OSPFv3-Based Home Networking – Report

Jari Arkko, Ericsson
Markus Stenberg, TMS/Cisco
Acee Lindem, Ericsson
Lorenzo Colitti, Google
Mark Townsley, Cisco
David Lamparter, ISC

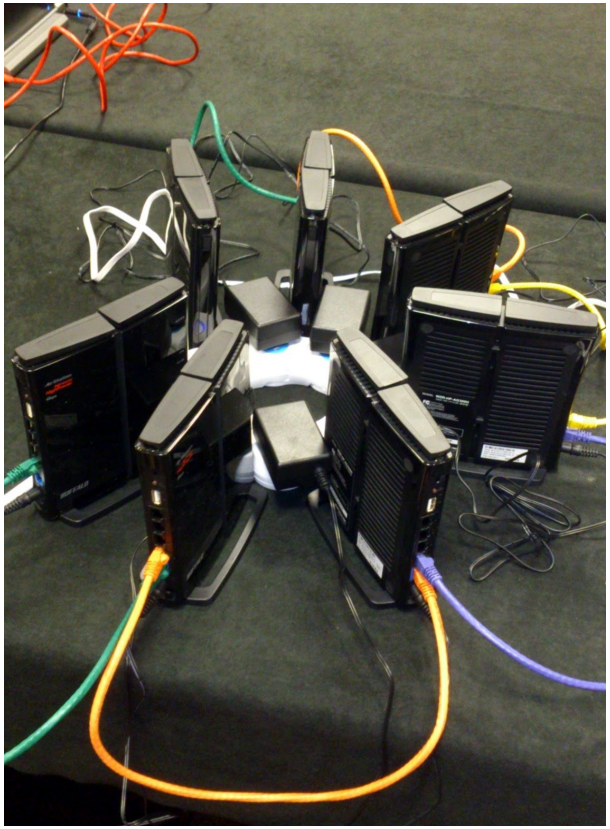

Based on implementations, and draft-ietf-ospf-ospfv3-autoconfig-00.txt and draft-arkko-homenet-prefix-assignment-03.txt

Outline

- Implementation & interop report
- What we learned
- Source + destination -based routing

Implementation Report

There are now three implementations!

- Including advanced features, such as DNS discovery, ULA-generation, source-based routing, and so on

Interop Report

- We have been testing our implementations here this week
  - Much of this is still getting the implementations to do the right thing on their own
  - But on complex setups, such as with 8 routers, 2 exit routers, and a number of hosts
  - And running quite advanced functionality
- There was also some interoperability testing
- Testing space organized by IPSO Alliance – thank you IPSO and Geoff Mulligan!

Interop Observations

- The system works!
- The protocols overall seem to be OK
- First interoperability achieved during WG session on Wednesday
- Brittle timer defaults found to be problematic
- Many implementation issues, byte-order, etc.
- OSPF prefix compression rule from RFC 2328 found to be unclear

Interop Observations 2

Across-the-subnets DNS support:

- Zero-config DNS support within a home (mDNS)
- But also zero-configuration to give hosts access to the Internet DNS
  - Requirement for either discovering DNS servers or providing a default server
  - Ensures that hosts can resolve DNS names
  - Current implementation: each router makes a decision on its own – could also coordinate

Interop Observations 3

- The routing daemon becomes connected to many other things – DHCP, RAs, DNS, …
  - Increases complexity
  - Implementation choices include being integrated, started via routing daemon, IPC interfaces, …
  - Many commonly available components are not so well suited to be included in the above (too integrated, too big, missing some functionality)
- Some issues in other components
  - For instance, RADVD timer defaults are too long

Things to Consider – Brittle Timer Defaults

- OSPF RFCs give sample values *(HI=10s, RDI=40)*
- The autoconfig draft says these MUST be used
  - Unfortunately, implementations have (a) widely varying default values *(9..20s, 40..120s)* and (b) are universally picky with deviation
  - As a result, autoconfiguration fails
- Wait for implementations to heed to the draft?
- Or support dynamically agreed values?
  - One implementation can adjust itself to values advertised in peer's hello messages

Source + Destination-Based Routing

The Scenario

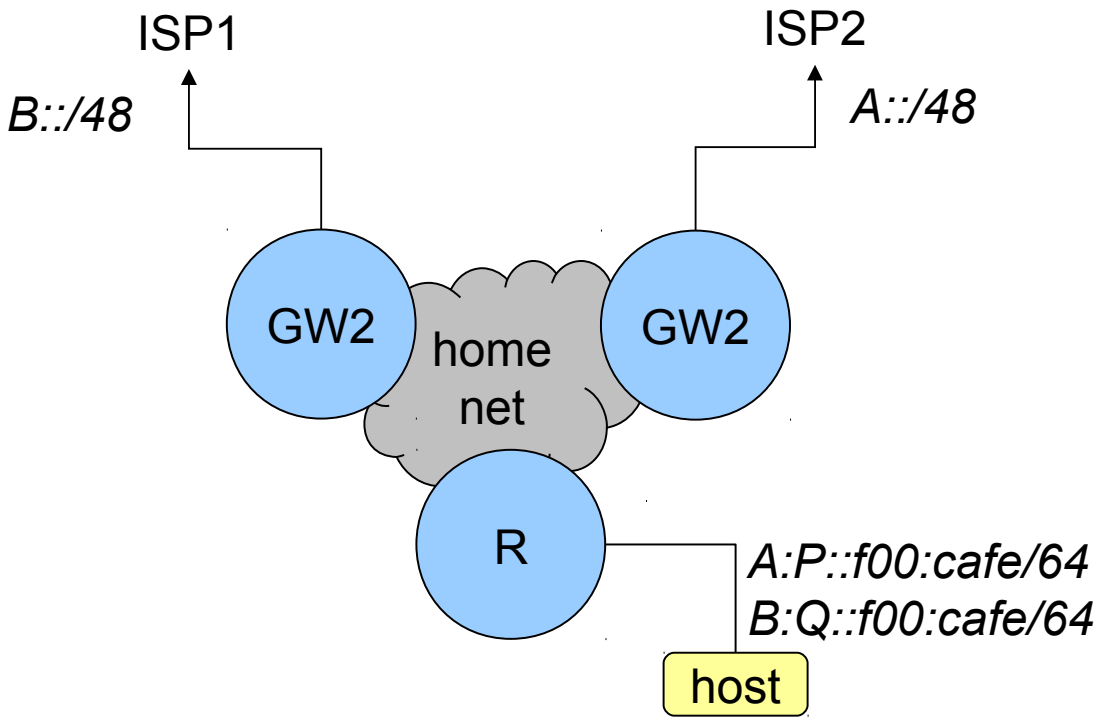

Assume that you have two exit routers from different ISPs, and that both are doing ingress filtering on traffic coming from you

Goals for an "Ingress Filtering Avoidance Solution"

- Ensure that the packet is not dropped
- No modifications to hosts

- Non-goals:
  - Connection survivability across ISP changes (orthogonal to any mobility, locator-id, and happy eyeballs solutions)
  - Policy routing or controls
  - Solving the multi-interface host problem

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, …
- Linux support, one implementation on OSPFv3