Homenet working Group
Internet-Draft
Intended status: Informational
Expires: August 22, 2013

O. Troan
Cisco Systems
L. Colitti
Google

February 18, 2013

IPv6 Multihoming with Source Address Dependent Routing (SADR)
draft-troan-homenet-sadr-00

Abstract

A multihomed network using provider aggregatable addresses must send the packet out the right path to avoid violating the provider's ingress filtering. This memo suggests a mechanism called Source Address Dependent Routing to solve that problem.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

IPv6 is designed to support multiple addresses on an interface, and the intention was to use this feature to support multihoming with provider aggregatable addresses.

One difficulty of multihoming with provider-aggregatable space is that providers typically employ BCP38 [RFC2827] filtering. If a network sends traffic to its upstream provider using a source address that was not assigned by that provider, the traffic will be dropped. Thus, if a network is multihomed to multiple providers, it must ensure that traffic is sent out the correct exit for the packet's source address.

As long as upstream traffic is sent to the correct provider, hosts inside the network are free to use source addresses assigned by any of the network's upstream providers. In such a scenario, each host has multiple addresses, one or more from each provider the network is connected to. The network ensures that packets are sent to the correct upstream by forwarding packets based on the destination address and the source address. This we call source address dependent routing (SADR). This memo shows how SADR can be used to implement multihoming.

2. Conventions

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 RFC 2119 [RFC2119].

3. Terminology

Service Provider An entity that provides the network with external connectivity, e.g. to the Internet.

WAN Interface — An interface connected to a Service Providers. WAN interfaces may either be physical links or virtual interfaces such as tunnels. WAN interfaces are used to send ingress traffic from the Internet to the End-User, and egress traffic from the End-User network to the Internet. Ingress traffic may be received on any active interface at any time. Egress traffic follows a set of rules within the router in order to choose the proper WAN interface.

Border Router — A border router has one or more external interfaces connecting it to one or more Service Providers. The border router receives one or more delegated prefixes, each associated with one or more WAN interfaces.

External Route — A route that is learned from a Service Provider. Each External Route has an Acceptable Source Prefix which determines which source addresses may use that route.

Internal Router — A router that is not a Border Router.

Internal Route — A route to a destination inside the network.

4. Using SADR for multihoming

SADR is similar to policy based routing. This memo proposes a simple extension to the destination based longest match algorithm to constrain it to source address.

In order to support ingress filtering by upstream networks, the network MUST treat external routes specially. Ingress filtering MAY also be used internally, by installing (S,D) routes for locally assigned prefixes, where the source prefix would be the aggregatable prefix. If no ingress filtering is performed inside the network, then normal non-source constrained forwarding is used.

5. A Conceptual Forwarding Algorithm

This section describes a conceptual forwarding algorithm. An implementation might implement this differently, e.g. with multiple tables, as long as the external behaviour is as described.

First a longest match lookup is done in the routing table for the destination address, then for the resulting set a longest matching lookup is done for the source address.

In a destination based routing table, an entry in the routing table can be shown as "D -> NH". That is, to get to a destination D, use next-hop NH. For a source constrained routing table we propose the following notation. (Source Network, Destination network) -> Next-hop. (S, D) -> NH. A route that is not source constrained can be represented as (*, D) -> NH.

For convenience this document shows the routing table as a single destination based routing table, with source address constrained paths. This does not preclude other implementations, as long as the external behaviour is the same.

A router forwarding a packet does a longest match look-up on the destination address. If this is a (*, D) entry, it forwards the packet out the best next-hop as before (doing equal cost multi path load balancing etc). If the look-up results in a (S, D) entry, the look-up function does a longest match on the source address among the set of (S, D) paths. If there is a match the packet is forwarded out the given next-hop, if not an ICMP destination address unreachable message, code 5 is returned [RFC4443]. A routing entry may have both (S, D) paths and (*, D) paths. The longest match wins.

The following example show the routing table of a network connected to two ISPs, ISP A and ISP B. Both ISPs offer default connectivity and ISP B also offers a more specific route to a walled garden service.

```
(2001:db8::/56, ::/0) -> ISP_A          # Default route to ISP A
(2001:db9::/56, ::/0) -> ISP_B          # Default route to ISP B
(*, 2001:db8::/64) -> R1                # Internal network, prefix from A
(*, 2001:db8:1::/64) -> R2              # Internal network, prefix from A
(*, 2001:db9::/64) -> R1                # Internal network, prefix from B
(*, 2001:db9:1::/64) -> R2              # Internal network, prefix from B
(*, fd00::/64) -> R3                    # Internal network ULA
(2001:db9::/56, 2001:420::/32) -> ISP_B # Walled garden route from ISP B
```

Figure 1: Example Routing Table

A packet with the SA, DA of 2001:db8::1, 2001:dead:beef::1 would be forwarded to ISP A, likewise a packet with SA, DA 2001:db9::1, 2001:dead:beef::1 would be forward to ISP B. An packet with SA,DA 2001:db8::1, 2001:db9::1 would be forwarded using normal destination based routing. A packet to the walled garden SA,DA 2001:db9::1, 2001:420::1 would be sent to ISP B. A packet with SA,DA 2001:db8::1,2001:420::1 would be dropped with an ICMP unreachable message being sent back.

6. Routing considerations

Now that we have described the function of the source constrained routing table. How does the table get populated?

6.1. Routing Protocol extensions

The generic answer is that the routing protocol used in the network has to be extended to support (S, D) routes. Specifically, the routing protocol should distribute, for each External Route, the Acceptable Source Prefix(es) for that route. This may be done, for example, using [I-D.baker-ipv6-ospf-dst-src-routing] or [I-D.baker-ipv6-isis-dst-src-routing]. In the case of OSPFv3, for example, external routes are advertised in an AS-External-prefix LSA, [RFC5340]

6.2. Simplified SADR in home networks

In a home network using a dynamic prefix assignment mechanism such as [I-D.arkko-homenet-prefix-assignment] it may be known that a particular Border Router is announcing both an External Route and a Usable Prefix (for example, if the same router ID is announcing both). In this case, interior routers may assume that the Acceptable Source Prefix of the External Route announced by that Border Router is in fact the Usable Prefix announced by that Border Router.

An internal router when receiving a AS-External LSA route will install that in the routing table as normal. When the internal router receives a usable prefix as part of prefix assignment, the router shall add source constrained entries to all the AS-External routes received from the same border router (matching router-ID).

Routes that are not associated with a border router or are not AS-External do not have source constrained paths.

The routing protocol requirements for simplified SADR in the home network are:

1. Routing protocol must flood all information to all routers in the home network. (Single area).

2. Prefix assignment and unicast routing must be done in the same protocol.

3. A router must be uniquely identified (router-id) so that router advertisements and prefix assignment can be tied together

7. Interaction between routers and hosts

Generally, hosts need not be aware that SADR is in use in the network. Hosts simply choose source addresses and the network will deliver the traffic to the appropriate upstream. One exception is when an Acceptable Source Prefix becomes invalid (e.g., if the Border Router which announced it crashes, or its WAN link goes down). In this case, current hosts will continue to use source addresses in that Acceptable Source Prefix without knowing that all communication outside the network is likely to fail. In this case, interior routers can improve responsiveness by deprecating the addresses in that Acceptable Source Prefix.

ICMP [RFC4443] includes a Destination unreachable code 5 - "Source address failed ingress/egress policy". Hosts MUST adhered to this message, and based on the unreachable message try another source address.

8. IANA Considerations

This specification does not require any IANA actions.

9. Security Considerations

10. Acknowledgements

The authors would like to thank Jari Arkko and Andrew Yourtchenko for their ideas and review.

11. References

11.1. Normative References

[I-D.arkko-homenet-prefix-assignment]
Arkko, J., Lindem, A., and B. Paterson, "Prefix Assignment in a Home Network", draft-arkko-homenet-prefix-assignment-03 (work in progress), October 2012.

[I-D.ietf-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration", draft-ietf-ospf-ospfv3-autoconfig-00 (work in progress), October 2012.

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

[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.

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control

Message Protocol (ICMPv6) for the Internet Protocol
Version 6 (IPv6) Specification", RFC 4443, March 2006.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

[RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

11.2. 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-00 (work in progress), February 2013.

[I-D.baker-ipv6-ospf-dst-src-routing]
Baker, F., "IPv6 Source/Destination Routing using OSPFv3", draft-baker-ipv6-ospf-dst-src-routing-00 (work in progress), February 2013.

[I-D.ietf-homenet-arch]
Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "Home Networking Architecture for IPv6", draft-ietf-homenet-arch-06 (work in progress), October 2012.

Authors' Addresses

Ole Troan
Cisco Systems
Philip Pedersens vei 1
Lysaker 1366
Norway

Email: ot@cisco.com

Lorenzo Colitti
Google
Roppongi Hills Mori Tower PO box 22
Minato, Tokyo 106-6126
Japan

Email: lorenzo@google.com