

Multicast Routing in a home network

draft-pfister-homenet-multicast-00

draft-pfister-pim-ssbidir-00

draft-pfister-pim-border-proxy-00

Pierre Pfister

Reminder: What is a 'Home Network'

- Homenet is an IETF Working Group
- Focuses on tomorrow's IPv6 home networks
 - Multiple routers and links
 - Multiple service providers / uplinks
 - End-to-end IPv6 connectivity
 - Zeroconf !
- What does it mean so far ?
 - Unicast routing protocol
 - Configuration protocol (HNCP)
 - Automatic prefix assignment
 - mDNS/DNS-SD proxying
 - What about multicast ?

Check out hnetd implementation:
www.homewrt.org

Presentation Outline

1. Problem statement
2. Solution space
3. A proposal

Why do we want multicast routing ?

- Architecture RFC says so:

It is desirable that, subject to the capacities of devices on certain media types, multicast routing is supported across the homenet, including source-specific multicast (SSM).

- Service Discovery

- UPnP uses site-local multicast
- mDNS/DNS-SD ? -> Not really. Current approach is based on proxying
- Others (SAP announces, vendor specific, etc...)

- Media streaming

- TV over IP
- CCTV
- In-home streaming

- Data sharing (IoT, sensors, etc...)

- Who knows ?

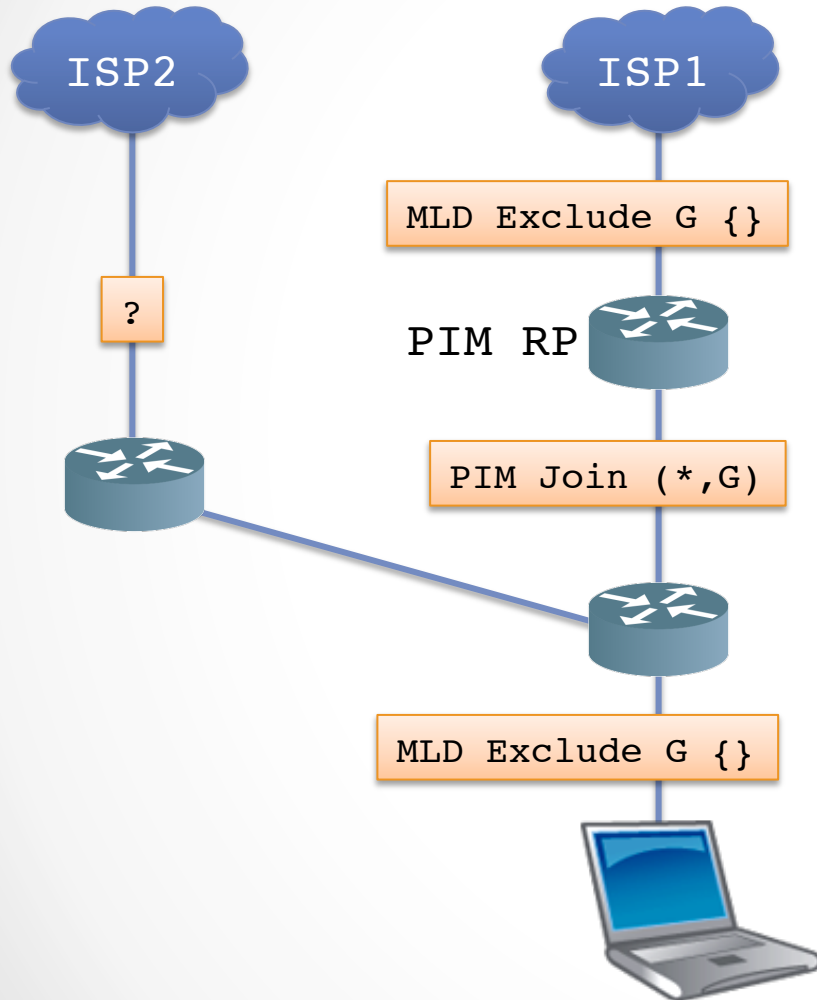
What do we want ?

- Receive ASM or SSM traffic from sources located:
 - On the same link (easy)
 - In the home (We know how to do that)
 - Outside the home (tricky)
 - On different uplinks (even trickier)
- Tricky parts:
 - Multi-homing (multiple default routes)
 - Interfacing with ISP (MLD/IGMP)
 - Zeroconf
 - Keep it Simple

What is the matter ? We have PIM !

- It's not so simple...
- Problem #1: Subscribe to ISP provided traffic.
- Problem #2: Source localization for SSM traffic.

Problem #1: Subscribe to ISP traffic



- PIM reacts to:
 - MLD/IGMP messages.
 - Multicast traffic reception.
- Border router needs to subscribe **first**
 - Needs to know to which group/source.
 - PIM RP at the border will know:
 - ASM subscriptions
 - SSM in singled home network
 - Multi-homing complicates things:
 - No single border (as PIM RP)
- No assumption on ISP interface protocol
 - MLD/IGMP is default.
 - Custom (PIM, ...) is possible.

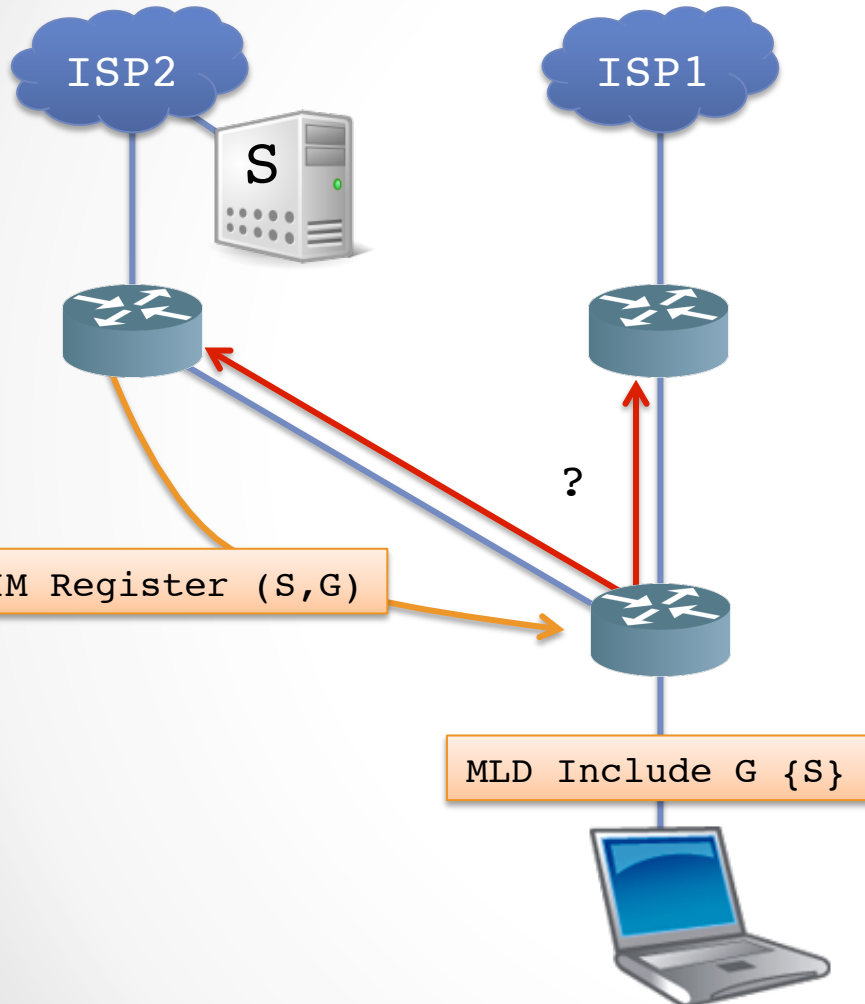
Solution space: Subscribe to ISP

Border routers need knowledge !

- Home-wide ASM and SSM subscription state
- PIM-SM RP does not know SSM state
- For global or multicast groups only

1. Send 'informational' J/Ps to Border Routers
 - Creates lots of state on intermediate routers
2. Single router gathers information and sends it to border routers.
 - 1 peering per border router
3. All routers send local subscriptions to border routers.
 - 1 peering per border router AND per link

Problem #2: Source Localization



- PIM uses the RIB for reverse path forwarding.
 - To the PIM RP (RP tree)
 - To the source (SSM)
- If multiple default routes
 - SSM will not work
 - ASM will be tunneled (PIM-SM)
- Send Join(S,G) to both ?
 - Routing not defined in PIM-SM
 - Possible routing loops

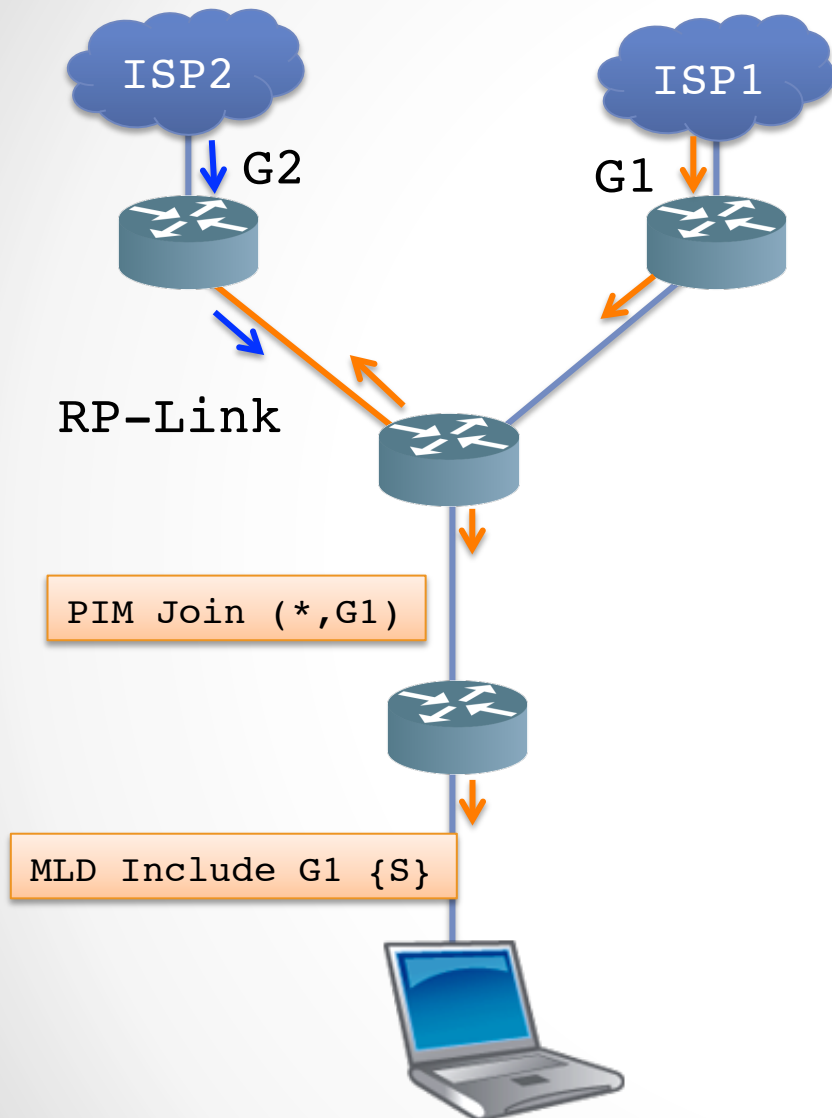
Solution space: Source localization

- RP can locate the source (Register source address)
 - Use RPF Vector to forward SSM to the first-hop-router
draft-ietf-pim-explicit-rpf-vector
 - Other routers could discover source location but first packets would be lost!
draft-ietf-pim-source-discovery-bsr
 - Send Join(S,G) to the RP, but conflicting RPF Vectors have unspecified behavior.
- Use PIM-BIDIR
 - Problem solved 😊 (Always forward traffic to the RPA)
 - This proposal uses PIM-SSBIDIR (draft-pfister-pim-ssbidir-00)
- Use MLD/IGMP proxying
 - Works in a tree. Needs upstream election, or asserts in arbitrary topology.
 - Similar to PIM-SSBIDIR with other messaging.

One way to skin this cat...

1. PIM-SSBIDIR draft-pfister-pim-ssbidir
 - No Source-localization problem
 - Similar to PIM-BIDIR but with Source-Specific state.
2. One single Proxy-Controller, on the RP-Link controls all border proxies draft-pfister-pim-border-proxy
 - Proxy-Controller knows the home-wide subscription state
3. Homenet glue draft-pfister-homenet-multicast
 - Use HNCP for RPA and controller election
 - Use HNCP for border-proxy discovery
 - Send Join/Prunes on the RP-Link

Reminder: PIM-BIDIR



- One single routing tree
 - Routed at the RP-Link
- Always forward toward the RP-Link
 - No need to know where the source is !
- All Source Multicast Only
- Designated-Forwarder Election

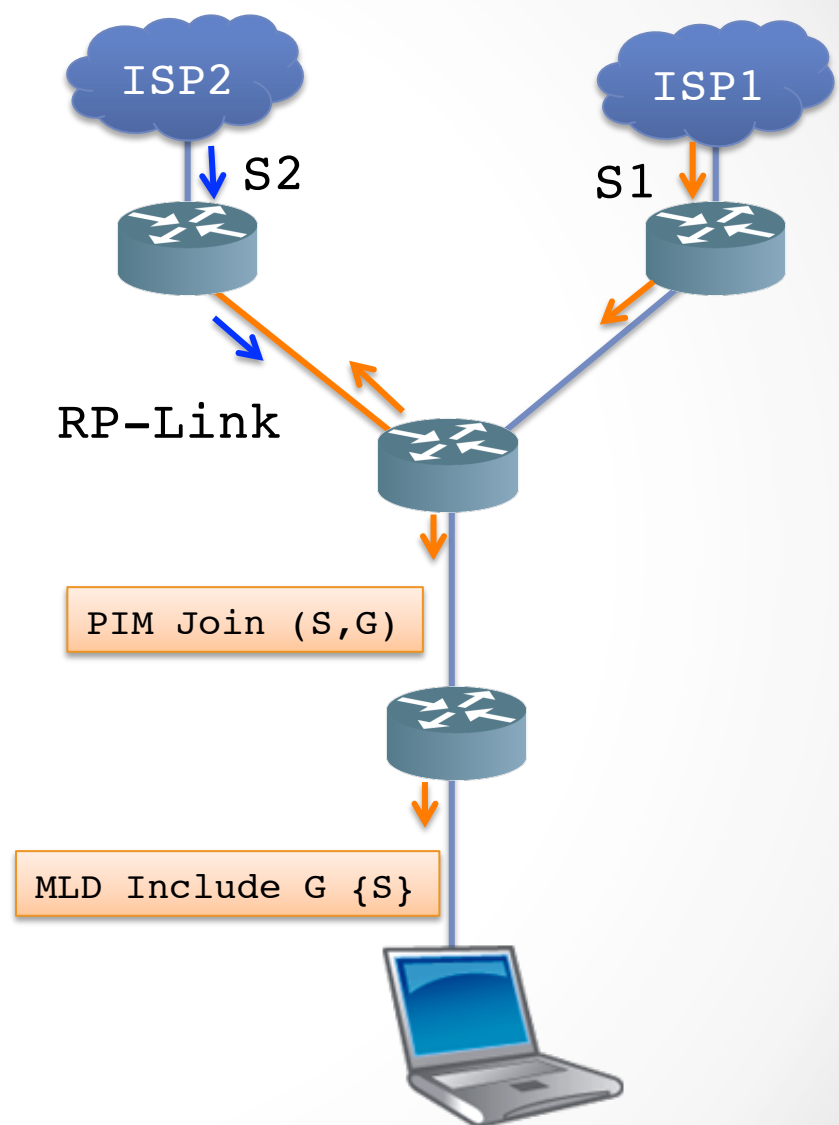
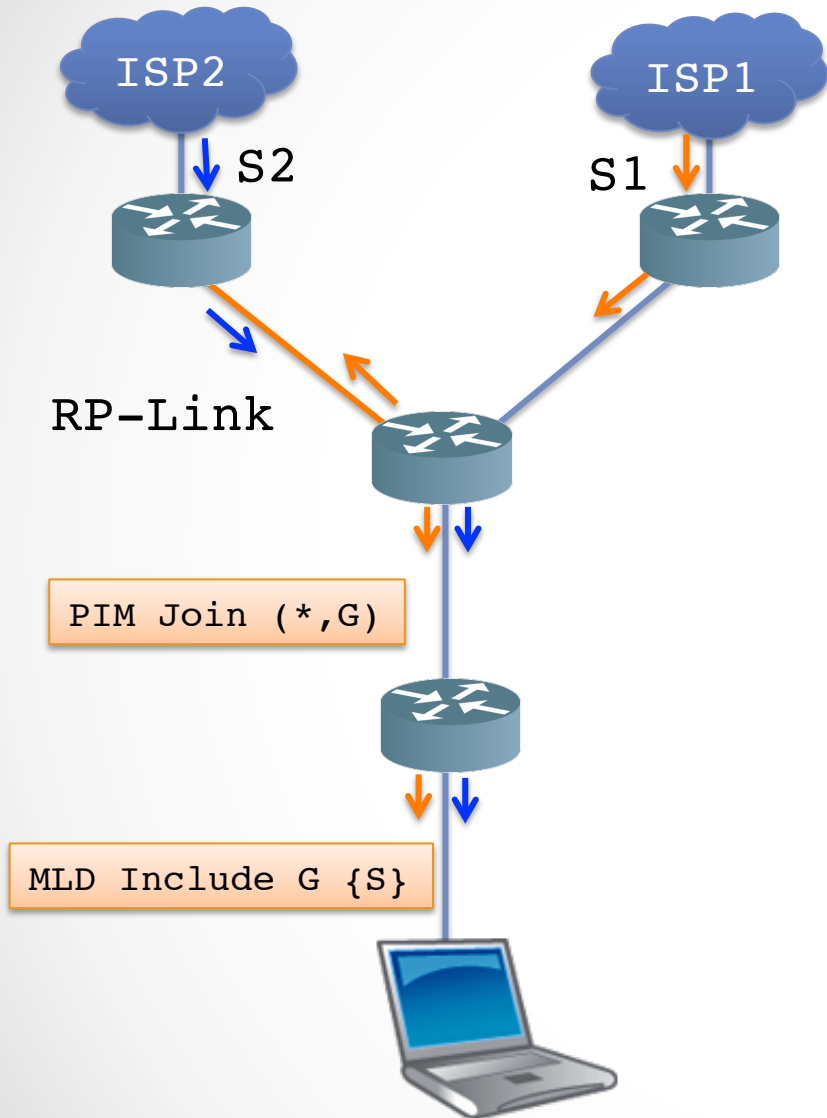
1. PIM-SSBIDIR

- Similar to PIM-BIDIR
 - One single tree is used
 - All traffic is always forwarded up to the RP-Link
 - No Source-Specific Path optimization
- With Source-Specific messaging
 - Join/Prune (S,G) used in 'include' mode
 - (*,G) and (S,G,rpt) used in 'exclude' mode
 - Traffic filtered based on source-specific state
- Backward compatible with PIM-BIDIR
 - PIM-SSBIDIR capable hello option
 - If a neighbor is BIDIR-only -> Do not use (S,G,rpt)
 - If the Designated Forwarder is BIDIR-only -> Do not use (S,G)

BIDIR

Vs

SSBIDIR



PIM-SSBIDIR Pros and Cons

With respect to PIM-SM

- Pros

- No need to locate the source
- Supports Source-Specific filtering
- Provides all subscription state to the RP-Link for free
- Simplicity

- Cons


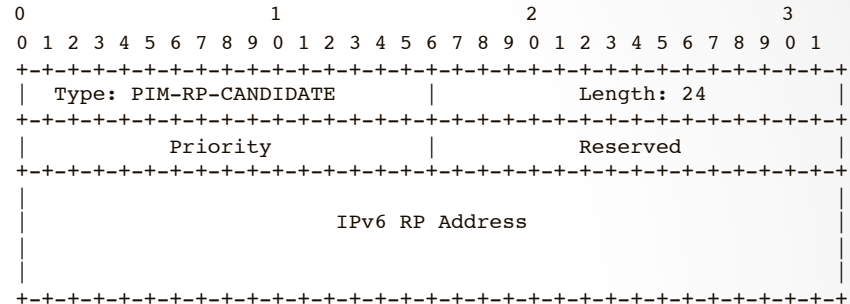
- No Source-Specific path optimization
- Relies on DF-Election (Less robust than asserts ?)

2. PIM Border Proxy - Controller

- PIM state replication
- Based on TCP
 - Reliability (Retries)
 - In order reception (Partial updates)
 - Session loss detection (State creation and destruction)
- Using PIM messaging (Can be used for any state)
- Uplink protocol agnostic (MLD/IGMP, PIM, logging, etc...)
- Multiple controllers and Multiple proxies

3. Homenet Glue

- HNCP is used for electing:
 - The RP Address
 - The Proxy-Controller
- Border routers are discovered using HCNP
- Join/Prune messages are sent on the RP-Link
- The Proxy-Controller sends subscription state to border routers



Summary

- Two homenet specific problems
 - Subscribe to ISP
 - Source localization
- Several possible solutions have been proposed
- One fully-functional proposal (from me)
 - draft-pfister-homenet-multicast-00
 - draft-pfister-pim-ssbidir-00
 - draft-pfister-pim-border-proxy-00
- We are just getting started.
 - Comments or other solutions ?