

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 18, 2016

M. Chen, Ed.
L. Zheng, Ed.
Huawei Technologies
G. Mirsky, Ed.
Ericsson
G. Fioccola, Ed.
Telecom Italia
T. Mizrahi, Ed.
Marvell
March 17, 2016

IP Flow Performance Measurement Framework
draft-chen-ippm-coloring-based-ipfpm-framework-06

Abstract

This document specifies a measurement method, the IP flow performance measurement (IPFPM). With IPFPM, data packets are marked into different blocks of markers by changing one or more bits of packets. No additional delimiting packet is needed and the performance is measured in-service and in-band without the insertion of additional traffic.

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 <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 September 18, 2016.

Copyright Notice

Copyright (c) 2016 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. Terminology	4
3. Overview and Concept	4
4. Consideration on Marking Bits	6
5. Reference Model and Functional Components	6
5.1. Reference Model	6
5.2. Measurement Control Point	7
5.3. Measurement Agent	7
6. Period Number	8
7. Re-ordering Tolerance	8
8. Packet Loss Measurement	9
9. Packet Delay Measurement	10
10. Synchronization Aspects	11
10.1. Synchronization for the Period Number	12
10.2. Synchronization for Delay Measurement	12
11. IANA Considerations	13
12. Security Considerations	13
13. Acknowledgements	14
14. Contributing Authors	14
15. References	15
15.1. Normative References	15
15.2. Informative References	15
Authors' Addresses	17

1. Introduction

Performance Measurement (PM) is an important tool for service providers, used for Service Level Agreement (SLA) verification, troubleshooting (e.g., fault localization or fault delimitation) and network visualization. Measurement methods could be roughly put into two categories - active measurement methods and passive measurement

methods. Active methods measure performance or reliability parameters by the examination of traffic (IP Packets) injected into the network, expressly for the purpose of measurement by the intended measurement points. In contrast, passive method measures some performance or reliability parameters associated with the existing traffic (packets) on the network. Both passive and active methods have their strengths and should be regarded as complementary. There are certain scenarios where active measurement alone is not enough or applicable and passive measurement is desirable[I-D.deng-ippm-passive-wireless-usecase].

With active measurement methods, the rate, numbers and interval between the injected packets may affect the accuracy of the results. Moreover, injected test packets are not always guaranteed to be in-band with the data traffic in the pure IP network due to Equal Cost Multi-Path (ECMP).

The Multiprotocol Label Switching (MPLS) PM protocol [RFC6374] for packet loss could be considered an example of a passive performance measurement method. By periodically inserting auxiliary Operations, Administration and Maintenance (OAM) packets, the traffic is delimited by OAM packets into consecutive blocks, and the receivers count the packets and calculate the packets lost in each block. However, solutions like [RFC6374] depend on the fixed positions of the delimiting OAM packets for packets counting, and thus are vulnerable to out-of-order arrival of packets. This could happen particularly with out-of-band OAM channels, but might also happen with in-band OAM because of the presence of multipath forwarding within the network. Out of order delivery of data and the delimiting OAM packets can give rise to inaccuracies in the performance measurement figures. The scale of these inaccuracies will depend on data speeds and the variation in delivery, but with out-of-band OAM, this could result in significant differences between real and reported performance.

This document specifies a different measurement method, the IP flow performance measurement (IPFPM). With IPFPM, data packets are marked into different blocks of markers by changing one or more bits of packets without altering normal processing in the network. No additional delimiting packet is needed and the performance can be measured in-service without the insertion of additional traffic. Furthermore, because marking-based IP performance measurement does not require extra OAM packets for traffic delimitation, it can be used in situations where there is packet re-ordering. IP Flow Information eXport (IPFIX) [RFC7011] is used for reporting the measurement data of IPFPM to a central calculation element for performance metrics calculation. Several new Information Elements of

IPFIX are defined for IPFPM. These are described in the companion document [I-D.chen-ippm-ipfpm-report].

2. Terminology

The acronyms used in this document will be listed here.

3. Overview and Concept

The concept of marking IP packets for performance measurement is described in [I-D.tempia-opsawg-p3m]. Marking of packets in a specific IP flow to different colors divides the flows into different consecutive blocks. Packets in a block have same marking and consecutive blocks will have different markings. This enables the measuring node to count and calculate packet loss and/or delay based on each block of markers without any additional auxiliary OAM packets. The following figure (Figure 1) is an example that illustrates the different markings in a single IP flow in alternate 0 and 1 blocks.

```
| 0 Block | 1 Block | 0 Block | 1 Block |
00000000000 11111111111 00000000000 11111111111
```

Figure 1: Packet Marking

For packet loss measurement, there are two ways to mark packets: fixed packet numbers or fixed time period for each block of markers. This document considers only fixed time period method. The sender and receiver nodes count the transmitted and received packets/octets based on each block of markers. By counting and comparing the transmitted and received packets/octets, the packet loss can be computed.

For packet delay measurement, there are three solutions. One is similar to the packet loss, it still marks the IP flows to different blocks of markers and uses the time of the marking change as the reference time for delay calculations. This solution requires that there must not be any out-of-order packets; otherwise, the result will not be accurate. Because it uses the first packet of each block of markers for delay measurement, if there is packet reordering, the first packet of each block at the sender will be probably different from the first packet of the block at the receiver. An alternate way is to periodically mark a single packet in the IP flow. Within a given time period, there is only one packet that can be marked. The sender records the timestamp when the marked packet is transmitted, and the receiver records the timestamp when receiving the marked packet. With the two timestamps, the packet delay can be computed.

An additional method consists of taking into account the average arrival time of the packets within a single block (i.e. the same block of markers used for packet loss measurement). The network device locally sums all the timestamps and divides by the total number of packets received, so the average arrival time for that block of packets can be calculated. By subtracting the average arrival times of two adjacent devices it is possible to calculate the average delay between those nodes. This method is robust to out of order packets and also to packet loss (only an error is introduced dependent from the number of lost packets).

A centralized calculation element Measurement Control Point (MCP) is introduced in Section 5.2 of this document, to collect the packet counts and timestamps from the senders and receivers for metrics calculation. The IP Flow Information eXport (IPFIX) [RFC7011] protocol is used for collecting the performance measurement statistic information [I-D.chen-ippm-ipfpm-report]. For the statistic information collected, the MCP has to know exactly what packet pair counts (one from the sender and the other is from the receiver) are based on the same block of markers and a pair of timestamps (one from the sender and the other is from the receiver) are based on the same marked packet. In case of average delay calculation the MCP has to know in addition to the packet pair counters also the pair of average timestamps for the same block of markers. The "Period Number" based solution Section 6 is introduced to achieve this.

For a specific IP flow to be measured, there may be one or more upstream and downstream Measurement Agents (MAs) (Section 5.3). An IP flow can be identified by the Source IP (SIP) and Destination IP (DIP) addresses, and it may combine the SIP and DIP with any or all of the Protocol number, the Source port, the Destination port, and the Type of Service (TOS) to identify an IP flow. For each flow, there will be a flow identifier that is unique within a certain administrative domain. To simplify the process description, the flows discussed in this document are all unidirectional. A bidirectional flow can be seen as two unidirectional flows.

IPPFM supports the measurement of a Multipoint-to-Multipoint (MP2MP) model, which satisfies all the scenarios that include Point-to-Point (P2P), Point-to-Multipoint (P2MP), Multipoint-to-Point (MP2P), and MP2MP. The P2P scenario is obvious and can be used anywhere. P2MP and MP2P are very common in mobile backhaul networks. For example, a Cell Site Gateway (CSG) that uses multi-homing to two Radio Network Controller (RNC) Site Gateways (RSGs) is a typical network design. When there is a failure, there is a requirement to monitor the flows between the CSG and the two RSGs hence to determine whether the fault is in the transport network or in the wireless network (typically called "fault delimitation"). This is especially useful in the

situation where the transport network belongs to one service provider and the wireless network belongs to other service providers.

4. Consideration on Marking Bits

The marking bits selection is encapsulation-related; different bits for marking should be allocated by different encapsulations. This document does not define any marking bits. The marking bits selection for specific encapsulations will be defined in the relevant documents. In general, at least one marking bit is required to support loss and delay measurement. Specifically, if the second delay measurement solution is used (see Section 3), then at least two marking bits are needed; one bit for packet loss measurement, the other for packet delay measurement.

In theory, so long as there are unused bits that could be allocated for marking purpose, the marking-based measurement mechanism can be applied to any encapsulation. It is relatively easier for new encapsulations to allocate marking bits. An example of such a case is Bit Indexed Explicit Replication (BIER). Two marking bits for passive performance measurement has been allocated in the BIER encapsulation [I-D.ietf-bier-mpls-encapsulation] (Section 3.). However, for sophisticated encapsulations, it is harder or even impossible to allocate bits for marking purpose. The IPv4 encapsulation is one of the examples. The IPv6 encapsulation is in a similar situation, but for IPv6, an alternative solution is to leverage the IPv6 extension header for marking.

Since marking will directly change some bits (of the header) of the real traffic packets, the marking operations MUST NOT affect the forwarding and processing of packets. Specifically, the marking bits MUST NOT be used for ECMP hashing. In addition, to increase the accuracy of measurement, hardware-based implementation is desired. Thus, the location of the marking bits SHOULD be easy for hardware implementation. For example, the marking bits would be best located at fixed positions in a packet header.

5. Reference Model and Functional Components

5.1. Reference Model

The outline of the measurement system of large-scale measurement platforms (LMAP) is introduced in [I-D.ietf-lmap-framework]. It describes the main functional components of the LMAP measurement system, and the interactions between the components. The Measurement Agent (MA) of IPFPM could be considered equivalent to the MA of LMAP. The Measurement Control Point (MCP) of IPFPM could be considered as the combined function of Controller and Collector. The IP Flow

Information eXport (IPFIX) [RFC7011] protocol is used for collecting the performance measurement data on the MAs and reporting to the MCP. The details are specified in the companion document [I-D.chen-ippm-ipfpm-report]. The control between MCP and MAs are left for future study. Figure 2 presents the reference model of IPFPM.




Figure 2: IPFPM Reference Model

5.2. Measurement Control Point

The Measurement Control Point (MCP) is responsible for collecting the measurement data from the Measurement Agents (MAs) and calculating the performance metrics according to the collected measurement data. For packet loss, based on each block of markers, the difference between the total counts received from all upstream MAs and the total counts received from all downstream MAs are the lost packet numbers. The MCP must make sure that the counts from the upstream MAs and downstream MAs are related to the same marking/packets block. For packet delay (e.g., one way delay), the difference between the timestamps from the downstream MA and upstream MA is the packet delay. Similarly to packet loss, the MCP must make sure the two timestamps are based on the same marked packet. This document introduces a Period Number (PN) based synchronization mechanism which is discussed in details in Section 6.

5.3. Measurement Agent

The Measurement Agent (MA) executes the measurement actions (e.g., marks the packets, counts the packets, records the timestamps, etc.), and reports the data to the Measurement Control Point (MCP). Each MA maintains two timers, one (C-timer, used at upstream MA) is for marking change, the other (R-timer, used at downstream MA) is for reading the packet counts and timestamps. The two timers have the same time interval but are started at different times. An MA can be either an upstream or a downstream MA: the role is specific to an IP flow to be measured. For a specific IP flow, the upstream MA will change the marking and read the packet counts and timestamps when the C-timer expires, the downstream MA just reads the packet counts and

timestamps when the R-timer expires. The MA may delay the reading for a certain time period when the R-timer expires, in order to be tolerant to a certain degree of packet re-ordering. Section 7 describes this in details.

For each Measurement Task (corresponding to an IP flow) [I-D.ietf-lmap-framework], an MA maintains a pair of packet counters and a timestamp counter for each block of markers. As for the pair of packet counters, one is for counting packets and the other is for counting octets.

6. Period Number

When data is collected on the upstream MA and downstream MA, e.g., packet counts or timestamps, and periodically reported to the MCP, a certain synchronization mechanism is required to ensure that the collected data is correlated. Synchronization aspects are further discussed in Section 10. This document introduces the Period Number (PN) to help the MCP to determine whether any two or more packet counts (from distributed MAs) are related to the same block of markers, or any two timestamps are related to the same marked packet.

Period Numbers assure the data correlation by literally splitting the packets into different measurement periods. The PN is generated each time an MA reads the packet counts or timestamps, and is associated with each packet count and timestamp reported to the MCP. For example, when the MCP sees two PNs associated with two packet counts from an upstream and a downstream MA, it assumes that these two packet counts correspond to the same measurement period by the same PN, i.e., that these two packet counts are related to the same block of markers. The assumption is that the upstream and downstream MAs are time synchronized. This requires the upstream and downstream MAs to have a certain time synchronization capability (e.g., the Network Time Protocol (NTP) [RFC5905], or the IEEE 1588 Precision Time Protocol (PTP) [IEEE1588]), as further discussed in Section 10. The PN is calculated as the modulo of the local time (when the counts or timestamps are read) and the interval of the marking time period.

7. Re-ordering Tolerance

In order to allow for a certain degree of packet re-ordering, the R-timer on downstream MAs should be started Δt (Dt) later than the C-timer is started. Dt is a defined period of time and should satisfy the following conditions:

$$(\text{Time-L} - \text{Time-MRO}) < \text{Dt} < (\text{Time-L} + \text{Time-MRO})$$

Where

Time-L: the link delay time between the sender and receiver;

Time-MRO: the maximum re-ordering time difference; if a packet is expected to arrive at t_1 but actually arrives at t_2 , then the Time-MRO = $|t_2 - t_1|$.

Thus, the R-timer should be started at " $t + Dt$ " (where t is the time at which C-timer is started).

For simplicity, the C-timer should be started at the beginning of each time period. This document recommends the implementation to support at least these time periods (1s, 10s, 1min, 10min and 1h). Thus, if the time period is 10s, then the C-timer should be started at the time of any multiples of 10 in seconds (e.g., 0s, 10s, 20s, etc.), and the R-timer should be started, for example, at $0s+Dt$, $10s+Dt$, $20s+Dt$, etc. With this method, each MA can independently start its C-timer and R-timer given that the clocks have been synchronized.

8. Packet Loss Measurement

To simplify the process description, the flows discussed in this document are all unidirectional. A bidirectional flow can be seen as two unidirectional flows. For a specific flow, there will be an upstream MA and a downstream MA, and for each of these MAs there will be corresponding packet counts/timestamp.

For packet loss measurement, this document defines the following counters and quantities:

U-CountP[n][m]: U-CountP is a two-dimensional array that stores the number of packets transmitted by each upstream MA in each marking time period. Specifically, parameter "n" is the "period number" of measured blocks of markers while parameter "m" refers to the m-th MA of the upstream MAs.

D-CountP[n][m]: D-CountP is a two-dimensional array that stores the number of packets received by each downstream MA in each marking time period. Specifically, parameter "n" is the "period number" of measured blocks of markers while parameter "m" refers to the m-th MA of the downstream MAs.

U-CountO[n][m]: U-CountO is a two-dimensional array that stores the number of octets transmitted by each upstream MA in each marking time period. Specifically, parameter "n" is the "period number" of measured blocks of markers while parameter "m" refers to the m-th MA of the upstream MAs.

D-CountO[n][m]: D-CountO is a two-dimensional array that stores the number of octets received by each downstream MA in each marking time period. Specifically, parameter "n" is the "period number" of measured blocks of markers while parameter "m" refers to the m-th MA of the downstream MAs.

LossP: the number of packets transmitted by the upstream MAs but not received at the downstream MAs.

LossO: the total octets transmitted by the upstream MAs but not received at the downstream MAs.

The total packet loss of a flow can be computed as follows:

$$\text{LossP} = \text{U-CountP}[1][1] + \text{U-CountP}[1][2] + \dots + \text{U-CountP}[n][m] - \text{D-CountP}[1][1] - \text{D-CountP}[1][2] - \dots - \text{D-CountP}[n][m'].$$

$$\text{LossO} = \text{U-CountO}[1][1] + \text{U-CountO}[1][2] + \dots + \text{U-CountO}[n][m] - \text{D-CountO}[1][1] - \text{D-CountO}[1][2] - \dots - \text{D-CountO}[n][m'].$$

Where the m and m' are the number of upstream MAs and downstream MAs of the measured flow, respectively.

9. Packet Delay Measurement

For packet delay measurement, there will be only one upstream MA and may be one or more (P2MP) downstream MAs. Although the marking-based IPFPM supports P2MP model, this document only discusses P2P model. The P2MP model is left for future study. This document defines the following timestamps and quantities:

U-Time[n]: U-Time is a one-dimension array that stores the time when marked packets are sent; in case the "average delay" method is being used, U-Time stores the average of the time when the packets of the same block are sent; parameter "n" is the "period number" of marked packets.

D-Time[n]: D-Time is a one-dimension array that stores the time when marked packets are received; in case the "average delay" method is being used, D-Time stores the average of the time when the packets of the same block are received; parameter "n" is the "period number" of marked packets. This is only for P2P model.

D-Time[n][m]: D-Time a two-dimension array that stores the time when the marked packet is received by downstream MAs at each marking time period; in case the "average delay" method is being used, D-Time stores the average of the times when the packets of the same block are received by downstream MAs at each marking time period. Here,

parameter "n" is the "period number" of marked packets while parameter "m" refers to the m-th MA of the downstream MAs. This is for P2MP model which is left for future study.

One-way Delay[n]: The one-way delay metric for packet networks is described in [RFC2679]. The "n" identifies the "period number" of the marked packet.

One-way Delay[1] = D-Time[1] - U-Time[1].

One-way Delay[2] = D-Time[2] - U-Time[2].

...

One-way Delay[n] = D-Time[n] - U-Time[n].

In the case of two-way delay, the delay is the sum of the two one-way delays of the two flows that have the same MAs but have opposite directions.

Two-way Delay[1] = (D-Time[1] - U-Time[1]) + (D-Time'[1] - U-Time'[1]).

Two-way Delay[2] = (D-Time[2] - U-Time[2]) + (D-Time'[2] - U-Time'[2]).

...

Two-way Delay[n] = (D-Time[n] - U-Time[n]) + (D-Time'[n] - U-Time'[n]).

Where the D-Time and U-Time are for one forward flow, the D-Time' and U-Time' are for reverse flow.

10. Synchronization Aspects

As noted in the previous sections, there are two mechanisms in IPFPM that require MAs to have synchronized clocks: (i) the period number (Section 6), and (ii) delay measurement.

This section elaborates on the level of synchronization that is required for each of the two mechanisms. Interestingly, IPFPM can be implemented even with very coarse-grained synchronization.

10.1. Synchronization for the Period Number

Period numbers are used to uniquely identify blocks, allowing the MCP to match the measurements of each block from multiple MAs.

The period number of each measurement is computed by the modulo of the local time. Therefore, if the length of the measurement period is L time units, then all MAs must be synchronized to the same clock reference with an accuracy of +/- L/2 time units. This level of accuracy guarantees that all MAs consistently match the color bit to the correct block. For example, if the color is toggled every second (L = 1 second), then clocks must be synchronized with an accuracy of +/- 0.5 second to a common time reference.

The synchronization requirement for maintaining the period number can be satisfied even with a relatively inaccurate synchronization method.

10.2. Synchronization for Delay Measurement

As discussed in Section 9, the delay between two MAs is computed by $D-Time[1] - U-Time[1]$, requiring the two MAs to be synchronized.

Notably, two-way delay measurement does not require the two MAs to be time synchronized. Therefore, a system that uses only two-way delay measurement does not require synchronization between MAs.




Figure 3: Two-way Delay Measurement

As shown in Section 9, the two way delay between two MAs is given by (see Figure 3):

$$(D-Time[1] - U-Time[1]) + (D-Time'[1] - U-Time'[1])$$

Therefore, the two-way delay is equal to:

$$(D-Time'[1] - U-Time[1]) - (U-Time'[1] - D-Time'[1])$$

The latter implies that the two-way delay is comprised of two time differences, $(D\text{-Time}'[1] - U\text{-Time}[1])$, and $(U\text{-Time}'[1] - D\text{-Time}'[1])$. Thus, the value of the clocks of MA A and MA B does not affect the computation, and synchronization is not required.

11. IANA Considerations

This document makes no request to IANA.

12. Security Considerations

This document specifies a passive mechanism for measuring packet loss and delay within a Service Provider's network where the IP packets are marked using unused bits in IP head field, thus avoiding the need to insert additional OAM packets during the measurement. Obviously, such a mechanism does not directly affect other applications running on the Internet but may potentially affect the measurement itself.

First, the measurement itself may be affected by routers (or other network devices) along the path of IP packets intentionally altering the value of marking bits of packets. As mentioned above, the mechanism specified in this document is just in the context of one Service Provider's network, and thus the routers (or other network devices) are locally administered and this type of attack can be avoided.

Second, one of the main security threats in OAM protocols is network reconnaissance; an attacker can gather information about the network performance by passively eavesdropping to OAM messages. The advantage of the methods described in this document is that the color bits are the only information that is exchanged between the MAs. Therefore, passive eavesdropping to data plane traffic does not allow attackers to gain information about the network performance. We note that the information exported from the MAs to the MCP can be subject to eavesdropping, and thus it should be encrypted.

Finally, delay attacks are another potential threat in the context of this document. Delay measurement is performed using a specific packet in each block, marked by a dedicated color bit. Therefore, a man-in-the-middle attacker can selectively induce synthetic delay only to delay-colored packets, causing systematic error in the delay measurements. As discussed in previous sections, the methods described in this document rely on an underlying time synchronization protocol. Thus, by attacking the time protocol an attacker can potentially compromise the integrity of the measurement. A detailed discussion about the threats against time protocols and how to mitigate them is presented in RFC 7384 [RFC7384].

13. Acknowledgements

The authors would like to thank Adrian Farrel for his review, suggestion and comments to this document.

14. Contributing Authors

Hongming Liu
Huawei Technologies

Email: liuhongming@huawei.com

Yuanbin Yin
Huawei Technologies

Email: yinyuanbin@huawei.com

Rajiv Papneja
Huawei Technologies

Email: Rajiv.Papneja@huawei.com

Shailesh Abhyankar
Vodafone
Vodafone House, Ganpat Rao kadam Marg Lower Parel
Mumbai 40003
India

Email: shailesh.abhyankar@vodafone.com

Guangqing Deng
CNNIC
4 South 4th Street, Zhongguancun, Haidian District
Beijing
China

Email: dengguangqing@cnnic.cn

Yongliang Huang
China Unicom

Email: huangyl@dipmt.com

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>>.

15.2. Informative References

- [I-D.chen-ippm-ipfpm-report]
Chen, M., Zheng, L., Liu, H., Yin, Y., Papneja, R., Abhyankar, S., Deng, G., and Y. Huang, "IP Flow Performance Measurement Report", draft-chen-ippm-ipfpm-report-00 (work in progress), July 2014.
- [I-D.deng-ippm-passive-wireless-usecase]
Lingli, D., Zheng, L., and G. Mirsky, "Use-cases for Passive Measurement in Wireless Networks", draft-deng-ippm-passive-wireless-usecase-01 (work in progress), January 2015.
- [I-D.ietf-bier-mpls-encapsulation]
Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., and S. Aldrin, "Encapsulation for Bit Index Explicit Replication in MPLS Networks", draft-ietf-bier-mpls-encapsulation-03 (work in progress), February 2016.
- [I-D.ietf-lmap-framework]
Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A framework for Large-Scale Measurement of Broadband Performance (LMAP)", draft-ietf-lmap-framework-14 (work in progress), April 2015.
- [I-D.tempia-opsawg-p3m]
Capello, A., Cociglio, M., Castaldelli, L., and A. Bonda, "A packet based method for passive performance monitoring", draft-tempia-opsawg-p3m-04 (work in progress), February 2014.
- [IEEE1588]
IEEE, "1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", March 2008.

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<http://www.rfc-editor.org/info/rfc2474>>.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, DOI 10.17487/RFC2679, September 1999, <<http://www.rfc-editor.org/info/rfc2679>>.
- [RFC3260] Grossman, D., "New Terminology and Clarifications for Diffserv", RFC 3260, DOI 10.17487/RFC3260, April 2002, <<http://www.rfc-editor.org/info/rfc3260>>.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, DOI 10.17487/RFC4656, September 2006, <<http://www.rfc-editor.org/info/rfc4656>>.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, DOI 10.17487/RFC5357, October 2008, <<http://www.rfc-editor.org/info/rfc5357>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<http://www.rfc-editor.org/info/rfc5905>>.
- [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>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<http://www.rfc-editor.org/info/rfc7011>>.
- [RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, DOI 10.17487/RFC7384, October 2014, <<http://www.rfc-editor.org/info/rfc7384>>.

Authors' Addresses

Mach(Guoyi) Chen (editor)
Huawei Technologies

Email: mach.chen@huawei.com

Lianshu Zheng (editor)
Huawei Technologies

Email: vero.zheng@huawei.com

Greg Mirsky (editor)
Ericsson
USA

Email: gregory.mirsky@ericsson.com

Giuseppe Fioccola (editor)
Telecom Italia
Via Reiss Romoli, 274
Torino 10148
Italy

Email: giuseppe.fioccola@telecomitalia.it

Tal Mizrahi (editor)
Marvell
6 Hamada st.
Yokneam
Israel

Email: talmi@marvell.com