

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 28, 2017

S. Bryant
M. Chen
Z. Li
Huawei
G. Swallow
S. Sivabalan
Cisco Systems
G. Mirsky
ZTE Corp.
G. Fioccola
Telecom Italia
April 26, 2017

RFC6374 Synonymous Flow Labels
draft-bryant-mpls-rfc6374-sfl-04

Abstract

This document describes a method of making RFC6374 performance measurements on flows carried over an MPLS Label Switched path. This allows loss and delay 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 October 28, 2017.

Copyright Notice

Copyright (c) 2017 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	4
3. RFC6374 Packet Loss Measurement with SFL	4
4. RFC6374 Single Packet Delay Measurement	4
5. Data Service Packet Delay Measurement	4
6. Some Simplifying Rules	6
7. Multiple Packet Delay Characteristics	6
7.1. Method 1: Time Buckets	7
7.2. Method 2 Classic Standard Deviation	9
7.2.1. RFC6374 Multi-Packet Delay Measurement Message Format	10
7.3. Per Packet Delay Measurement	11
7.4. Average Delay	11
8. Sampled Measurement	13
9. Carrying RFC6374 Packets over an LSP using an SFL	13
9.1. RFC6374 SFL TLV	15
10. Applicability to Pro-active and On-demand Measurement	16
11. RFC6374 Combined Loss-Delay Measurement	16
12. Privacy Considerations	16
13. Security Considerations	17
14. IANA Considerations	17
14.1. Allocation of PW Associated Channel Type	17
14.2. MPLS Loss/Delay TLV Object	17
15. References	17
15.1. Normative References	17
15.2. Informative References	18
Authors' Addresses	19

1. Introduction

[RFC6374] was originally designed for use as an OAM protocol for use with MPLS Transport Profile (MPLS-TP) [RFC5921] LSPs. MPLS-TP only supports point-to-point and point-to-multi-point LSPs. This document describes how to use RFC6374 in the general MPLS case, and also introduces a number of more sophisticated measurements of applicability to both cases.

[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 loss-measurement (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] describes 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 packet loss and RFC6374 delay 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 data service packets of the flow being instrumented are grouped into batches, and all the packets within a batch are marked with the SFL [I-D.ietf-mpls-flow-ident] corresponding to that batch. The sender counts the number of packets in the batch. When the batch has completed and the sender is confident that all of the packets in that batch will have been received, the sender issues an RFC6374 Query message to determine the number actually received and hence the number of packets lost. The RFC6374 Query message is sent using the same SFL as the co-responding batch of data service packets. The format of the Query and Response packet is described in Section 9.

4. RFC6374 Single Packet Delay Measurement

RFC6374 describes how to measure the packet delay by measuring the transit time of an RFC6374 packet over an LSP. Such a packet may not need to be carried over an SFL since the delay over a particular LSP should be a function of the TC bits.

However where SFLs are being used to monitor packet loss or where label inferred scheduling is used [RFC3270] then the SFL would be REQUIRED to ensure that the RFC6374 packet which was being used as a proxy for a data service packet experienced a representative delay. The format of an RFC6374 packet carried over the LSP using an SFL is shown in Section 9.

5. Data Service Packet Delay Measurement

Where it is desired to more thoroughly instrument a packet flow and to determine the delay of a number of packets it is undesirable to send a large number of RFC6374 packets acting as proxy data service packets Section 4. A method of directly measuring the delay characteristics of a batch of packets is therefore needed.

Given the long intervals over which it is necessary to measure packet loss, it is not necessarily the case that the batch times for the two measurement types would be identical. This it is proposed that the

two measurements are relatively independent. The notion that they are relatively independent arises for the potential for the two batches to overlap in time, in which case either the delay batch time will need to be cut short or the loss time will need to be extended to allow correct reconciliation of the various counters.

The problem is illustrated in Figure 1 below:

(1) AAAAAAAAAABBBBBBBBBBAAAAAAAAABBBBBBBBBB

SFL Marking of a packet batch for loss measurement

(2) AADDDDDAAAABBBBBBBBBBAAAAAAAAABBBBBBBBBB

SFL Marking of a subset if the packets for delay

(3) AAAAAAADDDDBBBBBBBBBBAAAAAAAAABBBBBBBBBB

SFL Marking of a subset of the packets across a packet loss measurement boundary

(4) AACDCDCDAABBBBBBBBBBAAAAAAAAABBBBBBBBBB

The case of multiple delay measurements within a packet loss measurement

Figure 1: RFC6734 Query Packet with SFL

In case 1 of Figure 1 we show the case were loss measurement alone is being carried out on the flow under analysis. For illustrative purposes consider that in the time interval being analyzed, 10 packets always flow.

Now consider case 2 of Figure 1 where a small batch of packets need to analyzed for delay. These are marked with a different SFL type indicating that they are to be monitored for both loss and delay. The SFL=A indicates loss batch A, SFL=D indicates a batch of packets that are to be instrumented for delay, but SFL D is synonymous with SFL A, which in turn is synonymous with the underlying FEC. Thus a packet marked D will be accumulated into the A loss batch, into the delay statistics and will be forwarded as normal. Whether the packet is actually counted twice (for loss and delay) or whether the two counters are reconciled during reporting is a local matter.

Now consider case 3 of Figure 1 where a small batch of packets are marked for delay across a loss batch boundary. These packets need to considered as part of batch A or a part of batch B, and any RFC6374

Query needs to take place after all the packets A or D (which ever option is chosen) have arrived at the receiving LSR.

Now consider case 4 of Figure 1. Here we have a case where it is required to take a number of delay measurements within a batch of packets that we are measuring for loss. To do this we need two SFLs for delay (C and D) and alternate between them (on a delay batch by delay batch basis) for the purposes of measuring the delay characteristics of the different batches of packets.

6. Some Simplifying Rules

It is possible to construct a large set of overlapping measurement type, in terms of loss, delay, loss and delay and batch overlap. If we allow all combination of cases, this leads to configuration, testing and implementation complexity and hence increased operation and capital cost. The following simplifying rules represent the default case:

1. Any system that needs to measure delay MUST be able to measure loss.
2. Any system that is to measure delay MUST be configured to measure loss. Whether the loss statistics are collected or not is a local matter.
3. A delay measurement MAY start at any point during a loss measurement batch, subject to rule 4.
4. A delay measurement interval MUST be short enough that it will complete before the enclosing loss batch completes.
5. The duration of a second delay (D in Figure 1 batch must be such that all packets from the packets belonging to a first delay batch (C in Figure 1) will have been received before the second delay batch completes.

Given that the sender controls both the start and duration of a loss and a delay packet batch, these rules are readily implemented in the control plane.

7. Multiple Packet Delay Characteristics

A number of methods are described. The expectation is that the MPLS WG possibly with the assistance of the IPPM WG will select one or maybe more than one of these methods for standardization.

Three Methods are discussed:

1. Time Buckets
2. Classic Standard Deviation
3. Average Delay

7.1. Method 1: Time Buckets

In this method the receiving LSR measures the inter-packet gap, classifies the delay into a number of delay buckets and records the number of packets in each bucket. As an example, if the operator were concerned about packets with a delay of up to 1us, 2us, 4us, 8us, and over 8us then there would be five buckets and packets that arrived up to 1us would cause the 1us bucket counter to increase, between 1us and 2us the 2us bucket counter would increase etc. In practice it might be better in terms of processing and potential parallelism if, when a packet had a delay relative to its predecessor of 2us both the up to 1us and the 2us counter were incremented and any more detailed information was calculated in the analytics system.

This method allows the operator to see more structure in the jitter characteristics than simply measuring the average jitter, and avoids the complication of needing to perform a per packet multiply, but will probably need to time intervals between buckets to be programmable by the operator.

The packet format of an RFC6374 Bucket Jitter Measurement Message is shown below:

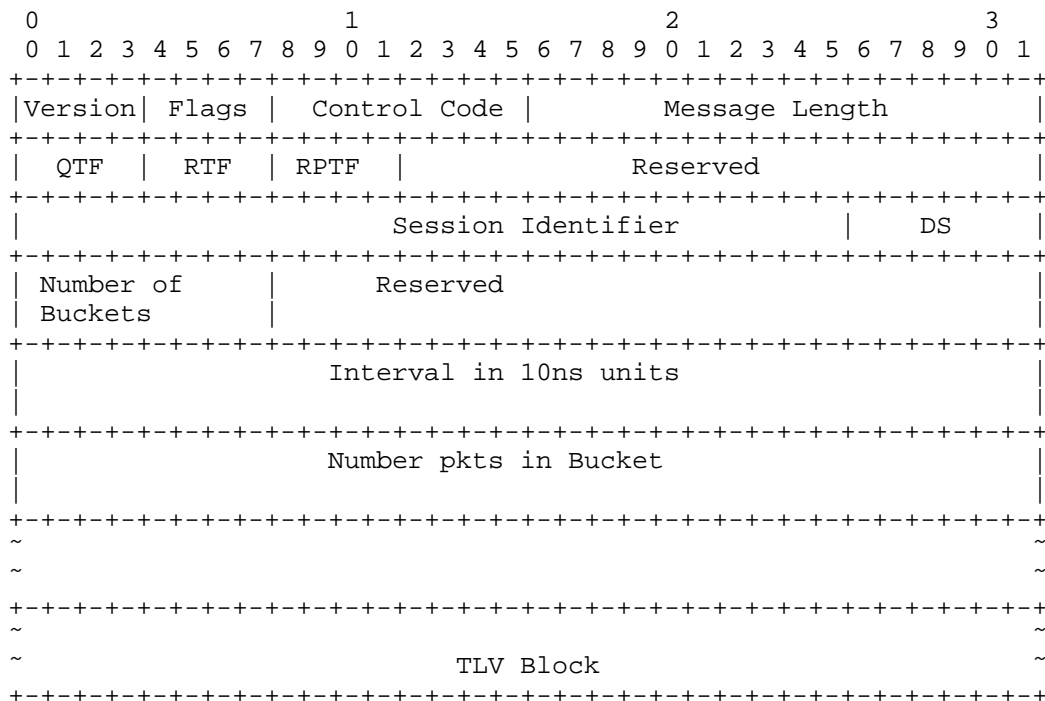


Figure 2: Bucket Jitter Measurement Message Format

The Version, Flags, Control Code, Message Length, QTF, RTF, RPTF, Session Identifier, and DS Fields are as defined in section 3.7 of RFC6374. The remaining fields are as follows:

- o Number of Buckets in the measurement
- o Reserved must be sent as zero and ignored on receipt
- o Interval in 10ns units is the inter-packet interval for this bucket
- o Number Pkts in Bucket is the number of packets found in this bucket.

There will be a number of Interval/Number pairs depending on the number of buckets being specified by the Querier. If an RFC6374 message is being used to configure the buckets, (i.e. the responder is creating or modifying the buckets according to the intervals in the Query message), then the Responder MUST respond with 0 packets in each bucket until it has been configured for a full measurement period. This indicates that it was configured at the time of the

last response message, and thus the response is valid for the whole interval. As per the [RFC6374] convention the Number of pkts in Bucket fields are included in the Query message and set to zero.

Out of band configuration is permitted by this mode of operation.

Note this is a departure from the normal fixed format used in RFC6374. We need to establish if this is problematic or not.

This RFC6374 message is carried over an LSP in the way described in [RFC6374] and over an LSP with an SFL as described in Section 9.

7.2. Method 2 Classic Standard Deviation

In this method, provision is made for reporting the following delay characteristics:

1. Number of packets in the batch (n).
2. Sum of delays in a batch (S)
3. Maximum Delay.
4. Minimum Delay.
5. Sum of squares of Inter-packet delay (SS).

Characteristic's 1 and 2 give the mean delay. Measuring the delay of each pair in the batch is discussed in Section 7.3.

Characteristics 3 and 4 give the outliers.

Characteristics 1, 2 and 5 can be used to calculate the variance of the inter-packet gap and hence the standard deviation giving a view of the distribution of packet delays and hence the jitter. The equation for the variance (var) is given by:

$$\text{var} = (SS - S*S/n)/(n-1)$$

There is some concern over the use of this algorithm for measuring variance, because SS and S*S/n can be similar numbers, particularly where variance is low. However the method commends it self by not requiring a division in the hardware. A future version of this document will look at using improved statistical methods such as the assumed mean.

7.2.1. RFC6374 Multi-Packet Delay Measurement Message Format

The packet format of an RFC6374 Multi-Packet Delay Measurement Message is shown below:

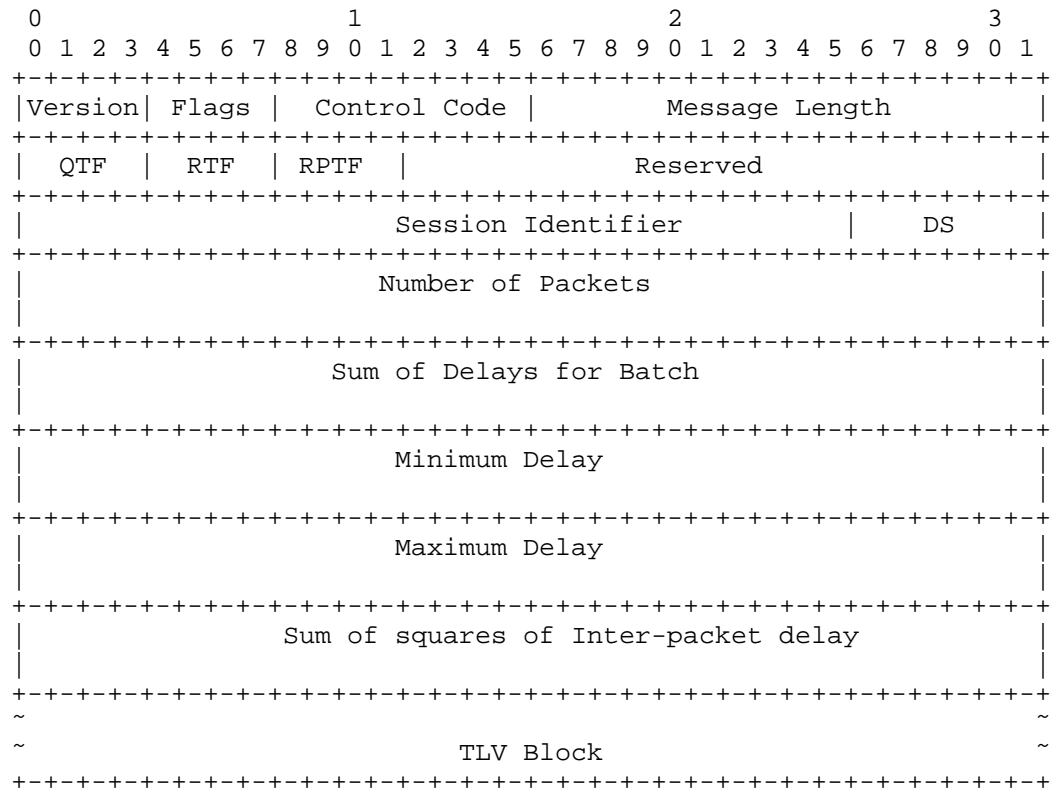


Figure 3: Multi-packet Delay Measurement Message Format

The Version, Flags, Control Code, Message Length, QTF, RTF, RPTF, Session Identifier, and DS Fields are as defined in section 3.7 of RFC6374. The remaining fields are as follows:

- o Number of Packets is the number of packets in this batch
- o Sum of Delays for Batch is the duration of the batch in the time measurement format specified in the RTF field.
- o Minimum Delay is the minimum inter-packet gap observed during the batch in the time format specified in the RTF field.
- o Maximum Delay is the maximum inter-packet gap observed during the batch in the time format specified in the RTF field.

This RFC6374 message is carried over an LSP in the way described in [RFC6374] and over an LSP with an SFL as described in Section 9.

7.3. Per Packet Delay Measurement

If detailed packet delay measurement is required then it might be possible to record the inter-packet gap for each packet pair. In other than exception cases of slow flows or small batch sizes, this would create a large demand on storage in the instrumentation system, bandwidth to such a storage system and bandwidth to the analytics system. Such a measurement technique is outside the scope of this document.

7.4. Average Delay

Introduced in [I-D.ietf-ippm-alt-mark] is the concept of a one way delay measurement in which the average time of arrival of a set of packets is measured. In this approach the packet is time-stamped at arrival and the Responder returns the sum of the time-stamps and the number of times-tamps. From this the analytics engine can determine the mean delay. An alternative model is that the Responder returns the time stamp of the first and last packet and the number of packets. This method has the advantage of allowing the average delay to be determined at a number of points along the packet path and allowing the components of the delay to be characterized.

The packet format of an RFC6374 Average Delay Measurement Message is shown below:

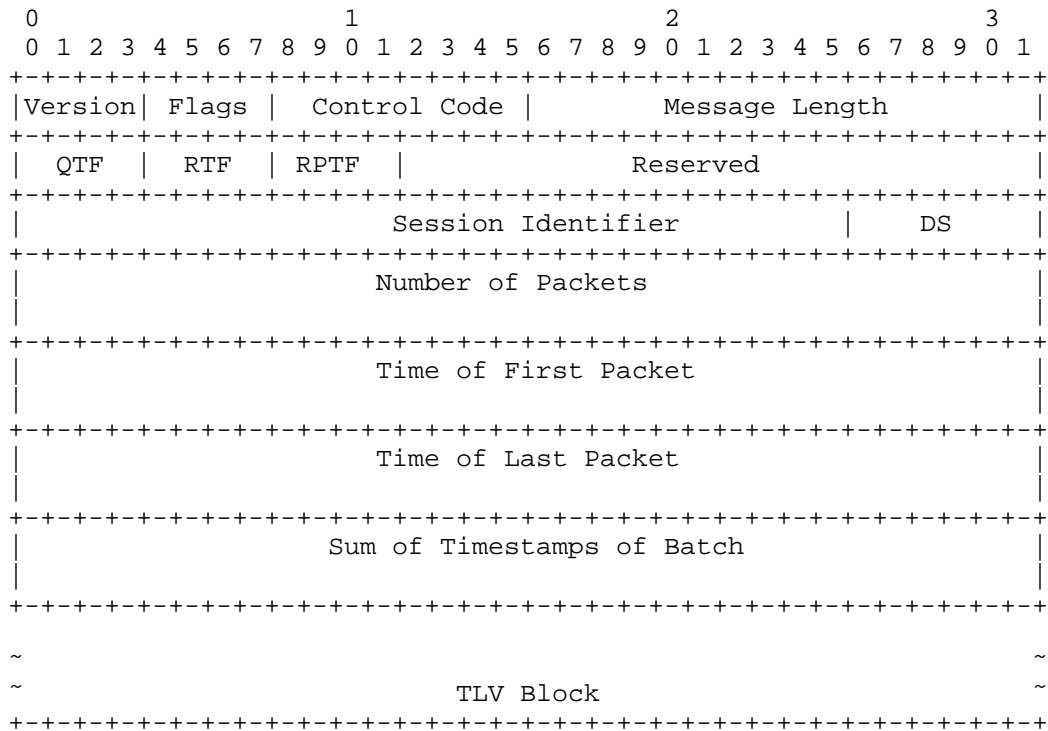


Figure 4: Average Delay Measurement Message Format

The Version, Flags, Control Code, Message Length, QTF, RTF, RPTF, Session Identifier, and DS Fields are as defined in section 3.7 of RFC6374. The remaining fields are as follows:

- o Number of Packets is the number of packets in this batch.
- o Time of First Packet is the time of arrival of the first packet in the batch.
- o Time of Last Packet is the time of arrival of the last packet in the batch.
- o Sum of Timestamps of Batch.

This RFC6374 message is carried over an LSP in the way described in [RFC6374] and over an LSP with an SFL as described in Section 9. As is the convention with RFC6374, the Query message contains placeholders for the Response message. The placeholders are sent as zero.

8. Sampled Measurement

In the discussion so far it has been assumed that we would measure the delay characteristics of every packet in a delay measurement interval defined by an SL of constant colour. In [I-D.ietf-ippm-alt-mark] the concept of a sampled measurement is considered. That is the Responder only measures a packet at the start of a group of packets being marked for delay measurement by a particular colour, rather than every packet in the marked batch. A measurement interval is not defined by the duration of a marked batch of packets but the interval between a pair of RFC6374 packets taking a readout of the delay characteristic. This approach has the advantage that the measurement is not impacted by ECMP effects.

9. Carrying RFC6374 Packets over an LSP using an SFL

The packet format of an RFC6374 Query message using SFLs is shown in Figure 5.

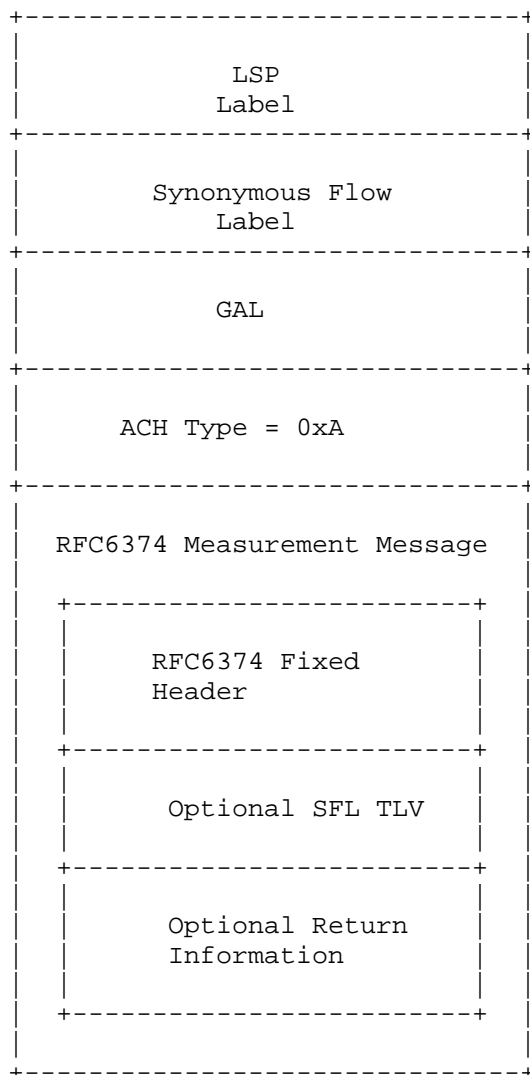


Figure 5: RFC6374 Query Packet with SFL

The MPLS label stack is exactly the same as that used for the user data service packets being instrumented except for the inclusion of the GAL [RFC5586] to allow the receiver to distinguish between normal data packets and OAM packets. Since the packet loss measurements are being made on the data service packets, an RFC6374 direct loss measurement is being made, and which is indicated by the type field in the ACH (Type = 0x000A).

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 9.1. 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].

9.1. 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 6. 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 6: 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.

10. Applicability to Pro-active and On-demand Measurement

A future version of the this document will discuss the applicability of the various methods to pro-active and on-demand Measurement.

11. RFC6374 Combined Loss-Delay Measurement

This mode of operation is not currently supported by this specification.

12. 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.

13. 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.

14. IANA Considerations

14.1. Allocation of PW Associated Channel Type

As per the IANA considerations in [RFC5586], IANA is requested to allocate the following Channel Type in the "PW Associated Channel Type" registry:

Value	Description	TLV Follows	Reference
TBD	RFC6374 Bucket Jitter Measurement	No	This
TBD	RFC6374 Multi-Packet Delay Measurement	No	This
TBD	RFC6374 Average Delay Measurement	No	This

14.2. MPLS Loss/Delay TLV Object

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.

15. References

15.1. 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>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, DOI 10.17487/RFC5586, June 2009, <<http://www.rfc-editor.org/info/rfc5586>>.
- [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>>.
- [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>>.

15.2. Informative References

- [I-D.bryant-mpls-sfl-control]
Bryant, S., Swallow, G., and S. Sivabalan, "MPLS Flow Identification Considerations", draft-bryant-mpls-sfl-control-01 (work in progress), March 2017.
- [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-04 (work in progress), April 2017.
- [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-ippm-alt-mark]

Fioccola, G., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate Marking method for passive performance monitoring", draft-ietf-ippm-alt-mark-04 (work in progress), March 2017.

[I-D.ietf-mpls-flow-ident]

Bryant, S., Pignataro, C., Chen, M., Li, Z., and G. Mirsky, "MPLS Flow Identification Considerations", draft-ietf-mpls-flow-ident-04 (work in progress), February 2017.

[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.

[RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, DOI 10.17487/RFC3270, May 2002, <<http://www.rfc-editor.org/info/rfc3270>>.

[RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, DOI 10.17487/RFC5921, July 2010, <<http://www.rfc-editor.org/info/rfc5921>>.

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.ietf@gmail.com

Siva Sivabalan
Cisco Systems

Email: msiva@cisco.com

Gregory Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

Giuseppe Fioccola
Telecom Italia

Email: giuseppe.fioccola@telecomitalia.it