Optimizing Traffic Distribution in Multicast Snooping LANs

Joseph Huang
David Vandewalle
Nate Karstens

Garmin

IETF 121 pim Working Group

Why IETF / pim?

- RFC 4541 – Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches
  - May 2006
  - Multicast & Anycast Group Membership (magma) Working Group
- Likely solutions involve revisions to RFC 4541 and controlling distribution of IGMP/MLD messages

IGMP/MLD

- Used by hosts to report multicast group membership to neighboring multicast routers
- Messages
  - Query (General and Group-Specific)
    - Used by multicast routers to solicit updates for multicast group membership
    - “There is normally only one Querier per physical network”
    - There is an algorithm to negotiate the Querier
    - Intercepted by multicast snooping switches and used to mark mrouter port
  - Report
    - Indicates group membership
    - IGMPv2/MLDv1 report suppression
    - Intercepted by multicast snooping switches and used to update port membership bit-field in ATU table

The Problem – RFC 4541 Forwarding Rules

§2.1.2 (Data Path) – Packets with a destination IP address outside 224.0.0.X which are not IGMP should be forwarded according to group-based port membership tables and **must also be forwarded on router ports**.

The Problem – RFC 4541 Forwarding Rules

§2.1.1 (Control Path) – A snooping switch should forward IGMP Membership Reports **only to those ports where multicast routers are attached**.

Alternatively stated: a snooping switch should not forward IGMP Membership Reports **to ports on which only hosts are attached**.

What about switches?

Problem Network

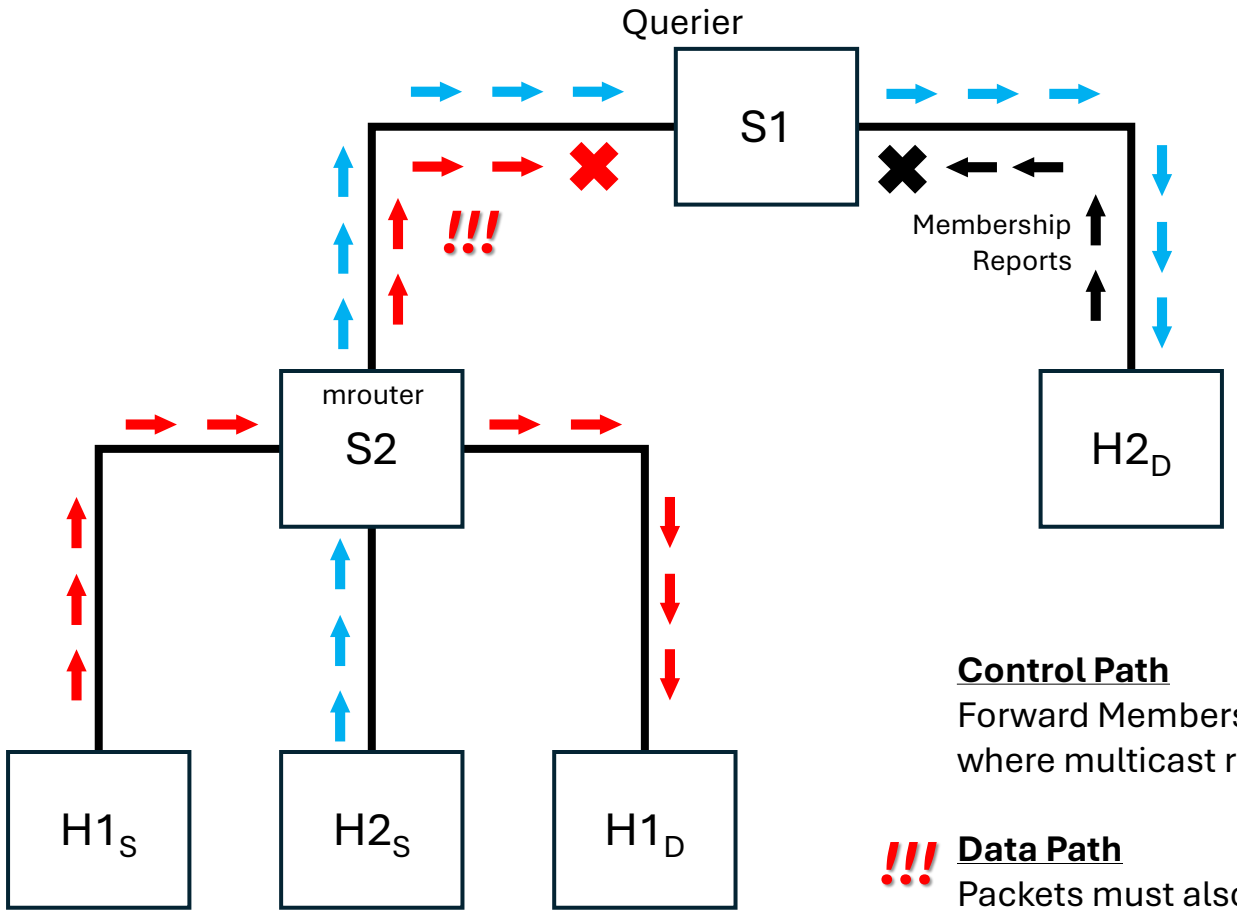

Control Path
Forward Membership Reports only to those ports where multicast routers are attached

!!! **Data Path**
Packets must also be forwarded on router ports

Fix #1

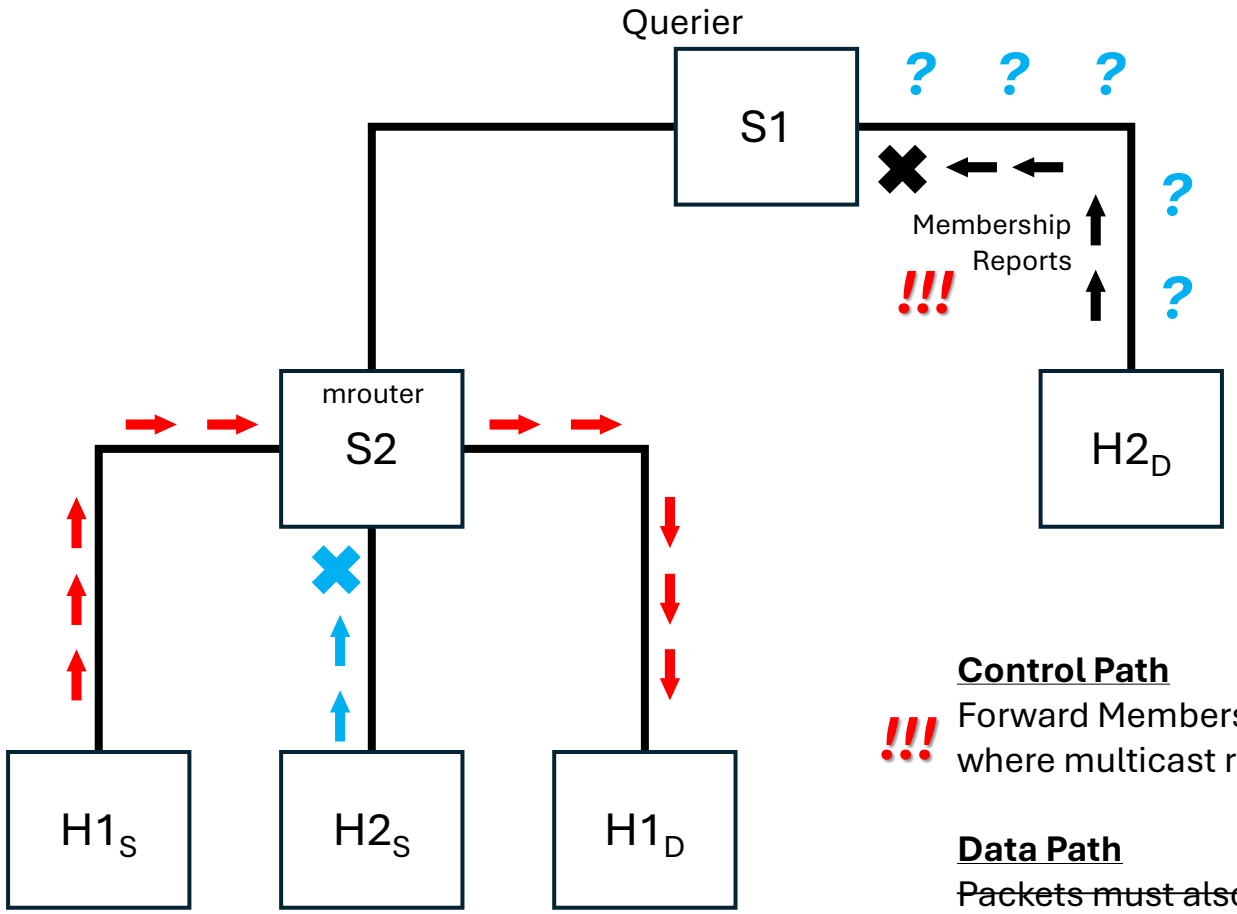

!!! **Control Path**

Forward Membership Reports only to those ports where multicast routers are attached

Data Path

~~Packets must also be forwarded on router ports~~

Fix #2

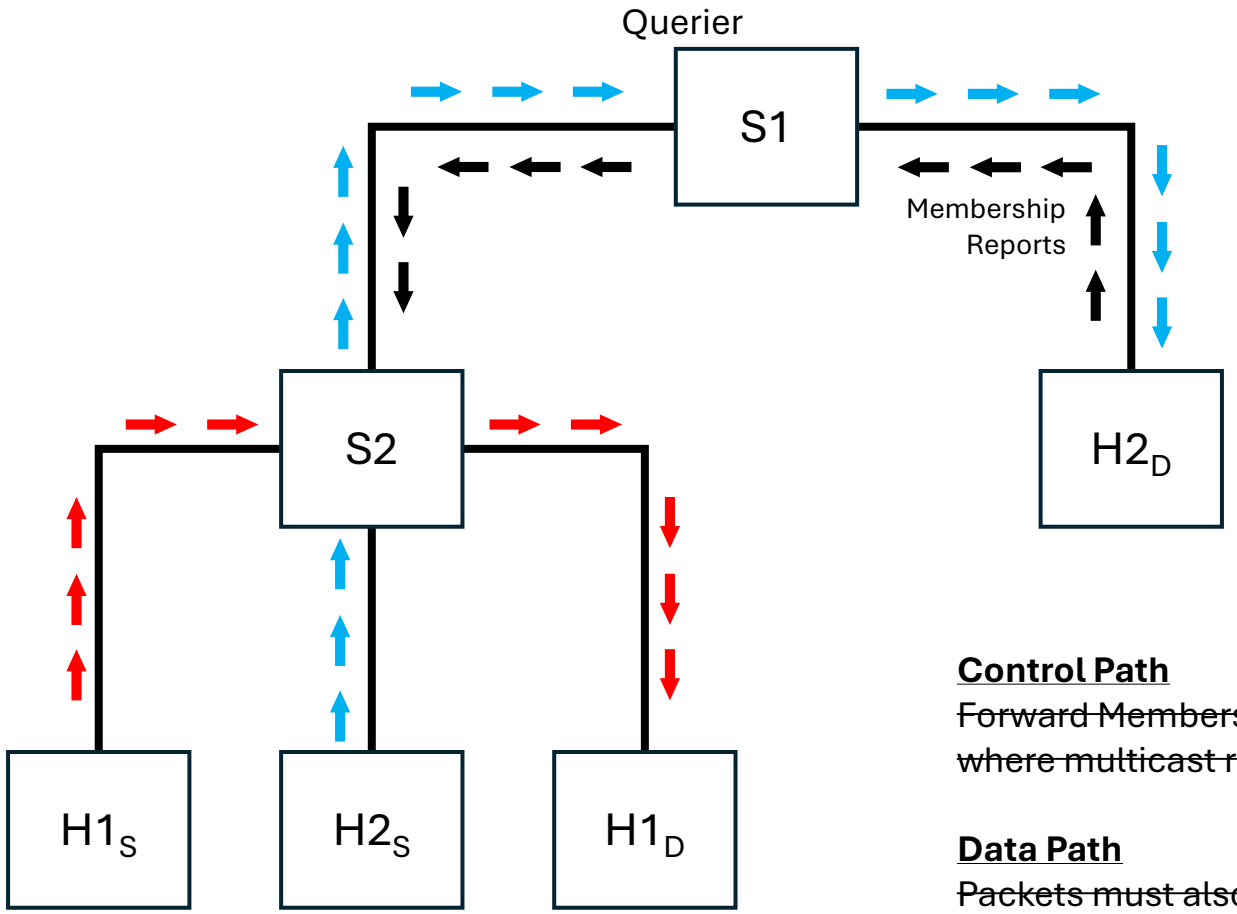

Control Path

~~Forward Membership Reports only to those ports where multicast routers are attached~~

Data Path

~~Packets must also be forwarded on router ports~~

What About Multicast Routing?

- Multicast routers no longer have all data forwarded to them ⇒ would need to take a more active role in requesting routable multicast traffic
- Suggestion: Limit these changes to local network only
  - IPv4 Local Subnet (224.0.0.X): Flood
  - IPv4 Administratively Scoped: Configurable
  - IPv6 All Nodes (ff02::1): Flood
  - IPv6 Link-Local (ffX2): Port-based membership only
  - IPv6 ffX3 through ffXd: Configurable
  - IPv6 Global (ffXe): Always forward to router

Another Problem: Report Suppression

From RFC 4541:

> When an IGMPv1 or IGMPv2 host receives a membership report for a group address that it intends to join, the host will suppress its own membership report for the same group. This join or message suppression is a requirement for IGMPv1 and IGMPv2 hosts.
>
> However, **if a switch does not receive a membership report from the host it will not forward multicast data to it**.

We still encounter IGMPv2 and MLDv1! Example: lwIP

Solution: Do not forward IGMP Membership Reports to ports on which **only hosts** are attached.

Questions?

Co-Authors?

Thank you!