Multicast Handling in 802.1 AVB Infrastructure

Gaël MACÉ
Thomson CR / CP&M Lab. (Rennes / France)

20th November 2007

Basic multicast handling at level 2

- Since a multicast group MAC address is never used as source MAC address for a packet, a 802.1 switch can not learn them by the classical MAC address learning method.

➔ Therefore, switches act for multicast as for broadcast: **forward multicast traffic to all ports**.

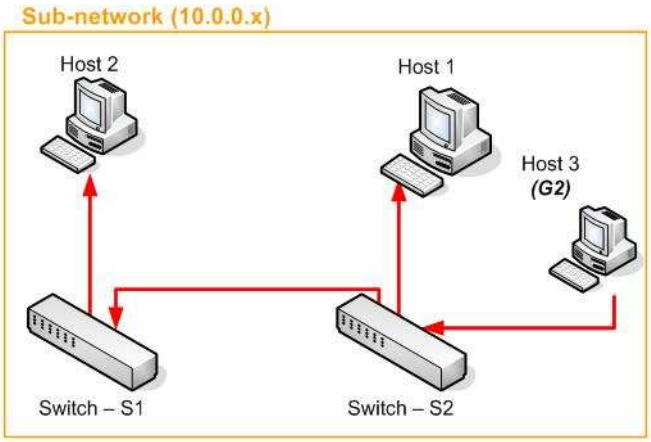

Assumption on strict multicast handling at level 2 (1/2)

- Any not-requested / unregistered multicast A/V traffic must be blocked by all switches

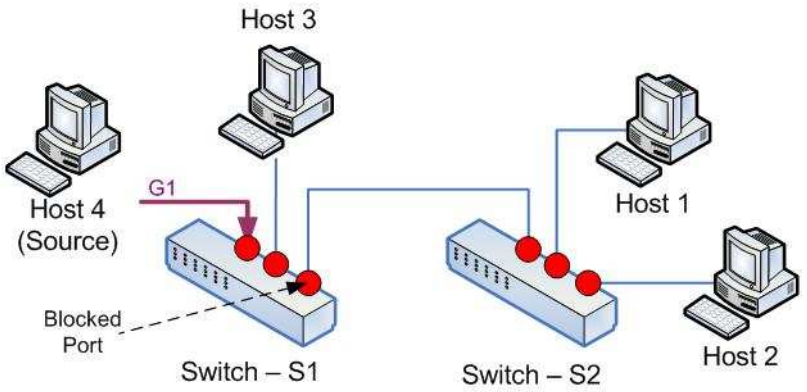

- Once a host wants to receive a multicast A/V traffic, it must subscribe for this multicast group using a “Join” like request that is propagated until the multicast source so as to ensure that every switch on the path is able to mark ports toward the host as forwarding for this multicast group.
- Port which won’t belong to this path, shall remain blocked to this given multicast A/V traffic.

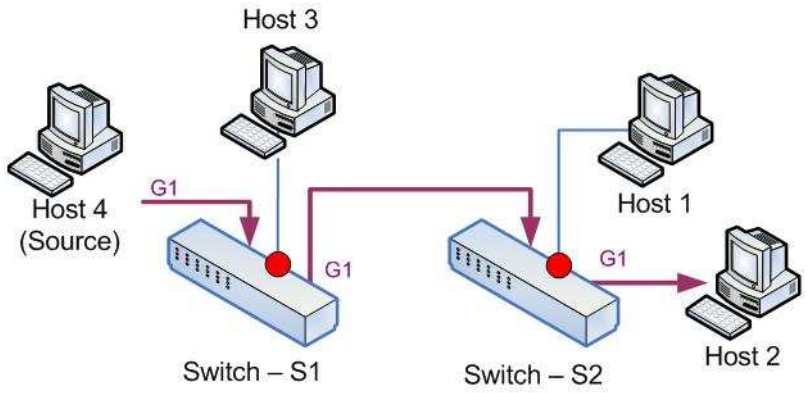

Assumption on strict multicast handling at level 2 (2/2)

- To release properly a subscription, a host shall use a “Leave” like request that is propagated at least toward the source, so that every switch knows that the host is not interested in the multicast group.

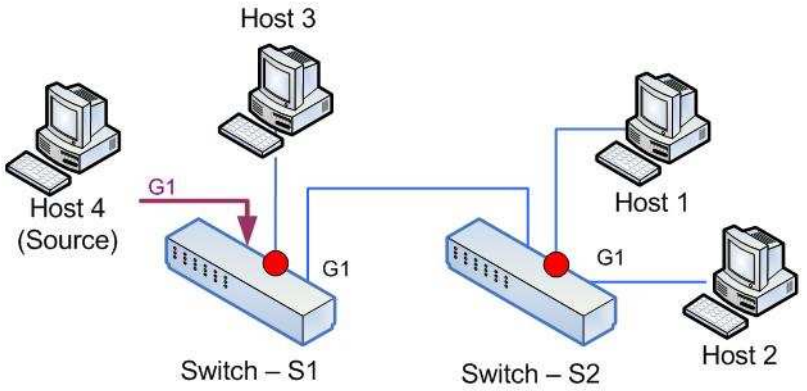

- A periodic “garbage collector” mechanism shall check the adequacy between subscriptions and switch’s port state, and block any port which is no more on a given multicast A/V traffic path.

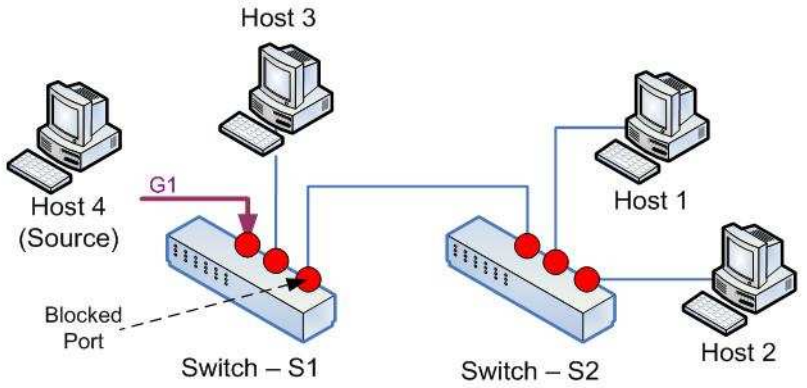

IGMP: introduction (1/4)

IGMP is a multicast group membership tracer protocol that enables routers to route multicast flows on right sub-networks.

•Version 1 defined in RFC 1112 : light version that does not handle leave request

•Version 2 defined in RFC 2236

•Version 3 defined in RFC 3376 : an extension of version 2 that adds the possibility of selecting multicast sources

The current presentation is based on version 2

- Without “subscription” information on given multicast flow, IGMP routers confine multicast traffic to sub-networks.
- All routers receive the complete IGMP traffic to maintain multicast group activity and to ensure multicast routing.

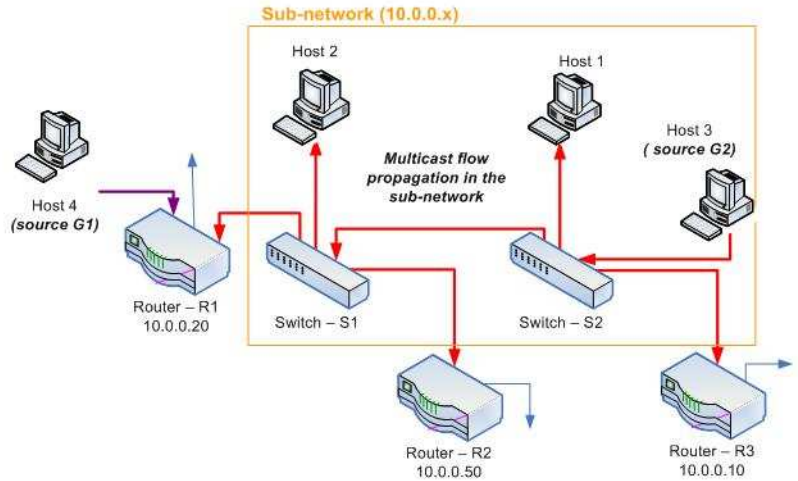

IGMP: *Join request* (2/4)

This request is sent to the multicast group address that the host wants to join. The TTL (Time To Live) of the frame set to 1 ensures that the request does not go beyond the sub-network but that all IGMP routers receive the “Join Request”

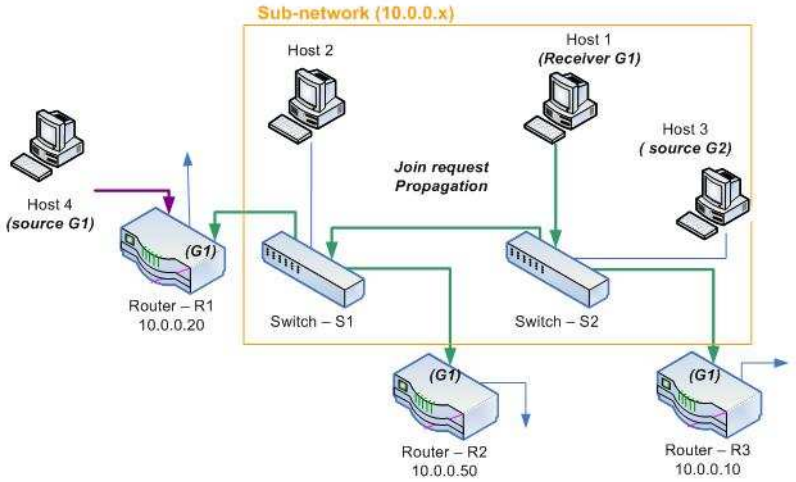

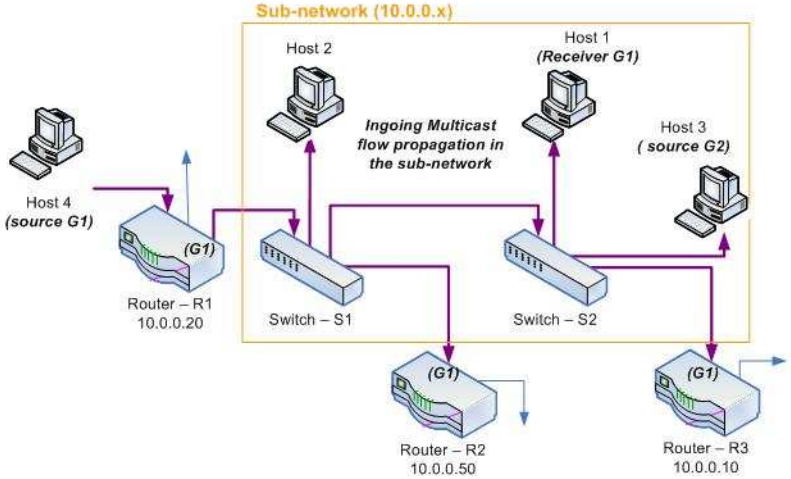

This request opens the sub-network to the requested source flow through the marked port of the IGMP router

IGMP: *Membership Query and Report* (3/4)

- A host belongs to a multicast group for fixed amount of time.
- The IGMP Querier is the router that has the lowest IP address.
- The IGMP Querier is in charge of sending regularly “General” and “Group” queries:
  - General Queries ask if some hosts still want to belong to a multicast group.
  - Group specific queries ask if some hosts still want to belong to the multicast group specified in the query.

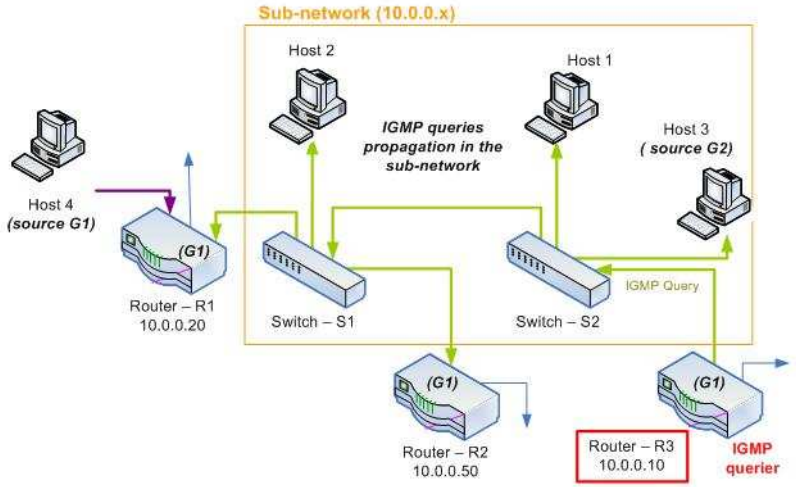

- To reduce IGMP traffic, a single IGMP querier is elected.

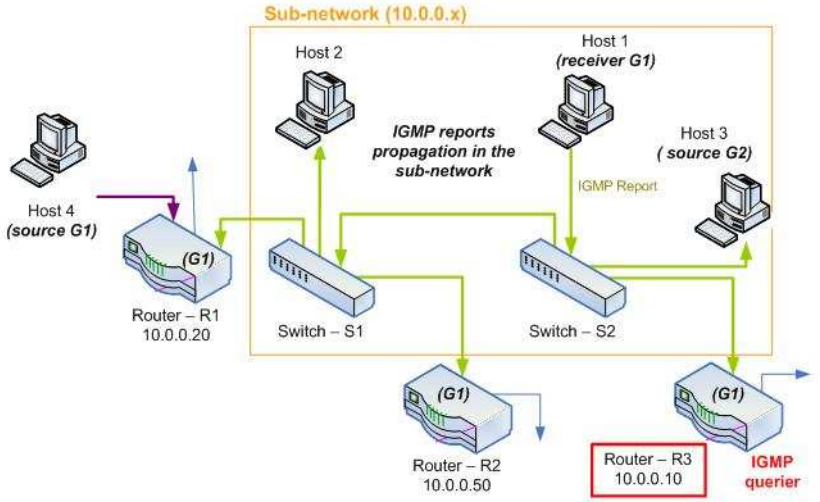

In response of a query, one host per multicast group must send a Membership Report specifying the belonging to this multicast group.

To reduce IGMP traffic, only one report has to be sent for a given group

IGMP: *Leave* and Version 3 (4/4)

- To leave a multicast group, a host has to send a “Leave Group” message
  - As IGMP v2 only maintains multicast group activity, the IGMP Querier does not know if there are hosts still interested in the multicast group notified with the “Leave Group” message. Therefore, it sends a Membership Query.

- IGMP v3 implements multicast source selection
  - a host belonging to a multicast group can choose to only receive this multicast flow from specific sources and not all as in IGMP v2. This requires that IGMP Queriers handle per-host multicast group belonging. As a host may select a set of multicast sources, routers must know the presence of each host involved in multicast traffic.
  - IGMP v3 protocol however works as the version 2, Membership Report are simply more complex to handle source selection
- No “Leave Group” message
  - The leave group functionality is ensured by the Membership Report message that not only handles leave functionality but also source selection changes.

IGMP Snooping: Introduction (1/3)

Enable to forward multicast traffic only to hosts that are interested in in a sub-network, by tracking IGMP traffic between hosts and IGMP routers.

- Defined in RFC 4541 (May 2006)
- **As long as the RFC is only informational, switch behavior could differ**
- Implemented in layer-2 switches by tracking IGMP exchanges and maintain registration of host for multicast groups.
- Differentiation between IGMP traffic forwarding and data multicast traffic forwarding

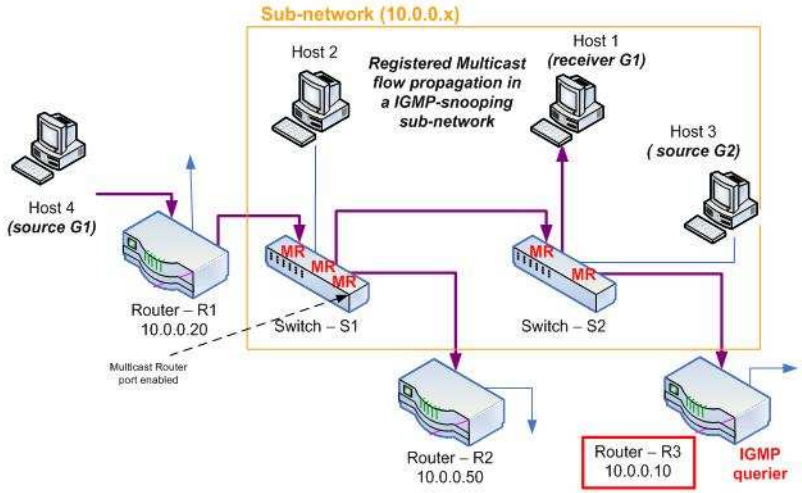

IGMP Snooping: IGMP Traffic forwarding (2/3)

- **IGMP Membership Reports must be forwarded only to multicast routers**
  - Implies that every switch knows where multicast routers are:
    - IGMP Multicast Router Discovery protocol
    - Tracking the IGMP Queries (during the IGMP querier election)
    - Explicit configuration of ports to be IGMP-forwarding
  - Timeout for Report membership after seeing an IGMP Query
- **IGMP Snooping must work with at least on Querier on the sub-network**
  - the case were none is present is not defined by the RFC.

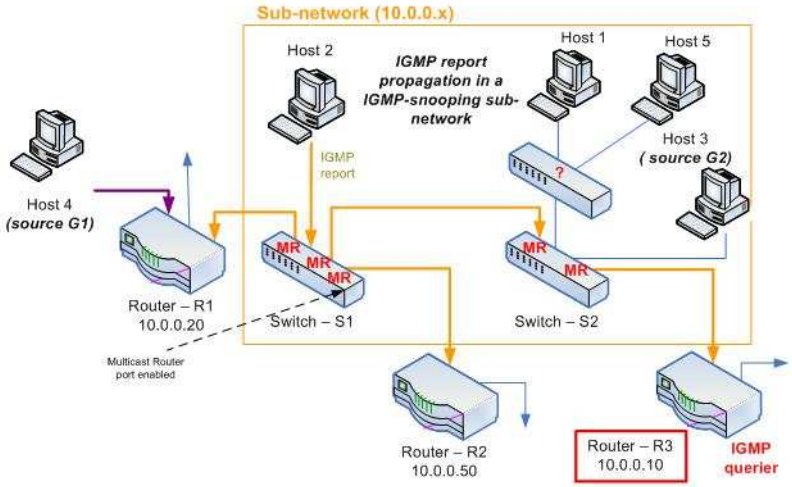

IGMP Snooping: Data multicast traffic forwarding (3/3)

- Multicast traffic for reserved IP addresses, i.e. 224.0.0.x, is propagated on all ports of switches
  - The need of host subscription for one of this multicast group depends on its definition

- **Registered multicast traffic** for classic group (outside 224.0.0.x IP addresses) must be forwarded to ports according to multicast membership information and also **must be forwarded to router ports**
  - A “registered” multicast group on a switch is a multicast group which has been tracked on the switch, i.e. which at least one Report for this group has been seen on a port

- **Unregistered multicast traffic must be forwarded on all ports of IGMP Snooping switches**

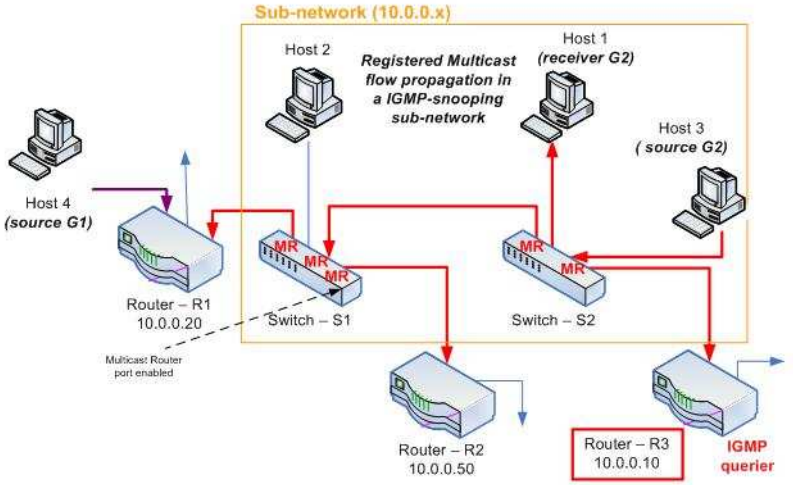

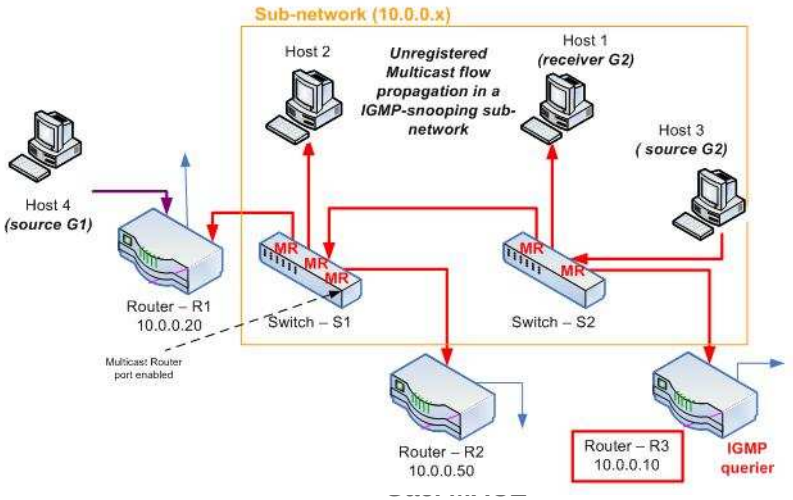

Conclusion

- **As defined nowadays, IGMP snooping doesn't not fit the requirements of strict multicast traffic inside a sub-network:**
  - RFC4541 is above all dedicated to ISP or multicast flow distributor in general. IGMP Snooping is efficient to route, on a LAN network, multicast traffic coming from outside the sub-network and entering it. However, IGMP Snooping does not conveniently convey multicast traffic targeted to remain inside the LAN network.
  - Because of some fuzziness in the RFC, current off-the-shelf switches behave greatly differently.

- **Potential Solutions**
  - *Use MMRP*
  - *RFC4541 amendments*:
    - Don't forward the multicast traffic to multicast router (or any equipment playing the role of IGMP querier)
    - Broadcast the IGMP reports over all the LAN network