



CHAPTER 20

Configuring Resilient Ethernet Protocol

This chapter describes how to use Resilient Ethernet Protocol (REP) on the Catalyst 3750 Metro switch. 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.

This chapter includes these sections:


- [Understanding REP, page 20-1](#)
- [Configuring REP, page 20-6](#)
- [Monitoring REP, page 20-12](#)

Understanding REP

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 switch can have no more than two ports that belong to the same segment, and each segment port can have only one external neighbor. A segment can go through a shared medium, but on any link only two ports can belong to the same segment. REP is supported only on Layer 2 trunk interfaces.

[Figure 20-1](#) 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, as shown in the diagram on the right, the blocked port returns to the forwarding state to minimize network disruption.

Figure 20-1 REP Open Segments




201888

The segment shown in Figure 20-1 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 switches inside the segment have two possible connections to the rest of the network through the edge ports, but only one connection is accessible at any time. If a failure causes a host to be unable to access its usual gateway, REP unblocks all ports to ensure that connectivity is available through the other gateway.

The segment shown in Figure 20-2, 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 20-2 REP Ring Segment



201889

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.

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. 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 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. 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 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 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 is less than 200 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 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** interface configuration 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.

Figure 20-3 shows neighbor offset numbers for a segment where E1 is the primary edge port and E2 is the secondary edge port. The red numbers inside the ring are numbers offset from the primary edge port; the black numbers outside of the ring show the offset numbers from the secondary edge port. Note that you can identify all 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 20-3 Neighbor Offset Numbers in a Segment



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** privileged EXEC command on the switch that has the primary edge port.

- You can 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 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 load balancing, you reconfigure the primary edge port. 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.

Spanning Tree Interaction

REP does not interact with STP or with the Flex Link feature, 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 PortFast is configured or if STP is disabled, the port goes into the forwarding state.

Configuring 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 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 topology change notices (STCNs) and VLAN load balancing.

This section includes this information:

- [Default REP Configuration, page 20-6](#)
- [REP Configuration Guidelines, page 20-6](#)
- [Configuring the REP Administrative VLAN, page 20-7](#)
- [Configuring REP Interfaces, page 20-8](#)
- [Setting Manual Preemption for VLAN Load Balancing, page 20-11](#)
- [Configuring SNMP Traps for REP, page 20-12](#)

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 IEEE 802.1Q or ISL 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 it, you might lose connectivity to the switch if you enable REP in a Telnet session that accesses the switch 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.
- You must configure all trunk ports in the segment with the same set of allowed VLANs, or a misconfiguration occurs.
- If REP is enabled on two ports on a switch, both ports must be either regular segment ports or edge ports. REP ports follow these rules:
 - 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, both ports must be edge ports or both ports must be regular segment ports.
 - 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 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 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.
- REP ports can not be configured as one of these port types:
 - SPAN destination port
 - Private VLAN port
 - Tunnel port
 - Access port
- 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 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.
- The administrative VLAN cannot be the RSPAN VLAN.

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

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	rep admin vlan <i>vlan-id</i>	Specify the administrative VLAN. The range is 2 to 4094. The default is VLAN 1. To set the admin VLAN to 1, enter the no rep admin vlan global configuration command.
Step 3	end	Return to privileged EXEC mode.
Step 4	show interface [<i>interface-id</i>] rep detail	Verify the configuration on one of the REP interfaces.
Step 5	copy running-config startup config	(Optional) Save your entries in the switch startup configuration file.

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:

```
Switch# configure terminal
Switch (conf)# rep admin vlan 100
Switch (conf-if)# end

Switch# 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: 0002001121A2D580E4D
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
```

Configuring REP Interfaces

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 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	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface, and enter interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 48.
Step 3	switchport mode trunk	Configure 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.
Step 4	rep segment <i>segment-id</i> [edge [primary]] [preferred]	Enable REP on the interface, and identify a segment number. The segment ID range is from 1 to 1024. These optional keywords are available. Note You must configure two edge ports, including one primary edge port for each segment. <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) 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. <ul style="list-style-type: none"> (Optional) 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 5	rep stcn { interface <i>interface-id</i> segment <i>id-list</i> stp }	(Optional) Configure the edge port to send segment topology change notices (STCNs). <ul style="list-style-type: none"> Enter interface <i>interface-id</i> to designate a physical interface or port channel to receive STCNs. Enter segment <i>id-list</i> to identify one or more segments to receive STCNs. The range is 1 to 1024. Enter stp to send STCNs to STP networks.

	Command	Purpose
Step 6	rep block port { <i>id port-id</i> <i>neighbor_offset</i> preferred } vlan { <i>vlan-list</i> all }	<p>(Optional) Configure VLAN load balancing on the primary edge port, identify the REP alternate port in one of three ways, and configure the VLANs to be blocked on the alternate port.</p> <ul style="list-style-type: none"> Enter the id <i>port-id</i> 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 <i>neighbor_offset</i> 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. See Figure 20-3 on page 20-4 for an example of neighbor offset numbering. <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 <i>vlan-list</i> 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>
Step 7	rep preempt delay <i>seconds</i>	<p>(Optional) You must enter this command and configure a preempt time delay 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 Enter this command only on the REP primary edge port.</p>
Step 8	end	Return to privileged EXEC mode.
Step 9	show interface [<i>interface-id</i>] rep [detail]	Verify the REP interface configuration.
Step 10	copy running-config startup config	(Optional) Save your entries in the switch startup configuration file.

Enter the **no** form of each command to return to the default configuration. Enter the **show rep topology** privileged EXEC command to see which port in the segment is the primary edge port.


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.

```
Switch# configure terminal
Switch (conf)# interface gigabitethernet0/1
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)# end
```

This example shows how to configure the VLAN blocking configuration shown in Figure 20-4. The alternate port is the neighbor with neighbor offset number 4. After manual preemption, VLANs 100-200 are blocked at this port and all other VLANs are blocked at the primary edge port E1 (Gigabit Ethernet port 0/1).

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

Figure 20-4 Example of VLAN Blocking



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 other segment configuration has been completed 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.

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	Purpose
Step 1	rep preempt segment <i>segment-id</i>	Manually trigger VLAN load balancing on the segment. You will need to confirm the command before it is executed.
Step 2	show rep topology	View 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	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>snmp mib rep trap-rate value</code>	Enable the switch to send REP traps, and set 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).
Step 3	<code>end</code>	Return to privileged EXEC mode.
Step 4	<code>show running-config</code>	Verify the REP trap configuration.
Step 5	<code>copy running-config startup config</code>	(Optional) Save your entries in the switch startup configuration file.

To remove the trap, enter the `no snmp mib rep trap-rate` global configuration command.

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

```
Switch(config)# snmp mib rep trap-rate 10
```

Monitoring REP

Use the privileged EXEC commands in [Table 20-1](#) to monitor REP.

Table 20-1 REP Monitoring Commands

Command	Purpose
<code>show interface [interface-id] rep [detail]</code>	Displays REP configuration and status for a specified interface or for all interfaces.
<code>show rep topology [segment segment_id] [archive] [detail]</code>	Displays REP topology information for a segment or for all segments, including the primary and secondary edge ports in the segment.