



VMware NSX-T Data Center 2.5.2 Release Notes

VMware NSX-T Data Center 2.5.2 | 30 July 2020 | Build 16615902

Check regularly for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

- [What's New](#)
- [Compatibility and System Requirements](#)
- [General Behavior Changes](#)
- [Available Languages](#)
- [API and CLI Resources](#)
- [Revision History](#)
- [Resolved Issues](#)
- [Known Issues](#)

What's New

Features, Functional Enhancements and Extensions

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

Compatibility and System Requirements

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

General Behavior Changes

Support for BFD Multihop on Management/TEP interfaces

Starting in NSX-T Data Center 2.5.2, multihop BFD is supported on management/TEP interfaces. When configuring the maximum allowed BFD hops in the edge-cluster profile as one (default), single-hop BFD is used. For any value greater than one, multi-hop BFD is used.

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.

NSX Intelligence

All NSX Intelligence known and resolved issues and detailed documentation to help you install, configure, update, use, and manage NSX Intelligence is now separately available at [NSX Intelligence Documentation](#).

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

30 July 2020. First edition.

17 August 2020. Second edition. Added resolved issue 2606608. Updated known issue 2590444 with additional workaround information.

21 August 2020. Third edition. Updated the workaround for known issue 2590444.

11 September 2020. Fourth edition. Added resolved issue 2586606.

24 September 2020. Fifth edition. Moved issue 2586606 to known issues. Added known issues 2621322, 2491206.

15 March 2021. Sixth edition. Added known issue 2730634.

Resolved Issues

- **Fixed 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.
- **Fixed Issue 2416130: No ARP proxy when Centralized Service Port (CSP) is connected to DR's downlink**
No ARP proxy when Centralized Service Port (CSP) is connected to DR's downlink causing no traffic to pass.
- **Fixed Issue 2462079: Some versions of ESXi hosts reboot during upgrade if there are stale DV filters present on the ESXi host.**
For hosts running ESXi 6.5-U2/U3 and/or 6.7-U1/U2, during maintenance mode upgrade to

NSX-T 2.5.1, the host may reboot if stale DV filters are found to be present on the host after VMs are moved out.

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

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

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

1. Identify the RPM for *nsx-exporter* and *nsx-aggservice* using the command `'rpm -qa | grep nsx'`
 2. Remove the RPM for *nsx-exporter* and *nsx-aggservice* using `'rpm -e nsx-exporter'` and `'rpm -e nsx-aggservice'`
 3. Download the *nsx-lcp* tar file on the server and untar it.
 4. Install the *nsx-aggservice* and *nsx-exporter* packages.
- **Fixed Issue 2470210: DFW local address set not updated on the VNIC after Storage vMotion of a DFW protected Virtual Machine.**
During a storage vMotion, a race condition is triggered where the cfgAgent is observing two filters with the same Virtual Interface and Logical Switch Port for a brief period of time, resulting in an incorrect address set update on the VNIC resulting in a traffic drop.
 - **Fixed Issue 2498350: Gateway firewall rules are not applied in some instances causing traffic to hit the default drop rule.**
Traffic is dropped due to hitting the default drop rule.
 - **Fixed Issue 2509879: Reduce pressure on activity framework by moving Application Initialization operations away from using activity framework.**
Host to NSX Manager connectivity may be impacted due to a buildup of activity in the activity framework table.
 - **Fixed Issue 2512778: Route advertisement fails from T1->T0 due to backed up activities in the activity framework queue.**
Processing of new activities fails when activity framework is backed up with activities.
 - **Fixed Issue 2517232: Inventory Objects not loading up in NSX Manager UI.**
When logging in to the NSX Manager UI, inventory objects does not show up as elastic search runs out of memory while trying to index huge objects while loading inventory.
 - **Fixed Issue 2523475: PCF, Container APP not added dynamically to Security Group despite having matching tags.**
NSX objects like logical switches, logical ports or virtual machines are not dynamically added to NSGroup even though the membership criteria matches.
 - **Fixed Issue 2543353: NSX T0 edge calculates incorrect UDP checksum post-eSP encapsulation for IPsec tunneled traffic.**
Traffic is dropped due to bad checksum in UDP packet.
 - **Fixed Issue 2547983: NSGroups may not be cleaned up when deleted causing stale NSGroup entries in database.**

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

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

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

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

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

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

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

- **Fixed Issue 2346636: Fragmented packets with MF and DF flags set in the IP header are dropped by firewall.**

Fragmented packets with MF and DF flags set in the IP header were dropped by firewall.

- **Fixed Issue 2424331: Log files do not gets rotated after root password expires.**

Log files do not gets rotated and size of log files keeps increasing. This causes the log partition to fill up and some services start failing.

- **Fixed Issue 2456534: Following pre-emptive failback the new standby T0 router loses BGP peering for 20 minutes.**

When a failback occurs from the non-preferred node to the preferred node in a pre-emptive Active/Standby T0 deployment, the non-preferred node goes into standby and the BGP peerings on this standby node are stuck in Active state for 20 minutes. During this time BGP commands return no output. The issue resolves itself after a 20 minute timeout period and BGP session come back up in an Established state.

- **Fixed Issue 2468846: Upgrade doesn't work when Host is in "Install Failed" state.**

Upgrade doesn't work when Host is in "Install Failed" state.

- **Fixed Issue 2479735: Modification of Firewall bypass option from the NSX Manager UI is not processed.**

Modification of Firewall bypass option from the NSX Manager UI is not processed by the backend and you do not see the modification in API and on Edge CLI.

- **Fixed Issue 2482817: A CA signed certificate is rejected because the signing certificate is not RSA.**

You are unable to replace the API or VIP certificate because it is an EC certificate, not RSA.

- **Fixed Issue 2485039: Gateway firewall drops traffic that it should not.**

Gateway firewall drops traffic that it should not. This is because default stateful policy is created for Active-Active Tier0 logical router.

- **Fixed Issue 2488535: Host header cannot be updated by LB rule.**
Host header cannot be updated by LB rule. Even if you set the host header to another value, the change is not applied.
- **Fixed Issue 2490312: Alarms do not get deleted.**
Alarms for default rules do not clear.
- **Fixed Issue 2490695 / 2481033: Any change to the transport node profile fails to apply on an ESXI Transport Node, if there are any running VMs on that Host.**
Any change to the transport node profile fails to apply on an ESXI Transport Node, if there are any running VMs on that Host.
- **Fixed Issue 249177: LB returns "500 Internal Server Error" instead of serving the actual page.**
If LB rule-match condition uses capture groups, and the matched content has certain special characters, then LB returns "500 Internal Server Error".
- **Fixed Issue 2500256: Configuring a VLAN on an out-of-band management interface does not work correctly.**
When a VLAN is configured on an out-of-band bond management interface the configuration is persisted incorrectly, resulting in the management interface failing to come up correctly after a reboot.
- **Fixed Issue 2502877: BFD session is not formed between edges of the same cluster over the management interface.**
The UI for NSX Edge states that Edge is in degraded state when one BFD channel is using multihop BFD while the other is using single hop. You receives false messages about the health of the Edges in the cluster.
- **Fixed Issue 2507474: FILE_INTEGRITY_CHECK fails for python files.**
FILE_INTEGRITY_CHECK fails for python files.
- **Fixed Issue 2508326: If the network address of a T1 segment overlaps LB VIP, NSX Manager incorrectly validates all virtual servers connected to the segment, causing a failure.**
Segment creation fails if the network address overlaps LB VIP.
- **Fixed Issue 2509162: Objects, such as firewall rules, fail to realize on edge node.**
When gateway firewall policy is created on the gateway, the realization of firewall rule fails to publish to the edge node.
- **Fixed Issue 2511654: Syching large AD domain fails.**
You see the error "Out of disk space" on/*config* partition leading to NSX Manager not functioning.
- **Fixed Issue 2512094: A system crash experienced while using grouping objects on the host.**
A system crash experienced while using grouping objects on the host.
- **Fixed Issue 2513835: Incorrect compute members are shown in a group on the UI when you try to edit another group while the current group edit operation is in-**

progress.

You see incorrect compute members in the UI for a group that you are currently editing, if you also start editing a different group simultaneously.

- **Fixed Issue 2513842: While upgrading, MUB upload fails because of renaming of MUB before upload.**
Upgrading NSX-T Data Center fails if you rename the MUB before uploading.
- **Fixed Issue 2513848: CPU usage for dhcp-backend process is 100%.**
CPU usage for dhcp-backend process reaches 100% and renders the DHCP server unusable.
- **Fixed Issue 251391: If the management interface is on anything other than vmk0, an exception occurs in host pre-checks.**
Host pre-checks fail during an upgrade, if the management interface is on anything other than vmk0.
- **Fixed Issue 2513920: A superuser other than 'policy' does not have privileges to deploy an E-W Service.**
You are unable to deploy an E-W service without the "policy" superuser privileges.
- **Fixed Issue 2515554: System crash caused by double freeing of fastslab.**
System crashes by the double freeing of *fastslab*.
- **Fixed Issue 2518312: Installation of NSX-T Data Center fails because there is no support for kernels later than 4.15.0-76.**
If you try to deploy NSX-T Data Center on kernel versions later than 4.15.0-76, installation fails.
- **Fixed Issue 2518415: Upgrade gets stuck because some files aren't being copied over to NSX Manager.**
Upgrade gets stuck because some files aren't being copied over to NSX Manager.
- **Fixed Issue 2526373: NSX edge datapath fails to start.**
NSX edge datapath fails to start on bare metal edge when it has more than 32G hugepage memory configured on a CPU with single numa node.
- **Fixed Issue 2525781: DFW filters are applied on logical segment ports consumed by the NSX edge.**
NSX edge VMs whose network interfaces are edited to use NSX logical segments have DFW rules and filters applied impacting traffic flow and possibly causing latency.
- **Fixed Issue 2523397: An NSX-T Data Center-prepared ESXi host might crash during vMotion.**
An NSX-T Data Center-prepared ESXi host might experience a crash during vMotion.
- **Fixed Issue 2520658: The reverse-proxy service crashes and does not restart automatically. The customer needs to manually restart the reverse-proxy.**
The reverse-proxy service crashes and does not restart automatically. You must manually restart the reverse-proxy service.

- Fixed Issue 2544127: Transport Node fails to synch because of invalid configuration and the the NSX edge cannot be added to an edge cluster.**
 Because of the error that a MAC address for a VNIC null is not found on the NSX edge node, the NSX edge cannot be added to an edge cluster.
- Fixed Issue 2543581: The system might crash during vMotion export if the number of active states increases significantly during the process.**
 The system might crash during vMotion export if the number of active states increases significantly during the process.
- Fixed Issue 2541552: You might experience 100% disk usage.**
 Disk compaction keeps continuing and makes the disk config exceed 100% usage.
- Fixed Issue 2539526: You are unable to upgrade the NSX Manager because data migration fails due to a corrupted database.**
 Database might get corrupted if NSX Manager services are started before completing the NSX Manager's upgrade.
- Fixed Issue 2537112: Transport Node state shows RPC timeout.**
 Transport Node state shows RPC timeout.
- Fixed Issue 2535682: Azure VNET onboarding fails if there are Network Security Groups with rules that have protocols other than Any/TCP/UDP.**
 You are unable to onboard new Azure VNETs in to NSX-T Data Center and therefore not able to manage VMs within those VNETs.
- Fixed Issue 2533267: When trying to retrieve LB configuration statistics, the process named nsx-edge-exporter crashes and keeps restarting.**
 When trying to retrieve LB configuration statistics, the process named *nsx-edge-exporter* crashes and keeps restarting.
- Fixed Issue 2530312: You might experience LIFs of Logical Routers not getting processed with messages noting pending realization and connectivity issues on the logical router.**
 You might experience delayed processing of LIFs with the *nsxapi.log* file noting the message: "Delaying processing for new logical router link port."
- Fixed Issue 2528314: MAC addresses move between physical switch ports.**
 When NSX edge comes out of maintenance mode, MAC addresses learned from NSX edge L2 bridge are sent back to the physical network through RARP requests, causing MAC addresses to move between physical switch ports.
- Fixed Issue 2527921: High memory utilization by BFDD process on NSX Edge.**
 Memory leaks can be caused when BFD is enabled for BGP neighbors and BFD events are generated for those neighbors. This results in excess memory consumption by the routing stack's BFDD module.
- Fixed Issue 2526083: Some NSX services might not function properly when the NSX Manager becomes disconnected from the NSX Intelligence appliance.**
 In the System > Appliances page of the NSX Manager UI, the NSX Intelligence Appliance card displays an error or shows a status that the appliance appears to be stuck in the data

fetching state.

- **Fixed Issue 2548935: Spoofguard on ARP packets may not work when ARP Snooping is enabled in the IP Discovery profile.**

There is a possibility that the ARP cache entries of a guest VM could be incorrect even when spoofguard and ARP snooping are enabled in the IP Discovery profile. The spoofguard functionality will not work for ARP packets.

- **Fixed Issue 2572394 / 2574635: Unable to take backup when using SFTP server, where "keyboard-interactive" authentication is enabled but "password" authentication is disabled.**

User is unable to use SFTP servers, where "keyboard-interactive" authentication is enabled, but "password" authentication is disabled.

- **Fixed Issue 2572116: After taking an NSX edge node out of maintenance mode, T0 SR HA status takes a few minutes to appear.**

After taking an NSX edge node out of maintenance mode, T0 SR HA status takes a few minutes to appear. If the other NSX edge node is not in active state, this NSX edge node will not be able to serve traffic until its T0 SR HA status is not active.

- **Fixed Issue 2568794: There are continuous log messages when NSX Manager is disconnected from NSX Intelligence.**

The system displays log messages continuously when NSX Manager is disconnected from NSX Intelligence.

- **Fixed Issue 2568617: While upgrading, you might have to manually check the bootbank space to ignore false positive check results due to the df timing out.**

Bootbank pre-upgrade checks fail during execution because *df* times out. You have to manually check the bootbank space to ignore false positive check results.

- **Fixed Issue 2562949: The Network_Engineer role is not usable in many scenarios.**

The Network_Engineer role doesn't have permissions to reload Enforcement Point, making this role unusable in many scenarios.

- **Fixed Issue 2548030: Packet capture during heavy traffic causes datapath to crash and the edge node fails over to the standby node.**

Packet capture during heavy traffic causes datapath to crash and the edge node fails over to the standby node.

- **Fixed Issue 2545412: Groups get deleted even when they are used in security policies as scope, and you can no longer modify such security policies.**

Groups get deleted even when they are used in security policies as scope, blocking the modification of such security policies where these deleted groups were used.

- **Fixed Issue 2580550: In place upgrade is not supported with L7 firewall enabled.**

During in-place upgrade, the L7 attribute of new connections are not classified. But as soon as the upgrade is complete, the L7 rules start functioning normally.

- **Fixed Issue 2581156: Unable to set up DHCP relay for a VLAN segment.**

You are either unable to set up DHCP relay for a VLAN segment or it is very complex to set up.

- Fixed Issue 2582543: Load Balancer persistence doesn't work and virtual server not responding for some traffic.**
 Load Balancer persistence doesn't work and virtual server not responding for some traffic.
- Fixed Issue 2584230: Traffic loss of 1-3 seconds on creation of logical router ports for Tier0/Tier1 gateways.**
 You might experience momentary N-S traffic loss when you create logical router ports for Tier0/Tier1 gateways.
- Fixed Issue 2585286: Duplicate IP alarm on double TEP NSX edge node.**
 You might experience duplicate IP alarms on double TEP NSX edge nodes.
- Fixed Issue 2494047: Creation of Transport Nodes on the hosts fails if the cluster has a powered-on VM.**
 Creation of Transport Nodes on the hosts fails if the cluster has a powered-on VM.
- Fixed Issue 2507291: Deleted segments continue to be displayed on the user interface.**
 Deleted segments continue to be displayed on the user interface.
- Fixed Issue 2549959: VM replication from the Site Recovery Manager fails, leaving some VMs unprotected.**
 VM replication from the Site Recovery Manager fails, leaving some VMs unprotected.
- Fixed Issue 2478390: LB stops working sporadically.**
 When persistence is used, on vip no memory is pre-allocated for ssl session resulting in LB not working.
- Fixed Issue 2593826: Route advertised by Tier-1 router is not learned by Tier-0 router.**
 Route advertised by Tier-1 router is not learned by Tier-0 router.
- Fixed Issue 2514657: Realization of services might get delayed or even fail.**
 When you configure a network address with small prefixes on services (for example, NAT rule configured with network address 10.0.0.0/8), service realization might fail, and services might not be accessible from external networks.
- Fixed Issue 2591998: Service reference realization is changed to error and is never corrected.**
 A realized service reference changes to unrealized state during a brief period when NSX Manager is rebooted and never returns to the realized state.
- Fixed Issue 2533630: CentOS host was moved to Failed state after host upgrade.**
 CentOS host was upgraded to NSX-T Data Center 2.5.1 release but after some time the host moved to "Install failed" state.
- Fixed Issue 2606608: When the host config is updated and there are multiple PNICs used in the host-switch defining the TEP, the NSX Edge node state may briefly change to DOWN and back to UP immediately.**
 When host-config is updated, NSX Edge node state may briefly change to DOWN and

back to UP immediately when there are multiple PNICs used in the host-switch defining the TEP. This causes HA state flips for all HA resources. The config flap can cause BGP flap if BGP is configured on T0 logical router. In addition, failover may happen on some of the services such as logical router, L2-bridge, or DHCP.

Known Issues

The known issues are grouped as follows.

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

General Known Issues

- **Issue 2320529: "Storage not accessible for service deployment" error thrown after adding third-party VMs for newly added datastores.**

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

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

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

where *<nsx-manager>* is the IP address of the NSX manager where the service deployment API has failed, and *<CC Ext ID>* is the identifier in NSX of the cluster where the deployment is being attempted.

- **Issue 2328126: Bare Metal issue: Linux OS bond interface when used in NSX uplink profile returns error.**

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

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

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

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

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

started on vCenter.

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

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

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

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

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

Workaround: None.

- **Issue 2275708: Unable to import a certificate with its private key when the private key has a passphrase.**

The message returned is, "Invalid PEM data received for certificate. (Error code: 2002)". Unable to import a new certificate with private key.

Workaround:

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

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

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

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

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

Workaround: None

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

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

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

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

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

Workaround: Use the following procedure to resolve this issue:

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

- **Issue 2484006: Protected VMs lose network connectivity.**

SRM Protected VMs in an NSX-T Data Center environment lose network connectivity despite being configured on a different logical network, when placeholder VMs in the secondary site are powered on. This issue occurs because the same VIF UUID is applied to both the protected and the placeholder VMs.

Workaround: None.

- **Issue 2549175: Searching in policy fails with the message: "Unable to resolve with 'start search resync policy'".**

Searching in policy fails with because search is out of sync. when the NSX Manager nodes are provided with new IPs.

Workaround: Ensure that the DNS PTR records (IP to hostname mappings in the DNS server) for all the NSX Managers are correct.

- **Issue 2572052: Scheduled backups might not get generated.**

In some corner case, scheduled backups are not generated.

Workaround: Restart all NSX Manager appliances.

- **Issue 2589694: A few seconds of IPv6 traffic loss might be observed when VM failover takes place.**

A few seconds of IPv6 traffic loss might be observed when VM failover takes place. This happens when the IPv6 address of a workload VM is ported to another workload VM which is communicating with a different workload VM on a different L2 segment. The two isolated L2 segments are connected by the VDR.

The two workload vms communicating also needs to be in two different ESX TNs for the problem to be seen.

Workaround: None.

- **Issue 2555333: "nsxuser" fails to get created during host prep.**

During the host prep lifecycle (install/uninstall/upgrade) 'nsxuser' is created internally in ESXi hosts managed by vCenter Server for managing NSX VIBs. This user creation fails intermittently because of the ESXi password requirements.

Workaround: Retry the task involving host prep.

- **Issue 2486119: PNICs are migrated from NVDS back to VDS uplinks with mapping that is different from the original mapping in VDS.**

When a Transport Node is created with a Transport Node Profile that has PNIC install and uninstall mappings, PNICs are migrated from VDS to NVDS. Later when NSX-T Data Center is removed from the Transport Node, the PNICs are migrated back to the VDS, but the mapping of PNIC to uplink may be different from the original mapping in VDS.

Workaround: Go to the vCenter Server UI to change the PNIC-to-uplink assignment in the VDS in the host.

- **Issue 2569691: Ping between external network and logical switch/segment does not work in specific cases.**

Consider the following configuration:

- 1) Create an uplink with x.x.x.x network.
- 2) The default route creation for nexthop is: x.x.x.y
- 3) Now update the connected IP for uplink to: x.x.x.y

This is a misconfiguration and causes pings to fail from the external network to the logical switch or segment.

Workaround: Either configure the gateway address as the IP present on the nexthop's interface or provide a gateway as interface, for example:

IP route 0.0.0.0/0 <uplink_id>

Caution: If you provide gateway as interface, bear in mind that traffic is always routed through the specified uplink.

- **Issue 2607651: NSX Manager does not reflect users from vIDM, if First Name Attribute is missing.**

If a vIDM user is created in AD with no First Name / Last Name / Email ID Attribute then the user is not reflected in NSX Manager.

Workaround: Configure the vIDM User with the required attributes.

- **Issue 2605659: Packets are not getting forwarded to the pool members on correct port when NSGroup for server pool is not statically configured, rule action is "select pool" in forwarding phase and there is no default pool for virtual server. The matched packets after the first non-matched packet will be forwarded to backend**

server on port 80.

Packets are set to incorrect port.

Workaround: None.

- **Issue 2607918: SRM works only if both protected and recovery VMs are connected to logical switches that are in the same Transport Zones.**

SRM works only if both protected and recovery VMs are connected to logical switches that are in the same Transport Zones.

Workaround: None.

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

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

Workaround: None.

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

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

Workaround: None.

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

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

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

Installation Known Issues

- **Issue 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 2577028: Host Preparation might fail.**

Host prep might fail due to config hash mismatch leading to discovery loop.

Workaround: Try one of the following options:

- Make FQDN false and restart *nsx-proxy* in the host. This will force host and NSX Manager to not use FQDN.

OR

- If you want to use FQDN mode, make sure to deploy the NSX Manager appliance using the FQDN for hostname, and ensure that the case sensitive spelling matches both the forward and reverse DNS lookup for the NSX Manager IP address. This setting has to be consistent throughout all the NSX Manager nodes.

Upgrade Known Issues

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

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

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

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

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

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

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

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

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

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

- **Issue 2519300: NSX Manager upgrade fails with no clear errors.**

Upgrading NSX Manager might fail because of the Upgrade Coordinator providing a message: "This page is only available on the NSX Manager where Upgrade Coordinator is running." or there are no clear errors.

Workaround:

1. Run the following command: `/opt/vmware/nsx-mpa/mpaconfigrestore.sh`
2. Restart napi: `/etc/init.d/nsx-mp-api-server restart`

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.

Security Services Known Issues

- **Issue 2448006: Querying a Firewall Section with inconsistencies in rule-mapping fails.**

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

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

- **Issue 2590444: VM tags are deleted when ESXi hosts disconnect from vCenter Server for longer than 30 minutes.**

VM tags are deleted when ESXi hosts disconnect from vCenter Server for longer than 30 minutes, causing DFW rules based on VM tags to stop working as expected.

Workaround:

Try one of the following options:

- Re-apply the tags or reconnect the host to the vCenter Server within 30 minutes.
- Before disconnecting hosts, increase the timeout setting from 30 minutes to up to 72 hours. Timeout increase can be made with the help of VMware support.
- **Issue 2569153: DHCP DFW allow rule with 0.0.0.0 and 255.255.255.255 as SRC/DST will drop DHCP UDP 67/78 packets.**

You may not be able to filter or honor DHCP traffic within or into the firewall enabled deployments which will impact IP allocations of VMs or other NSX-T Data Center resources.

Workaround: Configure one rule for DHCP requests and another rule for DHCP response. Each rule must honor the set of DHCP server IPs and the keyword "Any" in the source and destination fields.

- **Issue 2557166: Distributed Firewall rules using context-profiles (layer 7) are not working as expected when applied to Kubernetes pods.**

After configuring L7 rules on Kubernetes pods, traffic that is supposed to match L7 rules is hitting the default rule instead.

Workaround: Use Services instead of Context-profiles.

Copyright © 2022 VMware, Inc. All rights reserved.