

Chapter 23

Layer 2 Circuit Configuration Guidelines

To configure a Layer 2 circuit, include the `l2circuit` statement:

```
l2circuit {
  local-switching {
    interface interface-name {
      description text;
    }
    end-interface {
      interface interface-name;
      protect-interface interface-name;
    }
  }
  neighbor address {
    interface interface-name {
      community community-name;
      (control-word | no-control-word);
      description text;
      mtu mtu-number;
      protect-interface interface-name;
      virtual-circuit-id identifier;
    }
  }
  traceoptions {
    file filename <replace> <size size> <files number> <nostamp>;
    flag flag <flag-modifier> <disable>;
  }
}
```

You can include the `l2circuit` statement at the following hierarchy levels:

[edit protocols]

[edit logical-routers *logical-router-name* protocols]

This chapter describes how to configure Layer 2 circuits, discussing the following topics:

- Configuring Interfaces for Layer 2 Circuits on page 460
- Configuring Local Interface Switching on page 464
- Configuring the Virtual Circuit ID on page 465
- Configuring LDP for Layer 2 Circuits on page 465
- Configuring Layer 2 Circuit Policies on page 465
- Configuring the Control Word for Layer 2 Circuits on page 469
- Configuring ATM Trunking on Layer 2 Circuits on page 470
- Configuring Bandwidth Allocation and Call Admission Control on page 472
- Tracing Layer 2 Circuit Creation and Changes on page 474

Configuring Interfaces for Layer 2 Circuits

The following sections describe how to configure interfaces for Layer 2 circuits:

- Configuring the Neighbor and Interface on page 461
- Configuring the Protect Interface on page 461
- Configuring the Interface Encapsulation Type for Layer 2 Circuits on page 462
- Configuring ATM2 IQ Interfaces for Layer 2 Circuits on page 462
- Configuring the MTU Advertised for a Layer 2 Circuit on page 463

Configuring the Neighbor and Interface

Each Layer 2 circuit is represented by the logical interface connecting the local provider edge (PE) router to the local customer edge (CE) router. All the Layer 2 circuits using a particular remote PE router designated for remote CE routers are listed under the neighbor statement (“neighbor” designates the PE router). Each neighbor is identified by its IP address and is usually the end-point destination for the label-switched path (LSP) tunnel transporting the Layer 2 circuit.

To configure a PE router as a neighbor, include the neighbor statement:

```
neighbor address { ... }
```

You can include the neighbor statement at the following hierarchy levels:

```
[edit protocols l2circuit]
```

```
[edit logical-routers logical-router-name protocols l2circuit]
```

To configure an interface, include the interface statement:

```
interface interface-name { ... }
```

You can include the interface statement at the following hierarchy levels:

```
[edit protocols l2circuit neighbor address]
```

```
[edit logical-routers logical-router-name protocols l2circuit neighbor address]
```

Configuring the Protect Interface

You can configure a protect interface for the logical interface linking a virtual circuit to its destination, whether the destination is remote or local. A protect interface provides a backup for the protected interface in case of failure. Network traffic uses the primary interface only so long as the primary interface functions. If the primary interface fails, traffic is switched to the protect interface. The protect interface is optional.

To configure the protect interface, include the protect-interface statement:

```
protect-interface interface-name;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

For an example of how to configure a protect interface for a Layer 2 circuit, see “Layer 2 Circuits Example” on page 475.

Configuring the Interface Encapsulation Type for Layer 2 Circuits

The Layer 2 encapsulation type is carried in the Label Distribution Protocol (LDP) forwarding equivalence class (FEC). You can configure either circuit cross-connect (CCC) or translational cross-connect (TCC) encapsulation types for Layer 2 circuits. For more information, see the *JUNOS MPLS Applications Configuration Guide*.

To configure the interface encapsulation for a Layer 2 circuit, include the encapsulation statement:

```
encapsulation encapsulation-type;
```

You can include the encapsulation statement at the following hierarchy levels:

```
[edit interfaces interface-name]
```

```
[edit logical-routers logical-router-name interfaces interface-name]
```

Configuring ATM2 IQ Interfaces for Layer 2 Circuits

You can configure Asynchronous Transfer Mode 2 (ATM2) intelligent queuing (IQ) interfaces for Layer 2 circuits by using Layer 2 circuit ATM Adaptation Layer 5 (AAL5) transport mode, Layer 2 circuit ATM cell relay mode, and the Layer 2 circuit ATM trunk mode.

The configuration statements are as follows:

```
atm-l2circuit-mode aal5
```

```
atm-l2circuit-mode cell
```

```
atm-l2circuit-mode trunk
```

For more information on these statements and configuring ATM2 IQ interfaces, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

The JUNOS software implementation of sequence number processing for Layer 2 circuit ATM cell relay mode and Layer 2 circuit AAL5 mode differs from that described in the Internet draft *draft-martini-l2circuit-encap-mpls-version.txt*, *Frame Relay Encapsulation over Pseudo-Wires*.

The JUNOS software implementation has the following differences:

1. A packet with a sequence number of 0 is treated as out of sequence.
2. A packet that does not have the next incremental sequence number is considered out of sequence.

When out-of-sequence packets arrive, the expected sequence number for the neighbor is set to the sequence number in the Layer 2 circuit control word.

Configuring the MTU Advertised for a Layer 2 Circuit

By default, the MTU used to advertise a Layer 2 circuit is determined by taking the interface MTU for the associated physical interface and subtracting the encapsulation overhead for sending IP packets based on the encapsulation.

However, encapsulations that support multiple logical interfaces (and multiple Layer 2 circuits) rely on the same interface MTU (since they are all associated with the same physical interface). This can prove to be a limitation for VLAN Layer 2 circuits using the same Ethernet interface or for Layer 2 circuit DLCIs using the same Frame Relay interface.

This can also affect multivendor environments. For example, if you had three PE devices supplied by different vendors and one of the devices could only support an MTU of 1500, even if the other devices could support larger MTUs you would have to configure the MTU as 1500 (the smallest MTU of the three PE devices).

You can explicitly configure what MTU is advertised for a Layer 2 circuit, even if the Layer 2 circuit is sharing a physical interfaces with other Layer 2 circuits. When you explicitly configure an MTU for a Layer 2 circuit, be aware of the following:

An explicitly configured MTU is signalled to the remote PE device. The configured MTU is also compared to the MTU received from the remote PE device. If there is a conflict, the Layer 2 circuit is taken down.

If you configure an MTU for an ATM cell relay interface on an ATM II PIC, the configured MTU is used to compute the cell bundle size advertised for that Layer 2 circuit rather than using the default interface MTU.

A configured MTU is used only in the control plane. It is not enforced in the data plane. You need to ensure that the CE device for a given Layer 2 circuit uses the correct MTU for data transmission.

To configure the MTU for a Layer 2 circuit, include the `mtu` statement:

```
mtu mtu-number;
```

You can include this statement at the following hierarchy levels:

```
[edit logical-routers logical-router-name protocols l2circuit neighbor address
interface interface-name]
```

```
[edit protocols l2circuit neighbor address interface interface-name]
```

Configuring Local Interface Switching

You can configure a virtual circuit entirely on the local router, terminating the circuit on a local interface. Possible uses for this feature include being able to enable switching between Frame Relay DLCIs.

To configure a virtual circuit to terminate locally, include the local-switching statement:

```
local-switching {
  interface interface-name {
    description text;
    end-interface {
      interface interface-name;
      protect-interface interface-name;
    }
    protect-interface interface-name;
  }
}
```

You can include the local-switching statement at the following hierarchy levels:

[edit logical-routers *logical-router-name* protocols l2circuit],

[edit protocols l2circuit]

Local interface switching requires you to configure at least two interfaces:

Starting interface—Configure using the interface statement at the [edit protocols l2circuit local-switching] hierarchy level.

Ending interface—Configure using the end-interface statement at the [edit protocols l2circuit local-switching interface *interface-name*] hierarchy level.

You can also configure virtual circuit interface protection for each local interface:

Protect interface for the starting interface—Configure using the protect-interface statement at the [edit protocols l2circuit local-switching interface *interface-name*] hierarchy level.

Protect interface for the ending interface—Configure using the protect-interface statement at the [edit protocols l2circuit local-switching interface *interface-name* end-interface] hierarchy level.

For more information on how to configure protect interfaces, see “Configuring the Protect Interface” on page 461.

Configuring the Virtual Circuit ID

You configure a virtual circuit ID on each interface. Each virtual circuit ID uniquely identifies the Layer 2 circuit among all the Layer 2 circuits to a specific neighbor. The key to identifying a particular Layer 2 circuit on a PE router is the neighbor address and the virtual circuit ID. An LDP-FEC-to-label binding is associated with a Layer 2 circuit based on the virtual circuit ID in the FEC and the neighbor that sent this binding. The LDP-FEC-to-label binding enables the dissemination of the VPN label used for sending traffic on that Layer 2 circuit to the remote CE router.

To configure the virtual circuit ID, include the `virtual-circuit-id` statement:

```
virtual-circuit-id identifier;
```

You can include the `virtual-circuit-id` statement at the following hierarchy levels:

```
[edit protocols l2circuit neighbor address interface interface-name]
```

```
[edit logical-routers logical-router-name protocols l2circuit neighbor address
interface interface-name]
```

Configuring LDP for Layer 2 Circuits

Use LDP as the signaling protocol to advertise ingress labels to the remote PE routers. When configured, LDP examines the Layer 2 circuit configuration and initiates extended neighbor discovery for all the Layer 2 circuit neighbors (for example, remote PEs). This process is similar to how LDP works when tunneled over Resource Reservation Protocol (RSVP). You must run LDP on the `lo0.0` interface for extended neighbor discovery to function correctly.

For detailed information about how to configure LDP, see the *JUNOS MPLS Applications Configuration Guide*.

Configuring Layer 2 Circuit Policies

You can configure JUNOS routing policies to control the flow of packets over Layer 2 circuits. This capability allows you to provide different level of service over a set of equal-cost Layer 2 circuits. For example, you can configure a circuit for high-priority traffic, a circuit for average-priority traffic, and a circuit for low-priority traffic. By configuring Layer 2 circuit policies, you can ensure that higher-value traffic has a greater likelihood of reaching its destination.

The following sections explain how to configure Layer 2 circuit policies:

Configuring the Layer 2 Circuit Community on page 466

Configuring the Policy Statement for the Layer 2 Circuit Community on page 467

Verifying the Layer 2 Circuit Policy Configuration on page 469

Configuring the Layer 2 Circuit Community

To configure a community for Layer 2 circuits, include the community statement.

```
community name {
  members [ community-ids ];
}
```

You can include the community statement at the following hierarchy levels:

```
[edit policy-options]
```

```
[edit logical-routers logical-router-name policy-options]
```

name identifies the community or communities.

community-ids identifies the type of community or extended community:

A normal community uses the following community ID format:

```
as-number:community-value
```

as-number is the autonomous system (AS) number of the community member.

community-value is the identifier of the community member. It can be a number from 0 through 65,535.

An extended community uses the following community ID format:

```
type:administrator:assigned-number
```

type is the type of target community. The target community identifies the route's destination.

administrator is either an AS number or an Internet Protocol version 4 (IPv4) address prefix, depending on the type of community.

assigned-number identifies the local provider.

You need to associate the communities with the appropriate Layer 2 circuits. To associate a community with a Layer 2 circuit, include the virtual-circuit-id and community statements:

```
virtual-circuit-id number;
community community-name;
```

You can include these statements at the following hierarchy levels:

```
[edit protocols l2circuit neighbor address interface interface-name]
```

```
[edit logical-routers logical-router-name protocols l2circuit neighbor address
interface interface-name]
```

Configuring the Policy Statement for the Layer 2 Circuit Community

To configure a policy to send community traffic over a specific LSP, include the policy-statement statement:

```

policy-statement policy-name {
  term term-name {
    from community community-name;
    then {
      install-nexthop (except | lsp lsp-name | lsp-regex lsp-regular-expression);
      accept;
    }
  }
}

```

You can include the policy-statement statement at the following hierarchy levels:

[edit policy-options]

[edit logical-routers *logical-router-name* policy-options]

To prevent the installation of any matching next hops, include the install-nexthop statement with the except option:

```
install-nexthop except;
```

You can include the install-nexthop statement at the following hierarchy levels:

[edit policy-options policy-statement *policy-name* term *term-name* then]

[edit logical-routers *logical-router-name* policy-options policy-statement *policy-name* term *term-name* then]

To assign traffic from a community to a specific LSP, include the install-nexthop statement with the lsp *lsp-name* option and the accept statement:

```
install-nexthop lsp lsp-name;
accept;
```

You can include these statements at the following hierarchy levels:

[edit policy-options policy-statement *policy-name* term *term-name* then]

[edit logical-routers *logical-router-name* policy-options policy-statement *policy-name* term *term-name* then]

You can also use a regular expression to select an LSP from a set of similarly named LSPs for the `install-nexthop` statement. To configure a regular expression, include the `install-nexthop` statement with the `lsp-regex` option and the `accept` statement:

```
install-nexthop lsp-regex lsp-regular-expression;  
accept;
```

You can include these statements at the following hierarchy levels:

```
[edit policy-options policy-statement policy-name term term-name then]
```

```
[edit logical-routers logical-router-name policy-options policy-statement policy-name  
term term-name then]
```

Example: Configuring a Policy for a Layer 2 Circuit Community

The following example illustrates how you might configure a regular expression in a Layer 2 circuit policy. You create three LSPs to handle gold-tier traffic from a Layer 2 circuit. The LSPs are named `alpha-gold`, `beta-gold`, and `delta-gold`. You then include the `install-nexthop` statement with the `lsp-regex` option with the LSP regular expression `.*-gold` at the `[edit policy-options policy-statement policy-name term term-name then]` hierarchy level:

```
[edit policy-options]  
policy-statement gold-traffic {  
  term to-gold-LSPs {  
    from community gold;  
    then {  
      install-nexthop lsp-regex .*-gold;  
      accept;  
    }  
  }  
}
```

The community gold Layer 2 circuits can now use any of the `-gold` LSPs. Given equal utilization across the three `-gold` LSPs, LSP selection is made at random.

You need to apply the policy to the forwarding table. To apply a policy to the forwarding table, configure the `export` statement at the `[edit routing-options forwarding-table]` hierarchy level:

```
[edit routing-options forwarding-table]  
export policy-name;
```

Verifying the Layer 2 Circuit Policy Configuration

To verify that you have configured a policy for the Layer 2 circuit, issue the `show route table mpls detail` command. It should display the community for ingress routes that corresponds to the Layer 2 circuits, as shown by the following example:

```
user@host> show route table mpls detail
so-1/0/1.0 (1 entry, 1 announced)
  *L2VPN Preference: 7
    Next hop: via so-1/0/0.0 weight 1, selected
    Label-switched-path to-community-gold
    Label operation: Push 100000 Offset: -4
    Next hop: via so-1/0/0.0 weight 1
    Label-switched-path to-community-silver
    Label operation: Push 100000 Offset: -4
    Protocol next hop: 10.255.245.45
    Push 100000 Offset: -4
    Indirect next hop: 85333f0 314
    State: <Active Int>
    Local AS: 100
    Age: 22
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: 100:1
```

For more information on how to configure routing policies, see the *JUNOS Policy Framework Configuration Guide*.

Configuring the Control Word for Layer 2 Circuits

To emulate the virtual circuit (VC) encapsulation for Layer 2 circuits, a 4-byte control word is added between the Layer 2 protocol data unit (PDU) being transported and the VC label that is used for demultiplexing.

The following sections discuss how to configure the control word for Layer 2 circuits:

Configuring the Control Word for Frame Relay Interfaces on page 469

Disabling the Control Word for Layer 2 Circuits on page 470

Configuring the Control Word for Frame Relay Interfaces

On interfaces with Frame Relay CCC encapsulation, you can configure Frame Relay control bit translation to support Frame Relay services over IP and Multiprotocol Label Switching (MPLS) backbones by using CCC, Layer 2 VPNs, and Layer 2 circuits. When you configure translation of Frame Relay control bits, the bits are mapped into the Layer 2 circuit control word and preserved across the IP or MPLS backbone.

For information on how to configure the control bits, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

Disabling the Control Word for Layer 2 Circuits

To emulate the VC encapsulation for Layer 2 circuits, a 4-byte control word is added between the Layer 2 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.

The JUNOS software supports the control word for Frame Relay. However, it does not support the control word for any other 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, the JUNOS software supports a null control word (a control word of all zeros). If the JUNOS software receives a packet with a control word attached, the control word is discarded before the packet is forwarded to its destination.

The JUNOS software can typically determine whether a neighboring router supports the control word. However, if you want to explicitly disable its use on a specific interface, include the `no-control-word` statement:

```
no-control-word;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring ATM Trunking on Layer 2 Circuits

You can configure Layer 2 circuits to transport ATM traffic from directly connected ATM switches across an MPLS core network. Traffic from an ATM switch is received on the local PE router. The ATM cells are given an MPLS label and then sent across the MPLS network to the remote PE router. The receiving router removes the MPLS label from the ATM cell and then forwards the cell the receiving ATM switch.



NOTE: ATM trunking on Layer 2 circuits is supported only on T-series and M320 routers and ATM2 IQ PICs.

Figure 49: ATM Trunking on Layer 2 Circuits




Figure 49 illustrates how ATM switches could be linked together by a Layer 2 circuit. The PE1 Router is configured to receive ATM trunk traffic from ATM Switch 1. As each ATM cell is received on the PE1 Router, it is classified by means of the class-of-service (CoS) information in the cell header and then encapsulated as a labeled packet. The CoS information and cell loss priority (CLP) of the ATM cell are copied into the experimental (EXP) bits of the MPLS label. The labeled packet is then transported across the service provider network to the PE2 Router by means of a Layer 2 circuit.

On the PE2 Router, the label is removed and the plain ATM cell is forwarded to ATM Switch 2. The CoS and CLP are extracted from the EXP bits and are then used to select the correct output queue and determine whether the ATM cell should be dropped.

The ATM physical port on the router can support 32 logical trunks when network-to-network interface (NNI) is used and 8 logical trunks when user-to-network interface (UNI) is used. A trunk can carry traffic on 32 virtual path identifiers (VPIs), numbered 0 through 31. Each ATM trunk is associated with an MPLS label and a logical interface. On the ingress router, one or more of these trunks are mapped to a Layer 2 circuit.

The configuration for the Layer 2 circuit between PE routers is conventional. Follow the procedures outlined in this chapter for configuring the circuit. However, there is some specific configuration you need to complete for the Layer 2 circuit to carry traffic from an ATM trunk.

First, enable ATM trunking for Layer 2 circuits. To enable ATM trunking for Layer 2 circuits, specify the trunk option for the atm-l2circuit-mode statement:

```
atm-l2circuit-mode trunk (uni | nni);
```

You can include the atm-l2circuit-mode statement at the [chassis fpc *number* pic *number*] hierarchy level.

Specify the uni option for UNI trunks and the nni option for NNI trunks. The default option is uni.

You also need to configure each ATM trunk for a specific logical interface. Each ATM trunk has a trunk identifier in the range from 0 to 31. This configuration step is in addition to the typical configuration steps you follow related to configuring interfaces for Layer 2 circuits, as described in “Configuring Interfaces for Layer 2 Circuits” on page 460.

To associate a specific trunk identifier with a logical interface, include the trunk-id statement:

```
trunk-id number;
```

You can include the trunk-id statement at the following hierarchy levels:

```
[edit interfaces interface-name unit number]
```

```
[edit logical-router logical-router-name interfaces interface-name unit number]
```

Since ATM trunking is supported on ATM2 IQ PICs only, the only value you can configure for the pic-type statement is atm2. If you do not configure the pic-type statement but you do configure the trunk option for the atm-l2circuit-mode statement (at the [chassis fpc *number* pic *number*] hierarchy level), the pic-type statement defaults to atm2.

Configuring Bandwidth Allocation and Call Admission Control

You can configure bandwidth allocation and call admission control (CAC) on Layer 2 circuits. This feature is available for RSVP-signaled LSPs traversing an MPLS network.

When you enable bandwidth allocation on a Layer 2 circuit, attempts to establish an RSVP-signaled LSP are preceded by a check of the available bandwidth on the network. This check is the CAC. The available bandwidth is compared to the bandwidth requested by the LSP. If there is insufficient bandwidth, the Layer 2 circuit is not established and an error message is generated. To apply CAC to a Layer 2 circuit, a bandwidth constraint must be configured.

You can specify the bandwidth for a Layer 2 circuit without configuring a bandwidth for each class type (queue). To specify the bandwidth allocation for a Layer 2 circuit, include the bandwidth statement:

```
bandwidth bandwidth;
```

Specify the bandwidth in bits per second.

You can include the bandwidth statement at the following hierarchy levels:

```
[edit protocols l2circuit neighbor address interface interface-name]
```

```
[edit logical-routers logical-router-name protocols l2circuit neighbor address  
interface interface-name]
```

Alternatively, you can configure the bandwidth for each class type on a Layer 2 circuit. If you use this type of configuration, you cannot simultaneously configure the nonclass type of bandwidth configuration for the Layer 2 circuit (the commit fails).

To configure the bandwidth for each class type on an Layer 2 circuit, include the bandwidth statement:

```
bandwidth {
  ct0 bandwidth;
  ct1 bandwidth;
  ct2 bandwidth;
  ct3 bandwidth;
}
```

You can include the bandwidth statement at the following hierarchy levels:

```
[edit protocols l2circuit neighbor address interface interface-name]
```

```
[edit logical-routers logical-router-name protocols l2circuit neighbor address
interface interface-name]
```



NOTE: It is mandatory to configure a bandwidth model when the bandwidth allocation is per class. For information on how to configure a bandwidth model, see the *JUNOS MPLS Applications Configuration Guide*.

Specify the bandwidth for each class type in bits per second. It is not necessary to specify a bandwidth for all four class types.

Tracing Layer 2 Circuit Creation and Changes

To trace the creation of and changes to Layer 2 circuits, you can specify options in the traceoptions statement:

```
traceoptions {  
    file filename <replace> <size size> <files number> <nostamp>  
    <no-world-readable> <world-readable>;  
    flag flag <flag-modifier> <disable>;  
}
```

You can include the traceoptions statement at the following hierarchy levels:

[edit protocols l2circuit]

[edit logical-routers *logical-router-name* protocols l2circuit]

The following tracing flags display the operations associated with Layer 2 circuits:

connections—Layer 2 circuit connections (events and state changes)

error—Error conditions

FEC—Layer 2 circuit advertisements received or sent using LDP

topology—Layer 2 circuit topology changes caused by reconfiguration or advertisements received from other PE routers