Network
Internet-Draft
Intended status: Standards Track
Expires: 24 July 2026

Shaofu. Peng
ZTE Corporation

20 January 2026

Control Word Option
draft-peng-6man-cw-option-01

Abstract

This document introduces new IPv6 options for DOH, to carry flow identifier, sequence number, and other customer service mapped information that is encapsulated by the provider network, to support flow-specific treatment, such as statistics, monitoring, QoS, redundancy elimination and reordering, etc.

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 24 July 2026.

Table of Contents

1. Introduction

[RFC4385] defines Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for use over an MPLS packet switched network (PSN). It explicitly indicate that the payload behind the MPLS label stack is non-IP, to avoid intermediate nodes always treating payload as IP payload, e.g., in the case of hash function in a load distribution scheme. PW MPLS Control Word (PWMCW) is used to encapsulate PW data packets. PW label is used for flow identification. PWMCW includes Sequence Number field for out of order checking and reordering functions, that is suitable for circuits sensitive to packet out of order, such as Time Division Multiplexed (TDM) circuits. [RFC8964] also defines DetNet Control Word (d-CW) in MPLS data plane. S-label is used for flow identification. d-CW includes Sequence Number field for out of order checking and reordering functions for DetNet flows. The reason for out of order is multi-path transmission, which may be intentional path planning or forced path switching during network failures.

Some provider networks are migrating from MPLS to IPv6. Customer services (including out of order sensitive services) will be uniformly encapsulated in IPv6. The customer services may not be aware of this migrating. However, the service requirement should be smoothly met. The current IPv6 standards lack a unified encapsulation method for the identification and sequence number of original customer flows. Although, an IPv6 flow can be typically identified by 5-tuple (source address, destination address, source port, destination port, and the transport protocol type), some of these fields may be unavailable due to either fragmentation or encryption, or locating them past a chain of IPv6 extension headers may be inefficient. [RFC6437] defines Flow Label, which, combined with Source Address and Destination Address fields, is a more efficient IPv6 flow classification. However, the purpose of flow classification is mainly used for load distribution, instead of a discriminator for the original customer flows. Multiple customer

flows may be encapsulated in outer IPv6 packets that use the same Flow Label, regardless of the Flow Label of the encapsulated packets. [RFC9343] defines the AltMark option, which includes FlowMonID field to identify the monitored flow, but without enough space to define the sequence number. [RFC9566] defines MPLS based d-CW over IPv6, which is high cost and requires the IPv6 data plane to support additional MPLS forwarding logic beyond pure IPv6 forwarding.

This document defines Conctrol Word (CW) option in Destination Options Header that includes flow identity, sequence number, and other customer service mapped information to facilitate support for flow-specific treatment, such as statistics, monitoring, QoS, redundancy elimination and reordering, etc. The CW option is not intended to be used across the global Internet, but within a limited domain [RFC8799].

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. Control Word Option

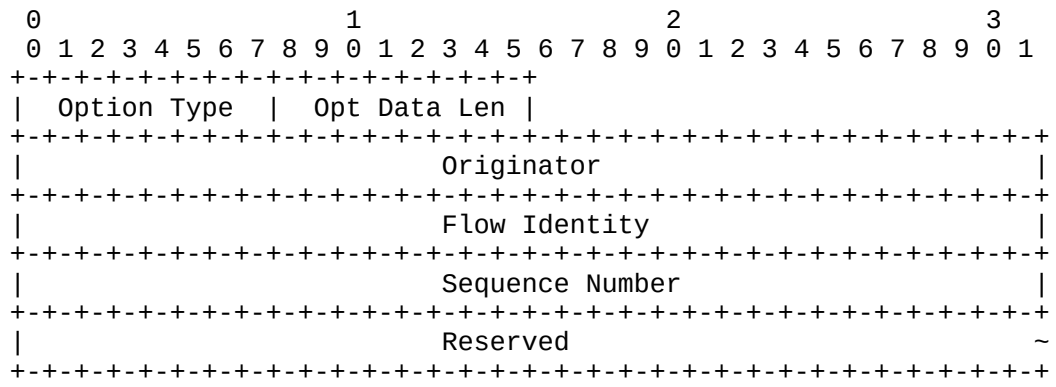

Figure 1

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 variable, may set to 12, or other larger values if the Reserved field has been defined in future.

Originator: 32-bit identifier of the originator that specify control word for the customer flow. In general, the originator is the flow entrance node. Note that some intentionally defined forwarding methods may frequently remove and add IPv6 header, resulting in the Source Address field no longer containing the original source address (i.e., the address of the flow entrance node).

Flow Identity: 32-bit identifier of the customer flow, allocated by the originator. It is used to mark packets of a given flow. The value of zero is to indicate unmarked packets.

Sequence Number: 32-bit unsigned integer, represents the sequence number of a packet in a flow, increasing by 1 with each newly sent packet of the same flow. The circular unsigned 32-bit number space excludes the value zero.

Reserved: If Opt Data Len is set to 12, the Reserved field does not exist. The actual length of field Reserved is equal to Opt Data Len minus 12.

3. Encapsulation of CW Options On Ingress Node

The flow entrance node, when encapsulating the customer flow with an outer IPv6 header, can explicitly insert a DOH contains CW option in the outer IPv6 header according to the flow states. The DOH must be inserted before the Routing Header (RH), if RH also needs to be inserted.

The flow entrance node can use local algorithm to assign different flow identities to different customer flows. The algorithm can check the 5-tuple of the customer flow to ensure that the generated flow identity value has local uniqueness. Although flow aggregation can map multiple flows to the same traffic class, it is still recommended to assign different flow identities to these member flows.

For a given cutomer flow, the sequence number assigned to the first received packet is 1. For each new packet received, the sequence number increases by 1, until it reaches the maximum value and then cycles back to 1. For consecutive packets of a given flow, their sequence number must be continuous.

The Originator field is set to the unique ID of the flow entrance node within the network.

For MPLS and SRv6 interworking case, the border node should copy Control Word information from the receiving header to the sending header, e.g, from MPLS CW to IPv6 CW.

4. Operations of CW Options On Destination

When the packet reaches the node identified in the Destination Address field of the outer IPv6 header, CW option is read and used for flow-specific treatment, such as packet elimination and reordering. The destination node may be each segment of Routing Header (RH) or final destination. How to config flow-specific treatment on the destinaiton node and trigger this treatment is out the scope of this document. Note that some processing may need flow states maintained on the node.

The content of CW option must not be modified en route. If the outer IPv6 header is not removed, the DOH with CW option is also not removed. Some intentionally defined forwarding methods may frequently remove and add outer IPv6 header en route, in this case the DOH with CW option should also be removed and added. If there are further outer IPv6 header encapsulated on the outer IPv6 header, e.g., an underlay traffic engineering path, the DOH with CW option is generally not necessary to copy to the further outer IPv6 header, since the flow-specific treatment is not usually configured on nodes along the underlay traffic engineering path to avoid too many flow states on intermediate nodes.

5. 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
TBD	00	1	00000	Control Word Option	This document

6. Security Considerations

Malicious tampering with any field in the CW option can result in corresponding processing exceptions. For example, tampering with the Flow Identity may result in packets being incorrectly identified as unrelated flows, and tampering with the Sequence Number may result in packets being mistakenly discarded as redundant packets. Therefore, the CW option must only be applied in a limited domain, where all nodes in the domain are trusted. Other common security considerations are described in [RFC8200], [RFC9099].

7. Acknowledgements

The authors would like to thank Brian Carpenter for his review and valuable comments.

8. Normative References

[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>.

[RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <https://www.rfc-editor.org/info/rfc4385>.

[RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <https://www.rfc-editor.org/info/rfc6437>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <https://www.rfc-editor.org/info/rfc8799>.

[RFC8964] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: MPLS", RFC 8964, DOI 10.17487/RFC8964, January 2021, <https://www.rfc-editor.org/info/rfc8964>.

[RFC9099] Vyncke, ., Chittimaneni, K., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", RFC 9099, DOI 10.17487/RFC9099, August 2021, <https://www.rfc-editor.org/info/rfc9099>.

[RFC9343] Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate-Marking Method", RFC 9343, DOI 10.17487/RFC9343, December 2022, <https://www.rfc-editor.org/info/rfc9343>.

[RFC9566] Varga, B., Farkas, J., and A. Malis, "Deterministic Networking (DetNet) Packet Replication, Elimination, and Ordering Functions (PREOF) via MPLS over UDP/IP", RFC 9566, DOI 10.17487/RFC9566, April 2024, <https://www.rfc-editor.org/info/rfc9566>.

Author’s Address

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