



CHAPTER 14

Configuring 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, to respond to link failures, and to 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.

The following sections describe how to configure REP:

- [Understanding Resilient Ethernet Protocol, page 14-1](#)
- [Configuring Resilient Ethernet Protocol \(REP\), page 14-6](#)
- [Configuration Examples for REP, page 14-16](#)

Understanding Resilient Ethernet Protocol


The following sections provide further information about REP:

- [Overview](#)
- [Link Integrity](#)
- [Fast Convergence](#)
- [VLAN Load Balancing](#)
- [REP Ports](#)

Overview


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 switch can have only two ports belonging 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 Layer 2 trunk interfaces.

[Figure 14-1](#) 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. When there is a network failure, as shown in the diagram on the right, the blocked port returns to the forwarding state to minimize network disruption.

Figure 14-1 REP Open Segments

The segment shown in Figure 14-1 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 host cannot access its usual gateway because of a failure, REP unblocks all ports to ensure that connectivity is available through the other gateway.

The segment shown in Figure 14-2, with both edge ports located on the same switch, is a ring segment. In this configuration, there is connectivity between the edge ports through the segment. With this configuration, you can create a redundant connection between any two switches in the segment.

Figure 14-2 REP Ring Segment


REP segments have these characteristics:

- If all ports in the 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, causing a link failure, all ports forward traffic on all VLANs to ensure connectivity.
- In case of a link failure, the alternate ports are unblocked as quickly as possible. When the failed link comes back 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, controlled by the primary edge port but occurring at any port in the segment.

In access ring topologies, the neighboring switch might not support REP, as shown in [Figure 14-3](#). In this case, you can configure the non-REP facing ports (E1 and E2) as edge no-neighbor ports. These ports inherit all 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 case the STP topology change notice (TCN) that is sent is a multiple spanning-tree (MST) STP message.

Figure 14-3 No-neighbor Topology



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.

Link Integrity

REP does not use an end-to-end polling mechanism 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 VLANs are blocked on an interface until it detects the neighbor. 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 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.
- The neighbor does not acknowledge the local port as a peer.

Each port creates an adjacency with its immediate neighbor. After the neighbor adjacencies are created, the ports negotiate to determine one blocked port for the segment, the alternate port. All other ports become unblocked. By default, REP packets are sent to a BPDU class MAC address. The packets can also be sent to the 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 devices not running REP.

Fast Convergence

Because REP runs on a physical link basis and not a per-VLAN basis, only one hello message is required for all VLANs, reducing the load on the protocol. We recommend that you create VLANs consistently on all 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 whole 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 a dedicated administrative VLAN for the whole domain.

The estimated convergence recovery time on fiber interfaces is less than 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; the other as the secondary edge port. The primary edge port 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 other VLANs at the primary edge port. When you configure VLAN load balancing, you can specify the alternate port in one of three ways:

- Enter 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.
- Enter 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 identify the secondary edge port (offset number -1) and its downstream neighbors.



Note


You configure offset numbers on the primary edge port by identifying the downstream position from the primary (or secondary) edge port. You would never enter an offset value of 1 because that is the offset number of the primary edge port itself.

Figure 14-4 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 the ring show the offset numbers from the secondary edge port. Note that you can identify all 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.

- 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.

Figure 14-4 Neighbor Offset Numbers in a Segment



201890

When the REP segment is complete, all 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 preempt 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 a message to alert all interfaces in the segment about the preemption. When the secondary port receives the message, it is reflected into 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 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 again 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.

**Note**

Do not configure VLAN load balancing on an interface that carries Ethernet over multiprotocol label switching (EoMPLS) traffic. VLAN load balancing across the REP ring might prevent forwarding some of the EoMPLS traffic.

Spanning Tree Interaction

REP does not interact with STP or with 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.

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 to the edge ports, you then configure the edge ports.

REP Ports

Ports in REP segments are Failed, Open, or Alternate.

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

For instructions on how to configure REP, see [Configuring Resilient Ethernet Protocol \(REP\)](#), page 14-6.

Configuring Resilient Ethernet Protocol (REP)

A segment is a collection of ports connected one to the other in a chain and configured with a segment ID. To configure REP segments, you 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 the segment, one as the primary edge port and the other, by default, the secondary edge port. A segment has only one primary edge port. If you configure two ports in a segment as the primary edge port, for example ports on different switches, the REP selects one to serve as the segment primary edge port. You can also optionally configure where to send segment topology change notices (STCNs) and VLAN load balancing messages.

The following sections describe how to configure REP on the Cisco MWR 2941:

- [Default REP Configuration, page 14-7](#)
- [REP Configuration Guidelines, page 14-7](#)
- [Configuring the REP Administrative VLAN, page 14-8](#)
- [Configuring REP Interfaces, page 14-10](#)
- [Setting Manual Preemption for VLAN Load Balancing, page 14-13](#)
- [Configuring SNMP Traps for REP, page 14-14](#)
- [Monitoring REP, page 14-15](#)

Default REP Configuration

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

When REP is enabled, the sending of segment topology change notices (STCNs) is disabled, all 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 VLANs at 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 the 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** privileged EXEC 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 election mechanism.
- REP ports must be Layer 2 trunk ports.
- Be careful when configuring REP through a Telnet connection. Because REP blocks all VLANs until another REP interface sends a message to unblock the VLAN, you might lose connectivity to the router if you enable REP in a Telnet session that accesses the router through the REP 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 the 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.
- 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 and remain in a blocked state until notified that it is safe to unblock. You need to be aware of this 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 three. In normal operation, three LSL hellos are sent before the age timer on the peer switch expires and searches for hello messages.
- 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 three. In normal operation, three LSL hellos are sent before the age timer on the peer switch expires and searches for hello messages.
- REP ports cannot be configured as one of these port types:
 - SPAN destination port
 - Private VLAN
 - Tunnel port
 - Access port
- There is a maximum of 64 REP segments per switch.

Configuring the REP Administrative VLAN

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

Follow these guidelines when configuring the REP administrative VLAN:

- If you do not configure an administrative VLAN, the default is VLAN 1.
- There can be only one administrative VLAN on a switch and on a segment. However, this is not enforced by software.

Beginning in privileged EXEC mode, follow these steps to configure the REP administrative VLAN:

	Command	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	rep admin vlan <i>vlan-id</i> Example: Router(config)# rep admin vlan 1	Configures a REP administrative VLAN. <ul style="list-style-type: none">• Specify the administrative VLAN. The range is 1–4094. The default is VLAN 1.
Step 4	end Example: Router(config)# end	Returns to privileged EXEC mode.
Step 5	show interface [<i>interface-id</i>] rep [<i>detail</i>] Example: Router# show interface gigabitethernet0/1 rep detail	Displays the REP configuration and status for a specified interface. <ul style="list-style-type: none">• Enter the physical interface.
Step 6	copy running-config startup config Example: Router# copy running-config startup config	(Optional) Saves your entries in the router startup configuration file.

Configuring REP Interfaces

For REP operation, you need to enable it on each segment interface and to identify the segment ID. This step is required and must be done before other REP configuration. You must also configure a primary and secondary edge port on each segment. All other steps are optional.

Beginning in privileged EXEC mode, follow these steps to enable and configure REP on an interface:

	Command	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	interface interface-id Example: Router(config)# interface gigabitethernet0/1	Specifies the interface, and enters interface configuration mode. <ul style="list-style-type: none"> Enter the interface ID. The interface can be a physical Layer 2 interface.
Step 4	switchport mode trunk Example: Router(config-if)# switchport mode trunk	Configures the interface as a Layer 2 trunk port.

	Command	Purpose
Step 5	<p>rep segment <i>segment-id</i> [edge [no-neighbor] [primary]] [preferred]</p> <p>Example: Router(config-if)# rep segment 1 edge preferred</p>	<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"> Enter edge to configure the port as an edge port. Entering edge without the primary keyword configures the port as the secondary edge port. Each segment has only two edge ports. (Optional) Enter no-neighbor to configure a port with no external REP neighbors as an edge port. The port inherits all properties of edge ports, and you can configure them the same as any edge port. On an edge port, enter primary to configure the port as the primary edge port, the port on which you can configure VLAN load balancing. <p>Note Although each segment can have only one primary edge port, if you configure edge ports on two different switches and enter the primary keyword on both switches, the configuration is allowed. 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.</p> <ul style="list-style-type: none"> Enter preferred to indicate 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 it a slight edge among equal contenders. The alternate port is usually a previously failed port.</p>
Step 6	<p>rep lsl-retries <i>number-of-retries</i></p> <p>Example: Router(config-if)# rep lsl-retries 4</p>	<p>Use the rep lsl-retries command to configure the REP link status layer (LSL) number of retries before the REP link is disabled.</p>
Step 7	<p>rep stcn {interface <i>interface-id</i> segment <i>id-list</i> stp}</p> <p>Example: Router(config-if)# rep stcn segment 2-5</p>	<p>(Optional) Configures the edge port to send segment topology change notices (STCNs).</p> <ul style="list-style-type: none"> Enter interface <i>interface-id</i> to designate a physical interface to receive STCNs. Enter segment <i>id-list</i> to identify one or more segments to receive STCNs. The range is from 1–1024. Enter stp to send STCNs to STP networks.

Command	Purpose
<p>Step 8 <code>rep block port {id port-id neighbor-offset preferred} vlan {vlan-list all}</code></p> <p>Example: Router(config-if)# <code>rep block port 0009001818D68700 vlan all</code></p>	<p>(Optional) Configures VLAN load balancing on the primary edge port, identifies the REP alternate port in one of three ways, and configures the VLANs to be blocked on the alternate port.</p> <ul style="list-style-type: none"> Enter the id port-id to identify 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 interface-id rep [detail] privileged EXEC command. Enter a 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. <p>Note Because you enter this command at the primary edge port (offset number 1), you would never enter an offset value of 1 to identify an alternate port.</p> <ul style="list-style-type: none"> Enter preferred to select the regular segment port previously identified as the preferred alternate port for VLAN load balancing. Enter vlan vlan-list to block one VLAN or a range of VLANs. Enter vlan all to block all VLANs. <p>Note Enter this command only on the REP primary edge port.</p>
<p>Step 9 <code>rep preempt delay seconds</code></p> <p>Example: Router(config-if)# <code>rep preempt delay 60</code></p>	<p>(Optional) Configures a preempt time delay. Use this command if you want VLAN load balancing to automatically trigger after a link failure and recovery. The time delay range is 15 to 300 seconds. The default is manual preemption with no time delay.</p> <p>Note Use this command only on the REP primary edge port.</p>
<p>Step 10 <code>rep lsl-age-timer value</code></p> <p>Example: Router(config-if) <code>rep lsl-age-timer 5000</code></p>	<p>(Optional) Configure 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).</p>
<p>Step 11 <code>end</code></p> <p>Example: Router(config-if)# <code>end</code></p>	<p>Returns to privileged EXEC mode.</p>
<p>Step 12 <code>show interface [interface-id] rep [detail]</code></p> <p>Example: Router(config-if)# <code>show interface gigabitethernet0/1 rep detail</code></p>	<p>Verifies the REP interface configuration.</p> <ul style="list-style-type: none"> Enter the interface ID and the optional detail keyword, if desired.

Command	Purpose
<p>Step 13 <code>show rep topology [segment <i>segment-id</i>]</code> <code>[archive] [detail]</code></p> <p>Example:</p> <pre>Router# show rep topology segment 1 REP Segment 1 BridgeName PortName Edge Role ----- sw1_multseg_3750 Gi1/1/1 Pri Alt sw3_multseg_3400 Gi0/13 Open sw3_multseg_3400 Gi0/14 Alt sw4_multseg_3400 Gi0/13 Open sw4_multseg_3400 Gi0/14 Open sw5_multseg_3400 Gi0/13 Open sw5_multseg_3400 Gi0/14 Open sw2_multseg_3750 Gi1/1/2 Open sw2_multseg_3750 Gi1/1/1 Open sw1_multseg_3750 Gi1/1/2 Sec Open</pre>	<p>Indicates which port in the segment is the primary edge port.</p>
<p>Step 14 <code>copy running-config startup config</code></p> <p>Example:</p> <pre>Router(config-if)# copy running-config startup config</pre>	<p>(Optional) Saves your entries in the router startup configuration file.</p>

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 to complete all other segment configuration before manually preempting VLAN load balancing. When you enter the **rep preempt segment** *segment-id* command, a confirmation message appears before the command is executed because preemption can cause network disruption.



Note

Do not configure VLAN load balancing on an interface that carries Ethernet over MultiprotocolLabel Switching (EoMPLS) traffic. VLAN load balancing across the REP ring might prevent forwarding some of the EoMPLS traffic.

Beginning in privileged EXEC mode, follow these steps on the switch that has the segment primary edge port to manually trigger VLAN load balancing on a segment:

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	rep preempt segment <i>segment-id</i> Example: Router(config)# rep preempt segment 1	Manually triggers VLAN load balancing on the segment. <ul style="list-style-type: none">Enter the segment ID. Note You will be asked to confirm the action before the command is executed.
Step 4	end Example: Router(config)# end	Returns to privileged EXEC mode.
Step 5	show rep topology Example: Router# show rep topology	Views the REP topology information.

Configuring SNMP Traps for REP

You can configure the switch to send REP-specific traps to notify the SNMP server of link operational status changes and port role changes. Beginning in privileged EXEC mode, follow these steps to configure REP traps:

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 3	snmp mib rep trap-rate value Example: Router(config)# snmp mib rep trap-rate 500	Enables the router to send REP traps, and sets the number of traps 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 imposed; a trap is sent at every occurrence). Note To remove the traps, enter the no snmp mib rep trap-rate command.
Step 4	end Example: Router(config)# end	Returns to privileged EXEC mode.
Step 5	show running-config Example: Router# show running-config	(Optional) Displays the running configuration, which you can use to verify the REP trap configuration.
Step 6	copy running-config startup config Example: Router# copy running-config startup config	(Optional) Saves your entries in the router startup configuration file.

Monitoring REP

To monitor the REP configuration, complete the following steps:

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted.
Step 2	show interface [interface-id] rep [detail] Example: Router# show interface gigabitethernet0/1 rep detail	(Optional) Displays the REP configuration and status for a specified interface. <ul style="list-style-type: none"> Enter the physical interface and the optional detail keyword.
Step 3	show rep topology [segment segment-id] [archive] [detail] Example: Router# show rep topology	(Optional) Displays REP topology information for a segment or for all segments, including the primary and secondary edge ports in the segment. <ul style="list-style-type: none"> Enter the optional keywords and arguments, as desired.

Configuration Examples for REP

- [Configuring the REP Administrative VLAN: Example, page 14-16](#)
- [Configuring a REP Interface: Example, page 14-16](#)
- [Setting the Preemption for VLAN Load Balancing: Example, page 14-17](#)
- [Configuring SNMP Traps for REP: Example, page 14-17](#)
- [Monitoring the REP Configuration: Example, page 14-17](#)
- [Sample MWR 2941 Topology: Example, page 14-18](#)

Configuring the REP Administrative VLAN: Example

This example shows how to configure the administrative VLAN as VLAN 100.

```
Router# configure terminal
Router(config)# rep admin vlan 100
Router(config-if)# end
```

Configuring a REP Interface: Example

This example shows how to configure an interface as the primary edge port for segment 1, to send Spanning Tree Topology Changes Notification (STCNs) to segments 2 through 5, and to configure the alternate port as the port with port ID 0009001818D68700 to block all VLANs after a preemption delay of 60 seconds after a segment port failure and recovery.


```
Router# configure terminal
Router(config)# interface gigabitethernet0/1
Router(config-if)# rep segment 1 edge primary
Router(config-if)# rep stcn segment 2-5
Router(config-if)# rep block port 0009001818D68700 vlan all
Router(config-if)# rep preempt delay 60
Switch (config-if)# rep lsl-age-timer 6000
Router(config-if)# end
```

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

```
Router# configure terminal
Router(conf)# interface gigabitethernet0/1
Router(config-if)# rep segment 1 edge no-neighbor primary
Router(config-if)# rep stcn segment 2-5
Router(config-if)# rep block port 0009001818D68700 vlan all
Router(config-if)# rep preempt delay 60
Router(config-if)# rep lsl-age-timer 6000
```

This example shows how to configure the VLAN blocking configuration shown in [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 other VLANs are blocked at the primary edge port E1 (Gigabit Ethernet port 0/1).

```
Router# configure terminal
Router(config)# interface gigabitethernet0/1
Router(config-if)# rep segment 1 edge primary
Router(config-if)# rep block port 4 vlan 100-200
Router(config-if)# end
```


Figure 5 Example of VLAN Blocking

Setting the Preemption for VLAN Load Balancing: Example

The following is an example of setting the preemption for VLAN load balancing on a REP segment.

```
Router> enable
Router# configure terminal
Router(config)# rep preempt segment 1
Router(config)# end
```

Configuring SNMP Traps for REP: Example

This example configures the router to send REP traps at a rate of 10 traps per second:

```
Router> enable
Router# configure terminal
Router(config)# snmp mib rep trap-rate 10
Router(config)# end
```

Monitoring the REP Configuration: Example

The following is sample output of the **show interface rep detail** command. Use the **show interface rep detail** command on one of the REP interfaces to monitor and verify the REP configuration.

```
Router# show interface gigabitethernet0/1 rep detail

GigabitEthernet0/1 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
Load-balancing block port: none
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

```

Sample MWR 2941 Topology: Example

The following configuration example shows two Cisco MWR 2941 routers and two Cisco 7600 series routers using a REP ring.



Note

This section provides partial configurations intended to demonstrate a specific feature.

2941_1

```

interface GigabitEthernet0/0
  switchport trunk allowed vlan 1,2
  switchport mode trunk
  rep segment 1
!
interface GigabitEthernet0/1
  switchport trunk allowed vlan 1,2
  switchport mode trunk
  rep segment 1
!
interface GigabitEthernet0/3
  switchport access vlan 3
!
interface GigabitEthernet0/4
  switchport access vlan 4
!
interface Vlan1
  ip address 172.18.40.70 255.255.255.128
  no ptp enable
!
interface Vlan2
  ip address 1.1.1.1 255.255.255.0
  no ptp enable
!
interface Vlan3
  ip address 2.2.2.2 255.255.255.0
  no ptp enable
!
interface Vlan3
  ip address 4.4.4.2 255.255.255.0
  no ptp enable
!
ip route 3.3.3.0 255.255.255.0 1.1.1.4
ip route 5.5.5.0 255.255.255.0 1.1.1.4

```

2941_2

```

interface GigabitEthernet0/0
  switchport trunk allowed vlan 1,2
  switchport mode trunk

```

```
rep segment 1
!
interface GigabitEthernet0/1
 switchport trunk allowed vlan 1,2
 switchport mode trunk
 rep segment 1
!
interface Vlan1
 ip address 172.18.44.239 255.255.255.0
 no ptp enable
!
interface Vlan2
 ip address 1.1.1.2 255.255.255.0
 no ptp enable
```

7600_1

```
interface Port-channel69
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
!
interface GigabitEthernet3/25
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
 channel-group 69 mode on
!
interface GigabitEthernet3/26
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
 channel-group 69 mode on
!
interface GigabitEthernet3/35
 ip address 3.3.3.2 255.255.255.0
!
interface GigabitEthernet3/36
 ip address 5.5.5.2 255.255.255.0
!
interface GigabitEthernet5/2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
 rep segment 1 edge
!
interface Vlan1
 no ip address
!
interface Vlan2
 ip address 1.1.1.4 255.255.255.0
!
ip route 2.2.2.0 255.255.255.0 1.1.1.1
ip route 4.4.4.0 255.255.255.0 1.1.1.1
```

7600_2

```
interface Port-channel69
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
!
interface GigabitEthernet5/2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
 rep segment 1 edge
!
interface GigabitEthernet7/25
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
 channel-group 69 mode on
!
interface GigabitEthernet7/26
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,2
 switchport mode trunk
 channel-group 69 mode on
!
interface Vlan1
 no ip address
!
interface Vlan2
 ip address 1.1.1.3 255.255.255.0
!
```