

Generalized SRv6 for Compression

Cheng Li
c.l@Huawei.com

Introduction



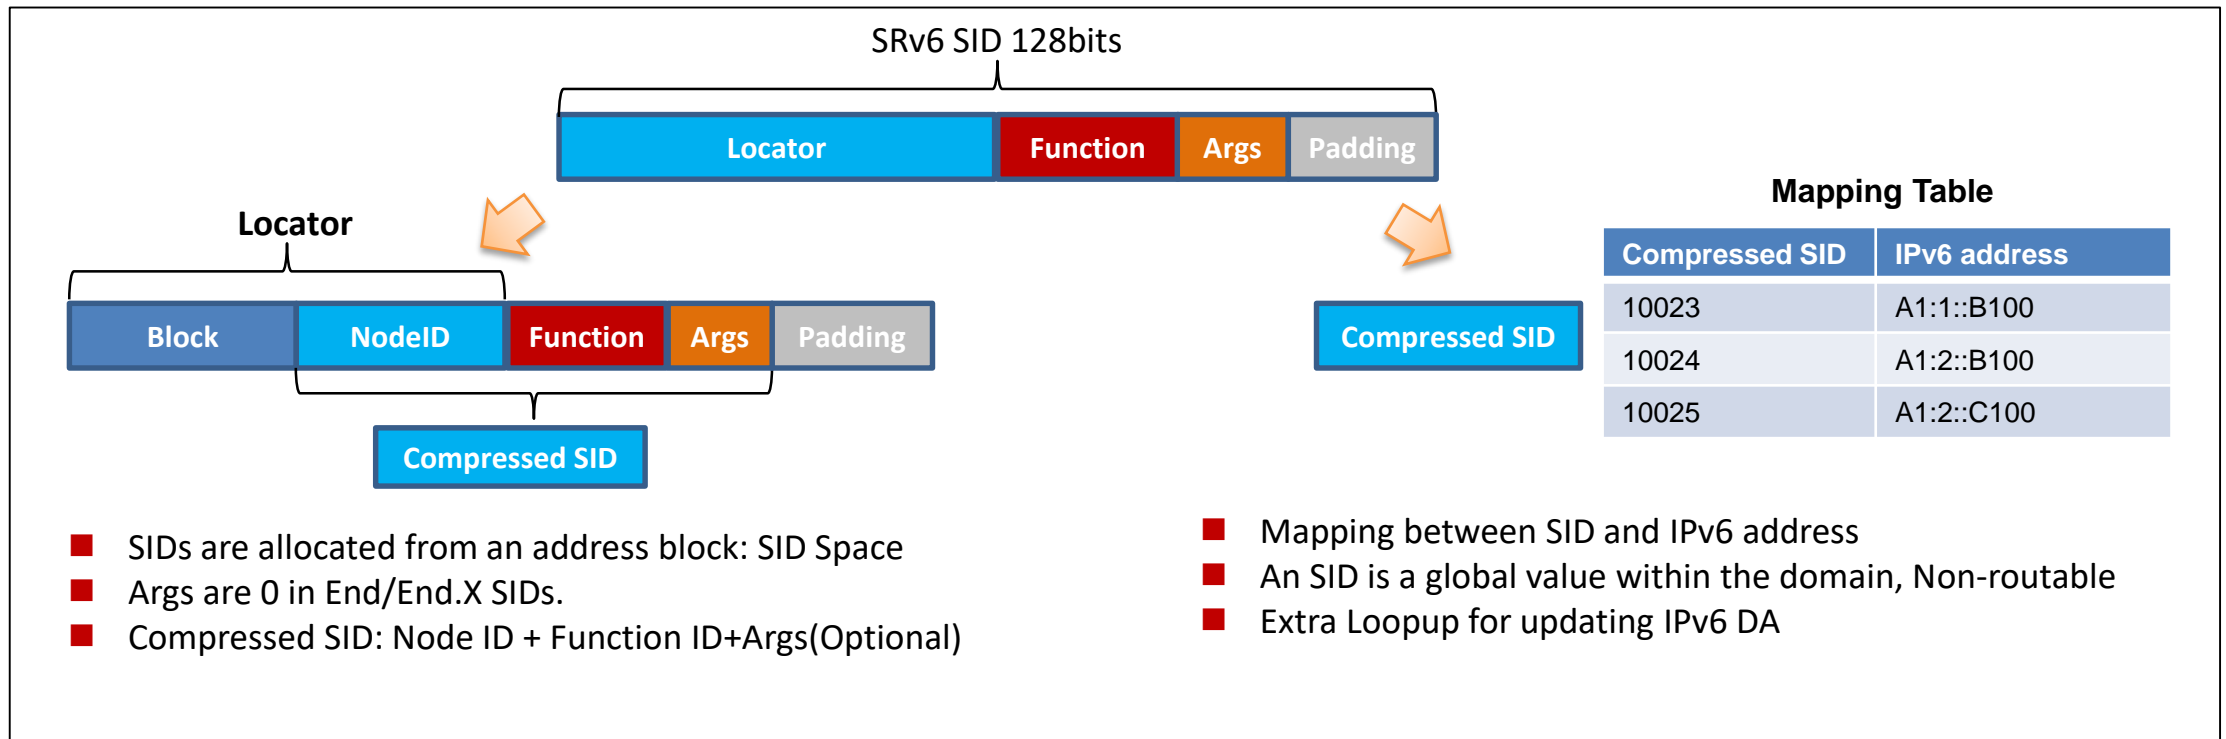
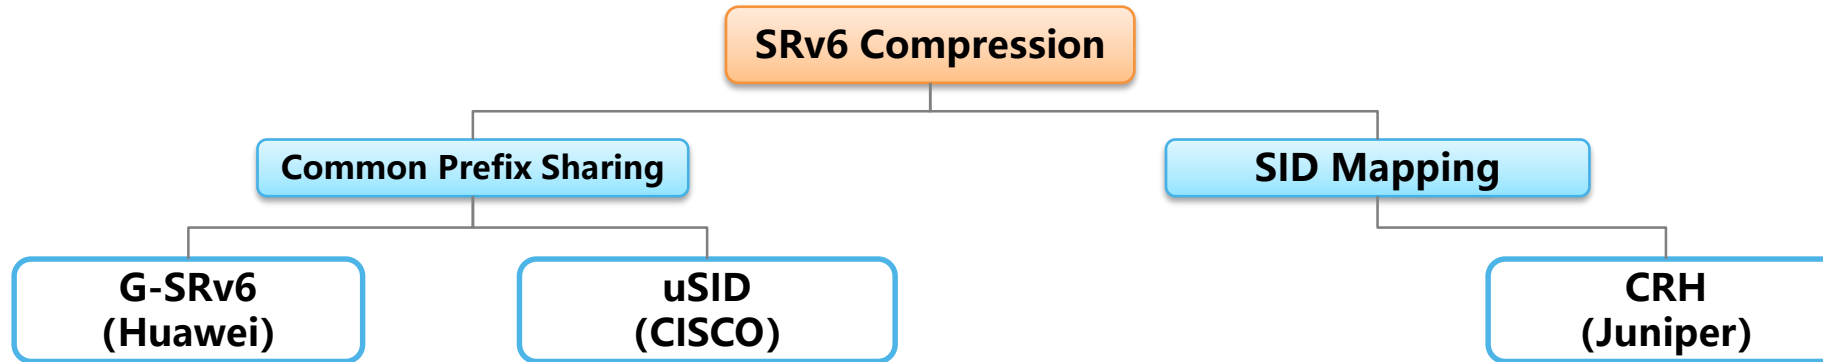
Cheng Li

Huawei IP Standard Representative

- 30+ IETF drafts, 10 + WG drafts, 1 RFC
- Currently focus on G-SRv6/SRv6, SFC, OAM
- Author of books
 - “SRv6 Network Programming - Ushering in a New Era of IP Networks”
 - “Refactoring Network: Architecture and Implementation of SDN”
- Paper: “ Application-aware G-SRv6 network enabling 5G services ” , INFOCOM 2021




Overview of SRv6 Compression Solutions




Design considerations

Compression efficiency



- Efficiency, Scalabilities, and Aligning should be considered.
- **32 bits is the ideal length**, 16 bits is not scalable

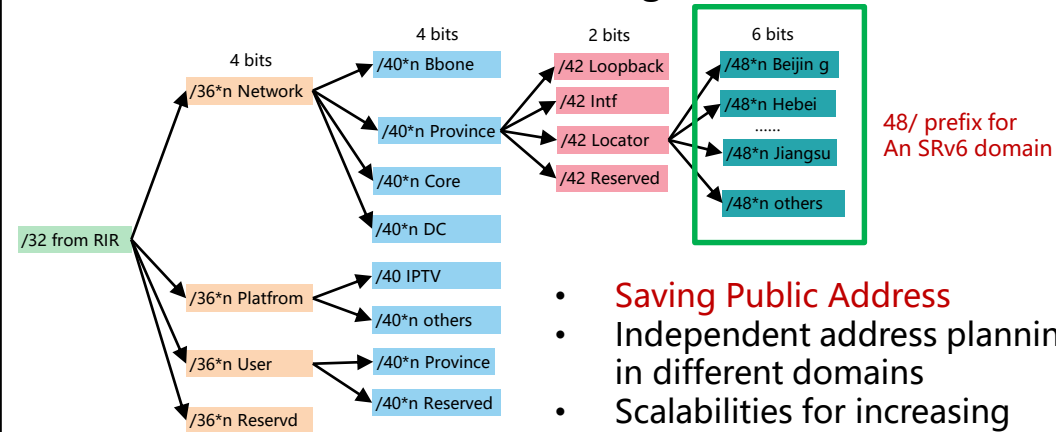
Native IPv6



Native IPv6 Routing

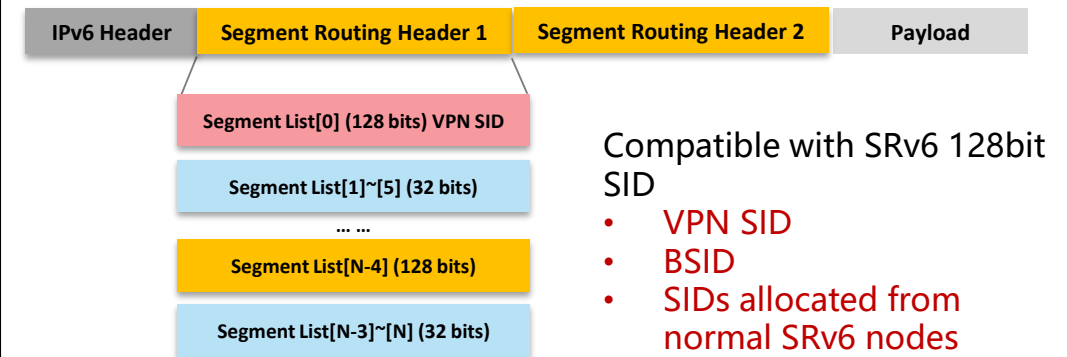
- Based on IP reachability, overlay routing on **cmpr disable nodes**
- Route aggregation, support super scale networking
- Failure domain isolation

Address Planning



- **Saving Public Address**
- Independent address planning in different domains
- Scalabilities for increasing Interfaces and nodes

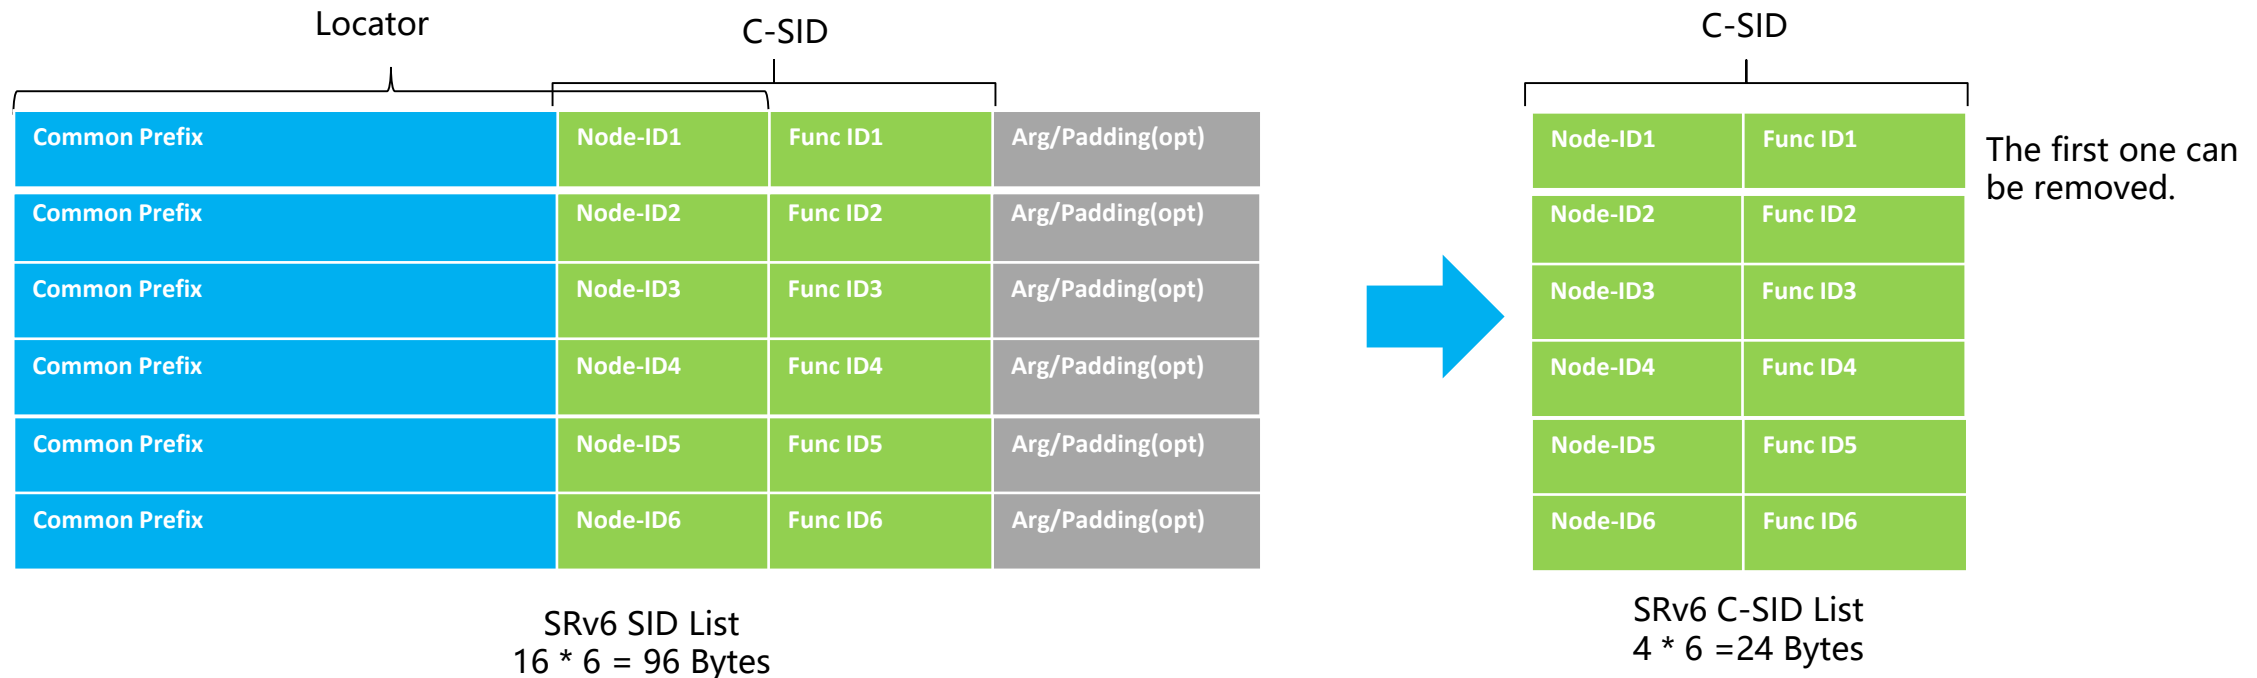
Compatible with SRv6



Generalized SRv6

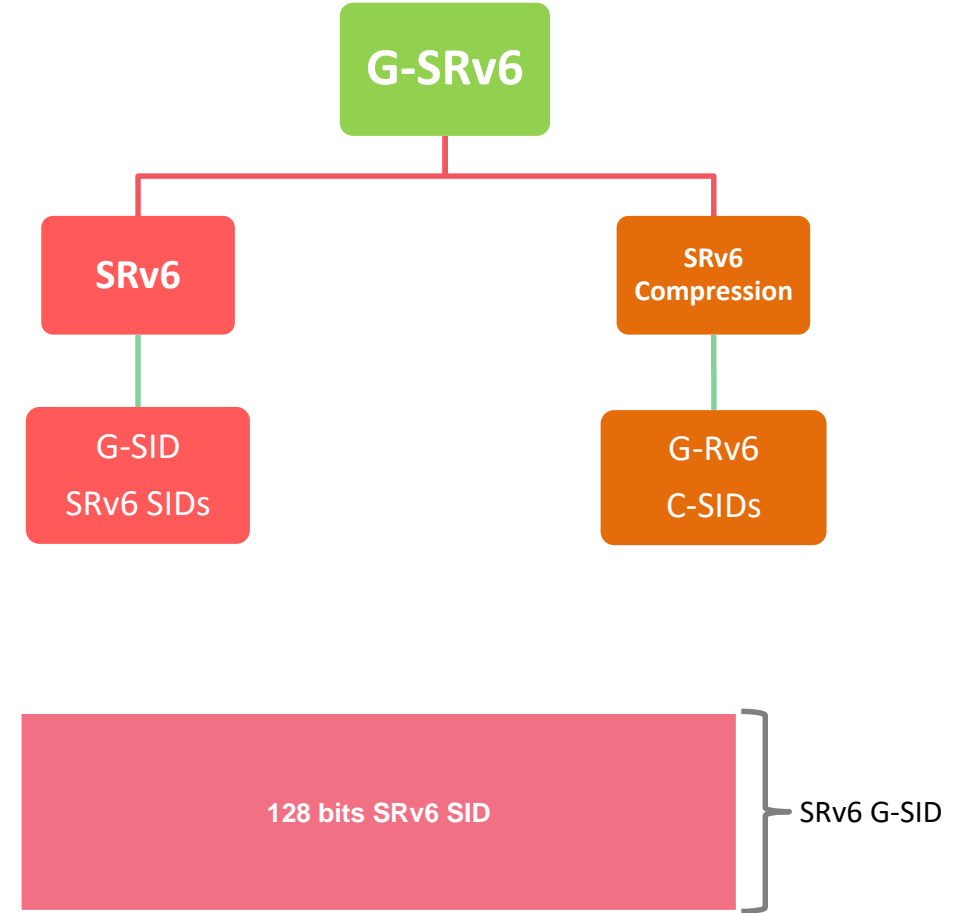
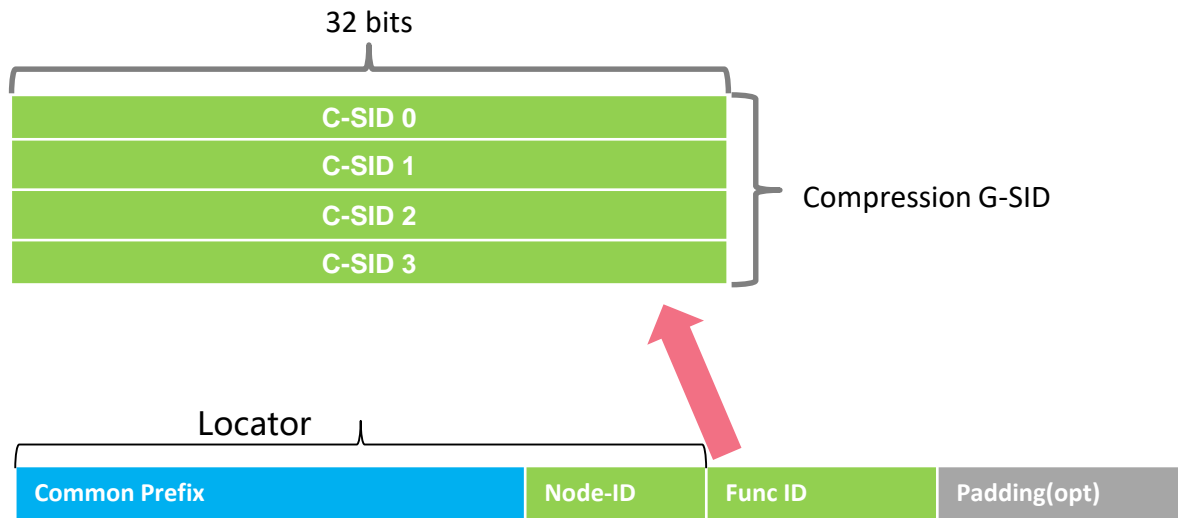
Before G-SRv6: SRv6 Compressed SID

- A normal SRv6 SID is a 128 bits IPv6 address allocated from an address block, called SID Space.
- For the SIDs in the SID list within an SRH, they may share the common prefix, and the common prefix is redundant that can be deleted to reduce the overhead.
- Each SRv6 SID has the format shown below, we called the different part of the SRv6 SID is compressed SID(C-SID), and the SID is a Compressible SRv6 SID.
- The prefix can be managed according to the real network address planning.
- Common Prefix is included in the first SID in the IPv6 Destination address.



G-SRv6: Compatible and Scalable

- Generalized SRv6 supports to encode multiple types of Segments in an enhanced SRH. G-SRv6 is compatible with SRv6 and uSID as well.
- These Segments can be called Generalized Segment. G-SID(Generalized Segment Identifier) is a 128-bits value, and it may contain:
 - an SRv6 SID(can be a Micro SID carrier)
 - a compression G-SID(4 32 bits C-SIDs at most)



G-SRH: Compatible with SRv6, Incremental Deployment, Hardware Friendly

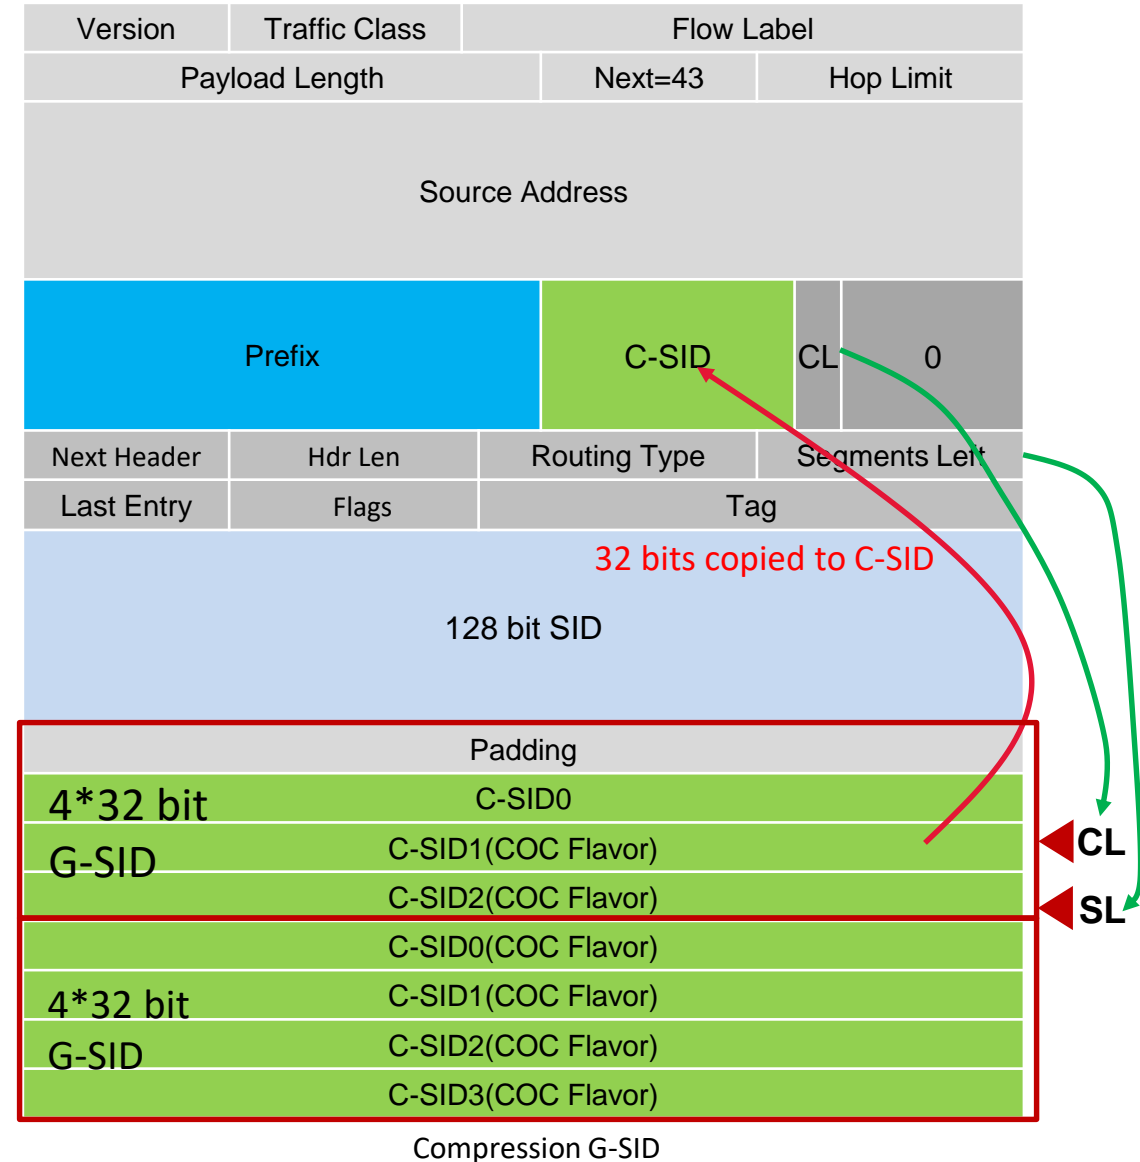
Solution: use SL to index a 128 bit G-SID, use CL to index C-SID inside this G-SID!

- C-flag in control plane: indicates the format of the SRv6 SID is compressible. The SID can be encoded as 128 or 32 bits in SRH
- COC(Continuation of Compression) flavor indicate the next SID is a 32-bits Compressed SID(C-SID)
- CL (Compressed SID left, the args of the compressible SRv6 SID) indicates the location of C-SID within the G-SID
 - Update C-SID from SRH[SL][CL] to IPv6 DA[CP: CP+31]



Pros

1. Fully compatible with SRH, NO modification of SRH
2. Fully compatible with SRv6, add COC Flavor endpoint behaviors, no affect of existing SIDs
3. Fully compatible with SRv6 control plane: (Can be) No modification of Control Plane
4. Address saving & easy to deploy:
 1. Flexible address planning, does not require for a short common prefix
 2. No new address required when reusing the Locator
 3. No new route, no modification of routing scheme(can share the same locator with normal SRv6 SIDs)
 4. Compressible SRv6 SID can be used as 128 bits or 32 bits. Reduce the number of SIDs.
5. Less overhead: A common prefix for a compressed sub-path instead of per 128 bits SID
6. Smooth upgrade/Incremental deployment: encode SRv6 SIDs and C-SIDs in a G-SRH
7. Hardware Friendly: No index mapping table
8. Compatible with Micro SID




Pseudo code: Only add code for COC Flavor SIDs, no Affection on Existing SIDs

Version	Traffic Class	Flow Label	
Payload Length		Next=43	Hop Limit
Source Address			
Prefix		C-SID	CL 0
Next Header	Hdr Len	Routing Type	Segments Left
Last Entry	Flags	Tag	
128 bit SID			
0	C-SID	C-SID(COC)	C-SID(COC)
Prefix		C-SID(COC)	0 (Padding)
128 SID			
C-SID	C-SID(COC)	C-SID(COC)	C-SID(COC)
C-SID(COC)	C-SID(COC)	C-SID(COC)	C-SID(COC)
Payload			

```

if DA is a COC Flavor SID:
    //update 32bits C-SID to DA
    if DA.CL = 0:
        //first C-SID in next 128 bits
        SL--
        DA.CL = 3;
    Else
        //next C-SID in current 128 bits
        DA.CL--
    DA[CP..CP+31] = SRH[SL][DA.CL]
    Forward the packet based on new DA
Else
    //update 128 bits SID to DA, original SRv6 Processing
    SRv6 processing
    
```



**CL is a location argument of the Compressible SID,
And it is the last 2 bits in Arguments**

PS. For easy understanding , the length of a row in SID list is 128bit

C-SID List + 128 VPN SID, 64 CP + 32 C-SID+32 Argument

SID List: 10 SIDs:

- A:1:1::, A:2:1::, A:3:1::, A:4:1::, A:5:1::, A:6:1::, A:7:1::, A:8:1:: are End.X with COC Flavor SIDs
- A:9:2:: is an End.X SID(C-flag=1, Without COC flavor)
- **A:10:10:: is an End.DT4 VPN SID**

Initialization: SL=3, CL=0, **Reduced mode.**

10 * 128 bits to 3 * 128 bits including a 128bit VPN SID. 70% overhead off.

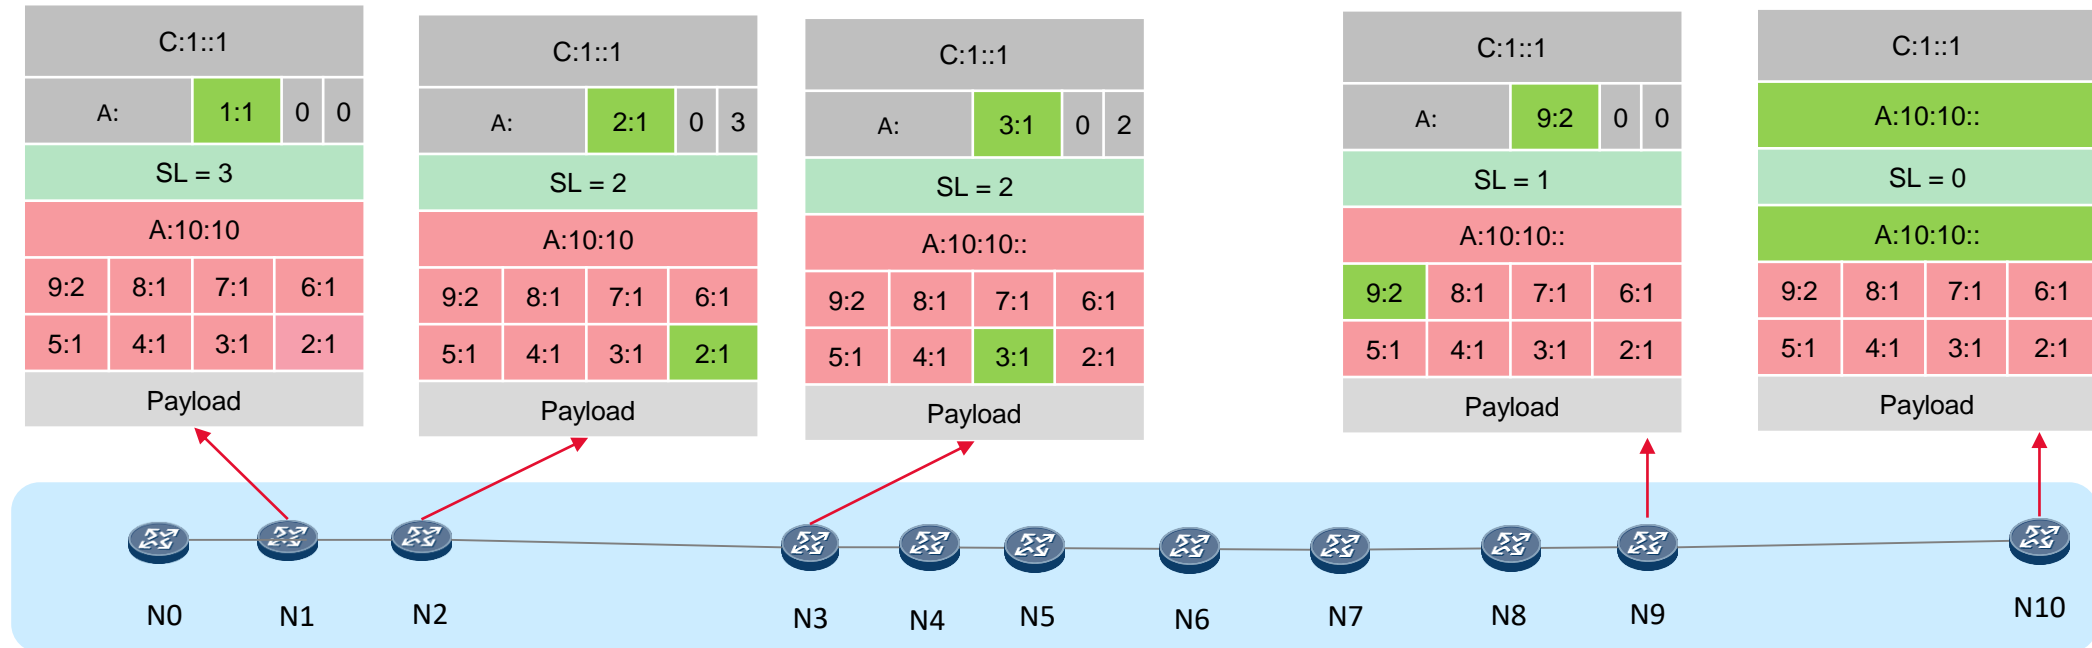
Compressible SRv6 SID and normal SRv6 SID use the same Locator, no new route is created!



Compressible SID: Locator A:1::/80 C-SID: 1:1 Argument 32bits 0



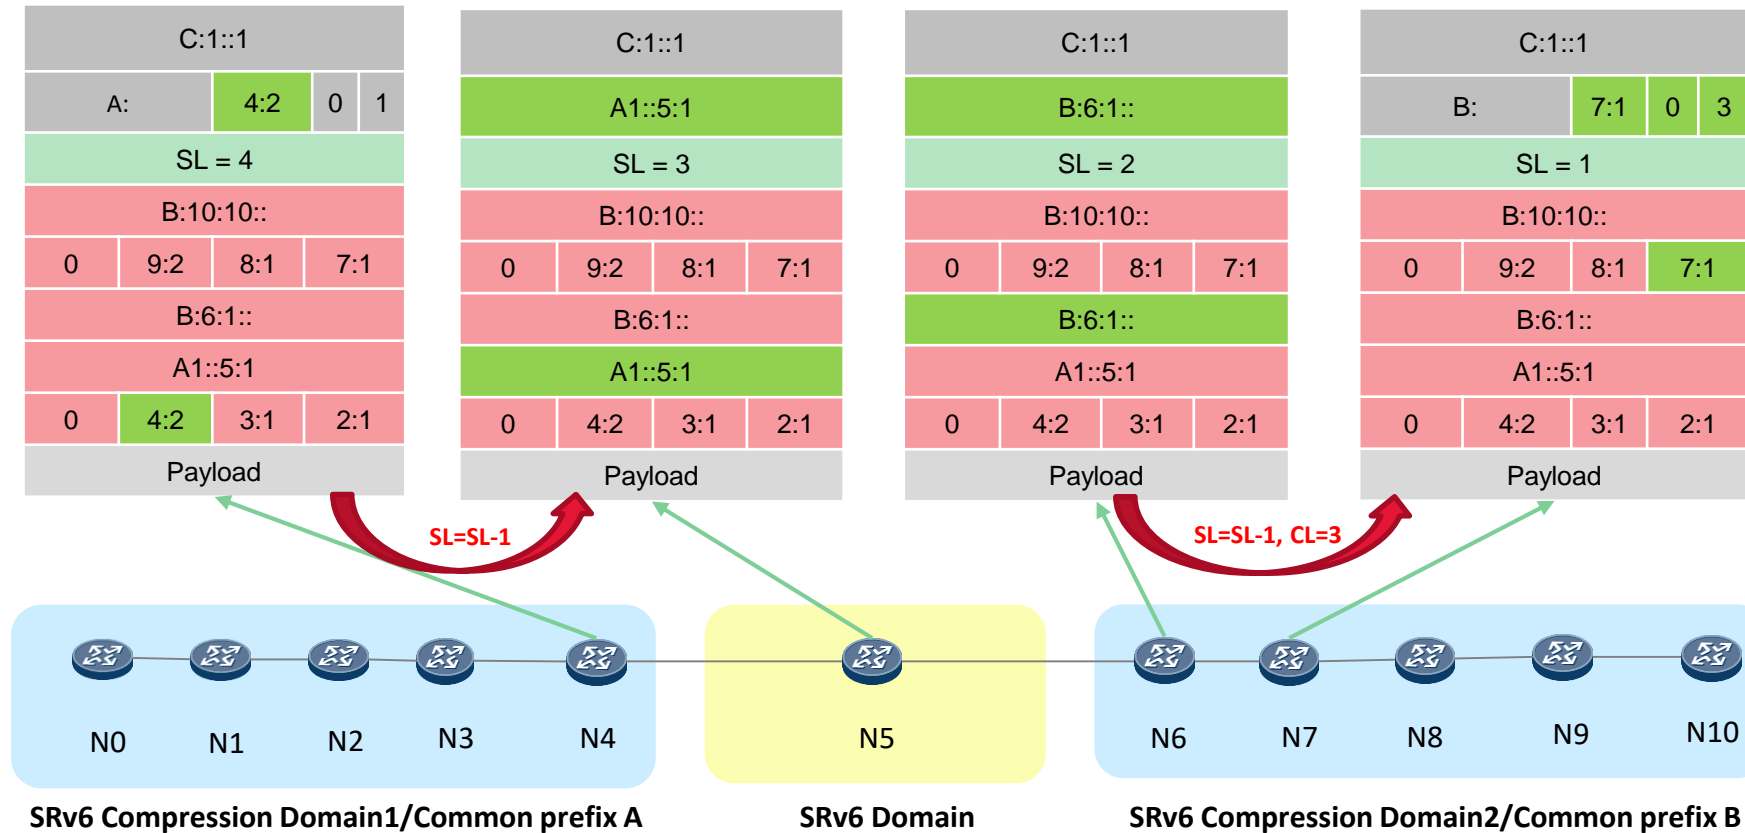
Normal SID: **Same Locator** A:1::/80 Function 1:1:1:1



Mixed Encoding with SRv6 SID for incremental deployment

SID List: 10 SIDs:

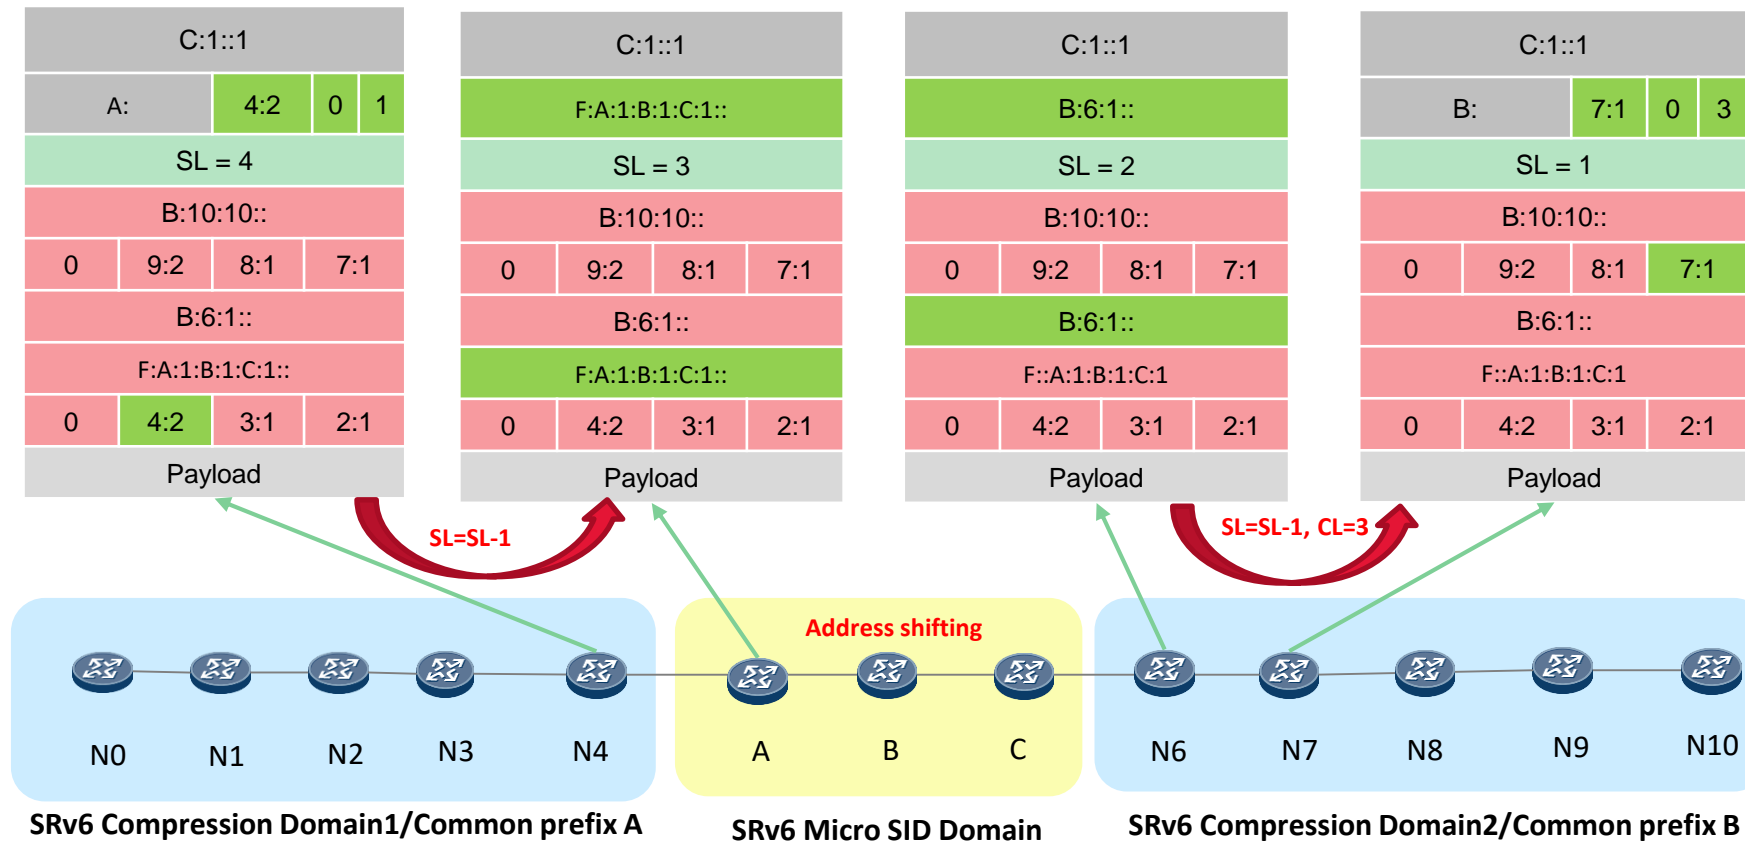
- A:1:1::, A:2:1::, A:3:1::, B:6:1:: , B:7:1::, B:8:1:: are End.X with COC Flavor SIDs
- A1::5:1 End.X does not support SRv6 compression.
- A:4:2:: , B:9:2:: are End.X SID(Without COC flavor)
- **B:10:10:: is an End.DT4 VPN SID**



Mixed Encoding with uSID

SID List: 10 SIDs:

- A:1:1::, A:2:1::, A:3:1::, B:6:1::, B:7:1::, B:8:1:: are End.X with COC Flavor SIDs
- F::A:1:B:1:C:1 is an uSID carrier, F is a 32/ prefix, A:1, B:1, C:1 is the uSID allocated by Node A, B and C.
- A:4:2::, B:9:2:: are End.X SID(Without COC flavor)
- B:10:10:: is an End.DT4 VPN SID



12 Vendors/10+ Customers support, CMCC Live Network trial done.

[Docs] [txt|pdf|xml|html] [Tracker] [Email] [Diff1] [Diff2] [Nits]

Versions: [00](#) [01](#)

SPRING Working Group Z. Li
 Internet-Draft C. Li
 Intended status: Standards Track Huawei Technologies
 Expires: February 15, 2021 W. Cheng
 China Mobile
 C. Xie
 C. Li
 China Telecom
 H. Tian
 F. Zhao
 CAICT
 August 14, 2020

Generalized Segment Routing Header
 draft-lc-6man-generalized-srh-01

[Docs] [txt|pdf] [Tracker] [Email] [Diff1] [Diff2] [Nits]


SPRING Working Group W. Cheng
 Internet-Draft China Mobile
 Intended status: Standards Track Z. Li
 Expires: November 21, 2020 C. Li
 Huawei Technologies
 F. Clad
 Cisco Systems, Inc
 A. Liu
 ZTE Corporation
 C. Xie
 China Telecom
 Y. Liu
 China Mobile
 S. Zadok
 Broadcom
 May 20, 2020

Generalized SRv6 Network Programming for SRv6 Compression
 draft-cl-spring-generalized-srv6-for-cmpr-01



12+ Vendors have PASSED Interop-test

Progress of G-SRv6: Interop-test and trial deployment plan



Conclusion

- **G-SRv6 is fully compatible with SRv6,**
 - **No SRH encapsulation modification**
 - **No new address consumption:** allocated SIDs from the Locator/ allocated to the node.
 - **No new route creation:** share the same locator with the normal SRv6 SID.
 - **No control plane modification:** Controller can install the SR policy with 128 bit G-SIDs, endpoint nodes understand the COC Flavor behaviors, Compression disable SRv6 nodes are unaware of Compression.
 - **No security policy modification.**
- **G-SRv6 has less overhead**
 - Each compression sub-path has only one common prefix, instead of for each 128 bits.
- **G-SRv6 has efficient address consumption**
 - It is **not** required to allocate a short common prefix for better compression.
- **G-SRv6 supports incremental deployments, which can be deployed on demand.**

G-SRv6 Info: INFOCOM2021 paper, SRv6 books, IPv6+ video, G-SRv6 Community

Application-aware G-SRv6 network enabling 5G services

Cheng Li, Jianwei Mao, Shuping Peng, Yang Xia, Zhibo Hu, Zhenbin Li
Huawei Technologies, Beijing, China
{c.l.maojianwei, pengshuping, yolanda.xia, huzhibo, lizhenbin}@huawei.com

Abstract—This demo showcased how application-aware G-SRv6 network provides fine-grained traffic steering with more economical IPv6 source routing encapsulation, effectively supporting 5G eMBB, mMTC and uRLLC services. G-SRv6, a new IPv6 source routing paradigm, introduces much less overhead than SRv6 and is fully compatible with SRv6. Up to 75 percent overhead of an SRv6 SID List can be reduced by using 32-bit compressed SID with G-SRv6, allowing most merchant chipsets to support up to 10 SIDs processing without introducing packet recirculation, significantly mitigating the challenges of SRv6 hardware processing overhead and facilitating large-scale SRv6 deployments. Furthermore, for the first time, by integrating with Application-aware IPv6 networking (APN6), the G-SRv6 network ingress node is able to steer a particular application flow into an appropriate G-SRv6 TE policy to guarantee its SLA requirements and save the transmission overhead in the meanwhile.

Keywords—SRv6 Compression, G-SRv6, APN6


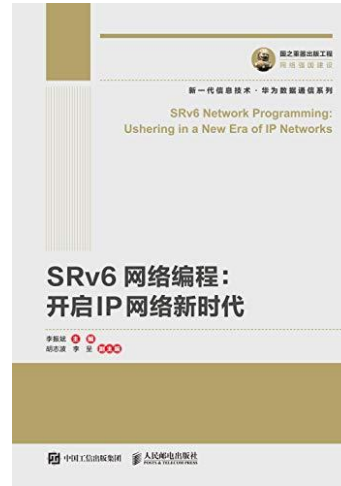
I. INTRODUCTION

As 5G and industry verticals evolve, ever-emerging new services with diverse but demanding requirements such as low latency and high reliability are accessing to the network. Different applications have differentiated network Service Level Agreement (SLA). For instance, on-line gaming has highly demanding requirements on latency, live video streaming has high requirements on both latency and bandwidth, while backup traffic mainly requires more bandwidth but is less sensitive of latency. However, in current networks, the operators remain unaware of the traffic type traversing their network, making the network infrastructure essentially dumb pipes and

recirculation. This has become a big obstacle for SRv6 deployment in practice.

We proposed Generalized Segment Routing over IPv6 (G-SRv6) [3][4][5] to address the challenges of SRv6 overhead. While compatible with SRv6, G-SRv6 provides a mechanism to encode Generalized SIDs (G-SID) in the Generalized SRH (G-SRH), where a G-SID can be a 128-bit SRv6 SID, a 32-bit compressed SID (C-SID) or some other types. A 32-bit C-SID saves 75% overhead of the SID, so that the size of SRH can be significantly compressed. It also supports incremental upgrade from SRv6 by encoding both SRv6 SIDs and C-SIDs in the SRH. With G-SRv6, most the merchant chipsets can support up to 10 SIDs processing without packet recirculation so that the challenges of SRv6 hardware processing is mitigated, facilitating the large-scale SRv6 deployment. So far, G-SRv6 has been implemented in Linux Kernel, and hardware devices from more than 10 vendors.

This demo showcases that APN6 over G-SRv6 enables fine-grained traffic scheduling and efficient IPv6 source routing encapsulation for services in 5G scenarios, and what benefits G-SRv6 can provide over SRv6. Using APN6, the eMBB, mMTC, and uRLLC traffic is forwarded following the high-bandwidth path, the Service Function Chain (SFC) path, and the lowest latency path, respectively. Using APN6 over G-SRv6, over 50% transmission overhead is reduced, and the Flow-Completion Time (FCT) is shortened from 923s to 102s. Comparing to SRv6 (with 10 SIDs in SRH), the forwarding rate of an SRv6 endpoint node is raised by 55% from 400Mpps to 620Mpps. In summary, the application-aware G-SRv6 helps network operators to reduce the cost and generate more revenue in the 5G area.



INFOCOM 2021: <https://www.youtube.com/watch?v=BI0r16XH4go>

IETF Hackathon: <https://www.youtube.com/watch?v=qbryDg8fXRM>

CAICT IPv6plus: https://mp.weixin.qq.com/s/PIrJyiVEQu2qln6vIgY_Zw

Huawei Support: <https://support.huawei.com/enterprise/zh/doc/EDOC1100183731>

G-SRv6 Community: <https://github.com/G-SRv6>

More G-SRv6 Info: <https://www.ipv6plus.net/Phase2/Generalized-SRv6/>

Book: SRv6 Network Programming Ushering in a New Era of IP Networks Chapter 13

IPv6+ G-SRv6: <https://www.ipv6plus.net/Phase2/Generalized-SRv6/>

Thanks

Huawei Live

