

# Chapter 5

## Layer 2 VPN Configuration Guidelines

To configure Layer 2 virtual private network (VPN) functionality, you must enable Layer 2 VPN support on the provider edge (PE) router. You must also configure PE routers to distribute routing information to the other PE routers in the VPN and configure the circuits between the PE routers and the customer edge (CE) routers.

Each Layer 2 VPN is configured under a routing instance of type l2vpn. An l2vpn routing instance can transparently carry Layer 3 traffic across the service provider's network. As with other routing instances, all logical interfaces belonging to a Layer 2 VPN routing instance are listed under that instance.

The configuration of the CE routers is not relevant to the service provider. The CE routers only need to provide appropriate Layer 2 circuits (with appropriate circuit identifiers, such as data-link connection identifier [DLCI], virtual path identifier/virtual channel identifier [VPI/VCI], or virtual local area network identifier [VLAN ID]) to send traffic to the PE router.

To configure Layer 2 VPNs, you include statements at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit]
routing-instances {
  routing-instance-name {
    description text;
    instance-type l2vpn;
    interface interface-name;
    route-distinguisher (as-number:id | ip-address:id);
    vrf-export [ policy-name ];
    vrf-import [ policy-name ];
    protocols {
      l2vpn {
        (control-word | no-control-word);
        encapsulation type;
        traceoptions {
          file filename <replace> <size size> <files number> <nostamp>;
          flag flag <flag-modifier> <disable>;
        }
        site site-name {
          site-identifier identifier;
          interface interface-name {
            remote-site-id remote-site-id;
          }
        }
      }
    }
  }
}
```

For Layer 2 VPNs, only some of the statements in the [edit routing-instances] hierarchy are valid. For the full hierarchy, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

In addition to these statements, you must configure Multiprotocol Label Switching (MPLS) label-switched paths (LSPs) between the PE routers, internal Border Gateway Protocol (IBGP) sessions between the PE routers, and an interior gateway protocol (IGP) on the PE and provider routers.

By default, Layer 2 VPNs are disabled.

Many of the configuration procedures for Layer 2 VPNs are identical to the procedures for Layer 3 VPNs and VPLS. These procedures are described in detail in Chapter 5, “VPN Configuration Guidelines,” on page 9 and include the following:

- Enable a Signaling Protocol on the PE Routers on page 10
- Configure an IGP on the PE and Provider Routers on page 13
- Configure an IBGP Session between PE Routers on page 14
- Configure a VPN Routing Instance on the PE Routers on page 15
- Configure Graceful Restart on page 27

This chapter describes the following tasks that are specific to configuring Layer 2 VPNs:

- Configure the Connections to the Local Site on page 41
- Configure CCC Encapsulation on Interfaces on page 45
- Configure TCC Encapsulation on Interfaces on page 46
- Configure Layer 2 VPN Policing on Interfaces on page 47
- Disable the Control Word for Layer 2 VPNs on page 48

## Configure the Connections to the Local Site

For each local site, the PE router advertises a set of VPN labels to the other PE routers servicing the Layer 2 VPN. The VPN labels constitute a single block of contiguous labels; however, to allow for reprovisioning, more than one such block can be advertised. Each label block consists of a label base, a range (the size of the block), and a remote site ID that identifies the sequence of remote sites that connect to the local site using this label block (the remote site ID is the first site identifier in the sequence). The encapsulation type is also advertised along with the label block.

To configure the connections to the local site on the PE router, perform the following tasks:

Configure the Site on page 41

Configure the Remote Site ID on page 42

Configure the Encapsulation Type on page 43

Trace Layer 2 VPN Traffic and Operations on page 44

### **Configure the Site**

All the Layer 2 circuits provisioned for a local site are listed as the set of logical interfaces (using the interface statement) within the site statement.

On each PE router, you must configure each site that has a circuit to the PE router. To do this, include the site statement at the [edit routing-instances *routing-instance-name* protocols l2vpn] hierarchy level:

```
[edit routing-instances routing-instance-name protocols l2vpn]
site site-name {
  site-identifier identifier;
  interface interface-name {
    remote-site-id remote-site-ID;
  }
}
```

You must configure the following for each site:

site—Name of the site.

site-identifier—Unsigned 16-bit number greater than zero that uniquely identifies the site. The site identifier should correspond to a remote site ID configured on another site within the same VPN.

interface—The name of the interface and, optionally, a remote site ID for remote site connections. See “Configure the Remote Site ID” on page 42.

## Configure the Remote Site ID

The remote site ID allows you to configure a sparse Layer 2 VPN topology. This means that each site does not have to connect to all the other sites in the VPN, making it unnecessary to allocate circuits for all the remote sites. Remote site IDs are particularly important if you configure a topology more complicated than full-mesh, such as a hub-and-spoke topology.

The remote site ID (configured with the `remote-site-id` statement) corresponds to the site ID (configured with the `site-identifier` statement) configured at a separate site.

For example, a configuration for Router PE1 connected to Router CE1 is as follows:

```
site-identifier 1;
interface so-0/0/0 {
  remote-site-id 2;
}
```

The configuration for Router PE2 connected to Router CE2 is then as follows:

```
site-identifier 2;
interface so-0/0/1 {
  remote-site-id 1;
}
```

Configure the `remote-site-id` statement at the `[edit routing-instances routing-instance-name protocols l2vpn site site-name interface interface-name]` hierarchy level:

```
[edit routing-instances routing-instance-name protocols l2vpn]
site site-name {
  interface interface-name {
    remote-site-id remote-site-ID;
  }
}
```

If you do not explicitly configure the `remote-site-id` statement for the interface configured at the `[edit routing-instances routing-instance-name protocols l2vpn site site-name]` hierarchy level, a remote site ID is assigned to that interface.

The remote site ID for an interface is automatically set to 1 higher than the remote site ID for the previous interface. The order of the interfaces is based on their `site-identifier` statements. For example, if the first interface in the list does not have a remote site ID, its ID is set to 1. The second interface in the list has its remote site ID set to 2, and the third has its remote site ID set to 3. The remote site IDs of any interfaces that follow are incremented in the same manner if you do not explicitly configure them.

## Configure the Encapsulation Type

The encapsulation type you configure at each Layer 2 VPN site varies depending on which Layer 2 protocol you choose to configure. You need to use the same protocol at each Layer 2 VPN site if you configure ethernet-vlan as the encapsulation type. You do *not* need to use the same protocol at each Layer 2 VPN site if you configure any of the following encapsulation types:

atm-aal5—ATM Adaptation Layer (AAL/5)

atm-cell—ATM cell relay

atm-cell-port-mode—ATM cell relay port promiscuous mode

atm-cell-vc-mode—ATM virtual circuit (VC) cell relay non-promiscuous mode

atm-cell-vp-mode—ATM virtual path (VP) cell relay promiscuous mode

cisco-hdlc—Cisco Systems-compatible High-level Data Link Control (HDLC)

ethernet—Ethernet

ethernet-vlan—Ethernet VLAN

frame-relay—Frame Relay

interworking—Layer 2.5 interworking VPN

ppp—Point-to-Point Protocol (PPP)

If you configure different protocols at your Layer 2 VPN sites, you need to configure a TCC encapsulation type. For more information, see “Configure TCC Encapsulation on Interfaces” on page 46.



**Note**

Any Layer 2 VPNs configured with the atm-cell-port-mode, atm-cell-vc-mode, or atm-cell-vp-mode encapsulations on a router with JUNOS 5.6 later cannot interoperate with a router running JUNOS 5.5 or earlier.

To configure the Layer 2 protocol accepted by the PE router, specify the encapsulation type by including the encapsulation statement at the [edit routing-instances *routing-instance-name* protocols l2vpn] hierarchy level:

```
[edit routing-instances routing-instance-name protocols l2vpn]
encapsulation type
```

## Trace Layer 2 VPN Traffic and Operations

To trace Layer 2 VPN protocol traffic, you can specify options in the Layer 2 VPN traceoptions statement at the [edit routing-instances *routing-instance-name* protocols l2vpn] hierarchy level:

```
[edit routing-instances routing-instance-name protocols l2vpn]
traceoptions {
  file filename <replace> <size size> <files number> <no-stamp>;
  flag flag <flag-modifier> <disable>;
}
```

The following trace flags display the operations associated with Layer 2 VPNs:

all—All Layer 2 VPN tracing options.

connections—Layer 2 VPN connections (events and state changes).

error—Error conditions.

nri—Layer 2 VPN advertisements received or sent using BGP.

route—Trace routing information.

topology—Layer 2 VPN topology changes caused by reconfiguration or advertisements received from other PE routers using BGP.

## Disable Normal TTL Decrementing for VPNs

To diagnose networking problems related to VPNs, it can be useful to disable normal time-to-live (TTL) decrementing. In JUNOS, you can do this with the no-propagate-ttl and no-decrement-ttl statements. However, when tracing VPN traffic, only the no-propagate-ttl statement is effective.

For the no-propagate-ttl statement to have an effect on VPN behavior, you need to clear the PE-router-to-PE-router BGP session, or disable and then enable the VPN routing instance.

For more information about the no-propagate-ttl and no-decrement-ttl statements, see the *JUNOS Internet Software Configuration Guide: MPLS Applications*.

## Configure CCC Encapsulation on Interfaces

You need to specify a circuit cross-connect (CCC) encapsulation type for each PE-router-to-CE-router interface running a Layer 2 VPN. This encapsulation type should match the encapsulation type configured under the routing instance. See “Configure the Encapsulation Type” on page 43 for information about how to configure the encapsulation type under the routing instance.

To configure the CCC encapsulation type, include the following statements at the [edit interfaces] hierarchy level:

```
[edit]
interfaces {
  interface name {
    encapsulation ccc-encapsulation-type;
    unit unit number {
      encapsulation ccc-encapsulation-type;
    }
  }
}
```

You configure the encapsulation type at the [edit interfaces] hierarchy level differently than you do at the [edit routing-instance] hierarchy level. For example, you specify the encapsulation as frame-relay at the [edit routing-instances] hierarchy level and as frame-relay-ccc at the [edit interfaces] hierarchy level.

You can run both standard Frame Relay and CCC Frame Relay on the same device. If you specify Frame Relay encapsulation (frame-relay-ccc) for the interface, you should also configure the encapsulation at the [edit interfaces *interface name* unit *unit-number*] hierarchy level as frame-relay-ccc. Otherwise, the logical interface unit defaults to standard Frame Relay.

For more information on how to configure interfaces and interface encapsulations, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

## Configure TCC Encapsulation on Interfaces

Also known as Layer 2.5 VPNs, the translation cross-connect (TCC) encapsulation types allow you to configure different encapsulation types at the ingress and egress of a Layer 2 VPN. For example, a CE router at the ingress of a Layer 2 VPN circuit can send traffic as Frame Relay. A CE router at the egress of that circuit can receive the traffic as ATM.

The configuration for TCC encapsulation types is similar to the configuration for CCC encapsulation types. Specify a TCC encapsulation type for each PE-router-to-CE-router interface running a Layer 2 VPN. This encapsulation type should match the encapsulation type configured under the routing instance. See “Configure the Encapsulation Type” on page 43 for information about how to configure the encapsulation type under the routing instance.

To configure the TCC encapsulation type, include the following statements at the [edit interfaces] hierarchy level:

```
[edit]
interfaces {
  interface name {
    encapsulation tcc-encapsulation-type;
    unit unit number {
      encapsulation tcc-encapsulation-type;
    }
  }
}
```

You configure the encapsulation type at the [edit interfaces] hierarchy level differently from the [edit routing-instance] hierarchy level. For example, you specify the encapsulation as frame-relay at the [edit routing-instances] hierarchy level and as frame-relay-tcc at the [edit interfaces] hierarchy level.

For Layer 2.5 VPNs employing an Ethernet interface as the TCC router, you can configure an Ethernet translational cross-connect (TCC) or an extended VLAN TCC.

To configure an Ethernet TCC or an extended VLAN TCC, include the proxy and remote statements at the [edit interfaces *interface-name* unit *logical-unit-number* family tcc] hierarchy level:

```
[edit interfaces interfaces interface-name unit logical-unit-number family tcc]
proxy {
  inet-address address;
}
remote {
  (inet-address | mac-address) address;
}
```

The proxy inet-address address statement defines the IP address for which the TCC router is proxying.

The remote (inet-address | mac-address) statement defines the location of the remote router.

Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on one-port Gigabit Ethernet, two-port Gigabit Ethernet, four-port Gigabit Ethernet, and four-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 routing node.

For more information on how to configure interfaces and interface encapsulations, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

## Configure Layer 2 VPN Policing on Interfaces

You can use policing to control the amount of traffic flowing over the interfaces servicing a Layer 2 VPN. If policing is disabled on an interface, all the available bandwidth on a Layer 2 VPN tunnel can be used by a single CCC or TCC interface.

For more information about the policer statement, see the *JUNOS Internet Software Configuration Guide: Policy Framework*.

If you configure CCC encapsulation, then include the policer statement at the [edit interfaces *interface-name* unit *unit-number* family ccc] hierarchy level to enable Layer 2 VPN policing on an interface:

```
[edit]
interfaces interface-name {
  encapsulation encapsulation-type;
  unit 0 {
    family ccc {
      policer {
        input policer-template-name;
        output policer-template-name;
      }
    }
  }
}
```

If you configure TCC encapsulation, then include the policer statement at the [edit interfaces *interface-name* unit *unit-number* family tcc] hierarchy level to enable Layer 2 VPN policing on an interface:

```
[edit]
interfaces interface-name {
  encapsulation encapsulation-type;
  unit 0 {
    family tcc {
      policer {
        input policer-template-name;
        output policer-template-name;
      }
    }
  }
}
```

For information about how to configure the encapsulation type, see “Configure the Encapsulation Type” on page 43.

## Disable the Control Word for Layer 2 VPNs

The emulated VC encapsulation for Layer 2 VPNs is accomplished by adding a 4-byte control word between the Layer 2 protocol data unit (PDU) being transported and the VC label that is used for demultiplexing. Various networking formats (ATM, Frame Relay, Ethernet, and so on) use the control word in a variety of ways.

JUNOS software does not support the control word for any networking format, meaning that it is not fully compliant with the Internet draft in cases where the control word is mandatory. To be minimally compliant with the Internet draft, JUNOS supports a null control word (a control word of all zeros). This null control word is configured by default. If JUNOS receives a packet with a control word attached, the control word is discarded before the packet is forwarded to its destination.

JUNOS 5.5 and earlier releases do not support the control word at all. If you have configured Layer 2 VPNs on a network where some routers are running the current JUNOS release and some routers are running JUNOS 5.5 or earlier releases, you need to disable the control word on the routers running JUNOS 5.6 and later releases. To disable the control word, include the `no-control-word` statement at the [edit routing-instances *routing-instance-name* protocols l2vpn] hierarchy level on the routers running the current JUNOS release:

```
[edit routing-instances routing-instance-name protocols l2vpn]
no-control-word;
```

This is not necessary when configuring Layer 2 circuits. For more information, see “Disable the Control Word for Layer 2 Circuits” on page 372.