

# Multi-Chassis Synchronization for IGMP Snooping

---

## In This Chapter

This section provides information about multi-chassis synchronization (MCS) for IGMP snooping.

Topics in this section include:

- [Applicability on page 100](#)
- [Overview on page 101](#)
- [Configuration on page 103](#)
- [Conclusion on page 115](#)

## **Applicability**

This example is applicable to all 7x50 and 7710 platforms and was tested on release 12.0.R1.

There are no hardware dependencies for MCS.

## Overview

Multi-Chassis Synchronization (MCS) is a proprietary protocol used for synchronizing state information between two 7x50/7710 SR peers. MCS can be used for the following applications:

- IGMP
- IGMP snooping
- MLDP snooping
- IPsec
- DHCP server
- Subscriber management
- Subscriber Router Redundancy Protocol (SRRP)

The focus of this example is on using MCS to synchronize IGMP snooping state information, that is, Layer 2 multicast forwarding entries called IGMP snooping entries, between two peer nodes.

MCS also supports synchronization of IGMP state information on IGMP-enabled router interfaces, but that is outside the scope of this example.

Standard behavior in an IGMP snooping-enabled Layer 2 domain is for the system to send the IGMP reports from the IGMP receivers towards the IGMP querier along the multicast router (Mrouter) ports. The Mrouter ports are either dynamically elected (because they are on the Layer 2 logical path to the IGMP querier) or statically configured on ports in the Layer 2 domain (either on the VPLS SAP or on the SDP).

As IGMP reports are forwarded towards the IGMP querier through the Mrouter ports, multicast forwarding entries are created in the data plane for each (S,G), or (\*,G) separately. This data is stored as multicast forwarding information base (mFIB) entries in the IGMP snooping-enabled VPLS service.

Because the Layer 2 path between the IGMP querier and the IGMP receivers can change over time due to Layer 2 network failures, the multicast service (IGMP traffic and multicast streams) can be disrupted. To minimize the outage of the multicast service when the path between the IGMP receivers and IGMP querier changes between redundant nodes, MCS copies and maintains the multicast forwarding entries in the data plane between redundant nodes. It then activates the entries on the standby node when it becomes active in the event of a Layer 2 network failure. This ensures that IGMP and multicast traffic continues to flow with minimal loss. Traffic converges faster and independently from the recovery mechanism provided by the IGMP querier, which is controlled through the IGMP query timer value.

## Overview

It is worth noting that a very similar configuration (but using IGMP instead of IGMP snooping as an MCS application) can also be used to synchronize IGMP states on a Layer 3 multicast interface. This can be one way of speeding up multicast convergence; for example, if there is a failure of a PIM designated router (which is also an IGMP querier).


## Configuration

The benefits of MCS for IGMP snooping are demonstrated by comparing two scenarios: one without and one with MCS for IGMP snooping.

The redundant Layer 2 path between the IGMP querier (and the multicast source) and the IGMP receiver for both scenarios can be managed through a number of mechanisms, for example Spanning Tree Protocol (STP), G.8032 or LAG.

Configuring redundant Layer 2 paths is outside the scope of this example.

For both scenarios, the same spanning tree and VPLS snooping configuration is used.



**Figure 15: Configuration without MCS for IGMP Snooping**

The baseline configuration common to both scenarios is shown in [Figure 15](#).

The nodes PE-10, PE-20, PE-30 and PE-90 are part of the Layer 2 domain, and all have a VPLS with a set of SAPs. PE-10 is the root of the STP topology, and all the SAPs in this topology are in the forwarding state except for SAP 1/1/2:800 on PE-20. This SAP is pruned from the Layer 2 topology and is therefore in the blocked state as the following output shows.

## Configuration

```
*A:PE-20# show service id 800 sap
=====
SAP(Summary), Service 800
=====
PortId          SvcId
  Ing.  Ing.    Egr.  Egr.  Adm  Opr
                               QoS  Fltr  QoS  Fltr
-----
1/1/1:800          800    1   none  1   none  Up  Up
1/1/2:800          800    1   none  1   none  Up  Prun
-----
```

VPLS service 800 is configured on all Layer 2 nodes with IGMP snooping enabled, with the following properties:

- The multicast source is connected to PE-90.
- The multicast receiver is connected to PE-30, SAP 1/1/3:800.
- IGMP snooping is enabled in all VPLSs on PE-10, PE-20, PE-30, and PE-90.
- PE-90 sends IGMP queries with source-address 10.1.1.90, as it is enabled on the VPLS SAP.

Configuration of PE-90 with VPLS 800 sending queries on SAP 1/1/3 towards PE-10:

```
*A:7750-90>config>service>vpls# info
-----
stp
  shutdown
exit
igmp-snooping
  query-src-ip 10.1.1.90
  no shutdown
exit
sap 1/1/3:800 create
  igmp-snooping
    send-queries
  exit
exit
no shutdown
```

## Configuration of MCS without IGMP snooping

For verification purposes, the IGMP querier status is checked on every node that is part of the Layer 2 domain.

PE-10 VPLS 800 IGMP querier status:

```
*A:PE-10# show service id 800 igmp-snooping querier
```

```
=====
IGMP Snooping Querier info for service 800
=====
```

```
Sap Id           : 1/1/3:800
IP Address       : 10.1.1.90
Expires         : 171s
Up Time         : 19d 21:59:26
Version         : 2
```

```
General Query Interval : 125s
Query Response Interval : 10.0s
Robust Count          : 2
=====
```

PE-20 VPLS 800 IGMP querier status:

```
*A:PE-20# show service id 800 igmp-snooping querier
```

```
=====
IGMP Snooping Querier info for service 800
=====
```

```
Sap Id           : 1/1/1:800
IP Address       : 10.1.1.90
Expires         : 132s
Up Time         : 19d 22:00:51
Version         : 2
```

```
General Query Interval : 125s
Query Response Interval : 10.0s
Robust Count          : 2
=====
```

PE-30 VPLS 800 IGMP querier status:

```
*A:PE-30# show service id 800 igmp-snooping querier
```

```
=====
IGMP Snooping Querier info for service 800
=====
```

```
Sap Id           : 1/1/2:800
IP Address       : 10.1.1.90
Expires         : 154s
Up Time         : 0d 00:05:50
Version         : 2
```

```
General Query Interval : 125s
Query Response Interval : 10.0s
```

## Configuration of MCS without IGMP snooping

```
Robust Count          : 2
```

```
=====
```

Because of the placement of the IGMP queriers, as well as the IGMP receiver and the STP forwarding status of SAP, the mFIB of the three PEs is as shown below (when MCS IGMP snooping is not configured).

### PE-10 VPLS 800 mFIB:

```
*A:PE-10# show service id 800 mfib
=====
Multicast FIB, Service 800
=====
Source Address  Group Address          Sap/Sdp Id              Svc Id  Fwd/Blk
-----
*               *                sap:1/1/3:800           Local   Fwd
*               239.255.1.1      sap:1/1/2:800           Local   Fwd
*               *                sap:1/1/3:800           Local   Fwd
-----
```

### PE-20 VPLS 800 mFIB:

```
*A:PE-20# show service id 800 mfib
=====
Multicast FIB, Service 800
=====
Source Address  Group Address          Sap/Sdp Id              Svc Id  Fwd/Blk
-----
*               *                sap:1/1/1:800           Local   Fwd
-----
Number of entries: 1
=====
```

### PE-30 VPLS 800 mFIB:

```
*A:PE-30>show>service>id>igmp-snooping# show service id 800 mfib
=====
Multicast FIB, Service 800
=====
Source Address  Group Address          Sap/Sdp Id              Svc Id  Fwd/Blk
-----
*               *                sap:1/1/2:800           Local   Fwd
*               239.255.1.1      sap:1/1/3:800           Local   Fwd
*               *                sap:1/1/2:800           Local   Fwd
-----
```



The mFIB for VPLS 800 on PE-10 shows the following:

- The (\*,\*) multicast entry is present in the mFIB for SAP 1/1/3:800 because IGMP queries (from PE-90) are received on that SAP (which now is an Mrouter port).
- The (\*,239.255.1.1) multicast entry is present in the mFIB for SAP 1/1/2:800 because an IGMP report has been received for that group on that SAP from the IGMP receiver.
- The (\*,239.255.1.1) multicast entry is also present in the mFIB for SAP 1/1/3:800 because this is the Mrouter port for the IGMP querier connected to that SAP.

The mFIB for VPLS 800 on PE-20 shows the following:

- The (\*,\*) multicast entry is present in the mFIB on SAP 1/1/1:800 because IGMP queries (from PE-10) are received on that SAP (which now is an Mrouter port).

The mFIB for VPLS 800 on PE-30 shows the following:

- The (\*,\*) multicast entry is present in the mFIB for SAP 1/1/2:800 because IGMP queries (from PE-10) are received on that SAP (which now is an Mrouter port).
- The (\*,239.255.1.1) multicast entry is present in the mFIB for SAP 1/1/3:800 because an IGMP report has been received for that group on that SAP (connects to the receiver).
- The (\*,239.255.1.1) multicast entry is also present in the mFIB for SAP 1/1/2:800 because that SAP is an Mrouter port.

If a link failure occurs between PE-10 and PE-30, the spanning tree converges and SAP 1/1/2:800 of VPLS 800 on PE-20 transitions to the forwarding state (“up”).

```
*A:PE-20# show service id 800 sap
=====
SAP(Summary), Service 800
=====
```

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
1/1/1:800	800	20	none	1	none	Up	Up
1/1/2:800	800	20	none	1	none	Up	Up

```
-----
```

Immediately after the link failure, while the spanning tree is converging, the mFIBs on PE-10, PE-20 and PE-30 are exactly the same as before the failure.

The mFIB for VPLS 800 on PE-10 looks as follows:

```
*A:PE-10# show service id 800 mfib
=====
Multicast FIB, Service 800
=====
```

Source Address	Group Address	Sap/Sdp Id	Svc Id	Fwd/Blk
----------------	---------------	------------	--------	---------

## Configuration of MCS without IGMP snooping

```
-----  
*          *          sap:1/1/3:800          Local   Fwd  
*          239.255.1.1      sap:1/1/2:800          Local   Fwd  
*          *          sap:1/1/3:800          Local   Fwd  
-----
```

The mFIB for VPLS 800 on PE-20 looks as follows:

```
*A:PE-20# show service id 800 mfib  
=====
```

Source Address	Group Address	Sap/Sdp Id	Svc Id	Fwd/Blk
*	*	sap:1/1/1:800	Local	Fwd

```
-----  
Number of entries: 1  
=====
```

The mFIB for VPLS 800 on PE-30 looks as follows:

```
*A:PE-30>show>service>id>igmp-snooping# show service id 800 mfib  
=====
```

Source Address	Group Address	Sap/Sdp Id	Svc Id	Fwd/Blk
*	*	sap:1/1/2:800	Local	Fwd
*	239.255.1.1	sap:1/1/3:800	Local	Fwd
		sap:1/1/2:800	Local	Fwd

```
-----
```

This means that the multicast stream has not recovered yet.

For the multicast stream to reconverge, spanning tree must re-converge and an IGMP query must be received on PE-30 (sourced by the PE-90 querier) via the new Layer 2 forwarding path between the source and the receiver, which now runs via PE-20.

When PE-30 receives an IGMP query on the new Layer 2 forwarding path (via SAP 1/1/1:800) it makes this SAP an Mrouter port and this triggers an IGMP report to be sent across this newly elected Mrouter port to the querier.

The IGMP report results in a change to the mFIBs along the new Layer 2 multicast forwarding path. At that time, the mFIBs in all three PEs will look as shown below.

The mFIB for VPLS 800 on PE-10 looks as follows:

```
*A:PE-10>config>port# show service id 800 mfib  
=====
```

Source Address	Group Address	Sap/Sdp Id	Svc Id	Fwd/Blk
----------------	---------------	------------	--------	---------

```
-----
```

## Multi-Chassis Synchronization for IGMP Snooping

```
-----  
*          *          sap:1/1/3:800          Local    Fwd  
*          239.255.1.1      sap:1/1/1:800          Local    Fwd  
*                                     sap:1/1/3:800          Local    Fwd  
-----
```

The mFIB for VPLS 800 on PE-20 looks as follows:

```
*A:PE-20# show service id 800 mfib  
=====
```

Source Address	Group Address	Sap/Sdp Id	Svc Id	Fwd/Blk
*	*	sap:1/1/1:800	Local	Fwd
*	239.255.1.1	sap:1/1/1:800	Local	Fwd
		sap:1/1/2:800	Local	Fwd

```
-----
```

The mFIB for VPLS 800 on PE-30 looks as follows:

```
*A:PE-30>show# show service id 800 mfib  
=====
```

Source Address	Group Address	Sap/Sdp Id	Svc Id	Fwd/Blk
*	*	sap:1/1/1:800	Local	Fwd
*	239.255.1.1	sap:1/1/1:800	Local	Fwd
		sap:1/1/3:800	Local	Fwd

```
-----
```

The presence of an mFIB entry for group (\*,239.255.1.1) on SAP 1/1/1:800 of PE-10 and SAP 1/1/2:800 of PE-20 means that the multicast stream has now recovered.

Because the recovery of multicast streams (IGMP query triggered) depend highly on the IGMP query timer, recovery is slow.

The main purpose of multi-chassis IGMP snooping synchronization is to reduce recovery time.

## Configuration of MCS with IGMP Snooping

Figure 16 shows the MCS configuration used for IGMP snooping; the only difference from Figure 15 is that now MCS is configured between PE-10 and PE-20.

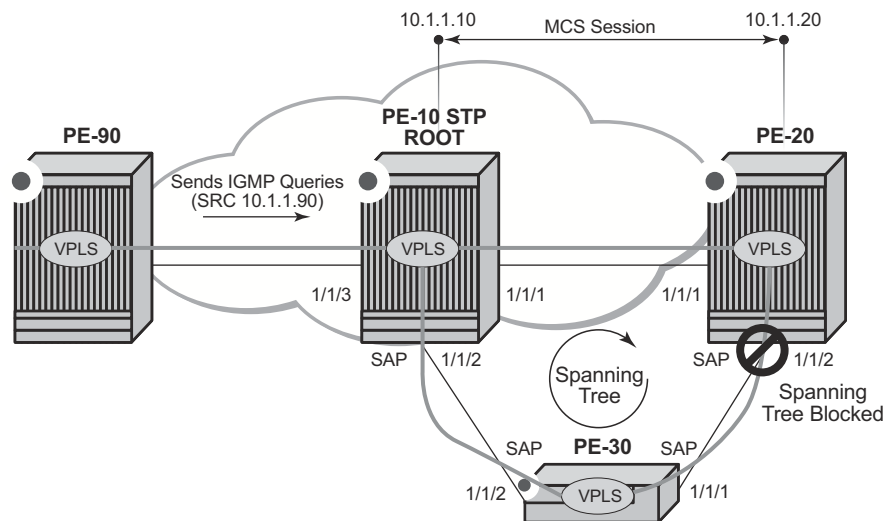


Figure 16: Configuration with MCS for IGMP Snooping

The MCS IGMP snooping configuration on PE-10 is shown below.

```
A:PE-10>config>redundancy# info
-----
multi-chassis
  peer 10.1.1.20 create
  sync
    igmp-snooping
    port 1/1/2 create
    range 800-800 sync-tag "igmp800"
  exit
  no shutdown
exit
no shutdown
exit
exit
```

The MCS IGMP snooping configuration on PE-20 is shown below.

```
*A:PE-20>config>redundancy>multi-chassis# info
-----
peer 10.1.1.10 create
  sync
  igmp-snooping
  port 1/1/2 create
    range 800-800 sync-tag "igmp800"
  exit
  no shutdown
exit
no shutdown
exit
```

This configuration results in IGMP snooping states being synchronized between PE-10 and PE-20 for port 1/1/2 VLAN 800 (VPLS 800) on PE-10 and for port 1/1/2 VLAN 800 (VPLS 800) on PE-20.

Synchronization happens for qtag 800 between these 2 ports (1/1/2 of PE10 and 1/1/2 of PE10) because both are using same sync-tag igmp800. The sync-tag parameter allows to synchronize IGMP contexts between 2 chassis independently from port and Qtag numbering.

The MCS session state can be verified on PE-10 using the following command:

```
*A:PE-10>show>redundancy>multi-chassis# all
=====
Multi-Chassis Peers
=====
Peer IP          Peer Admin   Client   Admin   Oper   State
  Src IP          Auth
-----
10.1.1.20        Enabled      MC-Sync: Enabled  Enabled  Enabled  inSync
10.1.1.10        None         MC-Ring:  --      --      --      --
                  MC-Endpt: --      --      --      --
                  MC-Lag:   Disabled Disabled  --
                  MC-IPsec: --      --      Disabled
```

Because the operational state is Enabled and State is inSync, the IGMP states for port 1/1/2 VLAN 800 on PE-20 are now synchronized with port 1/1/2 VLAN 800 on PE-10.

This can be observed on PE-20:

```
*A:PE-20# show redundancy multi-chassis sync peer 10.1.1.10 detail
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 10.1.1.10
Description          : (Not Specified)
Authentication       : Disabled
```

## Configuration of MCS with IGMP Snooping

```
Source IP Address      : 10.1.1.20
Admin State           : Enabled
-----
Sync-status
-----
Client Applications    : IGMPsnooping
Sync Admin State      : Up
Sync Oper State       : Up
Sync Oper Flags       :
DB Sync State         : inSync
Num Entries           : 1
Lcl Deleted Entries   : 0
Alarm Entries         : 0
OMCR Standby Entries  : 0
OMCR Alarm Entries    : 0
Rem Num Entries       : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries : 0
=====
..snip ..
```

The contents of the MCS database for the IGMP snooping application can be displayed on PE-20 using the following command:

```
*A:PE-20# tools dump redundancy multi-chassis sync-database application igmp-snooping
detail
```

If no entries are present for an application, no detail will be displayed.

```
FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete;
              oal - omcr alarmed; ost - omcr standby
```

```
Peer Ip 10.1.1.10
```

```
Application IGMPsnooping
Sap-id      Client Key
SyncTag     DLen  Flags      timeStamp
deleteReason code and description
-----
1/1/2:800   Group=239.255.1.1
igmp800     10    -- -- -- --- --- 05/22/2014 13:57:28
0x0
```

The following totals are for:

```
peer ip ALL, port/lag ALL, sync-tag ALL, application IGMPsnooping
Valid Entries:          1
Locally Deleted Entries: 0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries: 0
Omcr Alarmed Entries:   0
Omcr Standby Entries:   0
```

## Multi-Chassis Synchronization for IGMP Snooping

Note that an IGMP-snooping MCS entry for group 239.255.1.1 on SAP 1/1/2:800 exists, which means that there is also an IGMP mFIB entry in the data plane on PE-20, as shown in the output below.


```
*A:PE-20# show service id 800 mfib
=====
Multicast FIB, Service
800=====
Source Address  Group Address          Sap/Sdp Id              Svc Id  Fwd/Blk
-----
*               *                       sap:1/1/1:800           Local    Fwd
*               239.255.1.1          sap:1/1/1:800           Local    Fwd
*               *                       sap:1/1/2:800           Local    Fwd
-----
Number of entries: 2
=====
```

As well as an mFIB entry for group 239.255.1.1 on SAP 1/1/2:800 being created, PE-20 also automatically sends IGMP reports for that group towards the IGMP querier. This will create an additional mFIB entry for group 239.255.1.1 on SAP 1/1/1:800 on PE-10 (unlike the non-MCS-enabled scenario), as shown in the output below.

```
*A:PE-10# show service id 800 mfib
=====
Multicast FIB, Service 800
=====
Source Address  Group Address          Sap/Sdp Id              Svc Id  Fwd/Blk
-----
*               *                       sap:1/1/3:800           Local    Fwd
*               239.255.1.1          sap:1/1/1:800           Local    Fwd
*               *                       sap:1/1/3:800           Local    Fwd
*               *                       sap:1/1/2:800           Local    Fwd
-----
```

This means that the multicast stream not only is forwarded to PE-30 (via SAP 1/1/2:800 on PE-10) but also to PE-20 (via PE-10 SAP 1/1/1:800). However, PE-20 discards the stream as SAP 1/1/2:800 is blocked by the spanning tree.

## Configuration of MCS with IGMP Snooping



**Figure 17: Multicast Stream Forwarded to PE-20**

In the first scenario where MCS is not enabled, if port 1/1/2 on PE-10 fails, then SAP 1/1/2:800 on PE-20 will transition to the forwarding state once the spanning tree reconverges. However, it is only after the reception of an IGMP report from PE-30, which can take some seconds, that PE-20 sends the multicast traffic to PE-30.

Enabling MCS for IGMP snooping results in snooping entries being populated in the mFIB of PE-20 (notably on SAP 1/1/2:800), which means that as soon as SAP 1/1/2:800 transitions to the forwarding state, the multicast stream is forwarded immediately to PE-30's multicast receiver without the need for an IGMP query from the IGMP querier and without an IGMP report for the group from PE-30.



## Conclusion

Enabling MCS for IGMP snooping speeds up multicast convergence in the event of Layer 2 failures.

Conclusion