T Data Center 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.

- **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."

- **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.

- **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.

API Resolved Issues

- **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.

- **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.

- **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.

NSX Container Plug-in (NCP) Resolved Issues

- **Issue 2167491: NCP fails to start if the NSX-T load balancer has the maximum number of virtual servers**
  In the ConfigMap for NCP, you can set the size of the NSX-T load balancer to small, medium or large. The maximum number of virtual servers is 10 for a small load balancer, 100 for a medium, and 1000 for a large. If the load balancer has the maximum number of virtual servers, NCP will not start. To see if the load balancer has the maximum number of virtual servers, from the NSX-T Manager GUI, find the load balancer (it has a tag with the cluster's name) and count the number of virtual servers.

- **Issue 2160806: Updating the TLS spec of an active Ingress when NCP is not running is not supported**
  If NCP has assigned an external IP to an Ingress resource, and you update the TLS spec of the Ingress when NCP is not running, for example, by removing or changing the parameter secretName, NCP will not be aware of the changes. When NCP runs again, the certificate corresponding to the old secret still exists and will not be deleted.

Known Issues

The known issues are grouped as follows.

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 pnics 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 pnics in the template must be free on all hosts for TN configuration or the TN configuration might fail on those hosts whose pnics 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 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 2030784: Cannot log in to NSX Manager with remote username that contains non-ASCII characters.**
  You cannot log in to the NSX Manager appliance as a remote user with username containing non-ASCII characters.

  Workaround: Remote username should have ASCII characters when you log in to the NSX Manager appliance.
  Non-ASCII characters can be used if the remote username is set with non-ASCII characters in the Active Directory server.

- **Issue 2111047: Application Discovery not supported on VMware vSphere 6.7 hosts in the NSX-T 2.2 release**
  Running application discovery on a security group which has VMs running on a vSphere 6.7 host causes the discovery session to fail.

  Workaround: None

- **Issue 2157370: When configuring L3 Switched Port Analyzer (SPAN) with truncation, specific physical switch drops mirrored packets**
  When configuring L3 SPAN which includes GRE/ERSPAN with truncation, truncated mirrored packets are dropped because of the physical switch policy. A possible cause might be that the port is receiving packets where the number of bytes in the payload are not equal to type length field.

  Workaround: Remove the L3 SPAN truncation configuration.

- **Issue 216992: Mirrored packets with destination MAC address 02:50:56:56:44:52 from other hosts are dropped by the vSphere ESXi uplink**
  When host receives mirrored packets with destination MAC address 02:50:56:56:44:52 from other hosts, the vSphere ESXi uplink drops these mirrored packets.

  Workaround: None

- **Issue 2174583: In the Getting Started wizard, the Set Up Transport Nodes button does not work properly on the Microsoft Edge browser**
  In the Getting Started wizard, after you click the **Set Up Transport Nodes** button the Microsoft Edge web browser fails with a JavaScript error.

  Workaround: Use the Firefox or Google Chrome browser instead.

Installation Known Issues

- **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 vibs added to the cluster when trying to unprep**
  If vibs are not removed completely before adding the hosts to the cluster, the host unprep operations fails.

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

- **Issue 2106956: Joining two NSX Controllers of the same cluster to two different NSX Managers causes undefined datapath issues**
  Joining two NSX Controllers of the same NSX Controller cluster to two different NSX Managers causes undefined datapath issues.

  Workaround: Use the detach CLI command on the NSX Manager to remove the NSX Controller from the NSX Controller cluster. Reconfigure the NSX Controller cluster so that all the NSX Controllers in a cluster are registered with the same NSX Manager.

  See the NSX Controller Installation and Clustering section of the NSX-TData Center Installation Guide.

- **Issue 2106973: Initializing the NSX Controller cluster on all the NSX Controllers causes each of the NSX Controller to become a one node NSX Controller cluster resulting in undefined datapath connectivity issues**
  Avoid initializing the NSX Controller cluster on all the NSX Controllers which causes each of the NSX Controller to become a one node NSX Controller cluster resulting in undefined datapath connectivity issues. Initialize the NSX Controller cluster on only the first NSX

Controller and join the other NSX Controllers to the cluster by running the join control-cluster CLI command on the first NSX Controller.

Workaround: Reconfigure your NSX Controller cluster as described in the NSX Controller Installation and Clustering section of the NSX-T Data Center Installation Guide.

- **Issue 2114756: In some cases, VIBs are not removed when a host is removed from the NSX-T Data Center prepared cluster**
  When a host is removed from the NSX-T Data Center prepared cluster, some VIBs might remain on the host.

  Workaround: Manually uninstall VIBs from the host.

- **Issue 2059414: RHEL LCP bundle installation fails due to older version of python-gevent RPM**
  If a RHEL host contains a newer version of the python-gevent RPM, then the RHEL LCP bundle installation fails because NSX-T Data Center RPM contains an older version of python-gevent RPM.

  Workaround: Manually install the LCP bundle on the RHEL host if the host contains the latest version of python-gevent RPM.

  Complete the following steps:

  1. Extract the RHEL LCP bundle.
  2. Navigate to the LCP bundle folder.
  3. Delete libev, python-greenlet, and python-gevent RPMs from the LCP folder.
  4. Install the remaining RPMs. See the NSX-T Data Center Installation Guide.
- **Issue 2142755: OVS kernel modules fail to install depending on which minor RHEL 7.4 kernel version is running**
  OVS kernel modules fail to install on a RHEL 7.4 host running a minor kernel version 17.1 or above. The installation failure causes the kernel data paths to stop working which leads the appliance management console to become unavailable.

  Workaround: Upgrade the RHEL 7.4 kernel version. With admin privileges, run the script /usr/share/openvswitch/scripts/ovs-kmod-manage.sh on the host and reload the OVS kernel modules.

NSX Manager Known Issues

- **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 2128476: Scale setup with an inventory of more than 500 hosts, 1000 VMs, and 10000 VIFs might take about 30 minutes for full sync after hard reboot**
  After NSX Manager is rebooted, each host is synchronized with NSX manager so that the NSX Manager receives the latest data on the host, which includes information regarding VMs present on the host and VIFs present on the VMs. For a scale setup with an inventory that contains more than 500 hosts, 1000 VMs, and 10000 VIFs, full sync requires about 30 minutes.

  Workaround: Wait for the latest information to appear in the NSX Manager after hard reboot.

  Use the API api/v1/fabric/nodes/<nodeid>/status to check the last_sync_time property which indicates the latest sync time for a particular node.

- **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 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 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 2125725: After restoring large topology deployments, the search data becomes out of sync and several NSX Manager pages are unresponsive**
  After restoring NSX Manager with large topology deployments, the search data becomes out of sync and several NSX Manager pages display the error message, An unrecoverable error has occurred.

  Workaround: Complete the following steps:

  1. Log in to the NSX Manager CLI as administrator.
  2. Restart the search service.
     restart service search
     Wait for at least 15 minutes while the search service finishes fixing data discrepancies in the background.
- **Issue 2128361: CLI command to set the log level of the NSX Manager to the debug mode not working properly**
  Using the CLI command set service manager logging-level debug to set the log level of the NSX Manager to debug mode not collecting debugging log information.

  Workaround: Complete the following steps:

  1. Log in to the NSX Manager CLI as administrator.
  2. Run the command st e to switch to root user.
  3. Copy the log4j2.xml.default and log4j2.xml files.
     cp /opt/vmware/proton-tomcat/conf/log4j2.xml.default /opt/vmware/proton-tomcat/conf/log4j2.xml
  4. Change the ownership of the log4j2.xml file.
     chown uproton:uproton /opt/vmware/proton-tomcat/conf/log4j2.xml
- **Issue 1964681: The Hosts tab in Manager UI shows the status of a transport node host as Delete In Progress even after the host was deleted**
  From the Fabric > Nodes > Transport Nodes tab in Manager UI, after you successfully delete a transport node host, the Hosts tab still shows the status of the host as Delete In Progress.

  Workaround: Refresh the browser.

- **Issue 2169998: With the Chrome browser, clearing the browsing data when you are logged in to NSX Manger causes the manager UI to stop working**
  After you log in to NSX Manager using the Chrome browser, if you go to browser settings and clear all browsing data, including all basic and advanced, the browser will lose its connection to NSX Manager.

  Workaround: Do not clear the browsing data when logged in to NSX Manager.

NSX Edge Known Issues

- **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.

- **Issue 2122332: In some cases, SSH log in to a Bare Metal Edge does not work**
  Occasionally, the SSH log in to a Bare Metal Edge does not work.

  Workaround: Open a command prompt and navigate to the iLO driver. Restart the Edge SSH service.

- **Issue 2187888: Automatically deployed NSX Edge from the NSX Manager user interface remains in Registration Pending state indefinitely**
  Automatically deployed NSX Edge from the NSX Manager user interface remains in Registration Pending state indefinitely. This state causes the NSX Edge to become unavailable for further configuration.

  Workaround: Use CLI to manually register the NSX Edge with the NSX Manager.

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 flow.

  Workaround: None.

- **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 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 0xFFFFFFFFF**
  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 2075246: Moving a tier-1 logical router from one tier-0 logical router to another is not supported.**
  Moving a tier-1 logical router from one tier-0 logical router to another logical router causes the tier-1 logical router to lose downlink port route connection.

  Workaround: Complete the following steps:

  1. Detach the tier-1 logical router from the tier-0 logical router.
  2. Wait for about 20 minutes for the tier-1 logical router to completely detach from the tier-0 logical router.
  3. Attach the tier-1 logical router to another tier-0 logical router.
     The downlink port route connection is restored.
- **Issue 2077145: Attempting to forcefully delete the transport node in some cases might cause orphaned transport nodes**
  Attempting to forcefully delete the transport node using an API call where for example, there is a hardware failure and the hosts become irretrievable, changes the transport node state to Orphaned.

  Workaround: Delete the fabric node with the orphaned transport node.

- **Issue 2099530: Changing the bridge node VTEP IP address causes traffic outage**
  When the bridge node VTEP IP address is changed, the MAC table from VLAN to the overlay is not updated on the remote hypervisors causing traffic outage up to 10 minutes.

  Workaround: Initiate traffic changes from the VLAN so that the overlay MAC table on hypervisors is refreshed.

- **Issue 2106176: NSX Controller automatic installation stalls during the Waiting to Register step of the installation**
  During the automatic installation of NSX Controllers using either the NSX Manager API or UI, the status of one of the in-progress NSX Controllers stalls and shows as **Waiting to Register** indefinitely.

  Workaround: Complete the following steps:

  1. Send an API request to find the VM ID associated with the stalled NSX Controller.

     https://<nsx-mgr>/api/v1/cluster/nodes/deployments

  2. Send an API request to delete the stalled NSX Controller.

https://<nsx-mgr>/api/v1/cluster/nodes/deployments/<Controller id>?action=delete

- **Issue 2112459: Replacing a single node in the bridge cluster causes traffic drop**
  When you replace a single node in the bridge cluster, the bridged traffic flows to the old node which cause traffic drops until the forwarding entries in the remote hypervisors are updated or aged out.

  Workaround: Complete the following steps:

  1. Put the replacement node in the bridge cluster.
  2. Allow for HA to be established.
  3. Remove the old node.
- **Issue 216992: Using custom logical port MTU setting might result in packet drop**
  When using custom MTU setting on logical ports such as, logical router uplink port non-conforming values or certain configuration of the tier-0 and tier-1 logical routers can result in a packet drop. Default MTU setting is 1500.

  Workaround: Use the default MTU setting.

  Otherwise, the MTU applied on different logical ports must conform to the following relationship:

  1. Set the tier-0 logical router uplink MTU to 8900.
  2. Set the NSX Edge VTEP MTU to 9000.
  3. Set VM MTU to 8900.

  The tier-0 logical router and all the tier-1 logical routers connected to the tier-0 logical router must be collocated on the same NSX Edge nodes.

- **Issue 2125514: After layer 2 bridge failover, the logical switch on some NSX Edge VMs might do BUM replication of every single packet until the MAC is relearnt**
  After layer 2 bridge failover, the logical switch on some NSX Edge VMs might do BUM replication of every single packet for almost 10 minutes until the MAC is relearnt for the endpoint. The system recovers itself after the endpoints generate the next ARP.

  Workaround: None

- **Issue 2113769: DHCP relay not supported on the NSX Edge VLAN Layer 2 bridging**
  Connecting a VLAN host to the logical switch VNI through a Layer 2 bridging port on NSX Edge causes the DHCP relay agent on the logical router port to not provide an IP address to the VLAN host.

  Workaround: Complete the following steps:

  1. Configure the VLAN host manually.
  2. Move the Layer 2 bridging port to the ESXi host.
- **Issue 2183549: When editing a centralized service port, not able to view a newly created VLAN logical switch**
  In Manager UI, after you create a centralized service port and a new VLAN logical switch, if you edit the centralized service port, you cannot see the newly created VLAN logical switch.

Workaround: Use the API to edit the port.

- **Issue 2160634: Changing the IP address on a loopback can change the IP address of the router ID on an uplink**
  If the IP address on the loopback is changed, the NSX Edge selects the IP address on the uplink as the router ID. The IP address of the uplink which is assigned as the router ID cannot be changed.

  Impact to customer: 1. An expected side-effect of the Router-ID is that all the BGP sessions will flap.
  2. Real impact is the change of Router-ID, which can make the debugging BGP harder and can lead to confusion.

  Workaround: Disable the BGP configuration and change the IP address on the loopback.

- **Issue 2186040: If a transport node is not among the top 250 uplink profiles in the system, the physical NICs' uplink drop-down is disabled in the user interface**
  If a transport node is not among the top 250 uplink profiles in the system, the physical NICs' uplink drop-down is disabled in the user interface. Saving the Transport Node results in the removal of the uplink name from the transport node.

  Workaround: Reselect the uplink profile and the uplink name for that transport node.

- **Issues 2106635: During the static routes creation, changing the admin distance of the NULL routes causes the next-hop NULL setting to disappear from the user interface**
  During the static routes creation, when the you set the Next Hop to NULL and change the admin distance of the NULL routes, the next-hop NULL setting disappears from the user interface.

  Workaround: Reselect the next hop.

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: Complete the following steps:
  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
```

- **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 2108290: Bare Metal servers as transport nodes cannot assure NSX-T Data Center security features**
  Bare Metal servers as a new type of transport node does not offer the same level of security assurance such as, micro-segmentation as other hypervisor workloads. This is because, a reliable trust boundary is not enforced between the application workloads and the NSX agent.

  Workaround: For security reasons, do not assign tenant VMs the root privilege for Bare Metal servers or run applications as root. If tenant VMs have such access, a compromised tenant account or application might perform malicious activity on the Bare Metal server and introduce problems in the NSX-T Data Center network.

- **Issue 2162722: Popularity index is not applicable to DROP or REJECT rules and stateless rules**
  When traffic is hitting a rule with DROP/REJECT action or a stateless rule, the session count for the rule is not incremented as "session" is only applicable to a stateful ALLOW rule. Popularity Index is using the session count as a key parameter and thus it does not change for such rules.

  Workaround: None

- **Issue 2170512: CLI command to get firewall rules fails if an interface has more than 1,000 rules**
  If an interface has more than 1,000 rules, the CLI command `get firewall <VIF_ID> ruleset rules` will return an empty string.

  Workaround: There are 2 workarounds:

  - Run the command "nsxcli -c get firewall <VIF_ID> ruleset rules | json" instead.
  - Run the following raw CLI command. The name of a file containing the result will be displayed.
    ```
    ovs-appctl -t /var/run/vmware/nsx-agent/nsxa-ctl dfw/rules
    ```

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 Data Center 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-TData Center from the Network Manager.

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

- **Issue 2186045: On KVM, by default, logrotate runs on a daily basis instead of every minute**
  On KVM, if the size of a log file exceeds the size limit defined in its size-based rotation policy within one day, it will not get rotated until the end of that day when logrotate runs. Therefore, the log file sizes might be larger than the defined size limit.

  Workaround: Perform the following steps:

  1. Create a new directory /etc/cron.minutes.
  2. Create the /etc/cron.minutes/logrotate script with the following content:
     #!/bin/sh
     /usr/sbin/logrotate /etc/logrotate.conf
  3. Change the permission of /etc/cron.minutes/logrotate:
     chmod 755 /etc/cron.minutes/logrotate
  4. Append cron.minutes as an entry in /etc/crontab:
     echo "* * * * * root cd / && run-parts --report /etc/cron.minutes" >>/etc/crontab

Load Balancer Known Issues

- **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 2115545: When a load balancer health check is enabled, direct connectivity to the backend server pool members might fail**
  If a load balancer is attached to a logical router, then a client connected to the downlink of the logical router cannot access the pool members using the same protocol as the health check if the pool members are reachable using the uplink of the logical router.

  For example, if a load balancer is attached to logical router LR1 and has ICMP health check enabled for pool members reachable via LR1 uplink, then a client on the LR1 downlink cannot ping those pool members directly. However, the same client can use other protocols such as, SSH or HTTP to communicate with the server.

  Workaround: Use a different health check type on the load balancer. For example, to be able to ping the backend server, use the TCP or UDP health check instead of the ICMP health check.

- **Issue 2128560: Configuring both load balancer SNAT automap and health check might lead to occasional health check or connection failures**
  Configuring both the load balancer SNAT automap and health check such as, TCP, HTTP, HTTPS, or UDP for the same server pool might cause occasional health check or connection failures to that server pool.

  Workaround: Use SNAT IP list instead of SNAT automap.

  Note: SNAT IP addresses specified in the SNAT IP list mode should not include the logical router uplink IP address.

  For example, if a load balancer is attached to tier-1 logical router, LR1, then the configured SNAT IP range should not include LR1 uplink IP address.

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 Data Center relies on another user, nsx-user, to perform all NSX-T Data Center related tasks on the host.

Workaround: Do not use Lockdown mode. See Lockdown Mode in the vSphere documentation.

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 Data Center 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 Data Center 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.

- **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 2081979: Transport node host cannot connect to any controller**
  The NSX proxy log shows the following. A "certificate validation" message is expected but not present.

  TCP connection started: 10.171.0.73:0::3a4de8a2-3bc1-41ea-a94d-c1427d8cd757:1234
  Doing SSL handshake
  TCP connection established: 10.171.0.73:0::3a4de8a2-3bc1-41ea-a94d-c1427d8cd757, local addr: 10.171.0.59:36048, remote addr: 10.171.0.73

  Workaround: Log in to a controller as an administrator and run the following commands:

  set debug
  get mediator forcesync

Upgrade Known 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 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 2101728: Occasionally, the NSX Edge upgrade process is paused after a successful upgrade of an NSX Edge group**
  An NSX Edge group upgrade was successful however, during the second NSX Edge group upgrade the process is paused.

  Workaround: Click **Continue** to proceed with the NSX Edge group upgrade.

- **Issue 2106257: EULA API acceptance workflow changes for upgrade from NSX-T 2.1 to NSX-T 2.2**

EULA API acceptance should be called after updating the upgrade coordinator and prior to upgrading the existing hosts.

Workaround: None

- **Issue 2108649: Upgrade fails if there are files or directories open in the partition where upgrade is to occur**
  Avoid keeping files or directories open in the partition such as the NSX Manager or NSX Controller that are going to be upgraded, which causes the upgrade process to fail.

  Workaround: Reboot the appliance where the failure happened and restart the upgrade process.

- **Issue 2116020: After upgrade from NSX-T 2.1 to NSX-T 2.2, some Ubuntu KVM deprecated packages are not removed**
  After upgrade from NSX-T 2.1 to NSX-T 2.2, the following Ubuntu KVM deprecated packages are not removed.

  - nsx-host-node-status-reporter
  - nsx-lldp
  - nsx-logical-exporter
  - nsx-netcpa
  - nsx-support-bundle-client
  - nsx-transport-node-status-reporter
  - nsxa

  Workaround: Complete the following steps.

  1. Create a temporary file in the /etc/vmware/nsxa/ directory.
     cd /etc/vmware/nsxa
     touch temp.txt
  2. List all the nsxa package directory and files.
     dpkg -L nsxa
     /etc/vmware/nsxa# ls
  3. Remove the following packages.
     a) dpkg --purge nsx-lldp
     b) dpkg --purge nsx-support-bundle-client
     c) dpkg --purge nsx-transport-node-status-reporter
     d) dpkg --purge nsx-logical-exporter
     e) dpkg --purge nsx-netcpa
     f) dpkg --purge nsxa
     g) dpkg --purge nsx-host-node-status-reporter
  4. Verify that the following directory is available.
     /etc/vmware/nsxa/
  5. Remove the temp.txt file from the /etc/vmware/nsxa/ directory.
     rm -f temp.txt

- **Issue 2164930: Management plane upgrade finishes and shows a paused status when an empty host upgrade unit group is present**
  The overall Management plane upgrade status appears as paused and the host upgrade status is not marked as 100% when an empty host upgrade unit group is present.

Impact to customer: If the customer has empty host groups during upgrade then upgrade status is shown as PAUSED after MP upgrade is complete.

Workaround: Delete the empty host upgrade unit group before you upgrade the Management plane.
If the Management plane is upgraded, delete the empty host upgrade unit group and restart the install-upgrade service using CLI.

- **Issue 2097094: Cancelling an upgrade bundle upload in the midst of uploading is not supported**
  You cannot cancel the upload operation when the upgrade bundle .mub file is uploading.

  Workaround: Wait for the upgrade bundle .mub file to finish uploading.

- **Issue 2122242: Upgrading an Ubuntu KVM host from NSX-T 2.1 to 2.2 or NSX-T Data Center 2.3 does not remove the nsx-support-bundle-client package**
  When upgrading an Ubuntu KVM host from the NSX-T 2.1 release to a newer release (NSX-T 2.2 or NSX-T Data Center 2.3), the nsx-support-bundle-client package is still installed even though it is no longer used. Users can see that the package is still installed by invoking commands such as /usr/bin/dpkg -l.

  Workaround: Log in as root and run the following command to manually remove the package:

  ```
  # /usr/bin/dpkg --purge nsx-support-bundle-client
  ```

- **Issue 2186957: ESXi host not exiting from maintenance mode after upgrade**
  An ESXi host does not exit from maintenance mode after upgrade if the cluster has only one host and if the previous attempt by the upgrade coordinator to put it in maintenance mode failed.

  Workaround: Manually exit the host from maintenance mode or ensure that the host can enter maintenance mode (you must have at least 2 hosts per cluster).

- **Issue 2166207: During the upgrade from NSX-T Data Center 2.2 to NSX-T Data Center 2.3 with 500 hypervisors, the overall upgrade process might remain in the IN_PROGRESS state indefinitely**
  During the upgrade from NSX-T Data Center 2.2 to NSX-T Data Center 2.3 with 500 hypervisors, the overall upgrade process might remain in the IN_PROGRESS state indefinitely after clicking Pause followed by multiple web browser refresh.

  Workaround: Log in to the NSX-T Data Center CLI on the NSX Manager. Type the command, install-upgrade to restart the service.

- **Issue 2113681: If a KVM host becomes unreachable and fails after the NSX Edge upgrade, the Upgrade Coordinator attempts to upgrade the failed host instead of proceeding to upgrade the NSX Controller nodes**
  After you upgrade the KVM host and NSX Edge, and uninstall the new RPM and install an old RPM on the host, the host becomes unavailable in the Upgrade Coordinator. Therefore, the Upgrade Coordinator attempts to upgrade the KVM host instead of

proceeding to upgrade the NSX Controllers nodes.

Workaround: Refresh the Upgrade Coordinator user interface, click the **Hosts** tab, and attempt to upgrade the KVM host.

You can also skip the KVM host upgrade, open a command prompt and type the command, curl -i -k -u admin -X POST https://<nsx-manager-ip-address>/api/v1/upgrade/plan?action=continue\&skip=true

API Known Issues

- **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 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 1963850: The GET API displays items that are sorted in a case-sensitive manner**
  When a GET API returns items that are sorted by the display name, the sorting is case-sensitive.

  Workaround: None.

- **Issue 2070136: A distributed firewall API that processes a large amount of data fails**

A distributed firewall API that must create or update more than 100 MB of data fails with error code 500 and a message indicating a failed transaction. The API typically involves a section with more than 1000 rules, with each rule involving many sources, destinations, and applied-to objects.

Workaround: Create or update the rules incrementally.

- **Issue 1895497: The load balancer algorithm SRCDESTMACIPPORT in the API does not work**
  Calling an API to create a transport node's uplink profile with LAG that has source and destination MAC address, IP address and TCP/UDP port" will fail.

  Workaround: None

NSX Policy Manager Known Issues

- **Issue 2057616: During the NSX Policy Manager upgrade from NSX-T 2.1 to NSX-T 2.2, unsupported NSServices and NSGroups do not transfer**
  Unsupported NSService with Ether type and NSGroups with MAC Set and logical port membership criteria are not transferred during the NSX Policy Manager upgrade from NSX-T 2.1 to NSX-T 2.2.

  Workaround: Complete the following steps.

  1. In NSX-T 2.1, remove and modify NSServices with Ether type used in any communication entries.
  2. Remove and modify NSGroups with MAC Set and logical port membership criteria used in any communication entries.
  3. Upgrade the NSX Manager from NSX-T 2.1 to NSX-T 2.2.
  4. Upgrade the NSX Policy Manager using CLI.
- **Issue 2116117: NSX Policy Manager topology tab in the UI shows, Data connections failed**
  NSX Policy Manager topology tab in the UI shows, Data connections failed because groups in the policy domain contain VMs that are hosted on the ESXi 6.7 version, which is not supported.

  Workaround: None

- **Issue 2126647: Concurrent updates to the NSX Policy Manager distributed firewall causes override**
  When two users simultaneously edit the NSX Policy Manager distributed firewall section, the last user's change overrides the edits made by the other user that were made before.

  Workaround: Reinstate the distributed firewall changes made by the first user. After the changes are saved, the second user can make changes.

NSX Cloud Known Issues

- **Issue 2112947: While upgrading the NSX Agents in the Cloud Service Manager (CSM), some of the instances might appear as Failed**

While upgrading the NSX Agents in the CSM, some of the instances might appear as Failed because of an unresponsive user interface.

Workaround: Refresh the user interface.

- **Issue 2111262: While deploying PCG, you might see the error: "Gateway deployment failed: [Errorcode: 60609] Async operation failed with provisioning state: Failed." Or "Failed to create gateway virtual machine with name nsx-gw, Gateway deployment failed."**
  This is a rare event and occurs due to the Microsoft Azure infrastructure.

  Workaround: Redeploy the failed Public Cloud Gateway (PCG).

- **Issue 2110728: If you are using HA, but installed the NSX agent on VMs by specifying only one PCG's DNS name using the --gateway option, failover to the secondary PCG does not work.**
  Workload VMs are not able to connect to the PCG after failover and therefore the PCG won't be able to enforce/realize any logical state on the VM.

  Workaround: Don't use --gateway option at all when installing agents on the workload VMs. Use the value from the Gateway screen of the VPC or VNet. See **Installing NSX Agent** in the NSX-T Data Center Administration Guide for details.

- **Issue 2071374: Harmless error messages regarding "nscd" may appear when installing NSX Agent on certain Linux VM instances**
  Description: On VMs that have "nscd" running, you may see error messages like: "sent invalidate(passwd) request, exiting" while installing NSX agent. This occurs on VMs running, for example, Ubuntu 14.04 or 16.04

  Workaround: The messages appear because of a known bug with the Linux distribution. These messages are harmless and do not affect NSX agent installation.

- **Issue 2010739: Both Public Cloud Gateways (PCGs) shown as standby**
  If the primary PCG is not able to connect to controller during gateway on-boarding, then both gateways -- primary and secondary -- will be in standby mode until the connection between controller and gateway is restored.

- **Issue 2121686: CSM displays the exception "Server failed to authenticate the request."**
  You may see this error in CSM and the reason is that the CSM appliance time is not in sync with Microsoft Azure Storage Server or NTP. In this case, Microsoft Azure throws the exception "Server failed to authenticate the request." which is ambiguous but the same error is displayed in CSM.

  Workaround: Sync CSM appliance time with NTP or Microsoft Azure Storage server time.

- **Issue 2092378: Deploying PCG in HA mode shows both PCGs in standby mode, and Cloud Sync shows the primary PCG as active**
  After HA deployment of PCG through CSM in a private network, standby/standby or active/active status is seen on the deployed PCGs for upto 1 hour. During this interval, it seems to the user that there is some problem with PCG deployed and may be an unclear

state to proceed.

Workaround: Do the following:

1. Resync the account from UI after PCG deployment through which CSM can fetch the latest data and display in CSM.
2. If after resync, CSM still shows PCGs in wrong state, check the connectivity status of PCG in NSX Manager.
3. If connection shows as UP, and still the states are incorrect, proceed with debugging PCG.

- **Issue 2119726: While deploying PCG in a Microsoft Azure VNet, public IPs which were previously associated with VMs may get erroneously listed as free to use.**
  If the VMs to which public IPs were assigned previously are now powered off, they no longer have those public IPs associated with them. This is because Microsoft Azure dissociates public IPs associated with VMs after they are powered off for a certain period of time. This time period is not specifically defined by Microsoft Azure.

  Workaround: Do not power off PCGs in your VNet. This prevents the public IP dissociation with the uplink interface of the primary PCG. If you must power off the PCGs, then ensure that the PIP associated with the PCGs are not reused and as soon as PCGs are powered back on, they get the same PIP.

- **Issue 2165915: NSX Cloud support for Red Hat Enterprise Linux 7.4 with kmod.x86_64 0:20-15.el7_4.6**
  NSX Cloud does not support VM instances which run Red Hat Enterprise Linux 7.4 with kmod-20-15.el7_4.6. This is caused by a bug reported by Red Hat: https://bugzilla.redhat.com/show_bug.cgi?id=1522994.

  Workaround: Update to the kmod version where this bug is fixed, for NSX agent installation to succeed.

- **Issue 2102828: In Microsoft Azure deployments, during and after being upgraded from NSX-T 2.2 to NSX-T Data Center 2.3, the Public Cloud Gateway (PCG) may appear to be non-functional.**
  In Microsoft Azure deployments where the system has been upgraded from NSX-T 2.2 to NSX-T Data Center 2.3, in rare cases, the Public Cloud Gateway (PCG) may fail to obtain IP addresses on its interfaces. This may be experienced during an upgrade of the PCG step where the upgrade process of the PCG appears to hang. This issue may also manifest itself as a non-operational PCG if the administrator restarts the PCG appliance from the Microsoft Azure portal. This issue does not apply to new systems installing NSX-T Data Center 2.3 for the first time.

  Workaround: From the Microsoft Azure portal, restart the PCG you are upgrading, then in the Cloud Service Manager (CSM), verify that the status of the PCGs and the VM instances is valid.

- **Issue 2180531: NSX Agent is supported for Ubuntu 16.04 VM instances that have kernel 4.14 and lower**
  NSX Agent is supported for Ubuntu 16.04 VM instances that have kernel 4.14 and lower. NSX Agent will **not** work for Ubuntu 16.04 VM instance with kernel 4.15 and higher.

No workaround exists for this issue

- **Issue 2170445: After upgrading PCG from NSX-T Data Center 2.2. to NSX-T Data Center 2.3, PCG HA state will not be correctly set for Microsoft Azure PCGs**
  After upgrading Microsoft Azure PCGs from NSX-T 2.2 to NSX-T Data Center 2.3, the HA state of the PCGs does not become Active-Standby as expected. The preferred PCG HA state displays as SYNC and the non-preferred PCG HA state shows up as Active. Because of this, in the event of an HA failover after the upgrade, only one of the PCGs has a valid state.

  Workaround: In NSX-T 2.2, update the MTU in PCG's uplink host switch profile to 1500 before starting the upgrades to NSX-T Data Center 2.3.
  This can be either done using the NSX Manager UI or NSX Manager REST APIs.

  Through the UI, do the following:

  1. Go to **Fabric > Profiles**
  2. Select the profile with name "PCG-Uplink-HostSwitch-Profile" and description "PublicCloudGateway Uplink HostSwitch Profile"
  3. Click **EDIT** and modify **MTU** value to 1500 and click **SAVE**
  4. Start the upgrade from NSX-T 2.2 to NSX-T Data Center 2.3.

  Through the REST API, do the following:

  1. GET all Host Switch Profiles using:

```
curl -X GET \
  https://<NSX-Manager-URL>/api/v1/host-switch-profiles \
  -H 'authorization: Basic <AUTH ID>' \
  -H 'content-type: application/json'
```

  2. Identify the host switch profile with name "PCG-Uplink-HostSwitch-Profile" and description "PublicCloudGateway Uplink HostSwitch Profile" and get the ID of that profile:

```
curl -X PUT \
  https://<NSX-Manager-URL>/api/v1/host-switch-profiles/<host-switch-profile-id> \
  -H 'authorization: Basic <AUTH ID>' \
  -H 'content-type: application/json' \
  -d '{
          "resource_type": "UplinkHostSwitchProfile",
          "description": "PublicCloudGateway Uplink HostSwitch Profile",
          "id": "<host-switch-profile-id>",
          "display_name": "PCG-Uplink-HostSwitch-Profile",
          "tags": [
             {
                "scope": "CrossCloud",
                "tag": "public-cloud-manager"
             },
             {
                "scope": "PcmId",
                "tag": "<Existing PCM ID>"
```

```
            },
            {
                "scope": "EntityType",
                "tag": "default"
            },
            {
                "scope": "CloudScope",
                "tag": "<Existing VPC/VNET name>"
            },
            {
                "scope": "CloudType",
                "tag": "<Existing cloud type>"
            },
            {
                "scope": "CloudVpcId",
                "tag": "<Existing Vpc/Vnet id>"
            }
        ],
        "transport_vlan": 0,
        "teaming": {
            "active_list": [
                {
                    "uplink_type": "PNIC",
                    "uplink_name": "uplink-1"
                }
            ],
            "policy": "FAILOVER_ORDER"
        },
        "overlay_encap": "GENEVE",
        "mtu": 1500,
        "_revision": 1
    }'
```

- **Issue 2174725: Managed VPC/VNet with PCGs deployed is shown as unmanaged in CSM.**
  Managed AWS VPC or Microsoft Azure VNet with PCGs deployed is shown as unmanaged in CSM.

  Workaround: Restarting CSM should resolve the issue.

- **Issue 2162856: Azure PCGs have an invalid HA state (both active or both standby)**
  When you deploy a pair of PCGs in AWS and then deploy another pair of PCGs for Azure, the Azure PCGs will have an invalid HA state (both active or both standby).

  Workaround: Update the MTU in PCG's uplink host switch profile created by PCM to 1500 before starting the Cross cloud upgrade to NSX-T Data Center 2.3. From the manager UI, perform the following steps:

  - Go to Fabric > Profiles.
  - Select the profile with the name "PCG-Uplink-HostSwitch-Profile" and description "PublicCloudGateway Uplink HostSwitch Profile".
  - Click "EDIT", modify the "MTU" value to 1500 and click "SAVE".
  - Start the upgrade workflow.
- **Issue 2102321: Some NSX Cloud operations may be slow on Microsoft Azure during**

high-traffic periods.
NSX Cloud relies on Microsoft Azure ARM API for certain operations like managing VMs or withdrawing them from NSX-management; or taking quarantine actions on a VM. During peak periods, Microsoft Azure may hit API limits for given subscriptions, in which case, it will start throttling all API requests for that subscription. During this time, the above mentioned NSX operations may not complete in time. These operations will eventually be completed when Microsoft Azure stops throttling the requests. PCM logs on the Public Cloud Gateway will have logs like the following indicating that throttling is currently happening "*Azure Resource Manager read/write per hour limit reached.Will retry in: x seconds"*

WORKAROUND: Wait until Microsoft Azure throttling stops.

- **Issue 2189738: AWS workload VMs cannot be reached after Quarantine Policy is disabled for an onboarded VPC, when it was previously enabled.**
  If a PCG is deployed with Quarantine Policy enabled, and later if you disable quarantine mode, some NSX-managed AWS workload VMs in this VPC fail to communicate with PCG.

  Workaround: Add the following inbound rules to the NSX Cloud Security Group in the AWS VPC: **gw-mgmt-sg**:
  Note: Remove these rules when you enable Quarantine Policy again for security reasons.

TYPE	Protocol	Port	Source
CUSTOM-TCP	TCP	8080	VPC-CIDR
CUSTOM-TCP	TCP	5555	VPC-CIDR

- **Issue 2188950: You see the error: "No VNet found for specified ID." when using the API to retrieve a list of PCGs.**
  You see this error if an account associated with PCGs deployed is deleted from CSM.

  Workaround: Add the Microsoft Azure account in CSM upon which the PCGs were deployed.

- **Issue 2191571: PCG deployment does not start if the SSH public key for PCG deployment does not end with an email ID.**
  SSH public key must end with an email ID or PCG deployment will not start and display an error.

  Workaround: Ensure the SSH key ends with an email ID.

- **Issue 2092073: On Windows workload VMs, IPFIX templates are not received correctly.**
  On Windows workload VMs, logical switch and firewall IPFIX templates are not sent out immediately when the IPFIX collector is configured in the same subnet as the VM. This is because Windows Socket expects an ARP entry for the IPFIX collector's IP address before sending out the UDP packet. If an ARP entry is missing, it silently drops all the UDP packets except the last one. As a result, on the IPFIX collector, the data packet is received

with no template information.

Workaround: Do one of the following:

- Add a static ARP entry for the IPFIX collector using the command:

netsh interface ipv4 add neighbors “<Interface name>” <collector IP> <physical address of collector>

For example:

netsh interface ipv4 add neighbors “Ethernet 3" 172.26.15.7 12-34-56-78-9a-bc

- Configure the IPFIX collector on a different subnet than the workload VMs.

- **Issue 2210490: If you add a proxy profile in CSM, the password will be visible to all CSM API users with any of these roles assigned to them: Cloud Service Auditor or Cloud Service Administrator.**
  If you create a proxy profile in CSM and provide a username and password, even though you cannot view the password in the CSM UI, it will be visible in response to the following APIs:

  - /csm/proxy-server-profiles
  - /csm/proxy-server-profiles/<profile-id>
- **Issue 2039804: PCG deployment fails but the PCG instance is not terminated in AWS.**
  If you are deploying PCG but the deployment fails, you will still see the PCG instance(s) in your AWS VPC and auto-created logical entities in NSX Manager.

  Workaround: Manually delete the auto-created NSX Manager entities and terminate the PCG instance in your AWS VPC.

NSX Container Plug-in (NCP) Known Issues

- **PAS 2.1.0 CNI change**
  Due to the CNI plugin change in PAS 2.1.0, no NSX-T Tile, regardless of version, will work with PAS 2.1.0. This is fixed in PAS 2.1.1.

- **Issue 2118515: In a large-scale setup, NCP takes a long time to create firewalls on NSX-T**
  In a large-scale setup (for example, 250 Kubernetes nodes, 5000 pods, 2500 network policies), it can take NCP a few minutes to create the firewall sections and rules in NSX-T.

  Workaround: None. After the firewall sections and rules are created, performance should be back to normal.

- **Issue 2125755: A StatefullSet could lose network connectivity when performing canary updates and phased rolling updates**
  If a StatefulSet was created before NCP was upgraded to the current release, the StatefullSet could lose network connectivity when performing canary updates and phased rolling updates.

  Workaround: Create the StatefulSet after NCP is upgraded to the current release.

- **Issue 2131494: NGINX Kubernetes Ingress still works after changing the Ingress class from nginx to nsx**
  When you create an NGINX Kubernetes Ingress, NGINX create traffic forwarding rules. If you change the Ingress class to any other value, NGINX does not delete the rules and continues to apply them, even if you delete the Kubernetes Ingress after changing the class. This is a limitation of NGINX.

  Workaround: To delete the rules created by NGINX, delete the Kubernetes Ingress when the class value is nginx. Than re-create the Kubernetes Ingress.

- **Issue 2194845: The PAS Cloud Foundry V3 API feature "multiple processes per app" is not supported**
  When using the PAS Cloud Foundry V3 API `v3-push` to push an app with multiple processes, NCP does not create logical switch ports for the processes except the default one. This issue exists in NCP 2.3.0 and earlier releases.

  Workaround: None

- **Issue 2193901: Multiple PodSelectors or multiple NsSelectors for a single Kubernetes network policy rule is not supported**
  Applying multiple selectors allows only incoming traffic from specific pods.

  Workaround: Use matchLabels with matchExpressions in a single PodSelector or NsSelector instead.

- **Issue 2194646: Updating network policies when NCP is down is not supported**
  If you update a network policy when NCP is down, the destination IPset for the network policy will be incorrect when NCP comes back up.

  Workaround: Recreate the network policy when NCP is up.

- **Issue 2192489: After disabling 'BOSH DNS server' in PAS director config, the Bosh DNS server (169.254.0.2) still appears in the container's resolve.conf file.**
  In a PAS environment running PAS 2.2, after you disable 'BOSH DNS server' in PAS director config, the Bosh DNS server (169.254.0.2) still appears in the container's resove.conf file. This causes a ping command with a fully qualified domain name to take a long time. This issue does not exist with PAS 2.1.

  Workaround: None. This is a PAS issue.

- **Issue 2194367: NSX-T Tile does not work with PAS Isolation Segments which deploy their own Routers**
  NSX-T Tile does not work with Pivotal Application Service (PAS) Isolation Segments which deploy their own GoRouters and TCP Routers. This is because NCP cannot get the IP addresses of the Router VMs and create NSX firewall rules to allow traffic from the Routers to the PAS App containers.

  Workaround: None.

- **Issue 2199504: The display name of NSX-T resources created by NCP is limited to 80 characters**

When NCP creats an NSX-T resouce for a resource in the container environment, it generates the display name of the NSX-T resource by combining the cluster name, namespace or project name, and the name of the resource in the container environment. If the display name is longer than 80 characters, it is truncated to 80 characters.

Workaround: None

- **Issue 2199778: With NSX-T 2.2, Ingress, Service and Secrets with names longer than 65 characters are not supported**
  With NSX-T 2.2, when `use_native_loadbalancer` is set to `True`, the names of Ingresses, Secrets and Services referenced by the Ingress, and Services of type LoadBalancer, must be 65 characters or less. Otherwise, the Ingress or Service will not work properly.

  Workaround: When configuring an Ingress, Secret, or Service, specify a name that is 65 characters or less.

- **Issue 2065750: Installing the NSX-T CNI package fails with a file conflict**
  In a RHEL environment with kubernetes installed, if you install the NSX-T CNI Package using `yum localinstall` or `rpm -i`, you get an error indicating a conflict with a file from the kubernetes-cni package.

  Workaround: Install the NSX-T CNI package with the command`rpm -i --replacefiles nsx-cni-2.3.0.xxxxxxxx-1.x86_64.rpm`.

- **For a Kubernetes service of type ClusterIP, Client-IP based session affinity is not supported**
  NCP does not support Client-IP based session affinity for a Kubernetes service of type ClusterIP.

  Workaround: None

- **For a Kubernetes service of type ClusterIP, the hairpin-mode flag is not supported**
  NCP does not support the hairpin-mode flag for a Kubernetes service of type ClusterIP.

  Workaround: None

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.