SPRING WG

Segment Routing IPv6

Authors and contributors:

Stefano Previdi (sprevidi@cisco.com)
Clarence Filfsils (cfilsfil@cisco.com)
Brian Field (Brian_Field@cable.comcast.com)
John Brzozowski (john_brzozowski@cable.comcast.com)
John Leddy (John_Leddy@cable.comcast.com)
Ida Leung (Ida.Leung@rci.rogers.com)
Roberta Maglione (robmgl@cisco.com)
Eric Vyncke (evyncke@cisco.com)
Dave Barach (dbarach@cisco.com)
Mark Townsley (townsley@cisco.com)
Chris Martin (martincj@cisco.com)
Nagendra Kumar (naikumar@cisco.com)
David Lebrun (david.lebrun@uclouvain.be)
Pierre Francois (pierre.francois@imdea.org)
James Connolly (jconnolly@libertyglobal.com)

IETF91 – Honolulu, November 2014

Terminology

- SR-MPLS
  - Segment Routing applied to the MPLS dataplane
- SR-MPLS-IPv4
  - Segment Routing applied to the MPLS dataplane with an IPv4 control plane
- SR-MPLS-IPv6
  - Segment Routing applied to the MPLS dataplane with an IPv6 control plane
- SR-IPv6
  - Segment Routing applied to the IPv6 dataplane with an IPv6 control plane

Segment Routing in IPv6 Networks

- Source based routing model where the source chooses a path and encodes it in the packet header as an ordered list of segments
- Segment Routing leverages the source routing architecture defined in RFC2460 for IPv6
  - By “source we intend the originator of the packet or the ingress node of the SR domain”
- A segment is an instruction applied to the packet. Segments can represent any type of instruction:
  - IGP-based forwarding construct
  - BGP-based forwarding construct
  - local adjacency
  - service/application
  - location,
  - context, …
- The Segment Routing architecture is applicable to both MPLS and IPv6 dataplanes
  - In SR-MPLS a node has a segment identifier which is mapped to a (e.g.: local) label.
  - In SR-IPv6 a node has a segment identifier which is in fact it’s loopback IPv6 address

Current SR-IPv6 Drafts

- draft-ietf-spring-ipv6-use-cases (SPRING WG)
  - describes the SR-IPv6 use cases
- draft-previdi-6man-segment-routing-header (6MAN WG)
  - describes a new type of the Routing Header (SRH)
  - Will be presented in 6man wg meeting this friday
- draft-vyncke-6man-segment-routing-security (6MAN WG)
  - describes the security mechanisms applied to the SRH
  - Will be presented 6man wg meeting this friday

Segment Routing Header (SRH)

Segment Routing for IPv6 Dataplane

- A Segment is identified through its IPv6 address
  - No mapping needed between SIDs and node's addresses
  - Simplifies signaling, address == SID
- New Routing Extensions Header type
  - Segment Routing Header (SRH)
  - Contains Segment List, Policy List, and a few other bits…

SRH

- Next Header: 8-bit selector. Identifies the type of header immediately following the SRH
- Hdr Ext Len: 8-bit unsigned integer. Defines the length of the SRH header in 8-octet units, not including the first 8 octets
- Type: TBD (SRH)
- Next Segment: index, in the Segment List, of the next active segment in the SRH
- Last Segment: index, in the Segment List, of the last segment of the path
- Flags: 16 bits of flags. Following flags are defined:
  - bit-0: cleanup
  - bit-1: rerouted packet
  - bits 2 and 3: reserved
  - bits 4 to 15: policy flags. Define the type of the IPv6 addresses encoded into the Policy List (each address is described by 3 bits):
    - 0x0: Not present
    - 0x1: ingress SR
    - 0x2: egress SR
    - 0x3: original source address

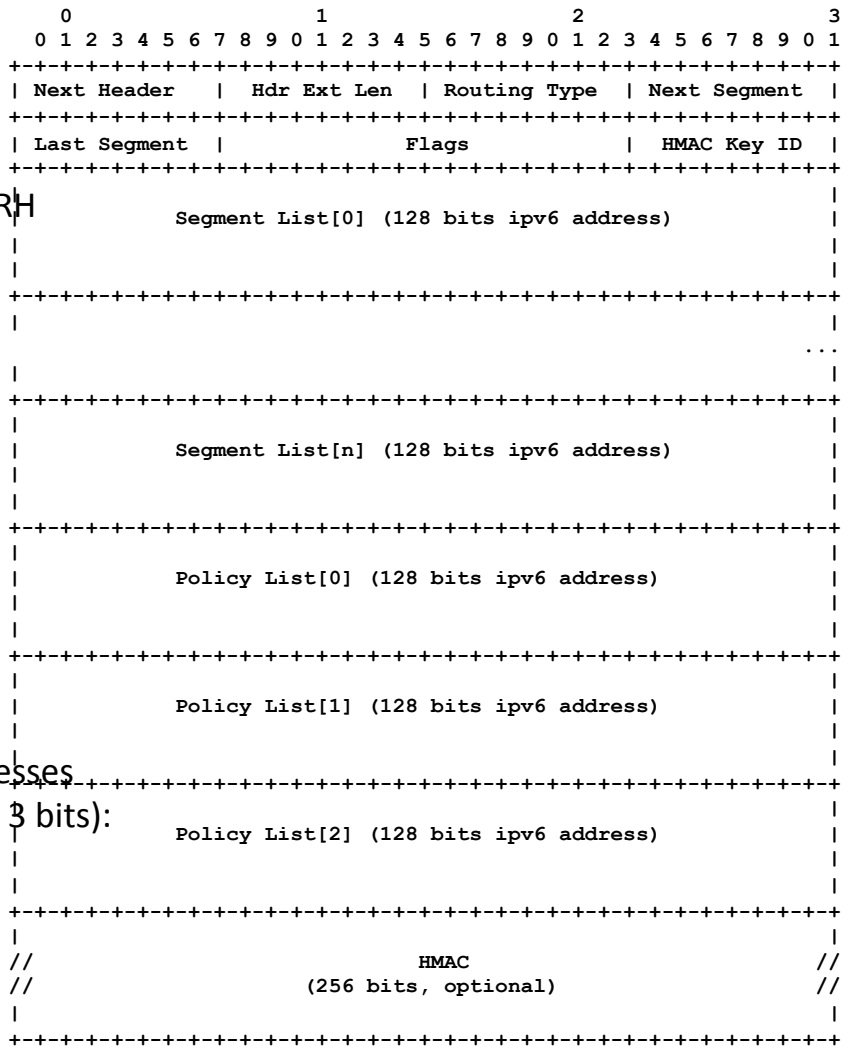

SRH

- Segment List[n]: 128 bit IPv6 addresses representing each segment of the path
- Policy List[n]: Specific nodes in the SR path:
  Ingress SR: 128 bit identifier representing the ingress in the SR domain
  Egress SR: 128 bit identifier representing the egress in the SR domain
  Original Source Address: IPv6 address originally present in the SA field of the packet
- HMAC: SRH security (optional)

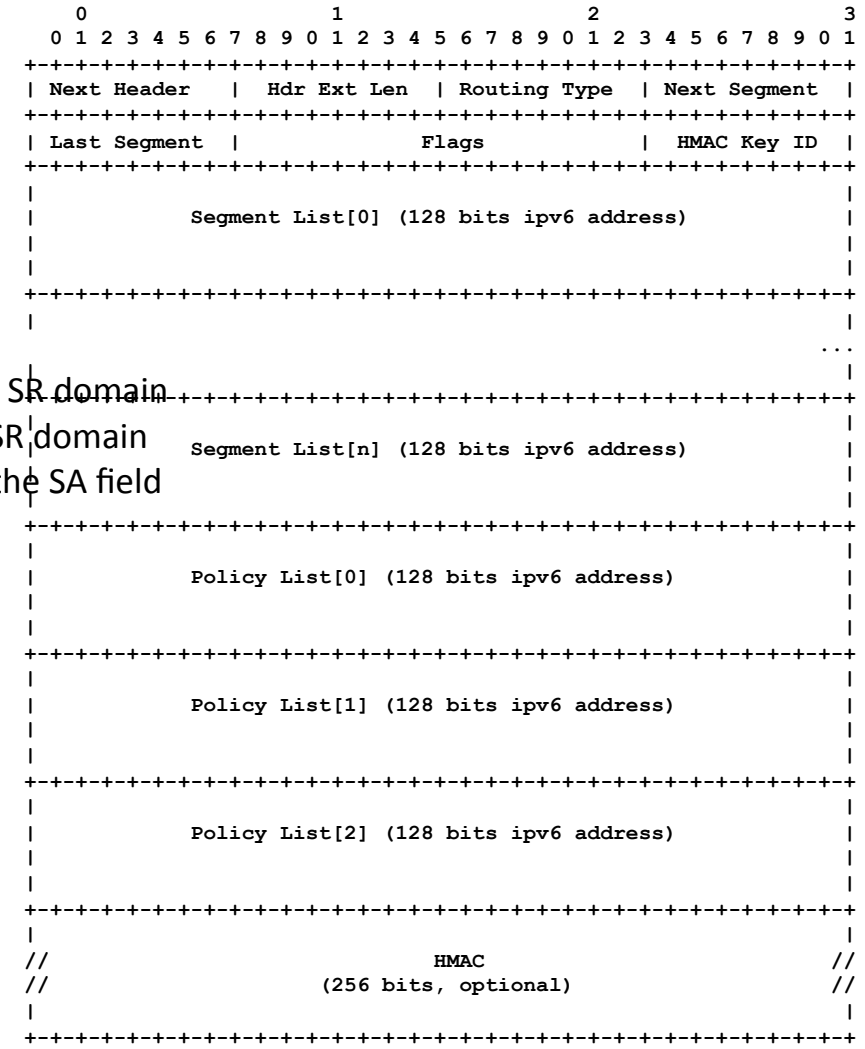

Segment Routing Header

- SRH is a new type of the existing routing header. Therefore, it inherits routing header properties:
  - Can only appear once
  - It is inspected by the node in the DA
  - It MUST be ignored by any other node
- SRH format is similar to RH0 that has been deprecated
  - Carry ipv6 addresses
  - Segments (SL and PL)
  - Security: HMAC
  - "Last Segment" field
- Deprecation has been motivated by security concerns
  - SRH address them through deployment guidelines and HMAC
  - draft-vyncke-6man-segment-routing-security

Examples of SRH use

SR-IPv6 Example

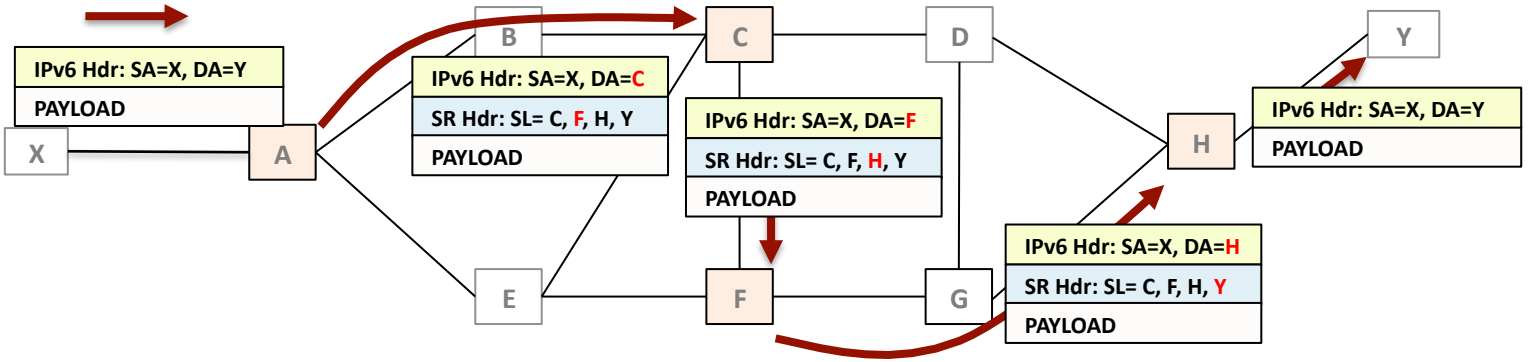

- SRH is inspected at the segment endpoint
- Non SR nodes fully interoperate (plain ipv6 routing)
- SRH is removed prior to deliver the packet to a non-SR destination
- A segment can represent a portion of shortest path as well as a hop in a service chain

Overlays with SR-IPv6

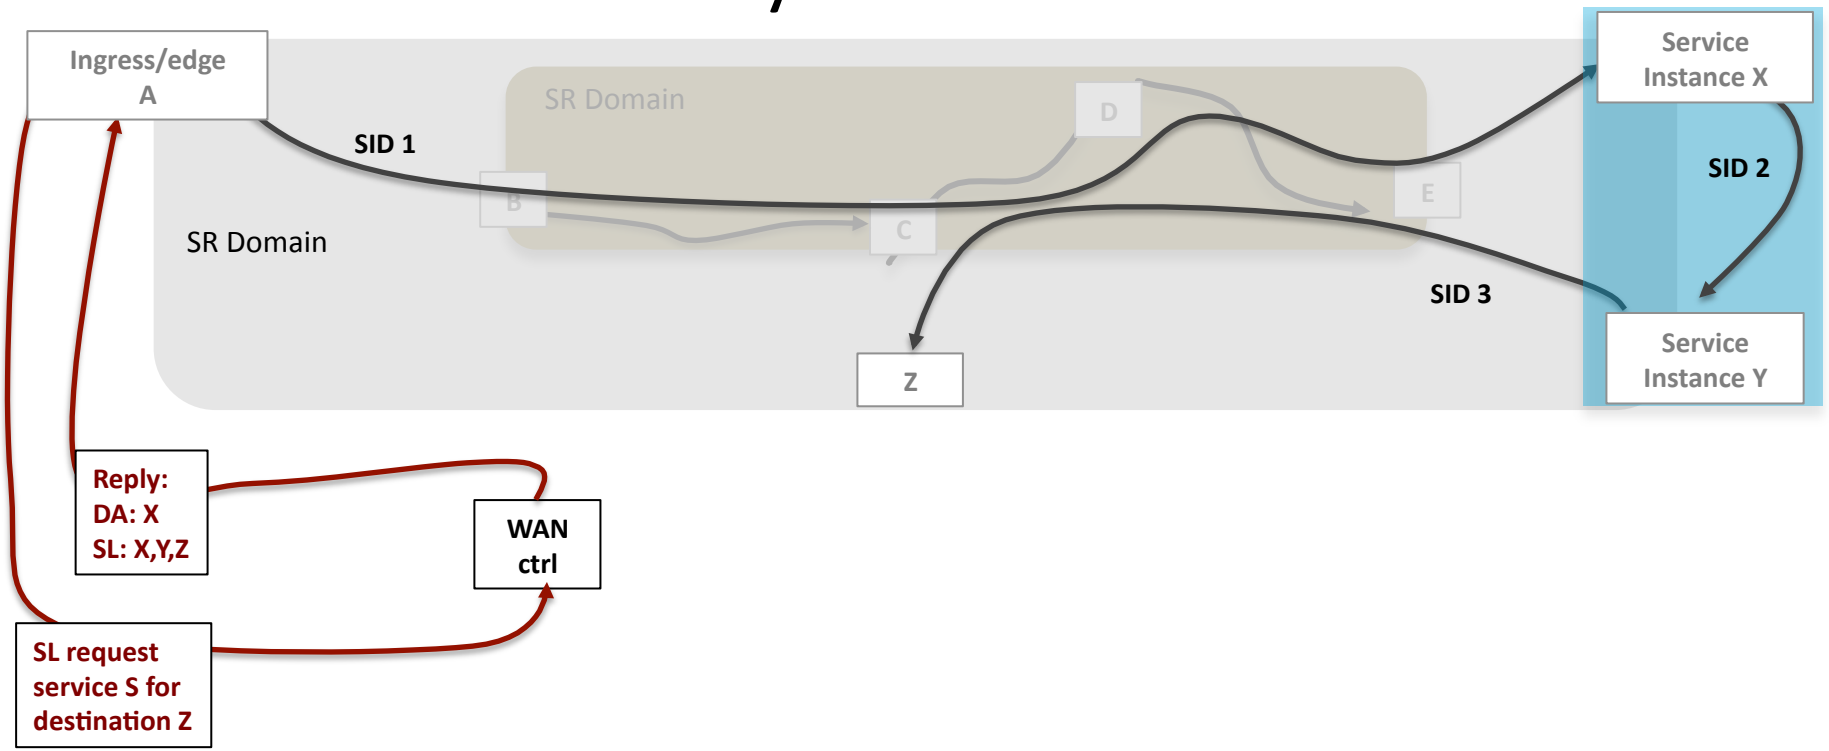

- X, Y, Z are instances of service S
- Host A requests service S for traffic destined to Z
- Ctrl returns the SL according to the definition of service S for host A

Overlays SR-IPv6

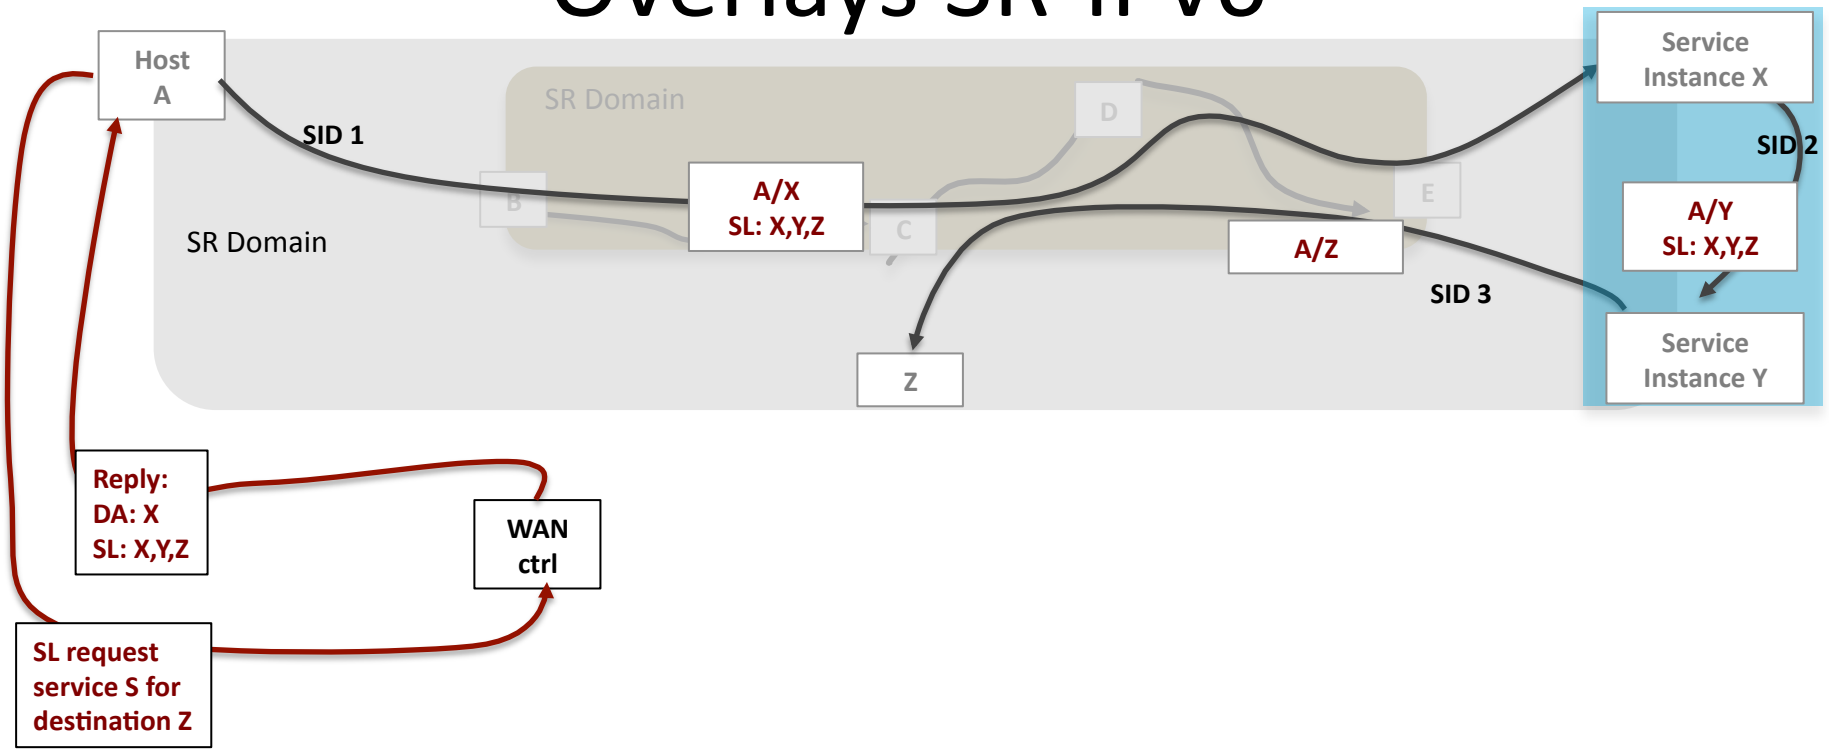

- X, Y, Z are instances of service S
- Host A requests service S for traffic destined to Z
- Ctrl returns the SL according to the definition of service S for host A
- SL reflects the instances of service S (X,Y) and destination (Z)

Implementations

- Multiple implementations exist and interoperability has been demonstrated during IETF90
  - Based on draft-previdi-6man-segment-routing-header-03.txt
  - Cisco,
    Comcast,
    Ecole Polytechnique (Paris),
    UCLouvain (LLN, Belgium)
- Demonstrated interoperability between multiple, independent IPv6 Segment Routing implementations (routers and hosts)
- Illustrate interoperability between SR and non-SR capable routers and hosts
- Illustrate how SR can be leveraged for video content delivery through SR capable caches

Implementations

- SR-IPv6 for video content delivery

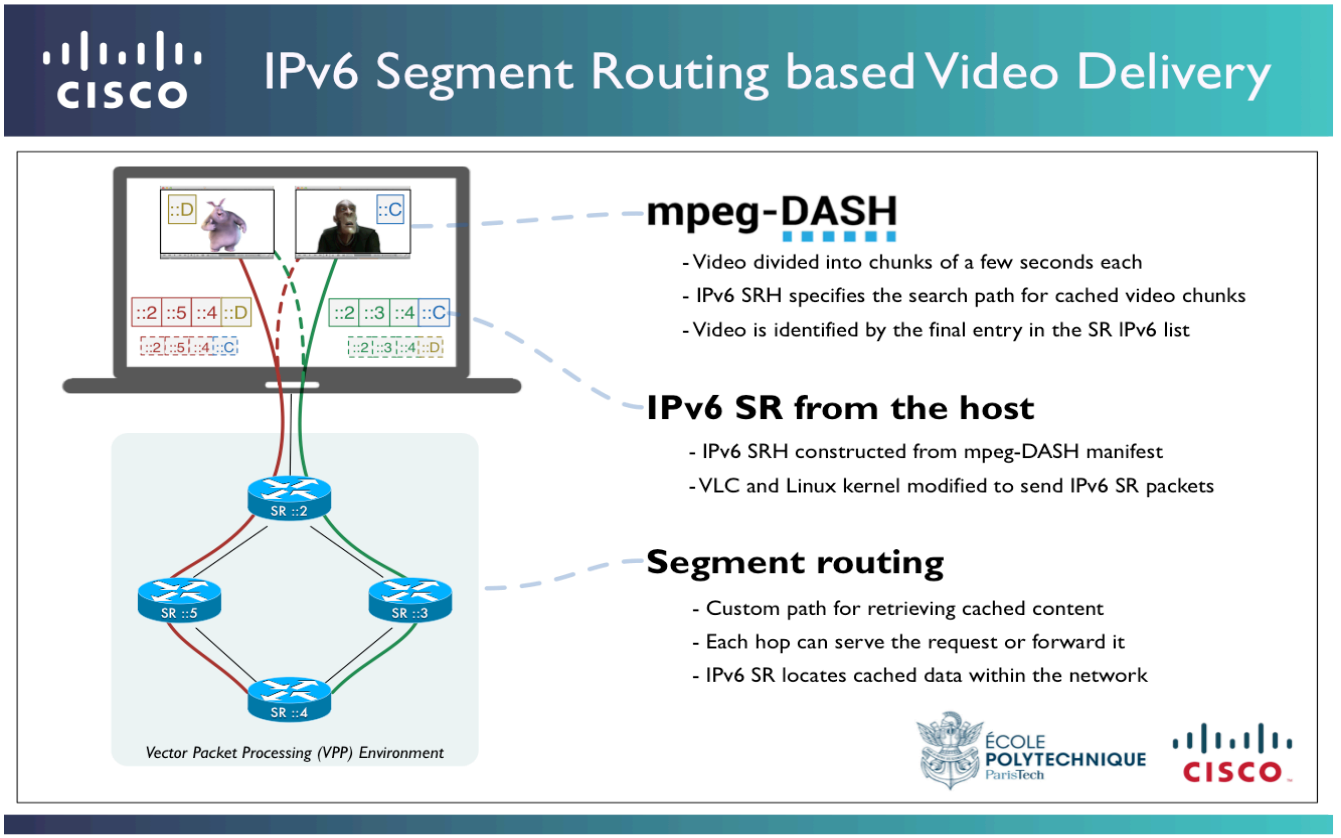

Open Stack VM Isolation in IPv6 (3 choices, one with IPv6 SR)

Within a data center, tenants/projects are given group numbers

Senders have packets labeled with their group identity

Receivers exclude packets they are not authorized to receive

1. IPv6 flow label used as a flat space of 2^{20} group identities
2. **Packets labeled with the IPv6 SR Ingress ID (policy-list), directed via SR to Egress ID and validated accordingly**
3. Use a destination or hop-by-hop option to express a potentially-federated group identity

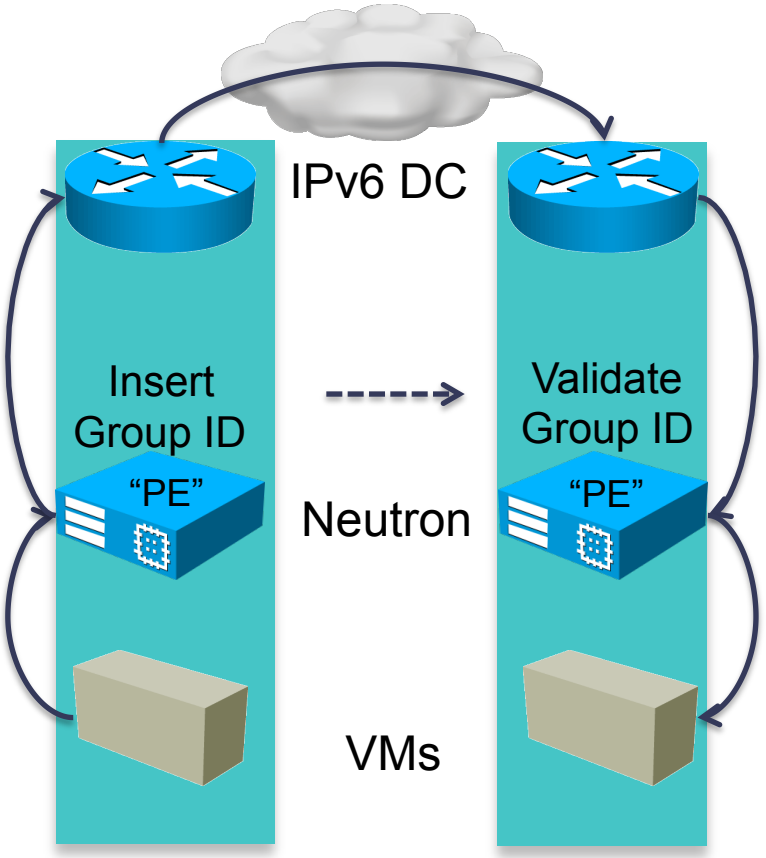

Ref: draft-baker-openstack-ipv6-model-00

Questions?

Thanks!