

Internet Engineering Task Force (IETF)
Request for Comments: 7775
Updates: 5308
Category: Standards Track
ISSN: 2070-1721

L. Ginsberg
Cisco Systems
S. Litkowski
Orange Business Service
S. Previdi
Cisco Systems
February 2016

IS-IS Route Preference for Extended IP and IPv6 Reachability

Abstract

In existing specifications, the route preferences for IPv4/IPv6 Extended Reachability TLVs are not explicitly stated. There are also inconsistencies in the definition of how the up/down bit applies to route preference when the prefix advertisement appears in Level 2 Link State Protocol Data Units (LSPs). This document addresses these issues.

This document updates RFC 5308.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7775>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. Use of the Up/Down Bit in Level 2 LSPs	3
3. Types of Routes in IS-IS Supported by Extended Reachability TLVs	4
3.1. Types of Routes Supported by TLVs 135 and 235	4
3.2. Types of Routes Supported by TLVs 236 and 237	6
3.3. Order of Preference for All Types of Routes Supported by TLVs 135 and 235	7
3.4. Order of Preference for All Types of Routes Supported by TLVs 236 and 237	8
4. Security Considerations	8
5. References	8
5.1. Normative References	8
5.2. Informative References	9
Appendix A. Example Interoperability Issue	10
Acknowledgements	11
Authors' Addresses	11

1. Introduction

[RFC5302] defines the route preference rules as they apply to TLVs 128 and 130. [RFC5305] introduced the IP Extended Reachability TLV 135 but did not explicitly adapt the route preference rules defined in [RFC5302] for the new TLV. [RFC5308] defines the IPv6 Reachability TLV 236 and does include an explicit statement regarding route preference -- but the statement introduces use of the up/down bit in advertisements that appear in Level 2 LSPs, which is inconsistent with statements made in [RFC5302] and [RFC5305]. This document defines explicit route preference rules for TLV 135, revises the route preference rules for TLV 236, and clarifies the usage of the up/down bit when it appears in TLVs in Level 2 LSPs. This document is a clarification (NOT a correction) of [RFC5302] and [RFC5305]; it is a correction of the route preference rules defined in [RFC5308] to be consistent with the rules for IPv4. It also makes explicit that the same rules apply to the Multi-Topology (MT) equivalent TLVs 235 and 237.

1.1. 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].

2. Use of the Up/Down Bit in Level 2 LSPs

The up/down bit was introduced in support of leaking prefixes downwards in the IS-IS level hierarchy. Routes that are leaked downwards have the bit set to 1. Such prefixes MUST NOT be leaked upwards in the hierarchy. So long as we confine ourselves to a single IS-IS instance and the current number of supported levels (two), it is impossible to have a prefix advertised in a Level 2 LSP and have the up/down bit set to 1. However, because [RFC5302] anticipated a future extension to IS-IS that might support additional levels, it allowed for the possibility that the up/down bit might be set in a Level 2 LSP and supported easy migration in the event such an extension was introduced. Section 3.3 of [RFC5302] states:

...it is RECOMMENDED that implementations ignore the up/down bit in L2 LSPs, and accept the prefixes in L2 LSPs regardless of whether the up/down bit is set.

[RFC5305] addressed an additional case wherein an implementation included support for multiple virtual routers running IS-IS in different areas. In such a case, it is possible to redistribute prefixes between two IS-IS instances in the same manner that prefixes are redistributed from other protocols into IS-IS. This introduced

the possibility that a prefix could be redistributed from Level 1 to Level 1 (as well as between Level 2 and Level 2), and in the event the redistributed route was leaked from Level 1 to Level 2, two different routers in different areas would be advertising the same prefix into the Level 2 sub-domain. To prevent this, Section 4.1 of [RFC5305] specifies:

If a prefix is advertised from one area to another at the same level, then the up/down bit SHALL be set to 1.

However, the statement in [RFC5302] that the up/down bit is ignored in Level 2 LSPs is not altered by [RFC5305].

The conclusion then is that there is no "L2 inter-area route"; indeed, no such route type is defined by [RFC5302]. However, [RFC5308] ignored this fact and introduced such a route type in Section 5 when it specified a preference for "Level 2 down prefix". This is an error that this document corrects. As changing the use of the up/down bit in TLVs 236 and 237 may introduce interoperability issues, implementors may wish to support transition mechanisms from the behavior described in [RFC5308] to the behavior described in this document.

3. Types of Routes in IS-IS Supported by Extended Reachability TLVs

[RFC5302] is the authoritative reference for the types of routes supported by TLVs 128 and 130. However, a number of attributes supported by those TLVs are NOT supported by TLVs 135, 235, 236, and 237. Distinction between internal/external metrics is not supported. In the case of IPv4 TLVs (135 and 235), the distinction between internal and external route types is not supported. However, the Prefix Attribute Flags sub-TLV defined in [PFXATTR] reintroduces the distinction between internal and external route types. The definitions below include references to the relevant attribute bits from [PFXATTR].

3.1. Types of Routes Supported by TLVs 135 and 235

This section defines the types of route supported for IPv4 when using TLV 135 [RFC5305] and/or TLV 235 [RFC5120]. The text follows as closely as possible the original text from [RFC5302].

L1 intra-area routes: These are advertised in L1 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 0. These IP prefixes are directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, both the X-Flag and the R-Flag are set to 0.

- L1 external routes: These are advertised in L1 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 0. These IP prefixes are learned from other protocols and are usually not directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, the X-Flag is set to 1, and the R-Flag is set to 0.
- L2 intra-area routes: These are advertised in L2 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 0. These IP prefixes are directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, both the X-Flag and the R-Flag are set to 0.
- L1->L2 inter-area routes: These are advertised in L2 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 0. These IP prefixes are learned via L1 routing and were derived during the L1 Shortest Path First (SPF) computation from prefixes advertised in L1 LSPs in TLV 135 or TLV 235. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 1.
- L2->L2 inter-area routes: These are advertised in L2 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 1 but is ignored and treated as if it were set to 0. These IP prefixes are learned from another IS-IS instance usually operating in another area. If the Prefix Attribute Flags sub-TLV is included, the X-Flag is set to 1, and the R-Flag is set to 0.
- L2 external routes: These are advertised in L2 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 0. These IP prefixes are learned from other protocols and are usually not directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, the X-Flag is set to 1, and the R-Flag is set to 0.
- L2->L1 inter-area routes: These are advertised in L1 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 1. These IP prefixes are learned via L2 routing and were derived during the L2 SPF computation from prefixes advertised in TLV 135 or TLV 235. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 1.
- L1->L1 inter-area routes: These are advertised in L1 LSPs, in TLV 135 or TLV 235. The up/down bit is set to 1. These IP prefixes are learned from another IS-IS instance usually operating in another area. If the Prefix Attribute Flags sub-TLV is included, the X-Flag is set to 1, and the R-Flag is set to 0.

3.2. Types of Routes Supported by TLVs 236 and 237

This section defines the types of route supported for IPv6 when using TLV 236 [RFC5308] and/or TLV 237 [RFC5120].

L1 intra-area routes: These are advertised in L1 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 0. The external bit is set to 0. These IPv6 prefixes are directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 0.

L1 external routes: These are advertised in L1 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 0. The external bit is set to 1. These IPv6 prefixes are learned from other protocols and are usually not directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 0.

L2 intra-area routes: These are advertised in L2 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 0. The external bit is set to 0. These IPv6 prefixes are directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 0.

L1->L2 inter-area routes: These are advertised in L2 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 0. The external bit is set to 0. These IPv6 prefixes are learned via L1 routing and were derived during the L1 Shortest Path First (SPF) computation from prefixes advertised in L1 LSPs in TLV 236 or TLV 237. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 1.

L2 external routes: These are advertised in L2 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 0. The external bit is set to 1. These IPv6 prefixes are learned from other protocols and are usually not directly connected to the advertising router. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 0.

L1->L2 external routes: These are advertised in L2 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 0. The external bit is set to 1. These IPv6 prefixes are learned via L1 routing and were derived during the L1 Shortest Path First (SPF) computation from L1 external routes advertised in L1 LSPs in TLV 236 or TLV 237. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 1.

L2->L2 inter-area routes: These are advertised in L2 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 1 but is ignored and treated as if it were set to 0. The external bit is set to 1. These IP prefixes are learned from another IS-IS instance usually operating in another area. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 0.

L2->L1 inter-area routes: These are advertised in L1 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 1. The external bit is set to 0. These IPv6 prefixes are learned via L2 routing and were derived during the L2 SPF computation from prefixes advertised in TLV 236 or TLV 237. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 1.

L2->L1 external routes: These are advertised in L1 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 1. The external bit is set to 1. These IPv6 prefixes are learned via L2 routing and were derived during the L2 SPF computation from prefixes advertised in TLV 236 or TLV 237. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 1.

L1->L1 inter-area routes: These are advertised in L1 LSPs, in TLV 236 or TLV 237. The up/down bit is set to 1. The external bit is set to 1. These IP prefixes are learned from another IS-IS instance usually operating in another area. If the Prefix Attribute Flags sub-TLV is included, the R-Flag is set to 0.

3.3. Order of Preference for All Types of Routes Supported by TLVs 135 and 235

This document defines the following route preferences for IPv4 routes advertised in TLVs 135 or 235. Note that all types of routes listed for a given preference are treated equally.

1. L1 intra-area routes; L1 external routes
2. L2 intra-area routes; L2 external routes; L1->L2 inter-area routes; L2-L2 inter-area routes
3. L2->L1 inter-area routes; L1->L1 inter-area routes

3.4. Order of Preference for All Types of Routes Supported by TLVs 236 and 237

This document defines the following route preferences for IPv6 routes advertised in TLVs 236 or 237. Note that all types of routes listed for a given preference are treated equally.

1. L1 intra-area routes; L1 external routes
2. L2 intra-area routes; L2 external routes; L1->L2 inter-area routes; L1-L2 external routes; L2-L2 inter-area routes
3. L2->L1 inter-area routes; L2->L1 external routes; L1->L1 inter-area routes

4. Security Considerations

This document raises no new security considerations. Security considerations for the IS-IS protocol are covered in [ISO10589], [RFC5304], and [RFC5310].

5. References

5.1. 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 Standard 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [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, <<http://www.rfc-editor.org/info/rfc5120>>.
- [RFC5302] Li, T., Smit, H., and T. Przygienda, "Domain-Wide Prefix Distribution with Two-Level IS-IS", RFC 5302, DOI 10.17487/RFC5302, October 2008, <<http://www.rfc-editor.org/info/rfc5302>>.

- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<http://www.rfc-editor.org/info/rfc5304>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<http://www.rfc-editor.org/info/rfc5308>>.
- [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, <<http://www.rfc-editor.org/info/rfc5310>>.

5.2. Informative References

- [PFXATTR] Ginsberg, L., Ed., Decraene, B., Previdi, S., Xu, X., and U. Chunduri, "IS-IS Prefix Attributes for Extended IP and IPv6 Reachability", Work in Progress, draft-ietf-isis-prefix-attributes-04, January 2016.

Appendix A. Example Interoperability Issue

This example documents a real-world interoperability issue that occurs because implementations from different vendors have interpreted the use of the up/down bit in Level 2 LSPs inconsistently.




Figure 1

In Figure 1, both R0 and R4 are advertising the prefix 10/8. Two IS-IS Level 2 instances are running on R3 to separate the network into two areas. R3 is performing route leaking and advertises prefixes from R4 to the other Level 2 process. The network is using extended metrics (TLV 135 defined in [RFC5305]). R0 advertises 10/8 with metric 2000, and R3 advertises 10/8 with metric 100. All links have a metric of 1. When advertising 10/8 in its Level 2 LSP, R3 sets the down bit as specified in [RFC5305].

R1, R2, and R3 are from three different vendors (R1->Vendor1, R2->Vendor2, R3->Vendor3). During interoperability testing, routing loops are observed in this scenario.

- o R2 has two possible paths to reach 10/8: Level 2 route with metric 2002 and up/down bit set to 0 (from R0) and Level 2 route with metric 101 and up/down bit set to 1 (from R3). R2 selects R1 as the next hop to 10/8 because it prefers the route that does NOT have the up/down bit set.
- o R3 has two possible paths to reach 10/8: Level 2 route with metric 2003 and up/down bit set to 0 (from R0) and Level 2 route with metric 101 and up/down bit set to 0 (from R4). R3 selects R4 as the next hop due to lowest metric.
- o R1 has two possible paths to reach 10/8: Level 2 route with metric 2001 and up/down bit set to 0 (from R0) and Level 2 route with metric 102 and up/down bit set to 1 (from R3). R1 selects R2 as the next hop due to lowest metric.

When R1 or R2 try to send traffic to 10/8, packets loop due to inconsistent routing decisions between R1 and R2.

Acknowledgements

The authors wish to thank Ahmed Bashandy for his insightful review.

Authors' Addresses

Les Ginsberg
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
United States

Email: ginsberg@cisco.com

Stephane Litkowski
Orange Business Service

Email: stephane.litkowski@orange.com

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

