

Extending IP Route Lookup & Dst/Src Routing

draft-baker-rtgwg-src-dst-routing-use-cases-02
draft-lamparter-rtgwg-routing-extra-qualifiers-00
draft-lamparter-rtgwg-dst-src-routing-00

Fred Baker
David Lamparter

IETF 91, Honolulu

Context

- ▶ homenet multiple uplinks vs. BCP 38 filtering
cf. draft-baker-rtgwg-src-dst-routing-use-cases
(IETF88)
- ▶ draft-baker-ipv6-ospf-dst-src-routing &
draft-baker-ipv6-isis-dst-src-routing
overlap in describing forwarding behaviour
- ▶ implementations exists (HNCP or IS-IS control
plane, Linux kernel forwarding)
- ▶ need to document to ensure clarity / interop

Untangling things

draft-lamparter-rtgwg-routing-extra-qualifiers:

- ▶ order-only IANA registry for lookup keys
- ▶ table has 1 initial entry: “destination LPM”
 - ▶ future drafts (e.g. flowlabel routing) may add entries in any place (front, middle, behind)

draft-lamparter-rtgwg-dst-src-routing:

- ▶ adds “source LPM” behind destination LPM
 - ▶ this is not policy routing

Lookup behaviour

drafts specify 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 2001:db8:1234:5678::/64 src 2001:db8:ab:cd::/64

2 dst 2001:db8:1234:5678::/64 src 2001:db8:cd::/48



lookup progresses to next dst prefix

3 dst 2001:db8:1234::/48 src 2001:db8:ab:cd::/64

4 dst 2001:db8:1234::/48 src 2001:db8:ef::/48

dst 2001:db8:1234::/48 src ::/0

Lookup behaviour

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

from 2001:db8:ef::1

less specific
↓

1 dst 2001:db8:1234:5678::/64 src 2001:db8:ab:cd::/64

2 dst 2001:db8:1234:5678::/64 src 2001:db8:cd::/48

3 *dst 2001:db8:1234:5678::/64 src ::/0 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

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?

Open topics: Unicast RPF checks

- ▶ Filtering incoming packets based on routes with dst/src reversed
- ▶ draft: “flip source too, or don’t check it”
 - ▶ other new qualifiers (e.g. flowlabel) may need more thought

Open topics: Multicast RPF lookup

- ▶ no source address available when doing MRIB lookups
- ▶ draft currently says
“completely ignore D/S routes”
- ▶ may want pim WG feedback

Further path?

- ▶ WG-adopt use-case document?
- ▶ proceed with routing-extra-qualifiers & dst-src-routing in rtgwg?
 - ▶ this is on the charter borderline (extending hop-by-hop routing)