

---

## Release Notes for SD-WAN Release 17.1.x

These release notes accompany Viptela Software Release 17.x, for Releases 17.1.0 through 17.1.5. The Viptela software runs on all Viptela devices, including vSmart controllers, vEdge routers, vBond orchestrators, and vManage NMSs.

Cisco SD-WAN Software Release 17.1.x

April 2019

---

## Product Features

Below are the main product features in Viptela Software Release 17.1:

- **AS number for OMP**—You can assign an AS number to OMP itself. For vEdge routers running BGP, this overlay AS number is included in the AS path of BGP route updates. See [Configuring Unicast Overlay Routing](#) and [overlay-as](#).
- **BGP AS path propagation**—When BGP advertises routes into OMP, it can include the prefix's AS path in the advertisements. See [Configuring Unicast Overlay Routing](#) and [propagate-aspath](#).
- **CHAP and PAP authentication for PPPoE**—You can configure both CHAP and PAP authentication on a PPP interface. See [Configuring PPPoE](#) and [ppp](#).
- **Gateway sites in CloudExpress service**—You can configure CloudExpress service on sites that access the Internet through another site in the overlay network, called a gateway site. See [Using CloudExpress Service](#).
- **Match OSPF tag in localized control policy**—In the match conditions for localized control policy, you can match on an OSPF tag. See [Configuring Localized Control Policy](#) and [match](#).
- **Multiple preferred colors for application-aware routing**—In the action of application-aware routing policy, you can set multiple tunnels to use when data traffic matches an SLA class. Traffic is load-balanced across all the tunnels. See [Configuring Application-Aware Routing](#) and [action](#).
- **Policy configuration using vManage templates**—You can configure centralized control and data policy using vManage templates. See [Configuring Application-Aware Routing](#), [Configuring Centralized Control Policy](#), and [Configuring Centralized Data Policy](#).
- **REST API**—The common vManage REST API calls have been documented. See [vManage REST APIs](#).



- **Secondary IP addresses**—On vEdge routers, you can configure up to four secondary IPv4 addresses for a service-side interface. See [Configuring Network Interfaces](#) and [secondary-address](#).
- **Static NAT for service-side LANs**—In Releases 17.1.1 and later, you can perform 1:1 static NAT, to translate service-side source addresses before sending packets out to another service-side LAN connected to the same vEdge router. See [Configuring Service-Side NAT](#).
- **vEdge Cloud router lifecycle management**—You can use vManage NMS to sign certificates and generate bootstrap configurations for vEdge Cloud routers, and to decommission them. See [Install Signed Certificates on vEdge Cloud Routers](#).
- **vManage client user session idle timeout**—You can configure how long a vManage client session is inactive before a user is logged out. See [Settings](#).
- **vManage cluster statistics database memory usage**—You can display the actual memory usage for the vManage cluster's statistics database. See [Cluster Management](#).
- **vManage user interface refresh**—The vManage GUI has been refreshed. Changes include an updated color palette, new icons, updated typography, and the addition of animation and interactive screens.

## Command Changes

### New and Modified Configuration Commands

Command	Hierarchy	New	Modified	Comments
<a href="#">action</a>	policy data-policy vpn-list sequence		X	Add <b>set next-hop</b> option.
<a href="#">admin-tech-on-failure</a>	system	X		
<a href="#">authentication</a>	vpn 0 interface ppp		X	Support for configuring both CHAP and PAP authentication.
<a href="#">block-non-source-ip</a>	vpn interface	X		On vEdge routers. In Releases 17.1.1 and later.
<a href="#">default-information originate</a>	vpn router ospf		X	Remove default value for <b>originate metric</b> .
<a href="#">keepalive</a>	vpn interface		X	GRE interfaces behind a NAT send keepalive messages.



Command	Hierarchy	New	Modified	Comments
<a href="#">logs</a>	system aaa	X		
<a href="#">match ospf-tag</a>	policy route-policy sequence	X		On vEdge routers.
<a href="#">node-type</a>	vpn cloudexpress		X	Support for gateway sites. Note: To ensure that CloudExpress service is set up properly, configure it in vManage NMS, not using the CLI.
<a href="#">overlay-as</a>	omp	X		On vEdge routers.
<a href="#">propagate-aspath</a>	vpn router bgp	X		On vEdge routers.
<a href="#">qos-map</a>	policy		X	Can no longer configure on VLAN interfaces.
<a href="#">route-consistency-check</a>	system	X		On vEdge routers.
<a href="#">secondary-address</a>	vpn interface	X		On vEdge routers.
<a href="#">shaping-rate</a>	vpn interface		X	Can no longer configure on VLAN interfaces.
<a href="#">sp-organization-name</a>	system	X		
<a href="#">user</a>	system aaa		X	Increase allowable username to 128 characters.
<a href="#">usergroup</a>	system aaa		X	Increase allowable group name to 128 characters.

## New and Modified Operational Commands

Command	New	Modified	Comments
<a href="#">clear cloudexpress computations</a>		X	
<a href="#">clear cloudinit data</a>	X		On vEdge Cloud routers.
<a href="#">clear wlan radius-stats</a>	X		On vEdge routers.
<a href="#">request admin-tech</a>		X	
<a href="#">show app dpi applications</a>		X	Modify output fields. In Releases 17.1.2 and later.
<a href="#">show bfd summary</a>		X	Display configured app-route poll interval.
<a href="#">show cloudexpress</a>	X		Display loss and latency on each gateway exit for applications configured with



Command	New	Modified	Comments
<a href="#">gateway-exits</a>			CloudExpress service.
<a href="#">show hardware environment</a>		X	Display status of hardware router LEDs.
<a href="#">show interface detail</a>		X	Rename tx-tail-drops field to tx-tail/red-drops.
<a href="#">show interface port-stats</a>		X	Rename tx-tail-drops field to tx-tail/red-drops.
<a href="#">show log</a>	X		
<a href="#">show omp cloudexpress</a>	X		Display OMP routes for applications configured with CloudExpress service.
<a href="#">show policy data-policy-filter</a>		X	Add OOS Bytes field to command output.
<a href="#">show system status</a>		X	Add CPU-reported reboot field.
<a href="#">show wlan radius</a>	X		On vEdge routers.
<a href="#">tools iperf</a>	X		
<a href="#">tools minicom</a>	X		On vEdge 1000 and vEdge 2000 routers.

## Upgrade to Release 17.1

For details on upgrading the Viptela software, see [Software Installation and Upgrade](#).

Note: It is recommended that all Viptela devices run the same software version. If this is not possible, ensure that the vManage software version is not lower than that of the other controllers and is not lower than that of the vEdge routers. That is, the vManage server software must be at least the same as the highest software version running on the controllers and the routers; it can also be higher. Also ensure that the vBond and vSmart software version is not lower than that of the vEdge routers. That is, the vBond and vSmart software must be at least the same as the highest software version running on the routers; it can also be higher.

If the vEdge router configuration includes commands that configure logging, NTP, RADIUS, or TACACS servers, when you upgrade from any Release 16.2 or Release 16.3 to Release 17.1.0, the upgrade might fail, and the router rolls its software version back to the earlier release. This situation occurs when you have configured the source interface to use to reach the logging, NTP, RADIUS, or TACACS server, and the source interface is not in VPN 0. To address this issue, before you upgrade to Release 17.1.0, configure the VPN to which the source interface connects, which is the same VPN in which the server is located, as described in the [logging server](#), [ntp](#), [radius](#), and [tacacs](#) articles. This issue has been resolved in Release 17.1.1, and you do not need to modify the router's configuration before upgrading to Release 17.1.1.

For vEdge 100 series routers, starting with Release 17.1.0, you must include VPN 512 in the configuration. Before you upgrade, ensure that the configuration includes this VPN. To add VPN 512 to an existing device configuration template, in the vManage Configuration ► Templates ► Device screen, select the device template, click the More Actions icon to



the right of the row, and click Edit. In the Transport & Management VPN section, set VPN 512 to Factory\_Default\_vEdge\_VPN\_512\_Template, and click Update. To add VPN 512 from the CLI, simply add the **vpn 512** configuration command. Configuring an interface in this VPN is optional.

To upgrade to Release 17.1:

1. In vManage NMS, select the Maintenance ► Software Upgrade screen.
2. Upgrade the controller devices to Release 17.1 in the following order:
  1. First, upgrade the vManage NMSs in the overlay network.
  2. Then, upgrade the vBond orchestrators.
  3. Next, upgrade the vSmart controllers.
3. Select the Monitor ► Network screen.
4. Select the devices you just upgraded, click the Control Connections tab, and verify that control connections have been established.
5. Select the Maintenance ► Software Upgrade screen, and upgrade the vEdge routers.

**Note:** After you upgrade software on a vManage NMS to any major release, you can never downgrade it to a previous major release. For example, if you upgrade the vManage NMS to Release 17.1, you can never downgrade it to Release 16.2 or to any earlier software release.

The major release number consists of the first two numbers in the software release number. For the Viptela software, 17.1 and 16.2 are examples of major releases. Releases 17.1.0 and 16.2.0 denote the initial releases, and Releases 17.1.5 and 16.2.1 are maintenance releases.

---

## Upgrade from Release 16.2 and Earlier Software Releases

Because of software changes in Release 16.3, you must modify the router configuration as follows before you upgrade from Release 16.2 or earlier to Release 17.1:

- You can no longer configure RED drops on low-latency queuing (LLQ; queue 0). That is, if you include the **policy qos-scheduler scheduling llq** command in the configuration, you cannot configure **drops red-drop** in the same QoS scheduler. If your vEdge router has this configuration, remove it before upgrading to Release 17.1. If you do not remove the RED drop configuration, the configuration process (confd) will fail after you perform the software upgrade, and the Viptela devices will roll back to their previous configuration.
- For vEdge 2000 routers, you can no longer configure interfaces that are not present in the router. That is, the interface names in the configuration must match the type of PIM installed in the router. For example, if the PIM module in slot 1 is a 10-Gigabit Ethernet PIM, the configuration must refer to the proper interface name, for example, **10ge1/0**, and not **ge1/0**. If the interface name does not match the PIM type, the software upgrade will fail. Before you upgrade from Release 16.2 or earlier to Release 17.1, ensure that the interface names in the router configurations are correct.



---

## Caveats

---

### Hardware Caveats

The following are known behaviors of the Viptela hardware

- On vEdge 1000 routers, support for USB controllers is disabled by default. To attach an LTE USB dongle to a vEdge 1000 router, first attach the dongle, and then enable support for USB controllers on the vEdge router, by adding the [system usb-controller](#) command to the configuration. When you enter this command in the configuration, the router immediately reboots. Then, when the router comes back up, continue with the router configuration. Also for vEdge 1000 routers, if you plug in an LTE USB dongle after you have enabled the USB controller, or if you hot swap an LTE USB dongle after you have enabled the USB controller, you must reboot the router in order for the USB dongle to be recognized. For information about enabling the USB controller, see [USB Dongle for Cellular Connection](#).
- For vEdge 2000 routers, if you change the PIM type from a 1-Gigabit Ethernet to a 10-Gigabit Ethernet PIM, or vice versa, possibly as part of an RMA process, follow these steps:
  1. Delete the configuration for the old PIM (the PIM you are returning as part of the RMA process).
  2. Remove the old PIM, and return it as part of the RMA process.
  3. Insert the new PIM (the PIM you received as part of the RMA process).
  4. Reboot the vEdge 2000 router.
  5. Configure the interfaces for the new PIM.

---

### Software Caveats

The following are known behaviors of the Viptela software:

#### Cellular Interfaces

- The vEdge 100wm router United States certification allows operation only on non-DFS channels.
- When you are configuring primary and last-resort cellular interfaces with high control hello interval and tolerance values, note the following caveats:
  1. When you configure two interfaces, one as the primary interface and the other as the last-resort interface, and when you configure a high control hello interval or tolerance values on the last-resort interface (using the [hello-interval](#) and [hello-tolerance](#) commands, respectively, the OMP state indicates init-in-gr even though it shows that the control connections and BFD are both Up. This issue was resolved in Release 16.2.3. However, the following caveats exist:
    - You can configure only one interface with a high hello interval and tolerance value. This interface can be either the primary or the last-resort interface.
    - In certain cases, such as when you reboot the router or when you issue shutdown and no shutdown commands on the interfaces, the control connections might take longer than expected to establish. In this case, it is



recommended that you issue the [request port-hop](#) command for the desired color. You can also choose to wait for the vEdge router to initiate an implicit port-hop operation. The **request port-hop** command or the implicit port hop initiates the control connection on a new port. When the new connection is established, the stale entry is flushed from the vSmart controllers.

2. If the primary interface is Up, as indicated by the presence of a control connection and a BFD session, and if you configure a last-resort interface with higher values of hello interval and tolerance than the primary interface, if you issue a **shutdown** command, followed by a **no shutdown** command on the last-resort interface, the last-resort interface comes up and continuously tries to establish control connections. Several minutes can elapse before the operational status of the last-resort interfaces changes to Down. If this situation occurs, it is recommended that you issue a **request port-hop** command for the desired color.

3. If you have configured a primary interface and a last-resort interface that has higher hello interval and tolerance values than the primary interface, and if the last-resort interface has control connections to two vSmart controllers, if you issue a **shutdown** command, followed by a **no shutdown** command on the last-resort interface, a control connection comes up within a reasonable amount of time with only one of the vSmart controllers. The control connection with the second vSmart controller might not come up until the timer value configured in the hello tolerance has passed. If this situation occurs, it is recommended that you issue a **request port-hop** command for the desired color.

- For cellular interface profile, the profile number can be 0 through 15. Profile number 16 is reserved, and you cannot modify it.

## Configuration

- When you issue the **request reset configuration** command on a vEdge Cloud router, a vManage NMS, or a vSmart controller, the software pointer to the device's certificate might be cleared even though the certificate itself is not deleted. When the device reboots and comes back up, installation of a new certificate fails, because the certificate is already present. To recover from this situation, issue the **request software reset** command.

## Control and BFD Connections

- When a vBond orchestrator, vManage NMS, or vSmart controller goes down for any reason and the vEdge routers remain up, when the controller device comes back up, the connection between it and the vEdge router might shut down and restart, and in some cases the BFD sessions on the vEdge router might shut down and restart. This behavior occurs because of port hopping: When one device loses its control connection to another device, it port hops to another port in an attempt to reestablish the connection. For more information, see the [Firewall Ports for Viptela Deployments](#) article. Two examples illustrate when this might occur:
  - When a vBond orchestrator goes down for any reason, the vManage NMS might take down all connections to the vEdge routers. The sequence of events that occurs is as follows: When the vBond orchestrator crashes, the vManage NMS might lose or close all its control connections. The vManage NMS then port hops, to try to establish



connections to the vSmart controllers on a different port. This port hopping on the vManage NMS shuts down and then restarts all its control connections, including those to the vEdge routers.

- All control sessions on all vSmart controllers go down, and BFD sessions on the vEdge routers remain up. When any one of the vSmart controllers comes back up, the BFD sessions on the routers go down and then come back up because the vEdge routers have already port hopped to a different port in an attempt to reconnect to the vSmart controllers.
- When a vEdge router running Release 16.2 or later is behind a symmetric NAT device, it can establish BFD sessions with remote vEdge routers only if the remote routers are running Release 16.2 or later. These routers cannot establish BFD sessions with a remote vEdge router that is running a software release earlier than Release 16.2.0.
- When you add or remove an IPv4 address on a tunnel interface (TLOC) that already has an IPv6 address, or when you add or remove an IPv6 address on a TLOC that already has an IPv4 address, the control and data plane connections for that interface go down and then come back up.
- When you add or remove an IPv4 address on a tunnel interface (TLOC) that already has an IPv6 address, or when you add or remove an IPv6 address on a TLOC that already has an IPv4 address, the control and data plane connections for that interface go down and then come back up.
- Release 16.3 introduces a feature that allows you to configure the preferred tunnel interface to use to exchange traffic with the vManage NMS. In the vManage NMS, you configure this on cellular, Ethernet, and PPP Interface feature templates, in the vManage Connection Preference field under Tunnel Interface. In the CLI, you configure this with the [vmanage-connection-preference](#) command. The preference value can be from 0 through 8, with a lower number being more preferred. The default value is 5. If you set the preference value to 0, that tunnel interface is never used to exchange traffic with the vManage NMS, and it is never able to send or receive any overlay network control traffic. With this configuration option, there is one situation in which you can accidentally configure a device such that it loses all its control connections to all Viptela controller devices (the vManage NMSs and the vSmart controllers). If you create feature templates and then consolidate them into a device template for the first time, the NMS software checks whether each device has at least one tunnel interface. If not, a software error is displayed. However, when a device template is already attached to a device, if you modify one of its feature templates such that the connection preference on all tunnel interfaces is 0, when you update the device with the changes, no software check is performed, because only the configuration changes are pushed to the device, not the entire device template. As a result, these devices lose all their control connections. To avoid this issue, ensure that the vManage connection preference on at least one tunnel interface is set either to the default or to a non-0 preference value.

## Interfaces

- On virtual interfaces, such as IRB, loopback, and system interfaces, the duplex and speed attributes do not apply, and you cannot configure these properties on the interfaces.
- When a vEdge router has two or more NAT interfaces, and hence two or more DIA connections to the internet, by default, data traffic is forwarding on the NAT interfaces using ECMP. To direct data traffic to a specific DIA interface,





configure a centralized data policy on the vSmart controller that sets two actions—**nat** and **local-tloc** color. In the **local-tloc** color action, specify the color of the TLOC that connects to the desired DIA connection.

## IPv6

- You can configure IPv6 only on physical interfaces (**ge** and **eth** interfaces), loopback interfaces (**loopback0**, **loopback1**, and so on), and on subinterfaces (such as **ge0/1.1**).
- For IPv6 WAN interfaces in VPN 0, you cannot configure more than two TLOCs on the vEdge router. If you configure more than two, control connections between the router and the Viptela controllers might not come up.
- IPv6 transport is supported over IPsec encapsulation. GRE encapsulation is not supported.
- You cannot configure NAT and TLOC extensions on IPv6 interfaces.

## IRB

- On integrated routing and bridging (IRB) interfaces, you cannot configure [autonegotiation](#).

## NAT

- When you reboot a vSmart controller, the BFD sessions for all symmetric NAT devices go down and come back up. This is expected behavior.

## Security

- In Releases 17.1 and earlier, you cannot disable the SSH HMAC-MD5 algorithm and other weaker algorithms.

## SNMP

- When you configure an SNMP trap target address, you must use an IPv4 address.
- The Viptela interface MIB supports both 32-bit and 64-bit counters, and by default sends 64-bit counters. If you are using an SNMP monitoring tool that does not recognize 64-bit counters, configure it to read 32-bit MIB counters.
- On a vEdge router, if you perform an snmpwalk getnext request for an OID for which there is no information, the response that is returned is the next available instance of that OID. This is the expected behavior.

## System

- In Releases 17.1 and earlier, you cannot disable the SSH HMAC, MD5, and other weaker algorithms.



- In Releases 17.1 and later, you cannot change the personality of a vEdge router to be a vBond controller, or vice versa, after certificates have been installed on the device.

## Virtual Machines

- For a vEdge Cloud VM instance on the KVM hypervisor, for Viptela Releases 16.2.2 and later, it is recommended that you use virtio interfaces. For software versions earlier than Release 16.2.2, if you are using the Ubuntu 14.04 or 16.04 LTS operating system, you can use IDE, virtio, or virtio-scsi interfaces.

## vManage NMS

- On a Viptela device that is being managed by a vManage NMS system, if you edit the device's configuration from the CLI, when you issue the **commit** command, you are prompted to confirm the commit operation. For example:  

```
vEdge(config-banner)# commit
```

The following warnings were generated:

```
'system is-vmanaged': This device is being managed by the vManage. Any configuration changes to this device will be overwritten by the vManage.
```

Proceed? [yes,no]

You must enter either **yes** or **no** in response to this prompt.

During the period of time between when you type commit and when you type either **yes** or **no**, the device's configuration database is locked. When the configuration database on a device is locked, the vManage NMS is not able to push a configuration to the device, and from the vManage NMS, you are not able to switch the device to CLI mode.
- The members of a vManage cluster rely on timestamps to synchronize data and to track device uptime. For this time-dependent data to remain accurate, do not change the clock time on any one of the vManage servers of the cluster after you create the cluster.
- When you use the vManage Maintenance ► Software Upgrade screen to set the default software version for a network device, that device must be running Release 16.1 or later at the time you set the default software version. If the network device is running Release 15.4 or earlier, use the CLI **request software set-default** command to set the default software version for that device.
- When you are using a vManage cluster, when you are bringing up new vManage NMS in the cluster, use an existing vManage NMS to install the certificate on the new vManage NMS.
- In vManage feature configuration templates, for the passwords listed below, you cannot enter a cleartext password that starts with \$4 or \$8. You can, however, use such passwords when you are configuring from the CLI.
  - Neighbor password, in the BGP feature configuration template



- User password, in the Cellular Profile feature configuration template
- Authentication type password and privacy type password, in the SNMP feature configuration template
- RADIUS secret key and TACACS+ secret key, in the System feature configuration template
- IEEE 802.1X secret key, in the VPN Interface Ethernet feature configuration template
- IPsec IKE authentication preshared key, in the VPN Interface IPsec feature configuration template
- CHAP and PAP passwords, in the VPN Interface PPP Ethernet feature configuration template
- Wireless LAN WPA key, in the WiFi SSID feature configuration template

---

## Outstanding Issues

The following are outstanding issues in Viptela Software Release 17.1. The number following each issue is the bug number in the Viptela bug-tracking database.

### AAA

- The Viptela software does not send a TACACS vendor-specific "service argument" field. [VIP-25629]

### Cellular Interfaces

- If you configure IPv6 on a cellular interface, the control connections might go down and come back up continuously. [VIP-21970]
- On a vEdge 100m-NA router, when you configure profile 1 for a wireless WAN, you might see the error "Aborted: 'vpn 0 interface cellular0 profile': Invalid profile 1 : APN missing". [VIP-31721]
- You cannot configure profile 16 in the [interface cellular0 profile](#) command. [No bug number.]

### Configuration and Command-Line Interface

- An IRB interface might remain up even if all interfaces in that bridge are in the link-down state. [VIP-23307]
- When you issue the **show vrrp interfaces** command from the vEdge router's CLI, the CLI might not recognize the command and might show a "syntax error: unknown argument" error message. [VIP-23918]



- If a physical interface is part of a bridge, you cannot adjust the MTU on the interface. As a result, the 802.1x interface's MTU has to be lowered to 1496. If the interface needs to also run OSPF, this MTU size can cause an MTU mismatch with other interfaces that have an MTU of 1500. [VIP-26759]
- On cellular interfaces, you might not be able to modify the maximum segment size (MSS) of TCP SYN packets. [VIP-28033]
- The **traceroute** command might not work if you specify an IPv6 address for the host. [VIP-30833]
- When two routes exist to the same neighbor, if you specify a single IP address in the **show ip routes** command, the command might return only one of the routes, but if you specify an IPv4 prefix and prefix length, the command returns both routes. [VIP-32736]

## DHCP

- When a vEdge 1000 router is acting as a DHCP server, the performance of the router might diminish. [VIP-34874]

## Forwarding

- For IEEE 802.1x, you cannot configure a RADIUS server for MAC authentication bypass (MAB). [VIP-18492]
- In application-aware routing policy, the `salesforce_chatter`, `oracle_rac`, and `google_photos` applications might not be classified properly. [VIP-21866]
- If you add a large number of bridge tagged interfaces in a single commit operation, these interfaces are listed in the output of the **show interface** command for a few minutes, even though this command is not supposed to list bridge tagged interfaces. After this time has elapsed, the bridge interfaces no longer show up in the command the output, which is the expected behavior. [VIP-25715]
- When you configure a weight on a TLOC that is also being used as a split tunnel, the weight is not used for weighted ECMP across the NATs. [VIP-27534]
- When you switch data traffic from one tunnel to another (for example, from a biz-ethernet to an lte tunnel), a small amount of traffic might be lost. [VIP-27992]
- For a source and destination NAT, return traffic might not be able to reach the VPN that originates the session.



[VIP-31299]

- You might not be able to perform a traceroute operation on an IRB interface. [VIP-31563]
- When you configure **policy cloud-qos** on a vEdge Cloud router, a TLOC from the remote site might go down and then come back up when multiple traffic flows are present on the TLOC. [VIP-32369]
- When you configure inbound and outbound port mirroring on the same interface, traffic might be mirrored only in one direction. [VIP-33247]
- When you have a localized data policy (ACL) that mirrors traffic on an interface in both directions, if you change the IP address of the interface and the mirror destination but do not remove the ACL, the outbound mirroring continues to work but the inbound mirroring stops working. If you then remove and reapply, the ACL, the mirroring again works in both directions. [VIP-33275]
- If you enable TCP optimization on a vEdge 1000 router, the router might drop ARP responses. [VIP-33507]
- When the output of the **show ipsec outbound-connections** command shows that tunnel MTU is 1441 bytes, a router fragments packets with the size (iplen) of 1438 bytes, but 1437-byte are not fragmented. There seems to be a 4-byte gap between tunnel MTU and the size at which the router actually starts fragmenting a packet. Also the TCP MSS seems to be 40 bytes smaller than expected for IPv4 packets and 60 bytes smaller for IPv6. [VIP-33527]
- When multiple NAT interfaces are present in VPN 0, port forwarding might not work. [VIP-34086]
- If you disable deep packet inspection (DPI) on a vEdge router, traffic directed towards queue 0 (LLQ) might become bursty or might be dropped. [VIP-34211]
- When you configure a cellular interface as a last-resort interface, the cellular interface might remain up at all times. [VIP-34495]
- Traffic might be blackholed because of stale BFD sessions. This happens in a scenario when there are two vEdge routers at a site, both configured with TLOC extension between them, and the circuit that they are connected to goes down. One router clears all its BFD sessions, but the second one does not, so all traffic is sent to the uncleared BFD sessions and is blackholed. [VIP-35113]

## Interfaces

- When a vEdge VRRP master is connected to a Cisco switch, the switch might report error messages indicating that the



source MAC address is invalid. [VIP-28922]

- Traffic flow on IPsec tunnels might be interrupted when you configure only tunnel interface parameters, such as MTU and dead-peer detection. [VIP-31426]
- When a VRRP backup vEdge router that has been promoted to a master again becomes a backup, other devices continue to point to the MAC address for the backup router, and traffic is blackholed until ARP cache on the other devices expires and is updated with the correct MAC address of master vEdge, a process that typically takes a few minutes. [VIP-33722]
- When a BGP route has the same administrative distance as a static route, the BGP route might not be installed in the forwarding table (FIB). [VIP-34428]
- On a vEdge router that has two TLOCs, one on a loopback interface and the second on a physical interface, when the physical interface goes down, the loopback interface might not be able to forward traffic. [VIP-34646]

## Policy

- QoS shaping rates might be inaccurate for rates less than 2 Mbps. [VIP-3860]
- If you have configured a policer for the LLQ, the output of the **show interface queue** command shows all queue statistics except those for the LLQ. As a workaround, use the **show policer** command to display the LLQ statistics. [VIP-4529]
- A centralized policy that is pushed from the vSmart controller to the vEdge routers might not be applied on the routers. [VIP-27046]
- On vEdge routers, the **show policy access-list-counters** command might not display any values in the Bytes column. [VIP-28890]
- The vSmart controller might not push a policy to the vEdge routers. [VIP-33016]
- After you change a policy on the vSmart controller, the OMP process (ompd) process might fail and the vSmart controller might crash. [VIP-34098]
- When the vSmart controller pushes a policy to a vEdge router and the push fails, no alarm or trap records the failure.



[VIP-34131]

- When you enable app-visibility and flow-visibility on a vEdge router, the vManage Dashboard might not display any any cflowd or DPI flows. [VIP-34406]

## Routing Protocols

- When the OSPF external distance is set to 254, an IP prefix learned first from OMP and then from OSPF as an type E2 route, the route might be redistributed into OMP. [VIP-20542]
- When you are upgrading vEdge routers to Release 16.2.12, the BGP process (bgpd) during the reboot process, when the router is shutting down. [VIP-29523]
- On a vEdge 1000 router, the OSPF process (ospfd) might fail and cause the router to crash. [VIP-30239]
- When you enable VRRP, multicast routing might not work. [VIP-34431]
- In a topology with one multicast replicator in the cloud and no PIM router, when you configure IGMP-Join-group on the vEdge router, multicast routing might not work. [VIP-34449]
- A vEdge router might not be able to establish OSPF point-to-point interfaces with a Juniper EX4200 device. [VIP-34936]
- Routes might be installed in the routing table with the incorrect color. [VIP-35088]

## Security

- If an IPv6 address for the IPsec tunnel source interface, the IPsec tunnel does not come up. [VIP-29912]

## SNMP

- When traffic exceeds 85% of the bandwidth configured on a transport interface, SNMP traps might not get triggered. [VIP-33435]



## System

- vBond orchestrators might report a large number of control-connection-auth-fail events. [VIP-22976]
- On a vEdge 100b router, upgrading from Release 15.4.1 to Release 16.2.2 might fail silently, because the uboot file is incorrect. As a workaround, copy the correct uboot file to the router before performing the upgrade. [VIP-23083]
- When a task stops and a vEdge router reboots, the router might no longer reboot. This problem occurs after the router reboots three times within 20 minutes, five times within 60 minutes, or seven times within the last 24 hours. However, the control plane on the router remains up, so traffic continues to be sent to the node. [VIP-23106]
- When you shut down a subinterface, the output of the **show interface** command might show that the interface is administratively down but operationally up. [VIP-23829]
- In an overlay network with three vSmart controllers, if a controller group list configured on a 100 vEdge router contains two vSmart controllers, the maximum number of controllers that the router can connect to is set to two, and the maximum number of OMP sessions on the router is set to two, 50 routers connect to each of two vSmart controllers. If you bring these two controllers down, all 100 connections then move to the third vSmart controller. However, if you then bring up one of the other vSmart controllers, 50 connections move to that controller, but the third controller might still have 100 connections. [VIP-27955]
- When the configuration process (confd) on a vEdge router crashes, the router might not reboot as expected. Instead, it remains at the Linux Bash shell. [VIP-28441]
- The vManage server might not process events received from vEdge routers. [VIP-28673]
- On a cellular vEdge router that has a PPP interface and a tunnel interface that is configured as the circuit of last resort, when you upgrade the software on the router and it reboots, the push of the configuration template from the vManage NMS to the router might fail because the configuration process on the router takes longer than 30 seconds to commit the pushed configuration. [VIP-28797]
- When a certificate for controllers is about to expire, no syslog message is generated. [VIP-28960]
- A vEdge router might try connect to a vSmart controller that has been invalidated and deleted from the vManage NMS. [VIP-30002]
- When a last-resort interface has been initiated and connections on that interface are being brought up, the value of the last-resort hold-down timer might be shown incorrectly in syslog files. [VIP-30423]





- When NAT is configured between in a service-side VPN, a ping operation between a vEdge router in that service VPN and another vEdge router reachable through the transport network might be successful even though it should be blocked because of the NAT. [VIP-31078]
- When a vEdge router is unable to reach one of the controllers in a controller, it might not try to reach other controllers in the same group. [VIP-31882]
- A vEdge router might choose to establish its control connection to the vManage NMS using an interface on which a tunnel interface is not configured even though an interface with a tunnel interface is operational. [VIP-32011]
- Pushing a device configuration template to a vEdge router might fail because of a bridge configuration validation failure. This issue occurs when a bridge with VLAN and interfaces is already configured on the router and the template being pushed modifies these parameters. As a workaround, copy the template, delete the entire bridge configuration, and push the template to the router. Then add the original bridge configuration to the template, and push that template to the router. [VIP-33204]
- The vdebug log file might contain no entries. [VIP-33662]
- A vSmart controller might crash and create the core file /rootfs.rw/var/crash/core.vtracker.vSmart, indicating an issue with the vtracker process, which pings the vBond orchestrator every second. [VIP-33719]

## vEdge Hardware

- When a vEdge 2000 router reboots, the reboot reason field might show only a value of 0. [VIP-23941]
- On a vEdge 100m router, after you execute the **request software reset** command, the router might reboot continuously. [VIP-24149]
- Hardware vEdge routers might categorize error packets incorrectly. [VIP-26039]
- On a vEdge 100 router, when you enable or disable debugging, a Forwarding Process (fp) core file might be created. [VIP-26965]
- A vEdge 2000 router physical interface might drop packets larger than 1480 bytes that are sent on loopback interfaces. [VIP-27216]



- On a vEdge100m router, the output of the **show interface** command might show the same interface in two different VPNs. [VIP-29069]
- On a vEdge 100b, when you change the IP address on an interface, that IP address might not be detached from the interface. [VIP-35047]

## vManage NMS

- In the vManage Monitor ► Network screen, the detailed device information might be difficult to read because of how it is formatted. [VIP-11612]
- On a vManage NMS running Release 16.2.2.1, you might not be able to push a template to a vSmart controller that includes a TLOC list in its configuration. [VIP-19483]
- If you try to configure a vEdge router using vManage configuration templates, you might see errors related to lock-denied problems. As a workaround, reboot the router. [VIP-23826]
- On vManage NMS, when you display interface queue statistics in real time, statistics for only one of the eight possible queues might be displayed. [VIP-23898]
- If you use the CLI to modify the organization name, this change might not be reflected on the vManage screens. [VIP-24343]
- When the majority of vManage cluster members are down, you can make changes to the device configuration templates on one of the cluster members that is up, and you can then push these changes when the cluster members come back up. This might lead to a situation in which the configuration templates on the vManage NMSs in the cluster are out of sync. [VIP-26016]
- The vManage dashboard does not automatically refresh the state of the members of the vManage cluster even when their state changes. [VIP-26017]
- A vManage NMS might not be able to synchronize its configuration with a vSmart controller. [VIP-26270]
- When you upgrade a software image on a vEdge router and then, in a separate action, activate the image, the new software image is not activated. As a workaround, when you upgrade the software image, check the Activate option. [VIP-27275]
- A vManage serve might continue to attempt to fetch certificates even though all certificates are installed. [VIP-27416]



- When you upgrade a vManage cluster from Release 16.3.2 to Release 17.1.0, the configuration database might become incorrect. [VIP-27951]
- The vManage server might not process events received from vEdge routers. [VIP-28312]
- In the vManage Monitor ► Network ► Troubleshooting Ping pane, when you enter an IPv6 destination address, the ping operation might fail. [VIP-30720]
- When you use the vManage NMS and the CLI **show system status** command, the reboot reason is incorrect; it is shown as unknown. Looking in the /var/log/tmplog/vdebug logs shows that the system reboot happened because of a user-initiated upgrade to Release 17.1.3. [VIP-31222]
- In certain situations, such as when the control plane has gone down and come back up or when you specify an invalid destination IPv6 address, the Simulated Flows option in vManage Monitor ► Network ► Troubleshooting might not work. [VIP-31576]
- In the vManage AAA feature template, you might not be able to enter the RADIUS secret key even though you can enter that same key in the CLI. [VIP-31856]
- When you push a policy that contains an error to the vSmart controller, the error message might not correctly indicate the cause of the error. [VIP-32253]
- You might not be able to push a configuration template to a vEdge router. [VIP-32277]
- When you are using a vManage cluster, pushing policies to vSmart controllers might time out. [VIP-32630]
- From a vManage server, you might not be able to SSH into a vEdge router that is in staging mode. [VIP-33119]
- Pushing a configuration template to a vEdge router might time out if the configuration has only one interface and that interface is configured as a last-resort interface. [VIP-33157]
- In the vManage Monitor ► Network ► Real Time screen, the output of the Interface Queue Stats command might show information for queue 0 only, showing no information about queues 1 through 7. [VIP-33508]
- When you copy the configuration database from the primary vManage NMS to bring up a secondary vManage NMS, the certificates for vEdge Cloud routers are not included, and the control plane and data plane for these routers do not



come up. [VIP-34085]

- When the vManage NMS experiences a kernel panic and reboots, the /var/crash/crash.dump file might be deleted. [VIP-34248]
- In a vManage cluster, the vManage server might run slowly, and the vmanage-server.log file might contain "Reached maximum number of concurrent connections" exception messages. [VIP-34594]
- You might not be able to push configuration templates to vEdge routers. [VIP-34886]

---

## Fixed Issues

---

### Issues Fixed in Release 17.1.5

The following issues have been fixed in Viptela Software Release 17.1.5. The number following each issue is the bug number in the Viptela bug-tracking database.

#### Forwarding

- On a vEdge Cloud router, traffic shaping through QoS might not work properly. [VIP-26557: This issue has been resolved.]

#### Interfaces

- When you configure VRRP on an interface that is operationally down, that interface might become the VRRP master. [VIP-33505: This issue has been resolved.]

#### Policy

- When vEdge hardware and software routers are using application-aware routing, they may no longer honor the preferred action when all paths comply with the SLA and ECMP is used. [VIP-33650: This issue has been resolved.]

#### Security

- The DSCP in data traffic sent over a TLS connection from a server to a vSmart controller might not be set properly. [VIP-30056: This issue has been resolved.]
- When you uninstall the certificate from a vEdge Cloud router, the root certificate might also be removed. [VIP-34092: This issue has been resolved.]



## System

- On a vEdge 2000 router, BFD sessions may go down and then come back up every 10-60 minutes. As a workaround, disable cflowd and DPI. [VIP-33784: This issue has been resolved.]

## vManage NMS

- A vManage authentication failure returns the HTTP 200/OK code, with the HTML body containing "Invalid User or Password", instead of the standard HTTP 401/Unauthorized code. [VIP-32082: This issue has been resolved.]
- The vManage NMS might not display statistics for some tunnels. [VIP-34054: This issue has been resolved.]
- On the Configuration ► Templates page, you might not be able to scroll horizontally. [VIP-34205: This issue has been resolved.]
- After the ZTP process succeeds, pushing a configuration template might fail. [VIP-34288: This issue has been resolved.]

---

## Issues Fixed in Release 17.1.4

The following issues have been fixed in Viptela Software Release 17.1.4. The number following each issue is the bug number in the Viptela bug-tracking database.

### CloudExpress Service

- When an upstream router fails, the CloudExpress service might take up to 30 minutes to switch to the overlay network. [VIP-28136: This issue has been resolved.]

### Forwarding

- When the options field in the TCP packet header includes the TCP MSS and other options, if the other options precede the TCP MSS option in the field, the TCP MSS adjustment might not take effect. [VIP-31509: This issue has been resolved.]

### OMP

- Routes from a staged vEdge router might be advertised to non-staged vEdge routers. [VIP-31295: This issue has been resolved.]

## Security



- The **show tunnel statistics ipsec** command displays no information about the count of inbound decrypted packets. [VIP-20637: This issue has been resolved.]

## SNMP

- After you upgrade from Release 15.4.x to a Release 16.x release, you can no longer use a VRRP physical interface IP address for snmpget and snmpwalk operations, because the SNMP listener starts on VRRP virtual IP address instead of on the physical interface IP address. This issue has been resolved in Release 17.x releases. [VIP-29085: This issue has been resolved.]

## vEdge Routers

- A vEdge 100 router might stop processing packets, and rebooting the router might not solve the problem. [VIP-31300: This issue has been resolved.]

## vManage NMS

- With Japanese versions of the Chrome and IE browsers, the vManage NMS might not display some tables and other screen fields. [VIP-28870: This issue has been resolved.]
- The vManage NMS servers might reboot twice with the errors "'ncs failed..rebooting after ncs cdb cleanup" and "Daemon 'ncs' failed". After the reboots, the vManage servers might stop functioning. [VIP-28896: This issue has been resolved.]
- The vManage NMS should allow only one upgrade operation to be performed at a time, but sometimes it might attempt to perform two at the same time. [VIP-32084: This issue has been resolved.]

---

## Issues Fixed in Release 17.1.3

The following issues have been fixed in Viptela Software Release 17.1.3. The number following each issue is the bug number in the Viptela bug-tracking database.

### Configuration and Command-Line Interface

- If the default gateway next hop does not respond to an ICMP request, the vSmart controller does not install the next hop unless you remove the **track-default-gateway** command from the configuration. This problem is not seen on the vBond orchestrator. [VIP-25986: This issue has been resolved.]

## Forwarding



- If you configure a translation rule on a NAT pool and then issue a **show policy service-path** command, the forwarding process (fp) might crash. [VIP-24417: This issue has been resolved.]
- When all the vSmart controllers connected to a TLOC fails, the IPsec connections between the vEdge routers might go down and come back up after control connections with the vSmart controllers are re-established. This happens only when port-hop is enabled. [VIP-27324: This issue has been resolved.]
- The traffic flow on LTE interfaces might remain very high even after you adjust timers on the interfaces. [VIP-27971: This issue has been resolved.]
- On vEdge 1000 and vEdge 2000 routers, when you enable fast deep packet inspection, which is used by the CloudExpress platform, DPI performance might suffer. [VIP-28008: This issue has been resolved.]
- On hardware vEdge routers, when packet fragmentation of data traffic sent over an IPsec tunnel occurs, a buffer leak might occur. [VIP-29314: This issue has been resolved.]

## Interfaces

- When you upgrade a vEdge router to Release 17.1.0, PPP interfaces might reset often, in the process losing their IP addresses. When the interfaces come back up, they need to reconnect and authenticate via PPP, and it can take about 2 minutes to bring up the GRE tunnels. [VIP-28616: This issue has been resolved.]

## Policy

- When you configure a centralized data policy that fails when it is applied on a vSmart controller, the failure error message might not clearly describe the problem. [VIP-27898: This issue has been resolved.]

## Routing Protocols

- When you configure BGP in a dual-homed setup that has EBGP peering sessions to both BGP neighbors in the same ASN, the vEdge router might install routes only from one of the peers even though it is receiving routes from both of them. [VIP-28144: This issue has been resolved.]

## Security

- If a vSmart controller becomes unreachable, the vEdge router might still show that controller as being reachable. [VIP-22262: This issue has been resolved.]
- During a POODLE attack against a TLS connection, all the device's CPU might be used or the vdaemon process might



crash. [VIP-29376: This issue has been resolved.]

## System

- The output of the **traceroute** command on a vEdge router might be incorrect. [VIP-23072: This issue has been resolved.]
- GRE tunnels might not leverage the interface MTU to automatically adjust their TCP MSS size. [VIP-28625: This issue has been resolved.]

## vManageNMS

- In vManage NMS, you might not be able to delete a beta version of the Viptela software. [VIP-28672: This issue has been resolved.]
- After you upgrade the vManage NMS software, you might not be able to add the vAnalytics feature to an existing user group in the vManage Administration ► Manage Users ► User Groups screen. As a workaround, first create a new user group and then add the vAnalytics feature to that group. [VIP-29407: This issue has been resolved.]

## Wireless WANs

- On a cellular interface, the **show cellular radio** command might display the wrong SNR classification, showing Poor instead of Excellent. [VIP-28847: This issue has been resolved.]

---

## Issues Fixed in Release 17.1.2

Release 17.1.2 was not released.

---

## Issues Fixed in Release 17.1.1

The following issues have been fixed in Viptela Software Release 17.1.1. The number following each issue is the bug number in the Viptela bug-tracking database.

### CloudExpress Service

- You cannot direct CloudExpress traffic over a GRE tunnel. [VIP-26954: This issue has been resolved.]

## Forwarding

- If you misconfigure a vSmart controller, the Forwarding Table Management process (ftmd) on the vEdge routers might crash when the routers receive OMP updates from the controller. [VIP-22116: This issue has been resolved.]





- When a Cisco phone is connected to a PoE interface on a vEdge 100wm router and the interface receives its IP address via DHCP, when DHCP renews its lease, the interface might periodically lose its assigned IP address. [VIP-26216: This issue has been resolved.]

## Policy

- Application-aware routing policy might not take effect, causing data traffic to be directed to the incorrect tunnel interface. [VIP-26356: This issue has been resolved.]

## Security

- On a vEdge router that has two Internet-facing interfaces, one of the interfaces might not be able to establish control connections or tunnels, and on this interface, many IPsec rekeying events might be occurring. [VIP-23870: This issue has been resolved.]

## System

- If you do not use ZTP and if the Organization Name is not set on the vManage NMS and vEdge routers, control connections between the two devices might come up. [VIP-24246: This issue has been resolved.]

## vManage NMS

- The vManage NMS might stop receiving event messages from the vEdge routers in the network. [VIP-21973: This issue has been resolved.]
- The Transport Interface pane on the vManage Dashboard might show information for non-transport interfaces. [VIP-21989: This issue has been resolved.]
- vEdge routers that have been marked as invalid might be listed in vManage device inventory API calls. [VIP-22410: This issue has been resolved.]
- When an admin user edits an existing CLI template, its device type is sometimes Null, so the user is unable to select the device type. [VIP-24182: This issue has been resolved.]
- On a standalone vManage NMS, the messaging bus might stop operating and then restart, but when it restarts the vManage application server does not restart. [VIP-25149: This issue has been resolved.]
- For an overlay network with a single vManage NMS, if the vManage NMS is stopped and move to another devices, the vManage services might not restart properly. [VIP-25692: This issue has been resolved.]



- Pushing a device configuration template from one of the vManage servers in a cluster might fail, with the "Failing with Server Not Leader" partition exception error. [VIP-25771: This issue has been resolved.]
- On the vManage Monitor ► Network screen, the real-time statistics displayed might be inaccurate. [VIP-26040: This issue has been resolved.]
- You might not be able to export the bootstrap configuration from the vManage server. [VIP-26209: This issue has been resolved.]
- For a vManage server using the Firefox browser, when you go to the Monitor ► Network screen, this screen might not display unless you refresh the page. [VIP-26755: This issue has been resolved.]
- If you edit an existing centralized policy on the vManage NMS and then re-apply the policy to vEdge routers, the policy might not work. As a workaround, copy the entire policy, delete the existing policy, and then paste the copied policy before re-applying it. [VIP-27048: This issue has been resolved.]
- If CSR generation on a vManage NMS fails, the log files provide no details about the reason for the failure. [VIP-27379: This issue has been resolved.]
- The vManage Configuration ► Certificates screen might not display the certificate icon for a router in the Valid state. [VIP-27380: This issue has been resolved.]
- You might not be able to save the vManage WiFi SSID feature configuration template. [VIP-27895: This issue has been resolved.]

---

## Issues Fixed in Release 17.1.0

The following issues have been fixed in Viptela Software Release 17.1.0. The number following each issue is the bug number in the Viptela bug-tracking database.

### Cellular Interfaces

- On a vEdge router, the **show cellular status** command does not reflect the current operational status of the modem. [VIP-13808: This issue has been resolved.]
- When you try to apply a device template that includes a cellular interface configuration to a device that does not already have a cellular profile defined, you might not be able to apply the template. [VIP-19372: This issue has been resolved.]
- If you have configured a profile for a cellular interface, when the interface is active, you cannot delete the profile. [VIP-20327: This issue has been resolved.]



- When you configure a cellular interface, you must configure a profile (with the **cellular cellular0 profile** command) before you can associate the profile with an interface (with the **interface cellular profile** command). [VIP-21569: This issue has been resolved.]

## Configuration and Command-Line Interface

- On a vEdge router, when you issue the "**show dhcp server bindings** command that contains a syntax error, such as **show dhcp server bindings h**, the CLI might display the message, "Error: application communicate failure", and the router might crash. [VIP-25467: This issue has been resolved.]

## Forwarding

- The flow collector on a vEdge router might fail to collect cflowd packet statistics. [VIP-11088: This issue has been resolved.]
- When there is conflicting decision regarding the path that a particular application needs to take, the CloudExpress configuration is overriding the decision made by data or application-aware routing policy. This is opposite of the expected behavior. [VIP-21863: This issue has been resolved.]
- When you are sending cflowd data to an external collector and the collector becomes unreachable, non-reachability messages are generated and written to the logs for every packet sent out for every flow. This process generates a huge number of log messages and might overwrite all other legitimate log messages. [VIP-22604: This issue has been resolved.]
- On vEdge routers, the routing table might not validate that a route is present in the forwarding table. [VIP-23091: This issue has been resolved.]
- On vEdge routers, the routing table might not validate that a route is present in the forwarding table. [VIP-23411: This issue has been resolved.]
- If a NAT-translated rule on a vEdge router receives an ICMP "error/unsupported proto" message, the forwarding process (fp) on the router might crash. [VIP-24220: This issue has been resolved.]
- When you disable BFD-based PMTU discovery on a vEdge router, the router might continue to receive PMTU packets. [VIP-25343: This issue has been resolved.]
- On a vEdge 2000 router, when there are a large number of cflowd flows, the Forwarding Table Management process (ftmd) might have high CPU utilization when writing flow data to the /tmp/xml files, and BFD sessions might go down and then come back up. [VIP-26346: This issue has been resolved.]



## System

- The **show bfd sessions** command might indicate that BFD sessions are down when they are actually up. [VIP-12058: This issue has been resolved.]
- Copying cflowd statistics from vEdge routers to the vManage NMS might take a long time, especially when there are large amounts of statistics. [VIP-23519: This issue has been resolved.]

## vEdge Hardware

- When you are using a fiber SFP, you are able to configure **no autonegotiate** on the interface even though autonegotiation for 1000BaseT fiber is always enforced at the register level. [VIP-24490: This issue has been resolved.]
- On a vEdge 2000 router, if you enable DPI and if the flow creation and deletion rate is very high, the forwarding table manager (FTM) might use large amounts of memory while performing flow visibility, and eventually, an out-of-memory condition might occur. [VIP-25005: This issue has been resolved.]

## vManage NMS

- vManage NMS might fail to make a Netconf connection to devices when attempting to collect statistics. [VIP-11087: This issue has been resolved.]
- When you upgrade from Release 16.1.2.3 to Release 16.2.2.1, the vManage NMS might lose all its stored network configurations and might then have to re-create them after the upgrade completes. [VIP-19454: This issue has been resolved.]
- When you attempt to push a software upgrade from the vManage NMS to a few vEdge routers, the vManage Dashboard might display the message "This server is not the leader for that partition exception", and the vManage server might not receive any event notifications. [VIP-21062: This issue has been resolved.]
- In the vManage Monitor ► Device Real-Time Data drop-down, the interface queue command is missing. [VIP-21796: This issue has been resolved.]
- The vManage server log files might contain a large number of duplicate BFD record messages. [VIP-22098: This issue has been resolved.]
- The vManage NMS might create alarms that are duplicates of previously created alarms. [VIP-22936: This issue has been resolved.]



- After you upgrade the vManage NMSs and vSmart controllers to Release 16.2.7, the vSmart controllers might not have the correct policy. As a workaround, deactivate and then reactivate the policy. [VIP-22980: This issue has been resolved.]
- The vManage dashboard might show the incorrect certificate expiration even after certificate has been renewed. [VIP-23779: This issue has been resolved.]
- When you push a configuration template from the vManage server, the headers on the template screen might become misaligned. [VIP-23881: This issue has been resolved.]
- In the VPN-Interface-NAT-Pool feature configuration template, you cannot create a variable for static NAT. [VIP-25563: This issue has been resolved.]
- You might not be able to edit existing user groups, to add new users to a group. [VIP-24312: This issue has been resolved.]
- A user associated with a group does not have permission to read and write cannot change their own password in the vManage Administration ► Manage Users screen. [VIP-24740: This issue has been resolved.]

---

## YANG Files for Netconf and Enterprise MIB Files

Netconf uses YANG files to install, manipulate, and delete device configurations, and Viptela supports a number of enterprise MIBs. Both are provided in a single tar file. Click the filename below to download the file.

- [YANG and Enterprise MIB files for Release 17.1.0](#)
- [YANG and Enterprise MIB files for Releases 17.1.1 and later](#)

---

## Using the Product Documentation

The Viptela product documentation is organized into seven modules:

Module	Description
Getting Started	Release notes for Viptela software releases, information on bringing up the Viptela overlay network for the first time, quick starts for vEdge routers, software download and installation, and an overview of the Viptela solution.
vEdge Routers	How to install, maintain, and troubleshoot vEdge routers and their components. Provides hardware server recommendations for the controller devices—vManage NMS, vSmart controller, and vBond orchestrator servers.



Module	Description
Software Features	Overview and configuration information for software features, organized by software release.
vManage How-Tos	Short step-by-step articles on how to configure, monitor, maintain, and troubleshoot Viptela devices using the vManage NMS.
Command Reference	Reference pages for CLI commands used to configure, monitor, and manage the Viptela devices. Includes reference pages for Viptela software REST API, a programmatic interface for controlling, configuring, and monitoring the Viptela devices in an overlay network.
vManage Help	Help pages for the vManage screens. These pages are also accessible from the vManage GUI.

---

## Tips

- To create a PDF of an article or a guide, click the PDF icon located at the top of the left navigation bar.
- To find information related to an article, see the Additional Information section at the end of each article.
- To help us improve the documentation, click the Feedback button located in the upper right corner of each article page and submit your comments.

---

## Using the Search Engine

- To search for information in the documentation, use the TechLibrary Search box located at the top of each page.
- On the Help results page, you can narrow down your search by selecting the appropriate documentation module at the top of the page. If, for example, you are searching for power supply information for your vEdge router model, select the Hardware module and then select your vEdge router model.
- When a search returns multiple entries with the same title, check the URL to select the article for your hardware platform or software release.
- When the search string is a phrase, the search engine prioritizes the individual words in a phrase before returning results for the entire phrase. For example, the search phrase *full-cone NAT* places links to "NAT" at the top of the search results. If such a search request does not return relevant results, enclose the entire search string in quotation marks (here, for example, "*full-cone NAT*").

---

## Issues

- The maximum PDF page limit is 50 pages.
- It is recommended that you use the Chrome browser when reading the production documentation. Some of the page



elements, such as the PDF icon, might not display properly in Safari.

---

## Requesting Technical Support

To request technical support, send email to [support@viptela.com](mailto:support@viptela.com).

To provide documentation feedback or comments, send email to [docs@viptela.com](mailto:docs@viptela.com).

---

## Revision History

Revision 1—Release 17.1.0, April 28, 2017

Revision 2—Release 17.1.1, June 5, 2017

Release 17.1.2 was not released.

Revision 3—Release 17.1.3, August 13, 2017

Revision 4—Release 17.1.4, October 10, 2017

Revision 5—Release 17.1.5, February 21, 2018

