



Configuring Resilient Ethernet Protocol

- [Information About Resilient Ethernet Protocol](#), on page 1
- [How to Configure Resilient Ethernet Protocol](#), on page 6
- [Monitoring Resilient Ethernet Protocol Configurations](#), on page 15
- [Additional References for Resilient Ethernet Protocol](#), on page 16
- [Feature History for Resilient Ethernet Protocol](#), on page 16

Information About Resilient Ethernet Protocol


Resilient Ethernet Protocol (REP) is a Cisco proprietary protocol that provides an alternative to Spanning Tree Protocol (STP) to control network loops, handle link failures, and improve convergence time. REP controls a group of ports connected in a segment, ensures that the segment does not create any bridging loops, and responds to link failures within the segment. REP provides a basis for constructing more complex networks and supports VLAN load balancing.



Note The feature is supported on Cisco Catalyst Series Switches with the Network Essentials license.


REP segment is a chain of ports connected to each other and configured with a segment ID. Each segment consists of standard (non-edge) segment ports and two user-configured edge ports. A switch can have no more than two ports that belong to the same segment, and each segment port can have only one external neighbor. A segment can go through a shared medium, but on any link, only two ports can belong to the same segment. REP is supported only on Trunk ports.

The figure below shows an example of a segment consisting of six ports spread across four switches. Ports E1 and E2 are configured as edge ports. When all ports are operational (as in the segment on the left), a single port is blocked, shown by the diagonal line. This blocked port is also known as the Alternate port (ALT port). When there is a failure in the network, the blocked port returns to the forwarding state to minimize network disruption.

Figure 1: REP Open Segment

The segment shown in the figure above is an open segment; there is no connectivity between the two edge ports. The REP segment cannot cause a bridging loop, and you can safely connect the segment edges to any network. All hosts connected to switches inside the segment have two possible connections to the rest of the network through the edge ports, but only one connection is accessible at any time. If a failure occurs on any segment or on any port on a REP segment, REP unblocks the ALT port to ensure that connectivity is available through the other gateway.

The segment below is a closed segment, also known as Ring Segment, with both edge ports located on the same router. With this configuration, you can create a redundant connection between any two routers in the segment.


Figure 2: REP Ring Segment

REP segments have the following characteristics:

- If all ports in a segment are operational, one port (referred to as the ALT port) is in the blocked state for each VLAN. If VLAN load balancing is configured, two ALT ports in the segment control the blocked state of VLANs.
- If a port is not operational, and cause a link failure, all ports forward traffic on all VLANs to ensure connectivity.
- In case of a link failure, alternate ports are unblocked as quickly as possible. When the failed link is restored, a logically blocked port per VLAN is selected with minimal disruption to the network.

You can construct almost any type of network based on REP segments.

In access ring topologies, the neighboring switch might not support REP as shown in the figure below. In this case, you can configure the non-REP facing ports (E1 and E2) as edge no-neighbor ports. The edge no-neighbor port can be configured to send an STP topology change notice (TCN) towards the aggregation switch.

Figure 3: Edge No-Neighbor Ports

REP has these limitations:

- You must configure each segment port; an incorrect configuration can cause forwarding loops in the networks.
- REP can manage only a single failed port within the segment; multiple port failures within the REP segment cause loss of network connectivity.
- You should configure REP only in networks with redundancy. Configuring REP in a network without redundancy causes loss of connectivity.
- REP is not supported in a Cisco StackWise Virtual solution.

Link Integrity

REP does not use an end-to-end polling function between edge ports to verify link integrity. It implements local link failure detection. The REP Link Status Layer (LSL) detects its REP-aware neighbor and establishes connectivity within the segment. All the VLANs are blocked on an interface until the neighbor is detected. After the neighbor is identified, REP determines which neighbor port should become the alternate port and which ports should forward traffic.

Each port in a segment has a unique port ID. The port ID format is similar to that used by the spanning tree algorithm: a port number (unique on the bridge) associated to a MAC address (unique in the network). When a segment port is coming up, its LSL starts sending packets that include the segment ID and the port ID. The port is declared as operational after it performs a three-way handshake with a neighbor in the same segment.

A segment port does not become operational if:

- No neighbor has the same segment ID.
- More than one neighbor has the same segment ID.
- A neighbor does not acknowledge a local port as a peer.

Each port creates an adjacency with its immediate neighbor. After the neighbor adjacencies are created, the ports negotiate with each other to determine the blocked port for the segment, which will function as the alternate port. All the other ports become unblocked. By default, REP packets are sent to a bridge protocol data unit-class MAC address. The packets can also be sent to a Cisco multicast address, which is used only

to send blocked port advertisement (BPA) messages when there is a failure in the segment. The packets are dropped by the devices not running REP.

Fast Convergence

REP runs on a physical link basis and not on a per-VLAN basis. Only one hello message is required for all the VLANs, and this reduces the load on the protocol. We recommend that you create VLANs consistently on all the switches in a given segment and configure the same allowed VLANs on the REP trunk ports. To avoid the delay introduced by relaying messages in software, REP also allows some packets to be flooded to a regular multicast address. These messages operate at the hardware flood layer (HFL) and are flooded to the entire network, not just the REP segment. Switches that do not belong to the segment treat them as data traffic. You can control flooding of these messages by configuring an administrative VLAN for the entire domain or for a particular segment.

VLAN Load Balancing

One edge port in the REP segment acts as the primary edge port; and another as the secondary edge port. It is the primary edge port that always participates in VLAN load balancing in the segment. REP VLAN balancing is achieved by blocking some VLANs at a configured alternate port and all the other VLANs at the primary edge port. When you configure VLAN load balancing, you can specify the alternate port in one of three ways:


- By entering the port ID of the interface. To identify the port ID of a port in the segment, enter the **show interface rep detail** interface configuration command for the port.
- By entering the **preferred** keyword to select the port that you previously configured as the preferred alternate port with the **rep segment segment-id preferred** interface configuration command.
- By entering the neighbor offset number of a port in the segment, which identifies the downstream neighbor port of an edge port. The neighbor offset number range is -256 to $+256$; a value of 0 is invalid. The primary edge port has an offset number of 1; positive numbers above 1 identify downstream neighbors of the primary edge port. Negative numbers indicate the secondary edge port (offset number -1) and its downstream neighbors.



Note Configure offset numbers on the primary edge port by identifying a port's downstream position from the primary (or secondary) edge port. Never enter an offset value of 1 because that is the offset number of the primary edge port.

The following figure shows neighbor offset numbers for a segment, where E1 is the primary edge port and E2 is the secondary edge port. The red numbers inside the ring are numbers offset from the primary edge port; the black numbers outside of the ring show the offset numbers from the secondary edge port. Note that you can identify all the ports (except the primary edge port) by either a positive offset number (downstream position from the primary edge port) or a negative offset number (downstream position from the secondary edge port). If E2 became the primary edge port, its offset number would then be 1 and E1 would be -1.

Figure 4: Neighbor Offset Numbers in a Segment



201890

When the REP segment is complete, all the VLANs are blocked. When you configure VLAN load balancing, you must also configure triggers in one of two ways:

- Manually trigger VLAN load balancing at any time by entering the **rep preempt segment** *segment-id* privileged EXEC command on the switch that has the primary edge port.
- Configure a preempt delay time by entering the **rep preempt delay** *seconds* interface configuration command. After a link failure and recovery, VLAN load balancing begins after the configured preemption time period elapses. Note that the delay timer restarts if another port fails before the time has elapsed.



Note When VLAN load balancing is configured, it does not start working until triggered by either manual intervention or a link failure and recovery.

When VLAN load balancing is triggered, the primary edge port sends out a message to alert all the interfaces in the segment about the preemption. When the secondary port receives the message, the message is sent to the network to notify the alternate port to block the set of VLANs specified in the message and to notify the primary edge port to block the remaining VLANs.

You can also configure a particular port in the segment to block all the VLANs. Only the primary edge port initiates VLAN load balancing, which is not possible if the segment is not terminated by an edge port on each end. The primary edge port determines the local VLAN load-balancing configuration.

Reconfigure the primary edge port to reconfigure load balancing. When you change the load-balancing configuration, the primary edge port waits for the **rep preempt segment** command or for the configured preempt delay period after a port failure and recovery, before executing the new configuration. If you change an edge port to a regular segment port, the existing VLAN load-balancing status does not change. Configuring a new edge port might cause a new topology configuration.

Spanning Tree Interaction

REP does not interact with STP, but it can coexist. A port that belongs to a segment is removed from spanning tree control and STP BPDUs are not accepted or sent from segment ports. Therefore, STP cannot run on a segment.

To migrate from an STP ring configuration to REP segment configuration, begin by configuring a single port in the ring as part of the segment and continue by configuring contiguous ports to minimize the number of segments. Each segment always contains a blocked port, so multiple segments means multiple blocked ports and a potential loss of connectivity. When the segment has been configured in both directions up to the location of the edge ports, you then configure the edge ports.

Resilient Ethernet Protocol Ports

REP segments consist of Failed, Open, or Alternate ports:

- A port configured as a regular segment port starts as a failed port.
- After the neighbor adjacencies are determined, the port transitions to alternate port state, blocking all the VLANs on the interface. Blocked-port negotiations occur, and when the segment settles, one blocked port remains in the alternate role and all the other ports become open ports.
- When a failure occurs in a link, all the ports move to the Failed state. When the Alternate port receives the failure notification, it changes to the Open state, forwarding all the VLANs.

A regular segment port that is converted to an edge port, or an edge port that is converted to a regular segment port, does not always result in a topology change. If you convert an edge port into a regular segment port, VLAN load balancing is not implemented unless it has been configured. For VLAN load balancing, you must configure two edge ports in the segment.

A segment port that is reconfigured as a spanning tree port restarts according to the spanning tree configuration. By default, this is a designated blocking port. If PortFast is configured or if STP is disabled, the port goes into the forwarding state.

How to Configure Resilient Ethernet Protocol

A segment is a collection of ports that are connected to one another in a chain and configured with a segment ID. To configure REP segments, configure the REP administrative VLAN (or use the default VLAN 1) and then add the ports to the segment, using interface configuration mode. You should configure two edge ports in a segment, with one of them being the primary edge port and the other the secondary edge port by default. A segment should have only one primary edge port. If you configure two ports in a segment as primary edge ports, for example, ports on different switches, the REP selects one of them to serve as the segment's primary edge port. If required, you can configure the location to which segment topology change notices (STCNs) and VLAN load balancing are to be sent.

Default Resilient Ethernet Protocol Configuration

REP is disabled on all the interfaces. When enabled, the interface is a regular segment port unless it is configured as an edge port.

When REP is enabled, the task of sending segment topology change notices (STCNs) is disabled, all the VLANs are blocked, and the administrative VLAN is VLAN 1.

When VLAN load balancing is enabled, the default is manual preemption with the delay timer disabled. If VLAN load balancing is not configured, the default after manual preemption is to block all the VLANs in the primary edge port.

Resilient Ethernet Protocol Configuration Guidelines

Follow these guidelines when configuring REP:

- REP is supported on 1-Gigabit Ethernet and 10-Gigabit Ethernet interfaces.
- We recommend that you begin by configuring one port and then configure contiguous ports to minimize the number of segments and the number of blocked ports.
- If more than two ports in a segment fail when no external neighbors are configured, one port goes into a forwarding state for the data path to help maintain connectivity during configuration. In the **show rep interface** command output, the Port Role for this port shows as “Fail Logical Open”; the Port Role for the other failed port shows as “Fail No Ext Neighbor”. When the external neighbors for the failed ports are configured, the ports go through the alternate port state transitions and eventually go to an open state or remain as the alternate port, based on the alternate port selection mechanism.
- REP ports must be Layer 2 IEEE 802.1Q or Trunk ports.
- We recommend that you configure all trunk ports in the segment with the same set of allowed VLANs.
- Be careful when configuring REP through a Telnet connection. Because REP blocks all VLANs until another REP interface sends a message to unblock it. You might lose connectivity to the router if you enable REP in a Telnet session that accesses the router through the same interface.
- You cannot run REP and STP on the same segment or interface.
- If you connect an STP network to a REP segment, be sure that the connection is at the segment edge. An STP connection that is not at the edge could cause a bridging loop because STP does not run on REP segments. All STP BPDUs are dropped at REP interfaces.
- You must configure all trunk ports in the segment with the same set of allowed VLANs, or a misconfiguration occurs.
- If REP is enabled on two ports on a switch, both ports must be either regular segment ports or edge ports. REP ports follow these rules:
 - There is no limit to the number of REP ports on a switch; however, only two ports on a switch can belong to the same REP segment.
 - If only one port on a switch is configured in a segment, the port should be an edge port.
 - If two ports on a switch belong to the same segment, they must be both edge ports, both regular segment ports, or one regular port and one edge no-neighbor port. An edge port and regular segment port on a switch cannot belong to the same segment.
 - If two ports on a switch belong to the same segment and one is configured as an edge port and one as a regular segment port (a misconfiguration), the edge port is treated as a regular segment port.
- REP interfaces come up in a blocked state and remain in a blocked state until they are safe to be unblocked. You need to be aware of this status to avoid sudden connection losses.
- REP sends all LSL PDUs in untagged frames on the native VLAN. The BPA message sent to the Cisco multicast address is sent on the administration VLAN, which is VLAN 1 by default.
- You can configure how long a REP interface remains up without receiving a hello from a neighbor. You can use the **rep lsl-age-timer** value interface configuration command to set the time from 120 ms to

10000 ms. The LSL hello timer is then set to the age-timer value divided by 3. In normal operation, three LSL hellos are sent before the age timer on the peer switch expires and checks for hello messages.

- EtherChannel port channel interfaces do not support LSL age-timer values less than 1000 ms. If you try to configure a value less than 1000 ms on a port channel, you receive an error message and the command is rejected.
- REP ports cannot be configured as one of the following port types:
 - Switched Port Analyzer (SPAN) destination port
 - Tunnel port
 - Access port
- REP is supported on EtherChannels, but not on an individual port that belongs to an EtherChannel.
- There can be a maximum of 26 REP segments per switch.
- If private VLANs are configured in a REP topology for an isolated trunk port, the secondary VLAN type must be mapped to the private VLAN isolated type only. A secondary VLAN cannot be mapped to a private VLAN community type for an isolated trunk port.

Configuring Resilient Ethernet Protocol Administrative VLAN

To avoid the delay created by link-failure messages, and VLAN-blocking notifications during load balancing, REP floods packets to a regular multicast address at the hardware flood layer (HFL). These messages are flooded to the whole network, and not just the REP segment. You can control the flooding of these messages by configuring an administrative VLAN.

Follow these guidelines when configuring the REP administrative VLAN:

- If you do not configure an administrative VLAN, the default is VLAN 1.
- The administrative VLAN cannot be the RSPAN VLAN.

If the REP domain has multiple REP segments that are mutually exclusive, it is difficult to maintain a loop-less single admin VLAN across the entire REP domain. Starting with the Cisco IOS XE 17.2.1 release, you can configure multiple REP VLANs and manage multiple REP segments which are mutually exclusive to each other. You can configure additional admin VLANs by specifying a segment id in the rep admin vlan command.

Existing configurations that use a single, global REP admin VLAN will continue working as earlier. REP segments that have not been assigned a specific admin VLAN will use the global admin VLAN. HFL packets will be flooded into the admin VLAN configured for the segment.

To configure the REP administrative VLAN for an REP segment, follow these steps:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.

	Command or Action	Purpose
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.
Step 3	rep admin vlan <i>vlan-id</i> segment <i>segment-id</i> Example: Device(config)# <code>rep admin vlan 2 segment 4</code>	Specifies the administrative VLAN for the segment. The VLAN range is from 2 to 4094. The segment id number ranges from 1-1024. To set the admin VLAN to 1, which is the default, enter the no rep admin vlan global configuration command.
Step 4	end Example: Device(config)# <code>end</code>	Exits global configuration mode and returns to privileged EXEC mode.
Step 5	show interface [<i>interface-id</i>] rep detail Example: Device# <code>show interface gigabitethernet1/1 rep detail</code>	(Optional) Verifies the configuration on a REP interface.
Step 6	copy running-config startup config Example: Device# <code>copy running-config startup config</code>	(Optional) Saves your entries in the switch startup configuration file.

Configuring a REP Interface

To configure REP, enable REP on each segment interface and identify the segment ID. This task is mandatory, and must be done before other REP configurations. You must also configure a primary and secondary edge port on each segment. All the other steps are optional.

Follow these steps to enable and configure REP on an interface:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> <code>enable</code>	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.

	Command or Action	Purpose
Step 3	interface <i>interface-id</i> Example: Device(config)# interface gigabitethernet1/1	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface).
Step 4	switchport mode trunk Example: Device(config-if)# switchport mode trunk	Configures the interface as a Layer 2 trunk port.
Step 5	rep segment <i>segment-id</i> [edge [no-neighbor] [primary]] [preferred] Example: Device(config-if)# rep segment 1 edge no-neighbor primary	<p>Enables REP on the interface and identifies a segment number. The segment ID range is from 1 to 1024.</p> <p>Note You must configure two edge ports, including one primary edge port, for each segment.</p> <p>These optional keywords are available:</p> <ul style="list-style-type: none"> • (Optional) edge—Configures the port as an edge port. Each segment has only two edge ports. Entering the keyword edge without the keyword primary configures the port as the secondary edge port. • (Optional) primary—Configures the port as the primary edge port, the port on which you can configure VLAN load balancing. • (Optional) no-neighbor—Configures a port with no external REP neighbors as an edge port. The port inherits all the properties of an edge port, and you can configure the properties the same way you would for an edge port. <p>Note Although each segment can have only one primary edge port, if you configure edge ports on two different switches and enter the keyword primary on both the switches, the configuration is valid. However, REP selects only one of these ports as the segment primary edge port. You can identify the primary edge port for a segment by entering the show rep topology command in privileged EXEC mode.</p>

	Command or Action	Purpose
		<ul style="list-style-type: none"> • (Optional) preferred—Indicates that the port is the preferred alternate port or the preferred port for VLAN load balancing. <p>Note Configuring a port as preferred does not guarantee that it becomes the alternate port; it merely gives the port a slight edge over equal contenders. The alternate port is usually a previously failed port.</p>
<p>Step 6</p>	<p>rep stcn {<i>interface interface id</i> segment id-list stp}</p> <p>Example:</p> <pre>Device(config-if)# rep stcn segment 25-50</pre>	<p>(Optional) Configures the edge port to send segment topology change notices (STCNs).</p> <ul style="list-style-type: none"> • interface interface-id—Designates a physical interface or port channel to receive STCNs. • segment id-list—Identifies one or more segments to receive STCNs. The range is from 1 to 1024. • stp—Sends STCNs to STP networks. <p>Note Spanning Tree (MST) mode is required on edge no-neighbor nodes when rep stcn stp command is configured for sending STCNs to STP networks.</p>
<p>Step 7</p>	<p>rep block port {<i>id port-id</i> <i>neighbor-offset</i> preferred} vlan {<i>vlan-list</i> all}</p> <p>Example:</p> <pre>Device(config-if)# rep block port id 0009001818D68700 vlan 1-100</pre>	<p>(Optional) Configures VLAN load balancing on the primary edge port, identifies the REP alternate port in one of three ways (id port-id, neighbor_offset, preferred), and configures the VLANs to be blocked on the alternate port.</p> <ul style="list-style-type: none"> • id port-id—Identifies the alternate port by port ID. The port ID is automatically generated for each port in the segment. You can view interface port IDs by entering the show interface type number rep [detail] privileged EXEC command. • neighbor_offset—Number to identify the alternate port as a downstream neighbor from an edge port. The range is from -256 to 256, with negative numbers indicating the downstream neighbor from the secondary edge port. A value of 0 is

	Command or Action	Purpose
		<p>invalid. Enter -1 to identify the secondary edge port as the alternate port.</p> <p>Note Because you enter the rep block port command at the primary edge port (offset number 1), you cannot enter an offset value of 1 to identify an alternate port.</p> <ul style="list-style-type: none"> • preferred—Selects the regular segment port that is previously identified as the preferred alternate port for VLAN load balancing. • vlan <i>vlan-list</i>—Blocks one VLAN or a range of VLANs. • vlan all—Blocks all the VLANs. <p>Note Enter this command only on the REP primary edge port.</p>
Step 8	<p>rep preempt delay <i>seconds</i></p> <p>Example:</p> <pre>Device(config-if)# rep preempt delay 100</pre>	<p>(Optional) Configures a preempt time delay.</p> <ul style="list-style-type: none"> • Use this command if you want VLAN load balancing to be automatically triggered after a link failure and recovery. • The time delay range is between 15 to 300 seconds. The default is manual preemption with no time delay. <p>Note Enter this command only on the REP primary edge port.</p>
Step 9	<p>rep lsl-age-timer <i>value</i></p> <p>Example:</p> <pre>Device(config-if)# rep lsl-age-timer 2000</pre>	<p>(Optional) Configures a time (in milliseconds) for which the REP interface remains up without receiving a hello from a neighbor.</p> <p>The range is from 120 to 10000 ms in 40-ms increments. The default is 5000 ms (5 seconds).</p> <p>Note</p> <ul style="list-style-type: none"> • EtherChannel port channel interfaces do not support LSL age-timer values that are less than 1000 ms. • Both the ports on the link should have the same LSL age that is configured in order to avoid link flaps.

	Command or Action	Purpose
Step 10	end Example: Device(config-if)# end	Exits global configuration mode and returns to privileged EXEC mode.
Step 11	show interface [interface-id] rep [detail] Example: Device# show interface gigabitethernet1/1 rep detail	(Optional) Displays the REP interface configuration.
Step 12	copy running-config startup-config Example: Device# copy running-config startup-config	(Optional) Saves your entries in the router startup configuration file.

Setting Manual Preemption for VLAN Load Balancing

If you do not enter the **rep preempt delay seconds** interface configuration command on the primary edge port to configure a preemption time delay, the default is to manually trigger VLAN load balancing on the segment. Be sure that all the other segment configurations have been completed before manually preempting VLAN load balancing. When you enter the **rep preempt delay segment segment-id** command, a confirmation message is displayed before the command is executed because preemption might cause network disruption.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	rep preempt segment segment-id Example: Device(config)# rep preempt segment 100 The command will cause a momentary traffic disruption. Do you still want to continue? [confirm]	Manually triggers VLAN load balancing on the segment. You need to confirm the command before it is executed.
Step 4	end Example: Device# end	Exits global configuration mode and returns to privileged EXEC mode.

	Command or Action	Purpose
Step 5	show rep topology segment <i>segment-id</i> Example: Device# <code>show rep topology segment 100</code>	(Optional) Displays REP topology information.
Step 6	end Example: Device# <code>end</code>	Exits privileged EXEC mode.

Configuring Simple Network Management Protocol Traps for Resilient Ethernet Protocol

You can configure a router to send REP-specific traps to notify the Simple Network Management Protocol (SNMP) server of link-operational status changes and port role changes.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> <code>enable</code>	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.
Step 3	snmp mib rep trap-rate <i>value</i> Example: Device(config)# <code>snmp mib rep trap-rate 500</code>	Enables the switch to send REP traps, and sets the number of traps that are sent per second. <ul style="list-style-type: none"> Enter the number of traps sent per second. The range is from 0 to 1000. The default is 0 (no limit is imposed; a trap is sent at every occurrence).
Step 4	end Example: Device(config)# <code>end</code>	Returns to privileged EXEC mode.
Step 5	show running-config Example: Device# <code>show running-config</code>	(Optional) Displays the running configuration, which can be used to verify the REP trap configuration.
Step 6	copy running-config startup-config Example:	(Optional) Saves your entries in the switch startup configuration file.

	Command or Action	Purpose
	Device# <code>copy running-config startup-config</code>	

Monitoring Resilient Ethernet Protocol Configurations



Note When the port on the peer side is down the `show rep topology` command displays both the primary and the secondary ports as secondary ports.

This is an example of the output for the `show interface [interface-id] rep [detail]` command. For this display, the REP configuration and status on an uplink port is shown.

```
Device# show interfaces TenGigabitEthernet4/1 rep detail
```

```
TenGigabitEthernet4/1 REP enabled
Segment-id: 3 (Primary Edge)
PortID: 03010015FA66FF80
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 02040015FA66FF804050
Port Role: Open
Blocked VLAN: <empty>
Admin-vlan: 1
Preempt Delay Timer: disabled
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 999, tx: 652
HFL PDU rx: 0, tx: 0
BPA TLV rx: 500, tx: 4
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 6, tx: 5
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 135, tx: 136
```

This is an example of the output for the `show interface [interface-id] rep [detail]` command. For this display, the REP configuration and status on a downlink port is shown.

```
Device#show interface TenGigabitEthernet5/0/27 rep detail
```

```
TenGigabitEthernet5/0/27 REP enabled
Segment-id: 1 (Segment)
PortID: 019B380E4D9ACAC0
Preferred flag: No
Operational Link Status: NO_NEIGHBOR
Current Key: 019B380E4D9ACAC0696B
Port Role: Fail No Ext Neighbor
Blocked VLAN: 1-4094
Admin-vlan: 1
Preempt Delay Timer: 100 sec
LSL Ageout Timer: 2000 ms
LSL Ageout Retries: 5
```

```

Configured Load-balancing Block Port: 09E9380E4D9ACAC0
Configured Load-balancing Block VLAN: 1-100
STCN Propagate to: segment 25
LSL PDU rx: 292, tx: 340
HFL PDU rx: 0, tx: 0
BPA TLV rx: 0, tx: 0
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 0, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 0, tx: 0

```

This is an example for the **show rep topology** [**segment** *segment-id*] [**archive**] [**detail**] command. For this display, the REP topology information for all the segments is shown.

```
Device# show rep topology
```

```

REP Segment 1
BridgeName      PortName      Edge Role
-----
10.64.106.63    Te5/4         Pri  Open
10.64.106.228  Te3/4         Open
10.64.106.228  Te3/3         Open
10.64.106.67   Te4/3         Open
10.64.106.67   Te4/4         Alt
10.64.106.63   Te4/4         Sec  Open

REP Segment 3
BridgeName      PortName      Edge Role
-----
10.64.106.63    Gi50/1        Pri  Open
SVT_3400_2      Gi0/3         Open
SVT_3400_2      Gi0/4         Open
10.64.106.68   Gi40/2        Open
10.64.106.68   Gi40/1        Open
10.64.106.63   Gi50/2        Sec  Alt

```

Additional References for Resilient Ethernet Protocol

Related Documents

Related Topic	Document Title
REP commands	See the Layer 2/3 Commands section of the <i>Command Reference (Catalyst 9500 Series Switches)</i>

Feature History for Resilient Ethernet Protocol

This table provides release and related information for features explained in this module.

These features are available on all releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.5.1a	Resilient Ethernet Protocol	REP is a Cisco proprietary protocol that provides an alternative to STP to control network loops, handle link failures, and improve convergence time. On Cisco Catalyst 9500 Series Switches, support was introduced only on the C9500-40X model.
Cisco IOS XE Fuji 16.8.1a	Resilient Ethernet Protocol	Support for the feature was introduced on all switch models with 1-Gigabit Ethernet and 10-Gigabit Ethernet interfaces.
Cisco IOS XE Fuji 16.9.1	Resilient Ethernet Protocol Support on Downlink Ports	Support for configuring REP on downlink ports was introduced.
Cisco IOS XE Amsterdam 17.2.1	Multiple Administrative VLANs for Resilient Ethernet Protocol	Support for configuring multiple administrative VLANs in REP was introduced.

Use Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.

