

VMware NSX for vSphere 6.3.2 Release Notes

VMware NSX for vSphere 6.3.2 | Released 08 June 2017 | Build 5672532

What's in the Release Notes

The release notes cover the following topics:

- [What's New in NSX 6.3.2](#)
- [Versions, System Requirements, and Installation](#)
- [Deprecated and Discontinued Functionality](#)
- [Upgrade Notes](#)
- [Known Issues](#)
- [Resolved Issues](#)
- [Document Revision History](#)

What's New in 6.3.2

NSX for vSphere 6.3.2 is a bug-fix release addressing specific customer bugs. See [Resolved in 6.3.2](#) for more information.

This version fully supports the Controller Disconnected Operation (CDO) Mode. For more information, see the section [Controller Disconnected Operation \(CDO\) Mode](#) in the *NSX Administration Guide*.

Starting in NSX 6.3.2, Open VM Tools is supported with Distributed Firewall.

See what's new and changed in NSX [6.3.1](#) and 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.

Product or Component

Recommended Version

NSX for vSphere

VMware recommends the latest NSX 6.3 release for new deployments and when upgrading from 6.1.x.

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.

vSphere

- 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.

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)

vRealize Orchestrator

NSX-vRO plugin 1.1.0 or later.

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

Deleting firewall configuration or default section:

- Request to delete firewall section is now rejected if the default section is specified: DELETE /api/4.0/firewall/globalroot-0/config/layer2sections|layer3sections/sectionId
- New method introduced to get default configuration. Use 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 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

Upgrade Notes

- [Upgrade Notes related to NSX and vSphere](#)
- [Upgrade Notes related to NSX Components](#)
- [Upgrade Notes related to FIPS](#)

Note: 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.

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

Upgrade Notes related to NSX and vSphere

- 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.

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
- **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.

- **Upgrading to vSphere 6.5a:** When upgrading from vSphere 5.5 or 6.0 to vSphere 6.5a, 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.

- **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.
- **Reset vSphere web client:** After upgrading NSX Manager, you must reset the vSphere web client server as explained in the [NSX Upgrade documentation](#). Until you do this, the **Networking and Security** tab may fail to appear in the vSphere web client. You also may need to clear your browser cache or history.
- **Stateless environments:** NSX upgrades in a stateless host environment use new VIB URLs: 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:

1. Manually download the latest NSX VIBs from NSX Manager from a fixed URL.
2. Add the VIBs to the host image profile.

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:

- Find the new VIB URL from `https://<NSX-Manager-IP>/bin/vdn/nwfabric.properties`.
 - Fetch VIBs of required ESX host version from corresponding URL.
 - Add them to host image profile.
- **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.

Upgrade Notes related to NSX Components

- **Upgrading to NSX 6.3.1**
 - If you are upgrading from NSX 6.3.0 to NSX 6.3.1: Only NSX Manager and Host Clusters need to be updated. NSX 6.3.1 shows only these components for upgrade.
 - If you are upgrading from NSX 6.2.x to NSX 6.3.1: You will need to perform a full upgrade, upgrading all NSX components, however, note that NSX Edge and Guest Introspection will show version 6.3.0, not 6.3.1.
- **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:

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 below for defaults).

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

- **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.
- **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).
- **Controller disk layout:** Upgrades from 6.2.2 and earlier will not receive the new disk layout introduced in 6.2.3 which provides separate disk partitions for data and logs to improve controller stability.
- **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"
}
```

- **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.

Upgrade Notes related to 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-SHA384, 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.

Known Issues

Known issues are grouped as follows:

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

General Known Issues

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 1558285: Deleting cluster with Guest Introspection from vCenter results in null pointer exception

Services such as Guest Introspection must be removed first before a cluster is removed from vCenter

Workaround: Delete the EAM Agency for the service deployment with no associated cluster.

Issue 1629030: The packet capture central CLI (debug packet capture and show packet capture) requires vSphere 5.5U3 or higher

These commands are not supported on earlier vSphere 5.5 releases.

Workaround: VMware advises all NSX customers to run vSphere 5.5U3 or higher.

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.

New 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.

New 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.

New 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.

New 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.

New Issue 1834285: HTTP/HTTPS Transaction fails on NSX Load Balancer after upgrading to NSX 6.1.5 or later.

Starting in NSX 6.1.5, when enabling x-forwarded-for, the HTTP connection mode has been changed from passive close (option `httpclose`) to the default HTTP server-close (option `http-server-close`) mode, keeping the client-facing connection open while the server-facing connection is closed after receiving a response from the server. This causes problems in some applications because in NSX versions earlier than 6.1.5, the Load Balancer did not close the connection proactively, but inserted the "Connection:close" header in both directions to indicate to the client or server to close the connection.

Workaround: Add an application rule with the script option `httpclose` and associate it with the broken virtual server.

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 API POST `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 1791371: When upgrading ESXi hosts to vSphere 6.5a, if Guest Introspection and VXLAN VIBs are upgraded in parallel, an alarm is raised

The Guest Introspection and VXLAN VIBs are different for vSphere 6.5a and when you upgrade these in parallel, the VXLAN VIB upgrade raises an alarm asking for a host reboot.

Workaround: First install the VXLAN VIBs and then the Guest Introspection VIBs when upgrading to vSphere 6.5a.

Issue 1805983: When you upgrade to NSX 6.2.5, 6.2.6 or 6.3.0, Virtual Servers do not work if they do not contain a server pool.

Virtual servers without server pools can only serve HTTP/HTTPS redirection. No other functionality works.

Workaround: Create a dummy pool without any members in it and assign it to the virtual server.

Issue 1797307: NSX Edge may run into split-brain after upgrade or redeploy

On the standby NSX Edge, the show service highavailability CLI command shows high availability status as "Standby" but the config engine status as "Active".

Workaround: Reboot the standby NSX Edge.

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.

Workaround: Instead of failing back as soon as the host comes back on, delay the failback to the newly prepped host of these VMs.

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>
</nwFabricFeatureConfig>
```

Issue 1802688: Upgrade from NSX 6.2.x to 6.3.0 does not reflect updated DFW enabled status

After you upgrade NSX from 6.2.x to 6.3.0 and cluster VIBs to 6.3.0 bits, when you add a new host to the upgraded cluster, the firewall status of the concerned host and the cluster keep spinning busy and the status does not get updated, even though the new VIBs have been installed on the new host.

Workaround: Do the following:

1. Run the Message Bus Sync API call for the host using the API POST <https://<NSX-IP>/api/2.0/nwfabric/configure?action=synchronize>. This will bring that host and cluster firewall status to "Disabled".

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

2. Now Enable Firewall for that cluster from UI Installation > Hostprep page. This should bring all the hosts in that cluster to DFW enabled mode.

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.^[1] 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 1660355: VMs which are migrated from 6.1.5 to 6.2.3 and later will not have support for TFTP ALG

Even though the host is enabled, VMs which are migrated from 6.1.5 to 6.2.3 and later will not have support for TFTP ALG.

Workaround: Add and remove the VM from the exclusion list or restart the VM, so that new 6.2.3 (and later) filter gets created which will have support for TFTP ALG.

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 1332563: 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 1473537: 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.

If not all clusters in your environment are 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 1215460: If a service group is modified after the upgrade to add or remove services, these changes are not reflected in the firewall table

User created service groups are expanded in the Edge Firewall table during upgrade - i.e., the Service column in the firewall table displays all services within the service group. If the service group is modified after the upgrade to add or remove services, these changes are not reflected in the firewall table.

Workaround: Create a new service group with a different name and then consume this service group in the firewall rule.

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 1266433: 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 1474066: The NSX REST API call to enable or disable IP detection seems to have no effect

If host cluster preparation is not yet complete, the NSX REST API call to enable or disable IP detection (<https://<nsxmgr-ip>/api/2.0/xvs/networks/universalwire-5/features>) has no effect.

Workaround: Before issuing this API call, make sure the host cluster preparation is complete.

Issue 1459032: Error configuring VXLAN gateway

When configuring VXLAN using a static IP pool (at **Networking & Security > Installation > Host Preparation > Configure VXLAN**) and the configuration fails to set an IP pool gateway IP on the VTEP (because the gateway is not properly configured or is not reachable), the VXLAN configuration status enters the Error (RED) state at for the host cluster.

The error message is VXLAN Gateway cannot be set on host and the error status is VXLAN_GATEWAY_SETUP_FAILURE. In the REST API call, GET <https://<nsxmgr-ip>/api/2.0/nwfabric/status?resource=<cluster-moid>>, the status of VXLAN is as follows:

```
<nwFabricFeatureStatus>
  <featureId>com.vmware.vshield.nsxmgr.vxlan</featureId>
  <featureVersion>5.5</featureVersion>
  <updateAvailable>false</updateAvailable>
  <status>RED</status>
  <message>VXLAN Gateway cannot be set on host</message>
  <installed>true</installed>
  <enabled>true</enabled>
  <errorStatus>VXLAN_GATEWAY_SETUP_FAILURE</errorStatus>
</nwFabricFeatureStatus>
```

Workaround: To fix the error, there are two options.

- Option 1: Remove VXLAN configuration for the host cluster, fix the underlying gateway setup in the IP pool by making sure the gateway is properly configured and reachable, and then reconfigure VXLAN for the host cluster.
- Option 2: Perform the following steps.
 1. Fix the underlying gateway setup in the IP pool by making sure the gateway is properly configured and reachable.
 2. Put the host (or hosts) into maintenance mode to ensure no VM traffic is active on the host.
 3. Delete the VXLAN VTEPs from the host.
 4. Take the host out of maintenance mode. Taking hosts out of maintenance mode triggers the VXLAN VTEP creation process on NSX Manager. NSX Manager will try to re-create the required VTEPs on the host.

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 1485862: Uninstalling NSX from a host cluster sometimes results in an error condition

When using the Uninstall action on the **Installation: Host Preparation** tab, an error might occur with the

eam.issue.OrphanedAgency message appearing in the EAM logs for the hosts. After using the Resolve action and rebooting the hosts, the error state continues even though the NSX VIBs are successfully uninstalled.

Workaround: Delete the orphaned agency from the vSphere ESX Agent Manager (**Administration: vCenter Server Extensions: vSphere ESX Agent Manager**).

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.

Service virtual machine deployed using the Service Deployments tab on the Installation page does not get powered on

Workaround: Follow the steps below.

1. Manually remove the service virtual machine from the ESX Agents resource pool in the cluster.
2. Click **Networking and Security** and then click **Installation**.
3. Click the **Service Deployments** tab.
4. Select the appropriate service and click the **Resolve** icon.
The service virtual machine is redeployed.

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.

Workaround: Click on the red warning icon and select Resolve.

NSX Manager Known Issues

New 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.

New Issue 1781080: DNS search configuration fails when more than one domain is added, using comma separation

When adding more than one domain search suffix to NSX Manager, all DNS lookups not using the fully qualified name will fail. This is due to an internal formatting issue of /etc/resolv.conf.

Workaround: Use a single DNS search suffix.

New 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.

New 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 1770436: Alerts generated even when duplicate IP is not present

Sometimes the arping command reports that the NSX Manager IP address is duplicated in the network even though that is not the case. This generates a false positive event.

Workaround: Contact VMware customer support.

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: Do the following:

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

Issue 1713669: NSX Manager fails due to full disk when database table `ai_useripmap` grows too large

This issue causes the NSX Manager appliance disk to become full, resulting in the failure of the NSX Manager. The postgres process cannot be started after a reboot. The `"/common"` partition is full. This occurs most commonly in sites that place a heavy load on the Event Log Server (ELS) and in sites with a large amount of Guest Introspection (GI) traffic. Sites that use Identity Firewall (IDFW) are frequently affected. See [VMware Knowledge Base article 2148341](#) for more information.

Workaround: Contact VMware customer support for help in recovering from this issue.

Issue 1787542: Exceptions in secondary NSX Manager(s) log after DB restore on primary NSX Manager

After restoring DB on primary, reinstated Universal DFW sections not seen on secondary NSX Manager(s).

Workaround: None. Reboot the secondary NSX Manager to recover.

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: Perform the following steps:

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 1534877: 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 1537258: 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 1386874: Networking and Security Tab not displayed in vSphere Web Client

After vSphere is upgraded to 6.0, you cannot see the Networking and Security Tab when you log in to the vSphere Web Client with the root user name.

Workaround: Log in as administrator@vsphere.local or as any other vCenter user which existed on vCenter Server before the upgrade and whose role was defined in NSX Manager.

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 1477041: NSX Manager virtual appliance summary page shows no DNS name

When you log in to the NSX Manager virtual appliance, the Summary page has a field for the DNS name. This field remains blank even though a DNS name has been defined for the NSX Manager appliance.

Workaround: You can view the NSX Manager's hostname and the search domains on the Manage: Network page.

Issue 1492880: 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 1468613: 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: Perform the following steps:

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 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 1489768: Change in behavior of NSX REST API call to add a namespace in a datacenter

In NSX 6.2, the POST `https://<nsxmgr-ip>/api/2.0/namespace/datacenter/` REST API call returns a URL with an absolute path, for example `http://198.51.100.3/api/2.0/namespace/api/2.0/namespace/datacenter/datacenter-1628/2`. In previous releases of NSX, this API call returned a URL with a relative path, for example: `/api/2.0/namespace/datacenter/datacenter-1628/2`.

Workaround: None.

Logical Networking Known Issues and NSX Edge Known Issues

New Issue 1878081: Some of the routes are flushed from forwarding table on the edge services gateway

For few instances, some of the routes are flushed from forwarding table. This leads to a traffic outage.

Workaround: Reboot the edge node.

New 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.

New 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.

New 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.

New Issue 1849760: 100% CPU usage during and after physical network outage. The routing process may hog CPU when certain prefixes are learned over IBGP network such that the next-hop belongs to the prefix subnet.

Workaround: If you encounter this issue, it is recommended to filter out prefixes for iBGP network between ESG and DLR and advertise only the default route to iBGP.

New 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.

New 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.

New 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.

New 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.

New Issue 1861814: SSL VPN fails when using Mac OS Sierra or El Capitan on NSX

On an NSX for vSphere 6.3.x environment with SSL VPN configured and enabled, SSL VPN fails on a Mac with OS Sierra or El Capitan with the error:

"Connection to the gateway has been terminated due to server certificate validation failure."

Windows based machines are not affected.

Workaround: See [VMware Knowledge Base article 2150054](#) for more information and workaround.

New 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.

New Issue 1852988: ESG advertises BGP routes with neighbor address as next-hop

ESG advertises BGP routes to neighbor with next-hop as neighbor address, instead of local address when neighbor has "default originate" option enabled. This happens when uplink BGP session flaps.

Workaround: Flap the BGP session with the neighbor to which the next-hop routes were advertised as the neighbor's address, not the local address. This will make the Routing Table Manager re-compute and re-advertise the correct routes.

New Issue 1849760: Routing process may hog CPU when certain prefixes are learned over IBGP network such that the next-hop belongs to the prefix subnet

The problem is seen in a setup where DLR NSX Edge Appliance (Control VM) is peering with ESGs in IBGP and ESGs are peering with physical routers over EBGP. Also, the DLR Control VM is running HA. Aggressive timers are in use for BGP session. When underlying network starts flapping thereby inducing flaps for BGP neighbors, the DLR Control VM also will re-learn the prefixes for the flapping neighbor and will need to resolve them. When the learned prefix subnet covers the IP of the nexthop (example 10.0.0.0/8 via 10.1.1.2), the routing process may get into a recursive loop.

Workaround: If you must advertise such prefixes towards NSX topology, then ESGs can enable default-originate and block default route towards physical routers using BGP neighbor filters and permit only default route towards the DLR. In essence, prevent learning of prefixes.

New Issue 1798847: In a Cross-vCenter NSX setup, update of VXLAN UDP port can get stuck forever.

If primary NSX Manager does not have any secondary NSX Managers configured, then updating VXLAN UDP port hangs forever.

Workaround: Use NSX Manager API to resume port update workflow on primary NSX Manager.

Issue 1791264: Double clicking on a Transport Zone fails to enable/disable CDO mode.

If you try to enable or disable CDO mode from the Summary page that you reach after double-clicking a Transport Zone from the vSphere web client, there is no effect.

Workaround: Do the following:

1. Go back to the Transport Zone listing page: **Installation > Logical Network Preparation > Transport Zones** and select the desired Transport Zone.
2. Select **Enable CDO mode/Disable CDO mode** from the **Actions** dropdown menu.
3. The selected action will take effect.

Issue 1773500: Invalid route (0.0.0.0/32) causes NSX to crash

If you push the route 0.0.0.0/32 on the NSX DLR, it does not support this route and rejects it. However, this still causes a crash (PSOD) when the associated LIF is deleted and re-added with an IP address on the same subnet.

Workaround: 0.0.0.0/32 is not a valid route. Do not configure it, or use routemap to reject it.

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 1771760: SNMP response packets containing OID type Counter64 will get dropped by NSX Edge when NAT is enabled.

SNMP ALG in NSX Edge is unable to process Counter64 types from SNMP response packet, and the packet will get dropped. As a result, the client does not get a response for the request.

Workaround: Contact VMware customer support if you encounter this problem.

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 eth0 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 forms 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 1773127: On setups with a significant number of hosts and logical switches, the screen that displays hosts related to a given logical switch fails to load correctly.

When you select Logical Switch > Related Objects > Hosts from your setup with a significant number of hosts, the vSphere Web Client fails to load after waiting for a couple of minutes and the following error appears: The data service timed out because a back-end task took more than 120 seconds. This occurs because the remote API call to NSX Manager takes too long to return.

Workaround: There are two ways to work around this problem:

- The first option: You can avoid this issue by increasing the API timeout as described in the [VMware Knowledge Base article 2040626](#). You may need to restart the vSphere Web Client after increasing this timeout. The likely outcome of

increasing the timeout is that there will be no error, but you will have to wait for about 2-4 minutes for the page to reload.

- The second option: If you only want to see the related hosts correctly, you can go to Home > Networking > Port group > Related Objects > Hosts to see the list of hosts associated with the logical switch.

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 1770114: Error message at cluster level is not cleared after successful host preparation.

When you assign an IP pool to a cluster that does not have enough IP addresses, and then try to add a host to this cluster, you get the error "Insufficient IP addresses". Even after you change this pool to add additional IP addresses, and you can successfully add hosts to this cluster, the error message remains at the cluster level.

Workaround: Contact VMware customer support.

Issue 1789088: NSX Edge stuck in the grub command line prompt

NSX Edge may fail to boot up and can get stuck at the grub command line prompt.

Workaround:

- First, investigate:
 1. Check the existing environment with the set command.
 2. Use the ls and cat commands to locate and dump the /boot/grub/grub.cfg file.

```
grub> ls /boot
grub> ls /boot/grub
grub> cat /boot/grub/grub.cfg
```

 3. Capture the host logs at this time (as close to the problem occurring as possible). There may be some NFS logs that indicate an NFS storage issue.
- Next, boot the NSX Edge manually. Try the following, in this order (try the next option only if the previous one does not successfully boot the Edge):

1. Reboot the Edge VM by selecting the Power Reset option in the vSphere Web Client.
2. OR specify the grub config file again, which should load up the menu being boot up the Edge immediately. Invoke the following command at the grub prompt:

```
grub> configfile /boot/grub/grub.cfg
```

3. OR use the following commands at the grub prompt :

```
grub> insmod ext2
grub> set root=(hd0,1)
grub> linux /boot/vmlinuz loglevel=3 root=/dev/sda1
grub> boot
```

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 1676085: Enabling Edge HA will fail if resource reservation fails

Starting with NSX for vSphere 6.2.3, enabling high availability on an existing Edge will fail when sufficient resources cannot be reserved for the second Edge VM appliance. The configuration will roll back to the last known good configuration. In previous releases, if HA is enabled after Edge deployment and resource reservation fails, the Edge VM still is created.

Workaround: This is an expected change in behavior.

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 1556924: L3 connectivity loss with VXLAN would block error

When DLR LIF's are configured on the host but underlying VXLAN layer is not fully prepared, connectivity through some of DLR LIF's may be affected. Some of the VMs belonging to DLR are not reachable. There might be "Failed to Create VXLAN trunk status: Would block" logs in /var/log/vmkernel.log file.

Workaround: You may delete the LIF's and recreate them. Another option is rebooting the affected ESX hosts.

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 1604514: Editing/Configuring default gateway on an unmanaged DLR fails after clicking Publish

When a default gateway is added to an unmanaged DLR, the publish will fail with error "Routing Distance is support only on NSX Edge version 6.2.0 and later with NSX Edge VMs deployed". This is due to the default admin distance "1" populated on the UI.

Workaround: Remove the admin distance "1" which is populated by default.

Issue 1642087: After modifying the securelocaltrafficbyip parameter value in the IPsec VPN Extension, forwarding to destination networks fails

When using an NSX Edge Services Gateway, you experience this symptom:

- After changing the securelocaltrafficbyip value to 0 in the NSX UI (Edit IPsec VPN screen), forwarding to a remote subnet of the IPsec VPN tunnel no longer works
- After changing this parameter, you no longer see the correct information for a remote subnet in the IP routing table

Workaround: Disable and re-enable the IPSec VPN service. Then validate that the expected routing information is shown in the CLI and the UI.

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 1534603: 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 1534799: Slow convergence when OSPF area border router with highest IP address is shut down

Convergence takes a long time when the NSX-based, OSPF area border router (ABR) with highest IP address is shut down or rebooted. If an ABR that does not have the numerically highest IP address is shut down or rebooted, traffic converges quickly to another path. However, if the ABR with the highest IP address is shut down or rebooted, a multi-minute re-convergence time is seen. The OSPF process can be cleared manually to reduce the convergence time.

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 1089745: 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 1498965: 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 1494025: 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 1288487: 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?vmob=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.

Security Services Known Issues

New 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.

New Issue 1881696: If a policy is created with some firewall actions associated with it and later the policy is edited to delete all actions, then after the NSX Manager is restarted, the policy shows incorrect status as "in progress"

If you create a policy with actions from one category and then later edit the policy by deleting the action, then when the NSX Manager is rebooted, the policy shows incorrect status as "in progress".

Workaround: There is no functional impact of this issue. However, if the UI appears to be slow, delete the policy and recreate it.

New Issue 1859531: No SNMP traps sent for certain Application Rule Manager system events

SNMP traps are not sent for the following Application Rule Manager system events:

- flow analysis started
- flow analysis failed
- flow analysis completed

These are visible on the NSX Manager systems events monitor tab.

Workaround: No workaround required. Analysis proceeds as usual but SNMP traps are not generated to mark these events.

New Issue 1846402: Multiple IP addresses on same vNIC cause delays in firewall publish operation

The problem occurs when a vNIC has multiple IP addresses and ARP snooping is enabled. Each packet out of vNIC may result in an ARP message from host to manager with the newly discovered IP. This is because the host is able to send only one ARP snooped IP address to manager. When the vNIC switches between the IPs, host detects the IP as new and sends the ARP snooped message to manager. This usually results in container messages churn on the NSX Manager, resulting in delays in realizing firewall configuration on host.

Workaround: Disable ARP snooping when vNICs have multiple IP addresses. Use DHCP snooping or VMTools instead.

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 1676043: VM removed from Exclusion List after two simultaneous additions

Two simultaneous additions of the same VM to the Exclude List by two users without refreshing the UI results in the already added VMs being removed from the excluded list.

Workaround: Refresh the UI before adding the VM to the Exclude List.

Issue 1770259: The appliedTo field for the DFW rule cannot be modified to have multiple appliedTo objects

When you apply the DFW rule to a set of vNICs or VMs, or clusters or datacenter, publish it and later want to modify by adding additional objects to the appliedTo field, the new changes will not be effective even though the publish succeeds.

Workaround: None.

Issue 1798779: After you upgrade NSX from 6.2.x to 6.3.0, the GUI of vSphere Web Client) erroneously allows you to add Universal Security Tag

6.3.0 introduces Universal Security Tags. When you try to add a Universal Security Tag to a Universal Security Group that was created on 6.2.x before the upgrade to NSX 6.3.0, the operation will fail with the error "The requested member is not a valid member". This error is correct because you cannot add a Universal Security Tag to an NSX 6.2.x Universal Security Group. The GUI is misleading.

Workaround: After upgrading, create an NSX 6.3.0 Universal Security Group and add the Universal Security Tags to that group.

Issue 1799543: After upgrading from NSX 6.2.x to NSX 6.3.0, the vSphere Web Client erroneously shows and allows you to select NSX 6.2.x Universal Security Groups and non Active-Standby Universal Security Groups when you are creating the first Active-Standby Universal Security Group.

When you create the very first ActiveStandby Universal Security Group, the vSphere Web Client UI shows and allows you to add a Universal Security Group which was created on NSX 6.2.x. The operation will fail with the error "The requested member is not a valid member".

Workaround: Create at least one Active-Standby Universal Security Group, then while creating the next ActiveStandby Universal Security Group, this problem will not occur.

Issue 1786780: Reordering/moving Policies in the Service Composer UI take a long time with high CPU utilization

The reordering or re-positioning of Policies from the Service Composer UI may take a very long time with high CPU utilization.

Workaround: The following steps are helpful:

- While creating the Policy, try to give the right precedence (weight) to the Policy, so that it lands at the right position in the first attempt, and you do not need to reorder Policies again.
- If you must move a Policy to another position, edit the Policy to be moved and change the precedence (weight) to an appropriate value. This will result in the modification of a single Policy and finish quickly.

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 1741844: Using ARP snooping to detect address of vNIC with multiple IP addresses results in 100% CPU consumption

This issue occurs when a virtual machine's vNIC is configured with multiple IP addresses and ARP snooping is enabled for IP detection. The IP discovery module keeps sending vNIC-IP updates to the NSX Manager continuously to change the vNIC-IP mapping for all VMs configured with multiple IP addresses.

Workaround: There is no workaround. Currently the ARP snooping feature supports only one IP address per vNIC. For more information, see the section [IP Discovery for Virtual Machines](#) in the *NSX Administration Guide*.

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 1462027: In cross vCenter NSX deployments, multiple versions of saved firewall configurations get replicated to secondary NSX Managers

Universal Sync saves multiple copies of universal configurations on secondary NSX Managers. The list of saved configurations contains multiple drafts created by the synchronizing across NSX Managers with the same name and at the same time or with a time difference of 1 second.

Workaround: Run the API call to delete duplicate drafts.

DELETE : <https://<nsxmgr-ip>/api/4.0/firewall/config/drafts/>

Find the drafts to be deleted by viewing all drafts:

GET: <https://<nsxmgr-ip>/api/4.0/firewall/config/drafts>

In the following sample output, drafts 143 and 144 have the same name and were created at the same time and are therefore duplicates. Likewise, drafts 127 and 128 have the same name are off by 1 second and are also duplicates.

```
<firewallDrafts>
  <firewallDraft id="144" name="AutoSaved_Wednesday, August 5, 2015 11:08:40 PM GMT" timestamp="1438816120917">
    <description>Auto saved configuration</description>
    <preserve>false</preserve>
    <user>replicator-1fd96022-db14-434d-811d-31912b1cb907</user>
    <mode>autosaved</mode>
  </firewallDraft>
  <firewallDraft id="143" name="AutoSaved_Wednesday, August 5, 2015 11:08:40 PM GMT" timestamp="1438816120713">
    <description>Auto saved configuration</description>
    <preserve>false</preserve>
    <user>replicator-1fd96022-db14-434d-811d-31912b1cb907</user>
    <mode>autosaved</mode>
  </firewallDraft>
  <firewallDraft id="128" name="AutoSaved_Wednesday, August 5, 2015 9:08:02 PM GMT" timestamp="1438808882608">
    <description>Auto saved configuration</description>
    <preserve>false</preserve>
    <user>replicator-1fd96022-db14-434d-811d-31912b1cb907</user>
    <mode>autosaved</mode>
  </firewallDraft>
  <firewallDraft id="127" name="AutoSaved_Wednesday, August 5, 2015 9:08:01 PM GMT" timestamp="1438808881750">
    <description>Auto saved configuration</description>
    <preserve>false</preserve>
    <user>replicator-1fd96022-db14-434d-811d-31912b1cb907</user>
    <mode>autosaved</mode>
  </firewallDraft>
</firewallDrafts>
```

Issue 1449611: When a firewall policy in the Service Composer is out of sync due to a deleted security group, the firewall rule cannot be fixed in the UI

Workaround: In the UI, you can delete the invalid firewall rule and then add it again. Or, in the API, you can fix the firewall rule by deleting the invalid security group. Then synchronize the firewall configuration: Select **Service Composer: Security Policies**, and for each security policy that has associated firewall rules, click **Actions** and select **Synchronize Firewall Config**. To prevent this issue, modify firewall rules so that they do not refer to security groups before deleting the security groups.

Issue 1557880: Layer 2 (L2) rules may be missing if the MAC address of a VM used in the rules is modified

Because L2 rule optimization is ON by default, L2 rules with both source and destination fields specified (other than "any") will be applied to vNICs(or filters) only if the vNIC MAC address matches the source or destination MAC address list. Hosts with VMs not matching the source or destination MAC addresses will not have those L2 rules applied.

Workaround: To have L2 rules applied to all vNICs(or filters), set one of the source or destination fields to "any".

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 1557924: Universal logical switch is allowed to be consumed in the appliedTo field of a local DFW rule

When a universal logical switch is used as a security group member, the DFW rule can use that security group in AppliedTo field. This indirectly applies the rule on the universal logical switch, which should not be allowed because it may cause unknown behavior of those rules.

Workaround: None.

Issue 1559971: Cross-vCenter NSX firewall exclude list not published if firewall is disabled on one cluster

In cross-vCenter NSX, firewall exclude list is not published to any cluster when the firewall is disabled on one of the clusters.

Workaround: Force sync the affected NSX Edges.

Issue 1407920: Firewall rule republish fails after DELETE API is used

If you delete the entire firewall configuration through the DELETE API method and then try to republish all the rules from a previously saved firewall rules draft, then the rule publish will fail.

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 1442379: Service composer firewall configuration out of sync

In the NSX service composer, if any firewall policy is invalid (for example of you deleted a security group that was currently in use in a firewall rule), deleting or modifying another firewall policy causes the service composer to become out of sync with the error message Firewall configuration is not in sync.

Workaround: Delete any invalid firewall rules and then synchronize the firewall configuration. Select **Service Composer: Security Policies**, and for each security policy that has associated firewall rules, click **Actions** and select **Synchronize Firewall Config**. To prevent this issue, always fix or delete invalid firewall configurations before making further firewall configuration changes.

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 1620491: Policy-level Sync status in Service Composer does not show publishing status of the rules within a policy

When a policy is created or modified, Service Composer will display a success status which indicates only the persistence state. It does not reflect whether the rules were published to the host successfully.

Workaround: Use the firewall UI to view publish status.

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.

Issue 1772504: Service Composer does not support Security Groups with MAC Set

Service Composer allows use of Security Groups in Policy configurations. In case there is a Security Group which contains

MAC Set, Service Composer accepts that Security Group without complaining, but fails to enforce rules for that specific MAC Set. This is because Service Composer works on Layer3 and does not support Layer2 constructs. Note that if a Security Group has IP Set and MAC Set both, the IP set will still be effective, but the MAC Set will be ignored. There is no harm in referencing a Security Group containing MAC Set - user must be aware that the MAC Set will be ignored.

Workaround: If the user's intent is to create Firewall rules using a MAC Set, then the user should use DFW Layer2/Ethernet configuration instead of Service Composer.

Issue 1718726: Cannot force-sync Service Composer after a user has manually deleted the Service Composer's policy section using DFW REST API

In a cross-vCenter NSX environment, a user's attempt to force sync NSX Service Composer configuration will fail if there was only one policy section and that policy section (the Service Composer-managed policy section) was deleted earlier via a REST API call.

Workaround: Do not delete the Service Composer-managed policy section via a REST API call. (Note that the UI already prevents deletion of this section.)

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.

Issue 1626233: When NetX service virtual machine (SVM) drops packets, traceflow does not generate dropped observation

The traceflow session exits after packet is sent to the NetX service virtual machine (SVM). When the SVM drops packets, traceflow does not generate dropped observation.

Workaround: There is no workaround. If the traceflow packet is not injected back, it can be assumed that the SVM dropped the packet.

Solution Interoperability Known Issues

Issue 1568861: The NSX Edge deployment fails during any edge deployment from a vCloud Director cell that does not own the vCenter listener

The NSX Edge deployment fails during any Edge deployment from a vCloud Director cell that does not own the vCenter listener. Also, NSX Edge actions, including a redeploy, fail from vCloud Director.

Workaround: Deploy an NSX Edge from the vCloud Director cell which owns the vCenter listener.

NSX Controller Known Issues

Issue 1765354: <deployType> is a required property but it is not used <deployType> is a required property but it is not used and does not mean anything.

Issue 1516207: Controller(s) may become isolated after IPsec communication is re-enabled on an NSX controller cluster

If an NSX controller cluster is set to allow controller-to-controller communications in the clear (IPsec is disabled), and IPsec-based communication is later re-enabled, one or more controllers may become isolated from the cluster majority due to a mismatched pre-shared key ("PSK"). When this occurs, the NSX API may become unable to change the IPsec settings of the controllers.

Workaround:

Follow these steps to address this issue:

1. Disable IPSec using the NSX API.

```
PUT /2.0/vdn/controller/node
```

```
<controllerNodeConfig>
  <ipSecEnabled>false</ipSecEnabled>
</controllerNodeConfig>
```

2. Re-enable IPsec using the NSX API.

```
PUT /2.0/vdn/controller/node
```

```
<controllerNodeConfig>
  <ipSecEnabled>true</ipSecEnabled>
</controllerNodeConfig>
```

Follow these best practices to avoid this issue:

- Always use the NSX API to disable IPsec. Using the NSX Controller CLI to disable IPsec is not supported.
- Always verify that all controllers are active before you use the API to change the IPsec setting.

Issue 1306408: NSX Controller logs must be downloaded sequentially

NSX Controller logs cannot be downloaded simultaneously. Even when downloading from multiple controllers, you must wait for the download from the current controller to finish before you start the download from the next controller. Note also that you cannot cancel a log download once it has started.

Workaround: Wait for the current controller log download to finish before starting another log download.

Resolved Issues

See what's resolved in [6.3.2](#) or in [6.3.1 and earlier](#).

New Issues Resolved in NSX 6.3.2

6.3.2 resolved issues are grouped as follows:

- [General Resolved Issues in NSX 6.3.2](#)
- [Installation and Upgrade Resolved Issues in NSX 6.3.2](#)
- [NSX Manager Resolved Issues in NSX 6.3.2](#)
- [Networking and Edge Services Resolved Issues in NSX 6.3.2](#)
- [Security Services Resolved Issues in NSX 6.3.2](#)

General Resolved Issues in NSX 6.3.2

Fixed Issue 1839275: VMs lose connectivity with physical servers behind the Hardware VTEP switches

When both data plane and control plane exist in host VMs and the control plane goes down by the physical unplugging of all network cables, the VMs lose connectivity with the physical servers located behind the Hardware VTEP switches. *Fixed in 6.3.2.*

Fixed Issue 1802366: Hosts and controllers out of sync for Logical (Distributed) Routers because of buffer overflow on the controllers

Sync status of one or more Logical (Distributed) Routers is shown as out-of-sync for a prolonged time(> 10 secs).

On controller CLI:

```
nsx-controller # show control-cluster logical-routers instance all
LR-Id LR-Name Universal Service-Controller Egress-Locale In-Sync Sync-
0x138b default+edge-21 false 10.165.26.22 local No NORMAL
0x138a default+edge-19 false 10.165.26.22 local No NORMAL
```

```
nsx-controller # show control-cluster logical-routers error-status 0x138b
Id IP Version Locale-Id Sync-State
67 50.6.0.21 6.2 00000000-0000-0000-0000-000000000000 SYNCHING
68 50.6.0.19 6.2 00000000-0000-0000-0000-000000000000 SYNCHING
69 50.6.0.22 6.2 00000000-0000-0000-0000-000000000000 SYNCHING
```

Fixed in 6.3.2.

Installation and Upgrade Resolved Issues in NSX 6.3.2

Fixed Issue 1811218: VMKernel port of non-NSX prepped vSphere Distributed Switches (VDS) lose connectivity after host reboot.

On hosts with multiple VDS, VMKernel port of non-NSX prepped dvSwitches lose connectivity after host reboot, if they

have the same dvPort ID as the VTEP VMKernel port of the NSX prepped VDS. See [VMware Knowledge Base article 2149116](#) for more information. *Fixed in 6.3.2.*

Fixed Issue 1850690: Guest Introspection USVM reports disk is full

In a VMware NSX for vSphere 6.3.x environment where Guest Introspection is deployed, you receive an alert that the /var/log disk space is full or almost full on one or more of the Guest Introspection USVMs. The same alert can be seen in vRealize Operations Manager or in any other monitoring tool that tracks virtual machine disk space. *Fixed in 6.3.2.*

Fixed Issue 1854519: VMs lose north to south connectivity after migration from a VLAN to a bridged VXLAN

As soon as the VM network is switched from a VLAN to a bridged VXLAN on DLR, the ingress traffic to the VM is lost. *Fixed in 6.3.2.*

NSX Manager Resolved Issues in NSX 6.3.2

Fixed Issue 1814683: NSX Manager become unresponsive when multiple HTTP threads stuck in authorization check

NSX Manager become unresponsive and needs to be restarted. *Fixed in 6.3.2.*

Networking and Edge Services Resolved Issues in NSX 6.3.2

Fixed Issue 1681063: VPN Tunnel Status for some Edge Gateways are not being accurately reflected in vCloud Director

vCloud Director (vCD) and NSX show different status for IPSec VPN Tunnel. vCD shows IPSec Tunnel as DOWN even when it is up and passing traffic. *Fixed in 6.3.2.*

Fixed Issue 1827643: Unicast flooding during controller down event

When there is a controller down event, all L3 data packets are sent to all the VTEPs on that logical switch shortly after the controllers go down. This is because the default L2 MAC aging time is lower than the L3 ARP timeout and controllers are not available for refreshing the in-use MAC table entries. The flooding can cause transmit resets in NICs which cannot sustain that much load. This can lead to packet drops of heartbeat protocols, affecting HA applications. For example, DLR with the NSX Edge appliance and HA enabled may get into a split brain situation leading to OSPF between DLR and ECMP ESG to flap. *Fixed in 6.3.2.*

Fixed Issue 1793575: All NSX traffic egresses out of one NIC when IP Hash teaming is selected

When NSX portgroups are configured with an IP Hash teaming policy (Static EtherChannel or LACP), all traffic from those portgroups will egress out of a single NIC. This does not cause traffic loss, but network load is not distributed across all available links. On failure of the NIC in use, all NSX portgroups will fail over to another active NIC in the team. Ingress traffic is not affected, and can be received on all available links. *Fixed in 6.3.2.*

Fixed Issue 1812445: L2VPN bridge table "poisoned" by DLR PMAC with duplicate ARP reply

The L2VPN server vxlan trunk port on the host does not drop the ARP reply coming from the client virtual machine with destination MAC as pMAC, which causes the MAC table on bridge getting poisoned resulting in traffic drop. *Fixed in 6.3.2.*

Fixed Issue 1806934: Openswan's CVE-2015-3240 causes Denial of Service (DoS)

IKE daemon in libreswan before 3.15 and Openswan before 2.6.45, when built with NSS, allow remote attackers to cause a denial of service (assertion failure and daemon restart) via a zero DH g^x value in a KE payload in an IKE packet. *Fixed in 6.3.2.*

Fixed Issue 1811884: Logical switch is not removed in Universal Synchronization Service (Replicator) host

It is not possible to remove a logical switch in Universal Synchronization Service (Replicator) host. *Fixed in 6.3.2.*

Fixed Issue 1803220: Loss of VXLAN connectivity to CDO-enabled hosts when controller to host connection goes down

Controller Disconnected Operation(CDO) feature ensures VXLAN connectivity when the whole Controller cluster is down/unreachable. However, in cases where the Controller cluster is Up, but a host loses connectivity with it, data plane traffic destined to that host from other hosts which are connected with the Controller may still be dropped. When this condition occurs, the host has been removed from the per-VNI VTEP list, and ARPs sent by remote hosts will be dropped. For traffic originating from the host which has lost connectivity with the Controller, CDO feature ensures that it will be able to reach the right destination. *Fixed in 6.3.2.*

Fixed Issue 1817673: ESXi host crashes and routing problems observed when DLR LIF IP is changed and the previous IP is assigned to a VM

The purple diagnostic screen is encountered on hosts where the bridge is active. DLR routing problems are faced on all hosts when communicating with the VM whose IP address is the previous DLR LIF's IP address. *Fixed in 6.3.2.*

Fixed Issue 1825416: Fenced vApps fail in vCloud Director 8.20 after upgrading to NSX for vSphere 6.3.x

After upgrading to NSX 6.3.x and NSX Edge Gateways to 6.3.x in vCloud Director 8.20, fenced vApps fail and Virtual

Machines in a fenced network fail to communicate to their gateway. *Fixed in 6.3.2.*

Fixed Issue 1842337: DLR ARP resolution for ECMP ESG fails.

ARP resolution at the DLR for the ECMP fails and impacts North-South traffic. *Fixed in 6.3.2.*

Fixed Issue 1834837: All vCNS Edges (version 5.5.4) and NSX Edges (6.1.x and 6.2.x) running on VIX communication mode get unmanageable after storage vMotion

NSX Edges running with MessageBus as communication mode are not impacted. *Fixed in 6.3.2.*

Fixed Issue 1798469: Some NSX Edges can get into split brain scenarios even if their Highavailability status is shown correctly

Due to a race condition in the HA mechanism, some NSX edges upgraded to NSX 6.2.5 and later could get into a split brain scenario even if their "Highavailability Status" is shown correctly. This could also happen after edge redeployment. *Fixed in 6.3.2.*

Fixed Issue 1828060: NSX Edge can experience connectivity issue under high traffic during force sync operation.

NSX Edge can experience connectivity issue under high traffic during the Force Sync operation. *Fixed in 6.3.2.*

Security Services Resolved Issues in NSX 6.3.2

Fixed Issue 1847753: Hosts fail with a purple diagnostic screen when retrieving flows for ALG enabled protocols

After upgrading NSX for vSphere 6.2.4 to 6.3.0 or 6.3.1 with Flow Monitoring enabled in the environment, the ESXi host experiences a purple diagnostic screen. *Fixed in 6.3.2.*

Fixed Issue 1798537: DFW controller process on ESXi (vsfwd) may run out of memory

In an environment with a large number of rules or large size security groups in DFW configuration, DFW controller process on ESXi (vsfwd) may run out of memory and fail to publish rules to datapath. *Fixed in 6.3.2.*

Fixed Issue 1853106: Error received on UI when the certification is unchecked or some changes on IPsec configuration are modified for IPsec VPN under PSK mode

You see the error "Failed to read certs & secrets. : 002 forgetting secrets" on the Global Configuration tab when PSK mode certification is unchecked or modified for IPsec VPN with certificate-based authentication. *Fixed in 6.3.2.*

Fixed Issue 1725524: Hostd restart could cause vsfwd to crash when it's sending data using RabbitMQ

Any operations to restart the hostd process will cause re-established connections for vsfwd. The vsfwd could crash when the main thread is clearing the old connection while some thread is still using the old channel to send data. *Fixed in 6.3.2.*

Fixed Issue 1800196: VMkernel logging stops if a large number of IP packets with broadcast MAC address match a Distributed Firewall reject rule

Distributed Firewall sends reject packets to unicast MAC addresses only. When IP packets with a broadcast MAC address match a reject rule, no reject packet is sent. However, this event is logged in vmkernel.log. If the network is flooded with this traffic, vmkernel.log drops messages due to log stress, and logging stops. Reject of packets with broadcast MAC address is now logged only if debug is enabled. *Fixed in 6.3.2. Rejection of packets with broadcast MAC address is now logged only if debug is enabled.*

Fixed Issue 1807152: When a VM is moved from one cluster to another newly prepared cluster, rules having Virtual Wire as appliedTo are not applied to the VM.

When a new cluster is added to the Datacenter, the cluster list for rules having Applied To as Virtual Wire is not updated. So, when a VM is migrated from the old cluster to the newly prepared cluster, the rules are not applied. *Fixed in 6.3.2.*

Fixed Issue 1812467: Container updates in large scale VMs with nested SG case will experience long delays for updates in the host

When the host has large scale VMs with nested SG organized, one single inventory change will cause a container update flush, resulting in a huge delay for the change to be applied eventually on the host. *Fixed in 6.3.2.*

Fixed Issue 1847880: Unloading VIBs from hosts in a given cluster fails with the error: "vmkload_mod: Can not remove module vsip: Consumed resource count of module is not zero"

After the uninstall of the VIBs is triggered and other NSX VIBs are uninstalled, the command `esxcli software vib list | grep vsip` shows the esx-vsip still installed. *Fixed in 6.3.2.*

Fixed Issue 1812792: DFW dropping packets related to split TCP handshake

Distributed Firewall does not handle the split TCP handshake properly. It ignores the Window Scale Option leading to future packet drops as the window calculation is incorrect. This was being caused by the first SYN ACK packet getting lost and the retransmission of the SYN was causing a Challenge ACK being generated which caused the split TCP Handshake. *Fixed in 6.3.2.*

The following issues were resolved in earlier 6.3.x releases:

General Resolved Issues in 6.3.1 and earlier

Fixed Issues 1575342, 1719402: In an NSX for vSphere 6.x environment, when a Service VM (SVM) is migrated (vMotion/SvMotion), there may be interruption in the service or the ESXi host may crash

Starting in 6.3.0, you cannot migrate a Service VM (SVM) using vMotion/SvMotion. SVMs must remain on the host on which they were deployed for correct operation.

Previously migration to another host was allowed, but not supported, and resulted in interrupted service and problems with the host.

See [VMware Knowledge Base article 2141410](#) for more information. *Fixed in 6.3.0.*

Fixed Issue 1708769: Increased latency on SVM (Service VM) after snapshot in NSX

This issue occurs because running a snapshot of a Service VM (SVM) can cause added network latency. Snapshot is sometimes invoked by backup applications running in the environment. *Fixed in 6.3.0.*

Fixed Issue 1760102: Virtual machines may fail to communicate after an NSX controller is deleted and redeployed to recover from a storage outage

An NSX Controller for the vSphere 6.2.4/6.2.5 environment may get into read-only mode in case of a storage outage and if you delete and redeploy the controller to recover from that state, some VMs may fail to communicate. Expected behavior in case of a storage outage on a controller is that rebooting of the controller should recover it from read-only mode, but currently that does not happen in NSX. *Fixed in 6.3.0.*

Fixed Issue 1662842: Guest Introspection: Connectivity lost between MUX and USVM when trying to resolve unresolvable Windows SIDs

Guest Introspection service will go into a warning state, with each Guest Introspection going in and out of a warning state. Until the Guest Introspection VM reconnects, network events will not be delivered to the NSX Manager. This will affect both Activity Monitoring and ID Firewall in the case that logon events are detected through the Guest Introspection path. *Fixed in 6.3.0.*

Issue 1752051: Service Status for Guest Introspection reported as "Not Ready" when NSX Manager to USVM communication times out

An error message similar to "PLAIN login refused: user 'usvm-admin-host-14' - invalid credentials" may be reported for the Guest Introspection Universal SVM when the expected password change process with NSX Manager on the internal message bus (rabbit MQ) does not succeed.

Workaround: To re-establish connectivity between the USVM and NSX Manager, restart the USVM or manually delete it and then select the Resolve button on the Service Composer UI to prompt a redeploy of the USVM for the affected host only.

Fixed Issue 1716328: Removing host that is in maintenance mode can result in later cluster preparation failure

If an administrator places an NSX-enabled ESXi host in maintenance mode and removes it from an NSX-prepared cluster, NSX fails to delete its record of the ID number of the removed host. After the installation is in this state, if there is another host with same ID in another cluster or if this host is being added to another cluster, the cluster preparation process will fail for that cluster.

Fixed in 6.3.0.

Fixed Issue 1710624: Windows 2008 event log server is added as "TYPE" of "WIN2K3", if serverType is not specified in REST API request body

If you create EventLog server API request, the server will be added as "TYPE" of "WIN2K3". If you use EventLog server only for IDFW, IDFW may not work correctly. *Fixed in 6.3.0.*

Installation and Upgrade Resolved Issues in 6.3.1 and earlier

Fixed Issue 1808478: vsfwd service fails to start if vmvisor memory cannot be allocated after upgrading from NSX 6.2.x to NSX 6.3.0

The vsfwd service fails to start if vmvisor memory cannot be allocated after upgrading from NSX 6.2.x to NSX 6.3.0. See [VMware knowledge base article 2148974](#) for more information. *Fixed in 6.3.1*

Fixed Issue 1818257: VTEP information is not reported to controllers when Enhanced LACP is used for VXLAN post host-upgrade from NSX 6.2.x to NSX 6.3.0 with ESXi 6.0

While upgrading from NSX 6.2.x to 6.3.0 with ESXi 6.0, after host-upgrade, VTEP information is not reported to controllers when Enhanced LACP is used. See [VMware knowledge base article 2149210](#) for more information. *Fixed in 6.3.1.*

Fixed Issue 1495307: During an upgrade, L2 and L3 firewall rules do not get published to hosts

After publishing a change to the distributed firewall configuration, the status remains InProgress both in the UI and the API indefinitely, and no log for L2 or L3 rules is written to the file vsfwd.log. *Fixed in 6.3.0.*

Fixed Issue 1491820: NSX Manager log collects WARN messagingTaskExecutor-7 messages after upgrade to NSX 6.2

After upgrading from NSX 6.1.x to NSX 6.2, the NSX Manager log becomes flooded with messages similar to: WARN messagingTaskExecutor-7 ControllerInfoHandler:48 - host is unknown: host-15 return empty list. There is no operational impact. *Fixed in 6.3.0.*

NSX Manager Resolved Issues in 6.3.1 and earlier**Fixed Issue 1671067: NSX Plugin does not appear in vCenter Web Client while ESXTOP plugin is also installed**

After deployment of NSX and successful registration with vCenter, NSX plugin does not appear in the vCenter Web Client. This issue is caused by conflict between NSX plugin and ESXTOP plugin. *Fixed in 6.3.0.*

Networking and Edge Services Resolved Issues in 6.3.1 and earlier**Fixed Issue 1740231: Unable to add IP address on the HA interface**

Starting in 6.3.0, you can add IP addresses on the DLR HA interface. This functionality was unavailable in some of the older NSX versions but has been reintroduced to match API behavior of the DLR HA management interface. *Fixed in 6.3.0*

Fixed Issue 1716333: Changing the Edge VM size or a placement parameter while enabling or disabling Edge HA may create extra Edge VMs

Simultaneous operations of changing the Edge VM size or a placement parameter (such as a datastore or resource pool) and enabling or disabling Edge HA may corrupt the NSX managed object database, leaving behind unusable Edge VMs. In addition, in an environment with cross-vCenter, the left-behind Edge VMs will be replicated to the secondary site. *Fixed in 6.3.0.*

Fixed Issue 1717369: When configured in HA mode, both active and standby Edge VMs may be deployed on the same host.

This issue results from anti-affinity rules not being created and applied on the vSphere hosts automatically during redeploy and upgrade operations. This issue will not be seen when HA is being enabled on existing Edge.

Fixed in 6.3.0. The following is the expected behavior:

- When vSphere HA is enabled, anti-affinity rules for Edge VMs of an HA pair will be created during redeploy, upgrade.
- When vSphere HA is disabled, anti-affinity rules for Edge VMs of an HA pair will not be created.

Fixed Issue 1675659: Floating Static Routes are preferred to OSPF Dynamic Routes

A backup Floating Static Route is incorrectly entered into an Edge's routing table when Route Redistribution is enabled even though an OSPF Route is available. *Fixed in 6.3.0.*

Fixed Issue 1733165: IPsec may cause removal of dynamic routes from NSX Edge forwarding table

If a subnet reachable via dynamic route is used as a remote subnet for IPsec configuration, NSX Edge removes this subnet from the forwarding table and does not reinstall it even after this subnet is deleted from the IPsec configuration. *Fixed in 6.3.0.*

Fixed Issue 1663902: Renaming an NSX Edge VM disrupts traffic flowing through the Edge

Renaming an NSX Edge VM disrupts traffic flowing through the Edge. *Fixed in 6.3.0.*

Fixed Issue 1624663: After clicking "Configure Advanced Debugging" refreshes the vCenter UI and the change does not persist

After clicking the specific edge ID > Configuration > Action > Configure Advanced Debugging causes the vCenter UI to refresh and the change does not persist. *Fixed in 6.3.0.*

Fixed Issue 1706429: Communication issues when enabling high availability (HA) after initial logical (distributed) router deployment might cause both logical router appliances to be active.

If you deploy a logical router without HA and then later enable HA (deploying a new logical router appliance), or if you disable and then re-enable HA, sometimes one of the logical router appliances is missing a connected route to the HA interface. This causes both appliances to be in the active state. *Fixed in 6.3.0.*

Fixed Issue 1542416: Data path not working for 5 min after edge re-deploy and HA failover with sub-interfaces

Redeploy or HA failover operation will see a five minute outage if sub-interfaces are used. Issue is not observed on interfaces. *Fixed in 6.3.0.*

Fixed Issue 1492547: Extended convergence time seen when NSX-based OSPF area border router with highest IP address is shut down or rebooted

If an NSSA area border router which does not have the highest IP address is shut down or rebooted, traffic converges rapidly to another path. If an NSSA area border router with the highest IP address is shut down or rebooted, a multi-minute re-convergence time is seen. The OSPF process can be cleared manually to reduce the convergence time. *Fixed in 6.3.0.*

Fixed Issue 1510724: Default routes do not populate on the hosts after creating a new Universal Distributed Logical Router (UDLR)

After changing NSX Manager from Standalone to Primary mode for the purpose of configuring Cross-vCenter in NSX for vSphere 6.2.x, you may experience these symptoms:

- When you create a new UDLR, the default routes are not populated on the host instance.
- Routes are populated on the UDLR NSX Edge Appliance (Control VM) but not on the host instance.
- Running the *show logical-router host host-ID dlr Edge-ID route* command fails to show default routes.

Fixed in 6.3.0.

Fixed Issue 1704540: High volume of MAC learning table updates with NSX L2 bridge and LACP may lead to out of memory condition

When an NSX L2 bridge sees a MAC address on a different uplink, it reports a MAC learning table change to controllers via the netcpa process. Networking environments with LACP will learn the same MAC address on multiple interfaces, resulting in a very high volume of table updates and potentially exhausting the memory needed by the netcpa process to do the reporting. See [VMware Knowledge Base article 2147181](#). *Fixed in 6.3.0.*

Fixed Issue 1716545: Changing appliance size of Edge does not affect standby Edge's CPU and Memory reservation

Only the first Edge VM created as part of an HA pair is assigned the reservation settings.

Fixed Issue 1772004: Edge HA Failover from node 0 to node 1 takes longer than expected

Failover from node 0 to node 1 takes longer than expected in Edge HA configured environment whereas failover of traffic from node 1 to node 0 is normal. *Fixed in 6.3.0.*

Fixed Issue 1726379: If IP multicast range has an upper bound value exceeding 99 in the last three octets, VXLAN trunk portgroup configuration fails.

While configuring segment ID, if you create a multicast IP range with an upper bound value exceeding 99 in the last three octets, for example, 1.100.100.100, and a multicast or hybrid logical switch with the same multicast IP range, VXLAN trunk portgroup configuration will fail. *Fixed in 6.3.0.*

Security Services Resolved Issues in 6.3.1 and earlier

Fixed Issue 1767402: DFW rules with "Applied To" set to a "Security Group" are not published to hosts

DFW rules with the "Applied To" field set to Security Group are not pushed to ESXi hosts in a new cluster. *Fixed in 6.3.0.*

Fixed Issue 1743366: NSX Threshold Monitoring is disabled by default to avoid a potential crash

When the Firewall module runs, NSX disables the threshold monitoring for memory to avoid a potential crash. When the host is running ESX 6.5P01 or ESX 6.0U3 or above, the memory threshold monitoring will be automatically enabled. *Fixed in 6.3.0.*

Fixed Issue 1491046: IPv4 IP address does not get auto approved

IPv4 IP address does not get auto approved when SpoofGuard policy is set to Trust On First Use (TOFU) in VMware NSX for vSphere 6.2.x. *Fixed in 6.3.0.*

Fixed Issue 1686036: Firewall rules cannot be added, edited, or removed when default section is deleted

If the default Layer2 or Layer3 section is deleted, publishing a firewall rule may fail. *Fixed in 6.3.0.*

Fixed Issue 1717994: Distributed Firewall (DFW) Status API query reports 500 internal server error intermittently

If the DFW status API query is issued while adding a new host into a host prepared cluster, the API query fails with 500 internal server error for few attempts, and then returns correct response once the host starts to get VIBs installed. *Fixed in 6.3.0.*

Fixed Issue 1717635: Firewall configuration operation fails if more than one cluster is present in environment and changes are done in parallel

In an environment with multiple clusters, if two or more users modify the firewall configuration continuously in a tight loop. (for example, Add/Delete sections or rules), some operations fail, and the user will see an API response similar to:

```
org.hibernate.exception.GenericJDBCException: Could not execute JDBC batch update; nested exception is  
javax.persistence.PersistenceException: org.hibernate.exception.GenericJDBCException: Could not execute JDBC  
batch update
```

Fixed in 6.3.0.

Issue 1707931: Order of distributed firewall rules changes when service policies defined in Service Composer are present, and a firewall rule is modified or published with a filter applied in the Firewall UI

Changing the order, adding or deleting service policies created in Service Composer after one or more publish operations are made from the Networking & Security > Firewall UI will cause the order of firewall rules to change and may result in unintended consequences. *Fixed in 6.3.0.*

Fixed Issue 1682552: Threshold events for CPU/Memory/CPS for Distributed Firewall (DFW) are not reported

Even when the DFW thresholds for CPU/Memory/CPS are set for reporting, the threshold events are not reported when the thresholds are crossed. *Fixed in 6.3.0.*

Fixed Issue 1620460: NSX fails to prevent users from creating rules in Service Composer rules section

In the vSphere Web Client, the Networking and Security: Firewall interface fails to prevent users from adding rules to the Service Composer rules section. Users should be permitted to add rules above/below the Service Composer section, but not inside it. *Fixed in 6.3.0.*

Fixed Issue 1445897: Publishing Distributed Firewall (DFW) rules fails after referenced object is deleted in VMware NSX for vSphere 6.1.x and 6.2.x *Fixed in 6.2.3.*

Fixed Issue 1704661, 1739613: VMs lose network connectivity with the error: "Failed to restore PF state: Limit exceeded"

VMs lose network connectivity with the error: "Failed to restore PF state: Limit exceeded." *Fixed in 6.3.0.*

Solution Interoperability Resolved Issues in 6.3.1 and earlier

Fixed Issue 1527402: Windows VM with NSX Network Introspection driver lose TCP connectivity

In VMware NSX for vSphere 6.x environment, Windows VM with NSX Network Introspection driver (vnetflt.sys) connected to USVM (Guest Introspection SVM) loses temporary TCP network connectivity. *Fixed in 6.3.0.*

Fixed Issue 1530360: After an NSX Manager VM has failed over, Site Recovery Manager (SRM) incorrectly reports a timeout error

When an NSX Manager VM is failed over, SRM incorrectly reports a timeout error waiting for VMware Tools. In this case, VMware Tools actually is up and running within the 300 second timeout. *Fixed in 6.3.0.*

Document Revision History

08 June 2017: First edition for NSX 6.3.2.

22 June 2017: Second edition for NSX 6.3.2. Added fixed issue 1847753.

11 August 2017: Third edition for NSX 6.3.2. Added note about Open VM Tools support.

18 August 2017: Fourth edition for NSX 6.3.2. Added 1878081.

21 August 2017: Fifth edition for NSX 6.3.2. Deleted few former issues.

02 October 2017: Sixth edition for NSX 6.3.2. Updated recommended versions.

27 October 2017: Seventh edition for NSX 6.3.2. Added fixed issue 1802366.