

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

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

DetNet Data Plane: MPLS
draft-ietf-detnet-mpls-03

Abstract

This document specifies the Deterministic Networking data plane when operating 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 April 29, 2020.

Copyright Notice

Copyright (c) 2019 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	3
2.2. Abbreviations	4
2.3. Requirements Language	5
3. DetNet MPLS Data Plane Overview	5
3.1. Layers of DetNet Data Plane	5
3.2. DetNet MPLS Data Plane Scenarios	6
4. MPLS-Based DetNet Data Plane Solution	8
4.1. DetNet Over MPLS Encapsulation Components	8
4.2. MPLS Data Plane Encapsulation	9
4.2.1. DetNet Control Word and the DetNet Sequence Number	10
4.2.2. S-Labels	11
4.2.3. F-Labels	14
4.3. OAM Indication	16
4.4. Flow Aggregation	17
4.4.1. Aggregation Via LSP Hierarchy	17
4.4.2. Aggregating DetNet Flows as a new DetNet flow	17
4.5. Service Sub-Layer Considerations	19
4.5.1. Edge Node Processing	19
4.5.2. Relay Node Processing	19
4.6. Forwarding Sub-Layer Considerations	20
4.6.1. Class of Service	20
4.6.2. Quality of Service	20
5. Management and Control Information Summary	21
5.1. Service Sub-Layer Information Summary	21
5.1.1. Service Aggregation Information Summary	22
5.2. Forwarding Sub-Layer Information Summary	23
6. Security Considerations	24
7. IANA Considerations	25
8. Acknowledgements	25
9. References	25
9.1. Normative References	25
9.2. Informative References	27
Authors' Addresses	29

1. Introduction

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

The DetNet Architecture models the DetNet related data plane functions 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. The forwarding sub-layer is used to provide forwarding assurance (low loss, assured latency, and limited reordering).

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 sub-layer 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 DetNet required level of service. However, the specific details of how DetNet MPLS is carried over different network technologies is out of scope of 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 of this document. An example of IP using DetNet MPLS sub-networks can be found in [I-D.ietf-detnet-ip]. 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 [I-D.ietf-detnet-data-plane-framework].

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 the DetNet Data Plane Framework [I-D.ietf-detnet-data-plane-framework]. The reader is assumed to be familiar with these documents and any terminology defined therein.

The following terminology is introduced in this document:

F-Label	A Detnet "forwarding" label that identifies the LSP used to forward 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 sub-layer functions. An S-Label is also used to identify a DetNet flow at DetNet service sub-layer.
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) 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.
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.

2.3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. DetNet MPLS Data Plane Overview

3.1. Layers of DetNet Data Plane

MPLS provides a wide range of capabilities that can be utilised by DetNet. A straight forward approach utilizing MPLS for a DetNet service sub-layer is based on the existing pseudowire (PW) encapsulations and by utilizing existing MPLS Traffic Engineering encapsulations and mechanisms. Background on PWs can be found in [RFC3985] and [RFC3031]. Background on MPLS Traffic Engineering can be found in [RFC3272] and [RFC3209].




Figure 1: DetNet Adaptation to MPLS Data Plane

The DetNet MPLS data plane representation is illustrated in Figure 1. The service sub-layer includes a DetNet control word (d-CW) and a identifying service label (S-Label). The DetNet control word (d-CW) conforms to the Generic PW MPLS Control Word (PWMCW) defined in [RFC4385]. An aggregation label (A-Label) is a special case of S-Label used for aggregation.

A node operating on a DetNet flow in the Detnet service sub-layer, uses the local context associated with that S-Label, provided by a received F-Label, to determine what local DetNet operation(s) are applied to that packet. An S-Label may be taken from the platform label space [RFC3031], making it unique, enabling DetNet flow identification regardless of which input interface or LSP the packet arrives on.

The DetNet forwarding sub-layer is supported by zero or more forwarding labels (F-Labels). MPLS Traffic Engineering encapsulations and mechanisms can be utilized to provide a forwarding sub-layer that is responsible for providing resource allocation and explicit routes.

3.2. DetNet MPLS Data Plane Scenarios




Figure 2: A DetNet MPLS Network

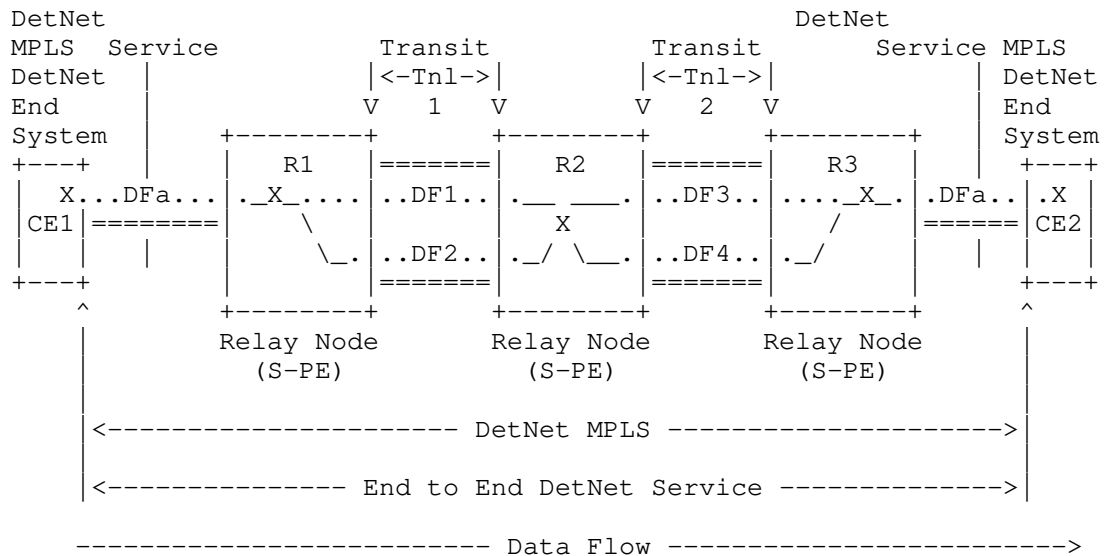
Figure 2 illustrates a hypothetical DetNet MPLS-only network composed of DetNet aware MPLS enabled end systems, operating over a DetNet aware MPLS network. In this figure, the relay nodes are PE devices that define the MPLS LSP boundaries and the transit nodes are LSRs.

DetNet end system 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 may be 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 plane/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 a more detail. In this figure, the customer end systems, 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 [I-D.ietf-detnet-mpls-over-tsn], or simply over interconnect network links.



X = Optional service protection (none, PRF, PREOF, PEF/POF)
 DFx = DetNet member flow x over a TE LSP

Figure 3: MPLS-Based Native DetNet

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 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 DetNet sequence number is carried in the DetNet Control word which 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 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 may be of any type (MPLS-LDP, MPLS-TE, MPLS-TP [RFC5921], or MPLS-SR [I-D.ietf-spring-segment-routing-mpls]). The LSP (F-Label) label and/or the S-Label may be used to indicate the 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 [I-D.ietf-detnet-data-plane-framework]. 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 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.
- o DetNet service Label (S-Label) that identifies a DetNet flow at the receiving DetNet service sub-layer processing node.
- o Zero or more Detnet MPLS Forwarding label(s) (F-Label) used to direct the packet along the label switched path (LSP) to the next service sub-layer processing node along the path. When Penultimate Hop Popping is in use there may be no label F-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 4: Encapsulation of a DetNet App-Flow in an MPLS PSN

4.2.1. DetNet Control Word and the DetNet Sequence Number

A DetNet control word (d-CW) conforms to the Generic PW MPLS Control Word (PVMCW) defined in [RFC4385]. The d-CW formatted as shown in Figure 5 MUST be present in all DetNet packets containing app-flow data.




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.

A separate sequence number space MUST be maintained by the node that adds the d-CW for each DetNet app-flow. 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 app-flow basis via configuration, i.e., via the controller plane described in [I-D.ietf-detnet-data-plane-framework].

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, were 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 Section 4.2.2.1 and Section 4.2.2.2.

4.2.2. S-Labels

App-flow identification at a DetNet service sub-layer is realized by an S-Label. MPLS-aware DetNet end systems and edge nodes, which are by definition MPLS ingress and egress nodes, MUST add and remove an app-flow specific d-CW and S-Label. Relay nodes MAY swap S-Label values when processing an app-flow.

The S-Label value MUST be provisioned per app-flow via configuration, e.g., via the controller plane described in [I-D.ietf-detnet-data-plane-framework]. Note that S-Labels provide app-flow 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 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 (e.g., a DetNet MPLS End System or a DetNet Relay or Edge Node and interpreted in the context of their received F-Label.

The 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 service sub-layer level OAM, 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/Entropy Label (ELI/EL) [RFC6790] MAY be carried below the S-Label in the label stack in networks where DetNet flows would otherwise received 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 outlines in [I-D.ietf-detnet-data-plane-framework] 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 flow, an implementation MUST identify the app-flow 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 app-flow. In the case where platform labels are not used, zero or more F-Labels and optionally, the incoming interface, proceeding the S-Label MUST be used together with the S-Label to uniquely identify the app-flows associated with a received packet. The incoming interface MAY also be used to together with any present F-Label(s) and the S-Label to uniquely identify an incoming app-flows, for example, to in the case where PHP is used. Note that choice to use platform label space for S-Label or S-Label plus one or more F-Labels to identify app flows is a local implementation choice, with one caveat. When one or more F-labels, or incoming interface, is needed together with an S-Label to uniquely identify, the controller plane MUST ensure that incoming DetNet MPLS packets arrive with the needed information (F-label(s) and/or incoming interface); the details of such are outside the scope of this document.

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

4.2.2.1. 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 app-flow MUST be provisioned via configuration, e.g., via the controller plane described in [I-D.ietf-detnet-data-plane-framework].

After an app-flow is identified for a received DetNet MPLS packet, as described above, an implementation MUST check if PEF is configured for that app-flow. When configured, the implementation MUST track the sequence number contained in received d-CWs and MUST ensure that duplicate (replicated) instances of a particular sequence number are 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.

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

4.2.2.2. 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 app-flow MUST be provisioned via configuration, e.g., via the controller plane described by [I-D.ietf-detnet-data-plane-framework]. Implementations MAY required that PEF and POF be used in combination. There is no requirement related to the order of execution of the Packet Elimination and Ordering Functions in an implementation.

After an app-flow is identified for a received DetNet MPLS packet, as described above, an implementation MUST check if POF is configured for that app-flow. When configured, the implementation MUST track the sequence number contained in received d-CWs and MUST ensure that packets are 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, is incremented by one (1) for each new app-flow packet sent, and the field value can wrap. The specific mechanisms used for an implementation to identify the order of received packets is an implementation choice.

Note that an implementation MAY wish to constrain the maximum number of out of order packets that can be processed, on platform-wide or per flow basis. Some implementations MAY support the provisioning of this number 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.

4.2.3. F-Labels

F-Labels are supported 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 and Packet Replication Function Processing

DetNet MPLS end systems, edge nodes and relay nodes may operate at the DetNet service sub-layer with understand of app-flows 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-Labels used for the app-flow 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-Labels to be pushed when sending a particular app-flow MUST be provisioned per app-flow via configuration, e.g., via the controller plane discussed in [I-D.ietf-detnet-data-plane-framework]. F-Labels can also provide context for an S-Label. To allow for the omission of F-Labels, an implementation SHOULD also allow an outgoing interface to be used.

The Packet Replication Function (PRF) function MAY be supported by an implementation for outgoing DetNet flows. When replication is supported, the same app-flow data will be sent over multiple outgoing forwarding sub-layer LSPs. To support PRF an implementation MUST support the setting of different sets of F-Labels. To allow for the omission of F-Labels, an implementation SHOULD also allow multiple outgoing interfaces to be provisioned. PRF MUST NOT be used with app-flows configured with a d-CW sequence number field length of 0 bits.

When a single set of F-Labels is provisioned for a particular outgoing app-flow, 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 F-Labels or interfaces are provisioned for a particular outgoing app-flow, a copy of the outgoing packet, including the pushed 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 app-flow 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 app-flow. In the case where platform labels are not used, F-Label(s) and, optionally, the incoming interface MUST also be provisioned for incoming app-flows. The provisioned information MUST then be used to identify incoming app-flows based on the combination of S-Label and F-Label(s) or incoming interface.

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 app-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, [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 LSP that is represented by each F-Label. Each LSP MAY be provisioned with a service parameters 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 includes the allocation of system resources (such as queues and buffers) and provisioning or related parameters (such as shaping, and policing). Support can also be provided via an underlying network technology such IEEE802.1 TSN [I-D.ietf-detnet-mpls-over-tsn]. The specific mechanisms used 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:

- o MUST defend the DetNet QoS by discarding or remarking (to an allocated DetNet flow or non-competing non-DetNet flow) packets received that are not associated with a completed resource allocation.
- o 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.
- o MUST ensure a QoS flow does not exceed its allocated resources or provisioned service level,
- o MUST ensure a CoS flow or service class does not impact the service delivered to other flows. This requirement is similar to requirement for MPLS LSRs to 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-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.

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 is out of scope of this documents. 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 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 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:

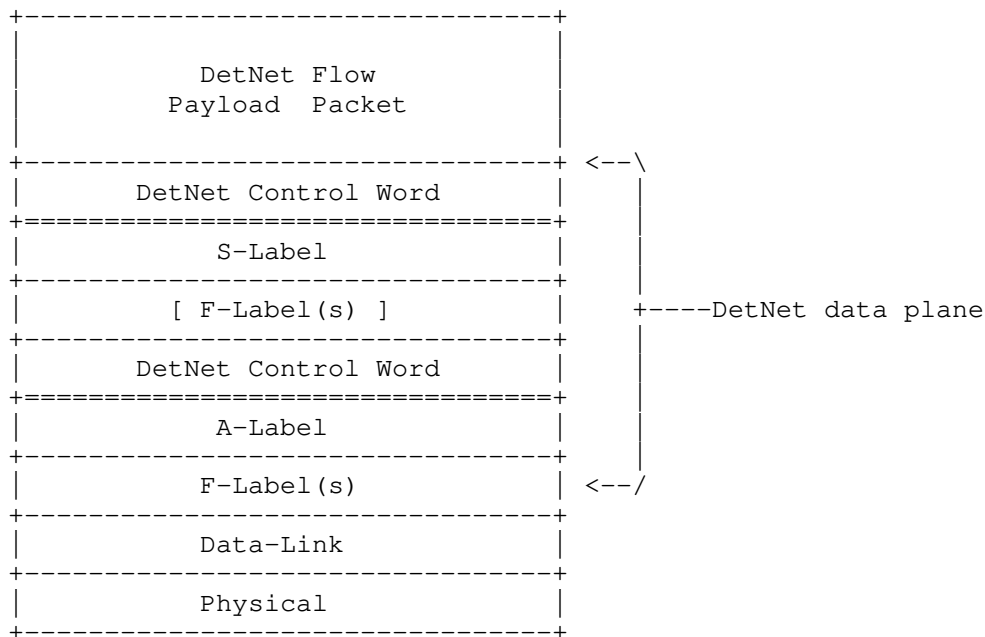


Figure 6: DetNet Aggregation as a new DetNet Flow

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 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 are properly provisioned.

4.5. Service Sub-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.

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

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 duplicate packet 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).

4.5.2. Relay Node Processing

A DetNet Relay node operates in the DetNet forwarding sub-layer . For DetNet using MPLS this processing is performed on the F-Label. 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 4.4. Again, this is entirely up to how the forwarding has been programmed.

4.6. Forwarding Sub-Layer Considerations

4.6.1. Class of Service

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].

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 [I-D.ietf-detnet-architecture], there are different mechanisms that maybe used separately or in combination to deliver a zero congestion loss service. This includes Quality of Service (QoS) mechanisms at the MPLS layer, that may be combined with the mechanisms defined by the underlying network layer such as 802.1TSN.

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.

A baseline set of QoS capabilities for DetNet flows carried in PWs and MPLS can be 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 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 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 NETCONF/YANG, BGP, 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 of this document. There are particular DetNet considerations and requirements that are discussed in [I-D.ietf-detnet-data-plane-framework].

5.1. Service Sub-Layer Information Summary

The following summarizes the information that is needed on service sub-layer aware nodes that send DetNet MPLS traffic, on a per service basis:

- o App-Flow identification information, e.g., an incoming service on a relay node or IP information as defined in [I-D.ietf-detnet-ip-over-mpls].
- o The sequence number length to be used for the service. Valid values included 0, 16 and 28 bits. 0 bits cannot be used when PRF is configured for the service.
- o The S-Label for the service.
- o If PRF is to be provided for the service.
- o The forwarding sub-layer information associated with the output of the service sub-layer. Note that when the PRF function 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 service sub-layer aware nodes that receives DetNet MPLS traffic, on a per service basis:

- o The forwarding sub-layer information associated with the input of the service sub-layer. Note that when the PEF function 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.
- o The S-Label for the received service.
- o If PEF or POF is to be provided for the service.
- o 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.

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 information that is needed on a node that sends aggregated flows using A-Labels:

- o The S-Labels or F-Labels that are to be carried over each aggregated service.
- o The A-Label associated with each aggregated service.
- o 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 appropriated. The related addition configuration that may be provided discussed elsewhere in this section.

5.2. Forwarding Sub-Layer Information Summary

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

- o The outgoing F-Label stack to be pushed. The stack may include H-LSP labels.
- o 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.
- o Outgoing interface and, for unicast traffic, the next hop information.
- o Sub-network specific parameters on a technology specific basis. For example, see [I-D.ietf-detnet-mpls-over-tsn].

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

- o The incoming interface. For some implementations and deployment scenarios, this information may not be needed.
- o The incoming F-Label stack to be popped. The stack may include H-LSP labels.
- o 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 provided the traffic treatment needed to meet each flow's service requirements. This applies for aggregated and individual flows.

6. Security Considerations

Security considerations for DetNet are described in detail in [I-D.ietf-detnet-security]. General security considerations are described in [I-D.ietf-detnet-architecture]. This section considers exclusively security considerations which are specific to the DetNet MPLS data plane.

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

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

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

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

To provide uninterrupted availability of the DetNet service, provisions can be made against DOS attacks and delay attacks. To protect against DOS attacks, excess traffic due to malicious or malfunctioning devices can be prevented or mitigated, for example through the use of existing mechanism such as policing and shaping applied at the input of a DetNet domain. To prevent DetNet packets from being delayed by an entity external to a DetNet domain, DetNet technology definition can allow for the mitigation of Man-In-The-Middle attacks, for example through use of authentication and authorization of devices within the DetNet domain.

7. IANA Considerations

This document makes no IANA requests.

8. Acknowledgements

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

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [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>>.

9.2. Informative References

- [I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas,
"Deterministic Networking Architecture", draft-ietf-
detnet-architecture-13 (work in progress), May 2019.
- [I-D.ietf-detnet-data-plane-framework]
Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A.,
Bryant, S., and J. Korhonen, "DetNet Data Plane
Framework", draft-ietf-detnet-data-plane-framework-02
(work in progress), September 2019.
- [I-D.ietf-detnet-ip]
Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A.,
Bryant, S., and J. Korhonen, "DetNet Data Plane: IP",
draft-ietf-detnet-ip-01 (work in progress), July 2019.
- [I-D.ietf-detnet-ip-over-mpls]
Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A.,
Bryant, S., and J. Korhonen, "DetNet Data Plane: IP over
MPLS", draft-ietf-detnet-ip-over-mpls-01 (work in
progress), July 2019.
- [I-D.ietf-detnet-mpls-over-tsn]
Varga, B., Farkas, J., Malis, A., Bryant, S., and J.
Korhonen, "DetNet Data Plane: MPLS over IEEE 802.1 Time
Sensitive Networking (TSN)", draft-ietf-detnet-mpls-over-
tsn-00 (work in progress), May 2019.
- [I-D.ietf-detnet-security]
Mizrahi, T., Grossman, E., Hacker, A., Das, S., Dowdell,
J., Austad, H., Stanton, K., and N. Finn, "Deterministic
Networking (DetNet) Security Considerations", draft-ietf-
detnet-security-05 (work in progress), August 2019.
- [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-22
(work in progress), May 2019.
- [IEEE802.1AE-2018]
IEEE Standards Association, "IEEE Std 802.1AE-2018 MAC
Security (MACsec)", 2018,
<<https://ieeexplore.ieee.org/document/8585421>>.

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

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

Authors' Addresses

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

Email: balazs.a.varga@ericsson.com

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

Email: janos.farkas@ericsson.com

Lou Berger
LabN Consulting, L.L.C.

Email: lberger@labn.net

Don Fedyk
LabN Consulting, L.L.C.

Email: dfedyk@labn.net

Andrew G. Malis
Independent

Email: agmalis@gmail.com

Stewart Bryant
Futurewei Technologies

Email: stewart.bryant@gmail.com

Jouni Korhonen

Email: jouni.nospam@gmail.com