

L
Internet-Draft
Intended status: Standards Track
Expires: March 25, 2021

P. Shaofu
C. Ran
ZTE Corporation
G. Mirsky
ZTE Corp.
September 21, 2020

**IGP Flexible Algorithm with L2bundles
draft-peng-lsr-flex-algo-l2bundles-03**

Abstract

IGP Flex Algorithm proposes a solution that allows IGP 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 March 25, 2021.

Copyright Notice

Copyright (c) 2020 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
2.	Requirements Language	3
3.	Color set on L2 Bundle Member	3
4.	Flex-algo plane with L2 link resource	3
4.1.	Best-effort	3
4.2.	Traffic Engineering	4
5.	Flex-algo L2bundles Use-cases	5
6.	IGP L2 Bundle Member EAG advertisement	5
6.1.	ISIS L2 Bundle Member EAG advertisement	5
6.2.	OSPF L2 Bundle Member EAG advertisement	5
7.	IANA Considerations	6
8.	Security Considerations	6
9.	Acknowledgements	6
10.	Normative References	6
	Authors' Addresses	7

[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 [[RFC8402](#)] 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 maintain the best path to other destination nodes. In the case of L2bundles scenario, each node need check the outgoing Layer 2 bundle interface, to see which L2 Bundle Member does exactly belong to the Flex-algo plane.

For the node who originate the l2-bundle interface, the forwarding information of the FIB entry with outgoing Layer 2 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 for all Flex-algo plane, but need continue to

select individual L2 Bundle Member for each 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.

The above processing is a local optimization for each node who originate l2-bundle interface.

In addition, for a remote node which received l2-bundle advertisement originated from other nodes, if that l2-bundle is in the flex-algo based path to a destination node, it must confirm which L2 Bundle Member belongs to the flex-algo plane and check that L2 Bundle Member really meets the constraints defined in the related FAD. This processing is necessary when Flex-algo is used to optimize SID stack depth for an SR-TE policy, e.g, the SR-TE policy defines TE affinity to select individual L2 Bundle Member and the SID list may contain Adjacency-SID for a specific L2 Bundle Member as described in [RFC8668] and [I-D.ketant-lsr-ospf-l2bundles]. Thus the flex-algo based path must be consistent with the original path of the optimized SR-TE policy, i.e, within the flex-algo plane when each node determine its next-hop towards a destination, the determination must be based on the above confirmation and check of L2 Bundle Members.

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 domains, a TI-LFA backup path within the Flex-algo plane, or an optimized TE path avoiding congested link within the Flex-algo plane. When the headend or controller compute these SR-TE paths within the specific flex-algo plane, in addition to the algorithm based Prefix-SID towards the loose node, an Adjacency-SID can also be used to strictly steer the packets along the expected L3 link. However, if the L3 link is a l2-bundle interface, it is necessary to see which L2 Bundle Member exactly belongs to the specific Flex-algo plane and use the Adjacency-SID for that member.

[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 can be contained in the SID list of the SR path within the flex-algo plane and steer the packets to that member.

5. Flex-algo L2bundles Use-cases

In some operator's networks, a large number of bundled links are deployed to improve the bandwidth. However, for a specific l2bundle, each member has different capabilities, such as different delay, bandwidth, AG/EAG, etc. When the path of an SR policy needs to go through an Layer 2 interface bundle, operators want to choose the individual member link to meet business requirements. Different SR policy may choose different member links, according to different set of constraints.

When Flex algorithm is enabled in the above networks, even all flex-algo planes share all Layer 2 interface bundles, i.e, all FA planes have the same structure, an important requirement to Flex-algo is that the constraint based computation of Flex-algo must consider how to select member links to meet service's criterias. In addition, different flex-algo planes can also have different structures, with different set of nodes and links, to meet more strict business requirements.

The extended behavior of flex-algo introduced in this document can meet the above requirement, and exactly it is independent with the structure of flex-algo plane.

6. IGP L2 Bundle Member EAG advertisement

6.1. ISIS L2 Bundle Member EAG advertisement

[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. If we want to advertise unique EAG values for each bundle member, we can use multiple L2 Bundle Attribute Descriptors with each specify a single bundle member. So it is sufficient to construct Flex-algo plane to select L2 link resource.

6.2. OSPF L2 Bundle Member EAG advertisement

[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 also sufficient to construct Flex-algo plane to select L2 link resource.

7. IANA Considerations

This document need not define new sub-TLV to IGP for Flex-algo combined with l2bundles.

8. Security Considerations

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

9. Acknowledgements

TBD

10. 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-11](#) (work in progress), September 2020.

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

Talaulikar, K. and P. Psenak, "Advertising L2 Bundle Member Link Attributes in OSPF", [draft-ketant-lsr-ospf-l2bundles-02](#) (work in progress), June 2020.

[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>>.

[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

