

IPV6 Destination/Source Routing

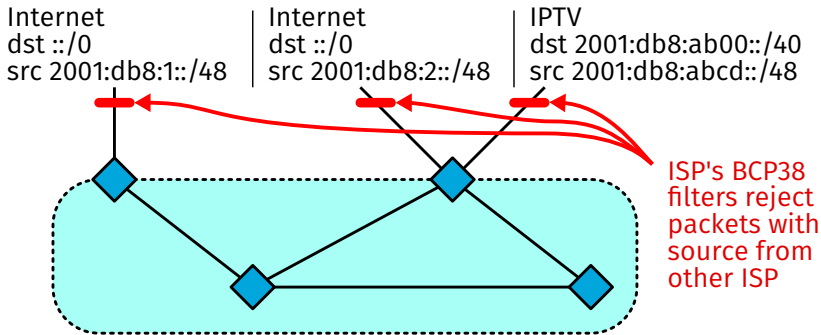
draft-lamparter-rtgwg-dst-src-routing-01

David Lamparter

IETF 93, July 2015, Prague

Context refresher

- ▶ homenet multiple uplinks vs. BCP 38 filtering
cf. draft-baker-rtgwg-src-dst-routing-use-cases
(IETF88)
- ▶ draft-baker-ipv6-ospf-dst-src-routing (exp) &
draft-baker-ipv6-isis-dst-src-routing
overlap in describing forwarding behaviour
- ▶ implementations exists (HNCP, BABEL & IS-IS
control planes, Linux kernel forwarding)



Changes since -00

Dropped “extra qualifiers” stuff.

No longer trying to create a generic framework for people changing the longest prefix match function.

Simply adding source LPM “after” dest LPM.

No more flowlabels in LPM

Lookup behaviour

A route with longer destination prefix match is always more specific than any other route with shorter destination prefix match, regardless of any source prefixes.

Only between the same destination prefix, source prefixes are longest-matched.

Route that doesn't match both is not a match.

Lookup behaviour

draft specifies continuing to less specific destination matches if no entry produces a source match (i.e. modeled as one integral lookup process, not 2 separate steps)

This is noted as general principle – stopping lookup can always be done by inserting an unreachable or blackhole route.

Lookup behaviour

packet to 2001:db8:1234:5678::1

from 2001:db8:ef::1

less specific
↓

- | | | |
|----------|---|--------------------------------|
| 1 | dst <u>2001:db8:1234:5678::/64</u> | src <u>2001:db8:ab:cd::/64</u> |
| 2 | dst <u>2001:db8:1234:5678::/64</u> | src <u>2001:db8:cd::/48</u> |
| | <i>↓ lookup progresses to next dst prefix</i> | |
| 3 | dst <u>2001:db8:1234::/48</u> | src <u>2001:db8:ab:cd::/64</u> |
| 4 | dst <u>2001:db8:1234::/48</u> | src <u>2001:db8:ef::/48</u> |
| | dst 2001:db8:1234::/48 | src ::/0 |

Not allowed to give up after step #2

Lookup behaviour

packet to 2001:db8:1234:5678::1

from 2001:db8:ef::1

less specific
↓

- | | | |
|----------|------------------------------------|--------------------------------|
| 1 | dst <u>2001:db8:1234:5678::/64</u> | src <u>2001:db8:ab:cd::/64</u> |
| 2 | dst <u>2001:db8:1234:5678::/64</u> | src <u>2001:db8:cd::/48</u> |
| 3 | dst <u>2001:db8:1234:5678::/64</u> | src <u>::/0</u> unreachable |
| | dst 2001:db8:1234::/48 | src 2001:db8:ab:cd::/64 |
| | dst 2001:db8:1234::/48 | src 2001:db8:ef::/48 |
| | dst 2001:db8:1234::/48 | src ::/0 |

Route #3 must be explicitly installed if “give up” desired

Open topics: recursive routes

Most dst-src work happened in homenet – little concern given to interop with non-homenet.

- ▶ Recursive routes where nexthop matches D/S route
 - ▶ 2001:db8::/32 via 2001:db8:abcd::1 recursive
 - ▶ 2001:db8:abcd::/48 src 2001:db8:1234::/48 via A
 - ▶ 2001:db8:abcd::/48 src 2001:db8:5678::/48 via B
- ▶ questionable relevance?
- ▶ multiple routes installed?

Other integration concerns

Unicast RPF:

- ▶ Filtering incoming packets based on route lookup with dst and src reversed
- ▶ previously: only check packet src \Leftrightarrow route dst
- ▶ draft says: allowed to ignore packet dst, or also check packet dst \Leftrightarrow route src

Multicast RPF:

- ▶ MRPF only ever uses multicast sender address
- ▶ draft says: “ignore D/S routes, not applicable”
- ▶ proper solution: separate multicast topology

Next steps

- ▶ looking to get this adopted as rtgwg WG document
- ▶ homenet WG needs this
 - ▶ independent of / applies to all choices of routing protocol
- ▶ have personally seen use case in service provider network (user class \Rightarrow BGP peers mapping)