Cisco SD-WAN (Viptela) Release Notes for Release 17.2

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Release Notes for Release 17.2

These release notes accompany Viptela Software Release 17.2, for Releases 17.2.0 through 17.2.10. The Viptela software runs on all Viptela devices, including vSmart controllers, vEdge routers, vBond orchestrators, and vManage NMSs.

Viptela Software Release 17.2
November 19, 2019
Revision 12

Product Features

Below are the main product features in Viptela Software Release 17.2:

- **AAA authentication order for console ports** —You can configure the order in which the software tries different authentication methods when verifying user access to an overlay network device for both SSH sessions and console ports. *See Configuring User Access and Authentication and auth-order .*
- **Application lists in centralized policy** —When you are configuring a centralized policy in vManage NMS, preset groups exist for Google and Office 365 applications. *See Policies .*
- **Automatically adjust the vManage data volume size** —The vManage NMS can automatically resize the third partition on the vManage NMS if the hypervisor has increased the size of this partition. *See request nms application-server .*
- **Cflowd traffic flows** —In Releases 17.2.2 and later, cflowd can track GRE, ICMP, IPsec, SCTP, TCP, and UDP flows. *See Traffic Flow Monitoring with Cflowd .*
- **Check for duplicate system IP addresses** —In Releases 17.2.2 and later, the vManage NMS checks that you have not configured duplicate system IP addresses for devices in the overlay network. You can modify the system IP addresses to remove duplicates during the process of attaching device configuration templates to the devices. *See Templates .*
- **Cloud OnRamp service** —Cloud OnRamp service for IaaS and SaaS extends the fabric of the Viptela overlay network into public cloud instances, allowing branches with vEdge routers to connect directly to public-cloud application providers. By eliminating the need for a physical data center, the Cloud OnRamp service improves the performance of SaaS applications. *See **Configuring Cloud OnRamp Service** and View Cloud OnRamp Instances .*
- **Collection of data streams from vEdge routers** —In Releases 17.2.2 and later, the vManage server can collect data streams from vEdge routers. *See Settings .*
- **Configuration audit log** —The vManage audit log shows the changes made to a configuration template. *See Audit Log .*
- **Configuration template rollback timer** —When you are attaching vEdge routers to a configuration template, you can configure the time interval at which the device rolls back to its previous configuration if the router loses its control connection to the overlay network. The default interval is 5 minutes. You can extend the interval to a maximum of 15 minutes. *See Templates .*
- **DHCP servers for DHCP helper interfaces** —In Releases 17.2.2 and later, for an interface acting as a DHCP helper, you can configure the addresses of up to eight DHCP servers. *See Configuring DHCP .*
- **Display status for controller devices** —In vManage NMS, you can display the system status for vBond orchestrator, vManage NMS, and vSmart controller devices. *See Determine the Status of a Network Device .*
- **Edit vManage configuration templates simultaneously** —You can edit templates simultaneously from one or more vManage servers. *See Templates .*
- **FIPS 140-2 level 1 certification for the crypto module on Viptela x86-based vEdge routers, vBond orchestrators, and vSmart controllers** —FIPS 140-2 is a U.S. Government security standard used to approve cryptographic modules. FIPS 140-2 certification means that Viptela products have implemented best practices for the use of cryptography to provide a more secure product. As part of the FIPS support, the SMVPv3 privacy type DES is no longer supported. For more details about FIPS 140-2, see the NIST website .

- **IKE** —The Viptela software supports IKE, Version 1, as defined in RFC 2409 , *Internet Key Exchange* , and IKE, Version 2, as defined in RFC 7296 , *Internet Key Exchange Protocol, Version 2* . With IKE, the Viptela software implements standard IPsec tunnels, which support AES-GCM, AES-CBC, and null encryption. Both AES-GCM and AES-CBC use 256-bit keys. In Release 17.2.3 and later, perfect forward secrecy (PFS) in enabled by on an IPsec tunnel that is being used for IKE key exchange, and you can disable PFS. *See Configuring IKE-Enabled IPsec Tunnels and ipsec .*

- **Ping host devices** —In Releases 17.2.2 and later, you can ping both the primary and secondary host IPv4 addresses in a service VPN. *See ping .*

- **Policy configuration wizard** —A vManage policy configuration wizard guides you through the creation of centralized control and data policies. *See Configure Policies and Policies .*

- **Queues for mapping forwarding classes** —In Releases 17.2.2 and later, vEdge Cloud routers support eight queues, with queue 0 reserved for LLQ. *See class-map .*

- **Rollback to previous configuration** —When you are configuring a vEdge router from vManage templates, you can configure the time interval at which the router rolls back to its previous configuration if the router loses its control connection to the overlay network. *See Templates .*

- **Split DNS** —You can create customized DNS-lookup architectures for devices in the overlay network. *See Configuring Split DNS .*

- **TCP optimization** —In service-side VPNs on vEdge routers, you can fine-tune TCP to decrease round-trip latency and improve throughput for TCP traffic. Optimizing TCP traffic can be useful for improving the performance of SaaS applications, transcontinental links, and high-latency transport devices such as VSAT satellite communications systems. *See Configuring TCP Optimization , action , and tcp-optimization .*

- **Track tunnels connected to the internet** —In Releases 17.2.2 and later, on NAT-enabled transport interfaces used for local internet exit, you can track the status of the internet connection. If the internet becomes unavailable, traffic is automatically redirected to the non-NATed tunnel on the transport interface. *See **Configuring Local Internet Exit** and tracker .*

- **Umbrella DNS security** —In Releases 17.2.2 and later, for DNS-layer security, you can configure a vEdge router to act as a DNS forwarder to Cisco Umbrella. *See Using Umbrella DNS Security .*

- **Use remote vManage NMS as software image repository** —You can store software images on a remote vManage NMS and retrieve these images when you are upgrading software images from another vManage NMS. *See Software Installation and Upgrade .*

- **vEdge 5000 router** —The vEdge 5000 router delivers highly secure site-to-site data connectivity to large enterprises. The vEdge 5000 router is a fixed-port-configuration router supporting up to eight 1-Gigabit Ethernet ports and up to two 10-Gigabit Ethernet ports. This router is available in Releases 17.2.2 and later. *See vEdge 5000 Router .*

- **vManage maintenance window** —You can configure a maintenance window for a vManage server. *See Dashboard and Settings .*

- **vManage multitenancy** —A service provider can manage multiple tenants from a single vManage NMS that is running in multitenant mode. *See Create a Multitenant vManage NMS and Use a Multitenant vManage NMS .*

- **vManage troubleshooting tools** —vManage NMS provides troubleshooting tools for checking the bringup status of a vEdge router, checking data traffic tunnel statistics, viewing application-aware routing traffic data, and generating network packets (nping). Release 17.2.2 adds troubleshooting tools for viewing control plane connections in real time, and displaying syslog messages. *See the vManage Troubleshooting How-Tos , including Check Application-Aware Routing Traffic , Ping a Viptela Device , Troubleshoot Device Bringup , Use Syslog Messages , View Control Connections , and View Tunnel Health .*

- **vManage web server certificate expiration** —Starting 60 days before the certificate expires, the vManage Dashboard displays a warning indicating that the certificate is about to expire. This warning is then redisplayed 30, 15, and 7 days before the expiration date, and then daily. *See Dashboard and Settings .*

Command Changes

New and Modified Configuration Commands

Command	Hierarchy	New	Modified	Comments
action	policy data-policy vpn-list sequence		X	Add **redirect-dns** option, for split DNS.
admin-auth-order	system aaa		X	Add support for console ports.
authentication-type	vpn interface ipsec ike	X		For IKE.
auth-fallback	system aaa		X	Add support for console ports.
auth-order	system aaa		X	Add support for console ports.
cipher-suite	vpn interface ipsec ike, vpn interface ipsec ipsec	X		For IKE.
class-map	policy		X	In Releases 17.2.2 and later.
dead-peer-detection	vpn interface ipsec	X		For IKE.
eco-friendly-mode	system	X		For vEdge Cloud routers.
group	vpn interface ipsec ike	X		For IKE.
dhcp-helper	system		X	In Releases 17.2.2 and later.
idle-timeout	system		X	
ike	vpn interface ipsec	X		For IKE.
interface ipsec	vpn	X		For IKE.
match	policy data-policy vpn-list sequence policy app-policy vpn-list sequence		X	Add **dns** and **dns-app-list** options, for split DNS.
mode	vpn interface ipsec ike	X		For IKE.
perfect-forward-secrecy	vpn interface ipsec ipsec	X		For IKE. In Release 17.2.3 and later.
port-hop	vpn 0 interface tunnel-interface	X		
priv	snmp user		X	Remove support for DES privacy.
rekey	vpn interface ipsec ike, vpn interface ipsec ipsec	X		For IKE.
replay-window	vpn interface ipsec ipsec	X		For IKE.
service	vpn		X	Add **TE** service.
tcp-optimization	vpn	X		
tcp-optimization-enabled	system	X		
tracker	system, vpn 0 interface	X		In Releases 17.2.2 and later.
tunnel-destination	vpn interface ipsec	X		For IKE.
tunnel-source	vpn interface ipsec	X		For IKE.
tunnel-source-interface	vpn interface ipsec	X		For IKE.

version	vpn interface ipsec ike	X		For IKE.

New and Modified Operational Commands

Command	New	Modified	Comments
clear ipsec ike sessions	X		
request ipsec ike-rekey	X		For IKE.
request ipsec ipsec-rekey	X		For IKE.
request nms application-server		X	Add **resize-data-partition** , **software reset** , and **software upgrade** options.
show app tcp-opt	X		
show bfd tloc-summary-list		X	Add **ipsec-ike** option.
show hardware real-time-information	X		
show ip routes		X	Rename **natpool-omp** and **netpool-service** options to **natpool-inside** and **natpool-outside** .
show ipsec ike inbound-connections	X		
show ipsec ike outbound-connections	X		
show ipsec ike sessions	X		
show security-info		X	Add FIPS status.
show system buffer-pool-status	X		
show system status		X	Add CPU allocation field to command output. Add FIPS status.

Upgrade to Release 17.2

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 controller devices 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 controller devices 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, and it can also be higher.

Note: To add the MIPS software image (for vEdge hardware routers) for Release 17.2.2 to the vManage NMS software repository, you must first upgrade the vManage NMS to Release 17.2.2. This same requirement holds true for later releases of Release 17.2: The vManage NMS must be running the newer software release before you can upload the MIPS software image for that newer software release to the vManage software repository.

Note: If your vManage NMS is running as a cluster, please contact Customer Support for assistance when upgrading to Release 17.2.0 or later.

To upgrade to Release 17.2:

1. In vManage NMS, select the Maintenance ► Software Upgrade screen.
2. Upgrade the controller devices to Release 17.2 in the following order:
    a. First, upgrade the vManage NMSs in the overlay network.
    b. Then, upgrade the vBond orchestrators.
    c. 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 17.2, you can never downgrade it to Release 17.1 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, 17.2 and 17.1 are examples of major releases. Releases 17.2.0 and 17.1.0 denote the initial releases, and Releases 17.2.7 and 17.1.5 are maintenance releases.

When you are upgrading to Release 17.2.5, if you have vEdge 5000 routers in your network, you must upgrade them manually to Release 17.2.5 first, before you upgrade the controller devices (the vBond orchestrator, vSmart controller, and vManage NMS) to Release 17.2.5. To upgrade the vEdge 5000 routers manually:

1. From the CLI, install the software on the vEdge 5000 router:
   vEdge-5000# **request software install http://** *vmanage-system-ip* **:8080/software package/viptela-17.2.5-x86_64.tar.gz**
2. Activate the software:
   vEdge-5000# **request software activate 17.2.5**
3. After the router reboots, confirm that the software upgrade was successful:
   vEdge5000# **request software upgrade-confirm**
4. Follow the standard procedure above to upgrade the controller devices to Release 17.2.5.

You cannot upgrade a multitenant vManage NMS from Release 17.2.0 or Release 17.2.1 to Release 17.2.2. Instead, you must freshly install the multitenant vManage NMS:

1. Upgrade the software on the vManage NMS to Release 17.2.2.
2. In the vManage Maintenance ► Software Upgrade screen, set the vManage default software version to 17.2.2.
3. In the vManage Tools ► SSH Terminal screen, open an SSH session to the vManage server itself.
4. Issue the **request software reset** command to return the NMS to its default configuration.
5. Re-create your configurations.

Note that you cannot install a Release 17.2 or earlier image on a vEdge router that is running Release 18.2.0 or later. This is the result of security enhancements implemented in Release 18.2.0. Note that if a Release 17.2 or earlier image is already present on the router, you can activate it.

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 17.2:

- Use **max-control-connections 0** instead of the **no control-connections** command in **tunnel-interface** configuration mode. The **no control-connections** command has been deprecated and has no effect on releases 17.2 and later.

- 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 17.2. 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 17.2, 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.

- On all vEdge hardware and vEdge Cloud routers, the reported CPU utilization might be 20 to 30 percent higher than in previous software releases. This increase occurs because of changes in how the software computes CPU utilization, not because of any underlying changes to the software itself. To display router CPU utilization, in vManage NMS, in the Monitor ► Network screen, select a router and then select the System Status command. In the CLI, check the Load average field in output of the **show system status** command.

- On a vEdge 5000 router, you cannot enable TCP optimization by configuring the **tcp-optimization-enabled** command.

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:

  - 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.

  - 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.

  - 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.

- When you activate the configuration on a router with cellular interfaces, the primary interfaces (that is, those interfaces not configured as circuits of last resort) and the circuit of last resort come up. In this process, all the interfaces begin the process of establishing control and BFD connections. When one or more of the primary interfaces establishes a TLOC connection, the circuit of last resort shuts itself down because it is not needed. During this shutdown process, the circuit of last resort triggers a BFD TLOC Down alarm and a Control TLOC Down alarm on the vEdge router. These two alarms are cleared only when all the primary interfaces lose their BFD connections to remote nodes and the circuit of last resort activates itself. This generation and clearing of alarms is expected behavior.

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

Configuration and Command-Line Interface

- When upgrade to Release 17.2 from any prior Viptela software release, the CLI history on the Viptela device is lost. The CLI history is the list of commands previously entered at the CLI prompt. You typically access the history using the up and down arrows on the kayboard or by typing Ctrl-P and Ctrl-N. When you upgrade from Release 17.2 to a later software release, the CLI history will be maintained.

- 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.

- 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.

- DHCPv6 returns only an IPv6 address. No default information is accepted. IPv6 router solicitation and router advertisement messages are not processed.

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

- It is recommended that you use IKE Version 2 only with Palo Alto Networks and Ubuntu strongSwan systems only. Viptela has not tested IKE Version 2 with other systems.

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

- When a task stops and a vEdge router reboots, the router might no longer reboot. This situation occurs after the router reboots three times within 20 minutes, five times within 60 minutes, or seven times within the last 24 hours. During this time, the control plane on the router remains up, so traffic continues to be sent to the node. To override this behavior, recover the router via the console port.
- The Viptela software includes a version of OpenSSH that is affected by the vulnerabilities identified by the following Common Vulnerability and Exposures (CVE) IDs: CVE-2016-10009 and CVE-2016-10012.
- The vEdge 100m and vEdge Cloud routers use an outdated and known vulnerable version of the OpenSSL library.

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 17.2. The number following each issue is the bug number in the Viptela bug-tracking database.

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.
- After you upgrade from Release 17.1.4 to Release 17.2.5, the cellular process (wwand) might not start. [VIP-40991]

Configuration and Command-Line Interface

- 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]
- 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]
- With the **ping source ip-address** command, if you type it as **ping so ip-address** , the CLI does not autocomplete **so** and the command fails. You must type out the keyword **source** . [VIP-36087]
- When you use the **show ip route** command to query a route that is not present in the route table, the command might return no output or no failure message. [VIP-36725]
- The **show omp routes** command output might display the incorrect tag for OMP routes. [VIP-39722]
- When you attempt to activate a vEdge cloud router on which the organization name is not configured, the router is not activated and no error message is displayed. [VIP-40503]
- A template push from the vSmart controller might fail with an application communication error. [VIP-41576]

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]

- 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]
- 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 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]
- Routes might be installed in the routing table with the incorrect color. [VIP-35088]
- Traffic might be discarded 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 discarded. [VIP-35113]
- On a vBond orchestrator, if you configure **allow-service netconf** , the vBond orchestrator does not open TCP port 830 and thus cannot connect to the vManage NMS. As a workaround, configure **allow-service all** . [VIP-35916]
- An implicit or a configured ACL might not work on loopback interfaces. [VIP-38731]
- On a router with multiple tunnels, one of the tunnels might stop forwarding traffic. As a workaround, reboot the router. [VIP-41305]

Interfaces

- 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 primary again becomes a backup, other devices continue to point to the MAC address for the backup router, and traffic is discarded until ARP cache on the other devices expires and is updated with the correct MAC address of primary vEdge router, a process that typically takes a few minutes. [VIP-33722]
- 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]
- You might not be able to configure the Cloud Onramp VPC even when vEdge routers are present. [VIP-34655]
- On a vEdge 100b router running Release 17.2.6, when you enable transport interface tracking on an interface, if you bring the interface down, the tracker might fail and cause the router to crash. [VIP-40006]

Policy

- A centralized policy that is pushed from the vSmart controller to the vEdge routers might not be applied on the routers. [VIP-27046]
- In vManage NMS, when you use the policy configuration wizard to create policies for a mesh topology, you might need to create an additional policy using a CLI template for the mesh policy to work. This situation is known to occur in a network that has two regions, where each region is mesh that is a subset of the entire network, where each region has its own data center, and where the branch vEdge routers in one region communicate with branch routers in the other region through the data centers. We will call these Region 1 and Region 2. Assume that Region 1 has a control policy that advertises its TLOCs to the data center in Region 2, and Region 2 has a control policy that prevents the spokes and data center in Region 2 from advertising TLOCs to the spokes in Region 1. The result is that the data center in Region 2 repeatedly attempts to form control tunnels to the data center in Region 1, but these attempts fail. As a workaround, you must a policy using a CLI template that allows the data center in Region 2 to exchange TLOCs with the data center in Region 1 and then attach that policy to the vEdge routers. [VIP-29933]

- In the vManage Centralized Policy UI Builder, the Membership has no options to accept or reject and no way to change the default action. These options are all available in CLI. [VIP-38730]
- If you enable NAT and apply a localized data policy on a transport interface, the control connection on that interface might not come up. [VIP-41968]
- On vEdge Cloud routers running Release 17.2, the **show policy service-path** command might not work. [VIP-34184]

Routing Protocols

- If two IPv6 interfaces on a vEdge router are active at the same time, basic packet forwarding (such as ping) might fail. Only one interface works at a time, and the router switches back and forth between the two interfaces. [VIP-37810]
- When you configure VPN-specific OMP route advertisement parameters to two VPNs on a vEdge router, they might not be applied properly in the VPNs. [VIP-37860]
- The **show omp tlocs advertised** command might not properly filter TLOC-advertised routes. [VIP-39400]

Security

- If an IPv6 address for the IPsec tunnel source interface, the IPsec tunnel does not come up. [VIP-29912]
- A vBond orchestrator might crash in response to a denial-of-service (DOS) attack. [VIP-37947]
- When you specify a DNS destination for IPsec on a tunnel in VPN 0, the Forwarding Table Management process (ftmd), the OMP process (ompd), or the IKE process (iked) might crash. [VIP-38888]
- When you send 1500-bytes UDP packets with a destination port of 1234 at 35 Mbps, the control and BFD connections to the far-end router might do down. [VIP-39440]
- From an SSH session on a vManage NMS running Release 17.2.7, you might not be able to display information about a router's control connections. [VIP-42217]

SNMP

- When traffic exceeds 85% of the bandwidth configured on a transport interface, SNMP traps might not get triggered. [VIP-33435]
- When you poll the VIPTELA-OPER-VPN MIB, interface descriptions are limited to 32 characters. [VIP-35787]
- An snmpwalk operation for table ompRoutesTableFamilyEntriesReceivedAttributesOriginMetric in the viptela-omp MIB might time out. [VIP-42604]

System

- vBond orchestrators might report a large number of control-connection-auth-fail events. [VIP-22976]
- 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]
- The vManage server might not process events received from vEdge routers. [VIP-28673]
- When a certificate for controllers is about to expire, no syslog message is generated. [VIP-28960]
- On vEdge routers, when you issue an **nping** command for IPv6, the command might fail, and a core file might be created on the router. From vManage NMS, you issue this command from the Monitor ► Network ► Troubleshooting ► Ping pane. From the CLI, you use the **tools nping** command, specifying **options "--ipv6"** . [VIP-31924]
- 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]
- A standalone vManage NMS deployed on ESXi might become inaccessible because the server_config.json file gets corrupted. [VIP-36282]
- The TCP optimization process might consume a large amount of CPU even though TCP optimization is not configured. [VIP-36675]
- When you change the negotiated interface speed on a vEdge router, the buffer allocation also changes. [VIP-37238]
- The gzip process on a router might consume a lot of CPU. [VIP-39228]
- A vEdge router might crash with the error message, "FTMD-3-ERRO-1000011: FP Core 1 Died. Core file recorded at /var/crash/core.fp1.5310". [VIP-38318]
- If you are using ZTP or attaching a template with bridge configuration to a new vEdge router, pushing the device configuration template to the router might fail because of a bridge configuration. As a workaround, remove the bridge configuration from feature template, apply the IRB configuration, and then apply the bridge configuration. [VIP-39525]
- When you use VLAN tagged interfaces for transport interfaces, the interface throughput might drop. [VIP-39647]
- You might not be able to change the AWS instance from C3 to C4. [VIP-40092]
- When you create a new vManage VM instance and are formatting the third partition, there might be no indication that the formatting is progressing. [VIP-40254]
- The output of the **show bridge interface** command might not list all the configured bridges. [VIP-40299]
- In Releases 17.2.8 and later, activating software on a platform hosted in AWS might fail. If you encounter this issue, contact Customer Support for assistance. [VIP-42195]
- You might not be able to reboot a vEdge 1000 router by issuing reboot commands from the CLI. [VIP-42265]
- You might not be able to open an SSH connection from a vEdge router to a vManage NMS. [VIP-42465]
- A vEdge 5000 router might experience application performance issues, timeouts, and lower throughput that a vEdge 2000. [VIP-42625]

vEdge Hardware

- On a vEdge 100m router, after you execute the **request software reset** command, the router might reboot continuously. [VIP-24149]
- 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]
- On a vEdge 2000 router, weighted round-robin queues might favor queues with larger packets. [VIP-41883]

vManage NMS

- 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]
- A vManage NMS might not be able to synchronize its configuration with a vSmart controller. [VIP-26270]
- 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 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 indicated the cause of the error. [VIP-32253]
- You might not be able to push a configuration template to a vEdge router. [VIP-32277]

- 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]
- 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]
- The vManage configuration database might not be able to process records. [VIP-34251]
- You might not be able to push configuration templates to vEdge routers. [VIP-34886]
- When you change the names of the route policies in a localized policy, the modified policy might not work as expected. [VIP-35026]
- After a vManage NMS silently reboots, it might be out of sync with the vManage cluster. [VIP-35891]
- When a vBond orchestrator is unreachable or has wrong credentials configured, pushing a vEdge list to it fails with the message “File /home/ *user* /vedge_serial_numbers must be in home directory”, which does not provide any useful information to the user to understand what is wrong. [VIP-36285]
- The default VPN 512 management feature template is named "Transport VPN", which is confusing because VPN 0 is the transport VPN. [VIP-36771]
- In a vManage cluster running Release 17.2.3, one of the vManage servers might not be able to connect to the configuration database. As a workaround, issue the **request nms all restart** command on the vManage server to restart all NMS services. [VIP-36805]
- If NMS services are down on one of the servers in a vManage cluster, you might not be able to perform an CLI operations from the vManage NMS. [VIP-37672]
- In the vManage policy builder, you can configure a site list name that is longer that is allowed on the vSmart controller. When you attempt to activate the policy, an error occurs on the vSmart controller, [VIP-37859]
- The vManage interface feature configuration templates do not have drop-downs for selecting interface speed and duplex settings. [VIP-37973]
- In the BGP feature template, the BGP neighbor route policy variable name is displayed as the default name (bgp_neighbor_policer_out_pol_name) rather than the name you enter. However, the name you enter does show up when you are attaching the template to the device. [VIP-37988]
- After a CSR request is sent to Symantec but before the certificate has been approved, the vManage request to retrieve the certificate might die, and so the vManage NMS might not be able to retrieve the certificate even after it has been approved. [VIP-38093]
- In a vManage cluster with two servers, if both servers go down, you might need to manually restart the vManage services to return the vManage servers to an operational state. To do this, issue the **request nms configuration-db restart** command from the vManage server. One way to determine whether you need to restart the services is to check the /var/log/nms/debug.log file on the vManage server for a message indicating that neo4j needs to be restarted. [VIP-38228]
- If you reboot one of vManage servers in a cluster while the vManage NMS is downloading a software image to a vEdge router, the cluster might report server errors and might stop operating properly. [VIP-38556]
- A color that you configure using a vManage configuration template might not be applied correctly on a vEdge router. [VIP-38735]
- The vManage DPI screen might display a DIA graph for a vEdge router on which local internet exit is not configured. [VIP-38987]
- The vManage GUI might not start, and the vmanage-server.log file might contain a lot of exceptions. [VIP-39644]
- When you bring up a router from the vManage NMS, duplicate entries for these devices might show up in the Configuration ► Device and Configuration ► Certificates screens. [VIP-39692]
- The vManage dashboard might show an incorrect number of vBond orchestrators. [VIP-40033]
- On the vManage Monitor ► Alarms screen, the OMP site down alarm might not clear after the OMP connection is restored. [VIP-40200]

- After you upgrade to Release 17.2.5, the web server certificate might no longer be valid. [VIP-40889]
- After you upgrade to Release 17.2.7, pushing configuration templates to a router might take a long time. [VIP-41233]
- The vManage NMS might show the incorrect number of control connections. [VIP-41289]
- You might not be able to remove a vManage NMS from a vManage cluster. [VIP-41880]
- After you upgrade vManage NMSs to Release 17.2.7, authentication timeouts and API errors might occur. [VIP-42168]
- When under a heavy load, the vManage NMS might become slow to respond, and control connections might go down and come back up. [VIP-42319]
- When you push the vEdge router list from the vManage NMS, some vEdge routers might become unreachable. [VIP-42567]
- To activate a vEdge Cloud router, entering the UUID of the device in upper case letters might not work. As a workaround, enter the UUID in lower case letters. [VIP-42660]
- When you upgrade the vManage NMS software to Release 17.2.0, errors might occur on the vManage server, including graphs being out of sync, tasks still appearing to be running ever after the server reboots, and incorrect status, such as a failure status, ever after the server is up and running. [VIP-31458]

Fixed Issues

Issues Fixed in Release 17.2.10

The following issues have been fixed in Viptela Software Release 17.2.10. The number following each issue is the bug number in the Viptela bug-tracking database.

Forwarding

- On vEdge 1K, race condition BFDs may result in out-of-sync when unpinned flows are configured. [VIP-38117: This issue has been resolved.]
- Configuring "icmp-error-pps 0" under **System ► Configuration** crashed on vEdge-Cloud with an arithmetic exception. [VIP-41653: This issue has been resolved.]
- vEdge 5000 frequently crashes with FP core watchdog failure. [VIP-48186: This issue has been resolved.]
- vEdge should update ARP table upon receipt of ARP. [VIP-50688: This issue has been resolved.]

Hardware

- pimd daemon terminated due to signal 11. [VIP-48959: This issue has been resolved.]
- pimd daemon crash. [VIP-49929: This issue has been resolved.]
- Match U-boot environment to Linux expectations. [VIP-50014: This issue has been resolved.]
- vEdge-100B/M/WM power-on issue with GE0/4. [VIP-51609: This issue has been resolved.]

Interface

- Decouple buffer allocation for egress queues from the negotiated interface speed. [VIP-37238: This issue has been resolved.]
- Pings above 1472 on vE5K fail on 10G interfaces. [VIP-46716: This issue has been resolved.]
- Packets sourced with the loopback interface exceeding mtu on the service side are not fragmented, which results in drops. [VIP-49531: This issue has been resolved.]

- After an upgrade from 17.2.7 to 18.3.5, cannot ping packets with a 1518 length unless we do a **shut and no shut** command. [VIP-51906: This issue has been resolved.]

Security

- STD_IPSEC: Charon daemon uses extremely high memory. [VIP-31686: This issue has been resolved.]
- Control connections to the controller take greater than 10 minutes in some situations. [VIP-48060: This issue has been resolved.]

System

- vEdge5k control not coming up with vBond control connections if there is tab for it. [VIP-47532: This issue has been resolved.]

Issues Fixed in Release 17.2.9

The following issues have been fixed in Viptela Software Release 17.2.9. The number following each issue is the bug number in the Viptela bug-tracking database.

Forwarding

- On a vEdge 100m router, the Forwarding Table Management process (ftmd) crashes when you configure **request admin-tech.** [VIP-39984: This issue has been resolved.]
- On a vEdge 100m router with the ftmd, memory leakage happens leading to segmentation fault. [VIP-45371: This issue has been resolved.]
- Traffic fails on WAN SRIOV - i350 driver when combination of Virtio and WAN SRIOV (i350 driver) is used to deploy vEdge on ENCS - NFVIS. [VIP-34528: This issue has been resolved.]

vEdge Hardware

- vdaemon crashes on vEdge 5000 after executing continuous control connections flaps. [VIP-45672: This issue has been resolved.]
- vEdge 5000 control connections are not coming up after a reboot. [VIP-44670: This issue has been resolved.]
- vEdge 5000 vDaemon crashes while attempting to establish control connections at bootup. [VIP-46135: This issue has been resolved.]

Issues Fixed in Release 17.2.8

The following issues have been fixed in Viptela Software Release 17.2.8. The number following each issue is the bug number in the Viptela bug-tracking database.

Cellular Interfaces

- On a vEdge 100m, the cellular process (wwand) might crash. [VIP-41798: This issue has been resolved.]

Forwarding

- On a vEdge 100m router, the Forwarding Table Management process (ftmd) might crash. [VIP-39984: This issue has been resolved.]

OMP

- On a vSmart controller, the OMP process (ompd) might crash. [VIP-41807: This issue has been resolved.]

QoS

- On x86 vEdge routers (Cloud vEdge routers and vEdge 5000 routers), when you configure QoS schedulers, any traffic that queued in a queue to which no bandwidth or buffer is assigned will be dropped. VIP-38008: This issue has been resolved.]

Routing Protocols

- When you configure **omp overlay-as** to have OMP advertise a BGP AS number, BGP might go down and then come back up. [VIP-33587: This issue has been resolved.]

vEdge Hardware

- vEdge 1000 and vEdge 2000 routers might experience a high CPU usage. [VIP-41438: This issue has been resolved.]

vManage NMS

- For some routers, the Assigned Template field on the Configuration ► Devices screen might be empty. [VIP-41246: This issue has been resolved.]
- Staging routers might take a long time if they are using the same site ID. [VIP-41554, VIP-41558: These issues have been resolved.]
- In vManage Policies ► Custom Options ► List ► TLOC, the scroll bar might not be able to scroll down to display all the TLOCs in the list. [VIP-41770: This issue has been resolved.]
- vManage NMS might not report router alarms. [VIP-42026: This issue has been resolved.]

Wireless WANs

- The **show hardware inventory** command does not display the AT&T SKUs 100m-AT and 100wm-AT. [VIP-41718: This issue has been resolved.]

Issues Fixed in Release 17.2.7

The following issues have been fixed in Viptela Software Release 17.2.7. The number following each issue is the bug number in the Viptela bug-tracking database.

DHCP

- When the software reads the DHCP lease file, the read might fail, resulting in a vdhcpsd core. [VIP-39037: This issue has been resolved.]

Forwarding

- A vEdge 2000 router physical interface might drop packets larger than 1480 bytes that are sent on loopback interfaces. [VIP-27216: This issue has been resolved.]
- On vEdge routers, the **show policy access-list-counters** command might not display any values in the Bytes column. [VIP-28890: This issue has been resolved.]
- 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: This issue has been resolved.]
- If the vEdge routers in your overlay network are running Release 17.2, you cannot add routers to the network that are running Release 15.4. [VIP-35084: This issue has been resolved.]
- A GRE tunnel might negotiate an MTU size of 512 bytes, so packets larger than about 500 bytes cannot be sent over the tunnel. [VIP-38791: This issue has been resolved.]

Interfaces

- When you configure interface tracking on two interfaces in a vEdge router, the router might crash. [VIP-38829: This issue has been resolved.]

Policy

- After you change a policy on the vSmart controller, the OMP process (ompd) process might fail and the vSmart controller might crash. [VIP-34098: This issue has been resolved.]

Routing Protocols

- When you are upgrading vEdge routers to Release 16.2.12, the BGP process (bgpd) might crash during the reboot process, when the router is shutting down. [VIP-29523: This issue has been resolved.]
- When OMP redistributes BGP routes, it might not include in the origin metric. [VIP-36580: This issue has been resolved.]

System

- 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: This issue has been resolved.]
- When you upgrade the vManage NMS to Release 17.2.2, the vManage server might boot continuously and login screen might not never be displayed. [VIP-34441: This issue has been resolved.]
- After you upgrade the vManage NMS to Release 17.2.2.1, you might not be able to log into the NMS. The error message displayed is "Server is initializing, please wait". [VIP-34697: This issue has been resolved.]
- In Viptela Software Release 17.1.4, when you use OpenStack Heat to create a vManage NMS, the /dev/vdb disk is recognized and mounts. However, in Release 17.2.3, the disk volume might not be automatically recognized and mounted. [VIP-35907: This issue has been resolved.]
- After a DOS attack on a vBond orchestrator or a vSmart controller, the Viptela software process (vdaemon) might crash. [VIP-37947: This issue has been resolved.]

vEdge Hardware

- In a vEdge 5000 router, if you remove and replace an SFP without powering off the router, the router is unable to detect the new SFP module details. For example, the **show hardware inventory** command will not list the SFP module. However, the ports on the SFP still connect and work as expected. [VIP-37562: This issue has been resolved.]
- A vEdge 5000 router might not be able to start because of issues with its internal verification hardware. [VIP-38219: This issue has been resolved.]

vManage NMS

- The vManage server might not process events received from vEdge routers. [VIP-28312: This issue has been resolved.]
- 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: This issue has been resolved.]
- The vManage NMS might not be able to retrieve records from the configuration database. When this occurs, the vManage NMS displays exception errors. [VIP-35702: This issue has been resolved.]
- In a vManage cluster, when you try to display the Maintenance ► Software Upgrade ► vEdge screen, the screen might not display, and the vManage-Server.Log shows an exception error. [VIP-35926: This issue has been resolved.]
- In Release 17.2.3, when a user who does not have permission to view certificate information logs in to a vManage server, the Dashboard displays an error. This behavior did not occur in Release 17.2.0. [VIP-36755: This issue has been resolved.]
- In the vManage SNMP feature configuration template, when you try to add trap types to a trap group, the ADD button might not work. [VIP-37376: This issue has been resolved.]

- If you try to push a Service VPN configuration template, the operation might fail with the error "Failed to create input variables". [VIP-37705: This issue has been resolved.]
- The vManage NMS might occasionally indicate that the software upgrade on a vEdge router has succeeded even though it has failed. This has been observed only rarely. [VIP-39004, VIP-39041: This issue has been resolved.]
- Autoformatting of a vManage VM partition might not work. [VIP-39137: This issue has been resolved.]
- The vManage dashboard might show an incorrect number of vBond orchestrators. [VIP-40033: This issue has been resolved.]
- When you issue the **show system buffer-pool-status** command, the following error might occur: "Error: application communication failure". [VIP-40073: This issue has been resolved.]
- After you reboot a large number of routers, the vManage cluster performance might degrade. [VIP-40330: This issue has been resolved.]

Issues Fixed in Release 17.2.6

The following issues have been fixed in Viptela Software Release 17.2.6. The number following each issue is the bug number in the Viptela bug-tracking database.

Security

- A vEdge 5000 router running Release 17.2.5 might crash with CRYPTO_free and EVP_CipherInit_ex, and create a core file. [VIP-38083: This issue has been resolved.]

System

- A vSmart controller running Release 17.2.5 might be missing the root chain certificate for a vEdge 5000 router. [VIP-38223: This issue has been resolved.]

Issues Fixed in Release 17.2.5

The following issues have been fixed in Viptela Software Release 17.2.5. 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: This issue has been resolved.]

Configuration and Command-Line Interface

- On cellular interfaces, you might not be able to modify the maximum segment size (MSS) of TCP SYN packets. [VIP-28033: This issue has been resolved.]

Forwarding

- 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: This issue has been resolved.]
- If you enable TCP optimization on a vEdge 1000 router, the router might drop ARP responses. [VIP-33507: This issue has been resolved.]
- On a vEdge router, when you configure a single interface in the transport VPN, the router might crash immediately and create a forwarding process (fp) core file. [VIP-34201: This issue has been resolved.]
- A vEdge 1000 router running Release 17.2.2 might reboot and might create a crash file that contains the message "Software initiated - FP core watchdog fail". [VIP-34846: This issue has been resolved.]

- When you enable VRRP, auto-RP messages might not be received consistently and so multicast routing might not work. [VIP-35416: This issue has been resolved.]
- A vEdge 5000 router might not be able to forward data traffic. [VIP-36850: This issue has been resolved.]
- When you are forwarding traffic or enable cFlowd on a vEdge 5000 router, the router might crash and create a core file. [VIP-36951, VIP-37218: This issue has been resolved.]
- A bridge might drop large-frame traffic. [VIP-36960: This issue has been resolved.]

Policy

- When you issue the **request admin-tech** command, the Forwarding Policy Manager process (fpmd) might crash. [VIP-34031: This issue has been resolved.]
- On vEdge Cloud routers running Release 17.2.1, the **show policy service-path** command might not work. [VIP-34184: This issue has been resolved.]
- 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: This issue has been resolved.]

Routing Protocols

- When you have configured a BGP peering session to restart after receiving a more than a set number of prefixes from its neighbor, the session might not restart when the number of prefixes is exceeded. [VIP-33780: This issue has been resolved.]
- A vEdge router might not be able to establish OSPF point-to-point interfaces with a Juniper EX4200 device. [VIP-34936: This issue has been resolved.]
- When a BGP route's AP PATH attribute is redistributed to OMP, OMP might advertise it only on one transport even when multiple transports are present. [VIP-36578: This issue has been resolved.]
- If there are multiple GRE TLOCs, OMP might not advertise the routes for some of them. [VIP-36658: This issue has been resolved.]

Security

- The connection between a vEdge router and a vBond orchestrator might not come up, and debugging messages indicate that the certificate validation failed. As a workaround, regenerate the certificate. [VIP-35350: This issue has been resolved.]
- After you upgrade to Release 17.2.3, the vBond orchestrator might take a long time to stabilize and to establish connections to all the vSmart controllers. [VIP-35514: 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.]
- When the vManage NMS experiences a kernel panic and reboots, the /var/crash/crash.dump file might be deleted. [VIP-34248: This issue has been resolved.]
- On a vManage NMS, when you issue the **request nms configuration-db backup** command, the following error message might be displayed: sh: tput: command not found. [VIP-35177: This issue has been resolved.]
- After you upgrade a vManage NMS to Release 17.2.3, the vManage server might not be able to push configuration templates to the vEdge routers. [VIP-35324: This issue has been resolved.]
- In vManage Monitor ► Network ► System Status, the CPU utilization reported for a vEdge router might be higher than the value reported when you log in to the router itself. [VIP-35342: This issue has been resolved.]
- When you access the vManage server from the Mozilla Firefox browser, if you update the interface in area 0 in the OSPF feature template, you might not be able to save the template. As a workaround, use the Chrome browser. [VIP-35584: This issue has been resolved.]

- The vManage Tenant Dashboard might display the name of the service provider instead of the name of the tenant. [VIP-35862: This issue has been resolved.]
- SSH forwarding from Viptela components was enabled by default. With this fix, SSH forwarding is now disabled. [VIP-37085: THis issue has been resolved.]

vEdge Hardware

- When a vEdge 2000 router reboots, the reboot reason field might show only a value of 0. [VIP-23941: This issue has been resolved.]
- On a vEdge 100 router, when you enable or disable debugging, a Forwarding Process (fp) core file might be created. [VIP-26965: This issue has been resolved.]

vManage NMS

- 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: This issue has been resolved.]
- A vManage serve might continue to attempt to fetch certificates even though all certificates are installed. [VIP-27416: This issue has been resolved.]
- From a vManage server, you might not be able to SSH into a vEdge router that is in staging mode. [VIP-33119: This issue has been resolved.]
- 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: This issue has been resolved.]
- The vManage NMS might not display statistics for some tunnels. [This issue has been resolved.]
- When you upgrade the vManage server from Release 17.1.4 to Release 17.2, you might see certificate null pointer exceptions. [VIP-34901: This issue has been resolved.]
- The vManage latency and jitter graphs might not match the reported values. [VIP-36729: This issue has been resolved.]
- In the vManage Configuration ► Policies ► Centralized Data Policies screen, a user who does not have policy write permission might see the copy, edit, and delete actions in the More Actions icon to the right of a policy listed in the policy table. [VIP-36770: This issue has been resolved.]

Issues Fixed in Release 17.2.4

The following issues have been fixed in Viptela Software Release 17.2.4. The number following each issue is the bug number in the Viptela bug-tracking database.

Forwarding

- In some scenarios, packets in certain queues might not be dequeued. [VIP-35222: This issue has been resolved.]
- On vEdge Cloud routers, the vdebug file might contain a large number of "stray: USER1: RED-DROP: queue 3, hw_queue 3, queue_depth 286, max_depth 286'' messages. [VIP-35301: This issue has been resolved.]

Issues Fixed in Release 17.2.3

The following issues have been fixed in Viptela Software Release 17.2.3. The number following each issue is the bug number in the Viptela bug-tracking database.

Cloud Express Service

- In the vManage Configure ► Policies screen, when you include the Microsoft_Apps application list in a policy, the policy might not work for some Office 365 applications. [VIP-34227: This issue has been resolved.]

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: This issue has been resolved.]

Forwarding

- 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: This issue has been resolved.]
- On a vEdge Cloud router, traffic shaping through QoS might not work properly. [VIP-26557: This issue has been resolved.]
- IRB interface queue statistics might display only 0 values. [VIP-31646: This issue has been resolved.]
- After a BFD session on a vEdge router has gone done and come back up multiple times, the Viptela software process (vdaemon) might stop operating. [VIP-31808: This issue has been resolved.]
- When a vEdge router receives a burst of TCP-optimized traffic before the receiver can revise its congestion window, the packet buffers on the router might be completely depleted. [VIP-31994: This issue has been resolved.]
- For a vEdge Cloud VM instance on Azure, when using IKE version 2, the IPsec tunnel to the vEdge Cloud router might go operationally down every hour or so, and you have to shut down the interface and then bring it back up. [VIP-33425: 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.]

SNMP

- For a vEdge Cloud VM instance on Azure, when using IKE Version 2, the IPsec tunnel to the vEdge Cloud router might go operationally down every hour or so, and you have to shut down the interface and then bring it back up. [VIP-33976: This issue has been resolved.]

System

- Syslog files might contain the names of unknown users. [VIP-34043: This issue has been resolved.]

vManage NMS

- On the Configuration ► Templates page, you might not be able to scroll horizontally. [VIP-34205: This issue has been resolved.]
- After you upgrade a vManage NMS to 17.2.0, events but no alarms might be generated, and using API call to the retrieve alarm statistics might return an internal server error. [VIP-34228: This issue has been resolved.]
- After you refresh the vManage real-time application-aware routing statistics screen, the data displayed might be inaccurate. [VIP-34229: This issue has been resolved.]
- After you upgrade the vManage NMS to software release 17.2.2.1, the NMS performance might be slow. [VIP-34338: This issue has been resolved.]
- When you try to attach an SNMP feature template to a configuration template that is attached to 1181 devices, the vManage message bus software might return a RecordTooLargeException error message. [VIP-34523: This issue has been resolved.]
- 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-34288: This issue has been resolved.]
- The vManage database of vEdge chassis numbers and UUIDs might contain entries for duplicate system IP addresses that are not associated with any chassis number or UUID. [VIP-35544: This issue has been resolved.]

Issues Fixed in Release 17.2.2.1

The following issues have been fixed in Viptela Software Release 17.2.2.1. This release is for vManage NMS only, so all the fixed issues address vManage NMS problems. The number following each issue is the bug number in the Viptela bug-tracking database.

vManage NMS

- In vManage NMS, the AAA feature template sends AAA user passwords in cleartext (ASCII) format. [VIP-34203: This issue has been resolved.]
- In the vManage REST APIs, the scrollId field is not included in the pageInfo section of the API output. [VIP-34222: This issue has been resolved.]

Issues Fixed in Release 17.2.2

The following issues have been fixed in Viptela Software Release 17.2.2. The number following each issue is the bug number in the Viptela bug-tracking database.

Configuration

- If you issue the **request daemon ncs restart** command from the vManage NMS, the command might not restart the network configuration (NCS) process. [VIP-31328: This issue has been resolved.]

Forwarding

- When you configure a TLOC extension and map BFD to queue 0, policer drops might occur on queue 0. [VIP-27601: This issue has been resolved.]
- When a DHCP server receives a large number of asynchronous DHCP_DISCOVER packets, there might be a delay before the server sends DHCP_OFFER packets. [VIP-28125: This issue has been resolved.]
- You might not be able to perform a traceroute operation on an IRB interface. [VIP-31563: This issue has been resolved.]
- On a vEdge 5000 router, you cannot enable TCP optimization by configuring the **tcp-optimization-enabled** command. [VIP-33011: This issue has been resolved.]

Interfaces

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

Policy

- In a data policy, when you set the action to "tloc", you may no longer be able to edit the policy. [VIP-33819: This issue has. been resolved.]

Routing Protocols

- 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.]
- 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 a collision occurs between IKE IPsec rekeying packets, the IPsec tunnel might go down and then come back up and traffic flow is interrupted. [VIP-30371: This issue has been resolved.]

- When a vEdge router has two TLOCs, a control to the vSmart controller might not be established. [VIP-30902: This issue has been resolved.]
- The output of the **show ipsec ike sessions** command might show the incorrect session state. [VIP-31590: This issue has been resolved.]
- A BGP peering session might not come up over the IKE/IPsec tunnel between a vEdge Cloud router on Azure and VNET. [VIP-33808: This issue has been resolved.]

System

- 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

- In the vManage Monitor ► Network screen, the detailed device information might be difficult to read because of how it is formatted. [VIP-11612: This issue has been resolved.]
- The vManage dashboard does not automatically refresh the state of the members of the vManage cluster even when their state changes. [VIP-26017: This issue has been resolved.]
- In the vManage Monitor ► Network ► Troubleshooting Ping pane, when you enter an IPv6 destination address, the ping operation might fail. [VIP-30720: This issue has been resolved.]
- When you upgrade the vManage NMS software to Release 17.2.0, errors might occur on the vManage server, including graphs being out of sync, tasks still appearing to be running ever after the server reboots, and incorrect status, such as a failure status, ever after the server is up and running. [VIP-31458: This issue has been resolved.]
- 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 Montior ► Network ► Troubleshooting might not work. [VIP-31576: This issue has been resolved.]
- The vManage configuration database might run out of memory. [VIP-32656: This issue has been resolved.]
- In vManage NMS, the validation failure message displayed when you try to update the software on a vEdge router does not provide sufficient detail about the cause of the failure. [VIP-32716: This issue has been resolved.]
- In vManage NMS, when you are configuring a collector in a cflowd policy, you cannot configure a transport type of transport_udp. [VIP-33046: This issue has been resolved.]
- The vManage web server may become unresponsive. This can be because the configuration database runs out of memory. [VIP-33649: This issue has been resolved.]
- When you invalidate a vSmart controller from the vManage server, an entry for the controller might still be present on the Configuration ► Devices screen. [VIP-33720: This issue has been resolved.]
- For Viptela controller devices are running Release 17.2.0, ZTP might not work. [VIP-34132: This issue has been resolved.]

Issues Fixed in Release 17.2.1

Viptela Software Release 17.2.1 was not released.

Issues Fixed in Release 17.2.0

The following issues have been fixed in Viptela Software Release 17.2.0. 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.]

DHCP

- An IP DHCPv6 client might not be able to get an IP address. [VIP-31583: THis issue has been resolved.]

Forwarding

- 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: This issue has been resolved.]
- 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.]

Interfaces

- A vEdge router acting as a VRRP primary might become a VRRP backup after receiving its own prior advertisement. [VIP-30956: This issue has been resolved.]

Security

- You might not be able to add a vSmart controller to the vManage server unless Netconf is enabled. [VIP-11688: This issue has been resolved.]
- The **show tunnel statistics ipsec** command displays no information about the count of inbound decrypted packets. [VIP-20637: This issue has been resolved.]
- In Releases 17.1 and earlier, the Viptela software negotiates the use of the SSH HMAC-MD5 algorithm and other weaker algorithms with its peers. [VIP-29170: 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.]

System

- When the vManage NMS is sending the vEdge list to the vSmart controller and vBond orchestrator, if an issue occurs, the vManage NMS provides no information to help pinpoint the problem. [VIP-11695: This issue has been resolved.]
- The **show log tail** command might not work. [VIP-30216: This issue has been resolved.]

vAnalytics Platform

- 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: This issue has been resolved.]

vManage NMS

- The vManage NMS might not validate or invalidate a certificate the first time you attempt the operation. [VIP-11809: This issue has been resolved.]
- The vManage NMS might not update the control connection status graph on the Monitor ► Network ► Control Connections screen after you add or remove a TLOC from the vEdge router. [VIP-17537: This issue has been resolved.]

- 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: This issue has been resolved.]
- When you upgrade a vManage cluster from Release 16.3.2 to Release 17.1.0, the configuration database might become incorrect. [VIP-27951: This issue has been resolved.]
- 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.]

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 17.2.0**
- **YANG and Enterprise MIB files for Releases 17.2.2 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.
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.
- The screenshots for the vManage NMS screens that are included in the vManage help files and other documentation articles might not match the vManage NMS software screens. We apologize for the inconvenience.

Requesting Technical Support

To request technical support, send email to support@viptela.com .

To provide documentation feedback or comments, send email to docs@viptela.com .

Revision History

Revision 1—Release 17.2.0, October 20, 2017
Release 17.2.1 was not released.
Revision 2—Release 17.2.2, December 19, 2017
Revision 3—Release 17.2.2.1, December 21, 2017
Revision 4—Release 17.2.3, February 15, 2018
Revision 5—Release 17.2.4, March 7, 2018
Revision 6—Release 17.2.5, April 24, 2018
Revision 7—Release 17.2.6, May 9, 2018
Revision 8—Release 17.2.7, June 29, 2018
Revision 9—Release 17.2.8, August 10, 2018

Revision 10—Release 17.2.9, November 10, 2018
Revision 11—Update, February 17, 2019

Revision 12—Update, November 19, 2019