

MPLS Working Group
Internet-Draft
Intended status: Informational
Expires: 27 October 2022

R. Raszuk, Ed.
NTT NI
25 April 2022

Framework of MPLS Reference Augmented Forwarding
draft-raszuk-mpls-raf-fwk-00

Abstract

This document specifies an architectural framework for enabling MPLS based forwarding with optional reference based packet processing in transit network elements.

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](#) [[RFC2119](#)].

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 27 October 2022.

Copyright Notice

Copyright (c) 2022 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

Internet-Draft

MPLS RAF

April 2022

extracted from this document must include Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	3
2.1.	Definitions	3
2.2.	Abbreviations	3
3.	Encoding	4
4.	Processing	5
5.	Reference to Special Instructions Mapping	5
6.	Operational Considerations	6
7.	Capability Advertisement	7
8.	IANA Considerations	7
9.	Security Considerations	7
10.	Acknowledgements	7
11.	References	7
11.1.	Normative References	7
11.2.	Informative References	8
	Author's Address	8

[1.](#) Introduction

Growth of network services results in increased demand for custom transit packet processing. Traditional per destination based best effort or constrained based forwarding is no longer sufficient. Hop by hop switching in addition to label or destination based LSP switching can be augmented with additional packet processing at all or at selective transit network elements.

Such additional processing can be triggered in a number of ways. Today some networks can apply local policies which enable selective processing on subset of packets based on the header's elements. Alternative to such filtering is to embed additional information in the packet header itself to indicate implicitly or explicitly what additional processing is required.

Examples of explicit encoding of such network actions can be found in SRv6 Network Programming [[RFC8986](#)] document. For MPLS data plane an analogy to SRv6 has been recently proposed in [draft-andersson-mpls-mna-fwk](#) [[I-D.andersson-mpls-mna-fwk](#)].

In this document authors propose an implicit model. Instead of explicitly encoding all required actions as a variable size ordered list in every packet this document proposes to insert fixed size 20 bit reference identifier. Such value will be used to mark groups of flows with identical custom forwarding behaviour.

By forwarding behaviour (further abbreviated as FB) this document assumes additional network actions which may exclude packet from default processing, may include additional security screening, OAM triggering actions or any other special handling including, but not limited to rate limited, queue mapping, redirection etc ...

[2.](#) Terminology

[2.1.](#) Definitions

- * Network Action aka Forwarding Behaviour: An operation to be performed on a packet or be triggered based on given packet's arrival without affecting the packet switching decision. A network action may affect forwarding decisions, queuing classification, OAM measurement and reporting etc Network Actions are not carried in packets. They are distributed by configuration or control plane.
- * Reference Augmented Forwarding: Packet forwarding in addition to or instead of normal lookup (by address or label) based on a set of Network Actions or Forwarding Behaviours indicated by Reference Identifier.
- * Reference Forwarding Value: A 20 bit value carried in a packet used to identify a set of Network Actions defining given packet's special handling.
- * Reference Forwarding Indicator: A base Special Purpose Label carried in a packet used to indicate to the packet processor that the following LSE is containing Reference Forwarding Value.

2.2. Abbreviations

- * FB - Network Actions or Forwarding Behaviours
- * LSE - Label Stack Entry
- * RAF - Reference Augmented Forwarding
- * bSPL - Base Special Purpose Label
- * ECMP - Equal Cost Multipath

Raszuk

Expires 27 October 2022

[Page 3]

Internet-Draft

MPLS RAF

April 2022

- * RFI - Reference Forwarding Indicator
- * RFV - Reference Forwarding Value
- * TBA - To Be Allocated
- * SDN - Software Defined Network

3. Encoding

A Reference Forwarding Value MUST be clearly distinguished from any forwarding label. LSE immediately preceding label stack entry containing RFV is called Reference Forwarding Indicator. RFI is REQUIRED to use Special Purpose Label value (TBA by IANA).

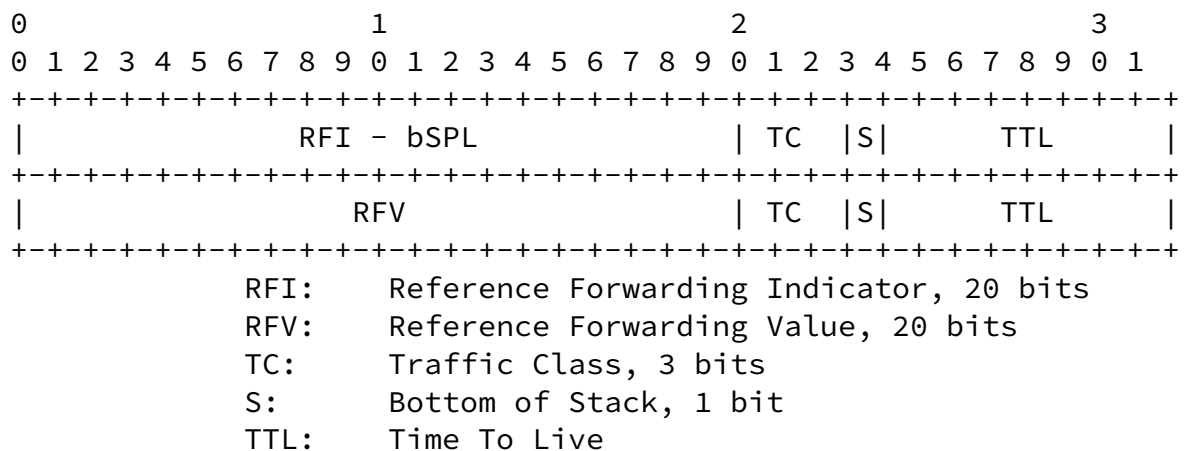


Figure 1: RFI and RFV Tuple

RFI and RFV tuple is always of fixed size of 8 octets. It also should occur only once in a given packet. It is optional. If a network uses the concept of LSPs it MUST be placed after the topmost label. If LSP is not setup and the network uses the concept of SDN based path computation it can be preset as topmost LSE.

How to set values of the TTL, TC, and Bottom of Stack (S bit) fields [RFC3032] for the RFI and RFV entries should be the same as described in [RFC6790] Section 4.2. Ingress LSR SHOULD put the same TTL and TC fields for the RFI as it does for transport label. Such ingress LSR MAY choose different values for the TTL and TC fields if it is known that the RFI will not be exposed as the top label at any point along the LSP (as may happen in cases where PHP is used and the RFI and RFV are not stripped at the penultimate hop. The BoS bit for the RFI MUST be set to zero (i.e., BoS is not set). The TTL for the RFV MUST be zero to ensure that it is not used inadvertently for forwarding. The TC for the RFV may be any value. The BoS bit for the RFV depends on whether or not there are more labels in the label stack.

[4.](#) Processing

Any network element can insert, delete or modify RFV or RFI-RFV tuple if instructed to do so by special action instructions.

If a network element understands RFI and recognizes based on the top most label value special handling requirement it MAY direct packet for special processing. MAY or MUST processing of all requested actions (or subset of those actions) really depend on the special action instructions.

[5.](#) Reference to Special Instructions Mapping

As the proposed architecture is based on indirection, what is present in the packet is only a reference to special handling instructions. Such instructions are not to be explicitly carried in the packet. As each network operator has a different set of operational preferences, embedding insertion of actions into a data plane parsing of each packet is very operationally limiting and inefficient.

Special Network Actions or Forwarding Behaviours are to be distributed by configuration or by control plane. Details of their

distribution are outside of scope of this framework document. However, it is important to recognize that this model in its roots allows open innovation for vendors in developing accelerated data plane action dictionaries as mapping and execution has only a local scope.

It needs to be further observed that format of such configuration or control plane (including PUB-SUB model) distributed Forwarding Behaviours MAY have a TLV/sub-TLV structure as illustrated on Figure 2 and 3 below:

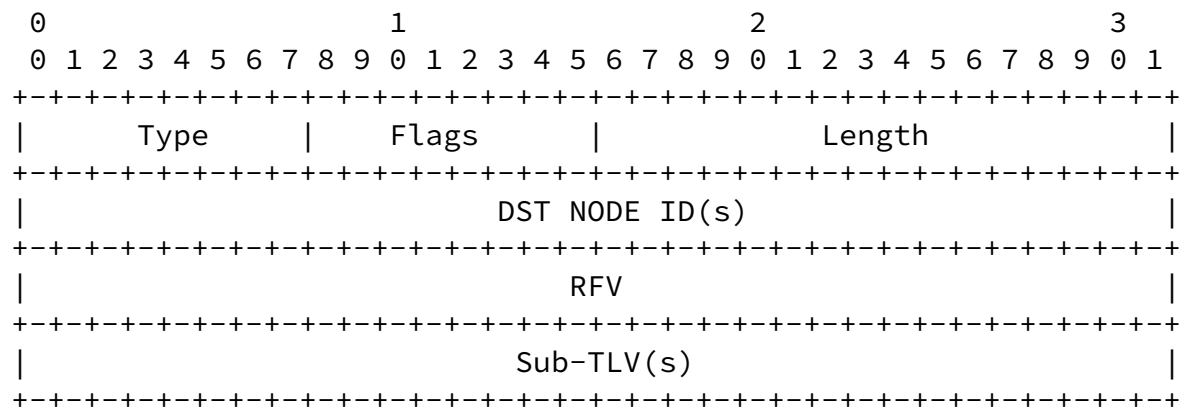


Figure 2: FB TLV




Figure 3: FB sub-TLV

With such encoding option it needs to be observed that for a given RFV only nodes listed in the TLV will accept and install special handling instructions.

6. Operational Considerations

The proposed model is designed to operate both in the networks using traditional MPLS LSP (with local labels) as well with SR-MPLS (domain wide labels).

Nodes which utilize MPLS LSPs and did not receive any special handling instructions via control plane are NOT REQUIRED to inspect anything above the top label and will continue to perform basic label switching. If a node is enabled to perform additional actions on the packets it should leave the RFI/RFV tuple as received immediately after the top label swap on the stack. The PHP action may take place as usual exposing RFI LSE. In such cases egress LSR MUST be able to understand bSPL and either discard RFI/RFV tuple or if configured so execute special actions on those packets before further forwarding it.

[Discussion point for WG: Should nodes inserting additional labels on the stack for example during FRR should copy the RFI/RFV to enable it to be executed on the repair path or not ?].

Nodes which utilize the concept of SR-MPLS and use global labels as a replacement for use of LDP can apply the same set of procedures as nodes using MPLS-LSP.

Nodes which utilize the concept of SR-MPLS and use global labels with segment endpoints encoded in the label stack MUST understand RFI bSPL in order to correctly copy the tuple to always place it immediately behind the top most segment endpoint during label stack modification.

To further simplify the concept of MPLS RAF deployment networks which utilize concept of domain wide labels can allocate two label values from each LSR. One would indicate non RAF forwarding and the other one RAF enhanced forwarding. With such allocation only nodes which recognize that arriving packets contain RAF aware label value and which received any special handling instructions may need to perform additional RFI/RFV lookup and processing. Such allocation may be unified to differentiate normal vs RAF enabled labels with common

label block offset or selected label bit set.

7. Capability Advertisement

This document does not require any new capability to be defined.

8. IANA Considerations

This document defines a new bSPL label called Reference Forwarding Indicator and is requesting IANA to allocate one from Base Special-Purpose MPLS Label Values registry.

9. Security Considerations

This document does not introduce any new security risks.

10. Acknowledgements

Author would like to thank Tony Li and Jeff Tantsura for encouraging me to write this up.

11. References

11.1. Normative References

- [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>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", [RFC 3032](#), DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [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>>.

11.2. Informative References

[I-D.andersson-mpls-mna-fwk]

Andersson, L., Bryant, S., Bocci, M., and T. Li, "MPLS Network Actions Framework", Work in Progress, Internet-Draft, [draft-andersson-mpls-mna-fwk-00](https://www.ietf.org/archive/id/draft-andersson-mpls-mna-fwk-00), 4 April 2022, <<https://www.ietf.org/archive/id/draft-andersson-mpls-mna-fwk-00.txt>>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", [RFC 8986](https://www.rfc-editor.org/info/rfc8986), DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

Author's Address

Robert Raszuk (editor)
NTT NI
Email: robert@raszuk.net