IETF 109
draft-srcompdt-spring-compression-requirement

Members:
Cheng Li, Chongfeng Xie, Darren Dukes, Peng Shaofu,
Ron Bonica, Sander Steffan, Wim Hendrickx

Presenter and Chair:
Weiqiang Cheng

The design team is to produce (rough) consensus (of the DT) outputs to the WG on two related topics:

1) What are the requirements for solutions to compressing segment routing information for use over IPv6;

2) An analysis of proposed approaches to compressing segment routing information for use over IPv6.

Long SRv6 SID lists = larger overhead

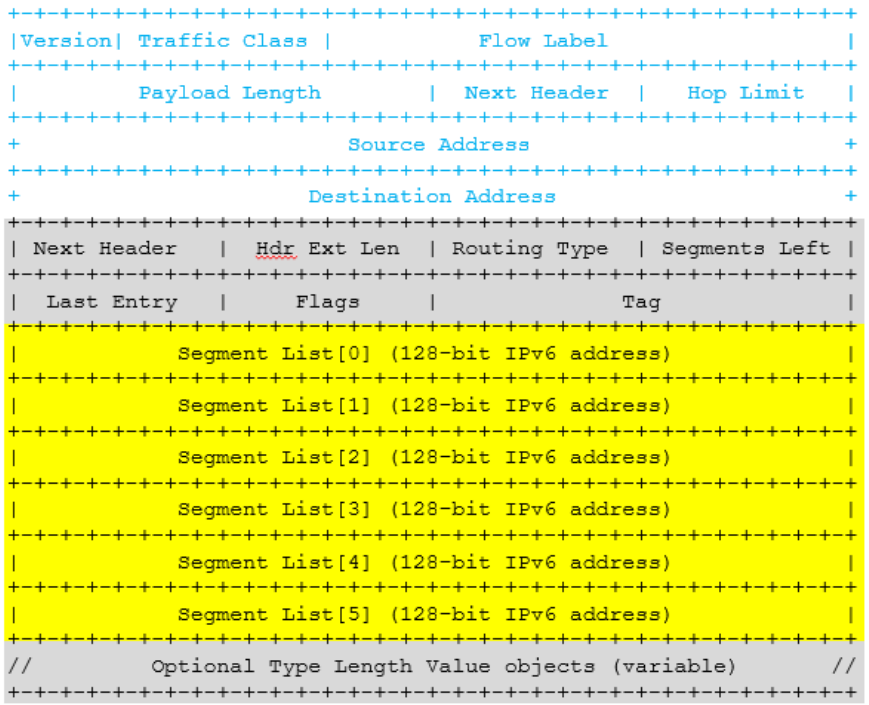

IPv6 Header
Fixed 40 octets

SRH 16 octets

+16 octets per SID

..
..
..
..
..
..
..

Strict path TE requires long SRv6 SID lists, therefore compressed solution is necessary.

Overview of Requirements draft

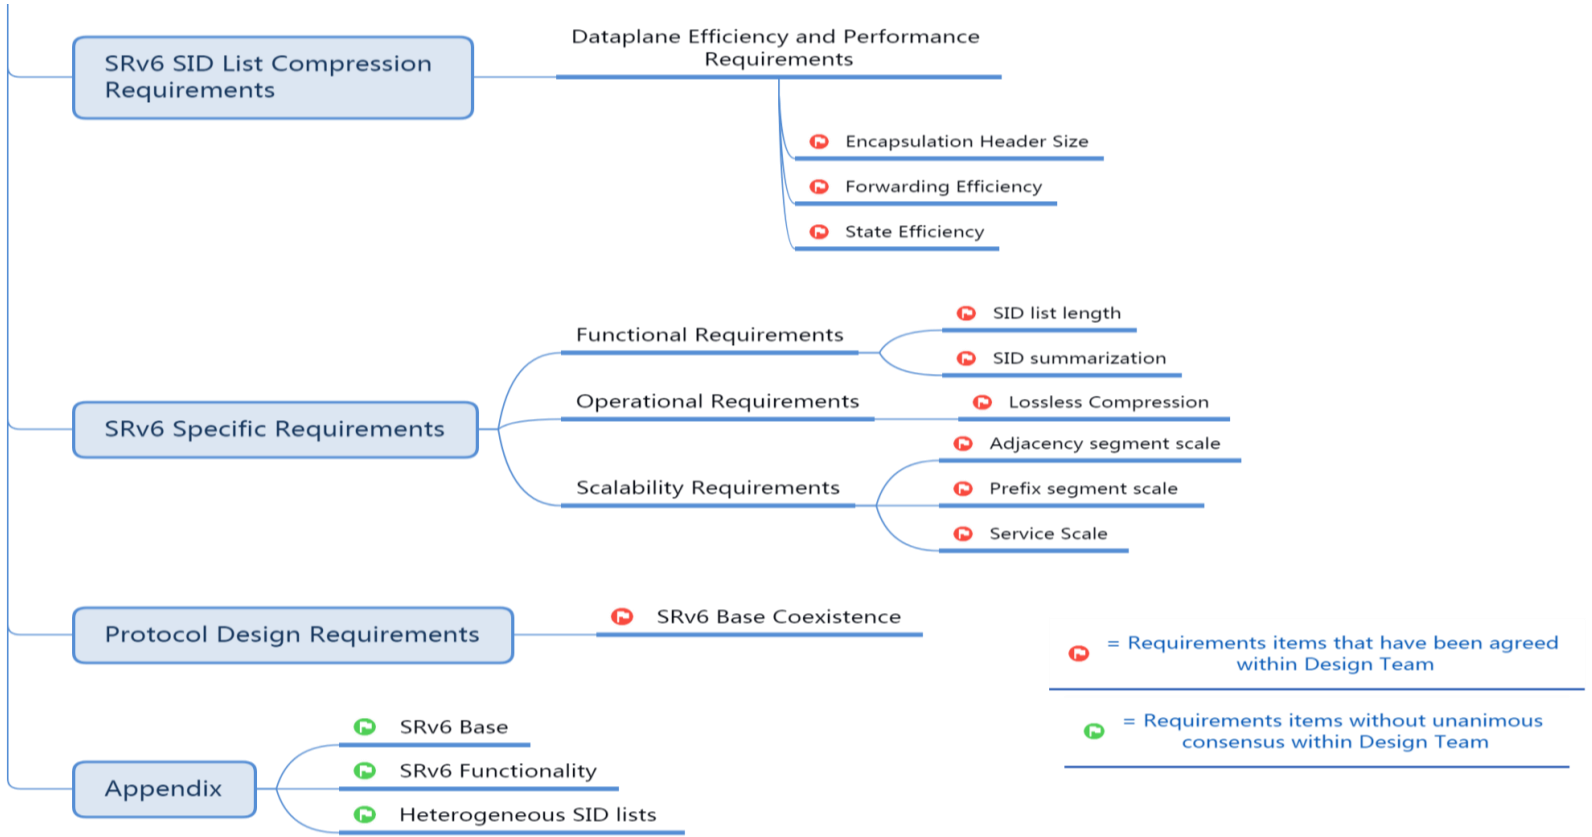

Requirements With Consensus

Encapsulation Header Size

Description: The compression solution MUST reduce the size of the SRv6 encapsulation header.

Notes: Given a specific segment list (realistic segment list specified during analysis)

- Records encapsulation size (E).
- Records encapsulation savings vs RFC8402 SRv6 (ES) .
- ES(proposal, segment list) = 1 - E(proposal, segment list)/ E(SRv6 base, segment list).

Forwarding Efficiency

Description: The compression solution SHOULD minimize the number of required hardware resources accessed to process a segment.

Notes: Given a specific segment list (realistic segment list specified during analysis)

- Records the number of headers parsed (D.PRS) during processing of the segment list.

- Records the number of FIB lookups (D.LKU) during processing of the segment list.

State Efficiency

Description: The compression solution SHOULD minimize the amount of additional forwarding state stored at a node

Notes: Given a set of node parameters (realistic parameters specified during analysis)

- Records the number of FIB entries (S) at a node.

SID List Length

Description: The compression mechanism MUST be able to represent SR paths that contain up to 16 segments.

Rationale: Strict TE paths require SID list lengths proportional to the diameter of the SR domain.

SID Summarization

Description: The solution MUST be compatible with segment summarization.

Rationale: Summarization of segments is a key benefit of SRv6 vs SR MPLS. In interdomain deployments, any node can reach any other node via a single prefix segment. Without summarization, border router SIDs must be leaked, and an additional global prefix segment is required for each domain border to be traversed.

Lossless Compression

Description: The segments of the compressed SID list MUST be equivalent to the original SID List.

For example, a strict path TE SID List is not compressed to a loose path TE SID list.

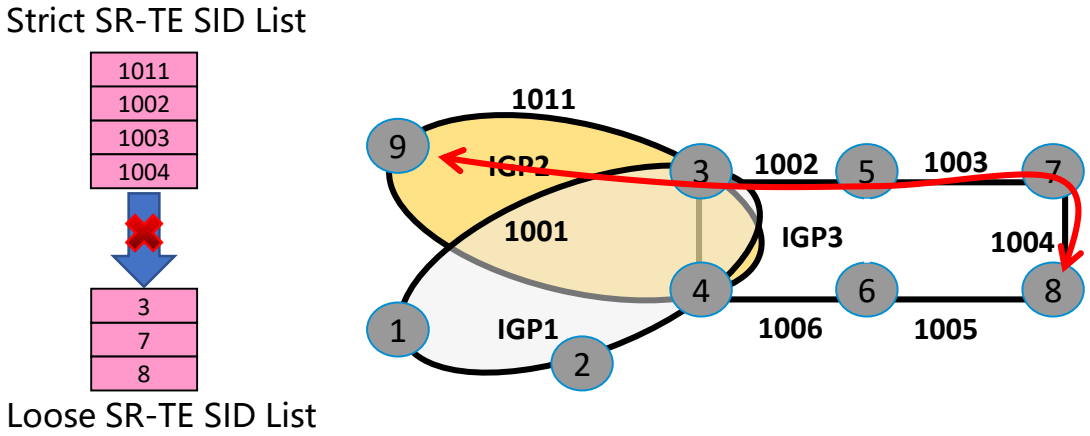

Adjacency Segment Scale

Description: The compression mechanism MUST be capable of representing 65000 adjacency segments per node

Rationale: Typically, network operators deploy networks with tens or hundreds of adjacency segments per node, but some network operators may deploy networks that use more adjacency segments per node.

Prefix Segment Scale

Description: The compression mechanism MUST be capable of representing 1 million prefix segments per SID numbering space.

Rationale: Typically, network operators deploy networks with thousands of prefix segments per SID numbering space, but some network operators may deploy networks that use more prefix segments per SID numbering space.

Service Scale

Description: The compression mechanism MUST be capable of representing 1 million services per node.

Rationale: Typically, network operators deploy networks with tens to hundreds of thousands of services per node, but some network operators may deploy networks that use more services per node.

SRv6 Base Coexistence

Description: The compression solution MUST support deployment in existing (RFC8402) SRv6 networks.

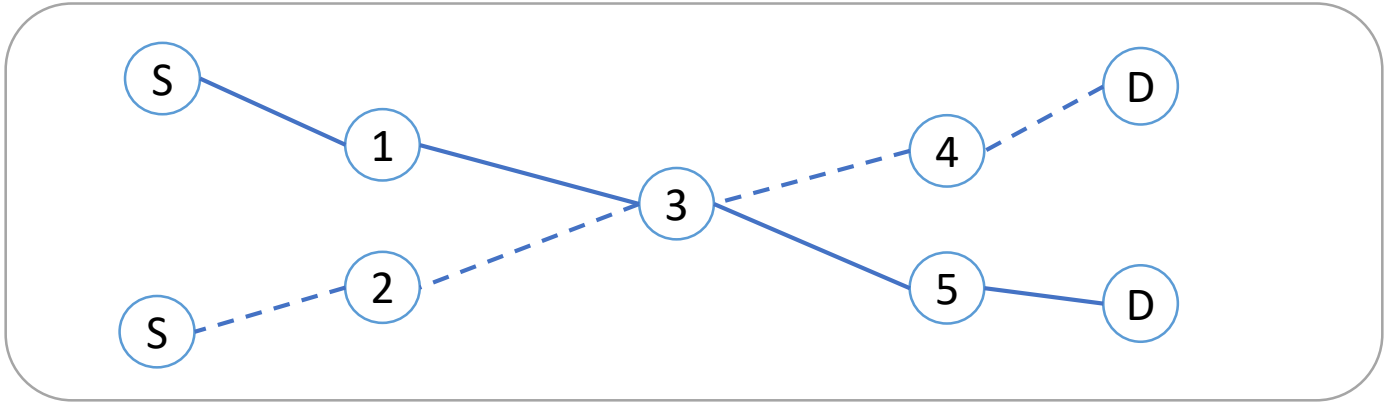

Requirements with rough but not unanimous consensus

Feedback appreciated

Introduction

It is a goal of the design team to identify solutions to SRv6 SID list compression that are based on the SRv6 standards. As such, this document provides requirements for SRv6 SID list compression solutions that utilize the existing SRv6 data plane and control plane.

It is also a goal of the design team to consider proposals that are not based on the SRv6 data plane and control plane. As such, this document includes requirements to evaluate whether a compression proposal provides all the functionality of SRv6 (section "SRv6 Functionality") in addition to satisfying compression specific requirements.

SRv6 Based

"**Description**: A solution to compress SRv6 SID Lists SHOULD be based on the SRv6 architecture, control plane and data plane"

*"**Rationale**: A compression solution built on existing IETF standards is preferable to creating new standards with equivalent functionality and performance."*

Here Preferred relationship

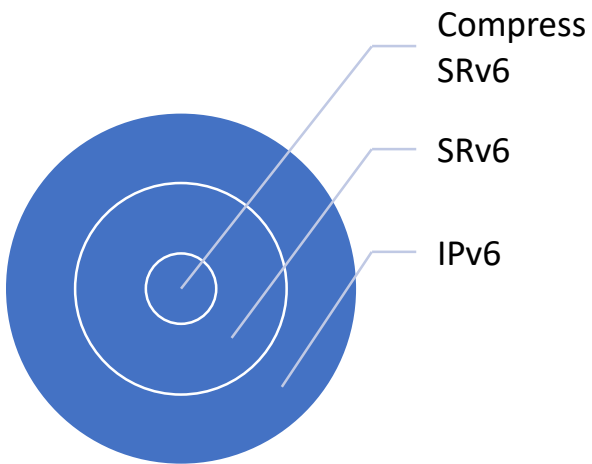

SRv6 Functionality

“Description: A solution to compress an SRv6 SID list MUST support the functionality of SRv6. This requirement and set of metrics, is meant to assess whether a proposal that is not fully SRv6 based, as evaluated in section "SRv6 Based", provides equivalent functionality to SRv6 Such a proposal may utilize different control plane and/or data planes.”

*“**Rationale**: Operators require SRv6 functionality. Evaluating the extent to which a proposal supports SRv6 functionality is important for operators and implementors to understand the impact on network operations.”*

Heterogeneous SID lists

“**Description**: The compression solution SHOULD support a combination of compressed and non-compressed segments in a single path.”

*“**Rationale**: Support of SID lists with compressed and non-compressed SIDs reduces encapsulation size when not all SRv6 nodes deploy the compression solution or 128-bit SIDs are required.”*

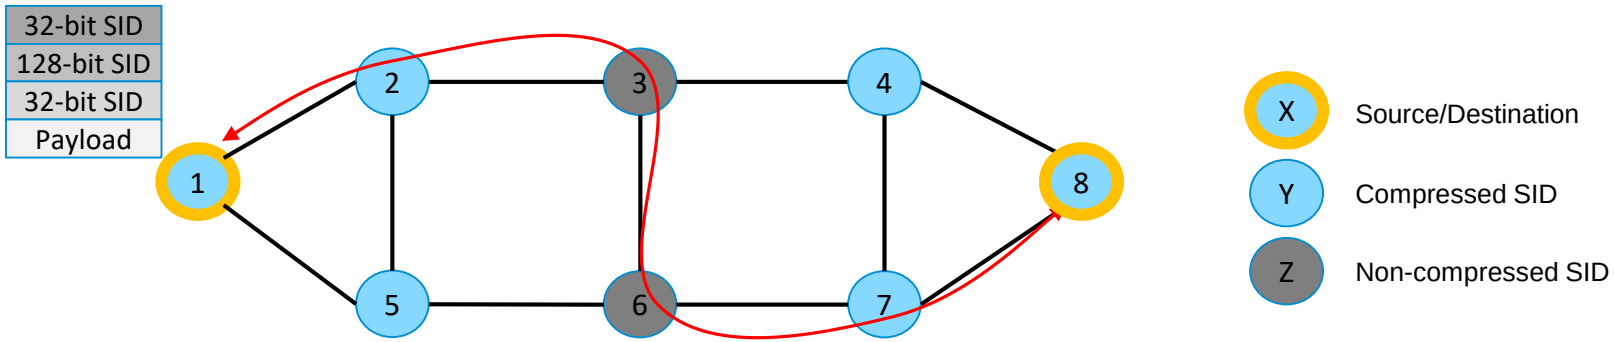

Next Steps

Requirements Completion

1 week:
gather feedback on requirements

End of December:
Decide on rough consensus requirements
Publish a new revision

Analysis Phase

January:
The SRCOMP team will provide a list of proposals to analyze

IETF 110:
Analysis draft revision 00

Questions?