SRv6 uSID in SONiC - AI Backend and Beyond

Ahmed Abdelsalam, Cisco Systems

Changrong Wu, Microsoft

IP: Best-Effort for > 40 years

- Destination Address only localize an endpoint
- Development of a myriad of per-domain shim layers to enrich IP
  - MPLS in the core
  - VxLAN in the DC
  - GTP for mobility
- Low scale
  - MPLS cannot summarize => BGP used as an IGP
- Poor End-to-End Services and High Operation Cost
  - Neither TE nor stateless service chaining in VxLAN
  - Translation gateways between domain (VxLAN/MPLS)

IPv6 becomes Programmable

- 128-bit DA becomes:
  - ordered list of uSID instructions (uSID)
  - uSID instruction can bound to anything .. Underlay TE, overlay, service chaining
  - 6 uSID instructions in DA
  - Most deployments requires less than 6 uSID
  - SRH is there for ultra-scale TE
- SRv6 uSID outperforms any IPv4 shim layer: MPLS, VxLAN, GTP

SRv6 uSID – Fully Standard

- SR Architecture – RFC8402
- IPv6 SR header – RFC 8754
- Network Programming – RFC 8986
- **SRv6 uSID – RFC 9800**
- BGP Services – RFC 9252
- ISIS - RFC9352
- OAM – RFC 9259
- Policy Architecture – RFC 9256
- PCEP – RFC 9603
- BGP-LS – RFC 9514
- Flex Algo – RFC 9350
- Performance Management – RFC 9503
- SRv6 SID Block – RFC 9602
- SR Policy Provisioning via BGP [link] – RFC XXXX
- SR Policy Advertisement in BGP-LS [link] – RFC XXXX
- BGP Color-Aware Routing [link] – RFC XXXX

SRv6 uSID - Rich Ecosystem

Network Equipment Manufacturers

Cisco, Huawei, Nokia, Ciena, Juniper Networks, ZTE, Arrcus, Ericsson, H3C

Merchant Silicon

Broadcom, Marvell, NVIDIA

Open-Source Applications

FRRouting, Cilium, GoBGP, ExaBGP, eBPF, Wireshark, Snort, Pyroute2, iptables SERA, GoBMP

SRv6 uSID

Open-Source Networking Stacks

SONiC, DENT, P4, OPEN Compute Project, OpenStack, Kubernetes, OvS Open vSwitch, VyOS

Smart NIC / DPU

Intel, NVIDIA

Partners

kaloom, Trend Micro, ENEA Qosmos Division, NoviFlow, ixia A Keysight Business, spirent Promise. Assured., Red Hat, F5, 6WIND, Keysight Technologies

SRv6 uSID in SONiC – Mature, Deployed, Rich Ecosystem

- Mature:
  - Since 2021: 350+ PR and 1600+ commits across SAI, SONiC, and FRR.
  - Static SRv6 uSID – Static Programming of SRv6 uSID from SDN controller “AI Backend”
  - L3VPN (IPv4/IPv6): Static & BGP (RFC9252)
  - Underlay Traffic Engineering: Static & IS-IS (RFC9352)
  - Static steering of IPv4/IPv6 traffic over SRv6 uSID
  - SRv6 SID Manager (F3216/GIB/LIB/WLIB)
- Deployed/Widely adopted:
  - Alibaba – deployed in production across all China
  - Microsoft – AI backend
- Rich Ecosystem:
  - Key contributors: Cisco, Microsoft, Alibaba, NVIDIA, Intel, Broadcom, LINE, and 6WIND

SRv6 uSID in SONiC – AI Backend

- Oct 2024 - May 2025 “SONiC 202505”
  - A total of 122 PRs have been merged across the SAI, SONiC, and FRR mainline repositories.
  - Cisco and Microsoft have been leading contributors, with 37 and 33 PRs respectively.
  - Alibaba has also made significant contributions with 23 PRs
  - NVIDIA has remained actively involved, contributing 15 PRs.
- SONiC 202505:
  - Static Programming of SRv6 uSID from SDN controller
  - Upgrade to FRR 10.3
  - SONiC Community blog: LINK

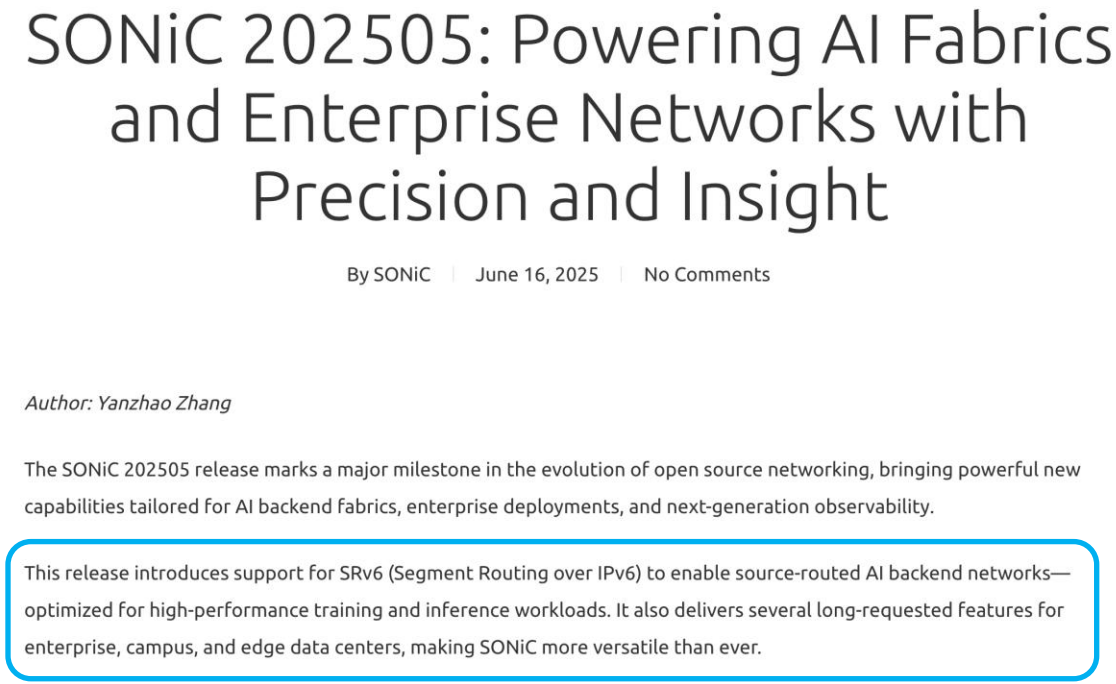

SONiC 202505: Powering AI Fabrics and Enterprise Networks with Precision and Insight

By SONiC | June 16, 2025 | No Comments

Author: Yanzhao Zhang

The SONiC 202505 release marks a major milestone in the evolution of open source networking, bringing powerful new capabilities tailored for AI backend fabrics, enterprise deployments, and next-generation observability.

This release introduces support for SRv6 (Segment Routing over IPv6) to enable source-routed AI backend networks—optimized for high-performance training and inference workloads. It also delivers several long-requested features for enterprise, campus, and edge data centers, making SONiC more versatile than ever.

SRv6 uSID – Deterministic Path Placement in AI Backend

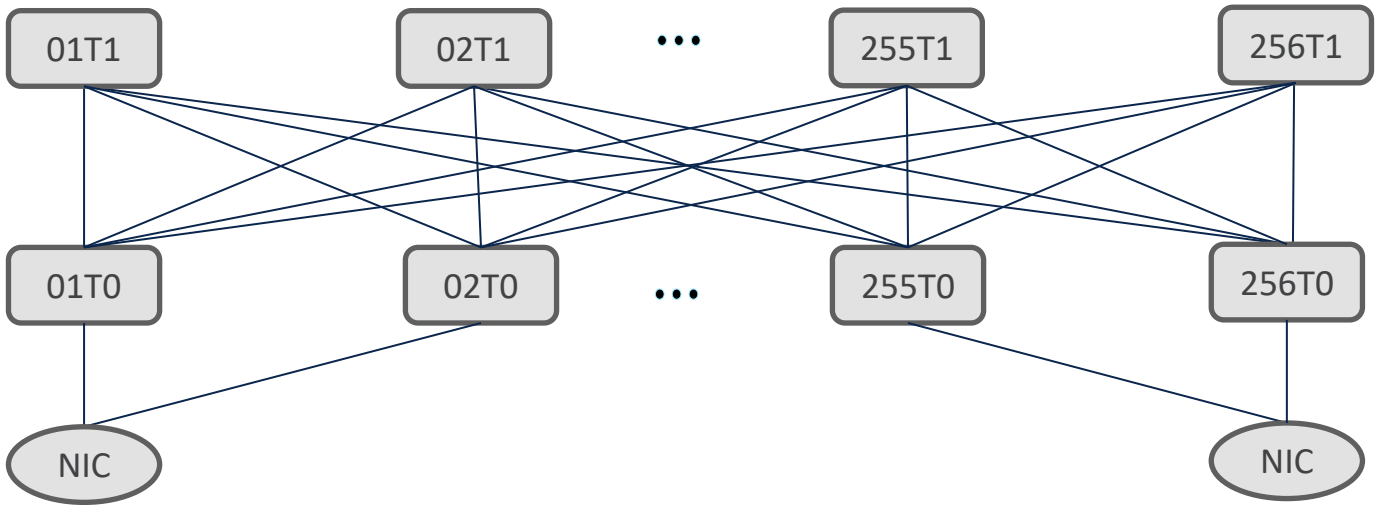

SRv6 uSID – Deterministic Path Placement in AI Backend

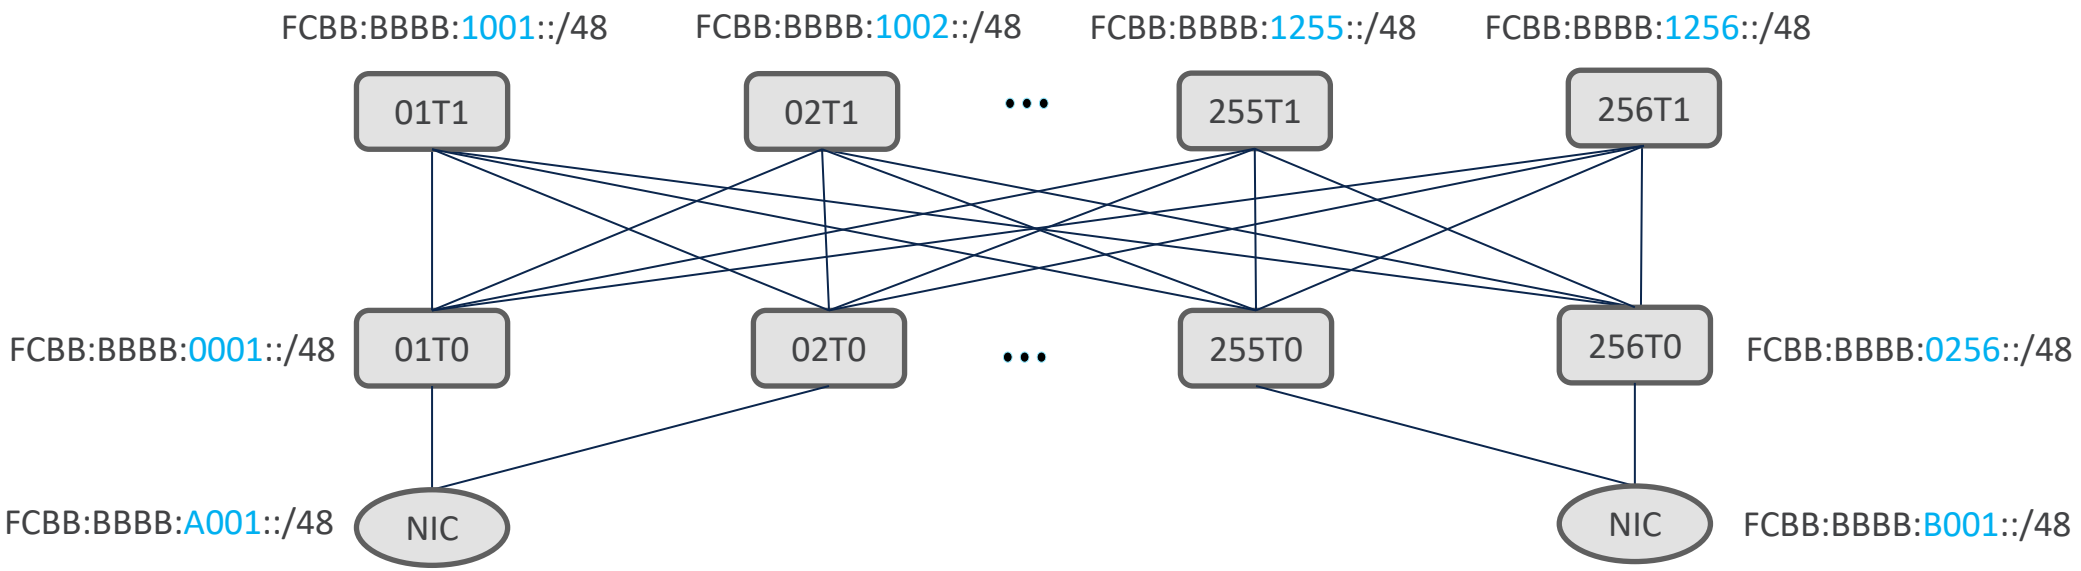

- uSID Block - FCBB:BBBB::/32
- uSID Locator - FCBB:BBBB:NNNN::/48
  - 01T0 => FCBB:BBBB:0001::/48

SRv6 uSID – Deterministic Path Placement in AI Backend

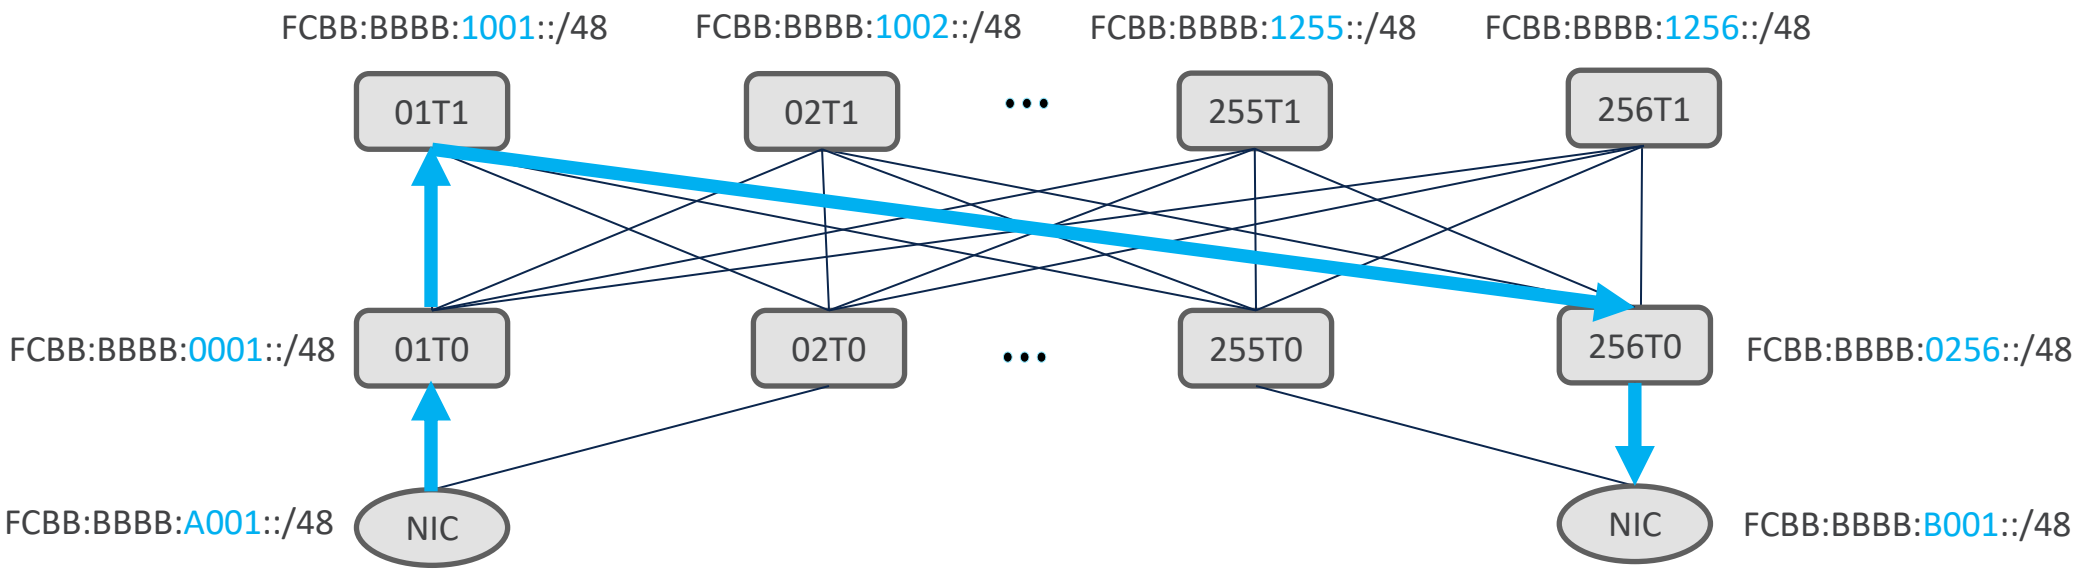

- NIC A to NIC B
  - Via FCBB:BBBB:0001:1001:0256:B001::
  - Go to 01T0 then 01T1 then 256T0 then NIC B

SRv6 uSID – Deterministic Path Placement in AI Backend

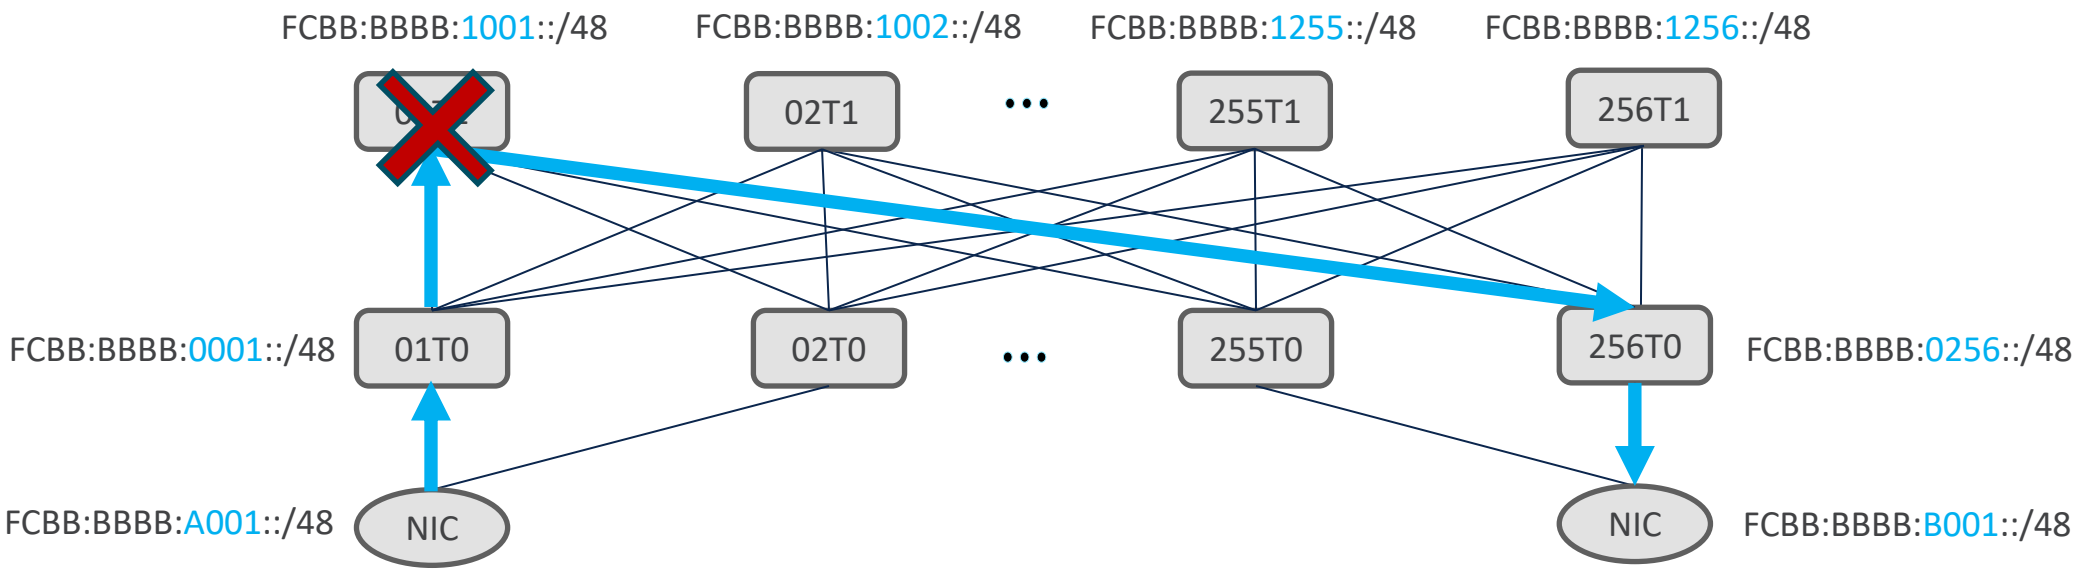

- NIC A to NIC B
  - Via FCBB:BBBB:0001:1001:0256:B001::
  - Go to 01T0 then 01T1 then 256T0 then NIC B

Congestion/Failure " 01T1"

- NIC A to NIC B
  - via FCBB:BBBB:0001:1002:0256:B001::
  - Go to 01T0 then 01T1 then 256T0 then NIC B

SRv6 uSID – Deterministic Path Placement in AI Backend

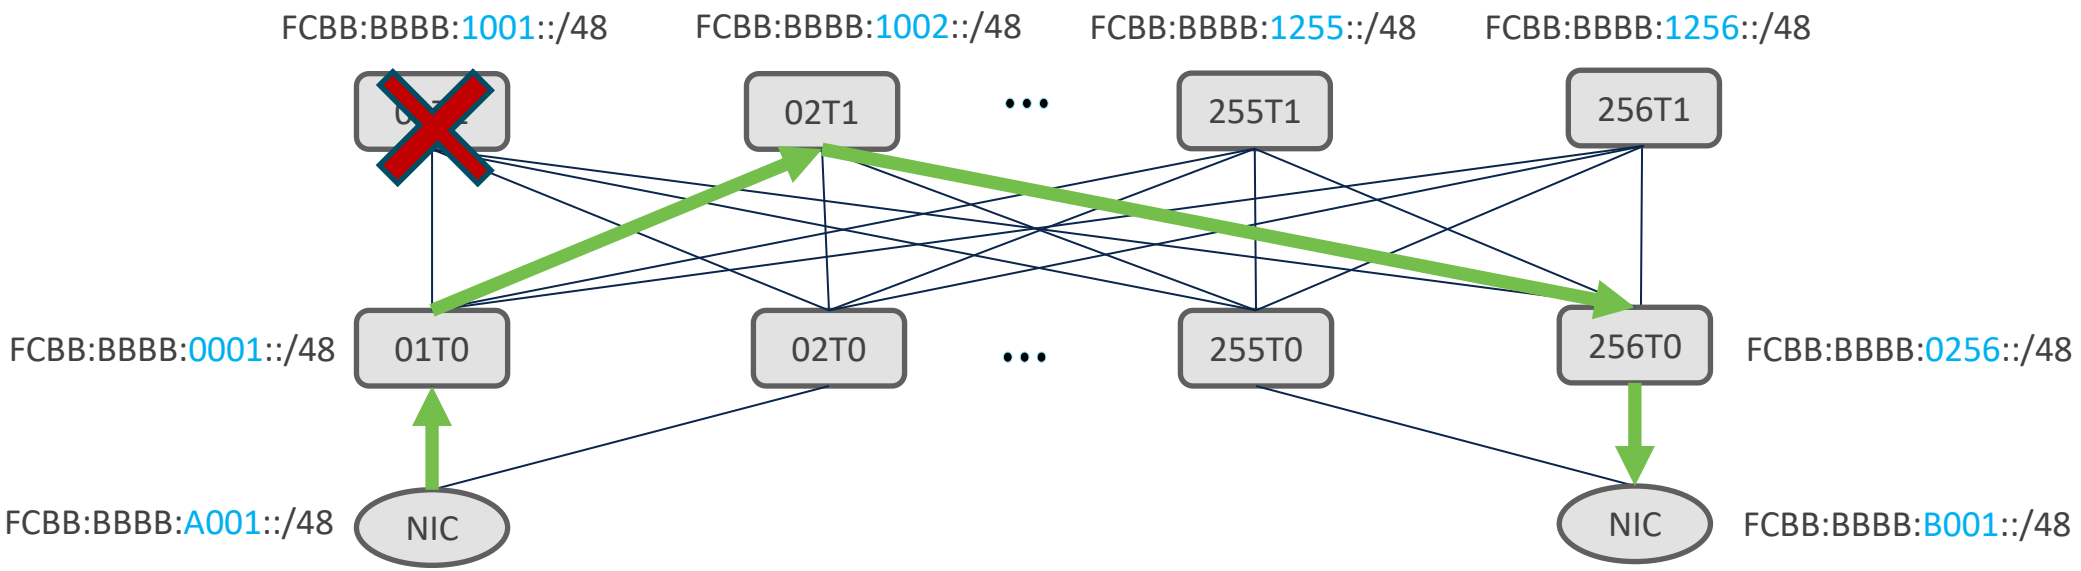

- NIC A to NIC B
  - via FCBB:BBBB:0001:1001:0256:B001::
  - Go to 01T0 then 01T1 then 256T0 then NIC B

Congestion/Failure " 01T1"

- NIC A to NIC B
  - via FCBB:BBBB:0001:1002:0256:B001::
  - Go to 01T0 then 01T1 then 256T0 then NIC B

Path control via SRv6 maximize utilization
(Using incast as an example)

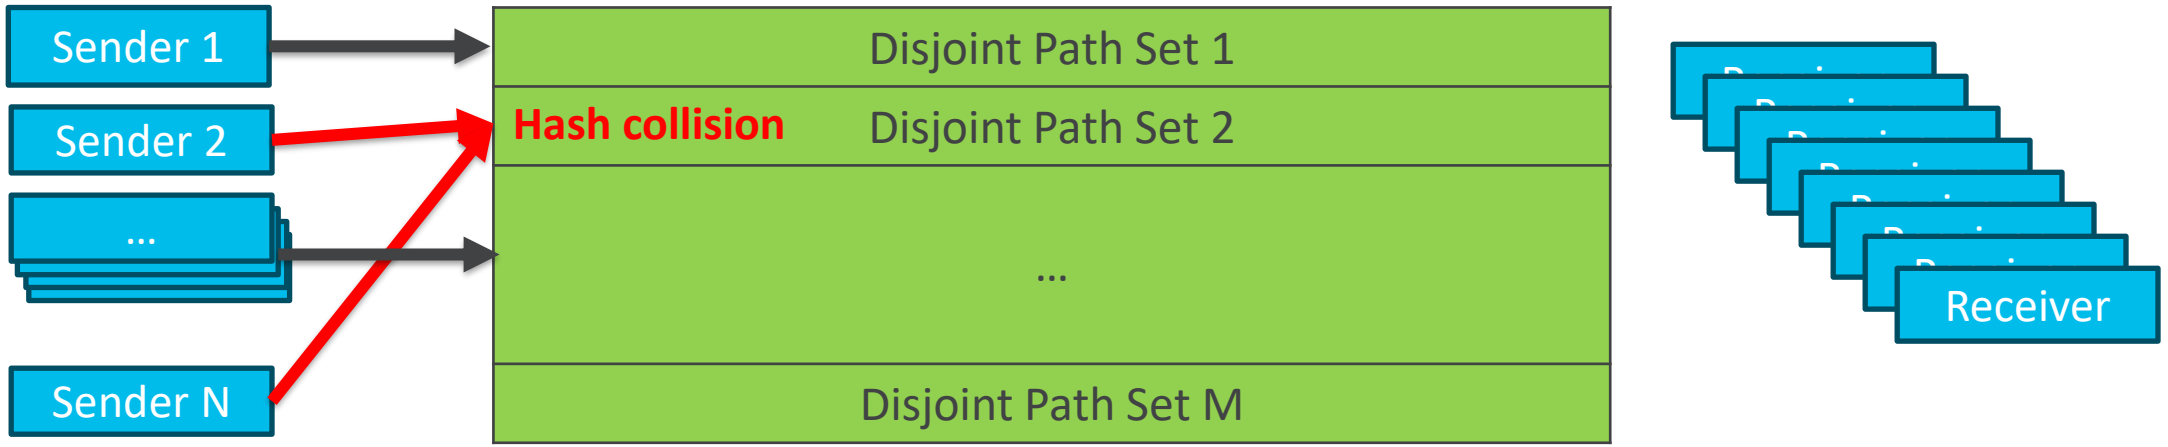

Case #1: using ECMP to distribute traffic

Passive hash-based ECMP cannot guarantee all disjoint paths are utilized and the distribution across paths is uneven most of the time.

Path control via SRv6 maximize utilization
(Using incast as an example)

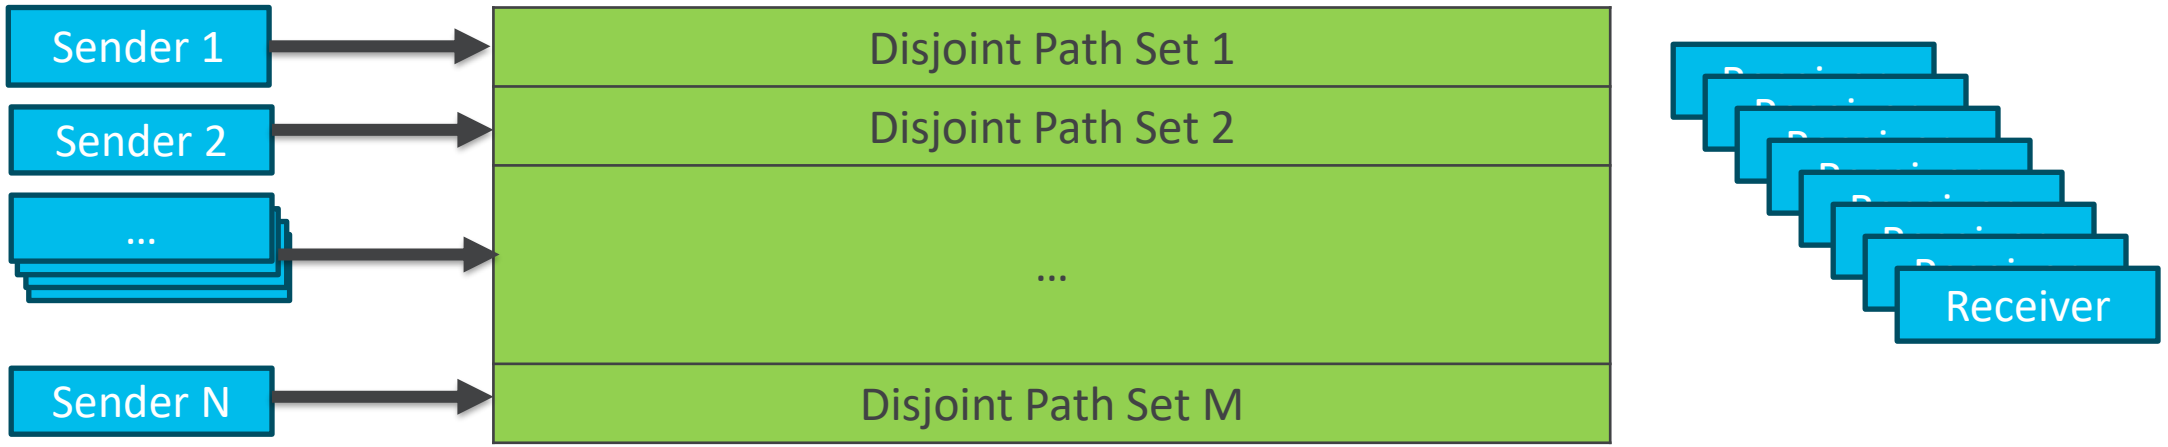

Case #2: using SRv6 to distribute traffic

SRv6 can deterministically control the path a flow takes, so we can achieve that ***all disjoint path sets are utilized,*** *and* ***each path set is not loaded more than a threshold***

Demo: SRv6 uSID – AI Backend use-case

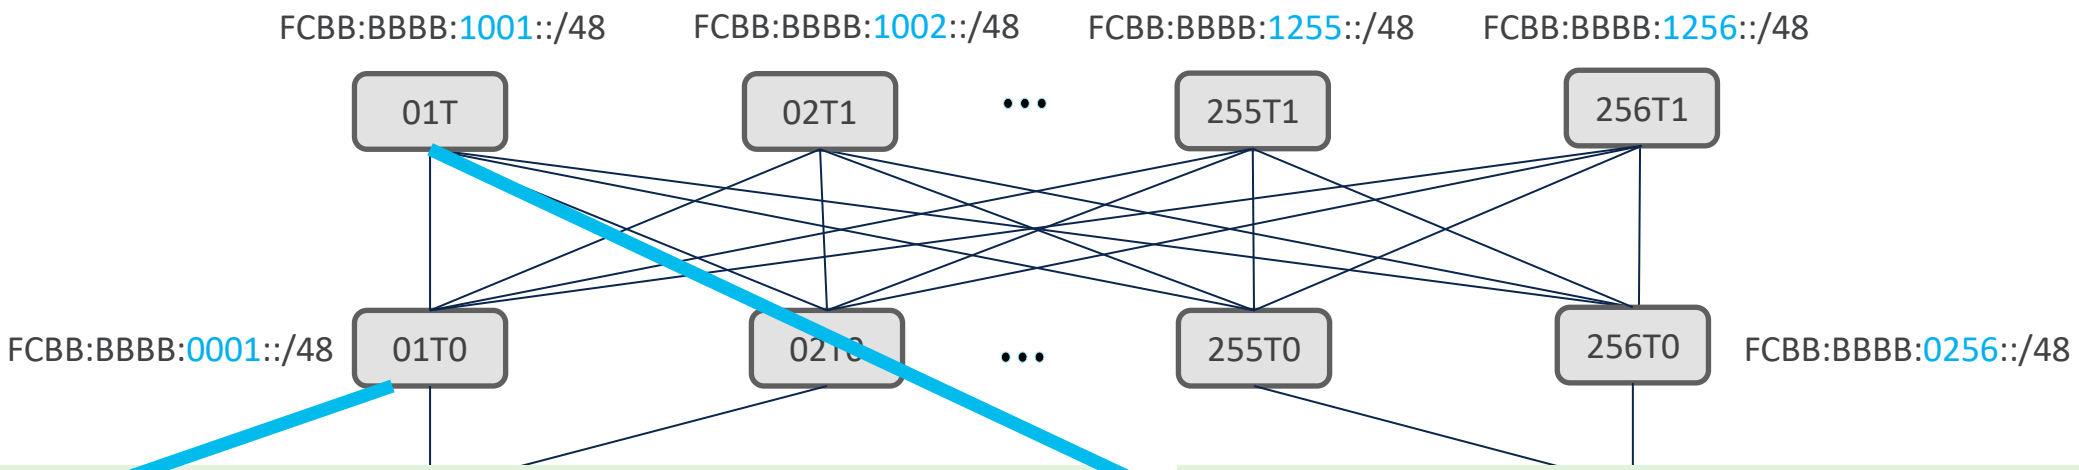

01T0

```
…
ipv6 route fcbb:bbbb:1001::/48 Ethernet01
ipv6 route fcbb:bbbb:1002::/48 Ethernet02
…
segment-routing
 srv6
  locators
   locator MAIN
    prefix fcbb:bbbb:0001::/48
    format usid-f3216
  static
   sid fcbb:bbbb:0011::/48 locator MAIN behavior uN
…
```

01T1

```
…
ipv6 route fcbb:bbbb:0001::/48 Ethernet01
ipv6 route fcbb:bbbb:0002::/48 Ethernet02
…
segment-routing
 srv6
  locators
   locator MAIN
    prefix fcbb:bbbb:1001::/48
    format usid-f3216
  static
   sid fcbb:bbbb:1001::/48 locator MAIN behavior uN
 …
```

Demo: SRv6 uSID – BGP L3VPN use-case

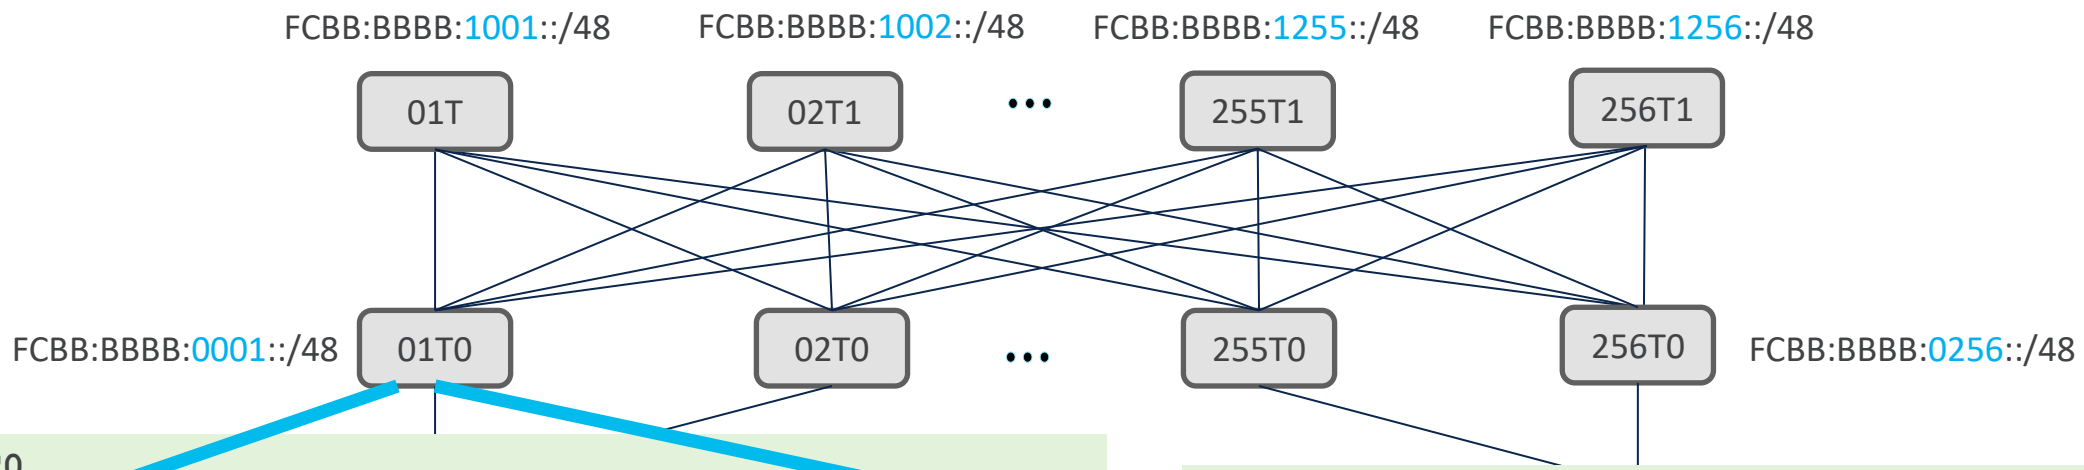

01T0

```
…
router bgp 65001
 bgp router-id 1.1.1.1
 neighbor 2001:db8:1:1001::1001 remote-as 65011

 address-family ipv6 unicast
  network fcbb:bbbb:1::/48
  network fcbb:bbbb:1::1/128

 address-family ipv6 vpn
   segment-routing srv6
     locator MAIN
     encap-behavior H_Encaps_Red
```

```
segment-routing
 srv6
  locators
   locator MAIN
    prefix fcbb:bbbb:1::/48
    format usid-f3216
  encapsulation
   source-address fcbb:bbbb:1::1
```

Do you want to try it yourself ?

- All config are available on GitHub
  - https://github.com/srv6-sonic/oss-eu-2025
- Two demo scenarios:
  - SRv6 AI backend
  - SRv6 L3VPN


CISCO