

DetNet
Internet-Draft
Intended status: Informational
Expires: 5 August 2022

B. Varga
J. Farkas
Ericsson
A. Malis
Malis Consulting
1 February 2022

Deterministic Networking (DetNet): DetNet PREOF via MPLS over UDP/IP
draft-varga-detnet-ip-preof-02

Abstract

This document describes how DetNet IP data plane can support the Packet Replication, Elimination, and Ordering Functions (PREOF) built on the existing MPLS PREOF solution [RFC8939] and the mechanisms defined in [RFC9025].

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 5 August 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

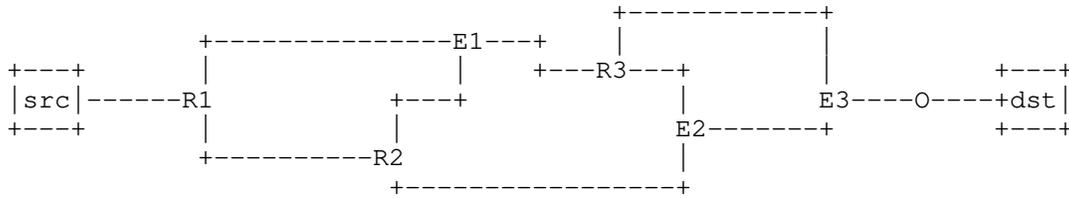
This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this 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
2.	Terminology	3
2.1.	Terms Used in This Document	3
2.2.	Abbreviations	3
2.3.	Requirements Language	4
3.	Requirements for adding PREOF to DetNet IP	4
4.	Adding PREOF to DetNet IP	4
4.1.	Solution Basics	4
4.2.	Encapsulation	5
4.3.	Packet Processing	6
4.4.	Flow Aggregation	6
4.5.	PREOF Procedures	7
4.6.	PREOF capable DetNet IP domain	8
5.	Control and Management Plane Parameters	8
6.	Security Considerations	10
7.	IANA Considerations	10
8.	References	10
8.1.	Normative References	10
8.2.	Informative References	11
	Authors' Addresses	11

1. Introduction

The DetNet Working Group has defined packet replication (PRF), packet elimination (PEF) and packet ordering (POF) functions 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 the PRF and PEF functions is described in [IEEE8021CB] and the related YANG data model is defined in [IEEEP8021CBcv]. A possible implementation of POF function is described in [I-D.varga-detnet-pof]. Figure 1 shows a DetNet flow on which PREOF functions 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 functions require sequencing information to be included in the packets of a DetNet compound flow. This may be done by adding a sequence number or time stamp as part of DetNet encapsulation. Sequencing information is typically added once, at or close to the source.

The DetNet MPLS data plane [RFC8939] 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 header. This document describes a DetNet IP encapsulation that includes sequencing information based on the DetNet MPLS over UDP/IP data plane [RFC9025], i.e., leveraging the MPLS-over-UDP technology.

2. Terminology

2.1. Terms Used in This Document

This document uses the terminology established in the DetNet architecture [RFC8655], and the reader is assumed to be 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.

2.3. 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.

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]).
- * to allow the DetNet service sub-layer for IP packet switched networks with minimal implementation effort.

The described solution practically gains from MPLS header fields without adding MPLS protocol stack complexity to the nodal requirements.

4. Adding PREOF to DetNet IP

4.1. Solution Basics

The DetNet IP encapsulation supporting 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 an UDP packet containing the sequence number used by PREOF functions 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.

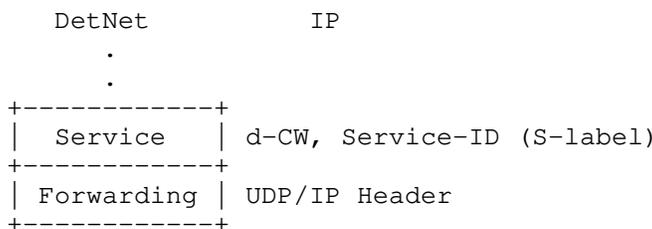


Figure 2: PREOF capable DetNet IP data plane

4.2. Encapsulation

The PREOF capable DetNet IP encapsulation builds on encapsulating DetNet 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) as 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.

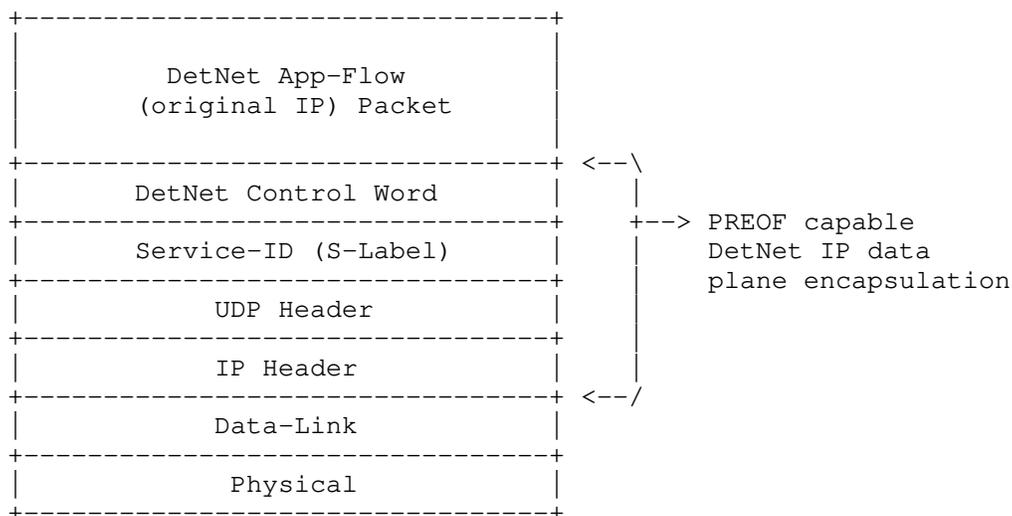


Figure 3: PREOF capable DetNet IP encapsulation

4.3. Packet Processing

IP ingress and egress nodes of the PREOF capable DetNet IP domain MUST add and remove a DetNet service-specific d-CW and Service-ID (i.e., S-Label). Relay nodes MAY 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 MUST be provisioned per DetNet service via configuration, i.e., 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 may need to use different Service-ID values for the different member flows for the same DetNet service.

Note, that Service-IDs provide identification at the downstream DetNet service sub-layer receiver, not the sender.

4.4. Flow Aggregation

Two methods can be used for flow aggregation:

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

In the first case, the different DetNet PWs use the same UDP tunnel, so they are treated as a single (aggregated) flow on all transit nodes.

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

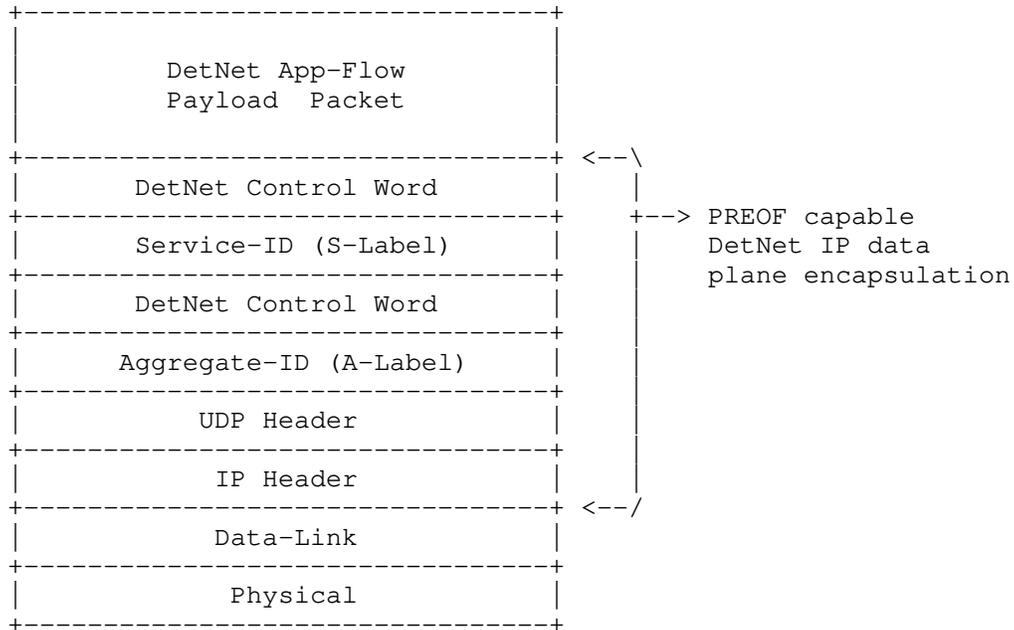


Figure 4: Aggregating DetNet flows as a new DetNet flow

4.5. PREOF Procedures

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 received packet. A Service-ID may be allocated to be unique and enabling DetNet flow identification regardless of which input interface or UDP tunnel the packet is received. 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.

To support outgoing PREOF capable DetNet IP encapsulation, an implementation MUST support 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 their own Service-ID, UDP and IP header information. The headers for each outgoing packet MUST be formatted according to the configuration information, and the UDP Source Port value MUST be set to uniquely identify the DetNet flow. The packet MUST then be handled as a PREOF capable DetNet IP packet.

To support the receive processing, an implementation MUST also support the provisioning of received Service-ID, UDP and IP header information. The provisioned information MUST be 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.

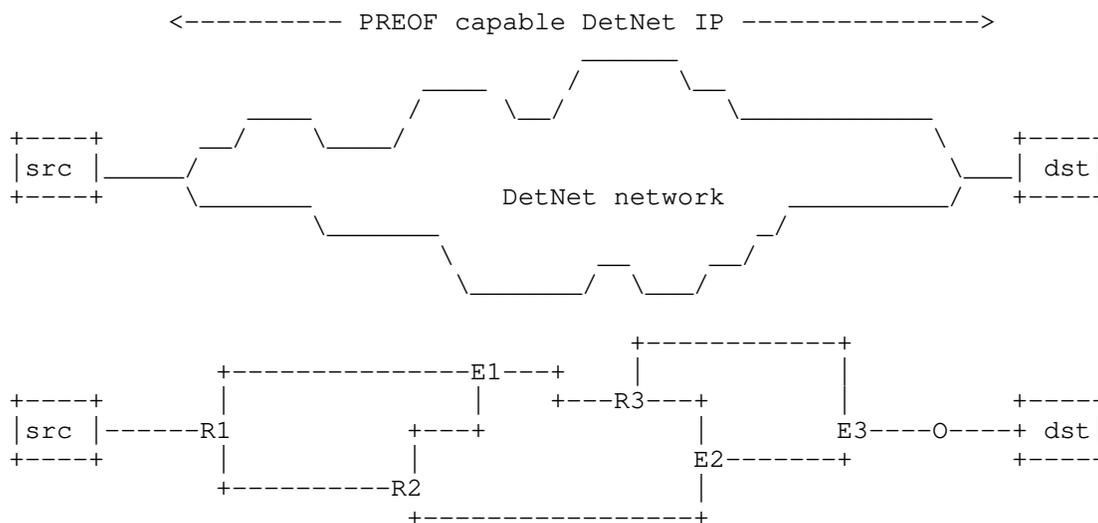


Figure 5: PREOF capable DetNet IP domain

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.
- * IPv4 protocol field set to "UDP".
- * IPv6 next header field set to "UDP".
- * For the IPv4 Type of Service and IPv6 Traffic Class Fields:
 - Whether or not the 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 MUST be provisioned per DetNet flow via configuration, e.g., via the controller plane.

An implementation MUST support ordering of the set of information used to identify an individual DetNet flow. This 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.
- * PREOF + related Service-ID(s).
- * 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.

6. Security Considerations

There are no new DetNet related security considerations introduced by this solution.

7. IANA Considerations

This document makes no IANA requests.

8. References

8.1. 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>>.
- [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>>.
- [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>>.

[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>>.

8.2. Informative References

[I-D.varga-detnet-pof]
Varga, B., Farkas, J., Kehrer, S., and T. Heer,
"Deterministic Networking (DetNet): Packet Ordering
Function", Work in Progress, Internet-Draft, draft-varga-
detnet-pof-02, 22 October 2021,
<<https://www.ietf.org/archive/id/draft-varga-detnet-pof-02.txt>>.

[IEEE8021CB]
IEEE, "IEEE Standard for Local and metropolitan area
networks -- Frame Replication and Elimination for
Reliability", DOI 10.1109/IEEESTD.2017.8091139, October
2017,
<https://standards.ieee.org/standard/802_1CB-2017.html>.

[IEEEP8021CBcv]
Kehrer, S., "FRER YANG Data Model and Management
Information Base Module", IEEE P802.1CBcv
/D1.2 P802.1CBcv, March 2021,
<<https://www.ieee802.org/1/files/private/cv-drafts/d1/802-1CBcv-d1-2.pdf>>.

Authors' Addresses

Balázs Varga
Ericsson
Budapest
Magyar Tudosok krt. 11.
1117
Hungary

Email: balazs.a.varga@ericsson.com

János Farkas
Ericsson
Budapest
Magyar Tudosok krt. 11.
1117
Hungary

Email: janos.farkas@ericsson.com

Andrew G. Malis
Malis Consulting

Email: agmalis@gmail.com