VMware NSX-T Data Center 2.4.1 Release Notes

VMware NSX-T Data Center 2.4.1 | 21 May 2019 | Build 13716575

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

VMware HCX now supports NSX-T for virtual machine migration to on-premises NSX-T based deployments. This enables customers to mass migrate virtual machines from NSX Data Center for vSphere to NSX-T, from NSX-T to NSX-T cross-site migrations, and non-NSX vSphere environments to NSX-T based SDDC environments.

Password policy enhancements were added to release 2.4.0, which enforce a minimum password length of 12 characters for default passwords and introduced the ability to set password expiration times. By default, passwords expire after 90 days. See Knowledge Base article 70691 for instructions on resetting passwords and adjusting password expiration.

Compatibility and System Requirements

For compatibility and system requirements information, see theNSX-T Data Center Installation Guide.

API and CLI Resources

See code.vmware.com to use the NSX-T Data Center APIs or CLIs for automation.

The API documentation is available from the **API Reference** tab. The CLI documentation is available from the **Documentation** tab.

Document Revision History

21 May 2019. First edition.
03 June 2019. Second edition. Added resolved issue 2339832.
20 June 2019. Third edition. Added known issues 2261818, 2334442.
21 June 2019. Fourth edition. Moved issue 2304571 to resolved status.
23rd August 2019. Fifth edition. Added known issues 2362688, 2395334, and 2392093.
12 November 2019. Sixth edition. Moved issue 2295470 to resolved status.

Resolved Issues

- **Fixed Issue 2248345: After installation of the NSX-T Edge, the machine boots up with blank black screen.**
  Unable to install NSX-T Edge on HPE ProLiant DL380 Gen9 machine.

- **Fixed Issue 2264386: Transport Node deletion takes place even though the Transport Node is a part of NS Group**
  Transport Node deletion is permitted even when the node is part of an NS Group. Deletion should be prevented. If you encounter this issue, you must recreate the NS Groups and rebuild the relationships with its Transport Nodes.

- **Fixed Issue 2275869: cfgAgent log rolling over in <1 minute on ESXi host if rules on the host have tags longer than 31 characters**
  Frequent log rollings may lead to loss of useful information in cfgAgent.log for debugging and troubleshooting on host. Log location on ESXi host: `/var/log/cfgAgent.log`

- **Fixed Issue 2288872: Install state shown as "Node not ready"**
  The Edge node is not getting onboarded. The Transport Node configuration state is Pending, and so cannot be added to an Edge cluster. Log location: `/var/log/proton/nsxapi.log`

- **Fixed Issue 2291267: PCM-created default gateway policy section has no sequence number assigned, so policy defaults it to 0**
  If a user creates gateway policies without sequence numbers or insert_top options, a policy conflict results. Log location: `/var/log/policy/policy.log`

- **Fixed Issue 2292995: The realization status is set to error even though all the configured rules are programmed in OVS**
  The API gives a false negative impression even when the DFW rules are programmed in the data plane.

- **Fixed Issue 2292997: Certain logical router interfaces may fail to create for Linux network stack**
  Certain logical router interfaces may fail to create for Linux network stack, returning the following error: `errorCode="EDG0100002", Operation failed creating sub-interface: max sub-interface exceeded`. As a result, traffic forwarded by tier0 service router (T0 SR) may be dropped due to missing routes.

- **Fixed Issue 2295470 - Firewall filters not present after migrating to NSX-T from NSX for vSphere.**

If services are used in many firewall rules, frequent updates on services may be caused during the migration process. As a result, firewall filters are not installed on the ESXi host. This may cause traffic disruption.

- **Fixed Issue 2285117: Kernel upgrade on NSX Managed VMs is not supported**
  On some Linux Ubuntu marketplace images, the kernel automatically upgrades itself on reboot of the VM. As a result, the nsx-agent does not function as expected. Although the NSX agent may appear to be functional, there will be some unrealized networking policies which affect the nsx-agent. The agent retries realizing these policies over and over, causing high CPU usage.
- **Fixed Issue 2252776: Transport Node Profile fails to be applied on one of the cluster member hosts even if validation error that has occurred previously on the host is now resolved**
  TNP is applied on the cluster. But TNP cannot be applied on one of the cluster member hosts because one of the validations could not be passed (e.g., VMs are powered on on the host). User resolves the issue, but validation is still shown on UI and TNP is not automatically applied on that host.
- **Fixed Issue 228688: BGP neighbor should be deleted first while deleting IPsec Route base session if BPG is configured over VTI**
  If BGP is configured over VTI and you delete the IPsec session, both SR will be in a down state which in turns blocks the traffic. In order to resume the traffic, the BGP neighbor configured for VTI should be deleted. In this scenario, only the BGP configured is over VTI.
- **Fixed Issue 2288509: MTU property is not supported for Tier0/Tier1 service interface (central service port)**
  MTU property is not supported for Tier0/Tier1 service interface (central service port).
- **Fixed Issue 2266553: In NSX appliance, a service may fail to initialize at its first boot**
  The deployed node is unable to serve requests, or unable to form a cluster.
- **Fixed Issue 2267632: Loss of GI Protection configuration**
  Guest protection rule published on Policy UI shows SUCCESS. Corresponding change in the behavior is not reflected on the guest VM. OpsAgent logs at the same time shows restart. Loss of guest VM protection.
- **Fixed Issue 2288773: Old TLS protocol API still available, gets overwritten**
  NSX-T has new API for setting NSX TLS protocol versions and cipher suites, which updates all nodes in an NSX-T cluster. However, the old API is still available. This can be used but the new settings will be overwritten by the global settings.
- **Fixed Issue 2269901: vmk interface is not included in packet capture CLI**
  This command cannot be issued.
- **Fixed Issue 2304571 - Critical error (PSOD) may occur when running L3 traffic using VDR.**
  Pending arp(ND) entry is not properly protected in some cases which may cause critical error (PSOD).
- **Fixed Issue 2275985: Vnics not connected to logical switch are listed as options for**

NSGroup direct members
A vnic that is not connected to a logical switch is added as direct member of the NSGroup. Operation succeeds but the policies applied on that group do not get enforced on the vnic.

- **Fixed Issue 2279973: If a blank group is created and upgrade proceeds, after MP upgrade, that blank group shows as not started**
  This occurs If a blank group is created and upgrade proceeds.
- **Fixed Issue 2282389: UC upgrade plan not in sync with VC cluster membership if ESX is moved across clusters**
  When ESX is moved from one cluster to another in VC, the change is not reflected in UC upgrade plan. This may lead to more than one HOST entering maintenance mode at the same time if user has selected "Parallel Upgrade" across groups.
- **Fixed Issue 2288921: Upgrade status goes out of synch when old version Edge nodes are added**
  Upgrade status goes out of synch if the user adds Edge nodes of an older version following the Edge upgrade. This causes issues in continuing the upgrade call.
- **Fixed Issue 2289278: Policy API throws error but allows configuration of multiple Virtual Servers with same pool with different persistence profile**
  The system does not support configuration of conflicting persistence types for the same pool for different LbVirtualServers. However, the Policy fails to properly validate/reject the conflicting input and allows the configuration. Subsequently, the Policy shows an alarm with the error message.
- **Fixed Issue 2289984: mux_connectivity_status shows as CONNECTED even after nsx-context-mux service is stopped on host**
  When nsx-context-mux or nsx-opsagent is not running on the host, the system (NSX interface or service instance API) incorrectly shows Solution status and GI agent status as running with an unchanged timestamp. As a result, the guest VMs may lose AV protection.
- **Fixed Issue 2290083: Validation missing when creating VLAN based segment**
  When you specify a VLAN transport zone with a VLAN ID property, the system fails to validate and identify the error. As a result, the intent will fail during realization and raise an error.
- **Fixed Issue 2290669: As the number of virtual servers increases, the configuration time for each increases**
  As the number of virtual servers increases, the configuration time for each increases due to large numbers of validation. For the first 100 virtual servers, the average response time is around 1 second. After 250 virtual servers, the average response time increases to 5-10 seconds. After 450 virtual servers, the response time increases to around 30 seconds.
- **Fixed Issue 2291625: PCG upgrade status changes from SUCCESS to NOT_STARTED after upgrade plan synch**
  This issue is only encountered if the user upgrades the PCG and then tries to upgrade more Agents/PCG afterward.
  In the recommended workflow, after PCG upgrade there are no more cross-cloud components to upgrade via the UC interface.

This is not impacting any functionality. The status of the previously successfully completed PCG upgrade is shown as "None" on the upgrade UI.

- **Fixed Issue 2291872: Log message shows a warning message when TFTP service is used in firewall rule**
  Log message shows irrelevant warning message when TFTP service is used in firewall rule.Log location on the ESXi node: /var/log/cfgAgent.log.

- **Fixed Issue 2292096: CLI command "get service router config route-maps" returns an empty output**
  CLI command "get service router config route-maps" returns an empty output even when route-maps are configured. This is a display issue only.

- **Fixed Issue 2292526: "Host not reachable” message shown when adding host**
  When adding an ESXi host, the “Host not reachable” message shows but does specify reason. The likely cause is incorrect credentials.

- **Fixed Issue 2292701: User unable to update the sequence number in a binding map**
  The user cannot change the ordering or precedence of the profiles applied to an entity by updating the sequence number.

- **Fixed Issue 2293227: After upgrade to 2.4, IDFW rules are not applied for VMs running VMTools 10.3.5**
  After performing a live NSX-T upgrade, IDFW rules are not applied for VMs running VMTools 10.3.5, resulting in possible loss of AV protection for those VMs.

- **Fixed Issue 2994002: Tier1 not listed in Tier0/ Tier1 Gateway dropdown list for selection in DNS forwarder creation**
  On a large scale deployment with thousands of records, Tier1 is not listed in Tier0/ Tier1 Gateway dropdown list for selection in the DNS forwarder creation workflow. As a result, you must use the API to configure DNS forwarder creation.

- **Fixed Issue 2294345: Running Application Discovery Classification on a group with both ESXi-hosted and KVM-hosted VMs can fail**
  Application Discovery feature is only supported on ESXi hypervisors. For groups of VMs that are on mixed hosts that include unsupported hosts, Application Discovery Classification results are not guaranteed.

- **Fixed Issue 2294821: NSX appliance information displays in the cluster monitoring dashboard with error "failure to delete node" with no guidance for user to handle the situation.**
  This issue has been observed after the user has tried to delete the auto-deployed node via the interface, and the powering off of the node has failed. If the cluster loses a node, you must manually add a new node and clean up the configuration states using the workaround below.

- **Fixed Issue 2281095: When host which has svm deployed was re-added to same cluster, no callback triggered from EAM**
  All guest VMs could be unprotected. NSX UI will not come out of in-progress state.

- **Fixed Issue 2295564: Edge node controller connectivity may go down after upgrading from 2.3 to 2.4**
  This is an intermittent issue that will impact some north-south traffic.
- **Fixed Issue 2296888: The Transport Node (TN)/Transport Node Profile (TNP) configuration cannot have both PNIC Only Migration flag set to true and VMK Mappings for Install populated across host switches**
  When giving a mismatch of configuration (both PNIC only migration flag set to true and VMK Mappings for Install populated across host switches) during CREATE, the following exception is seen:

  VMK migration for host b17afc36-bbdc-491a-b944-21f73cf91585 failed with error [com.vmware.nsx.management.switching.common.exceptions.SwitchingException: TransportNode [TransportNode/b17afc36-bbdc-491a-b944-21f73cf91585] can not be updated or deleted while migrating ESX vmk interface null to [null].]. (Error code: 9418)

  When giving a mismatch of configuration during UPDATE, the following exception is seen: General error (Error code: 400)

  An exception is seen when applying the TN/TNP configuration which contains both the PNIC Only Migration flag set to true and a vmk migration mapping.
- **Fixed Issue 2287124: After deploying PCG on a Microsoft Azure VNet, the VNet's tile in CSM erroneously reports a warning**
  After deploying PCG on a Microsoft Azure VNet, in CSM the VNet reports a warning sign (yellow triangle with an exclamation point). If you hover over the warning icon, CSM reports that the status of MP (Management Plane) and CCP (Control Plane) is Unknown. However, there may not be any problem with connectivity and the warning is displayed in error.
- **Fixed Issue 2273651 - After deleting transport node, user is unable to SSH into the host.**
  Observed in KVM implementations. The user deletes a transport node and receives a message that the deletion was successful. However, afterward the user is unable to access the same host via SSH. The issue is likely caused by the presence of an open virtual switch (OVS) that is not managed by NSX-T and was likely pre-installed as part of the KVM template.
- **Fixed Issue 2297157 - Load Balancing HTTPS performance is impacted by the FIPS mode.**
  Load balancing performance can be negatively affected when the default FIPS mode is enabled.
- **Fixed Issue 2290688 - Upgrading Windows 2016 VMs in AWS fails.**
  Upgrade of multiple Windows workload VMs fail in AWS. VM upgrade status displays in AWS portal as stuck at "1/2 Check." Retrying fails also. This issue is observed only within same NSX-T version upgrades.
- **Fixed Issue 2203863 - Identity firewall rules are not supported for UDP and ICMP traffic.**

Identity firewall rules do not work with ping testing The current functionality is supported for TCP traffic only.

- **Fixed Issue 2248186 -The BGP router installs IPv6 routes from its neighbor with its own interface as next hop.**
  As a result, IPv6 forwarding for the installed route may fail and cause a forwarding loop.
- **Fixed Issue 2281537 - Post-migration, the ESXi transport node with multi-VTEP fails to start BFD session.**
  After migrating a NSX-V node to NSX-T, the ESXi transport node with multi-VTEP fails to start BFD session on all VTEPs to Edge nodes.
- **Fixed Issue 2297918 - Post-upgrade from 2.3.1 to 2.4, unable to remove NSX from cluster.**
  After upgrading a cluster from 2.3.1 to 2.4, NSX-T cannot be removed and fails with the following message: "Failed to remove NSX on the cluster: Related transport node template or transport node collection exists for this fabric template. Transport node template or transport node collection must be deleted before a delete/disable on this fabric template."
- **Fixed Issue 2298499 - VPN fails between public cloud gateway and peer host if gateway is not deployed with public IP.**
  The VPN tunnel between the public cloud gateway (PCG) and the peer host cannot be established if the PCG is deployed without a public IP address on the uplink. The reason is that the PCG by default is performing SNAT on the VPN traffic by default.
- **Fixed Issue 2316831 - IPv6 traffic always get load shared, even though ECMP is disabled.**
  Disabling ECMP from Policy is not effective for IPv6 Unicast address family. (Although it is effective for IPv4 Unicast address family.)
- **Fixed Issue 2334515 - Use of IPv4 link local range (169.254.0.0/16) for T0-T1 router link port does not work.**
  Use of IPv4 link local range (169.254.0.0/16) for T0_t1 router link port does not work. However, using an IP range which is not IPv4 link-local address range (169.254.0.0/16) for T0_t1 router link does work.
- **Fixed Issue 2339832: Unable to apply node certificate or set cluster certificate with a CA-signed certificate.**
  This results in messages "Error updating certificate usage." or "An error occurred setting the cluster certificate."
  This may be caused after repeatedly applying certificates, alternating between node certificates and cluster certificate. The certificate will not be applied properly and REST API calls through the VIP may no longer work. Log location: /var/log/proton/nsxapi.log.

  If you encounter this error before upgrading to 2.4.1, use self-signed certificates instead.

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
- Load Balancer Known Issues
- Solution Interoperability Known Issues
- Operations and Monitoring Services Known Issues
- Upgrade Known Issues
- API Known Issues
- NSX Cloud Known Issues

General Known Issues

- **Issue 2239365: "Unauthorized" error is thrown**
  This error may result because the user attempts to open multiple authentication sessions on the same browser type. As a result, login will fail with above error and cannot authenticate. Log location: /var/log/proxy/reverse-proxy.log /var/log/syslog

  Workaround: Close all open authentication windows/tabs and retry authentication.

- **Issue 2252487: Transport Node Status is not saved for BM edge transport node when multiple TN is added in parallel**
  The transport node status is not shown correctly in MP UI.

  Workaround:

  1. Reboot the proton, all transport node status can be updated correctly.
  2. Or, use the API https://<nsx-manager>/api/v1/transport-nodes/<node-id>/status?source=realtime to query the transport node status.
- **Issue 2256709: Instant clone VM or VM reverted from a snapshot loses AV protection briefly during vMotion**
  Snapshot of a VM is reverted and migrates the VM to another host. Partner console doesn't show AV protection for migrated instant clone VM. There is a brief loss of AV protection.

  Workaround: None.

- **Issue 2261431: Filtered list of datastores is required depending upon the other deployment parameters**
  Appropriate error on UI seen if incorrect option was selected. Customer can delete this deployment and create a new one to recover from error.

  Workaround: Select shared datastore if you are creating a clustered deployment.

- **Issue 2274988: Service chains do not support consecutive service profiles from the same service**
  Traffic does not traverse a service chain and it gets dropped whenever the chain has two consecutive service profiles belonging to the same service.

Workaround: Add a service profile from a different service to make sure that no two consecutive service profiles belong to the same service. Alternatively, define a third service profile which will perform the same operations of the two original ones concatenated, then use this third profile alone in the service chain.

- **Issue 2275285: A node makes a second request to join the same cluster before the first request is complete and the cluster stabilized**
  The cluster may not function properly and the CLI commands get cluster status, get cluster config could return an error.

  Workaround: Do not issue any new join command within 10 minutes to join the same cluster after the first join request.

- **Issue 2275388: Loopback interface/connected interface routes could get redistributed before filters gets added to deny the routes**
  Unnecessary routes updates could cause the diversion on traffic for few seconds to min.

  Workaround: None.

- **Issue 2275708: Unable to import a certificate with its private key when the private key has a passphrase**
  The message returned is, "Invalid PEM data received for certificate. (Error code: 2002)". Unable to import a new certificate with private key.

  Workaround:

  1. Create a certificate with private key. Do not enter a new passphrase when prompted; press Enter instead.
  2. Select "Import Certificate" and select the certificate file and the private key file.

  Verify by opening the key-file. If a passphrase was entered when generating the key, the second line in the file will show something like "Proc-Type: 4,ENCRYPTED".

  This line is missing if the key-file was generated without passphrase.

- **Issue 2277742: Invoking PUT https://<MGR_IP>/api/v1/configs/management with a request body that sets publish_fqdns to true can fail if the NSX-T Manager appliance is configured with a fully qualified domain name (FQDN) instead of just a hostname**
  PUT https://<MGR_IP>/api/v1/configs/management cannot be invoked if a FQDN is configured.

  Workaround: Deploy the NSX Manager using a hostname instead of a FQDN.

- **Issue 2279249: Instant clone VM loses AV protection briefly during vMotion**
  Instant clone VM migrated from one host to another. Immediately after migration, eicar file is left behind on the VM. Brief loss of AV protection.

  Workaround: None.

- **Issue 2292116: IPFIX L2 applied to with CIDR-based group of IP addresses not listed on UI when group is created via the IPFIX L2 page**
  If you try to create a group of IP addresses from Applied to dialog and enter wrong IP

address or CIDR in the Set Members dialog box, those members are not listed under groups. You have to edit that group again to enter valid IP addresses.

  Workaround: Go to the groups listing page and add IP addresses in that group. Then that group can start populating in the Applied to dialog.

- **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 2268406 - Tag Anchor dialog box doesn't show all tags when maximum number of tags are added.**
  Tag Anchor dialog box doesn't show all tags when maximum number of tags are added, and cannot be resized or scrolled through. However, the user can still view all tags in the Summary page. No data is lost.

  Workaround: View the tags in the Summary page instead.

- **Issue 2310650 - Interface shows "Request timed out" error message.**
  Multiple pages on interface shows the following message: “Request timed out. This may occur when system is under load or running low on resources”

  Workaround: Using SSH, log into the NSX Manager VM and run the "start search resync manager" CLI command.

- **Issue 2320529 - "Storage not accessible for service deployment" error thrown after adding third-party VMs for newly added datastores.**
  "Storage not accessible for service deployment" error thrown after adding third-party VMs for newly added datastores even though the storage is accessible from all hosts in the cluster. This error state persists for up to thirty minutes.

  Workaround: Retry after thirty minutes. As an alternative, make the following API call to update the cache entry of datastore:
  https://{{NsxMgrIP}}/api/v1/fabric/compute-collections/<CC Ext ID>/storage-resources?uniform_cluster_access=true&source=realtime
  where NsxMgrIP is the IP address of the NSX manager where the service deployment API has failed, and CC Ext ID is the identifier in NSX of the cluster where the deployment is

being attempted.

- **Issue 2320855 - New VM security tag is not created if user doesn't click Add/Check button.**
  Interface issue. If a user adds a new security tag to a policy object or inventory and clicks **Save** without first clicking the **Add/Check** button next to the tag-scope pair field, the new tag pair is not created.

  Workaround: Be sure to click the **Add/Check** button before clicking **Save**.

- **Issue 2328126 - Bare Metal issue: Linux OS bond interface when used in NSX uplink profile returns error.**
  When you create a bond interface in the Linux OS and then use this interface in the NSX uplink profile, you see this error message: "Transport Node creation may fail." This issue occurs because VMware does not support Linux OS bonding. However, VMware does support Open vSwitch (OVS) bonding for Bare Metal Server Transport Nodes.

  Workaround: If you encounter this issue, see Knowledge Article 67835Bare Metal Server supports OVS bonding for Transport Node configuration in NSX-T.

- **Issue 2334442: User does not have permission to edit or delete created objects after admin user renamed**
  User does not have permission to edit or delete created objects after admin user is renamed. Unable to rename admin/auditor users.

  Workaround: Restart policy after rename by issuing the command "service nsx-policy-manager restart"

- **Issue 2261818: Routes learned from eBGP neighbor are advertised back to the same neighbor**
  Enabling bgp debug logs will indicate packets being received back and packet getting dropped with error message. BGP process will consume additional cpu resources in discarding the update messages sent to peers. If there are large number of routes and peers this can impact route convergence.

  Workaround: None.

Installation Known Issues

- **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.

NSX Manager Known Issues

- **Issue 2282798 - Host registration may fail when too many requests/hosts try to register with the NSX Manager simultaneously.**

This issue causes the fabric node to be in a FAILED state. The Fabric node status API call shows “Client has not responded to heartbeats yet”. The /etc/vmware/nsx-mpa/mpaconfig.json file on the host is also empty.

Workaround: Use the following procedure to recover from this issue.

1. Use the Resolver option.
2. Delete the FN from NSX.
3. Re-add the FN again manually through CLI command "join management-plane".

NSX Edge Known Issues

- **Issue 2283559: /routing-table and /forwarding-table MP APIs return an error if the edge has 65k+ routes for RIB and 100k+ routes for FIB**
  If the edge has 65k+ routes for RIB and 100k+ routes for FIB, the request from MP to Edge takes more than 10 seconds and results in a timeout. This is a read-only API and has an impact only if they need to download the 65k+ routes for RIB and 100k+ routes for FIB using API/UI.

  Workaround: There are two options to fetch the RIB/FIB.

  - These APIs support filtering options based on network prefixes or type of route. Use these options to download the routes of interest.
  - CLI support in case the entire RIB/FIB table is needed and there is no timeout for the same.
- **Issue 2204932 - Configuring BGP Peering can delay HA failover recovery.**
  When Dynamic-BGP-Peering is configured on the routers that peer with the T0 Edges and a failover event occurs on the Edges (active-standby mode), BGP neighborship may take up to 120 seconds.

  Workaround: Configure specific BGP peers to prevent the delay.

- **Issue 2285650 - BGP route tables populated with unwanted routes.**
  When the allowas-in option is enabled as part of the BGP configuration, routes advertised by Edge nodes are received back and installed in the BGP route table. This results in excess memory consumption and routing calculation processing. If higher local preference is configured for the excess routes, this forwarding loop may result in the route table on some routers being populated with redundant routes.

  For example, route X originates on router D, which is adertised to routers A and B. Router C, on which allowas-in is enabled, is peered with B, so it learns route X and installs it in its route table. As a result, there are now two paths for route X to be advertised to router C, resulting in the problem.

  Workaround: You can prevent forwarding loops by configuring the problematic router (or its peer) to block routes being advertised back to it.

Logical Networking Known Issues

- **Issue 2243415: Customer unable to deploy EPP service using Logical Switch (as a management network)**

On the EPP deployment screen, the user cannot see a logical switch in the network selection control. If the API is used directly with logical switch mentioned as management network, user will see the following error: "Specified Network not accessible for service deployment."

Workaround: Deploy using other type of switches like local or distributed.

- **Issue 2288774: Segment port gives realization error due to tags exceeding 30 (erroneously)**
  User input incorrectly tries to apply more than 30 tags. However, the Policy workflow does not properly validate/reject the user input and allows the configuration. Then Policy then shows an alarm with the proper error message that the user should not use more than 30 tags. At that point the user can correct this issue.

  Workaround: Correct the configuration after the error displays.

- **Issue 2275412: Port connection doesn't work across multiple TZs**
  Port connection can be used only in single TZ.

  Workaround: None.

- **Issue 2320147 - VTEP missing on the affected host.**
  If a LogSwitchStateMsg is removed and added in the same transaction and this operation is processed by the central control plane before management plane has sent the Logical Switch, the Logical switch state will not be updated. As a result, traffic cannot flow to or from the missing VTEP.

  Workaround: If you encounter this issue, restart the central control plane.

- **Issue 2327904 - After using pre-created Linux bond interface as an uplink, traffic is unstable or fails.**
  NSX-T does not support pre-created Linux bond interfaces as uplink.

  Workaround: For uplink, use OVS native bond configuration from uplink profile.

- **Issue 2295819 - L2 bridge stuck in "Stopped" state even though Edge VM is Active and PNIC is UP.**
  L2 bridge may be stuck in "Stopped" state even though the Edge VM is Active and the PNIC that backs the L2 bridge port is UP. This is because the Edge LCP fails to update the PNIC status in its local cache, thereby assuming that the PNIC is down.

  Impact to customer:
  Traffic outage for VMs that are reachable by edge l2bridge port

  Workaround: Restart the local-control agent on the affected Edge VM.

- **Issue 2392093: Traffic drops due to RPF-Check**
  RPF-Check may result in dropped traffic if traffic is hair-pinned through a T0 downlink, and Tier0 and Tier1 routers are on the same Edge Node.

  Workaround: None.

Security Services Known Issues

- **Issue 2395334 - (Windows) Packets wrongly dropped due to stateless firewall rule conntrack entry.**
  Stateless firewall rules are not well supported on Windows VMs.

  Workaround: Add a stateful firewall rule instead.

- **Issue 2458384 - NSX-T Manager interface pages fail to load with 403 error.**
  Observed in release versions 2.4.0 and 2.4.1. This issue affects both admin and Identity Manager logins. The FQDN of the NSX-T Manager uses the *.SLD.TLD format. For example: *.co.uk, *.co.il, *.com.au and so on.

  Workaround: Access the NSX-T Manager UI using shortname or IP instead of FQDN. See Knowledge Base article 71217.

- **Issue 2296430 - NSX-T Manager API does not provide subject alternative names during certificate generation.**
  NSX-T Manager API does not provide subject alternative names to issue certificates, specifically during CSR generation.

  Workaround: Create the CSR using an external tool that supports the extensions. After the signed certificate is received from the Certificate Authority, import it into NSX-T Manager with the key from the CSR.

- **Issue 2294410 - Some Application IDs are detected by the L7 firewall.**
  The following L7 Application IDs are detected based on port, not application: SAP, SUNRPC, and SVN. The following L7 Application IDs are unsupported: AD_BKUP, SKIP, and AD_NSP.

  Workaround: None. There is no customer impact.

- **Issue 2314537 - Connection status is down after vCenter certificate and thumbprint update.**
  No new updates from vCenter sync with NSX and all on-demand queries to fetch data from vCenter will fail. Users cannot deploy new Edge/Service VMs. Users cannot prepare new clusters or hosts added in the vCenter. Log locations: /var/log/cm-inventory/cm-inventory.log and /var/log/proton/nsxapi.log on the NSX Manager node.

  Workaround: Log in to each NSX Manager VM and switch to root user. Run the "/etc/init.d/cm-inventory restart" command on each Manager node.

Load Balancer Known Issues

- **Issue 2290899: IPSec VPN does not work, control plane realization for IPSec fails**
  IPSec VPN (or L2VPN) fails to comes up if more than 62 LbServers are enabled along with IPSec service on Tier-0 on the same Edge node.

  Workaround: Reduce the number of LbServers to fewer than 62.

- **Issue 2318525 - The next-hop IPv6 routes as the eBGP peer's IP address gets changed to its own IP.**

In case of eBGP IP4 sessions, advertised IPv4 routes that have their eBGP peer as the next hop, the next hop of the route is NOT changed at the sender side to its own IP address. This works for IPv4, but for IPv6 sessions, the next hop of the route is changed at the sender side to its own IP address. This behavior can result in route loops.

Workaround: None.

- **Issue 2362688: If some pool members are DOWN in a load balancer service, the UI shows the consolidated status as UP**
  When a pool member is down, there is no indication on the Policy UI where the Pool status is green and Up.

  Workaround: None.

Solution Interoperability Known Issues

- **Issue 2289150: PCM calls to AWS start to fail**
  If you update the PCG role for an AWS account on CSM from*old-pcg-role* to *new-pcg-role*, CSM updates the role for the PCG instance on AWS to *new-pcg-role*. However, the PCM does not know that the PCG role has been updated and as a result continues to use the old AWS clients it had created using *old-pcg-role*. This causes the PCM AWS cloud inventory scan and other AWS cloud calls to fail.

  Workaround: If you encounter this issue, do not modify/delete the old PCG role immediately after changing to new role for at least 6.5 hours. Restarting the PCG will re-initialize all AWS clients with new role credentials.

Operations and Monitoring Services Known Issues

- **Issue 2316943 - Workload unprotected briefly during vMotion.**
  VMware tools takes a few seconds to report correct computer name for VM after vMotion. As a result, VMs added to NSGroups using computer name are unprotected for a few seconds after vMotion.

  Workarround: For groups to be used in DFW rules, use VM name-based criteria instead of computer name-based criteria.

- **Issue 2331683 - Add-Load-balancer form on Advance UI not showing updated capacity of version 2.4.**
  When add-load-balancer form is opened, the form-factor-capacity shown on the Advance UI is not updated as per 2.4 version. The capacity shown is from the previous version.

  Workaround: None.

Upgrade Known Issues

- **Issue 2286030 - Transport node configuration displays as in failed state when upgrading from NSX-T 2.3.x and earlier to 2.4.x.**
  Transport node configuration goes into failed state when upgrading from NSX-T 2.3.x and earlier to 2.4.x due to a Null Pointer Exception. When you have ESXi transport node with vmkernel adapters migrated to N-VDS VLAN logical-switch and then upgrade from NSX-T

2.3.x to NSX-T 2.4.x, a race condition may cause the ESXi transport node configuration status to display as failed. However, the ESXi transport node connectivity with the NSX Manager and controllers is intact during upgrade even after the node is marked for configuration state failed.

Workaround: Update or resend the transport node to reset the configuration state to success.

1. From the NSX Manager, edit the ESXi transport node which displays as failed.
2. On the ESXi transport node configuration pop-up, click **Save**.
   This action resets the status. You do not have to modify the configuration.

- **Issue 2288549: RepoSync fails with checksum failure on manifest file**
  Observed in deployments recently upgraded to 2.4. When an upgraded setup is backed up and restored on a fresh deployed manager, the repository manifest checksum present in the database and the checksum of actual manifest file do not match. This causes the RepoSync to be marked as failed after backup restore.

  Workaround: To recover from this failure, perform the following steps:

  1. Run CLI command `get service install-upgrade`
     Note the IP of "Enabled on" in the results.
  2. Log in to the NSX manager IP shown in "Enabled on" return of the above command.
  3. Navigate to **System** > **Overview**, and locate the node with the same IP as "Enabled on" return.
  4. Click **Resolve** on that node.
  5. After the above resolve operation succeeds, click **Resolve** on all nodes from the same interface.
     All three nodes will now show RepoSync status as **Complete**.
- **Issue 2277543 - Host VIB update fails during in-place upgrade with 'Install of offline bundle failed on host' error.**
  This error may occur when storage vMotion was performed on the host before doing an in-place upgrade from NSX-T 2.3.x to 2.4 and hosts running ESXi-6.5P03 (build 10884925). The switch security module from 2.3.x is not get removed if storage vMotion was performed just before the host upgrade. The storage vMotion triggers a memory leak causing the switch security module unload to fail.

  Workaround: See Knowledge Base article 67444 Host VIB update may fail when upgrading from NSX-T 2.3.x to NSX-T 2.4.0 if VMs are storage vMotioned before host upgrade.

- **Issue 2276398 - When an AV Partner Service VM is upgraded using NSX, there may be up to twenty minutes of protection loss.**
  When a Partner SVM is upgraded, the new SVM is deployed and old SVM is deleted. SolutionHandler connection errors may appear on the host syslog.

  Workaround: Delete the ARP cache entry on the host after upgrade and then ping the Partner Control IP on the host to solve this issue.

- **Issue 2330417 - Unable to proceed with upgrade for non-upgraded transport nodes.**
  When upgrading, the upgrade is marked as successful even though some transport nodes

are not upgraded. Log location: /var/log/upgrade-coordinator/upgrade-coordinator.log.

Workaround: Restart the upgrade-coordinator service.

API Known Issues

- **Issue 2260435 - Stateless redirection policies/rules are created by default by API, which is not supported for east-west connections.**
  Stateless redirection policies/rules are created by default by API, which is not supported for east-west connections. As a result, traffic is not't get redirected to partners.

  Workaround: When creating redirection policies using the policy API, create a stateful section.

- **Issue 2332397 - API allows creation of DFW policies in nonexistent domain.**
  After creating such a policy on a nonexistent domain, the interface becomes unresponsive when user opens up a DFW security tab. The relevant log is /var/log/policy/policy.log.

  Workaround: Create the domain, with the same ID, on which the policy was created. This permits the validation to succeed.

NSX Cloud Known Issues

- **Issue 2275232: DHCP would not work for VMs on cloud if DFWs Connectivity_statregy is changed from BLACKLIST to WHITELIST**
  All the VMs requesting for new DHCP leases would lose IPs. Need to explicitly allow DHCP for cloud VMs in DFW.

  Workaround: Explicitly allow DHCP for cloud VMs in DFW.

- **Issue 2277814: VM gets moved to vm-overlay-sg on invalid value for nsx.network tag**
  VM tagged with invalid nsx.network tag will get moved to vm-overlay-sg.

  Workaround: Remove invalid Tag.