DetNet
Internet-Draft
Intended status: Standards Track
Expires: September 11, 2019

J. Korhonen, Ed.

B. Varga, Ed.
Ericsson
March 10, 2019

DetNet IP Data Plane Encapsulation
draft-ietf-detnet-dp-sol-ip-02

Abstract

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 11, 2019.

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 4
3. DetNet IP Data Plane Overview 5
4. DetNet IP Data Plane Considerations 7
4.1. End-System Specific Considerations 8
4.2. DetNet Domain-Specific Considerations 9
4.2.1. DetNet Routers 10
4.3. Networks With Multiple Technology Segments 11
4.4. OAM 12
4.5. Class of Service 12
4.6. Quality of Service 13
4.7. Cross-DetNet Flow Resource Aggregation 14
4.8. Time Synchronization 14
5. Management and Control Considerations 15
5.1. Flow Identification and Aggregation 15
5.2. Explcit Routes 16
5.3. Contention Loss and Jitter Reduction 16
5.4. Bidirectional Traffic 17
5.5. DetNet Controller (Control and Management) Plane Requirements 17
6. DetNet IP Data Plane Procedures 19
6.1. DetNet IP Flow Identification Procedures 19
6.1.1. IP Header Information 19
6.1.2. Other Protocol Header Information 21
6.1.3. Flow Identification Management and Control Information 22
6.2. Forwarding Procedures 23
6.3. DetNet IP Traffic Treatment Procedures 23
6.4. Aggregation Considerations 23
7. IP over DetNet MPLS 24
7.1. IP Over DetNet MPLS Data Plane Scenarios 24
7.2. DetNet IP over DetNet MPLS Encapsulation 27
7.3. DetNet IP over DetNet MPLS Flow Identification Procedures 29
7.4. DetNet IP over DetNet MPLS Traffic Treatment Procedures . 29
8. Mapping DetNet IP Flows to IEEE 802.1 TSN 29
8.1. TSN Stream ID Mapping 31
8.2. TSN Usage of FRER 33
8.3. Procedures 34
8.4. Management and Control Implications 34
9. Security Considerations 36
10. IANA Considerations 36
11. Contributors 36

12. Acknowledgements 38
13. References 38
13.1. Normative references 38
13.2. Informative references 40
Appendix A. Example of DetNet Data Plane Operation 43
Appendix B. Example of Pinned Paths Using IPv6 43
Authors' Addresses 43

1. Introduction

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

This document specifies the DetNet data plane operation for IP hosts and routers that provide DetNet service to IP encapsulated data. No DetNet specific encapsulation is defined to support IP flows, rather existing IP and higher layer protocol header information is used to support flow identification and DetNet service delivery.

The DetNet Architecture decomposes the DetNet related data plane functions 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 is used to provides congestion protection (low loss, assured latency, and limited reordering). As no DetNet specific headers are added to support DetNet IP flows, only the forwarding sub-layer functions are supported using the DetNet IP defined by this document. Service protection can be provided on a per sub-net basis using technologies such as MPLS [I-D.ietf-detnet-dp-sol-mpls] and IEEE802.1 TSN.

This document provides an overview of the DetNet IP data plane in Section 3, considerations that apply to providing DetNet services via the DetNet IP data plane in Section 4 and Section 5. Section 6 provides the procedures for hosts and routers that support IP-based DetNet services. Finally, Section 8 provides rules for mapping IP-based DetNet flows to IEEE 802.1 TSN streams.

2. Terminology

2.1. Terms Used In This Document

This document uses the terminology and concepts established in the DetNet architecture [I-D.ietf-detnet-architecture], and the reader is assumed to be familiar with that document and its terminology.

2.2. Abbreviations

The following abbreviations used in this document:

CE	Customer Edge equipment.
CoS	Class of Service.
DetNet	Deterministic Networking.
DF	DetNet Flow.
L2	Layer-2.
L3	Layer-3.
LSP	Label-switched path.
MPLS	Multiprotocol Label Switching.
OAM	Operations, Administration, and Maintenance.
PE	Provider Edge.
PREOF	Packet Replication, Ordering and Elimination Function.
PSN	Packet Switched Network.
PW	Pseudowire.
QoS	Quality of Service.
TE	Traffic Engineering.
TSN	Time-Sensitive Networking, TSN is a Task Group of the IEEE 802.1 Working Group.

2.3. Requirements Language

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

3. DetNet IP Data Plane Overview

This document describes how IP is used by DetNet nodes, i.e., hosts and routers, to identify DetNet flows and provide a DetNet service. From a data plane perspective, an end-to-end IP model is followed. As mentioned above, existing IP and higher layer protocol header information is used to support flow identification and DetNet service delivery.

DetNet uses "6-tuple" based flow identification, where "6-tuple" refers to information carried in IP and higher layer protocol headers. General background on the use of IP headers, and "5-tuples", to identify flows and support Quality of Service (QoS) can be found in [RFC3670]. [RFC7657] also provides useful background on the delivery differentiated services (DiffServ) and "6-tuple" based flow identification.

DetNet flow aggregation may be enabled via the use of wildcards, masks, prefixes and ranges. IP tunnels may also be used to support flow aggregation. In these cases, it is expected that DetNet aware intermediate nodes will provide DetNet service assurance on the aggregate through resource allocation and congestion control mechanisms.

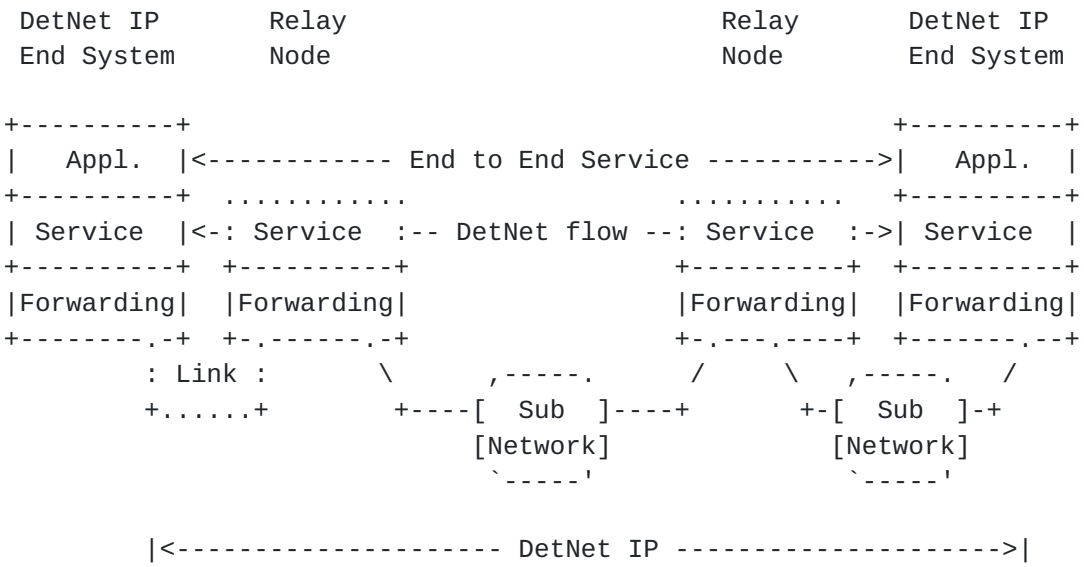

Figure 1: A Simple DetNet (DN) Enabled IP Network

Figure 1 illustrates a DetNet enabled IP network. The DetNet enabled end systems originate IP encapsulated traffic that is identified as DetNet flows, relay nodes understand the forwarding requirements of the DetNet flow and ensure that node, interface and sub-network resources are allocated to ensure DetNet service requirements. The

dotted line around the Service component of the Relay Nodes indicates that the transit routers are DetNet service aware but do not perform any DetNet service sub-layer function, e.g., PREOF. IEEE 802.1 TSN is an example sub-network type which can provide support for DetNet flows and service. The mapping of DetNet IP flows to TSN streams and TSN protection mechanisms is covered in Section 8.

Note: The sub-network can represent a TSN, MPLS or IP network segment.

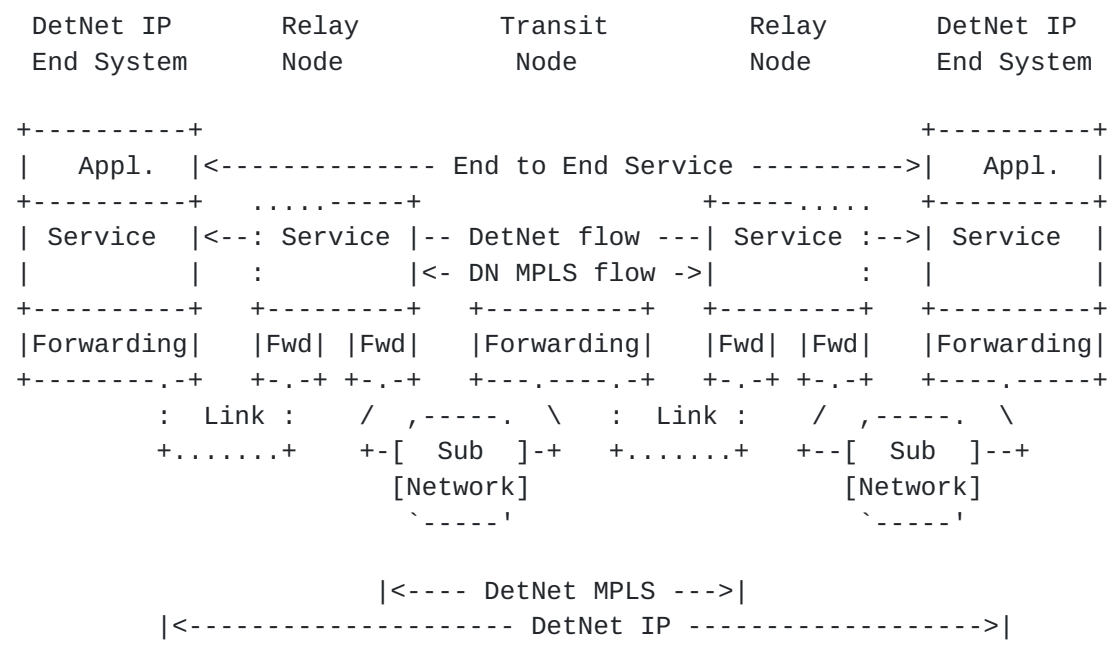

Figure 2: DetNet IP Over DetNet MPLS Network

Figure 2 illustrates a variant of Figure 1, with an MPLS based DetNet network as a sub-network between the relay nodes. It shows a more complex DetNet enabled IP network where an IP flow is mapped to one or more PWs and MPLS (TE) LSPs. The end systems still originate IP encapsulated traffic that is identified as DetNet flows. The relay nodes follow procedures defined in Section 7 to map each DetNet flow to MPLS LSPs. While not shown, relay nodes can provide service sub-layer functions such as PREOF using DetNet over MPLS, and this is indicated by the solid line for the MPLS facing portion of the Service component. Note that the Transit node is MPLS (TE) LSP aware and performs switching based on MPLS labels, and need not have any specific knowledge of the DetNet service or the corresponding DetNet flow identification. See Section 7 for details on the mapping of IP flows to MPLS, and [I-D.ietf-detnet-dp-sol-mpls] for general support of DetNet services using MPLS.

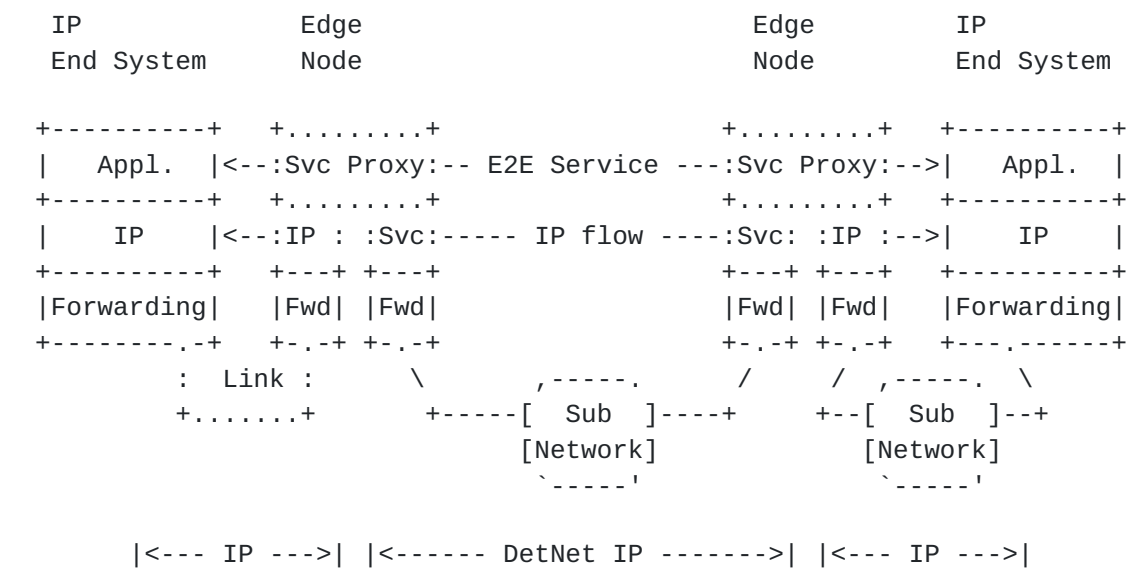

Figure 3: Non-DetNet aware IP end systems with DetNet IP Domain

Figure 3 illustrates another variant of Figure 1 where the end systems are not DetNet aware. In this case, edge nodes sit at the boundary of the DetNet domain and provide DetNet service proxies for the end applications by initiating and terminating DetNet service for the application's IP flows. The existing header information or an approach such as described in Section 4.7 can be used to support DetNet flow identification.

Non-DetNet and DetNet IP packets are identical on the wire. From data plane perspective, the only difference is that there is flow-associated DetNet information on each DetNet node that defines the flow related characteristics and required forwarding behavior. As shown above, edge nodes provide a Service Proxy function that "associates" one or more IP flows with the appropriate DetNet flow-specific information and ensures that the receives the proper traffic treatment within the domain.

Note: The operation of IEEE802.1 TSN end systems over DetNet enabled IP networks is not described in this document. While TSN flows could be encapsulated in IP packets by an IP End System or DetNet Edge Node in order to produce DetNet IP flows, the details of such are out of scope of this document.

4. DetNet IP Data Plane Considerations

This section provides informative considerations related to providing DetNet service to flows which are identified based on their header information. At a high level, the following are provided on a per flow basis:

Congestion protection and latency control:

Usage of allocated resources (queuing, policing, shaping) to ensure that the congestion-related loss and latency/jitter requirements of a DetNet flow are met.

Explicit routes:

Use of a specific path for a flow. This limits misordering and can improve delivery of deterministic latency.

Service protection:

Which in the case of this document translates to changing the explicit path after a failure is detected in order to restore delivery of the required DetNet service characteristics. Path changes, even in the case of failure recovery, can lead to the out of order delivery of data.

Note: DetNet PREOF is not provided by the mechanisms defined in this document.

Load sharing:

Generally, distributing packets of the same DetNet flow over multiple paths is not recommended. Such load sharing, e.g., via ECMP or UCMP, impacts ordering and end-to-end jitter.

Troubleshooting:

For example, to support identification of misbehaving flows.

Recognize flow(s) for analytics:

For example, increase counters.

Correlate events with flows:

For example, unexpected loss.

4.1. End-System Specific Considerations

Data-flows requiring DetNet service are generated and terminated on end systems. This document deals only with IP end systems. The protocols used by an IP end system are specific to an application and end systems peer with end systems using the same application encapsulation format. This said, DetNet's use of 6-tuple IP flow identification means that DetNet must be aware of not only the format

of the IP header, but also of the next protocol carried within an IP packet.

When IP end systems are DetNet aware, no application-level or service-level proxy functions are needed inside the DetNet domain. For DetNet unaware IP end systems service-level proxy functions are needed inside the DetNet domain.

End systems need to ensure that DetNet service requirements are met when processing packets associated with a DetNet flow. When forwarding packets, this means that packets are appropriately shaped on transmission and received appropriate traffic treatment on the connected sub-network, see Section 4.6 and Section 4.2.1 for more details. When receiving packets, this means that there are appropriate local node resources, e.g., buffers, to receive and process a DetNet flow packets.

4.2. DetNet Domain-Specific Considerations

As a general rule, DetNet IP domains need to be able to forward any DetNet flow identified by the IP 6-tuple. Doing otherwise would limit end system encapsulation format. From a practical standpoint this means that all nodes along the end-to-end path of a DetNet flows need to agree on what fields are used for flow identification, and the transport protocols (e.g., TCP/UDP/IPsec) which can be used to identify 6-tuple protocol ports.

From a connection type perspective two scenarios are identified:

1. DN attached: end system is directly connected to an edge node or end system is behind a sub-network. (See ES1 and ES2 in figure below)

2. DN integrated: end system is part of the DetNet domain. (See ES3 in figure below)

L3 (IP) end systems may use any of these connection types. DetNet domain allows communication between any end-systems using the same encapsulation format, independent of their connection type and DetNet capability. DN attached end systems have no knowledge about the DetNet domain and its encapsulation format. See Figure 4 for L3 end system connection scenarios.

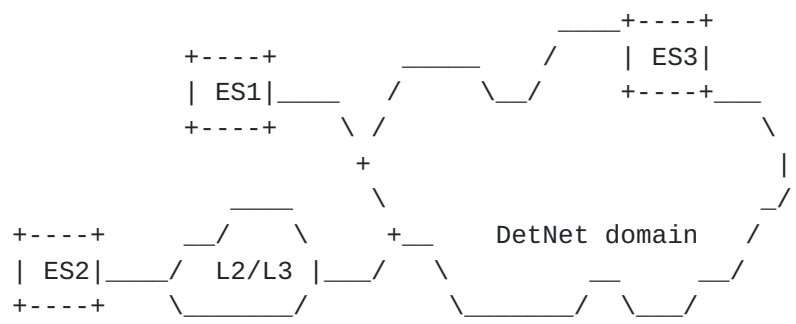

Figure 4: Connection types of L3 end systems

4.2.1. DetNet Routers

Within a DetNet domain, the DetNet enabled IP Routers interconnect links and sub-networks to support end-to-end delivery of DetNet flows. From a DetNet architecture perspective, these routers are DetNet relays, as they must be DetNet service aware. Such routers identify DetNet flows based on the IP 6-tuple, and ensure that the DetNet service required traffic treatment is provided both on the node and on any attached sub-network.

This solution provides DetNet functions end to end, but does so on a per link and sub-network basis. Congestion protection and latency control and the resource allocation (queuing, policing, shaping) are supported using the underlying link / sub net specific mechanisms. However, service protections (packet replication and packet elimination functions) are not provided at the DetNet layer end to end. But such service protection can be provided on a per underlying L2 link and sub-network basis.

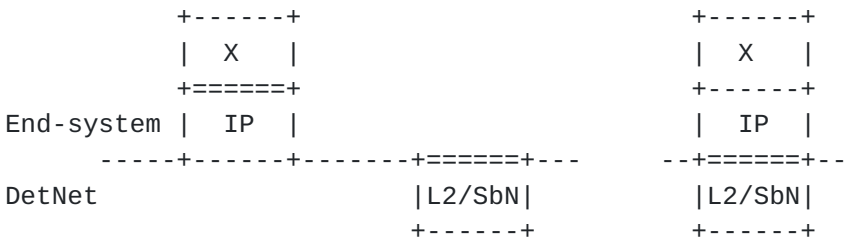

Figure 5: Encapsulation of DetNet Routing in simplified IP service L3 end-systems

The DetNet Service Flow is mapped to the link / sub-network specific resources using an underlying system specific means. This implies each DetNet aware node on path looks into the forwarded DetNet

Service Flow packet and utilize e.g., a 5- (or 6-) tuple to find out the required mapping within a node.

As noted earlier, the Service Protection is done within each link / sub-network independently using the domain specific mechanisms (due the lack of a unified end to end sequencing information that would be available for intermediate nodes). Therefore, service protection (if any) cannot be provided end-to-end, only within sub-networks. This is shown for a three sub-network scenario in Figure 6, where each sub-network can provide service protection between its borders.

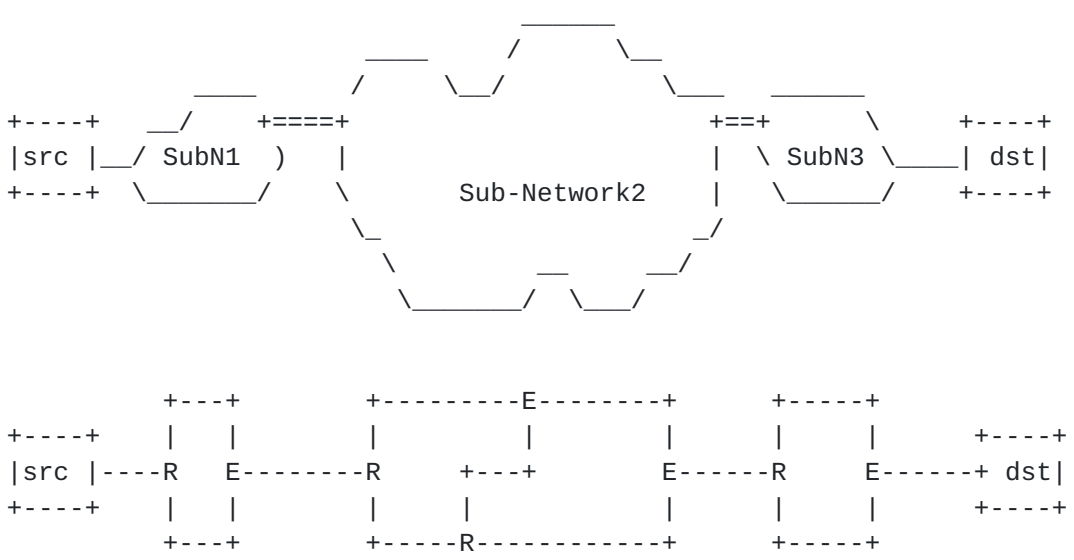

Figure 6: Replication and elimination in sub-networks for DetNet IP networks

If end to end service protection is desired that can be implemented, for example, by the DetNet end systems using Layer-4 (L4) transport protocols or application protocols. However, these are out of scope of this document.

4.3. Networks With Multiple Technology Segments

There are network scenarios, where the DetNet domain contains multiple technology segments (IEEE 802.1 TSN, MPLS) and all those segments are under the same administrative control (see Figure 7). Furthermore, DetNet nodes may be interconnected via TSN segments.

DetNet routers ensure that detnet service requirements are met per hop by allocating local resources, both receive and transmit, and by mapping the service requirements of each flow to appropriate sub-

network mechanisms. Such mapping is sub-network technology specific. The mapping of DetNet IP Flows to MPLS is covered Section 7. The mapping of IP DetNet Flows to IEEE 802.1 TSN is covered in Section 8.

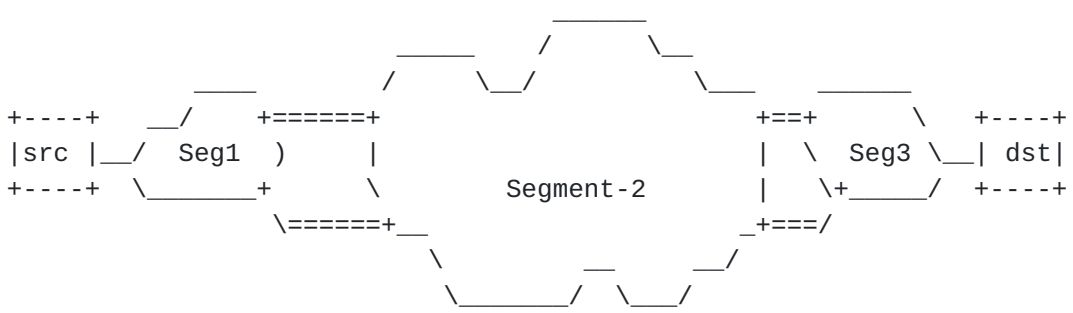

Figure 7: DetNet domains and multiple technology segments

4.4. OAM

[Editor's note: This section is TBD. OAM may be dropped from this document and left for future study.]

4.5. Class of Service

Class and quality of service, i.e., CoS and QoS, are terms that are often used interchangeably and confused. 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 IPv6 is provided using the standard differentiated services code point (DSCP) field [RFC2474] and related mechanisms. The 2-bit explicit congestion notification (ECN) [RFC3168] field MAY also be used.

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

4.6. Quality of Service

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

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

A baseline set of QoS capabilities for DetNet flows carried in PWs and MPLS can provided by MPLS with Traffic Engineering (MPLS-TE) [RFC3209] and [RFC3473]. TE LSPs can also support explicit routes (path pinning). Current service definitions for packet TE LSPs can be found in "Specification of the Controlled Load Quality of Service", [RFC2211], "Specification of Guaranteed Quality of Service", [RFC2212], and "Ethernet Traffic Parameters", [RFC6003]. Additional service definitions are expected in future documents to support the full range of DetNet services. In all cases, the existing label-based marking mechanisms defined for TE-LSPs and even E-LSPs are use to support the identification of flows requiring DetNet QoS.

QoS for DetNet service flows carried in IP MUST be provided locally by the DetNet-aware hosts and routers supporting DetNet flows. Such support will leverage the underlying network layer such as 802.1TSN. The traffic control mechanisms used to deliver QoS for IP encapsulated DetNet flows are expected to be defined in a future document. From an encapsulation perspective, the combination of the "6 tuple" i.e., the typical 5 tuple enhanced with the DSCP code, uniquely identifies a DetNet service flow.

Packets that are marked with a DetNet Class of Service value, but that have not been the subject of a completed reservation, can disrupt the QoS offered to properly reserved DetNet flows by using

resources allocated to the reserved flows. Therefore, the network nodes of a DetNet network must:

- Defend the DetNet QoS by discarding or remarking (to a non-DetNet CoS) packets received that are not the subject of a completed reservation.

- Not use a DetNet reserved resource, e.g. a queue or shaper reserved for DetNet flows, for any packet that does not carry a DetNet Class of Service marker.

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

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

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

4.8. Time Synchronization

While time synchronization can be important both from the perspective of operating the DetNet network itself and from the perspective of DetNet-based applications, time synchronization is outside the scope of this document. This said, a DetNet node can also support time synchronization or distribution mechanisms.

For example, [RFC8169] describes a method of recording the packet queuing time in an MPLS LSR on a packet by per packet basis and forwarding this information to the egress edge system. This allows compensation for any variable packet queuing delay to be applied at the packet receiver. Other mechanisms for IP networks are defined based on IEEE Standard 1588 [IEEE1588], such as ITU-T [G.8275.1] and [G.8275.2].

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

5. Management and Control Considerations

While management plane and control planes are traditionally considered separately, from the Data Plane perspective there is no practical difference based on the origin of flow provisioning information, and the DetNet architecture [I-D.ietf-detnet-architecture] refers to these collectively as the 'Controller Plane'. This document therefore does not distinguish between information provided by distributed control plane protocols, e.g., IGP routing protocols, or by centralized network management mechanisms, e.g., RestConf [RFC8040], YANG [RFC7950], and the Path Computation Element Communication Protocol (PCEP) [RFC8283] [I-D.ietf-teas-pce-native-ip] or any combination thereof. Specific considerations and requirements for the DetNet Controller Plane are discussed in Section 5.5.

5.1. Flow Identification and Aggregation

Section 3 introduces the use of the IP "6-tuple" for flow identification, and Section 4.6 goes on to discuss how identified flows use specific QoS mechanisms for flow-specific traffic treatment, including path control and resource allocation. Section 6.1 contains detailed DetNet IP flow identification procedures. Flow identification will play an important role for the DetNet controller plane.

Section 4.7 and Section 6.4 discuss the use of flow aggregation in DetNet. Flow aggregation can be accomplished using any of the 6-tuple fields defined in Section 6.1, using a DSCP identified traffic class or other field. It will be the responsibility of the DetNet controller plane to be able to properly provision the use of these aggregation mechanisms. These requirements are included in Section 5.5.

5.2. Explcit 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 nodal resources to provide the appropriate traffic treatment for the flow, and also to include 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 Section 8). Further considerations and requirements for the DetNet Controller Plane are discussed in Section 5.5.

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

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

- The path could be explicitly set up by a controller which calculates the path and explicitly configures each node along that path with the appropriate forwarding and resource allocation information.

- The path could be used 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 5.5 for further discussion of these alternatives. In addition, [RFC2386] contains useful background information on QoS-based routing, and [RFC5575] discusses a specific mechanism used by BGP for traffic flow specification and policy-based routing.

5.3. Contention Loss and Jitter Reduction

As discussed in Section 1, this document does not specify the mechanisms needed to eliminate contention loss or 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 will be a necessary part of the DetNet controller plane. It will also be necessary to be able to control the required queuing mechanisms used to provide these functions along a flow's path

through the network. See Section 6.3 and Section 5.5 for further discussion of these requirements.

5.4. Bidirectional Traffic

Some DetNet applications generate bidirectional traffic. Although this document discusses the DetNet IP data plane, MPLS definitions [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 which satisfies the additional constraint that its two unidirectional component LSPs follow the same path (in terms of both nodes and links) in both directions. An important property of co-routed bidirectional LSPs is that their unidirectional component LSPs share fate. In both types of bidirectional LSPs, resource 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 that bidirectional DetNet flows are solely represented at the management and control plane levels, without specific support or knowledge within the DetNet data plane. Fate sharing and associated vs. co-routed bidirectional flows can be managed at the control level.

Control and management mechanisms will need to support bidirectional flows, but the specification of such mechanisms are out of scope of this document. An example control plane solution for MPLS can be found in [RFC7551].

This is further discussed in Section 5.5.

5.5. DetNet Controller (Control and Management) Plane Requirements

While the definition of controller plane for DetNet is out of the scope of this document, there are particular considerations and requirements for such that result from the unique characteristics of the DetNet architecture [I-D.ietf-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 include some or all of 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.

- The ability to 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, and it may differ depending on the location in the network and the DetNet functionality.

These requirements, as stated earlier, could be satisfied using distributed control protocol signaling, centralized network management provisioning mechanisms, or hybrid combinations of the two, and could also make use of IPv6-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 IP data plane.

However, from a practical and implementation standpoint, they are not equivalent at all. Some approaches are more scalable than others in terms of signaling load on the network. Some can take advantage of global tracking of resources in the DetNet domain for better overall network resource optimization. Some 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 the scope of this document, the requirements from this document can be used as the basis of a later analysis of the alternatives.

6. DetNet IP Data Plane Procedures

This section provides DetNet IP data plane procedures. These procedures have been divided into the following areas: flow identification, forwarding and traffic treatment. Flow identification includes those procedures related to matching IP and higher layer protocol header information to DetNet flow (state) information and service requirements. Flow identification is also sometimes called Traffic classification, for example see [RFC5777]. Forwarding includes those procedures related to next hop selection and delivery. Traffic treatment includes those procedures related to providing an identified flow with the required DetNet service.

DetNet IP data plane procedures also have implications on the control and management of DetNet flows and these are also covered in this section. Specifically this section identifies a number of information elements that will require support via the management and control interfaces supported by a DetNet node. The specific mechanism used for such support is out of the scope of this document. A summary of the management and control related information requirements is included. Conformance language is not used in the summary as it applies to future mechanisms such as those that may be provided in YANG models [YANG-REF-TBD].

6.1. DetNet IP Flow Identification Procedures

IP and higher layer protocol header information is used to identify DetNet flows. All DetNet implementations that support this document MUST identify individual DetNet flows based on the set of information identified in this section. Note, that additional flow identification requirements, e.g., to support other higher layer protocols, may be defined in future.

The configuration and control information used to identify an individual DetNet flow MUST be ordered by an implementation. Implementations MUST support a fixed order when identifying flows, and MUST identify a DetNet flow by the first set of matching flow information.

Implementations of this document MUST support DetNet flow identification when the implementation is acting as a DetNet end systems, a relay node or as an edge node.

6.1.1. IP Header Information

Implementations of this document MUST support DetNet flow identification based on IP header information. The IPv4 header is defined in [RFC0791] and the IPv6 is defined in [RFC8200].

6.1.1.1. Source Address Field

Implementations of this document MUST support DetNet flow identification based on the Source Address field of an IP packet. Implementations SHOULD support longest prefix matching for this field, see [RFC1812] and [RFC7608]. Note that a prefix length of zero (0) effectively means that the field is ignored.

6.1.1.2. Destination Address Field

Implementations of this document MUST support DetNet flow identification based on the Destination Address field of an IP packet. Implementations SHOULD support longest prefix matching for this field, see [RFC1812] and [RFC7608]. Note that a prefix length of zero (0) effectively means that the field is ignored.

Note: any IP address value is allowed, including IP multicast destination address.

6.1.1.3. IPv4 Protocol and IPv6 Next Header Fields

Implementations of this document MUST support DetNet flow identification based on the IPv4 Protocol field when processing IPv4 packets, and the IPv6 Next Header Field when processing IPv6 packets. An implementation MUST support flow identification based based the next protocol values defined in Section 6.1.2. Other, non-zero values, MUST be used for flow identification. Implementations SHOULD allow for these fields to be ignored for a specific DetNet flow.

6.1.1.4. IPv4 Type of Service and IPv6 Traffic Class Fields

These fields are used to support Differentiated Services [RFC2474] and Explicit Congestion Notification [RFC3168]. Implementations of this document MUST support DetNet flow identification based on the IPv4 Type of Service field when processing IPv4 packets, and the IPv6 Traffic Class Field when processing IPv6 packets. Implementations MUST support bimask based matching, where one (1) values in the bitmask indicate which subset of the bits in the field are to be used in determining a match. Note that a zero (0) value as a bitmask effectively means that these fields are ignored.

6.1.1.5. IPv6 Flow Label Field

Implementations of this document SHOULD support identification of DetNet flows based on the IPv6 Flow Label field. Implementations that support matching based on this field MUST allow for this fields to be ignored for a specific DetNet flow. When this fields is used to identify a specific DetNet flow, implementations MAY exclude the

IPv6 Next Header field and next header information as part of DetNet flow identification.

6.1.2. Other Protocol Header Information

Implementations of this document MUST support DetNet flow identification based on header information identified in this section. Support for TCP, UDP and IPsec flows are defined. Future documents are expected to define support for other protocols.

6.1.2.1. TCP and UDP

DetNet flow identification for TCP [RFC0793] and UDP [RFC0768] is done based on the Source and Destination Port fields carried in each protocol's header. These fields share a common format and common DetNet flow identification procedures.

6.1.2.1.1. Source Port Field

Implementations of this document MUST support DetNet flow identification based on the Source Port field of a TCP or UDP packet. Implementations MUST support flow identification based on a particular value carried in the field, i.e., an exact. Implementations SHOULD support range-based port matching. Implementation MUST also allow for the field to be ignored for a specific DetNet flow.

6.1.2.1.2. Destination Port Field

Implementations of this document MUST support DetNet flow identification based on the Destination Port field of a TCP or UDP packet. Implementations MUST support flow identification based on a particular value carried in the field, i.e., an exact. Implementations SHOULD support range-based port matching. Implementation MUST also allow for the field to be ignored for a specific DetNet flow.

6.1.2.2. IPsec AH and ESP

IPsec Authentication Header (AH) [RFC4302] and Encapsulating Security Payload (ESP) [RFC4303] share a common format for the Security Parameters Index (SPI) field. Implementations MUST support flow identification based on a particular value carried in the field, i.e., an exact. Implementation SHOULD also allow for the field to be ignored for a specific DetNet flow.

6.1.3. Flow Identification Management and Control Information

The following summarizes the set of information that is needed to identify an individual DetNet flow:

- IPv4 and IPv6 source address field.
- IPv4 and IPv6 source address prefix length, where a zero (0) value effectively means that the address field is ignored.
- IPv4 and IPv6 destination address field.
- IPv4 and IPv6 destination address prefix length, where a zero (0) effectively means that the address field is ignored.
- IPv4 protocol field. A limited set of values is allowed, and the ability to ignore this field, e.g., via configuration of the value zero (0), is desirable.
- IPv6 next header field. A limited set of values is allowed, and the ability to ignore this field, e.g., via configuration of the value zero (0), is desirable.
- IPv4 Type of Service and IPv6 Traffic Class Fields.
- IPv4 Type of Service and IPv6 Traffic Class Field Bitmask, where a zero (0) effectively means that theses fields are ignored.
- IPv6 flow label field. This field can be optionally used for matching. When used, can be exclusive of matching against the next header field.
- TCP and UDP Source Port. Exact and wildcard matching is required. Port ranges can optionally be used.
- TCP and UDP Destination Port. Exact and wildcard matching is required. Port ranges can optionally be used.

This information MUST be provisioned per DetNet flow via configuration, e.g., via the controller plane described in Section 5.

Information identifying a DetNet flow is ordered and implementations use the first match. This can, for example, be used to provide a DetNet service for a specific UDP flow, with unique Source and Destination Port field values, while providing a different service for all other flows with that same UDP Destination Port value.

6.2. Forwarding Procedures

General requirements for IP nodes are defined in [RFC1122], [RFC1812] and [RFC6434], and are not modified by this document. The typical next-hop selection process is impacted by DetNet. Specifically, implementations of this document SHALL use management and control information to select the one or more outgoing interfaces and next hops to be used for a packet belonging to a DetNet flow.

The use of multiple paths or links, e.g., ECMP, to support a single DetNet flow is NOT RECOMMENDED. ECMP MAY be used for non-DetNet flows within a DetNet domain.

The above implies that management and control functions will be defined to support this requirement, e.g., see [YANG-REF-TBD].

6.3. DetNet IP Traffic Treatment Procedures

Implementations if this document MUST ensure that a DetNet flow receives the traffic treatment that is provisioned for it via configuration or the controller plane, e.g., via [YANG-REF-TBD]. General information on DetNet service can be found in [I-D.ietf-detnet-flow-information-model]. 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 as MPLS Section 7 and IEEE802.1 TSN Section 8. Other than in the TSN case, 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.

6.4. Aggregation Considerations

The use of prefixes, wildcards, bitmasks, and port ranges allows a DetNet node to aggregate DetNet flows. This aggregation can take place within a single node, when that node maintains state about both the aggregated and component flows. It can also take place between nodes, where one node maintains state about only flow aggregates while the other node maintains state on all or a portion of the component flows. In either case, the management or control function that provisions the aggregate flows must ensure that adequate resources are allocated and configured to provide combined service requirements of the component flows. As DetNet is concerned about latency and jitter, more than just bandwidth needs to be considered.

7. IP over DetNet MPLS

This section defines how IP encapsulated flows are carried over a DetNet MPLS data plane as defined in [I-D.ietf-detnet-dp-sol-mpls]. As Non-DetNet and DetNet IP packets are identical on the wire, this section is applicable to any node that supports IP over DetNet MPLS, and this section refers to both cases as DetNet IP over DetNet MPLS.

7.1. IP Over DetNet MPLS Data Plane Scenarios

This section provides example uses of IP over DetNet MPLS for illustrative purposes.

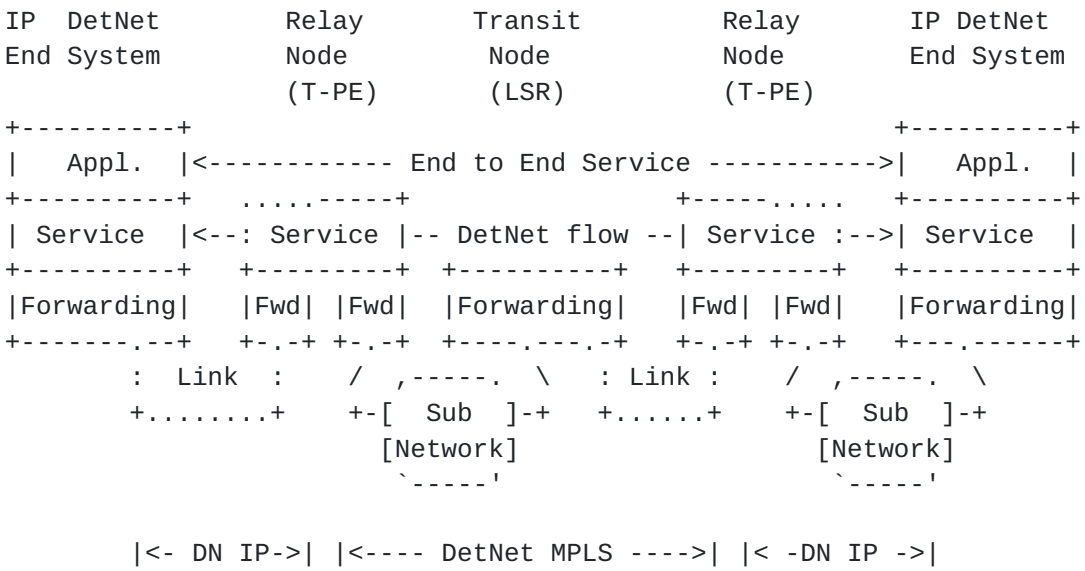

Figure 8: DetNet IP Over MPLS Network

Figure 8 illustrates DetNet enabled End Systems (hosts), connected to DetNet (DN) enabled IP networks, operating over a DetNet aware MPLS network. In this figure, Relay nodes sit at the boundary of the MPLS domain since the non-MPLS domain is DetNet aware. This figure is very similar to the DetNet MPLS Network figure in [I-D.ietf-detnet-dp-sol-mpls]. The primary difference is that the Relay nodes are at the edges of the MPLS domain and therefore function as T-PEs, and that service sub-layer functions are not provided over the DetNet IP network. The transit node functions show above are identical to those described in [I-D.ietf-detnet-dp-sol-mpls].

Figure 9 illustrates how relay nodes can provide service protection over an MPLS domain. In this case, CE1 and CE2 are IP DetNet end systems which are interconnected via a MPLS domain such as described in [I-D.ietf-detnet-dp-sol-mpls]. Note that R1 and R3 sit at the

edges of an MPLS domain and therefore are similar to T-PEs, while R2 sits in the middle of the domain and is therefore similar to an S-PE.

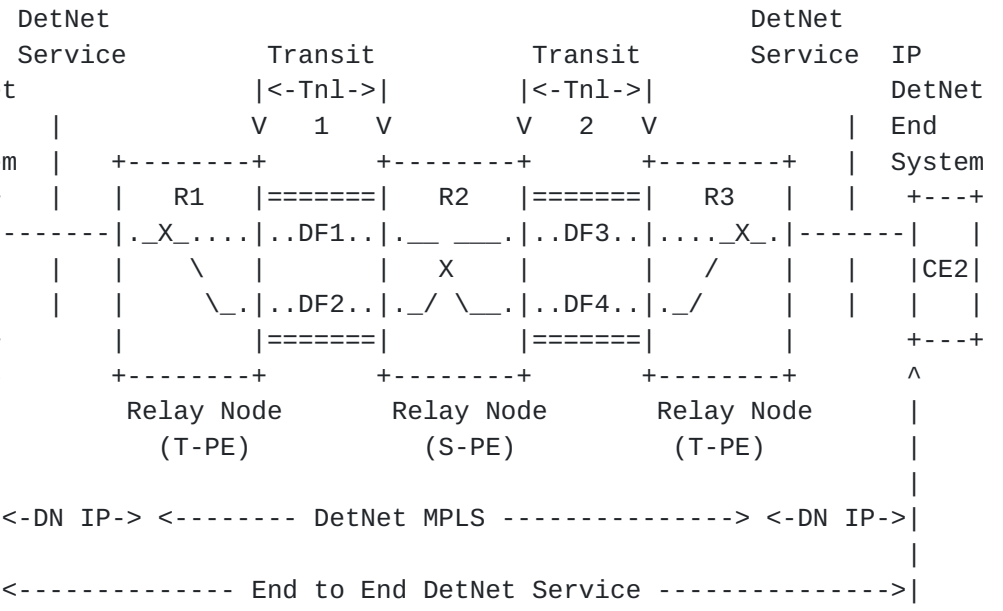

Figure 9: DetNet IP Over DetNet MPLS Network

[Editor's note: Text below in this sub-section is rather DetNet MPLS related, therefore candidate to be deleted in future versions.]

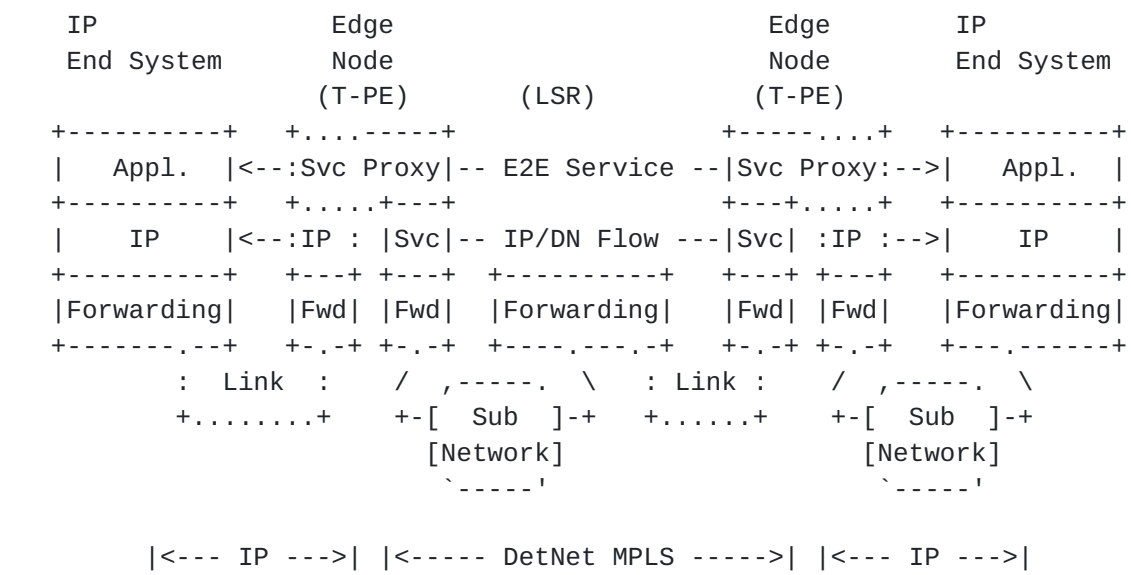

Figure 10: Non-DetNet Aware IP Over DetNet MPLS Network

Figure 10 illustrates non-DetNet enabled End Systems (hosts), connected to DetNet (DN) enabled MPLS network. It differs from Figure 8 in that the hosts and edge IP networks are not DetNet aware. In this case, edge nodes sit at the boundary of the MPLS domain since it is also a DetNet domain boundary. The edge nodes provide DetNet service proxies for the end applications by initiating and terminating DetNet service for the application's IP flows. While the node types differ, there is essentially no difference in data plane processing between relay and edges. There are likely to be differences in controller plane operation, particularly when distributed control plane protocols are used.

Figure 11 illustrates how it is still possible to provided DetNet service protection for non-DetNet aware end systems. This figures is basically the same as Figure 9, with the exception that CE1 and CE2 are non-DetNet aware end systems and E1 and E3 are edge nodes that replace the relay nodes R1 and R3.

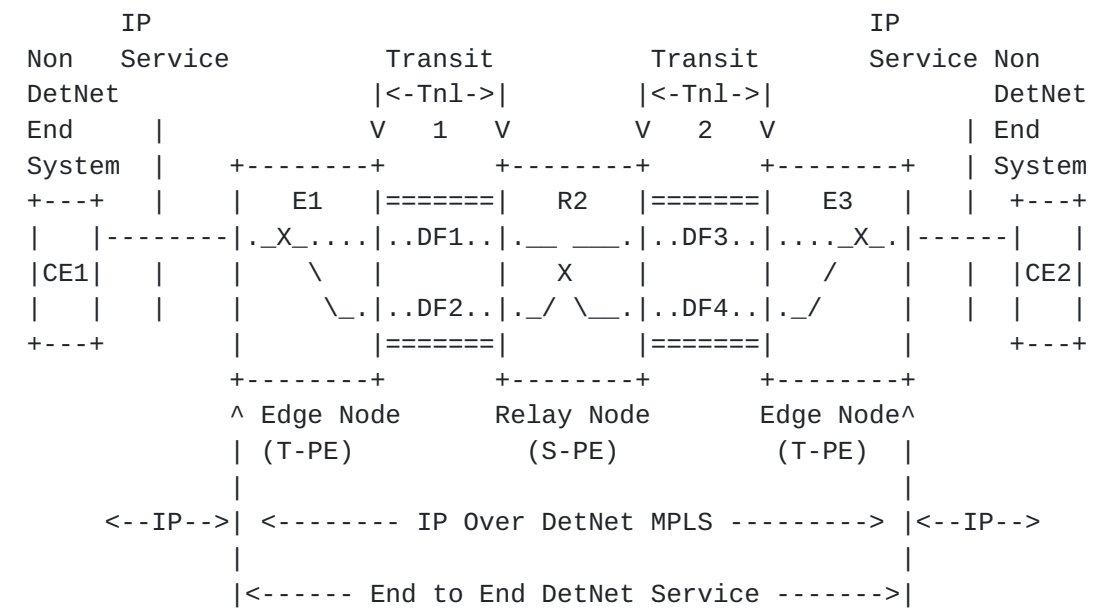

X = Optional service protection (none, PRF, PREOF, PEF/POF)
DFx = DetNet member flow x over a TE LSP

Figure 11: MPLS-Based DetNet (non-MPLS End System)

[Editor's note: End of text being rather DetNet MPLS related.]

7.2. DetNet IP over DetNet MPLS Encapsulation

The basic encapsulation approach is to treat a DetNet IP flow as an app-flow from the DetNet MPLS app perspective. The corresponding example DetNet Sub-Network format is shown in Figure 12.

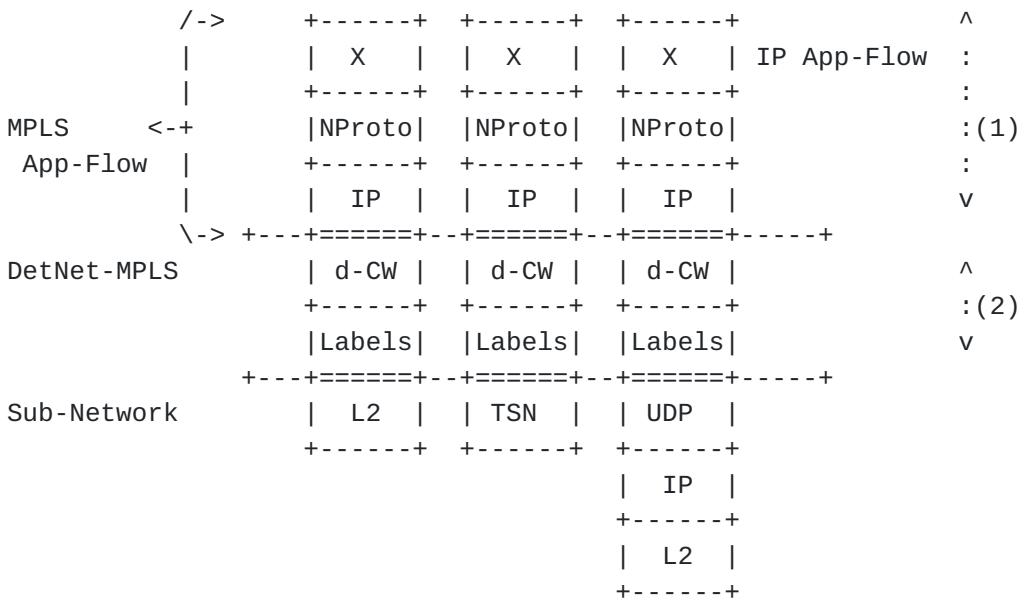

(1) DetNet IP Flow
(2) DetNet MPLS Flow

Figure 12: Example DetNet IP over MPLS Sub-Network Formats

In the figure, "IP App-Flow" indicates the payload carried by the DetNet IP data plane. "IP" and "NProto" indicate the fields described in Section 6.1.1 and Section 6.1.2, respectively. "MPLS App-Flow" indicates that an individual DetNet IP flow is the payload from the perspective of the DetNet MPLS data plane defined in [I-D.ietf-detnet-dp-sol-mpls].

Per [I-D.ietf-detnet-dp-sol-mpls], the DetNet MPLS data plane uses a single S-Label to support a single app flow. Section 6.1 states that a single DetNet flow is identified based on IP, and next level protocol, header information. It also defines that aggregation is supported (Section 6.4) through the use of prefixes, wildcards, bimasks, and port ranges. Collectively, this results in the fairly straight forward procedures defined in this section.

As shown in Figure 2, DetNet relay nodes are responsible for the mapping of a DetNet flow, at the service sub-layer, from the IP to MPLS DetNet data planes and back again. Their related DetNet IP over DetNet MPLS data plane operation is comprised of two sets of procedures: the mapping of flow identifiers; and ensuring proper traffic treatment.

7.3. DetNet IP over DetNet MPLS Flow Identification Procedures

A relay node that sends a DetNet IP flows over a DetNet MPLS network MUST map a single DetNet IP flow into a single app-flow and MUST process that app-flow in accordance to the procedures defined in [I-D.ietf-detnet-dp-sol-mpls] Section 6.2. PRF MAY be supported for DetNet IP flows sent over an DetNet MPLS network. Aggregation as defined in Section 6.4 MAY be used to identify an individual DetNet IP flow. The provisioning of the mapping of DetNet IP flows to DetNet MPLS app-flow information MUST be supported via configuration, e.g., via the controller plane described in Section 5.

A relay node MAY be provisioned to handle packets received via the DetNet MPLS data plane as DetNet IP flows. A single incoming MPLS app-flow MAY be treated as a single DetNet IP flow, without examination of IP headers. Alternatively, packets received via the DetNet MPLS data plane MAY follow the normal DetNet IP flow identification procedures defined in Section 6.1. An implementation MUST support the provisioning of handling of received DetNet MPLS data plane as DetNet IP flows via configuration. Note that such configuration MAY include support from PEOF on the incoming DetNet MPLS flow.

7.4. DetNet IP over DetNet MPLS Traffic Treatment Procedures

The traffic treatment required for a particular DetNet IP flow is provisioned via configuration or the controller plane. When an DetNet IP flow is sent over DetNet MPLS a relay node MUST ensure that the provisioned DetNet IP traffic treatment is provided at the forwarding sub-layer as described in [I-D.ietf-detnet-dp-sol-mpls] Section 6.2. Note that the PRF function can also be used when sending over MPLS.

Traffic treatment for DetNet IP flows received over the DetNet MPLS data plane MUST follow Section 6.3.

8. Mapping DetNet IP Flows to IEEE 802.1 TSN

[Authors note: how do we handle control protocols such as ICMP, IPsec, etc.]

This section covers how DetNet IP flows operate over an IEEE 802.1 TSN sub-network. Figure 13 illustrates such a scenario, where two IP (DetNet) nodes are interconnected by a TSN sub-network. Node-1 is single homed and Node-2 is dual-homed. IP nodes can be (1) DetNet IP End System, (2) DetNet IP Edge or Relay node or (3) IP End System.

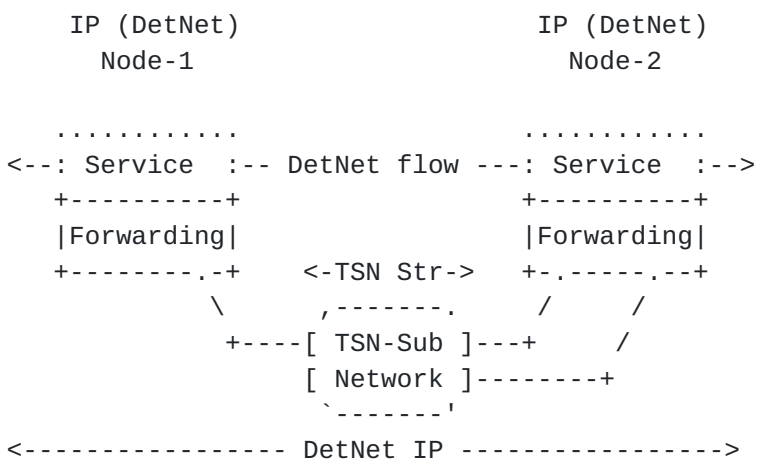

Figure 13: DetNet (DN) Enabled IP Network over a TSN sub-network

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

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

TSN capabilities of the TSN sub-network are made available for IP (DetNet) flows via the protocol interworking function defined in IEEE 802.1CB [IEEE8021CB]. For example, applied on the TSN edge port connected to the IP (DetNet) node it can convert an ingress unicast IP (DetNet) flow to use a specific multicast destination MAC address and VLAN, in order to direct the packet through a specific path inside the bridged network. A similar interworking pair at the other end of the TSN sub-network would restore the packet to its original destination MAC address and VLAN.

Placement of TSN functions depends on the TSN capabilities of nodes. IP (DetNet) Nodes may or may not support TSN functions. For a given TSN Stream (i.e., DetNet flow) an IP (DetNet) node is treated as a Talker or a Listener inside the TSN sub-network.

8.1. TSN Stream ID Mapping

DetNet IP Flow and TSN Stream mapping is based on the active Stream Identification function, that operates at the frame level. IEEE 802.1CB [IEEE8021CB] defines an Active Destination MAC and VLAN Stream identification function, what can replace some Ethernet header fields namely (1) the destination MAC-address, (2) the VLAN-ID and (3) priority parameters with alternate values. Replacement is provided for the frame passed down the stack from the upper layers or up the stack from the lower layers.

Active Destination MAC and VLAN Stream identification can be used within a Talker to set flow identity or a Listener to recover the original addressing information. It can be used also in a TSN bridge that is providing translation as a proxy service for an End System. As a result IP (DetNet) flows can be mapped to use a particular {MAC-address, VLAN} pair to match the Stream in the TSN sub-network.

[Editor's note: there are no requirement on IP DetNet nodes in case of "IP (DetNet) node without TSN functions" scenarios. Paragraph and figure beow are candidates to be deleted in future versions.]

From the TSN sub-network perspective DetNet IP nodes without any TSN functions can be treated as TSN-unaware Talker or Listener. In such cases relay nodes in the TSN sub-network MUST modify the Ethernet encapsulation of the DetNet IP flow (e.g., MAC translation, VLAN-ID setting, Sequence number addition, etc.) to allow proper TSN specific handling of the flow inside the sub-network. This is illustrated in Figure 14.

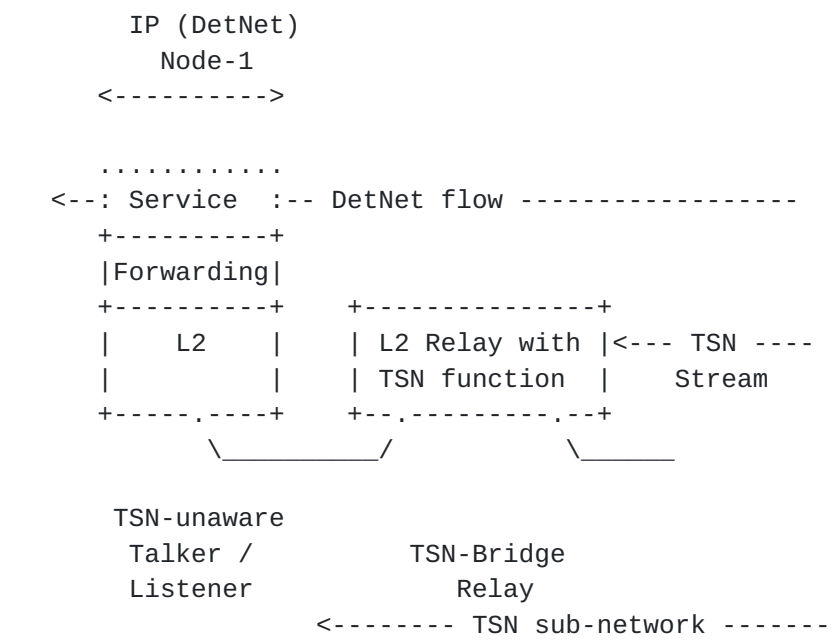

Figure 14: IP (DetNet) node without TSN functions

IP (DetNet) nodes being TSN-aware can be treated as a combination of a TSN-unaware Talker/Listener and a TSN-Relay, as shown in Figure 15. In such cases the IP (DetNet) node MUST provide the TSN sub-network specific Ethernet encapsulation over the link(s) towards the sub-network. An TSN-aware IP (DetNet) node MUST support the following TSN components:

1. For recognizing flows:

    * Stream Identification

2. For FRER used inside the TSN domain, additionally:

    * Sequencing function

    * Sequence encode/decode function

3. For FRER when the node is a replication or elimination point, additionally:

    * Stream splitting function

    * Individual recovery function

[Editor's note: Should we added here requirements regarding IEEE 802.1Q C-VLAN component?]

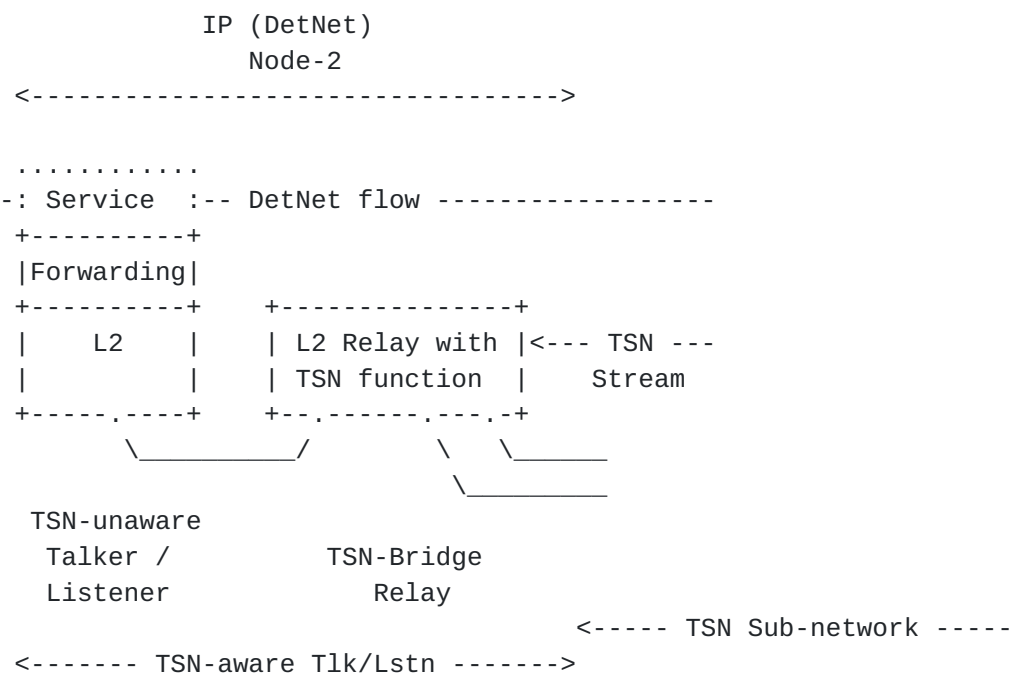

Figure 15: IP (DetNet) node with TSN functions

A Stream identification component MUST be able to instantiate the following functions (1) Active Destination MAC and VLAN Stream identification function, (2) IP Stream identification function and (3) the related managed objects in Clause 9 of IEEE 802.1CB [IEEE8021CB]. IP Stream identification function provides a 6-tuple match.

The Sequence encode/decode function MUST support the Redundancy tag (R-TAG) format as per Clause 7.8 of IEEE 802.1CB [IEEE8021CB].

8.2. TSN Usage of FRER

TSN Streams supporting DetNet flows may use Frame Replication and Elimination for Redundancy (FRER) [802.1CB] based on the loss service requirements of the TSN Stream, which is derived from the DetNet service requirements of the DetNet mapped flow. The specific operation of FRER is not modified by the use of DetNet and follows IEEE 802.1CB [IEEE8021CB].

FRER function and the provided service recovery is available only within the TSN sub-network (as shown in Figure 6) as the Stream-ID and the TSN sequence number are not valid outside the sub-network. An IP (DetNet) node represents a L3 border and as such it terminates all related information elements encoded in the L2 frames.

8.3. Procedures

[Editor's note: This section is TBD - covers required behavior of a TSN-aware DetNet node using a TSN underlay.]

This section provides DetNet IP data plane procedures to interwork with a TSN underlay sub-network when the IP (DetNet) node acts as a TSN-aware Talker or Listener (see Figure 15). These procedures have been divided into the following areas: flow identification, mapping of a DetNet flow to a TSN Stream and ensure proper TSN encapsulation.

Flow identification procedures are described in Section 6.1. A TSN-aware IP (DetNet) node SHALL support the Stream Identification TSN components as per IEEE 802.1CB [IEEE8021CB].

Implementations of this document SHALL use management and control information to map a DetNet flow to a TSN Stream. N:1 mapping (aggregating DetNet flows in a single TSN Stream) SHALL be supported. The management or control function that provisions flow mapping SHALL ensure that adequate resources are allocated and configured to provide proper service requirements of the mapped flows.

For proper TSN encapsulation implementations of this document SHALL support active Stream Identification function as defined in chapter 6.6 in IEEE 802.1CB [IEEE8021CB].

A TSN-aware IP (DetNet) node SHALL support Ethernet encapsulation with Redundancy tag (R-TAG) as per chapter 7.8 in IEEE 802.1CB [IEEE8021CB].

Depending whether FRER functions are used in the TSN sub-network to serve the mapped TSN Stream, a TSN-aware IP (DetNet) node SHALL support Sequencing function and Sequence encode/decode function as per chapter 7.4 and 7.6 in IEEE 802.1CB [IEEE8021CB]. Furthermore when a TSN-aware IP (DetNet) node acting as a replication or elimination point for FRER it SHALL implement the Stream splitting function and the Individual recovery function as per chapter 7.7 and 7.5 in IEEE 802.1CB [IEEE8021CB].

8.4. Management and Control Implications

[Editor's note: This section is TBD Covers Creation, mapping, removal of TSN Stream IDs, related parameters and,when needed, configuration of FRER. Supported by management/control plane.]

DetNet flow and TSN Stream mapping related information are required only for TSN-aware IP (DetNet) nodes. From the Data Plane

perspective there is no practical difference based on the origin of flow mapping related information (management plane or control plane).

TSN-aware DetNet IP nodes are member of both the DetNet domain and the TSN sub-network. Within the TSN sub-network the TSN-aware IP (DetNet) node has a TSM-aware Talker/Listener role, so TSN specific management and control plane functionalities must be implemented. There are many similarities in the management plane techniques used in DetNet and TSN, but that is not the case for the control plane protocols. For example, RSVP-TE and MSRP behaves differently. Therefore management and control plane design is an important aspect of scenarios, where mapping between DetNet and TSN is required.

In order to use a TSN sub-network between DetNet nodes, DetNet specific information must be converted to TSN sub-network specific ones. DetNet flow ID and flow related parameters/requirements must be converted to a TSN Stream ID and stream related parameters/ requirements. Note that, as the TSN sub-network is just a portion of the end2end DetNet path (i.e., single hop from IP perspective), some parameters (e.g., delay) may differ significantly. Other parameters (like bandwidth) also may have to be tuned due to the L2 encapsulation used in the TSN sub-network.

In some case it may be challenging to determine some TSN Stream related information. For example on a TSN-aware IP (DetNet) node that acts as a Talker, it is quite obvious which DetNet node is the Listener of the mapped TSN stream (i.e., the IP Next-Hop). However it may be not trivial to locate the point/interface where that Listener is connected to the TSN sub-network. Such attributes may require interaction between control and management plane functions and between DetNet and TSN domains.

Mapping between DetNet flow identifiers and TSN Stream identifiers, if not provided explicitly, can be done by a TSN-aware IP (DetNet) node locally based on information provided for configuration of the TSN Stream identification functions (IP Stream identification and active Stream identification function).

Triggering the setup/modification of a TSN Stream in the TSN sub-network is an example where management and/or control plane interactions are required between the DetNet and TSN sub-network. TSN-unaware IP (DetNet) nodes make such a triggering even more complicated as they are fully unaware of the sub-network and run independently.

Configuration of TSN specific functions (e.g., FRER) inside the TSN sub-network is a TSN specific decision and may not be visible in the DetNet domain.

9. Security Considerations

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

10. IANA Considerations

TBD.

11. Contributors

RFC7322 limits the number of authors listed on the front page of a draft to a maximum of 5, far fewer than the 20 individuals below who made important contributions to this draft. The editor wishes to thank and acknowledge each of the following authors for contributing text to this draft. See also Section 12.

Loa Andersson
Huawei
Email: loa@pi.nu

Yuanlong Jiang
Huawei
Email: jiangyuanlong@huawei.com

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

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

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

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

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

Andrew G. Malis
Huawei Technologies
Email: agmalis@gmail.com

12. Acknowledgements

The author(s) ACK and NACK.

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

Jouni Korhonen

Janos Farkas

Norman Finn

Balazs Varga

Loa Andersson

Tal Mizrahi

David Mozes

Yuanlong Jiang

Andrew Malis

Carlos J. Bernardos

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

Lou Berger

Pat Thaler

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

13. References

13.1. Normative references

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <https://www.rfc-editor.org/info/rfc768>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/info/rfc791>.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <https://www.rfc-editor.org/info/rfc793>.

[RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <https://www.rfc-editor.org/info/rfc1812>.

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

[RFC2211] Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, DOI 10.17487/RFC2211, September 1997, <https://www.rfc-editor.org/info/rfc2211>.

[RFC2212] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, DOI 10.17487/RFC2212, September 1997, <https://www.rfc-editor.org/info/rfc2212>.

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

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <https://www.rfc-editor.org/info/rfc3168>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, DOI 10.17487/RFC3270, May 2002, <https://www.rfc-editor.org/info/rfc3270>.

[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>.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <https://www.rfc-editor.org/info/rfc4302>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/info/rfc4303>.

[RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, DOI 10.17487/RFC5462, February 2009, <https://www.rfc-editor.org/info/rfc5462>.

[RFC6003] Papadimitriou, D., "Ethernet Traffic Parameters", RFC 6003, DOI 10.17487/RFC6003, October 2010, <https://www.rfc-editor.org/info/rfc6003>.

[RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <https://www.rfc-editor.org/info/rfc7608>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

13.2. Informative references

[G.8275.1] International Telecommunication Union, "Precision time protocol telecom profile for phase/time synchronization with full timing support from the network", ITU-T G.8275.1/Y.1369.1 G.8275.1, June 2016, <https://www.itu.int/rec/T-REC-G.8275.1/en>.

[G.8275.2]
International Telecommunication Union, "Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network", ITU-T G.8275.2/Y.1369.2 G.8275.2, June 2016, <https://www.itu.int/rec/T-REC-G.8275.2/en>.

[I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", draft-ietf-detnet-architecture-11 (work in progress), February 2019.

[I-D.ietf-detnet-dp-sol-mpls]
Korhonen, J. and B. Varga, "DetNet MPLS Data Plane Encapsulation", draft-ietf-detnet-dp-sol-mpls-01 (work in progress), October 2018.

[I-D.ietf-detnet-flow-information-model]
Farkas, J., Varga, B., Cummings, R., and Y. Jiang, "DetNet Flow Information Model", draft-ietf-detnet-flow-information-model-03 (work in progress), March 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-04 (work in progress), March 2019.

[I-D.ietf-teas-pce-native-ip]
Wang, A., Zhao, Q., Khasanov, B., Chen, H., and R. Mallya, "PCE in Native IP Network", draft-ietf-teas-pce-native-ip-02 (work in progress), October 2018.

[IEEE1588]
IEEE, "IEEE 1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems Version 2", 2008.

[IEEE8021CB]
Finn, N., "Draft Standard for Local and metropolitan area networks - Seamless Redundancy", IEEE P802.1CB /D2.1 P802.1CB, December 2015, <http://www.ieee802.org/1/files/private/cb-drafts/d2/802-1CB-d2-1.pdf>.

[IEEE8021Q] IEEE 802.1, "Standard for Local and metropolitan area networks--Bridges and Bridged Networks (IEEE Std 802.1Q-2014)", 2014, <http://standards.ieee.org/about/get/>.

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/info/rfc1122>.

[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>.

[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>.

[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>.

[RFC5777] Korhonen, J., Tschofenig, H., Arumaithurai, M., Jones, M., Ed., and A. Lior, "Traffic Classification and Quality of Service (QoS) Attributes for Diameter", RFC 5777, DOI 10.17487/RFC5777, February 2010, <https://www.rfc-editor.org/info/rfc5777>.

[RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <https://www.rfc-editor.org/info/rfc6434>.

[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>.

[RFC8169] Mirsky, G., Ruffini, S., Gray, E., Drake, J., Bryant, S., and A. Vainshtein, "Residence Time Measurement in MPLS Networks", RFC 8169, DOI 10.17487/RFC8169, May 2017, <https://www.rfc-editor.org/info/rfc8169>.

[RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", RFC 8283, DOI 10.17487/RFC8283, December 2017, <https://www.rfc-editor.org/info/rfc8283>.

Appendix A. Example of DetNet Data Plane Operation

[Editor's note: Add a simplified example of DetNet data plane and how labels etc work in the case of MPLS-based PSN and utilizing PREOF. The figure is subject to change depending on the further DT decisions on the label handling..]

Appendix B. Example of Pinned Paths Using IPv6

TBD.

Authors' Addresses

Jouni Korhonen (editor)

Email: jouni.nospam@gmail.com

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

Email: balazs.a.varga@ericsson.com