VMware NSX for vSphere 6.2.8 Release Notes

VMware NSX for vSphere 6.2.8 | Released 06 July 2017 |
Build 5901733

See the Revision History of this document.

What's in the Release Notes

The release notes cover the following topics:

- What's New
- Versions, System Requirements, and Installation
- Deprecated and Discontinued Functionality
- Upgrade Notes
- Revision History
- Resolved Issues
- Known Issues

What's New

Critical bug fixes: This release includes a number of critical bug fixes and security updates.Also, the logging information related to VTEP is improved. For details, see the Resolved Issues section.

See what's new and changed in NSX6.2.7, 6.2.6, 6.2.5, 6.2.4, 6.2.3, 6.2.2, 6.2.1 and 6.2.0.

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
	VMware recommends the latest NSX for vSphere release for new deployments. 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.

NSX for vSphere	information on specific issues before planning an upgrade. • NSX for vSphere 6.2.4 resolves a known issue with SSL VPN. For more information, see CVE-2016-2079. Customers running 6.2.2 or earlier are strongly advised to upgrade. • NSX for vSphere 6.2.4 with vSphere 6.0 Update 3 resolves the issue of duplicate VTEPs in ESXi hosts after rebooting vCenter server. See *VMware Knowledge Base article 2144605* for more information.
vSphere	VMware recommends a minimum of 5.5U3 and 6.0U3 in NSX environments. See also the VMware Product Interoperability Matrix for interoperability information. **Notes:** • NSX 6.2.x is not compatible with vSphere 6.5. • vSphere 6.0U3 with NSX for vSphere 6.2.4 resolves the issue of duplicate VTEPs in ESXi hosts after rebooting vCenter server. See *VMware Knowledge Base article 2144605* for more information. • For Distributed Service Insertion, ESXi versions 5.5 Patch 10 and ESXi 6.0U3 or later are recommended. See *VMware Knowledge Base article 2149704* for more information.
Guest Introspection	Guest Introspection-based features in NSX are compatible with specific VMware Tools (VMTools) versions. To enable the optional Thin Agent Network Introspection Driver component packaged with VMware Tools, you must upgrade to one of: • 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* • VMware Tools 10.0.9 and later for Windows 10 support
vRealize Orchestrator	NSX-vRO plugin 1.0.3 or later
VMware Integrated Openstack (VIO)	VIO 2.5.1 or later
vCloud Director (vCD)	• vCD 8.0 or later if migrating from vCNS to NSX • vCD 8.10.1 or later if already on NSX

System Requirements and Installation

For the complete list of NSX installation prerequisites, see theSystem 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*. Upcoming end-of-support dates include:

- vCloud Networking and Security has reached End of Availability (EOA) and End of General Support (EOGS) on September 19, 2016. (See also *VMware knowledge base article 2144733*.) (See also *VMware knowledge base article 2144620*.)
- NSX for vSphere 6.1.x has reached End of Availability (EOA) and End of General Support (EOGS) on January 15, 2017. (See also *VMware knowledge base article 2144769*.)
- As of NSX 6.2.3, the NSX Data Security feature has been deprecated. In NSX 6.2.3, you can continue to use this feature at your discretion, but be aware that this feature will be removed from NSX in a future release.
- Web Access Terminal (WAT) is being deprecated and will not be included in a future maintenance release. VMware recommends using the full access client with SSL VPN deployments for improved security.

Unsupported controller commands are no longer shown

Please review the CLI guide for the complete list of supported commands. You should only use commands which are documented in this guide. The join control-cluster command is not a supported command on NSX for vSphere. See also *VMware knowledge base article 2135280.*

TLS 1.0 support has been deprecated as of NSX 6.2.3

In the NSX VPN, IPsec, and Load Balancer cipher suites, TLS 1.0 support has been deprecated as of NSX 6.2.3. For information on changes in cipher support, see *VMware knowledge base article 2147293*.

Upgrade Notes

- **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 upgrade to NSX 6.2.4 or later**, you must perform a full NSX upgrade including host cluster upgrade (which upgrades the host VIBs to 6.2.4). For instructions, see the *NSX Upgrade Guide*

  including the Upgrade Host Clusters to NSX 6.2 section.
- **Minimum of 8GB memory requirement** Upgrade to NSX 6.2.3 or later may fail on hosts with less than 8 GB of memory.
- **Upgrading Edge Services Gateway (ESG):**
  Starting in 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:

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

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

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

- **Controller disk layout**: New installations of NSX 6.2.3 or later will deploy NSX Controller appliances with updated disk partitions to provide extra cluster resiliency. In previous releases, log overflow on the controller disk might impact controller stability. In addition to adding log management enhancements to prevent overflows, the NSX Controller appliance has separate disk partitions for data and logs to safeguard against these events. If you upgrade to NSX 6.2.3 or later, the NSX Controller appliances will retain their original disk layout.
- **Upgrade paths:**

- **Upgrade paths:**
  - 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.
  - Upgrade path from vCNS 5.5.x:
    - Using the NSX upgrade bundle you may upgrade directly from VMware vCloud Networking and Security (vCNS) 5.5.x to NSX 6.2.x. For instructions, see the *NSX Upgrade Guide*, in the section, vCloud Networking and Security to NSX Upgrade. This section also includes instructions for upgrading vCNS 5.5.x to NSX in a vCloud Director environment. A separate guide, the NSX Upgrade Guide for vShield Endpoint, includes instructions for upgrading vCNS 5.5.x to NSX 6.2.x if you are using vShield Endpoint for anti-virus protection only.
    - If you have virtual wires in your environment, you must update your host clusters. Once this is done, the virtual wires are renamed to logical switches. See Update Host Clusters for instructions.
- **To validate** that your upgrade to NSX 6.2.x was successful seeknowledge base article 2134525.
- **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.
- **Known issues affecting upgrades**: See the section, Installation and Upgrade Known Issues, later in this document, for a list of known upgrade-related issues.
- **New system requirements**: For the memory and CPU requirements while installing and upgrading NSX Manager, see the System Requirements for NSX section in the NSX 6.2 documentation.
- **Set Correct Cipher version for Load Balanced Clients that use TLS 1.0:**This affects vROPs pool members that use TLS version 1.0. If the servers, whose traffic is being load-balanced, use this

  version, you must set a monitor extension value explicitly with the "ssl-version=10" setting in the NSX Load Balancer. See the NSX Administration Guide.

  ```
  {
     "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"
  }
  ```

- **Maximum number of NAT rules**: For NSX Edge versions prior to 6.2, a user could configure 2048 SNAT and 2048 DNAT rules separately, giving a total limit of 4096 rules. Since NSX Edge version 6.2 onwards, a limit is enforced for the maximum allowed NAT rules, based on the NSX Edge appliance size:

  1024 SNAT and 1024 DNAT rules for a total limit of 2048 rules for COMPACT edge.

  2048 SNAT and 2048 DNAT for a total limit of 4096 rules for LARGE edge and QUADLARGE edge.

  4096 SNAT and 4096 DNAT rules for a total limit of 8192 rules for XLARGE edge.

  During an NSX Edge upgrade to version 6.2, any existing COMPACT edge whose total NAT rules (sum of SNAT and DNAT) exceeds the limit 2048 will fail validation, resulting in an upgrade failure. In this scenario, the user will need to change the appliance size to LARGE, QUADLARGE and retry the this scenario, the user will need to change the appliance size to LARGE, QUADLARGE and retry the upgrade.

- **Behavior change in redistribution filters** on distributed logical router and Edge Services Gateway: Starting in the 6.2 release, redistribution rules in the DLR and ESG work as ACLs only. That is, if a rule is an exact match, the respective action is taken.
- **VXLAN tunnel ID**: Before upgrading to NSX 6.2.x, you must make sure your installation is not using a VXLAN tunnel ID of 4094 on any tunnels. VXLAN tunnel ID 4094 is no longer available for use. To assess and address this follow these steps:
  1. In vCenter, navigate to **Home** > **Networking and Security** > **Installation** and select the **Host Preparation** tab.
  2. Click **Configure** in the VXLAN column.
  3. In the Configure VXLAN Networking window, set the VLAN ID to a value between 1 and 4093.
- **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.
- **Autosaved drafts and Service Composer**: In NSX 6.2.3 and later, you can disable the auto draft feature by setting *autoDraftDisabled* to true. Manually configured settings are maintained during the upgrade. Disabling the auto drafts feature before making large numbers of changes to the firewall rules might improve performance, and will prevent previously saved drafts from being overwritten. You can use the following API call to set the property *autoDraftDisabled* to true in the global configuration:

  1. Get the existing global firewall configuration (GlobalConfiguration):
     GET https://NSX-Manager-IP-Address/api/4.0/firewall/config/globalconfiguration
     Note that a GET will not show the autoDraftDisabled field.
  2. Use a PUT call to set the property autoDraftDisabled to true in the global configuration:
     PUT https://NSX-Manager-IP-Address/api/4.0/firewall/config/globalconfiguration
     with a request body that includes:

     ```
     <globalConfiguration>
        <layer3RuleOptimize>...</layer3RuleOptimize>
        <layer2RuleOptimize>...</layer2RuleOptimize>
        <tcpStrictOption>...</tcpStrictOption>
        <autoDraftDisabled>true</autoDraftDisabled>
     </globalConfiguration>
     ```
- **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*: 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.

Document Revision History

06 Jul 2017: First edition.
21 Aug 2017: Second edition. Added 1904842, 1878081, 1910593.
02 Oct 2017: Third edition. Updated minimum recommend version for NSX.

Resolved Issues

The resolved issues are grouped as follows.

- Installation and Upgrade Resolved Issues in NSX 6.2.8
- Networking and Edge Services Resolved Issues in NSX 6.2.8
- NSX Manager Resolved Issues in NSX 6.2.8
- Security Services Resolved Issues in NSX 6.2.8
- Solution Interoperability Resolved Issues in NSX 6.2.8

Installation and Upgrade Resolved Issues in NSX 6.2.8

- **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 6.2.8.*

Networking and Edge Services Resolved Issues in NSX 6.2.8

- **Fixed Issue 1849037: NSX Manager API threads get exhausted when communication link with NSX Edge is broken**
  If communication channel between NSX Manager and NSX Edge VM is broken, then status/statistics requests to Edge VM get hung and block the API thread. Multiple such requests may end up exhausting all API threads. *Fixed in 6.2.8.*

- **Fixed Issue 1865394: Traffic Shaping Policy never deleted from the dvPortGroup port**
  Traffic Shaping Policy (TSP) is set on the dvPortGroup port connected to the NSX Edge vNIC when you configure the port. Expected behavior is that when a vNIC is disconnected from this port, TSP is deleted (cleared). The following cases triggered by the NSX Manager lead to the change/disconnect of dvPortGroup port connected to the NSX Edge vNIC:
  - NSX Edge is Redeployed
  - NSX Edge is Upgraded
  - NSX Edge is disconnected and reconnected
  - NSX Edge applicance size is changed
  - NSX Edge interface is deleted
  - NSX Edge is deleted

  Fixed in 6.2.8.

- **Fixed Issue 1704940: You may encounter the purple diagnostic screen on the ESXi host if the pCPU count exceeds 256**
  NSX DLR has one flow table per pCPU and the number of pCPUs is restricted to 256. This causes a crash if the pCPU count exceeds 256. *Fixed in 6.2.8.*

- **Fixed Issue 1892265: NSX Edge file bundle does not get deleted from `/common/tmp` directory after**

every publish and fills up the /common directory
/common directory gets full and NSX Manager runs out of space because NSX Edge file bundle (sslvpn-plus config) does not get deleted from /common/tmp. *Fixed in 6.2.8.*

- **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 6.2.8.*

- **Fixed Issue 1849760: Routing process may hog CPU when certain prefixes are learned over IBGP network such that the nexthop belongs to the prefix subnet**
  The problem is seen in a setup where DLR 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. *Fixed in 6.2.8.*

NSX Manager Resolved Issues in NSX 6.2.8

- **Fixed Issue 1892208: NSX-Manager database is showing only one datastore (vmx file location) but UI is showing two (Current and Configured).**
  Fixed in 6.2.8.

- **Fixed Issue 1861785: 100% CPU usage observed on NSX manager with vCloud Director (vCD) and vRealize Operations (vROps)**
  If your database has a significant number of orphan records in system_event_message_params and system_event_event_metadata tables, APIs take longer to respond to requests from vCD and in the meantime vCD client gets disconnected resulting in calling the same APIs again. This leads to high CPU usage. Removing orphan records resolves this problem. *Fixed in 6.2.8.*

- **Fixed Issue 1760940: NSX Manager High CPU triggered by many simultaneous vMotion tasks**
  DFW is configured with dynamic security groups which have criteria based on VM name or VM guest OS. See *VMware Knowledge Base article 2150668* for more information. *Fixed in 6.2.8.*

Security Services Resolved Issues in NSX 6.2.8

- **Fixed Issue 1836322: "Flash Error" when editing NSX Firewall Rules**
  Sometimes, the vSphere web client UI displays an error related to the Flash plug-in when editing security groups. *Fixed in 6.2.8.*

- **Fixed Issue 1832912: When a host is put in maintenance mode some virtual machines from that host lose firewall rules applied previously**
  When a host is put in maintenance mode and DRS moves all VMs to another host, some virtual machines from that host lose firewall rules applied previously. The workaround is to avoid creating Security Groups using dynamic criteria based on Computer Name or Computer OS name. *Fixed in 6.2.8.*

- **Fixed Issue 1818550: Grouping objects updates are delayed on hosts that have a large number of VMs and nested security groups**
  On hosts that have a large number of VMs and nested security groups, a single inventory change will cause a grouping objects update flush, causing long delays in grouping object updates to the hosts. *Fixed in 6.2.8.*

- **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 6.2.8.*

- **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 6.2.8.*

Solution Interoperability Resolved Issues in NSX 6.2.8

- **Fixed Issue 1838742: NSX Edge version may differ between displayed UI / API version and deployed version in VCD environment**
  During a vCNS to NSX upgrade, when the NSX Manager is upgraded, all vCNS Edges must be upgraded or redeployed explicitly before performing any other edit operation on vCloud Director. If the upgrade is not initiated, the NSX Edge version displayed in the UI and API may differ from the deployed version. *Fixed in 6.2.8.*

Known Issues

The 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 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.

  Workaround: Refer to VMware knowledge base article 2146769.

- **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.

  Workaround: Restart NSX Manager or run the following API to get rid of the extra entry. Perform a PUT of the API method,
  `https://nsx-manager-address/api/internal/firewall/updatestatus`

- **Issue 1659043: 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.

- **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.
  Workaround: To return Guest Introspection to a stable state, Guest Introspection VMs must be configured to ignore lookups for these well-known SIDs. This is achieved by updating a configuration file on each Guest Introspection VM and then restarting the service. In addition, Active Directory log scraping can be used as a workaround for detecting logon events for ID Firewall.

  Steps to ignore SID lookups for unresolvable SIDs:
  1. Login to Guest Introspection VM.
  2. Edit the file at /usr/local/usvmmgmt/config/ignore-sids.lst.
  3. Append the following 2 lines:
     S-1-18-1
     S-1-18-2
  4. Save and close the file.
  5. Restart the Guest Introspection service with command:
     rcusvm restart.

- **Issue 1558285: Deleting cluster with Guest Introspection from Virtual Center results in null pointer exception**
  Services such as Guest Introspection must be removed first before a cluster is removed from VC
  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.

- **Issue 1477280: Cannot create hardware gateway instances when no controller is deployed**
  Controllers must be deployed before any hardware gateway instances are configured. If controllers are not deployed first, the error message "Failed to do the Operation on the Controller" is shown.

  Workaround: None.

- **Issue 1491275: NSX API returns JSON instead of XML in certain circumstances**

  On occasion, an API request will result in JSON, not XML, being returned to the user.

  Workaround: Add `Accept: application/xml` to the request header.

Installation and Upgrade Known Issues

- **New Issue 1905064: Some hosts in the cluster lose TCP connections to NSX Manager and controllers during a host upgrade**
  If an ESXi host loses TCP connection to the NSX manager during an upgrade, it cannot get controller information. This prevents the controller from pushing NSX related configs to this host causing all the VMs in this host to lose connectivity.

  Workaround: Reboot the affected host.

- **New Issue 1910593: Answer parameters are case sensitive in the NSX Manager upgrade API**
  If you use the NSX API to upgrade NSX Manager, and you want to enable SSH or join the VMware CEIP program, you must specify "Yes" (not "YES") for the answer parameter.

  Workaround: See the NSX API Guide for more information on upgrading using the API.

- **Issue 1838229: 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 scriptoption httpclose and associate it with the virtual server.

- **Issue 1820723: Unable to see filters on ESXi after upgrade from 6.2.x to 6.2.7 because of failure to connect to host**
  After you upgrade from NSX 6.2.x to 6.2.7 and cluster VIBs to 6.2.7 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 the failure of Firewall rules publish and Security Policy publish, and VXLAN configuration not being sent down to hosts.

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

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

- **Issue 1435504: HTTP/HTTPS Health Check appears DOWN after upgrading from 6.0.x or 6.1.x to 6.2.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 6.2.x, the additional double quotes result in the pool members shown as DOWN in Health Check.

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

- **Issue 1768144: Old NSX Edge appliance resource reservations that exceed new limits may cause failure during upgrade or redeployment**
  In NSX 6.2.4 and earlier, you could specify an arbitrarily large resource reservation for an NSX Edge appliance. NSX did not enforce a maximum value. After NSX Manager is upgraded to 6.2.5 or later, if an existing Edge has resources reserved (especially memory) that exceed the newly enforced maximum value imposed for the chosen form factor, it would fail during Edge upgrade or redeploy (which would trigger an upgrade). For example, if the user has specified a memory reservation of 1000MB on a pre-6.2.5 LARGE Edge and, after upgrade to 6.2.5 or later, 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 1730017: Upgrades from 6.2.3 to 6.2.7 do not show a version change for Guest Introspection**
  As the 6.2.3 Guest Introspection module is the latest version available, the version after a 6.2.7 upgrade remains unchanged. Note that upgrades from earlier NSX releases may show a version change to 6.2.7

  Workaround: This does not affect any functionality.

- **Issue 1683879: Upgrade to NSX 6.2.3 or later may fail on hosts with less than 8 GB of memory**
  NSX 6.2.3 and later requires a minimum of 8 GB of memory on prepared hosts running networking and security services. The minimum ESXi 6.0 memory requirement of 4 GB is not sufficient to run NSX.

  Workaround: None.

- **Issue 1673626: Modifying tcpLoose via /api/3.0/edges is not allowed after upgrading from vCloud Networking and Security to NSX**
  After upgrading from vCloud Networking and Security to NSX, you will see an error if you try to modify the tcpLoose setting in this API request: /api/3.0/edges

  Workaround: Use tcpPickOngoingConnections setting in the globalConfig section in the API request /api/4.0/firewall/config instead.

- **Issue 1658720: Enabling DFW for a given cluster would fail for a vCNS to NSX upgrade scenario where the cluster has VXLAN installed and vShield App not installed (or removed before upgrade) in vCNS deployment**
  This issue occurs because the cluster sync status is not invoked when the hosts are upgraded.

  Workaround: Restart NSX Manager.

- **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 Virtual Center 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 1465249: Installation status for Guest Introspection shows Succeeded even though the host is offline**
  After installing Guest Introspection on the cluster that has one host offline, the host that is offline shows Installation Status as Succeeded and Status Unknown.
  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 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 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 filter gets created which will have support for TFTP ALG.

- **Issue 1394287: Adding or removing VMs from a virtual wire does not update IP address set in vShieldApp rules**
  If an existing vCNS vShield App firewall installation is not upgraded to the NSX distributed firewall in enhanced mode, new VMs with firewall rules based on virtual wires will not have an updated IP address. As a result, these VMs are not protected by the NSX firewall This issue is only seen in the following scenarios:
  - After upgrading the Manager from vCNS to NSX and not switching to DFW Enhanced mode.
  - If you add new VMs to a virtualWire with vshield App rules consuming those virtualwires, these rules will not have the new IP Address set for the new VMs.
    This will cause the new VMs not protected by vShieldApp.

  Workaround: Publish the rule again which will set the new address.

- **Issue 1386874: 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 admininstrator@vsphere.local user name instead of root.

- **Issue 1375794: Shutdown Guest OS for agent VMs (SVA) before powering OFF**
  When a host is put into maintenance mode, all service appliances are powered-off, instead of shutting down gracefully. This may lead to errors within third-party appliances.
  Workaround: None.

- **Issue 1112628: Unable to power on the Service appliance that was deployed using the Service Deployments view**
  Workaround: Before you proceed, verify the following:
  - The deployment of the virtual machine is complete.
  - No tasks such as cloning, reconfiguring, and so on are in progress for the virtual machine displayed in VC task pane.
  - In the VC 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 1088913: vSphere Distributed Switch MTU does not get updated**
  If you specify an MTU value lower than the MTU of the vSphere distributed switch when preparing a cluster, the vSphere Distributed Switch does not get updated to this value. This is to ensure that existing traffic with the higher frame size isn't unintentionally dropped.
  Workaround: Ensure that the MTU you specify when preparing the cluster is higher than or matches the current MTU of the vSphere distributed switch. The minimum required MTU for VXLAN is 1550.
- **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 1288506: After upgrading from vCloud Networking and Security 5.5.3 to NSX vSphere 6.0.5 or later, NSX Manager does not start up if you are using an SSL certificate with DSA-1024 keysize**
  SSL certificates with DSA-1024 keysize are not supported in NSX vSphere 6.0.5 onwards, so the upgrade is not successful.
  Workaround: Import a new SSL certificate with a supported keysize before starting the upgrade.

- **Issue 1263858: SSL VPN does not send upgrade notification to remote client**
  SSL VPN gateway does not send an upgrade notification to users. The administrator has to manually communicate that the SSL VPN gateway (server) is updated to remote users and they must update their clients.
  Workaround: Users need to uninstall the older version of client and install the latest version manually.

- **Issue 1402307: If vCenter is rebooted during NSX vSphere upgrade process, incorrect Cluster Status is displayed**
  When you do host prep in an environment with multiple NSX prepared clusters during an upgrade and the vCenter Server gets rebooted after at least one cluster has been prepared, the other clusters may show Cluster Status as Not Ready instead of showing an Update link. The hosts in vCenter may also show Reboot Required.
  Workaround: Do not reboot vCenter during host preparation.

- **Issue 1487752: Momentary loss of third-party anti-virus protection during upgrade**
  When upgrading from NSX 6.0.x to NSX 6.1.x or 6.2.x, you might experience momentary loss of third-party anti-virus protection for VMs. This issue does not affect upgrades from NSX 6.1.x to NSX 6.2.
  Workaround: None.

- **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.

  Workaround: None.

- **Issue 1284735: After upgrade from vCNS, cannot place new grouping objects into some upgraded grouping objects**
  vCNS 5.x supported creation of grouping objects at scopes below GlobalRoot (below the NSX-wide scope). For example, in vCNS 5.x you could create a grouping object as the DC or PG level. By contrast, the NSX 6.x user interface creates such objects under the GlobalRoot, and these newly created grouping objects cannot be added to existing grouping objects that were created at a lower scope (DC or PG) in your pre-upgrade vCNS installation.
  Workaround: See VMware knowledge base article 2117821.

- **Issue 1495969: After upgrading from vCNS 5.5.4 to NSX 6.2.x, the firewall on the Host Preparation tab remains disabled**
  After upgrading from vCNS 5.5.x to NSX 6.2.x and upgrading all the clusters, the firewall on the Host Preparation tab remains disabled. In addition, the option to upgrade the firewall does not appear in the UI. This happens only when there are hosts that are not part of any prepared clusters in the datacenter, because the VIBs will not be installed on those hosts.
  Workaround: To resolve the issue, move the hosts into an NSX 6.2 prepared cluster.

- **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.
  Workaround: During an NSX upgrade, do not publish changes to the distributed firewall configuration. To exit from the `InProgress` state and resolve the issue, reboot the NSX Manager virtual appliance.

- **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 1479314: NSX 6.0.7 SSL VPN clients cannot connect to NSX 6.2 SSL VPN gateways**
  In NSX 6.2 SSL VPN gateways, the SSLv2 and SSLv3 protocols are disabled. This means the SSL VPN gateway only accepts the TLS protocol. The SSL VPN 6.2 clients have been upgraded to use the TLS protocol by default during connection establishment. In NSX 6.0.7, the SSL VPN client uses an older version of OpenSSL library and the SSLv3 protocol to establish a connection. When an NSX 6.0.x client tries to connect to an NSX 6.2 gateway, the connection establishment fails at the SSL handshake level.
  Workaround: Upgrade your SSL VPN client to NSX 6.2 after you have upgraded to NSX 6.2. For upgrade instructions, see the NSX Upgrade documentation.

- **Issue 1434909: Must create a segment ID pool for new or upgraded logical routers**
  In NSX 6.2, a segment ID pool with available segment IDs must be present before you can upgrade a logical router to 6.2 or create a new 6.2 logical router. This is true even if you have no plans to use NSX logical switches in your deployment.
  Workaround: If your NSX deployment does not have a local segment ID pool, create one as a prerequisite to NSX logical router upgrade or installation.

- **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 1463767: In a cross vCenter deployment, a universal firewall configuration section might be under (subordinate to) a local configuration section**

If you move a secondary NSX Manager to the standalone (transit) state and then change it back to the secondary state, any local configuration changes that you made while it was temporarily in the standalone state might be listed above the replicated universal configuration sections inherited from the primary NSX Manager. This produces the error condition universal section has to be on top of all other sections on secondary NSX Managers.

Workaround: Use the UI option to move sections up or down so that the local section is below the universal section.

- **Issue 1289348: After an upgrade, firewall rules and network introspection services might be out of sync with NSX Manager**
  After upgrading from NSX 6.0 to NSX 6.1 or 6.2, the NSX firewall configuration displays the error message: synchronization failed / out of sync . Using the **Force Sync Services**: **Firewall** action does not resolve the issue.

  Workaround: In NSX 6.1 and NSX 6.2, either security groups or dvPortgroups can be bound to a service profile, but not both. To resolve the issue, modify your service profiles.

- **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.

- **Issue 1393889: Data Security service status is shown as UP even when IP connectivity is not established**
  Data Security appliance may have not received the IP address from DHCP or is connected to an incorrect port group.
  Workaround: Ensure that the Data Security appliance gets the IP from DHCP/IP Pool and is reachable from the management network. Check if the ping to the Data Security appliance is successful from NSX/ESX.

- **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 1904842: NSX Manager is not registering with vCenter or Platform Service Controller**
  NSX Manager is not appearing on UI and any REST call to the NSX Manager fails.

  Workaround: Restart the NSX management service or reboot NSX manager appliance.

- **New Issue 1826225: Service status for partner service VMs reported as "Unknown" in NSX Manager**
  Service status for partner service VMs is reported as "Unknown" in NSX Manager. This issue occurs when there are stale database entries for partner VMs.

  Workaround: Contact VMware customer support.

- **New Issue 1713669: NSX Manager disk becomes full with IDFW data**
  Whether IDFW rules are used or not, Guest Introspection and Event Log Servers detected login events are stored in the NSX Manager database and retained in the database for 30 days after they expire. In environments with a large volume of login activity, the database may grow and impact space on the NSX Manager disk.

  Workaround: There is no workaround. Contact VMware Support if you encounter this issue.

- **Issue 1806368: In case of a cross-VC Failover, the DLR config is not pushed to all hosts if old controllers are reused for a previously failed primary NSX Manager that is made primary again after a failover.**
  In a cross-VC setup, when the primary NSX Manager fails, a secondary NSX Manager is promoted to primary and a new controller cluster is deployed to be used with the newly promoted secondary (now primary) NSX Manager. When the primary NSX Manager is back on, the secondary NSX Manager is demoted and the primary NSX Manager is restored. In this case, if you reuse the existing controllers that were deployed on this primary NSX Manager before the failover, the DLR config is not pushed to all hosts. This issue does not arise if you create a new controller cluster instead.

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

- **Issue 1831131: Connection from NSX Manager to SSO fails when authenticated using the LocalOS user**
  Connection from NSX Manager to SSO fails when authenticated using the LocalOS user with the error: "Could not establish communication with NSX Manager. Please contact administrator."

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

- **Issue 1772911: NSX Manager disk space consumption increasing rapidly, task and job tables size increasing with high CPU usage**
  NSX Manager CPU consistently is at 100% or is regularly spiking to 100%. Running the show process monitor command in NSX Manager Command Line Interface (CLI) displays the Java process that is consuming the highest CPU cycles. Disk space consumption growing rapidly with increase in DB size, resulting in slow performance of NSX Manager.

  Workaround: Contact VMware technical support.

- **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 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.

  Workaround: Remove ESXTOP plugin with the following procedure:

  1. Make sure there is a recent backup of vCenter Snapshot vCenter VM (without quiesce)
  2. Delete /usr/lib/vmware-vsphere-client/plugin-packages/esxtop-plugin
     rm -R /usr/lib/vmware-vsphere-client/plugin-packages/esxtop-plugin
  3. Delete /usr/lib/vmware-vsphere-client/server/work
     rm -R /usr/lib/vmware-vsphere-client/server/work
  4. Restart the web client
     service vsphere-client restart
  5. (Optional) Ensure that there is no output from the following command: "tail -f /var/log/vmware/vsphere-client/logs/eventlog.log | grep esx"
  6. Make sure to consolidate vCenter snapshot if that is the preferred method of roll back option

- **Issue 1486403: NSX Manager does not accept DNS search strings with a space delimiter**
  NSX Manager does not accept DNS search strings with a space delimiter. You may only use a comma as a delimiter. For example, if the DHCP server advertises eng.sample.com and sample.com for the DNS search list, NSX Manager is configured with eng.sample.com sample.com.

  Workaround: Use comma separators. NSX Manager only accepts comma separator as DNS search strings.

- **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 1534606: Host Preparation Page fails to load**
  When running Virtual Center in linked mode, each VC 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.

Network page.

- **Issue 1460766: NSX Manager UI do not automatically log out after changing password using NSX Command Line Interface**
  If you are logged in to NSX Manager and recently changed your password using CLI, you might continue to stay logged in to the NSX Manager UI using your old password. Typically, NSX Manager client should automatically log you out if the session times out for being inactive.
  Workaround: Log out from the NSX Manager UI and log back in with your new password.

- **Issue 1467382: Unable to edit a network host name**
  After you login to NSX Manager virtual appliance and navigate to the Appliance Management, click Manage Appliance Settings, and click Network under Settings to edit the network host name, you might receive an invalid domain name list error. This happens when the domain names specified in the Search Domains field are separated with whitespace characters, instead of commas. NSX Manager only accepts domain names that are comma separated.
  Workaround: 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.

- **Issue 1798847: In a Cross VC 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 1698286: Hardware VTEP in a cross-vCenter NSX environment only supported on the primary NSX Manager**
  In a cross-vCenter NSX environment, hardware gateway switch configurations and operations are supported only on the primary NSX Manager. Hardware gateway switches must be bound to non-universal logical switches. Hardware gateway configurations are not supported on secondary NSX Managers.

  Workaround: In a cross-vCenter NSX environment it is recommended to use L2 bridging to connect logical switches to physical networks.

- **Issue 1844966: NSX Edge file system becomes Read-only**
  If there are any storage connectivity issues where an edge is located, edge file system may enter read-only state to protect the OS file system. This is expected behaviour in a Linux server. See *VMware Knowledge Base article 2146870* for more information.

  Workaround: Do the following:

  1. Re-deploy the edge.
  2. Reboot the edge(Both edges in HA pair).
  3. Force sync the edge.

- **Issue 1799261: NSX Edge may run into split-brain condition after upgrade or redeploy**
  On the standby NSX Edge, the *show service highavailability* CLI command show high availability status as "Standby" but the config engine status as "Active".

  Workaround: Reboot the standby NSX Edge.

- **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 and 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 or FQDN in IPSec VPN configuration instead of ANY.

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

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

- **Issue 1758500: Static route with multiple next-hops does not get installed in NSX Edge routing and forwarding tables if at least one of the next-hop configured is the Edge's vNIC IP address**
  With ECMP and multiple next-hop addresses, NSX allows the Edge's vNIC's IP address to be configured as next-hop if at least one of the next-hop IP addresses is valid. This is accepted without any errors or warnings but route for the network is removed from the Edge's routing/forwarding table.

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

- **Issue 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.

  Workaround: Enable/Disable routing protocol or clear routing neighborship.

- **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.

  Workaround: To work around this issue, disable the route redistribution of static into OSPF.
  Note: This issue does not impact the data path. SeeVMware knowledge base article 2147998.

- **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 1776073: 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 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.
  To configure the same CPU/Memory reservation on both Edge VMs:

  - Use the PUT API https:// <NSXManager>/api/4.0/edgePublish/tuningConfiguration to set explicit values for both Edge VMs.
    or
  - Disable and re-enable Edge HA, which will delete the second Edge VM and redeploy a new one with the default reservations.

  Workaround: None.

- **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 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.

  Workaround: To recover from this issue, refer to VMware knowledge base article 2145959.

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

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.
Workaround: See VMware knowledge base article 2127369.

- **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.
  Workaround: No workaround.

- **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.
  Workaround: On the logical router appliance that is missing the connected route for the HA interface, either disconnect and then reconnect the vNIC of the logical router appliance, or reboot the logical router appliance.

- **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 1624663: After clicking "Configure Advanced Debugging" refreshes the VC UI and the change does not persist**
  After clicking the specific edge ID > Configuration > Action > Configure Advanced Debugging causes the VC UI to refresh and the change does not persist

  Workaround: Go directly to the Edge list menu, highlight the edge, and click Action > Configure Advanced Debugging to continue with the changes.

- **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 *https://NSX-Manager-IP-Address/api/4.0/edges/edgeId/status* API with *detailedStatus=true*.

- **Issue 1647657: Show commands on an ESXi host with VDR display no more than 2000 routes per VDR instance**

  Show commands on an ESXi host with VDR enabled will not show more than 2000 routes per VDR 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: No workaround.

- **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 1663902: Renaming an NSX Edge VM disrupts traffic flowing through the Edge**
- **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 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 1483426: IPsec and L2 VPN service status shows as down even when the service is not enabled**
  Under the Settings tab in the UI, the L2 service status is displayed as down, however the API shows the L2 status as up. L2 VPN and IPsec service always shows as down in the Settings tab unless the UI page is refreshed.
  Workaround: Refresh the page.

- **Issue 1446327: Some TCP-based applications may time out when connecting through NSX Edge**
  The default TCP established connection inactivity timeout is 3600 seconds. The NSX Edge deletes any connections idle for more than the inactivity timeout and drops those connections.

  Workaround:
  1. If the application has a relatively long inactivity time, enable TCP keepalives on the hosts with keep_alive_interval set to less than 3600 seconds.
  2. Increase the Edge TCP inactivity timeout to greater than 2 hours using the following NSX REST API. For example, to increase the inactivity timeout to 9000 seconds. NSX API URL: /api/4.0/edges/{edgeId}/systemcontrol/config PUT Method <systemControl> <property>sysctl.net.netfilter.nf_conntrack_tcp_timeout_established=9000</property> </systemControl>

- **Issue 1534602: UI does not display Edge management plane mode (VIX/MSGBUS), and does not provide the option to change from VIX to MSGBUS**
  When an Edge appliance is in VIX mode, it is not eligible to be selected for inclusion in DFW, and centralized CLI commands take much longer to run compared to MSGBUS mode
  Workaround: Make sure that the cluster where the Edge is deployed is prepared for NSX and its "NSX Manager to Firewall Agent" is in "Up" state, and redeploy the Edge.

- **Issue 1498243: Distributed logical router advertises incorrect next hop for default route when BGP neighbor filter is set to "DENY, ANY, OUT"**
  With 'default originate' enabled on an NSX distributed logical router (DLR), setting a BGP neighbor filter of "DENY, ANY, OUT" on the DLR causes the DLR to advertise an incorrect next hop address for the default route. This error occurs only when a BGP neighbor filter is added with the following attributes:
    - Action: DENY
    - Network: ANY
    - Direction: OUT

  Workaround: None.

- **Issue 1471561: BGP/OSPF neighbor relationship is not established with directly connected routers**

Dynamic routing does not work as expected with directly connected routers when ECMP routes exist for that directly connected network.
Workaround: Reboot Edge OR delete and re-create the associated vNIC interface.

- **Issue 1089238: Logical router LIF routes are advertised by upstream Edge Services Gateway even if logical router OSPF is disabled**
  Upstream Edge Services Gateway will continue to advertise OSPF external LSAs learned from logical router connected interfaces even when logical router OSPF is disabled.
  Workaround: Disable redistribution of connected routes into OSPF manually and publish before disabling OSPF protocol. This ensures that routes are properly withdrawn.
- **Issue 1499978: Edge syslog messages do not reach remote syslog server**
  Immediately after deployment, the Edge syslog server cannot resolve the hostnames for any configured remote syslog servers.
  Workaround: Configure remote syslog servers using their IP address, or use the UI to Force Sync the Edge.

- **Issue 1489829: Logical router DNS Client configuration settings are not fully applied after updating REST Edge API**
  Workaround: When you use REST API to configure DNS forwarder (resolver), perform the following steps:

  1. Specify the DNS Client XML server's settings so that they match the DNS forwarder setting.

  2. Enable DNS forwarder, and make sure that the forwarder settings are same as the DNS Client server's settings specified in the XML configuration.

- **Issue 1243112: Validation and error message not present for invalid next hop in static route, ECMP enabled**
  When trying to add a static route, with ECMP enabled, if the routing table does not contain a default route and there is an unreachable next hop in the static route configuration, no error message is displayed and the static route is not installed.
  Workaround: None.

- **Issue 1281425: If an NSX Edge virtual machine with one sub interface backed by a logical switch is deleted through the vCenter Web Client user interface, data path may not work for a new virtual machine that connects to the same port**
  When the Edge virtual machine is deleted through the vCenter Web Client user interface (and not from NSX Manager), the VXLAN trunk configured on dvPort over opaque channel does not get reset. This is because trunk configuration is managed by NSX Manager.
  Workaround: Manually delete the VXLAN trunk configuration by following the steps below:

  1. Navigate to the vCenter Managed Object Browser by typing the following in a browser window:
     https://*<vc-ip>*/mob?vmodl=1
  2. Click **Content**.
  3. Retrieve the dvsUuid value by following the steps below.
     a. Click the rootFolder link (for example, group-d1(Datacenters)).
     b. Click the data center name link (for example, datacenter-1).
     c. Click the networkFolder link (for example, group-n6).
     d. Click the DVS name link (for example, dvs-1)
     e. Copy the value of uuid.
  4. Click **DVSManager** and then click **updateOpaqueDataEx**.
  5. In *selectionSet*, add the following XML.

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

```
</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

- **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.

  Workaround: Change the Action on the DFW Firewall rule from "Reject" to "Block".

- **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.

- **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**

- **Issue 1686036: 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 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.

  Workaround: Do not use the "+" button at the global rule level to add rules to the Service Composer rules section.

- **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.

  Workaround:

  - Login to each ESXi host and restart the DFW controlplane process by running the following command:
    /etc/init.d/vShield_Stateful_Firewall restart
  - Verify the status using the following command:
    /etc/init.d/vShield_Stateful_Firewall status
  - The result similar to following is displayed:
    "vShield-Stateful-Firewall is running"

  Note: You should be cautious while performing this operation as this will push all DFW rules to all the filters again. If there are lot of rules, it might take some time to enforce them on all the filters.

- **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:

  ```
  <?xml version="1.0" encoding="UTF-8"? >
  neutron-server.log.1:70282:2016-08-23 17:58:23.429 30787 ERROR
  vmware_nsx.plugins.nsx_v.plugin
  <error>
  <details> 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 </details>
  <errorCode>258
  </errorCode>
  </error>
  ```

  Workaround: Avoid concurrent modification of the firewall configuration.

- **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.
  Workaround: Do not use the DFW status API query until the new host is prepared successfully.

- **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.
  Workaround: Do not delete the default rule. If the configuration with default rule was saved in draft, perform the following steps:

  1. Delete the complete firewall configuration using following DELETE API call.
     https://<NSX Manager IP>/api/4.0/firewall/globalroot-0/config
     This will restore the default section on the firewall.
  2. Load the saved draft of firewall rules with default section to the firewall.

- **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 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.

- **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:

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 1491046: 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**

  Workaround: See VMware knowledge base article 2144649.

- **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 1534585: Publishing Distributed Firewall (DFW) rules fails after referenced object is deleted in VMware NSX for vSphere 6.1.x and 6.2.x**
  Workaround: If this occurs, see knowledge base article 2126275.

- **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 1655593: Missing status on NSX Dashboard when logging in as Auditor or Security Admin roles**
  When viewing NSX Dashboard as Auditor or Security Admin, an error message "User is not authorized to access object ... and feature ... Please check object access scope and feature permissions for the user" appears. For example, Auditor may not be able to see "Logical Switch Status" from the Dashboard.

  Workaround: None.

- **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.

Solution Interoperability Known Issues

- **Issue 1840744: VMware ESXi 6.0.0 host experiences a purple diagnostic screen after a virtual machine is stuck in a reboot loop**
  This issue occurs due to race condition in dvfilter create/destroy events created by the virtual machine stuck in a reboot loop. - For more information, see VMware Knowledge Base article 2149782.

  Workaround: This issue is resolved in VMware ESXi 6.0 Patch 03 and later releases, available at VMware Downloads.
  To work around this issue if you do not want to upgrade, power off the affected virtual machine.

- **Issue 1568861: The NSX Edge deployment fails during any edge deployment from a vCD cell that does not own the VC listener**

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

  Workaround: Deploy an NSX Edge from the vCD cell which owns the VC listener.

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

  When a 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.

  Workaround: None.

NSX Controller Known Issues

- **Issue 1845087: NSX Controller API will be adversely affected if disk latency is very high**
  The NSX controller API may not respond within NSX manager's time limit if I/O latency is too high for the storage used by NSX controller. This may in turn affect upgrade and other functionalities of the NSX controller. The Network and Security plugin of the vSphere Web Client will show "High controller disk latency" as an error if the severity exceeds a certain limit.

  Workaround: To resolve this issue, VMware recommends using a dedicated local hard drive and SSD.

- **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 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.x environment may get into read-only mode in case of a storage outage. 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.

  Workaround: Restart the NSX Management Service.

- **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.