Basic Data Switch Setup (Cisco and Brocade)

Ericsson Compression Systems Manual Collection (Internal)

Copyright

Disclaimer

No part of this document may be reproduced in any form without the written permission of the copyright owner.

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Contents

List of Figures

List of Tables

1 Read This First!

Read and understand the following iSIS Manual Collection documents before proceeding with this document:

- Cisco Switch Fundamentals
- Brocade Switch Fundamentals

Many of the commands required in this document are detailed in the above two documents. The fundamentals documents also contain prerequisite setup detail.

2 What the Document Describes

This document is intended to provide information regarding set up of a basic data network carrying multicast traffic from input sources (receivers or encoders, for example) to the MX8400 and / or multicast traffic from an MX8400 to another device or customer network.

There is details of both Cisco and Brocade setup methods with example configurations for both provided towards the end of the document

3 Overview

A typical, basic, Ericsson Television headend includes a 1 + 1 data network supplying the inputs of the MX8400 with the outputs of Ericsson encoders and/or receivers in addition to other (possibly 3rd party) dual output sources. Many systems also include a 1 + 1 pair of output switches. Sometimes the same physical pair of switches is used in both roles. It is important that the configuration of these switches is correct and consistent from system to system.

A basic data network is assumed where all sources are local to the switches; the switches are mirrored with identical configurations. An example of a system with a basic data network is shown in *Figure 3.1*.

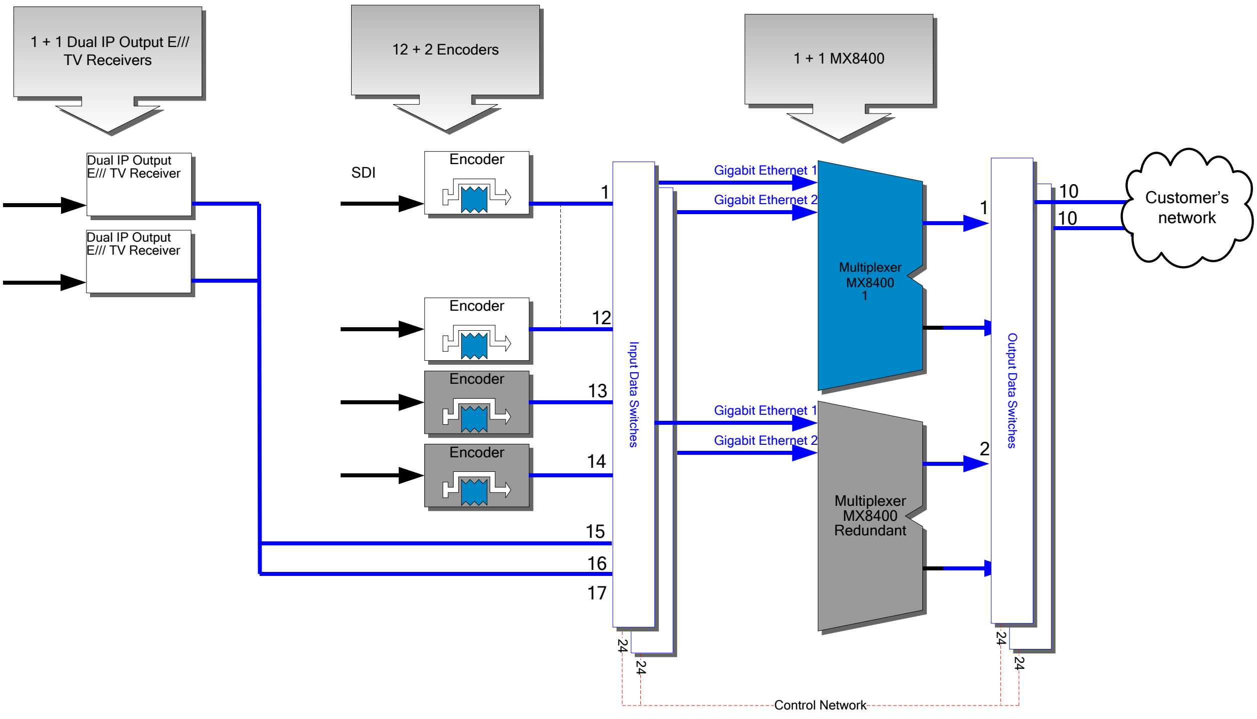

Figure 3.1 Typical IP system showing the use of data switch.

4 MX8400 Input and Output Ports

Refer to the MX8400 Static Parameters document in the iSIS Manual Collection for details on the MX8400 data port setup and behaviour.

5 VPC Setup

Refer to the VPC Setup and Optimization document in the iSIS Manual Collection for details on the VPC setup requirements.

6 Legacy Encoder Setup

6.1 E57xx, EN80xx

6.1.1 Transport Stream Output Bitrate

Ericsson Television encoders (EN80xx, E57xx) support the installation of a dual output Gigabit Ethernet Network Interface Card (GigE NIC). When configured, these cards will output an IP encapsulated transport stream. This stream can then be setup to unicast or multicast out of one or both GigE NIC outputs. However, iSIS systems only support multicast. The ASI outputs will remain unaffected and will continue to output a transport stream over ASI.

The diagram below shows where the total output rate of the Encoder is configured.

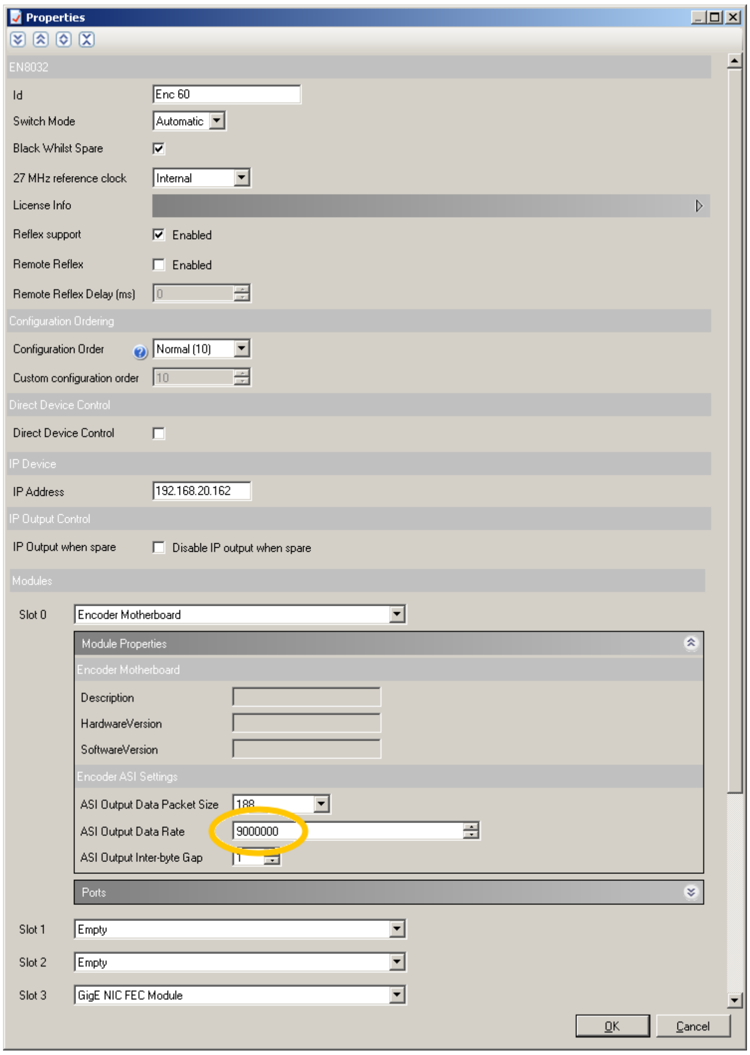

Id: Enc 60
Switch Mode: Automatic
Black Whilst Spare: ☑
27 MHz reference clock: Internal
License Info
Reflex support: ☑ Enabled
Remote Reflex: ☐ Enabled
Remote Reflex Delay (ms): 0
Configuration Order: Normal (10)
Custom configuration order: 10
Direct Device Control: ☐
IP Address: 192.168.20.162
IP Output when spare: ☐ Disable IP output when spare
Slot 0: Encoder Motherboard
Description:
HardwareVersion:
SoftwareVersion:
ASI Output Data Packet Size: 188
ASI Output Data Rate: 9000000
ASI Output Inter-byte Gap: 1
Slot 1: Empty
Slot 2: Empty
Slot 3: GigE NIC FEC Module

Figure 6.1 Encoder output transport stream rate set within Equipment Setup.

NOTE The ASI Output Data Rate also sets the IP output rate

It is essential that the output transport rate of the Encoder is correctly set, allowing for the sum of all the components that the Encoder could be configured to carry plus an overhead.

Having a large Encoder transport rate will cause problems when Encoders are used in IP systems as the sum of all the active Encoder transport streams must not exceed the bandwidth of the link connecting the MX8400 to the network.

By default, 50Mbit/s is configured. This is clearly excessive in a system where the sum of the components from the Encoder will never exceed 10Mbit/s.

Correctly setting the output rate can also have an impact on the Encoder’s PCR output accuracy. Though the PCR jitter from an ERICSSON Encoder is well within the specifications, regardless of the output bitrate set up, it is possible for the PCR jitter to be perfect.

The rate should be set to integer devisors of 27 MHz. For example, 27 Mbit/s, 13.5 Mbit/s, 9 Mbit/s, 6.75 Mbit/s or 4.5 Mbit/s. The illustration below shows the effect correct selection will have.

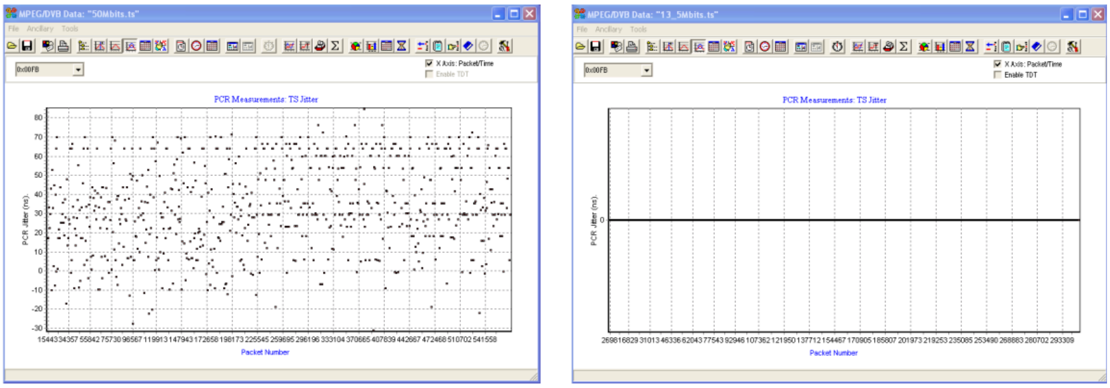

MPEG/DVB Data: "50Mbits.ts"

0x00FB

☑ X Axis: Packet/Time

□ Enable TDT

MPEG/DVB Data: "13_5Mbits.ts"

0x00FB

☑ X Axis: Packet/Time

□ Enable TDT

Figure 6.2: Encoder PCR output jitter. TS rate set to 50Mbit/s on the left. 13.5 Mbit/s on the right.

6.1.2 GigE NIC Module Settings via nCC

The diagram below shows the Gigabit Ethernet module properties which require editing.

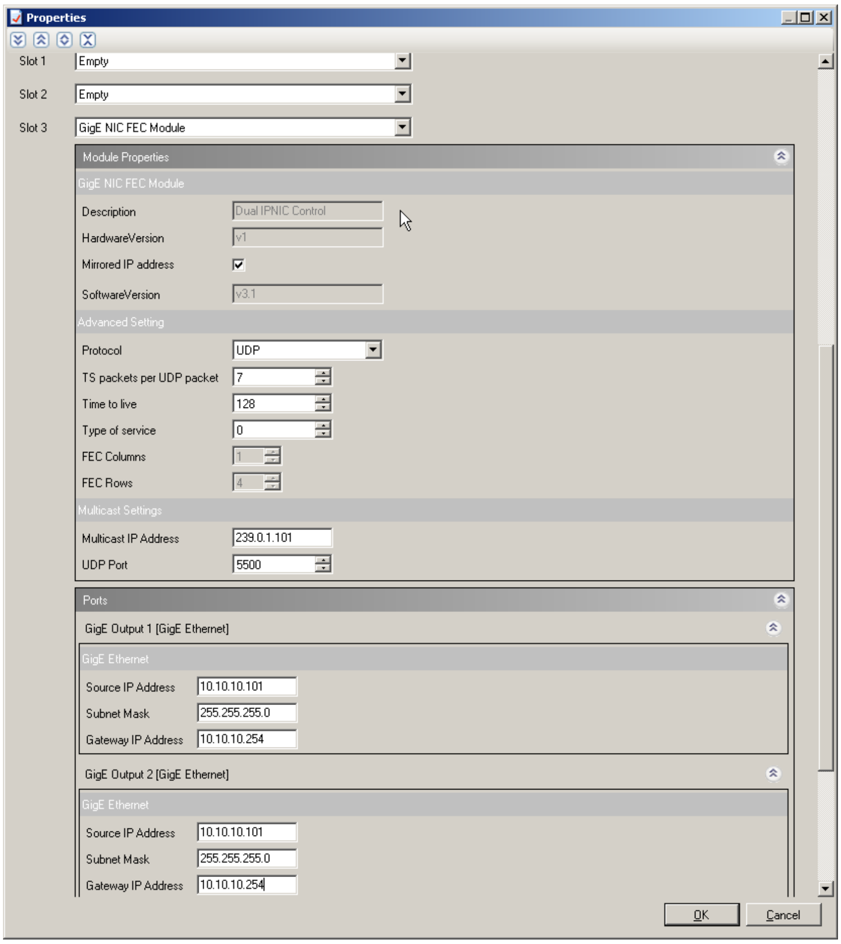

Slot 1: Empty
Slot 2: Empty
Slot 3: GigE NIC FEC Module

GigE NIC FEC Module
Description: Dual IPNIC Control
HardwareVersion: v1
Mirrored IP address: ☑
SoftwareVersion: v3.1

Advanced Setting
Protocol: UDP
TS packets per UDP packet: 7
Time to live: 128
Type of service: 0
FEC Columns: 1
FEC Rows: 4

Multicast Settings
Multicast IP Address: 239.0.1.101
UDP Port: 5500

GigE Output 1 [GigE Ethernet]
GigE Ethernet
Source IP Address: 10.10.10.101
Subnet Mask: 255.255.255.0
Gateway IP Address: 10.10.10.254

GigE Output 2 [GigE Ethernet]
GigE Ethernet
Source IP Address: 10.10.10.101
Subnet Mask: 255.255.255.0
Gateway IP Address: 10.10.10.254

Figure 6.3 Encoder GigE Module Properties via Equipment Setup.

Care should be taken so that the Multicast address, Gigabit Ethernet NIC address and the Encoder control IP address are related. The following table describes a numbering convention. This way, it becomes easy to deduce which device requires attention if an alarm is raised because of a multicast address, or IP address, or Encoder NIC.

The following table shows example IP addresses.

Table 6.1: Example IP addresses. It is important that all addresses are related.

	Encoder 1	Encoder 2	Encoder 3	…	Encoder n
Unit IP address	192.168.1.101	192.168.1.102	192.168.1.103	…	192.168.1.(100+n)
GigE NIC IP address	10.10.10.101	10.10.10.102	10.10.10.103	…	10.10.10. (100+n)
Multicast IP address	239.255.0.101	239.255.0.102	239.255.0.103	..	239.255.0.(100+n)

Ensure that **Mirrored IP address** is checked. This will cause both outputs of the GigE NIC to share the same multicast source and destination addresses. When **Mirrored IP address** is checked, identical streams should be present on both outputs at all times.

6.1.3 GigE NIC Module Speed and Duplex Settings

It is strongly recommended to fix the speed and duplex of the GigE NIC to 100Mbit/s Full duplex. The port on the data switch will also require manual speed and duplex configuration. This avoids the potential for sub-optimal settings being selected as a result of auto-negotiation.

This can be achieved via the Telnet or web browser under the Dual IP NIC control menu. When fixing the duplex on the GigE NIC then it must also be fixed on the corresponding switch port.

7 VLAN Managed Switches

7.1 Why VLAN Management?

VLAN Management is necessary for certain types of redundancy mechanisms. It is a requirement when the output of the spare device mirrors the output of the main device, when it takes over as the result of a redundancy switch. In this scenario it is important that the output of the failed device is turned off by nCC. As a backup mechanism in case of loss of communications to a failed device, the ports for ‘inactive’ devices are switched to an unused VLAN in the switch, hence the term VLAN management. This is done to protect against the possibility that the nCC cannot contact the failed device to turn its output off.

7.2 When is VLAN management required?

VLAN management is required in the following scenarios:

- Any iPlex system
- MX8400 outputs in R2C or earlier systems
- Any R2C device profile system involving redundancy

7.3 VLAN Management uses Cisco Switches Only

nCC does not have the ability to control the VLANs of Brocade switches. Any system requiring VLAN management has to use Cisco switches.

7.4 Multicast Guard Protocol (MGP)

Multicast Guard Protocol (MGP) is an Ericsson Television proprietary protocol that removes the need to VLAN manage switches.

It is a protocol that is based on in-band signalling to enable a failed device, in the instance of loss of communications with nCC, to turn off its output after detecting another device which has begun outputting the same multicast and source details.

MGP is used in the SPICA release and the R2D system release.

8 Switch setup

8.1 Setting up a Data VLAN

Care must be taken that multicast (data) traffic does not enter the control network, therefore the data ports are separated into a different VLAN to the control ports.

When new - out of the box, all the ports on the switch are in VLAN 1. This default VLAN will not be explicitly displayed when viewing the configuration; it is therefore easy to forget that the port resides in a VLAN. It is helpful to put the required ports into a known, manually specified VLAN that will be displayed when viewing the switch configuration.

Ensure the Data VLAN has an IP address on the same subnet as the data ports of the equipment.

For details on how to create VLAN, including adding ports and an IP address, refer to the relevant Cisco or Brocade Switch Fundamentals document in the iSIS Manual Collection. The Data VLAN will require an IP address.

If the same switches are to be used for both the input and outputs of an MX8400, then you need to create two Data VLANs; one for the input to the MX8400 and one for the output.

8.2 Creating an Unused VLAN (VLAN Management Only)

For nCC to be able to put ports into the Unused VLAN during VLAN management actions, the VLAN must exist on the switch.

The VLAN needs to be created and named but there is no need to have an IP address associated with it. For details on how to create VLAN, refer to the relevant Cisco or Brocade Switch Fundamentals document in the iSIS Manual Collection.

8.3 IGMP Snooping

For details of the function of IGMP Snooping please refer to Chapter 12 at the end of this document.

IGMP Snooping must be used, and is enabled by default on both Cisco and Brocade switches. However, for IGMP Snooping to work there must be an IGMP Querier running on the switch. The following section details the how to ensure an IGMP Querier is running.

8.4 IGMP Querier

For details of the function of IGMP please refer to 11.1 at the end of this document.

A basic data network requires an IGMP Snooping Querier to be enabled on all switches. The examples below show how to turn on IGMP Snooping and how to setup a Querier on both Cisco and Brocade switches.

This command is the same for layer 2 (SMI / IP Base) and layer 3 (EMI / IP Services) Cisco Catalyst switches:

```
SwitchName(config)#ip igmp snooping
SwitchName(config)#ip igmp snooping querier
```

This Command is for Brocade FastIron WS (FWS) switches:

```
SwitchName(config)#ip multicast active
```

Details on the function of IGMP can be found in Chapter 11.

8.5 Connecting the Data Switch to the Control Network

8.5.1 Why Connect to the Control Network?

nCC will need to communicate with the switch to enable it to health poll it through SNMP and in the case of a VLAN managed Cisco also control it through SNMP.

8.5.2 Using the Correct Method

There are two ways of configuring a Data Switch to be able to connect it to a control network. The table below is designed to explain which mechanism should be used in various scenarios:

Table 8.1 Control Network Connection Method

Type of Control Network	Cisco	Brocade
Basic	Layer 3 Port	Control VLAN
Redundancy	Control VLAN	Control VLAN

8.5.3 Creating a Layer 3 Port on a Cisco Data Switch

The following commands show how to turn a port into a Layer 3 port and add an address.

```
SwitchName(config)#interface fa0/24
SwitchName(config-if)#no switchport
SwitchName(config-if)#ip address 192.168.1.252 255.255.255.0
```

8.5.4 Adding a Control VLAN on a Data Switch

The layer 3 approach cannot be used when using a Brocade switch or when any switch type is to be part of a redundant control network. In these instances, another VLAN for control must be created in addition to the Data VLAN. If the control network uses Brocade or Cisco switches, then this Control VLAN must be configured on the same VLAN number as the VLAN on the control switches.

For details on how to create a Control VLAN including adding ports and an IP address, refer to the relevant Cisco or Brocade Switch Fundamentals document in the iSIS Manual Collection under the creating a VLAN section. An IP address will be required for the VLAN. This IP address will be the one entered into nCC for the switch.

9 Example Cisco Configuration

The following example shows how the configuration of a Cisco 3560 will look following the successful setup following the instructions contained in this document.

```
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
no service dhcp

hostname Primary_Data_1

enable secret 5 $1$fYKE$tgoxoOAAK/LHYG8PieLAo/

ip subnet-zero

ip igmp snooping querier

no ip domain-lookup
vtp mode transparent


spanning-tree mode rapid-pvst
spanning-tree extend system-id
spanning-tree vlan 11 priority 45056

interface FastEthernet0/1
 description E5770 1 NIC 1
 switchport mode access
 switchport access vlan 10
 spanning-tree portfast

interface FastEthernet0/2
 description E5770 2 NIC 1
 switchport mode access
 switchport access vlan 10
 spanning-tree portfast

*** lots more interfaces, each correctly setup, not shown for
clarity***

interface FastEthernet0/24
 description Connected to Backup_Ctrl_Network
 switchport access vlan 11
 duplex full
 speed 100


interface GigabitEthernet0/1
```

```
 description MX8400 1 Data port 1
 switchport mode access
 switchport access vlan 10
 spanning-tree portfast


interface GigabitEthernet0/2
description MX8400 2 Data port 1
 switchport mode access
 switchport access vlan 10
 spanning-tree portfast

interface Vlan1

interface Vlan10
 description Data_VLAN
 ip address 10.10.10.254 255.255.255.0

interface Vlan11
 description Control_VLAN
 ip address 192.168.1.252 255.255.255.0


ip default-gateway 192.168.1.254
ip classless
no ip http server

no cdp run
snmp-server community private RW
snmp-server community public RO

line con 0
line vty 0 4
 password 7 00071A150754
 login
line vty 5 15
 password 7 045802150C2E
```

10 Example Brocade Configuration

The following example shows how the configuration of a Brocade FWS will look following the successful setup following the instructions contained in this document.

```
ver 04.3.03T7e1

module 1 fws1g-24-port-copper-base-module


vlan 1 name DEFAULT-VLAN by port

vlan 2 name Control_Vlan by port
 untagged ethe 0/1/24
 spanning-tree 802-1w
 router-interface ve 2

vlan 10 name Data_Vlan by port
 untagged ethe 0/1/1 to 0/1/23
 spanning-tree 802-1w
 router-interface ve 10

enable telnet password 8 $1$Ya0..dT.$LnTr/MyDEqN2DaUg12AxH1
enable super-user-password 8 $1$bj2..GQ1$onZYFHEe/QyTG5Alccaw50
hostname DataSwitch_1


ip multicast active
snmp-server community 1 $Si2^=d ro
snmp-server community 1 $SU=r!on rw
interface ethernet 0/1/1
 port-name MX8400 1 Port 1


interface ethernet 0/1/2
 port-name MX8400 2 Port 1
 spanning-tree 802-1w admin-edge-port

interface ethernet 0/1/3
 port-name VPC 1 Port 1
 spanning-tree 802-1w admin-edge-port

interface ethernet 0/1/4
 port-name VPC 2 Port 1
 spanning-tree 802-1w admin-edge-port

*** lots more interfaces, each correctly setup, not shown for
clarity***

interface ethernet 0/1/23
 port-name VPC 21 Port 1
spanning-tree 802-1w admin-edge-port
```

```
interface ethernet 0/1/24
 port-name Connected to Control Network

interface ve 2
 port-name -name Control VLAN IP Address
 ip address 172.23.38.15 255.255.255.224

interface ve 10
 port-name Data VLAN IP Address
 ip address 172.23.38.94 255.255.255.224

end
```

11 Background Information

11.1 IGMP Overview

11.1.1 Introduction

IGMP stands for Internet Group Management Protocol. It has been developed over time as a way of handling multicasts around a network. For the purpose of this description, a network is defined as a single subnet. It is usually run on the edge of a routed network:

The purpose of IGMP is to keep track of which multicasts are requested within a network. It is used along side PIM for routing multicasts through multiple networks.

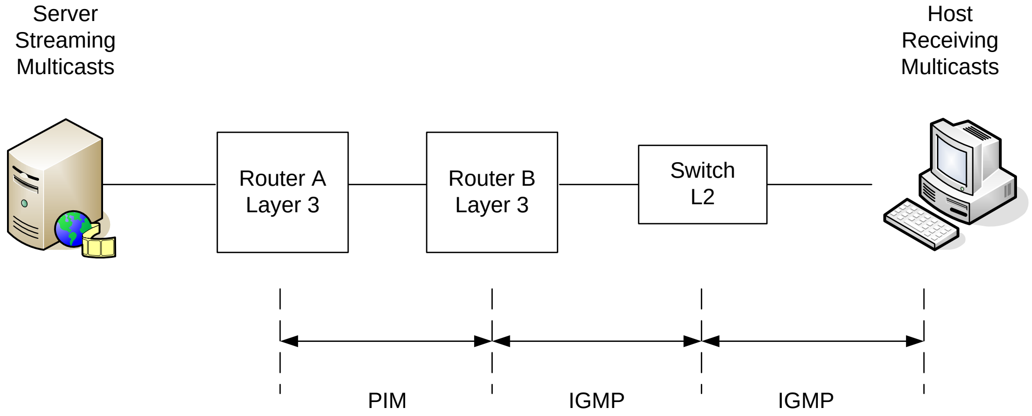

Figure 11.1 IGMP being used on the edge of a layer 3 network.

There are three versions of IGMP: v1, v2 and v3

11.1.2 IGMP v1

IGMP v1 was developed in the early stages of multicasts on networks. It has a simple structure and only two message types:

- IGMP Report
- IGMP Query

These two message types exist in the other versions and will be explained later. It was soon discovered that there were problems with this version and so v2 was developed. Our equipment supports V1 although we do not recommend we run this version on our data networks.

11.1.3 IGMP v2

One of the faults of IGMP v1 was that it had large timers for stopping multicasts when they were no longer needed. To rectify this problem the concept of ‘leave messages’ was added to IGMP v2.

In IGMP v2 the messages that are used are:

- IGMP Report
- IGMP Query
- IGMP Leave

11.1.4 IGMP v3

IGMP v3 was a further rewrite of the specification, with the addition of source specific multicasts and a new way of sending membership reports.

11.1.5 IGMP Messages Overview

There are 5 types of IGMP messages used in the 3 versions of IGMP. These are:

- 0x11: membership query
- 0x12: version 1 membership report
- 0x16: version 2 membership report
- 0x17: version 2 leave group
- 0x22: version 3 membership report

This chapter will only deal with IGMP v2 and v3. The description of the v2 and v3 messages is covered in this section.

11.1.6 0x11 Membership Query

The membership query is sent by the IGMP querier to a multicasts address of 224.0.0.1. Any multicast capable device will automatically be listening for this address.

NOTE: All multicasts addresses in the 224.0.0.1 – 244.0.0.255 range are reserved and must not be used for multicast addresses within a system

A Wireshark capture of a Membership Query is shown below:

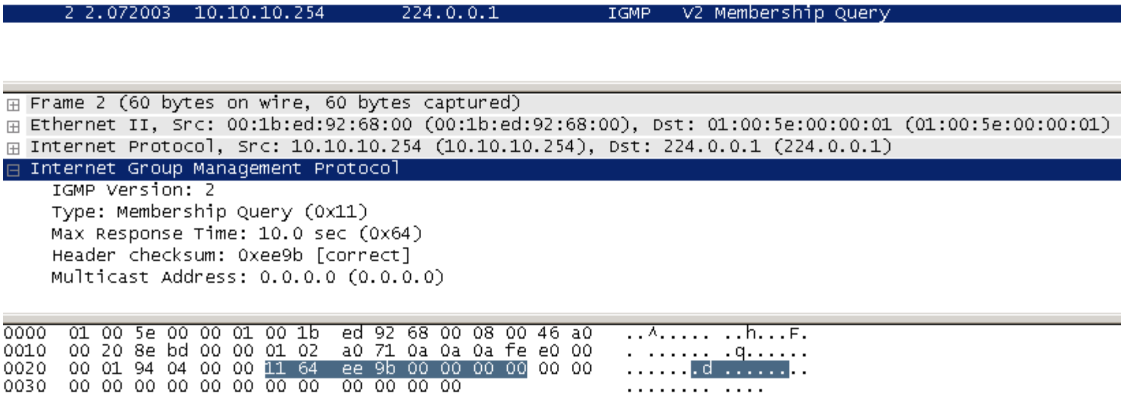

```
    2 2.072003   10.10.10.254        224.0.0.1           IGMP    V2 Membership Query

⊞ Frame 2 (60 bytes on wire, 60 bytes captured)
⊞ Ethernet II, Src: 00:1b:ed:92:68:00 (00:1b:ed:92:68:00), Dst: 01:00:5e:00:00:01 (01:00:5e:00:00:01)
⊞ Internet Protocol, Src: 10.10.10.254 (10.10.10.254), Dst: 224.0.0.1 (224.0.0.1)
⊟ Internet Group Management Protocol
    IGMP Version: 2
    Type: Membership Query (0x11)
    Max Response Time: 10.0 sec (0x64)
    Header checksum: 0xee9b [correct]
    Multicast Address: 0.0.0.0 (0.0.0.0)

0000   01 00 5e 00 00 01 00 1b  ed 92 68 00 08 00 46 a0   ..^..... ..h...F.
0010   00 20 8e bd 00 00 01 02  a0 71 0a 0a 0a fe e0 00   . ...... .q......
0020   00 01 94 04 00 00 11 64  ee 9b 00 00 00 00 00 00   .......d ........
0030   00 00 00 00 00 00 00 00  00 00 00 00               ........ ....
```

Figure 11.2: Wireshark capture of an IGMP Membership Query

As can been seen in the above figure, the Membership Query is sent from an IGMP querier of address 10.10.10.254. This is the address of the VLAN on the switch. It is sent to 224.0.0.1.

The decode window in the middle of the illustration shows that it is a Membership query. It also shows that there is a maximum response time of 10.0 sec. This is not used by hosts as a latest value that it can respond by. It is actually used as the maximum value for a randomly generated timer. When this timer reaches 0 it sends a response to the Membership query. To ensure that all devices on a network do not respond to the query at the same time, this value is used as the maximum value of a randomly generated timer. When the timer reaches zero the response is sent.

If the multicast address is set to 0.0.0.0 then it is a general query. A general query is asking what multicasts are needed. If the address is set to a multicast address range, then it is a query asking if that particular multicast group is needed

The messages are sent out periodically. On a Cisco this is every 60 seconds by default.

11.1.7 0x16 v2 Membership Report

A Membership report is sent in response to a Membership query. It is also sent when a device is configured to receive a multicast. When the latter happens it is known as ‘sending a join message’. It is in fact just a Membership report that is sent.

The figure below shows a Wireshark capture of a v2 Membership report.

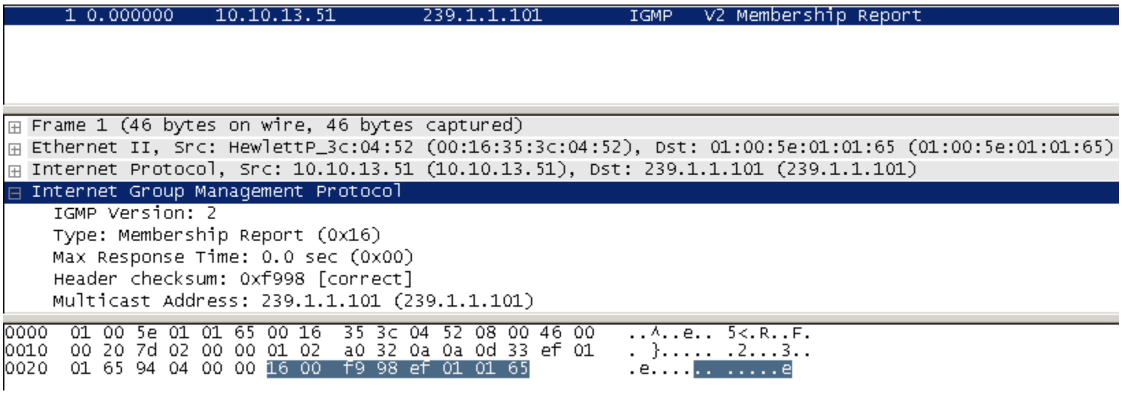

```
    1 0.000000    10.10.13.51         239.1.1.101         IGMP    V2 Membership Report

⊞ Frame 1 (46 bytes on wire, 46 bytes captured)
⊞ Ethernet II, Src: HewlettP_3c:04:52 (00:16:35:3c:04:52), Dst: 01:00:5e:01:01:65 (01:00:5e:01:01:65)
⊞ Internet Protocol, Src: 10.10.13.51 (10.10.13.51), Dst: 239.1.1.101 (239.1.1.101)
⊟ Internet Group Management Protocol
    IGMP Version: 2
    Type: Membership Report (0x16)
    Max Response Time: 0.0 sec (0x00)
    Header checksum: 0xf998 [correct]
    Multicast Address: 239.1.1.101 (239.1.1.101)

0000  01 00 5e 01 01 65 00 16  35 3c 04 52 08 00 46 00   ..^..e.. 5<.R..F.
0010  00 20 7d 02 00 00 01 02  a0 32 0a 0a 0d 33 ef 01   . }..... .2...3..
0020  01 65 94 04 00 00 16 00  f9 98 ef 01 01 65         .e...... ......e
```

Figure 11.3 Wireshark Capture of a v2 Membership Report

In a v2 membership report, the message is sent to the multicast group address that the device is configured to join. In the above example, the device with IP address 10.10.13.51 has been configured to receive multicast address 239.1.1.101. It sends a membership report to 239.1.1.101. The maximum response time is left at zero because it is only relevant to membership query messages and so is not sent in a membership report.

In IGMP v2 a host has to send one membership report for every multicast group it needs to join. If an MX8400 is being fed by 80 Encoders, then every 60 seconds it will be sending 80 membership reports in response the Cisco’s membership query.

11.1.8 0x17 v2 Leave Message

When a host using IGMP v2 stops receiving a multicast, it sends out a leave message.

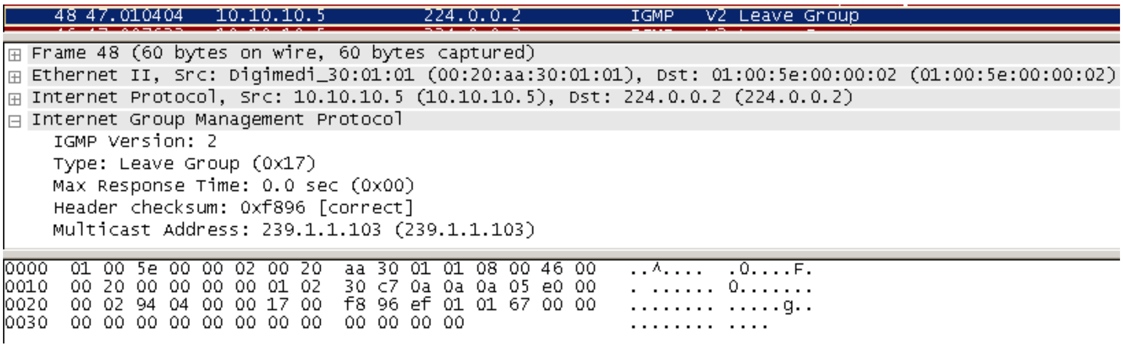

```
   48 47.010404    10.10.10.5          224.0.0.2           IGMP    V2 Leave Group

⊞ Frame 48 (60 bytes on wire, 60 bytes captured)
⊞ Ethernet II, Src: Digimedi_30:01:01 (00:20:aa:30:01:01), Dst: 01:00:5e:00:00:02 (01:00:5e:00:00:02)
⊞ Internet Protocol, Src: 10.10.10.5 (10.10.10.5), Dst: 224.0.0.2 (224.0.0.2)
⊟ Internet Group Management Protocol
    IGMP Version: 2
    Type: Leave Group (0x17)
    Max Response Time: 0.0 sec (0x00)
    Header checksum: 0xf896 [correct]
    Multicast Address: 239.1.1.103 (239.1.1.103)

0000  01 00 5e 00 00 02 00 20  aa 30 01 01 08 00 46 00   ..^.... .0....F.
0010  00 20 00 00 00 00 01 02  30 c7 0a 0a 0a 05 e0 00   . ...... 0.......
0020  00 02 94 04 00 00 17 00  f8 96 ef 01 01 67 00 00   ........ .....g..
0030  00 00 00 00 00 00 00 00  00 00 00 00               ........ ....
```

Figure 11.4 Wireshark Capture of a v2 Leave Message

In the above example, the host is configured to no longer receive the multicast 239.1.1.103. It sends a leave message to the address 224.0.0.2. One leave message is sent for each multicast it is wishes to leave.

11.1.9 0x22 v3 Membership Report

IGMP v3 includes the concept of source specific multicasts. This means a device can specify a list of sources it wants to receive a multicast from, or alternatively, a list of address it does not want to receive a multicast from.

The second major change in v3 Membership Reports, is the ability to send one membership report containing all the required multicasts. This reduced the amount of traffic on the network.

To be able to send one message, a v3 membership report is sent to a multicast address within the reserved range 224.0.0.1 – 224.0.0.255.

In the *Figure 11.5* example, the device with address 10.10.13.51 has been configured to receive the following multicasts:

239.1.1.101
239.1.1.102
239.1.1.103
239.1.1.105
239.1.1.106
239.1.1.107

To join all these multicasts groups, the device sends one v3 membership report containing details of all the relevant multicasts. In the above example, the entry for 239.1.1.101 has been expanded to show the information. It can bee seen that there is one source (src) address listed which is 10.10.10.101.

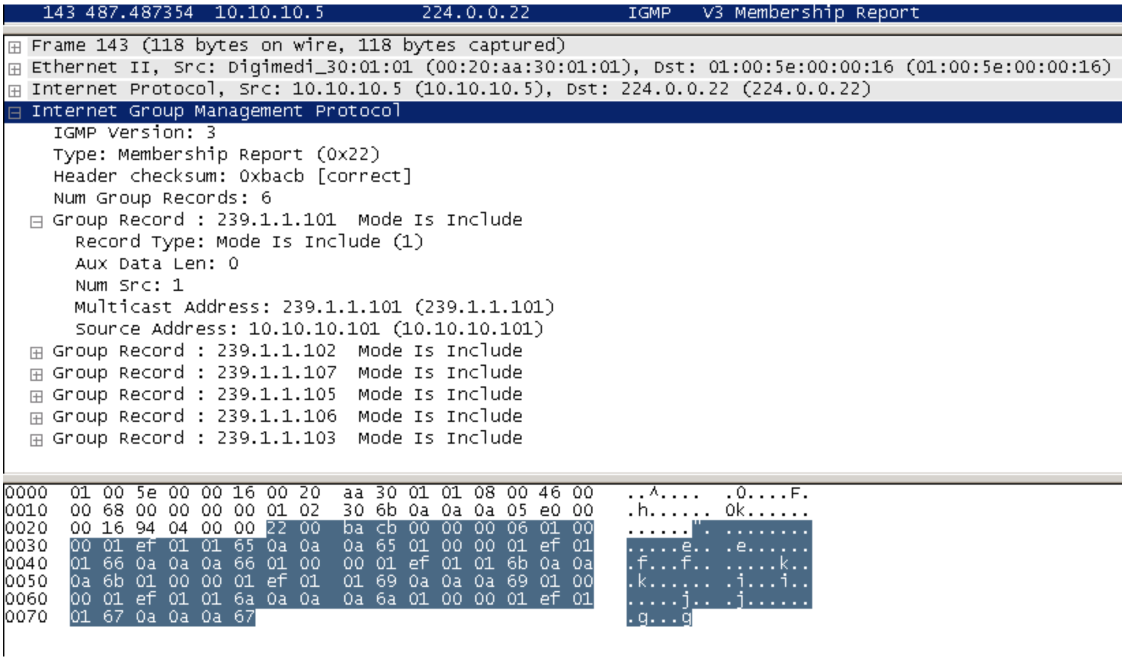

```
   143 487.487354  10.10.10.5           224.0.0.22          IGMP   V3 Membership Report
⊞ Frame 143 (118 bytes on wire, 118 bytes captured)
⊞ Ethernet II, Src: Digimedi_30:01:01 (00:20:aa:30:01:01), Dst: 01:00:5e:00:00:16 (01:00:5e:00:00:16)
⊞ Internet Protocol, Src: 10.10.10.5 (10.10.10.5), Dst: 224.0.0.22 (224.0.0.22)
⊟ Internet Group Management Protocol
    IGMP Version: 3
    Type: Membership Report (0x22)
    Header checksum: 0xbacb [correct]
    Num Group Records: 6
  ⊟ Group Record : 239.1.1.101  Mode Is Include
      Record Type: Mode Is Include (1)
      Aux Data Len: 0
      Num Src: 1
      Multicast Address: 239.1.1.101 (239.1.1.101)
      Source Address: 10.10.10.101 (10.10.10.101)
  ⊞ Group Record : 239.1.1.102  Mode Is Include
  ⊞ Group Record : 239.1.1.107  Mode Is Include
  ⊞ Group Record : 239.1.1.105  Mode Is Include
  ⊞ Group Record : 239.1.1.106  Mode Is Include
  ⊞ Group Record : 239.1.1.103  Mode Is Include

0000  01 00 5e 00 00 16 00 20  aa 30 01 01 08 00 46 00   ..^....  .0....F.
0010  00 68 00 00 00 00 01 02  30 6b 0a 0a 0a 05 e0 00   .h...... 0k......
0020  00 16 94 04 00 00 22 00  ba cb 00 00 00 06 01 00   ......". ........
0030  00 01 ef 01 01 65 0a 0a  0a 65 01 00 00 01 ef 01   .....e.. .e......
0040  01 66 0a 0a 0a 66 01 00  00 01 ef 01 01 6b 0a 0a   .f...f.. .....k..
0050  0a 6b 01 00 00 01 ef 01  01 69 0a 0a 0a 69 01 00   .k...... .i...i..
0060  00 01 ef 01 01 6a 0a 0a  0a 6a 01 00 00 01 ef 01   .....j.. .j......
0070  01 67 0a 0a 0a 67                                  .g...g
```

Figure 11.5 Wireshark Capture of a v3 Membership Report

The mode for all the multicasts is set to **include**. There are two options, **include** and **exclude**. If an entry is set to include, the IGMP querier will forward all multicasts from sources listed in the membership report. If the mode is set to exclude, the IGMP querier will forward all multicasts from sources not listed in the membership report.

A non-source specific v3 join is a message with an empty exclude list.

11.1.10 v3 Leaves

There is no leave message in IGMP v3. A specific membership report is sent to leave multicasts when using IGMP v3

A membership report with an empty **include** list for a multicast is a leave.

11.2 Multiple IGMP Queriers

In a single network there can be only one IGMP querier running at a time. More than one can be configured but only one of them will be running. Consider the following scenario:

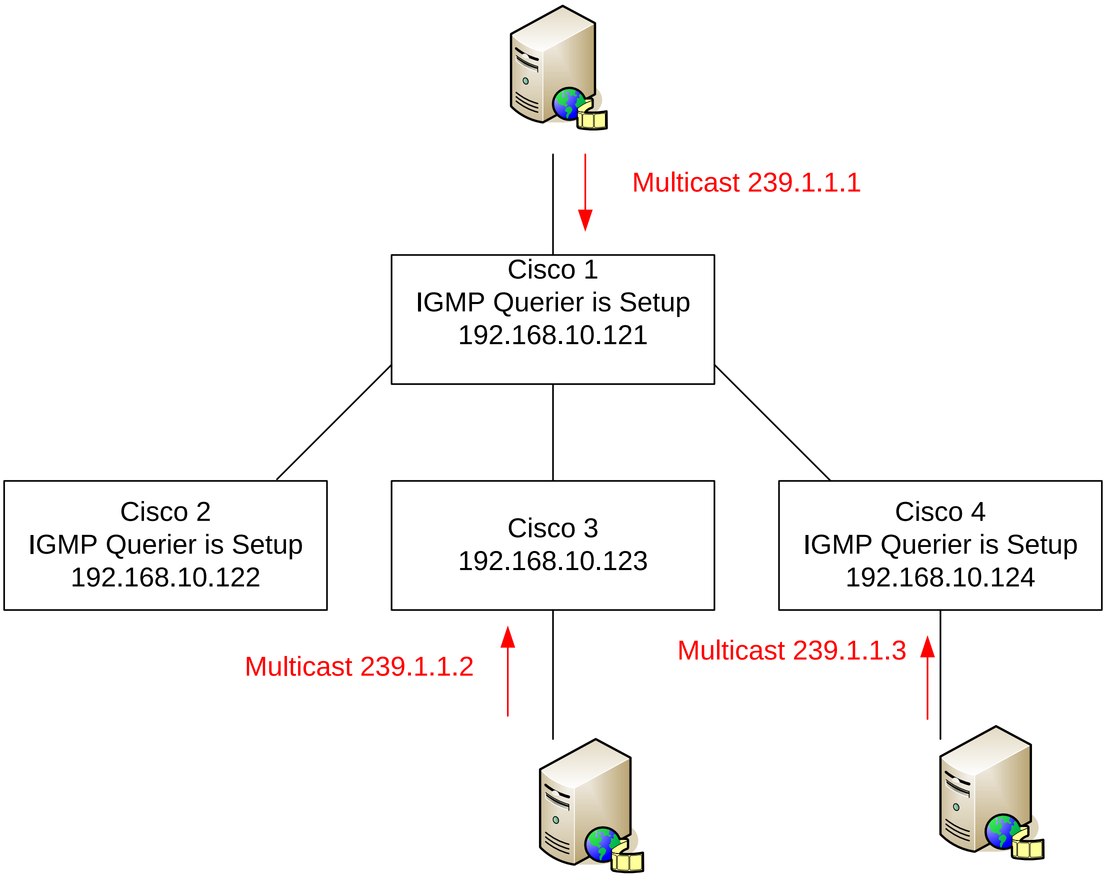

Figure 11.6 Example of multiple IGMP Querier Configuration

In this scenario, all the Ciscos are in the same network. Three of them all have IGMP querier configured, but only one of them will be the querier. This is determined but the lowest IP address. So in the above example, Cisco 1 will be the querier. All multicasts will be forwarded to it by Ciscos 3 and 4.

Because of the above fact, ensure that all Ciscos that are configured to have a querier have the same querier query interval set. This is the interval at which a Membership Query is sent. On a Cisco, the default is 60 seconds. If Cisco 4 has a query interval of 30 seconds configured this will cause a problem in the network: Cisco 1 will send out a query and will be the querier for the network. After 30 seconds Cisco 4 will have reached its interval and sends out a query because no other switches with a lower ip address have done so within the 30 second interval time. It will then be the querier for the network. Another 30 seconds later Cisco 1 will come to the end of its 60 second interval and send out another query, overriding Cisco 4 and become the querier for the network. The IGMP querier will continue to toggle between Cisco 1 and Cisco 4 every 30 seconds.

This example is only true for switches connected together without a PIM router configured. PIM routers run a querier and always take precedence over layer 2 switches running a querier, regardless of the IP addresses.

12 IGMP Snooping

IGMP snooping is separate to IGMP. IGMP is used to track multicast group memberships throughout a network.

An IGMP querier tracks what multicasts are requested, not who has requested them. Therefore on a network that is running IGMP, if a multicast is requested on one of a switch's interfaces, it will be available on all the interfaces on that switch (flooding).

IGMP snooping is used to prevent multicasts from leaving a switch when they have not been requested. IGMP Snooping runs on any switch even if it is not the querier.

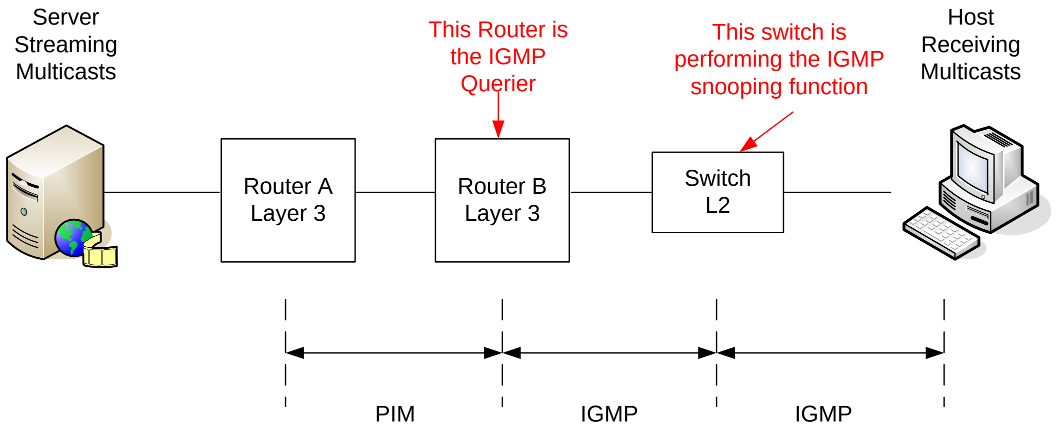
Server Streaming Multicasts
This Router is the IGMP Querier
This switch is performing the IGMP snooping function
Host Receiving Multicasts
Router A Layer 3
Router B Layer 3
Switch L2
PIM
IGMP
IGMP

Figure 12.1: IGMP Snooping being used on the edge switch of a routed network

IGMP Snooping is a layer 2 function. This means it uses only MAC addresses. It snoops the layer 3 query and membership IGMP messages to work out which hosts are requesting multicasts on which interfaces. It then only allows these multicasts out of the relevant interfaces.

Although IGMP snooping is active on Cisco 35xx switches by default. It will only work if there is a querier running.

Because IGMP snooping works on MAC addresses and the IGMP messages are sent to multicast addresses, IGMP snooping will still work even if the querier and host are on different subnets. However, it is best practice to have them on the same network.

BLANK