High-Availability Seamless Redundancy in the Factory Network

Design Guide

First Published: July 2018

Cisco Systems, Inc. www.cisco.com

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies are considered un-Controlled copies and the original on-line version should be referred to for latest version.

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

Contents

High-Availability Seamless Redundancy in the Factory Network

Introduction

In today's Internet of Things (IoT) world, manufacturing industries are adopting the vision of connecting devices and implementing a converged plant floor to reduce operational costs and increase productivity. Highly-available equipment requires a resilient network that minimizes downtime for the mission critical production cycle.

The High-availability Seamless Redundancy (HSR) protocol that is defined in IEC 62439-3 is designed to achieve zero recovery time in ring topologies, making it suitable for automation networks where the downtime of Industrial Automation and Control System (IACS) applications should be kept to a minimum.

HSR supports single fault in the ring with no downtime by sending each packet in both directions in the ring. The receiving node accepts the first packet and discards the second. In contrast, other ring redundancy protocols block one path and use only the primary one, dropping some packets during convergence.

The key advantages of HSR are:

- Zero convergence time when a node, link, or switch fails
- No single point of failure
- Easy to configure

This *High-Availability Seamless Redundancy in the Factory Network Cisco Validated Design* (CVD) provides a tested and validated architecture to deploy HSR in ring topologies in the factory Cell/Area zone, which is the area where the IACS and end devices (Levels 0-2) connect to the network. This document guides the reader through the high-level technology and architecture.

Organization

This guide includes the following sections:

Cisco Connected Factory and Converged Network Architecture Overview

The constantly evolving manufacturing industry has massive challenges because of the constant demand to increase productivity and decrease operational costs year after year. The industry is also facing challenges such as increased regulation and compliance measures becoming mandatory on the plant floor. This means the industry needs to have complete and real-time visibility to the very end of plant floor devices (sensors and actuators) and network components.

Realizing or achieving these goals requires not only automation in the industries' plant floor, but also a strong spinal network that runs through the different zones of the manufacturing industry, connecting all the devices and machines and data to provide real-time visibility of each component.

Manufacturing industries are in a much better place to adopt Industrial Ethernet, since the following capabilities and some of the problems it solves are known:

- Increased speed and distance (using Fiber Media)
- Better interoperability
- Master-slave architecture instead of peer-to-peer
- Resilient network for equipment availability, minimizing downtime for the mission critical production cycle

The Connected Factory architecture is designed from the ground up to leverage Cisco's industry-leading networking expertise and extend it to the factory floor where industrial machines, controls, and other devices require maximum resiliency, extremely low latency, compliance with industry standards, and ease of use for operations personnel. The architecture connects the plant floor network with the existing enterprise network, while the industrial DMZ protects mission critical data and machinery. This solution is designed to be easily scalable and repeatable so it can be easily deployed to a wide variety of manufacturing applications.

Using a converged network architecture based on the industry-standard Purdue Model for Control Hierarchy, Cisco addresses connectivity, security, and resiliency based on a layered approach. Cisco's Industrial Ethernet (IE) switches provide plant floor level access to IACS, while Wireless Access Points (WAPs) and Wireless LAN Controllers (WLCs) provide mobile connectivity for other devices where a hardwired connection is not feasible or required. Multiple levels of redundancy are implemented using multiple physical links and aggregation switches, as well as industrial protocols like HSR. Security is also critical for industrial networks because a malicious or accidental change to an industrial device can cause considerable damage or downtime.

Figure 1 Connected Factory Architecture

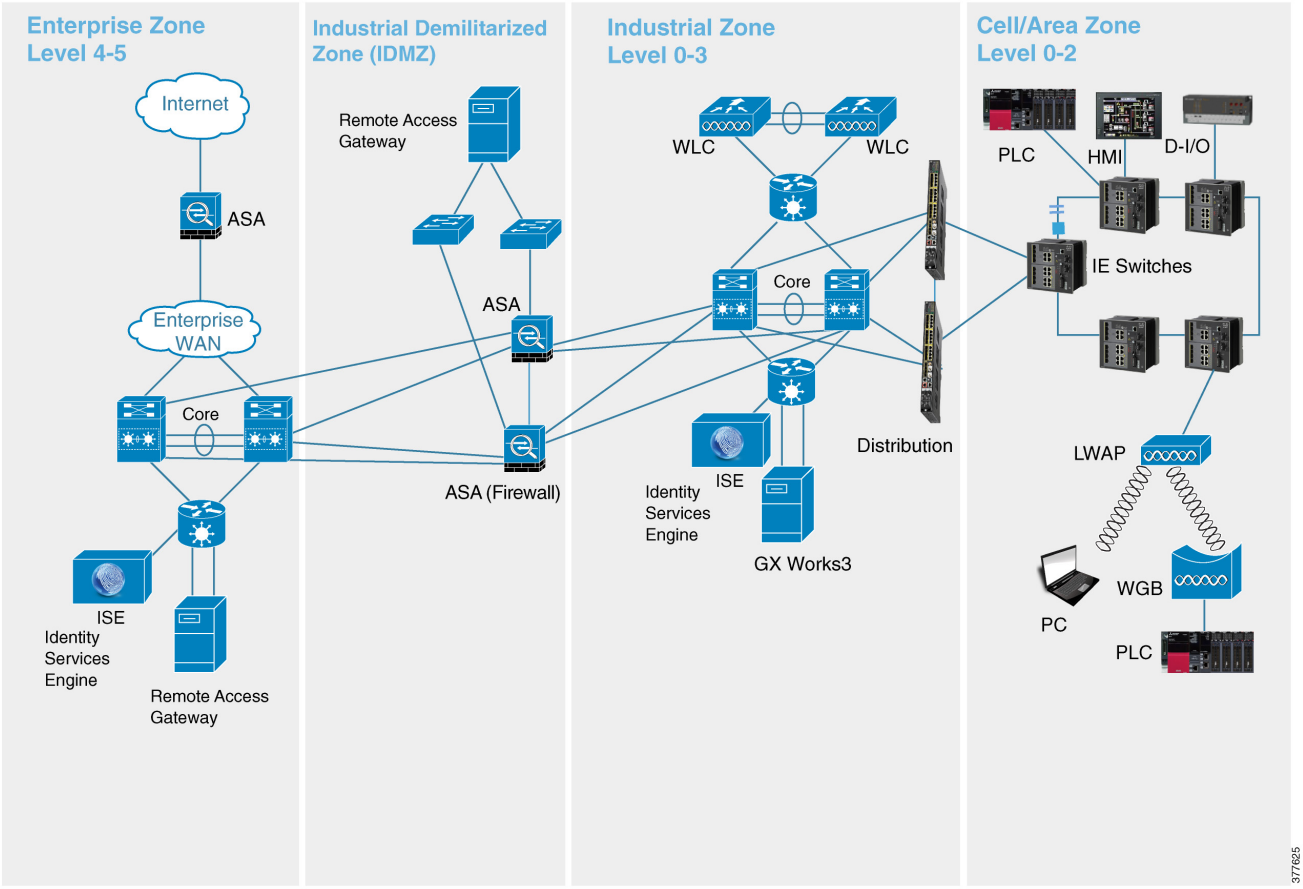

Purdue Model for Control Hierarchy

The Cisco Connected Factory solution employs the commonly used industry-standard Purdue Model for Control Hierarchy to divide the plant into a logical framework, as shown in Figure 2.

Figure 2 Purdue Model for Control Hierarchy

Enterprise Zone	Enterprise Network	Level 5
	Site Business Planning and Logistics Network	Level 4
IDMZ	Industrial Demilitarized Zone — Shared Access	
Industrial Zone	Site Manufacturing Operations and Control	Level 3
Cell/Area Zone	Area Supervisory Control	Level 2
	Basic Control	Level 1
	Process	Level 0

377621

Starting at the bottom of the model, the Cell/Area Zone contains three levels of equipment:

- **Level 0—Process**—Industrial sensors, drives, actuators, and similar devices that interact with the physical environment by taking measurements or performing actions like starting a motor or moving a robot arm.
- **Level 1—Basic Control**—Controllers, such as programmable logic controllers, distributed control system, and generically programmable automation controller, that communicate directly with the Level 0 devices, other controllers, and higher level control systems.
- **Level 2—Area Supervisory Control**—Operator interfaces including Human Machine Interface (HMI), alarm systems, and control room workstations.

The Industrial Zone contains (Level 0-3) systems that maintain site level control of the lower level IACS systems including reporting, scheduling, file and patch servers, and network services such as NTP, DNS, DHCP, and AD. One or more of the Cell/Area Zones (described above) actually reside within the Industrial Zone, as depicted in Figure 2.

Sitting between the Industrial and Enterprise Zones is the Industrial Demilitarized Zone (IDMZ), which provides a layer of separation between the traditional Information Technology (IT) and Operational Technology (OT)-operated areas of the network, allowing only the absolutely required traffic to traverse the zone.

The Enterprise Zone, containing Level 4 and Level 5, provides access to the Internet and higher-order network applications including email, database, Business-to-Business (B2B) and Business-to-Consumer (B2C) applications, and other non-critical resources. This area is often seen as a source of security threats to the lower level resources and is typically managed by the IT department.

As shown in Figure 1, the Cell/Area Zone is at the access layer of the network. In an industrial environment due to the physical layout of the plant floor, a ring topology is typically used to connect a series of access switches together. A ring topology allows every point in the ring to have multiple potential paths out of the ring, which is critical if a cable or device fails. If a failure in the ring occurs, the network must be able to keep forwarding traffic. In this deployment, Cisco IE switches use the HSR protocol to provide redundancy.

Cisco Industrial Network Portfolio

Cisco Industrial Ethernet Switches, which are designed from the ground up for the demanding environments and requirements of industrial networks, are fully tested and validated in IACS environments for the top protocols in use today, including Ethernet/IP, PROFINET, and CC-Link. Building on the years of class-leading experience in enterprise switching, these industrial-focused switches include additional software features and ruggedized hardware to deliver resilient, low-latency, and high speed performance. Cisco Industrial Ethernet Switches are available in a range of models, each with a variety of port configurations and sizes, including Cisco IE 1000, Cisco IE 2000, Cisco IE 3000, Cisco IE 4000, Cisco IE 4010, and Cisco IE 5000. Highlighted features of the switches include the following (not all features are available on all models):

- Industry-leading security including SUDI, ACT2, FIPS 140-2, TrustSec, and SGT/SGACL
- Port speeds ranging from 100 Mbps to 10 Gbps (copper and fiber available)
- DIN rail mount
- DC power supplies
- Alarm relays
- Industrial PoE and PoE+
- NetFlow Lite
- Support for industrial protocols including CIP, EtherNet I/P, and PROFINET
- Redundancy features including REP, PRP, MRP, EtherChannel, and Flex Links
- Horizontal stacking

The current implementation contains the following specific products:

- **Cisco IE 4000 Series Switches**—The industry's first DIN rail-mounted 40 Gigabit Ethernet switch platform that offers high bandwidth and low latency. The Cisco IE 4000 is available in various models with up to 20 Gigabit Ethernet interfaces.
- **Cisco IE 5000 Series Switches**—A 19-inch, one-rack unit multi-10 Gbps aggregation switch equipped with 24 Gigabit Ethernet ports plus four 10-Gigabit or four 1-Gigabit ports, making it ideal for the aggregation layer or backbone in large-scale industrial networks.

Design Considerations

Resiliency Protocol Considerations

Resiliency protocols provide redundant paths in the network while avoiding loops to maximize uptime for devices running on the plant floor, such as Programmable Logic Controllers (PLCs), Remote I/O, Human Machine Interfaces (HMI), PC-Supervisor, sensors, actuators, and drives.

Resiliency protocols can be classified using the following criteria:

- **Network Layer**—Table 1 and Table 2 show a comparison of redundancy protocols by network layer (Layer 2 or Layer 3 redundancy).
- **Topology**—Some redundancy protocols are suited for ring or star topologies.
- **Industry Standard**—These resiliency protocols are based on industry standards and are ideal for multi-vendor deployments.
- **Convergence Time**—A protocol that meets the maximum convergence time for the application should be selected.

Table 1 Layer 2 Redundancy Protocols Comparison

Resiliency Protocol	Industry Standard	Topology	Convergence Time
STP (802.1D)	Yes	Ring/Star	Above 250ms
RSTP (802.1w)	Yes	Ring/Star	Above 250ms
MSTP (802.1s)	Yes	Ring/Star	Above 250ms
RPVST+	Cisco proprietary	Ring/Star	Above 250ms
REP	Cisco proprietary	Ring	50-150 ms
HSR	X	Ring	Zero
PRP	Yes	Any	Zero
EtherChannel (LACP 802.3ad)	Yes	Star	150-250ms
Flex Links	Cisco proprietary	Star	50-150 ms
StackWise	Cisco proprietary	Ring/Star	150-250ms

Table 2 Layer 3 Redundancy Protocols Comparison

Resiliency Protocol	Industry Standard	Topology	Convergence Time
StackWise	Cisco proprietary	Ring/Star	150-250ms
HSRP	Cisco proprietary	Ring/Star	Above 250ms
GLBP	Cisco proprietary	Ring/Star	Above 250ms
VRRP (IETF RFC 3768)	Yes	Ring/Star	Above 250ms

As Table 1 and Table 2 show, HSR is suitable for applications in ring topologies that require fast network convergence such as process automation. HSR is supported on Cisco IE 4000, Cisco IE 4010, and Cisco IE 5000.

Use Case Overview

HSR is used in HSR-SAN mode to provide redundancy in the access layer where IACS devices are connected. The access layer is aggregated by distribution switches that also provide a Layer 2/Layer 3 boundary and default gateway for the Layer 2 domain. At the distribution layer, Hot Standby Router Protocol (HSRP) is used to provides stateless redundancy for IP routing. The following are two possible HSR implementations in the Cell Area network.

Single HSR Ring Connected to Distribution Layer with Resilient Ethernet Protocol Segment

Figure 3 Single HSR Ring Connected to Distribution Layer with REP Segment

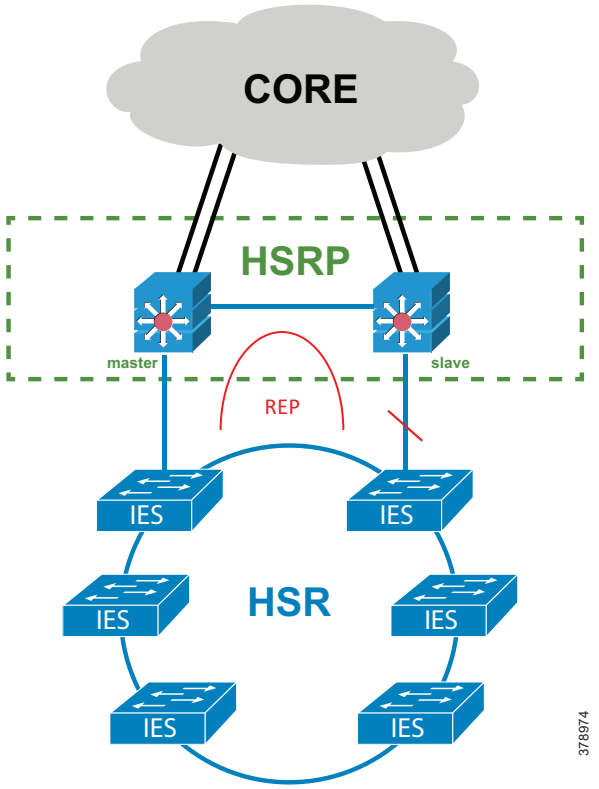

This topology combines hitless HSR topology in the Cell Area network with Resilient Ethernet Protocol (REP) in the uplinks to the aggregation to minimize the probability of HSRP and routing protocol convergence in the event of a failure in the uplinks. The only single event that would trigger an HSRP failover is the failure of the HSRP master node.

Single HSR Ring Connected to Isolated Distribution Switches

Figure 4 Single HSR Ring Connected to Isolated Distribution Switches

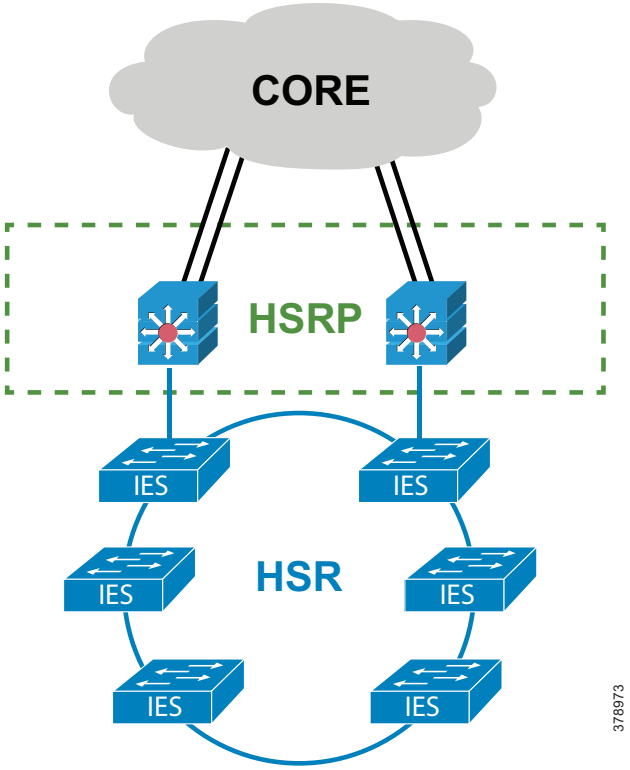

This topology has two distribution switches not directly connected to each other to achieve a loop free topology. Any failure in the uplink to the HSRP master results in an HSRP failover.

Choosing an HSR Implementation

Connecting the HSR to the distribution layer using REP, as shown in Single HSR Ring Connected to Distribution Layer with Resilient Ethernet Protocol Segment, page 6, is the recommended implementation since it provides the highest availability. Single HSR Ring Connected to Isolated Distribution Switches, page 7 is given as an alternative for environments with mixed platforms that don't support REP.

Technology Overview

High-Availability Seamless Redundancy

High-availability Seamless Redundancy (HSR) is defined in International Standard IEC 62439-3-2016 clause 5. HSR is designed to work in a ring topology. The same standard defines PRP used in star topologies. HSR defines a ring with traffic in opposite directions. Port-A sends traffic counter clockwise in the ring and Port-B sends traffic clockwise in the ring.

To allow the switch to determine and discard duplicate packets, additional protocol specific information is sent with the data frame as part of what is called the HSR header. The HSR header contain a sequence number, which is the primary data used to determine if the received frame is the first instance or a duplicate instance.

The nodes connecting to the HSR ring are referred to as Doubly Attached Nodes implementing HSR (DANHs). Similar to PRP, Singly Attached Nodes (SANs) are attached to the HSR ring through a device called a Redundancy Box (RedBox). The RedBox acts as a DANH for all traffic for which it is the source or the destination. The Cisco IE 4000 switch implements RedBox functionality using Gigabit Ethernet port connections to the HSR ring.

Figure 5 shows an example of an HSR ring as described in IEC 62439-3.

Figure 5 Example of HSR Ring Carrying Unicast Traffic

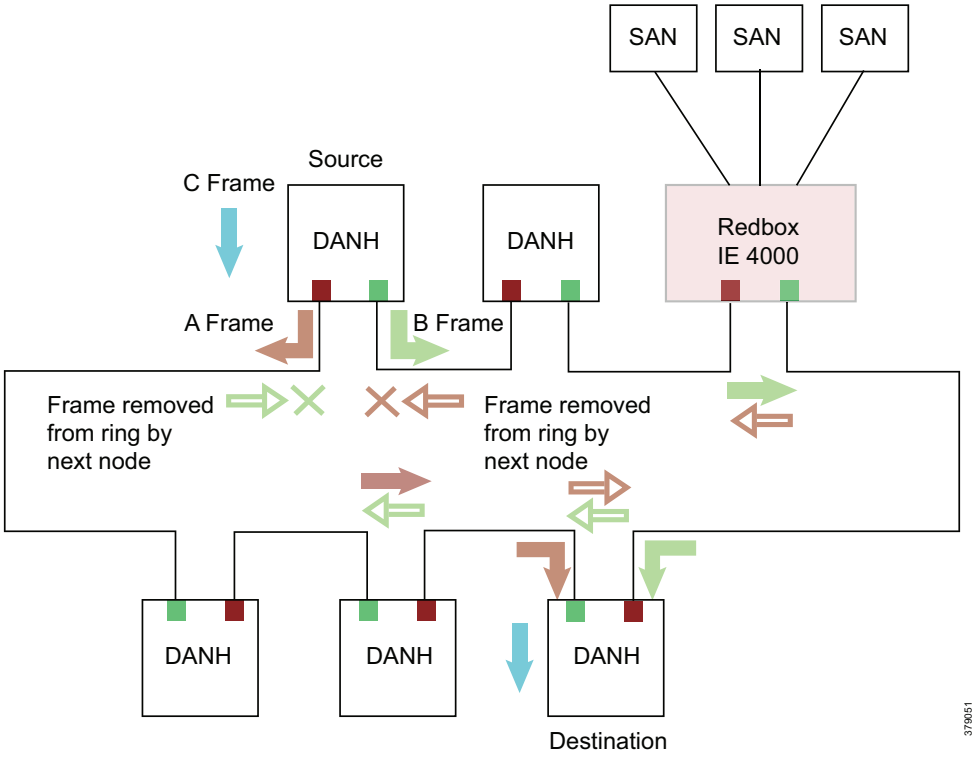

Devices that do not support HSR out of the box (for example, laptops and printers) cannot be attached to the HSR ring directly because an HSR device needs two ports. These nodes are attached to the HSR ring through a RedBox. As shown in Figure 3, the RedBox has two ports on the DANH side. Non-HSR SAN devices are attached to the upstream switch ports. The RedBox generates the supervision frames on behalf of these devices so that they are seen as DANH devices on the ring. Because the RedBox emulates these as DANH, they are called Virtual Doubly Attached Nodes (VDAN).

Loop Avoidance

To avoid loops and use network bandwidth effectively, the RedBox does not transmit frames that are already transmitted in the same direction. When a node injects a packet into the ring, the packet is handled as follows to avoid loops:

- **Unicast packet with destination inside the ring**—When the unicast packet reaches the destination node, the packet is consumed by the respective node and is not forwarded.
- **Unicast packet with destination not inside the ring**—Because this packet does not have a destination node in the ring, it is forwarded by every node in the ring until it reaches the originating node. Because every node has a record of the packet it sent, along with the direction in which it was sent, the originating node detects that packet has completed the loop and drops the packet.
- **Multicast packet**—A multicast packet is forwarded by each node because there can be more than one consumer of this packet. For this reason, a multicast packet always reaches the originating node. However, every node will check whether it has already forwarded the received packet through its outgoing interface. Once the packet reaches the originating node, the originating node determines that it already forwarded this packet and drops the packet instead of forwarding it again.

HSR RedBox Modes of Operation

An HSR RedBox can operate in one of the following modes that define how HSR handles packets in different scenarios:

- **HSR-SAN**—This is the most basic mode, in which the RedBox connects SAN devices to an HSR Ring. No other PRP or HSR network is involved in this configuration. In this mode, the traffic on the upstream switch port (the interlink port) does not have HSR/PRP tags and the RedBox represents the SAN device as a VDAN in the ring.
- **HSR-PRP**—This configuration is used to bridge HSR and PRP networks. Traffic on the interlink port (the RedBox interface that connects to the PRP network) is PRP tagged in this configuration. The RedBox extracts the data from the PRP frame and generates the HSR frame using this data and it performs the reverse operation for packets in the opposite direction. HSR-PRP mode is not used in this design. For more information, refer to High-Availability Seamless Redundancy (HSR) for Cisco IE 4000: https://www.cisco.com/c/en/us/td/docs/switches/connectedgrid/cg-switch-sw-master/software/configuration/guide/hsr/b_hsr_ie4k.html

HSR-SAN Mode

In HSR-SAN mode, the RedBox inserts the HSR tag on behalf of the host and forwards the ring traffic, except for frames sent by the node itself, duplicate frames, and frames for which the node is the unique destination. In this mode, packets are handled as follows:

- A source DANH sends a frame passed from its upper layers (“C” frame), prefixes it with an HSR tag to identify frame duplicates, and sends the frame over each port (“A” frame and “B” frame).
- A destination DANH receives two identical frames from each port within a certain interval. The destination DANH removes the HSR tag of the first frame before passing it to its upper layers and discards any duplicate.
- Each node in the HSR ring forwards frames received from one port to the other port of the HSR pair. A node will not forward frames received on one port to the other under the following conditions:
  - The received frame returns to the originating node in the ring.
  - The frame is a unicast frame with a destination MAC address of a node upstream of the receiving node.
  - The node had already sent the same frame in the same direction. This rule prevents a frame from spinning in the ring in an infinite loop.

Cisco Discovery Protocol and Link Layer Discovery Protocol for HSR

HSR supports the Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP), which are Layer 2 neighbor discovery protocols. Both CDP and LLDP can provide information about nodes directly connected to the device. They also provide additional information such as the local and remote interface and device names.

When CDP or LLDP is enabled, you can use the CDP or LLDP information to find the adjacent nodes on an HSR ring and their status. You can then use the neighbor information from each node to determine the complete HSR network topology and debug and locate ring faults.

CDP and LLDP are configured on physical interfaces only.

HSR Alarms

An HSR ring can generate the following two alarms:

- **Partial Ring Fault**—This fault is generated by an HSR RedBox when one of its physical ring ports/links is down. Because the packets can be sent using the redundant path, this is considered as a partial fault. However, this fault still requires user intervention to restore the ring. This is a minor fault and cannot be associated with an external hardware alarm relay.
- **Full Ring Fault**—This fault is generated by an HSR RedBox when both of its physical ring ports/links are down. This is a catastrophic failure and needs immediate attention. This is a major fault and can be associated with an external hardware alarm relay.

When an event that raises an alarm is generated, it can be associated with one or more of the following actions to notify the user:

- **Syslog**—A syslog is generated when the Alarm is raised/cleared.
- **SNMP Notification**—SNMP notification is sent when the alarm is raised/cleared.
- **Relay Output**—External relay contacts can be asserted/de-asserted in response to the alarm. Relays are activated by major faults only.

HSR Guidelines and Limitations

- HSR is supported on the Cisco IE 4000, Cisco IE 4010, and Cisco IE 5000.
- A maximum of four ports support HSR/PRP. Using these four ports, you can create one of the following configurations:
  - HSR rings—Maximum of two
  - PRP channel—Maximum of two
- One HSR ring and one PRP channel
- The maximum number of nodes in the node table is 512. Nodes are all the DANH and VDAN devices that can be connected to the ring at same time.
- HSR ring ports can only be configured in Layer 2 mode.
- Ring 1 and Ring 2 can both be configured in HSR-SAN mode.
- MTU sizes up to 1998 are supported.

- The following protocols and features are mutually exclusive with HSR on the same port:
  - PRP
  - EtherChannels
  - Link Aggregation Control Protocol (LACP)
  - Port Aggregation Protocol (PAgP)
  - Resilient Ethernet Protocol (REP)
- STP and PTP are not supported on the HSR ring.
- By default, all modes of Spanning Tree Protocol (STP) will be disabled on the ring ports.
- PTP must be disabled manually. It is recommended to disable PTP manually before enabling the ring ports.
- Once a port is part of a ring, the media-type, speed, and duplex settings of the port cannot be changed. It is recommended to apply those settings before configuring ring membership.
- Once a port is part of ring, the port cannot be shut down.
- VLAN configuration such as trunk and access mode must be the same on both of the ports participating in the ring.
- After an interface is added in the HSR ring, only the primary interface counters are updated. You should not need to configure and check the status of individual physical interfaces after they are added to the HSR ring.
- As soon as you configure an HSR ring on two ports of a switch, MAC flaps will be observed on other switches where HSR configuration is yet to be applied. We recommend that you shut down the ports before configuring the ring on all switches and then re-enable them one by one as shown below.
- Physical interfaces are predefined for the rings and ports in HSR-SAN and HSR-PRP modes and cannot be changed. Port assignments for Cisco IE 4000 HSR-SAN mode are shown in Table 3. For other devices or modes, refer to the relevant product documentation.

Table 3 Port Assignments for HSR Modes

SKU	HSR Mode	Port Type	Interface Number
Cisco IE 4000	HSR-SAN	Ring 1, Port 1	GE 1/1
		Ring 1, Port 2	GE 1/2
		Ring 2, Port-A	GE 1/3
		Ring 2, Port-B	GE 1/4

Note: Information in this guide is based on Cisco IE 4000 release 15.2.6E2a.

Hot Standby Redundancy Protocol

In the factory network architecture, the distribution switches are the Layer 2/ Layer 3 boundary and serve as gateways. Hot Standby Redundancy Protocol (HSRP) enables gateway redundancy that can be tuned to a failover time to the standby router in less than a second.

HSRP is Cisco's standard method of providing high network availability by providing first-hop redundancy for IP hosts on an IEEE 802 LAN configured with a default gateway IP address. HSRP routes IP traffic without relying on the availability of any single router. It enables a set of router interfaces to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN. When HSRP is configured on a network or segment, it provides a virtual

Media Access Control (MAC) address and an IP address that is shared among a group of configured routers. HSRP allows two or more HSRP-configured routers to use the MAC address and IP network address of a virtual router. The virtual router does not exist; it represents the common target for routers that are configured to provide backup to each other. One of the routers is selected to be the active router and another to be the standby router, which assumes control of the group MAC address and IP address should the designated active router fail.

Resilient Ethernet Protocol

REP is a Cisco proprietary protocol that provides an alternative to Spanning Tree Protocol (STP). REP provides a way to avoid network loops, handle link failures, and improve convergence time. It 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.

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 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.

Figure 6 shows an example of a segment consisting of eight ports spread across five switches. Ports A and B are configured as edge ports. When all ports are operational, a single port is blocked, as shown by the diagonal line. When a failure exists in the network, the blocked port returns to the forwarding state to minimize network disruption.

Figure 6 REP Segment

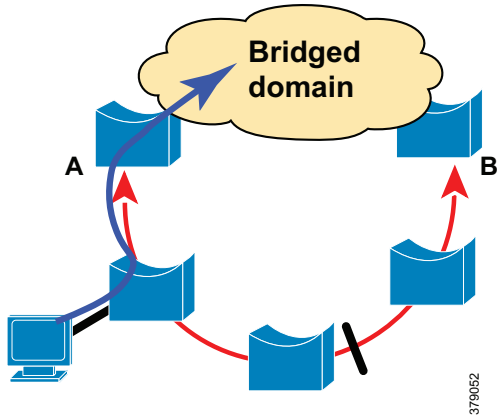

With REP, at least one port is always blocked in any given segment: that is, the alternate port. The blocked port helps ensure that the traffic within the segment is loop-free by requiring traffic flow to exit only one of the edge ports and not both. So when a failure occurs in the segment, REP opens the alternate port so traffic can reach the edge of the segment.

Internet Group Management Protocol Snooping

Internet Group Management Protocol (IGMP) snooping software examines multi-cast group messages within a VLAN to discover which interfaces are connected to hosts or other devices interested in receiving this traffic. Using the interface information, IGMP snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding multicast traffic to the entire VLAN. The IGMP snooping feature tracks which ports are attached to multicast-capable routers to help it manage the forwarding of IGMP membership reports. As the name implies, IGMP snooping requires the LAN switch to snoop on the IGMP transmissions between the host and the router and to keep track of multicast groups and member ports. When the switch receives an IGMP report from a host for a particular multicast group, the switch adds the host port number to the forwarding table entry; when it receives an IGMP Leave Group message from a host, it removes the host port from the table entry. It also periodically deletes entries if it does not receive IGMP membership reports from the multicast clients.

The multicast router sends out periodic general queries to all VLANs. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry. The switch creates one entry per VLAN in the IGMP snooping IP multicast forwarding table for each group from which it receives an IGMP join request.

Deploying High-Availability Seamless Redundancy in the Factory Network

The factory network follows a hierarchical design that includes three main layers:

- **Access Layer**—Provides endpoints and users with direct access to the network. IACS devices are connected to this layer. In the current implementation, this layer consists of Cisco IE 4000 switches in a ring topology deployed in HSR-SAN mode. The HSR ring can be attached to the distribution switch in two ways:
  - Using REP as shown in Figure 3. A REP segment would be formed by two ring switches and the distribution switches.
  - Without a Layer 2 redundancy protocol as shown in Figure 4. In this scenario, the distribution switches are not connected directly to each other.
- **Distribution Layer**—Aggregates access layers and provides connectivity to services. It acts as a Layer 2/Layer 3 boundary and default gateway for the Layer 2 domains. In this implementation, this layer consists of Cisco IE 5000 switches running HSRP for Layer 3 redundancy.
- **Core Layer**—Provides connectivity between distribution layers for large LAN environments.

Access to Distribution Connectivity

As mentioned before, the difference between deployments shown in Figure 3 and Figure 4 is the way the access layer is connected to the distribution layer. In the first scenario a REP segment is formed with two switches in the HSR ring and the distribution as shown in Figure 7. It is recommended to use REP to connect to the distribution layer since it minimizes the probability of HSRP failover, Deployment without REP should be considered only if interoperability is a concern.

Figure 7 REP Segment to Distribution

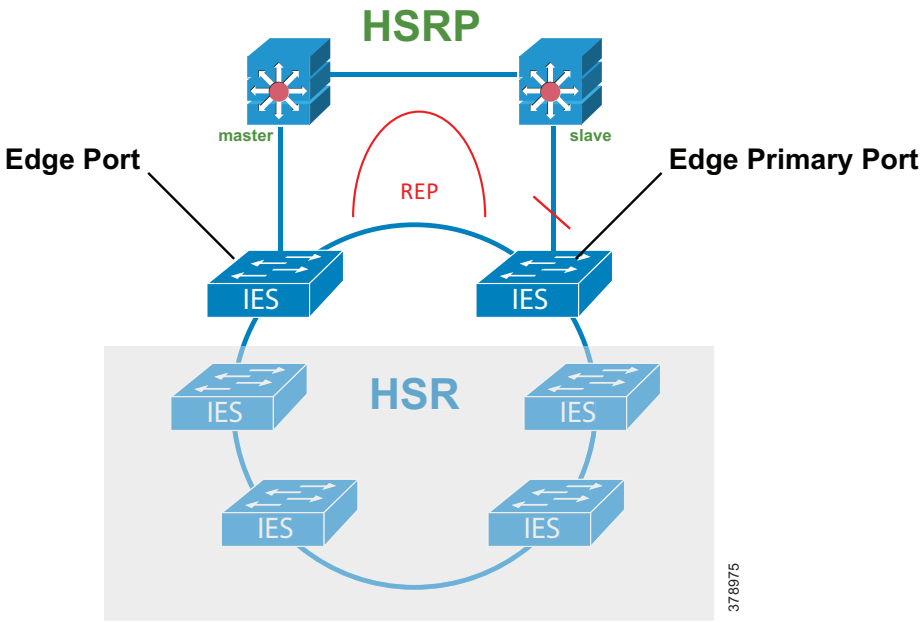

This topology combines hitless HSR topology in the cell area network with REP in the uplinks to minimize the probability of HSRP and routing protocol convergence in the event of a failure in the uplinks. The only single event that would trigger an HSRP failover would be the failure of the HSRP master node.

Configuration

This section describes how to configure HSR and related features:

- Configuring an HSR ring
- Configuring HSRP gateways
- Configuring REP segment

Configuring an HSR Ring

HSR configuration applies to the Cisco IE 4000 switches in the ring.

Before configuring HSR, check if HSR is enabled; newer versions have it enabled by default.

```
Switch# show version | inc Feature
Feature Mode                    : 0x25 Enabled: HSR (Disabled: MRP TSN)
```

If HSR is enabled, skip this step; otherwise, use the following command to enable:

```
Switch# license right-to-use activate hsr
```

For the change to take effect, the switch must be reloaded. Confirm the reload when prompted and wait for the switch to reload and boot. Verify that the HSR feature is activated.

Ensure that the member interfaces of a HSR ring are not participating in any redundancy protocols such as FlexLinks, EtherChannel, and REP before configuring a HSR ring.

Follow these steps to configure HSR:

1. Shut down the ports before configuring the HSR ring:

```
interface range GigabitEthernet1/1-2
shutdown
```

2. Configure switch port and VLANs as desired:

```
switchport mode trunk
switchport trunk allowed vlan 10,20,900 switchport trunk native vlan 900
```

3. Disable PTP. As explained in Choosing an HSR Implementation, page 7, PTP is not supported:

```
no ptp enable
```

4. Create the HSR ring interface and assign the ports to the HSR ring. This command should be issued in the interface configuration. The two interfaces will be bundled in a HSR interface:

```
hsr-ring 1
```

5. Turn on the HSR interface:

```
no shutdown
```

6. Make sure the enable DualUplinkEnhancement feature is not disabled. This feature is required to support the connectivity to a dual router (HSRP in this case) on the distribution layer:

```
Show run | include fpgamode-DualUplinkEnhancement
```

If the output shows *no hsr-ring 1 fpgamode-DualUplinkEnhancement*, issue the following command.

```
hsr-ring 1 fpgamode-DualUplinkEnhancement
```

Follow these optional steps to configure CDP and LLDP to provide information about HSR ring nodes:

7. Enable LLDP globally:

```
lldp run
```

8. Enable LLDP on the ports to be assigned to the HSR ring:

```
interface range GigabitEthernet1/1-2
lldp transmit
lldp receive
```

9. Enable CDP on the ports to be assigned to the HSR ring:

```
interface range GigabitEthernet1/1-2
cdp enable
```

Follow these optional steps to enable HSR alarms:

10. Enable the HSR alarm facility:

```
alarm facility hsr enable
```

11. Enable SNMP notification for HSR alarms:

```
alarm facility hsr notifies
```

12. Associate HSR alarms with the Major Relay:

```
alarm facility hsr relay major
```

Configure a REP Segment

REP configuration applies to the switches on the REP segment, as shown in Figure 7.

Configure Administrative VLAN

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

Follow these guidelines when configuring the REP administrative VLAN:

- Only one administrative VLAN can exist on a router and on a segment. However, the software does not enforce this.
- If you do not configure an administrative VLAN, the default is VLAN 1.
- If you want to configure REP on an interface, ensure that the REP administrative VLAN is part of the Trunk EFP encapsulation list:

```
vlan <vlanID>
name REP_Admin_VLAN
rep admin vlan <vlanID>
```

Enable REP on Interfaces

For the REP operation, you must enable REP on each segment interface and identify the segment ID. This task 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.

Edge Ports

To configure a port as an edge port, use the following command in interface configuration mode:

```
rep segment <ID> edge (primary)
```

The **primary** keyword is optional and allows for manual selection of the primary edge. If the primary keyword is used, the other edge port becomes the secondary edge port (no keyword required). To configure the secondary edge port, omit the primary keyword as shown:

```
rep segment <ID> edge
```

Non-Edge Ports

To configure a port as a member of the REP segment, use the following command in interface configuration mode:

```
rep segment <ID>
```

Preemption

Preemption is done either manually with the **rep preempt segment** <*ID*> command, or automatically if you configure **rep preempt delay** <*seconds*> under the primary edge port.

When a segment heals after a link failure, one of the two ports adjacent to the failure comes up as the ALT port. Then, after preemption, the location of the ALT ports become the primary edge port unless additional configuration is done for load balancing and alternate port, which is not covered in this document. For more information, refer to *Cisco Industrial Ethernet 4000, 4010 and 5000 Switch Software Configuration Guide*: https://www.cisco.com/c/en/us/td/docs/switches/lan/cisco_ie4010/software/release/15-2_4_EC/configuration/guide/scg-ie4010_5000.html

Example of automatic preemption:

```
interface GigabitEthernet1/1
rep segment 30 edge primary
rep preempt delay 30
```

Example of manual preemption:

```
SWITCH#rep preempt segment 30
The command will cause a momentary traffic disruption.
Do you still want to continue? [confirm]

Proceeding with Manual Preemption
```

Configuring HSRP

HSRP configuration applies to the two distribution switches (Cisco IE 5000s). The **standby ip interface configuration** command activates HSRP on the configured interface. If an IP address is specified, that address is used as the designated address for the Hot Standby group. If no IP address is specified, the address is learned through the standby function. You must configure at least one Layer 3 port on the LAN with the designated address. Configuring an IP address always overrides another designated address currently in use. It is recommended to configure the lowest IP in the network as standby IP to guarantee that the master router will become the IGMP snooping querier.

In the current implementation, HSRP is configured in a Switch Virtual Interface (SVI). To configure HSRP, assign a virtual IP and group number to the interface. The following is an example of HSRP configuration in master peer:

```
interface Vlan10
ip address 10.17.10.2 255.255.255.0
standby 1 ip 10.17.10.1
```

The following is an example of the standby peer:

```
interface Vlan10
ip address 10.17.10.3 255.255.255.0
standby 1 ip 10.17.10.1
```

Note that virtual IP is the same while physical IP varies per peer.

Configuring HSRP Priority

When configuring HSRP priority, follow these guidelines:

- Assigning a priority allows you to select the active and standby routers. If preemption is enabled, the router with the highest priority becomes the active router again after recovering from a failure. If priorities are equal, the current active router does not change.
- The highest number (1 to 255) represents the highest priority (most likely to become the active router).
- When routing is first enabled for the interface, it does not have a complete routing table. If it is configured to preempt, it becomes the active router, even though it is unable to provide adequate routing services. To solve this problem, configure a delay time to allow the router to update its routing table.

To configure priority in the desired active peer, add this line to the interface configuration (since default priority is 100, the configured number should be higher):

```
standby 1 priority 254
```

Configure the router to preempt, which means that when the local router has a higher priority than the active router, it assumes control as the active router. As an option a delay can be configured, which will cause the local router to postpone taking over the active role for the number of seconds shown:

```
standby 1 preempt delay minimum
```

HSRP uses two timers: hello interval and hold time. Hello interval defines the frequency that hello packets are sent to the other peer. Hold time indicates the amount of time to wait before marking the peer as down. The hold time should be three or more times greater than the hello interval. To configure those timers:

```
standby 1 timers msec 200 msec 750
```

Internet Group Management Protocol Design Considerations

IGMP snooping should be configured to route multicast traffic only to those hosts that request traffic from the specific multicast group. IGMP snooping is configured by default in Cisco IE switches, but IGMP snooping querier should be configured in the distribution switches (Cisco IE 5000s) using the following command:

```
ip igmp snooping querier
```

IGMP selects the querier with the lowest IP in the network, hence the importance of configuring the HSRP IP to be the lowest in the network.

Best Practices

- Configure preemption in HSRP for deterministic routing.
- If REP preemption is required, it is recommended to do manual preemption to avoid an unplanned downtime. REP preemption could cause a multicast tree re-convergence that affects nodes attached to the REP segment.
- For REP segment, the edge port in the Cisco IE 4000 connected directly to HSRP slave should be primary so it gets blocked by default in preemption.

- Enable BPDU filtering in ports connecting to end devices and distribution on the Cisco IE 4000 participating in HSR ring to avoid ports getting into a blocked state after topology changes.
- Avoid using access ports on the distribution switch for VLANs being used in the ring to avoid a HSRP split brain scenario. If connecting devices directly to the distribution switches, use a different VLAN.

Network and Ring Size Considerations

- The maximum number of nodes in the node table is 512. Nodes are all the DANH and VDAN devices that can be connected to the ring at same time.
- A maximum ring size of 24 switches is recommended. Given that HSR gives protection for a single failure in the ring, increasing the size would also increase the probability of more concurrent failures.

Validation Results

This section describes the convergence values observed during validation efforts.

Validation was done for a single HSR ring with the following characteristics:

- Single HSR ring with 24 nodes
- MAC scale up to 250 simulated devices connected to a single node
- 200 multicast groups
- Two VLANs in the ring
- Concurrent Multicast and Unicast traffic
- Intra VLAN traffic (Layer 2) and Inter VLAN traffic
- HSRP timers used (Hello interval = 200ms and hold time = 750ms)
- REP native VLAN

Test Platforms and Software Versions

Table 4 Test Platforms and Software Versions

Product/Platform	Software Release	Role
Cisco IE 4000	15.2.6E2a	Access
Cisco IE 5000	15.2.6E2a	Distribution

Convergence Values

The following tables show convergence times for traffic inside the HSR ring. Table 5 refers to a disruption inside the ring. Table 6 shows how a disruption on the REP segment affects the traffic inside the ring. Table 7 shows how an HSRP failover affects the traffic inside the ring.

Table 5 Single Failure in HSR Ring

Traffic Type	Disruption Type	Convergence Time	
		Max	Average
Layer 2 unicast	Link	0	0
	Switch	0	0
Layer 2 multicast	Link	0	0
	Switch	0	0
Layer 3 unicast (Inter VLAN traffic)	Link	0	0
	Switch	0	0

Table 6 Single Failure in REP Segment

Traffic Type	Disruption Type	Convergence Time During Failure (ms)		Convergence Time During REP Preemption (ms)	
		Max	Average	Max	Average
Layer 2 unicast	Link	0	0	0	0
	Switch	0	0	0	0
Layer 2 multicast	Link	0	0	0	0
	Switch	9102	468	5404	116
Layer 3 unicast (Inter VLAN traffic)	Link	276	144	244	137
	Switch	1424	779	214	116

Table 7 HSRP Failover

Traffic Type	Disruption Type	Convergence Time During Failure (ms)		Convergence Time During HSRP Preemption (ms)	
		Max	Average	Max	Average
Layer 2 unicast	Link	NA	NA	NA	NA
	Switch	0	0	0	0
Layer 2 multicast	Link	NA	NA	NA	NA
	Switch	20	1	8000	127
Layer 3 unicast (Inter VLAN traffic)	Link	NA	NA	NA	NA
	Switch	6446	1151	4012	168

Troubleshooting

HSR Troubleshooting

The following CLI can be used to troubleshoot HSR:

- **show hsr ring** {**1** | **2**} [**detail**] displays configuration details and current state for the specified HSR ring:

```
IE4000-1# sh hsr ring 2 detail
HSR-ring: HS2
------------
 Layer type = L2
 Operation Mode = mode-H
 Ports: 2       Maxports = 2
 Port state = hsr-ring is Inuse
 Protocol = Enabled  Redbox Mode = hsr-san
Ports in the ring:
  1) Port: Gi1/3
   Logical slot/port = 1/3      Port state = Inuse    ' Port is up
        Protocol = Enabled
  2) Port: Gi1/4
   Logical slot/port = 1/4      Port state = Inuse    ' Port is up
        Protocol = Enabled

Ring Parameters:
 Redbox MacAddr: f454.3365.8a84
 Node Forget Time: 60000 ms
 Node Reboot Interval: 500 ms
 Entry Forget Time: 400 ms
 Proxy Node Forget Time: 60000 ms
 Supervision Frame COS option: 0
 Supervision Frame CFI option: 0
 Supervision Frame VLAN Tag option: Disabled
 Supervision Frame MacDa: 0x00
 Supervision Frame VLAN id: 0
 Supervision Frame Time: 3 ms
 Life Check Interval: 2000 ms
 Pause Time: 25 ms
fpgamode-DualUplinkEnhancement: Enabled
```

- **show hsr statistics** displays statistics for HSR components. To clear HSR statistics information, enter the command **clear hsr statistics**.
- **show hsr node-table** displays all MAC addresses accessible to the switch using the HSR interface, including other nodes in the ring as well as devices attached to other nodes.
- **show hsr vdan-table** displays the HSR Virtual Doubly Attached Node (VDAN) table, which contains devices directly connected to the switch for whom the switch acts as proxy. This table is also known as the Proxy node table.
- **show cdp neighbors** and **show lldp neighbors** displays neighbor information for the switch, which is useful when troubleshooting connectivity issues.

- **show alarm settings | begin hsr** displays HSR alarm configuration. Table 8 lists the HSR events and their representations.

Table 8 HSR Events

Event Number	Event Description	System Log (Level)	Alert/Alarm Log	Alarm LED and Output Relay
1	Ring goes from UP to DOWN state.	2	2	Major Alarm/Assert
2	Ring goes from DOWN to UP state.	6	6	De-assert
3	One ring port goes DOWN and the other ring port and the ring itself are UP.	3	3	
4	Both ring ports are UP again.	6	6	

- You can view currently active alarms using the **show facility alarm status** command. The following example shows alarm status for minor and major HSR alarms:

```
Switch#show facility-alarm status
Source          Severity    Description                      Relay   Time
Switch            MINOR    34 HSR ring is partially down        MAJ    Oct 24 2017 10:16:10
-------
Switch# show facility-alarm status
Source          Severity    Description                   Relay     Time
Switch            MAJOR    33 HSR ring is down              MAJ      Oct 24 2017 10:17:07
```

REP Troubleshooting

- Enter this command in order to see the status of a REP adjacency:

```
SWITCH#show int gi1/7 rep
Interface                Seg-id Type            LinkOp      Role
------------------------ ------ --------------- ----------- ----
GigabitEthernet1/7       10     Primary Edge    TWO_WAY     Alt
```

- Use the **show rep topology** command on any router on the segment to see the current topology:

```
sh rep topology
```

HSRP Troubleshooting

- The commands **show standby** and **show standby brief** provide configuration and current status details.

```
IE5K-3#show standby
Vlan10 - Group 1
  State is Active
    7 state changes, last state change 2w1d
  Virtual IP address is 10.17.10.1
  Active virtual MAC address is 0000.0c07.ac01 (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
 Hello time 200 msec, hold time 750 msec
    Next hello sent in 0.144 secs
  Preemption enabled
  Active router is local
  Standby router is 10.17.10.3, priority 170 (expires in 0.736 sec)
```

```
  Priority 200 (configured 200)
  Group name is "hsrp-Vl10-1" (default)
Vlan20 - Group 1
  State is Active
    7 state changes, last state change 2w1d
  Virtual IP address is 10.17.20.1
  Active virtual MAC address is 0000.0c07.ac01 (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 200 msec, hold time 750 msec
    Next hello sent in 0.160 secs
  Preemption enabled
  Active router is local
  Standby router is 10.17.20.3, priority 170 (expires in 0.656 sec)
  Priority 200 (configured 200)
  Group name is "hsrp-Vl20-1" (default)
IE5K-3#
IE5K-3#sh standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Vl10        1    200 P Active  local           10.17.10.3      10.17.10.1
Vl20        1    200 P Active  local           10.17.20.3      10.17.20.1
```

- If HSRP does not recognize its HSRP peers, verify physical layer connectivity and configuration. Make sure **fpgamode-DualUplinkEnhancement** is configured in the HSR ring switches connected to the distribution layer.

Related Documents

- Cisco Industrial Ethernet 4000 Series Switches
  http://www.cisco.com/c/en/us/support/switches/industrial-ethernet-4000-series-switches/tsd-products-support-series-home.html
- IEC 62439-3, Industrial communication networks-High availability automation networks-Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)
  https://webstore.iec.ch/publication/24447

Glossary

Term	Definition
CDP	Cisco Discovery Protocol
CIP	Common Industrial Protocol
DANH	Doubly Attached Nodes
GLBP	Gateway Load Balancing Protocol
HSR	High-availability Seamless Redundancy
HSRP	Hot Standby Router Protocol
IACS	Industrial Automation and Control System
IDMZ	Industrial Demilitarized Zone
IE	Industrial Ethernet
IGMP	Internet Group Management Protocol
IoT	Internet of Things
IT	Information Technology
LLDP	Link Layer Discovery Protocol
MSTP	Multiple Spanning Tree Protocol
OT	Operational Technology
PLC	Programmable Logic Controllers
PRP	Parallel Redundancy Protocol
PTP	Precision Time Protocol
RedBox	Redundancy Box
REP	Resilient Ethernet Protocol
RSTP	Rapid Spanning Tree Protocol
SAN	single attached node
STP	Spanning Tree Protocol
VRRP	Virtual Router Redundancy Protocol

Glossary