

IP Multicasting at Layer 2

Module Agenda

- **MAC Layer Multicast Addresses**
- **IGMPv2**
- **IGMPv3**
- **L2 Multicast Frame Switching**
 - **IGMP Snooping**
 - **CGMP**
 - **PIM Snooping**


MAC Layer Multicast Addresses

Ethernet multicast cím

- **0. byte 0. bit**

Layer 2 Multicast Addressing

IP Multicast MAC Address Mapping (FDDI and Ethernet)



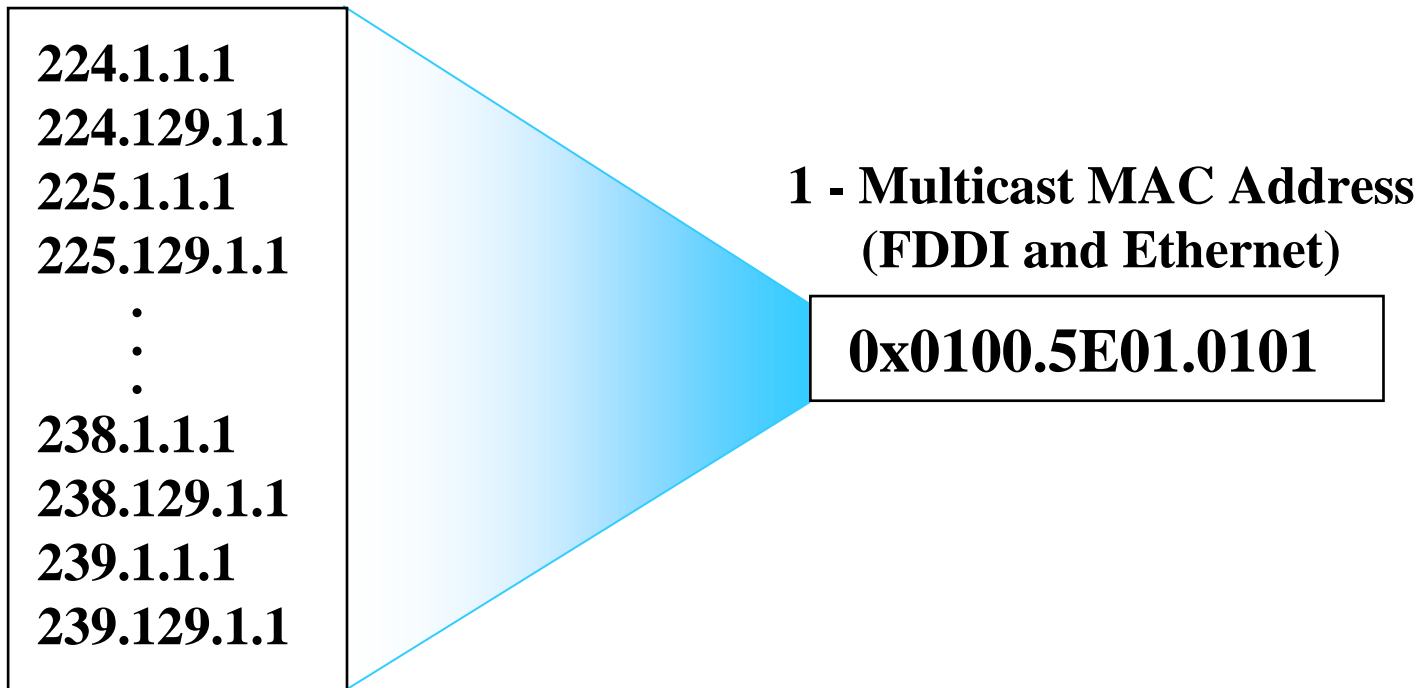
- 16 consecutive Organizational Unique Identifiers (OUIs) (IEEE)
- 1 OUI/1000 USD
- one OUI contains 24 bits worth of address space

Layer 2 Multicast Addressing

IP Multicast MAC Address Mapping (FDDI & Ethernet)

Be Aware of the 32:1 Address Overlap

32 - IP Multicast Addresses



Layer 2 Multicast Addressing

IP Multicast MAC Address Mapping (Token Ring)

A Layer 3 IPmc Address Maps to a single Token Ring Functional Address or the all ones' Broadcast address:

224.x.x.x



c0-00-00-04-00-00

(Shown in Token Ring, non-canonical format)

224.x.x.x



ff-ff-ff-ff-ff-ff

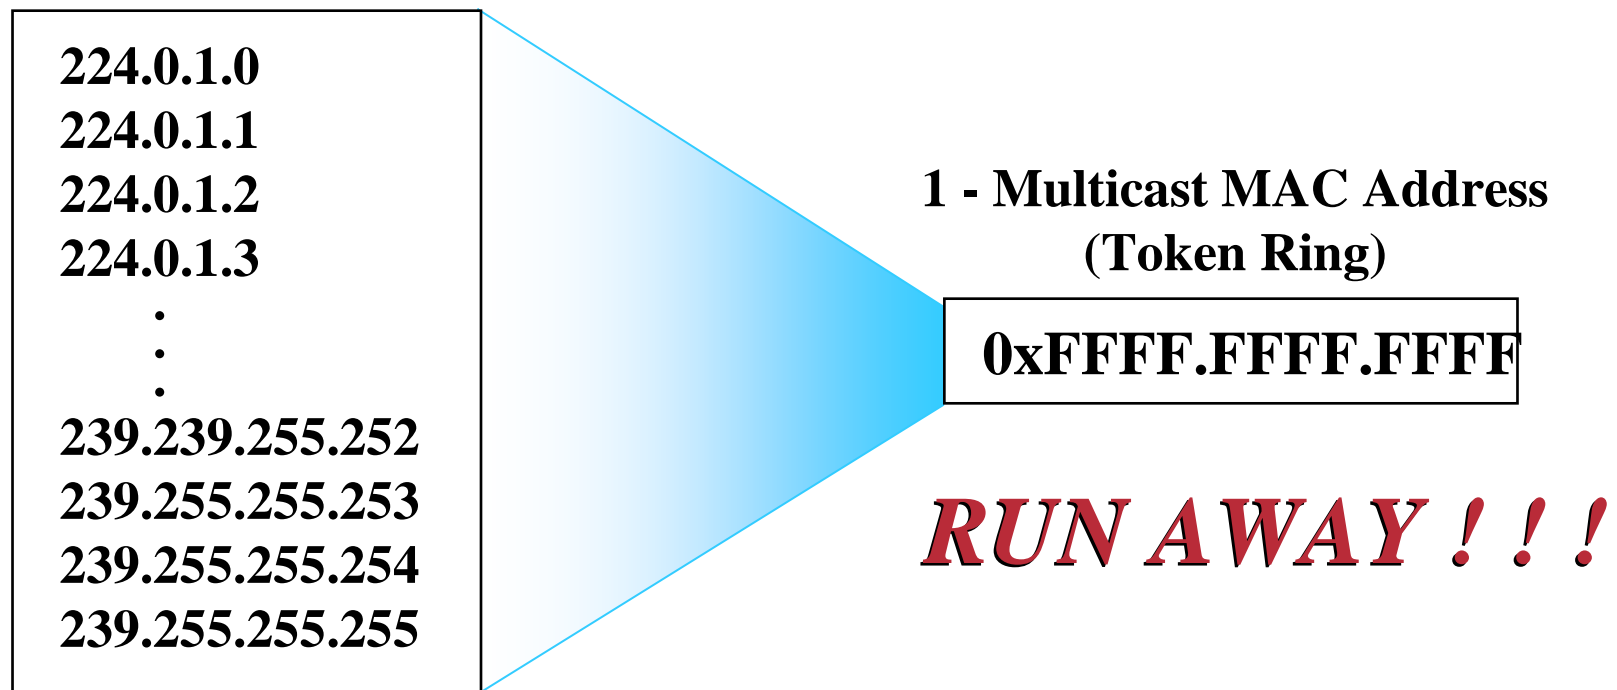
**Results in high levels of unwanted
interrupts for non-interested Hosts**

Layer 2 Multicast Addressing

IP Multicast MAC Address Mapping (Token Ring)

Be Aware of the 268,435,200:1 Address Overlap

ALL 268,435,200 - IP Multicast Addresses



IGMPv2

IGMP

- **How hosts tell routers about group membership**
- **Routers solicit group membership from directly connected hosts**
- **RFC 1112 specifies first version of IGMP**
- **RFC 2236 specifies current version of IGMP**
- **IGMP v3 enhancements**
- **Supported on UNIX systems, PCs, and MACs**

IGMPv2

- **RFC 2236**
 - **Membership Queries**
 - Queries sent to 224.0.0.1 with ttl = 1
 - One router on LAN is elected to send queries
 - Query interval 60–120 seconds
 - **Membership Reports**
 - IGMP report sent by one host suppresses sending by others
 - Restrict to one report per group per LAN
 - Unsolicited reports sent by host, when it first joins the group


IGMPv2

- **RFC 2236**
 - **Group-specific query**
 - Router sends Group-specific queries to make sure there are no members present before stopping to forward data for the group for that subnet
 - **Leave Group message**
 - Host sends leave message if it leaves the group and is the last member (reduces leave latency in comparison to v1)

IGMPv2


- **Querier election mechanism**
 - **On multi-access networks, an IGMP Querier router is elected based on lowest IP address. Only the Querier router sends Queries.**
- **Query-Interval Response Time**
 - **General Queries specify “Max. Response Time” which inform hosts of the maximum time within which a host must respond to General Query. (Improves burstiness of the responses.)**
- **Backward compatible with IGMPv1**

IGMPv2 – Joining a Group



- **Joining member sends report to 224.1.1.1 immediately upon joining (same as IGMPv1)**

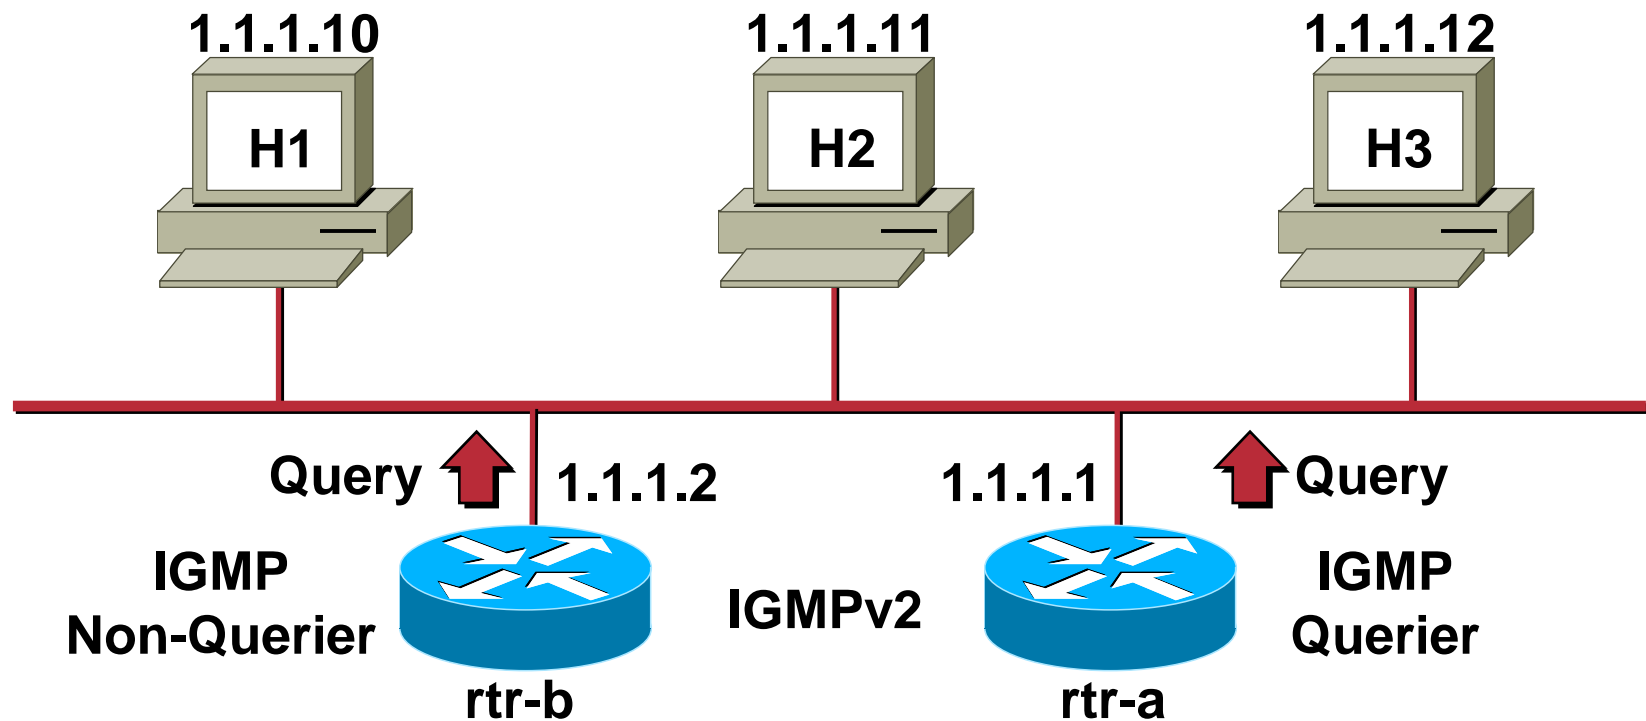
IGMPv2 – Joining a Group



IGMP State in “rtr-a”

```
rtr-a>show ip igmp group
IGMP Connected Group Membership
Group Address      Interface      Uptime        Expires       Last Reporter
224.1.1.1         Ethernet0     6d17h         00:02:31     1.1.1.11
```


IGMPv2 – Querier Election



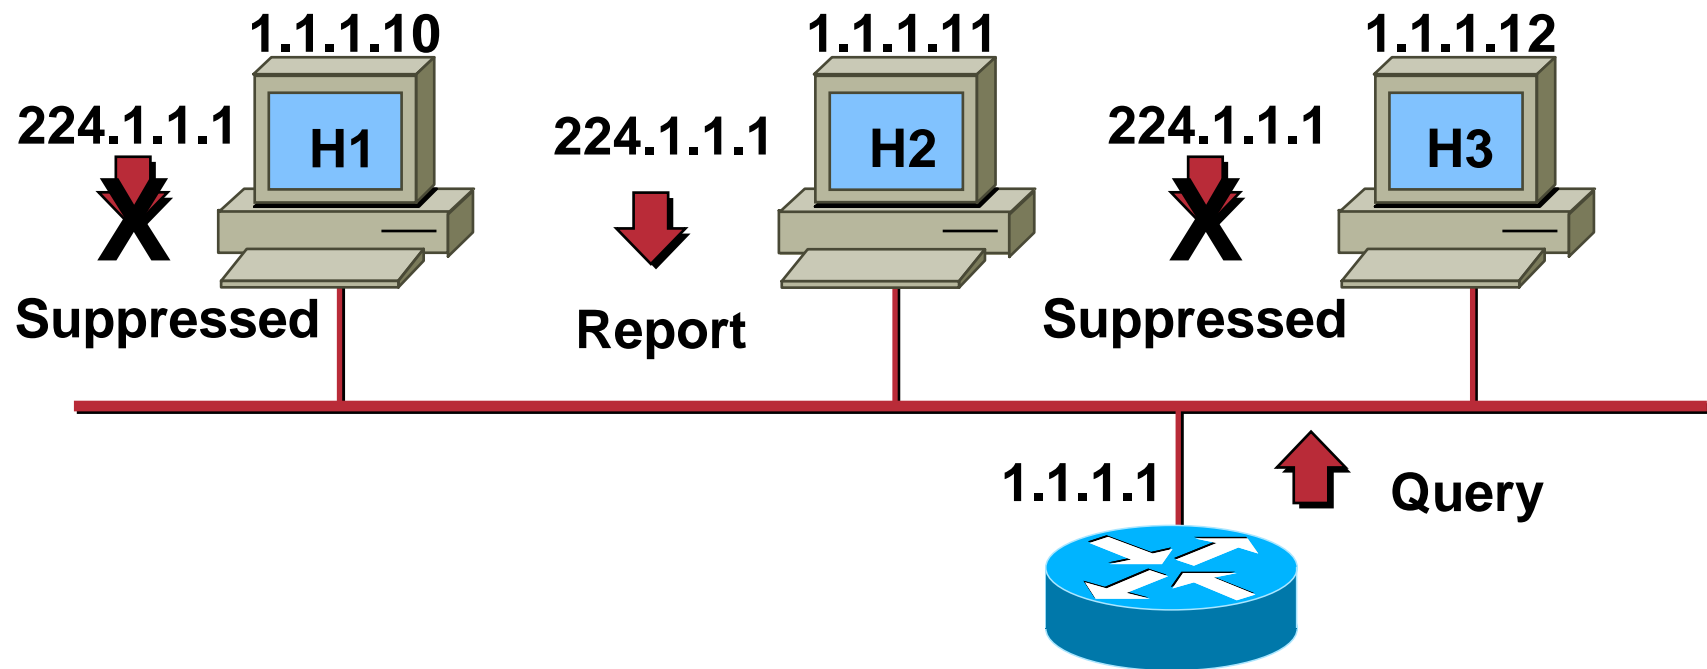
- Initially all routers send out a Query
- Router w/lowest IP address “elected” querier
- Other routers become “Non-Queriers”

IGMPv2 – Querier Election

Determining which router is the IGMP Querier


```
rtr-a>show ip igmp interface e0
Ethernet0 is up, line protocol is up
  Internet address is 1.1.1.1, subnet mask is 255.255.255.0
  IGMP is enabled on interface
  Current IGMP version is 2
  CGMP is disabled on interface
  IGMP query interval is 60 seconds
  IGMP querier timeout is 120 seconds
  IGMP max query response time is 10 seconds
  Inbound IGMP access group is not set
  Multicast routing is enabled on interface
  Multicast TTL threshold is 0
  Multicast designated router (DR) is 1.1.1.1 (this system)
  IGMP querying router is 1.1.1.1 (this system)
  Multicast groups joined: 224.0.1.40 224.2.127.254
```

IGMPv2 – Maintaining a Group



- Router sends periodic queries
- One member per group per subnet reports
- Other members suppress reports

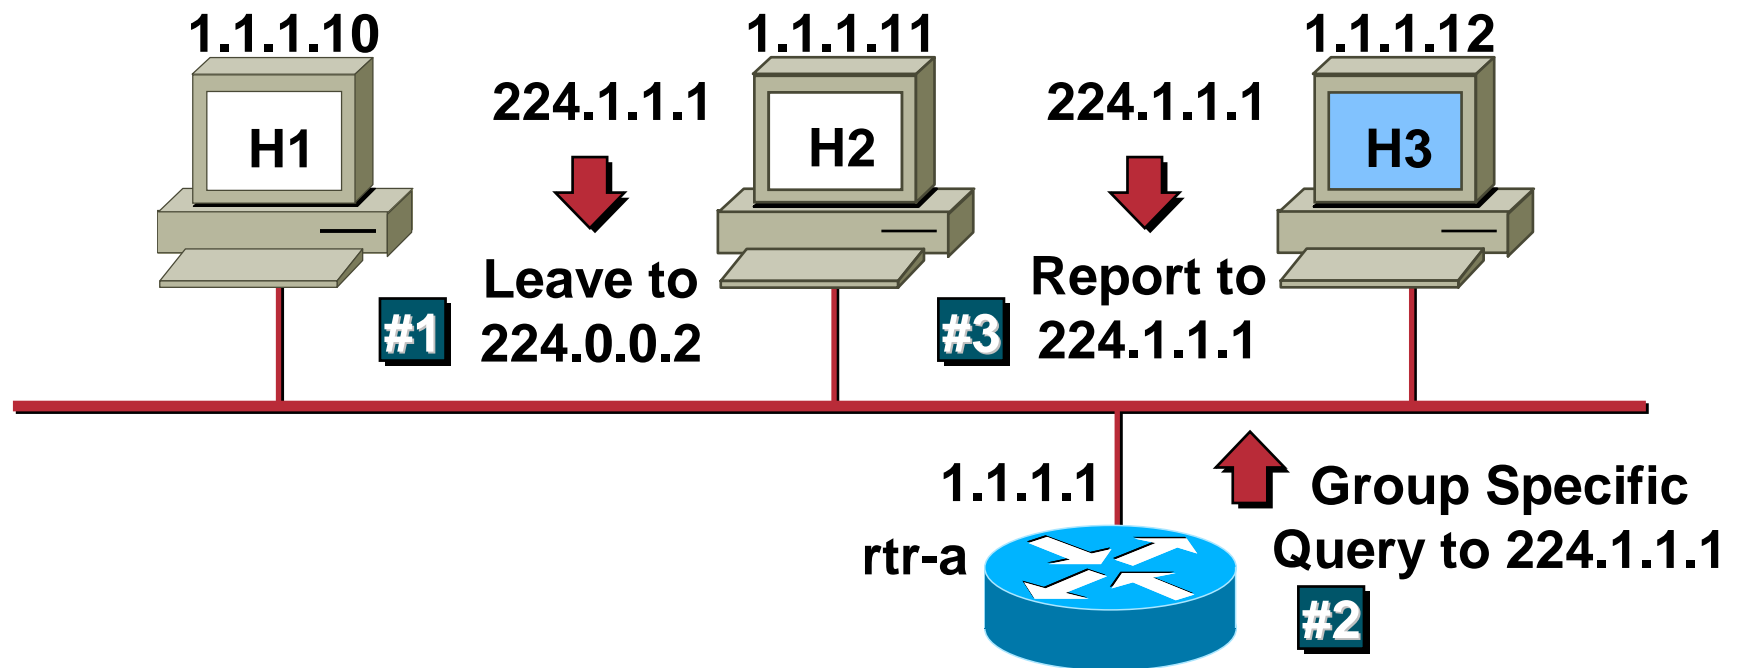
IGMPv2 – Leaving a Group



IGMP State in “rtr-a” before Leave


```
rtr-a>sh ip igmp group
IGMP Connected Group Membership
Group Address      Interface      Uptime        Expires       Last Reporter
224.1.1.1         Ethernet0     6d17h         00:02:31     1.1.1.11
```

IGMPv2 – Leaving a Group



- H2 leaves group; sends Leave message
- Router sends Group specific query
- A remaining member host sends report
- Group remains active


IGMPv2 – Leaving a Group



IGMP State in “rtr-a” after H2 Leaves


```
rtr-a>sh ip igmp group
IGMP Connected Group Membership
Group Address      Interface      Uptime        Expires       Last Reporter
224.1.1.1         Ethernet0     6d17h         00:01:47     1.1.1.12
```

IGMPv2 – Leaving a Group



- Last host leaves group; sends Leave message
- Router sends Group specific query
- No report is received
- Group times out

IGMPv2 – Leaving a Group



IGMP State in “rtr-a” after H3 Leaves

```
rtr-a>show ip igmp group
IGMP Connected Group Membership
Group Address      Interface      Uptime      Expires      Last Reporter
```


IGMPv3

IGMPv3

- **RFC 3376**
 - **Adds Include/Exclude Source Lists**
 - **Enables hosts to listen only to a specified subset of the hosts sending to the group**
 - **Requires new 'IPMulticastListen' API**
 - **New IGMPv3 stack required in the O/S.**
 - **Apps must be rewritten to use IGMPv3 Include/Exclude features**
 - **Available in IOS 12.2, 12.1(3)T and 12.0(15)S.**

IGMPv3

- **RFC 3376**
 - **New Membership Report address**
 - **224.0.0.22 (All-IGMPv3-Routers)**
 - All IGMPv3 Hosts send reports to this address
 - » Instead of the target group address as in IGMPv1/v2
 - All IGMPv3 Routers listen to this address
 - Hosts do not listen or respond to this address
 - **No Report Suppression**
 - All Hosts on wire respond to Queries
 - Response Interval may be tuned over broad range
 - » Useful when large numbers of hosts reside on subnet

IGMPv3 – Query Packet Format

Type = 0x11
IGMP Query

Max. Resp. Time

Max. time to send a response
if < 128, Time in 1/10 secs
if > 128, FP value (12.8 - 3174.4 secs)

Group Address:

Multicast Group Address
(0.0.0.0 for General Queries)

S Flag

Suppresses processing by routers

QRV (Querier Robustness Value)

Affects timers and # of retries

QQIC (Querier's Query Interval)

Same format as Max. Resp. Time

Number of Sources (N)

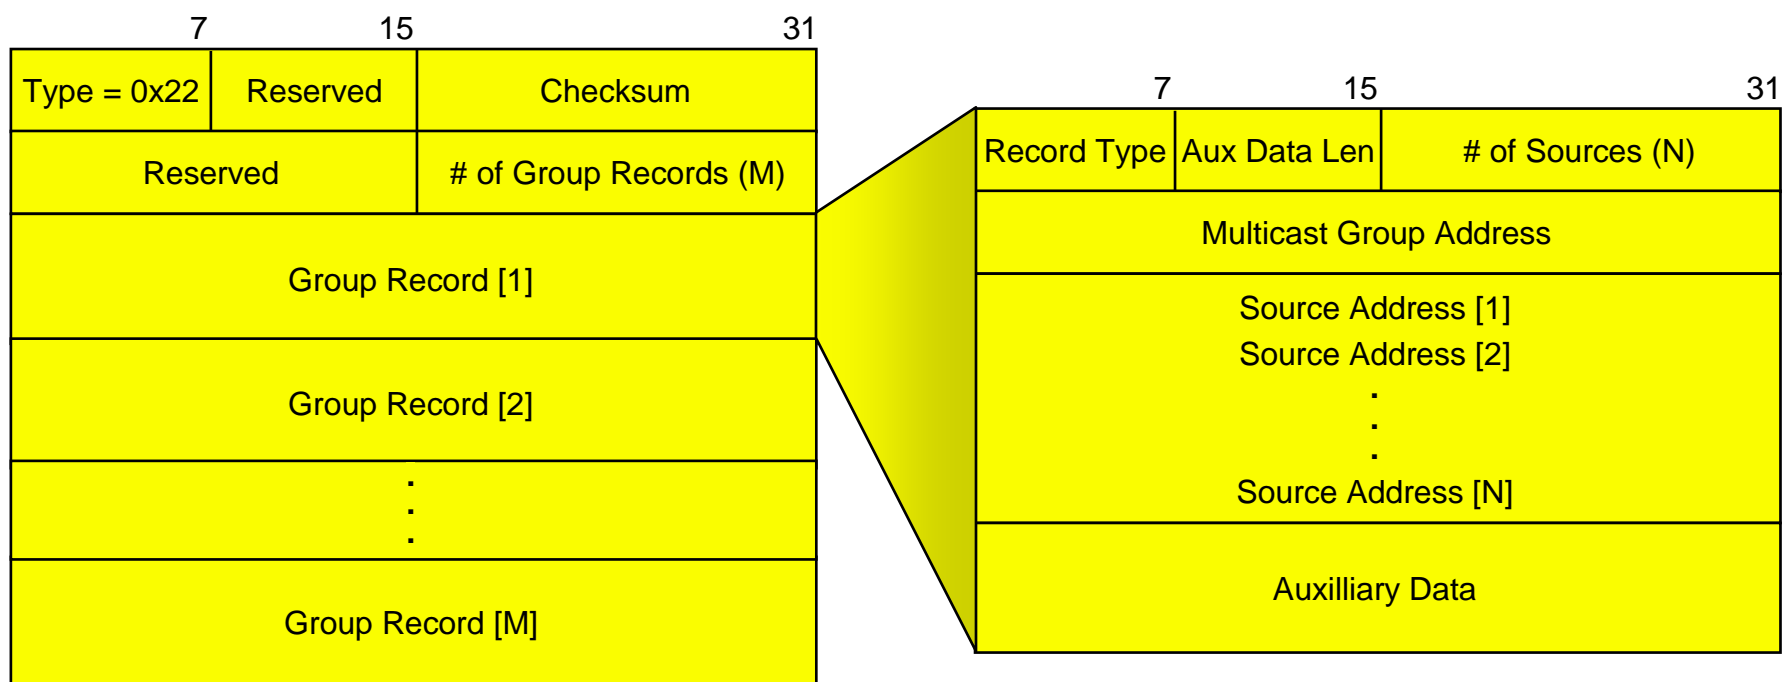
(Non-zero for Group-and-Source Query)

Source Address

Address of Source

| | | | | | |
|--------------------|-----|-----------------|--|-----------------------|--|
| 7 | | 15 | | 31 | |
| Type = 0x11 | | Max. Resp. Code | | Checksum | |
| Group Address | | | | | |
| S | QRV | QQIC | | Number of Sources (N) | |
| Source Address [1] | | | | | |
| Source Address [2] | | | | | |
| . | | | | | |
| . | | | | | |
| . | | | | | |
| Source Address [N] | | | | | |

IGMPv3 – Report Packet Format



of Group Records (M)
Number of Group Records in Report

Group Records 1 - M
Group address plus list of zero or more sources to Include/Exclude (See Group Record format)

Record Type
Include, Exclude, Chg-to-Include, Chg-to-Exclude, Add, Remove

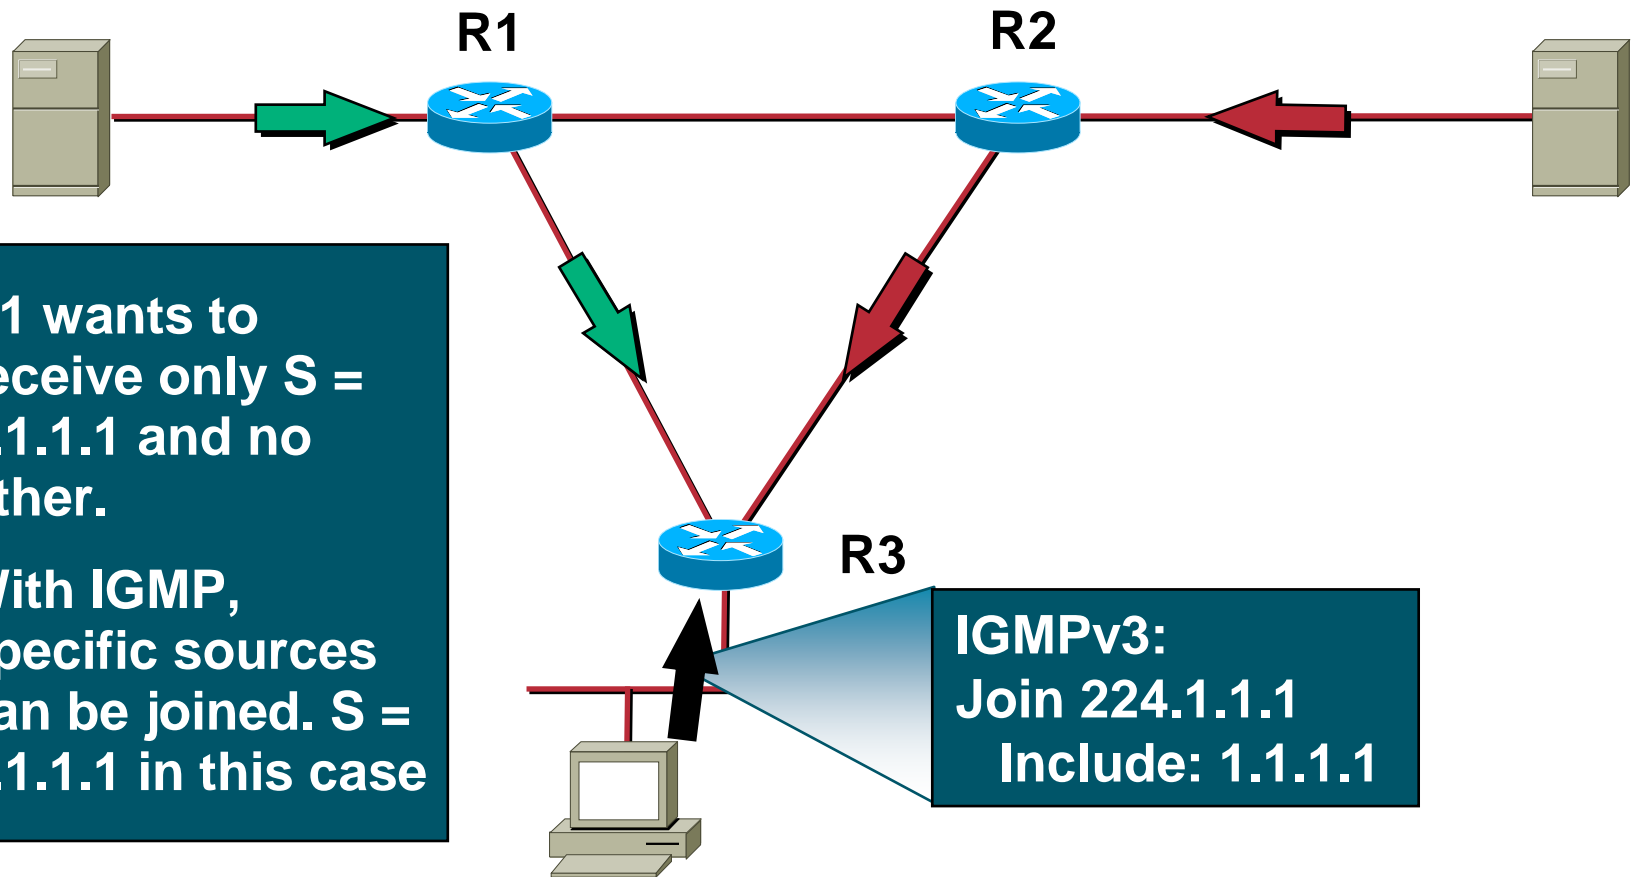
of Sources (N)
Number of Sources in Record

Source Address 1- N
Address of Source

IGMPv3 Example

Source = 1.1.1.1
Group = 224.1.1.1

Source = 2.2.2.2
Group = 224.1.1.1




- H1 wants to receive only S = 1.1.1.1 and no other.
- With IGMP, specific sources can be joined. S = 1.1.1.1 in this case

IGMPv3:
Join 224.1.1.1
Include: 1.1.1.1


H1 - Member of 224.1.1.1

IGMPv3 – Joining a Group



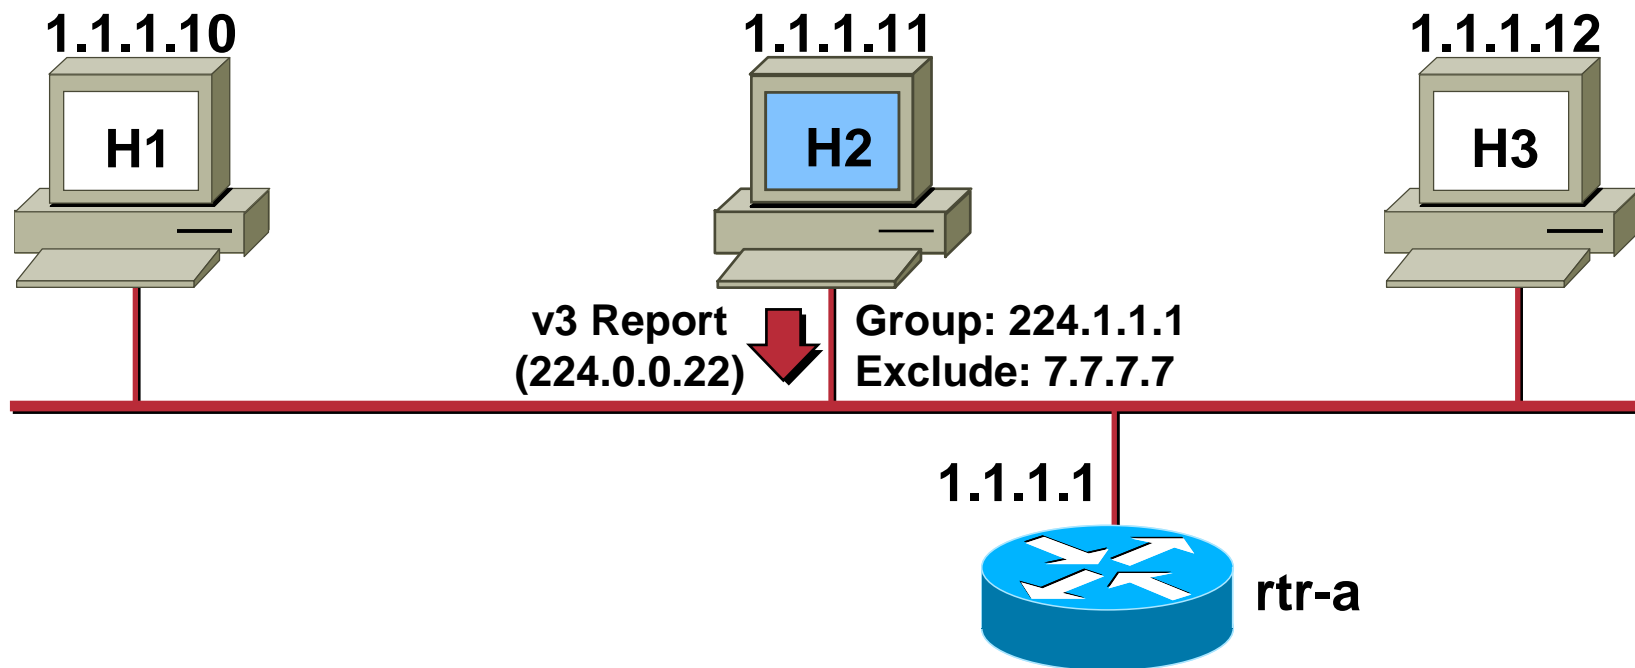
- **Joining member sends IGMPv3 Report to 224.0.0.22 immediately upon joining**

IGMPv3 – Joining specific Source(s)




- **IGMPv3 Report contains desired source(s) in the Include list.**
- **Only “Included” source(s) are joined.**

IGMPv3 – Excluding specific Source(s)



- IGMPv3 Report contains undesired source(s) in the Exclude list.
- All sources except “Excluded” source(s) are joined.

IGMPv3 – Maintaining State




- Router sends periodic queries
- All IGMPv3 members respond
 - Reports contain multiple Group state records

L2 Multicast Frame Switching

L2 Multicast Frame Switching

Problem: Layer 2 Flooding of Multicast Frames

- Typical L2 switches treat multicast traffic as unknown or broadcast and must “flood” the frame to every port
- Static entries can sometimes be set to specify which ports should receive which group(s) of multicast traffic
- Dynamic configuration of these entries would cut down on user administration



IGMP Snooping

L2 Multicast Frame Switching

Solution 1: IGMP Snooping

- Switches become “IGMP” aware
- IGMP packets intercepted by the NMP or by special hardware ASICs
- Switch must examine contents of IGMP messages to determine which ports want what traffic
 - IGMP membership reports
 - IGMP leave messages
- Impact on switch:
 - Must process ALL Layer 2 multicast packets
 - Admin. load increases with multicast traffic load
 - Requires special hardware to maintain throughput

