Workgroup: Network Working Group
Internet-Draft:
draft-raviolli-intarea-trusted-domain-srv6-03
Published: 9 April 2024
Intended Status: Standards Track
Expires: 11 October 2024
Authors: A. Alston
Liquid Intelligent Technologies
T. Hill
British Telecom
A. Przygienda
Juniper
L. Jalil
Verizon

Trusted Domain SRv6

Abstract

SRv6 as designed has evoked interest from various parties, though its deployment is being limited, amongst other things, by known security problems in its architecture. This document specifies a standard way to create a solution that closes some of the major security concerns, while retaining the tenants of the SRv6 protocol.

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 RFC 2119 [RFC2119].

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 11 October 2024.

Copyright Notice

Table of Contents

1. Description

SRv6 as designed has evoked interest from various parties, though its deployment is being limited by known security problems in its architecture. This document specifies a standard way to create a solution that closes some of the major security concerns, while retaining the basis of the SRv6 protocol.

2. Glossary

Fail-Closed Domain:
synonymous with a Trusted Domain.

Trusted Domain (TD):
A domain that prevents processing of a protocol without explicit configuration, defined in detail in Section 4. This document is

limited to treatment of deployment of SRv6 in the context of a trusted domain only.

Fail-Closed Protocol (FCP):
A protocol that can be deployed by establishing a fail-closed domain.

TD-SRv6:
SRv6 modified to become a FCP and with that allowing for easy deployment in a trusted domain.

3. The SRv6 Security Problems

SRv6 [RFC8402] relies on the concept of limited domain. The application of this concept in the context of the draft however, suffers from a lack of security that is easily deployable in an economi and scalable fashion.

Limited domains without very careful deployment will invariably leak beyond the domain and allow untrusted traffic to enter the domain and terminate on any arbitrary node.

As per RFC 8402 [RFC8402]RFC8402 Section 8, SRv6 that leaks beyond the border of a trusted domain creates a security violation.

An established and proven solution is to create a trusted domain that has a default fail-closed approach and a well-defined trusted/ untrusted boundary.

Examples of fail-closed protocols include:

*mpls

*clns

*bier

4. Characteristics of a Fail-Closed Domain

A fail-closed domain is determined by following properties:

Processing of the protocol packet on an interface requires explicit configuration. Otherwise, due to lack of packet classification, further processing and forwarding cannot be achieved. In practical terms the behavior used most often is a drop of the offending packet.

In a fail-closed protocol, leaking beyond the boundary of the domain requires explicit config.

Fail-closed protocols are easily identifiable by their top level (e.g. link layer) encoding early in the packet formats and often by fields at a fixed offset. In another words either their encoding or encapsulation allows such packets to be easily distinguished from other traffic.

Classification of the protocol packets is completely deterministic.

Confining the protocol to the trusted domaim does not require complex processing in either hardware or software to allow for scalable and economical deployment.

The boundary of a trusted domain consists of a set of interfaces that exhibit default behavior.

5. SRv6 in the context of a trusted domain - an objective analysis

It is impossible to differentiate SRv6 and IPv6 at the link-layer or easily at network layer by e.g. a reserved protocol number the way IPSec does since SRv6 and IPv6 share the same ethernet types and IP protocol numbers.

Hence, in the event of a packet being sent into a trusted domain, either accidentally or by a malicious actor, it is possible to send the frame to a node binding the specific SID, and have the packet processed, irrespective of the content of the underlying (encapsulated) packet.

The security proposals in RFC8402 section 8.2 is based on the application of filters preventing ingress traffic at the boundary routers destined towards a SID within the domain. Such filtering is prone to configuration errors and in addition, has significant impact on fast matching hardware utilization on devices that have large numbers of ingress points into the domain. The matching itself, due to the complexity and numerous possibilities of expressing a set of SIDs will likely necessitate a complete semantic parsing of such list to guarantee fully precise matching including wildcarding in different forms.

In the context of a trusted domain, anything outside of the operators control should not be considered trusted. This means applying filters to prevent leakage into the domain at every customer port, every server, and every cloud stack. The scale and complexity of maintaining such a "shorewall" is daunting and at large scale will not be likely to keep up with the timing necessary in case of attacks mounted and metamorphosing in short time intervals. An attack avoiding the filter wall may evade discovery for a long time in the absence of sophisticated traffic analyis and analytics tools.

6. Trusted-Domain Implementation

To implement SRv6 in the context of a trusted domain, it is necessary to modify it to allow deployment in a fail-closed boundary efficiently. This requires changes to the protocol encapsulation at both the boundary routers and the transit nodes. This document introduces a distinct ethertype to be used for TD-SRv6 packets.

6.1. Boundary routers

Trusted Domain boundary routers form the point at which the new ethertype is imposed on interfaces configured to represent such boundary. Imposition of the ethertype happens on packet ingress, at the same point as SRv6 header imposition is performed.

Boundary interfaces will, by default behavior and unless configured otherwise, drop packets containing the TD-SRv6 ethertype already and MUST drop packets containing an SRH (or otherwise being classified clearly as SRv6 frame) if received on any ethertype except TD-SRv6.

6.2. Transit and egress routers

In the case of a transit or egress router, should a frame not be marked with the TD-SRv6 ethertype, the frame will be treated as a standard IPv6 packet for the purposes of handling and forwarding. Even if an SRv6 packet is introduced into such domain with an ethertype different from TD-SRv6, the according SRv6 packet handling will not occur. Hence the resulting handling of the packet is indistinguishable from standard IPv6 processing.

A router configured to process TD-SRv6 MUST drop packets containing an SRH (or otherwise being classified clearly as SRv6 frame) if received on any ethertype except TD-SRv6 and MUST apply SRv6 processing if and only if the frame is marked as TD-SRv6 ethertype.

6.3. Transit and egress routers not using TD-SRv6

It cannot be excluded that deployment of TD-SRv6 are using TD-SRv6 on only a subset of external interfaces and/or choose to revert to standard IPv6 ether type for SRv6 packets within some or all interfaces facing the internal domain. The mechanisms required to realize such a deployment and risks incurred are outside the scope of this document.

7. Registry Considerations

7.1. IANA Considerations

No IANA Considerations

7.2. IEEE Considerations

TD-SRv6 Ethertype: TBD0

8. Security Considerations

This draft enhances the security mechanisms required by section 8 of RFC8402, and does not impose any further security considerations of its own.

9. Applicability Considerations

TD-SRv6 is applicable in situations where the transport domain using SRv6 is not considered a fully trusted closed user group, i.e. not every participant can be trusted to not accept IPv6 frames from other domains or issue IPv6 frames within the domain using some mechanism. In the latter case the attack surface to craft malicious SRv6 frames looking potentially like innocuous IPv6 frames is open. A good example being servers. On the other hand, a fully trusted user group can be assumed e.g. in overlay situation, i.e. a transport provider offering VPN service where IPv6 framed are neither injected or accepted from the overlay. In a sense, the VPN tunnel encapsulation acts as security mechanism preventing the closed user group from injecting IPv6 frames carried on the tunnel into the transport domain.

10. Contributors

Weiqiang Cheng

chengweiqiang@chinamobile.com

Anthony Somerset

anthony.somerset@liquid.tech

11. References

11.1. Informative References

11.2. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment

Routing Architecture", RFC 8402, DOI 10.17487/RFC8402,
July 2018, <https://www.rfc-editor.org/info/rfc8402>.

Authors' Addresses

Andrew Alston
Liquid Intelligent Technologies

Email: andrew-ietf@liquid.tech

Tom Hill
British Telecom

Email: tom@ninjabadger.net

Tony Przygienda
Juniper
United States of America

Email: prz@juniper.net

Luay Jalil
Verizon
United States of America

Email: luay.jalil@verizon.com