Internet Control Message Protocol (ICMPv6) Loopback

draft-mcb-6man-icmpv6-loopback

Tal Mizrahi, Tianran Zhou, Shahar Belkar, Reuven Cohen, Justin Iurman

IETF 120
July 2024

ICMPv6 Loopback

- Two new ICMPv6 types:
  Loopback Request, Loopback Reply.
- Similar to ICMPv6 Echo, except that The Loopback Reply includes the Loopback Request in its payload.
- The sender receives:
  - IPv6 options of the original request.
  - IPv6 header fields, e.g., hop limit.

- This is how ICMPv6 error messages work today.
- This is how Traceroute works today. (ICMPv6 Time Exceeded).

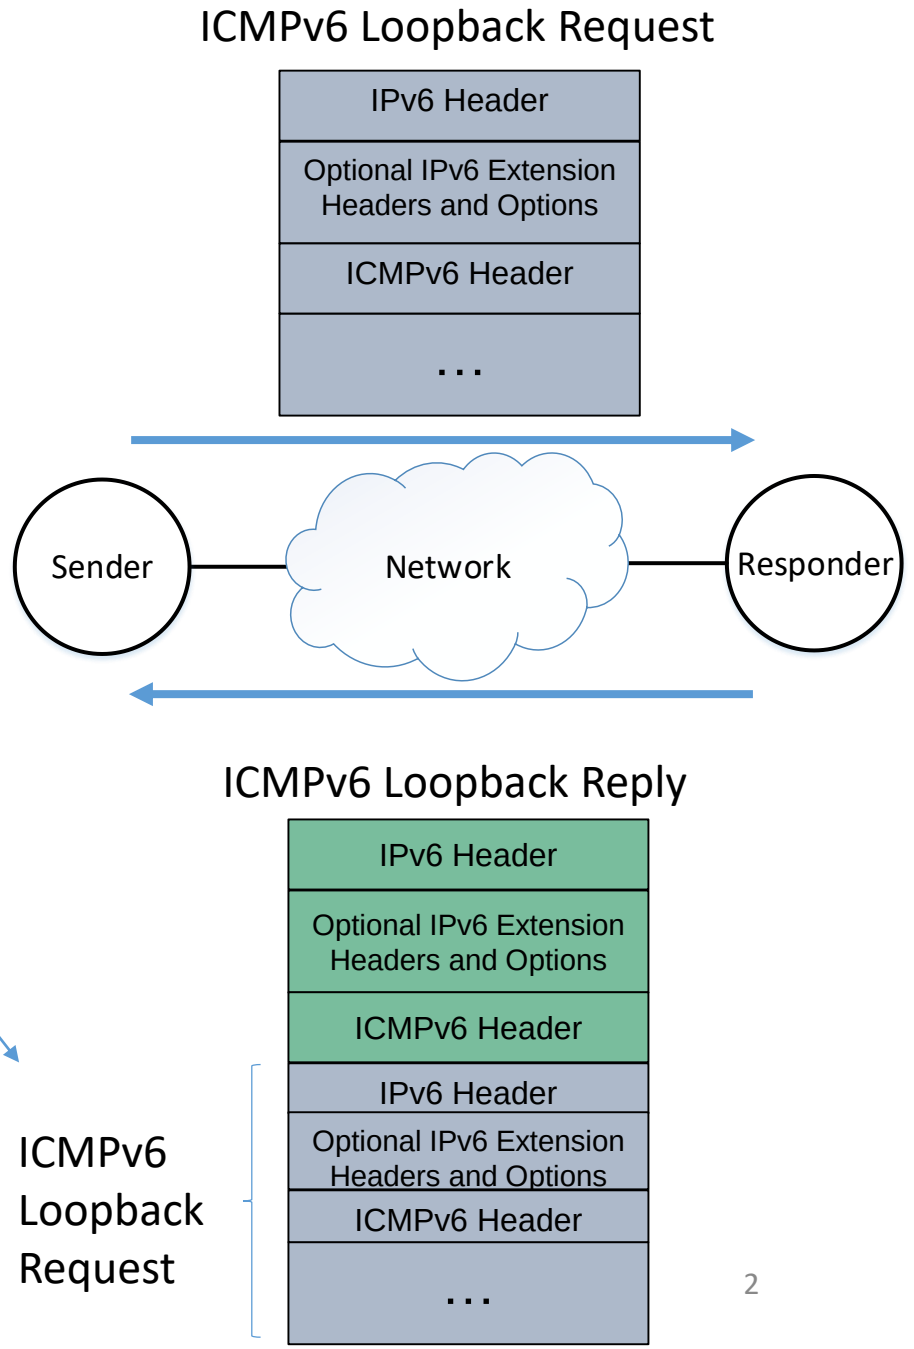

The IOAM Use Case

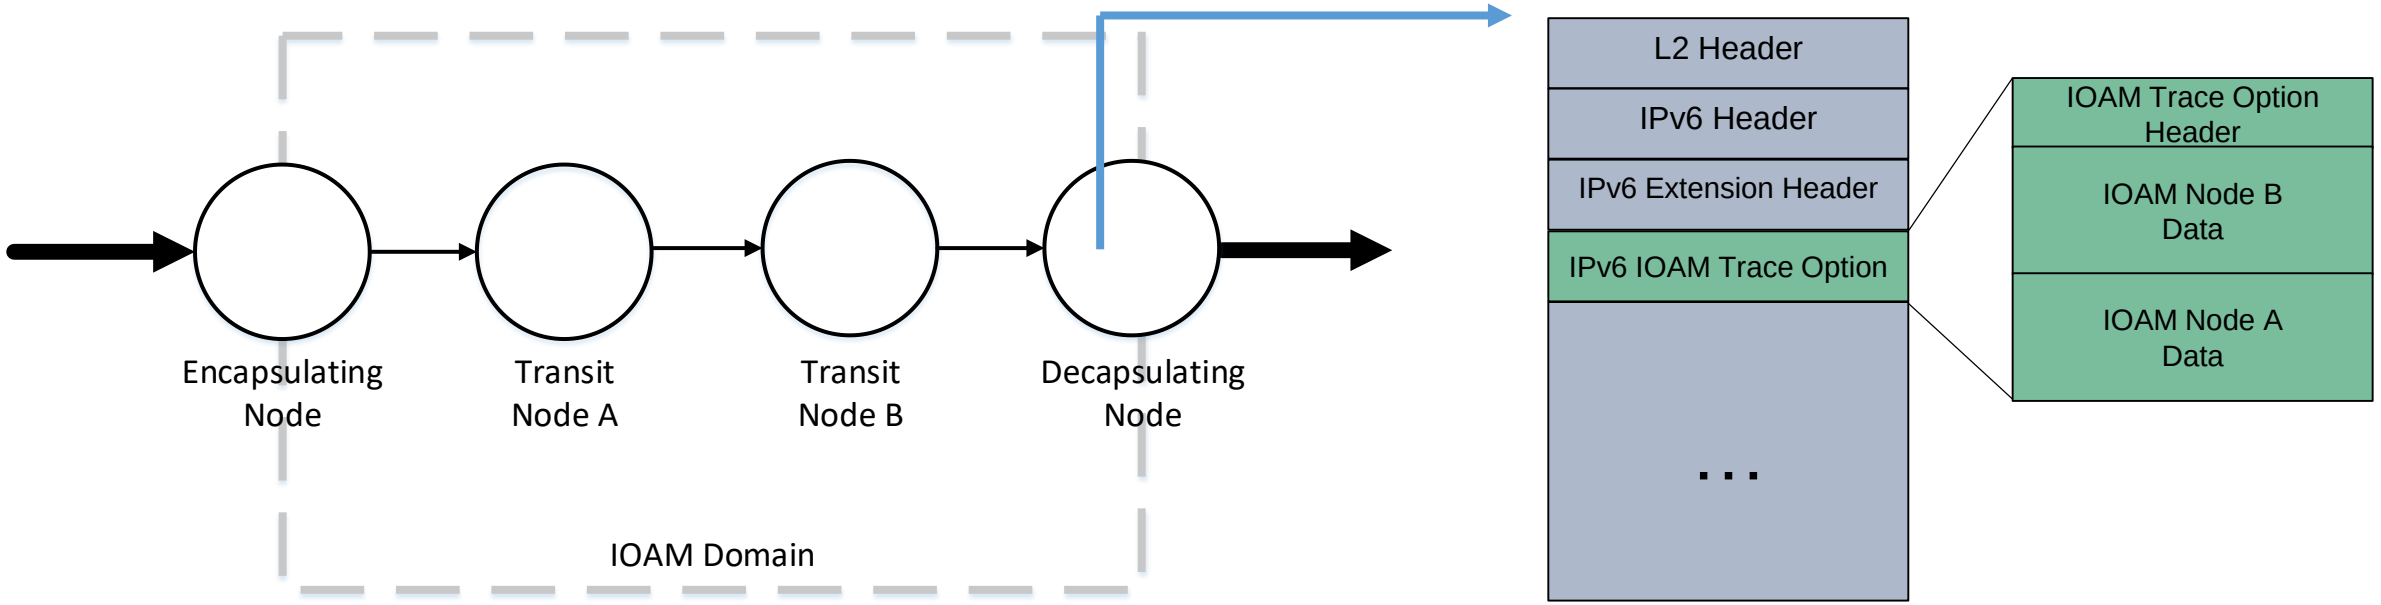

- IOAM [RFC 9197] can be used to measure and monitor an IOAM domain.
- The IPv6 IOAM Trace Option includes per-hop information, e.g., node ID, timestamp, queue depth.
- **Use case**: we want the encapsulating node to invoke a Ping-like request and receive the IOAM information back from the decapsulating node.

Other Use Cases

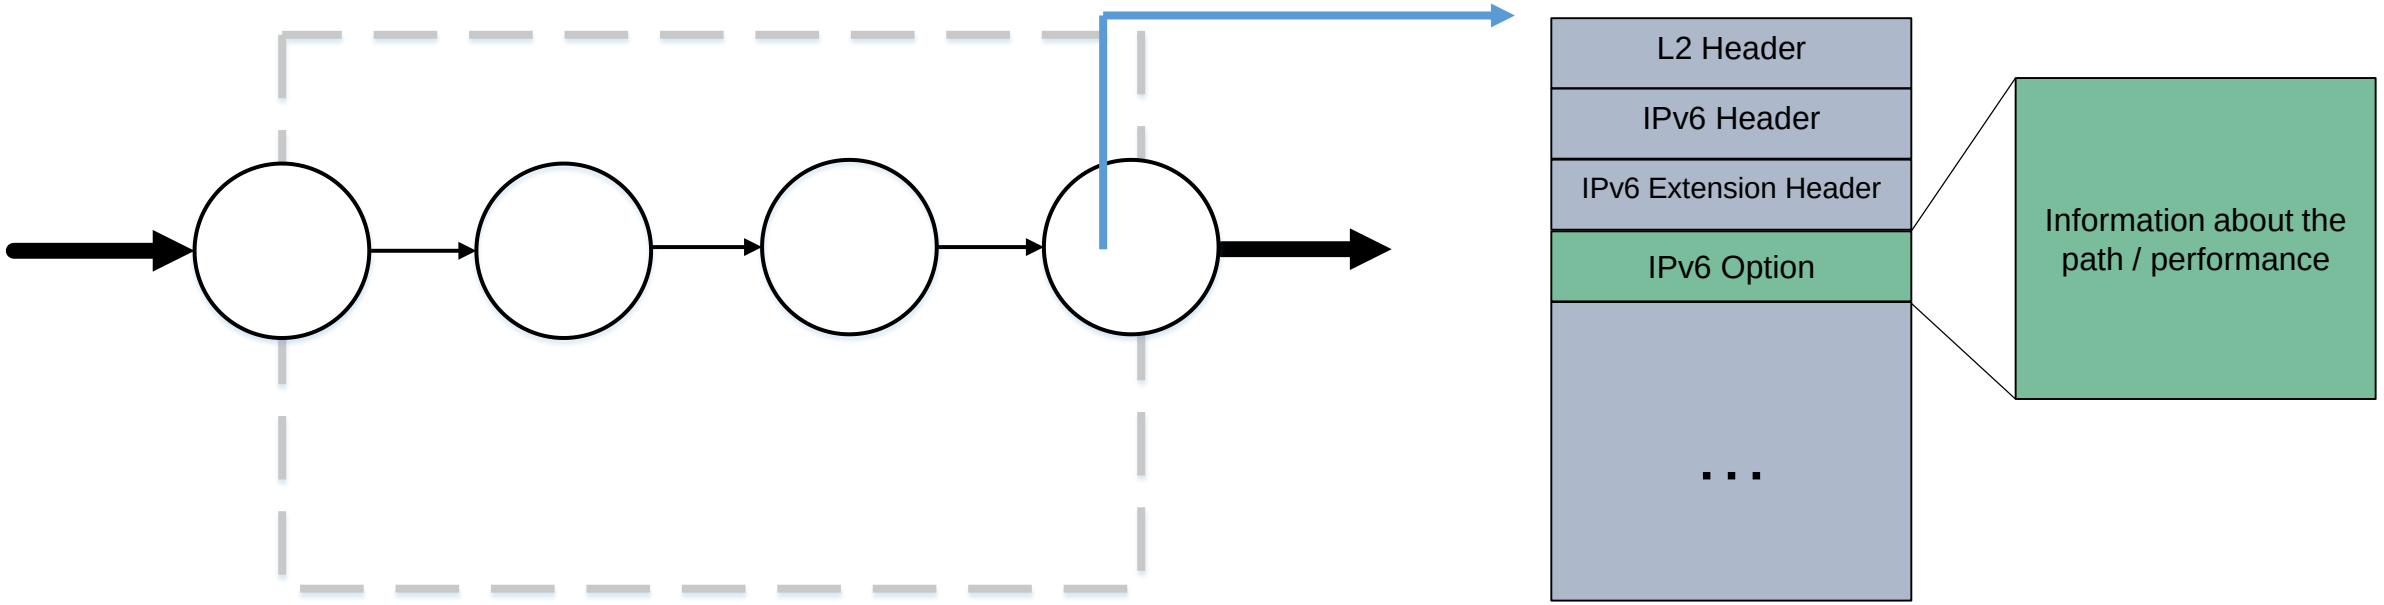

- Use case: we want the sender to invoke a Ping-like request and receive the **IPv6 options of the forward path** in the reply.
- Many in-progress and future protocols may benefit from this approach:
  - draft-filsfils-ippm-path-tracing
  - draft-kumar-ippm-ifa
  - “Follow up” of hybrid two-step measurement (draft-ietf-ippm-hybrid-two-step)
  - SRH reflection
  - Minimum MTU reporting (as in [RFC9268])
  - …

Demo: IOAM Ping

- We have prepared a demo utility of Ping that displays per-hop IOAM information.
- Uses **ICMPv6 Loopback**.
- Code is on Github: https://github.com/talmi/IOAM-Ping-Demo

```
user@ubuntu22:~/IOAM-Ping-Demo$ sudo ./ioam-ping.sh
tcpdump: listening on veth0, link-type EN10MB (Ethernet), snapshot length 262144 bytes
PING db01::1 (db01::1) 56 data bytes
152 bytes from db01::1: icmp_seq=1 ttl=64 time=0.124 ms IOAM: NodeID=1,HopLim=64,RcvTime=NA NodeID=2,HopLim=63,RcvTime=1721554104.948
152 bytes from db01::1: icmp_seq=2 ttl=64 time=0.087 ms IOAM: NodeID=1,HopLim=64,RcvTime=NA NodeID=2,HopLim=63,RcvTime=1721554105.983
152 bytes from db01::1: icmp_seq=3 ttl=64 time=0.086 ms IOAM: NodeID=1,HopLim=64,RcvTime=NA NodeID=2,HopLim=63,RcvTime=1721554106.993
152 bytes from db01::1: icmp_seq=4 ttl=64 time=0.099 ms IOAM: NodeID=1,HopLim=64,RcvTime=NA NodeID=2,HopLim=63,RcvTime=1721554108.33
152 bytes from db01::1: icmp_seq=5 ttl=64 time=0.109 ms IOAM: NodeID=1,HopLim=64,RcvTime=NA NodeID=2,HopLim=63,RcvTime=1721554109.45

--- db01::1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4097ms
rtt min/avg/max/mdev = 0.086/0.101/0.124/0.014 ms
12 packets captured
12 packets received by filter
0 packets dropped by kernel
```

- Thanks to: Ben Ben Ishay, Amit Geron, Justin Iurman and Tal Mizrahi.

Summary of Discussions on the Mailing Lists

- Considering which working group this belongs in:
  - 6MAN / IPPM / INT-AREA
- Defining two new ICMPv6 types vs. Defining two new codes for ICMPv6 Echo.
  - Detailed discussion in the draft.
- Reply is longer than Request vs. symmetric Request/Reply.
  - Current draft defines symmetric mode to mitigate amplification.
- Should Loopback be defined for ICMPv4 as well?

Next Steps

- Potentially merge with draft-he-6man-icmpv6-extensions-ipv6-ext-header.
- WG adoption call.