IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 21, 2018

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

December 18, 2017

Signalling ERLD using BGP-LS
draft-ietf-idr-bgp-ls-segment-routing-rld-01

Abstract

This document defines the attributes to use for BGP-LS to expose ERLD "Entropy capable Readable Label Depth" from a node or link 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 June 21, 2018.

Copyright Notice

Table of Contents

1. Introduction 2
2. Conventions used in this document 3
2.1. Terminology 3
3. Problem Statement 3
4. ERLD support by a node 3
5. ERLD support by a link 4
6. Security Considerations 5
7. Acknowledgements 5
8. IANA Considerations 5
9. References 5
9.1. Normative References 5
9.2. Informative References 5
Authors' Addresses 6

1. Introduction

When Segment Routing tunnels are computed by a centralised controller, it is beneficial that the controller knows the ERLD (Entropy capable 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 and link 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 and link 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

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Protocol

SID: Segment Identifier

SR: Segment routing

3. Problem Statement

In existing technology both ISIS [4] and OSPF [3] have proposed extensions to signal the RLD (Readable Label Depth) and ELC (Entropy Label Capability) of a node or link. However, if a network SDN controller is connected to the network through a BGP-LS session and not through ISIS or OSPF technology, then both RLD and ELC needs to be signaled in BGP-LS accordingly. This document describes the extension BGP-LS requires to transport the combination of RLD and ELC into according ERLD attributes for nodes and links.

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

4. ERLD support by a node

Node ERLD is encoded in a new Node Attribute TLV, as defined in RFC7752 [2].

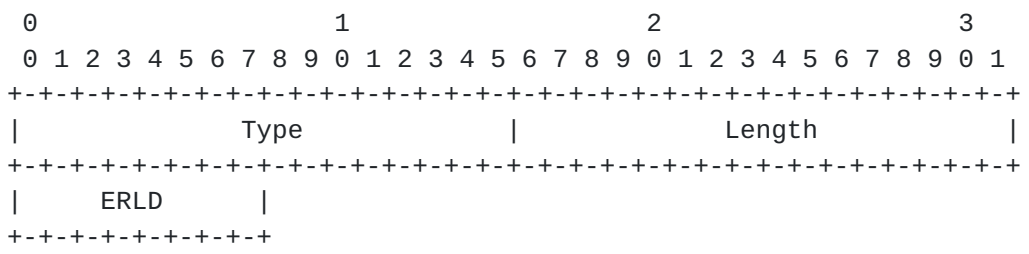

Figure 1

Type : A 2-octet field specifying code-point of the new TLV type. Code-point: TBA from BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs registry

Length: A 2-octet field that indicates the length of the value portion

ERLD: Node ERLD is a number in the range of 0-254. 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. ERLD support by a link

Link ERLD is encoded in a new Link Attribute TLV, as defined in RFC7752 [2].

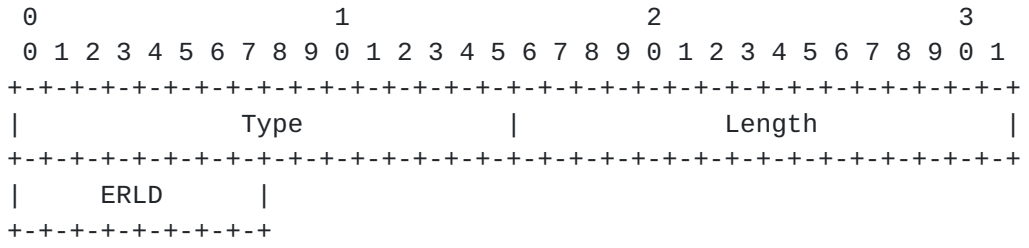

Figure 2

Type : A 2-octet field specifying code-point of the new TLV type. Code-point: TBA from BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs registry

Length: A 2-octet field that indicates the length of the value portion

ERLD: Link ERLD is a number in the range of 0-254. 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 link.

6. Security Considerations

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

7. 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.

8. IANA Considerations

This document requests assigning 2 new code-points from the BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs registry as specified in sections 4 and 5.

9. References

9.1. 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>.

9.2. Informative References

[3] Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "draft-ietf-ospf-mpls-elc", October 2016.

[4] Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "draft-ietf-isis-mpls-elc", October 2016.

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