

IETF 111

SRv6 Compression Design Team

Status Report

draft-srcompdt-spring-compression-requirement
draft-srcompdt-spring-compression-analysis

Members:

Ron Bonica, Darren Dukes, Wim Hendrickx, Cheng Li, Peng Shaofu, Chongfeng Xie

Presenter and Chair:

Weiqiang Cheng

The design team's goal 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;

<https://datatracker.ietf.org/doc/draft-srcompdt-spring-compression-requirement/>

Complete and Stable revision-07

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

<https://datatracker.ietf.org/doc/draft-srcompdt-spring-compression-analysis/>

Complete and Stable revision-02

Update to the requirement draft

Status of Requirements draft

- ❑ The latest revision is -07, which included all the requirements we've received.
- ❑ Compared to -06, we just deleted section "5.2 PS or BCP Compliance"

3. SRv6 SID List Compression Requirements

3.1. Dataplane Efficiency and Performance Requirements

3.1.1. Encapsulation Header Size

3.1.2. Forwarding Efficiency

3.1.3. State Efficiency

4. SRv6 Specific Requirements

4.1. SRv6 Based

4.2. Functional Requirements

4.2.1. SRv6 Functionality

4.2.2. Heterogeneous SID lists

4.2.3. SID list length

4.2.4. SID summarization

4.3. Operational Requirements

4.3.1. Lossless Compression

4.3.2. Preservation of non-routing information

4.3.3. Address Planning

4.4. Scalability Requirements

4.4.1. Adjacency segment scale

4.4.2. Prefix segment scale

4.4.3. Service Scale

4.4.4. Compression Levels

5. Protocol Design Requirements

5.1. SRv6 Base Coexistence

5.2. PS or BCP Compliance --deleted

6. Security Requirements

6.1. Security Mechanisms

6.2. SR Domain Protection

Appendix A. Proposed Requirements

A.1. IPv6 Based

A.2. Point to Multipoint

A.3. Parsability

Update of Requirements after IETF110

Delete PS or BCP Compliance.

The members agreed that this requirement is a normal IETF rule that should be satisfied by default.

5.2. PS or BCP Compliance

Description: The compression mechanism SHOULD comply with any proposed standard or BCP. If it does not comply with any PS or BCP it SHOULD update the related document.

Rationale: Compliance with existing standards makes the internet more robust.

Metric: A solution that complies with PS or BCP, or updates PS or BCP satisfies this requirement. The PS or BCP document may be updated at any time before the publication of the document specifying the compression mechanism.

6. Security Requirements

6. Security Requirements

Update to the analysis draft

status of analysis draft

- ❑ The latest revision is -02, which covered all the requirements analysis .
- ❑ Compared to -00, we just updated all the sections and now we've post the stable text

2. SRv6 Compression Requirements

2.1. Encapsulation Header Size

2.1.1. Reference Scenarios

2.2. Forwarding Efficiency

2.2.1. Headers Parsed

2.2.2. Lookups Performed (LKU)

2.3. State Efficiency

3. SRv6 Specific Requirements

3.1. SRv6 Based

3.2. Functional Requirements

3.2.1. SRv6 Functionality

3.2.2. Heterogeneous SID Lists

3.2.3. SID List Length

3.2.4. SID Summarization

3.3. Operational Requirements

3.3.1. Lossless Compression

3.3.2. Preservation of non-routing information

3.3.3. Address Planning .

3.4. Scalability Requirements

3.4.1. Compression Levels

4. Protocol Design Requirements

4.1. SRv6 Base Coexistence

5. Security Requirements

5.1. Security Mechanisms

5.2. SR Domain Protection

6. Conclusions

7. Normative References

Appendix A. Encapsulation analysis

A.1. CRH note

A.2. Analysis results

Authors' Addresses

Introduction

An analysis of each mechanism against the requirements.

“The following mechanisms are proposed to compress the SRv6 SID list.”

CSID	<i>I-D.filsfilscheng-spring-srv6-srh-comp-sl-enc</i>	Describes two new SRv6 SIDs, a combination of SIDs from [I-D.filsfils-spring-net-pgm-extension-srv6-usid] and [I-D.cl-spring-generalized-srv6-for-cmpr]
CRH	<i>I-D.bonica-6man-comp-rtg-hdr</i>	Requires two new routing header types and a label mapping technique
VSID	<i>I-D.dekraene-spring-srv6-vlsid</i>	Defines a set of SID behaviors to access smaller SIDs within the SR header
UID	<i>I-D.mirsky-6man-unified-id-sr</i>	Extends the SRH to carry MPLS labels or IPv4 addresses

CSID

A compressed SRv6 Segment List Encoding in the SRH.

- Does not require any SRH data plane change.
- Does not require SRv6 control plane change.
- Leverages the SRv6 Network Programming model.

Define two new SID flavors:

- NEXT-C-SID
- REPLACE-C-SID

Merges SID behaviors from uSID (draft-filsfils-spring-net-pgm-extension-srv6-usid) and GSID (draft-cl-spring-generalized-srv6-for-cmpr)

CRH

Two new IPv6 Routing Headers (CRH-16 and CRH-32)

- Next Header, Ext Hdr Len, Routing Type, Segments Left
- SID List (16 or 32-bit SIDs)

Each SID maps to a CRH-FIB entry

- IPv6 address or SRv6 SID
- Topological function plus optional arguments
- Service function plus optional arguments
- Flags

No change to IPv6 forwarding plane or addressing model

Minimal change to SRv6 control plane

vSID

Generalize the SRH for any size of SIDs (≤ 128 bits)

- 128-bit SIDs becomes a specific case
- Does not require any SRv6 control plane change.
- Leverages the SRv6 Network Programming model.

Defines one new SID flavor.

Builds on a common SRv6 locator prefix:

- $SID := \text{prefix} + \text{vSID}$
- Encodes only the vSID in the SRH. Not the redundant prefix.
- Everything else uses the regular 128-bits SID

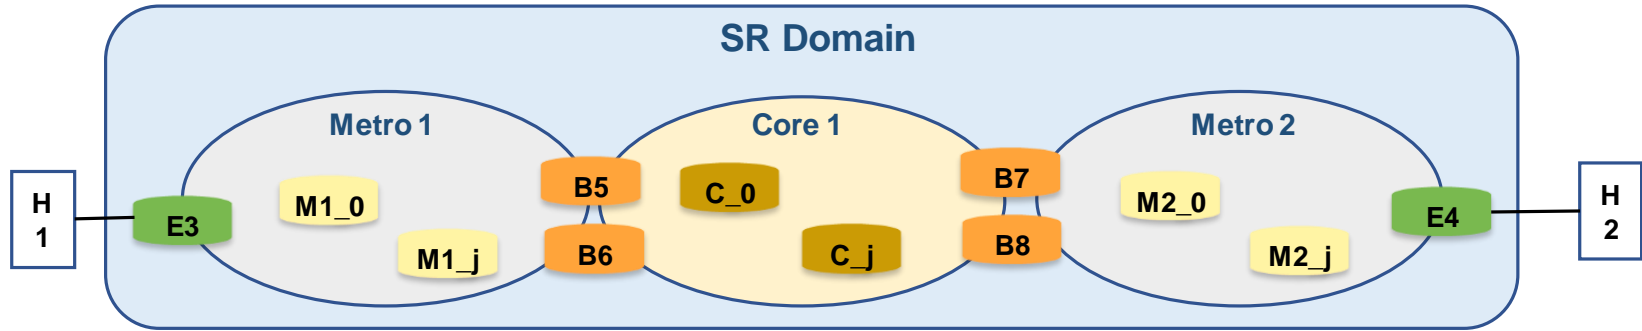
UID

A compressed SRv6 Segment List Encoding in the SRH (suggested) or other type of Routing Header.

- Introduce UET Flags to unify traditional SRv6 SID and U-SID forwarding behaviors, no compatibility issues.
 - 00: classical 128-bits IPv6 address
 - 01: 32-bits truncated piece of IPv6 address
 - 10: 32-bits index (MPLS label suggested)
 - 11: 16-bits truncated piece of IPv6 address
- Support MAPPING and STICHING mode, The former is used for disorderly IP address planning scenarios, while the latter is used for scenarios with common prefix.
- For MAPPING mode, index to IPv6 address mapping need to be advertised, MPLS prefix-SID can be reused; for STICHING mode, UET-32/16/etc flavors need to be advertised with the endpoint behavior of SRv6 SID.

Base-line topo for analysis

An SR domain consisting of 3 sub-domains is shown to illustrate the scenarios associated with encapsulation header size, forwarding efficiency and state efficiency.



H1、H2 hosts outside the SR domain

E3、E4 SR domain edge routers

B5、B6 border routers between the Metro 1 and Core

B7、B8 border routers between the Metro 2 and Core

Metro 1、Core、Metro 2 sub-domains with independent IGP instances

M1_1..M1_i routers in Metro 1

C_1..C_j routers in Core

M2_1..M2_k routers in Metro 2

Analysis Results: Conclusions summary-1/4

Encapsulation Header Size

- All proposals meet the requirement to reduce the size of the SRv6 encapsulation header. Variances between proposals are negligible.

Forwarding Efficiency

- Overall, the CSID parses the fewest headers. When per packet state is processed per segment, CSID, VSID and UIDSR proposals may include it in the routing header, CRH may include it in a destination option preceding the CRH.
- CSID, VSID, and UIDSR require a single lookup to process an adjacency or VPN segment. CRH always requires 2 lookups for VPN segments, and 2 and sometimes 3 lookups for adjacency segments. All proposals require two lookups to process a prefix segment and the next segment.

State Efficiency

- CSID, VSID and UIDSR minimize forwarding state stored at a node. CRH moves per segment state from the packet to the FIB.

Analysis Results: Conclusions summary-2/4

SRv6 Based

- CSID is SRv6 based, requiring no updates to existing SRv6 standards, VSID and UIDSR require updates. CRH is not strictly based on SRv6 but is able to provide equivalent functionality.

SRv6 Functionality

- CSID supports SRv6 functionality. CRH VSID and UID support SRv6 functionality or equivalent with some new specifications. Heterogeneous SID lists o All proposals support heterogeneous SID lists. CSID and UIDSR support heterogeneous SID lists in the SRH, while CRH and VSID require installation of binding SIDs at midpoint nodes

Heterogeneous SID lists

- All proposals support heterogeneous SID lists. CSID and UIDSR support heterogeneous SID lists in the SRH, while CRH and VSID require installation of binding SIDs at midpoint nodes.

SID List Length

- All proposals support segment lists of at least 16 segments.

Analysis Results: Conclusions summary-3/4

SID Summarization

- VSID, CSID and UIDSID support segment summarization, CRH does not.

Operational Requirements

- All proposals provide lossless compression.
- All proposals preserve non-routing information.
- All proposals support flexible IPv6 planning.

Scalability Requirements

- All proposals meet scalability requirements.
- All proposals support 16-bit and 32-bit SID variants.

Protocol Design Requirements

- All proposals can be deployed simultaneously with the SRv6 base solution.

Security Requirements

- All proposals address security issues they may introduce with existing security mechanisms.
- All proposals protect SIDs within the SR domain.

Analysis Results: Conclusions summary-4/4

	CRH	CSID	UIDSR	VSID
Encapsulation Header Size				
Forwarding Efficiency				
State Efficiency				
SRv6 Based				
SRv6 Functionality				
Heterogeneous SID Lists				
SID List Length				
SID Summarization				
Operational Requirements				
Scalability Requirements				
Protocol Design Requirements				
Security Requirements				

Legend

Requirement is met
Shade shows relative degree based on conclusion text.

Most Least

Requirement is not met

All requirements are not equally important. The working group must decide which are more significant.

Next Steps

- **WG review**
- **WG adoption**
- **adopt related solution drafts?**

Questions?

Analysis Results: Encapsulation Header Size

16-bit SIDs	CSID	CRH	CRH+TPF	VSID	UIDSR
Average ES	54.3%	54.2%	50.4%	51.6%	49.2%

Table 1: Average ES, 16-bit SIDs, 48B.0-15T.V

21-bit SIDs	CSID	CRH	CRH+TPF	VSID	UIDSR
Average ES	42.5%	45.5%	43.25	45.5%	42.5%

Table 2: Average ES, 32-bit SIDs, 48B.0-15T.V

E and ES are also evaluated for 32bit and 64bit SRv6 block sizes. The CSID 16-bit ES averages 57.4% for 32-bit blocks and 49.9% for 64-bit blocks, other proposals are unchanged.

Conclusion: All proposals meet the requirement to reduce the size of the SRv6 encapsulation header. Variances between proposals are negligible.

Analysis Results: Forwarding Efficiency.Headers Parsed-1/2

16-bit SIDs	CSID	CRH	CRH+TPF	VSID	UIDSR
PRS(48B.0T).V	IPv6	IPv6	IPv6	IPv6	IPv6
PRS(48B.1-4T).V	IPv6	IPv6 CRH	IPv6 CRH	IPv6 SRH	IPv6 SRH
PRS(48B.5-15T).V	IPv6 SRH	IPv6 CRH	IPv6 CRH	IPv6 SRH	IPv6 SRH

Table 3: Headers parsed on non-decapsulating SR segment endpoint nodes, 16-bit SIDs, 48B.0-15T.V

16-bit SIDs	CSID	CRH	CRH+TPF	VSID	UIDSR
PRS(48B.0T).V	IPv6	IPv6	IPv6	IPv6	IPv6
PRS(48B.1-4T).V	IPv6	IPv6 CRH	IPv6 CRH TPF	IPv6 SRH	IPv6 SRH
PRS(48B.5-15T).V	IPv6 SRH	IPv6 CRH	IPv6 CRH TPF	IPv6 SRH	IPv6 SRH

Table 4: Headers parsed on decapsulating SR segment endpoint nodes, 16-bit SIDs, 48B.0-15T.V

Analysis Results: Forwarding Efficiency.Headers Parsed-2/2

32-bit SIDs	CSID	CRH	CRH+TPF	VSID	UIDSR
PRS(48B.0T).V	IPv6	IPv6	IPv6	IPv6	IPv6
PRS(48B.1-15T).V	IPv6 SRH	IPv6 CRH	IPv6 CRH	IPv6 SRH	IPv6 SRH

Table 5: Headers parsed on non-decapsulating SR segment endpoint nodes, 32-bit SIDs, 48B.0-15T.V

32-bit SIDs	CSID	CRH	CRH+TPF	VSID	UIDSR
PRS(48B.0T).V	IPv6	IPv6	IPv6	IPv6	IPv6
PRS(48B.1-15T).V	IPv6 SRH	IPv6 CRH	IPv6 CRH TPF	IPv6 SRH	IPv6 SRH

Table 6: Headers parsed on decapsulating SR segment endpoint nodes, 32-bit SIDs, 48B.0-15T.V

Conclusion: Overall, the CSID parses the fewest headers. When per packet state is processed per segment, CSID, VSID and UIDSR proposals may include it in the routing header, CRH may include it in a destination option preceding the CRH.

Analysis Results: Forwarding Efficiency. Lookups Performed

	CSID	CRH	VSID	UIDSR
Adj and VPN Segment	LPM(a)	LPM(a) EM(b) EM(b,c)	LPM(a)	LPM(a)
Prefix Segment	LPM(a) LPM(d)	LPM(a) EM(b)	LPM(a) LPM(d)	LPM(a) LPM(d)

- [a] On active SID, appearing in the IPv6 Destination address
- [b] On SID in CRH header
- [c] This lookup is required only when the IPv6 next hop node is not non-CRH aware
- [d] On next SID, appearing in the IPv6 destination address

Conclusion: CSID, VSID, and UIDSR require a single lookup to process an adjacency or VPN segment. CRH always requires 2 lookups for VPN segments, and 2 and sometimes 3 lookups for adjacency segments. All proposals require two lookups to process a prefix segment and the next segment.

Analysis Results: State Efficiency

16-bit and 32-bit SIDs	CSID	CRH	VSID	UIDSR
S(N1000,I2,A100,D10)	102 A.1:112 A.2:102	2100 B.1:3300 C.1:2120 C.2:10120	102 A.1:112 A.2:102	102 A.1:112 A.2:102
S(V1000)	1000	2000	1000	1000

- N: the number of SRv6 nodes in the sub-domain
- I: the number of IGP algorithms configured
- A: the number of local adjacency SIDs at a node
- D: the number of attached SR sub-domains at a border node
- V: the number of VPN services at edge nodes

For a sub-domain consisting of:

- 1000 SRv6 nodes (N=1000) with some non-SRv6 nodes
- 2 IGP algorithms (I=2)
- 100 adjacencies per SRv6 node (A=100)
- up to 10 attached sub-domains per border node (D=10)
- 1000 VPN service segments per edge (V=1000)

Conclusion: CSID, VSID and UIDSR minimize forwarding state stored at a node. CRH moves per segment state from the packet to the FIB.

Analysis Results:SRv6 Based

	CSID	CRH	VSID	UIDSR
U.RFC8402	Y	Y update required for SRv6 data plane	Y	Y
U.RFC8754	Y	N	Y update SL	Y update for flags and SL
U.PGM	Y	N	Y update SID behaviors	Y
U.IGP	Y	N	Y	Y additional ext
U.BGP	Y	N	Y	Y
U.POL	Y	N	Y	Y
U.BLS	Y	N	Y	Y additional ext
U.SVC	Y	N	Y	Y
U.ALG	Y	Y adds IP flexalgo	Y	Y
U.OAM	Y	N	Y	Y

Conclusion: CSID is SRv6 based, requiring no updates to existing SRv6 standards, VSID and UIDSR require updates. CRH is not strictly based on SRv6 but is able to provide equivalent functionality.

Analysis Results:SRv6 Functionality-1/2

	CSID	CRH	VSID	UIDSR
F.SID	RFC8402	CRH	RFC8402	RFC8402 1
F.Scope	RFC8402	CRH	RFC8402	RFC8402 1
F.PFX	RFC8402, RFC8986 CSID adds End Flavor	CRH	RFC8402, RFC8986 CSID updates End behavior	RFC8402, RFC8986 With new flavor 1
F.ADJ	RFC8402, RFC8986 CSID adds End.X Flavor	CRH	RFC8402, RFC8986 CSID updates End.X behavior	RFC8402, RFC8986 With new flavor 1
F.BIND	RFC8402, RFC8986	CRH	RFC8402, RFC8986 VSID updates End.B behavior	RFC8402, RFC8986 With new flavor 1
F.PEER	RFC8402, RFC8986 CSID adds End.X Flavor	CRH	RFC8402, RFC8986 VSID updates End.X behavior	RFC8402, RFC8986 With new flavor 1,2
F.SVC	RFC8986	CRH	RFC8402, RFC8986 VSID updates service seg behavior	RFC8986 1
F.ALG	SRv6 Flexaglo	IPv6 Flexaglo	SRv6 Flexaglo	SRv6 Flexaglo

Analysis Results:SRv6 Functionality-2/2

	CSID	CRH	VSID	UIDSR
F.TILFA	SRv6 TILFA	SRv6 TILFA	SRv6 TILFA	SRv6 TILFA 3
F.SEC	RFC8754	CRH-EXT	RFC8754	RFC8754
F.IGP	SRv6 ext	CRH-EXT	SRv6 ext	SRv6 ext 1,4
F.BGP	SRv6 BGPSVC	SRv6 BGPSVC	SRv6 BGPSVC	SRv6 BGPSVC
F.POL	SRV6SRPOL	UPDATES SRV6SRPOL	SRV6SRPOL	SRV6SRPOL
F.BLS	SRV6BGPLS	EXT required	SRV6BGPLS add VSID length	SRV6BGPLS 5
F.SFC	SRv6SVCP	CRH	SRv6SVCP	SRv6SVCP 1
F.PING	SRv6OAM	CRH	SRv6OAM	SRv6OAM

Conclusion: CSID supports SRv6 functionality. CRH VSID and UID support SRv6 functionality or equivalent with some new specifications.

Analysis Results: Heterogeneous SID Lists

	CSID	CRH	VSID	UIDSR
Heterogeneous SID Lists	Yes	Yes	Yes	Yes

Conclusion: All proposals support heterogeneous SID lists. CSID and UIDSR support heterogeneous SID lists in the SRH, while CRH and VSID require installation of binding SIDs at midpoint nodes.

Analysis Results: SID List Length

	CSID	CRH	VSID	UIDSR
16 Segments	Yes	Yes	Yes	Yes

Conclusion: All proposals support segment lists of at least 16 segments.

Analysis Results: SID Summarization

	CSID	CRH	VSID	UIDSR
SID Summarization	Yes	No	Yes	Yes

Conclusion: CSID, VSID and UIDSR support segment summarization, CRH does not.

Analysis Results: Lossless Compression

	CSID	CRH	VSID	UIDSR
Lossless Compression	Yes	Yes	Yes	Yes

Conclusion: All proposals provide lossless compression.

Analysis Results: Preservation of non-routing information

	CSID	CRH	VSID	UIDSR
Preservation of non-routing information	Yes	Yes	Yes	Yes

Conclusion: All proposals preserve of non-routing information

Analysis Results: Address Planning

	CSID	CRH	VSID	UIDSR
Flexible Address Planning	Yes	Yes	Yes	Yes

Conclusion: All proposals support flexible IPv6 planning.

Analysis Results: Scalability Requirements

	CSID	CRH	VSID	UIDSR
Adjacency Segment Scale 65000	Yes	Yes	Yes	Yes
Prefix Segment Scale 1000000	Yes	Yes	Yes	Yes
Service Scale 1000000	Yes	Yes	Yes	Yes

Conclusion: All proposals meet scalability requirements.

Analysis Results: Compression Levels

	CSID	CRH	VSID	UIDSR
Multiple compression Levels	Yes	Yes	Yes	Yes

Conclusion: All proposals support 16-bit and 32-bit SID variants.

Analysis Results: SRv6 Base Coexistence

	CSID	CRH	VSID	UIDSR
SRv6 Base Coexistence	Yes	Yes	Yes	Yes

Conclusion: All proposals can be deployed simultaneously with the SRv6 base solution.

Analysis Results: Security Mechanisms

	CSID	CRH	VSID	UIDSR
Security Mechanisms	Yes	Yes	Yes	Yes

Conclusion: All proposals address security issues they may introduce with existing security mechanisms.

Analysis Results: SR Domain Protection

	CSID	CRH	VSID	UIDSR
SR Domain Protection	Yes	Yes	Yes	Yes

Conclusion: All proposals protect SIDs within the SR domain.