

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 21, 2020

G. Van de Velde, Ed.
W. Henderickx
M. Bocci
Nokia
K. Patel
Arrcus
August 20, 2019

Signalling ERLD using BGP-LS
draft-ietf-idr-bgp-ls-segment-routing-rld-05

Abstract

This document defines the attribute encoding to use for BGP-LS to expose ERLD "Entropy capable Readable Label Depth" from a node to a centralised controller (PCE/SDN).

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 [1].

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 February 21, 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

1. Introduction	2
2. Conventions used in this document	2
2.1. Terminology	3
3. Problem Statement	3
4. Advertising of ERLD in BGP-LS	3
5. Security Considerations	4
6. Acknowledgements	4
7. IANA Considerations	4
8. References	4
8.1. Normative References	4
8.2. Informative References	5
Authors' Addresses	5

1. Introduction

When Segment Routing tunnels are computed by a centralised controller, it is beneficial that the controller knows the ERLD (Entropy Readable Label Depth) of each node or link a tunnel traverses. A network node signalling an ERLD MUST support the ability to read the signalled number of labels before any action is done upon the packet and SHOULD support entropy awareness found within the signalled ERLD depth.

ERLD awareness of each node will allow a network SDN controller to influence the path used for each tunnel. The SDN controller may for example only create tunnels with a label stack smaller or equal as the ERLD of each node on the path. This will allow the network to behave accordingly (e.g. make use of Entropy Labels to improve ECMP) upon the imposed Segment Routing label stack on each packet.

This document describes how to use BGP-LS to expose the ERLD of a node.

2. Conventions used in this document

2.1. Terminology

BGP-LS: Distribution of Link-State and TE Information using Border Gateway Protocol

ERLD: Entropy capable Readable Label Depth

ELC: Entropy Label Capability

MSD: Maximum SID Depth

SID: Segment Identifier

3. Problem Statement

For Segment Routing technology both ISIS [7] and OSPF [6] have proposed extensions to signal the ERLD (Entropy Readable Label Depth) and ELC (Entropy Label Capability) of a node. However, if a network SDN controller is connected to the network through a BGP-LS session and not through either ISIS or OSPF technology, then both ERLD and ELC needs to be signalled using BGP-LS encoding. This document describes the extension BGP-LS requires to signal ERLD.

A network SDN controller having awareness of the ERLD can for example use it as a constraint on path computation to make sure that high bandwidth LSPs are not placed on LAG (Link Aggregation Group), containing links with smaller member bandwidth, if they know the entropy label cannot be processed by the node at the ingress to the link.

4. Advertising of ERLD in BGP-LS

Both ISIS [7] and OSPF [6] have proposed extensions to signal the ERLD (Entropy Readable Label Depth) and ELC (Entropy Label Capability) using new MSD-type of the Node MSD sub-Type TLV RFC8491 [4] or RFC8476 [3].

This document defines a new node BGP MSD sub-type TLV from draft-ietf-idr-bgp-ls-segment-routing-msd [5] to signal the ERLD.

A BGP-LS router exporting the IGP LSDB, MUST NOT encode the IGP ERLD value in an BGP-LS ERLD attribute, if the associated ELC is not signalled.




Figure 1

The BGP-LS ERLD is encoded as a node MSD sub-type defined in the IANA registry titled "IGP MSD-Types" under the "Interior Gateway Protocol (IGP) Parameters" registry created by RFC8491 [4].

The ERLD-Value field in the range between 0 to 255 is set to the BGP-LS imported IGP ERLD. The value of 0 represents lack of ability to read a label stack of any depth, any other value represents the readable label depth of the node.

5. Security Considerations

This document does not introduce security issues beyond those discussed in RFC7752 [2]

6. Acknowledgements

Thanks to discussions with Acee Lindem, Jeff Tantsura, Stephane Litkowski, Bruno Decraene, Kireeti Kompella, John E. Drake and Carlos Pignataro to bring the concept of combining ELC and RLD into a single ERLD signalled parameter more suitable for SDN controller based networks.

7. IANA Considerations

This document requests assigning new BGP MSD sub-TLV code-points as described in section 4.

Note: placeholder IANA request

8. References

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://xml.resource.org/public/rfc/html/rfc2119.html>>.

- [2] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [3] Tantsura, J., Chunduri, U., Aldrin, S., and P. Psenak, "Signaling Maximum SID Depth (MSD) Using OSPF", RFC 8476, DOI 10.17487/RFC8476, December 2018, <<https://www.rfc-editor.org/info/rfc8476>>.
- [4] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018, <<https://www.rfc-editor.org/info/rfc8491>>.
- [5] Tantsura, J., Chunduri, U., Talaulikar, K., Mirsky, G., and N. Triantafyllis, "draft-ietf-idr-bgp-ls-segment-routing-msd", June 2019.

8.2. Informative References

- [6] Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "draft-ietf-ospf-mpls-elc", May 2019.
- [7] Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "draft-ietf-isis-mpls-elc", May 2019.

Authors' Addresses

Gunter Van de Velde (editor)
Nokia
Antwerp
BE

Email: gunter.van_de_velde@nokia.com

Wim Henderickx
Nokia
Belgium

Email: wim.henderickx@nokia.com

Matthew Bocci
Nokia
Shoppenhangers Road
Maidenhead, Berks
UK

Email: matthew.bocci@nokia.com

Keyur Patel
Arrcus
USA

Email: keyur@arrcus.com