

OPSAWG
Internet-Draft
Intended status: Standards Track
Expires: 23 January 2025

M. Boucadair
Orange
T. Reddy.K
Nokia
22 July 2024

Export of UDP Options Information in IP Flow Information Export (IPFIX)
draft-ietf-opsawg-tsvwg-udp-ipfix-14

Abstract

This document specifies new IP Flow Information Export (IPFIX) Information Elements for UDP options.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Operations and Management Area Working Group Working Group mailing list (opsawg@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/opsawg/>.

Source for this draft and an issue tracker can be found at <https://github.com/boucadair/udp-ipfix>.

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 23 January 2025.

Copyright Notice

Copyright (c) 2024 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions and Definitions	3
3. UDP Options at a Glance	3
4. New UDP IPFIX Information Elements	5
4.1. udpSafeOptions	5
4.2. udpUnsafeOptions	6
4.3. udpExID	7
4.4. udpSafeExIDList	7
4.5. udpUnsafeExIDList	8
5. Examples	8
5.1. Reduced-size Encoding	8
5.2. SAFE Experimental Option	9
5.3. ExIDs and Reduced-size Encoding	9
6. Security Considerations	10
7. IANA Considerations	10
7.1. IPFIX Information Elements	10
8. References	11
8.1. Normative References	11
8.2. Informative References	12
Acknowledgments	13
Authors' Addresses	13

1. Introduction

IP Flow Information Export (IPFIX) [RFC7011] is a protocol that is widely deployed in networks for traffic management purposes (Section 2 of [RFC6632]). The protocol specifies the encoding of a set of basic data types and how the various Information Elements (IEs) are transmitted. In order to support the export of new flow-related measurement data, new IEs can be defined and registered in a dedicated IANA registry [IANA-IPFIX] for interoperability.

This document specifies new IPFIX Information Elements for UDP options (Section 4). A brief overview of UDP options is provided in Section 3.

The IE specified in Section 4.1 uses the new abstract data type ("unsigned256") defined in [I-D.ietf-opsawg-ipfix-tcpo-v6eh].

Transport (including MTU) considerations are discussed in Section 10 of [RFC7011].

Examples to illustrate the use of the new IPFIX Information Elements are provided in Section 5.

2. Conventions and Definitions

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.

This document uses the IPFIX-specific terminology (e.g., Flow) defined in Section 2 of [RFC7011]. As in the base IPFIX specification [RFC7011], these IPFIX-specific terms have the first letter of a word capitalized.

The document adheres to the naming conventions for Information Elements per Section 2.3 of [RFC7012].

Also, this document uses the terms defined in Section 3 of [I-D.ietf-tsvwg-udp-options], especially "datagram" and "surplus area".

3. UDP Options at a Glance

UDP [RFC0768] does not support an extension mechanism similar to the options supported by other transport protocols, such as TCP [RFC9293], SCTP [RFC9260], or DCCP [RFC4340]. Such a mechanism can be useful for various applications, e.g., to discover a path MTU or share timestamps. To fill that void, [I-D.ietf-tsvwg-udp-options] extends UDP with a mechanism to insert extensions in datagrams. To do so, and unlike the conventional approach that relies upon transport headers, [I-D.ietf-tsvwg-udp-options] uses trailers. Concretely, UDP options are placed in the surplus area (that is, the area of an IP payload that follows a UDP packet). See Figure 1. An example of the use of UDP options for Datagram Packetization Layer Path Maximum Transmission Unit Discovery (DPLPMTUD) is described in [I-D.ietf-tsvwg-udp-options-dplpmtud].




Figure 1: Surplus Area

Sections 4.1 and 4.2 introduce new IEs to export the observed UDP options.

UDP options are unambiguously identified by means of a 1-byte field, called "Kind".

Options indicated by Kind values in the range 0-191 are called SAFE options. Such options can be silently ignored by legacy receivers because they do not alter the UDP user data (Section 11 of [I-D.ietf-tsvwg-udp-options]). SAFE options are exported using the IE defined in Section 4.1.

Options indicated by Kind values in the range 192-255 are called UNSAFE options. Such options are not safe for legacy receivers to ignore because they alter the UDP user data (Section 12 of [I-D.ietf-tsvwg-udp-options]). UNSAFE options are exported using the IE defined in Section 4.2.

UDP options occur per-packet within a Flow and can be inserted at any time in the Flow.

[I-D.ietf-tsvwg-udp-options] reserves two options for experiments: the Experimental option (EXP, Kind=127) for SAFE options and the UNSAFE Experimental option (UEXP, Kind=254). For both options, Experiment Identifiers (ExIDs) are used to differentiate concurrent use of these options. Known ExIDs are expected to be registered within IANA. Section 4.4 specifies a new IPFIX IE to export observed ExIDs in the EXP options. Also, Section 4.5 specifies a new IPFIX IE to export observed ExIDs in the UEXP options. Only 16-bit ExIDs are supported in [I-D.ietf-tsvwg-udp-options].

This document does not intend to elaborate operational guidance/implications of UDP options. The document focuses exclusively on exporting observed UDP options in datagrams.

4. New UDP IPFIX Information Elements

RFC Editor Note: Please update "URL_IANA_UDP_OPTIONS" reference with the URL of the "UDP Option Kind Numbers" registry group and "URL_IANA_UDP_ExIDs" with the URL of the "UDP Experimental Option Experiment Identifiers (UDP ExIDs)" registry that will be created by IANA as per Section 25 of [I-D.ietf-tsvwg-udp-options].

Given the Kind structure of SAFE and UNSAFE UDP options, using one single IE that would multiplex both types of option will limit the benefits of reduced-size encoding in the presence of UNSAFE options. For example, at least 24 octets would be needed to report mandatory SAFE options that are observed in a Flow. In order to use less bits to report observed UDP options, distinct IEs are thus defined to report SAFE (Section 4.1) and UNSAFE (Section 4.2) UDP options. As further detailed in Section 5.1, only one octet is needed to report mandatory SAFE options.

4.1. udpSafeOptions

Name: udpSafeOptions

ElementID: TBD1

Description: Observed SAFE UDP options in a Flow. The information is encoded in a set of bit fields.

Options are mapped to bits according to their option numbers. UDP option Kind 0 corresponds to the least-significant bit in the udpSafeOptions IE while Kind 191 corresponds to the 65th most-significant bit of the IE. The bit is set to 1 if the corresponding SAFE UDP option is observed at least once in the Flow. The bit is set to 0 if the option is never observed in the Flow. The 64 most-significant bits MUST be set to 0.

The reduced-size encoding per Section 6.2 of [RFC7011] is followed whenever fewer octets are needed to report observed SAFE UDP options. For example, if only option Kinds ≤ 31 are observed, then the value of the udpSafeOptions IE can be encoded as unsigned32, or if only option Kinds ≤ 63 are observed, then the value of the udpSafeOptions IE can be encoded as unsigned64.

The presence of udpSafeExIDList is an indication that the SAFE

Experimental option is observed in a Flow. The presence of `udpSafeExIDList` takes precedence over setting the corresponding bit in the `udpSafeOptions` IE for the same Flow. In order to optimize the use of the reduced-size encoding in the presence of `udpSafeExIDList` IE, the Exporter MUST NOT set to 1 the EXP flag of the `udpSafeOptions` IE that is reported for the same Flow.

Abstract Data Type: `unsigned256`

Data Type Semantics: `flags`

Additional Information: See the "UDP Option Kind Numbers" registry at `[URL_IANA_UDP_OPTIONS]`.

See `[I-D.ietf-tsvwg-udp-options]` for more details about UDP options.

Reference: This-Document

4.2. `udpUnsafeOptions`

Name: `udpUnsafeOptions`

ElementID: TBD2

Description: Observed UNSAFE UDP options in a Flow. The information is encoded in a set of bit fields.

Options are mapped to bits according to their option numbers. UDP option Kind 192 corresponds to the least-significant bit in the `udpUnsafeOptions` IE while Kind 255 corresponds to the most-significant bit of the IE. The bit is set to 1 if the corresponding UNSAFE UDP option is observed at least once in the Flow. The bit is set to 0 if the option is never observed in the Flow.

The reduced-size encoding per Section 6.2 of `[RFC7011]` is followed whenever fewer octets are needed to report observed UNSAFE UDP options.

The presence of `udpUnsafeExIDList` is an indication that the UNSAFE Experimental option is observed in a Flow. The presence of `udpUnsafeExIDList` takes precedence over setting the corresponding bit in the `udpUnsafeOptions` IE for the same Flow. In order to optimize the use of the reduced-size encoding in the presence of `udpUnsafeExIDList` IE, the Exporter MUST NOT set to 1 the UEXP flag of the `udpUnsafeOptions` IE that is reported for the same Flow.

Abstract Data Type: unsigned64

Data Type Semantics: flags

Additional Information: See the "UDP Option Kind Numbers" registry at [URL_IANA_UDP_OPTIONS].

See [I-D.ietf-tsvwg-udp-options] for more details about UDP options.

Reference: This-Document

4.3. udpExID

Name: udpExID

ElementID: TBD3

Description: Observed ExID in an Experimental option (EXP, Kind=127) or an UNSAFE Experimental option (UEXP, Kind=254).

A basicList of udpExID is used to report udpSafeExIDList and udpUnsafeExIDList values.

Abstract Data Type: unsigned16

Data Type Semantics: identifier

Additional Information: See the "UDP Experimental Option Experiment Identifiers (UDP ExIDs)" registry at [URL_IANA_UDP_ExIDs].

See [I-D.ietf-tsvwg-udp-options] for more details about ExIDs.

Reference: This-Document

4.4. udpSafeExIDList

Name: udpSafeExIDList

ElementID: TBD4

Description: Observed ExIDs in the Experimental option (EXP, Kind=127).

A basicList of udpExID Information Elements in which each udpExID Information Element carries the ExID observed in an EXP option.

Abstract Data Type: basicList

Data Type Semantics: list

Additional Information: See the "UDP Experimental Option Experiment Identifiers (UDP ExIDs)" registry at [URL_IANA_UDP_ExIDs].

See [I-D.ietf-tsvwg-udp-options] for more details about ExIDs.

Reference: This-Document

4.5. udpUnsafeExIDList

Name: udpUnsafeExIDList

ElementID: TBD5

Description: Observed ExIDs in the UNSAFE Experimental option (UEXP, Kind=254).

A basicList of udpExID Information Elements in which each udpExID Information Element carries the ExID observed in an UEXP option.

Abstract Data Type: basicList

Data Type Semantics: list

Additional Information: See the "UDP Experimental Option Experiment Identifiers (UDP ExIDs)" registry at [URL_IANA_UDP_ExIDs].

See [I-D.ietf-tsvwg-udp-options] for more details about ExIDs.

Reference: This-Document

5. Examples

5.1. Reduced-size Encoding

Given the UDP Kind allocation in Section 10 of [I-D.ietf-tsvwg-udp-options] and the option mapping defined in Section 4.1 of this document, fewer octets are likely to be used for Flows with mandatory UDP options.

Figure 2 shows an example of the Kind/bit mappings in the udpSafeOptions IE for a Flow in which End of Options List (EOL, Kind=0) and Alternate payload checksum (APC, Kind=2) options are observed. Only the bits that corresponds to EOL and APC options are set to 1.


```

MSB                                     LSB
      1                               25
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 ... 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 2: An Example of udpSafeOptions IE with EOL and APC Options

One octet is sufficient to report these observed options because the leading zeros are dropped per the reduced-size encoding guidance. Concretely, the reported udpSafeOptions IE will be set to 0x05 (Figure 3).

```

MSB                                     LSB
      0 1 2 3 4 5 6 7
+---+---+---+---+---+---+---+---+
|0|0|0|0|0|0|1|0|1|
+---+---+---+---+---+---+---+---+

```

Figure 3: An Example of the Wire udpSafeOptions IE Value with EOL and APC Options

5.2. SAFE Experimental Option

Let us now consider a UDP Flow in which SAFE Experimental options are observed. If a udpSafeOptions IE is exported for this Flow, then that IE will have the EXP bit set to 1 (Figure 4). This example does not make any assumption about the presence of other UDP options ("X" can be set to 0 or 1).

```

MSB                                     LSB
      12                               25
0 1 2 3 ... 7 8 9 0 1 2 3 4 5 6 7 8 9 ... 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|X|X|X|X|    |X|X|X|X|X|X|X|X|X|X|X|X|1|X|X|    |X|X|X|X|X|X|X|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 4: An Example of udpSafeOptions with EXP Option

5.3. ExIDs and Reduced-size Encoding

Now assume that EOL, APC, EXP, and UEXP options are observed in a Flow. Let us also consider that the observed SAFE Experimental options have ExIDs set to 0x9858 and 0xE2D4, and UNSAFE Experimental options have ExIDs set to 0xC3D9 and 0x1234. Figure 5 shows an excerpt of the Data Set encoding with a focus on SAFE Experimental options have ExIDs. The meaning of the fields is defined in

[RFC6313].

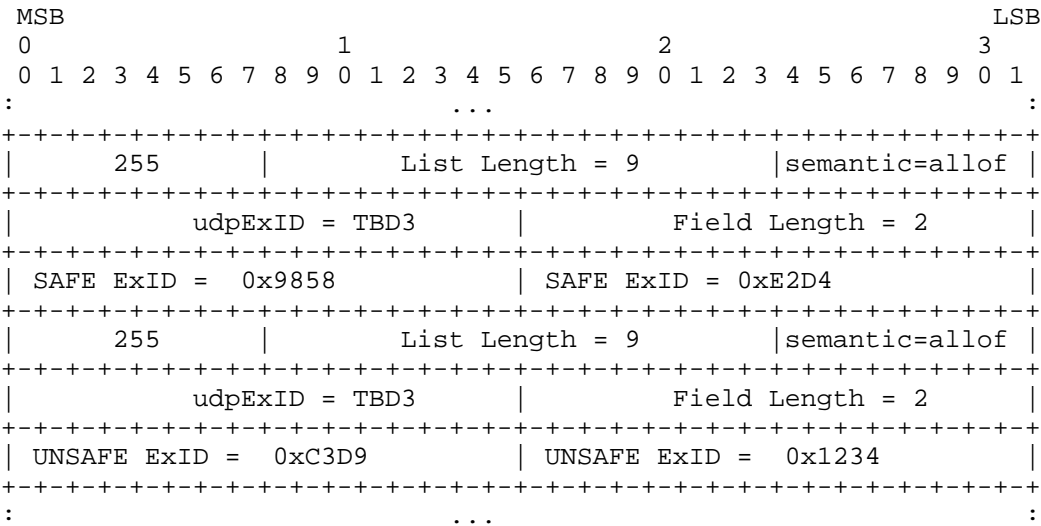


Figure 5: Example of UDP Experimental Option ExID IEs

Following the guidance in Section 4.1, the reported udpSafeOptions IE will be set to 0x05 even in the presence of EXP options.

6. Security Considerations

This document does not introduce new security considerations other than those already discussed in Section 11 of [RFC7011] and Section 8 of [RFC7012].

The reader may refer to Section 24 of [I-D.ietf-tsvwg-udp-options] for the security considerations related to UDP options.

7. IANA Considerations

7.1. IPFIX Information Elements

This document requests IANA to add the following new IEs to the "IPFIX Information Elements" registry under the "IP Flow Information Export (IPFIX) Entities" registry group [IANA-IPFIX]:

ElementID	Name	Specification
TBD1	udpSafeOptions	Section 4.1 of This-Document
TBD2	udpUnsafeOptions	Section 4.2 of This-Document
TBD3	udpExID	Section 4.3 of This-Document
TBD4	udpSafeExIDList	Section 4.4 of This-Document
TBD5	udpUnsafeExIDList	Section 4.5 of This-Document

Table 1: New IPFIX Information Elements

udpSafeOptions uses the abstract data type ("unsigned256") defined in [I-D.ietf-opsawg-ipfix-tcpo-v6eh].

Note to IANA: The "Specification" column points to the sections with the required information to register each IE.

Note to the RFC Editor: Please remove the IANA note once IANA actions are implemented.

8. References

8.1. Normative References

[I-D.ietf-opsawg-ipfix-tcpo-v6eh]

Boucadair, M. and B. Claise, "Extended TCP Options and IPv6 Extension Headers IPFIX Information Elements", Work in Progress, Internet-Draft, draft-ietf-opsawg-ipfix-tcpo-v6eh-17, 5 July 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-ipfix-tcpo-v6eh-17>>.

[I-D.ietf-tsvwg-udp-options]

Touch, J. D., "Transport Options for UDP", Work in Progress, Internet-Draft, draft-ietf-tsvwg-udp-options-32, 21 March 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-udp-options-32>>.

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/rfc/rfc768>>.

- [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/rfc/rfc2119>>.
- [RFC6313] Claise, B., Dhandapani, G., Aitken, P., and S. Yates, "Export of Structured Data in IP Flow Information Export (IPFIX)", RFC 6313, DOI 10.17487/RFC6313, July 2011, <<https://www.rfc-editor.org/rfc/rfc6313>>.
- [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, <<https://www.rfc-editor.org/rfc/rfc7011>>.
- [RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <<https://www.rfc-editor.org/rfc/rfc7012>>.
- [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/rfc/rfc8174>>.

8.2. Informative References

- [I-D.ietf-tsvwg-udp-options-dplpmtud] Fairhurst, G. and T. Jones, "Datagram PLPMTUD for UDP Options", Work in Progress, Internet-Draft, draft-ietf-tsvwg-udp-options-dplpmtud-12, 7 May 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-udp-options-dplpmtud-12>>.
- [IANA-IPFIX] "IP Flow Information Export (IPFIX) Entities", n.d., <<https://www.iana.org/assignments/ipfix/ipfix.xhtml>>.
- [RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, DOI 10.17487/RFC4340, March 2006, <<https://www.rfc-editor.org/rfc/rfc4340>>.
- [RFC6632] Ersue, M., Ed. and B. Claise, "An Overview of the IETF Network Management Standards", RFC 6632, DOI 10.17487/RFC6632, June 2012, <<https://www.rfc-editor.org/rfc/rfc6632>>.

- [RFC9260] Stewart, R., Tテシxen, M., and K. Nielsen, "Stream Control Transmission Protocol", RFC 9260, DOI 10.17487/RFC9260, June 2022, <<https://www.rfc-editor.org/rfc/rfc9260>>.
- [RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <<https://www.rfc-editor.org/rfc/rfc9293>>.
- [URL_IANA_UDP_ExIDs]
"UDP Experimental Option Experiment Identifiers (UDP ExIDs)", n.d., <<https://www.iana.org/assignments/url2>>.
- [URL_IANA_UDP_OPTIONS]
"UDP Option Kind Numbers", n.d., <<https://www.iana.org/assignments/url1>>.

Acknowledgments

Thanks to Benoテヨt Claise for the discussion on the ordering of IPFIX IEs. Thanks to Paul Aitken for the review and comments.

Thanks to Tommy Pauly for the tsvar review, Joe Touch for the intdir review, Robert Sparks for the genart review, Watson Ladd for the secdir review, and Jouni Korhonen for the opsdireview.

Thanks to Thomas Graf for the Shepherd review.

Thanks to Mahesh Jethanandani for the AD review.

Thanks to テ詠ic Vyncke, Roman Danyliw, and Zahed Sarker for the IESG review.

Authors' Addresses

Mohamed Boucadair
Orange
35000 Rennes
France
Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy.K
Nokia
India
Email: kondtir@gmail.com