IPFIX IPv6/TCP/UDP I-Ds Set: Updates & Next Steps

IETF#118, Prague
November 2023

M. Boucadair (Orange), B. Claise (Huawei), T. Reddy (Nokia)

IETF#116 "What's Next?" Plan

- Request WGLC for draft-ietf-opsawg-rfc7125-update

- Request WG adoption for the following I-D set:
  - *draft-boucla-opsawg-ipfix-fixes*
  - *draft-boucadair-opsawg-ipfix-tcpo-v6eh*
  - *draft-boucadair-opsawg-tsvwg-udp-ipfix*

- The last two documents may be merged, but we prefer to keep them separate because of the dependency on the UDP Options spec (*tsvwg*)

Update Since IETF#116

- draft-ietf-opsawg-rfc7125-update
  - Status: Passed the WGLC and IETF LC
- Adoption of the following I-Ds set:
  - *draft-ietf-opsawg-ipfix-fixes-03*
  - *draft-ietf-opsawg-ipfix-tcpo-v6eh-05*
  - *draft-ietf-opsawg-tsvwg-udp-ipfix-03*
- Seeked for cross-WGs reviews:
  - Sent messages to the following WGs
    - 6man: draft-ietf-opsawg-ipfix-tcpo-v6eh
    - tsvwg: draft-ietf-opsawg-tsvwg-udp-ipfix
  - …but no follow-up unfortunately
- However, we received good reviews from IPFIX IE Doctors and Éric Vyncke

Focus on IPv6 EH IEs (1)

- 4 IEs
- Specify how to report
  - Multiple EH chains in a Flow
  - Length of EH chains
  - Occurrences and order of Ehs
  - Non-consecutive EHs; not aggregates
  - Whether reported EHs are constrained by a HW/SW limit
  - Optimize the encoding
- Specify the dependency between the various IEs
- Add Examples

Focus on IPv6 EH IEs (2)

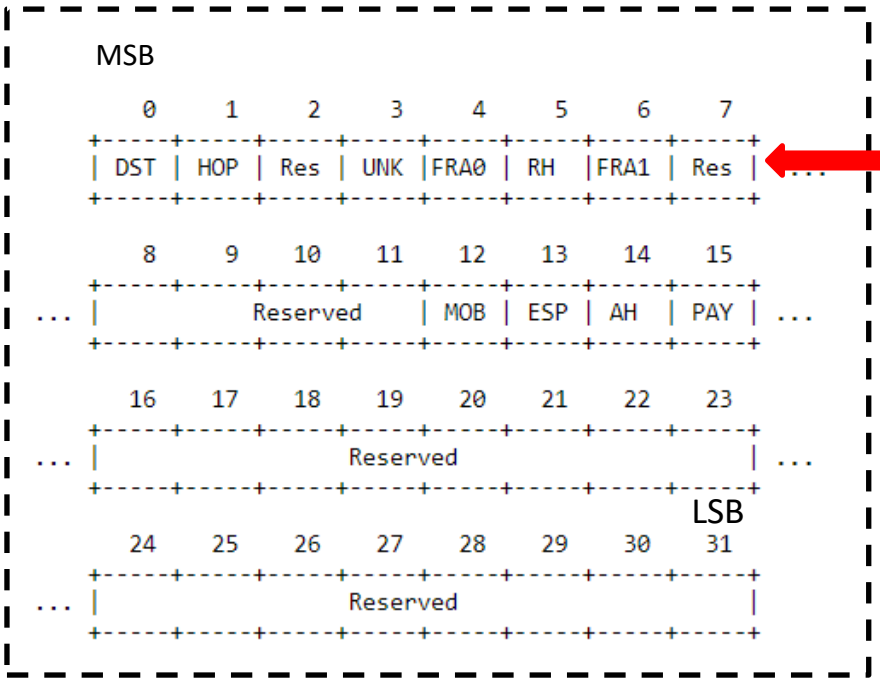

ipv6ExtensionHeaders (Existing IE)

Reduced-encoding (rfc7011#section-6.2) is unlikely because these EHs are the ones that are likely to be observed

Fix

draft-ietf-opsawg-ipfix-tcpo-v6eh

"Bit 0 corresponds to the least-significant bit in the ipv6ExtensionHeadersFull IE while bit 255 corresponds to the most-significant bit of the IE. In doing so, few octets will be needed to encode common IPv6 extension headers when observed in a Flow."

Focus on IPv6 EH IEs (3)

Bit	Label	Protocol Number	Description
0	DST	60	Destination Options for IPv6
1	HOP	0	IPv6 Hop-by-Hop Options
2	NoNxt	59	No Next Header for IPv6
3	UNK		Unknown Layer 4 header (compressed, encrypted, not suppor
4	FRA0	44	Fragment header - first fragment
5	RH	43	Routing header
6	FRA1	44	Fragmentation header - not first f
7 to 11			Unassigned
12	MOB	135	Mobility Header
13	ESP	50	Encapsulating Security Payload
14	AH	51	Authentication Header
15			Unassigned
16	HIP	139	Host Identity Protocol
17	SHIM6	140	Shim6 Protocol
18		253	Use for experimentation and testing
19		254	Use for experimentation and testing
20 to 255			Unassigned

Registry created by draft-ietf-opsawg-ipfix-fixes

These are not EHs per se, but:

- UNK was already assigned in the existing ipv6ExtensionHeaders
- Added NoNxt as per a comment from Éric Vyncke (better observability)

The value was selected to minimize the implications on the use of reduced-encoding (rfc7011#section-6.2)

Focus on IPv6 EH IEs (4)

- Exporting Destination Options and Hop-by-Hop Options and Routing Types
  - Left out of scope
  - If there is a need to export specific options/type, we suggest to follow an approach similar to
    - draft-ietf-opsawg-ipfix-srv6-srh-14 (RFC-to-be 9487)
- Are you OK with this approach?

Focus on TCP (1)

Options are mapped to bits according to their option numbers. Option number X is mapped to bit X. TCP option numbers are maintained by IANA.

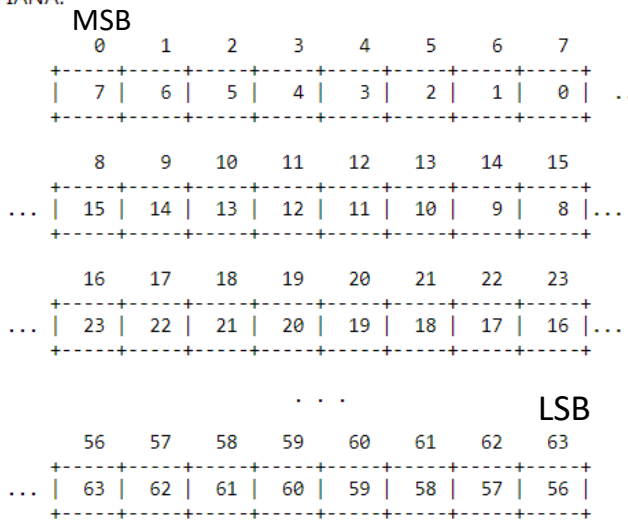

tcpOptions (Existing IE)

- "weird" mapping
- Inconsistency between the description vs. drawing
- Suboptimal encoding

Fix

draft-ietf-opsawg-ipfix-fixes

"~~*Option number X is mapped to bit X.*~~"

Fix

draft-ietf-opsawg-ipfix-tcpo-v6eh

"TCP option kind 0 corresponds to the least-significant bit in the tcpOptionsFull IE while kind 255 corresponds to the most-significant bit of the IE. This approach allows an observer to export any observed TCP option even if it does support that option and without requiring updating a mapping table."

Focus on TCP (2)

- TCP uses Experiments IDs (ExIDs) to disambiguate between shared TCP options
  - Two ExID flavors can be allocated: 2-byte or 4-byte ExIDs
  - A mix of ExIDs may be observed in a Flow
- Two IEs are defined to easily identify ExIDs
  - tcpSharedOptionExID16: List of 2-byte ExIDs
  - tcpSharedOptionExID32: List of 4-byte ExIDs
- We considered relaxing tcpSharedOptionExID32 to include both 2-byte and 4-byte ExIDs but this induces extra overhead
  - We decided to not include such a mention in the text

Misc.

- Éric Vyncke raised a comment during the call for adoption and also recently in the list
  - Split draft-ietf-opsawg-ipfix-tcpo-v6eh into two I-Ds: One of TCP and another one for IPv6 EHs
- The authors prefer to proceed with the current approach
- Thoughts?

Next Steps

- Request the WGLC for the set of I-Ds
  - with tcpm, tsvwg, 6man, and ipfix cced
- Consider early directorate reviews before or in // of the WGLC
  - simple-fixes: genart, opsdir
  - tcpo-v6eh: intdir, tsvart, opsdir
  - udp: tsvart, opsdir

Appendix: Example of Shared TCP Option

tcpSharedOptionExID16 IE:

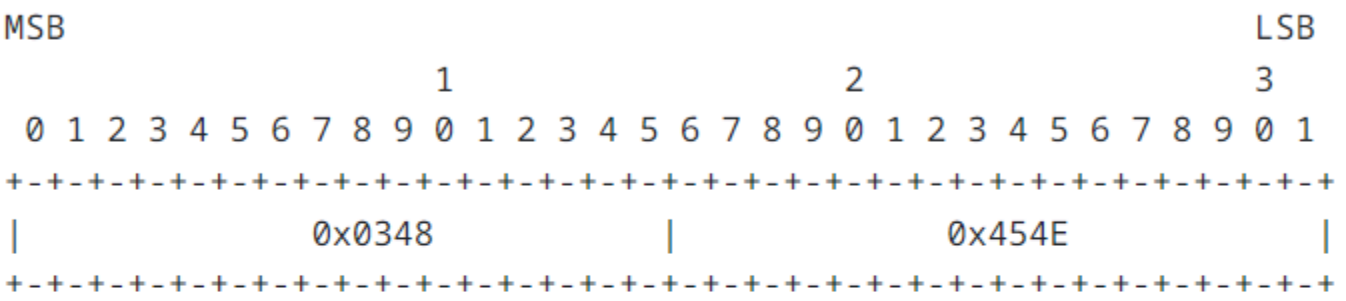

tcpSharedOptionExID32 IE: