DetNet
Internet-Draft
Intended status: Standards Track
Expires: April 29, 2020

B. Varga, Ed.
J. Farkas
Ericsson
L. Berger
D. Fedyk
LabN Consulting, L.L.C.
A. Malis
Independent
S. Bryant
Futurewei Technologies
J. Korhonen
October 27, 2019

DetNet Data Plane: IP over MPLS
draft-ietf-detnet-ip-over-mpls-03

Abstract

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

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 April 29, 2020.

Copyright Notice

Table of Contents

1. Introduction

Deterministic Networking (DetNet) is a service that can be offered by a network to DetNet flows. DetNet provides these flows extremely low packet loss rates and assured maximum end-to-end delivery latency. General background and concepts of DetNet can be found in the DetNet Architecture [I-D.ietf-detnet-architecture].

This document specifies use of the IP DetNet encapsulation over an MPLS network. It maps the IP data plane encapsulation described in [I-D.ietf-detnet-ip] to the DetNet MPLS data plane defined in [I-D.ietf-detnet-mpls].

2. Terminology

2.1. Terms Used In This Document

This document uses the terminology and concepts established in the DetNet architecture [I-D.ietf-detnet-architecture] and [I-D.ietf-detnet-data-plane-framework], and 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 [I-D.ietf-detnet-architecture] and [I-D.ietf-detnet-data-plane-framework]. This document uses the following abbreviations:

CE	Customer Edge equipment.
DetNet	Deterministic Networking.
DF	DetNet Flow.
DN	DetNet.
L2	Layer-2.
L3	Layer-3.
LSP	Label-switched path.
MPLS	Multiprotocol Label Switching.
PE	Provider Edge.
PREOF	Packet Replication, Ordering and Elimination Function.
PSN	Packet Switched Network.
PW	Pseudowire.
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. It shows a more complex DetNet enabled IP network where an IP flow is mapped to one or more PWs and MPLS (TE) LSPs. The end systems still originate IP encapsulated traffic that are 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, and 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 [I-D.ietf-detnet-mpls] for general support of DetNet services using MPLS.

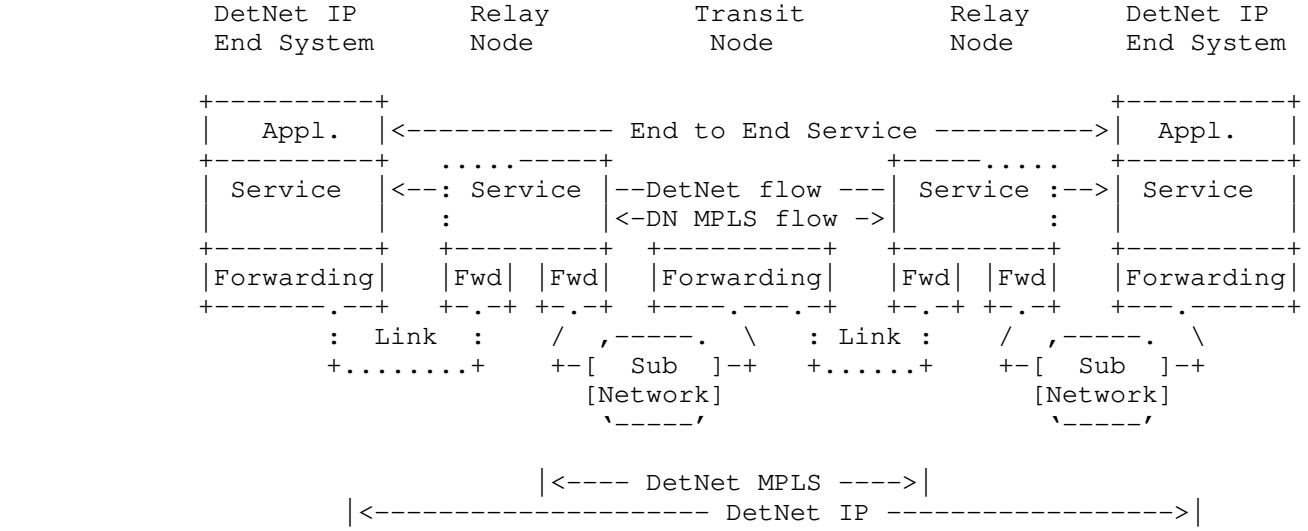

Figure 1: DetNet IP Over DetNet MPLS Network

4. IP over DetNet MPLS

This section defines how IP encapsulated flows are carried over a DetNet MPLS data plane as defined in [I-D.ietf-detnet-mpls]. Since both Non-DetNet and DetNet IP packet 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. IP Over DetNet MPLS Data Plane Scenarios

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

Figure 1 illustrated DetNet enabled End Systems (hosts), connected to DetNet (DN) enabled IP networks, operating over a DetNet aware MPLS network. Using this figure we can 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 [I-D.ietf-detnet-mpls]. However in [I-D.ietf-detnet-mpls] figure 2 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 show above are identical to those described in [I-D.ietf-detnet-mpls].

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 which are interconnected via a MPLS domain such as described in [I-D.ietf-detnet-mpls]. 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.

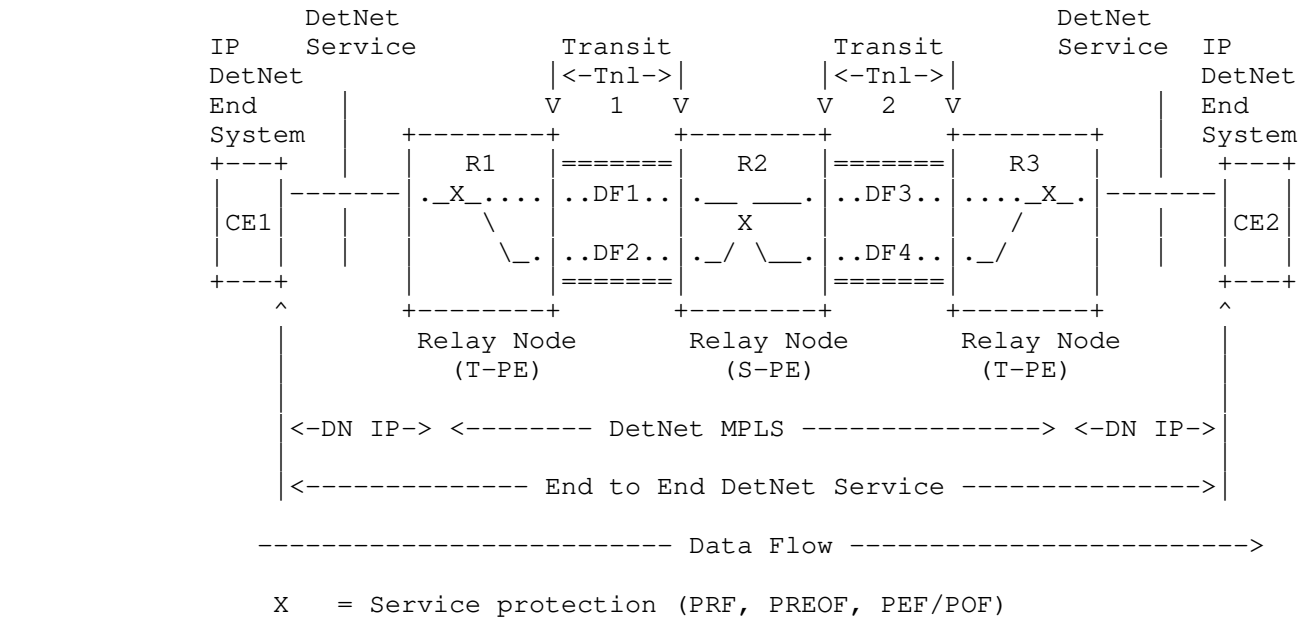

Figure 2: DetNet IP Over DetNet MPLS Network

Figure 1 illustrates DetNet enabled End Systems, connected to DetNet (DN) enabled MPLS network. A similar situation occurs when end systems are 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 relay 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 provided 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.

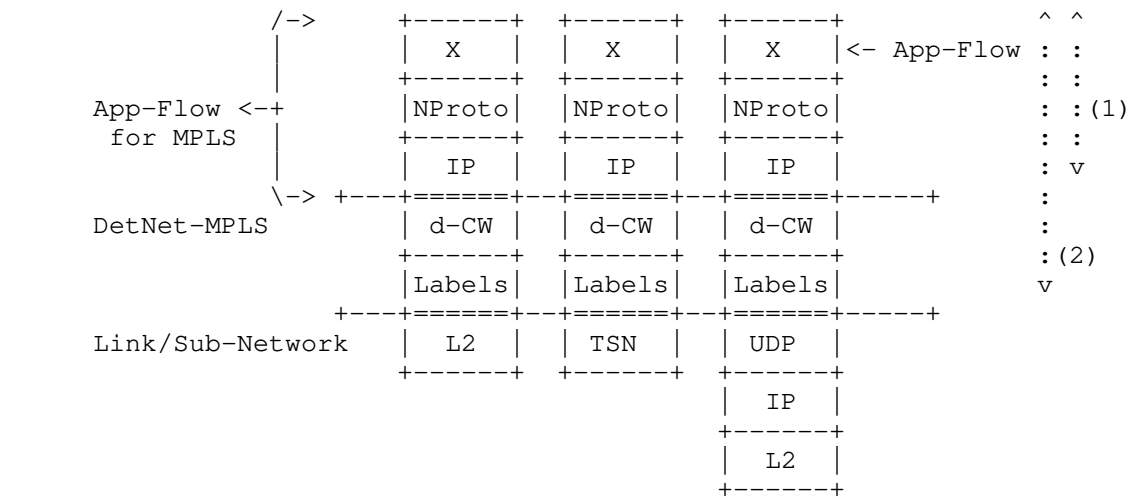

(1) DetNet IP Flow (or simply IP flow)
(2) DetNet MPLS Flow

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 Section 5.1.1. IP Header Information and Section 5.1.2. Other Protocol Header Information in [I-D.ietf-detnet-ip], 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 [I-D.ietf-detnet-mpls].

Per [I-D.ietf-detnet-mpls], the DetNet MPLS data plane uses a single S-Label to support a single app flow. Section 5.1. DetNet IP Flow Identification Procedures in [I-D.ietf-detnet-ip] states that a single DetNet flow is identified based on IP, and next level protocol, header information. Section 4.4. Aggregation Considerations in [I-D.ietf-detnet-ip] 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 this 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. IP over DetNet MPLS Procedures

5.1. DetNet IP over DetNet MPLS Flow Identification 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 [I-D.ietf-detnet-ip] into a single MPLS DetNet flow and MUST process it in accordance to the procedures defined in [I-D.ietf-detnet-mpls] Section 6.1. PRF MAY be supported at the MPLS level for DetNet IP flows sent over an DetNet MPLS network. Aggregation MAY be supported as defined in [I-D.ietf-detnet-mpls] Section 5.4. Aggregation considerations in [I-D.ietf-detnet-ip] 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 [I-D.ietf-detnet-ip] Section 7.1.

An implementation MUST support the provisioning for handling any received 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.

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 [I-D.ietf-detnet-mpls] Section 5.2. Note that the PRF function 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 DetNet IP Traffic Treatment Procedures in [I-D.ietf-detnet-ip].

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 identified using the IP flow identification information as defined in [I-D.ietf-detnet-ip]. The 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 [I-D.ietf-detnet-mpls] Section 5.1, 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:

- S-Label values that are carrying MPLS over IP encapsulated traffic.

- For each S-Label, how the received traffic is to be handled. The traffic may be processed according as any other DetNet IP traffic as defined in this document or in [I-D.ietf-detnet-ip], or the traffic may be directly treated as an MPLS App-flow for additional processing according to [I-D.ietf-detnet-mpls].

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

7. Security Considerations

This draft does not have additional security considerations. Security considerations for DetNet are described in detail in [I-D.ietf-detnet-security]. General security considerations are described in [I-D.ietf-detnet-architecture]. MPLS and IP specific considerations are described in [I-D.ietf-detnet-mpls] and [I-D.ietf-detnet-ip].

Security aspects which are unique to DetNet are those whose aim is to provide the 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 is 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 mechanism 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 definition 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 makes no IANA requests.

9. Acknowledgements

The authors wish to thank Pat Thaler, Norman Finn, Loa Anderson, 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.

10. References

10.1. Normative references

[I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", draft-ietf-detnet-architecture-13 (work in progress), May 2019.

[I-D.ietf-detnet-ip]
Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A., Bryant, S., and J. Korhonen, "DetNet Data Plane: IP", draft-ietf-detnet-ip-01 (work in progress), July 2019.

[I-D.ietf-detnet-mpls]
Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A., Bryant, S., and J. Korhonen, "DetNet Data Plane: MPLS", draft-ietf-detnet-mpls-01 (work in progress), July 2019.

[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>.

10.2. Informative references

[I-D.ietf-detnet-data-plane-framework]
Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A., Bryant, S., and J. Korhonen, "DetNet Data Plane Framework", draft-ietf-detnet-data-plane-framework-02 (work in progress), September 2019.

[I-D.ietf-detnet-security]
Mizrahi, T., Grossman, E., Hacker, A., Das, S., Dowdell, J., Austad, H., Stanton, K., and N. Finn, "Deterministic Networking (DetNet) Security Considerations", draft-ietf-detnet-security-05 (work in progress), August 2019.

[IEEE802.1AE-2018]
IEEE Standards Association, "IEEE Std 802.1AE-2018 MAC Security (MACsec)", 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, December 2005, <https://www.rfc-editor.org/info/rfc4301>.

Authors' Addresses

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

Email: balazs.a.varga@ericsson.com

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

Email: janos.farkas@ericsson.com

Lou Berger
LabN Consulting, L.L.C.

Email: lberger@labn.net

Don Fedyk
LabN Consulting, L.L.C.

Email: dfedyk@labn.net

Andrew G. Malis
Independent

Email: agmalis@gmail.com

Stewart Bryant
Futurewei Technologies

Email: stewart.bryant@gmail.com

Jouni Korhonen

Email: jouni.nospam@gmail.com