Network
Internet-Draft
Intended status: Standards Track
Expires: 31 July 2026

Shaofu. Peng
ZTE Corporation

27 January 2026

Delay Options
draft-peng-6man-delay-options-01

Abstract

This document introduces new IPv6 options for HBH or DOH Options header, to carry delay related information for deterministic forwarding.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 31 July 2026.

Table of Contents

1. Introduction

[RFC8655] describes the architecture of deterministic network and defines the QoS goals of deterministic forwarding: Minimum and maximum end-to-end latency from source to destination, timely delivery, and bounded jitter (packet delay variation); packet loss ratio under various assumptions as to the operational states of the nodes and links; an upper bound on out-of-order packet delivery. In order to achieve these goals, deterministic networks use resource reservation, explicit routing, service protection and other means. In general, a deterministic path is a strictly explicit path calculated by a centralized controller, and resources are reserved on the nodes along the path.

To provide deterministic forwarding service, the scheduling mechanisms applied in the network generally require application flows to comply with predefine constraints, such as a token bucket specification consisting of a "token rate" r and a "bucket size" b. This can be achieved by configuring regulators with parameter (r, b) and states per flow on each node, however, the cost is too high. Another more feasible way is to carry the states in the packet, and the scheduling mechanism automatically regulate and sorts the packet based on the states read from the packet.

There are some common states that are used by multiple scheduling mechanisms. For example, the latency deviation (E) defined in EDF [I-D.ietf-detnet-deadline-based-forwarding] or TQF [I-D.ietf-detnet-packet-timeslot-mechanism], and the damping delay defind in gLBF [I-D.ietf-detnet-glbf], are actually the same thing and can be considered as the latency compensation used for the forwarded path. Another example is that the endpoint damping delay defined in [I-D.peng-detnet-policing-jitter-control] that can be combined with any on-time scheduling mechanisms to further avoid jitter caused by policing.

This document introduces new IPv6 options for HBH or DOH Options header, to carry common scheduling parameters for deterministic forwarding. Note that the motivation of the common scheduling parameters defined in this document is to be shared and used by multiple scheduling mechanisms, rather than a container that includes some different parameters.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Path Latency Deviation Option

The path latency deviation is used to characterize the deviation between the delay budget (such as a planned residence time or estimated worst-case delay) and the actual delay of a packet at each hop. Each hop along the path can use this information to shape or sort arrived packets, to ensure that the flow conforms to predefined constraints. Examples include the latency deviation (E) defined in EDF or TQF mechanism, and damping delay defined in gLBF mechanism.

Strictly speaking, the path latency deviation should be a cumulative value, that is, it accumulates the latency deviation of all upstream nodes. However, if the scheduling mechanism used in the network can ensure that the path latency deviation is cleared and reset at each hop, then this cumulative value actually only includes the latency deviation generated by a single hop and updated again on the next hop.

The scope of use of the path latency deviation is the forwarded path. It is generally included in HBH to eliminate jitter at each hop (such as each transit node in EDF or gLBF technology domain). It may also be included in DOH, eliminating jitter only at the last hop (such as the egress node in TQF technology domain).

The path latency deviation option has the following format:

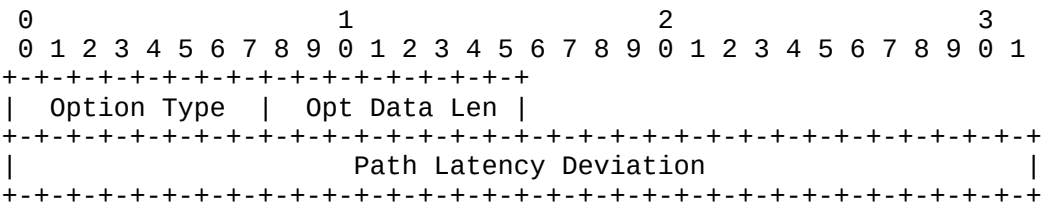

Figure 1: Path Latency Deviation Option

Option Type: 8-bit identifier of the type of option. Value TBD by IANA; the highest-order 3 bits of thie field is 001 to skip over this option and continue processing the header if the processing IPv6 node does not recognize the Option Type and to permit the Option Data to be changed en route to the packet’s final destination.

Opt Data Len: 8-bit unsigned integer. Length of the Option Data field of this option, in octets. It is set to 4.

Path Latency Deviation: 32-bit signed integer, represents the the deviation between delay budget and actual delay on each hop (such as EDF, gLBF), or the deviation between the ideal arrival/departure time and the actual arrival/departure time (such as TQF). It is the holding time imposed on the downstream node before the packet is further scheduled.

3. endpoint damping delay Option

The endpoint damping delay is used to characterize the necessary holding time of the packet on the endpoint of the path, i.e., the final destination, to avoid jitter caused by policing delay. Please refer to [I-D.peng-detnet-policing-jitter-control] for more details.

The scope of use of the endpoint damping delay is the final destination of the egress domain. It is recommoned to be included in DOH.

The endpoint damping delay option has the following format:

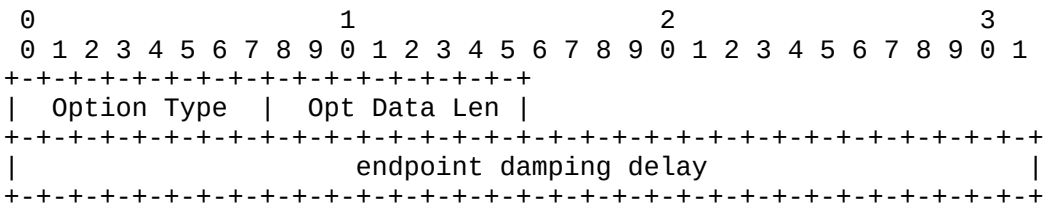

Figure 2: endpoint damping delay Option

Option Type: 8-bit identifier of the type of option. Value TBD by IANA; the highest-order 3 bits of thie field is 001 to skip over this option and continue processing the header if the processing IPv6 node does not recognize the Option Type and to permit the Option Data to be changed en route to the packet's final destination.

Opt Data Len: 8-bit unsigned integer. Length of the Option Data field of this option, in octets. It is set to 4.

endpoint damping delay: 32-bit signed integer, represents the the holding time imposed on the endpoint before the packet is delivered to the application destination. Note that the holding time may also be imposed on the downstream node.

4. Process of Path Latency Deviation Option

For a specific queueing technology domain that guarantee low jitter, the ingress node of the domain will encapsulte outer IPv6 header for the application flow. There are metadata related with the queueing technology contained in the outer IPv6 header. The ingress node may check the latency deviation between the delay budget and the actual delay expierenced at this node, and put a path latency deviation optioin in HBH or DOH within the outer IPv6 header. For those mechanisms that use the path latency deviation to eliminate jitter at transit nodes, the option is included in HBH. For mechanisms that do not require path latency deviation to eliminate jitter at transit nodes, the option is included in DOH.

For a transit node that received an IPv6 packet with HBH and the path latency deviation option, it should firstly impose the corresponding holding time for that packet. If the packet has undergone the holding time, the path latency deviation option is reset to zero; otherwise, the path latency deviation option remains unchanged. Then the packet continue to be shceduled at this node. The transit node will also check the latency deviation between the delay budget and the actual delay expierenced at this node, and adds this latency deviation to the path latency deviation optioin in HBH and then sends to the next hop.

For the egress node of the domain that received an IPv6 packet with HBH or DOH and the path latency deviation option, it should firstly impose the corresponding holding time for that packet. If the packet has undergone the holding time, the path latency deviation option is reset to zero; otherwise, the path latency deviation option remains unchanged. Then the packet continue to be shceduled at this node. The egress node will also check the latency deviation between the delay budget and the actual delay expierenced at this node, and adds this latency deviation to the path latency deviation optioin in HBH or DOH. The outer IPv6 header is removed and the remaining latency devication may be informed to the next hop.

5. Process of Endpoint Damping Delay Option

The ingress node of the ingress domain that received an application packet may calculate the the endpoint damping delay for the packet, that equals to edge-to-edge policing delay budget minus the headend runtime policing delay. An outer IPv6 header corresponding to end-to-end path is encapsulted and contains DOH with endpoint damping delay option. Note that the DOH is closed to the upper-layer header.

The egress node of the egress domain that received an IPv6 packet with DOH and the endpoint damping delay option, should impose the corresponding holding time for that packet. This may be happened after the packet is shceduled at this node, or the remaining damping delay may be informed to the next hop.

6. IANA Considerations

This document updates the "Destination Options and Hop-by-Hop Options" under the "Internet Protocol Version 6 (IPv6) Parameters" registry:

Hex Value	act	chg	rest	Description	Reference
TBD1	00	1	00000	Path Latency Deviation	This document
TBD2	00	1	00000	Endpoint Damping Delay	This document

7. Security Considerations

TBD

8. Acknowledgements

TBD

9. Normative References

[I-D.ietf-detnet-deadline-based-forwarding]
Peng, S., Du, Z., Basu, K., cheng, C., Yang, D., and C. Liu, "Deadline Based Deterministic Forwarding", Work in Progress, Internet-Draft, draft-ietf-detnet-deadline-based-forwarding-00, 16 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-deadline-based-forwarding-00>.

[I-D.ietf-detnet-glbf]
Eckert, T. T., Clemm, A., Bryant, S., and S. Hommes, "Deterministic Networking (DetNet) Data Plane - guaranteed Latency Based Forwarding (gLBF) for bounded latency with low jitter and asynchronous forwarding in Deterministic Networks", Work in Progress, Internet-Draft, draft-ietf-detnet-glbf-00, 16 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-glbf-00>.

[I-D.ietf-detnet-packet-timeslot-mechanism]
Peng, S., Liu, P., Basu, K., Liu, A., Yang, D., Peng, G., and J. Zhao, "Timeslot Queueing and Forwarding Mechanism", Work in Progress, Internet-Draft, draft-ietf-detnet-packet-timeslot-mechanism-00, 16 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-packet-timeslot-mechanism-00>.

[I-D.peng-detnet-policing-jitter-control]
Peng, S., Liu, P., and K. Basu, "Mechanism to control jitter caused by policing in Detnet", Work in Progress, Internet-Draft, draft-peng-detnet-policing-jitter-control-01, 8 October 2024, <https://datatracker.ietf.org/doc/html/draft-peng-detnet-policing-jitter-control-01>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

Author's Address

Shaofu Peng
ZTE Corporation
China
Email: peng.shaofu@zte.com.cn