



Deploying the Resilient Ethernet Protocol (REP) in a Converged Plantwide Ethernet System (CPwE) Design Guide

March 23, 2015

Rockwell Automation and Cisco Four Key Initiatives:

- **Common Technology View:**
A single system architecture, using open, industry standard networking technologies, such as Ethernet and IP, is paramount for achieving the flexibility, visibility, and efficiency required in a competitive manufacturing environment.
- **Converged Plantwide Ethernet Architectures:**
These manufacturing focused reference architectures, comprised of the Rockwell Automation Integrated Architecture™ and Cisco's Ethernet to the Factory, provide users with the foundation for success to deploy the latest technology by addressing topics relevant to both engineering and IT professionals.
- **Joint Product and Solution Collaboration:**
Stratix 5700™ and 8000™ Industrial Ethernet switches incorporating the best of Cisco and the best of Rockwell Automation.
- **People and Process Optimization:**
Education and services to facilitate Operational Technology (OT) and IT convergence and allow successful architecture deployment and efficient operations allowing critical resources to focus on increasing innovation and productivity.



Converged Plantwide Ethernet System (CPwE)

This CPwE-REP Cisco® Validated Design (CVD) describes the implementation of Resilient Ethernet Protocol (REP) for a switch ring topology in the CPwE system.

"CPwE-REP System Introduction" section on page 2

- [CPwE Overview, page 4](#)
- [Key Concepts, page 5](#)
- [System Features, page 5](#)
- [CPwE-REP Use Case Overview, page 6](#)

"System Design Considerations" section on page 7

- [Choosing the Right Resiliency Protocol for Your Application, page 7](#)
 - [REP Technology Overview, page 9](#)
 - [REP Terminology, page 9](#)
 - [REP Operation, page 10](#)
- [Applying REP in IACS Applications, page 13](#)
- [REP Design Considerations, page 13](#)
 - [Network and Ring Size Considerations, page 14](#)
 - [Distribution Switch, page 15](#)
 - [Single Segment versus Multiple Segments, page 17](#)

"Configuring the Infrastructure" section on page 18

- [Configuring REP Ring, page 18](#)
 - [Native VLAN Implementation, page 19](#)
 - [Admin VLAN for REP, page 19](#)
 - [REP Segment Configuration, page 20](#)
- [Stack or Standalone, page 21](#)
- [Device Manager, page 21](#)

"Test Objectives and Setup" section on page 22

- [Test Objectives, page 22](#)
- [Lab Setup, page 23](#)
 - [Test Topology, page 23](#)
 - [Test Platforms and Software Versions, page 23](#)
 - [Test Configurations, page 24](#)
 - [IACS Setup, page 24](#)
 - [Test Tools, page 25](#)
- [Test Setup, page 25](#)

"REP Troubleshooting Tips" section on page 26

CPwE-REP System Introduction

The CPwE-REP system establishes a resilient network architecture that goes beyond previous CPwE-explored resiliency protocols (such as Rapid Per-VLAN Spanning Tree+ (RPVST+), Multiple Spanning Trees Protocol (MSTP), EtherChannel, and Flex Links) within Industrial Automation and Control System (IACS) applications. All other CPwE design recommendations such as segmentation, data prioritization, and security, still apply.

REP delivers better convergence characteristics (described later in this document), and functions as an alternative to resiliency protocols described in the CPwE Design and Implementation Guide (DIG).

The key benefits of REP include:

1. REP provides easier migration from Spanning Tree Protocol (e.g., RSTP and MSTP).
2. REP is simple and easy to configure.
3. REP provides faster convergence (recovery from a failure) than RSTP or MSTP for a switch ring topology.
4. When repairing from a fault within the ring, REP will not initiate another convergence event (like other resiliency protocols).

An IACS is deployed in a wide variety of industries such as automotive, pharmaceuticals, consumer goods, pulp and paper, oil and gas, and energy. IACS applications are also deployed in a wide variety of manufacturing disciplines such as batch, discrete, process, and hybrid manufacturing.

As noted in the CPwE DIG, the Cell/Area Zone (Figure 2) is where the IACS end devices (Levels 0-2) connect into the Cell/Area Zone local area network (LAN).



Note

To achieve optimal design and performance of the Cell/Area Zone LAN and IACS devices, careful planning is required. This release of the CPwE architecture focuses on EtherNet/IP™, which is driven by the ODVA Common Industrial Protocol (CIP). Refer to the IACS Communication Protocols section of the CPwE Design and Implementation Guide.

http://literature.rockwellautomation.com/idc/groups/literature/documents/td/enet-td001_-en-p.pdf or http://www.cisco.com/en/US/docs/solutions/Verticals/CPwE/CPwE_DIG.html

This CPwE-REP CVD outlines the key requirements and technical considerations for REP within the Cell/Area Zone:

- Converged Plantwide Ethernet (CPwE) Overview
- CPwE-REP Use Case Overview
- REP Overview—Availability with Flexibility
- Key Design Considerations—Reliability with Performance

Seamless and resilient convergence between IACS applications within the Cell/Area Zone and the Level 3 Site Operations within the Industrial Zone requires a resiliency protocol that is applicable to both industrial and IT technologies. REP is supported in the Rockwell Automation Stratix™ Industrial Ethernet Switches (IES), the Cisco Industrial Ethernet switches, and the Cisco Catalyst 3750X Distribution switch (stack and non-stack).

REP is a well-established and proven resiliency protocol that has been deployed in both enterprise and service provider (SP) applications for many years. Although the published network convergence recovery time on fiber interfaces is less than 200ms, testing and validation of REP by Cisco and Rockwell Automation demonstrates that REP is suitable for many IACS applications that

require a switch ring topology to meet application availability requirements (Figure 1). REP is suitable for IACS applications that can tolerate up to a 100 ms network convergence recovery time on fiber interfaces.

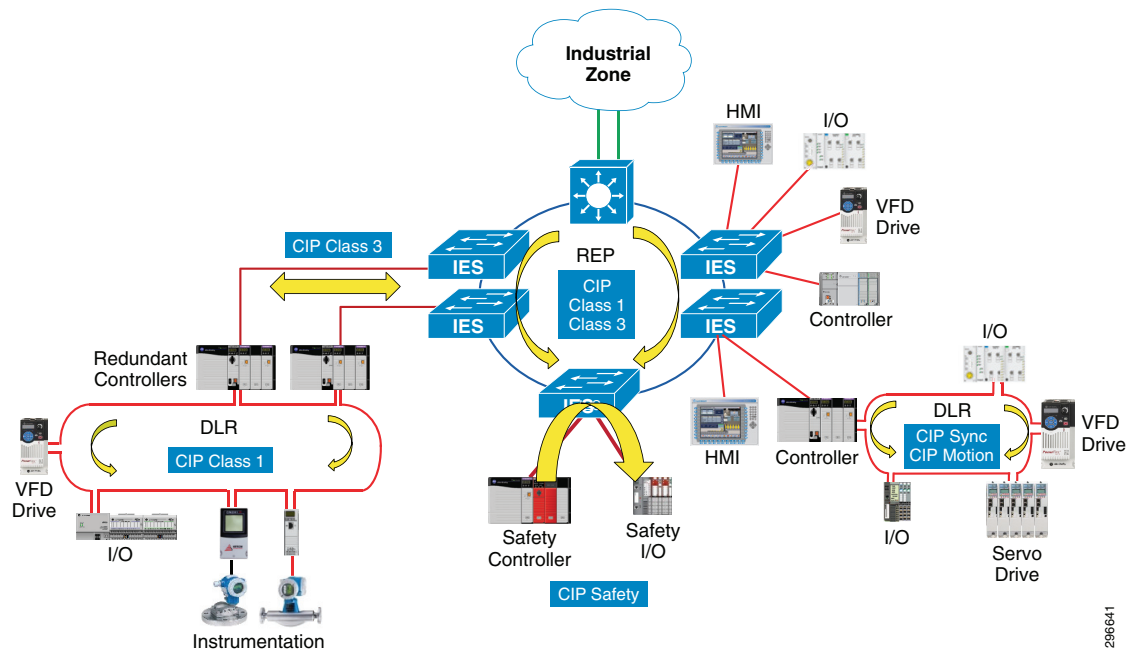
- Controller to HMI
- Controller to Controller
- Controller to I/O
- Controller to VFDs
- Controller to MCCs

For IACS applications that require a faster network convergence recovery time, Cisco and Rockwell Automation recommend either a redundant star switch topology with the Flex Links resiliency protocol, or a device-level ring topology, such as the ODVA DLR (Figure 1):

- Integrated motion applications utilizing CIP Motion™
- Applications with a high degree of multicast traffic
- Integrated safety applications utilizing CIP Safety™

Refer to [Rockwell Automation Embedded Switch Technology Reference Architectures](#) publication for more information.

Figure 1 REP IACS Applications



The test and validation of REP within the CPwE system builds upon established standards and partnerships, while expanding network resiliency functionality. The Cell/Area Zone LAN-switched architecture has been enhanced to include hardware models and software environments, and test and validation of network resiliency models (e.g., REP for ring topologies) available for an IACS application.

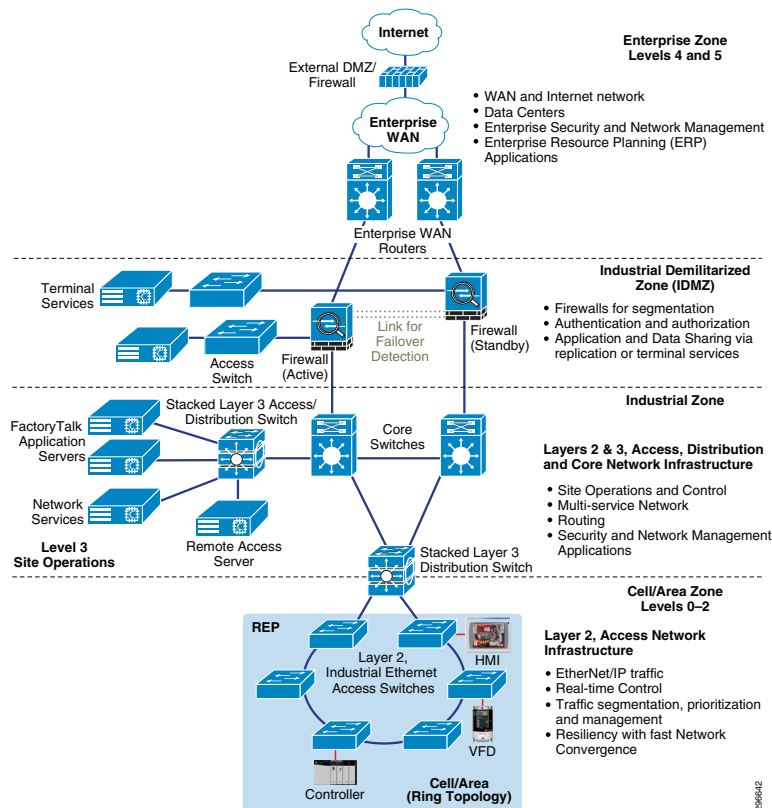
CPwE Overview

CPwE is an architecture that provides network services to IACS devices and equipment, and securely integrates those into the wider enterprise network. CPwE, which is a collaborative effort of Cisco Systems, Inc. and Rockwell Automation, reflects the IT and IACS knowledge and expertise of both companies.

CPwE defines a framework for the IACS devices, equipment, and basic network services that form the system architecture. An existing framework that identifies the levels of an IACS network is the Purdue Model for Control Hierarchy (ISBN 1-55617-265-6), an industry reference. The CPwE framework incorporates this model since defines different levels of operations.

The CPwE system architecture implements strict traffic segmentation to protect IACS applications from external and internal interruptions. Disruptions in the IACS create the greatest impact to the functionality of the production facility and are the primary consideration in this CPwE system architecture. Because of the different security requirements of the different levels, security technology limitations, different access requirements, and implications of failure in the Industrial Zone, an Industrialized Demilitarized Zone (IDMZ) was established. Levels 4-5 of the Enterprise Zone are similar to traditional enterprise networks and have similar availability requirements although access to the IDMZ and below is highly controlled. The CPwE-REP architecture is consistent with previous CPwE systems from an overall architecture framework perspective and with general design and implementation recommendations in order to continue to align with industry standards.

Figure 2 CPwE Overall Architecture



This release of CPwE-REP refreshes, expands, and enhances the previous CPwE program through the integration of the REP for switch ring topologies.

Key Concepts

Key CPwE concepts are included in the CPwE-REP system for fiber network-interconnected Industrial Ethernet Switches (IES), which include the following:

- **Availability**—The choice of LAN topology plays a pivotal role in determining overall IACS application uptime and productivity. IACS application requirements such as availability, performance, and geographic dispersion of equipment drive the choice of topology. For critical operations where uptime is crucial, a redundant path network topology, enabled by managed IES, helps provide maximum network robustness and availability. Whether deploying a ring or redundant star topology, a resiliency protocol, such as REP, is required to prevent Layer 2 loops while maintaining the redundant path topology. Without REP, a redundant path LAN would cause Ethernet frames to loop for an indefinite period of time, thus affecting performance and reliability of that Cell/Area Zone LAN.
- **Predictable Performance**—Meeting the predictable, reliable, real-time traffic requirements of IACS applications is a fundamental requirement for any successful CPwE deployment and the highest priority consideration and concern from manufacturing customers (process and discrete). Network convergence is also important to realize value from a redundant network, but predictable, real-time traffic performance is the highest priority requirement.
- **Efficiency and Ease-of-Use**—Non-IT personnel with limited Ethernet and IP networking skills, especially for IES access layer switches, are often responsible for deploying, configuring, and managing IACS networks. Ease-of-use, replacement and overall system simplicity, therefore, are key considerations
- **Industrial Protocols**—Manufacturing networking equipment needs to support industrial protocols (CIP, etc.) from a management perspective, and the network design and configuration needs to be optimized to support industrial protocol traffic patterns (e.g., I/O) and configurations (e.g., port restrictions for access control lists, etc.). Only standard Ethernet and IP protocol suite-based protocols are considered.

System Features

IACS network environments have evolved over the years, driven by a number of key design features. These features are not specific to industrial Ethernet, but to networking for industrial automation and control systems in general. In the move towards ruggedized industrial Ethernet, and ruggedized industrial wireless infrastructure in the future, many of these design features still apply although their importance sometimes shifts.

The CPwE system follows seven key features that the industry expects as best practices:

- Real-Time Communication and Performance
- Availability (low MTTR with high OEE)
- Traffic Segmentation
- Physicality (Physical Layout, Infrastructure, Robustness and Topology)
- Application Technology Coexistence
- Security (holistic defense-in-depth)
- Scalability (from OEM equipment to plant-wide architectures)

CPwE-REP Use Case Overview

The CPwE-REP use cases tested and validated by Cisco and Rockwell Automation demonstrate the scalability of REP to support a variety of IACS applications sizes:

Figure 3 Single REP Switch ring, All IES Solution, Single VLAN per REP Segment

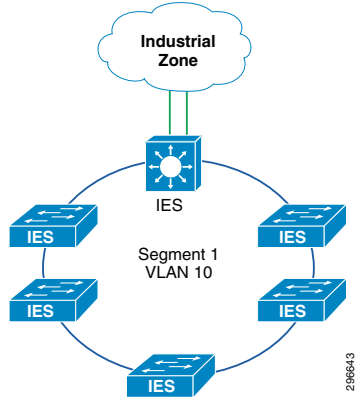


Figure 4 Multiple REP Switch Rings, All IES Solution, Single VLAN per REP Segment

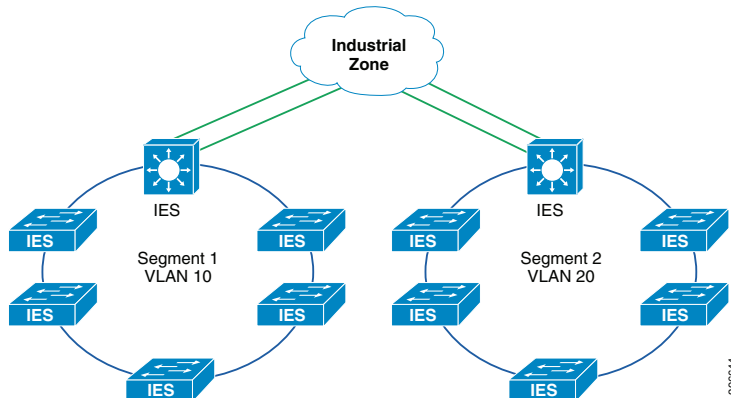


Figure 5 Multiple REP Switch Rings, IES and Distribution Switch (non-stacked) Solution, Single VLAN per REP Segment

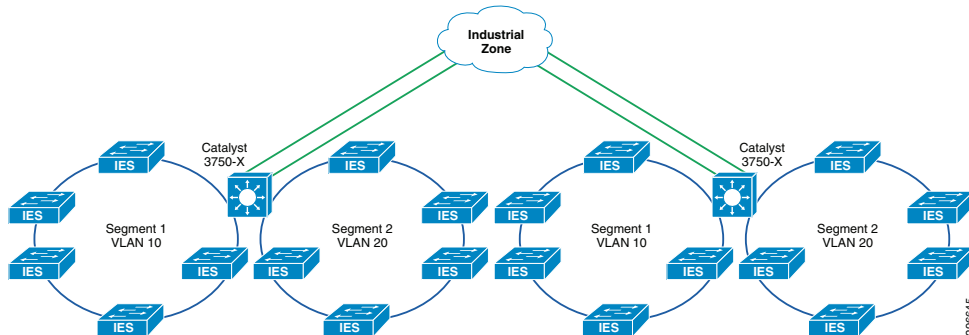



Figure 6 Multiple REP Switch Rings, IES and Distribution Switch (Stacked) Solution, Single VLAN per REP Segment



System Design Considerations

This section describes system design considerations for adding REP to CPwE.

Choosing the Right Resiliency Protocol for Your Application

The first step in selecting the proper resiliency protocol for the IACS should be assessing application and hardware requirements. You should consider failover time requirements, architectural limitations such as geographic dispersion, location within the hierarchal architecture, legacy connectivity, and the requirement for standard protocol support.

The adoption of industrial Ethernet has caused the creation of large Layer 2 domains requiring fast convergence. In particular, manufacturers require fast convergence to support their manufacturing automation deployments. It also should support fast convergence when scaling the number of MAC addresses. REP meets these requirements for fast convergence in Layer 2 ring topologies.

Table 1 provides guidance for the resiliency protocol that provides the best solution based on all requirements.

Table 1 Resiliency Protocol Selection Table Based on Application Requirements

Resiliency Protocol	Mixed Vendor	Ring	Redundant Star	Network Convergence > 250 ms	Network Convergence Sub 250 ms	Network Convergence 50 - 150 ms	Network Convergence 1 - 3 ms	L3	L2
STP (802.1D)	X	X	X						X
RSTP (802.1w)	X	X	X	X					X
MSTP (802.1s)	X	X	X	X					X
RPVST+		X	X	X					X
REP		X				X			X
EtherChannel (LACP 802.3ad)	X		X		X				X
Flex Links			X			X			X
DLR (IEC & ODVA)	X	X					X		X
StackWise		X	X		X			X	X
HSRP		X	X	X				X	

Table 1 Resiliency Protocol Selection Table Based on Application Requirements (continued)

Resiliency Protocol	Mixed Vendor	Ring	Redundant Star	Network Convergence > 250 ms	Network Convergence Sub 250 ms	Network Convergence 50 - 150 ms	Network Convergence 1 - 3 ms	L3	L2
GLBP		X	X	X				X	
VRRP (IETF RFC 3768)	X	X	X	X				X	

The next step in resiliency selection is looking at the desired topology. The major drivers in topology choice are the equipment that you plan to deploy and the level of redundancy you desire. Table 2 provides the equipment and the resiliency provided by the equipment in specific topologies.

Table 2 Hardware Topology & Resiliency Options

Topology / Resiliency Protocol	Redundant Star – Switch-level (MSTP, EtherChannel, Flex Links)	Star – Switch-level (None)	Ring – Switch-level (MSTP, REP)	Ring – Device-level (Device Level Ring Protocol - DLR)	Linear (None)
Stratix 5700 or IE 2000	X	X	X		X
Stratix 8300/8000 or IE 3000	X	X	X		X
Embedded 2 Port Switch				X	X
Catalyst 3750-X	X	X	X		X

Resiliency Protocol Selection Considerations:

- Use fiber media and SFPs for all inter-switch links – ring and redundant star switch-level topologies.
- Use MSTP for multi-vendor switch deployment, redundant star or ring switch-level topologies, with CIP explicit messaging such as HMI, or unicast CIP implicit I/O applications with an RPI of greater than or equal to 100 ms.
- Use Flex Links for Cisco/Rockwell Automation switch deployment, redundant star switch-level topology, with unicast or multicast CIP implicit I/O applications.
- Use REP or DLR for ring topology, with CIP implicit I/O applications.
- Use DLR for ring device-level topology, for applications such as CIP Safety, ControlLogix Redundancy, multicast CIP I/O applications, and CIP Motion.



Note

The fiber media recommendation is based on the testing and validation of REP by Cisco and Rockwell Automation, with the Allen-Bradley programmable automation controllers (PACs) configured with the default requested packet interval (RPI) value of 20 ms for standard I/O.

In many MCC-based applications, device and I/O RPI settings as fast as the default 20 ms are not required. For applications that can perform appropriately with device and I/O RPI settings of 100 ms or greater, a Fast Ethernet (100 Mbps) copper inter-switch link may provide sufficient convergence in a REP topology.

- The default RPI within the PAC must be changed to 100 ms or greater to support copper inter-switch links.
- Only fiber inter-switch links were tested as part of this CVD.

REP Technology Overview

REP is a technology implemented on Cisco Distribution switches and Cisco and Rockwell Automation IES. This software enhancement for Cisco Distribution and Cisco and Rockwell Automation IES extends network resiliency across Cell/Area Zone LAN designs. Requiring no hardware upgrades, REP is designed to provide fast network and application convergence in case of a media or network failure, without a negative impact on most network applications.


REP is a segment protocol that integrates easily into existing CPwE Cell/Area Zone LANs. It does not replace Spanning Tree Protocol (STP), but can coexist as part of the same Cell/Area Zone LAN. Since REP can also notify the STP about potential topology changes, it allows for interoperability with Spanning Tree. REP can be positioned as a migration strategy from legacy-spanning tree domains.

REP is a distributed and secure control plane protocol, it does not rely on a master node controlling the status of the ring. Hence, failures can be detected locally, either through loss of signal (LOS) or loss of connectivity to a neighboring switch. By default, REP elects an alternate port (the switch port being blocked). Any REP port within the REP topology can initiate a switchover to unblock the alternate port.

REP Terminology

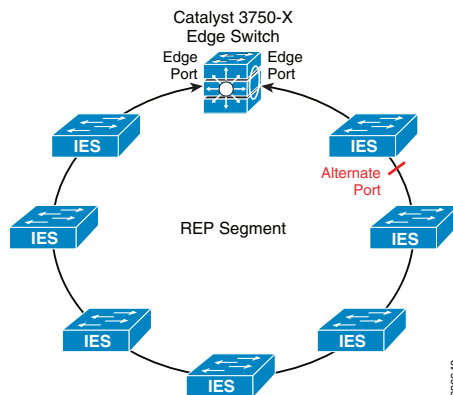
A REP segment is a chain of switch ports connected to each other and configured with the same segment ID. Each end of a segment terminates on what is called the “edge port” of an edge switch. [Figure 7](#) shows a REP segment. This basic element makes REP extremely flexible because you can plug this REP segment into an existing ring topology.

Figure 7 A REP Segment



[Figure 8](#) shows how REP wraps into a ring topology. Note that each node in the segment has exactly two REP-enabled ports.

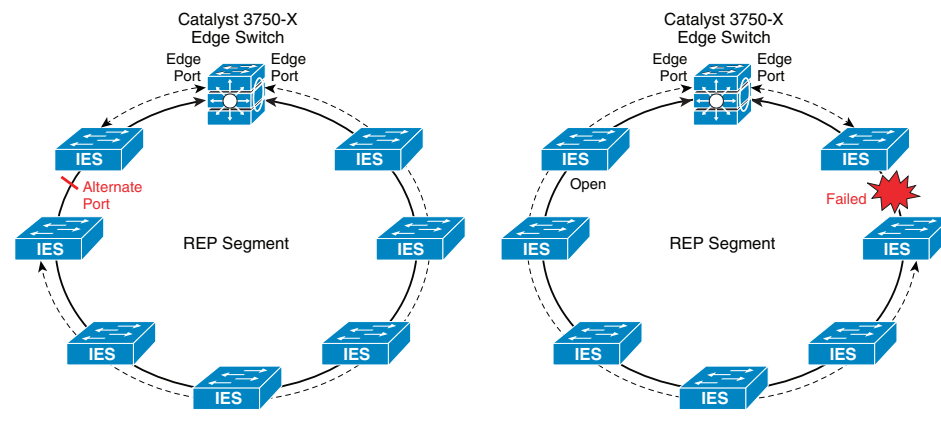
Figure 8 REP Ring Topology



REP Operation

With REP, in order to prevent a loop in the network, one switch port (the alternate port) is always blocked in any given segment. 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. Therefore, when a failure occurs in the segment, REP opens the alternate port so traffic can reach the edge of the segment (Figure 9).

Figure 9 REP Basic Operation



REP Fault Detection

REP, which relies primarily on loss of signal (LOS) to detect fiber link failure, can always learn the location of the failure within the ring. When a failure occurs, the failed ports immediately send link failure notifications to all REP peers. The failure notification has two purposes:

- Instruct the alternate port to unblock immediately because the segment is broken.
- Flush MAC table entries on all switches within the REP segment.

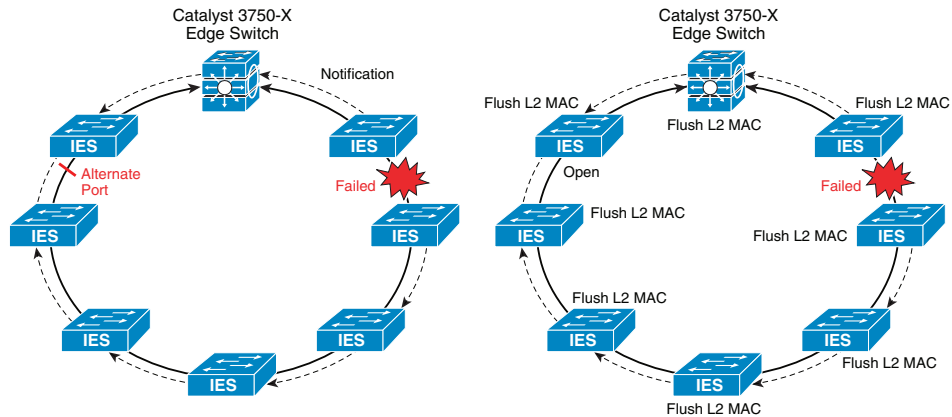
A REP node maintains neighbor adjacencies and continuously exchanges hello packets with its neighbors. In scenarios where LOS is not detected, the loss of a REP adjacency also triggers a switchover. Neighbor adjacency awareness is unique to REP and has advantages over alternate polling mechanisms that require centralized management from a master node. Note that the Unidirectional Link Detection Protocol (UDLD) can be enabled on REP interfaces to detect unidirectional failures, and this is enabled by default with IES Express Setup.

REP Failure Notification

Fast failure notification is critical for accomplishing fast convergence for IACS application. To ensure reliable and fast notification, REP propagates the notifications using the following two methods:

- **Fast Notification**—Using a Multicast MAC address, the notification is forwarded in hardware so that each node in the segment is notified immediately without software involvement from any node.
- **Reliable Notification**—Distributed through the REP Adjacency Protocol if lost REP retransmits the notification. The protocol uses sequence numbering and relies on packet acknowledgment. Upon receiving the notification, each REP node flushes MAC address entries learned on these REP ports and the alternate port then begins forwarding traffic. Because REP sends the notification through a reserved multicast address, the MAC addresses flushing can proceed in parallel on each REP node (Figure 10).

Figure 10 REP Link Fault Notification



REP Distributed and Secure

REP is a distributed and secure control plane protocol that does not rely on a master node monitoring the health of the ring. REP provides an additional layer of security, which protects the reliability and availability of the REP segment with the use of a 9-byte word generated by the alternate port and that is unique to that REP segment. The primary edge port is responsible only for initiating topology collection. Failure can be detected locally either through LOS or loss of neighbor adjacency. Any REP port can initiate a switchover as long as it has acquired a secure key to unblock the alternate port.

The secure key consists of a 9-byte length word that identifies each port. It is a combination of the port ID and a random number generated when the port activates. The alternate port key is secure because it is distributed only within a specific segment.


The REP alternate port generates and distributes its key to all other ports within the segment (Figure 11). Each port on the segment can use that key to unblock the alternate port. With this mechanism, users or attackers cannot unblock the alternate port unless they learn the key. This mechanism protects against potential security attacks; it also avoids problems with overlapping segment IDs.



Note

With 1024 segment IDs available, overlapping most likely will not occur, but misconfiguration could lead to such a scenario.

Figure 11 Alternate Port Key Distribution



Ease of Configuration and Management

REP configuration requires very few steps, and every switch in the segment is aware of the topology. The toolset includes a command line interface (CLI) and Device Manager.

The CLI topology reporting function displays the current topology as shown in [Figure 12](#).

Figure 12 REP Topology Command

```
3750-1 # show rep topology
REP Segment 1
BridgeName      PortName  Edge Role
-----
C3750X-1        Gi1/1/1   Open
IES-1           Gi1/1     Open
IES-1           Gi1/2     Open
IES-2           Gi1/1     Open
IES-2           Gi1/2     Open
IES-3           Gi1/1     Open
IES-3           Gi1/2     Open
IES-4           Gi1/1     Alt
IES-4           Gi1/2     Open
C3750X-1        Gi1/1/2   Open
```

A REP MIB is also available for SNMP management purposes.

To review the REP topology for one or all the network segments using Device Manager, choose Monitor > REP Topology from the Device Manager Web interface.

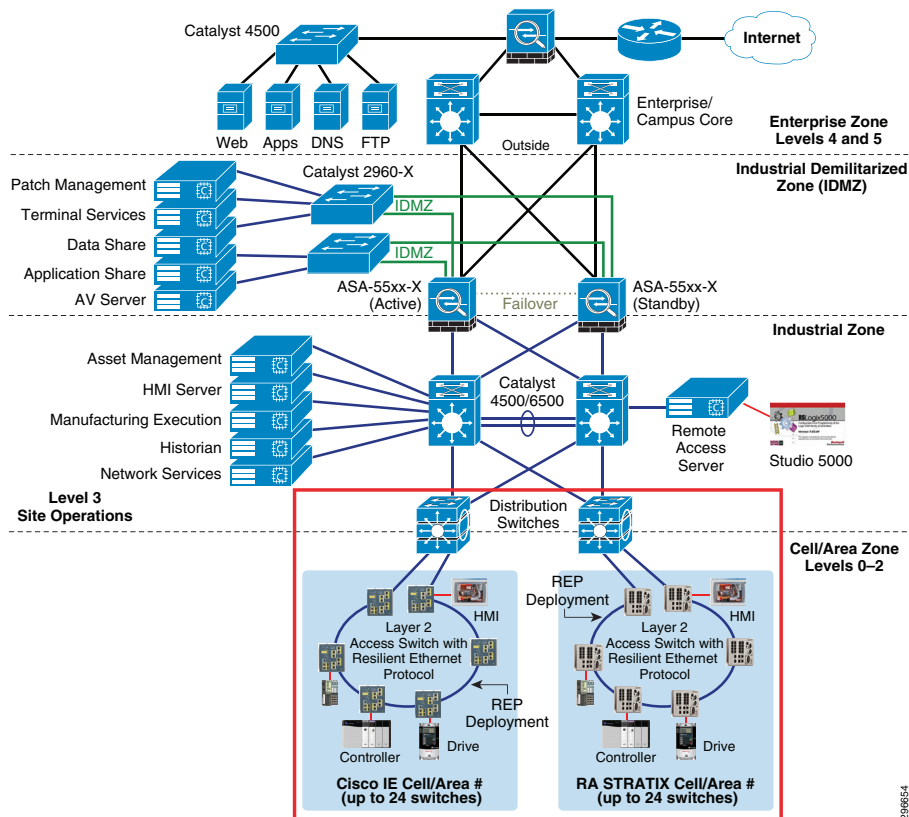
Figure 13 REP Deployment Location

Switch Name	Port Name	Edge	State
cz-3750	Gi2/0/8	Primary	Open
IES-8	Gi1/2	Transit	Open
IES-8	Gi1/1	Transit	Open
IES-7	Gi1/2	Transit	Alternate
IES-7	Gi1/1	Transit	Open
IES-6	Gi1/2	Transit	Open
IES-6	Gi1/1	Transit	Open
IES-5	Gi1/2	Transit	Open
IES-5	Gi1/1	Transit	Open
IES-4	Gi1/2	Transit	Open
IES-4	Gi1/1	Transit	Open
IES-3	Gi1/2	Transit	Open
IES-3	Gi1/1	Transit	Open
IES-2	Gi1/2	Transit	Open
IES-2	Gi1/1	Transit	Open
IES-1	Gi1/2	Transit	Open
IES-1	Gi1/1	Transit	Open
cz-3750	Gi1/0/1	Secondary	Open

Applying REP in IACS Applications

In the CPwE architecture, REP is targeted specifically within the Cell/Area Zone in the network architecture. As displayed in Figure 14, REP is attached to the Cell/Area Zone and Level 3 distribution switch directly, which in this architecture is the Catalyst 3750-X.

Figure 14 REP Deployment Location



298654

REP Design Considerations

REP is a Cisco protocol that provides an alternative to STP to control network loops and handle link failures, and to improve convergence time significantly. 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:

- **Faster Convergence**—REP can provide faster convergence because it runs on a physical link and not on a per-VLAN basis; only one hello message is required for all VLANs, reducing the load on the protocol.
 - 1 Gbps fiber inter-switch links are recommended to provide optimum convergence in REP topologies
 - Single REP segment with single CIP VLAN (1 CIP VLAN per segment)
 - Unicast traffic only

- Configuration of REP Admin VLAN
- Tested and validated up to 24 switches per segment (not including distribution switches)
- Up to 2 3750-X distribution / aggregation switches in a stack
- Up to 200 MACs (IACS devices) per VLAN
- Usage of the default Logix Requested Packet Intervals (RPI) is recommended - connection timeout values too close to the REP convergence time should be avoided
- **Link Integrity**—REP does not use an end-to-end polling function between edge ports to verify link integrity. It implements local link failure detection.

**Note**

The fiber media recommendation is based on the testing and validation of REP by Cisco and Rockwell Automation, with the Allen-Bradley programmable automation controllers (PACs) configured with the default requested packet interval (RPI) value of 20 ms for standard I/O.

In many MCC-based applications, device and I/O RPI settings as fast as the default 20 ms are not required. For applications that can perform appropriately with device and I/O RPI settings of 100 ms or greater, a Fast Ethernet (100 Mbps) copper inter-switch link may provide sufficient convergence in a REP topology.

- The default RPI within the PAC must be changed to 100 ms or greater to support copper inter-switch links.
- Only fiber inter-switch links were tested as part of this CVD.

Network and Ring Size Considerations

When determining where and how to deploy REP segments in an IACS architecture, consideration must be given to the number of devices and/or switches attached to the REP segment, the number of VLANs configured within the REP segment, and the number of MAC addresses that will be utilized in the REP segment. The combination of factors affect the recovery time of a REP segment during failover.

Another important factor to consider when designing a REP ring is latency. Since latency is a cumulative process based on the number of switches and/or other devices between the start and end points of a packet, limiting the size of the ring also becomes extremely important if latency is a determining factor in the design. Every device that the packet must pass through, such as a switch, adds latency to the data path.

The number of MAC addresses in an industrial Ethernet switch's Layer 2 forwarding table affects the performance of the REP segment. Since all packet lookups are performed in hardware, the number of addresses does affect the ultimate recovery time of a REP segment.

- When a failure occurs in the REP segment that causes a topology change, all of the IES on that specific REP segment **MUST** flush their Layer 2 Forwarding tables.
- When REP re-convergence of the segment has completed, **ALL** of the switches **MUST** re-learn all of the MAC Addresses and populate their respective Layer 2 forwarding tables.

During the learning process, all of the traffic in a layer 2 REP segment is flooded out of all ports until learning is complete, thus affecting all devices on the REP segment.

When the above characteristics are considered as part of the whole REP solution, the considered tolerances for network sizing include the following:

- Up to 24 IES per segment (not including the distribution switch)

- Up to 2 3750-X distribution / aggregation switches in a stack
- Up to 200 MACs (IACS devices) per VLAN
- 1 VLAN per segment

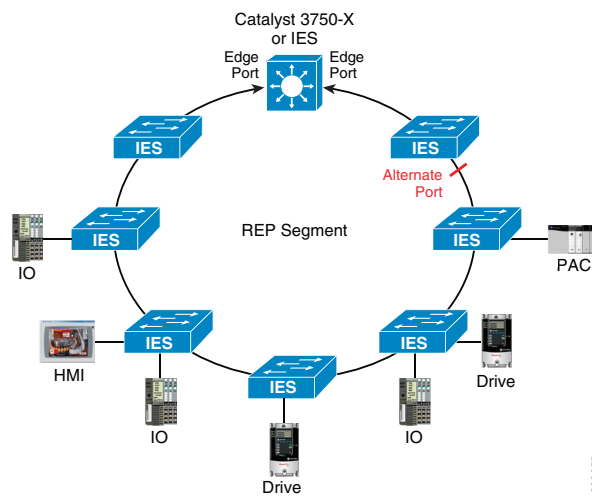
Distribution Switch

When deciding upon the distribution switch configuration to use with the REP architecture, several approaches should be considered. Most importantly, the distribution switch that sits between the Cell/Area Zone and Level 3 Site Operations within the CPwE architecture. This switch, which can be deployed as a single switch, or stacked switches, provides full support for both Layer 2 and Layer 3 and for hardware and network redundancy. This switch can also provide connectivity to multiple Cell/Area Zones.

Single Distribution Switch

In an REP design where distribution switch redundancy is not a requirement and the fiber ring begins and ends in the same location, utilization of a single distribution switch is the preferred design. [Figure 15](#) illustrates this design.

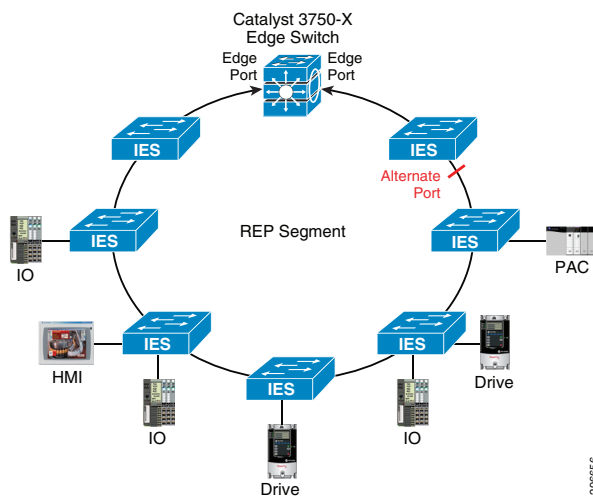
Figure 15 Single Distribution Switch



Stacked Distribution Switch


In some installations where switch redundancy is required and the fiber ring begins and ends in the same location, using stackable switches is the preferred solution. This architecture also provides full Layer 2 and Layer 3 coupled with hardware redundancy, as displayed in [Figure 16](#).

Figure 16 Layer 2/Layer 3 Distribution Stackable Switch



The Catalyst 3750-X is able to support redundant power supplies and redundant power sourcing (east-west grid power) by using Cisco StackPower. Each switch is only required to be configured with a single power supply and a StackPower cable between the two switches, as displayed in Figure 17.

Figure 17 Distribution Stack Switch Cable Configuration



Single Switch versus Stacked Switches

In a single switch configuration, REP segments are terminated on the same physical switch. While the single switch can provide power redundancy as an option, it does not protect against an aggregation/distribution switch failure.

In a stacked configuration, the Catalyst 3750-X switches overcome the hardware redundancy shortfall of the single switch by creating a single logical switch by interconnecting multiple switches using a stack cable (Figure 17).

Switch stacking simplifies configuration since all members in the stack are controlled by a single configuration. By having only a single configuration to manage, a reduction or elimination of configuration errors is achieved. Also, a switch stack appears as a single point of management in the network architecture: again, exactly the same as a single switch configuration.

The next advantage of stacking is that the customer can achieve both Layer 2 Stateful Switch Over (SSO) and Layer 3 Nonstop Forwarding (NSF) without complicated configurations. Another advantage of stack switching over single switch is the ability to terminate REP segments on different

member switches within the stack, thereby eliminating single points of failure. No segment should be terminated on the same physical switch, since this would eliminate physical layer termination protection in the instance of a switch failure.

When utilizing stacking with the Catalyst 3750-X, a customer can leverage the StackPower capability where two switches in a stack can utilize a single power supply each, while providing power redundancy for each other. If a power supply were to fail, the remaining power supply can provide sufficient power for both switches to operate normally. StackPower is only supported by Catalyst 3750-X switches operating in stack mode configuration.

Table 3 REP Feature Comparison

	Centralized Edge Ports	Distributed Edge Ports	Ring Redundancy
Single Switch	X		
Stacked Switch	X	X (different stack members)	X

Distribution Switch Conclusion

As shown, the distribution switch can be deployed in numerous ways depending upon the particular network requirements and cable/fiber plant limitations. [Table 4](#) provides a quick reference of basic feature functionality.

Table 4 Distribution Options

	Central Termination of Ring	Distributed Termination of Ring	Full Layer-2 Redundancy	Full Layer-3 Redundancy	Power Redundancy	Stack Power
Single Switch	X				X (Optional)	
Stacked Switch	X	X	X	X	X (Optional)	X (Optional)

Single Segment versus Multiple Segments

REP can be configured in various scalable styles of ring architectures. These architectures include:

- Single or multiple segments with a single VLAN per segment

A segment is a ring with its own unique identifier. Since all REP ports are trunk ports, multiple VLANs can traverse a single segments.


Trunk Port

A trunk is a point-to-point link between Ethernet switches. Ethernet trunks carry the traffic of multiple VLANs over a single link and by default is a member of all VLANs in the VLAN database.

Access Port

An access port belongs to and carries the traffic of only one VLAN. Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port.

Figure 18 Multiple Segments with Single VLAN per Segment



Configuring the Infrastructure

This section describes how to configure REP and related features in the CPwE-REP system based on the design considerations of the previous section. During the testing effort, we have validated the configurations, which are based on Cisco external and internal documentation. This section covers the following topics:

- Configuring the REP ring
- IES access switch considerations
- Distribution switch considerations (stacked and non-stacked)

Configuring REP Ring

This section describes the basic configurations necessary to implement REP in a ring-based Cell/Area Zone LAN. It is assumed that Express Setup and other Smart Port macro configurations for IES have already been applied, so the details of those configurations are not covered in this document (refer to previous CPwE design guides for these details). This section covers the following topics:

- 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.
- You cannot run REP and STP or REP and Flex Links on the same segment or interface.
- You must configure all trunk ports in the segment with the same set of allowed VLANs, or a misconfiguration occurs.
- Native VLAN implementation for REP control messages.

- REP administrative VLAN implementation for fast failure notifications.
- REP segment and edge configuration.
- 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. Configuration tips (use of console, or direct Ethernet connection to IES).
- REP is supported on EtherChannels, but not on an individual port that belongs to an EtherChannel.

Native VLAN Implementation

REP uses the native VLAN configured on the trunk interfaces of a network segment to establish and maintain connectivity across the segment, as well as reliably informing all nodes of any topology changes using Link Status Layer (LSL) frames. This behavior is similar to other Layer 2 control plane protocols such as Cisco Discovery Protocol (CDP) and VLAN Trunking Protocol (VTP). Best practices for configuring the native VLAN on the trunk interfaces include the following:

- The native VLAN on a trunk is 1 by default. For security purposes, select a different VLAN to configure as the native VLAN.
- When selecting the native VLAN, use a VLAN that is separate from the one carrying IACS traffic to prevent any interaction between the two domains.
- When pruning unused VLANs from the trunk, be sure to include the native VLAN (along with the IACS VLAN) as allowed.

If the native VLAN has not already been configured on the uplink ports using a Smart Port macro (Device Manager or CLI), it can be configured using the following command in interface configuration mode:

```
switchport trunk native vlan <VALUE>
```

In addition, ensure that the VLAN has been added to the global database using the following commands in global configuration mode:

```
vlan <VALUE>
name Native_VLAN
```

Admin VLAN for REP

In addition to the reliable notifications sent on the native VLAN after a topology change, REP also uses Hardware Flood Layer (HFL) notifications that are immediately sent out as multicast frames by the switch hardware. Because these frames are hardware switched by each device in the path, rather than relayed hop-by-hop, they can be received across the segment very quickly. This behavior allows REP to converge quickly following a failure and limit IACS device timeouts for many applications.

A REP administrative VLAN is configured globally on each switch within a segment to control the VLAN onto which the HFL frames are forwarded. In addition, since HFL frames are flooded as data traffic only on ports belonging to that VLAN, the scope of this traffic can be confined to the Cell/Area Zone LAN. Best practices for configuring the REP administrative VLAN include the following:

- As with the native VLAN, for security purposes change the REP administrative VLAN (via CLI or Device Manager) to a different value from its default of 1. Similarly, do not choose the VLAN carrying IACS traffic.

- Be sure to include the administrative VLAN as allowed when pruning unused VLANs from the trunk.

**Note**

If the administrative VLAN is not allowed across the entire REP ring, both within and outside the segment, the HFL frames will be dropped and network convergence will be dependent on the slower LSL mechanism. While Link Status Layer (LSL) frames are considered control traffic and are therefore relayed across the trunk regardless of pruning, HFL frames are considered data traffic and must be explicitly allowed across the trunk.

- Since the administrative VLAN has similar constraints to the native VLAN, it makes sense to assign the two as the same VLAN. In addition, most Cell/Area Zones will be separated by Layer 3 (distribution switch) domains, so constraining the HFL flooding does not need to be a significant consideration.

To configure the REP administrative VLAN, use the following command in global configuration mode:

```
rep admin vlan <VALUE>
```

In addition, ensure that the VLAN has been added to the global database using the following commands in global configuration mode:

```
vlan <VALUE>
name REP_Admin_VLAN
```

REP Segment Configuration

REP is configured on both IES and distribution switches simply by enabling it on each interface that will be part of the segment and including a segment ID to identify to which segment the port belongs. At each end of the segment, primary and secondary edge ports are configured. The purpose of the primary edge port is to initiate topology discovery and communicate special configurations for the segment. The secondary edge port has no special function beyond terminating the segment.

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

```
rep segment <ID>
```

To configure a port as an edge port (typically on a distribution switch), 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
```

If neither edge port has this designation, REP will elect one as the primary edge based on which has the best port ID.

Stack or Standalone

The 3750-X distribution switches can be made redundant through stacking or, if redundancy is not required, a standalone 3750-X may be used. Both of these scenarios were tested as part of the CPwE-REP test effort.

In a stacked 3750-X configuration, each of the edge ports for the REP segment should be placed on a different switch to provide redundancy in case one of the switches fails. The following configuration snippet provides an example of this setup:

```
interface GigabitEthernet1/1/1
  rep segment 10 edge primary
interface GigabitEthernet2/1/1
  rep segment 10 edge
```

Apart from this consideration, all REP-related configuration remains the same whether the distribution switch is a single node or stacked.

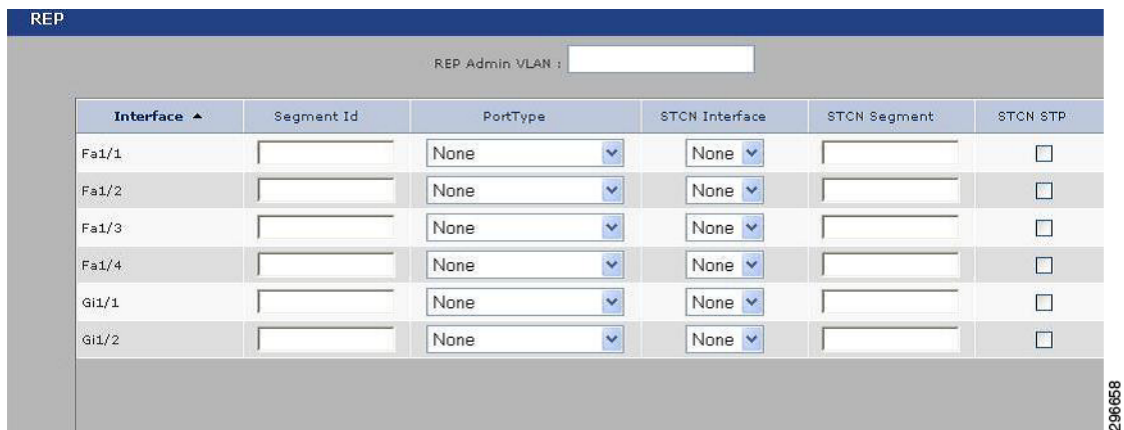
Device Manager

Device Manager provides a graphical user interface to configure REP (vs. CLI). This section shows the general steps for using Device Manager to configure REP. The actual configuration information is the same as with CLI as shown previously.

Use the REP window to configure REP segments.

- Step 1** To create a REP segment, set a segment ID and port type on the desired ports.
- Step 2** To display this window, choose Configure > REP from the Device Manager Web interface.

Figure 19 Reference Topology



- Step 3** Enter settings defined in [Table 5](#).

296658

Table 5 Reference Settings

Setting	Description
REP Admin VLAN	The administrative VLAN. The range is 2 -4094. The default is 1, and it should be changed as described previously. REP ports are assigned to the same REP Admin VLAN. If the REP Admin VLAN changes, all REP ports are automatically assigned to the REP Admin VLAN.
Interface	The number of the switch port, including port type (such as Fa for Fast Ethernet and Gi for Gigabit Ethernet), the base switch number (1), and the specific port number. For example: Fa1/1 is Fast Ethernet port 1 on the base switch.
Segment ID	The ID of the segment. The segment ID range is from 1-1024. If no segment ID is set, REP is disabled.
Port Type	The REP port type of the port can be: Primary, Edge, Transit, No-neighbor Primary, No-neighbor, and None. The default is None. Following are Port Type definitions: Primary —This port is a primary edge port. Edge —This port is a secondary edge port. Transit —This port is a non-edge port in the REP segment. No-Neighbor Primary —This port is a primary edge port connected a non-REP switch. No-Neighbor —This port is a secondary edge port connected to a non-REP switch. None —This port is not part of the REP segment.

Test Objectives and Setup

This section describes how the CPwE-REP solution, after being designed and configured as mentioned in the previous sections, was validated. This section covers the following topics:

- Test Objectives
- Use Cases
 - [Figure 3, “Single REP Switch ring, All IES Solution, Single VLAN per REP Segment”](#)
 - [Figure 4, “Multiple REP Switch Rings, All IES Solution, Single VLAN per REP Segment”](#)
 - [Figure 5, “Multiple REP Switch Rings, IES and Distribution Switch \(non-stacked\) Solution, Single VLAN per REP Segment”](#)
 - [Figure 6, “Multiple REP Switch Rings, IES and Distribution Switch \(Stacked\) Solution, Single VLAN per REP Segment”](#)
- Test Setup

Test Objectives

The major objectives of the CPwE-REP test effort were to validate the solution architecture and characterize its performance in a Cell/Area Zone LAN. The test setup was designed to ensure that the Cell/Area Zone LAN functioned as expected under a variety of conditions, as well as providing design and implementation guidance that could be utilized for real-world CPwE deployments. Testing focused around the following key network characteristics:

- Switch ring redundant path topology with REP resiliency protocol
- Performance - convergence time
- Scalability (e.g., number of IACS devices and IES)
- Adaptability (e.g., segment topology)

Aspects of the network not covered by testing include the following:

- IACS applications or devices themselves
- Adverse environmental conditions (temperature, shock, etc.)
- Full performance characterization of network infrastructure (IACS applications generally do not approach these thresholds in actual deployments)

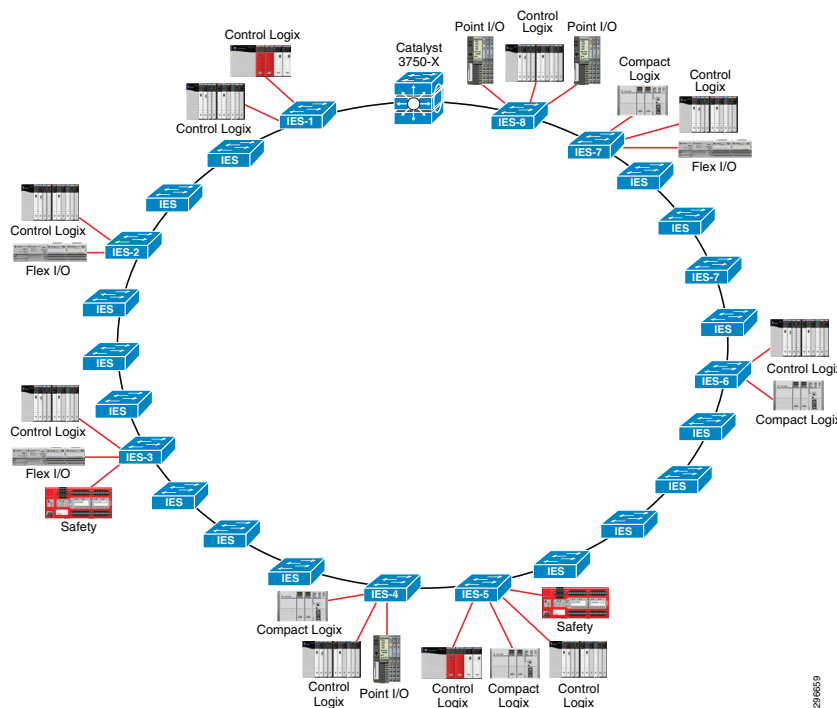
Lab Setup

The test network topology, platforms, and configurations used in the validation are described next, as well as the IACS devices, and test tools used to evaluate system performance.

Test Topology

This testing effort focused on switch ring topologies, consisting of IES access switch (with connected IACS devices) and terminated by a set of 3750-X distribution switches, with REP providing failure resiliency. Figure 20 shows the reference topology, which represents the base upon which the other topology variations for the REP testing were built. Switch name/number and IACS device connections remained the same, even with different topologies and other variables.

Figure 20 Reference Topology



Test Platforms and Software Versions

Table 6 shows the platforms tested, along with the validated software versions.

Table 6 Platforms and Software

Product/Platform	Software Release	Role
Stratix 5700 or IE 2000	15.0(2)EA1	Access
Stratix 8000 or IE 3000	15.0(2)EY3	Access
Catalyst 3750-X	15.0(2)SE6	Distribution

Test Configurations

Each IES in the ring was configured using standard industrial Smart Port macros and other recommendations given in the CPwE CVD. They were configured with VLAN 10 as the Cell/Area Zone VLAN for IACS device traffic and VLAN 900 as native VLAN for the uplink/trunk connections. Each Stratix IES was also configured with an IP address on the Cell/Area Zone VLAN (as shown in the reference topology) for access via the common industrial protocol (CIP). This supports Stratix IES manageability by the Rockwell Automation's Studio 5000® software.

Whether in a stacked or standalone configuration, the 3750-X distribution switch was configured with the following functions:

- Default gateway for all IES access switches (and connected devices)
- Primary and secondary REP edge ports
- Layer 2/3 demarcation (ports connected to the ring are Layer 2)
- Routing for traffic between Cell/Area Zone LAB and other areas within the Industrial Zone (e.g., HMI)

IACS Setup

The IACS devices were configured to monitor the status of the I/O connections in the system. The I/O connections were setup to ensure that IACS traffic is flowing across the failure points in the test. This ensures that IACS traffic is in the path of the network disruption. These traffic flows are used to measure the impact of a topology change on the IACS traffic.

IACS Single VLAN Application

The IACS Single VLAN test uses an updated version of the test software used for CPwE DIG. The test application records the impact of a topology change on both standard and safety I/O traffic. The standard I/O traffic used the default RPI of 20ms. The safety I/O had a connection reaction time limit of 120ms. Both the standard and safety I/O tests use a mix of I/O adaptors and produce/consume traffic between controllers.

The IACS application employs:

- 14 controllers, each with a dedicated network interface card (NIC)
- 6 I/O modules

The test application monitors and logs faults in the controller I/O tree. The data is summarized and any faults are reported.

Test Tools

An Ixia test chassis, running IxNetwork software, was used to generate network traffic and measure network performance. Traffic flows were designed to create worst-case failure scenarios by originating and terminating their traffic on opposite sides of a failure point. These Ixia connection points were located on the IES-4, IES-5, IES-7, and IES-8 switches.

Each Ixia port sent and received UDP unicast streams with a variety of MAC addresses. Each port sent packets of varying sizes at a specific rate that was high enough to ensure accurate convergence calculations without overwhelming the link capacity. Though recent Ixia versions now perform these calculations automatically, it should be noted that the value is determined manually using the following formula:

$$\text{Convergence} = [(\text{Tx packets} - \text{Rx packets}) / \text{Packet rate}] * 1000 \text{ ms/s}$$

Test Setup

Test cases for the CPwE-REP test effort were generally divided based on topology (single or multiple rings) and the number of VLANs carrying IACS traffic (one per ring). The test effort focused on new devices and features in the network. Each test scenario covered a subset of the key requirements, executed a separate set of test cases, and measured distinct characteristics of the network based on that scenario.

Table 7 shows a summary of the number and type of test cases that were executed for each topology and VLAN quantity variation. Link and Node Disruptions are forced events to enable measuring convergence time of the network.

Table 7 Test Cases

Topology	Description	Resiliency Method	VLANs per Ring	Link Disruption Locations on Ring	Link Recovery Locations on Ring	Node Disruption Locations on Ring	Node Recovery Locations on Ring	Number of Test Iterations
REP1	Single segment Single VLAN	Stack	1	4	4	2	2	10
REP2	Multiple segment Single VLAN per segment	Stack	1	5	5	2	2	5

Constant parameters for each test included the following:

- **Resiliency Protocol**—REP was used for Layer 2 resiliency for all tests.
- **Number of Ring Switches**—24 switches, excluding the distribution switch, were used in each test. For multiple ring scenarios, the nodes (and IACS devices) were divided equally between rings.
- **MAC Scale**—The system was scaled to 200 MAC addresses.
- **Uplink Medium**—Fiber links were used in the ring for all tests.



Note

Copper inter-switch links were used in some preliminary testing, but they produced convergence values greater than those acceptable by IACS applications, with the default PAC RPI of 20 ms. Therefore, only fiber inter-switch links were tested.

Figure 21 and Figure 22 show the arrangement of each test topology, as well as each failure point.

Figure 21 REP1 Topology with Failure Locations


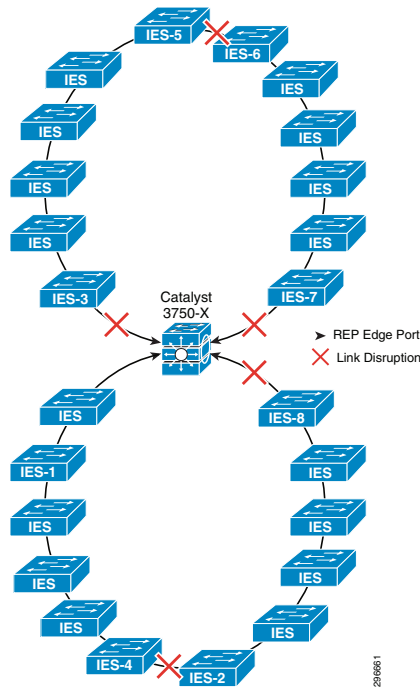


Figure 22 REP2 Topology with Failure Locations



REP Troubleshooting Tips

REP has two basic commands that can be used to troubleshoot any problems with an incomplete segment:

- show rep topology
- show interfaces rep

The first command, “show rep topology”, gives an overall view of the segment, including the locations of the primary and secondary edge ports and alternate (blocking) port. It shows all ports that belong to the segment in a linear fashion, which helps to pinpoint which device and port might be causing an issue. Typical output for a fully functional segment looks like the following:

```

IES-13#show rep topology
REP Segment 10
BridgeName      PortName      Edge Role
-----
D3750X          Gi1/1/1      Pri  Open
IES-11          Gi1/1         Open
IES-11          Gi1/2         Open
IES-10          Gi1/2         Open
IES-10          Gi1/1         Open
IES-12          Gi1/1         Open
IES-12          Gi1/2         Open
IES-13          Gi1/2         Open
IES-13          Gi1/1         Alt
IES-14          Gi1/1         Open
IES-14          Gi1/2         Open
IES-15          Gi1/2         Open

```

More detailed information about port status and identifiers can be found by adding “detail” to the command, as shown in the following output:

```

IES-13#show rep topology detail
REP Segment 10
D3750X, Gi1/1/1 (Primary Edge)
  Open Port, all vlans forwarding
  Bridge MAC: 0007.7d5c.6300
  Port Number: 019
  Port Priority: 000
  Neighbor Number: 1 / [-50]
IES-11, Gi1/1 (Intermediate)
  Open Port, all vlans forwarding
  Bridge MAC: 4c00.8254.de80
  Port Number: 001
  Port Priority: 000
  Neighbor Number: 2 / [-49]
IES-11, Gi1/2 (Intermediate)
  Open Port, all vlans forwarding
  Bridge MAC: 4c00.8254.de80
  Port Number: 002
  Port Priority: 000
  Neighbor Number: 3 / [-48]
<output omitted>

```

Finally, by adding “archive” to the command, the output that would have resulted before the last event (e.g., a failure) within the segment is displayed.

A more detailed view of REP-enabled ports on a particular switch within the segment is provided by the “show interfaces rep” command. Typical output for a switch with two REP-enabled uplinks is shown below:

```

IES-13#show interfaces rep
Interface      Seg-id Type      LinkOp      Role
-----
GigabitEthernet1/1      10      TWO_WAY      Alt
GigabitEthernet1/2      10      TWO_WAY      Open

```

Most of the fields are self-explanatory, but the LinkOp field indicates whether a full REP adjacency has been formed with the device connected to that port. When the port is first configured for REP, it will begin in a WAIT state. Next, it will send a Hello packet to the neighbor and change its state to ONE_WAY. If the adjacency fails, the port will likely remain in either this or another failed state (e.g., NO_NEIGHBOR). Reasons for a failed adjacency could include the opposite port not being configured for REP, REP traffic not being allowed on the trunk, or the REP process failing on the connected switch. Once a full adjacency is established, the state is changed to TWO_WAY.

Once again, adding “detail” to the command will give a much more detailed view of the REP port characteristics, as shown below:

```

IES-13#show interfaces rep detail
GigabitEthernet1/1  REP enabled
Segment-id: 10 (Segment)
PortID: 0001F84F575EBA00
Preferred flag: No
Operational Link Status: TWO_WAYf
Current Key: 0001F84F575EBA0011DD
Port Role: Alternate
Blocked VLAN: 1-4094
Admin-vlan: 900
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: 1563198, tx: 1830473
HFL PDU rx: 1139, tx: 948
BPA TLV rx: 551026, tx: 1078849
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 22649, tx: 25342
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 422937, tx: 422832

GigabitEthernet1/2  REP enabled
Segment-id: 10 (Segment)
PortID: 0002F84F575EBA00
Preferred flag: No
Operational Link Status: TWO_WAY
Current Key: 0001F84F575EBA0011DD
Port Role: Open
Blocked VLAN: <empty>
Admin-vlan: 900
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: 1330531, tx: 2110526
HFL PDU rx: 1087, tx: 0
BPA TLV rx: 28423, tx: 1601021
BPA (STCN, LSL) TLV rx: 0, tx: 0
BPA (STCN, HFL) TLV rx: 0, tx: 0
EPA-ELECTION TLV rx: 32022, tx: 22649
EPA-COMMAND TLV rx: 0, tx: 0
EPA-INFO TLV rx: 429606, tx: 429756

```

Significant fields from this output include:

- **PortID**—The full REP port identifier, formed by appending the port priority and port number to the bridge MAC address (these values can be seen in the output of “show rep topology detail”).

- **Current Key**—Indicates the key for the current alternate port in the segment. All segment ports should have synchronized keys.
- **Blocked VLAN**—Any VLANs blocked by this port for load balancing purposes.
- **Admin-VLAN**—Configured REP administrative VLAN.
- Statistics for LSL and HFL packets, as well as other REP-related messaging.

Useful debug commands for troubleshooting REP issues include the following:

- **Debug Rep Failure-Recovery**—Shows failure detection and HFL/LSL packets sent to inform the segment of the failure.

Cisco is the worldwide leader in networking that transforms how people connect, communicate and collaborate. Information about Cisco can be found at www.cisco.com. For ongoing news, please go to <http://newsroom.cisco.com>. Cisco equipment in Europe is supplied by Cisco Systems International BV, a wholly owned subsidiary of Cisco Systems, Inc.

www.cisco.com

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND 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.

© 2015 Cisco Systems, Inc. All rights reserved.

Rockwell Automation is a leading provider of power, control and information solutions that enable customers to get products to market faster, reduce their total cost of ownership, better utilize plant assets, and minimize risks in their manufacturing environments.

www.rockwellautomation.com

Americas:
Rockwell Automation
1201 South Second Street
Milwaukee, WI 53204-2496 USA
Tel: (1) 414.382.2000, Fax: (1) 414.382.4444

Asia Pacific:
Rockwell Automation
Level 14, Core F, Cyberport 3
100 Cyberport Road, Hong Kong
Tel: (852) 2887 4788, Fax: (852) 2508 1846

Europe/Middle East/Africa:
Rockwell Automation
Vorstlaan/Boulevard du Souverain 36
1170 Brussels, Belgium
Tel: (32) 2 663 0600, Fax: (32) 2 663 0640

Integrated Architecture, Stratix 8000, Stratix 5700, ControlLogix, FLEX I/O, POINT I/O, CompactLogix, Studio 5000, RSLogix 5000 are trademarks of Rockwell Automation, Inc. Trademarks not belonging to Rockwell Automation are property of their respective companies.