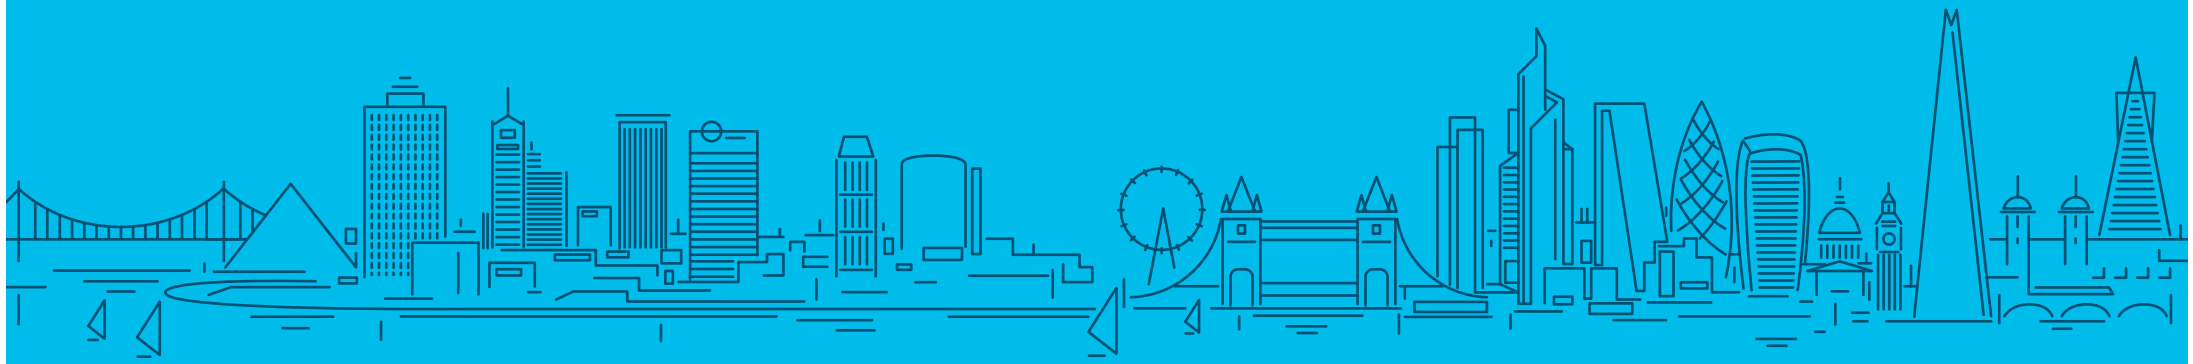

Engineering the new IP Transport

Derek Ong

Systems Engineer

April 2019

Is your current network ready to support these?

Enhanced Broadband Speed
1000x BW, 10-100x

Super Fast Response
50 msec – 300 msec

Ultra Reliable Low Latency Communication
~1 – 25 msec Latency

Massive Machine Type Communications
1000x Density

Secure Remote Car Software Update

10 – 100M lines of code and hundreds of subsystems

Vehicle updates, telematics, and infotainment

Network-as-a-Service Public Sector Private Networks

Police, fire, hospitals with strict SLAs and Security

Smart City Enablement

Smart-sensor enabled

Waste and recycling,, parking, smart grid, homes

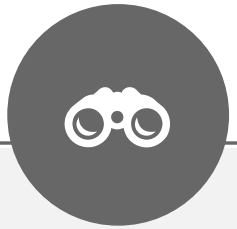

Augmented / Virtual Reality Delivery

Augmented, virtual, and mixed reality for learning, gaming, 4K/8K

Video enablement required

Today: Operational Domains and Layers of Protocols

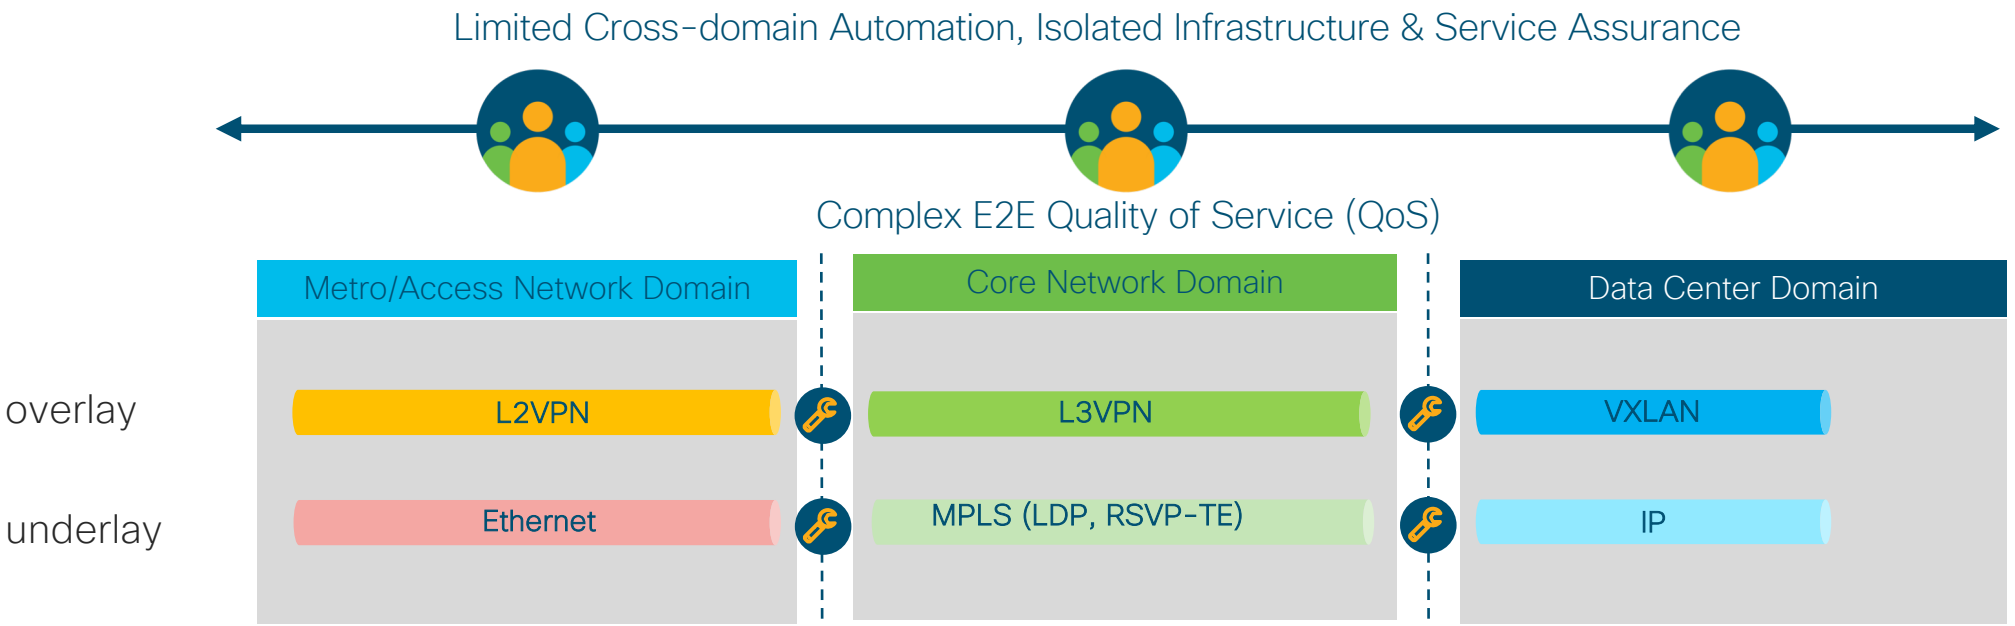

Automated end-to-end service provisioning can be tricky

- Multiple network domains under different management teams
- Manual operations
- Different types of underlay and overlay networks

Target: End-to-End and Domain automation

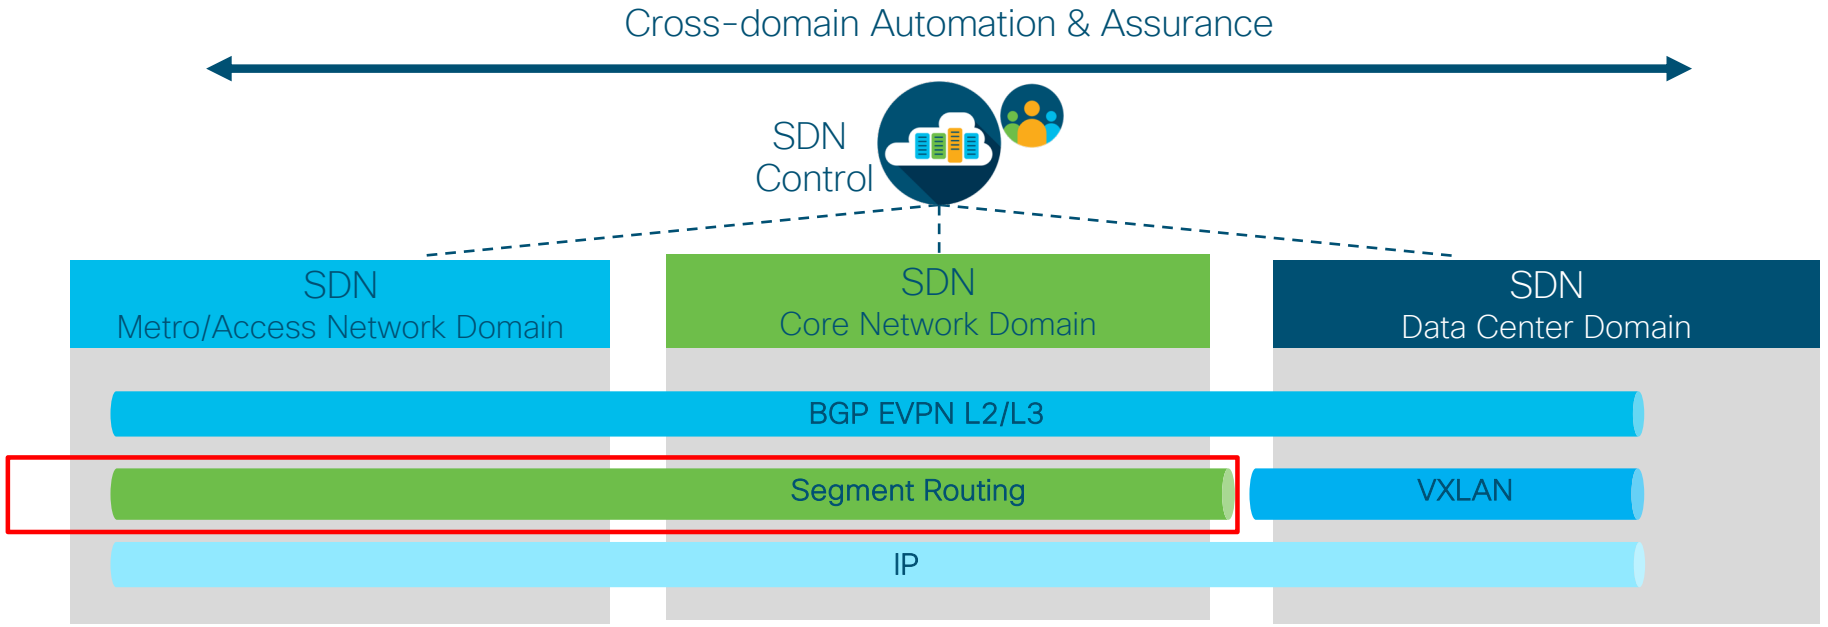

End-to-end service provisioning is more easily automated

- Common underlay and overlay network technologies
- Automated underlay and overlay operations
- Seamless integration between domains

Segment Routing Basics

An LS IGP protocol extension bringing network simplification/optimization

- No LDP
  - Lighter protocol suite
    - Less adjacencies, less states to maintain
- No IGP to LDP synchronization
  - Eliminates delays in activating a path
- Topology independent fast reroute using post convergence back up path
  - 50 ms protection
  - no microloops
  - 100% coverage of network topologies

Easy troubleshooting

An IP/MPLS architecture designed with SDN in mind

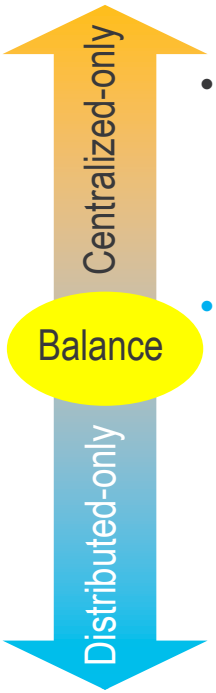

- Right balance between distributed intelligence and centralized optimization and programming
  - SR-TE
- Wide applications
  - (SP, OTT/Web, GET) across (WAN, Metro/Agg, DC)
  - MPLS and IPv6 dataplanes
  - SDN controller

Segment Routing Standardization

- IETF standardization in SPRING working group
- Protocol extensions progressing in multiple groups
  - IS-IS
  - OSPF
  - PCE
  - IDR
  - 6MAN
- Broad vendor and customer support

Sample IETF Documents
Segment Routing Architecture RFC 8402
Source Packet Routing in Networking (SPRING) Problem Statement and Requirements RFC 7855
Use Cases for IPv6 Source Packet Routing in Networking (SPRING) RFC 8354
Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane RFC 8287
Packet Loss and Delay Measurement for MPLS Networks RFC 6374
IS-IS Traffic Engineering (TE) Metric Extensions RFC 7810
OSPF Traffic Engineering (TE) Metric Extensions RFC 7471
UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks RFC 7876

Close to 30 IETF drafts in progress
(http://www.segment-routing.net/ietf/)

Segment Routing Protocol Efficiency

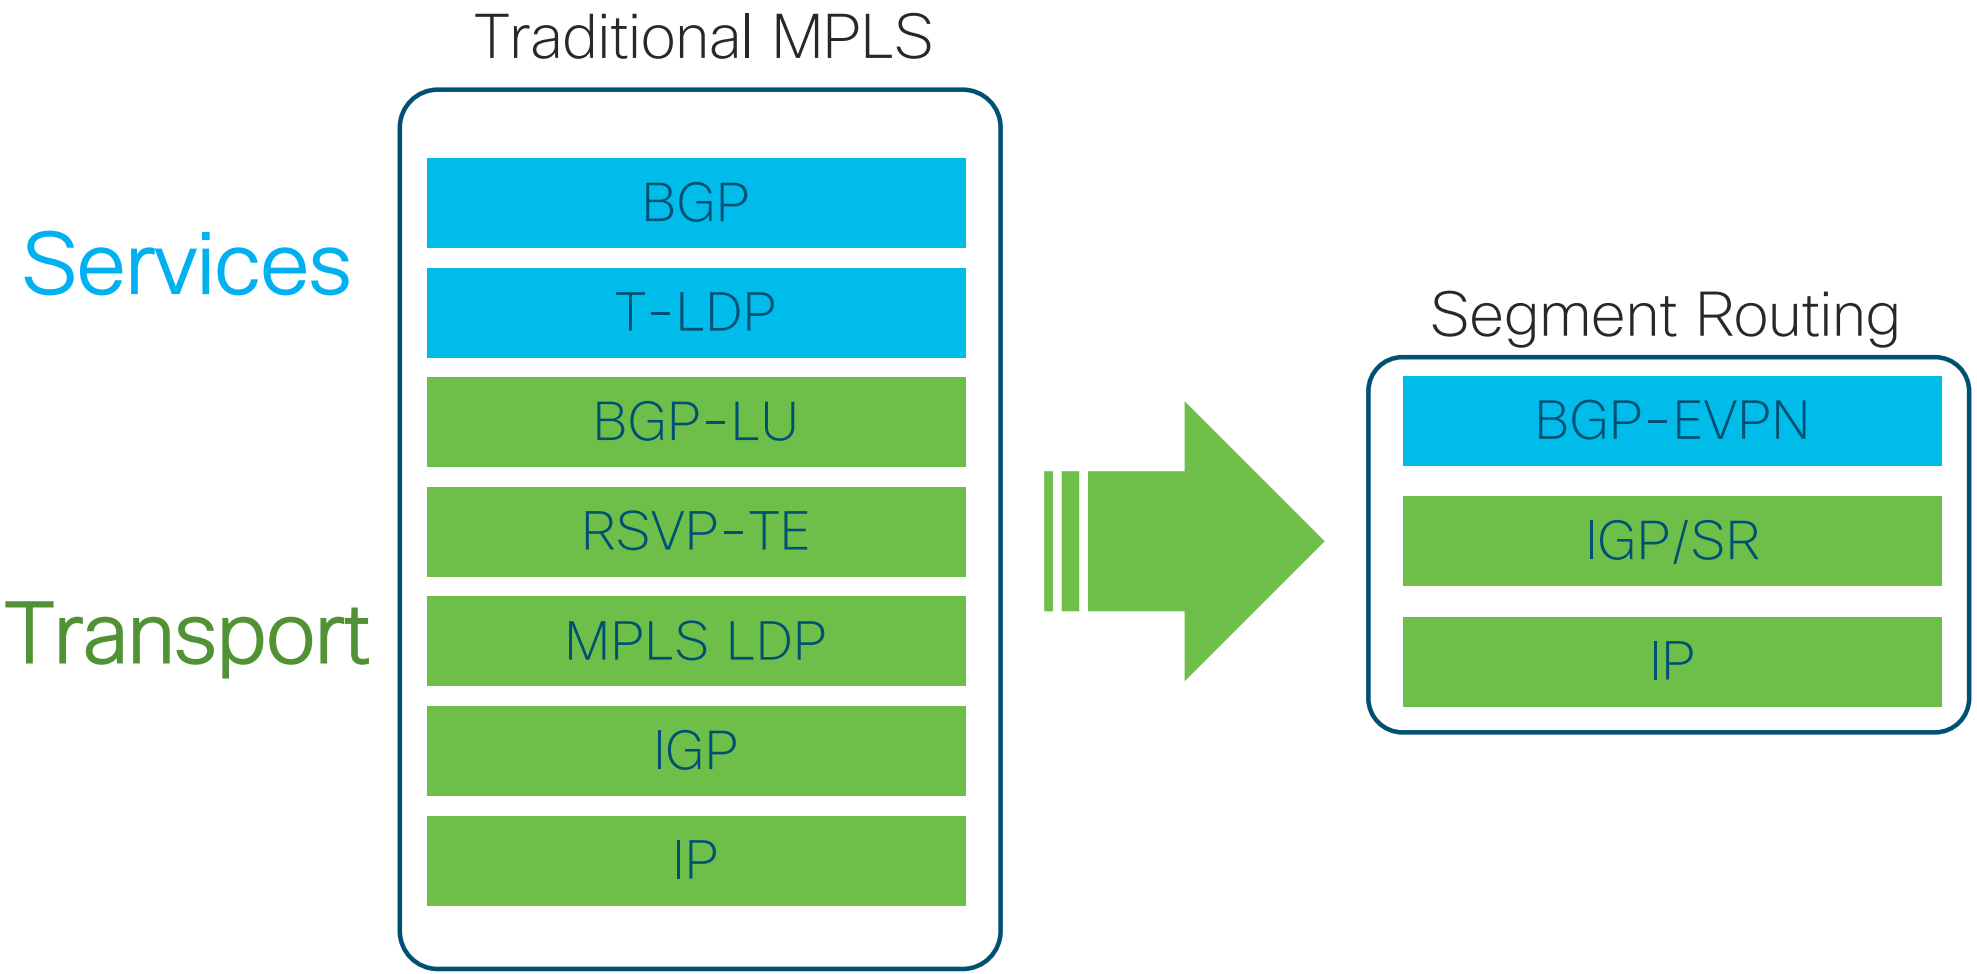

Segment Routing – Source Routing

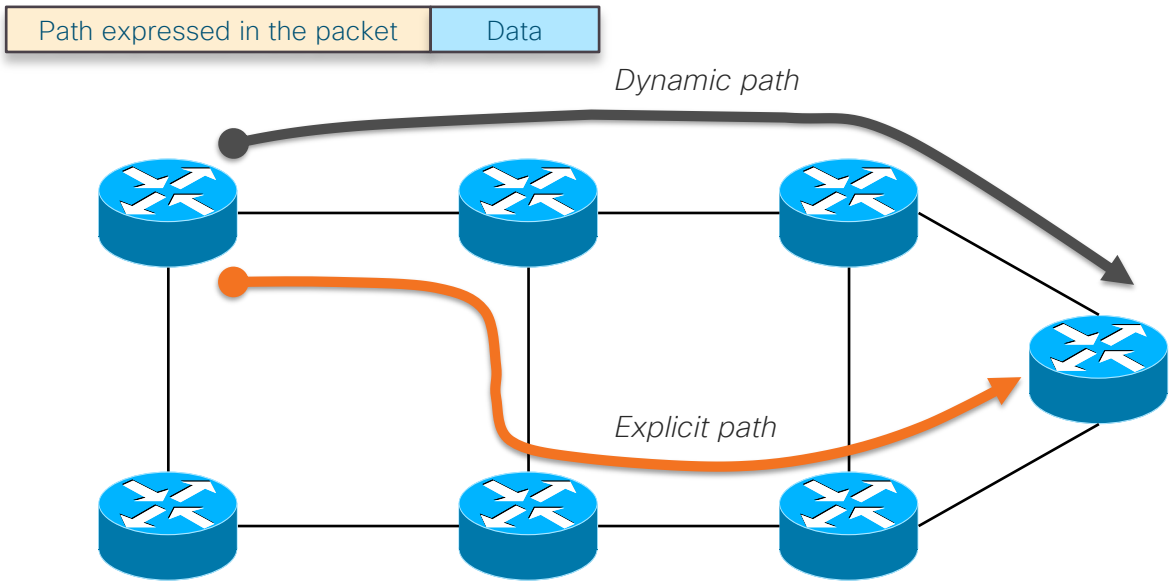

Data Plane

- MPLS (segment labels)
- IPv6 (+ SR extension header)

Control Plane

- Routing protocols with extensions (IS-IS,OSPF, BGP)
- SDN controller (BGP , PCEP, NETCONF/YANG)

Path options

- Dynamic (SPF)
- Explicit (expressed in the packet)

Global Segments – Global Label Indexes

- Global Segments always distributed as a label range (SRGB) + Index
  - Index must be unique in Segment Routing Domain
- Best practice: same SRGB on all nodes
  - “Global model”, requested by all operators
  - Global Segments are global label values, simplifying network operations
  - Default SRGB: 16,000 – 23,999
    - Other vendors also use this label range

IGP Prefix Segment

- Shortest-path to the IGP prefix
  - Equal Cost MultiPath (ECMP)-aware
- Global Segment
- Label = 16000 + Index
  - Advertised as index
- Distributed by ISIS/OSPF

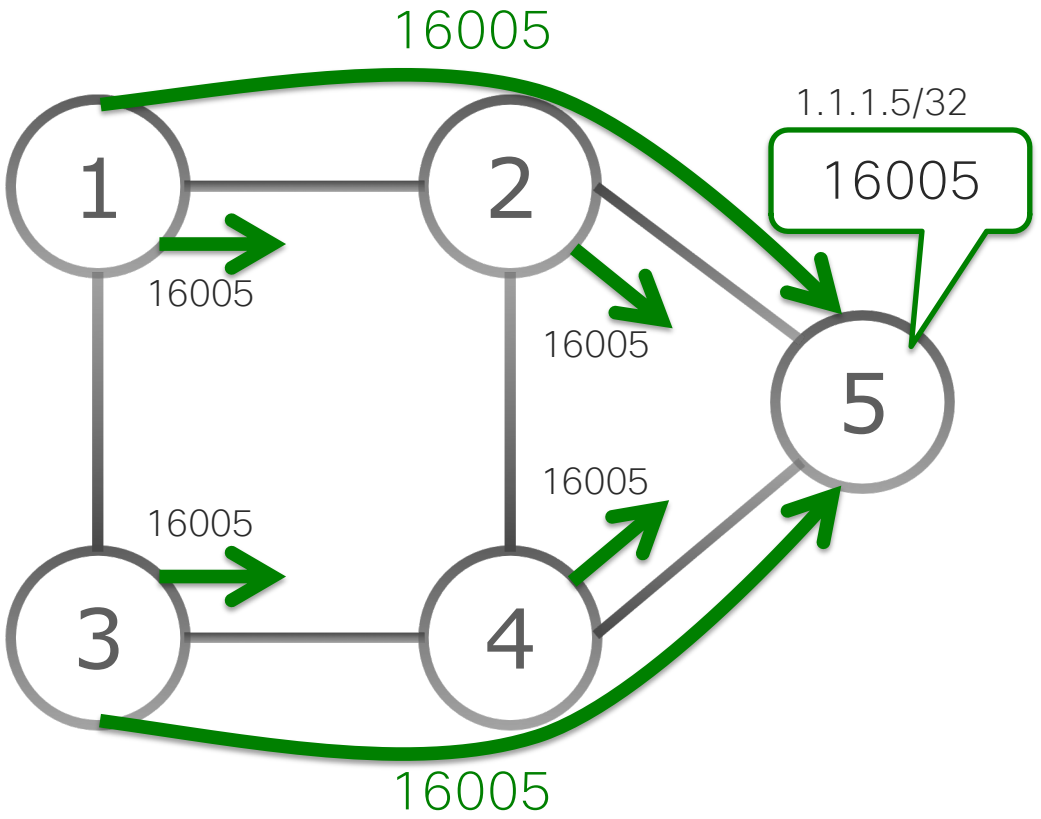

All nodes use default SRGB
16,000 – 23,999

IGP Adjacency Segment

- Forward on the IGP adjacency
- Local Segment
- Advertised as label value
- Distributed by ISIS/OSPF

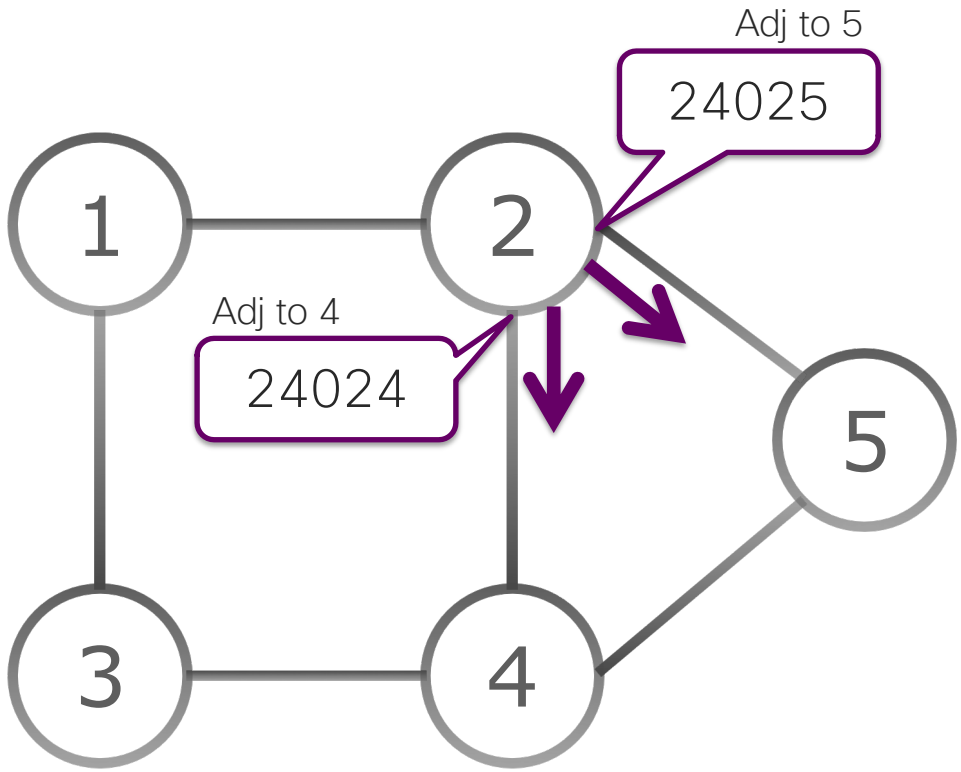

All nodes use default SRGB
16,000 – 23,999

Combining IGP Segments

All nodes use default SRGB
16,000 – 23,999

- Steer traffic on any path through the network
- Path is specified by list of segments in packet header, a stack of labels
- No path is signaled
- No per-flow state is created
- Single protocol: IS-IS or OSPF

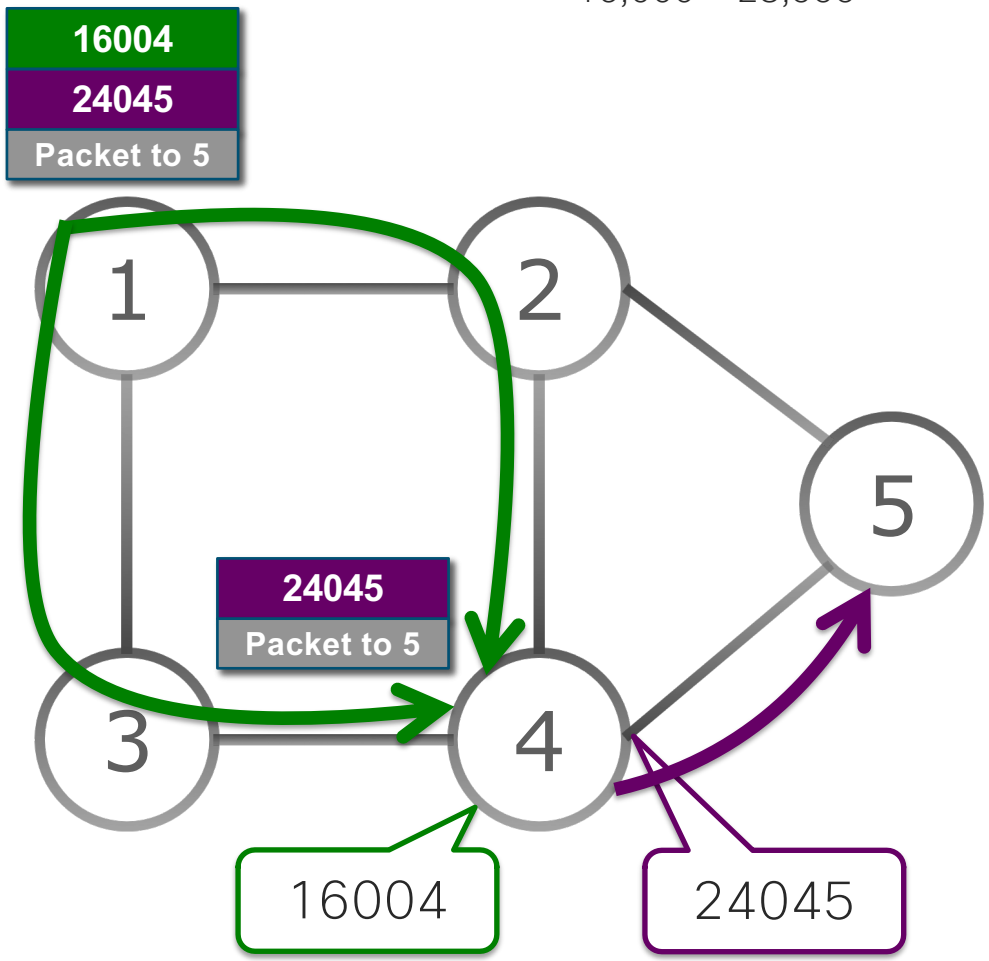

Topology-Independent LFA (TI-LFA FRR)

- 50msec FRR in any topology
  - Link, Node, or SRLG
- IGP Automated
  - No LDP, no RSVP-TE
- Optimum
  - Post-convergence path
- No midpoint backup state

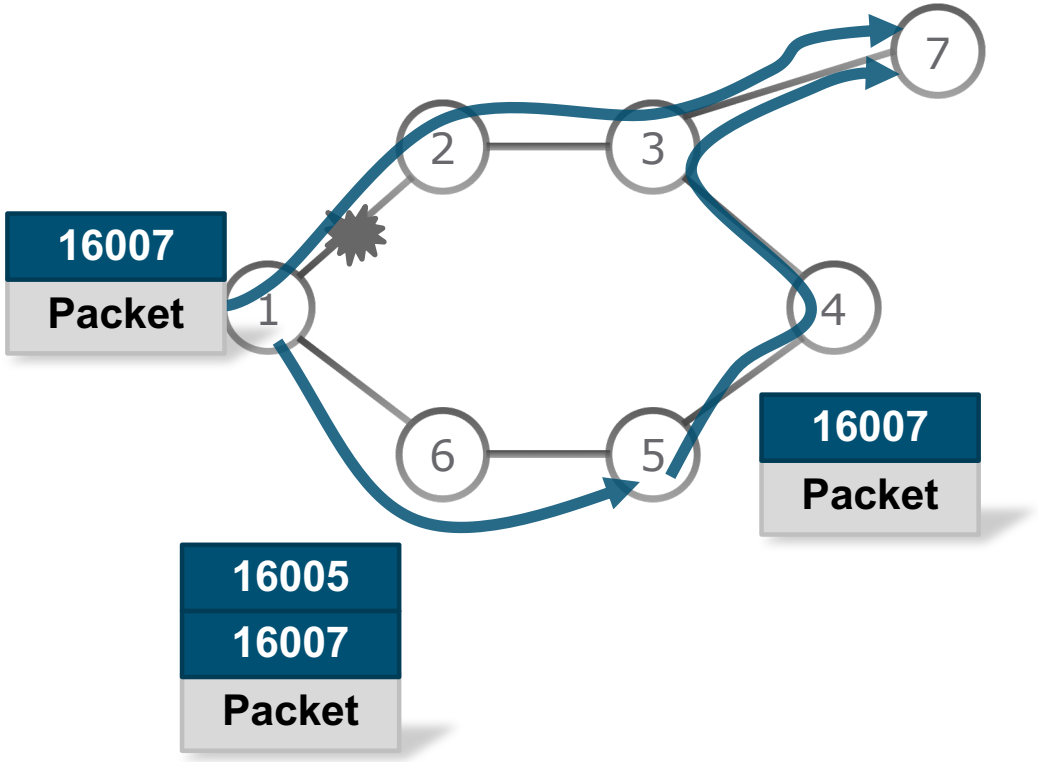

Multi-Domain Topology

- SR Path Computation Element (PCE)
- PCE collects via BGP-LS
  - IGP segments
  - BGP segments
  - Topology

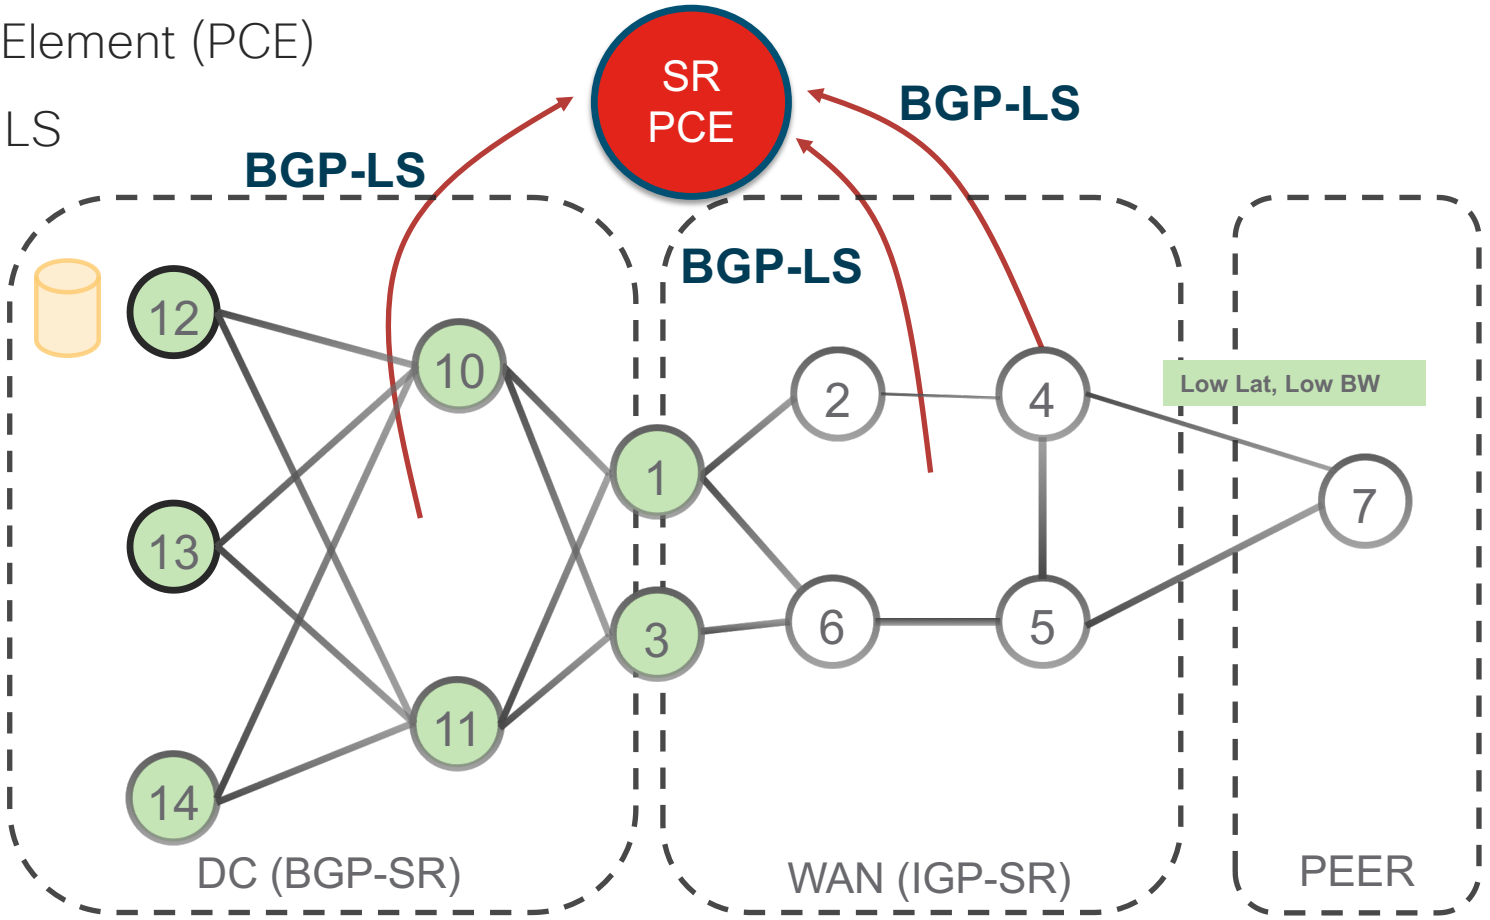

Creating a BGP SR-TE Policy

- Controller programs an SR TE Policy at ingress
- SR TE Policy defines the set of explicit paths from ingress to policy endpoint
  - Weighted ECMP among multiple paths

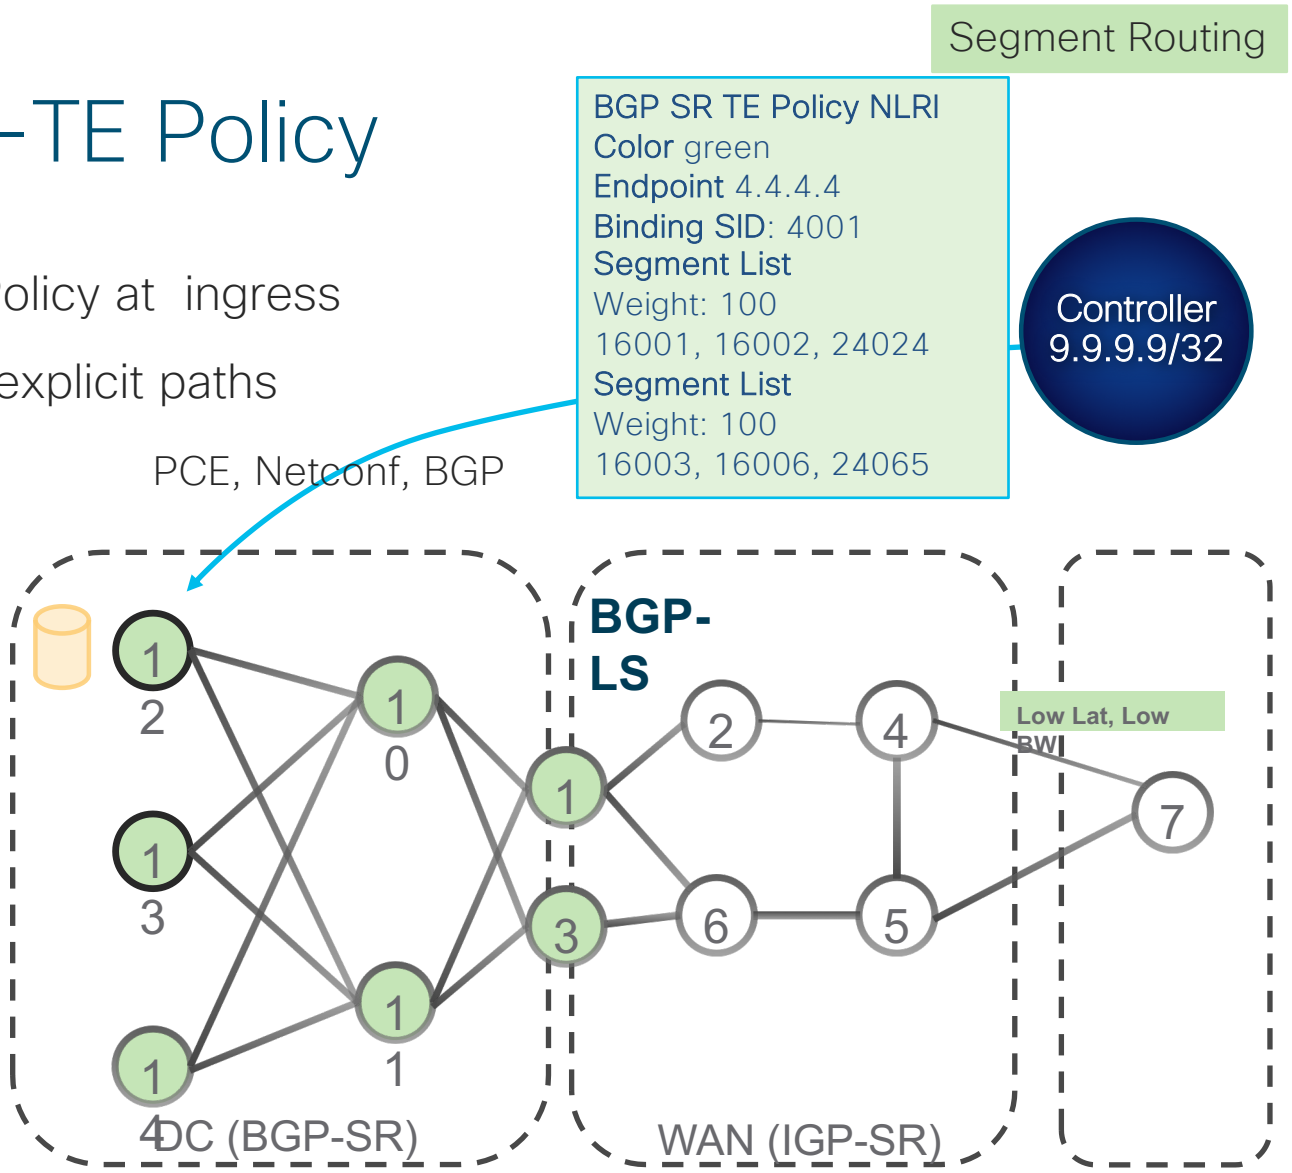

Slicing in transport network

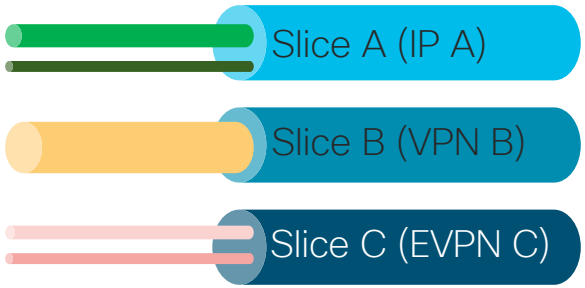

Traffic isolation:
Virtual Forwarding Topology for different slices

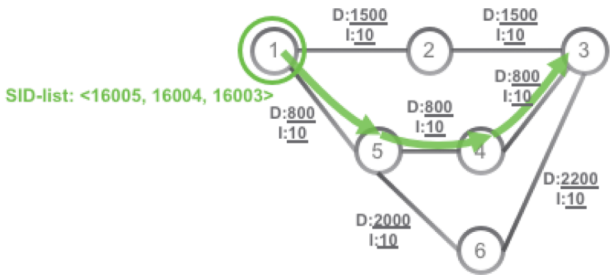

Low-delay path:
Segment routing head-end computes a SID-list of the shortest-path according to delay

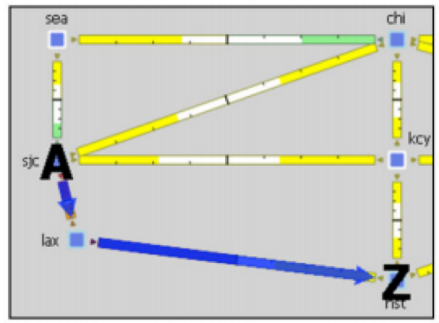

Bandwidth Optimization:
Segment routing traffic engineering can find best path to meet bandwidth requirements

Flex-Algo: SR IGP Flexible Algorithms

(draft-ietf-lsr-flex-algo-01.txt)
https://datatracker.ietf.org/doc/draft-ietf-lsr-flex-algo/

Segment Routing

What it Does

Customized IGP algorithms defined by operator for intent-based instantiation of traffic Engineering

Minimization of metrics: IGP, delay.
Exclusion of properties: link-affinity, SRLG.

Benefits

Simplicity and Automation

IGP-computed TE-path from anywhere to anywhere
Sub-50msec protection (TILFA) optimized per Flex-Algorithm plane.

Scalability

Single SID (instead of label stack) to enforce TE path.
Single prefix segment can participate in many Flex-Algos.

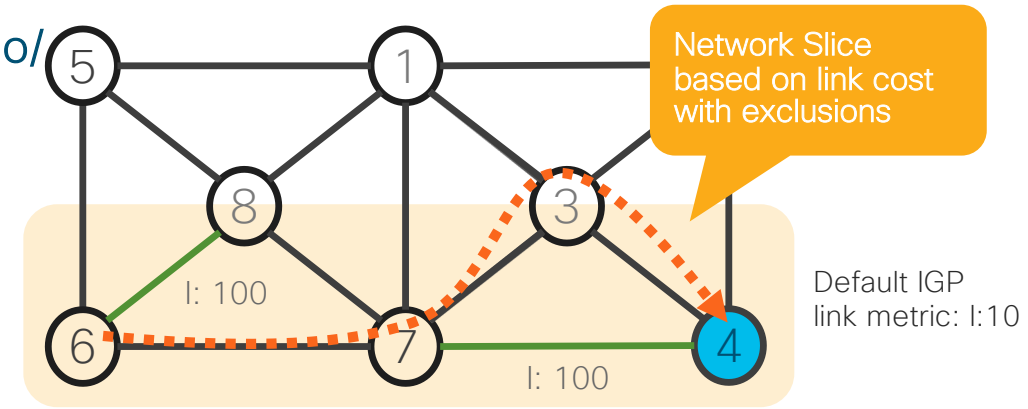

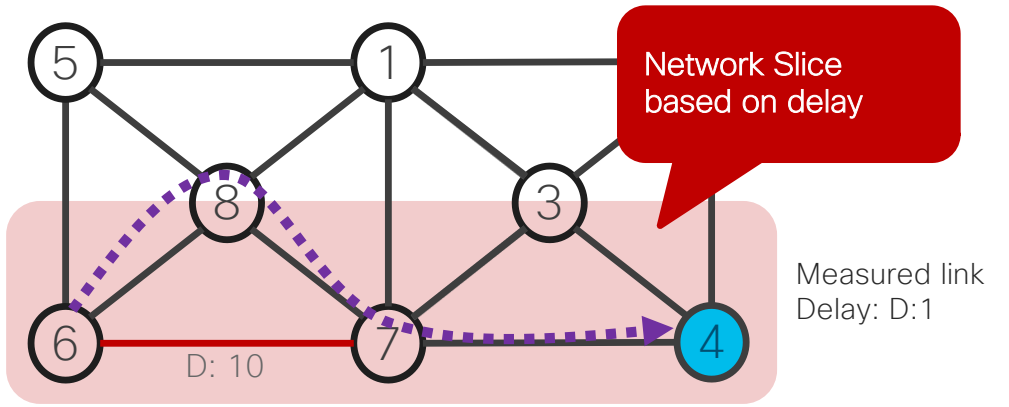

Same destination, different colour, different optimisation goals and constraints

Flex-Algo: Multi-Plane Topologies

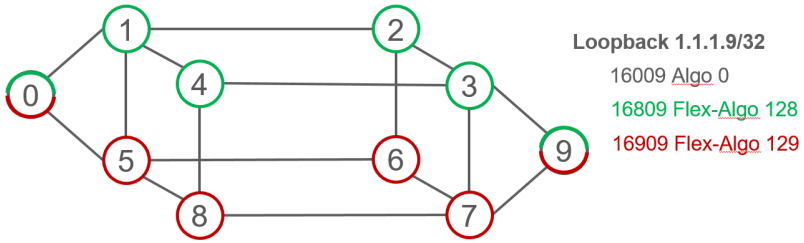

- All the nodes support Algo 0: minimize IGP metric
- Green nodes also support 128: minimize IGP metric
- Red nodes also support 129: minimize Delay

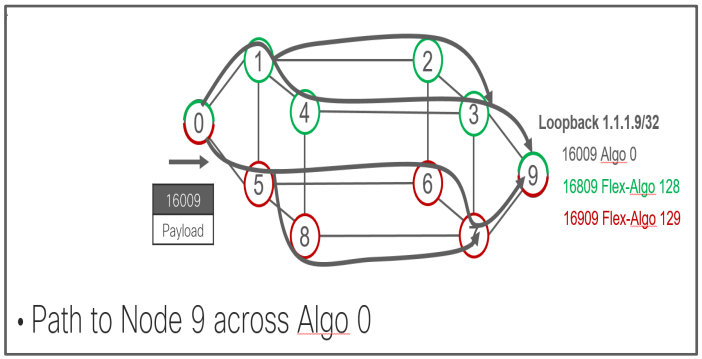

- Path to Node 9 across Algo 0

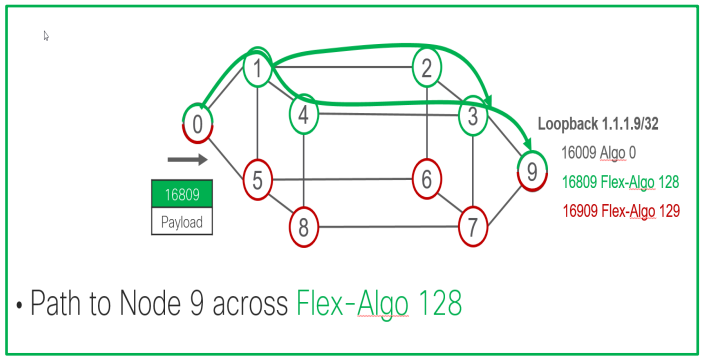

- Path to Node 9 across Flex-Algo 128

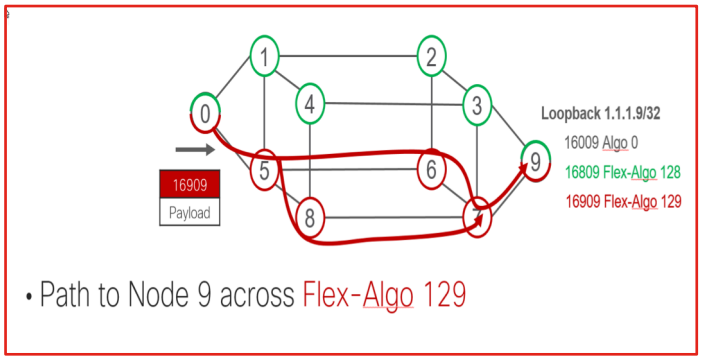

- Path to Node 9 across Flex-Algo 129

Segment Routing: Key capabilities

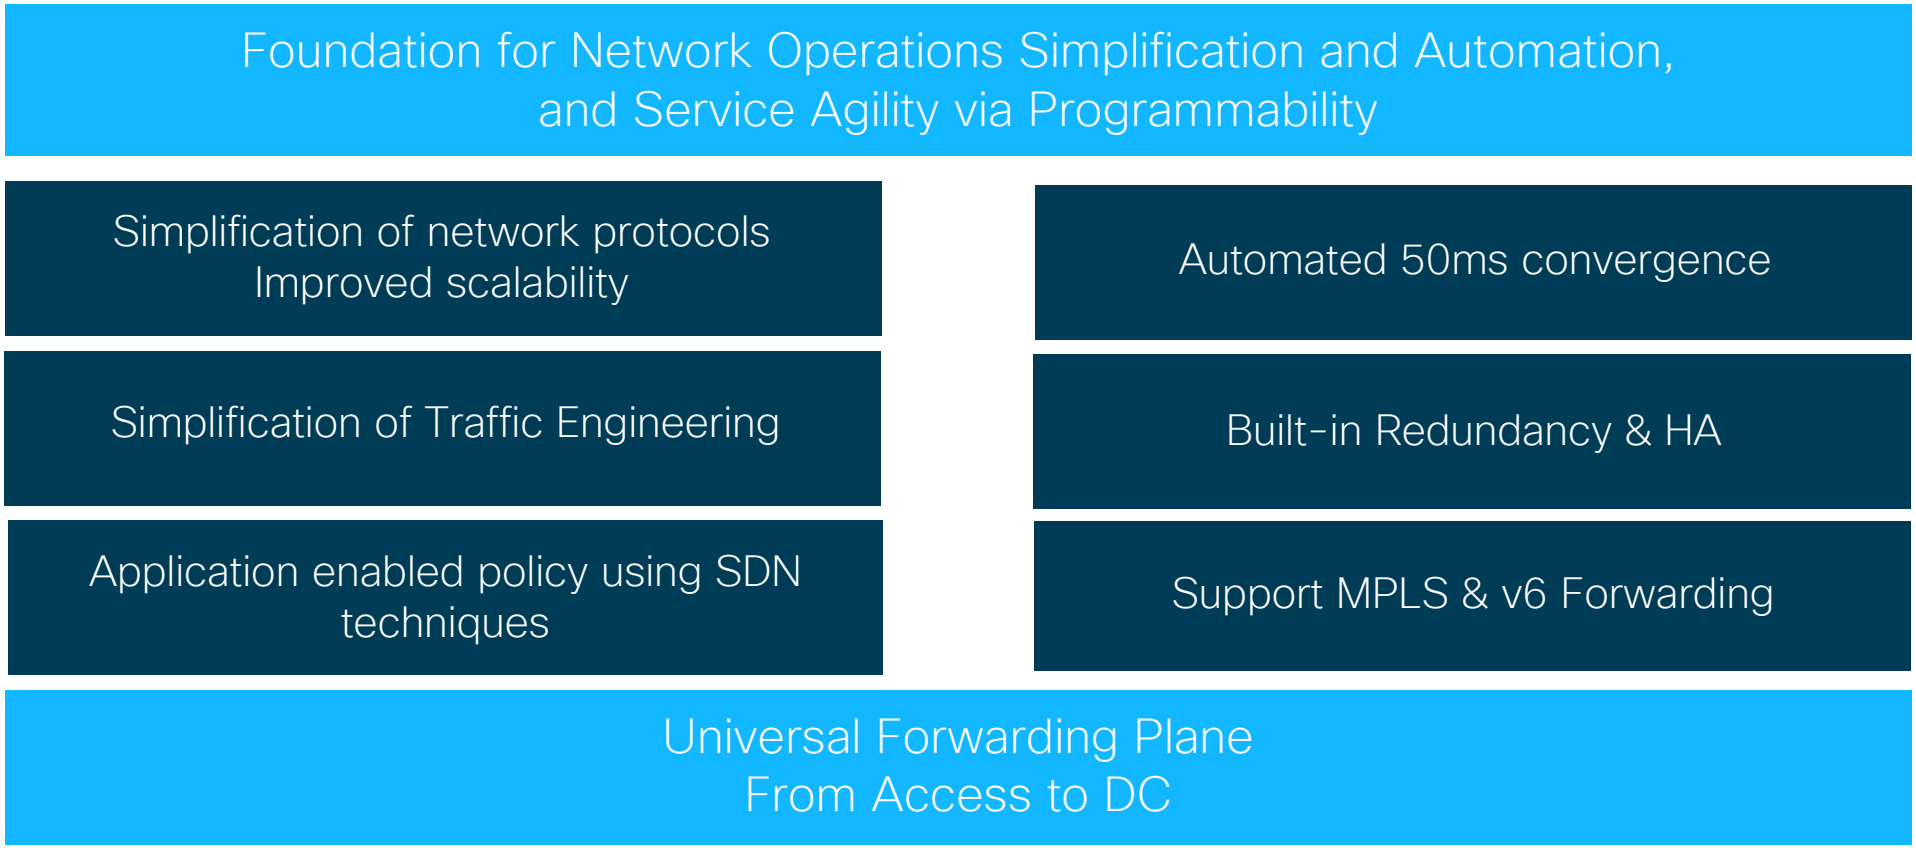

Summary

- IP Transport requires evolution to support new services
- **Network Slice = efficient use in a shared transport**
- **Segment Routing = Simplicity = Automation = Scale**

Thank you. Enjoy SGNOG. :-)