Separating Routing Planes using Segment Routing

draft-gulkohegde-spring-separating-routing-planes-using-sr-00

IETF 98 – Chicago, USA

Shraddha Hegde shraddha@juniper.net
Arkadiy Gulko arkadiy.gulko@thomsonreuters.com

Use Case

- MPLS based Dual Plane with common control plane -single IGP
- To support the following L3 VPN Unicast Services:
  - Optimized based on SPF loosed path per plane
    - No rerouting via inter- plane link (shunt)
    - Specific Plane traffic MUST be dropped if plane is partitioned.

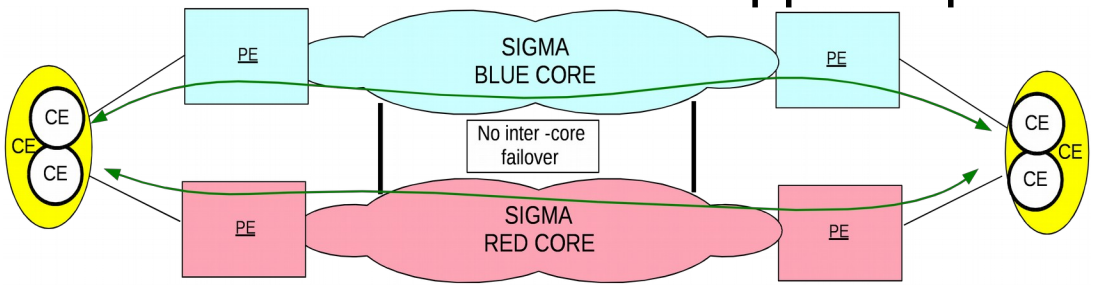

  - Optimized based on SPF loosed path via any planes
    - Traffic stays within a plane and
    - Allowed to be rerouted via shunt links if the plane is partitioned.

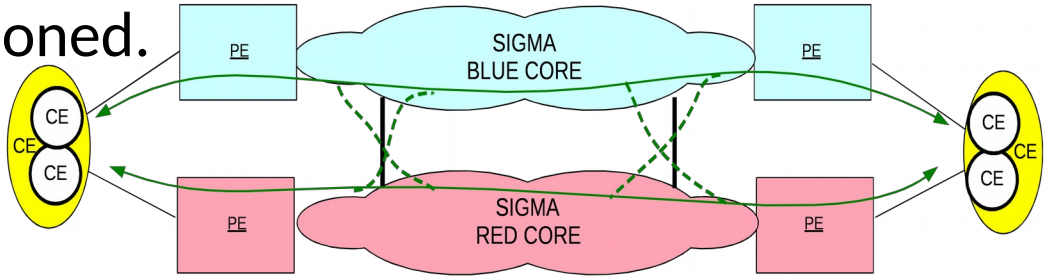

- Ability to fast-recover (TI-LFA) from a network node or link failure within a Service constrains above.

Requirements

- Maintain strict routing within routing planes.
- Allow traffic to failover within routing plane and do not allow traffic to failover to other planes.
- Achieve ease of configuration and operational management.
- ~50ms recovery from a network node or link failure.

Problem Statement

- None of the currently available techniques fully meet all of the requirements.
  - Using Node-Admin tags to color each of the planes separately,
  - Using Separate Anycast-SIDs - one per plane,
  - Multi-Topology (MT)-SIDs.
- Worth to mention about end to end ERO as another alternative, however deep label stack is a problem for current hardware and software. Thus it is out of scope for further discussion.

Node-Admin tags issues

- Node –Admin tags approach is part of SR-TE Policy-based Routing
- PCE or Head-end conducts path validation based on defined constraints, in dual plane design it is based on specific resource avoidance.
- **Problems**:
  - **TI-LFA will take traffic to alternative plane if plane got partitioned. High bandwidth flows saturate the alternative plane until path invalidation takes effect.**
  -

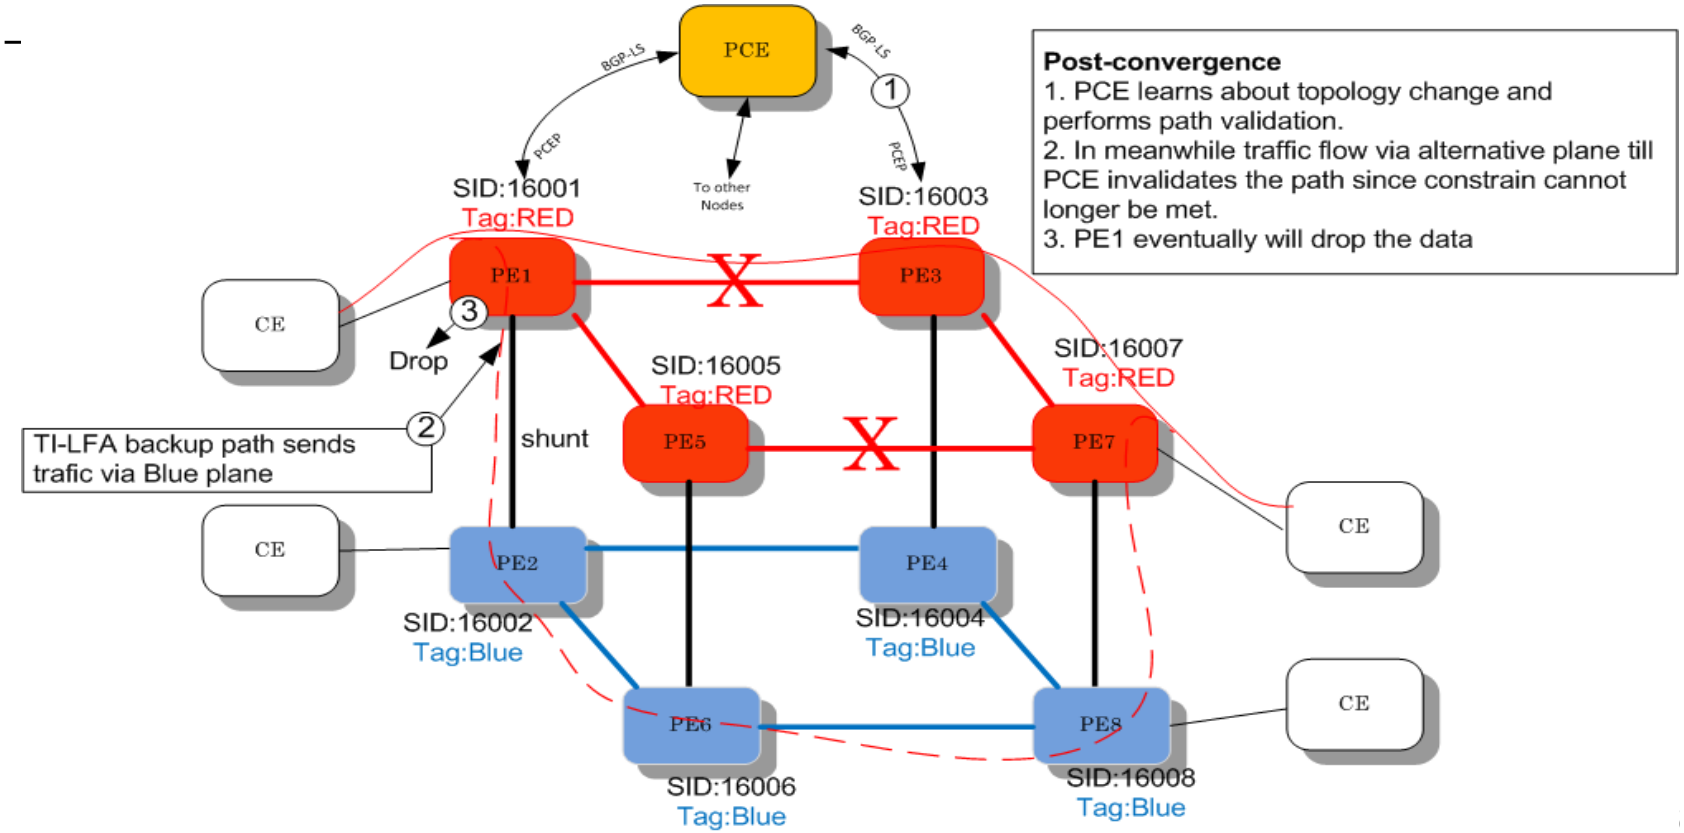

Anycast-SID issues

- Anycast SID is primarily used to steer traffic via shortest-path towards the topologically nearest node in a group of potential receiving devices.
- The Anycast SID label has served its purpose once the packet has reached any node in the plane with anycast SID. Anycast label popped and next label will be used to reach the destination.
- In the dual plane design Anycast SID inherit absolutely the same function as Node-Admin tag, just instead of 32bit color tag PCE would use per color SID label to perform path validation.
- **Problems: Exactly the same as for Node-Admin tag.**

Multi-topology-SID issues

- Multi topology Routing defines mechanisms to support multiple topologies in a single physical network.
- All the nodes in the network compute separate SPF per MT-ID and program the forwarding planes with MT-SIDs accordingly.
- This technology meets majority of dual plane requirements and provides additional benefits to assign different IGP costs to links for different MTs.
- **Problems: MT associated with operational overhead. Need to maintain separate IGP topology and complexity of mapping services to different topologies. Additional protocol overheads to advertise MT related information.**

Proposed Solution

- To introduce new SID -Routing Plane (RP)- SID
- RP- SID is defined and associated with new algorithm values.
- This document proposes 4 new algorithm values which represent SPF in routing-planes.
- OSPF Router Information (RI) TLVs Registry
  - 8 (IANA Preallocated) - SR-Algorithm TLV
  - Algorithm 2 -5 : SPF in routing plane
- ISIS Sub TLVs for Type 242
  - Type: TBD (suggested value 19)
  - Description: Segment Routing Algorithm
  - Algorithm 2-5 : SPF in Routing Plane

Proposed Solution

- Routing Plane (RP)- SID solution-based on per plane SPF algorithm for reachability within a plane.
- Specific RP-SID would never be known via alternative plane, based on principle that any node will ignore the RP-SID received from remote node with an algorithm value that such remote node has not advertised.
- Example below based on ISIS (node SIDs and default Algo advertisements are not shown)

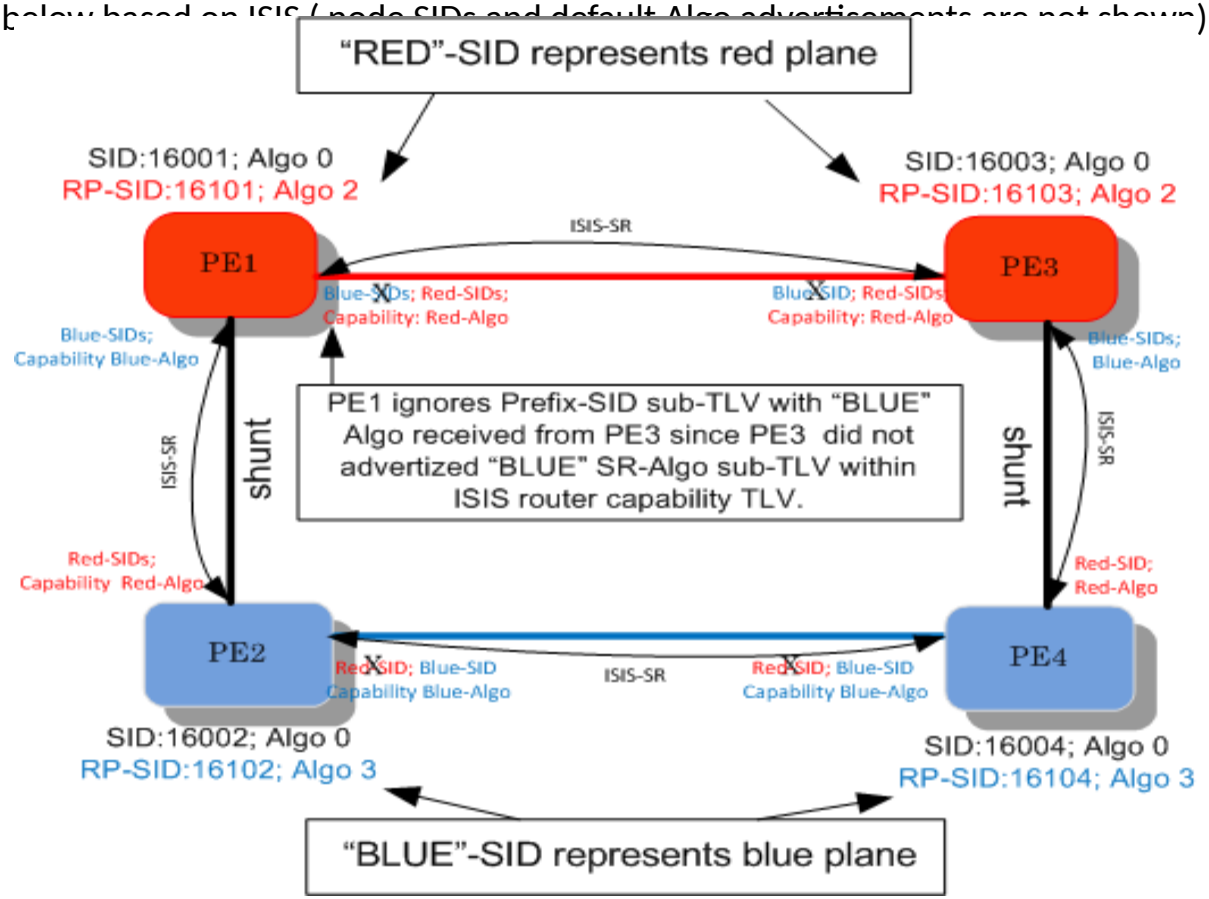

Mutli-Level/Area Support

- The Routing Plane SIDs MAY be re-originated from one IGP domain into the other domain by the border routers.
- The border IGP routers MUST re-advertise the Routing-Plane SIDs with its related Algorithms if they belong to the corresponding Routing plane and has advertised the algorithm corresponding to the routing-plane.

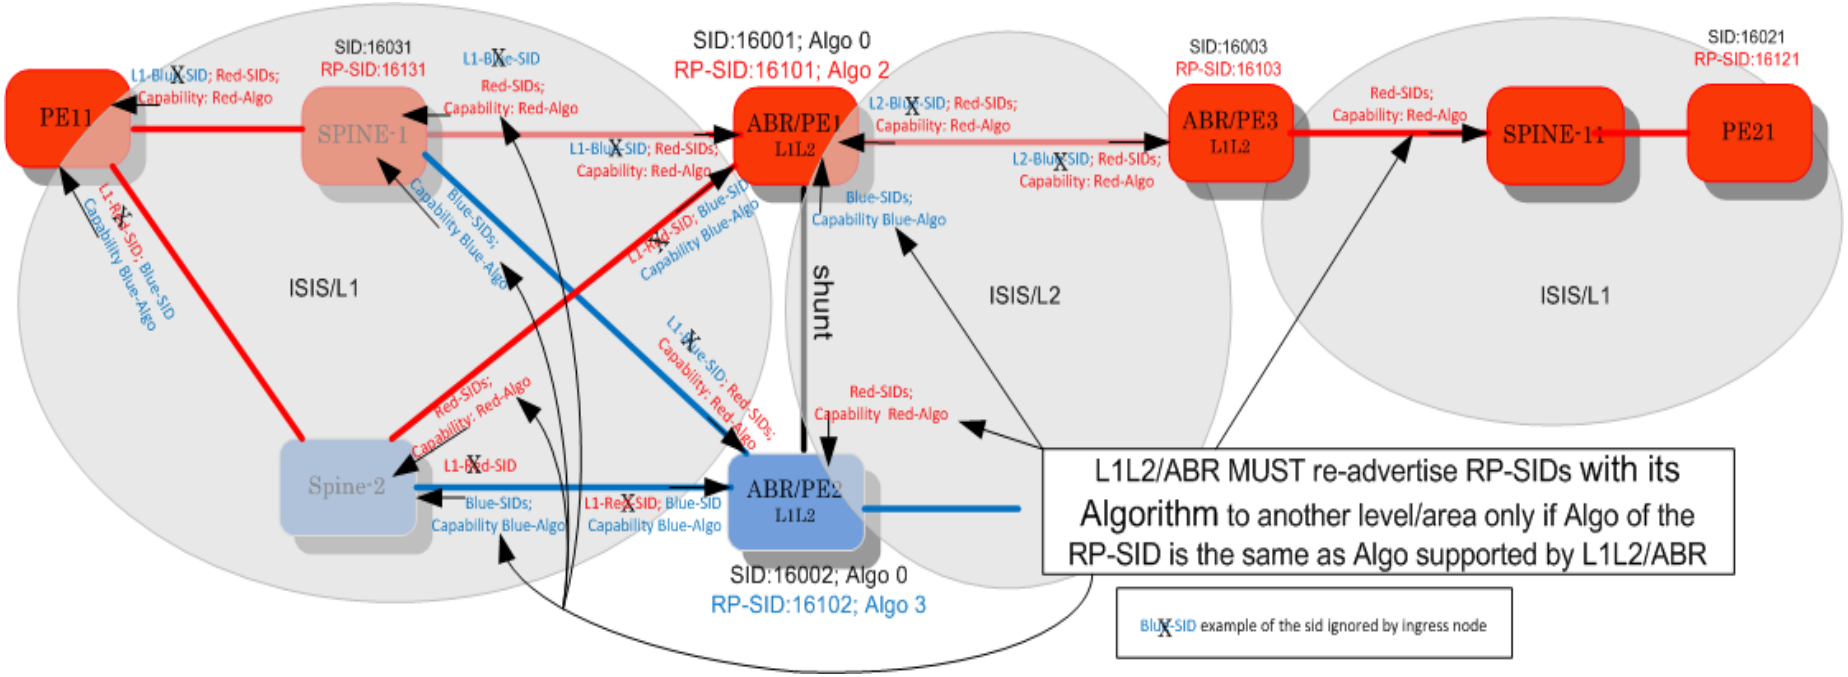

Service Provisioning

Control and Data Planes
Native SR forwarding- NO NEED for SR-TE!
One transport label
Next hop per tenant vpn or prefix range would be used to choose specific plane

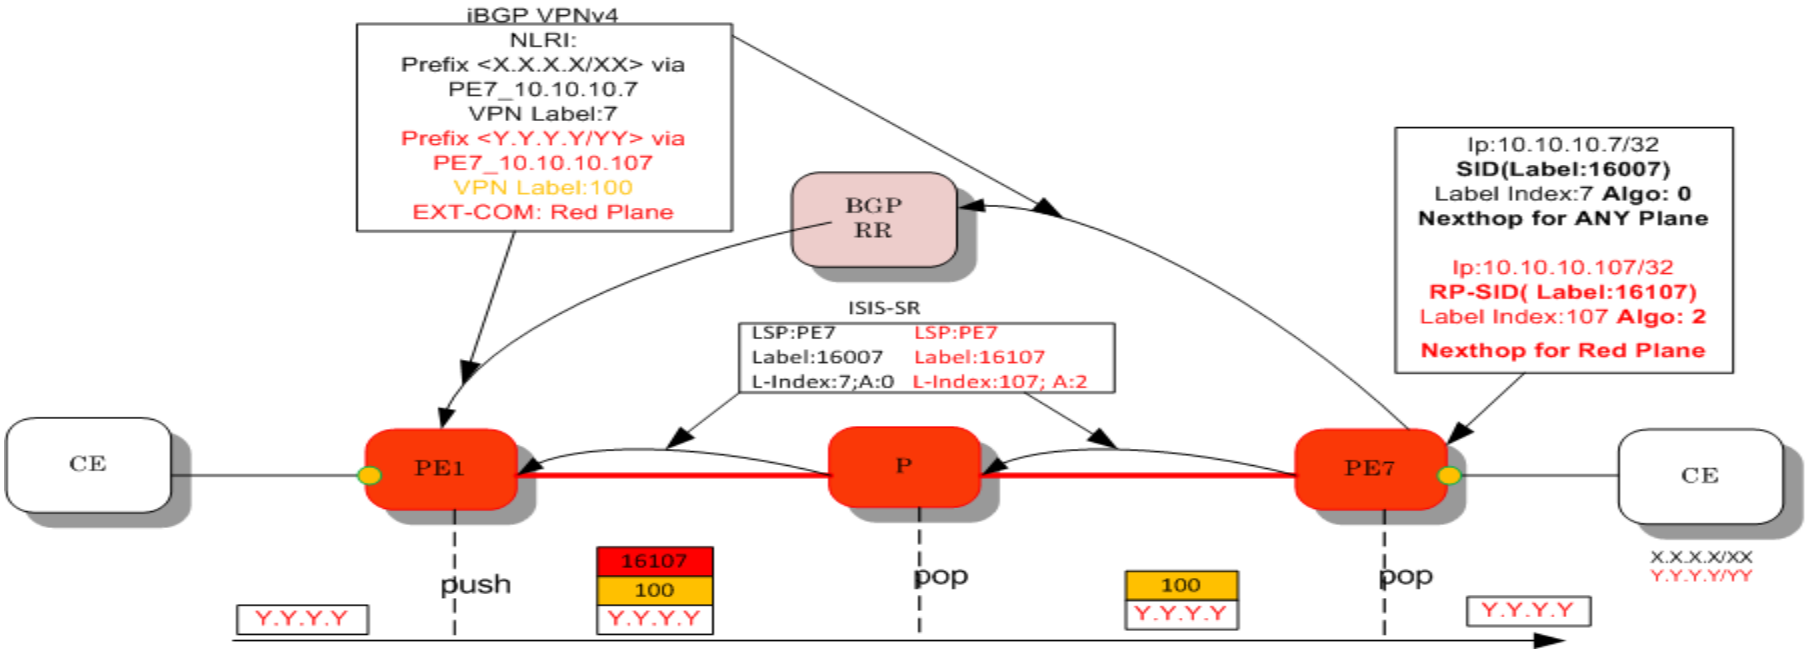

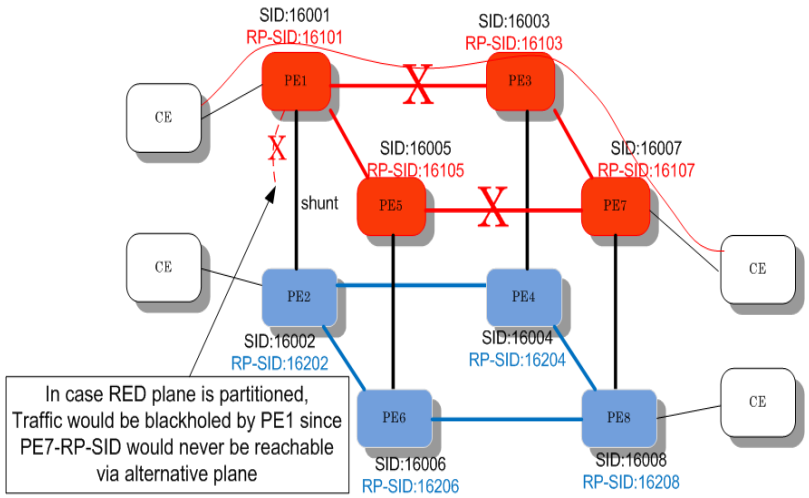

BGP DC Use Case

- Support for the advertisement and consistent filtering of RP- SIDs via BGP-LU for dual plane end to end provisioning. (BGP-LU within a DC-Fabric and ISIS within a core).
- It is desired to constrain LSPs to a sub-set of Spine switches (e.g., only those Spine switches which are 'coloured' RED).
- No SR-TE required to support end to end diverse paths.
- Segmented LSP over RED plane use case is only shown below; Single Transport Label!

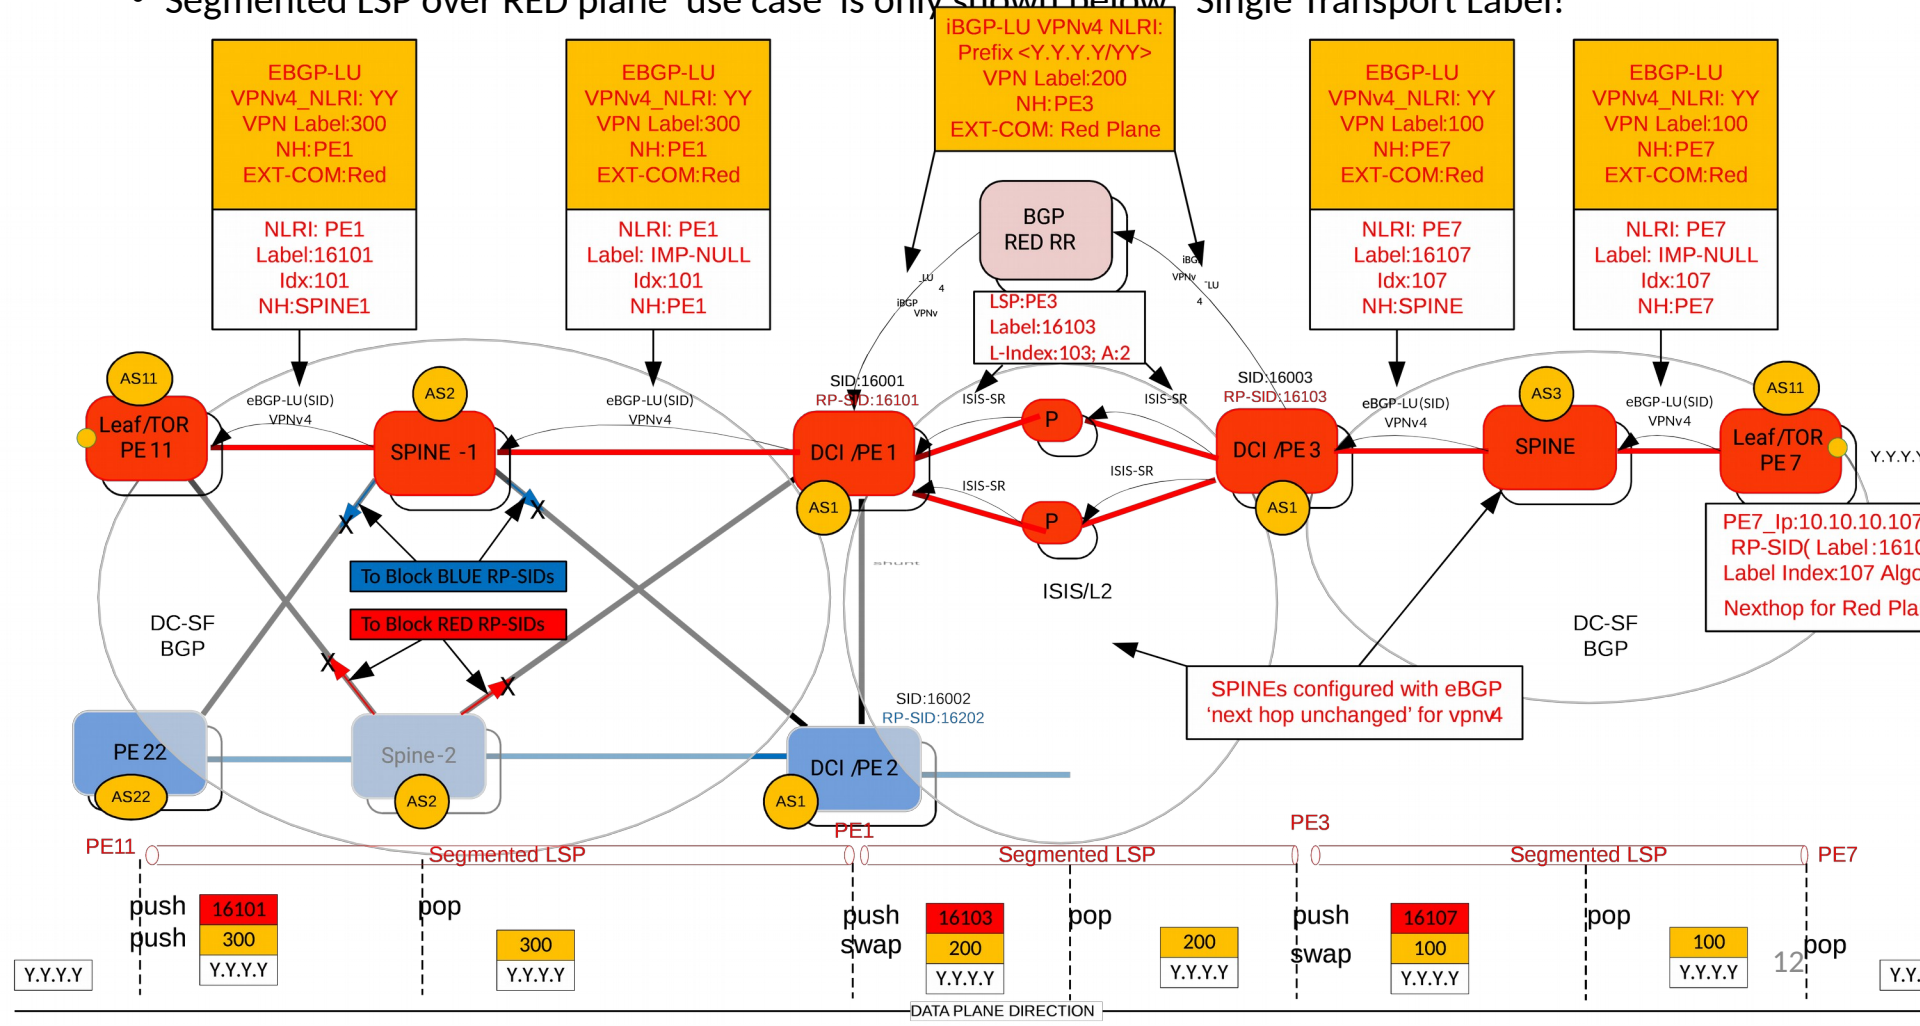

RP-SID Limitations

- Additional SID need to be advertised
- The solution is restricted to this use-case with routing-planes and will not accommodate other TE constraints like link colors/te-metric etc

Summary/ Conclusion

	RP-SID	MT-SID	Anycast SID	Node-admin tag
Prevention of traffic failover to different plane	Yes	Yes	No	No
Additional SID	Yes	Yes	Yes	No
No Additional protocol overheads	Yes	No	NO (BGP-SR-TE, PCEP needed for e-2-e solution)	No (BGP-SR-TE, PCEP needed for e-2-e solution)
No Operational overheads	Yes	No	No	No

Thank you!

Questions/Comments?