

Chapter 13

Multicast Data MDT

Protocol Independent Multicast (PIM) version 2 supports multicast over Layer 3 virtual private networks (VPNs) based on RFC 2547, *BGP/MPLS VPNs* and Option 2 (Multicast Domains) of Internet draft draft-rosen-vpn-mcast-06.txt, *Multicast in MPLS/BGP IP VPNs*. This implementation does not require the provider (P) routers to maintain any VPN-specific PIM information, but this lack of VPN-specific information is not optimal. The issue is that a single multicast group is defined for each VPN to carry multicast control and data traffic inside the provider core and all VPNs are mapped to this single group in the provider's space. This mapping results in the delivery of packets to each provider edge (PE) router attached to the P router even if the PE router has no receivers for traffic from a multicast group in that VPN. Each PE router must process the encapsulated VPN traffic even if the multicast packets are then dropped. This is a waste of resources, especially in environments characterized by low bandwidth links in the core or a multicast source in the VPN sending a very high volume of information (for example, high-definition television [HDTV] packets) through the core.

A data multicast distribution tree (MDT), based on section 7 of Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP IP VPNs*, solves the problem of P routers flooding unnecessary multicast information to PE routers that have no interested receivers for a particular VPN multicast group. The multicast data MDT solution requires the creation of a new tunnel by the PE router if the source exceeds a configured rate threshold parameter. All other PE routers join the new tunnel only if the PE router has receivers in the VPN for that multicast group.

This chapter provides the following information about data MDTs:

Data MDT Creation Overview on page 110

Data MDT Characteristics on page 110

Configuring Data MDTs on page 111

Examples: Configuring Data MDTs on page 113

Data MDT Configuration Statements on page 114

Data MDT Creation Overview

Initially, the PE routers discover each other in a VPN routing and forwarding (VRF) instance using the default MDT. Each PE router configuration includes in its VRF instance various parameters to control the creation of a data MDT, such as when the source traffic in the VRF instance exceeds the configured threshold rate. The PE router monitors the rate during its periodic statistics-collection cycles. If the source locally attached to the PE router in the VPN exceeds this limit, the source PE advertises the new data MDT group and new MDT with a User Datagram Protocol (UDP) type-length-vector (TLV) packet called an *MDT join TLV*. The MDT join TLV describes the source and group pair (S,G) in the VRF instance and the new data MDT group address used in the provider space. The source PE periodically announces the MDT join TLV over the default MDT for that VRF instance as long as the source is active.

All PE routers receive the MDT join TLV because it is sent over the default MDT. Only the PE routers with receivers in the VRF instance for that multicast group can join the new group, and the PE routers must join the new group to receive the multicast traffic now sent over the new MDT by the source PE. PE routers without interested receivers listed in the VRF instance ignore the MDT join TLVs.

When remote PE routers join the new data MDT group, they send a PIM join message for the new group in the provider space. The PIM join message for the new group is sent directly to the source PE router from the remote PE routers by means of PIM source-specific multicast (SSM). SSM using (S,G) is possible with data MDT instead of the default MDT any-source multicast (ASM) (*,G) because the source address is known from the UDP signaling used with data MDT.

The source PE router starts encapsulating the multicast traffic for the entries in the VRF instance using the new data MDT group after 3 seconds, allowing time for the remote PE routers to switch to the new group. The source PE router then halts the flow of multicast packets over the default MDT and the packet flow for the entries in the VRF instance source shifts to the newly created data MDT joined only by PE routers with interested receivers.

When the preconfigured conditions, such as the rate threshold, are no longer met by the source because the source stops sending or the rate falls below the threshold, the source PE stops announcing the MDT join TLVs and the PE router switches to sending on the default MDT for that VRF instance again.

Data MDT Characteristics

The maximum number of data MDTs for all VPNs on a PE router is limited to 8000, and the maximum number of data MDTs for a VRF instance is 1024. The configuration of a VRF instance can limit the number of MDTs possible. No new MDTs can be created after this limit is reached in the VRF instance, and all traffic for other sources that exceed the configured limits is still sent on the default MDT.

Creation of data MDTs depends on the monitoring of the multicast source data rate. This rate is checked once per minute, so the creation of data MDTs can be delayed up to 1 minute after a source exceeds a configured limit. In the same way, if the source data rate falls below the configured value, data MDT deletion can be delayed for up to 1 minute until the next statistics monitoring collection cycle.

Changes to the configured MDT limit value do not affect existing tunnels that exceed the new limit. MDTs that are already active remain in place until the threshold conditions are no longer met.

To remove active MDTs no longer included in a newly configured group address range, you must restart the PIM routing instance. This restart clears all remnants of the former group addresses but disrupts routing and therefore requires a maintenance window for the change.

Multicast tunnel (mt) interfaces created because of exceeded thresholds are not recreated if the routing process crashes. Therefore, graceful restart does not automatically reinstate the data MDT state. However, as soon as the periodic statistics collection reveals that the threshold condition is still exceeded, the tunnels are quickly recreated.

Configuring Data MDTs

To configure multicast data MDTs, include the `mdt` statement:

```
mdt {
  group-range multicast-prefix;
  threshold {
    group group-address {
      source source-address {
        rate threshold-rate;
      }
    }
  }
  tunnel-limit limit;
}
```

You can include the statement at the following hierarchy levels:

```
[edit routing-instances routing-instance-name protocols pim]
```

```
[edit logical-routers logical-router-name routing-instances routing-instance-name
protocols pim]
```



NOTE: Because MDT applies to VPNs and VRF instances, you cannot configure MDT statements in the master routing instance. If you configure MDT in the master routing instance, the configuration commit fails.

For an overview of routing instances and a detailed example of routing instance configuration, see the routing instances chapter of the *JUNOS Feature Guide*.

By default, creation of data MDTs is disabled.

This chapter describes the following tasks for configuring data MDTs:

Configuring the Data MDT Group Range on page 112

Configuring the Data MDT Threshold Parameters on page 112

Configuring the Data MDT Limit on page 113

Data MDTs require a correctly configured Layer 3 VPN for multicast. For more information about configuring Layer 3 VPNs for multicast, see “Configuring Multicast for Layer 3 VPNs” on page 210.

Configuring the Data MDT Group Range

The PE router implementing data MDTs for a local multicast source must establish the group range to use for data MDTs created in this VRF instance. This address range cannot overlap with any of the default MDT addresses for all VPNs on the router. If you configure overlapping group ranges, the configuration commit fails.

To configure the data MDT group range, include the `group-range` statement:

```
group-range multicast-prefix;
```

You can include this statement at the following hierarchy levels:

```
[edit logical-routers logical-router-name routing-instances routing-instance-name
 protocols pim mdt]
```

```
[edit routing-instances routing-instance-name protocols pim mdt]
```

Any multicast address range can be used as the multicast prefix, for example, 227.0.0.0/8.

Configuring the Data MDT Threshold Parameters

The PE router implementing data MDTs for a local multicast source must establish threshold limits for a multicast group and source. A multicast group can have more than one source of traffic.

To configure the data MDT threshold, include the `threshold` statement:

```
threshold {
  group group-address {
    source source-address {
      rate threshold-rate;
    }
  }
}
```

You can include this statement at the following hierarchy levels:

```
[edit logical-routers logical-router-name routing-instances routing-instance-name
 protocols pim mdt]
```

```
[edit routing-instances routing-instance-name protocols pim mdt]
```

The group address is the multicast group address to which the threshold limits apply. This could be a well-known address for a certain type of multicast traffic.

The source address is the unicast address of source of multicast information. This could be a source locally attached to or reached through the PE router. A group can have more than one source.

The rate is the threshold applied to multicast source to create a data MDT. The range is from 10 kilobits per second (Kbps) (the default) to 1 gigabit per second (Gbps) (1000000 Kbps).

Configuring the Data MDT Limit

The PE router implementing a data MDT for a local multicast source must establish a limit for the number of data MDTs created in this VRF instance. If the limit is 0 (the default), then no data MDTs are created for this VRF instance.

To configure the data MDT limit, include the tunnel-limit statement:

```
tunnel-limit limit;
```

You can include this statement at the following hierarchy levels:

```
[edit logical-routers logical-router-name routing-instances routing-instance-name
protocols pim mdt]
```

```
[edit routing-instances routing-instance-name protocols pim mdt]
```

The valid range is from 0 to 1024 for a VRF instance. There is a limit of 8000 tunnels for all data MDTs in all VRF instances on a PE router.

Examples: Configuring Data MDTs

Configure routing instance VPN-A on a PE router to use tunnel identifiers taken from the 227.0.0.0/8 multicast address range. Create a data MDT when traffic for the multicast group 224.4.4.4 from local source 10.10.20.43 exceeds a threshold rate of 10 Kbps. Only 10 tunnels can be in use for this routing instance at any one time:

```
[edit routing-instances VPN-A protocols pim]
mdt {
  group-range 227.0.0.0/8;
  threshold {
    group 224.4.4.4 {
      source 10.10.20.43 {
        rate 10;
      }
    }
  }
  tunnel-limit 10;
}
```

No tunnels are created if 10 tunnels already exist for this routing instance on the PE router. Tunnels are deleted when the rate of traffic from the source falls below 10 Kbps, as determined by the normal, 60-second multicast statistics-collection cycle.

Data MDT Configuration Statements

The following sections explain each data MDT configuration statement. The statements are organized alphabetically.

group

Syntax	<pre>group <i>group-address</i> { source <i>source-address</i> { rate <i>threshold-rate</i>; } }</pre>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols pim mdt threshold], [edit routing-instances <i>routing-instance-name</i> protocols pim mdt threshold]
Description	The multicast group address to which the threshold limits apply. This is typically a well-known address for a certain type of multicast traffic.
Options	<i>group-address</i> —Group address to limit. The remaining statements are explained separately.
Usage Guidelines	See “Configuring the Data MDT Threshold Parameters” on page 112.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

group-range

Syntax	<pre>group-range <i>multicast-prefix</i>;</pre>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols pim mdt], [edit routing-instances <i>routing-instance-name</i> protocols pim mdt]
Description	Establish the group range to use for data MDTs created in this VRF instance. This address range cannot overlap the default MDT addresses of any other VPNs on the router. If you configure overlapping group ranges, the configuration commit fails.
Options	<i>multicast-prefix</i> —Multicast address range to identify data MDTs. Range: Any valid, nonreserved multicast address range Default: None
Usage Guidelines	See “Configuring the Data MDT Group Range” on page 112.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

mdt

Syntax	<pre>mdt { group-range <i>multicast-prefix</i>; threshold { group <i>group-address</i> { source <i>source-address</i> { rate <i>threshold-rate</i>; } } } tunnel-limit <i>limit</i>; }</pre>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
Description	Establish group range for data MDTs and threshold limits for a multicast group and source. A multicast group can have more than one source of traffic.
Options	The remaining statements are explained separately.
Usage Guidelines	See “Configuring Data MDTs” on page 111.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

rate

Syntax	<i>rate threshold-rate</i> ;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols pim mdt threshold group <i>group-address</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim mdt threshold group <i>group-address</i> source <i>source-address</i>]
Description	Rate threshold applied to multicast source to create a data MDT.
Options	<i>threshold-rate</i> —Rate in Kbps to apply to source. Range: 10 Kbps through 1 Gbps (1000000 Kbps) Default: 10 Kbps
Usage Guidelines	See “Configuring the Data MDT Threshold Parameters” on page 112.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

source

Syntax	source <i>source-address</i> { rate <i>threshold-rate</i> ; }
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols pim mdt threshold group <i>group-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim mdt threshold group <i>group-address</i>]
Description	Unicast address of source of multicast information.
Options	<i>source-address</i> —Unicast address of source. The remaining statement is explained separately.
Usage Guidelines	See “Configuring the Data MDT Threshold Parameters” on page 112.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

threshold

Syntax	threshold { group <i>group-address</i> { source <i>source-address</i> { rate <i>threshold-rate</i> ; } } }
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols pim mdt], [edit routing-instances <i>routing-instance-name</i> protocols pim mdt]
Description	Establish threshold limits for a multicast group and source. A multicast group can have more than one source of traffic.
Options	The remaining statements are explained separately.
Usage Guidelines	See “Configuring the Data MDT Threshold Parameters” on page 112.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

tunnel-limit

Syntax	tunnel-limit <i>limit</i> ;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols pim mdt], [edit routing-instances <i>routing-instance-name</i> protocols pim mdt]
Description	Limit the number of data MDTs created in this VRF instance. If the limit is 0, then no data MDTs are created for this VRF instance.
Options	<i>limit</i> —Maximum number of data MDTs for this VRF instance. Range: 0 through 1024 Default: 0 (No data MDTs are created for this VRF instance.)
Usage Guidelines	See “Configuring the Data MDT Limit” on page 113.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

