# SRv6

Clarence Filsfils
Cisco Fellow – cf@cisco.com

# Industry at large backs up SR

Strong customer adoption

WEB, SP, DC, Metro, Enterprise

De-facto SDN Architecture

Standardization

IETF

Multi-vendor Consensus

Open Source

Linux, VPP

Bell

SoftBank

vodafone

orange™

Alibaba.com

Google

Walmart

# Segment Routing

- Source Routing
  - the topological and service (NFV) path is encoded in the packet header
- Scalability
  - the network fabric does not hold any per-flow state for TE or NFV
- Simplicity
  - automation: TILFA
  - protocol elimination: LDP, RSVP-TE, NSH…
- 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 RFC8200 provision for source routing extension header
- 1 segment = 1 address
- a segment list = an address list in the SRH

# Objective of SRv6

# IPv6 provides reachability

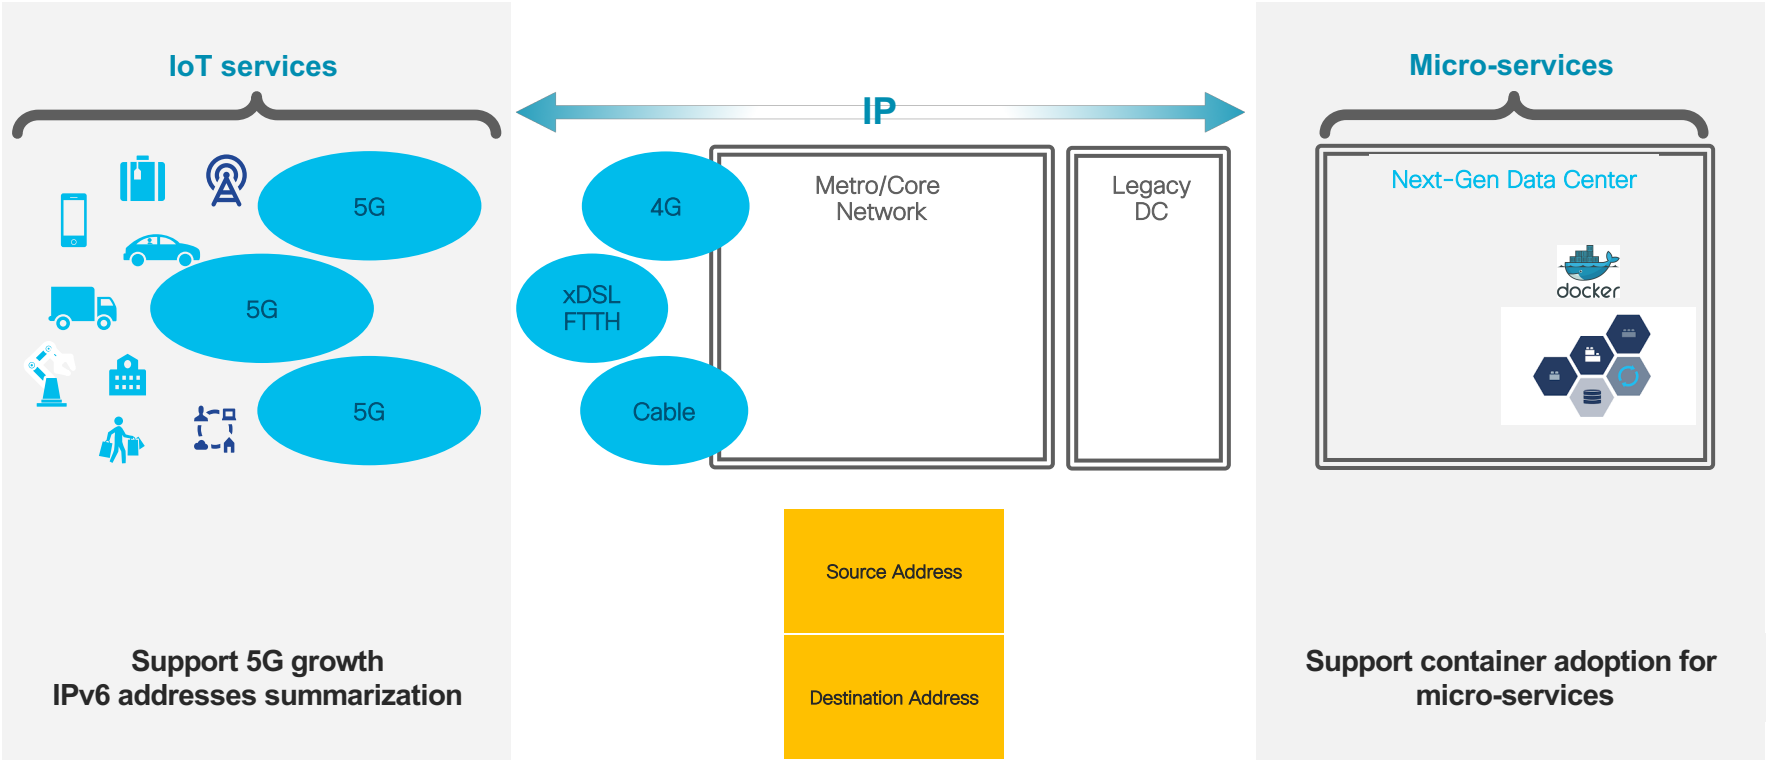

# SRv6 for underlay

SRv6 for Underlay

IPv6 for reachability

Simplification, FRR, TE, SDN

Non-redundant states scaling in k*N^2

# SRv6 for underlay and overlay

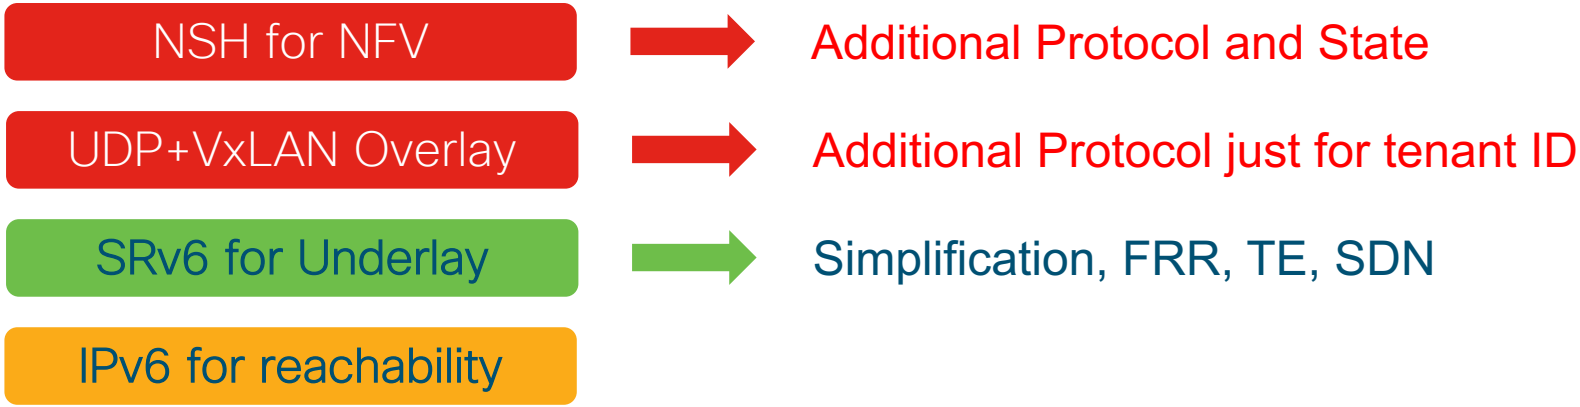


Multiplicity of protocols and states hinder network economics

# SRv6 for underlay and overlay

# Network Programming

Network as a computer – SR for anything

# 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 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
  - Arguments: optional argument bits to be used only by that SID
  - Flexible bit-length selection

# Network Program in the Packet Header

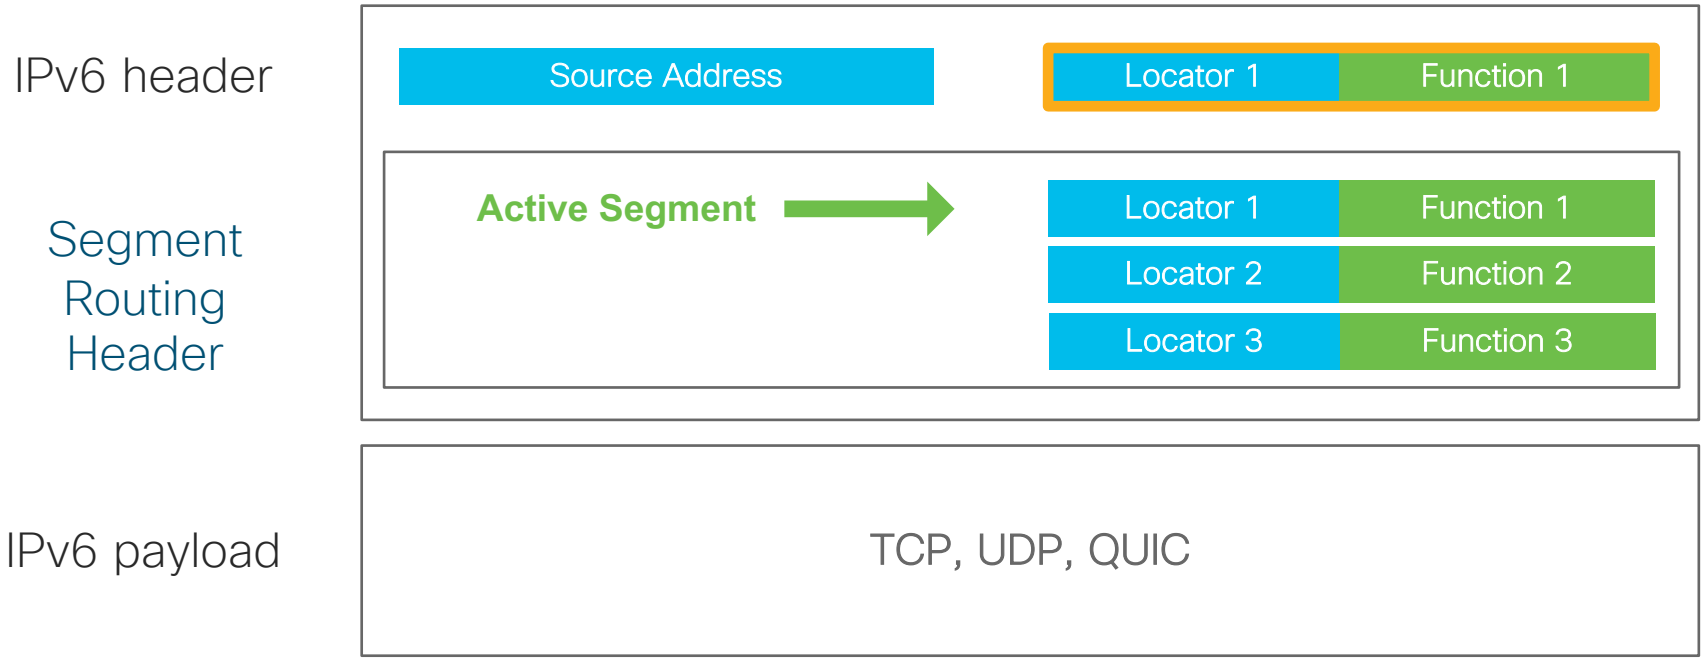

# Network Program

**Next Segment** →

| Locator 1 | Function 1 |
| --- | --- |
| Locator 2 | Function 2 |
| Locator 3 | Function 3 |

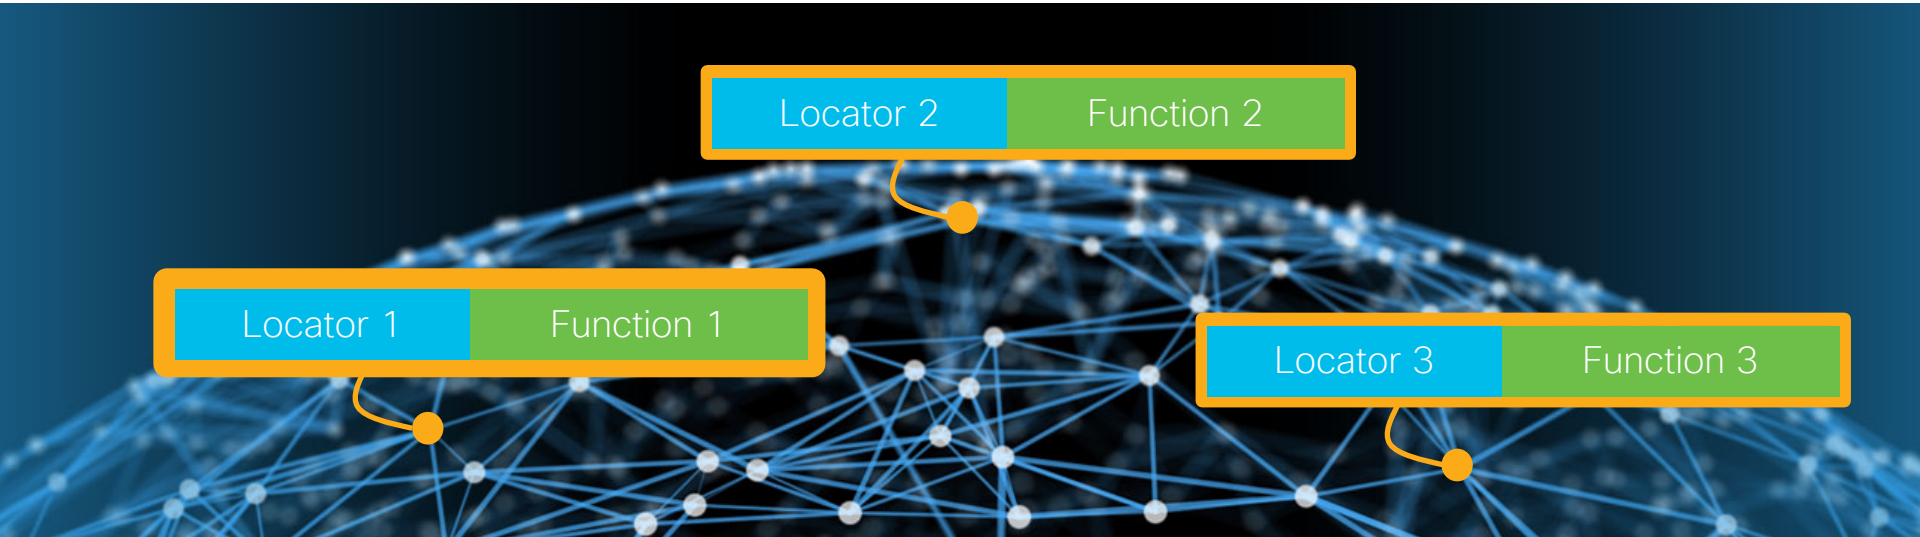

# Network Program

**Next Segment**

| Locator 1 | Function 1 |
| --- | --- |
| Locator 2 | Function 2 |
| Locator 3 | Function 3 |

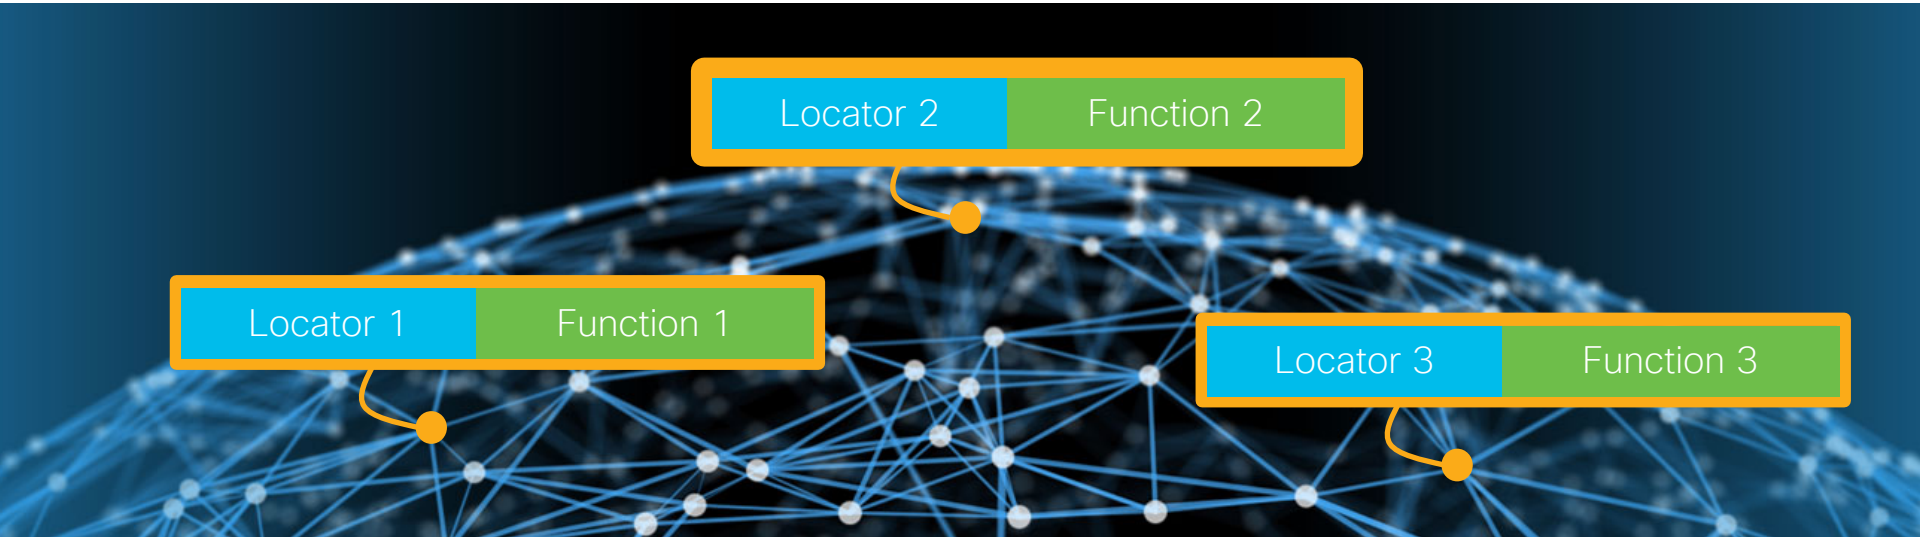

# Network Program

| Locator 1 | Function 1 |
| --- | --- |
| Locator 2 | Function 2 |
| Locator 3 | Function 3 |

**Next Segment** →

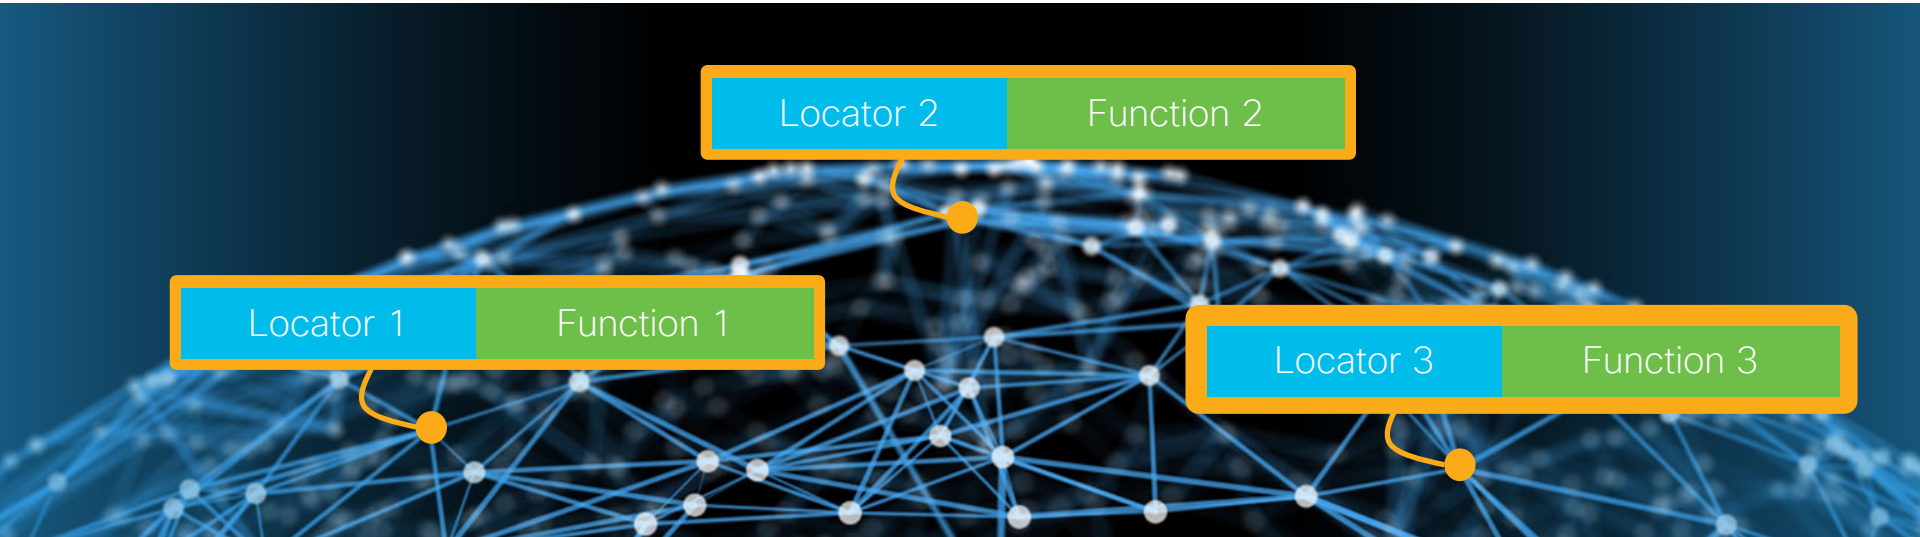

# Argument shared between functions

“Global” Argument

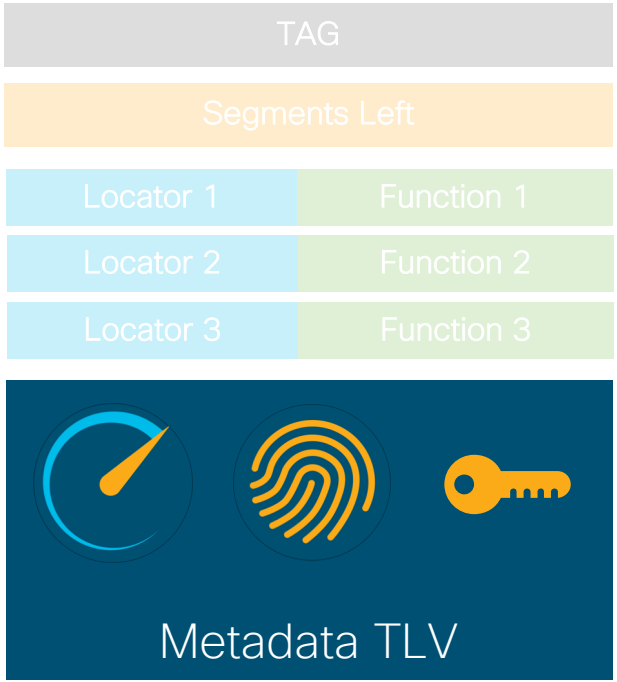

# Group-Based Policy

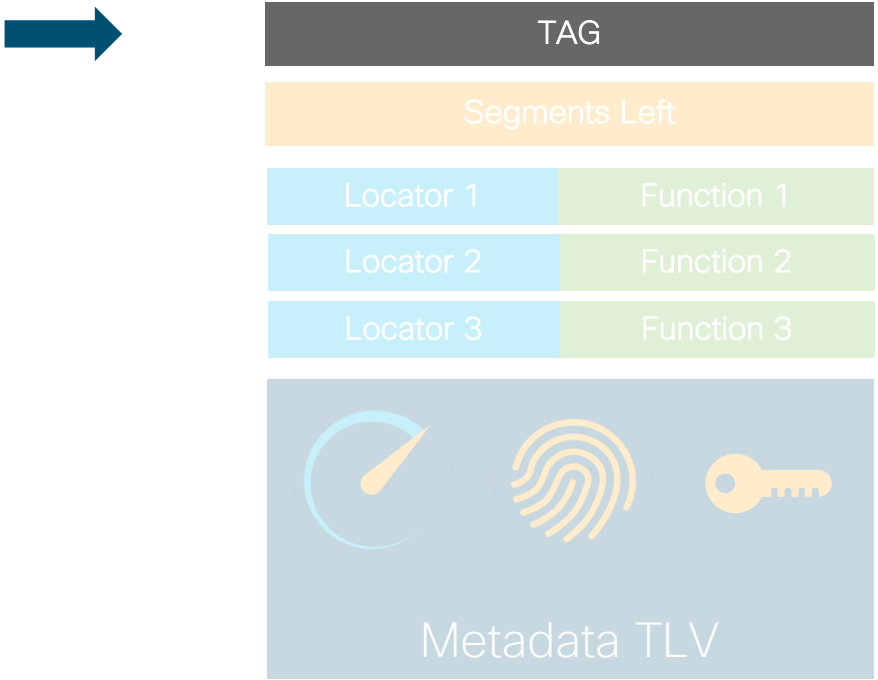

# SRv6 Header

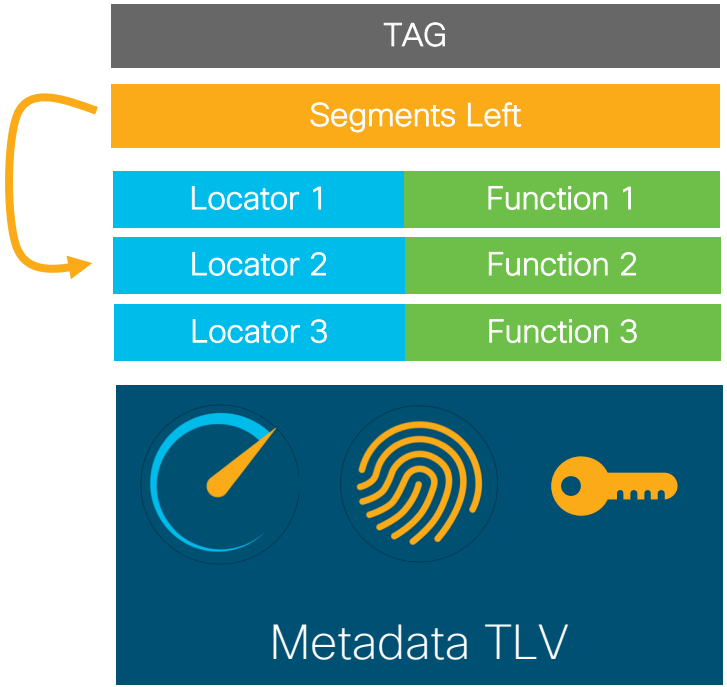
TAG
Segments Left
Locator 1
Function 1
Locator 2
Function 2
Locator 3
Function 3
Metadata TLV

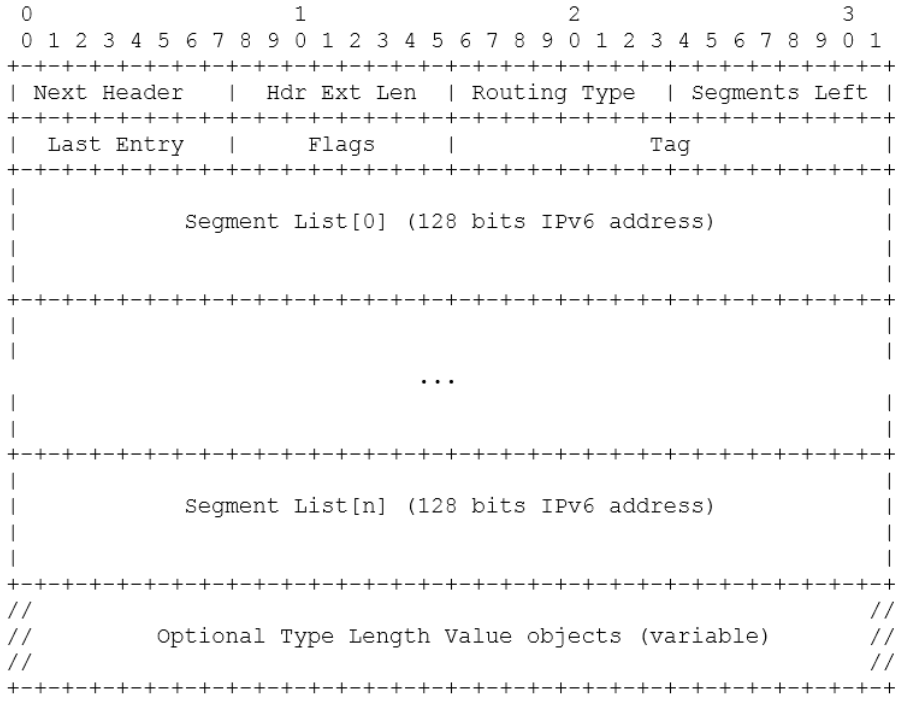
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Next Header
Hdr Ext Len
Routing Type
Segments Left
Last Entry
Flags
Tag
Segment List[0] (128 bits IPv6 address)
...
Segment List[n] (128 bits IPv6 address)
Optional Type Length Value objects (variable)

# Load-balancing

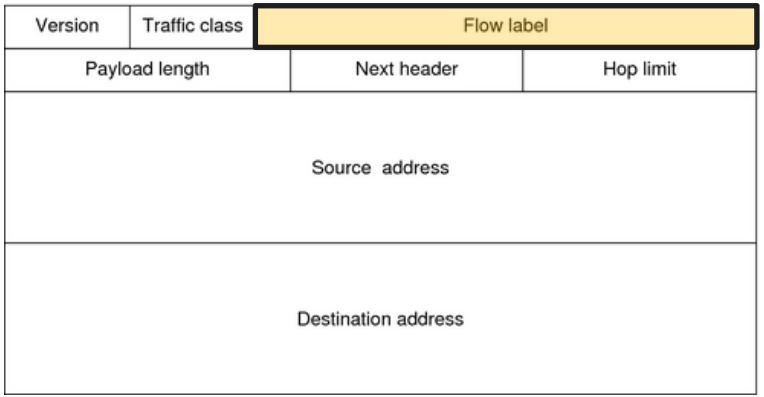


- 20-bit entropy
- No additional protocol
  - infamous mpls entropy label

# SRv6 for anything

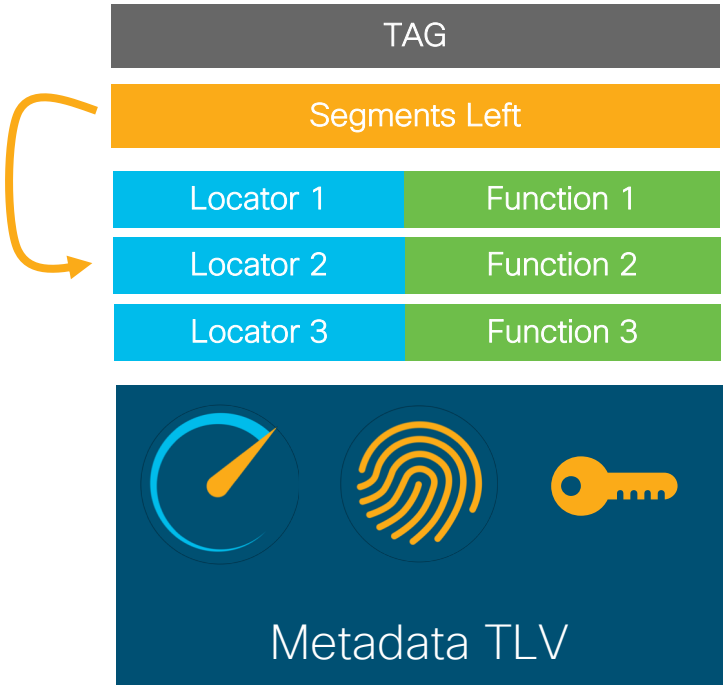


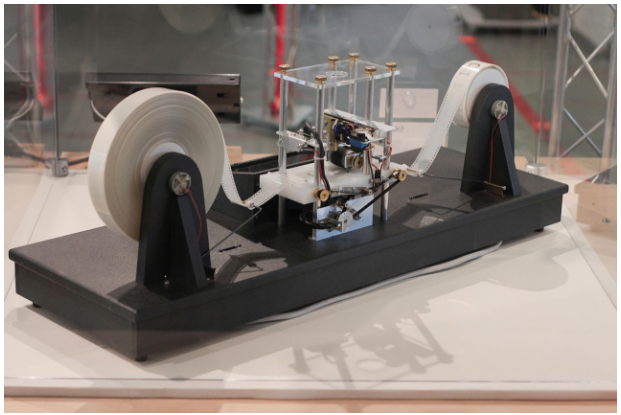

Turing

# SRv6 for anything


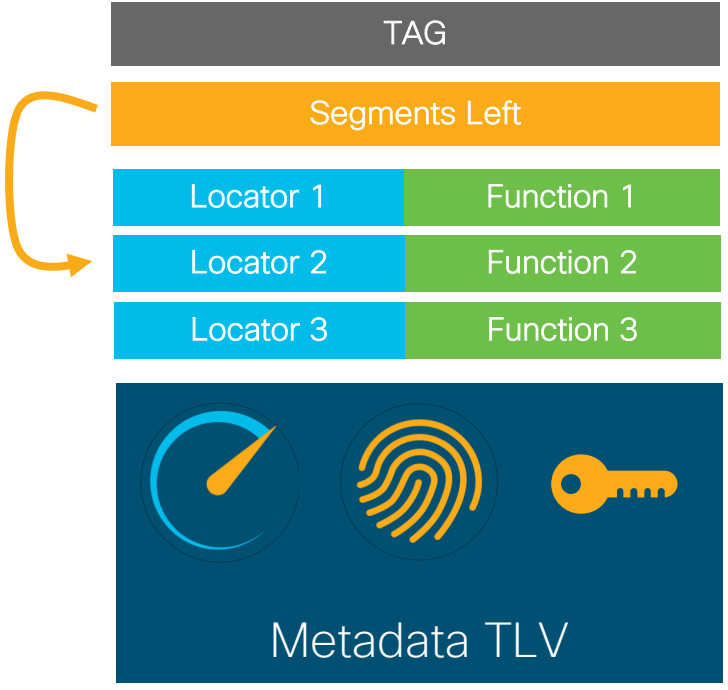


Optimized for HW processing
e.g. Underlay & Tenant use-cases

Optimized for SW processing
e.g. NFV, Container, Micro-Service

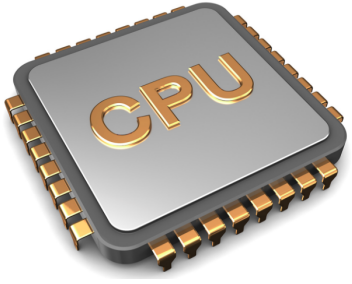

# Lead Operators

- Standardization
- Multi-Vendor Consensus


SPRING
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2017

C. Fisfils
Cisco Systems, Inc.
J. Leddy
Comcast
D. Voyer
D. Bernier
Bell Canada
D. Steinberg
Steinberg Consulting
R. Raszuk
Bloomberg LP
S. Matsushima
SoftBank Telecom
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
Dell Force10 Networks
A. Bashandy
K. Raza
D. Dukes
F. Clad
P. Camarillo, Ed.
Cisco Systems, Inc.
March 9, 2017

SRv6 Network Programming
draft-filsfils-spring-srv6-network-programming-00

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: September 13, 2017

G. Dawra, Ed.
C. Filsfils
D. Dukes
P. Brissette
P. Camarilo
Cisco Systems
J. Leddy
Comcast
D. Voyer
D. Bernier
Bell Canada
D. Steinberg
Steinberg Consulting
R. Raszuk
Bloomberg LP
B. Decraene
Orange
S. Matsushima
SoftBank Telecom Japan

March 12, 2017

**BGP Signaling of IPv6-Segment-Routing-based VPN Networks**
**draft-dawra-bgp-srv6-vpn-00.txt**

Network Working Group
Internet Draft
Intended status: Standard Track
Expires: September 2017

A. Bashandy, Ed.
C. Filsfils
L. Ginsberg
Cisco Systems
Bruno Decraene
Orange

March 10, 2017

IS-IS Extensions to Support Segment Routing over IPv6 Dataplane
draft-bashandy-isis-srv6-extensions-00

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

S. Previdi, Ed.
C. Filsfils
K. Raza
D. Dukes
Cisco Systems, Inc.
J. Leddy
B. Field
Comcast
D. Voyer
D. Bernier
Bell Canada
S. Matsushima
Softbank
I. Leung
Rogers Communications
J. Linkova
Google
E. Aries
Facebook
T. Kosugi
NTT
E. Vyncke
Cisco Systems, Inc.
D. Lebrun
Universite Catholique de Louvain
D. Steinberg
Steinberg Consulting
R. Raszuk
Bloomberg

March 13, 2017

IPv6 Segment Routing Header (SRH)
draft-ietf-6man-segment-routing-header-06

# SRv6 for Next-generation Mobile

SPRING and DMM
Internet-Draft
Intended status: Standards Track
Expires: January 18, 2018

S. Matsushima
SoftBank
C. Filsfils
Cisco Systems, Inc.
July 17, 2017

SRv6 for Mobile User-Plane
draft-matsushima-spring-dmm-srv6-mobile-uplane-01

# Use-Cases

# SID allocation for illustration purpose

- For simplicity: B:K:F::
  - SID Block B
  - Node K
  - Function F

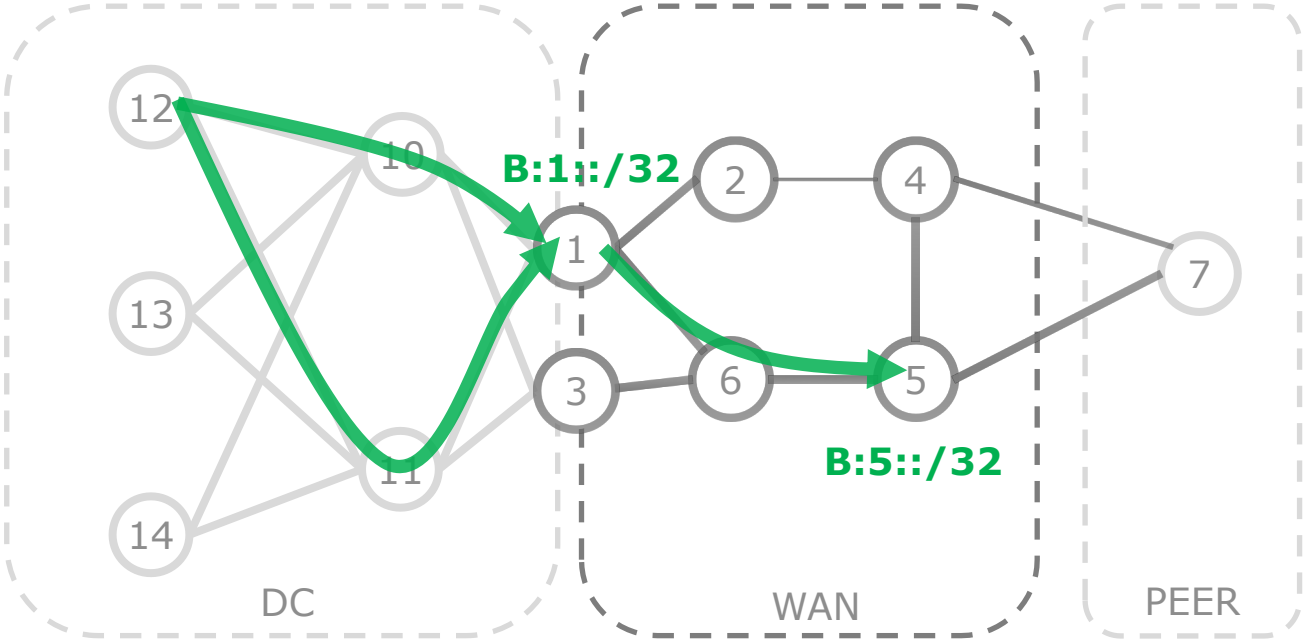

# Endpoint

- For simplicity
- Function 1 denotes the most basic function
- Shortest-path to the Node

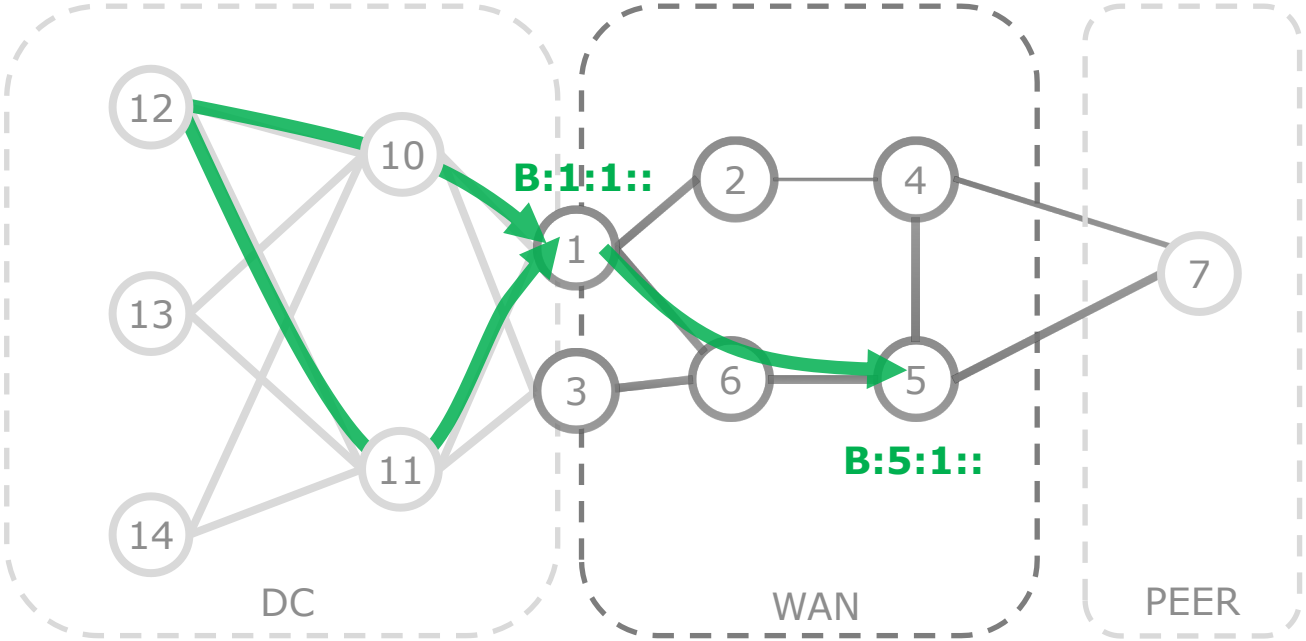

# B1::1 and then B5::1

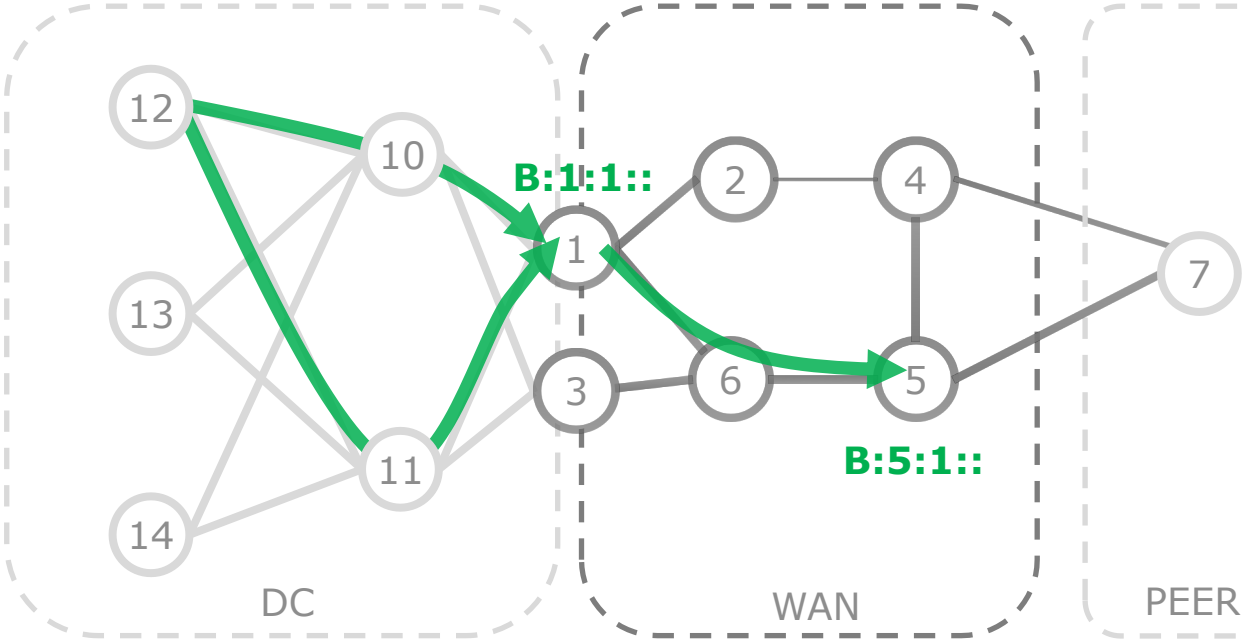

# Endpoint then xconnect to neighbor

- For simplicity
- BK::CJ denotes

Shortest-path to the Node K and then x-connect (function C) to the neighbor J

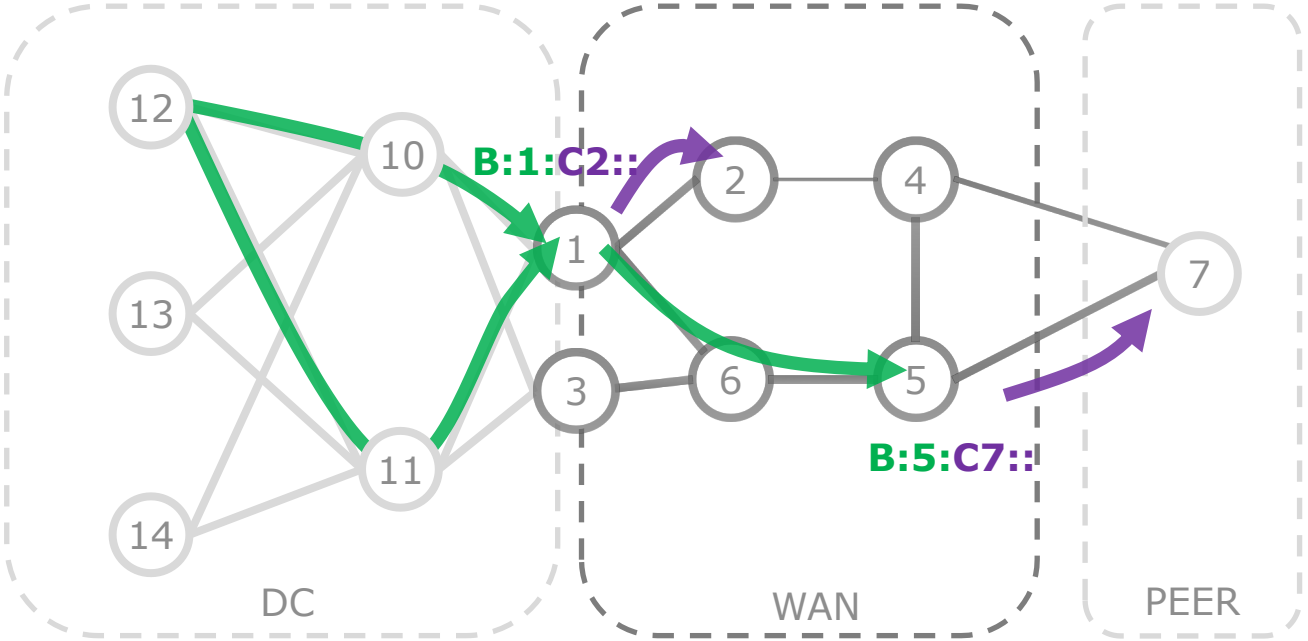

# B1::1 and then B5::C7

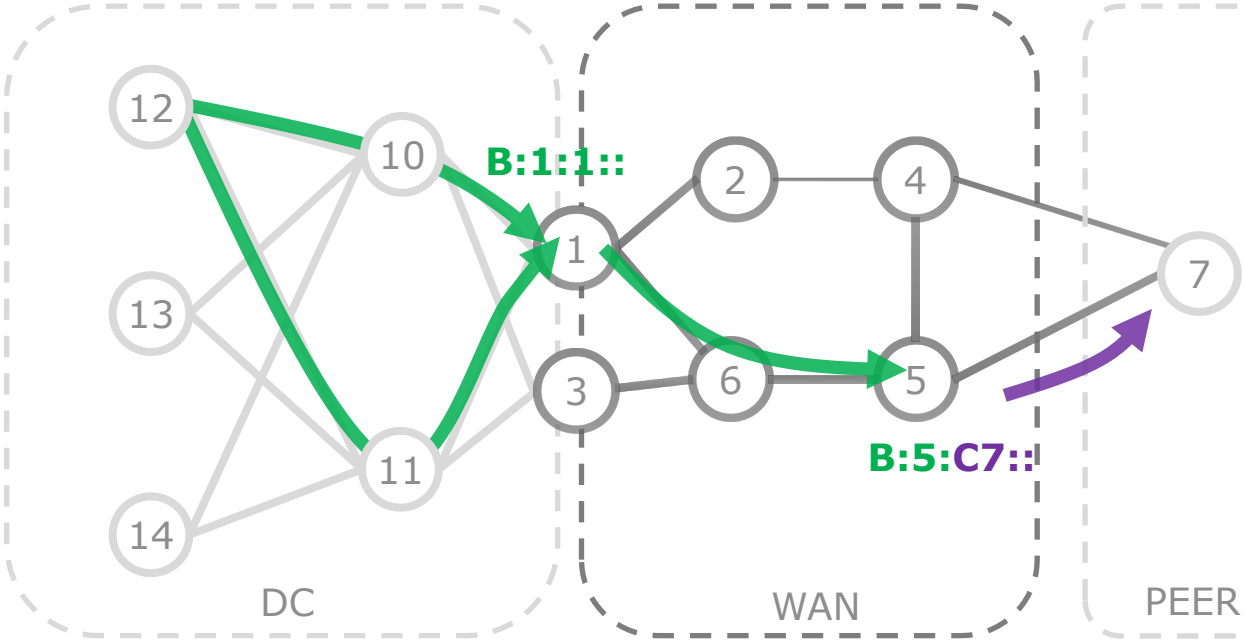

# TILFA

- 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, planned to carry the traffic
  - avoid any intermediate flap via alternate path
- Incremental deployment
- Distributed and Automated Intelligence

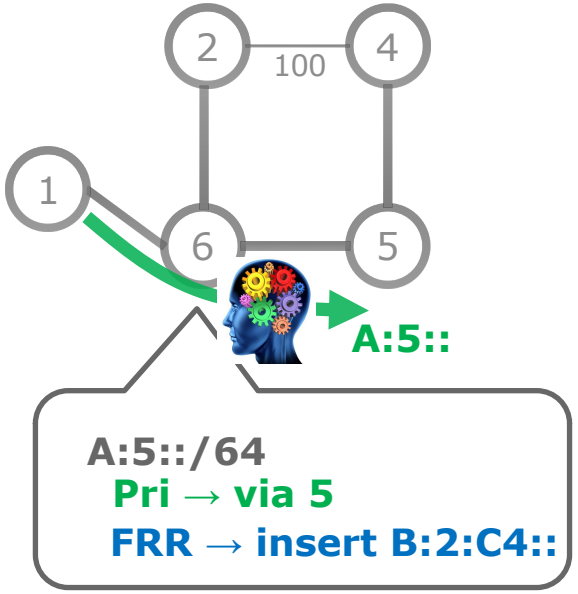

# TILFA

- 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, planned to carry the traffic
  - avoid any intermediate flap via alternate path
- Incremental deployment
- Distributed and Automated Intelligence

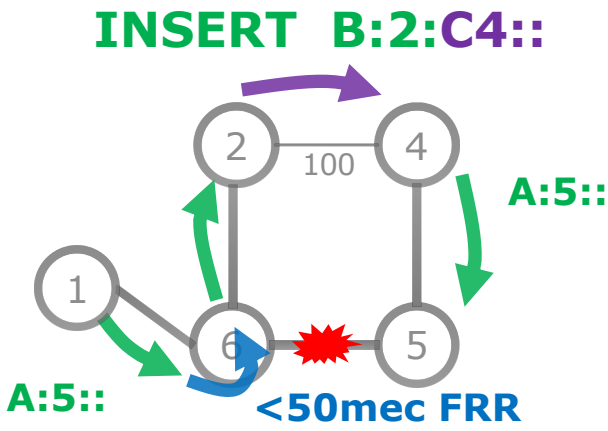

# Overlay

- Simple
  - Protocol elimination
- Automated
  - No tunnel to configure
- Efficient
  - SRv6 for everything
  - Reuse BGP/VPN signaling

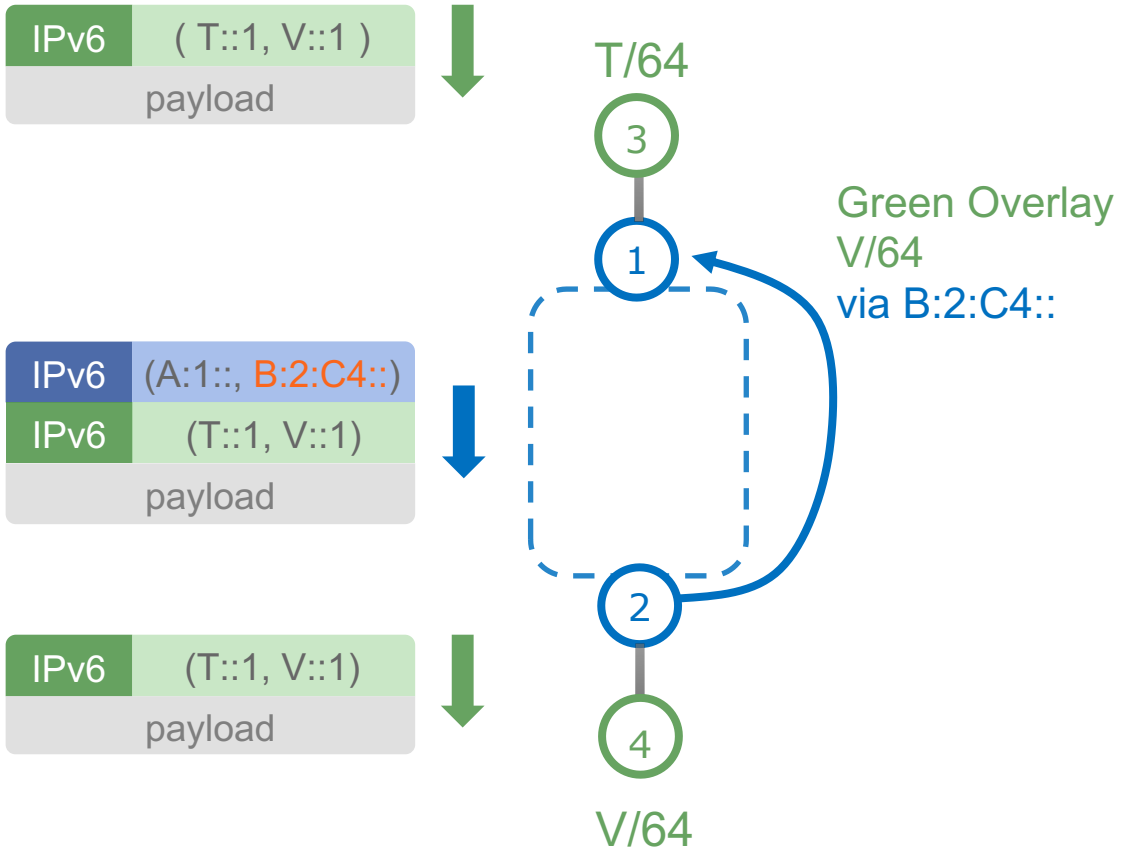

# Overlay with Underlay SLA

- SRv6 does not only eliminate unneeded overlay protocols
- SRv6 solves problems that these protocols cannot solve
- Also support IPv4 and Ethernet VPN's

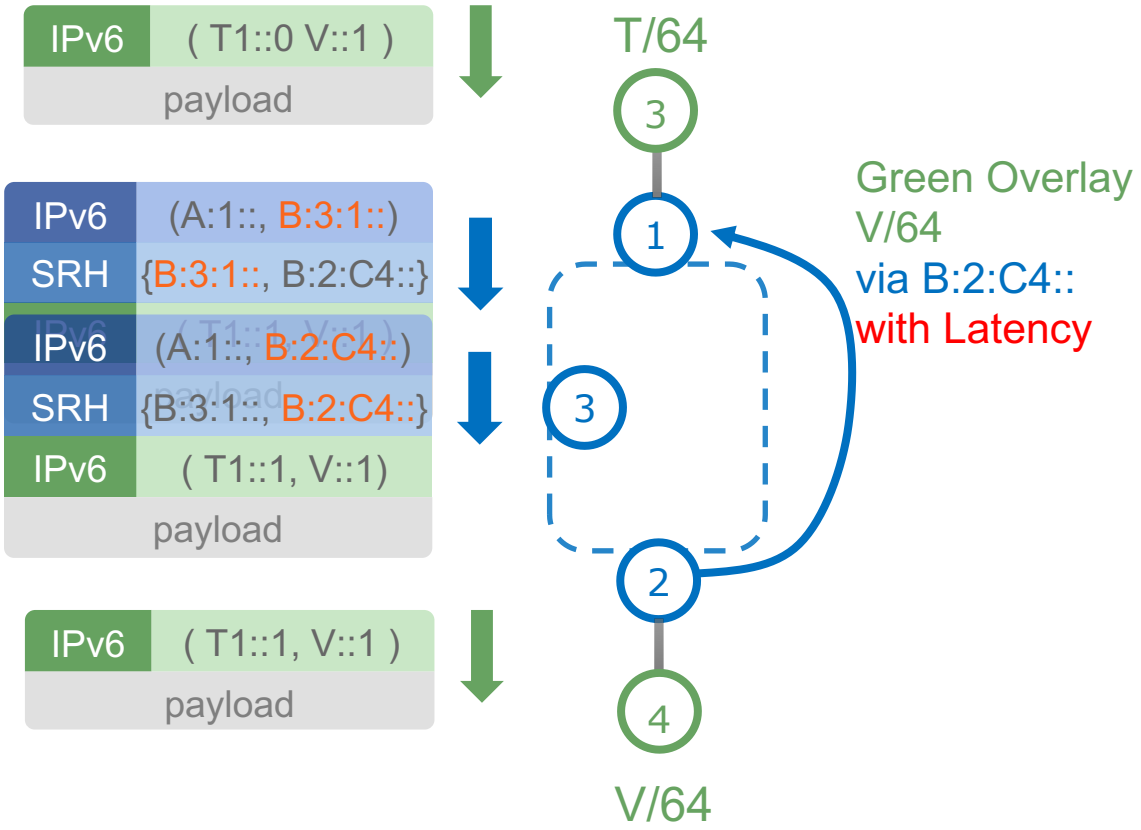

# Integrated NFV

- Stateless Service Chaining
  - NSH creates per-chain state in the fabric
  - SR does not
- App is SR aware or not
- App can work on IPv6 or IPv4 inner packets

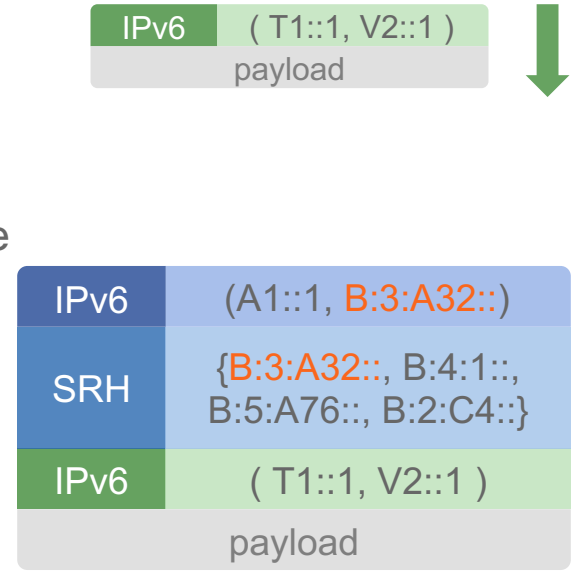


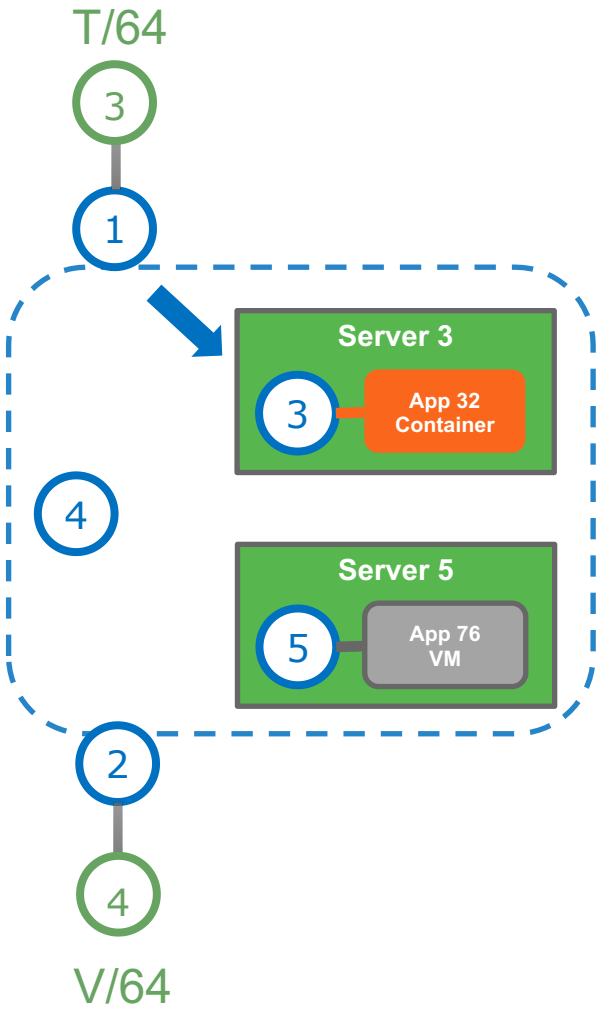

# Integrated NFV

- Integrated SLA

| | |
|---|---|
| IPv6 | (A1::1, B:4:1::) |
| SRH | {B:3:A32::, B:4:1::, B:5:A76::, B:2:C4::} |
| IPv6 | (T1::1, V2::1) |
| payload | |

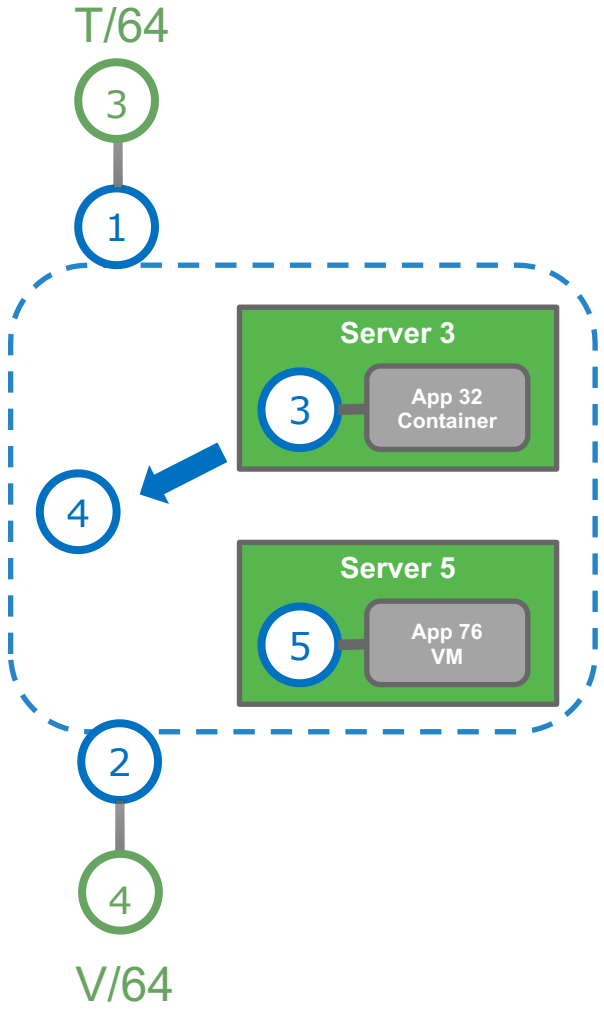

# Integrated NFV

- Stateless Service Chaining
  - NSH creates per-chain state in the fabric
  - SR does not
- App is SR aware or not
- App can work on IPv6 or IPv4 inner packets

| | |
|---|---|
| IPv6 | (A1::1, B:5:A76::) |
| SRH | {B:3:A32::, B:4:1::, B:5:A76::, B:2:C4::} |
| IPv6 | ( T1::1, V2::1) |
| payload | |

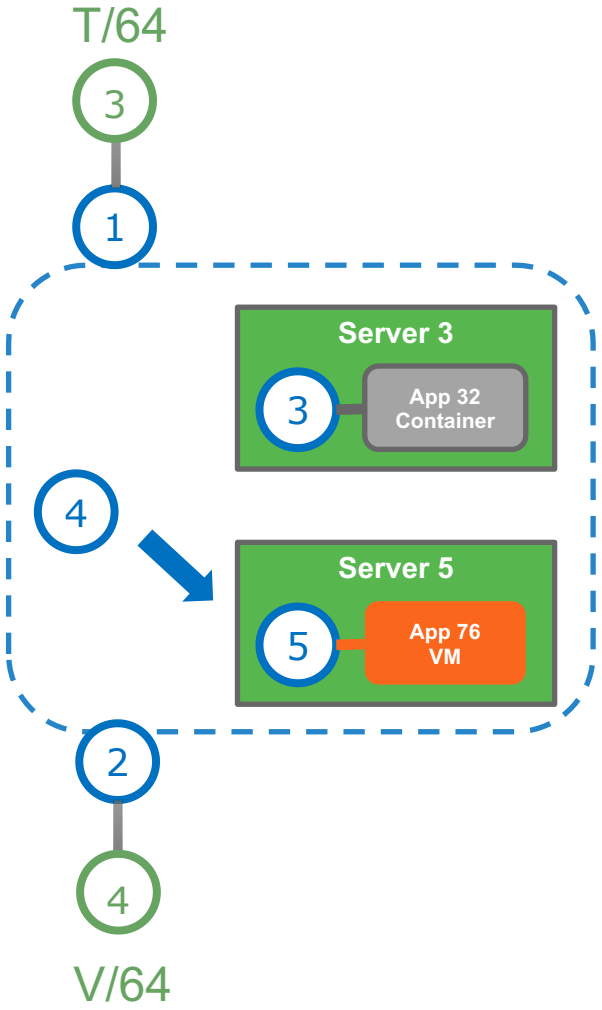

# Integrated NFV

- Integrated with Overlay

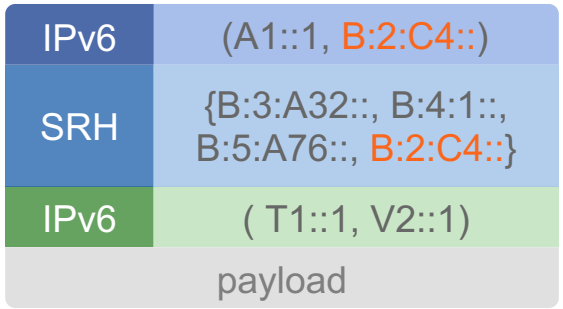



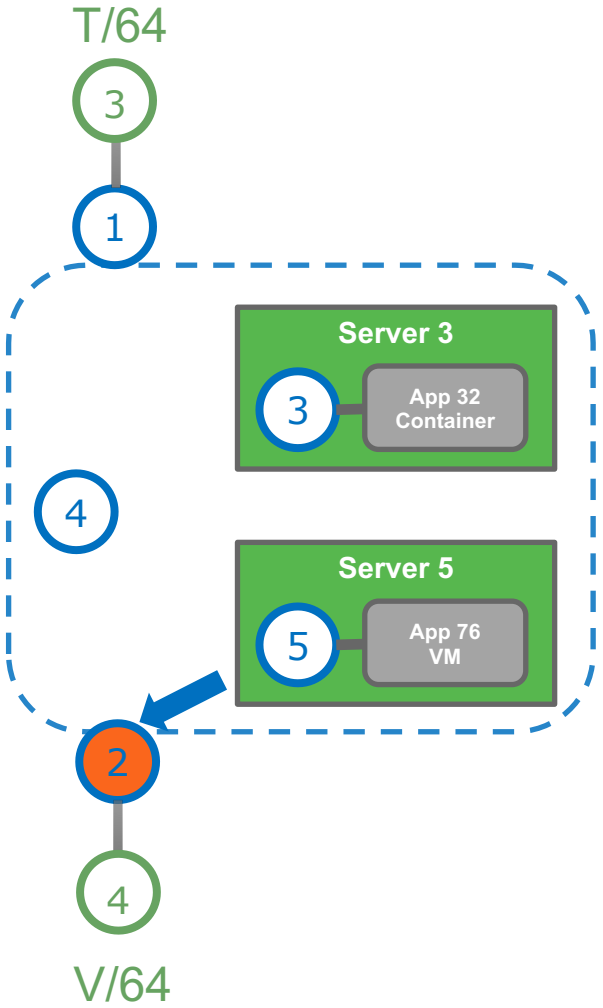

# Simplified and Scaled mobility

- Underlay: Traffic Engineering / End-to-end Network slicing
- Overlay: Efficient protocol replacement to GTP-U
- GiLAN: Scalable and flexible Service Programming

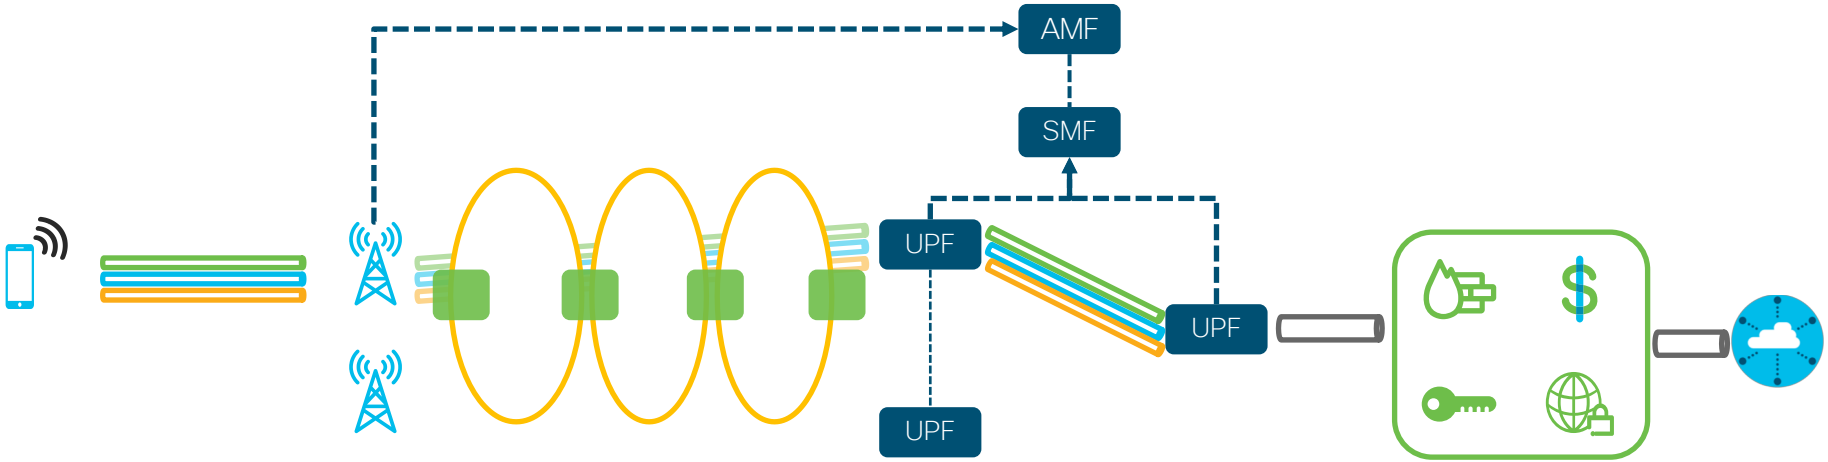

# More use-cases

- 6CN: enhancing IP to search for Content
- 6LB: enhancing load-balancers
  - Better flow stickiness and load distribution
- Video Pipeline

# Implementations

- Cisco HW
  - NCS5k - XR
  - ASR9k - XR
  - ASR1k - XE
- Open-Source
  - Linux 4.10
  - FD.IO
- Barefoot HW



# SRv6 Interoperability Demonstration

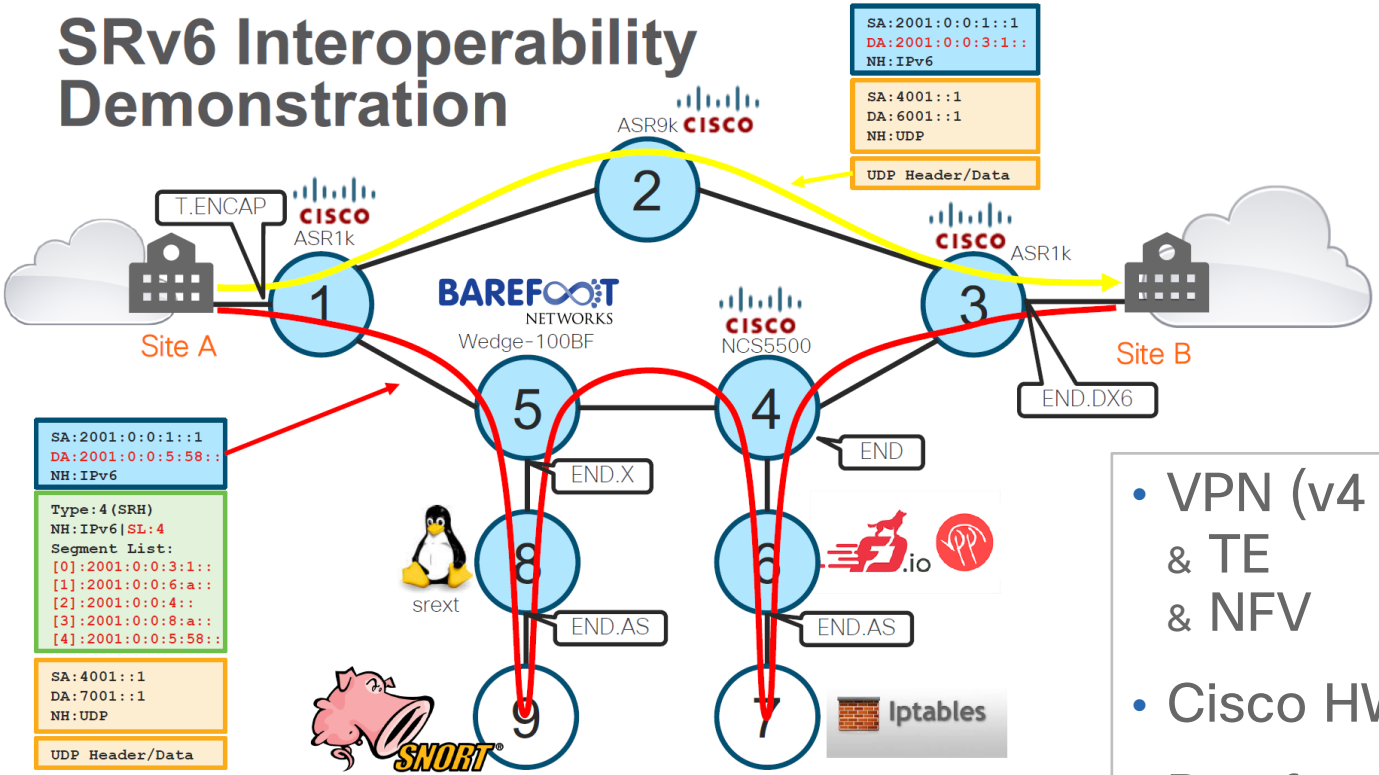


August 2017

- VPN (v4 and v6) & TE & NFV
- Cisco HW with XR and XE
- Barefoot HW with P4 code
- FD.IO
- Linux

blogs.cisco.com/sp/segment-routing-ipv6-interoperability-demo-is-already-there

# Academic Contribution

# The Segment Routing Architecture

C. Filsfils, N. K. Nainar, C. Pignataro, J. C. Cardona and P. Francois; **GLOBECOM 2015**

- Introduces the base building blocks and use-cases
- MPLS and IPv6 data planes
- Distributed and centralized path calculation
- Traffic engineering, service chaining and FRR use-cases

# Efficient label encoding in segment-routing enabled optical networks

F. Lazzeri, G. Bruno, J. Nijhof, A. Giorgetti and P. Castoldi; ONDM 2015

- SR-native path calculation algorithm
- Node SIDs represented as virtual links in the topology
- Adjacency SIDs mapped to physical links
- Provide ECMP-enabled SR paths with least number of SIDs

# Traffic duplication through segmentable disjoint paths

F. Aubry, D. Lebrun, Y. Deville and O. Bonaventure; IFIP Networking 2015

- Traffic replication and steering over disjoint SID lists
- Leverage SRv6 implementation in Linux kernel
- SR-native disjoint path calculation heuristic

# A Declarative and Expressive Approach to Control Forwarding Paths in Carrier-Grade Networks

R. Hartert, S. Vissicchio, P. Schaus, O. Bonaventure, C. Filsfils, T. Telkamp and P. Francois; SIGCOMM 2015

- Constraint programming-based SRTE controller
- Translation of high-level network intents into optimization objective and contraints
- Fast calculation of ECMP-enabled SR paths
- Dynamic reaction to network events

# A Linux kernel implementation of Segment Routing with IPv6

David Lebrun; **INFOCOM Workshops 2016**

- Native SRv6 dataplane support in Linux kernel
- SRv6 endpoint behaviors: SID functions
- SRv6 transit behaviors: traffic steering in SR policies
- Compliant with IETF specification

# SERA: SEgment Routing Aware Firewall for Service Function Chaining scenarios

A. Abdelsalam, S. Salsano, F. Clad, P. Camarillo and C. Filsfils; IFIP Networking 2018

- Process SRv6 traffic with existing iptables configurations
- SR-specific filtering rules (e.g., argument, segment list, tag)
- SR-specific actions (e.g., skip next segment)

# Software Resolved Networks: Rethinking Enterprise Networks with IPv6 Segment Routing

D. Lebrun, M. Jadin, F. Clad, C. Filsfils and O. Bonaventure; SOSR 2018

- End host may include path requirements in DNS queries
- DNS server communicates with controller to provision an suitable SR policy (match path req. and enterprise policy)
- DNS response contains a Binding SID to this SRv6 policy and a pre-computed HMAC value

# 6LB: Scalable and Application-Aware Load Balancing with Segment Routing

Y. Desmouceaux, P. Pfister, J. Tollet, M. Townsley and T. Clausen; IEEE/ACM Transactions on Networking 2018

- Leverage SRv6 for high performance load balancing
- Multiple load placing options encoded in a segment list
- In-band signaling of flow-server mapping

# SR+SFC workshop at CNSM 2018

- Flexible failure detection and fast reroute using eBPF and SRv6
  - SRv6 TI-LFA implementation using eBPF
  - BFD-like mechanism for fast failure detection with SRv6
- Performance of IPv6 Segment Routing in Linux Kernel
  - Performance evaluation framework
  - Compare SRv6 performances across implementations / platforms

# SR+SFC workshop at CNSM 2018

- Zero-Loss Virtual Machine Migration with IPv6 Segment Routing
  - Leverage SRv6 network programming to preserve VM reachability during its migration
- Proposal and Investigation of a Scalable NFV Orchestrator Based on Segment Routing Data/Control Plane
  - Leverage SR to drastically reduce resource allocation and routing complexity in datacenter and cloud environments

# Conclusion

# Segment Routing

- Strong industry support
- Fantastic deployment rate
- Bold architecture: network programming
- Numerous use-cases
- Feel free to join the lead-operator team!

# Stay uptodate

amzn.com/B01I58LSUO

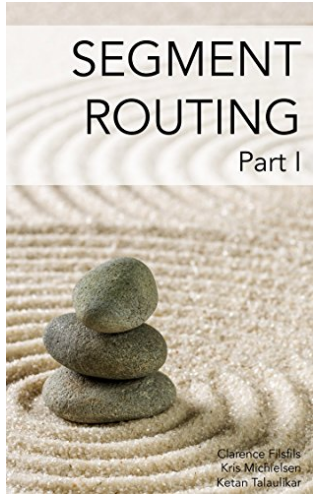


SR
Segment Routing

segment-routing.net

linkedin.com/groups/8266623

twitter.com/SegmentRouting

facebook.com/SegmentRouting/

Thank you