In-situ OAM (IOAM) in VXLAN-GPE

draft-brockners-ioam-vxlan-gpe-00

Frank Brockners, Shwetha Bhandari, Vengada Prasad Govindan, Carlos Pignataro (Cisco)
Hannes Gedler (rtbrick), Steve Youell (JPMC), John Leddy (Comcast)
David Mozes (Mellanox), Tal Mizrahi (Marvell), Petr Lapukhov (Facebook)
Remy Chang (Barefoot), Daniel Bernier (Bell Canada)

IETF 100 – LISP; November, 2017

In-situ OAM in a nutshell

- Gather telemetry and OAM information along the path within the data packet, (hence “in-situ OAM”) as part of an existing/additional header
  - No extra probe-traffic (as with ping, trace, ..)
  - “Hybrid, Type-1 OAM” per RFC 7799
- Generic, Transport independent data-fields for IOAM
  - Scope: Per-hop, specific-hops only, end-to-end
  - Data fields include: Node IDs, interface IDs, timestamps, sequence numbers, ...
- Encapsulation
  - IOAM data fields can be embedded into a variety of transports, including: IPv6, SRv6, NSH, GRE, Geneve, VXLAN-GPE ...
- *Base IOAM document adopted by IPPM!*
  - *draft-ietf-ippm-ioam-data-01*

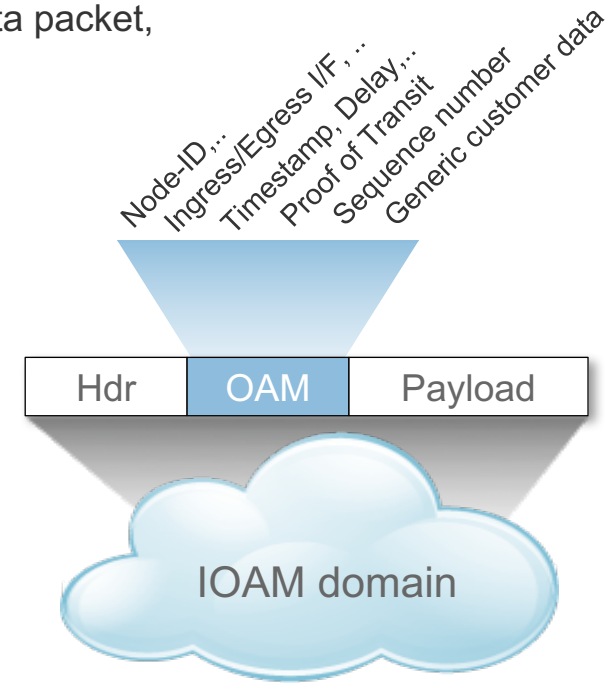

IOAM over VXLAN-GPE

(draft-brockners-ioam-vxlan-gpe-00)

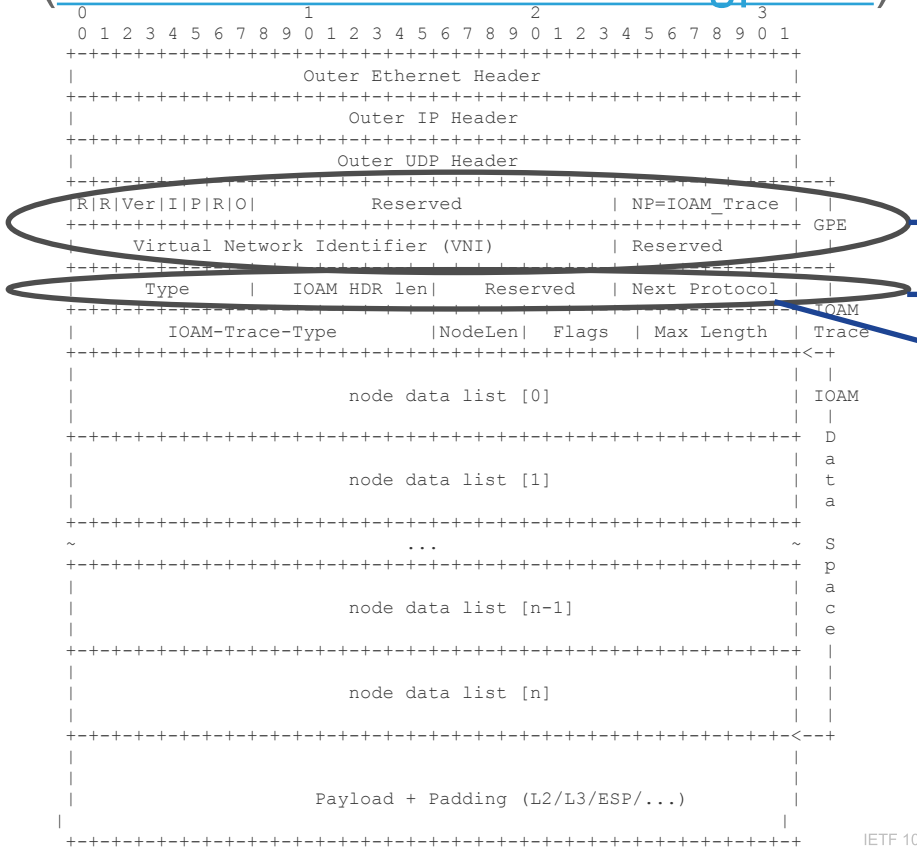

VXLAN-GPE header

IOAM shim header for VXLAN-GPE

Specifies the next protocol
Allows multiple data field

Open Questions

- Section 4: Discussion of the encapsulation approach
  - An encapsulation of IOAM data fields in VXLAN-GPE should be friendly to an implementation in both hardware as well as software forwarders and support a wide range of deployment cases, including large networks that desire to leverage multiple IOAM data fields at the same time.
- "To nest TLVs or not to"?, Single GPE protocol type for all IOAM data categories (trace, proof-of-transit, and edge-to-edge) or one GPE protocol type per IOAM data category?
  - "Next header" approach: Single GPE protocol type for all IOAM types: Results in serial list of IOAM data categories. Avoids iterative lookups. Finding the L4 header requires parsing each header of the list. Current approach in -00 version of the draft.
  - "TLV" approach: One GPE protocol type per IOAM category: Results in nested TLVs. Requires iterative lookups. Single length field for all options allows nodes not interested in the IOAM information to skip the information easily.

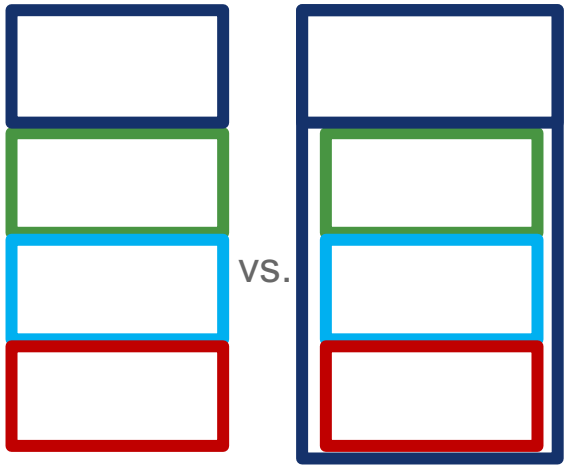

Status and next steps

- Data Fields for In-situ OAM
  - Adopted by IPPM
  - draft-ietf-ippm-ioam-data
- IOAM encapsulations into protocols
  - draft-brockners-nvo3-ioam-geneve-00
  - draft-brockners-ioam-vxlan-gpe-00 (this discussion)
  - draft-brockners-sfc-ioam-nsh-00
  - … more to come…
- IOAM in VXLAN-GPE in LISP WG
  - Feedback from LISP WG appreciated, especially on open questions
  - WG adoption?