How OAM Identified in Overlay Protocols

draft-mirsky-rtgwg-oam-identify

Greg Mirsky

IETF-102 July 2018, Montreal

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)

Overlay Tunnels and OAM

draft-ietf-nvo3-geneve:

O (1 bit): OAM packet. This packet contains a control message instead of a data payload. Endpoints MUST NOT forward the payload and transit devices MUST NOT attempt to interpret or process it. Since these are infrequent control messages, it is RECOMMENDED that 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.”

SFC NSH and OAM

RFC 8300 Network Service Header:

O bit: Setting this bit indicates an OAM packet.

draft-ietf-sfc-oam-framework:

While the presence of OAM marker in the overlay header (e.g., O bit in the NSH header) indicates it as OAM packet, it is not sufficient to indicate what OAM function the packet is intended for.

draft-ietf-sfc-ioam-nsh:

... the O bit MUST NOT be set for regular customer traffic which also carries IOAM data and the O bit MUST be set for OAM packets which carry only IOAM data without any regular data payload.

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:

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

Dual OAM Identification

- What is the definition of OAM packet?
  - OAM commands or information present in the header as meta-data
  - OAM commands or information immediately follow the header
- What are the issues with using the flag field to identify “OAM packet”?
  - The ambiguity of O-bit being set and a non-zero value in the Next Protocol field
  - RFC 8393 Operating the Network Service Header (NSH) with Next Protocol "None“ suggests the optional use of O-bit set and the value None in the Next Protocol field as active OAM with OAM commands or information in MD Type=2. The problem is that the Length is limited to 512 octets (Geneve limits TLV to 128 octets of payload). That will affect applicability in Service Activation Testing and OAM methods that must generate synthetic traffic with a wider range of packets (up to Jumbo frame size)
- . How to identify OAM command or information immediately follow the header?

Recommendations

OAM control commands and data may be present as part of the overlay encapsulation header or as a payload that follows the overlay network header.

The recommendations:

- OAM in the overlay header, if supported by the overlay network, identified by the dedicated flag. Use of this method as active OAM is possible but functionality is limited.
- The scope of the OAM flag is the meta-data included in the header.
- OAM that follows the overlay header identified as payload type, e.g. by the value of the Next Protocol field.

Active OAM: Source Identifier

- OAM packet source identifier is required for single-ended two-way methods, e.g., echo request/reply
- IP underlay natively provides the Source ID
- When the underlay is MPLS data plane options are:
  - use, if available, Sender ID in the overlay, e.g., BFR ID in BIER
  - use IP/UDP encapsulation with the destination IP address from the “martian” range, e.g., from 127/8 range for IPv4 and from the 0:0:0:0:0:FFFF:7F00/104 range for IPv6
  - non-IP encapsulation to avoid the overhead of extra IP/UDP encapsulation

On-path OAM Identification

OAM toolset may include methods that don't use specially constructed and injected in the network test packets. RFC 7799 defines OAM methods that are neither entirely active nor passive but are combine both as hybrid methods. Examples:

- The Alternate Marking Method (AMM) (RFC 8321) uses the dedicated field in the header to trigger a performance measurement (packet loss, delay, delay variation, residence time, etc.). Applicability of AMM discussed in the number of drafts: BIER, SFC, NVO3, IPv6 (V6OPS WG)
- In-situ OAM (iOAM) uses OAM flag in the header and/or values of the Next Protocol field to identify iOAM commands and/or data
- Hybrid Two-Step (HTS) collects and transports the measurement results and/or the telemetry information in the payload immediately following the header. Uses the value of the Next Protocol field.

Next steps

- Non-IP encapsulation of OAM packets over MPLS underlay
- Update the document on GUE (thanks to Tom Herbert)
- 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)