

Internet Engineering Task Force (IETF)
Request for Comments: 9056
Category: Standards Track
ISSN: 2070-1721

B. Varga, Ed.
Ericsson
L. Berger
D. Fedyk
LabN Consulting, L.L.C.
S. Bryant
Futurewei Technologies
J. Korhonen
October 2021

Deterministic Networking (DetNet) Data Plane: IP over MPLS

Abstract

This document specifies the Deterministic Networking data plane when encapsulating IP over an MPLS packet-switched network.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc9056>.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
2. Terminology
 - 2.1. Terms Used in This Document
 - 2.2. Abbreviations
 - 2.3. Requirements Language
3. DetNet IP Data Plane Overview
4. DetNet IP over DetNet MPLS
 - 4.1. DetNet IP over DetNet MPLS Data Plane Scenarios
 - 4.2. DetNet IP over DetNet MPLS Encapsulation
5. DetNet IP over DetNet MPLS Procedures
 - 5.1. DetNet IP over DetNet MPLS Flow Identification and Aggregation Procedures
 - 5.2. DetNet IP over DetNet MPLS Traffic Treatment Procedures
6. Management and Control Information Summary
7. Security Considerations
8. IANA Considerations
9. References
 - 9.1. Normative References

9.2. Informative References
Acknowledgements
Contributors
Authors' Addresses

1. Introduction

Deterministic Networking (DetNet) is a service that can be offered by a network to DetNet flows. DetNet provides a capability for the delivery of data flows with extremely low packet loss rates and bounded end-to-end delivery latency. General background and concepts of DetNet can be found in the DetNet architecture [RFC8655].

This document specifies use of the IP DetNet encapsulation over an MPLS network. It maps the IP data plane encapsulation described in [RFC8939] to the DetNet MPLS data plane defined in [RFC8964].

2. Terminology

2.1. Terms Used in This Document

This document uses the terminology and concepts established in the DetNet architecture [RFC8655] and in [RFC8938]. The reader is assumed to be familiar with these documents and their terminology.

2.2. Abbreviations

This document uses the abbreviations defined in the DetNet architecture [RFC8655] and in [RFC8938]. This document uses the following abbreviations:

| | |
|---------|--|
| CE | Customer Edge (equipment) |
| d-CW | DetNet Control Word |
| DetNet | Deterministic Networking |
| DF | DetNet Flow |
| DN | DetNet |
| L2 | Layer 2 |
| LSP | Label-Switched Path |
| MPLS | Multiprotocol Label Switching |
| PEF | Packet Elimination Function |
| PRF | Packet Replication Function |
| PREOF | Packet Replication, Elimination, and Ordering Functions |
| POF | Packet Ordering Function |
| PW | Pseudowire |
| S-Label | DetNet "service" Label |
| S-PE | Switching Provider Edge |
| T-PE | Terminating Provider Edge |
| TE | Traffic Engineering |
| TSN | Time-Sensitive Networking; TSN is a Task Group of the IEEE 802.1 Working Group |

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. DetNet IP Data Plane Overview

Figure 1 illustrates an IP DetNet with an MPLS-based DetNet network as a sub-network between the relay nodes. An IP flow is mapped to one or more PWs and MPLS (TE) LSPs. The end systems still originate IP-encapsulated traffic, identified as DetNet flows. The relay nodes follow procedures defined in Section 4 to map each DetNet flow to MPLS LSPs. While not shown, relay nodes can provide service sub-layer functions such as PREOF using DetNet over MPLS, and this is indicated by the solid line for the MPLS-facing portion of the Service component. Note that the Transit node is MPLS (TE) LSP aware and performs switching based on MPLS labels; it need not have any specific knowledge of the DetNet service or the corresponding DetNet flow identification. See Section 4 for details on the mapping of IP flows to MPLS, and [RFC8964] for general support of DetNet services using MPLS.

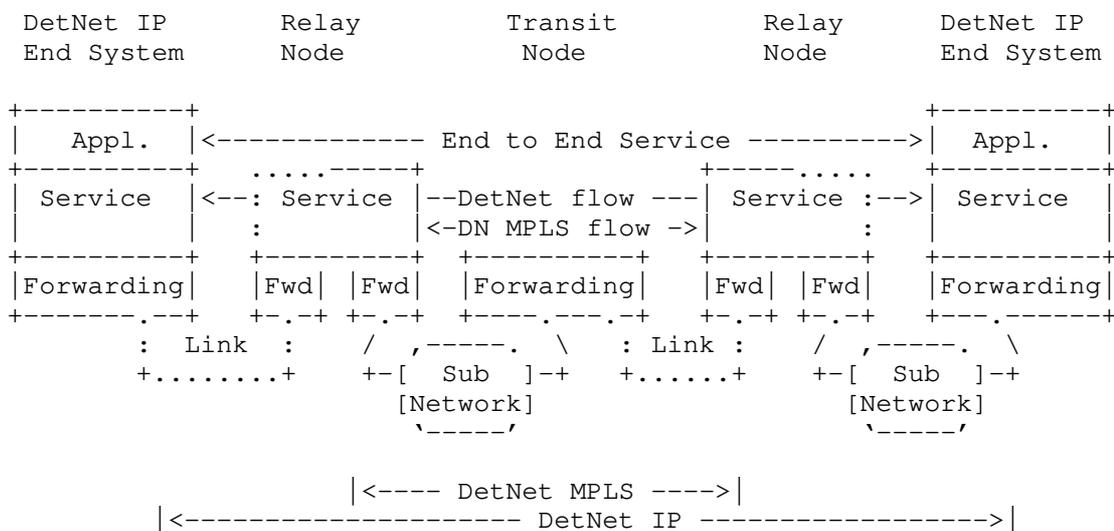


Figure 1: Architecture: DetNet IP over DetNet MPLS Network

4. DetNet IP over DetNet MPLS

This section defines how IP-encapsulated flows are carried over a DetNet MPLS data plane as defined in [RFC8964]. Since both non-DetNet and DetNet IP packets are identical on the wire, this section is applicable to any node that supports IP over DetNet MPLS, and this section refers to both cases as DetNet IP over DetNet MPLS.

4.1. DetNet IP over DetNet MPLS Data Plane Scenarios

An example use of DetNet IP over DetNet MPLS is presented here.

Figure 1 illustrates IP DetNet-enabled End Systems (hosts) connected to DetNet-enabled IP networks (DN IP), operating over a DetNet-aware MPLS network. In this figure, we have a case where the relay nodes act as T-PEs and sit at the boundary of the MPLS domain since the non-MPLS domain is DetNet aware. This case is very similar to the DetNet MPLS Network (Figure 2 in [RFC8964]). However, in Figure 2 of [RFC8964], the T-PEs are located at the end system and MPLS spans the whole DetNet service. The primary difference in this document is that the relay nodes are at the edges of the MPLS domain and therefore function as T-PEs, and that MPLS service sub-layer functions are not provided over the DetNet IP network. The transit node functions shown above are identical to those described in [RFC8964].

Figure 2 illustrates how relay nodes can provide service protection over an MPLS domain. In this case, CE1 and CE2 are IP DetNet end

systems that are interconnected via an MPLS domain such as that described in [RFC8964]. Note that R1 and R3 sit at the edges of an MPLS domain and therefore are similar to T-PEs, while R2 sits in the middle of the domain and is therefore similar to an S-PE.

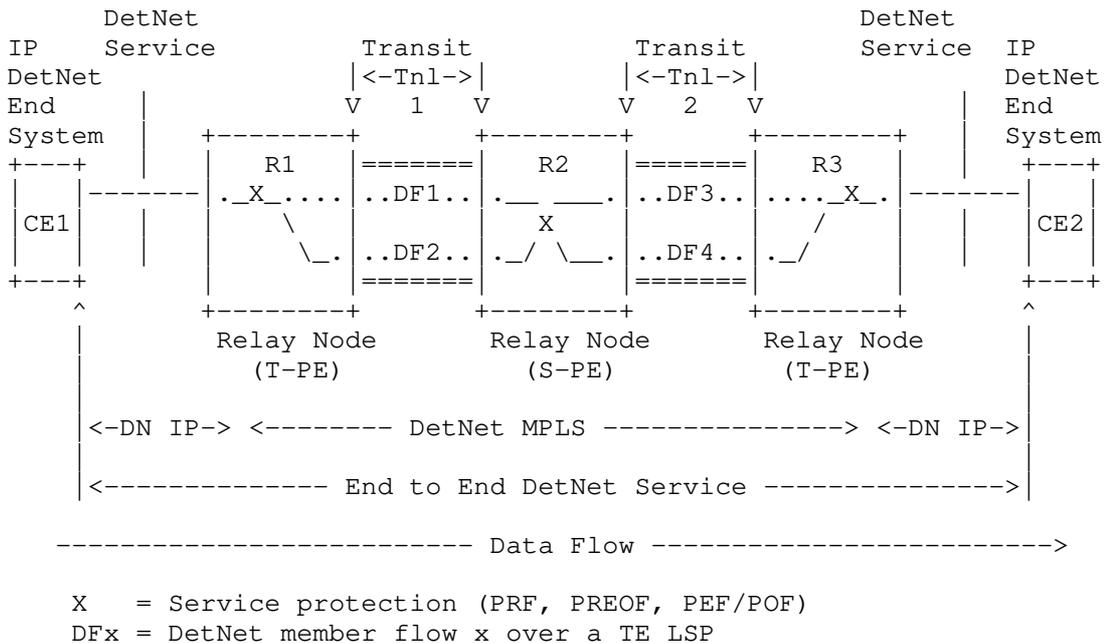


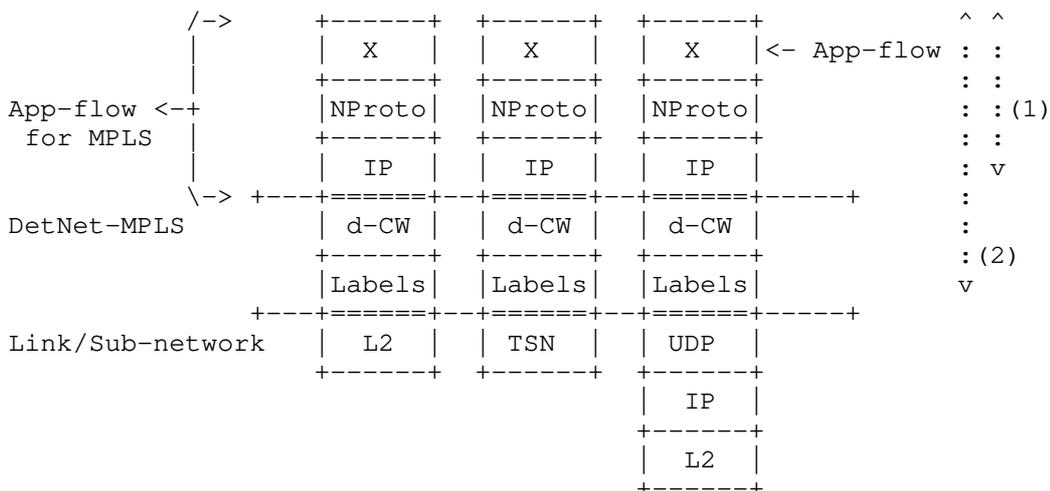
Figure 2: Service Protection over DetNet MPLS Network for DetNet IP

Figure 1 illustrates DetNet-enabled end systems connected to DetNet-enabled (DN) MPLS networks. A similar situation occurs when end systems are not DetNet aware. In this case, edge nodes sit at the boundary of the MPLS domain since it is also a DetNet domain boundary. The edge nodes provide DetNet service proxies for the end applications by initiating and terminating DetNet service for the application's IP flows. While the node types differ, there is essentially no difference in data plane processing between relays and edges. There are likely to be differences in Controller Plane operation, particularly when distributed control plane protocols are used.

It is still possible to provide DetNet service protection for non-DetNet-aware end systems. This case is basically the same as Figure 2, with the exception that CE1 and CE2 are non-DetNet-aware end systems and R1 and R3 become edge nodes.

4.2. DetNet IP over DetNet MPLS Encapsulation

The basic encapsulation approach is to treat a DetNet IP flow as an App-flow from the DetNet MPLS perspective. The corresponding example DetNet Sub-network format is shown in Figure 3.



(1) DetNet IP Flow (or simply IP flow)

(2) DetNet MPLS Flow

Figure 3: Example DetNet IP over MPLS Sub-network Formats

In Figure 3, "App-flow" indicates the payload carried by the DetNet IP data plane. "IP" and "NProto" indicate the fields described in Sections 5.1.1 (IP Header Information) and 5.1.2 (Other Protocol Header Information) of [RFC8939], respectively. "App-flow for MPLS" indicates that an individual DetNet IP flow is the payload from the perspective of the DetNet MPLS data plane defined in [RFC8964].

Per Section 5.1 of [RFC8964], the DetNet MPLS data plane uses a single S-Label to support a single App-flow. DetNet IP Flow Identification Procedures in Section 5.1 of [RFC8939] states that a single DetNet flow is identified based on IP- and next-level protocol header information. Section 4.4 of [RFC8939] (DetNet Flow Aggregation) defines the ways in which aggregation is supported through the use of prefixes, wildcards, lists, and port ranges. Collectively, this results in the fairly straightforward procedures defined in the next section.

As shown in Figure 2, DetNet relay nodes are responsible for the mapping of a DetNet flow, at the service sub-layer, from the IP to MPLS DetNet data planes and back again. Their related DetNet IP over DetNet MPLS data plane operation is comprised of two sets of procedures: the mapping of flow identifiers and ensuring proper traffic treatment.

Mapping of IP to DetNet MPLS is similar for DetNet IP flows and IP flows. The six-tuple of IP is mapped to the S-Label in both cases. The various fields may be mapped or ignored when going from IP to MPLS.

5. DetNet IP over DetNet MPLS Procedures

The main differences of mapping IP to DetNet MPLS (compared to plain MPLS) are that (1) there is a mandatory flow identification to make the forwarding decision (i.e., forwarding is not based on FEC), (2) the d-CW (DetNet Control Word) is mandatory for the MPLS encapsulation, and (3) during forwarding over the DetNet MPLS network, treatment specific to DetNet flows is needed.

5.1. DetNet IP over DetNet MPLS Flow Identification and Aggregation Procedures

A DetNet relay node (ingress T-PE) that sends a DetNet IP flow over a DetNet MPLS network MUST map a DetNet IP flow, as identified in [RFC8939], into a single MPLS DetNet flow and MUST process it in accordance to the procedures defined in [RFC8964]. PRF MAY be supported at the MPLS level for DetNet IP flows sent over a DetNet MPLS network. Aggregation MAY be supported as defined in Section 4.4 of [RFC8964]. Aggregation considerations in [RFC8939] MAY be used to identify an individual DetNet IP flow. The provisioning of the mapping of DetNet IP flows to DetNet MPLS flows MUST be supported via configuration, e.g., via the Controller Plane.

A DetNet relay node (egress T-PE) MAY be provisioned to handle packets received via the DetNet MPLS data plane as DetNet IP flows. A single incoming DetNet MPLS flow MAY be treated as a single DetNet IP flow, without examination of IP headers. Alternatively, packets received via the DetNet MPLS data plane MAY follow the normal DetNet IP flow identification procedures defined in Section 5.1 of [RFC8939].

An implementation MUST support the provisioning for handling any packet flows received via the DetNet MPLS data plane as DetNet IP flows via configuration. Note that such configuration MAY include support from PREOF on the incoming DetNet MPLS flow.

| Note: Using Layer 4 (L4) transport protocols (e.g., for
| multipath) are out of scope of this document both for a single

| flow and aggregate flows.

5.2. DetNet IP over DetNet MPLS Traffic Treatment Procedures

The traffic treatment required for a particular DetNet IP flow is provisioned via configuration or the Controller Plane. When a DetNet IP flow is sent over DetNet MPLS, a DetNet relay node MUST ensure that the provisioned DetNet IP traffic treatment is provided at the forwarding sub-layer as described in Section 5.2 of [RFC8964]. Note that PRF MAY be utilized when sending IP over MPLS.

Traffic treatment for DetNet IP flows received over the DetNet MPLS data plane MUST follow Section 5.3 of [RFC8939] (DetNet IP Traffic Treatment Procedures).

6. Management and Control Information Summary

The following summarizes the set of information that is needed to support DetNet IP over DetNet MPLS at the MPLS ingress node:

- * Each MPLS App-Flow is selected from the incoming IP traffic using the IP flow identification information defined in [RFC8939]. This information is summarized in Section 5.1 of that document and includes all wildcards, port ranges, and the ability to ignore specific IP fields.
- * The DetNet MPLS service that is to be used to send the matching IP traffic. This matching information is provided in Section 5.1 of [RFC8964] and includes both service and traffic delivery information.

The following summarizes the set of information that is needed to support DetNet IP over DetNet MPLS at the MPLS egress node:

- * The S-Label value that identifies the encapsulated App-flow traffic.
- * For each S-Label, how the received traffic is to be handled. The traffic may be processed as any other DetNet IP traffic as defined in this document or in [RFC8939], or the traffic may be directly treated as an MPLS App-flow for additional processing according to [RFC8964].

It is the responsibility of the DetNet Controller Plane to properly provision both flow identification information and the flow-specific resources needed to provide the traffic treatment to meet each flow's service requirements. This applies for aggregated and individual flows.

7. Security Considerations

General security considerations for DetNet are described in detail in [RFC9055]. DetNet MPLS and DetNet IP security considerations equally apply to this document and are described in [RFC8964] and [RFC8939].

Security aspects that are unique to DetNet are those whose aim is to protect the support of specific quality-of-service aspects of DetNet, which are primarily to deliver data flows with extremely low packet loss rates and bounded end-to-end delivery latency.

The primary considerations for the data plane are to maintain integrity of data and delivery of the associated DetNet service traversing the DetNet network. Application flows can be protected through whatever means is provided by the underlying technology. For example, encryption may be used, such as that provided by IPsec [RFC4301] for IP flows and/or by an underlying sub-net using MACsec [IEEE802.1AE-2018] for IP-over-Ethernet (Layer 2) flows.

From a data plane perspective, this document does not add or modify any header information.

At the management and control level, DetNet flows are identified on a per-flow basis, which may provide Controller Plane attackers with additional information about the data flows (when compared to Controller Planes that do not include per-flow identification). This is an inherent property of DetNet, which has security implications that should be considered when determining if DetNet is a suitable technology for any given use case.

To provide uninterrupted availability of the DetNet service, provisions can be made against DoS attacks and delay attacks. To protect against DoS attacks, excess traffic due to malicious or malfunctioning devices can be prevented or mitigated, for example, through the use of existing mechanisms such as policing and shaping applied at the input of a DetNet domain. To prevent DetNet packets from being delayed by an entity external to a DetNet domain, DetNet technology definitions can allow for the mitigation of man-in-the-middle attacks (for example, through use of authentication and authorization of devices within the DetNet domain).

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. 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>>.
- [RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <<https://www.rfc-editor.org/info/rfc9055>>.

9.2. Informative References

- [IEEE802.1AE-2018] IEEE, "IEEE Standard for Local and metropolitan area networks—Media Access Control (MAC) Security", IEEE 802.1AE-2018, DOI 10.1109/IEEESTD.2018.8585421, December 2018, <<https://ieeexplore.ieee.org/document/8585421>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301,

Acknowledgements

The authors wish to thank Pat Thaler, Norman Finn, Loa Andersson, David Black, Rodney Cummings, Ethan Grossman, Tal Mizrahi, David Mozes, Craig Gunther, George Swallow, Yuanlong Jiang, and Carlos J. Bernardos for their various contributions to this work.

Contributors

RFC 7322 limits the number of authors listed on the front page to a maximum of 5. The editor wishes to thank and acknowledge the following authors for contributing text to this document.

János Farkas
Ericsson

Email: janos.farkas@ericsson.com

Andrew G. Malis
Malis Consulting

Email: agmalis@gmail.com

János Farkas contributed substantially to the content of this document.

Authors' Addresses

Balázs Varga (editor)
Ericsson
Budapest
Magyar Tudosok krt. 11.
1117
Hungary

Email: balazs.a.varga@ericsson.com

Lou Berger
LabN Consulting, L.L.C.

Email: lberger@labn.net

Don Fedyk
LabN Consulting, L.L.C.

Email: dfedyk@labn.net

Stewart Bryant
Futurewei Technologies

Email: sb@stewartbryant.com

Jouni Korhonen

Email: jouni.nospam@gmail.com