



VMware NSX-T Data Center 2.5 Release Notes

VMware NSX-T Data Center 2.5 | 19 September 2019 | Build 14663974

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

What's New

NSX-T Data Center 2.5 provides a variety of new features to provide new functionality for virtualized networking and security for private, public, and hybrid clouds. Highlights include enhancements to intent-based networking user interface, context-aware firewall, guest and network introspection features, IPv6 support, highly-available cluster management, profile-based NSX installation for vSphere compute clusters, and enhancements to migration coordinator for migrating from NSX Data Center for vSphere to NSX-T Data Center.

NSX Intelligence

NSX-T Data Center 2.5 introduces NSX Intelligence v1.0, a new NSX analytics component. NSX Intelligence provides a user interface via a single management pane within NSX Manager, and provides the following features:

- Close to real-time flow information for workloads in your environment.
- NSX Intelligence correlates live or historic flows, user configurations, and workload inventory.
- Ability to view past information about flows, user configurations, and workload inventory.
- Automated micro-segmentation planning by recommending firewall rules, groups, and

services.

Container API Support

New API support is available for container inventory. See the API documentation.

L2 Networking

- **Enhancements for the Edge Bridge** - The Edge bridge now allows attaching the same segment to multiple bridge profiles, thus providing the ability to bridge a segment multiple times to VLANs in the physical infrastructure. This new functionality supersedes and deprecates the original ESXi bridge in previous versions of NSX-T Data Center. **Caution:** Use this feature at your risk. It introduces the risk of creating a bridging loop by bridging the same segment twice to the same L2 domain in the physical network. There is no loop mitigation mechanism.
- **MTU/VLAN Health Check** - From an operations point of view, network connectivity issues caused by configuration errors are often difficult to identify. Common scenarios include ones wherein virtual network admins using NSX Manager while physical network admins take management ownership of physical network switches.
 - **VLAN Health Check** - Checks whether N-VDS VLAN settings match trunk port configuration on the adjacent physical switch ports.
 - **MTU Health Check** - Checks whether the physical access switch port MTU setting based on per VLAN matches the N-VDS MTU setting.
- **Guest Inter-VLAN Tagging** - The Enhanced Datapath N-VDS enables users to map guest VLAN Tag to a segment. This capability overcomes the limitation of 10 vNICs per VM and allows guest VLAN tagged traffic (mapped to different segments) to be routed by the NSX infrastructure.

L3 Networking

- **Tier-1 Placement Inside Edge Cluster Based on Failure Domain**- Enables NSX-T to automatically place Tier-1 gateways based on failure domains defined by the user. This increases the reliability of Tier-1 gateways across availability zones, racks, or hosts, even when using automatic Tier-1 gateway placement.
- **Asymmetric Load Sharing After Router Failure in ECMP Topology**- On active/active Tier-0 gateway when one faulty service router was going down another router was taking over the faulty router traffic doubling the traffic going through the service router. After 30 minutes of a router failure, the faulty router IP address is removed from the list of next-hops avoiding the additional traffic to a specific router .
- **Get BGP Advertised and Received Routes Per Peer through API**- Simplifies BGP operations by avoiding CLI usage to verify the routes received and sent to BGP peers.
- **BGP Large Community Support** - Offers the option to use communities in conjunction with 4-byte ASN as defined in RFC8092.
- **BGP Graceful Restart Helper Mode Option Per Peer** - Offers the option for Tier-0 gateway to help maintain router for northbound physical routers with redundant control plane without compromising on the failover time across Tier-0 routers.
- **Bulk API to Create Multiple NAT Rules** - Enhances the existing NAT API to bundle the creation of a large number of NAT rules into a single API call.

Edge Platform

- **Support Mellanox ConnectX-4 and ConnectX-4 LX on Bare Metal Edge Node**- Bare Metal Edge nodes now support Mellanox ConnectX-4 and ConnectX-4 LX physical NICs in 10/25/40/50/100 Gbps.
- **Bare Metal Edge PNIC Management** - Provides the option to select the physical NICs to be used as dataplane NICs (fastpath). It also increases the number of physical NICs supported on the Bare Metal Edge node from 8 to 16 PNICs.

Enhanced IPv6 Support

NSX-T 2.5 continues to enhance the IPv6 routing/forwarding feature-set. This includes the support for:

- IPv6 SLAAC (Stateless Address Autoconfiguration), automatically providing IPv6 addresses to virtual machines.
- IPv6 Router Advertisement, NSX-T gateway provides IPv6 parameters through Router Advertisement.
- IPv6 DAD, NSX-T gateways detects duplicate IPv6 address allocation.

Firewall Improvements

Layer-7 AppID Support

NSX-T 2.5 adds more Layer-7 capabilities for distributed and gateway firewall. This includes the support for:

- Layer-7 AppID support for distributed firewall on KVM.
- Layer-7 AppID support for gateway firewall.
- Multiple Layer-7 AppID configuration in a single firewall rule.

FQDN/URL Filtering Enhancements

NSX-T 2.5 has minor enhancements to FQDN filtering support, including:

- Configuring TTL timers for DNS entries.
- Support for workloads running on KVM hypervisor.

Firewall Operations have been enhanced with the following features:

- **Autosave Configuration & Rollback Feature** - The system creates a copy of the configuration when published. This configuration can be re-deployed to rollback to an existing state.
- **Manual Drafts** - Users can now save drafts of their rules before they publish those rulesets for enforcement. Users can stage the rules in manual drafts. The system allows you to have multiple users work on the same draft with a locking mechanism to disable overriding of rules from different users.
- **Session Timers** - Users can configure session timers for TCP, UDP and ICMP sessions.
- **Flood Protection** - Both distributed firewall and gateway firewall can have SynFlood protection. Users can provide thresholds to alert, log and drop traffic to make it custom

workflows.

- **System auto-generates two groups** when NSX LoadBalancer is created and virtual servers are deployed. One group contains the server pool while the other group contains virtual server IP. These groups can be used on distributed firewall or gateway firewall to allow or deny traffic by firewall admins. These groups track the NSX load balancer config changes.
- **The number of IP addresses** detected per VM - vNIC has been increased from 128 to 256 IP addresses.

Identity Firewall

- With NSX-T 2.5, we support Active Directory Servers deployed on Windows 2016.
- We support the Identity Firewall for Windows Server workloads without Terminal Services enabled. This will allow customers to strictly control the lateral movement of administrators from one server to another.

Service Insertion

- **Packet Copy Support** - In addition to redirecting traffic through a service, NSX-T now supports the Network Monitoring use case, in which a copy of packets is forwarded to a partner Service Virtual Machine (SVM), allowing inspection, monitoring or collection of statistics while the original packet does not pass through the network monitoring service.
- **Automatic Host-based Partner SVM Deployment** - As of NSX-T 2.5, two modes of Partner SVM deployment are supported; clustered deployment in which Service Virtual Machines are deployed on a dedicated vSphere (Service) Cluster and Host-Based in which one Service Virtual Machine per service is deployed on each Compute Host in a particular cluster. In this mode, when a new compute host is added to a cluster, the appropriate SVMs are automatically deployed.
- **Notification Support for North-South Service Insertion**- NSX-T 2.4 introduced the notification framework for East-West Service Insertion, allowing partner services to automatically receive notifications upon relevant changes such as dynamic group updates. With NSX-T 2.5, this notification framework has also been extended to N-S Service Insertion. Partners can leverage this mechanism in order to allow customers to use dynamic NSX groups (i.e. based on Tags, OS, VM Name) in the partner policy.
- **Additional Troubleshooting and Visualization Features** - With NSX-T 2.5, several serviceability enhancements have been made to allow for better troubleshooting of Service Insertion related issues. This includes the ability to verify the runtime status of a Service Instance, the ability to fetch available Service Paths through the API and the inclusion of Service Insertion related logs in the support bundle.

Endpoint Protection (Guest Introspection)

- **Linux Support** - Support for Linux-based operating systems with Endpoint Protection. Please see the NSX-T Administration Guide for supported Linux operating systems for Guest Introspection.
- **Endpoint Protection Dashboard** - Endpoint Protection dashboard for visibility and monitoring the configuration status of protected and unprotected VMs, issues with Host agent and service VMs, and VMs configured with the file introspection driver that was installed as part of the VMware Tools installation.

- **Monitoring Dashboard** - To monitor the partner service deployment status across clusters in the system .

Load Balancing

- **API to Retrieve the Status on Edge Capacity for Load Balancers**- New API calls have been added to allow the admin to monitor the Edge capacity in terms of load balancing instances.
- **Intelligent Selection of Health Check IP Address**- When SNAT IP-list is configured, the first IP address in the list is going to be used for health monitoring instead of the uplink IP address of a Tier-1 Gateway. The IP address can be the same as the Virtual Server IP address. This enhancement allows the load balancer to use a single IP address for both source-nat and health monitoring.
- **Load Balancer Logging Enhancement** - With this enhancement, the load balancer can generate a rich log message per Virtual Server for monitoring. For example, the Virtual Server access log includes not only the client IP address but also a pool member IP address.
- **Persistent Enhancement in LB Rules** - A new action called "Persist" is introduced in LB rules. The Persist action enables the load balancer to provide application persistency based on a cookie set by a pool member.
- **LB Fits** - A small LB instance can fit into a small Edge VM. A medium LB instance can fit into a medium Edge VM. Previously, the small Edge VM did not support load balancing services because the size of an Edge VM had to be bigger than the size of an LB instance.
- **VS/Pool/Member Statistics** - All LB related statistics are available in simplified interface. Previously, the information was only available in Advanced Networking and Security interface.
- **ECC (Elliptical Curve Certificate) Support for SSL Termination**- EC certificates can be used for increased SSL performance.
- **FIPS Knob** - There is a global setting via API for FIPS compliance for load balancers. By default, the setting is turned off to improve performance.

VPN

- **IPsec VPN Support on Tier-1 Gateway** - IPsec VPN can be deployed and terminated on Tier-1 gateway for better tenant isolation and scalability. Previously, it was supported on only Tier-0 gateway.
- **VLAN Support for Layer-2 VPN on NSX-managed Edge**- With this enhancement, VLAN-backed segments can be extended. Previously, only logical segments were supported for Layer-2 extension. This includes VLAN Trunking support enabling multiple VLANs to be extended on one Edge Interface and Layer-2 VPN session.
- **TCP MSS Clamping for IPsec VPN** - TCP MSS Clamping allows the admin to enforce the MSS value of all TCP connections to avoid packet fragmentation.
- **ECC (Elliptical Curve Certificate) Support for IPsec VPN**- The EC certificate is required to enable various IPsec compliance suites, such as CNSA, UK Prime, etc.
- **Easy Button for Compliance Suite Configuration**- CNSA, Suite-B-GCM, Suite-B-GMAC, Prime, Foundation, and FIPS can be configured with a single click in the UI or a single API call.

Automation, OpenStack and other CMP

- **Expanded OpenStack Release Support** - Now includes the Stein and Rocky releases.
 - **OpenStack Neutron Plugin supporting Policy API** - In addition to existing plugin supporting management API, we now offer an OpenStack Neutron plugin consuming the new NSX-T Policy API. This plugin supports IPv6 for Layer-2, L3, firewall and SLAAC.
 - **OpenStack Neutron Router Optimization** - The plugin now optimizes the OpenStack Neutron Router by managing the creation/deletion of the service router dynamically. This allows a customer to have only a distributed router when no services are configured and one as soon as the services are added, all managed by the plugin.
 - **OpenStack Neutron Plugin Layer-2 Bridge** - The Layer-2 bridge configured from OpenStack is now configured on the Edge Cluster and not on the ESXi cluster.
 - **OpenStack Octavia Support** - In addition to LBaaSv2, the OpenStack Neutron Plugin supports Octavia as a way to support Load Balancing.
- For more details please see the VMware NSX-T Data Center 2.5 Plugin for OpenStack Neutron Release Notes.

NSX Cloud

- **Addition of a New Mode of Operation** - NSX Cloud will now have two modes of operation, this officially makes NSX Cloud the only Hybrid Cloud solution in the market to support agented and agentless modes of operation.
 - **NSX Enforced Mode (Agented)** - Provides a "Consistent" policy framework between on-premises and any public cloud. NSX Policy enforcement is done with NSX tools which are installed in every workload. This provides VM level granularity and all tagged VMs will be managed by NSX. This mode will overcome the differences/limitations of individual public cloud providers and provide a consistent policy framework between on-premises and public cloud workload.
 - **Native Cloud Enforced Mode (Agentless)** - Provides a "Common" policy framework between on-premises and any public cloud. This mode does not require the installation of NSX tools in the workloads. NSX Security Policies are converted into the Native Cloud providers security constructs. Hence, all the scale and feature limitations of the chosen public cloud are applicable. The granularity of control is at the VPC/VPNET level and every workload inside a managed VPC/VPNET will be managed by NSX unless it is whitelisted.

Both modes will provide Dynamic Group membership and a rich set of abstractions for nsx group membership criteria.
- **Support for Visibility and Security of Public Cloud Native Services from NSX Cloud** - From this release, it will be possible to program the security groups of Native SaaS services in Azure and AWS which have a local VPC/VPNET endpoint and a security group associated with it. The primary idea is to discover and secure cloud native service endpoints with user-specified rules on NSX policy. The following services will be supported in AWS (ELB, RDS & DynamoDB) and Azure (Azure Storage, Azure LB, Azure SQL Server & CosmosDB) in this release. Future NSX-T releases will add more support for more services.
- **New OS support:**
 - Support for Windows Server 2019
 - Windows 10 v1809

- Support for Ubuntu 18.04
- **Enhanced Quarantine Policy and VM White-listing**- Starting with NSX 2.5, NSX Cloud allows users to whitelist VMs from the CSM interface. Once whitelisted, cloud security groups of such VMs are not be managed by NSX, and users can put the VMs in whatever cloud security groups they want.
- **Enhanced Error Reporting on CSM Interface** - Enables quicker troubleshooting.

Operations

- **Support of vSphere HA for the NSX Manager(s)**- The NSX management cluster can now be protected by vSphere HA. This allows one node of the NSX management cluster to be recovered if the host running it fails. It also allows for the entire NSX management cluster to be recovered to an alternate site if there is a site-level failure. Please see the NSX-T Installation Guide for details on supported scenarios.
- **Capacity Dashboard Improvements** - New and improved metrics to the capacity dashboard show the number of objects a customer has configured relative to the maximum supported in the product. For a complete list of configuration maximums for NSX-T Data Center, see the VMware Configuration Maximums Tool.
- **Support for vSphere Lockdown Mode** - Enable more deployment options for customers by providing the ability to install, upgrade and operate NSX-T in a vSphere lockdown mode environment.
- **Logging Enhancement** - Reduce service impact during troubleshooting by enabling dynamic change of log levels via the NSX command line interface for NSX user space agents.
- **SNMPv3 Support** - Enhanced security compliance by adding support for configuring SNMPv3 for NSX Edge and Manager appliance.
- **New Traceflow Capability for Troubleshooting VM Address Resolution Issues**- Added support for injecting ARP/NDP packets via Traceflow to detect connectivity issues while doing address resolution for an IP destination.
- **Upgrade Order Change** - When upgrading to NSX-T 2.5, the new upgrade order is Edge-component upgrade before Host component upgrade. This enhancement provides significant benefits when upgrading the cloud infrastructure by allowing optimizations to reduce the overall maintenance window.
- **Log Insight Content Pack Enhancement**- Added support for out-of-box log alerts with the new NSX-T Content Pack compatible with NSX-T 2.5.

Platform Security

- **FIPS** - Users can now generate FIPS compliance reports. including the ability to configure and manage their NSX deployments in FIPS-compliant mode. Cryptographic modules are validated per the FIPS standards, offering security assurance for customers who want to be compliant per federal regulations or operate NSX in a secure manner that adheres to prescribed FIPS standards. With noted exceptions, all cryptographic modules in NSX-T 2.5 are FIPS certified. To view granted certifications for FIPS-validated modules, see <https://www.vmware.com/security/certifications/fips.html>.
- **Enhancements to Password Management** - Users can now extend the password expiry duration (day-count) since the last password change even after upgrade. Thirty-day expiry warnings and password expiry notifications now appear in the interface, CLI, and syslogs.

Support for Single Cluster Design

Support of single cluster designs with fully collapsed Edge+Management+Compute VMs, powered by a single N-VDS, in a cluster with a minimum of four hosts. The typical reference designs for VxRail and other cloud provider host solution prescribe 4x10G pNICs with two host switches. One switch is dedicated to Edge+Management (VDS), whereas the other one is dedicated to compute VMs (N-VDS). Two host-switches effectively separate the management traffic from the compute traffic. However, with the trending economics of 10 and 25G, many small data center and cloud provider customers are standardizing on two pNICs host. Using this form factor, small data centers and cloud provider customers can build an NSX-T based solution with single N-VDS, powering all the components with two pNICs.

NSX Data Center for vSphere to NSX-T Data Center Migration

- **Migration Coordinator Enhancements** - The Migration Coordinator has several usability enhancements that improve the workflow of the process required to migrate from NSX Data Center for vSphere to NSX-T Data Center, including improvements to providing user feedback during the migration.

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 with NSX-T Data Center 2.5, the NSX Messaging channel TCP port from all Transport and Edge nodes to NSX Managers has changed to TCP port 1234 from port 5671. With this change, make sure all NSX-T Transport and Edge nodes can communicate on both TCP ports 1234 to NSX Managers and TCP port 1235 to NSX Controllers before you upgrade to NSX-T Data Center 2.5. Also make sure to keep port 5671 open during the upgrade process.

L2 Networking

As a result of the enhancements for Layer-2 bridges, the ESXi bridge is deprecated. NSX-T was initially introduced with the capability of dedicating an ESXi host as a bridge to extend an overlay segment to a VLAN. This model is deprecated as of this release because the new Edge bridge supersedes it in term of features, does not require a dedicated ESXi host, and benefits from the optimized data path of the Edge node. See "What's New" for more information.

API Deprecations and Behavior Changes

Transport Node Template APIs are deprecated in this release. It is recommended that you use Transport Node Profiles APIs instead. See the [API Guide](#) for the list of deprecated types and methods.

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 September 2019. First edition.

23 September 2019. Added Known Issues 2424818 and 2419246. Added Resolved Issues 2364756, 2406018, and 2383328.

24 September 2019. Updated What's New items.

03 October 2019. Added Resolved Issue 2313673.

12 November 2019. Added Known Issues 2362688 and 2436302. Corrected Issue 2282798 by moving it to Resolved.

17 December 2019. Added Known Issue 2444170.

14 January 2020. Added Resolved Issue 2399994.

18 February 2020. Updated Known Issue 2436302 with link to KB article.

14 May 2020. Added Known Issue 2467479.

25 September 2020. Added Known Issue 2586606.

15 March 2021. Added Known Issue 2730634.

Resolved Issues

- **Fixed 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.
- **Fixed 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.
- **Fixed 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.

- **Fixed 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.
- **Fixed 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.
- **Fixed Issue 2277742 - Invoking PUT https://<nsx-manager>/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://<nsx-manager>/api/v1/configs/management cannot be invoked if a FQDN is configured.
- **Fixed 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.
- **Fixed 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.
- **Fixed 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.
- **Fixed 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.
- **Fixed Issue 2383867 - Log bundle collection fails for one of the Management Plane nodes.**
Log collection process experiences a failure when copying support bundle to remote server.
- **Fixed 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.

- **Fixed Issue 2410818 - After upgrading to 2.4.2, virtual servers created in NSX-T 2.3.x may stop working after more virtual servers are created.**
 In some deployments, virtual servers created in version 2.3.x stop working after upgrading to version 2.4.2 and after more virtual servers were created.
- **Fixed 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"
- **Fixed 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.
- **Fixed 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.
- **Fixed 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.
- **Fixed 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.
- **Fixed 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.
- **Fixed 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.
- **Fixed 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.

- **Fixed 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."
- **Fixed Issue 2364756 - Profile realization fails due to duplicate priority.**
 On scale setups, when user associated vRNI with NSX IPFIX, the profile would not realize on the management plane and would through realization errors.
- **Fixed 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.
- **Fixed issue 2307551 - NSX-T Host may lose management network connectivity when migrating all pNICs to N-VDS.**
 The issue results from the host migration retry removing all pNICs in the N-VDS that has vmk0 configured. The first host migration migrated all pNICs and vmk0 into the N-VDS but failed afterward. When you retry migration, all pNICs are removed from the N-VDS. As a result, users cannot access the host through the network; all VMs in the host also lose network connectivity, rendering their services unreachable.
- **Fixed Issue 2369792 - CBM process repeatedly crashes due to CBM process memory bloat.**
 CSM and CBM processes on Cloud Service Manager appliance fail database compaction. As a result, CBM process memory bloat causes the CBM process to repeatedly crash.
- **Fixed Issue 2361892 - The NSX Edge appliance experiences a memory leak, leading to process crash/restart.**
 Over an extended period of time, the NSX Edge appliance may experience a memory leak due to repeated rule lookup, leading to process crash/restarts. A memory leak was detected every time a rule lookup was executed. When the flow cache is cleared, the VIF interface is not removed, causing a building in memory.
- **Fixed Issue 2364529 - Load balancer memory leak after reconfiguration.**
 NSX Load balancer might leak memory upon consecutive/repetitive configuration events, resulting in nginx process core dump.
- **Fixed Issue 2378876 - PSOD on ESXi hosts with errors: "Usage error in dlmalloc" and "PF Exception 14 in world 3916803:VSIP PF Purg IP".**
 ESXi crashed (PSOD) after running traffic for a few days. No other symptoms were observed prior to the crash. Issue was ultimately identified in ALG traffic (FTP, Sunrpc, Oracle, Dcerpc, tftp) where nonatomized increment counter led to race conditions, corrupting the ALG tree structure.
- **Fixed Issue 2384922 - BGPD consumes 100% CPU usage on Edge node.**
 BGPD process on NSX-T Edge may consume 100% CPU when it has several open sessions with VTYSH.

- **Fixed Issue 2386738 - NAT rules ignored on traffic over LINKED port.**
NAT services are not enabled on LINKED router port type connecting Tier-0 and Tier-1 logical routers.
- **Fixed Issue 2363618 - VMware Identity Manager users unable to access Policy pages in NSX Manager dashboard.**
Users with roles assigned to Group permissions in VMware Identity Manager are unable to access Policy pages in the NSX Manager dashboard. Permissions from group assignment are ignored.
- **Fixed Issue 2298274 - Policy Group can be created/updated with invalid or partial domain name through REST API.**
The interface permitted creation of groups with identity expressions containing invalid Active Directory group or individual groups members for a single valid content. However, each member is valid only if it has exactly one LDAP group associated to the domain name. As a result, such groups created in a previous version of NSX-T, this error will not be flagged in the upgrade process, allowing the invalid groups to persist in subsequent releases. Fixed in 2.5.
- **Fixed Issue 2317147 - Users cannot see effective VMs for a group whose membership is based on IP or MAC addresses.**
If a user creates a group with only IP or MAC addresses in the group, no VMs are listed when effective membership for that group is called from the API. There is no functional impact. Policy properly creates an NSGroup on the management plane, and the list of IP and MAC addresses is directly sent to central control plane.
- **Fixed Issue 2327201 - Updates of VMs on KVM hypervisors not immediately synchronized.**
VM updates on KVM hypervisors may take a couple of hours to synchronize on NSX-T. As a result, new VMs created on KVM hypervisors cannot be added to NSGroups, no firewall rules can be applied on those VMs, upgrade of KVM hypervisor is not possible because the VM power status is not updated.
- **Fixed Issue 2329443 - Control cluster is not getting initialized due to forcesync timeout.**
The control cluster is not initializing due to a forcesync timeout when the IPV4 range in Ipset starts at 0.0.0.0, for example 0.0.0.0-1.1.1.20. This is caused by an issue in the IPSetFullSyncMessageProvider which becomes stuck in an infinite loop. Since the central control plane is not getting initialized, users can't deploy new workloads.
- **Fixed Issue 2337839 - NSX-T backup widgets display incorrect field names.**
Specifically, the NSX-T backup widgets are not displaying the correct number of backup errors. As a result, the customer needs to review the NSX Manager backup tab to see the accurate count of backup errors.
- **Fixed Issue 2341552 - Edge fails to boot when system has too many supported NICs present.**
No datapath service or connectivity can be seen, the datapath service is down, and Edge node is in a degraded state. This results in partial or total connectivity loss if the edge is required.

- **Fixed Issue 2390374 - NSX Manager becomes very slow or unresponsive, and logs show many corfu exceptions.**

NSX may also fail to start up. The corfu exceptions indicate that the scale of Active Directory members is too large, and above the tested limits.

- **Fixed Issue 2371150 - Unable to configure Layer-7 firewall rules on Bare Metal Edge nodes.**

Layer-7 firewall rules on Bare Metal Edge nodes are not supported in NSX-T 2.5. There is an internal command that enables this support but this is only available for proofs of concept.

- **Fixed Issue 2361238 - Downlink router doesn't pair with services router.**

NAT rules do not take affect on the downlink router after a services router which had been paired with a downlink router recreated after being deleted.

- **Fixed Issue 2363248 - Service Instance Health Status on interface appears down, though API call shows connected.**

This inconsistent reporting may cause a false alarm.

This issue and solution are described in greater detail in [Knowledge Base article 67165 - Service Instance status displays as "Down" when there are no VMs up to be protected in NSX-T.](#)

- **Fixed Issue 2359936 - Frequent cfgAgent log rolling on ESX Host.**

Frequent log rollings may cause loss of useful information in cfgAgent.log for debugging and troubleshooting on host.

- **Fixed Issue 2332938 - When the SYN Cache is enabled in the Flood Protection Security Profile, the actual TCP half-open connection limit can be larger than is configured on the NSX Manager.**

NSX-T auto-calculates an optimal TCP half-open connection limit, based on the configured limit. This calculated limit can be greater than the configured limit and is based on the formula $Limit = (PwrOf2 * Depth)$, where PwrOf2 is a power of 2 not less than 64, and Depth is an integer ≤ 32 .

- **Fixed Issue 2376336 - Address Family in route redistribution not supported by Policy and Edge.**

Address Family in Redistribution is not working or used in the application.

- **Fixed Issue 2412842 - Limit metrics logs to 40 MB on ESX to support hosts with ramdisk.**

This issue is addressed in detail by [Knowledge Base article 74574](#).

- **Fixed Issue 2385070 - IP discovery and DFW have opposite behaviors regarding IPv6 subnet.**

IP discovery considers 2001::1/64 as a host IP, while DFW considers it an IPv6 subnet.

- **Fixed Issue 2394896 - Host fails to upgrade from NSX-T Data Center 2.4.x to 2.5.**

The host fails to upgrade from NSX-T Data Center 2.4.0, 2.4.1 and 2.4.2 to 2.5. This may be due to KCP module unloading failure.

This issue is discussed in greater detail in [Knowledge Base article 74674](#).

- **Fixed Issue 2406018 - An event/alarm is triggered if password expiry is within 30 days.**
An event/alarm is triggered regarding password expiration if password expiry is within 30 days and even if password expiration is disabled.
- **Fixed Issue 2383328 - Feature request to provide utility that renders metrics data into human readable form.**
NSX-T Data Center collects and saves metrics data in a binary format; users have requested the ability to view this data in a human-readable format. This issue tracks that request.
- **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 2313673 - VM-based Edge transport nodes: users unable to connect uplinks to the NSX-T logical switches/segments.**
For VM-based Edge transport nodes, users are unable to connect the Edge transport node uplinks to the NSX-T logical switches/segments. They can connect them only to the vCenter's DVPGs. On the Configure NSX screen for VM-based Edge transport node's add/edit flows, the users are presented with the option to map the uplinks only with vCenter's DVPGs. The option to map the uplinks to the NSX-T logical switches/segments is missing.
- **Fixed Issue 2424394 - DHCP packets relayed by NSX-T DR cannot reach more than 10 hops.**
When DHCP server is more than 10 hops away, the relayed DHCP packets cannot reach the server.
- **Fixed Issue 2399994 - Redistributed routes missing intermittently.**
Network traffic may be impacted as route to T1 is unavailable for some time.

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](#)
- [NSX Intelligence Known Issues](#)
- [Operations and Monitoring Services Known Issues](#)
- [Upgrade Known Issues](#)
- [API Known Issues](#)

- [NSX Cloud Known Issues](#)

General Known Issues

- **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.

- **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 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 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 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 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 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://<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 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:

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

- **Issue 2377217 - After KVM host reboot, traffic flows between VMs may not work as expected.**

Rebooting the KVM host may result in reachability issues between VMs.

Workaround: After host reboot, restart the nsx-agent service with the following command:
systemctl restart nsx-agent.service

- **Issue 2371251 - Dashboard interface blinks when navigating to Backup & Restore page.**

This has been observed only in the Firefox browser and only in some deployments.

Workaround: Manually refresh the page or use another supported browser.

- **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.

Workaround: See [Knowledge Base article 70543](#) for details.

- **Issue 2267964 - If vCenter is removed, user is not warned about loss of services**

running on vCenter.

If a user removes the computer manager (vCenter) where services like Guest Introspection are deployed, the user is not notified about the potential loss of these services.

Workaround: This issue can be avoided if the user follows the correct procedure for adding a new vCenter as computer manager.

- **Issue 2444170: NSX CLI commands fail to uninstall datapath**

del nsx command does not uninstall NSX-T configuration and modules from host. This causes the installation or upgrade of NSX-T to fail.

Workaround: None.

- **Issue 2467479 - Once Firewall is set to Bypass for a SNAT rule, it cannot be blocked after change from Bypass to None.**

Once Firewall is set to Bypass for a SNAT rule, it cannot be blocked after change from Bypass to None.

Workaround: Delete and recreate the SNAT rule.

- **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 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 1957059 - Host unprep fails if host with existing vib is added to the cluster when trying to unprep.**

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

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

NSX Manager Known Issues

- **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 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 neighborhood 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 advertised 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.

- **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.

Workaround: Do not configure such VLAN ranges for Edge L2 Bridge and Point configuration.

Logical Networking Known Issues

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

A route map added to a redistribution rule from the Management Plane interface or API, may be removed if the same redistribution rule is subsequently modified through the Policy page interface or API. This is due to the Policy page interface or API do not support adding route-maps. This can result in advertisement of unwanted prefixes to the BGP peer.

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 2275412 - Port connection doesn't work across multiple TZs.**

Port connection can be used only in single TZ.

Workaround: None.

- **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 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).

Workaround: None.

- **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.

Workaround: None. The best option is to delete the logical router and re-create with the desired transit subnet configuration.

Security Services Known Issues

- **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 2395334 - (Windows) Packets wrongly dropped due to stateless firewall rule contrack entry.**

Stateless firewall rules are not well supported on Windows VMs.

Workaround: Add a stateful firewall rule instead.

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

If a VM uses an IPv6 address, but IPv6 snooping is not been 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 the IPv6 option in IPDiscovery profile is enabled at either the VIF or logical switch whenever IPv6 addresses are used.

- **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 2379632 - Multiple packets are logged when hitting Layer- 7 rule in classified stage.**

Multiple (2-3) packets are logged (dfwpktlogs) when hitting Layer- 7 rule in classified stage.

Workaround: None.

- **Issue 2368948 - Distributed firewall rules: Realized status for individual sections may not be current.**

Refreshing the DFW rule view doesn't update the realized status of individual sections in that view. As a result, the information may not be current.

Workaround: This affects only manual refreshing. Polling for realized status is periodic and will provide accurate updates. Users can also refresh individual sections for accurate status.

- **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 25 minutes to publish.

Workaround: None.

- **Issue 2424818 - Layer-2 and distributed firewall statuses not updated on NSX Manager interface.**

The status information produced by the logical exporter on workload VMs may not be forwarded to the management plane. As a result, the statuses displayed for these components are not correctly updated.

Workaround: None. The correct status information can be accessed via CLI on the corresponding VMs.

Load Balancer Known Issues

- **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 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 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.

- **Issue 2401715 - Error while updating the compute manager that thumbprint is invalid, even if correct thumbprint is provided.**

Observed when a vCenter v6.7U3 is added as compute manager in NSX-T manager. vSphere 6.7 supports changing PNID where FQDN or IP address can be changed. NSX-T 2.5 does not support this feature, hence the thumbprint issue.

Workaround: Delete the previously added vCenter and add the VC with newly changed FQDN. Adding registration may fail, as previous extension already exists on vCenter. Resolve the registration errors to get it successfully registered.

NSX Intelligence Known Issues

- **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 of IPAdressExpressions, MACAddressExpressions, paths in a PathExpression and external IDs in ExternalIDExpression should not exceed 500."

Workaround: If there are scenarios where 500-plus clients are connecting to the application VM or load balancer, you can create a rule to micro-segment access to the application load balancer, then select the application VMs to start recommendation discovery. In the alternative, you can subdivide the 500-plus member group into multiple, smaller groups.

- **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 the user reboots the NSX Intelligence appliance while a recommendations job is in progress, the job goes to a failed state. A user 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 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 2374231 - For SCTP, GRE and ESP protocol flows, Service is shown 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 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 anticipated amount within 30 days, the appliance might run out of disk space prematurely and become partially or completely non-operational.

Workaround: This can be prevented or mitigated by monitoring the disk usage of the NSX Intelligence appliance. If disk usage is being utilized at a high rate that indicates that space might run out, you can modify so the data retention period to a fewer number of days.

1. SSH into the NSX Intelligence appliance and access the `/opt/vmware/pace/druid-`

config/druid_data_retention.properties file.

2. Locate and change the correlated_flow setting to a value lower than 30 days. For example: correlated_flow=P14D
3. Save the file and apply the changes by running the following command:
/opt/vmware/pace/druid-config/druid-config-data-retention.sh

NOTE: It may require up to two hours for the data to be physically deleted.

- **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 then recommendation job and retry the publication.

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

After a transport node is disconnected or removed from the compute manager, NSX Intelligence shows the previous VMs 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 transport node disconnection from the compute manager. 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. Duplicate VMs are eventually removed from the interface depending on the selected time interval.

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

Customer sees additional flows from VM to IP-xxxx. This is due to the configuration data (Groups, VMs and services) from 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 2366630 - 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 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 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 2369802 - NSX Intelligence appliance backup excludes event datastore backup.**

This functionality is not supported in NSX 2.5.

Workaround: None.

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

If the user replaces 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 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 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 requires to access the interface.

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

- **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 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 on **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.

Operations and Monitoring Services Known Issues

- **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".

Workaround: Verify that the credentials for accessing the SFTP server are up to date.

Upgrade Known Issues

- **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.

- **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".

Workaround: Upgrade to ESXi 5 p04. Alternatively, put the host in maintenance mode, and reboot it. Retry the upgrade, and exit maintenance mode.

- **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 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.

Workaround: Use the API only to access these Policy groups post upgrade.

- **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 back by an NSX-T logical switch.

Workaround: Manually migrate all NSX-T backed workload VMs to an already upgraded host, then retry upgrade for the failed 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 2419246 - Ubuntu KVM upgrade fails.**

Upgrade of Ubuntu KVM nodes may fail due to nsx-vdpi service not running. However, the nsx-vdpi service depends on the nsx-agent, but at this point in the upgrade, the nsx-agent is not yet configured. The nsx-agent fails because the vm-command-relay component is not correctly started.

Workaround: Configure the incompletely installed nsx-agent. The following command reconfigures all unpacked or partially configured packages:

```
dpkg --configure -a
```

Or you can use the below commands to reconfigure only the nsx-agent and nsx-vdpi:

```
dpkg --configure nsx-agent
```

```
dpkg --configure nsx-vdpi
```

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 get redirected to partners.

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

- **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.

Workaround: Wait a few minutes, then retry.

- **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.

Workaround: Use the NSX Manager interface instead to perform this operation.

NSX Cloud Known Issues

- **Issue 2275232 - DHCP would not work for VMs on cloud if DFWs Connectivity_strategy 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.

- **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 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.

Workaround: None. After the delay period passes, the issue self-corrects.

- **Issue 2424818 - L2 and DFW stats are not updated on NSX Manager UI.**
All stats produced by logical exporter on workload VMs are not forwarded to MP. This causes a failure in displaying stats on the NSX Manager UI. There is no visibility of DFW statistics from the NSX Manager UI. Logical switch ports operational status will show up as DOWN and their corresponding stats will not work. This is only applicable for Cloud VMs.

Workaround: None. The stats can be seen via the CLI on the corresponding VMs.