How OAM Identified in Overlay Protocols

draft-mirsky-rtgwg-oam-identify

Greg Mirsky

IETF-104 March 2019, Prague

Problem statement

- How to achieve unambiguous identification of OAM?
- Active OAM uses specifically constructed packets – test packets.
  - Fault Management and Performance Monitoring (‘F’ and ‘P’ in FCAPS)
  - Single-ended vs. dual-ended, e.g., ping vs. BFD in Async mode
  - Two-way vs. one-way, e.g., Echo request/reply vs. BFD in Demand mode
- Hybrid OAM, according to RFC 7799, is an OAM method that combines properties of passive and active measurement methods:
  - Alternate Marking method triggers measurement
  - In-situ OAM triggers measurement, collects and transports the measurement results, network state information, a.k.a. telemetry information, in the data packet itself
  - The Hybrid Two-Step method collects and transports the telemetry information on-path in a follow-up packets
- Overlay network protocols use:
  - encapsulations that support optional meta-data, i.e., variable size headers (Geneve, SFC NSH, GUE)
  - encapsulations that use fixed-size headers (BIER, VXLAN-GPE)

SFC NSH and Active OAM

RFC 8300 Network Service Header:

O bit: Setting this bit indicates an OAM packet.

draft-ietf-sfc-multi-layer-oam :

O bit: Setting this bit indicates an OAM command and/or data in the NSH Context Header or packet payload.

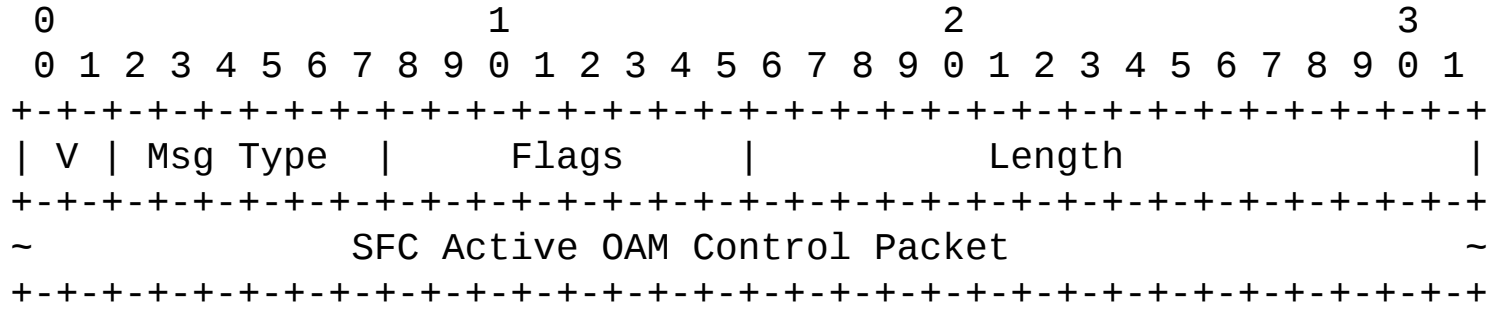

O bit and the Next Protocol interpretation I

- O bit set, and the Next Protocol value is not one of identifying active or hybrid OAM protocol (per [RFC7799] definitions), e.g., defined in this specification Active SFC OAM - a Fixed-Length Context Header or Variable-Length Context Header(s) contain OAM command or data, and the type of payload determined by the Next Protocol field

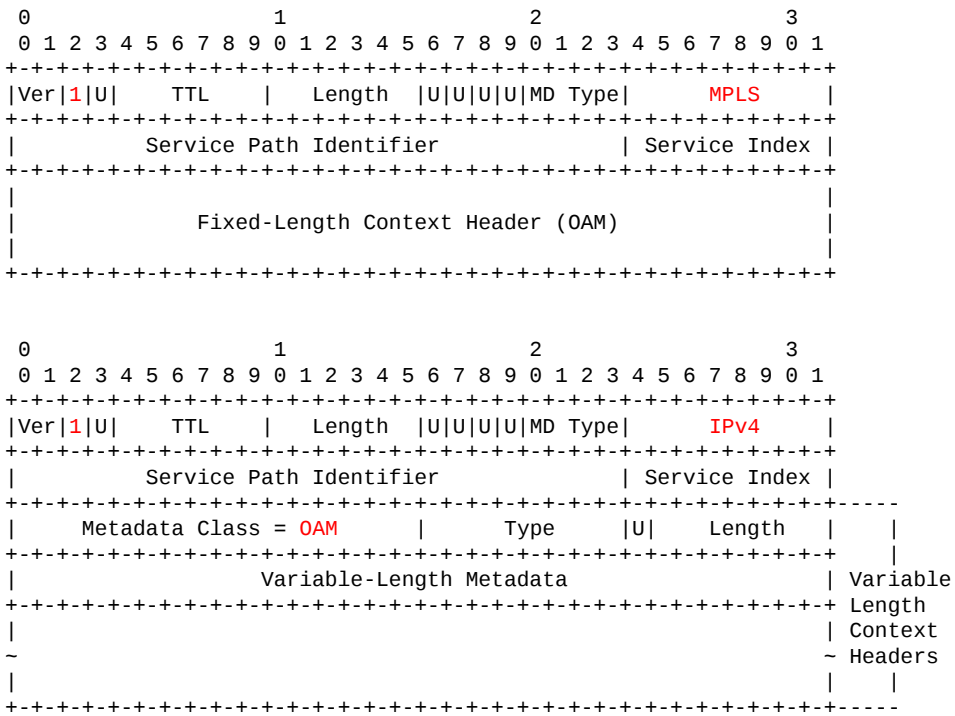

O bit and the Next Protocol interpretation II

- O bit set, and the Next Protocol value is one of identifying active or hybrid OAM protocol - the payload that immediately follows SFC NSH contains OAM command or data;

O bit and the Next Protocol interpretation III

- O bit is clear - no OAM in a Fixed-Length Context Header or Variable-Length Context Header(s) and the payload determined by the value of the Next Protocol field;

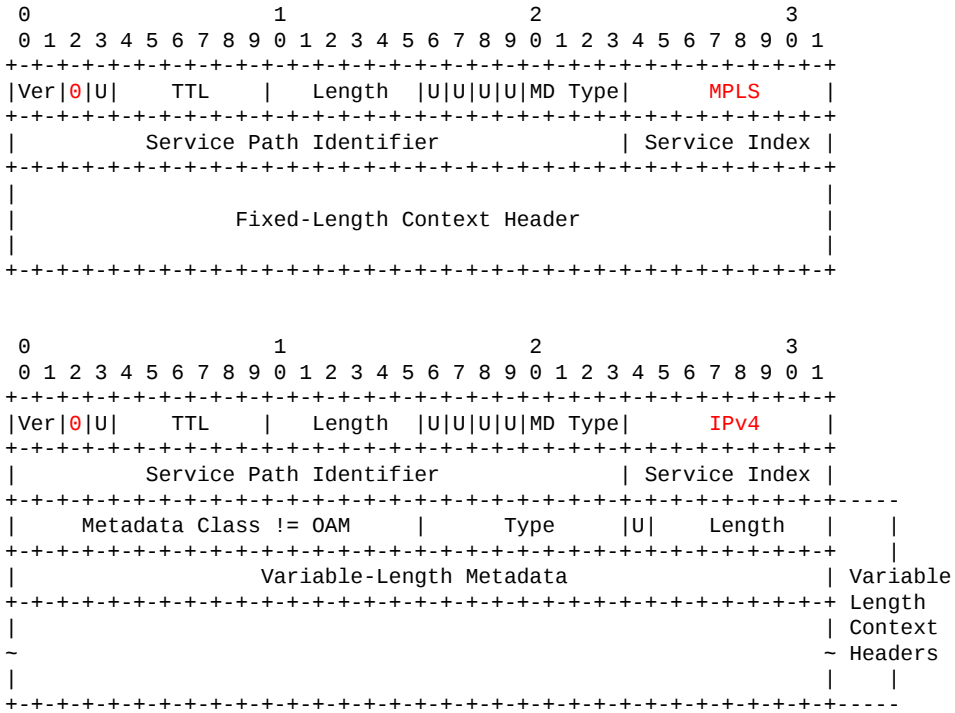

O bit and the Next Protocol interpretation IV

- O bit is clear and the Next Protocol value is one of identifying active or hybrid OAM protocol MUST be identified and reported as the erroneous combination. An implementation MAY have control to enable processing of the OAM payload. For example, in case the Fixed-Length Context Header being used:

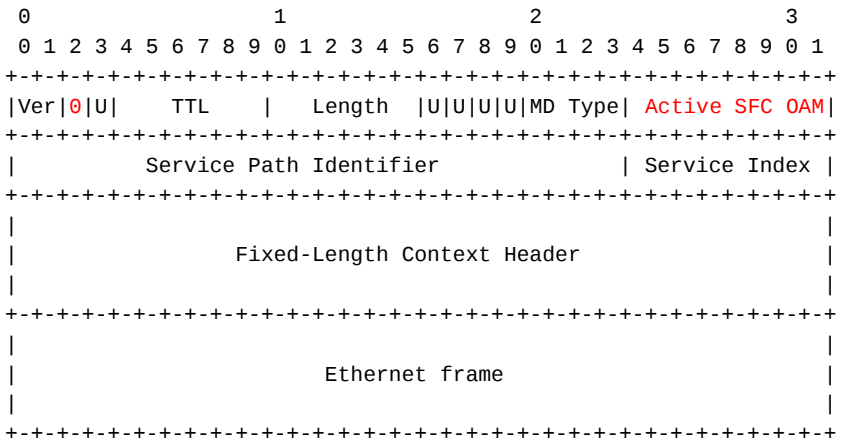

- … the recommendation to avoid combination of OAM in a Fixed-Length Context Header or Variable- Length Context Header(s) and in the payload immediately following the SFC NSH because there is no unambiguous way to identify such combination using the O bit and the Next Protocol field.

Overlay Tunnels and OAM

draft-ietf-nvo3-geneve:

- Definition of O bit has changed: s/OAM packet/Control packet/

  O (1 bit): Control packet. This packet contains a control message. Control messages are sent between tunnel endpoints. Tunnel Endpoints MUST NOT forward the payload and transit devices MUST NOT attempt to interpret it. Since these are infrequent control messages, it is RECOMMENDED that tunnel endpoints direct these packets to a high priority control queue (for example, to direct the packet to a general purpose CPU from a forwarding ASIC or to separate out control traffic on a NIC). Transit devices MUST NOT alter forwarding behavior on the basis of this bit, such as ECMP link selection.

  draft-ietf-intarea-gue:

  C-bit provides the separate namespace to “carry formatted data that are implicitly addressed to the decapsulator to monitor or control the state or behavior of a tunnel. … The payload is interpreted as a control message with type specified in the proto/ctype field. The format and contents of the control message are indicated by the type and can be variable length.”

Fixed-size header and OAM

RFC 8296 4 Encapsulation for BIER in MPLS and Non-MPLS Networks:

OAM packet identified by the value of the Next Protocol field. IANA BIER Next Protocol Identifiers registry includes the identifier for OAM (5).

draft-ietf-nvo3-vxlan-gpe (expired):

OAM Flag Bit (O bit): The O bit is set to indicate that the packet is an OAM packet.

Next steps

- Non-IP encapsulation of OAM packets over MPLS underlay
- Your comments, suggestions, questions always welcome and greatly appreciated
- WG adoption? (May not need to publish but it may serve to reflect on the discussion)