IPPM
Internet-Draft
Intended status: Standards Track
Expires: 6 January 2026

T. Mizrahi
Huawei
F. Brockners
Cisco
A. Clemm
Sympotech
J. Iurman
University of Liege
S. Bhandari
Databricks
T. Zhou
Huawei
5 July 2025

In Situ Operations, Administration, and Maintenance (IOAM) Template Option
draft-mbci-ippm-ioam-template-option-00

Abstract

In situ measurement is performed by incorporating performance related information into in-flight data packets. This document specifies a new IOAM Option-Type that has a fixed length and can be updated by transit nodes along the path. It enables lightweight monitoring while maintaining a constant length that is not changed in-flight and is not affected by the number of hops in the network.

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 https://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 6 January 2026.

Table of Contents

1. Introduction

In Situ Operations, Administration, and Maintenance (IOAM) [RFC9197] is used for measuring and monitoring a network by incorporating measurement and operational data into some or all of the data packets. [RFC9197] has defined several Option-Types, intended for different purposes.

This document introduces a new IOAM Option-Type that can be incorporated into data packets and updated by transit nodes along the path. Compared to existing IOAM Trace Option-Types, the new Option-Type provides performance information using data fields that have a constant length.

There are several in-progress proposals that use a fixed-size telemetry header, including [I-D.cxx-ippm-ioamaggr], [I-D.mzbc-ippm-transit-measurement-option], [I-D.xiao-ippm-ioam-trace-extensions], [I-D.filsfils-ippm-path-tracing], [I-D.ravi-ippm-csig], and [I-D.shi-ippm-congestion-measurement-data]. These proposals can potentially benefit from the IOAM Option-Type that is presented in this document.

2. Conventions

2.1. Requirement 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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Terminology

Abbreviations used in this document:

IOAM: In-situ Operations, Administration, and Maintenance

OAM: Operations, Administration, and Maintenance

The terms Option-Type, encapsulating node, decapsulating node, and transit node are defined in [RFC9197].

3. Use Cases

There is a set of use cases that the current IOAM Option-Types do not support. This section lists a set of use cases for the IOAM Template Option-Type.

* Aggregated information: aggregated information across several nodes, e.g., [I-D.cxx-ippm-ioamaggr]. Many applications interested in telemetry data across a path are not focused on individual node’s telemetry, but on an aggregated metric that can provide a more holistic picture. Aggregating IOAM data along a network path meets this requirement. IOAM nodes do not only retrieve information, but also perform functions such as sum, average, minimum, or maximum of a given data parameter and carry the result to the next IOAM node in an IOAM data field.

* Congestion information: information relating to the congestion status can be collected from nodes along the path, e.g., [I-D.ravi-ippm-csig] providing a finer level of granularity than conventional ECN while limiting the congestion status to a fixed size.

* Combined information: a combination of aggregated information and congestion-related information can be collected along the path, e.g., [I-D.mzbc-ippm-transit-measurement-option], while using a fixed size.

4. Requirements for the IOAM Template Option-Type

This section lists requirements for the IOAM Template Option-Type:

* Templates supported MUST have a fixed length and fixed structure. Fixed length and structure are to simplify parsing of the Option-Type.

* Each templates is a composition of one or more data fields.

* Data fields within a template can have different sizes (e.g., 8 bits, 16 bits, 32 bits).

* Templates MUST align to 4 octet boundaries. If necessary, padding fields are used to guarantee 4-octet alignment.

* Data fields within a Template can be read-only for IOAM nodes or read-write for IOAM nodes, depending on their use.

* The IOAM Template Option-Type SHOULD support IETF defined (IANA registry) Templates for use-cases which are defined by the IETF as well as custom/deployment specific templates, that are defined by the operator and are specific to a deployment.

5. In situ Template Option-Type

This document defines a new IOAM Option-Type, the Template Option-Type. The length of the Template Option-Type MUST NOT be modified by IOAM transit nodes. However, IOAM transit nodes MAY modify the option data in the Template Option-Type. Figure 1 presents the format of this Option-Type.

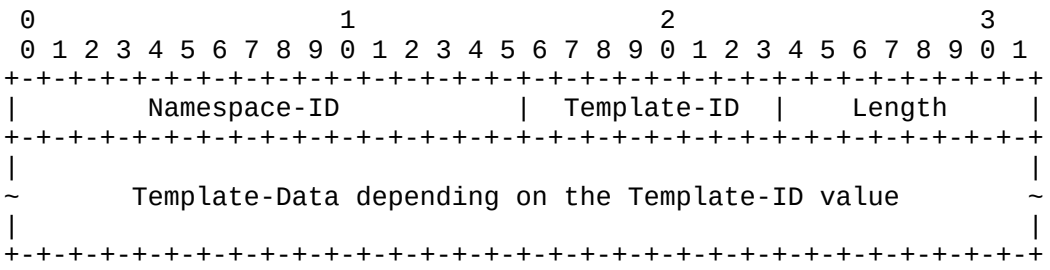

Figure 1: Template Option-Type

An IOAM node that complies to this draft MUST support the following fields, as depicted in Figure 1:

Namespace-ID: A 16-bit namespace identifier, as defined in [RFC9197].

Template-ID: An 8-bit identifier that specifies the template that follows. A new registry is defined for this field, as specified in Section 7. The value 0 has been assigned, indicating "No Option Data". Assignments of values 1 to 127 are controlled by IANA. Values 128 to 255 are can be defined by an operator for a specific deployment.

Length: An 8-bit length that specifies the size of the Template-Data in multiples of 4 octets.

Template-Data: The data that follows the Template-ID has a constant length. The semantics and length of the data are determined by the Template-ID. The option data might consist of more than one sub-field.

The specification of the Template-ID values and the corresponding option data formats is outside the scope of this document.

As in [RFC9197], the Template Option-Type can be incorporated into all or a subset of the traffic that is forwarded by the encapsulating node. Notably, this option adds a fixed and low overhead to data packets, which remains constant along the path.

6. Examples

The section lists examples of how the template option can be used.

6.1. IOAM Data Aggregation Along The Path

[I-D.cxx-ippm-ioamaggr] describes use cases to aggregate IOAM data along a network path. Rather than just collecting data at IOAM nodes, data is collected and processed by the IOAM nodes - using functions like sum, average, minimum, or maximum of a given data parameter - and the result stored in a data field called "Aggregate" (see below). The Template Option can be used to support this use-case with the following example template:

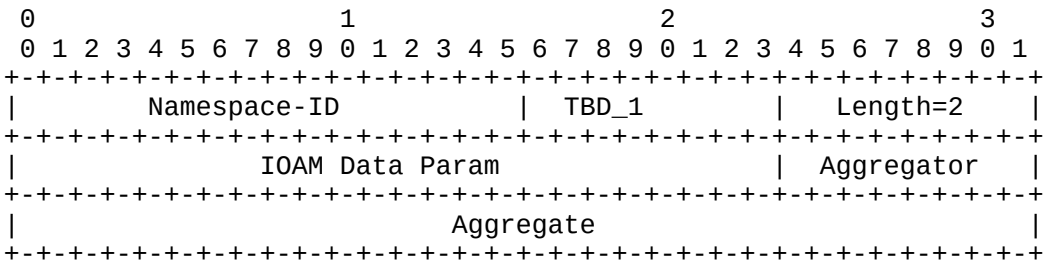

Figure 2: Aggregation Template

IOAM Data Param: This field identifies the data parameter that is to be aggregated across the nodes. It MUST be set by the IOAM encapsulating node. IOAM transit nodes MUST NOT change it.

Aggregator: This 8-bit field identifies the aggregation function that is to be applied. Its value MUST be set by the IOAM encapsulating node. IOAM transit nodes MUST not change it. The following aggregators are defined: Sum, Min, Max, Average.

Aggregate: This 32-bit field contains the aggregated value. Its value is initialized by the encapsulating node,in general by simply recording the value of its data parameter that is to be aggregated. The field is updated by each subsequent node pre the requested aggregation, including IOAM transit nodes as well as the IOAM decapsulating node (prior to performing decapsulation).

6.2. Transit Measurement Template

The use case that is presented in [I-D.mzbc-ippm-transit-measurement-option] provides aggregated transit delay information, as well as congestion status of transit nodes, as shown in the following template:

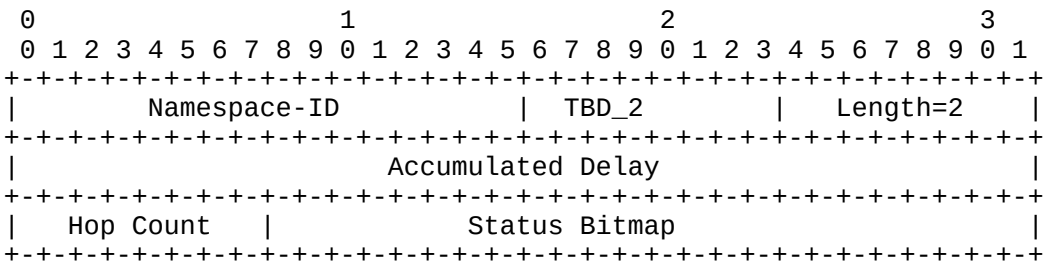

Figure 3: Transit Measurement Template

Accumulated Delay: represents the sum of the transit delay values in nanoseconds along the path of the packet, including the current node.

Hop Count/Status Bitmap: indicates the devices along the path that have experienced congestion. Hop Count is a one-octet field that indicates the number of hops since the encapsulating node, and is updated by each transit node. Status Bitmap is a three octet field that represents the congestion status of each transit node along the path. The value ’1’ indicates that the current packet was enqueued in a queue that is congested.

7. IANA Considerations

To be added to a future version of this document.

8. Security Considerations

The security considerations of IOAM in general are discussed in [RFC9197]. The Template Option-Type may be used for reconnaissance, which in turn can facilitate other types of attacks. As in other types of IOAM data fields, a malicious attacker can manipulate the field values in order to create a false illusion of nonexistent network issues or prevent the detection of actual ones.

9. References

9.1. 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, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC9197] Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <https://www.rfc-editor.org/info/rfc9197>.

9.2. Informative References

[I-D.cxx-ippm-ioamaggr]
Clemm, A., Metzger, Bister, R., and S. Dellsperger, "Aggregation Trace Option for In-situ Operations, Administration, and Maintenance (IOAM)", Work in Progress, Internet-Draft, draft-cxx-ippm-ioamaggr-03, 23 April 2025, <https://datatracker.ietf.org/doc/html/draft-cxx-ippm-ioamaggr-03>.

[I-D.filsfils-ippm-path-tracing]
Filsfils, C., Abdelsalam, A., Camarillo, P., Yufit, M., Graf, T., Su, Y., Matsushima, S., Valentine, M., and Dhamija, "Path Tracing in SRv6 networks", Work in Progress, Internet-Draft, draft-filsfils-ippm-path-tracing-03, 18 May 2025, <https://datatracker.ietf.org/doc/html/draft-filsfils-ippm-path-tracing-03>.

[I-D.mzbc-ippm-transit-measurement-option]
Mizrahi, T., Zhou, T., Belkar, S., and R. Cohen, "The Transit Measurement Option", Work in Progress, Internet-Draft, draft-mzbc-ippm-transit-measurement-option-06, 2 July 2025, <https://datatracker.ietf.org/doc/html/draft-mzbc-ippm-transit-measurement-option-06>.

[I-D.ravi-ippm-csig]
Ravi, A., Dukkipati, N., Mehta, N., and J. Kumar, "Congestion Signaling (CSIG)", Work in Progress, Internet-Draft, draft-ravi-ippm-csig-01, 2 February 2024, <https://datatracker.ietf.org/doc/html/draft-ravi-ippm-csig-01>.

[I-D.shi-ippm-congestion-measurement-data]
Fioccola, G., Zhou, T., Zhao, G., and Z. Li, "Data Fields for Congestion Measurement", Work in Progress, Internet-Draft, draft-shi-ippm-congestion-measurement-data-04, 20 June 2025, <https://datatracker.ietf.org/doc/html/draft-shi-ippm-congestion-measurement-data-04>.

[I-D.xiao-ippm-ioam-trace-extensions]
Min, X., Liu, Y., and C. Lin, "Extensions to IOAM Trace Option for Carrying Fixed-Size Data", Work in Progress, Internet-Draft, draft-xiao-ippm-ioam-trace-extensions-01, 18 April 2025, <https://datatracker.ietf.org/doc/html/draft-xiao-ippm-ioam-trace-extensions-01>.

Authors' Addresses

Tal Mizrahi
Huawei
Matam
Haifa 3190501
Israel
Email: tal.mizrahi.phd@gmail.com

Frank Brockners
Cisco Systems, Inc.
Hansaallee 249, 3rd Floor
40549 DUESSELDORF
Germany
Email: fbrockne@cisco.com

Alexander Clemm
Sympotech
Email: ludwig@clemm.org

Justin Iurman
University of Liege
10, Allee de la decouverte (B28)
4000 Sart-Tilman
Belgium
Email: justin.iurman@uliege.be

Shwetha Bhandari
Databricks
Angkor West Building Bagmane Capital Tech Park Ferns City Doddanekkundi
Mahadevpura Bengaluru, Karnataka 560048
India
Email: shwetha.bhandari@databricks.com

Tianran Zhou
Huawei
156 Beiqing Rd.
Beijing
100095
China
Email: zhoutianran@huawei.com