# Rail Communications-Based Train Control (CBTC) and Safety

## Implementation Guide

April 2024


Cisco Systems, Inc. www.cisco.com

# Contents

# Rail CBTC and Safety Implementation Guide

## Introduction

According to the UITP World Metro Figures 2021 report, approximately 3,300 km of new rail infrastructure was put in service between the start of 2018 and the end of 2020. During this timeframe, operational fleets worldwide increased by 28,000 vehicles to a total of 140,000 vehicles. In 2019, an average of 190 million passengers per day were taking the metro globally. Rail operators are constantly striving to keep their trains moving safely, providing superior and reliable services to the riders, and lowering their operational cost. A modern railway signaling system called "Communications-based train control (CBTC)" was introduced in the mid-1980s with the objective to achieve maximum capacity while maintaining the safety requirements.

On October 23, 2023, the Transportation Security Administration (TSA) renewed cybersecurity security directive – Enhancing Rail Cybersecurity -SD 1580/82-2022-01 to regulate passenger and freight railroad carriers through the implementation of layered cybersecurity measures, with the goal to reduce the risk that cybersecurity threats pose to critical railroad operations and infrastructures.

This implementation guide breaks down the overall rail CBTC and safety solution into sections based on the location and role in the overall network. This includes the core, backbone, wayside access, wayside wired, onboard train, and network management.

This implementation is based on the design recommendations in Rail Communications-Based Train Control (CBTC) and Safety Design Guide.

## Topology

The topology shown in the following figure represents the architecture used in the Rail CBTC and Safety solution. Sections in this document describe each part of the architecture in detail.

**Figure 1 Rail CBTC and Safety Architecture**

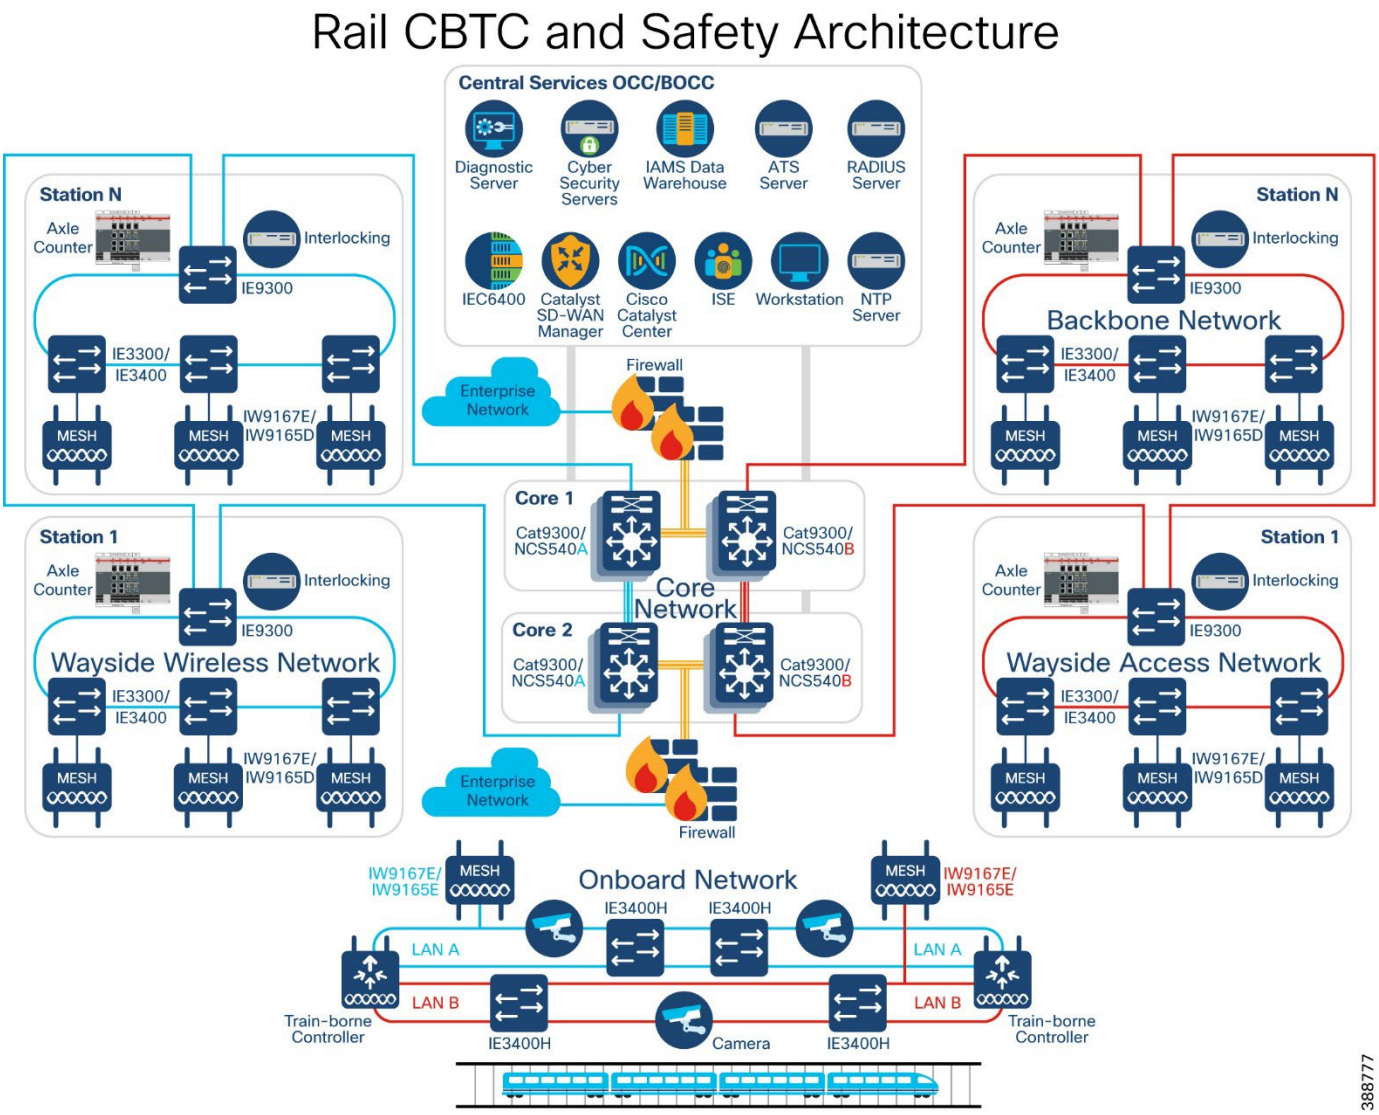


# Solution Components

This section describes the components of a CBTC and safety network. Several device models can be used at each layer of the network. The device models that are suitable for each role in the network and the corresponding CVD software versions are described in Solution Hardware and Software Compatibility.
The device model should be chosen based on specific deployment requirements such as network size, cabling and power options, and access requirements. Table 1 describes the device models that are used in this solution.

**Table 1 Components and Device Models in Rail CBTC and Safety Architecture**

| Component Role | Component | Description |
|---|---|---|
| **Wayside access switch** | Cisco Catalyst Industrial Ethernet (IE) 3400 Series Switch and/or<br>Cisco Catalyst Industrial Ethernet (IE) 3300 Series Switch | 1Gig REP access ring – IE3400<br>10Gig REP access ring – IE3300 |
| **Station Backbone switch** | Cisco Catalyst 9300 Series switch or Cisco Catalyst IE9300 Series switch | Connects to wayside REP access ring and core network. |
| **Core network switch – L3** | Cisco Catalyst 9300 Series switch | Connects to backbone switch and Operation Control Center (OCC)/Backup OCC (BOCC) |
| **Core network router – MPLS** | Cisco NCS540 | Connects to backbone switch and OCC/BOCC |
| **Next-Generation Firewall** | Firepower 2100 or 4100 Series | Next-generation firewall. |
| **OT network sensor** | Cisco Cyber Vision (CV) network sensor | CV network sensors on all IE switches in the wayside access, backbone, and core network switches. |
| **OT security dashboard** | Cisco Cyber Vision Center global and local virtual appliances | CVC deployed globally and locally in control center |
| **Ultra-reliable Wireless Backhaul (URWB) gateway** | IW9167E or IEC6400 Edge Compute Appliance | URWB wireless network mesh end. |
| **Network management** | Cisco Catalyst Center, Cisco Crosswork Network Controller | Network management application in datacenter. |
| **Authentication, authorization, and accounting (AAA)** | Cisco Identity Service Engine (ISE) | AAA and network policy administration. |
| **IT and OT security management** | Cisco Secure Network Analytics (Stealthwatch) Manager and Flow Collector Virtual Edition | Network flow analytics and security dashboard in control center. |
| **Train to wayside wireless (high performance)** | IW9167E | URWB Mesh radio on train and wayside |
| **Train to wayside wireless (lower cost)** | IW9165E (on train) / IW9165D (on wayside only) | URWB Mesh radio on train and wayside |
| **Onboard network** | Cisco Catalyst IE3400 Heavy Duty | Ruggedized network access switch |

# Solution Hardware and Software Compatibility

Table 2 lists the Cisco products and software versions that are validated in this CVD.

**Table 2 Cisco Hardware and Software Versions Validated in this CVD**

| Component Role | Hardware Model | Version |
|---|---|---|
| **Wayside network** | IE3400-8P2S, IE3300-8U2X | 17.13.1 |
| **Core network - L3** | C9300-24UX | 17.13.1 |
| **Core network - MPLS** | N540X-6Z18G-SYS-A | 7.9.2 |
| **Station backbone switch** | IE-9320-22S2C4X | 17.13.1 |
| **Network management application** | Cisco Catalyst Center Appliance<br>DN2-HW-APL | 2.3.7.0 |
| **Authentication, authorization, and accounting (AAA) server** | Cisco ISE Virtual Appliance | 3.2 |
| **URWB mesh point** | IW9167E, IW9165D, IW9165E | 17.13.1 |
| **URWB mesh end/ global gateway** | IW9167E or IEC6400 | 17.13.1 |
| **URWB IW Monitor** | IW Monitor VM | v2.0 |

# Network Management

This section includes the services needed to manage the rail network. These are shared services that are needed by all network sites but are not considered vital CBTC traffic that needs end-to-end duplication. Cisco Catalyst Center and ISE are specifically referred to in this implementation guide, but other services can be included such as DHCP, DNS, or other market specific applications.

## Cisco Catalyst Center

Cisco Catalyst Center is a dedicated hardware appliance that offers centralized management for designing, provisioning, and applying policies based on business intent for the entire network deployment. Cisco Catalyst Center does not support the NCS 540 for the MPLS core, but rather Cisco Crosswork Network Controller must be used. Other features include discovering and onboarding the physical network devices along with complete lifecycle management. This includes Day-0 templates for the initial onboarding and Day-N templates for extra configuration options. In the following sections, configuration snippets will be shown to enable certain features for the Catalyst 9300 and IE9300. It is recommended to configure these as part of Day-N templates to ensure reliable and automatable deployment of the features network wide.

For details about installing and configuring the hardware appliance, see Cisco DNA Center Second-Generation Appliance Installation Guide, Release 2.3.7.0 and 2.3.7.3:

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-3-7/install_guide/2ndgen/b_cisco_dna_center_install_guide_2_3_7_2ndGen.html

To configure specific features, see the Cisco DNA Center User Guide, Release 2.3.7.0 and 2.3.7.3:

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-3-7/user_guide/b_cisco_dna_center_ug_2_3_7.html

Other examples using Cisco Catalyst Center for onboarding and maintaining the network can be found here:

https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/Industrial_Automation/IA_Horizontal/IA_Networking/DNA_Center_IA_IG/DNA_Center_IA_IG.html

## Cisco ISE

Cisco ISE is a policy-based access control system that enables and enforces compliance and infrastructure security. It is the Authentication, Authorization, and Accounting (AAA) server that works with Catalyst Center for integrated device identity management, access control, and access policy enforcement.

With ISE and Catalyst Center working together, users and devices can be dynamically mapped to scalable groups to simplify the end-to-end security policy management. Macro and micro segmentation can be implemented at scale without relying on static access lists.

To install Cisco ISE, see the Cisco Identity Services Engine Installation Guide, Release 3.2:

https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/install_guide/b_ise_installationGuide32.html

After installing ISE, it must be integrated with Catalyst Center. For instructions, see “Cisco DNA Center and Cisco ISE Integration” in Cisco DNA Center Administrator Guide, Release 2.3.7.0 and 2.3.7.3.

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/2-3-7/admin_guide/b_cisco_dna_center_admin_guide_2_3_7/b_cisco_dna_center_admin_guide_2_3_7_chapter_010.html#id_54524

When completed, Catalyst Center and ISE will exchange information using PxGrid. The security policies can be defined and maintained using Catalyst Center which are then synced with ISE.

## Cisco Crosswork Network Controller

Cisco Crosswork Network Controller (CNC) is an integrated network automation solution for deploying and operating IP transport networks combining intent-based network automation with manual or automatic remediation. There are multiple components available in the full suite including Network Services Orchestrator (NSO), Segment Routing Path Computation Element (SR-PCE), and WAN Automation Engine (WAE).

In this solution, CNC is used to onboard and manage the NCS 540 along with the entire MPLS core. The installation and management of CNC is out of scope for this guide. See the Crosswork Network Controller Solution Workflow Guide for more details: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/crosswork-network-controller/6-0/Solution-Workflow-Guide/bk-crosswork-network-controller-6-0-solution-workflow-guide.html

## Industrial Wireless Service

Industrial Wireless Service (IW Service), an OT service in Cisco IoT Operations Dashboard, is used as the cloud-managed onboarding and management tool for the IW916x and IEC6400 products. It can operate in online or

offline mode depending on whether the devices have access to the Internet. In online mode, IW Service will compare the running configuration of the radio with the configuration in IW Service and inform the user if they are out of sync. IW Service can push the configuration from the tool as well as the latest firmware and keep all devices up to date. In offline mode, all the same configurations can be done but can't be pushed to the devices. The configuration must be downloaded and then added to the device using the local GUI configurator.

Detailed instructions on how to add new IW devices and manage current ones is here: https://developer.cisco.com/docs/iotod/industrial-wireless-overview-introduction/

The URWB devices can be managed individually or at a group level by using the **Groups** feature (https://developer.cisco.com/docs/iotod/create-groups/) . Placing similarly configured devices into groups reduces the time needed to configure the devices and reduces misconfigurations.

## Industrial Wireless Monitor

Industrial Wireless Monitor (IW Monitor) is an on-premises application that monitors the status of the wireless mesh devices and the mesh itself. It does not perform configuration or firmware management like IW Service but focuses on the telemetry from the IW devices to report link status, wireless mesh performance, device performance, and so on.

IW Monitor is distributed as a Docker container which enables it to run on a variety of computing platforms. The platform running the container must have IP reachability to all devices being monitored on port 6600. Only the IW916x products are supported as of this writing. The IEC6400 is not supported.

After installing the container and logging in for the first time, it will ask to add the available devices. This can also be performed any time in the future.

**Figure 2 IW Monitor Add Device**

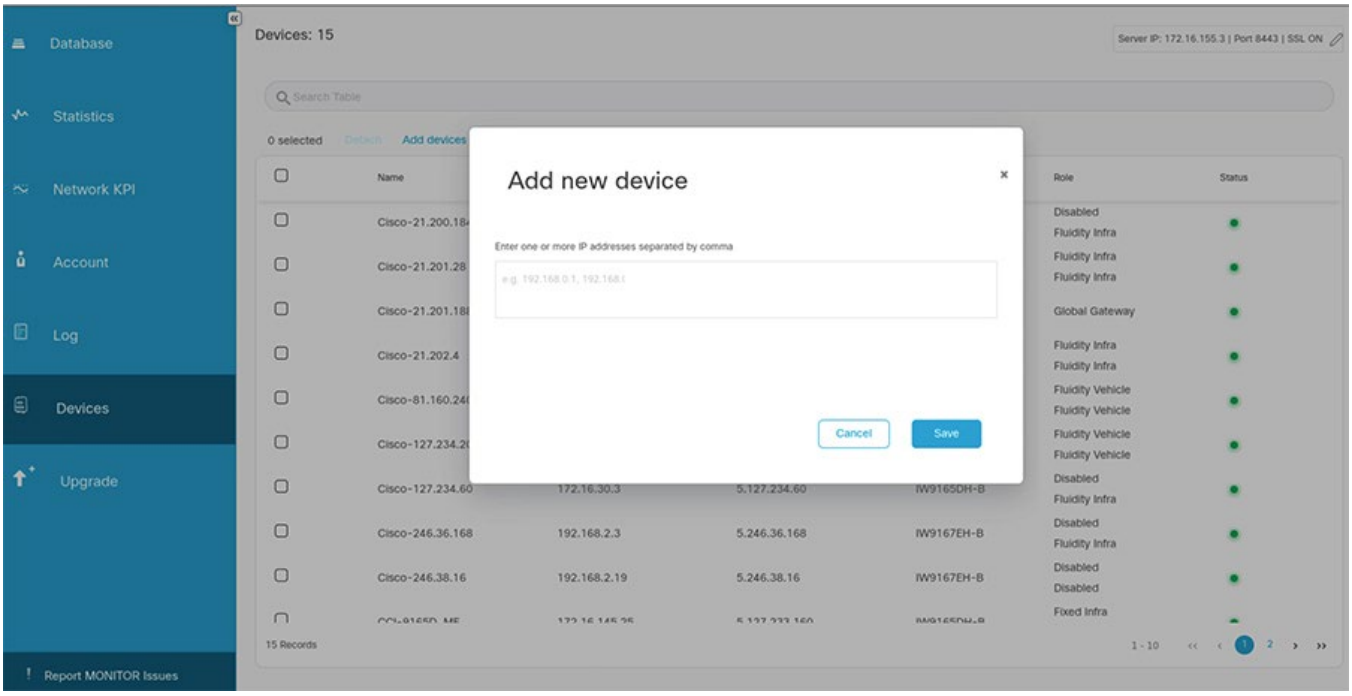


After devices have been added, the topology map is updated with the devices, telemetry information, and any links to other URWB devices.

For detailed information about IW Monitor, see the User Guide here: *https://www.cisco.com/c/en/us/td/docs/wireless/outdoor_industrial/IW-Monitor/b-iw-monitor-Configuration-Guide.html*

**Figure 3 IW Monitor Topology**

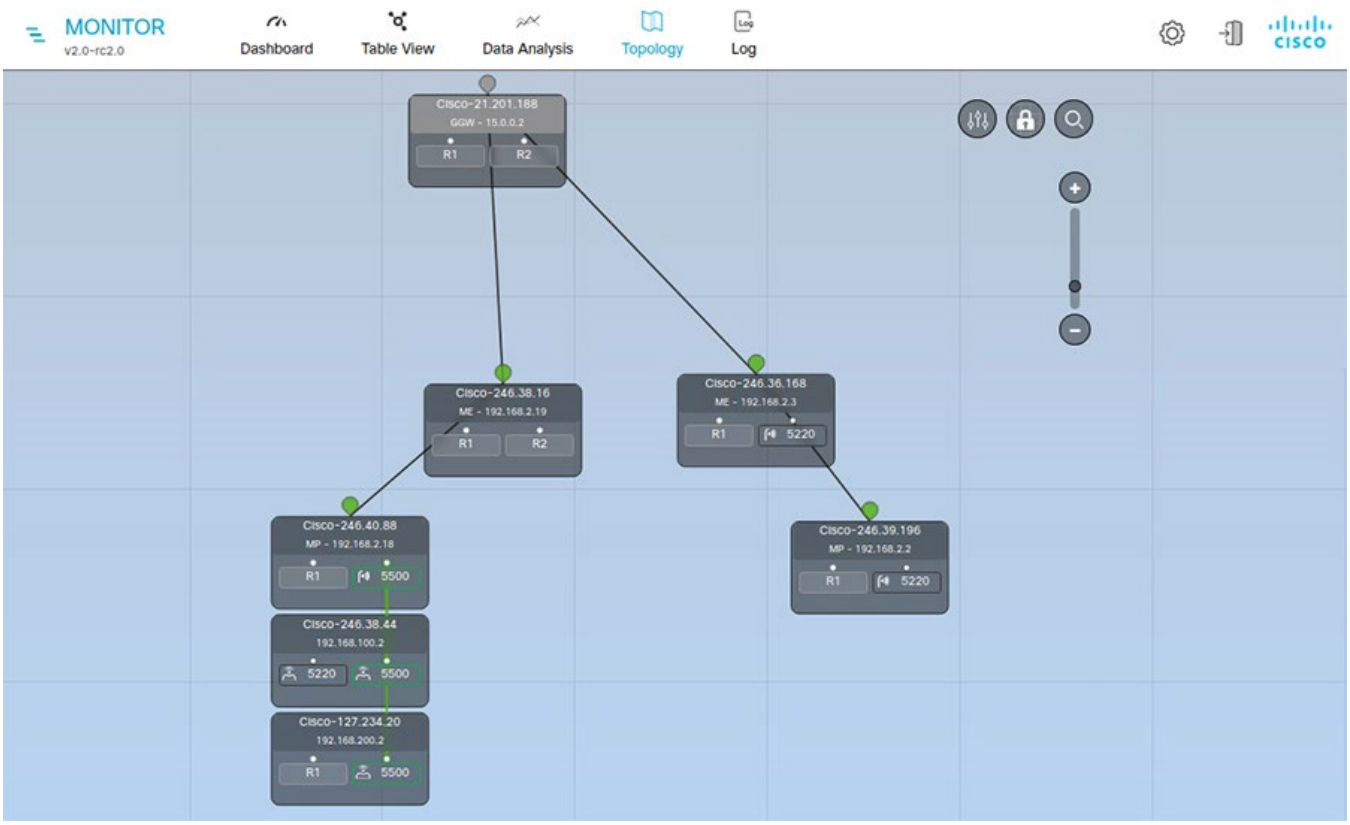


# Core Network Implementation

The core network is divided into multiple overlapping areas because of the redundancy requirements of the Rail CBTC and Safety solution. There are two separate LAN paths for the duplicated vital traffic to communicate between the wayside and control centers. There are also duplicate core networks that provide connectivity for both LAN paths if one of the core networks fails. Therefore, there are a minimum of four core nodes required to duplicate both the LAN paths and the core networks. The two core networks are connected to each other and the backbone nodes in their respective LAN path through a REP ring. The core nodes are also connected to the OCC/BOCC through dedicated switches that also provide connectivity to the enterprise network and Internet. The connection between the core and OCC/BOCC is out of scope for this document and is based on enterprise routing best practices.

## L3 Routing

The Catalyst 9300 core nodes are connected to the IE9300 backbone nodes as part of a REP ring which requires the ports to be configured as trunks. The core nodes are also the L3 gateway for all the connected backbone nodes. Since there are two core nodes per ring, a First Hop Redundancy Protocol (FHRP) like Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP) is used to provide a single default gateway

address to the backbone nodes. Static routing to the backbone node subnets with redistribution into the core routing protocol ensures full reachability. The example configuration used in this implementation guide is shown in the following figure. The second LAN path B is shown for completeness but not included in this implementation guide.

**Figure 4 L3 Core**

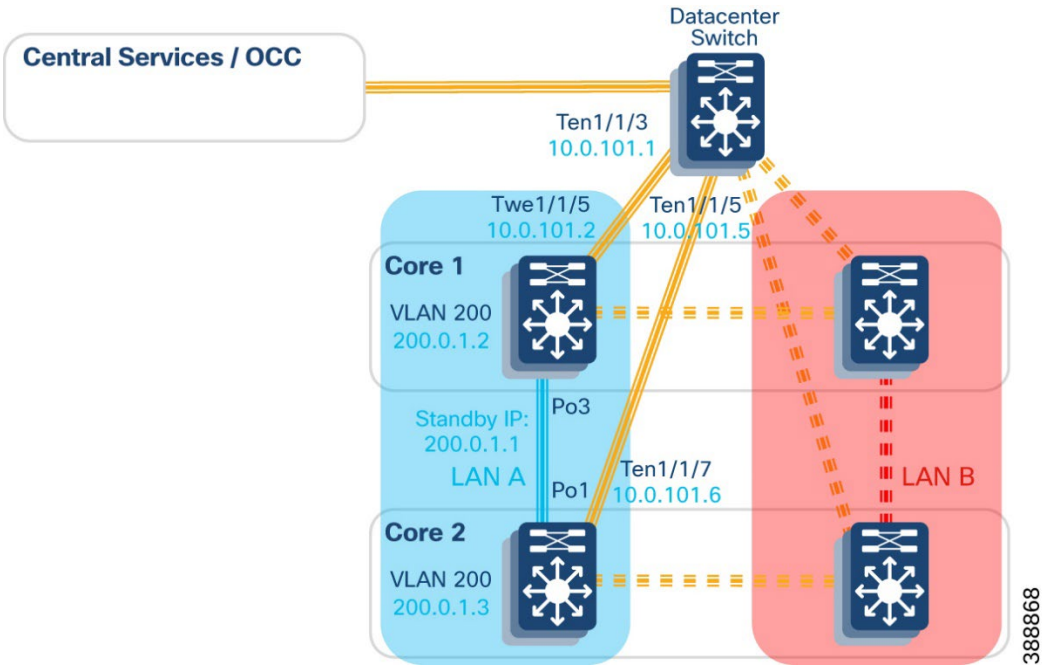


A sample configuration of the backbone facing interface follows.

Core 1

```
switchport mode trunk
carrier-delay msec 0
rep segment 3
```

The second core node is configured as the REP edge to minimize the chance of the link between the core nodes becoming the blocking link. This configuration follows.

Core 2

```
switchport mode trunk
carrier-delay msec 0
rep segment 3 edge primary
rep preempt delay 15
```

Setting the carrier-delay to 0ms gives the quickest reaction to a link down event to minimize convergence time. This command should be configured on all core facing links to minimize routing convergence times. The management SVI for the backbone ring is configured with HSRP in this example that follows.

Core 1

```
interface Vlan200
 ip address 200.0.1.2 255.255.255.0
 standby delay minimum 30 reload 60
```

```
 standby 1 ip 200.0.1.1
 standby 1 timers msec 200 msec 750
 standby 1 priority 200
 standby 1 preempt delay minimum 5
```

Core 2

```
interface Vlan200
 ip address 200.0.1.3 255.255.255.0
 standby delay minimum 30 reload 60
 standby 1 ip 200.0.1.1
 standby 1 timers msec 200 msec 750
```

It is important to modify the timeout value such that the backup does not try to take over the active role during certain convergence events. The convergence of a REP ring event or switch stack failover could cause the FHRP to inadvertently move from *backup* to *active* and then back to *backup* if the timers are too aggressive. In the example given, the hello interval is 200ms and the hold time, the time elapsed before the backup takes over, is 750ms.

## MPLS

Like the L3 routing core network, the NCS540 core nodes are connected to the IE9300 backbone nodes as part of a REP ring. The NCS540 is not a switch, however, and does not support trunks so the ports are configured with subinterfaces to permit the L2 traffic. Every VLAN that is passed between the backbone nodes must be configured as a subinterface on the NCS540 nodes. In this solution, the backbone nodes are the L3 gateway for all the end devices so there is only a single VLAN on the Core/backbone REP ring for node management. Another subinterface is used for the REP traffic that is untagged.

The example configuration used in this implementation guide follows. The second LAN path B is shown for completeness but not included in this implementation guide.

**Figure 5 MPLS Core**

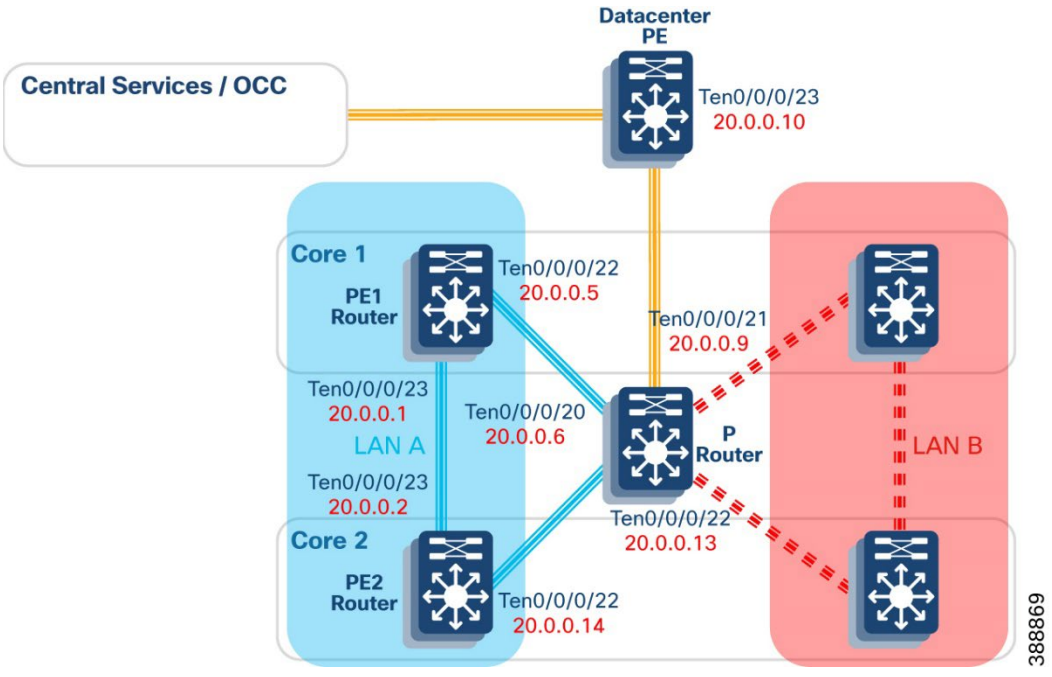

The connection between the NCS540 and backbone node follows.

Core1

```
interface Bundle-Ether1
!
interface Bundle-Ether1.1 l2transport
 encapsulation untagged
!
interface Bundle-Ether1.100 l2transport
 encapsulation dot1q 100
 rewrite ingress tag pop 1 symmetric
!
interface GigabitEthernet0/0/0/0
 bundle id 1 mode on
 cdp
 speed 1000
 carrier-delay down 0
!
```

To enable REP communication between the backbone nodes, the REP Access Gateway (REP-AG) feature is used on the NCS540 nodes. This configuration follows.

```
spanning-tree repag rep
 interface Bundle-Ether1.1
  instance 0
   priority 4096
   root-priority 4096
  !
 !
!
```

The backbone network is configured as an EVPN which enables L2 and L3 VPN services and supports multihoming for the backbone nodes. This means a VRF is needed for all traffic on this ring. The configuration example follows.

```
vrf red
 vpn id 200:200
 address-family ipv4 unicast
  import route-target
   200:200
  !
  export route-target
   200:200
  !
 !
!
interface BVI100
 host-routing
 vrf red
 ipv4 address 3.3.3.1 255.255.255.0
 mac-address aaa.aaaa.aaaa
```

This configuration is applied identically to both core nodes so the backbone nodes have a single default gateway to the rest of the core network.

Because the core nodes are redundant to each other, there must be a load balancing mode configured that allows multihoming while not creating a switching loop in the ring. This is accomplished using EVPN single-flow-active multihoming balancing mode.

More details on this feature can be found here:
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/73x/b-l2vpn-cg-ncs5500-73x/evpn-features.html#Cisco_Concept.dita_7c1cd1e2-2b0c-4b58-8515-a95426c327a7 .

The core nodes share the same Ethernet Segment ID number for the bundle interface connected to the directly connected backbone nodes. The core nodes also share the same VLANs on the REP ring and therefore have the same EVPN instance (EVI) for each VLAN. This configuration which is the same on both core nodes follows.

```
evpn
 evi 100
 !
 interface Bundle-Ether1
  ethernet-segment
   identifier type 0 11.11.11.11.11.11.11.11.11
   load-balancing-mode single-flow-active
   convergence
    mac-mobility
    nexthop-tracking
   !
  !
 !
!
```

L2VPN is then configured to enable communication between all the nodes on the configured VLANs by creating a bridge domain and associating the EVI to the backbone BVI interface. TCN-propagation ensures that the topology change notifications from REP and STP are propagated.

```
l2vpn
 tcn-propagation
 bridge group 100
  bridge-domain 100
   interface Bundle-Ether1.100
   !
   routed interface BVI100
   !
   evi 100
   !
  !
 !
!
```

MPLS with segment routing is then configured to complete the communication between all the core nodes in the network. This requires a supported IGP to carry the segment data between the nodes. In this solution, IS-IS is used between all the core nodes using a loopback for the prefix segment ID (SID).

Core 1

```
interface Loopback0
 ipv4 address 100.0.2.1 255.255.255.255
!
router isis 1
 is-type level-2-only
 net 49.0001.0000.0000.0001.00
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid index 1
  !
 !
 interface TenGigE0/0/0/22
  address-family ipv4 unicast
  !
 !
 interface TenGigE0/0/0/23
  address-family ipv4 unicast
  !
 !
!
```

BGP is then configured to enable communication across the core with L2VPN or L3VPN instances. In this example, the second core node is configured with the L2VPN EVPN address family while a core node connected to the datacenter is configured as a L3VPN node.

Core 1

```
router bgp 1
 bgp router-id 100.0.2.1
 address-family vpnv4 unicast
  vrf all
   label mode per-vrf
  !
 !
 address-family l2vpn evpn
  nexthop trigger-delay critical 0
 !
 neighbor 100.0.2.2
  remote-as 1
  update-source Loopback0
  address-family l2vpn evpn
  !
 !
 neighbor 100.0.2.4
  remote-as 1
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
 vrf red
  rd auto
  address-family ipv4 unicast
```

```
    redistribute connected
    redistribute static
   !
  !
!
```

The backbone nodes are configured with default routes to the core nodes, so the core nodes need static routes to all the networks behind the backbone nodes. These static routes are then redistributed into BGP for full reachability between the backbone and core. An example pointing to a backbone node follows.

```
router static
 vrf red
  address-family ipv4 unicast
   4.4.4.0/24 3.3.3.10
```

The other core node in the same LAN path must have the same routing statements to ensure consistent routing depending on which core node is the primary for a traffic flow.

# Backbone Network Implementation

## L3 Routing

The IE9300 backbone nodes are connected to the Catalyst 9300 core nodes and the other backbone nodes in a closed REP ring. They are also connected to the wayside access nodes in a different closed REP ring. Each REP segment has a different ID and admin VLAN. To ensure that traffic does not take suboptimal paths and to limit flooding, it is recommended to configure the trunk ports for each REP ring with only the VLANs that will be located there.

The REP ring interfaces can be configured using Cisco Catalyst Center with Day-N templates. For more information on configuring Day-N Templates for REP ports, see the Offshore Wind Farm Implementation Guide for TAN REP Ring Configuration.

An example topology used in this implementation guide follows.

**Figure 6 L3 Backbone**

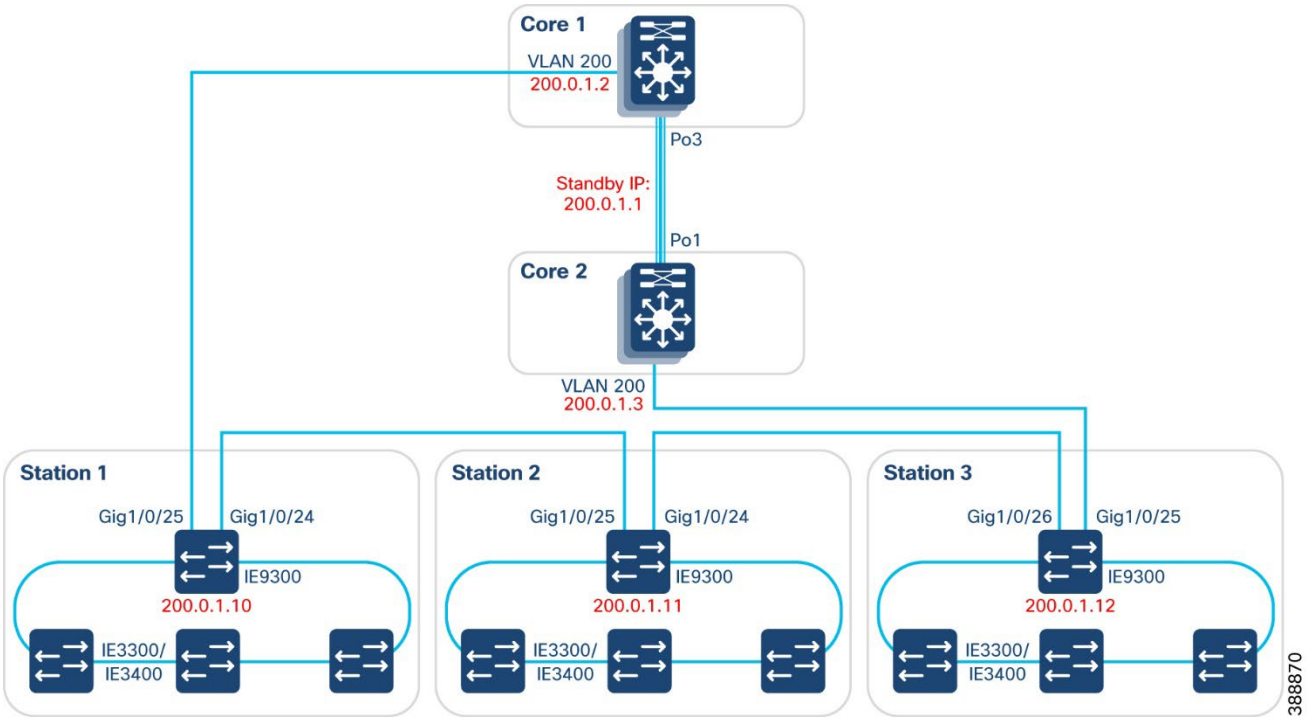


## MPLS

The IE9300 backbone nodes are connected to the NCS540 core nodes and the other backbone nodes in an open REP ring. The ring is open because the NCS540 does not support REP. The backbone nodes connected to the core nodes must configure the directly-connected interfaces as REP edge ports with the no-neighbor option because of this. Because the NCS540 does not support REP, the backbone node must advertise REP updates into STP. This is done using the topology change notification feature.

An example topology used in this implementation guide follows.

**Figure 7 Backbone with MPLS core**

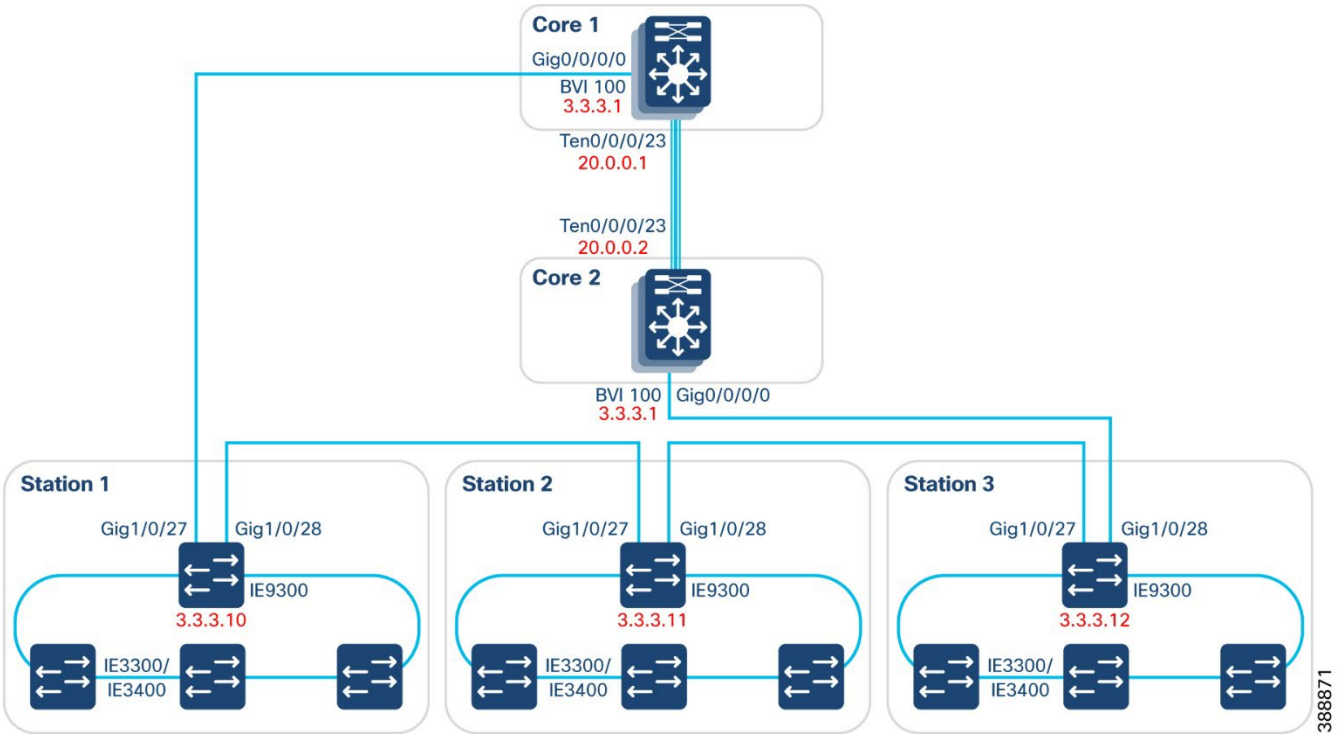


An example of this config follows.

```
switchport mode trunk
mtu 9198
rep segment 1 edge no-neighbor preferred
rep stcn stp
```

Because the backbone switches only have a single path to the rest of the core networks, a default route is configured to the BVI IP address given in the previous section.

```
ip route 0.0.0.0 0.0.0.0 3.3.3.1
```

# Wayside Access Network Implementation

The wayside access network composed of Catalyst IE3x00 switches is a closed REP ring connected to the Catalyst IE9300 backbone switch. The ring configuration is the same whether the core network is enterprise L3 routed or service provider MPLS. More details on REP and its configuration can be found here: https://www.cisco.com/c/en/us/td/docs/switches/lan/cisco_ie3X00/software/17_4/b_redundancy_17-4_iot_switch_cg/m_1610-layer2-rep-cg.html .

The REP ring interfaces can be configured using Cisco Catalyst Center with Day-N templates. For more information on configuring Day-N Templates for REP ports, see the Offshore Wind Farm Implementation Guide for TAN REP Ring Configuration.

The following figure shows an example wayside access network used in this implementation guide.

**Figure 8 Wayside Access Network**

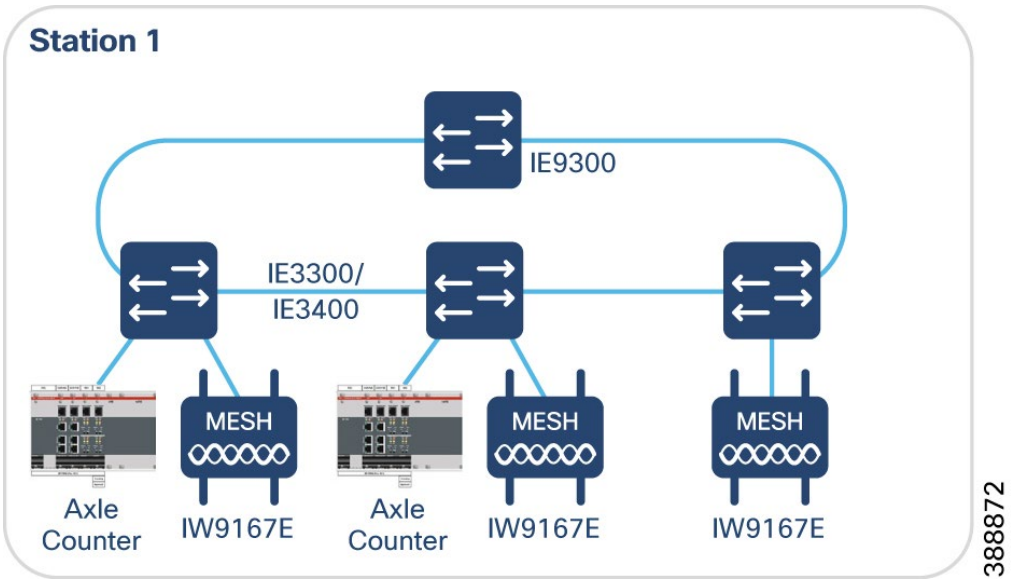


## Device Onboarding

When a wayside host device is connected to the network, the switchport must be manually configured with the correct VLAN or it can be automatically assigned using Cisco ISE with MAB or 802.1x. These devices can include the Frauscher axle counters, interlocking equipment, and other wayside network devices. The switchport must be configured using a Day-N template from Cisco Catalyst Center to authorize the connected host. When the host connects, the MAC address is sent to ISE where it is evaluated against the installed policies. If it matches, ISE replies back with the policies relevant to this host and the switchport will adjust accordingly. See the Industrial Automation Implementation Guide for more details: https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/Industrial_Automation/IA_Horizontal/IG/Industrial-AutomationIG/Industrial-AutomationIG.html .

# Wayside Wireless Network Implementation

The wayside wireless network includes the wayside radios connected to the wayside access network. The IEC6400 Edge Compute Appliance, Catalyst IW9167E, and IW9165D operate in URWB mode for this rail solution and can be configured using Industrial Wireless Service (IW Service) in IoT Operations Dashboard, the CLI, or the local GUI. The IW Service is recommended because of its centralized approach to configuring and deploying the IW products.

By default, the URWB are configured to connect to the IW Service. This requires the VLAN used for URWB management to have access to a DHCP server, DNS server, and access to the Internet. This configuration is outside the scope of this document.

All URWB networks have common elements regardless of the type of deployment. This includes Mesh Ends and Mesh Points. The primary role for the Mesh End is to connect the mesh network to the switched IP network. The Mesh Points connect the wireless clients to the mesh network whether they are configured on the wayside or the train.

# Layer 2 Fluidity

For Layer 2 Fluidity, all Mesh Ends and Mesh Points must be in the same broadcast domain and configured with the same wireless passphrase. The default gateway should be an SVI configured on the backbone switch to ensure network reachability from end to end.

An example L2 Fluidity network follows.

**Figure 9 L2 Fluidity Topology**

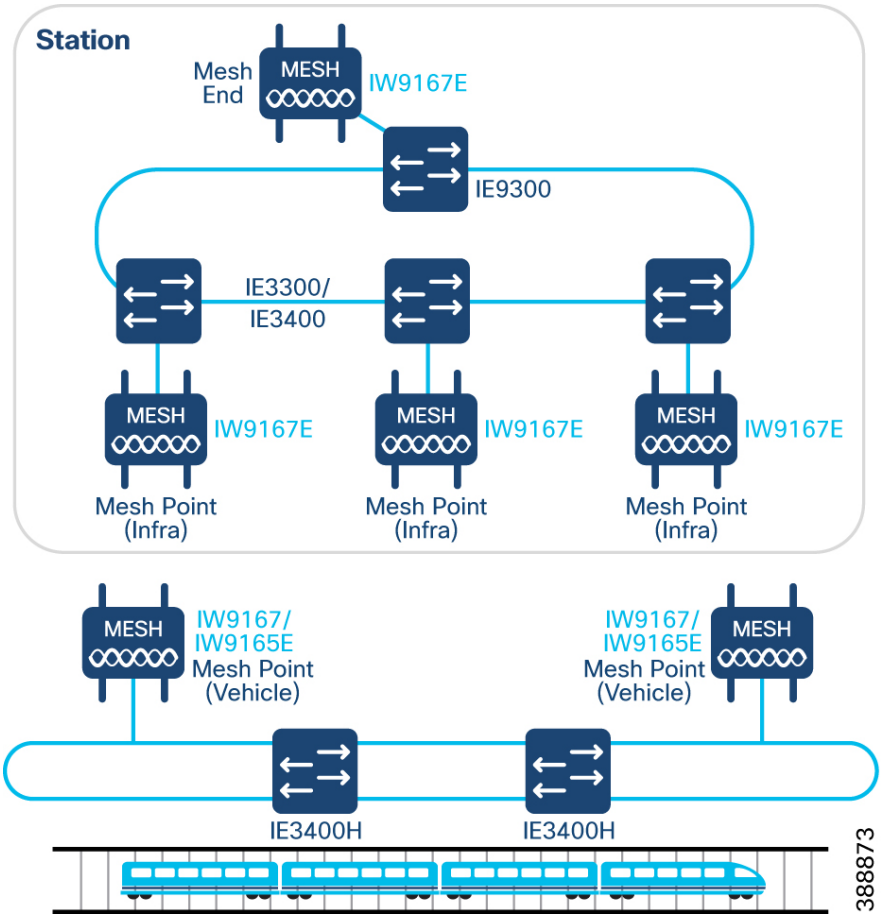


For networks with less than 50 IW devices or less than 2Gbps aggregate throughput, an IW9167E can perform the role of Mesh End. For deployments larger than this, it is recommended to configure an IEC6400 as the Mesh End. The Mesh Point role is always performed by the IW9167E or IW9165D on the wayside wireless network.

**Figure 10 Mesh Point General Mode**

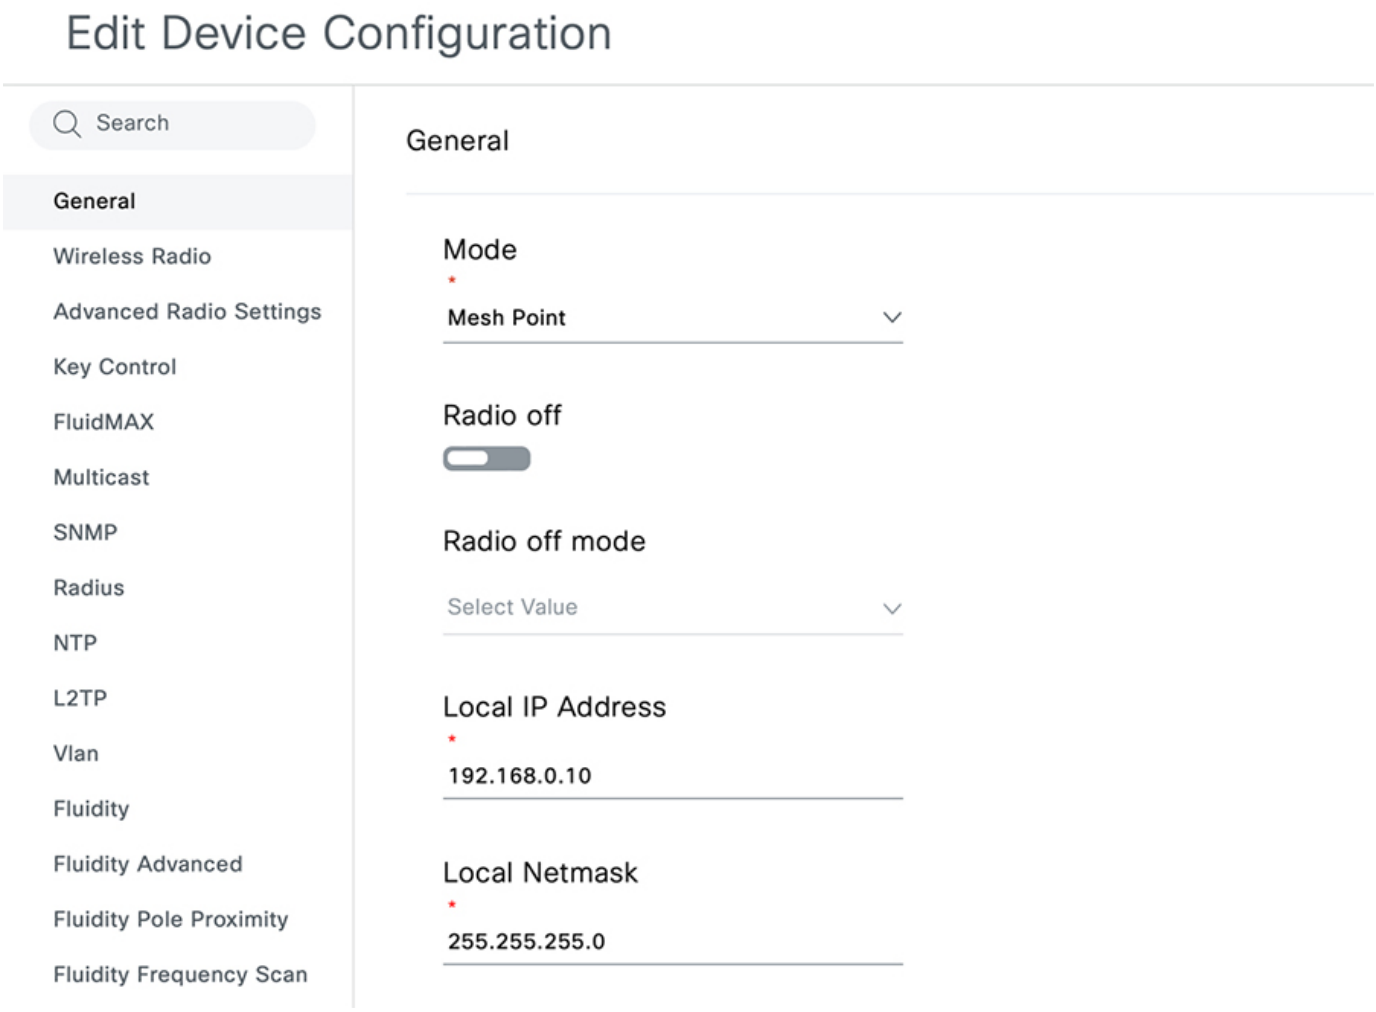


**Figure 11 Wireless Passphrase**

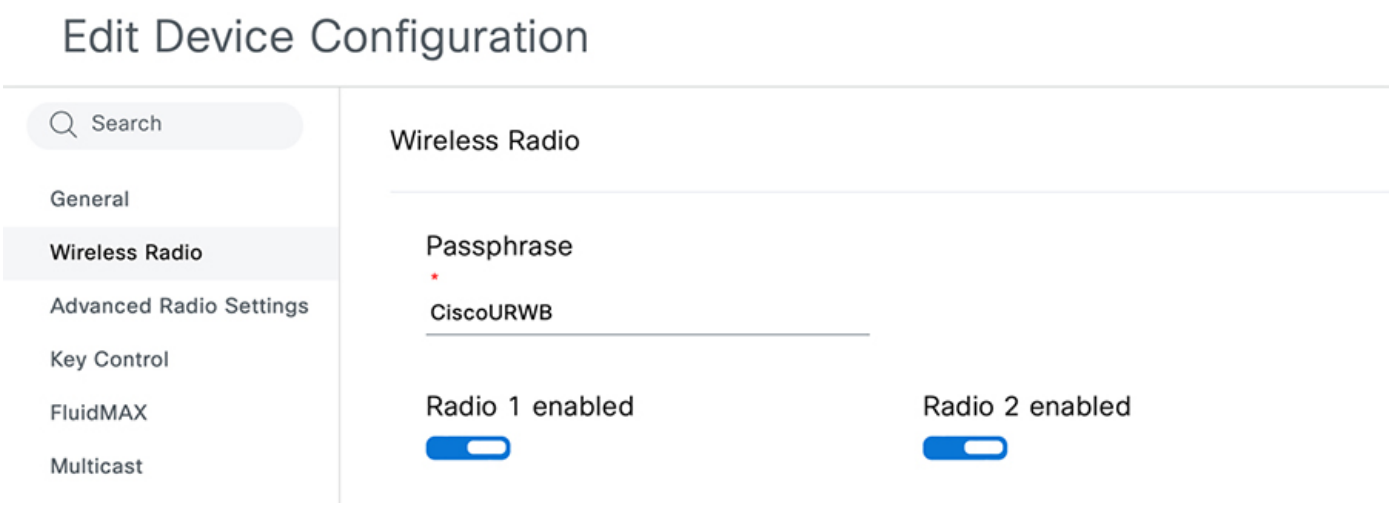


Other key configuration items are the radio settings which include which radios are enabled and the wireless parameters required.

**Figure 12 Fluidity Radio Settings**

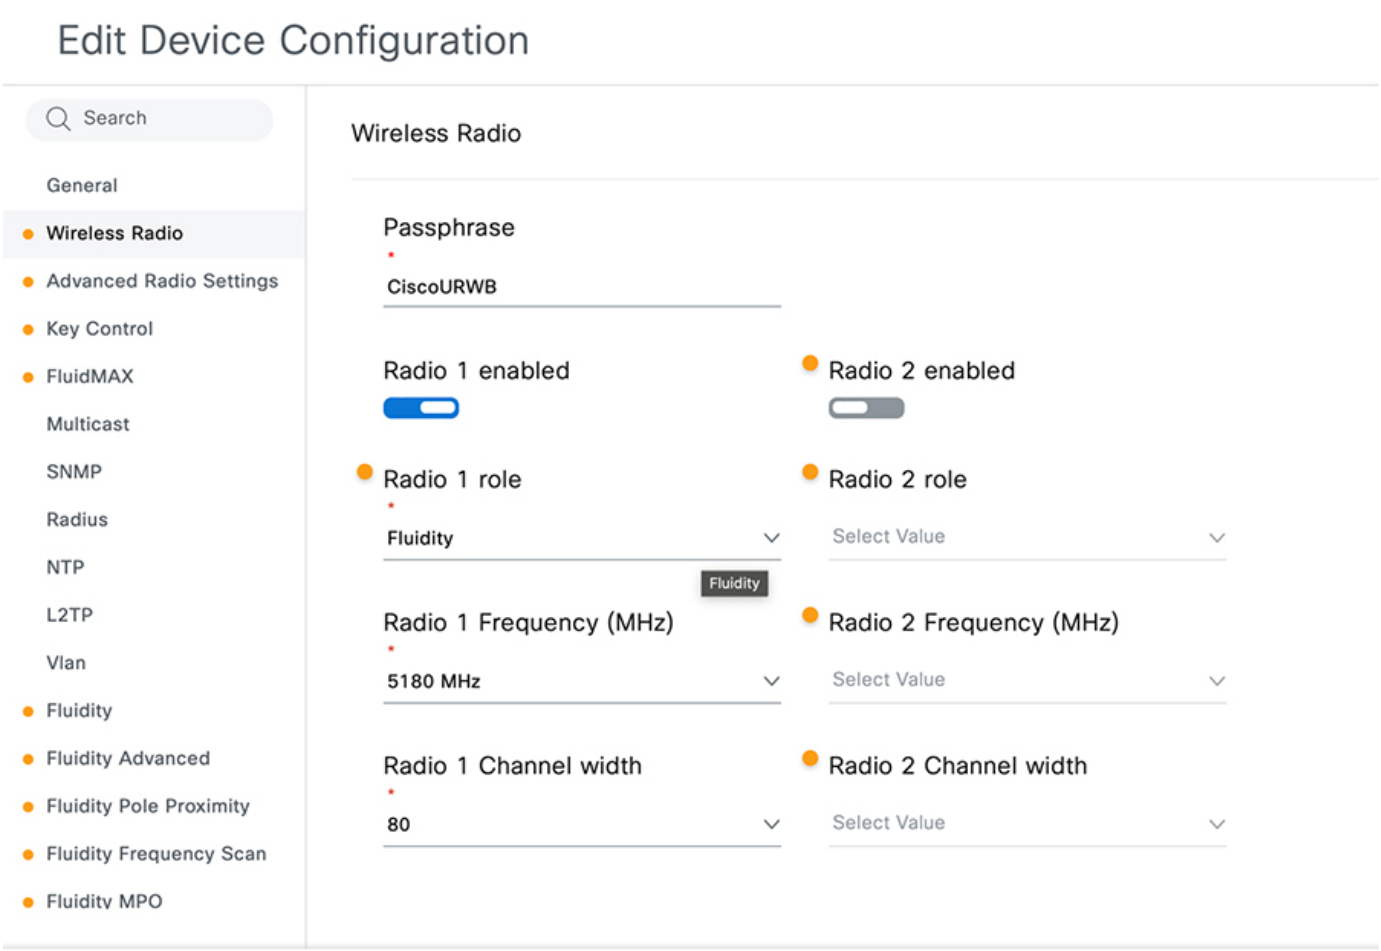


When enabling L2 Fluidity, the role must set to **Infrastructure** and the Network Type set to **Flat.**

**Figure 13 L2 Fluidity Role**

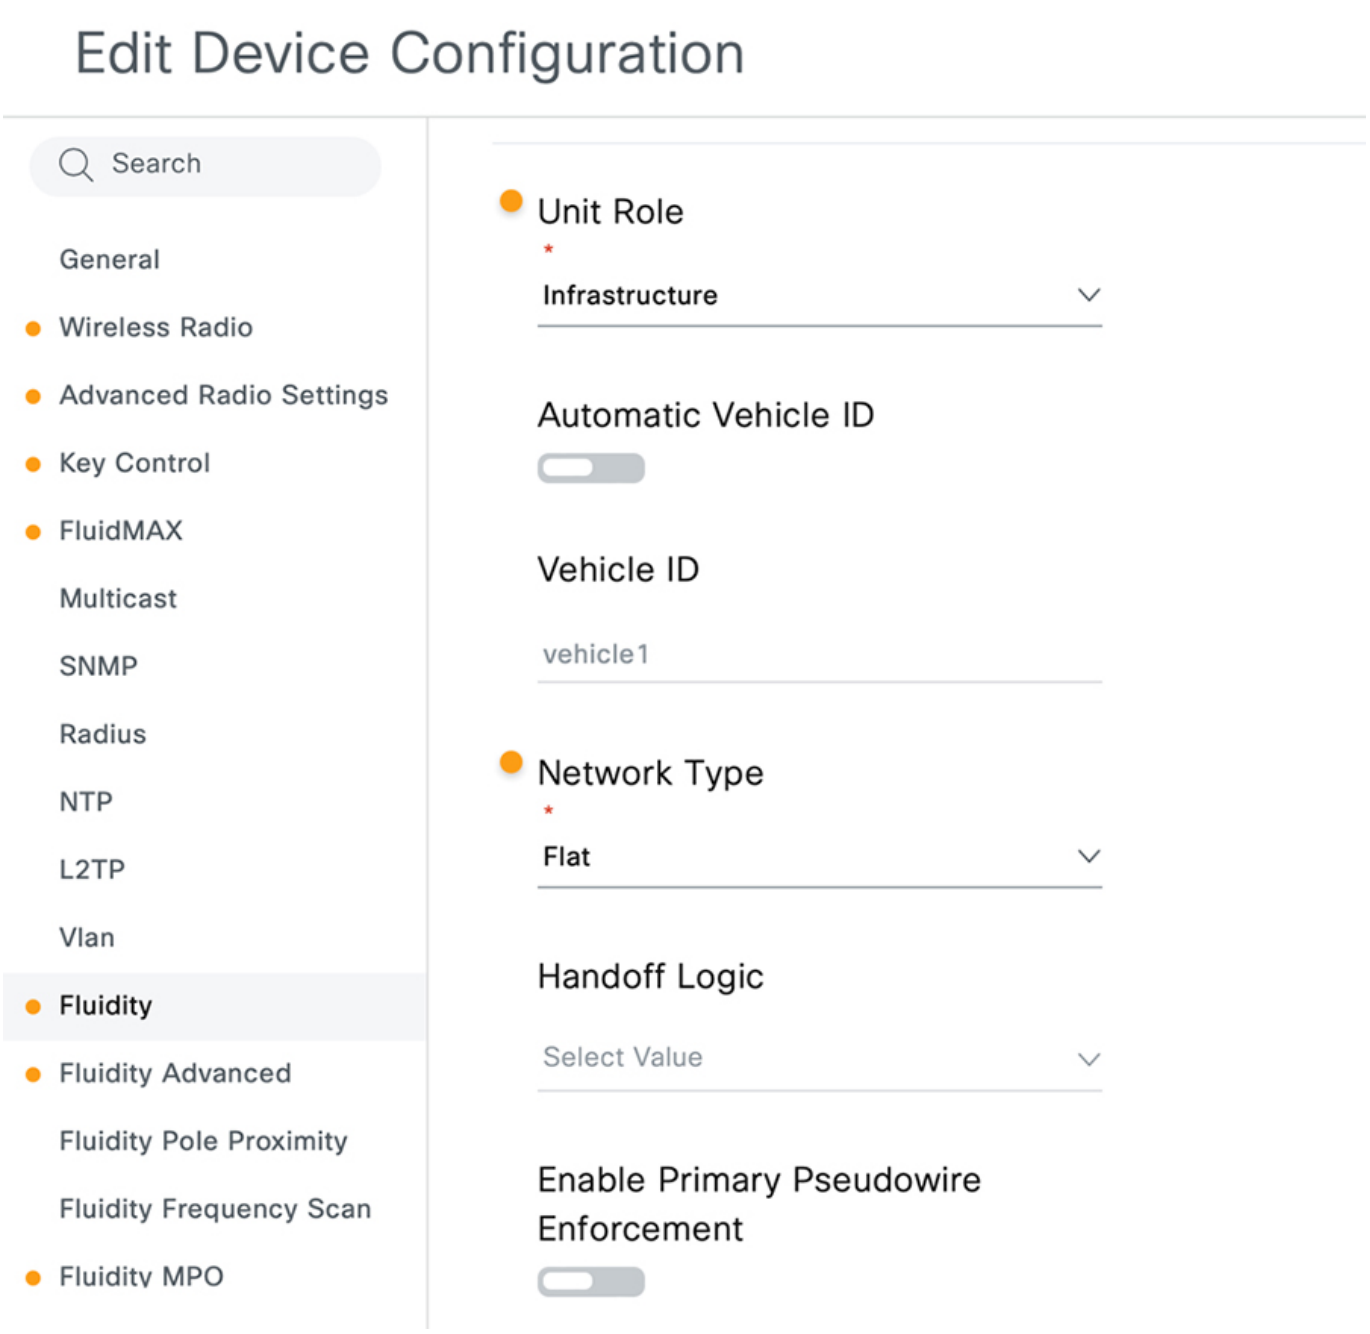


It is recommended to configure the Mesh Ends with redundancy using fast failover which enables a backup Mesh End to take over the primary role in the event of failure in less than 500ms.

**Figure 14 Fast Failover**

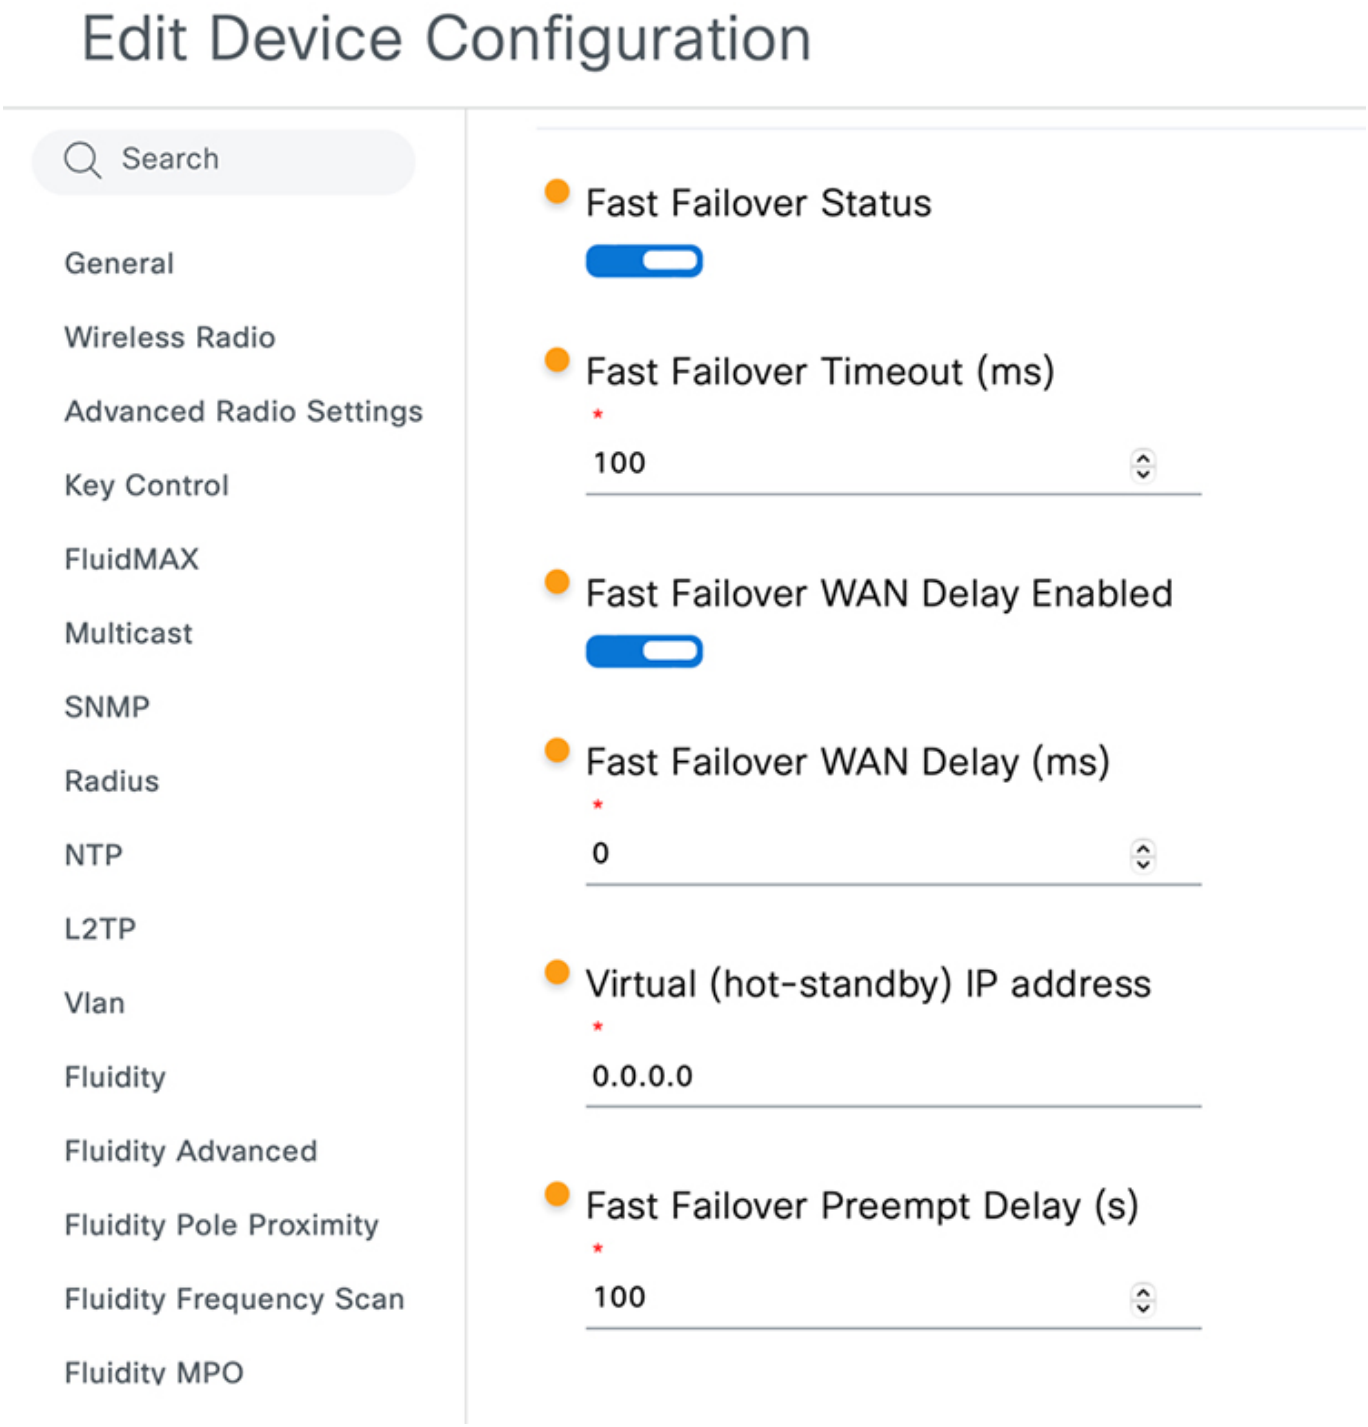


Layer 2 Fluidity also supports VLANs after enabling the feature. When enabled, the IW interfaces act as a trunk by retaining the VLAN tag from the client traffic. There are only two configuration options available, one is for the management VLAN ID for managing the radio and the other is the native VLAN for all untagged traffic.

**Figure 15 Configuring VLANs**

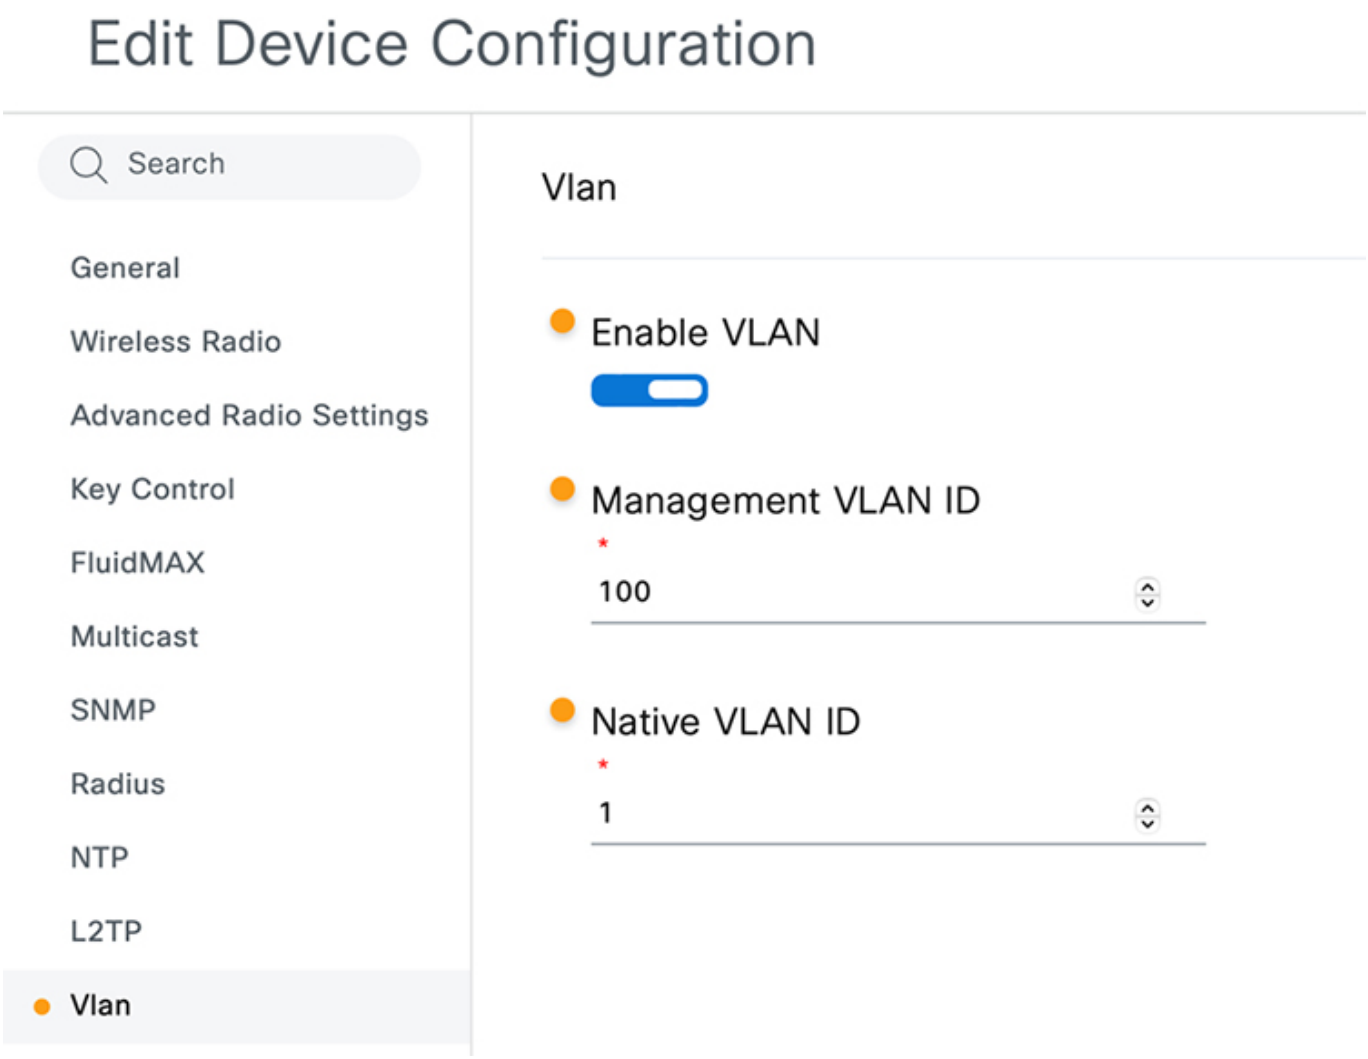


## Layer 3 Fluidity

L3 Fluidity is like L2 Fluidity in that there are Mesh Ends and Mesh Points for each mesh network. In L3 Fluidity, however, each grouping of Mesh Ends and Mesh Points are in a separate L3 domain from other groups of Mesh Ends and Mesh Points. Additionally, the train IW916x units are in a different L3 domain from the wayside units. To enable seamless roaming between the different mesh networks, a special Mesh End called the Global Gateway is configured to communicate with all the Mesh Ends. This mode is configured in the General settings of the IW9167E or IEC6400.

An example L3 Fluidity network follows.

**Figure 16 L3 Fluidity Topology**

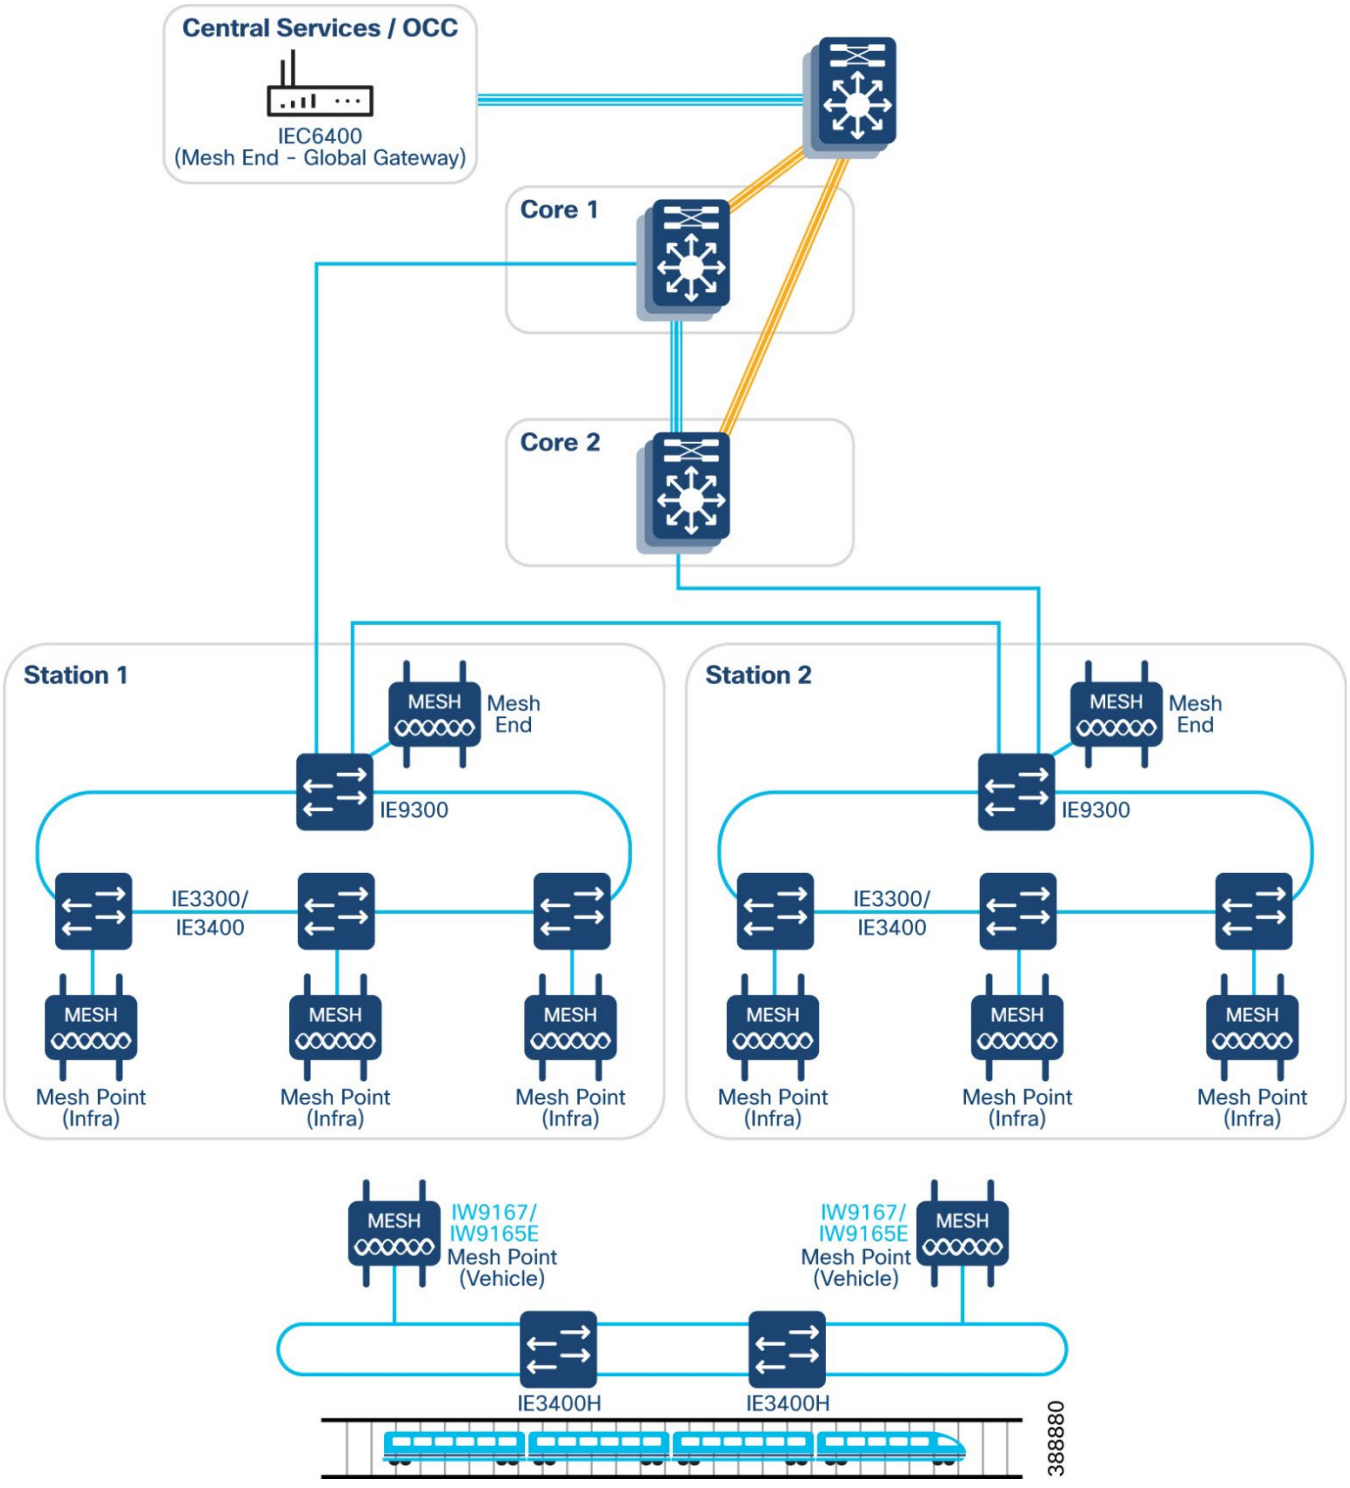


The default gateway for the wayside wireless devices should be an SVI on the backbone switch for full reachability end to end. The Global Gateway is installed at the OCC/BOCC and should have a default gateway pointing to the core switch in that location.

**Figure 17 Global Gateway IW**

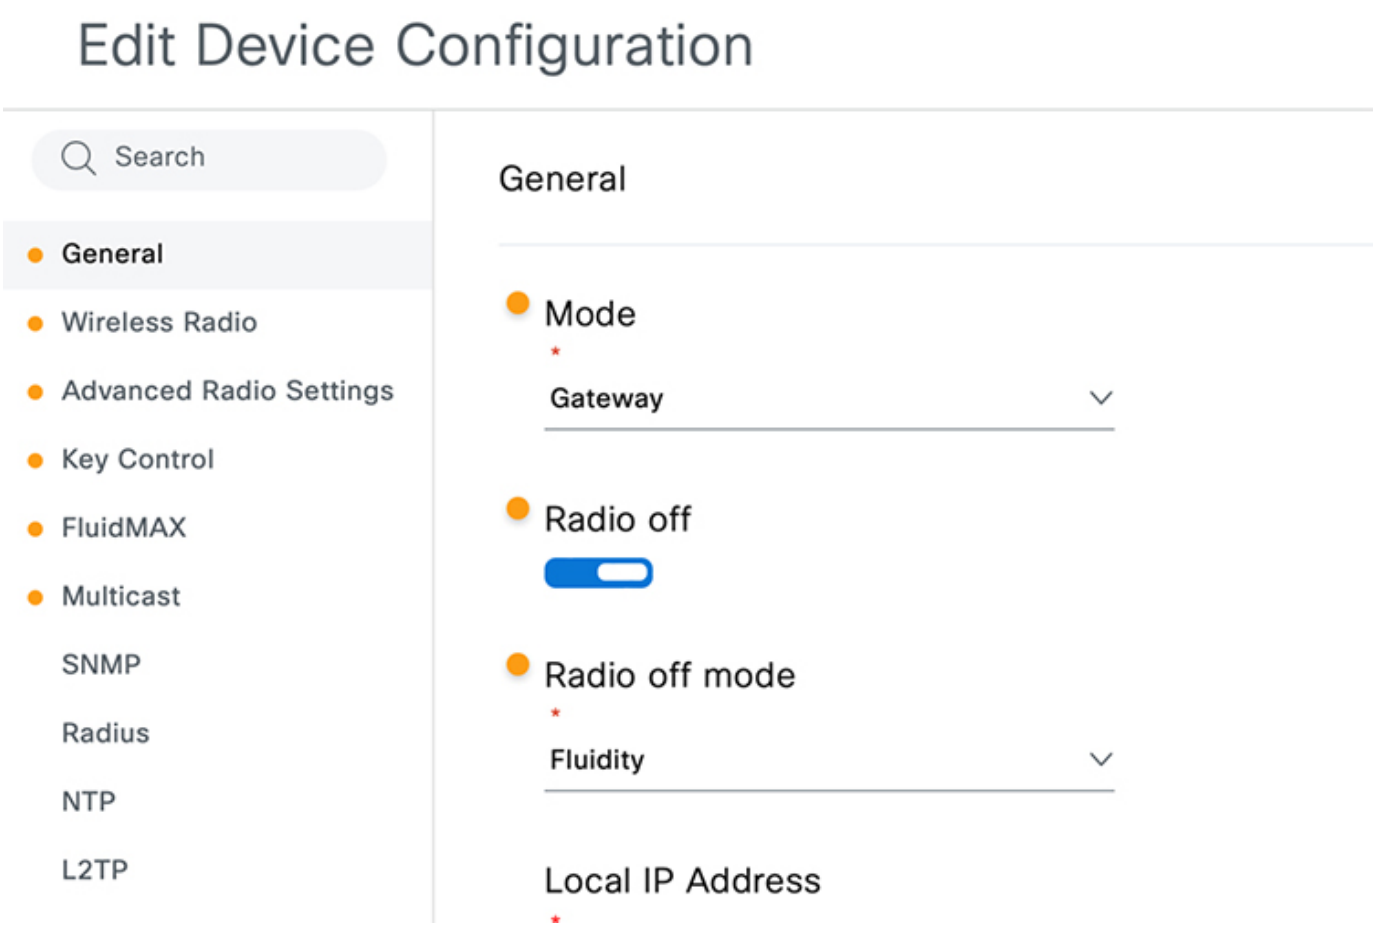


Because the IEC6400 has no radio interfaces, it has a different appearance in IW Service.

**Figure 18 Global Gateway IEC6400**

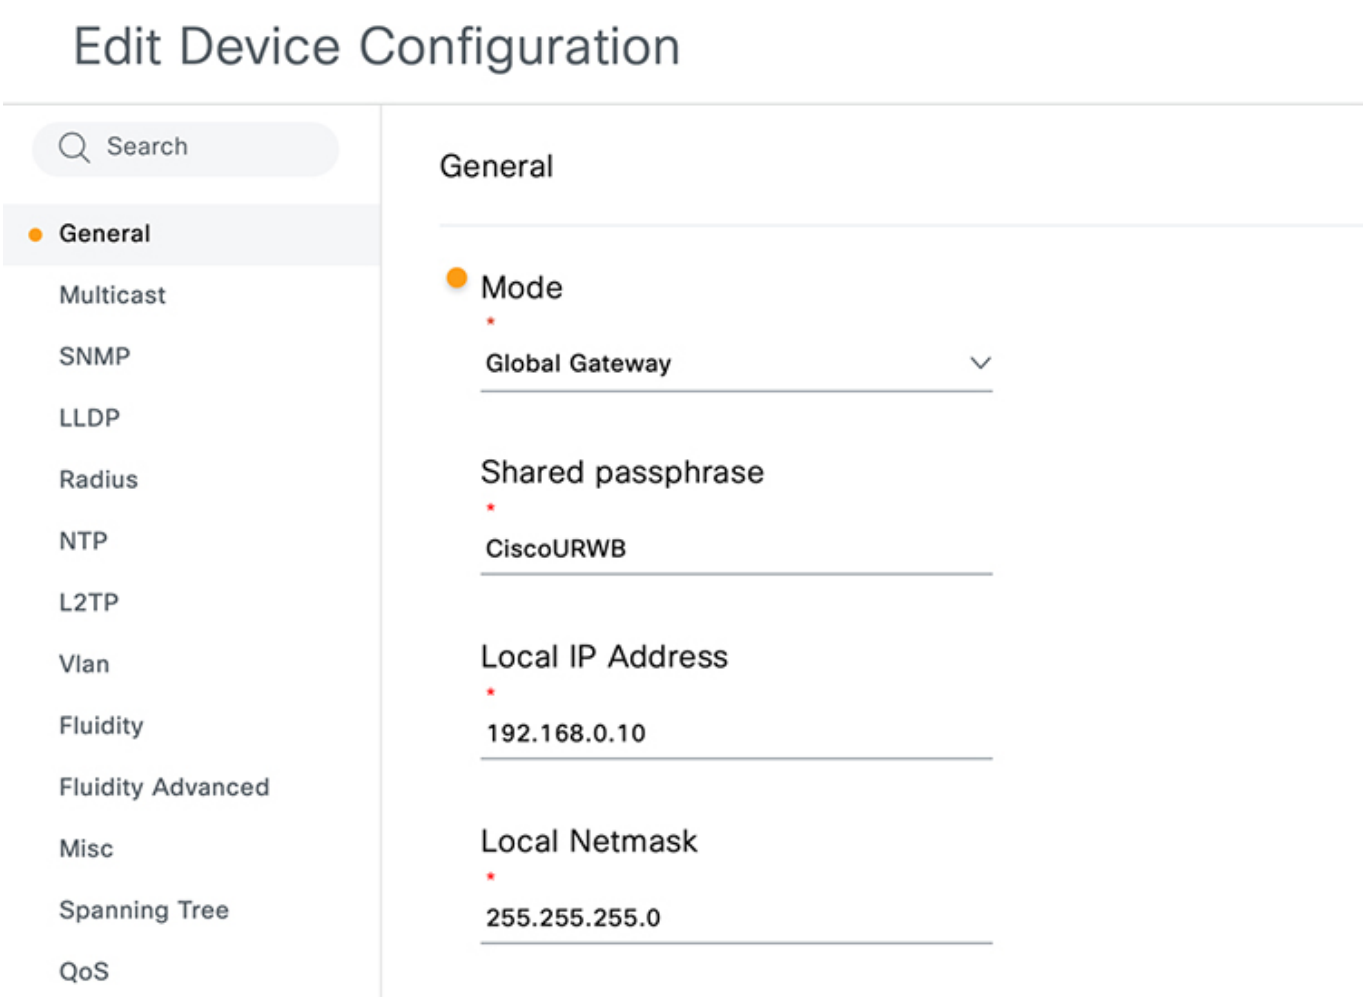


Configuring the Mesh Ends and Mesh Points in L3 Fluidity is like L2 Fluidity except that the Fluidity Network Type is **Multiple Subnet**.

**Figure 19 L3 Fluidity Role**

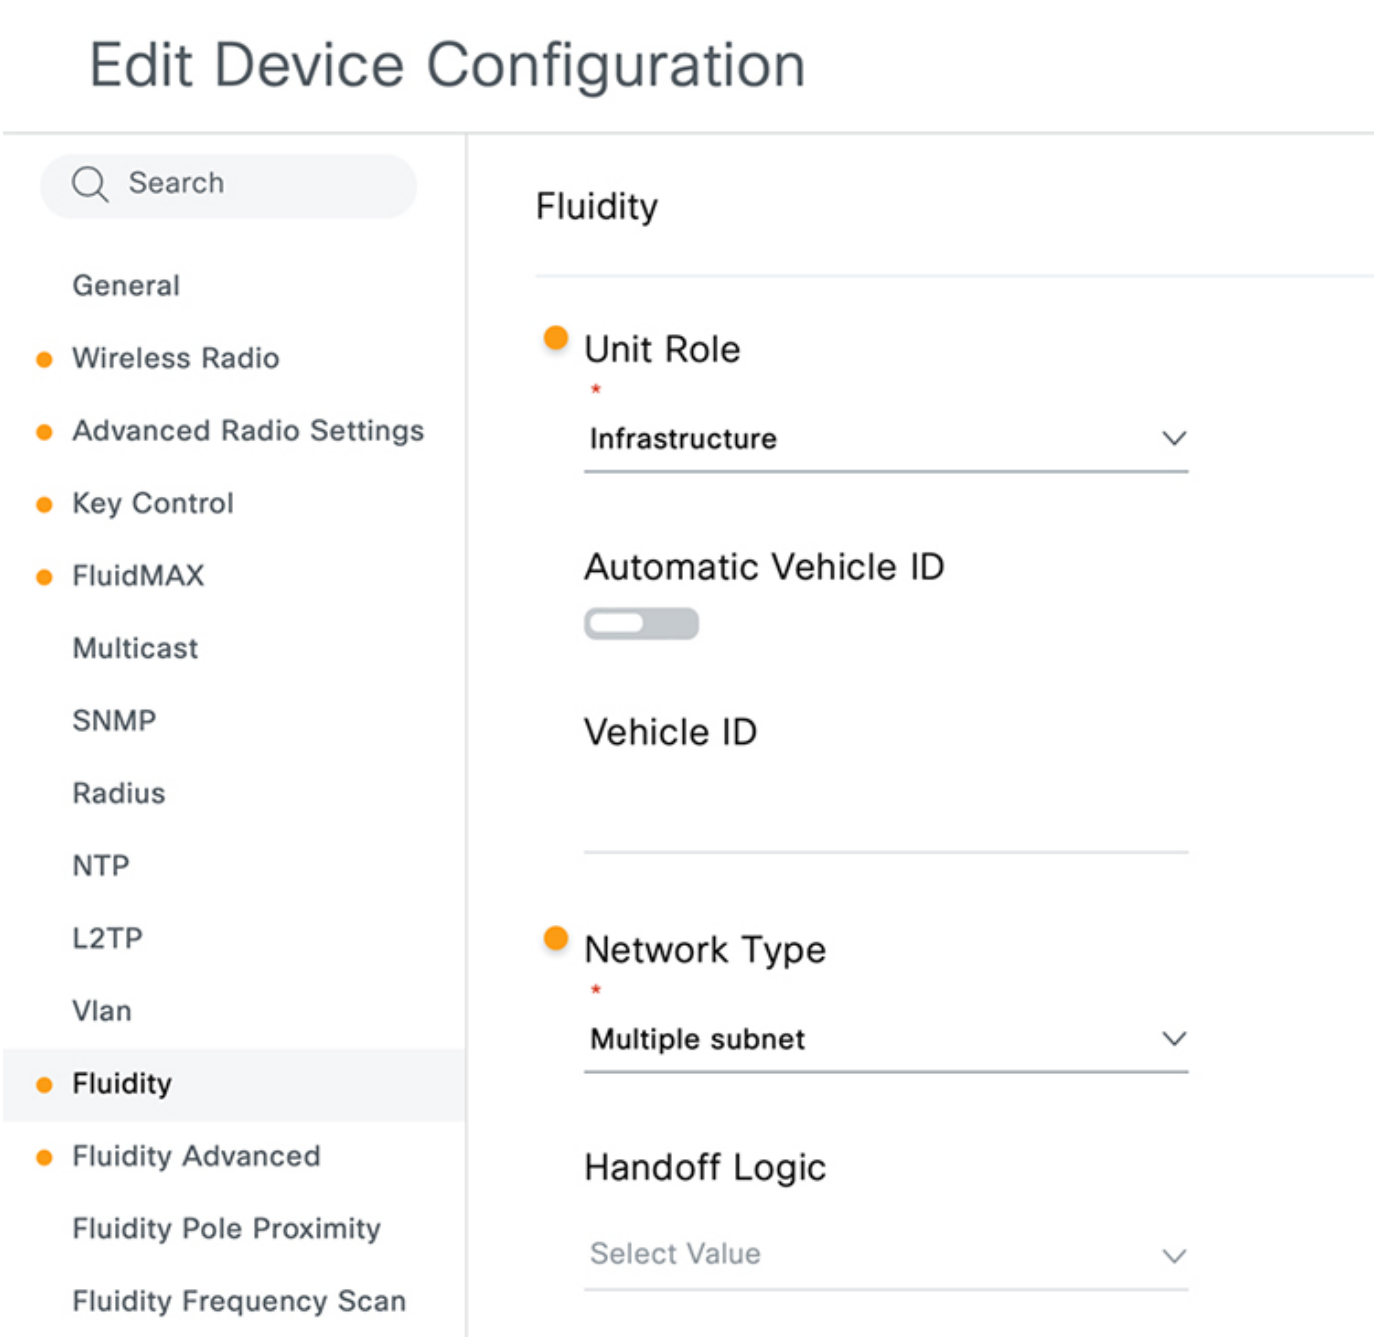


VLANs are also not supported in L3 Fluidity and must not be configured.

When configuring fast failover, the Mesh Ends and Global Gateway are configured the same as L2 Fluidity except that Global Gateways also use a virtual hot-standby IP address. This virtual IP address is used when the underlying network needs route reachability to the mobile networks inside the mesh network. This IP must be on the same subnet as the device IP and must be the same for both devices in the fast failover pair.

**Figure 20. Fast Failover Virtual IP**

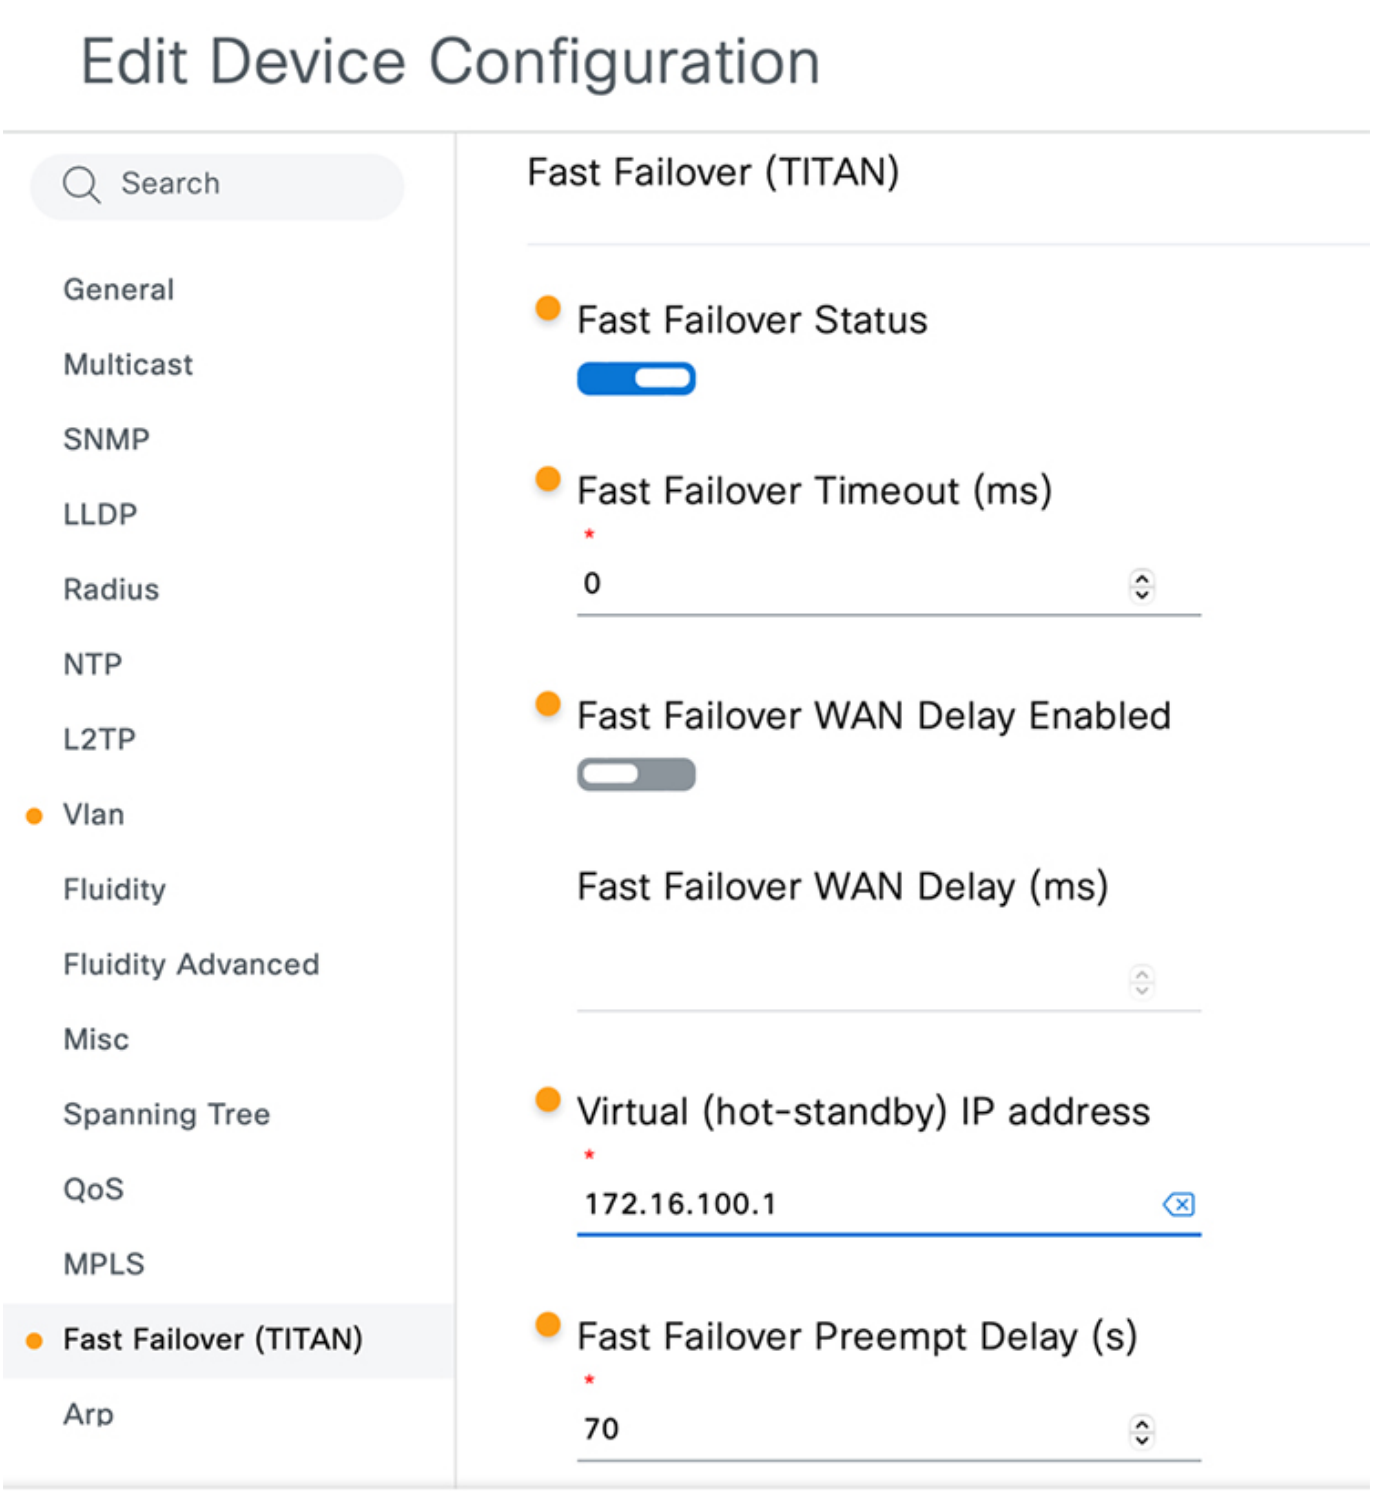


## L2TP

To enable the mesh network communication between the Mesh Ends through the Global Gateway, all the Mesh Ends must be configured with L2TP tunnels pointing to the Global Gateway. In addition to the device IP address, the Mesh Ends and Global Gateway are also configured with a WAN IP address that is used for L2TP communication. It must be on the same subnet as the device IP and not already assigned to another device.

After the WAN IP address is configured, the tunnels from Mesh End to Global Gateway are configured on the Mesh Ends and then the corresponding tunnel is configured on the Global Gateway to the Mesh Ends.

**Figure 21 L2TP**

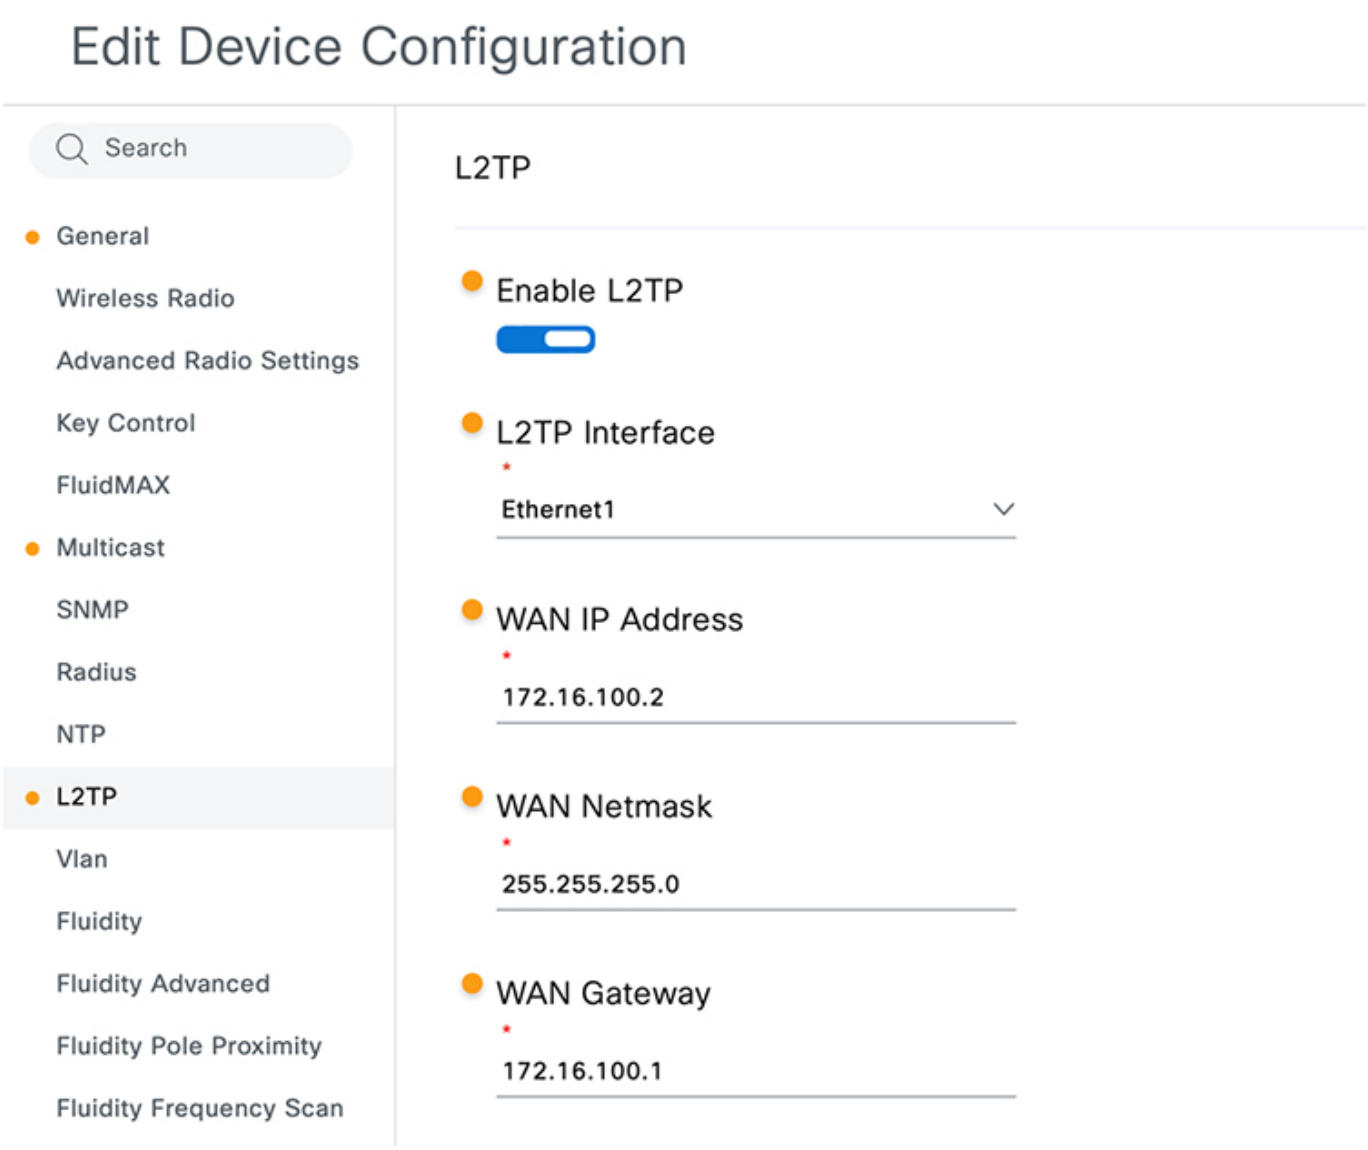


Figure 22 L2TP Tunnels

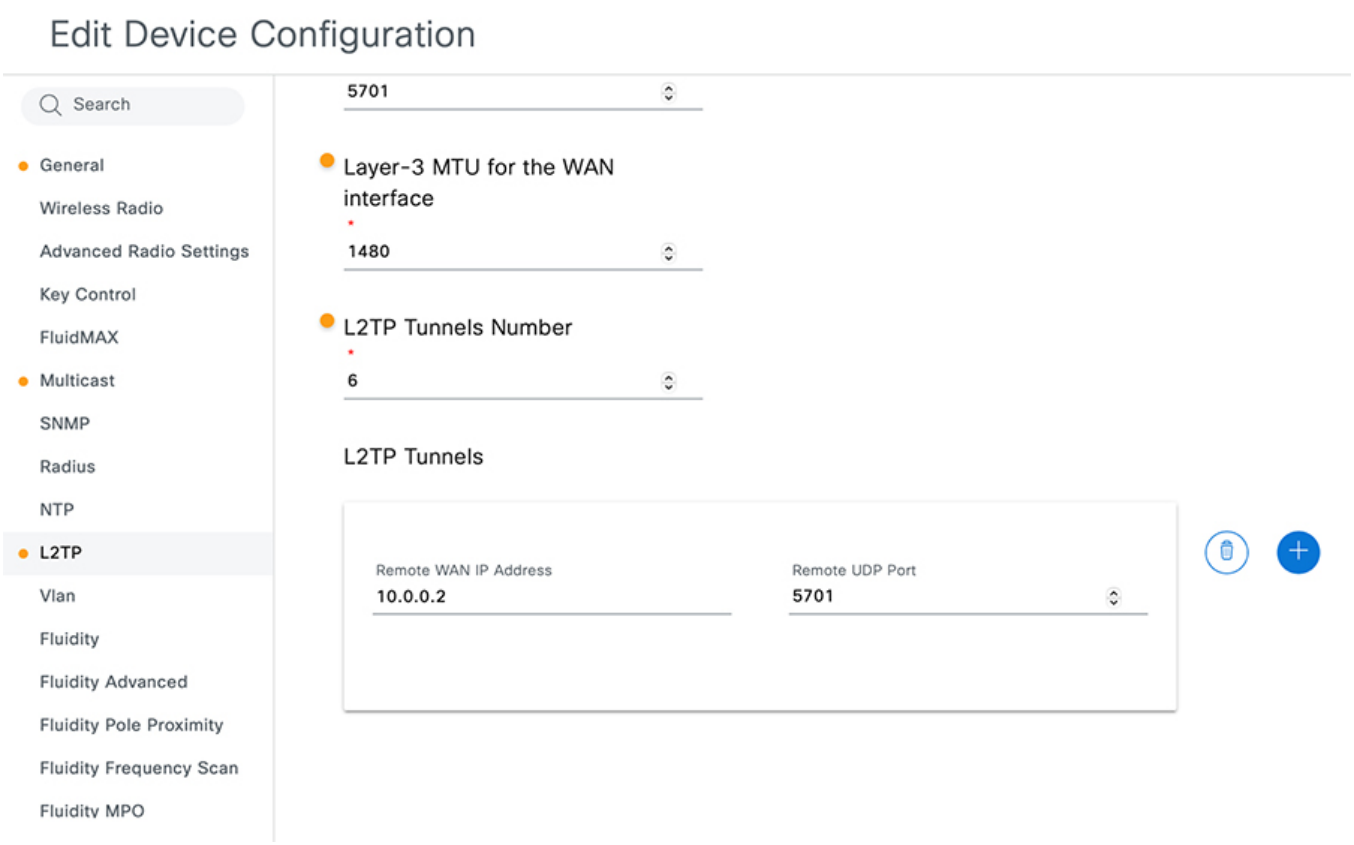


When configuring L2TP in addition to fast failover, the Mesh End procedure differs from the Global Gateway. The Mesh Ends are not configured with a virtual hot-standby address so the Global Gateway must have tunnels to both WAN IP addresses in the Mesh End pair. If the Global Gateways are configured with fast failover, the L2TP WAN IP address must be configured the same on both Global Gateways. The Mesh Ends then configure an L2TP tunnel to this one address. In the case of having fast failover configured on the Global Gateways, there will be two interface IP addresses, one virtual hot-standby IP address, and one L2TP WAN IP address for the pair of Global Gateways. With redundant Mesh Ends, there will be two interface IP addresses and two L2TP WAN IP addresses.

When all L2TP tunnels are created, the active Global Gateway will show the tunnels as **CONN** to all active Mesh Ends. The standby units will show the tunnels in **IDLE** state until they become active.

## Core Network

After the wireless mesh network is created for L3 Fluidity, the core network needs extra routing information to reach the mobile networks. Since the mobile networks and the train radio networks are not on the same subnet as the wayside mesh network devices, the core network needs static routing to reach those destinations. The static routes must use the Global Gateway IP address (or virtual hot-standby address when using redundant Global Gateways) as the next hop.

# Onboard Train Network Implementation

The onboard train network includes the part of the wayside wireless network that connects to the wayside in addition to the onboard switching network. The Catalyst IE3400 Heavy Duty switches can be configured as standalone for smaller deployments in a single car or as part of a REP ring for a large consist.

Integrating the switched network with the wireless wayside network depends on the type of Fluidity being deployed. L2 Fluidity is used when devices on the L2 wired network need L2 adjacency to devices on the train. In this case, traffic is trunked from the train to the wayside access network. L3 Fluidity is used when the train will roam between different L3 domains. In this scenario, the traffic from the train is routed through the wayside access network.

## Layer 2 Fluidity

### Wired Network

With L2 Fluidity, the switched network connects to the onboard IW916x with an L2 trunked port. After the traffic leaves the mesh network at the Mesh End, the VLAN header will still be maintained and switched to the destination.

### Wireless Network

The onboard IW916x is configured as a Mesh Point like the wayside Mesh Points. The device IP is configured on the same subnet as the wayside mesh devices with the same default gateway. The Fluidity role is configured as Vehicle instead of Infrastructure. The Vehicle ID can be assigned automatically or manually. When configuring a pair of IW916x radios for redundancy, the vehicle ID must be manually assigned to the same value. The network type is configured as **Flat** to support L2.

**Figure 23 L2 Fluidity Vehicle**

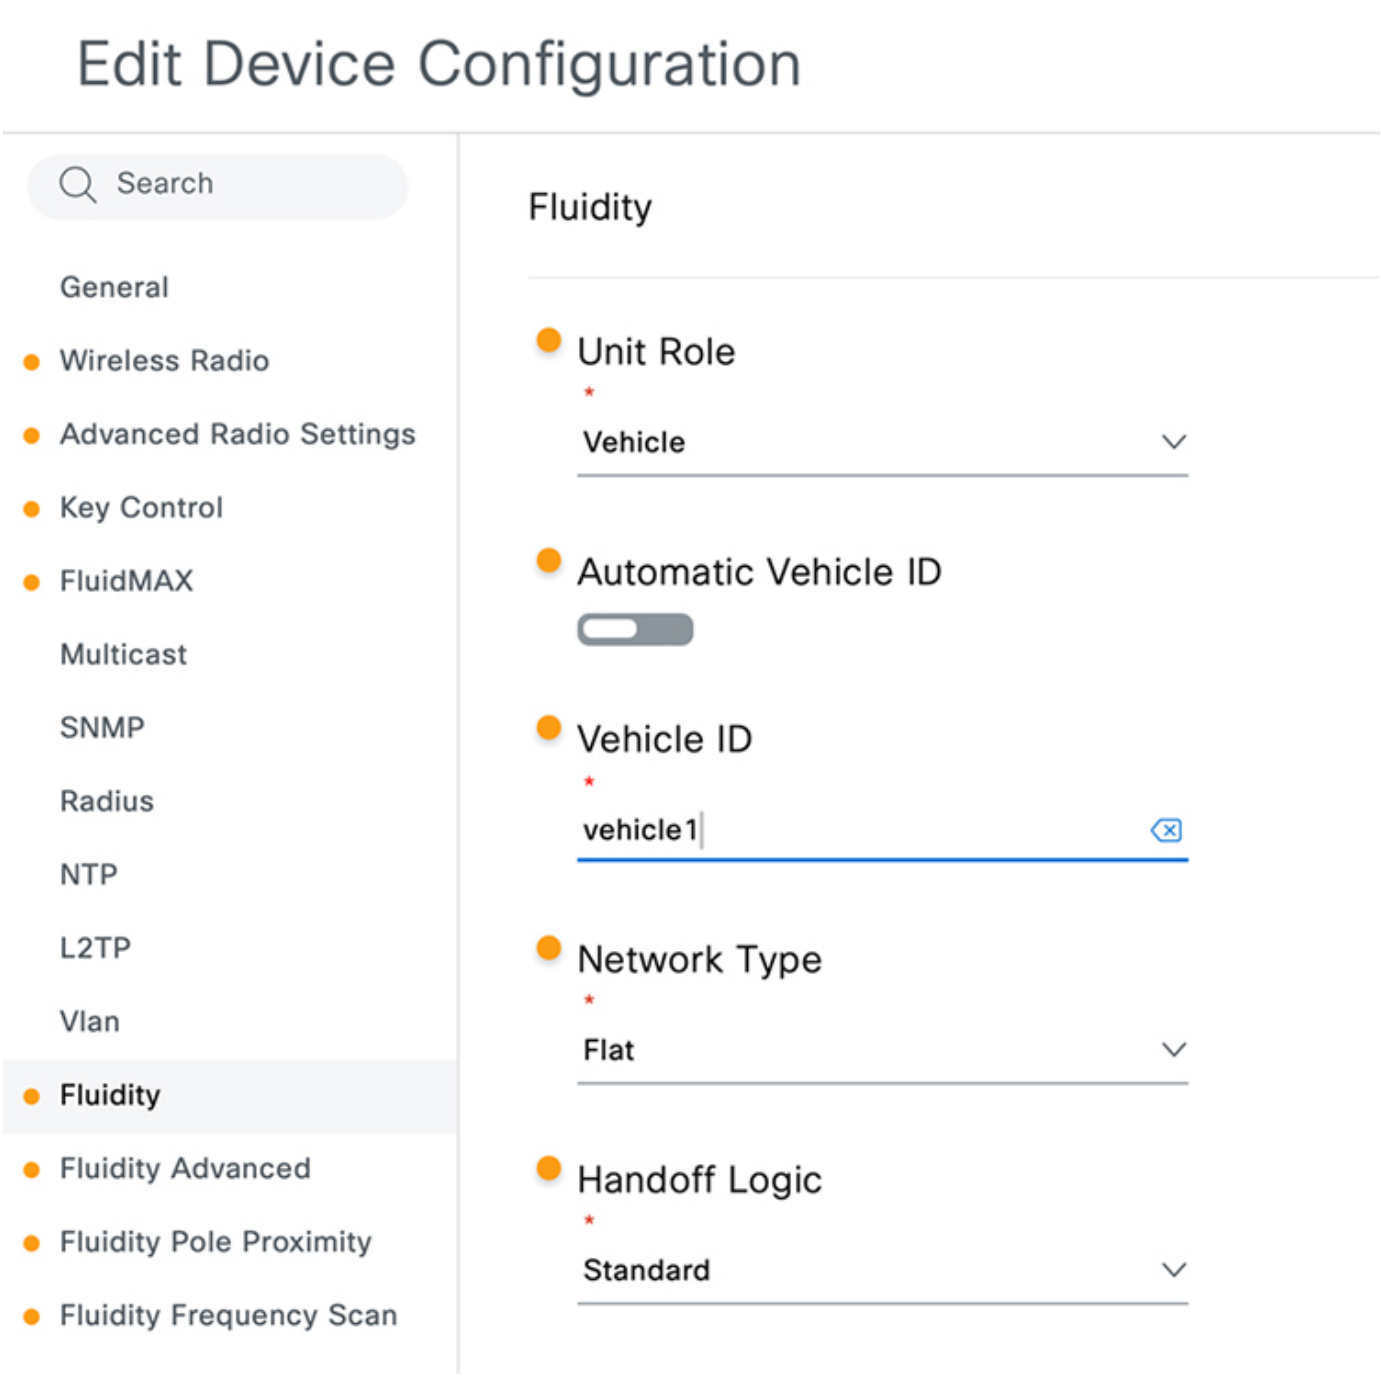


The VLAN feature must also be configured to support the trunked VLANs from the onboard switched network. This is configured the same as the Mesh Points and Mesh Ends on the wayside wired network.

As mentioned previously, when configuring the IW916x for fast failover, the vehicle ID on both must manually be configured to the same value. They each have a different device IP and do not use a virtual hot-standby address.

## Layer 3 Fluidity

## Wired Network

To support L3 Fluidity, the onboard IE3400 Heavy Duty switches are configured as the L3 gateway for the onboard train networks. The port connected to the IW916x radio can be configured as an L3 port or as an access port. If configured as an access port, the VLAN should be used solely for managing the IW916x. If a separate management VLAN is used for the IW916x radios, another IW916x could be configured on the same VLAN to support a redundant pair with fast failover enabled. The IP addresses used should be in a subnet dedicated to train management.

An SVI should be configured for each VLAN required by the train onboard network. These SVIs serve as the default gateway for each VLAN on the switched network. The switched network also requires a default route to send all traffic from the onboard network to the wayside access network. This default route uses the IW916x as the next hop. If the IW916x radios are configured with fast failover, a virtual hot-standby address is used and the switch used as the L3 gateway must use this virtual address as the next hop to reach the wayside access network.

## Wireless Network

The IW916x radios are configured in the same IP subnet as the switched train management network which is decoupled from the mesh network on the wayside access network. The default gateway for the IW916x radio is the L3 interface on the onboard switched network. Because the train management network and the client subnets on the train are decoupled from the wayside access network, the onboard IW916x radios must have static routes added that point to all the subnets onboard the train. The next hop for these static routes is the same L3 interface as the default gateway. Once these static routes are created, the mobile radio will advertise these subnets to the other mesh nodes it connects to. This enables the Global Gateway to send the return traffic to the correct train radio.

The Fluidity configuration is similar to the L2 Fluidity configuration except the Network Type is **Multiple Subnet** and the Vehicle ID must be set manually.

**Figure 24 L3 Fluidity Vehicle**

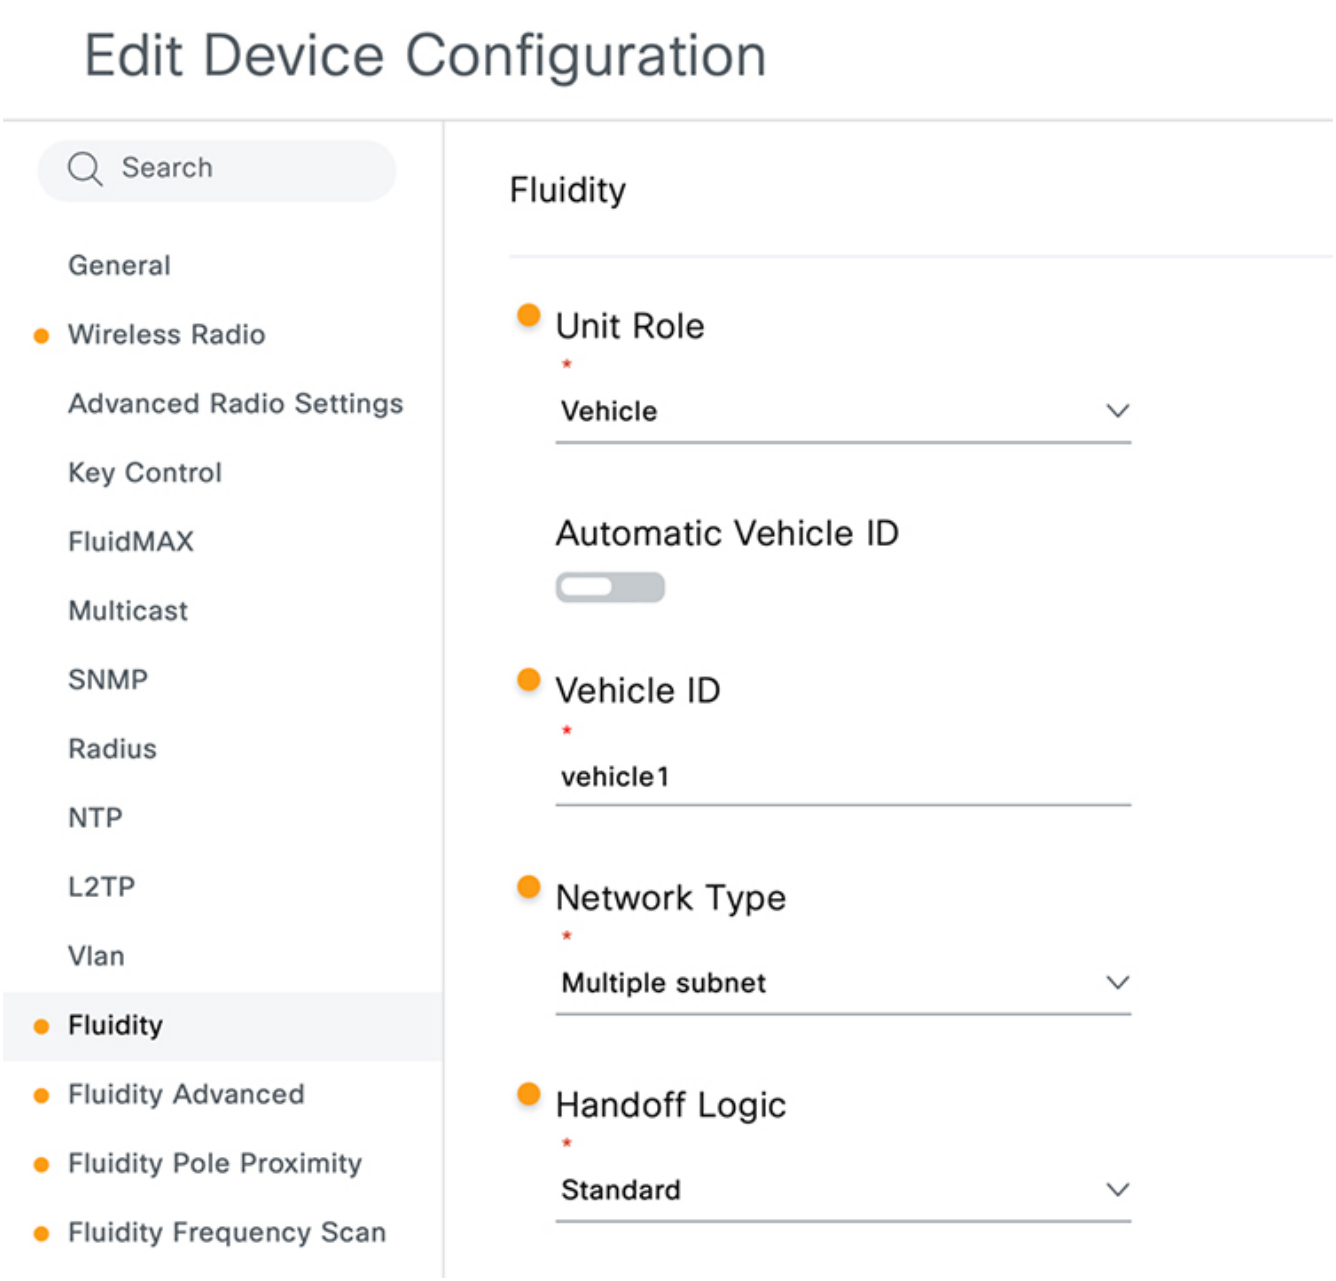


VLANs are also not supported in the L3 Fluidity mesh network and must not be configured.

Enabling fast failover on a pair of IW916x radios requires the use of a virtual hot-standby address which is configured on both radios. This virtual address is the default gateway for the onboard switched network.

# Quality of Service

In the Rail CBTC network, the vital traffic must be prioritized above all other traffic types. Each switch connected to vital equipment should have an access policy that matches that traffic as well as a class-map that matches on that access list. The policy-map should mark that traffic as **high priority** and put it in a priority queue while all other traffic is marked as **best effort**.

See the QoS Configuration Guide for more specifics.

The following is an example of matching on the MAC address for an IOT device used for vital traffic.

```
mac access-list extended axle-counter-1
 permit host 1111.2222.3333 any
```

A class-map for this ACL would look like this:

```
class-map match-all class-axle-counter
 match access-group name axle-counter-1
```

The input policy-map would look like this:

```
policy-map axle-counter-input_policy
 class class-axle-counter
  set cos 6
```

## Wireless

In the context of URWB, QoS only needs to be enabled on the radio for the QoS markings to be honored. As described in the Design Guide for URWB QoS, L2 Fluidity and L3 Fluidity require differing QoS strategies. L2 Fluidity translates the QoS marking of the original packet into the new VLAN header while in the mesh network. L3 Fluidity does not use a VLAN header while in the mesh network and instead copies the original QoS marking to the MPLS EXP section. Since the wayside access switches cannot match on this, the traffic from the onboard train network cannot be prioritized according to the original QoS markings. The wayside access switches can match on the MPLS ethertype which will give all traffic from the train the same priority treatment.

An access-list matching on the MPLS ethertype would look like this:

```
mac access-list extended l3_fluidity
 permit any any 0x8847 0x0
 permit any any 0x8848 0x0
```

# Acronyms and Initialisms

| Term | Definition |
|---|---|
| BGP | Border Gateway Protocol |
| CVD | Cisco Validated Design |
| DHCP | Dynamic Host Configuration Protocol |
| DNS | Domain Name System |
| HSRP | Hot Standby Router Protocol |
| IE | Industrial Ethernet |
| ISE | Identity Services Engine |
| L2TP | Layer 2 Tunneling Protocol |
| MAC | Media Access Control |
| NGFW | Next General Firewall |
| NGIPS | Next-Generation Intrusion Prevention System |
| NSF/SSO | Non-Stop Forwarding with Stateful Switchover |
| NTP | Network Time Protocol |
| OCC | Operational Control Center |
| pxGrid | Platform eXchange Grid |
| RADIUS | Remote Authentication Dial-In User Service |
| REP | Resilient Ethernet Protocol |
| SGTs | Scalable Group Tags |
| SGACL | Security Group-based Access Control List |
| UCS | Cisco Unified Computing System |
| VRF | virtual routing and forwarding |