draft-venaas-pim-ipv4-in-ipv6-core-00

Native IPv4 multicast in IPv6 Core using PIM

Stig Venaas, stig@cisco.com

Mankamana Mishra, mankamis@cisco.com

Salah M. Buraiky, salah.buraiky.1@aramco.com

Overview

- We want to allow forwarding of native IPv4 multicast through a core network that has no IPv4 addresses deployed, only IPv6.
- We show how this can be done using PIM messages with IPv6 headers containing IPv4 group and source addresses.
- We also show how one would do RPF lookup and find IPv6 PIM neighbors for an IPv4 (S,G) or (*,G).

IPv4/v6 network with IPv6 core

- Consider a unicast deployment where core routers don’t have IPv4 addresses, but they still have IPv4 RIB and route and forward IPv4 unicast packets.
- MP-BGP can be used for IPv4 prefix reachability with IPv6 next-hop, see RFC 5549.
- Routers R1 – R5 below all have IPv4 prefixes in the RIB to forward IPv4 unicast.
- Routers R2 – R4 only have IPv6 addresses configured.
- This is sufficient for building IPv4 unicast forwarding tables and native IPv4 unicast traffic can be forwarded.

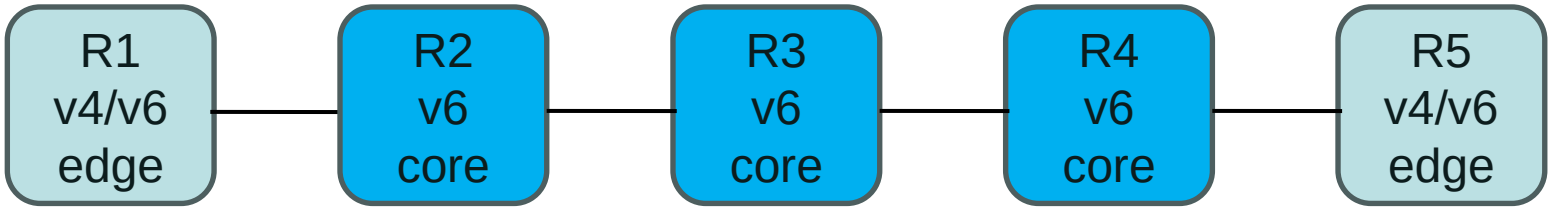

IPv4 multicast with IPv6 core

- Consider how we can do IPv4 multicast in the same network.
- IPv6 pim is enabled and routers are pim neighbors.
- Consider a source IPv4 S and a receiver R joining IPv4 (S,G).
- Routers R2 – R5 below all have IPv4 prefixes in the RIB with IPv6 next-hop, they should all have reachability to S.
- R2 – R5 all do RPF lookup for S and find IPv6 RPF pim neighbors.
- R2 – R5 send pim joins with IPv6 header, IPv6 target and IPv4 (S,G) in payload. This allows IPv4 multicast forwarding state to be created on all routers.

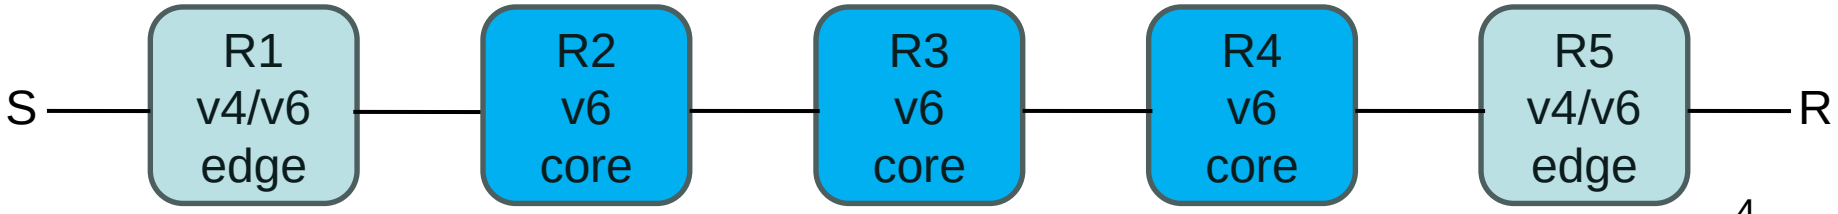

Join/prune format RFC 7761

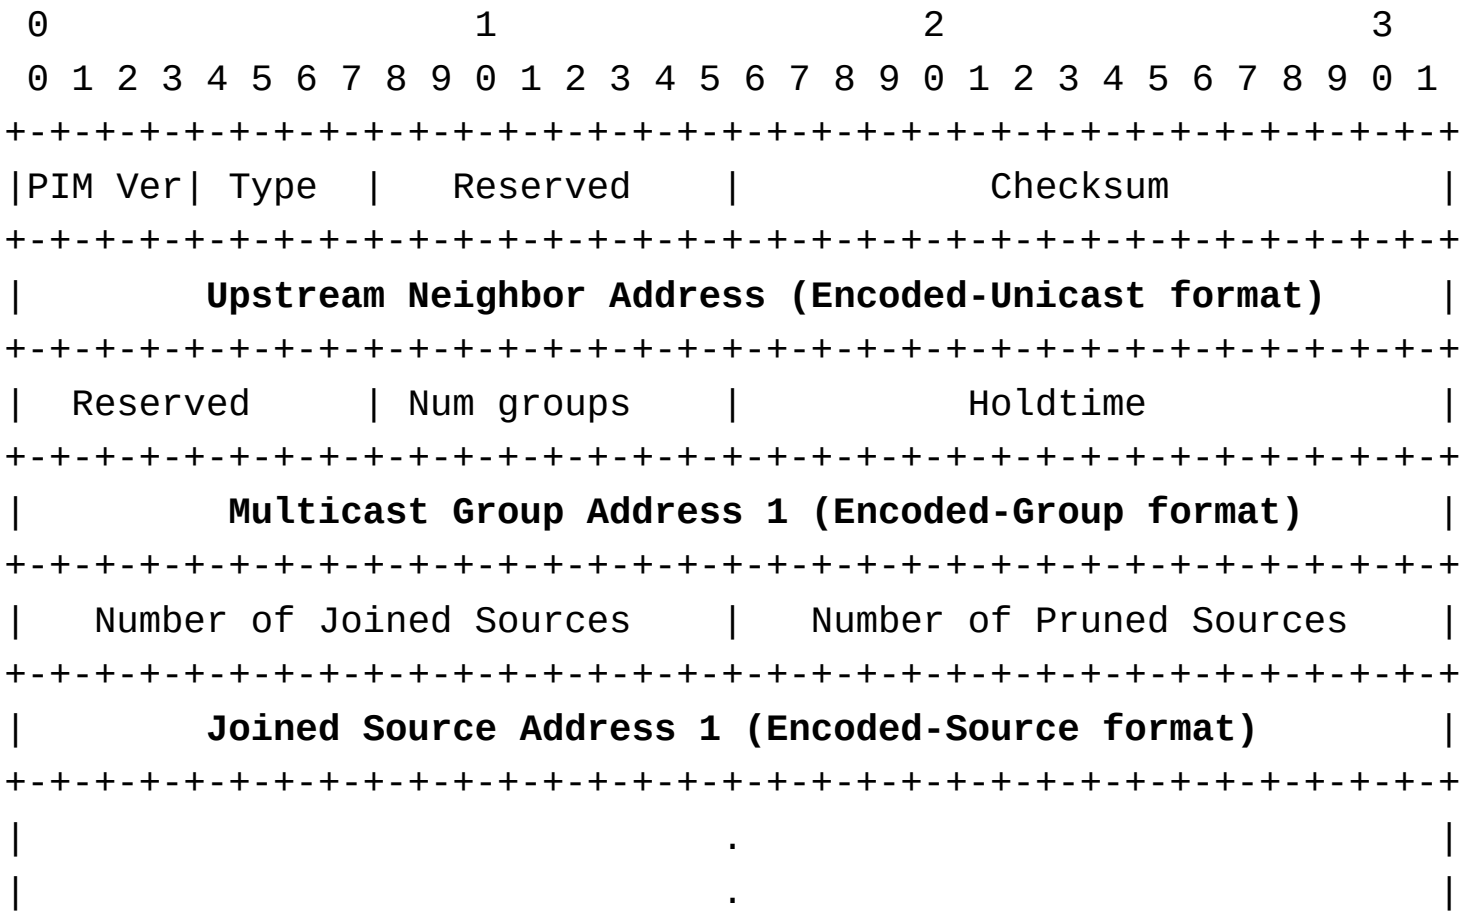

Address family encoding

- Note that all IP addresses are encoded in J/P and other PIM messages. Encoded source format:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Addr Family  | Encoding Type |     Unicast Address
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+...
```

- This allows encoding an IPv4 address in an IPv6 PIM message.
- How a PIM router treats addresses of a different family is unclear. From RFC 7761 section 4.9.5:

Within one PIM Join/Prune message, all the Multicast Group addresses, Joined Source addresses, and Pruned Source addresses MUST be of the same address family. It is NOT PERMITTED to mix IPv4 and IPv6 addresses within the same message. In addition, the address family of the fields in the message SHOULD be the same as the IP source and destination addresses of the packet. This permits maximum implementation flexibility for dual-stack IPv4/IPv6 routers. If a router receives a message with mixed family addresses, it SHOULD only process the addresses that are of the same family as the unicast upstream neighbor address.

Example join message

Standard IPv6 packet header followed by:

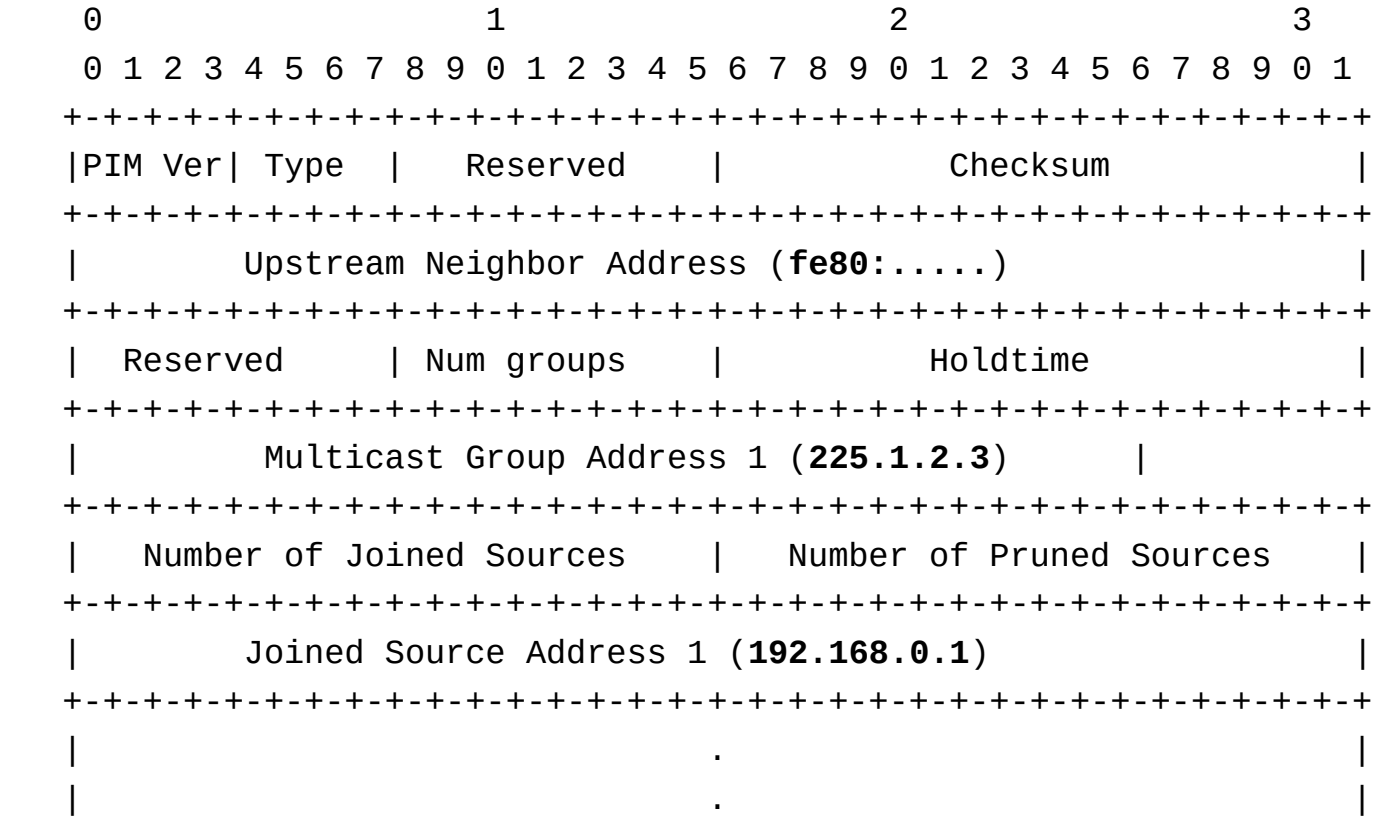

Hello option and other pim messages

- We propose a new hello option to indicate that router is capable of processing PIM messages that have IPv6 header and contain IPv4 source and group information.
- Assert messages may need to have IPv6 header and IPv4 source and group info.
- Propose keeping RP handling out of scope. Assuming RP is located outside of the core and has an IPv4 address. No special handling needed for asserts and BSR.

Avoiding IPv4 RIB in the core

- For IPv4 unicast forwarding core routers need an IPv4 RIB, and we also made use of that for RPF lookup previously.
- We can use RPF vector to avoid IPv4 RIB lookup in the core.
- If R5 determines that R1 is the core edge router to reach S, then R5 can specify the IPv6 address of R1 in an RPF vector in the join it sends upstream.
- R5 as well as R2, R3 and R4, each do an RPF lookup on the IPv6 address of R1 to determine the RPF pim neighbor.

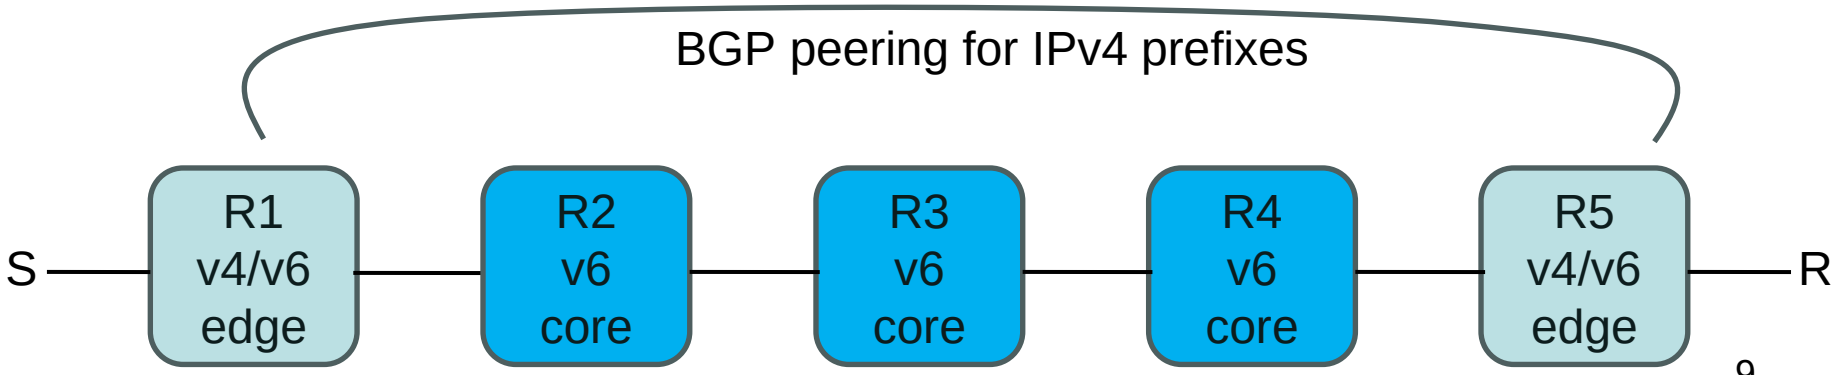

Next steps

- Feedback wanted
- Should the working group work on this?