

DetNet
Internet-Draft
Intended status: Standards Track
Expires: January 1, 2019

J. Korhonen, Ed.
B. Varga, Ed.
Ericsson
June 30, 2018

DetNet MPLS Data Plane Encapsulation
draft-ietf-detnet-dp-sol-mpls-00

Abstract

This document specifies Deterministic Networking data plane encapsulation solutions. The described data plane solutions is applied over an MPLS Packet Switched Networks.

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 January 1, 2019.

Copyright Notice

Copyright (c) 2018 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	3
2.	Terminology	3
2.1.	Terms used in this document	4
2.2.	Abbreviations	4
3.	Requirements language	6
4.	MPLS DetNet data plane overview	6
4.1.	DetNet data plane encapsulation requirements	12
5.	DetNet encapsulation	13
5.1.	End-system specific considerations	13
5.2.	DetNet domain specific considerations	15
5.2.1.	DetNet Layer Two Service	15
5.2.2.	DetNet Routing Service (IP over MPLS)	16
5.3.	DetNet Inter-Working Function (DN-IWF)	17
5.3.1.	Networks with multiple technology segments	17
5.3.2.	DN-IWF related considerations	18
6.	MPLS-based DetNet data plane solution	19
6.1.	DetNet over MPLS Encapsulation Components	19
6.2.	MPLS data plane encapsulation	21
6.3.	DetNet control word	22
6.4.	Flow Identification	23
6.5.	Indication of the DetNet Payload Type	23
6.6.	OAM Indication	24
6.7.	Flow Aggregation	24
6.7.1.	Aggregation at the LSP	25
6.7.2.	Aggregating DetNet flows as a new DetNet flow	25
6.7.3.	Simple Aggregation at the DetNet layer	26
6.8.	Service Layer Considerations	27
6.8.1.	Edge node processing	27
6.8.2.	Relay node processing	28
6.9.	Other DetNet data plane considerations	29
6.9.1.	Class of Service	29
6.9.2.	Quality of Service	30
6.9.3.	Cross-DetNet flow resource aggregation	31
6.9.4.	Layer 2 addressing and QoS Considerations	32
6.9.5.	Time Synchronization	32
7.	Management and control considerations	33
7.1.	MPLS-based data plane	33
7.1.1.	S-Label assignment and distribution	33
7.1.2.	Explicit routes	33
7.2.	Packet replication and elimination	34
7.3.	Congestion protection and latency control	35
7.4.	Bidirectional traffic	35
7.5.	Flow aggregation control	35
8.	DetNet IP Operation over DetNet MPLS Service	35
9.	IEEE 802.1 TSN Interconnection over DetNet MPLS Service	36
10.	DetNet MPLS Transport Layer Operation over IEEE 802.1 TSN	

Sub-Networks	36
11 . DetNet MPLS Transport Layer Operation over IP DetNet PSNs . .	36
12 . Security considerations	38
13 . IANA considerations	38
14 . Contributors	39
15 . Acknowledgements	41
16 . References	41
16.1 . Normative references	41
16.2 . Informative references	44
Appendix A . Example of DetNet data plane operation	45
Authors' Addresses	46

[1](#). Introduction

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

This document specifies the DetNet data plane and the on-wire encapsulation of DetNet flows over an MPLS-based Packet Switched Network (PSN). The specified encapsulation provides the building blocks to enable the DetNet service layer functions and allow flow identification as described in the DetNet Architecture.

The DetNet transport layer functionality that provides congestion protection for DetNet flows is assumed to be in place in a DetNet node.

Furthermore, this document also describes how DetNet flows are identified, and how a DetNet Relay/Edge/Transit nodes works. It also describes the function and operation of the Packet Replication (PRF) Packet Elimination (PEF) and Packet Ordering (POF) functions in the MPLS data plane.

This document does not define the associated control plane functions, or Operations, Administration, and Maintenance (OAM). It also does not specify traffic handling capabilities required to deliver congestion protection and latency control for DetNet flows at the DetNet transport layer.

[2](#). Terminology

2.1. Terms used in this document

This document uses the terminology established in the DetNet architecture [[I-D.ietf-detnet-architecture](#)] and the DetNet Data Plane Solution Alternatives [[I-D.ietf-detnet-dp-alt](#)].

T-Label	A label used to identify the LSP used to transport a DetNet flow across an MPLS PSN, e.g., a hop-by-hop label used between label switching routers (LSR).
S-Label	A DetNet "service" label that is used between DetNet nodes that implement also the DetNet service layer functions. An S-Label is also used to identify a DetNet flow at DetNet service layer.
PEF	A Packet Elimination Function (PEF) eliminates duplicate copies of packets received by an edge or a relay node to prevent excess packets flooding the network, or to prevent duplicate packets being sent out of the DetNet domain.
PRF	A Packet Replication Function (PRF) replicates DetNet flow packets in an edge or a relay node and forwards them to one or more next hops in the DetNet domain. The number of packet copies sent to each next hop is a DetNet Flow specific parameter at the node doing the replication.
POF	A Packet Order Function (POF) re-orders packets within a DetNet flow that are received out of order. This function may be implemented at an edge or a relay node.
PREOF	Collective name for Packet Replication, Elimination, and Ordering Functions.
d-CW	A DetNet Control Word (d-CW) is used for sequencing and identifying duplicate packets of a DetNet flow at the DetNet service layer.

2.2. Abbreviations

The following abbreviations used in this document:

AC	Attachment Circuit.
CE	Customer Edge equipment.
CoS	Class of Service.

CW	Control Word.
d-CW	DetNet Control Word.
DetNet	Deterministic Networking.
DF	DetNet Flow.
DN-IWF	DetNet Inter-Working Function.
L2	Layer 2.
L2VPN	Layer 2 Virtual Private Network.
L3	Layer 3.
LSR	Label Switching Router.
MPLS	Multiprotocol Label Switching.
MPLS-TE	Multiprotocol Label Switching - Traffic Engineering.
MPLS-TP	Multiprotocol Label Switching - Transport Profile.
MS-PW	Multi-Segment PseudoWire (MS-PW).
NSP	Native Service Processing.
OAM	Operations, Administration, and Maintenance.
PE	Provider Edge.
PEF	Packet Elimination Function.
PRF	Packet Replication Function.
PREOF	Packet Replication, Elimination and Ordering Functions.
POF	Packet Ordering Function.
PSN	Packet Switched Network.
PW	PseudoWire.
QoS	Quality of Service.
S-PE	Switching Provider Edge.

T-PE Terminating Provider Edge.

TSN Time-Sensitive Network.

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.

4. MPLS DetNet data plane overview

This document describes how DetNet flows are carried over MPLS networks. The DetNet Architecture, [[I-D.ietf-detnet-architecture](#)], decomposes the DetNet data plane into two layers: a service layer and a transport layer. The basic approach defined in this document supports the DetNet service layer based on existing pseudowire (PW) encapsulations and mechanisms, and supports the DetNet transport layer based on existing MPLS Traffic Engineering encapsulations and mechanisms. Background on PWs can be found in [[RFC3985](#)] and [[RFC3031](#)].




Figure 1: DetNet adaptation to MPLS data plane

The MPLS DetNet data plane approach defined in this document is shown in Figure 1. The service layer is supported by a DetNet control word (d-CW) which conforms to the Generic PW MPLS Control Word (PWMCW) defined in [[RFC4385](#)]. A d-CW identifying service label (S-Label) is also used. The transport layer is supported by one or labels (T-Labels).

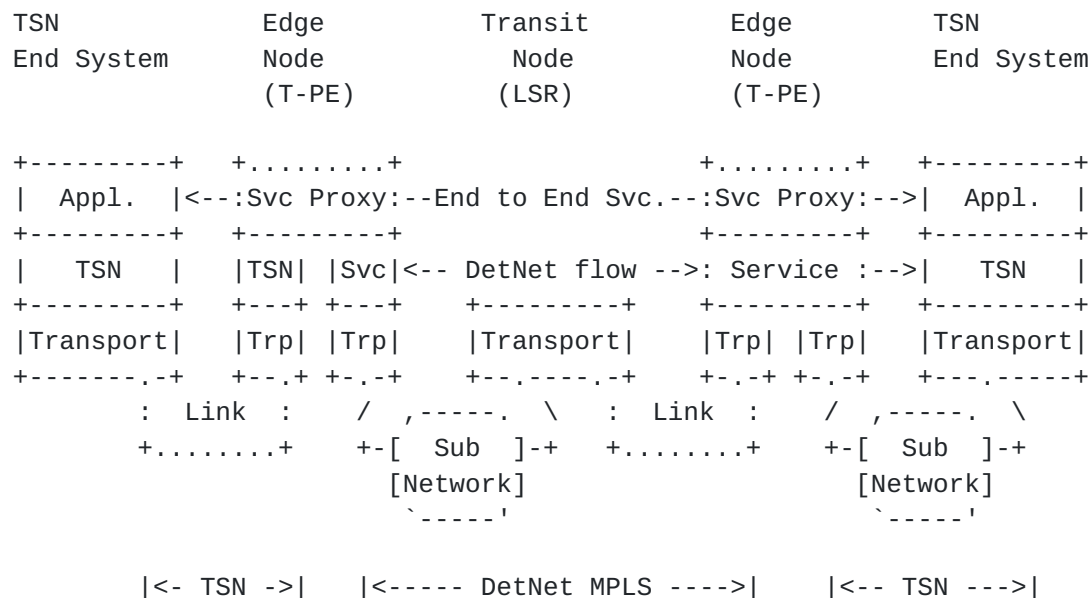


Figure 2: A TSN over DetNet MPLS Enabled Network

Figure 2 shows several node types defined in [\[I-D.ietf-detnet-architecture\]](#). DetNet Edge Nodes sit at the boundary of a DetNet domain. They are responsible for mapping non-DetNet aware traffic to DetNet services. They also support the imposition and disposition of the required DetNet encapsulation. These are functionally similar to pseudowire (PW) Terminating Provider Edge (T-PE) nodes which use MPLS-TE LSPs.

Transit nodes are normal MPLS Label Switching Routers (LSRs). They are generally unaware of the special requirements of DetNet flows, although they need to provide traffic engineering services and proper QoS to the LSPs associated with DetNet flows to enhance the prospect of the LSPs meeting the DetNet service requirements. Some implementations of transit nodes may be DetNet aware, but such nodes just support the DetNet transport layer.

The MPLS LSP may be provided by any MPLS method (provisioned, RSVP-TE, MPLS-TP, or MPLS Segment Routing (SR)).

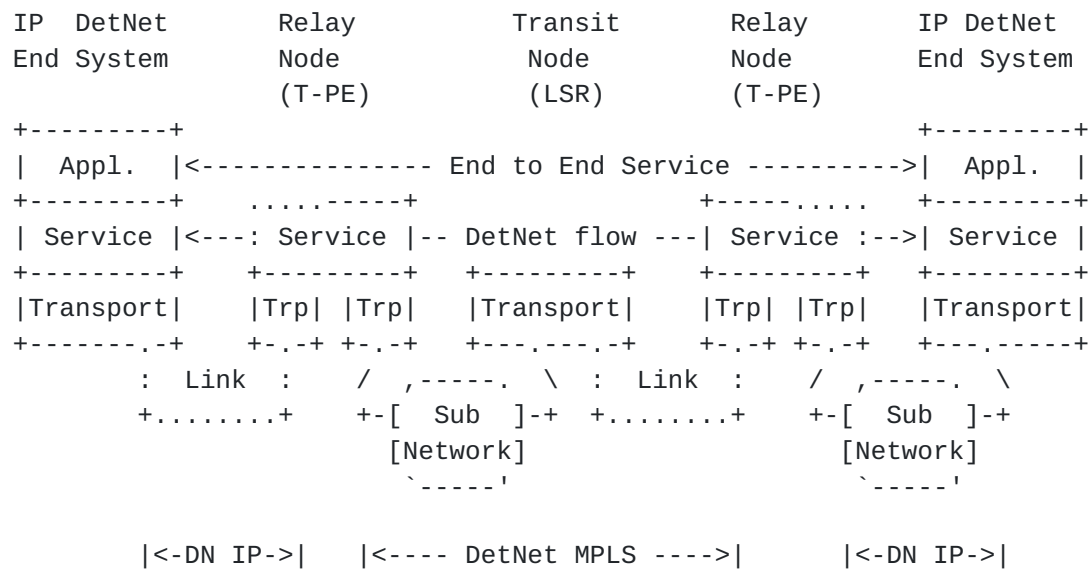


Figure 3: DetNet (DN) IP Over MPLS Network

Figure 3 and Figure 4, show different cases where relay nodes may be used. Relay nodes are similar to edge nodes in that both are aware of the needs of particular DetNet flows and take care to process them in accordance with the required performance needs. They differ in that relay nodes sit within a DetNet domain while edge nodes always sit at DetNet domain boundaries. Both node types can enhance the reliability of delivery by enabling the replication of packets so that multiple copies, possibly over multiple paths are forwarded through the DetNet domain. They also reduce the impact of replication by eliminating surplus copies of DetNet packets. Relay nodes may sit the boundary of an MPLS domain when the non-MPLS domain is DetNet aware. Relay nodes are functionally similar to PW S-PEs or, when at the edge of an MPLS network, T-PEs [RFC6073].

Figure 4 illustrates how DetNet can provide services for IEEE 802.1TSN end systems, CE1 and CE2, over a DetNet enabled network. The edge nodes, E1 and E2, insert and remove required DetNet data plane encapsulation. The 'X' in the edge nodes and relay node, R1, represent a potential DetNet flow packet replication and elimination point. This conceptually parallels L2VPN services, and could leverage existing related solutions as discussed below.

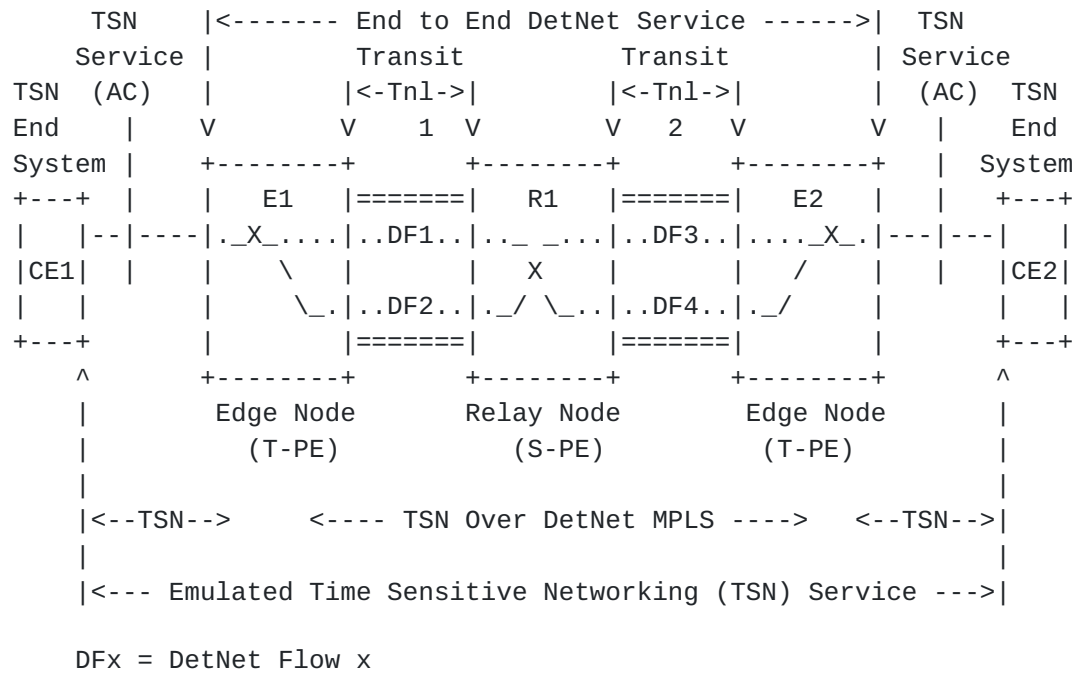


Figure 4: IEEE 802.1TSN over DetNet

Figure 5 illustrates how an end to end MPLS-based DetNet service is provided in a more detail. In this case, the end systems, CE1 and CE2, are able to send and receive DetNet flows, and R1 and R2 are relay nodes as they sit in the middle of a DetNet network. For example, an end system sends data encapsulated in MPLS. The 'X' in the end systems, and relay nodes represents potential DetNet flow packet replication and elimination points. Here the relay nodes may change the underlying transport, for example tunneling MPLS over IP [Section 11](#), or simply interconnect network segments.




Figure 5: MPLS-Based Native DetNet

Figure 6 illustrates how an end to end MPLS-based DetNet service is provided where the end systems are not able to send and receive DetNet flows. In this example, the nodes labeled CE1 and CE2 could be non-DetNet aware IP routers or hosts. Note that E1 and E2 are edge nodes as they sit boundaries of the DetNet enabled domain.




Figure 6: MPLS-Based DetNet (non-MPLS End System)

Figure 7 illustrates how end to end DetNet service is provided where the end systems are able to send and receive IP DetNet flows, e.g.,

DetNet				DetNet					
IP	Service	Transit		Transit		Service	IP		
DetNet		<-Tn1->		<-Tn1->			DetNet		
End		V	1	V	V	2	V		End
System		+-----+		+-----+		+-----+			System
+---+		R1	=====	R2	=====	R3		+---+	
		..X..	..DF1..	.._ _..	..DF3..	...X_			
CE1		\		X		/		CE2	
		_	..DF2..	../_ _	..DF4..	../_			
+---+		=====		=====				+---+	
^	+-----+		+-----+		+-----+		^		
	Relay Node		Relay Node		Relay Node				
	(T-PE)		(S-PE)		(T-PE)				
<-DN IP->	<----- DetNet MPLS ----->				<-DN IP->				
<-----	End to End DetNet Service							>----->	

Figure 7: DetNet IP over DetNet (DN) MPLS

```

      1      1.1      1.1      1.2.1      1.2.1      1.2.2
CE1-----EN1-----R1-----R2-----R3-----EN2-----CE2
      \           1.2.1 /
      \1.2      /-----+
      +-----R4-----+
              1.2.2

```

Figure 8: Example Packet flow in DetNet Enabled MPLS Network

Customer Equipment CE1 sends a packet into the DetNet enabled MPLS network. This is packet (1). Edge Node EN1 encapsulates the packet

as a DetNet Packet and sends it to Relay node R1 (packet 1.1). EN1 makes a copy of the packet (1.2), encapsulates it and sends this copy to Relay node R4.

Note that along the MPLS path from EN1 to R1 there may be zero or more LSRs which, for clarity, are not shown. The same is true for any other path between two DetNet entities shown in Figure 8.

Relay node R4 has been configured to send one copy of the packet to Relay Node R2 (packet 1.2.1) and one copy to Edge Node EN2 (packet 1.2.2).

R2 receives packet copy 1.2.1 before packet copy 1.1 arrives, and, having been configured to perform packet elimination on this DetNet flow, forwards packet 1.2.1 to Relay Node R3. Packet copy 1.1 is of no further use and so is discarded by R2.

Edge Node EN2 receives packet copy 1.2.2 from R4 before it receives packet copy 1.2.1 from R2 via relay Node R3. EN2 therefore strips any DetNet encapsulation from packet copy 1.2.2 and forwards the packet to CE2. When EN2 receives the later packet copy 1.2.1 this is discarded.

The above is of course illustrative of many network scenarios that can be configured. Between a pair of relay nodes there may be one or more transport nodes that simply forward the DetNet traffic, but these are omitted for clarity.

4.1. DetNet data plane encapsulation requirements

Two major groups of scenarios can be distinguished which require flow identification during transport:

1. DetNet function related scenarios:

- * Congestion protection and latency control: usage of allocated resources (queuing, policing, shaping).
- * Explicit routes: select/apply the flow specific path.
- * Service protection: recognize DetNet compound and member flows for replication and elimination.

2. OAM function related scenarios:

- * troubleshooting (e.g., identify misbehaving flows, etc.)

- * recognize flow(s) for analytics (e.g., increase counters, etc.)
- * correlate events with flows (e.g., volume above threshold, etc.)
- * etc.

The DetNet data plane allows for the aggregation of DetNet flows, e.g., via MPLS hierarchical LSPs, to improved scaling. When DetNet flows are aggregated, transit nodes may have limited ability to provide service on per-flow DetNet identifiers. Therefore, identifying each individual DetNet flow on a transit node may not be achieved in some network scenarios, but DetNet service can still be assured in these scenarios through resource allocation and control.

A node operating on a DetNet flow in the Detnet layer, i.e. a node processing a DetNet packet which has the S-label as top of stack uses the local context associated with that S-label to determine what local operation(s) are applied to that packet. The S-label has to be unique on each edge and relay node, which is achieved by using a label taken from the platform label space [[RFC3031](#)].

5. DetNet encapsulation

5.1. End-system specific considerations

Data-flows requiring DetNet service are generated and terminated on end-systems. Encapsulation depends on application and its preferences. In a DetNet (or even a TSN) domain the DN (TSN) functions use at most two flow parameters, namely Flow-ID and Sequence Number. However, an application may exchange further flow related parameters (e.g., time-stamp), which are not considered by DN functions.

Two types of end-systems are distinguished:

- o L2 (Ethernet) end-system: application directly over L2.
- o L3 (IP) end-system: application over L3.

In case of Ethernet end-systems the application data is encapsulated directly in L2. From the DN domain perspective no upper layer protocols are visible. The Data-flow uses only Ethernet tag(s) and further flow specific parameters (if needed) are hidden inside the protocol data unit (PDU).

The IP end-system scenario is different. Data-flows are encapsulated directly in L3 (i.e., IP) and the application may use further upper layer protocols (e.g., Real-time Transport Protocol (RTP)). Many valid combinations exist, and it may be application specific how the IP header fields are used. Also, usage of further upper layer protocols depends on application requirements (e.g., time-stamp). See [[I-D.ietf-detnet-dp-sol-ip](#)] more details.

[Editor's note: IP solution document does not really detail anything beyond 6-tuple.]

As a general rule, DetNet domains MUST be capable of forwarding any Data-flows and the DetNet domain MUST NOT mandate the end-system encapsulation format.

Furthermore, no application-level-proxy function is envisioned inside the DetNet domain, so end-systems peer with end-systems using the same application encapsulation format (see figure below):

- o L2 end-systems peer with L2 end-systems and
- o L3 end-systems peer with L3 end-systems.




Figure 9: End-systems and the DetNet domain

5.2. DetNet domain specific considerations

From a connection type perspective, three scenarios are distinguished:

1. Directly attached: end-system is directly connected to an edge node.
2. Indirectly attached: end-system is behind a (L2-TSN / L3-DetNet) sub-network.
3. DN integrated: end-system is part of the DetNet domain.

L3 end-systems may use any of these connection types, however L2 end-systems may use only the first two (directly or indirectly attached). DetNet domain MUST allow communication between any end-systems of the same type (L2-L2, L3-L3), independent of their connection type and DetNet capability. However directly attached and indirectly attached end-systems have no knowledge about the DetNet domain and its encapsulation format at all. See Figure 10 for L3 end-system scenarios.


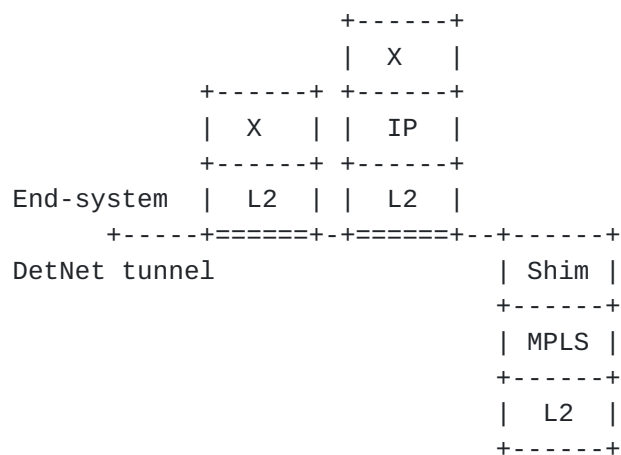


Figure 10: Connection types of L3 end-systems

5.2.1. DetNet Layer Two Service

The simplest DetNet service is to provide tunneling for layer two, where the connected hosts are in the same broadcast (BC) domain. Forwarding over the DetNet domain is based on L2 (MAC) addresses (i.e. dst-MAC), or on received interface [\[RFC3985\]](#). In both cases the L2 headers MUST either be kept, or provision must be made for their reconstruction at egress from the DetNet domain.



Examples:

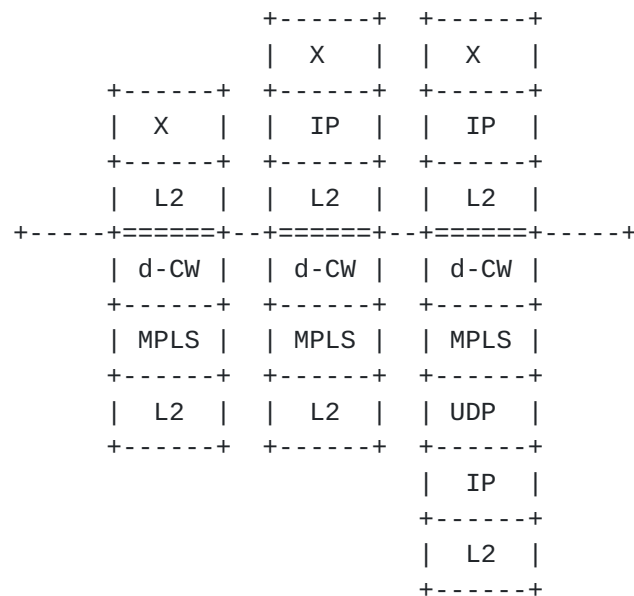



Figure 11: Encapsulation format for DetNet Layer Two Service

As shown in Figure 11 both L2 and L3 end-systems can be served by such a DetNet L2 encapsulation service. This encapsulation service may be carried over MPLS natively [Section 6.2](#), or over MPLS over IP [Section 11](#).

5.2.2. DetNet Routing Service (IP over MPLS)

IP traffic and IP DetNet flows, see [[I-D.ietf-detnet-dp-sol-ip](#)], can be carried over a DetNet MPLS domain. In such cases, the IP headers are modified per standard router behavior, e.g., TTL handling.

Figure 12 shows the encapsulation of an IP flow over MPLS as well as when MPLS is carried over an IP PSN, see [Section 11](#).



Examples:

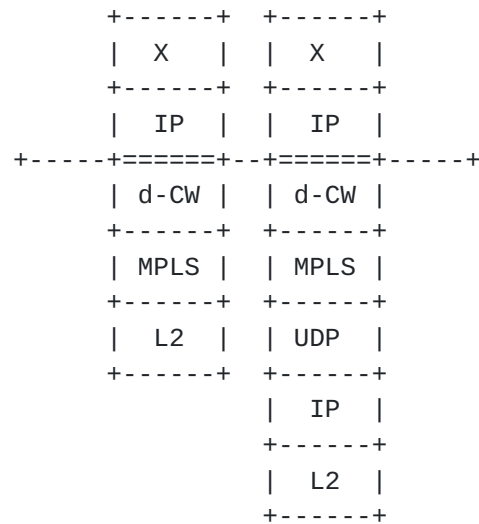


Figure 12: Encapsulation format for DetNet Routing in MPLS PSN for L3 end-systems

[5.3.](#) DetNet Inter-Working Function (DN-IWF)

[5.3.1.](#) Networks with multiple technology segments

There are networking scenarios, where the DetNet domain contains multiple technology segments (IP, MPLS, ..) and all those segments are under the same administrative control (see Figure 13). Furthermore, DetNet nodes may be interconnected via TSN segments.

An important aspect of DetNet network design is the placement of DetNet functions across the domain. Designs based on segment-by-segment optimization can provide only sub-optimal solutions. In order to achieve global optimized Inter-Working Functions (DN-IWF) can be placed at segment edge nodes, which stitch together DetNet flows across connected segments.

DN-IWF may ensure that flow attributes are correlated across segment edges. For example, there are two DetNet functions which require Sequence Numbers: (1) PEF: removes duplications from flows and (2) POF: ensures in-order-delivery of packet in a flow. Stitching flows together and correlating attributes means for example that replication of packets can happen in one segment and elimination of duplicates in a different one.

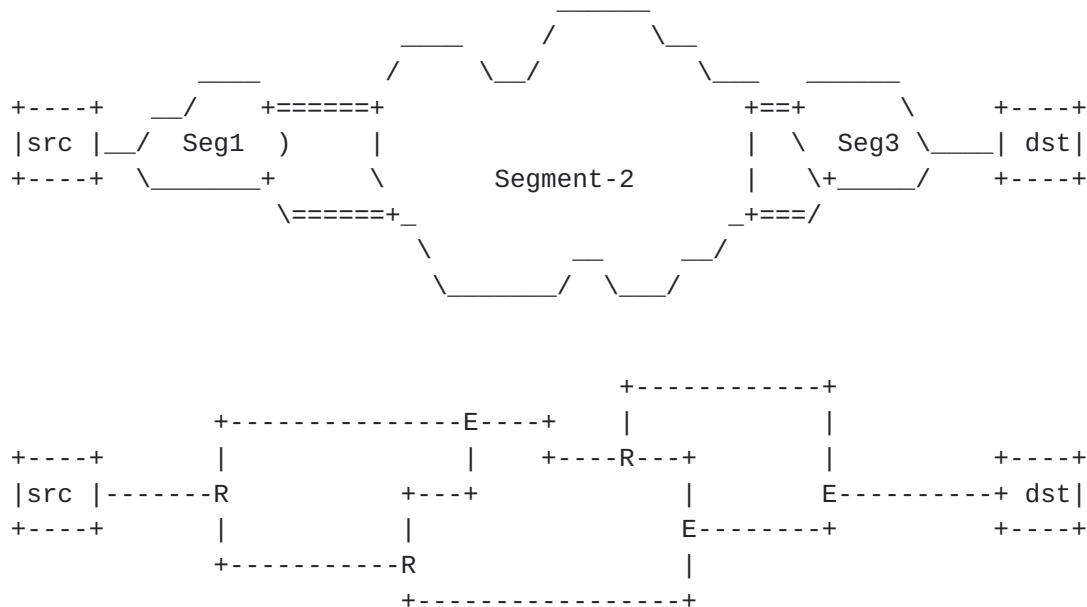


Figure 13: Optimal replication and elimination placement across technology segments example

5.3.2. DN-IWF related considerations

The goal of DN-IWF is to (1) match and (2) translate segment specific flow attributes. The DN-IWF ensures that segment specific attributes comprise per domain unique attributes for the whole DetNet domain. This characteristic can ensure that DetNet functions can be based on per domain attributes and not per segment attributes.

The two DetNet specific attributes have the following characteristics:

- o Flow-ID: it is same in all packets of a flow
- o Sequence Number: it is different packet-by-packet

For the Flow-ID the DN-IWF can implement a static mapping. The situation is more complicated for Sequence Number as it is different packet-by-packet, so it may need more sophisticated translation unless its format is exactly the same in the two technology segments. In this later case the DN-IWF can simple copy the Sequence Number field between the tunneling encapsulation of the two technology segments.

In case of three technology segments (IP, MPLS and TSN) three DN-IWF functions can be specified. In the rest of this section the focus is on the (1) IP - MPLS network scenario. Note: the use-cases are out-of-scope for (2) TSN - IP, (3) TSN - MPLS.

Simplest implementation of DN-IWF is provided if the flow attributes have the same format. Such a common denominator of the tunnel encapsulation format is the pseudowire encapsulation over both IP and MPLS.




Figure 14: FIGURE Placeholder PW over X

[Editor's note: Where is the text describing how 802.1 TSN Streams are mapping to DetNet services/flows. i.e., EVPN+]

6. MPLS-based DetNet data plane solution

6.1. 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 group 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), similar to a pseudowire (PW) label [[RFC3985](#)], is used to identify both the DetNet flow identity and the payload 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 sequence number is carried in the DetNet Control word which carries the Data/OAM discriminator. The LSP used to transport the DetNet packet may be of any type (MPLS-LDP, MPLS-TE, MPLS-TP [[RFC5921](#)], or MPLS-SR [[I-D.ietf-spring-segment-routing-mpls](#)]). The LSP (T-Label) label and/or the S-Label may be used to indicate the queue processing as well as the forwarding parameters.

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 ensure that in high speed networks the sequence number space does not wrap whilst packets are in flight. In addition it must be possible to send a packet with a zero length sequence number, to support the case where sequence numbers are not required by a particular DetNet flow.

Note that the concept of a zero length sequence number is not to be confused with a sequence number of zero. For example, were the sequence number size is 16 bits, the sequence will contain: 65535, 0, 1. In this case zero is an ordinary sequence number. Unlike [[RFC4448](#)] a sequence number of zero does not indicate that no sequence number is in use. Where sequence numbers are not in use, and thus a zero length sequence number is in used, the sequence number field in the packet is sent as zero. The DetNet packet forwarder knows which of these cases applies through configuration parameters associated with each specific DetNet flow.

Note that when the network consists only of DetNet enabled nodes with no aggregation, Penultimate Hop Popping (PHP) means that the only label in the label stack may be the S-label.

6.2. MPLS data plane encapsulation

Figure 15 illustrates a DetNet data plane MPLS encapsulation. The MPLS-based encapsulation of the DetNet flows is a good fit for the Layer-2 interconnect deployment cases (see Figure 4). Furthermore, end to end DetNet service i.e., native DetNet deployment (see Figure 5) is also possible if DetNet end systems are capable of initiating and termination MPLS encapsulated packets.

The MPLS-based DetNet data plane encapsulation consists of:

- o DetNet control word (d-CW) containing sequencing information for packet replication and duplicate elimination purposes, and the OAM indicator. There MUST be a separate sequence number space for each DetNet flow.
- o DetNet service Label (S-label) that identifies a DetNet flow to the peer node that is to process it. The S-Label is allocated from the platform label space [[RFC3031](#)].
- o Zero or more MPLS transport LSP label(s) (T-label) used to direct the packet along the label switched path (LSP) to the next peer node along the path. When Penultimate Hop Popping is in use there may be no label T-label in the protocol stack on the final hop.
- o The necessary data-link encapsulation is then applied prior to transmission over the physical media.

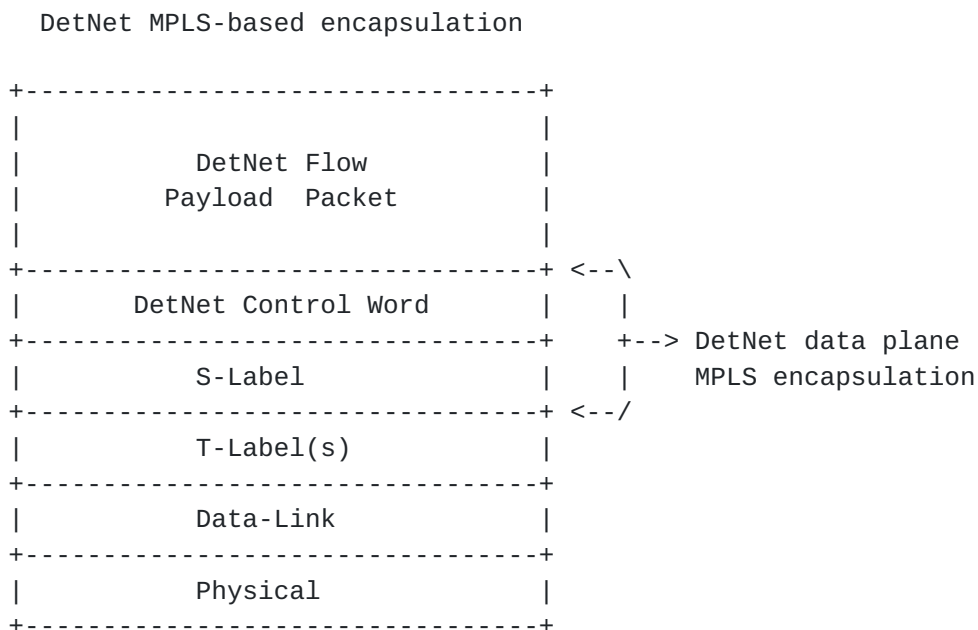


Figure 15: Encapsulation of a DetNet flow in an MPLS(-TP) PSN

6.3. DetNet control word

A DetNet control word (d-CW) conforms to the Generic PW MPLS Control Word (PWMCW) defined in [RFC4385] and is illustrated in Figure 16. The upper nibble of the d-CW MUST be set to zero (0). Two sequence number sizes are supported: 16 bits and 28 bits. The sequence number size in use for the d-CW associated with a DetNet flow (S-Label) is configured either by a control plane or manually for each DetNet flow. The sequence number is aligned to the right (least significant bits) and unused bits MUST be set to zero (0). Each DetNet flow MUST have its own sequence number counter. The sequence number is incremented by one for each new packet.

As discussed in [Section 6](#), zero is an ordinary sequence number with no special meaning. Also as discussed therein, where no sequence number is used by a particular DetNet flow, the sequence number field in the d-CW is set to zero.

The d-CW MUST always be present in a packet. In a case where the sequence number is not used (e.g., for DetNet-t-flows) a zero length sequence number is used and the sequence number MUST be set to zero (0).

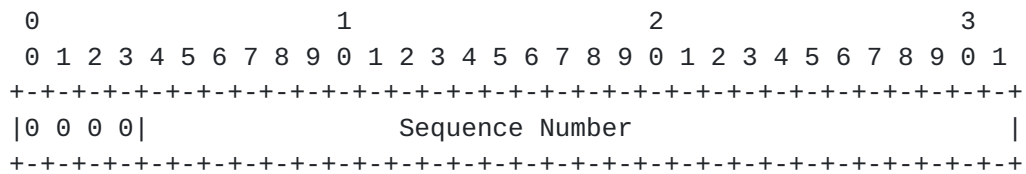


Figure 16: DetNet Control Word

6.4. Flow Identification

DetNet flow identification at a DetNet service layer is realized by an S-label. The S-label is allocated from the platform label space [RFC3031] which means that the DetNet flow is correctly identified and matched to the flow parameters, including the flow history, regardless of which input interface the packet arrives on. The S-label MUST be at the bottom label of the label stack for a DetNet-s- or DetNet-st-flow and MUST precede the d-CW.

The S-label for a specific DetNet flow is unique to that DetNet flow on a specific node, but is not required to be identical with the S-label for that DetNet flow in any other node within the DetNet domain. Thus the S-label can only be used to identify the DetNet flow at the intended receiving node.

6.5. Indication of the DetNet Payload Type

The only nodes that needs to know the payload type of a flow are the DetNet ingress node and the DetNet egress nodes. The ingress node has to know how to process the packet it receives from the ingress AC or IP flow, and the egress edge node has to know how to prepare the packet for transmission to the next hop.

On ingress a DetNet edge node has to classify the packets into those that are for transmission as Detnet packets and those that are for transmission as "normal" packets at one of more lower priorities. The packet type is indicated to the egress edge node through the value of the S-label. Thus, when the egress edge node looks up the S-label one of the parameters returned is the packet type which in turn tells the egress edge node how to prepare the packet for transmission to a next hop.

The consequence of this approach is that if multiple packet encapsulations are processed on a node pair, each encapsulation will need its own S-Label. That is not generally a problems, since it is anticipated that only one encapsulation type will be present for each DetNet flow. Of course, if for some reason the multiple encapsulations are needed to support a single DetNet service,

multiple S-labels will be required for that service. Note that in the unlikely case that IPv4 and IPv6 will map to the same DetNet flow, different S-labels will be needed to differentiate between the versions of IP.

6.6. 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-DW 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 for more information DetNet OAM.

6.7. Flow Aggregation

1. Aggregate at the LSP (Transport)
2. Aggregating DetNet flows as a new DetNet flow
3. Simple Aggregation at the DetNet layer

A further method of using SR to perform aggregation is for further study.

The resource control and management aspects of aggregation (including the queuing/shaping/ policing implications) will be covered in other documents.

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 three methods of introducing flow aggregation:

The following review comments were received when this section was committed to github.

General comment: We should points to the major issue of aggregation, namely the Seq.Num related problem. The aggregated flows have their

own Seq.Num and those are independent. We should consider to group the aggregation techniques as per their impact on what DetNet functions they allow on a DetNet flow. (E.g., aggregation without new Aggregate.Seq.Num would prohibit usage of FR, EF and in-order-delivery function on the aggregate flow).

SR based aggregation can be treated as a form of H-LSP aggregation. Should we differentiate them? What are the differences?

What are the issues when aggregating of different payload types? Should we add an editor note on this?

Simple-aggregation-at-the-detnet-layer: is this not the same as H-LSP? The A-label can be treated just as an additional T-label.

End of review comment.

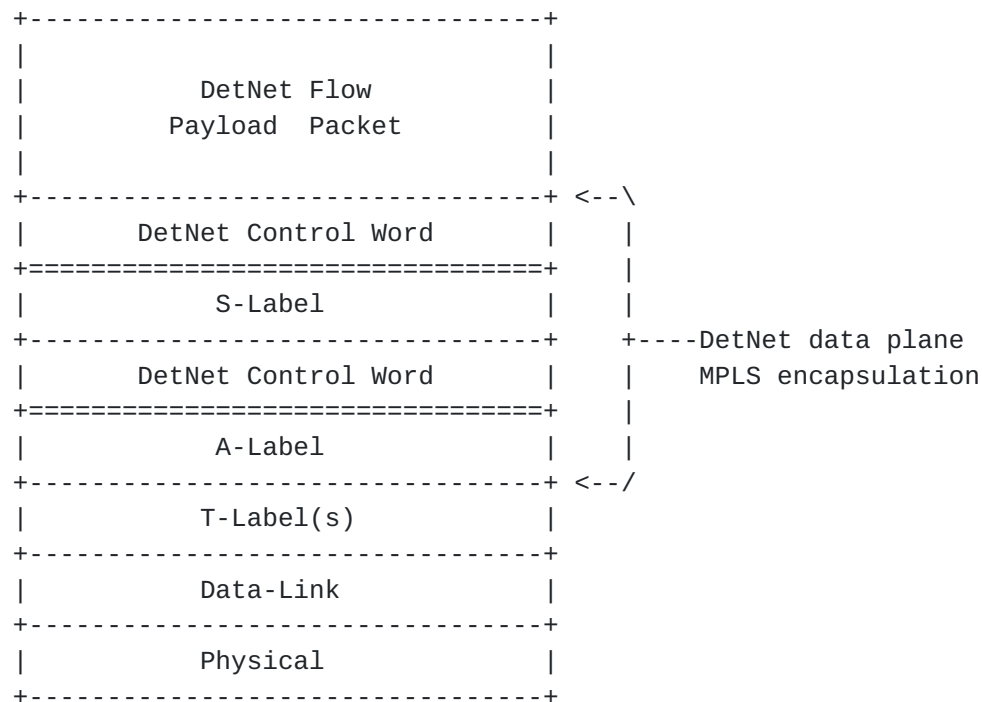
6.7.1. Aggregation at the LSP

DetNet flows transported 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 falls out of the definition for hierarchy and the MPLS label stack [[RFC3032](#)]. DetNet nodes which 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 TC field, i.e., L-LSPs or E-LSPs. Such nodes will need to ensure that traffic from aggregated LSPs are placed (shaped/policed/enqueued) onto the H-LSPs in a fashion that ensures 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 descriptions mentioned above or in separate future documents. Management and control plane mechanisms will also need to ensure that the service required on the aggregate flow (H-LSP or DSCP) are provided, which may include the discarding or remarking mentioned in the previous sections.

6.7.2. Aggregating DetNet flows as a new DetNet flow

An aggregate can be built by layering DetNet flows as shown below:




Both the Aggregation (A) label and the S-label have their MPLS S bit set indicating bottom of stack, and the d-CW allows the PREOF to work.

It is a property of the A-label that what follows is d-CW followed by an S-label. A relay node processing the A-label would not know the underlying payload type. 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.

6.7.3. Simple Aggregation at the DetNet layer

Another approach would be not to include a d-CW for the aggregated flow. This would be functionally similar to aggregation at the transport layer using H-LSPs, but would confine knowledge of the aggregation to the DetNet layer. Such an approach shares the disadvantage that PREOF operations would not be possible. OAM operation in this mode is for further study.



6.8. Service Layer Considerations

The edge and relay node internal procedures related to PREOF are implementation specific. The order of a packet elimination or replication is out of scope in this specification. However, care should be taken that the replication function does not actually loopback packets as "replicas". Looped back packets include artificial delay when the node that originally initiated the packet receives it again. Also, looped back packets may make the network condition to look healthier than it actually is (in some cases link failures are not reflected properly because looped back packets make the situation appear better than it actually is).

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, but given the sensitivity of applications to packet latency the impact on the DetNet application would be severe.

6.8.1. Edge node processing

An edge node is responsible for matching ingress packets to the service they require and encapsulating them accordingly. An edge node may participate in the packet replication and duplication elimination.

The DetNet-aware forwarder selects the egress DetNet member flow segment based on the flow identification. The mapping of ingress DetNet member flow segment to egress DetNet member flow segment may be statically or dynamically configured. Additionally the DetNet-aware forwarder does duplicate frame elimination based on the flow identification and the sequence number combination. The packet replication is also done within the DetNet-aware forwarder. During elimination and the replication process the sequence number of the DetNet member flow MUST be preserved and copied to the egress DetNet member flow.

The internal design of a relay node is out of scope of 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 the PREOF functions must be applied, and to maintain that state is such as 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).

[Editor's note: I think the rest of this section belongs in a new "802.1 TSN (island Interconnect) over MPLS DetNet" section.]

This may be done in the DetNet layer, or where the native service processing (NSP) [[RFC3985](#)] is IEEE 802.1CB [[IEEE8021CB](#)] capable, the packet replication and duplicate elimination MAY entirely be done in the NSP, bypassing the DetNet flow encapsulation and logic entirely. This enables operating over unmodified implementations and deployments. The NSP approach works only between edge nodes and cannot make use of relay nodes.

The NSP approach is useful end to end tunnel and for for "island interconnect" scenarios. However, when there is a need to do PREOF in a middle of the network, such plain edge to edge operation is not sufficient.

The extended forwarder MAY copy the sequencing information from the native DetNet packet into the DetNet sequence number field and vice versa. If there is no existing sequencing information available in the native packet or the forwarder chose not to copy it from the native packet, then the extended forwarder MUST maintain a sequence number counter for each DetNet flow (indexed by the DetNet flow identification).

[6.8.2. Relay node processing](#)

A DetNet Relay node operates in the DetNet transport layer . This processing is done within an extended forwarder function. Whether an ingress DetNet member flow receives DetNet specific processing

depends on how the forwarding is programmed. Some relay nodes may be DetNet service aware, while others may be unmodified LSRs that only understand how to switch MPLS-TE LSPs.

It is also possible to treat the relay node as a transit node, see [Section 6.9.3](#). Again, this is entirely up to how the forwarding has been programmed.

6.9. Other DetNet data plane considerations

6.9.1. Class of Service

[Editor's note: this section needs to be updated to discuss how DetNet service is mapped to E- and L-LSPs. Perhaps this gets merged with the aggregation section or dropped?]

Class and quality of service, i.e., CoS and 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 aggregate group basis and QoS is used to refer to mechanisms that provide traffic forwarding treatment based on a specific DetNet flow basis. Examples of existing network level CoS mechanisms include DiffServ which is enabled by 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 ECN MAY also be used as defined in ECN [[RFC5129](#)] and updated by [[RFC5462](#)].

CoS for DetNet flows carried in IPv6 is provided using the standard differentiated services code point (DSCP) field [[RFC2474](#)] and related mechanisms. The 2-bit explicit congestion notification (ECN) [[RFC3168](#)] field MAY also be used.

One additional consideration for DetNet nodes which support CoS services is that they MUST ensure that the CoS service classes do not impact the congestion protection and latency control mechanisms used to provide DetNet QoS. This requirement is similar to requirement for MPLS LSRs to that CoS LSPs do not impact the resources allocated to TE LSPs via [[RFC3473](#)].

6.9.2. Quality of Service

Quality of Service (QoS) mechanisms for flow specific traffic treatment typically includes 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 Resource ReSerVation Protocol (RSVP) [[RFC2205](#)] (RSVP) and RSVP-TE [[RFC3209](#)] and [[RFC3473](#)]. The existing MPLS mechanisms defined to support CoS [[RFC3270](#)] can also be used to reserve resources for specific traffic classes.

In addition to explicit routes, and packet replication and elimination, described in [Section 6](#) above, DetNet provides zero congestion loss and bounded latency and jitter. As described in [[I-D.ietf-detnet-architecture](#)], there are different mechanisms that maybe used separately or in combination to deliver a zero congestion loss service. These mechanisms are provided by the either the MPLS or IP layers, and may be combined with the mechanisms defined by the underlying network layer such as 802.1TSN.

A baseline set of QoS capabilities for DetNet flows carried in PWS and MPLS can provided by MPLS with Traffic Engineering (MPLS-TE) [[RFC3209](#)] and [[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 Quality of 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 use to support the identification of flows requiring DetNet QoS.

Packets that are marked with a DetNet Class of Service value, but that have not been the subject of a completed reservation, 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:

- o MUST defend the DetNet QoS by discarding or remarking (to a non-DetNet CoS) packets received that are not the subject of a completed reservation.
- o MUST NOT use a DetNet reserved resource, e.g. a queue or shaper reserved for DetNet flows, for any packet that does not carry a DetNet Class of Service marker.

6.9.3. Cross-DetNet flow resource aggregation

[Editor's NOTE: keep and extend this section.]

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. This document identifies the traffic identification related aspects of aggregation of DetNet flows. The resource control and management aspects of aggregation (including the queuing/shaping/policing implications) will be covered in other documents. The data plane implications of aggregation are independent for PW/MPLS and IP encapsulated DetNet flows.

DetNet flows transported 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 falls out of the definition for hierarchy and the MPLS label stack [[RFC3032](#)]. DetNet nodes which 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 TC field, i.e., L-LSPs or E-LSPs. Such nodes will need to ensure that traffic from aggregated LSPs are placed (shaped/policed/enqueued) onto the H-LSPs in a fashion that ensures the required DetNet service is preserved.

DetNet flows transported via IP have more limited aggregation options, due to the available traffic flow identification fields of the IP solution. One available approach is to manage the resources associated with a DSCP identified traffic class and to map (remark) individually controlled DetNet flows onto that traffic class. This approach also requires that nodes support aggregation ensure that traffic from aggregated LSPs are placed (shaped/policed/enqueued) in a fashion that ensures the required DetNet service is preserved.

In both the MPLS and IP cases, additional details of the traffic control capabilities needed at a DetNet-aware node may be covered in the new service descriptions mentioned above or in separate future documents. Management and control plane mechanisms will also need to ensure that the service required on the aggregate flow (H-LSP or DSCP) are provided, which may include the discarding or remarking mentioned in the previous sections.

6.9.4. Layer 2 addressing and QoS Considerations

[Editor's NOTE: review and simplify this section.]

The Time-Sensitive Networking (TSN) Task Group of the IEEE 802.1 Working Group have defined (and are defining) a number of amendments to IEEE 802.1Q [[IEEE8021Q](#)] that provide zero congestion loss and bounded latency in bridged networks. IEEE 802.1CB [[IEEE8021CB](#)] defines packet replication and elimination functions that should prove both compatible with and useful to, DetNet networks.

As is the case for DetNet, a Layer 2 network node such as a bridge may need to identify the specific DetNet flow to which a packet belongs in order to provide the TSN/DetNet QoS for that packet. It also will likely need a CoS marking, such as the priority field of an IEEE Std 802.1Q VLAN tag, to give the packet proper service.

Although the flow identification methods described in IEEE 802.1CB [[IEEE8021CB](#)] are flexible, and in fact, include IP 5-tuple identification methods, the baseline TSN standards assume that every Ethernet frame belonging to a TSN stream (i.e. DetNet flow) carries a multicast destination MAC address that is unique to that flow within the bridged network over which it is carried. Furthermore, IEEE 802.1CB [[IEEE8021CB](#)] describes three methods by which a packet sequence number can be encoded in an Ethernet frame.

Ensuring that the proper Ethernet VLAN tag priority and destination MAC address are used on a DetNet/TSN packet may require further clarification of the customary L2/L3 transformations carried out by routers and edge label switches. Edge nodes may also have to move sequence number fields among Layer 2, PW, and IPv6 encapsulations.

6.9.5. Time Synchronization

[Editor's Note: A detailed discussion of time synchronization is outside the scope of this document, and the production of a specialist text discussing this topic is encouraged. This section will be updated/removed if such a document is available before publication of this text.]

Time synchronization is important both from the perspective of operating the DetNet network itself and from the perspective of transferring time across the network between client applications. Some clients may be able to use the DetNet as their provider of time and frequency, others may require the DetNet to transfer time between a client clock source and a client clock user.

The reader's attention is drawn to [\[RFC8169\]](#) which describes a method of recording the packet queuing time in an MPLS LSR on a packet by per packet basis and forwarding this information to the egress edge system. This allows compensation for any variable packet queuing delay to be applied at the packet receiver. The mechanism described in [\[RFC8169\]](#) may have wider application than basic time transfer in a DetNet.

A more detailed discussion of time synchronization is outside the scope of this document.

[7.](#) Management and control considerations

[Editor's note: This section needs to be different for MPLS and IP solutions. Most solutions are technology dependant. Currently most text in this section is just a draft and may have bits that are already moved to other places/documents.]

While management plane and control planes are traditionally considered separately, from the Data Plane perspective there is no practical difference based on the origin of flow provisioning information. This document therefore does not distinguish between information provided by a control plane protocol, e.g., RSVP-TE [\[RFC3209\]](#) and [\[RFC3473\]](#), or by a network management mechanisms, e.g., RestConf [\[RFC8040\]](#) and YANG [\[RFC7950\]](#).

[Editor's note: This section is a work in progress. discuss here what kind of enhancements are needed for DetNet and specifically for PREOF and DetNet zero congest loss and latency control. Need to cover both traffic control (queuing) and connection control (control plane).]

[7.1.](#) MPLS-based data plane

[7.1.1.](#) S-Label assignment and distribution

[Editor's note: Outdated and needs more work.]

The DetNet S-Label distribution follows the same mechanisms specified for XYZ . The details of the control plane protocol solution required for the label distribution and the management of the label number space are out of scope of this document.

[7.1.2.](#) Explicit routes

It is necessary to consider explicit routes both at the DetNet layer and in the MPLS layer. In the DetNet layer the explicit route consists of the set of Relay Nodes that the DetNet flow must

traverse. In the MPLS layer the explicit route consists of the set of LSRs, links, and possibly link bundle members and queues that the DetNet packets of a flow must traverse between nodes in the DetNet layer (i.e. between a specific Edge Node and the next hop Relay Node, between specific Relay Nodes, and between a specific Relay node and the egress Edge Node. This detailed steering is needed to ensure that packets are routed through the resources that have been reserved for them, and hence provide the DetNet application with the required performance.

Whether configuring, calculating and instantiating this is a multi-stage process, or a single stage process is out of scope of this document.

The one method of explicitly setting up the explicit path at the DetNet layer is through the use of the management controller.

[Editor's note: a method of setting up a graph through the DetNet Nodes using the IGP has been proposed. A reference is needed to e.g., [RFC 7813](#) IS-IS Path Control and Reservation.]

There are a number of approaches that can be taken to provide explicit routes/paths in the MPLS layer:

- o The path can be explicitly set up by the management controller calculating the path and explicitly configuring each node along that path.
- o The LSP can be set up using RSVP-TE. Such an approach confines the packet to the explicit path.
- o The path can be implemented using segment routing.

Where the DetNet traffic is carried over IP [Section 11](#) explicit paths may need to be provided in the IP layer. This is for further study.

[7.2.](#) Packet replication and elimination

[Editor's note: Outdated and at the functional level technology independent.. but needs more work.]

The control plane protocol solution required for managing the PREOF processing is outside the scope of this document.

[7.3.](#) Congestion protection and latency control

[Editor's note: TBD]

[7.4.](#) Bidirectional traffic

[Editor's NOTE: this section needs to be updated to have its scope limited to management and control.]

Some DetNet applications generate bidirectional traffic. Using MPLS definitions [[RFC5654](#)] there are associated bidirectional flows, and co-routed bidirectional flows. MPLS defines a point-to-point associated bidirectional LSP as consisting of two unidirectional point-to-point LSPs, one from A to B and the other from B to A, which are regarded as providing a single logical bidirectional transport path. This would be analogous of standard IP routing, or PWs running over two reciprocal unidirectional LSPs. MPLS defines a point-to-point co-routed bidirectional LSP as an associated bidirectional LSP which satisfies the additional constraint that its two unidirectional component LSPs follow the same path (in terms of both nodes and links) in both directions. An important property of co-routed bidirectional LSPs is that their unidirectional component LSPs share fate. In both types of bidirectional LSPs, resource allocations may differ in each direction. The concepts of associated bidirectional flows and co-routed bidirectional flows can be applied to DetNet flows as well whether IPv6 or MPLS is used.

While the IPv6 and MPLS data planes must support bidirectional DetNet flows, there are no special bidirectional features with respect to the data plane other than need for the two directions take the same paths. Fate sharing and associated vs co-routed bidirectional flows can be managed at the control level. Note, that there is no stated requirement for bidirectional DetNet flows to be supported using the same IPv6 Flow Labels or MPLS Labels in each direction. Control mechanisms will need to support such bidirectional flows for both IPv6 and MPLS, but such mechanisms are out of scope of this document. An example control plane solution for MPLS can be found in [[RFC7551](#)].

[7.5.](#) Flow aggregation control

[TBD]

[8.](#) DetNet IP Operation over DetNet MPLS Service

[Editor's note: this is a place holder section. A standalone section on operation of IP flows over DetNet MPLS data plane. Includes [RFC2119](#) Language.]

9. IEEE 802.1 TSN Interconnection over DetNet MPLS Service

[Editor's note: this is a place holder section. A standalone section on TSN "island" interconnect over DetNet". Includes [RFC2119](#) Language.]

10. DetNet MPLS Transport Layer Operation over IEEE 802.1 TSN Sub-Networks

[Editor's note: this is a place holder section. A standalone section on MPLS over IEEE 802.1 TSN. Includes [RFC2119](#) Language.]

11. DetNet MPLS Transport Layer Operation over IP DetNet PSNs

This section specifies the DetNet encapsulation over an IP transport network. The approach is modeled on the operation of MPLS and PseudoWires (PW) over an IP Packet Switched Network (PSN) [[RFC3985](#)][[RFC4385](#)][[RFC7510](#)]. It is also based on the MPLS data plane encapsulation described in [Section 6.2](#).

To carry DetNet with full functionality at the DetNet layer over an IP transport network, the following components are required (these are a subset of the requirements for MPLS encapsulation listed in [Section 6.1](#)):

1. A method of identifying the DetNet flow group to the processing element.
2. A method of carrying the DetNet sequence number.
3. A method of distinguishing DetNet OAM packets from DetNet data packets.
4. A method of carrying queuing and forwarding indication.

These requirements are satisfied by the DetNet over MPLS Encapsulation described in [Section 6.2](#).

To simplify operations and implementations, rather than inventing a new encapsulation, the IP encapsulation takes advantage of the MPLS encapsulation. By using the specification of MPLS over UDP and IP in [[RFC7510](#)], the T-Label(s) shown in Figure 15 in [Section 6.2](#) can be replaced by UDP and IP, resulting in the following encapsulation:

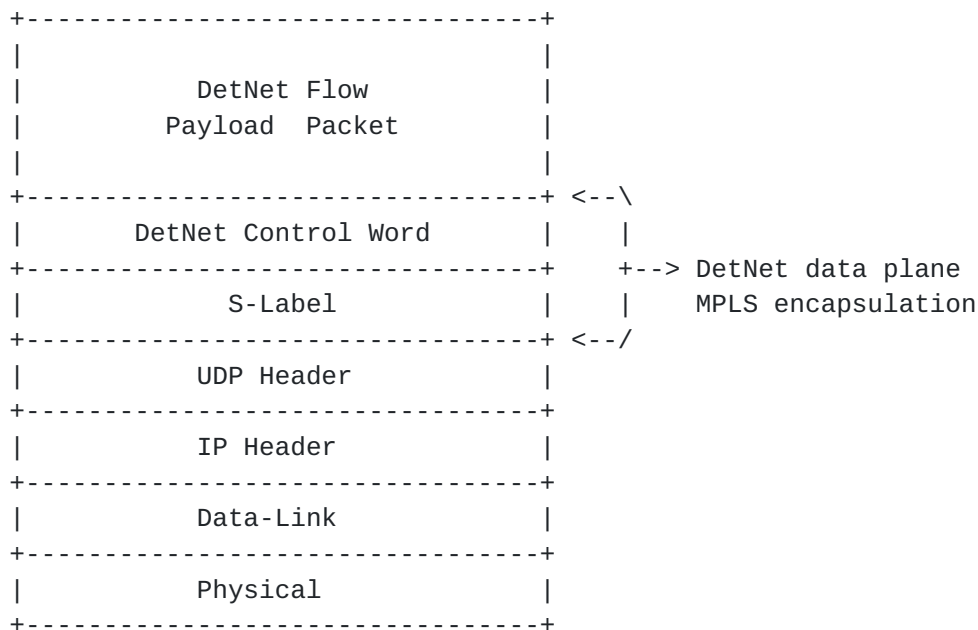


Figure 17: IP Encapsulation of DetNet

Where the UDP header is used as defined in [Section 3 of \[RFC7510\]](#).

As in [Section 6.2](#), the S-Label is used to identify a DetNet flow to the peer node that processes it, in this case the node addressed by the IP Header in Figure 17. The S-Label is allocated from the receiving node's platform label space [[RFC3031](#)].

In ingress Edge Nodes, the encapsulation in Figure 17 will be imposed on DetNet Flow Payload Packets as received from DetNet End Systems, and the encapsulation will be removed in egress Edge Nodes as they transmit the Payload Packets to the End Systems.

Note that this encapsulation works equally well with IPv4 and IPv6.

This encapsulation can also be used in conjunction with segment routing as specified in [[I-D.ietf-spring-segment-routing-mpls](#)]. In this case, the T-Label(s) in Figure 17 should be retained, and at each hop, the top T-label is popped and mapped to a corresponding UDP/IP tunnel, resulting in the following encapsulation:

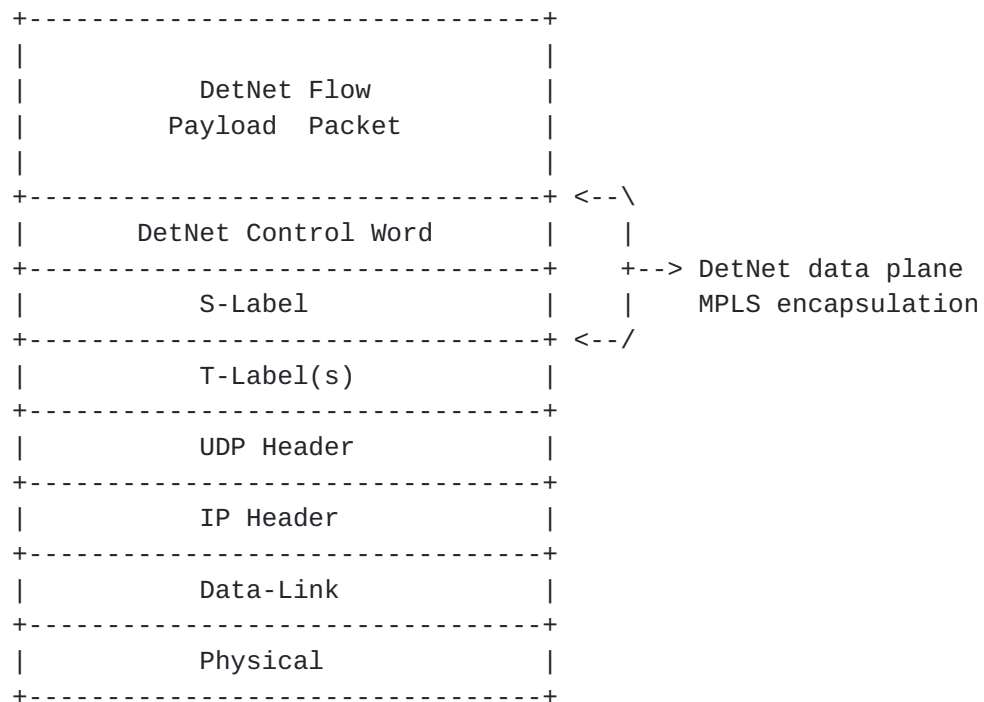


Figure 18: IP Encapsulation of DetNet with MPLS-SR

Again, the UDP header is used as defined in [Section 3 of \[RFC7510\]](#).

Note that if required in both the case of IP Encapsulation of DetNet Figure 17, and of IP Encapsulation of DetNet with MPLS-SR Figure 18, it is possible to omit the UDP header if required. Operation of MPLS directly over IP is described in [\[RFC4023\]](#). In this case DetNet Service can be provided on a per IP flow basis as described in [\[I-D.ietf-detnet-dp-sol-ip\]](#).

12. Security considerations

The security considerations of DetNet in general are discussed in [\[I-D.ietf-detnet-architecture\]](#) and [\[I-D.sdt-detnet-security\]](#). Other security considerations will be added in a future version of this draft.

13. IANA considerations

This document makes no IANA requests.

14. Contributors

[RFC7322](#) limits the number of authors listed on the front page of a draft to a maximum of 5, far fewer than the 20 individuals below who made important contributions to this draft. The editor wishes to thank and acknowledge each of the following authors for contributing text to this draft. See also [Section 15](#).

Loa Andersson
Huawei
Email: loa@pi.nu

Yuanlong Jiang
Huawei
Email: jiangyuanlong@huawei.com

Norman Finn
Huawei
3101 Rio Way
Spring Valley, CA 91977
USA
Email: norman.finn@mail01.huawei.com

Janos Farkas
Ericsson
Magyar Tudosok krt. 11.
Budapest 1117
Hungary
Email: janos.farkas@ericsson.com

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain
Email: cjbca@it.uc3m.es

Tal Mizrahi
Marvell
6 Hamada st.
Yokneam
Israel
Email: talmi@marvell.com

Lou Berger
LabN Consulting, L.L.C.
Email: lberger@labn.net

Stewart Bryant
Huawei Technologies
Email: stewart.bryant@gmail.com

Mach Chen
Huawei Technologies
Email: mach.chen@huawei.com

15. Acknowledgements

The author(s) ACK and NACK.

The following people were part of the DetNet Data Plane Solution Design Team:

Jouni Korhonen

Janos Farkas

Norman Finn

Balazs Varga

Loa Andersson

Tal Mizrahi

David Mozes

Yuanlong Jiang

Carlos J. Bernardos

The DetNet chairs serving during the DetNet Data Plane Solution Design Team:

Lou Berger

Pat Thaler

Thanks for Stewart Bryant for his extensive review of the previous versions of the document.

16. References

16.1. Normative references

[I-D.ietf-spring-segment-routing-mpls]

Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", [draft-ietf-spring-segment-routing-mpls-14](#) (work in progress), June 2018.

- [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>>.
- [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>>.
- [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>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), DOI 10.17487/RFC3168, September 2001, <<https://www.rfc-editor.org/info/rfc3168>>.
- [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>>.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", [RFC 4023](#), DOI 10.17487/RFC4023, March 2005, <<https://www.rfc-editor.org/info/rfc4023>>.
- [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>>.
- [RFC6003] Papadimitriou, D., "Ethernet Traffic Parameters", [RFC 6003](#), DOI 10.17487/RFC6003, October 2010, <<https://www.rfc-editor.org/info/rfc6003>>.

- [RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", [RFC 7510](#), DOI 10.17487/RFC7510, April 2015, <<https://www.rfc-editor.org/info/rfc7510>>.
- [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>>.

16.2. Informative references

- [I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", [draft-ietf-detnet-architecture-05](#) (work in progress), May 2018.
- [I-D.ietf-detnet-dp-alt]
Korhonen, J., Farkas, J., Mirsky, G., Thubert, P., Zhuangyan, Z., and L. Berger, "DetNet Data Plane Protocol and Solution Alternatives", [draft-ietf-detnet-dp-alt-00](#) (work in progress), October 2016.
- [I-D.ietf-detnet-dp-sol-ip]
Korhonen, J., Varga, B., "DetNet IP Data Plane Encapsulation", 2018.
- [I-D.sdt-detnet-security]
Mizrahi, T., Grossman, E., Hacker, A., Das, S., "Deterministic Networking (DetNet) Security Considerations", [draft-sdt-detnet-security](#), work in progress", 2017.
- [IEEE8021CB]
Finn, N., "Draft Standard for Local and metropolitan area networks - Seamless Redundancy", IEEE P802.1CB /D2.1 P802.1CB, December 2015, <<http://www.ieee802.org/1/files/private/cb-drafts/d2/802-1CB-d2-1.pdf>>.
- [IEEE8021Q]
IEEE 802.1, "Standard for Local and metropolitan area networks--Bridges and Bridged Networks (IEEE Std 802.1Q-2014)", 2014, <<http://standards.ieee.org/about/get/>>.
- [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>>.

- [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>>.
- [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>>.
- [RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", [RFC 5654](#), DOI 10.17487/RFC5654, September 2009, <<https://www.rfc-editor.org/info/rfc5654>>.
- [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>>.
- [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>>.
- [RFC7551] Zhang, F., Ed., Jing, R., and R. Gandhi, Ed., "RSVP-TE Extensions for Associated Bidirectional Label Switched Paths (LSPs)", [RFC 7551](#), DOI 10.17487/RFC7551, May 2015, <<https://www.rfc-editor.org/info/rfc7551>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", [RFC 7950](#), DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", [RFC 8040](#), DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8169] Mirsky, G., Ruffini, S., Gray, E., Drake, J., Bryant, S., and A. Vainshtein, "Residence Time Measurement in MPLS Networks", [RFC 8169](#), DOI 10.17487/RFC8169, May 2017, <<https://www.rfc-editor.org/info/rfc8169>>.

Appendix A. Example of DetNet data plane operation

[Editor's note: Add a simplified example of DetNet data plane and how labels etc work in the case of MPLS-based PSN and utilizing PREOF.]

The figure is subject to change depending on the further DT decisions on the label handling..]

Authors' Addresses

Jouni Korhonen (editor)

Email: jouni.nospam@gmail.com

Balazs Varga (editor)

Ericsson

Magyar Tudosok krt. 11.

Budapest 1117

Hungary

Email: balazs.a.varga@ericsson.com

