

Rebuilding the Multicast Entries with CGMP and Spanning Tree Topology Changes

Document ID: 24100

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

CGMP and Topology Changes

- Stable State
- During and After the Topology Change
- Two IGMP General Queries After the Topology Change Notification

CGMP Enhancements

- Communication Between the Switch and the Router
- Router Behavior
- Catalyst Switch Behavior

Related Information

Introduction

This document discusses how Cisco Group Management Protocol (CGMP) works on the Cisco Catalyst switches and the Cisco IOS® routers with regard to the rebuilding of the multicast entries for CGMP after a spanning tree topology change has occurred.

Prerequisites

Requirements

Cisco recommends you have knowledge of these topics:

- basic operation of switches, routers, and multicasting
- basic operation of the spanning tree, CGMP, and Internet Group Management Protocol (IGMP)

Components Used

The information in this document is based on these software and hardware versions:

- Catalyst 3550 version 12.1(9)EA1c
- Catalyst 2900/3500XL version 12.0(5)WC3b
- Catalyst 4000 Supervisor Engine III version 12.1(11b)EW
- Catalyst 4000 Supervisor Engine I/II version 7.2(2)
- Catalyst 6500 Supervisor Engine Cisco IOS Software Release 12.1(11b)EX
- Catalyst 6500 Catalyst OS (CatOS) version 7.2(2)
- Catalyst 5500 CatOS version 4.5(13a)


The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

CGMP and Topology Changes

This section describes step-by-step what occurs and what problems can arise when a spanning tree topology change is detected on a VLAN where CGMP is used in order to restrain multicast traffic from flooding on all ports. As this example shows, the network discussed in this document consists of one router, one switch, and four PCs:



- port 1 receiver PC 1
- port 2 receiver PC 2
- port 3 receiver PC 3
- port 4 not a receiver PC 4
- port 5 other switch (no receivers or routers on this switch)
- port 48 Cisco IOS router running IGMP and CGMP

For the purpose of this document, it is assumed that the receiver PCs use IGMP and the switch runs CGMP. The Cisco IOS router runs IGMP and CGMP, which receives a multicast stream from a video server on a different interface. This interface sends to IP multicast group 239.100.100.100.

Stable State

Once all the devices are booted and the receiver PCs have sent their IGMP join messages for group 239.100.100.100, they are all added by CGMP to the corresponding Layer 2 group represented by MAC address 01-00-5e-64-64-64.

This list shows which ports, highlighted in bold, on the switch receive the multicast stream that come through the Cisco IOS router.

- **port 1 receiver PC 1**
- **port 2 receiver PC 2**
- **port 3 receiver PC 3**
- port 4 not a receiver PC 4
- port 5 other switch (no receivers or routers on this switch)

- port 48 Cisco IOS router that runs IGMP and CGMP

Note: The Cisco IOS router is also added to the multicast group, but since it is the source, it does not receive its own packets.

At every query interval, the Cisco IOS router sends out an IGMP general query (which is sent to multicast group 224.0.0.1, and therefore flooded to all the other components). When this happens, all the receivers start to build an IGMP report for the 239.100.100.100 group. The receivers send this report back to IP multicast group 239.100.100.100, with a Layer 2 MAC address of 01-00-5E-64-64-64. Since this is sent to the group address, all receivers receive the reports that are sent by other receivers as well as the report sent back by the first receiver. This triggers the other receiver PCs to cancel their report for this group. This means that only one CGMP join message is sent for this group with the source MAC address of the PC that was the first to respond. This continues for a long period of time, and all the receiver PCs receive the video broadcast.

During and After the Topology Change

At this point, the other switch triggers a topology change in the network. As per the CGMP specification upon receiving the topology change, the switch clears all the multicast entries that it had learned through the CGMP. The multicast traffic from the router is flooded to all ports on the switch.

This list shows which ports, highlighted in bold, on the switch receive the multicast stream that come through the Cisco IOS router:

- **port 1 receiver PC 1**
- **port 2 receiver PC 2**
- **port 3 receiver PC 3**
- **port 4 not a receiver PC 4**
- **port 5 other switch (no receivers or routers on this switch)**
- port 48 Cisco IOS router that runs IGMP and CGMP

As traffic is flooded to all ports, the receiver PCs do not notice any difference, and they continue to receive the video broadcast. However, since traffic is flooded to all ports, PC 4, which is not a receiver, and the other switch now also receive the multicast stream, though they have not requested it. This continues until the Cisco IOS router sends out its periodic IGMP general query again. The default value for this is 60 seconds on Cisco IOS routers (configured with an IP IGMP query interval).

Two IGMP General Queries After the Topology Change Notification

When the Cisco IOS router sends out its first IGMP general query, all receiver PCs start to build their IGMP report for the 239.100.100.100 group. One of them (in this document, it is PC 3) is the first to send back its IGMP report. Since there is no multicast entry built on the switch yet, it is received by all PCs, and the other receiver PCs cancel their IGMP report. The Cisco IOS router receives the report and sends out the subsequent CGMP join message with the source address of receiver PC 3.

The switch builds a multicast entry again for group 01-00-5e-64-64-64 and adds port 3 to it, as this is the source address in the CGMP join packet. Since port 5 is the multicast router port, this is also added to the multicast group. Therefore, only receiver PC 3 receives the video stream, while the video stream on PC 1 and PC 2 stands still.

This list shows which ports, highlighted in bold, on the switch receive the multicast stream that comes through the Cisco IOS router:

- **port 1 receiver PC 1**

- port 2 receiver PC 2
- **port 3 receiver PC 3**
- port 4 not a receiver PC 4
- port 5 other switch (no receivers or routers on this switch)
- port 48 Cisco IOS router running IGMP and CGMP

At the end of an IGMP querying interval, the Cisco IOS router sends out another IGMP general query. Upon receiving the query, all the receiver PCs build a report for the 239.100.100.100 group. This time, however, the reports from the other PCs are only received by receiver PC 3 and the Cisco IOS router. (The router port is automatically added to every multicast group.)

Since receivers PC 1 and PC 2 do not see a report from any other receiver, they both send out their reports. The Cisco IOS router subsequently sends out a CGMP join message with the source MAC address of the respective PCs, and therefore, they are added and start receiving the multicast stream again through the Cisco IOS router.

This list shows which ports, highlighted in bold, on the switch receive the multicast stream that comes through the Cisco IOS router:

- **port 1 receiver PC 1**
- **port 2 receiver PC 2**
- **port 3 receiver PC 3**
- port 4 not a receiver PC 4
- port 5 other switch (no receivers or routers on this switch)
- port 48 Cisco IOS router running IGMP and CGMP

The configuration is back to the original stable state and everything works properly again. This is a breakdown of what has occurred:

1. A topology change occurs.

Tip: When portfast is not enabled in a host port, every time a host is rebooted, or connected/disconnected to/from the port, a change in the links status triggers a topology change notification in the VLAN. If the debugging of CGMP is enabled at the time of the topology change, this debug message is displayed:

```
CGMP SHIM: got short age timer
```

2. Flooding starts to all ports.
3. The first IGMP general query is sent out.
4. Flooding stops.
5. Not all receivers receive the multicast stream.
6. The second IGMP general query is sent out.
7. All receivers are added and receive the multicast stream again.

CGMP Enhancements

Since having a one-minute (the default IGMP querying interval) loss of a multicast stream for a PC is not always acceptable, there have been some enhancements made for both the routers and switches that run CGMP.

Communication Between the Switch and the Router

Since routers are Layer 3 devices and therefore do not generally know about spanning tree and topology

changes that occur, there is a need for the switches in the network to alert the router of this topology change. An IGMP global leave message is defined in order to handle this.

This IGMP global leave message is an IGMP leave that a switch can transmit, requesting to leave the group 0.0.0.0.

In order to ensure that the router is not overloaded with IGMP global leave messages, only the root switch in a spanning tree domain is responsible for sending this IGMP global leave message when the topology change is over.

Router Behavior

When the router receives this IGMP global leave message on an interface that runs Cisco IOS Software, it recognizes that a spanning tree topology change has occurred on that interface and takes these actions to try and limit the loss of multicast traffic for the multicast receivers:

1. Sends CGMP batch join messages after receiving the IGMP global leave message. The router sends out a CGMP join message with its own MAC address as the user source address for every multicast group it has in its IGMP cache for that interface. By sending these CGMP self-join messages, the CGMP switches automatically creates an entry for each group with only the router port in it.

This list shows the network used in this document, after the CGMP batch join. Only the Cisco IOS router has been added to multicast group, as shown in bold.

Note: While in previous examples in this document, the ports that receive traffic from the multicast router was shown in bold, this example shows all ports that are added on the switch to the multicast group.

- ◆ port 1 receiver PC 1
 - ◆ port 2 receiver PC 2
 - ◆ port 3 receiver PC 3
 - ◆ port 4 not a receiver PC 4
 - ◆ port 5 other switch (no receivers or routers on this switch)
 - ◆ **port 48 Cisco IOS router running IGMP and CGMP**
2. Sends out an IGMP general query. All receivers receive this IGMP general query, and build a report for every group that they have joined. Since the CGMP switch has already built a multicast entry for each of the groups with only the router as the receiver, all the reports are being sent to only the router. The router sends out subsequent CGMP join messages for adding all the receivers to the corresponding groups.

After all the receivers have sent back their IGMP report and the router has sent out the corresponding CGMP join messages, all the receivers should have been added back to the multicast group.

3. After 10 seconds (default IGMP max-response-time), another IGMP general query is sent out to make sure that all the receivers are added. This step is repeated a few times to make sure that all the receivers rejoin the multicast group.

All the ports that should have been added to the multicast group have been, as shown in bold in this example:

- ◆ **port 1 receiver PC 1**
- ◆ **port 2 receiver PC 2**
- ◆ **port 3 receiver PC 3**
- ◆ port 4 not a receiver PC 4
- ◆ port 5 other switch (no receivers or routers on this switch)

Catalyst Switch Behavior

Within the range of Catalyst switches, there are some differences in their behavior. Every switch that is CGMP-capable does as described in the CGMP and Topology Changes section of this document. The enhancements for CGMP, however, are not implemented on all platforms. This table provides a list of Catalyst switches and how they react to CGMP:

	CGMP Switch	CGMP Router	Sends Global Leave When Spanning Tree Protocol (STP) Root
Catalyst 6500 running Cisco IOS Software	N	Y	Y
Catalyst 6500 running CatOS	N	N	N
Catalyst 5500, Catalyst 2926/2926G	Y	N	Y
Catalyst 4000 Supervisor Engine I/II, Catalyst 2948G/2980G, Catalyst 4912G	Y	N	Y
Catalyst 4000/4500 Supervisor Engine III/IV	N	Y	Y
Catalyst 2900XL/3500XL	Y	N	Y
Catalyst 2940	N	N	N
Catalyst 2950	N	N	N
Catalyst 2970	N	N	N
Catalyst 3550	N	Y	Y
Catalyst 3750	N	Y	Y

Note: On the Catalyst 4000/4500 with a Supervisor Engine III/IV, the behavior with regard to topology changes and CGMP is configurable. Issue this command in order to configure the Catalyst 4000 to send or not send an IGMP global leave message when it is not the spanning tree root:

- **ip igmp snooping tcn query solicit**

Note: Issue this "no" form of the command in order to disable it:

- **no ip igmp snooping tcn query solicit**
-

Related Information

- [Understanding Spanning-Tree Protocol Topology Changes](#)
 - [Multicast in a Campus Network: CGMP and IGMP Snooping](#)
 - [Constraining Multicast Traffic with Source and Receivers on the Same VLAN on Catalyst Switches Running Catalyst OS](#)
 - [Configuration Guide for the Catalyst 4000 Cisco IOS Software: Understanding and Configuring IGMP Snooping](#)
 - [Spanning Tree Technical Support Page](#)
 - [LAN Product Support Pages](#)
 - [LAN Switching Support Page](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jul 20, 2006

Document ID: 24100
