

IPPM Working Group
Internet-Draft
Intended Status: Informational

Madhukar Anand
Sanjoy Bardhan
Radhakrishna Valiveti
Infinera Corporation

Ramesh Subrahmaniam
Individual

Carlos Pignataro
Shwetha Bhandari
Randy Zhang
Rajiv Asati
Cisco Systems, Inc

Expires: February 22, 2019

August 21, 2018

**Integrated Packet-Optical In-Situ OAM
draft-anand-ippm-po-ioam-01**

Abstract

This document proposes a way to extend in-situ OAM techniques to include operational data from multiple network layers with a view to create an integrated record of OAM information as the data flows between two network entities. An instance of this technique that is elaborated here focuses on packet-optical networks that are traditionally transport centric. The mechanisms described are general enough to allow future extensibility of in-situ OAM techniques into other non-packet domains.

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 to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2018 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	3
2	Reference Terminology	3
3	Packet Optical OAM Integration	4
4	Mechanism Assumptions and Overview	5
5	Packet-Optical IOAM Data Types and Formats	6
5.1	IOAM Packet-Optical Flow Tracing	6
5.2	IOAM Optical Path Characteristics	7
6	Summary	10
7	Security Considerations	10
8	IANA Considerations	10
9	References	10
9.1	Normative References	10
9.2	Informative References	11
	Authors' Addresses	11

1 Introduction

Packet and optical transport networks have evolved independently with different OAM mechanisms that have to be managed separately. Consequently, coordinating packet and optical networks for delivering services such as end-to-end traffic engineering has proved challenging. To address this challenge, a unified paradigm that provides an integrated record of OAM information as the data flows between two network entities that are connected by a multi-layer network is needed. Further, such a paradigm should provide an incremental path to complete packet-optical OAM integration while leveraging existing OAM mechanisms as much as possible in either domains. This document introduces such a paradigm by extending the mechanisms defined in [[I-D.ietf-ippm-ioam-data](#)].

The key construct to deliver an integrated IOAM across packet and optical network layers is that of a Multi-layer Border Device (MLBD). These are nodes in the network that map packet paths to the optical paths that originate and terminate at the MLBDs. The concept of MLBD allows for multiple realizations - depending on whether the packet and optical functions are integrated or not. In one case, the packet device is distinct from the optical transport device, and the MLBD is a logical entity that spans these two devices. In this case, the MLBD functionality is achieved with the help of external coordination between the packet and optical devices. In another case, the packet and optical components are integrated into one physical device, and the co-ordination required for functioning of the MLBD is performed by this integrated device. It must be noted that in either case, it is the packet/optical data plane that is either disaggregated or integrated. Control of the devices can be logically centralized or distributed in either scenario. The focus of this document is to define the logical functions of an MLBD without going into the exact instantiations of the concept.

2. Reference Terminology

IOAM: In-situ Operations, Administration, and Maintenance

MLBD: Multi-layer Border Device

POT: Proof of Transit

FEC: Forward Error Correction

BER: Bit Error Rate

SRLG: Shared Risk Link Group

3. Packet Optical OAM Integration

Many operators build and operate their networks that are both multi-layer and multi-domain and provide end-to-end services to their customers. Due to the nature of the different domains, the operations and management of the network has always been problematic and time consuming.




Figure 1: Representation of a packet-optical network

Figure 1 represents a packet optical network in which P1 and P2 are packet devices that are connected via two optical links comprising of devices 01 and 02 and with devices 03,...,06. Devices P1 and P2 are MLBDs that communicate with other packet devices and also with the devices in the optical transport domain. Each MLBD would maintain operational data such as path latency, Bit Error Rate (BER), pre-FEC counts, FEC corrected words, and Q-factor corresponding to active optical paths such as {P1, 01, 02, P2}. An MLBD, by the virtue of being a packet device, also participates in the in-situ OAM techniques as described in [[I-D.ietf-ippm-ioam-data](#)]. To facilitate the OAM integration of packet and optical network layers, an MLBD parses the IOAM data fields in a packet earmarked for optical OAM data and inserts the appropriate OAM information from the optical path corresponding to

that packet flow. The operations data corresponding to the optical paths may be obtained by leveraging supported optical OAM techniques.

It must be noted that the IOAM changes proposed in this document are limited to necessary optical characteristics that are of interest to the packet domain applications and have the potential to be used for routing/switching and traffic engineering. It is not intended to be a mechanism to obtain all supported optical OAM information from optical devices.

4. Mechanism Assumptions and Overview

The current proposal assumes that the packet and optical network layers use respective OAM techniques without any modification. There are also no modifications necessary to data forwarding across the layers. The only requirement is that an MLBD supports an embodiment of IOAM mechanisms. It is also assumed that the MLBD performs all functions of a regular packet IOAM device and there are specific data types allocated for querying optical OAM data (IOAM optical data type). The mechanism for supporting the multi-layer IOAM is as follows.

1. An MLBD participates in an in-situ OAM-domain and thus behaves either as an IOAM transit device, an IOAM encapsulating or an IOAM decapsulating device as described in [[I-D.ietf-ippm-ioam-data](#)].

2. An MLBD participates in all relevant optical OAM mechanisms and obtains operational data.

3. MLBDs interpret the IOAM optical data type in a data flow, map it to specific optical operational data, and attach the desired response to the data packet in an appropriate manner. Specific data formats for carrying the optical operational data are discussed in the subsequent sections. There are typically multiple MLBDs at the edges of a optical network that can process the IOAM optical data type, and depending on the use-case and the type of optical data requested, a specific MLBD or a set of MLBDs can respond to the requested data.

- (a) Mapping of an optical data type to a specific optical operational data is assumed to be programmed into an MLBD prior to its use.

- (b) It is possible that multiple packets share the same optical characteristics. In those circumstances, all the associated packet flows will share the optical OAM data. For example, if traffic from two VLANs is multiplexed into the same ODU, the ODU-level operational data is applicable to both the VLANs.

- (c) If a specific optical data type maps to multiple optical

operational data, an MLBD can determine how to communicate that suitably in response to the request. Options for data consolidation may include appending all responses, averaging the responses, or picking minimum or maximum value, or keeping them separate.

Considering the topology in Figure 1, if there is a query for the aggregate number of post-FEC errors in the optical path(s) between P1 and P2, P1 may choose to send the requested data as an appended list {FEC1, FEC2} in response, where FEC1 are the aggregate number of post-FEC errors along the optical path/circuit {P1,01,02,P2} and FEC2 are the aggregate number of post-FEC errors along the path/circuit {P1,03,04,05,06,P2}.

(d) If the requested data is not available immediately, the optical device can request a collection of the requested data upon parsing the request, and once it is available, the data could be included in a later packet carrying the same data request. One option to communicate that the requested data is delayed is by responding to the request with a standardized special value such as 0xFFFFFFFF that indicates this scenario. It is to be noted that this special value is something that is assumed to be standardized for the specific data queried and known in advance by all participating IOAM devices.

Considering the example in Figure 1, if there is a query for the optical path latency in the optical path(s) between P1 and P2, and if P1 only has the data LAT1 available for the optical path/circuit {P1,01,02,P2} and not for the path/circuit {P1,03,04,05,06,P2}, it can communicate {LAT1, 0xFFFFFFFF} in response till the path latency data LAT2 is available for the path {P1,03,04,05,06,P2}. Subsequently, it can communicate {LAT1, LAT2} in response.

4. Optical OAM data collected in the packet is decapsulated and exported using the same mechanism as that used in the packet IOAM domain.

5. Finally, if an IOAM optical data type field cannot be interpreted by a specific MLBD, it acts like a bypass for that data flow.

5. Packet-Optical IOAM Data Types and Formats

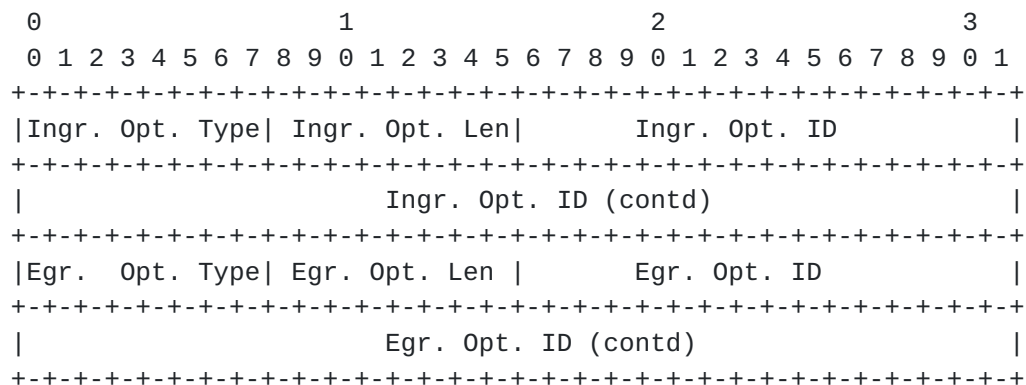
5.1 IOAM Packet-Optical Flow Tracing

Tracing of packet-optical flows can follow either the pre-allocated or incremental trace options as defined in [Section 4.1](#) of [I-D.ietf-ippm-ioam-data]. The data type and format follow that specification with the following modifications at an MLBD:

- o Identification of the interface that a packet was received on,

i.e. ingress interface or optical path/circuit/wavelength ID(s)

o Identification of the interface that a packet was sent out on,
i.e. egress interface or optical path/circuit/wavelength ID(s)



Ingr. Opt. Type : 8-bit identifier that identifies the specific type of ingress optical path/circuit/wavelength identifier(s).

Ingr. Opt. Len : Length of data field in bytes (4-byte aligned).

Ingr. Opt. ID : Unsigned integer. Interface identifier of the ingress optical path/circuit/wavelength identifier(s).

Egr. Opt. Type : 8-bit identifier that identifies the specific type of egress optical path/circuit/wavelength identifier(s).

Egr. Opt. Len : Length of data field in bytes (4-byte aligned).

Egr. Opt. ID : Unsigned integer. Interface identifier of the egress optical path/circuit/wavelength identifier(s).

Considering the topology in Figure 1, in response to a query for path tracing along the optical path/circuit {P1,01,02,P2}, the MLBD P1 could insert 0x01 as the egress optical type, and egress optical length of 4 and path identifier of 0x0A010101 corresponding to path {P1,01,02,P2}.

5.2 IOAM Optical Path Characteristics

An MLBD can also be queried for any optical path related operational data associated with a flow using a separate header. Optical data queried may include Bit error rates (BER), Q-factor, pre-FEC counts, FEC

corrected words, and fiber latency as measured from the ingress optical interface of a source MLBD to the egress optical interface on a destination MLBD. This document defines the following list of IOAM OPT types.

IOAM OPT Type	IOAM OPT Name	Reference
0	FEC corrected bits	This Draft
1	FEC corrected bytes	This Draft
2	FEC uncorrectable words	This Draft
3	Unavailable seconds	This Draft
4	Pre-FEC errors	This Draft
5	Post-FEC errors	This Draft
6	Q-value	This Draft
7	ESNR	This Draft
8	Path Latency	This Draft
9	Optical SRLG	This Draft
10	Wavelength	This Draft
11	List of L0 Optical Node IDs	This Draft
12	List of L1 Optical Node IDs	This Draft
13-127	Reserved	TBD
128-255	Private	This Draft

Where:

FEC corrected bits : The number of bits that were corrected by the FEC deployed along the optical path/circuit.

FEC corrected bytes : The number of bytes that were corrected by the FEC deployed along the optical path/circuit.

FEC uncorrectable words : The number of words that were uncorrectable by the FEC deployed along the optical path/circuit.

Unavailable seconds : The number of seconds that the path was unavailable due to a loss of sync, error, or loss of signal.

Pre-FEC errors : The BER before FEC has been applied along the optical path/circuit

Post-FEC errors : The BER after FEC has been applied along the optical path/circuit

Q-value : The Quality value (in dB) the optical channel

measured along the path/circuit.

ESNR : The electric signal-to-noise ratio of the channel measured along the path/circuit.

Path Latency : The overall latency to forward data along an optical path/circuit.

Optical SRLG	:	The SRLG identifier associated with the optical path/circuit.
--------------	---	---

Wavelength : The wavelength(s) information associated with the optical path/circuit.

List of L0 Optical Node IDs: List of all L0 device identifiers along the optical path/circuit.

List of L1 Optical Node IDs: List of all L1 device identifiers along the optical path/circuit.

Note that types 128-255 are designated as private for carrying any vendor-specific optical attributes.

The structure for carrying the OPT TLV is given below:

IOAM optical data option header:

[illegible]

IOAM OPT Type: 8-bit identifier that identifies the specific optical data type that has been requested.

IOAM OPT Subtype: 8-bit identifier that identifies the specific

subtype of the queried optical data type. A subtype is intended to qualify the data collected with a description of statistic used (min/max/instant/average/aggregate/list) in the data collection process.

IOAM OPT Len: Length of data field (4-byte aligned).

IOAM OPT Data: Free flowing data field carrying the data value corresponding to the type of data requested.

The Opt data field can be used to convey optical path characteristics (e.g., wavelength, optical SRLG, and geography information).

6. Summary

The motivation for introducing a multi-layer IOAM is to integrate transport and packet domain OAM data. This allows for operational simplicity when it comes to gathering and correlating OAM data across an end-to-end path. This document describes the mechanism to achieve this IOAM integration and also some sample multi-layer IOAM data types and formats.

7. Security Considerations

This document does not introduce any new security considerations.

8 IANA Considerations

A future revision of this document will lay down the exact IANA request to create a new registry for "In-Situ OAM (IOAM) Optical Interface Type" and "In-Situ OAM (IOAM) Optical Data type". This is the common registry that will include registrations for all optical IOAM types.

9 References

9.1 Normative References

[I-D.ietf-ippm-ioam-data]

Brockners, F., Bhandari, S., Pignataro, C., Gredler, H.,
Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov,

P., Chang, R., D. Bernier, and Lemon, J. "Data Fields for In-situ OAM", [draft-ietf-ippm-ioam-data-03](#) (work in progress), June 2018.

[9.2](#) Informative 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>>.

Authors' Addresses

Madhukar Anand
Infinera Corporation
[169](#) W Java Dr, Sunnyvale, CA 94089

Email: manand@infinera.com

Sanjoy Bardhan
Infinera Corporation
[169](#) W Java Dr, Sunnyvale, CA 94089

Email: sbardhan@infinera.com

Radhakrishna Valiveti
Infinera Corporation
[169](#) W Java Dr, Sunnyvale, CA 94089

Email: rvaliveti@infinera.com

Ramesh Subrahmaniam
Email: svr_fremont@yahoo.com

Carlos Pignataro
Cisco Systems, Inc
Email: cpignata@cisco.com

Shwetha Bhandari
Cisco Systems, Inc
Email: shwethab@cisco.com

Randy Zhang
Cisco Systems, Inc
Email: ranzhang@cisco.com

Rajiv Asati
Cisco Systems, Inc
Email: rajiva@cisco.com

