



# Segment Routing


## Deployment Experience and Technology Update

Clarence Filsfils  
Cisco Fellow  
[cf@cisco.com](mailto:cf@cisco.com)

# Deployed !

## Deployment

- In CY2015, SR will be deployed in all of these markets



WEB

SP Core/Edge

SP Agg/Metro

Large Enterprise



© 2014 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

- First deployments in 2015
- Strong start in 2016 with many new deployments

# Agenda

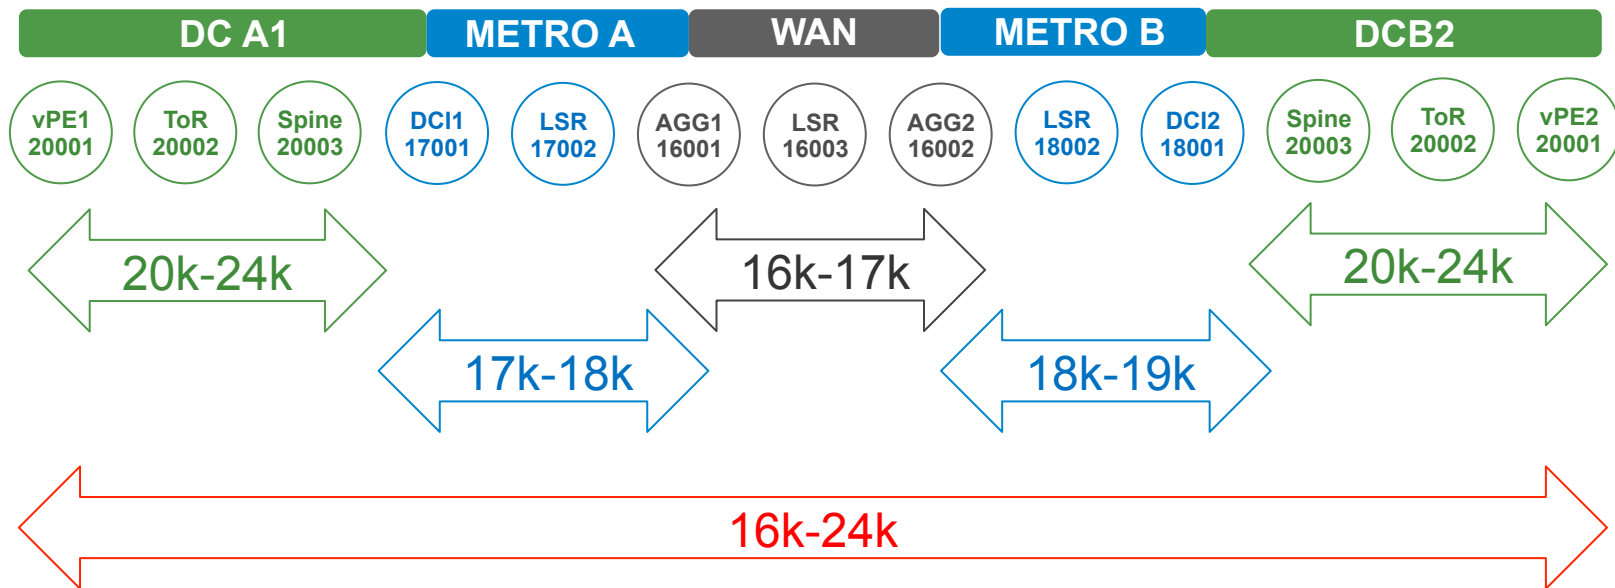
- Review advanced use-cases and related new technology
- Demo's are available at the booth
- Detailed SR tutorial
  - <http://www.segment-routing.net/home/tutorial>



# Inter-Domain Policy at Scale

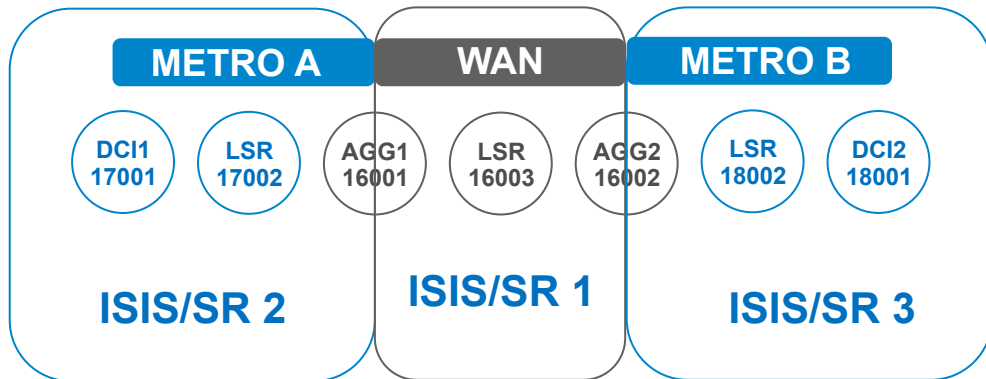
draft-filsfils-spring-large-scale-interconnect

# SRGB and SID allocation



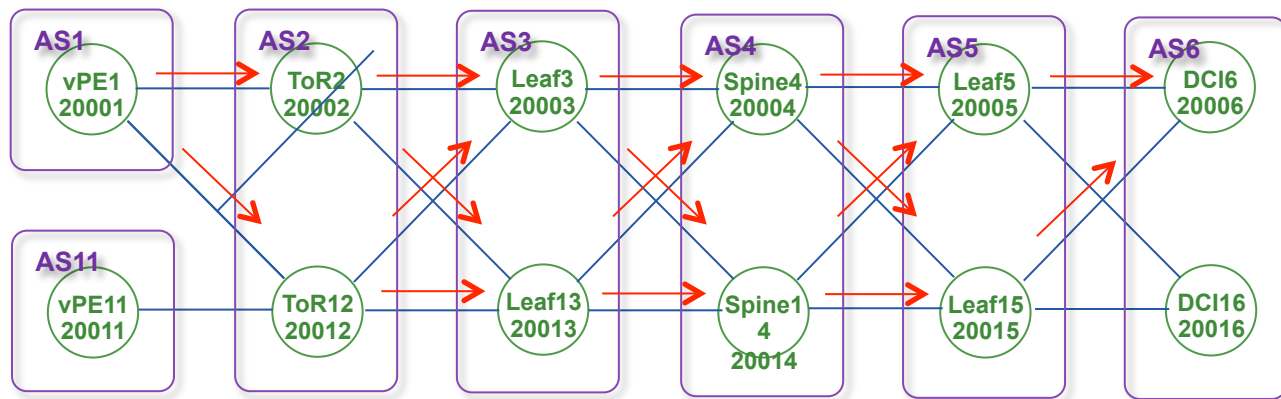
- Homogenous SRGB for simplicity

# IGP/SR within WAN and Metro



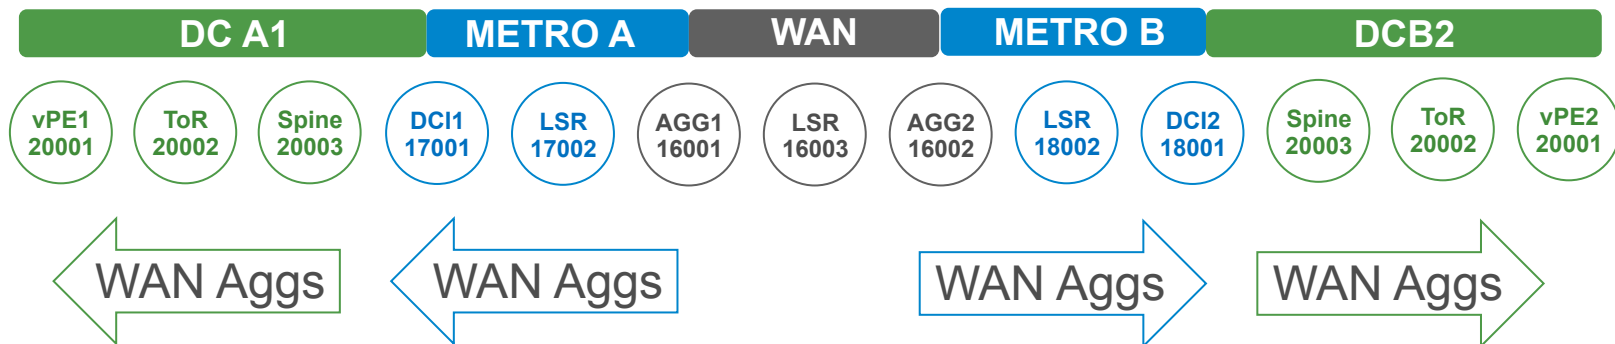
- Each domain runs ISIS/OSPF SR
- Technology available since June 2014
- Incremental deployment and seamless interworking with LDP

# SR in the DC



- 20006 is the BGP Prefix SID to DCI6
  - ECMP, simplicity (no LDP/RSVP) and policy
- Available on Nexus/XE and NCS5k/XR

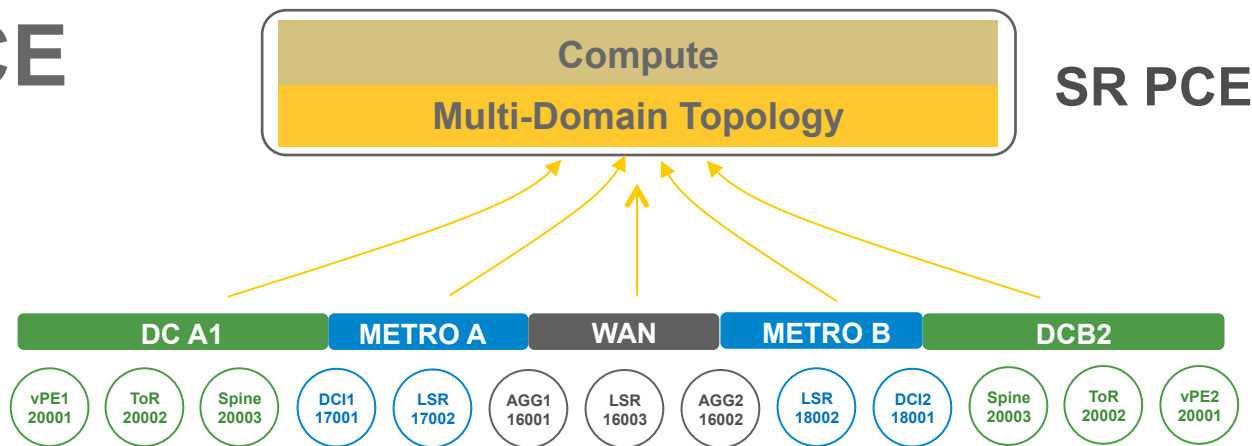
# Inter-Domain Routing



- WAN aggs are re-distributed down to Metro and DC
- Nothing is redistributed up
- How does vPE1 reaches vPE2?

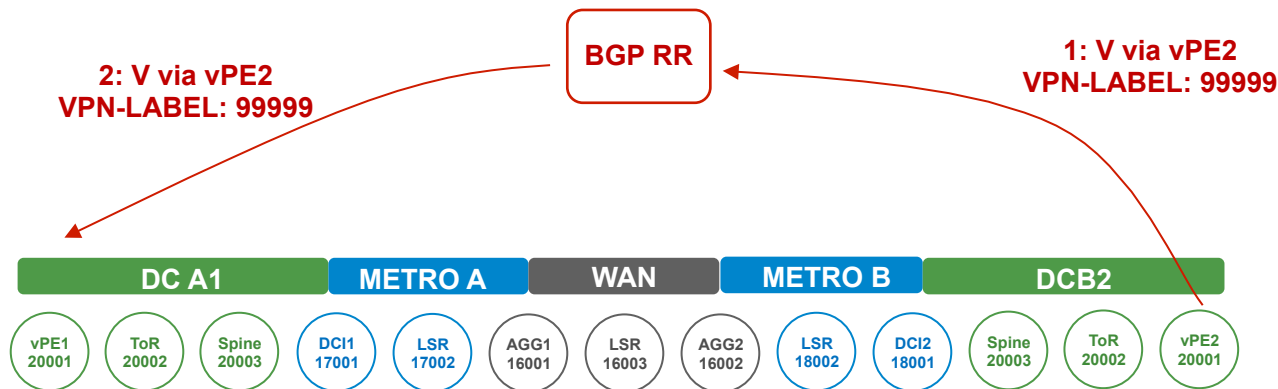


# SR PCE



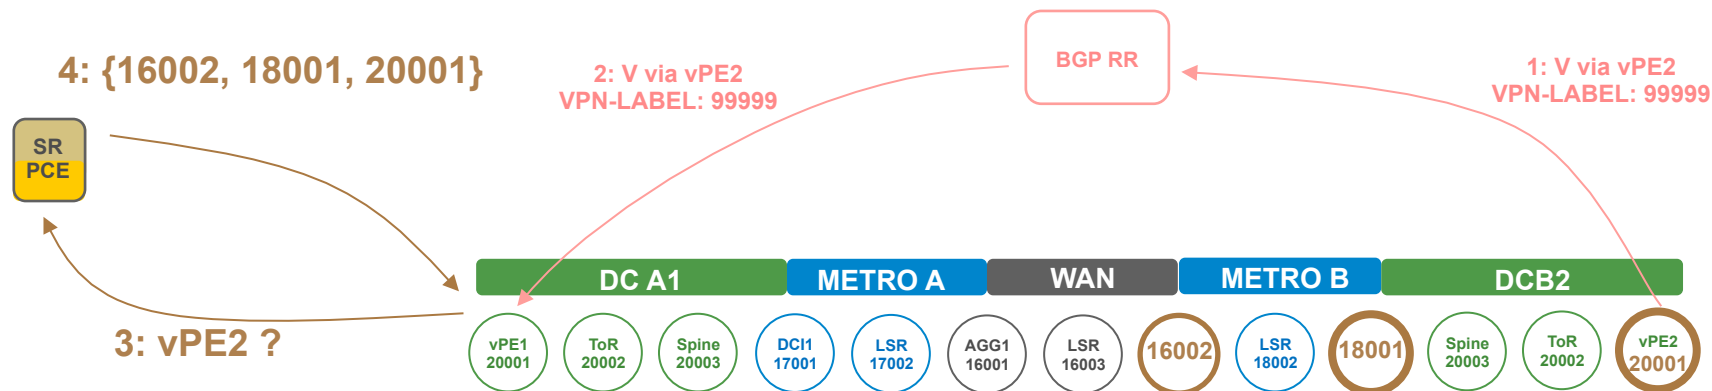
- Multi-Domain topology
  - **Realtime** reactive feed via BGP-LS/ISIS/OSPF from **multiple domains**
  - Including ip address and SID
- Compute: **stateful** with **native** SRTE algorithms

# Service Provisioning



- vPE1 learns about a service route with nhop vPE2
- How does vPE1 reach the nhop?
  - vPE1 only has routes within DC A1 and to the AGG's of the WAN domain
  - Solution: ODN (next slide)

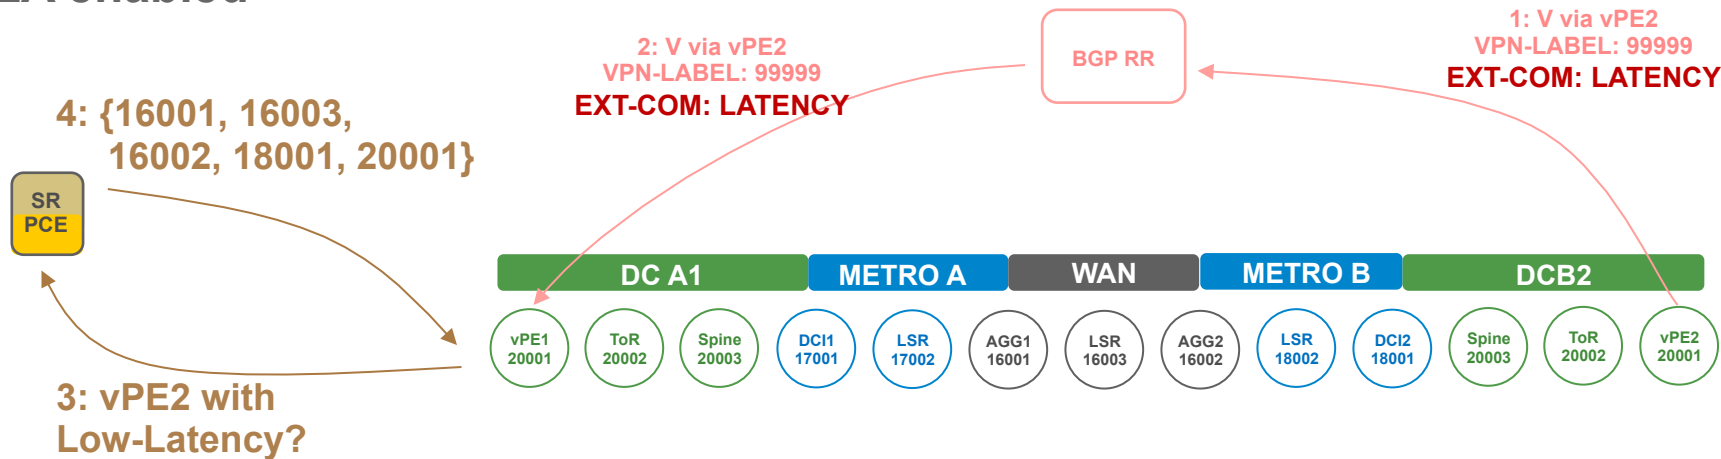
# On-Demand SR Next-Hop Reachability



- vPE1's ODN functionality automatically request a solution from SR-PCE
- Scalable: vPE1 only gets the inter-domain paths that it needs
- Simple: no BGP3107 pushing all routes everywhere

# On-Demand SR Next-Hop

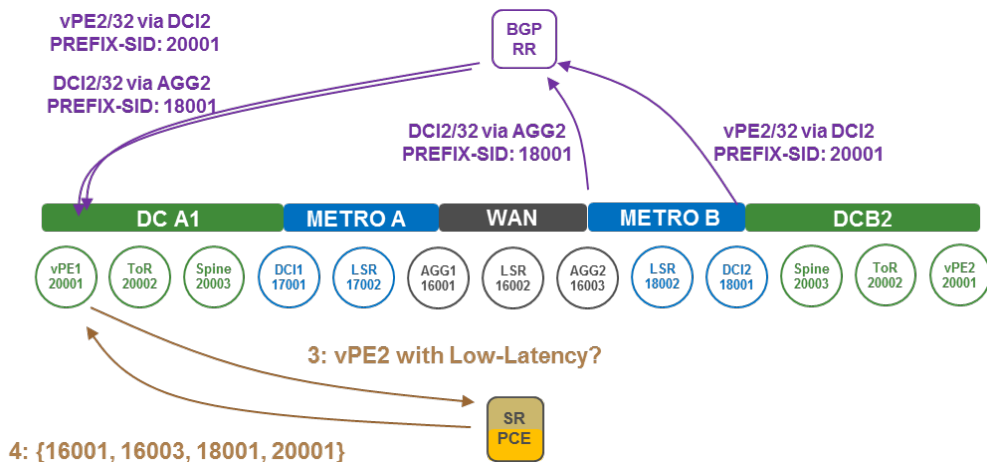
## SLA enabled



- Inter-domain SLA with scale and simplicity

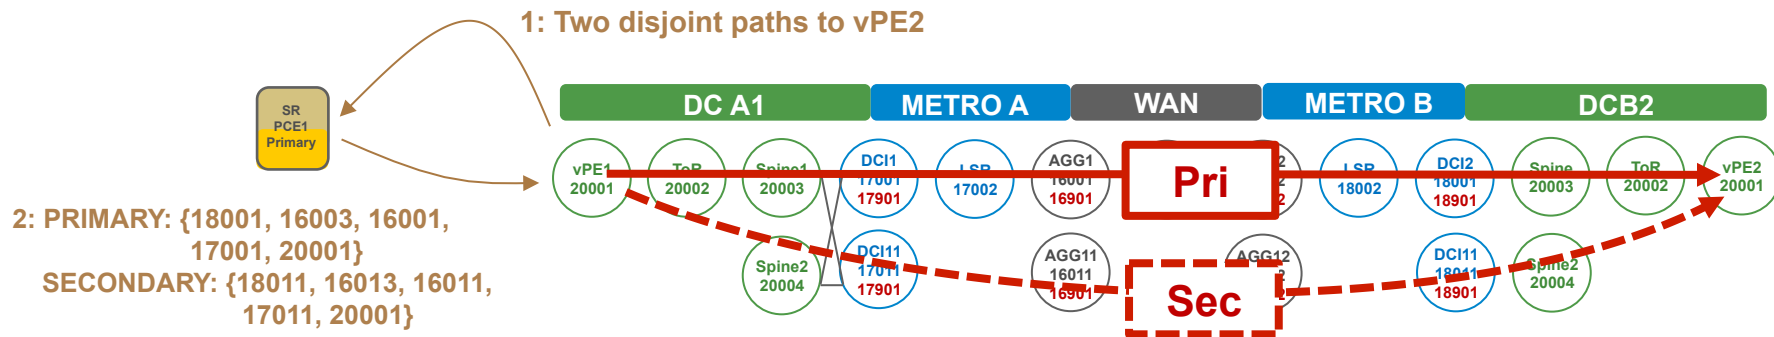
- No RSVP, no midpoint state, no tunnel to configure !!

# Seamless Transition



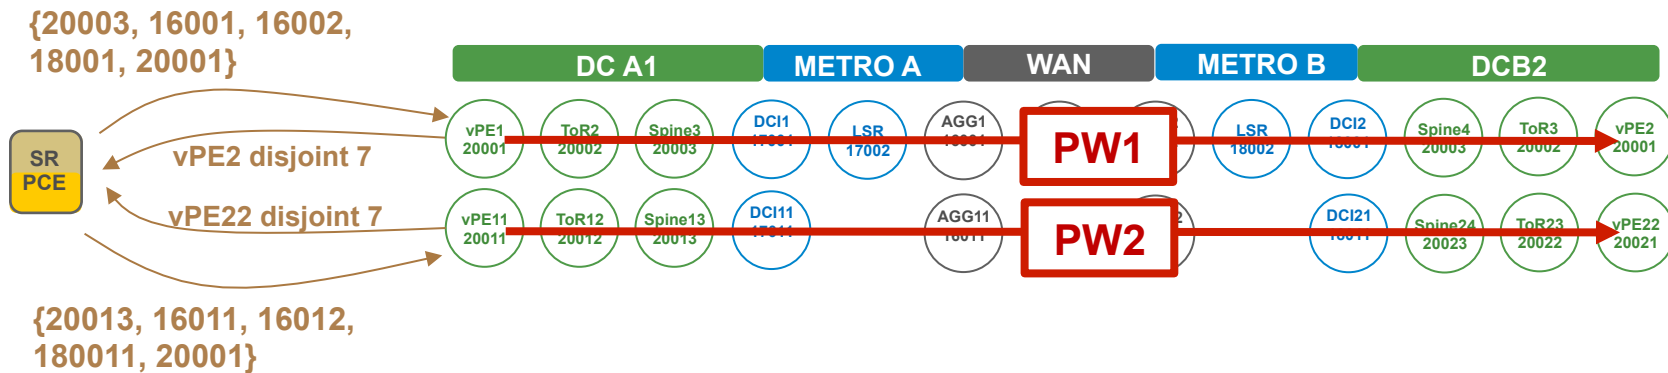
- Best-effort reachability is provided by BGP3107
- ODN and SRTE-PCE provides interdomain reachability with SLA
- Eventually, migration of more/all services over SR PCE

# Inter-Domain PW - Disjoint Primary/Backup



- ODN/SR-PCE automatically computes disjoint primary/sec paths for the PW
- sBFD runs at 3x50msec on each SRTE path
- Upon failure detection of the primary, the secondary SRTE Path is used
- **Inter-domain SLA with scale and simplicity**
  - No RSVP, no midpoint state, no tunnel to configure !!

# Two Disjoint Inter-domain PW's



- ODN/SR-PCE automated compute disjoint paths for PW1 and PW2
- PW1 and PW2 do not share the same headend, neither the same tailend
- **Inter-domain SLA with scale and simplicity**
  - No RSVP, no midpoint state, no tunnel to configure !!

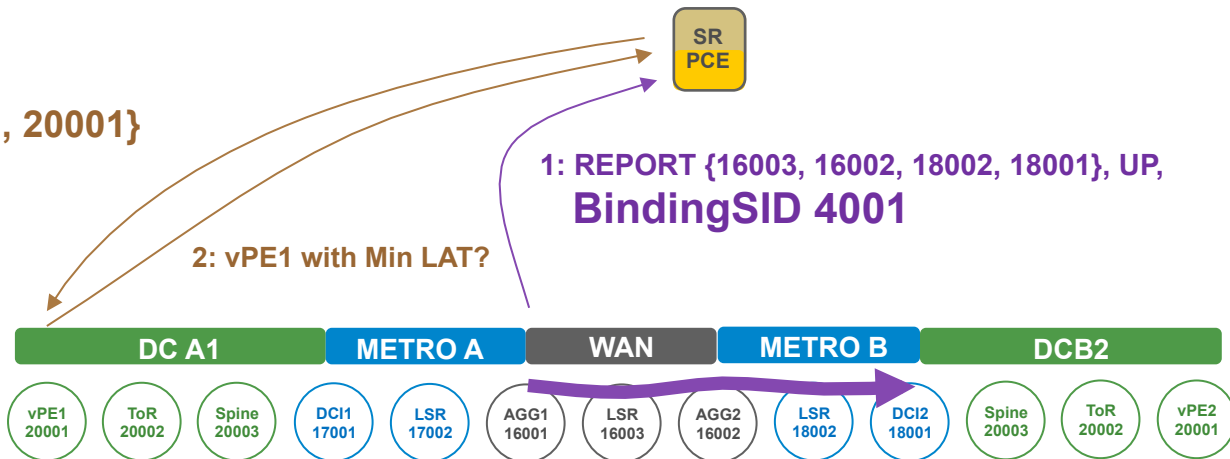
# Binding SID to stitch Policies

3: REPLY {16001, 4001, 20001}

instead of  
{16001,  
16003, 16002, 18002, 18001,  
20001}

2: vPE1 with Min LAT?


1: REPORT {16003, 16002, 18002, 18001}, UP,  
BindingSID 4001



- End-to-end policies can be composed from more basic ones
  - An SRTE policy is bound by default to a Binding SID
  - RSVP-TE tunnels can also be bound to a Binding SID and hence RSVP-TE tunnels can be used within an end-to-end SR policy
- **Shorter SID list and churn isolation between domains**
  - Even if the WAN-MetroA sub-path changes, the related Binding SID 4001 is constant

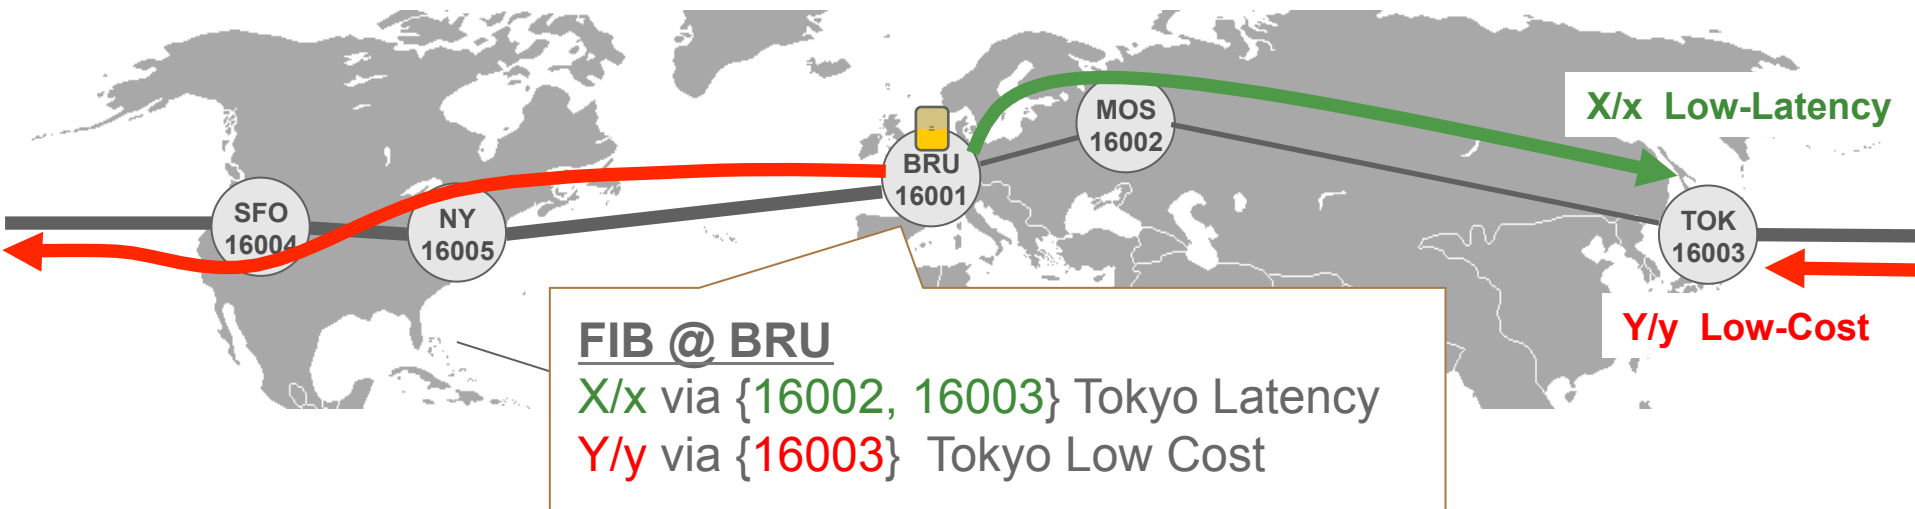


# Fundamentally Distributed



- SR-PCE not to be considered as a single “god” box
- SR-PCE deployment model more like RR
- Different vPE’s can use different pairs of SR PCE’s
- SR PCE preference can either be based on proximity or service

# Fully Distributed SRTE



- Bru learns the routes from Tok and dynamically compute the SRTE policy to nhop
  - Elimination of RSVP signalling, midpoint states and [headend tunnel configuration](#) !
  - [Commonality and Distribution](#): The router's SRTE functionality is same as SRTE-PCE
  - No PBR steering complexity
  - No PBR performance tax



# Topology Independent LFA (TI-LFA)

# TI-LFA - Benefits

- 50msec Protection upon local link, node or SRLG failure
- Simple to operate and understand
  - automatically computed by the router's IGP process (ISIS and OSPF)
  - 100% coverage across any topology
  - predictable (backup = postconvergence)
- Optimum backup path
  - leverages the post-convergence path, planned to carry the traffic
  - avoid any intermediate flap via alternate path
- Incremental deployment
  - also protects LDP and IP traffic



# uLoop Avoidance

# uLoop is a day-1 IP drawback

## Upon link down convergence




Illustration for the post-convergence uloop impacting traffic from 1 to 5 after link45 going down. Default link metric 10

- IP hop-by-hop routing may induce uloop at any topology transition
  - Link up/down, metric up/down

# SR uloop avoidance

- Prevent any uloop upon isolated convergence due to
  - link up/down event
  - metric increase/decrease event
- If multiple back-to-back convergences, fall back to native IP convergence

`microloop avoidance segment-routing`



# SRv6



# SRv6

- SR architecture thought since day 1 for IP
- All the control-plane benefits directly applicable
  - TILFA Link/Node/SRLG, uLoop Avoidance, SRTE, ODN, SRTE-PCE
- Data-plane requires SR extension header
- Smooth deployment plan
  - SRv6 only at a node inserting/removing/updating SRH
  - IPv6 classic at other nodes



# Conclusion

# Conclusion

- SR is fundamental architecture for modern IP network
- Unified Fabric with Policy through DC, Metro and WAN
- Functionality never seen before
- Simplification through Automation and protocol removal
- Strong operator endorsement
- Multi vendor consensus
- Impressive deployment and velocity

# Foundation for modern IP/MPLS networking

- **Simplicity**

- Unified fabric from DC to WAN including aggregation, metro and mobile backhaul
- Set of few well-chosen building blocks
  - > Lightweight extensions to core IP control-plane (BGP, ISIS, OSPF, PCEP)
  - > Remove protocols (LDP, RSVP-TE)
- Intelligent automation
  - > Automated Local protection, per-prefix, any topology, sub-50msec (TI-LFA)
  - > Automated uloop avoidance
  - > Automated traffic matrix
  - > On-Demand SR Next-Hop

# Foundation for modern IP/MPLS networking

- **Functionality: solves unsolved problems**
  - End-to-end policies through domains internal to the SP and external: disjointness, low-latency
  - SRTE algorithm: local to router or centralized PCE
  - Local protection, per-prefix, any topology, sub-50msec (TI-LFA)
  - uLoop avoidance
  - On-Demand SR Next-Hop
  - Application controls the network without the complexity/performance-impact of PBR/DPI
- **Scale**
  - No SRTE midpoint state
  - Not even any SRTE headend configuration (On-Demand SR Next-Hop)
  - Binding SID for compressed SID list length

# Foundation for modern IP/MPLS networking

- Seamless Deployment
  - SR/LDP interworking
  - SR/RSVP-TE interworking (binding SID)
  - Ship-in-the-night co-existence
  - SRTE policies progressively migrated from the edge to the DCI to the ToR to the vswitch to the host
  - SW upgrade (reuse existing platform)
- Decoupling data and control planes
  - Architecture natively thought to accommodate decoupled data and control planes
- Cost
  - Architecture natively thought to accommodate merchant silicon
  - Automation leads to opex saving
  - Tactical BW optimization

# Industry Consensus

- Huge operator adoption
  - WEB, SP, Enterprise
  - Core, Edge, Aggregation, DC
- Standardized
  - Cisco is leading all the innovation and productization and is committed since day1 on disclosing all the elements of the architecture for IETF standardization and as well in terms of very detailed public tutorial
    - > <http://www.segment-routing.net/home/ietf>
    - > <http://www.segment-routing.net/home/tutorial>
- Multi-vendor consensus
  - Nokia, Ericsson, Juniper, Huawei, Arista
  - Several rounds of proven interops
- Open
  - ONOS...

# Stay Informed

- <http://www.segment-routing.net/>
- [ask-segment-routing@cisco.com](mailto:ask-segment-routing@cisco.com)



Thank you





**CISCO**

*TOMORROW starts here.*