

Chapter 9

Layer 3 VPN Configuration Guidelines

To configure Layer 3 virtual private network (VPN) functionality, you must enable VPN support on the provider edge (PE) router. You must also configure any provider (P) routers that service the VPN, and you must configure the customer edge (CE) routers so that their routes are distributed into the VPN.

To configure Layer 3 VPNs, you include statements at the [edit routing-instances] hierarchy level:

```
[edit]
routing-instances {
  routing-instance-name {
    description text;
    interface interface-name;
    instance-type vrf;
    route-distinguisher ( as-number:number | ip-address:number );
    vrf-import [ policy-names ];
    vrf-export [ policy-names ];
    vrf-target ( community-name | export community-name | import community-name );
    vrf-table-label;
    protocols {
      bgp {
        bgp-configuration;
      }
      ospf {
        ospf-configuration;
      }
      pim {
        pim-configuration;
        vpn-group-address address;
      }
      rip {
        rip-configuration;
      }
    }
  }
}
```

```

routing-options {
  autonomous-system autonomous-system <loops number>;
  forwarding-table {
    export [ policy-names ];
  }
  interface-routes {
    rib-group group-name;
  }
  martians {
    destination-prefix match-type <allow>;
  }
  maximum-routes route-limit <log-only | threshold value>;
  options {
    syslog (level level | upto level);
  }
  rib routing-table {
    static {
      defaults {
        static-options;
      }
      route destination-prefix {
        next-hop;
        static-options;
      }
    }
  }
  martians {
    destination-prefix match-type <allow>;
  }
  static {
    defaults {
      static-options;
    }
    route destination-prefix {
      policy [ policy-names ];
      static-options;
    }
  }
}
router-id address;
static {
  defaults {
    static-options;
  }
  route destination-prefix {
    policy [ policy-names ];
    static-options;
  }
}
}
}
}

```

For Layer 3 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 enable a signaling protocol, 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 3 VPNs are disabled.

Many of the configuration procedures for Layer 3 VPNs are common to all types of VPNs. These procedures are described in detail in Chapter 2, “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 3 VPNs:

- Configure VPN Routing between the PE and CE Routers on page 103
- Filter Traffic Based on the IP Header on page 115
- Configure a VPN Tunnel for VRF Table Lookup on page 116
- Configure a Logical Unit on the Loopback Interface on page 116
- Configure Multicast over Layer 3 VPNs on page 117
- Configure Packet Forwarding for Layer 3 VPNs on page 118
- Configure a GRE Tunnel Interface for Layer 3 VPNs on page 119
- Configure an ES Tunnel Interface for Layer 3 VPNs on page 121
- Configure IPsec between PE Routers Instead of MPLS on page 123

For configuration examples, see “Layer 3 VPN Configuration Examples” on page 141 and “Layer 3 VPN Internet Access Examples” on page 233.

Configure VPN Routing between the PE and CE Routers

For the PE router to distribute VPN-related routes to and from connected CE routers, you must configure routing within the VPN routing instance. You can configure a routing protocol—BGP, OSPF, or RIP—or you can configure static routing. For the connection to each CE router, you can configure only one type of routing.

This section describes how to do the following tasks:

- Configure BGP between the PE and CE Routers on page 104
- Configure OSPF between the PE and CE Routers on page 104

Configure RIP between the PE and CE Routers on page 109

Configure Static Routes between the PE and CE Routers on page 110

Limit the Routes Accepted from a CE Router on page 110

Configure IPv6 between the PE and CE Routers on page 111

Configure EBGP or IBGP Multihop between PE and CE Routers on page 114

Configure BGP between the PE and CE Routers

To configure BGP as the routing protocol between the PE and the CE routers, include the `bgp` statement at the [edit routing-instances *routing-instance-name* protocols] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
bgp {
  group group-name {
    peer-as as-number;
    neighbor ip-address;
  }
}
```



Note

Route reflectors and cluster IDs are not supported on a routing instance. Do not configure the `cluster-id` statement at the [edit routing-instances *routing-instance-name* protocols `bgp` group *group-name*] hierarchy level. Doing so causes the configuration to fail.

Configure OSPF between the PE and CE Routers

You can configure OSPF (version 2 or version 3) to distribute VPN-related routes between PE and CE routers.

To configure OSPF version 2 as the routing protocol between a PE and CE router, include the `protocols ospf` statement at the [edit routing-instances *routing-instance-name* protocols] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
  area area {
    interface interface-name;
  }
}
```

To configure OSPF version 3 as the routing protocol between a PE and CE router, include the protocols ospf3 statement at the [edit routing-instances *routing-instance-name* protocols] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf3 {
  area area {
    interface interface-name;
  }
}
```

Configure an OSPF Domain ID

For most OSPF configurations involving Layer 3 VPNs, you do not need to configure an OSPF domain ID. However, for a Layer 3 VPN connecting multiple OSPF domains, configuring OSPF domain IDs can help you control link-state advertisement (LSA) translation (for Type 3 and Type 5 LSAs) between the OSPF domains and back-door paths. Each VRF table in a PE router associated with an OSPF instance is configured with the same OSPF domain ID. The default OSPF domain ID is 0.0.0.0. This is a null value for the domain ID. As shown in Table 2, a route with a null domain ID is handled differently from a route without any domain ID at all.

Table 2: How a PE Router Redistributes and Advertises Routes

| Route Received | Domain ID of the Route Received | Domain ID on the Receiving Router | Route Redistributed and Advertised As |
|----------------|---------------------------------|-----------------------------------|---------------------------------------|
| Type 3 route | A.B.C.D | A.B.C.D | Type 3 LSA |
| Type 3 route | A.B.C.D | E.F.G.H | Type 5 LSA |
| Type 3 route | 0.0.0.0 | 0.0.0.0 | Type 3 LSA |
| Type 3 route | Null | 0.0.0.0 | Type 3 LSA |
| Type 3 route | Null | Null | Type 3 LSA |
| Type 3 route | 0.0.0.0 | Null | Type 3 LSA |
| Type 3 route | A.B.C.D | Null | Type 5 LSA |
| Type 3 route | Null | A.B.C.D | Type 5 LSA |
| Type 5 route | - | - | Type 5 LSA |

You can configure an OSPF domain ID for both version 2 and version 3 of OSPF. The only difference in the configuration is that you include statements at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level for OSPF version 2 and at the [edit routing-instances *routing-instance-name* protocols ospf3] hierarchy level for OSPF version 3. The configuration descriptions that follow present the OSPF version 2 statement only. However, the substatements are also valid for OSPF version 3.

To configure an OSPF domain ID, include the domain-id statement at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
  domain-id domain ID;
}
```

You can set a VPN tag for the OSPF external routes generated by the PE router to prevent looping. By default, this tag is automatically calculated and needs no configuration. To configure the domain VPN tag for Type 5 LSAs, include the `domain-vpn-tag number` statement at the `[edit routing-instances routing-instance-name protocols ospf]` hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
  domain-vpn-tag number;
}
```

The range is 1 through 4,294,967,295. If you set VPN tags manually, you must set the same value for all PE routers in the VPN.

Hub-and-Spoke Layer 3 VPNs and OSPF Domain ID

The default behavior of an OSPF domain ID can cause the following problems for hub-and-spoke Layer 3 VPNs using OSPF between the PE and CE routers:

PE routers set the down (DN) bit on all OSPF summary LSAs originating from area 0. PE routers are designated as area 0 by default because of the OSPF domain ID. When a PE router receives a summary LSA with the DN bit set, the LSA is not used in the OSPF calculation. This is done to prevent routing loops.

For a hub-and-spoke Layer 3 VPN, when the hub PE router generates an OSPF summary LSA, it also sets the DN bit before sending it to the hub CE router. When the hub CE router sends the LSA back to the PE router, the PE router does not use the LSA in the OSPF calculation because the DN bit is set. Routes aggregated within the CE router are not affected.

PE routers generating external LSAs learned from BGP updates set the `vpn-route-tag` field to a value derived from the PE router's AS number and an arbitrary tag. When a PE router receives an external LSA with a `vpn-route-tag` field that matches its own `vpn-route-tag` field, the LSA is not used in the OSPF calculation. This is done to prevent routing loops.

For a hub-and-spoke Layer 3 VPN, an external LSA originated by a hub PE router is sent to the hub CE router, which then sends it back to the same PE router. Because the `vpn-route-tag` field matches the PE router's `vpn-route-tag` field, the LSA is not used in the OSPF calculation. Routes aggregated within the CE router are not affected.

For hub-and-spoke Layer 3 VPNs using OSPF between the PE and CE routers to work, you need to configure the following on the hub PE router:

Configure the `disable` statement at the `[edit routing-instances routing-instance-name protocols ospf domain-id]` hierarchy level on the routing instance for the hub CE router. This removes area 0 from the PE router, allowing the PE router to forward LSAs without setting the DN bit. When an LSA comes back from the hub CE router, the PE router can install it because the DN bit is not set.

Configure 0 for the `vpn-route-tag` statement at the `[edit routing-instances routing-instance-name protocols ospf]` hierarchy level on the routing instance for the spoke CE router. This removes any VPN route tags that are set on the external LSAs, preventing a VPN route tag match and allowing the PE router to install the LSA.

Compatibility with JUNOS Releases before 5.3

For JUNOS release 5.3, the format for domain-id, an extended community type defined in the BGP extended community attribute field, was modified to comply with the Internet Engineering Task Force (IETF) draft draft-rosen-vpns-ospf-bgp-mpls (available at <http://www.ietf.org/>). JUNOS releases prior to 5.3 continue to use the previously supported vendor-specific formats.

The OSPF domain ID format is incompatible between JUNOS 5.3 or later and JUNOS 5.2 or earlier. For OSPF domain IDs to function properly between a PE router running JUNOS 5.3 or later and a PE router running JUNOS 5.2 or earlier, you need to define the extended community type for the BGP extended community attribute field as domain-id-vendor (instead of as domain-id). This is part of the policy-options configuration for the OSPF domain ID configured at the [edit policy-options community vrf_export_attributes members] hierarchy level:

```
[edit policy-options community vrf_export_attributes members]
domain-id-vendor:ip-address
```

You also need to configure the route-type-community statement with the vendor option at the [edit routing-instances *routing-instance-name* protocols ospf] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
ospf {
  route-type-community vendor;
}
```

The default value for the route-type-community statement is iana. For OSPF version 3, you configure the route-type-community statement with the vendor option at the [edit routing-instances *routing-instance-name* protocols ospf3] hierarchy level

Example Configurations for Compatibility with JUNOS Releases before 5.3

The following example shows a configuration of the policy options for a PE router. The PE router has an OSPF domain ID configured.

It needs to be compatible with a router running a pre-5.3 version of JUNOS software. As a part of the community statement configuration, specify domain-id-vendor for the attribute that assigns the domain ID instead of domain-id:

```
[edit]
policy-options {
  policy-statement vrf_import_routes {
    term a {
      from {
        protocol bgp;
        community vrf_import_attributes;
      }
      then accept;
    }
    term b {
      then reject;
    }
  }
}
```

```

policy-statement vrf_export_routes {
  term a {
    from protocol ospf;
    then {
      community add vrf_export_attributes;
      accept;
    }
  }
  term b {
    then reject;
  }
}
community vrf_export_attributes members [ target:10.19.2.0:5
domain-id-vendor:1.2.3.4:0 ];
community vrf_import_attributes members target:10.19.1.0:5;
}

```

The following example shows a configuration for a routing instance on a PE router. The PE router has an OSPF domain ID configured. It needs to be compatible with a router running an earlier version of JUNOS software. The configuration includes the route-type-community statement with the vendor option. This is so the PE router receiving the route knows how to parse the incoming BGP attribute field containing the domain ID.

The example configuration follows:

```

[edit]
routing-instances {
  CE_A {
    instance-type vrf;
    interface fe-1/0/0.0;
    route-distinguisher 10.255.25.270:1;
    vrf-import vrf_import_routes;
    vrf-export vrf_export_routes;
    protocols {
      ospf {
        route-type-community vendor;
        domain-id 1.2.3.4;
        export vrf_import_routes;
        area 0.0.0.0 {
          interface fe-1/0/0.0;
        }
      }
    }
  }
}
}

```

Configure RIP between the PE and CE Routers

For a Layer 3 VPN, you can configure RIP on the PE router to learn the routes of the CE router or to propagate the routes of the PE router to the CE router. RIP routes learned from neighbors configured at any [edit routing-instances] hierarchy level are added to the routing instance's inet table (*instance_name.inet.0*).

To configure RIP as the routing protocol between the PE and the CE router, include the rip statement at the [edit routing-instances *routing-instance-name* protocols] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
rip {
  group group-name {
    neighbor interface-name;
  }
}
```

To install routes learned from a RIP routing instance to multiple routing tables, configure the rib-group statement at the [edit protocols rip] hierarchy level or at the [edit routing-instances *routing-instance-name* protocols rip] hierarchy level:

```
[edit protocols rip]
rib-group inet group-name;
group group-name {
  neighbor interface-name;
}
```

To configure a routing table group, configure the rib-group statement at the [edit routing-options] hierarchy level.

To add a routing table to a routing table group, you need to configure the import-rib statement at the [edit routing-options rib-groups *group-name*] hierarchy level. The first routing table name specified under the import-rib statement must be the name of the routing table you are configuring. See the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols* for more information about how to configure routing tables and routing table groups.

Configure the import-rib statement at the [edit routing-options rib-groups *group-name*] hierarchy level as follows:

```
[edit routing-options rib-groups group-name]
import-rib [group-name]
```

Configure Static Routes between the PE and CE Routers

To configure a static route between the PE and the CE routers, include the routing-options static statement at the [edit routing-instances *routing-instance-name* routing-options] hierarchy level:

```
[edit routing-instances routing-instance-name routing-options]
static {
  route destination-prefix {
    next-hop;
    static-options;
  }
}
```

For more information about configuring routing protocols and static routes, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Limit the Routes Accepted from a CE Router

A route limit sets an upper limit for the number of prefixes installed into routing tables. You can use route limits to curtail the number of routes received from a CE router in a VPN. A route limit applies only to dynamic routing protocols, and is not applicable to static or interface routes.

To limit the number of routes accepted by a PE router from a CE router, include the maximum-routes statement at the [edit routing-instances *routing-instance-name* routing-options] hierarchy level:

```
[edit routing-instances routing-instance-name routing-options]
maximum-routes route-limit <log-only | threshold value>;
```

There are two modes for route limits: advisory (set with the log-only option) and mandatory. An advisory limit triggers only warnings. The log messages are rate-limited to once every 30 seconds. A mandatory limit, in addition to triggering a warning message, rejects any additional routes after the threshold is reached. The threshold value is a percentage of the route limit at which warning messages are logged.



Note

Setting a route limit might result in unpredictable dynamic routing protocol behavior.

Configure IPv6 between the PE and CE Routers

You can configure IPv6 between the PE and CE routers of a Layer 3 VPN. The PE router must have the PE router to PE router BGP session configured with the family inet6-vpn statement. The CE router must be capable of receiving IPv6 traffic. You can configure BGP or static routes between the PE and CE routers.



Note

The vrf-table-label statement cannot be configured in an IPv6 Layer 3 VPN environment. If you configure a dual-stack VRF routing table (where both IPv4 and IPv6 routes are supported) and also configure the vrf-table-label statement for that VRF, the IPv4 traffic flows normally but the IPv6 traffic is dropped.

To configure IPv6 VPNS between the PE routers, complete the following steps:

Configure IPv6 on the PE Router on page 111

Configure BGP or Static Routes on the PE Router on page 112

Configure IPv6 on the Interfaces on page 114

Configure IPv6 on the PE Router

To configure IPv6 between the PE and CE routers, include the following statements at the [edit protocols bgp group *group-name*] hierarchy level on the PE router:

```
[edit protocols bgp group group-name]
family inet6-vpn {
  (unicast | multicast | any) {
    prefix-limit maximum prefix-limit;
    rib-group rib-group-name;
  }
}
```

Configure the ipv6-tunneling statement at the [edit protocols mpls] hierarchy level:

```
[edit protocols mpls]
ipv6-tunneling;
```

Configure BGP or Static Routes on the PE Router

You must configure either BGP or static routes for the connection between the PE and CE routers in the Layer 3 VPN. You can configure BGP to handle just IPv4 routes or both IPv4 and IPv6 routes.

For more information about IPv6, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

The following sections describe how to configure BGP and static routes:

Configure BGP on the PE Router to Handle IPv6 Routes on page 112

Configure BGP on the PE Router to Handle IPv4 and IPv6 Routes on page 113

Configure Static Routes on the PE Router on page 113

Configure BGP on the PE Router to Handle IPv6 Routes

Configure BGP in the Layer 3 VPN routing instance to handle IPv6 routes at the [edit routing-instances *routing-instance-name* protocols bgp] hierarchy level:

```
[edit]
routing-instances routing-instance-name {
  protocols {
    bgp {
      group group-name {
        local-address IPv6-address;
        family inet6 {
          unicast;
        }
        peer-as as-number;
        neighbor IPv6-address;
      }
    }
  }
}
```

Configure BGP on the PE R outer to Handle IPv4 and IPv6 R outes

Configure BGP in the Layer 3 VPN routing instance to handle both IPv4 and IPv6 routes at the [edit routing-instances *routing-instance-name* protocols bgp] hierarchy level:

```
[edit]
routing-instances routing-instance-name {
  protocols {
    bgp {
      group group-name {
        local-address IPv4-address;
        family inet {
          unicast;
        }
        family inet6 {
          unicast;
        }
        peer-as as-number;
        neighbor address;
      }
    }
  }
}
```

Configure Static Routes on the PE R outer

Configure a static route to the CE router in the Layer 3 VPN routing instance at the [edit routing-instances *routing-instance-name* routing-options rib *routing-table-group-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]
routing-options {
  rib routing-table-group-name.inet6.0 {
    static {
      defaults {
        static-options;
      }
    }
  }
}
```

Configure IPv6 on the Interfaces

You need to configure IPv6 on the PE router interfaces to the CE routers and on the CE router interfaces to the PE routers.

To configure the interface to handle IPv6 routes, include the family inet6 statement under the [edit interfaces *interface-name* unit *unit-number*] hierarchy level:

```
[edit]
interfaces {
  interface-name {
    unit unit-number {
      family inet6 {
        address IPv6-address;
      }
    }
  }
}
```

If you have configured the Layer 3 VPN to handle both IPv4 and IPv6 routes, you need to configure the interface to handle both IPv4 and IPv6 routes:

```
[edit]
interfaces {
  interface-name {
    unit unit-number {
      family inet {
        address IPv4-address;
      }
      family inet6 {
        address IPv6-address;
      }
    }
  }
}
```

Configure EBGP or IBGP Multihop between PE and CE Routers

You can configure an external BGP (EBGP) or internal BGP (IBGP) multihop session between the PE and CE routers of a Layer 3 VPN. This allows you to have one or more routers between the PE and CE routers. Using IBGP between PE and CE routers does not require the configuration of any additional statements. However, using EBGP between the PE and CE routers requires the configuration of the multihop statement.

To configure an external BGP multihop session for the connection between the PE and CE routers, include the multihop statement at the [edit routing-instances *routing-instance-name* protocols bgp], [edit routing-instances *routing-instance-name* protocols bgp group *group-name*], or [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*] hierarchy level:

```
multihop <ttl-value>;
```

Filter Traffic Based on the IP Header

The `vrf-table-label` statement makes it possible to map the inner label to a specific VRF and thus allow the examination of the encapsulated IP header at an egress VPN router. You might want to enable this functionality so you can do either of the following:

Forward traffic on a PE-router-to-CE-device interface, in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch).

The first lookup is done on the VPN label to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts on the shared medium.

Perform egress filtering at the egress PE router.

The first lookup on the VPN label is done to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to filter and forward packets. You can enable this functionality by configuring output filters on the VRF interfaces.

When you use the `vrf-table-label` statement to configure a VRF table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF.

Any routes configured in a VRF with the `vrf-table-label` statement are advertised with the LSI logical interface label allocated for the VRF. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table.

To filter traffic based on the IP header, include the `vrf-table-label` statement at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit routing-instances routing-instance-name]  
vrf-table-label;
```

You can configure the `vrf-table-label` statement for both IPv4 and IPv6 Layer 3 VPNs. If you configure the `vrf-table-label` statement for a dual-stack VRF routing table (where both IPv4 and IPv6 routes are supported), the `vrf-table-label` statement applies to both the IPv4 and IPv6 routes and the same label is advertised for both sets of routes.

Egress Filtering Options

You can enable egress filtering (which allows egress Layer 3 VPN PE routers to perform lookups on the VPN label and IP header at the same time) by including the `vrf-table-label` statement at the [edit routing-instances *instance-name*] hierarchy level. However, this feature works only for non-channelized Point-to-Point Protocol/High-level Data Link Control (PPP/HDLC) SONET core-facing interfaces and non-channelized Gigabit and Fast Ethernet core-facing interfaces. The `vrf-table-label` statement cannot be configured for the 10-port E1 Physical Interface Card (PIC) or for aggregated interfaces. There is no restriction on CE-router-to-PE-router interfaces.

You can also enable egress filtering by configuring a VPN tunnel (VT) interface on routers equipped with a Tunnel Services PIC. When you enable egress filtering this way, there is no restriction on the type of core-facing interface used. There is also no restriction on the type of CE-router-to-PE-router interface used.

Limitations

When you configure the `vrf-table-label` statement, be aware of the following limitations:

The `vrf-table-label` statement is supported on M-series platforms only. It is not supported on T-series platforms.

If you configure a virtual loopback tunnel interface and the `vrf-table-label` statement on the same routing instance, the `vrf-table-label` statement takes precedence over the virtual loopback tunnel interface.

Do not use the `vrf-table-label` statement for source class usage/destination class usage (SCU/DCU) configurations. For information on SCU/DCU configuration, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service*.

The `vrf-table-label` statement cannot be configured for the 10-port E1 Physical Interface Card (PIC) or for aggregated interfaces. There is no restriction on CE-router-to-PE-router interfaces.

You cannot configure the `vrf-table-label` statement when the encapsulation of the PE router to provider router interface is Multilink Point-to-Point Protocol (MLPPP).

Configure a VPN Tunnel for VRF Table Lookup

You can configure a VPN tunnel to facilitate VRF table lookup based on MPLS labels. You might want to enable this functionality to forward traffic on a PE-router-to-CE-device interface in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch), or to perform egress filtering at the egress PE router.

For more information on VPN tunnels and VT interfaces, see the *JUNOS Internet Software Configuration Guide: Services Interfaces*.

Configure a Logical Unit on the Loopback Interface

You can configure a logical unit on the loopback interface into each VRF routing instance you have configured on the router. This is possible only on Layer 3 VPNs (VRF routing instances). Associating a VRF routing instance with a logical unit on the loopback interface allows you to easily identify the VRF. This is useful for troubleshooting, allowing you to ping a remote CE router from a local PE router in a Layer 3 VPN. See “Ping a Remote CE Router from a PE Router” on page 137 for more information.

You can also configure a firewall filter for the logical unit on the loopback interface, allowing you to filter traffic for the VRF routing instance associated with it.

The following describes how firewall filters affect the VRF routing instance depending on whether they are configured on the default loopback interface, the VRF routing instance, or some combination of the two. The “default loopback interface” refers to `lo0.0` (associated with the default routing table) and the “VRF loopback interface” refers to `lo0.n`, which is configured in the VRF routing instance.

If you configure Filter A on the default loopback interface and Filter B on the VRF loopback interface, the VRF routing instance uses Filter B.

If you configure Filter A on the default loopback interface, but do not configure a filter on the VRF loopback interface, the VRF routing instance does not use a filter.

If you configure Filter A on the default loopback interface but do not even configure a VRF loopback interface, the VRF routing instance uses Filter A.

To configure a logical unit on the loopback interface, configure the unit statement at the [edit interfaces lo0] hierarchy level:

```
[edit interfaces]
lo0 {
  unit number {
    family inet {
      address address;
    }
  }
}
```

To associate a firewall filter with the logical unit on the loopback interface, include the following statements at the [edit interfaces lo0 unit *unit-number* family inet] hierarchy level:

```
[edit interfaces lo0 unit unit-number family inet]
filter {
  input filter-name;
}
```

You also need to include the lo0.*n* interface in the configuration for the VRF routing instance at the [edit routing-instances *routing-instance-name*] hierarchy level:

```
[edit]
routing-instances {
  routing-instance-name {
    interface lo0.n;
  }
}
```

For more information on how to configure firewall filters, see the *JUNOS Internet Software Configuration Guide: Policy Framework*.

Configure Multicast over Layer 3 VPNs

You can configure a Layer 3 VPN to support multicast traffic using the Protocol Independent Multicast (PIM) routing protocol. To support multicast, you need to configure PIM on routers within the VPN and within the service provider's network.

Each PE router configured to run multicast over Layer 3 VPNs must have a Tunnel PIC. A Tunnel PIC is also required on the provider routers that act as rendezvous points (RPs). Tunnel PICs are also needed on all the CE routers acting as designated routers (first-hop/last-hop routers) or as RPs, just as they are in non-VPN PIM environments.

Configure the master PIM instance at the [edit protocols pim] hierarchy level on the CE and PE routers. You also need to configure a PIM instance for the Layer 3 VPN at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level on the PE router. This creates a PIM instance for the indicated routing instance.

For information about how to configure PIM, see the *JUNOS Internet Software Configuration Guide: Multicast*.

The `vpn-group-address` statement is unique to a Layer 3 VPN PIM configuration. You use this statement to configure the group address for the VPN in the service provider's network. This address should be unique for each VPN. It ensures that multicast traffic is transmitted only to the specified VPN.

Configure the `vpn-group-address` statement at the [edit routing-instances *routing-instance-name* protocols pim] hierarchy level:

```
[edit routing-instances routing-instance-name protocols]
pim {
  vpn-group-address address;
}
```

The rest of the Layer 3 VPN configuration for multicast is conventional and is described in other sections of this manual. Most of the specific configuration tasks needed to activate multicast in a VPN environment involve PIM. For more information about how to configure PIM and multicast in JUNOS, including an example of how to configure multicast over Layer 3 VPNs, see the *JUNOS Internet Software Configuration Guide: Multicast*.

Configure Packet Forwarding for Layer 3 VPNs

You can configure the router to support packet forwarding for Internet Protocol version 4 (IPv4) traffic in Layer 2 and Layer 3 VPNs. Packet forwarding is handled in one of the following ways, depending on the type of helper service configured:

BOOTP service—Clients send Bootstrap Protocol (BOOTP) requests through the router configured with BOOTP service to a server in the specified routing instance. The server recognizes the client address and sends a response back to the router configured with BOOTP service. This router forwards the reply to the correct client address in the specified routing instance.

Other services—Clients send requests through the router configured with the service to a server in the specified routing instance. The server recognizes the client address and sends a response to the correct client address in the specified routing instance.

To enable packet forwarding for VPNs, configure the `helpers` statement at the [edit forwarding-options] hierarchy level as follows:

```
[edit forwarding-options]
helpers {
  service {
    description description-of-service;
    server {
      address address {
        routing-instance routing-instance-names;
      }
    }
  }
  interface interface-name {
    description description-of-interface;
    no-listen;
    server {
      address address {
        routing-instance routing-instance-names;
      }
    }
  }
}
```



You can enable packet forwarding for multiple VPNs. However, the client and server must be within the same VPN. Any Juniper Networks routers with packet forwarding enabled along the path between the client and server must also reside within the same VPN.

The address and routing instance together constitute a unique server. This has implications for routers configured with BOOTP service, which can accept multiple servers.

For example, a BOOTP service can be configured as follows:

```
[edit forwarding-options helpers bootp]
server address 1.2.3.4 routing-instance [instance-A instance-B];
```

Though the addresses are identical, the routing instances are different. A packet coming in for BOOTP service on instance-A is forwarded to 1.2.3.4 in the instance-A routing instance, while a packet coming in on instance-B is forwarded in the instance-B routing instance. Other services can only accept a single server, so this configuration does not apply in those cases.

For more information about the statements configured at the [edit forwarding-options] hierarchy level, see the *JUNOS Internet Software Configuration Guide: Policy Framework*.

Configure a GRE Tunnel Interface for Layer 3 VPNs

JUNOS software allows you to configure a generic routing encapsulation (GRE) tunnel between the PE and CE routers for a Layer 3 VPN. The GRE tunnel can have one or more hops.

For more information about how to configure tunnel interfaces, see the *JUNOS Internet Software Configuration Guide: Services Interfaces*.

To configure a GRE tunnel between the PE and CE routers for a Layer 3 VPN, complete the procedures in the following sections:

Configure the GRE Tunnel Interface on the PE Router on page 120

Configure the GRE Tunnel Interface on the CE Router on page 121

Configure the GRE Tunnel Interface on the PE Router

Configure the GRE tunnel interface on the PE router as follows:

```
[edit]
interfaces {
  interface-name {
    unit 0 {
      tunnel {
        source address;
        destination address;
      }
      family inet {
        address address;
      }
    }
  }
}
```

By default, the tunnel destination address is assumed to be in the default Internet routing table, inet.0. If the tunnel destination address is not in inet.0, you need to specify which routing table to search for the tunnel destination address by configuring the routing-instance statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

Configure the GRE tunnel interface on the PE router and specify the name of the routing instance:

```
[edit]
interfaces {
  interface-name {
    unit 0 {
      tunnel {
        source address;
        destination address;
        routing-instance {
          destination routing-instance-name;
        }
      }
      family inet {
        address address;
      }
    }
  }
}
```

To complete the GRE tunnel interface configuration, you need to configure the GRE interface at the [edit routing-instances *routing-instance-name*] hierarchy level under the appropriate routing-instance:

```
[edit]
routing-instances {
  routing-instance-name {
    interface interface-name;
  }
}
```

Configure the GRE Tunnel Interface on the CE Router

Configure the GRE tunnel interface on the CE router as follows:

```
[edit]
interfaces {
  interface-name {
    unit 0 {
      tunnel {
        source address;
        destination address;
      }
      family inet {
        address address;
      }
    }
  }
}
```

Configure an ES Tunnel Interface for Layer 3 VPNs

An ES tunnel interface allows you to configure an IP Security (IPSec) tunnel between the PE and CE routers of a Layer 3 VPN. The IPSec tunnel can include one or more hops.

To configure an ES tunnel interface between the PE and CE routers of a Layer 3 VPN, complete the procedures in the following sections:

Configure the ES Tunnel Interface on the PE Router on page 121

Configure the ES Tunnel Interface on the CE Router on page 123

Configure the ES Tunnel Interface on the PE Router

Configure the ES tunnel interface on the PE router as follows:

```
[edit]
interfaces {
  interface-name {
    unit 0 {
      tunnel {
        source address;
        destination address;
      }
      family inet {
        address address;
        ipsec-sa security-association-name;
      }
    }
  }
}
```

By default, the tunnel destination address is assumed to be in the default Internet routing table, inet.0. For IPSec tunnels using manual security association (SA), if the tunnel destination address is not in the default inet.0 routing table, you need to specify which routing table to search for the tunnel destination address by configuring the routing-instance statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

```
[edit]
interfaces {
  interface-name {
    unit 0 {
      tunnel {
        source address;
        destination address;
        routing-instance {
          destination routing-instance-name;
        }
        family inet {
          address address;
          ipsec-sa security-association-name;
        }
        family mpls;
      }
    }
  }
}
```



Note

For IPSec tunnels using dynamic SA, the tunnel destination address must be in the default Internet routing table, inet.0.

You also need to configure the ES interface at the [edit routing-instances *routing-instance-name*] hierarchy level for the appropriate routing instance:

```
[edit]
routing-instances {
  routing-instance-name {
    interface interface-name;
  }
}
```

Configure the ES Tunnel Interface on the CE Router

Configure the ES tunnel interface on the CE router as follows:

```
[edit]
interfaces {
  interface-name {
    unit 0 {
      tunnel {
        source address;
        destination address;
      }
      family inet {
        address address;
        ipsec-sa security-association-name;
      }
    }
  }
}
```

For more information about how to configure tunnel interfaces, see the *JUNOS Internet Software Configuration Guide: Services Interfaces*.

For more information about how to configure IPsec interfaces, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

Configure IPsec between PE Routers Instead of MPLS

A conventional Layer 3 BGP/MPLS VPN requires the configuration of MPLS LSPs between the PE routers. When a PE router receives a packet from a CE router, it performs a lookup in a specific VRF table for the IP destination address and obtains a corresponding MPLS label stack. The label stack is used to forward the packet to the egress PE router, where the bottom label is removed and the packet is forwarded to the specified CE router.

You can provide Layer 3 BGP/MPLS VPN service without an MPLS backbone. Instead of configuring MPLS LSPs between the PE routers, you configure GRE and IPsec tunnels between the PE routers. The MPLS information for the VPN (the VPN label) is encapsulated within an IP header and an IPsec header. The source address of the IP header is the address of the ingress PE router. The destination address has the BGP next hop, the address of the egress PE router.



Note

The IPsec tunnel requires the use of an ES PIC. The GRE tunnel requires the use of a Tunnel Services PIC.

To configure IPsec between PE routers, complete the following:

1. Configure an IPsec tunnel between the PE routers. The source address is that of the ingress PE router, and the destination address is that of the egress PE router:

```
[edit interfaces]
es-interface-name {
  unit unit-number {
    tunnel {
      source source-address;
      destination destination-address;
    }
    family inet {
      ipsec-sa sa-esp-dynamic;
      address address;
    }
    family mpls;
  }
}
```

2. Configure IPsec on the PE router. For information about how to configure IPsec, see the *JUNOS Internet Software Configuration Guide: Getting Started*.
3. Configure a GRE tunnel between the PE routers. Again, the source address is that of the ingress PE router, and the destination address is that of the egress PE router:

```
[edit interfaces]
gr-interface-name {
  unit unit-number {
    family inet {
      address address;
    }
    family mpls;
    tunnel {
      source source-address;
      destination destination-address;
    }
  }
}
```

4. Configure BGP between the PE routers:

```
[edit protocols]
bgp {
  group pe {
    type internal;
    local-address local-address;
    family inet {
      unicast;
    }
    family inet-vpn {
      unicast;
    }
    peer-as as-number;
    neighbor address;
  }
}
```

5. Configure the routing instance:

```
[edit]
routing-instances {
  routing-instance-name {
    instance-type vrf;
    interface interface-name;
    route-distinguisher address;
    vrf-import import-policy-name;
    vrf-export export-policy-name;
    protocols {
      bgp {
        group routing-instance-name {
          type external;
          peer-as as-number;
          as-override;
          neighbor address;
        }
      }
    }
  }
}
```

6. Configure the policy options:

```
[edit]
policy-options {
  policy-statement import-policy-name {
    term 1 {
      from {
        protocol bgp;
        community community-name;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
  policy-statement export-policy-name {
    term 1 {
      from protocol [ bgp direct ];
      then {
        community add community-name;
        accept;
      }
    }
    term 2 {
      then reject;
    }
  }
}
community community-name members target:target;
}
```

