SRV6+

Ron Bonica

NANOG 77, October 28, 2019

Juniper Networks | Engineering Simplicity

ALIAS

- Recently renamed in the IETF
- Also know as SR-mapped-6
  - Because Segment Identifiers (SIDs) are mapped to IPv6 addresses
- Also knows as SRm6
  - Because every technology needs an acronym !

SEGMENT ROUTING (SR) TODAY

- Source nodes steer packets through SR paths
  - The source node encodes the SR path in each packet
  - This alleviates the need for transit nodes to maintain path state
- SR-MPLS
  - Encodes the SR path in an MPLS label stack
  - Each label stack entry represents a segment in the SR path

SR AND IPV6

- SRv6 and SRv6+
  - Encode the SR path in an IPv6 Routing header
  - Each member of the Routing header's Segment List represents a segment in the SR path
- Heated debate between SRv6 and SRv6+ in IETF SPRING WG
  - Each has a unique encoding mechanism
  - Each has a unique approach to network programming and IPv6 extensibility

SRv6+ Preserves The Best of SRv6

TRAFFIC-ENGINEERED TUNNELS – PATHS AND SEGMENTS

- SR Path
  - Provides unidirectional connectivity from its ingress node to its egress node
  - Can follow least-cost or traffic-engineered path
  - Defined by the segments that it contains
- Segment
  - Provides unidirectional connectivity from its ingress node to its egress node
  - Defined by its ingress node and the topological instruction that controls it

TRAFFIC-ENGINEERED TUNNELS – TOPOLOGICAL INSTRUCTIONS

- Topological Instruction
  - Executes on segment ingress node
  - Defines segment egress node and method by which ingress node forwards packets to egress node
  - Types
    - Adjacency – forward packets through a specified interface
    - Node – forward packets through the least cost path to egress node
    - Binding – forwards packets through another tunnel (e.g., another SRv6+ path)

TRAFFIC-ENGINEERED TUNNELS – SEGMENT IDENTIFIERS (SID)

- Identifies a segment - 16 or 32-bit unsigned integer
  - Major difference between SRv6 and SRv6+
- Because there is a one-to-one relationship between a segment and the topological instruction that controls it, the SID that identifies a segment also identifies the topological instruction that controls it
- Node-local significance
- Distributed by IS-IS / OSPF

NETWORK PROGRAMMABILITY

- Service instructions augment SR paths
- Per segment service instructions
  - Execute on segment egress node
  - Examples
    - Expose a packet to a firewall policy
    - Expose a packet to a sampling policy

NETWORK PROGRAMMABILITY - CONTINUED

- Per Path
  - Executed on path egress node
  - Generally, override IPv6 default payload handling rules
  - Examples
    - Decapsulate payload and forward through a specified interface
    - Decapsulate payload and forward using a specified forwarding table
- Service Instruction augment, but do not define SR paths
  - All packets traversing an SR path do not need to be augmented by the same service instructions

PULLING IT ALL TOGETHER

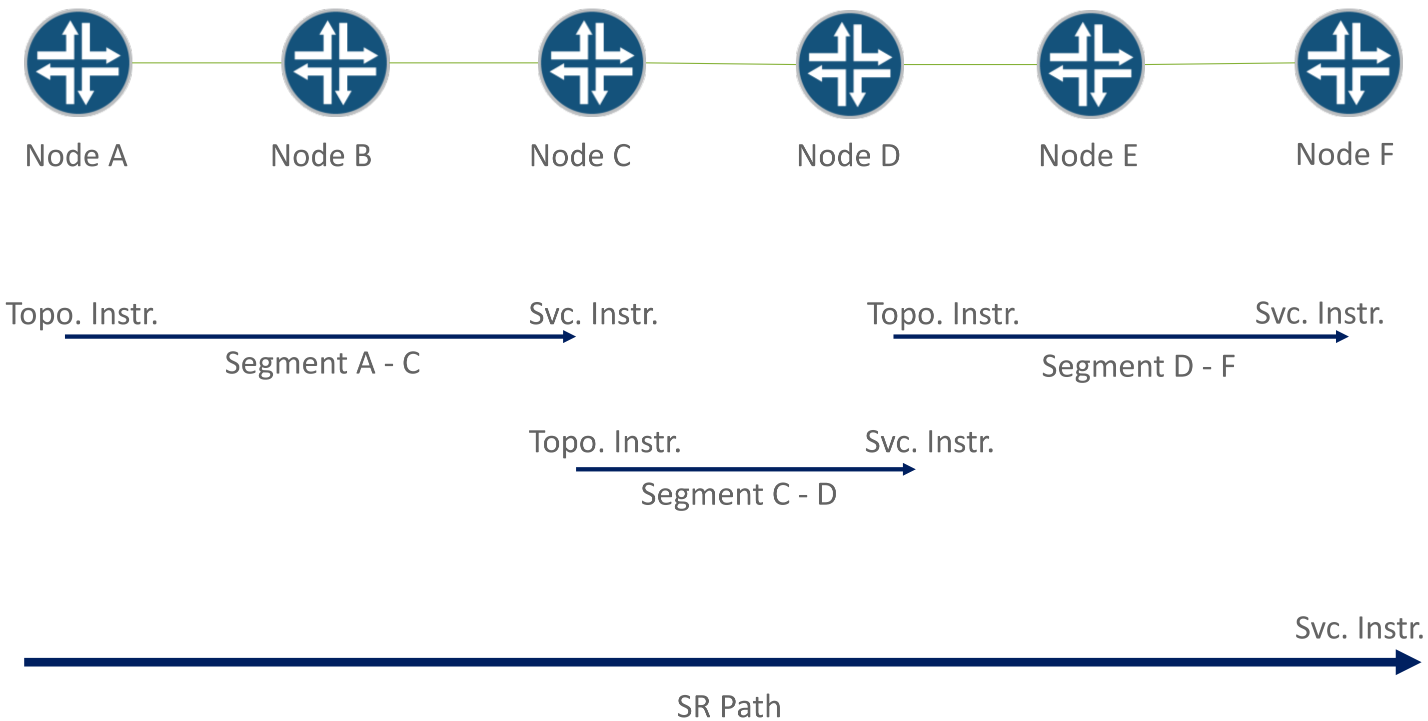

SRv6+ Leverages IPv6 Extensibility

IPV6 EXTENSION HEADERS

PER HOP (PROCESSED BY EVERY NODE)	PER SEGMENT (PROCESSED BY EVERY SEGMENT EGRESS NODE)	PER PATH (PROCESSED BY THE PATH EGRESS NODE)
• Hop-by-hop	• Destination Options • Routing	• Fragment, AH, ESP • Destination Options
Process These First	Process These Next	Process These Last

- SRv6+ packets can contain any valid combination of IPv6 extension headers
- SRv6+ packets leverage the Routing and Destination Options headers

IPV6 ROUTING HEADER [RFC 8200]

- "The Routing header is used by an IPv6 source to list one or more intermediate nodes to be "visited" on the way to a packet's destination."
- Fields
  - Next Header: Identifies the type of header immediately following the Routing header
  - Hdr Ext Len: Length of the Routing header
  - Routing Type: Identifies a Routing header variant
  - Segments Left: Number of segments still to be visited
  - Type Specific Data: Segment List

IPV6 ROUTING HEADER [RFC 8200] - CONTINUED

- Processing
  - If Segments Left is zero, skip and proceed to the next header
  - If Segments Left is greater than zero, decrement Segments Left, update IPv6 Destination Address to reflect next segment egress node and forward packet. Do not process next header.

SRV6+ COMPRESSED ROUTING HEADER

- SRv6+ defines two new Routing header types
  - Resembles the generic Routing header defined in RFC 8200
  - Same five fields
  - Same simple function
- CRH-16 & CRH-32
  - 16-bit SID versus 32-bit SID
- A local data structure, called the SID Forwarding Information Base (SFIB), maps 16 or 32-bit SIDs to IPv6 addresses and forwarding instructions

COMPRESSED ROUTING HEADER: 16-BIT VERSION

No Surprises Here

```
  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 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             SID[0]            |             SID[1]            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-|
|                           .........
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
```

Except That The SID Is

A Pointer Into the SFIB

COMPRESSED ROUTING HEADER: 32-BIT VERSION

No Surprises Here, Either

```
  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 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             SID[0]                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-|
                            .........
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             SID[N]                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-|
```

Except That The SID Is

A Pointer Into the SFIB

SID FORWARDING INFORMATION BASE (SFIB)

- Indexed by SID
- Contains the following information
  - The IPv6 address of an interface on the segment egress node
    - When CRH is processed, this address is copied to the IPv6 Dest Address
  - Segment Type
    - Adjacency – forward through a specified interface
    - Node – forward through the least cost path
    - Binding – forward through another tunnel (e.g., another SRv6+ path)
  - Type Specific Information
    - Adjacency: Interface though which packet is to be forwarded
    - Node: none

CRH ANALYSIS

- The CRH is used by an IPv6 source to list one or more intermediate nodes to be “visited” on the way to a packet’s destination
  - Nothing more, nothing less
- The CRH is relatively short because it represents intermediate nodes as 16 or 32-bit SIDs
  - Not 128-bit identifiers
- Because the CRH is relatively short, it can be supported in ASICS
- Because the CRH is relatively short, it does not consume excessive bandwidth

IPV6 DESTINATION OPTIONS HEADER [RFC 8200]

- “The Destination Options header is used to carry optional information that need be examined only by a packets destination node(s)”
- A packet may contain two Destination Options headers
  - Preceding the Routing header
    - Processed by every segment egress node including the path egress node
  - Preceding the upper-layer header
    - Processed by the path egress node

SRV6+ DEFINES TWO NEW IPV6 OPTIONS

- Per Segment Service Instruction Identifier
  - Encoded in a Destination Options header that precedes the Routing header
  - Identifies a Service Function Chain
    - Each node learns the function that it executes in this context from a controller
  - Network-wide significance
  - Optional
    - If processing node does not recognize the option, ignore and process next option or header

SRV6+ DEFINES TWO NEW IPV6 OPTIONS - CONTINUED

- Per Path Service Instruction Identifier
  - Encoded in a Destination Options header that precedes the upper-layer header
  - Identifies a per path instruction
  - Node-local significance
    - Allocated by processing node
    - Advertise by BGP
      - Replaces MPLS service label
  - Processing is not optional

Avoiding The Issues Of SRv6

SRV6 MODIFIES IPV6 ADDRESS SEMANTICS

- The IPv6 Destination Address field can contain either an IPv6 Address [RFC 4291] or an SRv6 SID
  - SIDs are copied from SRv6 Segment Routing Header (SRH) to IPv6 Destination Address
- SRv6 SID semantics differ from IPv6 address semantics and SRv6+ SID semantics
  - High-order bits: Locator
  - Low-order bits: Function and arguments
- Locator identifies instantiating node

SRV6 MODIFIES IPV6 ADDRESS SEMANTICS - CONTINUED

- Function and arguments define processing rules
  - If Segments Left > 0, the function and arguments determine how the Routing header is processed
  - If Segments Left == 0, the function and arguments determine how the upper-layer header is processed
- Many functions defined, more to be defined
  - Some functions are valid only when Segments Left equal to 0
  - Some functions are valid only when Segments Left is greater than 0

SRV6 COUPLES TOPOLOGICAL AND SERVICE INSTRUCTIONS

- Both are encoded in the SRv6 SRH
  - IPv6 provides the Routing header for topological instruction
  - IPv6 provides the Destination Options header for service instructions
- Routing header expands to carry flags, tags and TLVs
  - Mostly contain service instruction details
  - More flags, tags and TLV being proposed
  - Routing header scope creep / ASIC impact
- New security mechanisms required
  - Routing header cannot be protected by IPv6 Authentication or Encapsulating Security Payload headers
  - SRH HMAC TLV required

ROUTING HEADER SIZE – IMPEDIMENT TO SRV6 DEPLOYMENT

Number of SIDs	SRv6 SRH	SRv6+ CRH-32	SRv6+ CRH-16
1	24	8	8
2	40	16	8
3	56	16	16
4	72	24	16
5	80	24	16
6	104	32	16
7	120	32	32
8	136	40	32
9	152	40	32
10	168	48	32
11	184	48	40
12	200	56	40

SRV6 MODIFIES IPV6 ADDRESS SEMANTICS FURTHER TO ADDRESS ROUTING HEADER SIZE

- New SID type: uSID
- Network operator obtains a /32
  - Example: 2001:db8::/32
- Network operator assigns a 48-bit uSID to each node
  - Example: 2001:db8:n::/48 is assigned to Node n
- The path from any source to Node 1 via Nodes 3, and 2 is represented by 2001:db8:3:2:1::/128

SRV6 MODIFIES IPV6 ADDRESS SEMANTICS FURTHER TO ADDRESS ROUTING HEADER SIZE - CONTINUED

- Processing at each segment egress node
  - Node 3 receives Destination Address 2001:db8:3:2:1::/128 and maintains a uSID for 2001:db8:3/48
    - Right-shifts the lower 96 bits of Destination Address by 16 bits and forwards
  - Node 2 receives Destination Address 2001:db8:2:1::/128 and maintains a uSID for 2001:db8:2/48
    - Right-shifts lower 96 bits of Destination Address by 16 bits and forwards
  - Node 1 receives Destination Address 2001:db8:1::/128 and maintains a uSID for 2001:db8:1/48
    - Processes next header

USID IMPEDIMENTS TO DEPLOYMENT

- Network operator must obtain a /32
- Network operator must maintain a specific numbering policy in order to make the solution work
  - If a /48 represents a uSID, no more specific can be configured and advertised
- Further modifies the IPv6 address semantics

Implementations And Deployment

SRV6+ IMPLEMENTATION AND DEPLOYMENT STATUS

- PoC Implementations
  - LINUX (forwarding plane plus JUNOS cRPD)
  - JUNOS (MX-based POC may be complete by time of NANOG presentation)
  - Liquid Telecom Data Plane Developers Kit (DPDK)
- Interoperability demonstrated between LINUX and Liquid DPDK
- Experimental Deployments
  - Liquid Telecom
  - More to come in 4Q19

Conclusion

CONCLUSION

- SRv6+ preserves the best of SRv6
  - Traffic-engineering
  - Network programmability
- SRv6+ is deployable
  - Because it defines as little new protocol machinery as possible
  - Because it leverages all relevant and available IPv6 extensibility mechanisms
    - The Routing header and Destination Options header
  - Because it does not attempt to redefine fundamental IPv6 concepts
    - IPv6 address semantics
    - Acceptable use of the Routing header


THANK YOU
JUNIPER NETWORKS
Engineering
Simplicity