SRv6: Network as a Computer and Deployment use-cases

Gaurav Dawra, Technical Leader, Cisco Systems (gdawra@cisco.com)
John Brzozowski, Fellow and Chief Architect, Comcast
John Leddy, Network Engineering, Comcast
Clarence Filsfil, Fellow, Cisco Systems (cfilsfil@cisco.com)

NANOG, San Jose, Oct 2017

Agenda

1. SRv6 101
2. SRv6 LocalSID & Use-Cases
3. VPN Overlay
4. Service Chaining
5. VPP
6. SD-WAN
7. Spray
8. Mobile Network

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

Our commitment to **Lead Operators**

Strong customer adoption

WEB, SP, Enterprise

Orange

Standardization

IETF

De-facto SDN Architecture

Seamless deployment

Multi-vendor Consensus

Bloomberg

Bell COMCAST Proximus T-Mobile

Objectives of SRv6

IPv6 provides reachability

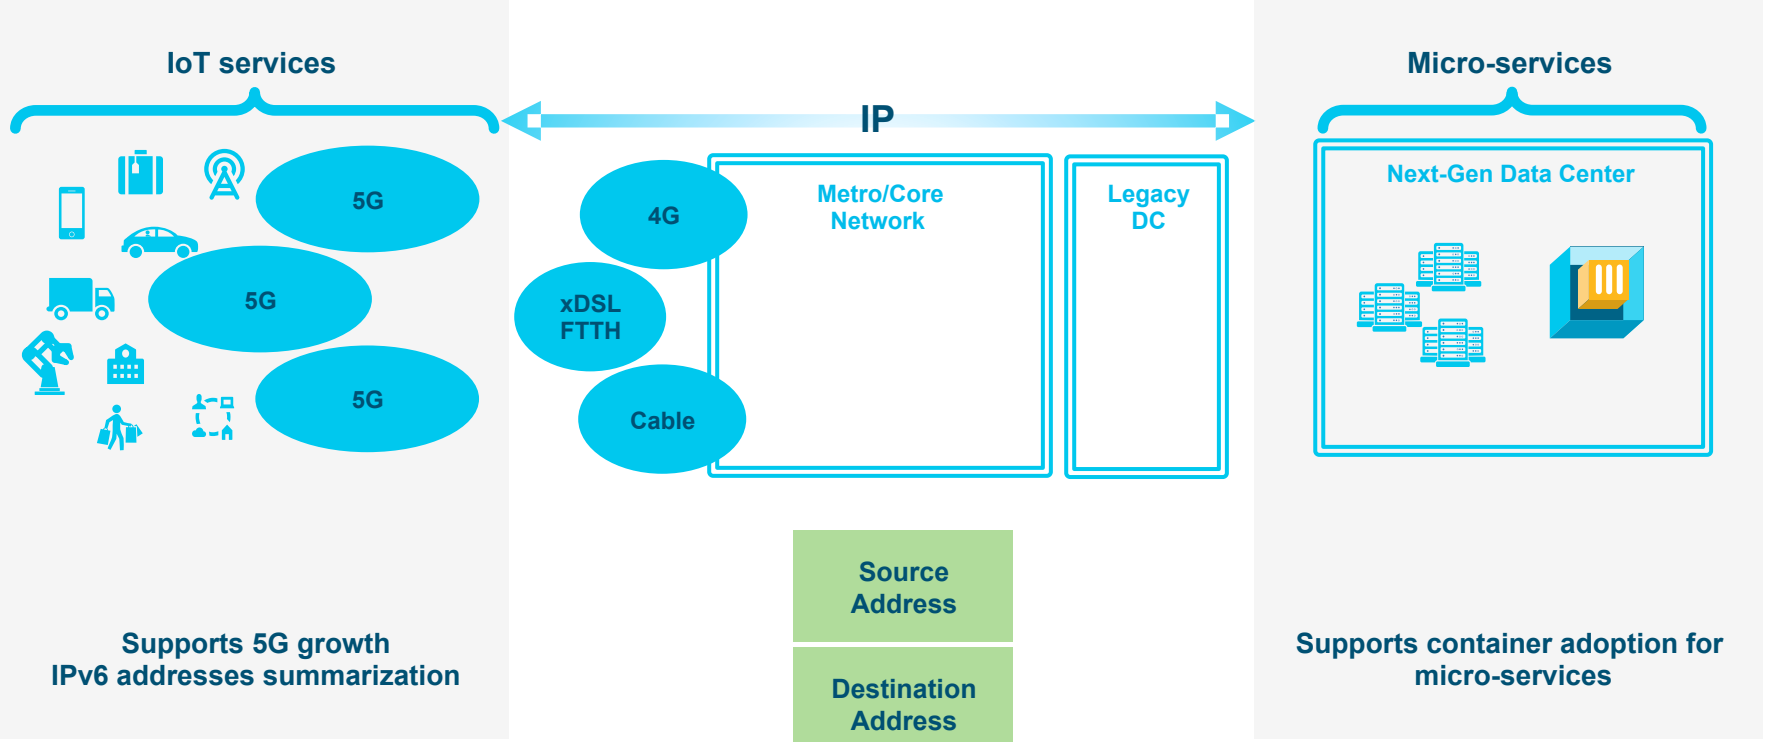

SRv6 for underlay

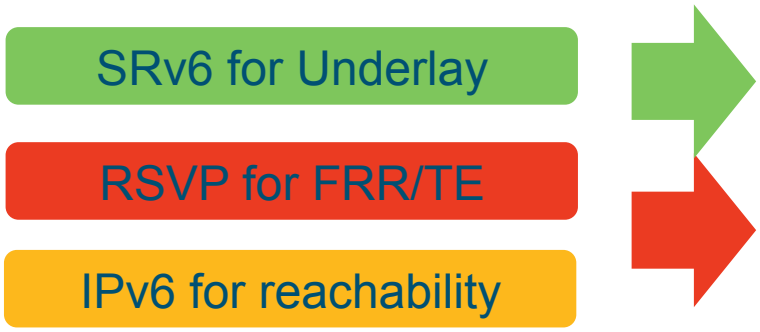

Simplification through protocol reduction

SLA through automated FRR and TE

De-facto SDN architecture

Scaling issue with k*N^2

SRv6 for underlay and overlay

Multiplicity of protocols and states hinder network economics

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 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

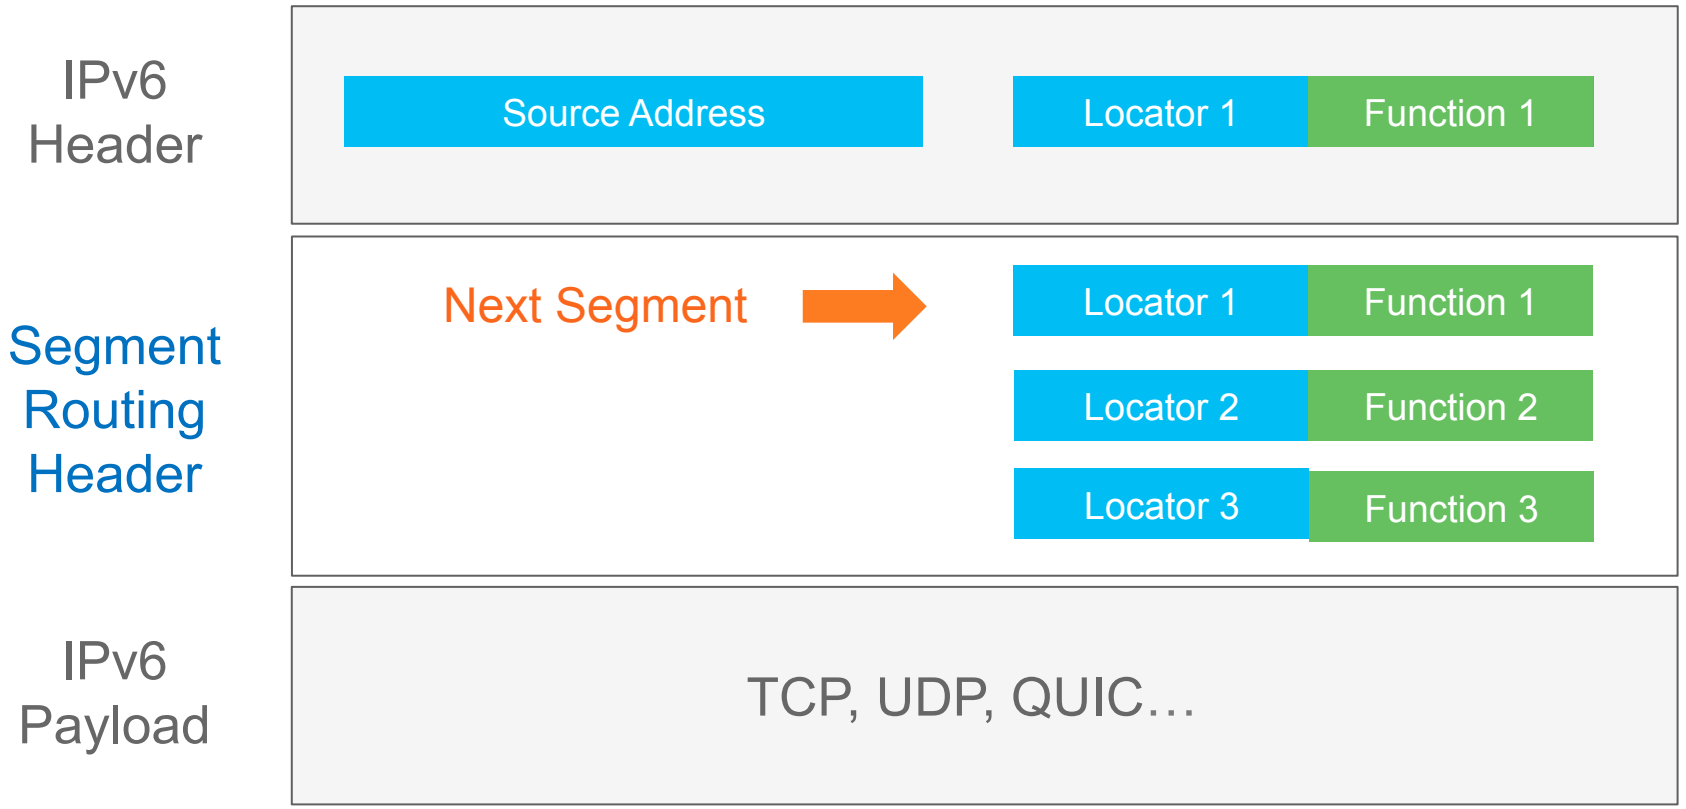

Network Program

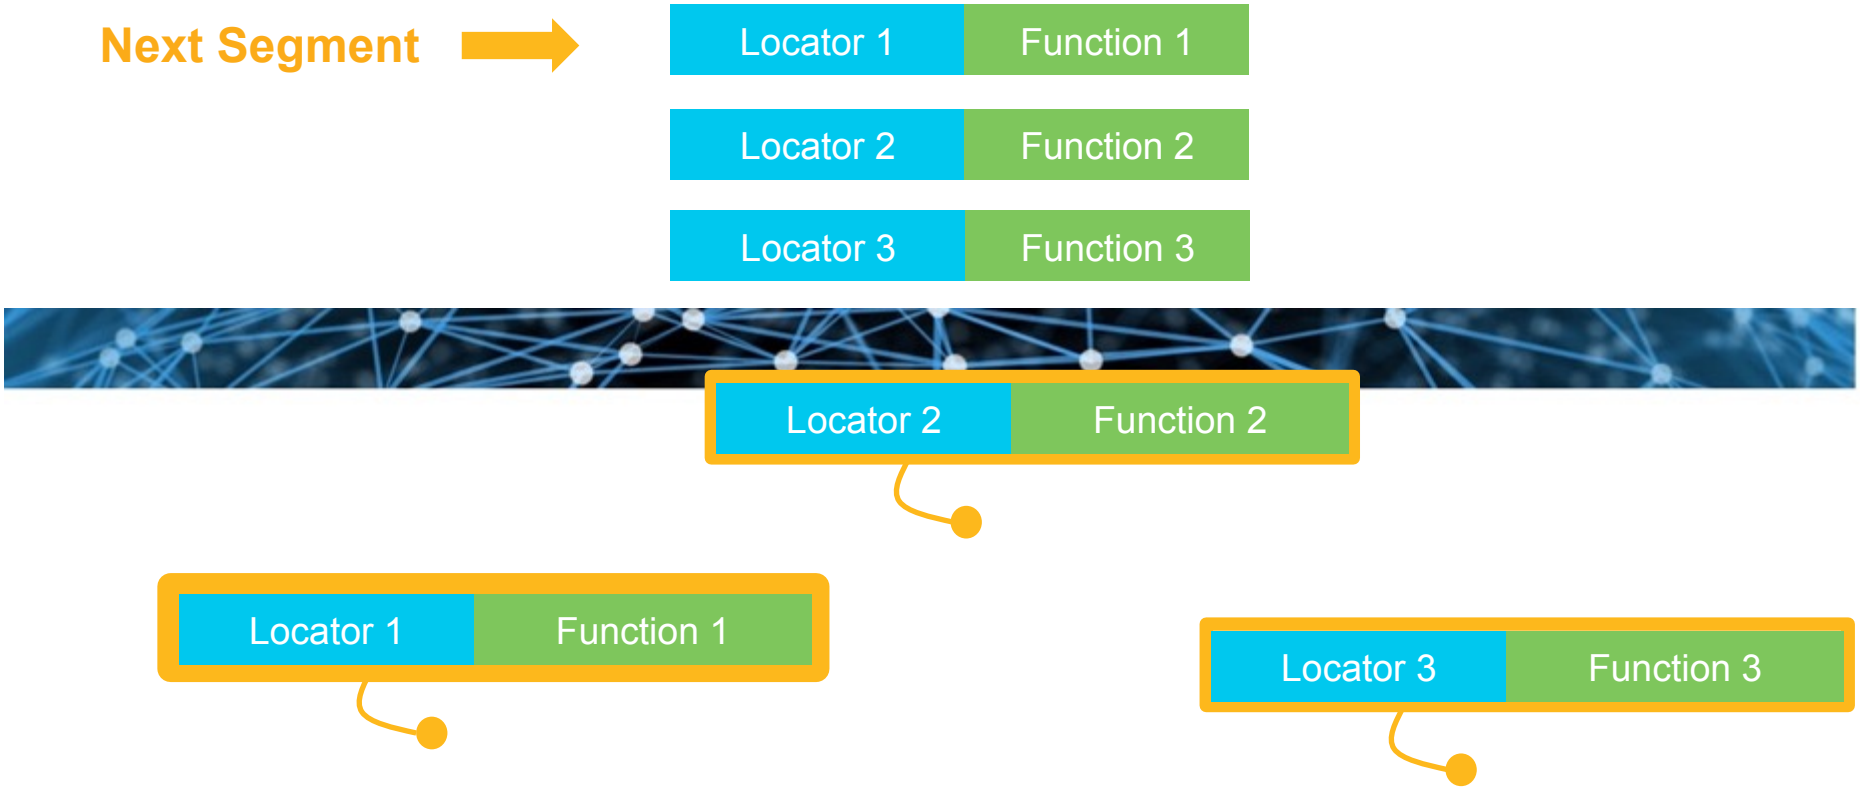

Network Program

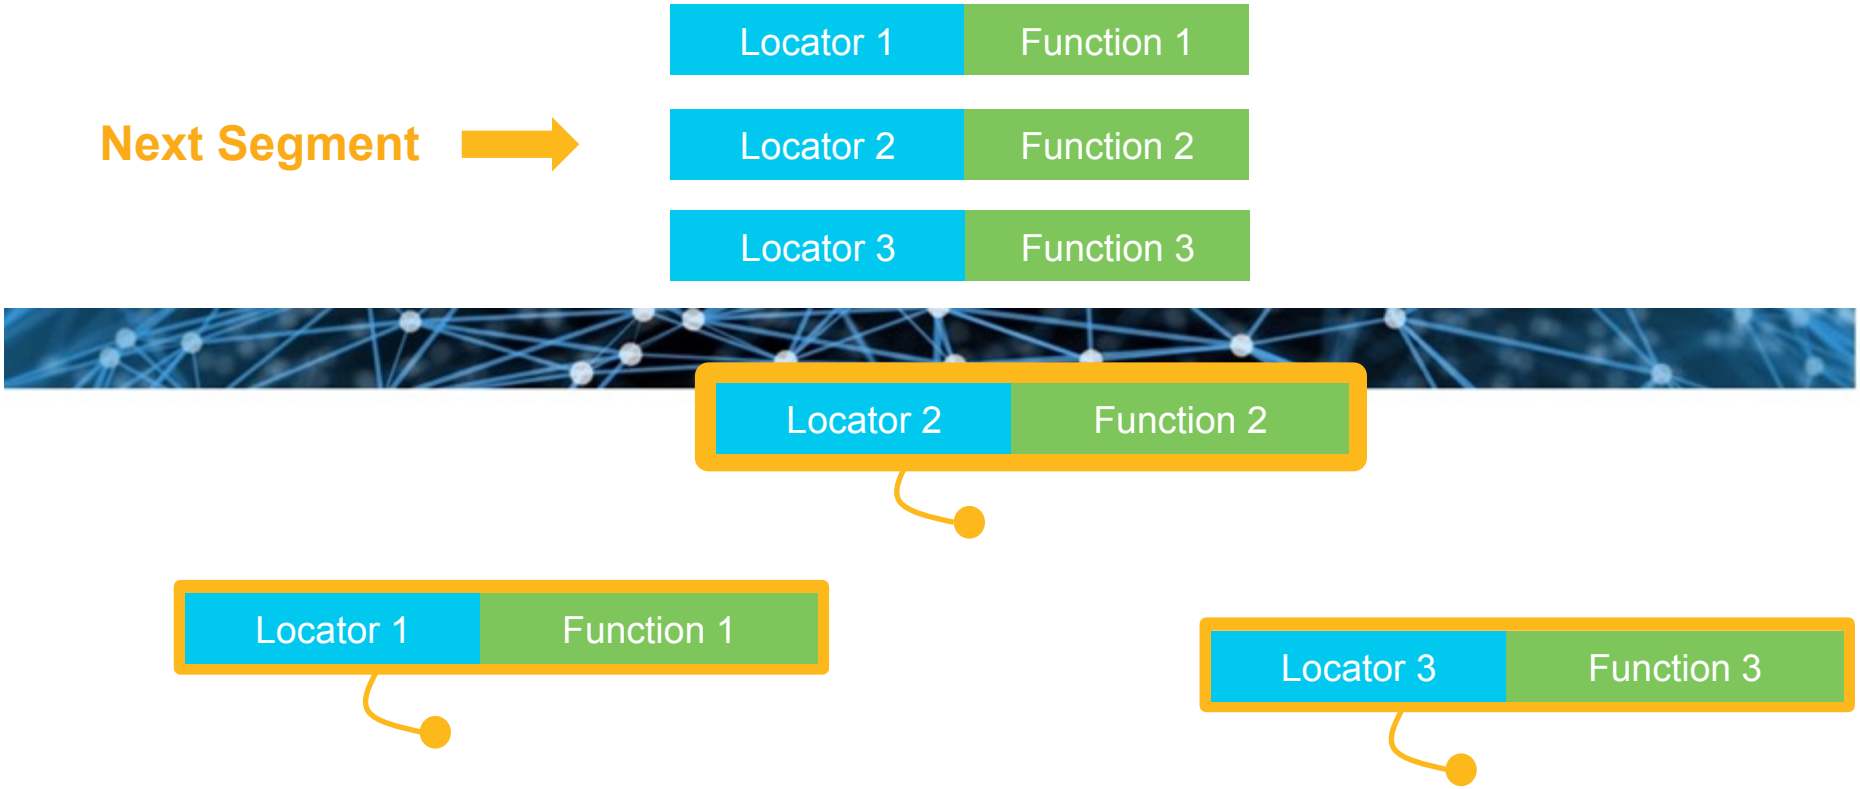

Network Program

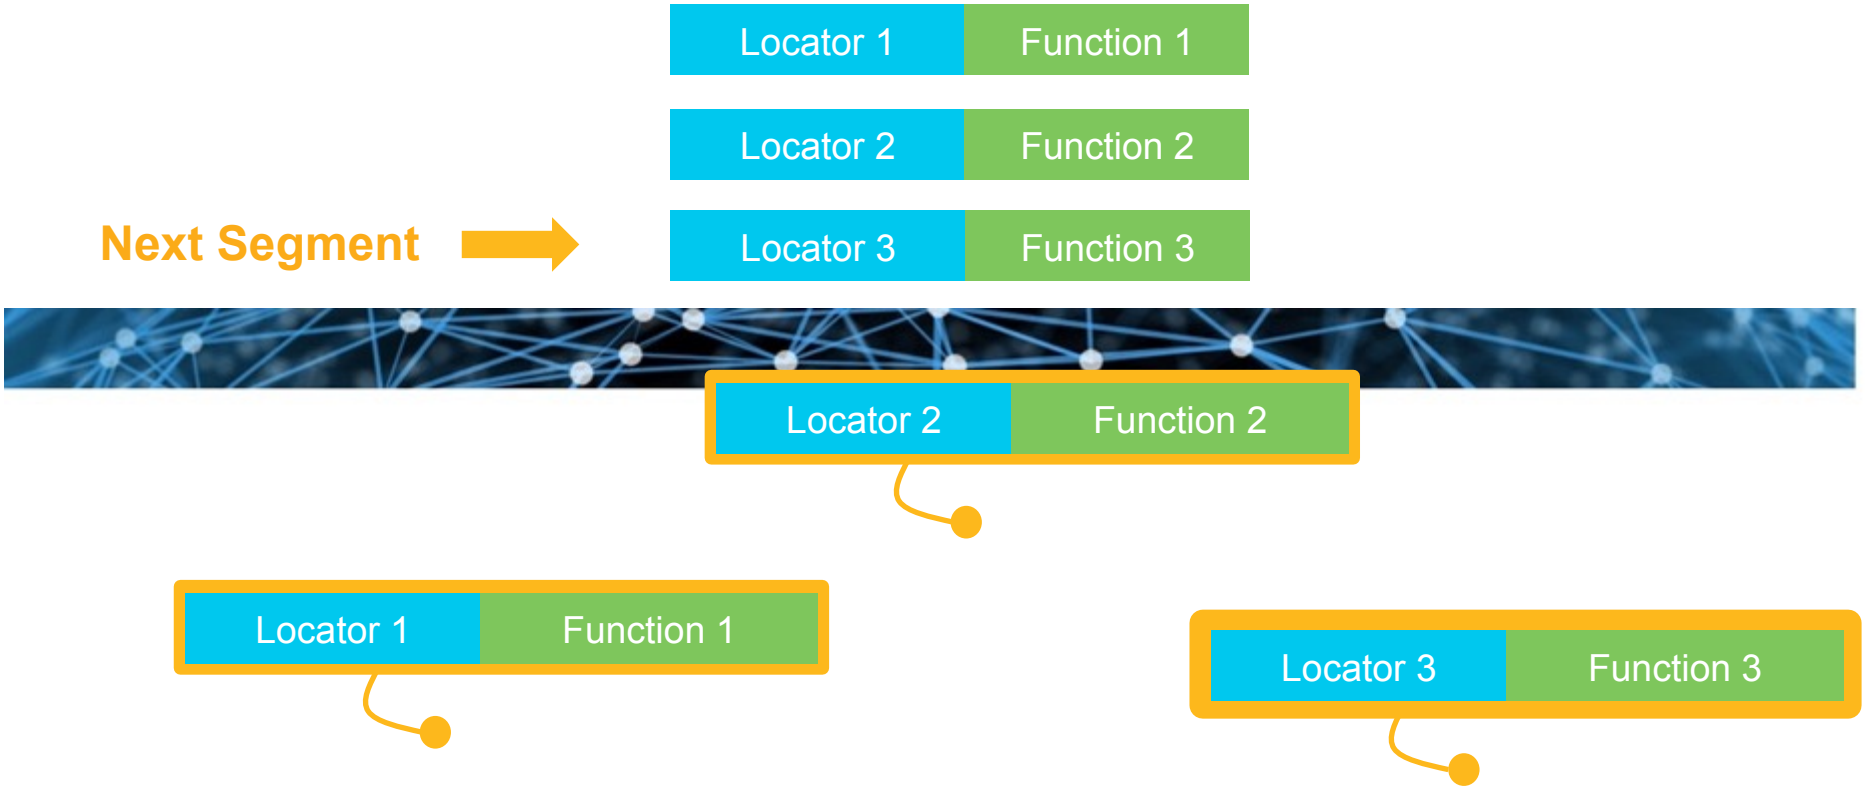

Argument shared between functions

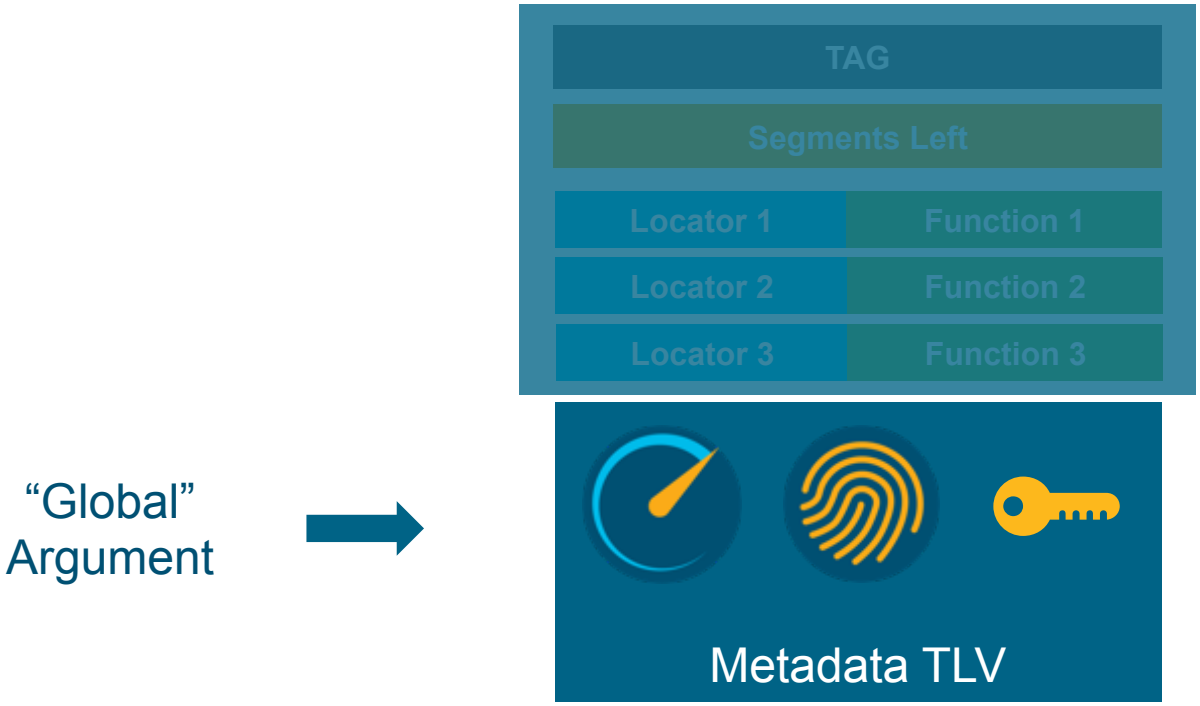

Group-Based Policy

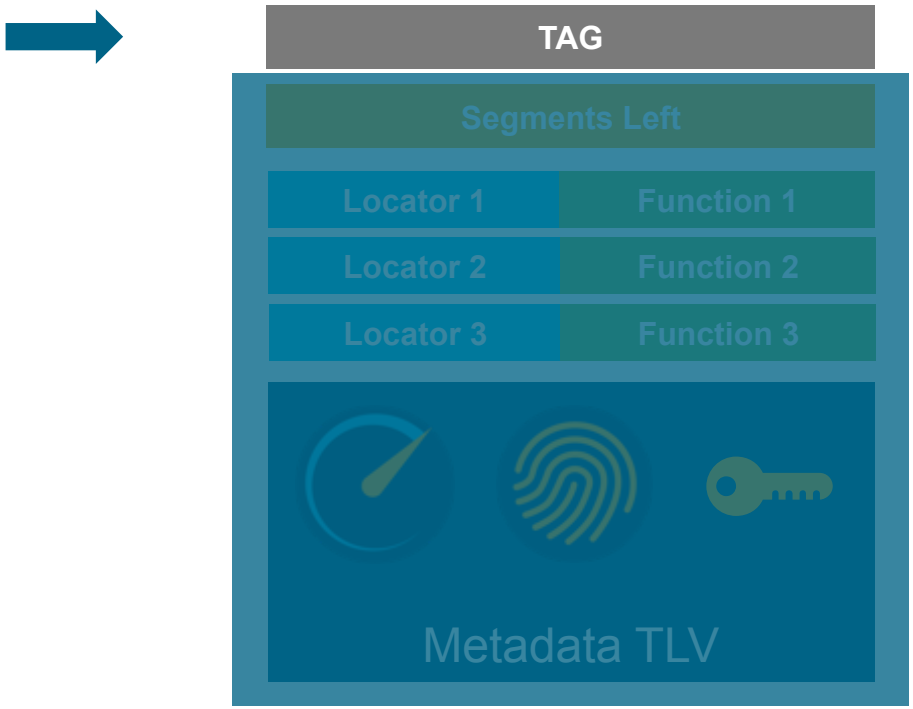

SR Header

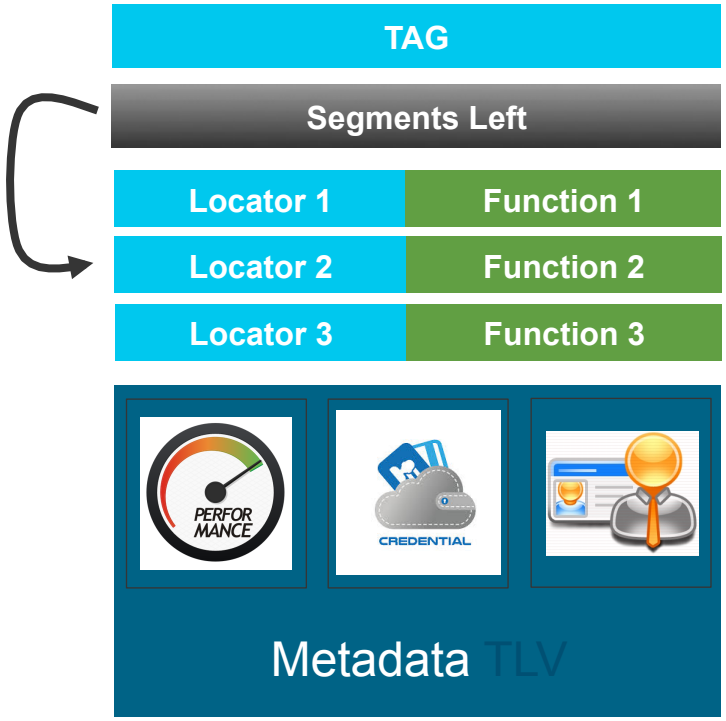


```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Next Header   |  Hdr Ext Len  | Routing Type  | Segments Left |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+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

SRv6 for anything

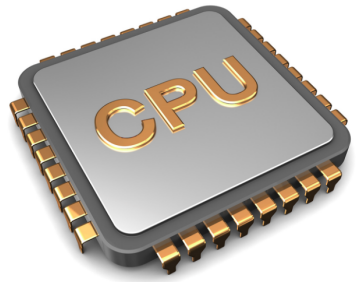

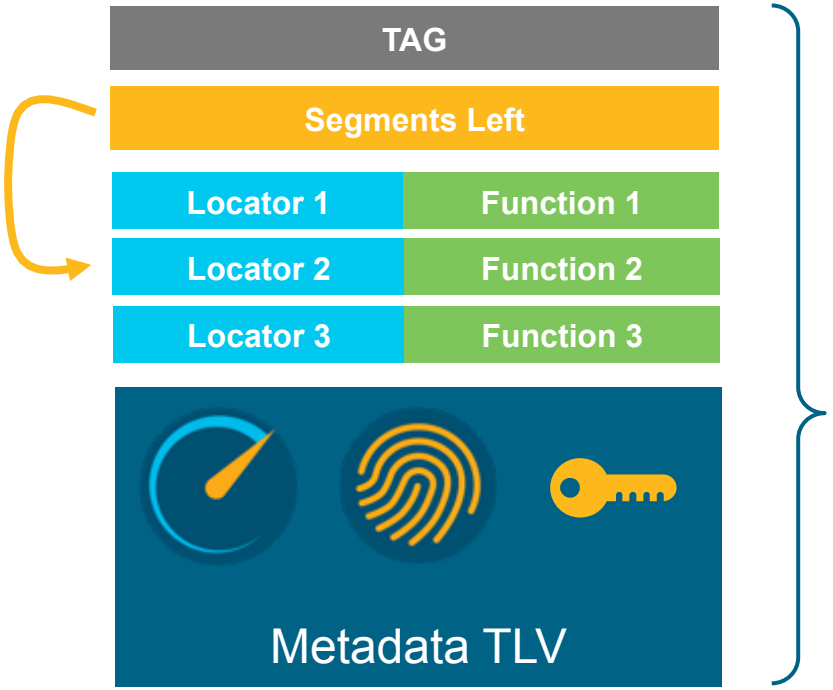

Optimized for HW processing
e.g. Underlay & Tenant use-cases

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

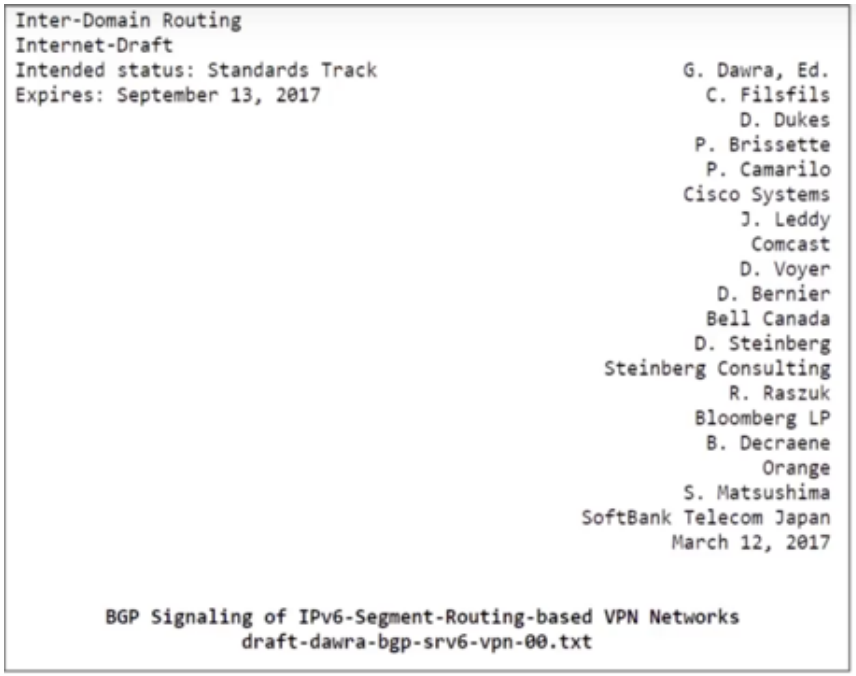

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: September 13, 2017

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

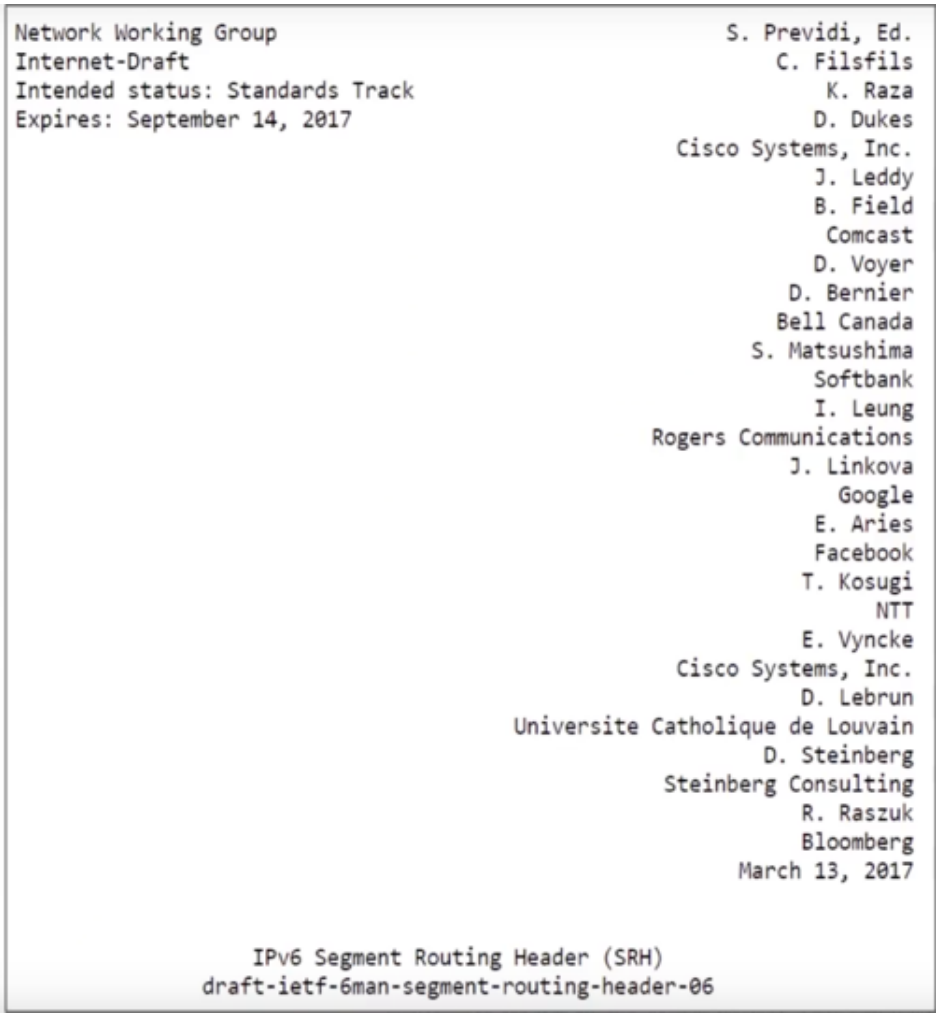

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

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

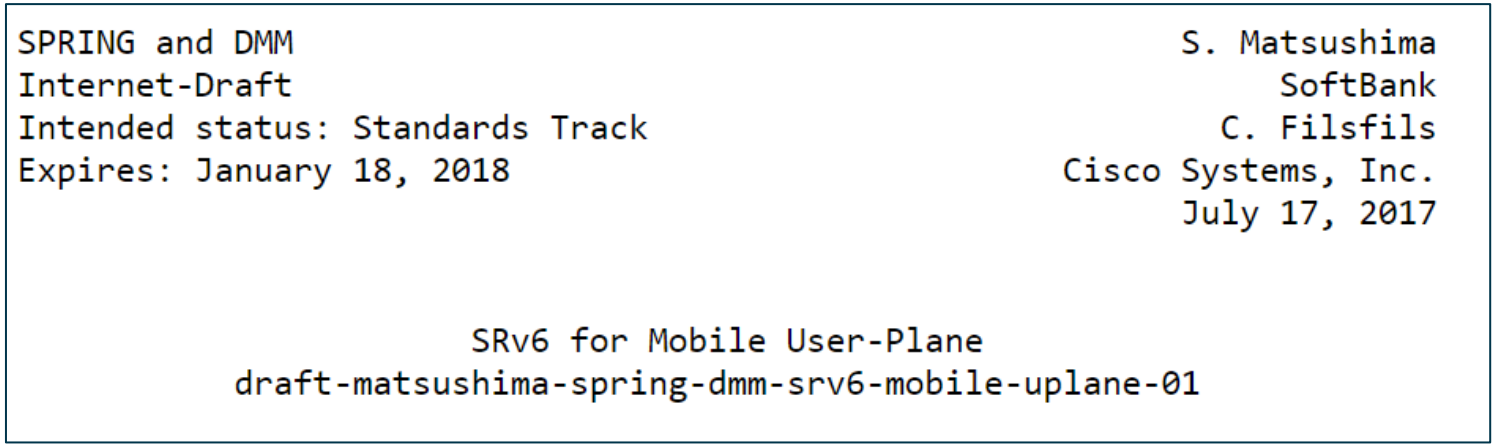

SPRING and DMM
Internet-Draft
Intended status: Standards Track
Expires: January 18, 2018

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
- Node K advertises prefix AK::/64
- The function is encoded in the last 64 bits

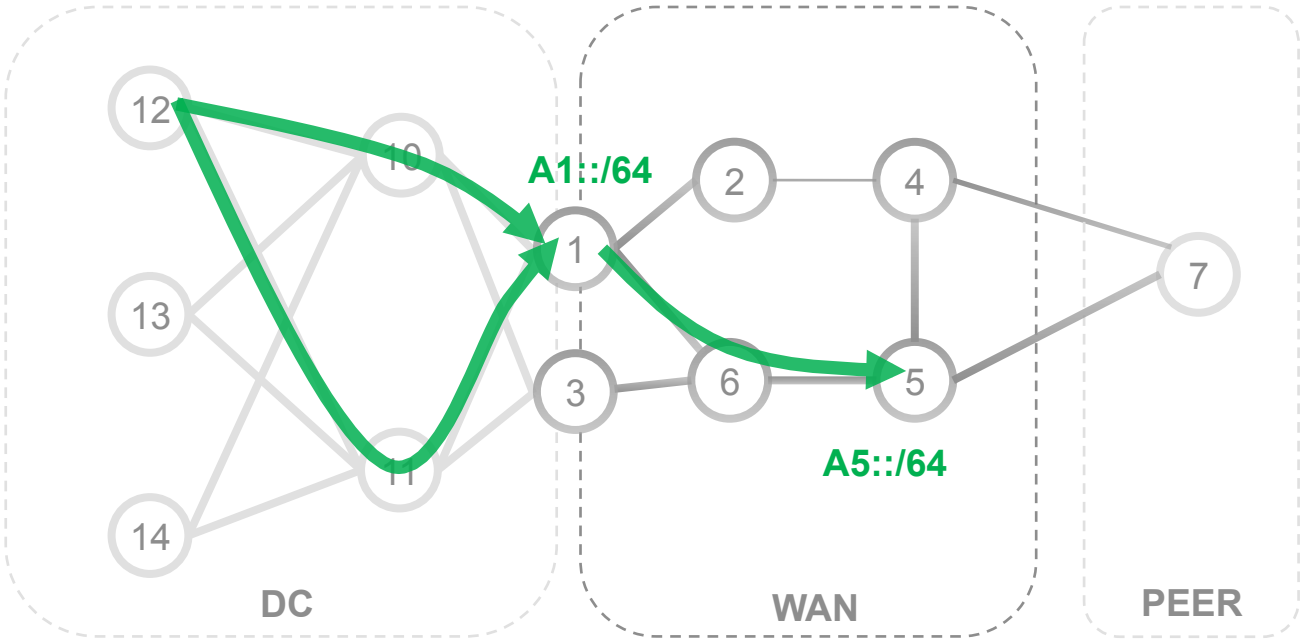

Endpoint

- For simplicity
- Function 0 denotes the most basic function
- Shortest-path to the Node

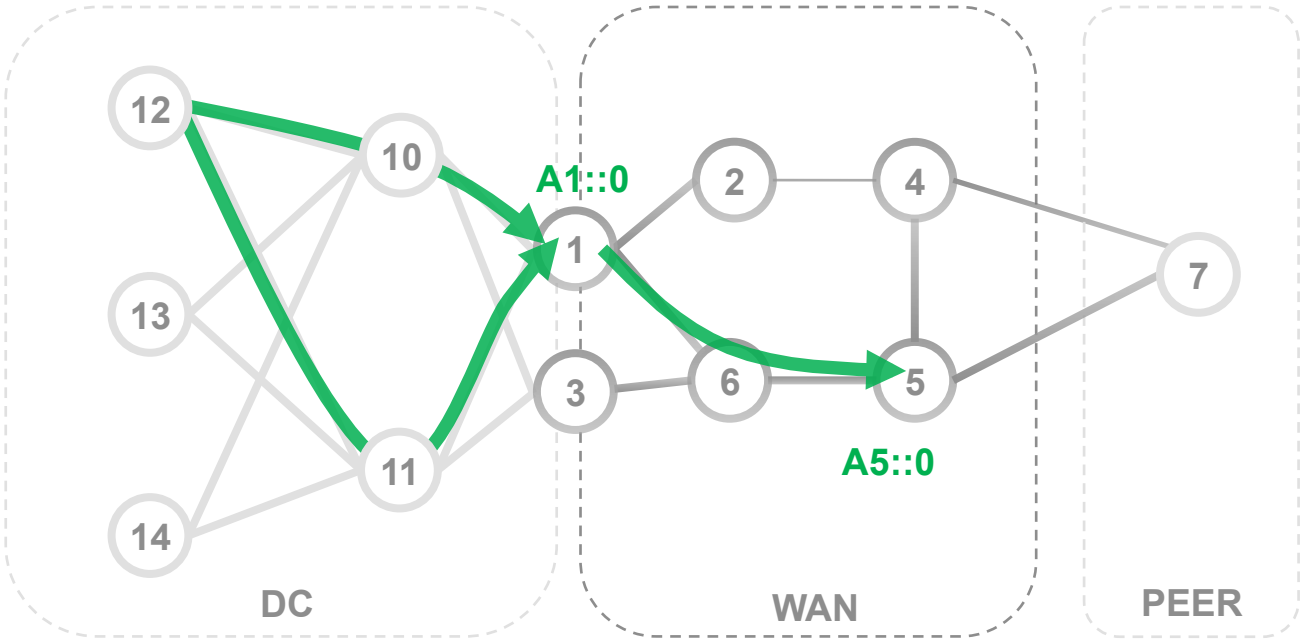

A1::0 and then A5::0

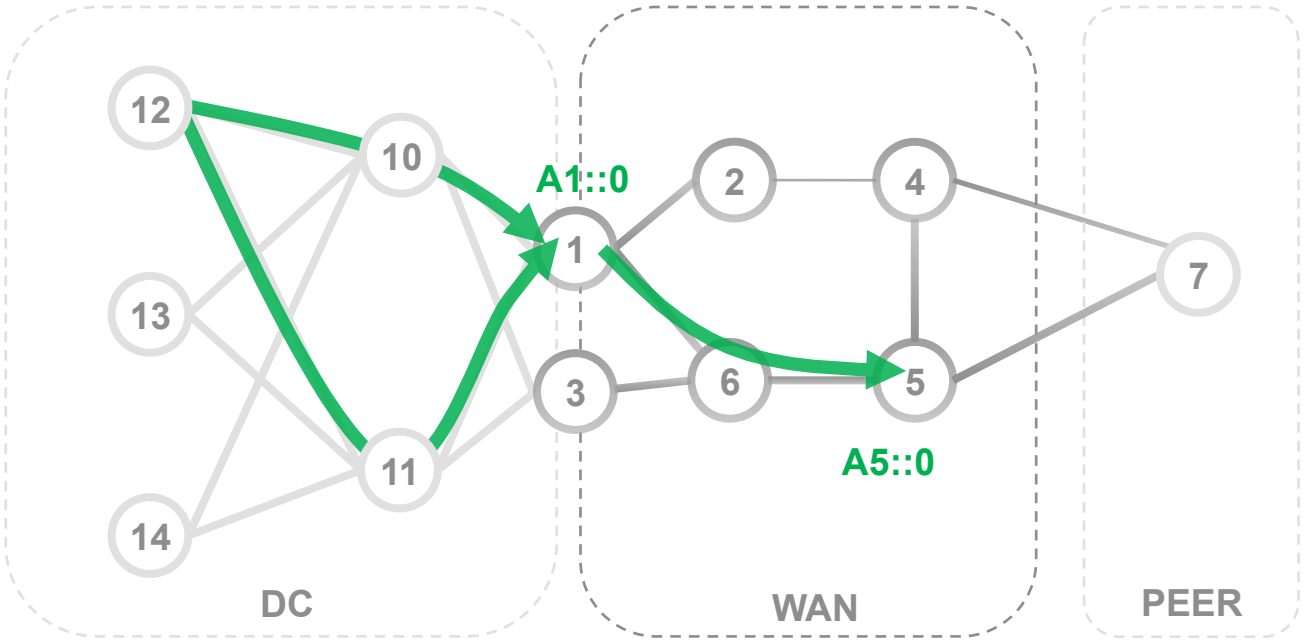

Endpoint then xconnect to neighbor

- For simplicity
- AK::CJ denotes

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

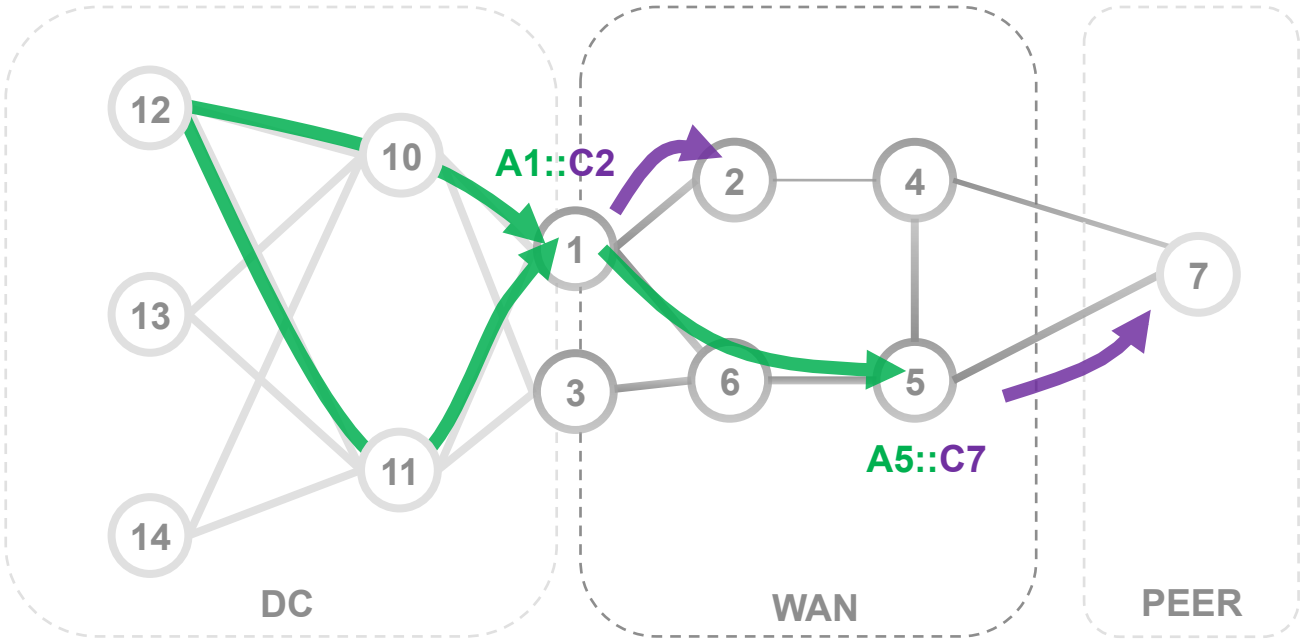

A1::0 and then A5::C7

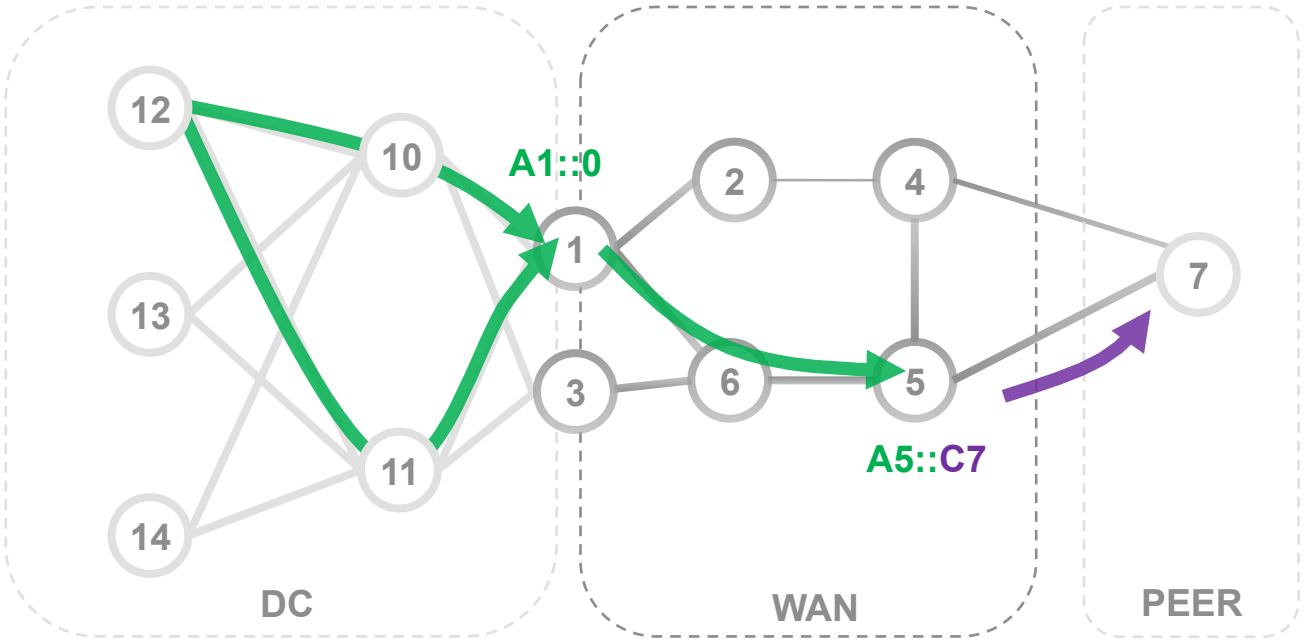

TILFA

- 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 = post convergence)
- 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

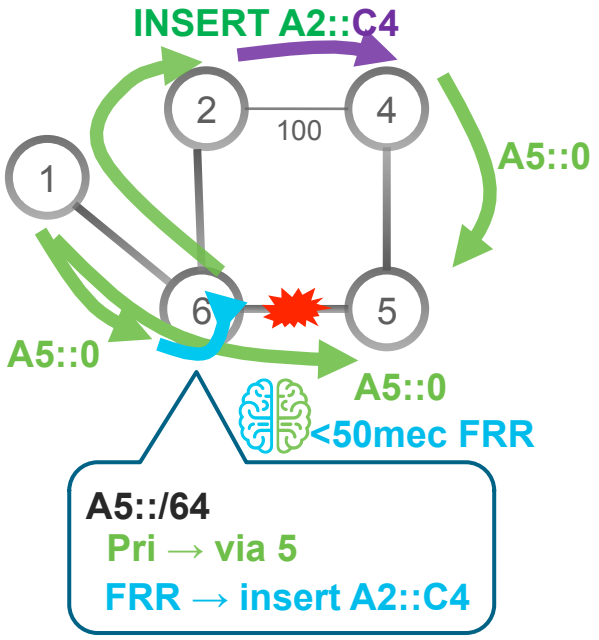

Overlay

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

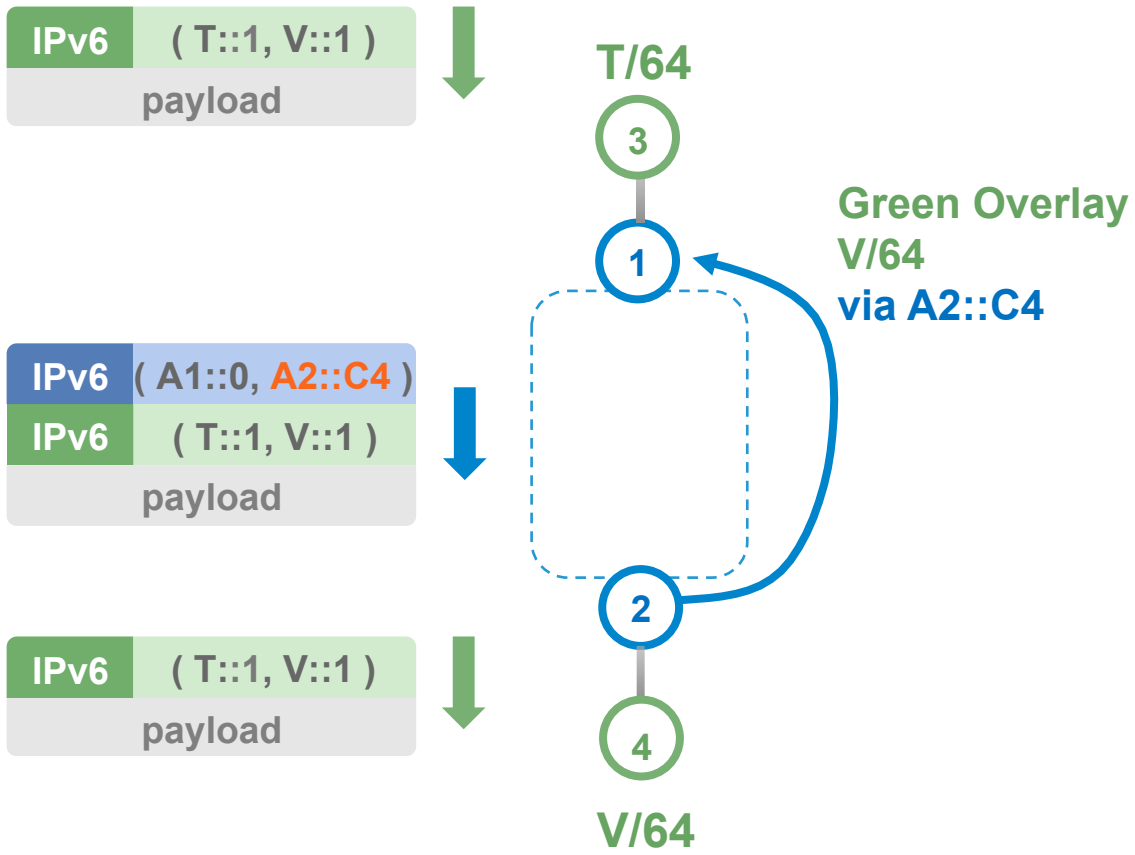

Overlay with Underlay SLA

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

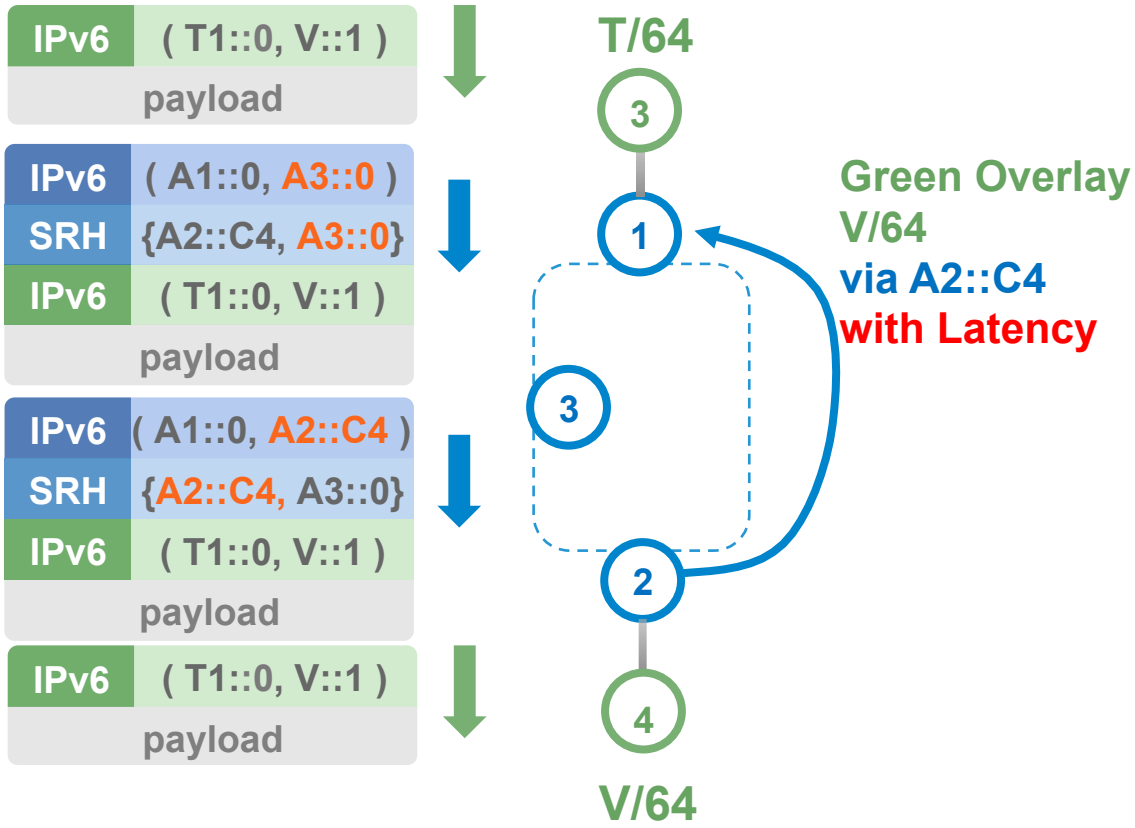

Integrated NFV

- 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

IPv6	(T1::0, V2::0)
payload	

IPv6	(A1::0, A3::A32)
SRH	{ A2::C4, A5::A76, A4::0, A3::A32 }
IPv6	(T1::0, V2::0)
payload	

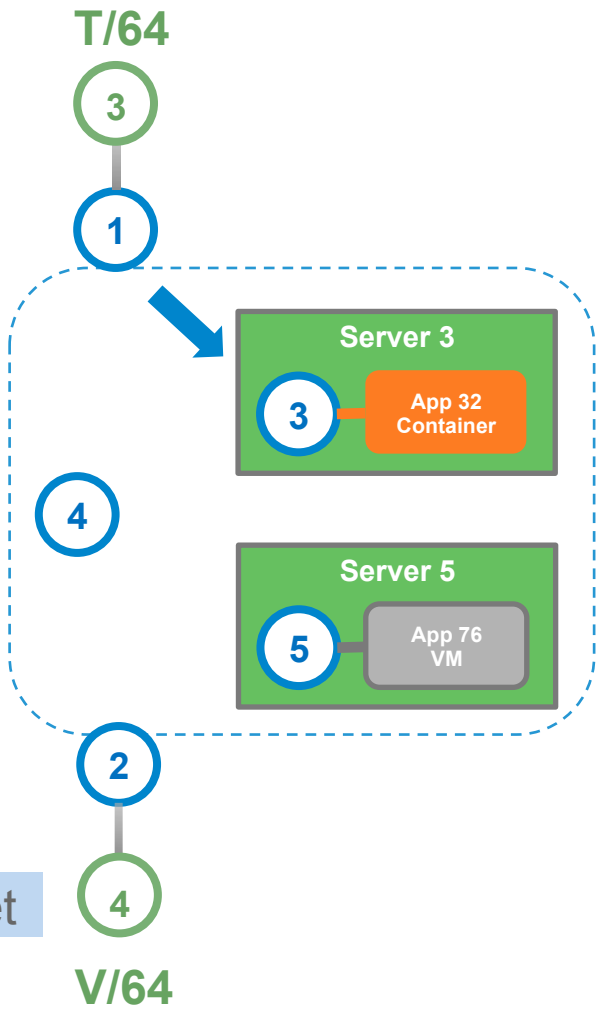

Inner header could be IPv4 or IPv6 or Ethernet

Integrated NFV

- Integrated SLA

IPv6	(A1::0, A4::0)
SRH	{ A2::C4, A5::A76, A4::0, A3::A32 }
IPv6	(T1::0, V2::0)
payload	

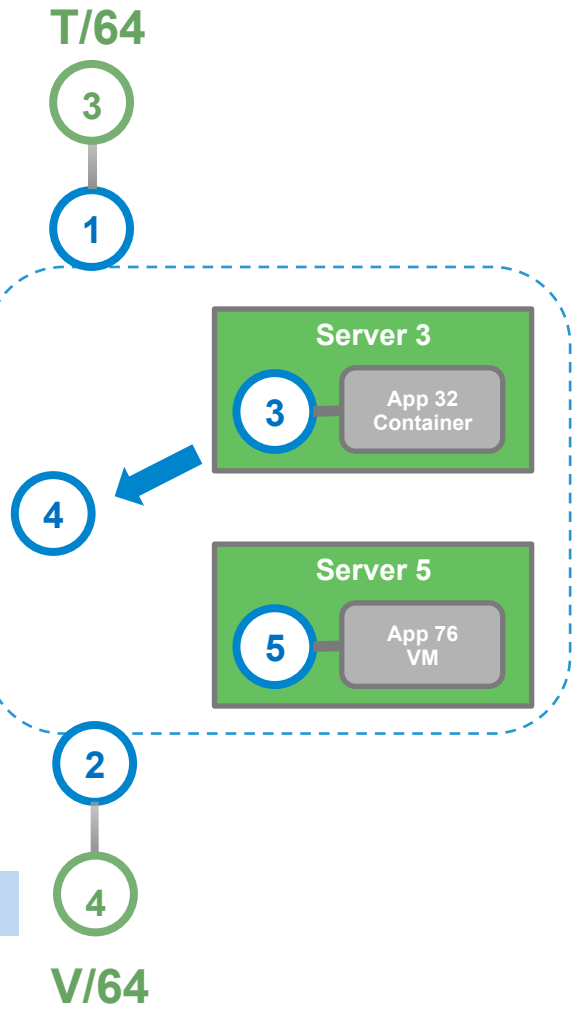

Inner header could be IPv4 or IPv6 or Ethernet

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::0, A5::A76)
SRH	{ A2::C4, A5::A76, A4::0, A3::A32 }
IPv6	(T1::0, V2::0)
payload	

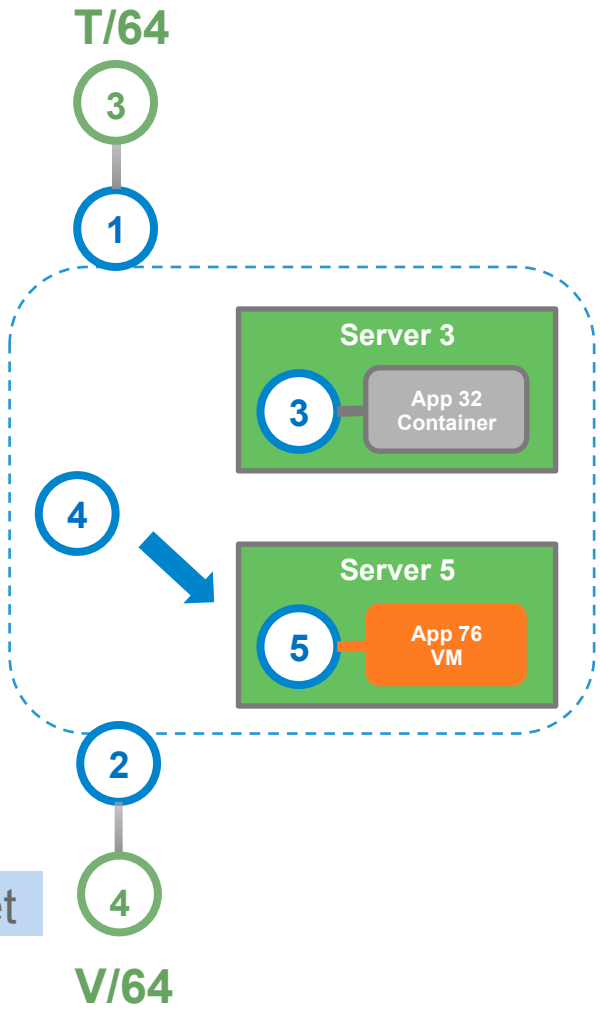

Inner header could be IPv4 or IPv6 or Ethernet

Integrated NFV

- Integrated with Overlay

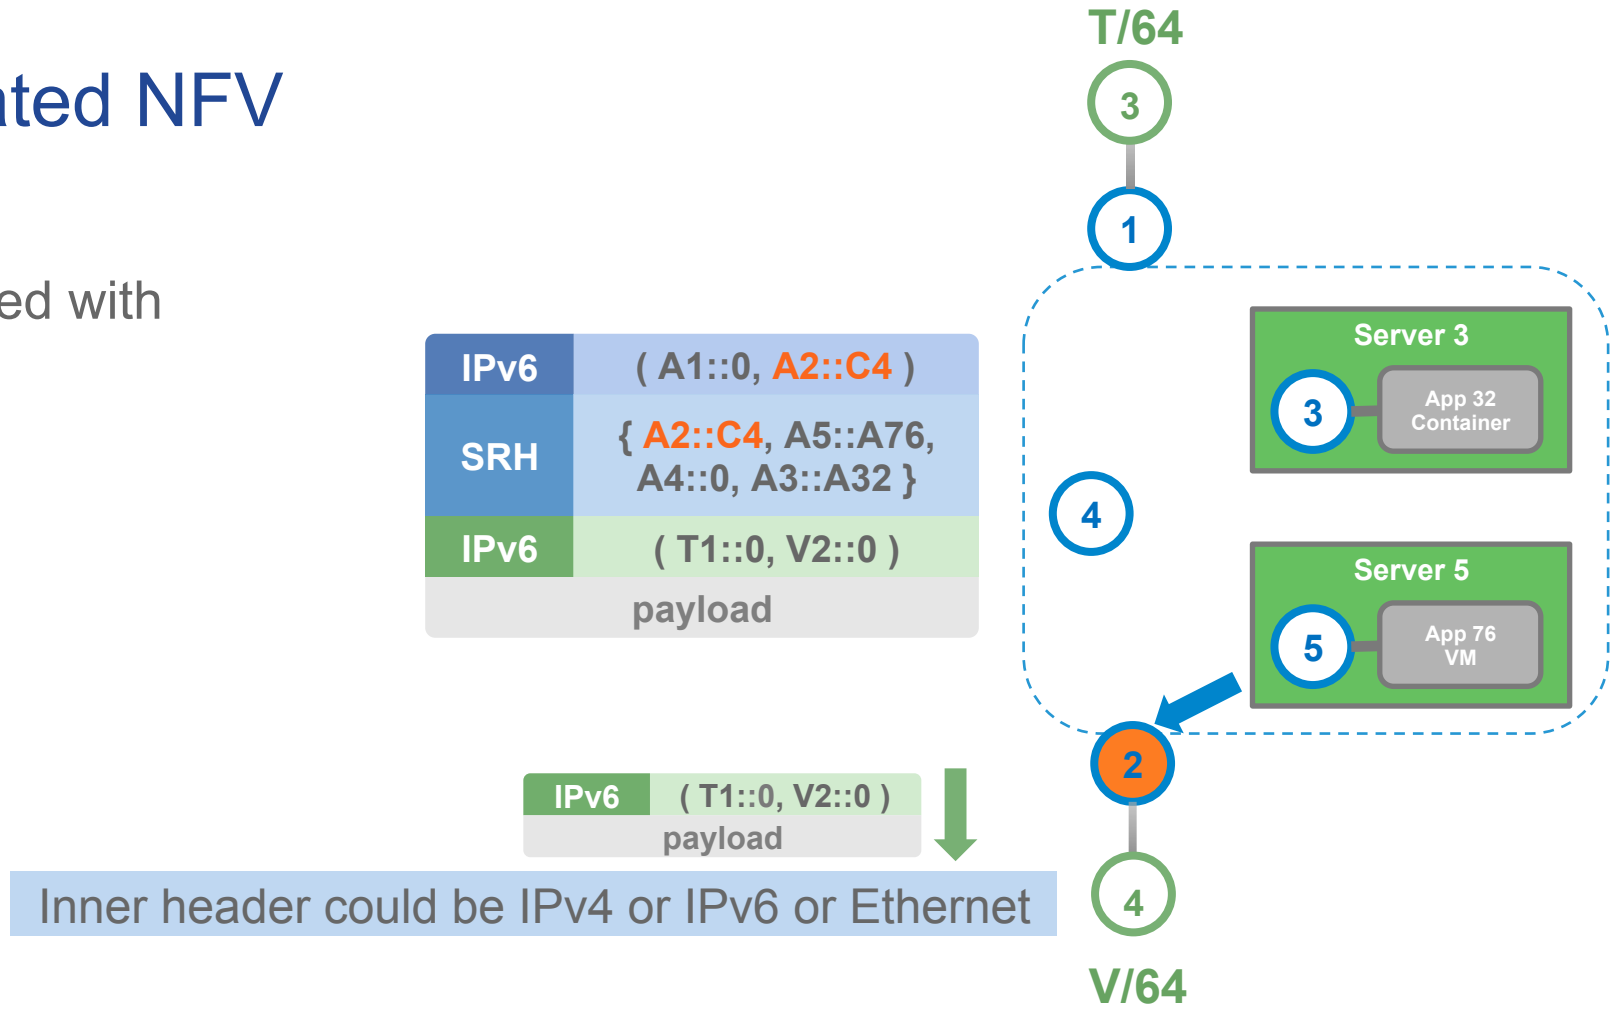

The VPP library

- FD.io project is a collection of several project/libraries to support flexible, programmable and composable services on generic hardware platforms
- VPP is the ‘core’ component of the project

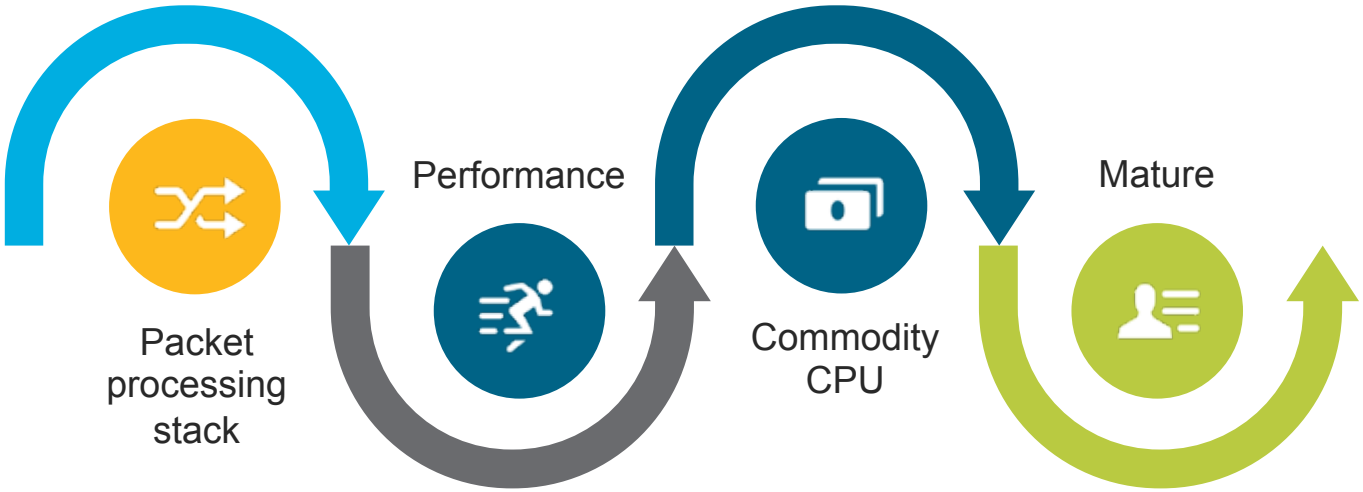

Spray

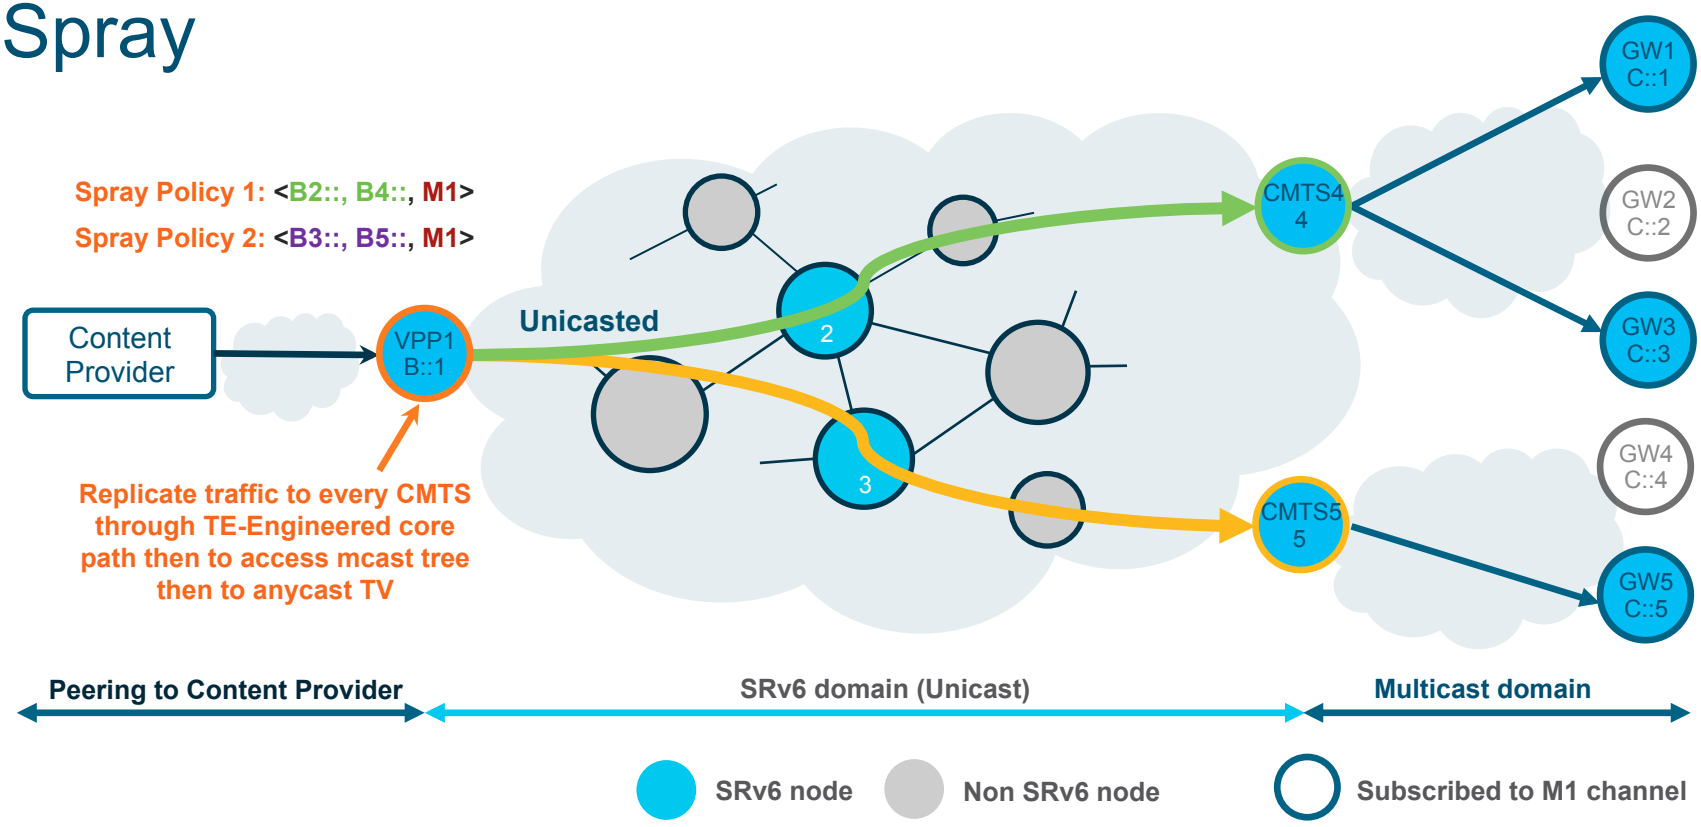

Flexible, SLA-enabled and efficient content injection without multicast core

SD-WAN: Default versus BW versus Latency

- Lisbon (1) to Athens (7)
- Default
  - <16007>
- BW: Guaranteed 50Mbps
  - <16010, 16011, 16007>
  - BSID: 24002
- Low-Latency
  - <16009, 16007>
  - BSID: 24003

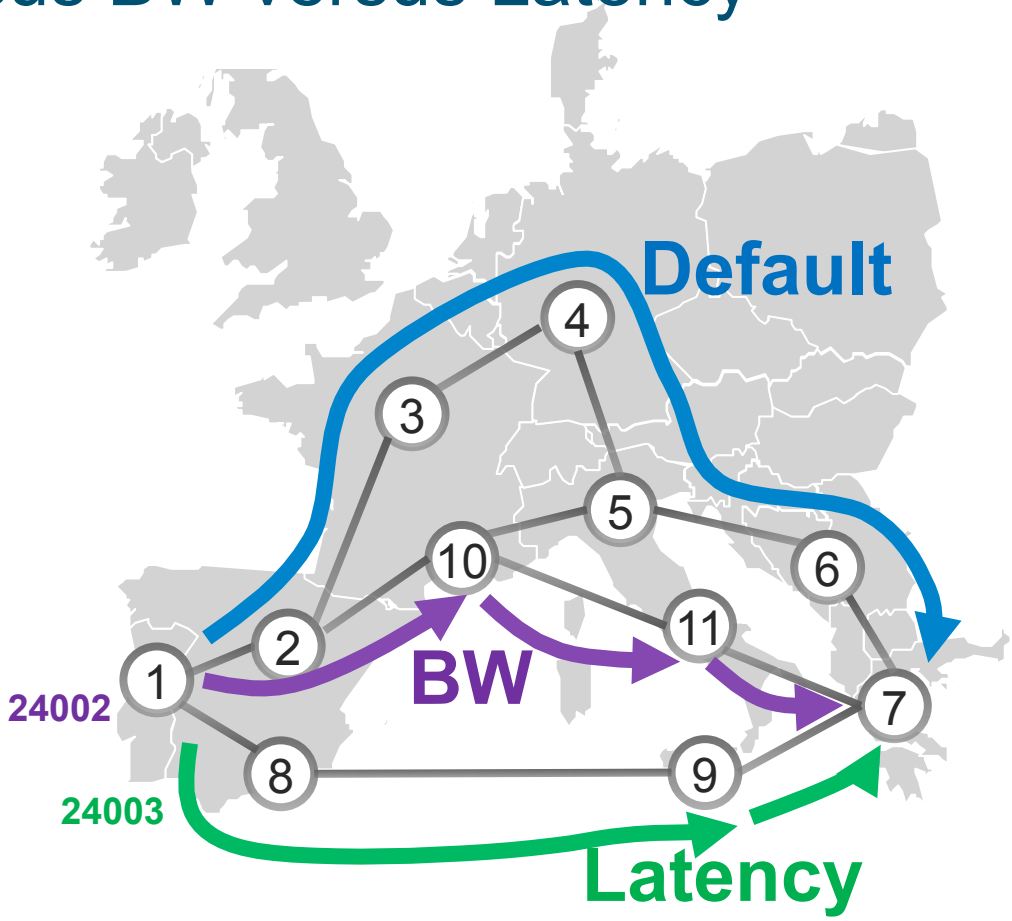

SD-WAN: App needs best-effort

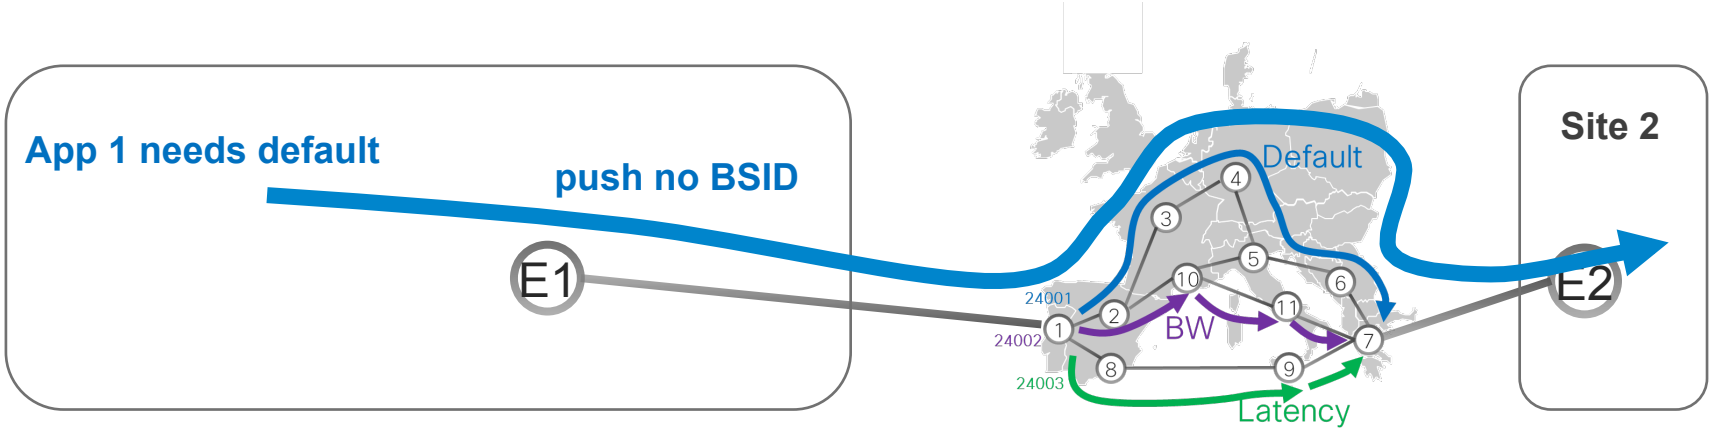

- E1 encrypts the inner packet and encapsulate in outer packet to E2
- E1 does not push any BSID

SD-WAN: App needs guaranteed BW

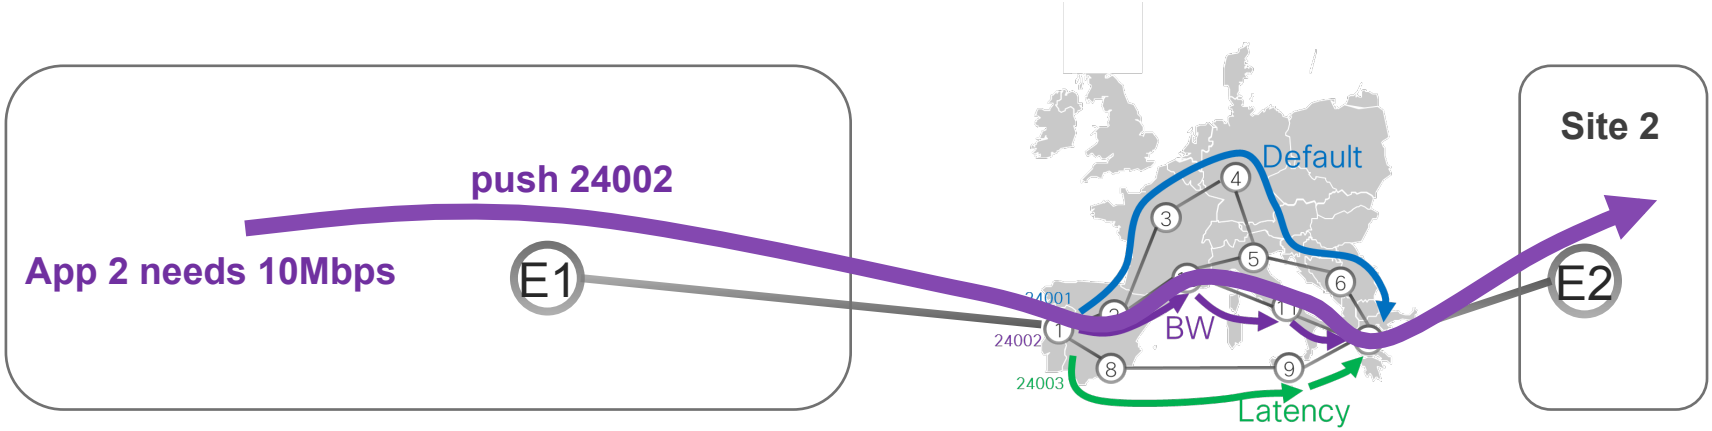

- E1 encrypts the inner packet and encapsulate in outer packet to E2
- E1 pushes 24002
- The network provides the guaranteed BW service to App2

SD-WAN: App needs low-latency

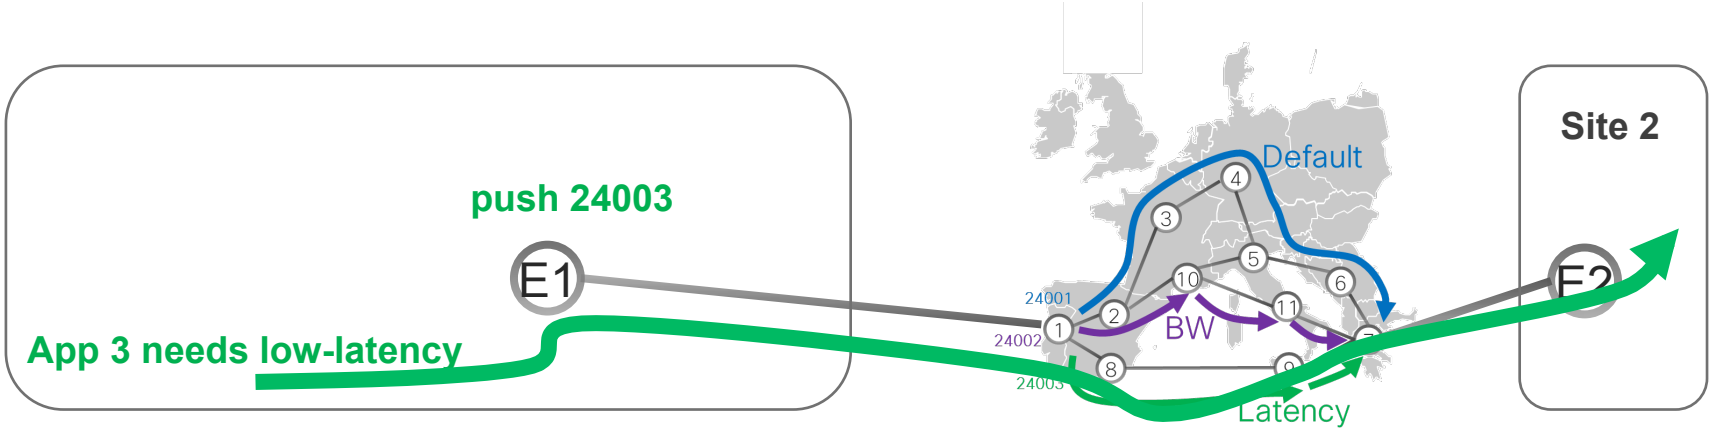

- E1 encrypts the inner packet and encapsulate in outer packet to E2
- E1 pushes 24003
- The network provides the low-latency service to App3

Binding SID is crucial for Enterprise

- Identifier for a customized SLA per application per Enterprise
- Secured
- Per-BSID counters for usage-based billing
- Delegates the application recognition and policy decision to the Enterprise who knows better **when** an application needs a non-default path and **which** non-default path is needed

Performance Monitoring – Enterprise-based

- Enterprise can easily monitors each individual service
- Simply sends the probes with the related BSID

Performance Monitoring – SP-based

- The SP can enable per-SR-policy perf monitoring
  - latency
  - loss
- These metrics can be leveraged by SDWAN controller and provided to the Enterprise
  - Simply reporting
  - Additional data to select which application to steer on which BSID

Current Mobile Network Example

- Well fragmented to RAN, EPC and SGi.
- Per-session tunnel creation and handling.
- Non-optimum data-path.

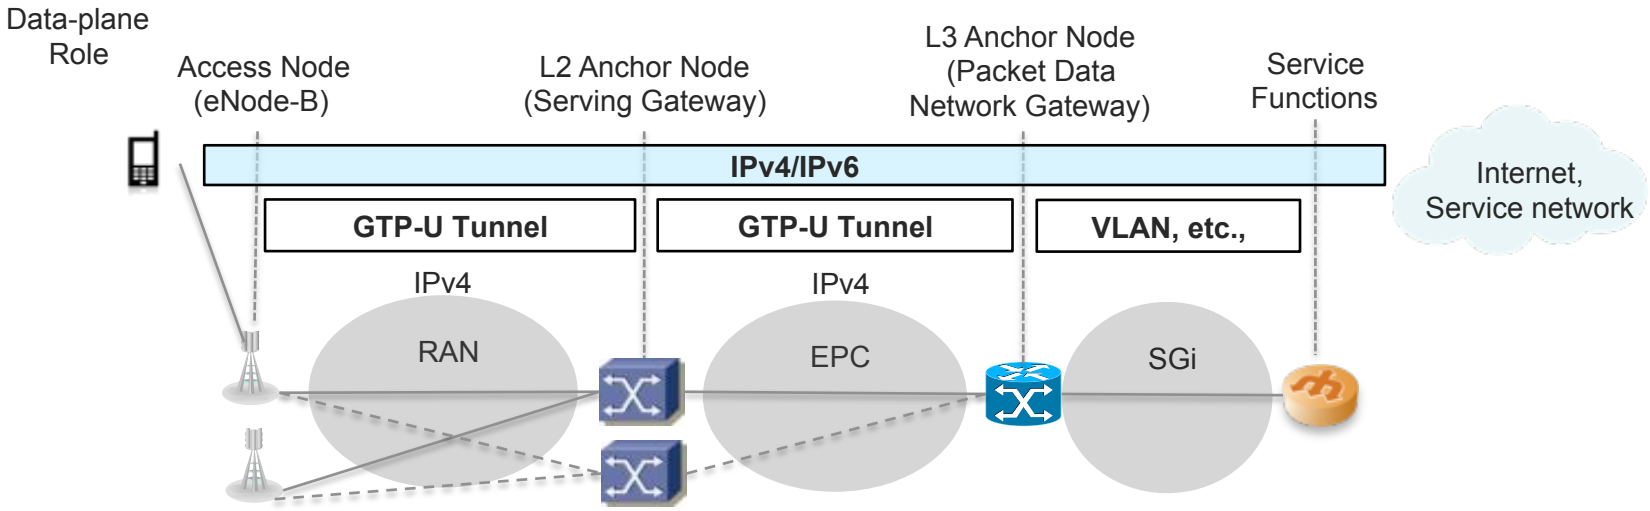

5G: SID Functions* for Mobile Data-Plane Roles

	Uplink	Downlink
Access Node	T.Insert, or T.Encaps	END.X/END.DX {6 | 4}
L2 Anchor Node	END or END.B6	END or END.B6
L3 Anchor Node	END.T/END.DT {6 | 4}	T.Insert, or T.Encaps

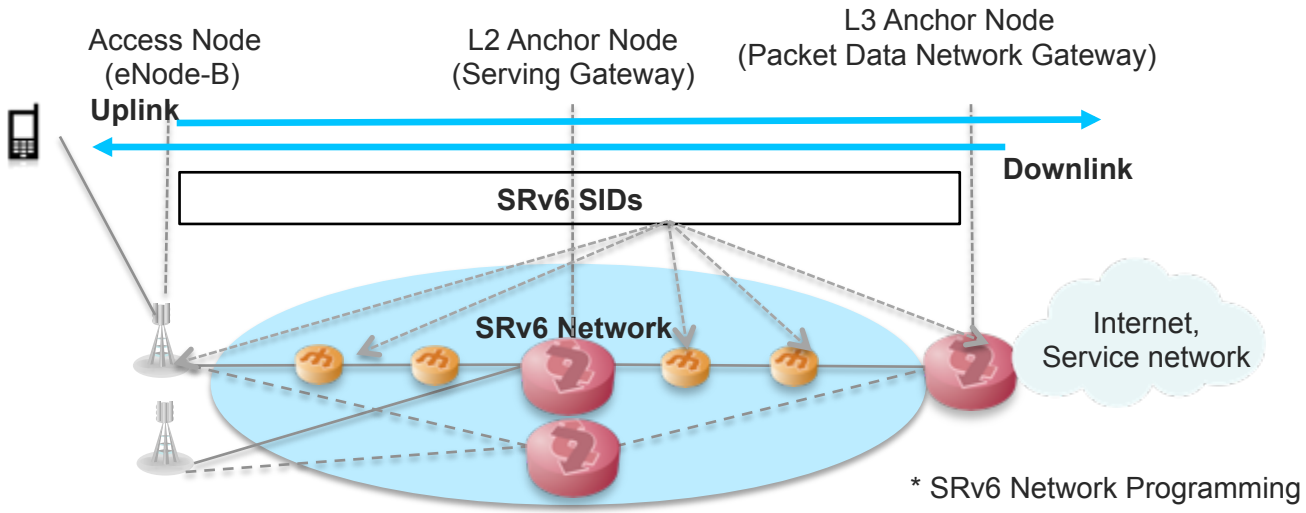

More use-cases

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

Conclusion

Homework time!

- Go on **www.segment-routing.net** and check the latest demo
- Read the IETF draft: **draft-filsfils-spring-srv6-network-programming**
- Read the IETF BGP draft:**draft-dawra-bgp-srv6-vpn-00**
- Read the IETF ISIS draft:**draft-bashandy-isis-srv6-extensions-00**