IPv6 Application of the Alternate Marking Method

draft-ietf-6man-ipv6-alt-mark-04

Online, Mar 2021, IETF 110

Giuseppe Fioccola (Huawei)
Tianran Zhou (Huawei)
Mauro Cociglio (Telecom Italia)
Fengwei Qin (China Mobile)
Ran Pang (China Unicom)

Alternate Marking at a glance

Alternate Marking methodology (**RFC 8321**) is an OAM Passive PM technique

- Batching packets based on time interval to measure **Packet Loss** by switching value of L flag.
- Use D flag to create a new set of marked packets: D-marked packets to calculate **more informative Packet Delay Metrics**

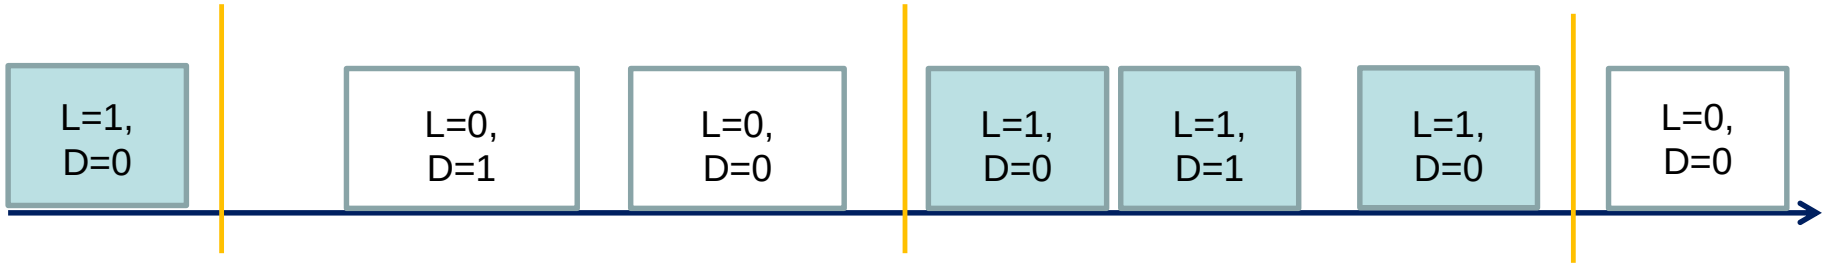

Multipoint Alternate Marking methodology (**RFC 8889**) generalizes the application for multipoint unicast flows and allows a flexible performance management approach

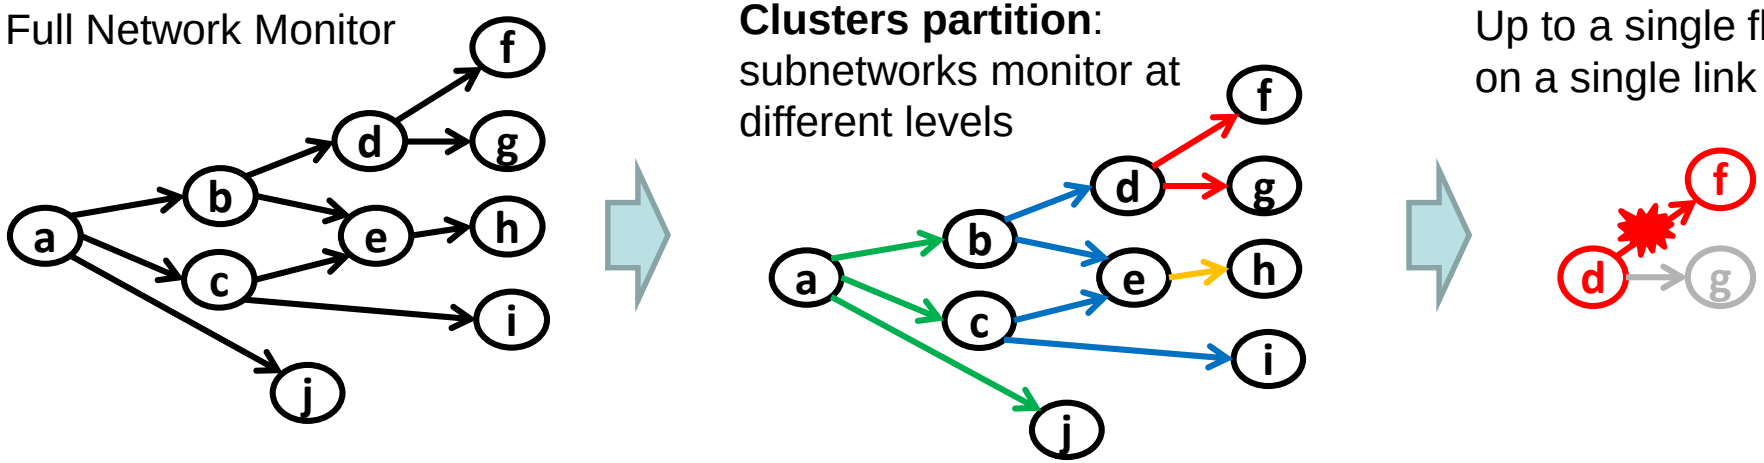

Alternate Marking Data Fields

- Definition of a new TLV to be encoded in the Options Header
- The **AltMark Option** is expected to be encapsulated as Hop-by-Hop Options Header or Destination Options Header.

Skip if do not recognize and data do not change en route

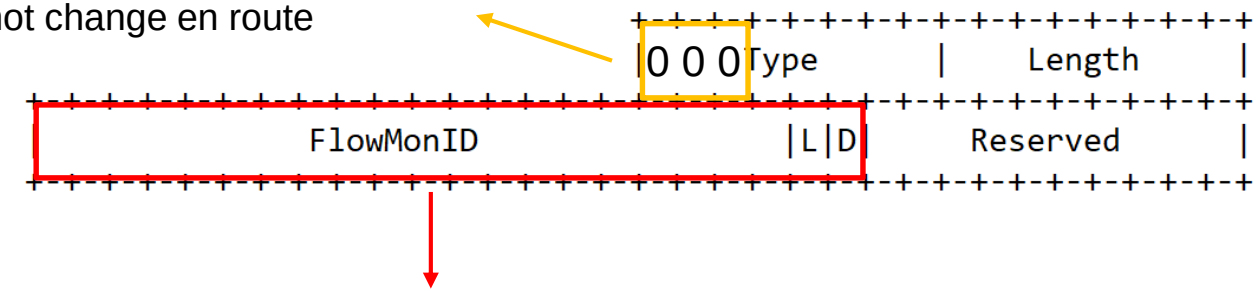

- **L** and **D** are the Marking Fields
- The Flow Monitoring Identification (**FlowMonID**) is required for specific deployment reasons (see next slide)

- The **source node** is the only one that writes the Option Header to mark alternately the flow (for both Hop-by-Hop and Destination Option).
- In case of **Hop-by-Hop Option Header**, it can only be read by the **intermediate nodes** along the path. The measurement is hop-by-hop.
- In case of **Destination Option Header**, it is not processed by any node until the packet reaches the **destination node**. The measurement is end-to-end.

Flow Monitoring Identification

The Flow Monitoring Identification (**FlowMonID**) is required for the following reasons:

- ✓ **It helps to reduce the per node configuration**. Using a flow identifier also allows a flexible granularity for the flow definition.
- ✓ **It simplifies the counters handling**. Hardware processing of flow tuples is challenging and often incurs into performance issues, especially in tunnel interfaces.
- ✓ **It eases the data export** encapsulation and correlation for the collectors.

How to allow disambiguation of the FlowMonID in case of collision.

1) In case of a **centralized controller**, it should set FlowMonID and instruct the nodes properly in order to guarantee its uniqueness.

2) FlowMonID can be **pseudo randomly generated by the source node**

- if the 20 bit FlowMonID is set independently and pseudo randomly there is a chance of collision (50% chance of collision for just 1206 flows!)
- For more entropy, FlowMonID can be combined with other identifying flow information in the packet (e.g. IP addresses and Flow Label)

AltMark EH Option alternatives

In summary, here are the alternative options based on the chosen type of PM:

- ✓ **Destination Option not preceding a Routing Header** => measurement only by node in Destination Address.
- ✓ **Hop-by-Hop Option** => every router on the path with feature enabled.
- ✓ **Destination Option preceding a Routing Header** => every destination node in the route list.

In many cases the end-to-end measurement is not enough and it could be required the hop-by-hop measurement.

- Nodes that do not support the Hop-by-Hop Option SHOULD ignore them. In this case, the measurement does not account for all links and nodes along a path.

Security Considerations

Security concerns:

- **Harm caused by the measurement**: Alternate Marking implies modifications on the fly to an Option Header by the source node
  - This must be performed in a way that does not alter the QoS experienced by the packets and that preserves stability of routers doing the measurements.

- **Harm to the Measurement**: Alternate Marking measurements could be harmed by routers altering the marking of the packets or by an attacker injecting artificial traffic.
  - In the context of a **controlled domain**, the network nodes are locally administered and this type of attack can be avoided
  - An **attacker cannot gain information** about network performance **from a single monitoring point** but it should be able to use multiple and synchronized monitoring points

Privacy concerns are limited because the method only relies on information contained in the Option Header without any release of user data.

  - The limited marking technique seems unlikely to substantially increase the existing privacy risks from header or encapsulation metadata.

Changes after Last Call (1/2)

Main inputs from Brian Carpenter

- ✓ Clarification about HbH processing.
  - The 3 high-order bits of the AltMark Option are 000 and this means "skip if do not recognize and data do not change en route“ (RFC8200 and draft-hinden-6man-hbh-processing)
  - RFC8200 also mentions that the nodes only examine and process the HbH Options header if explicitly configured to do so.

  Anyway, in practice, the things may be different and it can happen that packets with HbH are forced onto the slow path, and this is a general issue.

- ✓ It was suggested to include a proper sub-section "Controlled Domain“ in order to better highlight this important point. It answers the question of whether this new option is deployable (reference to RFC8799)

- ✓ Rewording of the description of Destination Option usage with Routing Header, to avoid ambiguity and be consistenbt with RFC8200

- ✓ Uniqueness of FlowMonID: it may not be a problem in some cases. But, for large scale measurements, the disambiguation of the FlowMonID is something to consider.

Changes after Last Call (2/2)

Inputs from Greg Mirsky

Alternate Marking bits are carried by the Options Header and it may have some impact on the path MTU.

- This point has been included in the Security Considerations.
- Anyway the relative small size (48 bit in total) of this HBH/DOH and its application to a controlled domain mitigate the problem.

Revisions from Bob Hinden and Ole Troan

- Several editorial comments addressed

Next Steps

- WGLC ended in January
- An agreed way to apply RFC 8321 and RFC 8889 to IPv6

- Welcome questions, comments

Thank you