



# Connected Communities Infrastructure – Rail Solution Design Guide

Passenger rail operators must ensure the safety of workers, passengers, and the public while concurrently keeping trains running on schedule, providing superior service, and minimizing operational costs. To achieve these business objectives, rail equipment and infrastructure owners can take advantage of increasingly advanced technology to address safety concerns, improve asset visibility, offer new value-added services, increase ridership, and reduce operating expenses.

Freight rail carriers use remote control, automation, and new Internet of Things (IoT) data to keep workers safe, streamline operations, and minimize freight cost per mile, so they remain competitive with other modes of overland freight transportation.

Passenger rail providers can offer services such as high-speed Internet access, infotainment, mobile ticketing, and security systems on trains and in stations to enhance the passenger experience and increase ridership. The critical signaling and control systems such as European Railway Traffic Management System (ERTMS), Communication-based train control (CBTC), and positive train control (PTC) are needed for safe operation.

All these capabilities can be key to business success for rail operations, and they all depend on a foundation of network capability that is secure, scalable, reliable, and resilient, and that satisfies an ever-increasing demand for data throughput. Cisco networking products and solutions provide the necessary foundation for securely solving these networking challenges. Cisco Industrial IoT products are proven to meet the unique demands of operating in a rail environment, and the Cisco Connected Rail solution provides a Cisco Validated Design for reliably interconnecting high-speed trains, trackside infrastructure, stations, and operations centers across all the rail operator's sites and regions.

The Cisco Connected Rail solution validates the architecture for high-speed, robust wireless connectivity between train and trackside as well as resilient, scalable access and backhaul transport infrastructure to interconnect wayside, stations, and operations centers across the rail operator regions. As listed in [Table 1](#), rail operators and Cisco solution partners use Connected Rail as a secure foundation on which to build their solutions to support use cases including passenger Wi-Fi, infotainment, video surveillance and analytics, operations management, maintenance, signaling, and control.

**Table 1 Cisco Connected Rail customer use cases and business outcomes**

Use case category	Definition	Business outcomes
<b>Safety and compliance</b>	?Connectivity for train control systems such as Communication-Based Train Control (CBTC), European Railway Traffic Management System (ERTMS), and Positive Train Control (PTC) ?Video surveillance for passenger safety at stations and onboard trains ?Faster emergency response time ?Real-time remote asset monitoring	?Improved passenger safety and security ?Improved situational awareness ?Reduced risk and more auditable compliance ?Optimized passenger operation ?Improved customer experience
<b>Passenger experience</b>	?Fast and reliable passenger Wi-Fi ?On-demand and bandwidth-efficient infotainment services ?Passenger interactive kiosks and wayfinding ?Real-time Estimated Time of Arrival (ETA)	?Increased ridership ?Improved customer service ?Revenue generation and new business models
<b>Operational efficiency</b>	?Fully resilient trackside infrastructure to support signaling and control applications ?High-speed train-to-ground wireless infrastructure to support service delivery even when a train moves at high speed ?Predictive maintenance and remote condition monitoring enabled through sensor data collection and correlation ?High-definition IP video surveillance for live monitoring of activity on trains, in stations, and at trackside ?Improved train headway and increased track utilization ?Support for better voice communication between train operators and passengers (VoIP)	?Higher utilization of track infrastructure ?Lowered OpEx and reduced CapEx ?Simplified maintenance ?Enhanced operational efficiency
<b>Data management</b>	?Reliable passenger Wi-Fi with seamless roaming between onboard and station locations ?User data management and targeted advertisements ?Ability to link current mode of transit to additional modes for seamless experience ?Video analytics for passenger counting and behavior monitoring	?Improved passenger experience ?Increased ridership ?Better customer service ?New revenue resource

The train-to-ground communication includes three major components:

- Onboard communication component

## Solution System Level Overview

This is the communication system on the train. It typically uses wireless and/or cellular technology to communicate between the train and ground network.

The onboard network, which includes network and safety equipment within the train, is out of scope for this guide.

### ■ Trackside communication component

This is the ground-based network to provide cellular and/or wireless coverage alongside the train track to communicate with the train.

- For cellular communication, it relies on cellular coverage along the train track.
- For a dedicated wireless trackside communication network, wireless radios are set up along the trackside to communicate with the wireless components on the train.

These trackside radios connect to the CCI network at an extended node or policy extended node within an Edge PoP. A cellular-based train to trackside solution is out of scope for this guide.

### ■ Network Backhaul and Datacenter component:

Each Edge Point of Presence (PoP) is connected to the datacenter/HQ PoP through an IP or SDA transit. Equipment in the datacenter provides services required to complete the end-to-end communication for the train.


This guide describes the design for enabling communication to the train, but not the services within the train. The challenge is to provide high bandwidth, low latency communication to a high-speed moving train.

Cisco Ultra Reliable Wireless Backhaul (CURWB) radios and technology are used. They are well suited to the rail environment, providing up to 500Mbps at speeds up to 225 MPH (360 KMH). The design and integration of the Cisco Ultra Reliable Wireless Backhaul (formerly Fluidmesh) technology within the CCI network is the central focus of this guide.

## Solution System Level Overview

The diagram below depicts a high-level system view of the Rail Solution in the CCI infrastructure.

Figure 1 CCI Rail Solution System Level Architecture



## Connected Trains

The train infrastructure consists of an onboard network and a train-to-trackside radio network. The train-to-trackside radio network connects the services supported on the train with systems and services in the centralized infrastructure.

For the dedicated train-to-trackside wireless communication, there is a CURWB radio on the train which communicates with the CURWB trackside radio. It supports high-speed seamless roaming between trackside radios while providing high throughput and low latency.

## Connected Trackside

The CCI network spreads across a large geographical area, logically divided into several Points of Presence (PoPs). Each Edge PoP has one or more Access Rings comprised of extended or policy extended node IE switches (maximum 30) in a Resilient Ethernet Protocol (REP) ring. The IE switch models include IE3300, IE3400, IE 4000, and IE 5000. Refer to the “Centralized WLC Deployment” section in the [CCI General Solution Design Guide](#) for more detail.

The CURWB trackside devices connect to the IE switches in the PoPs within a trackside virtual network (VN). A group of trackside radios in the same IP subnet forms a Trackside Radio Group (TRG). These trackside radio groups can span one or more Access Rings in the Edge PoP.

## Connected Station

A station network design must provide various passenger services and maintain the safety and security of the train station and its passengers. The network also must scale to meet the demand of the number of passengers at the station during peak hours.

## CURWB Overview

The station network can be an Edge PoP or connected to an Edge PoP in the CCI environment and provides connectivity to devices such as train schedule bulletin boards, ticketing kiosks, surveillance cameras, and passenger devices/mobiles via wired or wireless connections. These devices are either directly connected to the IE switches in the Access Ring or, more generally, connected to the Wireless Access Points using Wi-Fi technology. Refer to the “CCI Access Network” section in the [CCI General Solution Design Guide](#) for a complete list of Access Points supported in CCI to make your Access Point selection choices in the station network.

A Centralized Wireless LAN Controller (WLC) deployment model is recommended for Station Network. The Centralized WLC deployment model features a pair of WLCs in the Shared Services to serve wireless Access Points across all PoPs. This system is able to scale more efficiently to support the aggregated number of passengers during peak hours. Passengers typically enter the train station, travel from one station to another, then exit the train station. The WLC choice is dependent on the number of wireless users expected during peak travel times which is typically the morning and evening rush hours.

The Central WLC support is available in CCI infrastructure. The specific solution option for Station Network is documented in the “Centralized WLC Deployment” section of the [CCI General Solution Design Guide](#).

## Backhaul

The Rail solution on the CCI infrastructure supports two types of backhaul due to the low latency and high throughput requirements:

- transparent backhaul where traffic resides entirely within the SDA fabric using an SDA Transit (for example, routed over a private or dark fiber)
- opaque backhaul where traffic exits the SDA fabric domain to an IP transit network (for example, a Service Provider or private MPLS network) and returns to the SDA fabric at the other side of the transit network

Refer to the “Backhaul for Points of Presence” section [CCI General Solution Design Guide](#) for more information. Both types of backhaul are applicable for rail environment.

## Centralized Infrastructure

This is the area encompassing the CCI data center or headquarters PoP and the shared services. The servers and services supporting the trackside end-to-end solution reside here. This includes the CURWB gateway devices necessary to support the seamless roaming from train to trackside.

## CURWB Overview

The CURWB product line provides numerous benefits and features in the outdoor IoT space. They can provide network connectivity in a fixed infrastructure environment with point-to-point, point-to-multipoint, or bridge mode. They can also provide seamless network connectivity in a mobile environment such as a moving train.

What enables this flexible wireless deployment is a customized MPLS implementation that ensures an unbroken communication path which overcomes the limits of standard wireless protocols. This implementation acts as an overlay on the CCI network. It enables data throughput of up to 500Mbps at up to 225MPH (360 KMH) with optimal wireless conditions in a mobile environment. In a fixed infrastructure, it enables a flexible and resilient point-to-multipoint mesh network.

MPLS relies on label identifiers, rather than the network destination address as in traditional IP routing, to determine the sequence of nodes traversed to reach the end of the path. An MPLS-enabled device is also called a Label Switched Router (LSR). A sequence of LSR nodes configured to deliver packets from the ingress to the egress using label switching is denoted as a Label Switched Path (LSP), or “tunnel”. LSRs on the border of an MPLS-enabled network and / or other traditional IP-based devices are also called a Label Edge Router (LER).

## Solution Components

The following components are used as part of the wireless infrastructure, in addition to the components which are already part of the main CCI infrastructure.

**Table 2 Cisco Ultra Reliable Wireless Backhaul Solution Components**

Hardware	Software/ Firmware	Component Details / Role
FM 3500	9.3	Mobility: Trackside radio and Mesh Point / Mesh End Fixed Infrastructure: Mesh Point / Mesh End or Bridge
FM Tube Antenna	N/A	Trackside radio directional antenna, for rail trackside and tunnel deployment
FM Panel Antenna	N/A	Trackside radio antenna, for rail trackside deployment
FM 1000	1.3.1	Mobility: Mesh End / Global Gateway (no radio functions)
FM 10000	2.0.1	Mobility: Global Gateway
Configurator		Web-based interface built into each device to configure, monitor, and troubleshoot the radio network
RACER		Cloud-based tool for configuration building and firmware upgrading
FM-MONITOR		On-premises based monitoring and statistics tool
FM 4500	9.3	Mobility: Train Radio*

(\*) Train Radio is not part of the trackside infrastructure. The FM 4500 resides on the train to communicate with the FM 3500 on the trackside

Below is a brief description of terminologies frequently referred to in this document:

**Train radio** - The physical radio onboard the train that connects the Onboard Network (OBN) within the train to the trackside infrastructure. This is the demarcation point between the Ultra Reliable Wireless Backhaul network and the train network. The radio will impose an MPLS label on packets coming in from the train network or remove the label when packets are moving to the train network. A single train will typically have one or more train radios to communicate with the trackside infrastructure, for example at the front and at the back of the train.

**Trackside radio** - The physical radio installed along the trackside that communicates with the train radio and other trackside CURWB devices. It can operate as a Mesh Point, a Mesh Point Wireless Relay, or a Mesh End. A group of trackside radios with a common Mesh End will be called a Train Radio Group (TRG).

**Mesh Point** - A Mesh Point primarily serves to swap MPLS labels as traffic ingresses and egresses. This means all Mesh Points function as an LSR and act as a relay between the train radio or a host device and a Mesh End. When a Mesh Point is disconnected to the wired network, it is operating in infrastructure mode. A Mesh Point can also operate in wireless only mode to act as a wireless relay.

**Mesh End** - Whether used for mobility or fixed infrastructure, the Mesh End performs the same basic functionality. It is the logical demarcation point between Mesh Points and the L3 IP network. Using the MPLS terminology described before, all Mesh Ends function as LSRs and LERs. A Mesh End must have a wired connection and it must be in the same L2 broadcast domain as the Mesh Points.

Based on which version of Fluidity (L2 or L3) is being used, the Mesh End serves different purposes. In both versions, the Train Radio Group (which is a group of Mesh Points) communicates with the L3 IP network through the Mesh End. In a Layer 3 Fluidity network a Mesh End also serves to terminate L2TP tunnels connected to a Mesh End gateway

## CURWB Overview

in the datacenter. The traffic from Mesh Points enters the Mesh End and is then forwarded to the datacenter Mesh End through these L2TP tunnels. When traffic is received from the datacenter Mesh End, it is removed from the L2TP tunnels and forwarded to the train through the Mesh Points.

In a fixed infrastructure, the Mesh End enables communication between the hosts or switch behind the Mesh Point and the rest of the L3 IP network.

**Global Gateway** - A global gateway is a special type of Mesh End that enables seamless roaming between different Layer 3 domains. It resides in the datacenter as described above. A global gateway serves to anchor numerous Mesh Ends in different broadcast domains and provide seamless roaming across them. This is achieved by building L2TP tunnels between the Global Gateway and all Mesh End devices.

This fast MPLS label swapping between the above nodes along with L2TP tunnels between the Mesh Ends and Global Gateway enable seamless roaming at high speed and high throughput.

A Global Gateway is not used in a fixed infrastructure environment.

**Plug-ins** - CURWB features are dependent on software licenses called Plug-ins. There are plug-ins for maximum throughput, security, and other network features. The high availability feature, called TITAN and explained later in this document, also requires the appropriate plug-in.

## Mesh Point and Mesh End

In a TRG, there are Mesh Points and Mesh Ends and they must be in the same L2 broadcast domain.

The trackside radios are deployed along the rail track. For maximum performance, it is recommended to connect the mesh points wired to the IE switches in the Edge PoP access ring. Given the proper IE switch configuration, the mesh points can be powered through PoE.

When traffic enters the train radio, it will impose an MPLS label for that radio. As the train moves along the train track, the train radio associates and disassociates with the trackside radios (Mesh Point) along the track based on the radio coverage. As the train radio roams, it will change the MPLS label based on which trackside radio it is associated with. When the trackside Mesh Point receives this traffic, it will perform the function of a Label Switch Router (LSR) and do a lookup in its MPLS label table for the Mesh End and swap the labels. It will then send the packet onto the network with a destination address of the Mesh End.

The Mesh End unit functions as a Label Edge Router (LER) as well as an L2TP tunnel endpoint. When a packet arrives from a mesh point, the mesh end will add an L2TP header pointing to the Global Gateway in the datacenter PoP. It will then forward this packet to the Global Gateway. When L2TP traffic is received from the Global Gateway, it will remove the L2TP header and forward to the correct Mesh Point in the LSP. The FM 3500 radio can perform the role of a Mesh End or Mesh Point. When operating as a Mesh End, it can also process wireless traffic from the train radios.

A FM 3500 is suitable to serve as a Mesh End if the expected traffic does not exceed 500 Mbps. The FM 1000 is the recommended Mesh End unit when the traffic will not exceed 1 Gbps.

## Global Gateway

The Global Gateway enables seamless roaming when a train radio roams between different Train Radio Groups. This is necessary because the LSP from a mesh point terminates at the mesh end and does not go beyond. The Global Gateway overcomes this by using L2TP tunnels to every Mesh End in the Edge PoPs. The MPLS labeled packets from the Mesh Ends are encapsulated in these L2TP tunnels and the Global Gateway performs its role as a Mesh End by removing these headers and labels before forwarding these packets onto the CCI network. In the CCI network, the Global Gateway resides in the data center PoP in a Virtual Network for train and trackside communication. Because the Global Gateway is the head end for the radio network, all return traffic destined for the train must use the Global Gateway as the next hop. This traffic is then encapsulated with an MPLS label and L2TP header which is then forwarded to the appropriate Mesh End.

Both the FM 1000 and FM 10000 can perform as a Global Gateway. The selection criteria for choosing between them is based on the bandwidth requirements. The FM 1000 can process throughput of up to 1Gbps while the FM 10000 handles up to 10 Gbps.


## Fluidity

Fluidity is the technology that enables seamless roaming between a train radio and the trackside radio network. In this context, seamless roaming occurs when there is no disruption in the communication path as the train radio associates and disassociates with trackside radios.

Fluidity operates over a flat Layer 2 network or a routed Layer 3 network. In Layer 2 Fluidity, all trackside roaming occurs within a single subnet or broadcast domain. Layer 3 Fluidity supports roaming between L3 domains. As an Edge PoP is based on a Layer 3 network, it is required to deploy Layer 3 Fluidity to enable roaming when the train moves from one TRG (IP subnet) to another.

The diagram below depicts a Layer 3 Fluidity network to summarize the nodes and their placement in the network.

**Figure 2 Layer 3 Fluidity Network**



## Product Compliance and Physical Deployment

Products designed to support the Rail industry are routinely exposed to harsh conditions. When deploying a rail solution within the CCI context, care must be taken when choosing where to locate devices. Table 3 summarizes the location options given the temperature and mounting options.



**Table 3 Product Installation Summary**

Product	Installation Location Options	Temperature Rating	Note
FM 3500	Trackside	-40C to 75C	No M12 conn
FM 4500	Train/Trackside	-40C to 80C	M12 conn or Fiber
FM 1000	Datacenter or Field Center	-20C to 55C	Rack/VESA/DIN Rail/Wall mount
FM 10000	Datacenter	0C to 40C	Rack mount

Below is the Cisco Ultra Reliable Wireless Backhaul (CURWB) product compliance matrix for more detailed information.

**Table 4 CURWB Product Compliance Matrix**

	FM 4500	FM 3500	FM 1000	FM 10000
EN50155	✓			
EN45545	✓			
EN50121-4	✓	✓		
EN61373	✓			
EN55011				✓
EN55022			✓	✓
EN55024			✓	
EN55032			✓	
EN61000-3-2			✓	✓
EN61000-3-3			✓	✓
EN61000-6-2				✓
EN61000-6-4				✓

For a more in-depth look at how CURWB with mobility integrates into the CCI network, refer to [CURWB and CCI Network Integration and Considerations, page 11](#).

## Trackside Network Design and Considerations


Required trackside radio spacing is based on achieving the RSSI radio coverage for the targeted data rates. Typical distance is approximately 800 meters (0.5 mile) apart. Whenever possible, deploy radios along both sides of tracks in a zig-zag pattern for best coverage.

The radio placement can be a single radio or dual radios per pole as shown in the diagram below. The signal for single radio per pole will be split between two MIMO antennas, which results in a shorter trackside radio placement interval due to the RF power reduction. A typical splitter RF loss is -3 dB, resulting in half power. However, with a single radio per pole the handoff does not occur as the train passes the pole (unlike the two radios per pole option). While this is a cost saving measure, depending on the site survey, throughput requirements, and pole placement, a single radio per pole may not provide enough coverage.

CURWB System Components

The dual radios per pole configuration increases the allowable distance between poles for the same coverage requirement as shown in [Figure 3](#).

**Figure 3 Single Radio vs Dual Radio**



The dual radios per pole configuration enables multi-frequency support that allows multiple channels to be used to improve aggregate network throughput.

A dual radio deployment is recommended for better coverage over longer distances with more options in selecting frequencies.

As the typical distance between radios is approximately 800 meters (0.5 mile), an Access Ring with 30 IE switches covers approximately 15 miles in a linear setup along trackside. To cover the desired area, one can expand the nodes in a train radio group (TRG) into multiple Access Rings. Each TRG must be able to support the aggregate throughput desired for the train.

## CURWB System Components

### Management and Monitoring

For management and monitoring of the CURWB network, please refer to the Network Provisioning and Monitoring section of the [CCI General Solution Design Guide](#).

### Resiliency

TITAN is a software feature for failover technology that constantly tracks link status and network performance of a pair of Mesh Ends or Global Gateways configured in an active-standby role. In case of any failure of the primary unit, traffic is rerouted to the redundant secondary unit. The pair is configured with a single virtual IP address to appear as one unit.

Under the TITAN configuration, the pair of devices will fall into a primary or secondary role (based on the unit Mesh ID) and issue keepalives between them in a preconfigured interval (typically between 50 ms and 200 ms). The secondary unit becomes the new primary when it has not received a keepalive message within the predefined interval.

Simultaneously, the new primary issues commands to all other wireless devices in the domain to inform them of the change while updating its own tables and sending gratuitous Address Resolution Protocol requests (ARPs) out its Ethernet port to ensure new traffic will be forwarded properly to the new primary. This feature allows failure detection and recovery within 500ms.

When TITAN is configured on the Mesh Ends and Global Gateways, each device must be configured with two L2TP tunnels. Each Mesh End unit (Primary and Secondary) establishes a L2TP tunnel to each Global Gateway (Primary Global Gateway and Secondary Global Gateway with TITAN).

## CURWB and CCI Network Integration and Considerations

This is the expected result after configuring all the tunnels. Only one tunnel is in connected state and the other 3 are in idle state:

- L2TP tunnel between primary Global Gateway and primary Mesh End: *CONN*
- L2TP tunnel between primary Global Gateway and secondary Mesh End: *IDLE*
- L2TP tunnel between secondary Global Gateway and primary Mesh End: *IDLE*
- L2TP tunnel between secondary Global Gateway and secondary Mesh End: *IDLE*

If the primary Global Gateway fails, the L2TP tunnels between the primary Global Gateway and primary Mesh End become *IDLE*. The secondary Global Gateway will become the new elected primary and the L2TP tunnel between secondary Global Gateway to the primary Mesh End will become *CONN*.

Similarly, at the trackside network level, if the primary Mesh End fails, the L2TP tunnels between it and the primary Global Gateway will become *IDLE*. The L2TP tunnels between the secondary Mesh End (which will be elected the new primary) and the primary Global Gateway will become *CONN*.

Using TITAN on all Mesh End pairs and Global Gateway pairs is recommended.

## Quality of Service

Please see the “CURWB Quality of Service” section in the [CCI General Solution Design Guide](#).

# CURWB and CCI Network Integration and Considerations

## Cisco DNAC

### Virtual Network and Segmentation

When integrating CURWB into the CCI Network, it is recommended to set it up as a separate service in a Virtual Network dedicated to train to trackside communication. Alternatively, a virtual network with other shared train or station related services (such as signage, ticketing, etc.) can be used. Within the virtual network, micro-segmentation can then be used to prevent CURWB devices from communicating with these station devices. Note that in an Edge PoP, only Policy Extended Nodes support micro-segmentation.

Refer to the “Overlay Network” section in the [CCI General Solution Design Guide](#) for macro and micro segmentation information.

### IP Pool

The CURWB devices will only use a DHCP address if they are able to reach the RACER portal through the Internet. Otherwise, they must be statically addressed. Additionally, when a Mesh End is configured for L2TP and TITAN it requires more addresses. There will be one IP address for the interface, one IP address for the L2TP tunnel, and then a virtual IP address that is also configured on the secondary Mesh End. All of these addresses come out of the IP scope allocated to the Train Radio Group. If the Virtual Network is dedicated to the CURWB devices and RACER will not be used in Online mode, DHCP should be disabled for that Train Radio Group. Otherwise, the DHCP scope should be configured to exclude the number of addresses needed.

When using Layer 3 Fluidity, the IP addressing for the train radios and devices behind those radios is not related to or part of the IP addressing used for the trackside communication VN. It is recommended to create IP Pools for the trackside radios and gateway devices aboard the train as an administrative task.

## Host Onboarding

The CURWB devices do not support 802.1X, therefore MAB authentication is the only other option for secure onboarding. Once authenticated through MAB, the device can operate on the CCI network.

Refer to the “Extended Nodes and Policy Extended Nodes” section in the [CCI General Solution Design Guide](#) for more information.

## Datacenter PoP

As mentioned in the section on IP Pools, the IP addressing for the train radios and gateway devices is not part of the trackside communication VN. This means there must be an explicit route added for these networks with the Global Gateway as the next hop. This will ensure that all return traffic destined for the train will enter the Global Gateway and be tunneled to the appropriate Mesh End.

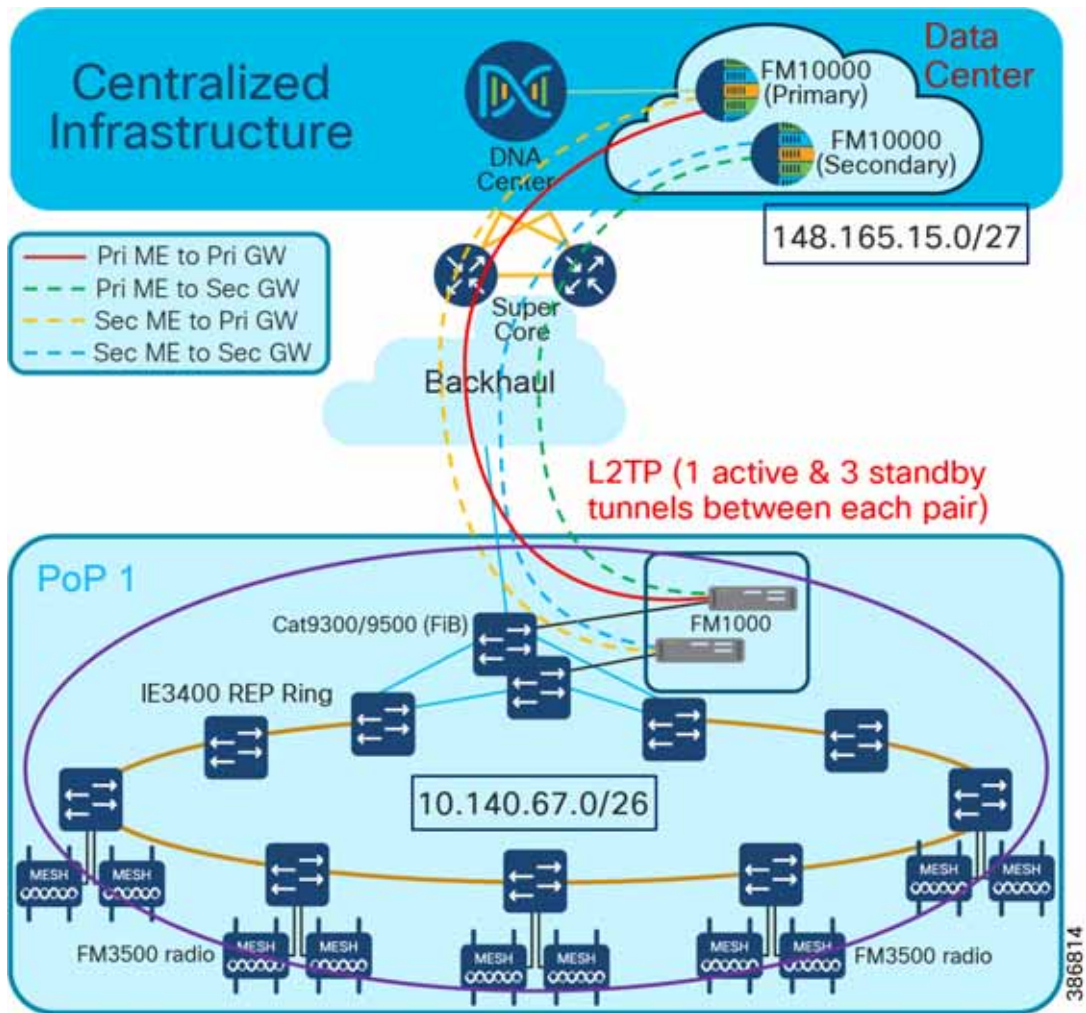
## Edge PoP

As discussed previously, Layer 3 Fluidity supports multiple Train Radio Groups, where each group is in a different IP subnet. There are multiple considerations and recommendations when planning this deployment.

- Each Mesh Point should be connected to a PoE capable IE switch in an access ring
- If the number of CURWB devices in the Train Radio Group exceeds the maximum number of nodes in an access ring, they can span across multiple access rings if that subnet is present in those rings. This should be balanced with the expected throughput in that Train Radio Group.
- The Redundant Mesh Ends should be connected to the Fabric in a box switch stack on different members to eliminate single points of failure. If the Mesh Ends are FM 3500s, they should be connected to different access ring switches.


See [Figure 4](#) for an example of a PoP with a pair of FM 1000s covering the entire PoP.

Figure 4 FM 1000 as Mesh End for a Single Ring in the Entire PoP



See Figure 5 for an example of a pair of FM 1000s covering two access rings within a PoP

Figure 5 Two FM 1000s as the Mesh End for two Access Rings within a PoP



See Figure 6 for an example of a pair of FM 1000s covering separate access rings. The standby links are not shown to improve clarity.

Figure 6 Edge PoPs with FM 1000 dedicated to a single Access Ring

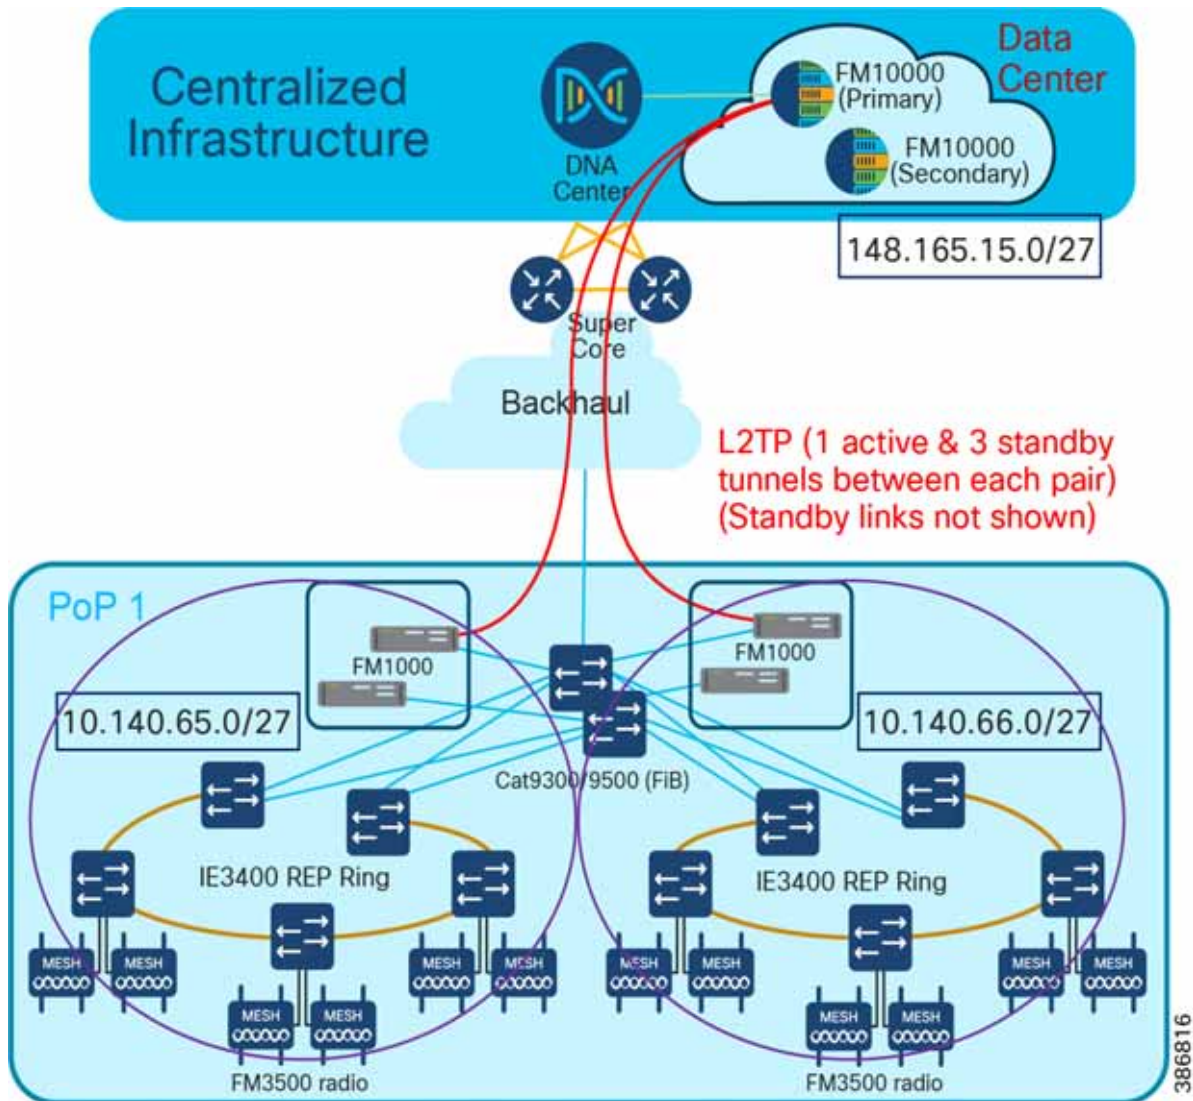
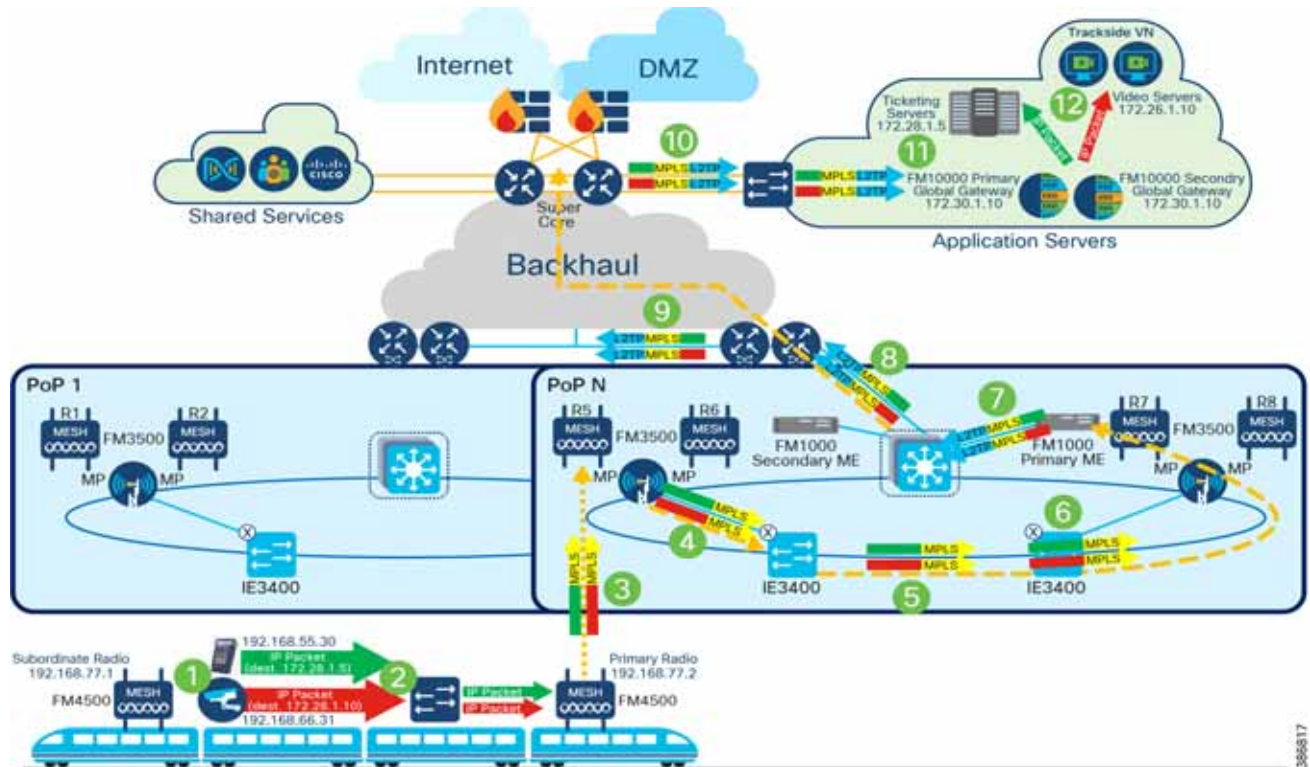


Figure 7 shows the sequence of MPLS tag handling and L2TP encapsulation events after the CURWB devices have been integrated into the CCI network.

Figure 7 CURWB L3 Fluidity and CCI Integration



The summary of the sequence of events follows.

1. IP devices on the train send packets to their destinations.
2. The packets are switched to the Train Radio (FM 4500) on the train.
3. The Train Radio adds MPLS tags to the packets, selects the best trackside radio and sends the packets over the wireless network to R5 on the trackside.
4. R5 on the trackside receives the packets. Since it is a Mesh Point, it will send the data toward the Mesh End. It looks up the destination in the label lookup table, swaps the label for the next Mesh End, and sends it out.
5. The packets are forwarded to the next IE switch.
6. The IE switch continues forwarding the packets to the Mesh End.
7. The FM 1000, which is the Primary Mesh End, encapsulates the packets into an L2TP tunnel connected to the Global Gateway and forwards to the Edge PoP Fabric-in-a-box.
8. The packets are forwarded from the Edge Pop Fabric-in-a-box to the Transit network.
9. The Transit network forwards the packets to the Datacenter PoP.
10. The packets arrive at the Datacenter PoP to be placed in the Trackside VN.
11. The packets are forwarded to the Primary Global Gateway.
12. The Global Gateway removes the L2TP and MPLS headers and forwards the original data packets to their destination.



## End-to-End QoS Integration

As described in the CURWB QoS section, when Mesh Points exchange packets with other Mesh Points and the Mesh End, they have an MPLS label where the EXP bits are set based on the inner payload DSCP. The IE switches are unable to match on the MPLS EXP bits and therefore within the access ring, a different QoS strategy is required. Using Day N templates, a MAC ACL configured on each switchport connected to the CURWB devices allows it to match on the MPLS ethertype of 0x8847. Using the MPLS ethertype means the templates can be generalized for each switch and not rely on knowing the device MAC address.

With this information, QoS can be performed based on those ACLs. Note that this does not allow differentiated service for the different types of traffic coming from the train, but rather all train traffic can be marked with a configurable level of service within the Edge PoP access ring.

When the packets reach the Mesh End, a L2TP header is attached with the DSCP value based on the inner payload. The Fabric in a box will then process the data according to the QoS policies.

