VMware NSX-T Data Center 2.4 Release Notes

VMware NSX-T Data Center 2.4 | 28 February 2019 | Build 12456646

Check regularly for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

What's New

NSX-T Data Center 2.4 provides a variety of new features to provide new functionality for virtualized networking and security for private, public, and hybrid clouds. Highlights include a new intent-based networking user interface, context-aware firewall, guest and network introspection features, IPv6, highly-available clustered management, profile-based NSX installation for vSphere compute clusters, rebootless maintenance upgrade mode of NSX for vSphere compute, new in-place upgrade mode for vSphere compute and a migration coordinator for migrating from NSX Data Center for vSphere to NSX-T Data Center.

The following new features and feature enhancements are available in the NSX-T Data Center 2.4 release.

Management Cluster

NSX-T Data Center 2.4 now supports the ability to create a cluster of managers for high availability of the user interface and API. This clustering supports both external balancers for redundancy and load distribution or a NSX provided virtual IP for redundancy. In addition, the management plane function and the central control plane function have been collapsed into this new management cluster to reduce the number of virtual appliances that need to be deployed and managed by the NSX administration. The NSX Manager appliance is available in three different sizes for different deployment scenarios. A small appliance for lab or proof-of-concept

deployments. A medium appliance for deployments to 64 hosts and a large appliance for customers who deploy to a large scale environment. For details on configuration maximums, see the VMware configuration maximums tool at: https://configmax.vmware.com

Support for Single Cluster Design

Support Single Cluster designs with collapsed Edge+Management+Compute VMs all powered by single N-VDS in a single physical host. Typical reference designs for VCF SP customers prescribe 4x10G pNICs with two host switches; one for Edge+Management and another for compute VMs. This effectively isolates the communication between edge VM and the compute VM so traffic leaves the host and comes back in. However, with the trending economics of 25G NICs, VCF SP customers are standardizing on 2x25G NIC hosts and with this design they will be able to move to a single N-VDS powering a host with 2pNICs. In this design, the Edge VM and compute VM belonging to the same subnet are able to communicate without traffic leaving host uplinks and coming back in.

Policy and UI

NSX Management and Automation

- **Declarative Policy Management** - Simplify and automate network and security configurations through outcome-driven policy statements. This new declarative policy API reduces the number of configuration steps by allowing users to describe desired end-goals, while letting the system figure out how best to achieve it. Define an entire network topology and deploy it all in one shot, in an order-independent, prescriptive manner.

User Interface Enhancements

- **Enhanced Navigation and Page Layouts:** improved navigation bar and page layouts to reduce the number of clicks to access critical information.
- **Internationalization:** improved handling of locale-specific items, such as date/time format, number format, time zone.

Note: The Network Topology Visualization feature for the NSX Policy Manager, introduced in version 2.3, is deprecated in this release.

Firewall

Distributed Firewall and Gateway Firewalls will support filtering of IPv6 traffic from NSX-T Data Center 2.4. Additionally, there are various operational features added to the product as listed below:

Publish/Revert Button

A single publish button is available for the entire firewall table. This will be available for both Distributed and Gateway Firewalls. Prior to NSX-T Data Center 2.4, the publish button was for each section separately. This will be available via API. Additionally, you will have the option to revert your changes. You will also have the option to lock sections when changes are updated.

Rule Statistics

Each rule will have the hit count, packet count, session count, byte count and popularity index. It will also have maximums seen versus current hit counts. This statistics can be reset by a button.

Grouping Enhancements

Additional grouping criteria based on operating system for the VM and Active Directory groups are available.

Rule Visibility per VM

List of Firewall rules for a particular VM is available by looking at the logical switch port associate for every virtual machine.

IP Discovery for Virtual Machines

The default IP Discovery profile is being updated to include VMTools based IP discovery in addition to ARP snooping and DHCP Snooping. Existing customers upgrading from prior releases will have to update the IP Discovery profile to enable VMTools based detection. Additionally, creation of a global IP discovery profile is supported with NSX-T 2.4. Additionally, there are the following changes:

1. IPv6 IP discovery based on DHCPv6 and Neighbor Discovery mechanisms are available.
2. IPv6 discovery is disabled by default
3. Auto-Discovered IP bindings can be either manually whitelisted or put in ignore list.
4. Local Link IPv4 addresses will be ignored by default.

Identity Firewall

NSX-T Data Center 2.4 introduces Identity (User ID) based rules for Distributed Firewall. Firewall Administrators can now configure distributed rules on virtual machines based on Active Directory based groups. This feature allows firewall administrators to provide firewall rules based on the users logged into the virtual machines. NSX will auto-detect users logged in / logged off and accordingly specific rules will be enabled for the users. Identity based Firewall can detect and enforce rules for a single user per VM or even track multiple users with particular sessions in the same VM. Firewall administrators will be creating NSX-T groups using Active Directory groups as a criteria. NSX-T manager will automatically retrieve a list of Active directory groups from the provide Domain Controllers. Firewall administrators can control east-west access of users especially in virtual desktop environments or remote desktop sessions with terminal services enabled.

L7 Application Signatures for Context-Aware Distributed Firewall

NSX-T Data Center 2.4 provides the ability to L7 based application signatures in Distributed Firewall rules. Users can either use a combination of L3/L4 rules with L7 Application Signatures or can just create L7 Application Signature-based rules only. We currently support Application signatures with various sub-attributes for Server-Server or Client-Server Communications only. In NSX-T Data Center 2.4, this will be available for ESXi based transport nodes only.

FQDN/URL Whitelisting for Context-Aware Distributed Firewall

NSX-T Data Center 2.4 introduces URL / FQDN whitelisting based rules in Distributed Firewall. NSX-T Data Center introduces an innovation using Distributed DNS snooping to allow each connection from each VM to have its own resolution of URL/FQDN. Firewall Administrators can use pre-canned URL domains and apply it to rules in distributed firewall. Applications that are hybrid in nature accessing SaaS services or cloud-based services can micro-segment based on the URLs accessed. Client applications or Browsers accessing SaaS applications can be given access on a granular basis. In NSX-T Data Center 2.4, this will be available for ESXi based transport nodes only.

Service Insertion

NSX-T Data Center 2.4 introduces a broad array of native security functionality such as Layer 7 Application Identity, FQDN Whitelisting and Identity Firewall, which enable even more granular micro-segmentation. Besides the native security controls delivered by the Distributed and Gateway Firewall, the NSX Service Insertion Framework allows various types of Partner Services like IDS/IPS, NGFW and Network Monitoring solutions to be inserted transparently into the data path and consumed from within NSX without making changes in the topology.

In NSX-T Data Center 2.4, Service Insertion now supports East-West Traffic (i.e., Traffic between the VMs in the Data Center). All the traffic between VMs in the Data Center can be redirected to a dynamic chain of partner services.

The E-W service plane provides its own forwarding mechanism that allows policy-based redirection of traffic along chains of services. Forwarding along the service plane is entirely automated by the platform: failures are detected, and existing/new flows redirected as appropriate, flow pinning is performed to support stateful services, and multiple path selection policies are available to optimize for throughput/latency or density.

Guest Introspection

NSX-T Data Center 2.4 introduces the Guest Introspection service platform for VMware partners to provide policy-based agent-less antivirus and anti-malware offload capabilities for Windows-based Guest VM workloads on vSphere ESXi hypervisors.

In NSX-T Data Center 2.4, the Guest Introspection platform provides:

- Simplified deployment and life-cycle management by consolidating the Guest Introspection deployment into the NSX Agent host preparation installation and no longer requiring the Guest Introspection Universal Service VM to be deployed on each ESXi hypervisor.
- Consistent policy-based services across multiple vCenters.
- VMware partner scale enhancements through sizing of partner SVM (i.e., "Small," "Medium, "Large" partner appliances).

L2 Networking

Multiple N-VDS per Host

Apart from providing flexibility to organize VM traffic, this new capability to support multiple N-VDS per host facilitates compliance with PCI regulation where strict isolation is required for VM traffic.

With the addition of this feature, one can now separate ENS uplinks from non-ENS uplinks; this is a useful functionality since ENS does not have feature parity with N-VDS today so ENS powered workloads will get fast path but go low on features.

N-VDS Visualization

This functionality provides the capability to manage the N-VDS as a standalone object with the ability to drill down to see connected hosts, etc. When looking at a specific host, one is able to see a UI grid that shows how it's connected to the N-VDS. Logical interfaces like VM kernel interfaces are also visible as part of the N-VDS. This is a dramatic improvement over the host view shows list of interfaces that includes all physical NICs, VM Kernel interfaces, and all OVS ports in one view.

LLDP Support for Physical NICs

This feature helps close the gaps in LLDP implementation for NSX. It provides debugability for physical switch connectivity. The ability to decipher which physical ports are connected to which interfaces on a host lends to easy troubleshooting of cabling problems. The scope of this feature applies to all physical hosts (ESXi, KVM, Baremetal Linux Hosts and Baremetal Edge) that participate in NSX dataplane.

Support for Proxy ARP on Edge Node

When external clients access services like LB, IKE etc., with same subnet addresses, they hit device routing. They send ARP queries for those addresses bound to loopback ports, however, LR loopback ports don't have the MAC addresses, so they do not respond to those ARP queries. This causes access problems.

Currently the workaround is to configure a /32 routing in those clients, like Loopback IP/32 → uplink/CSP, so the traffic can be forwarded to uplink/CSP ports, then it can go to the correct loopback port. ARP Proxy is the right solution to overcome this drawback.

L3 Networking

MTU Configuration Enhancements

NSX-T 2.4 offers two new MTU global parameters:

- Global Physical Uplink MTU, which configures the MTU for all the N-VDS instances in the NSX domain. This can be translated as the maximum frames size for the GENEVE encapsulated frames or TEP MTU.
    - Uplink Profile MTU can override the global physical uplink parameter on a specific host.
- Global Logical Interface MTU, which configures the MTU for all the logical router interfaces.
    - Logical Router Uplink MTU and CSP port MTU can override on a specific port the global logical interface MTU if needed.

This will allow end to end communications for VM configured with MTU larger than 1500 bytes for east-west and north-south traffic.

Inter SR Routing

Tier0 logical routers in active/active mode can now automatically establish full mesh iBGP peerings among all the Service Routers (SR) part of a given Tier0 logical router. This prevents traffic drops in case of SRs configured with multiple uplinks and a failure of only one of them. A SR in this failure scenario will now forward the traffic to another SR if the destination is not available on its own uplinks.

DNS Forwarder Enhancements

- The DNS forwarder function can now be enabled or disabled without losing its current configuration.
- The DNS forwarder function exposes also statistics, events and alarms through API and UI.

Support SNAT from Uplink to Uplink

NSX-T 2.4 introduces the support of SNAT (Source address Translation) for traffic entering a Tier0 logical router through an uplink and leaving the same logical router through another uplink. This feature is useful when multiple Tier0 logical routers are interconnected.

Support for Proxy ARP on Tier0 Logical Router

NSX-T 2.4 introduces the support of Proxy ARP on Tier0 logical router uplinks. This will allow the deployment of NSX-T in environments where routing cannot be configured on the routers northbound of the Tier0 logical router. With this feature, NAT, LB or any stateful services can be configured with an IP address that belong to the network of the Tier0 uplink.

Edge Node Enhancements

- NSX-T 2.4 introduces the option on Bare Metal Edge node to support the management on the fast path NICs, no longer requiring a dedicated management NIC.
- Bare Metal Edge node supports also 25 Gbps Intel NICs XXV710.

- Edge node supports multiple GENEVE Tunnel EndPoint (TEP). This allows edge nodes not to be forced to use LAG for high-availability of the overlay traffic.

BGP Enhancements

- Starting from NSX-T 2.4, Tier0 logical routers support iBGP peering with physical routers northbound.
- NSX-T 2.4 introduces the option to enable ECMP across eBGP peers in different ASN (as-path multipath relax) and also the support for the Tier0 logical router to allow its own ASN in the AS-path (allow-as in).

IPv6

NSX-T 2.4 introduces IPv6 routing/forwarding and security. This includes the support for:

- IPv6 static routing
- IPv6 Neighbor Discovery

- DHCPv6 relay
- IPv6 Distributed Firewall (DFW)
- IPv6 Edge Firewall
- IPv6 address-family for MP-BGP and associated prefix-list/route-map
- IPv6 switch security
- IPv6 address discovery
- IPv6 ops tools

Operations

Traceflow Enhancements

Traceflow adds support for even more troubleshooting and visualization capabilities. In NSX-T 2.4, Traceflow provides observations for centralized services such as Edge Firewall, Load Balancer, NAT and Route-based VPN.

Installation Enhancements

- NSX enables simplified deployments using a new profile-based installation of NSX components for vSphere compute clusters. This feature helps enable faster deployments, drives configuration consistency, avoids manual errors and provides a way to “define once and re-use multiple times.”
- Support for automated installation and clustering of NSX manager nodes from the UI.
- Support for more deployment configurations that can create multiple N-VDS switches, migrate VMKernel ports and physical adapters via profiles.

Upgrade Enhancements

- Enhancements to provide fully orchestrated upgrades of ESXi hosts without incurring the cost of a host reboot using the default Maintenance mode NSX upgrade.
- Introduced a new NSX upgrade mode called "in-place" upgrade. This features helps provide operational simplicity and enables faster upgrades. When using this mode, the NSX components on the ESXi hosts are upgraded without having to power-off workloads or migrate them to a different hypervisor.
- Introduced a new framework and provided out-of-the-box tests to do pre-checks and post-checks during the NSX upgrade that can help highlight dormant underlying issues before you begin the actual upgrade or immediately after the upgrade.

NSX Backup On Change Detection

NSX enhances its disaster recovery solution by providing the ability to detect configuration changes and proactively back them up to secure storage. This feature enables customers to have better SLA for configuration backups without incurring the cost of backing up unnecessary files to the storage server.

NFV

The N-VDS switch will now support the following enhancements in EDP mode.

- Distributed firewall

- IP Discovery
- Spoof guard
- IP Fix
- IPv6
- Enhanced performance for the Edge VM, which now provides up to five times more throughput in EDP mode.
- Path redundancy for multi-homed applications. Ability to pin a VM to a specific uplink allows the construction of a multi-homed redundant path today on NSX with VTEPs.

Operations - AAA/RBAC and Platform Security

Operations

- **Principal Identity Enhancements:** Enables Principal Identity users to register and install NSX components. Added UI support for creation of Principal Identity users and role assignment.
- **Password Policy Enhancements:** Enforces minimum password length of 12 characters for default passwords. Introduces ability to set password expiration times and generates alarms when password is about to expire. By default, passwords expire after 90 days. See Knowledge Base article 70691 for instructions on resetting passwords and adjusting password expiration.

- **Certificate Management:** Adds ability to check the revocation status of certificate.

VPN

NSX-T 2.4 has added the following capabilities for VPN services:

- Policy API and GUI are available for both L3 VPN and L2 VPN services.
- L3 VPN services support certificate-based authentication for better security management.
- The L2 VPN Client mode is available to support L2 extension from NSX-T SDDC to NSX-T SDDC.
- DH groups 19, 20, and 21 are available to meet high security requirements.

Load Balancing

NSX-T 2.4 has added the following capabilities for load balancing services:

- Policy API and new GUI are available. The old Load Balancer GUI is still available under the Advanced Networking & Security tab.
- VIPs on Standalone SR can belong to the same subnet as the Centralized Service Port or CSP. Prior to this release, if you wanted to create a VIP on the same subnet as the CSP network, the CSP IP address had to be used for the VIP. Otherwise, you had to create a VIP on a different network.
- DNAT and Edge Firewall are supported for load balancer traffic flows on the same Tier-1 gateway. Prior to this release, load balancer traffic flows bypassed the Edge firewall.
- LB rules support HTTP headers beginning with "_". With this enhancement, NSX Load Balancer can be deployed for vIDM and AirWatch.
- A VIP can be used as a source IP address for LB SNAT.

- The maximum HTTP response header size can be configured up to 64K Bytes. The default size stays the same as the previous release at 4K Bytes.
- A Large Edge VM supports a Large LB instance. Prior to this release, the Large Edge VM supported up to a Medium LB instance.

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

NSX-T 2.4 now has a Migration Coordinator which can be used to aid in the migration from NSX Data Center for vSphere to NSX-T Data Center. This feature is designed to migrate existing hosts without the use of vMotion. The Migration Coordinator supports migration of layer 2 networking, layer 3 networking, firewall, load balancing, and VPN. The *NSX-T Data Center Migration Coordinator Guide* provides details of the tool.

There is no need for additional compute resources beyond just the deployment of NSX-T Managers and Edge Nodes. Once the migration is complete, a customer can un-install NSX for vSphere and the associated Managers, Controllers and Edges. Please note that this migration does impact data plane traffic and is designed to be completed in a single change window.

Automation, OpenStack and other CMP

NSX-T 2.4 introduces the following capabilities for OpenStack consumption via its Neutron plugin:

- Support of Rocky and Queens
- Support of Management Plane Clustering
  The OpenStack Neutron Plugin takes advantage of the new ability to have a cluster of managers. It can consume the three managers REST API endpoints without an external VIP for additional performance and higher availability.
- Support of Barbican
  The OpenStack Neutron Plugin now supports Barbican. Barbican is a REST API designed for the secure storage, provisioning and management of secrets such as passwords, encryption keys and X.509 Certificates. This allows to manage certificate for the Load Balancer as a Service in order to do HTTPS termination. This is a feature currently supported in VIO environment only.

NSX-T Terraform Provider adds the following capabilities in NSX-T 2.4 to those already existing (Creation of logical switches, routers, firewall rules etc.):

- Ability to support CRUD on Load Balancer and Load Balancer configuration (Monitor, Pool etc.)
- Ability to support CRUD on DHCP Servers
- Ability to support CRUD on NSX-T IPAM (IP Block, IP pool)

NSX Cloud

NSX-T 2.4 for NSX Cloud has many new features to ease adoption/deployment for a customer, give more options with respect to how a customer can do service insertion, VPN termination, manage their VDI environments and thereby manage a true multi-region, multi-cloud hybrid deployment.

These are some of the key features in NSX Cloud with NSX-T 2.4:

- Shared Gateway in Transit VPC/VNET for simplified, faster onboarding and consolidation
- VPN for back-haul traffic back to on-premises DC
- Selective North-South Service Insertion & Partner Integration
- Micro-segmentation on Horizon Cloud for Azure
- Intent-Based Policy for Hybrid Workloads

Simplified Transit VPC/VNET architecture: Starting with 2.4, customers can install a single NSX Cloud gateway on a transit VPC/VNET and manage up to 10 compute VPCs/VNETs. This simplifies the hub and spoke transit/compute architecture and enables Transitive Routing between Compute VPCs even when they don't have a peering connection. Using NSX overlay tunneling, traffic between VPCs can now be sent in an overlay tunnel. Forwarding policies can be set up right down at the VM level to dictate if traffic should be Geneve encapsulated and sent in the Overlay, or if it should be sent in the public cloud provider's underlay network. All these features give more flexibility in terms of how users can route traffic within and outside of their public cloud network.

VPN for back-haul traffic: NSX Cloud now has built-in support to have VPN tunnels to back-haul traffic from public cloud to on-premises Data Center. VPNs from on-premises Data Center can now be directly terminated at the NSX Cloud Gateway in the public cloud. Customers don't need the VGW provided by public cloud vendors and this reduces cost. It also reduces the management overhead as NSX Cloud Gateway automatically propagates the routes over BGP. From a BW perspective, NSX Cloud gives a huge bump in the capacity as well: Inter-VPC traffic flows can be at 5Gbps over peered VPCs versus just 1Gbps offered over VGW.

Selective North-South Service Insertion & Partner Integration: Customers can deploy Partner Service directly from Public Cloud Marketplace in the Shared Services / Transit architecture. The NSX Cloud gateway present in the transit VPC/VNET can be programmed to selectively route traffic to partner service appliance based on NSX policies. This can be huge cost savings to a customer as they are not forced to direct all traffic through a virtual L7 firewall appliance that they have bought for the public cloud which is billed based on the traffic that passes through it. And if that wasn't enough, service insertion with NSX Cloud requires no VPNs to compute VPCs/VNETs. More cost savings and less operational.

Micro-segmentation on Horizon Cloud for Azure: NSX Cloud now has a combined solution with Horizon Cloud for Azure. For customers who choose to have a Horizon VDI environment deployed in Azure, NSX Cloud will provide the necessary micro-segmentation and secure the VDI env.

Intent-Based Policy for Hybrid Workloads: The Cloud Service Manager (CSM) is now integrated with the NSX manager. Customers can now define a single intent-based policy from the Policy Manager without worrying about where the workloads are deployed or where they will move in the future. NSX Cloud will realize this policy in a consistent manner across on-premises DC, Azure and AWS.

General Behavior Changes

Starting from NSX-T 2.4, by default, the as-path-multipath-relax option will be enabled. This means that routes learned from two BGP peers in different ASN will be considered as ECMP if the other BGP attributes are consider equals. Before NSX-T 2.4, only the routes from one the BGP peers would be installed in RIB. This option can be disabled through API or UI.

Compatibility and System Requirements

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

API and CLI Resources

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

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

Document Revision History

28th February 2019. First edition.
2nd April 2019. Second edition. Added known issues: 2273651, 2279326, 2281095, and 2296888. Added fixed issue: 2199785.
10th April 2019. Third edition. Added known issues: 2203863, 2248186, 2252738, 2277543, 2276398, 2279326, 2281537, 2287124, 2290688, 2294178, 2295592, 2296430, 2297157, 2297918, and 2298499. Updated What's New section to include Support for Single Cluster Design.
20th June 2019. Fourth edition. Added known issue 2261818. Added fixed issue 2182745.
23rd August 2019. Fifth edition. Added known issues 2362688, 2395334, and 2392093.
November 23, 2020. Added General Behavior Changes.

Resolved Issues

- **Fixed Issue 1842511: Multihop-BFD not supported for static routes**
  In NSX-T 2.0, BFD (Bi-Directional Forwarding Detection) can be enabled for a (MH-BGP) multihop BGP neighbor. The ability to back a multihop static route with BFD is not configurable in NSX-T 2.0, only BGP. Note that if you have configured a BFD backed multihop BGP neighbor and configure a corresponding multihop static route with the same nexthop as the BGP neighbor, the BFD session status affects both the BGP session as well as the static route.

- **Fixed Issue 2279326: No error is shown when creating an IPFIX L2 Collector with more than 4 IP:PORT combination**
  An error message is not shown for max no of IP:Port combinations allowed. There is no harm as the UI restricts tag creation if the maximum limit is exceeded.

- **Fixed Issue 1931707: Auto-TN feature requires all hosts in the cluster to have the same pnics setup**
  When the auto-TN feature is enabled for a cluster, a transport node template is created to apply to all hosts in this cluster. All pnics in the template must be free on all hosts for TN configuration or the TN configuration might fail on those hosts whose pnics were missing or occupied.

- **Fixed Issue 1909703: NSX admin is allowed to create new static routes, nat rules and ports in a router created by OpenStack directly from backend**
  As part of RBAC feature in NSX-T 2.0, resources like Switches, routers, Security Groups created by the OpenStack plugin cannot be deleted or modified directly by NSX admin from the NSX UI/API. These resources can only be modified/deleted by the APIs sent through the OpenStack plugin. There is a limitation in this feature. Currently NSX admin is only stopped from deleting/modifying the resources created by OpenStack, although admin is allowed to create new resources like static routes, nat rules inside the existing resources created by OpenStack.

- **Fixed Issue 1989407: vIDM users with the Enterprise Admin role cannot override object protection**
  vIDM user with the Enterprise Admin role cannot override object protection and cannot create or delete Principal Identities.

- **Fixed Issue 2030784: Cannot log in to NSX Manager with remote username that contains non-ASCII characters**
  You cannot log in to the NSX Manager appliance as a remote user with username containing non-ASCII characters.

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

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

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

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

- **Fixed Issue 2059414: RHEL LCP bundle installation fails due to older version of**

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

- **Fixed Issue 2142755: OVS kernel modules fail to install depending on which minor RHEL 7.4 kernel version is running**
  OVS kernel modules fail to install on a RHEL 7.4 host running a minor kernel version 17.1 or above. The installation failure causes the kernel data paths to stop working which leads the appliance management console to become unavailable.

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

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

- **Fixed Issue 2077145: Attempting to forcefully delete the transport node in some cases might cause orphaned transport nodes**
  Attempting to forcefully delete the transport node using an API call where for example, there is a hardware failure and the hosts become irretrievable, changes the transport node state to Orphaned.

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

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

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

- **Fixed Issue 2183549: When editing a centralized service port, not able to view a newly created VLAN logical switch**
  In Manager UI, after you create a centralized service port and a new VLAN logial switch, if you edit the centralized service port, you cannot see the newly created VLAN logical

switch.

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

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

- **Fixed Issue 1928376: Controller cluster member node degraded status after restoring NSX Manager**
  Controller cluster member node might become unstable and report degraded health status if the NSX Manager is restored to a backup image that was taken before this member node was detached from the cluster.

- **Fixed Issue 2128361: CLI command to set the log level of the NSX Manager to the debug mode not working properly**
  Using the CLI command set service manager logging-level debug to set the log level of the NSX Manager to debug mode not collecting debugging log information.

- **Fixed Issue 1940046: When the same static route is added and advertised in multiple Tier-1 logical routers, east-west traffic fails**
  If the same static route is added and advertised in multiple Tier-1 logical routers, east-west traffic fails.

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

- **Fixed Issue 2199785: NGINX Core observed when adding health monitor (without port number) to dynamic pool (having port number)**
  When configuring Load Balancing with server pool having dynamic members (with port number) and then trying to associate a health monitor that does not have any monitoring port configured, nginx may crash.

- **Fixed Issue 2182745: Previously le/ge in redistribution rules were not validated in the manager and were not functioning correctly**
  Redistribution rules support le/ge in prefixlists.

Known Issues

The known issues are grouped as follows.

General Known Issues

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

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

- **Issue 2287482: Auto-discovered bindings table may include bindings that are not currently discovered**
  Bindings that are marked "duplicate" in the auto-discovered bindings table might be no longer discovered.

  Workaround: None.

- **Issue 2278142: Switch IPFIX global profile is not editable**
  If Global profiles are available in the system, you cannot either modify or delete them through the interface as there is no workflow for Global profiles.

  Workaround: Delete such Global Profile using the API.

- **Issue 2292222: On Resolve Error screen, user is not notified if thumbprint is incorrect**
  If Host Preparation operation fails, the user can resolve the issue by clicking NSX Install Failed, in which case they need to provide username, password and thumbprint of host. If user gives incorrect thumbprint, the systems does not notify the user and the issue remains unresolved.

  There is no clear way to know that thumbprint was incorrect. Check log where this ThumbPrintValidationFailedException is logged.

  Workaround: Provide the correct thumbprint.

- **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 2285117: Kernel upgrade on NSX Managed VMs is not supported**
  On some Linux Ubuntu marketplace images, the kernel automatically upgrades itself on reboot of the VM. As a result, the nsx-agent does not function as expected. Although the NSX agent may appear to be functional, there will be some unrealized networking policies which affect the nsx-agent. The agent retries realizing these policies over and over, causing high CPU usage.

  Workaround: If a kernel upgrade is required, the appropriate Linux-headers for that newer kernel have to be downloaded first and the openvswitch data path dkms package needs to be recompiled.

- **Issue 2285544: MD5 hashes are no longer supported when invoking NSX APIs that require specifying a ssh_fingerprint value**
  NSX-T 2.4 no longer supports non-FIPS encryption algorithms, hashes, etc., which includes invoking the backup/restore, file-store and support bundle NSX APIs and specifying a MD5 hash for the ssh_fingerprint value. As a result, MD5 hashes are no longer supported.

  Workaround: Specify a different hash computed using a different hashing algorithm, for example, SHA256.

- **Issue 2256709: Instant clone VM or VM reverted from a snapshot loses AV protection briefly during vMotion**
  Snapshot of a VM is reverted and migrates the VM to another host. Partner console doesn't show AV protection for migrated instant clone VM. There is a brief loss of AV protection.

  Workaround: None.

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

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

- **Issue 2266553: In NSX appliance, a service may fail to initialize at its first boot**
  The deployed node is unable to serve requests, or unable to form a cluster.

  Workaround: Try to restart the failed service.

- **Issue 2267632: Loss of GI Protection configuration**
  Guest protection rule published on Policy UI shows SUCCESS. Corresponding change in the behavior is not reflected on the guest VM. OpsAgent logs at the same time shows

restart. Loss of guest VM protection.

Workaround: Replay the configuration change manually.

- **Issue 2269901: vmk interface is not included in packet capture CLI**
  This command cannot be issued.

  Workaround: Use packet capture uw to do the same.

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

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

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

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

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

  Workaround: None.

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

  Workaround:

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

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

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

- **Issue 2275985: Vnics not connected to logical switch are listed as options for NSGroup direct members**
  A vnic that is not connected to a logical switch is added as direct member of the NSGroup. Operation succeeds but the policies applied on that group do not get enforced on the vnic.

Workaround: None.

Check if a vnic is connected to a logical switch before adding it as a direct member of an NSGroup.

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

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

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

  Workaround: None.

- **Issue 2290669: As the number of virtual servers increases, the configuration time for each increases**
  As the number of virtual servers increases, the configuration time for each increases due to large numbers of validation. For the first 100 virtual servers, the average response time is around 1 second. After 250 virtual servers, the average response time increases to 5-10 seconds. After 450 virtual servers, the response time increases to around 30 seconds.

  Workaround: None. You may be able to configure Virtual Servers as multiple LbServices depending on topology, otherwise expect slower response times when configuring large scale setups with virtual servers.

- **Issue 2292116: IPFIX L2 applied to with CIDR-based group of IP addresses not listed on UI when group is created via the IPFIX L2 page**
  If you try to create a group of IP addresses from Applied to dialog and enter wrong IP address or CIDR in the Set Members dialog box, those members are not listed under groups. You have to edit that group again to enter valid IP addresses.

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

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

  Workaround: Once appliance deletion through API/UI failed, delete that appliance manually using force-delete API, as follows:
  POST api/v1/cluster/nodes/deployments/467a102d-472f-4f43-a93c-08b992b9f471?

action=delete&force_delete=true
Afterwards, destroy the VM from vCenter.

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

  Workaround: Remove SVM from host and then add it to cluster.

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

  Workaround: None.

Installation Known Issues

- **Issue 2238093: Resolver not supported if NSX packages are forcefully removed**
  For uninstalling NSX from the host, the NSX packages are forcefully removed. This may result in corrupt state of NSX packages. Resolver for NSX package installation may not work successfully, if prior to resolver the NSX packages are removed forcefully. Log location: /var/log/proton/nsxapi.log

  Workaround: None.

  Do not remove NSX packages forcefully. Uninstall the NSX components via standard steps described in NSX documentation.

- **Issue 2288872: Install state shown as "Node not ready"**
  The Edge node is not getting onboarded. The Transport Node configuration state is

Pending, and so cannot be added to an Edge cluster. Log location: /var/log/proton/nsxapi.log

Workaround: Retry Edge node registration. Alternatively, power off the Edge node. When it starts it will establish the MP-MPA channel.

- **Issue 2252776: Transport Node Profile fails to be applied on one of the cluster member hosts even if validation error that has occurred previously on the host is now resolved**
  TNP is applied on the cluster. But TNP cannot be applied on one of the cluster member hosts because one of the validations could not be passed (e.g., VMs are powered on on the host). User resolves the issue, but validation is still shown on UI and TNP is not automatically applied on that host.

  Workaround: Move host out of cluster and add it again. This will trigger the activity to apply Transport Node Profile on the host.

- **Issue 2284683: Not able to delete auto deployed appliance when a registered compute manager is deleted and added again**
  Deletion of appliance failed with error, "Failed to power-off" and stating compute manager is not found.

  Workaround: Once appliance deletion through API/UI failed, delete that appliance manually using force-delete API, as follows: POST api/v1/cluster/nodes/deployments/<node-id>?action=delete&force_delete=true . Destroy the VM from VC.

- **Issue 1957059: Host unprep fails if host with existing vibs added to the cluster when trying to unprep**
  If vibs are not removed completely before adding the hosts to the cluster, the host unprep operations fails.

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

- **Issue 2296888: The Transport Node (TN)/Transport Node Profile (TNP) configuration cannot have both PNIC Only Migration flag set to true and VMK Mappings for Install populated across host switches**
  When giving a mismatch of configuration (both PNIC only migration flag set to true and VMK Mappings for Install populated across host switches) during CREATE, the following exception is seen:

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

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

  An exception is seen when applying the TN/TNP configuration which contains both the PNIC Only Migration flag set to true and a vmk migration mapping.

Workaround: Each configuration sent to the host can either have PNIC only migration flag set to true or VMK Mappings for Install populated, but not both.

1. Send the TN configuration with the host switches that require the PNIC Only Migration set to true.
2. Update the TN configuration by setting all PNIC Only Migration flags to false and populate the VMK Mappings for Install as desired. In other words, ensure the configuration sent to the TN only has PNIC Only Migration flag set to true or VMK Mappings for Install populated across all host switches. Two separate configuration calls must be made for any configuration which requires both.

- **Issue 2273651 - After deleting transport node, user is unable to SSH into the host.**
  Observed in KVM implementations. The user deletes a transport node and receives a message that the deletion was successful. However, afterward the user is unable to access the same host via SSH. The issue is likely caused by the presence of an open virtual switch (OVS) that is not managed by NSX-T and was likely pre-installed as part of the KVM template.

  Workaround: Identify the problematic OVS before deleting the transport node.

  1. Run ovs-vsctl show to identify the OVS.
  2. Migrate any workload VM interfaces from the OVS to the Linux bridge.
  3. Delete the transport node as follows:
     DELETE api/v1/transport-nodes/<uuid>
- **Issue 2281537 - Post-migration, the ESXi transport node with multi-VTEP fails to start BFD session.**
  After migrating a NSX-V node to NSX-T, the ESXi transport node with multi-VTEP fails to start BFD session on all VTEPs to Edge nodes.

  Workaround: Restart the netcpa service.

NSX Manager Known Issues

- **Issue 2285306: Service deployment status for Guest Introspection services may remain in "Unknown" state until service VM gets powered on**
  After creating a Service Deployment and it is listed in the Service Deployment grid, the status may not immediately show as "In Progress" and may remain "Unknown" until the grid is refreshed.

  Workaround: None. Refresh the page after ten seconds. The status should update.

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

  Workaround: Review host configuration, redo the credentials, and retry adding the host.

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

  Workaround: Delete the binding map and recreate it with the new desired sequence

number.

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

  Workaround: None.

NSX Edge Known Issues

- **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.

  Workaround: Use a different machine or deploy NSX-T Edge as VM on a Hypervisor.

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

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

  - These APIs support filtering options based on network prefixes or type of route. Use these options to download the routes of interest.
  - CLI support in case the entire RIB/FIB table is needed and there is no timeout for the same.

Logical Networking Known Issues

- **Issue 2243415: Customer unable to deploy NXGI service using Logical Switch (as a management network)**
  On the NXGI deployment screen, the user cannot see a logical switch in the network selection control. If the API is used directly with logical switch mentioned as management network, user will see the following error: "Specified Network not accessible for service deployment."

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

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

  Workaround: To prevent this issue, manually verify if a Transport Node is associated with any NS Groups. In the Management Plane interface, navigate to **Advanced Networking**

Security > Inventory > Groups or to **System > Nodes > Transport Nodes > Related > NSGroup**.

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

  Workaround: Reboot the affected Edge node.

- **Issue 228688: BGP neighbor should be deleted first while deleting IPsec Route base session if BPG is configured over VTI**
  If BGP is configured over VTI and you delete the IPsec session, both SR will be in a down state which in turns blocks the traffic. In order to resume the traffic, the BGP neighbor configured for VTI should be deleted. In this scenario, only the BGP configured is over VTI.

  Workaround: Delete the BGP neighbor before deleting the IPsec session.

- **Issue 2288509: MTU property is not supported for Tier0/Tier1 service interface (central service port)**
  MTU property is not supported for Tier0/Tier1 service interface (central service port).

  Workaround: Configure MTU using management plane API even though CSP port is created by Policy workflow.

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

  Workaround: Correct the configuration after the error displays.

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

  Workaround: None.

- **Issue 2290083: Validation missing when creating VLAN based segment**
  When you specify a VLAN transport zone with a VLAN ID property, the system fails to validate and identify the error. As a result, the intent will fail during realization and raise an error.

  Workaround: See the Realization alarm error details for instructions to fix the input.

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

  Workaround: Perform the configuration using the API.

- **Issue 2298499 - VPN fails between public cloud gateway and peer host if gateway is not deployed with public IP.**
  The VPN tunnel between the public cloud gateway (PCG) and the peer host cannot be established if the PCG is deployed without a public IP address on the uplink. The reason is that the PCG by default is performing SNAT on the VPN traffic by default.

  Workaround: When deploying the public cloud gateway, enable public IP for the uplink interface.

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

  Workaround: None.

Security Services Known Issues

- **Issue 2288523: Uploading NSX Guest Introspection driver can lead to security problems**
  IDFW relies on User identity information from NSX Guest Introspection driver. Unloading the driver can lead to security problems for users logged in from the specific Guest. This presents as the following symptoms:

  - Firewall rules not enforced for users logged in from certain Guest VMs where Guest Introspection driver is unloaded.
  - IDFW component not logged in user details for users logging in from certain guest VMs where Guest Introspection driver is unloaded.
  - MUX logs don’t show any connections from these Guest VMs even though IDFW is enabled on the host.
  - MUX logs don’t show any network events from these Guest VMs even though IDFW is enabled on the host.

  As a result the default deny all rule can block access to logged in users from guest VMs where Guest Introspection driver is unloaded.

  Workaround: None. The IT Administrator should follow security best practices to ensure no users are given privileges to unload Guest Introspection drivers inside Guest VMs.

- **Issue 2288773: Old TLS protocol API still available, gets overwritten**

NSX-T has new API for setting NSX TLS protocol versions and cipher suites, which updates all nodes in an NSX-T cluster. However, the old API is still available. This can be used but the new settings will be overwritten by the global settings.

Workaround: Use the new API.

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

  Workaround: Create a new service for TFTP as L4PortSet service and use it in firewall rule.

- **Issue 2203863 - Identity firewall rules are not supported for UDP and ICMP traffic.**
  Identity firewall rules do not work with ping testing The current functionality is supported for TCP traffic only.

  Workaround: Use TCP for testing identity firewall rules. Never set ANY/UDP/ICMP in the service column when configuring identity firewall rules

- **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 2252738 - For fully qualified domain name (FQDN) rules, a packet not matching the rule is allow to reach the destination.**
  When a specific FQDN rule is created, the domain name associated to an IP address is add to the firewall database matching matching rule, and packets sent to that domain name are allow to reach the server. However, if a user changes the domain name associated with that IP address on the domain name server, the domain name entry is not updated in the firewall database (unless other another FQDN rules matching the new domain name exists). As a result, packets are sent to the new domain name even though should be dropped by the FQDN rule.

  Workaround: None.

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

  Workaround: Add a stateful firewall rule instead.

- **Issue 2458384 - NSX-T Manager interface pages fail to load with 403 error.**
  Observed in release versions 2.4.0 and 2.4.1. This issue affects both admin and Identity Manager logins. The FQDN of the NSX-T Manager uses the *.SLD.TLD format. For

example: *.co.uk, *.co.il, *.com.au and so on.

Workaround: Access the NSX-T Manager UI using shortname or IP instead of FQDN. See https://kb.vmware.com/s/article/71217.

KVM Networking Known Issues

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

  Workaround: An update to any DFW rule clears this error condition. For example just toggling the rule logging forces the KVM DFW module to clear the error condition.

Load Balancer Known Issues

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

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

- **Issue 2297157 - Load Balancing HTTPS performance is impacted by the FIPS mode.**
  Load balancing performance can be negatively affected when the default FIPS mode is enabled.

  Workaround: For a workaround, see Knowledge Base article 67400NSX-T 2.4.0 Load Balance Service may observe low performance on HTTPs.

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

  Workaround: None.

Solution Interoperability Known Issues

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

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

Operations and Monitoring Services Known Issues

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

  Workaround: None.

- **Issue 2289984: mux_connectivity_status shows as CONNECTED even after nsx-context-mux service is stopped on host**
  When nsx-context-mux or nsx-opsagent is not running on the host, the system (NSX interface or service instance API) incorrectly shows Solution status and GI agent status as running with an unchanged timestamp. As a result, the guest VMs may lose AV protection.

  Workaround: Try to manually start the mux and opsagent on the host if they are not already running.

  1. Log in to the host as root and execute following commands:
     /etc/init.d/nsx-opsagent start
     /etc/init.d/nsx-context-mux start
  2. After starting the agents, wait for a few minutes and verify that the health status timestamp on UI has updated.

Upgrade Known Issues

- **Issue 2273737: After upgrading from NSX-T 2.3 to 2.4, the vIDM server details are missing**
  If using vIDM, users where the vIDM server is configured only on the NSX Policy appliance, the vIDM server is migrated in the upgrade but the vIDM server will be missing from the converged appliance.

  Workaround: There are two options, depending on when the customer encounters this issue:

  - Before upgrading from version 2.3 to 2.4:
    Configure the same vIDM server details on both NSX Policy appliance and NSX Manager VM.
  - After upgrading from version 2.3 to 2.4:
    Reconfigure the same vIDM server details on the converged appliance.

- **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 2279973: If a blank group is created and upgrade proceeds, after MP upgrade, that blank group shows as not started**
  This occurs If a blank group is created and upgrade proceeds.

  Workaround: Do not create blank group.

  Do one of the following to proceed:

  - Delete the blank group
  - Click a resume button to finish the upgrade
  - Reset the plan
- **Issue 2282389: UC upgrade plan not in sync with VC cluster membership if ESX is moved across clusters**
  When ESX is moved from one cluster to another in VC, the change is not reflected in UC upgrade plan. This may lead to more than one HOST entering maintenance mode at the same time if user has selected "Parallel Upgrade" across groups.

  Workaround: On Host Upgrade page, click the "Reset" option to rebuild the plan so that UC upgrade plan is in sync with VC clusters.

- **Issue 2288921: Upgrade status goes out of synch when old version Edge nodes are added**
  Upgrade status goes out of synch if the user adds Edge nodes of an older version following the Edge upgrade. This causes issues in continuing the upgrade call.

  Workaround: First, avoid adding old version Edge nodes. If you do encounter this issue, restart the UC service.

- **Issue 2291625: PCG upgrade status changes from SUCCESS to NOT_STARTED after upgrade plan synch**
  This issue is only encountered if the user upgrades the PCG and then tries to upgrade more Agents/PCG afterward.
  In the recommended workflow, after PCG upgrade there are no more cross-cloud components to upgrade via the UC interface.

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

  Workaround: None. Functionality should not be affected.

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

Workaround: Restart the affected VMs.

- **Issue 2295564: Edge node controller connectivity may go down after upgrading from 2.3 to 2.4**
  This is an intermittent issue that will impact some north-south traffic.

  Workaround: Enable and disable maintenance mode on same edge node.

- **Issue 2294178 - Host VIB update fails during upgrade from 2.3.1 to 2.4**
  The upgrade process from the version 2.3.1 to 2.4 may fail with the error Install of offline bundle failed on host. More specifically, host VIB update fails because the switch security module fails to unload. The issue is known to occur if the IP discovery feature is enabled in the switching profile and when performing an in-place upgrade from NSX-T 2.3.1 to NSX-T 2.4 with a host running ESXi-6.7EP06 (build 11675023).

  Workaround: For a workaround, see the Knowledge Base article 67445With IP Discovery enabled, host VIB update may fail when upgrading from NSX-T 2.3.1 to NSX-T 2.4.

- **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 67444Host 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 2297918 - Post-upgrade from 2.3.1 to 2.4, unable to remove NSX from cluster.**
  After upgrading a cluster from 2.3.1 to 2.4, NSX-T cannot be removed and fails with the following message: "Failed to remove NSX on the cluster: Related transport node template or transport node collection exists for this fabric template. Transport node template or transport node collection must be deleted before a delete/disable on this fabric template."

  Workaround: Detach the transport node profile from the affected cluster, then use the "Remove NSX" workflow.

- **Issue 2286030 - Transport node configuration displays as in failed state when upgrading from NSX-T 2.3.x and earlier to 2.4.x.**
  Transport node configuration goes into failed state when upgrading from NSX-T 2.3.x and earlier to 2.4.x due to a Null Pointer Exception. When you have ESXi transport node with

vmkernel adapters migrated to N-VDS VLAN logical-switch and then upgrade from NSX-T 2.3.x to NSX-T 2.4.x, a race condition may cause the ESXi transport node configuration status to display as failed. However, the ESXi transport node connectivity with the NSX Manager and controllers is intact during upgrade even after the node is marked for configuration state failed.

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

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

API Known Issues

NSX Policy Manager Known Issues

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

  Workaround: Prevent this issue by always creating policies with appropriate sequence_numbers or using url parameters action=revise&operation=insert_top

- **Issue 2289278: Policy API throws error but allows configuration of multiple Virtual Servers with same pool with different persistence profile**
  The system does not support configuration of conflicting persistence types for the same pool for different LbVirtualServers. However, the Policy fails to properly validate/reject the conflicting input and allows the configuration. Subsequently, the Policy shows an alarm with the error message.

  Workaround: If you encounter this issue, you can correct it by changing the group setting on the LbVirtualServer.

- **Issue 2248186 -The BGP router installs IPV6 routes from its neighbor with its own interface as next hop.**
  As a result, IPV6 forwarding for the installed route may fail and cause a forwarding loop.

  Workaround: To avoid this issue, configure a route map to filter the IPv6 connected addresses as next hop in the BGP updates.

NSX Cloud Known Issues

- **Issue 2287884: Certain Centos marketplace images are not supported for NSX Cloud**
  Only the Centos marketplace images whose distribution versions match their expected minor kernel versions are supported for NSX Cloud.
  For example the distribution versions and their corresponding kernel versions are expected to be as follows:

  - RHEL 7.5 3.10.0-862

- RHEL 7.4 3.10.0-693
- RHEL 7.3 3.10.0-514

Workaround: Use only recommended Centos distributions as recommended in the documentation.

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

  Workaround: Explicitly allow DHCP for cloud VMs in DFW.

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

  Workaround: Remove invalid Tag.

- **Issue 2280663: Offboarding multiple VPCs in parallel could result in failures in rare cases**
  Offboarding of one of the compute VPCs would fail.

  Workaround: Manually clean up VPC and corresponding groups on Policy.

- **Fixed Issue 2287124: After deploying PCG on a Microsoft Azure VNet, the VNet's tile in CSM erroneously reports a warning**
  After deploying PCG on a Microsoft Azure VNet, in CSM the VNet reports a warning sign (yellow triangle with an exclamation point). If you hover over the warning icon, CSM reports that the status of MP (Management Plane) and CCP (Control Plane) is Unknown. However, there may not be any problem with connectivity and the warning is displayed in error.

- **Issue 2290688 - Upgrading Windows 2016 VMs in AWS fails.**
  Upgrade of multiple Windows workload VMs fail in AWS. VM upgrade status displays in AWS portal as stuck at "1/2 Check." Retrying fails also. This issue is observed only within same NSX-T version upgrades.

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

  1. Make sure PCG is upgraded on the affected hosts, so that latest host components can be downloaded by the VM.
  2. Reboot the VM to get into a good state.
  3. Manually run uninstall cmd.
  4. Manually run install cmd.