

Internet Engineering Task Force (IETF)
Request for Comments: 8687
Updates: 5786
Category: Standards Track
ISSN: 2070-1721

A. Smirnov
Cisco Systems, Inc.
A. Retana
Futurewei Technologies, Inc.
M. Barnes
November 2019

OSPF Routing with Cross-Address Family Traffic Engineering Tunnels

Abstract

When using Traffic Engineering (TE) in a dual-stack IPv4/IPv6 network, the Multiprotocol Label Switching (MPLS) TE Label Switched Path (LSP) infrastructure may be duplicated, even if the destination IPv4 and IPv6 addresses belong to the same remote router. In order to achieve an integrated MPLS TE LSP infrastructure, OSPF routes must be computed over MPLS TE tunnels created using information propagated in another OSPF instance. This issue is solved by advertising cross-address family (X-AF) OSPF TE information.

This document describes an update to RFC 5786 that allows for the easy identification of a router's local X-AF IP addresses.

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 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8687>.

Copyright Notice

Copyright (c) 2019 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 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

1. Introduction
2. Requirements Language
3. Operation
4. Backward Compatibility
 - 4.1. Automatically Switched Optical Networks
5. Security Considerations
6. IANA Considerations

7. References

7.1. Normative References

7.2. Informative References

Acknowledgements

Authors' Addresses

1. Introduction

TE extensions to OSPFv2 [RFC3630] and OSPFv3 [RFC5329] have been described to support intra-area TE in IPv4 and IPv6 networks, respectively. In both cases, the TE database provides a tight coupling between the routed protocol and advertised TE signaling information. In other words, any use of the TE database is limited to IPv4 for OSPFv2 [RFC2328] and IPv6 for OSPFv3 [RFC5340].

In a dual-stack network, it may be desirable to set up common MPLS TE LSPs to carry traffic destined to addresses from different address families on a router. The use of common LSPs eases potential scalability and management concerns by halving the number of LSPs in the network. Besides, it allows operators to group traffic based on business characteristics, class of service, and/or applications; the operators are not constrained by the network protocol used.

For example, an LSP created based on MPLS TE information propagated by an OSPFv2 instance can be used to transport both IPv4 and IPv6 traffic, as opposed to using both OSPFv2 and OSPFv3 to provision a separate LSP for each address family. Even if, in some cases, the address-family-specific traffic is to be separated, calculation from a common TE database may prove to be operationally beneficial.

During the SPF calculation on the TE tunnel head-end router, OSPF computes shortcut routes using TE tunnels. A commonly used algorithm for computing shortcuts is defined in [RFC3906]. For that or any similar algorithm to work with a common MPLS TE infrastructure in a dual-stack network, a requirement is to reliably map the X-AF addresses to the corresponding tail-end router. This mapping is a challenge because the Link State Advertisements (LSAs) containing the routing information are carried in one OSPF instance, while the TE calculations may be done using a TE database from a different OSPF instance.

A simple solution to this problem is to rely on the Router ID to identify a node in the corresponding OSPFv2 and OSPFv3 Link State Databases (LSDBs). This solution would mandate both instances on the same router to be configured with the same Router ID. However, relying on the correctness of configuration puts additional burden and cost on the operation of the network. The network becomes even more difficult to manage if OSPFv2 and OSPFv3 topologies do not match exactly, for example, if area borders are chosen differently in the two protocols. Also, if the routing processes do fall out of sync (e.g., having different Router IDs for local administrative reasons), there is no defined way for other routers to discover such misalignment and to take corrective measures (such as to avoid routing traffic through affected TE tunnels or alerting the network administrators). The use of misaligned Router IDs may result in delivering the traffic to the wrong tail-end router, which could lead to suboptimal routing or even traffic loops.

This document describes an update to [RFC5786] that allows for the easy identification of a router's local X-AF IP addresses. [RFC5786] defined the Node IPv4 Local Address and Node IPv6 Local Address sub-TLVs of the Node Attribute TLV for a router to advertise additional local IPv4 and IPv6 addresses. However, [RFC5786] did not describe the advertisement and usage of these sub-TLVs when the address family of the advertised local address differed from the address family of the OSPF traffic engineering protocol.

This document updates [RFC5786] so that a router can also announce one or more local X-AF addresses using the corresponding Local Address sub-TLV. Routers using the Node Attribute TLV [RFC5786] can include non-TE-enabled interface addresses in their OSPF TE advertisements and also use the same sub-TLVs to carry X-AF information, facilitating the mapping described above.

The method specified in this document can also be used to compute the X-AF mapping of the egress Label Switching Router (LSR) for sub-LSPs of a Point-to-Multipoint LSP [RFC4461]. Considerations of using Point-to-Multipoint MPLS TE for X-AF traffic forwarding is outside the scope of this document.

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. Operation

To implement the X-AF routing technique described in this document, OSPFv2 will advertise the Node IPv6 Local Address sub-TLV and OSPFv3 will advertise the Node IPv4 Local Address sub-TLV, possibly in addition to advertising other IP addresses as documented by [RFC5786].

Multiple instances of OSPFv3 are needed if it is used for both IPv4 and IPv6 [RFC5838]. The operation in this section is described with OSPFv2 as the protocol used for IPv4; that is the most common case. The case of OSPFv3 being used for IPv4 follows the same procedure as what is indicated for OSPFv2 below.

On a node that implements X-AF routing, each OSPF instance advertises, using the Node Local Address sub-TLV, all X-AF IPv6 (for OSPFv2 instance) or IPv4 (for OSPFv3) addresses local to the router that can be used by the Constrained Shortest Path First (CSPF) to calculate MPLS TE LSPs:

- * The OSPF instance MUST advertise the IP address listed in the Router Address TLV [RFC3630] [RFC5329] of the X-AF instance maintaining the TE database.
- * The OSPF instance SHOULD include additional local addresses advertised by the X-AF OSPF instance in its Node Local Address sub-TLVs.
- * An implementation MAY advertise other local X-AF addresses.

When TE information is advertised in an OSPF instance, both natively (i.e., as per RFC [RFC3630] or [RFC5329]) and as X-AF Node Attribute TLV, it is left to local configuration to determine which TE database is used to compute routes for the OSPF instance.

On Area Border Routers (ABRs), each advertised X-AF IP address MUST be advertised into, at most, one area. If OSPFv2 and OSPFv3 ABRs coincide (i.e., the areas for all OSPFv2 and OSPFv3 interfaces are the same), then the X-AF addresses MUST be advertised into the same area in both instances. This allows other ABRs connected to the same set of areas to know with which area to associate computed MPLS TE tunnels.

During the X-AF routing calculation, X-AF IP addresses are used to

map locally created LSPs to tail-end routers in the LSDB. The mapping algorithm can be described as:

Walk the list of all MPLS TE tunnels for which the computing router is a head end. For each MPLS TE tunnel T:

1. If T's destination address is from the same address family as the OSPF instance associated with the LSDB, then the extensions defined in this document do not apply.
2. Otherwise, it is a X-AF MPLS TE tunnel. Note the tunnel's destination IP address.
3. Walk the X-AF IP addresses in the LSDBs of all connected areas. If a matching IP address is found, advertised by router R in area A, then mark the tunnel T as belonging to area A and terminating on tail-end router R. Assign the intra-area SPF cost to reach router R within area A as the IGP cost of tunnel T.

After completing this calculation, each TE tunnel is associated with an area and tail-end router in terms of the routing LSDB of the computing OSPF instance and has a cost.

The algorithm described above is to be used only if the Node Local Address sub-TLV includes X-AF information.

Note that, for clarity of description, the mapping algorithm is specified as a single calculation. Implementations may choose to support equivalent mapping functionality without implementing the algorithm as described.

As an example, consider a router in a dual-stack network using OSPFv2 and OSPFv3 for IPv4 and IPv6 routing, respectively. Suppose the OSPFv2 instance is used to propagate MPLS TE information and the router is configured to accept TE LSPs terminating at local addresses 198.51.100.1 and 198.51.100.2. The router advertises in OSPFv2 the IPv4 address 198.51.100.1 in the Router Address TLV, the additional local IPv4 address 198.51.100.2 in the Node IPv4 Local Address sub-TLV, and other TE TLVs as required by [RFC3630]. If the OSPFv3 instance in the network is enabled for X-AF TE routing (that is, to use MPLS TE LSPs computed by OSPFv2 for IPv6 routing), then the OSPFv3 instance of the router will advertise the Node IPv4 Local Address sub-TLV listing the local IPv4 addresses 198.51.100.1 and 198.51.100.2. Other routers in the OSPFv3 network will use this information to reliably identify this router as the egress LSR for MPLS TE LSPs terminating at either 198.51.100.1 or 198.51.100.2.

4. Backward Compatibility

Only routers that serve as endpoints for one or more TE tunnels MUST be upgraded to support the procedures described herein:

- * Tunnel tail-end routers advertise the Node IPv4 Local Address sub-TLV and/or the Node IPv6 Local Address sub-TLV.
- * Tunnel head-end routers perform the X-AF routing calculation.

Both the endpoints MUST be upgraded before the tail end starts advertising the X-AF information. Other routers in the network do not need to support X-AF procedures.

4.1. Automatically Switched Optical Networks

[RFC6827] updates [RFC5786] by defining extensions to be used in an Automatically Switched Optical Network (ASON). The Local TE Router

ID sub-TLV is required for determining ASON reachability. The implication is that if the Local TE Router ID sub-TLV is present in the Node Attribute TLV, then the procedures in [RFC6827] apply, regardless of whether any X-AF information is advertised.

5. Security Considerations

This document describes the use of the Local Address sub-TLVs to provide X-AF information. The advertisement of these sub-TLVs, in any OSPF instance, is not precluded by [RFC5786]. As such, no new security threats are introduced beyond the considerations in OSPFv2 [RFC2328], OSPFv3 [RFC5340], and [RFC5786].

The X-AF information is not used for SPF computation or normal routing, so the mechanism specified here has no effect on IP routing. However, generating incorrect information or tampering with the sub-TLVs may have an effect on traffic engineering computations. Specifically, TE traffic may be delivered to the wrong tail-end router, which could lead to suboptimal routing, traffic loops, or exposing the traffic to attacker inspection or modification. These threats are already present in other TE-related specifications, and their considerations apply here as well, including [RFC3630] and [RFC5329].

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. Normative References

- [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>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, DOI 10.17487/RFC5329, September 2008, <<https://www.rfc-editor.org/info/rfc5329>>.
- [RFC5786] Aggarwal, R. and K. Kompella, "Advertising a Router's Local Addresses in OSPF Traffic Engineering (TE) Extensions", RFC 5786, DOI 10.17487/RFC5786, March 2010, <<https://www.rfc-editor.org/info/rfc5786>>.
- [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>>.

7.2. Informative References

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [RFC3906] Shen, N. and H. Smit, "Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels", RFC 3906, DOI 10.17487/RFC3906, October 2004, <<https://www.rfc-editor.org/info/rfc3906>>.

- [RFC4461] Yasukawa, S., Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, DOI 10.17487/RFC4461, April 2006, <<https://www.rfc-editor.org/info/rfc4461>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [RFC5838] Lindem, A., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, DOI 10.17487/RFC5838, April 2010, <<https://www.rfc-editor.org/info/rfc5838>>.
- [RFC6827] Malis, A., Ed., Lindem, A., Ed., and D. Papadimitriou, Ed., "Automatically Switched Optical Network (ASON) Routing for OSPFv2 Protocols", RFC 6827, DOI 10.17487/RFC6827, January 2013, <<https://www.rfc-editor.org/info/rfc6827>>.

Acknowledgements

The authors would like to thank Peter Psenak and Eric Osborne for early discussions and Acee Lindem for discussing compatibility with ASON extensions. Also, Eric Vyncke, Ben Kaduk, and Roman Danyliw provided useful comments.

We would also like to thank the authors of RFC 5786 for laying down the foundation for this work.

Authors' Addresses

Anton Smirnov
Cisco Systems, Inc.
De Kleetlaan 6a
1831 Diegem
Belgium

Email: as@cisco.com

Alvaro Retana
Futurewei Technologies, Inc.
2330 Central Expressway
Santa Clara, CA 95050
United States of America

Email: alvaro.retana@futurewei.com

Michael Barnes

Email: michael_barnes@usa.net