

Chapter 25

MSDP Configuration Guidelines

To configure the Multicast Source Discovery protocol (MSDP), include the `msdp` statement:

```
msdp {
  active-source-limit {
    maximum number;
    threshold number;
  }
  data-encapsulation (disable | enable);
  disable;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  rib-group group-name;
  traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
  peer address {
    active-source-limit {
      maximum number;
      threshold number;
    }
    authentication-key peer-key;
    default-peer;
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
    }
  }
  group group-name {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    mode (mesh-group | standard);
    traceoptions {
```


Configuring Multiple Rendezvous Points in a Domain on page 274

Configuring MSDP Data Encapsulation on page 276

Configuring the MSDP Active Source Limit on page 277

Configuring a Default MSDP Peer on page 279

Disabling MSDP on page 280

Tracing MSDP Protocol Traffic on page 280

For a configuration example, see “Example: Configuring MSDP” on page 282.

Minimum MSDP Configuration

To enable MSDP on the router, include at least the following statements:

```
msdp {
  local-address address;
  peer address;
}
```

You can include these statements at the following hierarchy levels:

[edit protocols]

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

You must configure at least one peer. The peer and the local-address statements are required. You should also configure the router to be a Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP). For more information about configuring PIM, see “PIM Configuration Guidelines” on page 189.

Enabling MSDP

To enable MSDP peering on the router, include the msdp statement:

```
msdp {
  local-address address;
  peer address;
  rib-group group-name;
}
```

You can include this statement at the following hierarchy levels:

[edit protocols]

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

To associate with MSDP a routing table group that imports and exports routes into the specified routing table group, include the `rib-group` statement. The routing table group is a group that you defined with the `rib-groups` statement at the [edit routing-options] hierarchy level. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring MSDP Peers

An MSDP router must know which routers are its peers. You define the peer relationships explicitly by configuring the neighboring routers that are the MSDP peers of the local router. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. You must configure at least one peer for MSDP to function.

To configure MSDP peers, include the `peer` statement:

```
peer address {
  active-source-limit {
    maximum number;
    threshold number;
  }
  authentication-key peer-key;
  default-peer;
  disable;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
}
```

The `peer` and the `local-address` statements are required.

You can configure MSDP peers globally or for a group:

Globally for all MSDP peers at the [edit protocols msdp] or [edit logical-routers *logical-router-name* protocols msdp] hierarchy level

In a group at the [edit protocols msdp group *group-name*] or [edit logical-routers *logical-router-name* protocols msdp group *group-name*] level

If you configure MSDP peers in a group, each individual peer in a group inherits all group-level options.

Configuring MSDP Groups

You can arrange MSDP peers into groups. Each group must contain at least one peer. Arranging peers into groups is useful if you want to block sources from some peers and accept them from others, or set tracing options on one group and not others.

To configure MSDP groups, include one or more of the following statements:

```
group group-name {
  disable;
  export [ policy-names ];
  import [ policy-names ];
  local-address address;
  mode <(mesh-group | standard)>;
  traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
  }
  peer address; {
    active-source-limit {
      maximum number;
      threshold number;
    }
    authentication-key peer-key;
    default-peer;
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

```
[edit protocols msdp]
```

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

The local-address and peer statements are mandatory. The individual statements are discussed in separate sections.

Configuring MSDP Mesh Groups

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if mesh-group is not specified.

To configure an MSDP mesh group, define a peer group, and include the mode mesh-group statement:

```
group group-name {
    local-address address;
    mode mesh-group;
    peer address;
}
```

You can include this statement at the following hierarchy levels:

```
[edit protocols msdp]
```

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



CAUTION: When configuring MSDP mesh groups, you must configure all members the same. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by the internal MSDP peer via the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.

Figure 35 illustrates source-active message flooding between different mesh groups and peers within the same mesh group. Table 7 explains how flooding is handled by peers in this configuration.

Figure 35: Source-Active Message Flooding

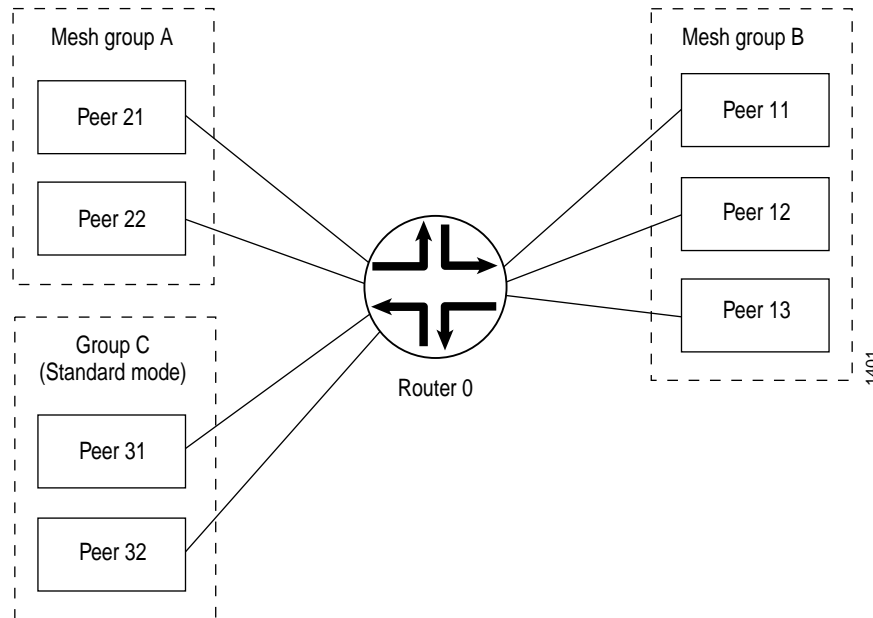


Table 7: Source-Active Message Flooding Explanation

Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message NOT Flooded To
Peer 21	Peer 11, Peer 12, Peer 13, Peer 31, Peer 32	Peer 22
Peer 11	Peer 21, Peer 22, Peer 31, Peer 32	Peer 12, Peer 13
Peer 31	Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32	

Configuring the MSDP Authentication Key

By default, multicast routers accept and process any properly formatted MSDP messages from the configured peer address. This default behavior might violate the security policies in many organizations because MSDP messages by definition come from another routing domain beyond the control of the security practices of the multicast router's organization.

The router can authenticate MSDP messages using the TCP message digest 5 (MD5) signature option for MSDP peering sessions. This authentication provides protection against spoofed packets being introduced into an MSDP peering session. Two organizations implementing MSDP authentication must decide on a human-readable key on both peers. This key is included in the MD5 signature computation for each MSDP segment sent between the two peers.

You configure an MSDP authentication key on a per-peer basis, whether the MSDP peer is defined in a group or individually. If you configure different authentication keys for the same peer at the [edit protocols msdp] and [edit protocols msdp group] hierarchy levels, the authentication key configured at the [edit protocols msdp] hierarchy level is used.

To configure MSDP authentication keys on the router, include the authentication-key statement:

```
authentication-key peer-key;
```

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

The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (,), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (" ").

The following example configures the MSDP authentication key grandmother for MSDP peer 10.0.0.1, and the MSDP authentication keys New York and phoenix5 for peers 172.16.0.1 and 192.168.0.1 in MSDP group msdp-one:

```
msdp {
  group msdp-one {
    peer 171.16.0.1 {
      authentication-key "New York";
      local-address 10.100.0.2;
    }
    peer 192.168.0.1 {
      authentication-key phoenix5;
      local-address 10.100.0.2;
    }
    peer 10.0.0.1 {
      authentication-key grandmother;
      local-address 10.100.0.2;
    }
  }
}
```

Adding, removing, or changing an MSDP authentication key in a peering session resets the existing MSDP session and establishes a new session between the affected MSDP peers. This immediate session termination prevents excessive retransmissions and eventual session timeouts due to mismatched keys.

Configuring MSDP Routing Policy

All routing protocols use the routing table to store the routes that they learn and to determine which routes they should advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in, and retrieve from, the routing table. For information about routing policy, see the *JUNOS Routing Protocols Configuration Guide*.

You can configure routing policy globally, for a group, or for an individual peer:

Globally for all MSDP peers at the [edit protocols msdp] or [edit logical-routers *logical-router-name* protocols msdp] hierarchy level.

For all peers in a group at the [edit protocols msdp group *group-name*] or [edit logical-routers *logical-router-name* protocols msdp group *group-name*] level.

For an individual peer at the [edit protocols msdp peer *address*] or [edit logical-routers *logical-router-name* protocols msdp peer *address*] level, or the [edit protocols msdp group *group-name* peer *address*] or [edit logical-routers *logical-router-name* protocols msdp group *group-name* peer *address*] level.

If you configure routing policy at the group level, each individual peer in a group inherits the group's routing policy.

To apply policies to source-active messages being imported into the source-active cache from MSDP, include the import statement, listing the names of one or more policy filters to be evaluated. See Table 8 for a list of match conditions.

Table 8: MSDP Source-Active Message Filter Match Conditions

Match Condition	Matches On
interface	Router interface or interfaces specified by name or IP address
neighbor	Neighbor address (the source address in the IP header of the source-active message)
route-filter	Multicast group address embedded in the source-active message
source-address-filter	Multicast source address embedded in the source-active message

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the route. If no match is found, MSDP shares with the routing table only those routes that were learned from MSDP routers.

```
import [ policy-names ];
```

To apply policies to source-active messages being exported from the source-active cache into MSDP, include the export statement, listing the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching policy is applied to the source-active cache entry. If no match is found, the default MSDP export policy is applied to entries in the source-active cache.

```
export [ policy-names ];
```

Configuring Multiple Rendezvous Points in a Domain

You can configure multiple RPs in a shared-tree PIM sparse-mode domain. You need to configure an MSDP local address to enable the RPs in the domain to maintain a consistent view of the active sources.

To configure a router to act as an RP in a domain with other RPs, do the following for each router in the domain that acts as an RP:

Create the router ID by configuring a unique and routable IP address on the loopback interface and setting the primary address flag.

Configure a non-unique, but routable, unicast address on the loopback interface.

Use the non-unique, routable unicast address to configure the PIM router to be the local RP.

Configure MSDP with the unique and routable address (router ID) as the local address of the peer.



NOTE: You cannot configure a local RP in a logical router. You can configure a static RP in a logical router only if the logical router is directly connected to the receiver and not connected to a source.

For a sample configuration of multiple RPs, see “Example: Configuring a Router to Use Anycast RP” on page 274. For more information about configuring interfaces, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

Example: Configuring a Router to Use Anycast RP

The following example configures a router to use anycast RP:

```
[edit]
interfaces {
  ...

  lo0 {
    unit 0 {
      family inet {
        unique routable address [and] router-id;
        address 10.1.1.1/32 {
          primary;
        }
        non-unique routable anycast RP address;
        address 10.10.10.10/32;
        address 127.0.0.1/32;
      }
    }
  }
}
```

```

routing-options {
  interface-routes {
    rib-group ifrg;
  }
  rib-groups {
    ifrg {
      import-rib [inet.0 inet.2];
    }
    mcrg {
      export-rib inet.2;
      import-rib inet.2;
    }
  }
  autonomous-system 1234;
}

protocols {
  bgp {
    group red {
      type internal;
      family inet any;
      neighbor 10.1.1.2 {
        local-address 10.1.1.1;
      }
    }
  }
  msdp {
    rib-group mcrg;
    group red {
      peer 10.1.1.2 {
        local-address 10.1.1.1;
      }
    }
  }
  pim {
    dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
    }
    rib-group mcrg;
    rp {
      local {
        address 10.10.10.10;
      }
    }
    interface all {
      mode sparse-dense;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

Configuring MSDP Data Encapsulation

MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have problems with the timeout of state relationships between sources and their multicast groups (S, G). Routers lose data while they attempt to reestablish (S, G) state tables. So multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

However, data encapsulation creates both a multicast forwarding cache entry in the inet.1 table (this is also the forwarding table) and a routing table entry in the inet.4 table. Without data encapsulation, MSDP creates only a routing table entry in the inet.4 table. In some circumstances, such as the presence of Internet worms or other forms of denial-of-service (DoS) attack, the router's forwarding table may fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation, multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

To configure MSDP data encapsulation on the router, include the data-encapsulation statement:

```
data-encapsulation (enable | disable);
```

You can include this statement at the following hierarchy levels:

```
[edit protocols msdp]
```

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

You should also configure the router to be a PIM sparse-mode RP. For more information about configuring PIM, see "PIM Configuration Guidelines" on page 189.

Configuring the MSDP Active Source Limit

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based DoS attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.

By default, the router accepts 25,000 source active messages before ignoring the rest to prevent a possible DoS attack. The limit can be from 1 to 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers. By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1,000 messages are screened by the RED profile and the accepted messages processed.

To configure the MSDP active source limit on the router, include the `active-source-limit` statement:

```
active-source-limit {  
    maximum number;  
    threshold number;  
}
```

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

The number configured for the threshold should be less than the number configured for the maximum number of active MSDP sources.

You can configure an active source limit at several levels of the MSDP hierarchy:

Configuring Global, Group, and Peer Active Source Limit on page 278

Configuring Per-Source Active Source Limit on page 278

Configuring Global, Group, and Peer Active Source Limit

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy, all are applied.

The following example applies a limit of 5,000 active sources to MSDP peer 10.0.0.1, a limit of 7,500 active sources to MSDP peer 10.10.10.10 in group MSDP-group, and a limit of 10,000 active sources to all others.

```
[edit protocols msdp]
active-source-limit {
  maximum 10000;
}
group MSDP-group {
  peer 10.10.10.10;
  active-source-limit {
    maximum 7500;
  }
  peer 10.10.10.11;
}
peer 10.0.0.1 {
  active-source-limit {
    maximum 5000;
  }
}
```

Configuring Per-Source Active Source Limit

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

```
[edit protocols msdp]
source 10.1.1.1/32 {
  active-source-limit {
    maximum 10000;
  }
}
source 10.1.0.0/16 {
  active-source-limit {
    maximum 500;
  }
}
source 0.0.0.0/0 {
  active-source-limit {
    maximum 5;
  }
}
```

In this example, the source 10.1.1.1 is allowed active sources for 10,000 groups. Any other source on the 10.1.0.0/16 network is allowed 500 groups. All other sources are allowed to source 5 active streams.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source
- Per-peer or group
- Per-instance

An active source message must “pass” all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is configured with a limit of 5,000 (and there are no other sources or limits configured), only 5,000 active source messages are accepted from this source.

Configuring a Default MSDP Peer

When a source-active message is received, a peer-RPF check is performed to make sure the peer is leading toward the originating RP and to decide whether the source-active message should be accepted. However, in networks with only one MSDP peer, especially stub networks, there is no question that the source-active message should be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check.

To establish an MSDP peer as the default peer, include the default-peer statement:

```
default-peer;
```

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

You can establish a default peer at the peer or group level. For more information about MSDP peers and peer-RPF checks, see “MSDP Overview” on page 263.

Disabling MSDP

To disable MSDP on the router, include the disable statement:

```
disable;
```

You can disable MSDP as follows:

Globally for all MSDP peers at the [edit protocols msdp] or [edit logical-routers *logical-router-name* protocols msdp] hierarchy level

For all peers in a group at the [edit protocols msdp group *group-name*] or [edit logical-routers *logical-router-name* protocols msdp group *group-name*] level

For an individual peer at the [edit protocols msdp peer *address*] or [edit logical-routers *logical-router-name* protocols msdp peer *address*] level, or the [edit protocols msdp group *group-name* peer *address*] or [edit logical-routers *logical-router-name* protocols msdp group *group-name* peer *address*] level

If you disable MSDP at the group level, each peer in the group is disabled.

Tracing MSDP Protocol Traffic

To trace MSDP protocol traffic, you can specify options in the global traceoptions statement at the [edit routing-options] or [edit logical-routers *logical-router-name* routing-options] hierarchy level, and you can specify MSDP-specific options by including the traceoptions statement. Options applied at the routing options level trace all packets, and options applied at the protocol level trace only IGMP traffic.

```
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

You can configure MSDP tracing as follows:

Globally for all MSDP peers at the [edit protocols msdp] hierarchy level

For all peers in a group at the [edit protocols msdp group *group-name*] level

For an individual peer at the [edit protocols msdp peer *address*] or the [edit protocols msdp group *group-name* peer *address*] level

If you configure tracing options at the group level, each peer in the group inherits the group's tracing options.

You can configure tracing options globally for all MSDP peers (at the [edit protocols msdp] hierarchy level), for all peers in a group (at the [edit protocols msdp group *group-name*] level), or for an individual peer (at the [edit protocols msdp peer *address*] or the [edit protocols msdp group *group-name* peer *address*] level). If you configure tracing options at the group level, each peer in the group inherits the group's tracing options.

You can specify the following MSDP-specific options in the flag statement:

- keepalive—Trace keepalive messages.
- packets—Trace all MSDP packets.
- route—Trace MSDP changes to the routing table.
- sa—Trace source-active packets.
- sa-request—Trace source-active request packets.
- sa-response—Trace source-active response packets.

For general information about tracing, see the *JUNOS System Basics Configuration Guide*.

Example: Tracing MSDP Protocol Traffic

Trace only unusual or abnormal operations to routing-log, and trace detailed information about all MSDP messages to msdp-log:

```
[edit]
routing-options {
  traceoptions {
    file routing-log;
    flag errors;
  }
}
protocols {
  msdp {
    peer 192.68.2.120; {
      local-address 192.68.1.200;
    }
    traceoptions {
      file msdp-log;
      flag packets;
    }
  }
}
```

Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
  interface-routes {
    rib-group ifrg;
  }
  rib-groups {
    ifrg {
      import-rib [inet.0 inet.2];
    }
    mcrg {
      export-rib inet.2;
      import-rib inet.2;
    }
  }
}
protocols {
  bgp {
    group lab {
      type internal;
      family any;
      neighbor 192.168.6.18 {
        local-address 192.168.6.17;
      }
    }
  }
}
pim {
  dense-groups {
    224.0.1.39/32;
    224.0.1.40/32;
  }
  rib-group mcrg;
  rp {
    local {
      address 192.168.1.1;
    }
  }
  interface all {
    mode sparse-dense;
    version 1;
  }
}
msdp {
  rib-group mcrg;
  group lab {
    peer 192.168.6.18 {
      local-address 192.168.6.17;
    }
  }
}
}
```