

OSPF
Internet-Draft
Updates: 5786 (if approved)
Intended status: Standards Track
Expires: July 14, 2017

A. Smirnov
A. Retana
M. Barnes
Cisco Systems, Inc.
January 10, 2017

OSPF Routing with Cross-Address Family MPLS Traffic Engineering Tunnels
draft-smirnov-ospf-xaf-te-07

Abstract

When using Traffic Engineering (TE) in a dual-stack IPv4/IPv6 network the Multiprotocol Label Switching (MPLS) TE Label Switched Paths (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 is solved by advertising cross-address family (X-AF) OSPF TE information.

This document describes an update to RFC5786 that allows for the easy identification of a router's local X-AF IP addresses.

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 July 14, 2017.

Copyright Notice

Copyright (c) 2017 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.

Table of Contents

1. Introduction	2
2. Requirements Language	3
3. Operation	3
4. Backward Compatibility	5
5. Security Considerations	6
6. IANA Considerations	6
7. Acknowledgements	6
8. References	6
8.1. Normative References	6
8.2. Informative References	7
Authors' Addresses	7

1. Introduction

TE Extensions to OSPFv2 [RFC3630] and to 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 TE signaling information in it. In other words, any use of the TE link state 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 and/or applications or class of service, not constrained by the network protocol which carries it.

For example, an LSP created based on MPLS TE information propagated by OSPFv2 instance can be defined to carry both IPv4 and IPv6 traffic, instead of having 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, the calculation from a common database may prove operationally beneficial.

A requirement when creating a common MPLS TE infrastructure is the ability to reliably map the X-AF family addresses to the

corresponding advertising tail-end router. This mapping is a challenge because the LSAs containing the routing information are carried in one OSPF instance while the TE calculation may be done using a TE database from a different instance.

A simple solution to this problem is to rely on the Router ID to identify a node in the corresponding OSPFv2 and OSPFv3 databases. This solution would mandate both instances on the same router to be configured with the same Router ID. However, relying on the correctness of the configuration puts additional burden on network management and adds cost to 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 drawn differently in the two protocols. Also, if the routing processes do fall out of sync (having different Router IDs, even if for local administrative reasons), there is no defined way for other routers to discover such misalignment and to take any corrective measures (such as to avoid routing through affected TE tunnels or issuing warning to network management). 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. 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 mentioned above.

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

2. 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].

3. Operation

[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. To solve the problem outlined in [RFC5786] OSPFv2 would advertise and use only IPv4 addresses and OSPFv3 would advertise and use only IPv6 addresses.

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. In other words, to implement the X-AF routing technique proposed 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].

A node that implements X-AF routing SHOULD advertise in the corresponding Node Local Address sub-TLV all X-AF IP addresses local to the router that can be used by Constrained SPF (CSPF) to calculate MPLS TE LSPs. In general, OSPF SHOULD advertise the IP address listed in the Router Address TLV of the X-AF instance maintaining MPLS TE database plus any additional local addresses advertised by the X-AF OSPF instance in its Node Local Address sub-TLV. Implementation MAY advertise other local X-AF addresses.

If the Node Attribute TLV carries both the Node IPv4 Local Address sub-TLV and the Node IPv6 Local Address sub-TLV, then the X-AF component must be considered for the consolidated calculation of MPLS TE LSPs. Both instances may carry the required information, it is left to local configuration to determine which database is used.

On Area Border Routers (ABR), each advertised X-AF IP address MUST be advertised into at most one area. If OSPFv2 and OSPFv3 area borders match (i.e. for each interface area number for OSPFv2 and OSPFv3 instances is numerically equal), 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 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 IP address is from the same address family as the computing OSPF instance, then the tunnel must have been signaled based on MPLS TE information propagated in the same OSPF instance. Process the tunnel as per [RFC3630] or [RFC5329].
2. Otherwise it is a X-AF MPLS TE tunnel. Note 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 an 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 metric.

Note that for clarity of description the mapping algorithm is specified as a single calculation. Actual implementations for the efficiency may choose to support equivalent mapping functionality without implementing the algorithm exactly as it is described.

As an example lets consider a router in dual-stack network running OSPFv2 and OSPFv3 for IPv4 and IPv6 routing correspondingly. Suppose 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. Then the router will advertise into OSPFv2 instance IPv4 address 198.51.100.1 in the Router Address TLV, additional local IPv4 address 198.51.100.2 in the Node IPv4 Local Address sub-TLV, plus other Traffic Engineering TLVs as required by [RFC3630]. If OSPFv3 instance in the network is enabled for X-AF TE routing (that is, to use for IPv6 routing MPLS TE LSPs computed by OSPFv2), then the OSPFv3 instance of the router will advertise the Node IPv4 Local Address sub-TLV listing 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 egress LSR for MPLS TE LSPs terminating at either 198.51.100.1 or 198.51.100.2.

4. Backward Compatibility

Node Attribute TLV and Node Local Address sub-TLVs and their usage are defined in [RFC5786] and updated by [RFC6827]. Way of using these TLVs as specified in this document is fully backward compatible with previous standard documents.

An implementation processing Node Attribute TLV MUST interpret its content as follows:

- o If the Node Attribute TLV contains Local TE Router ID sub-TLV then this Node Attribute TLV MUST be treated as carrying routing information for ASON (Automatically Switched Optical Network) and processed as specified in [RFC6827].
- o Otherwise Node Attribute TLV contains one or more instance(s) of Node IPv4 Local Address and/or Node IPv6 Local Address sub-TLVs.

Meaning of each Local Address sub-TLV has to be identified separately.

- * If Node Local Address sub-TLV belongs to the same address family as instance of OSPF protocol advertising it then address carried in the sub-TLV MUST be treated as described in [RFC5786].
- * Otherwise the address is used for X-AF tunnel tail-end mapping as defined by this document.

5. Security Considerations

This document introduces no new security concerns. Security considerations of using Node Attribute TLV are discussed in [RFC5786].

6. IANA Considerations

This document has no IANA actions.

7. Acknowledgements

The authors would like to thank Peter Psenak and Eric Osborne for early discussions and Acee Lindem for discussing compatibility with ASON extensions.

We would also like to thank the authors of RFC5786 for laying down the foundation for this work.

8. References

8.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [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, <<http://www.rfc-editor.org/info/rfc5786>>.

8.2. Informative References

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [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, <<http://www.rfc-editor.org/info/rfc4461>>.
- [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, <<http://www.rfc-editor.org/info/rfc5329>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<http://www.rfc-editor.org/info/rfc5340>>.
- [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, <<http://www.rfc-editor.org/info/rfc6827>>.

Authors' Addresses

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

Email: as@cisco.com

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA

Email: aretana@cisco.com

Michael Barnes
Cisco Systems, Inc.
510 McCarthy Blvd.
Milpitas, CA 95035
USA

Email: mjbarnes@cisco.com