

---

Stream: Internet Engineering Task Force (IETF)  
RFC: [8938](#)  
Category: Informational  
Published: November 2020  
ISSN: 2070-1721  
Authors:  
B. Varga, Ed. J. Farkas L. Berger A. Malis S. Bryant  
*Ericsson Ericsson LabN Consulting, L.L.C. Malis Consulting Futurewei Technologies*

# RFC 8938

## Deterministic Networking (DetNet) Data Plane Framework

---

### Abstract

This document provides an overall framework for the Deterministic Networking (DetNet) data plane. It covers concepts and considerations that are generally common to any DetNet data plane specification. It describes related Controller Plane considerations as well.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8938>.

### Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction
2. Terminology
  - 2.1. Terms Used in This Document
  - 2.2. Abbreviations
3. Overview of the DetNet Data Plane
  - 3.1. Data Plane Characteristics
    - 3.1.1. Data Plane Technology
    - 3.1.2. Encapsulation
  - 3.2. DetNet-Specific Metadata
  - 3.3. DetNet IP Data Plane
  - 3.4. DetNet MPLS Data Plane
  - 3.5. Further DetNet Data Plane Considerations
    - 3.5.1. Functions Provided on a Per-Flow Basis
    - 3.5.2. Service Protection
    - 3.5.3. Aggregation Considerations
    - 3.5.4. End-System-Specific Considerations
    - 3.5.5. Sub-network Considerations
4. Controller Plane (Management and Control) Considerations
  - 4.1. DetNet Controller Plane Requirements
  - 4.2. Generic Controller Plane Considerations
    - 4.2.1. Flow Aggregation Control
    - 4.2.2. Explicit Routes
    - 4.2.3. Contention Loss and Jitter Reduction
    - 4.2.4. Bidirectional Traffic
  - 4.3. Packet Replication, Elimination, and Ordering Functions (PREOF)
5. Security Considerations
6. IANA Considerations

## 7. References

### 7.1. Normative References

### 7.2. Informative References

## Acknowledgements

## Contributors

## Authors' Addresses

# 1. Introduction

DetNet (Deterministic Networking) provides the ability to carry specified unicast or multicast data flows for real-time applications with extremely low packet loss rates and assured maximum end-to-end delivery latency. A description of the general background and concepts of DetNet can be found in [\[RFC8655\]](#).

This document describes the concepts needed by any DetNet data plane specification (i.e., the DetNet-specific use of packet header fields) and provides considerations for any corresponding implementation. It covers the building blocks that provide the DetNet service, the DetNet service sub-layer, and the DetNet forwarding sub-layer functions as described in the DetNet architecture [\[RFC8655\]](#).

The DetNet architecture models the DetNet-related data plane functions as being decomposed into two sub-layers: a service sub-layer and a forwarding sub-layer. The service sub-layer is used to provide DetNet service protection and reordering. The forwarding sub-layer leverages traffic engineering mechanisms and provides congestion protection (low loss, assured latency, and limited out-of-order delivery). A particular forwarding sub-layer may have capabilities that are not available on other forwarding sub-layers. DetNet makes use of the existing forwarding sub-layers with their respective capabilities and does not require 1:1 equivalence between different forwarding sub-layer capabilities.

As part of the service sub-layer functions, this document describes typical DetNet node data plane operation. It describes the functionality and operation of the Packet Replication Function (PRF), the Packet Elimination Function (PEF), and the Packet Ordering Function (POF) within the service sub-layer. Furthermore, it describes the forwarding sub-layer.

As defined in [\[RFC8655\]](#), DetNet flows may be carried over network technologies that can provide service characteristics required by DetNet. For example, DetNet MPLS flows can be carried over IEEE 802.1 Time-Sensitive Networking (TSN) sub-networks [\[IEEE802.1TSNTG\]](#). However, IEEE 802.1 TSN support is not required in DetNet. TSN frame preemption is an example of a forwarding layer capability that is typically not replicated in other forwarding technologies.

Most of DetNet's benefits can be gained by running over a data-link layer that has not been specifically enhanced to support all TSN capabilities, but for such networks and traffic mixes, delay and jitter performance may vary due to the forwarding sub-layer's intrinsic properties.

Different application flows, such as Ethernet or IP, can be mapped on top of DetNet. DetNet can optionally reuse header information provided by, or shared with, applications. An example of shared header fields can be found in [\[RFC8939\]](#).

This document also covers basic concepts related to the Controller Plane and Operations, Administration, and Maintenance (OAM). Data plane OAM specifics are out of scope for this document.

## 2. Terminology

### 2.1. Terms Used in This Document

This document uses the terminology established in the DetNet architecture [\[RFC8655\]](#), and it is assumed that the reader is familiar with that document and its terminology.

### 2.2. Abbreviations

The following abbreviations are used in this document:

BGP	Border Gateway Protocol
CoS	Class of Service
d-CW	DetNet Control Word
DetNet	Deterministic Networking
DN	DetNet
GMPLS	Generalized Multiprotocol Label Switching
GRE	Generic Routing Encapsulation
IPsec	IP Security
L2	Layer 2
LSP	Label Switched Path
MPLS	Multiprotocol Label Switching
OAM	Operations, Administration, and Maintenance
PCEP	Path Computation Element Communication Protocol
PEF	Packet Elimination Function

POF	Packet Ordering Function
PREOF	Packet Replication, Elimination, and Ordering Functions
PRF	Packet Replication Function
PSN	Packet Switched Network
QoS	Quality of Service
S-Label	DetNet "service" label
TDM	Time-Division Multiplexing
TSN	Time-Sensitive Networking
YANG	Yet Another Next Generation

### 3. Overview of the DetNet Data Plane

This document describes how application flows, or App-flows [RFC8655], are carried over DetNet networks. The DetNet architecture [RFC8655] models the DetNet-related data plane functions as decomposed into two sub-layers: a service sub-layer and a forwarding sub-layer.

Figure 1, reproduced from [RFC8655], shows a logical DetNet service with the two sub-layers.

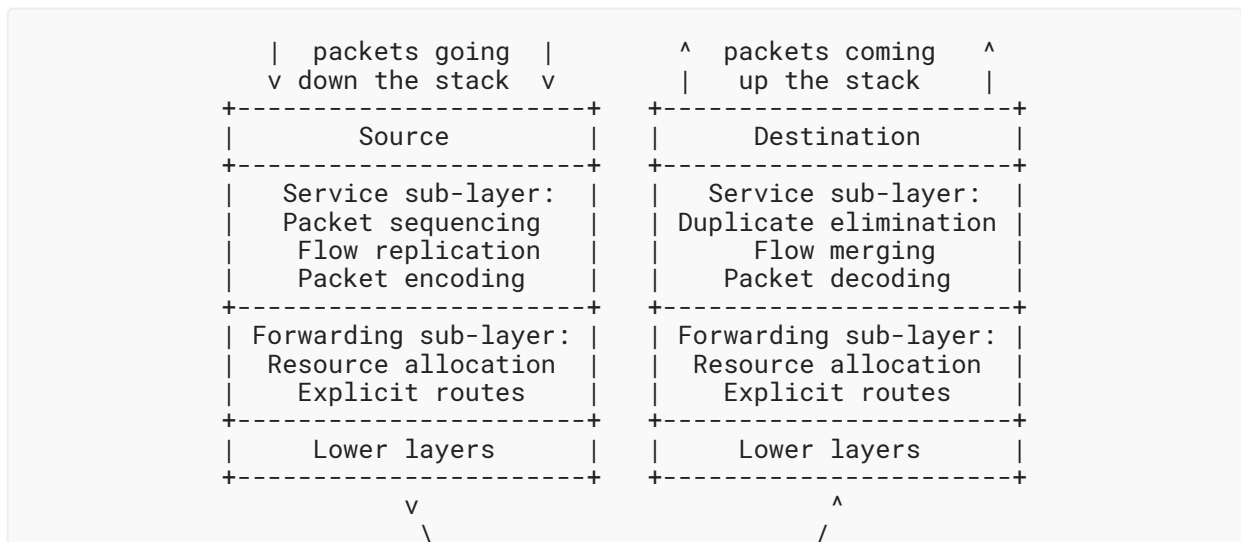


Figure 1: DetNet Data Plane Protocol Stack

The DetNet forwarding sub-layer may be directly provided by the DetNet service sub-layer -- for example, by IP tunnels or MPLS. Alternatively, an overlay approach may be used in which the packet is natively carried between key nodes within the DetNet network (say, between PREOF nodes), and a sub-layer is used to provide the information needed to reach the next hop in the overlay.

The forwarding sub-layer provides the QoS-related functions needed by the DetNet flow. It may do this directly through the use of queuing techniques and traffic engineering methods, or it may do this through the assistance of its underlying connectivity. For example, it may call upon Ethernet TSN capabilities defined in IEEE 802.1 TSN [IEEE802.1TSNTG]. The forwarding sub-layer uses buffer resources for packet queuing, as well as reservation and allocation of bandwidth capacity resources.

The service sub-layer provides additional support beyond the connectivity function of the forwarding sub-layer. See [Section 4.3](#) regarding PREOF. The POF uses sequence numbers added to packets, enabling a range of packet order protection from simple ordering and dropping out-of-order packets to more complex reordering of a fixed number of out-of-order, minimally delayed packets. Reordering requires buffer resources and has an impact on the delay and jitter of packets in the DetNet flow.

The method of instantiating each of the layers is specific to the particular DetNet data plane method, and more than one approach may be applicable to a given network type.

### 3.1. Data Plane Characteristics

The data plane has two major characteristics: the technology and the encapsulation, as discussed below.

#### 3.1.1. Data Plane Technology

The DetNet data plane is provided by the DetNet service and forwarding sub-layers. The DetNet service sub-layer generally provides its functions for the DetNet application flows by using or applying existing standardized headers and/or encapsulations. The DetNet forwarding sub-layer may provide capabilities leveraging that same header or encapsulation technology (e.g., DN IP or DN MPLS), or it may be achieved via other technologies, as shown in [Figure 2](#) below. DetNet is currently defined for operation over packet-switched (IP) networks or label-switched (MPLS) networks.

#### 3.1.2. Encapsulation

DetNet encodes specific flow attributes (flow identity and sequence number) in packets. For example, in DetNet IP, zero encapsulation is used, and no sequence number is available; in DetNet MPLS, DetNet-specific information may be added explicitly to the packets in the form of an S-Label and a d-CW [[DetNet-MPLS](#)].

The encapsulation of a DetNet flow allows it to be sent over a data plane technology other than its native type. DetNet uses header information to perform traffic classification, i.e., identify DetNet flows, and provide DetNet service and forwarding functions. As mentioned above, DetNet may add headers, as is the case for DN MPLS, or may use headers that are already present, as is the case for DN IP. [Figure 2](#) illustrates some relationships between the components.

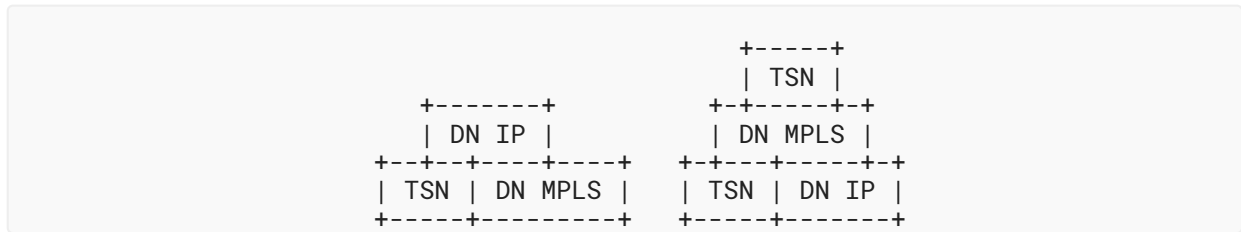


Figure 2: DetNet Service Examples

The use of encapsulation is also required if additional information (metadata) is needed by the DetNet data plane and either (1) there is no ability to include it in the client data packet or (2) the specification of the client data plane does not permit the modification of the packet to include additional data. An example of such metadata is the inclusion of a sequence number required by PREOF.

Encapsulation may also be used to carry or aggregate flows for equipment with limited DetNet capability.

### 3.2. DetNet-Specific Metadata

The DetNet data plane can provide or carry the following metadata:

1. Flow-ID
2. Sequence number

The DetNet data plane framework supports a Flow-ID (for identification of the flow or aggregate flow) and/or a sequence number (for PREOF) for each DetNet flow. The Flow-ID is used by both the service and forwarding sub-layers, but the sequence number is only used by the service layer. Metadata can also be used for OAM indications and instrumentation of DetNet data plane operation.

Metadata inclusion can be implicit or explicit. Explicit inclusions involve a dedicated header field that is used to include metadata in a DetNet packet. In the implicit method, part of an already-existing header field is used to encode the metadata.

Explicit inclusion of metadata is possible through the use of IP options or IP extension headers. New IP options are almost impossible to get standardized or to deploy in an operational network and will not be discussed further in this text. IPv6 extension headers are finding popularity in current IPv6 development work, particularly in connection with Segment Routing of IPv6 (SRv6) and IP OAM. The design of a new IPv6 extension header or the modification of an existing one is a technique available in the toolbox of the DetNet IP data plane designer.

Explicit inclusion of metadata in an IP packet is also possible through the inclusion of an MPLS label stack and the MPLS d-CW, using one of the methods for carrying MPLS over IP [[DetNet-MPLS-over-UDP-IP](#)]. This is described in more detail in [Section 3.5.5](#).

Implicit metadata in IP can be included through the use of the network programming paradigm [[SRv6-Network-Prog](#)], in which the suffix of an IPv6 address is used to encode additional information for use by the network of the receiving host.

An MPLS example of explicit metadata is the sequence number used by PREOF, or even the case where all the essential information is included in the DetNet-over-MPLS label stack (the d-CW and the DetNet S-Label).

### 3.3. DetNet IP Data Plane

An IP data plane may operate natively or through the use of an encapsulation. Many types of IP encapsulation can satisfy DetNet requirements, and it is anticipated that more than one encapsulation may be deployed -- for example, GRE, IPsec.

One method of operating an IP DetNet data plane without encapsulation is to use 6-tuple-based flow identification, where "6-tuple" refers to information carried in IP-layer and higher-layer protocol headers. General background on the use of IP headers and 6-tuples to identify flows and support QoS can be found in [[RFC3670](#)]. The extra field in the 6-tuple is the DSCP field in the packet. [[RFC7657](#)] provides useful background on differentiated services (Diffserv) and tuple-based flow identification. DetNet flow aggregation may be enabled via the use of wildcards, masks, prefixes, and ranges. The operation of this method is described in detail in [[RFC8939](#)].

The DetNet forwarding plane may use explicit route capabilities and traffic engineering capabilities to provide a forwarding sub-layer that is responsible for providing resource allocation and explicit routes. It is possible to include such information in a native IP packet either explicitly or implicitly.

### 3.4. DetNet MPLS Data Plane

MPLS provides a forwarding sub-layer for traffic over implicit and explicit paths to the point in the network where the next DetNet service sub-layer action needs to take place. It does this through the use of a stack of one or more labels with various forwarding semantics.

MPLS also provides the ability to identify a service instance that is used to process the packet through the use of a label that maps the packet to a service instance.

In cases where metadata is needed to process an MPLS-encapsulated packet at the service sub-layer, the d-CW [[DetNet-MPLS](#)] can be used. Although such d-CWs are frequently 32 bits long, there is no architectural constraint on the size of this structure -- only the requirement that it be fully understood by all parties operating on it in the DetNet service sub-layer. The operation of this method is described in detail in [[DetNet-MPLS](#)].

### 3.5. Further DetNet Data Plane Considerations

This section provides informative considerations related to providing DetNet service to flows that are identified based on their header information.



### 3.5.1. Functions Provided on a Per-Flow Basis

At a high level, the following functions are provided on a per-flow basis.

#### 3.5.1.1. Reservation and Allocation of Resources

Resources might be reserved in order to make them available for allocation to specific DetNet flows. This can eliminate packet contention and packet loss for DetNet traffic. This also can reduce jitter for DetNet traffic. Resources allocated to a DetNet flow protect it from other traffic flows. On the other hand, it is assumed that DetNet flows behave in accordance with the reserved traffic profile. It must be possible to detect misbehaving DetNet flows and to ensure that they do not compromise QoS of other flows. Queuing, policing, and shaping policies can be used to ensure that the allocation of resources reserved for DetNet is met.

#### 3.5.1.2. Explicit Routes

A flow can be routed over a specific, precomputed path. This allows control of network delay by steering the packet with the ability to influence the physical path. Explicit routes complement reservation by ensuring that a consistent path can be associated with its resources for the duration of that path. Coupled with the traffic mechanism, this limits misordering and bounds latency. Explicit route computation can encompass a wide set of constraints and can optimize the path for a certain characteristic, e.g., highest bandwidth or lowest jitter. In these cases, the "best" path for any set of characteristics may not be a shortest path. The selection of the path can take into account multiple network metrics. Some of these metrics are measured and distributed by the routing system as traffic engineering metrics.

#### 3.5.1.3. Service Protection

Service protection involves the use of multiple packet streams using multiple paths -- for example, 1+1 or 1:1 linear protection. For DetNet, this primarily relates to packet replication and elimination capabilities. MPLS offers a number of protection schemes. MPLS hitless protection can be used to switch traffic to an already-established path in order to restore delivery rapidly after a failure. Path changes, even in the case of failure recovery, can lead to the out-of-order delivery of data requiring POFs either within the DetNet service or at a high layer in the application traffic. Establishment of new paths after a failure is out of scope for DetNet services.

#### 3.5.1.4. Network Coding

Network Coding [[nwcrg](#)], not to be confused with network programming, comprises several techniques where multiple data flows are encoded. These resulting flows can then be sent on different paths. The encoding operation can combine flows and error recovery information. When the encoded flows are decoded and recombined, the original flows can be recovered. Note that Network Coding uses an alternative to packet-by-packet PREOF. Therefore, for certain network topologies and traffic loads, Network Coding can be used to improve a network's throughput, efficiency, latency, and scalability, as well as resilience to partition, attacks, and eavesdropping, as compared to traditional methods. DetNet could use Network Coding as an alternative to other means of protection. Network Coding is often applied in wireless networks and is being explored for other network types.

### 3.5.1.5. Load-Sharing

The use of packet-by-packet load-sharing of the same DetNet flow over multiple paths is not recommended, except for the cases listed above where PREOF are utilized to improve protection of traffic and maintain order. Packet-by-packet load-sharing, e.g., via Equal-Cost Multipath (ECMP) or Unequal-Cost Multipath (UCMP), impacts ordering and, possibly, jitter.

### 3.5.1.6. Troubleshooting

DetNet leverages many different forwarding sub-layers, each of which supports various tools to troubleshoot connectivity -- for example, identification of misbehaving flows. The DetNet service layer can leverage existing mechanisms to troubleshoot or monitor flows, such as those in use by IP and MPLS networks. At the Application layer, a client of a DetNet service can use existing techniques to detect and monitor delay and loss.

### 3.5.1.7. Flow Recognition for Analytics

Network analytics can be inherited from the technologies of the service and forwarding sub-layers. At the DetNet service edge, packet and bit counters (e.g., sent, received, dropped, out of sequence) can be maintained.

### 3.5.1.8. Correlation of Events with Flows

The provider of a DetNet service may provide other capabilities to monitor flows, such as more detailed loss statistics and timestamping of events. Details regarding these capabilities are out of scope for this document.

## 3.5.2. Service Protection

Service protection allows DetNet services to increase reliability and maintain a desired level of service assurance in the case of network congestion or network failure. DetNet relies on the underlying technology capabilities for various protection schemes. Protection schemes enable partial or complete coverage of the network paths and active protection with combinations of the PRF, PEF, and POF.

### 3.5.2.1. Linear Service Protection

An example DetNet MPLS network fragment and its packet flow are illustrated in [Figure 3](#).

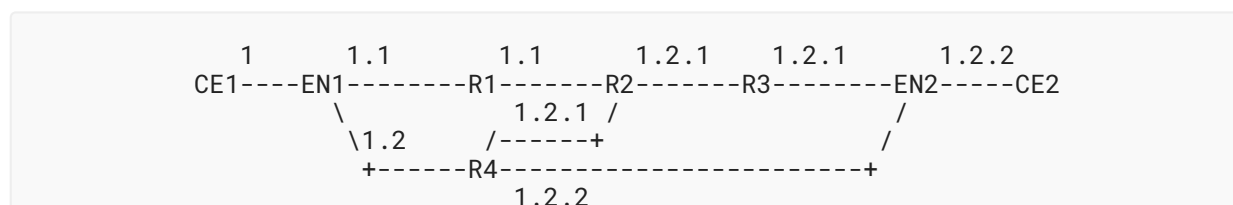


Figure 3: Example of Packet Flow Protected by DetNet

In [Figure 3](#), the numbers are used to identify the instance of a packet. Packet 1 is the original packet. Packets 1.1 and 1.2 are two first-generation copies of packet 1, packet 1.2.1 is a second-generation copy of packet 1.2, and so on. Note that these numbers never appear in the packet and

are not to be confused with sequence numbers, labels, or any other identifiers that appear in the packet. They simply indicate the generation number of the original packet so that its passage through the network fragment can be identified for the reader.

Customer Equipment device CE1 sends a packet into the DetNet-enabled 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 R1 may be directly attached to EN1, or there may be one or more nodes on the path that, for clarity, are not shown in [Figure 3](#). The same holds true for any other path between two DetNet entities as shown in the figure.

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 packet copy 1.2.1 later on, the copy is discarded.

The above is of course illustrative of many network scenarios that can be configured.

This example also illustrates a 1:1 protection scheme, meaning there is traffic over each segment of the end-to-end path. Local DetNet relay nodes determine which packets are eliminated and which packets are forwarded. A 1+1 scheme where only one path is used for traffic at a time could use the same topology. In this case, there is no PRF, and traffic is switched upon detection of failure. An OAM scheme that monitors the paths to detect the loss of a path or traffic is required to initiate the switch. A POF may still be used in this case to prevent misordering of packets. In both cases, the protection paths are established and maintained for the duration of the DetNet service.

### 3.5.2.2. Path Differential Delay

In the preceding example, proper operation of duplicate elimination and the reordering of packets are dependent on the number of out-of-order packets that can be buffered and the difference in delay of the arriving packets. DetNet uses flow-specific requirements (e.g., maximum number of out-of-order packets, maximum latency of the flow) for configuration of POF-related buffers. If the differential delay between paths is excessively large or there is excessive misordering of the packets, then packets may be dropped instead of being reordered. Likewise, the PEF uses the sequence number to identify duplicate packets, and large differential delays combined with high numbers of packets may exceed the PEF's ability to work properly.

### 3.5.2.3. Ring Service Protection

Ring protection may also be supported if the underlying technology supports it. Many of the same concepts apply; however, rings are normally 1+1 protection for data efficiency reasons. [RFC8227] provides an example of an MPLS Transport Profile (MPLS-TP) data plane that supports ring protection.

### 3.5.3. Aggregation Considerations

The DetNet data plane also allows for the aggregation of DetNet flows, which can improve scalability by reducing the per-hop state. How this is accomplished is data plane or control plane dependent. When DetNet flows are aggregated, transit nodes provide service to the aggregate and not on a per-DetNet-flow basis. When aggregating DetNet flows, the flows should be compatible, i.e., the same or very similar QoS and CoS characteristics. In this case, nodes performing aggregation will ensure that per-flow service requirements are achieved.

If bandwidth reservations are used, the reservation should be the sum of all the individual reservations; in other words, the reservations should not add up to an oversubscription of bandwidth reservation. If maximum delay bounds are used, the system should ensure that the aggregate does not exceed the delay bounds of the individual flows.

When an encapsulation is used, the choice of reserving a maximum resource level and then tracking the services in the aggregated service or adjusting the aggregated resources as the services are added is implementation and technology specific.

DetNet flows at edges must be able to handle rejection to an aggregation group due to lack of resources as well as conditions where requirements are not satisfied.

#### 3.5.3.1. IP Aggregation

IP aggregation has both data plane and Controller Plane aspects. For the data plane, flows may be aggregated for treatment based on shared characteristics such as 6-tuple [RFC8939]. Alternatively, an IP encapsulation may be used to tunnel an aggregate number of DetNet flows between relay nodes.

#### 3.5.3.2. MPLS Aggregation

MPLS aggregation also has data plane and Controller Plane aspects. MPLS flows are often tunneled in a forwarding sub-layer, under the reservation associated with that MPLS tunnel.

### 3.5.4. End-System-Specific Considerations

Data flows requiring DetNet service are generated and terminated on end systems. Encapsulation depends on the application and its preferences. For example, in a DetNet MPLS domain, the sub-layer functions use the d-CWs, S-Labels, and F-Labels [DetNet-MPLS] to provide DetNet services. However, an application may exchange further flow-related parameters (e.g., timestamps) that are not provided by DetNet functions.

As a general rule, DetNet domains are capable of forwarding any DetNet flows, and the DetNet domain does not mandate the encapsulation format for end systems or edge nodes. Unless some form of proxy is present, end systems peer with similar end systems using the same application encapsulation format. For example, as shown in [Figure 4](#), IP applications peer with IP applications, and Ethernet applications peer with Ethernet applications.

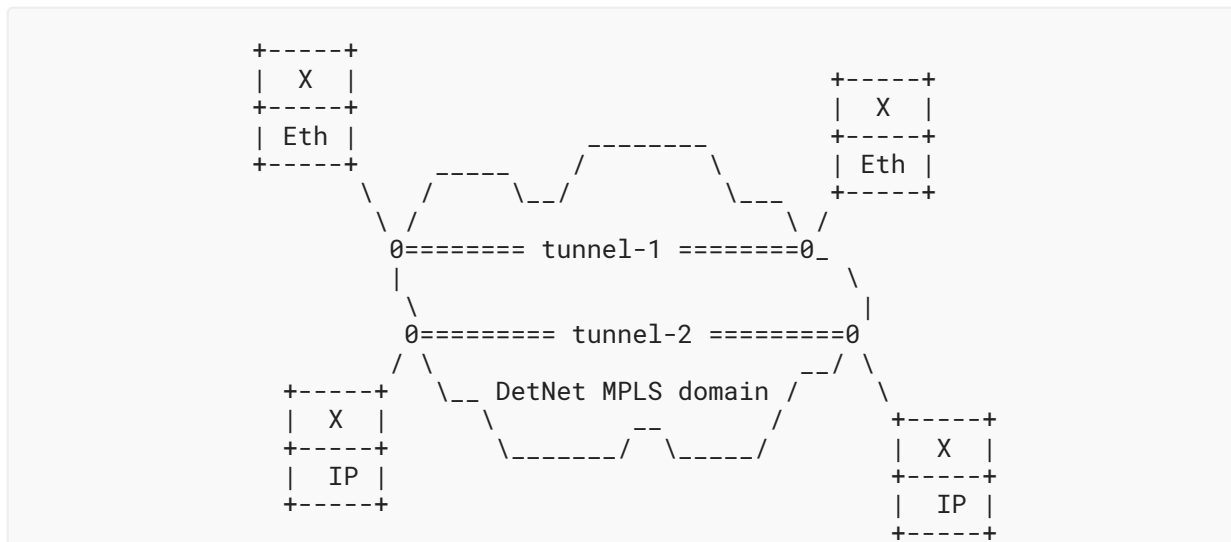


Figure 4: End Systems and the DetNet MPLS Domain

### 3.5.5. Sub-network Considerations

Any of the DetNet service types may be transported by another DetNet service. MPLS nodes may be interconnected by different sub-network technologies, which may include point-to-point links. Each of these sub-network technologies needs to provide appropriate service to DetNet flows. In some cases, e.g., on dedicated point-to-point links or TDM technologies, all that is required is for a DetNet node to appropriately queue its output traffic. In other cases, DetNet nodes will need to map DetNet flows to the flow semantics (i.e., identifiers) and mechanisms used by an underlying sub-network technology. [Figure 5](#) shows several examples of sub-network encapsulations that can be used to carry DetNet MPLS flows over different sub-network technologies. L2 represents a generic Layer 2 encapsulation that might be used on a point-to-point link. TSN represents the encapsulation used on an IEEE 802.1 TSN network, as described in [[DetNet-MPLS-over-TSN](#)]. UDP/IP represents the encapsulation used on a DetNet IP PSN, as described in [[DetNet-MPLS-over-UDP-IP](#)].

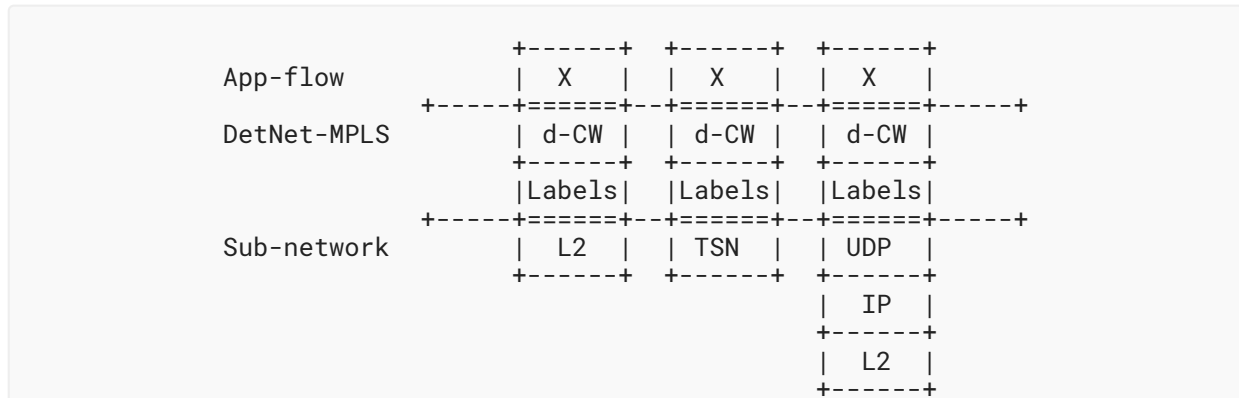


Figure 5: Example DetNet MPLS Encapsulations in Sub-networks

## 4. Controller Plane (Management and Control) Considerations

### 4.1. DetNet Controller Plane Requirements

The Controller Plane corresponds to the aggregation of the Control and Management Planes discussed in [RFC7426] and [RFC8655]. While more details regarding any DetNet Controller Plane are out of scope for this document, there are particular considerations and requirements for the Controller Plane that result from the unique characteristics of the DetNet architecture and data plane as defined herein.

The primary requirements of the DetNet Controller Plane are that it must be able to:

- Instantiate DetNet flows in a DetNet domain (which may, for example, include some or all of the following: explicit path determination, link bandwidth reservations, restricting flows to IEEE 802.1 TSN links, node buffer and other resource reservations, specification of required queuing disciplines along the path, ability to manage bidirectional flows, etc.) as needed for a flow.
- In the case of MPLS, manage DetNet S-Label and F-Label allocation and distribution. In cases where the DetNet MPLS encapsulation is being used, see [DetNet-MPLS].
- Support DetNet flow aggregation.
- Advertise static and dynamic node and link resources such as capabilities and adjacencies to other network nodes (for dynamic signaling approaches) or to network controllers (for centralized approaches).
- Scale to handle the number of DetNet flows expected in a domain (which may require per-flow signaling or provisioning).
- Provision flow identification information at each of the nodes along the path. Flow identification may differ, depending on the location in the network and the DetNet functionality (e.g., transit node vs. relay node).

These requirements, as stated earlier, could be satisfied using distributed control protocol signaling (such as RSVP-TE), centralized network management provisioning mechanisms (BGP, PCEP, YANG, [DetNet-Flow-Info], etc.), or hybrid combinations of the two, and could also make use of MPLS-based segment routing.

In the abstract, the results of either distributed signaling or centralized provisioning are equivalent from a DetNet data plane perspective – flows are instantiated, explicit routes are determined, resources are reserved, and packets are forwarded through the domain using the DetNet data plane.

However, from a practical and implementation standpoint, Controller Plane alternatives are not equivalent at all. Some approaches are more scalable than others in terms of signaling load on the network. Some alternatives can take advantage of global tracking of resources in the DetNet domain for better overall network resource optimization. Some solutions are more resilient than others if link, node, or management equipment failures occur. While a detailed analysis of the control plane alternatives is out of scope for this document, the requirements from this document can be used as the basis of a future analysis of the alternatives.

## 4.2. Generic Controller Plane Considerations

This section covers control plane considerations that are independent of the data plane technology used for DetNet service delivery.

While the management plane and the control plane are traditionally considered separately, from a data plane perspective, there is no practical difference based on the origin of flow-provisioning information, and the DetNet architecture [RFC8655] refers to these collectively as the "Controller Plane". This document therefore does not distinguish between information provided by distributed control plane protocols (e.g., RSVP-TE [RFC3209] [RFC3473]) or centralized network management mechanisms (e.g., RESTCONF [RFC8040], YANG [RFC7950], PCEP [PCECC]), or any combination thereof. Specific considerations and requirements for the DetNet Controller Plane are discussed in [Section 4.1](#).

Each respective data plane document also covers the control plane considerations for that technology. For example, [RFC8939] also covers IP control plane normative considerations, and [DetNet-MPLS] also covers MPLS control plane normative considerations.

### 4.2.1. Flow Aggregation Control

Flow aggregation means that multiple App-flows are served by a single new DetNet flow. There are many techniques to achieve aggregation. For example, in the case of IP, IP flows that share 6-tuple attributes or flow identifiers at the DetNet sub-layer can be grouped. Another example includes aggregation accomplished through the use of hierarchical LSPs in MPLS and tunnels.

Control of aggregation involves a set of procedures listed here. Aggregation may use some or all of these capabilities, and the order may vary:

Traffic engineering resource collection and distribution:

Available resources are tracked through control plane or management plane databases and distributed amongst controllers or nodes that can manage resources.

Path computation and resource allocation:

When DetNet services are provisioned or requested, one or more paths meeting the requirements are selected and the resources verified and recorded.

Resource assignment and data plane coordination:

The assignment of resources along the path depends on the technology and includes assignment of specific links, coordination of queuing, and other traffic management capabilities such as policing and shaping.

Assigned resource recording and updating:

Depending on the specific technology, the assigned resources are updated and distributed in the databases, preventing oversubscription.

#### 4.2.2. Explicit Routes

Explicit routes are used to ensure that packets are routed through the resources that have been reserved for them and hence provide the DetNet application with the required service. A requirement for the DetNet Controller Plane will be the ability to assign a particular identified DetNet IP flow to a path through the DetNet domain that has been assigned the required per-node resources. This provides the appropriate traffic treatment for the flow and also includes particular links as a part of the path that are able to support the DetNet flow. For example, by using IEEE 802.1 TSN links (as discussed in [[DetNet-MPLS-over-TSN](#)]), DetNet parameters can be maintained. Further considerations and requirements for the DetNet Controller Plane are discussed in [Section 4.1](#).

Whether configuring, calculating, and instantiating these routes is a single-stage or multi-stage process, or is performed in a centralized or distributed manner, is out of scope for this document.

There are several approaches that could be used to provide explicit routes and resource allocation in the DetNet forwarding sub-layer. For example:

- The path could be explicitly set up by a controller that calculates the path and explicitly configures each node along that path with the appropriate forwarding and resource allocation information.
- The path could use a distributed control plane such as [RSVP \[RFC2205\]](#) or [RSVP-TE \[RFC3473\]](#) extended to support DetNet IP flows.
- The path could be implemented using IPv6-based segment routing when extended to support resource allocation.

See [Section 4.1](#) for further discussion of these alternatives. In addition, [[RFC2386](#)] contains useful background information on QoS-based routing, and [[RFC5575](#)] (which will be updated by [[Flow-Spec-Rules](#)]) discusses a specific mechanism used by BGP for traffic flow specification and policy-based routing.



### 4.2.3. Contention Loss and Jitter Reduction

This document does not specify the mechanisms needed to eliminate packet contention or packet loss or to reduce jitter for DetNet flows at the DetNet forwarding sub-layer. The ability to manage node and link resources to be able to provide these functions is a necessary part of the DetNet Controller Plane. It is also necessary to be able to control the required queuing mechanisms used to provide these functions along a flow's path through the network. See [\[RFC8939\]](#) and [Section 4.1](#) for further discussion of these requirements. Some forms of protection may minimize packet loss or change jitter characteristics in the cases where packets are reordered when out-of-order packets are received at the service sub-layer.

### 4.2.4. Bidirectional Traffic

In many cases, DetNet flows can be considered unidirectional and independent. However, there are cases where the DetNet service requires bidirectional traffic from a DetNet application service perspective. IP and MPLS typically treat each direction separately and do not force interdependence of each direction. The IETF MPLS Working Group has studied bidirectional traffic requirements. The definitions provided in [\[RFC5654\]](#) are useful to illustrate terms such as 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 forwarding path. This is analogous to standard IP routing. MPLS defines a point-to-point co-routed bidirectional LSP as an associated bidirectional LSP that 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 reservations may differ in each direction. The concepts of associated bidirectional flows and co-routed bidirectional flows can also be applied to DetNet IP flows.

While the DetNet IP data plane must support bidirectional DetNet flows, there are no special bidirectional features with respect to the data plane other than the need for the two directions of a co-routed bidirectional flow to take the same path. That is to say, bidirectional DetNet flows are solely represented at the management plane and control plane levels, without specific support or knowledge within the DetNet data plane. Fate-sharing and associated or co-routed bidirectional flows can be managed at the control level.

DetNet's use of PREOF may increase the complexity of using co-routing bidirectional flows, because if PREOF are used, the replication points in one direction would have to match the elimination points in the other direction, and vice versa. In such cases, the optimal points for these functions in one direction may not match the optimal points in the other, due to network and traffic constraints. Furthermore, due to the per-packet service protection nature, bidirectional forwarding may not be ensured. The first packet of received member flows is selected by the elimination function independently of which path it has taken through the network.

Control and management mechanisms need to support bidirectional flows, but the specification of such mechanisms is out of scope for this document. Example control plane solutions for MPLS can be found in [RFC3473], [RFC6387], and [RFC7551]. These requirements are included in [Section 4.1](#).

### 4.3. Packet Replication, Elimination, and Ordering Functions (PREOF)

The Controller Plane protocol solution required for managing the processing of PREOF is outside the scope of this document. That said, it should be noted that the ability to determine, for a particular flow, optimal packet replication and elimination points in the DetNet domain requires explicit support. There may be existing capabilities that can be used or extended -- for example, GMPLS end-to-end recovery [RFC4872] and GMPLS segment recovery [RFC4873].

## 5. Security Considerations

Security considerations for DetNet are described in detail in [DetNet-Security]. General security considerations for the DetNet architecture are described in [RFC8655]. This section considers architecture-level DetNet security considerations applicable to all data planes.

Part of what makes DetNet unique is its ability to provide specific and reliable QoS (delivering data flows with extremely low packet loss rates and bounded end-to-end delivery latency), and the security-related aspects of protecting that QoS are similarly unique.

As for all communications protocols, the primary consideration 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-network using MACsec [IEEE802.1AE-2018] for Ethernet (Layer 2) flows.

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

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

In order to prevent or mitigate DetNet attacks on other networks via flow escape, edge devices can, for example, use existing mechanisms such as policing and shaping applied at the output of a DetNet domain.

## 6. IANA Considerations

This document has no IANA actions.

## 7. References

### 7.1. Normative References

- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.

### 7.2. Informative References

- [DetNet-Flow-Info] Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "DetNet Flow Information Model", Work in Progress, Internet-Draft, draft-ietf-detnet-flow-information-model-11, 21 October 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-flow-information-model-11>>.
- [DetNet-MPLS] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "DetNet Data Plane: MPLS", Work in Progress, Internet-Draft, draft-ietf-detnet-mpls-13, 11 October 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-mpls-13>>.
- [DetNet-MPLS-over-TSN] Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "DetNet Data Plane: MPLS over IEEE 802.1 Time Sensitive Networking (TSN)", Work in Progress, Internet-Draft, draft-ietf-detnet-mpls-over-tsn-04, 2 November 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-mpls-over-tsn-04>>.
- [DetNet-MPLS-over-UDP-IP] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "DetNet Data Plane: MPLS over UDP/IP", Work in Progress, Internet-Draft, draft-ietf-detnet-mpls-over-udp-ip-07, 11 October 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-mpls-over-udp-ip-07>>.
- [DetNet-Security] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", Work in Progress, Internet-Draft, draft-ietf-detnet-security-12, 2 October 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-security-12>>.
- [Flow-Spec-Rules] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", Work in Progress, Internet-Draft, draft-ietf-idr-rfc5575bis-27, 15 October 2020, <<https://tools.ietf.org/html/draft-ietf-idr-rfc5575bis-27>>.

- 
- [IEEE802.1AE-2018]** IEEE, "IEEE Standard for Local and metropolitan area networks-Media Access Control (MAC) Security", IEEE Std 802.1AE-2018, DOI 10.1109/IEEESTD.2018.8585421, December 2018, <<https://ieeexplore.ieee.org/document/8585421>>.
- [IEEE802.1TSNTG]** IEEE, "Time-Sensitive Networking (TSN) Task Group", <<https://1.ieee802.org/tsn/>>.
- [nwcrg]** IRTF, "Coding for efficient NetWork Communications Research Group (nwcrg)", <<https://datatracker.ietf.org/rg/nwcrg/about>>.
- [PCECC]** Li, Z., Peng, S., Negi, M. S., Zhao, Q., and C. Zhou, "PCEP Procedures and Protocol Extensions for Using PCE as a Central Controller (PCECC) of LSPs", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-extension-for-pce-controller-08, 1 November 2020, <<https://tools.ietf.org/html/draft-ietf-pce-pcep-extension-for-pce-controller-08>>.
- [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>>.
- [RFC2386]** Crawley, E., Nair, R., Rajagopalan, B., and H. Sandick, "A Framework for QoS-based Routing in the Internet", RFC 2386, DOI 10.17487/RFC2386, August 1998, <<https://www.rfc-editor.org/info/rfc2386>>.
- [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>>.
- [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>>.
- [RFC3670]** Moore, B., Durham, D., Strassner, J., Westerinen, A., and W. Weiss, "Information Model for Describing Network Device QoS Datapath Mechanisms", RFC 3670, DOI 10.17487/RFC3670, January 2004, <<https://www.rfc-editor.org/info/rfc3670>>.
- [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>>.
- [RFC4872]** Lang, J.P., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, DOI 10.17487/RFC4872, May 2007, <<https://www.rfc-editor.org/info/rfc4872>>.

- 
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, DOI 10.17487/RFC4873, May 2007, <<https://www.rfc-editor.org/info/rfc4873>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.
- [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>>.
- [RFC6387] Takacs, A., Berger, L., Caviglia, D., Fedyk, D., and J. Meuric, "GMPLS Asymmetric Bandwidth Bidirectional Label Switched Paths (LSPs)", RFC 6387, DOI 10.17487/RFC6387, September 2011, <<https://www.rfc-editor.org/info/rfc6387>>.
- [RFC7426] Haleplidis, E., Ed., Pentikousis, K., Ed., Denazis, S., Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-Defined Networking (SDN): Layers and Architecture Terminology", RFC 7426, DOI 10.17487/RFC7426, January 2015, <<https://www.rfc-editor.org/info/rfc7426>>.
- [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>>.
- [RFC7657] Black, D., Ed. and P. Jones, "Differentiated Services (Diffserv) and Real-Time Communication", RFC 7657, DOI 10.17487/RFC7657, November 2015, <<https://www.rfc-editor.org/info/rfc7657>>.
- [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>>.
- [RFC8227] Cheng, W., Wang, L., Li, H., van Helvoort, H., and J. Dong, "MPLS-TP Shared-Ring Protection (MSRP) Mechanism for Ring Topology", RFC 8227, DOI 10.17487/RFC8227, August 2017, <<https://www.rfc-editor.org/info/rfc8227>>.
- [RFC8939] Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939, November 2020, <<https://www.rfc-editor.org/info/rfc8939>>.
- [SRv6-Network-Prog] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-network-programming-26, 26 November 2020, <<https://tools.ietf.org/html/draft-ietf-spring-srv6-network-programming-26>>.

## Acknowledgements

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

## Contributors

The following people contributed substantially to the content of this document:

Don Fedyk

Jouni Korhonen

## Authors' Addresses

### **Balázs Varga (EDITOR)**

Ericsson

Budapest

Magyar Tudosok krt. 11.

1117

Hungary

Email: [balazs.a.varga@ericsson.com](mailto:balazs.a.varga@ericsson.com)

### **János Farkas**

Ericsson

Budapest

Magyar Tudosok krt. 11.

1117

Hungary

Email: [janos.farkas@ericsson.com](mailto:janos.farkas@ericsson.com)

### **Lou Berger**

LabN Consulting, L.L.C.

Email: [lberger@labn.net](mailto:lberger@labn.net)

### **Andrew G. Malis**

Malis Consulting

Email: [agmalis@gmail.com](mailto:agmalis@gmail.com)

### **Stewart Bryant**

Futurewei Technologies

Email: [sb@stewartbryant.com](mailto:sb@stewartbryant.com)