DetNet   B. Varga, Ed.
Internet-Draft   J. Farkas
Intended status: Standards Track   Ericsson
Expires: May 24, 2020   L. Berger
                                                              D. Fedyk
                                                  LabN Consulting, L.L.C.
                                                              A. Malis
                                                          Independent
                                                             S. Bryant
                                                 Futurewei Technologies
                                                           J. Korhonen
                                                    November 21, 2019

DetNet Data Plane: IP
draft-ietf-detnet-ip-04

Abstract

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

Status of This Memo

Copyright Notice

Table of Contents

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 [RFC8655].

   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, instead
the existing IP and higher layer protocol header information is used
to support flow identification and DetNet service delivery. Common
data plane procedures and control information for all DetNet data
planes can be found in the [I-D.ietf-detnet-data-plane-framework].

The DetNet Architecture models the DetNet related data plane
functions as two sub-layers: 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 (e.g., by packet replication and
packet elimination functions) and reordering. The forwarding sub-
layer is used to provides congestion protection (low loss, assured
latency, and limited out-of-order delivery). The service sub-layer
generally requires additional fields to provide its service; for
example see [I-D.ietf-detnet-mpls]. Since no DetNet-specific fields
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 Ethernet
as specified in the IEEE 802.1 TSN task group(referred to in this
document simply as 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. Section 5 provides the
procedures for hosts and routers that support IP-based DetNet
services. Section 6 summarizes the set of information that is needed
to identify an individual DetNet flow.

2. Terminology

2.1. Terms Used In This Document

This document uses the terminology and concepts established in the
DetNet architecture [RFC8655], and the reader is assumed to be
familiar with that document and its terminology.

2.2. Abbreviations

The following abbreviations used in this document:

CoS   Class of Service.

DetNet   Deterministic Networking.

DN   DetNet.

DiffServ   Differentiated Services

DSCP   Differentiated Services Code Point

L2   Layer-2.

L3   Layer-3.

LSP   Label-switched path.

MPLS   Multiprotocol Label Switching.

PREOF   Packet Replication, Elimination and Ordering Function.

QoS   Quality of Service.

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
using an IP data plane. 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. Common data plane
procedures and control information for all DetNet data planes can be
found in the [I-D.ietf-detnet-data-plane-framework].

The DetNet IP data plane primarily uses "6-tuple" based flow
identification, where 6-tuple refers to information carried in IP and
higher layer protocol headers. The 6-tuple referred to in this
document is the same as that defined in [RFC3290]. Specifically
6-tuple is (destination address, source address, IP protocol, source
port, destination port, and differentiated services (DiffServ) code
point (DSCP). 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 of DiffServ and "tuple" based flow identification.

The DetNet IP data plane also allows for optional matching on the
IPv6 flow label field, as defined in [RFC8200].

Non-DetNet and DetNet IP packets are identical on the wire.
Generally the fields used in flow identification are forwarded
unmodified, however modification of these fields is allowed, for
example to a DSCP value, when required by the DetNet service.

DetNet flow aggregation may be enabled via the use of wildcards,
masks, lists, 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 on the
aggregate through resource allocation and congestion control
mechanisms.

```
 DetNet IP       Relay                      Relay       DetNet IP
 End System      Node                       Node        End System


+----------+                                            +----------+
|   Appl.  |<------------ End to End Service ---------->|   Appl.  |
+----------+  ...........                  ...........  +----------+
| Service  |<-: Service  :-- DetNet flow --: Service  :->| Service  |
+----------+  +----------+                  +----------+  +----------+
|Forwarding|  |Forwarding|                  |Forwarding|  |Forwarding|
+--------.-+  +-.-------.-+                  +-.---.----+  +-------.--+
        : Link :        \        ,-----.      /     \   ,-----.   /
        +......+         +----[  Sub  ]----+     +-[  Sub  ]-+
                              [Network]            [Network]
                               `-----'              `-----'


         |<-------------------- DetNet IP -------------------->|
```
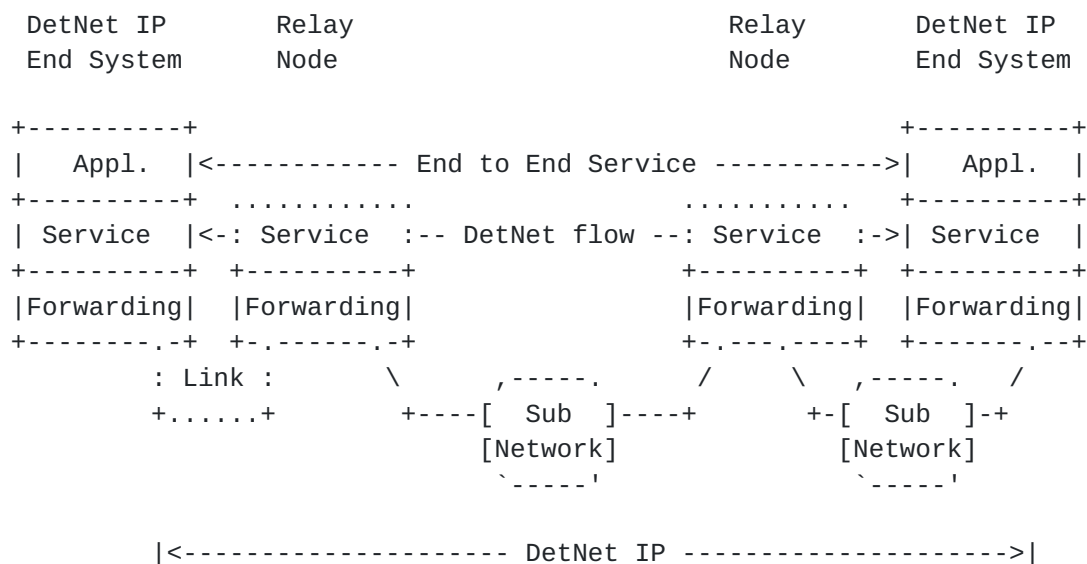
               Figure 1: A Simple DetNet (DN) Enabled IP Network

Figure 1 illustrates a DetNet enabled IP network. The DetNet enabled
end systems originate IP encapsulated traffic those are identified
within the DetNet domain 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.

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

```
  IP              Edge                        Edge         IP
  End System      Node                        Node         End System

  +----------+   +.........+              +.........+   +----------+
  |   Appl.  |<--:Svc Proxy:-- E2E Service---:Svc Proxy:-->|   Appl.  |
  +----------+   +.........+              +.........+   +----------+
  |   IP     |<--:IP : :Svc:---- IP flow ----:Svc: :IP :-->|   IP     |
  +----------+   +---+ +---+              +---+ +---+   +----------+
  |Forwarding|   |Fwd| |Fwd|              |Fwd| |Fwd|   |Forwarding|
  +---------.-+   +-.-+ +-.-+              +-.-+ +-.-+   +---.------+
         :  Link :      \      ,-----.      /     / ,-----.  \
         +.......+      +----[  Sub  ]----+     +--[  Sub  ]--+
                             [Network]            [Network]
                             `-----'              `-----'


     |<--- IP --->| |<------ DetNet IP ------>| |<--- IP --->|
```
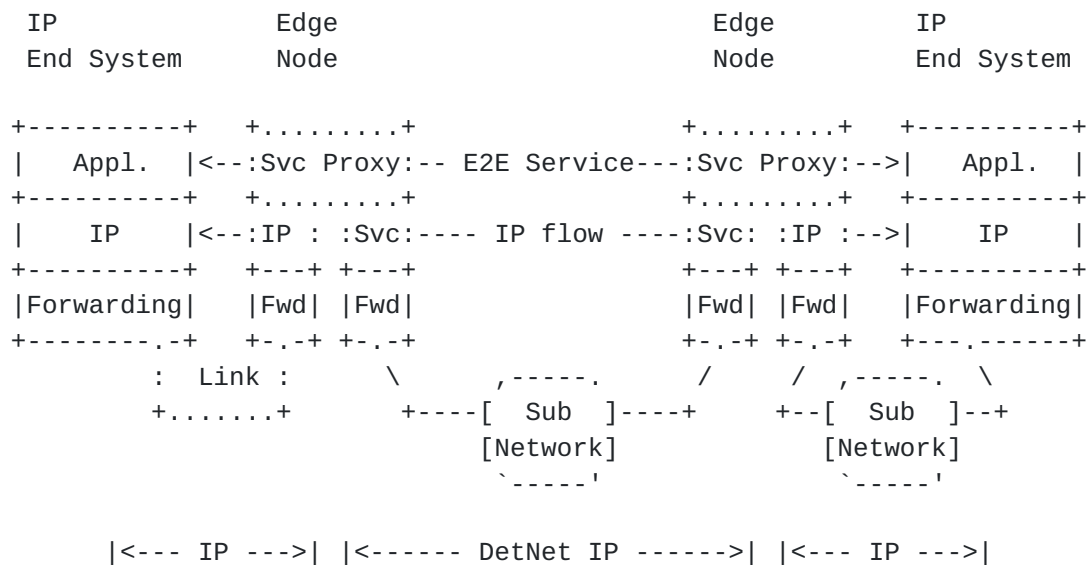
      Figure 2: Non-DetNet-aware IP end systems with DetNet IP Domain

   Figure 2 illustrates a 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 4.4 can be used to support
   DetNet flow identification.

   Note, that Figure 1 and Figure 2 can be collapsed, so IP DetNet End
   Systems can communicate over DetNet IP network with IP End System.

   As 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. TSN over MPLS is
   discribed in [I-D.ietf-detnet-tsn-vpn-over-mpls].

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.

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 other end systems. 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 (see Section 5.1.1.3).

   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 to a DetNet flow. When forwarding
   packets, this means that packets are appropriately shaped on
   transmission and receive appropriate traffic treatment on the
   connected sub-network, see Section 4.3.2 and Section 4.2 for more
   details. When receiving packets, this means that there are
   appropriate local node resources, e.g., buffers, to receive and
   process the packets of that DetNet flow.

   In order to maximize reuse of existing mechanisms, DetNet-aware
   applications and end systems SHOULD NOT mix DetNet and non-DetNet
   traffic within a single 5-tuple.

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 the number of 6-tuple flow ID combinations that could be used
   by the end systems. From a practical standpoint this means that all
   nodes along the end-to-end path of 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: the end system is directly connected to an edge
       node, or the end system is behind a sub-network (See ES1 and ES2
       in figure below)

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

L3 (IP) end systems may use any of these connection types. A 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 3 for L3 end
system connection examples.

```
                                          ____+----+
                   +----+          _____    /   | ES3|
                   | ES1|____    /      \__/     +----+___
                   +----+    \ /                          \
                             +                             |
                          ____       \                      _/
            +----+      __/    \      +__   DetNet IP domain  /
            | ES2|___ /   L2/L3 |__/    \             __       __/
            +----+     _____ /          _____/   \__/
```
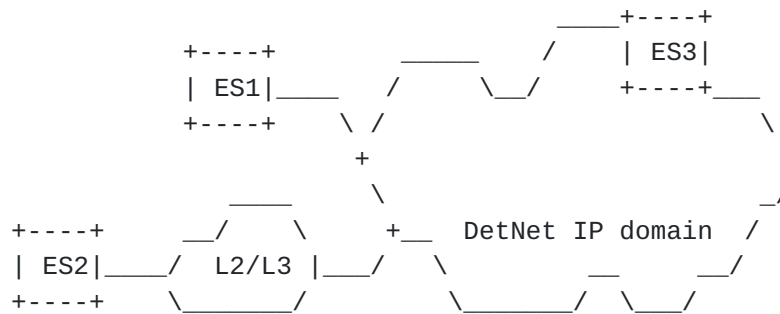
                  Figure 3: Connection types of L3 end systems

Within a DetNet domain, the DetNet-enabled IP Routers are
interconnected by 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 protection (packet replication and packet
elimination functions) is not provided at the DetNet layer end to
end. Instead service protection can be provided on a per underlying
L2 link and sub-network basis.

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 6-tuple to find out the
required mapping within a node.

As noted earlier, service protection must be implemented within each
link / sub-network independently, using the domain specific
mechanisms. This is due to the lack of unified end-to-end sequencing
information that could be used by the intermediate nodes. Therefore,

service protection (if enabled) cannot be provided end-to-end, only
within sub-networks. This is shown for a three sub-network scenario
in Figure 4, where each sub-network can provide service protection
between its borders. "R" and "E" denotes replication and elimination
points within the sub-network.

```
        <------------------- DenNet IP ----------------------->
                                    _____
                         ____      /      \__
                 ____        /     \__/          \___    _____
+----+    __/     +====+                     +==+      \       +----+
|src |__/ SubN1   )   |                      |  \ SubN3 \____| dst|
+----+  _____/   \           Sub-Network2   |   \_____/     +----+
                 \_                            _/
                   \               __    __/
                    _____/   \___/
```

```
         +---+          +---------E--------+      +-----+
+----+   |   |          |         |        |      |     |       +----+
|src |----R   E--------R      +---+         E------R     E------+ dst|
+----+   |   |          |     |             |      |     |       +----+
         +---+          +-----R------------+      +-----+
```
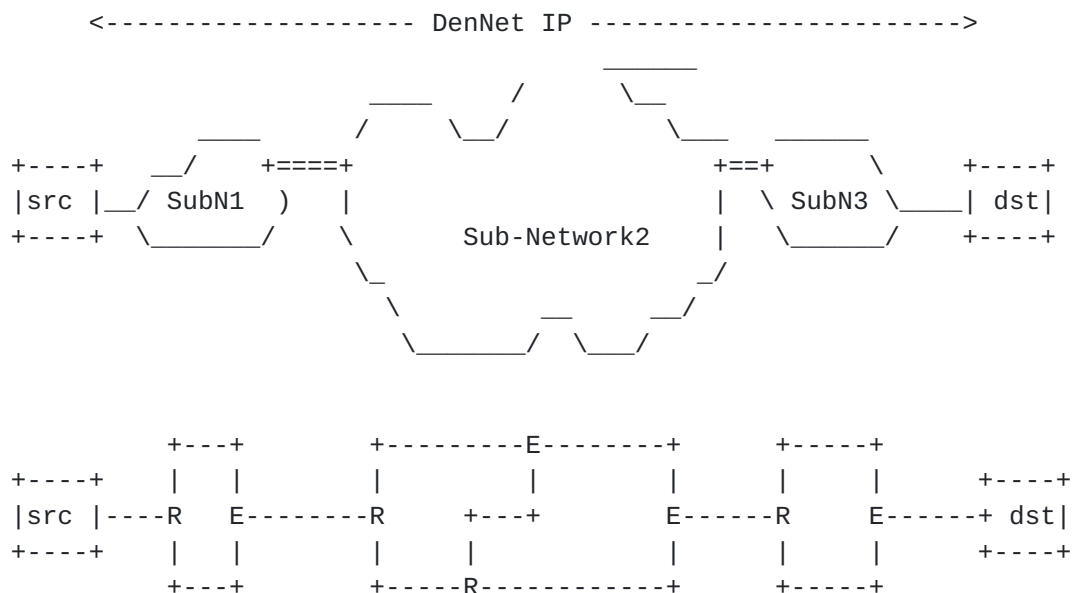
            Figure 4: Replication and elimination in sub-networks for DetNet IP
                                networks

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

Note that not mixing DetNet and non-DetNet traffic within a single
5-tuple, as described above, enables simpler 5-tuple filters to be
used (or re-used) at the edges of a DetNet network to prevent non-
congestion-responsive DetNet traffic from escaping the DetNet domain.

4.3. Forwarding Sub-Layer Considerations

4.3.1. Class of Service

Class of Service (CoS) for DetNet flows carried in IPv6 is provided
using the standard differentiated services code point (DSCP) field
[RFC2474] and related mechanisms.

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 the
requirement for MPLS LSRs that CoS LSPs cannot impact the resources
allocated to TE LSPs [RFC3473].

4.3.2. **Quality of Service**

Quality of Service (QoS) for DetNet service flows carried in IP MUST
be provided locally by the DetNet-aware hosts and routers supporting
DetNet flows. Such support leverages the underlying network layer
such as 802.1 TSN. 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 and
previously mentioned optional field, uniquely identifies a DetNet IP
flow.

Packets that are identified as part of a DetNet IP flow 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 ensure that no DetNet allocated resources, e.g.,
queue or shaper, is used by such flows. There are multiple methods
that MAY be used by an implementation to defend service delivery to
reserved DetNet flows, including but not limited to:

o Treating packets associated with an incomplete reservation as non-
   DetNet traffic.

o Discarding packets associated with an incomplete reservation.

o Remarking packets associated with an incomplete reservation.
   Remarking can be accomplished by changing the value of the DSCP,
   or optional, field to a value that results in the packet no longer
   matching any other reserved DetNet IP flow.

4.3.3. **Path Selection**

While path selection algorithms and mechanisms are out of scope of
the DetNet data plane definition, it is important to highlight the
implications of DetNet IP flow identification on path selection and
next hops. As mentioned above, the DetNet IP data plane identifies
flows using "6-tuple" header information as well as the additional
optional header field. DetNet generally allows for both flow-
specific traffic treatment and flow-specific next-hops.

In non-DetNet IP forwarding, it is generally assumed that the same
series of next hops, i.e., the same path, will be used for a
particular 5-tuple or, in some cases, e.g., [RFC5120], for a
particular 6-tuple. Using different next hops for different 5-tuples
does not take any special consideration for DetNet-aware
applications.

Care should be taken when using different next hops for the same
5-tuple. As discussed in [RFC7657], unexpected behavior can occur
when a single 5-tuple application flow experience reordering due to
being split across multiple next hops. Understanding of the
application and transport protocol impact of using different next
hops for the same 6-tuple is required. Again, this impacts path
selection for DetNet flows and this document only indirectly.

4.4. DetNet Flow Aggregation

As described in [I-D.ietf-detnet-data-plane-framework], the ability
to aggregate individual flows, and their associated resource control,
into a larger aggregate is an important technique for improving
scaling by reducing the state per hop. DetNet IP data plane
aggregation can take place within a single node, when that node
maintains state about both the aggregated and individual 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 individual flows. As
DetNet is concerned about latency and jitter, more than just
bandwidth needs to be considered.

From a single node perspective, the aggregation of IP flows impacts
DetNet IP data plane flow identification and resource allocation. As
discussed above, IP flow identification uses the IP "6-tuple" for
flow identification. DetNet IP flows can be aggregated using any of
the 6-tuple, and an additional optional field defined in Section 5.1.
The use of prefixes, wildcards, lists, and value ranges allows a
DetNet node to identify aggregate DetNet flows. From a resource
allocation perspective, DetNet nodes must provide service to an
aggregate and not on a component flow basis.

It is the responsibility of the DetNet controller plane to properly
provision the use of these aggregation mechanisms. This includes
ensuring that aggregated flows have compatible e.g., the same or very
similar QoS and/or CoS characteristics, see Section 4.3.2. It also
includes ensuring that per component-flow service requirements are
satisfied by the aggregate, see Section 5.3.

4.5. Bidirectional Traffic

   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 or co-
   routed bidirectional flows can be managed at the control level.

   Control and management mechanisms 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].

5. 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 establishment and operational procedures also
   have requirements on the control and management systems for DetNet
   flows and these are referred in this section. Specifically this
   section identifies a number of information elements that 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 requirements for
   management and control related information is included. Conformance
   language is not used in the summary since applies to future
   mechanisms such as those that may be provided in YANG models
   [I-D.ietf-detnet-yang].

5.1. 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 the 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.

5.1.1. 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].

5.1.1.1. 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.

5.1.1.2. 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 an IP multicast
destination address.

5.1.1.3. 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 on the next
protocol values defined in Section 5.1.2. Other, non-zero values,
MUST be used for flow identification. Implementations SHOULD allow
for these fields to be ignored for a specific DetNet flow.

5.1.1.4. **IPv4 Type of Service and IPv6 Traffic Class Fields**

   These fields are used to support Differentiated Services [RFC2474]
   [RFC2475]. Implementations of this document MUST support DetNet flow
   identification based on the DSCP field in the IPv4 Type of Service
   field when processing IPv4 packets, and the DSCP field in the IPv6
   Traffic Class Field when processing IPv6 packets. Implementations
   MUST support list based matching of DSCP values, where the list is
   composed of possible field values that are to be considered when
   identifying a specific DetNet flow. Implementations SHOULD allow for
   this field to be ignored for a specific DetNet flow.

5.1.1.5. **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 field
   to be ignored for a specific DetNet flow. When this field 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.

5.1.2. **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 is defined. Future
   documents are expected to define support for other protocols.

5.1.2.1. **TCP and UDP**

   DetNet flow identification for TCP [RFC0793] and UDP [RFC0768] is
   achieved 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.

5.1.2.1.1. **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 value.
   Implementations SHOULD support range-based port matching.
   Implementation MUST also allow for the field to be ignored for a
   specific DetNet flow.

5.1.2.1.2. **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 value.
   Implementations SHOULD support range-based port matching.
   Implementation MUST also allow for the field to be ignored for a
   specific DetNet flow.

5.1.2.2. **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 value. Implementation SHOULD also allow for the field
   to be ignored for a specific DetNet flow.

5.2. **Forwarding Procedures**

   General requirements for IP nodes are defined in [RFC1122], [RFC1812]
   and [RFC8504], 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 associated with 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
   [I-D.ietf-detnet-yang].

5.3. **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
   [I-D.ietf-detnet-yang]. 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 [I-D.ietf-detnet-ip-over-mpls] or IEEE802.1

TSN [I-D.ietf-detnet-ip-over-tsn]. 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. Management and Control Information Summary

The following summarizes the set of information that is needed to
identify individual and aggregated DetNet flows:

o IPv4 and IPv6 source address field.

o IPv4 and IPv6 source address prefix length, where a zero (0) value
   effectively means that the address field is ignored.

o IPv4 and IPv6 destination address field.

o IPv4 and IPv6 destination address prefix length, where a zero (0)
   effectively means that the address field is ignored.

o 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.

o 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.

o For the IPv4 Type of Service and IPv6 Traffic Class Fields:

   * If the DSCP field is to be used in flow identification.
      Ignoring the DSCP filed is optional.

   * When the DSCP field is used in flow identification, a list of
      field values that may be used by a specific flow.

o IPv6 flow label field. This field can be optionally used for
   matching. When used, can be used instead of matching against the
   Next Header field.

o TCP and UDP Source Port. Exact and wildcard matching is required.
   Port ranges can optionally be used.

o TCP and UDP Destination Port. Exact and wildcard matching is
   required. Port ranges can optionally be used.

o IPsec Header SPI field. Exact matching is required.

This information MUST be provisioned per DetNet flow via
configuration, e.g., via the controller or management plane.

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 the aggregate of all other flows with that same UDP Destination
Port value.

It is the responsibility of the DetNet controller plane to properly
provision both flow identification information and the flow specific
resources needed to provided the traffic treatment needed to meet
each flow's service requirements. This applies for aggregated and
individual flows.

7. Security Considerations

Security considerations for DetNet are described in detail in
[I-D.ietf-detnet-security]. General security considerations are
described in [RFC8655]. This section considers exclusively security
considerations which are specific to the DetNet IP data plane.

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

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

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

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

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

8. IANA Considerations

This document does not require an action from IANA.

9. Acknowledgements

The authors wish to thank Pat Thaler, Norman Finn, Loa Anderson,
David Black, Rodney Cummings, Ethan Grossman, Tal Mizrahi, David
Mozes, Craig Gunther, George Swallow, Yuanlong Jiang and Carlos J.
Bernardos for their various contributions to this work. David Black
served as technical advisor to the DetNet working group during the
development of this document and provided many valuable comments.

10. References

10.1. 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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

10.2. Informative references

   [I-D.ietf-detnet-data-plane-framework]
              Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A.,
              Bryant, S., and J. Korhonen, "DetNet Data Plane
              Framework", draft-ietf-detnet-data-plane-framework-03
              (work in progress), October 2019.

   [I-D.ietf-detnet-dp-sol-mpls]
              Korhonen, J. and B. Varga, "DetNet MPLS Data Plane
              Encapsulation", draft-ietf-detnet-dp-sol-mpls-02 (work in
              progress), March 2019.

   [I-D.ietf-detnet-flow-information-model]
              Farkas, J., Varga, B., Cummings, R., Jiang, Y., and D.
              Fedyk, "DetNet Flow Information Model", draft-ietf-detnet-
              flow-information-model-06 (work in progress), October
              2019.

   [I-D.ietf-detnet-ip-over-mpls]
              Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A.,
              Bryant, S., and J. Korhonen, "DetNet Data Plane: IP over
              MPLS", draft-ietf-detnet-ip-over-mpls-03 (work in
              progress), October 2019.

   [I-D.ietf-detnet-ip-over-tsn]
              Varga, B., Farkas, J., Malis, A., and S. Bryant, "DetNet
              Data Plane: IP over IEEE 802.1 Time Sensitive Networking
              (TSN)", draft-ietf-detnet-ip-over-tsn-01 (work in
              progress), October 2019.

   [I-D.ietf-detnet-mpls]
              Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A.,
              Bryant, S., and J. Korhonen, "DetNet Data Plane: MPLS",
              draft-ietf-detnet-mpls-03 (work in progress), October
              2019.

   [I-D.ietf-detnet-security]
              Mizrahi, T., Grossman, E., Hacker, A., Das, S., Dowdell,
              J., Austad, H., and N. Finn, "Deterministic Networking
              (DetNet) Security Considerations", draft-ietf-detnet-
              security-06 (work in progress), November 2019.

   [I-D.ietf-detnet-tsn-vpn-over-mpls]
              Varga, B., Farkas, J., Malis, A., Bryant, S., and D.
              Fedyk, "DetNet Data Plane: IEEE 802.1 Time Sensitive
              Networking over MPLS", draft-ietf-detnet-tsn-vpn-over-
              mpls-01 (work in progress), October 2019.

   [I-D.ietf-detnet-yang]
              Geng, X., Chen, M., Ryoo, Y., Li, Z., and R. Rahman,
              "Deterministic Networking (DetNet) Configuration YANG
              Model", draft-ietf-detnet-yang-03 (work in progress), July
              2019.

   [IEEE802.1AE-2018]
              IEEE Standards Association, "IEEE Std 802.1AE-2018 MAC
              Security (MACsec)", 2018,
              <https://ieeexplore.ieee.org/document/8585421>.

   [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>.

   [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z.,
              and W. Weiss, "An Architecture for Differentiated
              Services", RFC 2475, DOI 10.17487/RFC2475, December 1998,
              <https://www.rfc-editor.org/info/rfc2475>.

   [RFC3290] Bernet, Y., Blake, S., Grossman, D., and A. Smith, "An
              Informal Management Model for Diffserv Routers", RFC 3290,
              DOI 10.17487/RFC3290, May 2002,
              <https://www.rfc-editor.org/info/rfc3290>.

   [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>.

   [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi
              Topology (MT) Routing in Intermediate System to
              Intermediate Systems (IS-ISs)", RFC 5120,
              DOI 10.17487/RFC5120, February 2008,
              <https://www.rfc-editor.org/info/rfc5120>.

   [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>.

   [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>.

   [RFC8504]   Chown, T., Loughney, J., and T. Winters, "IPv6 Node
               Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504,
               January 2019, <https://www.rfc-editor.org/info/rfc8504>.

   [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>.

Authors' Addresses

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

   Email: balazs.a.varga@ericsson.com

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

   Email: janos.farkas@ericsson.com

   Lou Berger
   LabN Consulting, L.L.C.

   Email: lberger@labn.net

   Don Fedyk
   LabN Consulting, L.L.C.

   Email: dfedyk@labn.net

   Andrew G. Malis
   Independent

   Email: agmalis@gmail.com

   Stewart Bryant
   Futurewei Technologies

   Email: stewart.bryant@gmail.com

   Jouni Korhonen

   Email: jouni.nospam@gmail.com