

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 25 July 2022

P. Kaneriy
Tony. Li
Antoni. Przygienda
S. Hegde
C. Bowers
Juniper Networks
L. Ginsberg
Cisco Systems
21 January 2022

Multiple TLV Instances in IS-IS
draft-pkaneria-lsr-multi-tlv-00

Abstract

Emerging technologies are adding information into IS-IS TLVs at a steady pace while deployment scales are simultaneously increasing. This causes the contents of many critical TLVs to exceed the currently supported limit of 255 octets. Extensions such as [\[RFC7356\]](#) require significant IS-IS changes that could help address the problem, but a less drastic solution would be beneficial. This document codifies the common mechanism of extending the TLV space through multiple TLV instances.

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 25 July 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

Multi-TLV

January 2022

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	2
2.	Requirements Language	3
3.	Procedure for Advertising Multiple TLV Instances	3
4.	Procedure for Receiving Multiple TLV Instances	4
5.	Operational Considerations	4
6.	IANA Considerations	5
7.	Security Considerations	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	5
	Authors' Addresses	6

[1.](#) Introduction

The continued growth of the Internet has resulted in a commensurate growth in the scale of service provider networks and the amount of information carried in IS-IS [[IS010589](#)] Type-Length-Value (TLV) tuples. Simultaneously, new traffic engineering technologies are defining new attributes, further adding to the scaling pressures. The original TLV definition allows for 255 octets of payload, which is becoming increasingly inadequate.

Some TLV definitions have addressed this by explicitly stating that a TLV may appear multiple times inside of an LSP. However, this has not been done for many legacy TLVs, leaving the situation ambiguous. The intent of this document is to clarify the situation by explicitly defining the use of multiple instances of TLVs as the mechanism for scaling TLV contents, except where explicitly prohibited.

Today, as an example, the Extended IS Reachability TLV (22) [[RFC5305](#)] and MT Intermediate Systems TLV (222) [[RFC5120](#)] are TLVs where existing standards do not specify the behavior expected when multiple

copies of the TLV are present and no other mechanism for expanding the information carrying capacity of the TLV has been specified.

[RFC7356] has proposed a 16 bit length field for TLVs in flooding scoped Protocol Data Units (PDUs), but does nothing to address legacy implementations that do not yet implement the full breadth and scope of that RFC.

2. 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.

3. Procedure for Advertising Multiple TLV Instances

If the TLV contains information that specifies the applicability of its contents (i.e., a key), the key information MUST be replicated in additional TLV instances so that all contents specific to that key can be identified. This allows more sub-TLVs to be advertised for the same key information. The specification of the key for a given TLV is out of scope for this document.

This should be applied recursively. If a TLV is structured with sub-TLVs and sub-sub-TLVs, then replication of key information allows for the advertisement of additional sub-TLVs and sub-sub-TLVs for that key.

Note: If the fixed portion of a TLV includes information that is NOT part of the key and the non-key elements of the fixed portion of the TLV (metric and other bits in the control octet in the example below) differ, the values in the first TLV present in the lowest numbered LSP MUST be used.

As an example, consider the Extended IP Reachability TLV (type 135). A prefix in this TLV is specified by:

- * 4 octets of metric information
- * 1 octet of control information which includes 6 bits specifying the prefix length
- * 0-4 octets of IPv4 prefix

followed by up to 250 octets of sub-TLV information.

The key consists of the 6 bits of prefix length and the 0-4 octets of IPv4 prefix.

If this is insufficient sub-TLV space, then the node MAY advertise additional instances of the Extended IS Reachability TLV. The key information MUST be replicated identically and the additional sub-TLV space may be populated with additional information. The complete information for a given key in such cases is the joined set of all the carried information under the key in all the TLV instances.

[4.](#) Procedure for Receiving Multiple TLV Instances

A node that receives multiple TLV instances MUST accept all of the information in all of the instances. The order of arrival and placement of the TLV instances in LSP fragments MUST NOT change the behavior of the node once all the fragments have been received.

The contents of multiple TLV instances of the same type should be processed as if they were concatenated. If the internals of the TLV contain key information, then replication of the key information should be taken to indicate that subsequent data should be processed as if the subsequent data were concatenated.

Specifically, if a TLV is structured with sub-TLVs and key information is replicated, then the sub-TLVs should be processed as if they were concatenated.

This process should be applied recursively. If a TLV is structured with sub-TLVs and sub-sub-TLVs with replicated key information at all levels, then sub-sub-TLVs should be processed as if they were concatenated.

For example, suppose that a node received an LSP with multiple instances of the Extended IS Reachability TLV. The first instance contained key information K with sub-TLVs A, B, and C. The second instance contained key information K with sub-TLVs D, E, and F. The receiving node should then process this as having key information K and sub-TLVs A, B, C, D, E, F, or, because ordering is irrelevant, sub-TLVs D, E, F, A, B, C.

5. Operational Considerations

The generation of multiple TLVs for a given object is triggered by more information than will fit into a single TLV. If multiple TLVs are advertised in the presence of nodes which do not support multiple TLVs, operation of routing in the network is likely to be compromised.

6. IANA Considerations

This document makes no requests of IANA.

7. Security Considerations

This document creates no new security issues for IS-IS. Additional instances of existing TLVs expose no new information.

Security concerns for IS-IS are addressed in [[ISO10589](#)], [[RFC5304](#)], and [[RFC5310](#)].

8. References

8.1. Normative References

[ISO10589] ISO, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", August 1987, <ISO/IEC 10589:2002>.

[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>>.

- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", [RFC 5304](#), DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[8.2.](#) Informative References

- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.

- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", [RFC 7356](#), DOI 10.17487/RFC7356, September 2014, <<https://www.rfc-editor.org/info/rfc7356>>.

Authors' Addresses

Parag Kaneriya
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore 560103

Karnataka
India

Email: pkaneria@juniper.net

Tony Li,
Juniper Networks
1133 Innovation Way
Sunnyvale, California 94089
United States of America

Email: tony.li@tony.li

Antoni Przygienda,
Juniper Networks
1133 Innovation Way
Sunnyvale, California 94089
United States of America

Email: prz@juniper.net

Shraddha Hegde
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore 560103
Karnataka
India

Email: shraddha@juniper.net

Chris Bowers
Juniper Networks
1133 Innovation Way
Sunnyvale, California 94089
United States of America

Email: cbower@juniper.net

Les Ginsberg
Cisco Systems
821 Alder Drive
Milpitas, CA 95035
United States of America

Email: ginsberg@cisco.com