DetNet P. Thubert, Ed.
Internet-Draft Cisco Systems
Intended status: Standards Track F. Yang
Expires: 26 August 2022 Huawei Technologies
22 February 2022

IPv6 Options for DetNet
draft-pthubert-detnet-ipv6-hbh-07

Abstract

RFC 8938, the Deterministic Networking Data Plane Framework relies on the 6-tuple to identify an IPv6 flow. But the full DetNet operations require also the capabilities to signal meta-information such as a sequence within that flow, and to transport different types of packets along the same path with the same treatment, e.g., Operations, Administration, and Maintenance packets and/or multiple flows with fate and resource sharing. This document introduces new IPv6 options that signal that path and redundancy information to the intermediate DetNet relay and forwarding nodes.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Section 2 of the Deterministic Networking Problem Statement [DetNet-PBST] introduces the concept of Deterministic Networking (DetNet) to the IETF. DetNet extends the reach of lower layer technologies such as Time-Sensitive Networking (TSN) [IEEE 802.1 TSN] and Timeslotted Channel Hopping (TSCH) [IEEE Std. 802.15.4] over IPv6 and MPLS [RFC8938], to provide bounded latency and reliability guarantees over an end-to-end layer-3 nailed-down path.

The "Deterministic Networking Architecture" [DetNet-ARCH] details the contribution of layer-3 protocols, and defines three planes: the Application (User) Plane, the Controller Plane, and the Network Plane. [DetNet-ARCH] places an emphasis on the centralized model whereby a controller instantiates a DetNet state in the routers that is located based on matching information in the packet.

The "Deterministic Networking Data Plane Framework" [RFC8938] relies on the 6-tuple to identify an IPv6 flow. But the full DetNet operations require also the capabilities to signal meta-information such as a sequence within that flow, and to transport different types of packets along the same path with the same treatment. For instance, it is required that Operations, Administration, and Maintenance (OAM) [RFC6291] packets and/or multiple flows share the same fate and resource sharing over the same Track or the same Traffic Engineered (TE) [RFC3272] DetNet path. This document proposes a layer-3 signaling that is independent of the upper layer information, to locate the DetNet state and enable the same forwarding nehavior for the data flows and the OAM packets.

The "6TiSCH Architecture" [6TiSCH-ARCH] leverages RPL, the "Routing Protocol for Low Power and Lossy Networks" [RPL] and introduces concept of a Track as a highly redundant RPL Destination Oriented Directed Acyclic Graph (DODAG) rooted at the Track Ingress. The Track is indicative of a layer-3 forwarding behavior (e.g., next hops)as opposed to indicative of the upper layer content, so it is more in line with the DetNet needs than the 6-tuple.

A Track may for instance be installed using RPL route projection [RPL-PDAO]. In that case, the TrackId is an index from a namespace associated to one IPv6 address of the Track Ingress node, and the Track that an IPv6 packet follows is signaled by the combination of the source address (of the Track Ingress node), and the TrackID placed in a RPL Option [RFC6553] located in an IPv6 Hop-by-Hop (HbH) Options Header [IPv6] in the IPv6 packet.

The "Reliable and Available Wireless (RAW) Architecture/Framework" [RAW-ARCH], extends the DetNet Network Plane to accomodate one or multiple hops of homogeneous or heterogeneous wireless technologies, e.g. a Wi-Fi6 Mesh or parallel radio access links combining Wi-Fi and 5G. The RAW Architecture reuses the concept of Track and introduces a new dataplane component, the Path Selection Engine (PSE), to dynamically select a subpath and maintain the required quality of service within a Track in the face of the rapid evolution of the medium properties.

With [IPv6], the behavior of a router upon an IPv6 packet with a HbH Options Header has evolved, making the examination of the header by routers along the path optional, as opposed to previously mandatory. Additionally, the Option Type for any option in a HbH Options Header encodes in the leftmost bits whether a router that inspects the header should drop the packet or ignore the option when encountering an unknown option. Combined, these capabilities enable a larger use of the header beyond the boundaries of a limited domain, as examplified by the change of behavior of the RPL data plane, that was changed to allow a packet with a RPL option to escape the RPL domain in the larger Internet [RFC9008].

"IPv6 Hop-by-Hop Options Processing Procedures" [HbH-UPDT] further specifies the procedures for how IPv6 Hop-by-Hop options are processed to make their processing even more practical and increase their use in the Internet. In that context, it makes sense to consider Hop-by-Hop Options to transport the information that is relevant to DetNet.

As opposed to the HbH EH, the Destination Option Header (DOH) is only read by the destination of the packet, which can be one at a time the collection of nodes listed in a Routing Extension Header (RH) if the DOH is placed before the RH.

This document introduces new IPv6 Options, the DetNet Redundancy Information Option and the DetNet Path Options, that signal the DetNet information to the intermediate DetNet nodes in an abstract form, that is pure layer-3 and agnostic of the transport layer. The options are placed in either a HbH EH or in a DOH, which happens when the next node that needs to process the option is the IPv6 destination in the IPv6 header.

This pure layer-3 technique alines DetNet with the IPv6 architecture and opens to the progress / extensions done elsewhere for IPv6; e.g., if the DetNet path leverages Segment routing (SRv6) [RFC8402] for some reason - there are plausible ones in RAW -, the Segment Routing Header (SRH) [RFC8754] is inserted after the HbH and/or DOH by the PE and both are readily accessible for the on-path routers without the need of a deeper inspection of the packet (up to and beyond the transport header).

For instance, the DetNet Redundancy Information Option may be placed in a DOH before an SRH that signals the exhaustive list of the DetNet relays along the path of the packet, so every relay can process the redundancy information therein, while the DetNet Strict Path Option would be placed in an HbH EH to be read by every DetNet forwarding node, and intercepted should it strays away from its path.

2. Terminology

Timestamp semantics and timestamp formats used in this document are defined in "Guidelines for Defining Packet Timestamps" [RFC8877].

The Deterministic Networking terms used in this document are defined in the "Deterministic Networking Architecture" [DetNet-ARCH].

The terms Track and TrackID are defined in the "6TiSCH Architecture" [6TiSCH-ARCH].

3. Applicability

The "Deterministic Networking (DetNet) Data Plane: IP" [RFC8939] illustrates the need for DetNet Services at the IP network layer in conformance to the DetNet architecture [DetNet-ARCH] and data plane framework [RFC8938]). As the specification does not introduce a new header for DetNet, it also lacks the capability to signal PREOF at the network layer. Instead, the information of application flows (the 6-tuple), on which the network layer has no control, is used directly to makes network layer forwarding decisions. This confuses the signaling of the application-layer "water" that being transported and with the network-layer "pipe" that transports it and makes it problematic to aggregate applications flows and OAM for a equal treatment. It is thus desirable to introduce a new signaling to enable PREOF and flow aggregation independently of the application flow that is more appropriate for network layer operations.

[I-D.varga-detnet-ip-preof] provides a methods for signaling PREOF over UDP that combines the MPLS PREOF signaling defined in "DetNet Data Plane: MPLS" [RFC8964] and the mechanisms defined in "DetNet Data Plane: MPLS over UDP/IP" [RFC9025]. This approach provides IP version independence, allows to reuse MPLS structures and possibly to bridge MPLS DetNet flows onto IP with minimal mapping effort. It is thus specially valuable in environments where flows can be transported in any combination of MPLS and IP. OTOH, the signaled information inherits the limitations of MPLS in terms of stack structure, and is available after the transport header, which can be harder to reach for a hardware solution, in particular in the presence of a variable header chain at the network layer.

This draft proposes a native IPv6 solution that transports enriched DetNet PREOF information in IPv6 Extension Headers. This method reduces the encapsulation overhead, can be combined with the use of "IPv6 Segment Routing (SRv6) Header (SRH)" [RFC8754], and simplifies the access to the PREOF data by placing it early in the network header chain. The forwarding information (the path) is advertised independently of the application flow information (observable as the

6-tuple). This enables any mix and match of flows and OAM data over the same path with the same treatment. The method is thus suited for environments that are IPv6-only, where flows can be aggregated over larger pipes and disaggregated again, but it is harder to combine with MPLS and requires that all nodes on path can parse at least the IPv6 Extension Header that signal the path.

Transported in IPv6 Extension Headers, the DetNet options are easier to reach for a hardware or otherwise constrained implementation. A DetNet-aware end-system (see section 4.2 of [DetNet-ARCH]) may place the options in the header chain when constructing the packet, in which case there is no need of an encapsulation.

Alternatively, the source end system may signal the flow information some other way, or it may lack the full DetNet awareness; in that case the DetNet path endpoints are the provider Edge (PE) routers (see Figure 1 reproducing figure 5 of [DetNet-ARCH]) and the Ingress PE needs to encapsulate the packets to add the HbH options.

In Figure 1, the DetNet end systems may be f-aware and signal an IPv6 flow using the 6-tuple for the End-to-End service, but may not be s-aware, and may not sequence the packets for Packet Replication, Elimination, and Ordering Functions (PREOF), which operate at the detNet Service Layer. In that case, the Ingress PE will encapsulate the packets for this and possibly other flows to provide a common DetNet Service with OAM and PREOF, across the DetNet-1 service provider network, terminating the tunnel at the Egress PE router.

```
                                                          _
     / \     +----DetNet-UNI (U)                                   / \
    /App\    |                                                    /App\
   /-----\   |                                                   /-----\
   | NIC |   v         ________                                  | NIC |
   +--+--+   _____    /        \             DetNet-UNI (U) --+  +--+--+
      |     /     \__/          \                             |     |
      |    / +----+    +----+    \_____                       |     |
      |   /  |    |    |    |          \_______               |     |
      +------U PE +----+ P  +----+             \          _   v     |
          |  |    |    |    |    |              |     ___/ \        |
          |  +--+-+    +----+    |       +----+ |    /      \_      |
          \     |                |       |    | |   /         \     |
           \    |   +----+    +--+-+  +--+PE  |------         U-----+
            \   |   |    |    |    |  |  |    | |   \_      _/
             \  +---+ P  +----+ P  +--+  +----+ |     \____/
              \___  |    |    |    |           /
                  \ +----+__  +----+     DetNet-1    DetNet-2
      |            \_____/  \___________/                           |
      |                                                             |
      |      |     End-to-End Service         |     |         |     |
      <------------------------------------------------------------->
      |      |     DetNet Service             |     |         |     |
      |      <------------------------------------------------>     |
      |      |                                |     |         |     |
```

Figure 1: Figure 5 of RFC 8655, Reproduced

4. The DetNet Options

This document defines new IPv6 options for DetNet to signal path and a reliability information (e.g., sequencing) to the DetNet layers. Those options are to be placed in the IPv6 HbH Options Header, which is found right after the outer IPv6 header in the DetNet packet and immediately reachable for the forwarding engine. The format of the options follow the generic definition in section 4.2 of [IPv6]. For each tyoe of option, the draft allows to express the information in different fashions, depending on the use case, and possibly carrying an information that plays the same role at another layer, in which case the format of the information is opaque.

The reliability information may be inherited from another layer as long as the value is guaranteed to be unique within a reasonable set of sequential packet so all packets with the same value are redundant. Timestamping can be used as an alternate sequencing technique, that avoids maintaining per-path state at the path ingress, which is feasible for nodes that maintain a very precise sense of time (e.g., from GPS or PTP) for their DetNet operations.

As long as the time granularity is in the order of a few bytes transmission, the system timestamp provides an absolute sense of ordering over a very long period across all paths for which this node is ingress, and thus within any of those. Alternatively, the draft allows to combine a rough time stamp (e.g., from a system clock synchronized by NTP) and a sequence counter that differntiates the packets that are stamped within the timer resolution.

If a DetNet Path option (see Section 4.2), including the RPL Option, is present in the same HbH Option Header as a DetNet Redundancy Information option (see Section 4.1), then the redundancy information applies to the signaled path across all flows that traverse that path; else the redundancy information applies to the flow indicated by the 6-tuple [RFC8938].

4.1. DetNet Redundancy Information Option

The DetNet Redundancy Information Option helps discriminate copies of a same packet vs. different packets, and is useful for service-sublayer Packet Replication Elimination and Ordering Functions (PREOF). The option may be placed either in an HbH or a DoH EH, e.g., prior to a Segment Routing Header (SRH) [RFC8754] that lists the DetNet relays. A sequence counter is probably the most typical expression of the redundancy information, but it is not the only way to identify a packet and/or enable reordering, e.g., a timestamp can be seen as a large sequence counter with gaps.

It is also possible that a packet is divided in elements such as network-coded fragments. In that case, the pieces are discriminated with an opaque 8-bit fragment tag. The goal is to retain one copy of each fragment but not reorder them.

A packet sequence can be expressed uniquely as a wrapping counter, represented as an unsigned integer in the option. In that case, the size of the representation MUST be large enough to cover at least 3 times the upper bound on out-of-order packet delivery in terms of number of packets. The sequence counter may be copied from a field in another protocol, and it is possible that the value 0 is reserved when wrapping, to the option offers both possibilities, wrapping to either 0 or to 1.

This specification also allows to use a time stamp for the packet redundancy information, in conformance with the recommendations in [RFC8877]. This can be accomplished by utilizing the Precision Time Protocol (PTP) format defined in IEEE Std. 1588 [IEEE Std. 1588] or Network Time Protocol (NTP) [RFC5905] formats. In that case, the timestamp resolution at the origin node that builds the option MUST be fine enough to ensure that two consecutive packets are never stamped with the same value. There is no requirement for this particular stamping function that the sense of time at the origin node is synchronized with the rest of the DetNet network.

IEEEE TSN [IEEE 802.1 TSN] defined a redundancy tag (R-Tag) for the IEEE Std. 802.1CB Frame Replication and Elimination for Reliability (FRER). The R-Tag is a structured field and its content is subject to evolve; but the expectation for this specification is that the overall size remains 48 bits and that the 48-bit value is different for a large number of contiguous frames. When transporting TSN frames in a DetNet packet, it is possible to leverage the R-Tag as Redundancy information, though it cannot be assumed that the R-Tag is sequentially incremented; so it can be used for packet duplicate elimination but it is not suitable not for packet re-ordering.

This specification also allows for an hybrid model with a coarse grained packet sequence within a coarse grained time stamp. In that case, both a time stamp option and a wrapping counter options are found, and the counter is used to compare packets with the same time stamp and ignored otherwise In that case, the size of the representation of the counter MUST be large enough to cover at least 3 times the number of packets that may be sent with the same value of time stamp.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Option Type  |  Opt Data Len |  R.I. Type    | Fragment Tag  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.          Redundancy Information (variable Size)               .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Redundancy Information Option Format

Redundancy Information Option fields:

Option Type: 8-bit identifier of the type of option. Value TBD by IANA; if the processing IPv6 node does not recognize the Option Type it MUST skip over this option and continue processing the header (act =00); the Option Data of that option cannot change en route to the packet's final destination (chg=0). The

Opt Data Len: 8-bit length of the option data.

Fragment Tag: 8-bit field, set to 0 when the packet is sent in entirety; packets with the same Redundancy Information and different fragments tags MUST be considered as different by the elimination function and are not subject to ordering based on the Tag.

Redundancy Information Type: 8-bit identifier of the type of Redundancy information. Value to be confirmed by IANA.

Seq. Type Value	Category	Common Name	Redundancy Information Format
1	Wrapping Counter	Basic Sequence Counter	32-bit unsigned integer
2	Wrapping Counter	Zero-avoiding Sequence Counter	32-bit unsigned integer, wraps to 1
3	Wrapping Counter	RPL Sequence Counter	8-bit RPL sequence, see section 7. of [RPL]
11	Time Stamp	Fractional NTP	NTP 64-bit Timestamp Format, see section 4.2.1. of [RFC8877]
12	Time Stamp	Short NTP	NTP 32-bit Timestamp Format, see section 4.2.2. of [RFC8877]
13	Time Stamp	PTP	PTP 80-bit Timestamp Format, see [IEEE Std. 1588]
14	Time Stamp	Short PTP	PTP 64-bit Truncated Timestamp Format, see section 4.3. of [RFC8877]
24	Structured Unique Tag	TSN Redundancy Tag	48-bit opaque

Table 1: Redundancy Information Type values (suggested)

Redundancy Information: Variable size, as indicated in Table 1.

4.2. DetNet Path Options

The DetNet Architecture [DetNet-ARCH] assigns a DetNet flow "to specific paths through a network", but is not specific on how the path is then signaled in the packet. The DetNet Data Plane Framework [RFC8938] relies on the 6-tuple to identify an IPv6 flow and implicitely the path could be indexed by the flow identification. But this requires to maintain one path per flow and makes it difficult to assign other traffic such as OAM to the same path.

This draft provides aditional means to signal the path in which the flow is placed separately from the flow indentification, and independantly of the transport layer, so a path can be shared between one or more flows and OAM packets across IP address families. All the packets that are assigned to the same path are subject to the same DetNet forwarding treatment.

the DetNet expectation is that a PCE sets up a state at the DetNet forwarding sublayer to instruct each hop on how to process the DetNet flows. The DetNet Path Options when present contains information that MUST be used to select the DetNet state installed and if the DetNet state does not exist then the packet cannot be forwarded.

4.2.1. DetNet Strict Path Option

In complement to the RPL option, this specification defines a protocol-independent Strict Path Identifier, which is also taken from a namespace indicated by the IPv6 source address of the packet.

The DetNet Strict Path Option is to be used in a limited domain to indicate a routing state that must be present in all nodes to ensure that the packet is routed along a strictly predefined path, for instance pointing at a specific next hop with reserved resources for buffers and bandwidth. For that reason all the routers along the path are expected to support the option and own a state indexed by the Strict Path ID indicated therein.

The option is placed in an HbH EH to be seen by all routers on path. The path indicated therein may also be used by the service sublayer, to signal the scope where the redundancy information is unique across a number of packets large enough to ensure that a forwarding node never has to handle different packets with the same redundancy information, though the same value may be found for packets with a different path information.

The typical DetNet path is typically contained under a single administrative control or within a closed group of administrative control; these include campus-wide networks and private WANs

[DetNet-ARCH]. The typical expectation is that all nodes along a DetNet path are aware of the path and actively maintain a forwarding state for it. The DetNet Strict Path Option (see Section 4.2.1) is designed for that environment; if a packet escapes the local domain, a router that does not support the option will intercept it and return an error to the source.

In other environments such as RAW, it might be that the service-layer protection concentrates on just segments of the end-to-end path. In that case, the service-sublayer protection may require the signaling of both redundancy and path information, though the path information is potentially not used by some of the intermediate routers and may not be used for forwarding at all. The path information may also relate to segments that are installed along the path using a DetNet forwarding state as opposed to, say, source routing. In either case the DetNet Loose Path Option Section 4.2.2 can be used to signal the path without incurring an ICMP Error from an intermediate node.

An intermediate router that supports the DetNet Strict Path Option but is missing the necessary state to forward along the indicated path must drop the packet and return an ICMP error.code 0 pointing at the offset of the Strict Path ID in the DetNet Strict Path Option.

DetNet can also leverage the RPL Option that signals a Track in the RPL Packet Information (RPI) [RFC6553]. There are 2 versions of the RPL option, defined respectively in [RPL] with the act bits [IPv6] set to dropped the packet when the option is unknown, that defined in[RFC9008] which let the option be ignored.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Option Type  |  Opt Data Len |                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: DetNet Strict Path Option Format

Redundancy Option fields:

Option Type: 8-bit identifier of the type of option. Value TBD by IANA; if the processing IPv6 node does not recognize the Option Type it must discard the packet and send an ICMP Parameter Problem, Code 2, message to the packet's Source Address (act =10); the Option Data of that option cannot change en route to the packet's final destination (chg=0).

Opt Data Len: 8-bit length of the option data, set to 2.

Strict Path ID: 16-bit identifier of the DetNet Path, taken from a local namespace associated with the IPv6 source address of the packet.

4.2.2. DetNet Loose Path Option

The DetNet Loose Path Option transports a Loose Path identifier which is taken from a namespace indicated by the Origin Autonomous System (AS). When the DetNet path is contained within a single AS, the Origin Autonomous System field can be left to 0 indicating local AS. The option may be placed either in an HbH or a DoH EH, but the preferred method is a DOH that precedes an RH such as SRH.

The DetNet Loose Path Option is to be used to signal a path that may be loose and may exceed the boundaries of a local domain; a portion of the hops may traverse routers in the wider internet that will not leverage the option and are expected to ignore it. For instance, the path information may signal a specific topology in a multi-topology network and is only important for nodes that participate to more than one topology.

An intermediate router that supports the DetNet Loose Path Option but is missing the necessary state to forward along the indicated path must ignore the DetNet Loose Path Option, but it should raise a management alert as this is an unexpected situation with a limited chance that the packet may loop till TTL.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Option Type  |  Opt Data Len |    Origin Autonomous System   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Loose Path ID                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: DetNet Loose Path Option Format

Redundancy Option fields:

Option Type: 8-bit identifier of the type of option. Value TBD by IANA; if the processing IPv6 node does not recognize the Option Type it MUST skip over this option and continue processing the header (act =00); the Option Data of that option cannot change en route to the packet’s final destination (chg=0).

Opt Data Len: 8-bit length of the option data, set to 6.

Origin Autonomous System: 16-bit identifier of the Autonomous

Systems (AS) that originates the path. The value of 0 signals a DetNet path that is constrained within the local AS or the local administrative DetNet domain.

Loose Path ID: 32-bit identifier of the DetNet Path, taken from a local namespace associated with the origin AS of the DetNet path.

4.3. RPL Packet Information

6TiSCH [6TiSCH-ARCH] and RAW [RAW-ARCH] signal a Track using a RPL Option [RFC6553] with a RPLInstanceID used as TrackID. This specification reuses the RPL option as a method to signal a DetNet path. In that case, the Projected-Route 'P' flag [RPL-PDAO] MUST be set to 1, and the O, R, F flags, as well as the Sender Rank field, MUST be set to 0 by the originator, forwarded as-is, and ignored on reception.

5. Encapsulation of DetNet Options

In this section, encapsulations of three DetNet Options are specified separately in the scenarios of pure IPv6 and SRv6.

5.1. IPv6 Network

The DetNet Strict Path Option is intended to be placed in an IPv6 HbH EH since it must be processed by every DetNet forwarding node along the path. The DetNet Loose Path Option and the DetNet Redundancy Information Option may also carried in an IPv6 HbH Option header the case where the set of routers that need the information does not match the destinations along a source route path; those options are intended to be ignored by unaware intermediate routers.

In the specific case where path selection and PREOF are end-to-end performed between DetNet edge nodes, Redundancy Information Option can be alternatively placed in IPv6 Destination Option header. The encapsulation options are shown in Figure 5 and Figure 6.

```
     +-----------------------------------+
     |          DetNet App-Flow          |
     |        (original IP) Packet       |
     +-----------------------------------+
     |             other EHs             |
     +-----------------------------------+ <--\
     |        IPv6 Hop-by-Hop Ex Hdr     |    |
     |          (DetNet RI Option)       | DetNet Options
     | (DetNet Strict/Loose Path Option) |    |
     +-----------------------------------+ <--/
     |            IPv6 Header            |
     +-----------------------------------+
     |             Data-Link             |
     +-----------------------------------+
     |             Physical              |
     +-----------------------------------+
```

Figure 5: DetNet IPv6 Option Encapsulation Alternative 1

```
     +-----------------------------------+
     |          DetNet App-Flow          |
     |        (original IP) Packet       |
     +-----------------------------------+
     |        other EHs such as RH       |
     +-----------------------------------+ <--\
     |      IPv6 Destination Ext Hdr     |    |
     |         (DetNet RI Option)        |    |
     +-----------------------------------+ DetNet Options
     |      IPv6 Hop-by-Hop Ext Hdr      |    |
     | (DetNet Strict/Loose Path Option) |    |
     +-----------------------------------+ <--/
     |            IPv6 Header            |
     +-----------------------------------+
     |             Data-Link             |
     +-----------------------------------+
     |             Physical              |
     +-----------------------------------+
```

Figure 6: DetNet IPv6 Option Encapsulation Alternative 2

5.2. Segment Routing over IPv6 Network

In SRv6, partial or all of DetNet forwarding and relay nodes may be represented by SRv6 SIDs to determine a specific path for a DetNet flow. In the former case, DetNet Strict Path Option would be placed in an HbH EH to be read by every DetNet forwarding node, and intercepted should it strays away from its path. In the latter case, three DetNet Options can be placed either in an HbH EH or in a DOH EH before an SRH, as two encapsulation options are being functionally equivalent, as shown in Figure 7 .

```
+-----------------------------------+
|         DetNet App-Flow           |
|       (original IP) Packet        |
+-----------------------------------+
|       Segment Routing Header      |
+-----------------------------------+ <--\
|        IPv6 Hop-by-Hop Ex Hdr     |    |
|          (DetNet RI Option)       | DetNet Options
| (DetNet Strict/Loose Path Option) |    |
+-----------------------------------+ <--/
|            IPv6 Header            |
+-----------------------------------+
|             Data-Link             |
+-----------------------------------+
|             Physical              |
+-----------------------------------+
```

Figure 7: DetNet IPv6 Option Encapsulation in SRv6 Alternative 1

In the case where the SRv6 SRH signals the exhaustive list of the Detnet relays along the path, it is recommended to place the DetNet Redundancy Information Option in a DOH EH before the SRH, so that it is processed by every relay node therein without burdening the intermediate DetNet forwarding nodes, as illustrated in Figure 8 and Figure 9.

If all the nodes that process the loose path information are also listed in the SRH, then the DetNet Loose Path Option may also be placed in the DOH, as shown in Figure 8

```
            +-----------------------------------+
            |          DetNet App-Flow          |
            |        (original IP) Packet       |
            +-----------------------------------+
            |       Segment Routing Header      |
            +-----------------------------------+ <--\
            |        IPv6 Destination Ex Hdr    |    |
            |          (DetNet RI Option)       | DetNet Options
            |     (DetNet Loose Path Option)    |    |
            +-----------------------------------+ <--/
            |            IPv6 Header            |
            +-----------------------------------+
            |             Data-Link             |
            +-----------------------------------+
            |             Physical              |
            +-----------------------------------+
```

      Figure 8: DetNet IPv6 Option Encapsulation in SRv6 Alternative 2

Unless the SRH is a strict routing header indicating all the hops on the path, the DetNet Strict Path Option must remain separate in a HbH EH, to be observed by all routers on path, as shown in Figure 9

```
            +-----------------------------------+
            |          DetNet App-Flow          |
            |        (original IP) Packet       |
            +-----------------------------------+
            |       Segment Routing Header      |
            +-----------------------------------+ <--\
            |        IPv6 Destination Ex Hdr    |    |
            |          (DetNet RI Option)       |    |
            +-----------------------------------+ DetNet Options
            |       IPv6 Hop-by-Hop Ex Hdr      |    |
            |    (DetNet Strict Path Option)    |    |
            +-----------------------------------+ <--/
            |            IPv6 Header            |
            +-----------------------------------+
            |             Data-Link             |
            +-----------------------------------+
            |             Physical              |
            +-----------------------------------+
```

      Figure 9: DetNet IPv6 Option Encapsulation in SRv6 Alternative 3

6. Security Considerations

7. IANA Considerations

7.1. New Subregistry for the Redundancy Type

This specification creates a new Subregistry for the "Redundancy Type of the Redundancy Option" under the "Internet Protocol Version 6 (IPv6) Parameters" registry [IPV6-PARMS].

* Possible values are 8-bit unsigned integers (0..255).

* Registration procedure is "IETF Review" [RFC8126].

* Initial allocation is as Suggested in Table 2:

Suggested Value	Meaning	Reference
1	Basic Sequence Counter	THIS RFC
2	Zero-avoiding Sequence Counter	THIS RFC
3	RPL Sequence Counter	THIS RFC
11	Fractional NTP time stamp	THIS RFC
12	Short NTP time stamp	THIS RFC
13	PTP time stamp	THIS RFC
14	Short PTP time stamp	THIS RFC
24	TSN Redundancy Tag	THIS RFC

Table 2: Redundancy Information Type values

7.2. New Hop-by-Hop Options

This specification updates the "Destination Options and Hop-by-Hop Options" under the "Internet Protocol Version 6 (IPv6) Parameters" registry [IPV6-PARMS] with the (suggested) values below:

Hexa	act	chg	rest	Description	Reference
0x12	00	0	10010	DetNet Redundancy Information Option	THIS RFC
0x93	10	0	10011	DetNet Strict Path Option	THIS RFC
0x14	00	0	10100	DetNet Loose Path Option	THIS RFC

Table 3: DetNet Hop-by-Hop Options

8. Acknowledgments

TBD

9. References

9.1. Normative References

[RPL] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/info/rfc6550>.

[RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <https://www.rfc-editor.org/info/rfc6553>.

[IPv6] 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>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8877] Mizrahi, T., Fabini, J., and A. Morton, "Guidelines for Defining Packet Timestamps", RFC 8877, DOI 10.17487/RFC8877, September 2020, <https://www.rfc-editor.org/info/rfc8877>.

[HbH-UPDT] Hinden, R. M. and G. Fairhurst, "IPv6 Hop-by-Hop Options Processing Procedures", Work in Progress, Internet-Draft, draft-hinden-6man-hbh-processing-01, 2 June 2021, <https://datatracker.ietf.org/doc/html/draft-hinden-6man-hbh-processing-01>.

[DetNet-ARCH] 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>.

[RFC9008] Robles, M.I., Richardson, M., and P. Thubert, "Using RPI Option Type, Routing Header for Source Routes, and IPv6-in-IPv6 Encapsulation in the RPL Data Plane", RFC 9008, DOI 10.17487/RFC9008, April 2021, <https://www.rfc-editor.org/info/rfc9008>.

[6TiSCH-ARCH] Thubert, P., Ed., "An Architecture for IPv6 over the Time-Slotted Channel Hopping Mode of IEEE 802.15.4 (6TiSCH)", RFC 9030, DOI 10.17487/RFC9030, May 2021, <https://www.rfc-editor.org/info/rfc9030>.

[RAW-ARCH] Thubert, P., Papadopoulos, G. Z., and L. Berger, "Reliable and Available Wireless Architecture/Framework", Work in Progress, Internet-Draft, draft-pthubert-raw-architecture-09, 7 July 2021, <https://datatracker.ietf.org/doc/html/draft-pthubert-raw-architecture-09>.

9.2. Informative References

[I-D.varga-detnet-ip-preof] Varga, B., Farkas, J., and A. G. Malis, "Deterministic Networking (DetNet): DetNet PREOF via MPLS over UDP/IP", Work in Progress, Internet-Draft, draft-varga-detnet-ip-preof-02, 1 February 2022, <https://datatracker.ietf.org/doc/html/draft-varga-detnet-ip-preof-02>.

[RPL-PDAO] Thubert, P. and R. A. Jadhav, "Root initiated routing state in RPL", Work in Progress, Internet-Draft, draft-ietf-roll-dao-projection-23, 13 January 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-roll-dao-projection-23>.

[RFC3272] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, DOI 10.17487/RFC3272, May 2002, <https://www.rfc-editor.org/info/rfc3272>.

[RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <https://www.rfc-editor.org/info/rfc6291>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <https://www.rfc-editor.org/info/rfc5905>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

[DetNet-PBST]
Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <https://www.rfc-editor.org/info/rfc8557>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

[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/info/rfc8938>.

[RFC8939] Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939, November 2020, <https://www.rfc-editor.org/info/rfc8939>.

[RFC8964] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: MPLS", RFC 8964, DOI 10.17487/RFC8964, January 2021, <https://www.rfc-editor.org/info/rfc8964>.

[RFC9025] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: MPLS over UDP/IP", RFC 9025, DOI 10.17487/RFC9025, April 2021, <https://www.rfc-editor.org/info/rfc9025>.

[IEEE Std. 802.15.4] IEEE standard for Information Technology, "IEEE Std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks".

[IEEE 802.1 TSN] IEEE 802.1, "Time-Sensitive Networking (TSN) Task Group", <http://www.ieee802.org/1/pages/tsn.html>.

[IEEE Std. 1588] IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Standard 1588, <https://ieeexplore.ieee.org/document/4579760/>.

[IPV6-PARMS] IANA, "Internet Protocol Version 6 (IPv6) Parameters", <https://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml>.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
France
Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Fan Yang
Huawei Technologies
China
Email: shirley.yangfan@huawei.com