

I E T F

SPRING WG

Segment Routing Centralized Egress Peer Engineering *draft-filsfils-spring-segment-routing-central-epe-02*


C. Filsfils, S. Previdi, K. Patel (Cisco)

E. Aries, S. Shaw (Facebook)

D. Ginsburg, D. Afanasiev (Yandex)

Motivations


- draft-spring-problem-statement section 5.1.1.2.



- Requirement: steer traffic towards an egress point, an egress interface (or a set of), a peer (or a set of), a peer AS, or a combination of all the above

Motivations

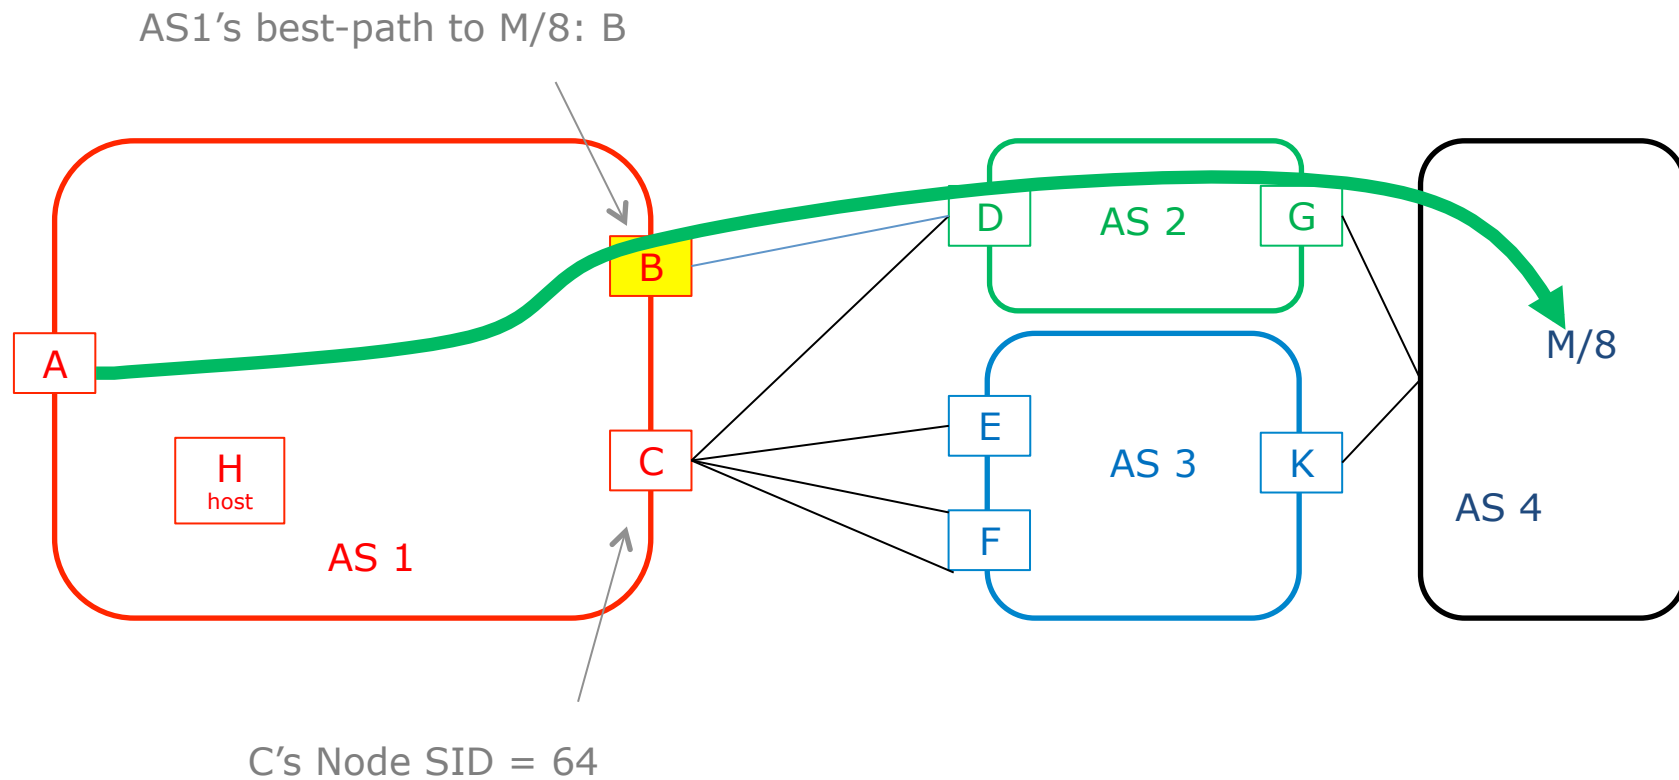
- Revised diagram in draft-filsfils-spring-segment-routing-central-epe




- Section 1.2 Problem Statement

A **centralized controller** should be able to instruct an ingress PE or a content source within the domain to use a specific egress PE and a specific external interface to reach a particular destination.

Reference Diagram




Objective: centralized egress peer engineering



- Per-Flow TE state only at the source node
 - Ingress router or directly at the source host

eBGP Peering Topology

BGP Peering Segments



Automated BGP Peering SID allocation


BGP Peering SID's in C's MPLS Dataplane

PeerNode SID's:

- 1012: pop and fwd to 1.0.1.2/32 (D)
- 1022: pop and fwd to 1.0.2.2/32 (E)
- 1052: pop and fwd to 1.0.5.2/32 (ecmp to F)


PeerAdj SID's:

- 1032: pop and fwd to 1.0.3.2/32 (upper link to F)
- 1042: pop and fwd to 1.0.4.2/32 (lower link to F)



BGP EPE Routes


- The controller learns the BGP Peering SID's and the external topology of the egress border router via BGP-LS EPE routes
 - draft-previdi-idr-bgpls-segment-routing-epe



Controller – Decision


- Collects valid internet routes from peers
- Collect performance information across peers
 - EPE solution allows to target probes across probed peer
- Based on business policy and performance information, decides to engineer a flow via an explicit peer different than the best-path
- Outside the scope of the IETF drafts

Controller Programming



- Draft does not define/mandate any specific programming method (out of scope)
- As an illustrative example:
 - PCEP extension to instantiate at A an SR Traffic Engineering Tunnel
 - Tunnel1: push {64, 1042}
 - PBR Policy: any traffic to M/8, set next-hop = tunnel T1
- Other methods: BGP-3107 policy route, Netconf...

Conclusion



- No assumption on the iBGP design with AS1 (nhop-self is fine)
- Integrated intra-domain and inter-domain TE
- EPE functionality only required at EPE egress border router and EPE controller
- Ability to deploy the same input policy across hosts connected to different routers
 - global property of the IGP prefix SID
- Per-flow TE state only at the source host or ingress border router

Questions?

Thanks!