

SEAMLESS SR

draft-hegde-spring-mpls-seamless-sr

IETF 108

Shraddha Hegde, Juniper Networks


Chris Bowers, Juniper Networks


Alex Bogdanov, Google

Arkadiy Gulko, Refinitiv

Xiaohu xu, Alibaba Inc.

Jim Uttaro, AT&T






Agenda

- Introduction
- Use cases
- Proposed Solution
- Examples
 - Data Sovereignty

Introduction to Seamless SR





- Based on Seamless MPLS architecture
- Accommodates heterogeneous networks
- Extends Seamless MPLS
 - Support for End-to-End Slicing
 - Native MPLS for IPv6 networks
 - Interworking with IPv6 Technologies
 - MPLSoUDP, MPLSoGRE, SRm6, SRv6

5G Transport Networks




- Massive bandwidth
 - (N*10GE, 25GE, 100GE, 400GE)
- Timing and E2E Latency
 - (Stringent Timing and latency requirements due to distributed/virtualized network functions)
- Operational Simplicity
 - (Fewer protocols, stateless/programmable)
- High availability
 - (Node/Link/Path protection)
- High scalability
 - (100xNode scale)
- E2E Service Differentiation
 - (Differentiated services with varied service mapping)
- Application aware routing

Large Scale Service Provider networks



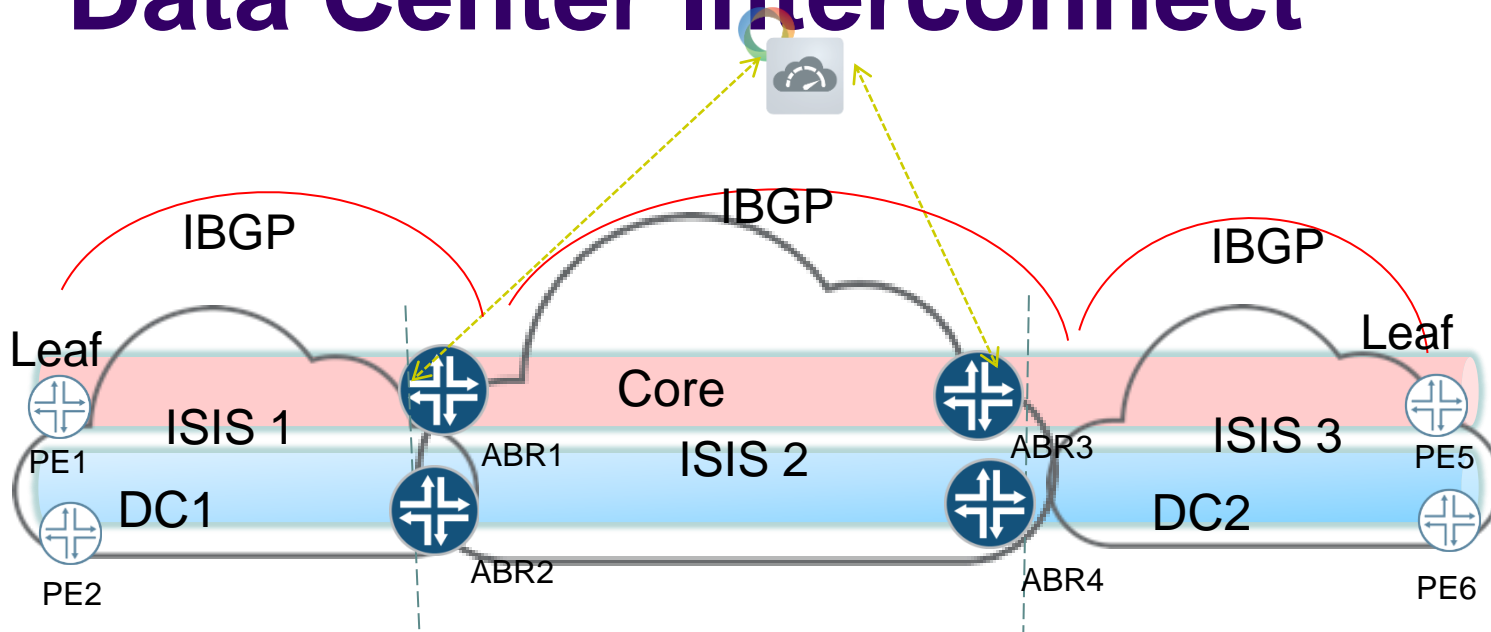
- Large number of nodes in the order of 100k
- Multiple domains with single ownership
- End-to-end SLO
- Low latency routing

LARGE SCALE WAN NETWORKS




- Smaller IGP domains
- Reduce overall state on the device
- Smaller fault domain
- E2E TE constraints
- Data Sovereignty

Data Center Interconnect

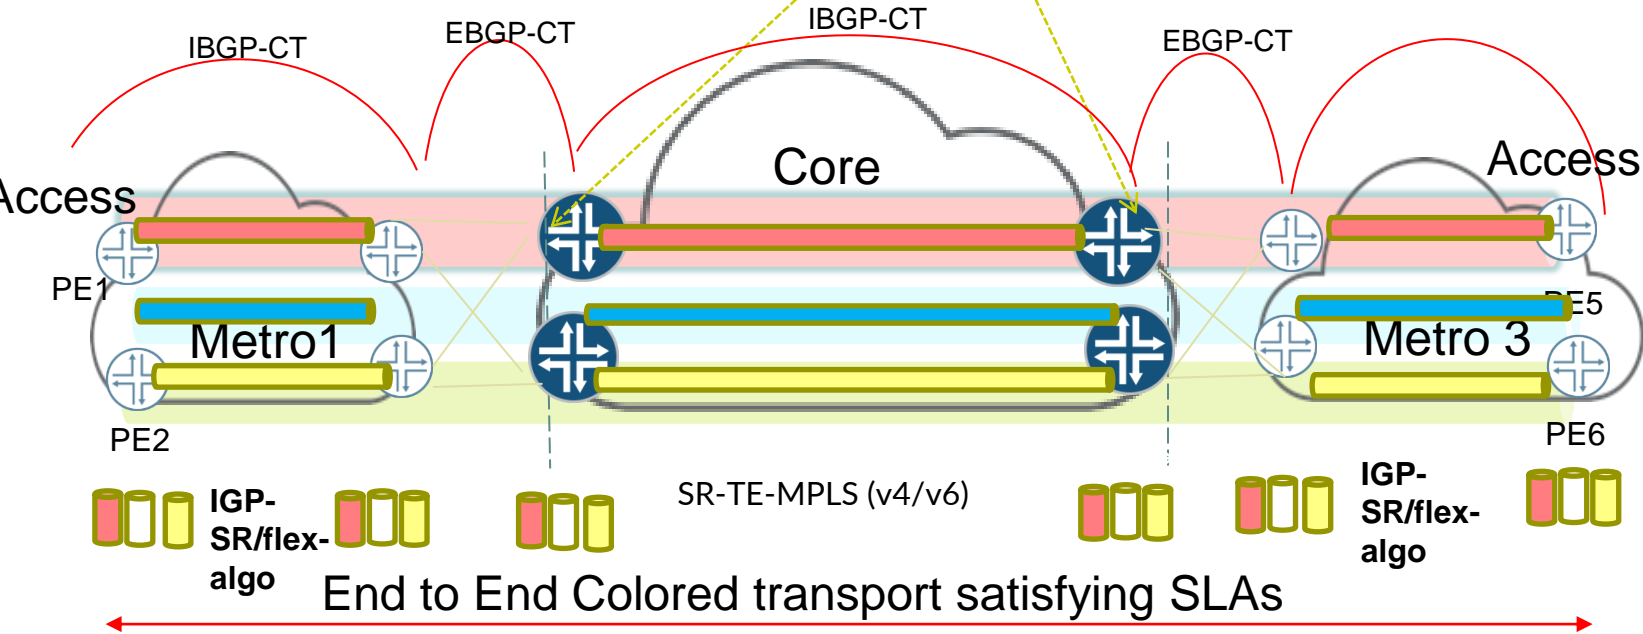


- E2E Path Diversity
- Low latency Paths
- Avoid service routes on ABRs

Seamless SR with BGP-Classful Transport

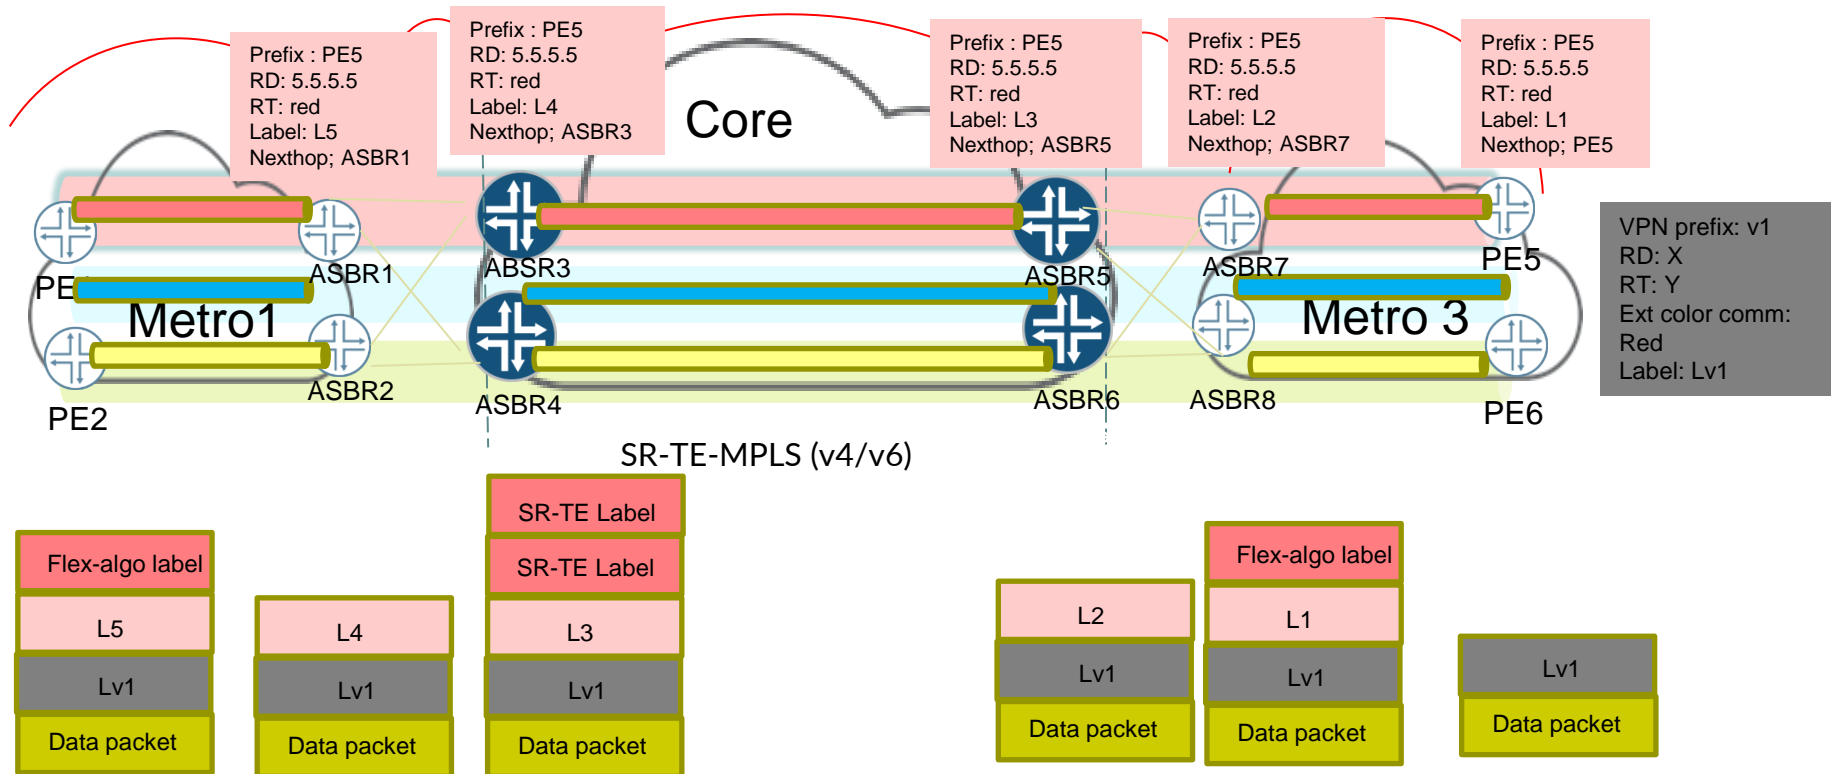



draft-kaliraj-idr-bgp-classful-transport-planes-00




- Transport class
- Transport RIB
- BGP-CT
- Route Distinguisher
- Route Target
- Mapping Community

Seamless SR with BGP-Classful Transport

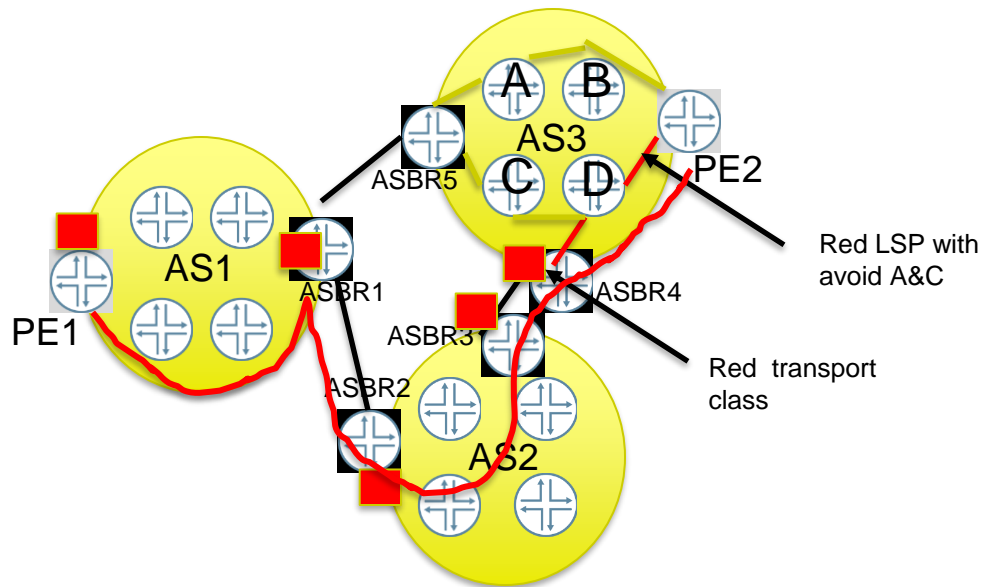


Data sovereignty Use case



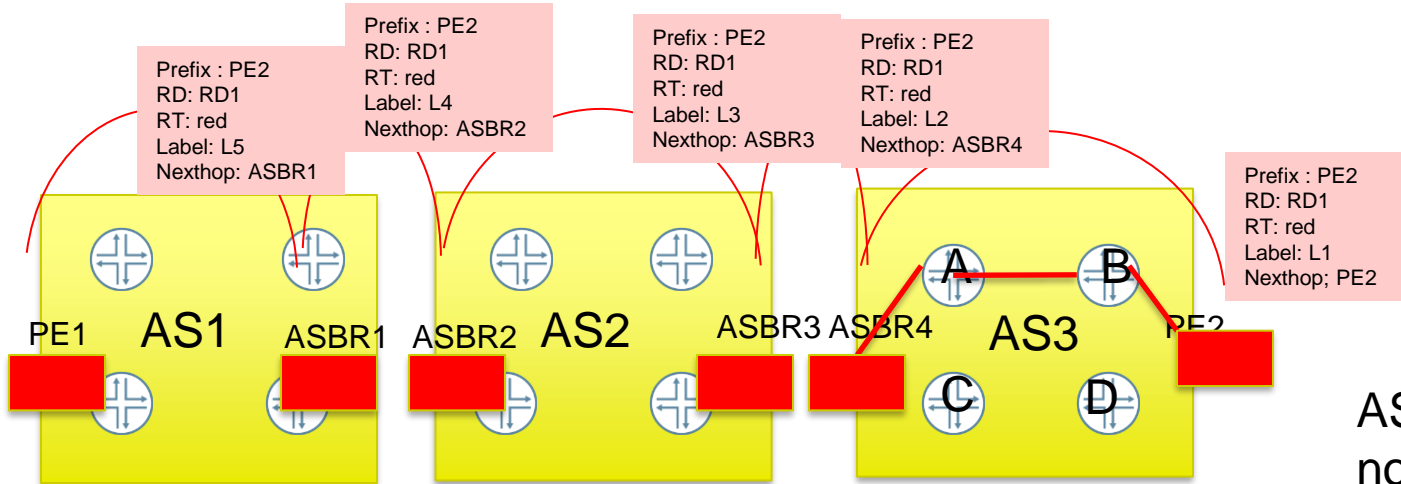
- Multiple AS
- Each AS represents a continent
- Data sovereignty
 - Avoid node A and C
- This “avoid node A& C” constraint is not applicable in AS1 and AS2

Data sovereignty Use case

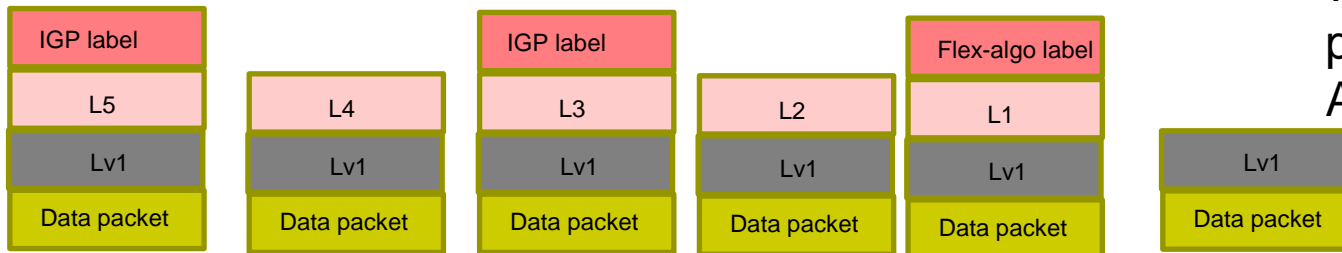


- Red LSP created on ASBR4
- Since ASBR5 is connected to A & C, Red LSP isn't created
- Red transport class created on all border nodes
- Resolution
 - AS3: Strict resolution
 - AS1 & AS2: Resolve on Red fallback on best

Solution with BGP-CT





AS1 and AS2 do not need to create Red Transport Tunnel. BGP-CT will use best effort paths in AS1 and AS2



Next steps

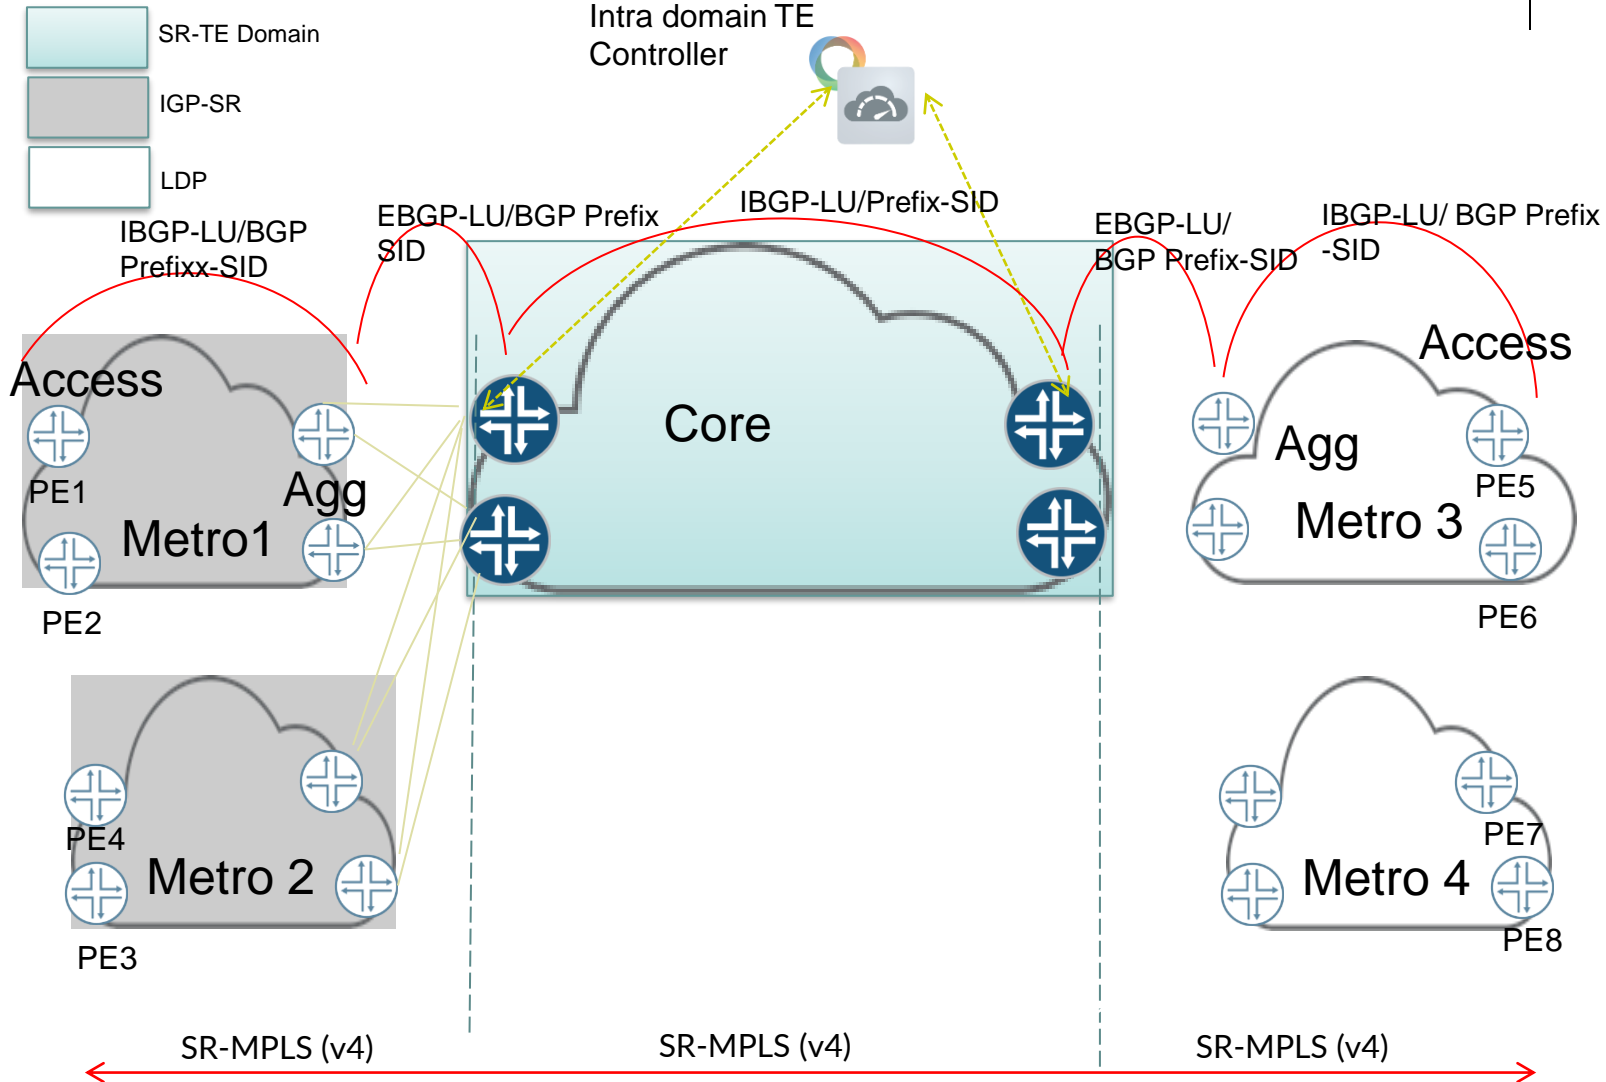
- Request review and comments






Thank you

Seamless MPLS






Shortcomings of BGP-LU

- Single path to remote loopback
- BGP best path selection applied to choose the path
- Possible options
 - Multiple loopbacks
 - BGP-LU with add-path

Both options are operationally very cumbersome.
For large networks, BGP based E2E TE solution needed



Solution Concepts

- **Transport Class:** Set of paths satisfying certain constraints
- **Transport RIB:** Collection of tunnels belonging to same Transport Class
- **BGP Classful Transport:** New NLRI for carrying E2E SLA paths
- **Route Distinguisher:** Used to distinguish multiple paths to same loopback
- **Route Target :** Carries the Transport class of the BGP-CT advertisement
- **Mapping Community :** Extended community as defined in RFC 4360 used for service mapping