

Chapter 4

Graceful Restart

With standard implementations of routing protocols, any service interruption requires an affected router to recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. As a result, an unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets.

In contrast, graceful restart allows a restarting router and its neighbors to continue forwarding packets without disrupting network performance. Because neighboring routers assist in the restart, the restarting router can quickly resume full operation without having to recalculate algorithms from scratch.

This feature guide covers these topics:

- Overview on page 148

- System Requirements on page 150

- Terms and Acronyms on page 151

- Configure Routing Protocol and Summaries Graceful Restart on page 151

- Configure Graceful Restart for an MPLS-Related Protocol on page 154

- Configure VPN Graceful Restart on page 156

 - Example: Layer 3 VPN Graceful Restart Configuration on page 157

 - Check Your Work on page 173

- For More Information on page 184

- Revision History on page 185

Overview

Three main types of graceful restart available on Juniper Networks platforms are:

Graceful restart for routing protocols and summaries—Provides protection for aggregate routes, Border Gateway Protocol (BGP), Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and static routes.

Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC).

Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

Graceful restart works similarly for routing protocols and MPLS protocols and combines components of these protocol types to enable graceful restart in VPNs. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart thus allows a router to pass through intermediate convergence states that are hidden from the rest of the network.

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state information so it can resume the forwarding of network traffic. The helper router assists the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router. A brief description of graceful restart for each supported protocol follows:

BGP—When a router enabled for BGP graceful restart restarts, it retains BGP peer routes in its forwarding table and marks them as stale. However, it continues to forward traffic to other peers (or receiving peers) during the restart. To re-establish sessions, the restarting router sets the “restart state” bit in the BGP OPEN message and sends it to all participating peers. The receiving peers reply to the restarting router with messages containing end-of-routing-table markers. When the restarting router receives all replies from the receiving peers, the restarting router performs route selection, the forwarding table is updated, and the routes previously marked as stale are discarded. At this point, all BGP sessions are re-established and the restarting peer can receive and process BGP messages as usual.

While the restarting router does its processing, the receiving peers also temporarily retain routing information. Once a receiving peer detects a TCP transport reset, it retains the routes received and marks the routes as stale. After the session is re-established with the restarting router, the stale routes are replaced with updated route information.

IS-IS—Normally, IS-IS routers move neighbor adjacencies to the down state when changes occur. However, a router enabled for IS-IS graceful restart sends out Hello messages with the Restart Request (RR) bit set in a restart type length value (TLV) message. This indicates to neighboring routers that a graceful restart is in progress and that the IS-IS adjacency should be left intact. For this to work, the neighboring routers must understand and implement restart signaling themselves. Besides maintaining the adjacency, the neighbors send complete sequence number PDUs (CSNPs) to the restarting router and flood their entire database.

The restarting router never floods any of its own link-state PDUs (LSPs), including pseudonode LSPs, to IS-IS neighbors while undergoing graceful restart. This allows neighbors to re-establish their adjacencies without transitioning to the down state and allows the restarting router to re-initiate a smooth database synchronization.

OSPF—When a router enabled for OSPF graceful restart restarts, it retains routes learned prior to the restart in its forwarding table. The router does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This router continues to forward traffic to other OSPF neighbors (or *helper routers*), and sends only a limited number of LSAs during the restart period. To re-establish OSPF adjacencies with neighbors, the restarting router must send a grace LSA to all neighbors. In response, the helper routers enter helper mode and send an acknowledgement back to the restarting router. Also, if there are no topology changes, the helper routers continue to advertise LSAs as if the restarting router had remained in continuous OSPF operation.

When the restarting router receives all replies from the helper routers, the restarting router performs route selection, the forwarding table is updated, and the routes previously retained are discarded. At this point, full OSPF adjacencies are re-established and the restarting router can receive and process OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting router or the topology of the network changes, the helper routers also resume normal operation.

RSVP—This protocol uses a field called the restart capabilities object. This object is sent in RSVP Hello messages to peers and is used to advertise a router's RSVP restart capabilities. When an RSVP-enabled router restarts, a Hello message is sent to neighbors (or helper routers) to indicate a restart is in progress. Helper routers reply to the restarting router with an RSVP PATH message that contains a recovery label. The recovery label contains the information from a previous label that was advertised by the restarting node before it restarted. After receiving the recovery label, the restarting router can restore its previous forwarding state and resume operation. Note that RSVP graceful restart is not supported on ingress provider edge routers.

CCC and TCC—These two protocols rely on RSVP for all graceful restart functionality. As a result, there are no configuration statements unique to CCC and TCC. However, you can use the `show connections` command to verify that CCC and TCC graceful restart is operating. CCC and TCC graceful restart is supported on label-switched path (LSP) switch and remote interface switch connections.

LDP—This protocol uses the Fault Tolerant (FT) Session TLV as an optional parameter in the LDP Initialization message. Routers exchange this TLV during session initialization to advertise their capability to perform graceful restart or act as a helper router. When an LDP router restarts, it sends an initialization message to neighbors. The message advertises the length of time the helper routers are requested to assist the restarting router (recovery time). During this time, both routers maintain MPLS forwarding states. The restarting router marks all routes as stale and discards the routes when full neighborhood is re-established and the restart is complete.

The neighbor (or helper router) of the restarting router marks all label bindings it received from the restarting router as stale and waits for them to be refreshed or to expire at the end of the recovery time. On the helper router, a local timer governs the maximum amount of time the helper router is willing to maintain forwarding states. LDP graceful restart can be configured in a master instance or in a routing instance and supports a carrier-of-carriers scenario.

Layer 2 and Layer 3 VPNs—VPN graceful restart uses three types of restart functionality:

1. BGP graceful restart functionality is used on all provider edge (PE) to PE BGP sessions. This affects sessions carrying any service signaling data for network layer reachability information (NLRI), for example, an IPv4 VPN or Layer 2 VPN NLRI.
2. OSPF, ISIS, LDP, or RSVP graceful restart functionality is used in all core routers. Routes added by these protocols are used to resolve Layer 2 and Layer 3 VPN NLRI.
3. Protocol restart functionality is used for any Layer 3 protocol (RIP, OSPF, LDP, and so on) used between the PE and customer edge (CE) routers. This does not apply to Layer 2 VPNs because Layer 2 protocols used between the CE and PE routers do not have graceful restart capabilities.

Before VPN graceful restart can work properly, all the above components should restart gracefully. In other words, the routers should preserve their forwarding states and request neighbors to continue forwarding to the router in case of a restart. If all the above conditions are satisfied, VPN graceful restart imposes the following rules on a restarting router:

The router must wait to receive all BGP NLRI information from other PE routers before advertising routes to the CE routers.

The router must wait for all protocols in all routing instances to converge (or complete the restart process) before it sends CE router information to other PE routers. In other words, the router must wait for all instance information (whether derived from local configuration or advertisements received from a remote peer) to be processed before it sends this information to other PE routers.

The router must preserve all forwarding state in the *instance.mpls.0* tables until the new labels and transit routes are allocated and announced to other PE routers (and CE routers in a carrier-of-carriers scenario).

If any condition is not met, VPN graceful restart will not succeed in providing uninterrupted forwarding between CE routers across the VPN infrastructure.

System Requirements

To implement graceful restart, your system must meet these minimum requirements:

JUNOS software Release 5.6 or later for the CCC, Layer 2 VPN, or Layer 3 VPN implementations of graceful restart

JUNOS software Release 5.5 or later for RSVP or LDP graceful restart

JUNOS software Release 5.3 or later for aggregate route, BGP, IS-IS, OSPF, RIP, RIPng, or static route graceful restart

Two or more Juniper Networks M-series or T-series platforms

Terms and Acronyms

CCC (Circuit cross-connect)—A Juniper Networks method of exchanging frames between one router interface running a Layer 2 protocol and another router interface using the same Layer 2 protocol. For more information about CCC, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service* or the *JUNOS Internet Software Configuration Guide: MPLS Applications*.

TCC (Translational cross-connect)—A Juniper Networks method of exchanging frames between two router interfaces running different Layer 2 protocols, such as ATM, Cisco HDLC, Ethernet, Ethernet Extended VLAN, Frame Relay, and PPP. For more information about TCC, see the *JUNOS Internet Software Configuration Guide: Network Interfaces and Class of Service* or the *JUNOS Internet Software Configuration Guide: MPLS Applications*.

Restarting router—A router that experiences protocol restart because of a routing protocol process (rpd) failure or similar event. In the case of dynamic protocols such as BGP, IS-IS, OSPF, RIP, and RIPng, the neighboring routers must also support graceful restart functionality.

Helper router—A router that peers with a restarting router and cooperates to restore connections and adjacencies with the restarting router.

Configure Routing Protocol and Summaries Graceful Restart

To implement graceful restart for a routing protocol or summary, you must configure the following:

Configure Graceful Restart as a Routing Option for All Routing Protocols on page 152

Configure Graceful Restart Options for BGP on page 152

Configure Graceful Restart Options for IS-IS on page 153

Configure Graceful Restart Options for OSPF on page 153

Configure Graceful Restart Options for RIP and RIPng on page 154

To apply your knowledge, visit these sections:

Example: Layer 3 VPN Graceful Restart Configuration on page 157

Check Your Work on page 173

Configure Graceful Restart as a Routing Option for All Routing Protocols

You enable graceful restart functionality with a global statement that applies to all routing protocols, including BGP, IS-IS, OSPF, RIP and RIPv6. Helper mode is enabled by default when graceful restart is enabled. To configure graceful restart, include the graceful-restart statement at the [edit routing-options] hierarchy level. To disable graceful restart globally, include the disable statement at the [edit routing-options graceful-restart] hierarchy level. To set the amount of time the router should wait after a restart before it selects routing paths, include the path-selection-defer-time-limit at the [edit routing-options graceful-restart] hierarchy level.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    path-selection-defer-time-limit seconds;
  }
}
```

Once graceful restart is enabled for all routing protocols at the [edit routing-options graceful-restart] hierarchy level, you can disable graceful restart on a per-protocol basis.

Configure Graceful Restart Options for BGP

BGP graceful restart can be disabled on a particular neighbor in a group, an entire group, or the entire BGP protocol itself. To disable BGP graceful restart capability for all BGP sessions, include the disable statement at the [edit protocols bgp graceful-restart] hierarchy level. There is no helper mode capability for BGP.

You can also set the amount of time a router will send graceful restart messages to peer routers or receive these messages from restarting neighbors before declaring them down. To configure, include the restart-time and stale-routes-time statements at the [edit protocols bgp graceful-restart] hierarchy level. The default value for both timers is 5 minutes.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
```



Note

To set BGP graceful restart properties or disable them for a group, configure the desired statements at the [edit protocols bgp group *group-name* graceful-restart] hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, configure the desired statements at the [edit protocols bgp group *group-name* neighbor *ip-address* graceful-restart] hierarchy level.

Configure Graceful Restart Options for IS-IS

To disable IS-IS graceful restart capability, include the `disable` statement at the `[edit protocols isis graceful-restart]` hierarchy level. To disable IS-IS graceful restart helper capability, include the `disable-helper` statement at the `[edit protocols isis graceful-restart]` hierarchy level.

You can also set the amount of time a helper router waits for a neighbor to restart before declaring the restarting router down. To configure, include the `restart-duration` statement at the `[edit protocols isis graceful-restart]` hierarchy level.

```
[edit]
  protocols {
    isis {
      graceful-restart {
        disable;
        helper-disable;
        restart-duration seconds;
      }
    }
  }
}
```

Configure Graceful Restart Options for OSPF

To disable OSPF graceful restart capability, include the `disable` statement at the `[edit protocols ospf graceful-restart]` hierarchy level. To disable OSPF graceful restart helper capability, include the `disable-helper` statement at the `[edit protocols ospf graceful-restart]` hierarchy level.

You can also set the amount of time a router will send graceful restart LSAs to helper routers or receive these LSAs from restarting neighbors before declaring them down. To configure, include the `notify-duration` and `restart-duration` statements at the `[edit protocols ospf graceful-restart]` hierarchy level. The default values for these statements are 30 seconds for notification and 60 seconds for restart.

```
[edit]
  protocols {
    ospf {
      graceful-restart {
        disable;
        helper-disable;
        notify-duration seconds;
        restart-duration seconds;
      }
    }
  }
}
```

Configure Graceful Restart Options for RIP and RIPng

To disable RIP or RIPng graceful restart capability, include the disable statement at the [edit protocols (rip | ripng) graceful-restart] hierarchy level.

You can also set the amount of time a router waits for restarting neighbors before declaring them down. To configure, include the restart-time statement at the [edit protocols (rip | ripng) graceful-restart] hierarchy level.

```
[edit]
  protocols {
    ( rip | ripng ) {
      graceful-restart {
        disable;
        restart-time seconds;
      }
    }
  }
}
```

Configure Graceful Restart for an MPLS-Related Protocol

To implement graceful restart for an MPLS-related protocol, you must configure the following:

Configure Graceful Restart for all MPLS-Related Protocols on page 154

Configure Graceful Restart Options for RSVP, CCC, and TCC on page 155

Configure Graceful Restart Options for LDP on page 155

To apply your knowledge, visit these sections:

Example: Layer 3 VPN Graceful Restart Configuration on page 157

Check Your Work on page 173

Configure Graceful Restart for all MPLS-Related Protocols

You enable graceful restart functionality with a global statement that applies to all MPLS-related protocols, including CCC, LDP, RSVP, and TCC. Helper mode is enabled by default when graceful restart is enabled. To configure graceful restart, include the graceful-restart statement at the [edit routing-options] hierarchy level. To disable graceful restart globally, include the disable statement at the [edit routing-options graceful-restart] hierarchy level. To set the amount of time the router should wait after a restart before it selects routing paths, include the path-selection-defer-time-limit at the [edit routing-options graceful-restart] hierarchy level.

```
[edit]
  routing-options {
    graceful-restart {
      disable;
      path-selection-defer-time-limit seconds;
    }
  }
}
```

Once graceful restart is enabled for all MPLS-related protocols at the [edit routing-options graceful-restart] hierarchy level, you can disable graceful restart for most protocols on a per-protocol basis.

Configure Graceful Restart Options for RSVP, CCC, and TCC

Because CCC and TCC rely on RSVP, you must modify these three protocols as a single group. To disable RSVP, CCC, and TCC graceful restart capabilities, include the disable statement at the [edit protocols rsvp graceful-restart] hierarchy level. To disable RSVP, CCC, and TCC graceful restart helper capability, include the disable-helper statement at the [edit protocols rsvp graceful-restart] hierarchy level. Note that RSVP graceful restart is not supported on ingress provider edge routers.

```
[edit]
  protocols {
    rsvp {
      graceful-restart {
        disable;
        helper-disable;
      }
    }
  }
}
```

Configure Graceful Restart Options for LDP

To disable LDP graceful restart capability, include the disable statement at the [edit protocols ldp graceful-restart] hierarchy level. To disable LDP graceful restart helper capability, include the disable-helper statement at the [edit protocols ldp graceful-restart] hierarchy level.

On the restarting router, you can configure a statement that suggests to the helper routers the amount of time they are required to maintain the old forwarding state during a restart. To configure, include the recovery-time statement at the [edit protocols ldp graceful-restart] hierarchy level.

On the helper router, you can configure a statement that overrides the request from the restarting router and sets the maximum amount of time the helper router will maintain the old forwarding state. To configure, include the maximum-recovery-time statement at the [edit protocols ldp graceful-restart] hierarchy level. The default value for both timers is 140 seconds.

```
[edit]
  protocols {
    ldp {
      graceful-restart {
        disable;
        helper-disable;
        maximum-recovery-time seconds;
        recovery-time seconds;
      }
    }
  }
}
```

Configure VPN Graceful Restart

To implement graceful restart for a Layer 2 VPN or Layer 3 VPN, configure the following:

Configure Graceful Restart as a Routing Option for All Routing and MPLS-Related Protocols on page 156

Enable Graceful Restart in the Routing Instance on page 156

To apply your knowledge, visit these sections:

Example: Layer 3 VPN Graceful Restart Configuration on page 157

Check Your Work on page 173

Configure Graceful Restart as a Routing Option for All Routing and MPLS-Related Protocols

To configure graceful restart for all routing and MPLS-related protocols, include the graceful-restart statement at the [edit routing-options] hierarchy level. To disable graceful restart globally, include the disable statement at the [edit routing-options graceful-restart] hierarchy level. To set the amount of time the router should wait after a restart before it selects routing paths, include the path-selection-defer-time-limit at the [edit routing-options graceful-restart] hierarchy level.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    path-selection-defer-time-limit seconds;
  }
}
```

Enable Graceful Restart in the Routing Instance

To configure graceful restart for all routing and MPLS-related protocols within a routing instance, include the graceful-restart statement at the [edit routing-instance *routing-instance-name* routing-options] hierarchy level. Because you can configure labeled BGP, multi-instance BGP, and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported. Also, you can disable graceful restart on a per-protocol basis with the disable statement at the [edit routing-instance *routing-instance-name* protocols *protocol-name* graceful-restart] hierarchy level (To review this statement for protocols, see “Configure Routing Protocol and Summaries Graceful Restart” on page 151).

```
[edit]
routing-instances {
  routing-instance-name {
    routing-options {
      graceful-restart {
        disable;
        path-selection-defer-time-limit seconds;
      }
    }
  }
}
```

Example: Layer 3 VPN Graceful Restart Configuration

Figure 22: Layer 3 VPN Graceful Restart Topology Diagram




Figure 22 shows a standard MPLS VPN network. Routers CE1 and CE2 are customer edge routers, PE1 and PE2 are provider edge routers, and P0 is a provider core router. Several Layer 3 VPNs are configured across this network, as well as one Layer 2 VPN. Interfaces are shown in the diagram and are not included in the configuration example below.

On router CE1, configure the following protocols on the logical interfaces of t3-3/1/0: OSPF on unit 101, RIP on unit 102, BGP on unit 103, and IS-IS on unit 512. Also configure graceful restart, BGP, IS-IS, OSPF and RIP on the main instance to be able to connect to the routing instances on router PE1.

```
CE1 [edit]
interfaces {
  t3-3/1/0 {
    encapsulation frame-relay;
    unit 100 {
      dlci 100;
      family inet {
        address 10.96.100.2/30;
      }
    }
    unit 101 {
      dlci 101;
      family inet {
        address 10.96.101.2/30;
      }
    }
    unit 102 {
      dlci 102;
      family inet {
        address 10.96.102.2/30;
      }
    }
  }
}
```

```

unit 103 {
  dlc1 103;
  family inet {
    address 10.96.103.2/30;
  }
}
unit 512 {
  dlc1 512;
  family inet {
    address 10.96.252.1/30;
  }
}
}
lo0 {
  unit 0 {
    family inet {
      address 10.245.14.172/32;
      primary;
    }
    address 10.96.110.1/32;
    address 10.96.111.1/32;
    address 10.96.112.1/32;
    address 10.96.113.1/32;
    address 10.96.116.1/32;
  }
  family iso {
    address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4172.00;
  }
}
}
routing-options {
  graceful-restart;
  autonomous-system 65100;
}
protocols {
  bgp {
    group CE-PE-INET {
      type external;
      export BGP_INET_LB_DIRECT;
      neighbor 10.96.103.1 {
        local-address 10.96.103.2;
        family inet {
          unicast;
        }
      }
      peer-as 65103;
    }
  }
}
isis {
  export ISIS_L2VPN_LB_DIRECT;
  interface t3-3/1/0.512;
}
ospf {
  export OSPF_LB_DIRECT;
  area 0.0.0.0 {
    interface t3-3/1/0.101;
  }
}
}

```

```

rip {
  group RIP {
    export RIP_LB_DIRECT;
    neighbor t3-3/1/0.102;
  }
}
}
policy-options {
  policy-statement OSPF_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.101.0/30 exact;
        route-filter 10.96.111.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement RIP_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.102.0/30 exact;
        route-filter 10.96.112.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement BGP_INET_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.103.0/30 exact;
        route-filter 10.96.113.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement ISIS_L2VPN_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.116.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
}
}

```

On router PE1, configure graceful restart in the master instance, along with BGP, OSPF, MPLS, and LDP. Next configure several protocol-specific instances of graceful restart. By including instances for BGP, OSPF, Layer 2 VPNs, RIP, and static routes, you can observe the wide range of options available to you when implementing graceful restart. Configure the following protocols in individual instances on the logical interfaces of t3-0/0/0: a static route on unit 100, OSPF on unit 101, RIP on unit 102, BGP on unit 103, and Frame Relay on unit 512 for the Layer 2 VPN instance.

```

PE1 [edit]
interfaces {
  t3-0/0/0 {
    dce;
    encapsulation frame-relay-ccc;
    unit 100 {
      dlcI 100;
      family inet {
        address 10.96.100.1/30;
      }
      family mpls;
    }
    unit 101 {
      dlcI 101;
      family inet {
        address 10.96.101.1/30;
      }
      family mpls;
    }
    unit 102 {
      dlcI 102;
      family inet {
        address 10.96.102.1/30;
      }
      family mpls;
    }
    unit 103 {
      dlcI 103;
      family inet {
        address 10.96.103.1/30;
      }
      family mpls;
    }
    unit 512 {
      encapsulation frame-relay-ccc;
      dlcI 512;
    }
  }
  t1-0/1/0 {
    unit 0 {
      family inet {
        address 10.96.0.2/30;
      }
      family mpls;
    }
  }
}

```

```

lo0 {
  unit 0 {
    family inet {
      address 10.245.14.176/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4176.00;
    }
  }
}
}
routing-options {
  graceful-restart;
  router-id 10.245.14.176;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  bgp {
    group PEPE {
      type internal;
      neighbor 10.245.14.182 {
        local-address 10.245.14.176;
        family inet-vpn {
          unicast;
        }
        family l2vpn {
          unicast;
        }
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface t1-0/1/0.0;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0 {
      passive;
    }
  }
}
ldp {
  interface all;
}
}
policy-options {
  policy-statement STATIC-import {
    from community STATIC;
    then accept;
  }
  policy-statement STATIC-export {
    then {
      community add STATIC;
      accept;
    }
  }
}

```

```

policy-statement OSPF-import {
  from community OSPF;
  then accept;
}
policy-statement OSPF-export {
  then {
    community add OSPF;
    accept;
  }
}
policy-statement RIP-import {
  from community RIP;
  then accept;
}
policy-statement RIP-export {
  then {
    community add RIP;
    accept;
  }
}
policy-statement BGP-INET-import {
  from community BGP-INET;
  then accept;
}
policy-statement BGP-INET-export {
  then {
    community add BGP-INET;
    accept;
  }
}
policy-statement L2VPN-import {
  from community L2VPN;
  then accept;
}
policy-statement L2VPN-export {
  then {
    community add L2VPN;
    accept;
  }
}
community BGP-INET members target:69:103;
community L2VPN members target:69:512;
community OSPF members target:69:101;
community RIP members target:69:102;
community STATIC members target:69:100;
}
routing-instances {
  BGP-INET {
    instance-type vrf;
    interface t3-0/0/0.103;
    route-distinguisher 10.245.14.176:103;
    vrf-import BGP-INET-import;
    vrf-export BGP-INET-export;
    routing-options {
      graceful-restart;
      autonomous-system 65103;
    }
  }
}

```

```

protocols {
  bgp {
    group BGP-INET {
      type external;
      export BGP-INET-import;
      neighbor 10.96.103.2 {
        local-address 10.96.103.1;
        family inet {
          unicast;
        }
      }
      peer-as 65100;
    }
  }
}
L2VPN {
  instance-type l2vpn;
  interface t3-0/0/0.512;
  route-distinguisher 10.245.14.176:512;
  vrf-import L2VPN-import;
  vrf-export L2VPN-export;
  protocols { # Note: there is no graceful restart statement for a Layer 2 VPN instance.
    l2vpn {
      encapsulation-type frame-relay;
      site CE1-ISIS {
        site-identifier 512;
        interface t3-0/0/0.512 {
          remote-site-id 612;
        }
      }
    }
  }
}
OSPF {
  instance-type vrf;
  interface t3-0/0/0.101;
  route-distinguisher 10.245.14.176:101;
  vrf-import OSPF-import;
  vrf-export OSPF-export;
  routing-options {
    graceful-restart;
  }
  protocols {
    ospf {
      export OSPF-import;
      area 0.0.0.0 {
        interface all;
      }
    }
  }
}

```

```

RIP {
  instance-type vrf;
  interface t3-0/0/0.102;
  route-distinguisher 10.245.14.176:102;
  vrf-import RIP-import;
  vrf-export RIP-export;
  routing-options {
    graceful-restart;
  }
  protocols {
    rip {
      group RIP {
        export RIP-import;
        neighbor t3-0/0/0.102;
      }
    }
  }
}
STATIC {
  instance-type vrf;
  interface t3-0/0/0.100;
  route-distinguisher 10.245.14.176:100;
  vrf-import STATIC-import;
  vrf-export STATIC-export;
  routing-options {
    graceful-restart;
    static {
      route 10.96.110.1/32 next-hop t3-0/0/0.100;
    }
  }
}
}

```

On router P0, configure graceful restart in the main instance, along with OSPF, MPLS, and LDP. This allows the protocols enabled on the PE routers to reach each other.

```

P0 [edit]
interfaces {
  t3-0/1/3 {
    unit 0 {
      family inet {
        address 10.96.0.5/30;
      }
      family mpls;
    }
  }
  t1-0/2/0 {
    unit 0 {
      family inet {
        address 10.96.0.1/30;
      }
      family mpls;
    }
  }
}

```

```

lo0 {
  unit 0 {
    family inet {
      address 10.245.14.174/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4174.00;
    }
  }
}
}
routing-options {
  graceful-restart;
  router-id 10.245.14.174;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  ospf {
    area 0.0.0.0 {
      interface t1-0/2/0.0;
      interface t3-0/1/3.0;
      interface fxp0.0 {
        disable;
      }
      interface lo0.0 {
        passive;
      }
    }
  }
  ldp {
    interface all;
  }
}

```

On router PE2, configure BGP, OSPF, MPLS, LDP and graceful restart in the master instance. Configure the following protocols in individual instances on the logical interfaces of t1-0/1/3: a static route on unit 200, OSPF on unit 201, RIP on unit 202, BGP on unit 203, and Frame Relay on unit 612 for the Layer 2 VPN instance. Also configure protocol-specific graceful restart in all routing instances, except the Layer 2 VPN instance.

```

PE2 [edit]
interfaces {
  t3-0/0/0 {
    unit 0 {
      family inet {
        address 10.96.0.6/30;
      }
      family mpls;
    }
  }
}

```

```

t1-0/1/3 {
  dce;
  encapsulation frame-relay-ccc;
  unit 200 {
    dlcI 200;
    family inet {
      address 10.96.200.1/30;
    }
    family mpls;
  }
  unit 201 {
    dlcI 201;
    family inet {
      address 10.96.201.1/30;
    }
    family mpls;
  }
  unit 202 {
    dlcI 202;
    family inet {
      address 10.96.202.1/30;
    }
    family mpls;
  }
  unit 203 {
    dlcI 203;
    family inet {
      address 10.96.203.1/30;
    }
    family mpls;
  }
  unit 612 {
    encapsulation frame-relay-ccc;
    dlcI 612;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.245.14.182/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4182.00;
    }
  }
}
}
routing-options {
  graceful-restart;
  router-id 10.245.14.182;
  autonomous-system 69;
}

```

```

protocols {
  mpls {
    interface all;
  }
  bgp {
    group PEPE {
      type internal;
      neighbor 10.245.14.176 {
        local-address 10.245.14.182;
        family inet-vpn {
          unicast;
        }
        family l2vpn {
          unicast;
        }
      }
    }
  }
  ospf {
    area 0.0.0.0 {
      interface t3-0/0/0.0;
      interface fxp0.0 {
        disable;
      }
      interface lo0.0 {
        passive;
      }
    }
  }
  ldp {
    interface all;
  }
}

policy-options {
  policy-statement STATIC-import {
    from community STATIC;
    then accept;
  }
  policy-statement STATIC-export {
    then {
      community add STATIC;
      accept;
    }
  }
  policy-statement OSPF-import {
    from community OSPF;
    then accept;
  }
  policy-statement OSPF-export {
    then {
      community add OSPF;
      accept;
    }
  }
  policy-statement RIP-import {
    from community RIP;
    then accept;
  }
}

```

```

policy-statement RIP-export {
  then {
    community add RIP;
    accept;
  }
}
policy-statement BGP-INET-import {
  from community BGP-INET;
  then accept;
}
policy-statement BGP-INET-export {
  then {
    community add BGP-INET;
    accept;
  }
}
policy-statement L2VPN-import {
  from community L2VPN;
  then accept;
}
policy-statement L2VPN-export {
  then {
    community add L2VPN;
    accept;
  }
}
community BGP-INET members target:69:103;
community L2VPN members target:69:512;
community OSPF members target:69:101;
community RIP members target:69:102;
community STATIC members target:69:100;
}}
routing-instances {
  BGP-INET {
    instance-type vrf;
    interface t1-0/1/3.203;
    route-distinguisher 10.245.14.182:203;
    vrf-import BGP-INET-import;
    vrf-export BGP-INET-export;
    routing-options {
      graceful-restart;
      autonomous-system 65203;
    }
    protocols {
      bgp {
        group BGP-INET {
          type external;
          export BGP-INET-import;
          neighbor 10.96.203.2 {
            local-address 10.96.203.1;
            family inet {
              unicast;
            }
          }
          peer-as 65200;
        }
      }
    }
  }
}

```

```

L2VPN {
  instance-type l2vpn;
  interface t1-0/1/3.612;
  route-distinguisher 10.245.14.182:612;
  vrf-import L2VPN-import;
  vrf-export L2VPN-export;
  protocols { # Note: there is no graceful restart statement for a Layer 2 VPN instance.
    l2vpn {
      encapsulation-type frame-relay;
      site CE2-ISIS {
        site-identifier 612;
        interface t1-0/1/3.612 {
          remote-site-id 512;
        }
      }
    }
  }
}
OSPF {
  instance-type vrf;
  interface t1-0/1/3.201;
  route-distinguisher 10.245.14.182:201;
  vrf-import OSPF-import;
  vrf-export OSPF-export;
  routing-options {
    graceful-restart;
  }
  protocols {
    ospf {
      export OSPF-import;
      area 0.0.0.0 {
        interface all;
      }
    }
  }
}
RIP {
  instance-type vrf;
  interface t1-0/1/3.202;
  route-distinguisher 10.245.14.182:202;
  vrf-import RIP-import;
  vrf-export RIP-export;
  routing-options {
    graceful-restart;
  }
  protocols {
    rip {
      group RIP {
        export RIP-import;
        neighbor t1-0/1/3.202;
      }
    }
  }
}

```

```

STATIC {
  instance-type vrf;
  interface t1-0/1/3.200;
  route-distinguisher 10.245.14.182:200;
  vrf-import STATIC-import;
  vrf-export STATIC-export;
  routing-options {
    graceful-restart;
    static {
      route 10.96.210.1/32 next-hop t1-0/1/3.200;
    }
  }
}

```

On router CE2, complete the Layer 2 and Layer 3 VPN configuration by mirroring the protocols already set on PE2 and CE1. Specifically, configure the following on the logical interfaces of t1-0/0/3: OSPF on unit 201, RIP on unit 202, BGP on unit 203, and IS-IS on unit 612. Finally, configure graceful restart, BGP, IS-IS, OSPF and RIP on the main instance to be able to connect to the routing instances on PE2.

```

CE2 [edit]
interfaces {
  t1-0/0/3 {
    encapsulation frame-relay;
    unit 200 {
      dlcI 200;
      family inet {
        address 10.96.200.2/30;
      }
    }
    unit 201 {
      dlcI 201;
      family inet {
        address 10.96.201.2/30;
      }
    }
    unit 202 {
      dlcI 202;
      family inet {
        address 10.96.202.2/30;
      }
    }
    unit 203 {
      dlcI 203;
      family inet {
        address 10.96.203.2/30;
      }
    }
    unit 512 {
      dlcI 512;
      family inet {
        address 10.96.252.2/30;
      }
    }
  }
}

```

```

lo0 {
  unit 0 {
    family inet {
      address 10.245.14.180/32 {
        primary;
      }
      address 10.96.210.1/32;
      address 10.96.111.1/32;
      address 10.96.212.1/32;
      address 10.96.213.1/32;
      address 10.96.216.1/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4180.00;
    }
  }
}
}
routing-options {
  graceful-restart;
  autonomous-system 65200;
}
protocols {
  bgp {
    group CE-PE-INET {
      type external;
      export BGP_INET_LB_DIRECT;
      neighbor 10.96.203.1 {
        local-address 10.96.203.2;
        family inet {
          unicast;
        }
      }
      peer-as 65203;
    }
  }
}
isis {
  export ISIS_L2VPN_LB_DIRECT;
  interface t1-0/0/3.612;
}
ospf {
  export OSPF_LB_DIRECT;
  area 0.0.0.0 {
    interface t1-0/0/3.201;
  }
}
rip {
  group RIP {
    export RIP_LB_DIRECT;
    neighbor t1-0/0/3.202;
  }
}
}

```

```

policy-options {
  policy-statement OSPF_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.201.0/30 exact;
        route-filter 10.96.211.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement RIP_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.202.0/30 exact;
        route-filter 10.96.212.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement BGP_INET_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.203.0/30 exact;
        route-filter 10.96.213.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement ISIS_L2VPN_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.216.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
}

```

Check Your Work

To verify proper operation of graceful restart, use the following commands:

```
show bgp neighbor (for BGP graceful restart)
traceoptions (for OSPF graceful restart)
show rsvp neighbor detail (for RSVP graceful restart—helper router)
show rsvp version (for RSVP graceful restart—restarting router)
show ldp session detail (for LDP graceful restart)
show connections (for CCC and TCC graceful restart)
show route instance detail (for Layer 3 VPN graceful restart and for any protocols using
graceful restart in a routing instance)
show route protocol l2vpn (for Layer 2 VPN graceful restart)
```

The following sections show the output of some of these commands as used with the configuration example:

Router PE1 Status Prior to a Restart on page 173

Router PE1 Status During a Restart on page 178

The remaining sections display output of some of these commands generated from other examples:

Verify OSPF Graceful Restart on page 183

Verify BGP Graceful Restart on page 183

Verify CCC and TCC Graceful Restart on page 184

Router PE1 Status Prior to a Restart

Before a restart happens, check your neighbor relationships on Router PE1:

```
user@PE1> show bgp neighbor
Peer: 10.96.103.2+3785 AS 65100 Local: 10.96.103.1+179 AS 65103
Type: External State: Established Flags: <>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Export: [ BGP-INET-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
```

```

PeerAS Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.96.110.1      Local ID: 10.96.103.1      Active Holdtime: 90
  Keepalive Interval: 30
  Local Interface: t3-0/0/0.103
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI peer can save forwarding state: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table BGP-INET.inet.0 Bit: 30001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 8      Sent 3      Checked 3
  Input messages: Total 15      Updates 0      Refreshes 0      Octets 321
  Output messages: Total 18      Updates 2      Refreshes 0      Octets 450
  Output Queue[2]: 0

Peer: 10.245.14.182+4701 AS 69      Local: 10.245.14.176+179 AS 69
  Type: Internal      State: Established      Flags: <>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
  Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 1
  Peer ID: 10.245.14.182      Local ID: 10.245.14.176      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of all end-of-rib markers sent: inet-vpn-unicast l2vpn
  Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Suppressed due to damping: 0

```

```

Table bgp.12vpn.0 Bit: 20000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:       1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Suppressed due to damping: 0
Table L2VPN.12vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:       1
  Suppressed due to damping: 0
Last traffic (seconds): Received 28   Sent 28   Checked 28
Input messages: Total 2   Updates 0   Refreshes 0   Octets 86
Output messages: Total 13  Updates 10  Refreshes 0   Octets 1073
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

```

user@PE1> show route instance detail
master:
  Router ID: 10.245.14.176
  Type: forwarding          State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.96.103.1
  Type: vrf                 State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.245.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.245.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0         : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.96.101.1
  Type: vrf                 State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0           : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete

```

```

RIP:
  Router ID: 10.96.102.1
  Type: vrf                               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.245.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0                : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.96.100.1
  Type: vrf                               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.245.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0             : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
__juniper_privatel__:
  Router ID: 0.0.0.0
  Type: forwarding             State: Active

user@PE1> show route protocol l2vpn

inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

OSPF.inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete

RIP.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete

STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

800003                *[L2VPN/7] 00:06:00
> via t3-0/0/0.512, Pop      Offset: 4
t3-0/0/0.512          *[L2VPN/7] 00:06:00
> via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4

```

```

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.245.14.176:512:512:611/96
          *[L2VPN/7] 00:06:01
          Discard

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

Router PE1 Status During a Restart

Before you can verify that graceful restart is working, you must simulate a router restart. To cause the routing process to refresh and simulate a restart, use the restart routing operational mode command:

```

user@PE1> restart routing
Routing protocol daemon started, pid 3558

```

Next is some output captured during the router restart:

```

user@PE1> show bgp neighbor
Peer: 10.96.103.2      AS 65100 Local: 10.96.103.1      AS 65103
  Type: External    State: Active      Flags: <ImportEval>
  Last State: Idle   Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
PeerAS Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0

Peer: 10.245.14.182+179 AS 69    Local: 10.245.14.176+2131 AS 69
  Type: Internal    State: Established  Flags: <ImportEval>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.14.182    Local ID: 10.245.14.176    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn

```

```

NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:       10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:       1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:       2
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:       2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:       2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:       1
  Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:       1
  Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

```

user@PE1> show route instance detail
master:
  Router ID: 10.245.14.176
  Type: forwarding          State: Active
  Restart State: Pending    Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.96.103.1
  Type: vrf                 State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.245.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn               State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.245.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
OSPF:
  Router ID: 10.96.101.1
  Type: vrf                 State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0           : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN

```

```

RIP:
  Router ID: 10.96.102.1
  Type: vrf                               State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.245.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0          : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.96.100.1
  Type: vrf                               State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.245.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
__juniper_privatel__:
  Router ID: 0.0.0.0
  Type: forwarding          State: Active

```

You might notice there was no significant difference between the first set of pre-restart output and the second set of mid-restart output. To verify the restart, you need more details.

```

user@PE1> show route instance summary

```

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding	inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		l3vpn.0	0/0/0
		inet6.0	2/0/0
		l2vpn.0	0/0/0
		l2circuit.0	0/0/0
BGP-INET	vrf	BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
L2VPN	l2vpn	L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.l2vpn.0	2/0/0
OSPF	vrf	OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0
RIP	vrf	RIP.inet.0	6/0/0
		RIP.iso.0	0/0/0
		RIP.inet6.0	0/0/0

```

STATIC          vrf
                STATIC.inet.0          4/0/0
                STATIC.iso.0           0/0/0
                STATIC.inet6.0         0/0/0
__juniper_privatel__ forwarding
                __juniper_priva.inet.0 0/0/0
                __juniper_privat.iso.0 0/0/0
                __juniper_priv.inet6.0 0/0/0

user@PE1> show route protocol l2vpn

inet.0: 16 destinations, 17 routes (15 active, 1 holddown, 1 hidden)
Restart Pending: OSPF LDP

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: OSPF LDP

BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Pending: VPN

OSPF.inet.0: 7 destinations, 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN

RIP.inet.0: 6 destinations, 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN

STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
Restart Pending: LDP VPN
+ = Active Route, - = Last Active, * = Both

800001          *[L2VPN/7] 00:00:13
                > via t3-0/0/0.512, Pop          Offset: 4
t3-0/0/0.512    *[L2VPN/7] 00:00:13
                > via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: VPN L2VPN
+ = Active Route, - = Last Active, * = Both

10.245.14.176:512:512:611/96
                *[L2VPN/7] 00:00:13
                Discard

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

```

Verify OSPF Graceful Restart

To view graceful restart information for OSPF, configure OSPF traceoptions. You can configure traceoptions for OSPF with the file and flag statements at the [edit protocols ospf traceoptions] hierarchy level.

Here is the output of a traceoptions log from a OSPF restarting router:

```
Oct  8 05:20:12 Restart mode - sending grace lsas
Oct  8 05:20:12 Restart mode - estimated restart duration timer triggered
Oct  8 05:20:13 Restart mode - Sending more grace lsas
```

Here is the output of a traceoptions log from an OSPF helper router:

```
Oct  8 05:20:14 Helper mode for neighbor 192.255.5.1
Oct  8 05:20:14 Received multiple grace lsa from 192.255.5.1
```

Verify BGP Graceful Restart

To view graceful restart information for BGP sessions, use the show bgp neighbor command:

```
user@PE1> show bgp neighbor 192.255.10.1
Peer: 192.255.10.1+179 AS 64595 Local: 192.255.5.1+1106 AS 64595
  Type: Internal      State: Established   Flags: <>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ static ]
  Options: <Preference LocalAddress HoldTime GracefulRestart Damping PeerAS
Refresh>
  Local Address: 192.255.5.1 Holdtime: 90 Preference: 170
  IPsec SA Name: hope
  Number of flaps: 0
  Peer ID: 192.255.10.1      Local ID: 192.255.5.1      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 180
  Stale routes from peer are kept for: 180
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 19 Sent 19 Checked 19
  Input messages: Total 2 Updates 1 Refreshes 0 Octets 42
  Output messages: Total 3 Updates 0 Refreshes 0 Octets 116
  Output Queue[0]: 0
```

Verify CCC and TCC Graceful Restart

To view graceful restart information for CCC and TCC connections, use the `show connections` command. This sample assumes four remote interface CCC connections between CE1 and CE2:

```

user@PE1> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP

CCC Graceful restart : Restarting

Connection/Circuit      Type  St      Time last up    # Up trans
CE1-CE2-0               rmt-if Restart -----          0
  fe-1/1/0.0            intf  Up
  PE1-PE2-0             tlsp  Up
  PE2-PE1-0             rlsp  Up
CE1-CE2-1               rmt-if Restart -----          0
  fe-1/1/0.1            intf  Up
  PE1-PE2-1             tlsp  Up
  PE2-PE1-1             rlsp  Up
CE1-CE2-2               rmt-if Restart -----          0
  fe-1/1/0.2            intf  Up
  PE1-PE2-2             tlsp  Up
  PE2-PE1-2             rlsp  Up
CE1-CE2-3               rmt-if Restart -----          0
  fe-1/1/0.3            intf  Up
  PE1-PE2-3             tlsp  Up
  PE2-PE1-3             rlsp  Up

```

For More Information

To view additional information about graceful restart, see the following:

JUNOS Internet Software Configuration Guide: Routing and Routing Protocols

JUNOS Internet Software Configuration Guide: MPLS Applications

JUNOS Internet Software Configuration Guide: VPNs

R. Braden, Editor, "Resource ReSerVation Protocol (RSVP)," RFC 2205, September 1997

L. Berger, editor, "Generalized Multi-Protocol Label Swit ching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions ," RFC 3473, January 2003. (only Section 9, "Fault Handling")

S. Sangli, et. al., "Graceful Restart Mechanism f or BGP," Internet draft, draft-ietf-idr-restart-06.txt, January 2003

M. Shand, “Restart signaling for IS-IS,” Internet draft, draft-ietf-isis-restart-03.txt, March 2003

M. Leelanivas, et. al., “Graceful Restart Mechanism for LDP,” Internet draft, draft-ietf-mpls-ldp-restart-06.txt, October 2002

J. Moy, “Hitless OSPF Restart,” Internet draft, draft-ietf-ospf-hitless-restart-07.txt, March 2003

Revision History

30 June 2003—6.0R1 Release. Richard Hendricks.

2 April 2003—5.7R1 Release. Richard Hendricks.

27 December 2002—5.6R1 Release. Richard Hendricks.

22 October 2002—Added LDP, RSVP, CCC, TCC, Layer 2 VPN, and Layer 3 VPN graceful restart information and reformatted the document as a Feature Guide for JUNOS software release 5.6. Richard Hendricks.

2002 February 17—Initial Quick Start Guide document written. Tony Sinopoli.

