

---

## Release Notes for SD-WAN Release 16.3

These release notes accompany Viptela Software Release 16.3, for Releases 16.3.0 through 16.3.3. The Viptela software runs on all Viptela devices, including vSmart controllers, vEdge routers, vBond orchestrators, and vManage NMSs.

Viptela Software Release 16.3

June 30, 2017

Revision 3

---

## Product Features

Below are the main product features in Viptela Software Release 16.3:

- **Access list support for multicast traffic**—On vEdge routers, you can create access lists that act on multicast traffic. You can configure multicast addresses in data prefix lists. You can apply all ACL actions except mirroring to multicast traffic. See [Localized Data Policy](#).
- **Advertise aggregate and specific routes into OMP**—You can aggregate routes before advertising them into OMP, and you can advertise a specific route into OMP instead of advertising all routes from a protocol. See [Configuring OMP and advertise](#).
- **Apply policies with overlapping site identifiers**—On the vSmart controller, you can apply site lists that contain overlapping site identifiers (site IDs) to different types of policies. Specifically, you can apply them to application-aware routing policy (**app-route-policy**), centralized control policy (**control-policy**), centralized data policy (**data-policy**), and centralized data policy used for cflowd flow monitoring (**data-policy** that includes a **cflowd** action and **apply-policy** that includes a **cflowd-template** command). However, site lists for a single one of these types of policy cannot contain overlapping site IDs. See [Configuring Application-Aware Routing](#), [Configuring Centralized Control Policy](#), [Configuring Centralized Data Policy](#), [Configuring Cflowd Traffic Flow Monitoring](#), and [apply-policy](#).
- **Cflowd sampling interval**—In a cflowd template, you can configure the sampling interval by setting how many packets to wait before creating a new flow. See [Configuring Cflowd Traffic Flow Monitoring and flow-sampling-interval](#).
- **Clearing vManage alarm improvements**—The vManage NMS leverages device state data to clear alarms periodically, to ensure that even if the vManage NMS misses the events due to control plane instability, any uncleared alarms are cleared.
- **CloudExpress service**—In vManage NMS, you can enable CloudExpress service to optimize access to cloud applications from Direct Internet Access (DIA) nodes in your Viptela overlay network, or nodes that reach the Internet through a secure web gateway. See [Enable and Configure CloudExpress Service](#), [View Application Performance with](#)



[CloudExpress Service](#), and [cloudexpress](#).

- **IEEE 802.1 port security**—IEEE 802.1x is a port-based network access control (PNAC) protocol that prevents unauthorized network devices from gaining access to wired networks (WANs), by providing authentication for devices that want to connect to a WAN. See [Configuring IEEE 801.1x and IEEE 80s.11i Authentication](#), [dot1x](#), and [radius](#).
- **IEEE 802.11i authentication**—IEEE 802.11i prevents unauthorized network devices from gaining access to wireless networks (WLANs). 802.11i implements Wi-Fi Protected Access II (WPA2) to provide authentication for devices that want to connect to a WLAN. See [Configuring IEEE 801.1x and IEEE 80s.11i Authentication](#), [radius](#), and [wlan](#).
- **Interface MTU maximum size**—The maximum MTU size for an interface is 2000 bytes. See [Configuring Network Interfaces](#) and [mtu](#).
- **IPv6 across transport networks**—On vEdge routers and vSmart controllers, you can configure IPv6 or dual stack (IPv4/IPv6) on WAN transport interfaces (in VPN 0) to allow the transport of IPv6 traffic. You can configure static IPv6 address or dynamic address retrieval from a DHCPv6 server. You can configure IPv6-specific access lists (ACLs) on vEdge routers to classify, count log, mirror, police, and set the traffic class on IPv6 traffic. See [Segmentation \(VPN\) Overview](#), [Configuring Segmentation \(VPNs\)](#), [Configuring Localized Data Policy for IPv6](#), [ipv6 address](#), [policy ipv6](#), and [Software Caveats](#), below.
- **Limit contents of admin-tech files**—You can limit the contents of admin-tech files, which you generate from the vManage Tools ► Operational Commands ► More Commands ► Admin Tech command or using the **request admin-tech** CLI command to help in troubleshooting and diagnoses issues on a Viptela device. Limiting the contents of this file can also decrease the time required to generate an admin-tech file. See [Operational Commands](#) and [request admin-tech](#).
- **Logging action in localized data policy (access lists)**—To log the packet headers for packets accepted or dropped by an ACL, include the **log** option in the **action** portion of the ACL. You can configure the frequency at which to log data packets, and you can configure the logging packets that are dropped because they do not match a service configured with an **allow-service** command. See [Configuring Localized Data Policy](#), [implicit-acl-logging](#), [log-frequency](#), [show app log flow-count](#), and [show app log flows](#).
- **Match packet loss priority (PLP) in policy**—You can match the PLP in application-aware routing policy, centralized data policy, and localized data policy (access lists). See [Configuring Application-Aware Routing](#), [Configuring Centralized Data Policy](#), [Configuring Localized Data Policy](#), [Forwarding and QoS Configuration Examples](#), and [match](#).
- **MIB changes**—The ifHighSpeed MIB object, in the standard interfaces MIB (IF-MIB) reports the speed for vEdge 2000 router interfaces with speeds greater than 1 Gigabit. In addition to reporting the link bandwidth, the ifSpeed and ifHighSpeed MIB objects also report the downstream bandwidth, if it is configured.



- **OMP route convergence improvements**—The convergence times for OMP routes have been improved so that OMP routes converge more quickly.
- **Perform parallel operations**—On vEdge routers you can perform the same operations, in parallel, from one or more vManage servers. You can perform the following operations in parallel: upgrade or activate a software image on a device, delete a software image from a device, set a software image to be the default image on a device, reboot a device, attach devices to a device template, detach devices from a device template, and change the variable values for a device template that has devices attached to it.
- **Pin application traffic to a TLOC**—For a TLOC action in a centralized data policy, if the configured TLOC is unavailable, you can configure the policy to drop data traffic which matches that policy rather than trying alternative TLOCs. See [Configuring Centralized Data Policy and action](#).
- **Prefer a path when tunnel interfaces do not meet application-aware routing SLA thresholds**—When no tunnel interface meets SLA thresholds, you can configure the path to use to send data traffic. See [Configuring Application-Aware Routing and action](#).
- **Preference for vManage connection**—On vEdge routers, you can configure the preferred tunnel interface to use to exchange traffic with the vManage NMS. See [vmanage-connection-preference](#) and *Software Caveats, below*.
- **Setting TLOC action in centralized control policy**—When you set a TLOC action in centralized control policy, you can configure the mechanism by which traffic is redirected to that TLOC. Configuring this mechanism also enables end-to-end tracking, which determines whether the ultimate traffic destination is reachable. When a sending router learns that a destination is unreachable, it can remove the corresponding route from its local route table. See [Configuration Centralized Control Policy and action](#).
- **Static NAT**—vEdge routers can run 1:1 static NAT on service-side interfaces. Please contact Customer Support before deploying this feature. See [Using a vEdge Router as a NAT Device](#), [Configuring Service-Side NAT](#), and [Service-Side NAT Configuration Example](#).
- **Stronger ciphers for vManage web servers**—vManage web server cipher have been strengthened. The servers support the following ciphers:
  - TLS\_DHE\_DSS\_WITH\_AES\_128\_GCM\_<wbr/>SHA256
  - TLS\_DHE\_DSS\_WITH\_AES\_256\_GCM\_<wbr/>SHA384
  - TLS\_DHE\_RSA\_WITH\_AES\_128\_GCM\_<wbr/>SHA256
  - TLS\_DHE\_RSA\_WITH\_AES\_256\_GCM\_<wbr/>SHA384
  - TLS\_ECDHE\_ECDSA\_WITH\_AES\_128\_<wbr/>GCM\_SHA256
  - TLS\_ECDHE\_ECDSA\_WITH\_AES\_256\_<wbr/>GCM\_SHA384



- TLS\_ECDHE\_RSA\_WITH\_AES\_128\_<wbr/>GCM\_SHA256
- TLS\_ECDHE\_RSA\_WITH\_AES\_256\_<wbr/>GCM\_SHA384

[See Deploy the vManage NMS.](#)

- **Suppress control traffic over low-bandwidth interfaces**—When the vEdge router with cellular (LTE) interface is deployed as a spoke and data tunnels are established in a hub-and-spoke manner, you can configure the cellular interface as a low-bandwidth interface. This configuration allows the vEdge router spoke site to synchronize all outgoing control packets. See [Configuring Cellular Interfaces](#) and [low-bandwidth-link](#).
- **vAnalytics platform**—From vManage NMS, you can access the vAnalytics platform for a high-level view of your entire overlay's network performance and availability, and application performance. vAnalytics platform also lets you drill down to data for a single carrier, tunnel, or application at a given moment in time. See [View Network Performance with vAnalytics Platform](#).
- **vEdge 100wm router**—The vEdge 100wm router delivers highly secure site-to-site data connectivity to small business and home offices. The vEdge 100wm router is a fixed-port-configuration router with a built-in LTE modem to connect to cellular networks and with a WLAN radio for access point functionality. The router supports IEEE 802.11a, 802.11b, 802.11g, 802.11n, and 802.11ac WLAN clients. See [vEdge 100wm Router](#).
- **vEdge Cloud multicore capability**—The vEdge Cloud virtualized vEdge router can use multiple cores to increase performance. With four vCPUs, the maximum tested throughput is 2 Gbps. No additional configuration is required to achieve this enhanced performance.
- **vEdge Cloud QoS**—You can configure egress scheduling and shaping for QoS on transport interfaces on vEdge Cloud routers. See [Configuring Localized Data Policy for IPv4](#), [class-map](#), [cloud-qos](#), and [cloud-qos-service-side](#).

## Command Changes

### New and Modified Configuration Commands

Command	Hierarchy	New	Modified	Comments
<a href="#">access-list</a>	policy ipv6, vpn 0 interface ipv6	X		
<a href="#">accounting-interval</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">acct-req-attr</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">action</a>	policy app-route-policy vpn-list		X	Add <b>backup-sla-preferred-color</b> action.



Command	Hierarchy	New	Modified	Comments
	sequence			
<a href="#">action</a>	policy data-policy sequence		X	Add <b>log</b> action. Add <b>strict</b> option to <b>set local-tloc</b> option.
<a href="#">action</a>	policy ipv6 access-list sequence	X		
<a href="#">advertise</a>	vpn omp	X		
<a href="#">allow-address</a>	container instance	X		On vContainer hosts.
<a href="#">allow-local-exit</a>	vpn cloudexpress	X		For <a href="#">CloudExpress service</a> .
<a href="#">allow-service</a>	vpn 0 interface tunnel-interface		X	Add support for DHCPv6.
<a href="#">applications</a>	vpn cloudexpress	X		For <a href="#">CloudExpress service</a> .
<a href="#">auth-fail-vlan</a>	vpn interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">auth-order</a>	vpn interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">auth-reject-vlan</a>	vpn interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">auth-req-attr</a>	vpn 0 interface dot1x			For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">channel</a>	wlan	X		For <a href="#">WLAN interfaces</a> .
<a href="#">channel-bandwidth</a>	wlan	X		For <a href="#">WLAN interfaces</a> .
<a href="#">class-map</a>	policy		X	Add support for multicast traffic and for vEdge Cloud routers.
<a href="#">cloud-qos</a>	policy	X		For vEdge Cloud routers.
<a href="#">cloud-qos-service-side</a>	policy	X		For vEdge Cloud routers.
<a href="#">cloudexpress</a>	vpn	X		For <a href="#">CloudExpress service</a> .
<a href="#">community</a>	snmp		X	Community names can include angle brackets (< and >).
<a href="#">control-direction</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">country</a>	wlan	X		For <a href="#">WLAN interfaces</a> .
<a href="#">das</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">data-security</a>	wlan interface	X		
<a href="#">default-vlan</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">description</a>	wlan interface	X		



Command	Hierarchy	New	Modified	Comments
<a href="#">direction</a>	vpn interface nat	X		On vEdge routers. For <a href="#">service-side static NAT</a> .
<a href="#">dns</a>	vpn		X	Add support for IPv6 DNS server addresses.
<a href="#">dot1x</a>	vpn interface	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">flow-sampling-interval</a>	policy cflowd-template	X		
<a href="#">guard-interval</a>	wlan	X		For <a href="#">WLAN interfaces</a> .
<a href="#">guest-vlan</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">host</a>	vpn		X	Add support for IPv6 addresses.
<a href="#">host-mode</a>	vp 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">host-policer-pps</a>	system		X	
<a href="#">implicit-acl-logging</a>	policy	X		
<a href="#">interface</a>	wlan	X		For <a href="#">WLAN interfaces</a> .
<a href="#">ipv6 address</a>	vpn 0 interface	X		
<a href="#">ipv6 dhcp-client</a>	vpn 0 interface	X		
<a href="#">ipv6 route</a>	vpn 0	X		
<a href="#">local-interface-list</a>	vpn cloudexpress	X		For <a href="#">CloudExpress service</a> .
<a href="#">log-frequency</a>	policy	X		
<a href="#">low-bandwidth-link</a>	vpn 0 interface tunnel-interface	X		
<a href="#">mac-authentication-bypass</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">match</a>	policy access-list, policy app-route-policy, policy data-policy		X	Add <b>plp</b> match condition.
<a href="#">match</a>	policy ipv6 access-list sequence	X		
<a href="#">max-clients</a>	wlan interface	X		For <a href="#">WLAN interfaces</a> .
<a href="#">mgmt-security</a>	wlan interface	X		For <a href="#">WLAN interfaces</a> .
<a href="#">mtu</a>	vpn interface		X	Maximum MTU size changed from 1804 bytes to 2000 bytes.
<a href="#">nas-identifier</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .



Command	Hierarchy	New	Modified	Comments
<a href="#">nas-ip-address</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">node-type</a>	vpn cloudexpress	X		
num-macs	bridge		X	Removed from CLI.
<a href="#">omp</a>	vpn	X		
<a href="#">overload</a>	vpn interface nat	X		On vEdge routers. For <a href="#">service-side static NAT</a> .
<a href="#">policer</a>	policy		X	Add support for multicast traffic.
<a href="#">policy ipv6</a>	Top level	X		
<a href="#">qos-map</a>	policy		X	Add support for multicast traffic.
<a href="#">qos-scheduler</a>	policy, policy qos-map		X	Add support for multicast traffic.
<a href="#">radius</a>	system		X	Add support for <a href="#">IEEE 802.1x and 802.11i authentication</a> .
<a href="#">radius-servers</a>	system radius server, vpn 0 interface dot1x, wlan interface		X	Add support for <a href="#">IEEE 802.1x authentication</a> and <a href="#">IEEE 802.11i authentication</a> .
<a href="#">reauthentication</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">redistribute</a>	vpn router bgp address-family		X	Add <b>natpool-outside</b> option.
<a href="#">redistribute</a>	vpn router ospf		X	Add <b>nat</b> and <b>natpool-outside</b> options.
<a href="#">rewrite-rule</a>	policy		X	Add support for multicast traffic.
<a href="#">set tloc-action</a>	policy control-policy sequence action accept	X		On vSmart controllers.
<a href="#">shaping-rate</a>	policy		X	Cannot configure on VLAN interfaces.
<a href="#">ssid</a>	wlan interface	X		For <a href="#">WLAN interfaces</a> .
<a href="#">static</a>	vpn interface nat	X		On vEdge routers. For <a href="#">service-side static NAT</a> .
<a href="#">tcp-mss-adjust</a>	vpn interface		X	Change maximum TCP MSS from 1460 bytes to 1960 bytes.
<a href="#">technology</a>	vpn interface cellular	X		On vEdge routers. In Releases 16.3.2 and later.
<a href="#">timeout inactivity</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">vbond</a>	system		X	
<a href="#">vbond-as-stun-server</a>	vpn 0 interface tunnel-interface	X		



Command	Hierarchy	New	Modified	Comments
<a href="#">vmanage-connection-preference</a>	vpn 0 interface tunnel-interface	X		
<a href="#">wake-on-lan</a>	vpn 0 interface dot1x	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">wlan</a>	Top level	X		For <a href="#">WLAN interfaces</a> and <a href="#">IEEE 802.11i authentication</a> .
<a href="#">wpa-personal-key</a>	wlan interface	X		For <a href="#">IEEE 802.11i authentication</a> .

## New and Modified Operational Commands

Command	New	Modified	Comments
<a href="#">clear app log flow-all</a>	X		
<a href="#">clear app log flows</a>	X		
<a href="#">clear cloudexpress computations</a>	X		For <a href="#">CloudExpress service</a> .
<a href="#">clear dot1x client</a>	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">clear ipv6 dhcp state</a>	X		
<a href="#">clear ipv6 neighbor</a>	X		
<a href="#">clear ipv6 policy</a>	X		
<a href="#">clear policer statistics</a>	X		
<a href="#">ping</a>		X	Add support for IPv6 host addresses.
<a href="#">request admin-tech</a>		X	
<a href="#">show app log flow-count</a>	X		
<a href="#">show app log flows</a>	X		
<a href="#">show cloudexpress applications</a>	X		For <a href="#">CloudExpress service</a> .
<a href="#">show cloudexpress local-exits</a>	X		For <a href="#">CloudExpress service</a> .
<a href="#">show control connections</a>		X	Add IPv6 addresses and ports to output.
<a href="#">show control local-properties</a>		X	Add IPv6 addresses and ports to output.





Command	New	Modified	Comments
<a href="#">show control summary</a>		X	Add IPv6 addresses and ports to output.
<a href="#">show dot1x clients</a>	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">show dot1x interfaces</a>	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">show dot1x radius</a>	X		For <a href="#">IEEE 802.1x authentication</a> .
<a href="#">show ip mfib stats</a>		X	
<a href="#">show ip nat interface</a>		X	Add output fields for static NAT.
<a href="#">show ip nat interface-statistics</a>		X	Add output fields for static NAT.
<a href="#">show ip routes</a>		X	Add <b>natpool-omp</b> and <b>natpool-service</b> options.
<a href="#">show ipv6 dhcp interface</a>	X		
<a href="#">show ipv6 fib</a>	X		
<a href="#">show ipv6 interface</a>	X		
<a href="#">show ipv6 neighbor</a>	X		
<a href="#">show ipv6 policy access-list-associations</a>	X		
<a href="#">show ipv6 policy access-list-counters</a>	X		
<a href="#">show ipv6 policy access-list-names</a>	X		
<a href="#">show ipv6 policy access-list-policers</a>	X		
<a href="#">show ipv6 routes</a>	X		
<a href="#">show orchestrator valid-ymanage-id</a>	X		On vBond orchestrators, in Releases 16.3.1 and later.
<a href="#">show omp tllocs</a>		X	Add IPv6 addresses to output.
<a href="#">show policer</a>		X	
<a href="#">show system statistics</a>		X	Add 802.1x statistics. In Releases 16.3.2 and later, display BFD PMTU statistics.
<a href="#">show system status</a>		X	In Release 16.3.2 and later, add System State field to command output.
<a href="#">show tunnel statistics</a>		X	Add <b>stats-ipsec</b> option. In Releases 16.3.2 and later, add <b>bfd</b> option, and display BFD PMTU statistics.



Command	New	Modified	Comments
<a href="#">show wlan clients</a>	X		For <a href="#">WLAN interfaces</a> .
<a href="#">show wlan interfaces</a>	X		For <a href="#">WLAN interfaces</a> .
<a href="#">show wlan radios</a>	X		For <a href="#">WLAN interfaces</a> .
<a href="#">tcpdump</a>		X	
<a href="#">traceroute</a>		X	Add support for IPv6 host addresses.

## Upgrade to Release 16.3

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, and it can also be higher.

Because of software changes in Release 16.3, you must modify the router configuration as follows before you upgrade:

- 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. 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, ensure that the interface names in the router configurations are correct.

To upgrade to Release 16.3 from Release 16.1 or Release 16.2:

1. In vManage NMS, select the Maintenance ► Software Upgrade screen.
2. Upgrade the controller devices to Release 16.3 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.

To upgrade to Release 16.3 from Release 15.4, you must first upgrade the vManage NMS to either Release 16.2.2 or Release 16.2.3.

When the upgrade is complete, clear the cache on the vManage browser so that you can use the Release 16.3 features.

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 16.3, 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, 16.3 and 16. are examples of major releases. Releases 16.3.0 and 16.2.0 denote the initial releases, and Releases 16.3.2 and 16.2.1 are maintenance releases.

If you are upgrading a Viptela controller—vManage NMS, vSmart controller, or vBond orchestrator—to Release 16.3.0 or later, and if you are switching from the E1000 to the VMXNET3 network adapter, you must perform the upgrade and switch in the following order:

1. Upgrade the software image to the desired version.
2. Shut down the Viptela controller.
3. Remove the existing E1000 network adapters.
4. Add the new VMXNET3 adapters to the controller, as explained in the appropriate [Create VM Instance on ESXi](#) article.

Note: You cannot add a mix of E1000 and VMXNET3 network adapters to any Viptela controller.

---

## Downgrade from Release 16.3 to Release 16.2

You cannot downgrade from Release 16.3 to any version of Release 16.2. If you do so, the vManage NMS reports a null pointer exception for statistics settings, and as a result, the vManage NMS collects no statistics from the vEdge routers in the network. If you need to downgrade to Release 16.2, contact Customer Support for assistance.

---

## 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 router, 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

- As of Release 16.3.0, the vEdge 100wm router is certified for use in the United States only. This 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.
- 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).
- For cellular interface profile, the profile number can be 0 through 15. Profile number 16 is reserved, and you cannot modify it.

## 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.

## Security

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

## SNMP

- In Release 16.3, 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.
- 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.
- 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.

## 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.

---

## Outstanding Issues

The following are outstanding issues in Viptela Software Release 16.3. 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

- On a vEdge router, the **show cellular status** command does not reflect the current operational status of the modem. [VIP-13808]
- If you have configured a profile for a cellular interface, when the interface is active, you cannot delete the profile. [VIP-20327]





- If you configure IPv6 on a cellular interface, the control connections might go down and come back up continuously. [VIP-21970]
- You cannot configure profile 16 in the [interface cellular0 profile](#) command.

## CloudExpress Service

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

## Configuration and Command-Line Interface

- The **show interface detail | include *string*** command attempts to paginate the output, displaying the More prompt, even when there is nothing to paginate. [VIP-2057]
- An IRB interface might remain up even if all interfaces in that bridge are in the link-down state. [VIP-23307]
- If you shut down and then reactivate an interface that has no connectivity to the Internet, the tunnels might not automatically come up. If you configure a TLOC preference on the interface, the tunnels come up. [VIP-23442]
- 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]

## Forwarding

- The flow collector on a vEdge router might fail to collect cflowd packet statistics. [VIP-11088]
- For IEEE 802.1x, you cannot configure a RADIUS server for MAC authentication bypass (MAB). [VIP-18492]
- 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]

- In application-aware routing policy, the `salesforce_chatter`, `oracle_rac`, and `google_photos` applications might not be classified properly. [VIP-21866]
- On vEdge routers, the routing table might not validate that a route is present in the forwarding table. [VIP-23411]
- After an interface goes down and then comes back up, dynamic ARP learning might fail. [VIP-26215]
- On a vEdge Cloud router, traffic shaping through QoS might not work properly. [VIP-26557]
- A centralized policy that is pushed from the vSmart controller to the vEdge routers might not be applied on the routers. [VIP-27046]
- 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]
- 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]

## Kernel

- When a vEdge router reboots and the control down flag is set, the `vdaemon` process might an Up event indicating that a configuration commit operation has occurred even though there was no configuration change. As a result, the vManage NMS pushes a device configuration to the vEdge router. [VIP-22395]

## 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.



## PPPoE

- When you are using PPPoE on a vEdge100m router, when you initially connect to the PPP WAN interface, the interface receives an IP address. But when you unplug the PPP interface and then plug it back in, we do not get an IP address. As a workaround, either reboot the DSL modem or reset the interface from the router's CLI. [VIP-23332]

## 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 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]

## Security

- The **show tunnel statistics ipsec** command displays no information about the count of inbound decrypted packets. [VIP-20637]
- When you use the TLS security protocol for the control plane, a control connection might not be established on a second TLOC when the vEdge router's TLOCs are in the same subnet. [VIP-27200]

## 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]



- If you use the loopback interface as the source interface for a cflowd collector, the Connection State on a vEdge router running Release 16.2.9 might show as false, while it shows true on a router running Release 16.2.10. [VIP-24770]
- After you configure NTP and TACACS from the CLI, the configuration process (confd) might crash. [VIP-28440]
- 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]

## vAnalytics Platform

- When you drill down from the vAnalytics dashboard, you might not be able to change the time history range that is set on the dashboard. [VIP-22749]
- The vAnalytics timestamps are in UTC, but the timezone is not indicated in the timestamp. [VIP-22750]

## vEdge Hardware

- On a vEdge 2000 router, when you remove an SFP, trace messages might be displayed. These messages do not affect the router, and you can ignore them. [VIP-3585]
- 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]

## vManage NMS

- vManage NMS might fail to make a Netconf connection to devices when attempting to collect statistics. [VIP-11087]



- The vManage Site Health pane might not list partially connected sites. [VIP-18957]
- In a three vManage cluster, if one vManage server is down for a prolonged time period (vManage1), and if the connection between the other two vManage servers goes down, it is possible that connection between the latter two vManage servers will not recover until the first vManage server (vManage1) is brought back up. [VIP-19373]
- From the CLI on the vManage NMS, when you issue the **request nms configuration-db restore** command, the output might indicate a failure in stopping NMS service when in fact it succeeded (the message "Could not stop NMS" followed by "Failed to restore the database"). If this occurs case, simply, re-issue the **request nms configuration-db restore** command. [VIP-19656]
- 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]
- In the vManage Monitor ► Device Real-Time Data drop-down, the interface queue command is missing. [VIP-21796]
- 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]
- In a vManage template, if you enter a password that starts with the string \$8, the password might not be pushed properly to the vEdge routers.[VIP-24131]
- 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]
- On a vManage NMS running Release 16.2.9.1, pushing device configuration templates to a large number of devices might take a long time. For example, with three vManage NMSs and about 800 devices, it might take 60 minutes to



push the templates to all the devices. [VIP-26185]

- 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]
- When you renew the vSmart certificate from the vManage NMS, the new signed certificate might not be installed automatically on the vSmart controller. [VIP-27791]
- The vManage server might not process events received from vEdge routers. [VIP-28312]
- The vManage NMS might show that a vEdge router is in sync-pending state even after the router is reachable on the network. [VIP-28663]
- The /client/activity/summary REST call might time out. [VIP-28737]

---

## Fixed Issues

---

### Issues Fixed in Release 16.3.3

---

The following issues have been fixed in Viptela Software Release 16.3.3. 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 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 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.]
- 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.]
- 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.]
- When you configure a tunnel interface to receive its address from a DHCP server, NAT port forwarding might not work on that interface. [VIP-28093: This issue has been resolved.]

## OMP

- When you push a large policy from the vManage NMS to a vSmart controller, it might take a long time for the policy to take effect. [VIP-22115: This issue has been resolved.]
- If a vEdge routers moves from one vSmart controller to another, and if a topology change occurs (such as a new or updated control policy or a route withdrawal), this change might not take effect on the vEdge router because it is in graceful restart mode with the previous vSmart controller. [VIP-23362: This issue has been resolved.]

## Policy

- In Release 16.3.2, you cannot apply two different data policies, one in each direction, for same site. [VIP-24010: This issue has been resolved.]
- 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.]

## 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.]
- When you invalidate a vSmart controller, connections between the controller and the vEdge routers might go down and



then come back up again. This flapping might interfere with OMP graceful restart. [VIP-22263: 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.]
- On a vBond orchestrator, an interface in the transport VPN (VPN 0) might not be able to renew its DHCP address. [VIP-23780: This issue has been resolved.]
- 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.]

## vEdge Hardware

- After you activate Release 16.2.10 on a vEdge router, the configuration database might become corrupted, and the router might continuously reboot. [VIP-25731: 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.]
- The vManage server log files might contain a large number of duplicate BFD record messages. [VIP-22098: 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.]
- The vManage dashboard might show the incorrect certificate expiration even after certificate has been renewed. [VIP-23779: This issue has been resolved.]





- When you attach a vEdge policy template that contains variables to a device template, during the CSV import action the policy variables might not be populated with the values from the CSV file. As a workaround, manually set the policy values and then import the CSV file. [VIP-23862: This issue has been resolved.]
- In the vManage NMS menu, the vAnalytics icon might not be displayed. [VIP-28891: This issue has been resolved.]

## Wireless WANs

- For a vEdge100m router, ZTP fails if the technology configured for the cellular interface is set to auto. To correct this problem, set the technology to be lte. Do this either on the vManage VPN-Interface-Cellular configuration template or with the **vpn 0 interface cellular technology** CLI command. [VIP-23988: This issue has been resolved.]

---

## Issues Fixed in Release 16.3.2

The YANG and enterprise MIB files have been modified for [Release 16.3.2](#).

The following issues have been fixed in Viptela Software Release 16.3.2. The number following each issue is the bug number in the Viptela bug-tracking database.

### AAA

- When you attach the same AAA feature template to two vEdge routers and then log in to each of the routers, on one router you might be placed into the correct usergroup, while on the second router you might be placed into a different usergroup. [VIP-22538: This issue has been resolved.]

### Cellular Interfaces

- When you enable a generic cellular interface, the **show cellular modem** command might report that the interface is in low-power mode (LPM). [VIP-21468: This issue has been resolved.]

### CloudExpress Platform

- CloudExpress platform does not generate SNMP traps. [VIP-21715: This issue has been resolved.]

### Configuration and Command-Line Interface

- When you delete a tunnel interface and then reconfigure it, the tunnel interface might not come up. [VIP-23069: This issue has been resolved.]

### Forwarding



- The **show policy service-path** command shows all paths even when you use data policy to restrict the path. [VIP-22451: This issue has been resolved.]
- When Viptela devices are establishing BFD connections, if you have configured the BFD hello interval to a value greater than the default value of 1000 milliseconds (1 second), it might take 3-4 minutes for the connections to come up. During this time, they remain in the "init" state. [VIP-22500: This issue has been resolved.]
- Ping attempts to the IP address of an IRB interface might randomly fail. [VIP-22721: This issue has been resolved.]
- When you apply vSmart policy from the vManage NMS and then update the policy, the new policy might not appear on the vEdge router. If you reboot the router to clear the old policy, when the vEdge router comes back up, the forwarding management process (fpm) might crash when it receives the updated policy. [VIP-23324: This issue has been resolved.]

## Kernel

- After you upgrade from Release 16.2.7 to Release 16.3.0, the vEdge router might report the following error when it is rebooting: "baudrate not defined, BusyBox v1.22.1 multi-call binary". [VIP-23146: This issue has been resolved.]

## OMP

- A vEdge router might continue to attempt to establish an OMP peering session with a vSmart controller that has been invalidated. [VIP-22385: This issue has been resolved.]

## Platform

- If you are using ESXi virtual machines (VMs) for the vManage NMSs in a cluster, some vManage operations might fail, because the time of the NMSs in the vManage cluster is not the same. This problem occurs because the VMs occasionally synchronize their clock with the ESXi host.

The following vSphere operations can reset the clock on a VM, moving the time either forwards or backwards, to match the clock on the VM host:

- Resume a VM from a suspended state.
- Take or restore a snapshot.
- Migrate a VM using vMotion.
- Reduce the size a virtual machine's disk.
- Reboot a VM.
- Restart the VMware Tools service on a VM. [VIP-22472: This issue has been resolved.]



## Policy

- The preferred-color configuration might be ignored even when the corresponding path meets the defined SLA class. [VIP-23445: This issue has been resolved.]

## System

- A kernel panic might cause a vEdge router to silently reboot. [VIP-16233: This issue has been resolved.]
- When one or more processes (daemons) on a vEdge router go down and the router reboots many times within a short period of time, the router sends the notification "process restarted" to the vManage NMS. Starting with Release 16.3.2, in this situation, the router sends the notification "Reboot (reason:...) aborted...too many reboots". To view the reboot status from the CLI, look at the System State field in the output of the [show system status](#) command. [VIP-23338: This issue has been resolved.]

## vEdge Hardware

- On vEdge 2000 routers, an unexpected reboot might occur, and the console has no information to indicate the cause of the reboot. [VIP-17514: This issue has been resolved.]
- If the /tmp/xml directory on a vEdge router contains too many statistics files, the router might run out of memory. [VIP-21300: This issue has been resolved.]
- Under normal operating conditions, a vEdge 100 router might show Yellow temperature alarms for CPU junction and Board temperature sensors. [VIP-22137: This issue has been resolved.]

## vAnalytics Platform

- If a user enters the incorrect login credentials for the vAnalytics platform, the user "admin" is unable to access the vAnalytics platform. [VIP-22633: This issue has been resolved.]
- If no DPI data is available, when you attempt to view the vAnalytics dashboard, the vAnalytics platform might crash. [VIP-22735: This issue has been resolved.]

## vManage NMS

- In the vManage template push page, you might not be able to click on the different device templates. [VIP-22732: 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.]

- In vManage NMS, the real-time display of tunnel interface statistics in octets might show the number of packets instead of the number of bytes (octets). [VIP-22990: This issue has been resolved.]
- The vManage NMS collects the statistics displayed in the real-time display of tunnel interface octet statistics pane every 10 seconds, but it might not display the statistics for up to 10-20 minutes. [VIP-23016: This issue has been resolved.]
- In vManage NMS, when you upload a CSV file to a list of variables, the input values are recorded and are displayed in the Edit Device Details pop-up. However, the table view of the template variable page shows many "--" characters, and exporting the variables file from this table view might result in empty cells. [VIP-23076: This issue has been resolved.]
- When you copy the interface configuration from a router running Release 16.2 to one running Release 16.3, the configuration might no longer work. [VIP-23109: This issue has been resolved.]
- In the vManage OSPF feature configuration template, you cannot configure an OSPF NSSA area. [VIP-23205: This issue has been resolved.]
- In the vManage OSPF feature configuration template, the OSPF area ID needs to be device specific. [VIP-23206: This issue has been resolved.]
- In the vManage logging feature template for a vSmart router, if you select a VPN other than 0 or 512, the error message refers to a RADIUS server rather than to a syslog server. [VIP-23246: This issue has been resolved.]
- In the vManage OMP feature template, the aggregate and network options are not device-specific. [VIP-23278: This issue has been resolved.]
- The vManage NMS does not support VPN 512 for the vEdge 100b router. [VIP-23279: This issue has been resolved.]
- The vManage NMS might not display statistics, and on the statistics screens it might show the message, 'Server:Unknown Error'. [VIP-23283: This issue has been resolved.]
- From a vManage NMS running Releases 16.2 or later, you cannot set the default software version on a Viptela device running Release 15.4. [VIP-23446: This issue has been resolved.]
- In the vManage Monitor ► Network screen, you might not be able to scroll to the bottom of the list of devices. [VIP-23446: This issue has been resolved.]



- In vManage NMS, the search function might not work if the search string contains digits. [VIP-23723: This issue has been resolved.]
- When you push a policy from the vManage NMS to the vSmart controller, an "unknown command" exception for NCS might occur on the vManage NMS. [VIP-23804: This issue has been resolved.]

## Wireless LANs

- When you configure wpa/wpa2-personal data security on a vap interface, the **show wlan clients** command might display the wrong data security value. [VIP-22497: This issue has been resolved.]

---

## Issues Fixed in Release 16.3.1

Release 16.3.1 was not released.

---

## Issues Fixed in Release 16.3.0

The following issues have been fixed in Viptela Software Release 16.3.0. The number following each issue is the bug number in the Viptela bug-tracking database.

### Cellular Interfaces

- When you perform a ZTP operation over a cellular interface, the state of the cellular interface might change to **no shutdown**. When this occurs, the control connections to the vManage NMS and the vSmart controllers, and BFD sessions come up. Then, when the device template is pushed to the router, the push operation fails because the template is out of sync with the state of the router. As a workaround, edit the configuration on the router to change the cellular interface's state to **shutdown**. This change allows the device template push operation to succeed. [VIP-22274: This issue has been resolved.]

### Configuration and Command-Line Interface

- The **ping** command might not work if you specify an IRB source interface as it appears in the interface table, for example irb2. The ping command works if you specify the interface name with a dot, for example, irb.2. [VIP-17075: This issue has been resolved.]

### DHCP

- Devices at a site behind a vEdge router might not be able receive IP addresses from a DHCP server, because their DHCP requests contain more options than the router supports. [VIP-20652: This issue has been resolved.]

### Forwarding



- The frequency for BFD-based PMTU discovery can use more than half of the bandwidth on low-speed links. [VIP-19887: This issue has been resolved.]
- On a vEdge router running Release 16.2.2, if you disable PMTU discovery, the **show system statistics** command output might show a large number of errors. [VIP-22145: This issue has been resolved.]
- For certain packet sizes, the vEdge router might send packets with incorrect ICMP checksums in ping requests. For example, packet sizes of 170 and 1300 bytes work fine, but packets of 171 and 1301 bytes end up with incorrect checksums in the response. [VIP-22174: This issue has been resolved.]

## Kernel

- When a user is connected via an SSH session to a vEdge router, they might be able enter a magic SysRq key or issue a Break sequence from the keyboard. Entering one of these sequences allows a user to reboot the router or perform other low-level operations on the router. [VIP-20386: This issue has been resolved.]

## Policy

- You cannot choose which TLOC is used to build the single connection to the vManage NMS. [VIP-8424: This issue has been resolved.]

## Routing Protocols

- In the **show ip routes** command, if you enter an IP prefix and a prefix length, the output displays all ECMP routes. If you enter just an IP prefix, the output displays only one route. [VIP-19872: This issue has been resolved.]

## SNMP

- When you configure loopback interfaces on a vSmart controller or a vManage NMS, performing an snmpwalk operation over VIPTELA-OPER-VPN.mib might fail. [VIP-16980: This issue has been resolved.]
- Log files do not hide TACACS secret keys or SNMP community strings. [VIP-19750: This issue has been resolved.]
- The ifLastChange MIB object returns how long before the present time an interface event occurred instead of the time between when the system came up and when the interface entered its current operational state. [VIP-19915: This issue has been resolved.]
- You cannot include the angle bracket characters (< and >) in SNMP community strings. [VIP-20399: This issue has been resolved.]



- In a vManage SNMP feature template, if you create a Trap Group but do not specify the list of modules, no error message is displayed but the deployment of the template fails. [VIP-20404: This issue has been resolved.]

## System

- A vEdge router might report that a fan has failed even when the fan is running. [VIP-12064: This issue has been resolved.]
- Moving a device from staging to valid might reset the control connections. [VIP-15009: This issue has been resolved.]
- The kernel might crash, with the message "Watchdog: Timer started; sleeping for 600 secs..." [VIP-16463: This issue has been resolved.]
- The Gigabit Ethernet and 10-Gigabit Ethernet interfaces on a vEdge router might place the following message in the /var/log/vdebug file: "Failed to set new phy settings (errno:122)". [VIP-19357: This issue has been resolved.]
- The **show interface** command lists GRE interfaces as being half-duplex. [VIP-19863: This issue has been resolved.]

## vEdge Hardware

- You might not be able to upgrade vEdge 1000 routers that are running Release 15.4.4 because of a firmware issue. [VIP-16902: This issue has been resolved.]
- A vEdge 2000 router might display the following error message on the console: "runsv: syslogd: fatal: cannot start ./run: Exec format error". [VIP-17788: This issue has been resolved.]
- On the vEdge 100m router, if you upgrade from Release 16.1.1 to Release 16.2.2 using the **request software install** command, the upgrade might fail and display the error "mount: mounting /dev/mmcblk0p1 on /mnt failed: Device or resource busy". [VIP-20387: This issue has been resolved.]
- On a vEdge 2000 router, if you install a 10-GB SFP into an 1-GB PIM, the router reboots continuously. [VIP-21888: This issue has been resolved.]

## vManage NMS

- In vManage ► Monitor ► Network ► Overview ► Environment, the Y axis on the CPU utilization graph shows the load average instead of the CPU utilization. [VIP-8114: This issue has been resolved.]
- Memory usage is shown as raw data, not in percentages. [VIP-8118: This issue has been resolved.]



- When you attach a template to the device, but the device rejects the template because of incorrect variable values and reverts to the previous configuration, the vManage NMS does not indicate the nature of the configuration problem. [VIP-8522: This issue has been resolved.]
- The vManage-server.log file in the nms directory contains a large amount of logging information, and some of the logs might not be useful for debugging issues. [VIP-11014: This issue has been resolved.]
- When the default value for a command option changes in a newer software version, the changes might not be reflected in configuration templates that you created on older software versions. However, the new default value is what is used in the newer software version. [VIP-12965: This issue has been resolved.]
- When displayed a DPI flow's packets, the vManage NMS might sort them incorrectly. [VIP-13832: This issue has been resolved.]
- In a vManage cluster that has two vManage servers, if you remove one of them, the remaining vManage server is no longer able to provide any management services. [VIP-14314: This issue has been resolved.]
- When you configure a vEdge router to be in more than one device group, the vManage screens that allow you to select by Device Groups might list all the groups as a single group instead of as separate groups. [VIP-14480: This issue has been resolved.]
- After you add a third vManage NMS to a cluster, the NMS services might become unresponsive and the web GUI might go down. [VIP-15455: This issue has been resolved.]
- After you upgrade the vManage NMS to Release 16.1.2, the Dashboard might not display any statistics data. [VIP-17462: This issue has been resolved.]
- When you use vManage templates to change the configuration, the configuration differences might not display when you click the Config diff tab. [VIP-17686: This issue has been resolved.]
- In the software version drop-down in vManage Maintenance ► Software Upgrade screen, the software versions are sorted by time, not by version number. [VIP-17942: This issue has been resolved.]
- The vManage NMS does not raise an alarm periodically prior to when a certificate is about to expire. [VIP-17957: This issue has been resolved.]
- The vManage NMS does not raise an alarm when a certificate expires. [VIP-17958: This issue has been resolved.]





- The information displayed by the vManage Monitor ► Network ► Real Time ► Cflowd Flows commands might not be consistent with the information displayed by the **show app cflowd flows** CLI command. [VIP-18133: This issue has been resolved.]
- After you upgrade to Release 16.2.1, the banner feature template might no longer recognize quotation marks that are included in the banner. [VIP-18509: This issue has been resolved.]
- A netadmin user might not have all administrative-level capabilities for operations such as API calls and setting the default software version. [VIP-18512: This issue has been resolved.]
- You might not be able to open log files on the vManage NMS because socket connections to the configuration database are closed. [VIP-18936: This issue has been resolved.]
- The vManage NMS might not recognize the serial number of a vEdge router that is already present in the network so might be unable to connect to the router. [VIP-18990: This issue has been resolved.]
- If a user tries add a vSmart controller that is not actually present in the network, the vManage NMS might add it to the network even after the vManage NMS displays an error message indicating that the device is not present. [VIP-19155: This issue has been resolved.]
- When you push a device template in which the chassis ID/UUID field contains a whitespace, an error message might incorrectly indicate that the device's personaility cannot be determined. [VIP-19810: This issue has been resolved.]
- Some vManage charts might display only a maximum of six data series. [VIP-19849: This issue has been resolved.]
- The **show policy data-policy-filter** CLI command shows the correct information for a policy in a VPN, but the vManage Policy Data Policy Filters real-time command on the Monitor ► Network screen might show incorrect information about the policy. [VIP-19862: This issue has been resolved.]
- In the vManage DPI flow output, the column "Entry Time" reports when the flow is logged, not when the flow started. [VIP-20497: This issue has been resolved.]
- On a vManage NMS running Release 16.2.3, you might not be able to activate policy changes. [VIP-20629: This issue has been resolved.]
- The vManage configuration templates might not be able to process CSV files. [VIP-22097: This issue has been resolved.]



- The vManage NMS might not be able push a template that contains a large number of variables to large number of devices.[VIP-22099: 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 16.3.0](#)
- [YANG and enterprise MIB files for Release 16.3.2](#)

---

## 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.

---

## 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 16.3.0, December 23, 2016

Release 16.3.1 was not released.

Revision 2—Release 16.3.2, February 23, 2017

Revision 3—Release 16.3.3, June 30, 2017

