MPLS   B. Decraene, Ed.
Internet-Draft   Orange
Updates: 6790 (if approved)   C. Filsfils
Intended status: Standards Track   Cisco Systems, Inc.
Expires: June 19, 2021   W. Henderickx
                                                                 Nokia
                                                               T. Saad
                                                              V. Beeram
                                                      Juniper Networks
                                                     December 16, 2020

       Using Entropy Label for Network Slice Identification in MPLS networks.
            draft-decraene-mpls-slid-encoded-entropy-label-id-00

Abstract

   This document defines a solution to encode a slice identifier in MPLS
   in order to distinguish packets that belong to different slices, to
   allow enforcing per network slice policies (.e.g, Qos).

   The slice identification is independent of the topology. It allows
   for QoS/DiffServ policy on a per slice basis in addition to the per
   packet QoS/DiffServ policy provided by the MPLS Traffic Class field.

   In order to minimize the size of the MPLS stack and to ease
   incremental deployment the slice identifier is encoded as part of the
   Entropy Label.

   This document also extends the use of the TTL field of the Entropy
   Label in order to provide a flexible set of flags called the Entropy
   Label Control field.

Status of This Memo

This Internet-Draft will expire on June 19, 2021.

Copyright Notice

Table of Contents

1. Introduction

Segment Routing (SR) [RFC8402] leverages the source-routing paradigm.
A node steers a packet through a controlled set of instructions,
called segments, by prepending the packet with an SR header. In the
SR-MPLS data plane [RFC8660], the SR header is instantiated through a
label stack.

This document defines a solution to encode a slice identifier in MPLS
in order to provide QoS on a per slice basis. It allows for QoS/
DiffServ policy on a per slice basis in addition to the per packet
QoS/DiffServ policy provided by the MPLS Traffic Class field. The
slice identification is independent of the topology and the QoS of
the network, thus enabling scalable network slicing.

This document encodes the slice identifier in a portion of the MPLS
Entropy Label (EL) defined in [RFC6790]. This has advantages in SR-
MPLS networks as it avoids the use of additional label which would
increase the size of the label stack. This also reuses the data
plane processing of the Entropy Label on the egress LSR, the
signaling of the Entropy Label capability from the egress to the
ingress [I-D.ietf-isis-mpls-elc] [I-D.ietf-ospf-mpls-elc], and the
signaling capability of transit routers to read this label [RFC8491]
which allows for an easier and faster incremental deployment.

2. Entropy Label Control field

[RFC6790] defines the MPLS Entropy Label. [RFC6790] section 4.2
defines the use of the Entropy Label Indicator (ELI) followed by the
Entropy Label (EL) and the MPLS header fields (Label, TC, S, TTL) in
each. [RFC6790] also specifies that the TTL field of the EL must be
set to zero by the ingress LSR.

Following the procedures of [RFC6790] EL is never used for forwarding
and its TTL is never looked at nor decremented:

o An EL capable Egress LSR performs a lookup on the ELI and as a
   result pop two labels: ELI and EL.

o An EL non-capable Egress LSR performs a lookup on the ELI and as a
   result must drop the packet as specified in [RFC3031] for the
   handling of an invalid incoming label.

Hence essentially the TTL field of the EL behaves as a reserved field
which must be set to zero when sent and ignored when received.

This documents extends the TTL field of the EL and calls it the
Entropy Label Control (ELC) field. The ELC is a set of eight flags:
ELC0 for bit 0, ELC1 for bit 1,..., ELC7 for bit 7.

Given that the MPLS header is very compact (32 bits) with no reserved
bits and that MPLS is used within a trusted administrative domain,
the semantic of these bits is not standardized but defined on a per
administrative domain basis. This allows for increased re-use and
flexibility of this scarce resource. As a consequence, an
application using one of those buts MUST allow the choice of the bit
by configuration by the network operator.

3. Slice Identifier

Each network slice in an MPLS domain is uniquely identified by a
Slice Identifier (SLID). This section proposes to encode the SLID in
a portion of the MPLS Entropy Label defined in [RFC6790].

The number of bits to be used for encoding the SLID in the EL is
governed by a local policy and uniform within a network slice policy
domain.

3.1. Ingress LSR

When an ingress LSR classifies that a packet belongs to the slice and
that the egress has indicated via signaling that it can process EL
for the tunnel, the ingress LSR pushes an Entropy Label with the:

o SLID encoded in the most significant bits of the Entropy Label.

o the entropy information encoded in the remaining lower bits of the
   Entropy Label as described in section 4.2 of [RFC6790].

o SPI bit (SLID Presence Indicator) set in one bit of the ELC field.

The choice of the ELC field used for SPI, and the number of bits to
be used for encoding the SLID MUST be configurable by the network
operator.

The slice classification method is outside the scope of this
document.

3.2. Transit LSR

Any router within the SR domain that forwards a packet with the SPI
bit set MUST use the SLID to select a slice and apply per-slice
policies.

There are many different policies that could define a slice for a
particular application or service. The most basic of these is
bandwidth-allocation, an implementation complying with this
specification SHOULD support the bandwidth-allocation slice as
defined in the next section.

3.3. Bandwidth-Allocation Slice

A per-slice policy is configured at each interface of each router in
the SR domain, with one traffic shaper per SLID. The bit rate of
each shaper is configured to reflect the bandwidth allocation of the
per-slice policy.

If shapers are not available, or desirable, an implementation MAY
configure one scheduling queue per SLID with a guaranteed bandwidth
equal to the bandwidth-allocation for the slice. This option allows
a slice to consume more bandwidth than its allocation when available.

Per-slice shapers or queues effectively provides a virtual port per
slice. This solution MAY be complemented with a per-virtual-port
hierarchical DiffServ policy. Within the context of one specific
slice, packets are further classified into children DiffServ queues
which hang from the virtual port. The Traffic Class value in the
MPLS header SHOULD be used for queue selection.

3.4. Backward Compatibility

The Entropy Label usage described in this document is consistent with
[RFC6790] as ingress LSRs freely chooses the EL of a given flow, and
transit LSRs treat the EL as an opaque set of bits.

As per [RFC6790] an ingress LSR that does not support this extension
has the SPI bit cleared, and thus does not enable the SLID semantic
of the Entropy bits. Hence, SLID-aware transit LSRs will not
classify these packets into a slice.

3.5. Benefits

From a Segment Routing architecture perspective, this network slice
identifier for SR-MPLS is inline with the network slice identifier
for SRv6 proposed in [I-D.filsfils-spring-srv6-stateless-slice-id].

From an SR-MPLS perspective, using the EL to carry the network slice
identifier has multiple benefits:

o This limits the number of labels pushed on the MPLS stack compared
   to using a pair of labels (ELI+EL) for flow entropy plus two or
   three labels for the slice indicator and the slice identifier.
   This is beneficial for the ingress LSR which may have limitations
   with regards to the number of labels pushed, for the transit LSR
   which may have limitations with regards to the label stack depth
   to be examined during transit in order to read both the entropy
   and the SLID. This presents additional benefit to network
   operators by reducing the packet overhead for traffic carried
   through the network;

o This avoids defining new extensions for the signaling of the
   egress capability to support the slice indicator and the slice
   identifier;

o This improves incremental deployment as all egress LSRs supporting
   EL can be sent the slice identifier from day one, allowing slice
   classification on transit LSRs.

4. End to end absolute loss measurements

This section describes the usage of a ELC flag to enable packet loss
measurements, as described in section 3.1 of [RFC8321], for SR-MPLS
networks.

TBD

5. Programmed sampling of packets

This section describes the usage of a ELC flag to detect end to end
packet loss.

TBD

6. References

6.1. Normative References

   [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and
              L. Yong, "The Use of Entropy Labels in MPLS Forwarding",
              RFC 6790, DOI 10.17487/RFC6790, November 2012,
              <https://www.rfc-editor.org/info/rfc6790>.

   [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>.

   [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S.,
              Decraene, B., Litkowski, S., and R. Shakir, "Segment
              Routing with the MPLS Data Plane", RFC 8660,
              DOI 10.17487/RFC8660, December 2019,
              <https://www.rfc-editor.org/info/rfc8660>.

6.2. Informative References

   [I-D.filsfils-spring-srv6-stateless-slice-id]
              Filsfils, C., Clad, F., Camarillo, P., and K. Raza,
              "Stateless and Scalable Network Slice Identification for
              SRv6", draft-filsfils-spring-srv6-stateless-slice-id-01
              (work in progress), July 2020.

   [I-D.ietf-isis-mpls-elc]
              Xu, X., Kini, S., Psenak, P., Filsfils, C., Litkowski, S.,
              and M. Bocci, "Signaling Entropy Label Capability and
              Entropy Readable Label Depth Using IS-IS", draft-ietf-
              isis-mpls-elc-13 (work in progress), May 2020.

[I-D.ietf-ospf-mpls-elc]
          Xu, X., Kini, S., Psenak, P., Filsfils, C., Litkowski, S.,
          and M. Bocci, "Signaling Entropy Label Capability and
          Entropy Readable Label Depth Using OSPF", draft-ietf-ospf-
          mpls-elc-15 (work in progress), June 2020.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol
          Label Switching Architecture", RFC 3031,
          DOI 10.17487/RFC3031, January 2001,
          <https://www.rfc-editor.org/info/rfc3031>.

[RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli,
          L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi,
          "Alternate-Marking Method for Passive and Hybrid
          Performance Monitoring", RFC 8321, DOI 10.17487/RFC8321,
          January 2018, <https://www.rfc-editor.org/info/rfc8321>.

[RFC8491] 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>.

Authors' Addresses

Bruno Decraene (editor)
Orange

Email: bruno.decraene@orange.com

Clarence Filsfils
Cisco Systems, Inc.
Belgium

Email: cf@cisco.com

Wim Henderickx
Nokia
Copernicuslaan 50
Antwerp 2018, CA 95134
Belgium

Email: wim.henderickx@nokia.com

   Tarek Saad
   Juniper Networks

   Email: tsaad@juniper.net

   Vishnu Pavan Beeram
   Juniper Networks

   Email: vbeeram@juniper.net