



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 BPDUs 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](#).



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.