draft-ietf-pim-ipv4-prefix-over-ipv6-nh-01
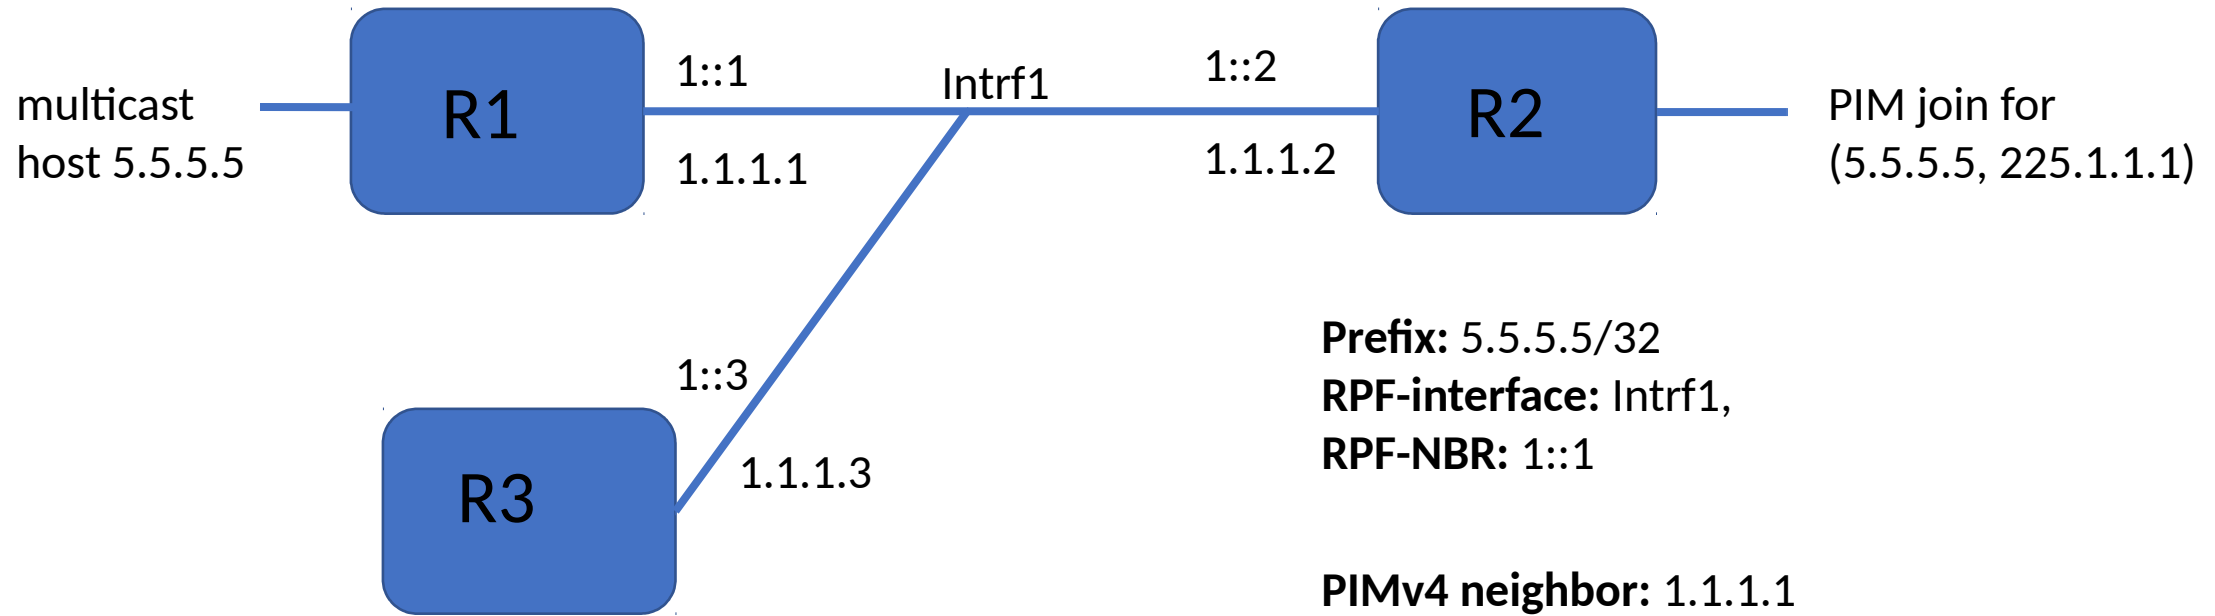
Ashutosh Gupta, ashugupt@cisco.com

Stig Venaas, stig@cisco.com

Example Topology

multicast host 5.5.5.5

R1

1::1

1.1.1.1

Intrf1

1::2

1.1.1.2

R2

PIM join for (5.5.5.5, 225.1.1.1)

1::3

1.1.1.3

R3

Prefix: 5.5.5.5/32
RPF-interface: Intrf1,
RPF-NBR: 1::1

PIMv4 neighbor: 1.1.1.1

Problem Statement

- Multicast routing needs a RPF tree to be formed in order to receive one copy of multicast data on lowest-cost, loop-free path.

- In case of PIMv4, it needs a valid PIMv4 neighbor to send PIMv4 join.

- When using RFC5549, a IPv4 prefix is reachable over IPv6 Next-hop [or vice-versa].

- If RPF-interface has more than 1 PIMv4 neighbor, then a new pim mechanism is needed to choose corresponding neighbor for IPv6 next-hop.

Solution

- Use of secondary address list option in PIM hello.
    As described in RFC 4601 section 4.3.4.
    IPv6 addresses are included in PIMv4 hello as secondary list and
        stored in neighbor cache.
    Downstream router look for IPv6 next-hop address in its local PIMv4
        nbr cache.
    If found, it sends sends join to chosen PIM nbr.

Status

Change in format of draft:

- Earlier it stated prev explained problem and its solution.

- Now It explores the use cases where addresses in secondary address list can be of different AF than primary address.

- PIM with RFC 5549 is one of such use case.

- RFC 4601 says that secondary neighbor address SHOULD be of same AF as primary address. But this draft presents a case where secondary address is from a different AF.