Network Working Group
Internet Draft
Intended status: Standards Track
Expires: August 15, 2025

W. Cheng
China Mobile
C. Lin
New H3C Technologies
S. Yue
China Mobile

February 15, 2025

Signaling SAVNET Capability Using IGP
draft-cheng-lsr-adv-savnet-capbility-01

Abstract

Existing intra-domain SAV solutions (e.g., BCP38 [RFC2827] and BCP84 [RFC3704]) have problems of high operational overhead or inaccurate validation (see [I-D.ietf-savnet-intra-domain-problem-statement]). To address these problems and guide the design of new intra-domain SAV solutions,[I-D.ietf-savnet-intra-domain-architecture] proposes the architecture of intra-domain SAVNET and introduces the use of SAV-specific information in intra-domain networks. This document defines a mechanism to signal the SAVNET capablity and the source prefix using IGP and BGP-LS.

Status of this Memo

This Internet-Draft will expire on August 15, 2025.

Table of Contents

1. Introduction

Existing intra-domain SAV solutions (e.g., BCP38 [RFC2827] and BCP84 [RFC3704]) have problems of high operational overhead or inaccurate validation (see [I-D.ietf-savnet-intra-domain-problem-statement]). ACL-based ingress filtering requires manual operations to configure and update the SAV rules, while uRPF-based solutions may improperly block legitimate data packets in the scenario of routing asymmetry. To address these problems and guide the design of new intra-domain

SAV solutions,[I-D.ietf-savnet-intra-domain-architecture] proposes the architecture of intra-domain SAVNET and introduces the use of SAV-specific information in intra-domain networks.

Following the intra-domain SAVNET architecture, this document defines a mechanism to signal the SAVNET capability and the source prefix using IGP and BGP-LS. This document focuses on announcing the SAVNET capability of the router. When the SAVNET capability is enabled, a flag can be used to indicate whether the prefix being announced is a SAVNET source prefix. Announcing the SAVNET capability and source prefix facilitates the operations and management of SAVNET and the implementation of specific schemes. The specific processing of SAVNET entries is not within the scope of this document.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Signaling SAVNET Capability Using IS-IS

2.1. SAVNET Capabilities Sub-TLV

A node indicates that it supports the SAVNET functionality as specified in [I-D.ietf-savnet-intra-domain-architecture] by advertising a new SAVNET Capabilities sub-TLV of the Router Capability TLV [RFC7981].

The SAVNET Capabilities sub-TLV may contain optional sub-sub-TLVs. No sub-sub-TLVs are currently defined.

The SAVNET Capabilities sub-TLV has the following format:

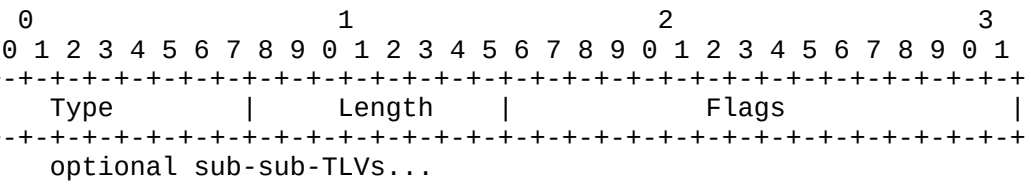

where:

- o Type: TBD. Single octet, as defined in Section 9 of [ISO10589].
- o Length: Single octet, as defined in Section 9 of [ISO10589]. The length value is 2 + length of sub-sub-TLVs.
- o Flags: 2 octets. The following flags are defined:

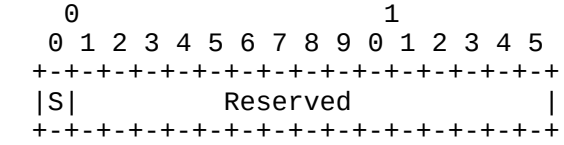

where:

S-flag: If set, the router supports SAVNET functionality. The remaining bits, including bit 0, are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

2.2. Source Prefix Flag

IPv4 SAVNET source prefixes are advertised using "IP Extended Reach TLV" (type 135), while IPv6 SAVNET source prefixes are advertised using "IPv6 Reachability TLV" (type 236, RFC5308).

A new bit in the IPv4/IPv6 Extended Reachability Attribute Flags [RFC7794] is defined:

S-Flag: Source Prefix Flag (Bit TBD)

When set, it indicates that the prefix is used for source address validation in the data plane.

3. Signaling SAVNET Capability Using OSPF

3.1. SAVNET Capabilities Sub-TLV

The SAVNET Capabilities TLV is used by an OSPF router to advertise its support for the SAVNET functionality. This is an optional top-level TLV of the OSPF Router Information LSA [RFC7770] that MUST be advertised by an SAVNET-enabled router. The RI LSA can be advertised at any of the defined opaque flooding scopes (link, area, or Autonomous System (AS)). For the purpose of SAVNET Capabilities TLV advertisement, area-scoped flooding is REQUIRED. Link and AS-scoped flooding is OPTIONAL.

The format of the OSPF SAVNET Capabilities TLV is shown below:

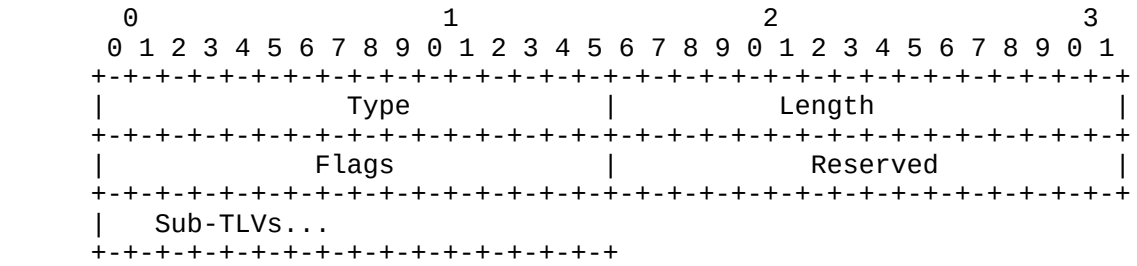

where:

- Type: 2-octet field. The value for this type is 20.
- Length: 2-octet field. The total length (in octets) of the value portion of the TLV, including nested sub-TLVs.
- Reserved: 2-octet field.It MUST be set to 0 on transmission and MUST be ignored on receipt.
- Flags: 2-octet field. The flags are defined as follows:

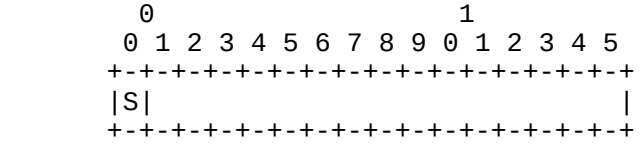

where:
S-flag: If set, the router supports SAVNET functionality.
The remaining bits, including bit 0, are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

The SAVNET Capabilities TLV may contain optional sub-TLVs. No sub-TLVs are defined in this specification.

3.2. Source Prefix Flag

SAVNET source prefixes are advertised using "OSPFv2 Extended Prefix Opaque LSA"[RFC7684].

A new bit in Flags field of the OSPFv2 Extended Prefix TLV [RFC7684] is defined:

S-Flag: Source Prefix Flag (Bit TBD)

When set, it indicates that the prefix is used for source address validation in the data plane.

4. Signaling SAVNET Capability Using OSPFv3

4.1. SAVNET Capabilities Sub-TLV

The SAVNET Capabilities TLV is used by an OSPFv3 router to advertise its support for the SAVNET functionality. This is an optional top-level TLV of the OSPFv3 Router Information LSA [RFC7770] that MUST be advertised by an SAVNET-enabled router.

This TLV MUST be advertised only once in the OSPFv3 Router Information LSA. When multiple SAVNET Capabilities TLVs are received from a given router, the receiver MUST use the first occurrence of the TLV in the OSPFv3 Router Information LSA. If the SAVNET Capabilities TLV appears in multiple OSPFv3 Router Information LSAs that have different flooding scopes, the TLV in the OSPFv3 Router Information LSA with the area-scoped flooding scope MUST be used. If the SAVNET Capabilities TLV appears in multiple OSPFv3 Router Information LSAs that have the same flooding scope, the TLV in the OSPFv3 Router Information LSA with the numerically smallest Link State ID MUST be used, and subsequent instances of the TLV MUST be ignored.

The OSPFv3 Router Information LSA can be advertised at any of the defined flooding scopes (link, area, or Autonomous System (AS)). For the purpose of SAVNET Capabilities TLV advertisement, area-scoped flooding is REQUIRED. Link and AS-scoped flooding is OPTIONAL.

The format of the OSPFv3 SAVNET Capabilities TLV is shown below:

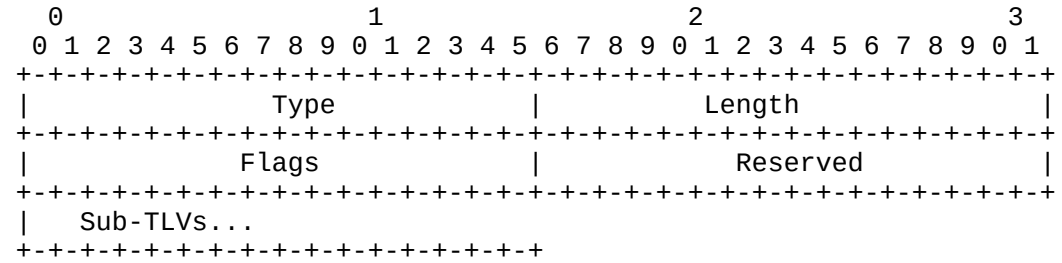

where:

- Type: 2-octet field. The value for this type is 20.
- Length: 2-octet field. The total length (in octets) of the value portion of the TLV, including nested sub-TLVs.
- Reserved: 2-octet field.It MUST be set to 0 on transmission and MUST be ignored on receipt.
- Flags: 2-octet field. The flags are defined as follows:

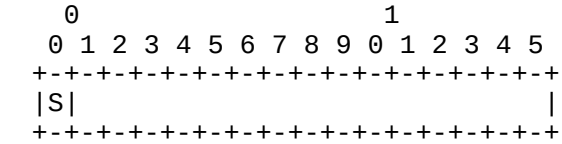

where:

S-flag: If set, the router supports SAVNET functionality. The remaining bits, including bit 0, are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

The SAVNET Capabilities TLV may contain optional sub-TLVs. No sub-TLVs are defined in this specification.

4.2. Source Prefix Flag

SAVNET source prefixes are advertised using "OSPFv3 Extended LSA", including "E-Intra-Area-Prefix-LSA", "E-Inter-Area-Prefix-LSA" and "E-AS-External-LSA".[RFC8362].

A new bit in the prefix Attribute Flags [I-D. draft-ietf-lsr-ospf-prefix-extended-flags-00] are defined:

S-Flag: Source Prefix Flag (Bit TBD)

When set, it indicates that the prefix is used for source address validation in the data plane.

5. Signaling SAVNET Capability in BGP-LS

The IGP extensions defined in this document can be advertised via BGP-LS (distribution of Link-State and Traffic Engineering information using BGP) [RFC7752] using existing BGP-LS TLVs.

This section defines the following Node Attribute TLV:

Type	Description
TBD	the SAVNET-Capabilities TLV

6. Security Considerations

TBD.

7. IANA Considerations

TBD

8. References

8.1. Normative References

[I-D.ietf-savnet-intra-domain-problem-statement] Li, D., Wu, J., Qin, L., Huang, M., and N. Geng, "Source Address Validation in Intra-domain Networks Gap Analysis,Problem Statement, and Requirements", Work in Progress,Internet-Draft, draft-ietf-savnet-intra-domain-problem-statement-03, 13 February 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-03>.

[I-D.ietf-savnet-intra-domain-architecture]Li, D., Wu, J., Qin, L., Geng, N., and L. Chen, "Intra-domain Source Address Validation (SAVNET) Architecture",Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-architecture-00, 12 April 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-architecture-00>.

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

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <https://www.rfc-editor.org/info/rfc7684>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

[RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 7770, DOI 10.17487/RFC7770, February 2016, <https://www.rfc-editor.org/info/rfc7770>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017

[RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <https://www.rfc-editor.org/info/rfc8362>.

8.2. Informational References

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <https://www.rfc-editor.org/info/rfc3704>.

Authors' Addresses

Weiqiang Cheng
China Mobile
China
Email: chengweiqiang@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com

Shengnan Yue
China Mobile
China
yueshengnan@chinamobile.com