

VMware NSX for vSphere 6.2.1 Release Notes

Document updated 20 May 2016

VMware NSX for vSphere 6.2.1 | Released 17 Dec 2015 | Build 3300239 |

What's in the Release Notes

The release notes cover the following topics:

- [What's New](#)
- [Recommended Versions, System Requirements and Installation](#)
- [Upgrade Notes](#)
- [Known Issues](#)
- [Resolved Issues](#)
- [Document Revision History](#)

What's New

See what's new and changed in NSX [6.2.1](#) and [6.2.0](#).

New in 6.2.1

The 6.2.1 release delivers a number of bug fixes that have been documented in the [Resolved Issues](#) section.

- **6.1.5 fixes:** Release includes the same critical fixes as NSX-vSphere 6.1.5 content.
- **Introduced new 'show control-cluster network ipsec status' command** that allows users to inspect the Internet Protocol Security (IPsec) state.
- **Connectivity status:** NSX Manager user interface now shows the connectivity status of the NSX Controller cluster.
- **Support for vRealize Orchestrator Plug-in for NSX 1.0.3:** With NSX 6.2.1 release, NSX-vRO plugin version 1.0.3 is introduced for use with vRealize Automation 7.0.0. This plugin includes fixes that improve performance when vRealize Automation 7.0 uses NSX for vSphere 6.2.1 as a networking and security end point.
- **Starting in 6.2.1, NSX Manager queries each Controller node in the cluster to get the connection information between that controller and the other controllers in the cluster.**
This is provided in the output of the NSX REST API ("GET https://[NSX-MANAGER-IP-ADDRESS]/api/2.0/vdn/controller" command), which now shows the peer connection status among the controller nodes. If NSX Manager finds the connection between any two controller nodes is broken, a system event is generated to alert the user.
- **Service Composer now exposes an API that enables users to configure auto creation of Firewall drafts for Service Composer workflows.**
This setting can be turned on/off using REST API and the changes can be saved across reboot. When disabled, no draft is created in the Distributed Firewall (DFW) for policy workflows. This limits the number of drafts that are auto-created in the system and provides better performance.

New in 6.2.0

NSX vSphere 6.2.0 included the following new and changed features:

- **Cross vCenter Networking and Security**
 - **NSX 6.2 with vSphere 6.0 supports Cross vCenter NSX** where logical switches (LS), distributed logical routers (DLR) and distributed firewalls (DFW) can be deployed across multiple vCenters, thereby enabling logical networking and security for applications with workloads (VMs) that span multiple vCenters or multiple physical locations.

- **Consistent firewall policy across multiple vCenters:** Firewall Rule Sections in NSX can now be marked as "Universal" whereby the rules defined in these sections get replicated across multiple NSX managers. This simplifies the workflows involving defining consistent firewall policy spanning multiple NSX installations
- **Cross vCenter vMotion with DFW:** Virtual Machines that have policies defined in the "Universal" sections can be moved across hosts that belong to different vCenters with consistent security policy enforcement.
- **Universal Security Groups:** Security Groups in NSX 6.2 that are based on IP Address, IP Set, MAC Address and MAC Set can now be used in Universal rules, whereby the groups and group memberships are synced up across multiple NSX managers. This improves the consistency in object group definitions across multiple NSX managers, and enables consistent policy enforcement
- **Universal Logical Switch (ULS):** This new functionality introduced in NSX 6.2 as a part of Cross vCenter NSX allows creation of logical switches that can span multiple vCenters, allowing the network administrator to create a contiguous L2 domain for an application or tenant.
- **Universal Distributed Logical Router (UDLR):** This new functionality introduced in NSX 6.2 as a part of Cross vCenter NSX allows creation of distributed logical routers that can span multiple vCenters. The universal distributed logical routers enable routing across the universal logical switches described earlier. In addition, NSX UDLR is capable of localized north-south routing based on the physical location of the workloads.

• Operations and Troubleshooting Enhancements

- **New traceflow troubleshooting tool:** Traceflow is a troubleshooting tool that helps identify if the problem is in the virtual or physical network. It provides the ability to trace a packet from source to destination and helps observe how that packet passes through the various network functions in the virtual network.
- **Flow monitoring and IPFIX separation:** In NSX 6.1.x, NSX supported IPFIX reporting, but IPFIX reporting could be enabled only if flow reporting to NSX Manager was also enabled. Starting in NSX 6.2.0, these features are decoupled. In NSX 6.2.0 and later, you can enable IPFIX independent of flow monitoring on NSX Manager.
- **New CLI monitoring and troubleshooting commands in 6.2:** See [knowledge base article 2129062](#) for more information.
- **Central CLI:** Central CLI reduces troubleshooting time for distributed network functions. Commands are run from the command line on NSX Manager and retrieve information from controllers, hosts, and the NSX Manager. This allows you to quickly access and compare information from multiple sources. The central CLI provides information about logical switches, logical routers, distributed firewall and edges.
- **CLI ping command adds configurable packet size and do-not-fragment flag:** Starting in NSX 6.2.0, the NSX CLI 'ping' command offers options to specify the data packet size (not including the ICMP header) and to set the do-not-fragment flag. See the [NSX CLI Reference](#) for details.
- **Show health of the communication channels:** NSX 6.2.0 adds the ability to monitor communication channel health. The channel health status between NSX Manager and the firewall agent, between NSX Manager and the control plane agent, and between host and the NSX Controller can be seen from the NSX Manager UI. In addition, the host command channel offers greater fault tolerance.
- **Standalone Edge L2 VPN client CLI:** Prior to NSX 6.2, a standalone NSX Edge L2 VPN client could be configured only by 'deploy OVF' settings provided to the virtual center. Commands specific to standalone NSX Edge have been added to allow configuration using the command line interface.

• Logical Networking and Routing

- **L2 Bridging Interoperability with Distributed Logical Router:** With VMware NSX for vSphere 6.2, L2 bridging can now participate in distributed logical routing. The VXLAN network to which the bridge instance is connected, will be used to connect the routing instance and the bridge instance together.
- **Support of /31 prefixes on ESG and DLR interfaces per RFC 3021.**
- **Enhanced support of relayed DHCP request on the ESG DHCP server.**
- **Ability to preserve VLAN IDs/headers inside NSX virtual networks.**
- **Exact Match for redistribution filters:** The redistribution filter has same matching algorithm as ACL, so exact prefix match by default (except if le or ge options are used).
- **Support of administrative distance for static route.**

- **Ability to enable, relax, or disable check per interface on Edge.**
 - **Display AS path in CLI command `show ip bgp`**
 - **HA interface exclusion** from redistribution into routing protocols on the DLR control VM.
 - **Distributed logical router (DLR) force-sync avoids data loss for east-west routing traffic across the DLR.** North-south routing and bridging may continue experience an interruption.
 - **View active edge in HA pair:** In the NSX 6.2 web client, you can find out if an NSX Edge appliance is the active or backup in an HA pair.
 - **REST API supports reverse path filter (rp_filter) on Edge:** Using the system control REST API, `rp_filter` sysctl can be configured through UI, and is also exposed through REST API for vNIC interfaces and sub-interfaces. See the [NSX API documentation](#) for more information.
 - **Behavior of the IP prefix GE and IP prefix LE BGP route filters:** In NSX 6.2, the following enhancements have been made to BGP route filters:
 - **LE / GE keywords not allowed:** For the null route network address (defined as ANY or in CIDR format 0.0.0.0/0), less-than-or-equal-to (LE) and greater-than-or-equal-to (GE) keywords are no longer allowed. In previous releases, these keywords were allowed.
 - **LE and GE values in the range 0-7 are now treated as valid.** In previous releases, this range was not valid.
 - **For a given route prefix, you can no longer specify a GE value that is greater than the specified LE value.**
- **Networking and Edge Services**
 - **The management interface of the DLR has been renamed to HA interface.** This has been done to highlight the fact that the HA keepalives travel through this interface and that interruptions in traffic on this interface can result in a split-brain condition.
 - **Load balancer health monitoring improvements:** Delivers granular health monitoring that reports information on failures, keeps track of last health check and status change, and reports failure reasons.
 - **Support VIP and pool port range:** Enables load balancer support for applications that require a range of ports.
 - **Increased maximum number of virtual IP addresses (VIPs):** VIP support rises to 1024.
 - **Security Service Enhancements**
 - **New IP address discovery mechanisms for VMs:** Authoritative enforcement of security policies based on VM names or other vCenter-based attributes requires that NSX know the IP address of the VM. In NSX 6.1 and earlier, IP address discovery for each VM relied on the presence of VMware Tools (vmtools) on that VM or the manual authorization of the IP address for that VM. NSX 6.2 introduces the option to discover the VM's IP address by doing discovery from the hypervisor. These new discovery mechanisms enable NSX to enforce object based distributed firewall rules on VMs that do not have VMware Tools installed.
 - **Solution Interoperability**
 - **Support for vSphere 6.0 Platform Services Controller topologies:** NSX now supports external Platform Services Controllers (PSC), in addition to the already supported embedded PSC configurations.
 - **Support for vRealize Orchestrator Plug-in for NSX:** NSX 6.2 supports the [NSX-vRO plug-in](#) for integration of NSX with vRealize Orchestrator.

System Requirements and Installation

Product or component	Recommended version
NSX for vSphere	6.2.1
vSphere	5.5U3 or 6.0U1
Guest Introspection	Guest Introspection and Network Introspection-based features in NSX are compatible with specific VMware Tools (VMTTools) versions. To enable the optional NSX Network

Introspection Driver component packaged with VMware Tools, you must upgrade to one of:

- VMware Tools 5.1 P07 and later
- VMware Tools 5.5 P04 and later
- VMware Tools 6.0 P01 and later
- VMware Tools 10.0 and later

vRealize Orchestrator

NSX-vRO plugin 1.0.3

For the complete list of NSX installation prerequisites, see the [System Requirements for NSX](#) section in the *NSX 6.2 Installation Guide*.

Upgrade Notes

- **Upgrade paths:**
 - **From NSX 6.x:** See the [VMware Interoperability Matrix: Upgrade Paths](#) for supported upgrade paths. Be aware that a higher numerical release number is not an indicator of a supported upgrade path.
 - **For cross-VC sites:** If you are upgrading a cross-vCenter NSX installation, please see the [Cross-vCenter Upgrades](#) section, later in this document.
 - **From vCNS 5.x:** There is no support for direct upgrades of vCNS 5.x to NSX 6.2.1. Instead, using the NSX 6.2.2 upgrade bundle posted on or after 31 March, 2016, you may upgrade directly from VMware vCloud Network and Security (vCNS) 5.1.x or 5.5.x to NSX 6.2.2. For instructions, see the *NSX Upgrade Guide*, in the section, [vCloud Networking and Security 5.5.x to NSX 6.2 Upgrade](#).
 - **From NSX 6.1.6:** There is no support for upgrades from NSX 6.1.6 to NSX 6.2.0, 6.2.1, or 6.2.2.
 - **From NSX 6.1.5:** Upgrades from NSX 6.1.5 to NSX 6.2.0 are not supported. Upgrades from NSX 6.1.5 to NSX 6.2.1 are not recommended. Instead, VMware recommends upgrading to 6.2.2 or later to get the latest security updates.
 - **To validate that your upgrade to 6.2.x was successful see [knowledge base article 2134525](#).**
 - **Upgrading as part of a wider VMware product upgrade:** When you are upgrading NSX in context with other VMware product upgrades, such as vCenter and ESXi, it is important to follow the supported upgrade sequence documented in [knowledge base article 2109760](#).
 - **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:** The memory and CPU requirements for installing and upgrading NSX Manager changed from NSX 6.1.x to NSX 6.2.x. See the [System Requirements for NSX](#) in the *NSX 6.2 Installation* or the *NSX 6.2 Upgrade* documentation.
 - **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 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 6.2.0, you must make sure your installation is not using a VXLAN tunnel ID of 4094 on any tunnels. VXLAN tunnel ID 4094 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.
- **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.
- **Cross-vCenter upgrades:** The NSX Upgrade Guide does not provide upgrade instructions for sites that run cross-vCenter NSX. Please follow the instructions below to upgrade your cross-vCenter NSX installation:
 1. Upgrade NSX Manager as explained in the NSX Upgrade Guide. Upgrade the primary NSX Manager first and then upgrade all secondary NSX Managers. You must upgrade all NSX Managers to the same version. VMware does not support mixed-version sets of NSX Managers in a single cross-vCenter NSX installation.
 2. Upgrade NSX Controllers as explained in the NSX Upgrade Guide. We recommend that you upgrade the controllers in the same maintenance window as NSX Manager upgrade.
 3. Complete the upgrade by following the NSX Upgrade Guide, resuming at the section, "Upgrade Host Clusters to NSX 6.2".
- **vCNS upgrades:** Before Upgrading VMware vCloud Network and Security 5.x to VMware NSX for vSphere 6.2: If you plan to upgrade to VMware NSX for vSphere 6.2 from VMware vCloud Network and Security 5.5.x, verify whether uplink port name information is missing from the tables by running the following REST API call:

```
GET https://<nsxmgr-IP>/api/2.0/vdn/switches
```

In the output, look for the uplinkPortName field. For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<vdsContexts>
  <vdsContext>
    <switch>
      <objectId>dvs-22</objectId>
      <objectTypeName>VmwareDistributedVirtualSwitch</objectTypeName>
      <nsxmgrUuid>4236F6CA-3B1A-56BE-4B55-1EF82B8CA12D</nsxmgrUuid>
      <revision>2</revision>
      <type>
        <typeName>VmwareDistributedVirtualSwitch</typeName>
      </type>
      <name>1-vds-20</name>
      <scope>
        <id>datacenter-3</id>
        <objectTypeName>Datacenter</objectTypeName>
        <name>datacenter-1</name>
      </scope>
      <clientHandle />
    </switch>
  </vdsContext>
</vdsContexts>
```

```
    <extendedAttributes />
  </switch>
  <mtu>1600</mtu>
  <teaming>FAILOVER_ORDER</teaming>
  <uplinkPortName>uplink2</uplinkPortName>
  <promiscuousMode>>false</promiscuousMode>
</vdsContext>
</vdsContexts>
```

If the output of this command contains at least one uplink port name for each vSphere distributed switch, you can proceed with the upgrade. If the uplink port name is missing from the output, see [knowledge base article 2129200](#).

Known Issues

Known issues are grouped as follows:

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

General Known Issues

Some Log Insight reports are not supported in NSX 6.2

Due to incompatibilities between NSX 6.2 and the vRealize Content Pack for NSX, NSX 6.2 does not support the following vRealize Log Insight Reports:

- In Dashboard: NSX-vSphere-Infrastructure: The Host - Controller: Communication Errors widget is not supported
- In Dashboard: Logical Switch-Overview, the following widgets are not supported:
 - Logical Switch create audit events
 - Logical Switch update audit events
 - Logical Switch delete audit events
- In Dashboard: Logical Router-Overview, the following widgets are not supported:
 - Logical Router create audit events
 - Logical Router update audit events

Workaround: None. See [VMware knowledge base article 2143058](#) for updated information.

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), then one of the source or destination field should be set to "any".

NSX has a URL length limit length of 16000 characters if users want to assign a single security tag to x VMs at a time in one API call

Users cannot assign a single security tag to x number of VMs at a time using one API call, if the URL length is more than 16,000 characters. The URL length must have a maximum of 16,000 characters.

Workaround:

1. The URL length should be less than 16,000 characters.
2. Performance is optimized when approximately 500 VMs are being tagged in a single call. Tagging more VMs in a single call may result in degraded performance.

Discrepancy between service status reported in UI and service status reported in API

Under the Settings tab in the UI, the L2 service status is displayed as down, however the API shows the L2 status as up.

Workaround: Refresh the page.

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.

User must download NSX Controller logs sequentially

NSX Controller logs can be downloaded for troubleshooting. Due to a known issue, you cannot download more than one controller log 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.

Log files exported as CSV from NSX Manager are timestamped with epoch not datetime

When you export log files as CSV from NSX Manager using the vSphere Web Client, you might notice that the log files are timestamped with the epoch time in milliseconds, instead of with the appropriate time based on the time zone.

Workaround: None.

Unable to choose VMs on bridged network using the NSX traceflow tool

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

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

Flow Monitoring drops flows that exceed a 2 million flows / 5 minutes limit

NSX Flow Monitoring retains up to 2 million flow records. If hosts generate more than 2 million records in 5 minutes, new flows are dropped.

Note that NSX Flow Monitoring is production ready, however, it is only applicable for scenarios with lower throughput/connection per sec requirements. It can be also be used for troubleshooting, as well as for creating rules for limited durations. High throughput scenarios can experience issues such as high CPU usage on the NSX Manager, RMQ message bus overflow and failure to deploy policy updates. Additionally, issues have been seen with large UDP workloads. For ongoing flow information collection at enterprise scale, IPFIX is the recommended approach. Please keep in mind that once flow monitoring is disabled it can take as long as 15 days to purge flow data in the database.

Workaround: None.

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.

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.

In cross vCenter NSX deployments, multiple versions of saved 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>

```

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.

Unable to power on guest virtual machine

When you power on a guest virtual machine, the error All required agent virtual machines are not currently deployed may be displayed.

Workaround: Perform the following steps:

1. In the vSphere Web Client, click **Home** and then click **Administration**.
2. In Solution, select **vCenter Server Extension**.
3. Click **vSphere ESX Agent Manager** and then click the **Manage** tab.
4. Click **Resolve**.

Installation and Upgrade Known Issues

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

During a Cross-vCenter Upgrade NSX is disconnected from Primary VC after upgrading from NSX 6.2 to 6.2.1

Workaround: Clear the NSX bundles and restart the VC Web Client.

Windows:

1. Navigate to C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity.
2. Remove the folder named com.vmware.vShieldManagerxxxxx (No need to take the backup).
3. Restart VC Web Client.

Unix

1. Navigate to /etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity.
2. Remove the folder named com.vmware.vShieldManagerxxxxx (No need to take the backup).
3. Restart VC Web Client.

After upgrading the NSX Edge from 6.1.x to 6.2.x, the NSX Manager vsm.log shows "INVALID DHCP CONFIG"

If you have an interface with an IPv6 subnet, DHCP generates an empty shared subnet and treats it as an invalid operation.

Workaround: Disable DHCP service and upgrade the NSX Edge. After upgrading, enable DHCP.

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.

Upgrading Edge Services Gateway fails with "Timed out waiting for Edge vm" message

Applying an IPv6 address to the NSX management interface causes NSX Manager to use the host name. The vsfwd proxy which connects the Edge VM to NSX Manager does not correctly handle a FQDN, resulting in a error similar to "ERROR TaskFrameworkExecutor-6 AbstractEdgeApplianceManager:185 - Timed out waiting for Edge vm {}. Vm took too long to boot and respond com.vmware.vshield.edge.exception.VshieldEdgeException".

Workaround: change the configuration with a command similar to "esxcfg-advcfg -q -s "10.20.233.160" /UserVars/RmqIpAddress" or remove the "ipv6 address" command from the interface management configuration and use IPv4 only.

NSX Manager certificate replacement requires restart of NSX Manager and may require restart of vSphere Web Client

After you replace the NSX Manager appliance certificate, you must always restart the NSX Manager appliance. In certain cases after a certificate replacement, the vSphere web client will not display the "Networking and Security" tab. If this occurs, follow the workaround below.

Workaround: Restart the NSX Manager appliance and then restart the vSphere Web Client.

To restart NSX Manager, perform the following steps:

1. Log in to the NSX Manager CLI.
2. Switch to enable/privileged mode by typing en.
3. Stop the web-manager service by typing no web-manager. Wait for the OK to confirm it has stopped.
4. Start NSX Manager by typing web-manager. Wait for the OK to confirm it has restarted.
5. To restart the vSphere web client, in vCenter 5.5, open <https://{vcenter-ip}:5480> and restart the Web Client server.
6. In the vCenter 6.0 appliance, log in to the vCenter Server shell as a root user and run the following commands:

```
shell.set --enabled True

shell

localhost:~ # cd /bin

localhost:~ # service-control --stop vsphere-client

localhost:~ # service-control --start vsphere-client
```

7. In vCenter Server 6.0, run the following commands:

```
cd C:\Program Files\VMware\vCenter Server\bin

service-control --stop vsphere-client

service-control --start vsphere-client
```

After vCenter upgrade, vCenter might lose connectivity with NSX

If you are using vCenter embedded SSO and you are upgrading vCenter 5.5 to vCenter 6.0, vCenter might lose connectivity with NSX. This happens if vCenter 5.5 was registered with NSX using the root user name. In NSX 6.2, vCenter registration with root is deprecated. NOTE: If you are using external SSO, no change is necessary. You can retain the same user name, for example admin@mybusiness.mydomain, and vCenter connectivity will not be lost.

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

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.

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.
- The 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
```

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.

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.

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.

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.

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.

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.

After upgrading NSX from version 6.0 to 6.0.x or 6.1, NSX Edges are not listed on the UI

When you upgrade from NSX 6.0 to NSX 6.0.x or 6.1, the vSphere Web Client plug-in may not upgrade correctly. This may result in UI display issues such as missing NSX Edges.

This issue is not seen if you are upgrading from NSX 6.0.1 or later.

Workaround: Follow the steps below.

1. In vCenter mob, click **Content**.
2. In the Value column, click **ExtensionManager**.
3. Look for extensionList property value (for example com.vmware.vShieldManager) and copy the string.
4. In the Methods area, click **UnregisterExtension**.
5. In the Value field, paste the string that you copied in step 3.
6. Click **Invoke Method**.

This ensures deployment of the latest plug-in package.

NSX Edge upgrade fails if L2 VPN is enabled on the Edge

L2 VPN configuration update from 5.x or 6.0.x to 6.1 is not supported. Hence, Edge upgrade fails if it has L2 VPN configured on it.

Workaround: Delete L2 VPN configuration before upgrading NSX Edge. After the upgrade, re-configure L2 VPN.

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.

Momentary loss of third-party anti-virus protection during upgrade

When upgrading from NSX 6.0.x to NSX 6.1.x or 6.2.0, 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.

Host error message appears while configuring distributed firewall

While configuring distributed firewall, if you encounter an error message related to the host, check the status of fabric feature com.vmware.vshield.nsxmgr.messagingInfra. If the status is Red, perform the following workaround.

Workaround: Use the following REST API call to reset communication between NSX Manager and a single host or all hosts in a cluster.

```
POST https://<NSX Manager IP>/api/2.0/nwfabric/configure?action=synchronize
```

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

Copy and paste of a firewall rule with negate source/destination enabled will list a new rule with Negate option disabled

If a firewall rule is copied and pasted with the negate source/destination option enabled, after the paste operation there will be a new firewall rule, however, the "negate source/destination" option is disabled.

Workaround: None.

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.

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](#)

After upgrading from vCNS 5.5.4 to NSX 6.2.0, the firewall on the Host Preparation tab remains disabled

After upgrading from vCNS 5.5.x to NSX 6.2.0 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.

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 in progress 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 in progress state and resolve the issue, reboot the NSX Manager virtual appliance.

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.

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](#).

PSOD during ESXi upgrade

When you are upgrading an NSX-enabled vSphere 5.5U2 host to vSphere 6.0, some of the ESXi host upgrades might halt with a purple diagnostic screen (also known as a PSOD).

Workaround: If this occurs, see [knowledge base article 2137826](#).

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.

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.

In a cross vCenter deployment, a universal 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.

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.

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.

Workaround: None.

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.

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**).

SSLv2 and SSLv3 deprecated in NSX 6.2

Starting in NSX 6.2, the SSL VPN gateway only accepts the TLS protocol. After the NSX upgrade, any new NSX 6.2 clients that you create automatically use the TLS protocol during connection establishment. When an NSX 6.0.x client tries to connect to an NSX 6.2 gateway, the connection establishment fails at the SSL handshake step.

Workaround: After the upgrade to NSX 6.2, uninstall your old SSL VPN clients and install the NSX 6.2 version of the SSL VPN clients.

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.

After upgrade to NSX 6.2, NSX Manager has more than 100 percentage of physical memory allocated

Starting in NSX 6.2, NSX Manager requires 16 GB of reserved memory. The former requirement was 12 GB.

Workaround: Increase the NSX Manager virtual appliance's reserved memory to 16 GB.

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.

NSX Manager Known Issues

Unable to add Secondary NSX manager if GUI is Japanese language on Firefox browser

While adding Secondary NSX Manager with German, Japanese, Korean, or French locale with a Firefox browser, the thumbprint dialog is not shown. Because of this, the user cannot configure a Secondary NSX Manager while using these locales.

Workaround: Use Internet Explorer or English locale.

NSX management service doesn't come up when the hostname's length is more than 64 characters.

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

NSX Manager list slow to display in Web Client

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

NSX controller shows as disconnected

NSX Manager logs report disconnection to controllers via a message similar to "ERROR http-nio-127.0.0.1-7441-exec-16908 BaseRestController:339 - Exception : 'I/O error: Read timed out; nested exception is java.net.SocketTimeoutException: Read timed out'". This condition occurs when an intermediate firewall on the network blocks the TCP/IP FIN message after the idle timeout value is reached. When this condition is occurring, the number of connections to the NSX Manager will increase.

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.

Previous backups are not displayed at the NSX Manager UI

Running a backup operation never shows a successful completion at the NSX Manager UI. Either one of these issues may manifest if a large number of backup files are stored in the destination folder. Each backup file has to be checked for compatibility before displaying the list on the same page. The current file list process can cause the page to timeout.

Workaround: Reduce the number files stored in the backup folder by periodic cleaning on your storage server or move older backups to another folder. To verify the backup completed, look for messages similar to the following in the NSX Manager log file: 2015-07-01 22:10:55.869 GMT INFO http-nio-443-exec-250 VsmServiceBackupRestoreExecutor:236 - Run backup script - Start 2015-07-01 22:14:35.992 GMT INFO http-nio-443-exec-250 VsmServiceBackupRestoreExecutor:278 - Run backup script - Completed

Service Composer goes out of sync when policy changes are made while one of the Service Managers is down.

This is related to instances of multiple Services/Service Managers registered and policies created referencing those services. If changes are made in Service Composer to such a policy when one of those Service Managers is down, the changes fail because of callback failure to the Service Manager that is down. As a result, Service Composer goes out of sync.

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

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.

After NSX Manager backup is restored, REST call shows status of fabric feature

com.vmware.vshield.nsxmgr.messagingInfra as Red

When you restore the backup of an NSX Manager and check the status of fabric feature

com.vmware.vshield.nsxmgr.messagingInfra using a REST API call, it is displayed as Red instead of Green.

Workaround: Use the following REST API call to reset communication between NSX Manager and a single host or all hosts in a cluster.

POST https://<NSX Manager IP>/api/2.0/nwfabric/configure?action=synchronize

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

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.

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.

Syslog shows host name of backed up NSX Manager on the restored NSX Manager

Suppose the host name of the first NSX Manager is A and a backup is created for that NSX Manager. Now a second NSX Manager is installed and configured to the same IP address as the old Manager according to backup-restore docs, but host name is B. Restore is run on this NSX Manager. The restored NSX Manager shows host name A just after restore and host name B again after reboot.

Workaround: Host name of second NSX Manager should be configured to be same as the backed up NSX Manager.

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.

NSX Manager UI does not automatically logs out users after changing password using NSX Command Line Interface

If 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.

Standalone NSX Manager incorrectly allows import of universal firewall configuration

Typically, NSX Manager running in standalone role should allow the import of local firewall rules only. Starting in NSX 6.2, NSX Manager can be run in standalone role (managing networks for one vCenter) or cross vCenter mode where it incorrectly allows you to import a universal firewall rule into an NSX Manager environment running in standalone role. Once imported you cannot delete the universal firewall rules either through REST API or the vSphere Web Client. As the manager is currently running in standalone role where the universal section is treated like a local section.

Workaround: If you are running NSX Manager in standalone role, do not import a firewall configuration that contains universal rules. If you have already imported a universal firewall rule into a standalone NSX Manager, fix the issue by importing a saved firewall configuration file that does not contain universal rules, and publish that configuration file by loading it in the Firewall table.

Perform the following steps:

1. Log in to the vSphere Web Client.
2. Click **Networking & Security** and then click **Firewall**.
3. Click the **Firewall** tab.

4. Click the **Saved Configurations** tab.
5. Click the **Import configuration (import)** icon.
6. Click **Browse** and select the file containing the configuration that you want to import.

Rules are imported based on the rule names. During the import, Firewall ensures that each object referenced in the rule exists in your environment. If an object is not found, the rule is marked as invalid. If a rule referenced a dynamic security group, the dynamic security group is created in NSX Manager during the import.

7. Add the node back as a secondary node. The synchronizing across NSX Managers automatically syncs up the universal section correctly performing any required cleanup.

Once you have successfully published the configuration file, the rules are pushed down to the host and impact the datapath. The system functions as expected.

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.

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.

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

Multicast traffic is dropped from some logical switches backed by Arista switch running EOS-4.12.7.1

In NSX installations where the physical network includes Arista switches running switch version EOS-4.12, logical multicast traffic may be dropped. This affects the case in which two ESX VTEPS are communicating over VxLAN port 8472 using an affected Arista switch as the physical underlay. All Arista 7150 products are affected when running any current major or maintenance versions of EOS 4.12.

Workaround: Customers with Arista switches can work around the problem two ways:

1. Upgrade your Arista switches to EOS 4.13.0 or newer.
2. Configure NSX to use a different UDP port using the NSX REST API. VMware recommends that you use port 4789 per the VXLAN specification.

Delays in connecting to NSX load balancer does not provide consistent connections over multiple VIPs

When the load balancer is configured to use Source IP Hash load balancing, a connected client session receives a consistent connection to a backend server. The load balancer should also be able to provide, for a given connected client, consistent connections over multiple VIPs if those VIPs are backed by the same server pool. That is, when one backend server serves multiple VIPs, a given client's connection to one of that backend server's VIPs should guarantee that that client will use the same backend server when connecting to other VIPs served by that backend server. A known issue prevents the NSX load balancer from providing such consistent connections over multiple VIPs.

Delays in connecting to RIB and FIB after assigning IP address to interface

When you attempt to assign an IP address to an interface, typically, the interface information is updated immediately. However, when the polling interval increases, you might experience delay in assigning IP address.

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

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

Total number of static bindings on a dhcp edge cannot go beyond 2048

If you exceed 2048, the following error message is appears "Total number of bindings and pools should not exceed 2,048."

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>`

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.

Distributed logical router advertises incorrect next hop for default route when BGP neighbor filter is set to "ANY , OUT , DENY"

With 'default originate' enabled on an NSX distributed logical router (DLR), setting a BGP neighbor filter of "ANY , OUT , DENY" 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:

- Direction: OUT
- Action: Deny
- Network: Any

Workaround: None.

Disabling a routing protocol on NSX Edge may result in temporary loss of data traffic

Disabling a routing protocol on NSX Edge does not send route-withdraw request to the peer, so traffic is black-holed until hold-down timer/dead timer expires.

Workaround: None.

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.

ESG syslog not able to send to remote server, saying it can't resolve hostname, however, DNS resolver is working
Immediately after deployment of an Edge, the syslog is unable to resolve 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. This issue was first seen in 6.2.

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.

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.

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

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

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

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

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

6. In *opaqueDataSpec*, add the following XML

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

7. Set **isRuntime** to false.
8. Click **Invoke Method**.
9. Repeat steps 5 through 8 for each trunk port configured on the deleted Edge virtual machine.

Security Services Known Issues

No connectivity on VLAN ID 0 in L2 VPN

NSX L2 VPN configuration incorrectly allows the user to configure an L2 VPN with VLAN ID 0. Once configured, no traffic can flow on this VPN.

Workaround: Workaround: Use a valid VLAN ID in the range from 1 to 4094.

There is no support for Cipher 3C (SHA-256) encryption algorithms for SSLVPN-Plus

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.

Using IPSet as source/destination in NetX rule displays the error Invalid container type: IPSet

Workaround: Instead of using IPSet as the source/destination for netX rules, create a security group and make ipset as a member of it.

This is a known issue in 6.2.1

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.

This is a known issue in 6.2.1

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.

This is a known issue in both 6.1.5 and 6.2.1

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](#).

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.

NSX UI takes 3 or more minutes to load Active Directory users/groups when creating Security Group

When configuring Security Groups and selecting "Directory Group", it takes the NSX UI 3 or more minutes to populate the list of Active Directory users and groups. There is no workaround.

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.

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.

Some versions of Palo Alto Networks VM-Series do not work with NSX Manager default settings

Some NSX 6.1.4 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.

In upgraded NSX installations, publishing a firewall rule may result in Null Pointer exception in Web Client

In upgraded NSX installations, publishing a firewall rule may result in a Null Pointer exception in the UI. The rule changes are saved. This is a display issue only.

If you delete the firewall configuration using a REST API call, you cannot load and publish saved configurations

When you delete the firewall configuration, a new default section is created with a new section ID. When you load a saved draft (that has the same section name but an older section ID), section names conflict and display the following error:

Duplicate key value violates unique constraint `firewall_section_name_key`

Workaround: Perform one of the following:

- Rename the current default firewall section after loading a saved configuration.
- Rename the default section on a loaded saved configuration before publishing it.

Monitoring Services Known Issues

Unable to add more than eight uplink interfaces during Distributed Logical Router (DLR) deployment using vSphere Web Client

Workaround: Wait for the DLR deployment to complete and then add additional interfaces to Distributed Logical Router.

Resolved Issues

The following issues were resolved in 6.2.0 and 6.2.1:

Resolved issues are grouped as follows:

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

General Resolved Issues

- **VMware ESXi 5.x and 6.x experiences a purple diagnostic screen when using IP discovery in VMware NSX for vSphere 6.2.0 (2134329)**
When using IP discovery on logical switches in VMware NSX for vSphere 6.2.0, the ESXi 5.x and 6.x host fails with a purple diagnostic screen as explained in [knowledge base article 2134329](#)

This has been fixed in NSX 6.2.1.

- **Manage option on Security Tag portlet is grayed out by default**
On a Virtual Machine's summary page, the "Manage" hyperlink on the security tag portlet remains grayed out till the user creates a new security tag.

This has been fixed in NSX 6.2.1.

- **Some Controller logs not available for syslog export.**
Controller logs, including Zookeeper clustering logs, are not part of syslog export.

This has been fixed in NSX 6.2.1.

- **ESXi 6.0 PSOD on vdl2 when pinging with data size higher than available data size for the MTU**
Starting ping from NSX host switch attached vmknics will lead to host PSOD if data size is greater than MTU.

This has been fixed in NSX 6.2.1.

- **Users needed to configure same IP address and port number for both TCP and UDP protocol**
This release resolves the following issues as well:

- UDP virtual server without pool configuration leads to configuration failure.
- Statistics shows incorrect data when UDP virtual server is not associated with any pool.

This has been fixed in NSX 6.2.1. With 6.2.1 release, users can use the same IP address and port number for both TCP and UDP with/without a pool associated.

Install and Upgrade Resolved Issues

- **SSL VPN-Plus client cannot be installed on Mac OS X Yosemite and higher**
Earlier versions of Mac OS X are supported.

This has been fixed in NSX 6.2.1.

- **After upgrading NSX vSphere from 6.0.7 to 6.1.3, vSphere Web Client crashes on login screen**
After upgrading NSX Manager from 6.0.7 to 6.1.3, you will see exceptions displayed on the vSphere Web Client UI login screen. You will not be able to login and perform operations on either vCenter or NSX Manager.

This has been fixed in NSX 6.2.0.

- **Guest Introspection installation fails with error**
When installing Guest Introspection on a cluster, the install fails with the following error:
Invalid format for VIB Module

This has been fixed in NSX 6.2.0.

- **DVPort fails to enable with "Would block" due to host prep issue**
On an NSX-enabled ESXi host, the DVPort fails to enable with "Would block" due to a host preparation issue. When this occurs, the error message first noticed varies (for example, this may be seen as a VTEP creation failure in VC/hostd.log, a DVPort connect failure in vmkernel.log, or a 'SIOCSIFFLAGS' error in the guest). This happens when VIBs are loaded after the vSphere Distributed Switch (vDS) properties are pushed by vCenter. This may happen during upgrade. See [knowledge base article 2107951](#).

This has been fixed in NSX 6.2.0.

- **Attempts to delete existing NSX Edge Gateway fail in an environment upgraded to NSX 6.1.4**
In NSX installations upgraded from 6.1.3 to 6.1.4, the existing NSX Edge Gateways cannot be deleted after the upgrade to 6.1.4. This issue does not affect new Edge Gateways created after the upgrade. Installations that upgraded directly from 6.1.2 or earlier are not affected by this issue.

This has been fixed in NSX 6.2.0.

- **The AES encryption unavailable when performing an NSX backup using third-party secured FTP backup**

This has been fixed in NSX 6.2.0.

- **NSX Manager UI does not display user-friendly error messages during host reboot**
In this 6.2 release, NSX Manager UI is updated to display detail error messages that describe the problems you might encounter during host reboot and provide possible solution.

This has been fixed in NSX 6.2.0.

- **Unable to install NSX VIB installation**
The installation of NSX VIB might not complete, as expected if the ixgbe driver fails to load from third-party module because it has been locked and prevents it from being used for installation.

This has been fixed in NSX 6.2.0.

- **Unable to start NSX Manager service after upgrading from vCloud Networking and Security (vCNS) 5.5.3**
After upgrading vCloud Networking and Security (vCNS) 5.5.3 to NSX 6.1.3, the NSX Manager service hangs and is unable to start successfully.

This has been fixed in NSX 6.2.0.

- **The message bus randomly does not start after NSX Edge reboot**
After restarting an Edge VM, the message bus often does not start after powering on, and an additional reboot is required.

This has been fixed in NSX 6.2.0.

NSX Manager Resolved Issues

- **Starting in 6.2.1, NSX Manager queries each controller node in the cluster to get the connection information between that controller and the other controllers in the cluster.**
This is provided in the output of the NSX REST API ("GET https://[NSX-MANAGER-IP-ADDRESS]/api/2.0/vdn/controller" command), which now shows the peer connection status between among the controller nodes. If NSX Manager finds the connection between any two controller nodes is broken, a system event is generated to alert the user.

This has been fixed in NSX 6.2.1.

- **Force-sync of controller is broken if backup-restore of manager is done on another appliance**
If an NSX Manager appliance is cloned and/or restored from a backup, a force-sync operation to an NSX controller cluster will fail. This issue does not occur for an NSX Manager deployed from scratch.

This has been fixed in NSX 6.2.1.

- **NSX logging heartbeat failures for hosts that are not part of the NSX installation**
When an NSX-prepared host is directly removed from the vCenter inventory (without first unpreparing it in NSX), NSX receives an unexpected 'Host Connected' DCN which causes partial removal of messaging infrastructure components from the host. As a result, the messaging link between NSX and the host may remain active when it should have been removed, and NSX may raise false 'Alert' SystemEvents for the host. This has been fixed in NSX 6.2.1.

This has been fixed in NSX 6.2.1.

- **NSX Manager is non-functional after running the write erase command**
When you restart the NSX Manager after running the write erase command, you might notice that the NSX Manager is not working as expected, such as the password to access the Linux shell has been reset, the setup command is missing, and so on.

This has been fixed in NSX 6.2.0.

- **Add Domain shows error at LDAP option with Use Domain Credentials**
In NSX 6.1.x, the user when trying to add an LDAP domain, the web client gave a User Name was not specified error, even when Username was provided in UI. This has been fixed in NSX 6.2.0.

This has been fixed in NSX 6.2.0.

- **CA signed certificate import needs an NSX Manager reboot before becoming effective**
When you import an NSX Manager certificate signed by CA, the newly imported certificate does not become effective until NSX Manager is rebooted.

This has been fixed in NSX 6.2.0.

- **Unable to import NSX Manager to LDAPS domain**
When you attempt to add NSX manager to LDAPS domain, the following error message appears.
Cannot connect to host <Server FQDN>
error message: simple bind failed: <Server FQDN:Number>

This has been fixed in NSX 6.2.0.

Logical Networking Resolved Issues and NSX Edge Resolved Issues

- **RADIUS authentication server configuration fails on NSX Edge**
In NSX 6.1.5 and earlier, the RADIUS server secret key string had a 32-character limit; if the string exceeded this character limit, the RADIUS server failed to connect with the NSX Edge. The limit is now 64 characters.

This was fixed in NSX 6.2.0.

- **VIO Heat stack deployment fails intermittently for the VMware NSX for vSphere 6.x Edge with the error: Cannot allocate memory**
Health monitoring memory usage increases over time, eventually causing edge failure.

This has been fixed in NSX 6.2.1.

- **BGP filters are taking approximately 40 seconds to be effectively applied.**
During this period all the redistribution policies are applied without filters. This delay applies only to NSX Distributed Logical Router (DLR) for OUT directions.

This has been fixed in NSX 6.2.0.

- **On NSX Edge subinterfaces, ICMP redirects are sent out, even when the Send ICMP redirect option is disabled**
By default, NSX Edge subinterfaces have Send ICMP redirect disabled. Although this option is disabled, ICMP redirects are sent out on edge subinterfaces.

This has been fixed in NSX 6.2.0.

- **Cannot add non-ASCII characters in bridge or tenant name for Logical Router**
NSX controller APIs do not support non-ASCII characters.

This has been fixed in NSX 6.2.0.

- **When a BGP neighbor filter rule is modified, the existing filters may not be applied for up to 40 seconds**
When BGP filters are applied to an NSX Edge running IBGP, it may take up to 40 seconds for the filters to be applied on the IBGP session. During this time, NSX Edge may advertise routes which are denied in the BGP filter for the IBGP peer.

This has been fixed in NSX 6.2.0.

- **One of the NSX Controllers does not hand over master role to other controllers when it is shut down**
Typically, when a controller assumes operations master role and is preparing to shut down, it automatically hands over the master role to other controllers. In this case, the controller fails to hand over the role to other controllers and the status becomes interrupted and then goes into disconnected mode.

This has been fixed in NSX 6.2.0.

- **Unable to pass VXLAN traffic between hosts with unicast or multicast**
When VMs are on the same host they can communicate across VXLAN with unicast or multicast, but cannot communicate when VMs are on different hosts.

This has been fixed in NSX 6.2.0.

- **Removing multiple BGP rules on NSX Edge/DLR at the same time causes web client to crash**

This has been fixed in NSX 6.2.0. You can now delete multiple BGP rules at a time.

- **Protocol address is briefly displayed after adding Border Gateway Protocol (BGP) deny rule**
You might notice that the protocol address is briefly displayed after adding Border Gateway Protocol (BGP) deny rule in NSX Edge services gateway.

This has been fixed in NSX 6.2.0.

- **VMs disconnect during vMotion**
You might notice that VMs disconnect during vMotion or you might receive alerts for VMs with disconnected NICs.

This has been fixed in NSX 6.2.0.

- **Unable to download controller snapshot**
When downloading controller snapshots, you might notice that you are unable to download snapshot for the last controller. For example, if you have three controllers, one can successfully download snapshots of the first two but you might fail to download snapshot of the third controller.

This has been fixed in NSX 6.2.0.

- **When configuring a virtual server, the previously selected IP address is applied**
When creating a new virtual server, you might notice that the IP address is automatically applied from the list of previously selected IP pool. This happens when you have previously selected an IP pool to derive the Virtual Server IP. When you attempt to edit the virtual server IP Pool information, the information is not automatically sent to the backend from the UI and previous IP address derived from the IP Pool is automatically applied.

This has been fixed in NSX 6.2.1.

- **When HA is enabled on Edge Services Gateway, OSPF hello and dead interval configured to values other than 30 seconds and 120 seconds respectively can cause some traffic loss during failover**

When the primary NSX Edge fails with OSPF running and HA enabled, the time required for standby to take over exceeds the graceful restart timeout and results in OSPF neighbors removing learned routes from their Forwarding Information Base (FIB) table. This results in dataplane outage until OSPF re-initiates converges.

This has been fixed in NSX 6.2.0.

- **VMs are unable to receive ping from Edge DHCP server**
VM's can ping the Edge gateway but unable to receive DHCP ping from an Edge gateway trunk over an overlay network. The Edge DHCP server is setup as a trunk port and fails to pass or receive any traffic. However, when the Edge Gateway and the DHCP Edge are on the same host they are able to ping each other. When the DHCP Edge is moved to another host, the DHCP Edge is unable to receive ping from the Edge Gateway.

This has been fixed in NSX 6.2.0.

- **Edge Load Balancer stats not correctly displayed in the vSphere Web Client**

The Load Balancer does not display the number of concurrent connection statistics in the chart in vSphere Web Client UI.

This has been fixed in NSX 6.2.0.

- **When the direct aggregate network in local and remote subnet of an IPsec VPN channel is removed, the aggregate route to the indirect subnets of the peer Edge also disappears**

When there is no default gateway on Edge and you remove all of the direct connect subnets in local subnets and part of the remote subnets at the same time when configuring IPsec, the remaining peer subnets become unreachable by IPsec VPN.

This has been fixed in NSX 6.2.0.

- **Unable to pass traffic through load balancer after upgrading to NSX 6.1.2 or later**

When using option Insert X-Forwarded-For on NSX Edge Load Balancer, traffic may not pass through the load balancer.

This has been fixed in NSX 6.2.0.

- **Running the clear ip ospf neighbor command returns a segmentation fault error**

This has been fixed in NSX 6.2.0.

- **Unable to process Kerberos requests**

Certain Kerberos requests are failing when being balanced with an NSX Edge.

This has been fixed in NSX 6.2.0.

Security Services Resolved Issues

- **If you rename an existing firewall draft, the operation will fail with the UI displaying "Internal Server Error."**

This has been fixed in NSX 6.2.1.

- **Some DFW central CLIs show "ERROR output 100" output**

In some situations, where a Virtual Network Adapter (vNIC) is disconnected, a discrepancy can arise between the vNIC state information in NSX Manager and the Host leading to an "ERROR output 100" in the centralized CLI.

This has been fixed in NSX 6.2.1.

- **Application Profile list is not sorted.**

The list of Application Profile names in NSX Edge when Service Insertion is enabled, is presented in an unordered fashion. This release incorporates the fix to present the Application Profile list in a sorted manner.

This has been fixed in NSX 6.2.1.

- **Central CLI commands that are run for a specific ESXi host time out on some setups.**

This has been fixed in NSX 6.2.1.

- **The vsfwd.log gets overwritten quickly with a large number of container updates**

After the SpoofGuard policy is changed the NSX Manager promptly sends the change to host but host takes longer to process the change and update the state of the virtual machine's SpoofGuard state.

This has been fixed in NSX 6.2.0.

- **Cannot configure NSX firewall using security groups or other grouping objects defined at global scope**

Administrator users defined at the NSX Edge scope cannot access objects defined at the global scope. For example, if user *abc* is defined at Edge scope and security group *sg-1* is defined at global scope, then *abc* will not be able to use *sg-1* in firewall configuration on the NSX Edge.

This has been fixed in NSX 6.2.0.

- **Delayed mouse movement when viewing FW rules**

In the NSX Networking and Security section of vSphere Web Client, moving the mouse over rows in the Firewall Rules display results in a 3 second delay.

This has been fixed in NSX 6.2.0.

- **UI shows error Firewall Publish Failed despite successful publish**
If Distributed Firewall is enabled on a subset of clusters in your environment and you update an application group that is used in one or more active firewall rules, any publish action on the UI will display an error message containing IDs of the hosts belonging to the clusters where NSX firewall is not enabled.
Despite error messages, rules will be successfully published and enforced on the hosts where Distributed Firewall is enabled.

This has been fixed in NSX 6.2.0.

- **Deleting security rules via REST displays error**
If a REST API call is used to delete security rules created by Service Composer, the corresponding rule set is not actually deleted in the service profile cache resulting in an `ObjectNotFoundException` error.

This has been fixed in NSX 6.2.0.

- **Firewall rules do not reflect newly added virtual machine**
When new VMs were added to the logical switch, firewall rules are not updated correctly to include the newly added VMs. If you make a change to the firewall and publish changes the new objects are added to the policy.

This has been fixed in NSX 6.2.0.

- **Cannot select Active Directory objects when configuring security groups**
In NSX 6.1.x, AD/LDAP Domain Objects took a long time to return in the Security Group Object selection screen.

This has been fixed in NSX 6.2.0.

- **Cannot add firewall rule with source/destination as multiple comma-separated IP addresses**

This has been fixed in NSX 6.2.0.

- **Unable to move NSX Distributed Firewall (DFW) section at the top of the list**
When using Service Composer to create a security group policy, the section created in the DFW table cannot be added to the top of the list. There is no way to move DFW section up or down.

This has been fixed in NSX 6.2.0.

- **Security policy configured as a port range causes firewall to go out of sync**
Configuring security policies as a port range (for example, "5900-5964") will cause the firewall to go out of sync with a `NumberFormatException` error.

This has been fixed in NSX 6.2.0.

Monitoring Services Resolved Issues

- **When reporting flow statistics, index 0 (bytes-in) and index 1 (bytes out) counts are sometimes reversed.**
Index 0 holds the counts for traffic for the origination direction and, and index 1 holds the counts for traffic in the reverse direction

This has been fixed in NSX 6.2.1.

- **The #show interface command does not display the bandwidth/speed of vNic_0 interface**
After running the "#show interface" command, a full duplex, "0 M/s" speed is displayed but not the bandwidth/speed of NSX Edge vNic_0 interface.

This has been fixed in NSX 6.2.0.

- **When IPFIX configuration is enabled for Distributed Firewall, firewall ports in the ESXi management interface for NetFlow for vDS or SNMP may be removed**
When a collector IP and port is defined for IPFIX, the firewall for ESXi management interface is opened up in the outbound direction for the specified UDP collector ports. This operation may remove the dynamic ruleset configuration on ESXi management interface firewall for the following services if they were previously configured on the ESXi host:
 - Netflow collector port configuration on vDS
 - SNMP target port configuration

This has been fixed in NSX 6.2.0.

- **Unable to process Denied/Block events through IPFIX protocol**
Typically, vsfwd user process handles the collection of flows, including dropped/denied ones and processes them for

IPFIX. This happens when the IPFIX Collector fails to see the Denied/Block events because the vSIP drop packet queue is either too narrow or is wrapped around by inactive flow events. In this release, the ability to send Denied/Block events using the IPFIX protocol is implemented.

This has been fixed in NSX 6.2.0.

Solution Interoperability Resolved Issues

- **Unable to set up organizational network**

When attempting to set up an organization-wide network, vCloud Director fails with an error message.

This has been fixed in NSX 6.2.0.

- **Unable to launch multiple VMs using VIO setup**

Users using VMware Integrated OpenStack were unable to launch large numbers of VMs or publish large numbers of firewall rules in a short period of time. This resulted in Error publishing ip for vnic messages in the log.

This has been fixed in NSX 6.2.0.

The following issues were resolved in 6.1.5 and 6.2.1:

Resolved issues are grouped as follows:

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

General Resolved Issues

- **One of the NSX Controllers does not hand over master role to other controllers when it is shut down**

Typically, when an NSX Controller assumes operations master role and is preparing to shut down, it automatically hands over the master role to other controllers. In this error case, the controller fails to hand over the role to other controllers, its status becomes interrupted, and it goes into disconnected mode.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Control plane connectivity fails for NSX Controller**

Control plane connectivity was seen to fail for a Controller, showing an error in netcpa related to txInProgress.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

Installation and Upgrade Resolved Issues

- **After NSX upgrade, guest introspection fails to communicate with NSX Manager.**

After upgrading from NSX 6.0.x to NSX 6.1.x or from NSX 6.0.x to NSX 6.2 and before the guest introspection service is upgraded, the NSX Manager cannot communicate with the guest introspection Universal Service Virtual Machine (USVM). The loss of communication between NSX Manager and guest introspection leads to a loss of protection for VMs in the NSX cluster when there is a change to the VMs, such as VM additions, vMotions, or deletions. The NSX Installation > **Service Deployments** tab shows the current version of guest introspection. When this issue is present, a warning appears in the Service Status column. The warning message includes the list of affected hosts and the error message, Guest Introspection is not ready.

Workaround: To resolve the issue, follow the procedure to upgrade guest introspection in the [NSX Upgrade documentation](#).

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

NSX Manager Resolved Issues

- **NSX Manager CPU utilization is high after adding it to Active Directory domain**

NSX Manager CPU utilization is high after adding it to Active Directory domain. In the system logs of the NSX Manager, multiple Postgres threads are seen as running.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Unable to register NSX Manager 6.1.4 with vCenter, gives error: NSX Management Service operation failed**

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **NSX manager web client displays error: Code 301002**

Description: When you navigate to NSX manager > Monitor > System Events, the web client displays the following message: Filter config not applied to vnic. Code 301002.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

Logical Networking Resolved Issues

- **Connectivity loss after removing a logical interface (LIF) in installations with dynamic routing**

A problem was identified in the NSX Logical Router (Edge/DLR) when using dynamic routing (OSPF & BGP) that will cause network connectivity loss after removing a LIF. This affects NSX versions 6.0.x through 6.1.4.

In NSX installations that use dynamic routing, each LIF has a redistribution rule index ID associated with it. When a user deletes a LIF in such installations, the index IDs assigned to some active LIFs may change. This index modification can result in a temporary loss of network connectivity for LIFs whose index IDs are changed. If the LIF deletion is serialized, you will see 5-30 seconds of disruption on affected LIFs after each LIF deletion. If the LIF deletion is done in bulk, you will see a total of 5-30 seconds of disruption on affected LIFs.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

Networking and Edge Services Resolved Issues

- **OSPF routes configured on NSX Edge Services Gateway (ESG) not honored in the logical router (DLR), and affected packets are dropped**

The problem occurs in cases when OSPF uses IP_HDRINCL option. On certain Linux kernels, when this option is present, it prevents the IP stack from fragmenting the packets. Hence, any packets greater than the interface MTU are dropped.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

Syslog shows host name of backed up NSX Manager on the restored NSX Manager

Suppose the host name of the first NSX Manager is A and a backup is created for that NSX Manager. Now a second NSX Manager is installed and configured to the same IP address as the old Manager according to backup-restore docs, but host name is B. Restore is run on this NSX Manager. The restored NSX Manager shows host name A just after restore and host name B again after reboot.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **ESXi host might lose network connectivity**

An ESXi host might lose network connectivity and experience stability issues when multiple error messages similar to the following are logged in:

WARNING: Heartbeat: 785: PCPU 63 didn't have a heartbeat for 7 seconds; *may* be locked up.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **VMs disconnect during vMotion**

VMs disconnect during vMotion on 6.0.8 with message, VISP heap depleted.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Cannot redeploy NSX Edge with L2VPN Service configured with CA-signed certificate**

Cannot redeploy or change size of NSX Edge with L2VPN Service configured with CA-signed or self-signed certificate.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **The message bus randomly does not start after NSX Edge reboot**

After restarting an Edge VM, the message bus often does not start after powering on, and an additional reboot is required.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

Security Services Resolved Issues

- **NSX for vSphere 6.x Controllers disconnect intermittently**

Due to an IPSEC bug in the StrongSWAN package that was shipping in 6.1.4 and earlier releases, the tunnels between controllers weren't established after IPSEC rekeying. This caused partial connectivity failures between controllers resulting in multiple different issues. For more information see [knowledge base article 2127655](#).

This has been fixed in NSX 6.1.5 and 6.2.1

- **LDAP Domain Objects take too long to return or fail to return in Security Group Object Selection screen**

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Delayed mouse movement when viewing FW rules**

In NSX Networking and Security section of vSphere Web Client, moving the mouse over rows in the Firewall Rules display results in a 3 second delay each time the mouse is moved.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Some IP Spoofguard rules in NSX-v are not applied correctly**

Some IP Spoofguard rules in NSX-v are not applied correctly. Instance is not present in the Security Group in NSX-v and needs to be manually added to the security group

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Bulk deletion in Service Composer user interface generates "between 0 to 0" message**

Bulk deletion of policies (~100) from the NSX Service Composer user interface generates a message, "It should be between 0 to 0". You may safely ignore this message.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Background operation for Policy deletion may take long time with high CPU utilization**

Deletion of a policy reevaluates all the remaining policies in background. This may take more than an hour on setups having large number of policies, large number of security groups, and/or large number of rules per policy.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **All queued publishable tasks are marked as failed after the default timeout of 20 minutes**

Queues are maintained per NSX Edge and can publish in parallel for different Edges. The queued up publishable tasks are executed sequentially where each task takes approximately 3-4 seconds, and 300-400 tasks are completed in 20 minutes. In situations where more than 400 publish tasks for an Edge are queued up in a short time and have exceeded the publish timeout limit of 20 minutes while waiting for execution, the tasks are automatically marked as failed. NSX Manager responds to the failure by reverting to the last known successful configuration where publication to the Edge has succeeded. Applications or plugins that are sending configuration updates for an Edge to the NSX Manager in a burst mode need to monitor the success and failure status of the task using the associated job id.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

Solution Interoperability Resolved Issues

- **Copy of VM via vCloud Connector fails when route traverses NSX Load Balancer**

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **In VIO Deployment, some newly deployed VMs appear to have valid port and IPs assigned but do not have access to the network**

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

- **Slow login to NSX tab of vSphere web client with AD-backed SSO**

In NSX for vSphere installations that use SSO for AD authentication, the user's initial login to the NSX Networking and Security section of the vSphere Web Client takes a long time.

This has been fixed in NSX 6.1.5 and NSX 6.2.1.

Document Revision History

20 Aug 2015: First edition for NSX 6.2.0.

04 Sept 2015: Second edition for NSX 6.2.0. Removed unneeded upgrade warning.

17 Dec 2015: First edition for NSX 6.2.1.

27 Feb 2016: Second edition for NSX 6.2.1. Clarified NSX-Log Insight compatibility.

20 May 2016: Third edition for NSX 6.2.2. Noted 6.1.x upgrade limitations.