Source + Destination-Based Routing

Jari Arkko, Ericsson
Lorenzo Colitti, Google
Markus Stenberg, Cisco
Mark Townsley, Cisco
David Lamparter, ISC

The Scenario

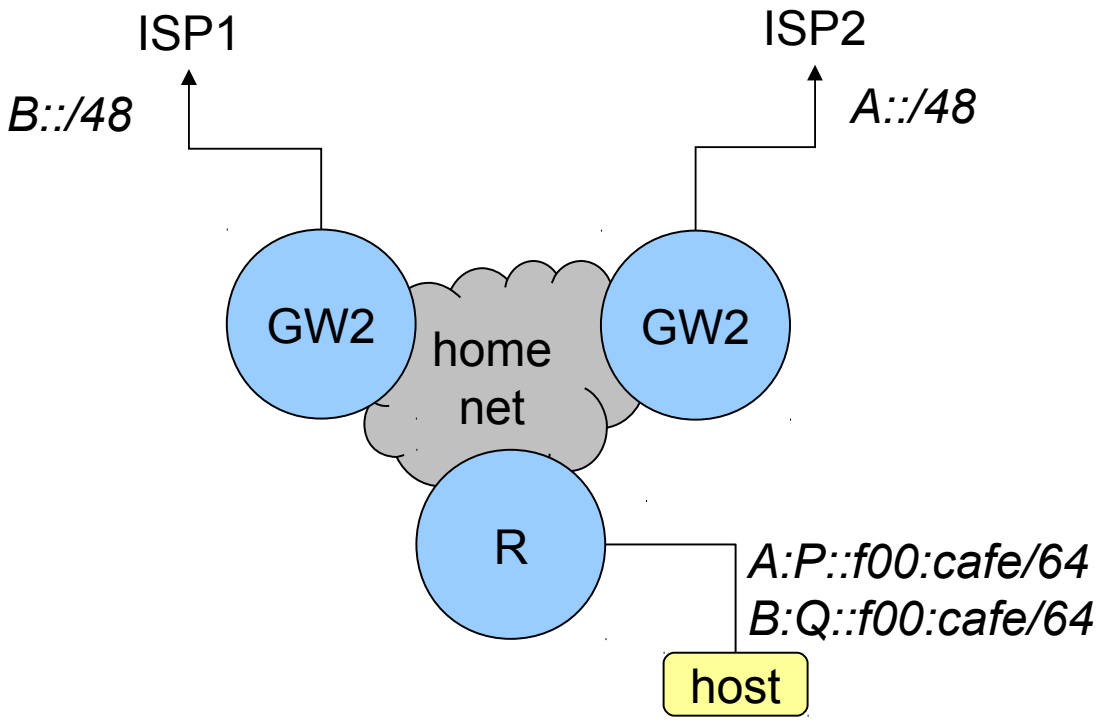

Assume that you have two exit routers from different ISPs, and that both are doing ingress filtering on traffic coming from you

The Scenario

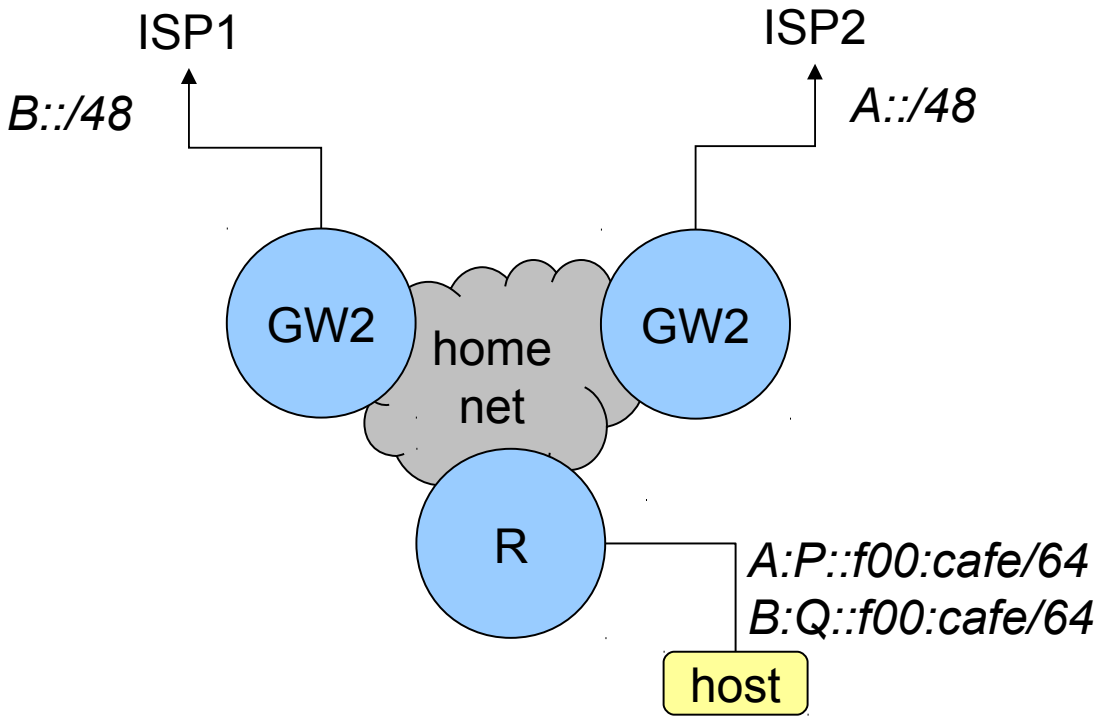

Host picks one source address

If network sends packet wrong way, it gets dropped

How to ensure that packet goes out to the right ISP?

Goals

- Ensure that the packet is not dropped
- No modifications to hosts
  - Including source address selection

- Non-goals:
  - Connection survivability across ISP changes
  - Policy routing or controls
  - Solving the multi-interface host problem

Some Questions...

- Can we do it?
  - Yes.
- Which of the many approaches to use?
- Do we want to do it?
  - Would help solve an important use case
  - May help avoid NAT66 in some cases

Solution Approaches

1) The border routers forward packets to each other if the source address requires it

   [only works if the routers are adjacent]

2) The border routers tunnel packets to each other

   An MTU and performance impat

3) Source + destination-based routing

Source + Destination-Based Routing

- Assume that all ISP links always filter
- Distribute information about what source addresses are legal on those links
  - "Acceptable source addresses"
- Compute routing tables such that for external routing entries, source address matters too
  - Default, VPN, ...
- If ISP goes down, deprecate its addresses
- Fits very well with prefix-autoconfig

Implementing Source + Destination -Based Routing

- Forwarding support – Linux kernels have supported multiple routing tables for a long time
  - Dst = usable prefix => use OSPF route table (1st prio)
  - Src = usable prefix => use route to GWs (2nd prio)
- Markus has an implementation that distributes information over OSPFv3 and allows multiple gateways to be used

Open questions

Metrics?

Which protocol to use?

What LSA to use (if any)?

Should prefix used for autoconfig always be the same as acceptable source address?