VMware NSX-T 2.2 Release Notes

VMware NSX-T 2.2 | 05 JUN 2018 | Build 8680772

Check regularly for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

- What's New
- Compatibility and System Requirements
- API Reference Information
- Resolved Issues
- Known Issues

What's New

What's New
NSX-T 2.2 is the incremental upgrade release that enhances the new multi-hypervisor platform delivered for cloud and containers.

The following new features, feature enhancements, and deprecated feature are available in the NSX-T 2.2 release.

New NSX-T Features

Automated NSX Controller Cluster Deployment

Automatically deploy NSX Controller cluster from the NSX Manager onto vSphere clusters discovered from a vCenter Server to simplify NSX-T installation in a vSphere environment.

NSX Management of Workloads in Azure

- NSX Datacenter and NSX Cloud with a single pane of glass for on-premise and Azure workloads.
- Single security policy across hybrid cloud, offering various set of attributes, including VM names, custom tags.
- Decouple workload deployment from security enforcement.

NIOC Version 3 Support

Network IO control (NIOC) allows configurable limits and shares on the network for both system-generated and user-defined network resource pools, based on the capacity of the physical adapters on an ESXi host.

- Provides a mechanism to reserve bandwidth for system traffic based on the capacity of the physical adapters on a host.
- Enables fine-grained resource control at the VM network adapter level to use for allocating CPU and memory resources.
- Allows to set up bandwidth allocation for virtual machines at the level of the entire distributed switch (N-VDS).

Guest VLAN Tagging

vSwitch, N-VDS in this case, port acts like a trunk and inspects incoming VLAN tags to ensure that they match the correct destination virtual port. However, the VLAN tags are left intact by the N-VDS.

Feature is applicable for both VLAN backed and overlay backed traffic and supports bridging only for forwarding packets based on guest VLAN tag. Routing based on guest VLAN tag within the hypervisor is not supported.

VPN Support

IPsec-based Layer 2 VPN and Layer 3 VPN support for site-to-site connectivity that can be configured using only NSX-T APIs.

Customer Experience Improvement Program

VMware Customer Experience Improvement Program (CEIP) support collects product usage information and reports to VMware to improve the quality of NSX-T. Customers can optionally disable this feature.

Terraform Support

Terraform provider support to automate NSX-T logical objects such as switches, routers, firewall rules, and grouping.

Cisco VIC 1387 Support on Bare Metal for NSX Edge

Support for NICs used in the Cisco UCS systems.

NSX Container Plug-in (NCP) Features

- TLS support for Kubernetes Ingress
  Support for HTTPS termination at an NSX-T load balancer with Kubernetes Ingress and Secret integration. All Kubernetes Ingresses with a TLS section will be hosted on an NSX-T load balancer dedicated to HTTPS termination on port 443.
- OpenShift Router support
  The NSX-T load balancer can work as the OpenShift Router to expose services to external layer 7 traffic via Route resources. Both HTTP traffic and HTTPS traffic with edge termination are supported.
- Support for longer names and values when creating tags

Tags on NSX-T objects now allow up to 30 characters for tag scopes and 65 characters for tag values.

- OpenShift installation improvements

NSX-T Enhancements

N-VDS Enhanced Datapath Mode

Support for a high-performance mode called enhanced datapath when used with vSphere 6.7. This mode implements some key DPDK features such as, Poll Mode Driver, Flow cache, optimized packet copy, and provides better performance for small and large packet sizes pertinent for Network Functions Virtualization (NFV) style workloads. Telecommunication operators can now have a highly performing virtual switch without sacrificing any benefits of virtualization such as, vMotion or Predictive DRS.

Note: Not all the features of N-VDS are available when operating in the enhanced datapath mode.

Monitoring and Troubleshooting Enhancements

- Traceflow API enhancements to troubleshoot IP address assignment using NSX-T DHCP service.
- New port mirroring type called, Logical SPAN to monitor source ports to a destination port on the same logical overlay switch.
- Enhanced IPFIX profiles to be applied to NSGroups.

VLAN Based Logical Switch Teaming Policy Support for ESXi Hosts

Enables association or pinning of logical switch traffic to a specific uplink. Configurable using teaming policy of Route based on the originating virtual port.

NSX Edge Firewall Interface

Layer 4 stateful firewall on a per uplink basis on the tier-1 or tier-0 logical routers to selectively filter traffic coming from various uplinks.

Distributed and NSX Edge Firewall Enhancement

Centralized place to view the status of a firewall. Use APIs to query the status of a firewall publish operation or retrieve information on whether a rule has been deployed on a particular VM.

Principle Identity Role Support

Configure principle identities with one of the default NSX-T roles.

Search Enhancement

Support for search auto-complete.

Backup Enhancements

Provides option to trust certificate thumbprints presented by the system where remote backup or restore archives are stored.

Support VLAN Backed Downlinks

Connecting VLAN-backed downlink to tier-0 or tier-1 logical router leverages centralized router port that is available only on the NSX Edge node.

Load Balancing Enhancements

- Load Balancer HTTPS Support
  - Support for HTTPS traffic with SSL termination on the load balancer.
  - SSL-Offload load balancing support for HTTPS from client to load balancer decrypted and HTTP from load balancer to server.
  - SSL End-To-End support for HTTPS from client to load balancer re-encrypted in new HTTPs from load balancer to server.
- Load balancer virtual server IP displays real-time graphics for Concurrent Connections, New Connection Rate, Throughput, and HTTP Request rate.
- Access Log granularity allows log setting for specific virtual server instead of a load balancer.
- Single API to download the entire load balancer configuration.
- WebSocket application support with enhanced HTTP protocol.
- Sorry server with the ability to define per virtual server a second server pool of sorry server to use in case all the members of the first server pool are down.
- New load balancer rule, Match cookie value and match value case insensitive.
- Layer 4 multiple port range support.
- New load balancer rule algorithm with Weighted Least Connection.
- Slow start enabled automatically for the load balancer algorithm Least Connection and Weighted Least Connection to prevent a new server added to an existing production server pool to be inundated by new connections.
- POST API request request_body_size can now be limited in size.

NSX Edge Layer 2 Bridge Enhancements

VLAN to overlay service hosted on the NSX Edge node for improved performance than ESXi-based Layer 2 bridge and Layer 3 firewall.

API Rate Limiting

Limit the number of transactions per second and concurrent transactions to the NSX-T REST API. This protects the system from being impacted when one or more API clients make API requests at a rate the API cannot process.

Deprecated NSX-T

Feature

Distributed Network Encryption feature is deprecated.

Compatibility and System Requirements

For compatibility and system requirement information, see theNSX-T Installation Guide.

NCP Compatibility Requirements:

Product	Version
NCP / NSX-T Tile for PAS	2.2.0
NSX-T	2.1, 2.2
Kubernetes	1.9, 1.10
OpenShift	3.7, 3.9
Kubernetes/OpenShift host VM OS	Ubuntu Xenial, RHEL 7.4, RHEL 7.5
PAS (PCF)	OpsManager 2.1.x + PAS 2.1.x (except PAS 2.1.0) OpsManager 2.2.0 + PAS 2.2.0

General Behavior Changes

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, both the TCP ports 1234 and 1235 must be open. After the upgrade is complete, the TCP 1235 port is in use.

API Reference Information

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

The latest API reference is located in the NSX-T 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 (Note: NCP 2.2 supports NSX-T 2.1. The following issues are resolved only if you run NSX-T 2.2.)

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.

  Workaround: Add X-XSRF-TOKEN header.

- **Issue 1989412: Domain deletion when NSX Manager is not reachable is not reflected when connectivity is restored**
  If a domain is deleted from Policy when the NSX manager is not reachable, after the connection is restored to the NSX Manager, the firewall and corresponding rules to the deleted domain still exist.

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

- **Issue 2018478: Attempting to remove a widget from the dashboard causes a crash with stack trace error**
  Custom dashboard user interface changes such as, removing a widget from multiple widgets causes the user interface to crash with a stack trace error.

  Workaround: Complete the following steps.

  1. Create a widget.
  2. Locate the multiple widgets you want to modify.
  3. Add a reference to the newly created widget in the multiple widget.
- **Issue 1959647: Using a database server alias name to create a DSN might cause the installation of vCenter Server to fail**
  When you use a database server alias name to create a DSN, the installation of vCenter Server with an external Microsoft SQL database fails. The following error appears during the installation of the inventory service: An error occurred while starting invsvc.

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

- **Issue 1998217: HyperBus interface vmk50 might be missing on vSphere ESXi causing container creation failure**
  Container is not created because the HyperBus interface vmk50 might be missing on vSphere ESXi.

  Workaround: Complete the following steps.

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

Installation Resolved Issues

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

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

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

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

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

  NSX-T Unified appliance

  Small - 2 CPU, 8GB memory virt-install ... --ram 8192 --vcpus 2
  Medium - 4 CPU, 16 GB memory virt-install ... --ram 16384 --vcpus 4
  Large - 8 CPU, 32 GB memory virt-install ... --ram 32768 --vcpus 8

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

  NSX Edge
  Small - 2 CPU, 4G memory virt-install ... --ram 4096 --vcpus 2
  Medium - 4 CPU, 8G memory virt-install ... --ram 8192 --vcpus 4
  Large - 8CPU, 16G memory virt-install ... --ram 16384 --vcpus 8

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

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

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

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

NSX Manager Resolved Issues

- **Issue 1932987: After restoring the Management Plane, the connection between Management Plane and Key Manager Server fails**
  When you detach the Key Manager Server from the Management Plane, restore the Management Plane, and attempt to reattach the Key Manager Server, the connection fails.

  Workaround: None.

- **Issue 1954293: vMotion of VMs connected to logical switches fails during Management Plane upgrade**
  While Management Plane is being upgraded, if you attempt to vMotion a VM connected to a logical switch, the vMotion fails.

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

- **Issue 1954297: If NSX Manager's restore is done, and any new non-VC managed ESX host is registered with NSX Manager and its VMs are connected to existing Logical Switches, then on ESX hosts' MOB, MAC address for VM becomes blank**
  If NSX Manager restore is done, and any new non-VC managed ESX host is registered with Management Plane and its VMs are connected to existing Logical Switches, then on ESX hosts' MOB, MAC address for VM becomes blank.
  This does not have any effect on VM's Inventory with respect to MAC on NSX Manager.

  Workaround: None.

- **Issue 1978104: Some pages in the NSX Manager user interface are not accessible on Internet Explorer 11**
  The Dashboard, Getting Started workflows, and load balancer pages in the NSX Manager user interface are not accessible on the Windows machine that is running Internet Explorer 11.

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

- **Issue 1954986: The license key is shown in the logs when the key is deleted from the UI**
  The NSX license key is shown in /var/log/syslog as follows:

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

```
status:'success', New value: [{"atomic":false} {"request":
[{"method":"DELETE","uri":"/v1/licenses/<license_key>"}]}]
```

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

Workaround: None.

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

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

- **Issue 1957165: Loading the last page in a search result set that includes 10,040 or more records yields a Search exception**
  In a large environment that could return 10,040 or more possible objects for a search query, you might see an exception when trying to load the last few records in the result set from the UI listing.

  Workaround: Narrow the search criteria.

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.

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

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

- **Issue 2015445: Firewall state on the active service router might not be duplicated on the newly active service router**
  Tenant logical router (TLR) might have multiple failovers from NSX Edge1 to NSX Edge2 and from NSX Edge2 to NSX Edge1. Firewall or NAT flow states are synchronized between active/standby TLR service routers. When the TLR is configured in a non-

preemptive failover mode, the synchronization occurs before the first failover, but does not occur between first and the subsequent failover. As a result, at the second failover, the TCP traffic can time out. This problem does not occur with TLR configured in preemptive mode.

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

- **Issue 2016629: RSPAN_SRC mirror session fails after migration**
  When a VM connected to a port assigned for RSPAN_SRC mirror session is migrated to another hypervisor, and there is no required pNic on the destination network of the destination hypervisor, then the RSPAN_SRC mirror session fails to configure on the port. This failure causes the port connection failure but the vMotion migration process succeeds.

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

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

  The mirror session fails to configure, but the port connection is restored.
- **Issue 1620144: NSX-T CLI, `get logical-switches` lists logical switches with status UP, even after the transport node is deleted**
  The NSX-T CLI might mislead the user that there is a functional logical switch. Even when logical switches are seen, they are not functional. The opaque switch is disabled when the transport node is deleted, thus no traffic gets through.

  Workaround: None.
- **Issue 1590888: Warning needed that logical ports selected in Ethernet section apply only within same L2 network**
  For the NSX-T distributed firewall, in the Ethernet section, when any logical port or MAC address is entered in the source/destination section, a warning should be displayed that MAC addresses or logical ports should belong to VM ports in same L2 network (attached to same Logical switch). Currently, there is no warning message.

  Workaround: None.
- **Issue 1763576: Hypervisors are allowed to be removed as transport nodes even when they have VMs on the NSX-T network**
  NSX-T does not prevent you from deleting a transport node even when there are VMs on the node that are part of the NSX-T network. The VMs lose connectivity after the transport node is deleted.

  Workaround: For both ESXi and KVM hosts, recreate the transport node again with the same host switch name.
- **Issue 1780798: In a large-scale environment, some hosts might get into a failed state**
  In a large-scale environment with 200 or more host nodes after running for some time, some hosts might lose connectivity with NSX Manager and the log contains error messages such as:

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

Workaround: Restart the MPA process on the failed hosts.

- **Issue 1954997: Transport Node deletion fails if VMs on the transport node are connected to Logical Switch at the time of deletion**
  1. Fabric Node and Transport Node are created.
  2. Attach VIFs to logical switch.
  3. Delete transport node without removing VIF attachments to Logical Switch fails.

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

- **Issue 1958041: BUM traffic might not work for Layer 3 flow across physical Layer 2 segments when ESX hypervisor has multiple uplinks**
  If all of the following conditions are met, it is possible that BUM traffic from source hypervisor across logical router does not reach the destination hypervisor.

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

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

  Workaround:

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

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.

  Workaround: None.

- **Issue 2008882: For Application Discovery to work properly, do not create a security group that spans multiple hosts**
  If one security group has VMs that span across multiple hosts, the Application Discovery session might fail.

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

Load Balancer Resolved Issues

- **Issue 195228: Weighted round-robin and weighted least connection algorithms might not distribute traffic properly after a configuration is changed and reloaded**
  Servers lose connection when a weighted round-robin or weighted least connection configuration is changed and reloaded. After the connectivity loss, the historical traffic distribution information is not preserved which leads to traffic being distributed improperly.

  Workaround: None.

- **Issue 2018629: Health check table not showing the updated monitor type for the NS group pool**
  When you create static and dynamic NS group pools with the same members with a monitor type and change that monitor type on dynamic pool, the dynamic pool health check does not appear in the health check table.

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

- **Issue 2020372: Passive health check does not consider the pool member down after the maximum fall count is reached**
  Passive health check requires additional fall count value than configured to consider the pool member down.

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. So the dashboards are executing old queries which contain fields such as *vmw_nsxt_comp*, *vmw_nsxt_subcomp*, and *vmw_nsxt_errorcode*, and these fields are named differently in the newer version of the query script. As a result, the queries will return empty results and the dashboards will be blank.

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

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.
  ```

  Workaround: Run the `initialize control-cluster` command again.

Upgrade Resolved Issues

- **Issue 1847884: Do not make NSX-T related changes until the upgrade process for the Management Plane has completed**
  Performing any changes such as, creating, updating, or deleting a transport zone, transport node, or logical switches during the Management Plane upgrade might corrupt the Management Plane, leading to NSX Edge, host, and data path connectivity failures.

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

- **Issue 2005709: Upgrade coordinator page becomes inaccessible when you use the NSX Manager FQDN**
  When you use the NSX Manager FQDN to open the NSX Manager user interface, the following error message appears in the Upgrade Coordinator page, This page is only available on the NSX Manager where Upgrade Coordinator is running. To enable the service, run the command "set service install-upgrade enabled" on the NSX Manager. If the install-upgrade service is already enabled, try disabling it using "clear service install-upgrade enabled" and then enable it again."

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

- **Issue 2022609: Managed hosts are treated as unmanaged host in the upgrade coordinator**
  If an environment has more than 128 managed hosts, during the upgrade process the hosts that were part of a cluster appear in the unmanaged ESXi group.

  Workaround:

  - Manually upgrade the unmanaged ESXi hosts above 128.
  - Navigate to the /opt/vmware/upgrade-coordinator-tc-server/webapps/upgrade-coordinator/WEB_INF/classes/config.properties file, change the value of upgrade.host.service.hostNodeListPageSize from 128 to 512, and restart the upgrade coordinator /etc/init.d/upgrade-coordinator restart.
- **Issue 1944731: DHCP leases might have conflicting records if numerous requests are served by the first upgraded NSX Edge during the upgrade of the second NSX Edge**
  If numerous requests are served by first upgraded NSX Edge during the upgrade of the second NSX Edge, then the DHCP leases might have conflict records.

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

API Resolved Issues

- **Issue 1619450: Test vertical is returned by polling frequency configuration APIGET /api/v1/hpm/features**
  GET /api/v1/hpm/features returns the list of all features for which polling frequency can be configured. This API returns some internal, test-only features. There is no functional impact on the user other than extra noise.

  Workaround: Ignore the extraneous API response.

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

  Workaround: None

- **Issue 1954990: Realization API inaccurate status return**
  If you use a Realization API to check the realization status for all APIs executed before a barrier, the return status by the Realization API can be misleading relative to the actual status. Because of the complexity of the execution of the DFW inside the Management Plane, DFW API can slip after the barrier they are supposed to follow which leads to this inaccuracy.

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

NSX Container Plug-in (NCP) Resolved Issues (Note: NCP 2.2 supports NSX-T 2.1. The following issues are resolved only if you run NSX-T 2.2.)

- **Issue 2083074: NSX Manager shows incorrect auto-discovered address bindings for a container logical port.**
  If ARP snooping is enabled (through the IP discobery switching profile) for a container logical port and the container host VM's logical port, the NSX Manager's GUI might show incorrect auto-discovered address bindings for the container logical port. This is a cosmetic error and can be ignored. The actual address bindings are displayed under Manual Bindings.

  Workaround: None needed.

- **Issue 2091962: Pivotal Bosh provisioning of virtual machines fails randomly on NSX-T**
  When Pivotal BOSH deploys new VMs on vSphere, the provisioning might randomly fail.

  Workaround: See https://kb.vmware.com/s/article/54139.

- **Issue 2094336: New virtual machines created not seen in NSX-T inventory**
  When Pivotal BOSH deploys new VMs on vSphere, the VMs might not appear in NSX-T's inventory.

  Workaround: See https://kb.vmware.com/s/article/54138.

- **Issue 2022750: BOSH VM cold migration issue**
  Newly provisioned PODs or containers in a container host VM in PKS, PAS, Openshift or other Kubernetes solutions will not have network connectivity under the following conditions:

  - When a container host VM is powered off and subject to cold migration.
  - In some cases, when a container host VM is powering on and is subject to DRS (Distributed Resource Scheduler) before the power-up is completed.
  - If the vNIC of the container host is detached and re-attached.

  Workaround: Delete the container host. In PKS and PAS scenarios, BOSH will recreate the container host.

- **Issue 1998217: HyperBus interface vmk50 might be missing on vSphere ESXi causing container creation failure**
  Container is not created because the HyperBus interface vmk50 might be missing on vSphere ESXi.

  Workaround: Complete the following steps:

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

Known Issues

The known issues are grouped as follows.

- General Known Issues
- Installation Known Issues
- NSX Manager Known Issues
- NSX Edge Known Issues
- Logical Networking Known Issues
- Security Services Known Issues
- KVM Networking Known Issues
- Load Balancer Known Issues
- Solution Interoperability Known Issues
- Operations and Monitoring Services Known Issues
- Upgrade Issues
- API Known Issues
- NSX Policy Manager
- NSX Cloud
- NSX Container Plug-in (NCP) Known Issues
- Documentation Errata and Additions

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

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-T 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 Installation Guide.

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

  Workaround: Manually uninstall VIBs from the host.

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
     ```

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.

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.

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.

KVM Networking Known Issues

- **Issue 1775916: The resolver API `POST /api/v1/error-resolver?action=resolve_error` does not resolve errors after a RHEL KVM host fails to be added to the fabric**
  After a RHEL KVM host fails to be added to the fabric and the NSX Manager user interface shows the installation status as failed, the resolver API `POST /api/v1/error-resolver?action=resolve_error` is run to resolve errors. However, adding the host to the fabric again results in the following error messages:
  ```
  Failed to install software on host. Un-handled deployment plug-in perform-action.
  Install command failed.
  ```

Workaround: Complete the following steps.

1. Manually remove the following packages.

   ```
   rpm -e glog-0.3.1-1nn5.x86_64
   rpm -e json_spirit-v4.06-1.el6.x86_64
   rpm -e kmod-openvswitch-2.6.0.4557686-1.el7.x86_64
   rpm -e nicira-ovs-hypervisor-node-2.6.0.4557686-1.x86_64
   rpm -e nsx-agent-1.1.0.0.0.4690847-1.el7.x86_64
   rpm -e nsx-aggservice-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-cli-1.1.0.0.0.4690892-1.el6.x86_64
   rpm -e nsx-da-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-host-1.1.0.0.0.4690932-1.x86_64 rpm -e nsx-host_node_status_reporter-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-lldp-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-logical_exporter-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-mpa-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-netcpa-1.1.0.0.0.4690924-1.el7.x86_64 rpm -e nsx-sfhc-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-support-bundle-client-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsx-transport_node_status-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e nsxa-1.1.0.0.0.4690845-1.el7.x86_64
   rpm -e openvswitch-2.6.0.4557686-1.x86_64
   rpm -e openvswitch-selinux-policy-2.6.0.4557686-1.noarch
   rpm -e python-simplejson-3.3.3-1.el7.x86_64
   ```

   If there are any error while running the `rpm -e` command, include the `--noscripts` flag to the command.
2. Run the resolver API `POST /api/v1/error-resolver?action=resolve_error`.
3. Add the KVM host to the fabric again.

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

- **Issue 1611154: VMs in one KVM transport node cannot reach VMs located in another transport node**
  When multiple IP pools are used for VTEPs that belong to different networks, the VM on the KVM host might not reach the VM deployed on other hosts that have VTEP IP addresses from a different IP pool.

  Workaround: Add routes so that the KVM transport node can reach all of the networks used for VTEP on other transport nodes.
  For example, if you have two networks 25.10.10.0/24 and 35.10.10.0/24 and the local VTEP has the IP address 25.10.10.20 with gateway 25.10.10.1, you can use the following command to add the route for another network:

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

- **Issue 1654999: Connection tracking of underlay traffic reduces available memory**
  When establishing a large number of connections between virtual machines, you might experience the following symptoms.
  In the /var/log/syslog or /var/log/messages file, you see entries similar to:

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

```
Apr 26 11:45:44 prmh-nsx-perf-server149 kernel: [1625289.958436] nf_conntrack: table full, dropping packet
```

The issue seems to manifest itself when default firewall rules have been configured. The issue does not manifest itself if firewall rules are not configured (For example: Logical switches are put in the firewall exclusion list).
Note: The preceding log excerpts are only examples. Date, time, and environmental variables might vary depending on your environment.

Workaround: Add a firewall rule to disable connection tracking for UDP on port 6081 on underlay devices.
Here is an example command:

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

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

- **Issue 2002353: Using Linux Network Manager to manage a KVM host uplinks is not supported**
  NSX-T manages all the NICs on KVM hosts that are used for N-VDS. Configuration error occurs when the Network Manager is also enabled for these uplinks.

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

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

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

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

Operations and Monitoring Services Known Issues

- **Issue 1749078: After deleting a tenant VM on an ESXi host and the corresponding host transport node, deleting the ESXi host fails**
  Deleting a host node involves reconfiguring various objects and can take several minutes or more.

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

- **Issue 1761955: Unable to connect a VM's vNIC to an NSX-T logical switch after registering the VM**
  If an existing vmx file is used to register a VM on an ESXi host, the register operation ignores following vNIC-specific errors:

  - vNICs that are configured with invalid network backing.

    - VIF attachment failures for vNICs that are connected to an NSX-T logical-switch.

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

- **Issue 1774858: On rare occasions, the NSX Controller cluster becomes inactive after running for multiple days**
  When the NSX Controller cluster becomes inactive, all transport and NSX Edge nodes lose connectivity to the NSX Controllers and changes to the configuration cannot be made. However, data traffic is unaffected.

  Workaround: Complete the following steps.

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

- **Issue 1576304: Dropped-byte count is not included as part of the Port Status and Statistics report**
  When using /api/v1/logical-ports/<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

Upgrade Issues

- **Issue 1930705: vMotion of VMs connected to the logical switches fails during the Management Plane upgrade**
  During the Management Plane upgrade, attempting to vMotion VMs connected to a logical switch fails.

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

- **Issue 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 change 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

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.

NSX Policy Manager

- **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

- **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: Failover to secondary PCG won't take effect if customer had installed agent on the VM by specifying only sing1le PCG's DNS name using the --gateway option.**
  Workload VM is not able to connect to the PCG after failover and hence PCG won't be able to enforce/realize any logical state on the VM.

  Recommendation: Don't use `--gateway` option at all when installing agents on the workload VMs.

- **Issue 2071374: A string of error messages appears while installing the NSX Agent on VMs with certain Linux distributions (such as Ubuntu 14.04, Ubuntu 16.04)**
  On VMs that have "nscd" running, you will see error messages like: "sent invalidate(passwd) request, exiting" while installing NSX agent.

  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.

- **NSX support for Red Hat Enterprise Linux 7.4 with kmod.x86_64 0:20-15.el7_4.6**

NSX does not support VM instances which run Red Hat Enterprise Linux 7.4 with kmod.x86_64 0:20-15.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.

- **Get realized state of a firewall rule in NSX Cloud is not supported**
  Getting the realized state of DFW firewall rule on VM instances using the NSX-T Manager API

  ```
  /api/v1/firewall/rules/<rule-id>/state
  ```

  is not fully supported in this release

  No workaround exists at this time

- **Issue 2066636: While undeploying PCGs from a VNET that has quarantine policy enabled, the last PCG undeployment will get stuck.**
  Last PCG depployment gets stuck. If CSM API is polled to check the gateways status, then it shows the PCG being stuck in deleting the resource group stage. On the Microsoft Azure portal, you see a few security groups in PCG's Resource Group being present while the Resource Group itself is marked for deletion.

  Workaround:

  - Disable the quarantine mode of the VNet being off-boarded.
  - Move the VMs to a user-created security group.
  - Trigger undeployment of the last PCG.

NSX Container Plug-in (NCP) Known Issues

- **PAS 2.1.0 CNI change**
  Due to the CNI plugin change in PAS 2.1.0, none of the NSX-T Tile 2.1.x and 2.2.x 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 2127607: After removing the TLS specification from Kubernetes Ingress, the default pool is not moved to the HTTP virtual server**
  A Kubernetes Ingress is configured with TLS and a default backend. After removing the TLS specification, the rules are transferred from the HTTPS virtual server to the HTTP virtual server but the default pool for the backend service remains in the HTTPS virtual server.

  Workaround: Delete the Kubernetes Ingress, restart NCP, and re-create the Kubernetes

Ingress.

- **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 2193714: Error in documentation regarding the rewrite annotation**
  In the NSX-T Container Plug-in for Kubernetes and Cloud Foundry - Installation and Administration Guide, an example illustrating how to specify a rewrite annotation (https://docs.vmware.com/en/VMware-NSX-T/2.2/com.vmware.nsxt.ncp_kubernetes.doc/GUID-FF86F9C2-E70D-418B-AFDB-B3CF8DE106C2.html) contains the following line:

  ```
  ncp/rewrite_target: "/"
  ```

  This line is incorrect. The correct annotation should be:

  ```
  ingress.kubernetes.io/rewrite-target: /
  ```

- **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.