Stream: Internet Engineering Task Force (IETF)
RFC: 8939
Category: Standards Track
Published: November 2020
ISSN: 2070-1721
Authors:
B. Varga, Ed. *Ericsson*
J. Farkas *Ericsson*
L. Berger *LabN Consulting, L.L.C.*
D. Fedyk *LabN Consulting, L.L.C.*
S. Bryant *Futurewei Technologies*

RFC 8939
Deterministic Networking (DetNet) Data Plane: IP

Abstract

This document specifies the Deterministic Networking (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-layer and higher-layer protocol header information is used to support flow identification and DetNet service delivery. This document builds on the DetNet architecture (RFC 8655) and data plane framework (RFC 8938).

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 with 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-layer 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 [RFC8938].

The DetNet architecture models the DetNet-related data plane functions as 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 the Packet Replication Function (PRF) and Packet Elimination Function (PEF)) and reordering. The forwarding sub-layer is used to provide congestion protection (low loss, assured latency, and limited out-of-order delivery). The service sub-layer generally requires additional header fields to provide its service; for example, see [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-network basis using technologies such as MPLS [DetNet-MPLS] and Ethernet, as specified by the IEEE 802.1 TSN (Time-Sensitive Networking) task group (referred to in this document simply as "IEEE 802.1 TSN"). See [IEEE802.1TSNTG].

This document provides an overview of the DetNet IP data plane in Section 3 and 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 it is assumed that the reader is familiar with that document and its terminology.

2.2. Abbreviations

The following abbreviations are used in this document:

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
PEF	Packet Elimination Function
PREOF	Packet Replication, Elimination, and Ordering Functions
PRF	Packet Replication 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. Overview of the DetNet IP Data Plane

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-layer 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 [RFC8938].

The DetNet IP data plane primarily uses 6-tuple-based flow identification, where "6-tuple" refers to information carried in IP-layer and higher-layer protocol headers. The 6-tuple referred to in this document is the same as that defined in [RFC3290]. Specifically, the 6-tuple is destination address, source address, IP protocol, source port, destination port, and 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. Note that a 6-tuple is composed of a 5-tuple plus the addition of a DSCP component.

For some of the protocols, 5-tuples and 6-tuples cannot be used, because the port information is not available (e.g., ICMP, IPsec, and Encapsulating Security Payload (ESP)). This is also the case for flow aggregates. In such cases, using fewer fields is appropriate, such as a 3-tuple (2 IP addresses, IP protocol) or even a 2-tuple (all IP traffic between two IP addresses).

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 have the same protocol header format on the wire. Generally, the fields used in flow identification are forwarded unmodified. However, standard modification of the DSCP field [RFC2474] is not precluded.

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.

The specific procedures that are required to be implemented by a DetNet node supporting this document can be found in Section 5. The DetNet Controller Plane, as defined in [RFC8655], is responsible for providing each node with the information needed to identify and handle each DetNet flow.

```
 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-Enabled IP Network

Figure 1 illustrates a DetNet-enabled IP network. The DetNet-enabled end systems originate IP-encapsulated traffic that is identified within the DetNet domain as DetNet flows based on IP header information. 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 network, or other network technology that can carry DetNet IP traffic.

```
 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 Figures 1 and 2 can be collapsed, so IP DetNet end systems can communicate over a DetNet IP network with IP end systems.

As non-DetNet and DetNet IP packets have the same protocol header format on the wire, from a 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 flow receives the proper traffic treatment within the domain.

> Note: The operation of IEEE 802.1 TSN end systems over DetNet-enabled IP networks is not described in this document. TSN over MPLS is described in [DetNet-TSN-over-MPLS].

4. DetNet IP Data Plane Considerations

This section provides considerations related to providing DetNet service to flows that 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 value carried within an IP packet (see Section 5.1.1.3).

For DetNet-unaware IP end systems, service-level proxy functions are needed inside the DetNet domain.

When IP end systems are DetNet aware, no application-level or 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 sending packets, this means that packets are appropriately shaped on transmission and receive appropriate traffic treatment on the connected sub-network; see Sections 4.3.2 and 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.

An important additional consideration for DetNet-aware end systems is avoiding IP fragmentation. Full 6-tuple flow identification is not possible on IP fragments, as fragments don't include the transport headers or their port information. As such, it is important that applications and/or end systems use an IP packet size that will avoid fragmentation within the network when sending DetNet flows. The maximum size can be learned via Path MTU Discovery [RFC1191] [RFC8201] or via the Controller Plane. Note that Path MTU Discovery relies on ICMP, which may not follow the same path as an individual 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. Possible consequences of not having such an agreement include some flows interfering with other flows, and the traffic treatment expected for a service not being provided.

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 3.)
2. DN integrated: the end system is part of the DetNet domain. (See ES3 in Figure 3.)

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 traffic treatment required by the DetNet service is provided on both the node and any attached sub-network.

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

The DetNet service flow is mapped to the link/sub-network-specific resources using an underlying system-specific means. This implies that each DetNet-aware node on the path looks into the forwarded DetNet service flow packet and utilizes, for example, 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 scenario with three sub-networks in Figure 4, where each sub-network can provide service protection between its borders. "R" and "E" denote replication and elimination points within the sub-network.

```
     <-------------------- DetNet IP ------------------------>
                                 ______
                       ____     /      \__
            ____      /     \__/          \___   ______
+----+   __/    +====+                       +==+      \     +----+
|src |__/ Sub-N1 )   |                       |  \ Sub-N3\____| dst|
+----+  \_______/    \      Sub-network 2    |   \______/    +----+
                      \_                    _/
                        \         __     __/
                         \_______/  \___/


          +---+        +---------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 the 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 reused) 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 IPv4 and IPv6 is provided using the standard DSCP field [RFC2474] and related mechanisms.

One additional consideration for DetNet nodes that 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 Label Switching Routers (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 node-internal traffic control mechanisms used to deliver

QoS for IP-encapsulated DetNet flows are outside the scope of this document. From an encapsulation perspective, the combination of the 6-tuple (the typical 5-tuple enhanced with the DSCP) and optionally the flow label 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 resource, 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:

- Treating packets associated with an incomplete reservation as non-DetNet traffic.
- Discarding packets associated with an incomplete reservation.
- Re-marking packets associated with an incomplete reservation. Re-marking can be accomplished by changing the value of the DSCP 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 the 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 optional (flow label) 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 experiences 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 5-tuple is required. Again, this only indirectly impacts path selection for DetNet flows and this document.

4.4. DetNet Flow Aggregation

As described in [RFC8938], 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, when 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 the 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 fields and optionally also by the flow label. 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 ought to provide service to an aggregate rather than 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.

The DetNet Controller Plane **MUST** ensure that non-congestion-responsive DetNet traffic is not forwarded outside a DetNet domain.

4.5. Bidirectional Traffic

While the DetNet IP data plane must support bidirectional DetNet flows, there are no special bidirectional features within the data plane. The special case of co-routed bidirectional DetNet flows is 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 is out of the 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-layer 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 to 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 it applies to future mechanisms such as those that may be provided in YANG models [DetNet-YANG].

5.1. DetNet IP Flow Identification Procedures

IP-layer 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 requirements for flow identification, 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 system, a relay node, or 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. This includes the next protocol values defined in Section 5.1.2 and any other value, including zero. 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 it 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, ICMP, 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.

The rules defined in this section only apply when the IPv4 Protocol or IPv6 Next Header field contains the IANA-defined value for UDP or TCP.

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. ICMP

DetNet flow identification for ICMP [RFC0792] is achieved based on the protocol number in the IP header. Note that ICMP type is not included in the flow definition.

5.1.2.3. 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. Implementations **SHOULD** also allow for the field to be ignored for a specific DetNet flow.

The rules defined in this section only apply when the IPv4 Protocol or IPv6 Next Header field contains the IANA-defined value for AH or ESP.

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. Specific management and control information will be defined in future documents, e.g., [DetNet-YANG].

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 [DetNet-YANG].

5.3. DetNet IP Traffic Treatment Procedures

Implementations of 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 [DetNet-YANG]. General information on DetNet service can be found in [DetNet-Flow-Info]. Typical mechanisms used to provide different treatment to different flows include the allocation of system resources (such as queues and buffers) and provisioning of related parameters (such as shaping and policing). Support can also be provided via an underlying network technology such as MPLS [DetNet-IP-over-MPLS] or IEEE 802.1 TSN [DetNet-IP-over-TSN]. Other mechanisms than the ones used in the TSN case are 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:

- IPv4 and IPv6 Source Address field.

- IPv4 and IPv6 source address prefix length, where a zero (0) value effectively means that the Source Address field is ignored.
- IPv4 and IPv6 Destination Address field.
- IPv4 and IPv6 destination address prefix length, where a zero (0) value effectively means that the Destination Address field is ignored.
- IPv4 Protocol field. A limited set of values is allowed, and the ability to ignore this field is desirable.
- IPv6 Next Header field. A limited set of values is allowed, and the ability to ignore this field is desirable.
- For the IPv4 Type of Service and IPv6 Traffic Class fields:
    - Whether or not the DSCP field is used in flow identification. Use of the DSCP field for flow identification is optional.
    - If the DSCP field is used to identify a flow, then the flow identification information (for that flow) includes a list of DSCPs used by that flow.
- IPv6 Flow Label field. This field can be optionally used for matching. When used, this field can be used instead of matching against the Next Header field.
- TCP and UDP Source Port. Support for both exact and wildcard matching is required. Port ranges can optionally be used.
- TCP and UDP Destination Port. Support for both exact and wildcard matching is required. Port ranges can optionally be used.
- IPsec Header SPI field. Exact matching is required. Support for wildcard matching is recommended.
- For end systems, an optional maximum IP packet size that should be used for that outgoing DetNet IP flow.

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

An implementation **MUST** support ordering of the set of information used to identify an individual DetNet flow. 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 provide the traffic treatment needed to meet each flow's service requirements. This applies for aggregated and individual flows.

7. Security Considerations

Detailed security considerations for DetNet are cataloged in [DetNet-Security], and more general security considerations are described in [RFC8655]. This section exclusively considers security considerations that are specific to the DetNet IP data plane.

Security aspects that are unique to DetNet are those whose aim is to provide the specific QoS aspects of DetNet, which are primarily to deliver data flows with extremely low packet loss rates and bounded end-to-end delivery latency. Achieving such loss rates and bounded latency may not be possible in the face of a highly capable adversary, such as the one envisioned by the Internet Threat Model of BCP 72 [RFC3552] that can arbitrarily drop or delay any or all traffic. In order to present meaningful security considerations, we consider a somewhat weaker attacker who does not control the physical links of the DetNet domain but may have the ability to control a network node within the boundary of the DetNet domain.

The primary consideration for the DetNet data plane is to maintain integrity of data and delivery of the associated DetNet service traversing the DetNet network. Since no DetNet-specific fields are available in the DetNet IP data plane, the integrity and confidentiality of application flows can be protected through whatever means are provided by the underlying technology. For example, encryption may be used, such as that provided by IPsec [RFC4301] for IP flows and/or by an underlying sub-network using MACsec [IEEE802.1AE-2018] for 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 that has security implications that should be considered when determining if DetNet is a suitable technology for any given use case.

To provide uninterrupted availability of the DetNet service, provisions can be made against DoS attacks and delay attacks. To protect against DoS attacks, excess traffic due to malicious or malfunctioning devices can be prevented or mitigated -- for example, through the use of existing mechanisms such as policing and shaping applied at the input of a DetNet domain or within an edge IEEE 802.1 TSN domain. To prevent DetNet packets from being delayed by an entity external to a DetNet domain, DetNet technology definitions can allow for the mitigation of man-in-the-middle attacks -- for example, through the use of authentication and authorization of devices within the DetNet domain.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. 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>.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <https://www.rfc-editor.org/info/rfc792>.

[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>.

[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>.

[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>.

[RFC8938] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <https://www.rfc-editor.org/rfc/rfc8938>.

9.2. Informative References

[DetNet-Flow-Info] Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "DetNet Flow Information Model", Work in Progress, Internet-Draft, draft-ietf-detnet-flow-information-model-11, 21 October 2020, <https://tools.ietf.org/html/draft-ietf-detnet-flow-information-model-11>.

[DetNet-IP-over-MPLS] Varga, B., Ed., Berger, L., Fedyk, D., Bryant, S., and J. Korhonen, "DetNet Data Plane: IP over MPLS", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-over-mpls-09, 11 October 2020, <https://tools.ietf.org/html/draft-ietf-detnet-ip-over-mpls-09>.

[DetNet-IP-over-TSN] Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "DetNet Data Plane: IP over IEEE 802.1 Time Sensitive Networking (TSN)", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-over-tsn-04, 2 November 2020, <https://tools.ietf.org/html/draft-ietf-detnet-ip-over-tsn-04>.

[DetNet-MPLS] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "DetNet Data Plane: MPLS", Work in Progress, Internet-Draft, draft-ietf-detnet-mpls-13, 11 October 2020, <https://tools.ietf.org/html/draft-ietf-detnet-mpls-13>.

[DetNet-Security] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", Work in Progress, Internet-Draft, draft-ietf-detnet-security-12, 2 October 2020, <https://tools.ietf.org/html/draft-ietf-detnet-security-12>.

[DetNet-TSN-over-MPLS] Varga, B., Ed., Farkas, J., Malis, A., Bryant, S., and D. Fedyk, "DetNet Data Plane: IEEE 802.1 Time Sensitive Networking over MPLS", Work in Progress, Internet-Draft, draft-ietf-detnet-tsn-vpn-over-mpls-04, 2 November 2020, <https://tools.ietf.org/html/draft-ietf-detnet-tsn-vpn-over-mpls-04>.

[DetNet-YANG] Geng, X., Chen, M., Ryoo, Y., Fedyk, D., Rahman, R., and Z. Li, "Deterministic Networking (DetNet) Configuration YANG Model", Work in Progress, Internet-Draft, draft-ietf-detnet-yang-09, 16 November 2020, <https://tools.ietf.org/html/draft-ietf-detnet-yang-09>.

[IEEE802.1AE-2018] IEEE, "IEEE Standard for Local and metropolitan area networks-Media Access Control (MAC) Security", IEEE 802.1AE-2018, DOI 10.1109/IEEESTD.2018.8585421, December 2018, <https://ieeexplore.ieee.org/document/8585421>.

[IEEE802.1TSNTG] IEEE, "Time-Sensitive Networking (TSN) Task Group", <https://1.ieee802.org/tsn/>.

[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>.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <https://www.rfc-editor.org/info/rfc1191>.

[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>.

[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>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/info/rfc3552>.

[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>.

[RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <https://www.rfc-editor.org/info/rfc8201>.

[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>.

Acknowledgements

The authors wish to thank Pat Thaler, Norman Finn, Loa Andersson, David Black, Rodney Cummings, Ethan Grossman, Tal Mizrahi, David Mozes, Craig Gunther, George Swallow, Yuanlong Jiang, and Carlos J. Bernardos for their various contributions to this work. David Black served as technical advisor to the DetNet working group during the development of this document and provided many valuable comments. IESG comments were provided by Murray Kucherawy, Roman Danyliw, Alvaro Retana, Benjamin Kaduk, Rob Wilton, and Érik Vyncke.

Contributors

The editor of this document wishes to thank and acknowledge the following people who contributed substantially to the content of this document and should be considered coauthors:

Jouni Korhonen
Email: jouni.nospam@gmail.com

Andrew G. Malis
Malis Consulting
Email: agmalis@gmail.com

Authors' Addresses

Balázs Varga (EDITOR)
Ericsson
Budapest
Magyar Tudosok krt. 11.
1117
Hungary
Email: balazs.a.varga@ericsson.com

János Farkas
Ericsson
Budapest
Magyar Tudosok krt. 11.
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

Stewart Bryant
Futurewei Technologies
Email: sb@stewartbryant.com