C-SID Address Allocation

How to Assign SRv6 Locators to Network Nodes for deployment

Presenting by Bell Canada
daniel.voyer@bell.ca
IETF-119 SRv6ops, March 2024

Motivation

- C-SID Domain → C-SID network

  →Identified by a C-SID block, e.g., **fccc:cc00::/32**

- Each node in the C-SID network → 16-bit Global C-SID, e.g., **0005**

  →Configured in the form of SRv6 locator, e.g., **fccc:cc00:0005::/48**

- Total number of global C-SIDs ~ 65K

- Network divided into multiple ISIS Areas (different size)

- **Developed an allocation method based on**:
  1. Area Classification (Small/Medium/Large)
  2. Concept of C-SID Sets
  3. Information Encoding

- **Automated the allocation through the development of a python tool**

Question: how can we **allocate SRv6 locator prefixes** from the /32 C-SID block to **ISIS Areas** efficiently?

Naïve way: allocate /48 locators to nodes on a first come first serve basis regardless of the ISIS Area → Issue: can't summarize at Island boundary

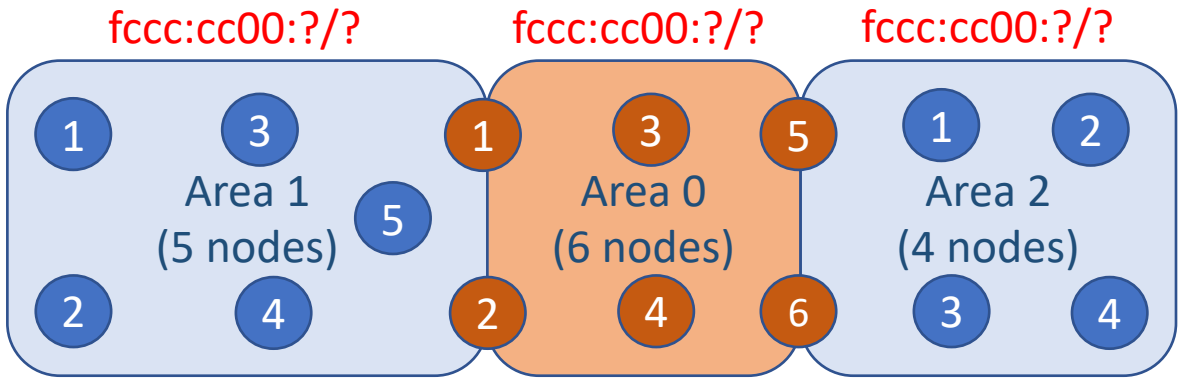

Outline

I. SRv6 C-SID Format Review
II. SRv6 C-SID Locator Format
III. Variables Encoded in the SRv6 Locator ("what is **encoded** in the **SRv6 Locator**?")
IV. Network Design Assumptions
V. C-SID Encoding Scheme ("how are the **variables encoded**?")
   - Flex Algo
   - SRv6 Site ID & Node ID
   - Sets -GIB & LIB
   - ISIS Level
VI. Sets to Srv6 Site ID Mapping ("how are **Sets mapped** to **SRv6 Sites**?")
   - Let's put this together – w/ Automated Sets Assignment in mind
VII. IPv6 loopback & BGP/ISIS Router ID Assignment

C-SID Format – Quick review

- Compressed SID. Can be of length 16/32 bits
- Configured on the nodes in the form SRv6 locators
- SRv6 Locator consists of two portions: C-SID block, and a C-SID.

  SRv6 Locator→ [**C-SID Block**]:[**C-SID**]::/**L**
- Four different formats defining the length of the uSID block & uSID:

  F1616: 16-bit C-SID block, 16-bit C-SID → **L** = 32

  F3216: 32-bit C-SID block, 16-bit C-SID → **L** = 48 "Industry Standard"

  F4816: 48-bit C-SID block, 16-bit C-SID → **L** = 64

  F3232: 32-bit C-SID block, 32-bit C-SID → **L** = 64
- **L** represents SRv6 locator length
- **F3216:** uSID Container with a **32 bit C-SID Block** + **16 bit C-SIDs** making up a TE path

  fccc:cc00:0001:0004:0002:0005:0003:0006/128
- Max. of 6 C-SIDs per container. More than 6 C-SIDs, use of SRH required.

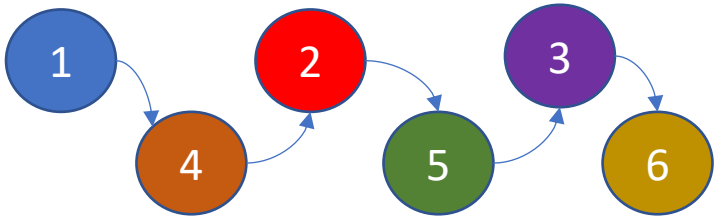

SRv6 C-SID Locator Format (F3216)

- At each SRv6 node, assign a SRv6 Locator (unique per node).
- **SRv6 locator** = **C-SID block** + **C-SID** (Global). F3216 32-bit C-SID Block, 16-bit C-SID. ~57K (Global C-SIDs)/nodes per C-SID block.
- **C-SID Block**: use **/24** SRv6 prefix to allow for information Encoding within the /32 uSID Block. e.g., **fccc:ccBB::/32.**
- **SRv6 prefix**: ULA /24.
  - **Lab**: fccc:cc::/24 (fc::/8)
  - **Prod:** fd01:00::/24 (fd::/8)
  - Adv. : make distinction between the two environments & be compliant with RFC 4193 (https://www.rfc-editor.org/rfc/rfc4193.html#section-3.2)
- Why ULA?
  - More space - can encode info in the C-SID block compared to GUA /32 prefix.
  - Have more /24 ULA prefixes compared to /32 GUA prefix.
  - More secured underlay - IPs can't be leaked over internet/across networks

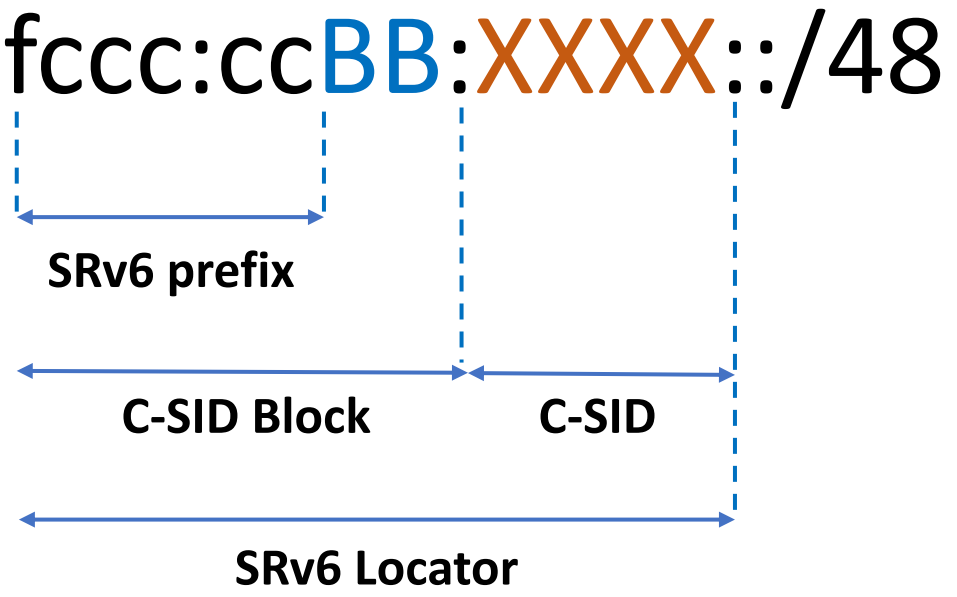

Network Arch. Assumptions

- *All ISIS Areas must lie within the same C-SID domain*→ inter-Area TE using a single C-SID container → C-SIDs coming from different Areas can be carried within the same C-SID container.

- *Max. Number of Algos to be supported*:|AlGo|≤ 4, (8 considering future expansion)

- *Max. Number of ISIS Areas*:|Area| ≤ 128 , (reality: we don't think to go beyond 99 areas)

- Max. Number of Large Areas ≤ 7

- *Max. Number of nodes/Large-Area* ≤ 1024, (a limit of 600 nodes)

- *Max. Number of nodes for ISIS level 2, Area 0* = ≤ 512

Variables Encoded in the SRv6 Locator

- F3216 C-SID:

  SRv6 locator = C-SID block (32 bits) + C-SID (16 bit),

  e.g., fccc:ccGG:XXXX::/48

- Variable encoded in the SRv6 locator
  - SRv6 Flexible Algorithm (Algo)
  - SRv6 Site ID (equivalent to ISIS Area ID)
  - Node ID
  - ISIS Level

For brownfield purposes, we ended up doing a mapping for SRv6 Site ID with ISIS level/”area ID”

C-SID Encoding: Flex Algo

- Flex. **A**lgo enables the use of user-defined link metrics and topology constraints (e.g. node/link exclusion) for IGP shortest path computations, e.g., **Algo 0** → uses default IGP metric; **Algo 128** → uses latency metric, **Algo 129** → uses default IGP metric & restricted links
- Encode the **A**lGo in the last two nibbles of the C-SID block

C-SID Block: fccc:ccGG::/32

  - Example:
    - **Locator 0:** fccc:cc00:*xxxx*:/48
    - **Locator 128:** fccc:cc01:*xxxx*:/48
    - **Locator 129:** fccc:cc02:*xxxx*:/48

  - Each **A**lgo represents a different C-SID domain. Only C-SIDs of the same **A**lgo could be carried by the same C-SID container/carrier.

C-SID Encoding: SRv6 Site ID & Node ID

- Encode the SRv6 Site ID & Node ID in the 16-bit C-SID.
- Encode SRv6 Site ID in the **first two** nibbles (8 bits) using the concept of Sets (**SS**), and the Node ID in the **last two** nibbles (8 bits).
- A Set is defined as a group of C-SID addresses.
  - Encoded in the first two nibbles of the C-SID.
  - **1 S**et contains **256** C-SIDs. Fixing the first two nibbles of the C-SID results in 256 unique C-SIDs
  - A SRv6 Site can be assigned **one**, **two** or **more** Sets depending on its size (number of nodes it contains)

C-SID: Set,Set, Node-id, Node-id (16-bit)

SRv6 Locator: fccc:ccGG:**SS**NN::/48

C-SID Encoding: Sets -GIB & LIB (1)

SRv6 Locator: fccc:ccGG:SSNN::/48

- Total number of Sets = 256 (2^8=256)
- Encode the **GIB** & the **LIB** within the sets (“how?” in next slide)
- What is GIB & LIB?
  - In SRv6, two types of C-SID (IETF Compress SRv6)
    - **Global C-SID** – globally unique across nodes. Associated with the uN function. Part of the C-SID Locator. Advertised by each node.
    - **Local C-SID** – locally unique within the node (i.e. same C-SID can be used across nodes). Associated with local functions on the node (RFC8986). Always preceded by a global C-SID in the C-SID container. Not Advertised.

      SRv6 Container: fccc:ccGG:[**Global_C-SID**:**Local_C-SID**:..:0000]/128

  - **GIB** (Global ID Block): Block containing all Global C-SIDs.
  - **LIB** (Local ID Block): Block containing all Local C-SIDs/node.

C-SID Encoding: Sets -GIB & LIB (2)

fcccc:cc**BB**:**SS**::/40

- How do we encode GIB & LIB within the Sets?
- Reserve the last **32** sets ({E0-FF}) for the **LIB** & the remaining **224** ({00-DF}) sets for the **GIB**

$S \in \{00, 01, 02, \ldots\ldots FF\}$ → Total of **256** Sets

$S\_GIB \in \{\mathbf{0}0, \mathbf{0}1, \mathbf{0}2, \ldots\ldots \mathbf{D}F\}$ → **224** Sets

$S\_LIB \in \{\mathbf{E}0, \mathbf{E}1, \mathbf{E}2, \ldots, \mathbf{F}F\}$ → **32** Sets

- GIB & LIB automatically encoded in the first nibble of the Set ID. Observing 0xE-0xF in the first nibble of the C-SID indicates a Local C-SID, and Global C-SID otherwise.

- Total number of Global C-SIDs in GIB= 224 x 256 = 57,344 ≅57 K C-SIDs (nodes)
- Total number of Local C-SIDs in LIB= 32 x 256 = 8,192 ≅8 K C-SIDs

C-SID Encoding: ISIS Level

- ISIS Level is automatically encoded in the SRv6 Site ID (**SRv6 Site ID=0** represents **level 2 Area 0**, and **SRv6 Site ID≠0** represents **level 1 ISIS Areas**)

- **SRv6 Locator:** fccc:ccGG:**SS**NN::/48

- Assign set **00** & **01** from 224 sets for Level 2 Area 0

  SRv6 Locator: fccc:ccGG:**00**NN::/48

  SRv6 Locator: fccc:ccGG:**01**NN::/48

  Locator Summary Prefix: fccc:ccGG::/32 -> for ISIS level 2

- For the remaining Areas, Sets assignment & SRv6 Site ID Mapping explained next

Sets to SRv6 Site ID Mapping

Two Questions

- How are Sets assigned to SRv6 Sites?
- How can we infer the SRv6 Site idx from the Set ID advertised in the Locator's Summary Prefix

E.g., SRv6 Site 1 → [2]

Locator Summary Prefix: fccc:ccGG:0200:/40

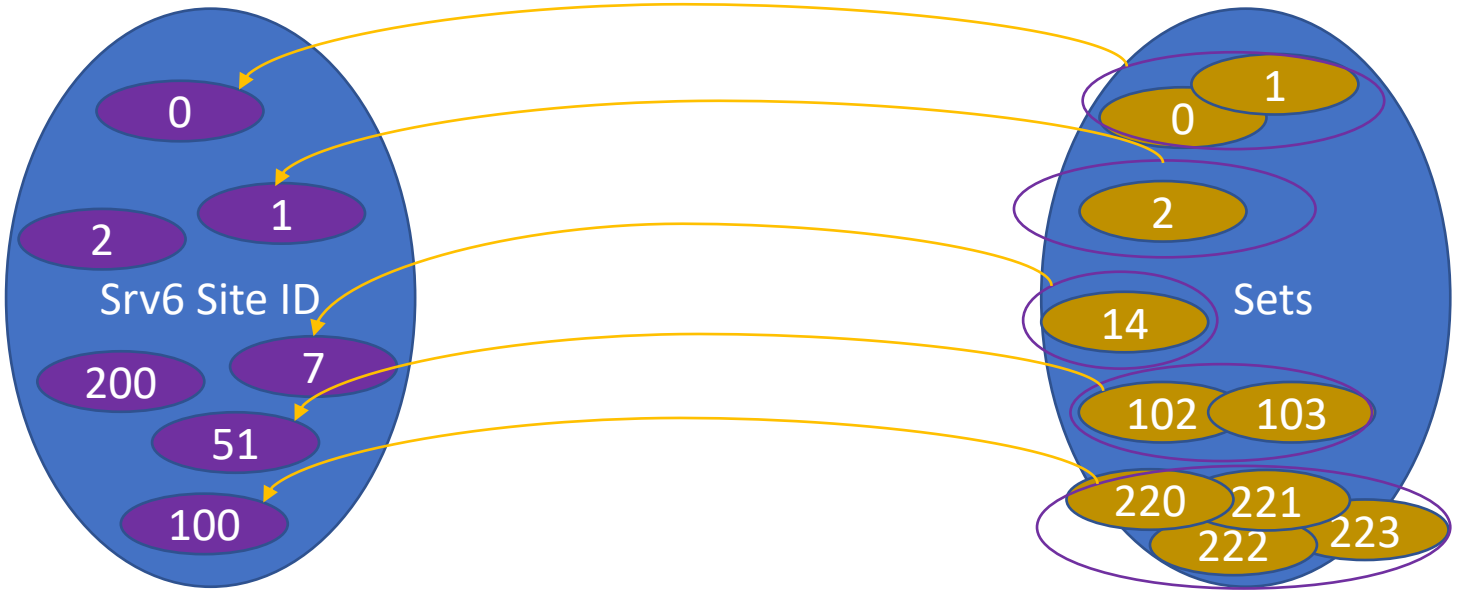

Let's put this together w/ Automated Sets Assignment in mind

- Developed a python tool that automates the Sets to SRv6 Site assignment
- **Tool Input** : Sites ISIS Level, Max. number of nodes, SRv6 Site ID (proposes one)
- **Tool Output**: Allocated Sets, & Assigned Locator Summary Prefixes

ISIS Level	Site Size *(no. of nodes)*	Srv6 Site ID (ISIS area ID)	Allocated Sets (dec)	Locator 0 Summary Prefix	Locator 128 Summary Prefix	Locator 129 Summary Prefix
2	300	**0**	**0,1**	fccc:cc00::/32	fccc:cc01::/32	fccc:cc02::/32
1	100	**1**	**2**	fccc:cc00:0200::/40	fccc:cc01:0200::/40	fccc:cc02:0200::/40
1	500	**2**	**4, 5**	fccc:cc00:0400::/39	fccc:cc01:0400::/39	fccc:cc02:0400::/39
1	1000	**110**	**220,221,222,223**	fccc:cc00:dc00::/38	fccc:cc01:dc00::/38	fccc:cc02:dc00::/38

2 Consecutive Sets

4 Consecutive Sets

IPv6 Loopback, ISIS NET-ID & BGP/ISIS Router ID

- **IPv6 Loopback:**

  SRv6_locator_0::1/128

  - Uses ULA addressing
  - Part of locator range - doesn't need to be advertised by IS-IS

- **ISIS NET-ID:**

  49.**ISIS_Area_ID**.**SRv6_locator_0**.00

  - 49—AFI
  - **ISIS_Area_ID**—Area ID
  - **SRv6_locator_0**—System identifier
  - 00—Selector

- **BGP Router ID:**
- For iBGP:

  6.0.**Set_id(dec)**.**Node_id(dec)**

  - 6 → ***IPv6-only*** node
  - 0→ Constant
  - **Set_id(dec)** → SRv6 Set index in decimal (first set assigned)
  - **Node_id(dec)**→ Node Index in decimal

Summary

Developing an IP address plan is a pivotal component for the deployment of network architecture and is essential for its subsequent daily operations.

The allocation of **C-SIDs must be unique** throughout the network and should be meticulously designed to ensure this uniqueness.

Charter Aspect Scope:

- **Deployment Scenarios:** This involves formulating strategies to address operational challenges within a variety of network settings and applications.
- **Key Deliverables:** Recommendations and guidance for IPv6 address planning for SRv6 SID, including SRv6 compression.