

ippm  
Internet-Draft  
Intended status: Standards Track  
Expires: August 25, 2022

B. Weis, Ed.  
Independent  
F. Brockners, Ed.  
C. Hill  
Cisco  
S. Bhandari  
Thoughtspot  
V. Govindan  
C. Pignataro, Ed.  
N. Nainar, Ed.  
Cisco  
H. Gredler  
RtBrick Inc.  
J. Leddy

S. Youell  
JMPC

T. Mizrahi  
Huawei Network.IO Innovation Lab

A. Kfir  
B. Gafni  
Nvidia

P. Lapukhov  
Facebook

M. Spiegel  
Barefoot Networks, an Intel company  
February 21, 2022

EtherType Protocol Identification of In-situ OAM Data  
draft-weis-ippm-ioam-eth-05

Abstract

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path between two points in the network. This document defines an EtherType that identifies IOAM data fields as being the next protocol in a packet, and a header that encapsulates the IOAM data fields.

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 August 25, 2022.

#### Copyright Notice

Copyright (c) 2022 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 (<https://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. Conventions . . . . .	3
2.1. Requirements Language . . . . .	3
2.2. Abbreviations . . . . .	3
3. IOAM EtherType . . . . .	3
4. Usage Examples of the IOAM EtherType . . . . .	4
4.1. Example: GRE Encapsulation of IOAM Data Fields . . . . .	5
4.2. Example: Geneve Encapsulation of IOAM Data Fields . . . . .	6
5. Security Considerations . . . . .	7
6. IANA Considerations . . . . .	8
7. Acknowledgements . . . . .	8
8. References . . . . .	8
8.1. Normative References . . . . .	8
8.2. Informative References . . . . .	9
Authors' Addresses . . . . .	9

#### 1. Introduction

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a particular network domain. The term "in-situ" refers to

the fact that the IOAM data fields are added to the data packets rather than being sent within packets specifically dedicated to OAM. This document proposes a new EtherType for IOAM and defines how IOAM data fields are carried as part of encapsulations where the IOAM data fields follows an encapsulation header that uses an EtherType to denote the type of protocol data unit. Examples of these protocols are GRE [RFC2784] [RFC2890] and Geneve [RFC8926]). This document outlines how IOAM data fields are encoded in these encapsulation headers.

## 2. Conventions

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

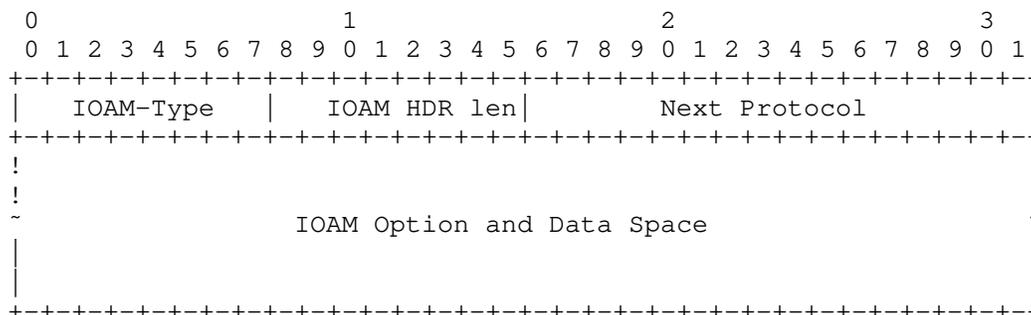
### 2.2. Abbreviations

Abbreviations used in this document:

E2E:           Edge-to-Edge  
Geneve:        Generic Network Virtualization Encapsulation  
GRE:           Generic Routing Encapsulation  
IOAM:          In-situ Operations, Administration, and Maintenance  
OAM:           Operations, Administration, and Maintenance  
POT:           Proof of Transit

## 3. IOAM EtherType

When the IOAM data fields are included within an encapsulation that identifies the next protocol using an EtherType (e.g., GRE or Geneve) the presence of IOAM data fields are identified with TBD\_IOAM. When this EtherType is used, an additional IOAM header is also included. This header indicates the type of IOAM data fields that follows, and the next protocol that follows the IOAM data fields.



The IOAM encapsulation is defined as follows.

**IOAM Type:** 8-bit field defining the IOAM Option type, as defined in Section 5.1 of [I-D.ietf-ippm-ioam-data].

**IOAM HDR Len:** 8 bit Length field contains the length of the IOAM header in 4-octet units.

**Next Protocol:** 16 bits Next Protocol Type field contains the protocol type of the protocol data unit following IOAM protocol header. Protocol Type is defined to be an EtherType value from [ETYPES]. An implementation receiving a packet containing a Protocol Type which is not listed in one of those registries SHOULD discard the packet.

**IOAM Option and Data Space:** IOAM option header and data is present as specified by the IOAM-Option-Type field, and is defined in Section 5 of [I-D.ietf-ippm-ioam-data].

Multiple IOAM options MAY be included within the encapsulation header. For example, if a GRE encapsulation contains two IOAM options before the data payload, the Next Protocol field of the first IOAM option will contain the value of TBD\_IOAM, while the Next Protocol field of the second IOAM option will contain the EtherType indicating the type of the data payload.

#### 4. Usage Examples of the IOAM EtherType

The IOAM EtherType can be used with any encapsulation that uses EtherType to denote the type of the protocol data unit. The following sections show how it can be used when GRE and Geneve are used as the encapsulation header.

4.1. Example: GRE Encapsulation of IOAM Data Fields

When IOAM data fields are carried in GRE, the IOAM encapsulation defined above follows the GRE header, as shown in Figure 1.

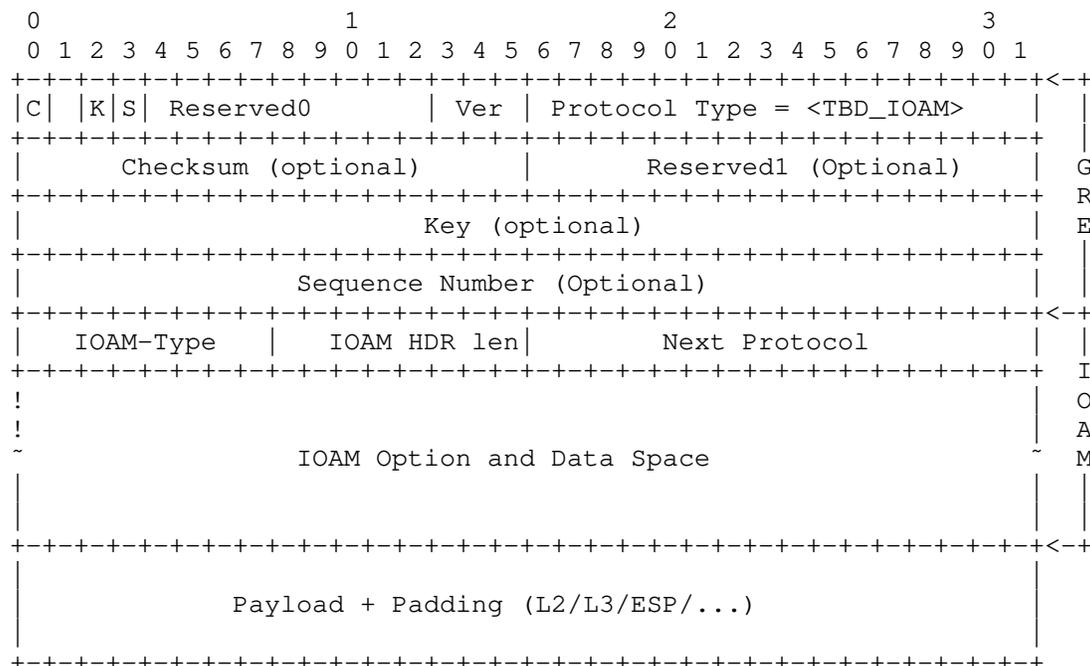


Figure 1: GRE Encapsulation Example

The GRE header and fields are defined in [RFC2890]. The GRE Protocol Type value is set to TBD\_IOAM.

Figure 2 shows two example protocol header stacks that use GRE along with IOAM. IOAM Option-Types (the below diagram uses "IOAM" as shorthand for IOAM Option-Types) are sequenced in behind the GRE header that follows the "outer" header of the next protocol unit.

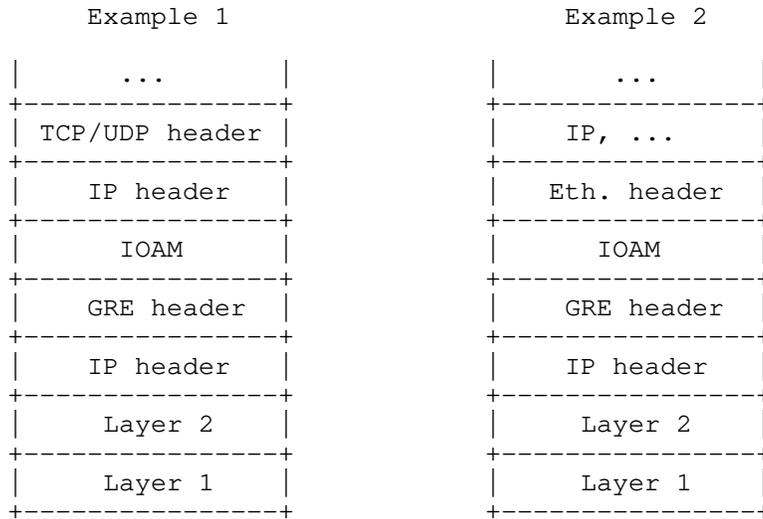


Figure 2: GRE with IOAM examples

4.2. Example: Geneve Encapsulation of IOAM Data Fields

When IOAM data fields are carried in Geneve, the IOAM encapsulation defined above follows the Geneve header, as shown in Figure 3.

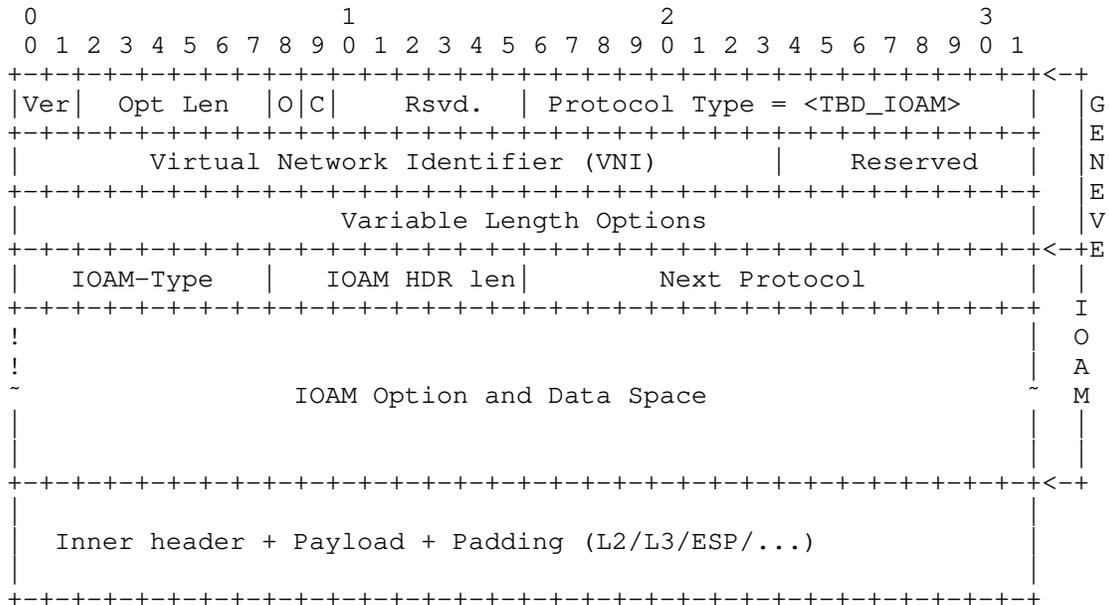


Figure 3: Geneve Encapsulation Example

The Geneve header and fields are defined in [RFC8926]. The Geneve Protocol Type value is TBD\_IOAM.

5. Security Considerations

This document describes the encapsulation of IOAM data fields in the encapsulation header such as GRE and Geneve that uses EtherType to denote the protocol data unit. Security considerations of the specific IOAM data fields for each case (i.e., Trace, Proof of Transit, and E2E) are described in [I-D.ietf-ippm-ioam-data].

As this document describes new protocol fields within the existing encapsulation, any security considerations of the respective encapsulation header is applicable. When the encapsulation is GRE, the security considerations of [RFC2890] is applicable. When the encapsulation is Geneve, the security considerations of [RFC8926] is applicable.

IOAM data fields SHOULD be integrity protected (e.g., with [I-D.ietf-ippm-ioam-data-integrity]) to detect changes made by a device between the IOAM encapsulating node and the IOAM decapsulating node.

## 6. IANA Considerations

A new EtherType value is requested to be added to the [ETYPES] IANA registry by IEEE Registration Authority. The description should be "In-situ OAM (IOAM)".

## 7. Acknowledgements

We would like to thank Nagendra Kumar Nainar for the contribution.

## 8. References

### 8.1. Normative References

- [ETYPES] "IANA Ethernet Numbers",  
<<https://www.iana.org/assignments/ieee-802-numbers/ieee-802-numbers.xhtml>>.
- [I-D.ietf-ippm-ioam-data]  
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-17 (work in progress), December 2021.
- [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>>.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", RFC 2890, DOI 10.17487/RFC2890, September 2000, <<https://www.rfc-editor.org/info/rfc2890>>.
- [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>>.
- [RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <<https://www.rfc-editor.org/info/rfc8926>>.

## 8.2. Informative References

[I-D.ietf-ippm-ioam-data-integrity]  
Brockners, F., Bhandari, S., and T. Mizrahi, "Integrity of  
In-situ OAM Data Fields", draft-ietf-ippm-ioam-data-  
integrity-00 (work in progress), October 2021.

## Authors' Addresses

Brian Weis (editor)  
Independent  
USA

Email: bew.stds@gmail.com

Frank Brockners (editor)  
Cisco Systems, Inc.  
Hansaallee 249, 3rd Floor  
DUESSELDORF, NORDRHEIN-WESTFALEN 40549  
Germany

Email: fbrockne@cisco.com

Craig Hill  
Cisco Systems, Inc.  
13600 Dulles Technology Drive  
Herndon, Virginia 20171  
United States

Email: crhill@cisco.com

Shwetha Bhandari  
Thoughtspot  
3rd Floor, Indiqube Orion, 24th Main Rd, Garden Layout, HSR Layout  
Bangalore, KARNATAKA 560 102  
India

Email: shwetha.bhandari@thoughtspot.com

Vengada Prasad Govindan  
Cisco Systems, Inc.

Email: venggovi@cisco.com

Carlos Pignataro (editor)  
Cisco Systems, Inc.  
7200-11 Kit Creek Road  
Research Triangle Park, NC 27709  
United States

Email: cpignata@cisco.com

Nagendra Kumar Nainar (editor)  
Cisco Systems, Inc.  
7200-11 Kit Creek Road  
Research Triangle Park, NC 27709  
United States

Email: naikumar@cisco.com

Hannes Gredler  
RtBrick Inc.

Email: hannes@rtbrick.com

John Leddy  
United States

Email: john@leddy.net

Stephen Youell  
JP Morgan Chase  
25 Bank Street  
London E14 5JP  
United Kingdom

Email: stephen.youell@jpmorgan.com

Tal Mizrahi  
Huawei Network.IO Innovation Lab  
Israel

Email: tal.mizrahi.phd@gmail.com

Aviv Kfir  
Nvidia  
350 Oakmead Parkway, Suite 100  
Sunnyvale, CA 94085  
U.S.A.

Email: avivk@nvidia.com

Barak Gafni  
Nvidia  
350 Oakmead Parkway, Suite 100  
Sunnyvale, CA 94085  
U.S.A.

Email: gbarak@nvidia.com

Petr Lapukhov  
Facebook  
1 Hacker Way  
Menlo Park, CA 94025  
US

Email: petr@fb.com

Mickey Spiegel  
Barefoot Networks, an Intel company  
4750 Patrick Henry Drive  
Santa Clara, CA 95054  
US

Email: mickey.spiegel@intel.com