

IETF-120



DetNet Data Plane: SRv6

draft-varga-detnet-srv6-data-plane
draft-varga-spring-preof-sid


Authors: Balázs Varga (A), Ferenc Fejes

2024-07-23

DetNet IP Data Plane History




- RFC8655: Deterministic Networking Architecture
 - DetNet provides a capability to carry specified unicast or multicast data flows for real-time applications with extremely low data loss rates and bounded latency within a network domain



- RFC8939: Deterministic Networking (DetNet) Data Plane: IP
 - Universal IP Data Plane for DetNet: both IPv4 and IPv6
 - Limitations: no support for DetNet Service sub-layer
- RFC9566: Deterministic Networking (DetNet): DetNet PREOF via MPLS over UDP/IP
 - Universal IP Data Plane for DetNet: both IPv4 and IPv6
 - Support for DetNet Service sub-layer (PREOF) using existing DetNet PW and MPLS-over-UDP technologies

Native support of PREOF for IPv6



- DetNet toolset (RFC8655):
 - (i) Resource allocation, (ii) Service protection, and (iii) Explicit routes
 - If service protection (i.e., PREOF) is used it is always combined with explicit routes
- IPv6:
 - Flexible & Extensible header format
 - Native PREOF support in IPv6 can mitigate the overhead of “PW + MPLS over UDP” solution
 - Task to be solved: adding Sequence Number to the IPv6 header
 - Option: Re-Use existing IPv6 header fields

Note on Explicit routes:

- Needed between DetNet Relay nodes (implementing service sub-layer (i.e., PREOF))
- Do not need to describe the whole PREOF graph, only the path to the next DetNet Relay node executing PREOF on the DetNet Flow

Proposals



Building blocks of the proposed DetNet SRv6 solution:

1. A DetNet-specific SID format, which contains the Sequence Number parameter used by PREOF instances.
2. A method to place Flow-ID and Sequence Number in the argument (ARG) part of the DetNet-specific SID.
3. A new set of SRv6 SID behaviors regarding the DetNet-specific SID
 - SR Endpoint Behavior
 - End.PREOF
 - SR Policy Headend Behaviors
 - H.Encaps.PREOF
 - H.Encaps.PREOF.Red
 - H.Encaps.PREOF.L2
 - H.Encaps.PREOF.L2.Red

DetNet-specific SID




- An SRv6 Endpoint behavior may require additional information for its processing (e.g., related to the flow or service). This information may be encoded in the ARG bits of the SID.
- For PREOF processing two arguments are needed
 - Flow-ID: defines which DetNet flow the packet belongs to (what is used to determine which PREOF instance has to be used), size: 20 bits for DetNet MPLS data plane [RFC8964]
 - Sequence Number: defines the sequencing information, it is used by PRF/PEF/POF, size: 0/16/28 bits [RFC8964]
 - **Required size: max. 48 bits**
- Recommendations for size of the SID parts
 - RFC8986: Section 3.2. – 64+16+**48 bits**

Proposed format of DetNet-specific SID



- Locator:
 - Specifies the node (allocation same as for any SID of the node)
- Function:
 - Single value is enough for all PREOF functions of the nodes
- Arguments
 - Contains the Flow-ID and SeqNum


Note: if Function=PREOF, Arg=0 is also meaningful



New Set of SRv6 Behaviors



- The DetNet-specific SID must be the last segment in an SR Policy!
- SR Endpoint Behavior
 - End.PREOF
- SR Policy Headend Behaviors
 - H.Encaps.PREOF
 - H.Encaps.PREOF.Red
 - H.Encaps.PREOF.L2
 - H.Encaps.PREOF.L2.Red



End.PREOF



When N receives a packet destined to S and S is a local End.DF.PPREOF SID, N does the following:

- S01. When an SRH is processed {
- S02. If (Segments Left != 0) {
- S03. Send an ICMP Parameter Problem to the Source Address
 with Code 0 (Erroneous header field encountered)
 and Pointer set to the Segments Left field,
 interrupt packet processing, and discard the packet.
- S04. }
- S05. Extract the ARG part of the SID
- S06. Remove the outer IPv6 header with all its extension headers
- S07. Forward the exposed payload and the ARG part to
 the PREOF functionality
- S08. }

H.Encaps.PREOF



Node N receives a packet P1=(A, B2) identified as a DetNet Flow. B2 is neither a local address nor SID of N. **It executes the DetNet Flow related PREOF functions**, resulting on one or more replica packets with related parameters (Flow-ID, SeqNum).

Node N is configured with an IPv6 address T (e.g., assigned to its loopback). N steers the egress of the DetNet packet P1' into an SRv6 Policy with a Source Address T and a segment list <S1, S2, S3>.

The H.Encaps.PREOF encapsulation behavior is defined as follows:

- S01. Push an IPv6 header with its own SRH
 - Set the ARG part of the LAST SID in the segment list**
- S02. Set outer IPv6 SA = T and outer IPv6 DA to the first SID in the segment list
- S03. Set outer Payload Length, Traffic Class, Hop Limit, and Flow Label fields
- S04. Set the outer Next Header value
- S05. Decrement inner IPv6 Hop Limit or IPv4 TTL
- S06. Submit the packet to the IPv6 module for transmission to S1

The H.Encaps.PREOF.Red behavior is an optimization of the H.Encaps.PREOF behavior.

H.Encap.PREOF.L2



The H.Encaps.PREOF.L2 encapsulation behavior is similar to H.Encaps.PREOF but sets an Ethernet specific outer Next Header and lacks the TTL/Hop Limit related action. H.Encaps.PREOF.L2 is defined as follows:

- S01. Push an IPv6 header with its own SRH
 - Set the ARG part of the LAST SID in the segment list
- S02. Set outer IPv6 SA = T and outer IPv6 DA to the first SID in the segment list
- S03. Set outer Payload Length, Traffic Class, Hop Limit, and Flow Label fields
- S04. Set the outer Next Header value
- S05. <N/A>
- S06. Submit the packet to the IPv6 module for transmission to S11

The H.Encaps.PREOF.L2.Red behavior is an optimization of the H.Encaps.PREOF.L2 behavior.

DetNet-specific SID related counters



- PREOF implementation may provide counters per DetNet flow.
- However, in order to be inline with the intention of RFC8986 (section 6. Counters), its recommendation may apply on the DetNet-specific SID and the above described set of SR Behaviors.
 - It means, a node supporting DetNet-specific SID should implement a pair of traffic counters (one for packets and one for bytes) per local SID entry, for traffic that matched that SID and was processed successfully (i.e., packets that generate ICMP Error Messages or are dropped are not counted).
 - The retrieval of these counters from MIB, NETCONF/YANG, or any other data structure is outside the scope of this document.