



Monitoring and Maintaining Multicast HA Operations

Last Updated: December 22, 2011

This module describes IPv4 and IPv6 multicast high availability (HA) support and the concepts and tasks necessary to monitor and maintain multicast HA operations.

Multicast HA capabilities enable Cisco nonstop forwarding (NSF) with stateful switchover (SSO) support for IPv4 and IPv6 multicast, which--following a Route Processor (RP) switchover--reduces the reconvergence time of the multicast control plane to a level that is transparent to most multicast-based applications and In Service Software Upgrade (ISSU) support for Protocol Independent Multicast (PIM).

- [Finding Feature Information, page 1](#)
- [Prerequisites for Monitoring and Maintaining HA Operations, page 1](#)
- [Restrictions for Monitoring and Maintaining HA Operations, page 2](#)
- [Information About Monitoring and Maintaining Multicast HA Operations, page 2](#)
- [How to Monitor and Maintain Multicast HA Operations, page 10](#)
- [Configuration Examples for Monitoring and Maintaining Multicast HA Operations, page 14](#)
- [Additional References, page 24](#)
- [Feature Information for Monitoring and Maintaining Multicast HA Operations, page 25](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Monitoring and Maintaining HA Operations

- This module assumes that your device is configured for IP multicast and is participating in an IP multicast network. For more information about configuring IP multicast using PIM sparse mode (PIM-



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

SM), Source Specific Multicast (PIM-SSM), or bidirectional PIM (bidir-PIM), see the “ Configuring a Basic IP Multicast Network ” module.

- SSO must be configured and working properly. If you do not have SSO enabled, see the “ Stateful Switchover ” module.
- This module assumes that you are familiar with NSF concepts. For more information about NSF, see the “ Cisco Nonstop Forwarding ” module.
- This module assumes that you are familiar with the ISSU process.

Restrictions for Monitoring and Maintaining HA Operations

- Multicast IPv6 multicast SSO is supported only for PIM-SSM mode and PIM sparse mode using static RP configuration. SSO for bidir-PIM is not supported for IPv6 multicast.

Information About Monitoring and Maintaining Multicast HA Operations

- [Multicast HA Support Differences from Other Routing Protocols, page 2](#)
- [Multicast Graceful Restart Overview, page 2](#)
- [Multicast Checkpointing for HA Operations, page 3](#)
- [PIM Triggered Joins, page 3](#)
- [Multicast HA Operations, page 5](#)
- [ISSU Support for IP Multicast, page 10](#)

Multicast HA Support Differences from Other Routing Protocols

Multicast HA support is different than HA support for other routing protocols because multicast routing (mroute) state is dynamic; that is, mroute state depends on the presence of sources and receivers. At the beginning of SSO, multicast state information known by downstream PIM neighbors is refreshed by the control plane. In addition, mroute state creation can be triggered by data driven events (DDEs) in the following cases:

- Mroute state creation triggered on the first hop designated router (DR) as a result of active source traffic.
- Shortest path tree (SPT) switchovers on the last hop DR; this occurs when traffic on the shared tree is detected on the last hop router.

Mroute states created in these data driven event cases are not learned from PIM join and prune messages from PIM neighbors.

Multicast Graceful Restart Overview

Multicast Graceful Restart (GR) is achieved with a combination of the NSF/SSO--IPv4 Multicast feature, the NSF/SSO--IPv6 feature, and the PIM Triggered Joins feature.

- [Multicast NSF, page 3](#)
- [Multicast SSO, page 3](#)

- [PIM Triggered Joins, page 3](#)

Multicast NSF

Multicast NSF ensures uninterrupted flow of multicast traffic during an RP failure.

Multicast SSO

Multicast SSO ensures that necessary information such as RP information, DDEs, and other multicast information is checkpointed to ensure the seamless takeover of the standby RP after an RP failover.

PIM Triggered Joins

The PIM Triggered Joins feature is used to trigger adjacent PIM neighbors on an interface to send PIM join messages for all (*, G) and (S, G) mroutes that use that interface as a reverse path forwarding (RPF) interface.

Multicast Checkpointing for HA Operations

The following multicast information is synchronized between the active and standby RPs:

- Dynamically learned group-to-RP mappings learned from either Auto-RP or bootstrap router (BSR) (IPv4 only).
- Bidir-PIM designated forwarder (DF) information and bidir-PIM RP route information (IPv4 only).
- Multicast Call Admission Control (MCAC) reservation synchronization information (IPv6 only).
- Multicast VPN (MVPN) tunnel information (IPv4 only).
- Multicast Distribution Tree (MDT) data group state information (IPv4 only).
- PIM register tunnel information (IPv4/IPv6).
- Multicast forwarding state created by DDEs (IPv4/IPv6).

PIM Triggered Joins

The PIM Triggered Joins features is a multicast HA enhancement that improves the convergence of mroutes after an RP switchover. In the event of an RP switchover, this feature utilizes the Generation ID (GenID) value (defined in RFC 4601) as a mechanism to trigger adjacent PIM neighbors on an interface to send PIM join messages for all (*, G) and (S, G) mroutes that use that interface as an RPF interface. This mechanism enables PIM neighbors to immediately reestablish those states on the newly active RP.

- [Generation IDs, page 3](#)
- [PIM Triggered Joins Functional Overview, page 4](#)
- [PIM Triggered Joins and Multicast Traffic Flow, page 4](#)

Generation IDs

A Generation ID (GenID) is a randomly generated 32-bit value that is regenerated each time PIM forwarding is started or restarted on an interface. In order to process the GenID value in PIM hello messages, PIM neighbors must be running software with an implementation of PIM that is compliant with RFC 4601.


**Note**

PIM neighbors that are not compliant with RFC 4601 and are unable to process GenID differences in PIM hello messages will ignore the GenIDs.

PIM Triggered Joins Functional Overview

The figure illustrates the operations that occur after an RP switchover on routers that support the PIM Triggered Joins feature.

Figure 1 Operation of PIM Triggered Joins During an RP Switchover



The mechanics of the PIM Triggered Joins feature are as follows:

- In steady state, PIM neighbors exchange periodic PIM hello messages.
- An active RP receives PIM joins periodically to refresh mroute states.
- When an active RP fails, the standby RP takes over to become the new active RP.
- The new active RP then modifies the GenID value and sends the new GenID in PIM hello messages to adjacent PIM neighbors.
- Adjacent PIM neighbors that receive PIM hello messages on an interface with a new GenID then send PIM triggered joins for all (*, G) and (S, G) mroutes that use that interfaces as an RPF interface.
- Those mroute states are then immediately reestablished on the newly active RP.

PIM Triggered Joins and Multicast Traffic Flow

Multicast traffic flow on the neighbors is not affected if it detects a PIM triggered join or PIM hello message from a node with the failing RP within the default PIM hello hold-time interval. Multicast traffic flow on a failing RP is not affected if it is NSF capable.

**Caution**

The default PIM hello hold-time interval is 3.5 x the PIM hello period. Multicast HA operations may not behave properly if the PIM hello interval is configured with a value lower than the default value of 30 seconds.

Multicast HA Operations

The following sections describe the multicast NSF/SSO operations that occur before, during, and following an RP switchover:

- [Dynamic Multicast SSO Synchronization Events That Occur During Normal Operation, page 5](#)
- [MFIB Interactions on the Active and Standby RPs Before an RP Switchover, page 6](#)
- [Unicast and Multicast NSF and SSO Hold-Off Period, page 6](#)
- [MFIB Interactions During an RP Switchover, page 7](#)
- [Unicast and Multicast NSF and SSO Events That Occur Following an RP Switchover, page 7](#)
- [PIM and MFIB Interactions Following an RP Switchover to Replay DDEs, page 8](#)
- [Operation After the RP Switchover, page 9](#)

Dynamic Multicast SSO Synchronization Events That Occur During Normal Operation

During normal operation (steady state), the software dynamically synchronizes information corresponding to events that modify the multicast forwarding state on the standby RP. Instead of performing periodic bulk synchronization updates, the software sends updates only for modified entities within internal databases. These updates are triggered by events that cause internal database changes related to the multicast forwarding state.

**Note**

This functionality applies only to the dynamic synchronization on the standby RP for updates to the multicast forwarding state that occur during steady state operation. Bulk synchronization updates, however, are required whenever a standby RP is inserted, reloaded, or reset.

In steady state, the following internal multicast forwarding databases are dynamically synchronized on the standby RP:

- RP Mapping--Internal database that stores group-to-RP mapping information (IPv4 only).
- Bidirectional Route Information--Internal database that stores bidir-PIM RP route information (IPv4 only).
- Bootstrap Cache--Internal database that stores BSR candidate information (IPv4 only).
- AutoRP Discovery IDB--Internal database that stores Auto-RP discovery message information (IPv4 only).
- RPDF--Internal database that stores the set of interfaces enabled for the reception of bidir-PIM packets for a given bidir-PIM RP (IPv4 only).
- MDT Tunnel--Internal database that stores MVPN MDT tunnel information (IPv4 only).
- PIM Register Tunnel--Internal database that stores PIM register tunnel information (IPv4/IPv6).

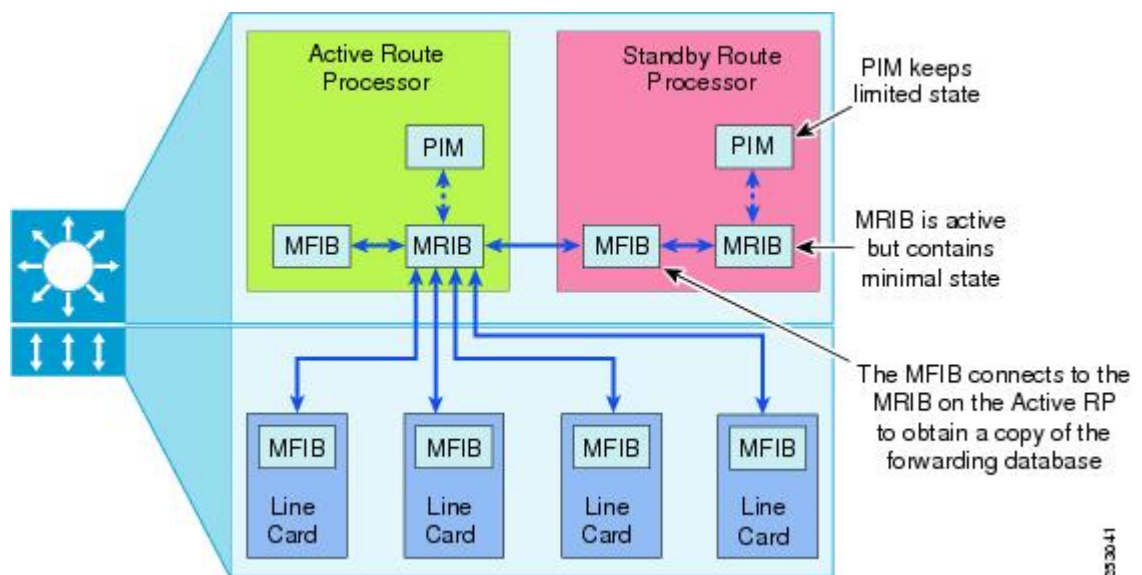
- MCAC Reservation--Internal database that stores the identity of IPv6 (S, G) multicast routes for which a MCAC cost is currently accrued for each interface on the active RP (IPv6 only).

MFIB Interactions on the Active and Standby RPs Before an RP Switchover

Before an RP switchover, each Multicast Forwarding Information Base (MFIB) instance keeps a permanent record of DDEs it generated that are passed through the Multicast Routing Information Base (MRIB) on the active RP to the MFIB on the standby RP.

The figure illustrates the multicast NSF/SSO interactions between the MFIB components on the active and standby RPs before a switchover.

Figure 2 MFIB Interactions on the Active and Standby RPs Before an RP Switchover



Unicast and Multicast NSF and SSO Hold-Off Period

Following an RP failure, data plane forwarding information is retained despite the fact that the new primary RP does not have a complete set of control plane information. The retention of this information enables forwarding to continue during unicast and multicast routing protocol reconvergence. While unicast and multicast routing protocol reconvergence is in progress, a hold-off period is observed during which no multicast forwarding updates are sent from the multicast routing protocol layer to the data plane layer. The hold-off period ends after unicast and multicast protocol convergence has completed.

Unicast routing protocol convergence begins before multicast protocol convergence. Multicast routing protocol (PIM) convergence does not begin until the multicast protocol layer receives explicit signaling that unicast routing protocol convergence has completed. Unicast protocols that are not SSO-aware are not covered by this signal and are not taken into account when waiting for convergence.



Note

Some SSO-aware routing protocols (for example, Border Gateway Protocol (BGP)) may generate logging messages indicating that the initial convergence has completed (based on an internal timer) before full convergence has occurred. PIM, however, does not provide any explicit indication of reconvergence.

The hold-off period may terminate before full convergence of unicast routing protocols, which will result in null RPF interfaces for any affected IP addresses. As additional unicast routing updates are received, the affected multicast routes are updated as needed. This is expected and acceptable behavior for SSO-aware routing protocols that are slower in converging.



Note


An RP switchover occurring on a system operating with unicast protocols that are not SSO-aware will cause undesirably long convergence times--but no routing loops--for multicast routes.

At the end of the hold-off period, the multicast data plane layer marks any existing data plane information as stale. That information is subsequently flushed if it is not refreshed through the downloading of the current reconverged control plane information.

MFIB Interactions During an RP Switchover

During an RP switchover, while the routing protocols are reconverging, no changes to the multicast tables will occur. All MFIB instances will enter NSF mode, as illustrated in the figure.

Figure 3 MFIB Interactions During an RP Switchover



Unicast and Multicast NSF and SSO Events That Occur Following an RP Switchover

In the event of an RP switchover, even with the continuous synchronization of unicast and multicast routing information from the primary to the standby RP, it is not possible to guarantee that the information most recently updated on the primary RP can be synchronized to the standby RP before a failure occurs on the

primary RP. For this reason, following an RP switchover, both unicast and multicast routing protocols trigger the retransmission of routing information from neighboring routers to ensure that the unicast and multicast routing information is current.

For multicast protocol retransmission, the software triggers a refresh of all multicast routing information available from PIM neighbors using the PIM GenID capability described in RFC 4601. GenID support enables fast mroute reconvergence after a switchover. A GenID is a randomly generated 32-bit value regenerated each time PIM forwarding is started or restarted on an interface. In the event of a switchover, the GenID value is used as a mechanism to trigger adjacent PIM neighbors on an interface to send PIM join messages for all (*, G) and (S, G) mroutes that use that interface as an RPF interface, immediately reestablishing those states on the new primary RP. Internet Group Management Protocol (IGMP) for IPv4 multicast and Multicast Listener Discovery (MLD) group membership information for IPv6 multicast is restored by executing IGMP/MLD queries on all IGMP/MLD interfaces.

**Note**

For more information about GenID support, see the [PIM Triggered Joins](#), page 3 section.

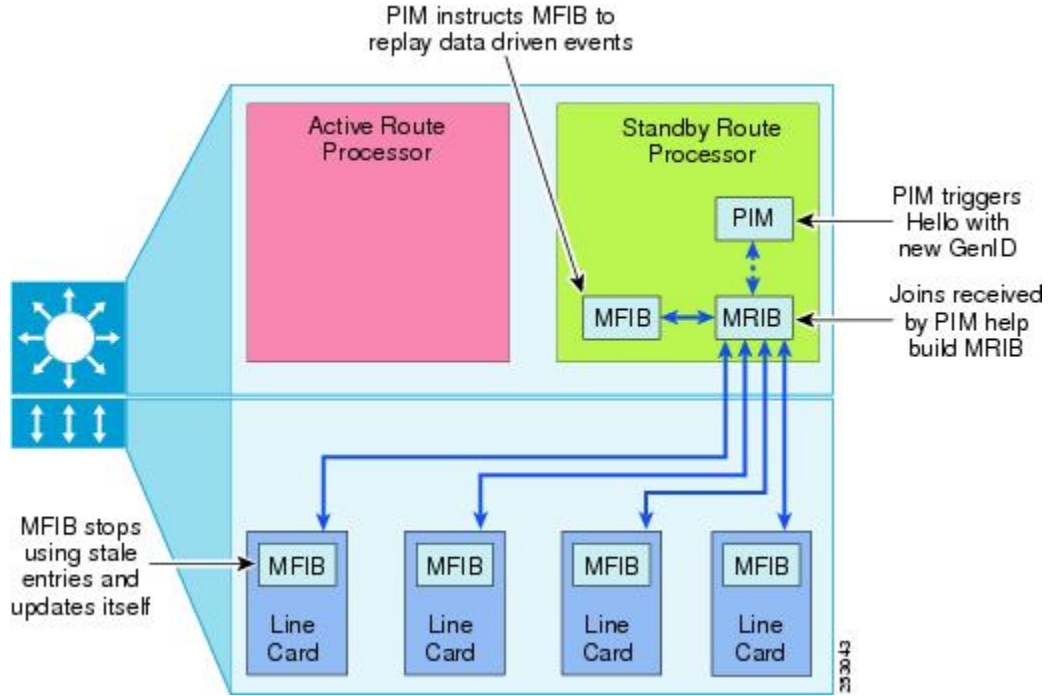
The following multicast NSF/SSO events occur in parallel following an RP switchover:

- The software empties the queue containing unprocessed synchronization messages for multicast sent by the previous primary RP and starts a *unicast IGP convergence fail-safe timer* to handle the possibility that unicast Interior Gateway Protocol (IGP) convergence never completes.
- As interfaces come up on the new primary RP, unicast routing protocol reconvergence processing proceeds.
- As each PIM-enabled interface comes up, PIM hello messages are sent out using a new GenID value for the interface. The modified GenID value triggers PIM join and prune messages from all adjacent PIM neighbors on the network to which the interface is attached. As these messages are received, information about mroute states that were missing on the new primary RP are restored except for last hop SPT (S, G) routes and mroutes associated with directly connected hosts with no other intermediate routers. Because this routing information begins to arrive before unicast IGP convergence has occurred, mroutes may initially have NULL RPF ingress interfaces. As this state information is learned, the multicast protocol layer sends the corresponding update messages to the MRIB.
- IGMP/MLD group membership information is restored by the execution of IGMP/MLD queries on all IGMP/MLD interfaces.
- Following IGMP/MLD reporting, the control plane then sends out requests for the MFIB replay of DDEs to retrigger multicast route information that cannot be obtained from PIM neighbors or directly connected hosts.
- After DDE replay, the hold-off period ends. At the end of the hold-off period, the multicast data plane layer marks any existing data plane information as stale and that information is subsequently flushed if it is not refreshed via the downloading of the current reconverged control plane information.

PIM and MFIB Interactions Following an RP Switchover to Replay DDEs

The underlying components that make up the MFIB infrastructure coordinate to ensure successful multicast NSF/SSO operations. In particular, the internal exchange of instructions between PIM and the MFIB, as illustrated in the figure, ensure error-free operation and the successful replay of DDEs.


Figure 4 PIM and MFIB Interactions Following an RP Switchover



Operation After the RP Switchover

The new RP (the previous active RP that went down) will work as the standby RP after the repair, reboot, or reinstallation, as shown in the figure.

Figure 5 PIM and MFIB Interactions Following an RP Switchover



ISSU Support for IP Multicast

The ISSU process allows software to be updated or otherwise modified while packet forwarding continues. In most networks, planned software upgrades are a significant cause of downtime. ISSU allows software to be modified while packet forwarding continues, which increases network availability and reduces downtime caused by planned software upgrades.

To provide the required ISSU and SSO support necessary for IP multicast, a PIM ISSU client is introduced. The PIM ISSU client resides on both the primary and the standby RPs and enables PIM synchronization message transmission between two RPs using different versions of software. The PIM ISSU client performs transformation of PIM dynamic state synchronization messages sent from or received by the RP having the most recent software version. If synchronization messages are sent to a RP not using the most recent software version, the messages are translated to the older format used by this RP. If messages are received from this RP, the messages are translated to the newer format used by the receiving RP before being passed to the PIM HA software for processing.

How to Monitor and Maintain Multicast HA Operations

- [Monitoring Multicast HA Events, page 10](#)
- [Modifying the Stale Mroute Flush Timeout Period for Multicast HA Operations, page 12](#)

Monitoring Multicast HA Events

Perform this optional task to monitor multicast HA NSF/SSO and ISSU events.

SUMMARY STEPS

1. `enable`
2. `debug ip multicast redundancy [verbose]`
3. `show ip pim neighbor`
4. `show ip multicast redundancy state [verbose]`
5. `show ip multicast redundancy statistics`
6. `clear ip multicast redundancy statistics`

DETAILED STEPS

-
- Step 1** `enable`
Enables privileged EXEC mode. Enter your password if prompted.

Example:

```
Router> enable
```

- Step 2** `debug ip multicast redundancy [verbose]`
Use this command to display IP multicast redundancy events.

This command logs events that are important in verifying the operation of NSF/SSO operation for IP multicast. The classes of events logged by **debug ip multicast redundancy** command include SSO events during an RP switchover and dynamic synchronization events that occur during steady state operation.

Use the optional **verbose** keyword to log events that may occur frequently during normal operation, but that may be useful for tracking in short intervals.

The following is output from the **debug ip multicast redundancy** command. The output displays the logging message that is displayed when the standby RP is recovered after a standby RP transition:

Example:

```
*Aug 7 02:36:07.843: MCAST-HA-RF: Status event: status=RF_STATUS_OPER_REDUNDANCY_MODE_CHANGE
Op=7 RFState=ACTIVE
```

Step 3

show ip pim neighbor

Use this command to display the PIM neighbors discovered by PIMv1 router query messages or PIMv2 hello messages that support the GenID capability.

The output of the **show ip pim neighbor** command displays the “G” flag to indicate GenID support status for each PIM neighbor. The “G” flag is displayed only if the neighbor supports the GenID capabilities provided by PIM.

GenID support enables fast mroute reconvergence after a switchover. A GenID is a randomly generated 32-bit value regenerated each time PIM forwarding is started or restarted on an interface. In the event of a switchover, the GenID value is used as a mechanism to trigger adjacent PIM neighbors on an interface to send PIM join messages for all (*, G) and (S, G) mroutes that use that interface as an RPF interface, immediately reestablishing those states on the newly active RP.

Example:

```
Router# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable
Neighbor      Interface      Uptime/Expires   Ver   DR
Address
192.168.10.5   GigabitEthernet0/1      00:01:35/00:01:37 v2    1 / DR B S P G
```

Step 4

show ip multicast redundancy state [verbose]

Use this command to display the current redundancy state for IP multicast.

The output displays information about the current multicast redundancy state of the RPs and the current synchronization state of the standby RP.

Example:

```
Router# show ip multicast redundancy state
Multicast IPv4 Redundancy Mode: SSO
Multicast IPv6 Redundancy Mode: Not enabled
Multicast IPv4 HA state machine status: Idle
Multicast IPv6 HA state machine status: Idle
Sync message epoch: 0
Sync message sequence number: 24
Stale NSF state flush timeout: 30000 ms
Current sync state: Synced
Multicast ISSU Client Status:
  PIM MIC client      ISSU compatible
  MRIB MIC client     ISSU compatible
  MFIB IPv4 MIC client ISSU compatible
  MFIB IPv6 MIC client No ISSU result reported
```

```

PLATFORM IPv4 MIC client  Unregistered - ignored
PLATFORM IPv6 MIC client  Unregistered - ignored
IPv4 SSO supported for:   PIM, MRIB, MFIBV4
IPv6 SSO blocked by:     MFIBV6

```

Step 5 **show ip multicast redundancy statistics**

Use this command to display IP multicast redundancy statistics.

The output displays the following information:

- A summary statistic showing the current number of synchronization messages awaiting transmission from the active RP to the standby RP. (This count is summed across all synchronization database types.)
- A summary statistic showing the current number of synchronization messages that have been sent from the active RP to the standby RP, but for which the active RP has not yet received acknowledgment from the standby for successful reception. (This count is summed across all synchronization database types.)
- The last two statistics, displaying the count of messages awaiting transmission or acknowledgment, provide a way to measure the load on the internal synchronization message-sending mechanism.

Example:

```

Router# show ip multicast redundancy statistics
Multicast Redundancy Statistics
Sync Type           Updates      Syncs        Sync failures
RP mapping           0             0             0
Bidir. RP route info 0             0             0
Bootstrap cache      4             4             0
Autorp discovery IDB 4             4             0
RPDF                 0             0             0
MDT tunnel           0             0             0
PIM register tunnel  13            13            0
MCAC Reservation     0             0             0
Requests Awaiting Sync Msg Transmission: 0
Requests Awaiting Sync Msg Acknowledgement: 0
Average Sync Wait Time = 1 ms
Average Sync Ack Time = 7 ms

```

Step 6 **clear ip multicast redundancy statistics**

Use this command to reset IP multicast redundancy statistics.

Example:

```

Router# clear ip multicast redundancy statistics

```

Modifying the Stale Mroute Flush Timeout Period for Multicast HA Operations

Perform this optional task to configure an additional timeout period before stale forwarding plane mroute information is flushed. This timeout period is added on to the default NSF route flush time as a delay between the downloading of refreshed multicast control plane route information to the forwarding plane and the flushing of “stale” NSF forwarding plane information retained from SSO before the RP switchover.



Caution

It is not recommended that you configure this additional delay unless it is specifically required for your topology because it could increase the risk of routing loops during NSF.



Note

You would need to perform this task only if you have a routing protocol that requires additional time to populate routing information after the signaling of unicast routing convergence (for example, BGP in a configuration with a large number of VPN routing and forwarding (VRF) instances). The need to configure this timeout period may be determined during predeployment SSO stress testing.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip multicast redundancy routeflush maxtime *seconds***
4. **end**
5. show ip multicast redundancy state

DETAILED STEPS

Command or Action	Purpose
<p>Step 1 enable</p> <p>Example:</p> <pre>Router> enable</pre>	<p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> • Enter your password if prompted.
<p>Step 2 configure terminal</p> <p>Example:</p> <pre>Router# configure terminal</pre>	<p>Enters global configuration mode.</p>
<p>Step 3 ip multicast redundancy routeflush maxtime <i>seconds</i></p> <p>Example:</p> <pre>Router(config)# ip multicast redundancy routeflush maxtime 900</pre>	<p>Configures an additional timeout period before stale forwarding plane mroute information is flushed following an RP switchover.</p> <ul style="list-style-type: none"> • The range is from 0 to 3600 seconds. The default is 30 seconds.
<p>Step 4 end</p> <p>Example:</p> <pre>Router(config)# end</pre>	<p>Ends the current configuration session and returns to privileged EXEC mode.</p>

Command or Action	Purpose
Step 5 show ip multicast redundancy state Example: Router# show ip multicast redundancy state	Displays the current redundancy state for IP multicast. <ul style="list-style-type: none"> Use this command to confirm the stale NSF state flush timeout period being used. The “Stale NSF state flush timeout” output field will display the timeout period setting.

Configuration Examples for Monitoring and Maintaining Multicast HA Operations

- [Example Monitoring Multicast NSF and SSO Events During an RP Switchover, page 14](#)
- [Example Monitoring the Transition from Standby RP to Active RP Following a Switchover, page 20](#)

Example Monitoring Multicast NSF and SSO Events During an RP Switchover

The following example shows how to monitor IP multicast NSF/SSO events during an RP switchover using the **debug ip multicast redundancy** command. The example shows IP multicast events occurring as a standby RP assumes the role of active RP during an SSO switchover. The events labeled “MCAST-HA” are logged by the IP multicast SSO debug facility.

- [Initial Switchover Detection, page 14](#)
- [Unicast Convergence Detection and Multicast Route Control Plane Convergence, page 15](#)
- [IGMP Queries DDE Replay Termination of the NSF Hold-Off Period and Flushing of Stale Forwarding Information, page 16](#)
- [Standby RP Bringup, page 17](#)

Initial Switchover Detection

The following output is from the **debug ip multicast redundancy** command. The output shows the initial logging messages that display when the system detects an RP switchover.

```
00:10:33: %REDUNDANCY-3-SWITCHOVER: RP switchover (PEER_DOWN_INTERRUPT)
00:10:33: %REDUNDANCY-5-PEER_MONITOR_EVENT: Standby received a switchover (raw-
event=PEER_DOWN_INTERRUPT(11))
*Aug 7 02:31:28.051: MCAST-HA: Received cf status CHKPT_STATUS_PEER_NOT_READY
*Aug 7 02:31:28.063: MCAST-HA: Received cf status CHKPT_STATUS_PEER_NOT_READY
*Aug 7 02:31:28.063: MCAST-HA-RF: Status event: status=RF_STATUS_PEER_COMM Op=0
RFState=STANDBY HOT
*Aug 7 02:31:28.063: MCAST-HA-RF: Status event:
status=RF_STATUS_OPER_REDUNDANCY_MODE_CHANGE Op=0 RFState=STANDBY HOT
*Aug 7 02:31:28.063: MCAST-HA-RF: Status event: status=RF_STATUS_REDUNDANCY_MODE_CHANGE
Op=0 RFState=STANDBY HOT
*Aug 7 02:31:28.063: MCAST-HA-RF: Status event: status=RF_STATUS_PEER_PRESENCE Op=0
RFState=STANDBY HOT
*Aug 7 02:31:28.063: MCAST-HA-RF: Status event: status=RF_STATUS_MAINTENANCE_ENABLE
Op=0 RFState=ACTIVE-FAST
```

```
*Aug 7 02:31:28.063: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_ACTIVE_FAST
RFState=ACTIVE-FAST
*Aug 7 02:31:28.091: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_ACTIVE_DRAIN
RFState=ACTIVE-DRAIN
*Aug 7 02:31:28.091: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_ACTIVE_PRECONFIG
RFState=ACTIVE_PRECONFIG
*Aug 7 02:31:28.091: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_ACTIVE_POSTCONFIG
RFState=ACTIVE_POSTCONFIG
*Aug 7 02:31:28.103: MCAST-HA: Received cf status CHKPT_STATUS_IPC_FLOW_ON
*Aug 7 02:31:28.103: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_ACTIVE
RFState=ACTIVE
```

Unicast Convergence Detection and Multicast Route Control Plane Convergence

The following output is from the **debug ip multicast redundancy** command. As interfaces come up on the new active RP, unicast convergence occurs in parallel with multicast route refresh from PIM neighbors. Unicast convergence is followed by RPF adjustments to the refreshed mroute information.

```
*Aug 7 02:31:28.107: MCAST-HA: Triggering unicast convergence notification process
handling for MVRF IPv4 default
*Aug 7 02:31:28.107: MCAST-HA: Triggering unicast convergence notification process
handling for MVRF mvrfl
*Aug 7 02:31:28.107: MCAST-HA: Triggering unicast convergence notification process
handling for MVRF mvrfl2
*Aug 7 02:31:28.107: MCAST-HA: Triggering unicast convergence notification process
handling for MVRF mvrfl3
*Aug 7 02:31:28.107: MCAST-HA: Triggering unicast convergence notification process
handling for all MVRFs
*Aug 7 02:31:28.111: MCAST-HA: Beginning unicast convergence notification process
handling.
*Aug 7 02:31:28.111: MCAST-HA: Unicast convergence completed for MVRF IPv4 default:
Triggering RPF updates
*Aug 7 02:31:28.111: MCAST-HA: Beginning unicast convergence notification process
handling.
*Aug 7 02:31:28.111: MCAST-HA: Unicast convergence completed for MVRF mvrfl: Triggering
RPF updates
*Aug 7 02:31:28.111: MCAST-HA: Beginning unicast convergence notification process
handling.
*Aug 7 02:31:28.111: MCAST-HA: Unicast convergence completed for MVRF mvrfl2: Triggering
RPF updates
*Aug 7 02:31:28.111: MCAST-HA: Beginning unicast convergence notification process
handling.
*Aug 7 02:31:28.111: MCAST-HA: Unicast convergence completed for MVRF mvrfl3: Triggering
RPF updates
*Aug 7 02:31:28.111: MCAST-HA: Unicast convergence notification has been received for
the only unconverged VRF.
Stopping the unicast routing convergence failsafe timer.
*Aug 7 02:31:28.111: MCAST-HA: Beginning unicast convergence notification process
handling.
*Aug 7 02:31:28.111: MCAST-HA: Unicast convergence notification received for the
wildcard tableid (all VRFs).
Triggering RPF updates for all MVRFs and stopping the unicast IGP convergence failsafe
timer.
00:10:34: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.1.1 on interface
Loopback0
00:10:34: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.31.10.1 on interface
Loopback1
00:10:35: %PIM-5-DRCHG: VRF mvrfl2: DR change from neighbor 0.0.0.0 to 172.16.1.1 on
interface Tunnell
00:10:35: %PIM-5-DRCHG: VRF red: DR change from neighbor 0.0.0.0 to 172.16.1.1 on
interface Tunnel2
00:10:35: %LINK-3-UPDOWN: Interface Null0, changed state to up
00:10:35: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
00:10:35: %LINK-3-UPDOWN: Interface Loopback1, changed state to up
00:10:35: %LINK-3-UPDOWN: Interface Tunnel0, changed state to up
00:10:35: %LINK-3-UPDOWN: Interface Tunnell, changed state to up
00:10:35: %LINK-3-UPDOWN: Interface Tunnel2, changed state to up
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet0/0, changed state to
administratively down
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet0/1, changed state to
```

```

administratively down
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet0/2, changed state to
administratively down
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet0/3, changed state to
administratively down
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet1/0, changed state to
administratively down
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet1/1, changed state to
administratively down
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet1/2, changed state to
administratively down
00:10:35: %LINK-5-CHANGED: Interface GigabitEthernet1/3, changed state to
administratively down
00:10:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface Null0, changed state to up
00:10:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed state to up
00:10:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback1, changed state to up
00:10:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to up
00:10:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state to up
00:10:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel2, changed state to up
00:10:38: %PIM-5-DRCHG: VRF mvrfl: DR change from neighbor 0.0.0.0 to 172.16.1.1 on
interface Tunnel0

```

IGMP Queries DDE Replay Termination of the NSF Hold-Off Period and Flushing of Stale Forwarding Information

The following output is from the **debug ip multicast redundancy** command. After the processing of unicast and multicast route convergence, time is allowed for IGMP reporting. Following IGMP reporting, the control plane then sends out requests for the MFIB replay of DDEs to retrigger multicast route information that cannot be obtained from PIM neighbors or directly connected hosts. After this processing completes, the control plane waits for the NSF hold-off time period to terminate. The refreshed multicast control plane information is then downloaded to the forwarding plane and when this is completed, the stale multicast forwarding plane information is subsequently flushed.

```

*Aug 7 02:31:43.651: MCAST-HA: IGMP response timer expired. Ready for DDE replay for
MVRF mvrfl3
*Aug 7 02:31:43.651: MCAST-HA: Sending DDE replay request for MVRF mvrfl3.
*Aug 7 02:31:43.651: MCAST-HA: MFIB DDE replay completed for mvrfl mvrfl3
*Aug 7 02:31:43.651: MCAST-HA: No NSF Holdoff extension requested for mvrfl mvrfl3 at
completion of DDE replay.
*Aug 7 02:31:43.651: MCAST-HA: Terminating multicast NSF holdoff for MVRF mvrfl3
*Aug 7 02:31:43.651: MCAST-HA: Still awaiting MFIB DDE replay for mvrfl mvrfl2
DDE replay: NOT COMPLETED, MRIB update: NOT PENDING
*Aug 7 02:31:43.651: MCAST-HA: IGMP response timer expired. Ready for DDE replay for
MVRF mvrfl2
*Aug 7 02:31:43.651: MCAST-HA: Sending DDE replay request for MVRF mvrfl2.
*Aug 7 02:31:43.651: MCAST-HA: MFIB DDE replay completed for mvrfl mvrfl2
*Aug 7 02:31:43.651: MCAST-HA: No NSF Holdoff extension requested for mvrfl mvrfl2 at
completion of DDE replay.
*Aug 7 02:31:43.651: MCAST-HA: Terminating multicast NSF holdoff for MVRF mvrfl2
*Aug 7 02:31:43.651: MCAST-HA: Still awaiting MFIB DDE replay for mvrfl mvrfl1
DDE replay: NOT COMPLETED, MRIB update: NOT PENDING
*Aug 7 02:31:43.651: MCAST-HA: IGMP response timer expired. Ready for DDE replay for
MVRF mvrfl1
*Aug 7 02:31:43.651: MCAST-HA: Sending DDE replay request for MVRF mvrfl1.
*Aug 7 02:31:43.651: MCAST-HA: MFIB DDE replay completed for mvrfl mvrfl1
*Aug 7 02:31:43.651: MCAST-HA: No NSF Holdoff extension requested for mvrfl mvrfl1 at
completion of DDE replay.
*Aug 7 02:31:43.651: MCAST-HA: Terminating multicast NSF holdoff for MVRF mvrfl1
*Aug 7 02:31:43.651: MCAST-HA: Still awaiting MFIB DDE replay for mvrfl IPv4 default
DDE replay: NOT COMPLETED, MRIB update: NOT PENDING
*Aug 7 02:31:43.651: MCAST-HA: IGMP response timer expired. Ready for DDE replay for
MVRF IPv4 default
*Aug 7 02:31:43.651: MCAST-HA: Sending DDE replay request for MVRF IPv4 default.
*Aug 7 02:31:43.651: MCAST-HA: MFIB DDE replay completed for mvrfl IPv4 default
*Aug 7 02:31:43.651: MCAST-HA: No NSF Holdoff extension requested for mvrfl IPv4 default
at completion of DDE replay.
*Aug 7 02:31:43.651: MCAST-HA: Terminating multicast NSF holdoff for MVRF IPv4 default
*Aug 7 02:31:43.651: MCAST-HA: MFIB DDE replay completed for all MVRFs.

```



```

*Aug 7 02:31:43.651: MCAST-HA: Stopping the MFIB DDE replay failsafe timer.
*Aug 7 02:32:13.651: MCAST-HA: Flush timer expired. Starting final RPF check for MVRF
IPv4 default
*Aug 7 02:32:13.651: MCAST-HA: Flush timer expired. Starting final RPF check for MVRF
mvrfl
*Aug 7 02:32:13.651: MCAST-HA: Flush timer expired. Starting final RPF check for MVRF
mvrfl2
*Aug 7 02:32:13.651: MCAST-HA: Flush timer expired. Starting final RPF check for MVRF
mvrfl3
*Aug 7 02:32:14.151: MCAST-HA: Flushing stale mcast state. RP failover processing
complete for MVRF IPv4 default.
*Aug 7 02:32:14.151: MCAST-HA: Flushing stale mcast state. RP failover processing
complete for MVRF mvrfl.
*Aug 7 02:32:14.151: MCAST-HA: Flushing stale mcast state. RP failover processing
complete for MVRF mvrfl2.
*Aug 7 02:32:14.151: MCAST-HA: Flushing stale mcast state. RP failover processing
complete for MVRF mvrfl3.
*Aug 7 02:32:14.151: MCAST-HA: RP failover processing complete for all MVRFs.

```

Standby RP Bringup

The following is sample output from the **debug ip multicast redundancy** command. This output shows events related to the reloading of the standby RP; in particular, events related to ISSU negotiation between the active and standby RP and events related to the synchronization of dynamic multicast forwarding information from the active RP to the standby RP. Synchronization events are also logged in steady state for events that occur that affect dynamic group-to-RP mapping information or dynamic tunnel state.

```

00:11:50: %HA-6-MODE: Operating RP redundancy mode is SSO
*Aug 7 02:32:45.435: MCAST-HA-RF: Status event:
status=RF_STATUS_OPER_REDUNDANCY_MODE_CHANGE Op=7 RFState=ACTIVE
*Aug 7 02:32:45.435: MCAST-HA-RF: Status event: status=RF_STATUS_REDUNDANCY_MODE_CHANGE
Op=7 RFState=ACTIVE
*Aug 7 02:32:45.435: MCAST-HA-RF: Status event: status=RF_STATUS_PEER_PRESENCE Op=1
RFState=ACTIVE
*Aug 7 02:32:45.463: MCAST-HA-RF: Status event: status=RF_STATUS_PEER_COMM Op=1
RFState=ACTIVE
*Aug 7 02:32:45.563: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_ISSU_NEGOTIATION
RFState=ACTIVE
*Aug 7 02:32:46.039: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_PLATFORM_SYNC
RFState=ACTIVE
*Aug 7 02:32:46.979: MCAST-HA: Received cf status CHKPT_STATUS_PEER_READY
*Aug 7 02:32:46.979: MCAST-ISSU Handling communication up transition for PIM HA
transport type 0, RF comm = TRUE, renegotiation NOT PENDING
*Aug 7 02:32:46.979: MCAST-HA: Received cf status CHKPT_STATUS_IPC_FLOW_ON
*Aug 7 02:32:47.043: MCAST-HA-RF: Progression event:
RF_Event=RF_PROG_STANDBY_ISSU_NEGOTIATION_LATE RFState=ACTIVE
*Aug 7 02:32:50.943: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_STANDBY_CONFIG
RFState=ACTIVE
*Aug 7 02:32:50.947: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.947: MCAST-HA-RF: Started PIM ISSU negotiation on the primary RP.
*Aug 7 02:32:50.947: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.947: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.951: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.951: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.951: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.951: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.955: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.955: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.955: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.959: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.959: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.959: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.959: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.959: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.963: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.963: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.963: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.963: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.967: MCAST-ISSU Negotiation message sent from primary, rc = 0

```

```

*Aug 7 02:32:50.967: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.967: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.967: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.967: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.971: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.971: MCAST-ISSU Negotiation message sent from primary, rc = 0
*Aug 7 02:32:50.971: MCAST-ISSU Negotiation completed for PIM Checkpoint Facility
client, negotiation rc = 4, negotiation result = COMPATIBLE
*Aug 7 02:32:59.927: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_STANDBY_FILESYS
RFState=ACTIVE
*Aug 7 02:32:59.963: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_STANDBY_BULK
RFState=ACTIVE
*Aug 7 02:32:59.963: MCAST-HA-RF: Starting Bulk Sync.
*Aug 7 02:32:59.963: MCAST-HA: Successfully created the bulk sync process
*Aug 7 02:32:59.963: MCAST-HA: Starting Bulk sync
*Aug 7 02:32:59.963: MCAST HA Executing RP mapping bulk sync.
*Aug 7 02:32:59.963: MCAST HA Executing Bidir RP route bulk sync.
*Aug 7 02:32:59.963: MCAST HA Executing BSR cache bulk sync.
*Aug 7 02:32:59.963: MCAST-HA BSR cache sync request received for mvrfl IPv4 default
*Aug 7 02:32:59.963: MCAST-HA: Creating Bootstrap cache sync request chunk size=112
max=585 align=8
*Aug 7 02:32:59.963: MCAST-HA: Allocating Bootstrap cache sync request sync request
*Aug 7 02:32:59.963: MCAST-HA Formatting BSR cache sync message:
search for mvrfl IPv4 default result is 0 mvrfl at 0x4A21680
*Aug 7 02:32:59.971: MCAST-HA BSR cache sync request received for mvrfl mvrfl1
*Aug 7 02:32:59.971: MCAST-HA: Allocating Bootstrap cache sync request sync request
*Aug 7 02:32:59.971: MCAST-HA Formatting BSR cache sync message:
search for mvrfl mvrfl1 result is 0 mvrfl at 0x50EE660
*Aug 7 02:32:59.983: MCAST-HA BSR cache sync request received for mvrfl mvrfl2
*Aug 7 02:32:59.983: MCAST-HA: Allocating Bootstrap cache sync request sync request
*Aug 7 02:32:59.983: MCAST-HA Formatting BSR cache sync message:
search for mvrfl mvrfl2 result is 0 mvrfl at 0x5103300
*Aug 7 02:32:59.991: MCAST-HA BSR cache sync request received for mvrfl mvrfl3
*Aug 7 02:32:59.991: MCAST-HA: Allocating Bootstrap cache sync request sync request
*Aug 7 02:32:59.991: MCAST-HA Formatting BSR cache sync message:
search for mvrfl mvrfl3 result is 0 mvrfl at 0x5135FE0
*Aug 7 02:33:00.003: MCAST HA Executing AutoRP discovery IDB bulk sync.
*Aug 7 02:33:00.003: MCAST-HA AutoRP discovery IDB sync request received for
mvrfl IPv4 default
*Aug 7 02:33:00.003: MCAST-HA: Creating Autorp discovery IDB sync request chunk size=112
max=585 align=8
*Aug 7 02:33:00.003: MCAST-HA: Allocating Autorp discovery IDB sync request sync request
*Aug 7 02:33:00.003: MCAST-HA Formatting AutoRP discovery IDB sync message:
search for mvrfl IPv4 default result is 0 mvrfl at 0x4A21680
*Aug 7 02:33:00.011: MCAST-HA AutoRP discovery IDB sync request received for
mvrfl mvrfl1
*Aug 7 02:33:00.011: MCAST-HA: Allocating Autorp discovery IDB sync request sync request
*Aug 7 02:33:00.011: MCAST-HA Formatting AutoRP discovery IDB sync message:
search for mvrfl mvrfl1 result is 0 mvrfl at 0x50EE660
*Aug 7 02:33:00.023: MCAST-HA AutoRP discovery IDB sync request received for
mvrfl mvrfl2
*Aug 7 02:33:00.023: MCAST-HA: Allocating Autorp discovery IDB sync request sync request
*Aug 7 02:33:00.023: MCAST-HA Formatting AutoRP discovery IDB sync message:
search for mvrfl mvrfl2 result is 0 mvrfl at 0x5103300
*Aug 7 02:33:00.031: MCAST-HA AutoRP discovery IDB sync request received for
mvrfl mvrfl3
*Aug 7 02:33:00.031: MCAST-HA: Allocating Autorp discovery IDB sync request sync request
*Aug 7 02:33:00.031: MCAST-HA Formatting AutoRP discovery IDB sync message:
search for mvrfl mvrfl3 result is 0 mvrfl at 0x5135FE0
*Aug 7 02:33:00.043: MCAST HA Executing dummy bulk sync function.
*Aug 7 02:33:00.043: MCAST HA Executing dummy bulk sync function.
*Aug 7 02:33:00.043: MCAST HA Executing dummy bulk sync function.
*Aug 7 02:33:00.043: MCAST HA Executing MDT tunnel bulk sync.
*Aug 7 02:33:00.043: MCAST-HA MDT tunnel sync request received for mvrfl mvrfl1
*Aug 7 02:33:00.043: MCAST-HA: Creating MDT tunnel sync request chunk size=112 max=585
align=8
*Aug 7 02:33:00.043: MCAST-HA: Allocating MDT tunnel sync request sync request
*Aug 7 02:33:00.043: MCAST-HA Formatting MDT tunnel sync message:
search for mvrfl mvrfl1 result is 0 mvrfl at 0x50EE660
*Aug 7 02:33:00.051: MCAST-HA MDT tunnel sync request received for mvrfl mvrfl2
*Aug 7 02:33:00.051: MCAST-HA: Allocating MDT tunnel sync request sync request
*Aug 7 02:33:00.051: MCAST-HA Formatting MDT tunnel sync message:
search for mvrfl mvrfl2 result is 0 mvrfl at 0x5103300

```

```
*Aug 7 02:33:00.063: MCAST-HA MDT tunnel sync request received for mvrf mvrf3
*Aug 7 02:33:00.063: MCAST-HA: Allocating MDT tunnel sync request
*Aug 7 02:33:00.063: MCAST-HA Formatting MDT tunnel sync message:
search for mvrf mvrf3 result is 0 mvrf at 0x5135FE0
*Aug 7 02:33:00.071: MCAST HA Executing Bidir RP DF bulk sync.
*Aug 7 02:33:00.071: MCAST HA Executing register tunnel bulk sync.
*Aug 7 02:33:00.071: MCAST-HA: Completed enqueueing of bulk sync messages.
*Aug 7 02:33:00.071: MCAST-HA: Bulk sync message queue has drained.
*Aug 7 02:33:00.071: MCAST-HA: Received acknowledgement from standby for all bulk sync
messages.
*Aug 7 02:33:00.071: MCAST-HA Creating bulk sync completion message for peer.
*Aug 7 02:33:00.071: MCAST-HA: Primary has notified standby of bulk sync completion.
Waiting for final bulk sync ACK from stby.
*Aug 7 02:33:00.075: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.075: MCAST-HA: Sent message type is 2
*Aug 7 02:33:00.075: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.075: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 2. Cleanup is complete.
*Aug 7 02:33:00.075: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.075: MCAST-HA: Sent message type is 2
*Aug 7 02:33:00.075: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.075: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 2. Cleanup is complete.
*Aug 7 02:33:00.075: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.075: MCAST-HA: Sent message type is 2
*Aug 7 02:33:00.075: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.075: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 2. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 2
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 2. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 3
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 3. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 3
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 3. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 3
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 3. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 3
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 3. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 8
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 8. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 8
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 8. Cleanup is complete.
```

```

*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 8
*Aug 7 02:33:00.087: MCAST-HA Searching for sync request corresponding to the
successfully received message.
*Aug 7 02:33:00.087: MCAST-HA Transmission from primary and reception by standby
confirmed for sync type 8. Cleanup is complete.
*Aug 7 02:33:00.087: MCAST-HA: Received cf status CHKPT_STATUS_SEND_OK
*Aug 7 02:33:00.087: MCAST-HA: Sent message type is 11
*Aug 7 02:33:00.087: MCAST-HA Process: Primary RP received standby ACK for reception of
bulk sync completion message.
*Aug 7 02:33:00.087: MCAST-HA Notifying RF to continue progression.
*Aug 7 02:33:00.087: MCAST-HA: Wakeup received for bulk sync completion.
major = 4, minor = 2.
*Aug 7 02:33:00.091: MCAST-HA Process: Primary RP received bulk sync completion
confirmation from standby.
*Aug 7 02:33:00.091: MCAST-HA RF notification previously sent.
*Aug 7 02:33:00.455: MCAST-HA-RF: Progression event: RF_Event=RF_PROG_STANDBY_HOT
RFState=ACTIVE
00:12:05: %HA_CONFIG_SYNC-6-BULK_CFGSYNC_SUCCEED: Bulk Sync succeeded
00:12:05: %HA-6-STANDBY_READY: Standby RP in slot 7 is operational in SSO mode
00:12:05: %RF-5-RF_TERMINAL_STATE: Terminal state reached for (SSO)

```

Example Monitoring the Transition from Standby RP to Active RP Following a Switchover

The following example shows how to monitor the transition from standby RP to active RP and confirm the IP multicast redundancy state and the status on the standby RP after it has resynchronized with the new active RP.



Note

In this example scenario, a router is configured for IPv4 multicast routing operation, but not for IPv6 multicast. As a result, some of the output fields that are specific to IPv6 multicast will indicate status such as “Not enabled” or “Idle” in the example outputs.

- [Initial State on Standby RP Before Switchover](#), page 20
- [Standby RP Transition to Active RP After an RP Switchover](#), page 21
- [Standby RP Resynchronization](#), page 23

Initial State on Standby RP Before Switchover

The following output is from the **show ip multicast redundancy state** command on a standby RP before an active RP goes down. In the sample output, notice that the “Current sync state” field displays “Not syncing,” indicating that the standby RP is not synchronizing data to the active RP. The standby RP serves only as a passive recipient of synchronization updates and does not initiate synchronization updates to the active RP.

```

Router_Standby# show ip multicast redundancy state
Multicast IPv4 Redundancy Mode: Not enabled
Multicast IPv6 Redundancy Mode: Not enabled
Sync message epoch: 0
Sync message sequence number: 24
Stale NSF state flush timeout: 30000 ms
Current sync state: Not syncing
Multicast ISSU Client Status:
  PIM MIC client      No ISSU result reported
  MRIB MIC client     Unregistered - ignored
  MFIB IPv4 MIC client Unregistered - ignored
  MFIB IPv6 MIC client Unregistered - ignored
  PLATFORM IPv4 MIC client Unregistered - ignored
  PLATFORM IPv6 MIC client Unregistered - ignored

```

```
IPv4 SSO blocked by:    PIM
IPv6 SSO blocked by:    PIM
```

The following output is unconditionally logged by the Redundancy Facility (RF) software when the standby RP detects that it has become the active RP due to a failure of the original active RP. The output shows the message used to indicate that an RP switchover has occurred:

```
00:00:49: %REDUNDANCY-3-SWITCHOVER: RP switchover (PEER_DOWN_INTERRUPT)
```

Standby RP Transition to Active RP After an RP Switchover

The following output is from the **show ip multicast redundancy state** command on the standby RP during its transition from standby RP to active RP. Notice that the “Multicast IPv4 HA state machine status” field displays “Unicast converging,” indicating that unicast convergence on the new active RP has begun. At this point in the RP switchover, the standby RP is waiting for unicast convergence.

```
Router-Standby# show ip multicast redundancy state
Multicast IPv4 Redundancy Mode:    Not enabled
Multicast IPv6 Redundancy Mode:    Not enabled
Multicast IPv4 HA state machine status: Unicast converging
Multicast IPv6 HA state machine status: Idle
Sync message epoch:                0
Sync message sequence number:      24
Stale NSF state flush timeout:      30000 ms
Current sync state: Not synching
Multicast ISSU Client Status:
  PIM MIC client                    No ISSU result reported
  MRIB MIC client                   No ISSU result reported
  MFIB IPv4 MIC client               No ISSU result reported
  MFIB IPv6 MIC client               No ISSU result reported
  PLATFORM IPv4 MIC client           Unregistered - ignored
  PLATFORM IPv6 MIC client           Unregistered - ignored
IPv4 SSO blocked by:                PIM, MRIB, MFIBV4
IPv6 SSO blocked by:                PIM, MRIB, MFIBV6
```

The following output from the **debug ip multicast redundancy state** command shows messages indicating that the interfaces on the new active RP are coming up. As interfaces come up on the new active RP, unicast convergence occurs in parallel with multicast route refresh from PIM neighbors. Unicast convergence is followed by RPF adjustments to the refreshed mroute information.

```
00:00:51: %LINK-3-UPDOWN: Interface Null0, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Loopback1, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel0, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel1, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel2, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel3, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel4, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel5, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel6, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel7, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel8, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel9, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel10, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel11, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel12, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel13, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel14, changed state to up
00:00:51: %LINK-3-UPDOWN: Interface Tunnel15, changed state to up
00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet0/0, changed state to
administratively down
00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet0/1, changed state to
administratively down
00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet0/2, changed state to
administratively down
00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet0/3, changed state to
administratively down
```

```

00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet1/0, changed state to
administratively down
00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet1/1, changed state to
administratively down
00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet1/2, changed state to
administratively down
00:00:51: %LINK-5-CHANGED: Interface GigabitEthernet1/3, changed state to
administratively down
00:00:51: %LINK-5-CHANGED: Interface Serial2/0, changed state to administratively down
00:00:51: %LINK-5-CHANGED: Interface Serial2/1, changed state to administratively down

```

The following is output from the **show ip multicast redundancy state** command during the transition from the standby RP to the new active RP. Notice that the “Multicast IPv4 HA state machine status” displays “DDE replaying,” indicating that the MFIB is replaying DDEs. After the processing of unicast and multicast route convergence, time is allowed for IGMP reporting. Following IGMP reporting, the control plane then sends out requests for the MFIB replay of DDEs to retrigger multicast route information that cannot be obtained from PIM neighbors or directly connected hosts.

```

Router# show ip multicast redundancy state
Multicast IPv4 Redundancy Mode: Not enabled
Multicast IPv6 Redundancy Mode: Not enabled
Multicast IPv4 HA state machine status: DDE replaying
Multicast IPv6 HA state machine status: Idle
Sync message epoch: 0
Sync message sequence number: 24
Stale NSF state flush timeout: 30000 ms
Current sync state: Not syncing
Multicast ISSU Client Status:
  PIM MIC client           No ISSU result reported
  MRIB MIC client          No ISSU result reported
  MFIB IPv4 MIC client     No ISSU result reported
  MFIB IPv6 MIC client     No ISSU result reported
  PLATFORM IPv4 MIC client Unregistered - ignored
  PLATFORM IPv6 MIC client Unregistered - ignored
  IPv4 SSO blocked by:    PIM, MRIB, MFIBV4
  IPv6 SSO blocked by:    PIM, MRIB, MFIBV6

```

After this processing completes, the control plane terminates the NSF hold-off or, if the platform multicast driver software requests an extension to the hold-off period, allows additional time for the platform multicast driver software to release the NSF hold-off extension.

The refreshed multicast control plane information is then downloaded to the forwarding plane. Although reconvergence is considered complete at this point, additional “refresh” updates may occur after this point in time. An additional time interval is provided for any remaining updates before stale multicast forwarding plane information is subsequently flushed.

The following is output from the **show ip multicast redundancy state** command. Notice that the “Multicast IPv4 HA state machine status” field displays, “Flush pending,” indicating that stale NSF data plane state is still being temporarily retained to allow for any additional refreshed multicast control plane information to be downloaded to the forwarding plane.

```

Router# show ip multicast redundancy state
Multicast IPv4 Redundancy Mode: Not enabled
Multicast IPv6 Redundancy Mode: Not enabled
Multicast IPv4 HA state machine status: Flush pending
Multicast IPv6 HA state machine status: Idle
Sync message epoch: 0
Sync message sequence number: 24
Stale NSF state flush timeout: 30000 ms
Current sync state: Not syncing
Multicast ISSU Client Status:
  PIM MIC client           No ISSU result reported
  MRIB MIC client          No ISSU result reported
  MFIB IPv4 MIC client     No ISSU result reported
  MFIB IPv6 MIC client     No ISSU result reported
  PLATFORM IPv4 MIC client Unregistered - ignored
  PLATFORM IPv6 MIC client Unregistered - ignored

```

```
IPv4 SSO blocked by:      PIM, MRIB, MFIBV4
IPv6 SSO blocked by:      PIM, MRIB, MFIBV6
```

The following is output from the **show ip multicast redundancy state** command after the refreshed multicast control plane information has been downloaded to the forwarding plane and the stale multicast forwarding plane information has been flushed. Notice that at this stage in the RP switchover the “Multicast IPv4 HA state machine status” field displays “Idle” because multicast IPv4 HA state machine operations have completed.

```
Router# show ip multicast redundancy state
Multicast IPv4 Redundancy Mode:  Not enabled
Multicast IPv6 Redundancy Mode:  Not enabled
Multicast IPv4 HA state machine status: Idle
Multicast IPv6 HA state machine status: Idle
Sync message epoch:              0
Sync message sequence number: 24
Stale NSF state flush timeout: 30000 ms
Current sync state: Not syncing
Multicast ISSU Client Status:
  PIM MIC client                  No ISSU result reported
  MRIB MIC client                  No ISSU result reported
  MFIB IPv4 MIC client             No ISSU result reported
  MFIB IPv6 MIC client             No ISSU result reported
  PLATFORM IPv4 MIC client         Unregistered - ignored
  PLATFORM IPv6 MIC client         Unregistered - ignored
  IPv4 SSO blocked by:             PIM, MRIB, MFIBV4
  IPv6 SSO blocked by:             PIM, MRIB, MFIBV6
```

Standby RP Resynchronization

The following is sample output from the **debug ip multicast redundancy** command. The output shows the messages used to indicate that a standby RP has been resynchronized.

```
00:25:42: %HA-6-MODE: Operating RP redundancy mode is SSO
00:26:04: %HA_CONFIG_SYNC-6-BULK_CFGSYNC_SUCCEED: Bulk Sync succeeded
00:26:04: %HA-6-STANDBY_READY: Standby RP in slot 7 is operational in SSO mode
00:26:04: %RF-5-RF_TERMINAL_STATE: Terminal state reached for
(SSO).
00:15:28: %RF-5-RF_TERMINAL_STATE: Terminal state reached for (SSO)
```

The following is output from the **show ip multicast redundancy state** command after the standby RP has completed resynchronization with the new active RP. Notice that the “Multicast IPv4 Redundancy Mode” field displays “SSO,” indicating that all information between the standby RP and active RP has been synchronized. Also, notice that the “Current sync state” field displays “Synched,” indicating that the standby has resynchronized with the new active RP.

```
Router# show ip multicast redundancy state
Multicast IPv4 Redundancy Mode:  SSO
Multicast IPv6 Redundancy Mode:  Not enabled
Multicast IPv4 HA state machine status: Idle
Multicast IPv6 HA state machine status: Idle
Sync message epoch:              0
Sync message sequence number: 24
Stale NSF state flush timeout: 30000 ms
Current sync state: Synched
Multicast ISSU Client Status:
  PIM MIC client                  ISSU compatible
  MRIB MIC client                  ISSU compatible
  MFIB IPv4 MIC client             ISSU compatible
  MFIB IPv6 MIC client             No ISSU result reported
  PLATFORM IPv4 MIC client         Unregistered - ignored
  PLATFORM IPv6 MIC client         Unregistered - ignored
  IPv4 SSO supported for:          PIM, MRIB, MFIBV4
  IPv6 SSO blocked by:             MFIBV6
```

Additional References

Related Documents

Related Topic	Document Title
Overview of the IP multicast technology area	“ IP Multicast Technology Overview ” module
Concepts, tasks, and examples for configuring an IP multicast network using PIM	“ Configuring a Basic IP Multicast Network ” module
ISSU concepts, tasks, and examples	“ Cisco IOS In Service Software Upgrade Process ” module
NSF concepts, tasks, and examples	“ Cisco Nonstop Forwarding ” module
Stateful switchover concepts, tasks, and examples	“ Stateful Switchover ” module
Multicast commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	<i>Cisco IOS IP Multicast Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	--

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature.	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 4601	Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	<p>http://www.cisco.com/cisco/web/support/index.html</p>

Feature Information for Monitoring and Maintaining Multicast HA Operations

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1 *Monitoring and Maintaining Multicast HA Operations*

Feature Name	Releases	Feature Information
ISSU--IPv4 Multicast	Cisco IOS XE Release 2.1	The ISSU--IPv4 Multicast feature enhances IPv4 multicast HA by providing support for ISSU.

Feature Name	Releases	Feature Information
NSF/SSO--IPv4 Multicast	Cisco IOS XE Release 2.1	<p>The NSF/SSO--IPv4 Multicast feature reduces the reconvergence time of the multicast control plane to a level that is transparent to most IPv4 multicast-based applications. Multicast NSF ensures uninterrupted flow of multicast traffic during an RP failure. Multicast SSO ensures that necessary information such as RP information and DDEs is checkpointed to ensure the seamless takeover of the standby RP after an RP failover.</p> <p>The following commands were introduced or modified: clear ip multicast redundancy statistics, debug ip multicast redundancy, ip multicast redundancy route flush maxtime, show ip multicast redundancy state, show ip multicast redundancy statistics, show ip pim neighbor.</p>
NSF/SSO--IPv6 Multicast	Cisco IOS XE Release 2.1	<p>The NSF/SSO--IPv6 Multicast feature reduces the reconvergence time of the multicast control plane to a level that is transparent to most IPv6 multicast-based applications. Multicast NSF ensures uninterrupted flow of multicast traffic during an RP failure. Multicast SSO ensures that necessary information such as RP information and DDEs is checkpointed to ensure the seamless takeover of the standby RP after an RP failover.</p> <p>The following commands were introduced or modified: clear ip multicast redundancy statistics, debug ip multicast redundancy, ip multicast redundancy route flush maxtime, show ip multicast redundancy state, show ip multicast redundancy statistics, show ip pim neighbor.</p>

Feature Name	Releases	Feature Information
PIM Triggered Joins	Cisco IOS XE Release 2.6	The PIM Triggered Joins feature is a multicast HA enhancement that improves the reconvergence of mroutes after an RP switchover. In the event of an RP switchover, this feature utilizes the PIM-SM GenID value as a mechanism to trigger adjacent PIM neighbors on an interface to send PIM join messages for all (*, G) and (S, G) mroutes that use that interface as an RPF interface, immediately reestablishing those states on the newly active RP.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2011 Cisco Systems, Inc. All rights reserved.