Cisco live!

January 28 - February 1, 2019 - Barcelona

cisco

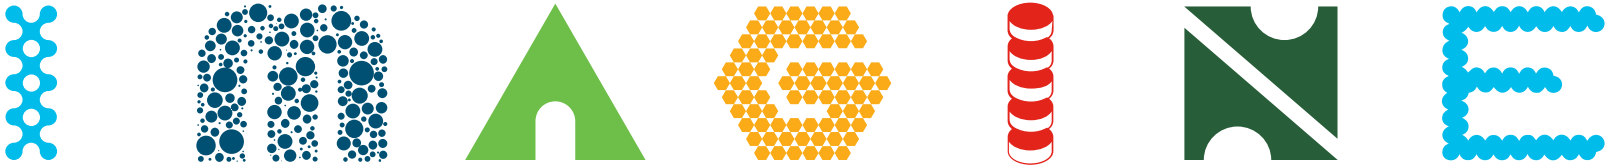

INTUITIVE

LTRMPL-2201

SP SDN – Segment Routing In Action

Jose Liste / Derek Tay / Josh Peters

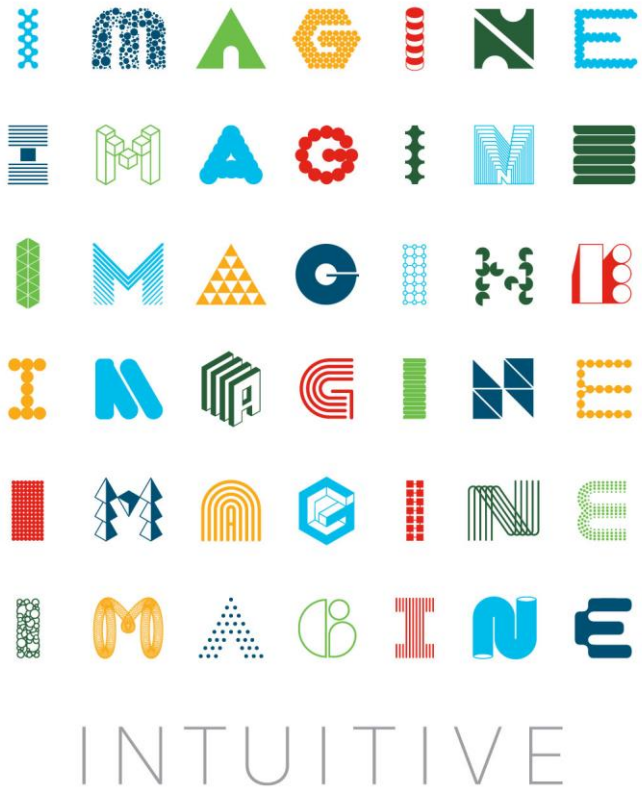

Cisco live!

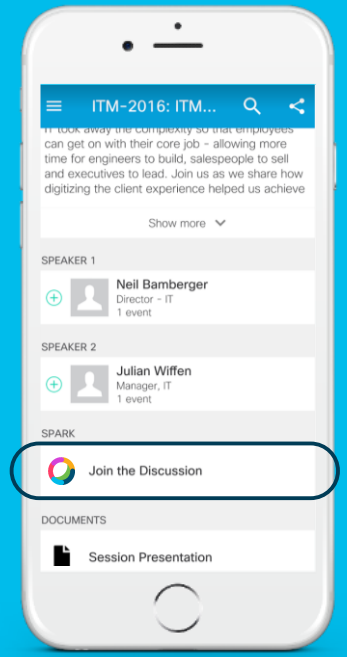

cs.co/ciscolivebot#LTRMPL-2201

Cisco Webex Teams

Questions?

Use Cisco Webex Teams (formerly Cisco Spark) to chat with the speaker after the session

How

1. Find this session in the Cisco Events Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

Lab Overview

Lab Workflow

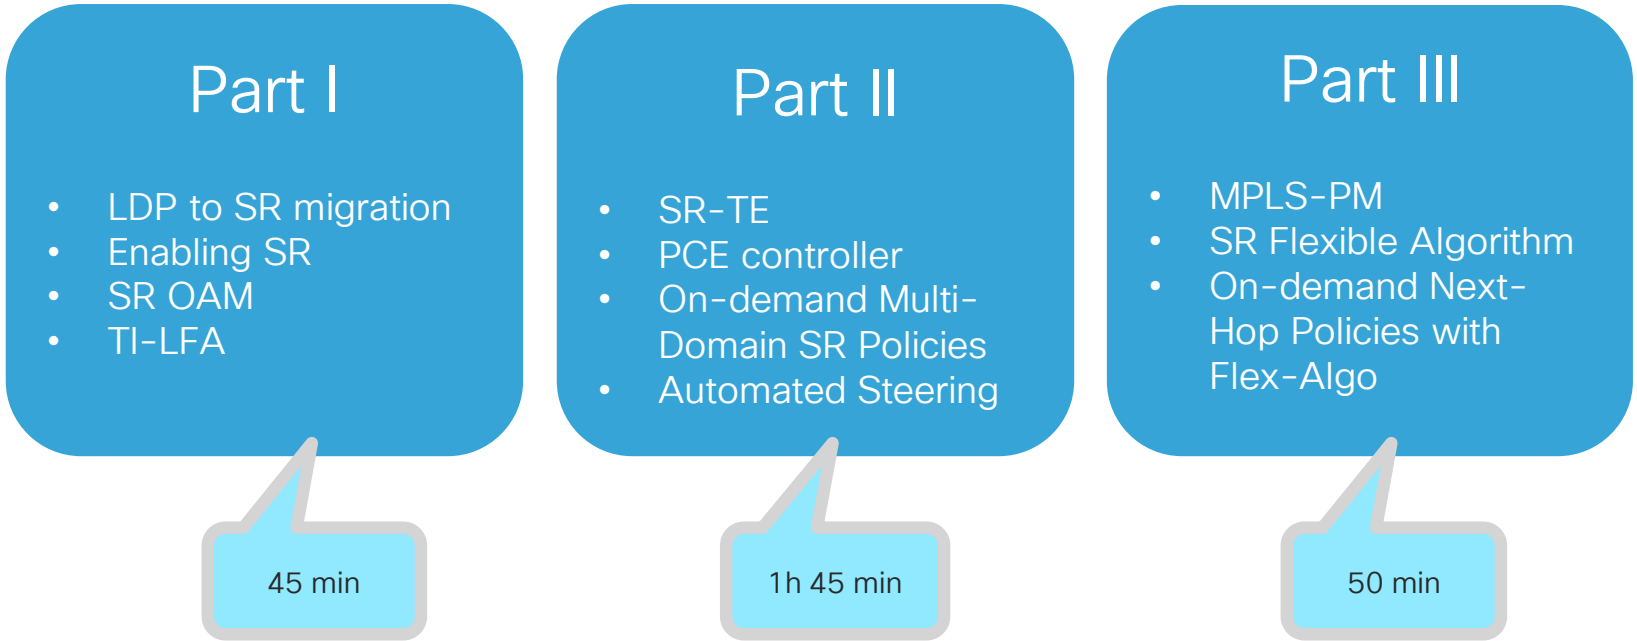

Topology

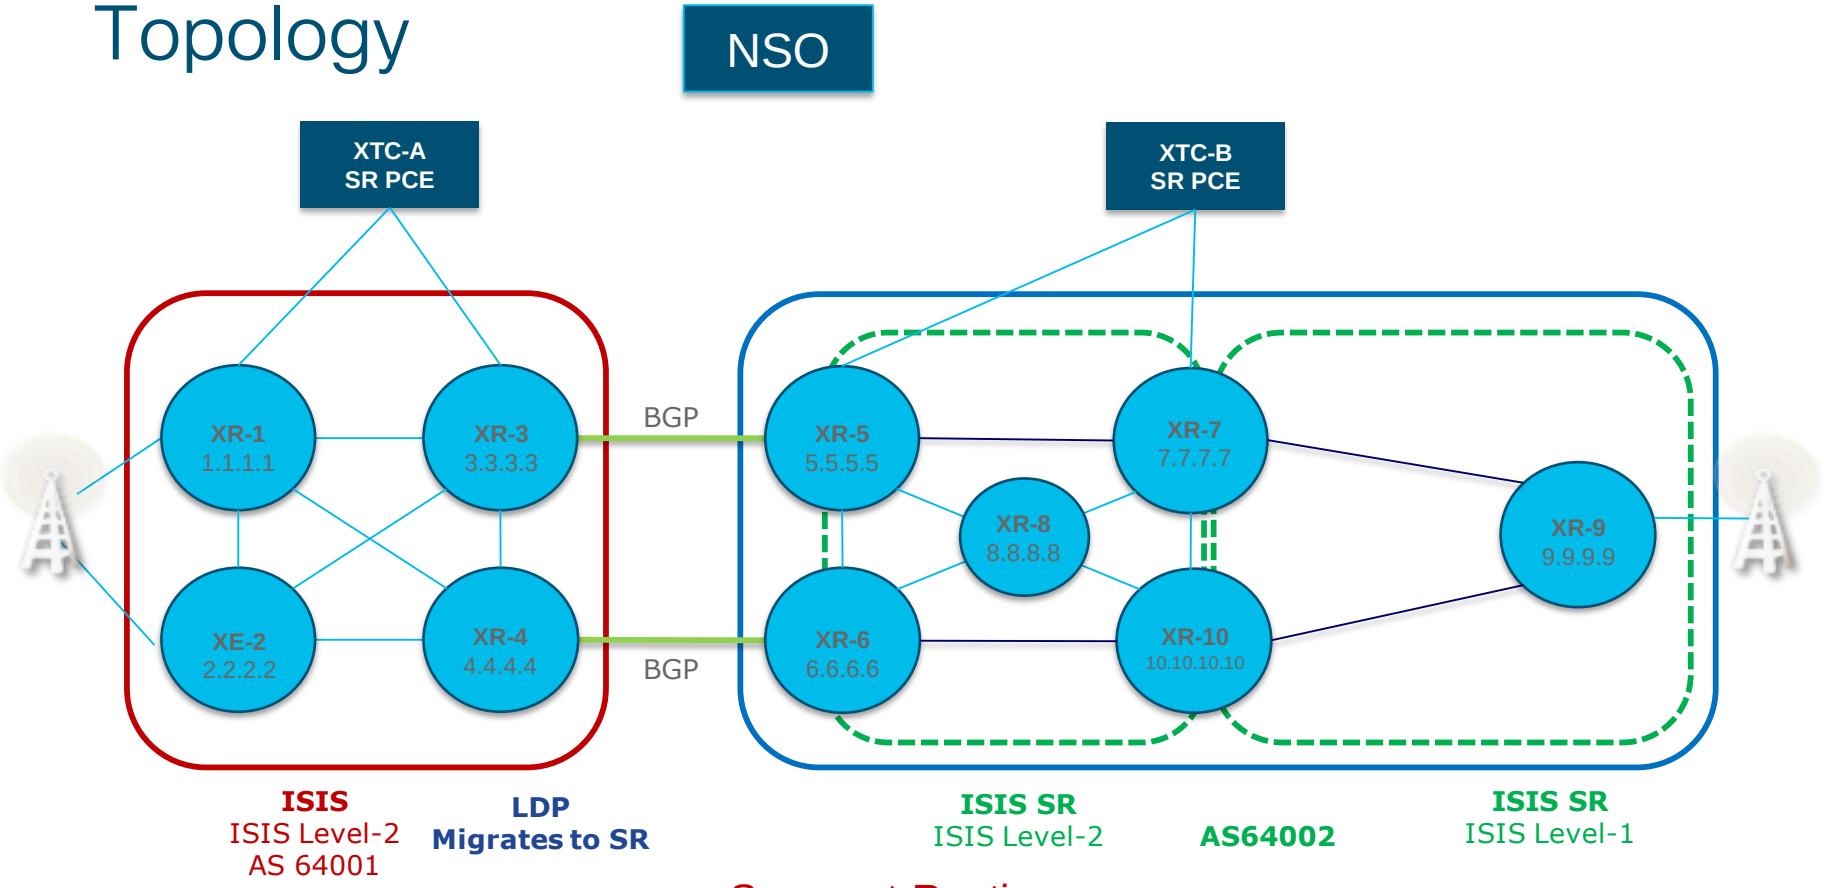

Lab Part I

Objectives – Part I

- In section you will learn:
  - Segment Routing (SR) configuration in IOS-XR and IOS-XE
  - LDP to SR migration steps
  - SR verification and monitoring
  - Topology-Independent Loop Free Alternate (TI-LFA) configuration and verification

Simplest migration LDP to SR

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

- Initial state: All nodes run LDP, not SR

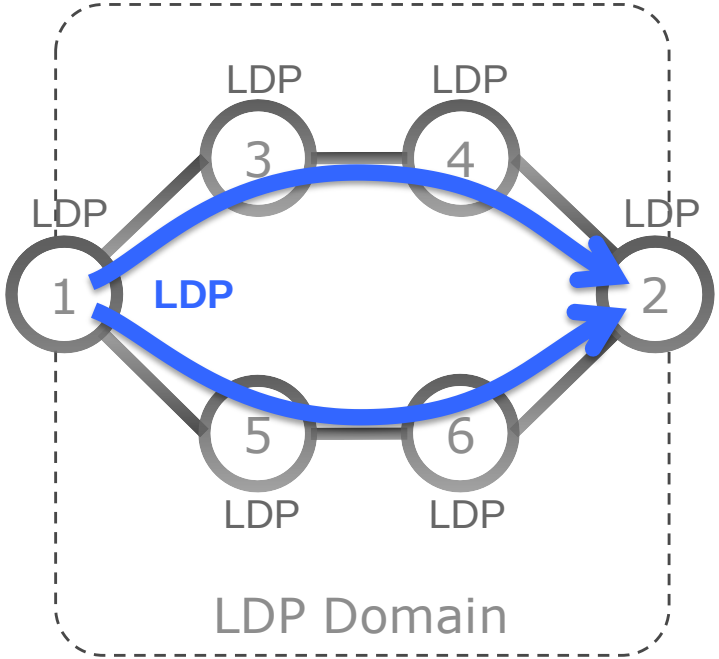

Simplest migration LDP to SR

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

- Initial state: All nodes run LDP, not SR
- Step1: All nodes are upgraded to SR
  - In no particular order
  - leave default LDP label imposition preference

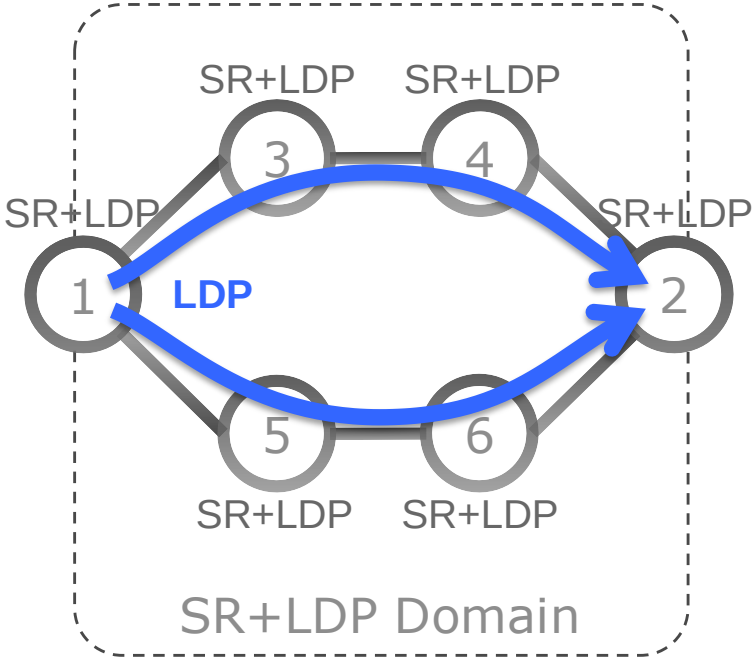

Simplest migration LDP to SR

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

- Initial state: All nodes run LDP, not SR
- Step1: All nodes are upgraded to SR
  - In no particular order
  - leave default LDP label imposition preference
- Step2: All PEs are configured to prefer SR label imposition
  - In no particular order

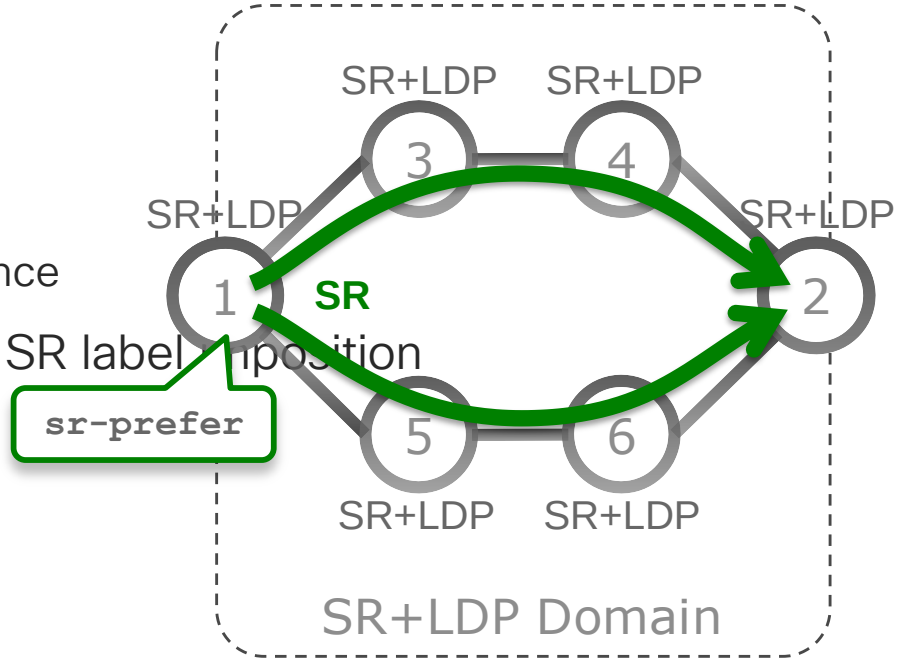

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

Simplest migration LDP to SR

- Initial state: All nodes run LDP, not SR
- Step1: All nodes are upgraded to SR
  - In no particular order
  - leave default LDP label imposition preference
- Step2: All PEs are configured to prefer SR label imposition
  - In no particular order
- Step3: LDP is removed from the nodes in the network
  - In no particular order

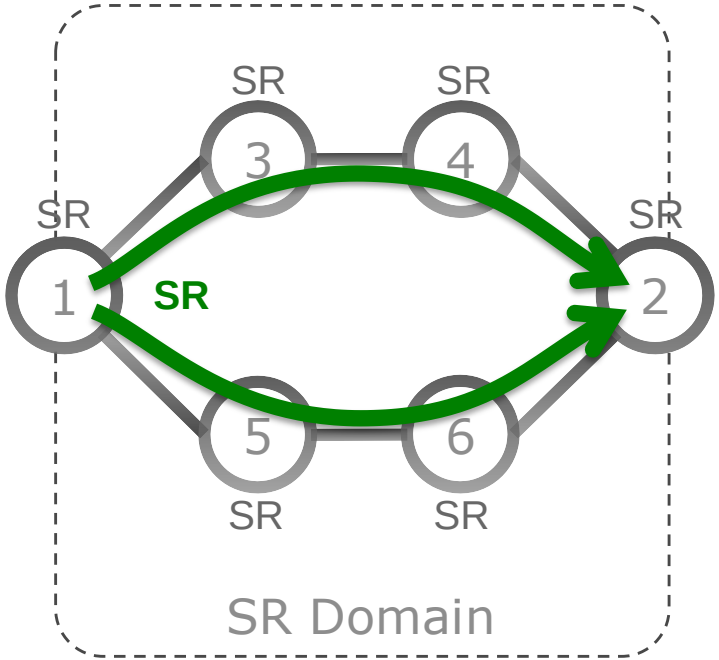

- Final state: All nodes run SR, not LDP

Exercises – Part I

Time Management: 45 min

- Exercise 1 – LDP to SR Migration
- Enable Segment Routing
- SR Verification
- Enabling forwarding to prefer SR LSP
- Enabling TI-LFA

SR Topology

Node X
Lo0 – 1.1.1.x/32
Link XY – 99.X.Y.0/24; X < Y
Prefix SID – 16000 + X

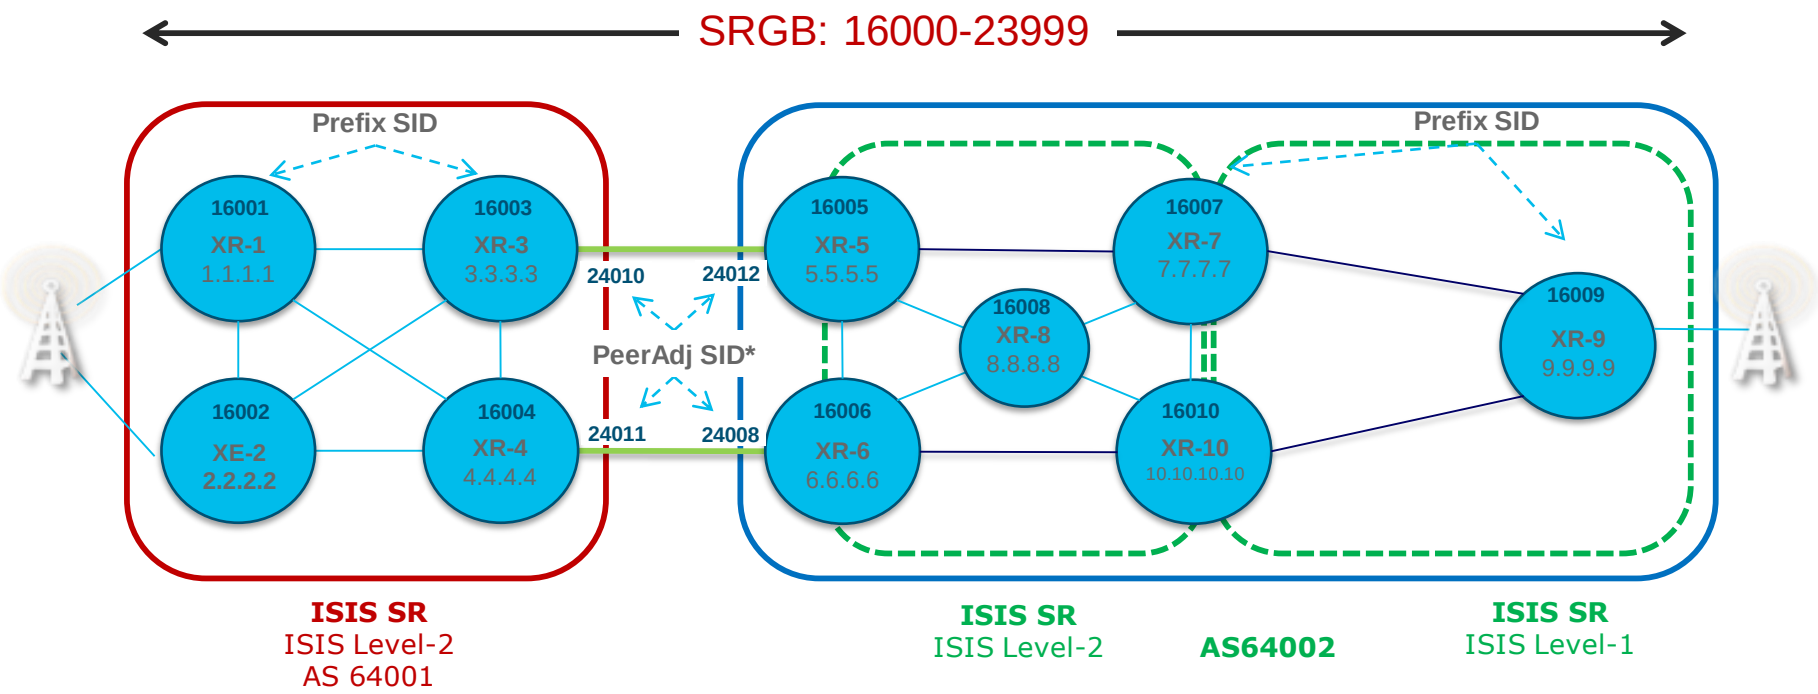

Note (*) = PeerAdj SID values are dynamically allocated

Lab Part II

Objectives – Part II

- In this section you will learn:
  - SR On-Demand Next-Hop (ODN) for on-demand instantiation of SR policies in IOS-XR and IOS-XE
  - Automated traffic steering onto SR policies without performance degradation
  - IOS XR's SR-PCE acting as stateful PCE for multi-domain SR-TE policies
  - NSO for faster and more reliable service orchestration

SR-PCE

- SR-PCE is an IOS XR multi-domain stateful SR Path Computation Element (PCE)
  - IOS XR: XTC functionality is available on any physical or virtual IOS XR node, activated with a single configuration command
  - Multi-domain: Real-time reactive feed via BGP-LS/ISIS/OSPF from multiple domains; computes inter-area/domain/AS paths
  - Stateful: takes control of SRTE Policies, updates them when required
  - SR PCE: native SR-optimized computation algorithms
- SR-PCE is fundamentally distributed
  - Not a single all-overseeing entity (“god box”), but distributed across the network; RR-alike deployment

SR-PCE Building Blocks

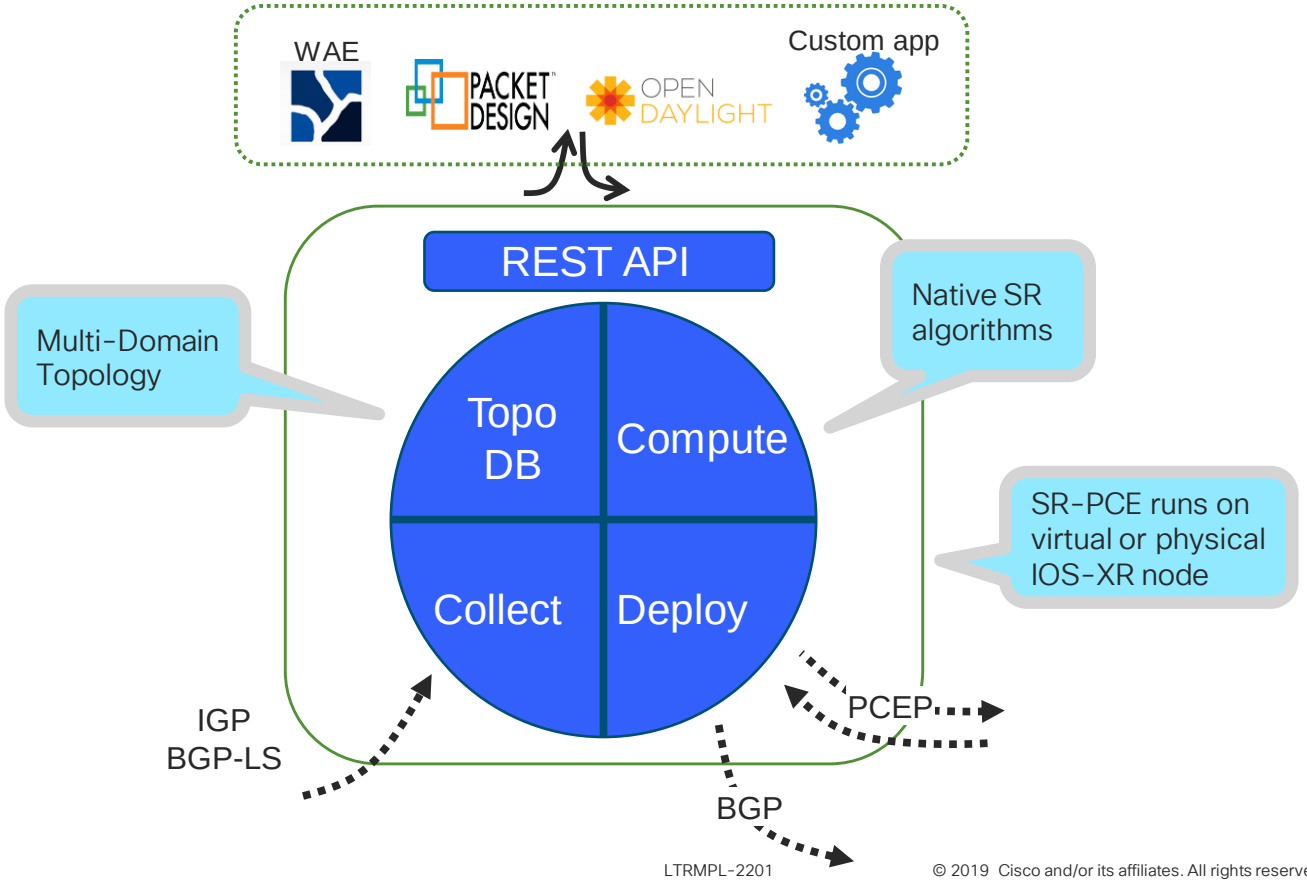

Exercises – Part II

Time Management: 1h 45 min

- Exercise 2 – BGP Peer-Adj SID on Inter-AS Links
- Exercise 3 – Enabling BGP-LS on SR-PCE
- Exercise 4 – Enabling PCE on SR-PCE
- Exercise 5 – SR On-Demand Next-Hop – day 0 provisioning
- Exercise 6 – Service provisioning (NSO)
- Exercise 7 – ODN operation & service verification
- Exercise 8 – Service decommissioning (NSO)

Topology Discovery

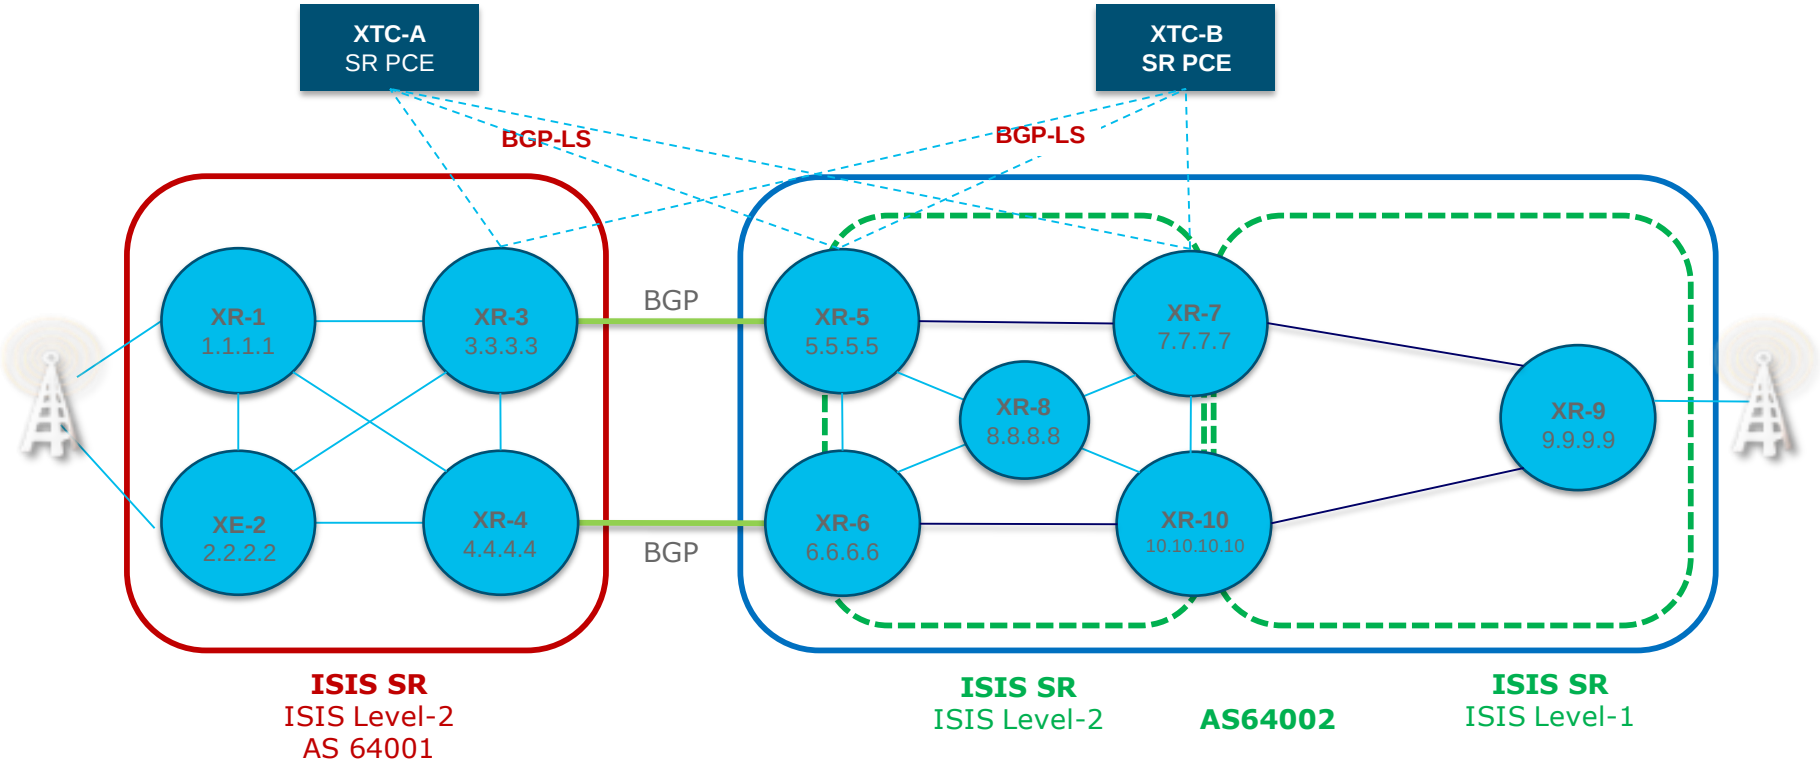

PCEP

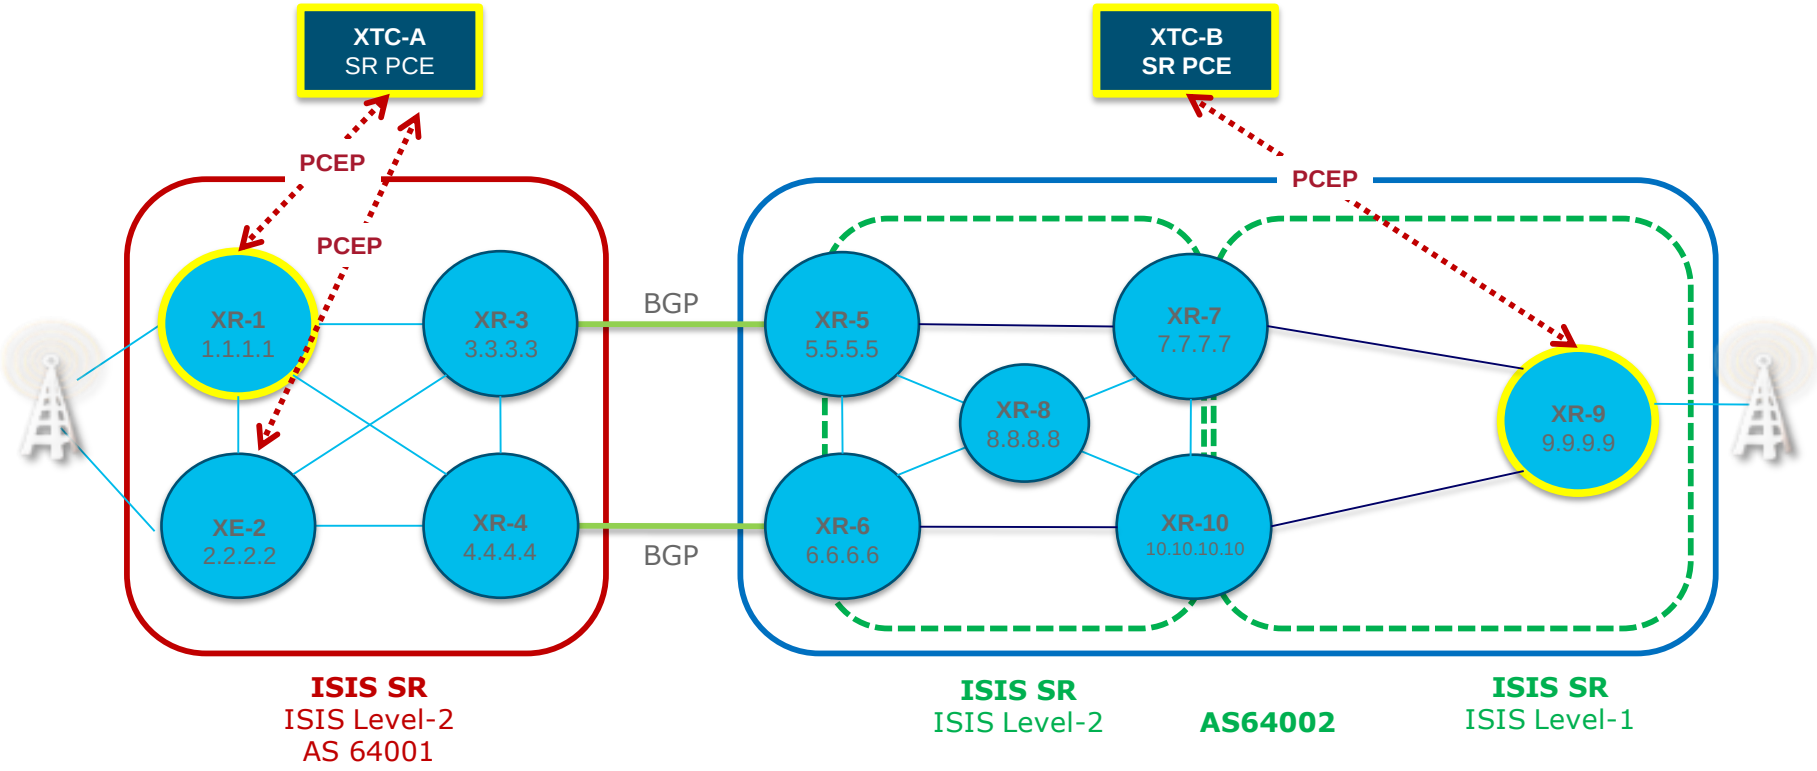

ODN Workflow

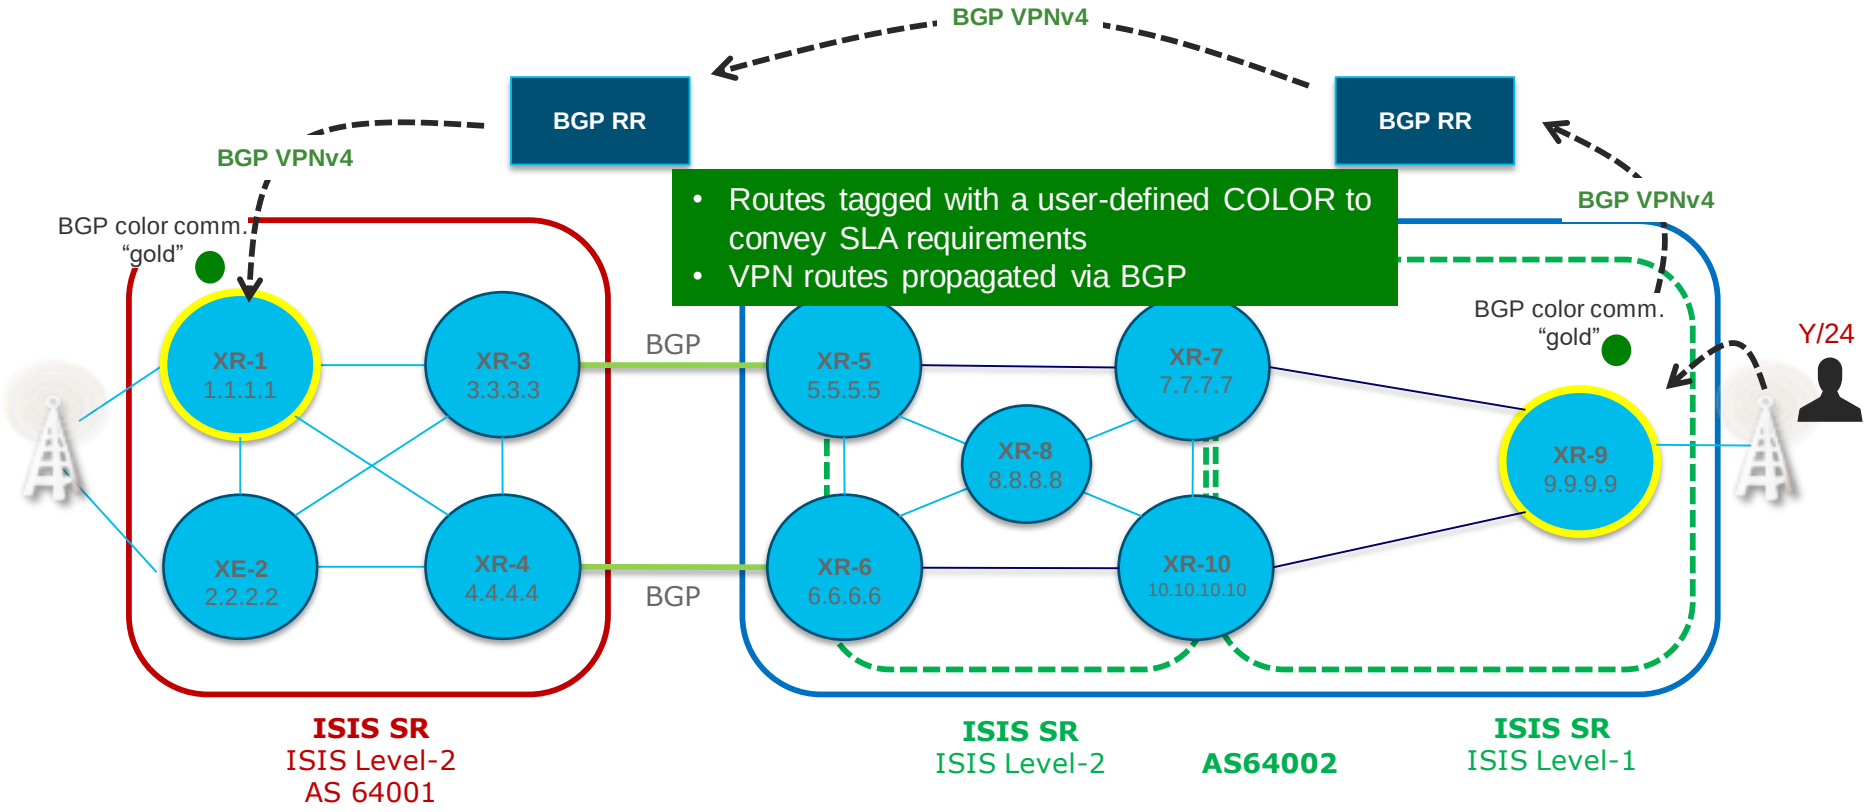

ODN Workflow

- Ingress PE matches on user-specified BGP “color” community
- Ingress PE enforces a “template” associated with the color community

SRTE
On-demand color “gold”
contact PCE
request path to BGP NH
minimize TE metric

XTC-A
SR PCE

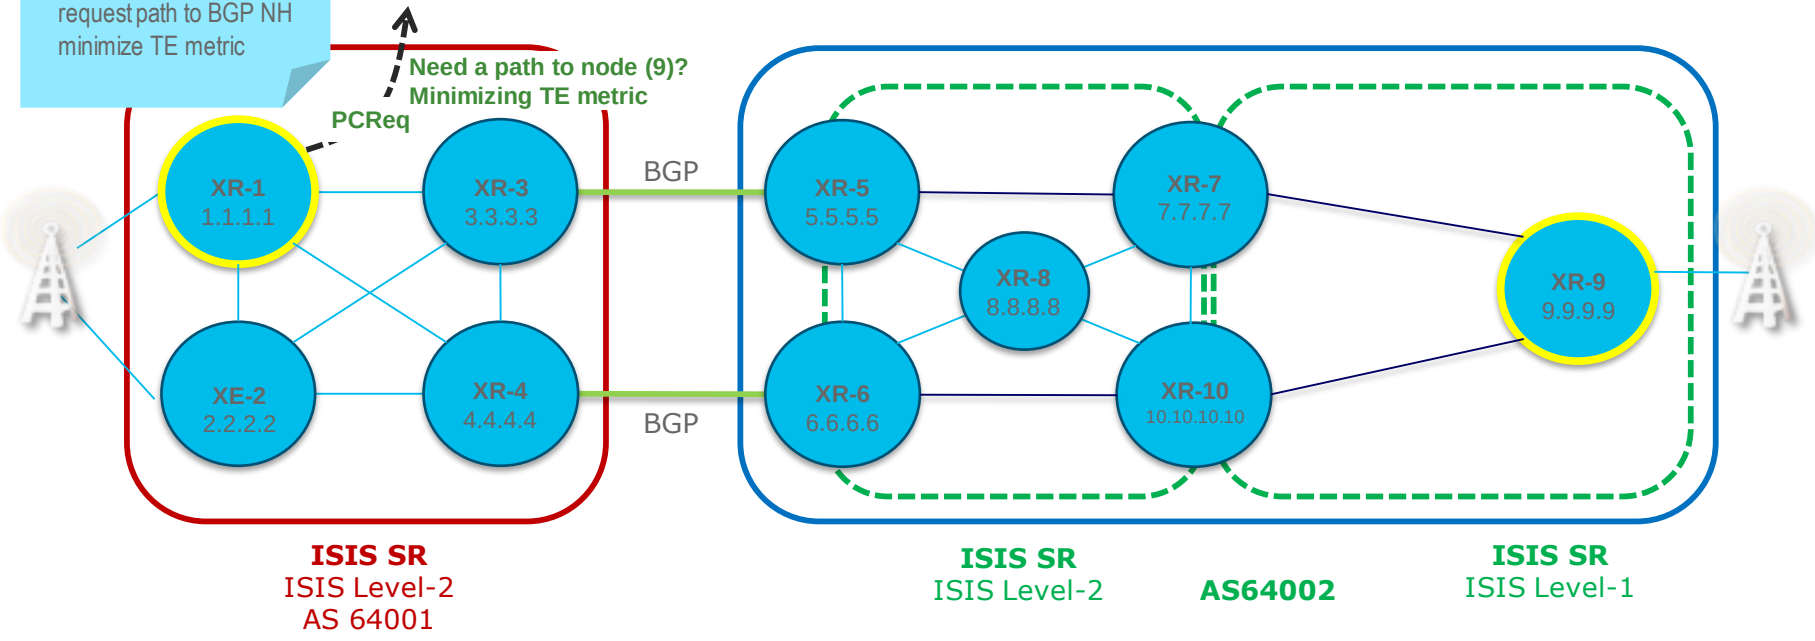

ODN Workflow

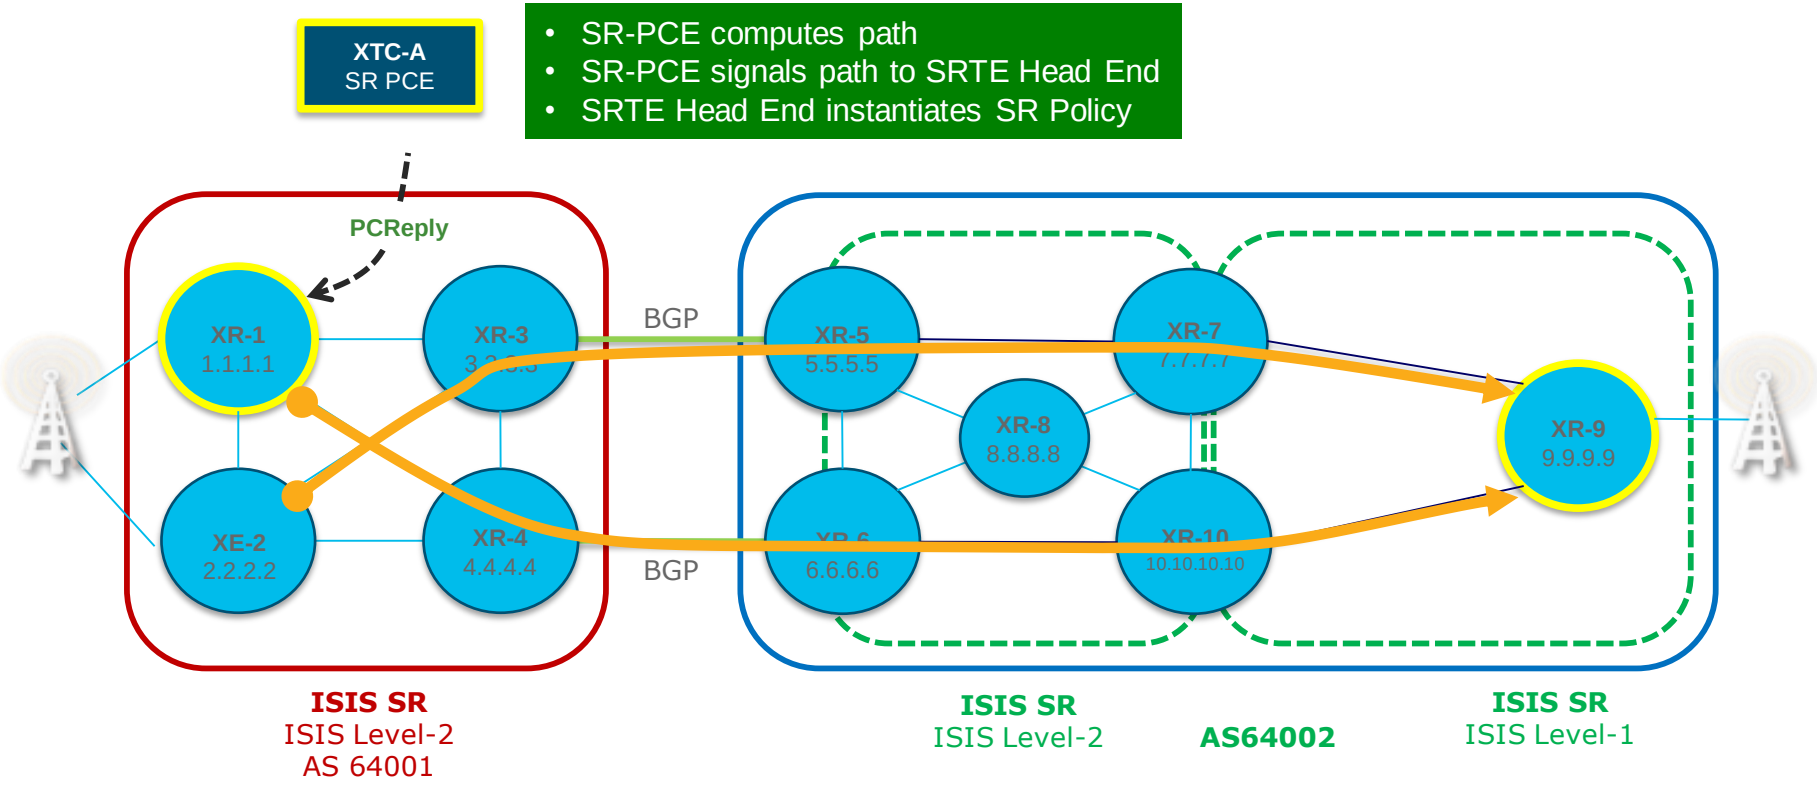

Lab Part III

Exercises – Part III

Time Management: 50 min

- Exercise 9 – Enable MPLS-PM
- Exercise 10 – Enable SR Flex-Algorithm
- Exercise 11 – SR ODN with Flex-Algo operation

MPLS PM Overview

Per-Link delay Measurement

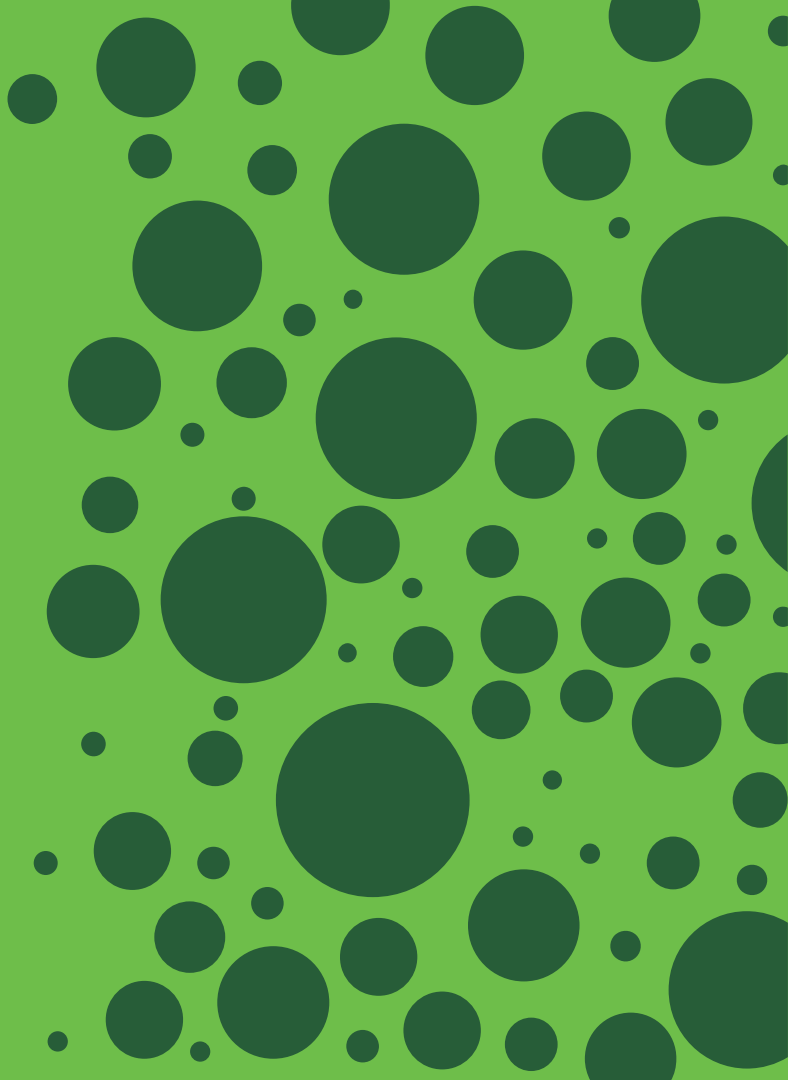

ISIS Signaling

```
Type   Description
-----------------------------------------------------
 33    Unidirectional Link Delay

 34    Min/Max Unidirectional Link Delay

 35    Unidirectional Delay Variation

?
                ISIS
```

- RFC 7810 (IS-IS Traffic Engineering (TE) Metric Extensions)
- Used to advertise extended TE metrics – e.g. link delay (in usec)

OSPF and BGP-LS

```
  Value  Sub-TLV
    27   Unidirectional Link Delay
    28   Min/Max Unidirectional Link Delay
    29   Unidirectional Delay Variation
?
               OSPF
```

- RFC 7471 (OSPF Traffic Engineering (TE) Metric Extensions)
- Used to advertise extended TE metrics – e.g. link delay (in usec)
- BGP-LS: draft-ietf-idr-te-pm-bgp

Leveraged by SRTE – SR Policy

- SR Policy for min delay

```
segment-routing
  traffic-eng
    policy FOO
      color 20 end-point ipv4 1.1.1.3
      binding-sid mpls 1000
      candidate-paths
        preference 100
          dynamic mpls
            metric
              type delay
```

Leveraged by SRTE – IGP Flex Algo

- IGP SR Flex Algo for minimum delay

```
router isis 1
  flex-algo 128
    metric-type delay
```

Per-link delay Measurement

- Over a measurement internal
  - minimum → Used as metric for SRTE (Policy or Flex-Algo)
  - average
  - maximum
  - variance

  → Not used by SRTE
- One-way or Two-way
  - one-way requires clock synchronization

Minimum delay is of interest for SRTE

- Minimum delay provides the propagation delay
  - fiber length / speed of light
- A property of the topology
  - with awareness of DWDM circuit change
- SRTE (Policy or Flex-Algo) can optimize on min delay

Link Delay - Configuration

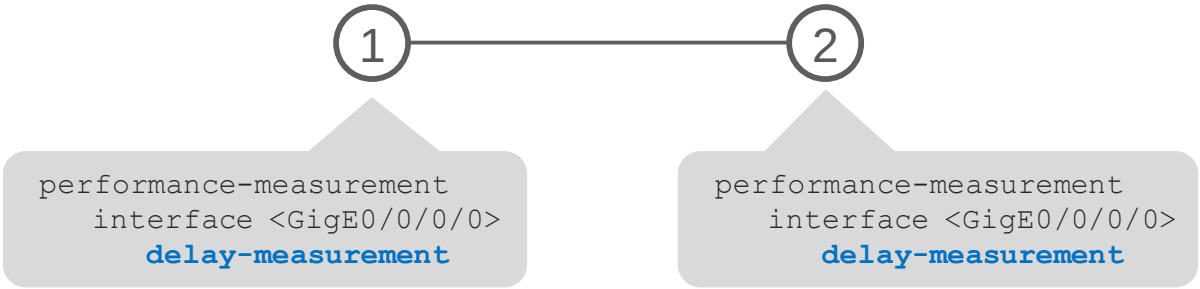

- If the link is enabled for an IGP, then this IGP automatically includes the delay TLV in its LSP/LSA

Link Delay – Probe Measurement

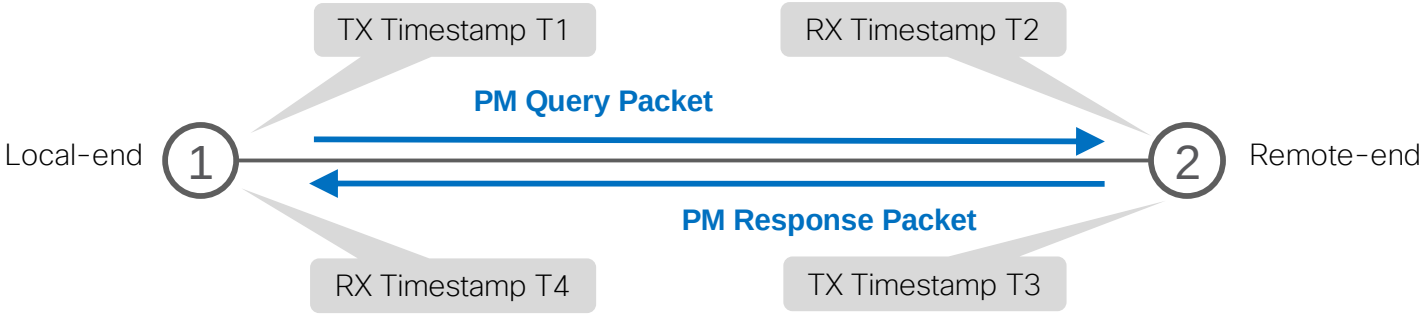

- One Way Delay = (T2 – T1)
- Two-Way Delay = (T2 – T1) + (T4 – T3)

Query Packet using RFC 6374 Packet Format

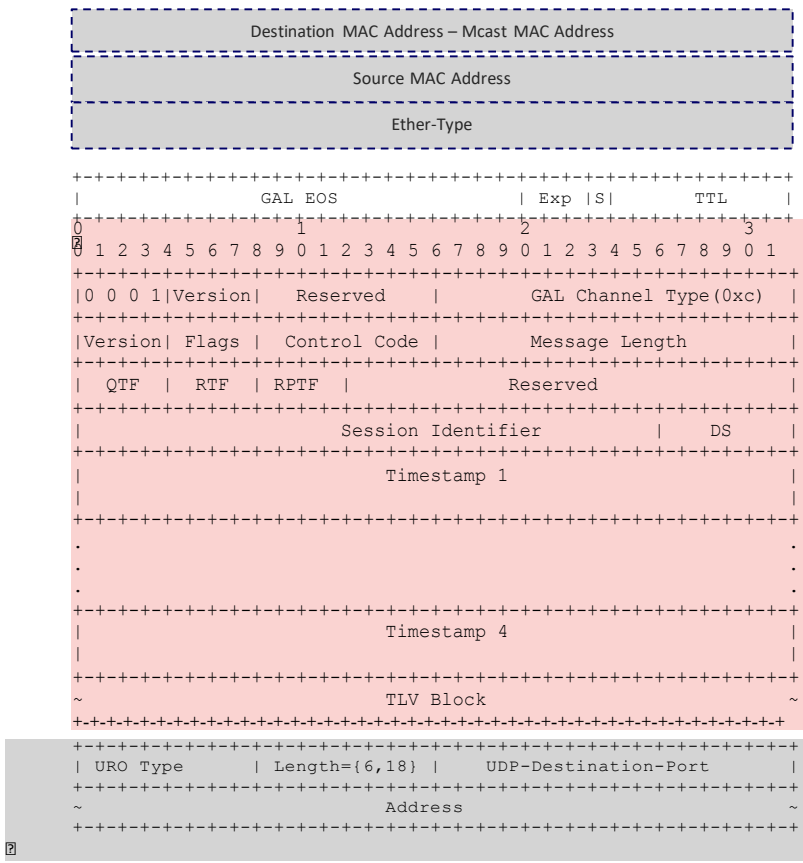

L2 Encap - RFC 7213

GAL/GAch

DM Packet - RFC 6374
Timestamp - 1588v1 (8 bytes)

IP/UDP Return TLV - RFC 7876
(for one-way mode)

Enable MPLS-PM

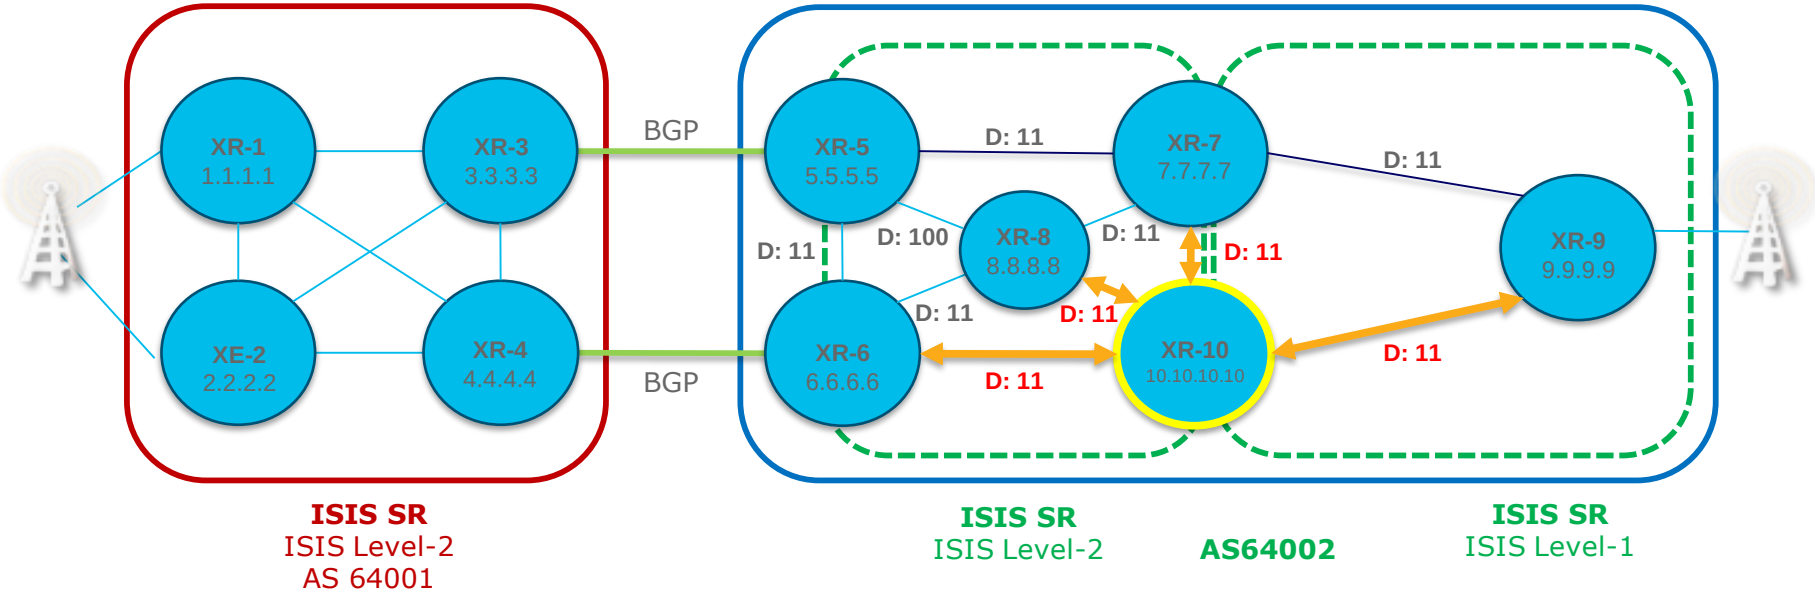

SR Flex Algo Overview

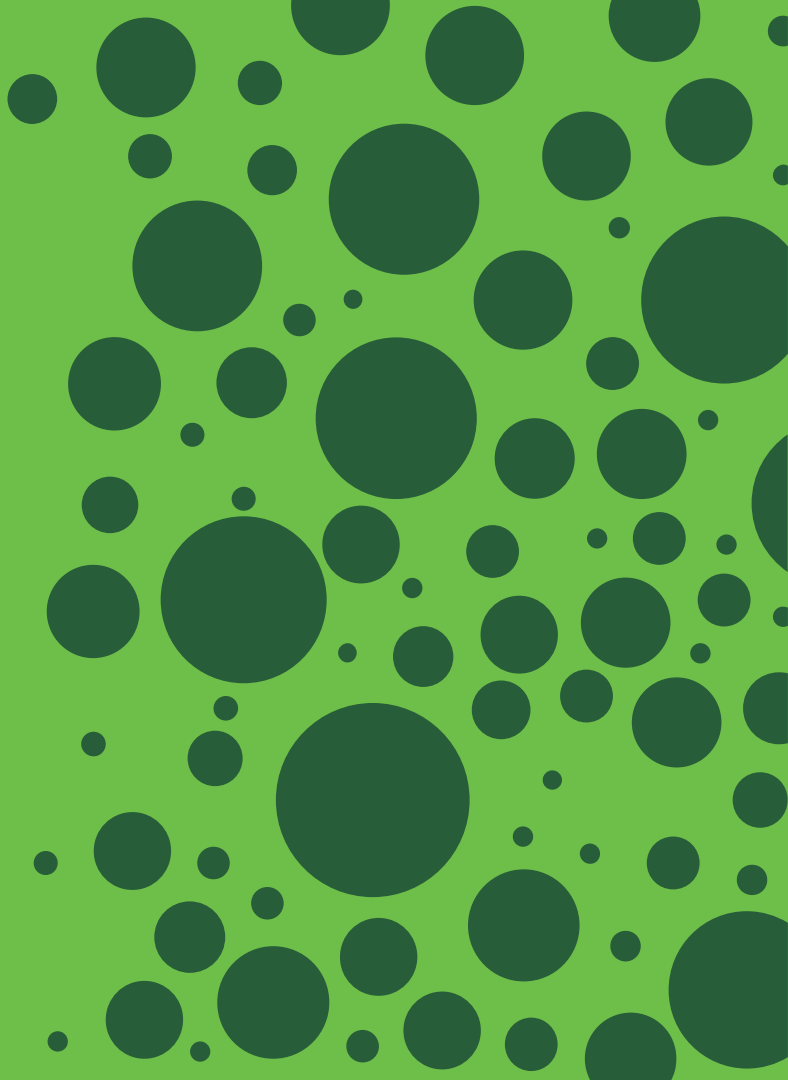

SR Flexible Algorithm

- We call “Flex-Algo”
  - The algorithm is defined by the operator, on a per-deployment basis
- Flex-Algo K is defined as
  - The minimization of a specified metric: IGP, delay, …
  - The exclusion of certain link properties: link-affinity, SRLG, …

SR IGP Flexible Algorithms

- Example
  - Operator1 defines Flex-Algo 128 as “minimize IGP metric and avoid link-affinity “brown”
  - Operator2 defines Flex-Algo 128 as “minimize delay metric”

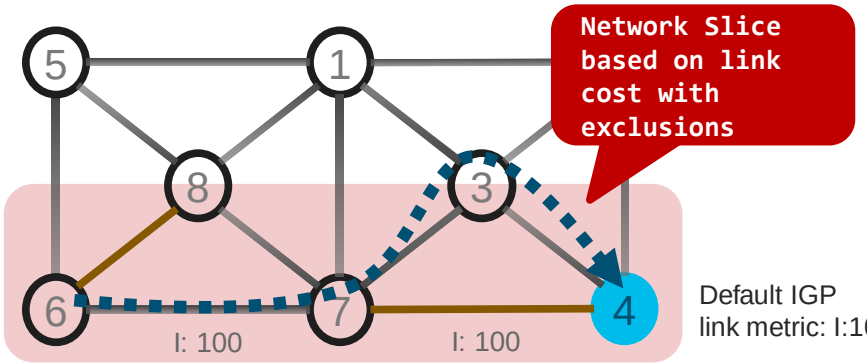

Default IGP link metric: I:10

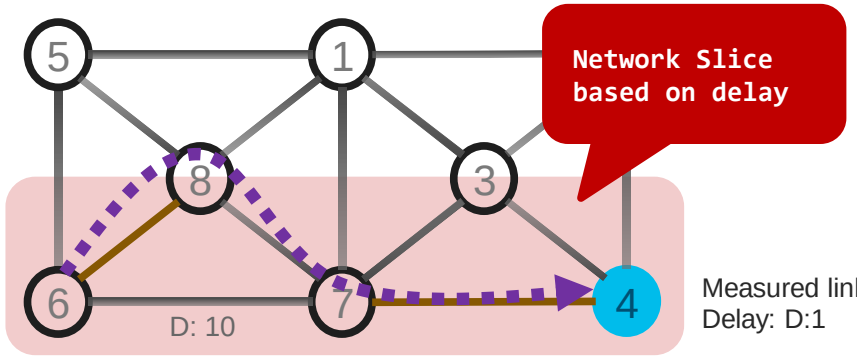

Measured link Delay: D:1

Flex-Algo Participation Advertisement

- Each node MUST advertise Flex-Algo(s) that it is participating in

Nodes 0 and 9 participate to Algo 0 and 128 and 129

Nodes 1/2/3/4 participate to Algo 0 and 128

Nodes 5/6/7/8 participate to Algo 0 and 129

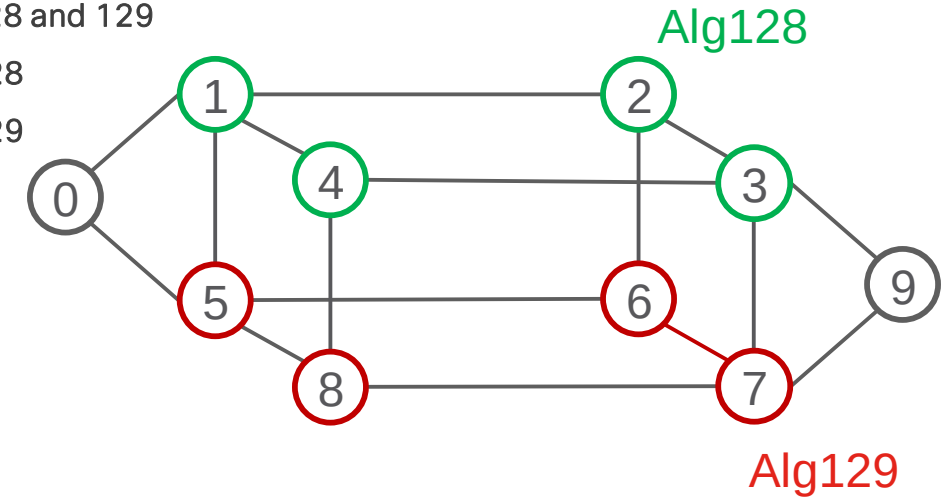

Prefix-SID for each Flex-Algo

- If a node advertises participation in a Flex-Algo likely it also advertises a prefix SID for that Flex-Algo

Node 9 advertises

Prefix SID 16009 for ALGO 0

Prefix SID 16809 for ALGO 128

Prefix SID 16909 for ALGO 129

Node 2 advertises

Prefix SID 16002 for ALGO 0

Prefix SID 16802 for ALGO 128

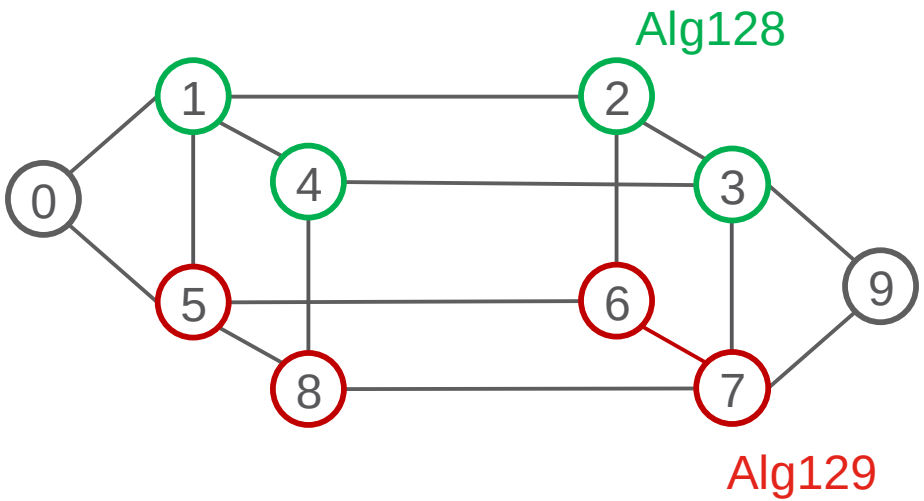

No additional loopback address

- Flex-Algo Prefix SID's can be advertised as additional prefix-SID's of the existing loopback address

Node 9 advertises loopback0 1.1.1.9/32 with

Prefix SID 16009 for ALGO 0

Prefix SID 16809 for ALGO 128

Prefix SID 16909 for ALGO 129

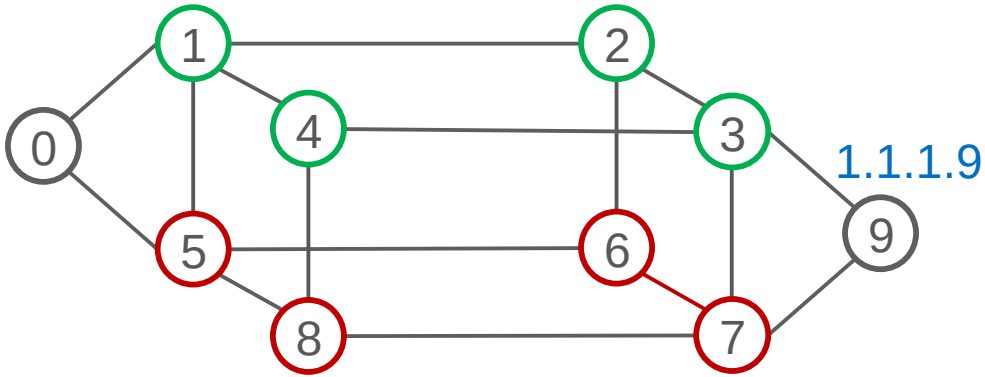

Multi-Plane Networks

- Powered by SR IGP Flex Algo

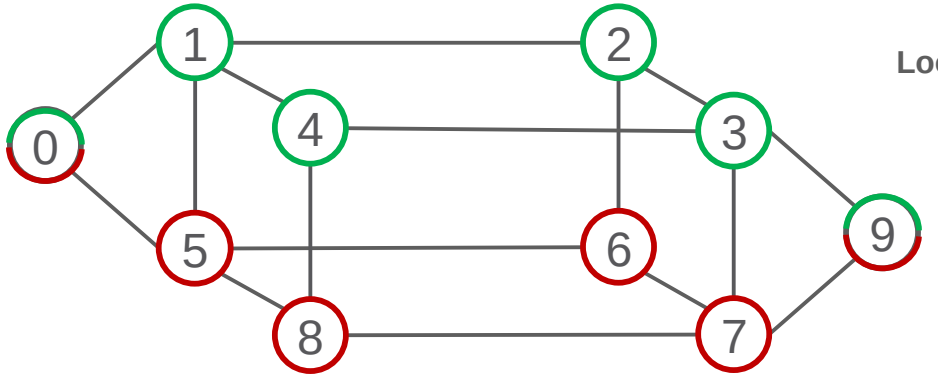

Loopback 1.1.1.9/32

16009 Algo 0

16809 Flex-Algo 128

16909 Flex-Algo 129

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

Multi-Plane Networks (cont.)

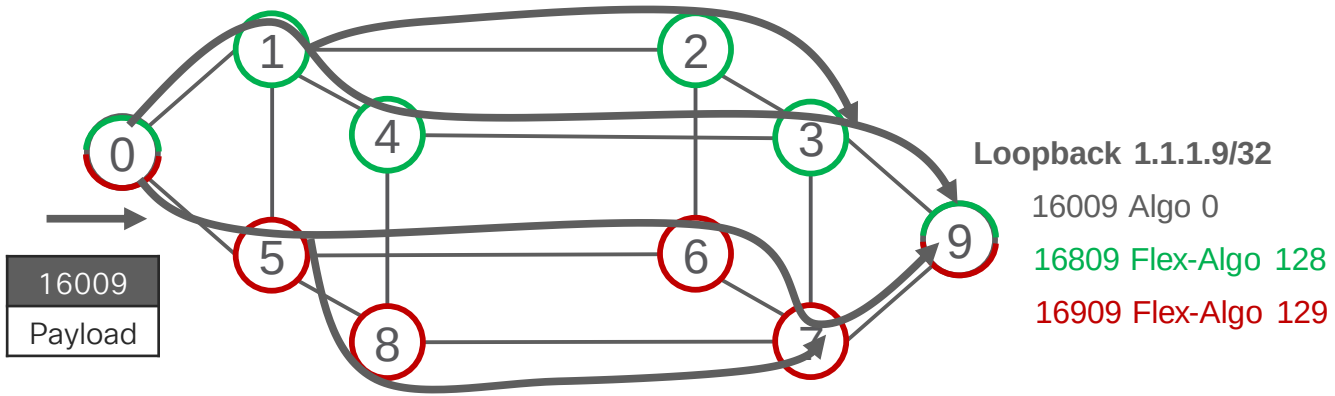

- Path to Node 9 across Algo 0

Multi-Plane Networks (cont.)

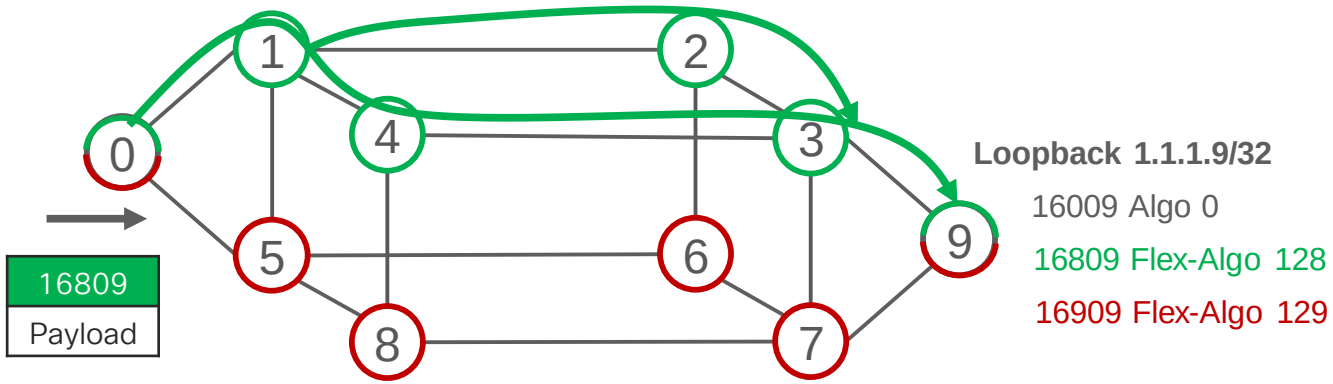

- Path to Node 9 across Flex-Algo 128

Multi-Plane Networks (cont.)

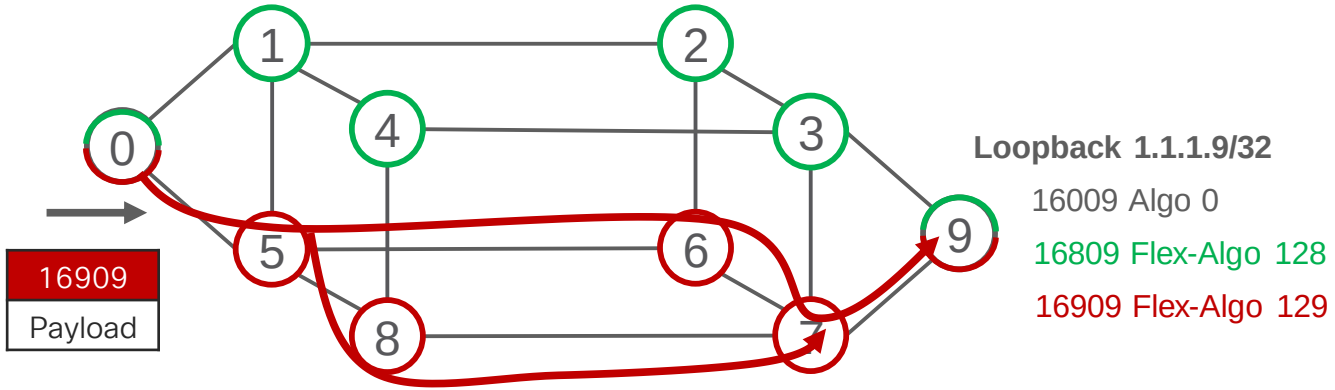

- Path to Node 9 across Flex-Algo 129

Automated Steering

- SRTE Automated Steering is leveraged for IGP Flex-Algo

```
segment-routing
  traffic-eng
    on-demand color 100
       dynamic mpls
          sid-algorithm 128
```

"Any 100-colored BGP route should be steered via the prefix-SID(ALGO 128) of the BGP nhop"

Automated Steering – Dual Plane

```
segment-routing
  traffic-eng
    on-demand color 100
      dynamic mpls
       sid-algorithm 128

router isis 1
  flex-algo 128
```

```
FIB

8/8: push 80000 16809
```

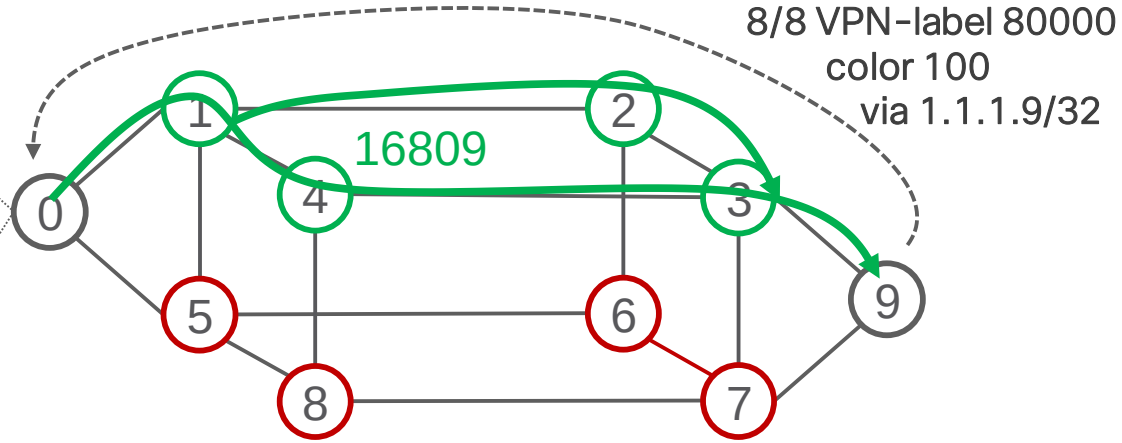

- Node 0 automatically steers any BGP route with color 100 from 9 via 16809 hence via the green plane only
- One single Flex-Algo Prefix-SID expresses the end-to-end SLA path

Automated Steering – Dual Plane

```
segment-routing
  traffic-eng
    on-demand color 100
      dynamic mpls
       sid-algorithm 128

router isis 1
  flex-algo 128
```

```
FIB
7/8: push 70000 16009
```

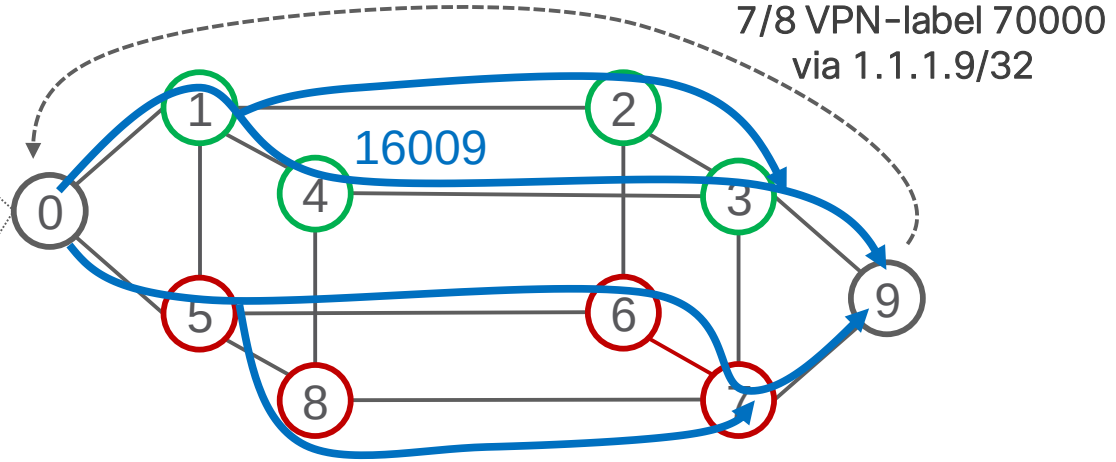

- Node 0 automatically steers any BGP route without color from 9 via 16009 (any plane)

Enable SR Flex-Algorithm

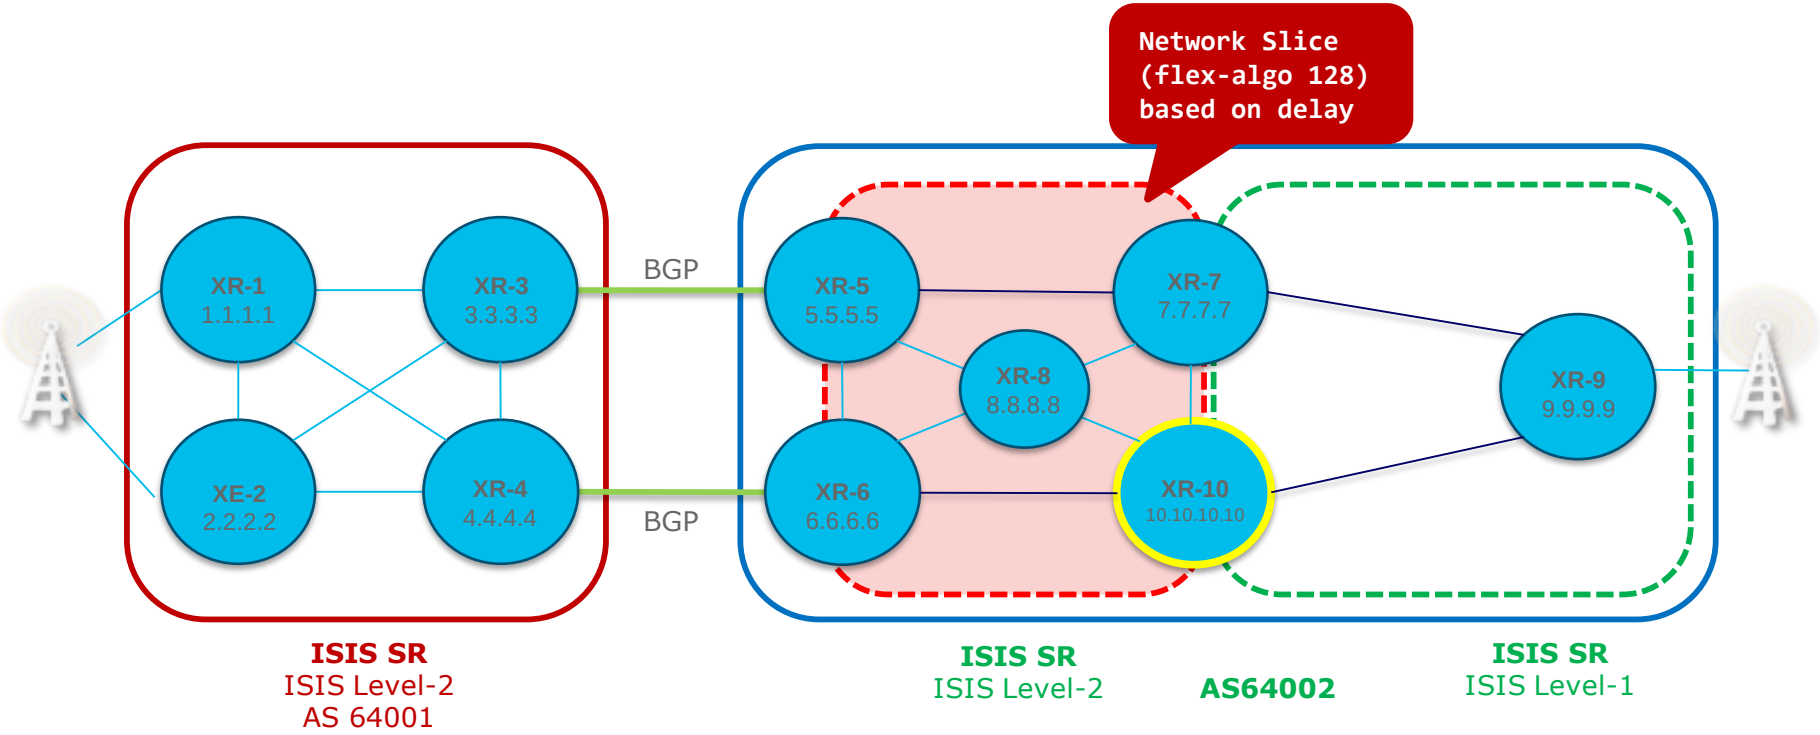

SR ODN with Flex-Algo

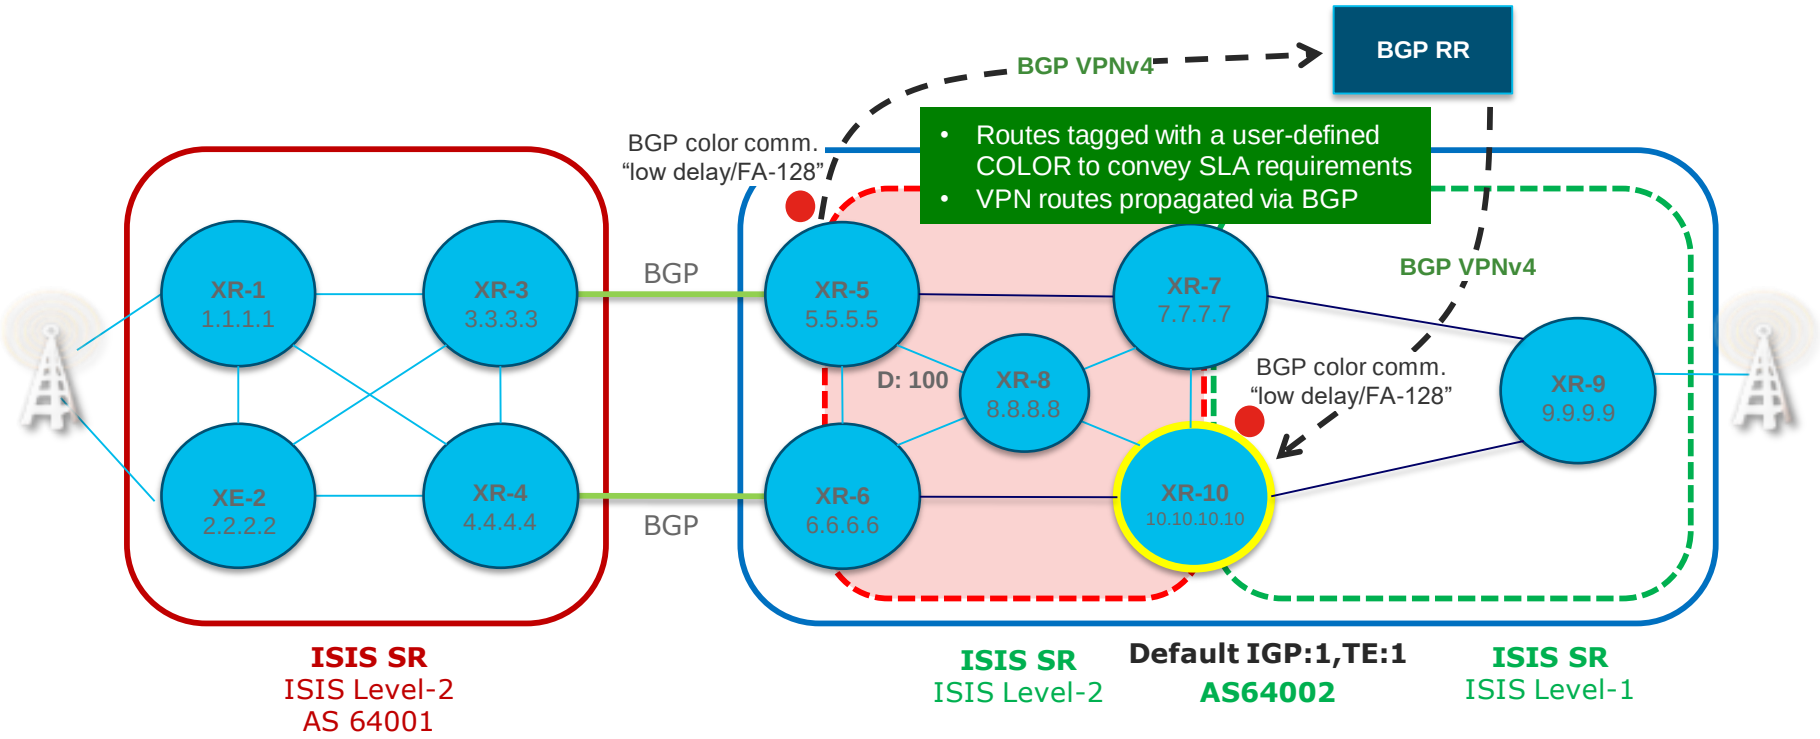

SR ODN with Flex-Algo

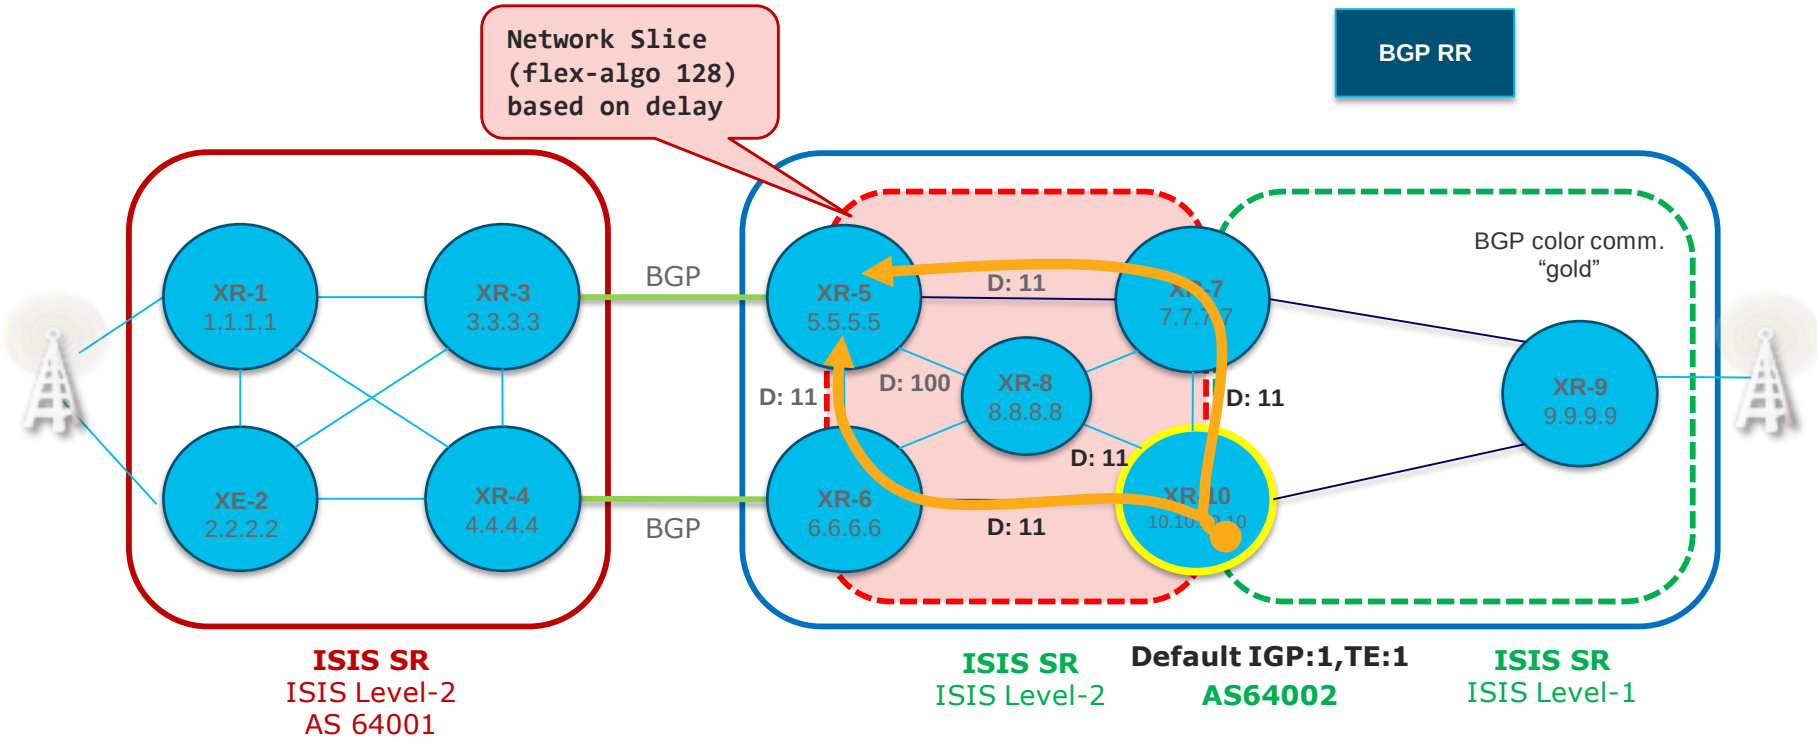

Conclusions

Conclusion

- In this lab you learnt:
  - Segment Routing (SR) configuration in IOS-XR and IOS-XE
  - LDP to SR migration steps
  - SR verification and monitoring
  - Topology-Independent Loop Free Alternate (TI-LFA) configuration and verification
  - SR On-Demand Next-Hop (ODN) for on-demand instantiation of SR policies
  - Automatic traffic steering onto SR policies without performance degradation
  - IOS XR's SR-PCE acting as stateful PCE for multi-domain SR-TE policies
  - New SR-TE infra for IOS-XR and IOS-XE head-end nodes
  - MPLS Performance Management
  - SR Flexible Algorithm
  - NSO for faster and more reliable service orchestration

Stay Up-To-Date

http://www.segment-routing.net/

https://www.linkedin.com/groups/8266623

https://twitter.com/SegmentRouting

https://www.facebook.com/SegmentRouting/

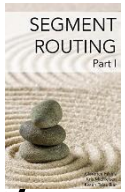

Segment Routing, Part I - Textbook

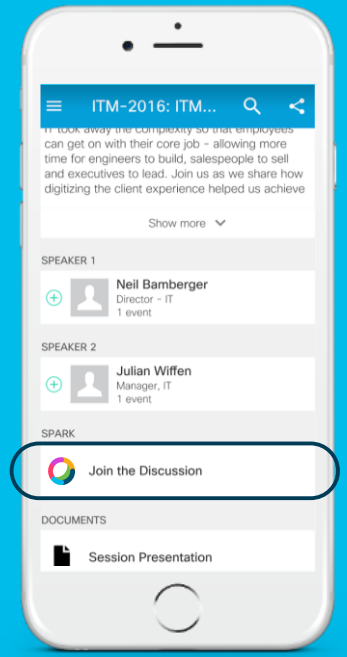

cs.co/ciscolivebot#LTRMPL-2201

Cisco Webex Teams

Questions?

Use Cisco Webex Teams (formerly Cisco Spark) to chat with the speaker after the session

How

1. Find this session in the Cisco Events Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

Complete your online session survey

- Please complete your Online Session Survey after each session
- Complete 4 Session Surveys & the Overall Conference Survey (available from Thursday) to receive your Cisco Live T-shirt
- All surveys can be completed via the Cisco Events Mobile App or the Communication Stations

Don't forget: Cisco Live sessions will be available for viewing on demand after the event at ciscolive.cisco.com

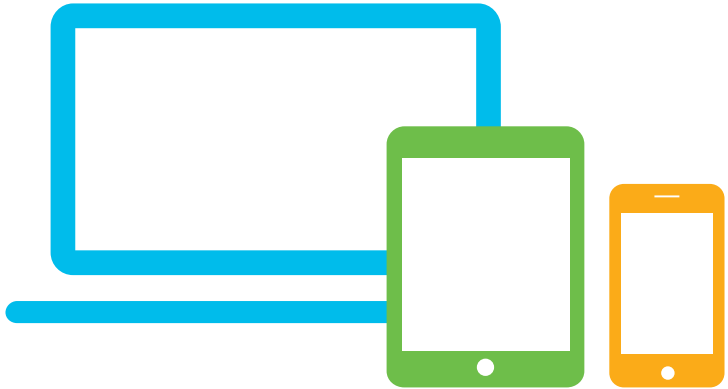

Cisco live!

Continue Your Education

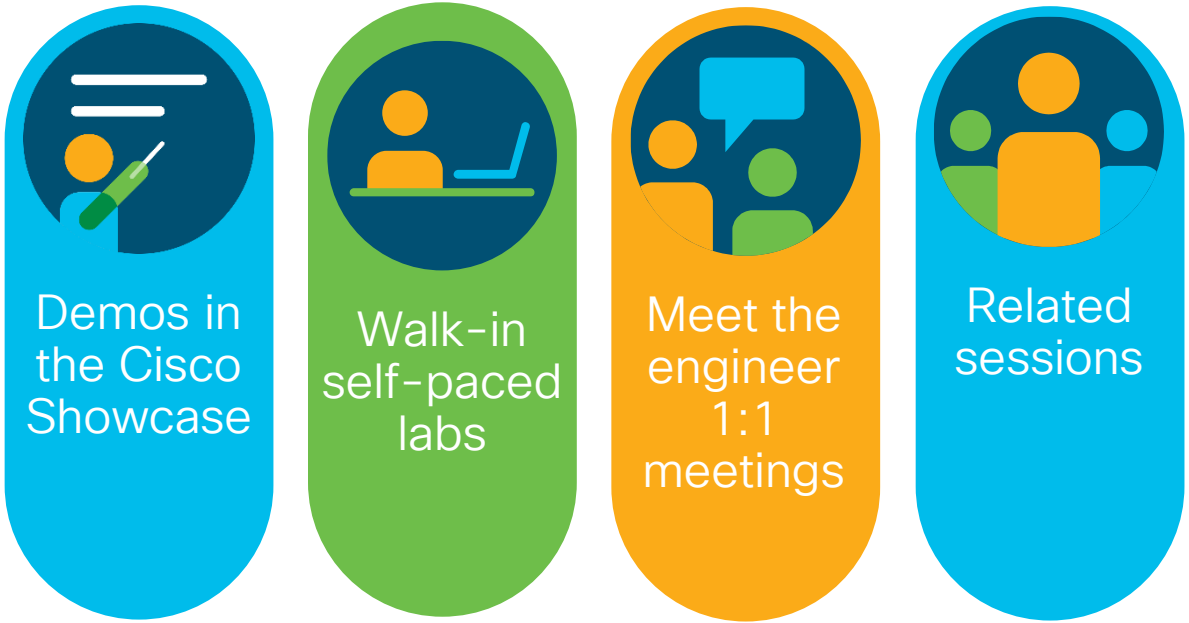

Thank you

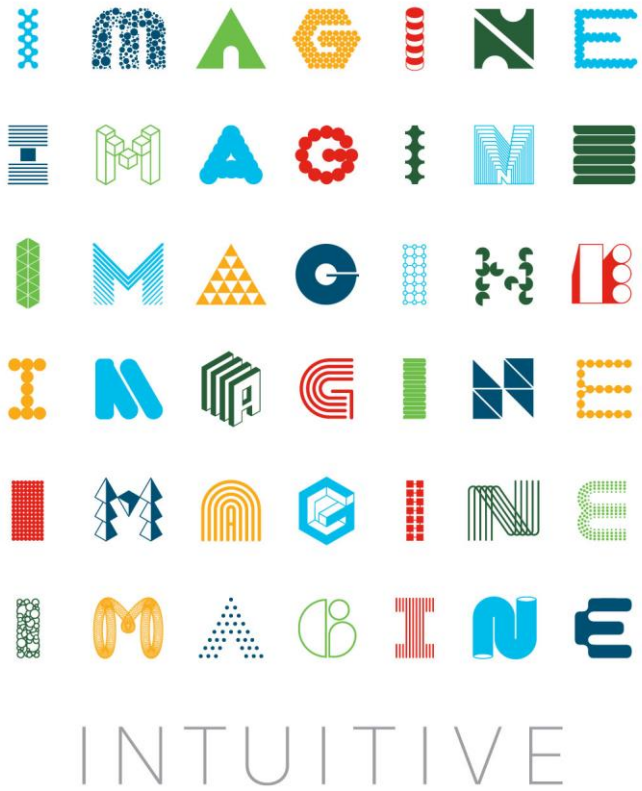

INTUITIVE