Configuring Resilient Ethernet Protocol

Overview of 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.

A REP segment is a chain of ports connected to each other and configured with a segment ID. Each segment consists of standard (nonedge) segment ports and two user-configured edge ports. A device 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 Ethernet Flow Point (EFP) interfaces.

The following figure 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 the ports are operational (as in the segment on the left), a single port is blocked, as shown by the diagonal line. 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

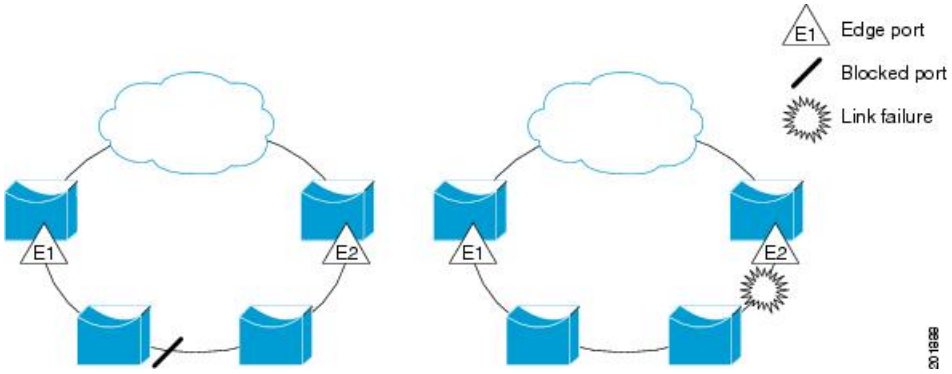

The segment shown in the figure above is an open segment; there is no connectivity between the two edge ports. The REP segment cannot cause a bridging loop, and you can safely connect the segment edges to any network. All the hosts connected to devices inside the segment have two possible connections 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 all the ports to ensure that connectivity is available through the other gateway.

The segment shown in the following figure is a ring segment, with both the edge ports located on the same device. With this configuration, you can create a redundant connection between any two devices in the segment.

Figure 2: REP Ring Segment

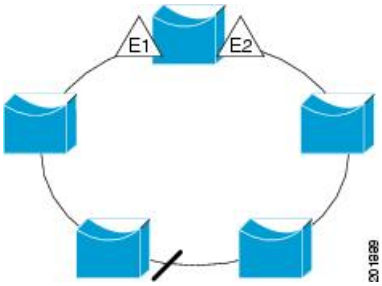

REP segments have the following characteristics:

- If all the ports in a segment are operational, one port (referred to as the *alternate* port) is in the blocked state for each VLAN. If VLAN load balancing is configured, two ports in the segment control the blocked state of VLANs.
- If one or more ports in a segment is not operational, and cause a link failure, all the ports forward traffic on all the VLANs to ensure connectivity.
- In case of a link failure, alternate ports are unblocked as quickly as possible. When the failed link is up, 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. REP also supports VLAN load balancing, which is controlled by the primary edge port (any port in the segment).

In access ring-topologies, the neighboring switch might not support REP as shown in the following figure. In this scenario, you can configure the non-REP-facing ports (E1 and E2) as edge no-neighbor ports. These ports inherit all the properties of edge ports, and you can configure them the same as any edge port, including configuring them to send STP or REP topology change notices to the aggregation switch. In this scenario, the STP topology change notice (TCN) that is sent is a multiple spanning-tree (MST) STP message.

Figure 3: Edge No-Neighbor Ports

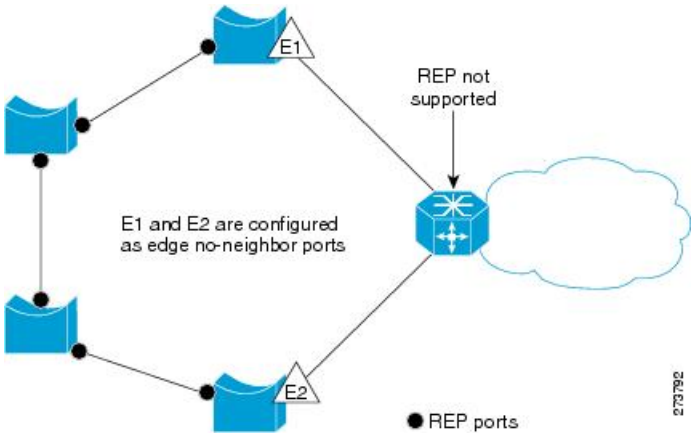

REP has these limitations:

- You must configure each segment port; an incorrect configuration might 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.

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.

The estimated convergence recovery time on fiber interfaces is between 50 ms and 200 ms for the local segment with 200 VLANs configured. Convergence for VLAN load balancing is 300 ms or less.

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 numbers inside the ring are numbers offset from the primary edge port; the 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

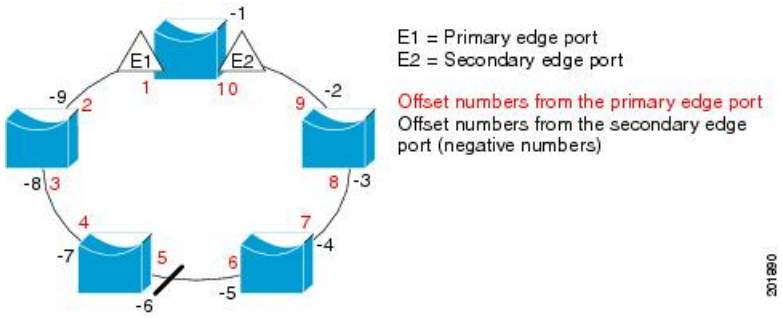

When the REP segment is complete, all the VLANs are blocked. When you configure VLAN load balancing, you must also configure triggers in one of two ways:

- 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 the STP or the Flex Link feature, but can coexist with both. 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 an 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. Since each segment always contains a blocked port, multiple segments means multiple blocked

ports and a potential loss of connectivity. After the segment is configured in both directions up to the location of the edge ports, configure the edge ports.

REP 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 converted to an edge port, or an edge port 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 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 REP 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.

REP Configuration Guidelines

Follow these guidelines when configuring REP:

- 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 is displayed as **Fail Logical Open**; the Port Role for the other failed port is displayed as **Fail No Ext Neighbor**. When the external neighbors for the failed ports are configured, the ports go through the alternate port 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 the trunk ports in a segment with the same set of allowed VLANs.
- Be careful when configuring REP through a Telnet connection because REP blocks all the 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 or REP and Flex Links on the same segment or interface.
- If you connect an STP network to an REP segment, be sure that the connection is at the segment edge. An STP connection that is not at the edge might cause a bridging loop because STP does not run on REP segments. All the STP BPDUs are dropped at REP interfaces.
- You must configure all the trunk ports in a segment with the same set of allowed VLANs. If this is not done, misconfiguration occurs.
- If REP is enabled on two ports on a switch, both the 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 both be edge ports, 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 must, therefore, be aware of the status of REP interfaces to avoid sudden connection losses.
- REP sends all the LSL PDUs in the untagged frames to the native VLAN. The BPA message sent to a Cisco multicast address is sent to the administration VLAN, which is VLAN 1 by default.
- You can configure the duration for which a REP interface remains up without receiving a hello from a neighbor. 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 64 REP segments per switch.

Configuring REP 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.
- You can configure one admin VLAN on the switch for all segments.
- The administrative VLAN cannot be the RSPAN VLAN.

To configure the REP administrative VLAN, follow these steps, beginning in privileged EXEC mode:

SUMMARY STEPS

1. **configure terminal**
2. **rep admin vlan** *vlan-id*
3. **end**
4. **show interface** [*interface-id*] **rep detail**
5. **copy running-config startup config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 2	**rep admin vlan** *vlan-id* **Example:** `Device(config)# rep admin vlan 2`	Specifies the administrative VLAN. The range is from 2 to 4094. To set the admin VLAN to 1, which is the default, enter the **no rep admin vlan** global configuration command.

	Command or Action	Purpose
Step 3	**end** **Example:** `Device(config)# end`	Exits global configuration mode and returns to privileged EXEC mode.
Step 4	**show interface** [*interface-id*] **rep detail** **Example:** `Device# show interface gigabitethernet1/1 rep detail`	(Optional) Verifies the configuration on a REP interface.
Step 5	**copy running-config startup config** **Example:** `Device# copy running-config startup config`	(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:

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *interface-id*
4. **switchport mode trunk**
5. **rep segment** *segment-id* [**edge** [**no-neighbor**] [**primary**]] [**preferred**]
6. **rep stcn** {**interface** *interface id* | **segment** *id-list* | **stp**}
7. **rep block port** {**id** *port-id* | *neighbor-offset* | **preferred**} **vlan** {*vlan-list* | **all**}
8. **rep preempt delay** *seconds*
9. **rep lsl-age-timer** *value*
10. **end**
11. **show interface** [*interface-id*] **rep** [**detail**]
12. **copy running-config startup-config**

DETAILED STEPS

	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# configure terminal`	Enters global configuration mode.
Step 3	**interface** *interface-id* **Example:** `Device# 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# switchport mode trunk`	Configures the interface as a Layer 2 trunk port.
Step 5	**rep segment** *segment-id* [**edge** [**no-neighbor**] [**primary**]] [**preferred**] **Example:** `Device# rep segment 1 edge no-neighbor primary`	Enables REP on the interface and identifies a segment number. The segment ID range is from 1 to 1024. **Note** You must configure two edge ports, including one primary edge port, for each segment. These optional keywords are available: • (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. **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** privileged EXEC command. • (Optional) **preferred**—Indicates that the port is the preferred alternate port or the preferred port for VLAN load balancing.

	Command or Action	Purpose
		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.
Step 6	**rep stcn** {**interface** *interface id* | **segment** *id-list* | **stp**} **Example:** `Device# rep stcn segment 25-50`	(Optional) Configures the edge port to send segment topology change notices (STCNs). • **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. **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.
Step 7	**rep block port** {**id** *port-id* | *neighbor-offset* | **preferred**} **vlan** {*vlan-list* | **all**} **Example:** `Device# rep block port id 0009001818D68700 vlan 1-100`	(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. • **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 invalid. Enter **-1** to identify the secondary edge port as the alternate port. **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. • **preferred**—Selects the regular segment port previously identified as the preferred alternate port for VLAN load balancing. • **vlan** *vlan-list*—Blocks one VLAN or a range of VLANs.

	Command or Action	Purpose
		• **vlan all**—Blocks all the VLANs. **Note** Enter this command only on the REP primary edge port.
Step 8	**rep preempt delay** *seconds* **Example:** `Device# rep preempt delay 100`	(Optional) Configures a preempt time delay. • Use this command if you want VLAN load balancing to be automatically triggered after a link failure and recovery. • The time delay range is between15 to 300 seconds. The default is manual preemption with no time delay. **Note** Enter this command only on the REP primary edge port.
Step 9	**rep lsl-age-timer** *value* **Example:** `Device# rep lsl-age-timer 2000`	(Optional) Configures a time (in milliseconds) for which the REP interface remains up without receiving a hello from a neighbor. The range is from 120 to 10000 ms in 40-ms increments. The default is 5000 ms (5 seconds). **Note** • 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 configured in order to avoid link flaps.
Step 10	**end** **Example:** `Device(config)# end`	Exits global configuration mode and returns to privileged EXEC mode.
Step 11	**show interface** [*interface-id*] **rep** [**detail**] **Example:** `Device(config)# show interface gigabitethernet1/1 rep detail`	(Optional) Displays the REP interface configuration.
Step 12	**copy running-config startup-config** **Example:** `Device(config)# 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.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **rep preempt segment** *segment-id*
4. **show rep topology segment** *segment-id*
5. **end**

DETAILED STEPS

	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# 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	**show rep topology segment** *segment-id* **Example:** `Device# show rep topology segment 100`	(Optional) Displays REP topology information.
Step 5	**end** **Example:** `Device# end`	Exits privileged EXEC mode.

Configuring SNMP Traps for REP

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.

SUMMARY STEPS

1. **configure terminal**
2. **snmp mib rep trap-rate** *value*
3. **end**
4. **show running-config**
5. **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 2	**snmp mib rep trap-rate** *value* **Example:** `Device(config)# snmp mib rep trap-rate 500`	Enables the switch to send REP traps, and sets the number of traps sent per second. • 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 3	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.
Step 4	**show running-config** **Example:** `Device# show running-config`	(Optional) Displays the running configuration, which can be used to verify the REP trap configuration.
Step 5	**copy running-config startup-config** **Example:** `Device# copy running-config startup-config`	(Optional) Saves your entries in the switch startup configuration file.

Monitoring Resilient Ethernet Protocol Configurations

You can display the rep interface and rep topology details using the commands in this topic.

SUMMARY STEPS

1. **show interface** [*interface-id*] **rep** [**detail**]
2. **show rep topology** [**segment** *segment-id*] [**archive**] [**detail**]

DETAILED STEPS

	Command or Action	Purpose

Step 1	**show interface** [*interface-id*] **rep** [**detail**]
	Example:
	`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`
		Displays REP configuration and status for an interface or for all the interfaces.
		• (Optional) **detail**—Displays interface-specific REP information.

Step 2	**show rep topology** [**segment** *segment-id*] [**archive**] [**detail**]
	Example:
	`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`
		Displays REP topology information for a segment or for all the segments, including the primary and secondary edge ports in the segment.
		• (Optional) **archive**—Displays the last stable topology.
		Note An archive topology is not retained when the switch reloads.
		• (Optional) **detail**—Displays detailed archived information.

	Command or Action	Purpose
	10.64.106.68 Gi40/1 Open 10.64.106.63 Gi50/2 Sec Alt	

Configuration Examples for Resilient Ethernet Protocol

This section provides the following configuration examples:

Example: Configuring the REP Administrative VLAN

This example shows how to configure the administrative VLAN as VLAN 100, and verify the configuration by entering the **show interface rep detail** command on one of the REP interfaces:

```
Device# configure terminal
Device(config)# rep admin vlan 100
Device(config)# end
Device# show interface gigabitethernet1/1 rep detail

GigabitEthernet1/1 REP enabled
Segment-id: 2 (Edge)
PortID: 00010019E7144680
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 0002001121A2D5800E4D
Port Role: Open
Blocked Vlan: <empty>
Admin-vlan: 100
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 3322, tx: 1722
HFL PDU rx: 32, tx: 5
BPA TLV rx: 16849, tx: 508
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 118, tx: 118
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 4214, tx: 4190
```

The following example shows how to create an administrative VLAN per segment. Here, VLAN 2 is configured as the administrative VLAN only for REP segment 2. All the remaining segments that are not configured have VLAN 1 as the administrative VLAN by default.

```
Device# configure terminal
Device(config)# rep admin vlan 2 segment 2
Device(config)# end
```

Example: Configuring a REP Interface

This example shows how to configure an interface as the primary edge port for segment 1, to send STCNs to segments 2 through 5, and to configure the alternate port as the port with port ID 0009001818D68700 to block

all the VLANs after a preemption delay of 60 seconds after a segment port failure and recovery. The interface is configured to remain up for 6000 ms without receiving a hello from a neighbor.

```
Switch# configure terminal
Switch (conf)# interface gigabitethernet1/1
Switch (conf-if)# rep segment 1 edge primary
Switch (conf-if)# rep stcn segment 2-5
Switch (conf-if)# rep block port 0009001818D68700 vlan all
Switch (conf-if)# rep preempt delay 60
Switch (conf-if)# rep lsl-age-timer 6000
Switch (conf-if)# end
```

This example shows how to configure the same configuration when the interface has no external REP neighbor:

```
Switch# configure terminal
Switch (conf)# interface gigabitethernet1/1
Switch (conf-if)# rep segment 1 edge no-neighbor primary
Switch (conf-if)# rep stcn segment 2-5
Switch (conf-if)# rep block port 0009001818D68700 vlan all
Switch (conf-if)# rep preempt delay 60
Switch (conf-if)# rep lsl-age-timer 6000
Switch (conf-if)# end
```

This example shows how to configure the VLAN blocking configuration shown in the Figure 5. The alternate port is the neighbor with neighbor offset number 4. After manual preemption, VLANs 100 to 200 are blocked at this port, and all the other VLANs are blocked at the primary edge port E1 (Gigabit Ethernet port 1/1).

Figure 5: Example of VLAN Blocking

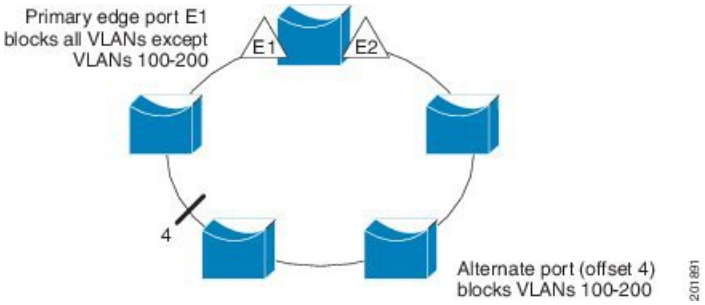


```
Switch# configure terminal
Switch (conf)# interface gigabitethernet1/1
Switch (conf-if)# rep segment 1 edge primary
Switch (conf-if)# rep block port 4 vlan 100-200
Switch (conf-if)# end
```