

L  
Internet-Draft  
Intended status: Standards Track  
Expires: July 2, 2020

P. Shaofu  
C. Ran  
ZTE Corporation  
G. Mirsky  
ZTE Corp.  
December 30, 2019

**IGP Flexible Algorithm with L2bundles  
draft-peng-lsr-flex-algo-l2bundles-00**

**Abstract**

IGP Flex Algorithm proposes a solution that allows IGPs themselves to compute constraint based paths over the network, and it also specifies a way of using Segment Routing (SR) Prefix-SIDs and SRv6 locators to steer packets along the constraint-based paths. This document describes how to create Flex-algo plane with L2bundles scenario.

**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 July 2, 2020.

**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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Requirements Language . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Color set on L2 Bundle Member . . . . .	<a href="#">3</a>
<a href="#">4.</a>	Flex-algo plane with L2 link resource . . . . .	<a href="#">3</a>
<a href="#">4.1.</a>	Best-effort . . . . .	<a href="#">3</a>
<a href="#">4.2.</a>	Traffic Engineering . . . . .	<a href="#">4</a>
<a href="#">5.</a>	IGP L2 Bundle Member EAG Extensions . . . . .	<a href="#">4</a>
<a href="#">5.1.</a>	ISIS L2 Bundle Member EAG Extension . . . . .	<a href="#">4</a>
<a href="#">5.2.</a>	OSPF L2 Bundle Member EAG Extension . . . . .	<a href="#">5</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">6</a>
<a href="#">6.1.</a>	ISIS IANA Considerations . . . . .	<a href="#">6</a>
<a href="#">7.</a>	Security Considerations . . . . .	<a href="#">6</a>
<a href="#">8.</a>	Acknowledgements . . . . .	<a href="#">6</a>
<a href="#">9.</a>	Normative References . . . . .	<a href="#">6</a>
	Authors' Addresses . . . . .	<a href="#">7</a>

## [1.](#) Introduction

IGP Flex Algorithm [[I-D.ietf-lsr-flex-algo](#)] proposes a solution that allows IGPs themselves to compute constraint based paths over the network, and it also specifies a way of using Segment Routing (SR) Prefix-SIDs and SRv6 locators to steer packets along the constraint-based paths. It specifies a set of extensions to ISIS, OSPFv2 and OSPFv3 that enable a router to send TLVs that identify (a) calculation-type, (b) specify a metric-type, and (c) describe a set of constraints on the topology, that are to be used to compute the best paths along the constrained topology. A given combination of calculation-type, metric-type, and constraints is known as an FAD (Flexible Algorithm Definition).

[RFC8668] and [[I-D.ketant-lsr-ospf-l2bundles](#)] introduces the ability for IS-IS and OSPF respectively to advertise the link attributes of Layer 2 (L2) Bundle Members. Especially, the link attribute "Administrative Group" and "Extended Administrative Group" could be individual to each L2 Bundle Member for purpose of Flex-algo plane construction, where multiple Flex-algo planes share the same Layer 3 parent interface and each Flex-algo plane has dedicated L2 Bundle Member.

This document describes how to create Flex-algo plane with L2bundles scenario.



## **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. Color set on L2 Bundle Member**

Traffic Engineering affinity (also termed as Color) is often to be set on the Layer 3 interface and be flooded by IGP-TE. However, when the Layer 3 interface is a Layer 2 interface bundle, operators can config individual color for each L2 Bundle Member. So that IGP link-state database will contain the TE affinity attribute of L2 Bundle Member, as well as Layer 3 parent interface.

Note that Layer 3 interface can join to IGP instance explicitly, but L2 Bundle Member not.

The TE affinity of the Layer 3 parent interface can be a combined value of all L2 Bundle Members. For example, if the Layer 3 parent interface contains three L2 Bundle Members, each with color "RED", "GREEN", "BLUE" respectively, the Layer 3 parent interface will have color "RED|GREEN|BLUE".

## **4. Flex-algo plane with L2 link resource**

### **4.1. Best-effort**

[I-D.ietf-lsr-flex-algo] defines the color-based link resource selection rules in FAD to construct the expected Flex-algo plane. Each node in the Flex-algo plane will establish the SPT with self as root node, to maintain the best path to other nodes and get the FIB entry based on that. The root node need check the outgoing Layer 2 interface bundle interface, to see which L2 Bundle Member does exactly belong to the Flex-algo plane. The forwarding information of the FIB entry with outgoing Layer 2 interface bundle interface will exactly select the L2 Bundle Member that belongs to the Flex-algo plane to forward packets.

For example, three Flex-algo plane share the same Layer 3 parent interface including three L2 Bundle Members each with color "RED", "GREEN", "BLUE" respectively, and each Flex-algo plane with link selection rule "Include-Any RED", "Include-Any GREEN", "Include-Any BLUE" respectively, Flex-algo SHOULD not simply select the Layer 3 parent interface to all Flex-algo plane, but need continue to select individual L2 Bundle Member to the specific Flex-algo plane. As a



result, the FIB entry within Flex-algo RED plane will exactly choose the L2 Bundle Members with color "RED" to forward packets, the FIB entry within Flex-algo GREEN plane will exactly choose the L2 Bundle Members with color "GREEN" to forward packets, and the FIB entry within Flex-algo BLUE plane will exactly choose the L2 Bundle Members with color "BLUE" to forward packets.

## **4.2. Traffic Engineering**

A segment list contains SIDs advertised specifically for the given algorithm is possible, such as an inter-domain path contains multiple Flex-algo planes, a TI-LFA backup path within the Flex-algo plane, or an optimized TE path avoiding congested link within the Flex-algo plane. In these cases, an Adjacency segment could be used to steer the packets along the expected L2 Bundle Member that belongs to the specific Flex-algo plane.

[RFC8668] and [[I-D.ketant-lsr-ospf-l2bundles](#)] have defined Adjacency-SID for each L2 Bundle Member, that can be used to isolate flows among multiple Flex-algo planes, when these Flex-algo planes share the same Layer 3 parent interface. A specific Adjacency-SID for a specific L2 Bundle Member will steer the packets to that member.

## **5. IGP L2 Bundle Member EAG Extensions**

### **5.1. ISIS L2 Bundle Member EAG Extension**

[RFC8668] defined TLV-25 for ISIS to advertise the link attributes of L2 Bundle Members, and mentioned that the traditional "Administrative group (color) Sub-TLV" and "Extended Administrative Group Sub-TLV" may appear in TLV-25 and MAY be shared by multiple L2 Bundle Members. For individual EAG attributes of L2 Bundle Member purpose, this document defines a new "ISIS L2 Bundle Member EAG sub-TLV". It has the following format:



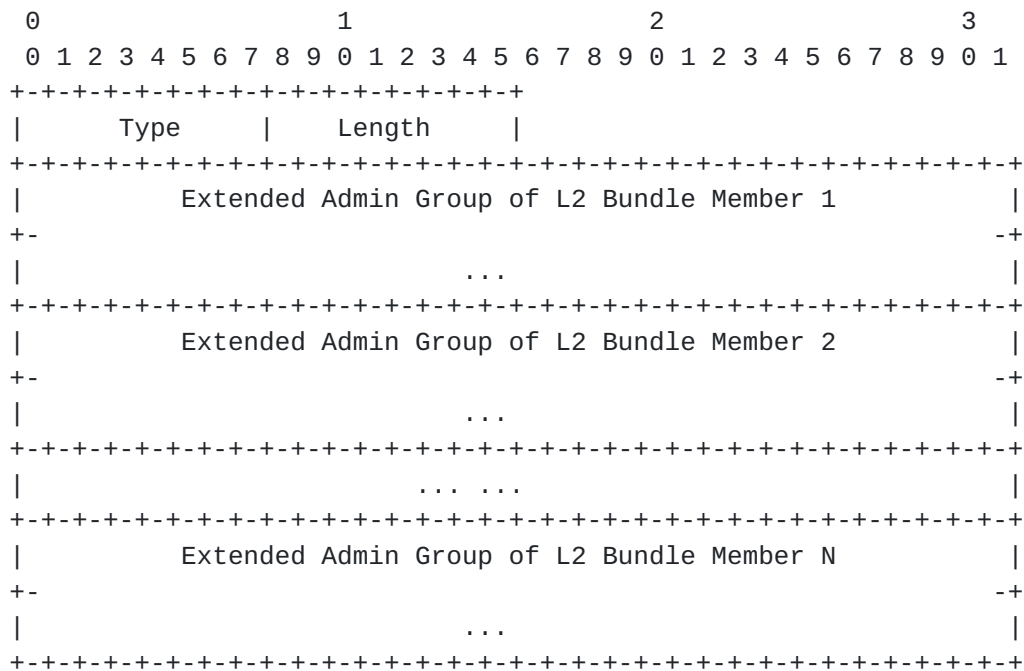


Figure 1: ISIS L2 Bundle Member EAG sub-TLV format

where:

Type: TBD1.

Length: 4 octets, it give the total length of EAG information of all members.

Extended Admin Group of L2 Bundle Member: for each L2 Bundle Member there is an EAG value. Note that "L2 Bundle Member Attributes sub-TLV" explicitly give how many L2 Bundle Members there are and list each L2 Bundle Member Identifier in order, thus in "ISIS L2 Bundle Member EAG sub-TLV" there MUST be an EAG value for each L2 Bundle Member in the same order.

"ISIS L2 Bundle Member EAG sub-TLV" MAY NOT appear more then once in TLV-25. If it appears more then once, the first one is valid.

## 5.2. OSPF L2 Bundle Member EAG Extension

[I-D.ketant-lsr-ospf-l2bundles] defined "L2 Bundle Member Attributes sub-TLV" for OSPF/OSPFv3 to advertise the link attributes of L2 Bundle Members, and mentioned that the traditional "Administrative group (color) Sub-TLV" and "Extended Administrative Group Sub-TLV" are applicable in "L2 Bundle Member Attributes sub-TLV". Because there is "L2 Bundle Member Attributes sub-TLV" per L2 Bundle Member,





it is sufficient to construct Flex-algo plane to select L2 link resource.

## **6. IANA Considerations**

### **6.1. ISIS IANA Considerations**

This document adds the following new sub-TLV to the ISIS TLV-25 registry.

Type: TBD1

Description: ISIS L2 Bundle Member Extended Administrative Group

## **7. Security Considerations**

There are no new security issues introduced by the extensions in this document.

## **8. Acknowledgements**

TBD

## **9. Normative References**

[I-D.ietf-lsr-flex-algo]

Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-05](#) (work in progress), November 2019.

[I-D.ketant-lsr-ospf-l2bundles]

Talaulikar, K. and P. Psenak, "Advertising L2 Bundle Member Link Attributes in OSPF", [draft-ketant-lsr-ospf-l2bundles-00](#) (work in progress), July 2019.

[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>>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](#), DOI 10.17487/RFC4915, June 2007, <<https://www.rfc-editor.org/info/rfc4915>>.



- [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>>.
- [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>>.
- [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>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8668] Ginsberg, L., Ed., Bashandy, A., Filsfils, C., Nanduri, M., and E. Aries, "Advertising Layer 2 Bundle Member Link Attributes in IS-IS", [RFC 8668](#), DOI 10.17487/RFC8668, December 2019, <<https://www.rfc-editor.org/info/rfc8668>>.

#### Authors' Addresses

Peng Shaofu  
ZTE Corporation  
No.68 Zijinghua Road, Yuhuatai District  
Nanjing  
China

Email: peng.shaofu@zte.com.cn

Chen Ran  
ZTE Corporation  
No.50 Software Avenue, Yuhuatai District  
Nanjing  
China

Email: chen.ran@zte.com.cn

Greg Mirsky  
ZTE Corp.

Email: gregimirsky@gmail.com

