IETF 124 – Montreal
Nov 2025
srv6ops Working Group

draft-horn-srv6ops-srv6addressing-00
SRv6 Addressing

Jakub Horn, Kris Michielsen (Cisco Systems)
Nicklous Morris (Verizon)
Martin Horneffer (DT)
Clayton Hassen (Bell Canada)
Sen Li (China Mobile)

Motivation

- SRv6 addressing is fully flexible
- Addressing must be easily summarizable
- Same is valid if CSID-NEXT (uSID) compression is used
- Different addressing schemes can result in different compression efficiency
- We need to describe best common practices for addressing to avoid common pitfalls

Principles

- Simplicity
- Summarization
- **TE Efficiency**
- Focus on F3216 format
  - But applicable to any other format

Addressing Ranges

- ULA Unique Local FC00:/7
  - Most deployments
  - Very flexible
  - Inherent Security (Still Infrastructure ACL required)
- GUA Global Unicast
  - Assigned by authority (IANA , AFRNIC, APNIC, ARIN, RIPE)
  - Consumes significant portion of public address space
- 5F00::/16- RFC9602
  - Same properties as ULA

CSID Format

- Currently all deployments and interop uses “**F3216**”
  - **32** bits uSID block
  - **16** bits uSID
- Addressing principles are applicable to any other format

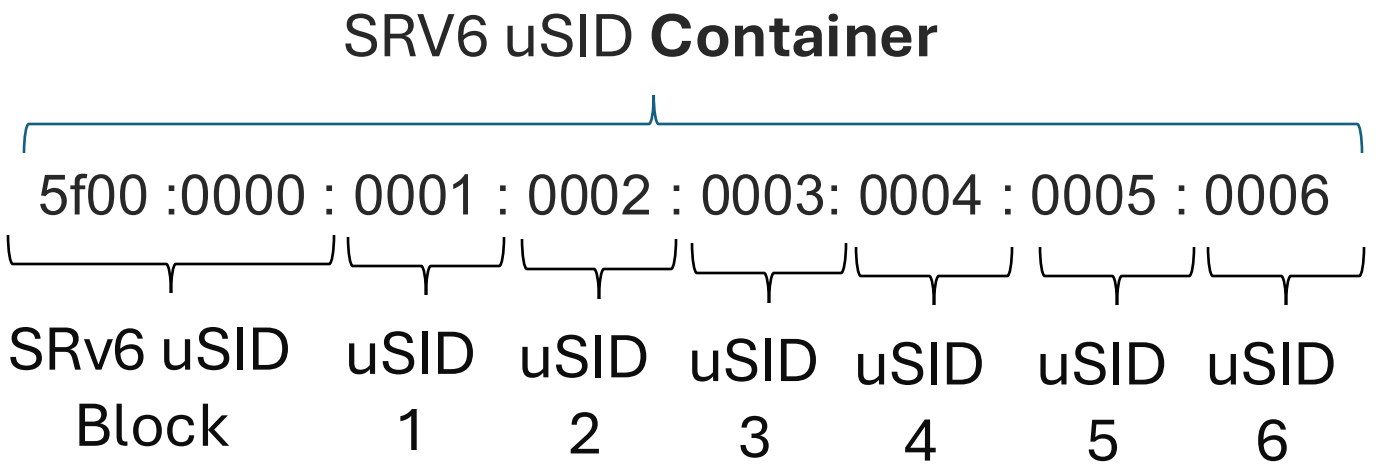

TE Efficiency

- NEXT CSID is the most efficient TE compression schema!
  - With proper addressing plan

Overhead

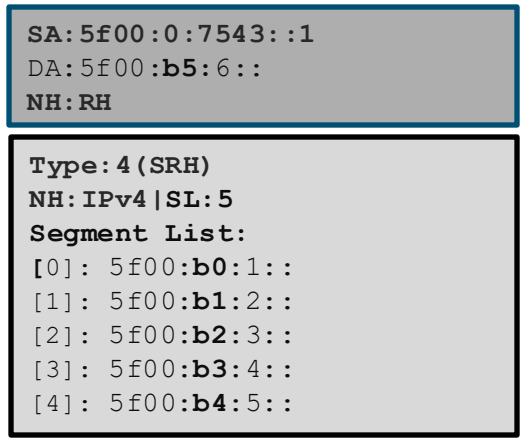

SIDs from different blocks!

5 Containers

Very Efficient

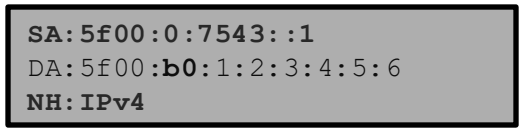

All SIDs from single block!

1 Container

Addressing plan must minimize number of blocks any TE policy will cross!!!

“Small Network” = Single Block

- How much devices fits single block?
- Each SRv6 device must have single uSID assigned
- uSIDs is16bits = 65535 unique values
- But default GIB/LIB split reduces it to 57k
- Summarization inefficiencies makes reality ~35k

SETs

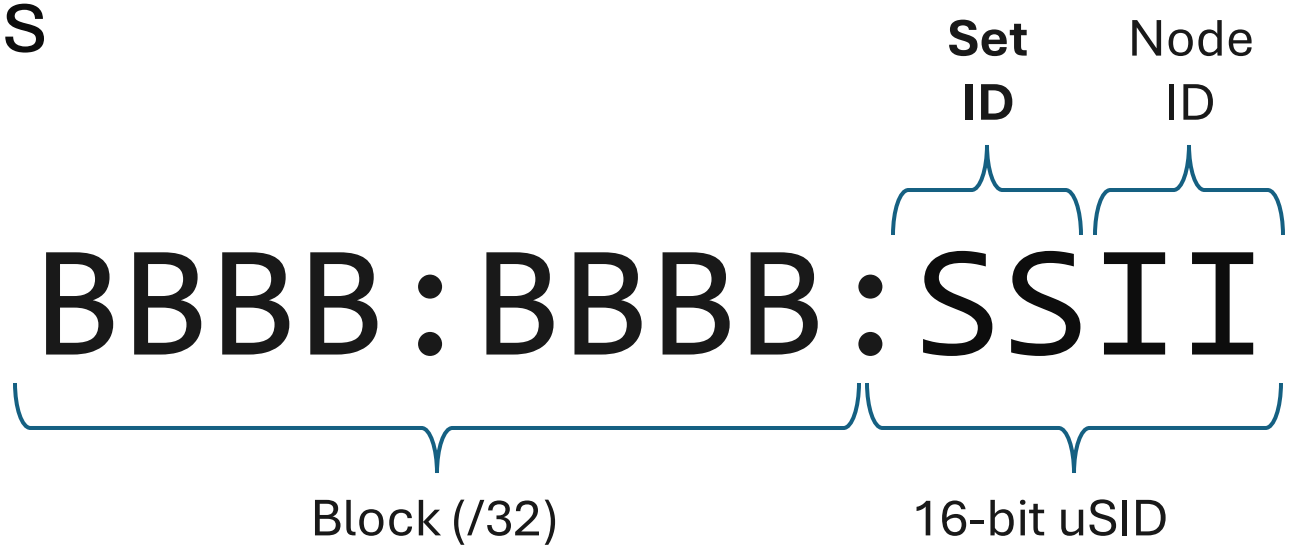

- Each block is split to certain number of equally sized Sets
  - /40 is very good example as it follows nibble boundary (/36 possible)
- Each area is assigned certain number of sets
- **Summarization is performed on set boundary only**
- Single block has 224 /40 sets (because of LIB)

Area addressing – Example with Sets

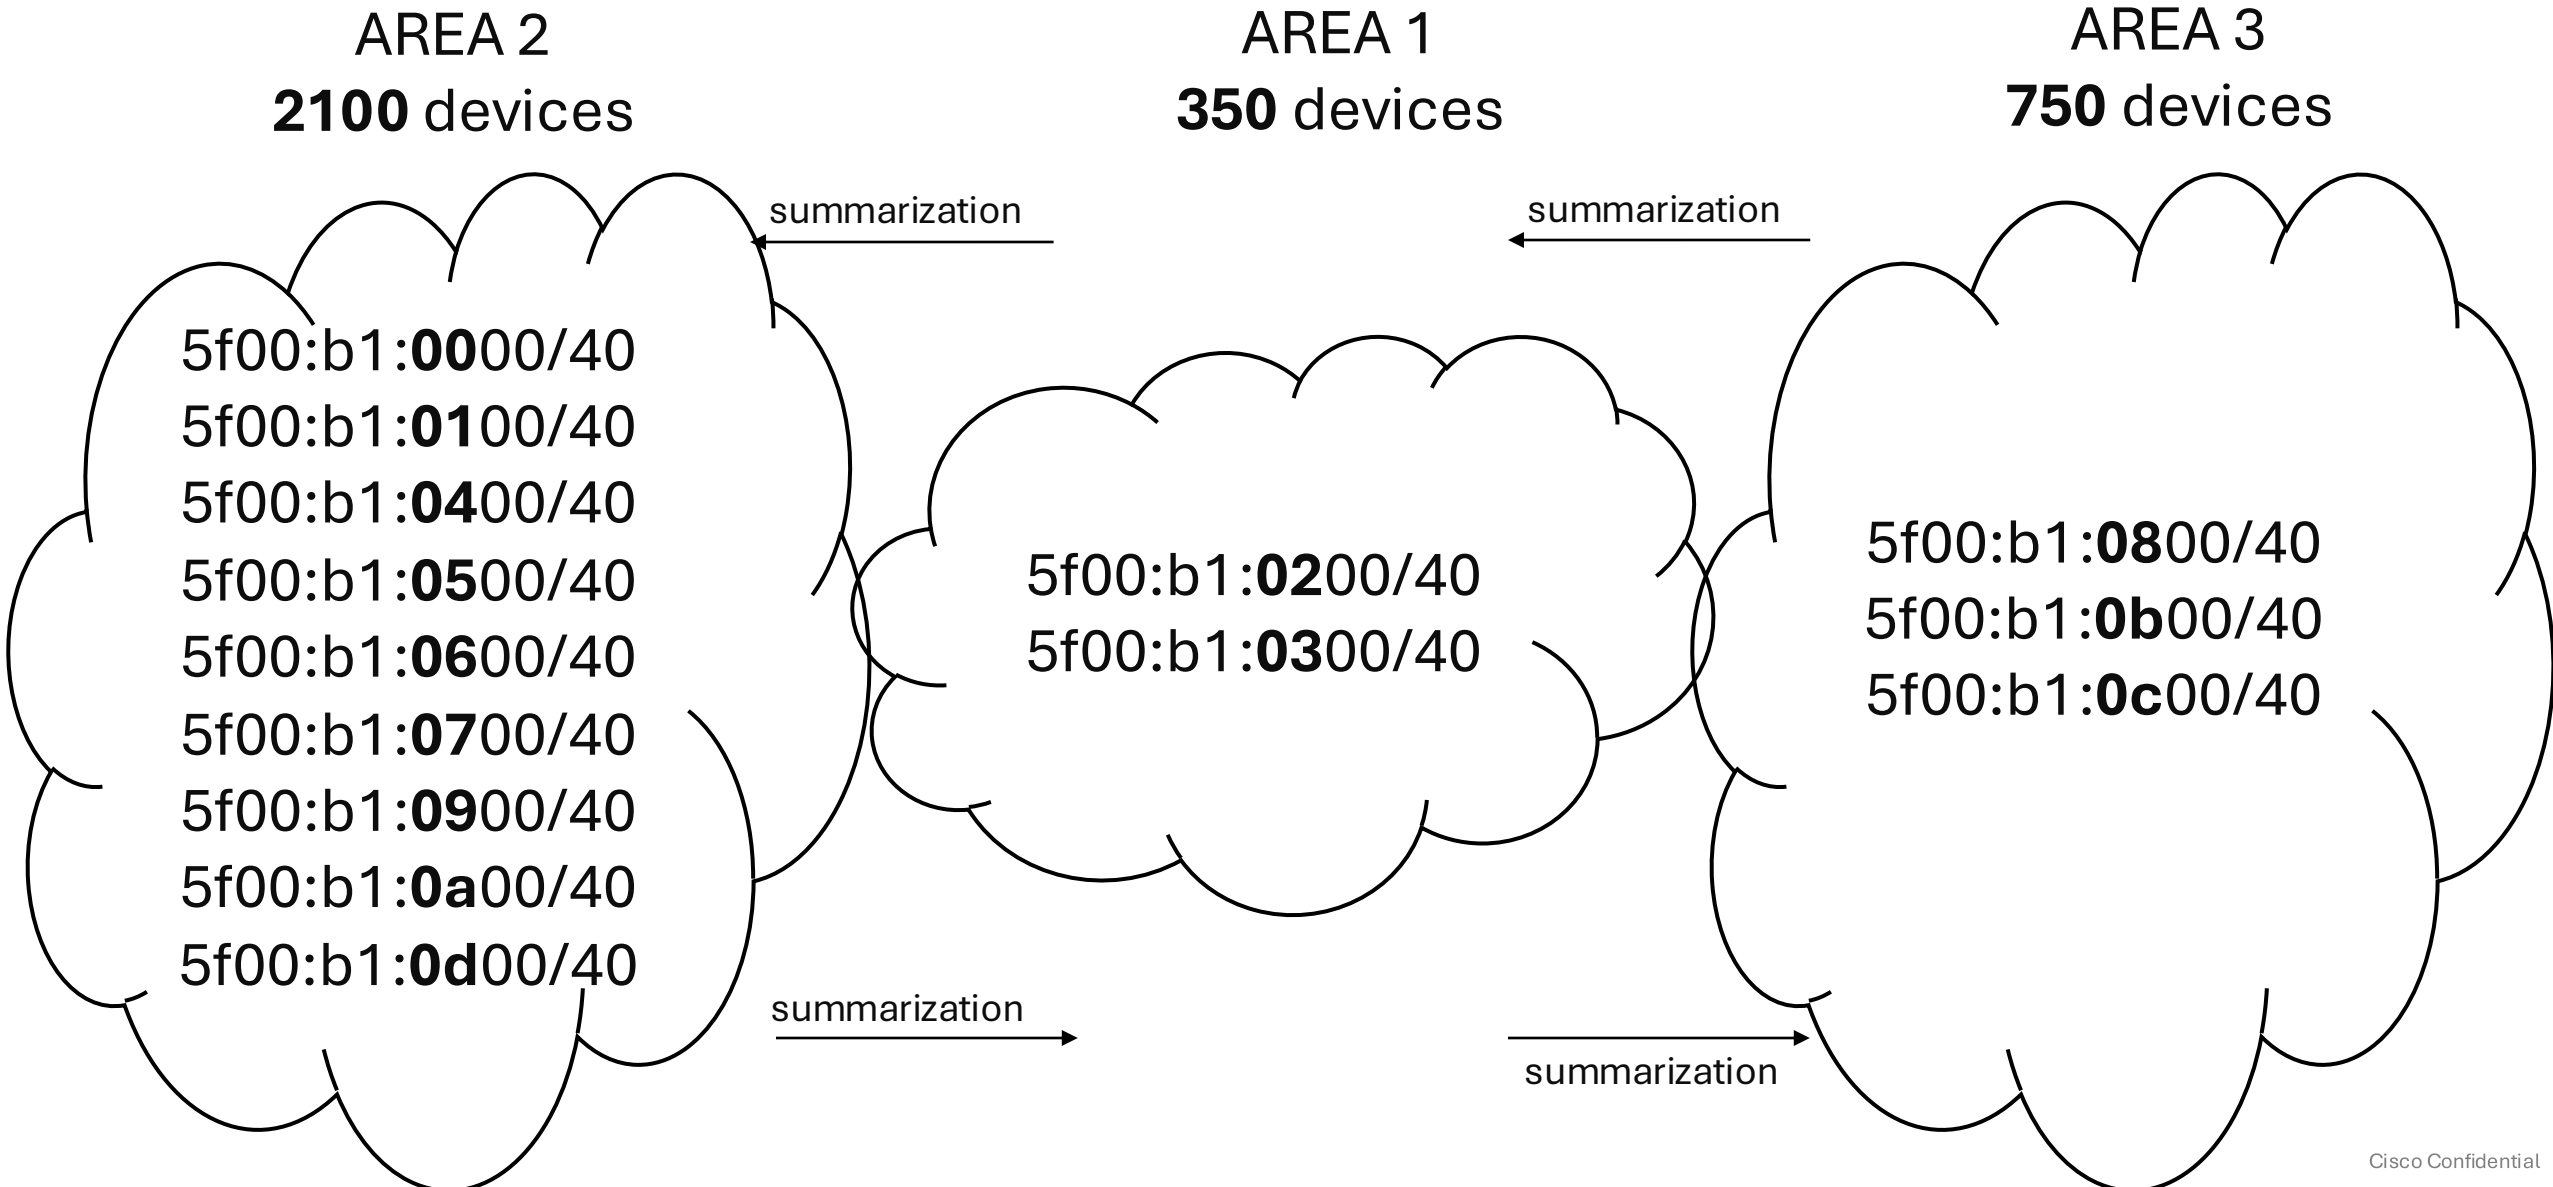

Flexible Algorithm locators

- Always use different block for different algorithm
  - Otherwise LIB is shared-> limited scalability
- Use same Set and Node ID across the algorithms
  - Improves Operational simplicity
  - Same summarization rules for all algorithms
- Number of algorithms in the network will be always limited
  - 1 nibble

BBBB:BBB**A**:SSII

Larger Networks

- Any network which cannot be addressed from single block – bigger than 35k
- We need multiple blocks- Block assigned to “Region”
- Number of regions should be size of network/35k
- Region must be encoded into the block
- Following nibble boundary
- Large Network `BBBB:BBA**R**:SSII – 16x35k=~500k devices`
- Very Large Network BBBB:BA**RR**:SSII 16x500k=8M devices
- Algorithm should be always encoded in higher order bits than Region

Loopbacks/Links

- Loopbacks
  - Independent
  - From locator
    - Improves Scalability
    - Simplifies operation
    - Simplifies summarization
- Links
  - Link Local only RFC7404
    - Simplifies operation
    - Remote reachability can be achieved using END.X of the interface

Asks

- Verify applicability to your network
- Provide feedback