

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 30, 2020

S. Bryant
Futurewei Technologies Inc.
G. Swallow
Southend Technical Center
S. Sivabalan
Cisco Systems
July 29, 2019

A Simple Control Protocol for MPLS SFLs
draft-bryant-mpls-sfl-control-04

Abstract

In [draft-ietf-mpls-sfl-framework](#) the concept of MPLS synonymous flow labels (SFL) was introduced. This document describes a control protocol that runs over an associated control header to request, withdrawn and extend the lifetime of such labels.

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 <http://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 January 30, 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 (<http://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.	Requirements Language	2
3.	SFL Control	2
3.1.	SFL Control Message	3
3.2.	SFL Control Procedures	6
3.2.1.	Request/Grant	6
3.2.2.	Refresh	8
3.2.3.	Withdraw	9
3.2.4.	Timer Accuracy	9
4.	Return Path	10
5.	Manageability Considerations	10
6.	Privacy Considerations	10
7.	Security Considerations	10
8.	IANA Considerations	10
9.	References	11
9.1.	Normative References	11
9.2.	Informative References	11
	Authors' Addresses	11

[1.](#) Introduction

In [[I-D.ietf-mpls-sfl-framework](#)] the concept of MPLS synonymous flow labels (SFL) was introduced. This document describes a simple control protocol that runs over an associated control header to request, withdrawn and extend the lifetime of such labels.

[2.](#) Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

[3.](#) SFL Control

This section describes the process by which the [RFC6374](#) Querier requests SFLs, the process by which the [RFC6374](#) Responder sends them to the Querier, and the process for managing the SFL lifetime. SFL Control Messages are carried over the SFL Control ACH. The SFL ACH is carried over a Pseudowire(PW) in place of the PW Control Word (CW), over an MPLS LSP using the GAL, or over some other mutually agreed path. Similarly the response may be returned over a PW, over

a bidirectional LSP or over some other mutually agreed path. See [Section 4](#).

3.1. SFL Control Message

The format of an SFL Control message, which follows the Associated Channel Header (ACH), is as follows:

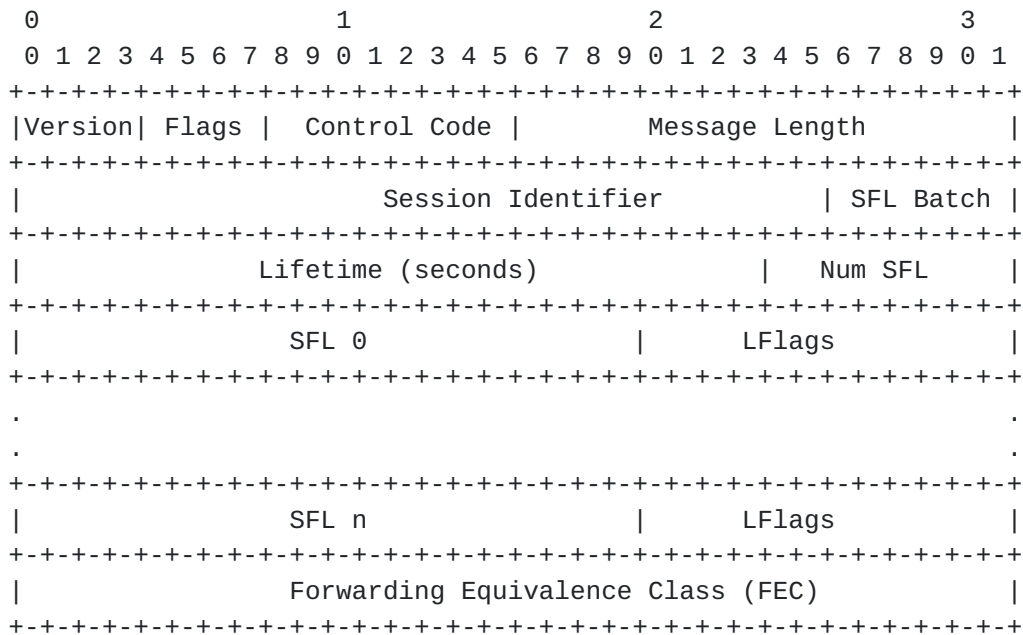


Figure 1: SFL Control Message Format

Reserved fields MUST be set to 0 and ignored upon receipt. The possible values for the remaining fields are as follows:

Version Protocol version. Set to zero in this specification.

Flags Message control flags.

Control Code Code identifying the query or response type.

Message Length Total length of this message in bytes.

Session Identifier Set arbitrarily by the querier and used as a message handle.

SFL Batch Used where the SFLs for this Session Identifier managed across multiple SFL Control Messages. A given set of SFLs MUST be retained in the same batch.

Lifetime The lifetime in seconds of the SFLs in this message.

In a Query message it is the requested lifetime. In a Response message it is the lifetime that the SFLs have been allocated for by the Responder. The Querier MUST NOT use an SFL after expiry of its lifetime, a Responder MUST make the SFL available for at least its lifetime.

Num SFL	The number of SFLs in this SFL Batch. This MUST be constant for the lifetime of the batch.
SFL n	The n'th SFL carried in this TLV. This is an MPLS label which is a component of a label stack entry as defined in Section 2.1 of [RFC3032] . The position of a label within a batch is constant for the lifetime of the batch. Enumeration starts at zero.
LFlags	The set of flags associated with the immediately preceding SFL. See below.
FEC	The Forwarding Equivalence Class that the SFLs in this TLV correspond to. This is encoded as per Section 3.4.1 of [RFC5036] .

Flags: The format of the Flags field is shown below.

```

+---+---+
|R|0|0|0|
+---+---+

```

SFL Control Message Flag

R: Query/Response indicator. Set to 0 for a Query and 1 for a Response.

0: Set to zero by the Sender and ignored by the Receiver.

Control Code: Set as follows according to whether the message is a Query or a Response as identified by the R flag.

For a Query:

Request This indicates that the responder is requested to allocate the set of SFLs marked with the R LFlag in this Message.

Refresh This indicates that the responder is requested to refresh the set of SFLs marked with the V LFlag in this Message.

Withdraw This indicates that the querier will no longer use the set of SFLs marked with the V Lflag and the responder may expire their lifetime.

For a Response:

Grant This indicates that the responder allocated the set of SFLs marked with the A LFlag in this Message.

Refresh-Ack This indicates that the responder has refreshed the set of SFLs marked with the V LFlag in this Message, and the lifetime is now as indicated by the lifetime field.

Withdraw-Ack This indicates that the responder has received the Withdraw message and will withdraw the SFLs

SFL-Unable The Responder was unable to satisfy the SFL Request. The details of the failure can be determined by comparing the Request and Grant messages.

Further error codes are for future study.

The LFlags field is defined as follows:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+
|0|1|2|3|          MBZ      |
+--+--+--+--+--+--+--+--+--+--+

```

Figure 2: LFlags Bit Definition

Where:

- 0 (Valid (V)) The Label value of the corresponding SFL is valid. In an SFL Request setting the V Lflag indicates a request for the specified label value. Where an SFL has a valid flag clear in a request message this indicates that any SFL value is acceptable.
- 1 (Request (R)) Indicates to the Querier that this member of the SFL batch is requested. Where a value is specified in the request, but the Responder is unable honour that request, no SFL is allocated and the corresponding A flag MUST be cleared.
- 2 (Allocated (A)) Indicates to the Querier that this SFL was allocated.
- 3 (Withdraw (W)) Indicates to the Responder that this SFL is to be withdrawn and to the Querier that the withdrawal has been carried out.

MBZ MUST be sent as zero and ignored on receive.

A flag value of one is true/set and a flag value of zero is false/clear. The use of these bits is described in more detail in the following sub-sections.

3.2. SFL Control Procedures

3.2.1. Request/Grant

To request a batch of SFLs the Querier constructs an SFL Control Request, encapsulates it in an SFL Control ACH and sends it to the Responder via an appropriate path. It sets the Control Message Flag to Query and the Control Code to Request. It chooses a session identifier as a handle for this transaction and as a way of binding this batch of SFLs to other operations that will use members of this SFL batch. Since members of the batch are treated as a group, the SFL Batch identifier is used to identify different SFL batches used in conjunction with the same session identifier.

The requested lifetime is set. This is the number of seconds from the time of the query to the time when the batch of SFLs will expire unless refreshed.

The Num SFL field is set to the SFL batch size.

Each SFL is set as follows: if a specific value is requested (for example for continuity across system restarts) this is written into the SFV n field and the V LFlag set. Otherwise, and including spare

SFLs where an allocation is not requested, the label value is set to zero and the V LFlag is cleared. For each SFL entry where an allocation is requested the R LFlag is set. All other LFlags are cleared.

The Forwarding Equivalence Class (FEC) is set to the FEC for which the SFLs are requested.

The Message Length is determined and filled in.

The Responder proceeds as follows:

It sets the control Message Flag to Response and initially sets the Control Code to Grant.

For each SFL with an R flag set, it determines whether it can honour the request, if so sets the A Lflag, and if the SFL value in the query was zero it overwrites it with the allocated SFL label value. In all other cases it leaves the SFL value and LFlag unchanged.

The lifetime field is updated with the lifetime of the SFLs if this is different from the requested lifetime.

All other fields in the Query message are left unchanged and the message is sent back to the Querier using the signaled or previously agreed message path.

Where the offered lifetime is other than the requested lifetime the Querier may accept the proposed value, or withdraw the SFLs and attempt to negotiate a new set of SFLs with a different lifetime.

If the Responder is unable to allocate all of the requested SFLs it MUST respond with a response code of SFL-Unable. The Querier MUST determine whether the allocated SFLs were adequate for its purposes and MUST send a withdraw if there are not adequate. A Querier MUST NOT attempt to hoard labels in the hope that the residual labels needed may become available in the future.

A Querier MUST wait a configured time (suggested wait of 60 seconds) before reattempting negotiation for a resource. Any failure to negotiate the required resources MUST be notified through the management interface of both Querier and Responder.

A Querier MUST NOT send an expired SFL to a Responder since to do so may invalidate another SFL operation.

3.2.2. Refresh

To request the lifetime refresh of a batch of SFLs the Querier constructs an SFL Refresh Request, encapsulates it in an SFL Control ACH and sends it to the Responder via an appropriate path. It sets the Control Message Flag to Query and the Control Code to Refresh. It uses the session identifier and the SFL Batch identifier that it used to request this SFL batch.

The requested lifetime is set. This is the number of seconds from the time of the query to the time when the batch of SFLs will expire unless refreshed.

The Num SFL field is set to the SFL batch size.

Each SFL is set as follows: the allocated SFL label value is written into the SFL n field and the V LFlag set. All other LFlags are cleared.

The Forwarding Equivalence Class (FEC) is set to the FEC for which the SFLs are requested.

The Message Length is determined and filled in.

The Responder proceeds as follows:

It sets the control Message Flag to Response and sets the Control Code to Refresh-Ack.

It sets the lifetime to the lifetime of the SFL.

All other fields in the Query message are left unchanged and the message is sent back to the Querier using the signaled or previously agreed message path.

Where the offered lifetime is other than the requested lifetime the Querier may accept the proposed value, or withdraw the SFLs and attempt to negotiate a new set of SFLs with a different lifetime.

A Querier **MUST** wait a configured time (suggested wait of 60 seconds) before reattempting negotiation for a resource. Any failure to negotiate the required resources **MUST** be notified through the management interface of both Querier and Responder.

3.2.3. Withdraw

To request the withdrawal of some or all of a batch of SFLs the Querier constructs an SFL Withdraw Request, encapsulates it in an SFL Control ACH and sends it to the Responder via an appropriate path. It sets the Control Message Flag to Query and the Control Code to Withdraw. It uses the session identifier and the SFL Batch identifier that it used to request this SFL batch.

The requested lifetime is set to zero.

The Num SFL field is set to the SFL batch size.

Each SFL being withdrawn is set as follows: the allocated SFL label value is written into the SFL n field and the V and W LFlags set. All other LFlags are cleared.

The Forwarding Equivalence Class (FEC) is set to the FEC for which the SFLs are requested.

The Message Length is determined and filled in.

The Responder proceeds as follows:

It sets the control Message Flag to Response and sets the Control Code to Withdraw-Ack.

All other fields in the Query message are left unchanged and the message is sent back to the Querier using the signaled or previously agreed message path.

A Querier MUST wait a configured time (suggested wait of 60 seconds) before reattempting a Withdraw request. No more than three Withdraw requests should be made.

3.2.4. Timer Accuracy

The lifetime of SFLs is expected to be sufficiently long that there are no significant constraints on timer accuracy. A node should be conservative in its assumptions concerning the lifetime of an SFL. A Querier MUST stop using a SFL significantly before the expiry of its lifetime and a Responder must maintain an SFL in active operation significantly beyond nominal expiry. A margin of the order of minutes is RECOMMENDED.

4. Return Path

Where the LSP is a multi-point to point, or multi-point to multi-point MPLS LSP (or other MPLS construct) the [RFC6374](#) Address TLV MUST be included in Query packet, even if the response is requested in-band, since this is needed to provide the necessary return address for this request.

5. Manageability Considerations

This may be provided in a future version of this memo.

6. Privacy Considerations

The inclusion of originating and/or flow information in a packet provides more identity information and hence potentially degrades the privacy of the communication. Whilst the inclusion of the additional granularity does allow greater insight into the flow characteristics it does not specifically identify which node originated the packet other than by inspection of the network at the point of ingress, or inspection of the control protocol packets. This privacy threat may be mitigated by encrypting the control protocol packets, regularly changing the synonymous labels and by concurrently using a number of such labels.

7. Security Considerations

It is assumed that this protocol is run in a well managed MPLS network with strict access controls preventing unwanted parties from generating MPLS OAM packets. The control protocol described in this memo thus introduced no additional MPLS security vulnerabilities.

8. IANA Considerations

As per the IANA considerations in [\[RFC5586\]](#), IANA is requested to allocate the following Channel Types in the "MPLS Generalized Associated Channel (G-ACh) Types" registry:

Value	Description	TLV Follows	Reference
0x0XXX	SFL Control	No	This

A value of 0x60 is suggested.

Force references to appear with mkd [\[RFC3032\]](#) [\[RFC5036\]](#)

9. References

9.1. 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>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", [RFC 5036](#), DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", [RFC 5586](#), DOI 10.17487/RFC5586, June 2009, <<https://www.rfc-editor.org/info/rfc5586>>.

9.2. Informative References

- [I-D.ietf-mpls-sfl-framework]
Bryant, S., Chen, M., Li, Z., Swallow, G., Sivabalan, S., and G. Mirsky, "Synonymous Flow Label Framework", [draft-ietf-mpls-sfl-framework-04](#) (work in progress), December 2018.

Authors' Addresses

Stewart Bryant
Futurewei Technologies Inc.

Email: stewart.bryant@gmail.com

George Swallow
Southend Technical Center

Email: swallow.ietf@gmail.com

Siva Sivabalan
Cisco Systems

Email: msiva@cisco.com