



Segment Routing with IPv6

Hernán Contreras G.

Consulting Systems Engineer – Cisco Systems

May 2015

Gianpietro Lavado

glch@cisco.com

 @gianpietro_lc

 /gianpietro

Segment Routing in Summary

The source chooses a path and encodes it in the packet header as an ordered list of segments

Packet is transmitted from source node with a list of segment ID (SID) inserted at header

List dictates the path to follow in the network

A segment is an instruction applied to the packet and represent any type of instruction:

- Nodal SID (global, IGP and BGP)
- Adjacency SID (local, IGP)
- Peer SID (local, BGP)

Segment Routing uses current Network Protocols

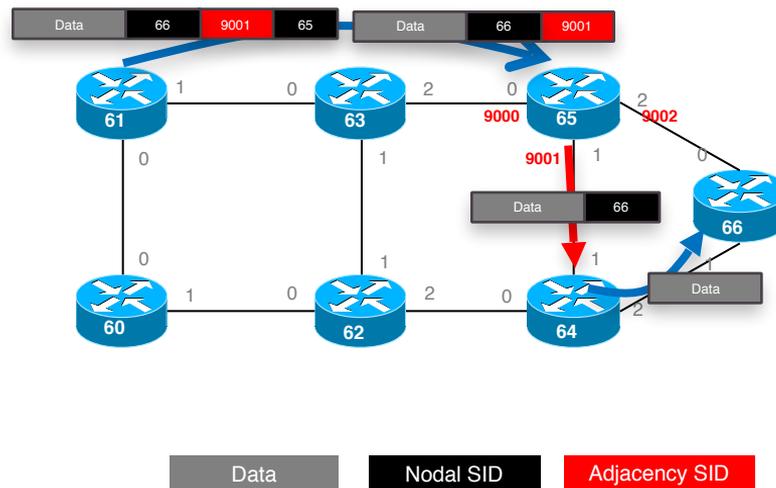
- SID are signaled in IGP: ISIS or OSPF
- SID are signaled in BGP

Segment Routing defines two data planes

- MPLS (Label = Segment ID)
- IPv6 + optional SR header (SID = 20bits flag)

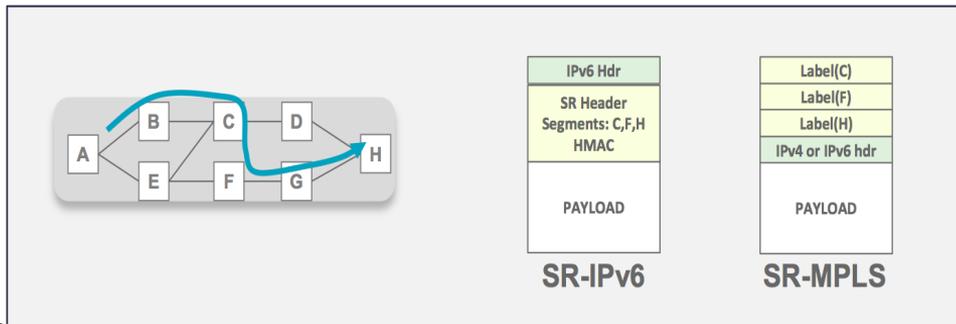
Segment Routing is currently being standardized at IETF

- SPRING, Segment Packet Routing in Networking
- <https://tools.ietf.org/wg/spring/>



Segment Routing for IPv6 Dataplane

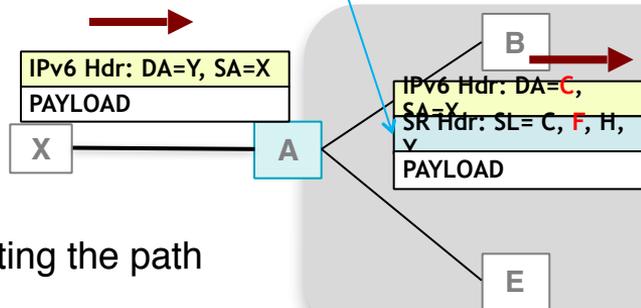
- SR-IPv6 allows Segment Routing to be deployed over non-MPLS networks and/or in areas of the network where MPLS is not present (e.g.: datacenters)
- Segment Routing for IPv6 Dataplane
 - Defines SID as 128 bit IPv6 addresses (no label or MPLS Data Plane)
 - Much simpler from a signaling perspective: no need to advertise anything else than IPv6 prefixes (the prefix is the SID)
 - draft-psenak-ospf-segment-routing-ospfv3-extension-02
- Segments are coded at Header of IPv6 Packet
 - Define a new Routing Extensions Header type
 - > Segment Routing Header (SRH)
 - > Contains Segment List
 - > VERY close to what already specified in RFC2460
 - > draft-previdi-6man-segment-routing-header-00
- SR nodes fully interoperate with non-SR nodes and could be used in combination with SR-MPLS (in networks where MPLS is used in the core and IPv6 is used at the edge, like home networks and datacenters).



IPv6 Segment Routing Header

SRH Contains the list of segments the packet should traverse

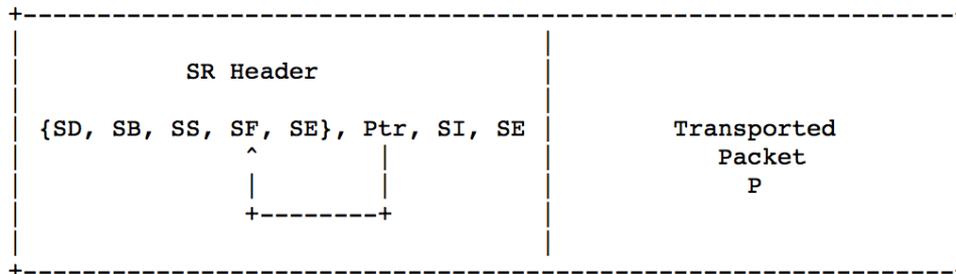
- At ingress:
 - Path is computed or received by a controller
 - Path is instantiated through a list of segments
 - A SRH is created with the segment list representing the path
- Segment Routing Header:
 - **Segment List** describes the path of the packet: list of segments (IPv6 addresses)
 - **Next Segment**: a pointer to the segment list element identifying the next segment
 - **HMAC**
 - **Flags** and optional policy information
- The Active Segment is set as the DA of the packet
 - At each segment endpoint, the DA is updated with the “Next Segment”
 - Compliant with RFC2460 rules for the Routing Header



IPv6 Segment Routing Header (SRH) draft-previdi-6man-segment-routing-header-05

- This draft describes the Segment Routing Extension Header Type and how it is used by SR capable nodes.
- The abstract SR header contains an ordered list of segments, a pointer identifying the next segment to process and the identifications of the ingress and egress SR edge routers on the path of this packet.
 - The pointer identifies the segment that **MUST** be used by the receiving router to process the packet. This segment is called the active segment.
 - A property of SR is that the entire source route of the packet, including the identity of the ingress and egress edge routers is always available with the packet.
 - Pointer is moved to the active segment over the list according to **PUSH, NEXT and CONTINUE** operations

Routers on the path that simply forward an IPv6 packet (i.e. the IPv6 destination address is none of theirs) will never inspect and process the SRH.

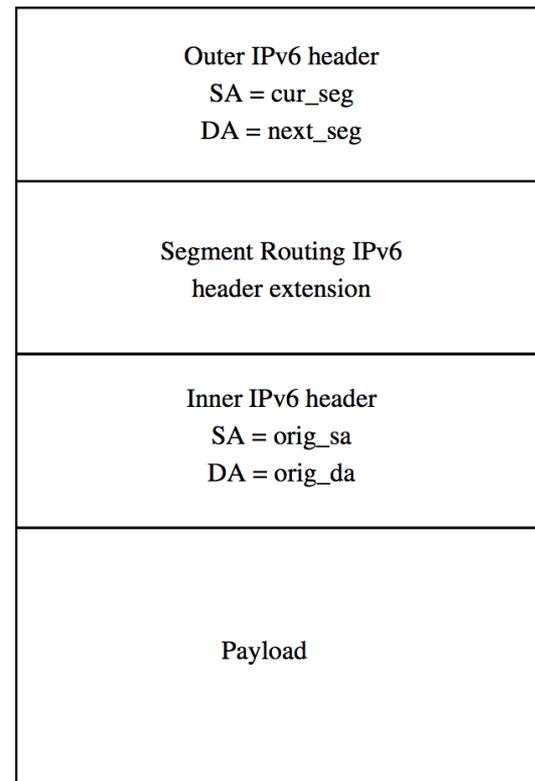


IPv6 Segment Routing Header (SRH)

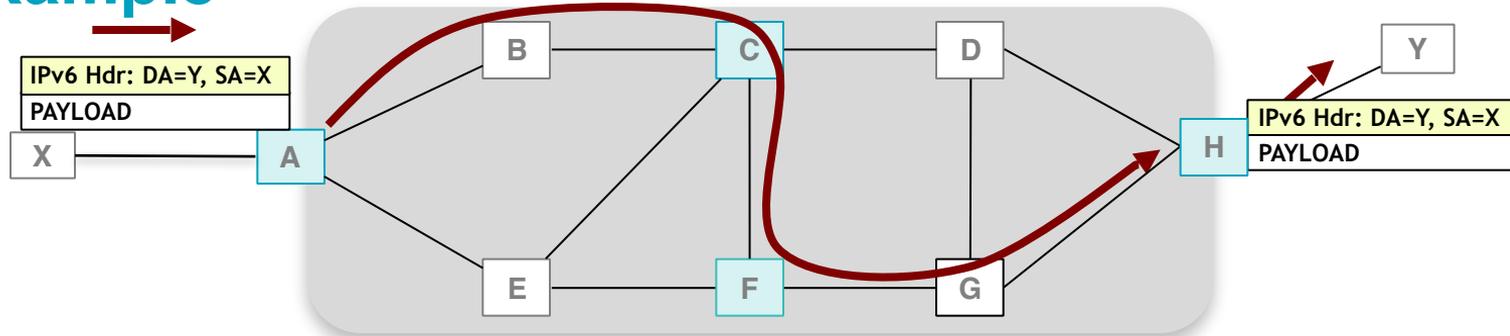
draft-previdi-6man-segment-routing-header-05

The active segment is in the destination address of the packet.

- Segment Routing in IPv6 networks implies that the segment identifier is moved into the DA of the packet.
 - The DA of the packet changes at each segment termination/completion and therefore the original DA of the packet MUST be encoded as the last segment of the path.
 - Nodes that are within the path of a segment will forward packets based on the DA of the packet without inspecting the SRH. This ensures full interoperability between SR-capable and non-SR-capable nodes.
- SR operations involves inspecting and updating the SRH and update the DA.
 - Next Segment field contains the offset of the SECOND segment of the path which is encoded in the FIRST position of the segment list
- Encapsulation can be realized in two different ways with SR-IPv6:
 - Outer encapsulation.
 - SRH with SA/DA original addresses.



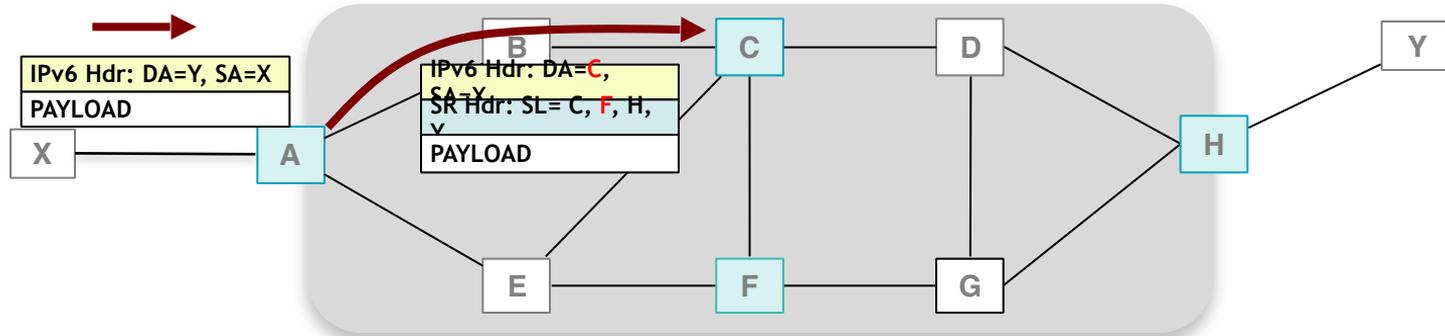
SR-IPv6 Example



- Example:

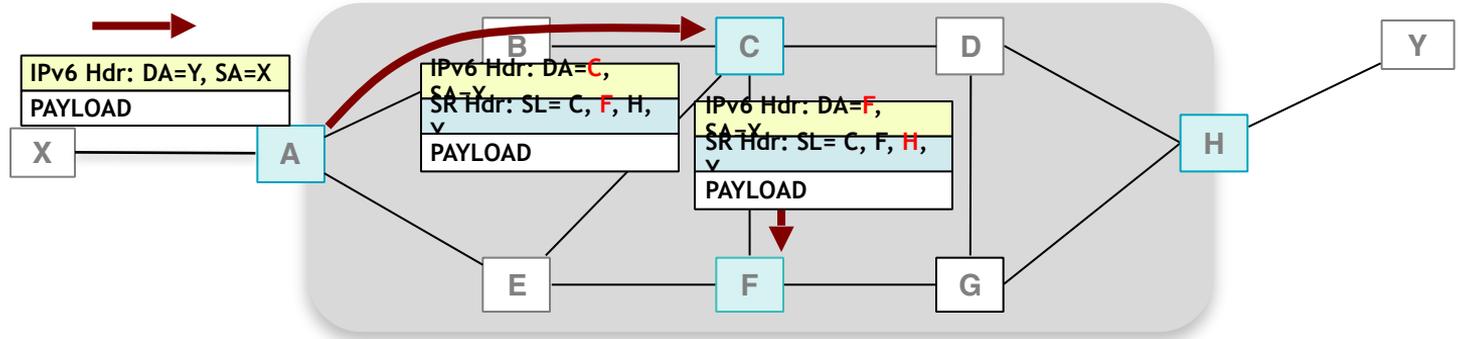
- Classify packets coming from X and destined to Y and forward them across A,B,C,F,G,H path
- Nodes A, C, F and H are SR capable

SR-IPv6 Example



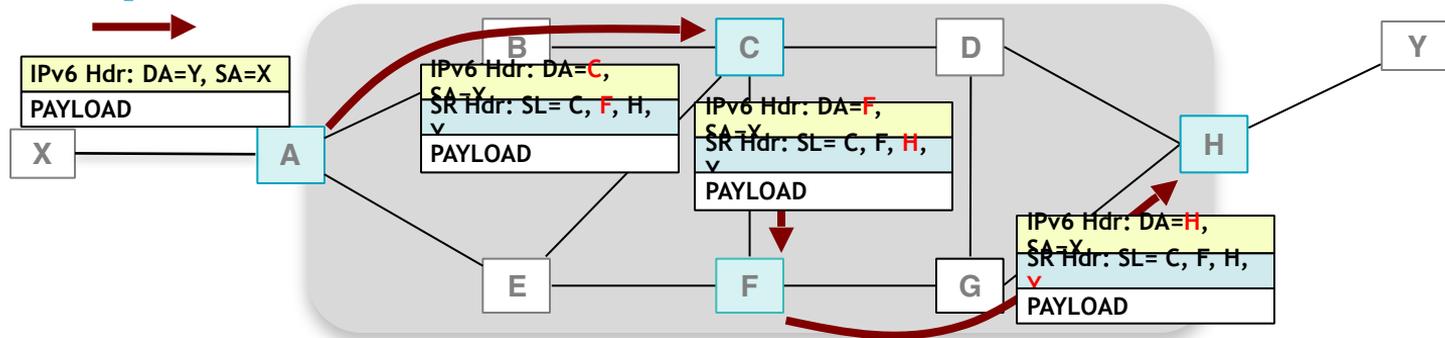
- At ingress, the Segment Routing Header (SRH) contains
 - **Segment List:** C,F,H,Y (original destination address is encoded as last segment of the path)
 - **Next Segment:** points to the next segment of the path (F)
 - **DA** is set as the address of the first segment: C
- Packet is sent towards its DA (C, representing the first segment)
 - Packet can travel across non SR nodes who will just ignore the SRH
 - RFC2460 mandates only the node in the DA must examine the SRH

SR-IPv6 Example



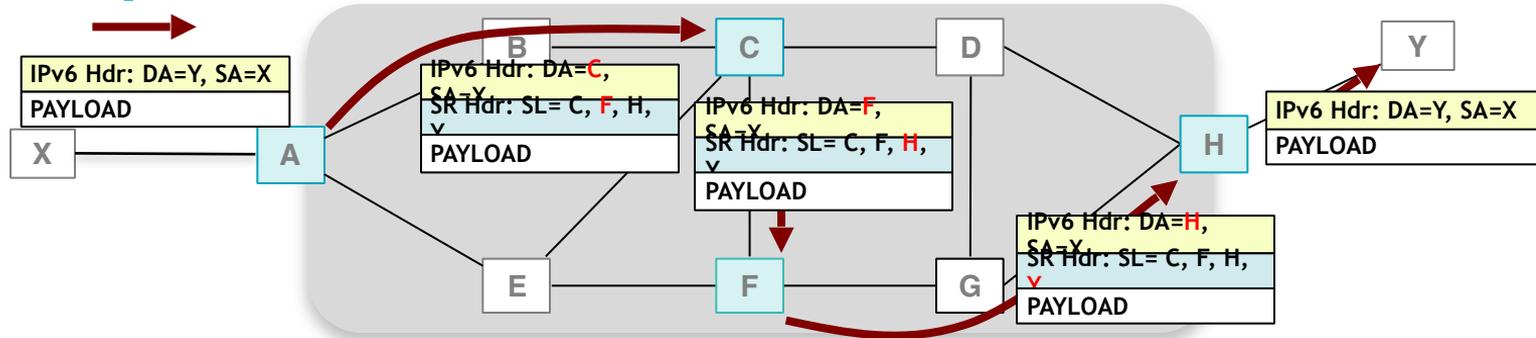
- When packet reaches the segment endpoint C
 - Next Segment is inspected and used in order to update the DA with the next segment address: F
 - Next Segment pointer is incremented: now points to H
 - Packet is sent towards its DA

SR-IPv6 Example



- When packet reaches the segment endpoint F the same process is executed:
 - Next Segment is inspected and used in order to update the DA with the next segment address: H
 - Next Segment pointer is incremented: now points to Y (the original DA)
 - Packet is sent towards its DA

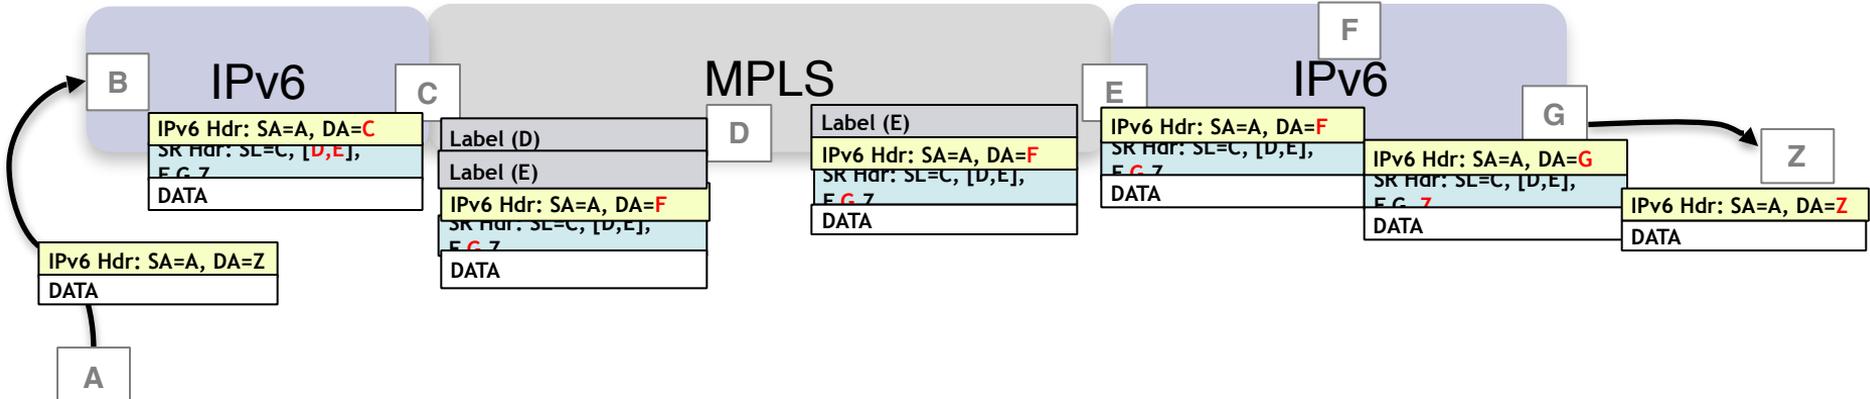
SR-IPv6 Example



- When packet reaches the segment endpoint H:
 - Next Segment is inspected and used in order to update the DA with the next segment address: Y
 - A flag (cleanup-flag) in SRH tells H to cleanup the packet and remove the SRH
 - Packet is sent towards its DA

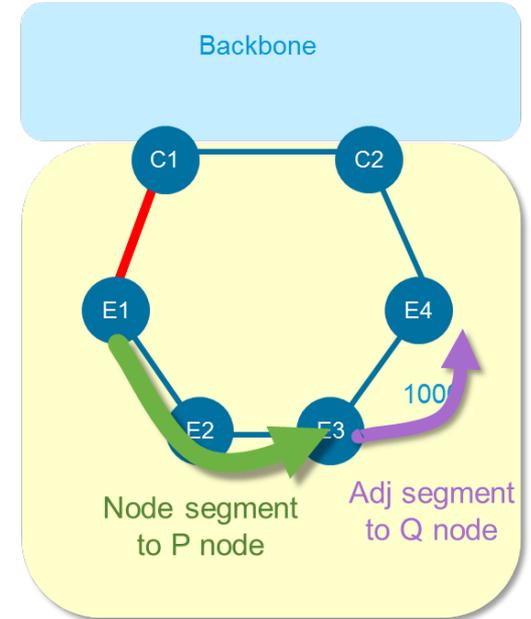
IPv6 SR6-LABEL Prefix

- MPLS segments are encoded in the special SR6-LABEL segment
 - Allow end to end path to include IPv6 and MPLS dataplane nodes
 - Up to 4 MPLS segments in one SR6-LABEL prefix that counts as a single segment
- Path is computed/expressed at ingress with both V6 and MPLS portion
- SRH contains both
 - V6 segments → IPv6 addresses
 - MPLS segments → encoded within a reserved/configured v6 prefix (up to 4 labels per v6 address)



Segment Routing Use Cases: Fast Reroute 50-msec Protection for IGP Segments

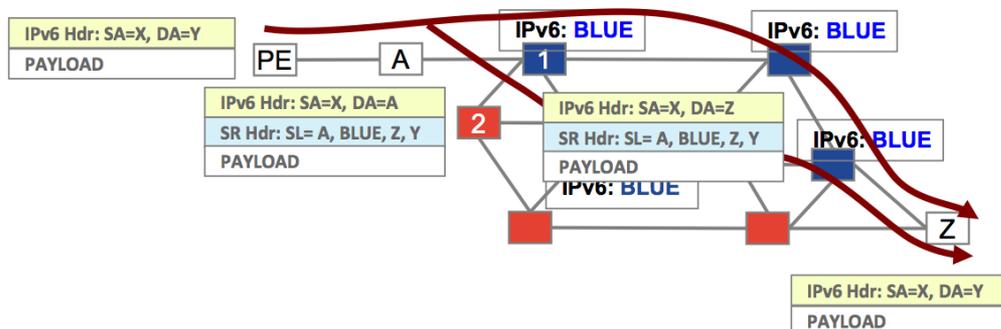
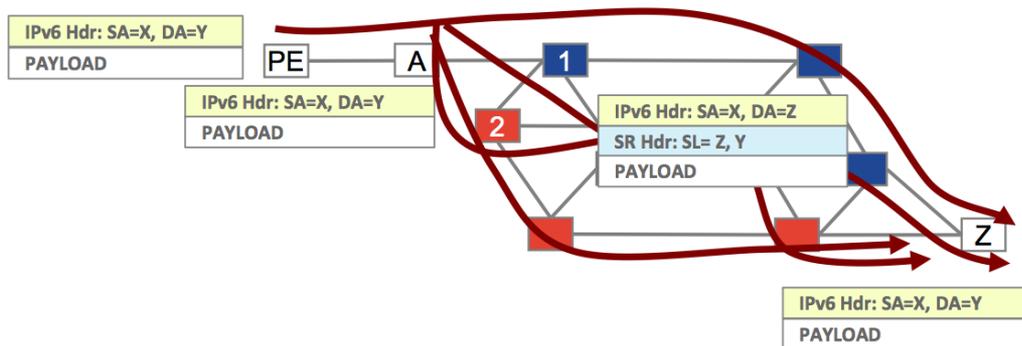
- Guaranteed Link/Node FRR in any topology
 - Repair/Backup path expressed as a list of segment and pre-installed in FIB
 - Upon failure, the protecting node reroute traffic according to new Segment List
- Simplicity
 - Entirely automated
 - No signaling
 - No intermediate state
- Incremental deployment – Applicable to all traffic
- Optimal backup path along post-convergence path
 - Prevents transient congestion and suboptimal routing



Default metric: 10

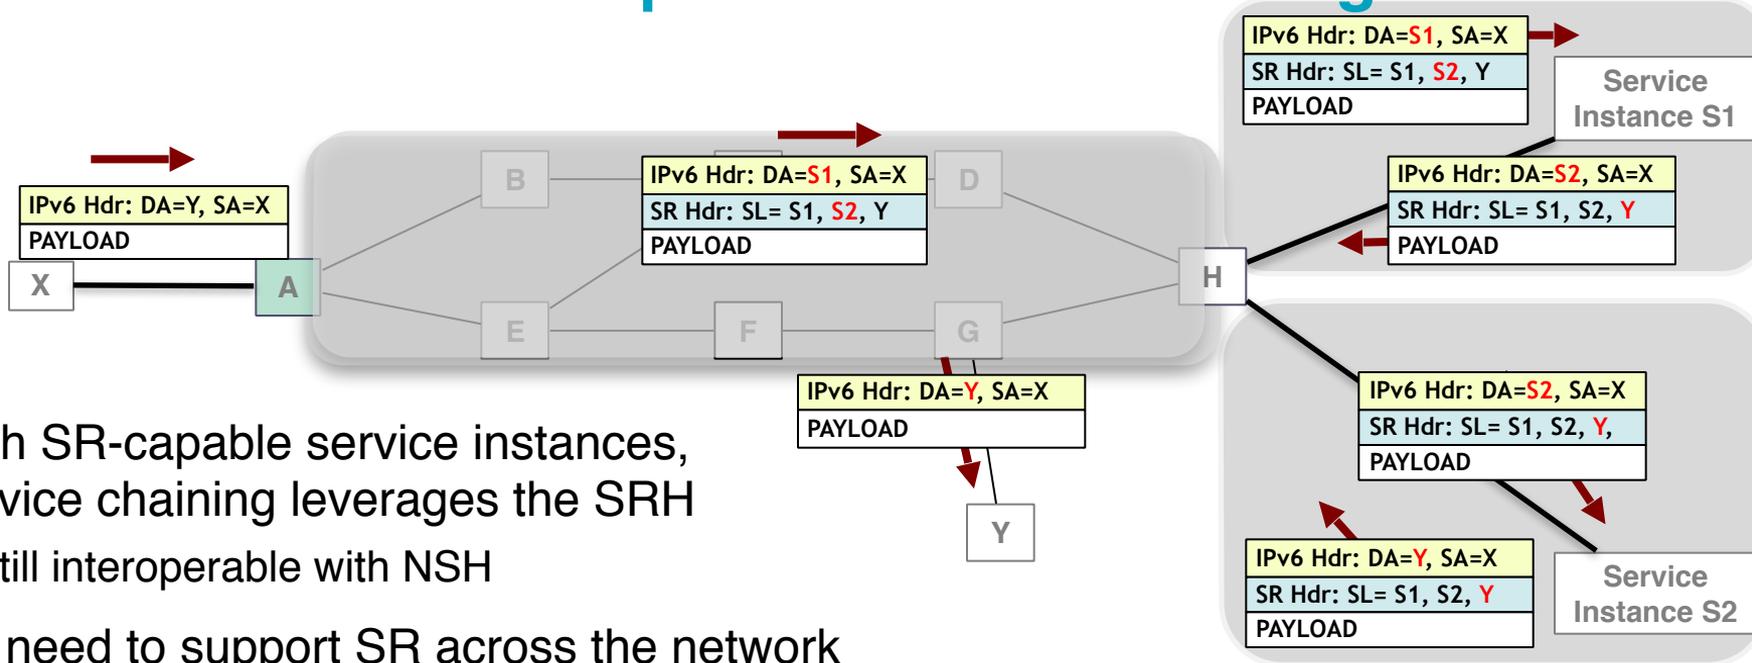
Segment Routing Use Cases Disjoint TE Service

- A to Z any plane
 - IGP shortest-path
 - PrefixSID of Z (65)
- A to Z via blue plane
 - SRTE policy pushes one additional segment “Blue Anycast” (111)
- Benefits
 - ECMP
 - No hop-by-hop signalling load and delay
 - No midpoint state



© 2014 Cisco and/or its affiliates. All rights reserved.

Use Cases: SR-IPv6 Capable Service Chaining



- With SR-capable service instances, service chaining leverages the SRH
 - Still interoperable with NSH
- No need to support SR across the network
 - Transparent to network infrastructure

Segment Routing Use Cases: Application driven traffic steering

- Impose source-routing semantics within an application or at the edge of a network (for example, a CPE or home gateway)
- CPE gets the SRH from a controller and impose it to outgoing traffic
- SRH includes HMAC that is going to be validated at ingress

IPv6 SPRING Use Cases

draft-ietf-spring-ipv6-use-cases-01

- Spring in the Home Network:
 - Address the problem of multiple egress points
- Spring in the Access Network:
 - Associate traffic flows to different capacity pipes
- Spring in the Data Center:
 - Build Service chaining
- Spring in the Content Delivery Network:
 - Handle Hierarchical caches
- Spring in the Core network:
 - Create Traffic Engineering paths (disjoint, latency, bandwidth, other constraints)

IPv6 Segment Routing Security

- Routing Type 0 (RH0) extension header has been deprecated by RFC5095
 - Reason: vulnerability (amplification attack) of RH0
- SRH defines an HMAC field to be used at ingress of a SR domain in order to validate the SRH
- Avoid malicious attempts to steer a packet out of its intended path
 - Amplification attack with RH-0
- Addresses concerns of RFC5095
- When used within the boundaries of a controlled domain, the HMAC is not necessary
- Similarly, IETF has standardized the Routing Extension Header type 3 (RPL) without any security mechanism
 - RH3 is assumed to be used within the boundaries of a private/controlled domain

IPv6-SR Summary

- Segment Routing implements the source routing model for both MPLS and IPv6
- IPv6 source routing model is already integrated in RC2460 and Segment Routing introduces minor changes through a new routing type header
 - Segment Routing Header
 - The SRH addresses the security requirements (HMAC)
- Segment Routing is very flexible and interoperable with non-SR nodes
 - Any IPv6 device can be enhanced to SR: a router, a server, any appliance, application, terminal,
- Segments are identified by IPv6 addresses, no specific signaling is needed
 - Router Segment: no need to advertise anything more in IGP – An IPv6 address may represent more than just a router, e.g.:
 - An interface, an appliance, a service, an application
 - A set of anything above

IPv6-SR References

- [1] Clarence Filsfils, et al. Segment Routing Architecture. draft-ietf-spring-segment-routing-02. IETF Draft, May 2015.
- [2] Clarence Filsfils, et al. Segment Routing with MPLS data plane. draft-ietf-spring-segment-routing-mpls-00. IETF Draft, Nov 2014.
- [3] Stefano Previdi, et al. IPv6 Segment Routing Header. draft-previdi-6man-segment-routing-header-06. IETF Draft, May 2015.
- [4] Stefano Previdi, et al. IS-IS Extensions for Segment Routing. draft-ietf-isis-segment-routing-extensions-04. IETF Draft, May 2015.
- [5] Peter Psenak, et al. OSPF Extensions for Segment Routing. draft-ietf-ospf-segment-routing-extensions-05. IETF Draft, June 2014.
- [6] Peter Psenak et al, OSPFv3 Extensions for Segment Routing. draft-psenak-ospf-segment-routing-ospfv3-extension-02, IETF Draft, July 2014
- [7] Clarence Filsfils, et al. Segment Routing Use Cases. draft-filsfils-spring-segment-routing-use-cases-01. IETF Draft, 2014
- [8] J. Brzozowski et al, IPv6 SPRING Use Cases. draft-ietf-spring-ipv6-use-cases-04, IETF Draft, March 2015
- [9] P. Francois et al, Topology Independent Fast Reroute using Segment Routing, draft-francois-spring-segment-routing-ti-lfa-01, IETF Draft, Oct 2014
- [10] E. Vyncke et al, IPv6 Segment Routing Security Considerations. draft-vyncke-6man-segment-routing-security-02, IETF Draft, Feb 2015
- [11] IPv6 Segment Routing implementation in Linux kernel,
<http://www.segment-routing.org/sr6-doc.pdf>
<https://github.com/segment-routing/sr-ipv6>

Thank you.

<http://www.segment-routing.net>

Gianpietro Lavado

glch@cisco.com

 @gianpietro_lc

 /gianpietro



Perú NOG

Estás en Perú y operas una red de telecomunicaciones?

Inscríbete en **el Perú Network Operators Group (Perú NOG)**

<https://mail.lacnic.net/mailman/listinfo/peru-nog>

Empecemos a intercambiar información y experiencias!