MPLS Working Group
Internet-Draft
Intended status: Informational
Expires: April 29, 2017

S. Bryant
M. Chen
Z. Li
Huawei
G. Swallow
S. Sivabalan
Cisco Systems
G. Mirsky
Ericsson

October 26, 2016

RFC6374 Synonymous Flow Labels
draft-bryant-mpls-rfc6374-sfl-02

Abstract

This document describes a method of providing flow identification information when making RFC6374 performance measurements. This allows RFC6374 measurements to be made on multi-point to point LSPs and allows the measurement of flows within an MPLS construct using RFC6374.

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 April 29, 2017.

Table of Contents

1. Introduction 2
2. Requirements Language 3
3. RFC6374 Packet Loss Measurement with SFL 3
3.1. RFC6374 SFL TLV 5
4. The Application of SFL to other PM Types 6
5. Privacy Considerations 6
6. Security Considerations 7
7. IANA Considerations 7
8. References 7
8.1. Normative References 7
8.2. Informative References 7
Authors' Addresses 8

1. Introduction

[I-D.ietf-mpls-flow-ident] describes the requirement for introducing flow identities when using RFC6374 [RFC6374] packet Loss Measurements (LM). In summary RFC6374 uses the LM packet as the packet accounting demarcation point. Unfortunately this gives rise to a number of problems that may lead to significant packet accounting errors in certain situations. For example:

1. Where a flow is subjected to Equal Cost Multi-Path (ECMP) treatment packets can arrive out of order with respect to the LM packet.

2. Where a flow is subjected to ECMP treatment, packets can arrive at different hardware interfaces, thus requiring reception of an LM packet on one interface to trigger a packet accounting action on a different interface which may not be co-located with it. This is a difficult technical problem to address with the required degree of accuracy.

3. Even where there is no ECMP (for example on RSVP-TE, MPLS-TP LSPs and PWs) local processing may be distributed over a number of processor cores, leading to synchronization problems.

4. Link aggregation techniques may also lead to synchronization issues.

5. Some forwarder implementations have a long pipeline between processing a packet and incrementing the associated counter again leading to synchronization difficulties.

An approach to mitigating these synchronization issue is described in [I-D.tempia-ippm-p3m] and [I-D.chen-ippm-coloring-based-ipfpm-framework] in which packets are batched by the sender and each batch is marked in some way such that adjacent batches can be easily recognized by the receiver.

An additional problem arises where the LSP is a multi-point to point LSP, since MPLS does not include a source address in the packet. Network management operations require the measurement of packet loss between a source and destination. It is thus necessary to introduce some source specific information into the packet to identify packet batches from a specific source.

[I-D.bryant-mpls-sfl-framework] specifies a method of encoding per flow instructions in an MPLS label stack using a technique called Synonymous Flow Labels (SFL) in which labels which mimic the behaviour of other labels provide the packet batch identifiers and enable the per batch packet accounting. This memo specifies how SFLs are used to perform RFC6374 performance measurements.

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. RFC6374 Packet Loss Measurement with SFL

The packet format of an RFC6374 Query message using SFLs is shown in Figure 1.

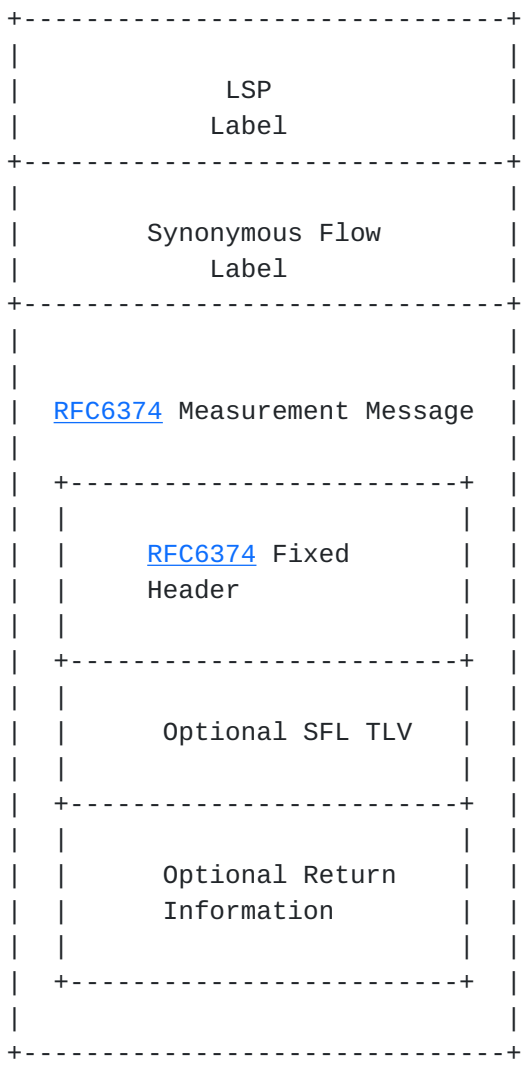

Figure 1: RFC6734 Query Packet with SFL

The MPLS label stack is exactly the same as that used for the user data service packets being instrumented except for the replacement of the appropriate label with an SFL . The RFC6374 measurement message consists of the three components, the RFC6374 fixed header as specified in [RFC6374] carried over the ACH channel type specified the type of measurement being made (currently: loss, delay or loss and delay) as specified in RFC6374.

Two optional TLVs MAY also be carried if needed. The first is the SFL TLV specified in Section 3.1. This is used to provide the implementation with a reminder of the SFL that was used to carry the RFC6374 message. This is needed because a number of MPLS implementations do not provide the MPLS label stack to the MPLS OAM handler. This TLV is required if RFC6374 messages are sent over UDP [RFC7876]. This TLV MUST be included unless, by some method outside

the scope of this document, it is known that this information is not needed by the RFC6374 Responder.

The second set of information that may be needed is the return information that allows the responder send the RFC6374 response to the Querier. This is not needed if the response is requested in-band and the MPLS construct being measured is a point to point LSP, but otherwise MUST be carried. The return address TLV is defined in RFC6378 and the optional UDP Return Object is defined in [RFC7876].

3.1. RFC6374 SFL TLV

[Editor's Note we need to review the following in the light of further thoughts on the associated signaling protocol(s). I am fairly confident that we need all the fields other than SFL Batch and SFL Index. The Index is useful in order to map between the label and information associated with the FEC. The batch is part of the lifetime management process.]

The required RFC6374 SFL TLV is shown in Figure 2. This contains the SFL that was carried in the label stack, the FEC that was used to allocate the SFL and the index into the batch of SLs that were allocated for the FEC that corresponds to this SFL.

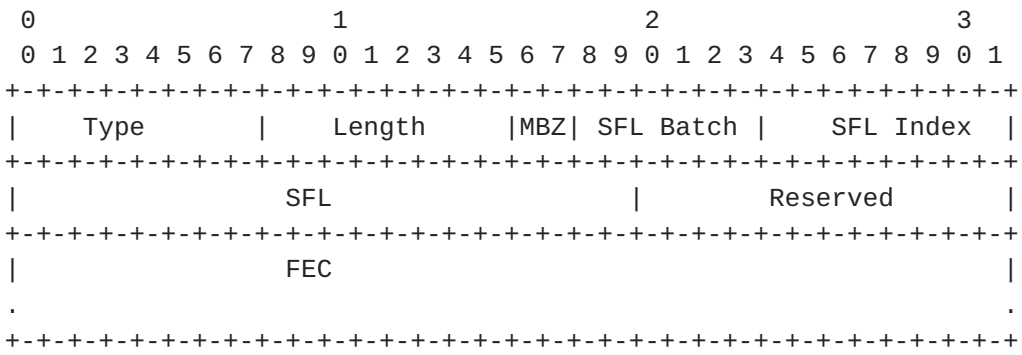

Figure 2: SFL TLV

Where:

Type Type is set to Synonymous Flow Label (SFL-TLV).

Length The length of the TLV as specified in RFC6374.

MBZ MUST be sent as zero and ignored on receive.

SFL Batch The SFL batch that this SFL was allocated as part of see [I-D.bryant-mpls-sfl-control]

SPL Index The index into the list of SFLs that were assigned against the FEC that corresponds to the SFL.

SFL The SFL used to deliver this packet. This is an MPLS label which is a component of a label stack entry as defined in Section 2.1 of [RFC3032].

Reserved MUST be sent as zero and ignored on receive.

FEC The Forwarding Equivalence Class that was used to request this SFL. This is encoded as per Section 3.4.1 of TBD

This information is needed to allow for operation with hardware that discards the MPLS label stack before passing the remainder of the stack to the OAM handler. By providing both the SFL and the FEC plus index into the array of allocated SFLs a number of implementation types are supported.

4. The Application of SFL to other PM Types

SFL can be used to enable other types of PM in addition to loss. Delay, Delay Variation and Throughput may be calculated based on measurement results collected through Loss and Delay Measurement test sessions. Further details will be provided in a future version of this draft.

5. 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.

6. Security Considerations

The issue noted in Section 5 is a security consideration. There are no other new security issues associated with the MPLS dataplane. Any control protocol used to request SFLs will need to ensure the legitimacy of the request.

7. IANA Considerations

IANA is request to allocate a new TLV from the 0-127 range on the MPLS Loss/Delay Measurement TLV Object Registry:

Type	Description	Reference
TBD	Synonymous Flow Label	This

A value of 4 is recommended.

8. References

8.1. 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, <http://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, <http://www.rfc-editor.org/info/rfc3032>.

[RFC7876] Bryant, S., Sivabalan, S., and S. Soni, "UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks", RFC 7876, DOI 10.17487/RFC7876, July 2016, <http://www.rfc-editor.org/info/rfc7876>.

8.2. Informative References

[I-D.bryant-mpls-sfl-control] Bryant, S., Swallow, G., and S. Sivabalan, "A Control Protocol for Synonymous Flow Labels", draft-bryant-mpls-sfl-control-00 (work in progress), March 2015.

[I-D.bryant-mpls-sfl-framework]
Bryant, S., Chen, M., Li, Z., Swallow, G., Sivabalan, S., and G. Mirsky, "Synonymous Flow Label Framework", draft-bryant-mpls-sfl-framework-02 (work in progress), October 2016.

[I-D.chen-ippm-coloring-based-ipfpm-framework]
Chen, M., Zheng, L., Mirsky, G., Fioccola, G., and T. Mizrahi, "IP Flow Performance Measurement Framework", draft-chen-ippm-coloring-based-ipfpm-framework-06 (work in progress), March 2016.

[I-D.ietf-mpls-flow-ident]
Bryant, S., Chen, M., Li, Z., Pignataro, C., and G. Mirsky, "MPLS Flow Identification Considerations", draft-ietf-mpls-flow-ident-02 (work in progress), August 2016.

[I-D.tempia-ippm-p3m]
Capello, A., Cociglio, M., Fioccola, G., Castaldelli, L., and A. Bonda, "A packet based method for passive performance monitoring", draft-tempia-ippm-p3m-03 (work in progress), March 2016.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <http://www.rfc-editor.org/info/rfc6374>.

Authors' Addresses

Stewart Bryant
Huawei

Email: stewart.bryant@gmail.com

Mach Chen
Huawei

Email: mach.chen@huawei.com

Zhenbin Li
Huawei

Email: lizhenbin@huawei.com

George Swallow
Cisco Systems

Email: swallow@cisco.com

Siva Sivabalan
Cisco Systems

Email: msiva@cisco.com

Gregory Mirsky
Ericsson

Email: gregory.mirsky@eicsson.com