

Internet Engineering Task Force (IETF)
Request for Comments: 8964
Category: Standards Track
ISSN: 2070-1721

B. Varga, Ed.
J. Farkas
Ericsson
L. Berger
LabN Consulting, L.L.C.
A. Malis
Malis Consulting
S. Bryant
Futurewei Technologies
J. Korhonen
January 2021

Deterministic Networking (DetNet) Data Plane: MPLS

Abstract

This document specifies the Deterministic Networking (DetNet) data plane when operating over an MPLS Packet Switched Network. It leverages existing pseudowire (PW) encapsulations and MPLS Traffic Engineering (MPLS-TE) encapsulations and mechanisms. This document builds on the DetNet architecture and data plane framework.

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/rfc8964>.

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. Overview of the DetNet MPLS Data Plane
 - 3.1. Layers of DetNet Data Plane
 - 3.2. DetNet MPLS Data Plane Scenarios

- 4. MPLS-Based DetNet Data Plane Solution
 - 4.1. DetNet over MPLS Encapsulation Components
 - 4.2. MPLS Data Plane Encapsulation
 - 4.2.1. DetNet Control Word and DetNet Sequence Number
 - 4.2.2. S-Labels
 - 4.2.3. F-Labels
 - 4.3. OAM Indication
 - 4.4. Flow Aggregation
 - 4.4.1. Aggregation via LSP Hierarchy
 - 4.4.2. Aggregating DetNet Flows as a New DetNet Flow
 - 4.5. Service Sub-Layer Considerations
 - 4.5.1. Edge Node Processing
 - 4.5.2. Relay Node Processing
 - 4.6. Forwarding Sub-Layer Considerations
 - 4.6.1. Class of Service
 - 4.6.2. Quality of Service
- 5. Management and Control Information Summary
 - 5.1. Service Sub-Layer Information Summary
 - 5.1.1. Service Aggregation Information Summary
 - 5.2. Forwarding Sub-Layer Information Summary
- 6. Security Considerations
- 7. IANA Considerations
- 8. References
 - 8.1. Normative References
 - 8.2. 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].

The purpose of this document is to describe the use of the MPLS data plane to establish and support DetNet flows. The DetNet architecture models the DetNet-related data plane functions as being decomposed into two sub-layers: a service sub-layer and a forwarding sub-layer. The service sub-layer is used to provide DetNet service functions, such as protection and reordering. At the DetNet data plane, a new set of functions (Packet Replication, Elimination and Ordering Functions (PREOF)) provide the tasks specific to the service sub-layer. The forwarding sub-layer is used to provide forwarding assurance (low loss, assured latency, and limited out-of-order delivery). The use of the functionalities of the DetNet service sub-layer and the DetNet forwarding sub-layer require careful design and control by the Controller Plane in addition to the DetNet-specific use of MPLS encapsulation as specified by this document.

This document specifies the DetNet data plane operation and the on-wire encapsulation of DetNet flows over an MPLS-based Packet Switched Network (PSN) using the service reference model. MPLS encapsulation already provides a solid foundation of building blocks to enable the DetNet service and forwarding sub-layer functions. MPLS-encapsulated DetNet can be carried over a variety of different network technologies that can provide the level of service required for DetNet. However, the specific details of how DetNet MPLS is carried over different network technologies are out of scope for this document.

MPLS-encapsulated DetNet flows can carry different types of traffic. The details of the types of traffic that are carried in DetNet are also out of scope for this document. An example of IP using DetNet

MPLS sub-networks can be found in [RFC8939]. DetNet MPLS may use an associated controller and Operations, Administration, and Maintenance (OAM) functions that are defined outside of this document.

Background information common to all data planes for DetNet can be found in the DetNet data plane framework [RFC8938].

2. Terminology

2.1. Terms Used in This Document

This document uses the terminology established in the DetNet architecture [RFC8655] and the DetNet data plane framework [RFC8938]. The reader is assumed to be familiar with these documents, any terminology defined therein, and basic MPLS-related terminologies in [RFC3031].

The following terminology is introduced in this document:

F-Label	A DetNet "forwarding" label that identifies the Label Switched Path (LSP) used to forward a DetNet flow across an MPLS PSN, i.e., a hop-by-hop label used between Label Switching Routers (LSRs).
S-Label	A DetNet "service" label that is used between DetNet nodes that implement the DetNet service sub-layer functions. An S-Label is used to identify a DetNet flow at the DetNet service sub-layer at a receiving DetNet node.
A-Label	A special case of an S-Label, whose aggregation properties are known only at the aggregation and deaggregation end points.
d-CW	A DetNet Control Word (d-CW) that is used for sequencing information of a DetNet flow at the DetNet service sub-layer.

2.2. Abbreviations

The following abbreviations are used in this document:

CoS	Class of Service
CW	Control Word
DetNet	Deterministic Networking
LSR	Label Switching Router
MPLS	Multiprotocol Label Switching
MPLS-TE	Multiprotocol Label Switching Traffic Engineering
MPLS-TP	Multiprotocol Label Switching Transport Profile
OAM	Operations, Administration, and Maintenance
PE	Provider Edge
PEF	Packet Elimination Function
PRF	Packet Replication Function
PREOF	Packet Replication, Elimination and Ordering Functions

for providing resource allocation and explicit routes.

3.2. DetNet MPLS Data Plane Scenarios

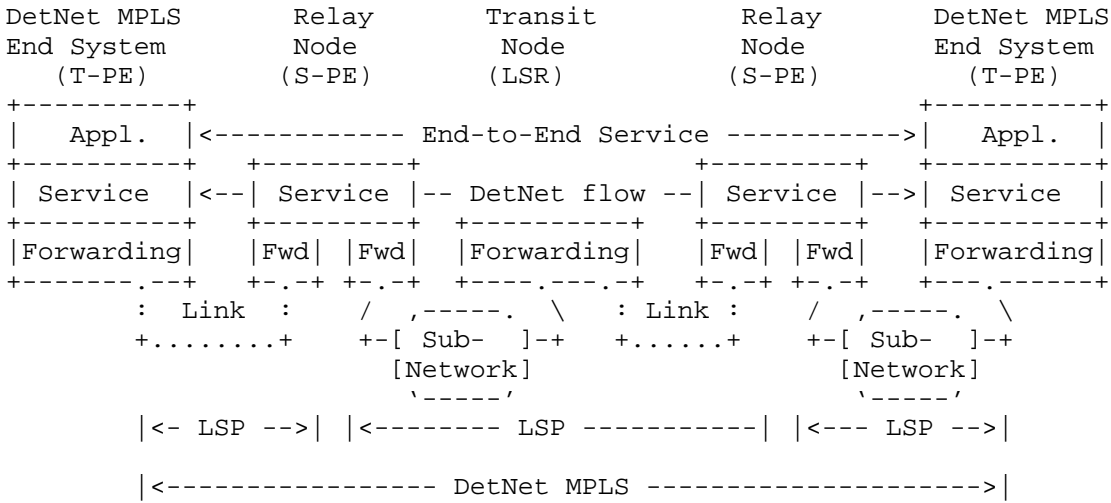


Figure 2: A DetNet MPLS Network

Figure 2 illustrates a hypothetical DetNet MPLS-only network composed of DetNet-aware MPLS-enabled end systems operating over a DetNet-aware MPLS network. In this figure, the relay nodes are PE devices that define the MPLS LSP boundaries, and the transit nodes are LSRs.

DetNet end systems and relay nodes understand the particular needs of DetNet flows and provide both DetNet service and forwarding sub-layer functions. In the case of MPLS, DetNet service-aware nodes add, remove, and process d-CWs, S-Labels, and F-Labels as needed. DetNet MPLS nodes provide functionality analogous to T-PEs when they sit at the edge of an MPLS domain and S-PEs when they are in the middle of an MPLS domain; see [RFC6073].

In a DetNet MPLS network, transit nodes may be DetNet-service-aware or DetNet-unaware MPLS Label Switching Routers (LSRs). In this latter case, such LSRs would be unaware of the special requirements of the DetNet service sub-layer but would still provide traffic engineering functions and the QoS capabilities needed to ensure that the (TE) LSPs meet the service requirements of the carried DetNet flows.

The application of DetNet using MPLS supports a number of control and management plane types. These types support certain MPLS data plane deployments. For example, RSVP-TE, MPLS-TP, or MPLS Segment Routing (when extended to support resource allocation) are all valid MPLS deployment cases.

Figure 3 illustrates how an end-to-end MPLS-based DetNet service is provided in more detail. In this figure, the Customer Edge (CE1 and CE2) are able to send and receive MPLS-encapsulated DetNet flows, and R1, R2, and R3 are relay nodes in the middle of a DetNet network. The 'X' in the end systems and relay nodes represents potential DetNet compound flow packet replication and elimination points. In this example, service protection is supported utilizing at least two DetNet member flows and TE LSPs. For a unidirectional flow, R1 supports PRF, and R3 supports PEF and POF. Note that the relay nodes may change the underlying forwarding sub-layer, for example, tunneling MPLS over IEEE 802.1 TSN [DetNet-MPLS-over-TSN] or simply over interconnected network links.



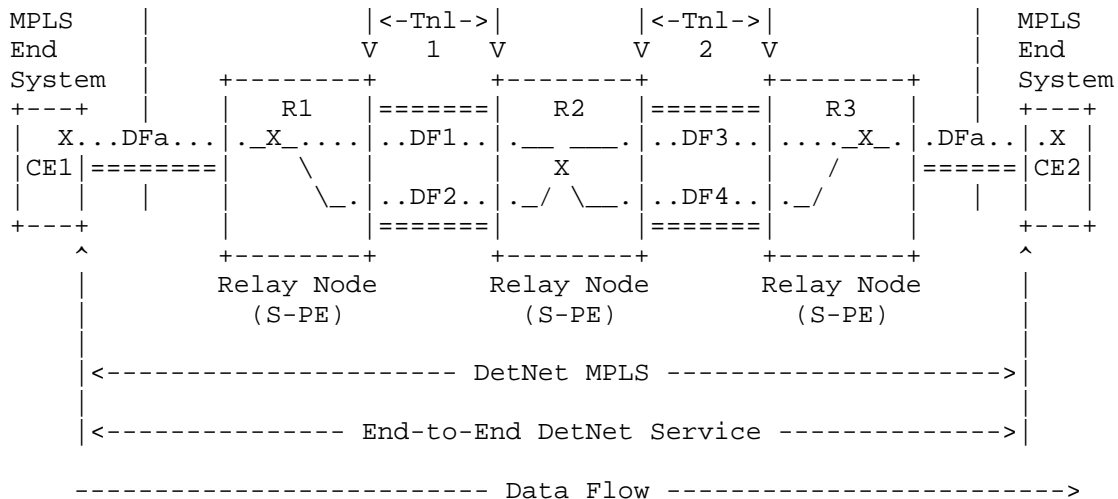


Figure 3: MPLS-Based Native DetNet

X - Optional service protection (none, PRF, PREOF, PEF/POF)

DFx - DetNet member flow x over a TE LSP

4. MPLS-Based DetNet Data Plane Solution

4.1. DetNet over MPLS Encapsulation Components

To carry DetNet over MPLS, the following is required:

1. A method of identifying the MPLS payload type.
2. A method of identifying the DetNet flow(s) to the processing element.
3. A method of distinguishing DetNet OAM packets from DetNet data packets.
4. A method of carrying the DetNet sequence number.
5. A suitable LSP to deliver the packet to the egress PE.
6. A method of carrying queuing and forwarding indication.

In this design, an MPLS service label (the S-Label) is similar to a pseudowire (PW) label [RFC3985] and is used to identify both the DetNet flow identity and the MPLS payload type satisfying (1) and (2) in the list above. OAM traffic discrimination happens through the use of the Associated Channel method described in [RFC4385]. The DetNet sequence number is carried in the DetNet Control Word, which also carries the Data/OAM discriminator. To simplify implementation and to maximize interoperability, two sequence number sizes are supported: a 16-bit sequence number and a 28-bit sequence number. The 16-bit sequence number is needed to support some types of legacy clients. The 28-bit sequence number is used in situations where it is necessary to ensure that, in high-speed networks, the sequence number space does not wrap whilst packets are in flight.

The LSP used to forward the DetNet packet is not restricted regarding any method used for establishing that LSP (for example, MPLS-LDP, MPLS-TE, MPLS-TP [RFC5921], MPLS Segment Routing [RFC8660], etc.). The F-Label(s) and the S-Label may be used alone or together to indicate the required queue processing as well as the forwarding parameters. Note that the possible use of Penultimate Hop Popping (PHP) means that the S-Label may be the only label received at the terminating DetNet service.

4.2. MPLS Data Plane Encapsulation

Figure 4 illustrates a DetNet data plane MPLS encapsulation. The MPLS-based encapsulation of the DetNet flows is well suited for the scenarios described in [RFC8938]. Furthermore, an end-to-end DetNet service, i.e., native DetNet deployment (see Section 3.2), is also possible if DetNet end systems are capable of initiating and terminating MPLS-encapsulated packets.

The MPLS-based DetNet data plane encapsulation consists of:

- * A DetNet Control Word (d-CW) containing sequencing information for packet replication and duplicate elimination purposes, and the OAM indicator.
- * A DetNet service label (S-Label) that identifies a DetNet flow at the receiving DetNet service sub-layer processing node.
- * Zero or more DetNet MPLS forwarding label(s) (F-Label) used to direct the packet along the Label Switched Path (LSP) to the next DetNet service sub-layer processing node along the path. When PHP is in use, there may be no F-Label in the protocol stack on the final hop.
- * The necessary data-link encapsulation is then applied prior to transmission over the physical media.

DetNet MPLS-based encapsulation

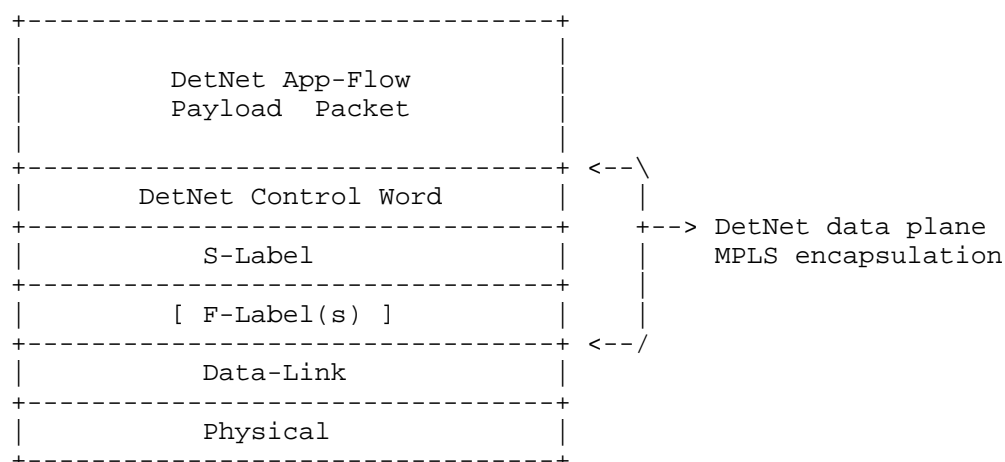


Figure 4: Encapsulation of a DetNet App-Flow in an MPLS PSN

4.2.1. DetNet Control Word and DetNet Sequence Number

A DetNet Control Word (d-CW) conforms to the Generic PW MPLS Control Word (PWMCW) defined in [RFC4385]. The d-CW formatted as shown in Figure 5 MUST be present in all DetNet packets containing App-flow data. This format of the d-CW was created in order to (1) allow larger sequence number space to avoid sequence number rollover frequency in some applications and (2) allow sequence numbering systems that include the value zero as a valid sequence number, which simplifies implementation.




Figure 5: DetNet Control Word

(bits 0 to 3)

Per [RFC4385], MUST be set to zero (0).

Sequence Number (bits 4 to 31)

An unsigned value implementing the DetNet sequence number. The sequence number space is a circular one with no restriction on the initial value.

A separate sequence number space MUST be maintained by the node that adds the d-CW for each DetNet App-flow, i.e., DetNet service. The following Sequence Number field lengths MUST be supported:

- * 0 bits
- * 16 bits
- * 28 bits

The sequence number length MUST be provisioned on a per-DetNet-service basis via configuration, i.e., via the Controller Plane described in [RFC8938].

A 0-bit Sequence Number field length indicates that there is no DetNet sequence number used for the flow. When the length is zero, the Sequence Number field MUST be set to zero (0) on all packets sent for the flow.

When the Sequence Number field length is 16 or 28 bits for a flow, the sequence number MUST be incremented by one for each new App-flow packet sent. When the field length is 16 bits, d-CW bits 4 to 15 MUST be set to zero (0). The values carried in this field can wrap, and it is important to note that zero (0) is a valid field value. For example, where the sequence number size is 16 bits, the sequence will contain: 65535, 0, 1, where zero (0) is an ordinary sequence number.

It is important to note that this document differs from [RFC4448], where a sequence number of zero (0) is used to indicate that the sequence number check algorithm is not used.

The sequence number is optionally used during receive processing, as described below in Sections 4.2.2.2 and 4.2.2.3.

4.2.2. S-Labels

A DetNet flow at the DetNet service sub-layer is identified by an S-Label. MPLS-aware DetNet end systems and edge nodes, which are by definition MPLS ingress and egress nodes, MUST add (push) and remove (pop) a DetNet service-specific d-CW and S-Label. Relay nodes MAY swap S-Label values when processing a DetNet flow, i.e., incoming and outgoing S-Labels of a DetNet flow can be different.

S-Label values MUST be provisioned per DetNet service via configuration, i.e., via the Controller Plane described in [RFC8938]. Note that S-Labels provide identification at the downstream DetNet service sub-layer receiver, not the sender. As such, S-Labels MUST be allocated by the entity that controls the service sub-layer receiving a node's label space and MAY be allocated from the platform label space [RFC3031]. Because S-Labels are local to each node, rather than being a global identifier within a domain, they must be advertised to their upstream DetNet service-aware peer nodes (i.e., a DetNet MPLS end system or a DetNet relay or edge node) and interpreted in the context of their received F-Label(s). In some PREOF topologies, the node performing replication will be sending to

multiple nodes performing PEF or POF and may need to send different S-Label values for the different member flows for the same DetNet service.

An S-Label will normally be at the bottom of the label stack once the last F-Label is removed, immediately preceding the d-CW. To support OAM at the service sub-layer level, an OAM Associated Channel Header (ACH) [RFC4385] together with a Generic Associated Channel Label (GAL) [RFC5586] MAY be used in place of a d-CW.

Similarly, an Entropy Label Indicator (ELI) and Entropy Label (EL) [RFC6790] MAY be carried below the S-Label in the label stack in networks where DetNet flows would otherwise receive ECMP treatment. When ELs are used, the same EL value SHOULD be used for all of the packets sent using a specific S-Label to force the flow to follow the same path. However, as outlined in [RFC8938], the use of ECMP for DetNet flows is NOT RECOMMENDED. ECMP MAY be used for non-DetNet flows within a DetNet domain.

When receiving a DetNet MPLS packet, an implementation MUST identify the DetNet service associated with the incoming packet based on the S-Label. When a node is using platform labels for S-Labels, no additional information is needed, as the S-Label uniquely identifies the DetNet service. In the case where platform labels are not used, zero or more F-Labels preceding the S-Label MUST be used together with the S-Label to uniquely identify the DetNet service associated with a received packet. The incoming interface MAY also be used together with any present F-Label(s) and the S-Label to uniquely identify an incoming DetNet service, for example, in the case where PHP is used. Note that the choice to use the platform label space for an S-Label or an S-Label plus one or more F-Labels to identify DetNet services is a local implementation choice, with one caveat. When one or more F-Labels, or the incoming interface, is needed together with an S-Label to uniquely identify a service, the Controller Plane must ensure that incoming DetNet MPLS packets arrive with the needed information (F-Label(s) and/or the incoming interface) and provision the needed information. The provisioned information MUST then be used to identify incoming DetNet service based on the combination of S-Label and F-Label(s) or the incoming interface.

The use of platform labels for S-Labels matches other pseudowire encapsulations for consistency, but there is no hard requirement in this regard.

Implementation details of PREOF are out of scope for this document. [IEEE802.1CB-2017] defines equivalent replication and elimination-specific aspects, which can be applied to PRF and PEF.

4.2.2.1. Packet Replication Function Processing

The Packet Replication Function (PRF) MAY be supported by an implementation for outgoing DetNet flows. The use of the PRF for a particular DetNet service MUST be provisioned via configuration, i.e., via the Controller Plane described in [RFC8938]. When replication is configured, the same App-flow data will be sent over multiple outgoing DetNet member flows using forwarding sub-layer LSPs. An S-Label value MUST be configured per outgoing member flow. The same d-CW field value MUST be used on all outgoing member flows for each replicated MPLS packet.

4.2.2.2. Packet Elimination Function Processing

Implementations MAY support the Packet Elimination Function (PEF) for received DetNet MPLS flows. When supported, use of the PEF for a particular DetNet service MUST be provisioned via configuration,

i.e., via the Controller Plane described in [RFC8938].

After a DetNet service is identified for a received DetNet MPLS packet, as described above, if PEF is configured for that DetNet service, duplicate (replicated) instances of a particular sequence number MUST be discarded. The specific mechanisms used for an implementation to identify which received packets are duplicates and which are new is an implementation choice. Note that, per Section 4.2.1, the Sequence Number field length may be 16 or 28 bits, and the field value can wrap. PEF MUST NOT be used with DetNet flows configured with a d-CW Sequence Number field length of 0 bits.

An implementation MAY constrain the maximum number of sequence numbers that are tracked on either a platform-wide or per-flow basis. Some implementations MAY support the provisioning of the maximum number of sequence numbers that are tracked on either a platform-wide or per-flow basis.

4.2.2.3. Packet Ordering Function Processing

A function that is related to in-order delivery is the Packet Ordering Function (POF). Implementations MAY support POF. When supported, use of the POF for a particular DetNet service MUST be provisioned via configuration, i.e., via the Controller Plane described by [RFC8938]. Implementations MAY require that PEF and POF be used in combination. There is no requirement related to the order of execution of the PEF and POF in an implementation.

After a DetNet service is identified for a received DetNet MPLS packet, as described above, if POF is configured for that DetNet service, packets MUST be processed in the order indicated in the received d-CW Sequence Number field, which may not be in the order the packets are received. As defined in Section 4.2.1, the Sequence Number field length may be 16 or 28 bits, the sequence number is incremented by one (1) for each new MPLS packet sent for a particular DetNet service, and the field value can wrap. The specific mechanisms used for an implementation to identify the order of received packets is an implementation choice.

An implementation MAY constrain the maximum number of out-of-order packets that can be processed on either a platform-wide or per-flow basis. The number of out-of-order packets that can be processed also impacts the latency of a flow.

The latency impact on the system resources needed to support a specific DetNet flow will need to be evaluated by the Controller Plane based on that flow's traffic specification. An example traffic specification that can be used with MPLS-TE can be found in [RFC2212].

DetNet implementations can use flow-specific requirements (e.g., maximum number of out-of-order packets and maximum latency of the flow) for configuration of POF-related buffers. POF implementation details are out of scope for this document, and POF configuration parameters are implementation specific. The Controller Plane identifies and sets the POF configuration parameters.

4.2.3. F-Labels

F-Labels support the DetNet forwarding sub-layer. F-Labels are used to provide LSP-based connectivity between DetNet service sub-layer processing nodes.

4.2.3.1. Service Sub-Layer-Related Processing

DetNet MPLS end systems, edge nodes, and relay nodes may operate at

the DetNet service sub-layer with understanding of DetNet services and their requirements. As mentioned earlier, when operating at this layer, such nodes can push, pop, or swap (pop then push) S-Labels. In all cases, the F-Label(s) used for a DetNet service are always replaced, and the following procedures apply.

When sending a DetNet flow, zero or more F-Labels MAY be pushed on top of an S-Label by the node pushing an S-Label. The F-Label(s) to be pushed when sending a particular DetNet service MUST be provisioned per outgoing S-Label via configuration, i.e., via the Controller Plane discussed in [RFC8938]. F-Label(s) can also provide context for an S-Label. To allow for the omission of F-Label(s), an implementation SHOULD also allow an outgoing interface to be configured per S-Label.

Note that when PRF is supported, the same App-flow data will be sent over multiple outgoing DetNet member flows using forwarding sub-layer LSPs. This means that an implementation may be sending different sets of F-Labels per DetNet member flow, each with a different S-Label.

When a single set of F-Labels is provisioned for a particular outgoing S-Label, that set of F-Labels MUST be pushed after the S-Label is pushed. The outgoing packet is then forwarded, as described below in Section 4.2.3.2. When a single outgoing interface is provisioned, the outgoing packet is then forwarded, as described below in Section 4.2.3.2.

When multiple sets of outgoing F-Labels or interfaces are provisioned for a particular DetNet service (i.e., for PRF), a copy of the outgoing packet, including the pushed member flow-specific S-Label, MUST be made per F-Label set and outgoing interface. Each set of provisioned F-Labels are then pushed onto a copy of the packet. Each copy is then forwarded, as described below in Section 4.2.3.2.

As described in the previous section, when receiving a DetNet MPLS flow, an implementation identifies the DetNet service associated with the incoming packet based on the S-Label. When a node is using platform labels for S-Labels, any F-Labels can be popped, and the S-Label uniquely identifies the DetNet service. In the case where platform labels are not used, incoming F-Label(s) and, optionally, the incoming interface MUST also be provisioned for a DetNet service.

4.2.3.2. Common F-Label Processing

All DetNet-aware MPLS nodes process F-Labels as needed to meet the service requirements of the DetNet flow or flows carried in the LSPs represented by the F-Labels. This includes normal push, pop, and swap operations. Such processing is essentially the same type of processing provided for TE LSPs, although the specific service parameters, or traffic specification, can differ. When the DetNet service parameters of the DetNet flow or flows carried in an LSP represented by an F-Label can be met by an existing TE mechanism, the forwarding sub-layer processing node MAY be a DetNet-unaware, i.e., standard, MPLS LSR. Such TE LSPs may provide LSP forwarding service as defined in, but not limited to, the following: [RFC3209], [RFC3270], [RFC3272], [RFC3473], [RFC4875], [RFC5440], and [RFC8306].

More specifically, as mentioned above, the DetNet forwarding sub-layer provides explicit routes and allocated resources, and F-Labels are used to map to each. Explicit routes are supported based on the topmost (outermost) F-Label that is pushed or swapped and the LSP that corresponds to this label. This topmost (outgoing) label MUST be associated with a provisioned outgoing interface and, for non-point-to-point outgoing interfaces, a next-hop LSR. Note that this information MUST be provisioned via configuration or the Controller

Plane. In the previously mentioned special case where there are no added F-Labels and the outgoing interface is not a point-to-point interface, the outgoing interface MUST also be associated with a next-hop LSR.

Resources may be allocated in a hierarchical fashion per each LSP that is represented by each F-Label. Each LSP MAY be provisioned with a service parameter that dictates the specific traffic treatment to be received by the traffic carried over that LSP. Implementations of this document MUST ensure that traffic carried over each LSP represented by one or more F-Labels receives the traffic treatment provisioned for that LSP. Typical mechanisms used to provide different treatment to different flows include the allocation of system resources (such as queues and buffers) and provisioning of related parameters (such as shaping and policing) that may be found in implementations of the Resource ReSerVation Protocol (RSVP) [RFC2205] and RSVP-TE [RFC3209] [RFC3473]. Support can also be provided via an underlying network technology, such as IEEE 802.1 TSN [DetNet-MPLS-over-TSN]. The specific mechanisms selected by a DetNet node to ensure DetNet service delivery requirements are met for supported DetNet flows is outside the scope of this document.

Packets that are marked in a way that do not correspond to allocated resources, e.g., lack a provisioned F-Label, can disrupt the QoS offered to properly reserved DetNet flows by using resources allocated to the reserved flows. Therefore, the network nodes of a DetNet network:

- * MUST defend the DetNet QoS by discarding or remarking (to an allocated DetNet flow or noncompeting non-DetNet flow) packets received that are not associated with a completed resource allocation.
- * MUST NOT use a DetNet allocated resource, e.g., a queue or shaper reserved for DetNet flows, for any packet that does match the corresponding DetNet flow.
- * MUST ensure a QoS flow does not exceed its allocated resources or provisioned service level.
- * MUST ensure a CoS flow or service class does not impact the service delivered to other flows. This requirement is similar to the requirement for MPLS LSRs that CoS LSPs do not impact the resources allocated to TE LSPs, e.g., via [RFC3473].

Subsequent sections provide additional considerations related to CoS (Section 4.6.1), QoS (Section 4.6.2), and aggregation (Section 4.4).

4.3. OAM Indication

OAM follows the procedures set out in [RFC5085] with the restriction that only Virtual Circuit Connectivity Verification (VCCV) type 1 is supported.

As shown in Figure 3 of [RFC5085], when the first nibble of the d-CW is 0x0, the payload following the d-CW is normal user data. However, when the first nibble of the d-CW is 0x1, the payload that follows the d-CW is an OAM payload with the OAM type indicated by the value in the d-CW Channel Type field.

The reader is referred to [RFC5085] for a more detailed description of the Associated Channel mechanism and to the DetNet work on OAM [DetNet-MPLS-OAM] for more information about DetNet OAM.

Additional considerations on DetNet-specific OAM are subjects for further study.

4.4. Flow Aggregation

The ability to aggregate individual flows and their associated resource control into a larger aggregate is an important technique for improving scaling of control in the data, management, and control planes. The DetNet data plane allows for the aggregation of DetNet flows to improved scaling. There are two methods of supporting flow aggregation covered in this section.

The resource control and management aspects of aggregation (including the configuration of queuing, shaping, and policing) are the responsibility of the DetNet Controller Plane and are out of scope for this document. It is also the responsibility of the Controller Plane to ensure that consistent aggregation methods are used.

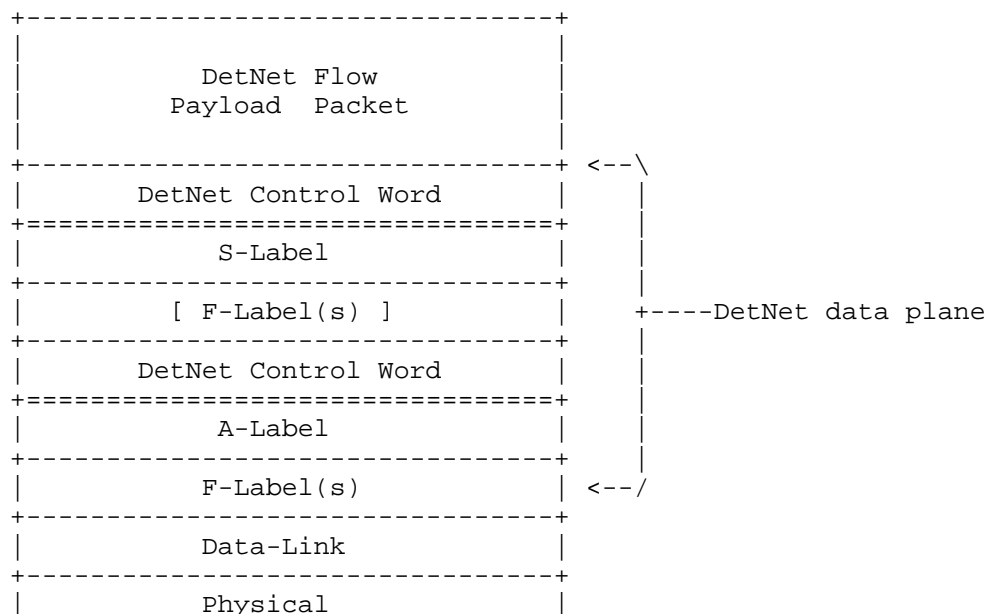
4.4.1. Aggregation via LSP Hierarchy

DetNet flows forwarded via MPLS can leverage MPLS-TE's existing support for hierarchical LSPs (H-LSPs); see [RFC4206]. H-LSPs are typically used to aggregate control and resources; they may also be used to provide OAM or protection for the aggregated LSPs. Arbitrary levels of aggregation naturally fall out of the definition for hierarchy and the MPLS label stack [RFC3032]. DetNet nodes that support aggregation (LSP hierarchy) map one or more LSPs (labels) into and from an H-LSP. Both carried LSPs and H-LSPs may or may not use the Traffic Class (TC) field, i.e., L-LSPs (Label-Only-Inferred-PSC LSPs) or E-LSPs (EXP-Inferred-PSC LSPs [RFC3270], which were renamed to "Explicitly TC-encoded-PSC LSPs" in Section 2.2 of [RFC5462]). Such nodes will need to ensure that individual LSPs and H-LSPs receive the traffic treatment required to ensure the required DetNet service is preserved.

Additional details of the traffic control capabilities needed at a DetNet-aware node may be covered in the new service definitions mentioned above or in separate future documents. Controller Plane mechanisms will also need to ensure that the service required on the aggregate flow are provided, which may include the discarding or remarking mentioned in the previous sections.

4.4.2. Aggregating DetNet Flows as a New DetNet Flow

An aggregate can be built by layering DetNet flows, including both their S-Label and (when present) F-Labels, as shown below:



+-----+

Figure 6: DetNet Aggregation as a New DetNet Flow

Both the aggregation label, which is referred to as an A-Label, and the individual flow's S-Label have their MPLS S bit set indicating the bottom of stack, and the d-CW allows the PREOF to work. An A-Label is a special case of an S-Label, whose properties are known only at the aggregation and deaggregation end points.

It is a property of the A-Label that what follows is a d-CW followed by an MPLS label stack. A relay node processing the A-Label would not know the underlying payload type, and the A-Label would be processed as a normal S-Label. This would only be known to a node that was a peer of the node imposing the S-Label. However, there is no real need for it to know the payload type during aggregation processing.

As in the previous section, nodes supporting this type of aggregation will need to ensure that individual and aggregated flows receive the traffic treatment required to ensure the required DetNet service is preserved. Also, it is the Controller Plane's responsibility to ensure that the service required on the aggregate flow is properly provisioned.

4.5. Service Sub-Layer Considerations

The internal procedures for edge and relay nodes related to PREOF are implementation specific. The order of a packet elimination or replication is out of scope for this specification.

It is important that the DetNet layer is configured such that a DetNet node never receives its own replicated packets. If it were to receive such packets, the replication function would make the loop more destructive of bandwidth than a conventional unicast loop. Ultimately, the TTL in the S-Label will cause the packet to die during a transient loop, but given the sensitivity of applications to packet latency, the impact on the DetNet application would be severe. To avoid the problem of a transient forwarding loop, changes to an LSP supporting DetNet MUST be loop-free.

4.5.1. Edge Node Processing

A DetNet edge node operates in the DetNet forwarding sub-layer and service sub-layer. An edge node is responsible for matching incoming packets to the service they require and encapsulating them accordingly. An edge node may perform PRF, PEF, and/or POF. Details on encapsulation can be found in Section 4.2; details on PRF can be found in Section 4.2.2.1; details on PEF can be found in Section 4.2.2.2; and details on POF can be found in Section 4.2.2.3.

4.5.2. Relay Node Processing

A DetNet relay node operates in the DetNet forwarding sub-layer and service sub-layer. For DetNet using MPLS, forwarding-related processing is performed on the F-Label. This processing is done within an extended forwarder function. Whether an incoming DetNet flow receives DetNet-specific processing depends on how the forwarding is programmed. Some relay nodes may be DetNet service aware for certain DetNet services, while, for other DetNet services, these nodes may perform as unmodified LSRs that only understand how to switch MPLS-TE LSPs, i.e., as a transit node; see Section 4.4. Again, this is entirely up to how the forwarding has been programmed.

During the elimination and replication process, the sequence number of an incoming DetNet packet MUST be preserved and carried in the

corresponding outgoing DetNet packet. For example, a relay node that performs both PEF and PRF first performs PEF on incoming packets to create a compound flow. It then performs PRF and copies the App-flow data and the d-CW into packets for each outgoing DetNet member flow.

The internal design of a relay node is out of scope for this document. However, the reader's attention is drawn to the need to make any PREOF state available to the packet processor(s) dealing with packets to which PREOF must be applied and to maintain that state in such a way that it is available to the packet processor operation on the next packet in the DetNet flow (which may be a duplicate, a late packet, or the next packet in sequence).

4.6. Forwarding Sub-Layer Considerations

4.6.1. Class of Service

Class of Service (CoS) and Quality of Service (QoS) are terms that are often used interchangeably and confused with each other. In the context of DetNet, CoS is used to refer to mechanisms that provide traffic-forwarding treatment based on non-flow-specific traffic classification, and QoS is used to refer to mechanisms that provide traffic-forwarding treatment based on DetNet flow-specific traffic classification. Examples of existing network-level CoS mechanisms include Differentiated Services (Diffserv), which is enabled by the IP header Differentiated Services Code Point (DSCP) field [RFC2474] and MPLS label Traffic Class field [RFC5462] and at Layer 2 by IEEE 802.1p Priority Code Point (PCP).

CoS for DetNet flows carried in PWs and MPLS is provided using the existing MPLS Differentiated Services (Diffserv) architecture [RFC3270]. Both E-LSP and L-LSP MPLS Diffserv modes MAY be used to support DetNet flows. The Traffic Class field (formerly the EXP field) of an MPLS label follows the definition of [RFC5462] and [RFC3270]. The Uniform, Pipe, and Short Pipe Diffserv tunneling and TTL processing models are described in [RFC3270] and [RFC3443] and MAY be used for MPLS LSPs supporting DetNet flows. MPLS Explicit Congestion Notification (ECN) MAY also be used, as defined in ECN [RFC5129] and updated by [RFC5462].

4.6.2. Quality of Service

In addition to explicit routes and packet replication and elimination (described in Section 4 above), DetNet provides zero congestion loss and bounded latency and jitter. As described in [RFC8655], there are different mechanisms that may be used separately or in combination to deliver a zero congestion loss service. This includes QoS mechanisms at the MPLS layer, which may be combined with the mechanisms defined by the underlying network layer, such as IEEE 802.1 TSN.

QoS mechanisms for flow-specific traffic treatment typically include a guarantee/agreement for the service and allocation of resources to support the service. Example QoS mechanisms include discrete resource allocation, admission control, flow identification and isolation, and sometimes path control, traffic protection, shaping, policing, and remarking. Example protocols that support QoS control include the Resource ReSerVation Protocol (RSVP) [RFC2205] and RSVP-TE [RFC3209] [RFC3473]. The existing MPLS mechanisms defined to support CoS [RFC3270] can also be used to reserve resources for specific traffic classes.

A baseline set of QoS capabilities for DetNet flows carried in PWs and MPLS can be provided by MPLS-TE [RFC3209] [RFC3473]. TE LSPs can also support explicit routes (path pinning). Current service definitions for packet TE LSPs can be found in "Specification of the Controlled-Load Network Element Service" [RFC2211], "Specification of

Guaranteed Quality of Service" [RFC2212], and "Ethernet Traffic Parameters" [RFC6003]. Additional service definitions are expected in future documents to support the full range of DetNet services. In all cases, the existing label-based marking mechanisms defined for TE LSPs and even E-LSPs are used to support the identification of flows requiring DetNet QoS.

5. Management and Control Information Summary

The specific information needed for the processing of each DetNet service depends on the DetNet node type and the functions being provided on the node. This section summarizes this information based on the DetNet sub-layer and if the DetNet traffic is being sent or received. All DetNet node types are DetNet forwarding sub-layer aware, while all but transit nodes are service sub-layer aware. This is shown in Figure 2.

Much like other MPLS labels, there are a number of alternatives available for DetNet S-Label and F-Label advertisement to an upstream peer node. These include distributed signaling protocols (such as RSVP-TE), centralized label distribution via a controller that manages both the sender and the receiver using the Network Configuration Protocol (NETCONF) and YANG, BGP, the Path Computation Element Communication Protocol (PCEP), etc., and hybrid combinations of the two. The details of the Controller Plane solution required for the label distribution and the management of the label number space are out of scope for this document. Particular DetNet considerations and requirements are discussed in [RFC8938]. Conformance language is not used in the summary, since it applies to future mechanisms, such as those that may be provided in signaling and YANG models, e.g., [DetNet-YANG].

5.1. Service Sub-Layer Information Summary

The following summarizes the information that is needed (on a per-service basis) on nodes that are service sub-layer aware and transmit DetNet MPLS traffic:

- * App-flow identification information, e.g., IP information as defined in [DetNet-IP-over-MPLS]. Note that this information is not needed on DetNet relay nodes.
- * The sequence number length to be used for the service. Valid values include 0, 16, and 28 bits. 0 bits cannot be used when PEF or POF is configured for the service.
- * If PRF is to be provided for the service.
- * The outgoing S-Label for each of the service's outgoing DetNet (member) flows.
- * The forwarding sub-layer information associated with the output of the service sub-layer. Note that when PRF is provisioned, this information is per DetNet member flow. Logically, the forwarding sub-layer information is a pointer to further details of transmission of DetNet flows at the forwarding sub-layer.

The following summarizes the information that is needed (on a per-service basis) on nodes that are service sub-layer aware and receive DetNet MPLS traffic:

- * The forwarding sub-layer information associated with the input of the service sub-layer. Note that when PEF is provisioned, this information is per DetNet member flow. Logically, the forwarding sub-layer information is a pointer to further details of the reception of DetNet flows at the forwarding sub-layer or A-Label.

- * The incoming S-Label for the service.
- * If PEF or POF is to be provided for the service.
- * The sequence number length to be used for the service. Valid values included 0, 16, and 28 bits. 0 bits cannot be used when PEF or POF are configured for the service.
- * App-flow identification information, e.g., IP information as defined in [DetNet-IP-over-MPLS]. Note that this information is not needed on DetNet relay nodes.

5.1.1. Service Aggregation Information Summary

Nodes performing aggregation using A-Labels, per Section 4.4.2, require the additional information summarized in this section.

The following summarizes the additional information that is needed on a node that sends aggregated flows using A-Labels:

- * The S-Labels or F-Labels that are to be carried over each aggregated service.
- * The A-Label associated with each aggregated service.
- * The other S-Label information summarized above.

On the receiving node, the A-Label provides the forwarding context of an incoming interface or an F-Label and is used in subsequent service or forwarding sub-layer receive processing, as appropriate. The related additional configuration that may be provided is discussed elsewhere in this section.

5.2. Forwarding Sub-Layer Information Summary

The following summarizes the information that is needed (on a per-forwarding-sub-layer-flow basis) on nodes that are forwarding sub-layer aware and send DetNet MPLS traffic:

- * The outgoing F-Label stack to be pushed. The stack may include H-LSP labels.
- * The traffic parameters associated with a specific label in the stack. Note that there may be multiple sets of traffic parameters associated with specific labels in the stack, e.g., when H-LSPs are used.
- * Outgoing interface and, for unicast traffic, the next-hop information.
- * Sub-network-specific parameters on a technology-specific basis. For example, see [DetNet-MPLS-over-TSN].

The following summarizes the information that is needed (on a per-forwarding-sub-layer-flow basis) on nodes that are forwarding sub-layer aware and receive DetNet MPLS traffic:

- * The incoming interface. For some implementations and deployment scenarios, this information may not be needed.
- * The incoming F-Label stack to be popped. The stack may include H-LSP labels.
- * How the incoming forwarding sub-layer flow is to be handled, i.e., forwarded as a transit node or provided to the service sub-layer.

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 needed to meet each flow's service requirements. This applies for aggregated and individual flows.

6. Security Considerations

Detailed security considerations for DetNet are cataloged in [DetNet-Security], and more general security considerations are described in [RFC8655]. This section exclusively considers security considerations that are specific to the DetNet MPLS data plane. The considerations raised related to MPLS networks in general in [RFC5920] are equally applicable to the DetNet MPLS data plane.

Security aspects that are unique to DetNet are those whose aim is to protect the support of specific QoS aspects of DetNet, which are primarily to deliver data flows with extremely low packet loss rates and bounded end-to-end delivery latency. Achieving such loss rates and bounded latency may not be possible in the face of a highly capable adversary, such as the one envisioned by the Internet Threat Model of BCP 72 [RFC3552] that can arbitrarily drop or delay any or all traffic. In order to present meaningful security considerations, we consider a somewhat weaker attacker who does not control the physical links of the DetNet domain but may have the ability to control a network node within the boundary of the DetNet domain.

An additional consideration for the DetNet 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 are 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-network using MACsec [IEEE802.1AE-2018] for IP over Ethernet (Layer 2) flows. MPLS doesn't provide any native security services to account for confidentiality and integrity.

From a data plane perspective, this document does not add or modify any application 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 that 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 is prevented or mitigated through the use of existing mechanisms, for example, by policing and shaping incoming traffic. To prevent DetNet packets from having their delay manipulated by an external entity, precautions need to be taken to ensure that all devices on an LSP are those intended to be there by the network operator and that they are well behaved. In addition to physical security, technical methods, such as authentication and authorization of network equipment and the instrumentation and monitoring of the LSP packet delay, may be used. If a delay attack is suspected, traffic may be moved to an alternate path, for example, through changing the LSP or management of the PREOF configuration.

7. IANA Considerations

This document has no IANA actions.

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>>.
- [RFC2211] Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, DOI 10.17487/RFC2211, September 1997, <<https://www.rfc-editor.org/info/rfc2211>>.
- [RFC2212] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, DOI 10.17487/RFC2212, September 1997, <<https://www.rfc-editor.org/info/rfc2212>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, DOI 10.17487/RFC3270, May 2002, <<https://www.rfc-editor.org/info/rfc3270>>.
- [RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, DOI 10.17487/RFC3443, January 2003, <<https://www.rfc-editor.org/info/rfc3443>>.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, DOI 10.17487/RFC3473, January 2003, <<https://www.rfc-editor.org/info/rfc3473>>.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, DOI 10.17487/RFC4206, October 2005, <<https://www.rfc-editor.org/info/rfc4206>>.
- [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>>.
- [RFC5085] Nadeau, T., Ed. and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, DOI 10.17487/RFC5085,

December 2007, <<https://www.rfc-editor.org/info/rfc5085>>.

- [RFC5129] Davie, B., Briscoe, B., and J. Tay, "Explicit Congestion Marking in MPLS", RFC 5129, DOI 10.17487/RFC5129, January 2008, <<https://www.rfc-editor.org/info/rfc5129>>.
- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, DOI 10.17487/RFC5462, February 2009, <<https://www.rfc-editor.org/info/rfc5462>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, DOI 10.17487/RFC5586, June 2009, <<https://www.rfc-editor.org/info/rfc5586>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.
- [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>>.

8.2. Informative References

- [DetNet-IP-over-MPLS]
Varga, B., Ed., Berger, L., Fedyk, D., Bryant, S., and J. Korhonen, "DetNet Data Plane: IP over MPLS", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-over-mpls-09, 11 October 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-ip-over-mpls-09>>.
- [DetNet-MPLS-OAM]
Mirsky, G. and M. Chen, "Operations, Administration and Maintenance (OAM) for Deterministic Networks (DetNet) with MPLS Data Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-mpls-oam-02, 15 January 2021, <<https://tools.ietf.org/html/draft-ietf-detnet-mpls-oam-02>>.
- [DetNet-MPLS-over-TSN]
Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "DetNet Data Plane: MPLS over IEEE 802.1 Time Sensitive Networking (TSN)", Work in Progress, Internet-Draft, draft-ietf-detnet-mpls-over-tsn-05, 13 December 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-mpls-over-tsn-05>>.
- [DetNet-Security]
Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", Work in Progress, Internet-Draft, draft-ietf-detnet-security-13, 11 December 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-security-13>>.

13>.

[DetNet-YANG]

Geng, X., Chen, M., Ryoo, Y., Fedyk, D., Rahman, R., and Z. Li, "Deterministic Networking (DetNet) Configuration YANG Model", Work in Progress, Internet-Draft, draft-ietf-detnet-yang-09, 16 November 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-yang-09>>.

[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>>.

[IEEE802.1CB-2017]

IEEE, "IEEE Standard for Local and metropolitan area networks-- Frame Replication and Elimination for Reliability", IEEE 802.1CB-2017, DOI 10.1109/IEEESTD.2017.8091139, October 2017, <<https://ieeexplore.ieee.org/document/8091139>>.

[RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <<https://www.rfc-editor.org/info/rfc2205>>.

[RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.

[RFC3272] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, DOI 10.17487/RFC3272, May 2002, <<https://www.rfc-editor.org/info/rfc3272>>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.

[RFC3985] Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, DOI 10.17487/RFC3985, March 2005, <<https://www.rfc-editor.org/info/rfc3985>>.

[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>>.

[RFC4448] Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, DOI 10.17487/RFC4448, April 2006, <<https://www.rfc-editor.org/info/rfc4448>>.

[RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<https://www.rfc-editor.org/info/rfc4875>>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440,

DOI 10.17487/RFC5440, March 2009,
<<https://www.rfc-editor.org/info/rfc5440>>.

- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010,
<<https://www.rfc-editor.org/info/rfc5920>>.
- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, DOI 10.17487/RFC5921, July 2010,
<<https://www.rfc-editor.org/info/rfc5921>>.
- [RFC6003] Papadimitriou, D., "Ethernet Traffic Parameters", RFC 6003, DOI 10.17487/RFC6003, October 2010,
<<https://www.rfc-editor.org/info/rfc6003>>.
- [RFC6073] Martini, L., Metz, C., Nadeau, T., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", RFC 6073, DOI 10.17487/RFC6073, January 2011,
<<https://www.rfc-editor.org/info/rfc6073>>.
- [RFC8306] Zhao, Q., Dhody, D., Ed., Palleti, R., and D. King, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", RFC 8306, DOI 10.17487/RFC8306, November 2017,
<<https://www.rfc-editor.org/info/rfc8306>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019,
<<https://www.rfc-editor.org/info/rfc8660>>.
- [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>>.

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, Jeong-dong Ryoo, and Carlos J. Bernardos for their various contributions to this work.

Contributors

The editor of this document wishes to thank and acknowledge the following person who contributed substantially to the content of this document and should be considered a coauthor:

Don Fedyk
LabN Consulting, L.L.C.

Email: dfedyk@labn.net

Authors' Addresses

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

Hungary

Email: balazs.a.varga@ericsson.com

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

Email: janos.farkas@ericsson.com

Lou Berger
LabN Consulting, L.L.C.

Email: lberger@labn.net

Andrew G. Malis
Malis Consulting

Email: agmalis@gmail.com

Stewart Bryant
Futurewei Technologies

Email: sb@stewartbryant.com

Jouni Korhonen

Email: jouni.nospam@gmail.com