Segment Routing

LACNOG 16 – San José, Costa Rica

Enrique Dávila
Services Technical Leader

Agenda

- Technology Overview
- Use Cases
- Control and Data Plane
- Traffic Protection
- Conclusions

Segment Routing

- **Source Routing**
  - the source chooses a path and encodes it in the packet header as an ordered list of segments
  - the rest of the network executes the encoded instructions without any further per-flow state
- **Segment**: an identifier for any type of instruction
  - forwarding or service

IGP Prefix Segment

- Shortest-path to the IGP prefix
- Global
- 16000 + Index
- Signaled by ISIS/OSPF

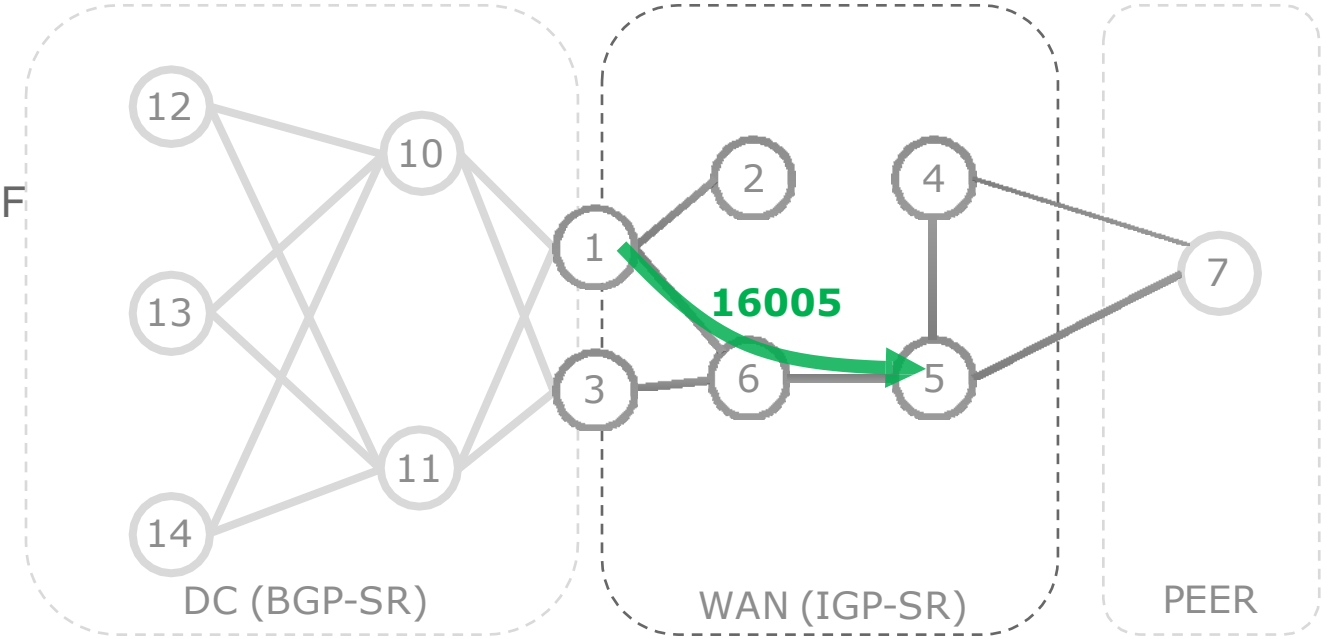

IGP Adjacency Segment

- Forward on the IGP adjacency
- Local
- 1XY
  - X is the “from”
  - Y is the “to”
- Signaled by ISIS/OSPF

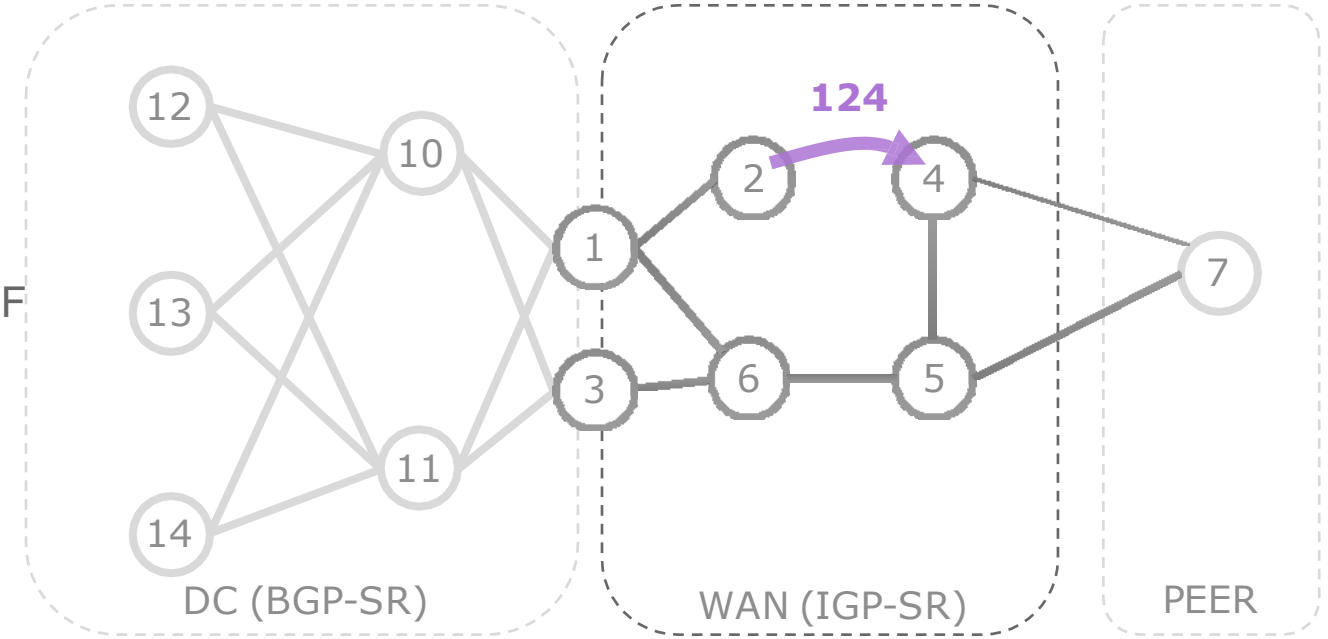

BGP Prefix Segment

- Shortest-path to the BGP prefix
- Global
- 16000 + Index
- Signaled by BGP

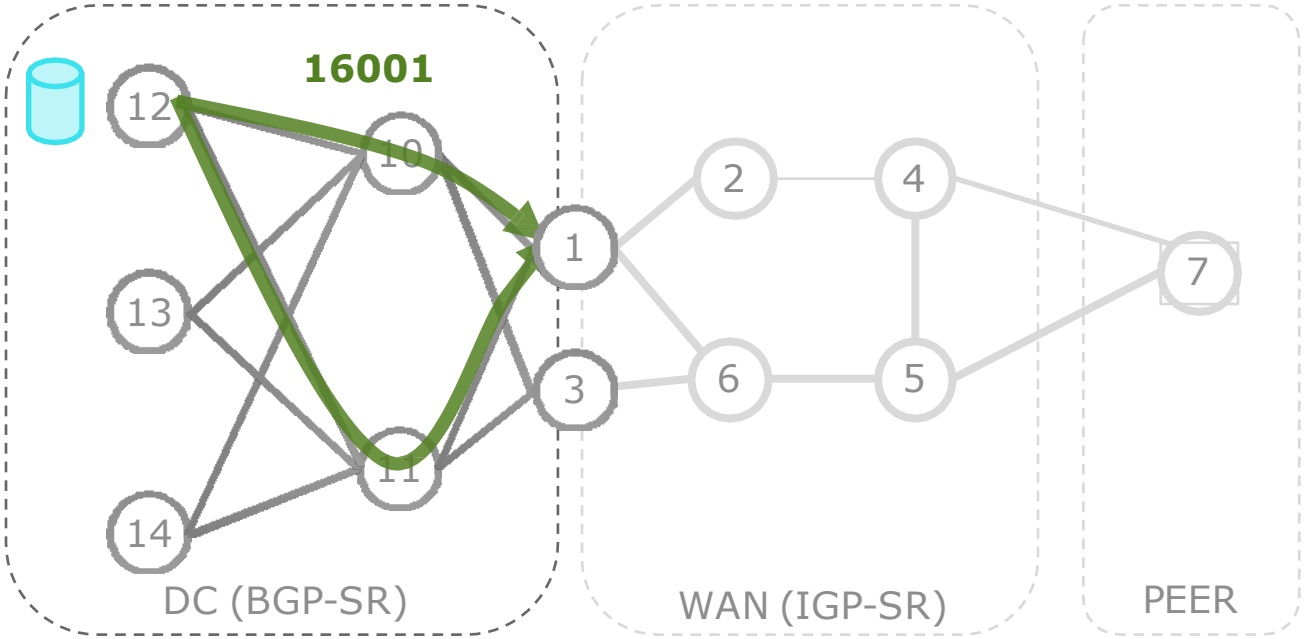

BGP Peering Segment

- Forward to the BGP peer
- Local
- 1XY
  - X is the “from”
  - Y is the “to”
- Signaled by BGP-LS (topology information) to the controller

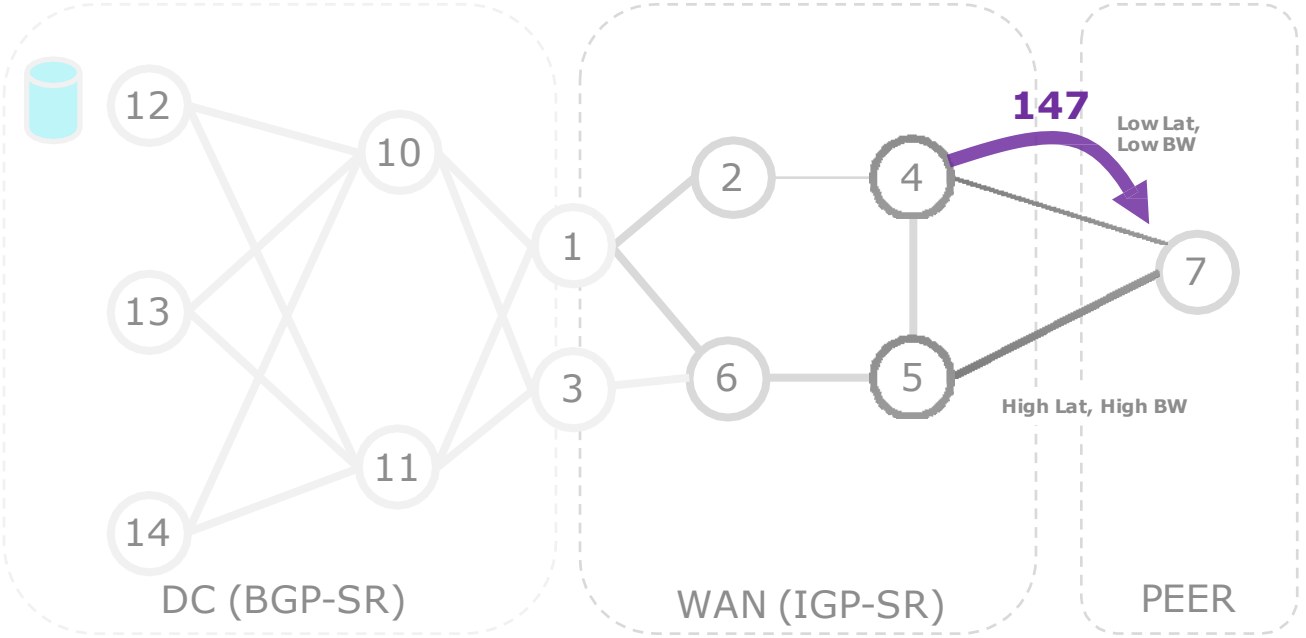

WAN Controller

- WAE collects via BGP-LS
  - IGP segments
  - BGP segments
  - Topology

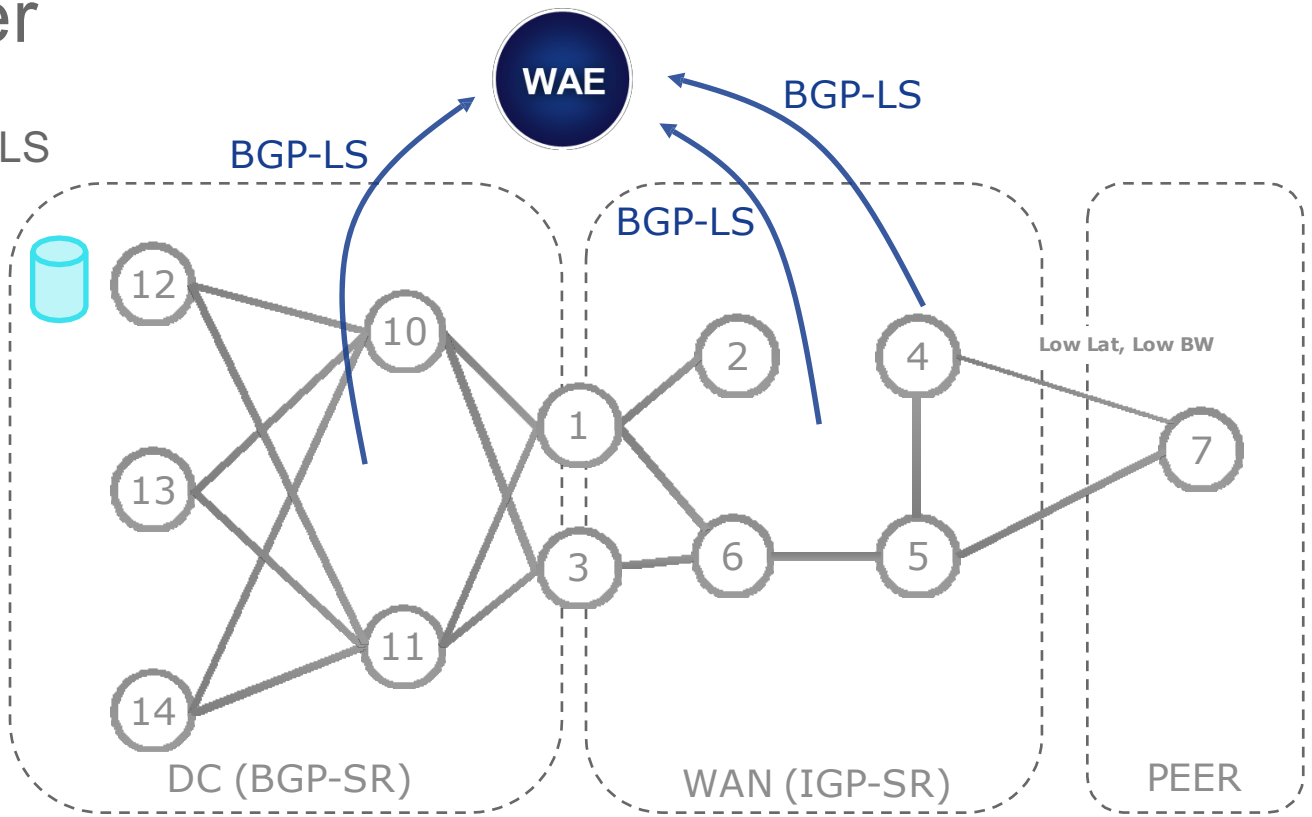

An end-to-end path as a list of segments

- WAE computes that the green path can be encoded as
  - 16001
  - 16002
  - 124
  - 147
- WAE programs a single per-flow state to create an application-engineered end-to-end policy

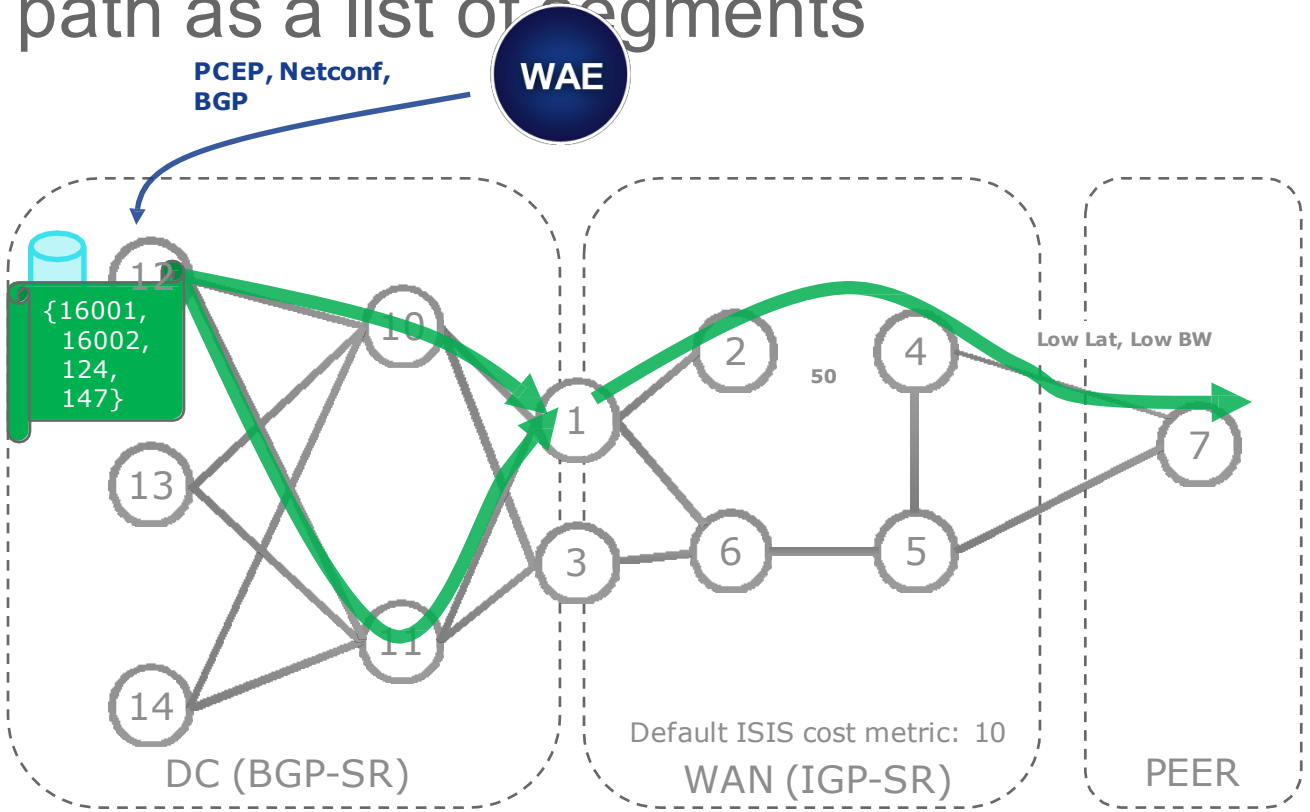

Segment Routing Standardization

- 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 (draft-ietf-spring-segment-routing)
Problem Statement and Requirements (draft-ietf-spring-problem-statement)
IPv6 SPRING Use Cases (draft-ietf-spring-ipv6-use-cases)
Segment Routing Use Cases (draft-filsfils-spring-segment-routing-use-cases)
Topology Independent Fast Reroute using Segment Routing (draft-francois-spring-segment-routing-ti-lfa)
IS-IS Extensions for Segment Routing (draft-ietf-isis-segment-routing-extensions)
OSPF Extensions for Segment Routing (draft-ietf-ospf-segment-routing-extensions)
PCEP Extensions for Segment Routing (draft-ietf-pce-segment-routing)

- Close to 30 IETF drafts in progress

Use Cases

IPv4/6 VPN/Service transport

- IGP only
  - No LDP, no RSVP-TE
- ECMP

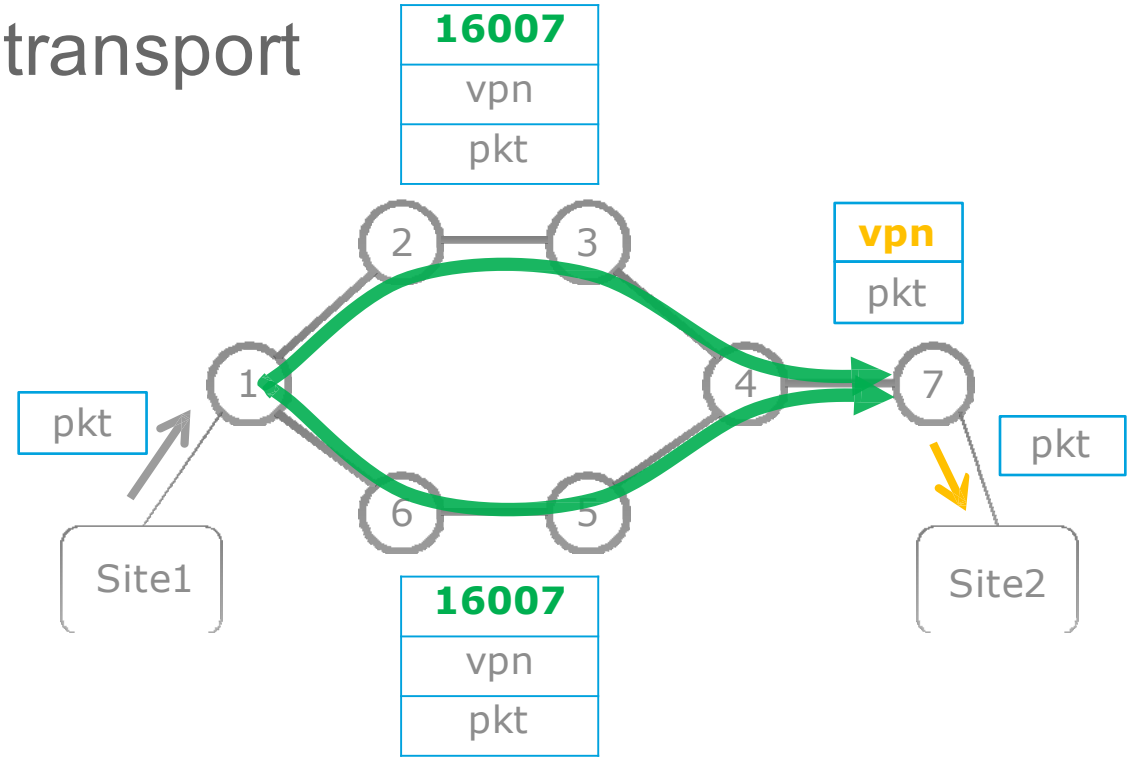

Seamless interworking with LDP

- Seamless deployment

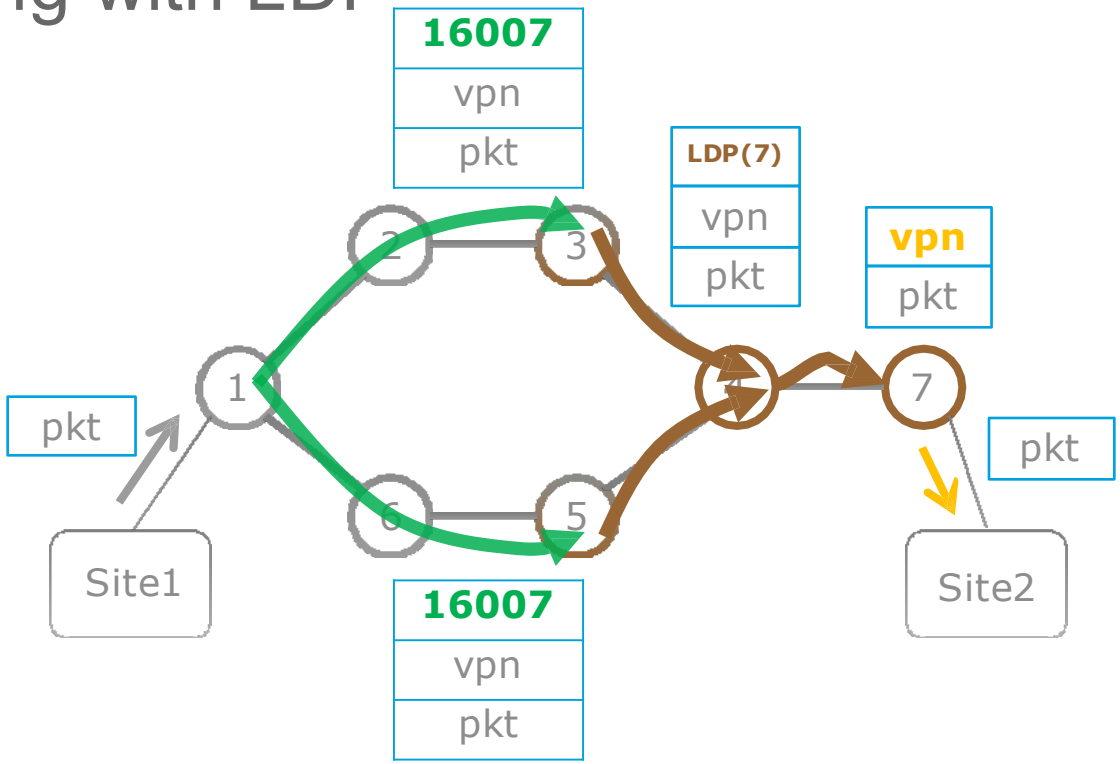

Topology-Independent LFA (TI-LFA FRR)

- 50msec FRR in any topology
- IGP Automated
  - No LDP, no RSVP-TE
- Optimum
  - Post-convergence path
- No midpoint backup state
- Detailed operator report
  - S. Litkowski, B. Decraene, Orange
- Mate Design
  - How many backup segments
  - Capacity analysis

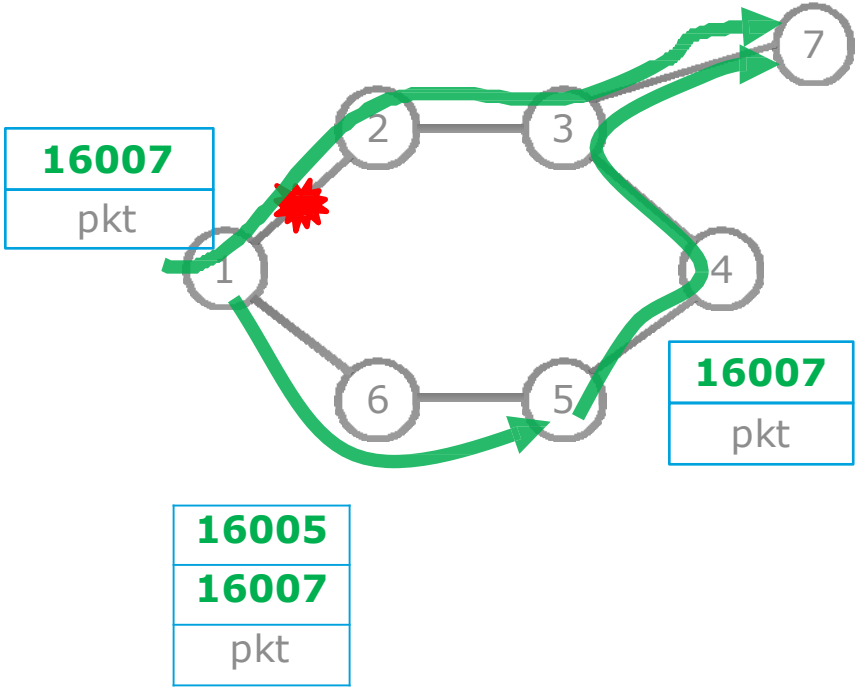

Automated Traffic Matrix Collection

- Traffic Matrix is fundamental for
  - capacity planning
  - centralized traffic engineering
  - IP/Optical optimization
- Most operators do not have an accurate traffic matrix
- With SR, the traffic matrix collection is automated

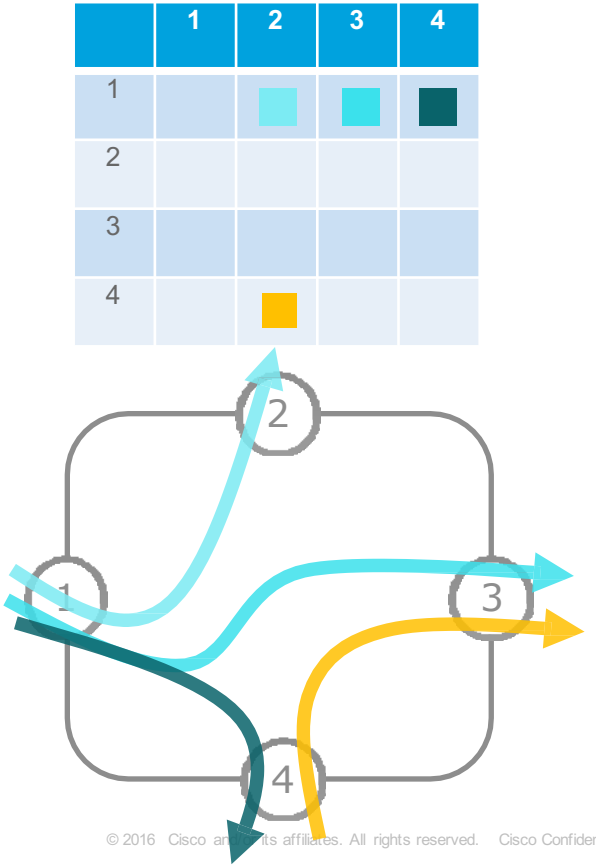

Optimized Content Delivery

- On a per-content, per-user basis, the content delivery application can engineer
  - the path within the AS
  - the selected border router
  - the selected peer
- Also applicable for engineering egress traffic from DC to peer
  - BGP Prefix and Peering Segments

16003
16002
126
pkt

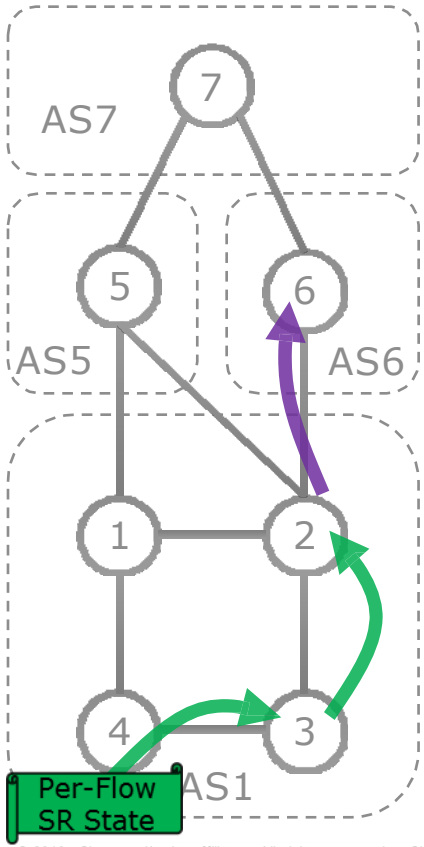

Application Engineered Routing

- Per-application flow engineering
- End-to-End
  - DC, WAN, AGG, PEER
- Millions of flows
  - No signaling
  - No midpoint state
  - No reclassification at boundaries

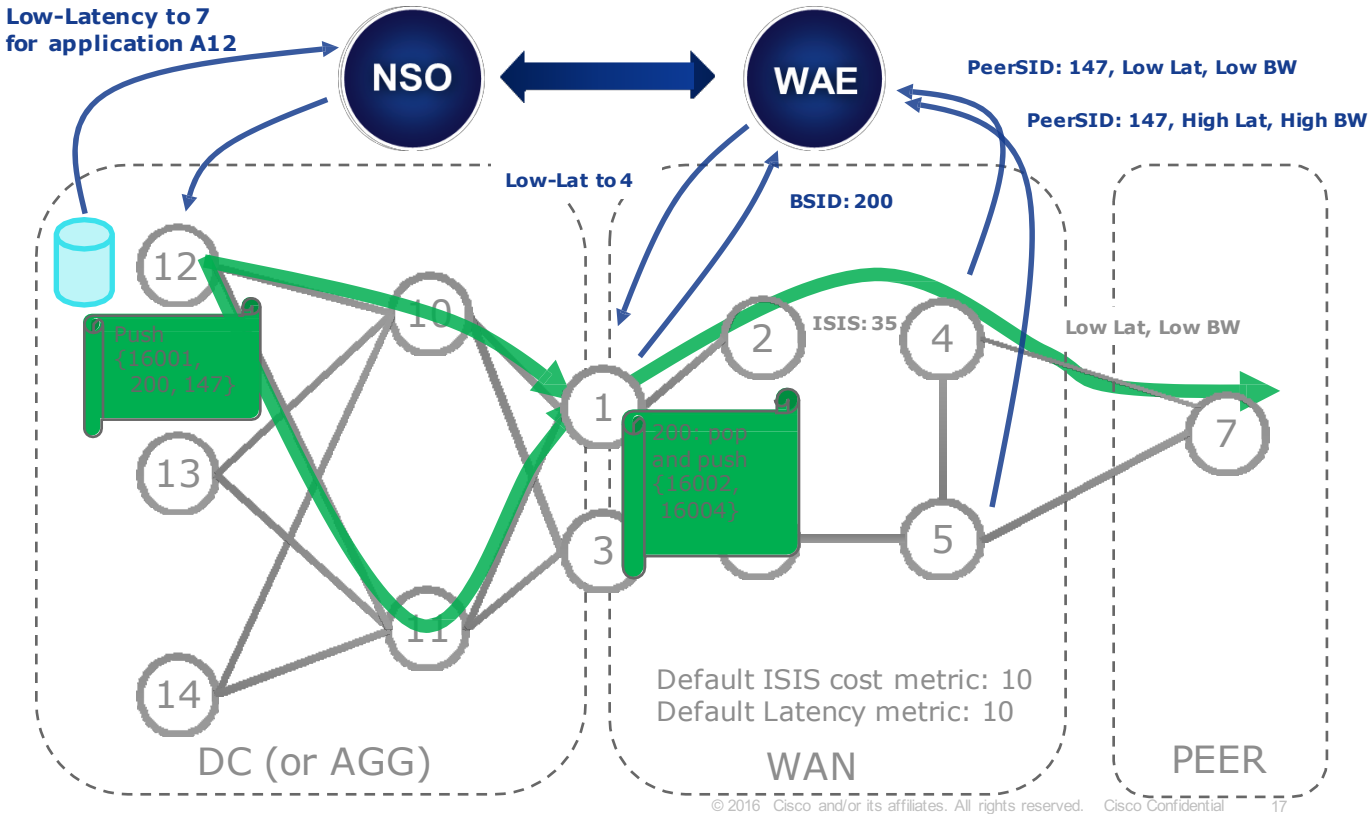

Application Engineered Routing

- Per-application flow engineering
- End-to-End
  - DC, WAN, AGG, PEER
- Millions of flows
  - No signaling
  - No midpoint state
  - No reclassification at boundaries

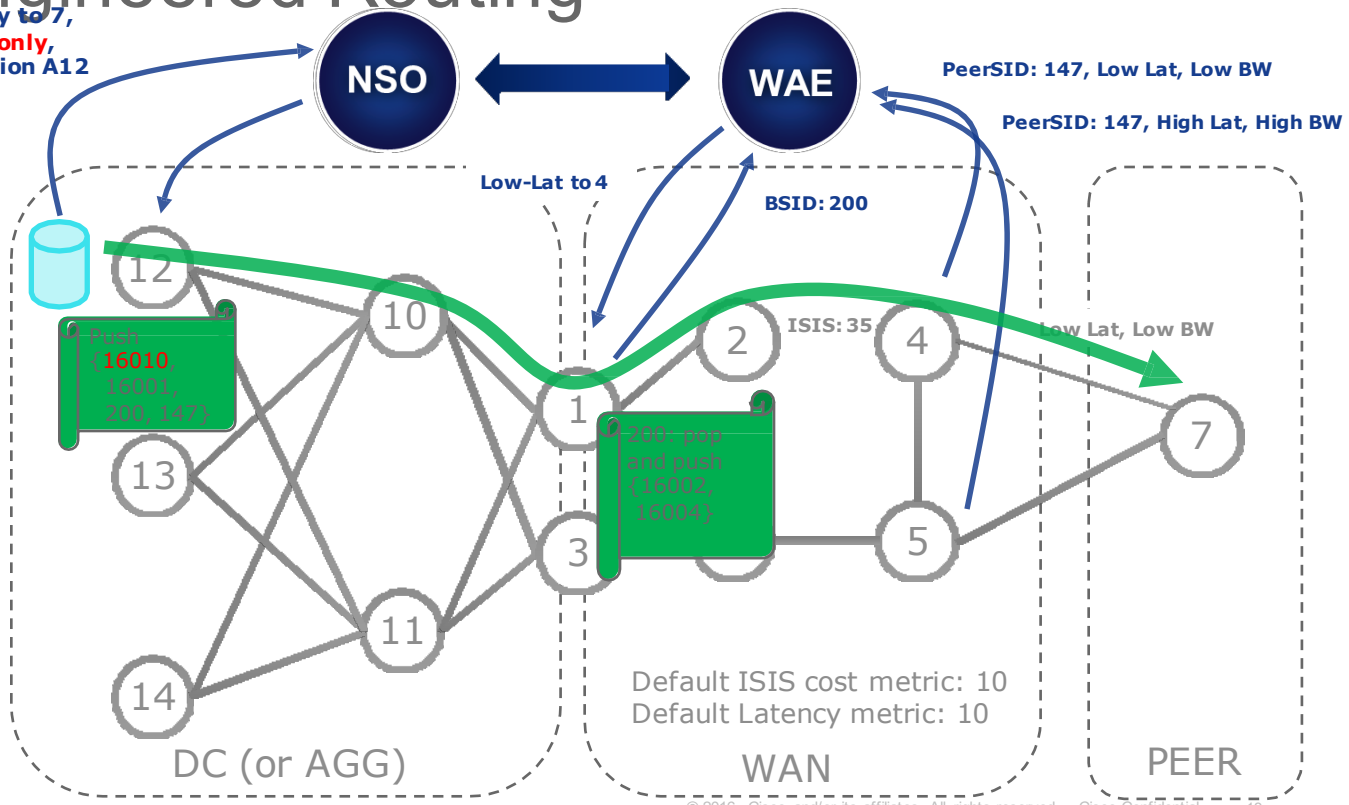

A Closer look to Control and Data Plane

MPLS Control and Forwarding Operation with Segment Routing

Services

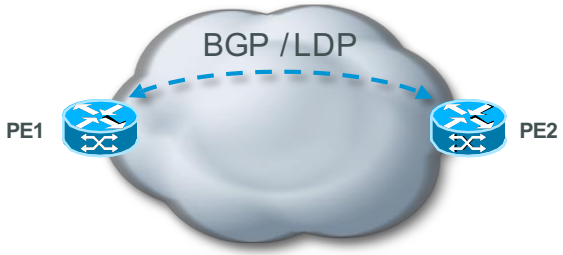

No changes to control or forwarding plane

Packet Transport

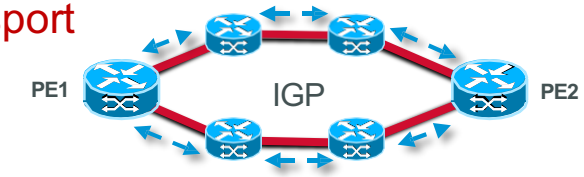

IGP label distribution for IPv4 and IPv6, same forwarding plane

SID Encoding

- Prefix SID
  - SID encoded as an index
  - Index represents an offset from SRGB base
  - Index globally unique
  - SRGB may vary across LSRs
  - SRGB (base and range) advertised with router capabilities
- Adjacency SID
  - SID encoded as absolute (i.e. not indexed) value
  - Locally significant
  - Automatically allocated for each adjacency

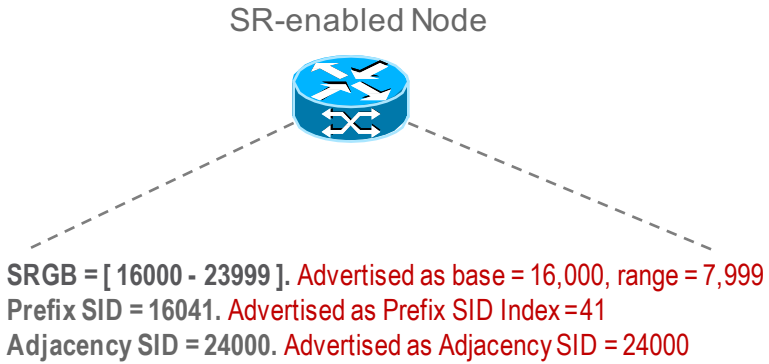

MPLS Data Plane Operation (Prefix SID)

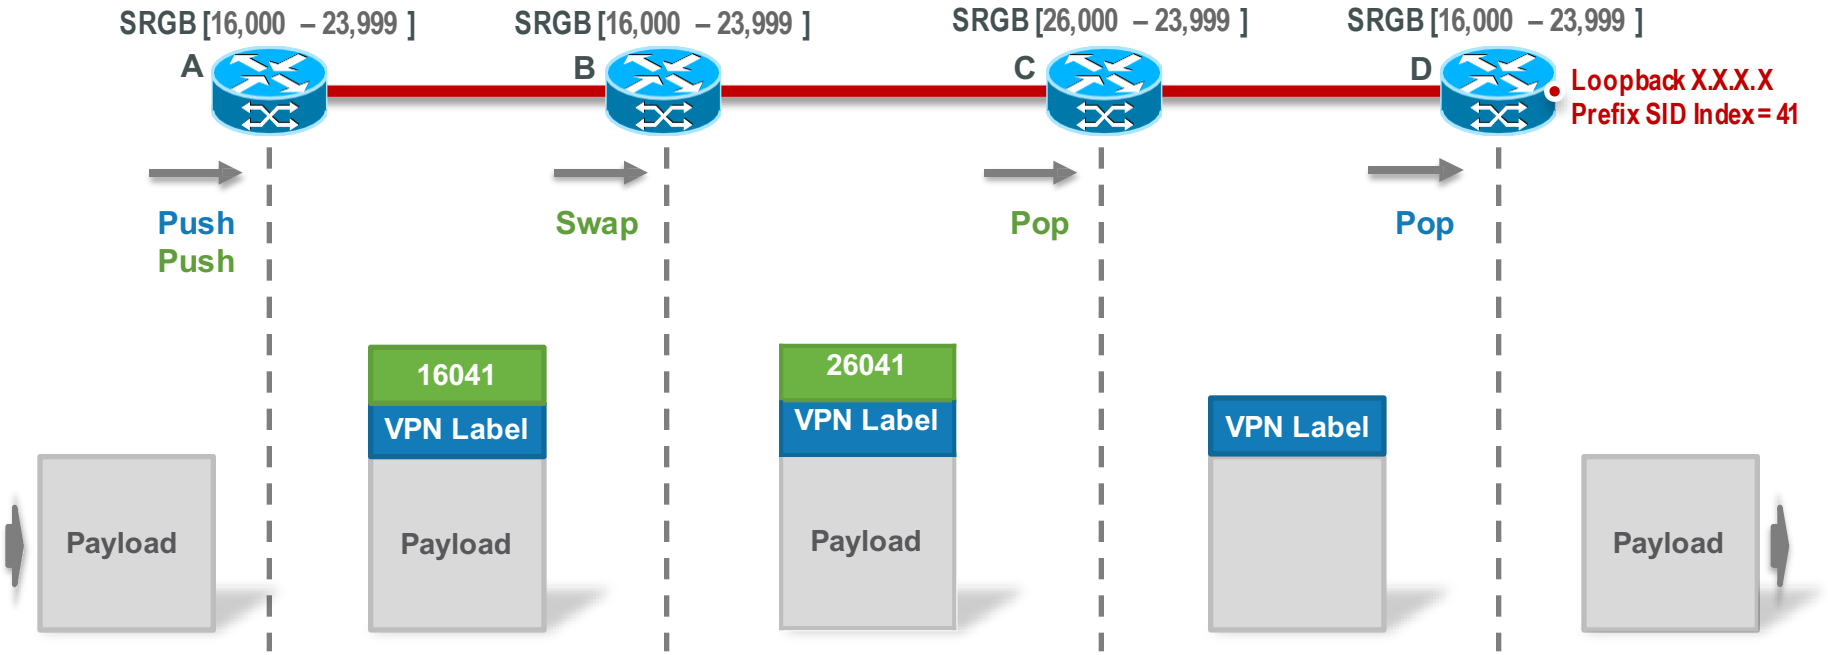

MPLS Data Plane Operation (Adjacency SIDs)

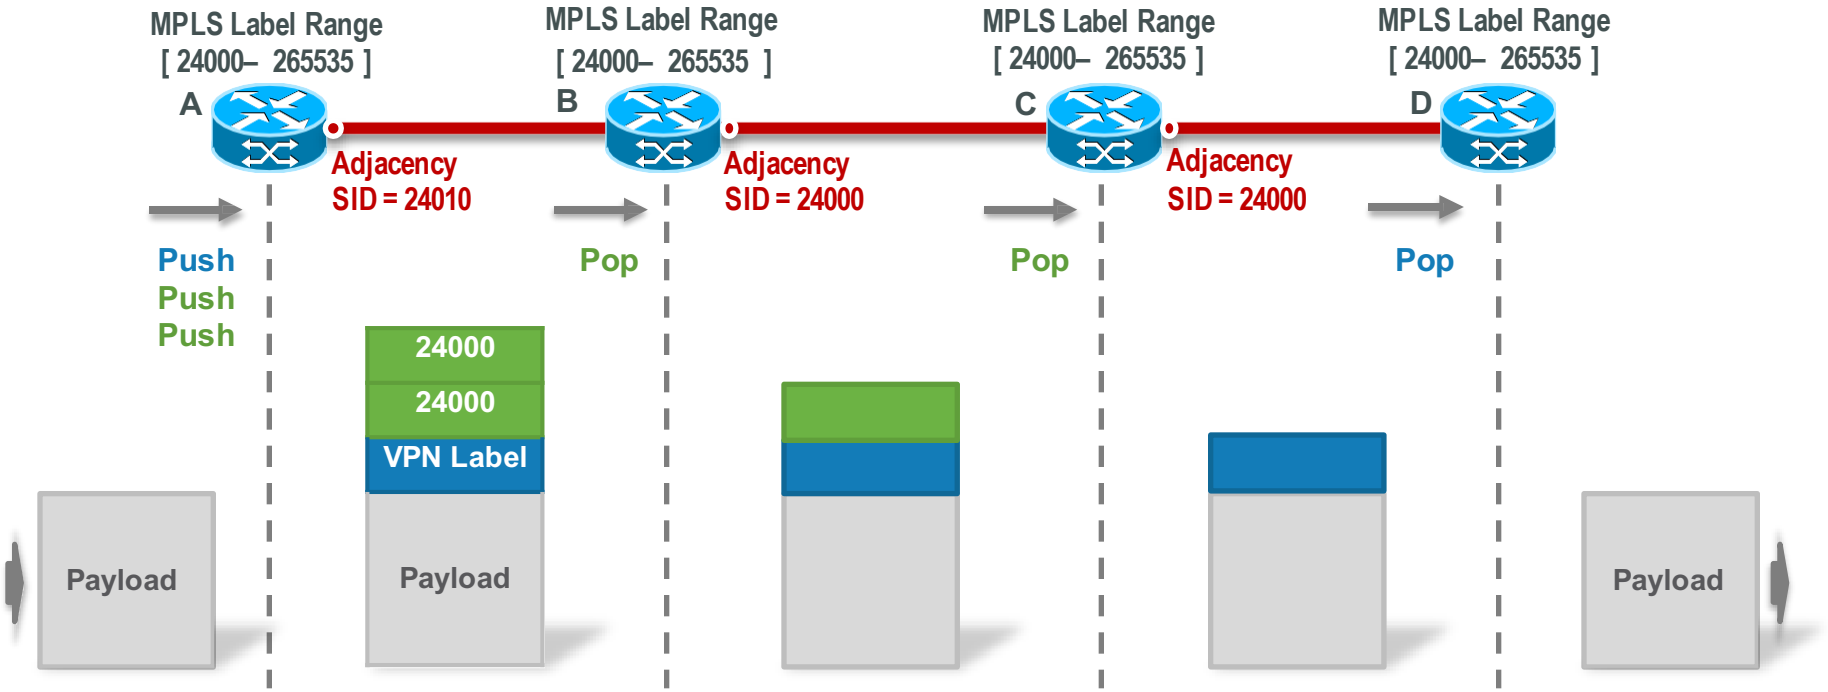

MPLS LFIB with Segment Routing

- LFIB populated by IGP (ISIS / OSPF)
- Forwarding table remains constant (Nodes + Adjacencies) regardless of number of paths
- Other protocols (LDP, RSVP, BGP) can still program LFIB

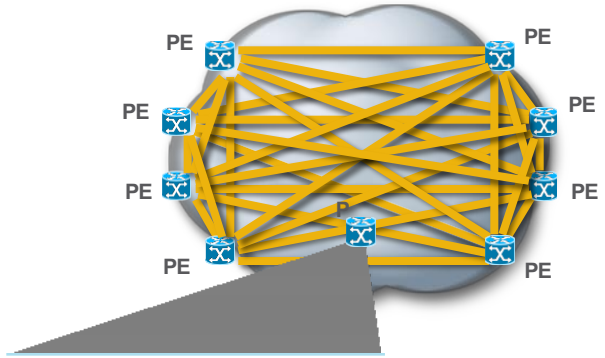

	In Label	Out Label	Out Interface
Network Node Segment Ids	L1	L1	Intf1
	L2	L2	Intf1
	…	…	…
	L8	L8	Intf4
	L9	L9	Intf2
Node Adjacency Segment Ids	L10	Pop	Intf2
	…	…	…
	Ln	Pop	Intf5

Forwarding table remains constant

Traffic Protection

Topology Independent LFA (TI-LFA) – Benefits

- **100%-coverage** 50-msec link and node protection
- **Simple** to operate and understand
  - automatically computed by the IGP
- **Prevents** transient **congestion** and suboptimal routing
  - leverages the post-convergence path, planned to carry the traffic
- Incremental deployment
  - also **protects LDP traffic**

Topology Independent LFA – Implementation

- Leverages existing and proven LFA technology
  - P space: set of nodes reachable from node S (PLR) without using protected link L
  - Q space: set of nodes that can reach destination D without using protected link L
- Enforcing loop-freeness on post-convergence path
  - Where can I release the packet?
    At the intersection between the post-convergence shortest path and the Q space
  - How do I reach the release point?
    By chaining intermediate segments that are assessed to be loop-free

Conclusion

Conclusion

- Simple routing extensions to implement source routing
- Packet path determined by prepended segment identifiers (one or more)
- Data plane agnostic (MPLS, IPv6)
- Network scalability and agility by reducing network state and simplifying control plane
- Traffic protection with 100% coverage with more optimal routing


CISCO