
# SRv6

Sai Nyan Lynn Swe

CCIE # 38501 (R&S , SP and DC)

March 15, 2024

# At record speed

- 9+ large-scale commercial deployments
  - Softbank, Iliad, China Telecom, LINE corporation, China Unicom, CERNET2, China Bank, MTN Uganda, NOIA Network, …
- 25+ HW linerate implementations
  - Cisco Systems, Huawei, Juniper
  - Arrcus, Broadcom, Barefoot, Intel, Marvell, Mellanox, Kaloom
  - Spirent, Ixia
  - Multiple Interop Reports
- 11+ open-source platforms/ Applications
  - Linux, FD.io VPP, P4, iptables, nftables, snort, SERA, ExaBGP, GoBGP, GoBMP, Contiv-VPP, …

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

# IPv6 provides reachability

# Network requirements for today & beyond

**High scale of network**
As 5G and cloud develops, IoT and virtual nodes bring in large number of network connections, which require high scale of IP addresses.

**High-quality Connections**
Low-latency & bandwidth guarantee enhanced user-experience. Such as cloud AR/VR services require low delay, which driving the demand of data path with traffic-engineering.

**Network programmable**
Smart and automate way to set up connection that allow service provision in hours instead of weeks. Also, easier to locate the faults in minutes instead of days.

# Network Scalability and Technology Shifting


IPv4 address space is exhausted

IPv6 is inevitable

IPv6 migration solutions are ready

Dual stack is recommended & widely deployed

IPv6 is growing faster than IPv4

Adopting IPv6 will show technology leadership & enable business innovation

IPv6 enables advanced innovations

SRv6, 5G, Cloud Computing & more futuristic technologies are fueled by IPv6

# Understanding Today's Service Creation

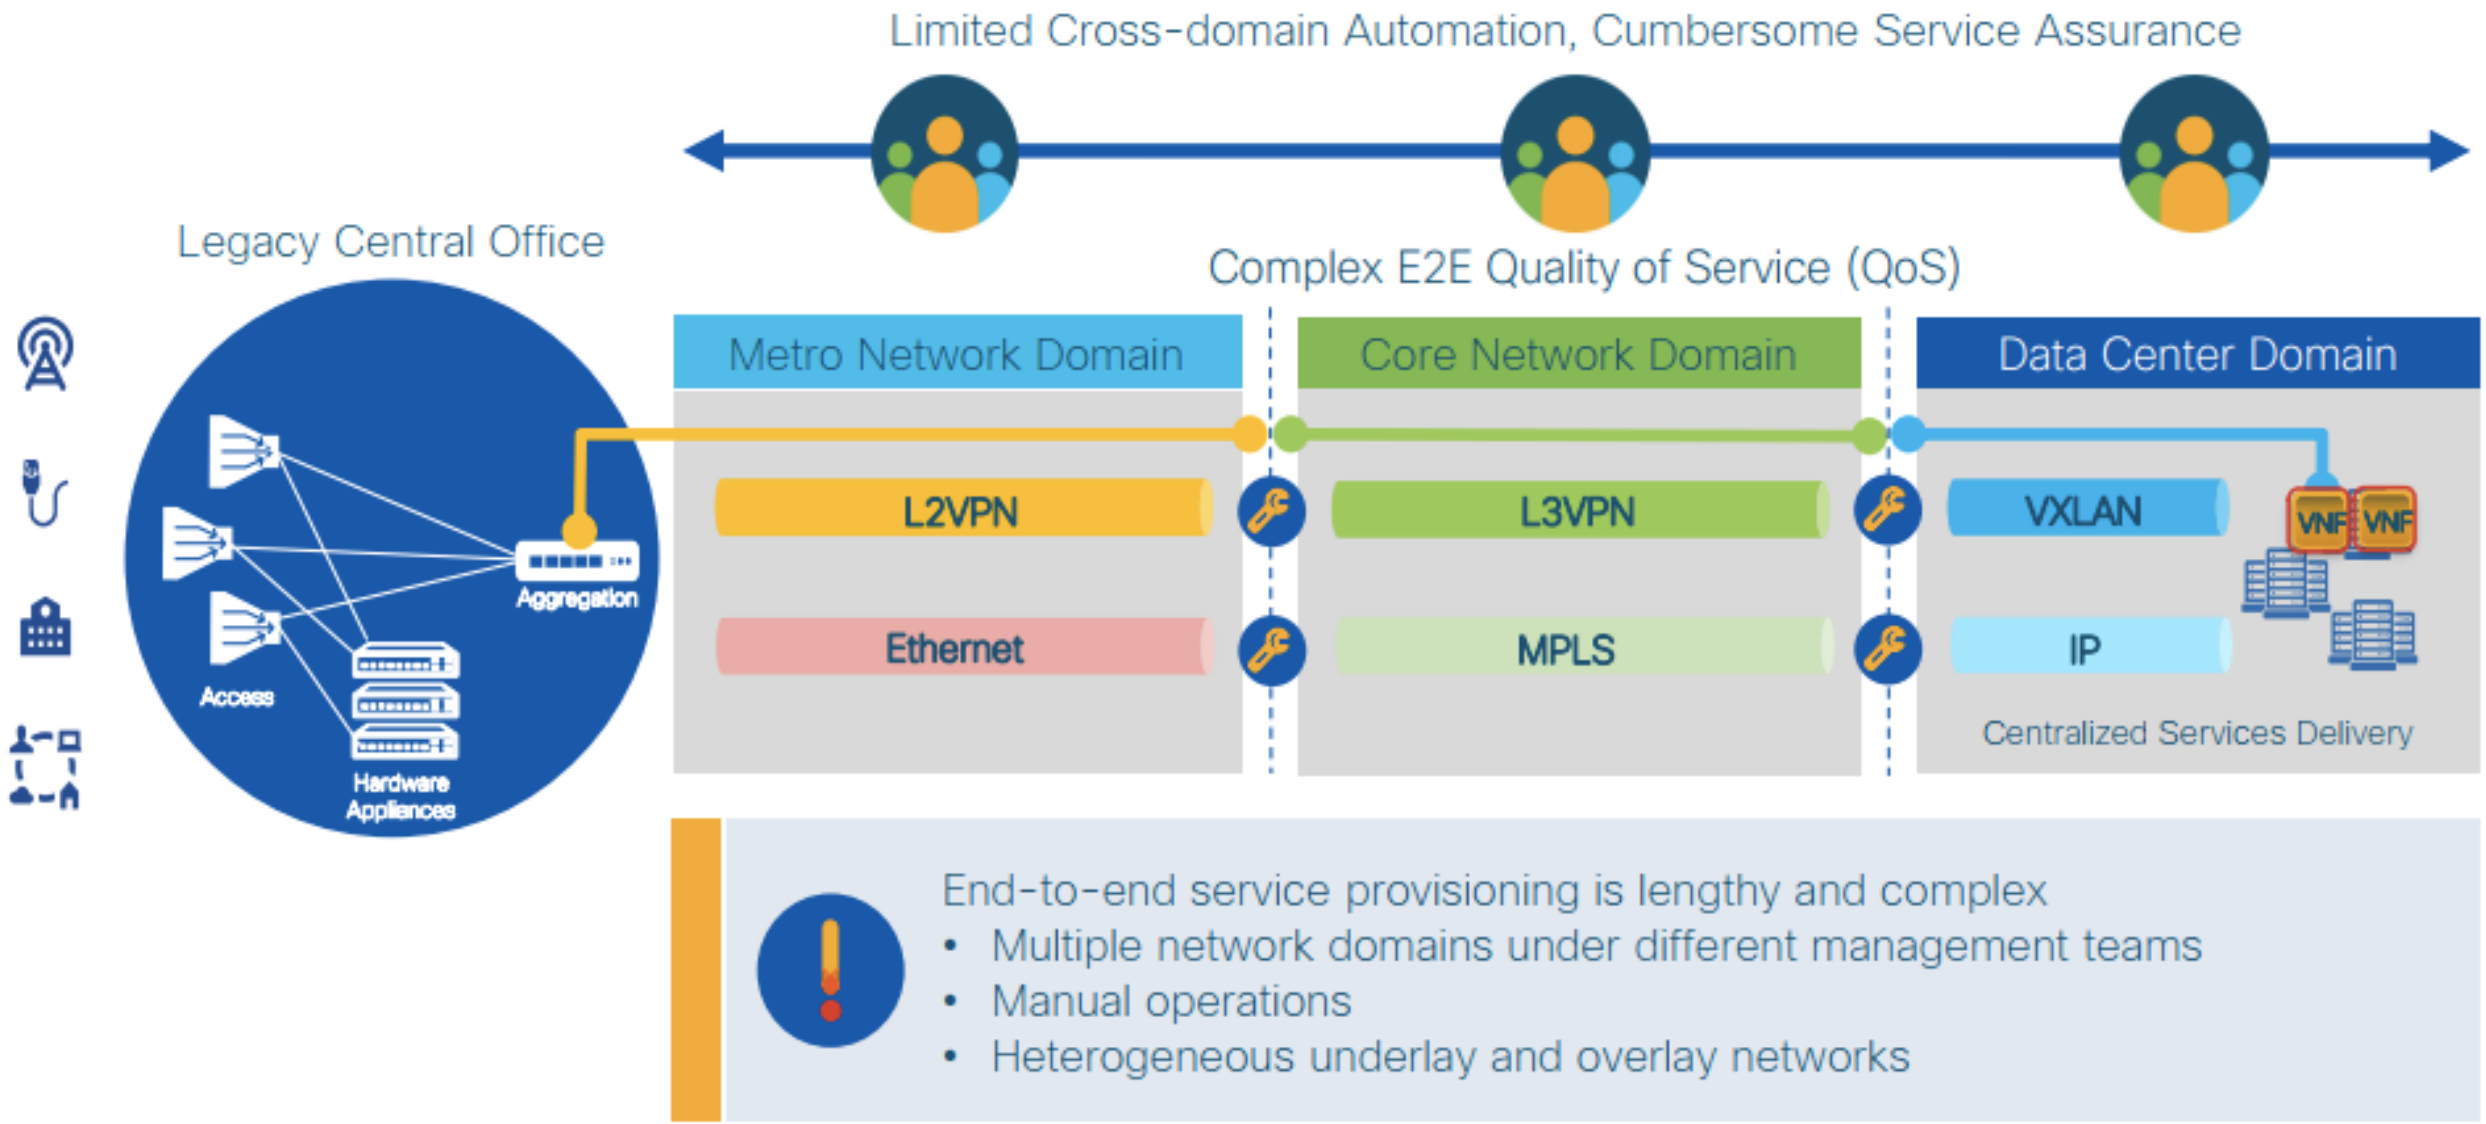

# SR-MPLS: SDN ready
# "Network as a Fabric" for Service Creation

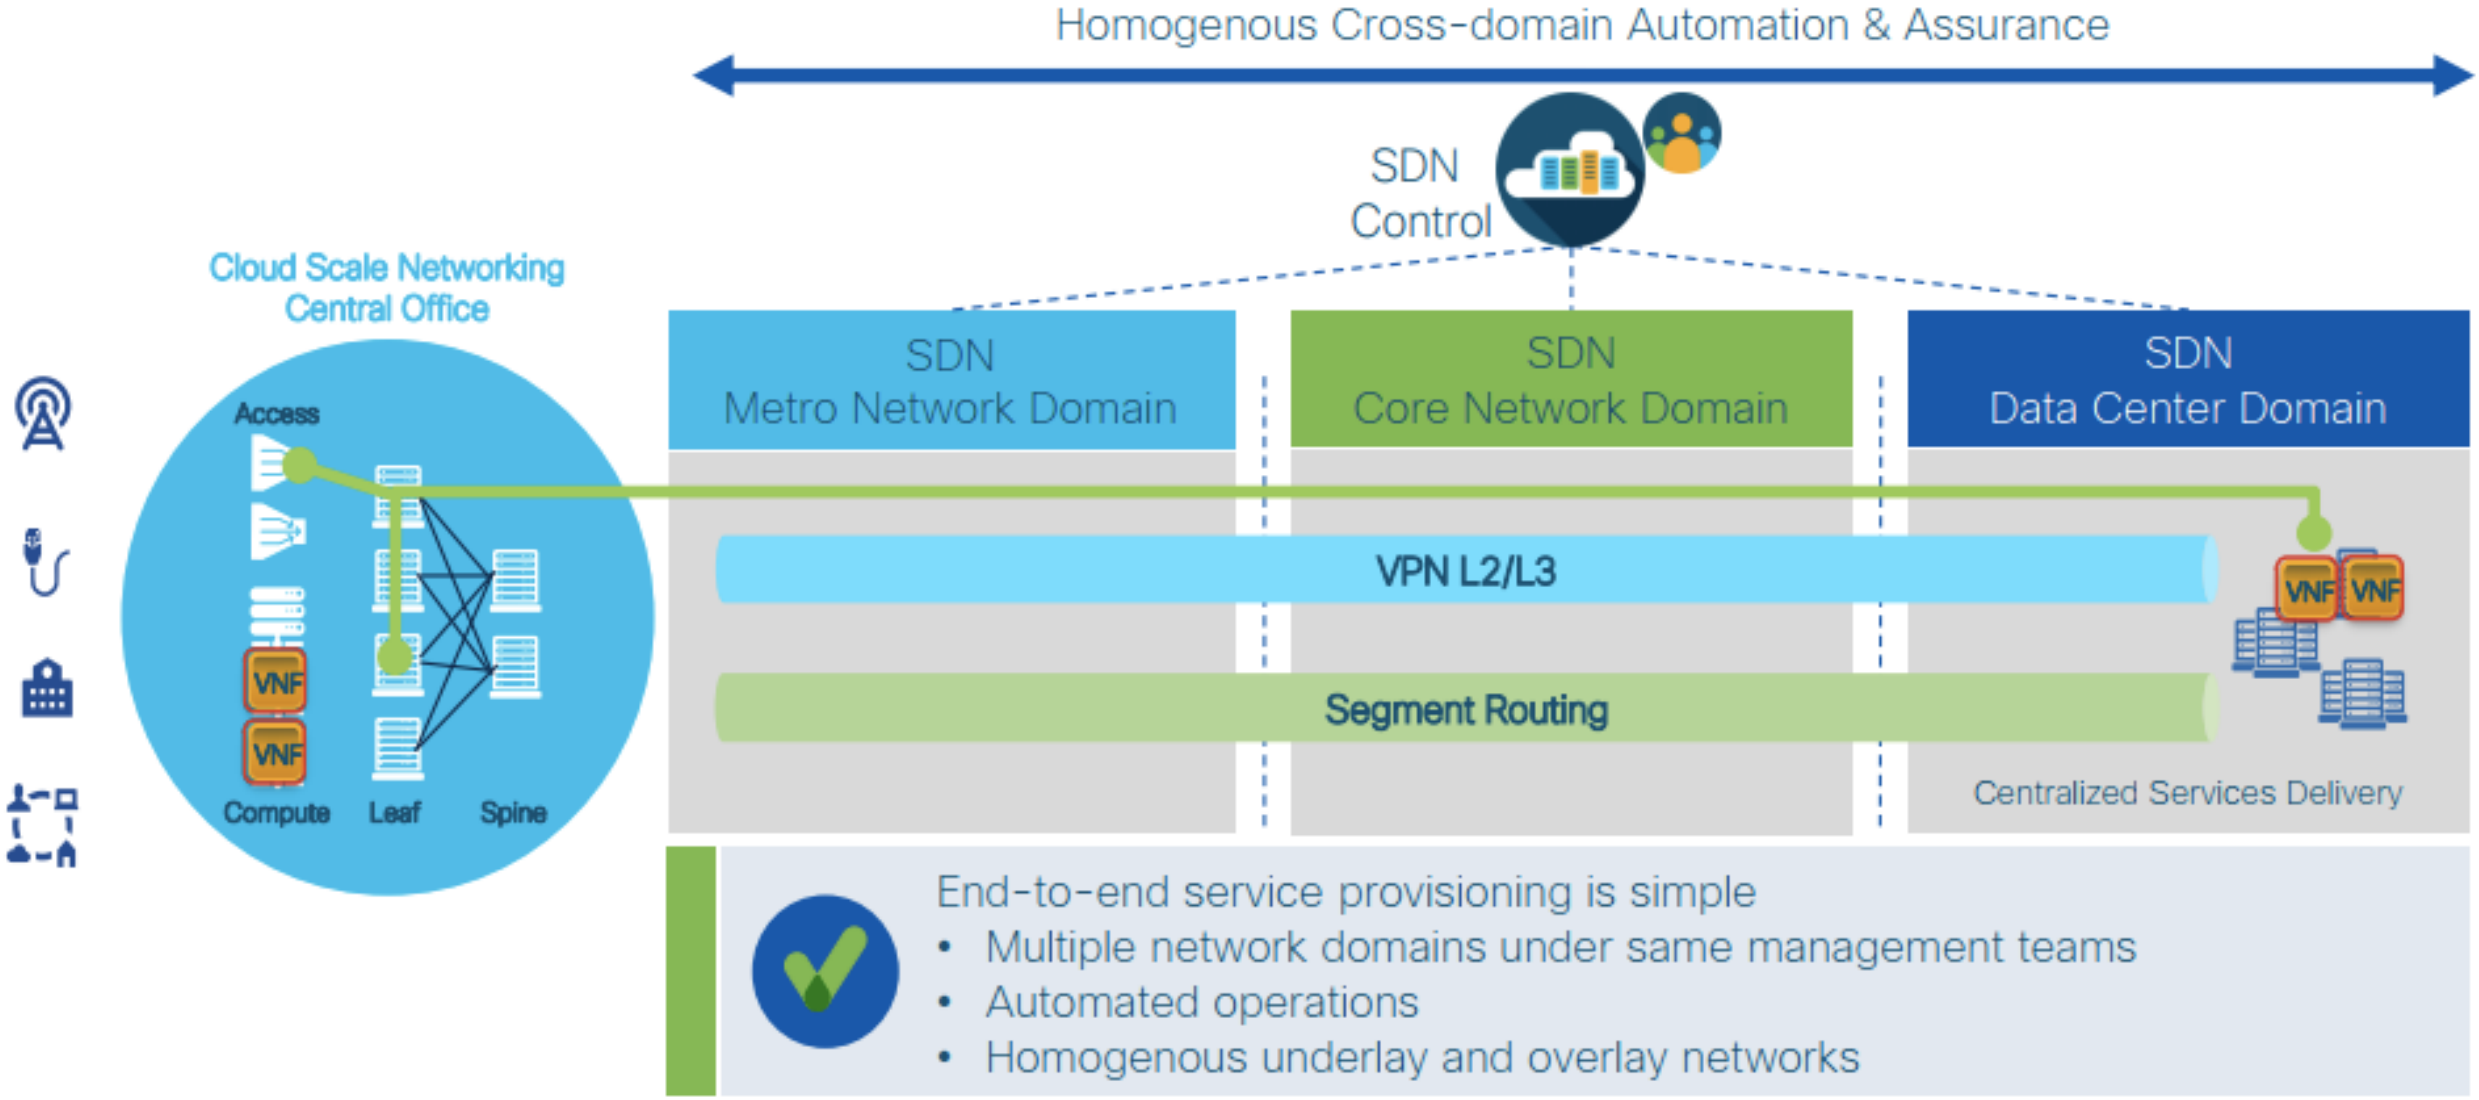

# SRv6: SDN, NfV, 5G ready
# “Network as an API” for Service Creation

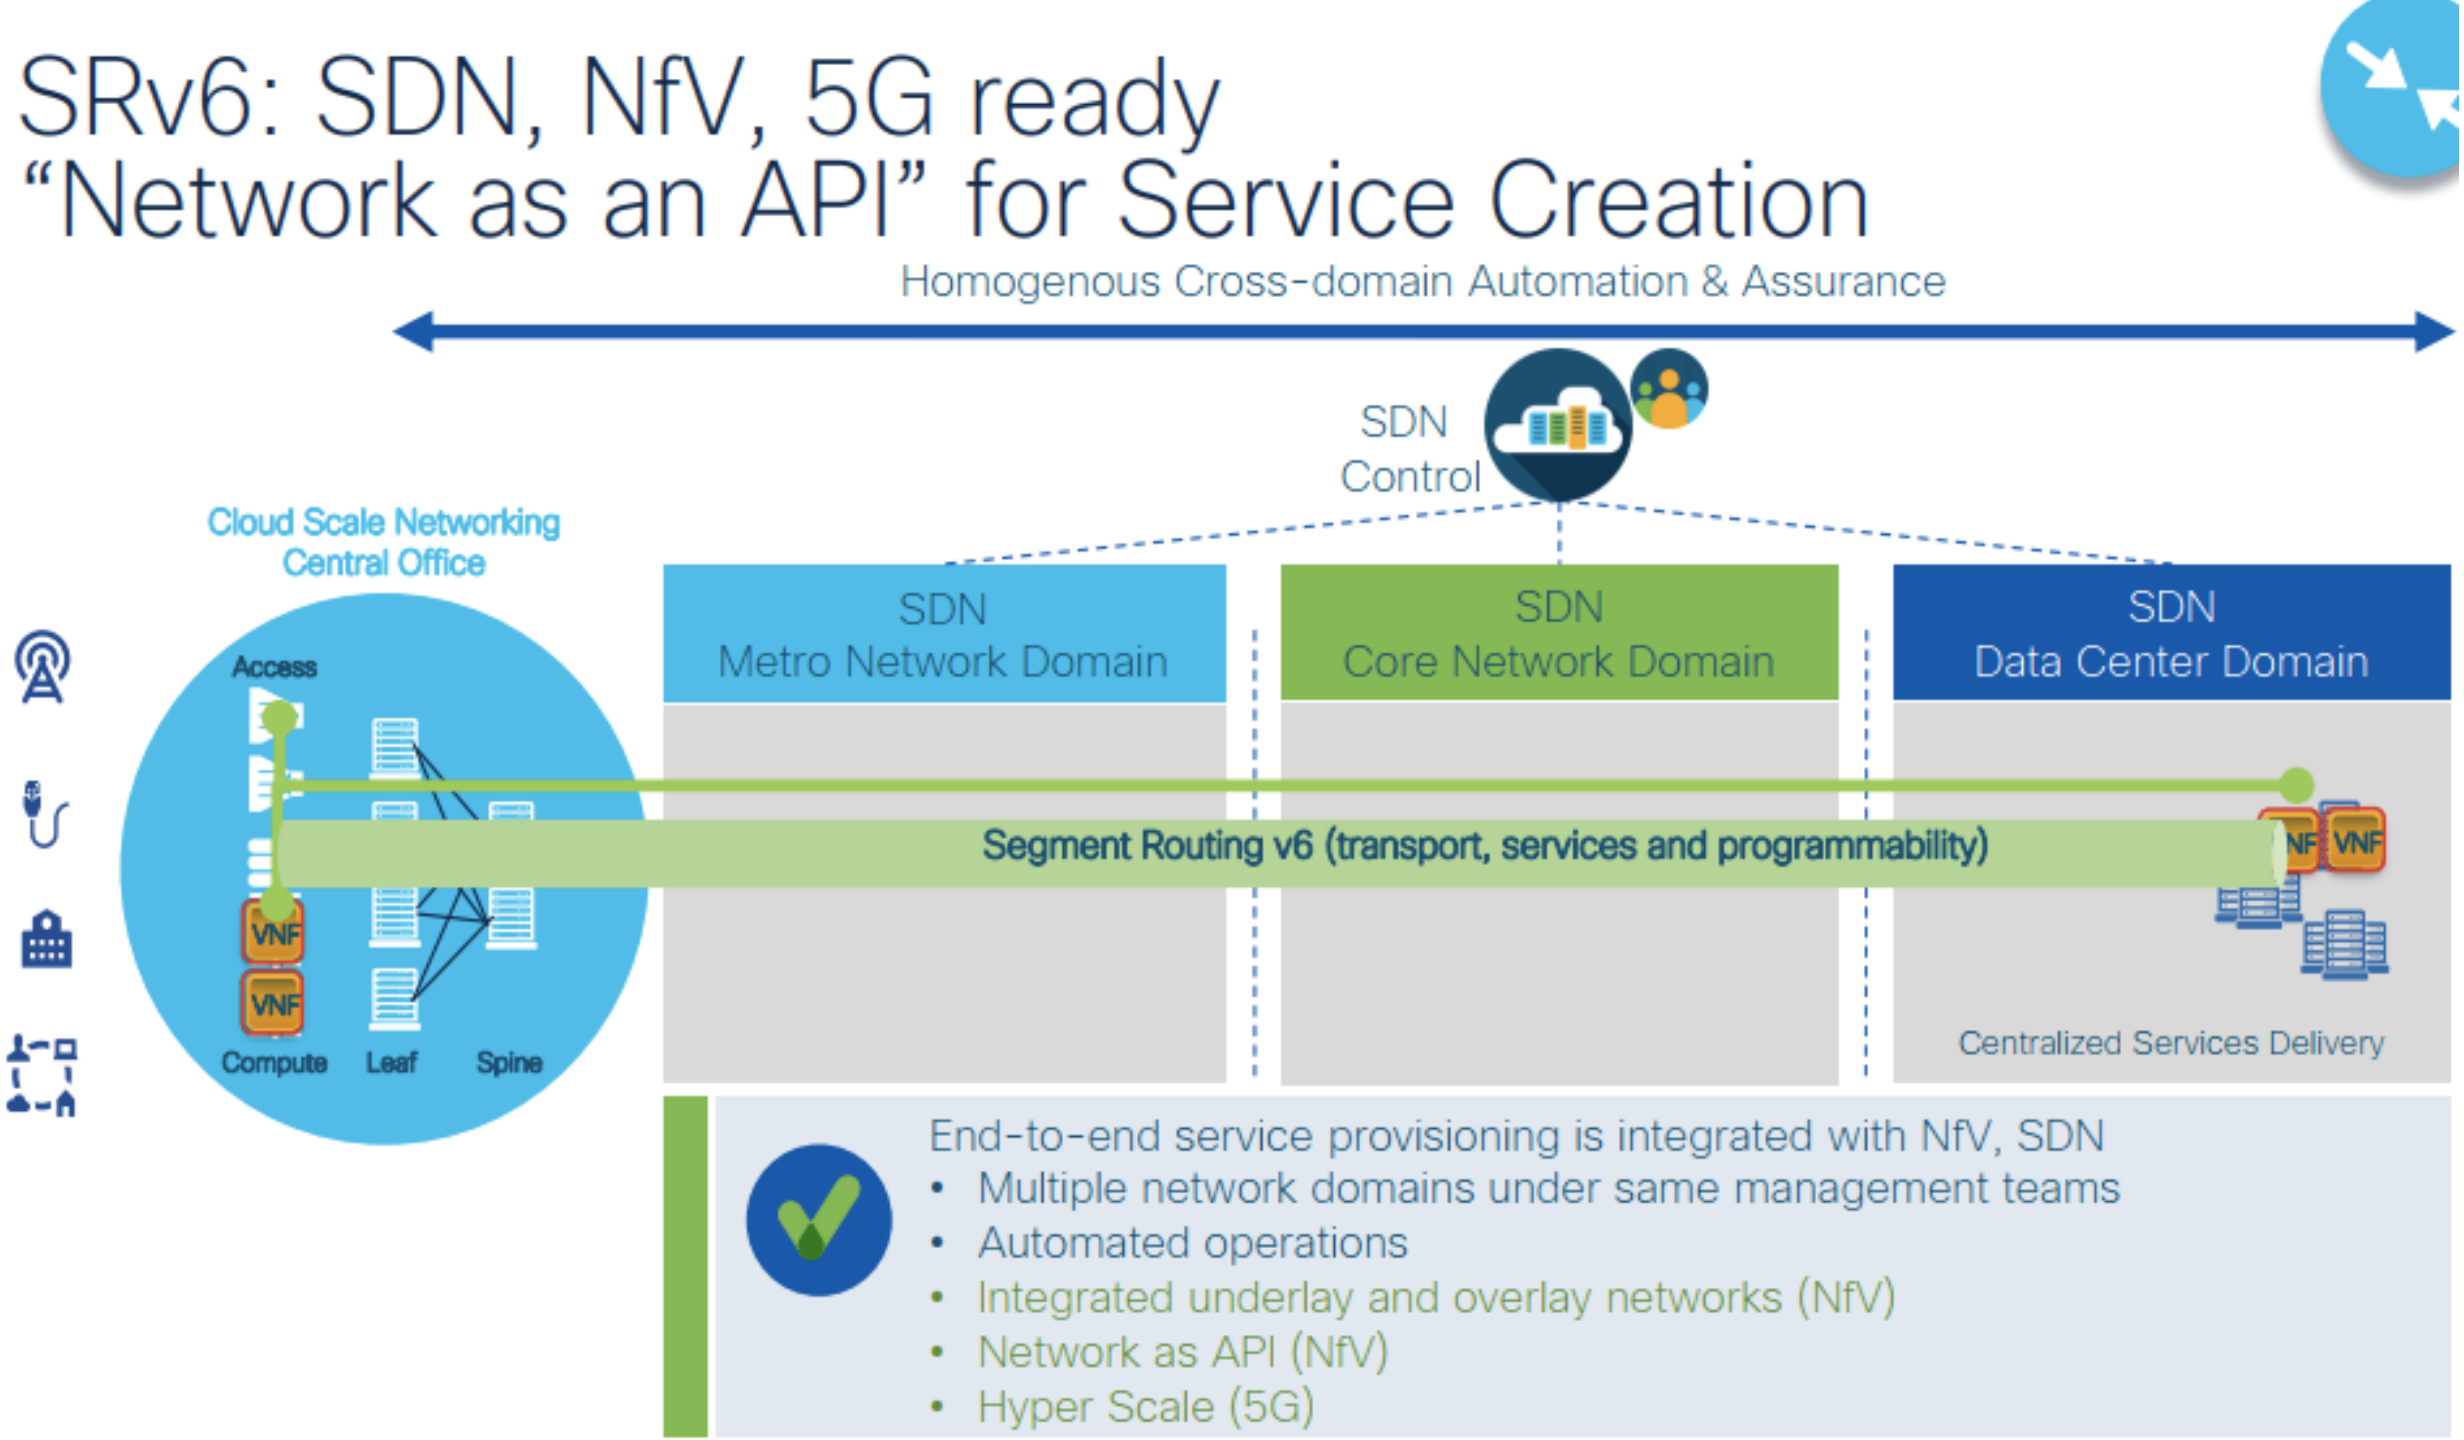

# L2/L3 Service over Classic MPLS + BGP-LU

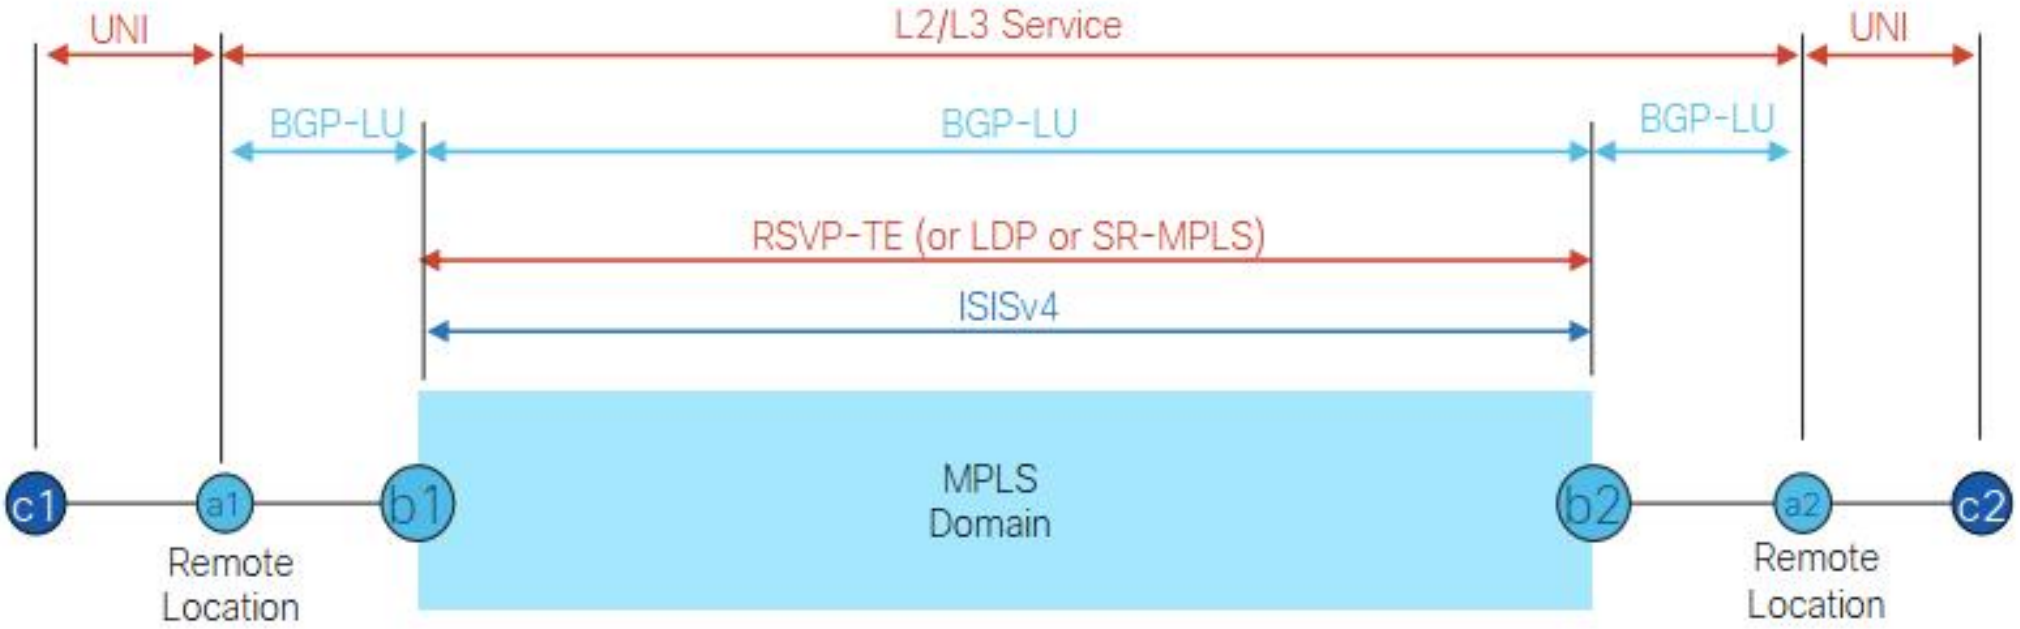


- This common design requires multiple protocols with significant complexity
- BGP-LU is used as a "shim layer" primarily to reduce the size of the IGP domain
- In many cases LDP or SR-MPLS are used instead of RSVP-TE, but overall, the picture doesn't change

Client device/router

MPLS Node

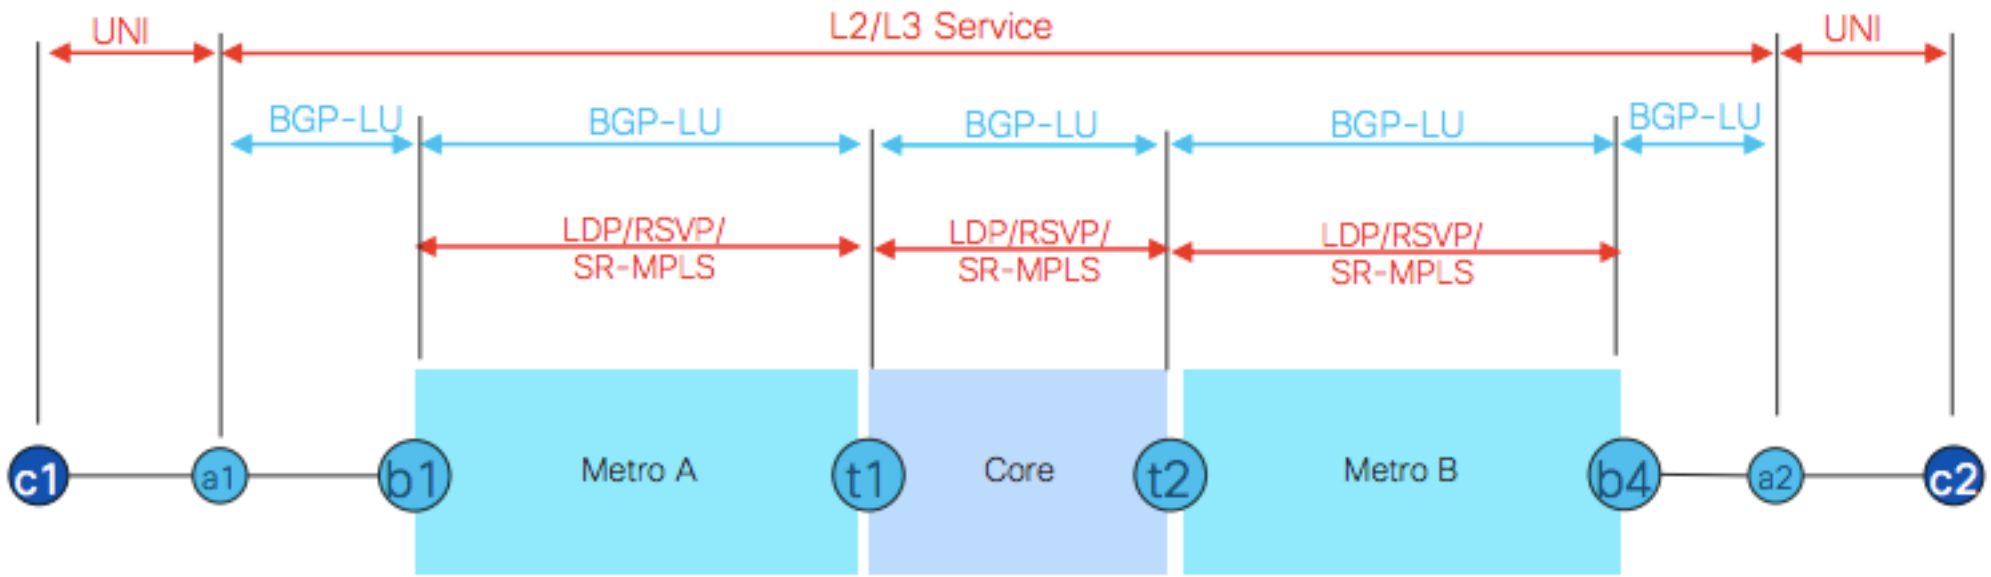


- This common design requires BGP-LU to stitch between domains within the SP
  - Also commonly used to stitch between sub-domains within a domain
- BGP-LU adds a substantial tax of complexity and limits scalability

Client device/router

MPLS Node

# Load-Balancing Challenges with MPLS

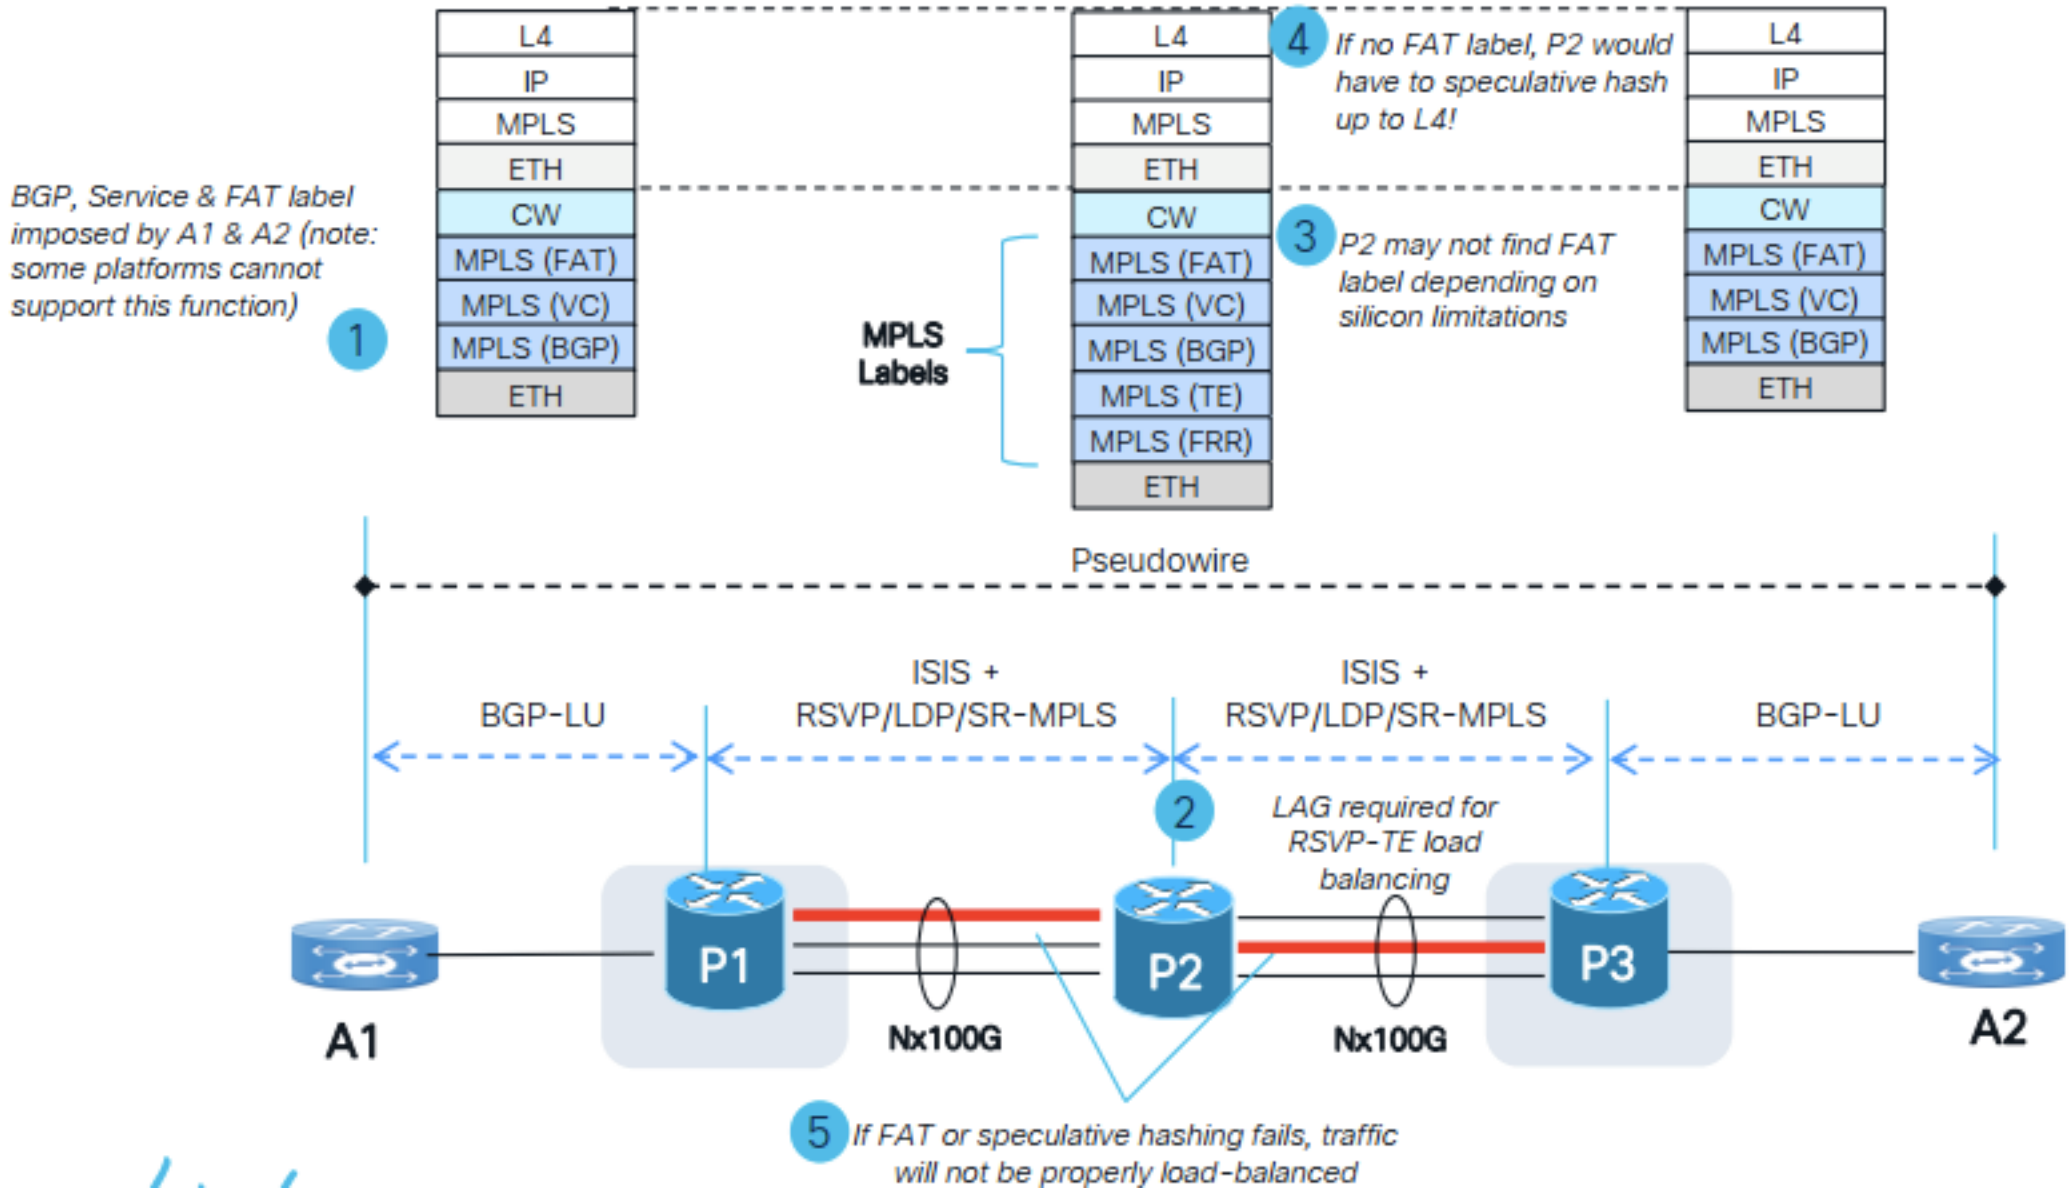

# IPv4 limitations & work-arounds

| Network Functions | IPv4 |
|---|---|
| Reachability | IPv4 Header |
| Engineered Load Balancing | MPLS Entropy Label, VxLAN UDP |
| VPN | MPLS VPN's, VxLAN |
| Traffic Engineering | RSVP-TE, SR-TE MPLS |
| Source Routing | SR-TE MPLS |
| Service Chaining | NSH |

Address space 32-bit limitation
No optional header
IPv4 header doesn't support
- VPN
- Traffic Engineer
- Service Chaining
- Engineered Flow optimization
- Source-Routing

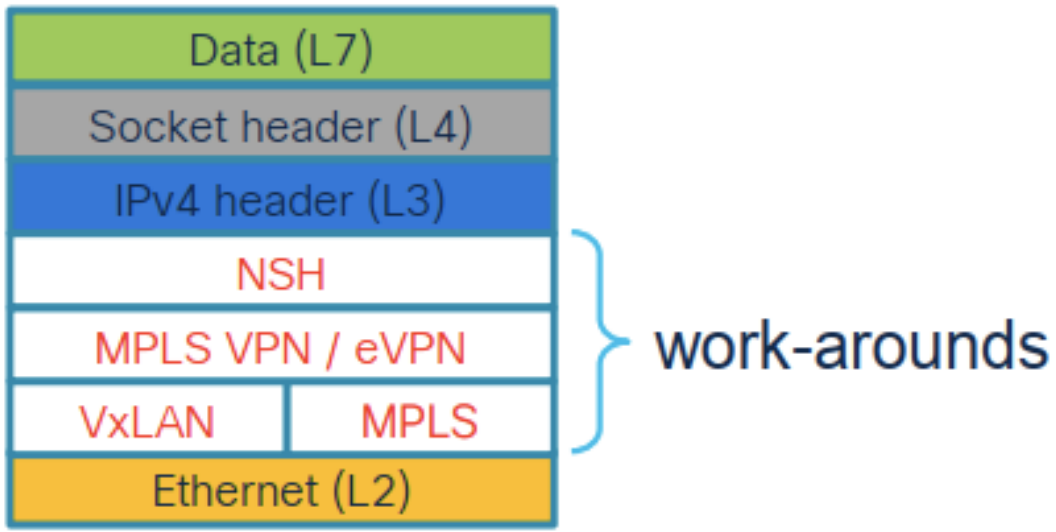

# SRv6 Solution

| Network Functions | IPv6 |
|---|---|
| Reachability | IPv6 Header |
| Engineered Load Balancing | IPv6 Header |
| VPN | IPv6 Header |
| Traffic Engineering | IPv6 Header |
| Source Routing | IPv6 Header |
| Service Chaining | IPv6 Header |

IPv6 Address 128bits
IPv6 Flow Header
- Engineered Flow optimization

SRv6 Header
- Source-Routing
- Traffic Engineering
- VPN
- Service Chaining

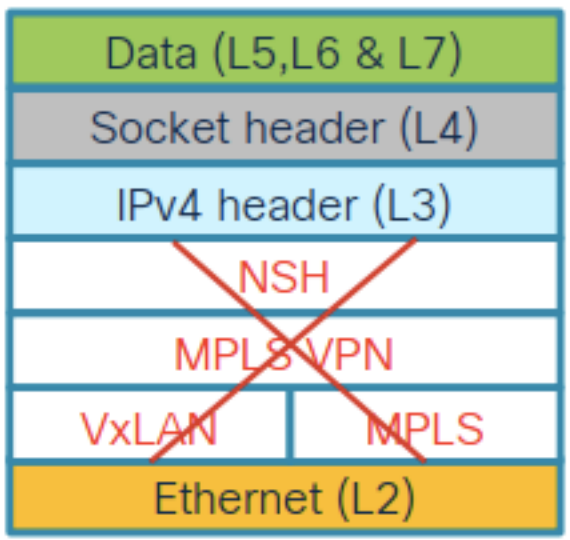


SR

Simplicity
(back to OSI model)

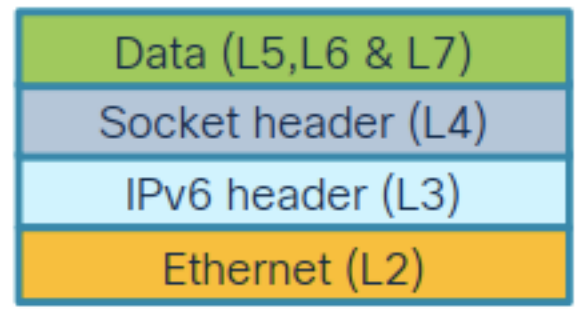

# Network Evolution

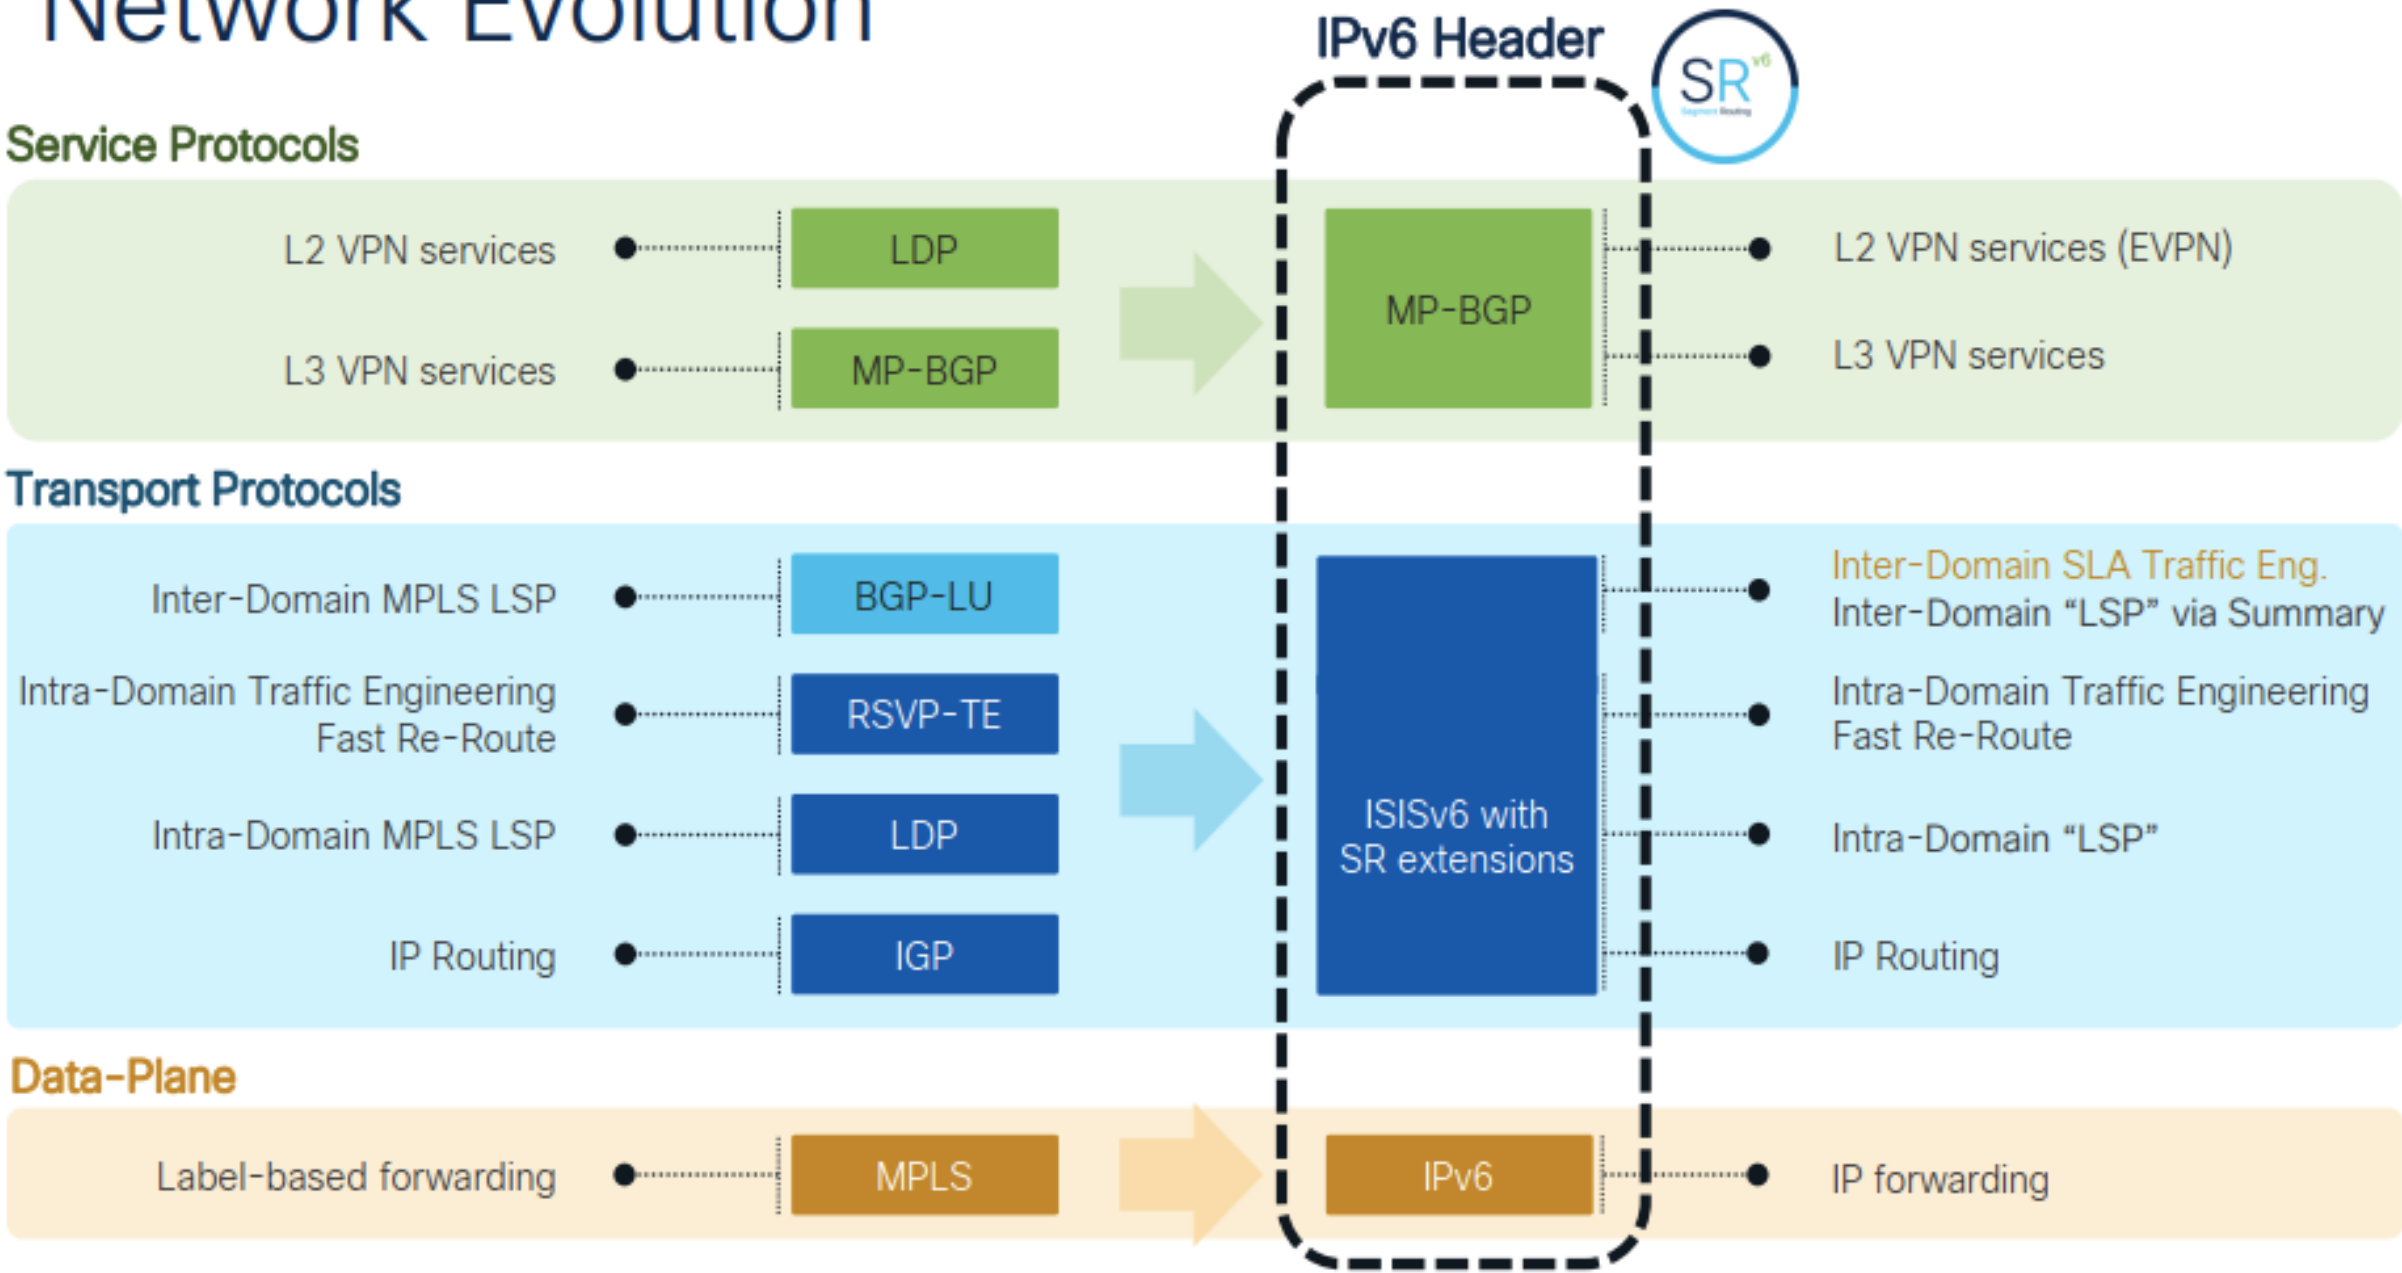


**LDP**: Label Distribution Protocol, **MP-BGP**: Multi-protocol BGP, **BGP-LU**: BGP Labeled-Unicast, **RSVP-TE**: Reservation Protocol Traffic Engineering

# L2/L3 Service over SRv6

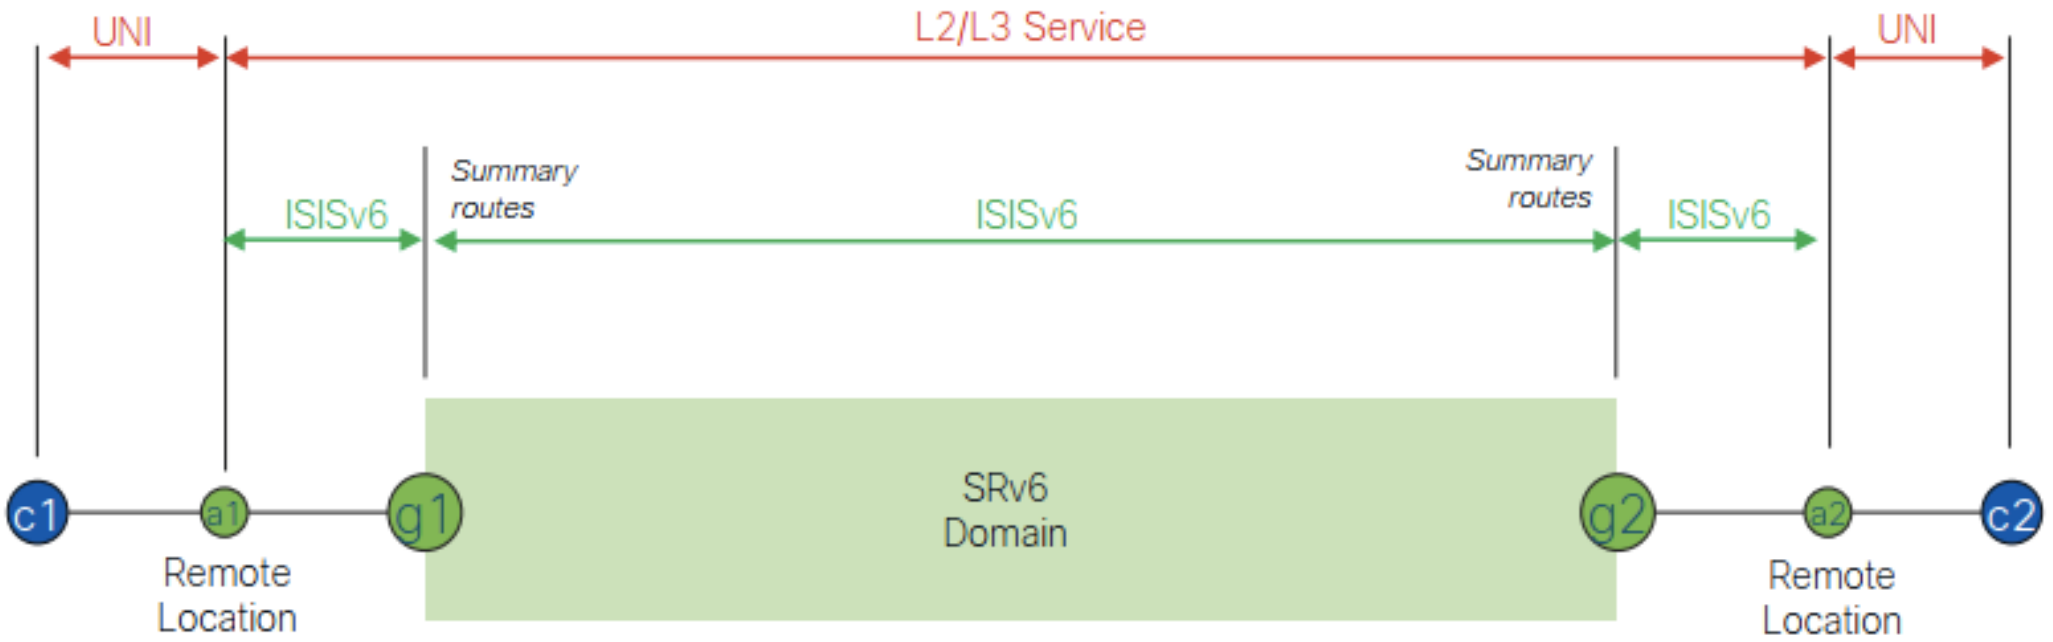


- SRv6 greatly simplifies the design
- Eliminates BGP-LU, RSVP-TE, LDP

Client device/router

SRv6 Node

# Load-Balancing with SRv6

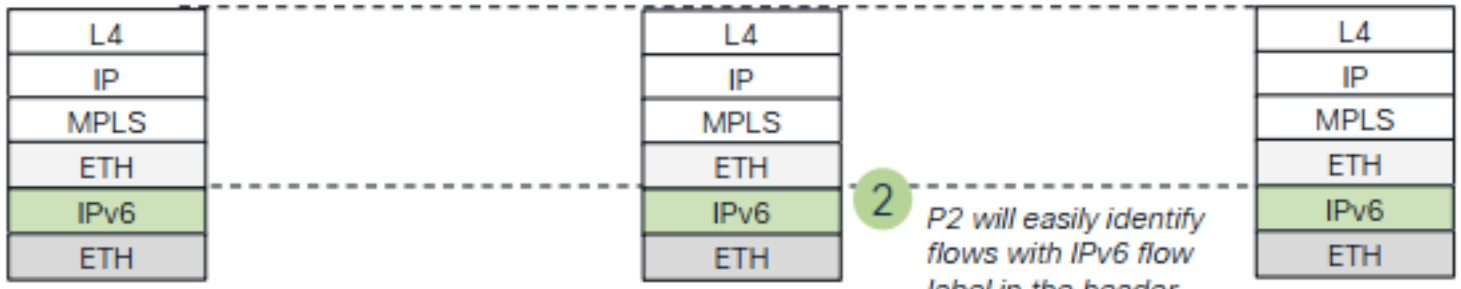


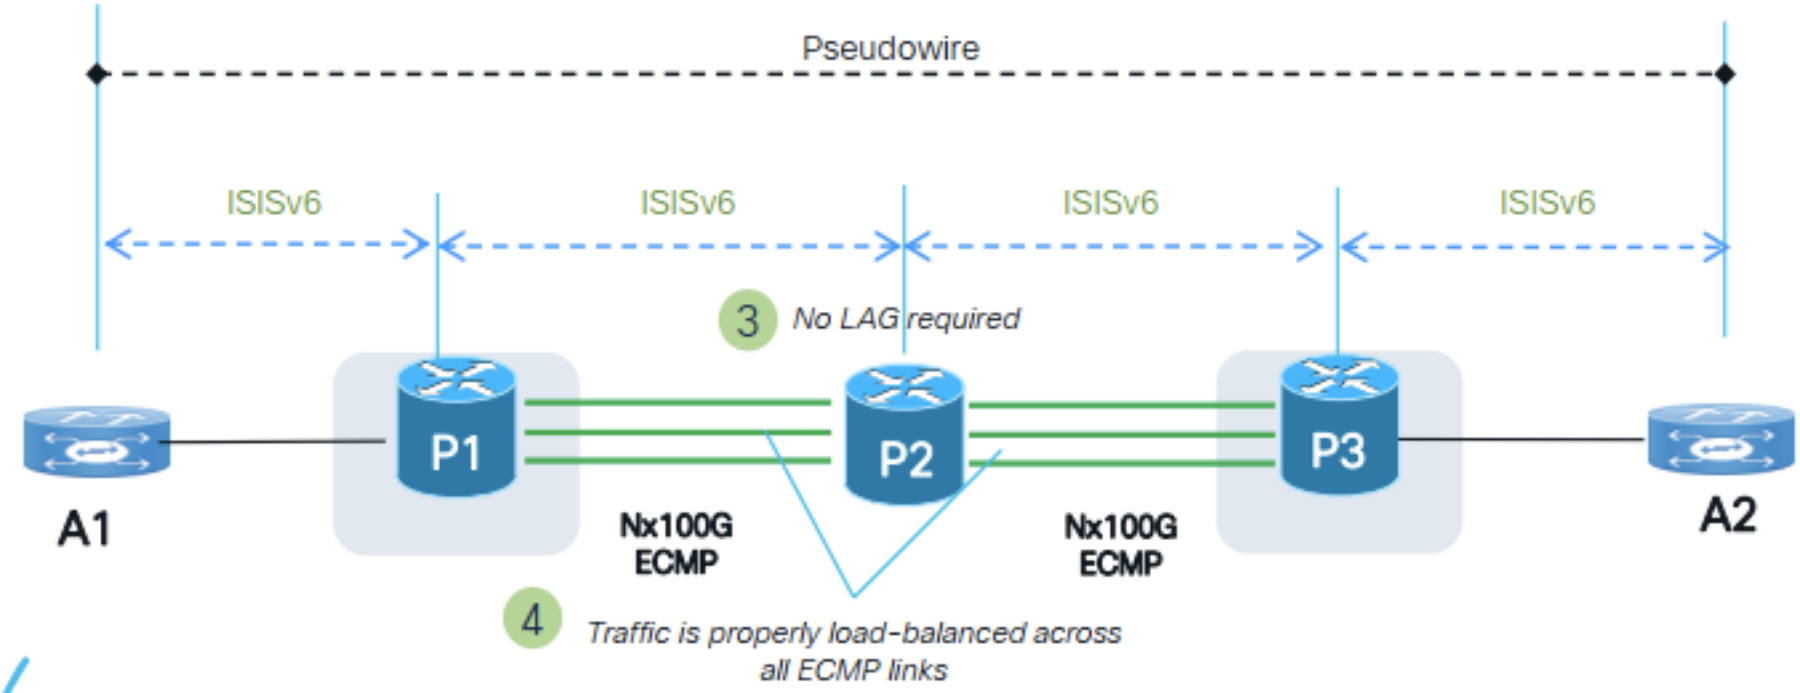

# Inter-Domain Using SRv6

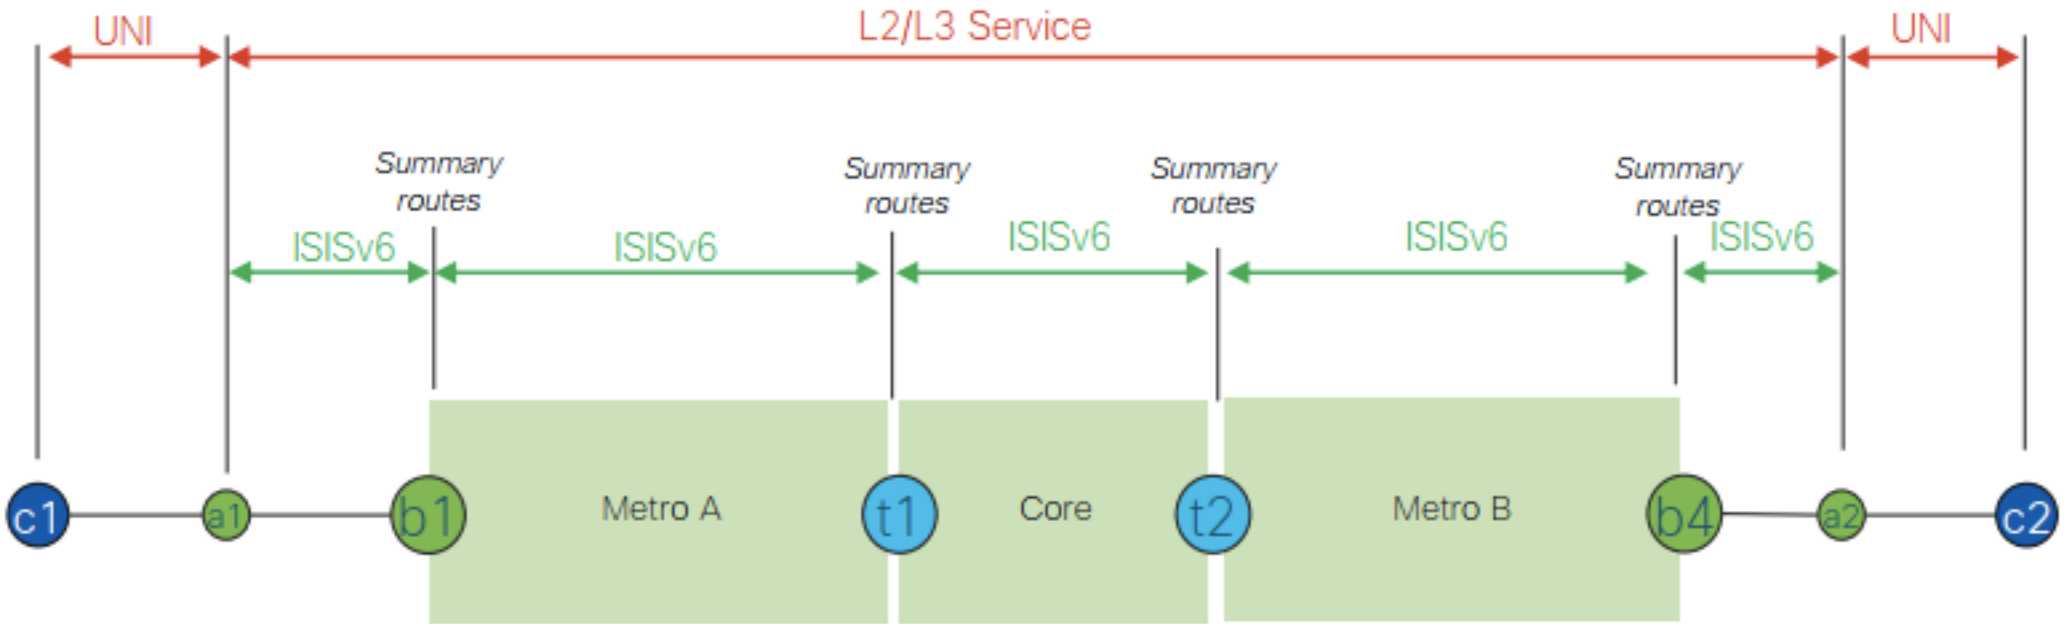


- SRv6 eliminates the BGP-LU shim layer and significantly improves scalability through summarization

Client device/router

SRv6 Node

# SRv6 deployment status

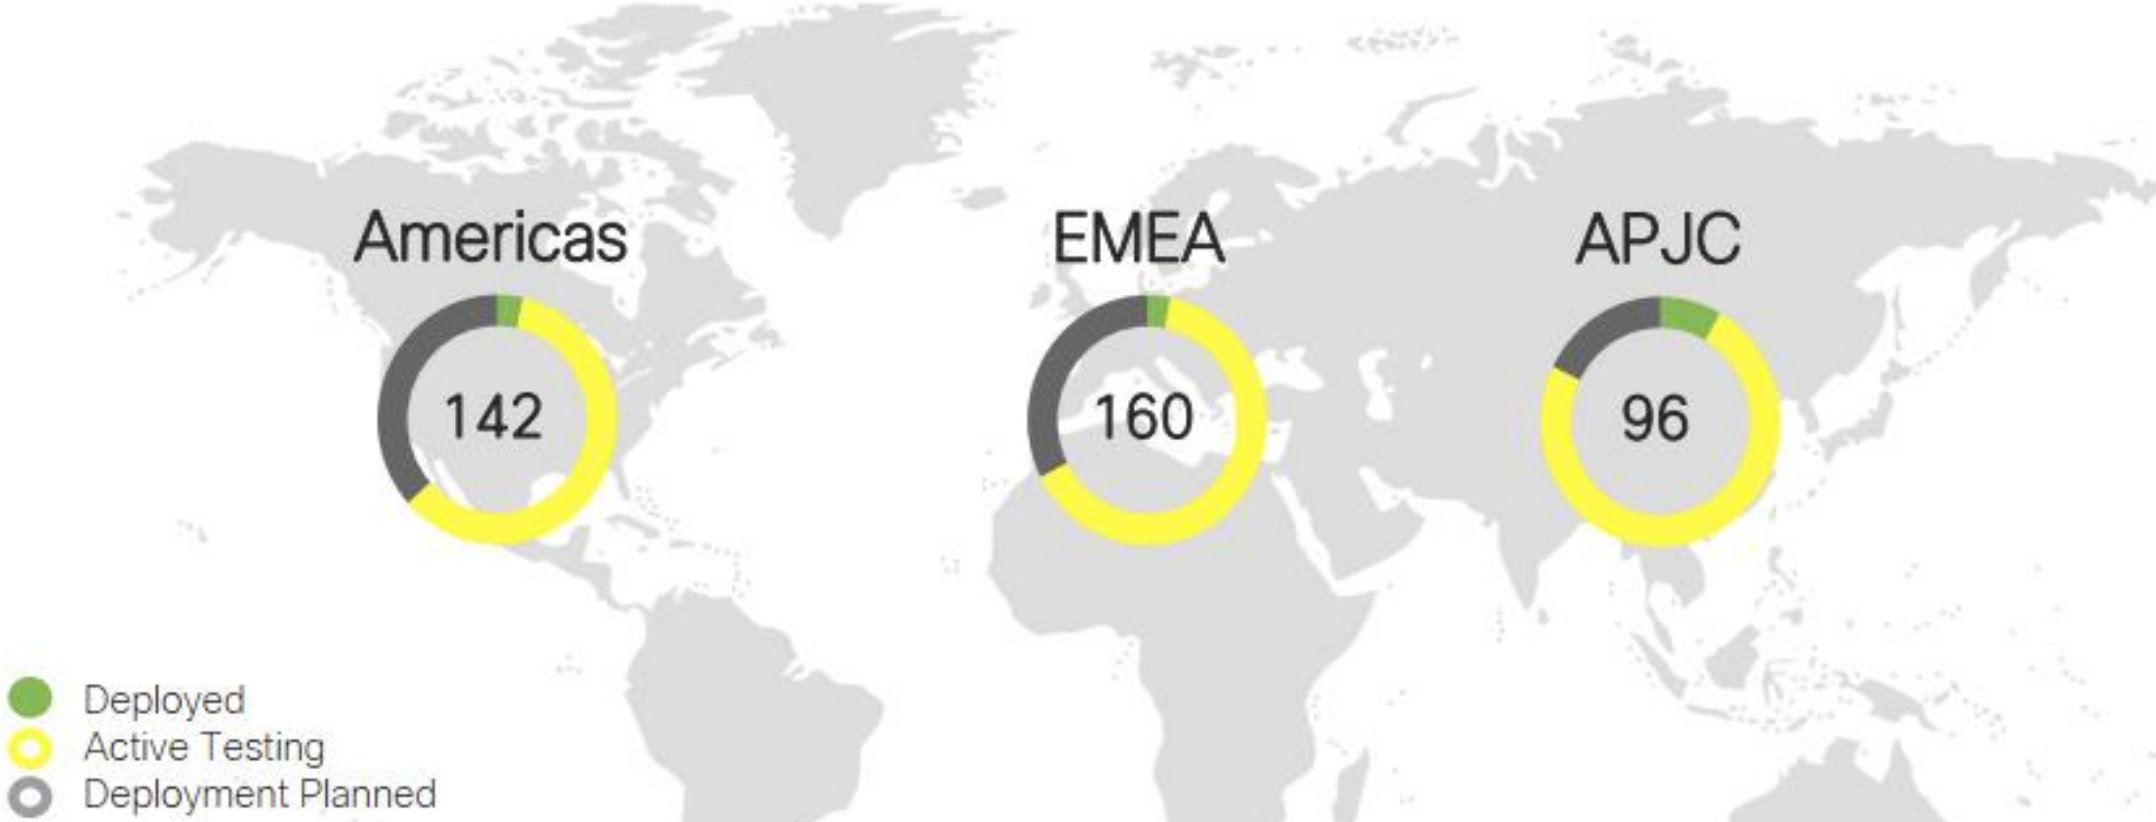


SoftBank free LINE indosat ooredoo Rakuten Bell Telefónica T-Mobile

iliad cesnet swisscom Alibaba.com SCHWARZ XL axiata tpg TELECOM SYMA

# SRv6
# IPv6 Segment Routing Header (SRH)

# SRv6

- IPv6 Header
- Destination IP address

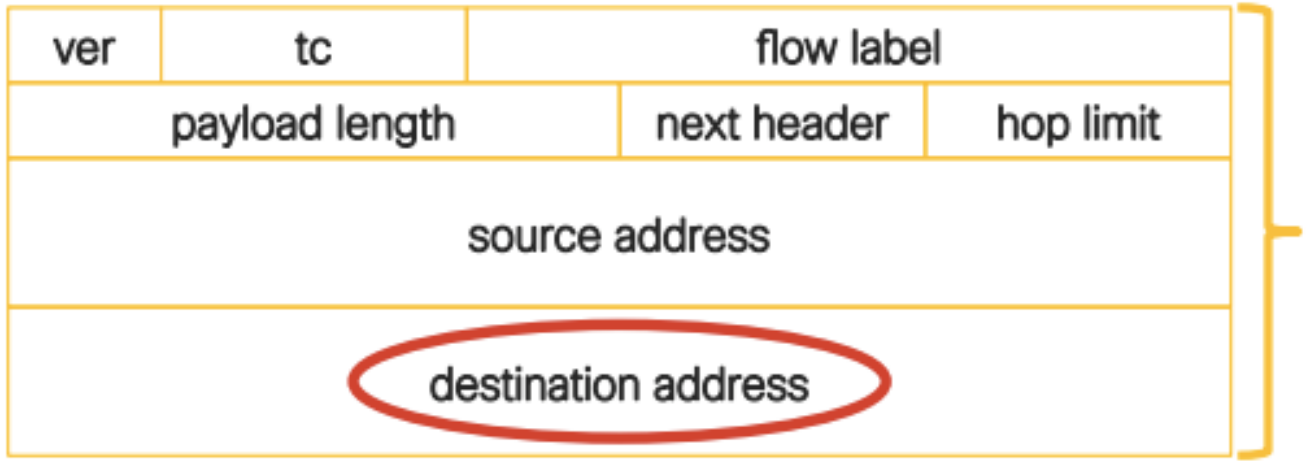

# Network instruction


- 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

# SRv6

- IPv6 Header
- Destination IP address
- Next header field:
  - TCP, UDP, ICMP....

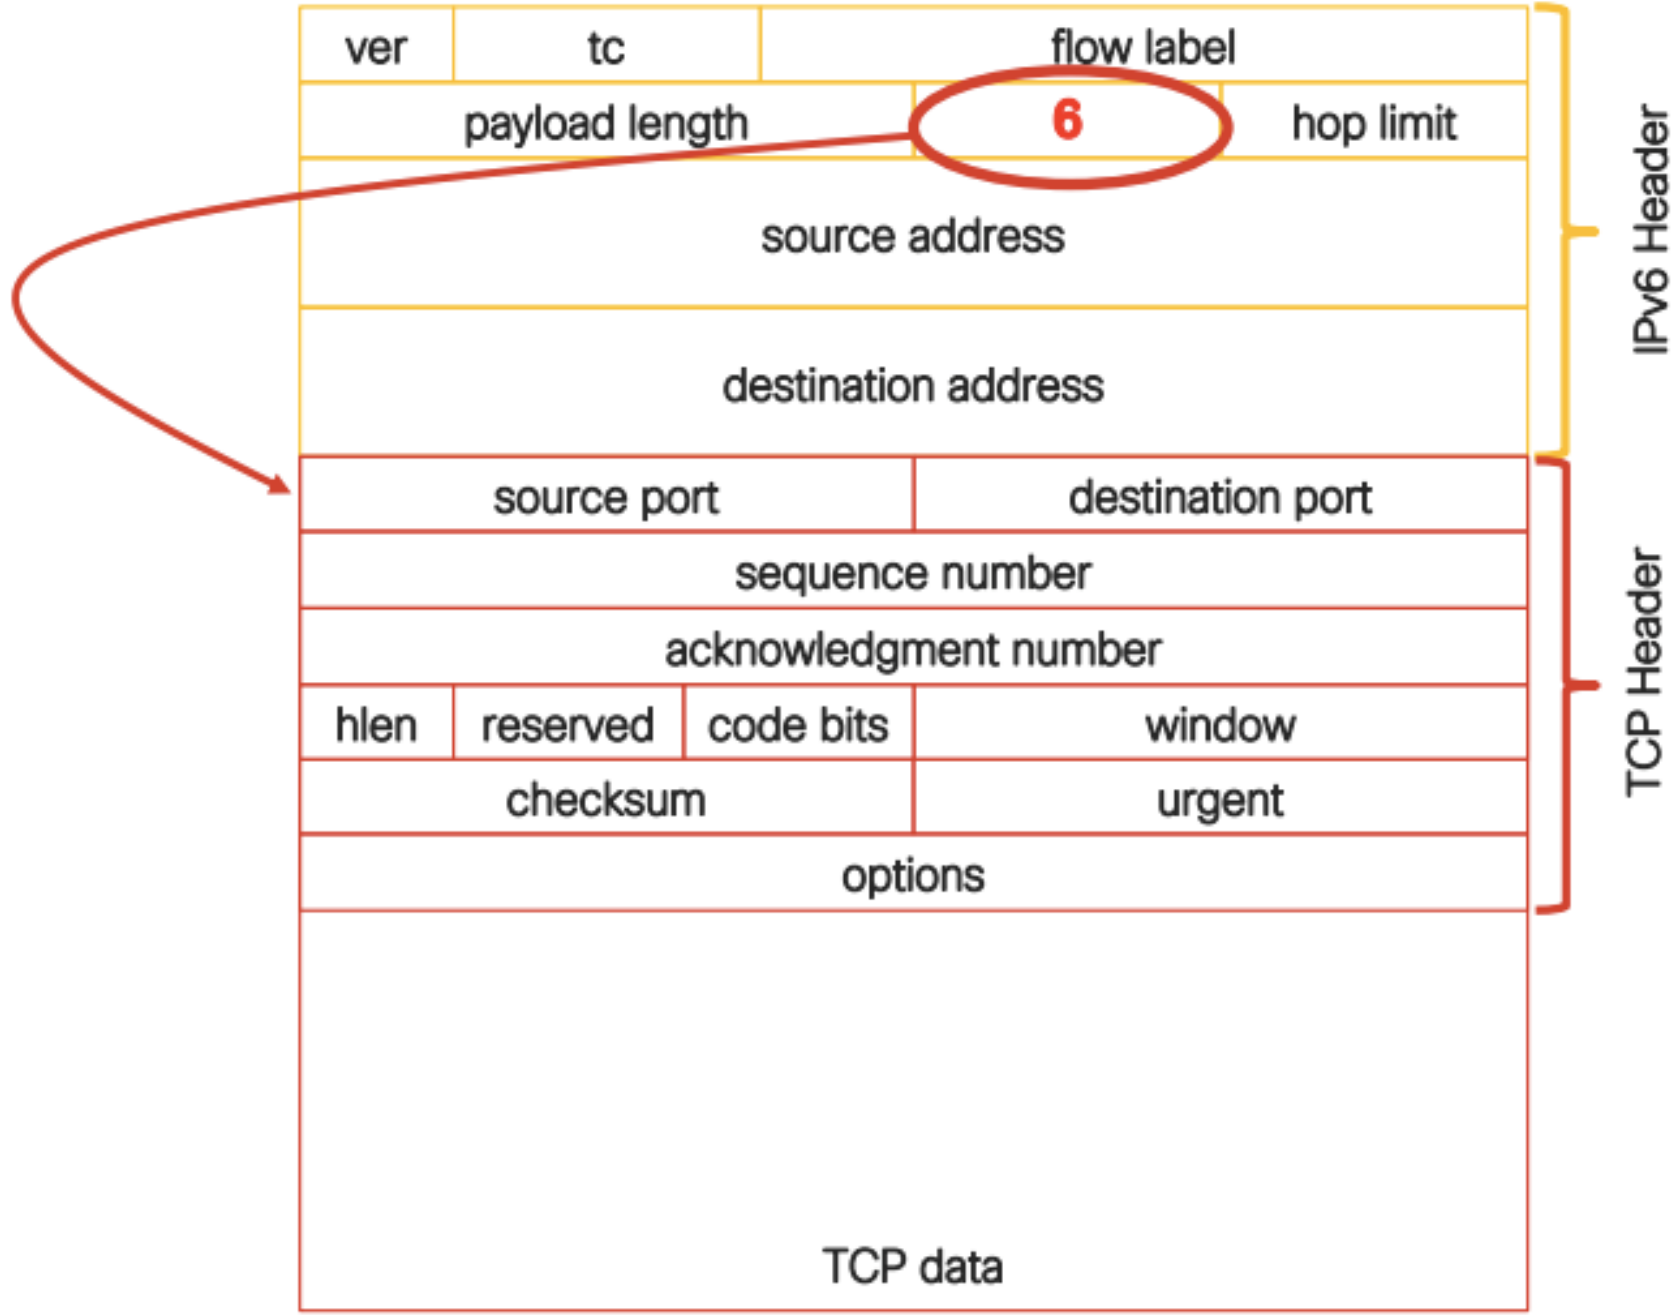

# SRv6

- IPv6 Header
- Destination IP address
- Next header field:
  - TCP, UDP, ICMP....
  - IPv4, IPv6

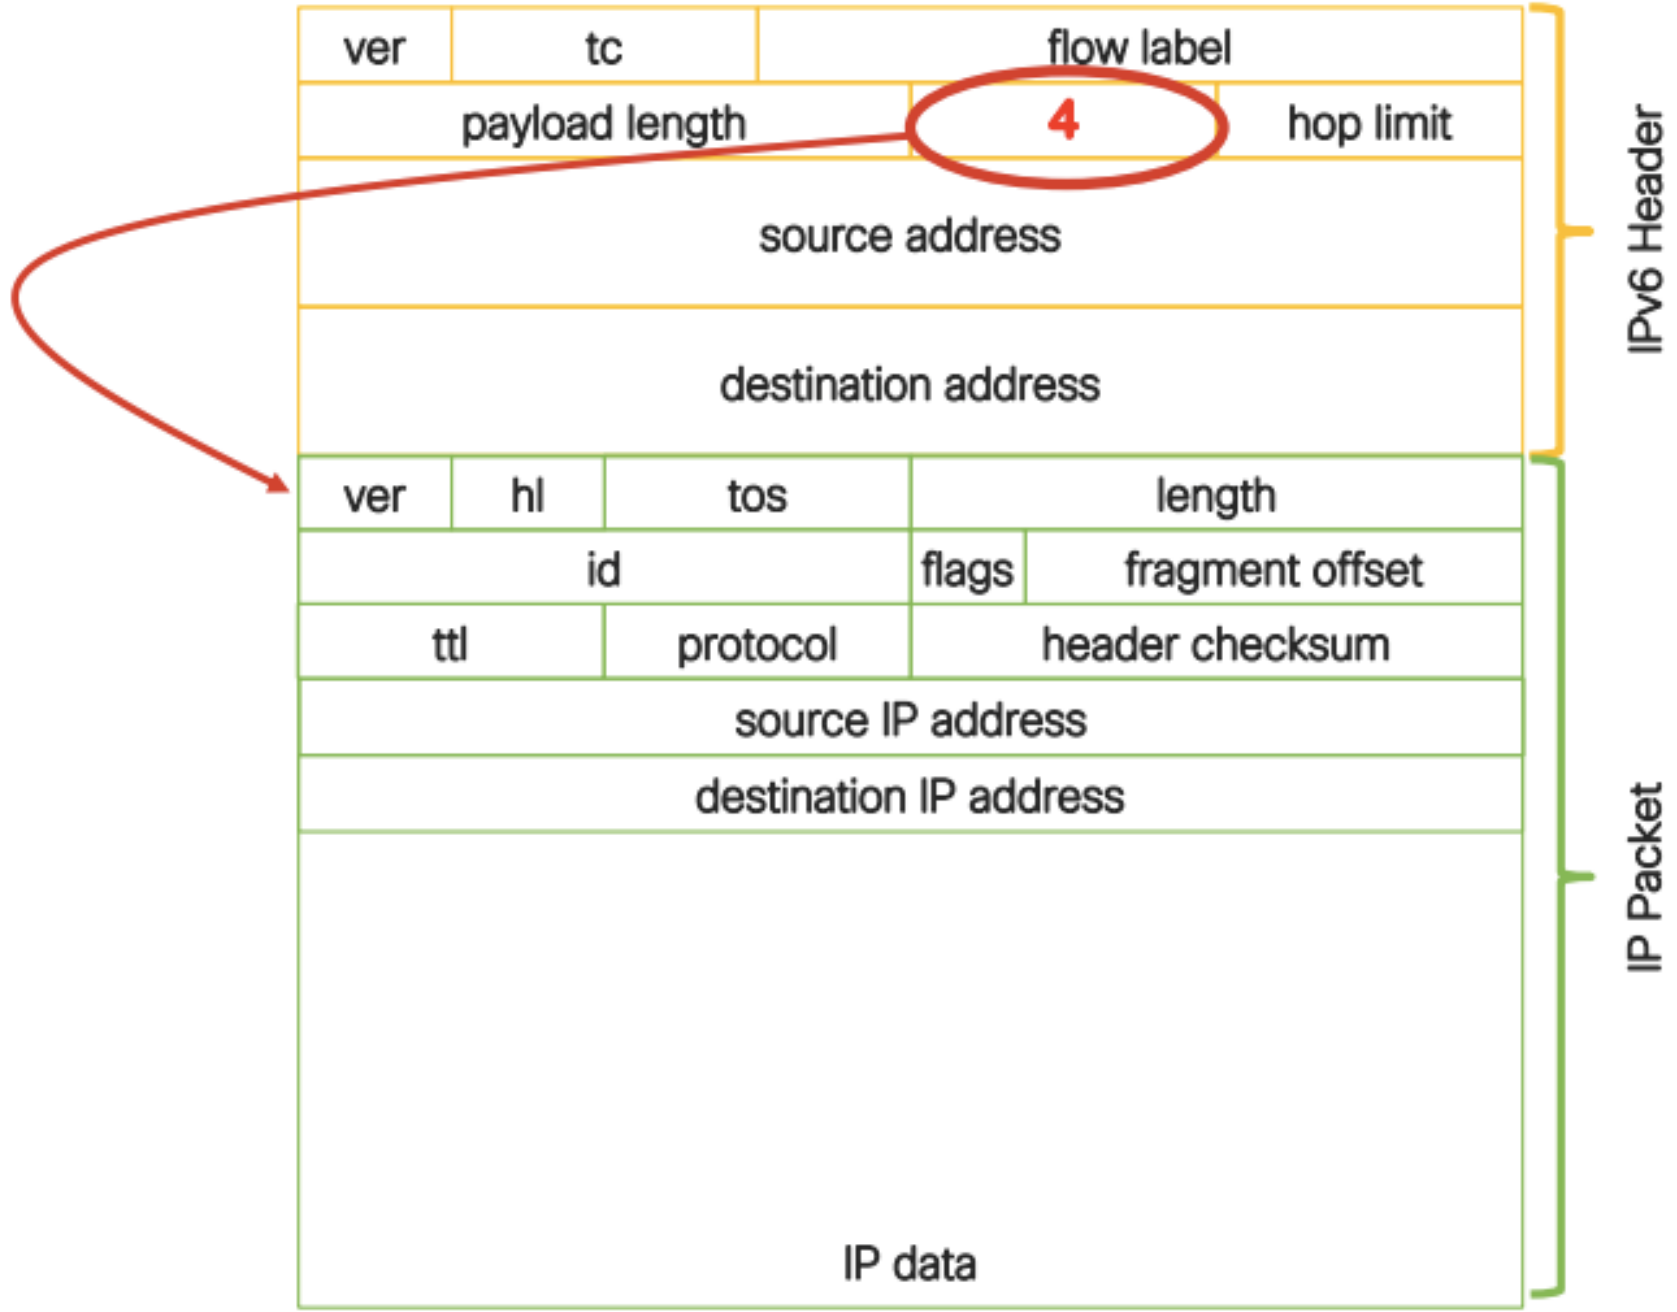

# SRv6

- IPv6 Header
- Destination IP address
- Next header field:
  - TCP, UDP, ICMP....
  - IPv4, IPv6, L2

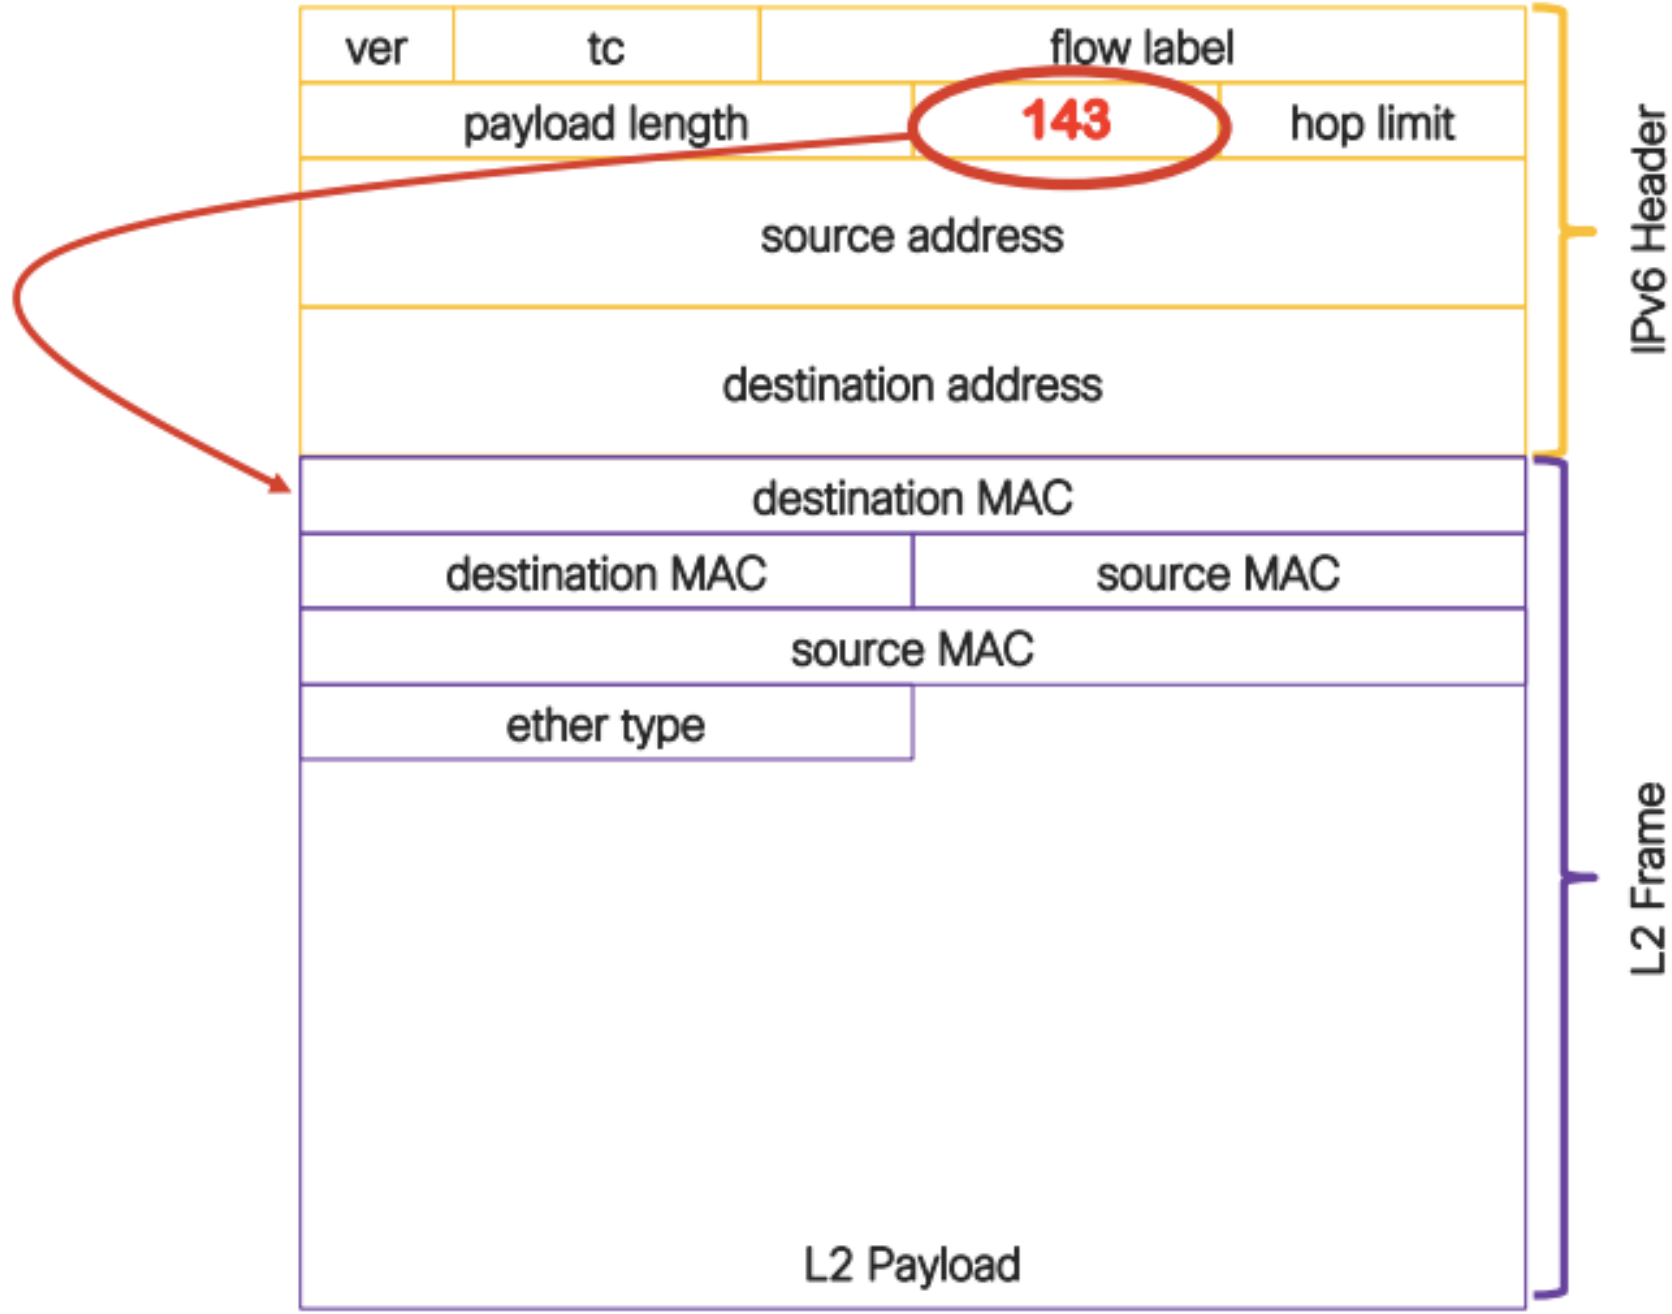

# SRv6

- IPv6 Header
- Destination IP address
- Next header field:
  - TCP, UDP, ICMP....
  - IPv4, IPv6, L2
  - Hop by Hop, Dest. Options, Fragmentation, Authentication Header ...

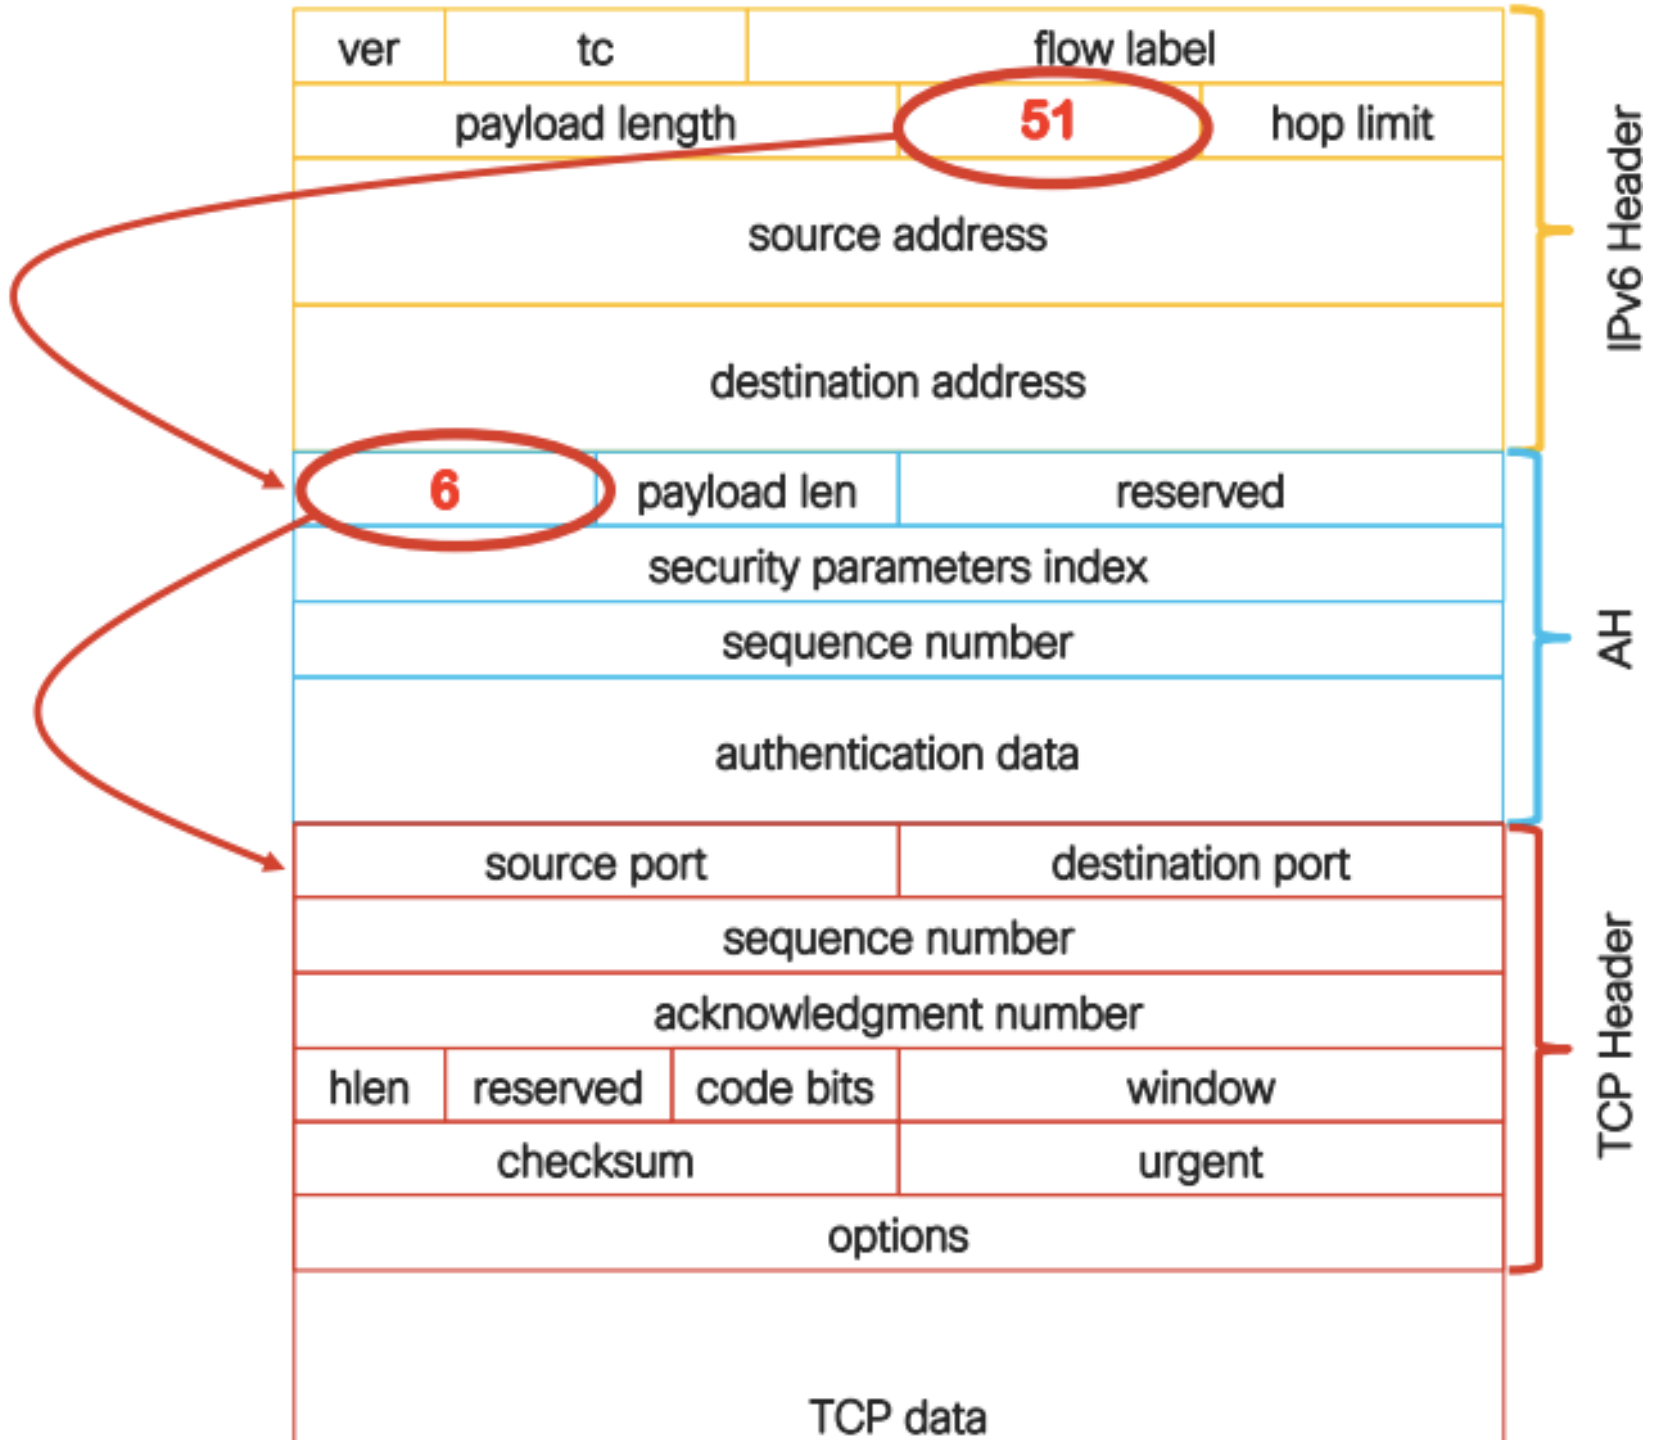

# SRv6

- IPv6 Header
- Destination IP address
- Next header field:
  - TCP, UDP, ICMP....
  - IPv4, IPv6, L2
  - Hop by Hop, Dest. Options, Fragmentation, Authentication Header ...
  - Routing Header

0 Source Route (deprecated)

1 Nimrod (deprecated)

2 Type 2 (RFC 6275)

3 RPL (RFC 6554)

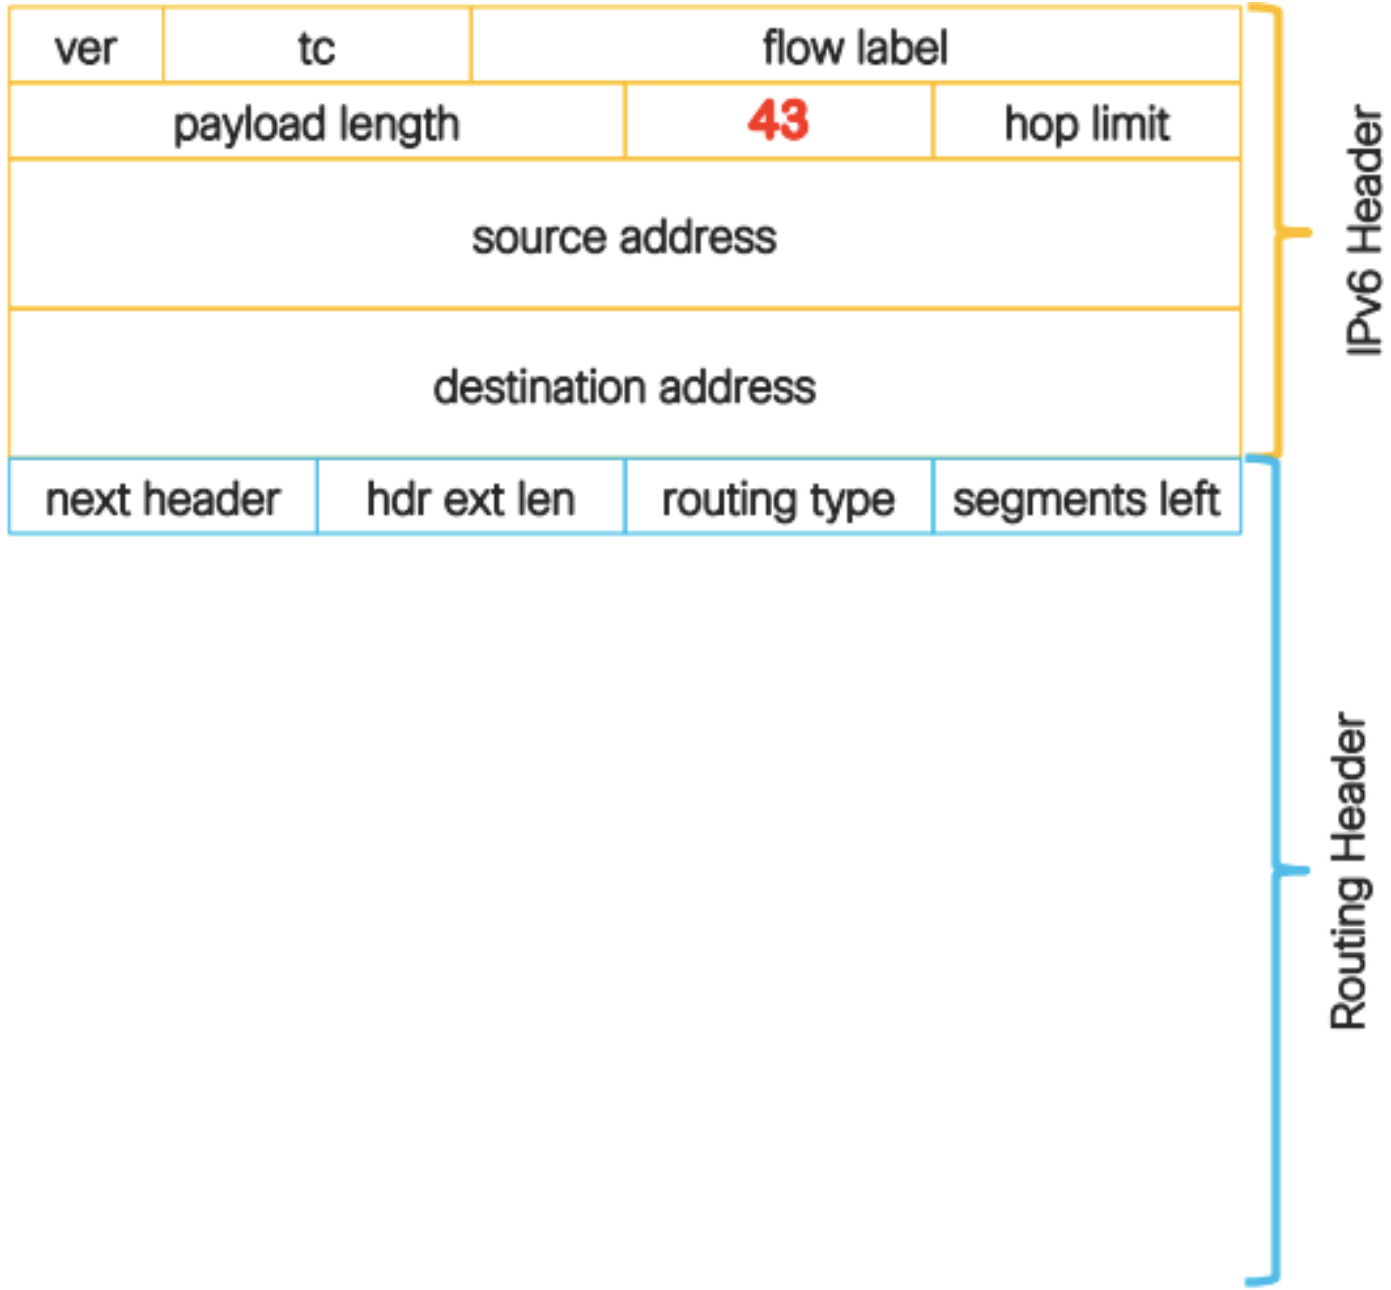

# SRv6

- IPv6 Header
- Destination IP address
- Next header field:
  - TCP, UDP, ICMP....
  - IPv4, IPv6, L2
  - Hop by Hop, Dest. Options, Fragmentation, Authentication Header ...
  - Routing Header

0 Source Route (deprecated)

1 Nimrod (deprecated)

2 Type 2 (RFC 6275)

3 RPL (RFC 6554)

**4 SRH (RFC 8754)**

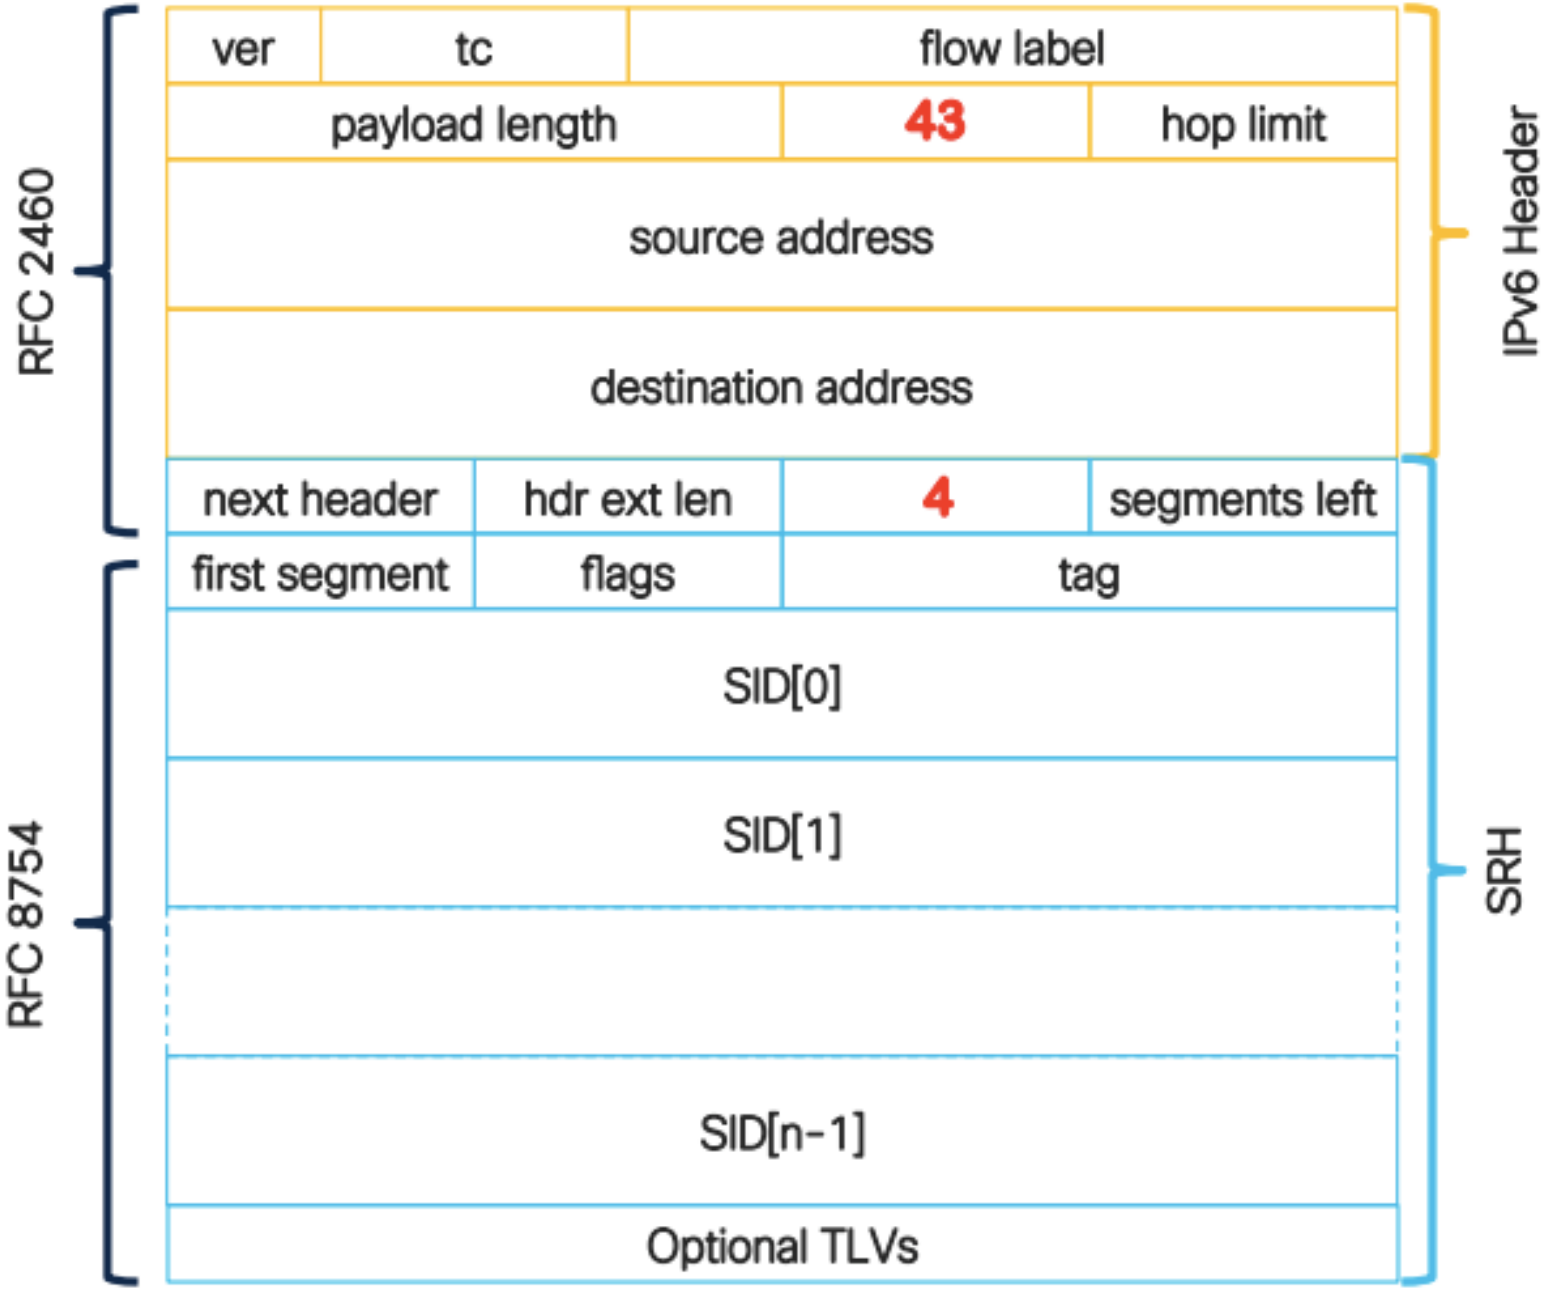

# SRH

- Segment Routing Header
- First Segment
  - Pointer to very first SID

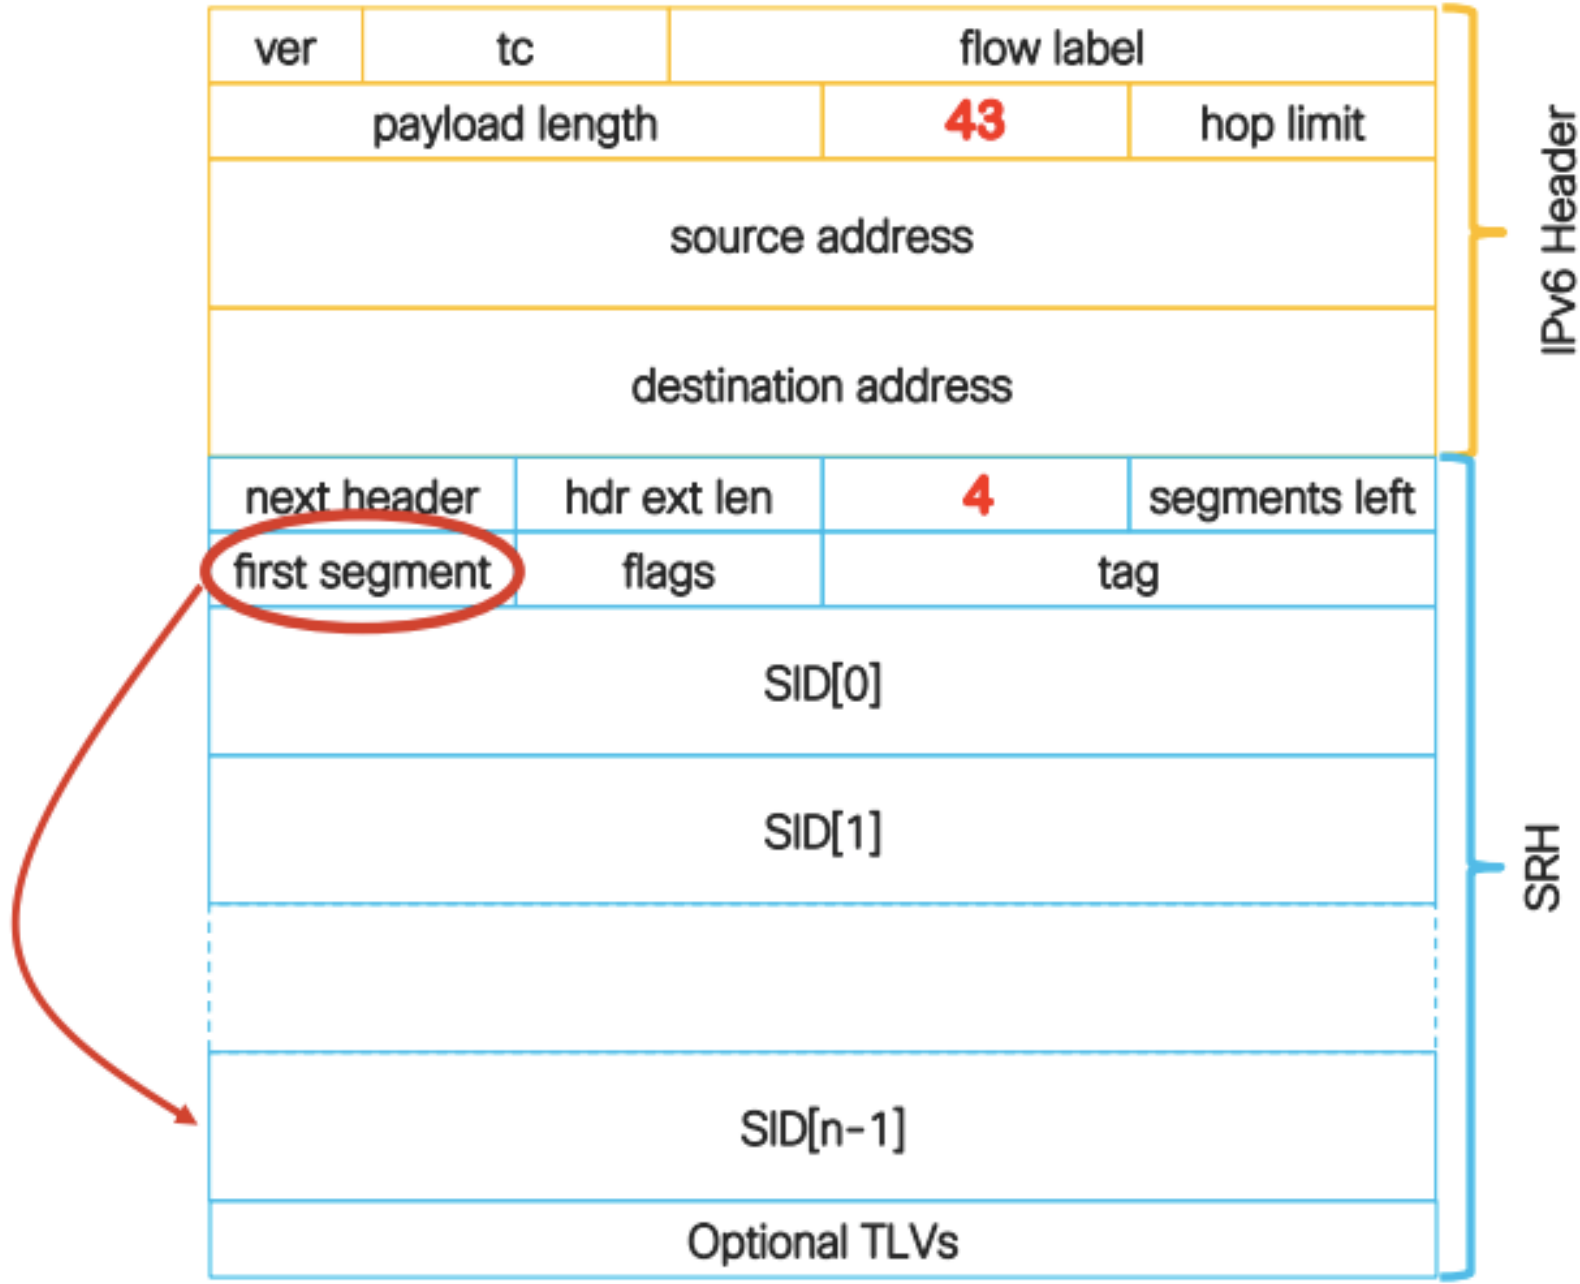

# SRH

- Segment Routing Header
- First Segment
  - Pointer to very first SID
- Segments left
  - Pointer to Active SID
  - Active SID always in destination addr

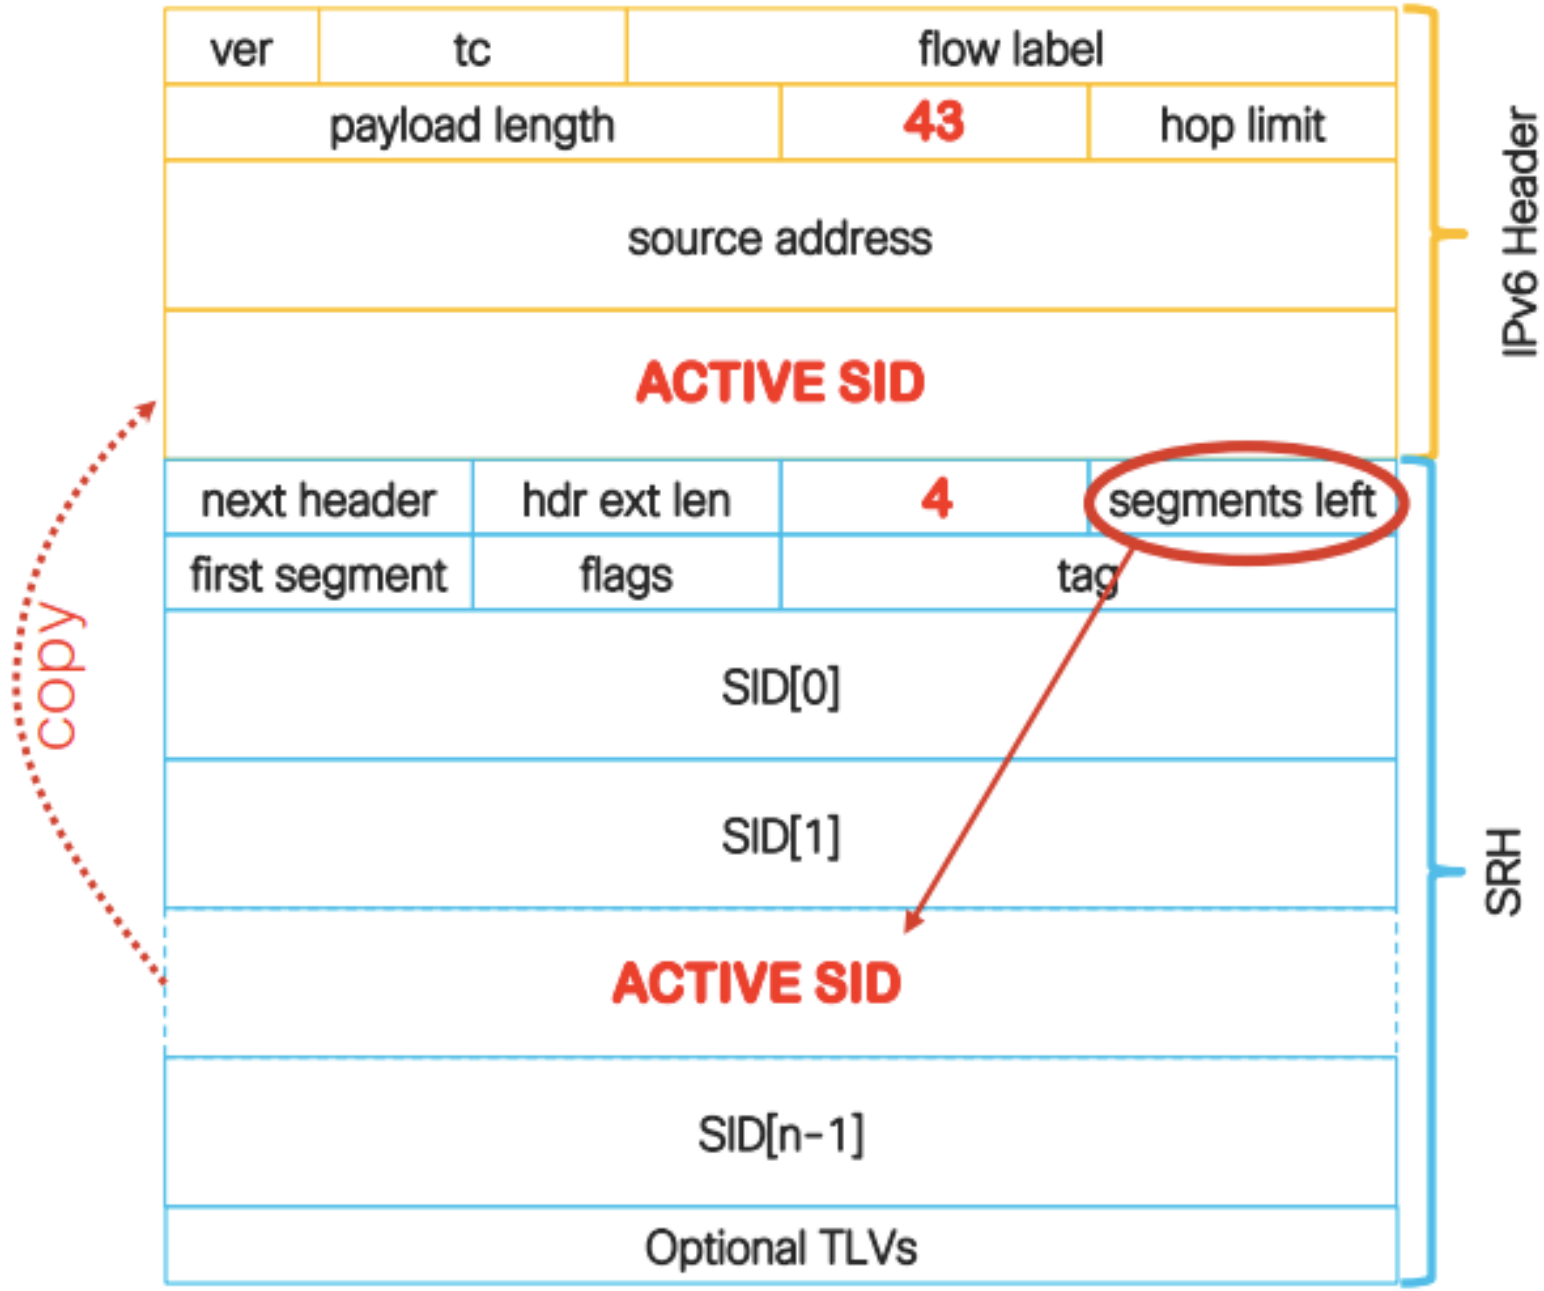

# SID Structure -Locator

128 Bits Like IPv6 address but different semantics

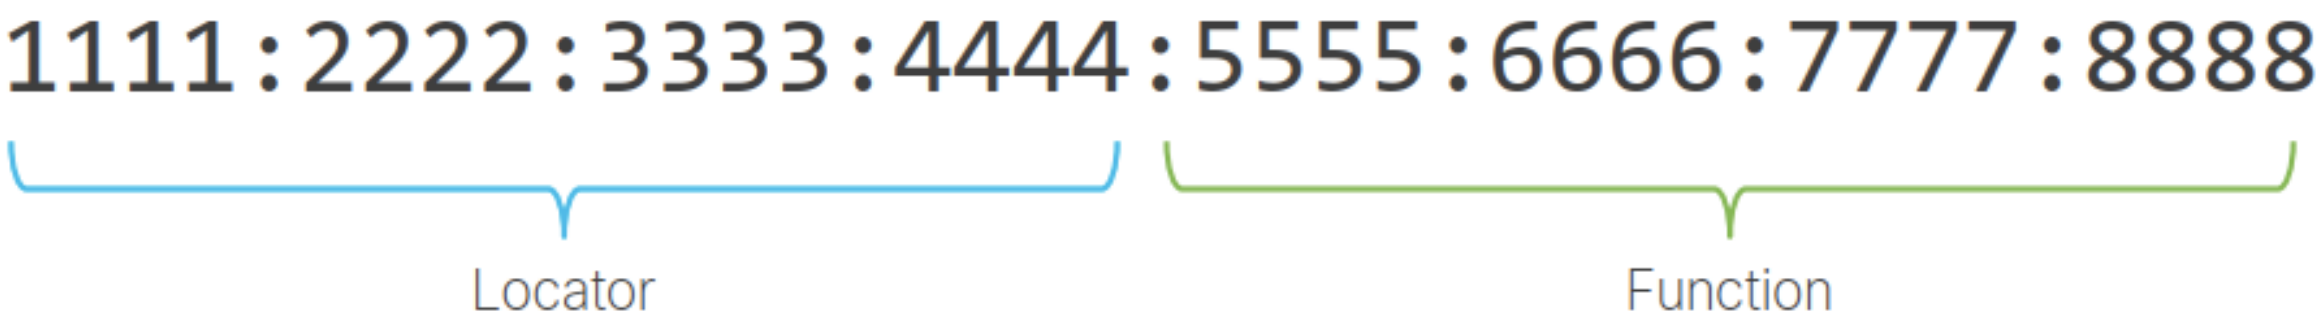

# SRv6 Full SID

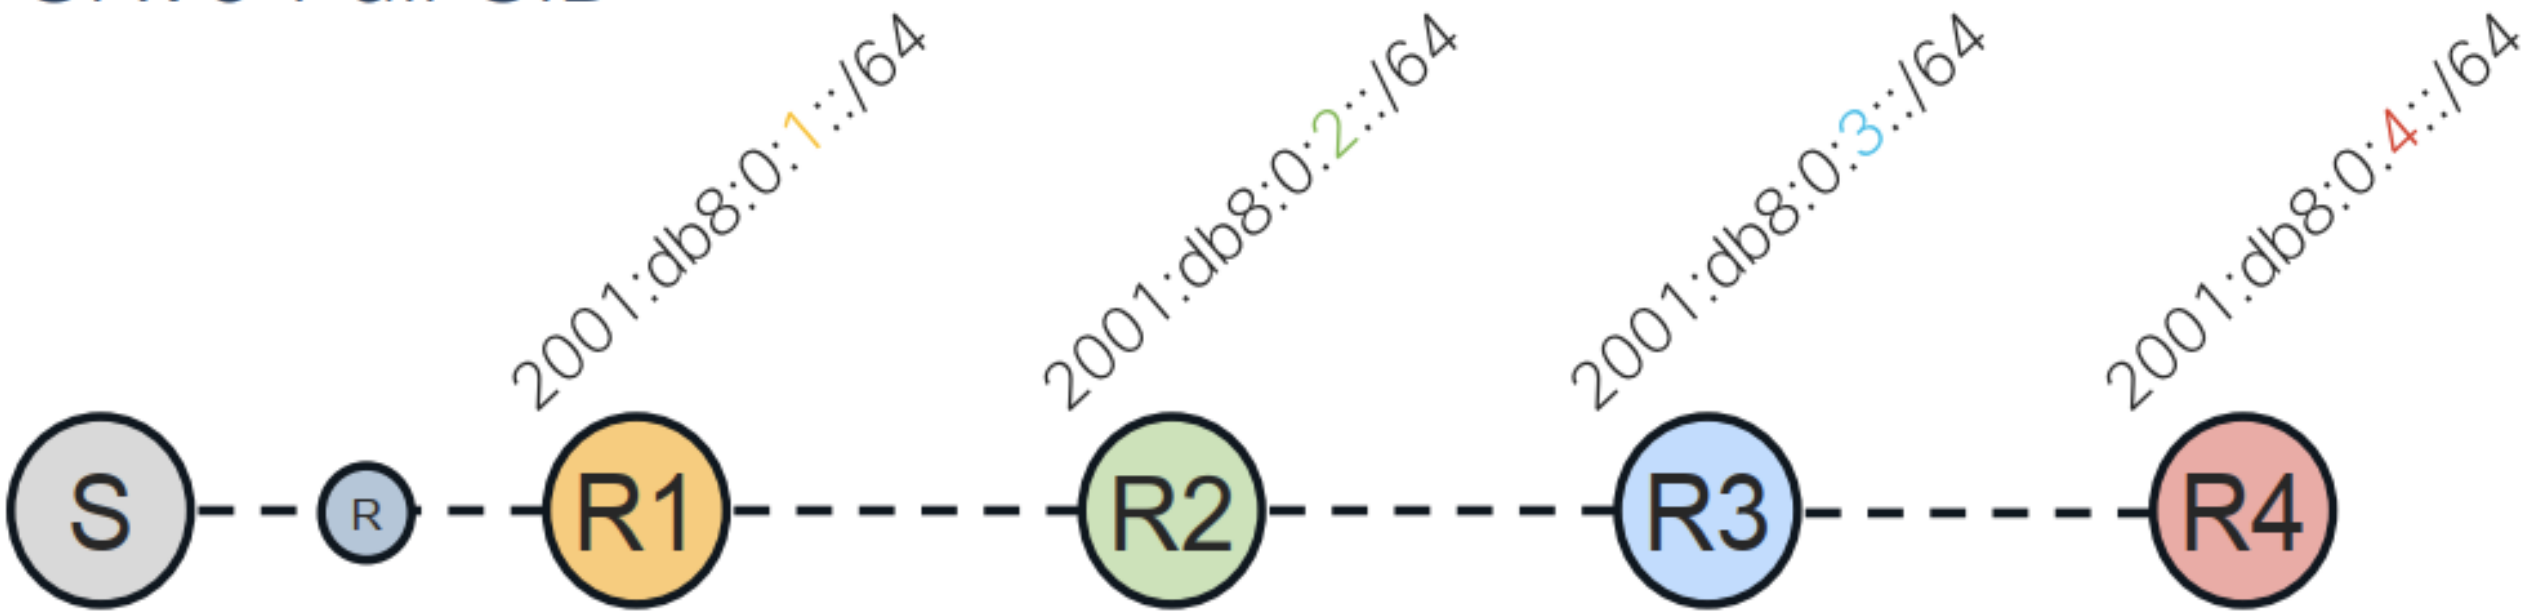


BGP:2001:db8:0:4:eeee::

```
SA:2001::1
DA:2001:db8:0:1:1::
NH:RH
Type:4(SRH)
NH:IPv4|SL:3
Segment List:
[0]:2001:db8:0:4:eeee::
[1]:2001:db8:0:3:48::
[2]:2001:db8:0:2:1::
[3]:2001:db8:0:1:1::
```

```
SA:2001::1
DA:2001:db8:0:2:1::
NH:RH
Type:4(SRH)
NH:IPv4|SL:2
Segment List:
[0]:2001:db8:0:4:eeee::
[1]:2001:db8:0:3:48::
[2]:2001:db8:0:2:1::
[3]:2001:db8:0:1:1::
```

```
SA:2001::1
DA:2001:db8:0:3:48::
NH:RH
Type:4(SRH)
NH:IPv4|SL:1
Segment List:
[0]:2001:db8:0:4:eeee::
[1]:2001:db8:0:3:48::
[2]:2001:db8:0:2:1::
[3]:2001:db8:0:1:1::
```

```
SA:2001::1
DA:2001:db8:0:4:eeee::
NH:IPv4
```

# SID Structure

128 Bits Like IPv6 address but different semantics

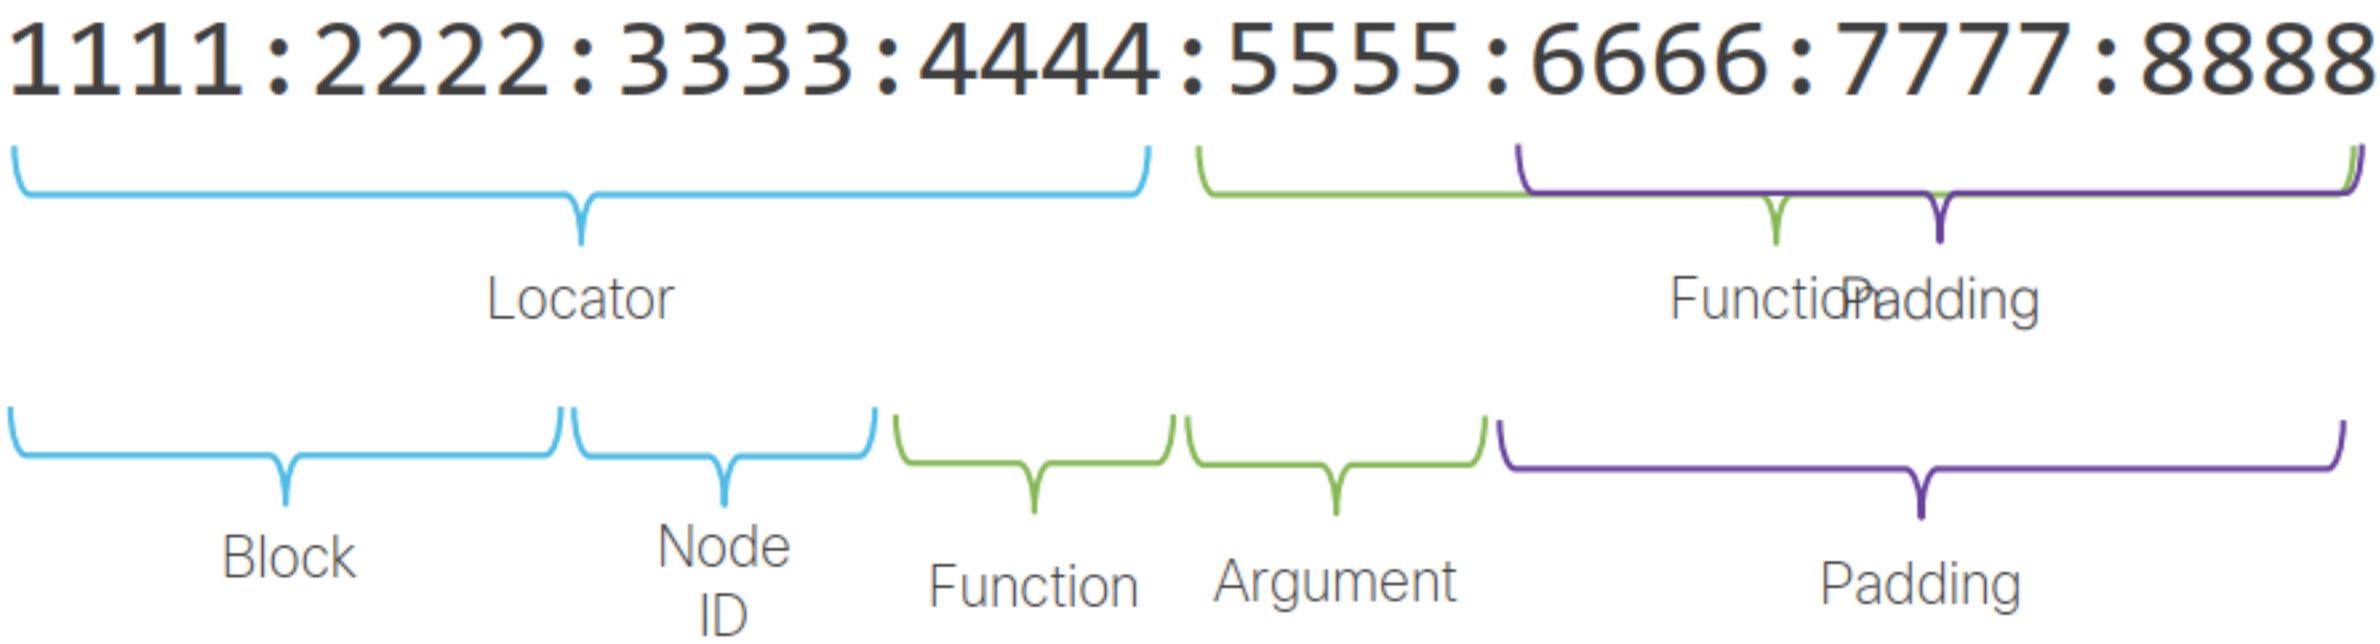

# SRv6 uSID

# SRv6 uSID format

: 0100 : =SRV6 uSID

16 bits here, but can be anything

SRV6 uSID Container

2001 :0db8 : 0100 : 0200 : 0300 : 0400 : 0500 : 0000

SRv6 uSID Block | uSID 1 | uSID 2 | uSID 3 | uSID 4 | uSID 5 | EoC 6

32 bits here,
but can be anything

SRV6 Encapsulation

```
SA:2001::1
DA:2001:db8:0:4:1:0:0:0
NH:RH

Type:4(SRH)
NH:IPv4|SL:1
Segment List:
[0]: 2001:db8:0:5:45:0:0:0
[1]: 2001:db8:0:4:1:0:0:0
[2]: 2001:db8:0:3:48:0:0:0
[3]: 2001:db8:0:2:1:0:0:0
[4]: 2001:db8:0:1:42:0:0:0

SA:7.5.4.3
DA:11.6.19.71
Port:UDP

UDP Header/Data
```

SRV6 uSID Encapsulation

```
SA:2001::1
DA:2001:db8:100:200:300:400:500::
NH:IPv4

SA:7.5.4.3
DA:11.6.19.71
Port:UDP

UDP Header/Data
```

# SRv6 uSID F3216

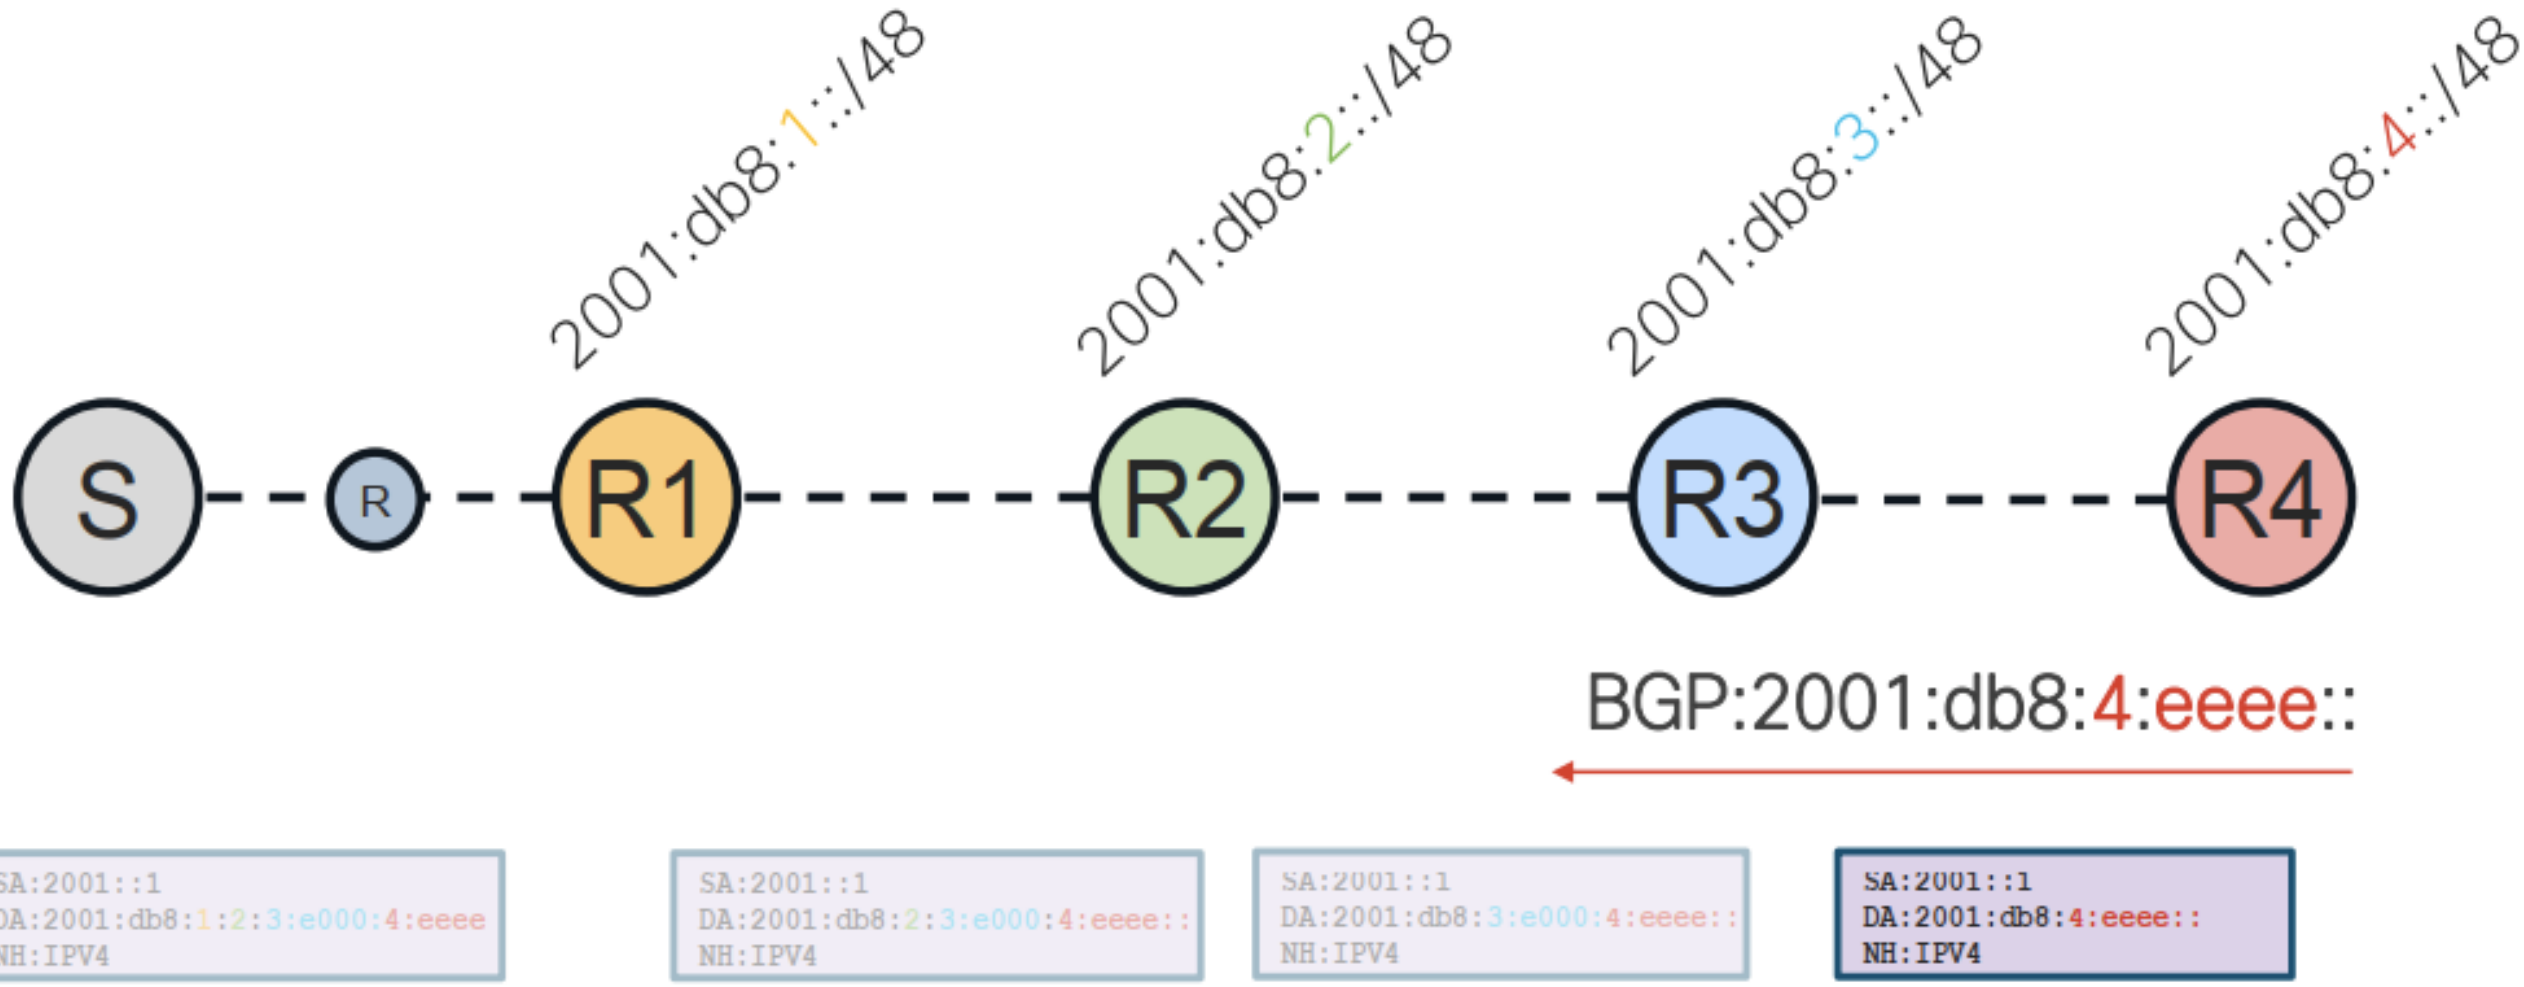

# SRv6 uSID More Than 6 SIDs?

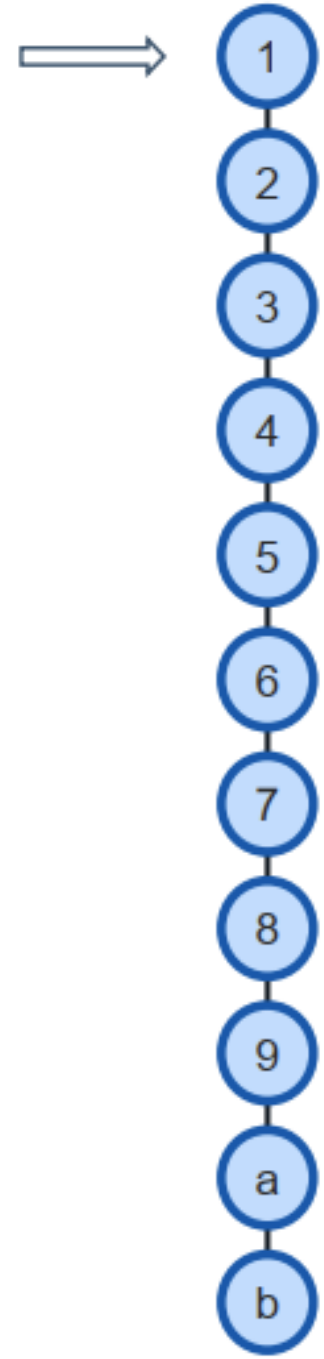


100->200->300->400->500->600->700->800->900->a00->b00

Carrier 1 2001 : 0db8 : 0100 : 0200 : 0300 : 0400 : 0500 : 0600

Carrier 2 2001 : 0db8 : 0700 : 0800 : 0900 : 0a00 : 0b00 : 0000

```
SA:2001::1
DA:2001:db8:[illegible]
NH:[illegible]
```

```
Type:4(SRH)
NH:IPv4|SL:[illegible]
Segment List:
[0]: 2001:db8:700:800:900:a00:b00::
```

```
SA:7.5.4.3
DA:11.6.19.71
Port:UDP
```

```
UDP Header/Data
```

Shift & Forward
END of Carrier
-> is there SRH?
Decrement SL
Copy New SID (Carrier)
PSP

# SRv6 uSID Configuration

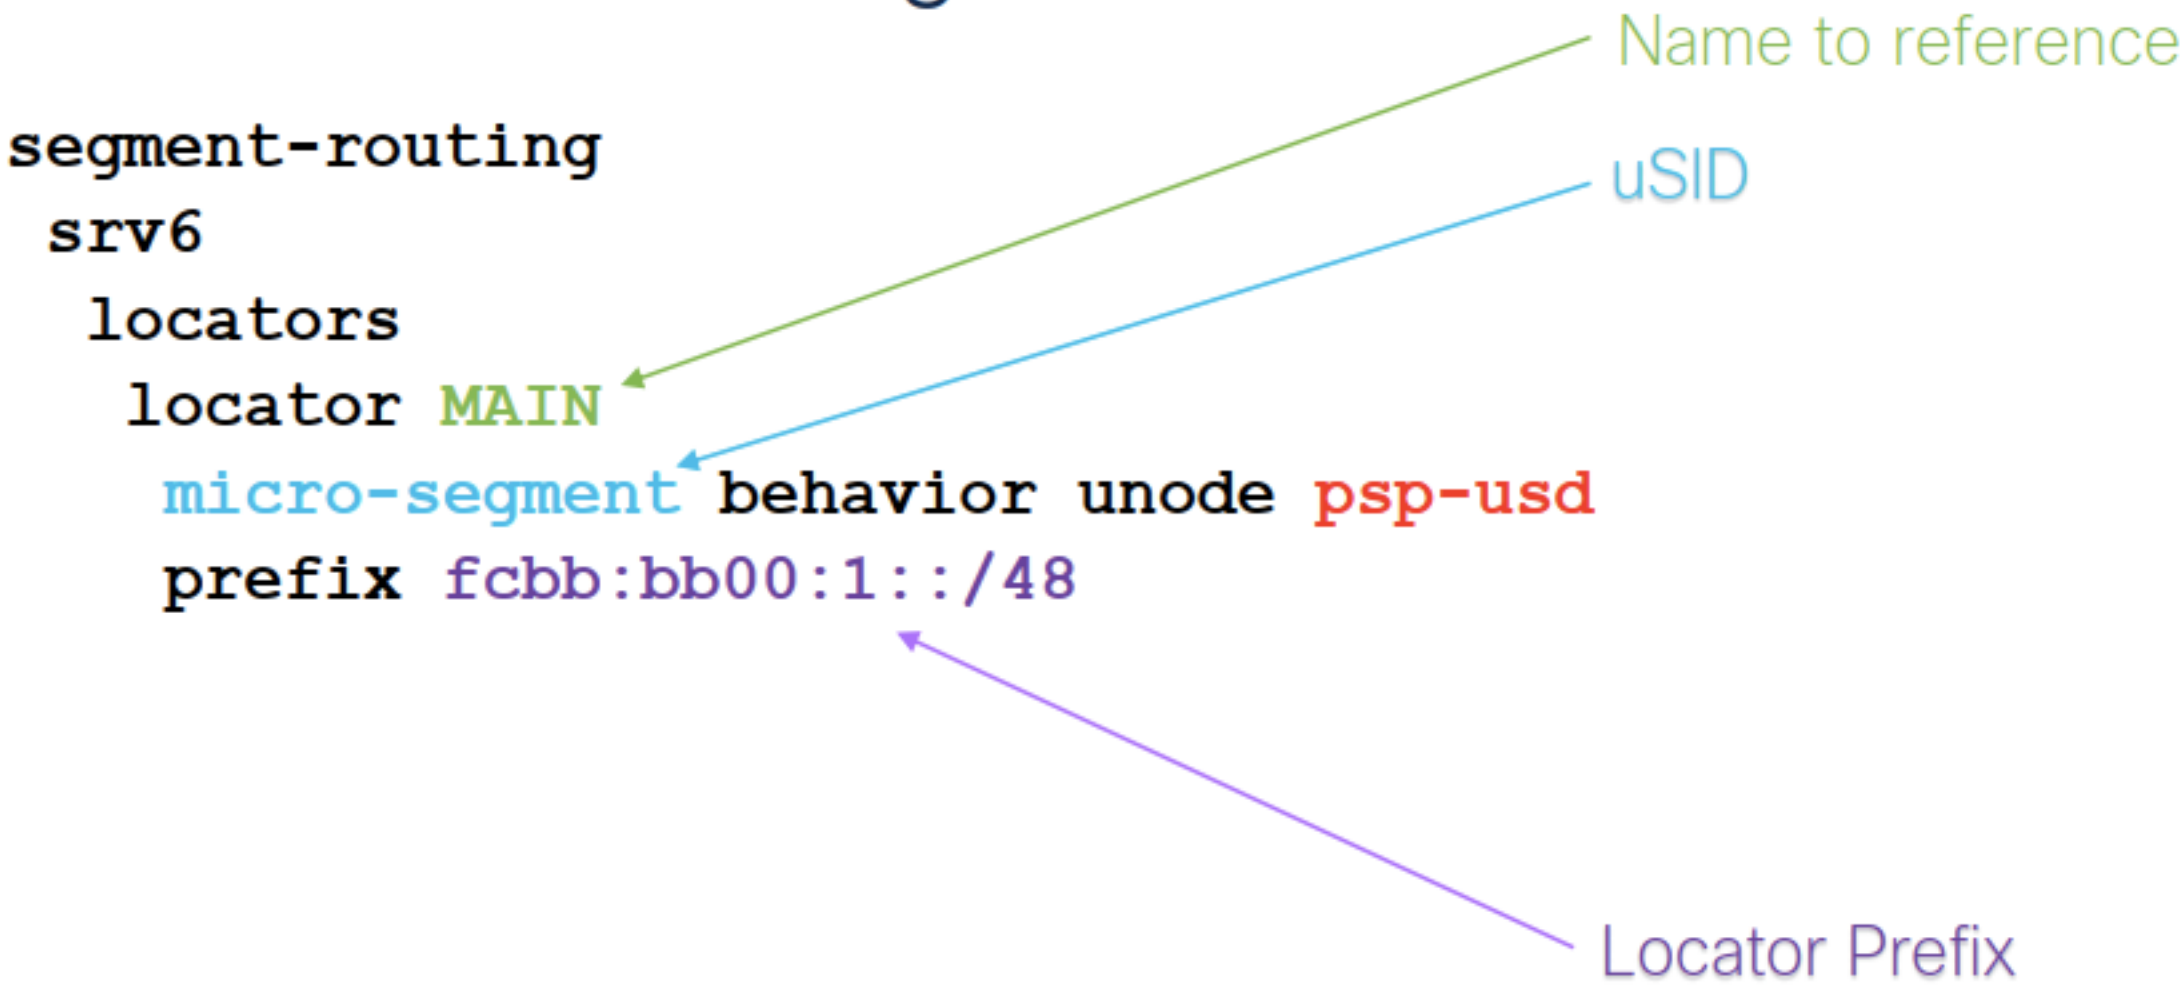

# ISIS Extensions

# SRv6 ISIS Configuration

```
router isis 1
 address-family ipv6 unicast
  segment-routing srv6
   locator MAIN
```

Name of the Locator

This will result in:

- Locator is advertised
- uN function is advertised
- uA for each ISIS interface is allocated and advertised

# BGP

- No Changes Required!

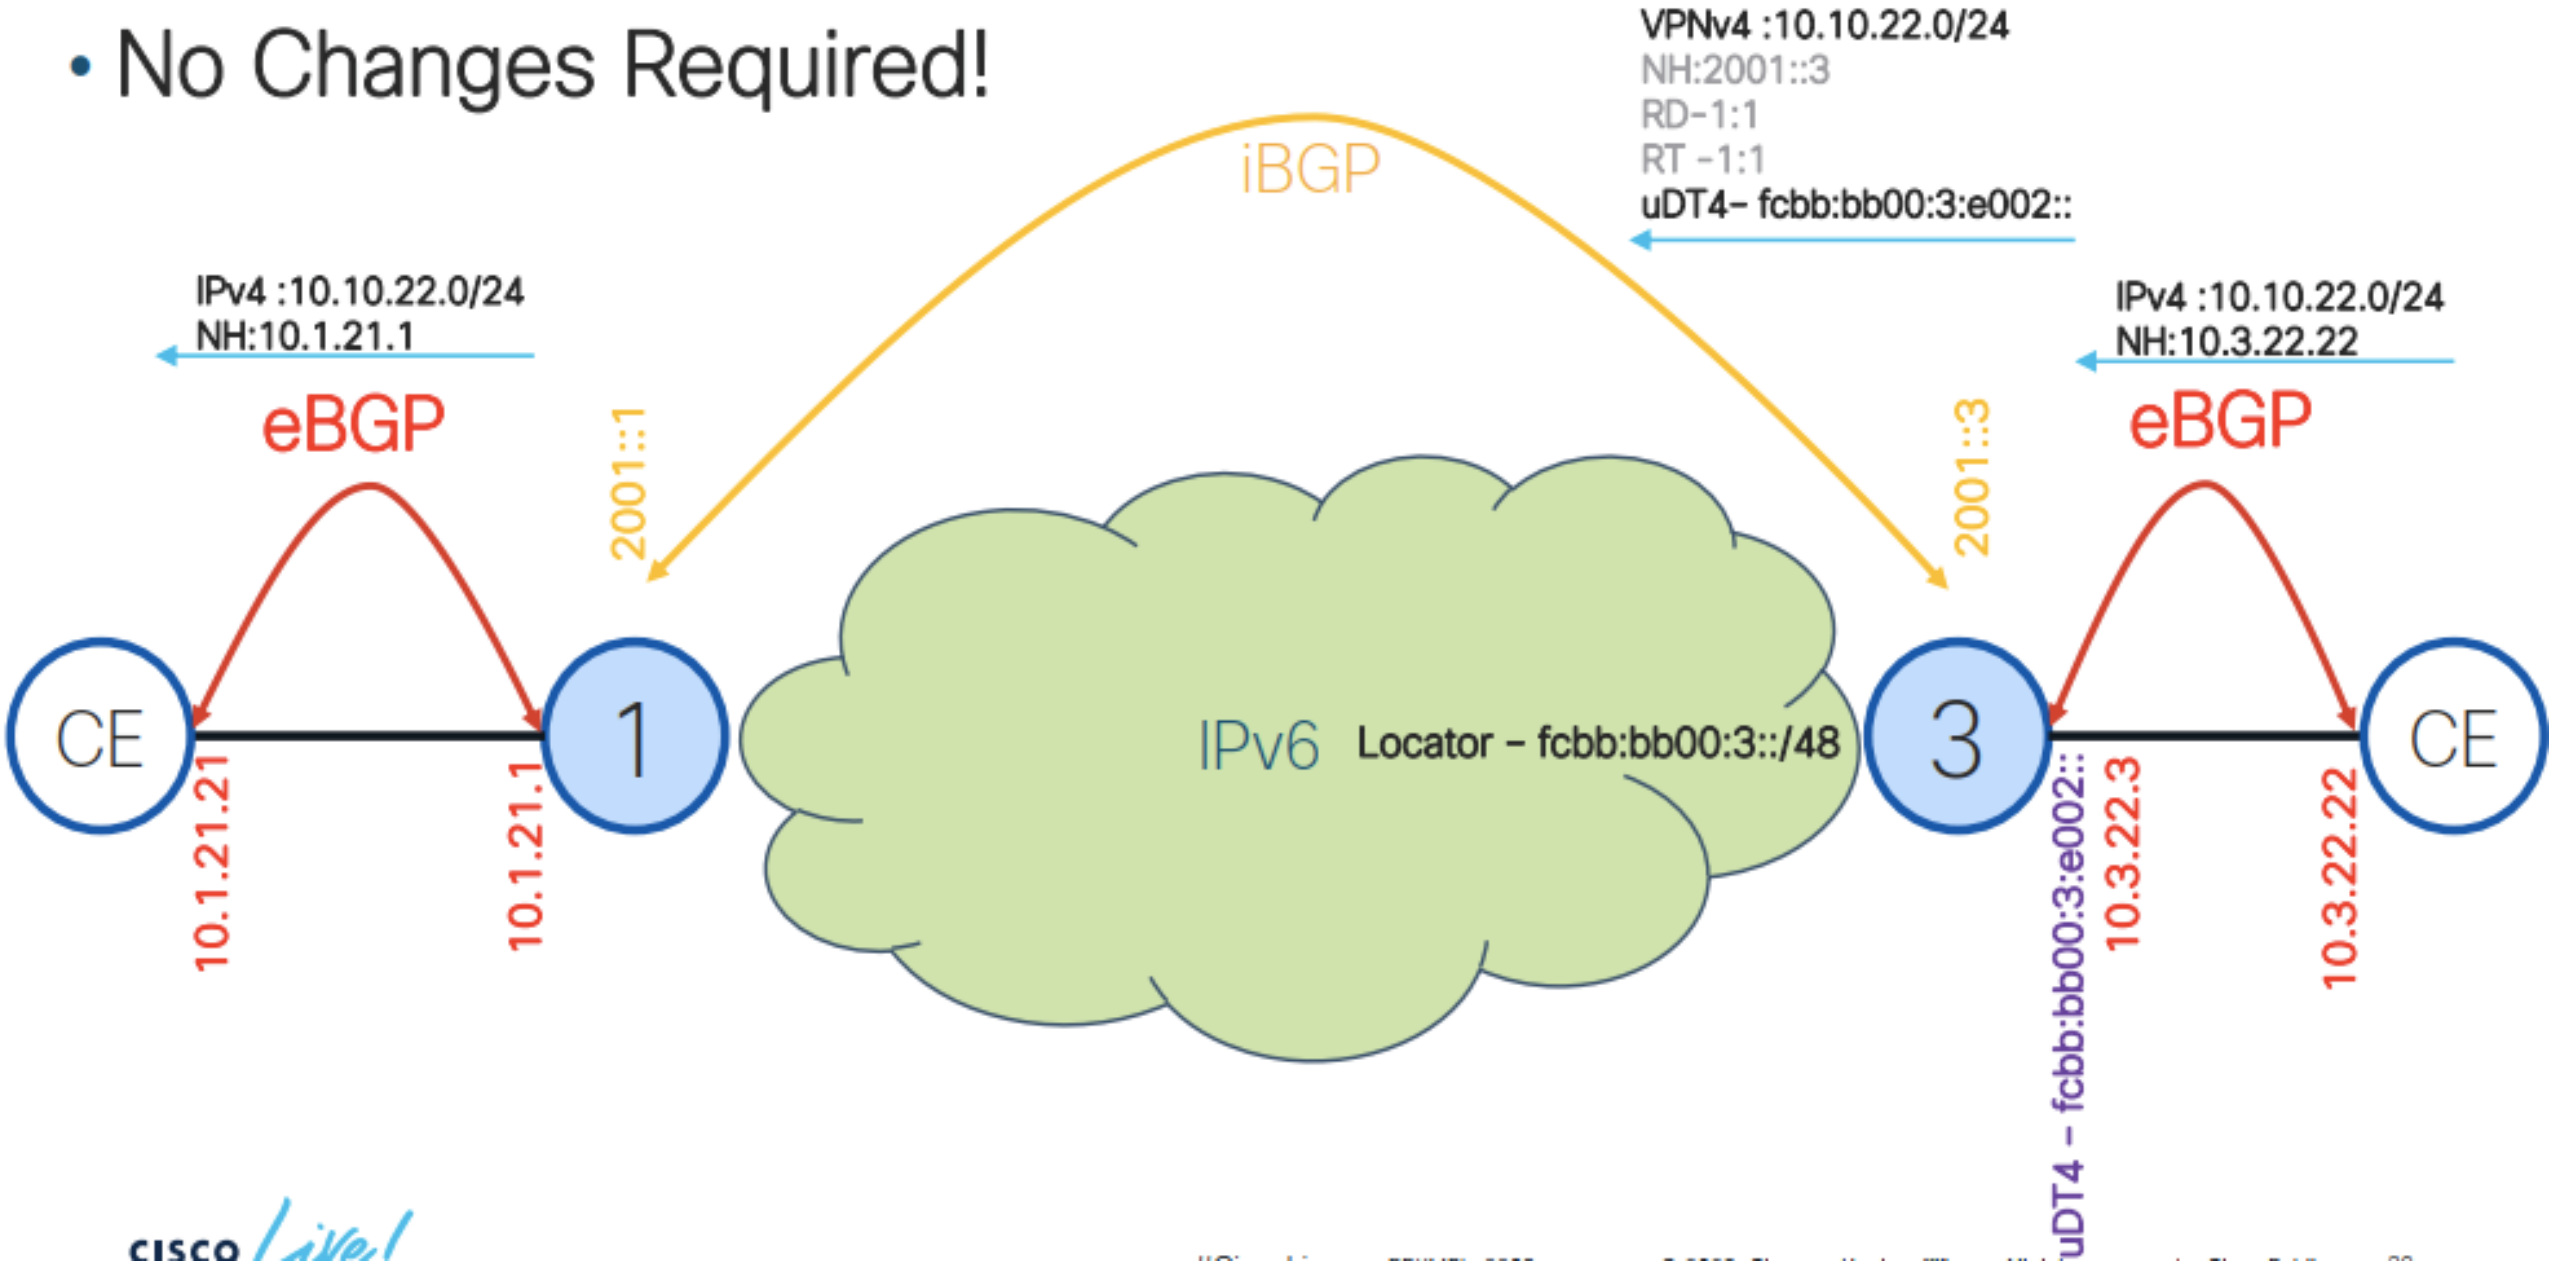

# L3 VPN Dataplane

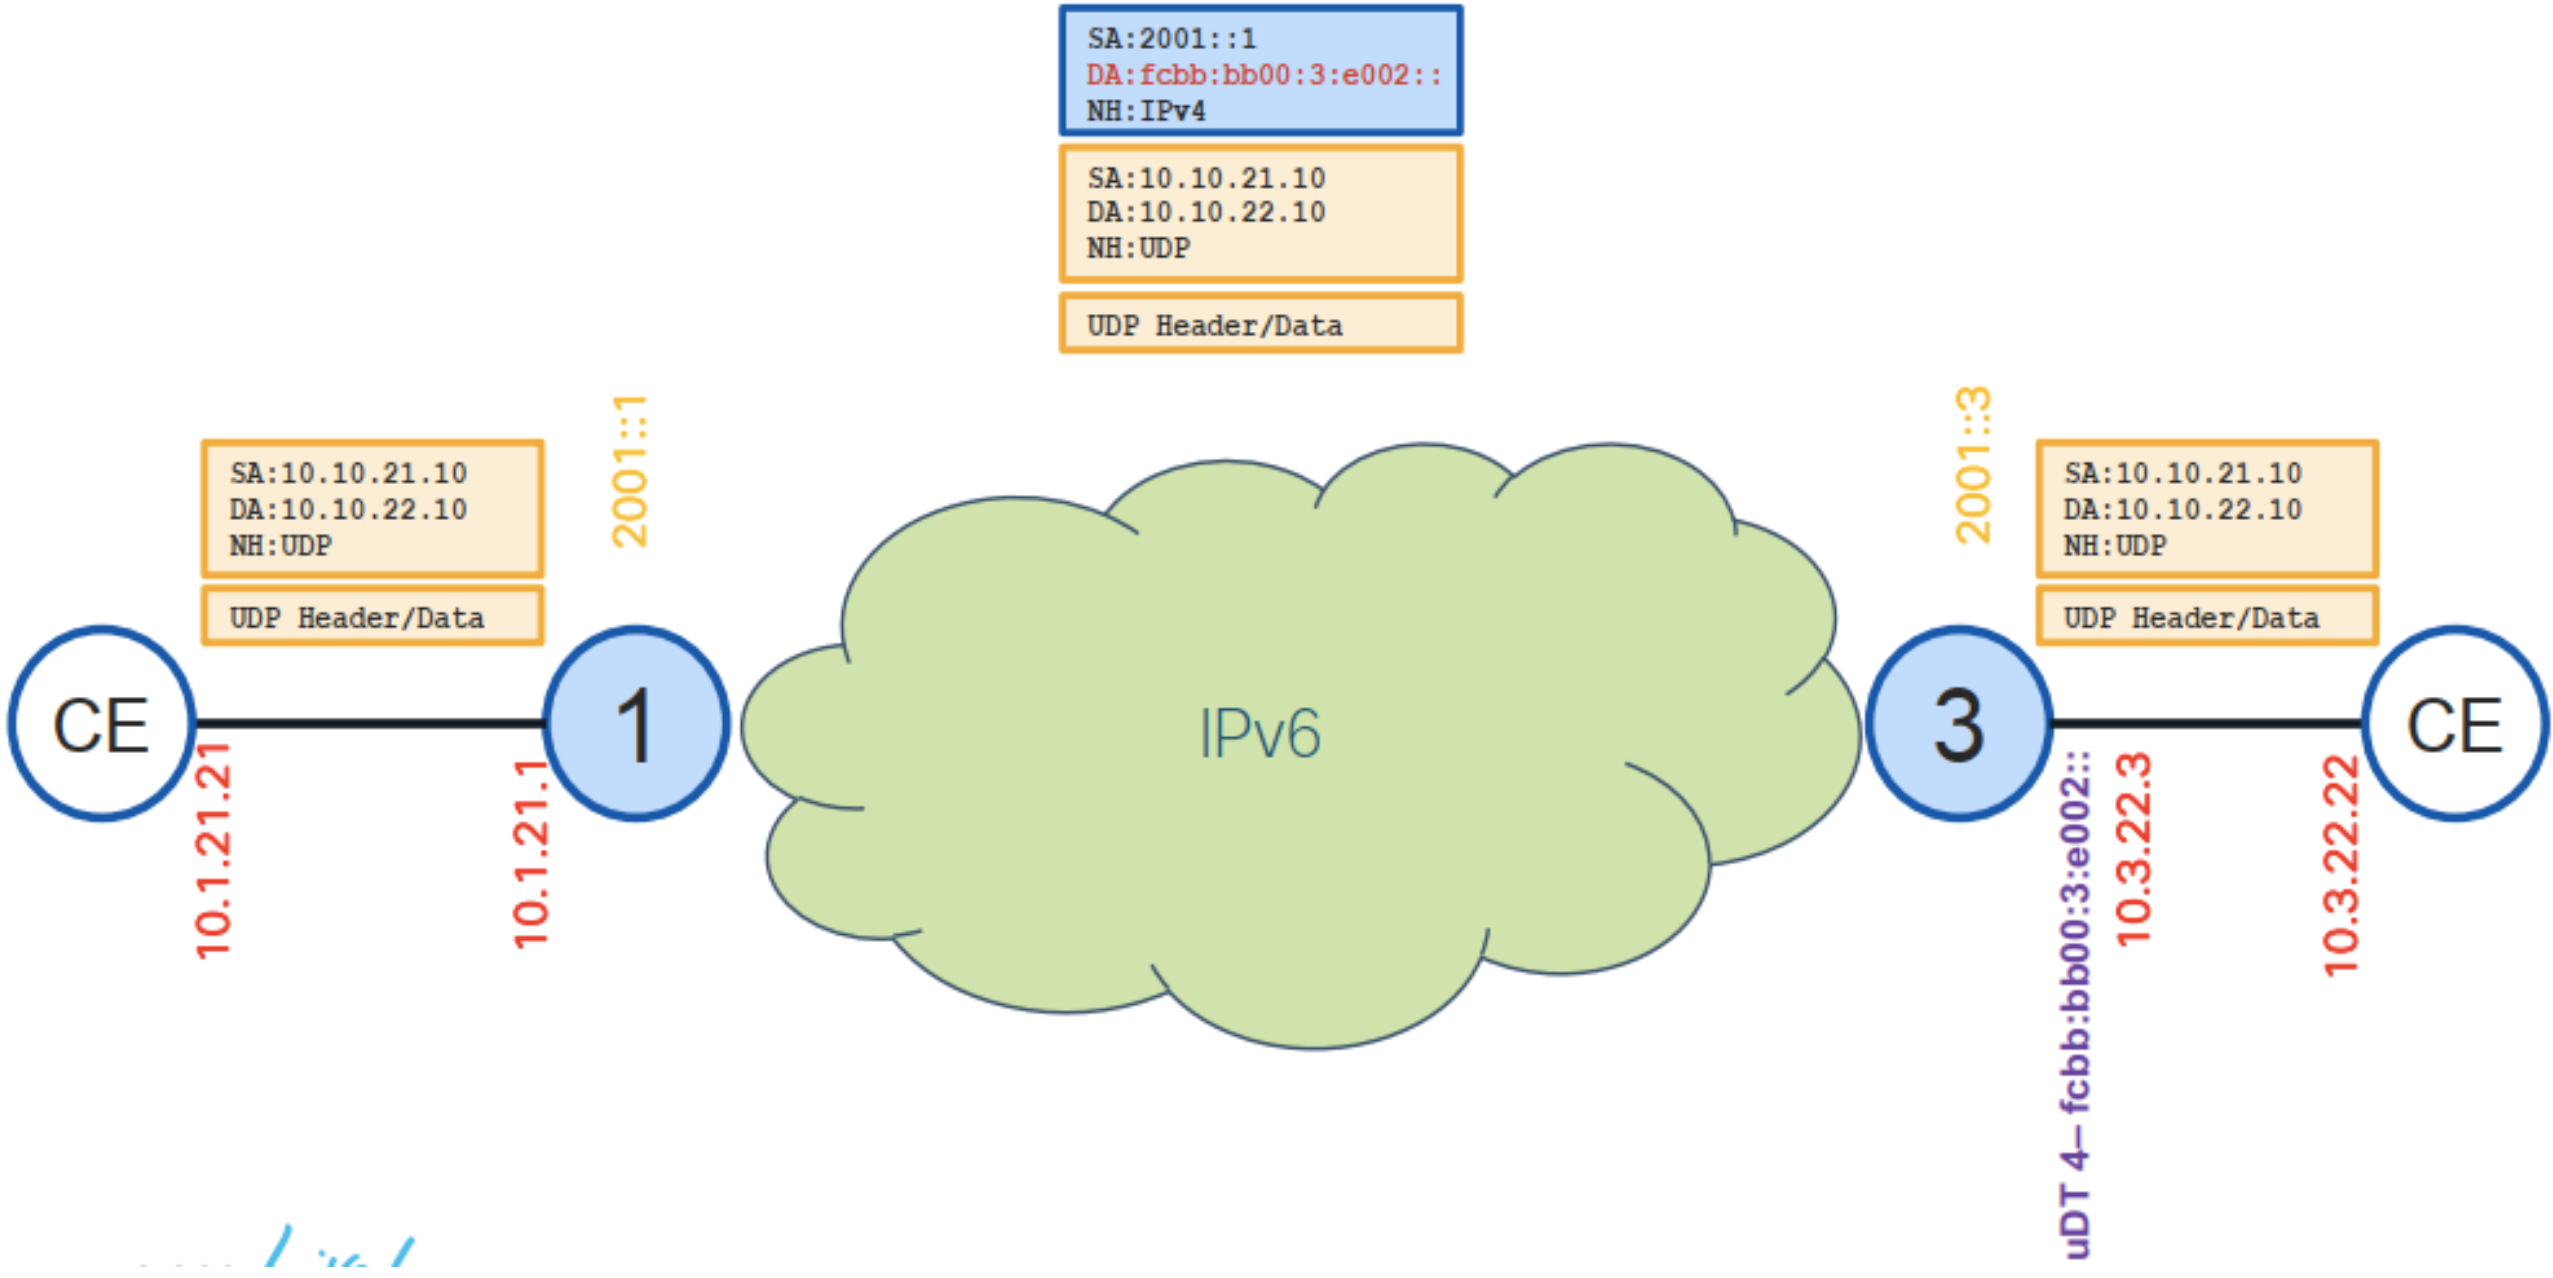

# SRv6 L3 VPN Configuration

```
router bgp 1
 address-family vpnv4 unicast
 vrf BestEffort
  rd 1:1
  address-family ipv4 unicast
   segment-routing srv6
    locator MAIN
    alloc mode per-vrf
```

Name of the Locator

Single DT function is allocated per VRF and AF

This will result in:

- uDT4 function is allocated
- All prefixes are advertised with uDT4 function

# SRv6 Feature Support for Cisco IOS-XR Platforms

| Feature name | NCS 5500<br>NCS 540 | NCS 560 | NCS 5700<br>NCS540-Q2A | ASR9K<br>(LSP) | 8000<br>(Q200) |
|---|---|---|---|---|---|
| LSR: ISIS (incl. Ti-LFA / uLoop / Flex-Algo) | Supported | | | | |
| OAM (Ping, Traceroute, SID Verification) | Supported | | | | |
| SRv6 PM (Delay, Loss, Liveness) | Supported | | | | |
| Seamless Migration (F1 -> uSID + Dual-mode) | Supported | | | | |
| L3 Services: VPNv4 / VPNv6 | Supported | | | | |
| L3 Services: IPv4 / IPv6 Internet (GRT) | Supported | | | | |
| L2 Services: EVPN-VPWS (ELINE P2P) | Supported | | | | Not supported |
| L2 Services: EVPN (ELAN BD) | Supported | | Roadmap | Supported | Not supported |
| SRv6TE: SRv6 PCE (ODN) | Supported | | | | |
| SRv6TE: Headend w/ Explicit Path | Supported | | | | |
| Path Tracing | Not supported | | Supported | | |

Reference
Cisco Live
https://www.ericsson.com/en/blog/2023/5/bright-future-of-srv6

# Thank You!!!

www.mm-ix.net

info@mm-ix.net