Multicast Notes

Esthetys Ltd 31 Prospect road Surbiton, Surrey KT6 5PY, United Kingdom fabien.sage@eyscube.com www.eyescube.com

Overview

Definition

Multicast traffic is the traffic generated from n source(s) sent to a multicast group IP.

The group IP can be allocated from the IPv4 class D range, 224.0.0.0 to 239.255.255.255 or from the IPv6 range ....

The multicast IPv4 addresses map to the Ethernet addresses 01-00-5e-00-00-00 to 01-00-5e-7f-ff-ff. This range is only composed of 23 bits instead of the 28 bits used to generate the layer 3 IP address. As a result from one multicast Ethernet address you can generate 32 different multicast IP address (28-23=5).

The multicast IPv6 addresses is a one to one map to the Ethernet addresses so there is no overlap issues.

IGMP

Internet Group Management Protocol is used by the multicast client application to inform the local router(s) of the Multicast Group the client wants to be part of.

The local routers use IGMP to track multicast membership of the multicast application clients directly connected to the router.

There is currently 3 version of this protocol.

IGMP v1

In this IGMP version there are only 2 messages. IGMP membership query and IGMP membership report. These packets are generated with a TTL value set to one as there are destined to the local subnet only.

Hosts can join a multicast group IP (membership report) but they can not explicitly leave a group they are a member of. For a router to know that hosts on the local subnet still want to be part of a multicast group it use membership queries. The router send membership queries to 224.0.0.1 (all devices on the subnet) for the multicast group every 60 seconds. If the router doesn't receive a membership report back for 3 consecutive minutes the router stop forwarding multicast traffic for this group to the local subnet.

To avoid the router to be flooded with membership report back from all the hosts attached to the local subnet, each host on the subnet will start a local timer before it replies. If it sees a membership report back from one of the other host before its timer expire it will not send the membership report.

The router only need to see one membership report to maintain the distribution of the multicast feed.

IGMP v2

This version of IGMP adds a leave message for the hosts who want to leave the group. The router will only send membership query if it has received a leave message from on the host attached to the local subnet. The router will then test if the multicast feed for the group is still required by sending a membership query.

This avoid unnecessary IGMP messages to be generated by the hosts on the network and will more promptly stop unrequired multicast feed to be maintain on the local subnet.

The IGMP leave message is only sent to the local routers as the message is sent to 224.0.0.2.

IGMP v3

This version of IGMP adds source filtering capabilities to the IGMP protocol with 2 new messages. INCLUDE and EXCLUDE, they specify the group IP and the source IP the multicast feed need to be composed of.

Multicast traffic distribution

The multicast traffic is distributed through the network with 2 types of trees. These trees are build on demand based on the IGMP request sent by the devices on the network and multicast traffic sent by the sources.

The first type of tree is called (S,G) or source tree. These trees are build with the source as the root with S as the IP of the source and G the IP of the multicast group.

The second type of tree is called (*,G) or shared tree. These trees are build with a router as the root. This router is called a Rendez-vous Point. The RP forward traffic to receivers on behalf of sources.

With shared tree, routers do not need to maintain (S,G) entries for all the sources of a multicast group, which save some resource.

PIM

Protocol Independent Multicast is the network protocol used by routers to route multicast traffic through the network. PIM base is routing reverse path selection on the unicast routing table generated on the router by any running routing protocol.

PIM would forward the multicast traffic to all the interfaces except the reverse path interface, this is the interface the closest to the source for (s,g) trees and the interface closest to the RP for (*,g) trees. This process is called the Reverse Path Forwarding check, (RPF).

With Cisco there is 3 PIM mode of operation.

PIM Dense Mode

In this mode the multicast traffic is initially forwarded to all the PIM enabled interfaces except the interface that fail the RPF check. The edge routers will then send PIM prune messages back if they don't have any receiver attached to the local subnets. Routers on the way back to the source will then update their outgoing multicast interfaces for this group accordingly. When the process is completed the multicast traffic will only be flooded to the network where it has been requested.

This mode is proved inefficient as this flood process is done every 3 minutes and if you do not have receivers on most of your subnets your are flooding your network with unnecessary traffic every 3 minutes.

In this mode only (s,g) trees are build, this also can use a lot of resource if you have a lot of multicast sources on your network.

PIM Spare Mode

In this mode the multicast traffic is only forwarded to the location where it has been requested by the applications through an IGMP join request.

From the receiver.

The IGMP join request trigger the edge router to send a PIM join message towards the RP via the RPF interface using the destination multicast IP 224.0.0.13 with a TTL value set to 1. The IGMP join request is also used by the routers on the client subnet to build the (*,g) tree from the RP to the multicast client.

From the sender.

The sender just start to send multicast traffic. The DR on the local subnet will register the multicast group IP to the RP on behalf of the sender. To do this the DR encapsulate the UDP multicast packet into a PIM register unicast packet and send it to the RP. This packet is also used to build the source tree (s,g) from the source to the RP.