Segment Routing
Centralized Egress Peer Engineering

Clarence Filsfils

cf@cisco.com

Segment Routing

- Implementation available
  - IOS-XR 5.2.0
- Efficient Standardization
  - 15 drafts
- Significant industry support
- Many use cases

www.segment-routing.net

EPE use-case

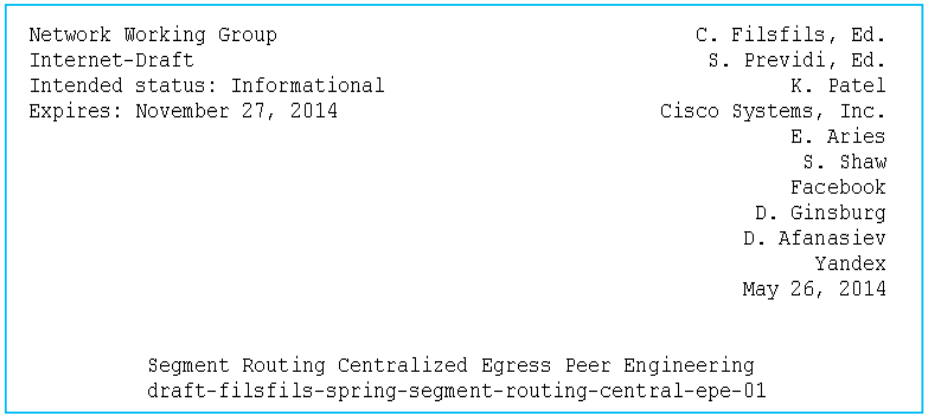

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 27, 2014

C. Filsfils, Ed.
S. Previdi, Ed.
K. Patel
Cisco Systems, Inc.
E. Aries
S. Shaw
Facebook
D. Ginsburg
D. Afanasiev
Yandex
May 26, 2014

Segment Routing Centralized Egress Peer Engineering
draft-filsfils-spring-segment-routing-central-epe-01

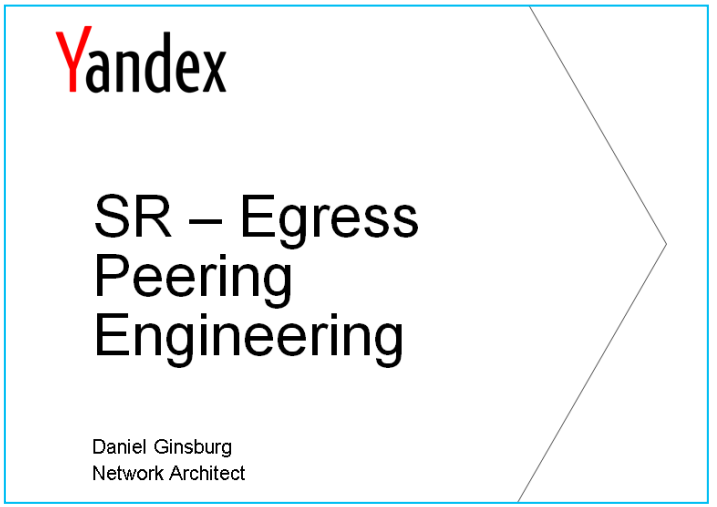

- IETF
  - draft-filsfils-spring-segment-routing-central-epe-01
  - draft-previdi-idr-bgpls-segment-routing-epe-00
- ENOG presentation
  - www.enog.org/presentations/enog-7/263-EPE-ENOG7.pptx

Reference Diagram

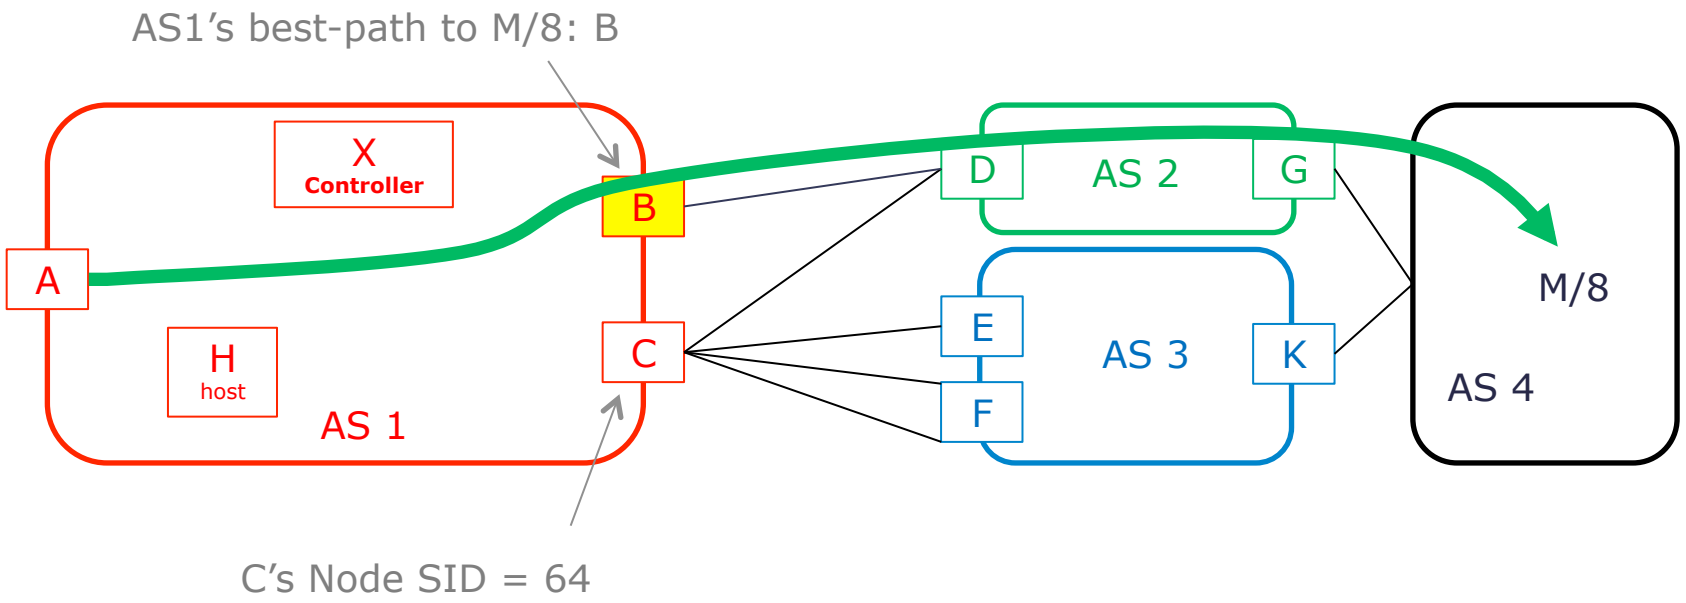

Objective: centralized egress peer engineering

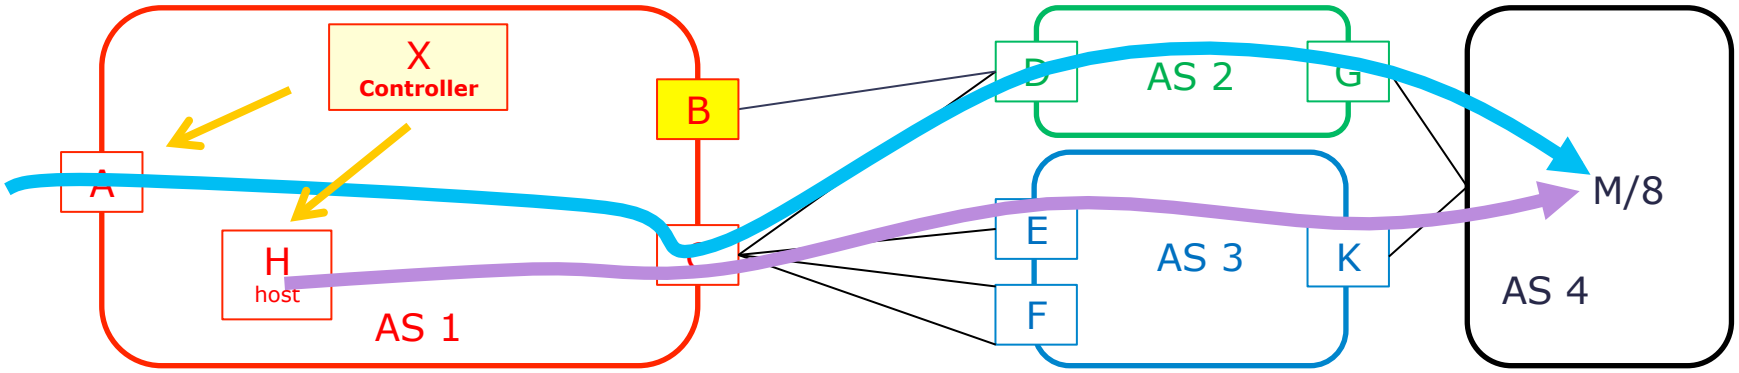

- Per-Flow TE state only at the source node
  - Ingress router or directly at the source host

eBGP Peering Topology

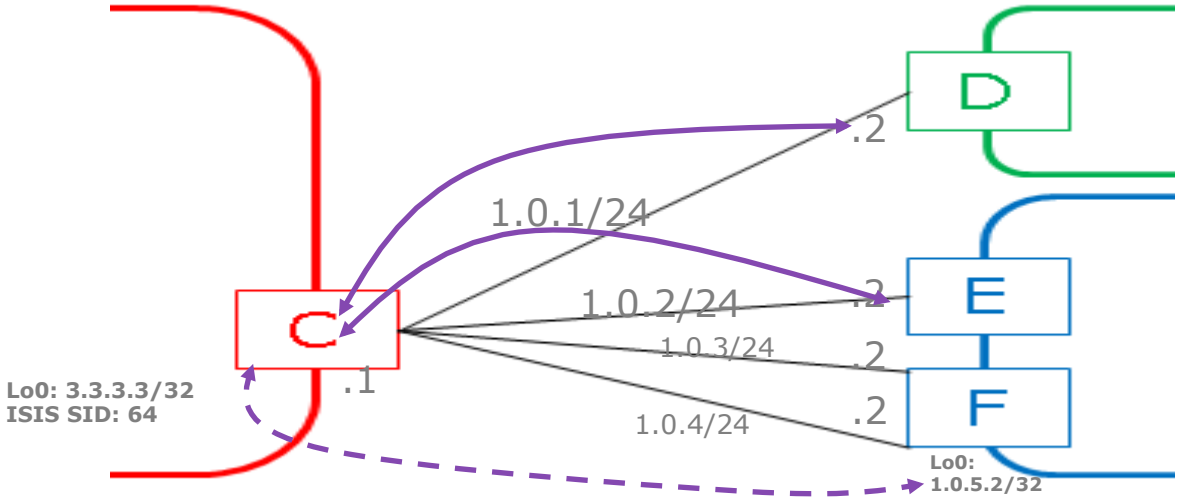

eBGP single-hop

eBGP multi-hop

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
1022: pop and fwd to 1.0.2.2/32
1052: pop and fwd to 1.0.5.2/32 (ecmp!)

PeerAdj SID's:
1032: pop and fwd to 1.0.3.2/32
1042: pop and fwd to 1.0.4.2/32

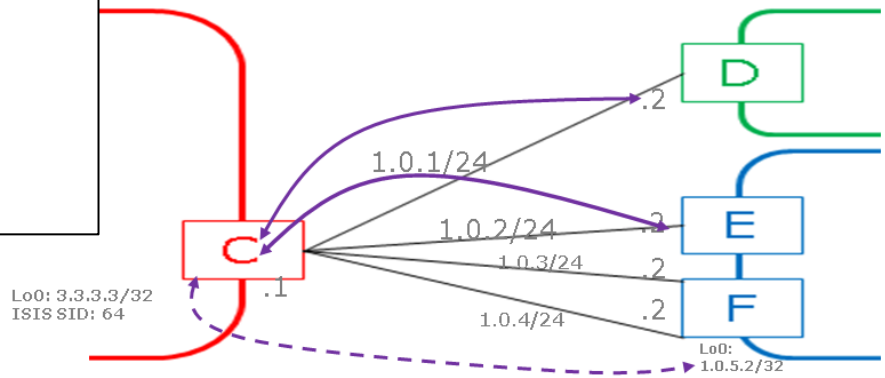

BGP EPE Routes

- The controller learns the BGP Peering SID's and the external topology of the egress border router via BGP-LS EPE routes

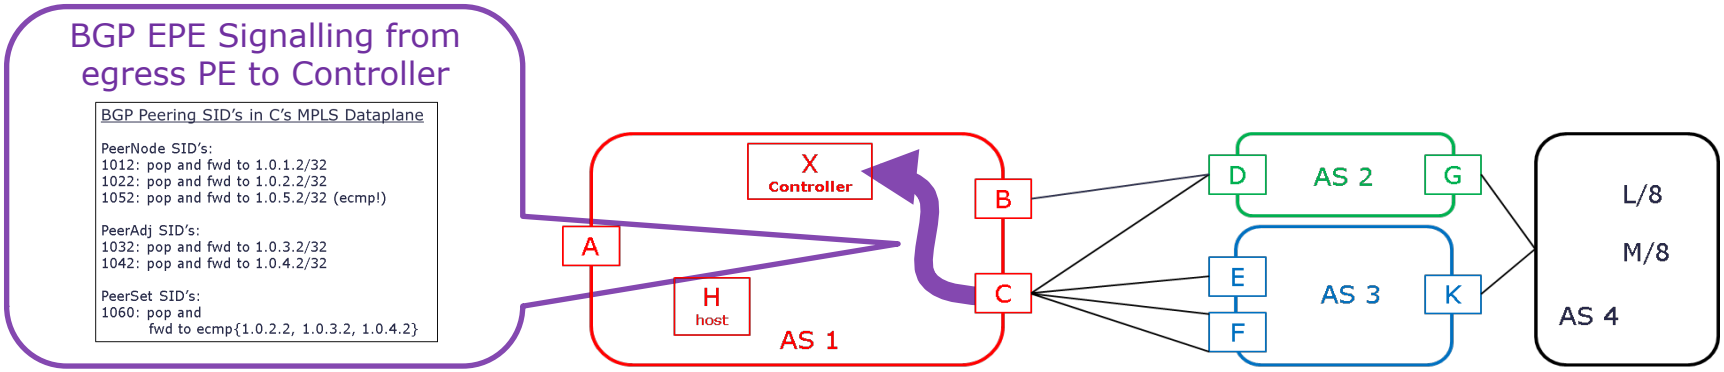

Controller – Decision

- 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

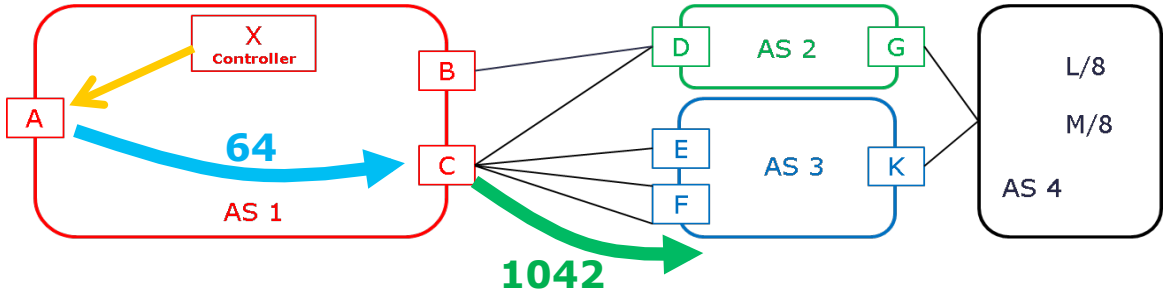

- PCEP extension to instantiate at A an SR Traffic Engineering Tunnel
  - Tunnel1: push {**64**, **1042**}
  - PBR Policy: any traffic to M/8, set next-hop = tunnel T1
- Other methods: BGP-3107 policy route, Netconf…

Conclusion

Centralized EPE
SR use-case

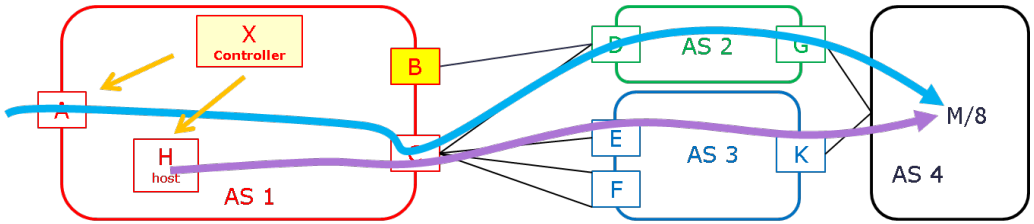

- No assumption on the iBGP design with AS1 (nhop-self is fine)
- Integrated intra-domain and inter-domain TE
- EPE functionality only required at EPE egress border router and EPE controller
- Ability to deploy the same input policy across hosts connected to different routers
  - global property of the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

Thank you.

Intra and Peer TE

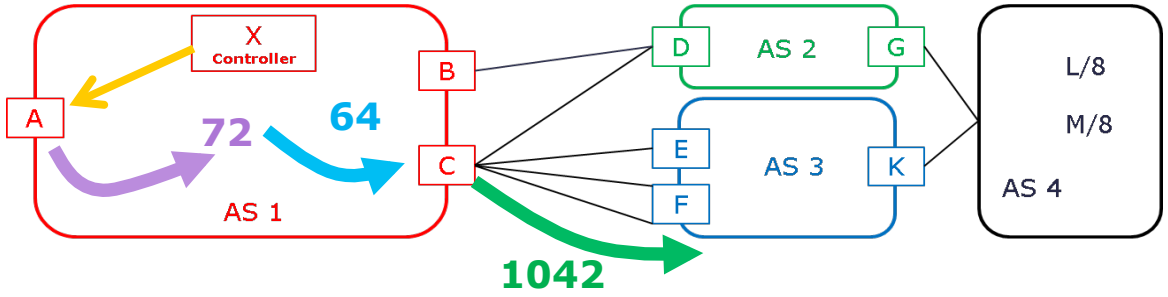

- PCEP extension to instantiate at A an SR Traffic Engineering Tunnel
  - Tunnel1: push {72, 64, 1042}
  - PBR Policy: any traffic to M/8, set next-hop = tunnel T1
- Other methods: BGP-3107 policy route, Netconf…