SRv6 Network Programming

Pablo Camarillo
pcamaril@cisco.com
Mayo 2019 – ESNOG23

Segment Routing

- Source Routing
  - the topological and service (NFV) path is encoded in packet header
- Scalability
  - the network fabric does not hold any per-flow state for TE or NFV
- Simplicity
  - automation: TILFA sub-50msec FRR
  - protocol elimination: LDP, RSVP-TE, VxLAN, NSH, GTP, ...
- End-to-End
  - DC, Metro, WAN

Two dataplane instantiations

Segment Routing

MPLS

- leverage the mature MPLS HW with only SW upgrade
- 1 segment = 1 label
- a segment list = a label stack

IPv6

- leverages ~~RFC2460~~ RFC8200 provision for source routing extension header
- 1 segment = 1 address
- a segment list = an address list in the SRH

IPv6 adoption is a reality

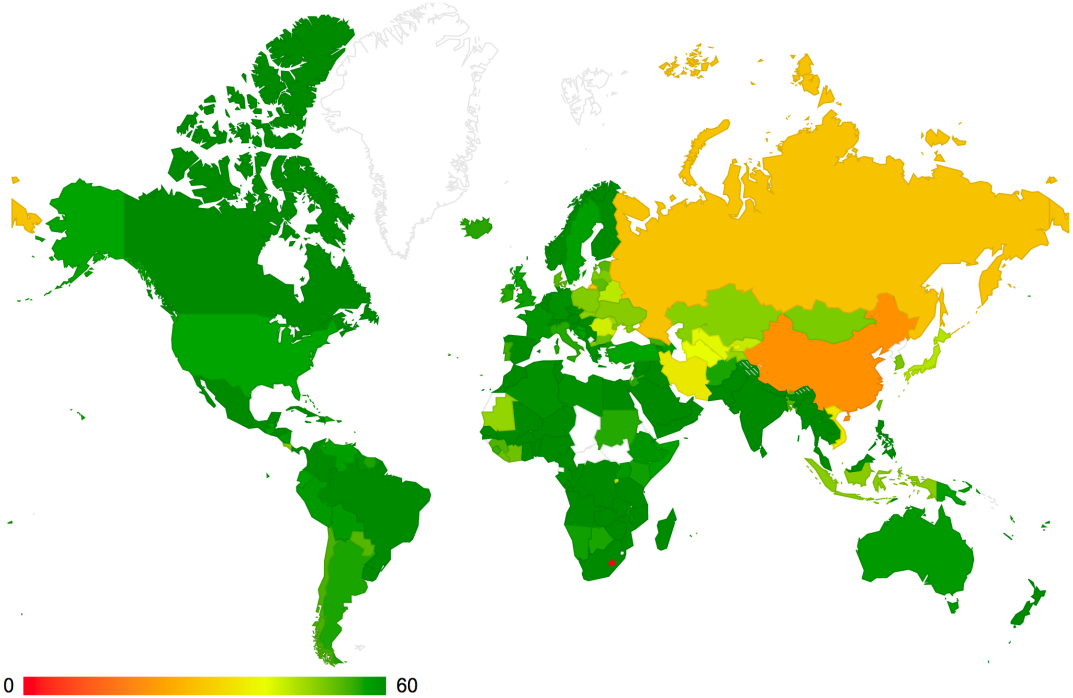

% Web pages available over IPv6

Sources: 6lab.cisco.com – Web content
Cisco VNI Global IP Traffic Forecast, 2017-2022

Global IPv6 traffic **grew 226%** in 2017

Globally IPv6 traffic **will grow 18-fold** from 2017 to 2022

IPv6 **will be 38%** of total Internet traffic in 2022

IPv6 provides reachability

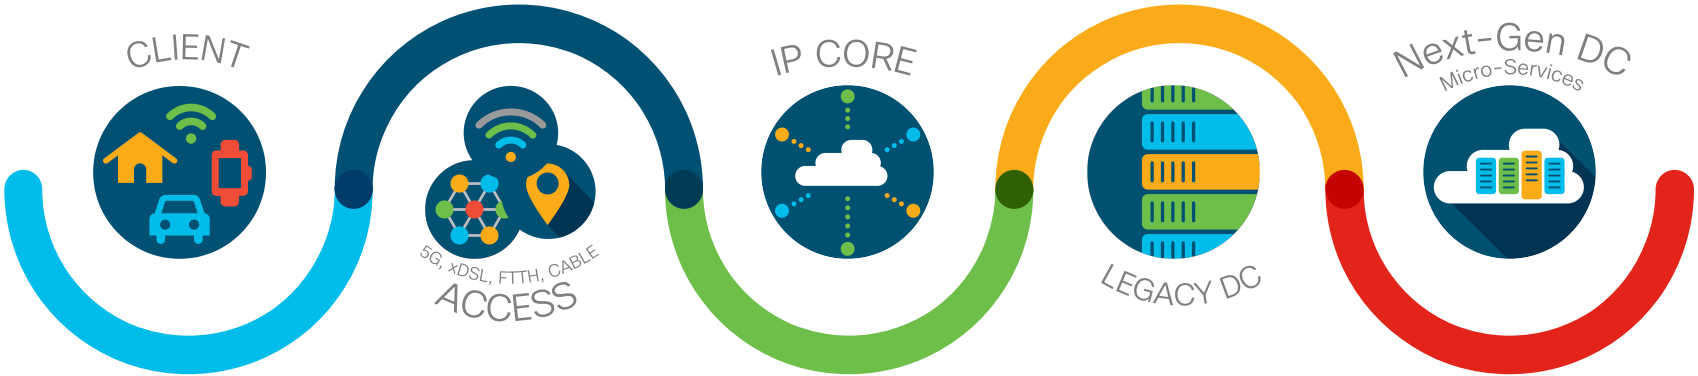

IPv4 limitations & work-arounds

- × Limited address space → NAT
- × No engineered Load Balancing → MPLS Entropy Label, VxLAN UDP
- × No VPN → MPLS VPN's, VxLAN
- × No Traffic Engineering → RSVP-TE, SR-TE MPLS
- × No Service Chaining → NSH

IPv4 limitations & work-arounds

- × Limited address space → NAT
- × No engineered Load Balancing → MPLS Entropy Label, VxLAN UDP
- × No VPN → MPLS VPN's, VxLAN
- × No Traffic Engineering → RSVP-TE, SR-TE MPLS
- × No Service Chaining → NSH

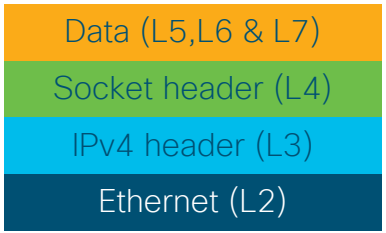

Simplicity (back to the OSI model)

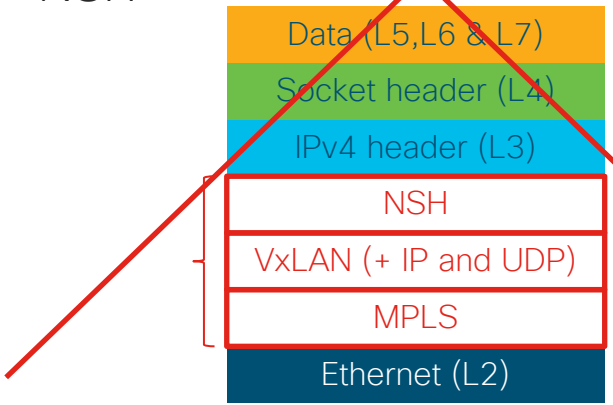

SRv6 unleashes IPv6 potential

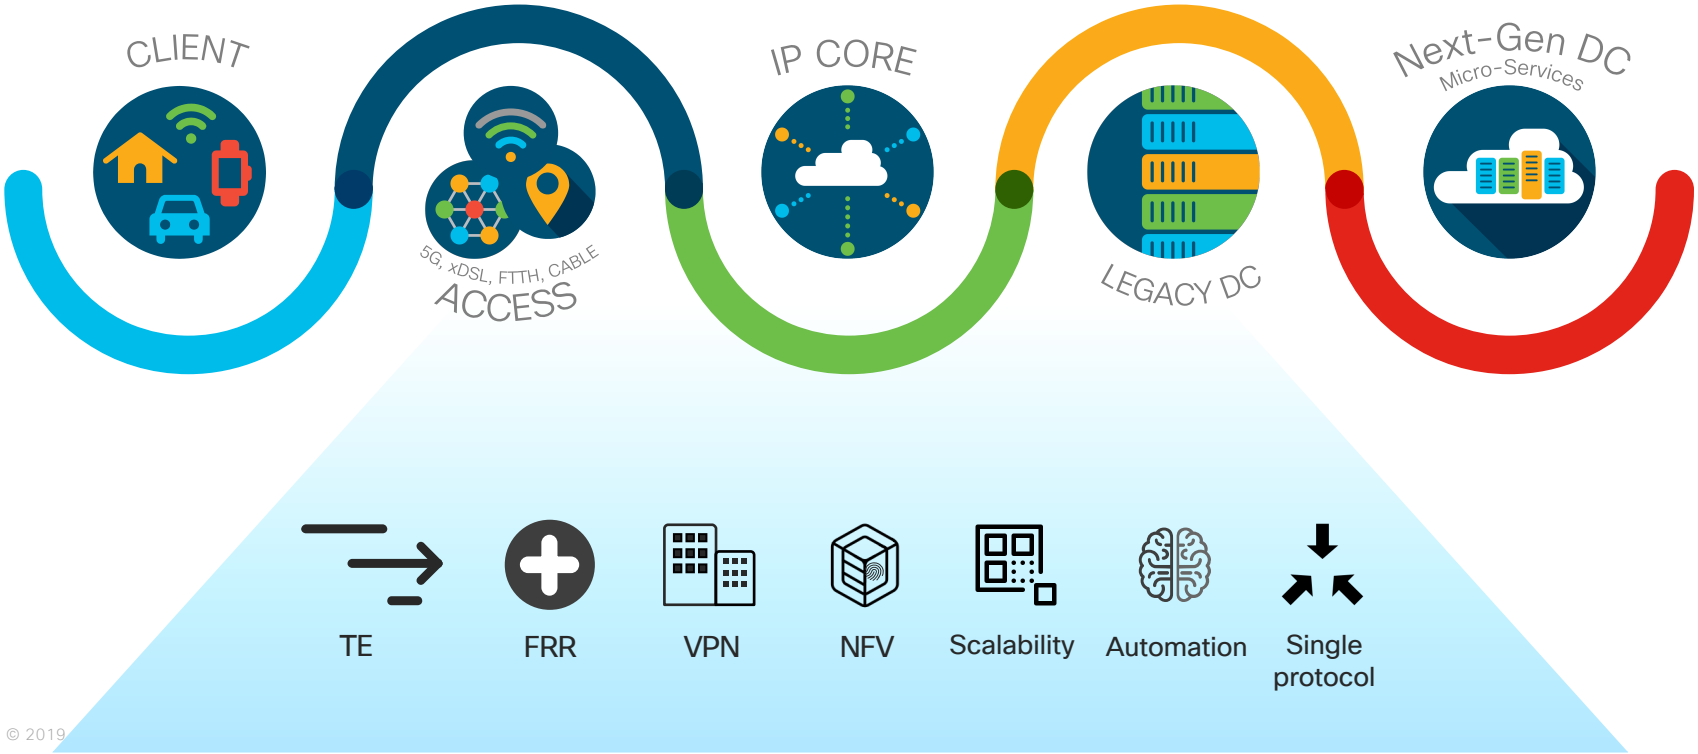

SR for anything:
Network as a Computer

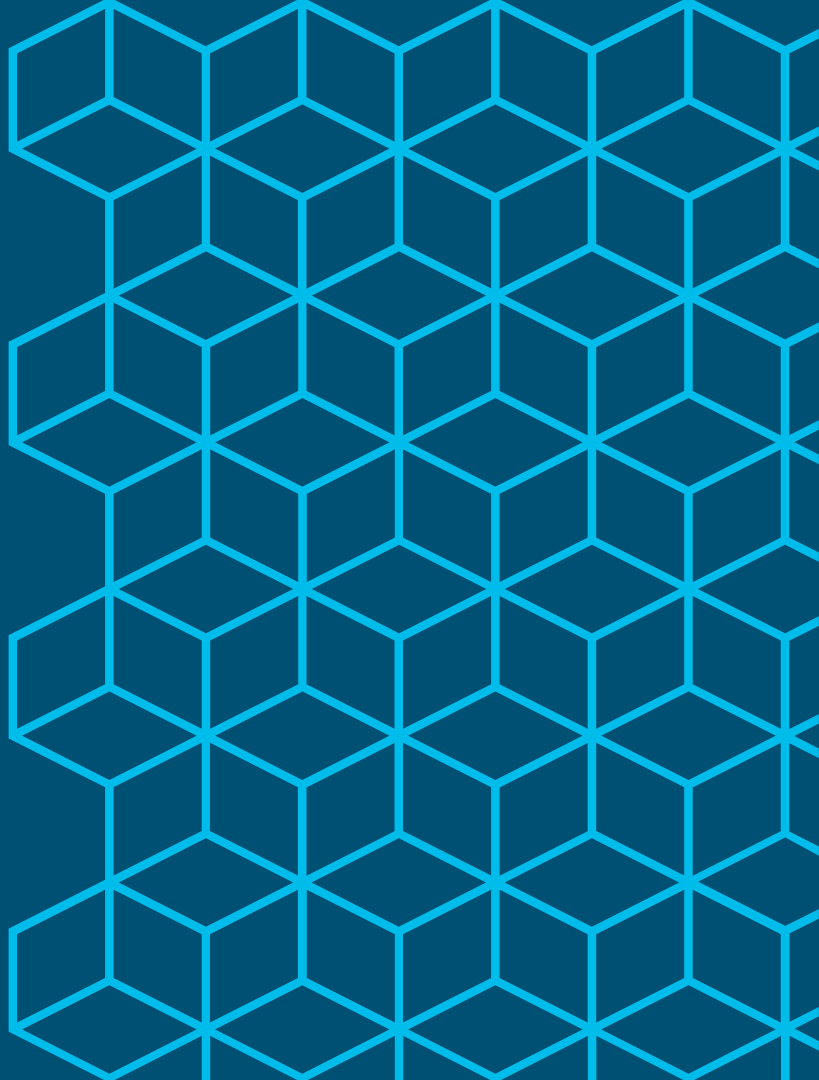

Network instruction

- 128-bit SRv6 SID
  - Locator: routed to the node performing the function
  - Function: any possible function
    either local to NPU or app in VM/Container
  - Flexible bit-length selection

Network Program

Next Segment →

Locator 1	Function 1
Locator 2	Function 2
Locator 3	Function 3

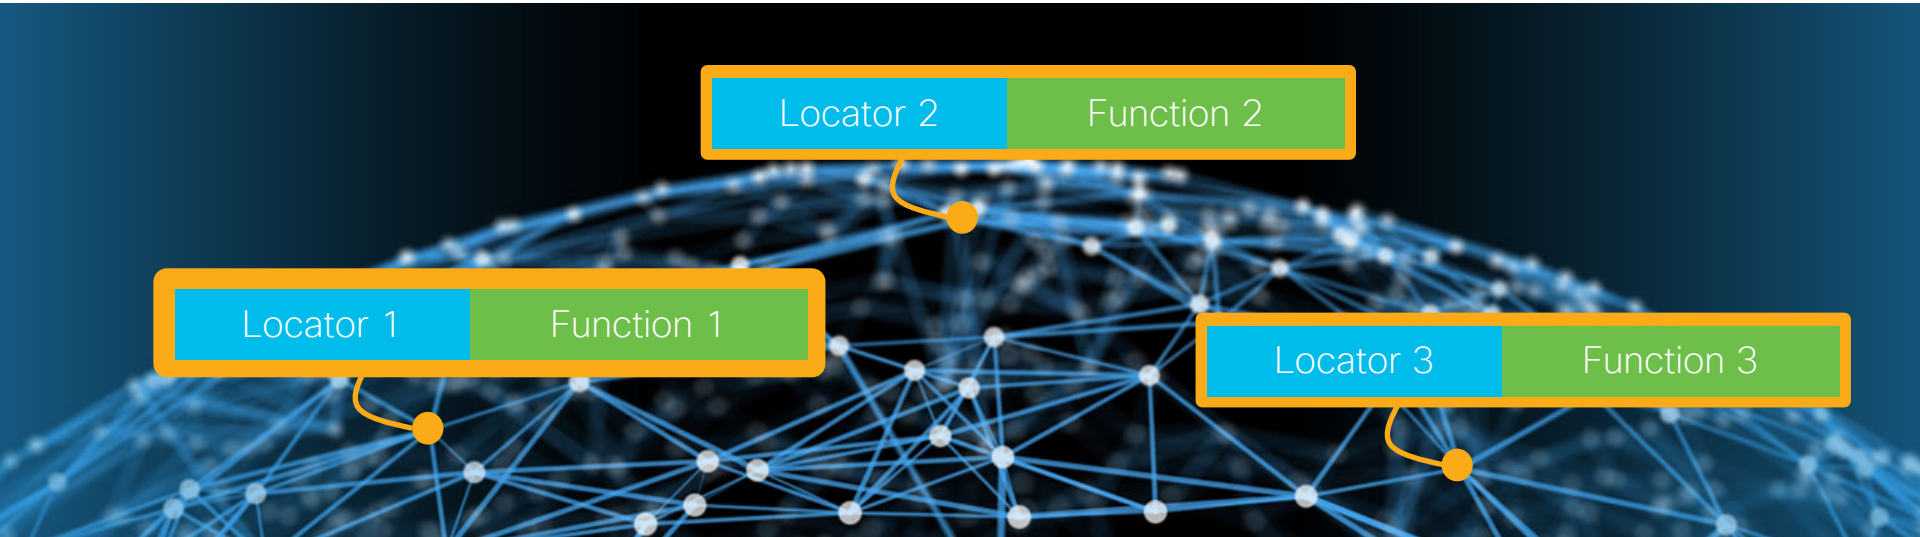

Network Program

Next Segment

Locator 1	Function 1
Locator 2	Function 2
Locator 3	Function 3

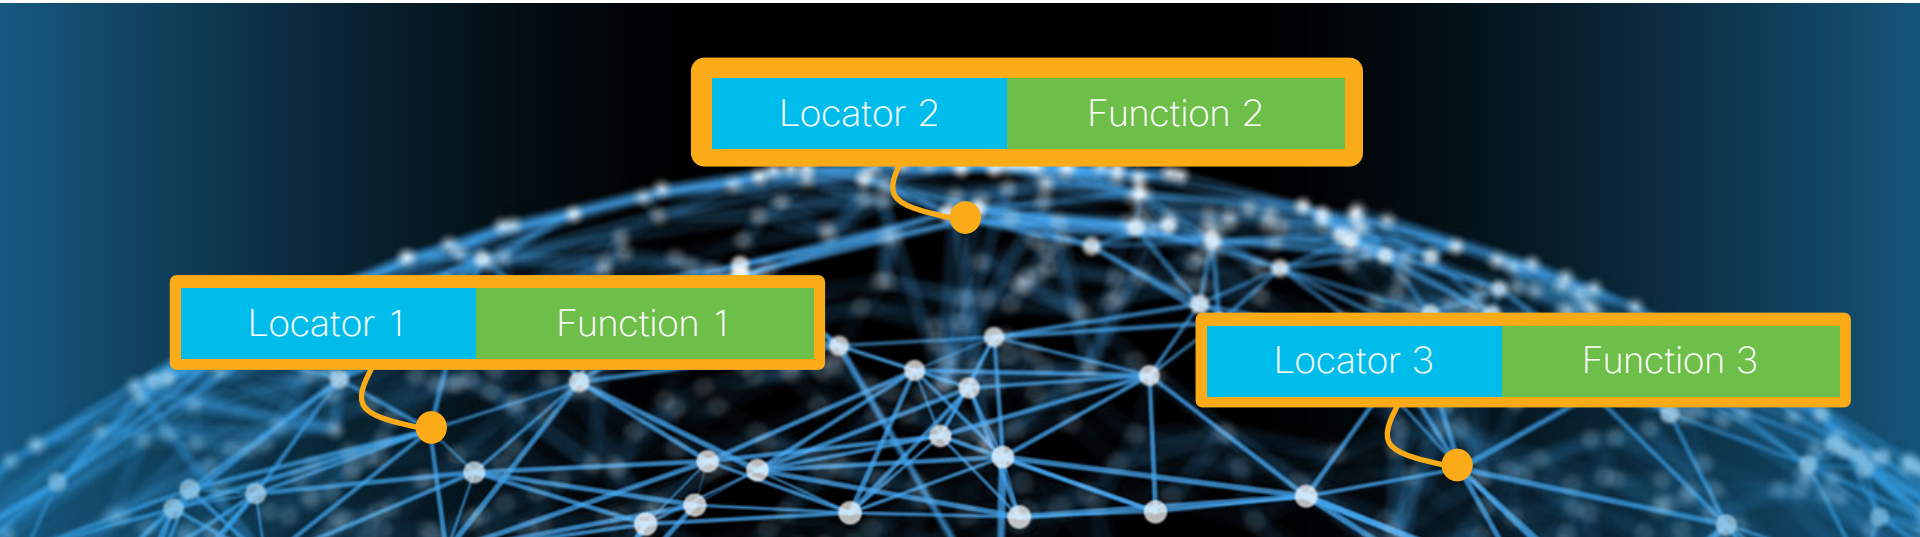

Network Program

Locator 1	Function 1
Locator 2	Function 2
Locator 3	Function 3

Next Segment →

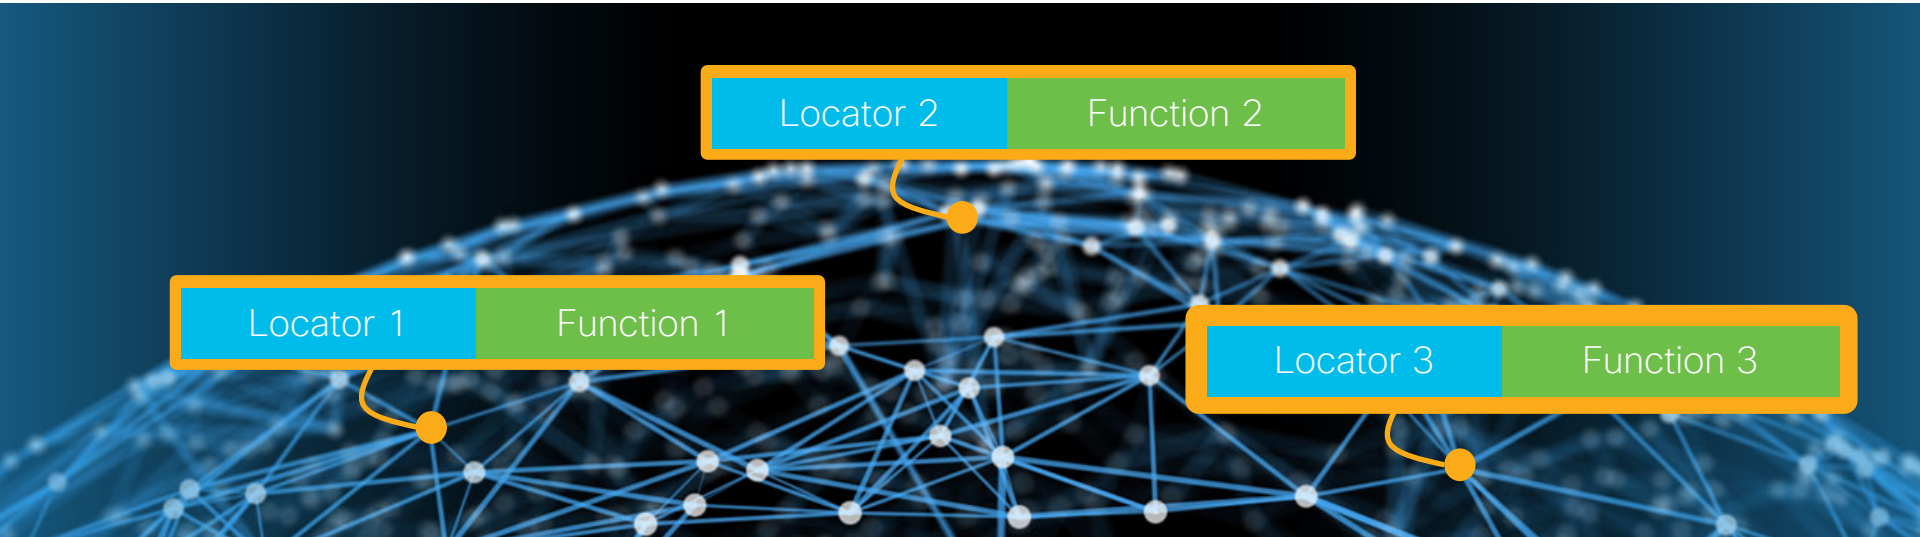

Network Program in the Packet Header

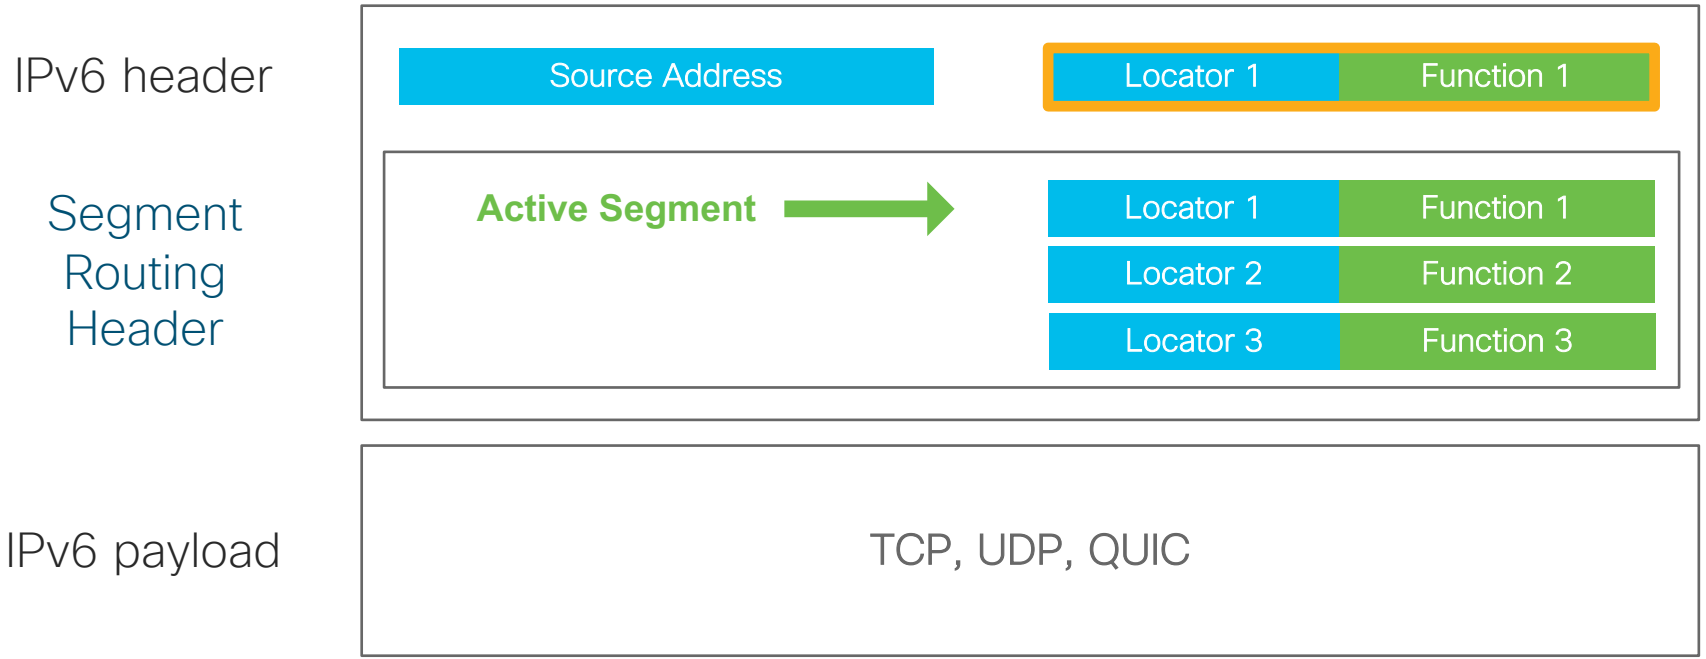

Network Program in the Packet Header

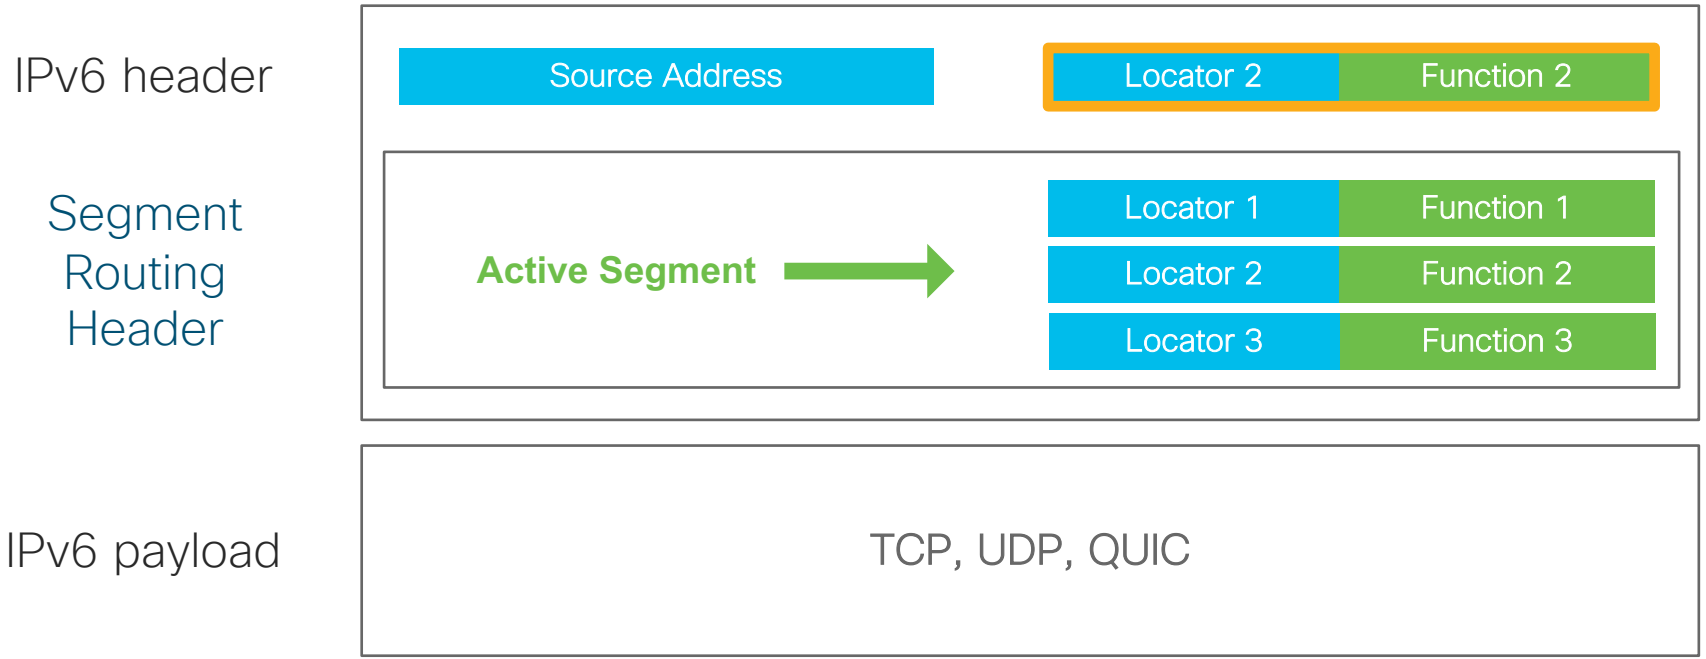

Argument shared between functions

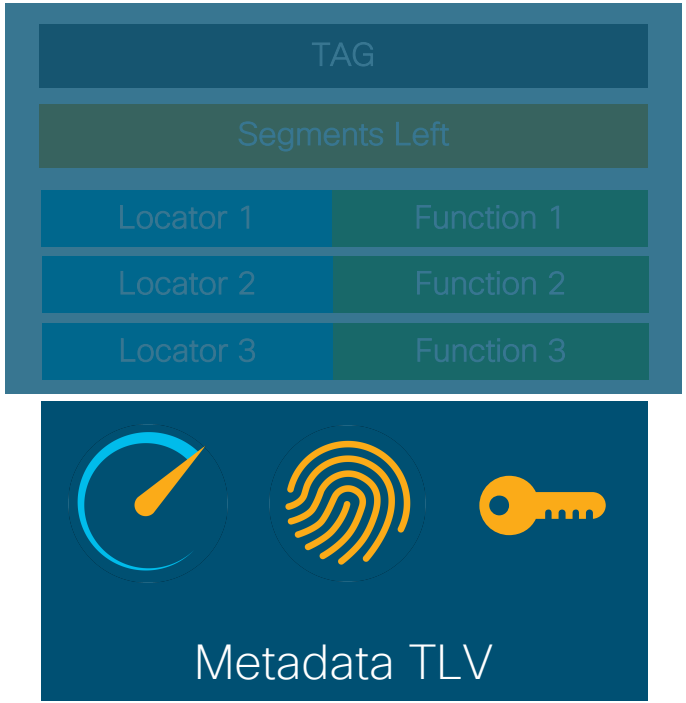

"Global"
Argument

Group-Based Policy

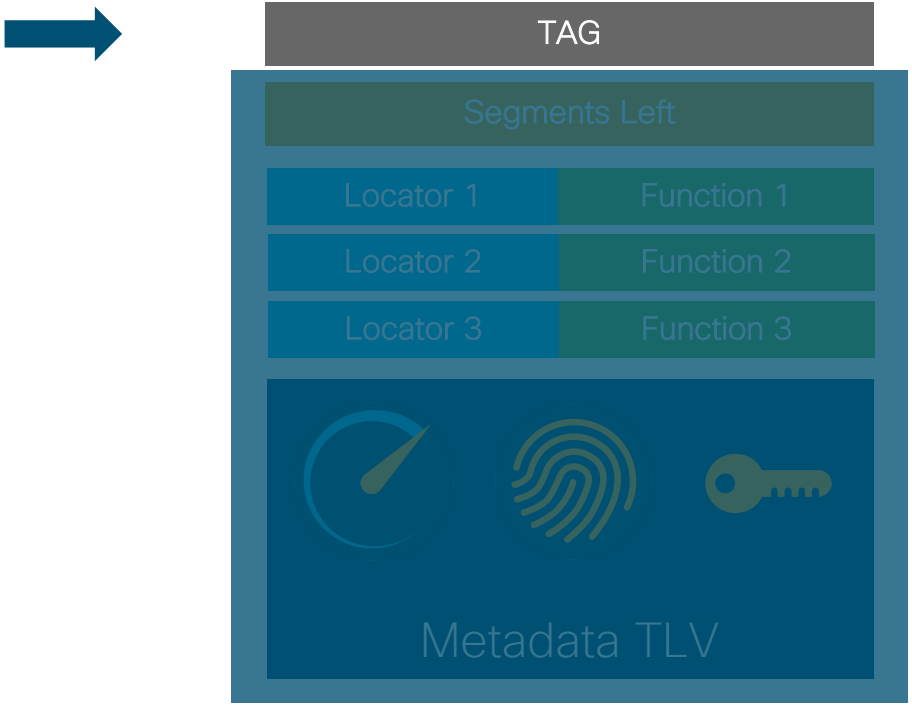

SRv6 Header

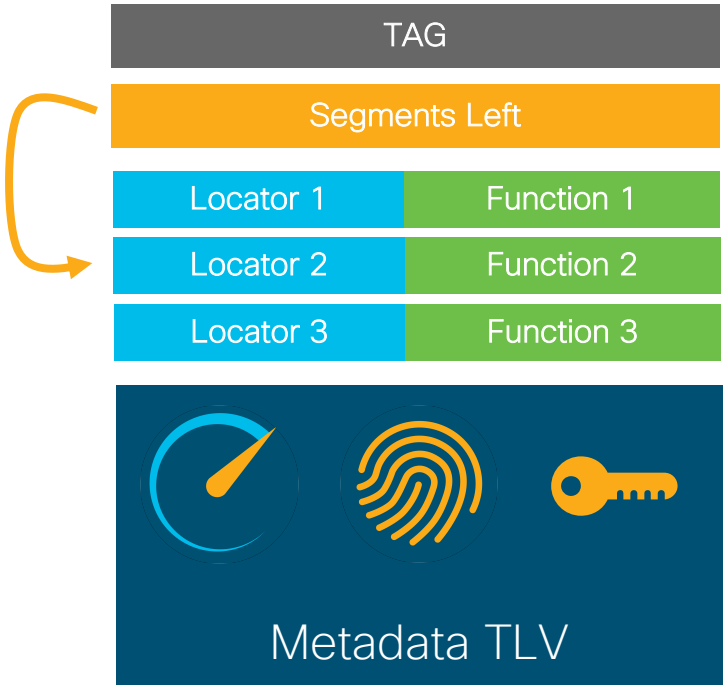

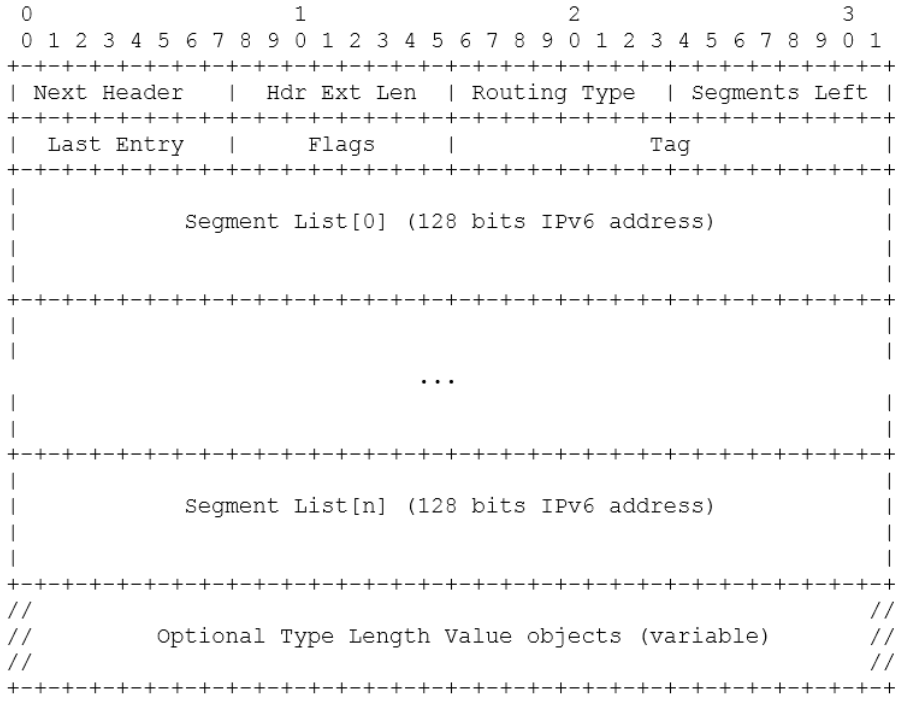

Lead Operators

- Standardization
- Multi-Vendor Consensus

SPRING
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2018

C. Filsfils
Cisco Systems, Inc.
Z. Li
Huawei Technologies
J. Leddy
Comcast
D. Voyer
D. Bernier
Bell Canada
D. Steinberg
Steinberg Consulting
R. Raszuk
Bloomberg LP
S. Matsushima
SoftBank
D. Lebrun
Universite catholique de Louvain
B. Decraene
Orange
B. Peirens
Proximus
S. Salsano
Universita di Roma "Tor Vergata"
G. Naik
Drexel University
H. Elmalky
Ericsson
P. Jonnalagadda
M. Sharif
Barefoot Networks
A. Ayyangar
Arista
S. Mynam
Innovium Inc.
W. Henderickx
Nokia
S. Ma
Juniper
A. Bashandy
K. Raza
D. Dukes
F. Clad
P. Camarillo, Ed.
Cisco Systems, Inc.
March 4, 2018

SRv6 Network Programming
draft-filsfils-spring-srv6-network-programming-04

SRv6 LocalSIDs

End and End.X SID behaviors

Illustration convention:
- IPv6 address of node k is A:<k>::
- SRv6 SID of node k is B:<k>:<function>::

- End – Default endpoint behavior
  - shortest-path to the SID's endpoint
  - endpoint updates DA with next SID
  - endpoint forwards according to updated DA

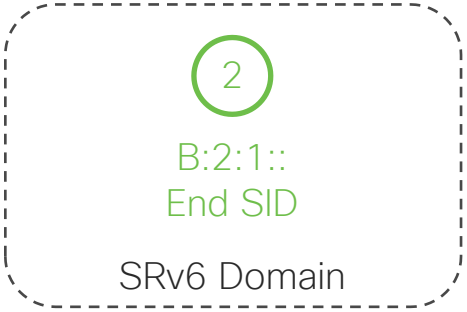

- End.X – Endpoint with cross-connect
  - shortest-path to SID's endpoint
  - endpoint updates DA with next SID
  - endpoint forwards to interface associated with SID

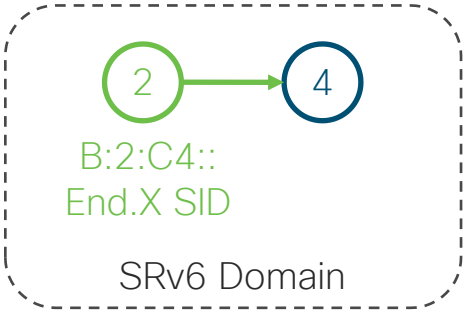

Endpoint behaviors illustration

SR: ⟨ B:4:1::, B:5:C6::, A:8:: ⟩

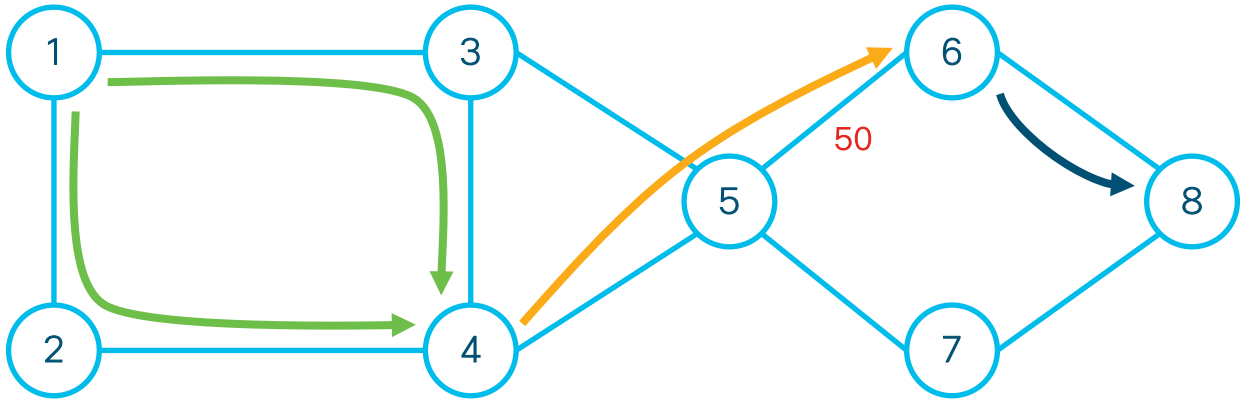

Default metric 10

- B:4:1:: shortest path to node 4
- B:5:C6:: shortest path to node 5, then cross-connect towards 6
- A:8:: regular IPv6 address of node 8

Deployment use-cases

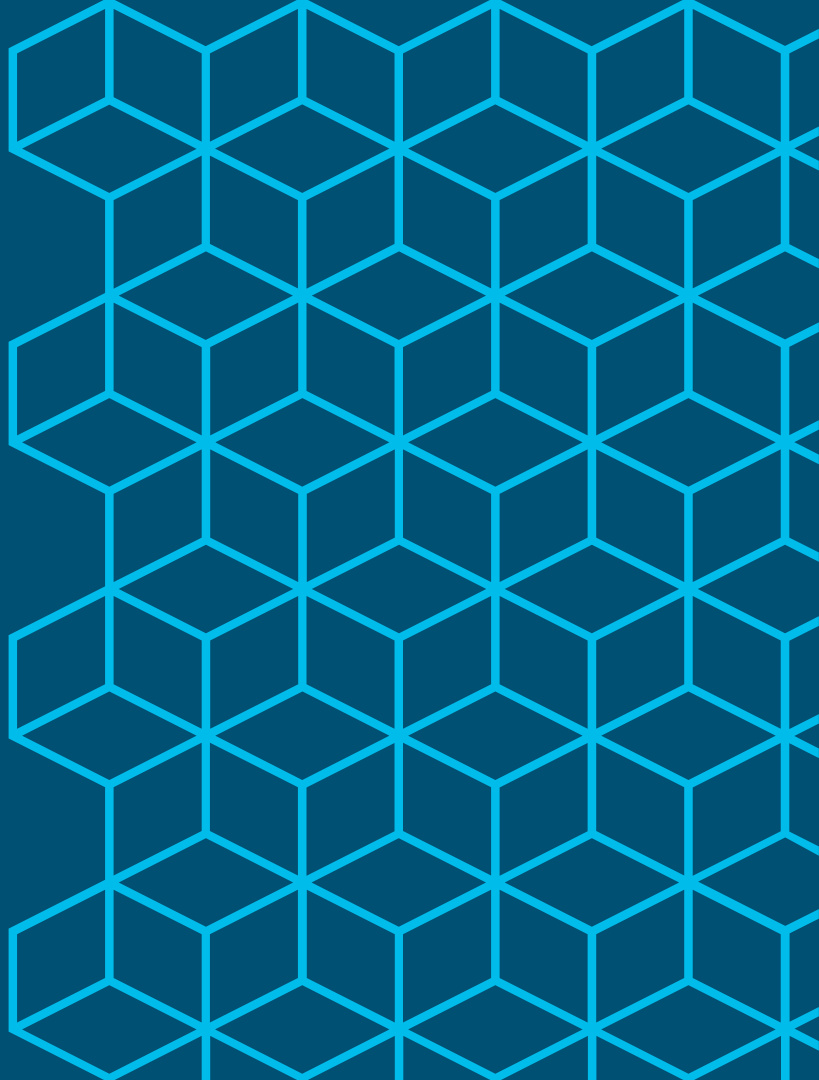

SR VPN

Overlay VPN

- Automated
  - No tunnel to configure
- Simple
  - Protocol elimination
- Efficient
  - SRv6 for everything
    - End.DT/DX behaviors

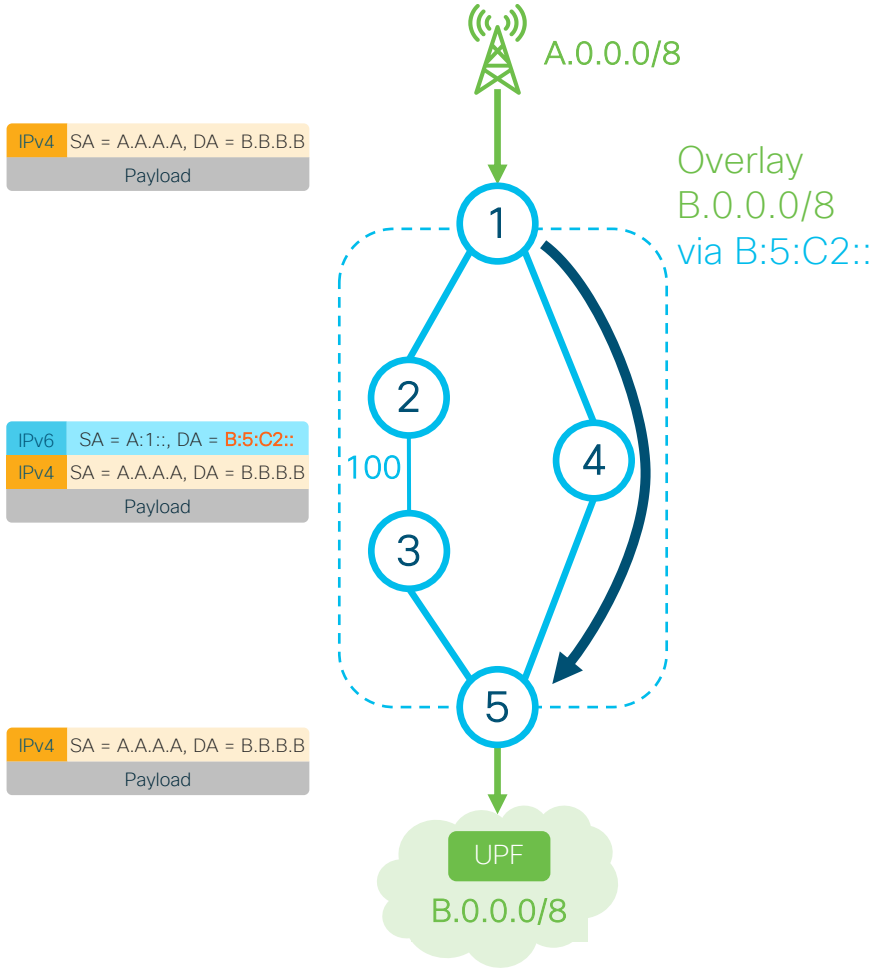

Overlay configuration

```
router bgp <inst>
  address-family vpnv4 unicast
    segment-routing srv6
      locator <name>
    !
  !
  neighbor <ipv6-addr>
    address-family vpnv4 unicast
    !
  !
  vrf <>
    address-family ipv4 unicast
      segment-routing srv6
        alloc mode {per-vrf | per-ce}
      !
  !
!
```

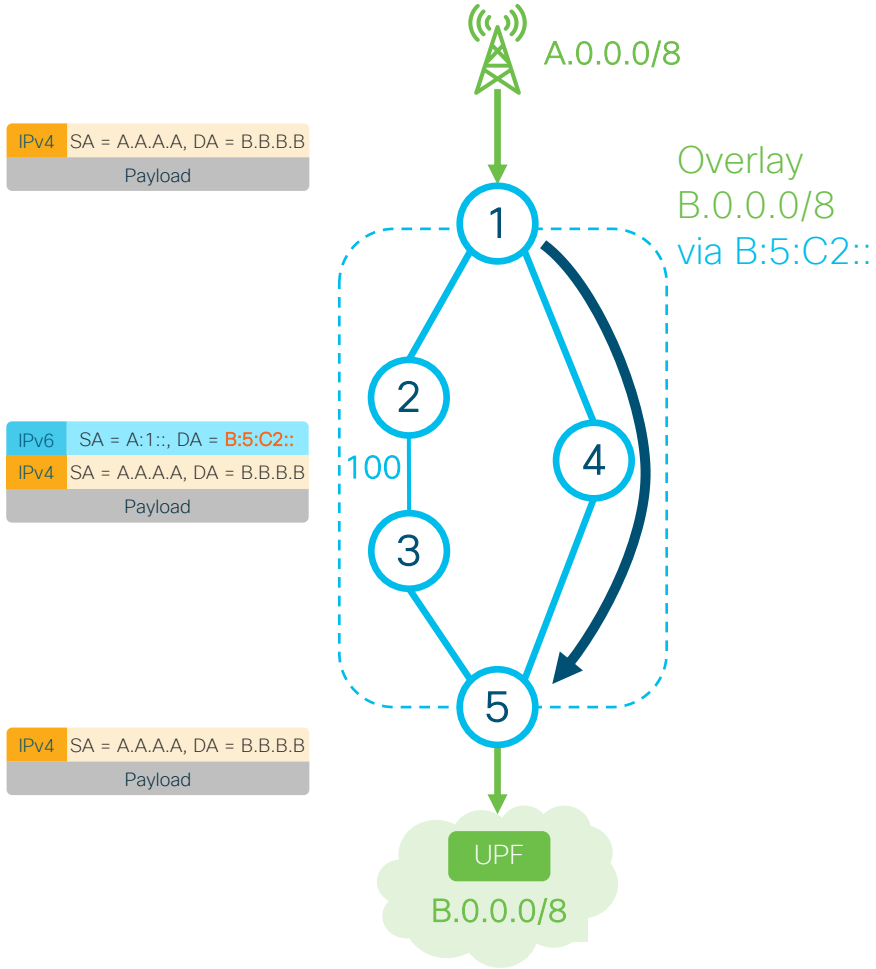

SR FRR – TILFA

TILFA

- 50msec Protection upon local link, node or SRLG failure
- Simple to operate and understand
  - automatically computed by the router's IGP process
  - 100% coverage across any topology
  - predictable (backup = postconvergence)
- Optimum backup path
  - leverages the post-convergence path
  - avoid any intermediate flap via alternate path
- Incremental deployment
- Distributed and Automated Intelligence

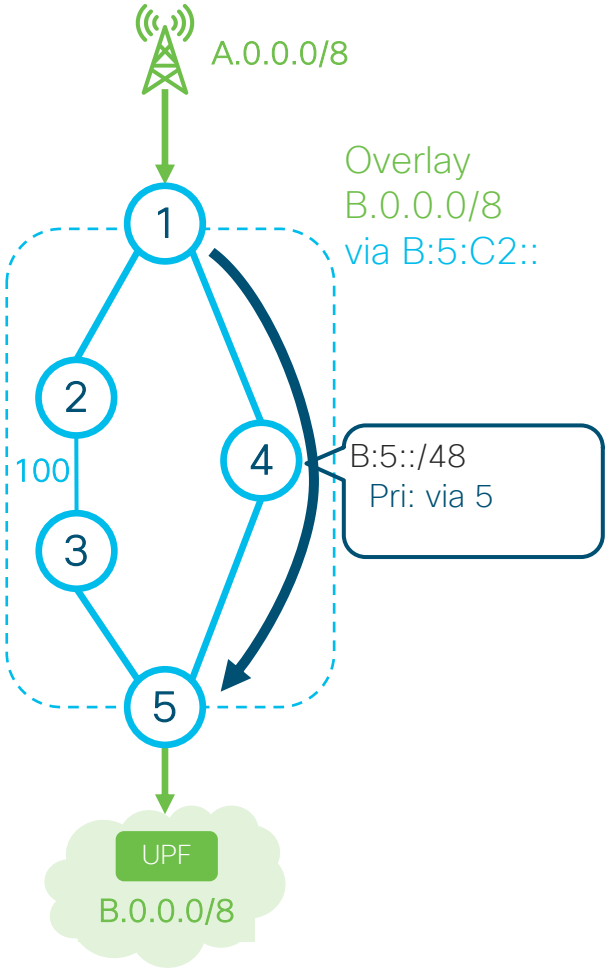

TILFA

- 50msec Protection upon local link, node or SRLG failure
- Simple to operate and understand
  - automatically computed by the router's IGP process
  - 100% coverage across any topology
  - predictable (backup = postconvergence)
- Optimum backup path
  - leverages the post-convergence path
  - avoid any intermediate flap via alternate path
- Incremental deployment
- Distributed and Automated Intelligence

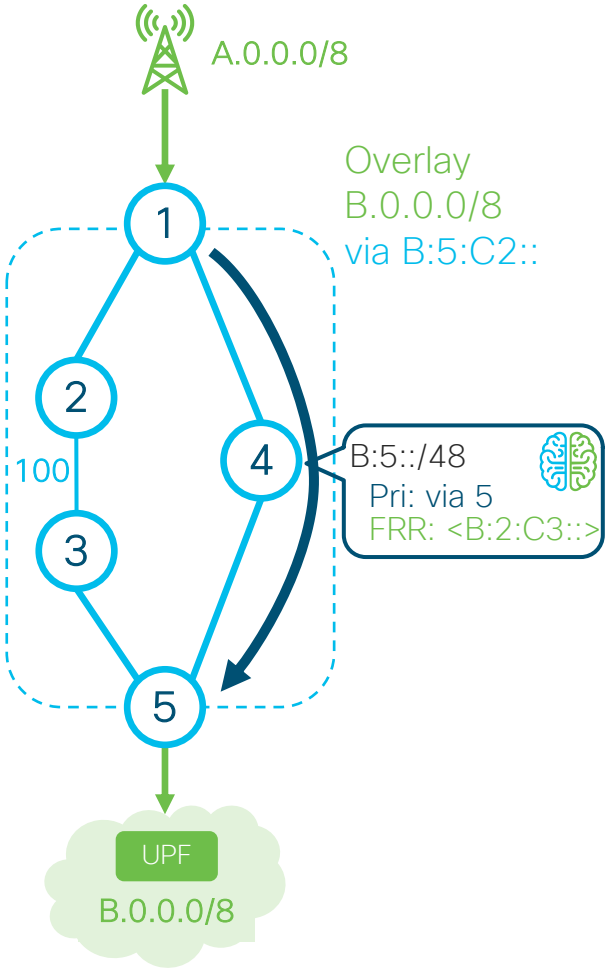

TILFA

- 50msec Protection upon local link, node or SRLG failure
- Simple to operate and understand
  - automatically computed by the router's IGP process
  - 100% coverage across any topology
  - predictable (backup = postconvergence)
- Optimum backup path
  - leverages the post-convergence path
  - avoid any intermediate flap via alternate path
- Incremental deployment
- Distributed and Automated Intelligence

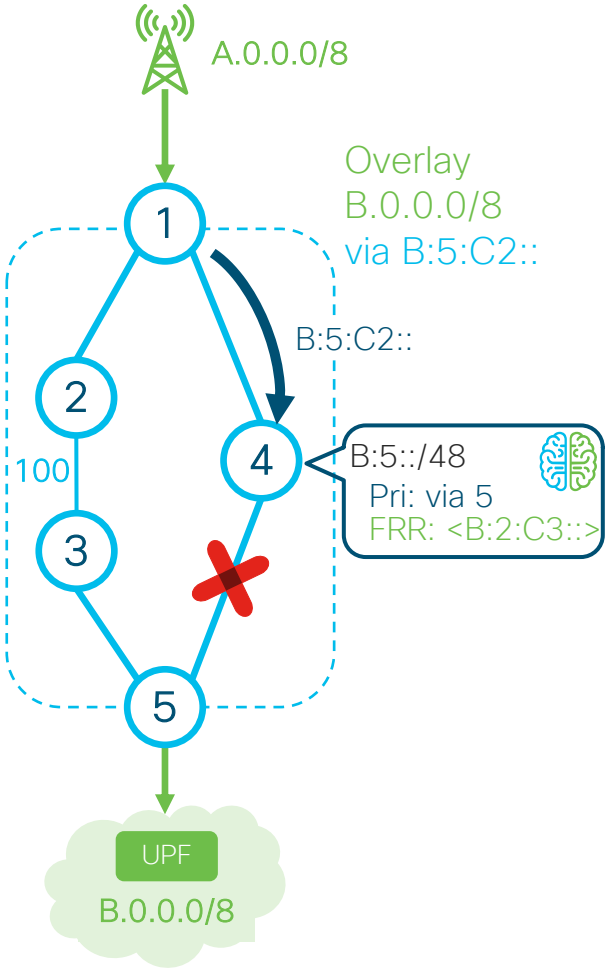

TILFA

- 50msec Protection upon local link, node or SRLG failure
- Simple to operate and understand
  - automatically computed by the router's IGP process
  - 100% coverage across any topology
  - predictable (backup = postconvergence)
- Optimum backup path
  - leverages the post-convergence path
  - avoid any intermediate flap via alternate path
- Incremental deployment
- Distributed and Automated Intelligence

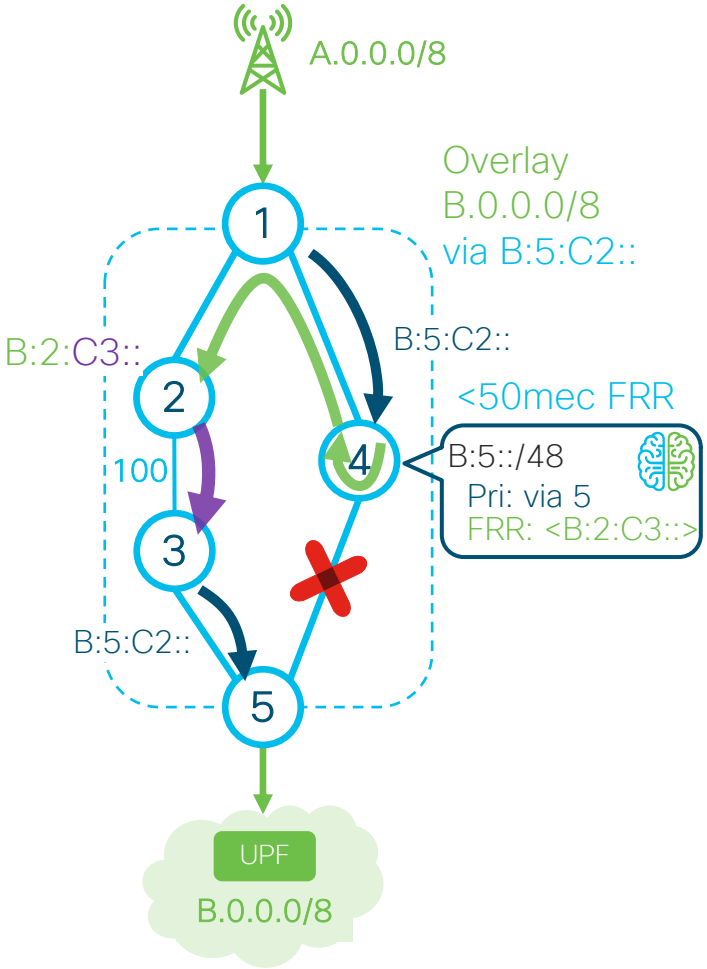

SR Control-Plane

- No new protocol
- Lightweight extension to ISIS and OSPF

Cisco Implementation

- Leveraging the existing TI-LFA SR-MPLS code
  - FCS since 2014
  - Numerous deployments

```
router isis <inst>
   interface <>
     address-family ipv6 unicast
       fast-reroute per-prefix ti-lfa
     !
   !
!
```

SRv6 VPN with integrated TE

SRv6 overlay with underlay optimization

- Automated SR TE policy triggered by SLA color of BGP route
  - No RSVP, no tunnel
- Automated Steering
  - No PBR steering complexity
- Inter-domain calculated by SR PCE
- SR native algorithms

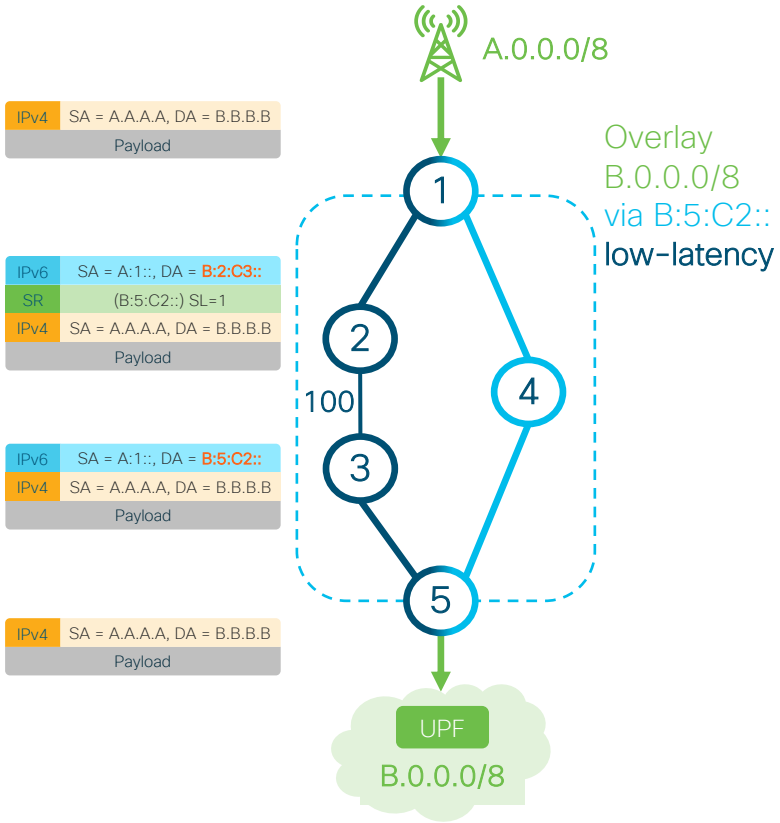

Distributed & Automated TE

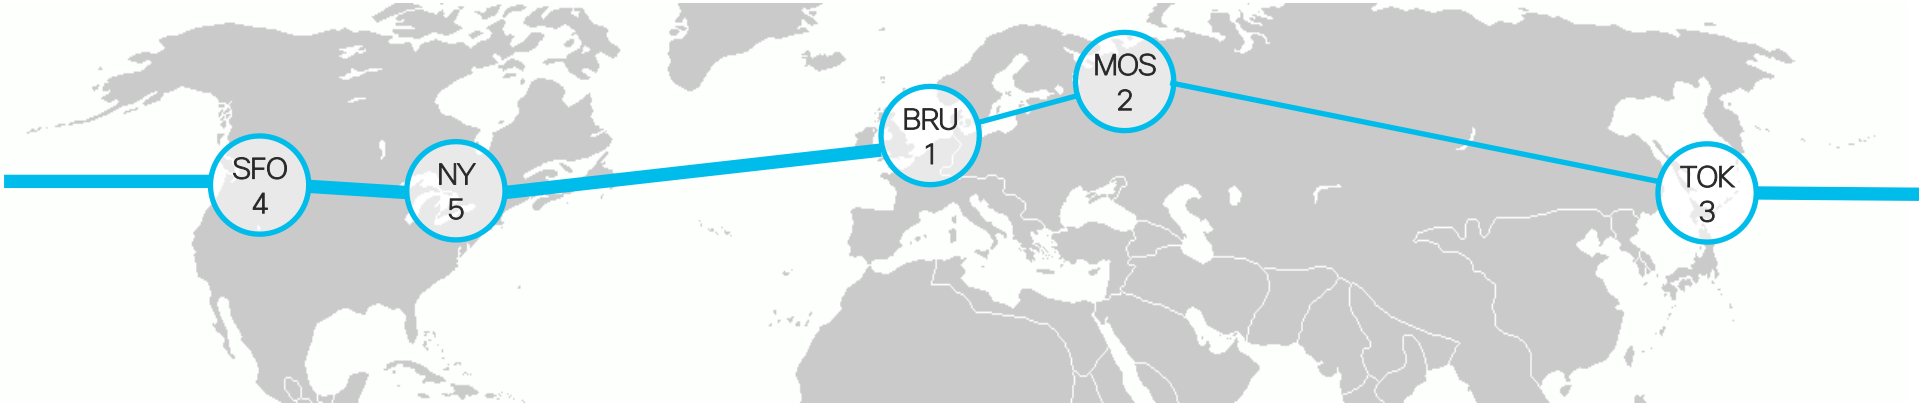

- IGP minimizes cost instead of latency

Distributed & Automated TE

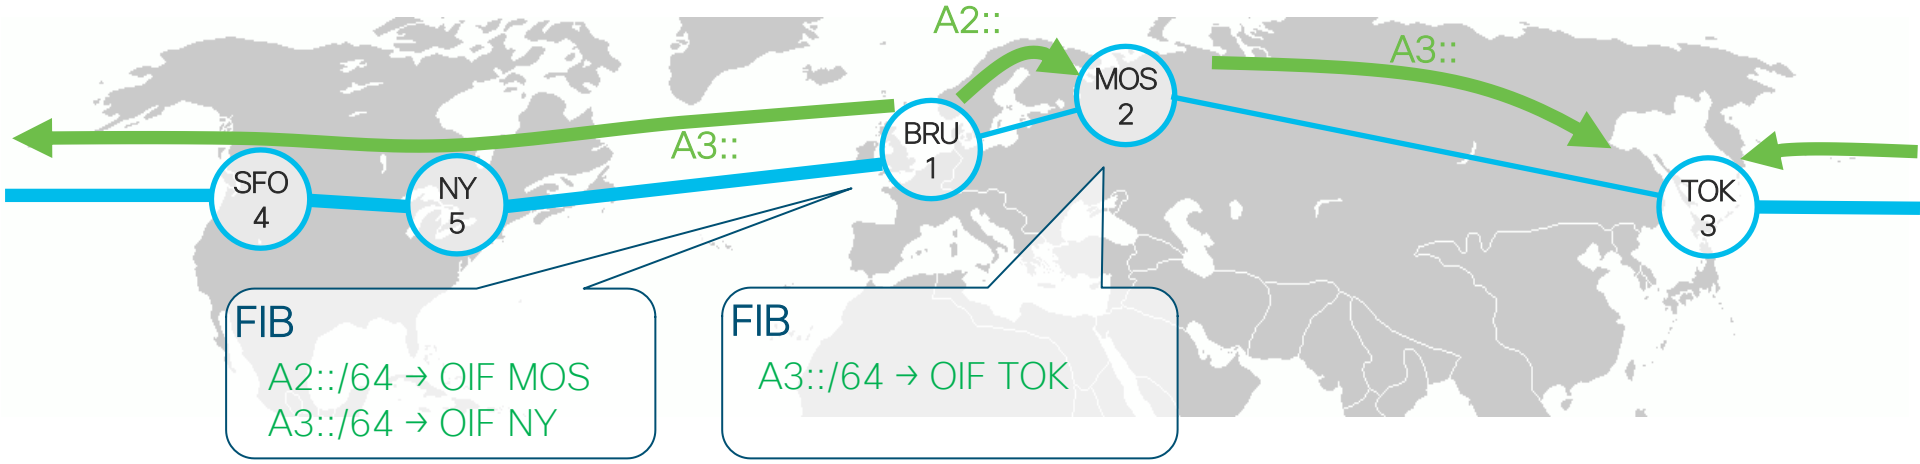

- IGP minimizes cost instead of latency

Distributed & Automated TE

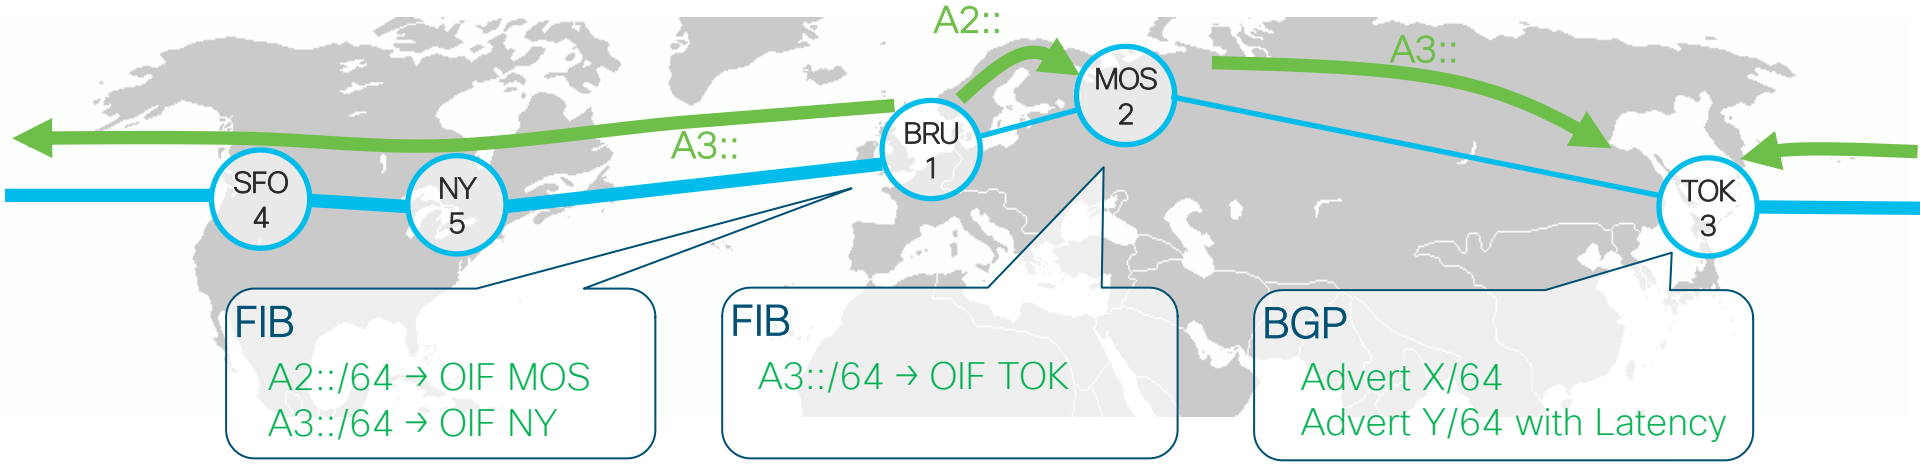

- IGP minimizes cost instead of latency

Distributed & Automated TE

On-Demand distributed TE

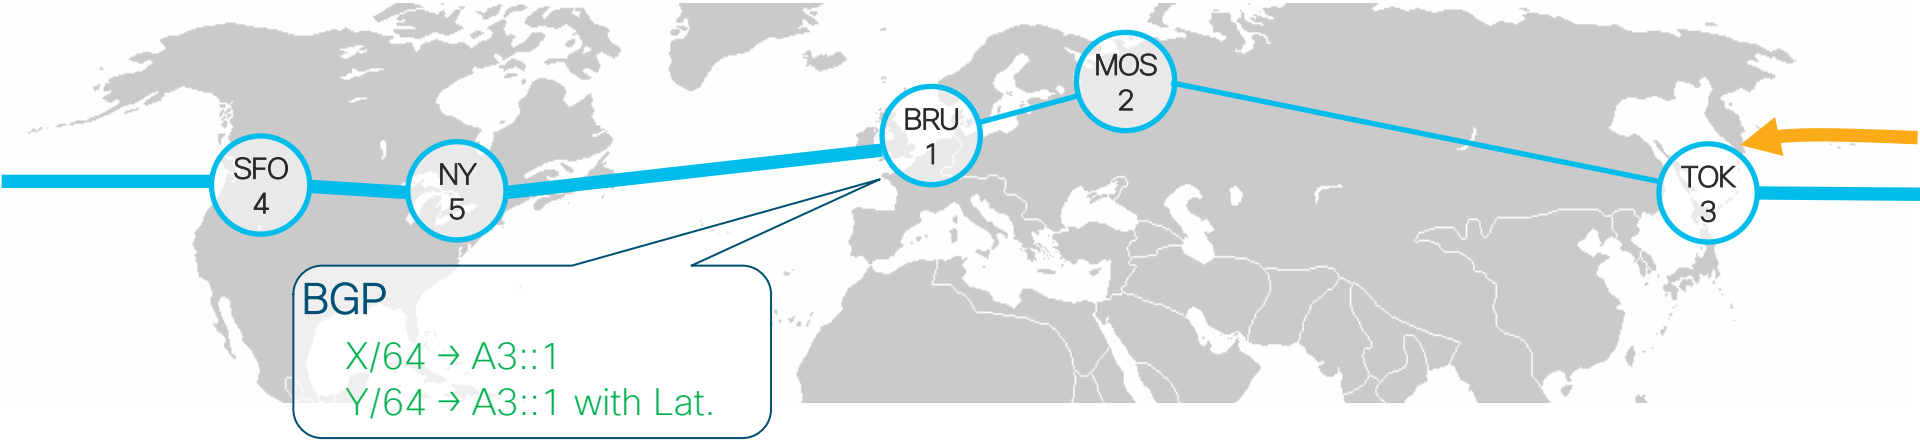

Distributed & Automated TE

On-Demand distributed TE

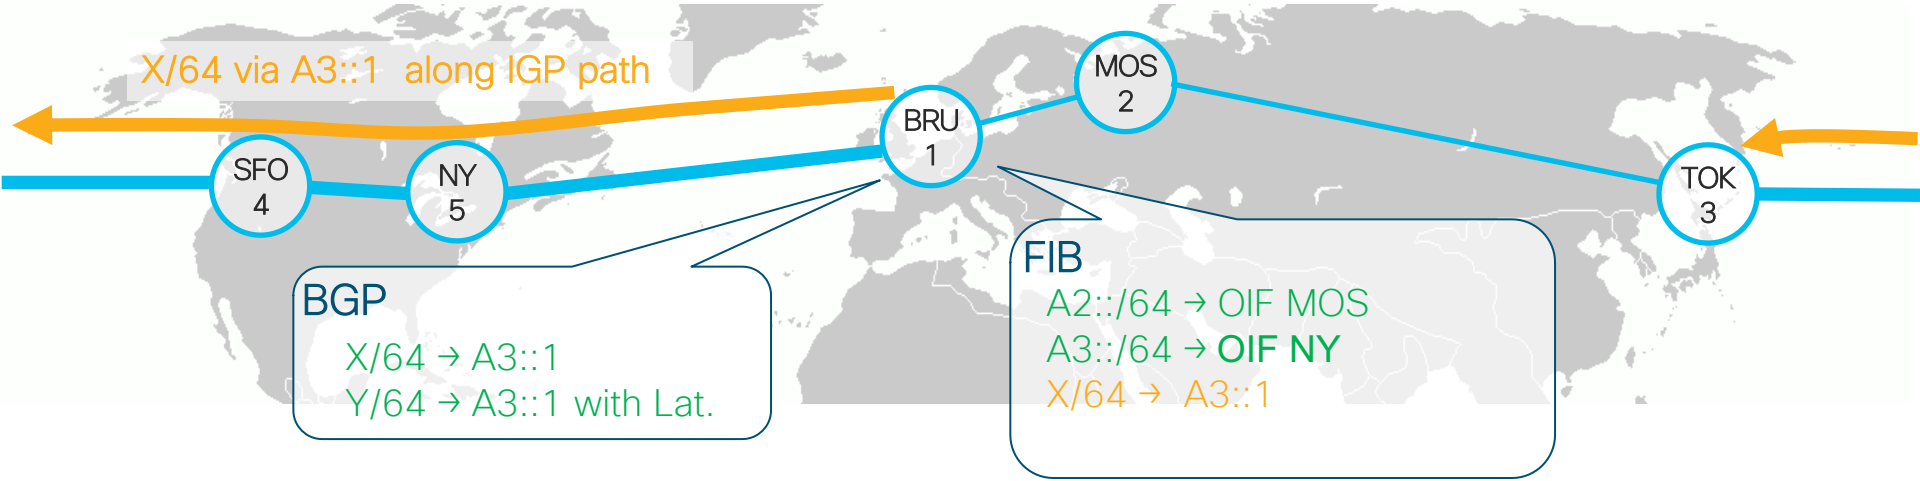

Distributed & Automated TE

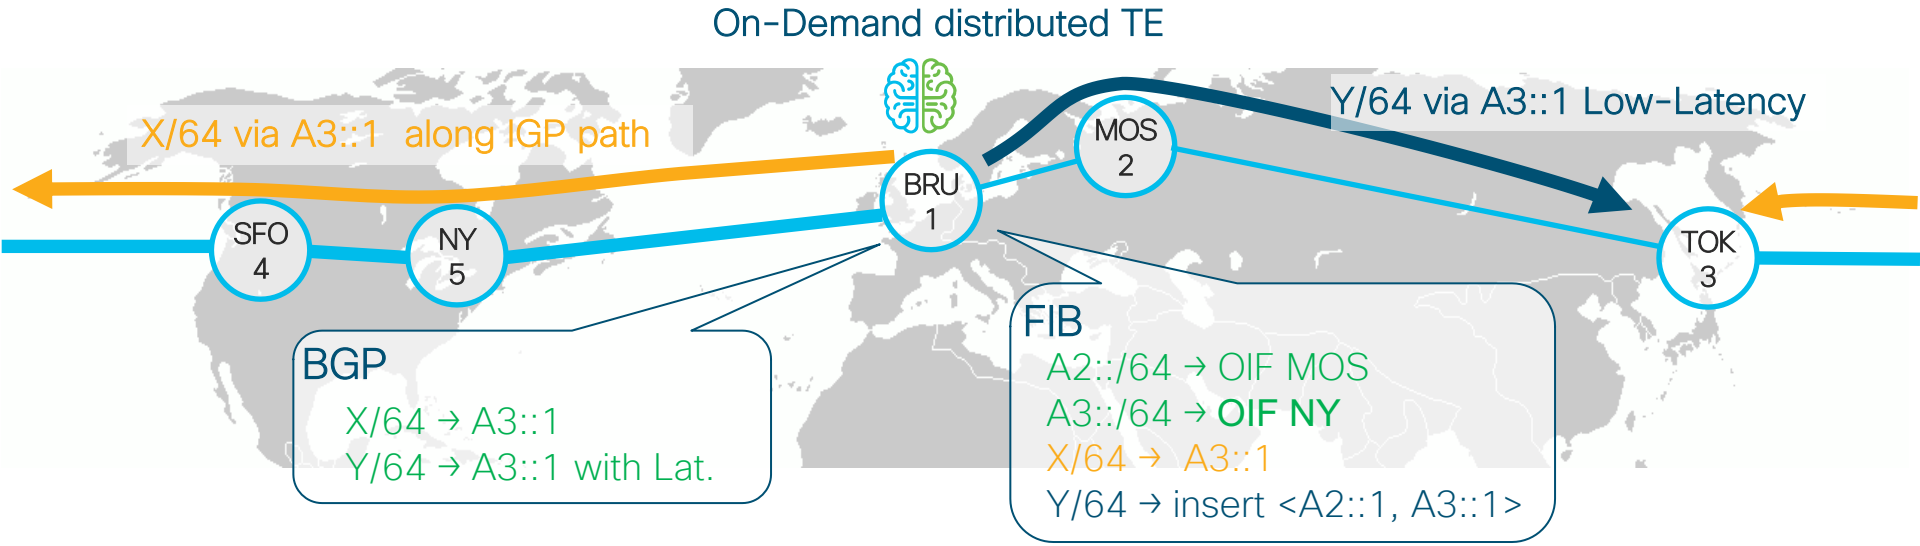

- Distributed and Automated Intelligence
- Dynamic SRTE Policy triggered by learning a BGP route with SLA contract
- No PBR steering complexity, No PBR performance tax, No RSVP, No tunnel to configure

Centralized TE

Input Acquisition

- BGP-LS
- Telemetry

Policy Instantiation

- PCEP
- BGP-TE
- Netconf / Yang

Algorithm

- SR native

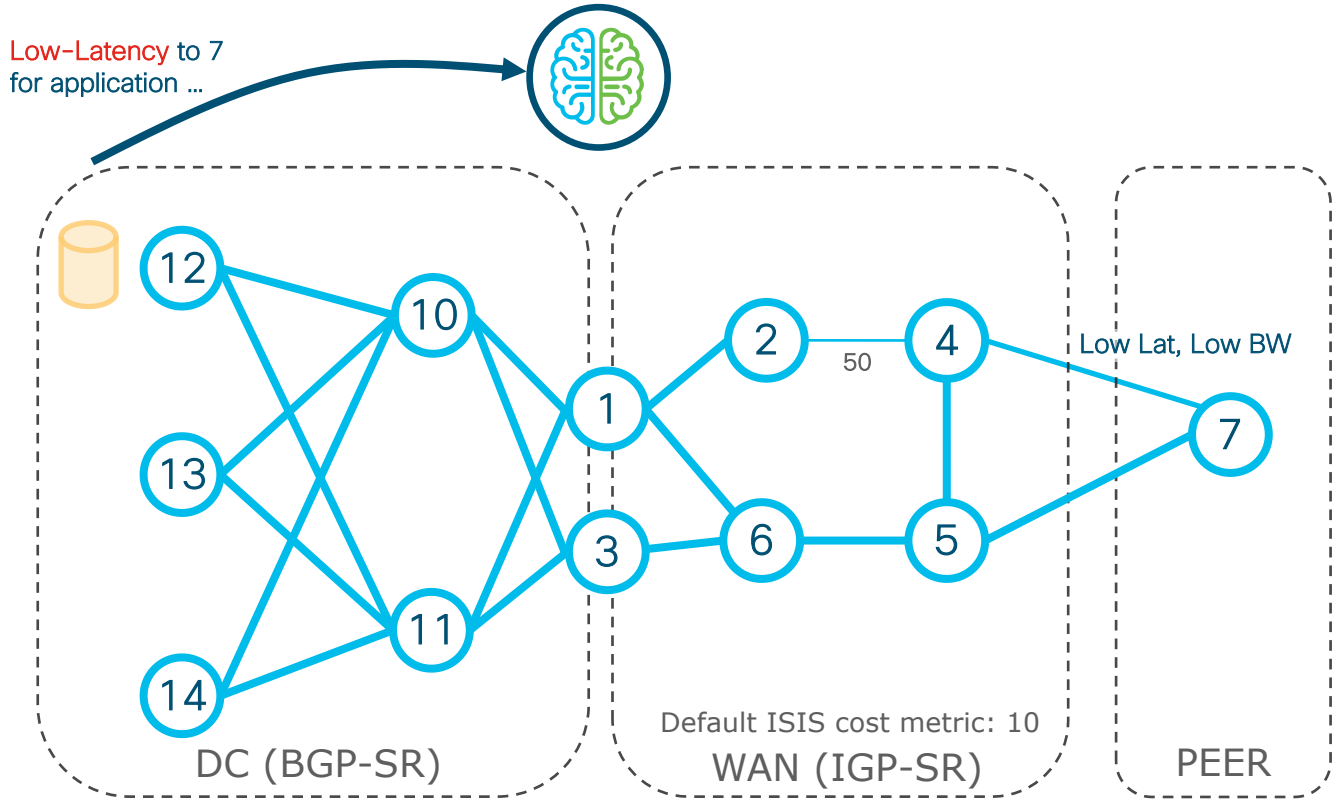

Centralized TE

Input Acquisition

- BGP-LS
- Telemetry

Policy Instantiation

- PCEP
- BGP-TE
- Netconf / Yang

Algorithm

- SR native

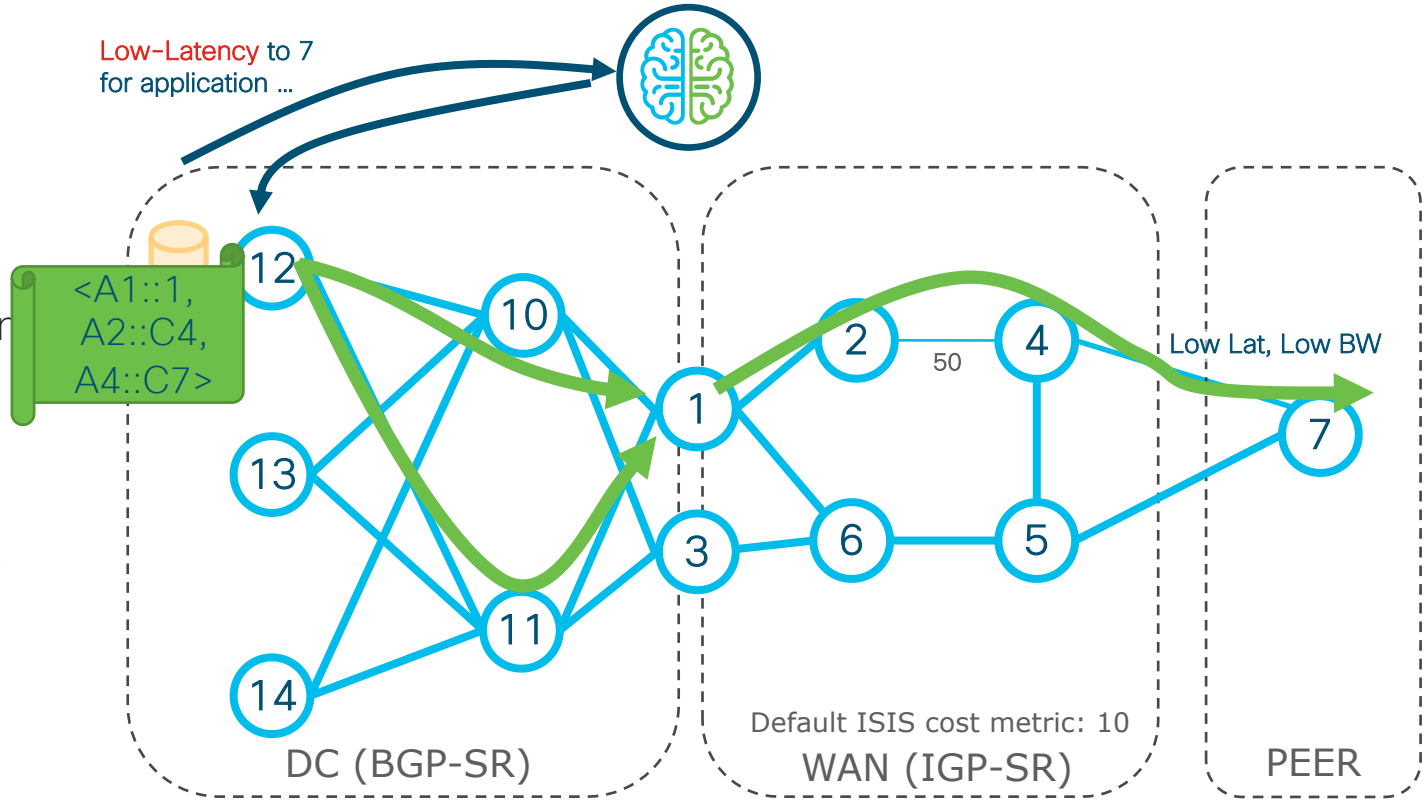

SRv6 VPN with integrated TE and NFV

SR-aware Network Functions

- SID bound to a Network Function
  - Just another type of segment
  - Stateless in the fabric
  - Seamless integration with VPN and TE
- NF can leverage the SRH
  - Implement branching operation
  - Read / write metadata

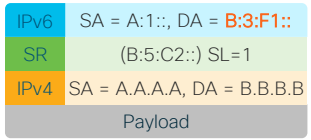

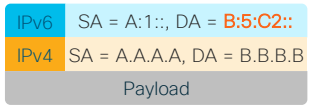

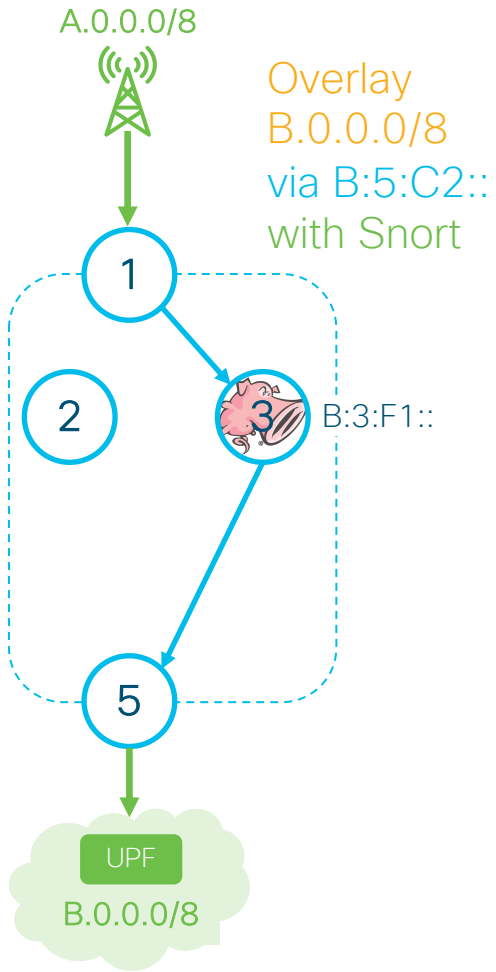

Integrated TE and NFV

- SRH may contain any combination of
  - VPN instructions
  - Topological instructions
  - SR-aware services
  - SR-unaware services
- SR for everything

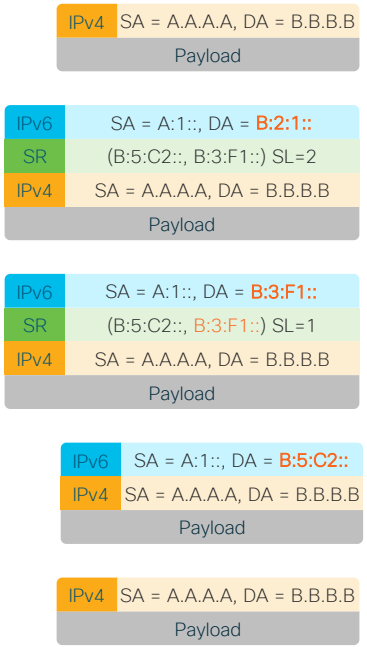

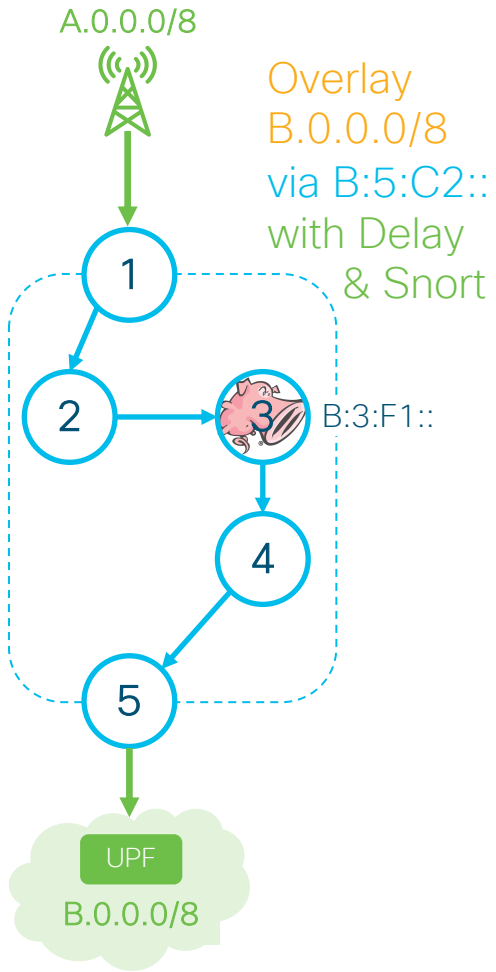

SRv6 service programming with metadata

- DPI inspects the traffic and indicates the classification result as a SID argument
- CPE may also perform some local service chaining (e.g. via Snort)
- Node 2 steers the traffic in an SR policy and copies the DA argument in an SRH TLV
- Firewall node 3 leverages this SRH metadata for smarter packet filtering

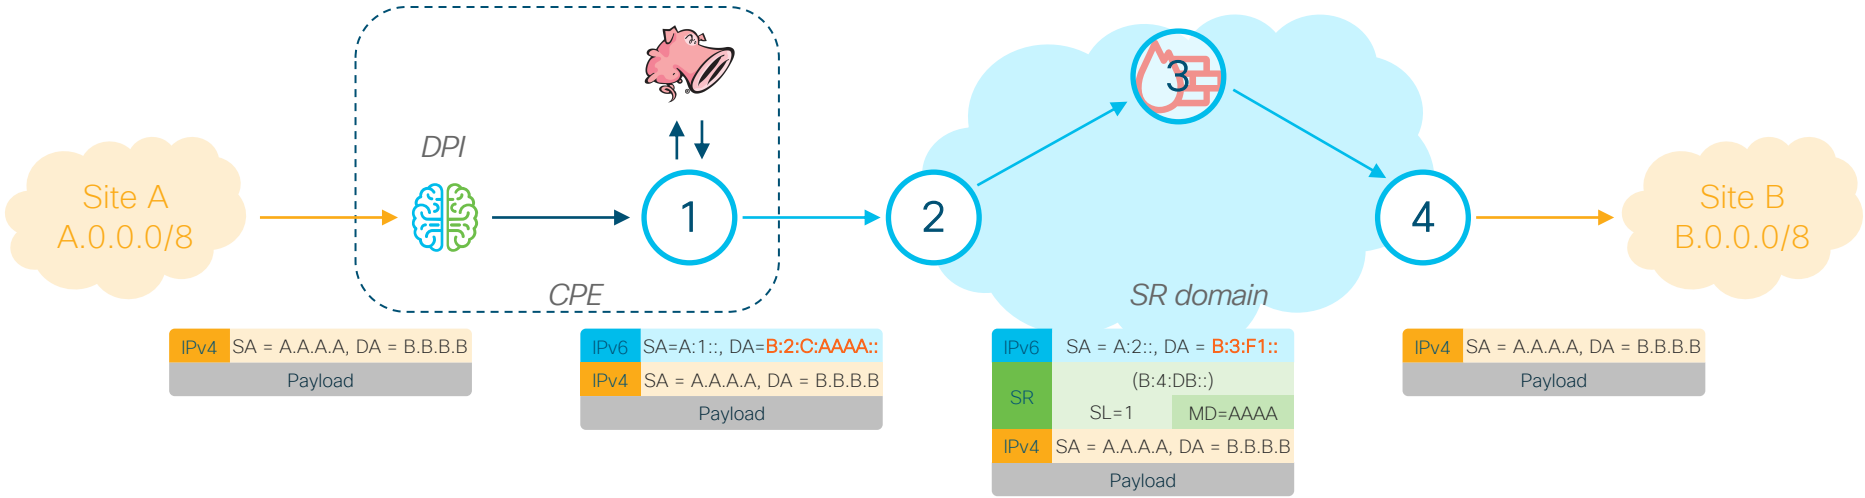

Custom SRv6 behaviors with eBPF (End.BPF)

- Associates local SRv6 SID with user-defined eBPF program
  - Leverage Extended Berkeley Packet Filter (eBPF) functionality of the Linux kernel
  - User-defined C function inserted into the networking pipeline at run-time
  - No kernel compilation required
  - Guaranteed stability
- Provides helper functions to
  - Apply basic SRv6 behaviors (End, End.X,…)
  - Steer traffic into an SR policy
  - Add, modify or delete TLVs
- Available in Linux kernel 4.18 (August 2018)

SRv6 Eco-System

Cisco FCS and in deployment

- SRv6 ISIS
- SRv6 TILFA
- SRv6 BGP L3-VPNv4
- SRv6 OAM

- More coming in CY19… ask us

Other HW

- Jericho1 and above
  - We have proven applicability by shipping it
- Barefoot
  - Interoperability shown @ Sigcomm 2017
- Huawei
  - Strong interest and declared product plan
- Intel SmartNic

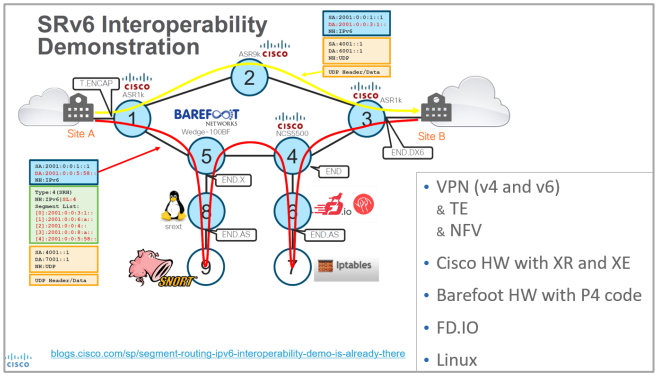

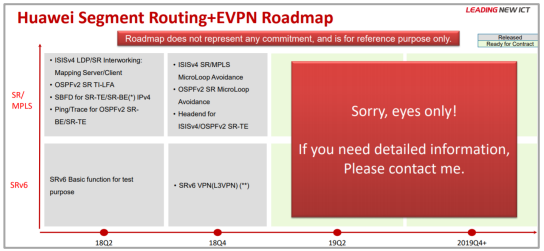

Other SW

- Linux Kernel since 4.10
  - Extensive implementation supported by Cisco Research
- FD.io VPP 17.04
  - Extensive implementation supported by Cisco
- P4 implementation
  - Extensive implementation supported by Barefoot/Toyota
- Container networking

Also in the DC – with linerate SRv6 @ 400G

- Amazing set of SRv6 network instructions @ 400G !

SRv6 is happening in 2019!

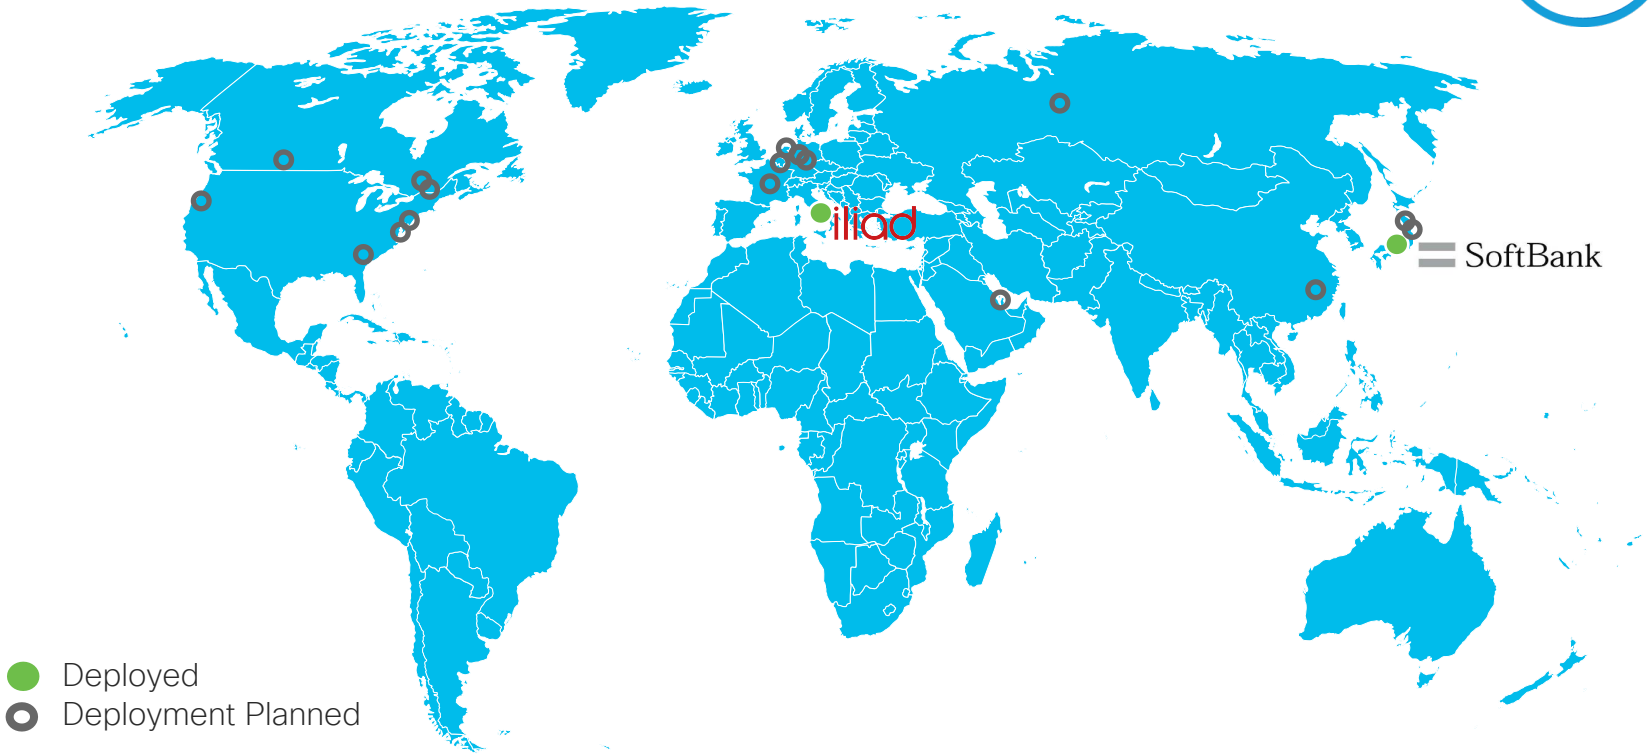

Cisco Supports SoftBank on First Segment Routing IPv6 Deployment in Prep for 5G

Link to PR - https://newsroom.cisco.com/press-release-content?type=webcontent&articleId=1969030

Cisco Supports SoftBank on First Segment Routing IPv6 Deployment in Prep for 5G

NEWS PROVIDED BY
Cisco Systems, Inc. →
Feb 24, 2019, 02:00 ET

SHARE THIS ARTICLE

BARCELONA, Spain, Feb. 24, 2019 /PRNewswire/ -- Mobile World Congress -- Cisco announced today its collaboration with SoftBank on the world's first Segment Routing IPv6 (SRv6) deployment.

With the anticipation of the coming 5G era, Cisco has been assisting SoftBank to deploy state-of-the-art Segment Routing IPv6 (SRv6) networks nationwide to build a future network architecture that is extremely scalable, with improved reliability, flexibility and agility, all while helping to reduce CapEx and OpEx.

Current mobile networks are deployed as divided networks, with several layers and complicated control plane processing, which makes it difficult to respond to strict quality requirements like in the case of 5G. Deploying SRv6 in a 5G mobile network simplifies network layers and integrates user plane functions from end-to-end with only IPv6 protocol, making things simple, controllable, and flexible.

"Converging 5G features into the end-to-end IPv6 layer with Segment Routing capabilities, is the key to embodying 5G in a simple, scalable architecture," **said Mr. Junichi Miyakawa, Representative Director & CTO for SoftBank.** "With the depth of portfolio and strong network knowledge that Cisco brings to the table, we knew together we could bring our vision to life."

"SoftBank has kept an intense focus on improving service quality for its customers, which can be challenging when trying to reduce costs," **said Sumeet Arora, Senior Vice President of Service Provider Networks, Cisco.** "With the launch of SRv6 network programming, it is pioneering the next phase of IP networking through automation, and championing

Thanks to SRv6 network programming capabilities, Iliad is set to further disrupt the mobile market by delivering truly innovative service offerings

Iliad's NodeBox is SRv6 enabled

https://newsroom.cisco.com/press-release-content?type=webcontent&articleId=1978361

iliad

CISCO

Conclusion

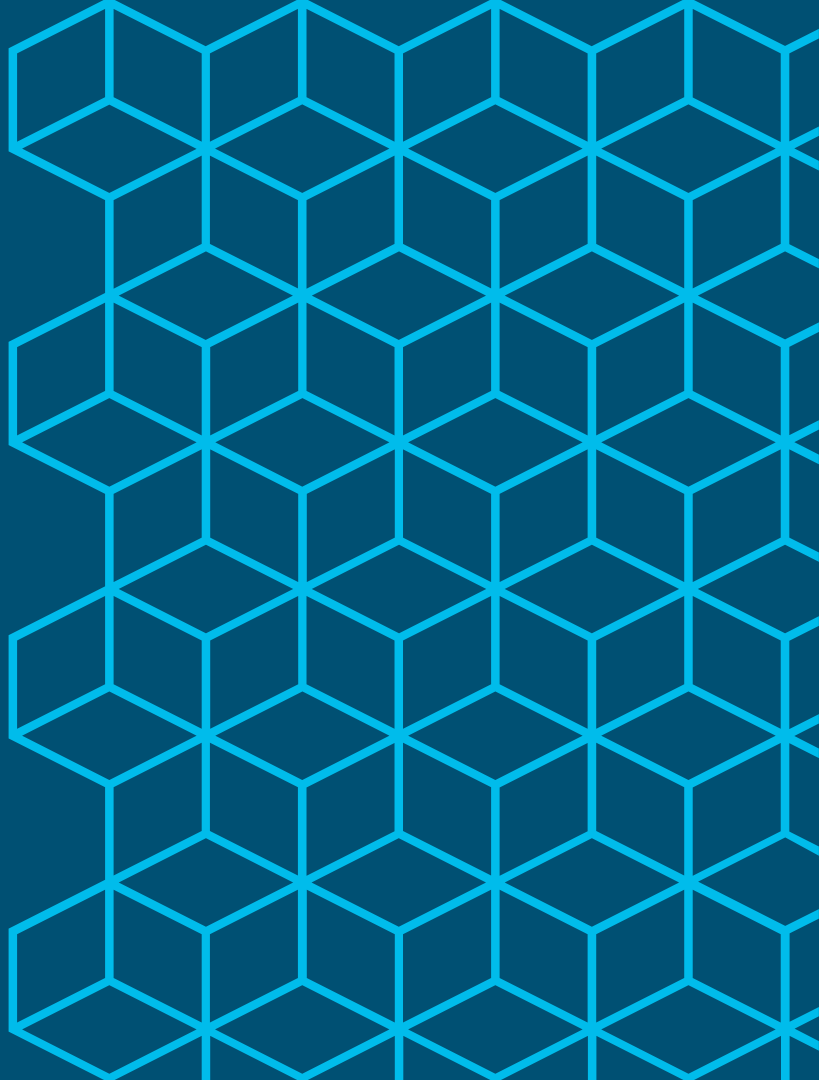

Simplicity always prevails

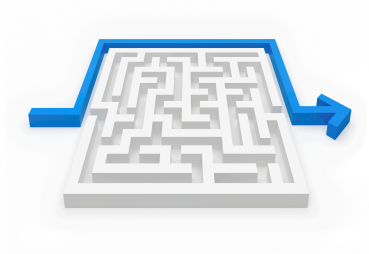

Furthermore with more scale and functionality

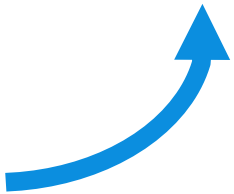

Segment Routing conclusion

- Strong industry support
- Fantastic deployment rate
- Bold architecture: network programming
- Numerous use-cases
  - FRR, TE, SDN, Overlay with SLA, NFV, Spray, SD-WAN, 5G & NS, ...
- Feel free to join the lead-operator team!

Partnering

- Track-record collaboration with operator
  - Focus on real operator needs
  - Seamless Deployment
  - Standardization
  - Multi-Vendor consensus
- Looking forward to working together

Stay up-to-date

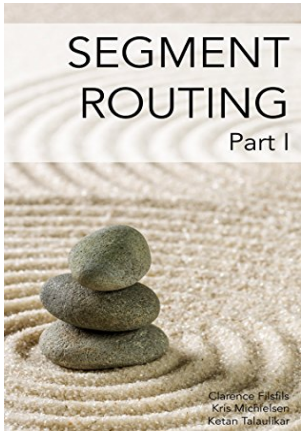

amzn.com/B01I58LSUO

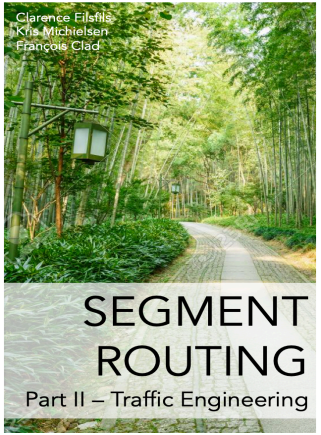

amazon.com/dp/B07N13RDM9

twitter.com/SegmentRouting

SR
Segment Routing

segment-routing.net

facebook.com/SegmentRouting/

linkedin.com/groups/8266623

Thank you!

ask-segment-routing@cisco.com