Should 6man do something about source address dependent routing for IPv6 hosts?

Brian Carpenter

IETF 93
July 2015

Scope

- Host problem statement & requirements only
  - NOT routing beyond the host
  - NOT solution details
  - ESPECIALLY NOT dhcpv6 vs RA
- The goal is to figure out whether 6man needs to work on this.

Previous work

- Discussions in various WGs, as far back as multi6, and including mif, homenet and v6ops (especially RFC 7157)
- Solutions drafts (off the table today)
- Routing aspects (off the table today, but draft-baker-rtgwg-src-dst-routing-use-cases is helpful)
- Analysis draft (draft-sarikaya-6man-sadr-overview)

Problem space

- A host with more than one global unicast address wants to send packets to a remote target with more than one global unicast address.
- The sending host's ISPs operate BCP 38 ingress filters.

General scenario 1: Single routing cloud

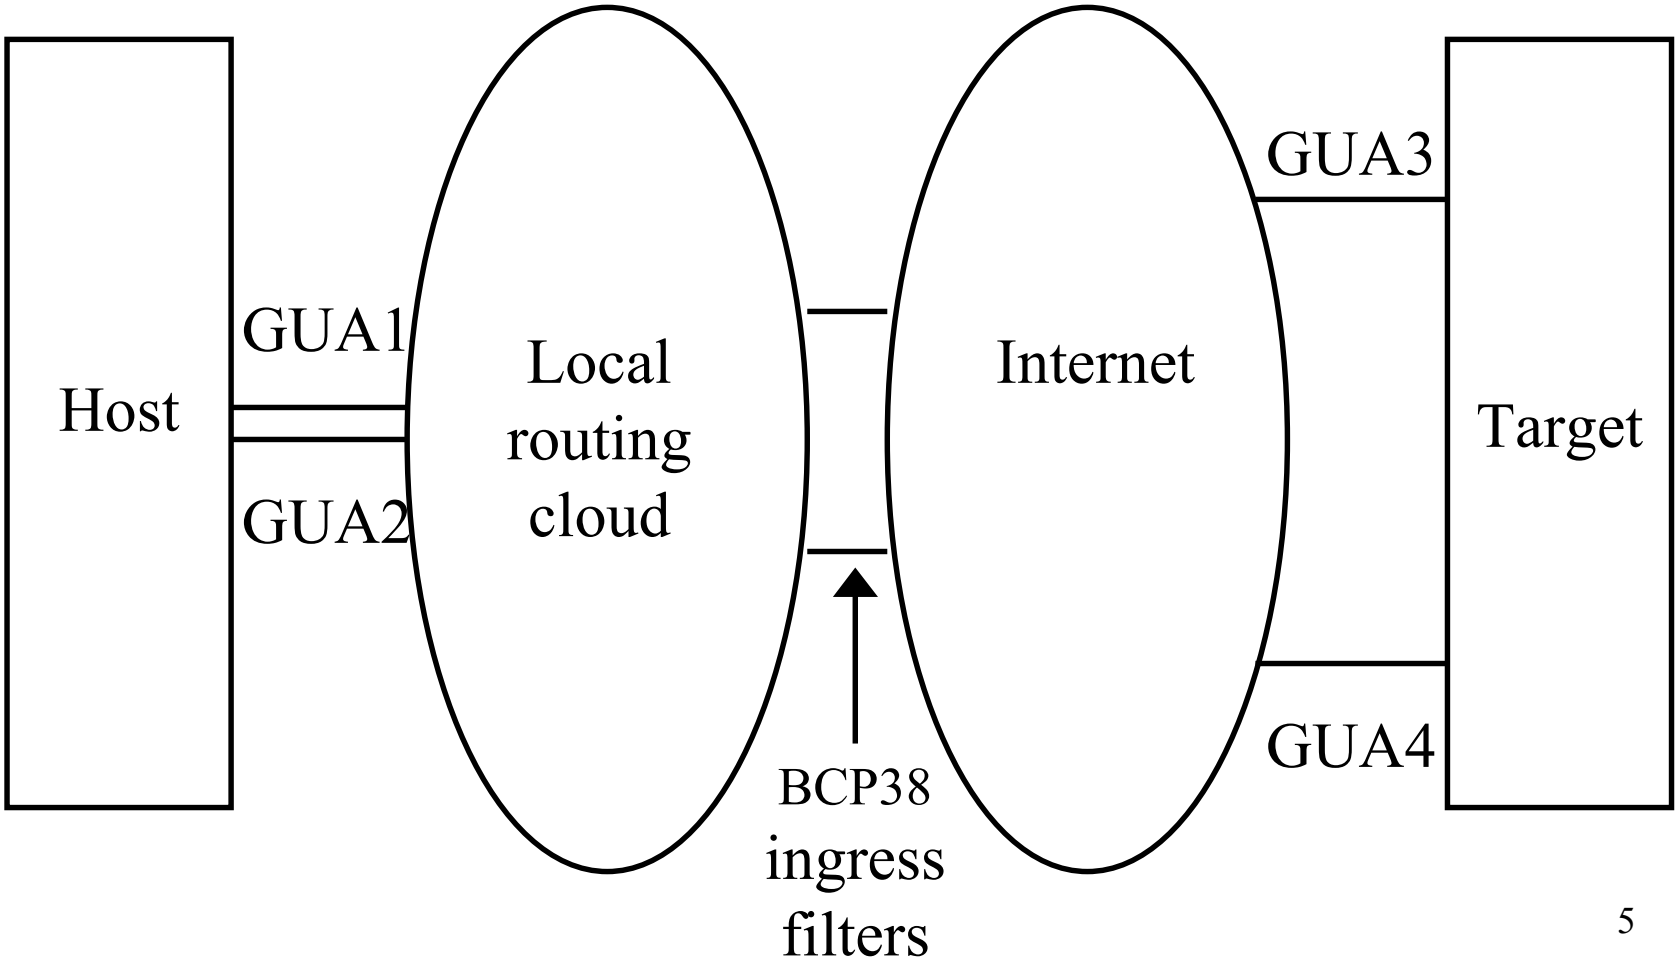

General scenario 2: Disjoint local routing

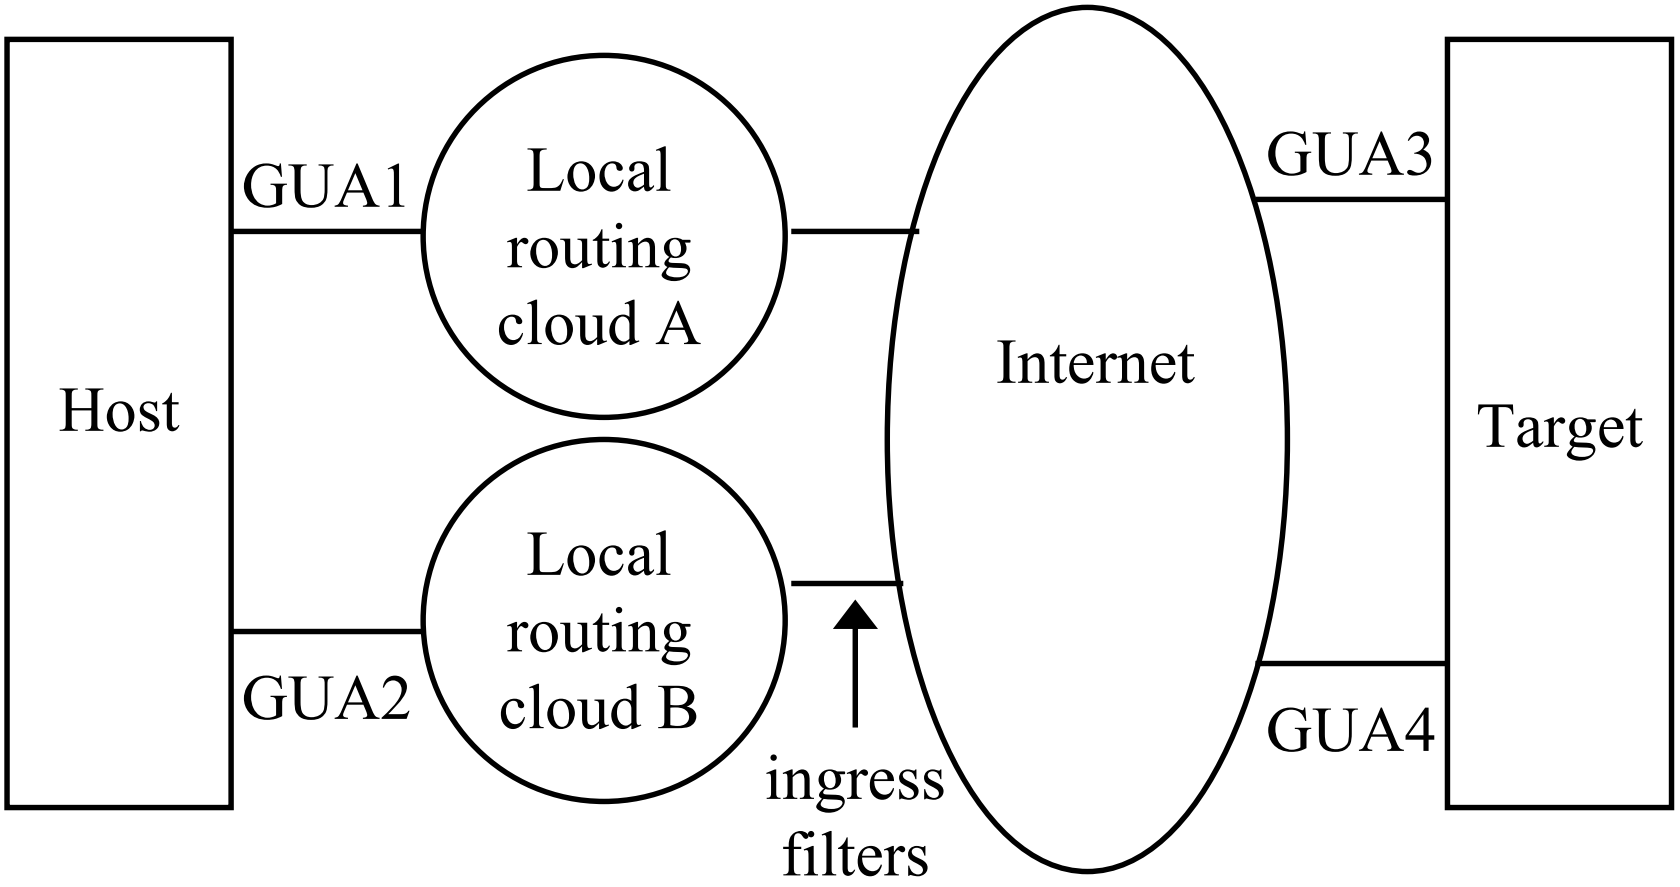

General scenario 2: Disjoint local routing

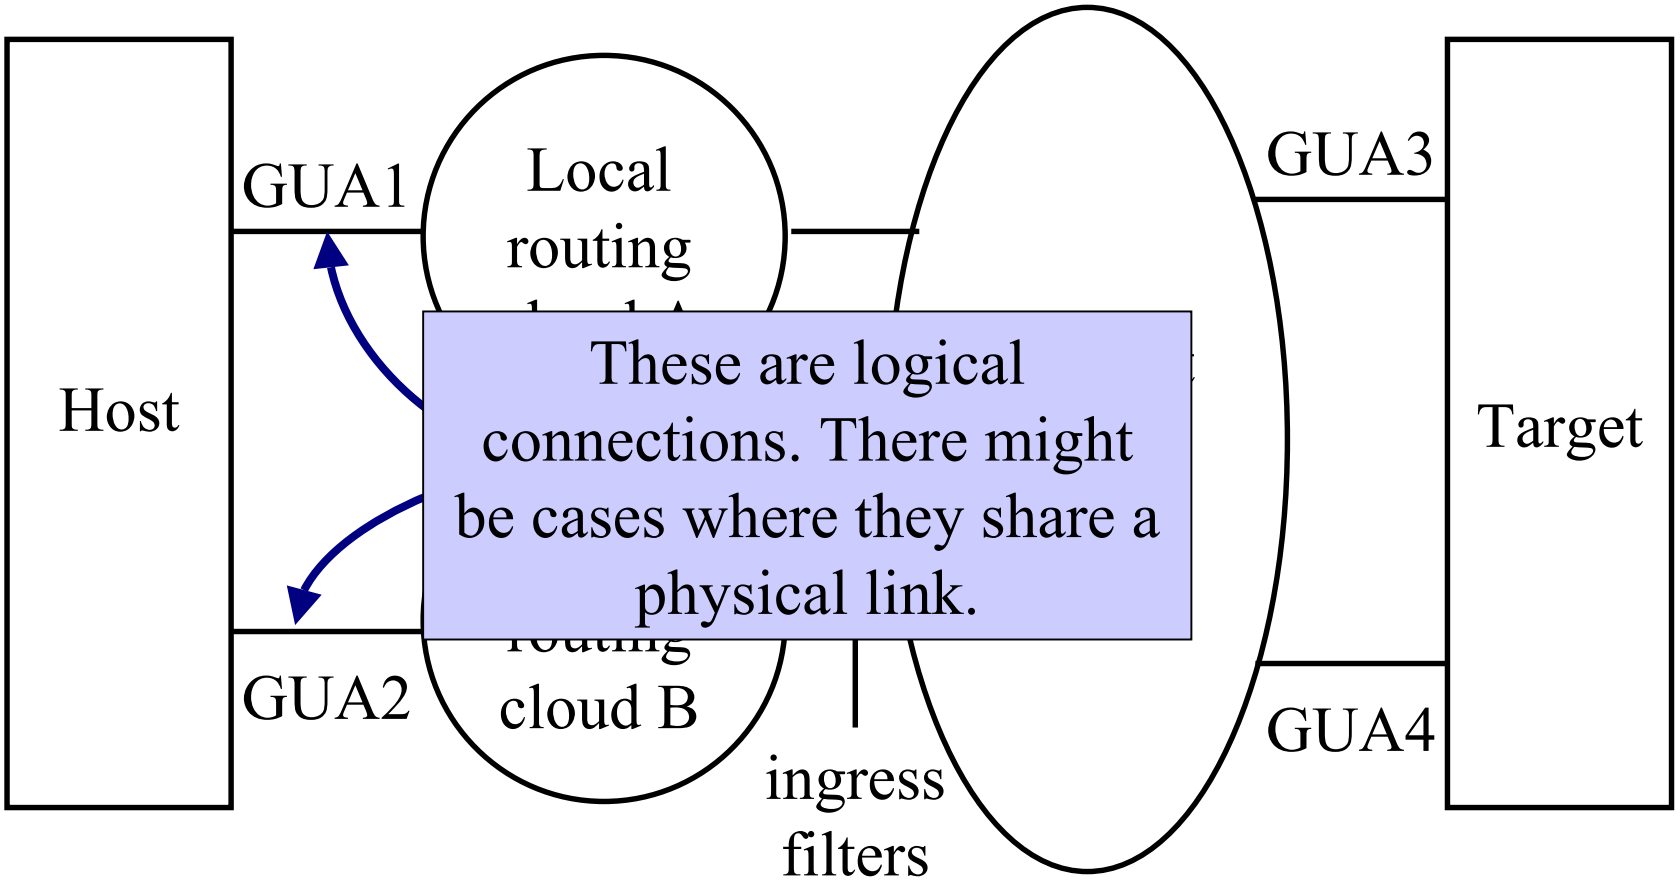

Failure cases

- A failure case is when a packet from (e.g.) GUA1 to GUA4 leaves the local network via a link to an ISP that filters and discards packets from the GUA1 prefix.
- If all such cases can be prevented by an existing mechanism, nothing new is strictly needed.

Not considered failures in the host stack

- Suboptimal routing
  - selecting a suboptimal first hop router is not a failure.
- Suboptimal performance
  - selecting a suboptimal address pair is not a stack failure; it's an application issue.

Scenario 1

- If the routing cloud is fully connected & capable of source and destination address based routing to select a valid exit, failure will not occur.

Scenario 2

- A failure will occur if a packet with source address GUA1 is sent from the interface that GUA2 belongs to.
  - Such address pair selection will not occur if rule 5.5 in RFC 6724 is applied.
  - This would need all stacks to remember which next-hops advertised which prefixes.
  - Also needs stacks to apply redirects per source prefix.

Personal conclusion

- IMHO these things (only) are strictly needed:
  - REQUIRE that all stacks remember which next-hops advertised which prefixes, to enable rule 5.5.
  - Small fix to redirect rules in RFC 4861
  - REQUIRE that RA/PIOs are sent and processed on all links (even with DHCPv6).
  - RECOMMEND that all routers in edge networks support source/destination routing throughout, configured to satisfy BCP38 filters (see draft-baker...).

Discussion