MPLS-based Service Function Path (SFP) Consistency Verification

draft-lm-mpls-sfc-path-verification-02

Yao Liu (ZTE)
Greg Mirsky(ZTE)

MPLS WG IETF#110 March, 2021

Requirements and Scope

Requirements:

- Leverage MPLS LSP ping mechanism to support verification between the control plane and the data plane for MPLS-based Service Function Path.
  - MPLS-SFC(RFC 8595)
  - SR-SFC(draft-ietf-spring-sr-service-programming)

Scope:

- MPLS-SFC OAM
  - SFFs process the MPLS echo request
  - Use GAL to identify the OAM packet
  - New FEC sub-TLVs
- SR-SFC OAM
  - Based on LSP Ping for SR-MPLS [RFC8287]

Updates Since IETF-109

- MPLS-SFC OAM and SR-SFC OAM are separated and discussed as two independent topics.
- More consistent with the LSP Ping mechanism
  - SFC Validation TLV and its sub-TLVs defined in the previous version are deleted.
  - SFC Basic Unit Nil FEC Sub-TLV is introduced (for MPLS-SFC OAM) to ensure that the proper validation in LSP ping can still be performed.
- Why SFFs are responsible for MPLS echo request processing in MPLS-SFC.
- Add a note that the draft updates RFC 8595
  - The paragraph in RFC8595 that this draft updates is changed.
- The terminology section is updated
- Various editorial changes

FEC sub-TLVs

Sub-TLVs for TLV Types 1(FEC TLV), 16, and 21 defined only for SFC-MPLS

SFC Basic Unit Sub-TLV

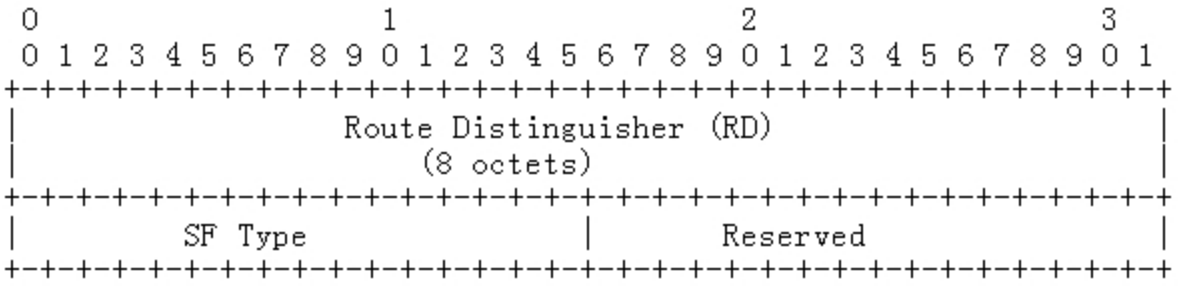

The RD and SF Type fields are taken from the BGP control plane for MPLS-SFC [draft-ietf-bess-nsh-bgp-control-plane] .

SFC Basic Unit Nil Sub-TLV

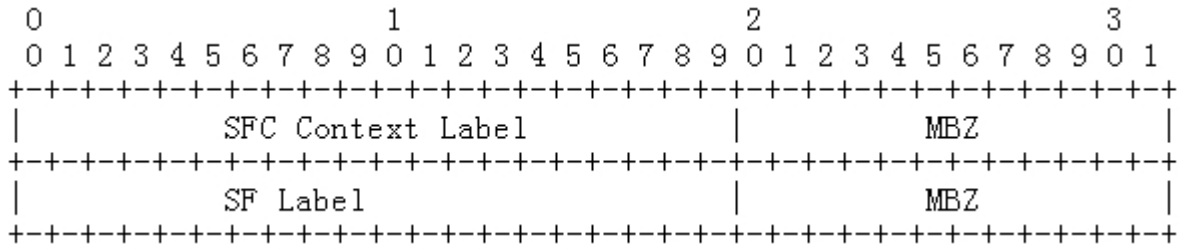

When present, MUST be immediately followed by an SFC Basic Unit sub-TLV.

During FEC validation, an SFF should skip the SFC Basic Unit Nil sub-TLV and use the SFC Basic Unit sub-TLV to validate the FEC of the basic unit.

Theory of Operation (SFC-MPLS)

Sending an MPLS Echo Request

- The echo request is sent with a label stack corresponding to the SFP being tested, the GAL(s) is also included.
- If FEC validation is required,
  - the SFC Basic Unit sub-TLV SHOULD be carried in the FEC stack of the request packet,
  - the SFC Basic Unit Nil sub-TLV MAY also be carried.
- A Downstream Detailed Mapping TLV MAY be included in the MPLS echo request of the SFP.
- The operator should manage the TTL field of SF Label properly .

Theory of Operation (SFC-MPLS)

Sending an MPLS Echo Request: Examples

Swapping Mode

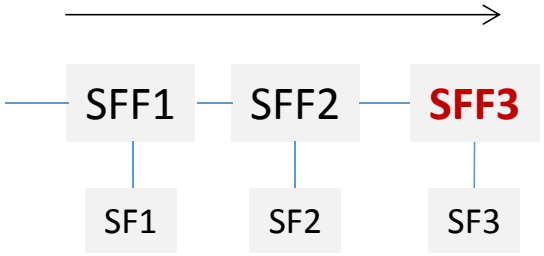

SFC Context Label
GAL
SF Label(TTL=3)

MPLS Label Stack

SFC Basic Unit Sub-TLV 3

FEC Stack

"traceroute" mode for Label Swapping: the TTL of the SF Label is set successively to 1, 2, ...

Stacking Mode

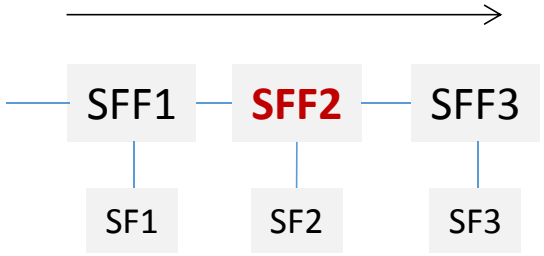

SFC Context Label 1
GAL
SF Label 1(TTL=255)
SFC Context Label 2
GAL
SF Label 2(TTL=1)
SFC Context Label 3
SF Label 3

MPLS Label Stack

SFC Basic Unit Sub-TLV 2
SFC Basic Unit Nil Sub-TLV 3
SFC Basic Unit Sub-TLV 3

FEC Stack

SFC Context Label 3
SF Label 3

Labels in Downstream Detailed Mapping TLV

Theory of Operation (SFC-MPLS)

Receiving an MPLS Echo Request

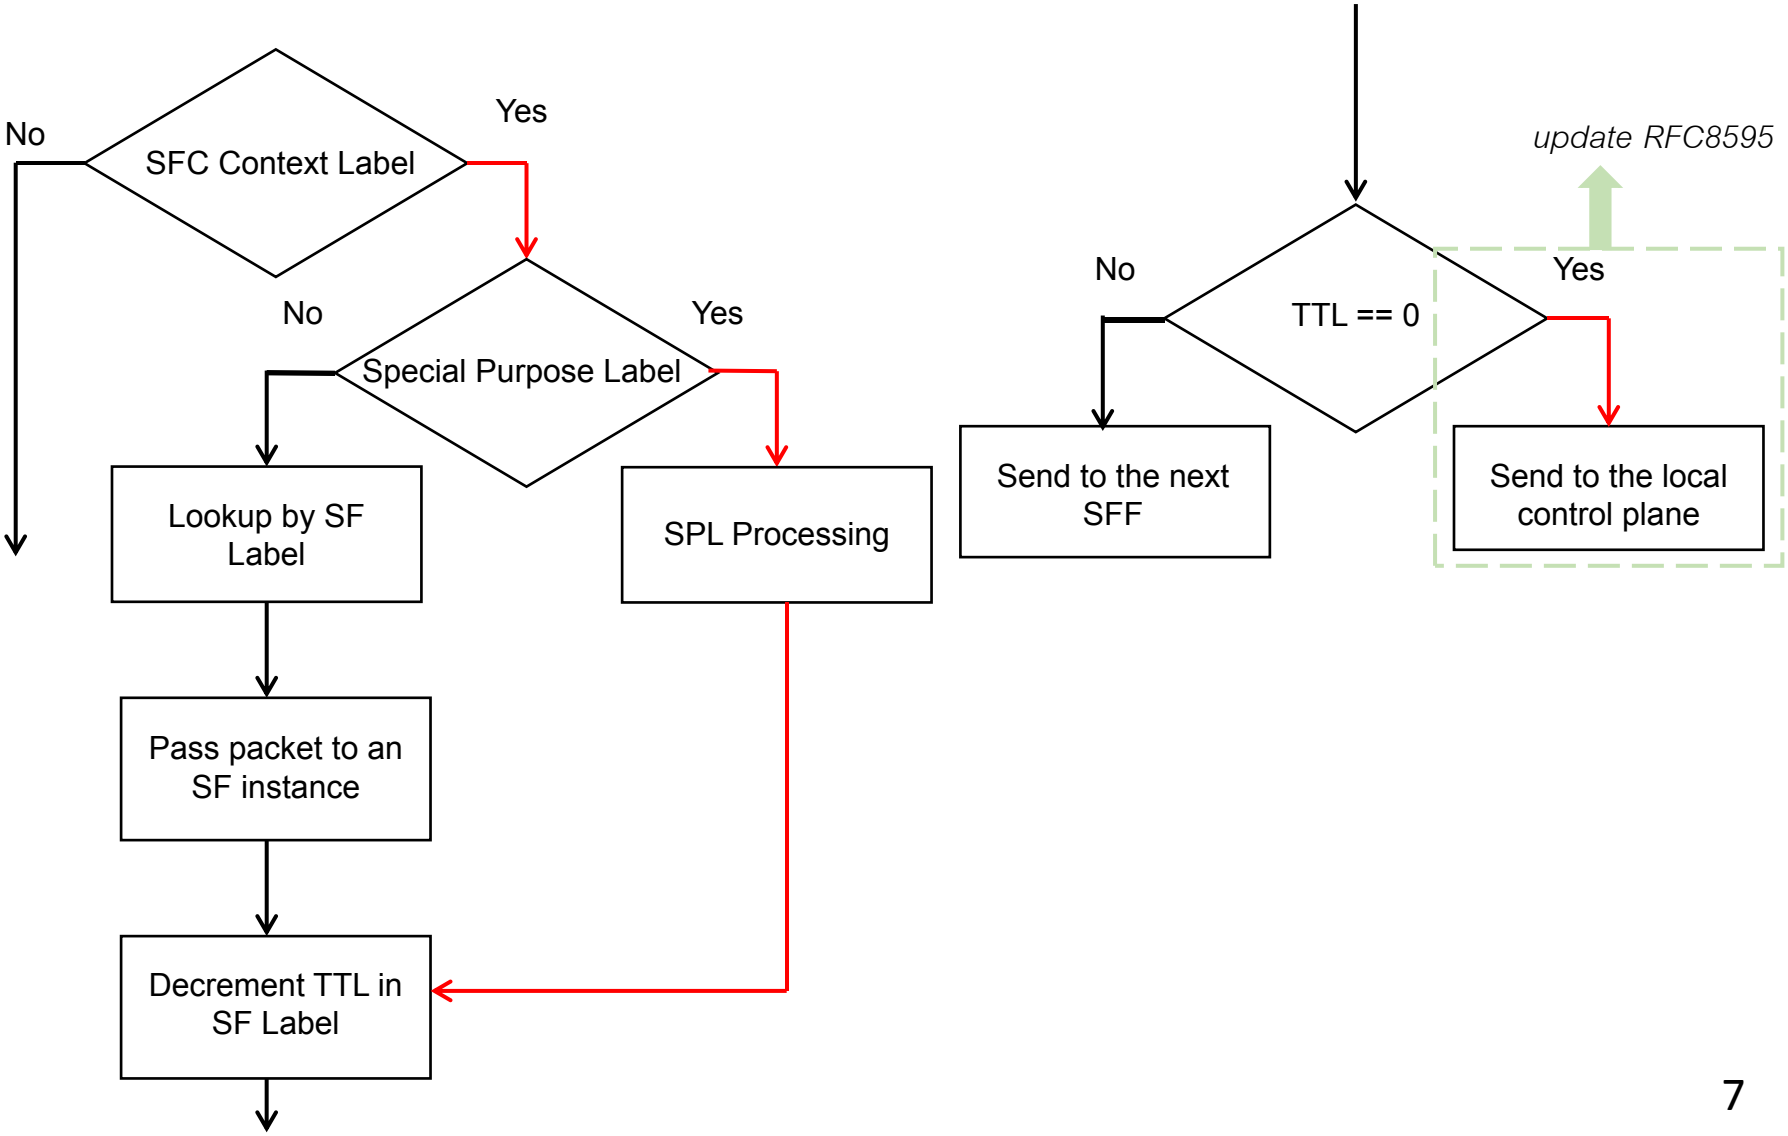

Theory of Operation (SFC-MPLS)

Receiving an MPLS Echo Request: Update of RFC8595

Original text:

If an SFF decrements the TTL to zero, it MUST NOT send the packet and MUST discard the packet.

Updated text:

If an SFF decrements the TTL to zero, an OAM packet MAY be sent to the control plane given it does not exceed the configured rate intended to protect the system from the possible denial-of-service attack.

Next Steps

- Your feedback and comments are welcome and much appreciated

Thank You !

Background

- MPLS-based Service Function Path (SFP)

- SR-MPLS Service Programming (SFC-SR) : each SF is associated with an MPLS label, an SFP can be encoded as a stack of MPLS labels and pushed on top of the packet. [draft-ietf-spring-sr-service-programming]

forwarding based on a single label

- MPLS-based Network Service Header (SFC-MPLS) : a basic unit of representation is used, which comprises two MPLS labels, one carries a label to provide a context within the SFC scope, and the other carries a label to show which SF is to be enacted. [RFC8595]

forwarding based on a basic unit

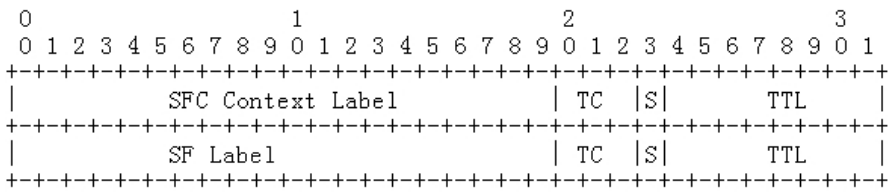

The Basic Unit of MPLS Label Stack for SFC

Special-purpose Label (SPL) in SFC-MPLS

An SFF needs to identify an OAM packet with the SFP scope because an SF may not be capable of processing the SFP OAM payload.

Using an SPL in the basic unit allows for a closer functional match between NSH and SFC-MPLS.

An SPL unit MAY be present in one or more basic units.

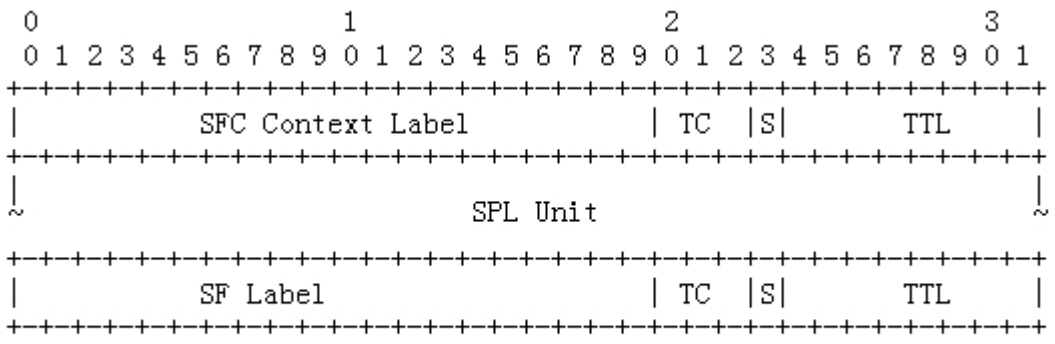

Special-purpose Label Unit with the Basic Unit of MPLS Label Stack for SFC

G-ACh over SFC-MPLS

G-ACh: Generic Associated Channel, over which OAM and other control messages can be exchanged [RFC5586].

GAL: G-ACh Label. If the GAL immediately follows the SFC Context label, then the packet is recognized as an SFP OAM packet.

Processing Rules:

- An SFF MUST NOT pass the OAM packet to a local SFI or SFC proxy.
- The SFF MUST decrement SF Label entry's TTL value. If the resulting value equals zero, the SFF MUST pass the SFP OAM packet to the control plane for processing.
- If the TTL value is not zero, the SFP OAM packet is processed as defined in [RFC8595], according to the type of MPLS forwarding used in the SFP.

GAL in SFC-MPLS Packet

GAL in Label Swapping Mode

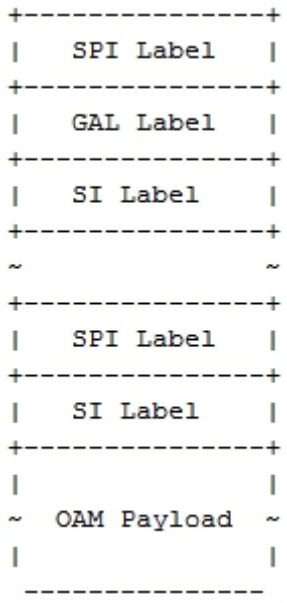

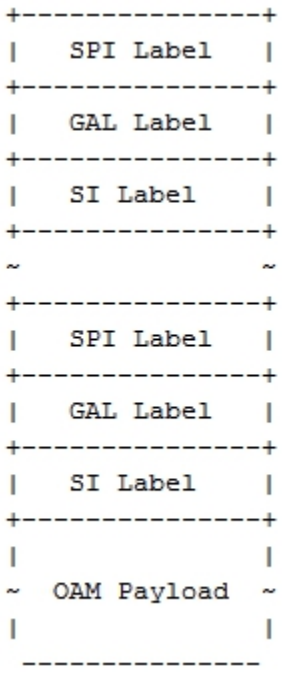

GAL in Label Stacking Mode or Mix Mode

FEC validation

FEC validation(Intermediate node)

```
First determine FEC-stack-depth from the Downstream Detailed
Mapping TLV. This is done by walking through Stack-D (the
Downstream labels) from the bottom, decrementing the number of
labels for each non-Implicit Null label, while incrementing
FEC-stack-depth for each label.
... ...

Set FEC-stack-depth to 0.  Set i to Label-stack-depth.

While (i > 0) do {

    ++FEC-stack-depth.
    if Stack-D [ FEC-stack-depth ] != 3 (Implicit Null)
    --i.
}

If the number of FECs in the FEC stack is greater than or equal
to FEC-stack-depth {
Perform the FEC Checking procedure (see Section 4.4.1).
... ...
}
Go to step 7 (Send Reply Packet).
```