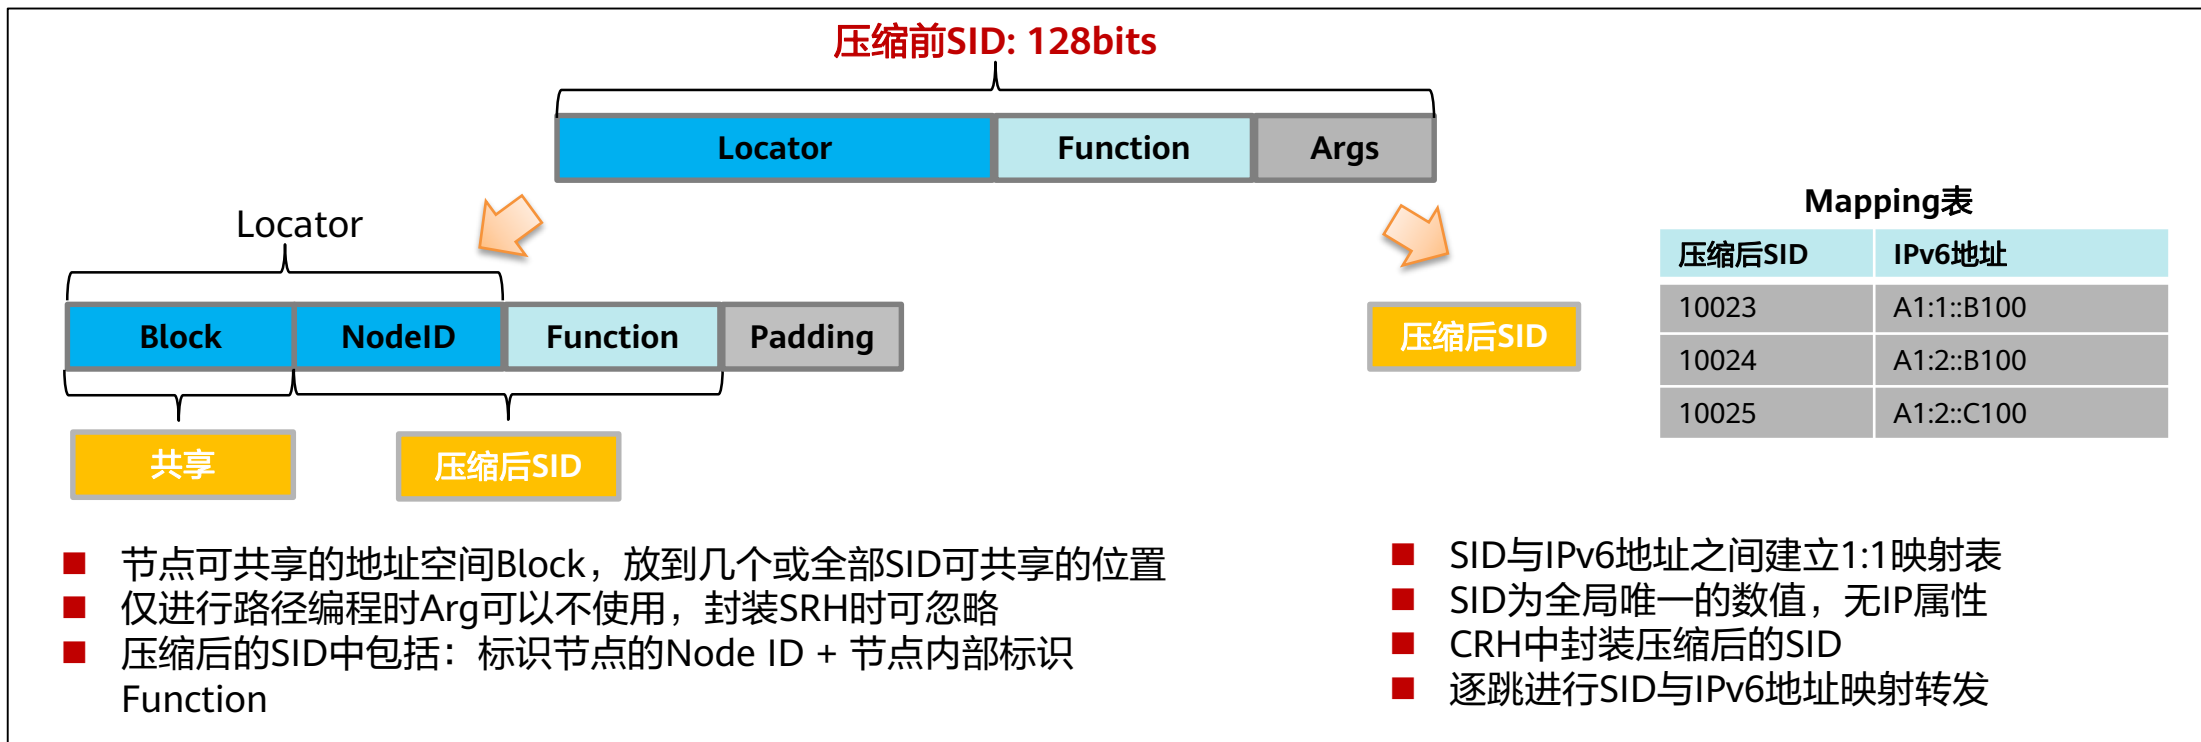
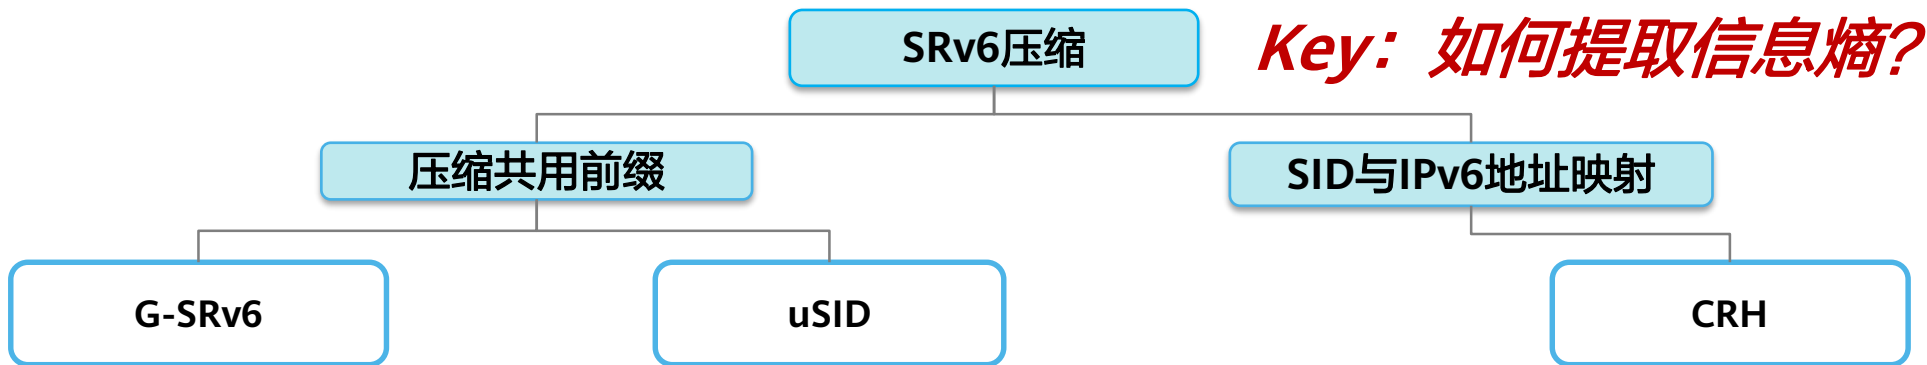


# G-SRv6技术介绍与方案对比



# 业界SRv6压缩方案总览




# 压缩方案考量因素



## 网络规模

### 压缩效率高，不降低网络规模


- ✓ 满足字节对齐的前提下，尽量提升压缩效率；
- ✓ 压缩方案需要考虑网络规模



## 设备兼容

### 满足不同设备，不同需求

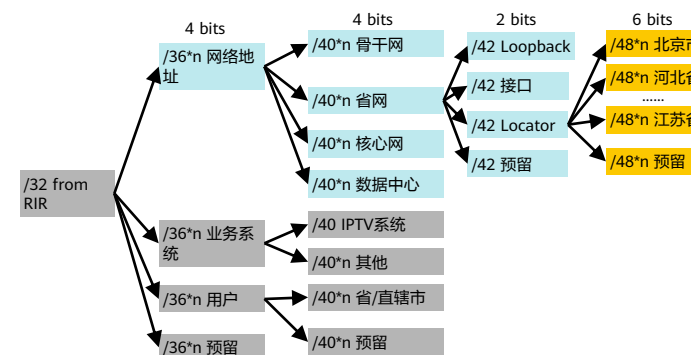
- ✓ 兼容现有SRv6部署；
- ✓ 保留Native IP价值属性，路由可达、可聚合、域间隔离



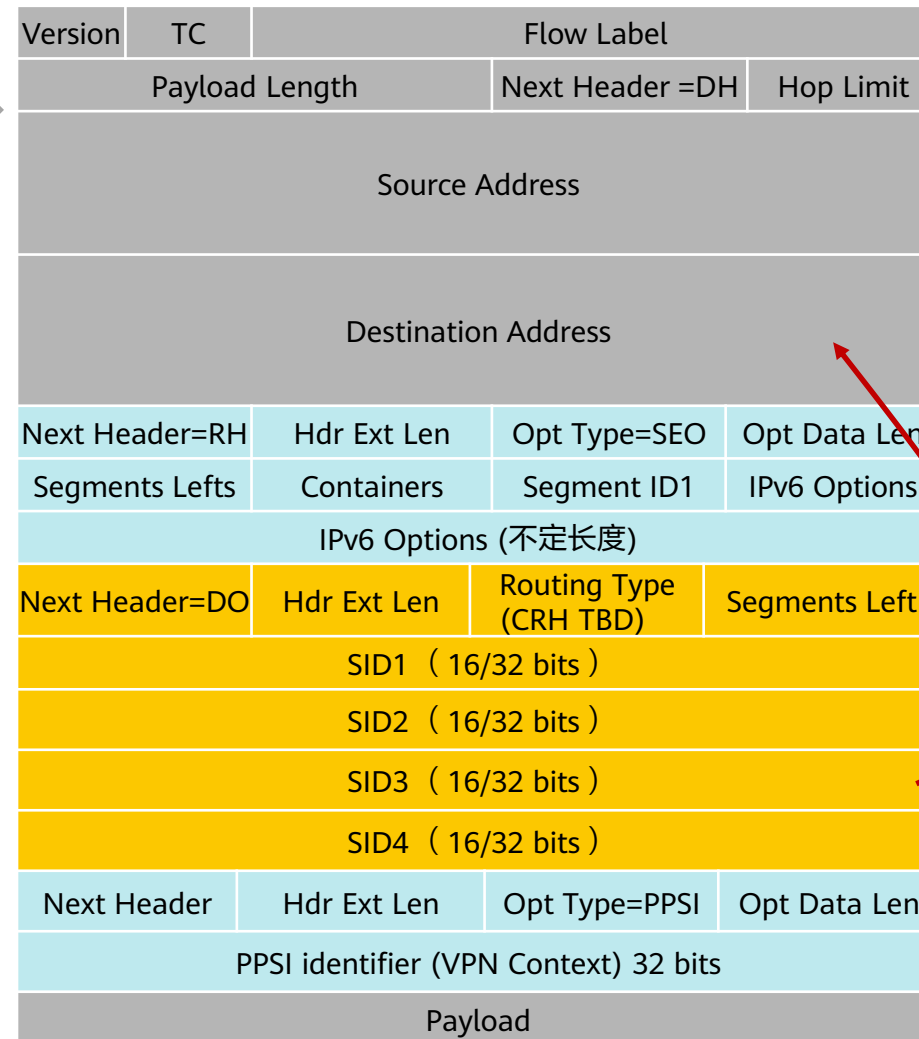
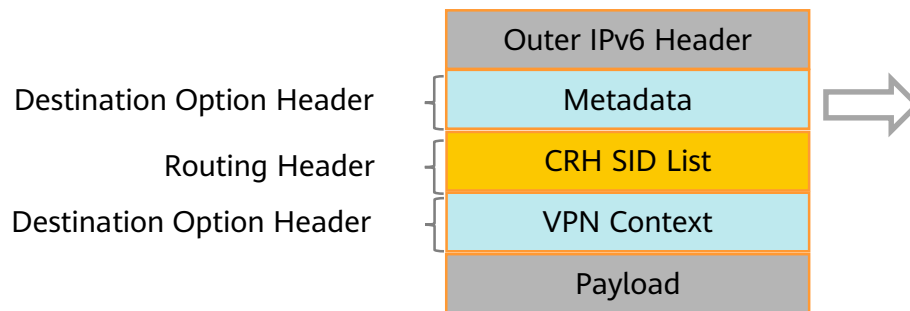
## 地址规划

### 整体不增加网络规划复杂度

- ✓ 尽可能节省公网空间；
- ✓ 不同管理域独立规划



# CRH ( SRm6 ) 方案 — 与SRv6不兼容



## 原理:

- 定义新的Routing Header类型CRH-16/CRH-32 ( TBD ) , 只携带路径SID, 采用16/32bits的SID;
- 建立SID与IPv6地址之间的映射表, 逐跳进行映射转发
- Routing Header中不携带VPN信息, 由Destination Option Header的IPv6 Option TLV携带 Option TLV由Destination Option Header的IPv6 Option TLV携带, 定义新的类型Segment Endpoint Option ( TBD )

## 压缩效率:

- 采用16 bits SID, 87.5% (较大规模网络无法使用16bits), 采用32 bits SID, 75%

## 缺点:

- 新定义RH类型与Option TLV类型, 不兼容原生SRv6的SRH
- SID不具备IP寻址能力, 也无汇聚大规模组网能力
- Node SID需全局唯一, 统一规划
- VPN信息采用Option TLV编码, 处理复杂, 效率低

Mapping  
表查询获  
取IPv6地  
址

SL

# uSID ( Micro SID ) 方案 32bit压缩效率不足

uSID处理前

uSID-Block	Active uSID	Next uSID	Last uSID
bbbb:bbbb:/32	0100:0001	0200:0002	0300:0003

Shifting

uSID处理后

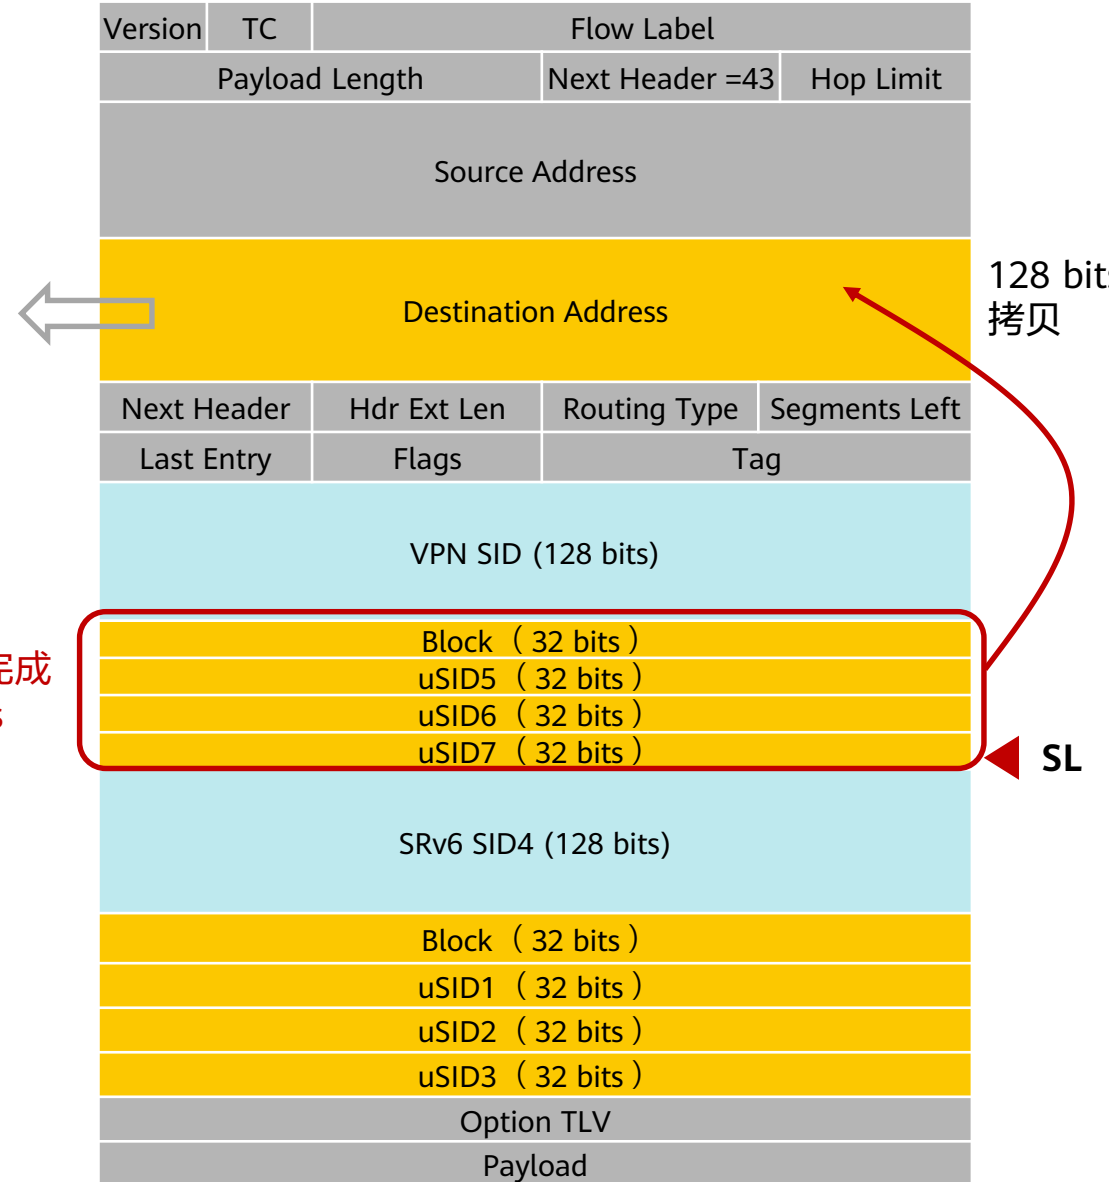
uSID-Block	Active uSID	Next uSID	EOC
bbbb:bbbb:/32	0200:0002	0300:0003	0000:0000

uSID-Block	Active uSID	Next uSID	EOC
bbbb:bbbb:bbbb:/48	0200:0002	0300:0003	0000:0000

一个uSID-Carrier 128bits处理完成后，SL减1，处理下一个128bits

## 原理

- 控制面发布特殊SID，标识为uSID，以及对应的前缀路由 Block : uSID
- 携带uSID的载体（uSID-Carrier）为128bit的SRv6 SID，其格式为：
  - **Active uSID**: 当前活跃的SID，举例中为32bits
  - **Next uSID**: 下一个 u SID，在处理后会偏移到Active uSID位置，然后在SID[96-127]的位置置<0000>EOC
  - **Last uSID**: 本128bits中最后一个uSID
  - **EOC**(End of Carrier): uSID的终止符，全0
  - 当检测Active uSID为EOC时，终止uSID的处理，执行END SID的操作，替换下一个128bit SRv6 SID到DA中。



# uSID ( Micro SID ) 16 bit 方案复杂，定义新的地址格式

uSID处理前	uSID-Block	uSID	uSID	uSID	uSID
	bbbb:bbbb::/64	1000	2000	3000	4000
uSID处理后	uSID-Block	uSID	uSID	uSID	uSID
	bbbb:bbbb::/64	2000	3000	4000	0000

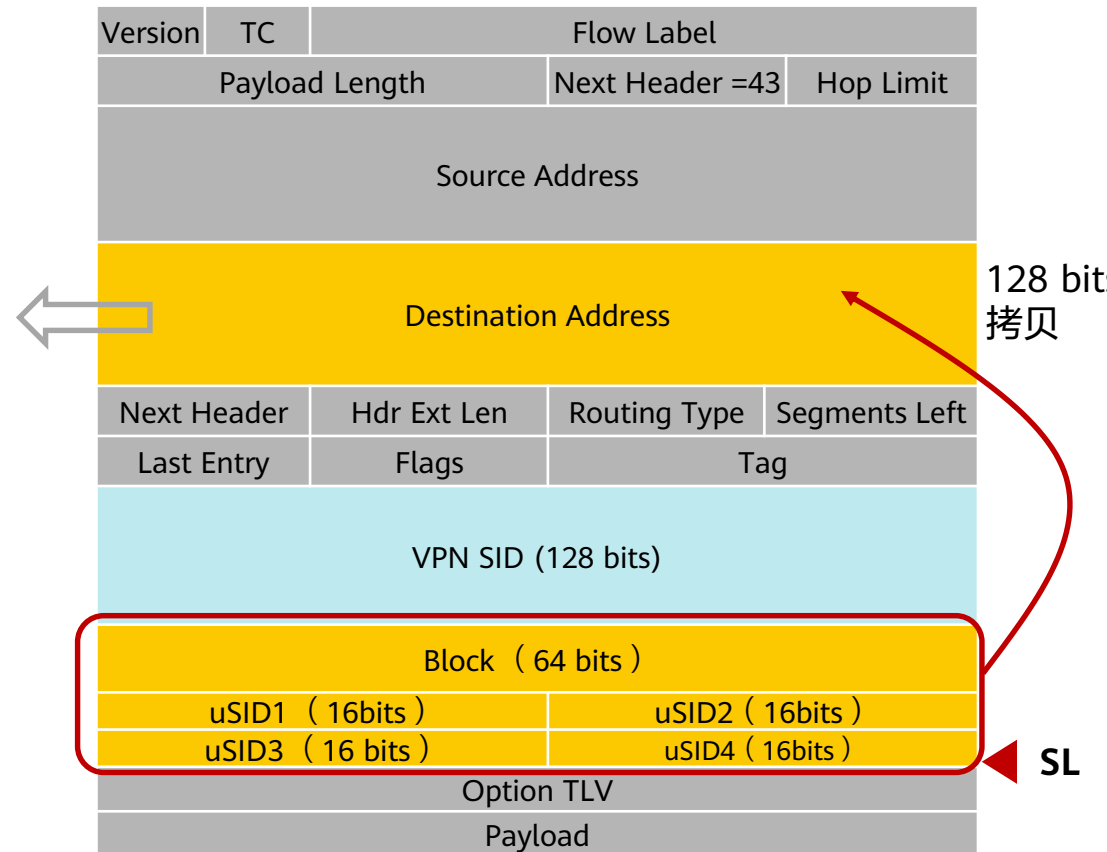
*Shifting*

引入全局SID和本地SID概念，在16bit uSID中使用高4位用于区分全局和本地SID：

- 前4bit 0-D：全局 14/16
- 前4bit E-F：本地 2/16

强制要求地址规划符合此概念，否则无法部署。地址强规划

所以一个Block中有  $14 \times 2^{12} = 57344$  个全局ID（大规模组网受限，分层划分编码需更多bit，大网络普遍选择20bit），8192个本地ID




一个uSID-Carrier 128bits处理完成后，SL减1，处理下一个128bits

# uSID ( Micro SID ) 方案

## 举例:

节点1到节点2之间的L3VPN业务, 需要经过节点8和节点7, 实现流量工程



**准备:** 为每个节点/业务分配唯一的uSID, 如节点8分配0800, uSID与公用的IPv6 block `bbbb:bbbb::/32`拼接后形成路由`bbbb:bbbb:0800::/48`, 通过IGP发布

- @8 advertises the IGP route `bbbb:bbbb:0800::/48`
- @7 advertises the IGP route `bbbb:bbbb:0700::/48`
- @2 advertises the IGP route `bbbb:bbbb:0200::/48`

## 转发:

- @1 encapsulates IPv4 packet from Site A and sends an IPv6 packet with DA = `bbbb:bbbb:0800:0700:0200:0000:0000:0000`
- @1 forwards to 4 (shortest-path to 8 (`bbbb:bbbb:0800::/48`))
- @4 & @5 IPv6 forwarding (shortest-path to 8 (`bbbb:bbbb:0800::/48`)), DA no change
- @8 SRv6 uN behavior: shift and forward  
Longest-Match local FIB `bbbb:bbbb:0800::/48`  
Rx'd DA: `bbbb:bbbb:0800:0700:0200:0000:0000:0000`  
**shift << 16**  
Tx'd DA `bbbb:bbbb:0700:0200:0000:0000:0000:0000`
- @7 SRv6 uN behavior: shift and forward, outer DA becomes `bbbb:bbbb:0200:0000:0000:0000:0000:0000`
- @6 & @3 IPv6 forwarding (shortest-path to 8 (`bbbb:bbbb:0200::/48`)), DA no change
- @2 SRv6 End.DX4: Decapsulate and cross-connect inner packet


# uSID: 扩展性不足, 需规划多套地址, 需要ULA才能部署

## 优点

- 兼容原生SRv6
- SID保留Native IPv6属性

## 缺点

- 在SRv6 Locator之外, 需单独规划uSID block地址块, 新增uSID路由表项
- 为支持压缩, 需要规划私网地址ULA部署uSID, 回落到SR-MPLS; 跨域, 公网业务发布均存在问题, 需再规划对外发布的Locator用于发布Service。需要规划三套地址:
  - SRv6公网Service SID: SD-WAN, 端到端, 跨域等场景
  - SRv6私网SID: 用于内部拓扑描述(可以合并到公网Service SID网段)
  - uSID私网SID: 用于支持压缩
- uSID block越短, 压缩效率越好, 但地址消耗越多, 压缩效果依赖于地址规划, 强规划
- 前缀长于48则最多一个128bit uSID carrier 能带4个 16 bit uSID, 等于2个32bit uSID。
- 当SID全为NodeID或Adj SID时才可以实现16bit压缩。松散路径时需要由一个16 bit 可路由的uSID+ 16 bit本地uSID组成一个32bit的可路由uSID指导数据包转发, 回落32bit uSID。



- @8 advertises the IGP route `bbbb:bbbb:0800::/48`
- @7 advertises the IGP route `bbbb:bbbb:0700::/48`
- @2 advertises the IGP route `bbbb:bbbb:0200::/48`

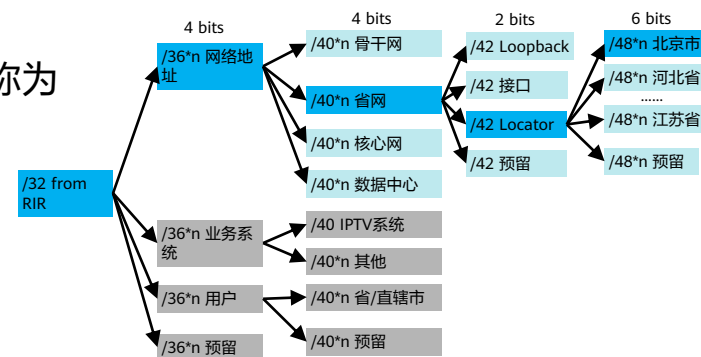
公网地址长度	uSID长度	扩展性	压缩效率
32	16	差	78%
32	32	中	66%
64	32	中	50%



# Generalized SRv6

# G-SRv6压缩原理

- 原生SRv6 SID为128bits IPv6地址，每个节点从自身的Locator地址空间中独立分配
- 而网络中节点的Locator绝大部分都是从**同一个大段的地址空间中逐级分配的**，该地址空间，称为Common Prefix
- 在一个SRH SR List中
  - Common Prefix在SRH中为冗余信息，可将其放到统一的位置：IPv6 DA**
  - NodeID + Function ID + Args (可选) 为有效信息，SRH中封装该信息熵即可，称为Compressed SID(C-SID)**
  - Padding字段通常为0，无用信息，可直接删除




Node Locator

Common Prefix	Node-ID1	Func + Args(opt)	Padding
Common Prefix	Node-ID2	Func + Args(opt)	Padding
Common Prefix	Node-ID3	Func + Args(opt)	Padding
Common Prefix	Node-ID4	Func + Args(opt)	Padding
Common Prefix	Node-ID5	Func + Args(opt)	Padding
Common Prefix	Node-ID6	Func + Args(opt)	Padding

SRv6 SID List  
16 Bytes \* 6 = 96 Bytes

信息熵



Common Prefix in IPv6 DA	
Node-ID1	Func + Args(opt)
Node-ID2	Func + Args(opt)
Node-ID3	Func + Args(opt)
Node-ID4	Func + Args(opt)
Node-ID5	Func + Args(opt)
Node-ID6	Func + Args(opt)

SRv6 C-SID List  
4 \* 6 = 24 Bytes

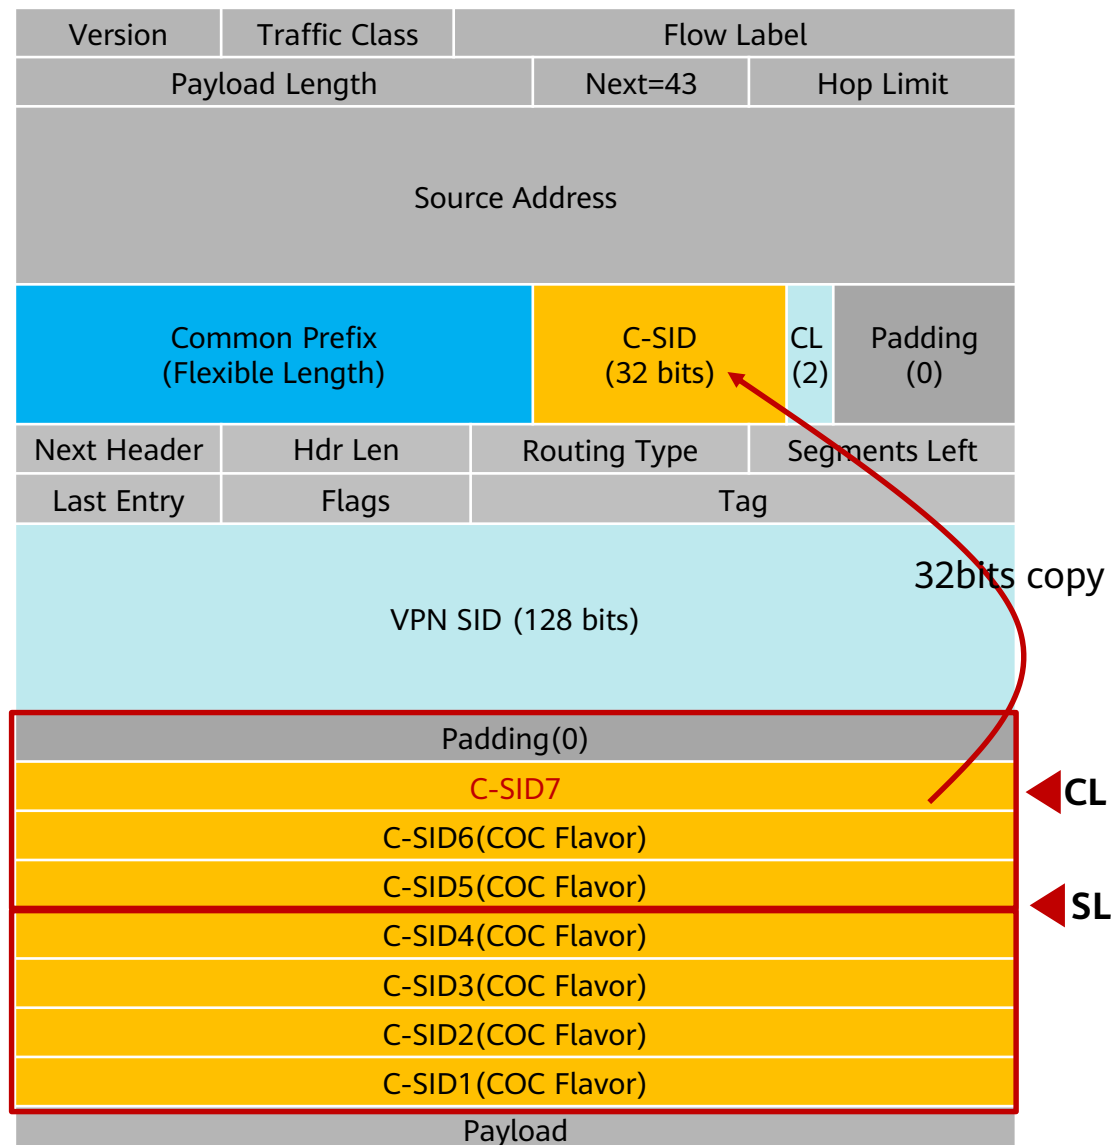
# G-SRH: 兼容原生SRv6，可增量部署，高扩展性

## Solution

- SRH中的128bits中可封装4 \* 32bits C-SID，通过**CL (Compressed SID left)** 标识 C-SID在128bits/32bits=4 SID小循环中的位置，取值0~3
- 更新后的**C-SID = SRH[SL][CL]**，将该32bits C-SID拷贝到IPv6 DA[CP: CP+31]
- 定义**COC(Continuation of Compression)** Flavor，标识下一个SID是压缩后的C-SID，如果没有COC Flavor，标识下一个SID为128bits SRv6 SID
- COC flavor类似于PSP flavor，在IGP/BGP分配SRv6 SID时，通过控制面发布

## Benefit

- 128bits->32bits，压缩效率75%
- 兼容现有SRv6**，支持SRv6 C-SID与128bits SRv6 SID混编，**可增量部署**
- 地址规划灵活，不浪费
  - 压缩SID和非压缩SID可共用地址空间，**不需要规划新的地址**
  - Common Prefix长度可灵活规划，长度<94即可，需要保证Node-ID+Function = 32bit
- 32bits C-SID内部可任意划分Node-ID+Function，**满足不同网络规模要求**，如：
  - 1) 16 bits Node ID + 16 bits Function，支持64K node + 64k function/node
  - 2) 18 bits Node ID + 14 bits Function，支持262K node + 16K function/node




# G-SRv6伪码: 只增加COC Flavor处理, 不影响原有SRv6 SID处理

Version	Traffic Class	Flow Label	
Payload Length		Next=43	Hop Limit
Source Address			
Common Prefix	C-SID	CL	Padding (0)
Next Header	Hdr Len	Routing Type	Segments Left
Last Entry	Flags	Tag	
128 bits SID			
0(Padding)	C-SID	C-SID(COC)	C-SID(COC)
Common Prefix		C-SID(COC)	0(Padding)
128 bits 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 local SID is a COC Flavor SID // Update 32bits C-SID to DA
    if DA.CL = 0 // First C-SID in next 128 bits
        SL--; CL = 3;
    else // Next C-SID in current 128 bits
        CL--;
    DA[CP..CP+31] = SRH[SL][DA.CL]; // CP: Common Prefix length
    Forward the packet based on new DA;

else
    // Original SRv6 processing, Update 128 bits SID to DA
    // Non COC Flavor SID or 128bits SID
    
```

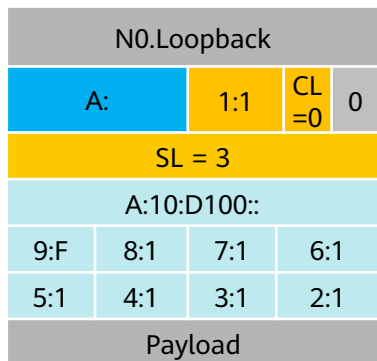


PS. For easy understanding , the length of a row in SID list is 128bit

# G-SRv6转发示例 — 纯压缩场景 1


## SID List: 严格显示路径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:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- **A:10:D100:: 为N10的End.DT4 VPN SID**



N0为头结点，进行SRH编排:

- A:1:1::为第一跳SID，进行Reduced操作，封装到IPv6 DA中
- 封装N2~N8的压缩SID (COC Flavor)
- 封装倒数第二跳N9的压缩SID (NO COC Flavor)
- 封装尾节点128bits VPN SID



64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀

# G-SRv6转发示例 — 纯压缩场景2

## SID List: 严格显示路径10 SIDs:

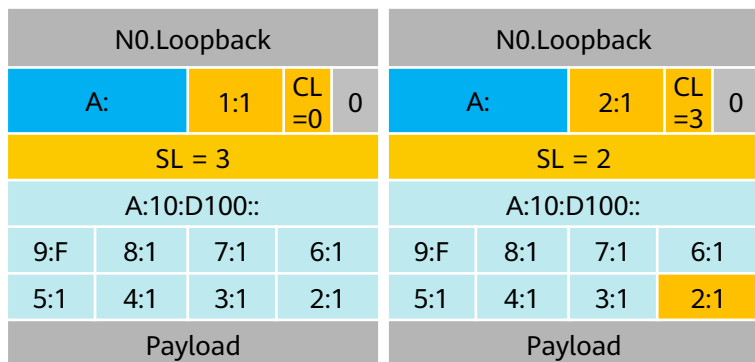
- A:1:1::, A:2:1::, A:3:1::, A:4:1::, A:5:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀




N1匹配到Local SID:

A:1:1::为COC Flavor&&CL==0

SL--;CL=3;

DA[CP..CP+31] = SRH[SL][CL];



# G-SRv6转发示例 — 纯压缩场景3

SID List: 严格显示路径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:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

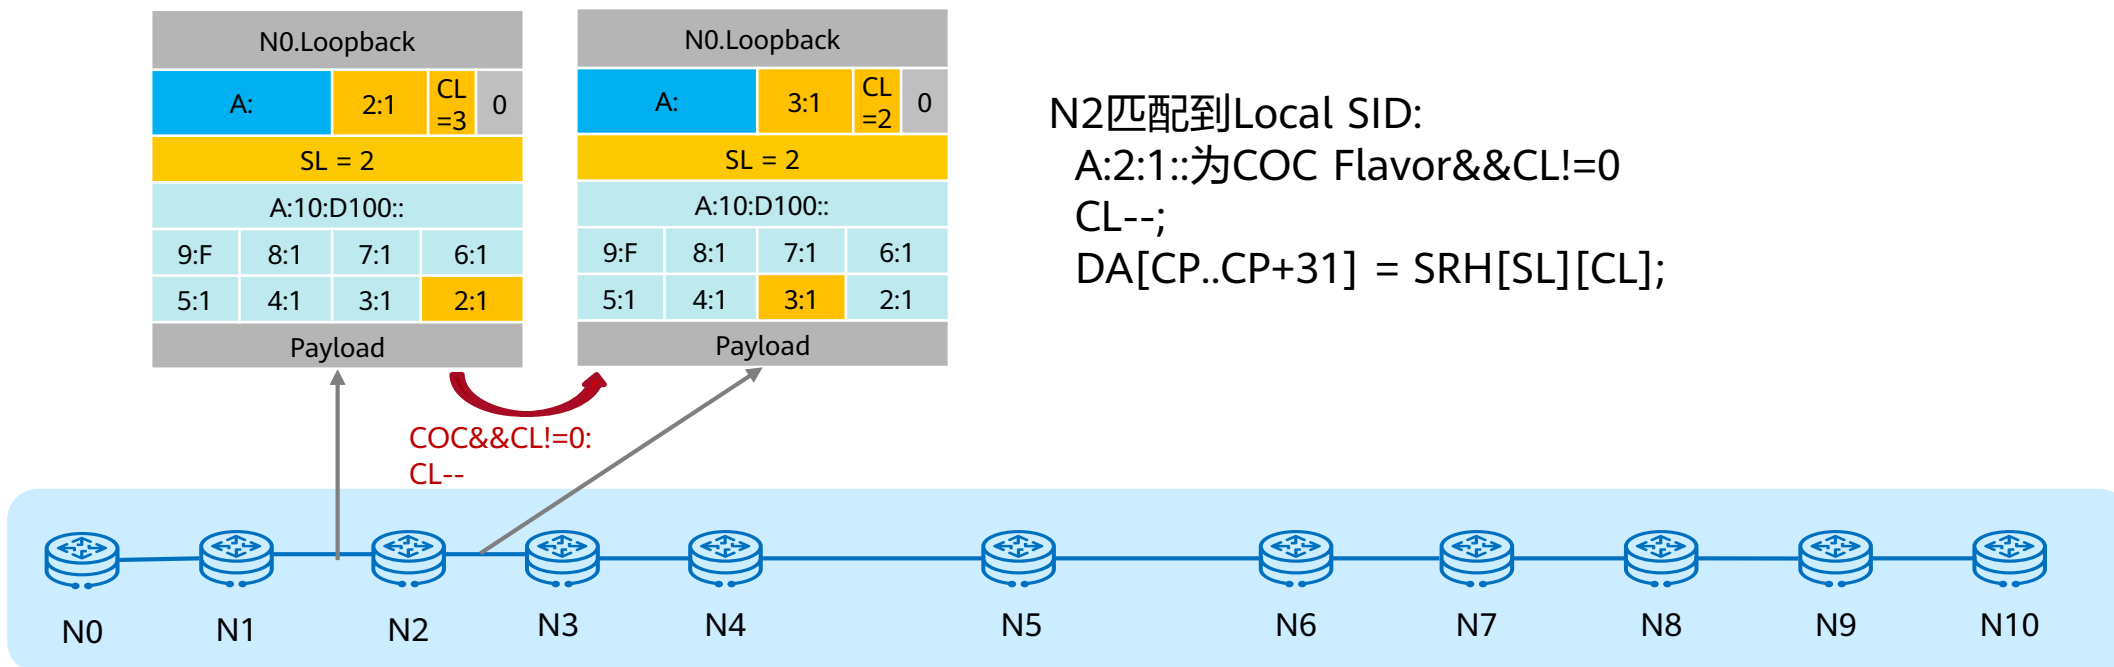
A指代一个64bits的前缀

N2匹配到Local SID:

A:2:1::为COC Flavor&&CL!=0

CL--;

DA[CP..CP+31] = SRH[SL][CL];



# G-SRv6转发示例 — 纯压缩场景4

SID List: 严格显示路径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:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

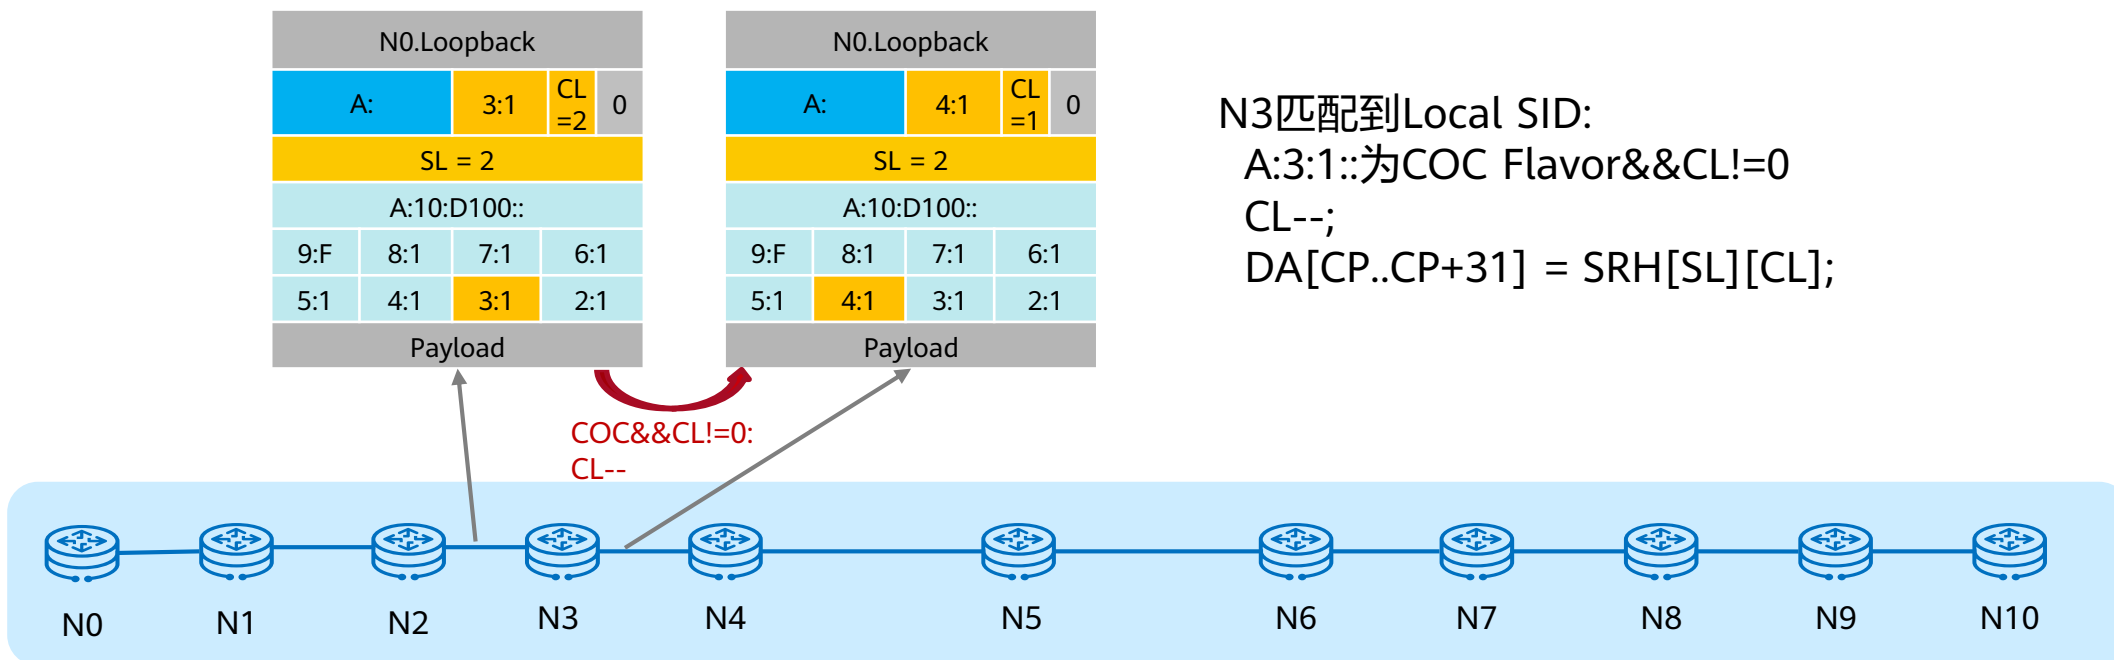
A指代一个64bits的前缀

N3匹配到Local SID:

A:3:1::为COC Flavor&&CL!=0

CL--;

DA[CP..CP+31] = SRH[SL][CL];





# G-SRv6转发示例 — 纯压缩场景5

SID List: 严格显示路径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:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1


A指代一个64bits的前缀

N4匹配到Local SID:


A:4:1::为COC Flavor && CL!=0

CL--;

DA[CP..CP+31] = SRH[SL][CL];



COC && CL!=0:  
CL--



# G-SRv6转发示例 — 纯压缩场景6

SID List: 严格显示路径10 SIDs:

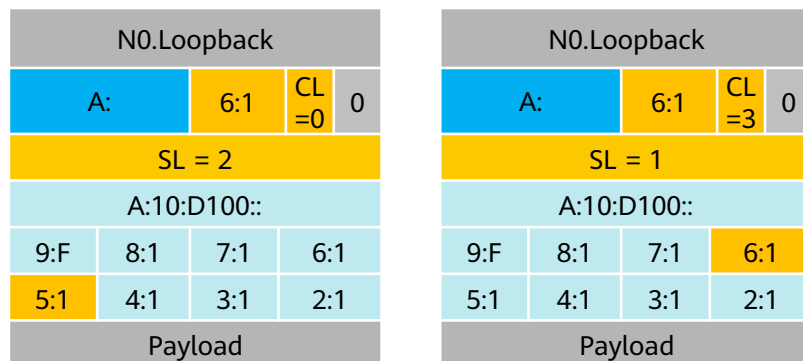
- A:1:1::, A:2:1::, A:3:1::, A:4:1::, A:5:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀




N5匹配到Local SID:

A:5:1::为COC Flavor && CL==0

SL--;CL=3;

DA[CP..CP+31] = SRH[SL][CL];

COC && CL=0:  
SL--;CL=3



# G-SRv6转发示例 — 纯压缩场景7

SID List: 严格显示路径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:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

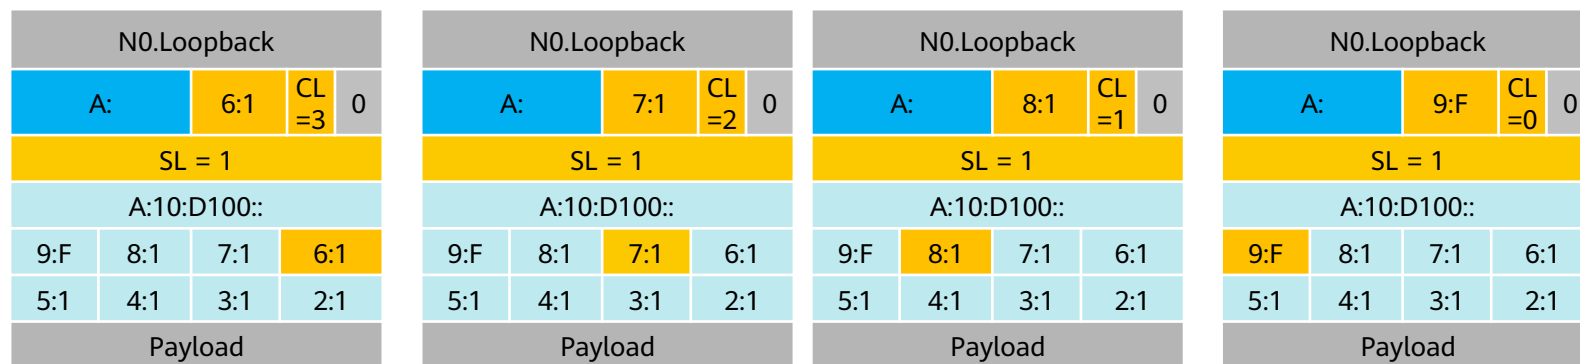
64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀


N6~N8匹配到Local SID:  
为COC Flavor & CL==0  
CL--;  
DA[CP..CP+31] = SRH[SL][CL];



COC & CL != 0:  
CL--

COC & CL != 0:  
CL--

COC & CL != 0:  
CL--




# G-SRv6转发示例 — 纯压缩场景8

SID List: 严格显示路径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:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- **A:10:D100::** 为N10的End.DT4 VPN SID

N9匹配到Local SID:  
A:9:F::为Non COC Flavor  
SL--;  
DA = SRH[SL];




64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀



Non COC:  
SL--;  
DA=128bits SID

# G-SRv6转发示例 — 纯压缩场景

SID List: 严格显示路径10 SIDs:

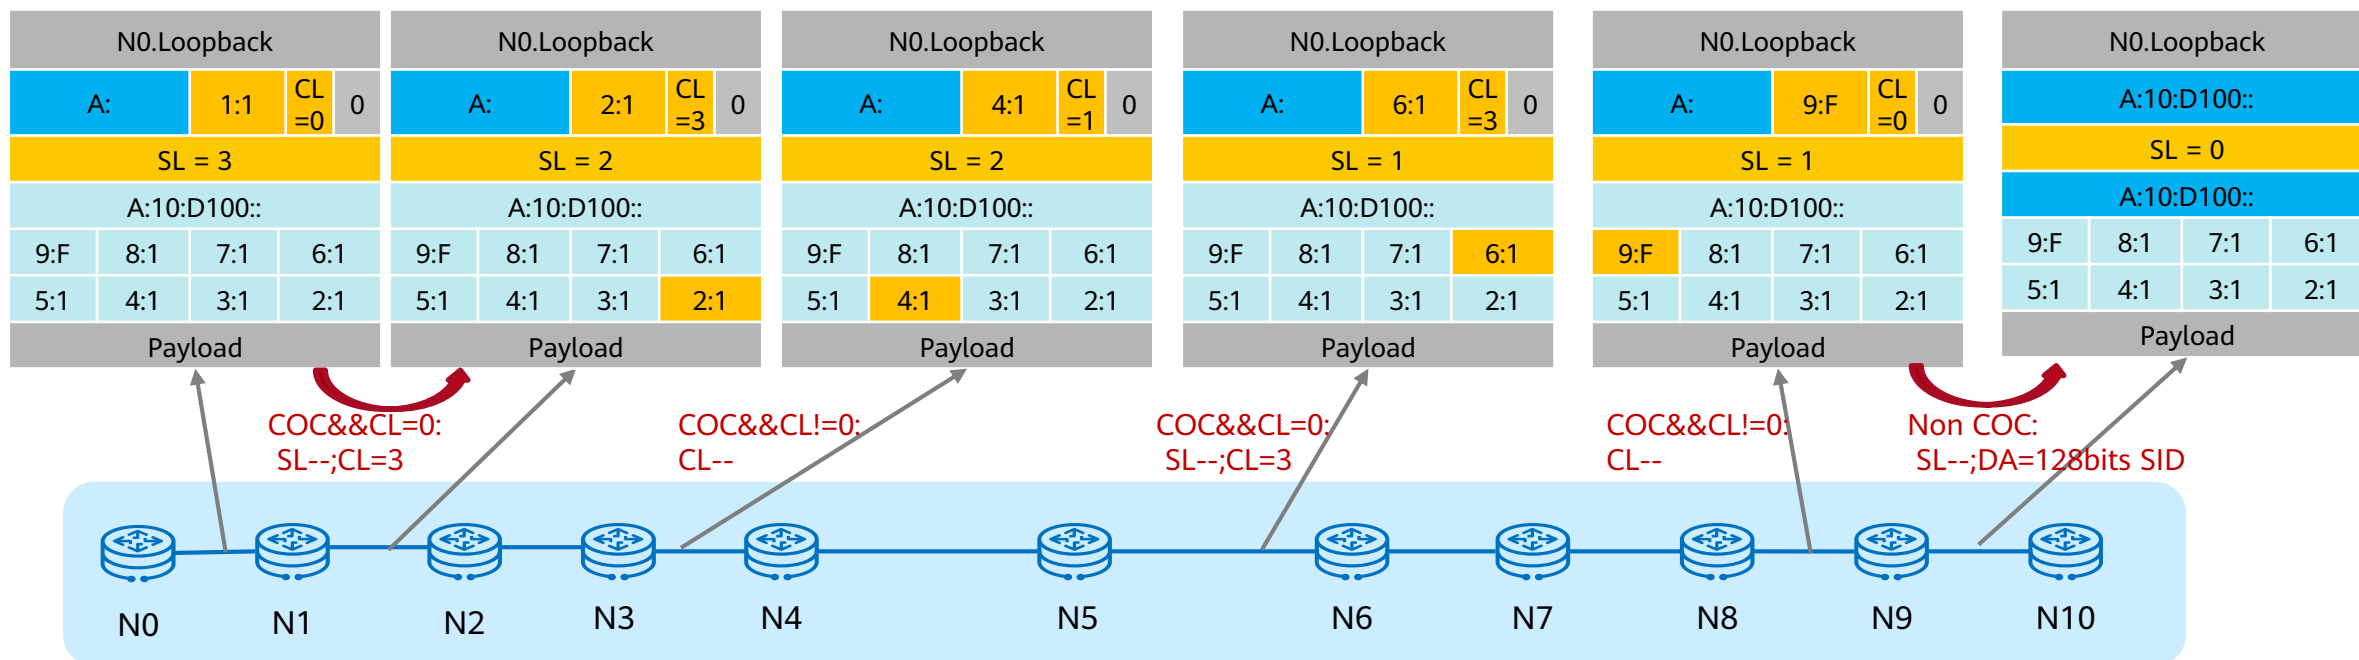
- A:1:1::, A:2:1::, A:3:1::, A:4:1::, A:5:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- A:9:F:: 为N9->N10的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀



# G-SRv6转发示例 — 压缩与非压缩混编场景1

SID List: 严格显示路径10 SIDs:

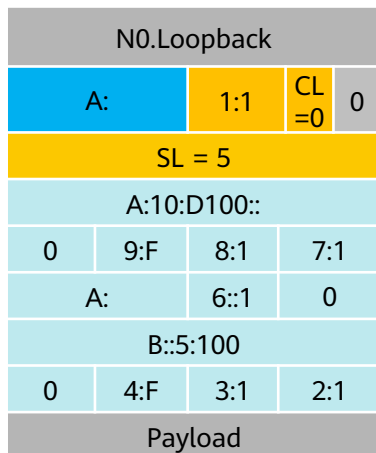
- A:1:1::, A:2:1::, A:3:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- B::5:100 为不支持压缩的N5节点分配的End.X SID
- A:4:F::, A:9:F:: 为N4与N9分配的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding




- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀



N0为头结点, 进行SRH编排封装:

- A:1:1::为第一跳SID, 进行Reduced操作, 封装到IPv6 DA中
- N2~N4的压缩SID (COC Flavor), 不足128bits, 补0对齐
- 不支持压缩的N5节点分片的SID
- Common Prefix + N6节点的压缩SID (COC Flavor), 不足128bits, 补0对齐
- N7~N8的压缩SID (COC Flavor) + N9的压缩SID (NO COC Flavor), 不足128bits, 补0对齐
- 尾节点128bits VPN SID



# G-SRv6转发示例 — 压缩与非压缩混编场景2

SID List: 严格显示路径10 SIDs:


- A:1:1::, A:2:1::, A:3:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- B::5:100 为不支持压缩的N5节点分配的End.X SID
- A:4:F::, A:9:F:: 为N4与N9分配的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀



# G-SRv6转发示例 — 压缩与非压缩混编场景3

SID List: 严格显示路径10 SIDs:

- A:1:1::, A:2:1::, A:3:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- B::5:100 为不支持压缩的N5节点分配的End.X SID
- A:4:F::, A:9:F:: 为N4与N9分配的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding

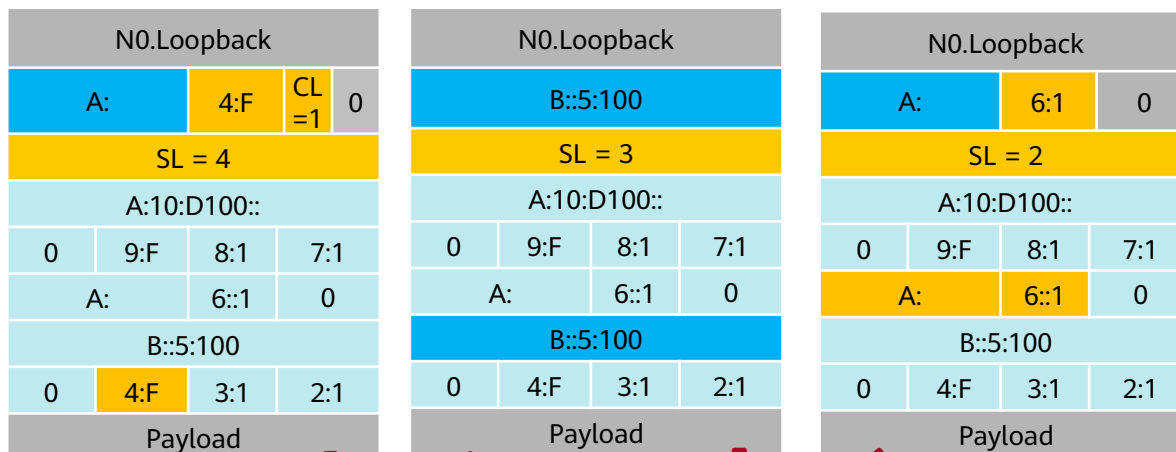


- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1


A指代一个64bits的前缀

N4匹配到Local SID:  
A:4:F::为Non COC Flavor  
SL--;  
DA = SRH[SL];

N5匹配到Local SID, 不支持压缩, 按照正常SRv6处理:  
SL--;  
DA = SRH[SL];



Non COC: SL--; DA=128bits SID  
Non SRv6 Compressed: SL--; DA=128bits SID





# G-SRv6转发示例 — 压缩与非压缩混编场景4

SID List: 严格显示路径10 SIDs:

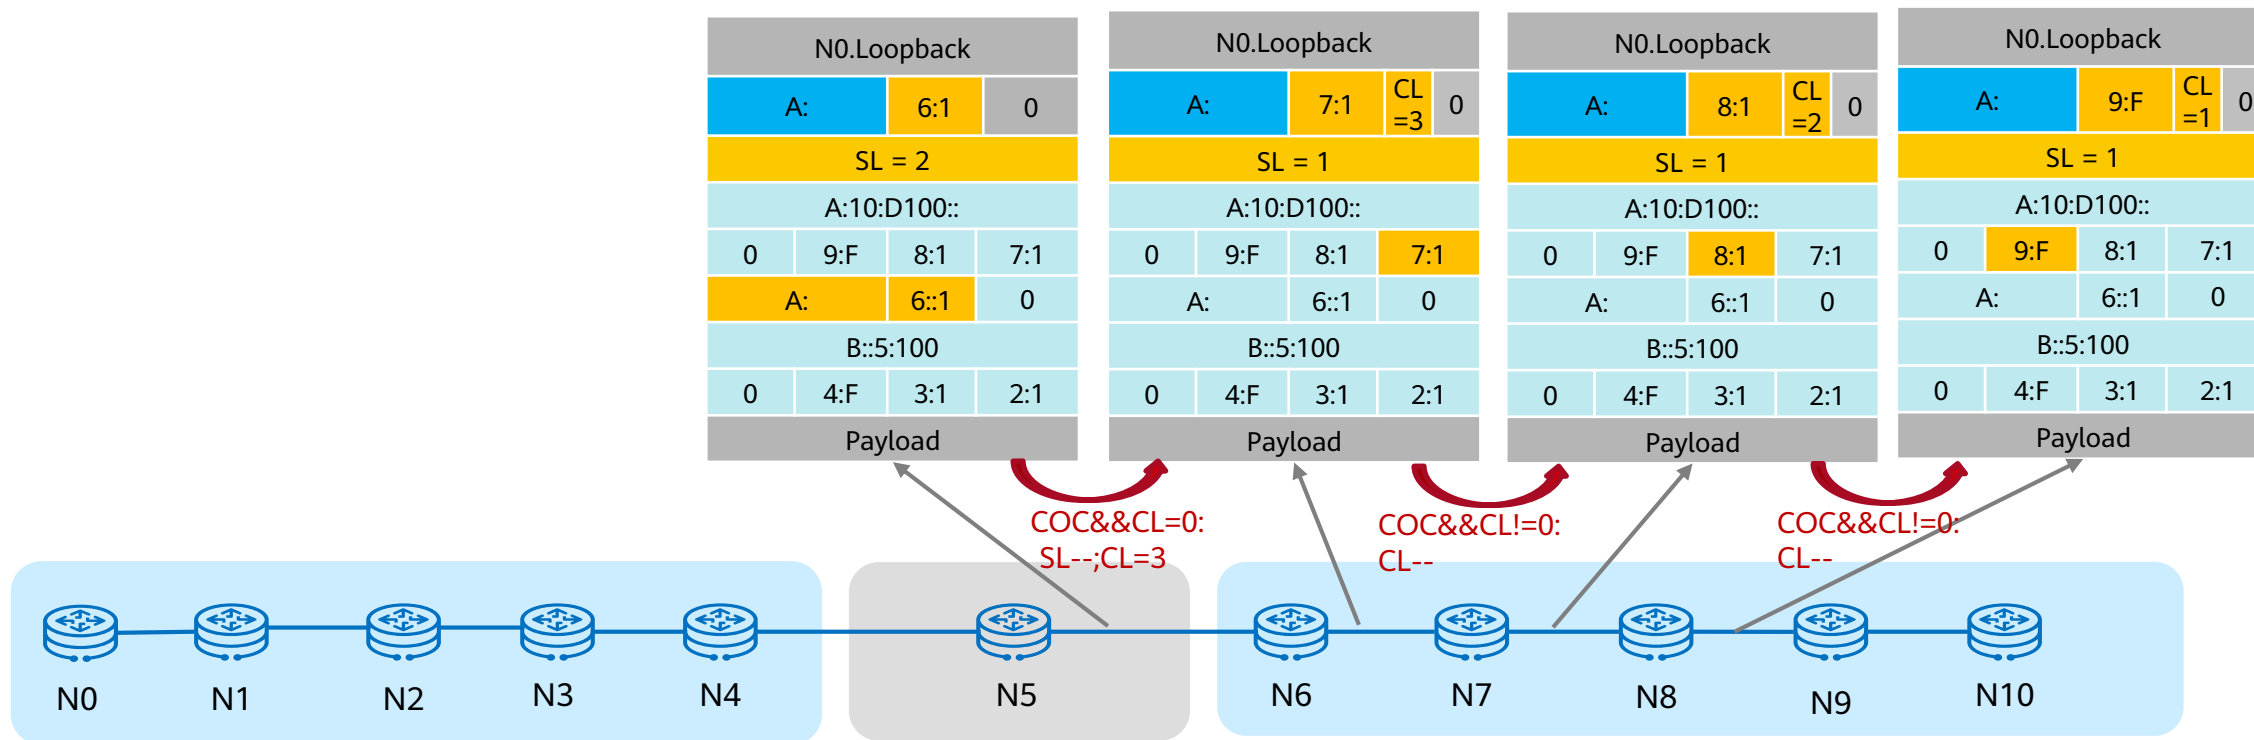
- A:1:1::, A:2:1::, A:3:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- B::5:100 为不支持压缩的N5节点分配的End.X SID
- A:4:F::, A:9:F:: 为N4与N9分配的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding



- Common Prefix A: /64
- C-SID: 1:1
- Locator A: 1/80
- Function 0x1

A指代一个64bits的前缀



# G-SRv6转发示例 — 压缩与非压缩混编场景5

SID List: 严格显示路径10 SIDs:

- A:1:1::, A:2:1::, A:3:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- B::5:100 为不支持压缩的N5节点分配的End.X SID
- A:4:F::, A:9:F:: 为N4与N9分配的End.X SID (without COC Flavor)
- A:10:D100:: 为N10的End.DT4 VPN SID

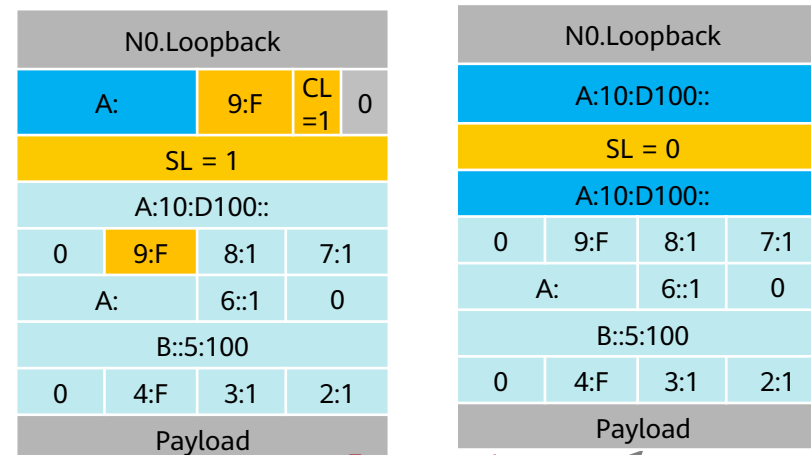
64 bits CP + 32 bits C-SID + 32 bits Padding




- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀

N9匹配到Local SID:  
A:9:F::为Non COC Flavor  
SL--;  
DA = SRH[SL];



Non COC:  
SL--;DA=128bits SID



# G-SRv6转发示例 — 压缩与非压缩混编场景

SID List: 严格显示路径10 SIDs:

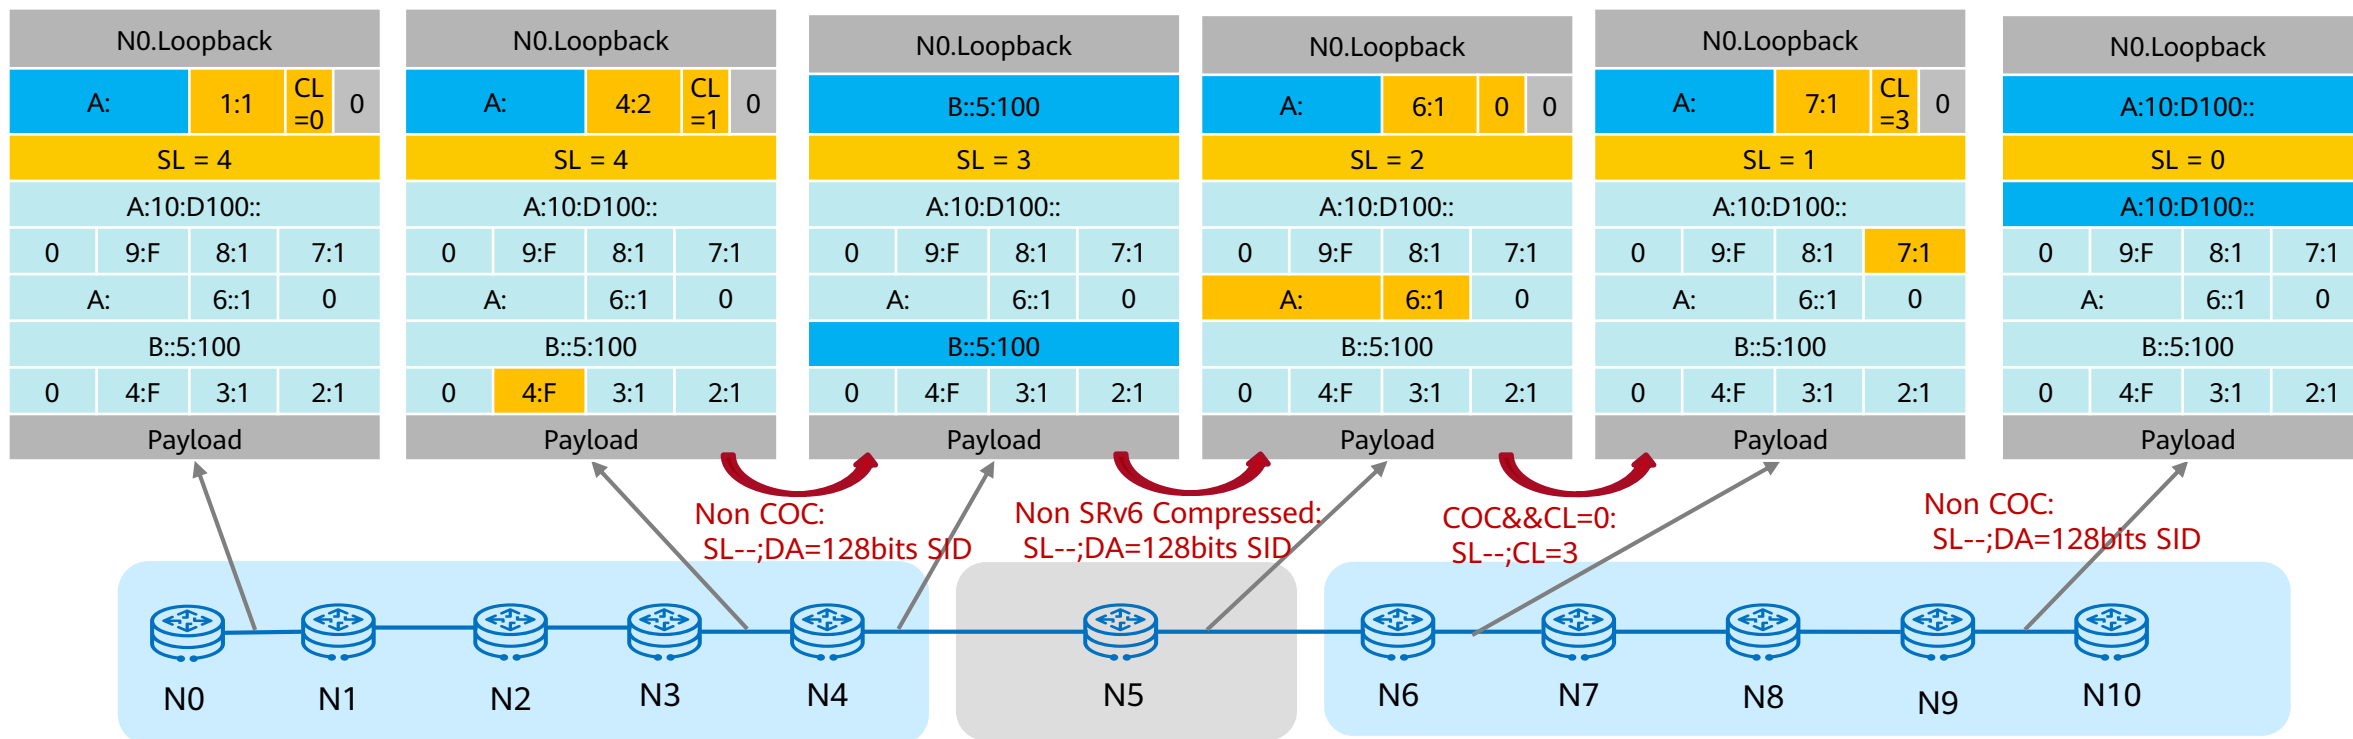
- A:1:1::, A:2:1::, A:3:1::, A:6:1::, A:7:1::, A:8:1:: 为各节点 End.X SID (with COC Flavor)
- B::5:100 为不支持压缩的N5节点分配的End.X SID
- A:4:F::, A:9:F:: 为N4与N9分配的End.X SID (without COC Flavor)
- **A:10:D100::** 为N10的End.DT4 VPN SID

64 bits CP + 32 bits C-SID + 32 bits Padding




- Common Prefix A::/64
- C-SID: 1:1
- Locator A::1/80
- Function 0x1

A指代一个64bits的前缀



# 方案对比总结：地址规划灵活，G-SRv6可少75%开销，10 SID转发性能提升55%+

对比项	子项	CRH	uSID	G-SRv6
压缩效率	压缩效率	4 SID/128bits 75%	(2~3 32bits uSID)/128bits /32 Block 66% /64 Block 50%	4 SID/128bits <b>75%</b> 与SR-MPLS一样
设备兼容	Native IP能力，路由可汇聚，跨域部署，故障隔离	SID无IP属性	保留	保留
	已部署的SRv6兼容	不兼容	兼容	兼容
	128bits/32bits SID混编	不兼容	支持同一个SRH中混编	支持同一个SRH中混编
地址规划	公网地址占用	节省	申请新的Address Block， <b>短前缀</b> 地址空间换压缩效率	节省，压缩与非压缩SID可共用地址空间， <b>零新增</b>
	地址类型	不限	<b>限制</b> 公网地址不足，只能使用私网地址	<b>不限</b>
	部署复杂度	IPv6地址+ Mapping SID (全局唯一)	<b>强制</b> 规划uSID前4bit 0-D为公网，E-F为私网，可能与现网规划冲突 <b>至少两次规划，甚至三次 (uSID,私网SRv6,公网SRv6)</b>	<b>兼容IPv6/SRv6地址规划</b> <b>无强制地址规划要求</b> <b>一次规划 (公网SRv6/G-SRv6)</b>



G-SRv6压缩后与SR-MPLS开销基本相同，支持公网/私网地址部署，一次地址规划，重用Locator

# SRv6压缩方案对比-地址规划

对比项	uSID	G-SRv6
地址规划	至少 <b>两次</b> 地址规划，可能需要 <b>三次</b> <ul style="list-style-type: none"><li>• ULA：原生SRv6域内SID（可合并到GUA）</li><li>• GUA：发布到域外的SRv6 SID</li><li>• ULA：用于压缩的域内uSID</li></ul>	<b>一次</b> 地址规划，兼容原有地址规划 <ul style="list-style-type: none"><li>• GUA：原生SRv6域内，发布到域外的服务SID以及G-SRv6 SID均可用一个Locator</li></ul>
	地址规划限制： <ul style="list-style-type: none"><li>• 最大只能用16bit描述节点，最多支持<b>56K</b></li><li>• 公网地址部署受限，<b>需使用私网地址部署</b></li><li>• 要求短前缀，才能支持较好压缩效果</li></ul>	地址规划灵活： <ul style="list-style-type: none"><li>• 可用20甚至更多bit描述节点，可支持<b>1M</b>以上节点，</li><li>• <b>公网地址均可部署，无限制</b></li><li>• 不要求短前缀，节省地址，前缀长度对压缩效果无影响</li></ul>

**条件：** uSID仅支持Block 48/， 更长前缀压缩效果差，仅能使用ULA部署。

**结论：** G-SRv6地址规划兼容现网，无需多套地址规划，运维简单。对地址前缀长度无要求，节省地质，更易部署维护

# G-SRv6: 产业生态支持好, 12+芯片商/设备商/运营商支持

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

Versions: [00\\_01](#)

SPRING Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: February 15, 2021

Z. Li  
C. Li  
Huawei Technologies  
W. Cheng  
China Mobile  
C. Xie  
C. Li  
China Telecom  
H. Tian  
F. Zhao  
CAICT  
August 14, 2020

Generalized Segment Routing Header  
draft-1c-6man-generalized-srh-01

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

SPRING Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: November 21, 2020

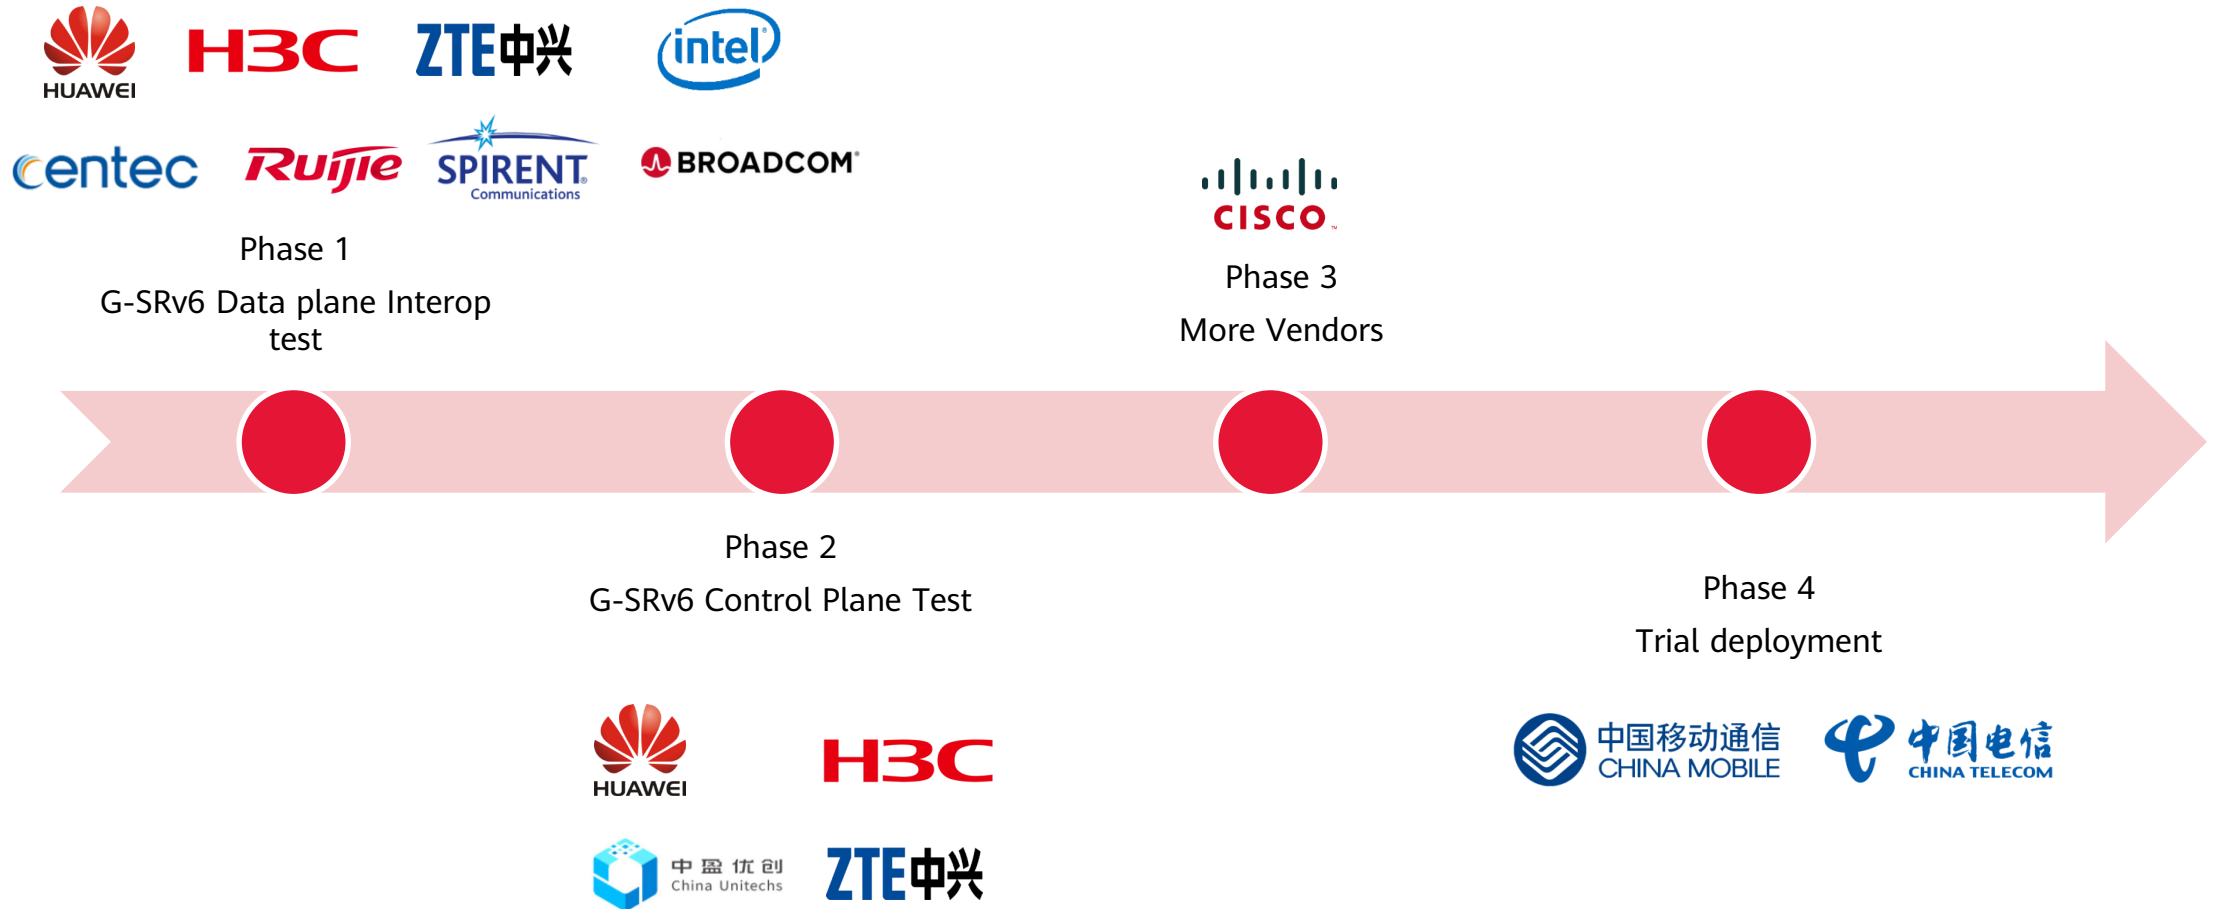
W. Cheng  
China Mobile  
Z. Li  
C. Li  
Huawei Technologies  
F. Clad  
Cisco Systems, Inc  
A. Liu  
ZTE Corporation  
C. Xie  
China Telecom  
Y. Liu  
China Mobile  
S. Zadok  
Broadcom  
May 20, 2020

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



10+设备商已完成G-SRv6互通, 11+设备商支持, 8+运营商/行业客户支持

# G-SRv6进展与计划：12家Vendor支持，现网试点已完成



# G-SRv6资料： IPv6+平台G-SRv6视频， INFOCOM2021论文， SRv6书籍

## 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, luzhibo, 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

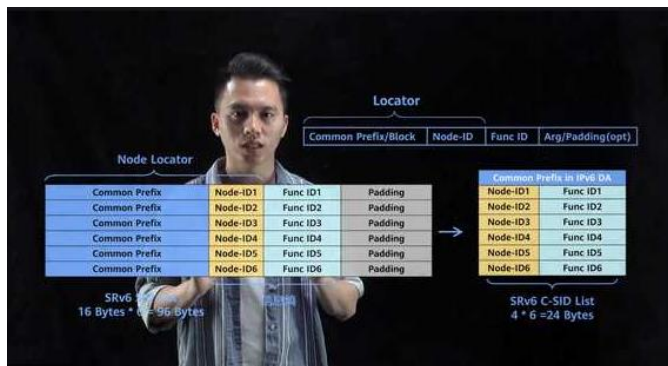

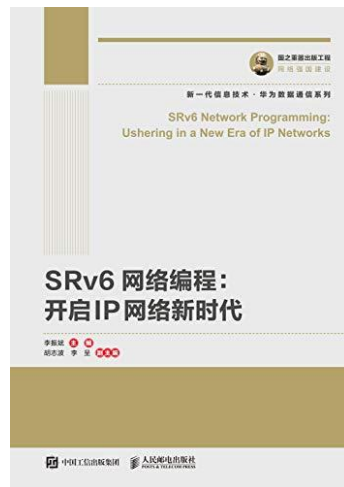
### I. INTRODUCTION

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 923 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 reduce the cost and generate more revenue in the 5G area.



**INFOCOM 2021:** <https://www.youtube.com/watch?v=Bl0r16XH4go>  
**IETF Hackathon:** <https://www.youtube.com/watch?v=qbryDg8fXRM>  
**信通院IPv6plus:** [https://mp.weixin.qq.com/s/PIrJyiVEQu2qIn6vIgY\\_Zw](https://mp.weixin.qq.com/s/PIrJyiVEQu2qIn6vIgY_Zw)  
**华为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/>  
**SRv6中英文书:** 《SRv6网络编程：开启IP网络新时代》第13章  
**IPv6+ G-SRv6:** <https://www.ipv6plus.net/Phase2/Generalized-SRv6/>



# 总结：G-SRv6高效压缩支持大网，兼容现网，节省地址，生态成熟



## 高效压缩

传输开销减少**50+%**  
转发速率提升**50+%**



## 节省地址

重用地址**无开销**  
**避免**地址重分配



## 高可扩展

32bit标识  
支持**超大规模**组网



## 平滑演进

**兼容**SRv6  
继承**全部**优点

# Thanks

Huawei Live

