

Advanced Networking Topics:  
BGP, BGP Hijacking, MPLS, MPLS-based VPNs,  
Segment Routing, and others

Segment Routing over IPv6 (SRv6)

Jorge Crichigno, Elie Kfoury  
University of South Carolina

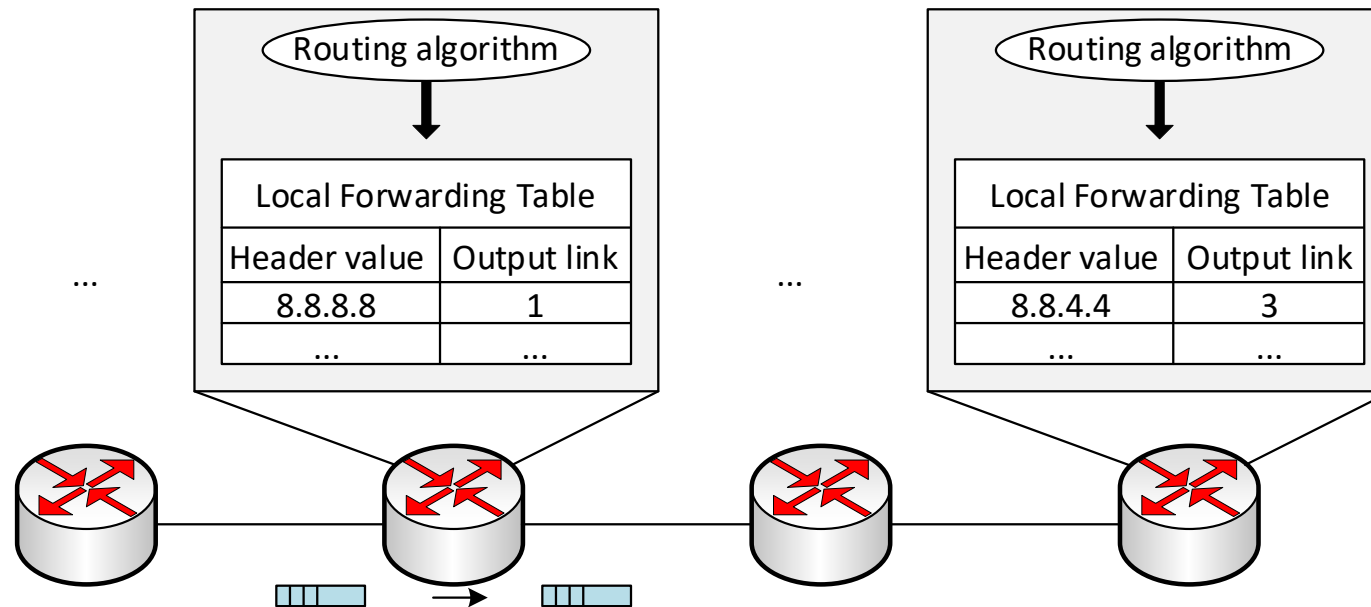
WASTC 2021 virtual Faculty Development Weeks (vFDW)

June 18, 2021

# Introduction to Segment Routing

# Conventional IP Routing


- When a packet arrives to a router, the router looks up in its forwarding table to match the incoming packet and determine the corresponding action (e.g., forward the packet to port 1)



J. Kurose, K. Ross, "Computer networking, a top-down approach," 7th Edition, Pearson, 2017

# Conventional IP Routing

- Every router in the network maintains a routing table
- Routing table lookup determines the appropriate output port for an incoming packet
- Conventional IP routing uses the packet's destination IP address



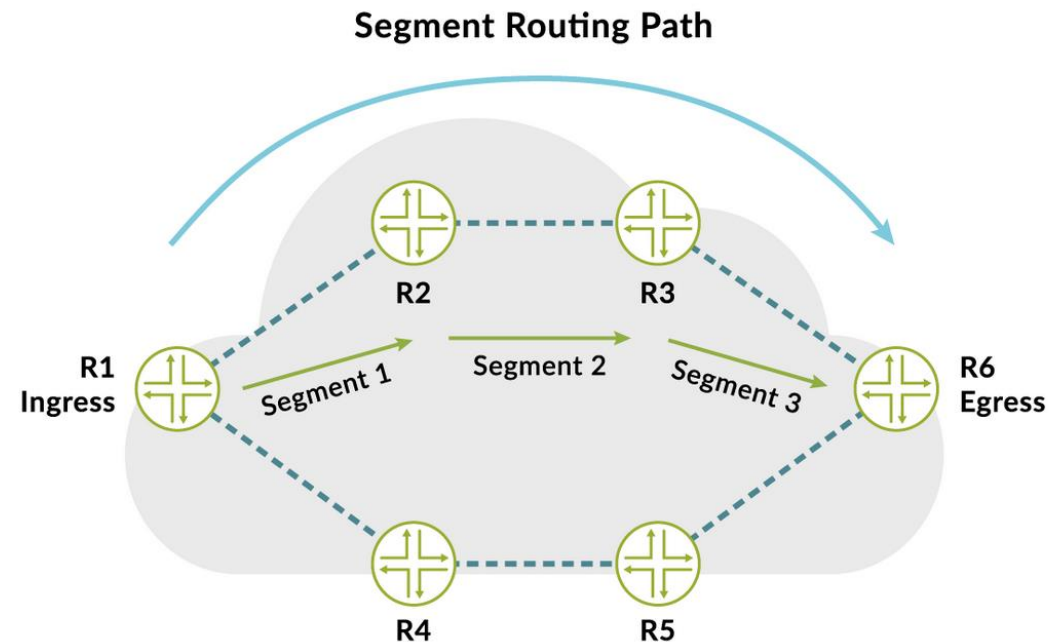
```
R1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

  10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D    10.1.1.0/24 [90/2170112] via 209.165.200.226, 00:00:05, Serial0/0/0
D    10.1.2.0/24 [90/2170112] via 209.165.200.226, 00:00:05, Serial0/0/0
  192.168.10.0/24 is variably subnetted, 2 subnets, 3 masks
C    192.168.10.0/24 is directly connected, GigabitEthernet0/0
L    192.168.10.1/32 is directly connected, GigabitEthernet0/0
C    192.168.11.0/24 is variably subnetted, 2 subnets, 3 masks
C    192.168.11.0/24 is directly connected, GigabitEthernet0/1
L    192.168.11.1/32 is directly connected, GigabitEthernet0/1
  209.165.200.0/24 is variably subnetted, 2 subnets, 3 masks
C    209.165.200.224/30 is directly connected, Serial0/0/0
L    209.165.200.225/32 is directly connected, Serial0/0/0
R1#
```

# Segment Routing

- Segment Routing (SR) is a form of source routing that encodes path information in the packet
- A node inserts a header that contains a list of segments to packets
- Subsequent nodes in the network execute the instructions in the segments



Juniper Networks, "What is Segment Routing?". Online: <https://tinyurl.com/2tp2njb5>

# Advantages of Segment Routing

- Scalability; SR removes network state information from transit routers and nodes
- Flexibility and agility; highly responsive to network changes
- Simplicity; SR eliminates protocols: LDP, RSVP-TE, VxLAN, NSH, GTP
- End-to-end; zero modification for the packet outside the segment routing domain

# Segment Routing Data Plane Implementations

---

- Segment routing can be implemented using two different data planes
- Segment routing over Multiprotocol Label Switching (SR-MPLS)
  - Each segment is a label
  - The segment list is a label stack
  - No change in the forwarding plane
- Segment routing over IPv6 (SRv6)
  - Each segment is an address
  - The segment list is an address list
  - New header Segment Routing Header (SRH) is defined in RFC8754

# Segment Routing Data Plane Implementations

---

- Segment routing can be implemented using two different data planes
- Segment routing over Multiprotocol Label Switching (SR-MPLS)
  - Each segment is a label
  - The segment list is a label stack
  - No change in the forwarding plane
- Segment routing over IPv6 (SRv6) ← Focus of this session
  - Each segment is an address
  - The segment list is an address list
  - New header Segment Routing Header (SRH) is defined in RFC8754



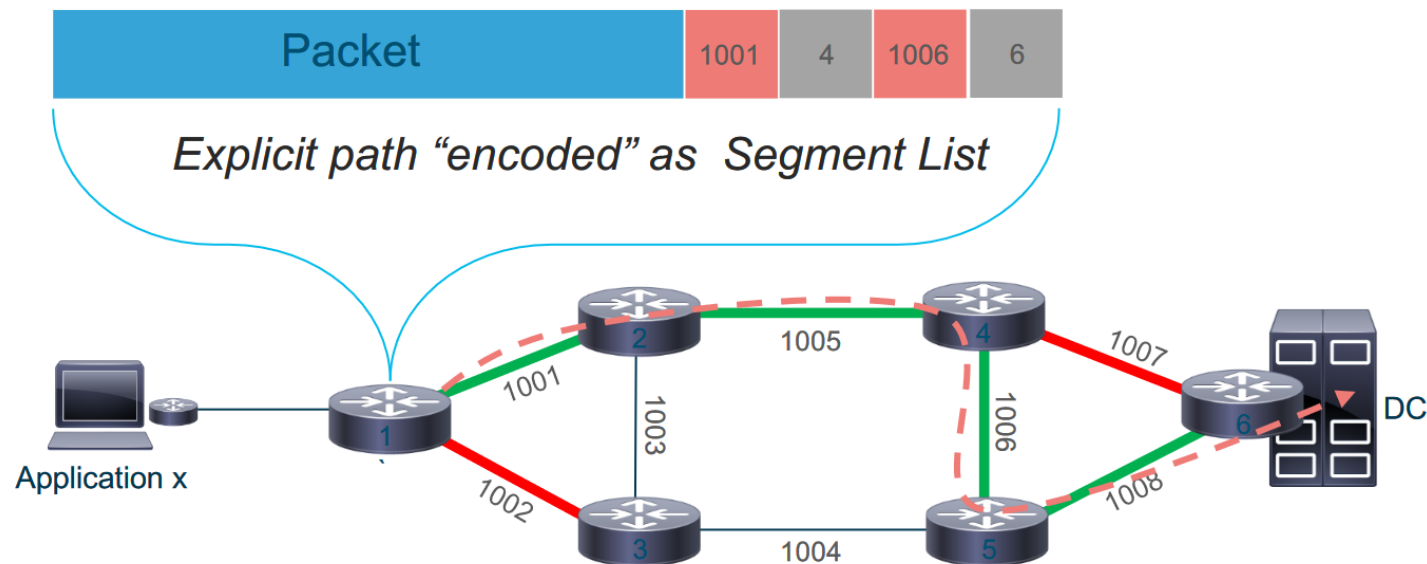
# SRv6 Features

---

- SRv6 leverages the IPv6 Extension Headers (RFC 8200)
- SRv6 benefits from all features deployed over the years on MPLS network
  - Traffic engineering
  - Fast reroute
  - etc.
- With SRv6, it is possible to pack more than IPv6 addresses into a segment ID
- SRv6 enables functionalities that go beyond routing and traffic steering
  - Service Function Chaining (SFC)
  - Virtual Private Networks (VPNs)
  - etc.
- The focus of this session is on SRv6 basic configuration for traffic steering

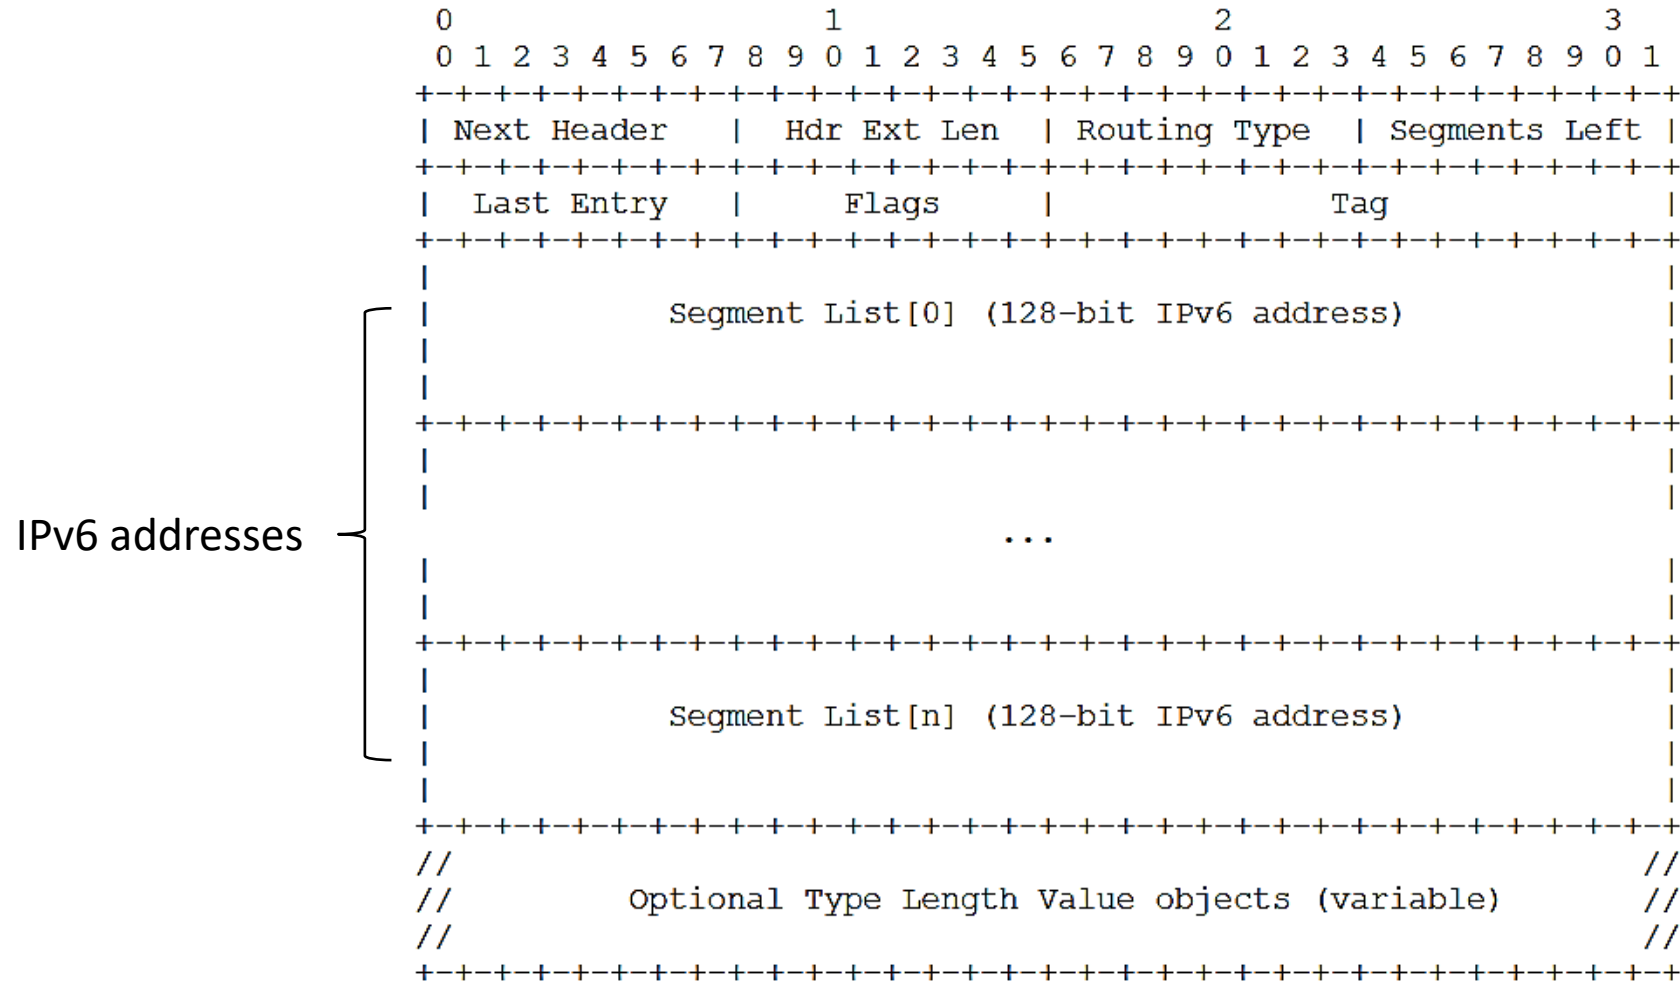
# Near-stateless Traffic Engineering

- SRv6 enables stateless traffic engineering (TE)
- SRv6 eliminates the need for complex RSVP-TE and complex TE configurations

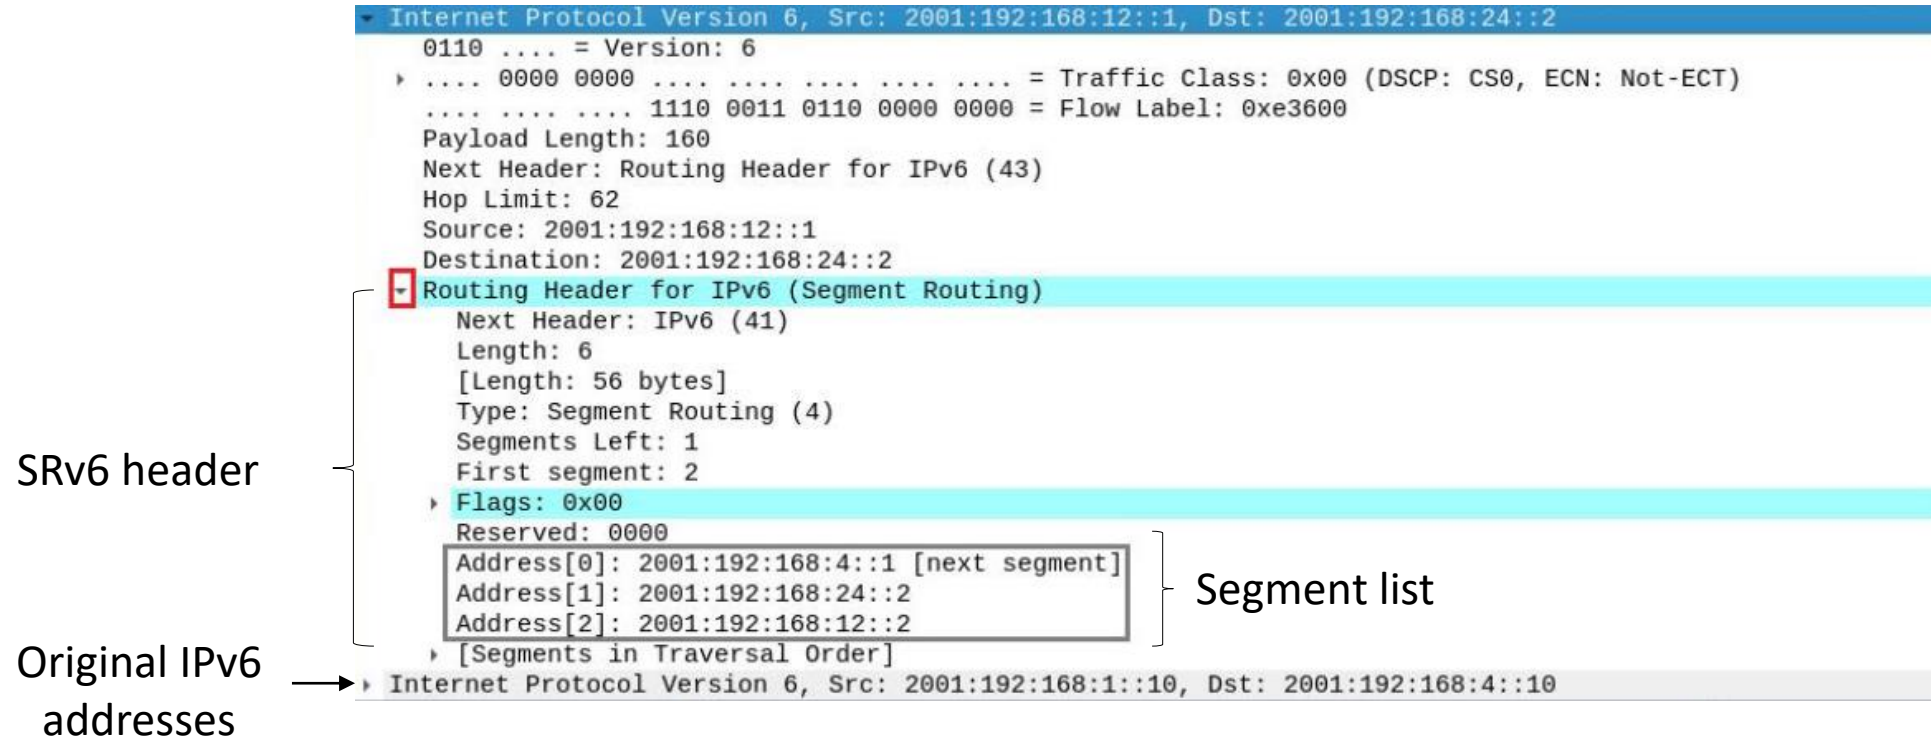


Hernán Contreras G, Cisco Systems. "Network Programming with SRv6". LACNOG, September 2018

# SRv6 Header Format




# SRv6 Header Format



# SRv6 Example


- SR path: R1-R2-R4



- › Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
- › Internet Control Message Protocol v6

# SRv6 Example


- SR path: R1-R2-R4



```
Internet Protocol Version 6, Src: 2001:192:168:12::1, Dst: 2001:192:168:12::2
  0110 .... = Version: 6
  ▶ .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 1011 1110 1101 0011 0111 = Flow Label: 0xbed37
  Payload Length: 144
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 63
  Source: 2001:192:168:12::1
  Destination: 2001:192:168:12::2
  ▶ Routing Header for IPv6 (Segment Routing)
  ▶ Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
```

# SRv6 Example

- SR path: R1-R2-R4



```


- Internet Protocol Version 6, Src: 2001:192:168:12::1, Dst: 2001:192:168:12::2
  0110 .... = Version: 6
  ▶ .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 1011 1110 1101 0011 0111 = Flow Label: 0xbed37
  Payload Length: 144
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 63
  Source: 2001:192:168:12::1
  Destination: 2001:192:168:12::2
  ▶ Routing Header for IPv6 (Segment Routing)
- Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
  
```

```

Routing Header for IPv6 (Segment Routing)
  Next Header: IPv6 (41)
  Length: 4
  [Length: 40 bytes]
  Type: Segment Routing (4)
  Segments Left: 1
  First segment: 1
  ▶ Flags: 0x00
  Reserved: 0000
  Address[0]: 2001:192:168:24::2 [next segment]
  Address[1]: 2001:192:168:12::2
  ▶ [Segments in Traversal Order]
  
```

# SRv6 Example

- SR path: R1-R2-R4



```

Internet Protocol Version 6, Src: 2001:192:168:12::1, Dst: 2001:192:168:12::2
  0110 .... = Version: 6
  ▶ .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 1011 1110 1101 0011 0111 = Flow Label: 0xbed37
  Payload Length: 144
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 63
  Source: 2001:192:168:12::1
  Destination: 2001:192:168:12::2
  ▶ Routing Header for IPv6 (Segment Routing)
  ▶ Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
  
```


```

Routing Header for IPv6 (Segment Routing)
  Next Header: IPv6 (41)
  Length: 4
  [Length: 40 bytes]
  Type: Segment Routing (4)
  Segments Left: 1
  First segment: 1
  ▶ Flags: 0x00
  Reserved: 0000
  Address[0]: 2001:192:168:24::2 [next segment]
  Address[1]: 2001:192:168:12::2
  ▶ [Segments in Traversal Order]
  
```



# SRv6 Example


- SR path: R1-R2-R4



```
Internet Protocol Version 6, Src: 2001:192:168:12::1, Dst: 2001:192:168:24::2
  0110 .... = Version: 6
  ▶ .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 1011 1110 1101 0011 0111 = Flow Label: 0xbed37
  Payload Length: 144
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 62
  Source: 2001:192:168:12::1
  Destination: 2001:192:168:24::2
  ▶ Routing Header for IPv6 (Segment Routing)
  ▶ Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
```

# SRv6 Example

- SR path: R1-R2-R4



```

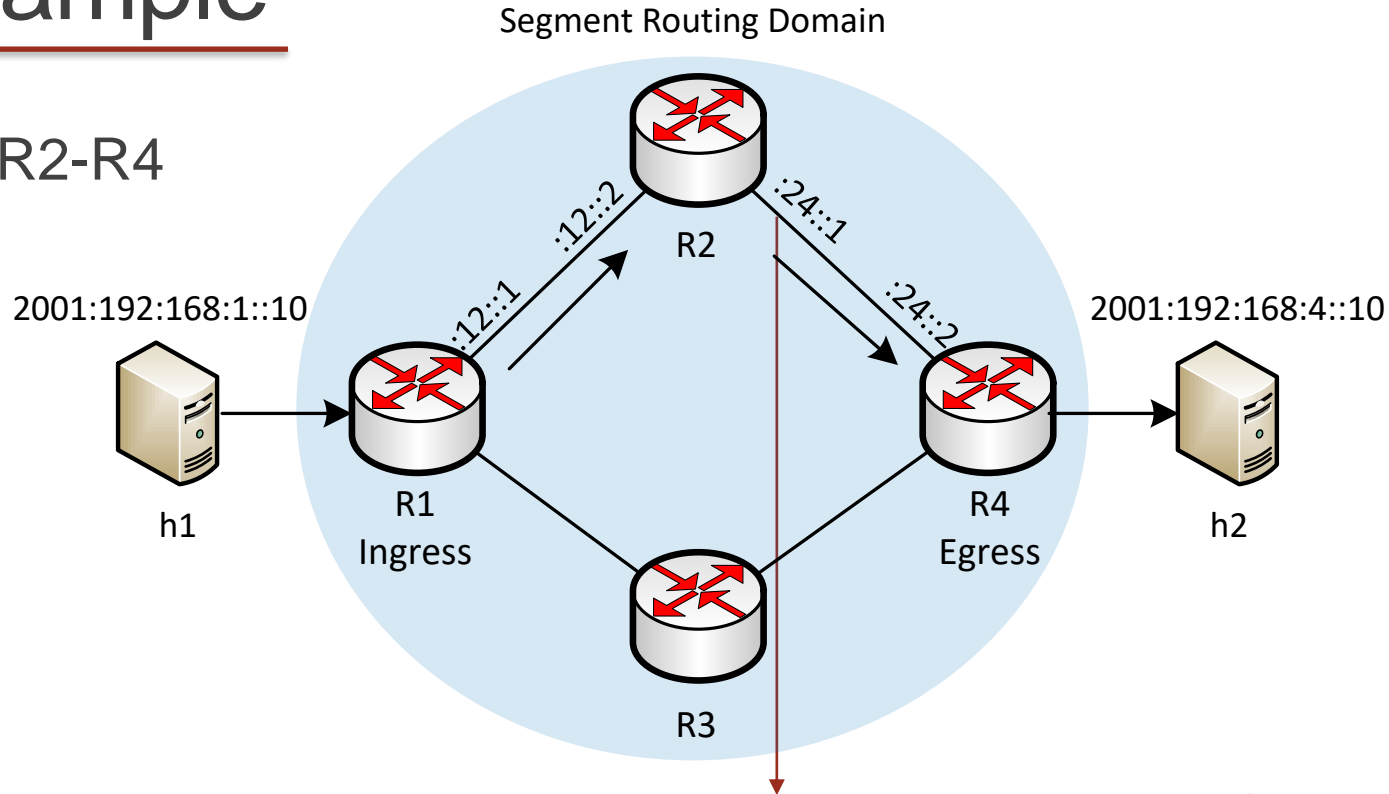
- Internet Protocol Version 6, Src: 2001:192:168:12::1, Dst: 2001:192:168:24::2
  0110 .... = Version: 6
  > .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 1011 1110 1101 0011 0111 = Flow Label: 0xbed37
  Payload Length: 144
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 62
  Source: 2001:192:168:12::1
  Destination: 2001:192:168:24::2
  > Routing Header for IPv6 (Segment Routing)
- Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
  
```

```

- Routing Header for IPv6 (Segment Routing)
  Next Header: IPv6 (41)
  Length: 4
  [Length: 40 bytes]
  Type: Segment Routing (4)
  Segments Left: 0
  First segment: 1
  > Flags: 0x00
  Reserved: 0000
  Address[0]: 2001:192:168:24::2
  Address[1]: 2001:192:168:12::2
  > [Segments in Traversal Order]
  
```

# SRv6 Example

- SR path: R1-R2-R4




```

- Internet Protocol Version 6, Src: 2001:192:168:12::1, Dst: 2001:192:168:24::2
  0110 .... = Version: 6
  > .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 1011 1110 1101 0011 0111 = Flow Label: 0xbed37
  Payload Length: 144
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 62
  Source: 2001:192:168:12::1
  Destination: 2001:192:168:24::2
  > Routing Header for IPv6 (Segment Routing)
- Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
- Routing Header for IPv6 (Segment Routing)
  Next Header: IPv6 (41)
  Length: 4
  [Length: 40 bytes]
  Type: Segment Routing (4)
  Segments Left: 0
  First segment: 1
  > Flags: 0x00
  Reserved: 0000
  Address[0]: 2001:192:168:24::2
  Address[1]: 2001:192:168:12::2
  > [Segments in Traversal Order]
  
```

# SRv6 Example

- SR path: R1-R2-R4




```

- Internet Protocol Version 6, Src: 2001:192:168:12::1, Dst: 2001:192:168:24::2
  0110 .... = Version: 6
  > .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 1011 1110 1101 0011 0111 = Flow Label: 0xbed37
  Payload Length: 144
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 62
  Source: 2001:192:168:12::1
  Destination: 2001:192:168:24::2
  > Routing Header for IPv6 (Segment Routing)
- Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
- Routing Header for IPv6 (Segment Routing)
  Next Header: IPv6 (41)
  Length: 4
  [Length: 40 bytes]
  Type: Segment Routing (4)
  Segments Left: 0
  First segment: 1
  > Flags: 0x00
  Reserved: 0000
  Address[0]: 2001:192:168:24::2
  Address[1]: 2001:192:168:12::2
  > [Segments in Traversal Order]
  
```

# SRv6 Example

- SR path: R1-R2-R4



- › Internet Protocol Version 6, Src: 2001:192:168:1::10, Dst: 2001:192:168:4::10
- › Internet Control Message Protocol v6

# SRv6 beyond Traffic Steering

# SRv6 Network Programming

---

- SRv6's original intention was only to steer traffic
- Unlike SR-MPLS, SRv6 was viewed as a more distant goal and received less attention
- SRv6 gained back the attention when the SRv6 Network Programming draft was submitted to the IETF

# SRv6 Network Programming

---

- Network instruction: 128-bit SRv6 SID
- Locator: routed to the node performing the function
- Function: either local to a Network Processing Unit (NPU) or app in VM/Container
- Flexible bit-length selection



---

Clarence Filsfils, Cisco "SRv6". Online: <https://www.segment-routing.net/tutorials/2017-12-05-srv6-introduction/>



# SRv6 Functions

---

**End** Endpoint function The SRv6 instantiation of a prefix SID

**End.X** Endpoint function with Layer-3 cross-connect The SRv6 instantiation of a Adj SID

**End.T** Endpoint function with specific IPv6 table lookup

**End.DX2** Endpoint with decapsulation and Layer-2 cross-connect L2VPN use-case

**End.DX2V** Endpoint with decapsulation and VLAN L2 table lookup EVPN Flexible cross-connect use-cases

**End.DT2U** Endpoint with decaps and unicast MAC L2 table lookup EVPN Bridging unicast use-cases

**End.DT2M** Endpoint with decapsulation and L2 table flooding EVPN Bridging BUM use-cases with ESI filtering

**End.DX6** Endpoint with decapsulation and IPv6 cross-connect IPv6 L3VPN use (equivalent of a per-CE VPN label)

**End.DX4** Endpoint with decapsulation and IPv4 cross-connect IPv4 L3VPN use (equivalent of a per-CE VPN label)

**End.DT6** Endpoint with decapsulation and IPv6 table lookup IPv6 L3VPN use (equivalent of a per-VRF VPN label)

**End.DT4** Endpoint with decapsulation and IPv4 table lookup IPv4 L3VPN use (equivalent of a per-VRF VPN label)

**End.B6** Endpoint bound to an SRv6 policy SRv6 instantiation of a Binding SID

**End.B6.Encaps** Endpoint bound to an SRv6 encapsulation Policy SRv6 instantiation of a Binding SID

**End.BM** Endpoint bound to an SR-MPLS Policy SRv6/SR-MPLS instantiation of a Binding SID

---

Clarence Filisfilis, Cisco "SRv6". Online: <https://www.segment-routing.net/tutorials/2017-12-05-srv6-introduction/>

# End SID

---

- Endpoint function ("End" for short) is the most basic function
- End function pseudocode

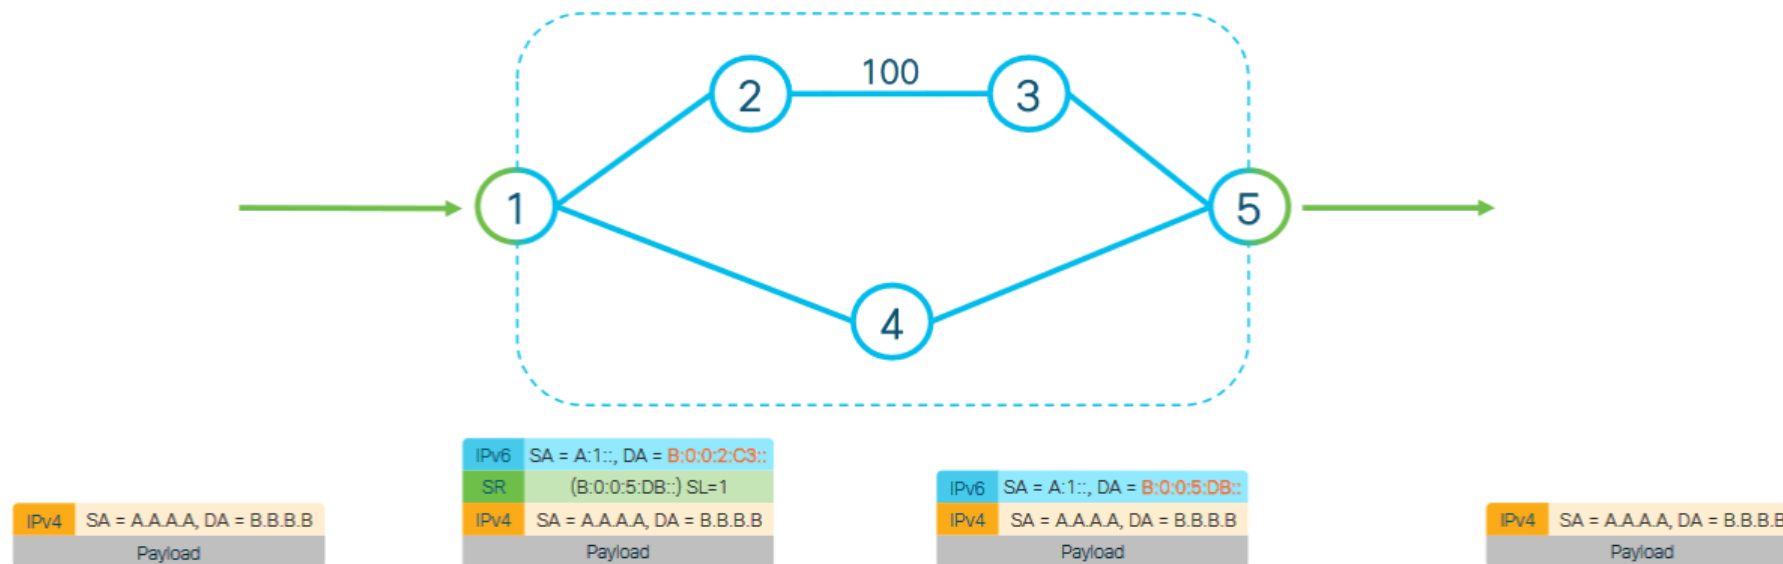
```
1.  IF NH=SRH and SL > 0
2.      decrement SL
3.      update the IPv6 DA with SRH[SL]
4.      FIB lookup on updated DA
5.      forward accordingly to the matched entry
6.  ELSE
7.      drop the packet
```

---

Clarence Filsfils, Cisco "SRv6". Online: <https://www.segment-routing.net/tutorials/2017-12-05-srv6-introduction/>

# IPv6-enabled Provider Infrastructure

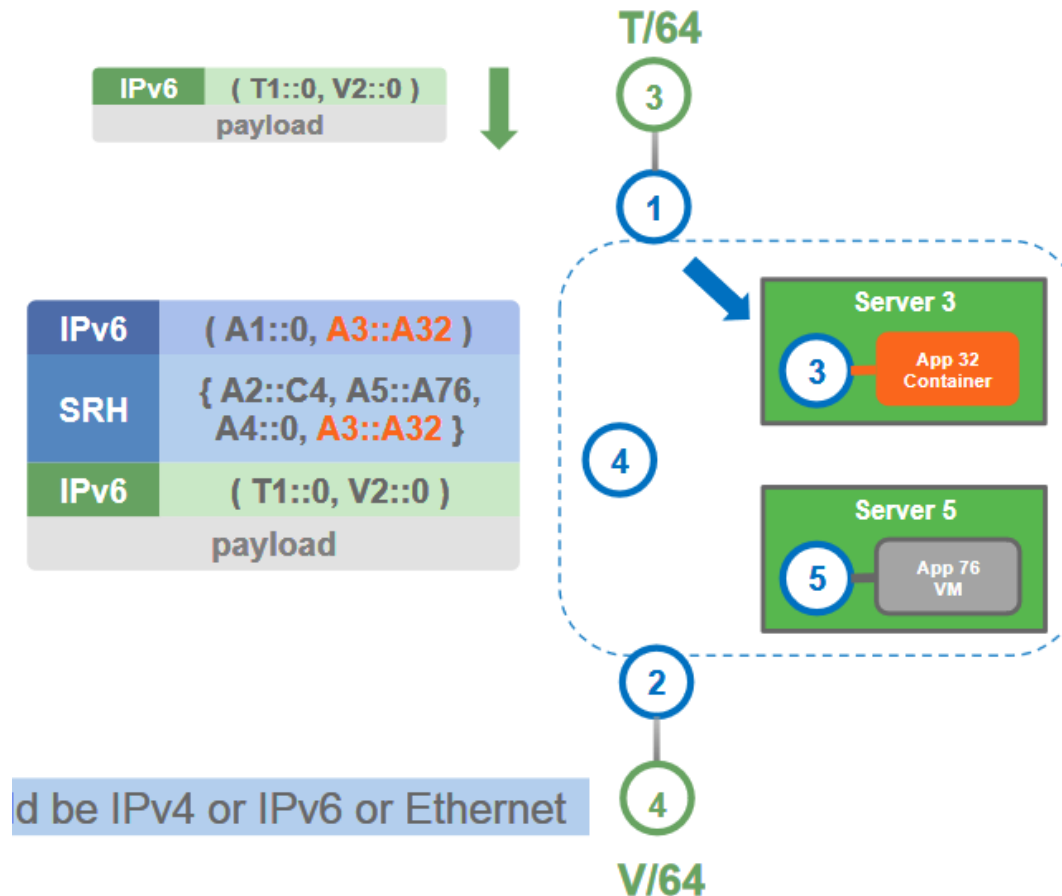
- IPv4, IPv6 or L2 frame is encapsulated within the SR Domain
- Outer IPv6 header includes an SRH with the list of segments



Clarence Filisfilis, Cisco "SRv6". Online: <https://www.segment-routing.net/tutorials/2017-12-05-srv6-introduction/>

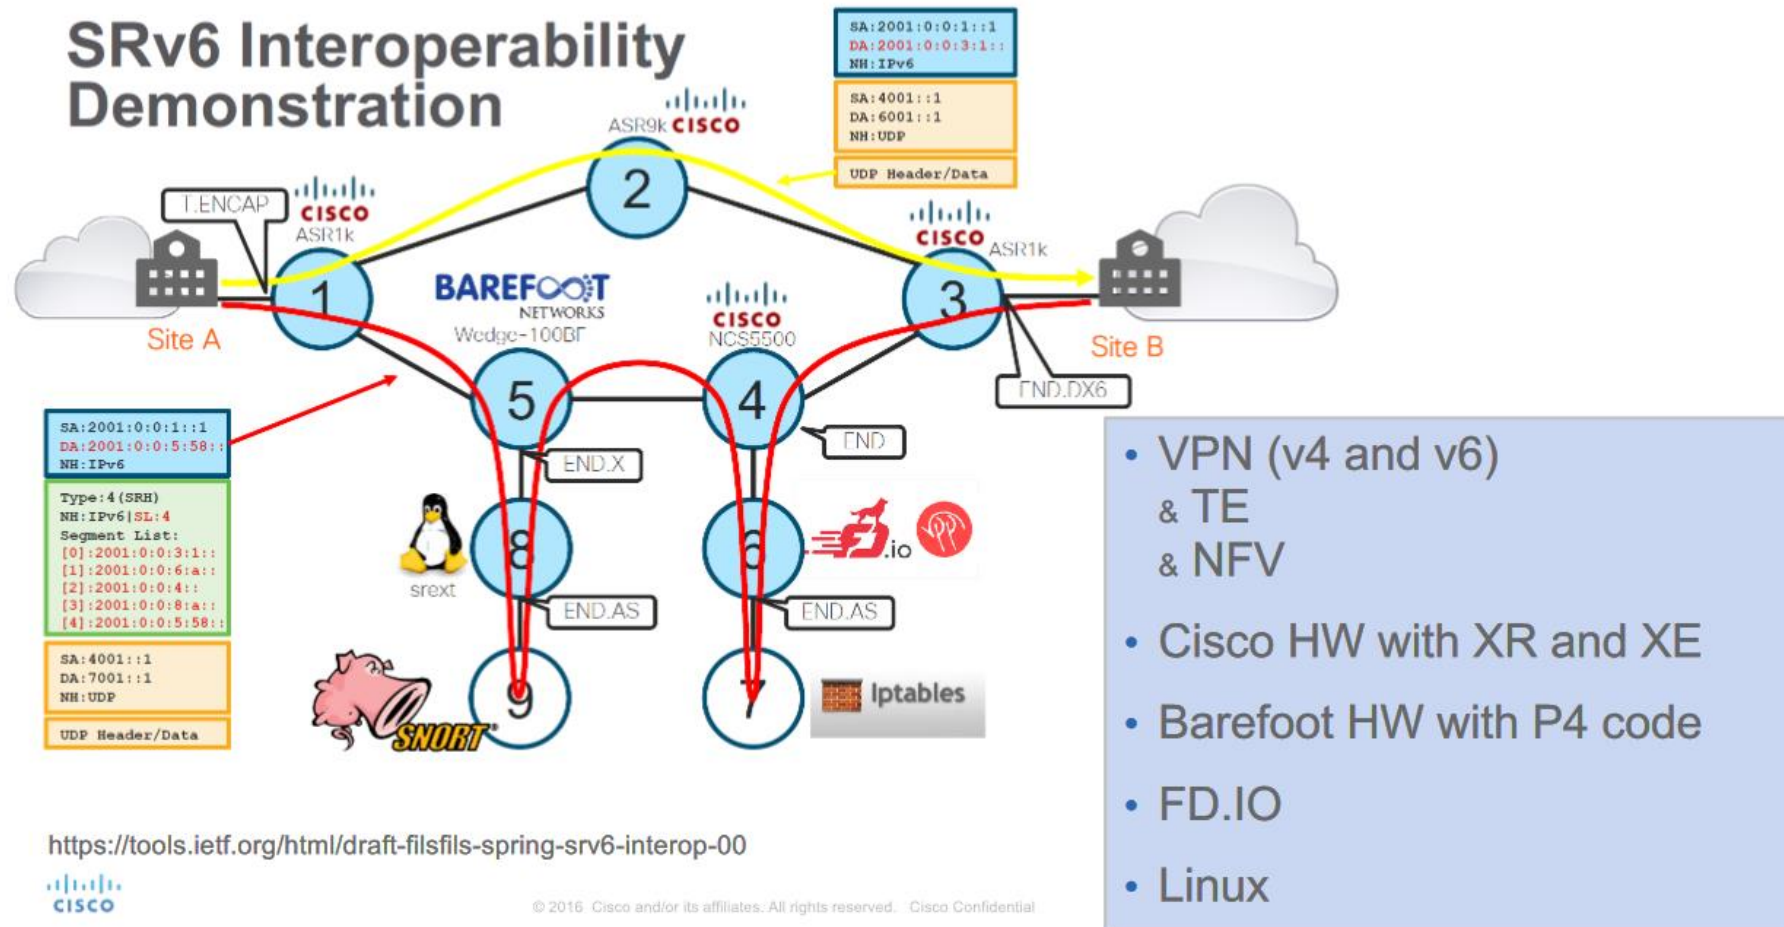
# Use Case: Integrated NFV

- Stateless Service Chaining
- App is SR aware or not



Clarence Filisfilis, Cisco "SRv6". Online: <https://www.segment-routing.net/tutorials/2017-12-05-srv6-introduction/>


# SRv6 Interoperability



Clarence Filsfils, Cisco “SRv6”. Online: <https://www.segment-routing.net/tutorials/2017-12-05-srv6-introduction/>

# Further Readings

- State of SR (Cisco-maintained website): <https://www.segment-routing.net/>
- SRv6 Linux implementation: <https://segment-routing.org/>



Segment Routing, Part I  
Clarence Filstils



Segment Routing, Part II  
Clarence Filstils



IP Network eBook Series

**SRv6**

Author: Lanjun Luo

Information Digitalization and Experience Assurance (IDEA) Department



IP Network eBook Series  
- SRv6, Huawei