SRv6 uSID

Ahmed Abdelsalam

Engineering Technical Leader – Cisco Systems

IP is better than ever

- IP is at the heart of our industry
- We got accustomed that IP cannot achieve anything alone
- Plethora of shim layers: MPLS, GTP, NSH, UDP/VxLAN…
- These shim layers cannot be combined and require expansive translations

SRv6 uSID

- Build Anything
  - Any combination of underlay, overlay, service chaining, security…
  - VPN, Slicing, Traffic Engineering, Green Routing, FRR, NFV
- Any Domain
  - Access, Metro, Core, DC, Host, Cloud
  - End-to-End Stateless Policy
  - No protocol conversion or gateways at domain boundaries
- Seamless Deployment in Brownfield
- Built day-1 for Automation
- Standardized, Rich Eco-system, Rich Open Source (Linux, SONiC, FRR)

Outperform MPLS/VxLAN

Outperform MPLS - Daniel Voyer (Bell Canada)

- Native Optimum Slicing
  - SLID is encoded in Flow Label
- HW Linerate Push: 3 times better
  - J2 uSID linerate push: 30 uSIDs >> 10 MPLS Labels
- HW Counter and FIB consumption: 4 times better
  - uSID requires 4 times less counters and FIB entries than MPLS
- Routing scale: 20 times better
  - uSID supports summarization. MPLS requires host routes.
- Lookup efficiency: 2 to 3 times better
  - uSID can process 2 to 3 SIDs in a single lookup (LPM nature)
- Load-balancing: optimum and deterministic
  - uSID provides HW friendly entropy (fixed offset, shallow)

Bell SRv6 uSID Deployment
Paris 2022

Outperforms VxLAN – Gyan Mishra (Verizon)

- Seamless Host support for Network Programming
  - 6 uSID's in outer DA: RFC2460 IPinIP with opaque DA
- TE in the DC
  - elephant flows exist, asymmetric fabrics exist, TE is needed
- TE in the Metro/Core from the host
  - An SRv6 uSID DC allows for the application to control the network program in the metro/core without complex DPI and protocol conversion at the DC boundary,
- uSID DC provides lower MTU overhead (~5%)
  - Lower MTU overhead means lower DC cost
- Vendor, Merchant and SONIC/SAI maturity
  - uSID support across DC vendor (Cisco), Merchant (Cisco, Broadcom, Marvell), Sonic/Sai (Alibaba deployment)

SRv6 uSID DC Use-Case
Paris 2023

SRv6 uSID Reminder

segment-routing.net

- https://www.segment-routing.net/conferences/2024-05-09-lacnic41-srv6-usid-and-ipm/

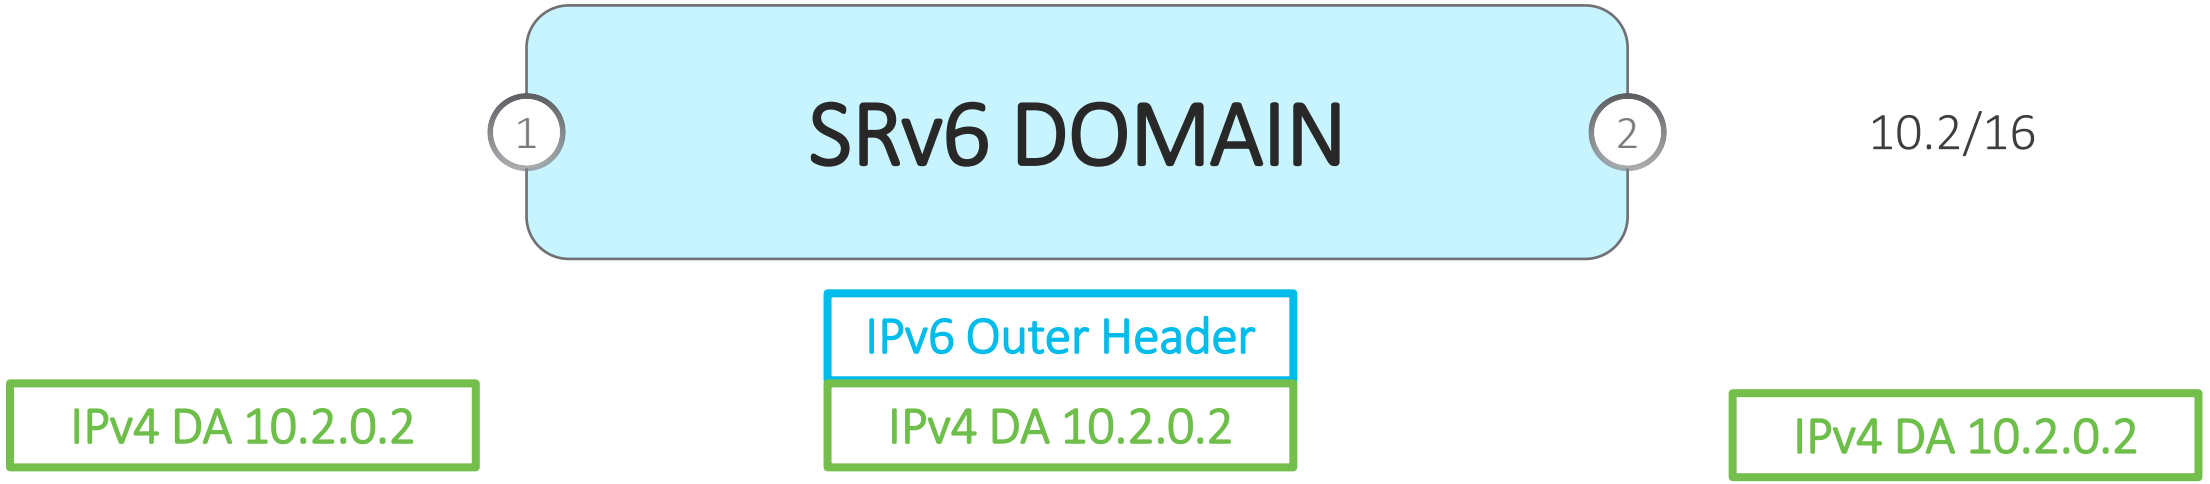
1
SRv6 DOMAIN
2
10.2/16
IPv6 Outer Header
IPv4 DA 10.2.0.2
IPv4 DA 10.2.0.2
IPv4 DA 10.2.0.2

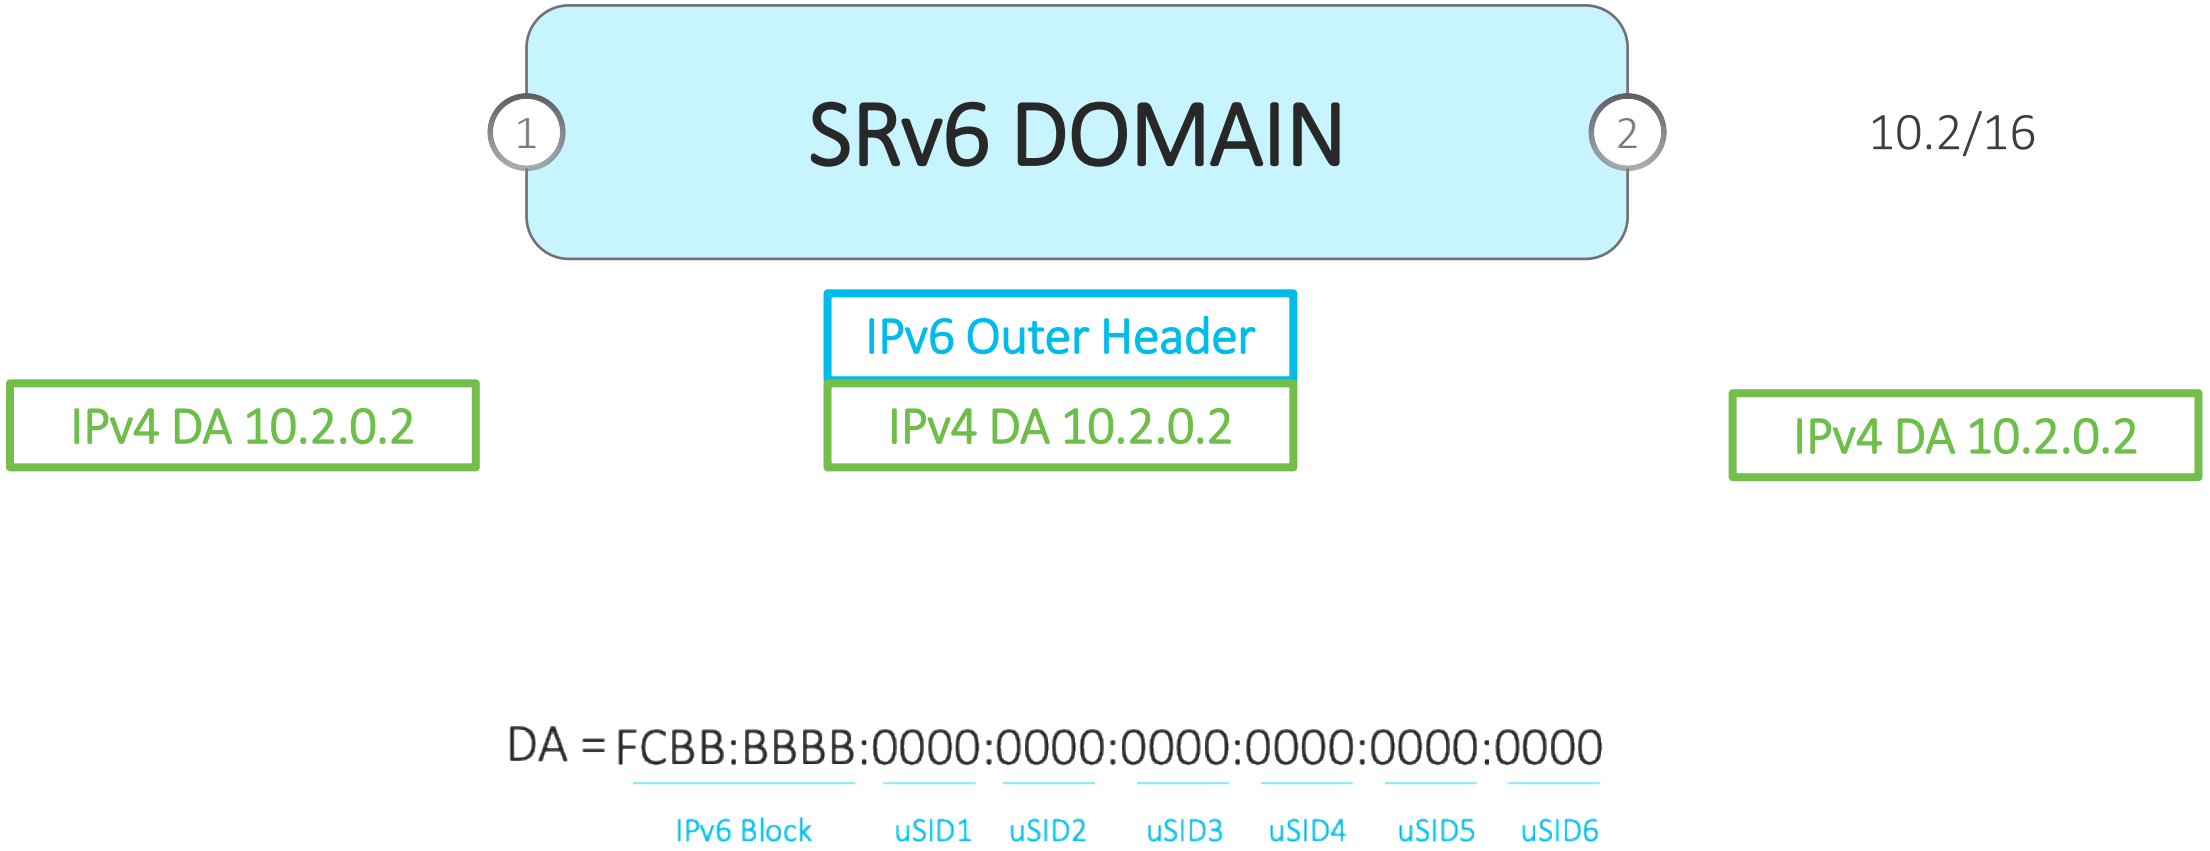
1
SRv6 DOMAIN
2
10.2/16
IPv6 Outer Header
IPv4 DA 10.2.0.2
IPv4 DA 10.2.0.2
IPv4 DA 10.2.0.2
DA = FCBB:BBBB:0000:0000:0000:0000:0000:0000
IPv6 Block
uSID1
uSID2
uSID3
uSID4
uSID5
uSID6

Transparent Service

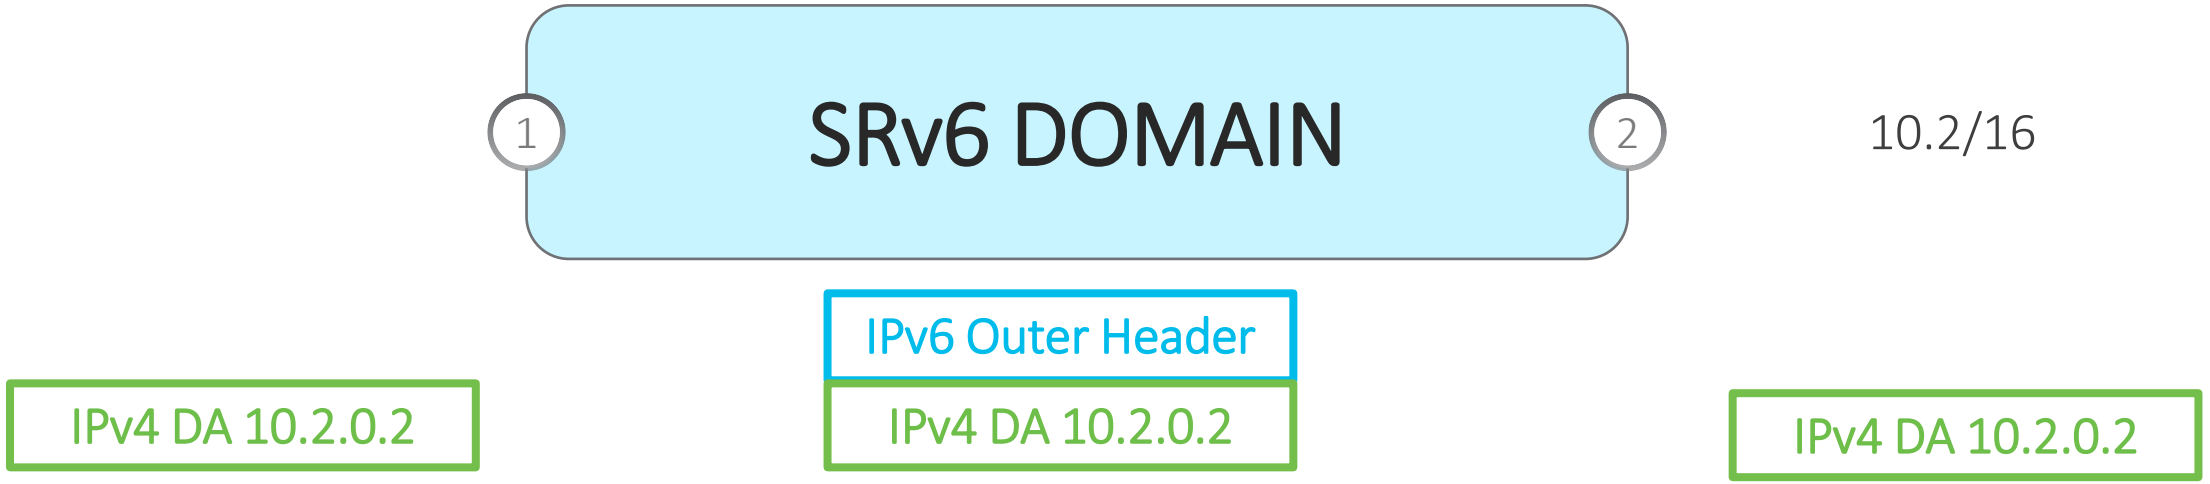

DA = FCBB:BBBB:0000:0000:0000:0000:0000:0000

IPv6 Block uSID1 uSID2 uSID3 uSID4 uSID5 uSID6

uSID can bound to any instruction (VPN, TE, NFV)

SRv6 uSID - VPN use-case (1)

- The VPN “inner” packet gets in the VRF of 1
- Inner packet can be IPv4, IPv6, Ethernet, …

SRv6 uSID - VPN use-case (2)

- An outer IPv6 header is added with Destination Address (DA)
- FD00:0000:0002:F009:0000:0000:0000:0000
- The outer DA holds the end-to-end stateless network program

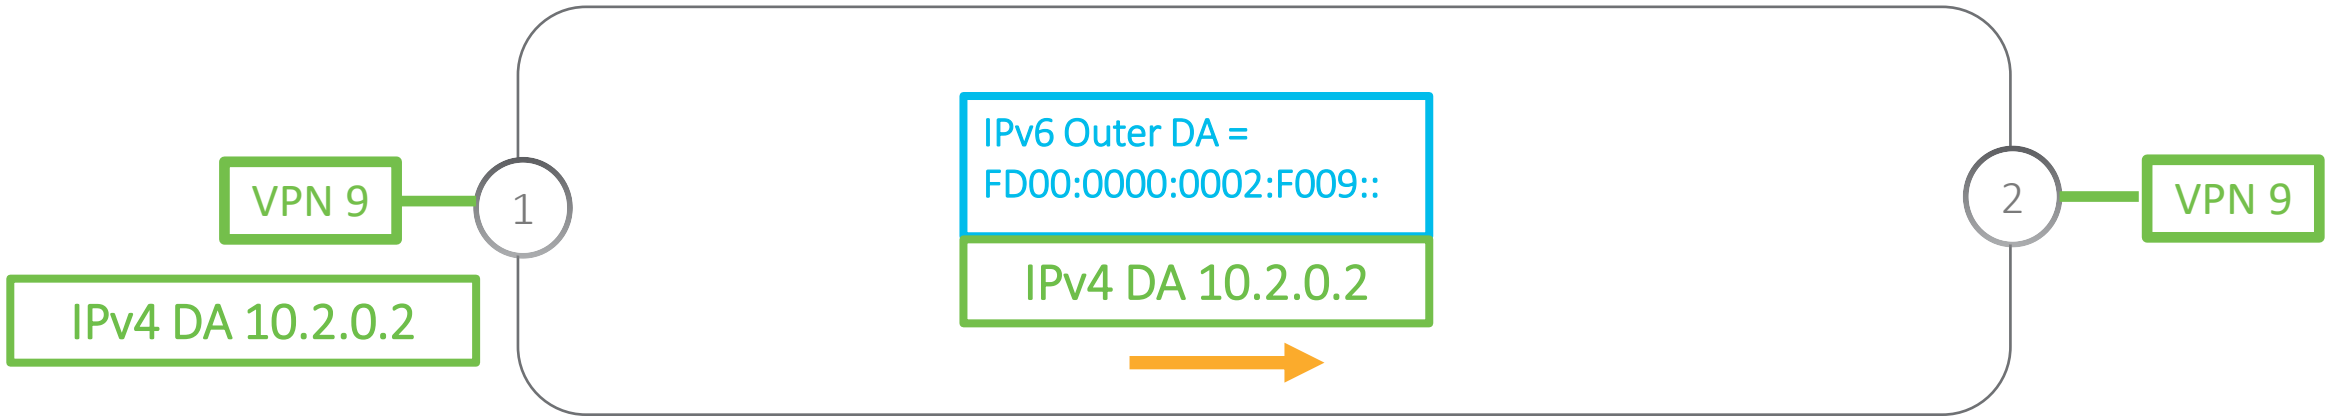

SRv6 uSID - VPN use-case (3)

- The outer header is popped
- The inner packet is sent to VPN 9 site
- The desired end-to-end behavior has been delivered without any per-flow state

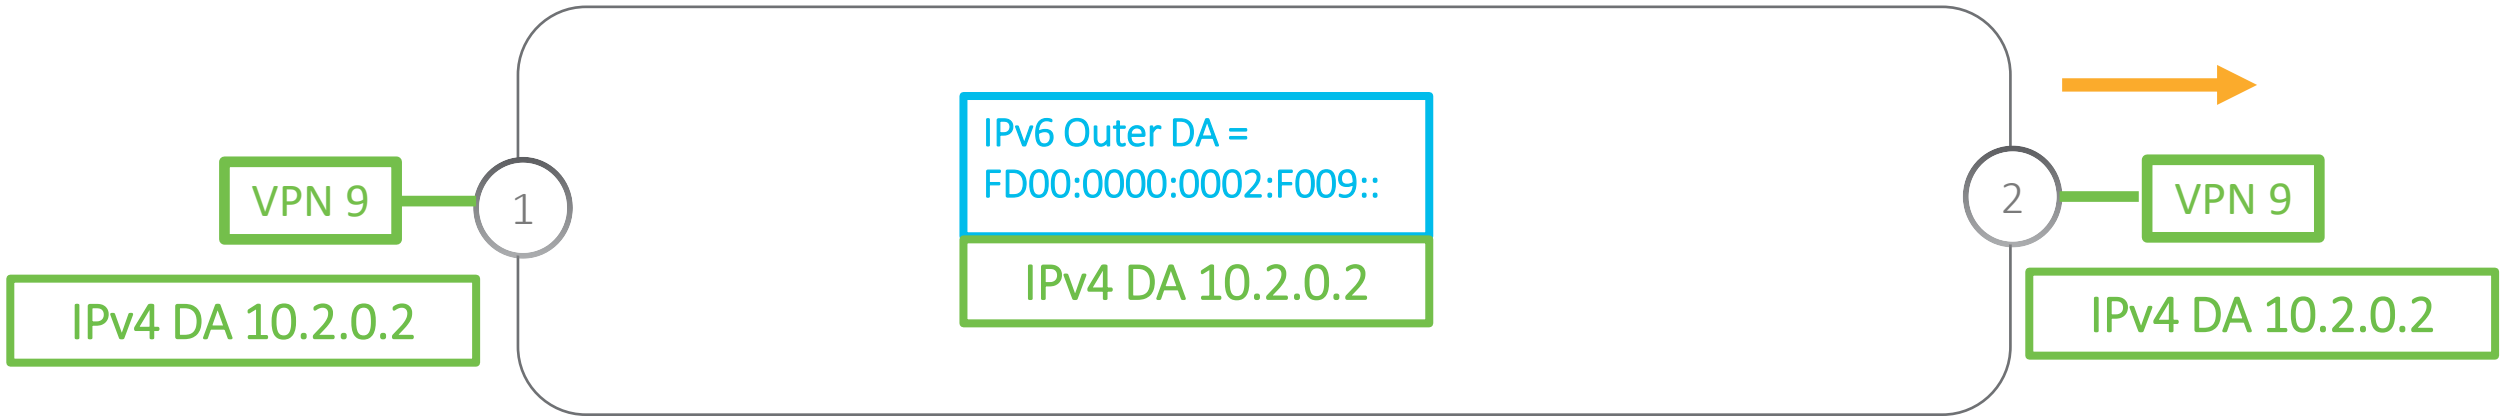

SRv6 uSID – Industry Update

Rich SRv6 uSID Ecosystem

Network Equipment Manufacturers

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

Merchant Silicon

Broadcom, Marvell, Barefoot Networks (Intel)

Open-Source Applications

Pyroute2, IPTables SERA, GoBGP, exa, BGP, eBPF, cilium, Wireshark, FRRouting, Snort, GoBMP

SRv6 uSID

Open-Source Networking Stacks

SONiC, Linux, fd.io, P4, OpenStack, Open Compute Project, 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 – Very robust support in Linux Kernel

- More than 8 years
  - Kernel 4.10 (Feb 2017)
- Most of the SRv6 behaviors are supported
  - Headend
  - Underlay
  - Overlay
  - Traffic steering
- Full commitment to SRv6 in Linux kernel
  - Adding new behaviors
  - Bug fixes
  - Performance improvements

SRv6 is Proposed Standard

Architecture

- SR Architecture – RFC 8402
- SRTE Policy Architecture – RFC 9256

Data Plane

- SRv6 Network Programming – RFC 8986
- IPv6 SR header – RFC 8754

Control Plane

- SRv6 BGP Services – RFC 9252
- SRv6 ISIS – RFC 9352
- SR Flex-Algo – RFC 9350

Operation & Management

- SRv6 OAM – RFC 9259
- Performance Management – RFC 5357

Strong Commitment and Leadership

Editor of 96% IETF RFCs
Co-author of 100% IETF RFCs

Interoperability test

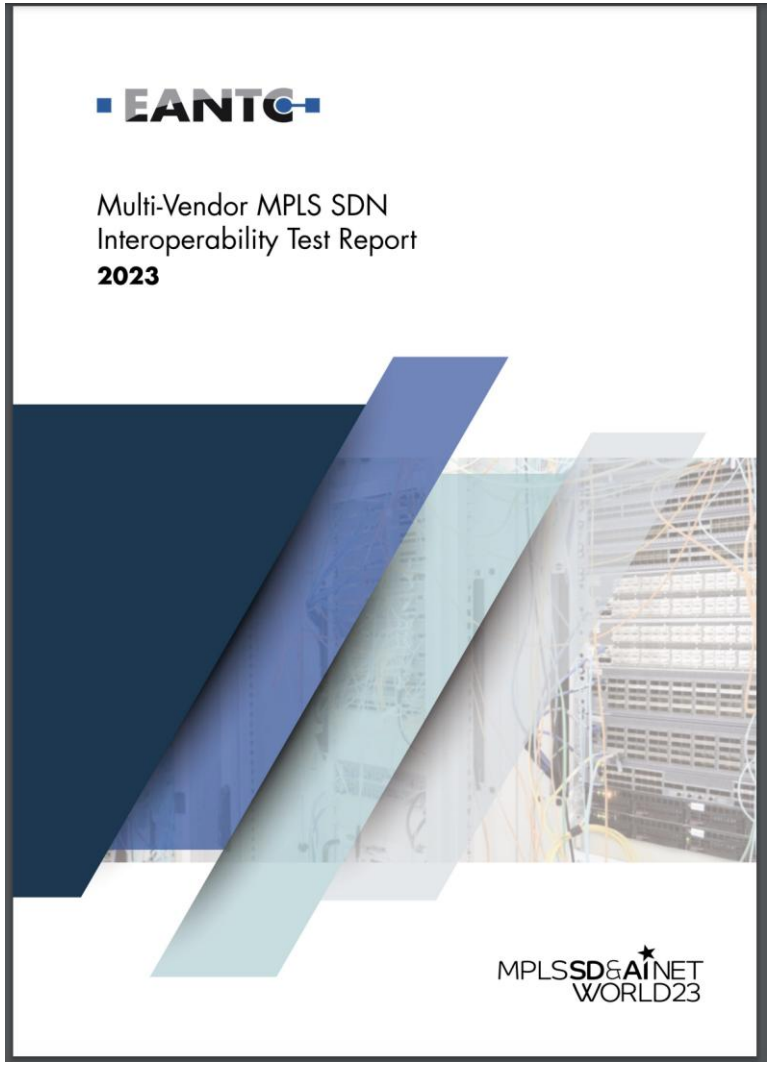

[EANTC-InteropTest2023]

[EANTC-InteropTest2024]

Interoperability test 2024

- Cisco, Arista, Juniper, Huawei, Nokia, Ericsson, ZTE, H3C

This year marked a significant milestone in our testing procedures as, for the first time, we exclusively utilized micro Segment ID (μSID) across all our tests. We had started testing multi-vendor interoperability of μSIDs already in 2023, and this year, the concept has been adopted by all participating vendors for all SRv6 test scenarios. This move underscores a notable industry trend towards embracing this method. Moreover, we explored using SRv6 argument signaling for BGP service routes within the ELAN multi-homing test. This year also saw the first verification of multicast functionality using Msr6, alongside presenting a case for link resource slicing within an SRv6 framework.

Segment Routing over IPv6 (SRv6) Test Results

SRv6 stands out in the evolving scene of network technologies for its innovative approach to simplifying network operations, enhancing programmability, and supporting the demands of modern network services, especially in the context of 5G and beyond.

This year marked a significant milestone in our testing procedures as, for the first time, we exclusively utilized micro Segment ID (μSID) across all our tests. We had started testing multi-vendor interoperability of μSIDs already in 2023, and this year, the concept has been adopted by all participating vendors for all SRv6 test scenarios. This move underscores a notable industry trend towards embracing this method. Moreover, we explored using SRv6 argument signaling for BGP service routes within the ELAN multi-homing test. This year also saw the first verification of multicast functionality using Msr6, alongside presenting a case for link resource slicing within an SRv6 framework.

It's also significant to highlight that our evaluations, including Layer 2/3 VPN services, SRv6 Locator Summarization, and Unreachable Prefix Announcement (UPA) frameworks and the rest of the tests, have drawn a wide array of participants this year. Such extensive involvement points to a considerable progression towards improved interoperability across the industry.

In all our testing scenarios, we connected every participating router to the traffic generator (Keysight IxNetwork) and initiated mesh traffic among all nodes. This was done over the service being evaluated or the service implemented, specifically to demonstrate the functionality of the feature under test.

L2/L3VPN Services over SRv6 Test

Layer 2/3 VPN services are basic constructs for transporting isolated and protected customer traffic across an SRv6 network – similar to other Segment Routing and MPLS network platforms.

Our testing began with constructing a comprehensive topology incorporating all participants, enabling us to conduct most of our tests using this setup. This topology featured two IS-IS levels linked by several Area Border Routers (ABRs). Every node was configured to establish a BGP connection with a route reflector and form IS-IS adjacencies with neighboring routers. The ABRs were set up to support both levels and were tasked with route leaking from level 2 to level 1, ensuring seamless connectivity across the entire network architecture.

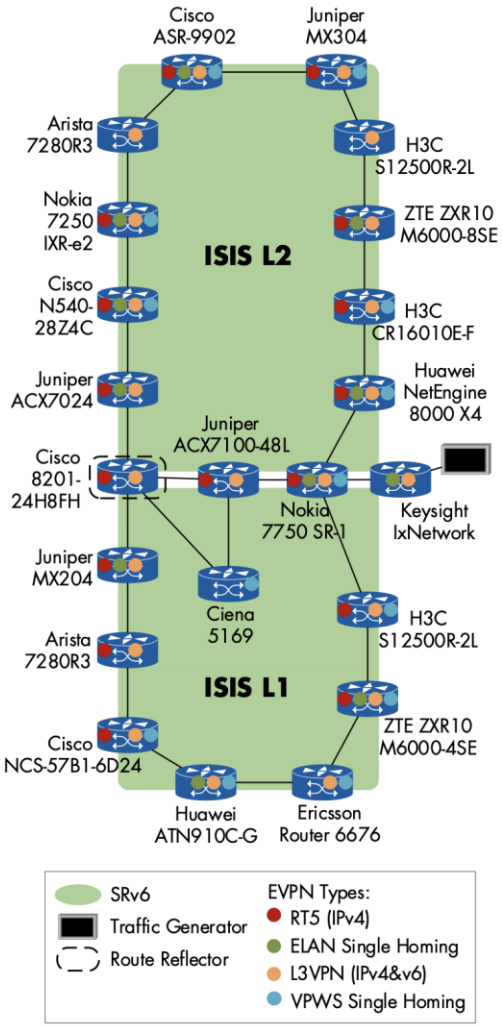

Figure 48: L2/L3 Services over SRv6

Configurations were established for SRv6 locators using 32-bit Locator Block and 16-bit Locator Node μSID format (F3216). Following IETF RFC 9252 (BGP Overlay Services Based on SRv6), vendors built various Layer 2 and Layer 3 VPN services. Some participants implemented the Transposition Scheme defined in RFC 9252 to increase the efficient packing of service routes. This method involves shifting the FUNCTION portion into the label field within a route's Network Layer Reachability Information (NLRI). Given that the remain-

Linux kernel since the very beginning

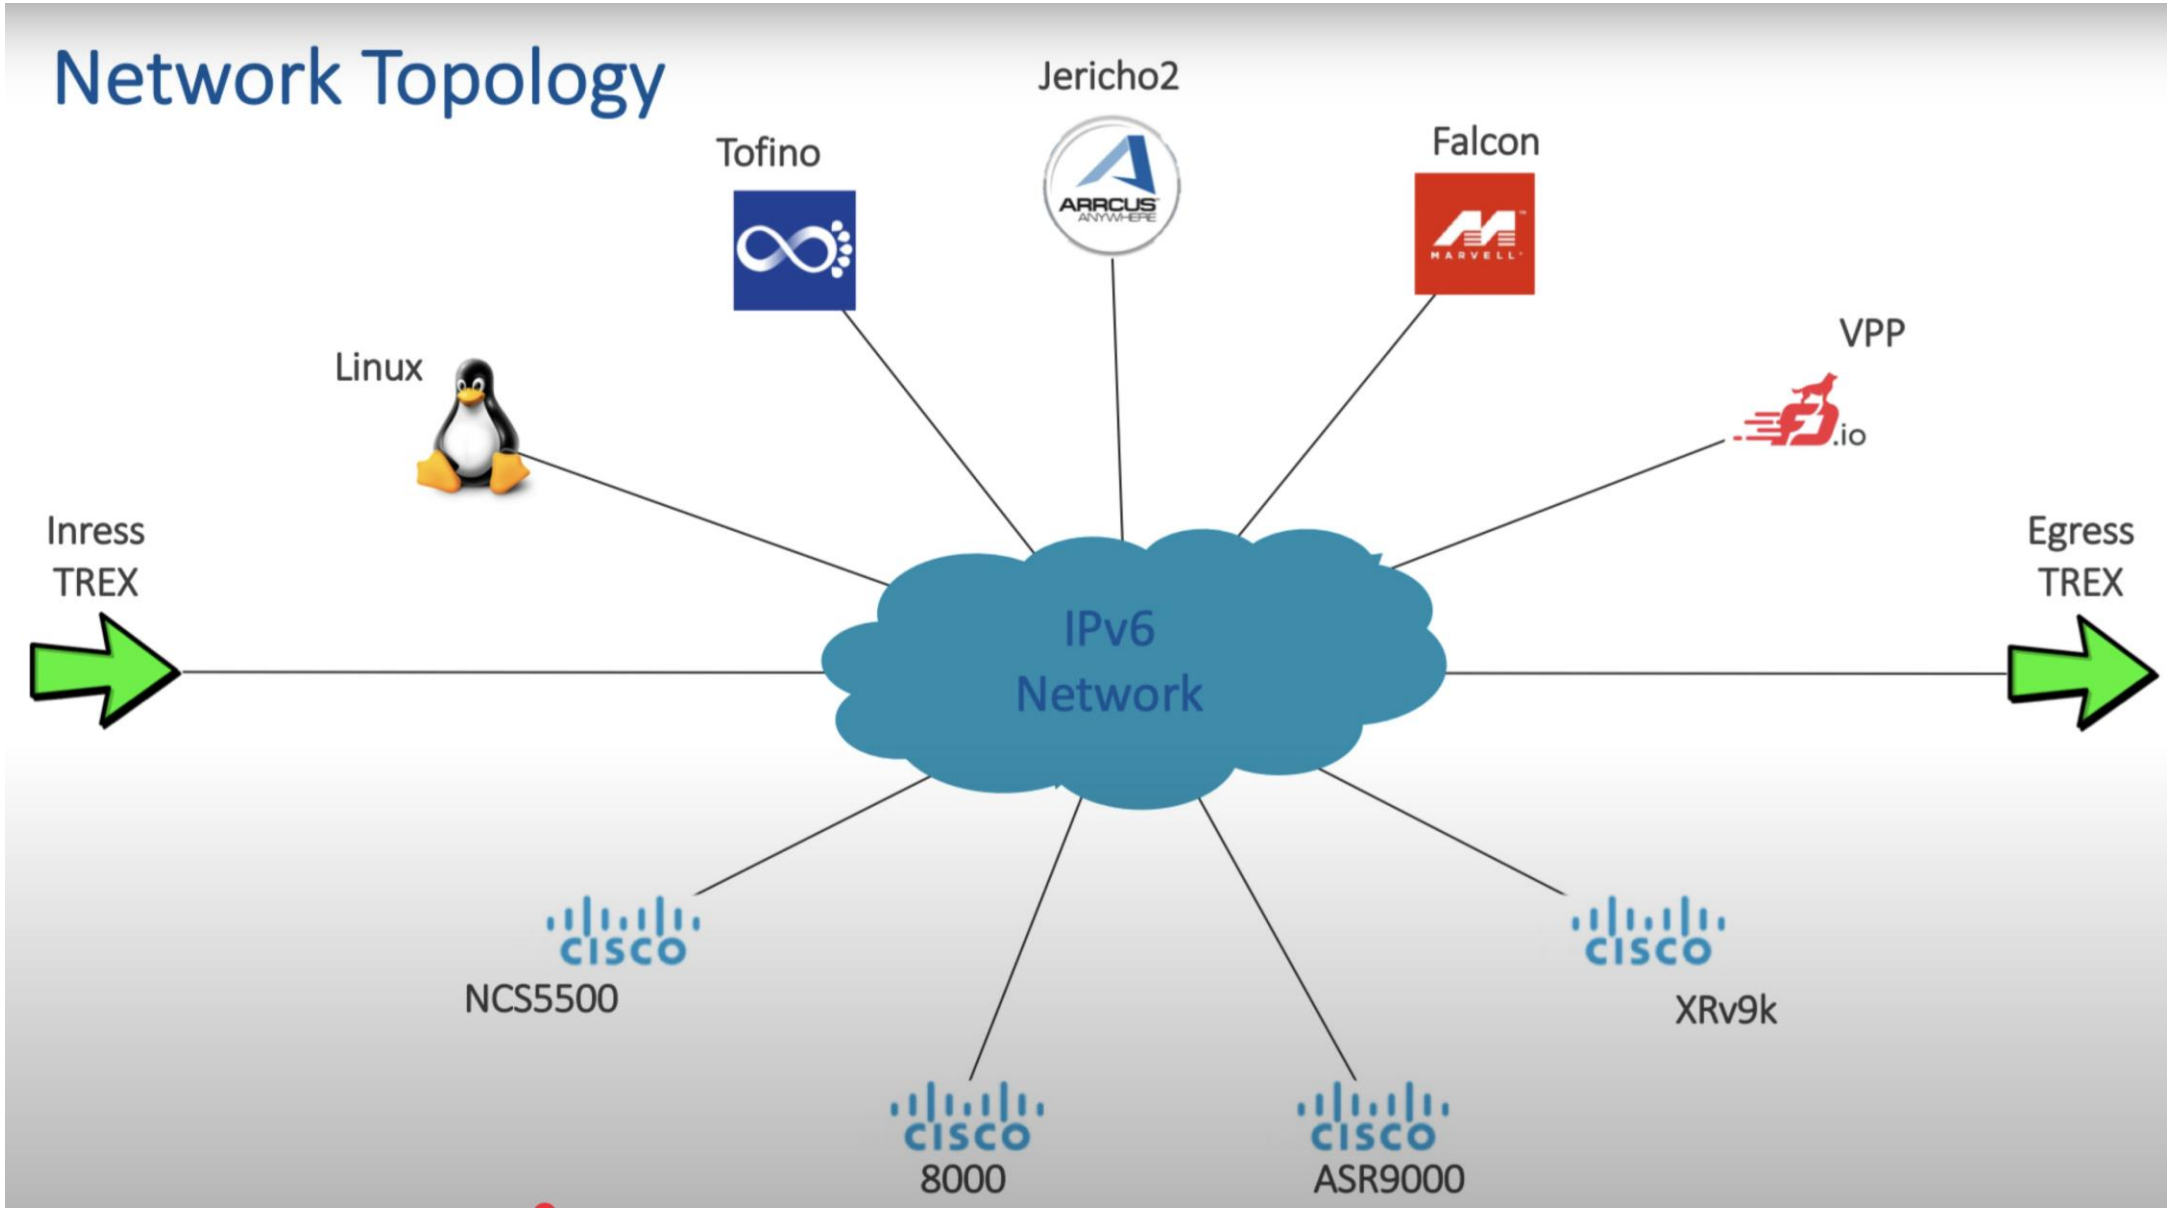

Over 85000 SRv6 uSID routers deployed

Inter-DC/Metro Traffic Engineering across all of China
Eddie Ruan

14k+ devices, 70% services on uSID
Akash Agrawal

TITAN: Swisscom's new converged IP Network replacing the many IP Networks they used to have.

Simplicity Always Prevails

Stay up-to-date

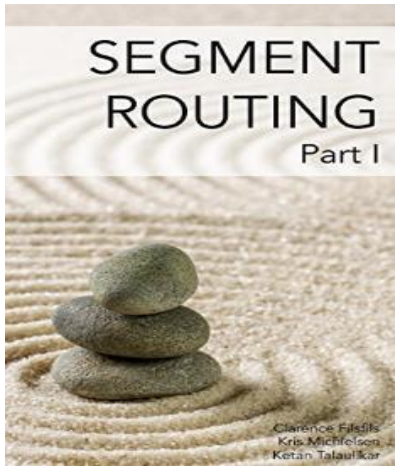

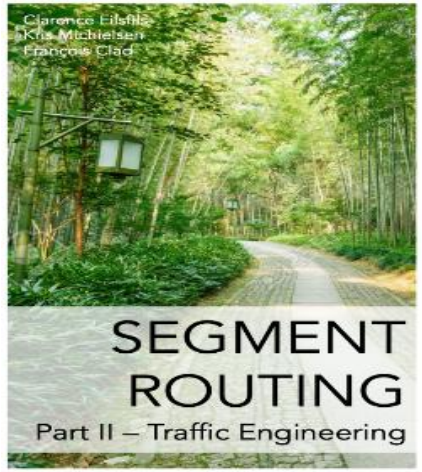

segment-routing.net

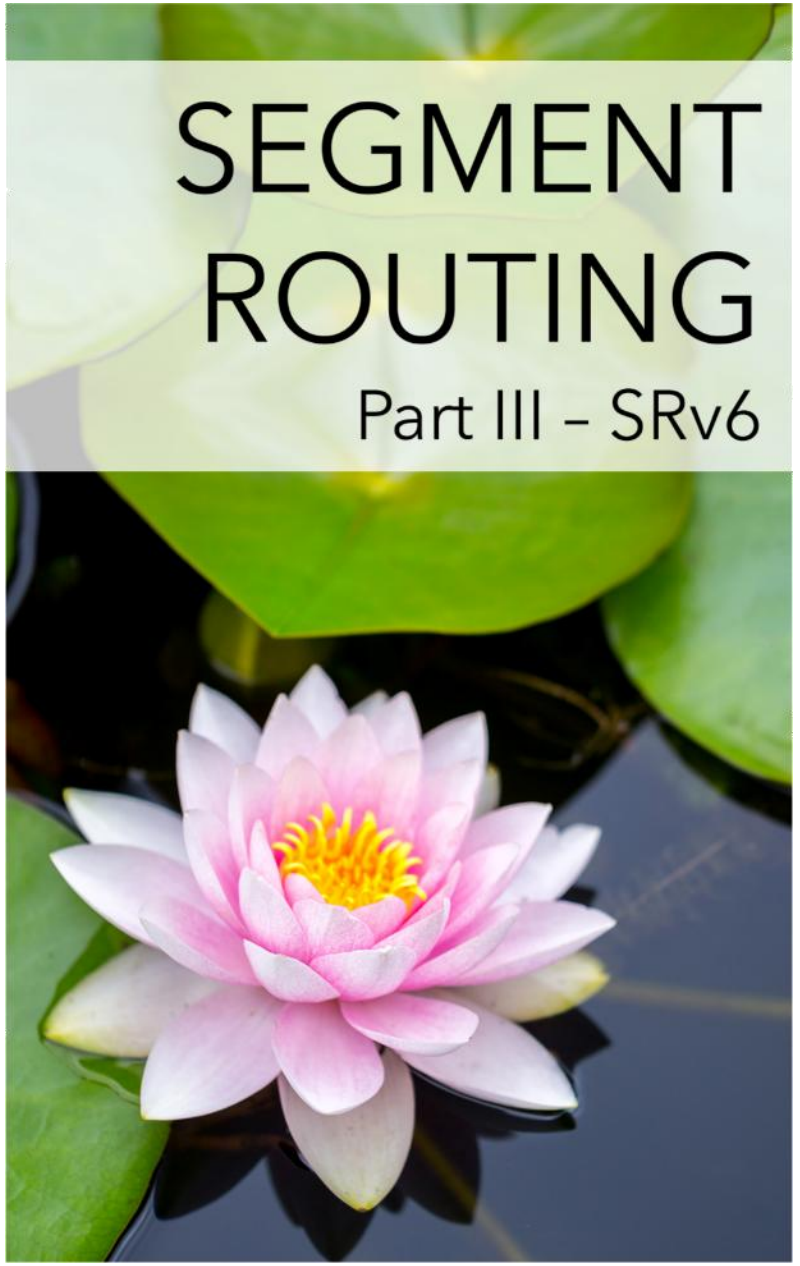

https://www.amazon.com/dp/B0D6GWWRWH

Conclusion

- IP is better than ever
- SRv6 uSID
  - Build Any service: any combination of underlay, overlay, service chaining, security…
  - Across any domain: Access, Metro, Core, DC, Host, Cloud (End-to-End)
  - No shim layers
  - No protocol conversion or gateways at domain boundaries
- Standardized, Rich Eco-system, Rich Open Source
  - Linux Kernel
  - SONiC
  - FRR


CISCO