

Cisco IOS "ip igmp join-group" and "ip igmp static-group" Command Use



Document ID: 119383

Contributed by Luc De Ghein, Cisco TAC Engineer.
Dec 02, 2015

Contents

Introduction

Statically Join the IGMP Group

Receiver is Active

IGMP Join Command

IGMP Static Command

PIM DR Role

Safe Use of the *ip igmp join-group* Command

Important Notes About Command Use

Introduction

This document describes how the **ip igmp join-group** and **ip igmp static-group** commands function within the Cisco IOS®.

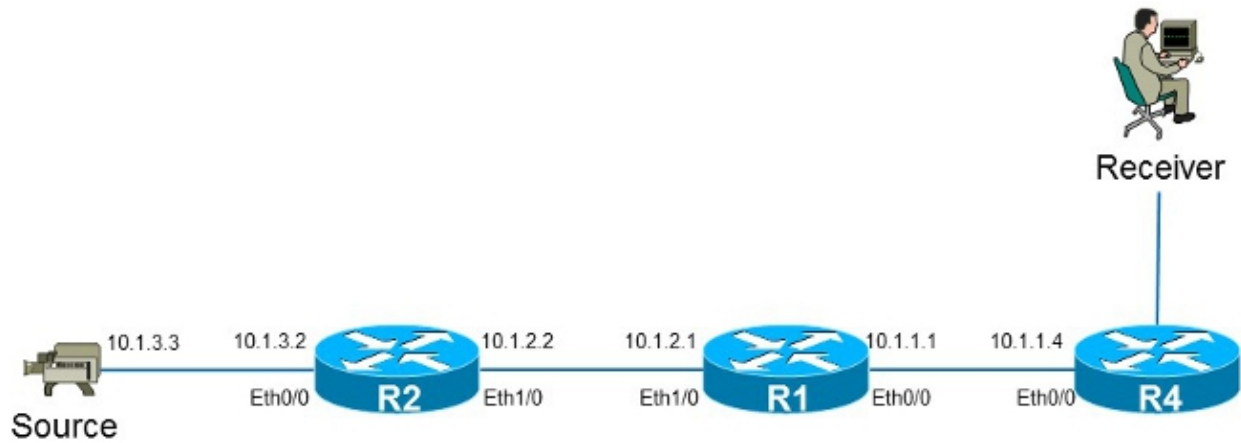
Statically Join the IGMP Group

If the router has the **ip igmp join-group** command on any of the interfaces, the router itself becomes a receiver for the multicast stream. This command is used in order to move multicast traffic to this router without a real directly-connected receiver or without a Protocol Independent Multicast (PIM) neighbor downstream that sends PIM Join requests for the multicast flow. However, because this router joins the multicast stream, all of the multicast packets are punted to the CPU. This can cause high CPU, or it can cause the rate-limiters (if any) or the Control Plane Protection (CoPP) to be hit.

A better alternative that you can use in order to attract the multicast stream for this router is to configure the **ip igmp static-group** interface command. With this command, the router can still attract the multicast stream and forward it out on the interface, but the router itself does not become a receiver for the stream.

Both the **ip igmp join-group** interface command and the **ip igmp static-group** command cause the PIM to send Join requests upstream towards the source or towards the Rendezvous Point (RP), but this only occurs if the router with this command is the PIM Designated Router (DR) on that interface. In order to ensure that the command takes effect and attracts the multicast traffic, use the command on the router that is the DR for that particular network. Alternatively, you can make the router that uses the command the PIM DR. In order to do this, configure the **ip pim dr-priority** command on the interface and ensure that it has the highest PIM DR priority value of any PIM router on that network.

Here is an example:



In this example, there is a source with IP address **10.1.3.3** and a receiver for the group **232.1.1.1**.

Receiver is Active

Here is the multicast forwarding entry on the router R1:

```
R1#show ip mroute 232.1.1.1 10.1.3.3
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.3.3, 232.1.1.1), 01:54:48/00:02:54, flags: sT
Incoming interface: Ethernet1/0, RPF nbr 10.1.2.2
Outgoing interface list:
    Ethernet0/0, Forward/Sparse-Dense, 01:54:48/00:02:54
```

As shown in the output, the interface **Ethernet0/0** is in the Outgoing Interface List (OIL), and the **(10.1.3.3, 232.1.1.1)** multicast traffic is forwarded to the interface Ethernet0/0.

This can also be observed in the Multicast Forwarding Information Base (MFIB) entry:

```
R1#show ip mfib 232.1.1.1 10.1.3.3
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
             ET - Data Rate Exceeds Threshold, K - Keepalive
             DDE - Data Driven Event, HW - Hardware Installed
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
               NS - Negate Signalling, SP - Signal Present,
               A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
               MA - MFIB Accept
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:  FS Pkt Count/PS Pkt Count
Default
(10.1.3.3,232.1.1.1) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
Ethernet1/0 Flags: A
Ethernet0/0 Flags: F NS
Pkts: 0/0
```

IGMP Join Command

If the router R1 does not receive a PIM Join request for the multicast stream from the router R4 (for any reason), then the multicast stream does not flow. One possible reason is that the PIM is not allowed to form a neighborhood between the routers R1 and R4 because the routers belong to a different administrative domain. A solution is to forward the traffic from the router R1 towards the router R4 in a static fashion.

The **ip igmp join-group** command is used on interface Ethernet0/0 on the router R1. This allows the router R1 to send a PIM Join request upstream (to the source or RP) and attract the multicast stream (10.1.3.3, 232.1.1.1). This traffic is then forwarded to the interface Ethernet0/0, as this interface is in the OIL. However, the traffic is also punted to the CPU.

```
R1#show running-config interface Ethernet 0/0
!
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip pim sparse-dense-mode
  ip igmp join-group 232.1.1.1 source 10.1.3.3
end
```

```
R1#show ip mroute 232.1.1.1 10.1.3.3
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
  L - Local, P - Pruned, R - RP-bit set, F - Register flag,
  T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
  X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
  U - URD, I - Received Source Specific Host Report,
  Z - Multicast Tunnel, z - MDT-data group sender,
  Y - Joined MDT-data group, y - Sending to MDT-data group,
  G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
  Q - Received BGP S-A Route, q - Sent BGP S-A Route,
  V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.3.3, 232.1.1.1), 00:09:30/00:02:19, flags: sLTI
Incoming interface: Ethernet1/0, RPF nbr 10.1.2.2
Outgoing interface list:
  Ethernet0/0, Forward/Sparse-Dense, 00:00:40/00:02:19
```

The **L** flag means that the multicast traffic is punted. The interface Ethernet0/0 is in the OIL, so the traffic is punted to the CPU and forwarded to the interface Ethernet0/0.

The MFIB entry shows the *Internal Copy* (IC) flag. This means that the packets for this flow are punted to the CPU.

```
R1#show ip mfib 232.1.1.1 10.1.3.3
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
  ET - Data Rate Exceeds Threshold, K - Keepalive
  DDE - Data Driven Event, HW - Hardware Installed
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
  NS - Negate Signalling, SP - Signal Present,
  A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
  MA - MFIB Accept
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts: Total/RPF failed/Other drops
```

```

I/O Item Counts:   FS Pkt Count/PS Pkt Count
Default
(10.1.3.3,232.1.1.1) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  Ethernet1/0 Flags: A
  Ethernet0/0 Flags: F IC NS
  Pkts: 0/0

```

Because all of the traffic for this multicast stream is punted, it can cause unwanted side effects, as previously described.

Warning: Do not use the `ip igmp join-group` command unless it is used in a test environment.

IGMP Static Command

The `ip igmp static-group` command is used as a solution in order to forward the traffic from the router R1 towards the router R4 in a static fashion. In this scenario, the router R1 sends a PIM Join request upstream (to the source or RP) and attracts the multicast stream (10.1.3.3, 232.1.1.1). This traffic is then forwarded to the interface Ethernet0/0, as this interface is in the OIL, but the traffic is *not* punted to the CPU.

```

R1#show running-config interface Ethernet 0/0
!
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip pim sparse-dense-mode
  ip igmp static-group 232.1.1.1 source 10.1.3.3
end

```

```

R1#show ip mroute 232.1.1.1 10.1.3.3
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
  L - Local, P - Pruned, R - RP-bit set, F - Register flag,
  T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
  X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
  U - URD, I - Received Source Specific Host Report,
  Z - Multicast Tunnel, z - MDT-data group sender,
  Y - Joined MDT-data group, y - Sending to MDT-data group,
  G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
  Q - Received BGP S-A Route, q - Sent BGP S-A Route,
  V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.3.3, 232.1.1.1), 00:07:41/stopped, flags: sTI
Incoming interface: Ethernet1/0, RPF nbr 10.1.2.2
Outgoing interface list:
  Ethernet0/0, Forward/Sparse-Dense, 00:05:06/00:00:53

```

The *L* flag no longer appears. The traffic is not punted on this router, but it is forwarded to the interfaces in the OIL.

Similarly, the MFB entry does not show the *IC* flag:

```

R1#show ip mfib 232.1.1.1 10.1.3.3
Entry Flags:   C - Directly Connected, S - Signal, IA - Inherit A flag,
  ET - Data Rate Exceeds Threshold, K - Keepalive
  DDE - Data Driven Event, HW - Hardware Installed
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
  NS - Negate Signalling, SP - Signal Present,
  A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,

```

```

MA - MFIB Accept
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:  FS Pkt Count/PS Pkt Count
Default
(10.1.3.3,232.1.1.1) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  Ethernet1/0 Flags: A
  Ethernet0/0 Flags: F NS
  Pkts: 0/0

```

PIM DR Role

Neither the **ip igmp static-group** command nor the **ip igmp join-group** command takes effect if the router R1 is not the PIM DR for the interface Ethernet0/0.

Here is an example:

```

R1#show running-config interface Ethernet 0/0
!
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip pim sparse-dense-mode
ip igmp static-group 232.1.1.1 source 10.1.3.3
end

```

Note: The command allows a source to be specified (PIM SSM), or no source to be specified (PIM Sparse Mode/PIM BiDIR mode).

```

R1#show ip mroute 232.1.1.1 10.1.3.3
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.3.3, 232.1.1.1), 00:00:30/00:02:29, flags: sPT
Incoming interface: Ethernet1/0, RPF nbr 10.1.2.2
Outgoing interface list: Null

```

The interface Ethernet0/0 is not in the OIL. This is because the router R1 is not the PIM DR on the link with the **ip igmp static-group** command:

```

R1#show ip pim interface ethernet 0/0

Address          Interface          Ver/  Nbr  Query DR    DR
                Interface          Mode  Count Intvl Prior
10.1.1.1         Ethernet0/0       v2/SD 1    30    1    10.1.1.4

```

The router R1 also does not send a PIM Join request upstream. This is evident on the router R2, as the multicast entry is missing:

```
R2#show ip mroute 232.1.1.1 10.1.3.3
Group 232.1.1.1 not found
```

Here is the output that can be observed as soon as the router R1 is the PIM DR on the interface Ethernet0/0:

```
R1#show ip pim interface ethernet 0/0
```

Address	Interface	Ver/ Mode	Nbr Count	Query Intvl	DR Prior	DR
10.1.1.1	Ethernet0/0	v2/SD	1	30	1	10.1.1.1

```
R1#show ip mroute 232.1.1.1 10.1.3.3
```

IP Multicast Routing Table

Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector

Outgoing interface flags: H - Hardware switched, A - Assert winner

Timers: Uptime/Expires

Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.3.3, 232.1.1.1), 00:02:39/00:02:55, flags: sTI

Incoming interface: Ethernet1/0, RPF nbr 10.1.2.2


Outgoing interface list:

Ethernet0/0, Forward/Sparse-Dense, 00:00:04/00:02:55

Safe Use of the *ip igmp join-group* Command

In order to troubleshoot issues, you might desire to perform a test with multicast, even outside of the lab. In such a case, ensure that you use the **ip igmp join-group** command in a safe manner. The reason that you should use the **ip igmp join-group** command over the **ip igmp static-group** command is because the multicast packets are punted. As such, if you perform a ping with a multicast destination, the router with the command is a receiver for the multicast flow and can reply to the ping.

Here is an example:



The source **10.1.3.3** is an IP address on the router R3. If you place the command on the Ethernet0/0 interface on the router R1 and ping from the router R3, then the router R1 can reply to the ping. As such, you can perform tests as if there was a directly-connected receiver on the router R1. The **ip igmp join-group** command is placed on the Ethernet0/0 interface on the router R1, and the source is specified in order to ensure

that the router R1 only punts traffic from that source (and responds to it).

```
R1#show running-config interface Ethernet 0/0
!  
interface Ethernet0/0  
ip address 10.1.1.1 255.255.255.0  
ip pim sparse-dense-mode  
ip igmp join-group 232.1.1.1 source 10.1.3.3  
end
```

```
R3#ping 232.1.1.1 source 10.1.3.3  
Type escape sequence to abort.  
Sending 1, 100-byte ICMP Echos to 232.1.1.1, timeout is 2 seconds:  
Packet sent with a source address of 10.1.3.3
```

```
Reply to request 0 from 10.1.1.1, 2 ms  
R3#
```

The **debug ip icmp** command on the router R1 indicates that the ping arrived and that the router R1 sends a reply:

```
R1#debug ip icmp  
ICMP packet debugging is on  
R1#
```

```
*Oct 30 11:35:41.133: ICMP: echo reply sent, src 10.1.1.1, dst 10.1.3.3,  
topology BASE, dscp 0 topoid 0
```

Important Notes About Command Use

The best practice is not to use the **ip igmp join-group** command unless it is for test purposes in the lab or a temporary test on a live network. Remove the command after all tests are complete. If the multicast traffic must be forwarded only statically, use the **ip igmp static-group** command instead.