



Configuring Resilient Ethernet Protocol for IE 4000, IE 4010, and IE 5000 Switches

First Published: 2023-05-12

Last Modified: 2024-02-22

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883



CONTENTS

CHAPTER 1

Information About Configuring REP 1

- REP Overview 1
- Link Integrity 3
- REP Negotiated 4
- Fast Convergence 4
- VLAN Load Balancing 4
- Spanning Tree Interaction 6
- REP Ports 6
- REP Zero Touch Provisioning 7
 - REP and Day Zero 7
 - REP ZTP Overview 9
- REP Segment-ID Autodiscovery 10
 - REP Segment-ID Autodiscovery Deployment 11
 - REP Segment-ID Autodiscovery Limitations 11

CHAPTER 2

Default REP Configuration 13

- Default REP Configuration 13
- REP Configuration Guidelines 13
- REP Administrative VLAN 15

CHAPTER 3

How to Configure REP 17

- Configuring the REP Administrative VLAN 17
- Configuring REP Interfaces 17
- Configuring REP Negotiated 19
- Setting Manual Preemption for VLAN Load Balancing 21
- Configuring SNMP Traps for REP 22

Configuring REP ZTP 22

Configuring REP Segment-ID Autodiscovery 23

 Enable REP Segment-ID Autodiscovery 23

 Configure the Interfaces 24

CHAPTER 4 **Monitoring and Maintaining REP 25**

 Monitoring and Maintaining REP 25

 Displaying REP ZTP Status 25

 View REP Segment ID Autodiscovery Status 29

CHAPTER 5 **Configuration Examples for Configuring REP 31**

 Configuring the Administrative VLAN: Example 31

 Configuring a Primary Edge Port: Examples 31

 Configuring VLAN Blocking: Example 32

CHAPTER 6 **Feature History 33**

 Feature History 33



CHAPTER 1

Information About Configuring REP

- [REP Overview, on page 1](#)
- [Link Integrity, on page 3](#)
- [REP Negotiated, on page 4](#)
- [Fast Convergence, on page 4](#)
- [VLAN Load Balancing, on page 4](#)
- [Spanning Tree Interaction, on page 6](#)
- [REP Ports, on page 6](#)
- [REP Zero Touch Provisioning, on page 7](#)
- [REP Segment-ID Autodiscovery, on page 10](#)

REP Overview


This document provides details about configuring Resilient Ethernet Protocol (REP) on the Cisco Industrial Ethernet 4000 Series, Cisco Industrial Ethernet 4010 Series, and Cisco Industrial Ethernet 5000 Series switches.

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

A REP segment is a chain of ports connected to each other and configured with a segment ID. Each segment consists of standard (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.

The following diagram 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 1: REP Open Segments



The segment shown above 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 the next example, 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 2: REP Ring Segment



REP segments have these characteristics:

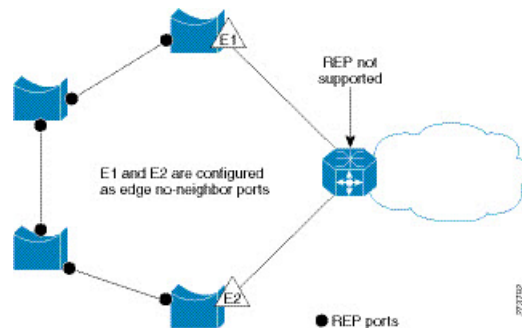
- If all ports in the segment are operational, one port (referred to as the *alternate* port) is in the blocked state for each VLAN. If VLAN load balancing is configured, two ports in the segment control the blocked state of VLANs.
- If one or more ports in a segment is not operational, causing a link 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 the following diagram. 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 3: Edge No-Neighbor Ports



REP has these limitations:

- You must configure each segment port; an incorrect configuration can cause forwarding loops in the networks.
- REP can manage only a single failed port within the segment; multiple port failures within the REP segment cause loss of network connectivity.
- You should configure REP only in networks with redundancy. Configuring REP in a network without redundancy causes loss of connectivity.

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 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 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.

REP Negotiated

Spanning Tree Protocol (STP) is enabled by default on Cisco switches. If a Cisco switch is inserted in an already running REP ring (for example, to add a new node or replace an existing node), the new switch running STP will cause a break in the REP ring and cannot communicate over the REP ring until it is configured to be part of the ring.

After a new switch is inserted in the ring, it is running STP, but the rest of the ring is running REP. Neither of these protocols can recognize a loop in the ring and keep both ends of the ring in the forwarding state, causing an endless loop. To address this problem, **rep bpduleak** should be configured on the new switch so that REP BPDUs are transparently forwarded between two ring ports on the switch when REP is not configured. This function, called BPDU leaking, causes the REP ring to converge but new devices will not be part of the ring nor be seen in or show the REP topology.

When the switch interfaces are configured with REP Negotiated, REP status is negotiated with the peers. If the peer supports REP, it is migrated to REP. If the peer does not support REP, it is migrated to STP. The peer is migrated to REP or STP using an Embedded Event Manager (EEM) macro.



Note REP Negotiated works only on uplink ports.

See [Configuring REP Negotiated](#) for information about configuring REP Negotiated.

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 50-200 ms for the local segment with 200 VLANs configured. When REP is configured on RJ45 Gigabit copper interfaces, the convergence time is 500-750 ms. 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. 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.

The figure below 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 4: Neighbor Offset Numbers in a Segment



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 preemption time period elapses. Note that the delay timer restarts if another port fails before the time has elapsed.



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

When VLAN load balancing is triggered, the primary edge port sends out a message to alert all 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.

Spanning Tree Interaction

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

To migrate from an STP ring configuration to 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 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 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.

REP Zero Touch Provisioning

Before a network device such as a router or a switch is deployed online and fully functional, a fair amount of manual configuration is required. Zero Touch Provisioning (ZTP) technologies automate these processes, bringing up network devices into a functional state with minimal to no manual configuration. The Cisco Network Plug and Play (PnP) and Autoinstall Day Zero solutions provide a simple, secure, unified, and integrated offering for enterprise and industrial network customers to ease device rollouts for provisioning updates to an existing network. However, PnP does not support Resilient Ethernet Protocol (REP) due to the way REP is designed. Prior to the REP ZTP feature, REP ring provisioning for Day Zero required manual intervention. The REP ZTP feature introduces a new type-length-value (TLV) extension into the REP LSL packets to support configuring REP rings with zero-touch technologies.



Note REP ZTP is supported in Cisco IOS release 15.2(8)E4 and later.


Cisco DNA Center does not support the use case of REP ZTP and classic Cisco IOS IE devices.

REP and Day Zero

In a typical switch deployment using ZTP, the switch, with no startup configuration in the NVRAM, triggers the Cisco Open Plug-n-Play (PnP) agent to initiate a DHCP discovery process. This process acquires the IP configuration required for the switch from the DHCP server. The DHCP server can be configured to insert additional information in a DHCP message using vendor specific option 43. After the DHCP server receives a DHCP DISCOVER message with option 60 and the string "cisco pnp" from the switch, the DHCP server sends the IP address or hostname of the PnP server to the requesting switch. When the switch receives the DHCP response, the PnP agent extracts the option 43 from the response to get the IP address or the hostname of the PnP server. The PnP agent on the switch then uses this IP address or hostname to communicate with the PnP server. Finally, the PnP server downloads the required Day Zero configuration to the switch to complete the provisioning.

The example shown in the following diagrams illustrates REP ring provisioning on Day Zero, prior to the introduction of REP ZTP.

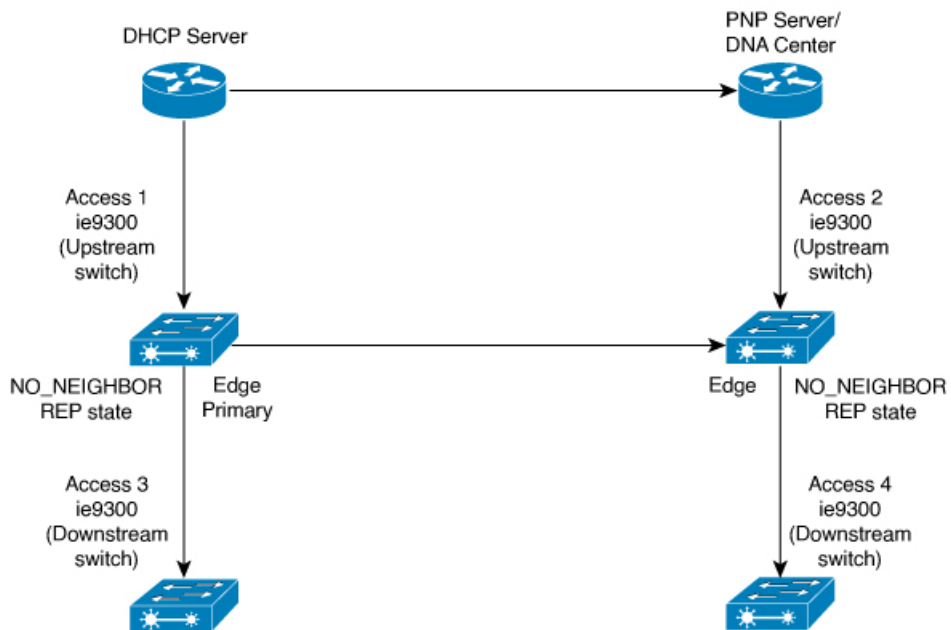
Figure 5: Adding Edge Nodes to the REP Ring



Note The DHCP Server and the PnP Server/DNA Center are not part of the REP ring.


The first set of nodes to be provisioned are Access 1 and Access 2 in the diagram. These are the 2 edge nodes of the REP ring. Note that PnP has configured the downlink port as primary edge on Access 1 and secondary edge on Access 2.

Figure 6: Adding Downstream Nodes



When either Access 3 or Access 4 are powered on, the REP edge primary port starts the REP protocol negotiation and discovers that the neighbor port is not a REP enabled port. (Recall that the switch will be added to the REP ring only after PnP provisioning, for which it needs to first contact the DHCP server as explained earlier.) When an upstream switch port has REP configured and a downstream switch is getting on-boarded with PnP, the REP port goes into the NO_NEIGHBOR state because it is not able to discover its REP peer. In the NO_NEIGHBOR state, REP blocks all the VLANs on that port. This means that the DHCP discovery message from the new switch on the PnP startup VLAN is dropped by the upstream switch because its REP state is NO_NEIGHBOR. The same sequence of blocked ports continues for all new switches added to the REP ring (see Access 5 in figure below).

Figure 7: NO_NEIGHBOR REP State



REP ZTP Overview

The REP ZTP enhancements require that both the upstream and the downstream switches support the feature. When the new downstream switch is powered on, it initiates PnP/autoinstall. The upstream switch's interface is configured for REP and blocks the interface to the downstream switch because the downstream switch is not REP by default (the upstream switch is in REP_NO_NEIGHBOR state).

Even though the interface on the upstream switch is blocked, it will transmit REP LSL packets to the downstream switch. This is normal. With the enhancement of the REP ZTP feature, the downstream switch will start transmitting REP LSL packets with a new TLV to inform the upstream switch that its neighbor is attempting PnP provisioning.

When the upstream switch reads this REP LSL with the new TLV, it will unblock the interface for the PnP startup VLAN only. All other VLANs for which the upstream interface is a member continue to be blocked. Because the upstream switch is forwarding packets on the PnP startup VLAN for this interface, the downstream switch can complete the PnP process.

The intent of this feature is to allow new switches to join a REP ring with no manual intervention. The interface on the upstream switch keeps the startup VLAN unblocked until the downstream switch has received its configuration and has configured its own interface for REP. If there's a failure in the PnP process, the interface on the upstream switch reverts to blocking on the PnP startup VLAN. If the configuration received by the downstream switch does configure the interface for REP, the upstream switch reverts to blocking the PnP startup VLAN.

The downstream behavior to transmit the REP LSL with new TLV to request the PnP startup VLAN be unblocked is the default behavior for switches with no startup configuration. For security purposes, the upstream switch must have the interface to the downstream switch explicitly enabled to put the PnP startup VLAN into unblocked state. The interface level command is **rep ztp-enable**. See [Configuring REP ZTP, on page 22](#).



Note The upstream switch can be part of multiple REP rings and thereby connected to multiple downstream neighbours. The PnP startup VLAN is unblocked only on the interfaces to which the downstream switch is connected.

REP Segment-ID Autodiscovery

Resilient Ethernet Protocol (REP) Segment-ID Autodiscovery enables automatic configuration and continued static configuration of segment IDs in REP segments. This feature is available on switches running Classic Cisco IOS as well as switches running Cisco IOS XE and is interoperable in these scenarios:

- Adding a Classic Cisco IOS switch into a REP ring of Cisco IOS XE switches
- Adding a Cisco IOS XE switch into a REP ring of Classic Cisco IOS switches

A REP segment is a chain of ports that are connected to each other and configured with a segment ID. Forming multiple REP segments statically by configuring each port of the device is a manual task, and any mismatch in configuring the segment ID leads to convergence issues. However, REP Segment-ID Autodiscovery adds new CLI commands to enable a switch to learn and retain segment ID information automatically.

You can use REP Segment-ID Autodiscovery in several different scenarios. You can insert a new switch into an existing REP segment or in a new REP segment that you build yourself. The feature is ideal for multiple REP ring deployments when incorrect REP Segment IDs might be entered manually. Such errors can occur when deploying multiple REP rings from the same REP seed node.

See the following sections in this guide for more information:

- [REP Segment-ID Autodiscovery Deployment](#)
- [Configuring REP Segment-ID Autodiscovery](#)

REP Segment-ID Autodiscovery Deployment

You can configure REP Segment-ID Autodiscovery when you add a switch to a REP segment or when you create a REP segment. In either case, the feature reduces the amount of manual configuration that you must do.

Adding a new Switch to an REP Segment

When you add a switch to an existing REP segment, you enable autodiscovery by entering the **rep autodisc** command on the switch interfaces connecting to the upstream and downstream switches.

When the new switch is connected to the upstream and downstream switches, the upstream and downstream switches send CDP packets with REP segment ID information to the new switch interfaces. You enter the command **rep segment auto** on the new switch interfaces so they can learn the segment ID.

Building a new REP Segment

When you build a closed REP segment, you must start with a static REP segment ID configuration from an edge device. The primary and secondary edge devices in a closed segment are on the same switch. When you build an open REP segment, you must start a static REP segment ID configuration from both primary and secondary edge devices.

The remaining steps are the same for both closed and open REP segments. You bring up the next node in the REP ring. You then add any next new node between these two switches for autodiscovery to work correctly.

Building a REP Segment with Uplinks

When you build a ring segment with uplinks (daisy chain), you must start with a static REP segment ID configuration from the REP edge node. Connect the next device to one of the uplinks to the edge node, and enable autodiscovery on the connected uplink. Because of port pairing support, the same REP configuration is duplicated on the paired uplink port.

When the next device is connected with the uplink, the process repeats to bring the REP segment in a daisy chain manner. Each new REP node automatically joins the ring by learning the REP Segment ID from the node above it. For a REP open ring, the last device on the segment is an edge device with static REP configuration.

REP Segment-ID Autodiscovery Limitations

The following are restrictions for the REP Segment-ID Autodiscovery feature:

- REP Segment-ID Autodiscovery feature is supported in Cisco IOS release 15.2(8)E4 and later.
- The only supported port-pairing is the first two uplink ports. No predefined port pairing is supported for downlinks.

If you configure a REP segment on a downlink port, the switch receives the segment ID from the upstream switch, and the partner downlink port is connected to the same segment. However, the switch does not pass the segment ID to its partner port. Instead, you must explicitly configure the partner port of the downlink pair.

- The REP Segment-ID Autodiscovery feature is not supported when you insert an edge node into the existing segment. You must configure static or manual REP segment ID on primary and secondary edge devices.

- If you insert a new switch between two switches that are part of a segment, you must connect the new switch interfaces to the interfaces of existing switches that transmit the same segment ID. Any incorrect connections to other interfaces of the existing switches leads to segment failure.

For example, assume gi1/1 of switch1 and gi1/2 of switch2 are connected as a part an existing segment, and switch3 is inserted between these two switches. In such a case, you must ensure that the interfaces are connected to gi1/1 of switch1 and gi1/2 of switch2 to include switch3 as a part of the same segment.

- If you configure REP automatically on an interface with the **rep segment auto** command, and you remove the REP configuration with the **no rep segment** command or overwrite it with the **rep segment <>** command, you cannot configure REP automatically again with the **rep segment auto** command. Instead, you must shut down the interface, bring it up, and then enter the **rep segment auto** command.
- REP Segment ID Autodiscovery depends on the CDP protocol. The feature does not support EtherChannel links.



CHAPTER 2

Default REP Configuration

- [Default REP Configuration, on page 13](#)
- [REP Configuration Guidelines, on page 13](#)
- [REP Administrative VLAN, on page 15](#)

Default REP Configuration

Default configuration for REP and REP features is as follows:

- REP is disabled on all interfaces. When enabled, the
- When REP is enabled:
 - The interface is a regular segment port unless it is configured as an edge port.
 - The sending of segment topology change notices (STCNs) is disabled.
 - All VLANs are blocked.

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.

- The administrative VLAN is VLAN 1.
- REP Segment-ID Autodiscovery is enabled.
- REP ZTP is enabled globally and disabled on interfaces.

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 changes to 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 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. It is not recommended to have a trunk port without explicitly specifying an allowed VLAN list for that trunk, to avoid flooding and impact to convergence.
- REP ports follow these rules:
 - There is no limit to the number of REP ports on a switch; however, only two ports on a switch can belong to the same REP segment.
 - If only one port on a switch is configured in a segment, the port should be an edge port.
 - If two ports on a switch belong to the same segment, they must be both edge ports, both regular segment ports, or one regular port and one edge no-neighbor port. An edge port and regular segment port on a switch cannot belong to the same segment.
 - If two ports on a switch belong to the same segment and one is configured as an edge port and one as a regular segment port (a misconfiguration), the edge port is treated as a regular segment port.
- REP interfaces come up in a blocked state and 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.
- You can configure how long a REP interface remains up without receiving a hello from a neighbor. You can use the **rep lsl-age-timer value** interface configuration command to set the time from 120 ms to 10000 ms. The LSL hello timer is then set to the age-timer value divided by 3. In normal operation, three LSL hellos are sent before the age timer on the peer switch expires and checks for hello messages.
 - In Cisco IOS Release 12.2(52)SE, the LSL age-timer range changed from 3000 to 10000 ms in 500-ms increments to 120 to 10000 ms in 40-ms increments. If the REP neighbor device is not running Cisco IOS release 12.2(52)SE or later, do not configure a timer value less than 3000 ms. Configuring a value less than 3000 ms causes the port to shut down because the neighbor switch does not respond within the requested time period.
 - EtherChannel port channel interfaces do not support LSL age-timer values less than 1000 ms. If you try to configure a value less than 1000 ms on a port channel, you receive an error message and the command is rejected.
- When configuring the REP LSL age timer, make sure that both ends of the link have the same time value configured. Configuring different values on ports at each end of the link results in a REP link flap.

- REP ports cannot be configured as one of these port types:
 - SPAN destination port
 - Tunnel port
 - Access port
- REP is supported on EtherChannels, but not on an individual port that belongs to an EtherChannel.
- In a stacked switch configuration, the maximum number of REP segments depends on the topology and the mix of features enabled on the switch and across the network. The recommended number of REP segments per switch for a stacked switch environment is 8, and the recommended maximum number of switches per segment is 24.
- REP Negotiated can be configured only on the first two uplink ports of the switch:
 - IE3000 and IE4000—GigabitEthernet 1/1 and GigabitEthernet 1/2
 - IE4010 and IE5000 (without 10G ports)—GigabitEthernet 1/25 and GigabitEthernet 1/26
 - IE5000 with 10G ports—TenGigabitEthernet 1/25 and TenGigabitEthernet 1/26

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.
- It is recommended to have a unique administrative VLAN per segment on a switch.
- The administrative VLAN cannot be the RSPAN VLAN.



CHAPTER 3

How to Configure REP

- [Configuring the REP Administrative VLAN, on page 17](#)
- [Configuring REP Interfaces, on page 17](#)
- [Configuring REP Negotiated, on page 19](#)
- [Setting Manual Preemption for VLAN Load Balancing, on page 21](#)
- [Configuring SNMP Traps for REP, on page 22](#)
- [Configuring REP ZTP, on page 22](#)
- [Configuring REP Segment-ID Autodiscovery, on page 23](#)

Configuring the REP Administrative VLAN

To configure the REP administrative VLAN, enter the following commands:

Step 1 Enter global configuration mode:

```
configure terminal
```

Step 2 Specify the administrative VLAN:

```
rep admin vlan vlan-id
```

The range is 2 to 4096. The default is VLAN 1. To set the admin VLAN to 1, enter the **no rep admin vlan** global configuration command.

Step 3 Return to privileged EXEC mode:

```
end
```

Configuring REP Interfaces

Before you begin

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.

Step 1 Enter global configuration mode:

```
configure terminal
```

Step 2 Specify the interface, and enter interface configuration mode:

```
interface interface-id
```

The interface can be a physical Layer 2 interface or a port channel (logical interface). The port-channel range is 1 to 10.

Step 3 Configure the interface as a Layer 2 trunk port:

```
switchport mode trunk
```

Step 4 Enable REP on the interface, and identify a segment number:

```
rep segment segment-id [edge [no-neighbor] [primary]] [preferred]
```

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.

- **edge** —Configures 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) **primary** — Configures the port as the primary edge port, the port on which you can configure VLAN load balancing.
- (Optional) **no-neighbor**— Configures a port with no external REP neighbors as an edge port. The port inherits all properties of edge ports, and you can configure them the same as any edge port.

Note Although each segment can have only one primary edge port, if you configure edge ports on two different switches and enter the **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.

- (Optional) **preferred** —Indicates 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 (Optional) Configure the edge port to send segment topology change notices (STCNs):

```
rep stcn {interface interface-id | segment id-list | stp }
```

- **interface interface-id** —Designates a physical interface or port channel to receive STCNs.
- **segment id-list**— Identifies one or more segments to receive STCNs. The range is 1 to 1024.
- **stp**— Sends STCNs to STP networks.

Step 6 (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.

```
rep block port {id port-id | neighbor_offset |preferred} vlan {vlan-list | all}
```

- **id port-id**—Identifies the alternate port by port ID. The port ID is automatically generated for each port in the segment. You can view interface port IDs by entering the **show interface interface-id rep [detail]** privileged EXEC command.
- **neighbor_offset number**—Identifies 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 [VLAN Load Balancing, on page 4](#) for an example of neighbor offset numbering.

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.

- **preferred**—Selects the regular segment port previously identified as the preferred alternate port for VLAN load balancing.
- **vlan vlan-list**—Blocks one VLAN or a range of VLANs.
- **vlan all**—Blocks all VLANs.

Note Enter this command only on the REP primary edge port.

Step 7 (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.

```
rep preempt delay seconds
```

The time delay range is 15 to 300 seconds. The default is manual preemption with no time delay.

Note Enter this command only on the REP primary edge port.

Step 8 (Optional) Configure a time (in milliseconds) for which the REP interface remains up without receiving a hello from a neighbor.

```
rep lsl-age-timer value
```

The range is from 120 to 10000 ms in 40-ms increments. The default is 5000 ms (5 seconds).

Note If the neighbor device is not running Cisco IOS Release 12.2(52)SE or later, it only accepts values from 3000 to 10000 ms in 500-ms intervals. EtherChannel port channel interfaces do not support LSL age-timer values less than 1000 ms.

Step 9 Return to privileged EXEC mode:

```
end
```

Configuring REP Negotiated

Use the following procedure to configure REP Negotiated in a REP network where a new switch is being inserted into the existing REP ring topology. The adjacent switches to this newly inserted switch are referred to as peer switches.

Procedure

	Command or Action	Purpose
Step 1	On the new switch, configure rep bpduleak in global configuration mode:	<code>rep bpduleak</code> !
Step 2	<p>Configure the EEM macros on the new switch as shown in the following example. This example assumes that the peer switches are configured with REP Segment 777 and the newly inserted switch has uplink ports GigabitEthernet 1/1 and GigabitEthernet 1/2.</p> <p>Example:</p> <pre>macro auto execute CISCO_REP_NEG_EVENT { config terminal no rep bpduleak interface GigabitEthernet 1/1 switchport mode trunk no rep negotiated rep segment 777 interface GigabitEthernet 1/2 switchport mode trunk no rep negotiated rep segment 777 exit } macro auto execute CISCO_REP_NONNEG_EVENT { config terminal no rep bpduleak interface GigabitEthernet 1/1 no rep negotiated interface GigabitEthernet 1/2 no rep negotiated } !</pre>	
Step 3	Insert the new switch into the existing REP Ring topology.	Note The newly inserted switch still does not have any REP Segment configurations.
Step 4	Check the output of show rep topology on the peer switches.	The output should show that rep bpduleak is in effect. The REP segment remains intact, but the newly inserted switch is not reflected in the topology. This indicates that the newly inserted switch is transparently forwarding the REP traffic between its uplink ports.
Step 5	<p>Configure rep negotiated on both the uplink interfaces of the newly inserted switch.</p> <p>Example:</p> <pre>interface range GigabitEthernet 1/1-2 rep negotiated !</pre>	
Step 6	<p>Use the show rep negotiated command on the newly inserted switch to verify the status.</p> <p>Example:</p>	

	Command or Action	Purpose
	<pre>Switch2 #show rep negotiated REP negotiation status : Fail Interface1: GigabitEthernet1/1 Status : enabled Rx State : fail, Segment-ID: 0 Interface2: GigabitEthernet1/2 Status : enabled Rx State : fail, Segment-ID: 0</pre>	
Step 7	<p>Configure rep negotiated on the connected uplink interfaces of both the peer switches and wait for the REP Negotiation to complete. The following console log message on the newly inserted switch indicates that an EEM event has been triggered by REP Negotiation.</p> <p>Example:</p> <pre>May 22 22:54:41.087: REP negotiated event generated, executed CISCO_REP_NEG for Segment 777</pre>	
Step 8	<p>Use the show rep negotiated command on the newly inserted switch to verify the status.</p> <p>Example:</p> <pre>Switch2#show rep negotiated REP negotiation status : Success Interface1: GigabitEthernet1/1 Status : disabled Interface2: GigabitEthernet1/2 Status : disabled</pre>	<p>REP is configured automatically in the newly inserted switch and it also appears in the show rep topology output on all the switches in the REP Ring.</p>

Setting Manual Preemption for VLAN Load Balancing

Before you begin

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.

Procedure

	Command or Action	Purpose
Step 1	Manually trigger VLAN load balancing on the segment:	<pre>rep preempt segment segment-id</pre> <p>You will need to confirm the command before it is executed.</p>
Step 2	Display REP topology information:	<pre>show rep topology</pre>

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.

Procedure

	Command or Action	Purpose
Step 1	Enter global configuration mode:	<code>configure terminal</code>
Step 2	Enable the switch to send REP traps, and set the number of traps sent per second:	<code>snmp mib rep trap-rate value</code> The range is from 0 to 1000. The default is 0 (no limit imposed; a trap is sent at every occurrence).
Step 3	Return to privileged EXEC mode:	<code>end</code>

Configuring REP ZTP

To configure REP ZTP, you enable or disable it at the global level and the interface level. The default states are:

- Global level: Enabled
- Interface level: Disabled

You must explicitly enable the feature at the interface level on the upstream device interface connected to the downstream device. When enabled, only that interface will receive notification from the downstream switch to block or unblock the PnP startup VLAN.



Note When applying configuration from DNAC or PNP server user must explicitly add this CLI configuration in the configuration template for the feature to be enabled.

Step 1 Enter global configuration mode:

```
Switch# configure terminal
```

Step 2 Globally enable REP ZTP:

```
Switch(config)# rep ztp
```

Use the no form of the command to disable REP ZTP: `Switch(config)# no rep ztp`

Step 3 Enter interface configuration mode on the upstream device interface that is connected to the downstream device:

```
Switch(config)# interface <interface-name>
```

Step 4 Enable REP ZTP on the interface:

```
Switch(config-if)#rep ztp-enable
```

Use the no form of the command to disable REP ZTP on the interface: `Switch(config-if)#no rep ztp-enable`

Example

The following example shows the minimum configuration required to enable the REP ZTP feature on the upstream device interface that is connected to a downstream device.

```
Switch#show running-config interface gigabitEthernet 1/2
Building configuration...

Current configuration : 93 bytes
!
interface GigabitEthernet1/2
  switchport mode trunk
  rep segment 100
  rep ztp-enable
end
```

Configuring REP Segment-ID Autodiscovery

You use CLI commands for REP Segment-ID Autodiscovery. One enables or disables autodiscovery on a REP switch, and one configures new interfaces so the switch learns the segment-ID. You also use CLI commands to view the status of the feature on the segment.

Enable REP Segment-ID Autodiscovery

REP Segment-ID Autodiscovery is enabled by default. However, you can re-enable it on the switch upstream and downstream interfaces.

Enable REP Segment-ID Autodiscovery on the switch.

Example:

```
switch(config)#rep autodisc
```

You disable REP Segment-ID Autodiscovery by entering the following command:

```
switch(config)#no rep autodisc
```

What to do next

You can check the status of REP Segment-ID Autodiscovery. See the section [View REP Segment ID Autodiscovery Status, on page 29](#) in this guide.

Configure the Interfaces

Configure the interface on the newly inserted switch to enable REP Segment-ID learning. This configuration command is available on all uplink and downlink ports.

Before you begin

Ensure that the REP segment ID is configured on the primary and secondary edge devices. You configure the segment ID by entering the command **rep segment *segment_id* edge**, in which *segment_id* is the segment ID of the ring to be propagated through CDP packet to the neighboring device when connected.

Enable the switch to learn the segment ID.

Example:

```
switch(config)#int gig1/1
switch(config-if)#rep seg auto
```

This command is not reflected in the "running-config" of the switch until the REP Segment-ID learning is completed successfully. The learned REP Segment-ID is cached internally. This cached information is used whenever available instead of learning the REP Segment-ID from the Peer switch. Any previously configured REP Segment-ID configuration automatically gets cached when `rep segment auto` is configured on an interface.

The following example shows the minimum configuration to enable the feature on an interface on the upstream device switch. The upstream device with an explicit REP segment is typically an edge switch.

```
switch#show running-config interface gigabitEthernet 1/3
Building configuration...
```

```
Current configuration : 93 bytes
!
interface GigabitEthernet1/3
 switchport mode trunk
 rep segment auto 1
```

The following example shows the minimum configuration to enable the feature on an interface on the downstream switch interface. Enter the command **show running-config interface *interface_id*** to confirm that the downstream switch knows to expect to receive its REP segment through CDP message.

```
switch#show running-config interface gigabitEthernet 1/2
Building configuration...
```

```
Current configuration : 93 bytes
!
interface GigabitEthernet1/2
 switchport mode trunk
 rep segment auto
end
```

You disable the ability of the switch to learn the segment ID by entering the following command:

```
switch(config-if)#no rep segment
```

What to do next

You can check the status of REP Segment-ID Autodiscovery. See the section [View REP Segment ID Autodiscovery Status, on page 29](#) in this guide.



CHAPTER 4

Monitoring and Maintaining REP

- [Monitoring and Maintaining REP, on page 25](#)
- [Displaying REP ZTP Status, on page 25](#)
- [View REP Segment ID Autodiscovery Status, on page 29](#)

Monitoring and Maintaining REP

Command	Purpose
<code>show interface [interface-id] rep [detail]</code>	Displays REP configuration and status for an 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.
<code>copy running-config startup config</code>	Saves your entries in the switch startup configuration file.

Displaying REP ZTP Status

Use the **show** command to identify the state of REP ZTP on an interface. In the following example, the feature is disabled on interface GigabitEthernet 1/1 and it is enabled on interface GigabitEthernet 1/2. The status of **pn_p_startup_vlan** is "Blocked".

Step 1 In privileged exec mode, enter:

```
show interfaces rep detail
```

Example:

```
Switch#show interfaces rep detail
GigabitEthernet1/1 REP enabled
Segment-id: 100 (Segment)
PortID: 00016C13D5AC4320
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 00026C13D5AC43209DAB
```

```

Port Role: Open
Blocked VLAN: <empty>
Admin-vlan: 1
REP-ZTP Status: Disabled
REP Segment Id Auto Discovery Status: Enabled
REP Segment Id Type: Manual
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
LSL Ageout Retries: 5
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 382, tx: 297
HFL PDU rx: 0, tx: 0
BPA TLV rx: 1, tx: 19
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 95, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 95, tx: 95

GigabitEthernet1/2 REP enabled
Segment-id: 100 (Segment)
PortID: 00026C13D5AC4320
Preferred flag: No
Operational Link Status: NO_NEIGHBOR
Current Key: 00026C13D5AC43209DAB
Port Role: Fail No Ext Neighbor
Blocked VLAN: 1-4094
Admin-vlan: 1
REP-ZTP Status: Enabled
REP-ZTP PnP Status: Unknown
REP-ZTP PnP Vlan: 1
REP-ZTP Port Status: Blocked
REP Segment Id Auto Discovery Status: Enabled
REP Segment Id Type: Manual
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
LSL Ageout Retries: 5
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 11, tx: 11
HFL PDU rx: 0, tx: 0
BPA TLV rx: 0, tx: 0
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 0, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 0, tx: 0

```

Step 2 Use the show command again to display the status of **pnp_startup_vlan**.

When the downstream device is booted up, it sends notification to the connected upstream switch interface to unblock the **pnp_startup_vlan** for it to get the DHCP IP address and further establish communication with the PNP server or DNAC. The show command indicates the status as "Unblocked".

The following syslogs on the upstream switch notify you about FWD and BLK of ports. There are no syslogs in the downstream switch as PnP takes control of the console and no syslogs can be printed on the console.

```
REP-6-ZTPPORTFWD: Interface GigabitEthernet1/2 moved to forwarding on ZTP notification
```

```
REP-6-ZTPPORTBLK: Interface GigabitEthernet1/2 moved to blocking on ZTP notification
```

Example:

```
Switch#show interfaces rep detail
GigabitEthernet1/1 REP enabled
Segment-id: 100 (Segment)
PortID: 00016C13D5AC4320
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 00026C13D5AC43209DAB
Port Role: Open
Blocked VLAN: <empty>
Admin-vlan: 1
REP-ZTP Status: Disabled
REP Segment Id Auto Discovery Status: Enabled
REP Segment Id Type: Manual
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
LSL Ageout Retries: 5
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 430, tx: 358
HFL PDU rx: 0, tx: 0
BPA TLV rx: 1, tx: 67
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 107, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 107, tx: 108

GigabitEthernet1/2 REP enabled
Segment-id: 100 (Segment)
PortID: 00026C13D5AC4320
Preferred flag: No
Operational Link Status: NO_NEIGHBOR
Current Key: 00026C13D5AC43209DAB
Port Role: Fail No Ext Neighbor
Blocked VLAN: 1-4094
Admin-vlan: 1
REP-ZTP Status: Enabled
REP-ZTP PnP Status: In-Progress
REP-ZTP PnP Vlan: 69
REP-ZTP Port Status: Unblocked
REP Segment Id Auto Discovery Status: Enabled
REP Segment Id Type: Manual
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
LSL Ageout Retries: 5
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 32, tx: 40
HFL PDU rx: 0, tx: 0
BPA TLV rx: 0, tx: 0
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 0, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 0, tx: 0

Switch#show interfaces rep detail
GigabitEthernet1/1 REP enabled
```

Displaying REP ZTP Status

```

Segment-id: 100 (Segment)
PortID: 00016C13D5AC4320
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 00026C13D5AC43209DAB
Port Role: Open
Blocked VLAN: <empty>
Admin-vlan: 1
REP-ZTP Status: Disabled
REP Segment Id Auto Discovery Status: Enabled
REP Segment Id Type: Manual
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
LSL Ageout Retries: 5
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 430, tx: 358
HFL PDU rx: 0, tx: 0
BPA TLV rx: 1, tx: 67
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 107, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 107, tx: 108

```

```

GigabitEthernet1/2 REP enabled
Segment-id: 100 (Segment)
PortID: 00026C13D5AC4320
Preferred flag: No
Operational Link Status: NO_NEIGHBOR
Current Key: 00026C13D5AC43209DAB
Port Role: Fail No Ext Neighbor
Blocked VLAN: 1-4094
Admin-vlan: 1
REP-ZTP Status: Enabled
REP-ZTP PnP Status: Completed
REP-ZTP PnP Vlan: 69
REP-ZTP Port Status: Blocked
REP Segment Id Auto Discovery Status: Enabled
REP Segment Id Type: Manual
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
LSL Ageout Retries: 5
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 32, tx: 40
HFL PDU rx: 0, tx: 0
BPA TLV rx: 0, tx: 0
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 0, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 0, tx: 0

```

Step 3 (Optional) Use the following debug commands to troubleshoot REP ZTP:

- **debug rep lsism**: This command helps you understand LSL state machine events in the NO_NEIGHBOR state.

- **debug rep packet:** Use this command to dump LSL packets with the REP ZTP LSL TLV to check the PnP status on the peer client node.

View REP Segment ID Autodiscovery Status

You can use the **show interfaces rep detail** CLI command to check the status of REP Segment-ID Autodiscovery on the segment.

Enter **show interfaces rep detail** to confirm that REP Segment-ID Autodiscovery is globally enabled on the switch and to see whether the segment ID on the interface is configured automatically or manually.

Example:

```
Switch3#show interface rep detail
GigabitEthernet1/10  REP enabled
Segment-id: 100 (Edge)
PortID: 000A54A274103B00
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 00062C3311D266001510
Port Role: Open
Blocked VLAN: <empty>
Admin-vlan: 2
REP Segment Id Type: Auto
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 953387, tx: 1056967
HFL PDU rx: 97, tx: 63
BPA TLV rx: 344515, tx: 446
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 0, tx: 0
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 260153, tx: 261511

Switch3#
```




CHAPTER 5

Configuration Examples for Configuring REP

- [Configuring the Administrative VLAN: Example, on page 31](#)
- [Configuring a Primary Edge Port: Examples, on page 31](#)
- [Configuring VLAN Blocking: Example, on page 32](#)

Configuring the Administrative VLAN: Example

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 GigabitEthernet1/17 rep detail
GigabitEthernet1/17 REP enabled
Segment-id: 2 (Edge)
PortID: 00010019E7144680
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 0002001121A2D5800E4D
Port Role: Open
Blocked Vlan: <empty>
Admin-vlan: 100
Preempt Delay Timer: disabled
LSL Ageout Timer: 5000 ms
Configured Load-balancing Block Port: none
Configured Load-balancing Block VLAN: none
STCN Propagate to: none
LSL PDU rx: 3322, tx: 1722
HFL PDU rx: 32, tx: 5
BPA TLV rx: 16849, tx: 508
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 118, tx: 118
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 4214, tx: 4190
```

Configuring a Primary Edge Port: Examples

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. The interface is configured to remain up for 6000 milliseconds without receiving a hello from a neighbor.

```
Switch# configure terminal
Switch (conf)# interface GigabitEthernet1/17
Switch (conf-if)# rep segment 1 edge primary
Switch (conf-if)# rep stcn segment 2-5
Switch (conf-if)# rep block port 0009001818D68700 vlan all
Switch (conf-if)# rep preempt delay 60
Switch (conf-if)# rep lsl-age-timer 6000
Switch (conf-if)# end
```

This example shows how to configure an interface as the primary edge port when the interface has no external REP neighbor:


```
Switch# configure terminal
Switch (conf)# interface GigabitEthernet1/17
Switch (conf-if)# rep segment 1 edge no-neighbor primary
Switch (conf-if)# rep stcn segment 2-5
Switch (conf-if)# rep block port 0009001818D68700 vlan all
Switch (conf-if)# rep preempt delay 60
Switch (conf-if)# rep lsl-age-timer 6000
```

Configuring VLAN Blocking: Example

This example shows how to configure the VLAN blocking configuration shown in the diagram below. 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 1/0/1).

```
Switch# configure terminal
Switch (conf)# interface GigabitEthernet1/17
Switch (conf-if)# rep segment 1 edge primary
Switch (conf-if)# rep block port 4 vlan 100-200
Switch (conf-if)# end
```

Figure 8: Example of VLAN Blocking





CHAPTER 6

Feature History

- [Feature History, on page 33](#)

Feature History

Table 1: Feature History

Feature	Release	Feature Information	Platform
REP Segment ID Autodiscovery - Interoperability with Cisco IOS XE	15.2(8)E4	This feature allows the REP Segment ID Autodiscovery feature on Cisco Classic IOS switches to interoperate with Cisco IOS XE based platforms.	IE 4000, IE 4010, and IE 5000
REP Zero Touch Provisioning	15.2(8)E4	The REP ZTP feature allows PnP to function on insertion of a new IE switch into an existing REP ring.	IE 4000, IE 4010, and IE 5000

Feature	Release	Feature Information	Platform
REP Negotiated and REP Segment ID Validation	15.2(8)E2	<p>When the switch interfaces are configured with REP Negotiated (see REP Negotiated), REP status is negotiated with the peers. If the peer supports REP, it is migrated to REP. If the peer does not support REP, it is migrated to STP. The peer is migrated to REP or STP using an Embedded Event Manager (EEM) macro.</p> <p>The REP node learns the Segment ID and maintains it on interfaces until the link goes down. The ring must be initially configured with a static REP Segment ID from the edge, but the rest of the REP ring can implement REP Segment ID autodiscovery to facilitate deployment and automation.</p>	IE 4000, IE 4010, and IE 5000