

LSR Working Group
Internet Draft
Intended status: Standards Track
Expires: March 02, 2026

D. Li
Tsinghua University
C. Lin
New H3C Technologies
A. Lindem
Arrcus, Inc.
L. Liu
Zhongguancun Laboratory
X. Song
ZTE Corporation
September 02, 2025

IGP Reverse Prefix Metric
draft-li-lsr-igp-reverse-prefix-metric-03

Abstract

This document defines a method for calculating reverse paths by advertising reverse prefix costs. This method aims to solve the problem of strict RPF (Reverse Path Forwarding) check failure caused by mismatched bidirectional path costs in multi-area IGP scenarios.

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 March 02, 2026.

Copyright Notice

Copyright (c) 2025 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 (<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
1.1. Requirements Language.....	3
2. Usecase.....	3
3. Solution.....	4
4. Protocol Extension.....	5
4.1. Extension of OSPFv2 Reverse Prefix Cost.....	5
4.2. Extension of OSPFv3 Reverse Prefix Cost.....	6
4.3. Extension of IS-IS Reverse Prefix Cost.....	6
5. Security Considerations.....	7
6. IANA Considerations.....	7
6.1. OSPFv2 Extended Prefix TLV Sub-TLVs Registry.....	7
6.2. OSPFv3 Extended-LSA Sub-TLVs Registry.....	8
6.3. IS-IS Reverse Prefix Cost Sub-TLV.....	8
7. References.....	8
7.1. Normative References.....	8
Contributors.....	9
Authors' Addresses.....	10

1. Introduction

The process of strict RPF (Reverse Path Forwarding) checks involves verifying that a packet is received on the interface that matches the router's reverse path to the source address. If the cost to the source is inconsistent between the forward and reverse directions, the strict RPF check fails, resulting in the packet being discarded.

Another scenario involves running IGP multi-topology, where multicast traffic is usually situated within a separate topology. In this case, multicast also requires reverse path calculation.

This document defines a method for calculating reverse paths by advertising reverse prefix costs. This method aims to solve the problem of strict RPF check failure caused by mismatched bidirectional path costs in multi-area IGP scenarios.

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. Usecase

The strict RPF check process involves verifying that a packet is received on the interface that matches the router's reverse path to the source. If the cost between the forward and reverse path to the source is inconsistent, the strict RPF check will fail, resulting in the packet being discarded.

Therefore, when performing strict RPF checks, in cases where forward and reverse path costs are inconsistent, it is necessary to calculate the optimal path based on reverse path costs. This allows strict RPF checks to be conducted using the reverse optimal path.

Another scenario involves running multicast in an IGP multi-topology scenarios, where the cost of the multicast topology differs from that of the base topology. In multi-area scenarios, when the multicast topology requires reverse path calculation, the reverse cost between areas must also be considered.

Typically, IGP can advertise link information through the link-state database, which provides knowledge of both forward and reverse path costs within the domain. However, in multi-area scenarios, the situation is different. The following describes the situation in multi-area scenarios in detail:

In large-scale networks, an AS may be divided into different areas to avoid the problems caused by too many nodes. As shown in the following figure, an AS divided into two areas, each router is connected to the corresponding subnet, R1 is connected to P1, and so on, and R7 is connected to P7.

Taking OSPFv2 as an example, R4 and R5, as ABRs, will convert the router LSA (type-1) of R6 and R7 in Area1 into network Summary LSA (type-3) and advertise it to the routers in Area0. Area0 to Area1 are also processed in the same way.

The type-3 route received by R3 from R4 will include the subnet of R6, with an originator of R4 and a cost of 10. According to the method described in section 4, R3 will calculate the valid incoming interface of P6 as intf1.

If the cost of the two directions of the link between R4 and R6 is different for some reason, for example, the cost from R4 to R6 is 10, but the cost in the reverse direction is 100, which will cause the packet sent by R6 to arrive at R3 from intf2 actually. But the type-3 route advertised by R4 to R3 has only one-way cost from R4 to R6, which cannot reflect the real situation.

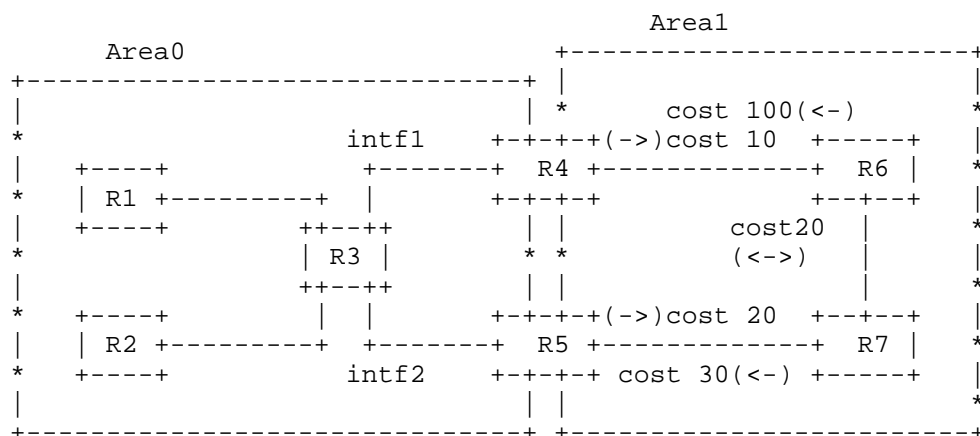


Figure 3: example topology of multi-area

3. Solution

In order to accurately calculate strict RPF in the scenarios of multi-area, it is necessary to expand the type-3 route and advertise the cost in reverse directions between ABR and a prefix at the same time. That is, when R4 advertises the prefix information of R6, it carries cost of 10 and reverse cost of 100 at the same time. Similarly, the cost of R6 network prefix information advertised by R5 in two directions is 40 and 50 respectively.

When ABR advertises network Summary LSA (type-3), ABR needs to calculate the total cost from the node where the prefix in LSA is located to this ABR, and advertises it together through protocol extension.

By extending the protocol, R3 can be aware that packets from P6 and P7 will arrive at R3 from R5, so the valid incoming interface of the two protected prefixes can be calculated as intf2.

After the AS is divided into different areas, in order to reduce routing messages, the ABR may aggregate the routing information with the same prefix and only publish one route to other areas. If the

forward or reverse path costs of the aggregated prefixes are different, after advertising the aggregated route, the ABR also needs to separately advertise a route for the prefixes with different costs, and advertise the forward and reverse costs corresponding to this prefix in this route.

4. Protocol Extension

4.1. Extension of OSPFv2 Reverse Prefix Cost

A sub-TLV called Prefix-Reverse-cost sub-TLV is defined to carry the total costs from the router where the prefix is located to reach ABR.

The Prefix-Reverse-cost Sub-TLV is a sub-TLV of the OSPF Extended Prefix TLV described in [RFC7684].

When the Route Type of OSPFv2 Extended Prefix TLV is Inter-Area (3), Prefix-Reverse-cost sub-TLV can be used.

For Multi-Topology support, the TOS field is redefined as MT-ID in the payload of Router, Summary, and Type-5 and Type-7 AS-external LSAs [RFC4915].

It SHOULD appear only once in the parent TLV and has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type																				Length																			
reverse metric																																							

where:

Type: TBD2.

Length: 4.

Reverse metric: Total cost value from the router where the prefix is located to ABR.

4.2. Extension of OSPFv3 Reverse Prefix Cost

A sub-TLV called Prefix-Reverse-cost sub-TLV is defined to carry the total cost from the router where the prefix is located to reach ABR.

The Prefix-Reverse-cost sub-TLV is a sub-TLV of the following OSPFv3 TLVs as defined in [RFC8362] and in Section 5:

Inter-Area Prefix TLV

It SHOULD appear only once in the parent TLV and has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type										Length																													
Reverse metric																																							

where:

Type: TBD4.

Length: 4.

Reverse metric: Total cost value from the router where the prefix is located to ABR.

4.3. Extension of IS-IS Reverse Prefix Cost

A sub-TLV called Prefix-Reverse-cost sub-TLV is defined to carry the total cost from the router where the prefix is located to reach ABR.

The Prefix-Reverse-cost sub-TLV is a sub-TLV of the following of the following IS-IS TLVs:

TLV-135 (Extended IPv4 reachability) defined in [RFC5305].

TLV-235 (Multi-topology IPv4 Reachability) defined in [RFC5120].

TLV-236 (IPv6 IP Reachability) defined in [RFC5308].

TLV-237 (Multi-topology IPv6 IP Reachability) defined in [RFC5120].

When the level 2 router leaks routes through the above TLVs, Prefix-Reverse-cost sub-TLV can be used to carry reverse total cost.

It SHOULD appear only once in the parent TLV and has the following format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Type   |      Length   |      Reserved   |
+-----+-----+-----+-----+-----+-----+-----+
|                                     Reverse metric                                     |
+-----+-----+-----+-----+-----+-----+-----+

```

where:

Type: TBD6.

Length: 6.

Reserved: SHOULD be set to 0 on transmission and MUST be ignored on reception

Reverse metric: Total cost value from the router where the prefix is located to ABR.

5. Security Considerations

This document does not introduce any new security consideration.

6. IANA Considerations

6.1. OSPFv2 Extended Prefix TLV Sub-TLVs Registry

The following values have been allocated:

Value	Description	Reference
TBD	Prefix-Reverse-cost	This document

Table 1: OSPFv2 Extended Prefix TLV Sub-TLVs

6.2. OSPFv3 Extended-LSA Sub-TLVs Registry

The following values have been allocated:

Value	Description	Reference
TBD	Prefix-Reverse-cost	This document

Table 2: OSPFv3 Extended Prefix TLV Sub-TLVs

6.3. IS-IS Reverse Prefix Cost Sub-TLV

This document makes the following registrations in the "Sub-TLVs for TLV 135, 235, 236, and 237" registry.

Type	Description	135	235	236	237
TBD	Prefix-Reverse-cost	Y	Y	Y	Y

Table 3: IS-IS Prefix-Reverse-cost Sub-TLVs

7. References

7.1. Normative 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>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [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>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.

[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>>.

[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>>.

Contributors

Yuanxiang Qiu
New H3C Technologies
China
Email: qiuyuanxiang@h3c.com

Authors' Addresses

Dan Li
Tsinghua University
Beijing
China
Email: toliidan@tsinghua.edu.cn

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

Acee Lindem
Arrcus, Inc.
United States of America
Email: acee.ietf@gmail.com

Libin Liu
Zhongguancun Laboratory
Beijing
China
Email: liulb@zgclab.edu.cn

Xueyan Song
ZTE Corporation
China
Email: song.xueyan2@zte.com.cn

