

Internet Engineering Task Force (IETF)
Request for Comments: 9566
Category: Informational
ISSN: 2070-1721

B. Varga
J. Farkas
Ericsson
A. Malis
Malis Consulting
April 2024

Deterministic Networking (DetNet) Packet Replication, Elimination, and Ordering Functions (PREOF) via MPLS over UDP/IP

Abstract

This document describes how the DetNet IP data plane can support the Packet Replication, Elimination, and Ordering Functions (PREOF) built on the existing MPLS PREOF solution defined for the DetNet MPLS data plane and the mechanisms defined by MPLS-over-UDP technology.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9566>.

Copyright Notice

Copyright (c) 2024 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.


Table of Contents

1. Introduction
2. Terminology
 - 2.1. Terms Used in This Document
 - 2.2. Abbreviations
3. Requirements for Adding PREOF to DetNet IP
4. Adding PREOF to DetNet IP
 - 4.1. Solution Basics
 - 4.2. Encapsulation
 - 4.3. Packet Processing
 - 4.4. Flow Aggregation
 - 4.5. PREOF Processing

4.6.	PREOF-Capable DetNet IP Domain
5.	Control and Management Plane Parameters
6.	Security Considerations
7.	IANA Considerations
8.	References
8.1.	Normative References
8.2.	Informative References
	Acknowledgements
	Authors' Addresses

1. Introduction

The DetNet Working Group has defined Packet Replication (PRF), Packet Elimination (PEF), and Packet Ordering (POF) Functions (represented as PREOF) to provide service protection by the DetNet service sub-layer [RFC8655]. The PREOF service protection method relies on copies of the same packet sent over multiple maximally disjoint paths and uses sequencing information to eliminate duplicates. A possible implementation of PRF and PEF is described in [IEEE8021CB], and the related YANG data model is defined in [IEEE8021CBcv]. A possible implementation of POF is described in [RFC9550]. Figure 1 shows a DetNet flow on which PREOF are applied during forwarding from the source to the destination.



R: Replication Function (PRF)
 E: Elimination Function (PEF)
 O: Ordering Function (POF)

Figure 1: PREOF Scenario in a DetNet Network

In general, the use of PREOF require sequencing information to be included in the packets of a DetNet compound flow. This can be done by adding a sequence number or timestamp as part of DetNet encapsulation. Sequencing information is typically added once, at or close to the source.

The DetNet MPLS data plane [RFC8964] specifies how sequencing information is encoded in the MPLS header. However, the DetNet IP data plane described in [RFC8939] does not specify how sequencing information can be encoded in the IP packet. This document provides sequencing information to DetNet IP nodes, so it results in an improved version of the DetNet IP data plane. As suggested by [RFC8938], the solution uses existing standardized headers and encapsulations. The improvement is achieved by reusing the DetNet MPLS-over-UDP/IP data plane [RFC9025] with the restriction of using zero F-Labels.

2. Terminology

2.1. Terms Used in This Document

This document uses the terminology established in the DetNet architecture [RFC8655], and it is assumed that the reader is familiar with that document and its terminology.

2.2. Abbreviations

The following abbreviations are used in this document:

DetNet	Deterministic Networking
PEF	Packet Elimination Function
POF	Packet Ordering Function
PREOF	Packet Replication, Elimination, and Ordering Functions
PRF	Packet Replication Function

3. Requirements for Adding PREOF to DetNet IP

The requirements for adding PREOF to DetNet IP are:

- * to reuse existing DetNet data plane solutions (e.g., [RFC8964], [RFC9025]), and
- * to allow the DetNet service sub-layer for IP packet-switched networks with minimal implementation effort.

The described solution leverages MPLS header fields without requiring the support of the MPLS forwarding plane.

4. Adding PREOF to DetNet IP

4.1. Solution Basics

The DetNet IP encapsulation supporting the DetNet service sub-layer is based on the "UDP tunneling" concept. The solution creates a set of underlay UDP/IP tunnels between an overlay set of DetNet relay nodes.

At the edge of a PREOF-capable DetNet IP domain, the DetNet flow is encapsulated in a UDP packet containing the sequence number used by PREOF within the domain. This solution maintains the 6-tuple-based DetNet flow identification in DetNet transit nodes, which operate at the DetNet forwarding sub-layer between the DetNet service sub-layer nodes; therefore, it is compatible with [RFC8939]. Figure 2 shows how the PREOF-capable DetNet IP data plane fits into the DetNet sub-layers.

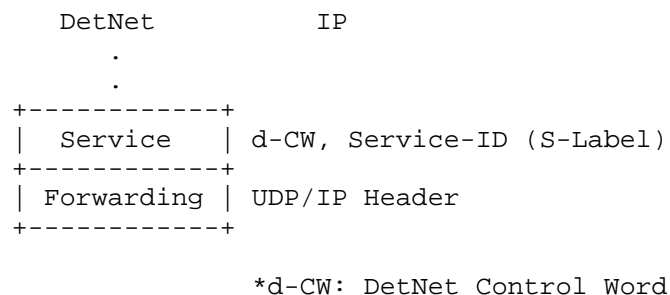


Figure 2: PREOF-Capable DetNet IP Data Plane

4.2. Encapsulation

The PREOF-capable DetNet IP encapsulation builds on encapsulating DetNet pseudowire (PW) directly over UDP. That is, it combines DetNet MPLS [RFC8964] with DetNet MPLS-in-UDP [RFC9025], without using any F-Labels, as shown in Figure 3. DetNet flows are identified at the receiving DetNet service sub-layer processing node via the S-Label and/or the UDP/IP header information. Sequencing information for PREOF is provided by the DetNet Control Word (d-CW) per [RFC8964]. The S-Label is used to identify both the DetNet flow and the DetNet App-flow type. The UDP tunnel is used to direct the

packet across the DetNet domain to the next DetNet service sub-layer processing node.

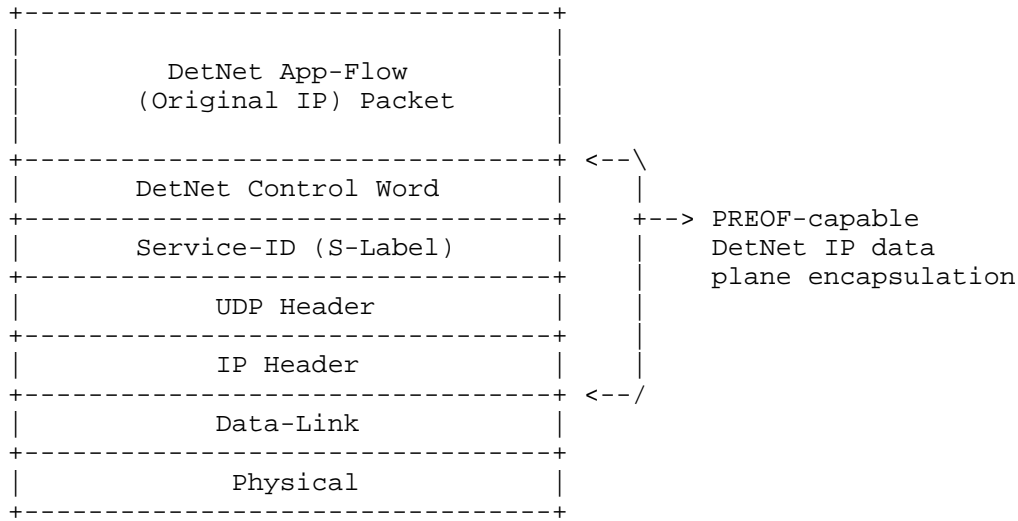


Figure 3: PREOF-Capable DetNet IP Encapsulation

4.3. Packet Processing

IP ingress and egress nodes of the PREOF-capable DetNet IP domain add and remove a DetNet service-specific d-CW and Service-ID (i.e., S-Label). Relay nodes can change Service-ID values when processing a DetNet flow, i.e., incoming and outgoing Service-IDs of a DetNet flow can be different. Service-ID values are provisioned per DetNet service via configuration, e.g., via the Controller Plane described in [RFC8938]. In some PREOF topologies, the node performing replication sends the packets to multiple nodes performing, e.g., PEF or POF, and the replication node can use different Service-ID values for the different member flows for the same DetNet service.

Note that the Service-ID is a local ID on the receiver side that identifies the DetNet flow at the downstream DetNet service sub-layer receiver.


4.4. Flow Aggregation

Two methods can be used for flow aggregation:

- * aggregation using same UDP tunnel, and
- * aggregation of DetNet flows as a new DetNet flow.

In the first method, the different DetNet pseudowires use the same UDP tunnel, so they are treated as a single (aggregated) flow at the forwarding sub-layer. At the service sub-layer, each flow uses a different Service-ID (see Figure 3).

For the second method, an additional hierarchy is created by adding an additional Service-ID and d-CW tuple to the encapsulation. The Aggregate-ID is a special case of a Service-ID, whose properties are known only at the aggregation and deaggregation end points. It is a property of the Aggregate-ID that it is followed by a d-CW followed by a Service-ID/d-CW tuple. Figure 4 shows the encapsulation in the case of aggregation.



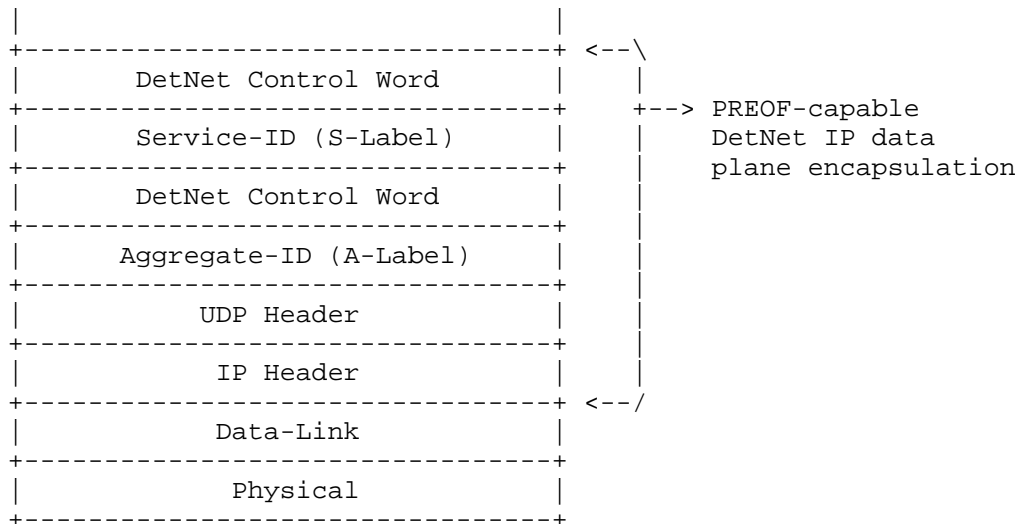


Figure 4: Aggregating DetNet Flows as a New DetNet Flow

The aggregation method is configured in the aggregation/deaggregation nodes.

If several DetNet flows are aggregated in a single UDP tunnel, they all need to follow the same path in the network.

4.5. PREOF Processing

A node operating on a received DetNet flow at the DetNet service sub-layer uses the local context associated with a received Service-ID to determine which local DetNet operation(s) are applied to the received packet. A unique Service-ID can be allocated and can be used to identify a DetNet flow regardless of which input interface or UDP tunnel receives the packet. It is important to note that Service-ID values are driven by the receiver, not the sender.


The DetNet forwarding sub-layer is supported by the UDP tunnel and is responsible for providing resource allocation and explicit routes.

The outgoing PREOF encapsulation and processing can be implemented via the provisioning of UDP and IP header information. Note, when PRF is performed at the DetNet service sub-layer, there are multiple member flows, and each member flow requires its own Service-ID, UDP header information, and IP header information. The headers for each outgoing packet are formatted according to the configuration information, and the UDP Source Port value is set to uniquely identify the DetNet flow. The packet is then handled as a PREOF-capable DetNet IP packet.

The incoming PREOF processing can be implemented by assigning a Service-ID to the received DetNet flow and processing the information in the UDP and IP headers. The provisioned information is used to identify incoming App-flows based on the combination of Service-ID and/or incoming encapsulation header information.

4.6. PREOF-Capable DetNet IP Domain

Figure 5 shows using PREOF in a PREOF-capable DetNet IP network, where service protection is provided end to end, and not only within sub-networks, as is depicted in Figure 4 <<https://www.rfc-editor.org/rfc/rfc8939#figure-4>> of [RFC8939].






Figure 5: PREOF-Capable DetNet IP Domain

5. Control and Management Plane Parameters

The information needed to identify individual and aggregated DetNet flows is summarized as follows:

- * Service-ID information to be mapped to UDP/IP flows. Note that, for example, a single Service-ID can map to multiple sets of UDP/IP information when PREOF is used.
- * IPv4 or IPv6 Source Address field.
- * IPv4 or IPv6 source address prefix length, where a zero (0) value effectively means that the address field is ignored.
- * IPv4 or IPv6 Destination Address field.
- * IPv4 or IPv6 destination address prefix length, where a zero (0) effectively means that the address field is ignored.
- * IPv6 Flow Label field.
- * IPv4 Protocol field being equal to "UDP".
- * IPv6 (last) Next Header field being equal to "UDP".
- * For the IPv4 Type of Service and IPv6 Traffic Class fields:
 - Whether or not the Differentiated Services Code Point (DSCP) field is used in flow identification, as the use of the DSCP field for flow identification is optional.
 - If the DSCP field is used to identify a flow, then the flow identification information (for that flow) includes a list of DSCPs used by the given DetNet flow.
- * UDP Source Port. Support for both exact and wildcard matching is required. Port ranges can optionally be used.
- * UDP Destination Port. Support for both exact and wildcard matching is required. Port ranges can optionally be used.
- * For end systems, an optional maximum IP packet size that should be used for that outgoing DetNet IP flow.

This information is provisioned per DetNet flow via configuration, e.g., via the Controller Plane.

Ordering of the set of information used to identify an individual

DetNet flow can, for example, be used to provide a DetNet service for a specific UDP flow, with unique Source and Destination Port field values, while providing a different service for the aggregate of all other flows with that same UDP Destination Port value.

The minimum set of information for the configuration of the DetNet service sub-layer is summarized as follows:

- * App-flow identification information
- * Sequence number length
- * Type of PREOF to be executed on the DetNet flow
- * Service-ID(s) used by the member flows
- * Associated forwarding sub-layer information
- * Service aggregation information

The minimum set of information for the configuration of the DetNet forwarding sub-layer is summarized as follows:

- * UDP tunnel-specific information
- * Traffic parameters

These parameters are defined in the DetNet flow and service information model [RFC9016] and the DetNet YANG model.

Note: this document focuses on the use of MPLS-over-UDP/IP encapsulation throughout an entire DetNet IP network, making MPLS-based DetNet Operations, Administration, and Maintenance (OAM) techniques applicable [RFC9546]. Using the described encapsulation only for a portion of a DetNet IP network that handles PREOF would complicate OAM.

6. Security Considerations

There are no new DetNet-related security considerations introduced by this solution. Security considerations of DetNet MPLS [RFC8964] and DetNet MPLS over UDP/IP [RFC9025] apply.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. Normative References

- [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>>.
- [RFC8938] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <<https://www.rfc-editor.org/info/rfc8938>>.
- [RFC8939] Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939, November 2020, <<https://www.rfc-editor.org/info/rfc8939>>.

- [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>>.
- [RFC9016] Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "Flow and Service Information Model for Deterministic Networking (DetNet)", RFC 9016, DOI 10.17487/RFC9016, March 2021, <<https://www.rfc-editor.org/info/rfc9016>>.
- [RFC9025] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: MPLS over UDP/IP", RFC 9025, DOI 10.17487/RFC9025, April 2021, <<https://www.rfc-editor.org/info/rfc9025>>.
- [RFC9546] Mirsky, G., Chen, M., and B. Varga, "Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet) with the MPLS Data Plane", RFC 9546, DOI 10.17487/RFC9546, February 2024, <<https://www.rfc-editor.org/info/rfc9546>>.

8.2. Informative References

- [IEEE8021CB]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Frame Replication and Elimination for Reliability", IEEE Std 802.1CB-2017, DOI 10.1109/IEEESTD.2017.8091139, October 2017, <<https://doi.org/10.1109/IEEESTD.2017.8091139>>.
- [IEEE8021CBcv]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Frame Replication and Elimination for Reliability - Amendment 1: Information Model, YANG Data Model, and Management Information Base Module", Amendment to IEEE Std 802.1CB-2017, IEEE Std 802.1CBcv-2021, DOI 10.1109/IEEESTD.2022.9715061, February 2022, <<https://doi.org/10.1109/IEEESTD.2022.9715061>>.
- [RFC9550] Varga, B., Ed., Farkas, J., Kehrler, S., and T. Heer, "Deterministic Networking (DetNet): Packet Ordering Function", RFC 9550, DOI 10.17487/RFC9550, March 2024, <<https://www.rfc-editor.org/info/rfc9550>>.

Acknowledgements

Authors extend their appreciation to Stewart Bryant, Pascal Thubert, David Black, Shirley Yangfan, and Greg Mirsky for their insightful comments and productive discussion that helped to improve the document.

Authors' Addresses

Balazs Varga
Ericsson
Budapest
Magyar Tudosok krt. 11.
1117
Hungary
Email: balazs.a.varga@ericsson.com

Janos Farkas
Ericsson
Budapest

Magyar Tudosok krt. 11.
1117
Hungary
Email: janos.farkas@ericsson.com

Andrew G. Malis
Malis Consulting
Email: agmalis@gmail.com