



I E T F[®]

IETF 98 - Chicago
Mar 2017

SRv6 Network Programming

(draft-filsfils-spring-srv6-network-programming-00)

C. Filsfils (Cisco)

J. Leddy (Comcast)

D. Voyer (Bell Canada)

D. Bernier (Bell Canada)

D. Steinberg (Steinberg Consulting)

R. Raszuk (Bloomberg LP)

S. Matsushima (SoftBank Telecom)

D. Lebrun (Universite catholique de Louvain)

B. Decraene (Orange)

B. Peirens (Proximus)

S. Salsano (Universita di Roma "Tor Vergata")

G. Naik (Drexel University)

H. Elmalky (Ericsson)

P. Jonnalagadda (Barefoot Network)

M. Sharif (Barefoot Networks)

A. Ayyangar (Arista)

S. Mynam (Dell Force10 Networks)

A. Bashandy (Cisco)

K. Raza (Cisco) >> Prsenter

D. Dukes (Cisco)

F. Clad (Cisco)

P. Camarillo, Ed. (Cisco)

Introduction

“SRv6 network programming” refers to the capability for an application to encode any complex program as a set of individual functions distributed through the SRv6 network.

Introduction (2)

- This draft is the “key” SRv6 document that describes SRv6 network programming concepts, its various functions, and their use cases:
 - Local-SID Functions, Transit Behavior
 - Control Plane
 - Counters, Security
 - Use case illustrations

- Status:
 - Larger community support (from vendors and operators)
 - Multiple interoperable implementations
 - Open Software Projects: <http://www.segment-routing.net/open-software/> (VPP 17.04 and Linux Kernel 4.10)

Local SID

- A local SID has a specific instruction bound to it.
- An SRv6-capable node N maintains a table containing all the local SRv6 segments explicitly instantiated at node N.
 - N is the parent node for these SIDs.
- A local SID of N could be routed to N but it does not have to be. Most often, it is routed to N via a shorter-mask prefix.

Local SID (2)

- SRv6 local SID is represented as LOC:FUNCT
 - LOC is the L most significant bits
 - FUNCT is the (128-L) least significant bits.
 - L is called the locator length and is flexible:
 - no assumption on size/length
- Most often the LOC part of the SID is routable and leads to the node which owns that SID.
- The FUNCT part of the SID is an opaque identification of a local function bound to the SID. Hence the name SRv6 “Local” SID.
 - LOC:FUNCT:ARGS if function requires argument(s)

Local SID Functions

- This draft defines a set of well-known functions that can be associated with a local SID.
 - For each function, packet processing algorithm is also documented at a high level

Local SID Functions (2)

Name	Forwarding	Use case
End *	Lookup	Prefix SID
End.X *	L3 Xconnect	Adj SID
End.T *	Lookup in table T	Multi-table operation in the core
End.DT6	Decap and IPv6 table T lookup	IPv6 L3VPN - Per-VRF
End.DT4	Decap and IPv4 table T lookup	IPv4 L3VPN - Per-VRF
End.DX6	Decap and IPv6 Xconn	IPv6 L3VPN - Per-CE
End.DX4	Decap and IPv4 Xconn	IPv4 L3VPN - Per-CE
End.DX2	Decap and L2 Xconn	L2VPN

*: With variants

Local SID Functions (3)

Name	Forwarding	Use case
End.B6	SRv6 policy	Binding SID
End.B6.Encaps	SRv6 policy (with encap)	Binding SID
End.BM	SR-MPLS policy	Binding SID
End.S	Search of a target (Locally forward or END behavior)	ICN
End.AS	Remove Outer IPv6 header and SRH, forward to interface	Service Chaining via an SR-unaware App
End.AM	Update Outer IPv6 header DA with LAST SID and forward to interface	Service Chaining via an SR-unaware App (with masquerade)

SRH Pop

- “SRH Pop” refers to removal (pop) of the “top” SRH in a received SRv6 packet at an endpoint.
- We define SRH popping for the following functions:
 - End, End.X, and End.T
- Flavors:
 - Two variants:
 - **Ultimate Segment Pop (USP)** : SRH Popped at last segment
 - **Penultimate Segment Pop (PSP)**: SRH Popped at penultimate segment
 - For each of the above End functions, these variants can be enabled or disabled either individually or together.

Transit Behaviors

- Transit node: A node that receives an IPv6/SRv6 packet whose DA is neither local address nor local SID

Namen	Behavior
T	Pure Transit
T.Insert	Insert an SRv6 policy
T.Encaps	Encap an SRv6 policy
T.Encaps.L2	Encap an SRv6 policy on L2 frame

Control Plane

- The following table summarizes which SID would be signaled in which signaling protocol

Name	IGP	BGP-IP/VPN	BGP-LS
End	X		X
End.X	X		X
End.T	X		X
End.DT6		X	X
End.DT4		X	X
End.DX6		X	X
End.DX4		X	X
End.DX2		X	X
End.BM			X
End.S			X
End.AS			X
End.AM			X
T			X
T.Insert			X
T.Encaps			X
T.Encaps.L2			X

Counters and Security

- Counters:
 - Local SID - Matched and processed correctly/incorrectly
 - SR policy - Steered into and processed correctly/incorrectly

- Security:
 - “How a domain of trust can operate SRv6-based services for internal traffic while preventing any external traffic from accessing these internal SRv6-based services.”
 - Some mechanisms:
 - ACL on the external interface to drop any traffic with SA or DA in the internal SID space
 - ACL to prevent access to local SIDs from outside the operator's infrastructure
 - An SRv6 router MUST only implement the End behavior on a local IPv6 address if that address has been explicitly enabled as a segment (local SID)
 - Support Unicast-RPF on source address on external interface

Use Case Illustrations

- Basic Security
- SR-L3VPN
- SR-L2VPN-VPWS
- SRTE for Underlay SLAs
 - Policy @ ingress PE
 - Policy @ mid
- End-to-end SRTE policy
- TI-LFA
- SRTE for Service Chaining

Draft: Next Steps

- Seeking WG input and feedback
- Comments and suggestions are welcomed !!!