



VMware NSX for vSphere 6.3.5 Release Notes

VMware NSX for vSphere 6.3.5 | Released November 16, 2017 | Build 7119875

See the [Revision History](#) of this document.

What's in the Release Notes

The release notes cover the following topics:

- [What's New in NSX 6.3.5](#)
- [Versions, System Requirements, and Installation](#)
- [Deprecated and Discontinued Functionality](#)
- [Upgrade Notes](#)
- [FIPS Compliance](#)
- [Revision History](#)
- [Resolved Issues](#)
- [Known Issues](#)

What's New in NSX 6.3.5

Important: If you are upgrading NSX 6.2.0, 6.2.1, or 6.2.2 to NSX 6.3.5, you must complete a workaround before starting the upgrade. See [VMware Knowledge Base article 000051624](#) for details.

NSX for vSphere 6.3.5 adds serviceability enhancements, and addresses a number of specific customer bugs. See [Resolved Issues](#) for more information.

- Guest Introspection service VM now ignore network events sent by guest VMs, unless Identify Firewall or Endpoint Monitoring is enabled
- You can also modify the threshold for CPU and memory usage system events with this API: PUT /api/2.0/endpointsecurity/usvmstats/usvmhealththresholds
- Serviceability enhancements to L2 VPN including:
 - Changing and/or enabling logging on the fly, without a process restart
 - Enhanced logging
 - Tunnel state and statistics
 - CLI enhancements
 - Events for tunnel status changes
- Forwarded syslog messages now include additional details previously visible only on the

vSphere Web Client

- Host prep now has troubleshooting enhancements, including additional information for "not ready" errors

View Release Notes for previous versions:

- NSX [6.3.4](#)
- NSX [6.3.3](#)
- NSX [6.3.2](#)
- NSX [6.3.1](#)
- NSX [6.3.0](#)

Versions, System Requirements, and Installation

Note:

- The table below lists recommended versions of VMware software. These recommendations are general and should not replace or override environment-specific recommendations.
- This information is current as of the publication date of this document.
- For the **minimum supported** version of NSX and other VMware products, see the [VMware Product Interoperability Matrix](#). VMware declares minimum supported versions based on internal testing.
 - **The minimum supported version of vSphere required for NSX interoperability changes between NSX 6.3.2 and NSX 6.3.3.** See the [VMware Product Interoperability Matrix](#) for details.

Product or Component	Recommended Version
NSX for vSphere	<p>VMware recommends the latest NSX 6.3 release for new deployments and when upgrading from 6.1.x.</p> <p>When upgrading existing deployments, please review the NSX Release Notes or contact your VMware technical support representative for more information on specific issues before planning an upgrade.</p>
vSphere	<ul style="list-style-type: none">• vSphere 5.5U3 and later• vSphere 6.0U3 and later. vSphere 6.0U3 resolves the issue of duplicate VTEPs in ESXi hosts after rebooting vCenter server. See VMware Knowledge Base article 2144605 for more information.

- vSphere 6.5U1 and later. vSphere 6.5U1 resolves the issue of EAM failing with OutOfMemory. See [VMware Knowledge Base Article 2135378](#) for more information.

All versions of VMware Tools are supported. Some Guest Introspection-based features require newer VMware Tools versions:

Guest Introspection for Windows

- Use VMware Tools 10.0.9 and 10.0.12 to enable the optional Thin Agent Network Introspection Driver component packaged with VMware Tools.
- Upgrade to VMware Tools 10.0.8 and later to resolve slow VMs after upgrading VMware Tools in NSX / vCloud Networking and Security (see [VMware knowledge base article 2144236](#)).
- Use VMware Tools 10.1.0 and later for Windows 10 support.
- Use VMware Tools 10.1.10 and later for Windows Server 2016 support.

This NSX version supports the following Linux versions:

Guest Introspection for Linux

- RHEL 7 GA (64 bit)
- SLES 12 GA (64 bit)
- Ubuntu 14.04 LTS (64 bit)

Note: VMware currently does not support NSX for vSphere 6.3.x with vRealize Networking Insight 3.2.

System Requirements and Installation

For the complete list of NSX installation prerequisites, see the [System Requirements for NSX](#) section in the *NSX Installation Guide*.

For installation instructions, see the [NSX Installation Guide](#) or the [NSX Cross-vCenter Installation Guide](#).

Deprecated and Discontinued Functionality

End of Life and End of Support Warnings

For information about NSX and other VMware products that must be upgraded soon, please consult the [VMware Lifecycle Product Matrix](#).

- **NSX for vSphere 6.1.x** reached End of Availability (EOA) and End of General Support (EOGS) on January 15, 2017. (See also [VMware knowledge base article 2144769](#).)
- **NSX Data Security removed:** As of NSX 6.3.0, the NSX Data Security feature has been removed from the product.
- **NSX Activity Monitoring (SAM) deprecated:** As of NSX 6.3.0, Activity Monitoring is no longer a supported feature of NSX. As a replacement, please use Endpoint Monitoring. For more information see [Endpoint Monitoring](#) in the *NSX Administration Guide*.
- **Web Access Terminal removed:** Web Access Terminal (WAT) has been removed from NSX 6.3.0. You cannot configure Web Access SSL VPN-Plus and enable the public URL access through NSX Edge. VMware recommends using the full access client with SSL VPN deployments for improved security. If you are using WAT functionality in an earlier release, you must disable it before upgrading to 6.3.0.
- **IS-IS removed from NSX Edge:** From NSX 6.3.0, you cannot configure IS-IS Protocol from the **Routing** tab.
- **vCNS Edges no longer supported.** You must upgrade to an NSX Edge first before upgrading to NSX 6.3.x.

API Removals and Behavior Changes

Changes in API error handling

NSX 6.3.5 introduces these changes in error handling:

- If an API request results in a database exception on the NSX Manager, the response is *500 Internal server error*. In previous releases, NSX Manager responded with *200 OK*, even though the request failed.
- If you send an API request with an empty body when a request body is expected, the response is *400 Bad request*. In previous releases NSX Manager responded with *500 Internal server error*.
- If you specify an incorrect security group in this API, GET `/api/2.0/services/policy/securitygroup/{ID}/securitypolicies`, the response is *404 Not found*. In previous releases NSX Manager responded with *200 OK*.

Changes in backup and restore API defaults

Starting in 6.3.3, the defaults for two backup and restore parameters have changed to match the defaults in the UI. Previously **passiveMode** and **useEPSV** defaulted to *false*, now they default to *true*. This affects the following APIs:

- PUT `/api/1.0/appliance-management/backuprestore/backupsettings`
- PUT `/api/1.0/appliance-management/backuprestore/backupsettings/ftpsettings`

Deleting firewall configuration or default section

- Starting in 6.3.0, this request is rejected if the default section is specified: DELETE `/api/4.0/firewall/globalroot-0/config/layer2sections|layer3sections/sectionId`

- A new method is introduced to get default configuration. Use the output of this method to replace entire configuration or any of the default sections:
 - Get default configuration with GET /api/4.0/firewall/globalroot-0/defaultconfig
 - Update entire configuration with PUT /api/4.0/firewall/globalroot-0/config
 - Update single section with PUT /4.0/firewall/globalroot-0/config/layer2sections|layer3sections/{sectionId}

defaultOriginate parameter:

Starting in NSX 6.3.0, the defaultOriginate parameter is removed from the following methods for logical (distributed) router NSX Edge appliances only:

- GET/PUT /api/4.0/edges/{edge-id}/routing/config/ospf
- GET/PUT /api/4.0/edges/{edge-id}/routing/config/bgp
- GET/PUT /api/4.0/edges/{edge-id}/routing/config

Setting defaultOriginate to true on an NSX 6.3.0 or later logical (distributed) router edge appliance will fail.

All IS-IS methods removed from NSX Edge routing

- GET/PUT/DELETE /4.0/edges/{edge-id}/routing/config/isis
- GET/PUT /4.0/edges/{edge-id}/routing/config

CLI Removals and Behavior Changes

Do not use unsupported commands on NSX Controller nodes

There are undocumented commands to configure NTP and DNS on NSX Controller nodes. These commands are not supported, and should not be used on NSX Controller nodes. You should only use commands which are documented in the NSX CLI Guide.

Upgrade Notes

- [General Upgrade Notes](#)
- [Upgrade Notes for NSX Components](#)
- [Upgrade Notes for FIPS](#)

Note: For a list of known issues affecting installation and upgrades see the section [Installation and Upgrade Known Issues](#).

General Upgrade Notes

- To upgrade NSX, you must perform a full NSX upgrade including host cluster upgrade (which upgrades the host VIBs). For instructions, see the [NSX Upgrade Guide](#) including the [Upgrade Host Clusters](#) section.
- **System Requirements:** For information on system requirements while installing and upgrading NSX, see the [System Requirements for NSX](#) section in NSX documentation.

- **Upgrade path from NSX 6.x:** The [VMware Product Interoperability Matrix](#) provides details about the upgrade paths from VMware NSX.
- **Cross-vCenter NSX upgrade** is covered in the [NSX Upgrade Guide](#).
- **Downgrades are not supported:**
 - Always capture a backup of NSX Manager before proceeding with an upgrade.
 - Once NSX has been upgraded successfully, NSX cannot be downgraded.
- **To validate** that your upgrade to NSX 6.3.x was successful see [knowledge base article 2134525](#).
- There is no support for upgrades from vCloud Networking and Security to NSX 6.3.x. You must upgrade to a supported 6.2.x release first.
- **Interoperability:** Check the [VMware Product Interoperability Matrix](#) for all relevant VMware products before upgrading.
 - **Upgrading to vSphere 6.5a or later:** When upgrading from vSphere 5.5 or 6.0 to vSphere 6.5a or later, you must first upgrade to NSX 6.3.x. See [Upgrading vSphere in an NSX Environment](#) in the *NSX Upgrade Guide*.
Note: NSX 6.2.x is not compatible with vSphere 6.5.
 - **Upgrading to NSX 6.3.3 or later:** The minimum supported version of vSphere for NSX interoperability changes between NSX 6.3.2 and NSX 6.3.3. See the [VMware Product Interoperability Matrix](#) for details.
- **Partner services compatibility:** If your site uses VMware partner services for Guest Introspection or Network Introspection, you must review the [VMware Compatibility Guide](#) before you upgrade, to verify that your vendor's service is compatible with this release of NSX.
- **Networking and Security plug-in:** After upgrading NSX Manager, you must log out and log back in to the vSphere Web Client. If the NSX plug-in does not display correctly, clear your browser cache and history. If the Networking and Security plug-in does not appear in the vSphere Web Client, reset the vSphere Web Client server as explained in the [NSX Upgrade Guide](#).
- **Stateless environments:** In NSX upgrades in a stateless host environment, the new VIBs are pre-added to the Host Image profile during the NSX upgrade process. As a result, NSX on stateless hosts upgrade process follows this sequence:
 Prior to NSX 6.2.0, there was a single URL on NSX Manager from which VIBs for a certain version of the ESX Host could be found. (Meaning the administrator only needed to know a single URL, regardless of NSX version.) In NSX 6.2.0 and later, the new NSX VIBs are available at different URLs. To find the correct VIBs, you must perform the following steps:
 1. Find the new VIB URL from `https://<NSXManager>/bin/vdn/nwfabric.properties`.
 2. Fetch VIBs of required ESX host version from corresponding URL.
 3. Add them to host image profile.

Upgrade Notes for NSX Components

NSX Manager Upgrade

- **Important:** If you are upgrading NSX 6.2.0, 6.2.1, or 6.2.2 to NSX 6.3.5, you must complete a workaround before starting the upgrade. See [VMware Knowledge Base article 000051624](#) for details.
- If you use SFTP for NSX backups, change to hmac-sha2-256 after upgrading to 6.3.x because there is no support for hmac-sha1. See [VMware Knowledge Base article 2149282](#) for a list of supported security algorithms in 6.3.x.
- If you want to upgrade from NSX 6.3.3 to NSX 6.3.4 or later you must first follow the workaround instructions in [VMware Knowledge Base article 2151719](#).

Controller Upgrade

- In NSX 6.3.3, the NSX Controller appliance disk size changes from 20GB to 28GB.
- The NSX Controller cluster must contain three controller nodes to upgrade to NSX 6.3.3. If it has fewer than three controllers, you must add controllers before starting the upgrade. See [Deploy NSX Controller Cluster](#) for instructions.
- In NSX 6.3.3, the underlying operating system of the NSX Controller changes. This means that when you upgrade from NSX 6.3.2 or earlier to NSX 6.3.3 or later, instead of an in-place software upgrade, the existing controllers are deleted one at a time, and new Photon OS based controllers are deployed using the same IP addresses.

When the controllers are deleted, this also deletes any associated DRS anti-affinity rules. You must create new anti-affinity rules in vCenter to prevent the new controller VMs from residing on the same host.

See [Upgrade the NSX Controller Cluster](#) for more information on controller upgrades.

Host Cluster Upgrade

- In NSX 6.3.3, NSX VIB names change. The esx-vxlan and esx-vsip VIBs are replaced with esx-nsxv if you have NSX 6.3.3 or later.
- **Rebootless upgrade and uninstall on hosts:** On vSphere 6.0 and later, once you have upgraded to NSX 6.3.x, any subsequent NSX VIB changes will not require a reboot. Instead hosts must enter maintenance mode to complete the VIB change.

A host reboot **is not required** during the following tasks:

- NSX 6.3.0 to NSX 6.3.x upgrades on ESXi 6.0 or later.
- The NSX 6.3.x VIB install that is required after upgrading ESXi from 6.0 to 6.5.0a or later.
Note: The ESXi upgrade still requires a host reboot.
- NSX 6.3.x VIB uninstall on ESXi 6.0 or later.

A host reboot **is required** during the following tasks:

- NSX 6.2.x or earlier to NSX 6.3.x upgrades (any ESXi version).
- NSX 6.3.0 to NSX 6.3.x upgrades on ESXi 5.5.
- The NSX 6.3.x VIB install that is required after upgrading ESXi from 5.5 to 6.0 or later.

- NSX 6.3.x VIB uninstall on ESXi 5.5.
- **Host may become stuck in the installing state** During large NSX upgrades, a host may become stuck in the installing state for a long time. This can occur due to issues uninstalling old NSX VIBs. In this case the EAM thread associated with this host will be reported in the VI Client Tasks list as stuck.

Workaround: Do the following:

- Log into vCenter using the VI Client.
- Right click on the stuck EAM task and cancel it.
- From the vSphere Web Client, issue a Resolve on the cluster. The stuck host may now show as InProgress.
- Log into the host and issue a reboot to force completion of the upgrade on that host.

NSX Edge Upgrade

- In NSX 6.3.0, the NSX Edge appliance disk sizes have changed:
 - **Compact, Large, Quad Large:** 1 disk 584MB + 1 disk 512MB
 - **XLarge:** 1 disk 584MB + 1 disk 2GB + 1 disk 256MB
- **Host clusters must be prepared for NSX before upgrading NSX Edge appliances**
Management-plane communication between NSX Manager and Edge via the VIX channel is no longer supported starting in 6.3.0. Only the message bus channel is supported. When you upgrade from NSX 6.2.x or earlier to NSX 6.3.0 or later, you must verify that host clusters where NSX Edge appliances are deployed are prepared for NSX, and that the messaging infrastructure status is GREEN. If host clusters are not prepared for NSX, upgrade of the NSX Edge appliance will fail. See [Upgrade NSX Edge](#) in the *NSX Upgrade Guide* for details.

- **Upgrading Edge Services Gateway (ESG):**

Starting in NSX 6.2.5, resource reservation is carried out at the time of NSX Edge upgrade. When vSphere HA is enabled on a cluster having insufficient resources, the upgrade operation may fail due to vSphere HA constraints being violated.

To avoid such upgrade failures, perform the following steps before you upgrade an ESG:

The following resource reservations are used by the NSX Manager if you have not explicitly set values at the time of install or upgrade.

NSX Edge Form Factor	CPU Reservation	Memory Reservation
COMPACT	1000MHz	512 MB
LARGE	2000MHz	1024 MB
QUADLARGE	4000MHz	2048 MB
X-LARGE	6000MHz	8192 MB

1. Always ensure that your installation follows the best practices laid out for vSphere

HA. Refer to document [Knowledge Base article 1002080](#).

2. Use the NSX tuning configuration API:

PUT `https://<NSXManager>/api/4.0/edgePublish/tuningConfiguration` ensuring that values for `edgeVCpuReservationPercentage` and `edgeMemoryReservationPercentage` fit within available resources for the form factor (see table above for defaults).

- **Disable vSphere's Virtual Machine Startup option where vSphere HA is enabled and Edges are deployed.** After you upgrade your 6.2.4 or earlier NSX Edges to 6.2.5 or later, you must turn off the vSphere Virtual Machine Startup option for each NSX Edge in a cluster where vSphere HA is enabled and Edges are deployed. To do this, open the vSphere Web Client, find the ESXi host where NSX Edge virtual machine resides, click Manage > Settings, and, under Virtual Machines, select VM Startup/Shutdown, click Edit, and make sure that the virtual machine is in Manual mode (that is, make sure it is not added to the Automatic Startup/Shutdown list).
- **Before upgrading to NSX 6.2.5 or later, make sure all load balancer cipher lists are colon separated.** If your cipher list uses another separator such as a comma, make a PUT call to `https://nsxmgr_ip/api/4.0/edges/EdgeID/loadbalancer/config/applicationprofiles` and replace each `<ciphers>` list in `<clientSsl>` and `<serverSsl>` with a colon-separated list. For example, the relevant segment of the request body might look like the following. Repeat this procedure for all application profiles:

```
<applicationProfile>
  <name>https-profile</name>
  <insertXForwardedFor>false</insertXForwardedFor>
  <sslPassthrough>false</sslPassthrough>
  <template>HTTPS</template>
  <serverSslEnabled>true</serverSslEnabled>
  <clientSsl>
    <ciphers>AES128-SHA:AES256-SHA:ECDHE-ECDSA-AES256-SHA</ciphers>
    <clientAuth>ignore</clientAuth>
    <serviceCertificate>certificate-4</serviceCertificate>
  </clientSsl>
  <serverSsl>
    <ciphers>AES128-SHA:AES256-SHA:ECDHE-ECDSA-AES256-SHA</ciphers>
    <serviceCertificate>certificate-4</serviceCertificate>
  </serverSsl>
  ...
</applicationProfile>
```

- **Set Correct Cipher version for Load Balanced Clients on vROPs versions older than 6.2.0:** vROPs pool members on vROPs versions older than 6.2.0 use TLS version 1.0 and therefore you must set a monitor extension value explicitly by setting "ssl-version=10" in the NSX Load Balancer configuration. See [Create a Service Monitor](#) in the *NSX Administration Guide* for instructions.

```
{
  "expected" : null,
  "extension" : "ssl-version=10",
  "send" : null,
  "maxRetries" : 2,
  "name" : "sm_vrops",
```

```
"url" : "/suite-api/api/deployment/node/status",
"timeout" : 5,
"type" : "https",
"receive" : null,
"interval" : 60,
"method" : "GET"
}
```

Guest Introspection Upgrade

- Guest Introspection VM's now contain additional host identifying information in an XML file on the machine. When logging in to the Guest Introspection VM, the file `/opt/vmware/etc/vami/ovfEnv.xml` should include host identity information.

Upgrade Notes for FIPS

When you upgrade from a version of NSX earlier than NSX 6.3.0 to NSX 6.3.0 or later, you must not enable FIPS mode before the upgrade is completed. Enabling FIPS mode before the upgrade is complete will interrupt communication between upgraded and not-upgraded components. See [Understanding FIPS Mode and NSX Upgrade](#) in the *NSX Upgrade Guide* for more information.

- **Ciphers supported on OS X Yosemite and OS X El Capitan:** If you are using SSL VPN client on OS X 10.11 (El Capitan), you will be able to connect using AES128-GCM-SHA256, ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-GCM-SHA38, AES256-SHA and AES128-SHA ciphers and those using OS X 10.10 (Yosemite) will be able to connect using AES256-SHA and AES128-SHA ciphers only.
- Do not enable FIPS before the upgrade to NSX 6.3.x is complete. See [Understand FIPS mode and NSX Upgrade](#) in the *NSX Upgrade Guide* for more information.
- Before you enable FIPS, verify any partner solutions are FIPS mode certified. See the [VMware Compatibility Guide](#) and the relevant partner documentation.

FIPS Compliance

- **NSS and OpenSwan:** The NSX Edge IPsec VPN uses the Mozilla NSS crypto module. Due to critical security issues, this version of NSX uses a newer version of NSS that has not been FIPS 140-2 validated. VMware affirms that the module works correctly, but it is no longer formally validated.
- **NSS and Password Entry:** The NSX Edge password hashing use the Mozilla NSS crypto module. Due to critical security issues, this version of NSX uses a newer version of NSS that has not been FIPS 140-2 validated. VMware affirms that the module works correctly, but it is no longer formally validated.
- **Controller and Clustering VPN:** The NSX Controller uses IPsec VPN to connect Controller clusters. The IPsec VPN uses the VMware Linux kernel crypto module (Photon 1 environment), which is in the process of being CMVP validated.

Document Revision History

16 Nov 2017: First edition.

17 Nov 2017: Second edition. Added known issue 2000749.

28 Nov 2017: Third edition. Adding upgrade information for NSX 6.2.0, 6.2.1, 6.2.2.

1 Dec 2017: Fourth edition. Added resolved issues 1937124 and 1976332.

8 Dec 2017: Fifth edition. Added resolved issues 1790951 and 1935204.

8 Jan 2018: Sixth edition. Added resolved issue 1920574.

29 Mar 2018: Seventh edition. Added resolved issues 1967608 and 1947687.

13 May 2019: Eighth edition. Updated Host Cluster Upgrade section.

Resolved Issues

The resolved issues are grouped as follows.

- [General Resolved Issues](#)
- [Logical Networking and Edge Resolved Issues](#)
- [NSX Manager Resolved Issues](#)
- [NSX Controller Resolved Issues](#)
- [Security Services Resolved Issues](#)
- [Installation and Upgrade](#)

General Resolved Issues

- **Fixed Issue 1293896: VMs migrated from 6.0.x can cause host PSOD**
When upgrading a cluster from 6.0.x to 6.2.3-6.2.8 or 6.3.x, the VM state exported can be corrupted and cause the receiving host to PSOD. *Fixed in 6.3.5*
- **Fixed Issue 1952277: GI USVM does not receive IP address from pool**
Dynamic IP assigned to GI-SVM eth0 nic given by manager from the ip pool causes error on GI-SVM boot with nic stating that the given address is already taken. *Fixed in 6.3.5*
- **Fixed Issue 1918023: Guest Introspection USVM consumes 100% memory**
Guest Introspection USVM consumes 100% memory and might result in guest VMs losing connectivity to Guest Introspection USVM. *Fixed in 6.3.5*
- **Fixed Issue 1912443: Potential security risk by not running latest Tomcat version.**
Older version of NSX Manager Tomcat Service has security vulnerabilities *Fixed in 6.3.5*
- **Fixed Issue 1920032: SynFlood protection causes TCP timestamp corruption leading to cloning and cold migration failures**
When you initiate cold migration of VMs between two ESXi hosts or create a full clone of the VM, SynFlood protection causes TCP timestamp corruption. *Fixed in 6.3.5.*
- **Fixed Issue 1920343: Server certificate can be created without a private key**
If the Private Key data is provided in the Certificate Content, the private key is ignored. *Fixed in 6.3.5.*
- **Fixed Issue 1874735: "Upgrade Available" link not shown if cluster has an alarm.**
Users are not be able to push the new service spec to EAM because the link is missing and the service will not be upgraded. *Fixed in 6.3.5*

- **Fixed Issue 1926060: The Negate source checkbox on the Firewall > Specify Source or Destination page gets selected even when you click outside**

The checkbox for Negate source gets selected when you move objects from the list of Available objects to the Selected objects. *Fixed in 6.3.5.*

- **Fixed Issue 1790951: Guest Introspection deployment might fail if the first 20 characters of the cluster names were identical**

Failure in Guest Introspection deployment at the same time on multiple cluster with identical names.

- **Fixed Issue 1947687: NSX VIB installation may fail if DVS config has non-ASCII characters**

Automatic host prep may fail for customers running ESX 6.5 or above if their DVS configuration (output of "net-dvs -l") contains non-ASCII characters.

Workaround: Avoid using non-ASCII characters for DVS port names. *Fixed in 6.3.4*

Logical Networking and Edge Resolved Issues

- **Issue 1967608: Brief interruption to the bridged data path may be observed when the toragent or NSX Controller restarts**

In a environment where a hardware gateway is deployed, a brief interruption to the bridged data path may be observed when the toragent or NSX Controller restarts. In an impacted environment, Controller restarts should only occur in a planned maintenance window.

Fixed in 6.3.5.

- **Fixed Issue 1976332: NSX controller does not store the MAC entry of the Workload VM on the ESXi where active L2-bridge control VM is running**

Traffic drops for all workloads installed on the Hypervisor where the Active-Bridging Control VM is running. *Fixed in 6.3.5.*

- **Fixed Issue 1937124: VXLAN-VLAN bridged traffic disrupted after vMotion of Distributed Logical Router appliance (control VM)**

Bridges do not update physical network with VXLAN VM MACs under a vMotion or failover event when multiple DLRs are configured to do bridging. See VMware Knowledge Base article [2151647](#) for more information and a workaround. *Fixed in 6.3.5.*

- **Fixed Issue 1922967: Edge IPsec VPN goes down when the peer end ip address keep changing.**

The problem is seen in deployments where VPN peer IP address keep changing (3G dongle / Dynamic Wan ip address). After the peer IP change, when the tunnel goes down, route for peer subnet is not flushed properly. Due to this when the tunnel is renegotiated IPSec SA installation fails at the Edge. *Fixed in 6.3.5.*

- **Fixed Issue 1957065: DNS forwarder address stripped from ESG configuration after REST API PUT configuration**

DNS forwarder address stripped from ESG configuration after REST API PUT configuration. *Fixed in 6.3.5*

- **Fixed Issue 1916360: HA failover may fail due to full disk when >100 routes are installed.**

When more than 100 routes are installed, the vmtools daemon on the standby Edge will submit 2 warning log messages every 30s. The logs are saved to a file called `/var/log/vmware-vmtoolsd.log`, which can eventually grow to completely fill the log partition. Log rotation is not configured for this log file. When this occurs, HA failover may fail. *Fixed in 6.3.5*

- **Fixed Issue 1916580: Edge does not allow any operation with an error statement "certificate-xx doesn't contain private key".**

When upgrading NSX Manager from 6.2.x (having wrongly configured service certificate) to 6.3.x, the edge cannot be upgraded with an error "certificate-xx doesn't contain private key". A service certificate is wrongly configured in NSX as a CA Certificate (the private key is added as part of Certificate). *Fixed in 6.3.5.*

- **Fixed Issue 1973130: North-South (N-S) traffic loss happens when distributed logical router control VM goes through an HA failover even if there is a floating static route installed on the Edge Services Gateway**

The N-S traffic loss happens during control VM HA failover, even if there is a floating static route installed on the Edge Services Gateway towards distributed logical router. *This is fixed in 6.3.5.*

- **Fixed Issue 1972659: NSX Edge Interface shows datastores are both null and valid on a single DLR appliance instance**

After deleting one of the DataStore the NSX Edges shows Configured DataStore as null. *Fixed in 6.3.5*

- **Fixed Issue 1983497: Purple screen appears when bridge failover and bridge configuration change happens at the same time**

When a bridge failover and bridge configuration change happens at the same time, it may result into deadlock and cause purple screen. The chances running into deadlock is low. *Fixed in 6.3.5.*

- **Fixed Issue 1897999: Traffic from VTEP fails when VTEP using LACP and first uplink in LAG goes down.**

Traffic from VTEP fails, for example, cannot get IP address if using DHCP, first uplink in LAG goes down. *Fixed in 6.3.5.*

- **Fixed Issue 1888743: Cannot set IPv6 default gateway on NSX Edge, if present in earlier NSX version, upgrade fails**

In NSX 6.3.3 and 6.3.4 you cannot create an IPv6 default gateway. If you upgrade to NSX 6.3.3 or NSX 6.3.4, and have an IPv6 default gateway set on NSX Edge, the upgrade will fail. Error message seen is "Failed to make IPv6 static routing changes". *Fixed in 6.3.5.*

- **Fixed Issue 1935204: DLR takes 1 to 1.5 secs for ARP resolution**

When ARP suppression fails, local DLR's ARP resolution for a VM running in a remote host takes around 1 to 1.5 sec. *Fixed in 6.3.5.*

- **Fixed Issue 1920574: Unable to configure sub interfaces for an Edge**

Creating sub interfaces on Edge fails with NSX for vSphere 6.3.2/6.3.3 (unable to publish the sub interface with IP).

NSX Manager Resolved Issues

- **Fixed Issue 1891547: NSX Manager does not connect to Event Log Servers after multiple reboots of Event log servers.**
Restarting the Event log server numerous times causes the NSX Manager not to reconnect to the event log server. *Fixed in 6.3.5*
- **Fixed Issue 1931288: NSX Manager crashes with high NSX Manager CPU**
NSX Manager has an OOM (out of memory) error and continuously restarts. *Fixed in 6.3.5*
- **Fixed Issue 1926309: NSX Manager Plug-in is not able to load and shows "Authentication Exception"**
Sometimes NSX Manager Plug-in is unable to load any pages and eventually displaying timeout error. *Fixed in 6.3.5.*
- **Fixed Issue 1904842: NSX Manager is not registering with vCenter or Platform Service Controller**
NSX Manager is not appearing on UI and any REST call to the NSX Manager fails.
- **Fixed Issue 1900144: Request to include specific changes made to SGs in audit logs**
NSX Manager can now successfully forward syslog messages to LogInsight and includes the details that are visible on the vSphere Web Client's GUI. *Fixed in 6.3.5.*

NSX Controller Resolved Issues

- **Fixed Issue 1898862: Controller High CPU Utilization caused by hardware VTEP**
User observes controller in high CPU and connection between Controller TOR Manager and Switch is down. When checking ToR agent log, the log file has multiple lines of following messages:
2017-05-03 17:13:18,991 | DEBUG | pool-9-thread-5 | OvsdbConnectionService | Handshake status NEED_UNWRAP. *Fixed in 6.3.5.*
- **Fixed Issue 1898937: ESXi Host does not have full vTEP lists for VNIs causing East-West connectivity issue**
Race condition which only happens when controller cluster has sharding change, and legacy master controller handles message slower than the session message comes down after new controller node complete state sync. This leads to unsynchronized changes between controllers for the affected VNI. *Fixed in 6.3.5.*
- **Fixed Issue 1965859: NSX Controller memory increases with hardware VTEP configuration causing high CPU usage**
A controller process memory increase is seen with hardware VTEP configurations running for few days. The memory increase causes high CPU usage that lasts for some time (minutes) while the controller recovers the memory. During this time the data path is affected. *Fixed in 6.3.5.*

Security Services Resolved Issues

- **Fixed Issue 1897878: ESXi tasks and events displays error "Lost Communication with ESX module"**
If the ESXi host Guest Introspection module (EPSec Mux) loses communication with the

ESX module, the error "Lost Communication with ESX module" appears on the ESXi hosts.
Fixed in 6.3.5.

- **Fixed Issue 1945954: When upgrading, DFW is auto-enabled on DFW disabled cluster**

Newly added hosts in a cluster have the firewall enabled, resulting in traffic disruption.
Fixed in 6.3.5

- **Fixed Issue 1951626: A user with the Auditor role can disable the DFW at cluster level on the Host Preparation tab**

A user with Auditor role should have read only permissions on NSX and not have the ability to disable the DFW. *Fixed in 6.3.5*

- **Fixed Issue 1944599: Translated IPs are not getting added to vNIC filters which is causing Distributed Firewall to drop traffic**

When new VMs are deployed, the vNIC filters do not get updated with the right set of IPs causing Distributed Firewall to block the traffic. *Fixed in 6.3.5.*

- **Fixed Issue 1958657: Services aren't seen in the servicegroup in firewall UI**

After creating a security service with ICMP and not defining a sub-protocol, service are not seen in the firewall UI. *Fixed in 6.3.5*

Installation and Upgrade

- **Fixed Issue 1789989: During a host cluster upgrade, you may experience packet loss in the data plane**

During the VIB upgrade, the password file of VSFWD (vShield Firewall Daemon), which is kept in the VIB is removed so VSFWD cannot use the old password to connect to the NSX Manager and has to wait until the new password is updated. This process takes some time to complete after host reboot, however, in a fully automated DRS cluster VMs are moved immediately once the prepped host comes up and because the VSFWD process is not ready at that time, there is a chance of packet loss in the data plane for a brief time. *Fixed in 6.3.5.*

Known Issues

The known issues are grouped as follows.

- [General Known Issues](#)
- [Installation and Upgrade Known Issues](#)
- [NSX Manager Known Issues](#)
- [NSX Controller Known Issues](#)
- [Logical Networking and NSX Edge Known Issues](#)
- [Security Services Known Issues](#)
- [Monitoring Services Known Issues](#)

General Known Issues

- **Issue 2003765: TOR Manager on NSX Controller fails to send update when the physical TOR device is reset/rebooted or power cycled**

VM remote macs are missing on TOR OVSDDB table if TOR is reloaded.

Workaround: Reboot all NSX Controllers. See VMware Knowledge Base article [52074](#) for more information.

- **Issue 1874863: Unable to authenticate with changed password after sslvpn service disable/enable with local authentication server**

When SSL VPN service is disabled and re-enabled and when using local authentication, users are unable to log in with changed passwords.

See [VMware Knowledge Base Article 2151236](#) for more information.

- **Issue 1702339: Vulnerability scanners might report a Quagga bgp_dump_routes vulnerability CVE-2016-4049**

Vulnerability scanners might report a Quagga bgp_dump_routes vulnerability CVE-2016-4049 in NSX for vSphere. NSX for vSphere uses Quagga, but the BGP functionality (including the vulnerability) is not enabled. This vulnerability alert can be safely disregarded.

Workaround: As the product is not vulnerable, no workaround is required.

- **Issue 1740625, 1749975: UI problems on Mac OS in Firefox and Safari**

If you are using Firefox or Safari in Mac OS, the Back navigation button will not work in NSX Edge from the Networking and Security page in the vSphere 6.5 Web Client, and sometimes the UI freezes in Firefox.

Workaround: Use Google Chrome on Mac OS or click on the Home button then proceed as required.

- **Issue 1700980: For security patch CVE-2016-2775, a query name which is too long can cause a segmentation fault in lwresd**

NSX 6.2.4 has BIND 9.10.4 installed with the product, but it does not use lwres option in *named.conf*, hence the product is not vulnerable.

Workaround: As the product is not vulnerable, no workaround is required.

- **Issue 1568180: Feature list incorrect for NSX when using vCenter Server Appliance (vCSA) 5.5**

You can view the features of a license in the vSphere Web Client by selecting the license and clicking **Actions > View Features**. If you upgrade to NSX 6.2.3, your license is upgraded to an Enterprise license, which enables all features. However, if NSX Manager is registered with vCenter Server Appliance (vCSA) 5.5, selecting **View Features** will display the list of features for the license used before the upgrade, not the new Enterprise license.

Workaround: All Enterprise licenses have the same features, even if they are not displayed correctly in the vSphere Web Client. See the [NSX Licensing Page](#) for more information.

Installation and Upgrade Known Issues

Before upgrading, please read the section [Upgrade Notes](#), earlier in this document.

- **Issue 2001988: During NSX host cluster upgrade, Installation status in Host**

Preparation tab alternates between "not ready" and "installing" for the entire cluster when each host in the cluster is upgrading

During NSX upgrade, clicking "upgrade available" for NSX prepared cluster triggers host upgrade. For clusters configured with DRS FULL AUTOMATIC, the installation status alternates between "installing" and "not ready", even though the hosts are upgraded in the background without issues.

Workaround: This is a user interface issue and can be ignored. Wait for the host cluster upgrade to proceed.

- **Issue 1932907: Upgrade of Guest Introspection SVM Failed**

While trying to upgrade the Guest Introspection SVM, the installation status for the GI SVM is 'Failed'. This might be applicable for GI-SVMs of one or multiple hosts in the cluster.

Workaround:

1. Delete the GI-SVM from the VC.
2. Click **Resolve** in the GI-SVM Service deployment pane. This will re-deploy the GI-SVM.

- **Issue 1848058: Upgrade of ESXi host VIBs to NSX 6.3.2 might fail**

In some cases, during an upgrade of ESXi host VIBs to NSX 6.3.2, the older VIB directory is deleted from NSX Manager, causing the upgrade to fail. Clicking the Resolve button does not fix the issue.

Workaround: To recover from this, use the upgrade API:

PUT <https://<nsx-mgr-ip>/api/2.0/nwfabric/configure>

```
<nwFabricFeatureConfig>
  <featureId>com.vmware.vshield.vsm.vxlan</featureId>
  <resourceConfig>
    <resourceId>domain-cXX</resourceId>
  </resourceConfig>
</nwFabricFeatureConfig>
```

where <nsx-mgr-ip> is the IP address of your NSX Manager, and domain-cXX is the domain id of the cluster.

- **Issue 1747217: Preparing ESXi hosts results in muxconfig.xml.bad and Guest Introspection does not function correctly**

If the "vmx path" is missing for one of the VMs in muxconfig.xml, when MUX tries to parse the config file and doesn't find the "xml path" property, it renames the config file as "muxconfig.xml.bad", sends the error "Error - MUX Parsing config" to the USVM and closes the config channel.

Workaround: Remove the orphaned VMs from the vCenter inventory.

- **Issue 1859572: During the uninstall of NSX VIBs version 6.3.x on ESXi hosts that are being managed by vCenter version 6.0.0, the host continues to stay in Maintenance mode**

If you are uninstalling NSX VIBs version 6.3.x on a cluster, the workflow involves putting the hosts into Maintenance mode, uninstalling the VIBs and then removing the hosts from Maintenance mode by the EAM service. However, if such hosts are managed by vCenter server version 6.0.0, then this results in the host being stuck in Maintenance mode post uninstalling the VIBs. The EAM service responsible for uninstalling the VIBs puts the host

in Maintenance mode but fails to move the hosts out of Maintenance mode.

Workaround: Manually move the host out of Maintenance mode. This issue will not be seen if the host is managed by vCenter server version 6.5a and above.

- **Issue 1435504: HTTP/HTTPS Health Check appears DOWN after upgrading from 6.0.x or 6.1.x to 6.3.x with failure reason "Return code of 127 is out of bounds - plugin may be missing"**

In NSX 6.0.x and 6.1.x releases, URLs configured without double quotes (") caused Health Check to fail with this error: "Return code of 127 is out of bounds - plugin may be missing". The workaround for this was to add the double quotes (") to the input URL (this was not required for send/receive/expect fields). However, this issue was fixed in 6.2.0 and as a result, if you are upgrading from 6.0.x or 6.1.x to 6.3.x, the additional double quotes result in the pool members shown as DOWN in Health Check.

Workaround: Remove the double quotes (") in the URL field from all the relevant Health Check configurations after upgrading.

- **Issue 1734245: Data Security causes upgrades to 6.3.0 to fail**
Upgrades to 6.3.0 will fail if Data Security is configured as part of a service policy. Ensure you remove Data Security from any service policies before upgrading.
- **Issue 1801685: Unable to see filters on ESXi after upgrade from 6.2.x to 6.3.0 because of failure to connect to host**

After you upgrade from NSX 6.2.x to 6.3.0 and cluster VIBs to 6.3.0 bits, even though the installation status shows successful and Firewall Enabled, the "communication channel health" will show the NSX Manager to Firewall Agent connectivity and NSX Manager to ControlPlane Agent connectivity as down. This will lead to Firewall rules publish, Security Policy publish failures and VXLAN configuration not being sent down to hosts.

Workaround: Run the message bus sync API call for the cluster using the APIPOST <https://<NSX-IP>/api/2.0/nwfabric/configure?action=synchronize>.

API Body:

```
<nwFabricFeatureConfig>
  <featureId>com.vmware.vshield.vsm.messagingInfra</featureId>
  <resourceConfig>
    <resourceId>{Cluster-MOID}</resourceId>
  </resourceConfig>
</nwFabricFeatureConfig>
```

- **Issue 1797929: Message bus channel down after host cluster upgrade**
After a host cluster upgrade, vCenter 6.0 (and earlier) does not generate the event "reconnect", and as a result, NSX Manager does not set up the messaging infrastructure on the host. This issue has been fixed in vCenter 6.5.

Workaround: Resync the messaging infrastructure as below:

POST <https://<ip>/api/2.0/nwfabric/configure?action=synchronize>

```
<nwFabricFeatureConfig>
  <featureId>com.vmware.vshield.vsm.messagingInfra</featureId>
  <resourceConfig>
    <resourceId>host-15</resourceId>
  </resourceConfig>
```

- **Issue 1768144: Old NSX Edge appliance resource reservations that exceed new limits may cause failure during upgrade or redeployment**

In NSX 6.2.4 and earlier, you could specify an arbitrarily large resource reservation for an NSX Edge appliance. NSX did not enforce a maximum value. After NSX Manager is upgraded to 6.2.5 or later, if an existing Edge has resources reserved (especially memory) that exceed the newly enforced maximum value imposed for the chosen form factor, it would fail during Edge upgrade or redeploy (which would trigger an upgrade). For example, if the user has specified a memory reservation of 1000MB on a pre-6.2.5 LARGE Edge and, after upgrade to 6.2.5, changes the appliance size to COMPACT, the user-specified memory reservation will exceed the newly enforced maximum value, in this case 512 for a COMPACT Edge, and the operation will fail.

See [Upgrading Edge Service Gateway \(ESG\)](#) for information on recommended resource allocation starting in NSX 6.2.5.

Workaround: Use the appliance REST API: PUT <https://<NSXManager>/api/4.0/edges/<edge-id>/appliances/> to reconfigure the memory reservation to be within values specified for the form factor, without any other appliance changes. You can change the appliance size after this operation completes.

- **Issue 1600281: USVM Installation Status for Guest Introspection shows as Failed in the Service Deployments tab**

If the backing datastore for the Guest Introspection Universal SVM goes offline or becomes inaccessible, the USVM may need to be rebooted or re-deployed to recover.

Workaround: Reboot or re-deploy USVM to recover.

- **Issue 1660373: vCenter enforces expired NSX license**

As of vSphere 5.5 update 3 or vSphere 6.0.x vSphere Distributed Switch is included in the NSX license. However, vCenter does not allow ESX hosts to be added to a vSphere Distributed Switch if the NSX license is expired.

Workaround: Your NSX license must be active in order to add a host to a vSphere Distributed Switch.

- **Issue 1569010/1645525: When upgrading from 6.1.x to NSX for vSphere 6.2.3 on a system connected to vCenter 5.5, the Product field in the "Assign License Key" window displays the NSX license as a generic value of "NSX for vSphere" and not a more specific version such as "NSX for vSphere - Enterprise."**

Workaround: None.

- **Issue 1636916: In a vCloud Air environment, when the NSX Edge version is upgraded from vCNS 5.5.x to NSX 6.x, Edge firewall rules with a source protocol value of "any" are changed to "tcp:any, udp:any"**

As a result, ICMP traffic is blocked, and packet drops may be seen.

Workaround: Before upgrading your NSX Edge version, create more specific Edge firewall rules and replace "any" with specific source port values.

- **Issue 1474238: After vCenter upgrade, vCenter might lose connectivity with NSX**

If you are using vCenter embedded SSO and you are upgrading vCenter 5.5 to vCenter 6.0, vCenter might lose connectivity with NSX. This happens if vCenter 5.5 was registered with NSX using the root user name. In NSX 6.2, vCenter registration with root is deprecated.

Note: If you are using external SSO, no change is necessary. You can retain the same user name, for example admin@mybusiness.mydomain, and vCenter connectivity will not be lost.

Workaround: Reregister vCenter with NSX using the administrator@vsphere.local user name instead of root.

- **Issue 1375794: Shutdown Guest OS for agent VMs (SVA) before powering OFF**

When a host is put into maintenance mode, all service appliances are powered-off, instead of shutting down gracefully. This may lead to errors within third-party appliances.

Workaround: None.

- **Issue 1112628: Unable to power on the Service appliance that was deployed using the Service Deployments view**

Workaround: Before you proceed, verify the following:

- The deployment of the virtual machine is complete.
- No tasks such as cloning, reconfiguring, and so on are in progress for the virtual machine displayed in vCenter task pane.
- In the vCenter events pane of the virtual machine, the following events are displayed after the deployment is initiated:

Agent VM <vm name> has been provisioned.

Mark agent as available to proceed agent workflow.

In such a case, delete the service virtual machine. In service deployment UI, the deployment is seen as Failed. Upon clicking the Red icon, an alarm for an unavailable Agent VM is displayed for the host. When you resolve the alarm, the virtual machine is redeployed and powered on.

- **Issue 1497101: If not all clusters and the hosts in your environment are not prepared, the Upgrade message for Distributed Firewall does not appear on the Host Preparation tab of Installation page**

When you prepare clusters for network virtualization, distributed firewall is enabled on those clusters. If not all clusters in your environment are prepared, the upgrade message for Distributed Firewall does not appear on the Host Preparation tab.

Workaround: Use the following REST call to upgrade Distributed Firewall:

PUT https://<nsxmgr-ip>/api/4.0/firewall/globalroot-0/state

- **Issue 1413125: SSO cannot be reconfigured after upgrade**

When the SSO server configured on NSX Manager is the one native on vCenter server, you cannot reconfigure SSO settings on NSX Manager after vCenter Server is upgraded to version 6.0 and NSX Manager is upgraded to version 6.x.

Workaround: None.

- **Issue 1263858: SSL VPN does not send upgrade notification to remote client**

SSL VPN gateway does not send an upgrade notification to users. The administrator has to manually communicate that the SSL VPN gateway (server) is updated to remote users and they must update their clients.

Workaround: Users need to uninstall the older version of client and install the latest version manually.

- **Issue 1462319: The esx-dvfilter-switch-security VIB is no longer present in the output of the "esxcli software vib list | grep esx" command.**

Starting in NSX 6.2, the esx-dvfilter-switch-security modules are included within the esx-vxlan VIB. The only NSX VIBs installed for 6.2 are esx-vsip and esx-vxlan. During an NSX upgrade to 6.2, the old esx-dvfilter-switch-security VIB gets removed from the ESXi hosts. Starting in NSX 6.2.3, a third VIB, esx-vdpi, is provided along with the esx-vsip and esx-vxlan NSX VIBs. A successful installation will show all 3 VIBs.

Workaround: None.

- **Issue 1481083: After the upgrade, logical routers with explicit failover teaming configured might fail to forward packets properly**

When the hosts are running ESXi 5.5, the explicit failover NSX 6.2 teaming policy does not support multiple active uplinks on distributed logical routers.

Workaround: Alter the explicit failover teaming policy so that there is only one active uplink and the other uplinks are in standby mode.

- **Issue 1411275: vSphere Web Client does not display Networking and Security tab after backup and restore in NSX vSphere 6.2**

When you perform a backup and restore operation after upgrading to NSX vSphere 6.2, the vSphere Web Client does not display the **Networking and Security** tab.

Workaround: When an NSX Manager backup is restored, you are logged out of the Appliance Manager. Wait a few minutes before logging in to the vSphere Web Client.

- **Issue 1764460: After completing Host Preparation, all cluster members appear in ready state, but cluster level erroneously appears as "Invalid"**

After you complete Host Preparation, all cluster members correctly appear in "Ready" state, but cluster level appears as "Invalid" and the reason displayed is that you need a host reboot, even though the host has already been rebooted. This can occur intermittently with vSphere 5.5 and 6.0, and is fixed in vSphere 6.5.

Workaround: In the vCenter ESX Agency Manager MOB https://VC_IP/eam/mob/ you can access the agencies associated with your host clusters. Click on one of the agencies, and click **config** to see the cluster details. Click **ResolveAll** for the affected clusters.

- **Issue 1979457: If GI-SVM is deleted or removed during the upgrade process and backward compatibility mode, then identity firewall through Guest Introspection (GI) will not work unless the GI cluster is upgraded.**

Identity firewall will not work and no logs related to identity firewall would be seen. Identity firewall protection will be suspended unless the cluster is upgraded.

Workaround: Upgrade the cluster so that all the hosts are running the newer version of GI-SVM.

-Or -

Enable Log scraper for identity firewall to work.

NSX Manager Known Issues

- **Issue 1892999: Cannot modify the Unique Selection Criteria even when no VMs are attached to the Universal Security Tag**

If a VM attached to a universal security tag gets deleted, an internal object representing the VM still remains attached to the universal security tag. This causes the universal selection criteria change to fail with error that universal security tags are still attached to VMs.

Workaround: Delete all the universal security tags and then change the universal selection criteria.

- **Issue 1801325: 'Critical' system events and logging generated in the NSX Manager with high CPU and/or disk usage**

You may encounter one or more of the following problems when you have high disk space usage, high churn in job data, or a high job queue size on the NSX Manager:

- 'Critical' system events in the vSphere web client
- High disk usage on NSX Manager for /common partition
- High CPU usage for prolonged periods or at regular intervals
- Negative impact on NSX Manager's performance

Workaround: Contact VMware customer support. See [VMware Knowledge Base article 2147907](#) for more information.

- **Issue 1806368: Reusing controllers from a previously failed primary NSX Manager which is made primary again after a failover causes the DLR config to not be pushed to all hosts**

In a cross-vCenter NSX setup, when the primary NSX Manager fails, a secondary NSX Manager is promoted to primary and a new controller cluster is deployed to be used with the newly promoted secondary (now primary) NSX Manager. When the primary NSX Manager is back on, the secondary NSX Manager is demoted and the primary NSX Manager is restored. In this case, if you reuse the existing controllers that were deployed on this primary NSX Manager before the failover, the DLR config is not pushed to all hosts. This issue does not arise if you create a new controller cluster instead.

Workaround: Deploy a new controller cluster for the restored primary NSX Manager.

- **Issue 1831131: Connection from NSX Manager to SSO fails when authenticated using the LocalOS user**

Connection from NSX Manager to SSO fails when authenticated using the LocalOS user with the error: "Could not establish communication with NSX Manager. Please contact

administrator."

Workaround: Add the Enterprise Admin role for `nsxmanager@localos` in addition to `nsxmanager@domain`.

- **Issue 1800820: UDLR interface update fails on secondary NSX Manager when the old UDLR interface is already deleted from the system**

In a scenario where the Universal Synchronization Service (Replicator) stops working on the primary NSX Manager, you have to delete the UDLR (Universal Distributed Logical Router) and ULS (Universal Logical Switch) interfaces on the primary NSX Manager and create new ones, and then replicate these on the secondary NSX Manager. In this case, the UDLR interface does not get updated in the secondary NSX Manager because a new ULS gets created on the secondary NSX Manager during replication and the UDLR is not connected with the new ULS.

Workaround: Ensure that the replicator is running and delete the UDLR interface (LIF) on the primary NSX Manager which has a newly created ULS as backing and recreate the UDLR interface (LIF) again with the same backing ULS.

- **Issue 1772911: NSX Manager performing very slowly with disk space consumption, and task and job table sizes increasing with close to 100% CPU usage**

You will experience the following:

- NSX Manager CPU is at 100% or is regularly spiking to 100% consumption and adding extra resources to NSX Manager appliance does not make a difference.
- Running the `show process monitor` command in the NSX Manager Command Line Interface (CLI) displays the Java process that is consuming the highest CPU cycles.
- Running the `show filesystems` command on the NSX Manager CLI shows the `/common` directory as having a very high percentage in use, such as `> 90%`.
- Some of the configuration changes time out (sometimes taking over 50 minutes) and are not effective.

See [VMware Knowledge Base article 2147907](#) for more information.

Workaround: Contact VMware customer support for a resolution of this problem.

- **Issue 1785142: Delay in showing 'Synchronization Issues' on primary NSX Manager when communication between primary and secondary NSX Manager is blocked.**

When communication between primary and secondary NSX Manager is blocked, you will not immediately see 'Synchronization Issues' on the primary NSX Manager.

Workaround: Wait for about 20 minutes for communication to be reestablished.

- **Issue 1786066: In a cross-vCenter installation of NSX, disconnecting a secondary NSX Manager may render that NSX Manager unable to reconnect as secondary**

In a cross-vCenter installation of NSX, if you disconnect a secondary NSX Manager, you may be unable to re-add that NSX Manager later as a secondary NSX Manager. Attempts to reconnect the NSX Manager as secondary will result in the NSX Manager being listed as "Secondary" in the Management tab of the vSphere Web Client, but the connection to the primary is not established.

Workaround:

1. Disconnect the secondary NSX Manager from the primary NSX Manager.
2. Add the secondary NSX Manager again to the primary NSX Manager.

- **Issue 1715354: Delay in availability of the REST API**

The NSX Manager API takes some time to be up and running after NSX Manager restarts when FIPS mode is toggled. It may appear as if the API is hung, but this occurs because it takes time for the controllers to re-establish connection with the NSX Manager. You are advised to wait for the NSX API server to be up and running and ensure all controllers are in the connected state before doing any operations.

- **Issue 1441874: Upgrading a single NSX Manager in a vCenter Linked Mode Environment displays an error message**

In an environment with multiple VMware vCenter Servers with multiple NSX Managers, when selecting one or more NSX Managers from the vSphere Web Client > Networking and Security > Installation > Host Preparation, you see this error:

"Could not establish communication with NSX Manager. Please contact administrator."

Workaround: See [VMware Knowledge Base article 2127061](#) for more information.

- **Issue 1696750: Assigning an IPv6 address to NSX Manager via PUT API requires a reboot to take effect**

Changing the configured network settings for NSX Manager via `https://{NSX Manager IP address}/api/1.0/appliance-management/system/network` requires a system reboot to take effect. Until the reboot, pre-existing settings will be shown.

Workaround: None.

- **Issue 1529178: Uploading a server certificate which does not include a common name returns an "internal server error" message**

If you upload a server certificate that does not have any common name, an "internal server error" message appears.

Workaround: Use a server certificate which has both a SubAltName and a common name, or at least a common name.

- **Issue 1655388: NSX Manager 6.2.3 UI displays English language instead of local language when using IE11/Edge browser on Windows 10 OS for JA, CN, and DE languages**

When you launch NSX Manager 6.2.3 with IE11/Edge browser on Windows 10 OS for JA, CN, and DE languages, English language is displayed.

Workaround:

1. Launch the Microsoft Registry Editor (regedit.exe), and go to **Computer > HKEY_CURRENT_USER > SOFTWARE > Microsoft > Internet Explorer > International**.
2. Modify the value of *AcceptLanguage* file to native language. For example, If you want to change language to **DE**, change value and make the **DE** show the first position.
3. Restart the browser, and log in to the NSX Manager again. Appropriate language is displayed.

- **Issue 1435996: Log files exported as CSV from NSX Manager are timestamped with epoch not datetime**

Log files exported as CSV from NSX Manager using the vSphere Web Client are timestamped with the epoch time in milliseconds, instead of with the appropriate time based on the time zone.

Workaround: None.

- **Issue 1644297: Add/delete operation for any DFW section on the primary NSX creates two DFW saved configurations on the secondary NSX**

In a cross-vCenter setup, when an additional universal or local DFW section is added to the primary NSX Manager, two DFW configurations are saved on the secondary NSX Manager. While it does not affect any functionality, this issue will cause the saved configurations limit to be reached more quickly, possibly overwriting critical configurations.

Workaround: None.

- **Issue 1477138: NSX management service doesn't come up when the hostname's length is more than 64 characters**

Certificate creation via OpenSSL library requires a hostname less than or equal to 64 characters.

- **Issue 1437664: NSX Manager list slow to display in Web Client**

In vSphere 6.0 environments with multiple NSX Managers, the vSphere web client may take up to two minutes to display the list of NSX Managers when the logged-in user is being validated with a large AD Group set. You may see a data service timeout error when attempting to display the NSX Manager list. There is no workaround. You must wait for the list to load/relogin to see the NSX Manager list.

- **Issue 1534606: Host Preparation Page fails to load**

When running vCenter in linked mode, each vCenter must be connected to an NSX Manager on the same NSX version. If the NSX versions differ, the vSphere Web Client will only be able to communicate with the NSX Manager running the higher version of NSX. An error similar to "Could not establish communication with NSX Manager. Please contact administrator," will be displayed on the Host Preparation tab.

Workaround: All NSX Managers should be upgraded to the same NSX software version.

- **Issue 1027066: vMotion of NSX Manager may display the error message, "Virtual ethernet card Network adapter 1 is not supported"**

You can ignore this error. Networking will work correctly after vMotion.

- **Issue 1460766: NSX Manager UI do not automatically log out after changing password using NSX Command Line Interface**

If you are logged in to NSX Manager and recently changed your password using CLI, you might continue to stay logged in to the NSX Manager UI using your old password.

Typically, NSX Manager client should automatically log you out if the session times out for being inactive.

Workaround: Log out from the NSX Manager UI and log back in with your new password.

- **Issue 1966681: Incorrect reporting of duplicate NSX Manager IP**

The log file get flooded with the duplicate NSX Manager IP and reports the incorrect information about the duplicate IP in the network.

- **Issue 1467382: Unable to edit a network host name**

After you login to NSX Manager virtual appliance and navigate to the Appliance Management, click Manage Appliance Settings, and click Network under Settings to edit

the network host name, you might receive an invalid domain name list error. This happens when the domain names specified in the Search Domains field are separated with whitespace characters, instead of commas. NSX Manager only accepts domain names that are comma separated.

Workaround:

1. Log in to the NSX Manager virtual appliance.
 2. Under **Appliance Management**, click **Manage Appliance Settings**.
 3. From the Settings panel, click **Network**.
 4. Click **Edit** next to DNS Servers.
 5. In the Search Domains field, replace all whitespace characters with commas.
 6. Click **OK** to save the changes.
- **Issue 1486193/1436953: False system event is generated even after successfully restoring NSX Manager from a backup**
After successfully restoring NSX Manager from a backup, the following system events appear in the vSphere Web Client when you navigate to **Networking & Security: NSX Managers: Monitor: System Events**.
 - Restore of NSX Manager from backup failed (with Severity=Critical) .
 - Restore of NSX Manager successfully completed (with Severity=Informational) .

Workaround: If the final system event message shows as successful, you can ignore the system generated event messages.

- **Issue 1783528: NSX Manager CPU Utilization spikes every Friday night / Saturday Morning**
NSX polls LDAP for full sync every Friday night. There is no option to configure specific Active Directory Organisational Unit or Container, therefore NSX pulls in all objects which are related to the domain provided.
Workaround: Increase NSX Manager vCPU from 4 to 6

Workaround: Increase NSX Manager vCPU from 4 to 6

NSX Controller Known Issues

- **Issue 1856465: If an ESXi host is down on one of the site in a NSX Cross-vCenter environment, the CDO mode does not get enabled on that site**
If an ESXi host is down on a site, enabling or disabling CDO mode will not be completely successfully on that site.
If the host is down on one of the Secondary site, the CDO mode operation will succeed on the Primary site. However the CDO mode operation will fail on the Secondary site. This may lead to inconsistent behavior.
Workaround: This issue impacts NSX 6.3.0 and above.

- Ensure that all the ESXi hosts are up before doing any CDO operations.
- In order to recover from an inconsistent state, remove the host from the vCenter inventory and add it again.

Logical Networking and NSX Edge Known Issues

- **Issue 1904612: Layer 2 VPN tunnel displays "up" on L2VPN server when client is powered off**

If you create a L2 VPN between two NSX Edges, then power down the client NSX Edge, the Server NSX Edge still displays that the VPN tunnel is up.

Workaround: None.

- **Issue 1242207: Changing router ID during the run time is not reflected in OSPF topology**

If you try to change router ID without disabling OSPF, new external link-state advertisements (LSAs) are not re-generated with this router ID causing loss of OSPF external routes.

Disable OSPF, change router ID and then enable OSPF again.

- **Issue 1894277: IPSec site configuration PSK is not retained when the local or peer subnet gets changed**

As the masked PSK gets saved in the database, tunnel between the peers won't come up because of the password mismatch.

Workaround: Reconfigure the IPSec configuration with a valid password.

- **Issue 1492497: Cannot filter NSX Edge DHCP traffic**

You cannot apply any firewall filters to DHCP traffic on an NSX Edge because the DHCP server on an NSX Edge utilizes raw sockets that bypass the TCP/IP stack.

Workaround: None.

- **Issue 1781438: On the ESG or DLR NSX Edge appliances, the routing service does not send an error message if it receives the BGP path attribute MULTI_EXIT_DISC more than once.**

The edge router or distributed logical router does not send an error message if it receives the BGP path attribute MULTI_EXIT_DISC more than once. AS per RFC 4271 [Sec 5], the same attribute (attribute with the same type) cannot appear more than once within the Path Attributes field of a particular UPDATE message.

Workaround: None.

- **Issue 1786515: User with 'Security Administrator' privileges unable to edit the load balancer configuration through the vSphere web client UI.**

A user with "Security Administrator" privileges for a specific NSX Edge is not able to edit the Global Load Balancer Configuration for that edge, using the vSphere web client UI.

The following error message is displayed: "User is not authorized to access object Global and feature si.service, please check object access scope and feature permissions for the user."

Workaround: None.

- **Issue 1849042/1849043: Admin account lockout when password aging is configured on the NSX Edge appliance**

If password aging is configured for the admin user on the NSX Edge appliance, when the password ages out there is a 7 day period where the user will be asked to change the password when logging into the appliance. Failure to change the password will result in the account being locked. Additionally, if the password is changed at the time of logging in at the CLI prompt, the new password may not meet the strong password policy enforced by the UI and REST.

Workaround: To avoid this problem, always use the UI or REST API to change the admin password before the existing password expires. If the account does become locked, also use the UI or REST API to configure a new password and the account will become unlocked again.

- **Issue 1711013: Takes about 15 minutes to sync FIB between Active/Standby NSX Edge after rebooting the standby VM.**

When a standby NSX Edge is powered off, the TCP session is not closed between active and standby mode. The active edge will detect that standby is down after keepalive (KA) failure (15 minutes). After 15 minutes, a new socket connection is established with the standby edge and FIB is synced between the active/standby edge.

Workaround: None.

- **Issue 1733282: NSX Edge no longer supports static device routes**

NSX Edge does not support configuration of static routes with NULL nexthop address.

Workaround: None.

- **Issue 1860583: Avoid using remote sysloggers as FQDN if DNS is not reachable.**

On an NSX edge, if the remote sysloggers are configured using FQDN and DNS is not reachable, then routing functionality might be impacted. The problem might not happen consistently.

Workaround: It is recommended to use IP addresses instead of FQDN.

- **Issue 1850773: NSX Edge NAT reports invalid configuration when multiple ports are used on the Load Balancer configuration**

This issue occurs every time you configure a Load Balancer virtual server with more than one port. Due to this, NAT becomes unmanageable while this configuration state exists for the affected NSX Edge.

Workaround: See [VMware Knowledge Base article 2149942](#) for more information and workaround.

- **Issue 1764258: Traffic blackholed for upto eight minutes post HA failover or Force-Sync on NSX Edge configured with sub-interface**

If an HA failover is triggered or you start a Force-Sync over a sub-interface, traffic is blackholed for upto eight minutes.

Workaround: Do not use subinterfaces for HA.

- **Issue 1767135: Errors when trying to access certificates and application profiles under Load Balancer**

Users with Security Admin privileges and Edge scope are unable to access certificates and application profiles under Load Balancer. The vSphere Web Client shows error messages.

Workaround: None.

- **Issue 1792548: NSX Controller may get stuck at the message: 'Waiting to join cluster'**

NSX Controller may get stuck at the message: 'Waiting to join cluster' (CLI command: `show control-cluster status`). This occurs because the same IP address has been configured for `foreth0` and `breth0` interfaces of the controller while the controller is coming up. You can verify this by using the following CLI command on the controller: `show network interface`

Workaround: Contact VMware customer support.

- **Issue 1747978: OSPF adjacencies are deleted with MD5 authentication after NSX Edge HA failover**

In an NSX for vSphere 6.2.4 environment where the NSX Edge is configured for HA with OSPF graceful restart configured and MD5 is used for authentication, OSPF fails to start gracefully. Adjacencies form only after the dead timer expires on the OSPF neighbor nodes.

Workaround: None

- **Issue 1804116: Logical Router goes into Bad State on a host that has lost communication with the NSX Manager**

If a Logical Router is powered on or redeployed on a host that has lost communication with the NSX Manager (due to NSX VIB upgrade/install failure or host communication issue), the Logical Router will go into Bad State and continuous auto-recovery operation via Force-Sync will fail.

Workaround: After resolving the host and NSX Manager communication issue, reboot the NSX Edge manually and wait for all interfaces to come up. This workaround is only needed for Logical Routers and not NSX Edge Services Gateway (ESG) because the auto-recovery process via force-sync reboots NSX Edge.

- **Issue 1783065: Cannot configure Load Balancer for UDP port along with TCP by IPv4 and IPv6 address together**

UDP only supports `ipv4-ipv4`, `ipv6-ipv6` (frontend-backend). There is a bug in NSX Manager that even IPv6 link local address is read and pushed as an IP address of the grouping object, and you cannot select IP protocol to use in LB configuration.

Here is an example LB configuration demonstrating the issue:

In the Load Balancer configuration, pool "vCloud_Connector" is configured with a grouping object (vm-2681) as pool member and this object contains both IPv4 and IPv6 addresses, which cannot be supported by LB L4 Engine.

```
{
  "algorithm" : {
    ...
  },
  "members" : [
    {
      ... ,
      ...
    }
  ],
  "applicationRules" : [],
  "name" : "vCloud_Connector",
  "transparent" : {
    "enable" : false
  }
}

{
  "value" : [
    "fe80::250:56ff:feb0:d6c9",
    "10.204.252.220"
  ],
  "id" : "vm-2681"
}
```

Workaround:

- Option 1: Enter the IP address of the pool member instead of grouping objects in pool member.
- Option 2: Do not use IPv6 in the VMs.
- **Issue 1777792: Peer Endpoint set as 'ANY' causes IPSec connection to fail**
When IPSec configuration on NSX Edge sets remote peer endpoint as 'ANY', the Edge acts as an IPSec "server" and waits for remote peers to initiate connections. However, when the initiator sends a request for authentication using PSK+XAUTH, the Edge displays this error message: "initial Main Mode message received on XXX.XXX.XX.XX:500 but no connection has been authorized with policy=PSK+XAUTH" and IPsec can't be established.

Workaround: Use specific peer endpoint IP address or FQDN in IPSec VPN configuration instead of ANY.

- **Issue 1741158: Creating a new, unconfigured NSX Edge and applying configuration can result in premature Edge service activation.**
If you use the NSX API to create a new, unconfigured NSX Edge, then make an API call to disable one of the Edge services on that Edge (for example, set dhcp-enabled to "false"), and finally apply configuration changes to the disabled Edge service, that service will be made active immediately.

Workaround: After you make a configuration change to an Edge service that you wish to keep in disabled state, immediately issue a PUT call to set the enabled flag to "false" for that service.

- **Issue 1758500: Static route with multiple next-hops does not get installed in NSX Edge routing and forwarding tables if at least one of the next-hop configured is the**

Edge's vNIC IP address

With ECMP and multiple next-hop addresses, NSX allows the Edge's vNIC's IP address to be configured as next-hop if at least one of the next-hop IP addresses is valid. This is accepted without any errors or warnings but route for the network is removed from the Edge's routing/forwarding table.

Workaround: Do not configure the Edge's own vNIC IP address as a next-hop in static route when using ECMP.

- **Issue 1716464: NSX Load Balancer will not route to VMs newly tagged with a Security tag.**

If we deploy two VMs with a given tag, and then configure an LB to route to that tag, the LB will successfully route to those two VMs. But if we then deploy a third VM with that tag, the LB only routes to the first two VMs.

Workaround: Click "Save" on the LB Pool. This rescans the VMs and will start routing to newly tagged VMs.

- **Issue 1753621: When Edge with private local AS sends routes to EBGp peers, all the private AS paths are stripped off from the BGP routing updates sent.**

NSX currently has a limitation that prevents it from sharing the full AS path with eBGP neighbors when the AS path contains only private AS paths. While this is the desired behavior in most cases, there are cases in which the administrator may want to share private AS paths with an eBGP neighbor.

Workaround: No workaround available to make the Edge announce all the AS paths in the BGP update.

- **Issue 1461421: "show ip bgp neighbor" command output for NSX Edge retains the historical count of previously established connections**

The "show ip bgp neighbor" command displays the number of times that the BGP state machine transitioned into the Established state for a given peer. Changing the password used with MD5 authentication causes the peer connection to be destroyed and re-created, which in turn will clear the counters. This issue does not occur with an Edge DLR.

Workaround: To clear the counters, execute the "clear ip bgp neighbor" command.

- **Issue 1656713: IPsec Security Policies (SPs) missing on the NSX Edge after HA failover, traffic cannot flow over tunnel**

The **Standby > Active** switchover will not work for traffic flowing on IPsec tunnels.

Workaround: Disable/Enable IPsec after the NSX Edge switchover.

- **Issue 1354824: When an Edge VM becomes corrupted or becomes otherwise unreachable due to such reasons as a power failure, system events are raised when the health check from NSX Manager fails**

The system events tab will report "Edge Unreachability" events. The NSX Edges list may continue to report a Status of Deployed.

Workaround: Use the following API to get detailed status information about an NSX Edge:

GET https://NSX-Manager-IP-Address/api/4.0/edges/edgeId/status?detailedStatus=true

- **Issue 1647657: Show commands on an ESXi host with DLR (Distributed Logical Router) display no more than 2000 routes per DLR instance**

Show commands on an ESXi host with DLR enabled will not show more than 2000 routes per DLR instance, although more than this maximum may be running. This issue is a display issue, and the data path will work as expected for all routes.

Workaround: None.

- **Issue 1634215: OSPF CLI commands output does not indicate whether routing is disabled**

When OSPF is disabled, routing CLI commands output does not show any message saying "*OSPF is disabled*". The output is empty.

Workaround: The `show ip ospf` command will display the correct status.

- **Issue 1647739: Redeploying an Edge VM after a vMotion operation will cause the Edge or DLR VM to be placed back on the original cluster.**

Workaround: To place the Edge VM in a different resource pool or cluster, use the NSX Manager UI to configure the desired location.

- **Issue 1463856: When NSX Edge Firewall is enabled, existing TCP connections are blocked**

TCP connections are blocked through the Edge stateful firewall as the initial three-way handshake cannot be seen.

Workaround: To handle such existing flows, do the following. Use the NSX REST API to enable the flag "tcpPickOngoingConnections" in the firewall global configuration. This switches the firewall from strict mode to lenient mode. Next, enable the firewall. Once existing connections have been picked up (this may take a few minutes after you enable the firewall), set the flag "tcpPickOngoingConnections" back to false to return the firewall to strict mode. (This setting is persistent.)

PUT /api/4.0/edges/{edgeId}/firewall/config/global

```
<globalConfig>
  <tcpPickOngoingConnections>true</tcpPickOngoingConnections>
</globalConfig>
```

- **Issue 1374523: Reboot ESXi, or run `[services.sh restart]` after installation of VXLAN VIB to make the VXLAN commands available using `esxcli`**

After installation of VXLAN VIB, you must reboot ESXi or run the `[services.sh restart]` command, so that the VXLAN commands becomes available using `esxcli`.

Workaround: Instead of using `esxcli`, use `localcli`.

- **Issue 1525003: Restoring an NSX Manager backup with an incorrect passphrase will silently fail as critical root folders cannot be accessed**

Workaround: None.

- **Issue 1637639: When using the Windows 8 SSL VPN PHAT client, the virtual IP is not**

assigned from the IP pool

On Windows 8, the virtual IP address is not assigned as expected from the IP pool when a new IP address is assigned by the Edge Services Gateway or when the IP pool changes to use a different IP range.

Workaround: This issue occurs only on Windows 8. Use a different Windows OS to avoid experiencing this issue.

- **Issue 1628220: DFW or NetX observations are not seen on receiver side**

Traceflow may not show DFW and NetX observations on receiver side if switch port associated with the destination vNIC changed. It will not be fixed for vSphere 5.5 releases. For vSphere 6.0 and up, there is no such issue.

Workaround: Do not disable vNIC. Reboot VM.

- **Issue 1483426: IPsec and L2 VPN service status shows as down even when the service is not enabled**

Under the Settings tab in the UI, the L2 service status is displayed as down, however the API shows the L2 status as up. L2 VPN and IPsec service always shows as down in the Settings tab unless the UI page is refreshed.

Workaround: Refresh the page.

- **Issue 1446327: Some TCP-based applications may time out when connecting through NSX Edge**

The default TCP established connection inactivity timeout is 3600 seconds. The NSX Edge deletes any connections idle for more than the inactivity timeout and drops those connections.

Workaround:

1. If the application has a relatively long inactivity time, enable TCP keepalives on the hosts with `keep_alive_interval` set to less than 3600 seconds.
2. Increase the Edge TCP inactivity timeout to greater than 2 hours using the following NSX REST API. For example, to increase the inactivity timeout to 9000 seconds.

NSX API URL:

```
/api/4.0/edges/{edgeId}/systemcontrol/config PUT Method <systemControl>
```

```
<property>sysctl.net.netfilter.nf_conntrack_tcp_timeout_established=9000</property> </systemControl>
```

- **Issue 1089238: Cannot configure OSPF on more than one DLR Edge uplink**

Currently it is not possible to configure OSPF on more than one of the eight DLR Edge uplinks. This limitation is a result of the sharing of a single forwarding address per DLR instance.

Workaround: This is a current system limitation and there is no workaround.

- **Issue 1499978: Edge syslog messages do not reach remote syslog server**

Immediately after deployment, the Edge syslog server cannot resolve the hostnames for any configured remote syslog servers.

Workaround: Configure remote syslog servers using their IP address, or use the UI to Force Sync the Edge.

- **Issue 1489829: Logical router DNS Client configuration settings are not fully applied after updating REST Edge API**

Workaround: When you use REST API to configure DNS forwarder (resolver), perform the

following steps:

1. Specify the DNS Client XML server's settings so that they match the DNS forwarder setting.
2. Enable DNS forwarder, and make sure that the forwarder settings are same as the DNS Client server's settings specified in the XML configuration.

- **Issue 1243112: Validation and error message not present for invalid next hop in static route, ECMP enabled**

When trying to add a static route, with ECMP enabled, if the routing table does not contain a default route and there is an unreachable next hop in the static route configuration, no error message is displayed and the static route is not installed.

Workaround: None.

- **Issue 1281425: If an NSX Edge virtual machine with one sub interface backed by a logical switch is deleted through the vCenter Web Client user interface, data path may not work for a new virtual machine that connects to the same port**

When the Edge virtual machine is deleted through the vCenter Web Client user interface (and not from NSX Manager), the VXLAN trunk configured on dvPort over opaque channel does not get reset. This is because trunk configuration is managed by NSX Manager.

Workaround: Manually delete the VXLAN trunk configuration by following the steps below:

1. Navigate to the vCenter Managed Object Browser by typing the following in a browser window:

`https://<vc-ip>/mob?vmodl=1`

2. Click **Content**.
3. Retrieve the dvsUuid value by following the steps below.
 - a. Click the rootFolder link (for example, group-d1(Datacenters)).
 - b. Click the data center name link (for example, datacenter-1).
 - c. Click the networkFolder link (for example, group-n6).
 - d. Click the DVS name link (for example, dvs-1)
 - e. Copy the value of uuid.
4. Click **DVSManager** and then click **updateOpaqueDataEx**.
5. In *selectionSet*, add the following XML.

```
<selectionSet xsi:type="DVPortSelection">
  <dvsUuid>value</dvsUuid>
  <portKey>value</portKey> <!--port number of the DVPG where trunk vnic got connected-->
</selectionSet>
```

6. In *opaqueDataSpec*, add the following XML

```
<opaqueDataSpec>
  <operation>remove</operation>
  <opaqueData>
    <key>com.vmware.net.vxlan.trunkcfg</key>
    <opaqueData></opaqueData>
  </opaqueData>
</opaqueDataSpec>
```

7. Set **isRuntime** to false.
8. Click **Invoke Method**.

9. Repeat steps 5 through 8 for each trunk port configured on the deleted Edge virtual machine.

- **Issue 1637939: MD5 certificates are not supported while deploying hardware gateways**

While deploying hardware gateway switches as VTEPs for logical L2 VLAN to VXLAN bridging, the physical switches support at minimum SHA1 SSL certificates for OVSDB connection between the NSX controller and OVSDB switch.

Workaround: None.

- **Issue 1637943: No support for hybrid or multicast replication modes for VNIs that have a hardware gateway binding**

Hardware gateway switches when used as VTEPs for L2 VXLAN-to-VLAN bridging support Unicast replication mode only.

Workaround: Use Unicast replication mode only.

- **Issue 1995142: Host is not removed from replication cluster after being removed from VC inventory**

If a user adds a host to a replication cluster and then removes the host from VC inventory before removing it from the cluster, the legacy host will remain in the cluster.

Workaround: Whenever removing a host, first make sure it has already been removed from replication cluster if any.

Security Services Known Issues

- **Issue 2000749: Distributed Firewall stays in Publishing state with certain firewall configurations**

Distributed Firewall stays in "Publishing" state if you have a security group that contains an IPSet with 0.0.0.0/0 as an EXCLUDE member, an INCLUDE member, or as a part of 'dynamic membership containing Intersection (AND)'.

Workaround: Use a subnet mask other than /0 in your IPSet configuration. You can define 0.0.0.0/0 as "0.0.0.0/1,128.0.0.0/1".

- **Issue 1854661: In a cross-VC setup, filtered firewall rules do not display the index value when you switch between NSX Managers**

After you apply a rule filter criteria to an NSX Manager and then switch to a different NSX Manager, the rule index shows as '0' for all filtered rules instead of showing the actual position of the rule.

Workaround: Clear the filter to see the rule position.

- **Issue 1474650: For NetX users, ESXi 5.5.x and 6.x hosts experience a purple diagnostic screen mentioning ALERT: NMI: 709: NMI IPI received**

When a large number of packets are transmitted or received by a service VM, DVFilter continues to dominate the CPU resulting in heartbeat loss and a purple diagnostic screen. See [VMware Knowledge Base article 2149704](#) for more information.

Workaround: Upgrade the ESXi host to any of the following ESXi versions that are the

minimum required to use NetX:

- 5.5 Patch 10
- ESXi 6.0U3
- ESXi 6.5

- **Issue 1787680: Deleting Universal Firewall Section fails when NSX Manager is in Transit mode**

When you try to delete a Universal Firewall Section from the UI of an NSX Manager in Transit mode, and publish, the Publish fails and as a result you are not able to set the NSX Manager to Standalone mode.

Workaround: Use the Single Delete Section REST API to delete the Universal Firewall Section.

- **Issue 1689159: The Add Rule feature in Flow Monitoring does not work correctly for ICMP flows.**

When adding a rule from Flow Monitoring, the Services field will remain blank if you do not explicitly set it to ICMP and as a result, you may end up adding a rule with the service type "ANY".

Workaround: Update the Services field to reflect ICMP traffic.

- **Issue 1632235: During Guest Introspection installation, network drop down list displays "Specified on Host" only**

When installing Guest Introspection with the NSX anti-virus-only license and vSphere Essential or Standard license, the network drop down list will display only the existing list of DV port groups. This license does not support DVS creation.

Workaround: Before installing Guest Introspection on a vSphere host with one of these licenses, first specify the network in the "Agent VM Settings" window.

- **Issue 1652155: Creating or migrating firewall rules using REST APIs may fail under certain conditions and report HTTP 404 error**

Adding or migrating firewall rules using REST APIs is not supported under these conditions:

- Creating firewall rules as a bulk operation when the autosavedraft=true is set.
- Adding firewall rules in sections concurrently.

Workaround: Set the autoSaveDraft parameter to false in the API call when performing bulk firewall rule creation or migration.

- **Issue 1509687: URL length supports up to 16000 characters when assigning a single security tag to many VMs at a time in one API call**

A single security tag cannot be assigned to a large number of VMs simultaneously with a single API if the URL length is more than 16,000 characters.

Workaround: To optimize performance, tag up to 500 VMs in a single call.

- **Issue 1662020: Publish operation may fail resulting in an error message "Last publish failed on host *host number*" on DFW UI in General and Partner Security Services sections**

After changing any rule, the UI displays "Last publish failed on host *host number*". The

hosts listed on the UI may not have correct version of firewall rules, resulting in lack of security and/or network disruption.

The problem is usually seen in the following scenarios:

- After upgrade from older to latest NSXv version.
- Move a host out of cluster and move it back in.
- Move a host from one cluster to another.

Workaround: To recover, you must force sync the affected clusters (firewall only).

- **Issue 1481522: Migrating firewall rule drafts from 6.1.x to 6.2.3 is not supported as the drafts are not compatible between the releases**

Workaround: None.

- **Issue 1628679: With identity-based firewall, the VM for removed users continues to be part of the security group**

When a user is removed from a group on the AD server, the VM where the user is logged-in continues to be a part of the security-group. This retains firewall policies at the VM vNIC on the hypervisor, thereby granting the user full access to services.

Workaround: None. This behavior is expected by design.

- **Issue 1496273: UI allows creation of in/out NSX firewall rules that cannot be applied to Edges**

The web client incorrectly allows creation of an NSX firewall rule applied to one or more NSX Edges when the rule has traffic traveling in the 'in' or 'out' direction and when PacketType is IPV4 or IPV6. The UI should not allow creation of such rules, as NSX cannot apply them to NSX Edges.

Workaround: None.

- **Issue 1494718: New universal rules cannot be created, and existing universal rules cannot be edited from the flow monitoring UI**

Workaround: Universal rules cannot be added or edited from the flow monitoring UI. EditRule will be automatically disabled.

- **Issue 1066277: Security policy name does not allow more than 229 characters**

The security policy name field in the Security Policy tab of Service Composer can accept up to 229 characters. This is because policy names are prepended internally with a prefix.

Workaround: None.

- **Issue 1443344: Some versions of 3rd-party Networks VM-Series do not work with NSX Manager default settings**

Some NSX 6.1.4 or later components disable SSLv3 by default. Before you upgrade, please check that all third-party solutions integrated with your NSX deployment do *not* rely on SSLv3 communication. For example, some versions of the Palo Alto Networks VM-series solution require support for SSLv3, so please check with your vendors for their version requirements.

- **Issue 1660718: Service Composer policy status is shown as "In Progress" at the UI and "Pending" in the API output**

Workaround: None.

- **Issue 1317814: Service Composer goes out of sync when policy changes are made while one of the Service Managers is down**

When a policy changes is made when one of multiple Service Managers is down, the changes will fail, and Service Composer will fall out of sync.

Workaround: Ensure the Service Manager is responding and then issue a force sync from Service Composer.

- **Issue 1070905: Cannot remove and re-add a host to a cluster protected by Guest Introspection and third-party security solutions**

If you remove a host from a cluster protected by Guest Introspection and third-party security solutions by disconnecting it and then removing it from vCenter Server, you may experience problems if you try to re-add the same host to the same cluster.

Workaround: To remove a host from a protected cluster, first put the host in maintenance mode. Next, move the host into an unprotected cluster or outside all clusters and then disconnect and remove the host.

- **Issue 1648578: NSX forces the addition of cluster/network/storage when creating a new NetX host-based service instance**

When you create a new service instance from the vSphere Web Client for NetX host-based services such as Firewall, IDS, and IPS , you are forced to add cluster/network/storage even though these are not required.

Workaround: When creating a new service instance, you may add any information for cluster/network/storage to fill out the fields. This will allow the creation of the service instance and you will be able to proceed as required.

Monitoring Services Known Issues

- **Issue 1466790: Unable to choose VMs on bridged network using the NSX traceflow tool**

Using the NSX traceflow tool, you cannot select VMs that are not attached to a logical switch. This means that VMs on an L2 bridged network cannot be chosen by VM name as the source or destination address for traceflow inspection.

Workaround: For VMs attached to L2 bridged networks, use the IP address or MAC address of the interface you wish to specify as destination in a traceflow inspection. You cannot choose VMs attached to L2 bridged networks as source. See the [knowledge base article 2129191](#) for more information.