In-situ OAM IPv6 Options

draft-ietf-ippm-ioam-ipv6-options-01

S. Bhandari, F. Brockners, C. Pignataro, H. Gredler, J. Leddy, S. Youell, T. Mizrahi,
A. Kfir, B. Gafni, P. Lapukhov, M. Spiegel, S. Krishnan, R. Asati

IPPM WG Meeting, IETF 108
July 31, 2020

Status of draft-ietf-ippm-ioam-ipv6-options-01

- Draft has been stable;
- Early allocation of v6 Option-Types has been completed

- Open Issue (which surfaced thanks to the work of Univ. of Liege IOAM Kernel Implementation):
  Forwarding behavior for nodes which do not support IOAM needs clarification; see email thread:
  https://mailarchive.ietf.org/arch/msg/ippm/F3Ew6hBp6xihMajU7iV0hTtqPYo/

Background: IOAM Deployment w/ IPv6

Deployment Options for IOAM with IPv6
(from draft-ioametal-ippm-6man-ioam-ipv6-deployment-03):

- IOAM domains bounded by hosts
- IOAM domains bounded by network devices
  - IPv6-in-IPv6 encapsulation
  - IP-in-IPv6 encapsulation with ULA
  - x-in-IPv6 encapsulation that is used Independently

➡ Encapsulation with another IPv6 header ensures that packets with IOAM option types will always stay within the IOAM domain.

Background: IOAM Deployment w/ IPv6

Assumptions about IOAM support by nodes within a domain

Different IOAM Option-Types have different assumptions about IOAM support in an IOAM Domain:

- IOAM Trace Option-Types:
  - Encap-node, decap-node, transit nodes which an operator desires to track
    (i.e. not all transit nodes might (or have to) support IOAM)
- IOAM POT Option-Type:
  - Encap-node, decap-node, transit nodes which are to “proof transit”
    (i.e. typically only a subset of all transit nodes support IOAM)
- IOAM DEX Option-Type:
  - Encap-node, decap-node, transit nodes which an operator desires to track
    (i.e. not all transit nodes might (or have to) support IOAM)
- IOAM E2E Option-Type:
  - Encap-node, decap-node
    (i.e. no need for transit nodes to support IOAM)

IOAM Option Type Assignments

Early Allocation of IOAM Option Types:

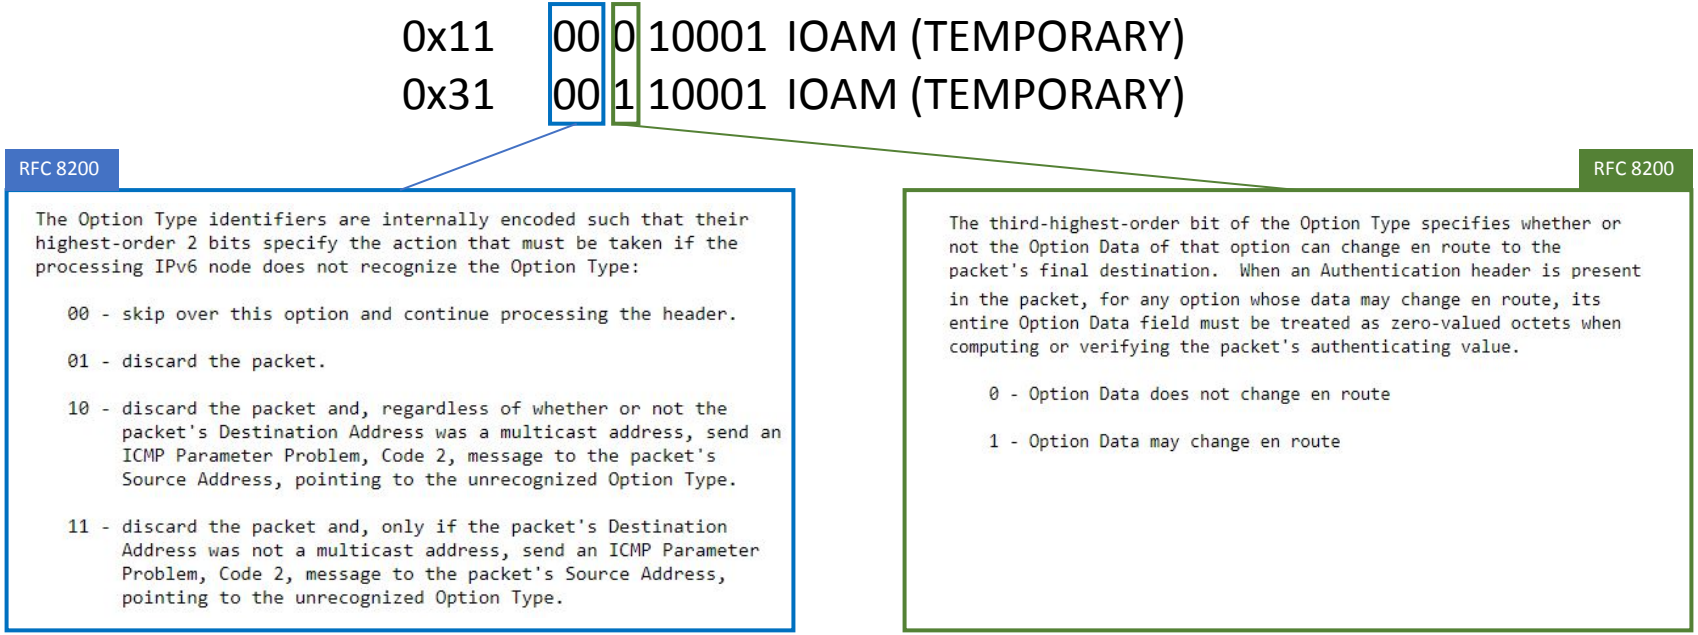

draft-ietf-ippm-ioam-ipv6-options-01 forwarding behavior

Forwarding behavior defined through Option-Types differs from that defined in draft-ietf-ippm-ioam-ipv6-options-01:

> In order for IOAM to work in IPv6 networks, IOAM MUST be explicitly enabled per interface on every node within the IOAM domain. Unless a particular interface is explicitly enabled (i.e. explicitly configured) for IOAM, a router MUST drop packets which contain extension headers carrying IOAM data-fields. This is the default behavior and is independent of whether the Hop-by-Hop options or Destination options are used to encode the IOAM data. This ensures that IOAM data does not unintentionally get forwarded outside the IOAM domain.

This text was originally introduced to ensure packets with IOAM options cannot “leak”;
Given that IOAM either is between hosts or uses encapsulation into IPv6 between edge routers of a domain, packets with IOAM options cannot leak:
Text is no longer required.

Suggested Approach / Next Steps

- Remove contradicting paragraph from draft-ietf-ippm-ioam-ipv6-options

- Adopt draft-ioametal-ippm-6man-ioam-ipv6-deployment-03 as WG document?

In-situ OAM Flags

draft-ietf-ippm-ioam-flags-01

Tal Mizrahi, Frank Brockners, Shwetha Bhandari, Ramesh Sivakolundu,
Carlos Pignataro, Aviv Kfir, Barak Gafni, Mickey Spiegel, Jennifer Lemon

IETF 108, IPPM
July 2020

Status of this Draft

- Version 01 (hopefully) addressed the main issue that were raised - amplification attacks.

- Discussed in the IPPM interim meeting in April 2020.

- Still expecting feedback from people who volunteered to review this issue.

- We still have an open issue about loopback on the reverse path – to be continued on the mailing list.

Open Issue – Loopback Flag

Loopback on the reverse path:

- Pushing IOAM data on the reverse path is not necessary.
- Problem: how do transit nodes know that a looped back packet is in transit on the reverse path?
  - New flag?
  - New IOAM type?
  - Clearing the RemainingLen field when the packet is looped back?

In-situ OAM Direct Exporting

draft-ietf-ippm-ioam-direct-export-00

Haoyu Song, Barak Gafni, Tianran Zhou, Zhenbin Li,
Frank Brockners, Shwetha Bhandari, Ramesh Sivakolundu, Tal Mizrahi

IETF 108, IPPM,
July 2020

Status of this Draft

- This draft is the product of a design team that worked on combining two documents (PBT-I and immediate exporting).

- One main open issue remains to be resolved – Hop Count field.

Open Issue – Hop Count

- Question: should the DEX option include an explicit Hop Count field, or is the Hop_Lim/Node_ID data field sufficient?

- No Hop Count:
  - Using existing functionality: Hop_Lim/Node_ID data field can be used, copied from the TTL/Hop Limit from the lower layer, and included in the exported packet.
  - The DEX option does not need to be modified by transit switches.

- Explicit Hop Count:
  - The lower layer TTL may not be accurate, e.g., L2 or hierarchical VPN.
  - Allows to detect IOAM-capable node that fails to export packets.