Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 4, 2015

M. Chen, Ed.
L. Zheng, Ed.
H. Liu
Y. Yin
R. Papneja
Huawei Technologies
G. Mirsky , Ed.
Ericsson
S. Abhyankar
Vodafone
G. Deng
CNNIC
Y. Huang
China Unicom

July 3, 2014

IP Flow Performance Measurement Report
draft-chen-ippm-ipfpm-report-00

Abstract

A "marking" based IP Flow Performance Measurement (IPFPM) framework is specified in draft-chen-ippm-coloring-based-ipfpm-framework. IP Flow Information eXport (IPFIX) is used for exporting the performance of IPFPM. Several new Information Elements of IPFIX are defined for IPFPM in this document.

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 January 4, 2015.

Table of Contents

1. Introduction

A "marking" based IP Flow Performance Measurement (IPFPM) framework is specified in [I-D.chen-ippm-coloring-based-ipfpm-framework]. This document is the companion document of [I-D.chen-ippm-coloring-based-ipfpm-framework]. The IPFPM framework describes a mechanism where data packets are marked so that they form blocks of data. No additional delimiting OAM is needed and the performance metrics 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 packets re-ordering. IP Flow Information eXport (IPFIX) [RFC7011] is used for exporting the performance of IPFPM. Several new Information Elements of IPFIX are defined for IPFPM in this document.

2. Measurement Data Report

In the IPFPM reference model, 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). The MCP is a centralized calculation element, responsible for collecting measurement data from MA and calculating the performance metrics according to the received measurement data from the MAs. During the performance measurement period, each MA reports performance measurementmeasurement data to the MCP, and the MCP will compute the performance measurement results according to the received measurement data. For a specific IP flow, for either packet delay or loss measurement, there will be at least one upstream and one downstream MA. For accurate measurements, time synchronization is required and the Period Number [I-D.chen-ippm-coloring-based-ipfpm-framework] is used by MCP to uniquely identify and correlate the packet counts/ timestamps between the upstream and downstream MAs for a specific block of markers or marked packet.

For packet loss measurement, the following information is required to report to the MCP:

MA Identifier

Flow Identifier (identify the flow to be measured)

Period Number

Packet Number Count

Packet Octets Count

For packet delay measurement, the following information is required to report to the MCP:

MA Identifier

Flow Identifier

Period Number

Timestamp

In addition, a MA may report some status (e.g., whether a MA is time synchronized) to the MCP, hence to help the MCP to determine whether measurement data from a MA is valid and can be used for computation. The following information may be included:

MA Identifier

MA Status

3. IPFIX Information Element for IPFPM

The IPFIX protocol [RFC7011] defines how IP Flow information can be exported from routers, measurement probes, or other devices. It defines many Information Elements [RFC7012] that can be used to carry and export the above information from MAs to MCP. Section 2 lists the statistic Information and status information need to be reported for IPFPM. Most of them can be identified with the existing Information Elements. New Information Element is defined for MA Identifer, Period Number and MA Status respectively in the following sections.

Flow Identifier: flowId (148)

Packet Number Count: packetTotalCount (86)

Packet Octets Count: octetTotalCount (85)

Timestamp: flowStartMicroseconds (154)

3.1. maIdentifier

Description: The maIdentifier is used to identify a MA. An maIdentifiler is unique within a specific administrative domain (e.g. within one MCP). The MA identifier can be generated and maintained by MCP. How to generate and maintain the maIdentifier is out the scope of this document.

Abstract Data Type: unsigned32

ElementId: TBD1

Status: current

3.2. periodNumber

Description: The periodNumber (PN) is used to identify a packet count or timestamp that belongs to a specific block of markers or marked packet. The MCP uses it 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. The PN is generated each time a MA reads the packet counts and timestamp, and is equal to the modulo of the local time (when the counts and timestamps are read) and the interval of the marking time period [I-D.chen-ippm-coloring-based-ipfpm-framework].

Abstract Data Type: unsigned32

ElementId: TBD2

Status: current

3.3. maStatus

Description: The maStatus is used to carry status information of a MA (For example, whether a MA has already time synchronized, whether is a upstream or downstream MA for a specific measured flow).

Abstract Data Type: unsigned16

ElementId: TBD3

Status: current

The maStaus is defined as follows:

Two bits are defined in this document. The T bit (Time synchronized bit) it is used to indicate whether a MA is time synchronized. When the T bit set, the MA is time synchronized; when the T bit is cleared, the MA is not time synchronized. The MCP MUST calculate the results when all related MAs of a flow are time synchronized, otherwise, the results will not correct. The U bit (Upstream MA bit) is used to indicate whether a MA is the upstream or downstream TLP for an IP flow. When the U bit set, the MA is the upstream MA of the flow; otherwise, the MA is the downstream MA of the flow.

4. Templates for IPFPM

4.1. MA Status Options Template

The MA Status Options Template is used to report the status of a MA; it SHOULD contain the following Information Elements:

meteringProcessId (scope) [IPFIX-IANA]

MA Identifier

The maIdentifier is as defined in Section 3.1 of this document.

MA Status

The maStatus is as defined in Section 3.3 of this document.

The Data Records specified by the MA Status Options Template SHOULD be exported once the IPFIX session established or when status changed.

An example of the MA Status Options Template Set is as follows:

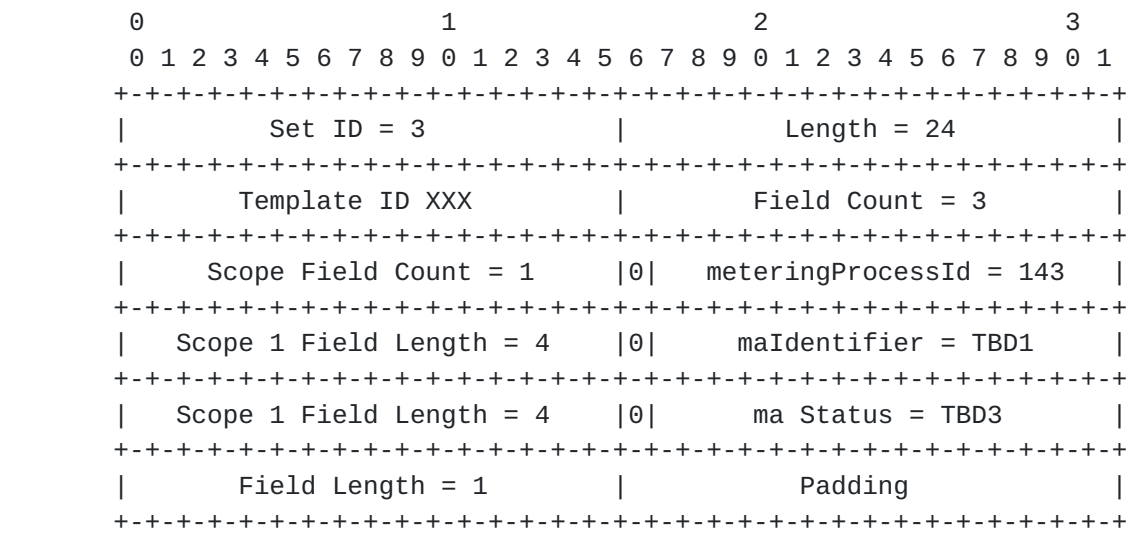

4.2. Flow Options Template

The Flow Options Template is used to report all the configured flows (to be measured) on a MA. It SHOULD include the following the Information Elements:

meteringProcessId (scope) [IPFIX-IANA]

MA Identifier

The maIdentifier is as defined in Section 3.1 of this document.

maStatus [Section 3.3]

flowId [IPFIX-IANA]

protocolIdentifier [IPFIX-IANA]

Source IP address

The source IP address or prefix of an IP flow, for this address, any of the following Information Elements can be used:

sourceIPv4Address [IPFIX-IANA]

sourceIPv6Address [IPFIX-IANA]

sourceIPv4Prefix [IPFIX-IANA]

sourceIPv6Prefix [IPFIX-IANA]

Source IP prefix length

The source IP prefix length of a prefix, any of the following Information Elements can be used:

sourceIPv4PrefixLength [IPFIX-IANA]

sourceIPv6PrefixLength [IPFIX-IANA]

Source port

The source port of an IP flow, any of the following Information Elements can be used:

udpSourcePort [IPFIX-IANA]

tcpSourcePort [IPFIX-IANA]

Destination IP address

The destination IP address or prefix of an IP flow, for this address, any of the following Information Elements can be used:

destinationIPv4Address [IPFIX-IANA]

destinationIPv6Address [IPFIX-IANA]

destinationIPv4Prefix [IPFIX-IANA]

destinationIPv6Prefix [IPFIX-IANA]

Destination IP prefix length

The destination IP prefix length of a prefix, any of the following Information Elements can be used:

destinationIPv4PrefixLength [IPFIX-IANA]

destinationIPv6PrefixLength [IPFIX-IANA]

Destination port

The destination port of an IP flow, any of the following Information Elements can be used:

udpDestinationPort [IPFIX-IANA]

tcpDestinationPort [IPFIX-IANA]

The Data Records specified by the Flow Options Template SHOULD be exported once the IPFIX session established or when the configured flows changed (e.g., a new flow is added for measurement or a flow deleted to stop the measurement).

An example of the Flow Options Template Set is as follows:

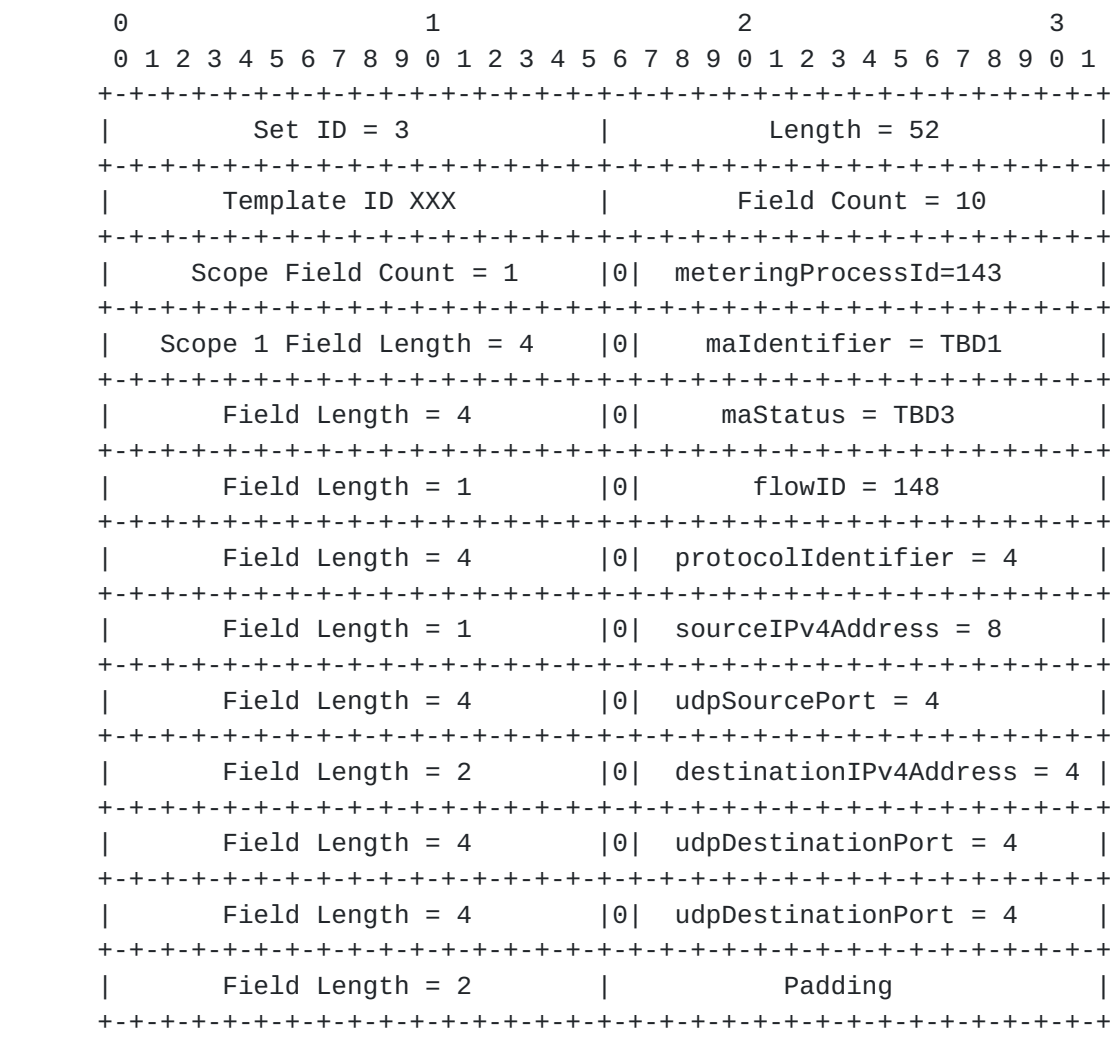

4.3. Packet Loss Template

The Packet Loss Template is used by a MA to report the packet loss measurement statistic of a flow to the MCP; it SHOULD contain the following Information Elements:

MA Identifier

The maIdentifier is as defined in Section 3.1 of this document.

flowId[IPFIX-IANA]

periodNumber [Section 3.2]

packetTotalCount[IPFIX-IANA]

octetTotalCount[IPFIX-IANA]

An example of the Packet Loss Data Set is as follows:

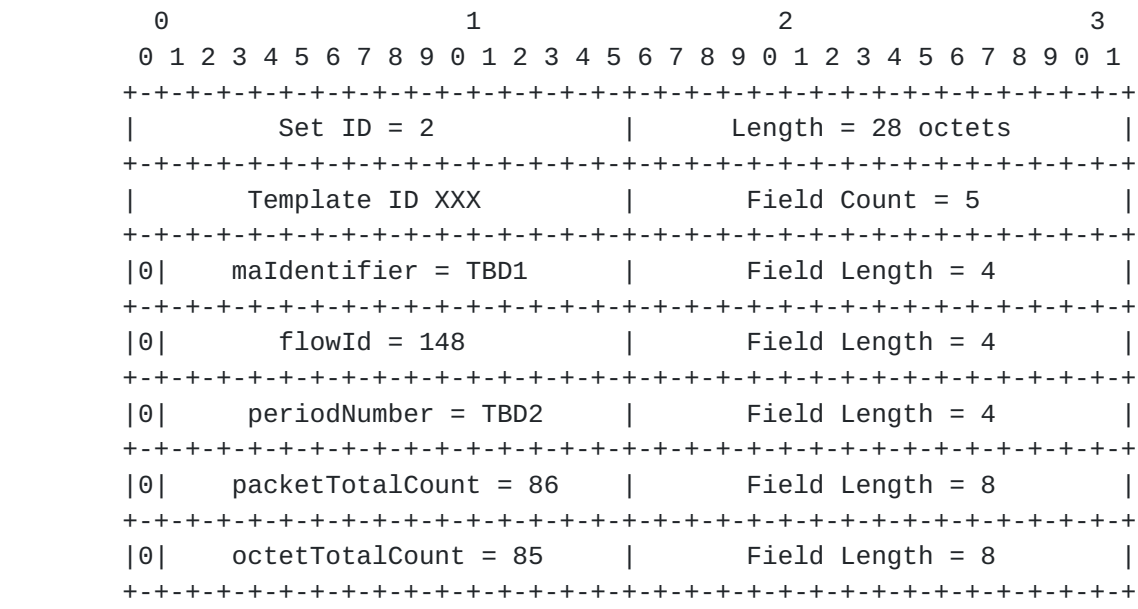

The maIdenfier is used to identify and MA.

The flowId is a identifier that is unique within a specific administrative domain (e.g., within one MCP). The MA and MCP have to agree a flow identifier related to a specific flow. For example, the flow identifier can be generated and maintained by a centralized element. How to generate and maintain the flowId is out the scope of this document.

The periodNumber is as defined in Section 3.2 of this document.

The packetTotalCount is used to carry the total transmitted/received packets of a flow since the measurement start.

The octetTotalCount is used to carry the total transmitted/received octets of a flow since the measurement start.

4.4. Packet Delay Template

The Packet Delay Template is used by a MA to report the packet delay measurement statistic of a flow to the MCP; it SHOULD contain the following Information Elements:

MA Identifier

The maIdentifier is as defined in Section 3.1 of this document.

flowId [IPFIX-IANA]

periodNumber [Section 3.2]

timestamp

The time when marked a packet, flowStartMicroseconds is used to carry the timestamp:

flowStartMicroseconds[IPFIX-IANA]

An example of the Packet Delay Data Set is as follows:

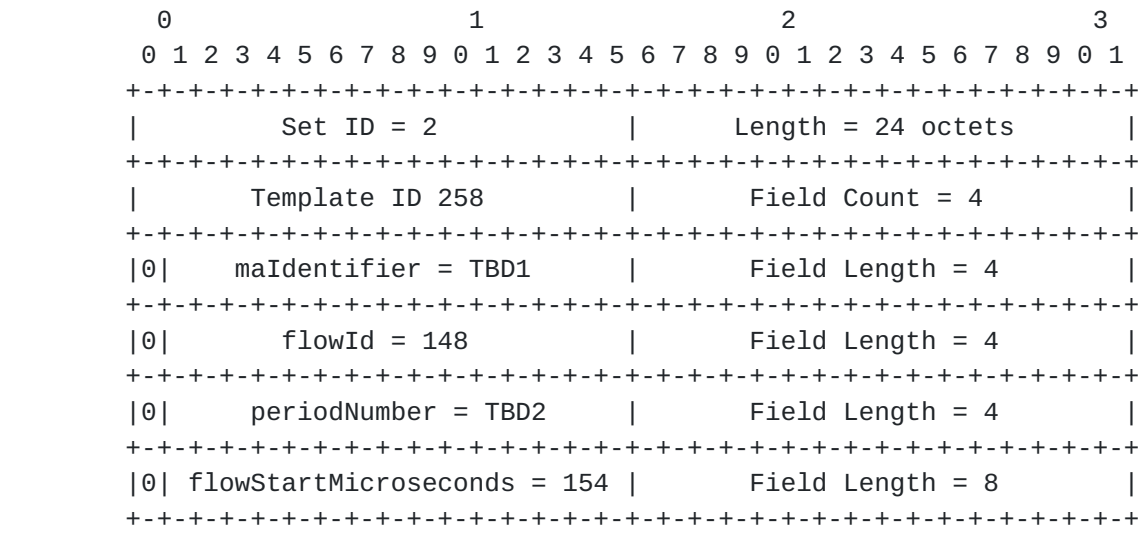

The maIdenfier is used to identify and MA.

The flowId is used to carry the flow identifier of a flow; the structure is defined in Section 4.3 of this document.

The periodNumber is as defined in Section 3.2 of this document.

The flowStartMicroseconds is used to carry the timestamp of a marked packet of a specific flow.

5. IANA Considerations

The IANA is requested to allocate 3new Information Elements codes for the Information Elements defined in Section 3 from the IPFIX Information Elements registry.

6. Security Considerations

This document does not bring new security issue to IPFIX.

7. Acknowledgements

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

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013.

[RFC7012] Claise, B. and B. Trammell, "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, September 2013.

8.2. Informative References

[I-D.chen-ippm-coloring-based-ipfpm-framework] Chen, M., Liu, H., Yin, Y., Papneja, R., Abhyankar, S., Deng, G., and Y. Huang, "Coloring based IP Flow Performance Measurement Framework", draft-chen-ippm-coloring-based-ipfpm-framework-01 (work in progress), October 2013.

[IPFIX-IANA] "http://www.iana.org/assignments/ipfix/ipfix.xml", .

Authors' Addresses

Mach(Guoyi) Chen (editor)
Huawei Technologies

Email: mach.chen@huawei.com

Lianshu Zheng (editor)
Huawei Technologies

Email: vero.zheng@huawei.com

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

Greg Mirsky (editor)
Ericsson
USA

Email: gregory.mirsky@ericsson.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