

# Segment Routing

IETF 87

Clarence Filsfils – [cf@cisco.com](mailto:cf@cisco.com)

C. Filsfils, S. Previdi, A. Bashandy, B. Decraene, S. Litkowski, M. Horneffer,  
I. Milojevic, R. Shakir, S. Ytti, W. Henderickx, J. Tantsura, Ericsson, E. Crabbe,  
H. Gredler, and a few other contributors...

# Technology

- Generality of a segment
  - Intra and inter domain
  - Forwarding construct
  - Service construct
  - virtualization
  - Abstract Routing Model (draft-filsfils-rtgwg-segment-routing), see nanog video
  - SR is not a label-in-IGP solution. Label-in-IGP is a subset of SR !
- Agnostic Control Plane
- Instantiation in two dataplanes
  - MPLS
  - IPv6

# Productization

- Wide and rapid industry adoption
- Committed deployments received from operators within 5 months of first review
  - MPLS
  - IPv6

# The last 9 months

- Oct: first SR presentation to operators
  - Lead Operator group formed, see co-authors, weekly meeting since then
  - Commitment to velocity, transparency and multi-vendor agreement
- Feb: initial implementation released as per commitment
- Mar: MPLS WC and IPv6 conference
- Mar: first draft submitted to IETF-86
  - Multi-vendor technology agreement and interoperability plans (cisco, Alcatel and Ericsson)
  - draft-gredler-... co-authors want more details as draft. We commit to detailed drafts by end of May
- May
  - Team shares 6 detailed drafts with draft-gredler-... co-authors and seek merge agreement
- June
  - Merge agreement on a subset of SR: MPLS/SR instantiation, use-cases, FRR, ISIS and OSPF: great collaborative work
- The point:
  - Detailed and thoughtful work
  - Velocity
  - Collaborative
  - Commitments met

# Next few months - IETF

Topic	IETF Reference	WG
Abstract Routing Model	draft-filsfils-rtgwg-segment-routing	RTGWG
MPLS Instantiation	New draft to be submitted <small>(based on section 5 of draft-filsfils-rtgwg-segment-routing)</small>	MPLS
IPv6 Instantiation	New draft to be submitted	IPv6
Use Cases	draft-filsfils-rtgwg-segment-routing-use-cases	RTGWG
Perf Eng. LSP with SR	draft-shakir-rtgwg-sr-performance-engineered-lsps	RTGWG
ISIS SR Extensions	draft-previdi-isis-segment-routing-extensions	ISIS
OSPF SR Extensions	draft-psenak-ospf-segment-routing-extensions	OSPF
FRR SR	draft-francois-sr-frr	RTGWG
PCEP SR Extensions	draft-sivabalan-pce-segment-routing	PCEP




- Positive collaboration and consensus

# Progress

- Agreement between draft-sr and draft-gredler co-authors on a subset of SR:
  - ISIS, OSPF: merge already submitted
  - MPLS/SR, FRR and use-case: work in progress
- Disagreement on IPv6/SR, should not be an issue
  - We thus accepted to organize the documents such that those that only want to support the MPLS instantiation can do so
  - We believe that there is a clear demand (e.g. as confirmed by feedback in the room) , we are committed to a positive collaboration process
  - For IPv6, together with operators (Comcast, Rogers...), Martha Steenstrup and academia, we will submit the IPv6/SR draft proposal for the next IETF and will work with the community to improve it as required

# Visually





# Next IETF

- While a new WG might be formed, we would like to be able to present and review our proposal in their home WG's
  - ISIS
  - OSPF
  - MPLS
  - 6man
  - PCEP
  - RTGWWG
- Significant operator support and vendor consensus

## draft-filsfils-rtgwg-segment-routing-use-cases-00

Sections	Presentations today
2. IGP-based MPLS Tunneling	Martin, Victor
3. FRR	Bruno
4.1.1 Disjointness in dual-plane networks	Martin
4.1.2. CoS-based Traffic Engineering	Martin
4.4. Deterministic non-ECMP Path	Rob
5.2. SDN /SR use-case	Victor
6.4. Leveraging SR benefits for LDP-based traffic	Bruno
7. OAM	Rudiger

There is a lot of requirements, and SR meets all of them, despite their variety, with few extensions to core protocols. No new protocol is added.

# Conclusion

- Multi-vendor/operator constructive collaboration
- Many requirements/use-cases supported by small extensions to well-established core protocols
  - ISIS, OSPF
  - LFA
  - PCEP
  - MPLS
  - IPv6
- Significant industry interest and contribution to SR
- Your feedback and contribution are welcome!

# SR documentation

- <http://www.segment-routing.net/>
  - Conferences
  - IETF links
  - ...

# Annex

# Abstract Routing Model

draft-filsfils-rtgwg-segment-routing-00

# SR Objectives

- Tackling issues reported by operators for years
  - IGP-based FRR for any topology
  - Simpler to operate, more scalable explicit routing
- Supporting “SDN”-based services
  - Provide a more responsive and scalable interaction between WAN orchestration, the applications and the network
- Evolution, no revolution
  - Must be simple to operate
  - Must support incremental deployment

# Segment Routing


- A 32-bit segment can represent any instruction
  - Service
  - Context
  - IGP-based forwarding construct
  - Locator
- Ordered list of segments
  - An ordered chain of topological and service instructions
- Per-flow state only at ingress SR edge node
  - Ingress edge node pushes the segment list on the packet



# IGP Segments

- Prefix Segment
  - Steers traffic along ECMP-aware shortest-path to the related IGP Prefix
  - Global segment within the SR IGP domain
  - Node Segment: a prefix segment allocated to a prefix that identifies a specific node (e.g. the prefix is its loopback)
- Adjacency Segment
  - Steers traffic onto an adjacency or a set of adjacencies
  - Local segment related to a specific SR node
- SR Global Block
  - A subset of the Segment space
  - All the global segments must be allocated from SRGB
  - Operator manages SRGB like an IP address block: it ensures unique allocation of a global segment within the SR domain


# IGP Prefix Segment



A packet injected anywhere with active segment 65 will reach Z via ecmp-aware shortest-path

- Z advertises its global prefix segment 65 with his loopback address Z/32
  - simple ISIS sub-TLV extension
  - simple OSPF Opaque sub-TLV extension
- All remote nodes install the prefix segment to Z in the SR dataplane along the shortest path to Z/32
- IPv4 and IPv6
  - draft-previdi-isis-segment-routing-extensions-00
  - draft-psenak-ospf-segment-routing-extensions-00

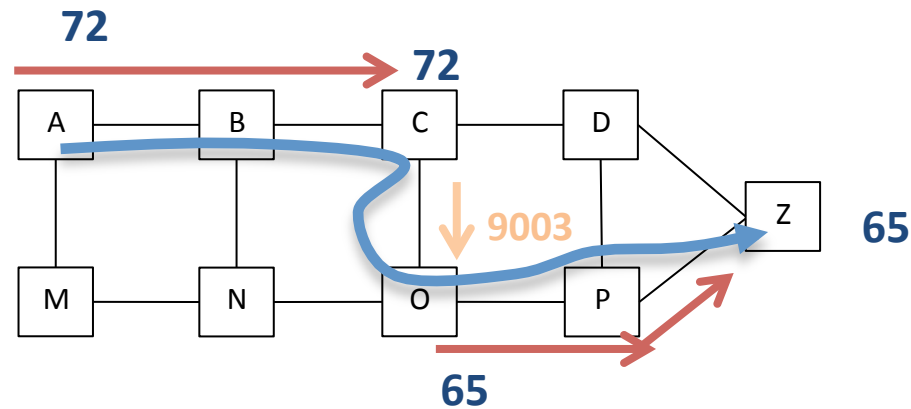
# IGP Adjacency Segment



A packet injected at node C with active segment 9003 is forced through datalink CO

- C allocates a local segment 9003 for its adjacency CO
- C advertises the adjacency segment in the IGP
  - Simple ISIS sub-TLV extension
  - simple OSPF Opaque sub-TLV extension
- C is the only node to install the adjacency segment in SR dataplane
- IPv4 and IPv6


# Combining Segments



- Source Routing
- ABCOPZ is expressed as {72, 9003, 65}

# Combining Segments

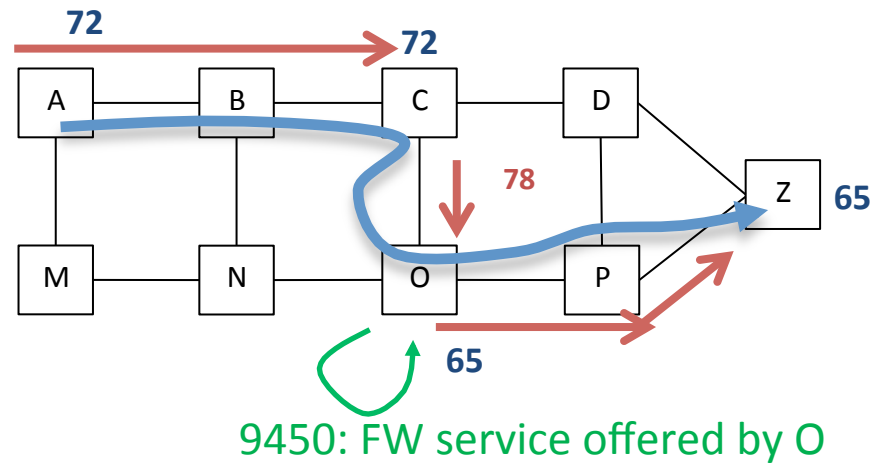
{72, 78, 65}



- Prefix Segment is at the heart of the proposal
  - ecmp multi-hop shortest-path
  - in most topologies, any path can be expressed as list of prefix segments

# Combining Segments

{72, 78, 9450, 65}



- Service Segments can be part of the source route

# SR Control-Plane

- Lightweight extension to ISIS/OSPF
- IPv4 and IPv6
- Agnostic to the dataplane
  - works with any dataplane that supports the encoding of a list of segments on the packet

# MPLS dataplane

- The 20 right-most bits of the segment are encoded as a label
- A list of segments is represented as a stack of labels
- The active segment is the top label
- The IGP Prefix segment stays on the top of the stack thanks to a SWAP operation where the ingress and egress label values are the same
- Transports IPv4 and IPv6
- No changes in the operations of the MPLS dataplane
- SR can co-exist and interwork with other MPLS control-plane protocols (LDP, RSVP)



# IPv6 dataplane

(without any MPLS dataplane)

- All the SR ISIS/OSPF Control Plane is dataplane agnostic and hence applies directly to IPv6
- Remaining work: detailing the IPv6 tunneling and new Routing Extension type header
  - High-level description provided at March IPv6 Conference
  - Detailed Draft should be available soon
    - We are working on this in close collaboration with Comcast and other SP/Enterprise operators and academia
    - Any contribution is welcome

# Annex