

A Policy Control Mechanism in IS-IS Using Administrative Tags

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document describes an extension to the IS-IS protocol to add operational capabilities that allow for ease of management and control over IP prefix distribution within an IS-IS domain. This document enhances the IS-IS protocol by extending the information that an Intermediate System (IS) router can place in Link State Protocol (LSP) Data Units for policy use. This extension will provide operators with a mechanism to control IP prefix distribution throughout multi-level IS-IS domains.

Table of Contents

| | |
|--|---|
| 1. Introduction | 2 |
| 2. Conventions Used in This Document | 2 |
| 3. Sub-TLV Additions | 2 |
| 3.1. 32-bit Administrative Tag Sub-TLV 1 | 3 |
| 3.2. 64-bit Administrative Tag Sub-TLV 2 | 3 |
| 4. Ordering of Tags | 3 |
| 5. Compliance | 4 |
| 6. Operations | 4 |
| 7. Security Considerations | 5 |
| 8. IANA Considerations | 5 |
| 9. Manageability Considerations | 5 |
| 10. Acknowledgements | 6 |
| 11. Contributors | 6 |
| 12. References | 6 |
| 12.1. Normative References | 6 |
| 12.2. Informative References | 6 |

1. Introduction

As defined in [RFC1195] and extended in [RFC3784], the IS-IS protocol [ISO10589] may be used to distribute IPv4 prefix reachability information throughout an IS-IS domain. In addition, thanks to extensions made in [RFC5120] and [ISIS-IPv6], IS-IS may be used to distribute IPv6 reachability information.

The IPv4 prefix information is encoded as TLV type 128 and 130 in [RFC1195], with additional information carried in TLV 135 as specified in [RFC3784] and TLV 235 as defined in [RFC5120]. In particular, the extended IP Reachability TLV (TLV 135) contains support for a larger metric space, an up/down bit to indicate redistribution between different levels in the hierarchy, an IP prefix, and one or more sub-TLVs that can be used to carry specific information about the prefix. TLV 235 is a derivative of TLV 135, with the addition of Multi-Topology membership information [RFC5120]. The IPv6 prefix information is encoded as TLV 236 in [ISIS-IPv6], and TLV 237 in [RFC5120].

This document defines 2 new sub-TLVs for TLV 135, TLV 235, TLV 236 and TLV 237 that may be used to carry administrative information about an IP prefix.

2. Conventions Used in This Document

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 BCP 14, [RFC2119].

3. Sub-TLV Additions

This document creates 2 new "Administrative Tag" sub-TLVs to be added to TLV 135, TLV 235, TLV 236 and TLV 237. These TLVs specify one or more 32- or 64-bit unsigned integers that may be associated with an IP prefix. Example uses of these tags include carrying BGP standard (or extended) communities and controlling redistribution between levels and areas, different routing protocols, or multiple instances of IS-IS running on the same router.

The methods for which their use is employed is beyond the scope of this document and left to the implementer and/or operator.

The encoding of the sub-TLV(s) is discussed in the following subsections.

3.1. 32-bit Administrative Tag Sub-TLV 1

The Administrative Tag SHALL be encoded as one or more 4-octet unsigned integers using Sub-TLV 1 in TLV 135 [RFC3784], TLV 235 [RFC5120], TLV 236 [ISIS-IPv6], and TLV 237 [RFC5120]. The Administrative Tag Sub-TLV has following structure:

- o 1 octet of type (value: 1)
- o 1 octet of length (value: multiple of 4)
- o one or more instances of 4 octets of administrative tag

On receipt, an implementation MAY consider only one encoded tag, in which case, the first encoded tag MUST be considered and any additional tags ignored. A tag value of zero is reserved and SHOULD be treated as "no tag".

3.2. 64-bit Administrative Tag Sub-TLV 2

The Administrative Tag SHALL be encoded as one or more 8-octet unsigned integers using Sub-TLV 2 in TLV 135 [RFC3784], TLV 235 [RFC5120], TLV 236 [ISIS-IPv6], and TLV 237 [RFC5120]. The 64-bit Administrative Tag Sub-TLV has following structure:

- o 1 octet of type (value: 2)
- o 1 octet of length (value: multiple of 8)
- o one or more instances of 8 octets of administrative tag

On receipt, an implementation MAY consider only one encoded tag; in which case, the first encoded tag MUST be considered and any additional tags ignored. A tag value of zero is reserved and SHOULD be treated as "no tag".

4. Ordering of Tags

The semantics of the tag order are implementation-dependent. That is, there is no implied meaning to the ordering of the tags that indicates a certain operation or set of operations need be performed based on the order of the tags. Each tag SHOULD be treated as an autonomous identifier that MAY be used in policy to perform a policy action. Whether or not tag A precedes or succeeds tag B SHOULD not change the meaning of the tag set. However, when propagating TLVs that contain multiple tags between levels, an implementation SHOULD

preserve the ordering such that the first tag remains the first tag, so that implementations that only recognize a single tag will have a consistent view across levels.

Each IS that receives an LSP with TLV(s) 135 and/or 235 and/or 236 and/or 237, that have associated sub-TLV(s) 1 and/or 2, MAY operate on the tag values as warranted by the implementation. If an implementation needs to change tag values, for example, when propagating TLVs between levels at an area boundary, then the TLV(s) SHOULD be copied to the newly generated Level-1 or Level-2 LSP. At that point, the contents of the sub-TLV(s) MAY change as dictated by the policy action. In the event that no change is required, the sub-TLV(s) SHOULD be copied in order into the new LSP, such that ordering is preserved.

5. Compliance

A compliant IS-IS implementation MUST be able to assign one tag to any IP prefix in any of the following TLVs: TLV 135, TLV 235, TLV 236, TLV 237. It MUST be able to interpret a single tag present in the sub-TLV, or the first tag where there is more than one tag present in the sub-TLV.

A compliant IS-IS implementation MAY be able to assign more than one tag to any IP prefix in any of the following TLVs: TLV 135, TLV 235, TLV 236, TLV 237. It MAY be able to interpret the second and subsequent tags where more than one tag is present in the sub-TLV.

When propagating TLVs between levels, a compliant IS-IS implementation MAY be able to rewrite or remove one or more tags associated with a prefix in any of the following TLVs: TLV 135, TLV 235, TLV 236, TLV 237.

6. Operations

An administrator associates an Administrative Tag value with some interesting property. When IS-IS advertises reachability for some IP prefix that has that property, it adds the Administrative Tag to the IP reachability information TLV for that prefix, and the tag "sticks" to the prefix as it is flooded throughout the routing domain.

Consider the network in Figure 1. We wish to "leak" L1 prefixes [RFC2966] with some property, A, from L2 to the L1 router R1. Without policy groups, there is no way for R2 to know property A prefixes from property B prefixes.




Figure 1: Example of usage

We associate Administrative Tag 100 with property A, and have R5 attach that value to the IP extended reachability information TLV for prefix 1.1.2.0/24. R2 has a policy in place to "match prefixes with Administrative Tag 100, and leak to L1".

The previous example is rather simplistic; it seems that it would be just as easy for R2 simply to match the prefix 1.1.2.0/24. However, if there are a large number of routers that need to apply some policy according to property A and a large number of "A" prefixes, this mechanism can be quite helpful.

Implementations that support only a single tag and those that support multiple tags may coexist in the same IS-IS domain. An implementation supporting multiple tags SHOULD therefore assign any tag that is required to be interpreted by all systems as the first tag in any set of multiple tags.

7. Security Considerations

This document raises no new security issues for IS-IS, as any annotations to IP prefixes should not pass outside the administrative control of the network operator of the IS-IS domain. Such an allowance would violate the spirit of Interior Gateway Protocols in general and IS-IS in particular.

8. IANA Considerations

IANA has assigned "1" as the type code of the 32-bit Administrative Tag Sub-TLV and "2" as the type code of the 64-bit Administrative Tag Sub-TLV.

9. Manageability Considerations

These extensions have been designed, developed, and deployed for many years and do not have any new impact on management and operation of the IS-IS protocol via this standardization process.

10. Acknowledgements

The authors would like to thank Henk Smit for clarifying the best place to describe this new information, Tony Li and Tony Przygienda for useful comments on this document, and Danny McPherson for some much needed formatting assistance.

11. Contributors

Brad Neal contributed portions of this document.

12. References

12.1. Normative References

- [ISO10589] International Organization for Standardization, "Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, Nov 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

12.2. Informative References

- [ISIS-IPv6] Hopps, C., "Routing IPv6 with IS-IS", Work in Progress, October 2007.
- [RFC2966] Li, T., Przygienda, T., and H. Smit, "Domain-wide Prefix Distribution with Two-Level IS-IS", RFC 2966, October 2000.
- [RFC3784] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, June 2004.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in IS-IS", RFC 5120, February 2008.

Authors' Addresses

Stefano Previdi
Cisco Systems
Via Del Serafico, 200
00142 Rome,
Italy

EMail: sprevidi@cisco.com

Mike Shand (editor)
Cisco Systems
250, Longwater Avenue.
Reading, Berks RG2 6GB
UK

Phone: +44 208 824 8690
EMail: mshand@cisco.com

Christian Martin
iPath Services

EMail: chris@ipath.net

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

