

Segment Routing with MPLS

- draft-filsfils-spring-segment-routing-mpls-00.txt
- draft-filsfils-spring-segment-routing-ldp-interop-00.txt

C. Filsfils
S. Previdi
Cisco Systems, Inc.

P. Francois
IMDEA Networks

B. Decraene
S. Litkowski
Orange

M. Horneffer
Deutsche Telekom

I. Milojevic
Telekom Srbija

R. Shakir
British Telecom

S. Ytti
TDC Oy

W. Henderickx
Alcatel-Lucent

J. Tantsura
S. Kini
Ericsson

E. Crabbe
Google, Inc.

SR for MPLS

- Segment Routing architecture leverages MPLS dataplane without any modification
- An MPLS segment is a 20-bit label
- Co-existence between SR and LDP/RSVP

SR for MPLS

- Co-existence between SR and LDP/RSVP
 - SR and RSVP
 - > ships in the night
 - > advertisement of TE-LSPs into ISIS/OSPF using Binding SID/Label TLV
 - SR and LDP: interoperability
 - > both can be used within same infrastructure

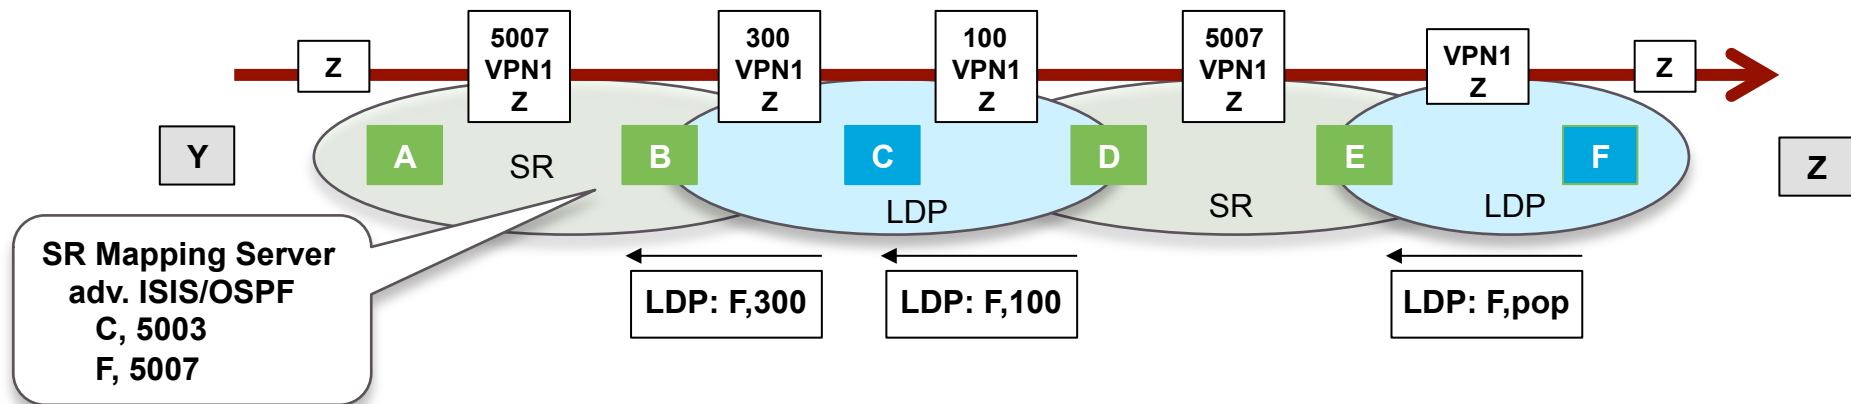
SR for MPLS

- SR over MPLS dataplane introduces global labels
 - The architecture has the notion of indexed labels which allow a node to use locally administered label space
- SR over MPLS supports explicit routing with the following properties
 - ECMP awareness
 - no signaling across the path
 - state is at ingress not in core

SR for MPLS

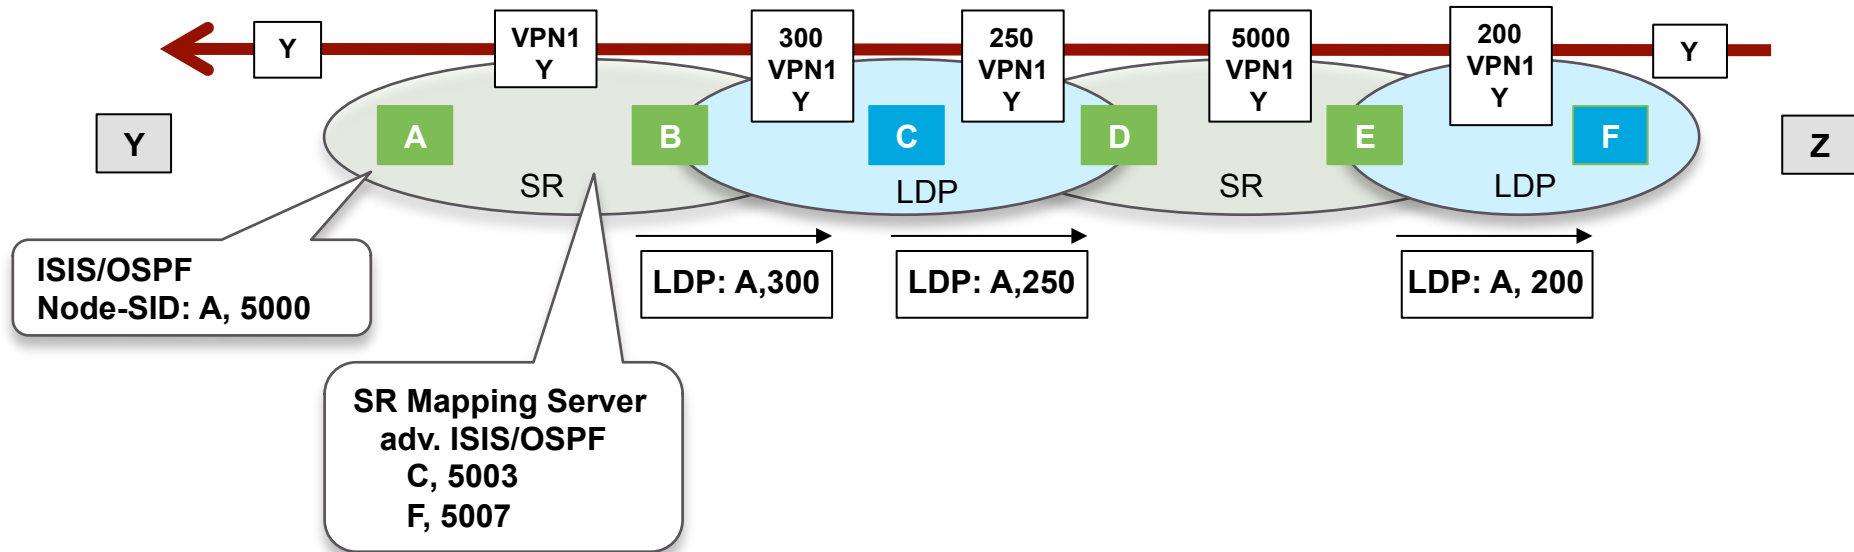
- When SR is used over MPLS:
 - The Segment List is represented by the label stack
 - The Active Segment is the top label
 - SR defines operations:
 - > CONTINUE: equivalent to MPLS swap
 - > NEXT: equivalent to MPLS pop
 - > PUSH: equivalent to MPLS push
- No changes in MPLS operations

SR/LDP Interoperability



- One or more nodes in the SR domain act as SR Mapping Server (SR-MS)
 - ISIS/OSPF SID/Label Binding TLV
 - SR-MS advertises SID mappings on behalf of non SR capable nodes
- Each SR capable node learns about SID assigned to non-SR capable nodes
 - no need to configure anything in each individual node
 - completely transparent to LDP nodes
- Example of mixed SR/LDP in MPLS-VPN case
 - PE A resolves the VPN destination and encapsulates traffic destined to Z into the SID learned by the SR-MS
 - Node B swaps the SID with the LDP learned label for FEC F
 - Node C does plain MPLS LDP-based label swap
 - Node D uses the SR-MS learned SID for F
 - Node E does plain LDP-learned PHP

SR/LDP Interoperability



- Mapping Server can be located anywhere in the ISIS/OSPF domain
- Mapping Server SID/Label Binding TLV is flooded throughout the domain
- Nodes can be gradually upgraded to SR