Resilient Ethernet Protocol

First Published: January 31, 2008
Last Updated: November 30, 2010

Resilient Ethernet Protocol (REP) is a Cisco proprietary protocol that provides an alternative to Spanning Tree Protocol (STP). REP provides a way 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.

The router supports REP only when the router is running the metro IP access or metro access image.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "Feature Information for Resilient Ethernet Protocol" section on page 22.

Use Cisco Feature Navigator to find information about platform support and software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Contents

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Information About REP

REP Segments

One 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 router 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 Layer 2 trunk interfaces.

Figure 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 failure in the network, the blocked port returns to the forwarding state to minimize network disruption.

Figure 1 REP Open Segments

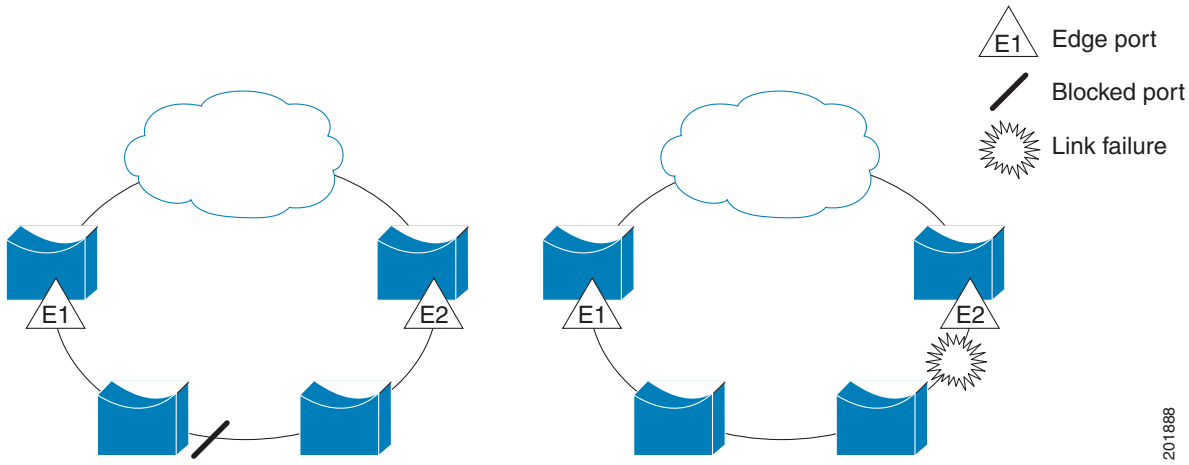

The segment shown in Figure 1 is an open segment; there is no connectivity between the two edge ports. The REP segment cannot cause a bridging loop and it is safe to connect the segment edges to any network. All hosts connected to routers inside the segment have two possible connections to the rest of the network through the edge ports, but only one connection is accessible at any time. If a failure occurs on any segment or any port on a REP segment, REP unblocks all ports to ensure that connectivity is available through the other gateway.

The segment shown in Figure 2, with both edge ports located on the same router, 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 routers in the segment.

Figure 2 REP Ring Segment

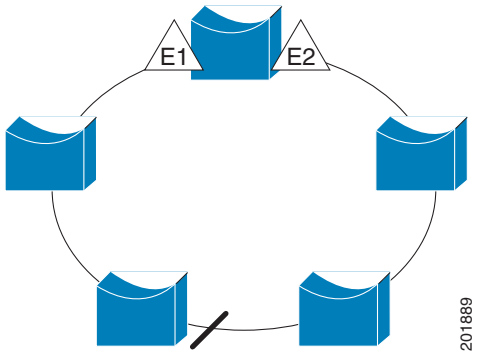

REP segments have the following characteristics:

- 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.

REP has the following 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 causes 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. When enabled on an interface, 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 as operational after it performs a three-way handshake with a neighbor in the same segment. A segment port does not become operational under the following conditions:

- 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. Once 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 PortFast Bridge Protocol Data Unit (BPDU) class MAC address. The packets can also be sent to the Cisco multicast address, which at present 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 the messages 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 is less than 200 milliseconds (ms) for the local segment.

VLAN Load Balancing

One edge port in the REP segment acts as the primary edge port; the other as the secondary edge port. It is the primary edge port that always participates in VLAN load balancing in the segment. REP VLAN load 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:

- 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** command for the port.
- 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 You configure offset numbers on the primary edge port by identifying a port's 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.

- 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** command.

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

- You can manually trigger VLAN load balancing at any time by entering the **rep preempt segment** *segment-id* command on the router that has the primary edge port.
- You can configure a preempt delay time by entering the **rep preempt delay** *seconds* 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 then sends out a message to alert all interfaces in the segment about the preemption. When the message is received by the secondary edge port, 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. VLAN load balancing is initiated only by the primary edge port and 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.

To reconfigure VLAN load balancing, you reconfigure the primary edge port. When you change the VLAN-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 VLAN load balancing 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 Protocol Interaction

REP does not interact with Spanning Tree Protocol (STP) or with Flex Links, but can coexist with both of them. 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.

REP Ports

Ports in REP segments take one of three roles or states: Failed, Open, or Alternate.

- A port configured as a regular segment port starts as a failed port.
- Once the neighbor adjacencies are determined, the port transitions 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 failed 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 to the spanning tree configuration. By default, this is a designated blocking port. If the PortFast BPDU Guard Enhancement feature is configured or if STP is disabled, the port goes into the forwarding state.

REP Integrated with VPLS

Normally, in a Virtual Private LAN Services (VPLS) network core, all nodes are connected in a full-mesh topology and each node has connectivity to all other nodes. In the full-mesh topology, there is no need for a node to retransmit data to another node. In Figure 3, the common ring provides a path where the packet could be forwarded to another network provider edge (N-PE) router, breaking the split horizon model.

Figure 3 REP Integrated with VPLS

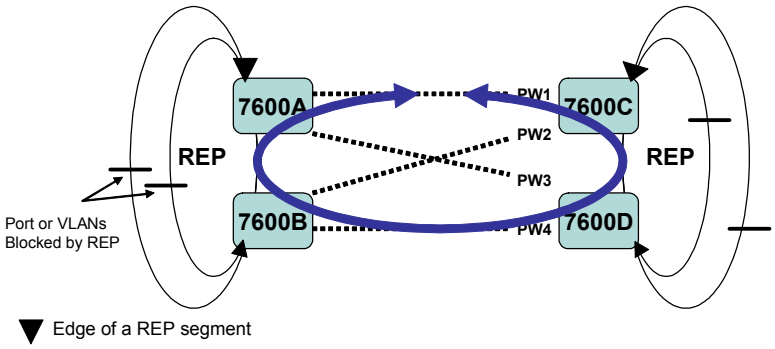

REP emulates a common link connection that so that the REP ring supports the VPLS full-mesh model, but maintains the split horizon properties so that the super-loop does not exist. The emulated common link uses the Clustering over the WAN (CWAN) line card, which is also used for the VPLS uplink. This emulated common link forwards data from the ring to either the VPLS uplink or to the other side of the ring; blocks data coming from the VPLS core network; handles access pseudo-wire for H-VPLS topologies.

REP Integrated with an EVC Port

REP can be integrated with an EVC port by using the REP over Ethernet Virtual Circuit (EVC) feature. The REP over EVC feature, intended for use on the Cisco 7600 series router, allows you to configure and manage ports at service level. An EVC port can have multiple service instances. Each service instance corresponds to a unique Event Flow Processor (EFP). By default, REP is disabled on all ports. Using the REP over EVC feature, you can:

- Control data traffic
- Configure VLAN load balancing at service instance level.

The ports on a Cisco 7600 series router are classified into three different types: switchports, routed ports, and EVC ports. By default, a port is a routed port. REP is not supported on routed ports. You need to configure a port to a switchport or EVC port to configure REP on it. A port that is configured with one or more service instances is called an EVC port.

This feature allows you to configure an EVC port to participate in a REP segment. REP can selectively block or forward data traffic on particular VLANs. For EVC, the VLAN Id refers to the outer tag of the dot1q encapsulation that is configured on a service instance. REP is supported on a bridge-domain service. If the **ethernet vlan color-block all** command is configured, REP is supported on connect and xconnect services.

For more information about this feature, see the "Configuring Layer 1 and Layer 2 Features" chapter in the *Cisco 7600 Series Ethernet Services Plus (ES+) and Ethernet Services Plus T (ES+T) Line Card Configuration Guide*.

REP-Configurable Timers

Cisco IOS Release 15.0(1)S includes the REP Configurable Timer (REP Fast Hellos) feature. This feature allows you to configure a failure detection time in the range of 120 ms to 10,000 ms. This feature is intended for use on the Cisco 7600 series router.

For more information about this feature, see the "Configuring Layer 1 and Layer 2 Features" chapter in the *Cisco 7600 Series Ethernet Services Plus (ES+) and Ethernet Services Plus T (ES+T) Line Card Configuration Guide*.

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 Segments and REP Administrative VLANs

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 should 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, with one of them 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 of them to serve as the segment primary edge port. You can also optionally configure where to send segment STCNs and VLAN load balancing. For more information about configuring REP Administrative VLANS, see the "Configuring the REP Administrative VLAN" section on page 10.

REP Configuration Guidelines

Follow these guidelines when configuring REP:

- Cisco recommends 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** 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 IEEE 802.1Q or ISL trunk ports.
- You must configure all trunk ports in the segment with the same set of allowed VLANs, or a misconfiguration occurs.
- 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 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.
- If REP is enabled on two ports on a router, both ports must be either regular segment ports or edge ports. REP ports follow these rules:
  - If only one port on a router is configured in a segment, the port should be an edge port.
  - If two ports on a router belong to the same segment, both ports must be edge ports or both ports must be regular segment ports.
  - If two ports on a router 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 remains 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 protocol data units (PDUs) in untagged frames on the native VLAN. The BPA message sent to the Cisco multicast address is sent on the administrative VLAN, which is VLAN 1 by default.
- REP ports can not be configured as one of these port types:
  - Switched Port Analyzer (SPAN) destination port
  - Private VLAN port
  - Tunnel port
  - Access port
- On a Cisco ME-3400 series router, REP ports must be network node interfaces (NNI). User-network interfaces (UNIs) cannot be REP ports.
- REP is supported on EtherChannels, but not on an individual port that belongs to an EtherChannel.
- There is a maximum of 64 REP segments per router.

How to Configure REP

Configuring the REP Administrative VLAN

To configure the REP administrative VLAN, complete the following steps:

Guidelines for Configuring the REP Administrative VLAN

To avoid the delay introduced by relaying messages related to link-failure or VLAN-blocking notification during VLAN 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 router and on a segment. However, this is not enforced by the software.
- The administrative VLAN cannot be the Remote SPAN (RSPAN) VLAN.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**rep admin vlan** *vlan-id* **Example:** `Router(config)# rep admin vlan 1`	Configures a REP administrative VLAN. • Specify the administrative VLAN. The range is 1 to 4094. The default is VLAN 1.

	Command or Action	Purpose
Step 4	**end** **Example:** `Router(config)# end`	Return to privileged EXEC mode.
Step 5	**show interface** [*interface-id*] **rep** [**detail**] **Example:** `Router# show interface gigabitethernet0/1 rep detail`	Displays the REP configuration and status for a specified interface. • Enter the physical interface or port channel ID.
Step 6	**copy running-config startup config** **Example:** `Router# copy running-config startup config`	(Optional) Save your entries in the router startup configuration file.

Configuring REP Interfaces

To enable and configure REP on an interface, complete the following steps:

Prerequisite

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

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *interface-id*
4. **port-type nni**
5. **switchport mode trunk**
6. **rep segment** *segment-id* [**edge** [**primary**]] [**preferred**]
7. **rep stcn** {**interface** *interface-id* | **segment** *id-list* | **stp**}
8. **rep block port** {**id** *port-id* | *neighbor-offset* | **preferred**} **vlan** {*vlan-list* | **all**}
9. **rep preempt delay** *seconds*
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:** `Router> enable`	Enables privileged EXEC mode. • 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. • Enter the interface ID. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port channel range is a number from 1 to 48.
Step 4	**port-type nni** **Example:** `Router(config-if)# port-type nni`	Configures the port as a network node interface (NNI).
Step 5	**switchport mode trunk** **Example:** `Router(config-if)# switchport mode trunk`	Configures the interface as a Layer 2 trunk port. On non-ES ports, you can also configure the encapsulation type by entering the switchport trunk encapsulation {isl | dot1q | negotiate} interface configuration command. ES ports support only IEEE 802.1Q encapsulation.

	Command or Action	Purpose
Step 6	**rep segment** *segment-id* [**edge** [**primary**]] [**preferred**] **Example:** `Router(config-if)# rep segment 1 edge preferred`	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. • 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. • 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. **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. • Enter **preferred** to indicate that the port is the preferred alternate port or the preferred port for VLAN load balancing. **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.
Step 7	**rep stcn** {**interface** *interface-id* | **segment** *id-list* | **stp**} **Example:** `Router(config-if)# rep stcn segment 2-5`	(Optional) Configures the edge port to send STCNs. • Enter **interface** *interface-id* to designate a physical interface or port channel to receive STCNs. • Enter **segment** *id-list* to identify one or more segments to receive STCNs. The range is 1 to 1024. • Enter **stp** to send STCNs to STP networks.

	Command or Action	Purpose
Step 8	**rep block port** {**id** *port-id* | *neighbor-offset* | **preferred**} **vlan** {*vlan-list* | **all**} **Example:** `Router(config-if)# rep block port 0009001818D68700 vlan all`	(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. • 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. **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. • 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. **Note** Enter this command only on the REP primary edge port.
Step 9	**rep preempt delay** *seconds* **Example:** `Router(config-if)# rep preempt delay 60`	(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. **Note** Use this command only on the REP primary edge port.
Step 10	**end** **Example:** `Router(config-if)# end`	Returns to privileged EXEC mode.
Step 11	**show interface** [*interface-id*] **rep** [**detail**] **Example:** `Router(config-if)# show interface gigabitethernet0/1 rep detail`	Verifies the REP interface configuration. • Enter the interface ID and the optional **detail** keyword, if desired.
Step 12	**copy running-config startup config** **Example:** `Router(config-if)# copy running-config startup config`	(Optional) Saves your entries in the router startup configuration file.

Setting the Preemption for VLAN Load Balancing

To set the preemption for VLAN load balancing, complete these steps on the router that has the segment with the primary edge port.

Restrictions

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. Use the **show rep topology** privileged EXEC command to see which port in the segment is the primary edge port.

Prerequisite

Be sure that all other segment configuration has been completed before setting preemption for VLAN load balancing. When you enter the **rep preempt segment** *segment-id* command, a confirmation message appears before the command is executed because preemption for VLAN load balancing can disrupt the network.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**rep preempt segment** *segment-id* **Example:** `Router(config)# rep preempt segment 1`	Manually triggers VLAN load balancing on the segment. • Enter the segment ID. **Note** You will be asked to confirm the action before the command is executed.

	Command or Action	Purpose
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 router to send REP-specific traps to notify the SNMP server of link operational status changes and any port role changes. To configure REP traps, complete the following steps:

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
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. • 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.

	Command or Action	Purpose
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 the REP Configuration

To monitor the REP configuration, complete the following steps:

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • 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. • Enter the physical interface or port channel ID, and the optional **detail** keyword, if desired.
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. • Enter the optional keywords and arguments, as desired.

Configuration Examples for REP

Example: Configuring the REP Administrative VLAN

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
```

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 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
Router(config-if)# end
```

This example shows how to configure the VLAN blocking configuration shown in Figure 4. 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 4 Example of VLAN Blocking

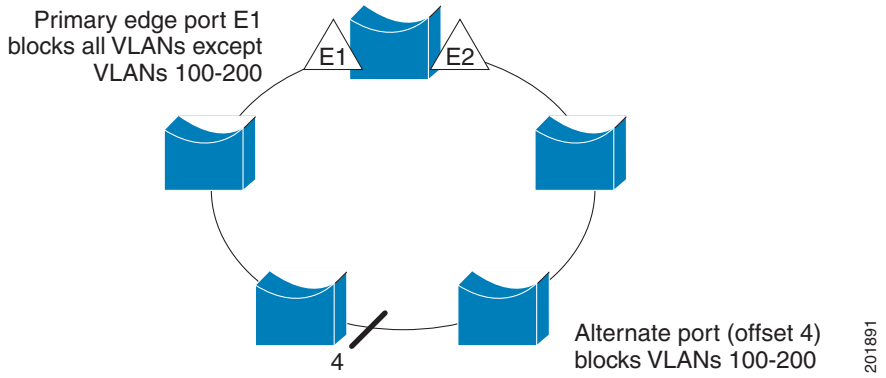

Example: Setting the Preemption for VLAN Load Balancing

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
```

Example: Configuring SNMP Traps for REP

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
```

Example: Monitoring the REP Configuration

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
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	*Cisco IOS Master Commands List, All Releases*
LAN Switching commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples.	*Cisco IOS LAN Switching Command Reference*
Flex links	"Configuring Flex Links" chapter of the *Catalyst 6500 Release 12.2SXH and Later Software Configuration Guide*
Introduction to spanning tree protocols	*Spanning Tree Protocol (STP)/802.1D*
Spanning Tree PortFast BPDU Guard Enhancement feature	*Spanning Tree PortFast BPDU Guard Enhancement*
Cisco IOS 7600 series routers	*Cisco 7600 Series Ethernet Services Plus (ES+) and Ethernet Services Plus T (ES+T) Line Card Configuration Guide*

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Resilient Ethernet Protocol

Table 1 lists the release history for this feature.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Note Table 1 lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Table 1 Feature Information for Resilient Ethernet Protocol

Feature Name	Releases	Feature Information
Resilient Ethernet Protocol	12.2(44)SE 12.2(33)SRC	REP is a Cisco proprietary protocol that provides an alternative to Spanning Tree Protocol (STP); it provides a way 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. In 12.2(44)SE, this feature was introduced on the Cisco 3750 series router. In 12.2(33)SRC, support was added for the Cisco 7600 series router and included integration with VPLS core networks. The following commands were introduced or modified by this feature: **rep admin vlan**, **rep block port**, **rep preempt delay**, **rep preempt segment**, **rep segment**, **rep stcn**, **show rep topology**.
REP Integration with EVC and VPLS	12.2(33)SRE	REP can be integrated with an EVC and VPLS by using the REP over Ethernet Virtual Circuit (EVC) feature. The REP over EVC feature, intended for use on the Cisco 7600 series router, allows you to configure and manage ports at the service level. For more information, see the following section: • REP Integrated with an EVC Port, page 7 For detailed information about this feature, see the "Configuring Layer 1 and Layer 2 Features" chapter in the *Cisco 7600 Series Ethernet Services Plus (ES+) and Ethernet Services Plus T (ES+T) Line Card Configuration Guide*. No commands were introduced or modified by this feature.

Table 1 Feature Information for Resilient Ethernet Protocol (continued)

Feature Name	Releases	Feature Information
REP Configurable Timer (REP Fast Hellos)	15.0(1)S	The REP Configurable Timer (REP Fast Hellos) feature allows you to configure a failure detection time in the range of 120 milliseconds (ms) to 10,000 ms. This feature is intended for use on the Cisco 7600 series router. For more information, see the following section: • REP-Configurable Timers, page 8 For detailed information about this feature, see the "Configuring Layer 1 and Layer 2 Features" chapter in the *Cisco 7600 Series Ethernet Services Plus (ES+) and Ethernet Services Plus T (ES+T) Line Card Configuration Guide*. The following commands were introduced by this feature: **rep lsl-age-timer**, **rep lsl-retries**.