

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 13, 2015

X. Xu
Huawei
S. Kini
Ericsson
S. Sivabalan
C. Filsfils
Cisco
S. Litkowski
Orange
October 10, 2014

**Signaling Entropy Label Capability Using OSPF
draft-xu-ospf-mpls-elc-01**

Abstract

Multi Protocol Label Switching (MPLS) has defined a mechanism to load balance traffic flows using Entropy Labels (EL). An ingress LSR cannot insert ELs for packets going into a given tunnel unless an egress LSR has indicated via signaling that it can process ELs on that tunnel. This draft defines a mechanism to signal that capability using OSPF. This mechanism is useful when the label advertisement is also done via OSPF.

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 April 13, 2015.

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 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](#)
 [1.1.](#) Requirements Language [3](#)
[2.](#) Terminology [3](#)
[3.](#) Advertising ELC Using OSPF [3](#)
[4.](#) Advertising RLSDC Using OSPF [3](#)
[5.](#) Acknowledgements [3](#)
[6.](#) IANA Considerations [4](#)
[7.](#) Security Considerations [4](#)
[8.](#) References [4](#)
 [8.1.](#) Normative References [4](#)
 [8.2.](#) Informative References [4](#)
 Authors' Addresses [5](#)

[1.](#) Introduction

Multi Protocol Label Switching (MPLS) has defined a method in [[RFC6790](#)] to load balance traffic flows using Entropy Labels (EL). An ingress LSR cannot insert ELs for packets going into a given tunnel unless an egress LSR has indicated that it can process ELs on that tunnel. [[RFC6790](#)] defines the signaling of this capability (a.k.a Entropy Label Capability - ELC) via signaling protocols. Recently, mechanisms are being defined to signal labels via link state Interior Gateway Protocols (IGP) such as OSPF [[I-D.ietf-ospf-segment-routing-extensions](#)]. In such scenario the signaling mechanisms defined in [[RFC6790](#)] are inadequate. This draft defines a mechanism to signal the ELC using OSPF. This mechanism is useful when the label advertisement is also done via OSPF. In addition, in the cases where stacked LSPs are used for whatever reasons (e.g., SPRING-MPLS [[I-D.gredler-spring-mpls](#)] [[I-D.filsfils-spring-segment-routing-mpls](#)]), it would be useful for ingress LSRs to know each LSR's capability of reading the maximum label stack depth. This capability, referred to as Readable Label Stack Depth Capability (RLSDC) can be used by ingress LSRs to determine whether it's necessary to insert an EL for a given LSP tunnel in the case where there has already been at least one EL in the label stack [[I-D.kini-mpls-spring-entropy-label](#)]. Of course,

even it has been determined that it's necessary to insert an EL for a given LSP tunnel, if the egress LSR of that LSP tunnel has not yet indicated that it can process ELs for that tunnel, the ingress LSR MUST NOT include an entropy label for that tunnel as well.

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. Terminology

This memo makes use of the terms defined in [[RFC6790](#)] and [[RFC4970](#)].

3. Advertising ELC Using OSPF

The OSPF Router Information (RI) Opaque LSA defined in [[RFC4970](#)] is used by OSPF routers to announce their capabilities. A new TLV within the body of this LSA, called ELC TLV is defined to advertise the capability of the router to process the ELs. It is formatted as described in [Section 2.1 of \[RFC4970\]](#). This TLV is applicable to both OSPFv2 and OSPFv3. The Type for the ELC TLV needs to be assigned by IANA and it has a Length of zero. The scope of the advertisement depends on the application but it is recommended that it SHOULD be AS-scoped.

4. Advertising RLSDC Using OSPF

A new TLV within the body of the OSPF RI LSA, called RLSDC TLV is defined to advertise the capability of the router to read the maximum label stack depth. It is formatted as described in [Section 2.1 of \[RFC4970\]](#) with a Type code to be assigned by IANA and a Length of one. The Value field is set to the maximum readable label stack depth in the range between 1 to 255. The scope of the advertisement depends on the application but it is RECOMMENDED that it SHOULD be domain-wide. If a router has multiple linecards with different capabilities of reading the maximum label stack depth, the router MUST advertise the smallest one in the RLSDC TLV. This TLV is applicable to both OSPFv2 and OSPFv3.

5. Acknowledgements

The authors would like to thank Yimin Shen and George Swallow for their comments.

6. IANA Considerations

This memo includes a request to IANA to allocate two TLV types from the OSPF RI TLVs registry.

7. Security Considerations

This document does not introduce any new security risk.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 4970](#), July 2007.

8.2. Informative References

- [I-D.filsfils-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing with MPLS data plane", [draft-filsfils-spring-segment-routing-mpls-03](#) (work in progress), August 2014.
- [I-D.gredler-spring-mpls]
Gredler, H., Rekhter, Y., Jalil, L., Kini, S., and X. Xu, "Supporting Source/Explicitly Routed Tunnels via Stacked LSPs", [draft-gredler-spring-mpls-06](#) (work in progress), May 2014.
- [I-D.ietf-ospf-segment-routing-extensions]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [draft-ietf-ospf-segment-routing-extensions-02](#) (work in progress), August 2014.
- [I-D.kini-mpls-spring-entropy-label]
Kini, S., Kompella, K., Sivabalan, S., Litkowski, S., Shakir, R., Xu, X., Henderickx, W., and J. Tantsura, "Entropy labels for source routed stacked tunnels", [draft-kini-mpls-spring-entropy-label-01](#) (work in progress), September 2014.

[RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", [RFC 6790](#), November 2012.

Authors' Addresses

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

Sriganesh Kini
Ericsson

Email: sriganesh.kini@ericsson.com

Siva Sivabalan
Cisco

Email: msiva@cisco.com

Clarence Filsfils
Cisco

Email: cfilsfil@cisco.com

Stephane Litkowski
Orange

Email: stephane.litkowski@orange.com

