6MAN
Internet-Draft
Intended status: Informational
Expires: December 14, 2020

D. Dukes, Ed.
C. Filsfils, Ed.
Cisco Systems

June 12, 2020

Segment Routing Traffic Engineering Leveraging Existing IPv6 Interface Addresses
draft-dukes-6man-sr-te-intf-address-00

Abstract

This document illustrates how an operator may re-use an existing IPv6 address allocation within its domain to deliver SR-based Traffic Engineering service.

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 https://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 December 14, 2020.

Table of Contents

1. Introduction

This document illustrates how an operator may re-use an existing IPv6 address allocation within its domain to deliver SR-based Traffic Engineering service by describing:

- A reference topology with IPv6 address allocation.
- Binding a SID behavior to existing IPv6 addresses.
- The life of a packet forwarded via an SR policy.
- Upper-layer header processing for a SID bound to an existing IPv6 address.

The illustrations cover traffic engineering (TE) SR policy between two border routers of the domain and two hosts of the domain.

2. Reference Topology

The reference topology is the same as Section 6.2 of [RFC8754].

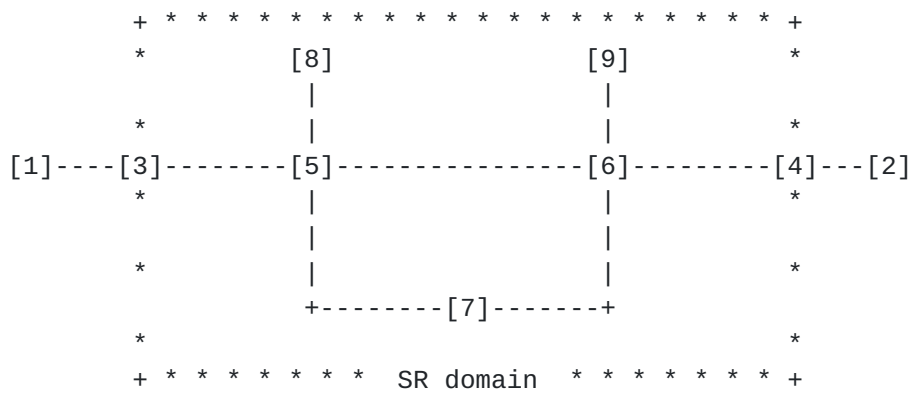

Figure 1: Reference topology

- 3 and 4 are SR domain edge routers
- 5, 6, and 7 are all SR domain routers
- 8 and 9 are hosts within the SR domain
- 1 and 2 are hosts outside the SR domain

3. Address Allocation

The operator has allocated 2001:db8:0::/48 to their domain.

A router K is sub-allocated 2001:db8:0:K::/64.

A router K has at least one loopback interface.

The first loopback interface of a router K's is assigned 2001:db8:0:K::1/128.

The interfaces of a router K attached to point to point links connected to other nodes within the domain are assigned link-local addresses.

4. SID Bound To Existing Interface Address

The operator enables SR segment endpoint node functionality on a few routers within the domain by binding the SID described in Section 4.3.1 of [RFC8754] to the IPv6 address assigned to the loopback interface of router 3 (2001:db8:0:3::1), router 4 (2001:db8:0:4::1) and router 7 (2001:db8:0:7::1).

Packet processing at these segment endpoint nodes follows that defined in Section 4.3 of [RFC8754].

5. Life Of A Packet

This section uses the abstract representation of an SRH as defined in Section 6.1 of [RFC8754].

It illustrates two examples from Section 6 of [RFC8754] for inter SR domain and intra SR domain packets and the processing at SR source nodes, transit nodes and SR segment endpoint nodes using the SIDs bound to interface addresses.

5.1. Inter SR Domain

Host 1 sends a packet (P1) to host 2

P1: (A1,A2)

The SR domain ingress router 3 receives P1 and steers it to SR domain egress router 4 via an SR Policy <2001:db8:0:7::1, 2001:db8:0:4::1>. Router 3 encapsulates the received packet P1 in an outer header with a reduced SRH and sends the packet

P2: (2001:db8:0:3::1, 2001:db8:0:7::1)(2001:db8:0:4::1; SL=1)(A1,A2)

Router 5 acts as a transit node for P2 and forwards it on the interface toward router 7.

Router 7 receives packet P2 and, using the logic in Section 4.3.1.1 of [RFC8754], decrements the Segments Left value and updates the Destination Address to 2001:db8:0:4::1. It sends the resulting packet

P3: (2001:db8:0:3::1, 2001:db8:0:4::1)(2001:db8:0:4::1; SL=0)(A1,A2)

on the interface toward router 6.

Router 6 acts as a transit node for packet P3 and forwards P3 on the interface toward router 4.

Router 4 receives packet P3 and, using the logic in Section 4.3.1.2 of [RFC8754], performs IPv6 decapsulation on P2 and forwards the inner packet P1: (A1,A2) on the interface toward host 2.

5.2. Intra SR Domain

When host 8 sends a TCP packet to host 9 via an SR Policy <2001:db8:0:7::1, 2001:db8:0:9::1> the packet is

P4: (2001:db8:0:8::1, 2001:db8:0:7::1)(2001:db8:0:9::1; SL=1) (TCP)

Processing of P4 is similar to P2 above; router 5 forwards while router 7 processes the SRH resulting in the following packet

P5: (2001:db8:0:8::1, 2001:db8:0:9::1)(2001:db8:0:9::1; SL=0) (TCP)

P5 is forwarded by router 6 to host 9 where the packet is consumed and its TCP payload is processed.

6. Upper-Layer Header Processing

The SID behavior described in [RFC8754] permits some upper-layer processing and blocks others. In some use-cases upper-layer processing may be limited when additional SID's are allocated independently of any existing interface address, and as a conservative security measure.

In this use-case the operator re-uses existing interface addresses for SIDs, it is expected that upper-layer processing is preserved and permitted for those addresses.

The following sections describe ping, ping via an SR policy and SSH session initiation for these SIDs.

6.1. ICMPv6 Echo Request and Reply

This section illustrates the life of an ICMPv6 echo request from router 3 (2001:db8:0:3::1) to router 4 (2001:db8:0:4::1) and of the corresponding ICMPv6 echo reply.

When router 3 sends an ICMPv6 echo request from 2001:db8:0:3::1 to 2001:db8:0:4::1 on router 4, the packet is

P6: (2001:db8:0:3::1, 2001:db8:0:4::1)(ICMPv6 echo request)

Router 4 receives packet P6 and follows Section 4.3.1 of [RFC8754]. Specifically, P6 does not contain an SRH and, since upper-layer header processing is permitted, router 4 processes packet P3 as per [RFC4443] and sends the response packet

P7: (2001:db8:0:4::1, 2001:db8:0:3::1)(ICMPv6 echo reply)

on the interface toward router 6.

Router 3 receives packet P7 and applies Section 4.3.1 of [RFC8754]. Specifically, P7 does not contain an SRH and, since upper-layer header processing is permitted, router 3 processes packet P4 as per [RFC4443].

6.2. ICMPv6 Echo Request via an SR Policy

This section illustrates the life of an ICMPv6 echo request from router 3 (2001:db8:0:3::1) to router 4 (2001:db8:0:4::1) via router 7 (2001:db8:0:7::1), and of the corresponding ICMPv6 echo reply.

When router 3 sends an ICMPv6 echo request from 2001:db8:0:3::1 to 2001:db8:0:4::1 via an SR Policy <2001:db8:0:7::1, 2001:db8:0:4::1> using a reduced SRH, the packet is

P8: (2001:db8:0:3::1, 2001:db8:0:7::1)(2001:db8:0:4::1; SL=1)(ICMPv6 echo request)

Router 7 eventually receives packet P8 and, using the logic in Section 4.3.1.1 of [RFC8754], decrements the Segments Left value and updates the Destination Address to 2001:db8:0:4::1. It sends the resulting packet

P9: (2001:db8:0:3::1, 2001:db8:0:4::1)(2001:db8:0:4::1; SL=0)(ICMPv6 echo request)

on the interface toward router 6.

Router 4 receives packet P9 and applies Section 4.3.1 of [RFC8754]. Specifically, it determines that packet P9 contains an SRH with Segments Left equal to 0 and proceeds to process the next header in the extension header chain, as per Section 4.3.1.1 of [RFC8754]. Since upper-layer header processing is permitted, router 4 processes packet P9 as per [RFC4443] and sends the response packet

P10: (2001:db8:0:4::1, 2001:db8:0:3::1)(ICMPv6 echo reply)

on the interface toward router 6.

Packet P10 follows the same return path as packet P7 above.

6.3. SSH Session Initiation

This section illustrates the initiation of a SSH session between router 3 (2001:db8:0:3::1) and router 4 (2001:db8:0:4::1).

SSH first establishes a TCP session between the two routers. Router 3 sends an TCP SYN packet from 2001:db8:0:3::1 to 2001:db8:0:4::1 on router 4, resulting in

P11: (2001:db8:0:3::1, 2001:db8:0:4::1)(TCP SYN)

Router 4 receives packet P11 and applies Section 4.3.1 of [RFC8754]. Specifically, it determines that packet P11 does not contain an SRH and, since upper-layer header processing is permitted, processes packet P11 as per [RFC0793] and sends the response packet

P12: (2001:db8:0:4::1, 2001:db8:0:3::1)(TCP SYN-ACK)

on the interface toward router 6.

The rest of the communication occurs as normal for SSH [RFC4253].

7. Security Considerations

The SR domain is secured via ingress filtering of packets as described in [RFC8754] Section 5.1. In this document packets entering the SR domain destined to infrastructure addresses are dropped at ingress edge nodes since the SID and infrastructure address prefixes are the same (eg. 2001:db8:0::/48).

When an SRv6-capable node receives an IPv6 packet, it performs a longest-prefix-match lookup on the packet's destination address. It processes any SRH in the packet only when the destination address is bound to a SID ([RFC8754] Section 4.3). This further limits the possible attack surface to a subset of the infrastructure address prefix protected by ingress filtering.

The SID behavior bound to an address may limit some upper-layer processing ([RFC8754] Section 4.3.1.2). In the use-case described in this document, upper-layer header processing is not limited for an address the SID behavior is bound to.

8. IANA Considerations

This document has no IANA actions.

9. Ecosystem

The use-case described in this document is supported on Arccus, Broadcom, Cisco, and Linux.

10. References

10.1. Normative References

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

10.2. Informative References

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <https://www.rfc-editor.org/info/rfc793>.

[RFC4253] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Transport Layer Protocol", RFC 4253, DOI 10.17487/RFC4253, January 2006, <https://www.rfc-editor.org/info/rfc4253>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

Authors' Addresses

Darren Dukes (editor)
Cisco Systems
Canada

Email: ddukes@cisco.com

Clarence Filsfils (editor)
Cisco Systems
Belgium

Email: cfilsfil@cisco.com