

DetNet
Internet-Draft
Intended status: Standards Track
Expires: 4 January 2026

B. Varga, Ed.
F. Fejes
Ericsson
3 July 2025

Deterministic Networking SRv6 Data Plane
draft-varga-detnet-srv6-data-plane-02

Abstract

This document specifies the Deterministic Networking (DetNet) data plane when operating over an IPv6/SRv6 Packet Switched Network. It leverages existing IPv6 encapsulations using DetNet specific SIDs and optionally the Traffic Engineering mechanisms provided by SRv6. This document builds on the DetNet architecture and data plane framework.

Status of This Memo

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

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

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

This Internet-Draft will expire on 4 January 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Terminology	3
2.1. Terms Used in This Document	3
2.2. Abbreviations	3
2.3. Requirements Language	4
3. DetNet SRv6 Data Plane Overview	4
3.1. DetNet sub-layers with SRv6 Data Plane	4
3.2. DetNet SRv6 Data Plane Scenarios	5
4. SRv6-Based DetNet Data Plane Solution	7
4.1. DetNet Over SRv6 Encapsulation Components	7
4.2. SRv6 Data Plane Encapsulation	7
4.2.1. DetNet specific SID	8
4.2.2. Flow-ID	9
4.2.3. SeqNum	10
4.3. Service Sub-Layer Related Processing	10
4.3.1. Packet Replication Function Processing	10
4.3.2. Packet Elimination Function Processing	11
4.3.3. Packet Ordering Function Processing	11
4.4. Forwarding Sub-Layer Related Processing	12
5. Security Considerations	12
6. IANA Considerations	12
7. Acknowledgements	12
8. Appendix A. Illustrations	12
9. Normative References	15
Authors' Addresses	16

1. Introduction

Deterministic Networking (DetNet) is a service that can be offered by a network to DetNet flows. DetNet provides a capability for the delivery of data flows with extremely low packet loss rates and bounded end-to-end delivery latency. General background and concepts of DetNet can be found in the DetNet Architecture [RFC8655].

The purpose of this document is to describe the use of the IPv6/SRv6 data plane to establish and support DetNet flows. The DetNet Architecture models the DetNet related data plane functions decomposed into two sub-layers: a service sub-layer and a forwarding sub-layer. The service sub-layer is used to provide DetNet service functions such as protection and reordering. At the DetNet data plane a new set of functions (PREOF) provide the service sub-layer specific tasks. The forwarding sub-layer is used to provide forwarding assurance (low loss, assured latency, and limited out-of-order delivery). The use of the functionalities of the DetNet service sub-layer and the DetNet forwarding sub-layer require careful design and control by the controller plane in addition to the DetNet specific use of SRv6 encapsulation as specified by this document.

This document specifies the DetNet data plane operation and the on-wire encapsulation of DetNet flows over an IPv6/SRv6-based Packet Switched Network (PSN) using the service reference model.

2. Terminology

2.1. Terms Used in This Document

This document uses the terminology established in the DetNet architecture [RFC8655] and in the SRv6 Network Programming [RFC8986]. The reader is assumed to be familiar with that document and its terminology.

The following terminology is introduced in this document:

Flow-ID	A DetNet "service" identifier that is used between DetNet nodes that implement the DetNet service sub-layer functions. A Flow-ID is used to identify a DetNet flow at DetNet service sub-layer at a receiving DetNet node.
SeqNum	A SeqNum is used for sequencing information of a DetNet flow at the DetNet service sub-layer.

2.2. Abbreviations

The following abbreviations are used in this document:

ARG	Arguments.
DetNet	Deterministic Networking.
FUNCT	Function.

LOC	Locator.
PEF	Packet Elimination Function.
POF	Packet Ordering Function.
PREOF	Packet Replication, Elimination and Ordering Functions.
PRF	Packet Replication Function.
SeqNum	Sequence Number.

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 SRv6 Data Plane Overview

3.1. DetNet sub-layers with SRv6 Data Plane

A straight forward approach utilizing SRv6 for a DetNet service sub-layer is based on the DetNet specific SID [I-D.varga-spring-preof-sid] and by optionally utilizing existing SRv6 Traffic Engineering encapsulations and mechanisms using SRH between the DetNet Relay nodes. Background on SRv6 Network Programming can be found in [RFC8986].




Figure 1: DetNet Adaptation to SRv6 Data Plane used in both sub-layers

The DetNet SRv6 data plane representation is illustrated in Figure 1. The service sub-layer includes a DetNet specific SID that contains the Flow-ID and the SeqNum. More details of DetNet specific SID behaviors are described in [I-D.varga-spring-preof-sid].

A node operating on a received DetNet flow at the Detnet service sub-layer terminates the SRv6 encapsulation, uses the local context associated with a Flow-ID, to determine which local DetNet operation(s) are applied to that packet. It is important to note that Flow-ID values are driven by the receiver, not the sender.

Optionally, the DetNet forwarding sub-layer is supported by the SRv6 tunnel header (e.g., SRH, TC, Flowlabel). SRv6 Traffic Engineering mechanisms may be utilized to provide a forwarding sub-layer that is responsible for providing resource allocation and explicit routes. Other forwarding sub-layer solutions (e.g., policy based routing) are not precluded.

3.2. DetNet SRv6 Data Plane Scenarios

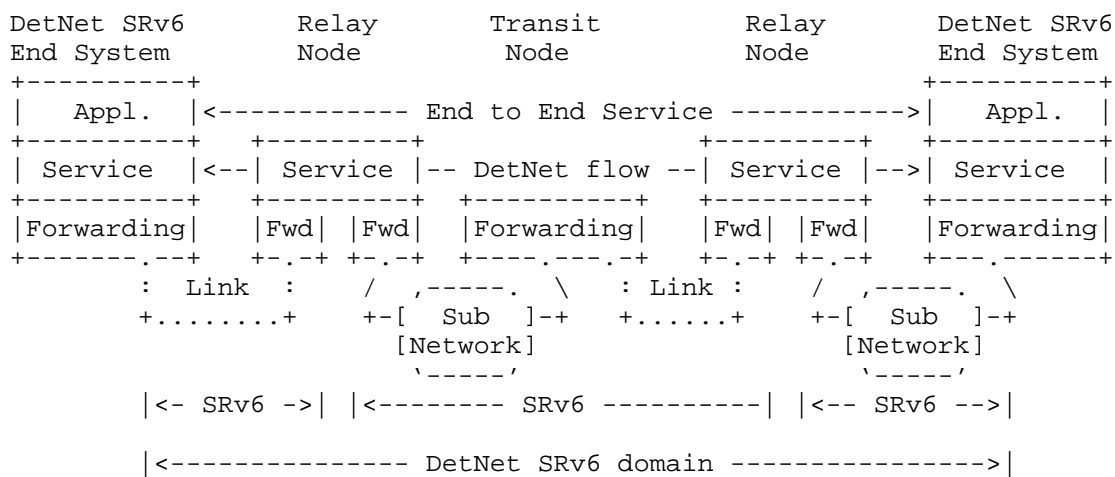


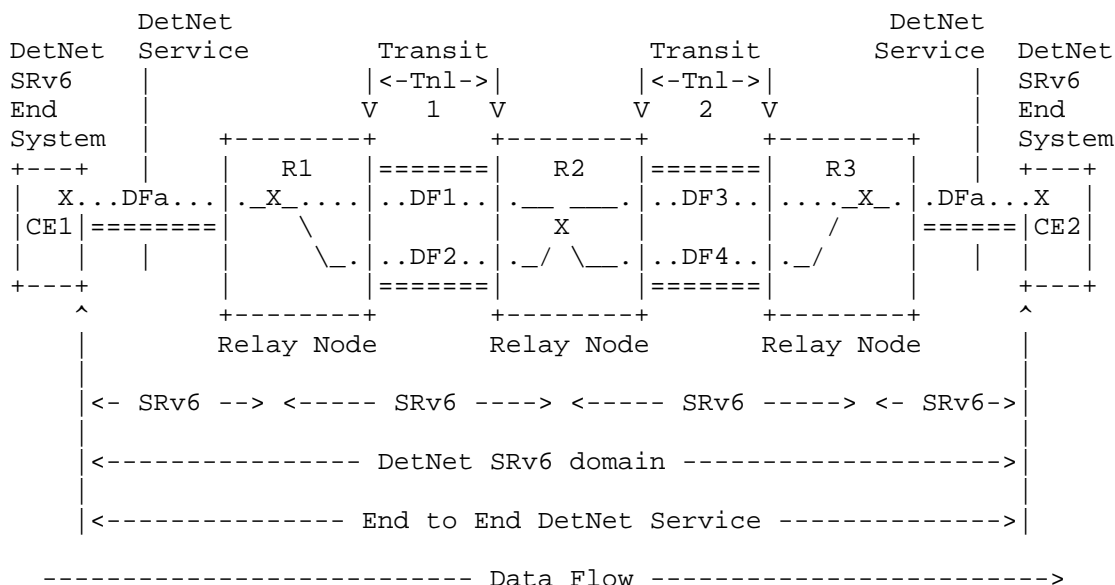
Figure 2: A DetNet SRv6 Network example

Figure 2 illustrates a hypothetical DetNet SRv6 network composed of DetNet aware SRv6 enabled end systems, operating over a DetNet aware SRv6 network. In this figure, SRv6 tunnels are used between the DetNet nodes implementing the service sub-layer.

DetNet end systems and relay nodes understand the particular needs of DetNet flows and provide both DetNet service and forwarding sub-layer functions. In this illustrative case, DetNet service-aware nodes creates/terminates the SRv6 tunnels and add/remove/process SeqNum and Flow-ID as needed.

In a DetNet SRv6 network, transit nodes may be DetNet service aware or may be DetNet unaware SRv6 Routers. In this latter case, such Routers would be unaware of the special requirements of the DetNet service sub-layer, but would still provide traffic engineering functions and the QoS capabilities needed to ensure that the SRv6 tunnels meet the service requirements of the carried DetNet flows.

Figure 3 illustrates how an end-to-end SRv6-based DetNet service is provided in more detail. In this figure, the customer end systems, CE1 and CE2, are able to send and receive SRv6 encapsulated DetNet flows, and R1, R2 and R3 are relay nodes in the middle of a DetNet network. The SRv6 tunnels between the Relay nodes may include transit nodes, which are not illustrated in the figure. The 'X' in the end systems, and relay nodes represents potential DetNet compound flow packet replication and elimination points. In this example, service protection is supported by utilizing at least two DetNet member flows and TE tunnels. For a unidirectional flow, R1 supports PRF and R3 supports PEF and POF.



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

Figure 3: SRv6 based DetNet Service

4. SRv6-Based DetNet Data Plane Solution

4.1. DetNet Over SRv6 Encapsulation Components

To carry DetNet over SRv6 the following is required:

1. A method of identifying the SRv6 payload type.
2. A method of identifying the DetNet flow(s) to the processing element.
3. A method of carrying the DetNet sequence number.
4. A suitable tunnel to deliver the packet to the egress node.
5. A method of carrying queuing and forwarding indication.

In this design the IPv6 NextHeader or the SRH refers to the payload type. The DetNet specific SID contains the Flow-ID and the SeqNum information. To simplify implementation and to maximize interoperability two sequence number sizes are supported: a 16 bit sequence number and a 28 bit sequence number.

The SRv6 encapsulation is used to forward the DetNet packet across the IPv6/SRv6 network between the DetNet Relay nodes and to indicate (e.g., by the SID(s), TC, Flowlabel) the required queue processing as well as the forwarding parameters.

4.2. SRv6 Data Plane Encapsulation

Figure 4 illustrates a DetNet data plane SRv6 encapsulation. The SRv6-based encapsulation of the DetNet flows is well suited for the scenarios described in [RFC8938].

The SRv6-based DetNet data plane encapsulation consists of:

- * DetNet specific SID containing flow identification and sequencing information for packet replication, duplicate elimination and ordering purposes.
- * Zero or more SID(s) used to direct the packet along the SRv6 tunnel to the next DetNet service sub-layer processing node.
- * The necessary data-link encapsulation is then applied prior to transmission over the physical media.

DetNet SRv6-based encapsulation example




Figure 4: Encapsulation of a DetNet App-Flow in an SRv6 PSN

Note: usage of SRv6 as the forwarding sub-layer between the DetNet Relay nodes is optional. In such cases the DetNet specific SID is set as the IPv6 destination address during the encapsulation.

4.2.1. DetNet specific SID

For PREOF processing, two arguments are needed:

1. Flow-ID: defines which DetNet flow the packet belongs to (what is used to determine which PREOF instance has to be used on a node). Its size is 20 bits for the DetNet MPLS data plane [RFC8986] and same size is appropriate for DetNet SRv6 data plane as well.
2. SeqNum: defines the sequencing information, it is created at the DetNet edge node (or by the first PRF node) and used by PEF/POF functionalities. Same sizes as for the DetNet MPLS data plane are defined for the SRv6 case: 0/16/28 bits [RFC8964].

The required size for these two arguments are maximum 48 bits. The explicit format (size of the three parts) of a DetNet specific SID is network addressing design specific. PREOF specific parameters are encoded as follows:

- * LOC: specifies the DetNet Relay node (same allocation rule applies as for any SRv6-enabled node).

- * FUNCT: a single value represents all PREOF instances of a DetNet Relay node.
- * ARG: Contains the Flow-ID and the SeqNum parameters.

The Flow-ID and SeqNum start at the MSB of the ARG. Unused bits (if any) MUST be set to zero (0).

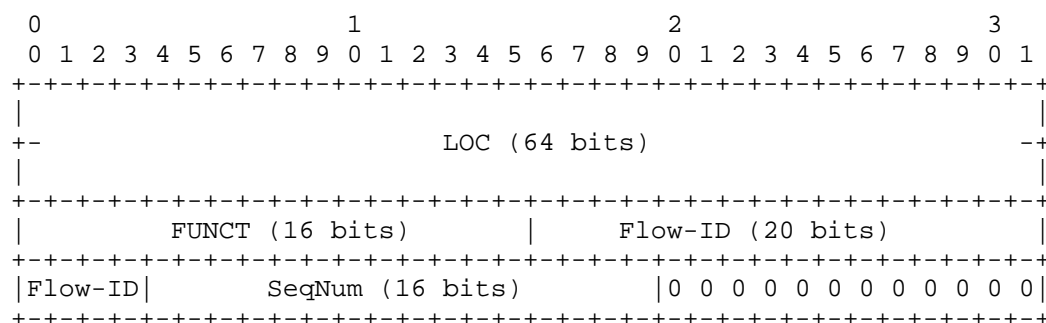


Figure 5: DetNet-specific SID example - LOC(64)+FUNCT(16)+ARG(Flow-ID(20)+SeqNum(16))

Note: if Function=PREOF, Argument=0 is also a meaningful value and does not refer to the lack of arguments.

The DetNet-specific SID MUST be the last segment in an SR Policy, if SRv6 TE tunnels are used between the DetNet Relay nodes.

4.2.2. Flow-ID

A DetNet flow at the DetNet service sub-layer is identified by a Flow-ID. SRv6-aware DetNet end systems and edge nodes, which are by definition SRv6 ingress and egress nodes, MUST add and remove a DetNet service-specific SID and the SRv6 tunnel information (SRH, if exists). Relay nodes MAY swap Flow-ID values when processing a DetNet flow, i.e., incoming and outgoing Flow-IDs of a DetNet flow can be different.

Flow-ID values MUST be provisioned per DetNet service via configuration, i.e., via the controller plane described in [RFC8938]. Note that Flow-IDs provide identification at the downstream DetNet service sub-layer receiver, not the sender. As such, Flow-IDs MUST be allocated by the entity that controls the service sub-layer receiving node's Flow-ID space. Because Flow-IDs are local to each node rather than being a global identifier within a domain, they MUST be advertised to their upstream DetNet service-aware peer nodes (i.e., a DetNet SRv6 End System or a DetNet Relay or Edge Node).

When receiving a DetNet SRv6 packet, an implementation MUST identify the DetNet service associated with the incoming packet based on the Flow-ID.

4.2.3. SeqNum

The sequence number characteristics MUST comply with the requirements specified in [RFC8964].

4.3. Service Sub-Layer Related Processing

DetNet SRv6 end systems, edge nodes and relay nodes may operate at the DetNet service sub-layer with understanding of DetNet services and their requirements. When operating at this layer such nodes can push, pop or swap Flow-IDs. Optionally, the SRH SID(s) can specify the SRv6 tunnel used by the DetNet flow.

Note, when PRF is supported, the same app-flow data will be sent over multiple outgoing DetNet member flows using e.g., forwarding sub-layer SRv6 tunnels. This means that implementation may send different sets of SRH SID(s) per DetNet member flow, each with a proper Flow-ID in the DetNet specific SID.

4.3.1. Packet Replication Function Processing

The Packet Replication Function (PRF) function MAY be supported by an implementation for outgoing DetNet flows. The use of the PRF for a particular DetNet service MUST be provisioned via configuration, i.e., via the controller plane described in [RFC8938].

When replication is configured, the same app-flow data will be sent over multiple outgoing DetNet member flows using forwarding sub-layer tunnels. An Flow-ID value MUST be configured per outgoing member flow. The same SeqNum field value MUST be used on all outgoing member flows for each replicated data packet.

4.3.2. Packet Elimination Function Processing

Implementations MAY support the Packet Elimination Function (PEF) for received DetNet SRv6 flows. When supported, use of the PEF for a particular DetNet service MUST be provisioned via configuration, i.e., via the controller plane described in [RFC8938].

After a DetNet service is identified for a received DetNet SRv6 packet, if PEF is configured for that DetNet service, duplicate (replicated) instances of a particular sequence number MUST be discarded. The specific mechanisms used for an implementation to identify which received packets are duplicates and which are new is an implementation choice. Note that per Section 4.2.3 the sequence number field length may be 16 or 28 bits, and the field value can wrap. PEF MUST NOT be used with DetNet flows configured with a sequence number field length of 0 bits.

An implementation MAY constrain the maximum number of sequence numbers that are tracked on either a platform-wide or per flow basis. Some implementations MAY support the provisioning of the maximum number of sequence numbers that are tracked on either a platform-wide or per flow basis.

4.3.3. Packet Ordering Function Processing

A function that is related to in-order delivery is the Packet Ordering Function (POF). Implementations MAY support POF. When supported, use of the POF for a particular DetNet service MUST be provisioned via configuration, i.e., via the controller plane described by [RFC8938].

Implementations MAY require that PEF and POF be used in combination. There is no requirement related to the order of execution of the Packet Elimination and Ordering Functions in an implementation.

After a DetNet service is identified for a received DetNet SRv6 packet, if POF is configured for that DetNet service, packets MUST be processed in the order indicated in the received SeqNum field, which may not be in the order the packets are received. As defined in Section 4.2.3 the sequence number field length may be 16 or 28 bits, is incremented by one (1) for each new data packet sent for a particular DetNet service, and the field value can wrap.

The specific mechanisms used for an implementation to identify the order of received packets is an implementation choice. Some possible implementations are described in [RFC9550]

4.4. Forwarding Sub-Layer Related Processing

Using SRv6 as a forwarding sub-layer between DetNet Relay nodes is optional. Other Traffic Engineering technologies (e.g., policy based routing) are NOT precluded.

If SRv6 is selected for TE, the SRv6 tunnel is used to provide connectivity between DetNet service sub-layer processing nodes. In such cases, the DetNet forwarding sub-layer provides explicit routes and allocated resources, and the SRv6 tunnel specific header (e.g., SRH, TC, Flowlabel) is used to map to each. Explicit routes are supported based on the SRH.

5. Security Considerations

DetNet PREOF related security considerations are described in section 3.3 of [RFC9055]. SRv6 Network Programming related security considerations are described in section 9 of [RFC8986]. There are no additional related security considerations originating from this document.

6. IANA Considerations


This document makes no IANA requests.

7. Acknowledgements

Authors extend their appreciation to Janos Farkas, Istvan Moldovan and Miklos Mate for their insightful comments and productive discussion that helped to improve the document.

8. Appendix A. Illustrations

This appendix shows how the described End.DPREOF mechanisms can be used in an SRv6 network.



N: non-SRv6 IPv6 node
 N: SRv6-capable node
 R: Node with Replication Function (PRF)
 E: Node with Elimination Function (PEF)
 L: Link between nodes

Figure 6: Example Topology

In the reference topology:

- * Nodes N3, R1, R2, E5 and E6 are SRv6-capable nodes.
- * Nodes R1, R2, E5 and E6 are PREOF nodes.
- * Nodes N4 is an IPv6 node that is not SRv6-capable.
- * Node j has an IPv6 loopback address 2001:db8:L:j::/128.
- * A SID at node j with locator block 2001:db8:K::/48 and function U is represented by 2001:db8:K:j:U::.
- * 2001:db8:K:j:P:: is explicitly allocated as the End.DPREOF SID at node j. For example, 2001:db8:K:2:P:: represents End.DPREOF at node R2.
- * 2001:db8:K:j:Xin:: is explicitly allocated as the End.X SID at node j towards neighbor node i via the nth link between nodes i and j. For example, 2001:db8:K:3:X51:: represents End.X at node N3 towards node E5 via link3 (the first link between nodes N3 and E5). Similarly, 2001:db8:K:3:X52:: represents the End.X at node N3 towards node E5 via link4 (the second link between nodes N3 and E5).

If the src node sends a packet to the dst node for which per packet redundancy is configured, then the nodes with PREOF functions provide the required replication or elimination functions. For instance, in the example in Figure 6:

- * Node src sends a UDP packet as follows: (2001:db8:src::1, 2001:db8:dst::1, NH = UDP)(UDP payload).
- * Node R1, which is an SRv6-capable PREOF node, identifies the flow the packet belongs to. As replication is configured for the given flow, R1 performs the replication action and intends to send the packet to the next PREOF nodes (E5 and R2). These nodes are reachable via SRv6, so R1 performs H.Encaps.PREOF(.Red) on the replicas with a path specific SRH. The argument part of the End.DPREOF SID involves the Flow-ID and the SeqNum. Specifically, one replica is sent on link-1 towards E5 (2001:db8:L:1::, 2001:db8:K:3:X51::) (2001:db8:K:5:P:arg::, 2001:db8:K:3:X51::, SL=1, NH = IPv6) (2001:db8:src::1, 2001:db8:dst::1, NH = UDP)(UDP payload) and the other replica is sent on link-2 towards R2 (2001:db8:L:1::, 2001:db8:K:2:P:arg::, NH = IPv6) (2001:db8:src::1, 2001:db8:dst::1, NH = UDP)(UDP payload).
- * Node N3, which is an SRv6-capable node, performs the standard SRH processing. Specifically, it executes the End.X behavior indicated by the 2001:db8:K:3:X51:: SID and forwards the packet on link3 to node E5.
- * Node N4, which is a non-SRv6-capable node, performs the standard IPv6 processing. Specifically, it forwards the UDP packet based on DA 2001:db8:K:2:P:arg:: in the IPv6 header towards node R2.
- * Node R2, which is an SRv6-capable PREOF node, identifies the packet as targeted to the local PREOF function. R2 performs the decapsulation and forwards the exposed payload and the ARG part to the PREOF functionality. The PREOF function identifies the flow the packet belongs to. As replication is configured for the given flow, R2 performs the replication action and intends to send the packet to the next PREOF nodes (E5 and E6). These nodes are reachable via SRv6, so R2 performs H.Encaps.PREOF(.Red) on the replicas with a path specific SRH. The argument part of the End.DPREOF SID involves the Flow-ID and the SeqNum. Specifically, one replica is sent on link-7 towards E5 (2001:db8:L:2::, 2001:db8:K:5:P:arg::, NH = IPv6) (2001:db8:src::1, 2001:db8:dst::1, NH = UDP)(UDP payload) and the other replica is sent on link-8 towards E6 (2001:db8:L:2::, 2001:db8:K:6:P:arg::, NH = IPv6) (2001:db8:src::1, 2001:db8:dst::1, NH = UDP)(UDP payload).

- * Node E5, which is an SRv6-capable PREOF node, identifies the packets as targeted to the local PREOF function. E5 performs the decapsulation and forwards the payload and the ARG part to the PREOF functionality. The PREOF function identifies the flow the packet belongs to. As elimination is configured for the given flow, the elimination action is performed on the packets received over Link3 and Link7. E5 intends to send the packet to the next PREOF node (E6), which is reachable via SRv6, so E6 performs H.Encaps.PREOF(.Red) with a path specific SRH. The argument part of the End.DPREOF SID involves the Flow-ID and the SeqNum. Specifically, the replica received first is sent on link-6 towards E6 (2001:db8:L:5::, 2001:db8:K:6:P:arg::, NH = IPv6) (2001:db8:src::1, 2001:db8:dst::1, NH = UDP)(UDP payload).
- * Node E6, which is an SRv6-capable PREOF node, identifies the packets as targeted to the local PREOF function. It performs the decapsulation and forwards the payload and the ARG part to the PREOF functionality. The PREOF function identifies the flow the packet belongs to. As elimination is configured for the given flow, the elimination action is performed on the packets received over Link6 and Link8. E6 is the last PREOF node, so after the PREOF function it send the UDP packet towards the destination. Specifically, the replica received first is sent towards the destination (2001:db8:src::1, 2001:db8:dst::1, NH = UDP)(UDP payload).

The example topology shown in Figure 6 is constructed to show the usage of PREOF-SID. Note that any of the links can be replaced with an SRv6 network segment. The above described principles are applicable to such more complex network topologies as well.

9. Normative References

- [I-D.varga-spring-preof-sid]
Varga, B. and F. Fejes, "Deterministic Networking specific SID", Work in Progress, Internet-Draft, draft-varga-spring-preof-sid-02, 3 July 2025, <<https://datatracker.ietf.org/doc/html/draft-varga-spring-preof-sid-02>>.
- [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>>.
- [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>>.

- [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>>.
- [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/info/rfc8938>>.
- [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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <<https://www.rfc-editor.org/info/rfc9055>>.
- [RFC9550] Varga, B., Ed., Farkas, J., Kehrner, S., and T. Heer, "Deterministic Networking (DetNet): Packet Ordering Function", RFC 9550, DOI 10.17487/RFC9550, March 2024, <<https://www.rfc-editor.org/info/rfc9550>>.

Authors' Addresses

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

Ferenc Fejes
Ericsson
Budapest
Magyar Tudosok krt. 11.
1117
Hungary
Email: ferenc.fejes@ericsson.com