

ISIS Working Group  
Internet-Draft  
Intended status: Informational  
Expires: December 29, 2014

S. Litkowski  
Orange Business Service  
June 27, 2014

IP Route preference specification issue  
draft-litkowski-isis-ip-route-preference-issue-00

## Abstract

This document details a potential specification issue in IP route preference in ISIS. As a consequence, some implementations does not interoperate correctly and leads to routing loops. Authors tries to analyse if we need to fix current specification.

## 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 [[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 <http://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 December 29, 2014.

## Copyright Notice

Copyright (c) 2014 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

Internet-Draft

isis-route-pref-pb

June 2014

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.

## Table of Contents

<a href="#">1.</a>	Problem statement . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Specification analysis . . . . .	<a href="#">3</a>
<a href="#">3.</a>	IPv6 and MT extensions cases . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Enhancing the <a href="#">RFC5305</a> . . . . .	<a href="#">5</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">5</a>
<a href="#">6.</a>	Acknowledgements . . . . .	<a href="#">5</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">5</a>
<a href="#">8.</a>	Normative References . . . . .	<a href="#">5</a>
	Author's Address . . . . .	<a href="#">6</a>

### [1.](#) Problem statement




Figure 1

Considering the figure 1, both R0 and R4 are advertising the prefix 10/8. Two ISIS L2 process are running on R3 to separate network. R3 is performing route-leaking and advertise prefix from R4 to the other L2 process. The network is using extended metrics (TLV135 defined in [\[RFC5305\]](#)). Considering that R0 is advertising 10/8 with metric 2000 and R3 with metric 100, and all the links have a metric of 1. When advertising route in L2, R3 set the down bit, according to [\[RFC5305\]](#).

During interoperability testings, authors observed some routing loops in this scenario.

R1, R2 and R3 are from three different vendors (R1->Vendor1, R2->Vendor2, R3->Vendor3).

- o R2 has two possible paths to reach 10/8, L2 up route with metric 2002 (from R0), and L2 down route with metric 101 (from R3). R2 selects R1 as nexthop to 10/8 because it is an Up route.

- o R3 has two possible paths to reach 10/8, L2 up route with metric 2002 from R1, L2 up metric 101 from R4. R3 selects R4 as nexthop due to lowest metric.

- o R1 has two possible paths to reach 10/8, L2 up route with metric 2001 from R0, L2 down metric 102 from R2. R1 selects R2 as nexthop due to lowest metric.

When R1 or R2 try to send traffic to 10/8, packets are looping due to inconsistent routing decision between R1 and R2.

## 2. Specification analysis

[RFC5305] defines IP extended reachability (TLV135), it also defines the notion of "Up/Down bit" that did not exist in [\[RFC1195\]](#) :

"The existing IP reachability TLVs (TLV type 128 and TLV type 130, defined in [\[RFC1195\]](#)) carry IP prefixes in a format that is analogous to the IS neighbor TLV from ISO 10589 [ISO-10589]. They carry four metrics, of which only the default metric is commonly used. The default metric has a possible range of 0-63. We would like to remove this restriction.

In addition, route redistribution (a.k.a. route leaking) has a key problem that was not fully addressed by the existing IP reachability TLVs. [\[RFC1195\]](#) allows a router to advertise prefixes upwards in the level hierarchy. Unfortunately, there were no mechanisms defined to advertise prefixes downwards in the level hierarchy.

To address these two issues, the proposed extended IP reachability TLV provides for a 32-bit metric and adds one bit to indicate that a prefix has been redistributed 'down' in the hierarchy."

[RFC5305] does not provide any rule for taking into account up/down bit in route preference.

[RFC5302] replaces [\[RFC2966\]](#) and defines extension to support optimal routing in multi-level environment. It especially defines up/down bit for IP prefix semantics defined in [\[RFC1195\]](#) (aka TLV 128, TLV130). [Section 3.2 of \[RFC5302\]](#) clearly specifies the order of

preference between IP route types in ISIS.

"Based on these assumptions, this document defines the following route preferences.

1. L1 intra-area routes with internal metric; L1 external routes with internal metric
2. L2 intra-area routes with internal metric; L2 external routes with internal metric; L1->L2 inter-area routes with internal metric; L1->L2 inter-area external routes with internal metric
3. L2->L1 inter-area routes with internal metric; L2->L1 inter-area external routes with internal metric
4. L1 external routes with external metric
5. L2 external routes with external metric; L1->L2 inter-area external routes with external metric
6. L2->L1 inter-area external routes with external metric"

It is quite clear that for IP Reachability defined in [[RFC1195](#)], up routes are preferred over down routes. It sounds that [[RFC5302](#)] does not apply to TLV defined in [[RFC5305](#)] : [section 5 of \[RFC5302\]](#) describes [[RFC5305](#)] as another proposal to deal with the issues described.

### [3.](#) IPv6 and MT extensions cases

[RFC5308] defines IPv6 Reachability extension for ISIS (TLV 236).

[Section 5](#) of the RFC clearly defines the order of preference between route types :

"The order of preference between paths for a given prefix MUST be modified to consider the up/down bit. The new order of preference is as follows (from best to worst).

1. Level 1 up prefix
2. Level 2 up prefix
3. Level 2 down prefix
4. Level 1 down prefix

If multiple paths have the same best preference, then selection occurs based on metric. Any remaining multiple paths SHOULD be considered for equal-cost multi-path routing if the router supports this; otherwise, the router can select any one of the multiple paths."

[RFC5120] defines Multitopology extension for ISIS and new IPv4 and IPv6 reachability TLVs (TLV 235 and 237). No guideline are provided in this RFC for route type preference but as MT extensions are based on basic TLVs (135 and 236), we expect the same behavior as for the associated TLVs.

#### [4.](#) Enhancing the [RFC5305](#)

[RFC5305] lacks of text regarding order of route preference compared to [[RFC5308](#)] and [[RFC5302](#)]. [[RFC5302](#)] does not seem to apply to TLV135 defined in [[RFC5305](#)]. As [[RFC5302](#)] and [[RFC5308](#)] are already aligned in term of behavior, authors propose to enhance [[RFC5305](#)] with a clear text stating the route preference order with the same behavior described in the two other specifications.

#### [5.](#) Security Considerations

There is no security consideration.

#### [6.](#) Acknowledgements

#### [7.](#) IANA Considerations

There is no IANA consideration.

## 8. Normative References

- [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.
- [RFC2966] Li, T., Przygienda, T., and H. Smit, "Domain-wide Prefix Distribution with Two-Level IS-IS", [RFC 2966](#), October 2000.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), February 2008.
- [RFC5302] Li, T., Smit, H., and T. Przygienda, "Domain-Wide Prefix Distribution with Two-Level IS-IS", [RFC 5302](#), October 2008.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), October 2008.

Litkowski

Expires December 29, 2014

[Page 5]

---

Internet-Draft

isis-route-pref-pb

June 2014

- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", [RFC 5308](#), October 2008.

## Author's Address

Stephane Litkowski  
Orange Business Service

Email: [stephane.litkowski@orange.com](mailto:stephane.litkowski@orange.com)

