rtgwg
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2015

D. Lamparter
NetDEF

October 20, 2014

Destination/Source Routing
draft-lamparter-rtgwg-dst-src-routing-00

Abstract

This note specifies using packets' source addresses in route lookups as additional qualifier per [extra-qualifiers], as applicable to IPv6 [RFC2460] without specific considerations for any routing protocol.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 23, 2015.

Copyright Notice

Table of Contents

1. Introduction 2
1.1. Requirements Language 2
2. Principle of operation 2
2.1. Lookup ordering and disambiguation 3
3. Applicability 3
3.1. Recursive Route Lookups 3
3.2. Unicast Reverse Path Filtering 4
3.3. Multicast Reverse Path Forwarding 5
4. IANA Considerations 5
5. Security Considerations 5
6. Privacy Considerations 5
7. Acknowledgements 5
8. Change Log 6
9. References 6
9.1. Normative References 6
9.2. Informative References 6
Author's Address 6

1. Introduction

Since connectivity providers generally secure their ingress along the lines of BCP 38 [RFC2827], small multihomed networks have a need to ensure their traffic leaves their network with a correct combination of source address and exit taken. This applies to networks of a particular pattern where the provider's default (dynamic) address provisioning methods are used and no fixed IP space is allocated, e.g. home networks, small business users and mobile ad-hoc setups.

While IPv4 networks would conventionally use NAT or policy routing to produce correct behaviour, this not desirable to carry over to IPv6. Instead, assigning addresses from multiple prefixes in parallel shifts the choice of uplink to the host. However, now for finding the proper exit the source address of packets must be taken into account.

For a general introduction and aspects of interfacing routers to hosts, refer to [I-D.sarikaya-6man-sadr-overview].

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Principle of operation

The mechanism in this document is such that a source prefix is added to all route entries. This document assumes all entries have a source prefix, with ::/0 as default value for entries installed without a specified source prefix. This need not be implemented in this particular way, however the system MUST behave exactly as if it were. In particular, a difference in behaviour between routes with a source prefix of ::/0 and routes without source prefix MUST NOT be visible.

For uniqueness considerations, the source prefix factors MUST be taken into account for comparisons. Two routes with identical information except the source prefix MAY exist and MUST be installed and matched.

2.1. Lookup ordering and disambiguation

As outlined in [extra-qualifiers], there is a fixed longest-match ordering dor route lookup.

For longest-match lookups, the source prefix is matched after the destination prefix. This is to say, first the longest matching destination prefix is found, then the table is searched for the route with the longest source prefix match, while only considering routes with exactly the destination prefix previously found. If and only if no such route exists (because none of the source prefixes match), the lookup moves to the next less specific destination prefix.

A router MUST continue to a less specific destination prefix if no route matches on the source prefix. It MUST NOT terminate lookup on such an event.

3. Applicability

3.1. Recursive Route Lookups

TBD, multiple possible approaches:

variant 1: ignore dst-src routes, only use routes with src ::/0

variant 2: exact-match src prefixes from resolvee to resolvent (will not work for a lot of cases)

variant 3: longer-match src prefixes from resolvee to resolvent (nexthop src may be superset of looked-up route)

variant 4: create multiple instances of the route whose nexthop is resolved, with different source prefixes

(Variant 4:)

When doing recursive nexthop resolution, the route that is being resolved is installed in potentially multiple copies, inheriting all possible more-specific routes that match the nexthop as destination. The algorithm to do this is:

1. form the set of attributes for lookup by using the (unresolved, recursive) nexthop as destination (with full host prefix length, i.e. /128), copy all other attributes from the original route

2. find all routes that overlap with this set of attributes (including both more-specific and less-specific routes)

3. order the result from most to less specific

4. for each route, install a route using the original route's destination and the "logical and" overlap of each extra match attribute with same attribute from the set. Copy nexthop data from the route under iteration. Then, reduce the set of extra attributes by what was covered by the route just installed ("logical AND NOT").

Example recursive route resolution

```
route to be resolved:
2001:db8:1234::/48, source 2001:db8:3456::/48,
                    recursive nexthop via 2001:db8:abcd::1

routes considered for recursive nexthop:
::/0,                                            via fe80::1
2001:db8:abcd::/48,                              via fe80::2
2001:db8:abcd::/48,   source 2001:db8:3456:3::/64, via fe80::3
2001:db8:abcd::1/128, source 2001:db8:3456:4::/64, via fe80::4

recursive resolution result:
2001:db8:1234::/48,   source 2001:db8:3456::/48,   via fe80::2
2001:db8:1234::/48,   source 2001:db8:3456:3::/64, via fe80::3
2001:db8:1234::/48,   source 2001:db8:3456:4::/64, via fe80::4
```

3.2. Unicast Reverse Path Filtering

Unicast reverse path filtering MUST use dst-src routes analog to its usage of destination-only routes. However, the system MAY match either only incoming source against routes' destinations, or it MAY match source and destination against routes' destination and source. It MUST NOT ignore dst-src routes on uRPF checks.

3.3. Multicast Reverse Path Forwarding

Multicast Reverse Path Lookups are used to find paths towards the (known) sender of multicast packets. Since the destination of these packets is the multicast group, it cannot be matched against the source part of a dst-src route. Therefore, dst-src routes MUST be ignored for Multicast RPF lookups.

4. IANA Considerations

This document updates the "Routing Qualifier Registry" described in [extra-qualifiers]. The following entry is added:

Name: IPv6 Source Address

Applicable Protocols: IPv6

Reference: this document

Position: after IPv6 Destination Address, before end of list

5. Security Considerations

Systems operating under the principles of this document can have routes that are more specific than the previously most specific, i.e. host routes. This can be a security concern if an operator was relying on the impossibility of hijacking such a route.

While source/destination routing could be used as part of a security solution, it is not really intended for the purpose. The approach limits routing, in the sense that it routes traffic to an appropriate egress, or gives a way to prevent communication between systems not included in a source/destination route, and in that sense could be considered similar to an access list that is managed by and scales with routing.

6. Privacy Considerations

If a host's addresses are known, injecting a dst-src route allows isolation of traffic from that host, which may compromise privacy. However, this requires access to the routing system. As with similar problems with the destination only, defending against it is left to general mechanisms protecting the routing infrastructure.

7. Acknowledgements

The base underlying this document was first outlaid by Ole Troan and Lorenzo Colitti in [I-D.troan-homenet-sadr] for application in the homenet area.

This document is largely the result of discussions with Fred Baker and derives from [I-D.baker-ipv6-isis-dst-src-routing].

8. Change Log

Initial Version: October 2014

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S.E. and R.M. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[extra-qualifiers] Lamparter, D., "Considerations and Registry for extending IP route lookup", .

9.2. Informative References

[I-D.baker-ipv6-isis-dst-src-routing] Baker, F., "IPv6 Source/Destination Routing using IS-IS", draft-baker-ipv6-isis-dst-src-routing-01 (work in progress), August 2013.

[I-D.sarikaya-6man-sadr-overview] Sarikaya, B., "Overview of Source Address Dependent Routing", draft-sarikaya-6man-sadr-overview-01 (work in progress), September 2014.

[I-D.troan-homenet-sadr] Troan, O. and L. Colitti, "IPv6 Multihoming with Source Address Dependent Routing (SADR)", draft-troan-homenet-sadr-01 (work in progress), September 2013.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

Author's Address

David Lamparter
NetDEF
Leipzig 04103
Germany

Email: david@opensourcerouting.org