



VMware NSX-T Data Center 2.5.1 Release Notes

VMware NSX-T Data Center 2.5.1 | 19 Dec 2019 | Build 15314288

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](#)
- [General Behavior Changes](#)
- [Available Languages](#)
- [API and CLI Resources](#)
- [Revision History](#)
- [Resolved Issues](#)
- [Known Issues](#)

What's New

Features, Functional Enhancements and Extensions

This release of NSX-T Data Center is a maintenance release and there are no major or minor features, functional enhancements or extensions.

Maintenance Corrections

NSX Edge VM Node now supports vMotion, DRS and vSphere HA

This release provides enhanced support for NSX Edge vNIC connected to trunk portgroup and as a result, vMotion, DRS and vSphere HA are now supported on NSX Edge VM nodes.

NSX Edge Node CPU monitoring enhancement

The NSX Manager API and CLI command for the Edge Node CPU monitoring include CPU utilization of L2-L4 services (or DPDK cores) and L7 services (or non-DPDK cores). In previous releases, the system provided CPU utilization of only L2-L4 services.

Resolved Issues

This release addresses issues documented in the Resolved Issues section.

Compatibility and System Requirements

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

General Behavior Changes

NSX-T Data Center System Communication Port Changes

Starting in NSX-T Data Center 2.5.1, the outbound port used by NSX Cloud for communicating with public clouds has changed from the non-standard port 7442 to port 80.

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.

Available Languages

NSX-T Data Center has been localized into multiple languages: English, German, French, Japanese, Simplified Chinese, Korean, Traditional Chinese, and Spanish. Because NSX-T Data Center localization utilizes the browser language settings, ensure that your settings match the desired language.

Document Revision History

19 Dec 2019. First edition.
22 Dec 2019. Second edition. Added resolved issue 2429162.
17 Jan 2020. Third edition. Added known issue 2481033.
22 Jan 2020. Fourth edition. Added resolved issue 2410596.
03 Feb 2020. Fifth edition. Updated the workaround for known issue 2481033.
18 Feb 2020. Sixth edition. Updated known issue 2436302 with link to KB article.
24 Feb 2020. Seventh edition. Updated known issue 2481033 with link to KB article. Added known issue 2483552.
28 Feb 2020. Eighth edition. Updated a note in the "What's New" section.
03 Mar 2020. Ninth edition. Added known issue 2508429.
01 June 2020. Tenth edition. Added known issues 2470210, 2498350, 2509879, 2512778, 2517232, 2522782, 2523475, 2543353, 2547983, 2561740, 2572505.
25 Sept 2020. Eleventh edition. Added known issues 2586606, 2621322, 2491206.
15 March 2021. Twelfth edition. Added known issue 2730634.

Resolved Issues

- Fixed Issue 2391231 - Detection of changes to Azure VMs might be delayed.**
 Intermittently, changes to Azure VMs on the cloud are detected with a slight delay. As a result, a corresponding delay might affect onboarding the VMs and creating logical entities for the VMs in NSX-T. The maximum delay observed was approximately eight minutes.
- Fixed Issue 2401164 - Backups incorrectly reported as successful despite SFTP server error.**
 If the password expires for the SFTP server used for backups, NSX-T reports the generic error "backup operation unknown error".
- Fixed Issue 2395334 - (Windows) Packets wrongly dropped due to stateless firewall rule conntrack entry.**
 Stateless firewall rules are not well supported on Windows VMs.
- Fixed Issue 2200856 - cloud-service-manager service restart fails.**
 Cloud-service-manager service restart can fail if the user tries it without waiting for the API service to come up for the first time.
- Fixed Issue 2388158 - User unable to edit transit subnet settings in Tier-0 logical router configuration.**
 After creating the Tier-0 logical router, the transit subnet configuration cannot be modified in the NSX Manager interface.
- Fixed Issue 2410806 - Publishing generated recommendation fails with exception citing 500 total limitation.**
 If the total number of members (IP addresses or VMs) in a recommended group exceeds 500, the publication of generated recommendation into a policy configuration will fail with an exception message such as:

"The total number of IPAdressExpressions, MACAddressExpressions, paths in a PathExpression and external IDs in ExternalIDExpression should not exceed 500."
- Fixed Issue 2408453 - VMware Tools 10.3.5 crashes when NSX Guest Introspection driver is installed.**
 VMware Tools 10.3.5 crashes irregularly on Windows VM, most noticeably when the remote session is disconnected or the guest VM is shutting down.
- Fixed Issue 2380833 - Publishing of policy draft with 8,000 or more rules requires a lot of time.**
 A policy draft containing 8,000 or more rules can take a considerable amount of time to publish. For example, a policy draft with 8,000 rules can take 25 minutes to publish.
- Fixed Issue 2343954 - Edge L2 bridge end point interface permits configuration of unsupported VLAN ranges.**
 The Edge L2 Bridge and Point configuration interface permits you to configure VLAN range and multiple VLAN ranges even though these are not supported.
- Fixed Issue 2408972 - During upgrade, vSphere Update Manager fails while remediating last host.**
 During upgrade, vSphere Update Manager remediation fails for the last host that has workloads backed by an NSX-T logical switch.

- **Fixed Issue 2378752 - API allows creation of multiple binding maps under segments or ports.**

Observed only on API. When a user creates multiple binding maps under a segment or port, no error is reported. The issue is seen when the user tries to bind multiple profiles on segment or port simultaneously.

- **Fixed Issue 2419246 - Ubuntu KVM upgrade fails.**

Upgrade of Ubuntu KVM nodes may fail due to *nsx-vdpi* service not running.

- **Fixed Issue 2410596 - "segfault" on NULL flow cache entry added due to flow cache masks running out.**

"segfault" on NULL flow cache entry added due to flow cache masks running out. Active edge will be failed over to standby edge.

- **Fixed Issue 2252487 - Transport Node Status is not saved for BM edge transport node when multiple transport nodes are added in parallel.**

The transport node status is not shown correctly in the NSX Manager UI or through APIs when multiple transport nodes are added in parallel.

- **Fixed 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.

- **Fixed Issue 2275232 - DHCP does not work for public cloud VMs if DFW Connectivity Strategy is changed from BLACKLIST to WHITELIST.**

All the VMs requesting new DHCP leases lose IP addresses when the connectivity strategy is changed from Blacklist to Whitelist. Need to explicitly allow DHCP service for public cloud VMs in DFW.

- **Fixed 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.

- **Fixed Issue 2290899 - IPSec VPN does not work, control plane realization for IPSec fails.**

IPSec VPN (or L2VPN) fails to come up if more than 62 LB servers are enabled along with IPSec service on Tier-0 on the same Edge node.

- **Fixed Issue 1957059 - Host unprep fails if host with existing VIBs are added to the cluster when trying to unprep.**

The host unprep operation fails if VIBs are not removed completely with a host reboot before adding the hosts to the cluster.

- **Fixed 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.

- Fixed Issue 2260435 - Redirect policies/rules created using API are stateless by default, which is not supported for east-west connections.**
 Redirect policies/rules created using API are stateless by default, which is not supported for east-west connections. As a result, traffic will not be redirected to partners.
- Fixed 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.
- Fixed Issue 2294410 - Some unsupported application IDs are detected based on ports by the L7 firewall.**
 The following L7 application IDs are unsupported and removed from NSX-T Data Center:
 - AD_BKUP
 - SKIP
 - AD_NSP
 - SAP
 - SUNRPC
 - SVN
- Fixed Issue 2330417 - Unable to proceed with upgrade of hosts after upgrade started and group changed for the host.**
 When upgrading, the upgrade is marked as successful even though some transport nodes are not upgraded as a result of group change.
- Fixed Issue 2304571 - PSOD may occur when running L3 traffic using VDR.**
 Pending ARP (ND) entry is not properly protected in some cases, which may cause PSOD.
- Fixed Issue 2348994 - Intermittent failure during upgrade of NSX VIBs on ESXi 6.5 p03 Transport Node.**
 Observed in some 2.4.x to 2.5 upgrades. When the NSX VIBs on an ESXi 6.5 p03 transport node are upgraded, the upgrade operation sometimes fails with the following error: "VI SDK invoke exception: Got no data from process: LANG=en_US.UTF-8".
- Fixed Issue 2401715 - Error while updating the compute manager that thumbprint is invalid, even if correct thumbprint is provided.**
 When PNID of vCenter v6.7U3 is changed and its connection with NSX Manager goes down, reconnecting with a new thumbprint fails.
- Fixed Issue 2372653 - Post-upgrade to 2.5, user unable to locate LogicalPort- and LogicalSwitch-based groups in earlier NSX-T versions.**
 After upgrading to 2.5, the LogicalPort- and LogicalSwitch-based groups created from Policy in previous NSX-T versions do not appear in the dashboard interface. However, they can still be located in the API. This is due to a name change caused by the upgrade process. In 2.5, LogicalPort- and LogicalSwitch-based groups appear as Segment- and SegmentPort-based groups.
- Fixed Issue 2337944 - The ESXi host cannot be accessed through the network.**

The management kernel interface in an ESXi host can be migrated to use a standby uplink in NVDS such that the host will lose network connectivity.

- **Fixed Issue 2395390 - RHEL LCP Bundle installation fails on KVM nodes where OpenStack (RHOSP13) is installed.**

RHEL LCP Bundle installation will fail on RHEL OpenStack KVM nodes (RHOSP13).

- **Fixed Issue 2434573 - Central Control Plane (CCP) node fails to join the CCP cluster.**
CCP UUID file `/config/vmware/node-uuid` is modified by node API and causes CCP node to fail to join the CCP cluster. As a result, configurations cannot be pushed to hosts.
- **Fixed Issue 2438674 - The command "nsx-cli" is not working.**
Unable to launch NSX CLI using the command `"nsx-cli"`, because `/scratch/log` is not present
- **Fixed Issue 2442933 - Intermittent network connectivity loss to VMs in multi-tenant environments with overlapping IP subnets.**
Network connectivity is lost until the Gateway MAC entry expires in the Guest VM's ARP cache.
- **Fixed Issue 2445682 - Unable to resync a transport-node to its transport-node profile.**
Unable to resync a transport-node to its transport-node profile.
- **Fixed Issue 2454034 - GRE traffic may not be able to pass through edge, or may be forwarded by the edge with wrong header info.**
The pass-through GRE traffic is handled by flow-cache which is designed to only handle UDP and TCP traffic. This causes GRE traffic to not pass through NSX Edge or be forwarded with incorrect header info.
- **Fixed Issue 2457498 - Missing data or Intermittent data as some nodes will get data and some do not.**
Missing data or Intermittent data as some nodes will get data and some don't when NSX Manager encounters a trim exception.
- **Fixed Issue 2289941 - Datapath soft-limits hit on certain large-scale baremetal deployments, preventing NSX Edge from running.**
Datapath soft-limits hit on certain large-scale baremetal deployments, preventing NSX Edge from running. NSX Edge appliances cannot run and network functions do not work if they depend on this failing deployment.
- **Fixed Issue 2347671 - BFD tunnel down between Edge and ESXi when Edge uses trunk logical switch.**
See [Knowledge Base article 70745](#) for details.
- **Fixed Issue 2364445/2439357 - Service Insertion created Logical Switches do not get cleaned up on failed N-S deployments.**
If an East-West service insertion is deployed in a setup with orphan logical switches from North-South deployment SPF port might end up inheriting the wrong VNI. You have to

clean up logical switches manually.

- **Fixed Issue 2387578 - BFD session is not formed between edges of the same cluster over the management interface.**

BFD traffic is dropped by the intermediary router configured with BFD ACL rules. You are unable to have HA between the edges of the same cluster unless those edges are located in the same L2 domain.

- **Fixed Issue 2392487 - The edge dataplane doesn't come up when the number of cores are increased beyond 20 cores.**

The edge dataplane doesn't come up when the number of cores are increased beyond 20 cores, due to mempool allocators hitting a hard limit because of a configuration parameter.

- **Fixed Issue 2396296 - In ESXi hosts without a scratch partition upgrade may fail with "tmp partition is 90 per cent full" error.**

Host upgrade fails since /tmp does not have sufficient space.

- **Fixed Issue 2408925 - Changing VLAN ID for Edge uplink makes packet forwarding stop.**

Changing VLAN ID for Edge uplink makes packet forwarding stop.

- **Fixed Issue 2411335 - NSX Manager UI is inconsistent across nodes.**

NSX Manager UI is inconsistent across NSX Manager nodes.

- **Fixed Issue 2412406 - Kernel crashes upon Bare Metal Edge installation.**

You may experience kernel crashing during deployment.

- **Fixed Issue 2413487 - VMs migrated using vMotion lose network connectivity.**

VMs migrated using vMotion lose network connectivity. See [Knowledge Base article 74767](#) for more information.

- **Fixed Issue 2415609 - KVM host installation fails intermittently during manual host installation.**

Host goes into install-failed state with error '*Failed to get response from NSX-SFHC component*'. Transport node configuration is not applied on host.

- **Fixed Issue 2418972 - During host migration, some hosts can fail to migrate.**

Host migration fails and cannot be recovered on retry.

- **Fixed Issue 2420763 - Core dump in edge with load balancer prevents new configurations from taking effect.**

A core dump in the *nginx* process can cause a deadlock requiring the load balancer to be detached and reattached to recover.

- **Fixed Issue 2421226 - Advanced UI allows network_operator to change routing properties.**

Users with *network_operator* role have extra privileges and they are able to change routing rules, for example.

- **Fixed Issue 2422111 - NSX-T Integrated with vIDM shows SQL string in username field.**

After integrating NSX-T with vIDM, the following message appears on screen when attempting to log into NSX-T using vIDM as authentication mechanism:

In place of a real username, the listed user is: x' and 1=2 union select '202cb962ac59075b964b07152d234b70','1

- **Fixed Issue 2424847 - A bond joining a LACP state machine might cause drops on slaves that were already up and running on the very same bond.**
You can have problems with traffic loss, BFD flaps, split brain in HA.
- **Fixed Issue 2426486 - PSOD on multiple CPU socket ESXi host having ENS enabled N-VDS with High latency sensitivity VM.**
ESXi host running NSX-T 2.5.0 crash in the following cases:
 - Enhance Network mode enabled on host switch.
 - More than 1 LCore are configured for the host switch.
 - HIGH latency sensitivity VM has NUMA affinity to a NUMA node.
 - No LCore configured at the NUMA node.
- **Fixed Issue 2429931 - vsip kernel module crashes on hosts with PCPU > 255.**
Hosts may experience PSOD if they have PCPU > 255.
- **Fixed Issue 2430585 - SPF did not handle the multiple uplink scenario correctly.**
Packet drops seen when using E-W Service Insertion on hosts with multiple uplinks.
- **Fixed Issue 2431227 - Escape characters not handled for JSON string in some fields.**
If there are some special characters in some fields (e.g., router name), the JSON strings are not handled correctly.
- **Fixed Issue 2434700 - Host DFW / VSIP kernel module out of memory when generating large amount of logs.**
Host DFW / VSIP kernel module out of memory when generating large amount of logs. This can impact traffic and debug commands can fail.
- **Fixed Issue 2435321 - Error messages not appearing in LDAPS UI.**
Help text is missing for some error messages.
- **Fixed Issue 2442095 - Cannot configure in-band management on Mellanox NIC on NSX Edge.**
You see an error while configuring in-band management on Mellanox NIC.
- **Fixed Issue 2442676 - Logical port mirror with egress direction may corrupt the MAC table if the mirror destination VM collector is located on the host.**
Traffic between VMs located on different hosts may have intermittent communication if the MAC table was corrupted.
- **Fixed Issue 2443118 - In NSX Manager, a CLI command fails intermittently.**
In NSX Manager, "get certificate cluster" CLI command fails intermittently with "An internal error occurred".
- **Fixed Issue 2446143 - Edge Transport Node configuration Edit workflow from NSX Manager does not work after upgrading to 2.5.**

After upgrading to 2.5, editing the Edge Transport Node configuration fails. For example, it's not possible to change the uplink profile, or the TEP pool / addressing.

- **Fixed Issue 2446618 - Opsagent crashed because VMs were migrated using storage vMotion.**
Opsagent crashed because VMs were migrated using storage vMotion.
- **Fixed Issue 2447899 - One of the Controller nodes may have conflicting or missing TN data, LSP-ip, IPV6 DAD data and CCP computed ServicePath data.**
The data plane may not follow the rules established by the management plane if a controller is disconnected from its database for a long time and then reconnected. Some VTEP/MAC information may be missing on hosts.
- **Fixed Issue 2449425 - Alarm is raised even if password-expiration is disabled for a user.**
Alarm is raised even if password-expiration is disabled for a user.
- **Fixed Issue 2450972 - Applying a TNP with migration on a host with a realized switch that has a portgroup with different teaming policies and no active uplinks causes an error.**
TNP application on the cluster fails until the user updates the configuration in VC to either set the portgroup to have at least one active uplink or to configure the portgroup and switch teaming policies to be the same.
- **Fixed Issue 2455489 - Restarting of Opsagent causes errors.**
You find that the hyperbus status is wrong when Opsagent is upgraded or restarted then removed and Transport Node added back to Transport Zone.
- **Fixed Issue 2425477 - AD Sync does not complete.**
You see one or more of these errors resulting in AD AD sync to not complete:
 - AD Group realization error.
 - Unable to realized intent.
 - Request has unknown parameters.
- **Fixed Issue 2424720 - Cisco VNIC devices can have UCS queue configuration that might not be supported by NSX Edge.**
Cisco VNIC devices can have UCS queue configuration that might not be supported by NSX Edge, causing the datapath to fail and the specific NSX Edge to become unusable.
- **Fixed Issue 2412487 - IXGBE DPDK driver takes a long time to detect link status, causing other protocols to failover or timeout before the link status notification can be seen.**
On a link-down on a bond slave, if the link status is slow, BFD or LACP might time out even before the link appears to be down. You may see a small traffic-loss period during failover due to unsynchronized network configuration with the actual physical network state.
- **Fixed Issue 2474534 - Duplicate IP detection not working correctly with manual bindings.**
Duplicate IP detection should occur between manual bindings, discovered bindings, and the mix of these two types. Duplicate IP detection does not work in the mixed case with

manual and discovered bindings.

- **Fixed Issue 2412696 - HTTP service not restarting after failure.**

The reverse-proxy (HTTP) service fails and does not automatically restart. You need to manually restart it.

- **Fixed Issue 2450768 - Unable to prioritize rules and sections from UI.**

Drag and drop to change the order for sections and rules is not working in some cases. As a result, you are unable to change the priorities of rules and sections from the UI.

- **Fixed Issue 2429162: Host loses network connectivity after the NSX CLI command "del nsx" is run.**

Note: For uninstalling NSX-T Data Center from a host, follow the correct and recommended procedure in the *NSX-T Data Center Installation Guide*. Do not use the "del nsx" command.

Known Issues

The known issues are grouped as follows.

- [General Known Issues](#)
- [Installation Known Issues](#)
- [Upgrade Known Issues](#)
- [NSX Manager Known Issues](#)
- [NSX Edge Known Issues](#)
- [Security Services Known Issues](#)
- [NSX Intelligence Known Issues](#)
- [NSX Cloud Known Issues](#)

General Known Issues

- **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.

Retry after thirty minutes. As an alternative, make the following API call to update the cache entry of datastore:

```
https://<nsx-manager>/api/v1/fabric/compute-collections/<CC Ext ID>/storage-resources?  
uniform_cluster_access=true&source=realtime
```

where *<nsx-manager>* 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 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 67835 [Bare Metal Server supports OVS bonding for Transport Node configuration in NSX-T](#).

- **Issue 2390624 - Anti-affinity rule prevents service VM from vMotion when host is in maintenance mode.**

If a service VM is deployed in a cluster with exactly two hosts, the HA pair with anti-affinity rule will prevent the VMs from vMotioning to the other host during any maintenance mode tasks. This may prevent the host from entering Maintenance Mode automatically.

Workaround: Power off the service VM on the host before the Maintenance Mode task is started on vCenter.

- **Issue 2389993 - Route map removed after redistribution rule is modified using the Policy page or API.**

If there is a route-map added using management plane UI/API in Redistribution Rule, it will get removed If you modify the same Redistribution Rule from Simplified (Policy) UI/API.

Workaround: You can restore the route map by returning the management plane interface or API to re-add it to the same rule. If you wish to include a route map in a redistribution rule, it is recommended you always use the management plane interface or API to create and modify it.

- **Issue 2275388 - Loopback interface/connected interface routes could get redistributed before filters are added to deny the routes.**

Unnecessary route updates could cause sub-optimal routing on traffic for a few seconds.

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 2329273 - No connectivity between VLANs bridged to the same segment by the same edge node.**

Bridging a segment twice on the same edge node is not supported. However, it is possible to bridge two VLANs to the same segment on two different edge nodes.

Workaround: None

- **Issue 2355113 - Unable to install NSX Tools on RedHat and CentOS Workload VMs with accelerated networking enabled in Microsoft Azure.**

In Microsoft Azure when accelerated networking is enabled on RedHat (7.4 or later) or CentOS (7.4 or later) based OS and with NSX Agent installed, the ethernet interface does not obtain an IP address.

Workaround: After booting up RedHat or CentOS based VM in Microsoft Azure, install the latest Linux Integration Services driver available at <https://www.microsoft.com/en-us/download/details.aspx?id=55106> before installing NSX tools.

- **Issue 2370555 - User can delete certain objects in the Advanced interface, but deletions are not reflected in the Simplified interface.**

Specifically, groups added as part of a distributed firewall exclude list can be deleted in the Advanced interface Distributed Firewall Exclusion List settings. This leads to inconsistent behavior in the interface.

Workaround: Use the following procedure to resolve this issue:

1. Add an object to an exclusion list in the Simplified interface.
2. Verify that it appears displayed in the Distributed Firewall exclusion list in the Advanced interface.
3. Delete the object from the Distributed Firewall exclusion list in the Advanced interface.
4. Return to the Simplified interface and add a second object to the exclusion list and apply it.
5. Verify that the new object appears in the Advanced interface.

- **Issue 2470210 - DFW local address set not updated on the VNIC after Storage vMotion of a DFW protected Virtual Machine.**

During a storage vMotion, a race condition is triggered where the cfgAgent is observing two filters with the same Virtual Interface and Logical Switch Port for a brief period of time, resulting in an incorrect address set update on the VNIC resulting in a traffic drop.

Workaround: None.

- **Issue 2498350 - Gateway firewall rules are not applied in some instances causing traffic to hit the default drop rule.**

Traffic is dropped due to hitting the default drop rule.

Workaround: None.

- **Issue 2509879 - Reduce pressure on activity framework by moving Application Initialization operations away from using activity framework.**

Host to NSX Manager connectivity may be impacted due to a buildup of activity in the activity framework table.

Workaround: None.

- **Issue 2512778 - Route advertisement fails from T1->T0 due to backed up activities in the activity framework queue.**

Processing of new activities fails when activity framework is backed up with activities.

Workaround: None.

- **Issue 2517232 - Inventory Objects not loading up in NSX Manager UI.**

When logging in to the NSX Manager UI, inventory objects does not show up as elastic search runs out of memory while trying to index huge objects while loading inventory.

Workaround: Reboot NSX Manager to recover from the error.

- **Issue 2523475 - PCF, Container APP not added dynamically to Security Group despite having matching tags.**

NSX objects like logical switches, logical ports or virtual machines are not dynamically added to NSGroup even though the membership criteria matches.

Workaround: None.

- **Issue 2543353 - NSX T0 edge calculates incorrect UDP checksum post-eSP encapsulation for IPsec tunneled traffic.**

Traffic is dropped due to bad checksum in UDP packet.

Workaround: None.

- **Issue 2547983 - NSGroups may not be cleaned up when deleted causing stale NSGroup entries in database.**

Due to a message size exception in database, NSGroup can get stale causing inconsistency in NSGroup membership.

Workaround: None.

- **Issue 2561740 - PAS Egress DFW rule not applied due to effective members not updated in NSGroup.**

Due to ConcurrentUpdateException a LogicalPort creation was not processed causing failure in updating the corresponding NSGroup.

Workaround: None.

- **Issue 2572505 - VM receives unintended traffic due to incorrect VLAN in the Geneve encapsulated packet.**

In an ENS Stack, Geneve UDP Source Port is incorrectly set to 0 and VLAN ID is not set for split packets, causing a failure to verify outer header, thus resulting in a packet drop.

Workaround: None.

- **Issue 2522782 - False-Positive alerts for NSX-T system Event when Service Router (SR) switches over from Down to Standby.**

Alarm is raised for SR in High Availability (HA) when its state is changed; however alarm is not cleared when its peer SR in HA becomes active.

Workaround: None.

- **Issue 2586606: Load balancer does not work when Source-IP persistence is**

configured on a large number of virtual servers.

When Source-IP persistence is configured on a large number of virtual servers on a load balancer, it consumes significant amount of memory and may lead to NSX Edge running out of memory. However the issue can reoccur with addition of more virtual servers.

Workaround: Disable source IP persistence or move VIPs with source IP persistence to different LB Services.

- **Issue 2621322: HTTP health check does not work when the HTTP content is in multi TCP segments.**

Load Balancer cannot check backend server status according to the HTTP content.

Workaround: None.

- **Issue 2491206: Load Balancer health check does not work well for body content match when there is chunk encoding in HTTP packet.**

There is CHUNK header in HTTP packet from backend server for health check. The pool member status cannot be up. The backend server is not down and available.

Workaround: None.

- **Issue 2730634: Post uniscale upgrade networking component page shows an "Index out of sync" error.**

Post uniscale upgrade networking component page shows an "Index out of sync" error.

Workaround: Log in to NSX Manager with admin credentials and run the "start search resync policy" command. It will take a few minutes to load the networking components.

Installation Known Issues

- **Issue 2481033: Updates to an ESXi Host Transport Node and Transport Node Profile attached to a host with powered on VMs fail with the error: "The host has powered on VMs which must be moved or powered off before transport node create/update/delete can continue".**

Updates to an ESXi host Transport Node (TN) will fail if it has VMK migration specified and there are any powered-on VMs on that ESXi host. Updates to a Transport Node Profile (TNP) attached to such TNs will fail regardless of the VMK migration setting on the TNP. This happens because powered-on VMs cause the migration validation to fail, preventing updates to the TN or TNP.

Workaround: See [Knowledge Base article 77123](#) for the workaround.

- **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.

Upgrade Known Issues

- **Issue 2475963 - NSX-T VIBs fail to install due to insufficient space.**

NSX-T VIBs fail to install due to insufficient space in bootbank on ESXi host, returning a BootBankInstaller.pyc: ERROR. Some ESXi images provided by third-party vendors may include VIBs which are not in use and can be relatively large in size. This can result in insufficient space in bootbank/alt-bootbank when installing/upgrading any VIBs.

Workaround: See Knowledge Base article 74864 [NSX-T VIBs fail to install, due to insufficient space in bootbank on ESXi host](#).

- **Issue 2400379 - Context Profile page shows unsupported APP_ID error message.**

The Context Profile page shows the following error message: "This context profile uses an unsupported APP_ID - [<APP_ID>]. Please delete this context profile manually after making sure it is not being used in any rule." This is caused by the post-upgrade presence of six deprecated APP_IDs (AD_BKUP, SKIP, AD_NSP, SAP, SUNRPC, SVN) that no longer work on the data path.

Workaround: After ensuring that they are no longer consumed, manually delete the six APP_ID context profiles.

- **Issue 2462079 - Some versions of ESXi hosts reboot during upgrade if there are stale DV filters present on the ESXi host.**

For hosts running ESXi 6.5-U2/U3 and/or 6.7-U1/U2, during maintenance mode upgrade to NSX-T 2.5.1, the host may reboot if stale DV filters are found to be present on the host after VMs are moved out.

Workaround: Upgrade to ESXi 6.7 U3 or ESXi 6.5 P04 prior to upgrading to NSX-T Data Center 2.5.1 if you want to avoid rebooting the host during the NSX-T Data Center upgrade. See [Knowledge Base article 76607](#) for details.

- **Issue 2441985 - Host Live upgrade from NSX-T Data Center 2.5.0 to NSX-T data Center 2.5.1 may fail in some cases.**

Host Live upgrade from NSX-T Data Center 2.5.0 to NSX-T Data Center 2.5.1 fails in some cases and you see the following error:

```
Unexpected error while upgrading upgrade unit: Install of offline bundle failed on host 34206ca2-67e1-4ab0-99aa-488c3beac5cb with error : [LiveInstallationError] Error in running ['/etc/init.d/nsx-datapath', 'start', 'upgrade']: Return code: 1 Output: ioctl failed: No such file or directory start upgrade begin Exception: Traceback (most recent call last): File "/etc/init.d/nsx-datapath", line 1394, in CheckAllFiltersCleared() File "/etc/init.d/nsx-datapath", line 413, in CheckAllFiltersCleared if FilterIsCleared(): File "/etc/init.d/nsx-datapath", line 393, in FilterIsCleared output = os.popen(cmd).read() File "/build/mts/release/bora-13885523/bora/build/esx/release/vmvisor/sys-boot/lib64/python3.5/os.py", line 1037, in popen File "/build/mts/release/bora-13885523/bora/build/esx/release/vmvisor/sys-boot/lib64/python3.5/subprocess.py", line 676, in __init__ File "/build/mts/release/bora-13885523/bora/build/esx/release/vmvisor/sys-boot/lib64/python3.5/subprocess.py", line 1228, in _execute_child OSError: [Errno 28] No space left on device It is not safe to continue. Please reboot the host immediately to discard the unfinished update. Please refer to the log file for more details..
```

Workaround: See [Knowledge Base article 76606](#) for details and workaround.

- **Issue 2477859 - In rare cases, NSX Manager upgrade may fail during the data migration task.**

When upgrading to NSX-T Data Center 2.5.1, in a very rare scenario where the deletion of a logical router in an earlier version did not process correctly, it is possible that NSX Manager upgrade may fail during the data migration task with this error: *NullPointerException*.

Workaround: Contact VMware support if you run into this issue.

- **Issue 2483552: After upgrading from 2.4.x to 2.5.x, "nsx-exporter" binary gets removed from the host**

After upgrading NSX-T Data Center from versions 2.4.x to versions 2.5.x, the binary of *nsx-exporter* (/opt/vmware/nsx-exporter) and *nsx-aggservice* (/opt/vmware/nsx-aggservice) get removed causing *nsx-exporter* to stop running.

Workaround: Reinstall the *nsx-exporter* and *nsx-aggregator* packages as follows:

1. Identify the RPM for *nsx-exporter* and *nsx-aggservice* using the command `'rpm -qa | grep nsx'`
2. Remove the RPM for *nsx-exporter* and *nsx-aggservice* using `'rpm -e nsx-exporter'` and `'rpm -e nsx-aggservice'`
3. Download the *nsx-lcp* tar file on the server and untar it.
4. Install the *nsx-aggservice* and *nsx-exporter* packages.

NSX Manager Known Issues

- **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.

Workaround: Use the CLI command `get service router config` command, which returns route-map configuration as a subset of entire output.

- **Issue 2378970 - Cluster-level Enable/Disable setting for distributed firewall incorrectly shown as Disabled.**

Cluster-level Enable/Disable setting for IDFW on Simplified UI may show as Disabled even though it is Enabled on the management plane. After upgrading from 2.4.x to 2.5, this inaccuracy will persist until explicitly changed.

Workaround: Manually modify the Enable/Disable setting for IDFW on Simplified UI to match the same on the management plane.

NSX Edge Known Issues

- **Issue 2283559 - `https://<nsx-manager>/api/v1/routing-table` and `https://<nsx-manager>/api/v1/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 2416130 - No ARP proxy when Centralized Service Port (CSP) is connected to DR's downlink**
No ARP proxy when Centralized Service Port (CSP) is connected to DR's downlink causing no traffic to pass.

Workaround: Place the static router (SR) on the same edge. Then hypervisor will not require ARP proxy.

Security Services Known Issues

- **Issue 2448006 - Querying a Firewall Section with inconsistencies in rule-mapping fails.**
Querying a Firewall Section with rule-mapping inconsistencies fails if you use the *GetSectionWithRules* API call. The UI is not impacted as it depends on *GetSection* and *GetRules* API calls.

Workaround: Fetch the Firewall Section using the APIs *GetSection* and *GetRules* or use the UI.

NSX Intelligence Known Issues

- **Issue 2362865 - Filter by Rule Name not available for default rule.**
Observed in the **Plan & Troubleshoot > Discover and Take Action** page and affects only rules created by connectivity strategy. This issue is caused by the absence of a default policy based on the connectivity strategy specified. A default rule may be created on the management plane, but with no corresponding default policy, the user cannot filter based on that default rule. (The filter for flows visualization uses the rule name to filter by flows that hit that rule.)

Workaround: Do not apply a rule name filter. Instead, check the Unprotected flag. This configuration will include flows hitting the default rule as well as any rule that has "any" source and "any" destination specified.

- **Issue 2368926 - Recommendations job fails if user reboots appliance while job is in progress.**
If you reboot the NSX Intelligence appliance while a recommendations job is in progress, the job goes to a failed state. You can start a recommendation job for a set of context VMs. The reboot deletes the context and the job fails as a result.

Workaround: After reboot, repeat the recommendations job for the same set of VMs.

- **Issue 2369802 - NSX Intelligence appliance backup excludes event datastore**

backup.

This functionality is not supported in NSX 2.5.

Workaround: None.

- **Issue 2389691 - Publish recommendation job fails with error "request payload size exceeds the permitted limit, max 2,000 objects are allowed per request."**

If you try to publish a single recommendation job that contains more than 2,000 objects, it will fail with error "request payload size exceeds the permitted limit, max 2,000 objects are allowed per request."

Workaround: Reduce the number of objects to fewer than 2,000 in the recommendation job and retry the publication.

- **Issue 2396630 - Delete transport node operation may fail when NSX intelligence appliance is deployed.**

If a transport node is being deleted while the NSX Intelligence appliance is being deployed, the deletion can fail because the transport node is referred by NSX-INTELLIGENCE-GROUP NSGroup. To delete a transport node, the force delete option is required when NSX Intelligence appliance is deployed.

Workaround: Use the force option to delete the transport node.

- **Issue 2393240 - Additional Flows are observed from VM to IP address.**

Additional flows from VM to IP-xxxx are seen. This is due to the configuration data (Groups, VMs and services) when the NSX Policy manager reaches the NSX Intelligence appliance after the flow is created. Therefore the (earlier) flow cannot be correlated with the configuration, because it is non-existent from the flow perspective. Since the flow cannot be normally correlated, it defaults to IP-xxxx for its VM during flow lookup. After the configuration is synchronized, the actual VM flow appears.

Workaround: Modify the time window to exclude the flow you do want to see.

- **Issue 2370660 - NSX Intelligence shows inconsistent data for specific VMs.**

This is likely caused by those VMs having the same IP address in the datacenter. This is not supported by NSX Intelligence in NSX-T 2.5.

Workaround: None. Avoid assigning the same IP address to two VMs in the datacenter.

- **Issue 2372657 - VM-GROUP relationship and GROUP-GROUP flow correlation temporarily display incorrectly.**

VM-GROUP relationship and GROUP-GROUP flow correlation temporarily display incorrectly if the NSX Intelligence appliance is deployed while there are ongoing flows in the datacenter. Specifically, the following elements may display incorrectly during this temporary period:

- VMs wrongly belong to Uncategorized group.
- VMs wrongly belong to Unknown group.
- Correlated flows between two groups can be shown wrongly.

These errors will self-correct after the NSX Intelligence appliance has been deployed longer than the user-selected visualization period.

Workaround: None. If the user moves out of the Visualization period during which the NSX Intelligence appliance was deployed, the issue will not appear.

- **Issue 2393142 - Logging in to NSX Manager with vIDM credentials may return a 403 unauthorized user error.**

This only affects users logging in as vIDM users, as opposed to a local user, on NSX Manager. vIDM login and integration are not supported in NSX-T 2.5 when interacting with the NSX Intelligence appliance.

Workaround: Log in as a local user by appending the NSX Manager IP/FQDN with the string 'login.jsp?local=true'.

- **Issue 2346545 - NSX Intelligence appliance: certificate replacement affects new flow information reporting.**

If you replace the principal identity certificate for the NSX Intelligence appliance with a self-signed certificate, processing of new flows is affected and the appliance will not show updated information that point forward.

Workaround: None.

- **Issue 2410224 - After completing NSX Intelligence appliance registration, refreshing view may return a 403 Forbidden error.**

After completing NSX Intelligence appliance registration, if you click **Refresh to View**, the system may return a 403 Forbidden error. This is a temporary condition caused by the time required for the NSX Intelligence appliance to access the interface.

Workaround: If you receive this error, wait a few moments and try again.

- **Issue 2436302 - After replacing the NSX-T unified appliance cluster certificate, NSX Intelligence cannot be accessed via API or the Manager interface.**

In the NSX-T Manager interface, go to the **Plan & Troubleshoot** tab and click **Discover & Take Action** or **Recommendations**. The interface will not load and will eventually return an error like: Failed to load requested application. Please try again or contact support if the problem persists.

Workaround: See [Knowledge Base article 76223](#) for more details and workaround.

- **Issue 2374229 - NSX Intelligence appliance runs out of disk space.**

The NSX Intelligence appliance has a default data retention period of 30 days. If the amount of flow data is larger than the anticipated amount within 30 days, the appliance might run out of disk space prematurely and become partially or completely non-operational.

Workaround: See [Knowledge Base article 76523](#) for more details and workaround.

- **Issue 2376389 - VMs are incorrectly marked as deleted in 'Last 24 hours' view on mid-scale setup.**

After a host is disconnected from compute managers, NSX Intelligence shows the previous VMs on the host as deleted, with new VMs in their place. This issue results from NSX Intelligence tracking inventory updates in the NSX database, and this behavior reflects how the inventory handles host disconnection from compute managers. This does not affect the total count of live VMs in NSX Intelligence, although you may see duplicate VMs in NSX

Intelligence.

Workaround: No action required. The duplicate VMs will stop appearing after approximately 24 hours.

- **Issue 2385599 - Groups of static IPs not supported in NSX-T Intelligence recommendations.**

VMs and workloads that are not recognized in the NSX-T inventory, if they have intranet IP addresses, may be still be subject to recommendation as a group of static IPs, including recommendation-define rules containing these groups. However, NSX Intelligence does not support such groups and as a result, visualization shows traffic sent to them as sent to "Unknown" instead of the recommended group.

Workaround: None. However, recommendation is functioning correctly. This is a display issue.

- **Issue 2407198 - VMs incorrectly appear in Uncategorized VMs group in NSX intelligence security posture.**

When ESXi hosts are disconnected from vCenter, VMs in those hosts can be shown in "Uncategorized VMs" group even if they belong to other groups. When the ESXi hosts reconnected with vCenter, the VMs will appear in their correct groups.

Workaround: Reconnect the hosts to vCenter.

- **Issue 2366599 - Rules for VMs with IPv6 addresses not enforced.**

If a VM uses an IPv6 address, but IPv6 snooping is not enabled for that VIF via the IP discovery profile, the IPv6 address is not populated in the rule for that VM in the data path. As a result, that rule is never enforced.

Workaround: Verify that IPv6 discovery profile is enabled at either the VIF or logical switch whenever IPv6 addresses are used.

- **Issue 2374231 - Port scan with nmap tool generates flow with service as UNKNOWN and port as 0.**

NSX Intelligence does not support source or destination port parsing for GRE, ESP, and SCTP protocol flows. NSX Intelligence provides full header parsing for TCP and UDP flows along with flow related statistics. For other supported protocols (such as GRE, ESP, and SCTP) NSX Intelligence can only provide IP information without protocol specific source or destination ports. For these protocols, the source or destination port will be zero.

Workaround: None.

- **Issue 2410096 - After rebooting the NSX Intelligence appliance, flows collected in the last 10 minutes prior to reboot may not be displayed.**

Caused by an indexing issue.

Workaround: None.

- **Issue 2357296 - Flows may not be reported to NSX Intelligence by some ESX hosts under certain scale and stress conditions.**

The NSX Intelligence interface may not show flows from certain VMs on certain hosts, and

fails to provide firewall rule recommendations for those VMs. As a result, firewall security could be compromised on some hosts. This is observed in deployments with vSphere versions below 6.7U2 and 6.5U3. The problem is identified as core ESX hypervisor VM filter creation and deletion out of order.

Workaround: Upgrade host to version vSphere 6.7U2 and above or vSphere 6.5U3 and above.

- **Issue 2456118 - Error accessing NSX Intelligence.**

When loading the "Plan & Troubleshoot" page in NSX-T Data Center, you may see one or more of the following:

- The Application server fails to fulfill request.
- The NSX-T Intelligence agent rejects any admin user attempts.
- You get the error: *Failed to load requested application. Please try refreshing the browser or contact support if the problem persists.*

Workaround: See [Knowledge Base 76223](#) for more details.

- **Issue 2508429: Only Base64-encoded certificate files are supported in NSX Intelligence 1.0.1. Extra attributes that are part of a PEM-encoded certificate are not allowed.**

"Bag attributes" in certificate files are not accepted in NSX Intelligence 1.0.1. Only Base64 encoding is supported in NSX Intelligence 1.0.1.

Workaround: See Knowledge Base article <https://kb.vmware.com/s/article/78048> for more details and workaround.

NSX Cloud Known Issues

- **Issue 2289150 - PCM calls to AWS start to fail.**

If a user updates 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 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.