PA Address Multihoming and Source/Destination Routing

Fred Baker

draft-baker-rtgwg-src-dst-routing-use-cases

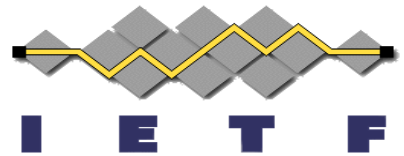

This talk triggered by operator request in v6ops

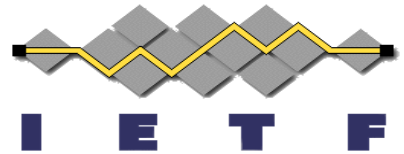

- A number of operators opined at IETF 94 that,
  - while PI multihoming is common and works well,
  - PA is difficult for enterprise to deploy without egress routing
- Those few networks using it resort to operational means such as
  - Flash renumbering
  - Using one ISP’s prefix in one place and another ISP’s in another
  - Forcing all traffic through a single egress router

Discussion of use cases

- General comment:
  - The chairs asked me to comment on source/destination routing in the context of PA Address Multihoming
    - This is a special case, in which the network routes toward a network egress appropriate to a source address
    - Source/Destination routing has other uses as well
- My biggest concern is that by focusing on a specialized (although common) use case, the tool will be limited in value
  - Other uses of source/destination routing could be described as an ACL embedded in routing

To give you an idea

- draft-xu-ospf-multi-homing-ipv6 uses same concepts as
  - draft-ietf-rtgwg-dst-src-routing
  - draft-baker-ipv6-isis-dst-src-routing
  - e.g., routing to a remote router advertising a source/destination LSA
- Currently deployed in CERNET2
  - Traffic engineering for three universities without MPLS
  - Load balancing application
  - (not egress routing, not homenet)
- Three vendors:
  - Huawei
  - ZTE
  - Bitway

Egress Routing impetus

- IETF generally recommends* use of provider-allocated prefixes in generalized multihoming for smaller networks
  - PI obviously works and is used by larger networks that use BGP and have AS numbers
  - The point is to minimize impact on the global route table by enabling ISPs to aggregate smaller multihomed customers into their own prefix
- Issue:
  - BCP 38 encourages ISPs to drop customer traffic that uses addresses they don’t know the customer to be using

History

- This came to a head in the IETF in 2004, when v6ops WG Chair asked me to write up a solution
  - RFC 3704
- Concept:
  - Destination route within a network
  - At the egress, wonder what source prefix is in use
    - If the correct one for upstream, send upstream
    - Else, re-route to the correct egress router
- My question:
  - Why not route it to the right router in the first place?

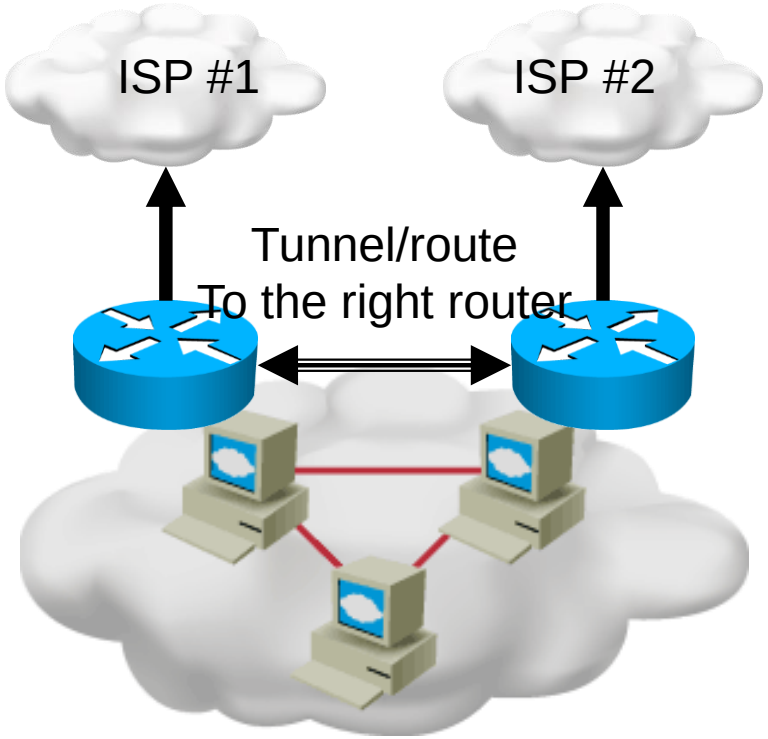

First use case: egress routing

- Which is routing from a prefix to ::/0 (default route)
  - Destination or ::/0=>Destination route within the enterprise network

Second use case: egress routing with a specialized external route

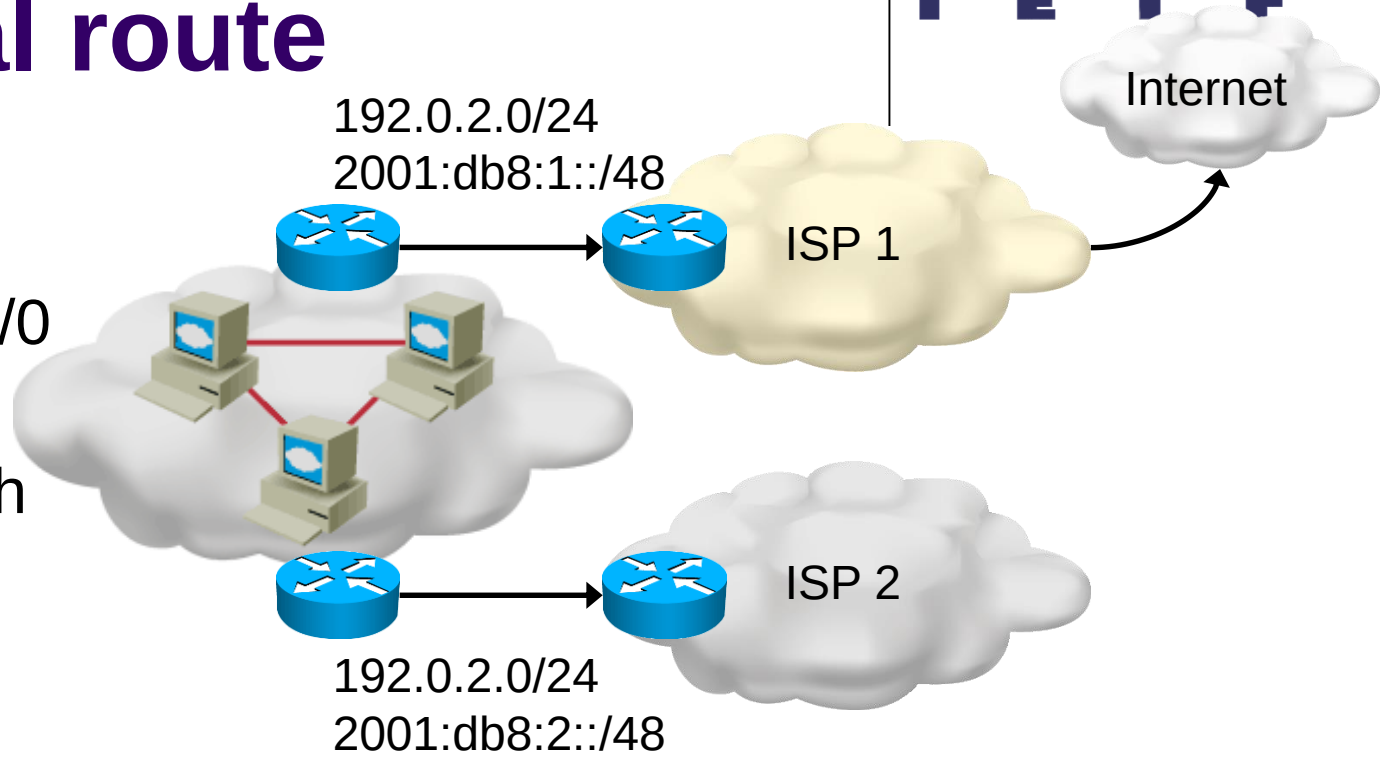

- Multiple ISPs
  - ISP 1: routing from a prefix to ::/0 (default route)
  - ISP 2: Specialized service (such as NTT BFLETS)
- Specialized ISP offers a destination route to its prefix, and requires network (home) to use its PA prefix when accessing it.

- Yes, you could use destination routing and let hosts learn which source address actually works. If they actually learn.