



# Overview of Circuit/VC Discovery and Provisioning

---

- [Circuits/VCs Provisioning Overview, page 1](#)
- [Supported Carrier Ethernet VCs , page 2](#)
- [Supported Network Structure for Provisioning EVCs, page 5](#)
- [Supported Optical Circuits , page 6](#)
- [Supported Circuit Emulation Services, page 10](#)
- [Supported L3VPN Services, page 11](#)
- [Supported MPLS Traffic Engineering Services, page 12](#)
- [Supported Serial Services, page 13](#)
- [Circuit/VC Discovery Overview, page 14](#)

## Circuits/VCs Provisioning Overview

Cisco EPN Manager supports provisioning of circuits/VCs for various technologies such as Carrier Ethernet (CE), Optical/DWDM, L3VPN, Circuit Emulation, and MPLS Traffic Engineering. Mostly, a circuit spans across multiple devices. You must make configuration changes across multiple devices to provision a circuit. Cisco EPN Manager provides a Provisioning Wizard that allows you to make the required configuration changes across multiple devices that participate in a circuit.

The Provisioning Wizard collects all the required information in a step-by-step approach and generates the required configuration for all the devices. You can review the configurations generated for each device, and then choose to either make any changes in the service parameters or deploy the configurations to the devices.

The configuration changes are deployed to the participating devices as an 'atomic' transaction. Cisco EPN Manager does a best-effort attempt to either carry out all these operations together or does none at all. To implement the concept of 'atomic' transaction, Cisco EPN Manager has the rollback feature, which helps to recover from failures during provisioning.

When configuring multiple devices, if the configuration fails in any of the devices, Cisco EPN Manager does a best-effort to rollback the configuration changes made so far in all the participating devices. The device

configuration states are restored to the same state, which was there before the provisioning operation was attempted.

## Supported Carrier Ethernet VCs

In a Carrier Ethernet (CE) network, data is transported across point-to-point and multipoint-to-multipoint Ethernet Virtual Connections (EVCs) and Operator Virtual Connections (OVCs) according to the attributes and definitions of the various service types—that is, E-Line, E-LAN, E-Tree, and E-Access.

Each EVC type has a port-based service and a VLAN-based service. These are differentiated by the method for service identification used at the UNIs. EVCs using all to one bundling UNIs (port-based) are referred to as 'Private', while EVCs using UNIs that are that are service multiplexed (VLAN-based), are referred to as 'Virtual Private'

For E-Line, E-LAN, and E-Tree services, each EVC carries data in the form of CE service frames from UNI (User Network Interface) to UNI, where the UNI is the physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber. E-Access Operator Virtual Circuits (OVCs) allow service provider interconnections at the ENNI (External Network Network Interface), which is the physical demarcation point between the responsibility of two interconnecting Service Providers.

Each EVC can be configured with a rich set of attributes that include bandwidth profiles (Committed Information Rate - CIR, Excess Information Rate - EIR, Committed Burst Size - CBS, Excess Burst Size - EBS), multiple classes of service, application-oriented performance objectives, traffic management, forwarding rules, and so on.

Cisco EPN Manager supports the discovery and provisioning of the following EVC types, which are described in these topics:

- [E-Line, on page 3](#):
  - MPLS to the edge
  - Single-segment pseudowire
  - Ethernet access—local, G.8032, ICCP-SM
- [E-LAN, on page 3](#)
  - MPLS to the edge
  - Single-segment pseudowire
  - VPLS/H-VPLS with redundant pseudowire
  - Ethernet access—VPLS-based
- [E-Tree, on page 4](#)—MPLS to the edge
- [E-Access, on page 4](#)—MPLS to the edge

## Core Technology for Multipoint EVCs

The core technology for the E-LAN or E-Tree EVC can be either VPLS (Virtual Private LAN Services) or H-VPLS (Hierarchical VPLS).

- VPLS—A Layer 2 VPN technology that provides Ethernet-based multipoint-to-multipoint communication over MPLS networks. VPLS allows geographically dispersed sites to share an Ethernet broadcast domain by connecting sites through pseudowires. The network emulates a LAN switch or bridge by connecting customer LAN segments to create a single bridged Ethernet LAN.
- H-VPLS—Partitions the network into several edge domains that are interconnected using an MPLS core. The edge devices learn only of their local U-PE devices and therefore do not need large routing table support. The H-VPLS architecture provides a flexible architectural model that enables Ethernet multipoint and point-to-point Layer 2 VPN services, as well as Ethernet access to Layer 3 VPN services, enabling service providers to offer multiple services across a single high-speed architecture.

In E-TREE EVCs, H-VPLS supports redundancy. Two hubs operate as connectors through which all traffic passes. If the primary hub fails, traffic is switched to the backup hub. With H-VPLS as the core technology, there is no direct connection between the E-tree root and leaf. H-VPLS is used together with split-horizon capabilities to prevent leaf to leaf communication.

If VPLS is used as the core technology, redundancy is not supported and there is a direct connection between root and leaves. The hub is located in the root, meaning that the root assumes the role of the hub.

## E-Line

E-Line refers to an Ethernet service that is based on a point-to-point EVC. There are two types of E-Line VCs:

- Ethernet Private Line (EPL), which has the following characteristics:
  - Port-based
  - Uses a point-to-point EVC between two UNIs to provide a high degree of transparency such that service frames, headers, and most Layer 2 protocols are identical at both the source and destination UNI.
  - All to one bundling where all CE-VLAN IDs are bundled to one EVC. No service multiplexing.
- Ethernet Virtual Private Line (EVPL), which has the following characteristics:
  - VLAN-based
  - Uses a point-to-point EVC between two UNIs, but does not provide full transparency as with the EPL, that is, all Layer 2 control protocols are discarded at the UNI.
  - Allows for service multiplexing, which means that more than one EVC can be supported at the UNI.

## E-LAN

E-LAN refers to an Ethernet service that is based on a multipoint-to-multipoint EVC. There are two types of E-LAN VCs:

- Ethernet Private LAN (EP-LAN), which has the following characteristics:
  - Port-based
  - All-to-one bundling at the UNI

- Very transparent, no manipulation of CE-VLAN IDs and PCP bits
- EP-LAN multiport transparency is more complex than EPL
- Ethernet Virtual Private LAN (EVP-LAN), which has the following characteristics:
  - VLAN-based
  - Allows for service multiplexing and bundling

## E-Tree

An E-Tree VC is a rooted multipoint VC that connects a number of UNIs providing sites with hub and spoke multipoint connectivity. Each UNI is designated as either *root* or *leaf*. A root UNI can communicate with any leaf UNI. A leaf UNI can communicate only with a root UNI, not with another leaf UNI.

E-Tree VCs provide the separation between UNIs required to deliver a single service instance in which different customers (each having a leaf UNI) connect to an ISP which has one or more root UNIs. Having more than one root UNI is useful for load sharing and resiliency schemes.

There are two types of E-Tree VCs:

- Ethernet Private TREE (EP-TREE), which has the following characteristics:
  - Rooted multipoint, port-based.
  - All-to-one bundling at the UNI.
  - Simpler than typical hub and spoke configuration using multiple EPLs. Hub function is performed by the root UNI.
  - Provides CE-VLAN tag preservation and tunneling of key Layer 2 Control Protocols.
  - Supports CE-VLAN CoS preservation.
- Ethernet Virtual Private TREE (EVP-TREE), which has the following characteristics:
  - Rooted multipoint, VLAN-based.
  - Provides an alternative to multiple EVPLs multiplexed at the hub site.
  - Used in cases where one or more of the subscriber's UNIs also supports other services, e.g., EVPL or EVP-LAN.

## E-Access

An Ethernet Access service allows a service provider to construct an Operator Virtual Connection (OVC) between two customer sites where one of the sites is located outside of the service provider's own network. In such cases a service provider uses an E-Access service offered by a local wholesale access provider to reach the out-of-franchise UNI. The service provider connects to the E-Access service at an ENNI, and traffic is forwarded between the ENNI and the out-of-franchise UNI across an Operator Virtual Connection (OVC).

E-Access definitions include attributes related to the external interfaces, in this case, the ENNI and the UNI, as well as attributes related to the virtual Ethernet connection associating these external interfaces. E-Access

services use a point-to-point OVC to associate one OVC endpoint at an ENNI and one OVC endpoint at a UNI.

There are two types of E-Access VCs:

- Access EPL, which has the following characteristics:
  - Private or port-based
  - One OVC per UNI
  - All CE-VLAN IDs are mapped to the OVC
- Access EVPL, which has the following characteristics:
  - VLAN-based
  - Can be multiple OVCs per UNI
  - Multiple but not all CE-VLAN IDs are bundled to one OVC

## Supported Network Structure for Provisioning EVCs

Cisco EPN Manager can provision EVCs and OVCs over a mix of access networks. The endpoints can be configured directly on an MPLS router, an Ethernet Access switch, or an nV satellite attached to a Cisco ASR 9000 router. The EVCs may have endpoints in different Ethernet Access networks, in the same network, or on the same device. Cisco EPN Manager will configure as much as is needed to create the connectivity.

EVCs can be provisioned over the following networks:

- MPLS Domain—Cisco EPN Manager assumes that the managed network contains a single MPLS domain. Any router can communicate with any other router via a targeted LDP session. This enables the implementation of a point-to-point service between two routers with a single-segment pseudowire.
- Ethernet Access Network—Cisco EPN Manager supports EVC provisioning over Ethernet Access networks attached to a central MPLS domain. The networks are discovered by the system. EVCs can be provisioned over a G.8032 access ring or over ICCP-SM links. The access network can be:
  - A G.8032 ring. This should include a router to enable the creation of EVCs that cross the MPLS domain.
  - A G.8032 open ring, which means a sequence of links.
- Cisco ASR 9000 nV Satellite Topology—Cisco EPN Manager can configure EVCs on single-homed nV satellite devices attached to an Cisco ASR 9000 host.

To support service discovery and provisioning, Cisco EPN Manager must discover the topology in the access network. For successful discovery, the following prerequisites must be fulfilled:

- For ICCP-SM, LAG must be configured with LACP.
- For G.8032, CDP must be configured on the ring ports.

## Supported Optical Circuits

A circuit represents an end-to-end connection between two or more connection termination points (CTPs). A circuit consists of an alternating series of cross-connections and link connections. In its simplest form, a circuit consists of a single cross-connection (if the circuit is defined between two CTPs on the same NE). A circuit can be bidirectional or unidirectional, point-to-point or point-to-multipoint, and protected or unprotected.

ess supports the provisioning of Dense Wavelength Division Multiplexing (DWDM) optical channel (OCH) circuit types and Optical Transport Network (OTN) circuit types. The DWDM optical technology is used to increase bandwidth over existing fiber optic backbones. It combines and transmits multiple signals simultaneously at different wavelengths on the same fiber. In effect, one fiber is transformed into multiple virtual fibers.

Cisco EPN Manager supports the following optical circuits types:

- [Dense Wavelength Division Multiplexing \(DWDM\) Optical Channel \(OCH\) Circuit, on page 6](#)
  - [Optical Channel Network Connection \(OCHNC\) WSON, on page 6](#)
  - [Optical Channel Client Connection \(OCHCC\) WSON, on page 6](#)
  - [Optical Channel \(OCH\) Trail WSON, on page 7](#)
  - [Optical Channel \(OCH\) Trail Connecting NCS 1002 Devices, on page 7](#)
  - [Optical Channel \(OCH\) Trail User-to-Network Interface \(UNI\), on page 7](#)
- [Optical Transport Network \(OTN\) Circuit, on page 8](#)
  - [Optical Channel Data Unit User-to-Network Interface \(ODU UNI\), on page 8](#)
  - [Optical Channel Data Unit \(ODU\) Tunnel, on page 9](#)
  - [Optical Channel Payload Unit \(OPU\) Over Optical Channel Data Unit \(ODU\), on page 9](#)

## Dense Wavelength Division Multiplexing (DWDM) Optical Channel (OCH) Circuit

The following topics describe the different optical channel circuit types.

### Optical Channel Network Connection (OCHNC) WSON

OCHNC WSON circuits establish connectivity between two optical nodes on a specified C-band wavelength. The connection is made through the ports present on the wavelength selective switches, multiplexers, demultiplexer, and add/drop cards. In an OCHNC WSON circuit, the wavelength from a source OCH port ingresses to a DWDM system and then egresses from the DWDM system to the destination OCH port.

### Optical Channel Client Connection (OCHCC) WSON

OCHCC WSON circuits extend the OCHNC WSON to create an optical connection from the source client port to the destination client port of the TXP/MXP cards. An OCHCC WSON circuit represents the actual

end-to-end client service passing through the DWDM system. Each OCHCC WSON circuit is associated to a pair of client or trunk ports on the transponder (TXP), muxponder (MXP), GE\_XP (in layer-1 DWDM mode), 10GE\_XP (in layer-1 DWDM mode), or ITU-T line card. The OCHCC WSON circuits can manage splitter protection as a single protected circuit. However, for the Y-Cable protection, two OCHCC WSON circuits and two protection groups are required.

## Optical Channel (OCH) Trail WSON

OCH trail WSON circuits transport the OCHCC WSON circuits. The OCH trail WSON circuit creates an optical connection from the source trunk port to the destination trunk port of the Transponder (TXP), Muxponder (MXP), GE\_XP, 10GE\_XP, or ITU-T line card. The OCH trail WSON represents the common connection between the two cards, over which all the client OCHCC WSON circuits, SVLAN circuits or STS circuits are carried. Once an OCHCC WSON is created, a corresponding OCH Trail is automatically created. If the OCHCC WSON is created between two TXP, MXP, GE\_XP, or 10GE\_XP cards, two circuits are created in the CTC. These are:

- One OCHCC WSON (at client port endpoints)
- One OCH trail WSON (at trunk port endpoints)

If the OCHCC WSON is created between two TXPP or two MXPP cards, three circuits are created in the CTC. These are:

- One OCHCC WSON (at client port endpoints)
- Two OCH Trails WSON (at trunk port endpoints). One for the working and other for the protect trunk.

## Optical Channel (OCH) Trail Connecting NCS 1002 Devices

Cisco EPN Manager can discover an OCH Trail circuit from the source trunk port of an NCS 1002 device to the destination trunk port of another NCS 1002 device. The trunk port of each NCS 1002 must be connected via a manual link to the passive units of NCS 2K devices. Where the manual links are terminated, and OCH-NC circuit must be created as a pre-requisite between the ports of the passive units of NCS 2K network.



---

**Note** Provisioning is not supported for this type of optical circuits.

---

## Optical Channel (OCH) Trail User-to-Network Interface (UNI)

An OCH trail UNI circuit establishes connectivity between Cisco NCS 2000 series devices and Cisco NCS 4000 series devices. It provides an end-to-end configuration of DWDM network that consists of Cisco NCS 2000 series devices and terminates on a Cisco NCS 4000 series device. When an OCH trail UNI circuit is created in a Cisco NCS 4016 network element, a corresponding OCHNC circuit is created in the Cisco NCS 2006 network element.



---

**Note** You will not be able to modify or delete the OCHNC circuit.

---

## Optical Transport Network (OTN) Circuit

OTN specifies a digital wrapper, which is a method of encapsulating an existing frame of data, regardless of the native protocol, to create an optical data unit (ODU), similar to that used in SDH/SONET. OTN provides the network management functionality of SDH/SONET, but on a wavelength basis. A digital wrapper, however, is flexible in terms of frame size and allows multiple existing frames of data to be wrapped together into a single entity that can be more efficiently managed through a lesser amount of overhead in a multi-wavelength system.

The OTN specification includes framing conventions, non-intrusive performance monitoring, error correction (FEC), rate adaption, multiplexing mechanisms, ring protection, and network restoration mechanisms operating on a wavelength basis.

A key element of a digital wrapper is the forward error correction (FEC) mechanism that provides performance gains for improved margins and extended optical reach.

The OTN architecture is compliant to ITU-T G.872. An OTN circuit can be established statically or dynamically between ingress and egress nodes using Resource Reservation Protocol (RSVP) signaling. An OTN circuit is established and maintained as a label switched path (LSP) between the ingress and egress Label Switched Routers (LSRs) switched through transit LSRs. An LSP can be established as a soft permanent connection (SPC) when the request comes from the user interface.

Following are the types of OTN circuits:

- [Optical Channel Data Unit User-to-Network Interface \(ODU UNI\), on page 8](#)
- [Optical Channel Data Unit \(ODU\) Tunnel, on page 9](#)
- [Optical Channel Payload Unit \(OPU\) Over Optical Channel Data Unit \(ODU\), on page 9](#)

### Optical Channel Data Unit User-to-Network Interface (ODU UNI)

ODU is the transport container defined to carry client signals from network ingress to egress. The ODU provides a payload area for client data along with performance monitoring and fault management. The payload area of an ODU may contain a single non-OTN signal as a client or may contain multiple lower rate ODUs as clients. An ODU UNI circuit represents the actual end-to-end client service passing through the OTN architecture.

#### Open Ended ODU UNI

In an open-ended ODU UNI circuit, one or both end points may be connected to ODU subcontrollers, instead of client payload controllers.

Cisco EPN Manager supports three types of open-ended ODU UNIs:

- Only the source interface is an ODU subcontroller
- Only the destination interface is an ODU subcontroller
- Both source and destination interfaces are ODU subcontrollers

To create an open-ended ODU UNI circuit, you must configure ODU subcontrollers on devices before adding the devices to Cisco EPN Manager. Use the **controller oduk** command to configure ODU subcontrollers on devices.

**Example: Configure an ODU subcontroller on a Cisco NCS 4000 Device**

In this example, two ODU0 subcontrollers are configured to an ODU1 controller.

```
RP/0/RP0:router#conf
RP/0/RP0:router(config)# controller ODU10/1/0/1
RP/0/RP0:router(config-odul)# tsg 1.25G
RP/0/RP0:router(config-odul)# ODU0 tpn 1 ts 1
RP/0/RP0:router(config-odul)# ODU0 tpn 2 ts 2
RP/0/RP0:router(config-odul)#commit
```

To verify that the ODU subcontrollers are configured correctly on the device:

```
RP/0/RP0:router#sh controllers ODU0 ?
 0/1/0/0      ODU0 Interface Instance
 0/1/0/1/10   ODU0 Interface Instance
 0/1/0/1/20   ODU0 Interface Instance
R/S/I/P      Forward interface in Rack/Slot/Instance/Port format
```

After configuring the ODU subcontrollers on the device, you must add the device to Cisco EPN Manager . You can then verify that the ODU0 0/1/0/1/10 and ODU0 0/1/0/1/20 subcontrollers are available in the inventory.

**Optical Channel Data Unit (ODU) Tunnel**

ODU tunnel circuits transport the ODU UNIs. The ODU tunnel represents the common connection between two Cisco NCS 4000 series devices that are connected with Traffic Engineering (TE) links. Once an ODU UNI circuit is created, a corresponding ODU tunnel is automatically created.

**Optical Channel Payload Unit (OPU) Over Optical Channel Data Unit (ODU)**

OPU over ODU circuits provide a high-bandwidth point-to-point connection between two customer designated premises. Client signals are mapped over an OTN framing structure with in-band management through GCC0. These circuits uses ODU UNI circuits to carry client signals through the network. You need to perform the following tasks to create and provision an OPU over ODU circuit:

- Using Cisco EPN Manager , create an ODU UNI circuit. For information about how to create ODU UNI circuits, see [Create and Provision an OTN Circuit](#).
- Using Cisco Transport Controller (CTC), create LMP links and enable the links on the devices that you want to use in the OPU over ODU circuits. For information about how to create LMP, see the "DLP-K27 Create an LMP Using CTC" section in the [OTN and DWDM Configuration Guide for Cisco NCS 4000 Series](#).
- Using Cisco EPN Manager , create an OPU over ODU circuit with the LMP links enabled devices. For information about how to create OPU over ODU circuits, see [Create and Provision an OTN Circuit](#).

**Optical Channel Data Unit User-to-Network Interface (ODU UNI) Hairpin**

An ODU UNI Hairpin circuit is similar to an ODU UNI circuit, but it is created in the management plane and it is an intra node circuit, that is, the source and destination is the same device but with different interfaces. In this type of circuit, the connection is established between two clients or two ODU subcontrollers.

Cisco EPN Manager supports the following types of ODU UNI Hairpin circuits:

- Circuits without open-ended cross-connect—In this type of circuits, the interfaces of both source and destination are not OTU interfaces.

- Circuits with open-ended cross-connect on one side—In this type of circuits, the interface of either source or destination is an OTU interface.
- Circuits with open-ended cross-connect on both sides—In this type of circuits, the interface of both source or destination are OTU interfaces.

## Supported Circuit Emulation Services

Circuit Emulation (CEM) provides a protocol-independent transport over IP networks. It enables proprietary or legacy applications to be carried transparently to the destination, similar to a leased line. In traditional TDM networks, numerous physical circuits are maintained between geographically diverse locations to provide TDM transport. CEM allows TDM endpoints to connect across an IP/MPLS core. In CEM, endpoints are connected to the TDM circuits, but the circuits terminate at each local router that has the IP/MPLS connectivity available. The router then transports those TDM frames across the IP/MPLS core via Circuit Emulation (CEM) pseudowires (PWs) to the remote endpoint that also has the IP/MPLS connectivity available. Thus, the TDM endpoints can communicate as if they were directly connected by physical circuits.

Cisco EPN Manager supports the following CEM modes:

- Structure-Agnostic time-division multiplexing (TDM) over Packet (SAToP)—This is the unstructured mode in which the incoming TDM data is considered as an arbitrary bit stream. It disregards any structure that may be imposed on the bit stream. SAToP encapsulates the TDM bit streams as pseudowire (PWs) over PSN.
- Circuit Emulation over Packet (CEP)—This mode is used to emulate Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) circuits and services over MPLS. To transport SONET/SDH circuits through a packet-oriented network, the Synchronous Payload Envelope (SPE) or Virtual Tributary (VT) is broken into fragments. A CEP header and optionally an RTP header are prepended to each fragment.
- Circuit Emulation Service over Packet Switched Network (CESoPSN)—This is the structured mode in which the structured TDM signals are encapsulated as PWs and transmitted over PSN. It selects only valid timeslots and disregards the idle timeslots for transmission. Thus, CESoPSN can save utilized bandwidth.

Cisco EPN Manager supports the following CEM service types depending on the rate at which the circuits can transmit data:

- DS0—A basic digital signal with transmission data rate of up to 64 Kbps.
- T1 and E1—Digital Signal (DS) is known as T-carrier in North America, South Korea, and Japan and as E-carrier in the rest of the world. The T1 circuit has a transmission data rate of up to 1.544 Mbps. The E1 circuit has a transmission data rate of up to 1.984 Mbps in framed mode and 2.048 Mbps in unframed mode. The T1 or E1 circuit can transport 24 DS0 level channels within its payload.
- T3 and E3—The T3 circuit has a transmission data rate of up to 44.736 Mbps. The E3 circuit has a transmission data rate of up to 34.368 Mbps. The T3 or E3 circuit can transport 672 DS0 level channels and 28 DS1 level channels within its payload.
- VT 1.5—A virtual tributary network line with transmission data rate of up to 1.728 Mbps.
- VT 2—A virtual tributary network line with transmission data rate of up to 2.304 Mbps.
- STS1—A synchronous transport signal with transmission data rate of up to 51.84 Mbps.

- STS3—A synchronous transport signal with transmission data rate of up to 155.52 Mbps.
- STS12—A synchronous transport signal with transmission data rate of up to 622.08 Mbps.
- STS48—A synchronous transport signal with transmission data rate of up to 2488.32 Mbps.
- STS192—A synchronous transport signal with transmission data rate of up to 9953.28 Mbps.

## Supported L3VPN Services

An MPLS Layer 3 VPN creates a private IP network. The customer connects to the network via customer edge (CE) routers, which act as IP peers of provider edge (PE) routers.

### Virtual Routing and Forwarding (VRFs)

On the PE, Virtual Routing and Forwarding (VRF) instances act as virtual IP routers dedicated to forwarding traffic for the L3VPN service. The VRFs learn the routes to each other via the Multi-Protocol Border Gateway Protocol (MP-BGP), and then forward traffic using MPLS.

A VPN is comprised of at least one but typically several VRFs. Cisco EPN Manager uses the VPN ID to discover which VRFs together form a single VPN. If Cisco EPN Manager discovers an existing network where no VPN ID has been provisioned, it takes all VRFs with the same name and associates them into one VPN. For VPNs created using Cisco Prime Provisioning, which uses a naming convention with version number prefixes and different suffixes, Cisco EPN Manager will recognize the different VRFs as belonging to one VPN.

In general there is a regular expression which can be configured to allow for varying naming convention.

### Route Targets (RTs)

The connections between VRFs are defined using Route Targets (RTs) that are imported and exported by the VRFs. Cisco EPN Manager makes it easy to set up a full mesh of connections, and automatically allocates the route target to be used. The route target consists of a prefix which is either an AS number or an IPv4 address, for example, a full mesh prefix, 100 [681682]. The prefix can be selected from the existing BGP autonomous system (AS) numbers in the network, or it can be entered manually. The second number following the prefix is allocated automatically by Cisco EPN Manager.

Alternatively or in addition to the full mesh, it is possible to manually select route targets. During VPN creation, there is an initial screen where you type in the route targets to be used within a VPN, and then for each VRF you can select which route targets you import and export. You also specify for which address family (IPv4 or IPv6) you will use the route target. This can be used for example to configure extranets, by importing route targets used in other VPNs.

### Route Redistribution

The routes that are exchanged between the PE and the CE have to be redistributed into the MP-BGP routing protocol so that remote endpoints can know which prefixes can be reached at each VRF. To control route redistribution, Cisco EPN Manager allows you to define the required protocol (Static, Connected, or RIP), the protocol's metric value, and optionally the applicable route policy.

### Endpoints

Cisco EPN Manager supports the creation of IP endpoints on Ethernet subinterfaces. It supports selecting untagged encapsulation, or specifying an outer and optionally an inner VLAN, with 802.1q or 802.1ad encapsulation. You can specify both IPv4 and IPv6 addresses at an endpoint. You can also specify the BGP neighbor details to provision BGP neighbors between the CE and PE.

For information on how to provision L3VPN service using Cisco EPN Manager, see, [Provision L3VPN Services](#).

## Supported MPLS Traffic Engineering Services

In traditional IP networks, packets are forwarded on a per-hop basis where a route lookup is performed on each router from source to destination. The destination-based forwarding mechanism leads to suboptimal use of available bandwidth between a pair of routers in the network. Mostly, the suboptimal paths are under-utilized in IP networks. To avoid packet drops due to inefficient use of available bandwidth and to provide better performance, traffic engineering (TE) is implemented. TE directs the traffic that is destined to follow the optimal path to a suboptimal path, thus enabling better bandwidth utilization between a pair of routers.

Multiprotocol Label Switching (MPLS) is an integration of Layer 2 and Layer 3 technologies. In an MPLS domain, unique labels are assigned to data packets and the packets are forwarded based on these labels. It avoids the complex lookup in a routing table. MPLS creates a VC switching function to provide similar performance on the IP-based network services as compared to those delivered over traditional networks such as Frame Relay or Asynchronous Transfer Mode (ATM).

By making traditional Layer 2 features available to Layer 3, MPLS enables traffic engineering. MPLS TE enables an MPLS backbone to replicate and expand the TE capabilities of Layer 2 over Layer 3.

MPLS TE uses Resource Reservation Protocol (RSVP) to establish and maintain label-switched path (LSP) across the backbone. The path that an LSP uses is based on the LSP resource requirements and network resources, such as bandwidth and link attributes. Available resources are flooded by means of extensions to a link-state-based Interior Gateway Protocol (IGP). Cisco EPN Manager supports OSPF as the IGP to flood the available bandwidth and link status information across the network. Based on this information, the ingress (headend) router gathers information on all the available resources in the network along with the topology to define tunnels through the network between a set of MPLS-enabled routers. This is called as constraint-based routing. When a shortest path is over-utilized, the IGP automatically routes the traffic to these LSPs. You can also provision explicit paths for MPLS TE tunnels.

Cisco EPN Manager provides full path protection mechanism for MPLS TE tunnels against path, link, and node failures. A secondary LSP is established to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the source router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.

Cisco EPN Manager supports the following MPLS TE service types:

- [Unidirectional TE Tunnel, on page 12](#)
- [Bidirectional TE Tunnel, on page 13](#)
- [Layer 3 Link, on page 13](#)

## Unidirectional TE Tunnel

MPLS TE tunnels are unidirectional tunnels that connect a pair of LSRs. Once the unidirectional tunnel is created, a label is assigned for the tunnel that corresponds to a specific path in a MPLS network. The traffic is routed through the tunnel. You must create another unidirectional tunnel between the same routers to route the return traffic. For example, router A is the head end and router B is the tail end of tunnel 1, which is a unidirectional tunnel. You must create another unidirectional tunnel, say tunnel 2 with router B as the head end and router A as the tail end.

## Bidirectional TE Tunnel

Two unidirectional TE tunnels established between a pair of LSRs that are connected to each other, can be bound together to form a bidirectional co-routed TE tunnel. The binding of unidirectional tunnels is based on the source and destination addresses, global ID, association ID, and association address of the tunnels. For example, router A and router B that are connected by two unidirectional tunnels, tunnel C and tunnel D, can be bound together to form a bidirectional TE tunnel only if the following conditions are met:

- The source address of tunnel C is the destination address of tunnel D and vice versa.
- The global ID, association ID, and association address of tunnel C and D are the same. The association ID and association address for the tunnels are system-defined and you need to assign a global ID for the tunnels.

Bidirectional TE tunnels inherit the security features of RSVP-TE.

## Layer 3 Link

To enable traffic engineering links between two devices, you need to configure the following on both ends of the devices:

- Loopback interface
- Ethernet interface
- BDI Interface
- OSPF, RSVP, and MPLS
- IS-IS and BGP

You can perform these configurations using the Layer 3 link provisioning feature in Cisco EPN Manager .

## Supported Serial Services

In a serial communication, the serial port sends and receives bytes of information one bit at a time. The serial communication can be used over longer distances. The cabling between devices can extend up to 1200 meters. Serial communication is used to transmit ASCII data. Communication is completed using three transmission lines—ground, transmit, and receive. Since serial is asynchronous, the port is able to transmit data on one line while receiving data on another. Other lines are available for handshaking, but are not required.

Cisco EPN Manager supports RS232 serial service type. RS-232 is a standard communication protocol that links devices in a network to allow serial data exchange. It defines the voltage for the path used for data exchange between the devices. It specifies common voltage, signal level, common pin wire configuration, and minimum amount of control signals.

## Circuit/VC Discovery Overview

Cisco EPN Manager uses the Service Discovery feature to automatically discover circuits/VCs existing in the network. Ensure that the Service Discovery feature is enabled under **Administration > System Settings**. See [Enable and Disable Service Discovery](#).

Circuit/VC discovery depends on device-level inventory discovery, and consists of two parts:

- Resource facing service (RFS) discovery—The RFS represents the relations between resources on different devices. During RFS discovery, the system creates device-level objects and network-level objects. Device-level RFS objects represent the circuit/VC configuration parts of the device-level configuration. Network-level RFS objects aggregate device or other network-level objects to represent network-level entities.
- Customer facing service (CFS) matching—The CFS represents the customer facing data for a circuit/VC. The CFS is derived from discovered RFS and represents the endpoints of the circuit/VC in the network. During CFS discovery, the system creates CFS objects for the discovered RFS objects.

Discovery is an ongoing process in Cisco EPN Manager. When you first start using Cisco EPN Manager, the circuits/VCs that exist in the network are discovered. Later, when you start provisioning circuits/VCs using the Provisioning Wizard, Cisco EPN Manager will discover the provisioned circuits/VCs and will search for a match between the resources used in the circuit/VC and the resources discovered from the network. When a match is found between a discovered circuit/VC and a provisioned circuit/VC, information from the provisioned CFS is copied into the discovered CFS.

Cisco EPN Manager allows you to compare between the provisioned and discovered versions to identify changes that might have been made in the device configurations and you can do a reconciliation, if necessary. See [Compare and Reconcile Provisioned and Discovered Versions of a Circuit/VC](#)