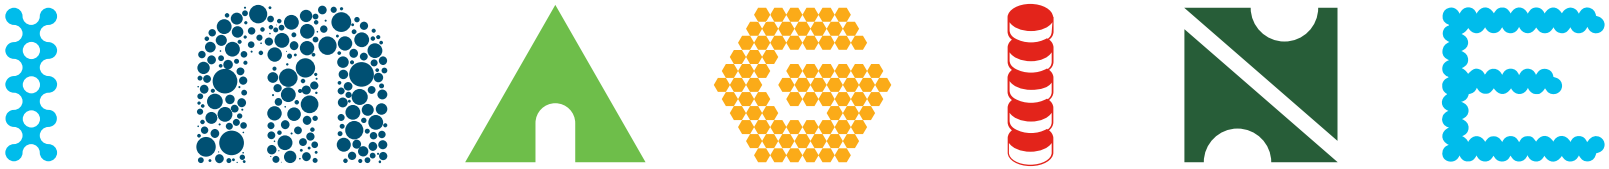

INTUITIVE

Cisco live!

June 10-14, 2018 • Orlando, FL

#CLUS

Troubleshooting IS-IS

Brad Edgeworth, Systems Engineer,
CCIE#31574 (R&S / SP)
@BradEdgeworth

BRKRST-3302

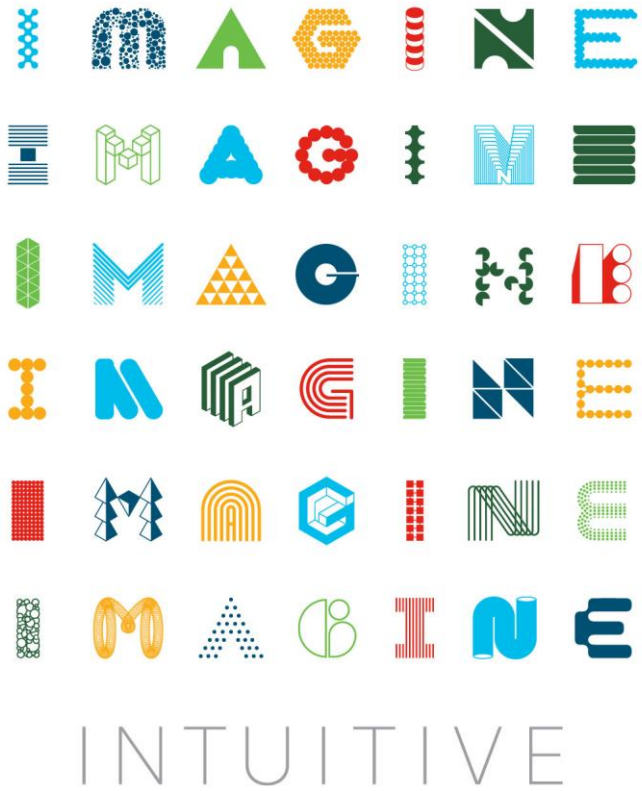

Agenda

- Introduction to IS-IS
  - Overview
  - Configuration
- Neighbor Adjacencies
  - Verification of Packets
- Viewing IS-IS Topologies
  - Building of Topologies
- Metrics
- Route Advertisements
- Sub-Optimal Routing
- Design Considerations
- IPv6 Topologies

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 Live 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

Webex Teams will be moderated by the speaker until June 18, 2018.

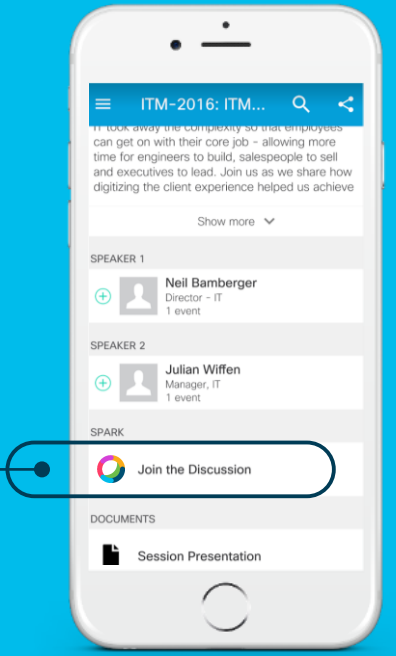

cs.co/ciscolivebot#BRKRST-3302

Setting Expectations

- This session covers IOS, IOS XR, and NX-OS
- This is a 3000 series Cisco Live course. You can learn IS-IS if you are a beginner, but we will go deep on some topics.
- There are some circular references, but we try to minimize those.
- There is a lot of content we are going to cover. At the end you will realize how easy IS-IS is to work with.
- This icon references a hidden slide, it will not be covered if you are following along with the PDF. 

IS-IS Trivia Question

- How do you make the traffic between R1 and R4, take R5?
- How do you make the traffic between R2 and R3, take the direct link?

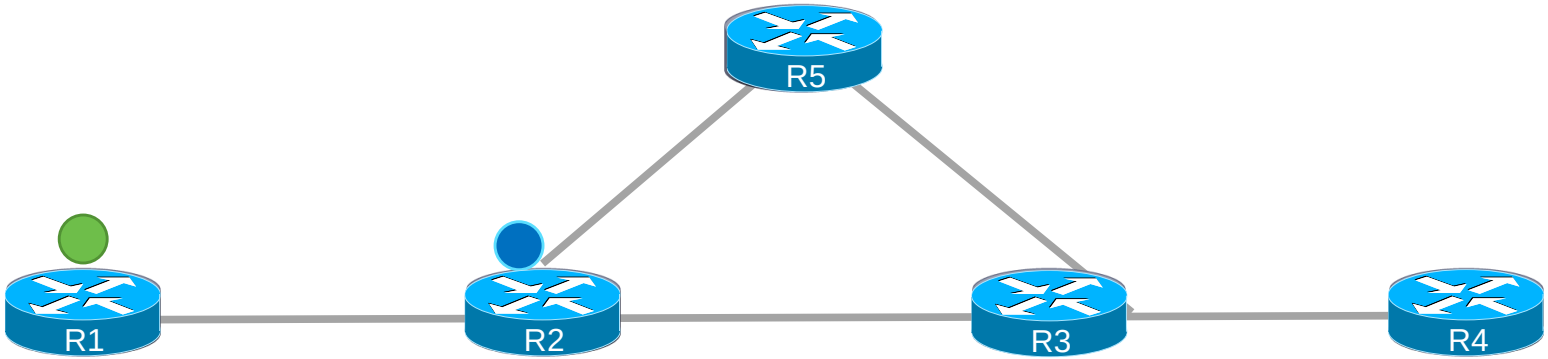

Introduction to IS-IS

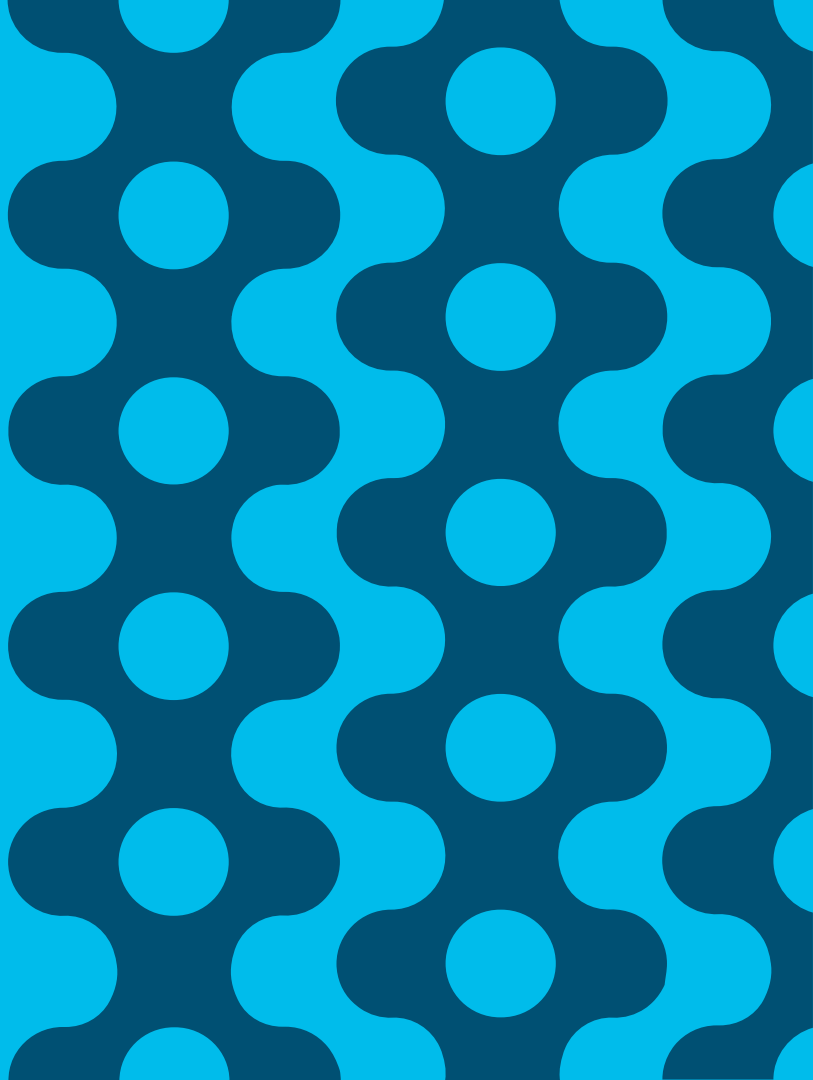

What is IS-IS?

Intermediate System-to-Intermediate System (IS-IS) Overview

- IS-IS is a link-state routing protocol;
  - Offer Fast convergence
  - Excellent scalability
  - Flexibility in terms of tuning
- Adopted and published by International Organization for Standardization (ISO)… The guys who gave us the OSI model
- Easily extensible with Type/Length/Value (TLV) extensions;
  - IPv6 Address Family support (RFC 2308)
  - Multi-Topology support (RFC 5120)
  - MPLS Traffic Engineering (RFC 3316)

Hierarchy Levels

- IS-IS presently has a two-layer hierarchy
  - The backbone (level 2)
  - Non-backbone areas (level 1)

Routers, not interfaces are associated to an area

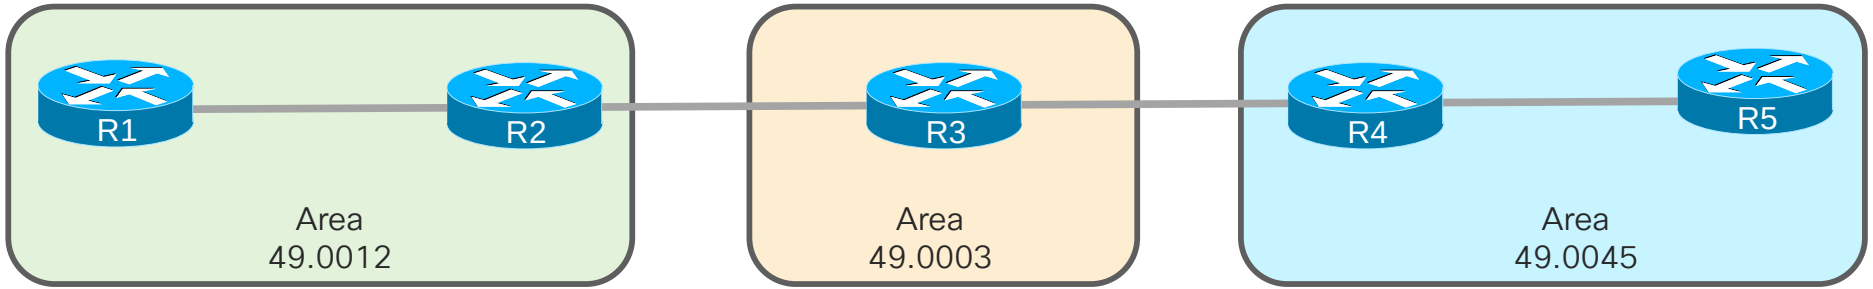

Hierarchy Levels

- IS-IS presently has a two-layer hierarchy
  - The backbone (level 2)
    - Formed between areas
    - Formed within an area

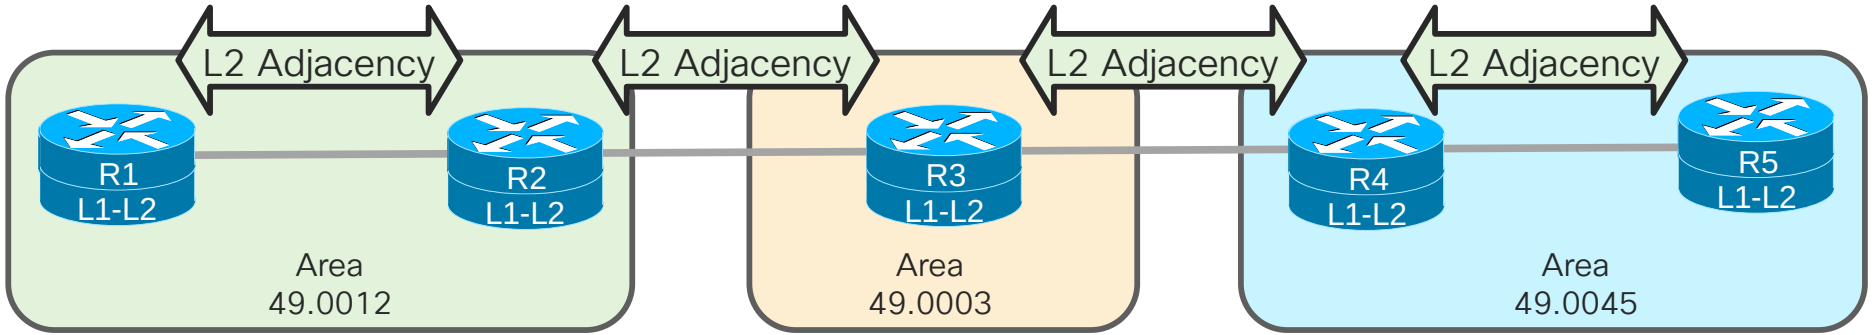

Hierarchy Levels

- IS-IS presently has a two-layer hierarchy
  - The backbone (level 2)
  - Non-backbone areas (level 1)
    - Formed within an area

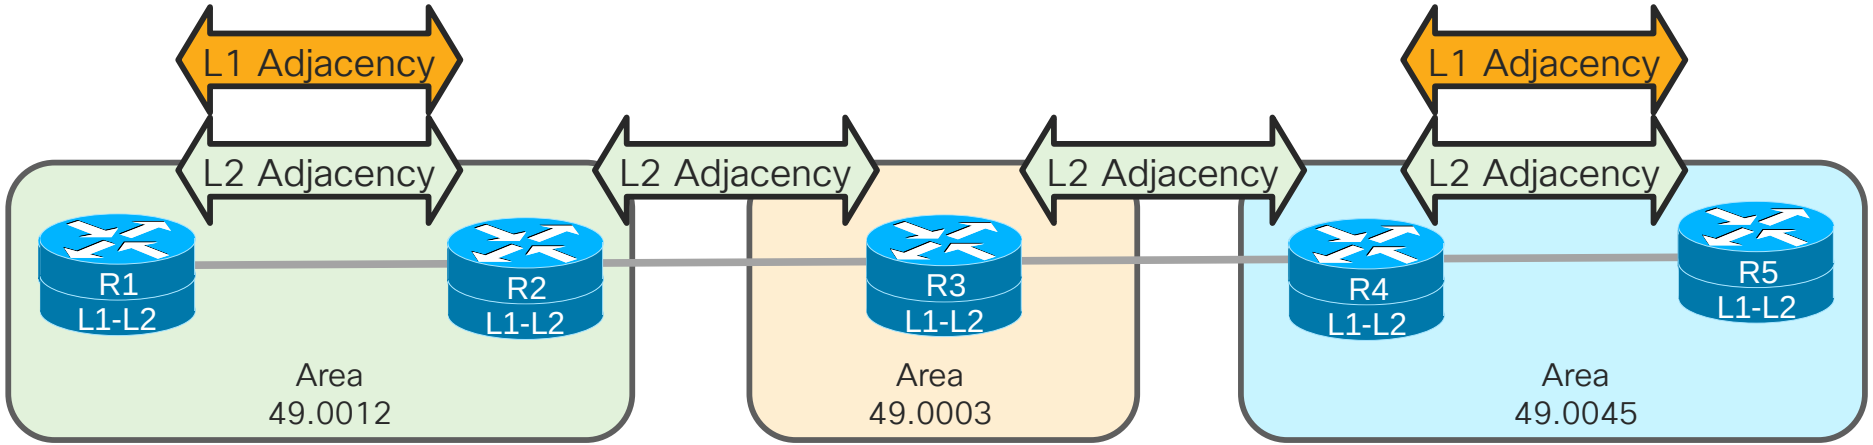

Hierarchy Levels (Routers)

A router can be a:

- L1 only router (i.e. R1 & R5)
- L2 only router (i.e. R3)
- L1-L2 routers (i.e. R2 & R4)

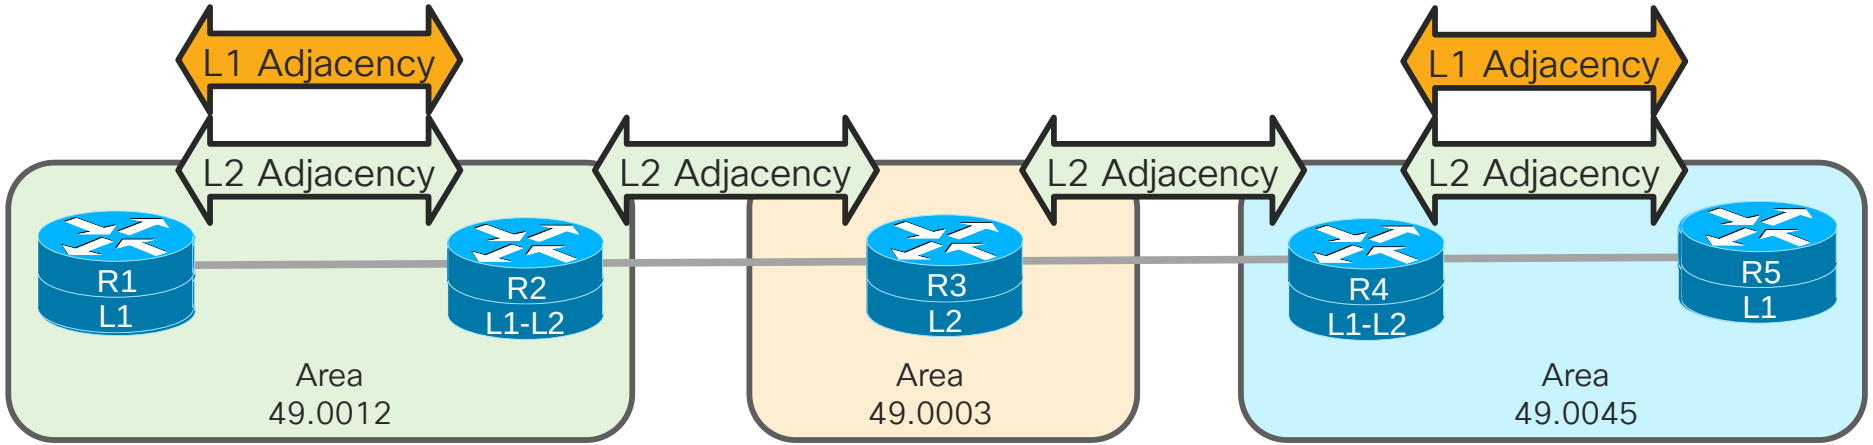

The Backbone

- Connects Areas
- Responsible for taking routes from L1 routes and advertising to other domains
- Can cross multiple areas

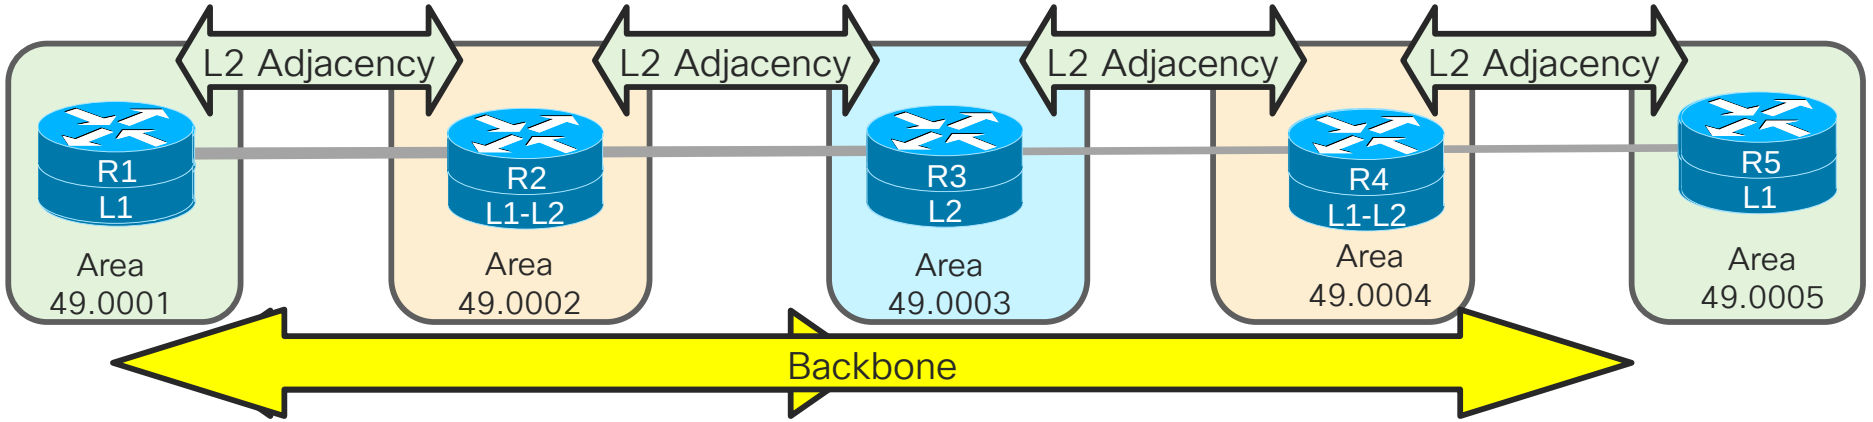

IS-IS Adjacencies with Hierarchy Levels

- Level 1-only routers
  - Can only form adjacencies with Level 1 or Level-1-2 routers in the same area
- Level-1-2 (default mode for Cisco devices)
  - Can form L1 adjacencies with Level 1 and Level-1-2 routers in the same area
  - Can form L2 adjacencies with Level-1-2 and Level 2 routers in the same or different area
- Level-2-only routers
  - Can form L2 adjacencies with other Level-1-2 and Level 2 routers in the same or different area

IS-IS Communication

Based on OSI Layer 2 addresses (MAC addresses for Ethernet).

Name	Destination MAC Address
All L1 IS Devices	0180.c200.0014
All L2 IS Devices	0180.c200.0015
All IS Devices	0900.2b00.0005

- Does not work on IP based tunnels
- Does work on GRE tunnels
- Does not work on DMVPN tunnels

IS-IS Packet Types

IS-IS has three types of PDUs (packets)

- **IS-IS Hello (IIH) Packets** – Used to establish/monitor neighbors
- **Link State Packets (LSPs)** – used to build a topology and share routes
- **Sequence Number Packets (SNPs)** – used to synchronize LSPs

Type, Length, Value (TLV) Tuples

- This is the true magic of IS-IS. Provides the ability to support multiple protocols in the same architecture.
- TLVs provide variable modules, and support nesting.
- A TLV is assigned a numerical value which directly correlates to a function.
- When a router receives a IS PDU and detects an unrecognizable TLV, it just skips the TLV and continues to the next TLV in that packet.
- TLVs are not modified in transit.

IS-IS Packet Structure

IIS-IS Hello (IIH), LSPs, SNPs all contain these fields in every packet:

- Protocol Descriptor – 0x83 for IS-IS
- PDU Length
- PDU Type – Defines if it is an IIH, LSP, or SNP
- Reserved – Identifies the level of a packet (L1 or L2)
- Max Areas – Maximum number of areas a router will support

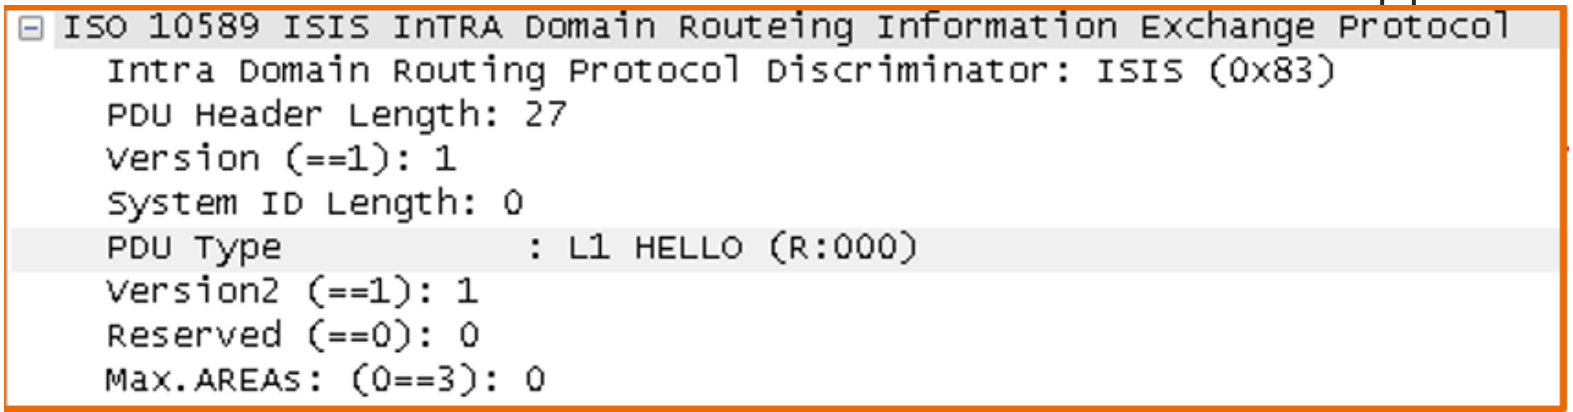

IS-IS Hello Structure

Type	Description
L1 IS-IS hello (IIH) PDU Type 15	Discovers, forms, and maintains L1 IS-IS neighbors
L2 IS-IS hello (IIH) PDU Type 16	Discovers, forms, and maintains 2 IS-IS neighbors
P2P hello (IIH) IS-IS PDU Type 17	Discovers, forms, and maintains P2P IS-IS neighbors
End system hello (ESH)	Used for end systems (ESs) to discover intermediate systems (ISs) and vice versa; similar to Internet Control Message Protocol (ICMP)
Intermediate system hello (ISH)	Used for ESs to discover ISs and vice versa for router selection

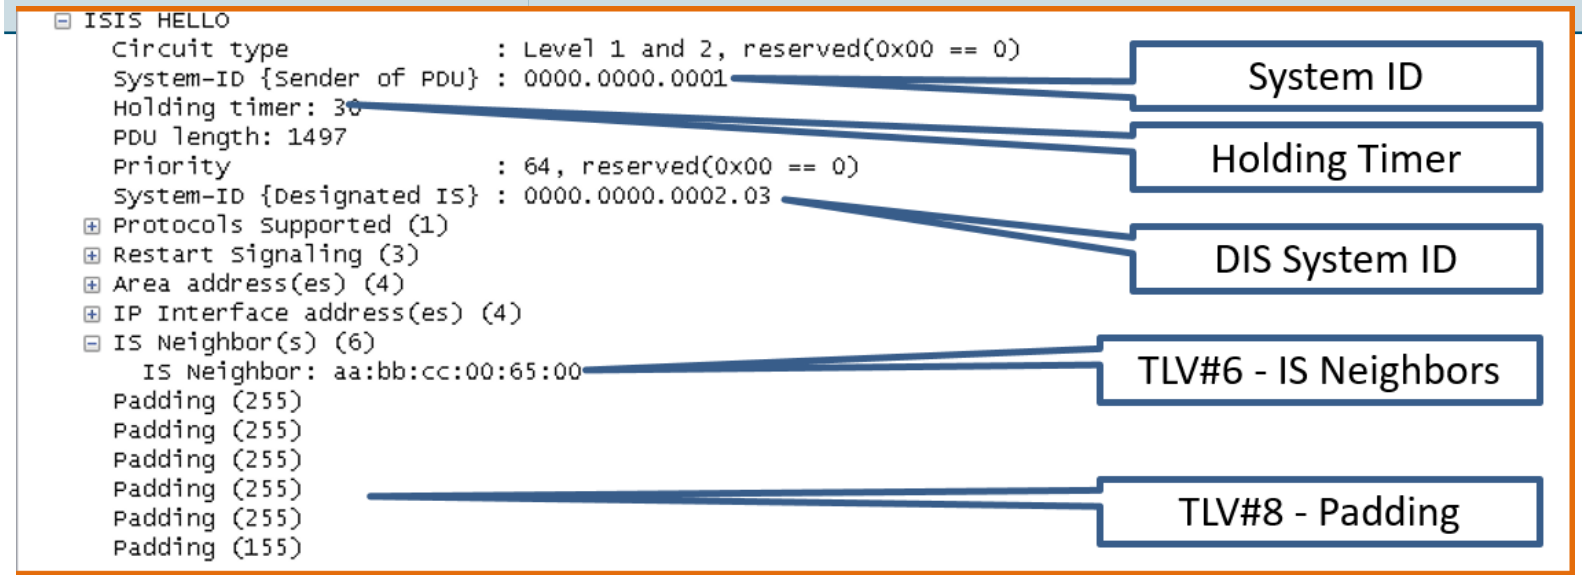

IS-IS LSP Structure

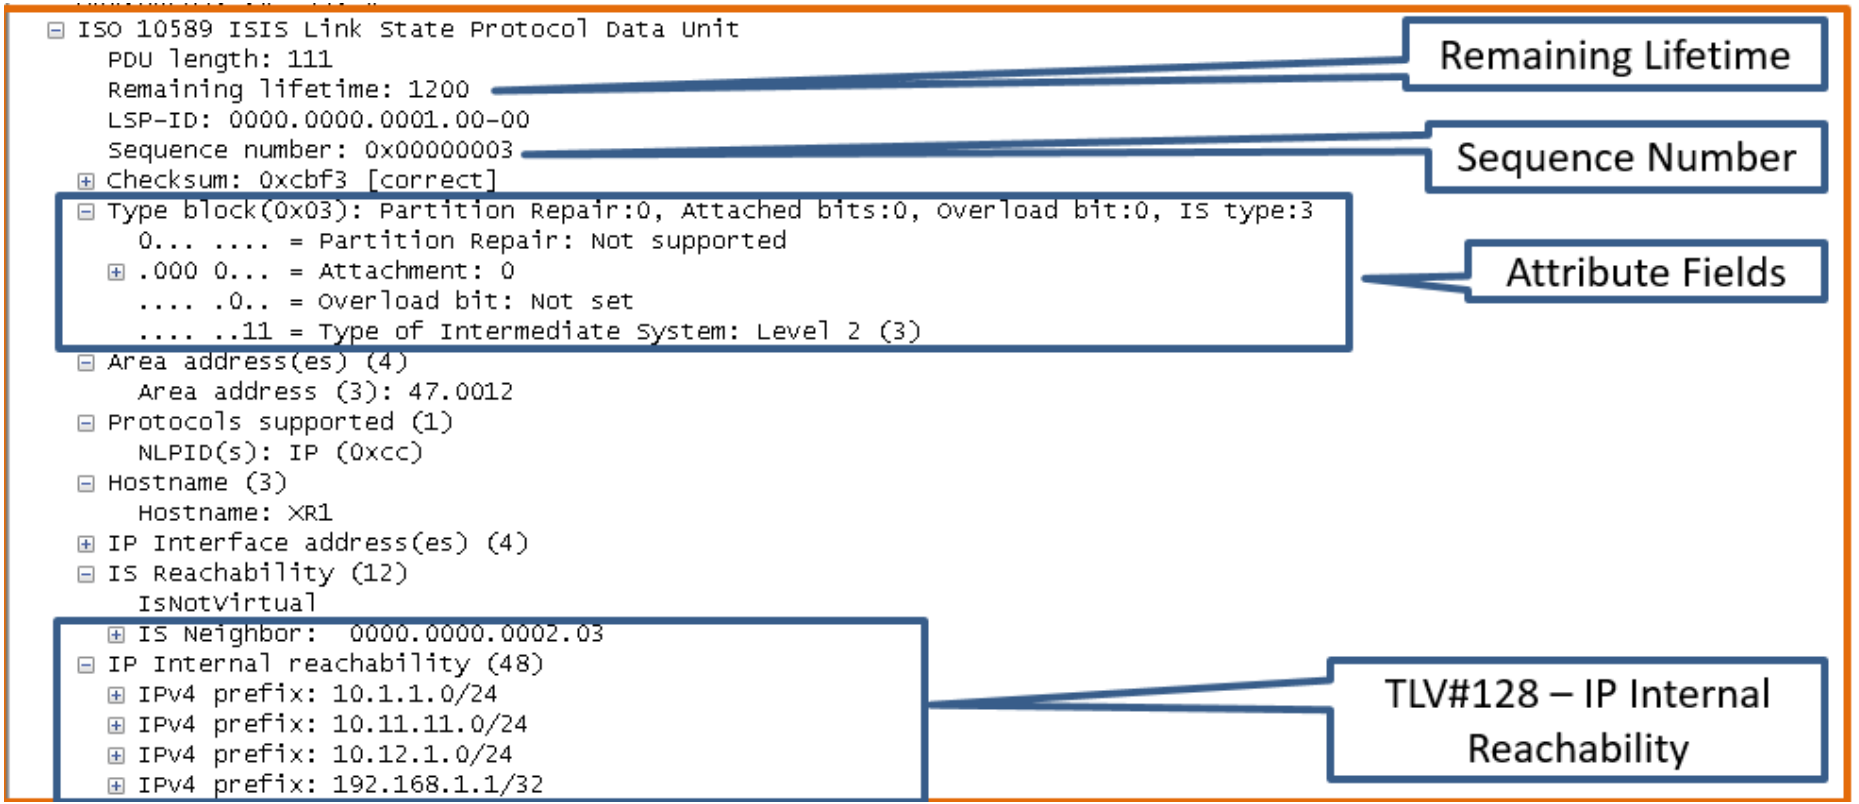

Common LSP TLVs

TLV #	Function
1	List of area addresses on router
2	List of IS Neighbors (Narrow Metrics)
10	Authentication
22	Extended IS Neighbors (Wide Metrics)
128	IP network and metric from advertising router (Narrow Metrics)
130	External networks and metrics when redistributed
132	IP Addresses on transmitting interface (includes secondary interfaces) (Narrow Metrics)
135	IP Addresses on transmitting interface (includes secondary interfaces) (Wide Metrics)
137	Router hostname (Allows correlation of name to System ID
232	IPv6 Interface Address
236	IPv6 Reachability Information
237	Multi Topology Reachable IPv6 Prefix

IS-IS Packet Structure

IIS-IS Hello (IIH), LSPs, SNPs all contain these fields in every packet:

- Protocol Descriptor – 0x83 for IS-IS
- PDU Length
- PDU Type – Defines if it is an IIH, LSP, or SNP
- Reserved – Identifies the level of a packet (L1 or L2)
- Max Areas – Maximum number of areas a router will support

```
ISO 10589 ISIS InTRA Domain Routeing Information Exchange Protocol
    Intra Domain Routing Protocol Discriminator: ISIS (0x83)
    PDU Header Length: 27
    Version (==1): 1
    System ID Length: 0
    PDU Type            : L1 HELLO (R:000)
    Version2 (==1): 1
    Reserved (==0): 0
    Max.AREAs: (0==3): 0
```

IS-IS Interfaces

There are two types of interface in IS-IS:

- **Broadcast** – This is the default. Allows for more than one neighbor to connect on this medium. Requires the election of a pseudonode called a Designated Intermediate System (DIS)
- **Point-to-Point** – Used to reduce some of the overhead mechanisms with broadcasts networks if only 2 devices exist on a segment.

IS-IS Addressing

Anatomy of a NET Address

- Each IS-IS router is identified with a Network Entity Title (NET)

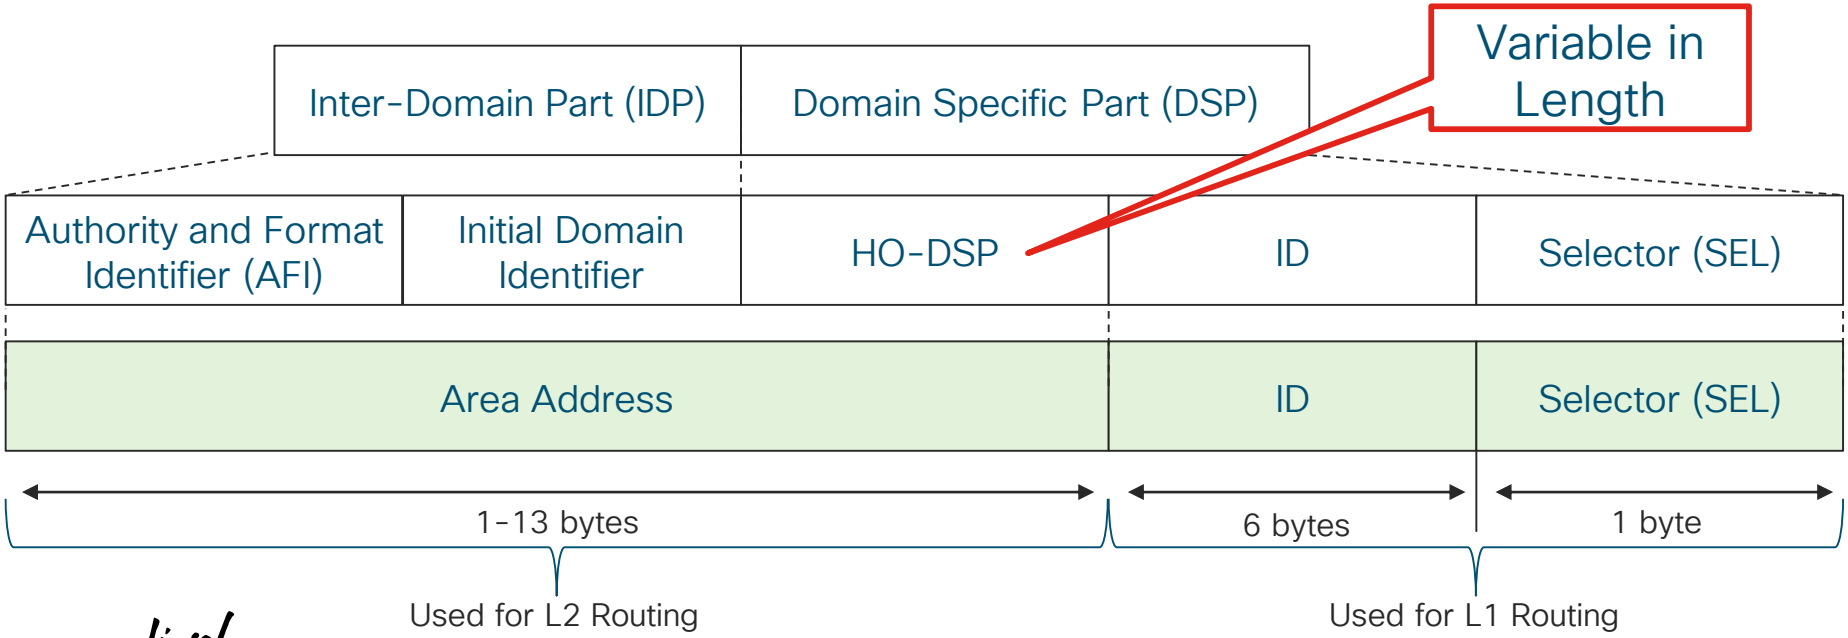

IS-IS Addressing

Reading the NET Address

- Start from Right work your way back:
  - Final 8 bits – zero
  - Next 48 bits – router identifiter
  - Next 16 bits – area
  - First 8 bits – pick a number (49 is the private AFI family)

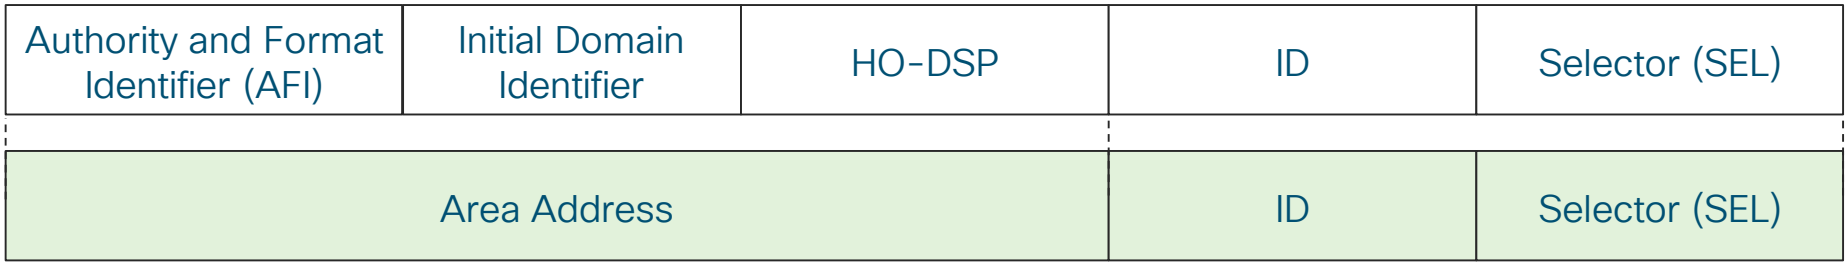

IS-IS Addressing

Sample NET Addresses

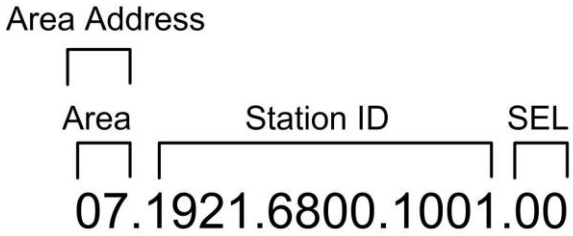

8-byte Area ID/System ID

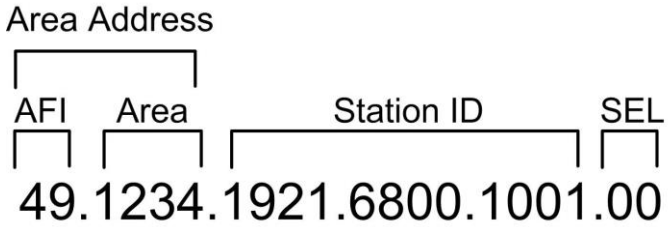

10-byte Private AFI/
Area ID/System ID

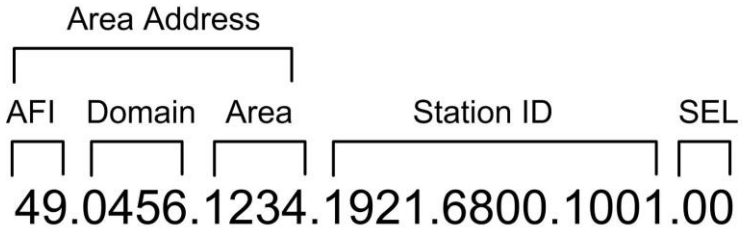

OSI NSAP Format

IS-IS Addressing

Sample NET Addresses

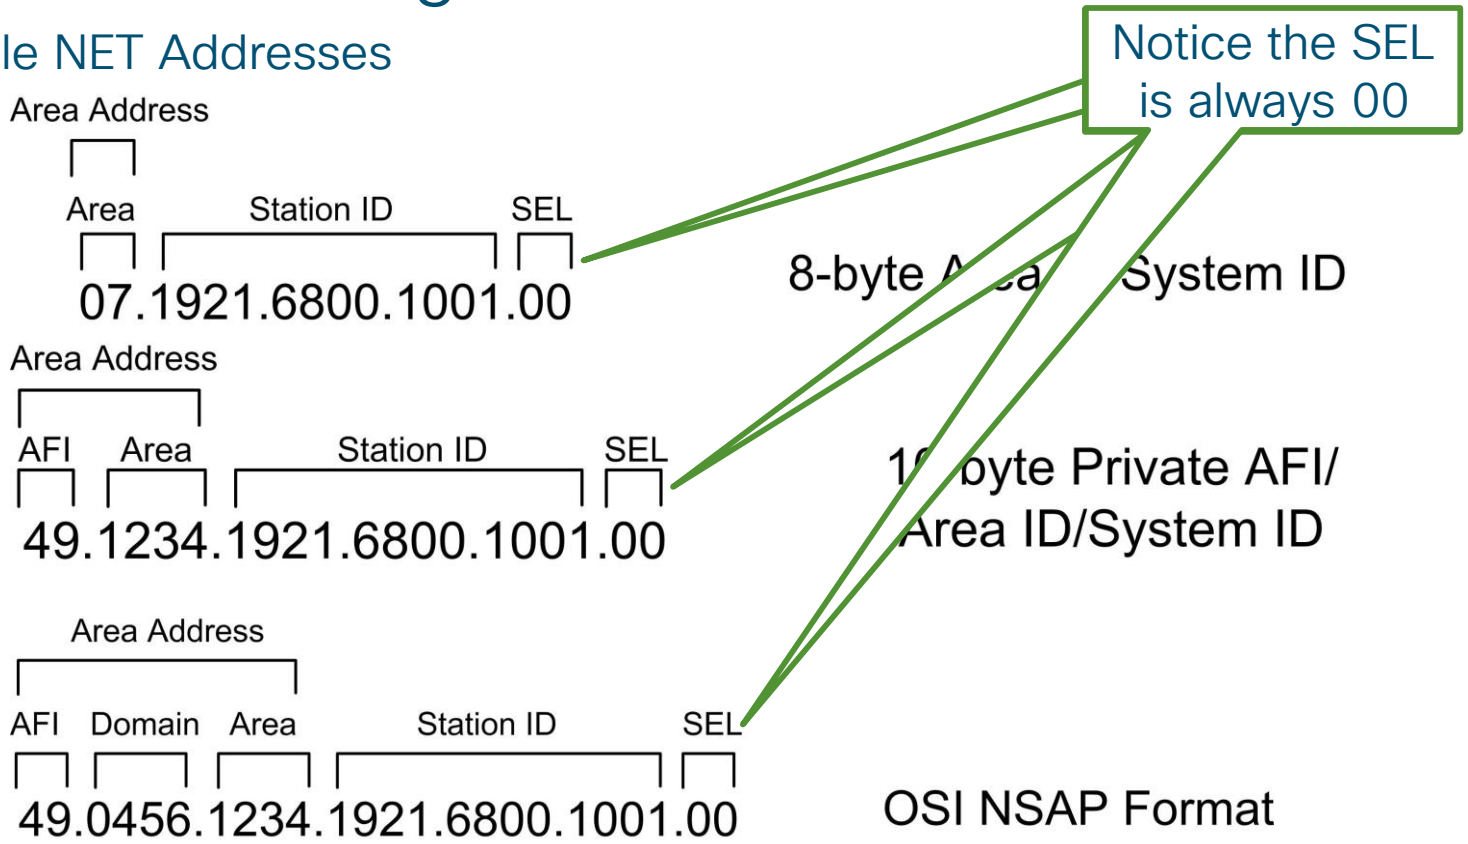

Designated Intermedate System (DIS)

- Broadcast networks support more than two routers which could cause scalability problems with IS-IS
- IS-IS overcomes this inefficiency by creating a pseudonode to manage synchronization issues that arise on the broadcast network segment. A DIS exist for each IS-IS level (L1 and L2).
- By inserting the logical pseudonode into a broadcast segment, the multi-access network segment is converted into multiple P2P networks in the LSPDB

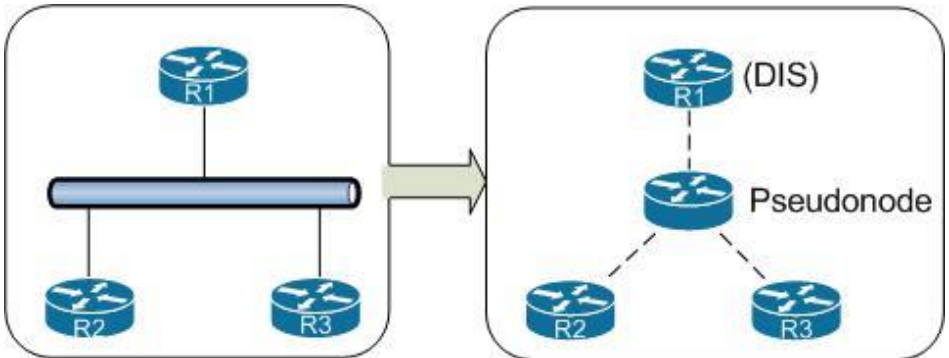

Corny Network Engineer Joke

- Three routers walk into a bar (R1, R2, and R3)
- They have a couple of shots…?
- Who Drives home?

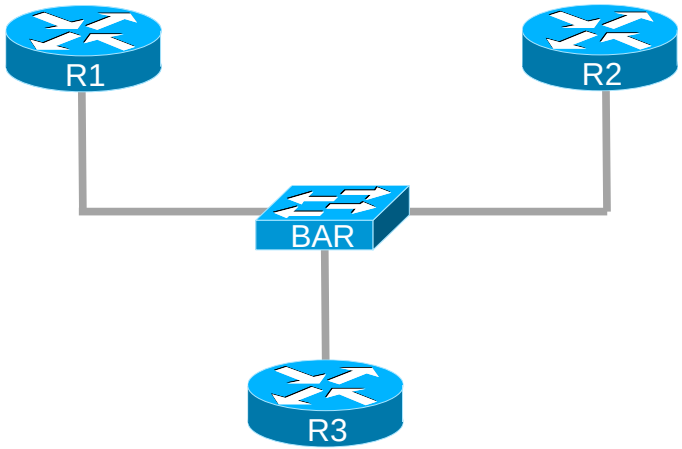

ISIS vs OSPF

Notable Similarities and Differences

* draft-bhatia-manral-diff-isis-c

IS-IS and OSPF are both link state protocols, there are similarities and differences

- Similarities:
  - Link-state representation, aging, and metrics
  - Use of Link-state databases and SPF algorithms
  - Update, routing decisions, and flooding processes similar
- Differences:
  - IS-IS organizes domain into two layers; OSPF designates backbone area (area 0)
  - IS-IS peering is more flexible than OSPF (hello time, dead intervals, and subnet mask need not match)
  - IS-IS selects single DIS which may be preempted; OSPF elects a DR/BDR which cannot be preempted,
  - IS-IS does not support NBMA, point-to-multipoint, or virtual links (it rides L2 directly)

IS-IS Configuration

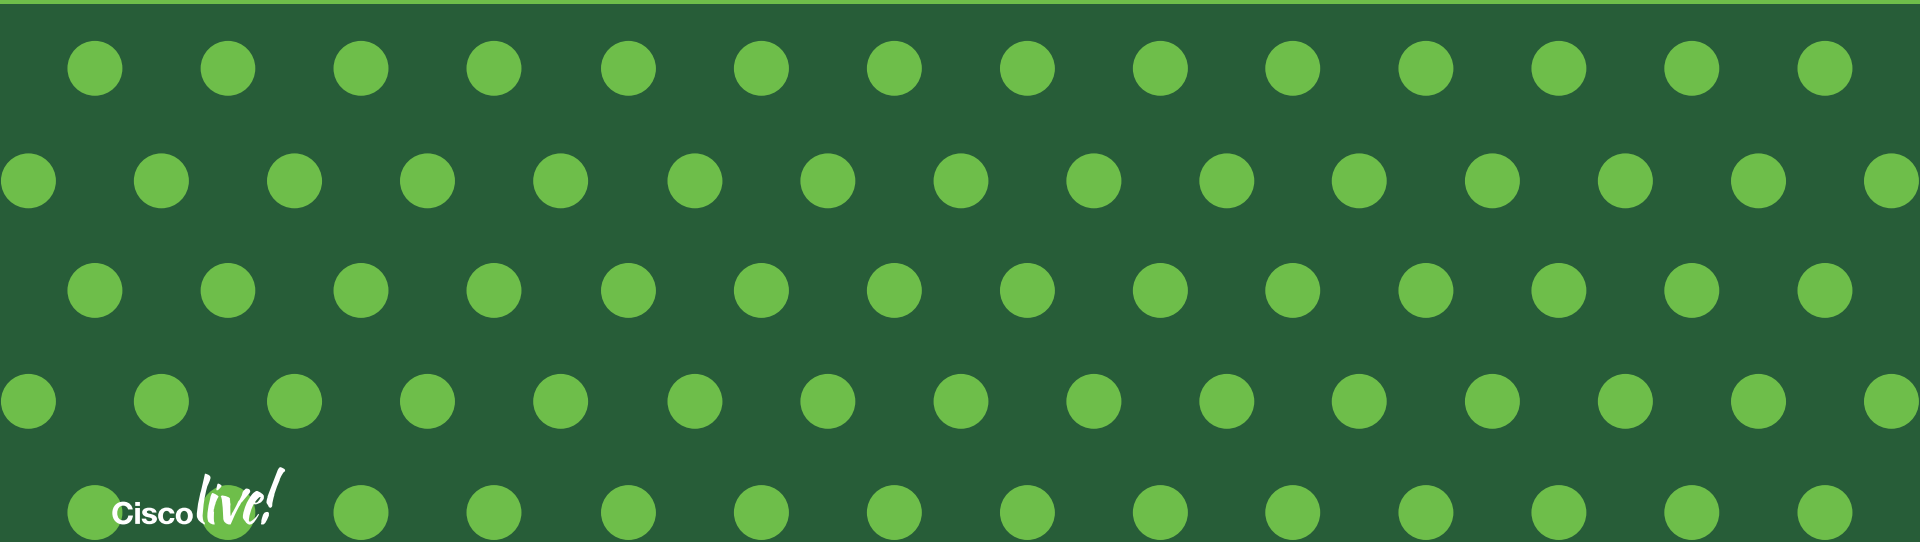

Topology for Configuration

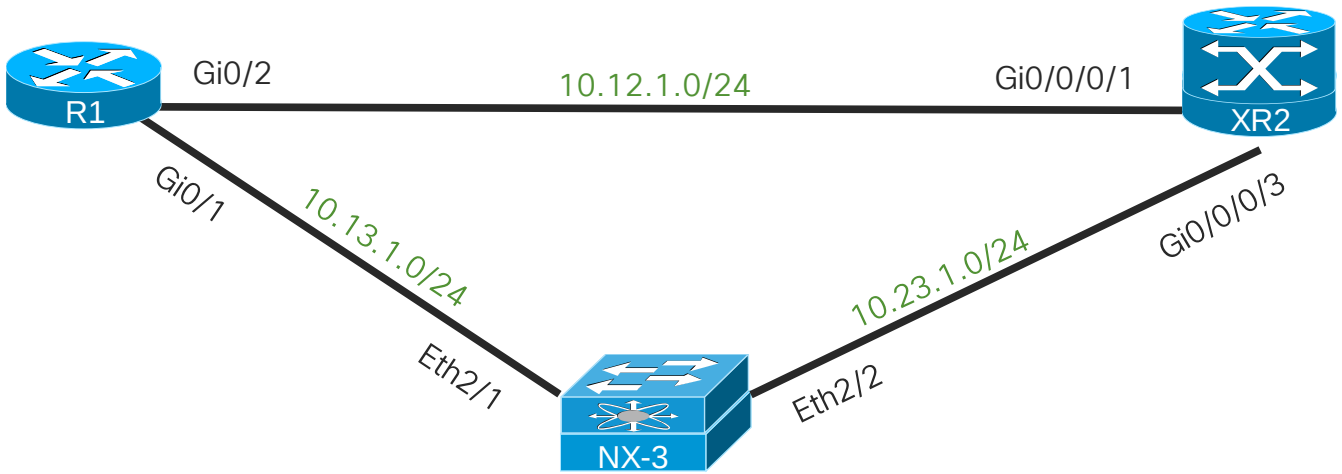

IS-IS Configuration: IOS

- Initialize the routing protocol

```
router isis [process-id]
```

- Enable Adjacency Logging (Optional)

```
log-adjacency-changes
```

- Define the NET Address

```
net area-systemid.sel
```

- Enable IS-IS on the interface

```
interface interface-id
  ip router isis [process-id]
  ipv6 router isis [process-id]
```

IS-IS Configuration: IOS

```
R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#router isis CISCOLIVE
R1(config-router)# log-adjacency-changes
R1(config-router)# net 49.0123.0001.0001.0001.00
R1(config-router)#interface gi0/2
R1(config-if)#  ip router isis CISCOLIVE
R1(config-if)#  ipv6 router isis CISCOLIVE
R1(config-if)#interface gi0/3
R1(config-if)#  ip router isis CISCOLIVE
R1(config-if)#  ipv6 router isis CISCOLIVE
03:38:39.967: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0002.0002.0002 (GigabitEthernet0/2)
Up, new adjacency
03:38:41.967: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0002.0002.0002 (GigabitEthernet0/2)
Up, new adjacency
```

IS-IS Configuration: IOS XR

- Initialize the routing protocol

```
router isis process-id
```

- Enable Adjacency Logging (Optional)

```
log-adjacency-changes
```

- Define the NET Address

```
net area-systemid.sel
```

- Initialize IPv6 Address family (optional)

```
address-family ipv6 unicast
```

- Enable IS-IS on the interface

```
interface interface-id

  address-family ipv4 unicast

  address-family ipv6 unicast
```

IS-IS Configuration: IOS XR

```
RP/0/0/CPU0:XR2#conf t
RP/0/0/CPU0:XR2(config)#router isis CISCOLIVE
RP/0/0/CPU0:XR2(config-isis)# log-adjacency-changes
RP/0/0/CPU0:XR2(config-isis)# net 49.0123.0002.0002.0002.00
RP/0/0/CPU0:XR2(config-isis)# interface gi0/0/0/1
RP/0/0/CPU0:XR2(config-isis-if)#  address-family ipv4 unicast
RP/0/0/CPU0:XR2(config-isis-if)#  address-family ipv6 unicast
RP/0/0/CPU0:XR2(config-isis-if-af)# interface gi0/0/0/3
RP/0/0/CPU0:XR2(config-isis-if)#  address-family ipv4 unicast
RP/0/0/CPU0:XR2(config-isis-if)#  address-family ipv6 unicast
RP/0/0/CPU0:XR2(config-isis-if-af)#commit
RP/0/0/CPU0:May        8:37.226 : isis[1010]: %ROUTING-ISIS-6-INFO_STARTUP_START : Cold
controlled start          ning

                          0.996 : isis[1010]: %ROUTING-ISIS-5-ADJCHANGE : Adjacency to
                          gabitEthernet0/0/0/1) (L1) Up, New adjacency
                          2.015 : isis[1010]: %ROUTING-ISIS-5-ADJCHANGE : Adjacency to
49.0123.0001.0001.0001 (GigabitEthernet0/0/0/1) (L2) Up, New adjacency
```

Notice the config is under isis process

IS-IS Configuration: NX-OS

- Enable the IS-IS feature

```
feature isis
```

- Initialize the routing protocol

```
router isis process-id
```

- Enable Adjacency Logging (Optional)

```
log-adjacency
```

- Define the NET Address

```
net area-systemid.sel
```

- Enable IS-IS on the interface

```
interface interface-id
  ip router isis process-id
  ipv6 router isis process-id
```

IS-IS Configuration: NX-OS

```
NX-3(config)# feature isis
NX-3(config)# router isis CISCOLIVE
NX-3(config-router)# net 49.0123.0003.0003.0003.00
NX-3(config-router)# log-adjacency
NX-3(config-router)# interface ethernet2/1
NX-3(config-if)# ip router isis CISCOLIVE
NX-3(config-if)# ipv6 router isis CISCOLIVE
NX-3(config-if)# interface ethernet2/2
NX-3(config-if)# ip router isis CISCOLIVE
NX-3(config-if)# ipv6 router isis CISCOLIVE
03:55:40 NX-3 %ISIS-5-ADJCHANGE:  isis-CISCOLIVE [9333]  LAN adj L1 0001.0001.0001 over
Ethernet2/1 - INIT (New) on MT--1
03:55:41 NX-3 %ISIS-5-ADJCHANGE:  isis-CISCOLIVE [9333]  LAN adj L2 0001.0001.0001 over
Ethernet2/1 - INIT (New) on MT--1
03:55:41 NX-3 %ISIS-5-ADJCHANGE:  isis-CISCOLIVE [9333]  LAN adj L2 0001.0001.0001 over
Ethernet2/1 - UP on MT-0
03:55:41 NX-3 %ISIS-5-ADJCHANGE:  isis-CISCOLIVE [9333]  LAN adj L1 0001.0001.0001 over
Ethernet2/1 - UP on MT-0
```

Troubleshooting IS-IS Connectivity

- Almost every network troubleshooting starts from one host to another. Start at the edge routers, and try to ping the far end.
- When that fails, we check to see somewhere in the middle.
- If that fails, then we check the router's routes made it to the middle and vice versa.
- We go one router away from the local router to see if we can see its routes
- Then we go to the edge of the area and see if we can see either route.

Trivia Question Hint

- How do you make the traffic between R1 and R4, take R5?
- How do[illegible]e direct link?

HINT!
One command on Three Routers

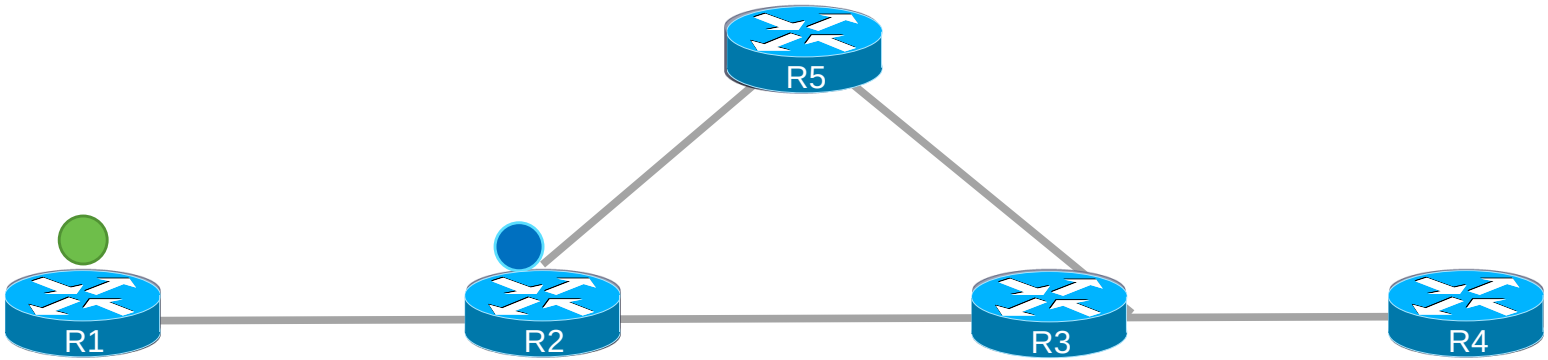

Neighbor Adjacencies

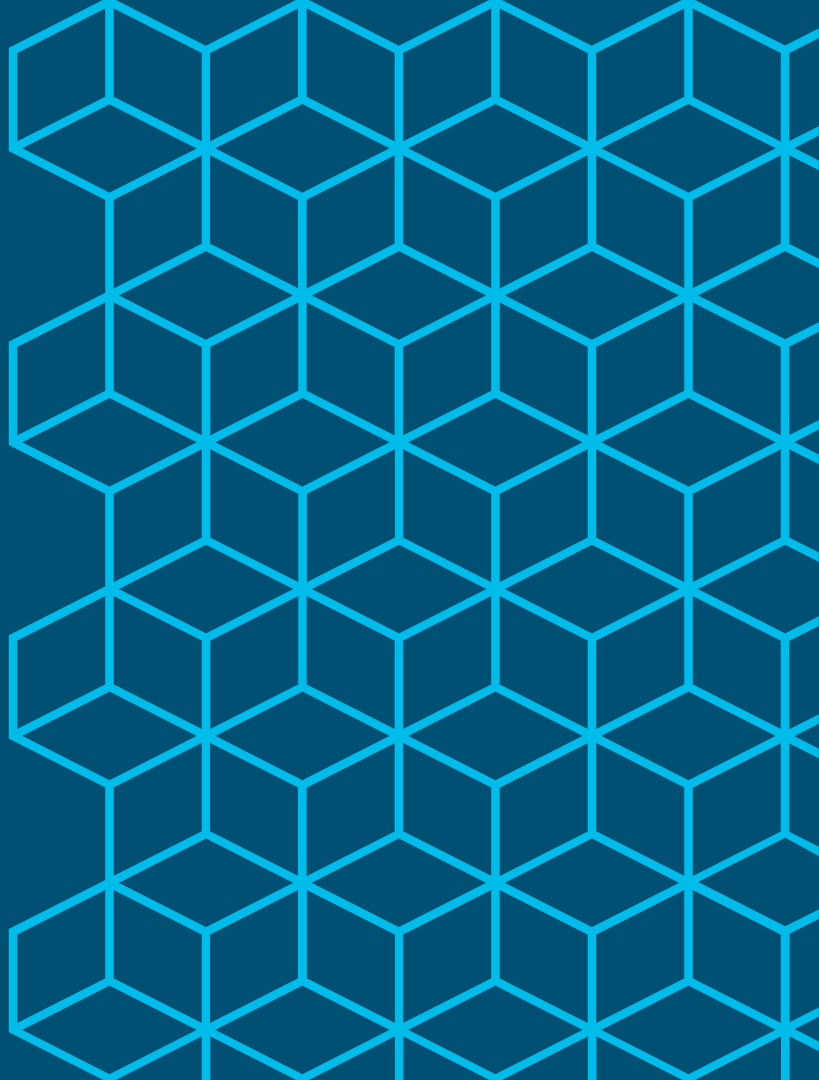

Checking for Neighbor Adjacencies

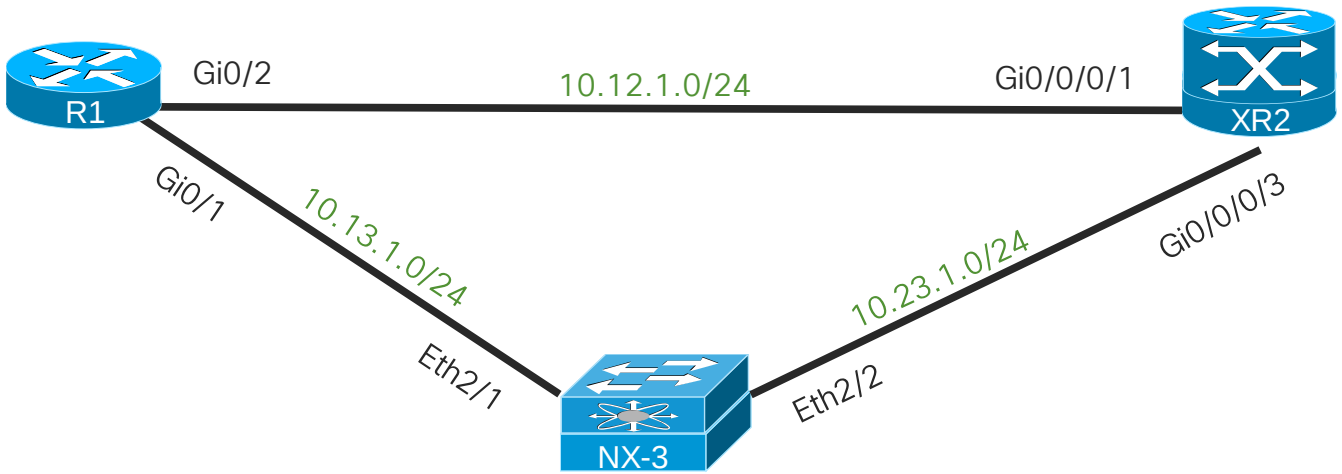

Viewing IS-IS Neighbors

IOS

```
R1#show isis neighbors
Tag CISCOLIVE:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
XR2             L1   Gi0/2         10.12.1.2       UP    22       R1.02
XR2             L2   Gi0/2         10.12.1.2       UP    26       R1.02
NX-3            L1   Gi0/3         10.13.1.3       UP    26       R1.01
NX-3            L2   Gi0/3         10.13.1.3       UP    27       R1.01
```

Viewing IS-IS Neighbors

IOS

```
R1#show isis neighbors
Tag CISCOLIVE:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
XR2             L1   Gi0/2         10.12.1.2       UP    22       R1.02
XR2             L2   Gi0/2         10.12.1.2       UP    26       R1.02
NX-3            L1   Gi0/3         10.13.1.3       UP    26       R1.01
NX-3            L2                 10.13.1.3       UP    27       R1.01
```

How did R1 find out the system ID for XR2 and NX-3?

Was it CDP?

Is it DNS?

Viewing IS-IS Neighbors

IOS

No! It is TLV#137

```
Tag CISCOLIVE:
System I
XR2
XR2
NX-3
NX-3
```

```
ISO 10589 ISIS InTRA Domain Routeing Information Exchange Protocol
    Intra Domain Routing Protocol Discriminator: ISIS (0x83)
    PDU Header Length: 27
    Version (==1): 1
    System ID Length: 0
    PDU Type           : L1 LSP (R:000)
    Version2 (==1): 1
    Reserved (==0): 0
    Max.AREAs: (0==3): 0
  ISO 10589 ISIS Link State Protocol Data Unit
      PDU length: 97
      Remaining lifetime: 1199
      LSP-ID: 0001.0001.0001.00-00
      Sequence number: 0x0000000b
    + Checksum: 0xec74 [correct]
    + Type block(0x03): Partition Repair:0, Attached bits:0, Overload bit:0, IS type:3
    + Area address(es) (4)
    + Protocols supported (1)
    - Hostname (2)
        Hostname: R1
    + IS Reachability (23)
    + IP Interface address(es) (4)
    + IP Internal reachability (24)
```

Hostname to LSP ID Conversion Can Be Disabled

IOS and NX-OS

- `no hostname dynamic`

IOS XR

- `hostname dynamic disable`

IOS

```
R1#show isis neighbors
Tag CISCOLIVE:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
0002.0002.0002  L1   Gi0/2         10.12.1.2       UP    22       R1.02
0002.0002.0002  L2   Gi0/2         10.12.1.2       UP    26       R1.02
0003.0003.0003  L1   Gi0/3         10.13.1.3       UP    26       R1.01
0003.0003.0003  L2   Gi0/3         10.13.1.3       UP    27       R1.01
```

Viewing IS-IS Neighbors

IOS

```
R1#show isis neighbors
Tag CISCOLIVE:
System Id       Type Interface     IP Address      State Holdtime Circuit Id
XR2             L1   Gi0/2         10.12.1.2       UP    22       R1.02
XR2             L2   Gi0/2         10.12.1.2       UP    26       R1.02
NX-3            L1   Gi0/3         10.13.1.3       UP    26       R1.01
NX-3            L2   Gi0/3         10.13.1.3       UP    27       R1.01
```

IOS XR

```
RP/0/0/CPU0:XR2#show isis neighbors
IS-IS CISCOLIVE neighbors:
System Id      Interface        SNPA           State Holdtime Type IETF-NSF
R1             Gi0/0/0/1        fa16.3eac.7a9b Up    9        L1L2 Capable
NX-3           Gi0/0/0/3        fa16.3e00.0002 Up    21       L1L2 Capable
```

Viewing IS-IS Neighbors

NX-OS

```
NX-3# show isis adjacency
IS-IS process: CISCOLIVE VRF: default
IS-IS adjacency database:
Legend: '!': No AF level connectivity in given topology
System ID       SNPA            Level  State  Hold Time  Interface
R1              fa16.3e69.d5fc  1      UP     00:00:10   Ethernet2/1
R1              fa16.3e69.d5fc  2      UP     00:00:10   Ethernet2/1
XR2             fa16.3e1f.787e  1      UP     00:00:08   Ethernet2/2
XR2             fa16.3e1f.787e  2      UP     00:00:07   Ethernet2/2
```

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
  - L1 adjacencies require the area address to matches
  - The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

Troubleshooting Adjacency:

Viewing IS-IS Interface State

IOS

```
R1#show clns interface
GigabitEthernet0/1 is up, line protocol is up
  CLNS protocol processing disabled
GigabitEthernet0/2 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  ..
  Next ESH/ISH in 13 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x1, local circuit ID 0x2
    Level-1 Metric: 10, Priority: 64, Circuit ID: R1.02
    DR ID: R1.02
    Level-1 IPv6 Metric: 10
    Number of active level-1 adjacencies: 1
```

IS-IS is not configured or the interface is passive

This repeats for L2 info

Troubleshooting Adjacency:

Viewing IS-IS Interface State

IOS XR

```
RP/0/0/CPU0:XR2#show isis interface
IS-IS CISCOLIVE Interfaces
GigabitEthernet0/0/0/0       Enabled
  Adjacency Formation:       Disabled (Passive in IS-IS cfg)
  Prefix Advertisement:      Enabled
  ..
  Circuit Type:              level-1-2
  Media Type:                LAN
  Circuit Number:            0
  IPv4 Unicast Topology:     Enabled
    Adjacency Formation:     Disabled (Intf passive in IS-IS cfg)
    Prefix Advertisement:    Running
    Metric (L1/L2):          0/0
    Weight (L1/L2):          0/0
```

IS-IS is configured and the interface is passive

Troubleshooting Adjacency:

Viewing IS-IS Interface State

NX-OS

```
NX-3# show isis interface
IS-IS process: CISCOLIVE VRF: default
Ethernet2/1, Interface status: protocol-up/link-up/a          ap
  IP address: 10.13.1.3, IP subnet: 10.13.1.0/2
  Index: 0x0001, Local Circuit ID: 0x01, C    uit Type: L1-2
  ..
  Passive level: level-1-2
  LSP interval: 33 ms, MTU: 1500
  Level    Metric-0    Metric-2    CSNP  Next CSNP  Hello    Multi    Next IIH
  1              40           0      10  Inactive       10    3        Inactive
  2              40           0      10  Inactive       10    3        Inactive
  Level  Adjs    AdjsUp Pri  Circuit ID          Since
  1          0         0  64  0000.0000.0000.00   00:04:30
  2          0         0  64  0000.0000.0000.00   00:04:29
```

IS-IS interface is passive

Troubleshooting Adjacency:

Looking for Passive Interface Configurations on IOS

```
R1#show run | s router isis
 ip router isis CISCOLIVE
router isis CISCOLIVE
 net 49.0123.0001.0001.0001.00
 log-adjacency-changes
 passive-interface GigabitEthernet0/1
```

Or `passive-interface default`

Be careful pulling off the `passive-interface`

```
R1#show run int gi0/1
Building configuration...
interface GigabitEthernet0/1
 description to NX-3
 ip address 10.13.1.1 255.255.255.0
end
```

What happened to the IS-IS configuration?

Troubleshooting Adjacency:

Looking for Passive Interface Configurations on IOS XR and NX-OS

```
RP/0/0/CPU0:XR2#show run router isis
router isis CISCOLIVE
 net 49.0123.0002.0002.0002.00
 log adjacency changes
 interface GigabitEthernet0/0/0/0
  passive
  address-family ipv4 unicast
```

```
NX-3# show run isis
router isis CISCOLIVE
..
interface Ethernet2/1
  ip router isis CISCOLIVE
  isis passive-interface level-1-2
```

Or `passive-interface default`

Checking for IIH Packets

Always three sides to the story

- Your side
- My side
- On the wire

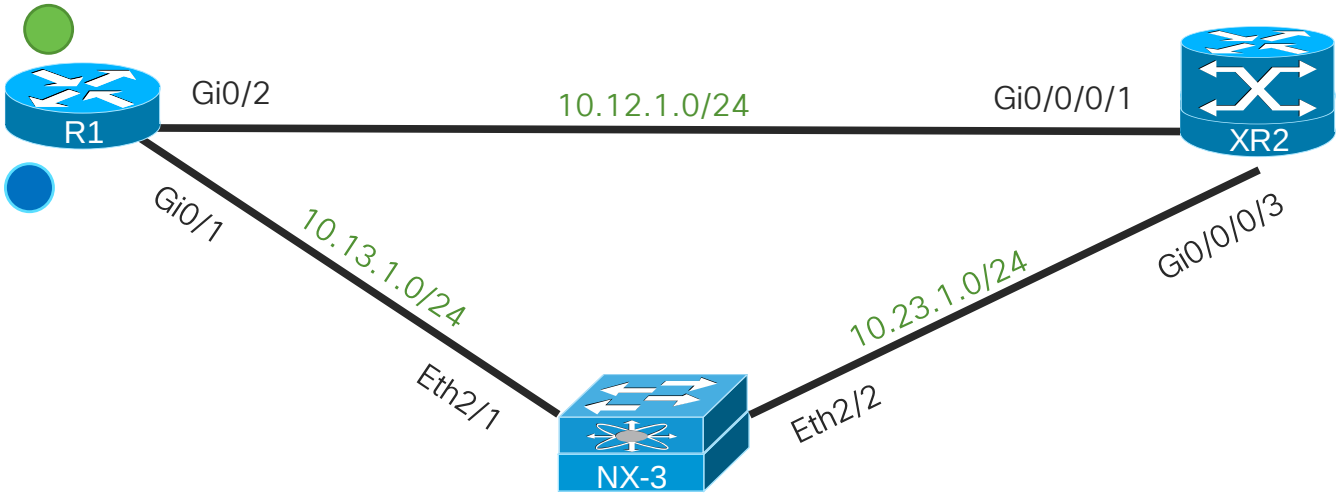

Checking for IIH Packets

Tools that we can use:

- Wireshark (NX-OS)
- Embedded Packet Captures (IOS, IOS XE)
- SPAN sessions
- Traces (IOS XR)
- Event-History (NX-OS)
- Debugs (IOS, IOS XR, NX-OS) – Hit to your CPU ☹
- Checking CoPP for Drops

ACLs cannot be used to check IS-IS traffic! Its based on L2.

Using Wireshark (NX-OS)

Capture traffic to Nexus Switches (Not through it)

```
NX-1# ethanalyzer local interface inband capture-filter "ether host 01:80:c2:00:00:15"
Capturing on inband
09:08:42.979127 88:5a:92:de:61:7c -> 01:80:c2:00:00:15 ISIS L2 HELLO,  System-ID: 0000.0000.0001
09:08:46.055807 88:5a:92:de:61:7c -> 01:80:c2:00:00:15 ISIS L2 HELLO,  System-ID: 0000.0000.0001
09:08:47.489024 88:5a:92:de:61:7c -> 01:80:c2:00:00:15 ISIS L2 CSNP, Source-ID: 0000.0000.0001.00,
Start LSP-ID: 0000.0000.0000.00-00, End LSP-ID: ffff.ffff.ffff.ff-ff
09:08:48.570401 00:2a:10:03:f2:80 -> 01:80:c2:00:00:15 ISIS L2 HELLO, System-ID: 0000.0000.0002
09:08:49.215861 88:5a:92:de:61:7c -> 01:80:c2:00:00:15 ISIS L2 HELLO, System-ID: 0000.0000.0001
09:08:52.219001 88:5a:92:de:61:7c -> 01:80:c2:00:00:15 ISIS L2 HELLO, System-ID: 0000.0000.0001
```

Using Wireshark (NX-OS)

Viewing an explicit packet

```
NX-1# ethanalyzer local interface inband capture-filter "ether host 01:80:c2:00:00:15" detail
Capturing on inband
Frame 1 (1014 bytes on wire, 1014 bytes captured)
    Arrival Time: May 22, 2017 09:07:16.082561000
    [Time delta from previous captured frame: 0.000000000 seconds]
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.000000000 seconds]
    Frame Number: 1
    Frame Length: 1014 bytes
    Capture Length: 1014 bytes
    [Frame is marked: False]
    [Protocols in frame: eth:llc:osi:isis]
IEEE 802.3 Ethernet
    Destination: 01:80:c2:00:00:15 (01:80:c2:00:00:15)
        Address: 01:80:c2:00:00:15 (01:80:c2:00:00:15)
        .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
```

Embedded Packet Capture (IOS & IOS XE)

IOS

```
R1#monitor capture buffer PACKETS
R1#monitor capture buffer PACKETS limit duration 1200
R1#monitor capture buffer PACKETS size 10240
R1#monitor capture point ip process-switched PACKETS both
R1#monitor capture point associate PACKETS CoPP
R1#monitor capture point start PACKETS
R1#monitor capture buffer PACKETS export tftp://192.168.0.1/R41.pcap
```

IOS XE

```
R1#monitor capture PACKETS control-plane both match any limit duration 600 buffer size 10
R1#monitor capture PACKETS start
R1#monitor capture PACKETS export tftp://192.168.0.1/R1-PACKETS.pcap
```

Traces (IOS XR)

Built-In, Always Running Debugs without the Performance Hit

```
RP/0/0/CPU0:XR2#show isis trace all reverse

--- Trace data for instance CISCOLIVE --

72678 wrapping entries (89344 possible, 88576 allocated, 0 filtered, 995940 total)

19:44:28.484 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_SCHED_LAN_SCHEDULE L1 ifh: 0x60

19:44:28.484 isis/CISCOLIVE/hlo 0/0/CPU0 t7  IO_PDU_OUTPUT L1 LAN IIH Gi0/0/0/1 0180.c200.0014

19:44:28.484 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_SEND_LAN_DETAILS L1 Gi0/0/0/1

19:44:28.484 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_SEND_LAN_ADD_NEIGHBOR L1 Gi0/0/0/1
0001.0001.0001 fa16.3e8f.e522

19:44:28.314 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_SCHED_LAN_SCHEDULE L2 Gi0/0/0/3

19:44:28.314 isis/CISCOLIVE/hlo 0/0/CPU0 t7  IO_PDU_OUTPUT L2 LAN IIH Gi0/0/0/3 0180.c200.0015

19:44:27.624 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_RECV_LAN_STATE_INIT L1 Gi0/0/0/1
fa16.3e8f.e522 L1 LAN IIH

19:44:27.624 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_RECV_BAD_IF_ADDRESS L1 Gi0/0/0/1
fa16.3e8f.e522 L1 LAN IIH IPv4

19:44:27.624 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_RECV_LAN_DETAILS L1 Gi0/0/0/1 fa16.3e8f.e522
L1 LAN IIH
```

IS-IS IIH Debugs

Built-In diagnostic tools

```
R1#debug isis adj-packets

IS-IS Adjacency related packets debugging is on for router process CISCOLIVE

01:04:18.503: ISIS-Adj: Sending L2 LAN IIH on GigabitEthernet0/3, length 1497

01:04:18.766: ISIS-Adj: Sending L1 LAN IIH on GigabitEthernet0/2, length 1497

01:04:19.695: ISIS-Adj: Rec L1 IIH from fa16.3e05.7eb4 (GigabitEthernet0/2), cir type L1L2,
cir id 0001.0001.0001.01, length 1497, ht(30)

01:04:21.066: ISIS-Adj: Rec L1 IIH from fa16.3e00.0001 (GigabitEthernet0/3), cir type L1, cir
id 0001.0001.0001.02, length 1497, ht(30)
```

```
NX-3# debug isis iih

01:11:30.305488 isis: CISCOLIVE [8723] Receive L2 LAN IIH over Ethernet2/1 from R1
(fa16.3e6b.c4fd) len 1497 prio 0 = FALSE

01:11:30.366816 isis: CISCOLIVE [8723] Receive L1 LAN IIH over Ethernet2/2 from XR2
(fa16.3e56.c1a5) len 1497 prio 0

01:11:30.366932 isis: CISCOLIVE [8723] Send L2 LAN IIH over Ethernet2/2 len 1497 prio 6,dmac
0180.c200.0015
```

Back to Troubleshooting IS-IS Adjacencies

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

Troubleshooting Adjacency:

IS-IS interface not on a common subnet

IOS XR

```
RP/0/0/CPU0: 14:55:29.161 : isis[1010]: %ROUTING-ISIS-6-IIH_IF_ADDRESS : IIH
received from GigabitEthernet0/0/0/1 SNPA fa16.3e47.0695 contains unusable IPv4
interface address: 10.1.12.1 not on same subnet as local interface
```

NX-OS

```
NX-3# show isis event-history iih

ISIS CISCOLIVE process

 iih Events for ISIS process

02:41:01.430468 isis CISCOLIVE [8723]: [8724]: Send L2 LAN IIH over Ethernet2/1
len 1497 prio 6,dmac 0180.c200.0015

02:41:01.430448 isis CISCOLIVE [8723]: [8724]: no adj ? 4

02:40:57.072088 isis CISCOLIVE [8723]: [8728]: isis_iih_find_ipv4_addr: Unable
to find IPv4 address for Ethernet2/1

02:40:57.072068 isis CISCOLIVE [8723]: [8728]: Fail to find usable IPv4 address
```

IOS requires checking interface addresses on all desired neighbors

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

Troubleshooting Adjacency:

Circuit types do not match

Point-2-Point links use the same concept as broadcast interfaces except:

- Three-way handshake is used versus a 2 way handshake used by broadcast interfaces. If a device does not support TLV240; a router can establish an adjacency using a 2-way handshake
- Each router sends a PSNP to acknowledge the receipt of an LSP.

Troubleshooting Adjacency:

Circuit types do not match

IOS

```
R1#debug isis adj-packets

IS-IS Adjacency related packets debugging is on for router process CISCOLIVE

02:49:24.920: ISIS-Adj: Rec L1 IIH from fa16.3e05.7eb4 (GigabitEthernet0/2), cir
type L1L2, cir id 0002.0002.0002.01, length 1497, ht(30)

02:49:24.920: ISIS-Adj: Multi-point IIH received on point-to-point interface:
ignored IIH
```

IOS XR

```
RP/0/0/CPU0:02:48:12.892 : isis[1010]: %ROUTING-ISIS-7-ERR_RCV_PAKTYPE : Invalid
IS-IS packet type 17 received on GigabitEthernet0/0/0/1 SNPA fa16.3e8f.e522
(inappropriate code)
```

NX-OS requires checking interface configs on both sides

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

IS-IS Adjacencies with Hierarchy Levels

- Level 1-only routers
  - Can only form adjacencies with Level 1 or Level-1-2 routers in the same area
- Level-1-2 (default mode for Cisco devices)
  - Can form L1 adjacencies with Level 1 and Level-1-2 routers in the same area
  - Can form L2 adjacencies with Level-1-2 and Level 2 routers in the same or different area
- Level-2-only routers
  - Can form L2 adjacencies with other Level-1-2 and Level 2 routers in the same or different area

No error messages.. Just missing IIH packets

IS-IS Router Type with Interface Type

	Router Set IS-IS Type L1	Router Set IS-IS Type L2	Router Set IS-IS Type L1-L2
Circuit Type L1	Level-1	No Adjacency Feasible	Level-1
Circuit Type L2	No Adjacency Feasible	Level-2	Level-2
Circuit Type L1-L2	Level-1	Level-2	Level-1 and Level-2

IS-IS Adjacencies with Hierarchy Levels

Checking Interface IS-Setting

IOS

```
R1#show clns interface gi0/2 | i Type
    Circuit Type: level-1-2
```

IOS XR

```
RP/0/0/CPU0:XR2#show isis interface gi0/0/0/3 | i Circuit Type
Circuit Type:              level-1-2
```

NX-OS

```
NX-3# show isis interface e2/1 | i Type
  Index: 0x0001, Local Circuit ID: 0x01, Circuit Type: L1-2
```

IS-IS Adjacencies with Hierarchy Levels

Checking Interface IS-Setting

Router Level Commands

IOS, IOS XR

```
is-type {level-1 | level-1-2 | level-2-only}
```

NX-OS

```
is-type {level-1 | level-1-2 | level-2}
```

Interface Level Commands

IOS

```
is-type {level-1 | level-1-2 | level-2-only}
```

IOS XR

```
circuit-type {level-1 | level-1-2 | level-2-only}
```

NX-OS

```
Isis circuit-type {level-1 | level-1-2 | level-2}
```

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

Checking for L1 Area Mismatch (IOS)

```
R1#debug isis adj-packets
IS-IS Adjacency related packets debugging is on for router process CISCOLIVE
05:00:22.958: ISIS-Adj: Rec L1 IIH from fa16.3e05.7eb4 (GigabitEthernet0/2), cir
type L1, cir id 0002.0002.0002.01, length 1497, ht(30)
05:00:22.958: ISIS-Adj: Area mismatch, level 1 IIH on GigabitEthernet0/2
```

```
R1#show isis protocol
Tag CISCOLIVE:
IS-IS Router: CISCOLIVE
  System Id: 0001.0001.0001.00  IS-Type: level-1
  Manual area address(es):
      49.0123
  Routing for area address(es):
      49.0123
```

Checking for L1 Area Mismatch (IOS XR)

```
RP/0/0/CPU0:XR2#show isis trace all reverse
--- Trace data for instance CISCOLIVE --
04:59:01.715 isis/CISCOLIVE/hlo 0/0/CPU0 t7  ADJ_SEND_LAN_DETAILS L1 Gi0/0/0/3
04:58:56.285 isis/CISCOLIVE/hlo 0/0/CPU0 t7 ADJ_RECV_LAN_STATE_DOWN_AREA_MISMATCH
L1 Gi0/0/0/1 fa16.3e8f.e522 L1 LAN IIH
```

```
RP/0/0/CPU0:XR2#show isis protocol
IS-IS Router: CISCOLIVE
  System Id: 0002.0002.0002
  Instance Id: 0
  IS Levels: level-1
  Manual area address(es):
    49.0002
  Routing for area address(es):
    49.0002
```

Checking for L1 Area Mismatch (NX-OS)

```
NX-3# show isis event-history iih

ISIS CISCOLIVE process

 iih Events for ISIS process

05:08:31.928651 isis CISCOLIVE [8723]: [8728]: Receive L1 LAN IIH over Ethernet2/2
from 0002.0002.0002 (fa16.3e56.c1a5) len 1497 prio 0

05:08:30.560626 isis CISCOLIVE [8723]: [8724]: Send L1 LAN IIH over Ethernet2/2
len 1497 prio 6,dmac 0180.c200.0014

05:08:24.358015 isis CISCOLIVE [8723]: [8728]: No common area
```

```
NX-3# show isis protocol | section Area

  Area address(es) :

    49.0123
```

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

Error messages from Duplicate System-ID

IOS

```
R1#
05:14:40.652: %CLNS-3-BADPACKET: ISIS: LAN L1 hello, Duplicate system ID
detected from fa16.3e05.7eb4 (GigabitEthernet0/2)
```

IOS XR

```
RP/0/0/CPU0:May 11 05:16:04.125 : isis[1010]: %ROUTING-ISIS-6-ERR_DUPID :
Duplicate System ID 0001.0001.0001 already used by Local System detected in IIH
received on GigabitEthernet0/0/0/1 SNPA fa16.3e8f.e522
```

NX-OS

```
NX-3#
05:21:59 NX-3 %ISIS-4-LAN_DUP_SYSID:  isis-CISCOLIVE [8723]  L1 LAN IIH -
Duplicate system ID 0001.0001.0001 detected over Ethernet2/1 from fa16.3e6b.c4fd
```

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

What Settings are required for an IS-IS Adjacency?

- IS-IS Interface is Active
- IS-IS Interface shares a common subnet
- Protocols match (IPv4, IPv6) per topology
- Circuit types match (General or Point to Point)
- Router levels are compatible
- L1 adjacencies require the area address to matches
- The system ID must be unique within the same area address
- IS-IS Hello Authentication matches
- MTU Matches

MTU Check During Neighbor Adjacency

- Padding the IIHs provides a mechanism for detect large frames or mismatched MTU.
- IS-IS hellos (IIHs) are padded with TLV #8 to reach the maximum transmission unit (MTU) size of the network interface.
- By default, the IS-IS MTU must match
- Neighbor with the higher MTU will show the neighbor with lower neighbor in 'INIT' state. The neighbor with lower will not show a neighbor

```
NX-3# show isis adjacency
IS-IS process: CISCOLIVE VRF: default
IS-IS adjacency database:
System ID       SNPA            Level  State  Hold Time  Interface
0001.0001.0001  fa16.3e6b.c4fd  1      INIT   00:00:27   Ethernet2/1
XR2             fa16.3e56.c1a5  1      UP     00:00:09   Ethernet2/2
```

MTU Mismatch Detection

- NX-OS can view the event viewer and check MTU on packets

```
NX-3# show isis event-history iih

ISIS CISCOLIVE process

 iih Events for ISIS process

18:14:31.856565 isis CISCOLIVE [8723]: [8728]:  Neighbor TLV missing in hello
from 0001.0001.0001 , hence adjacency in INIT state

18:14:31.856546 isis CISCOLIVE [8723]: [8728]: Fail to find iih nbr tlv

18:14:31.856311 isis CISCOLIVE [8723]: [8728]: Receive L1 LAN IIH over
Ethernet2/1 from 0001.0001.0001 (fa16.3e6b.c4fd) len 1297 prio 0

18:14:30.921832 isis CISCOLIVE [8723]: [8728]: Fail to find iih nbr tlv

18:14:30.200596 isis CISCOLIVE [8723]: [8724]: Send L1 LAN IIH over Ethernet2/1
len 1497 prio 6,dmac 0180.c200.0014
```

- IOS-XR and IOS devices can check MTU by viewing IS-IS Interface

MTU Mismatch Remediation

- Correct the MTU on the interfaces to be the same
- Disable MTU padding
  - IOS – **`no isis hello padding [always]`**
  - IOS XR – **`hello-padding {disable|sometimes} [level {1|2}]`**
  - NX-OS – **`no isis hello padding [always]`**

Trivia Question Hint

- How do you make the traffic between R1 and R4, take R5?
- How do you make the traffic between R2 and R3, take the direct link?

One command on R1

IS-IS Topologies

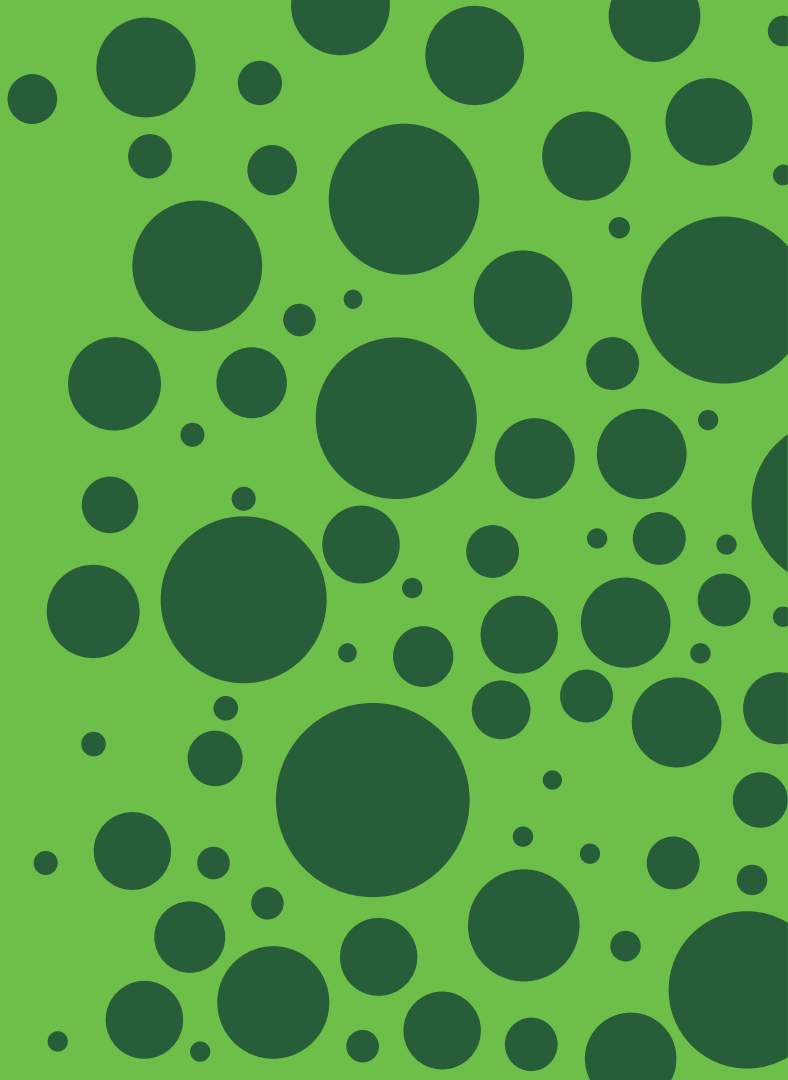

IS-IS Topologies

IS-IS maintains a copy of all the LSPs in a database for an Level

An LSP database per Level can be thought of as a topology.

How many topologies do you see?

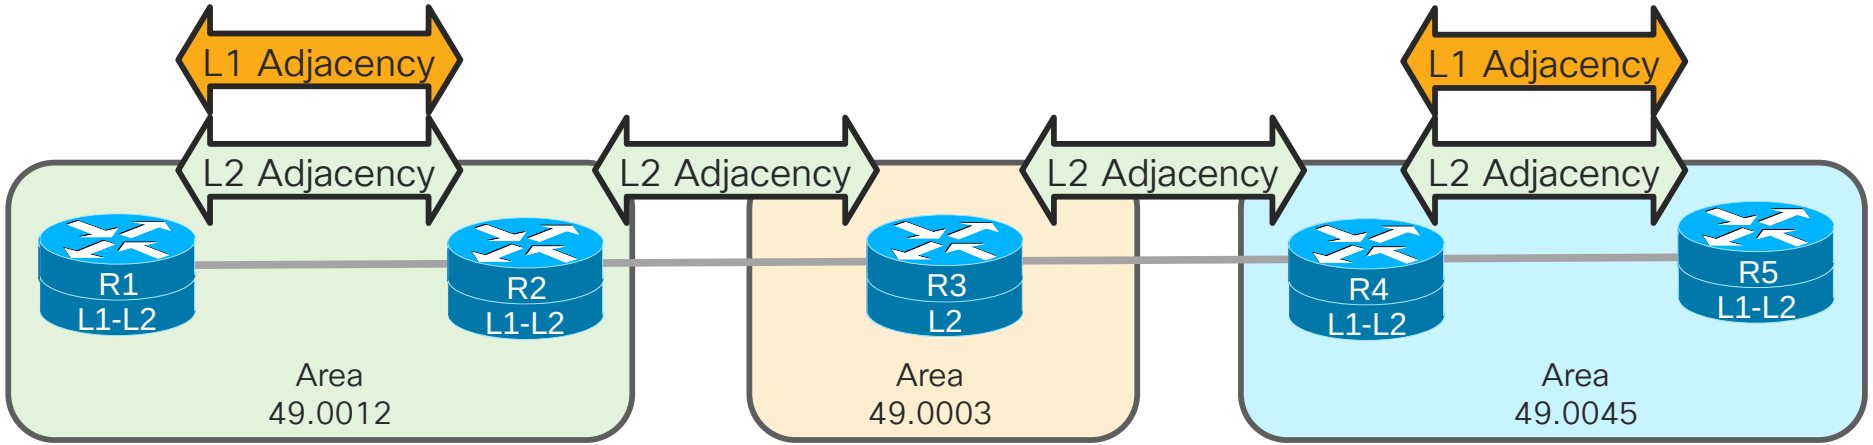

IS-IS Topologies

IS-IS maintains a copy of all the LSPs in a database for an Level

An LSP database per Level can be thought of as a topology.

How many topologies do you see?

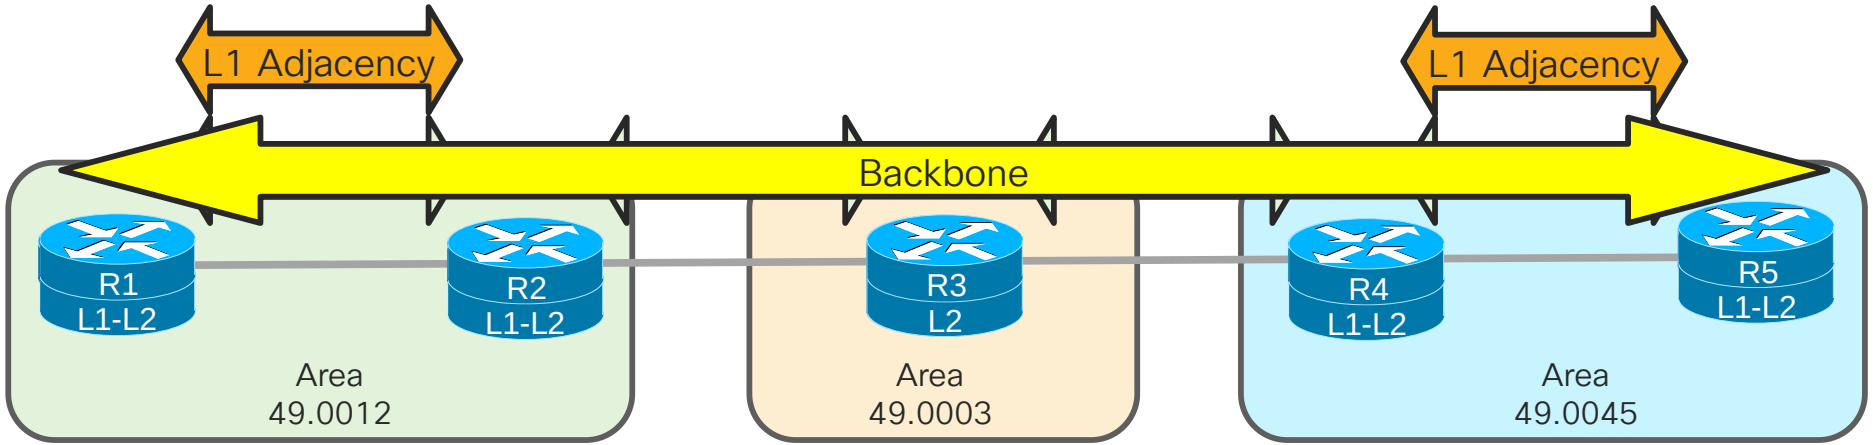

IS-IS Topologies

IS-IS maintains a copy of all the LSPs in a database for an Level

An LSP database per Level can be thought of as a topology.

How many topologies do you see?

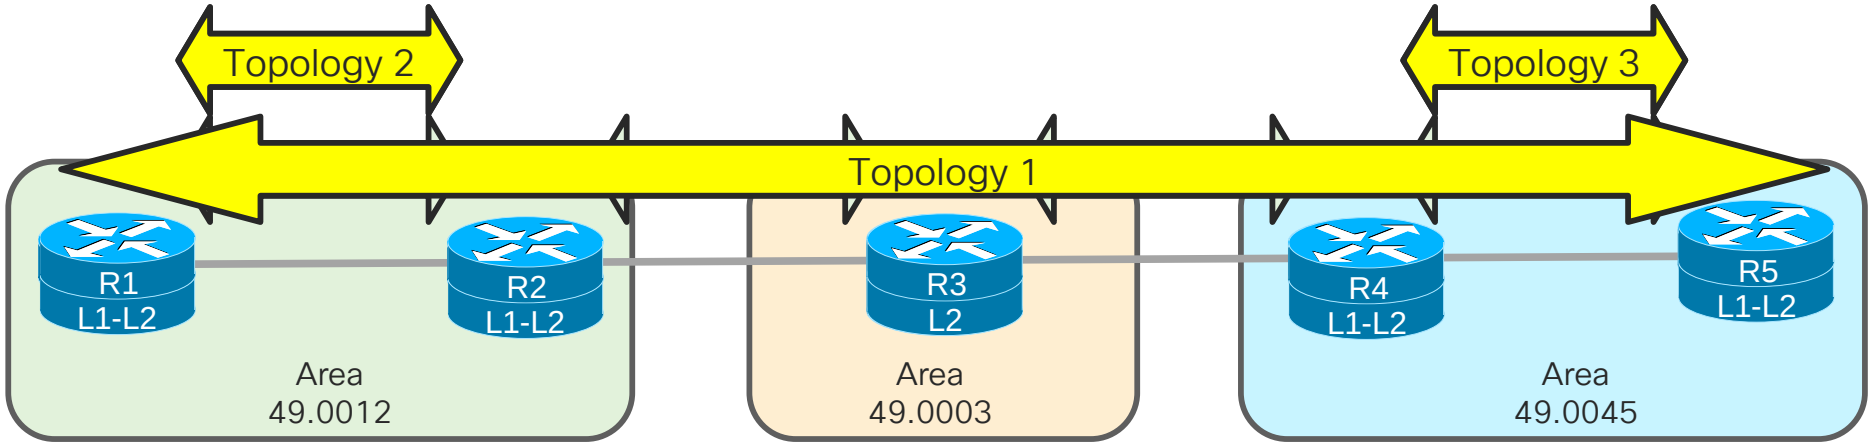

Viewing an IS-IS Topology (IOS)

```
R1#show isis topology

Tag CISCOLIVE:
IS-IS TID 0 paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2                   Gi0/2       fa16.3ed4.04f5

IS-IS TID 0 paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2                   Gi0/2       fa16.3ed4.04f5
R3                   20         R2                   Gi0/2       fa16.3ed4.04f5
R4                   30         R2                   Gi0/2       fa16.3ed4.04f5
R5                   40         R2                   Gi0/2       fa16.3ed4.04f5
```

Viewing an IS-IS Topology (IOS XR)

```
RP/0/0/CPU0:XR1#show isis topology

IS-IS CISCOLIVE paths to IPv4 Unicast (Level-1) routers
System Id       Metric  Next-Hop        Interface       SNPA
XR1             --
XR2             10      XR2             Gi0/0/0/2       *PtoP*

IS-IS CISCOLIVE paths to IPv4 Unicast (Level-2) routers
System Id       Metric  Next-Hop        Interface       SNPA
XR1             --
XR2             10      XR2             Gi0/0/0/2       *PtoP*
XR3             20      XR2             Gi0/0/0/2       *PtoP*
R4-XR           30      XR2             Gi0/0/0/2       *PtoP*
R5-XR           40      XR2             Gi0/0/0/2       *PtoP*
```

Viewing an IS-IS Topology (NX-OS)

```
NX-1# show isis topology
IS-IS process: CISCOLIVE
VRF: default
IS-IS Level-1 IS routing table
NX-2.00, Instance 0x00000006
   *via NX-2, Ethernet2/2, metric 40

IS-IS Level-2 IS routing table
NX-2.00, Instance 0x00000009
   *via NX-2, Ethernet2/2, metric 40
NX-3.00, Instance 0x00000009
   *via NX-2, Ethernet2/2, metric 80
R4-NX.00, Instance 0x00000009
   *via NX-2, Ethernet2/2, metric 120
R4-NX.01, Instance 0x00000009
   *via NX-2, Ethernet2/2, metric 120
R5-NX.00, Instance 0x00000009
   *via NX-2, Ethernet2/2, metric 130
```

Optimizing an Area

So currently R1 can reach R5 using just the L2 Backbone

What happens when the backbone shrinks between to R2 to R4?

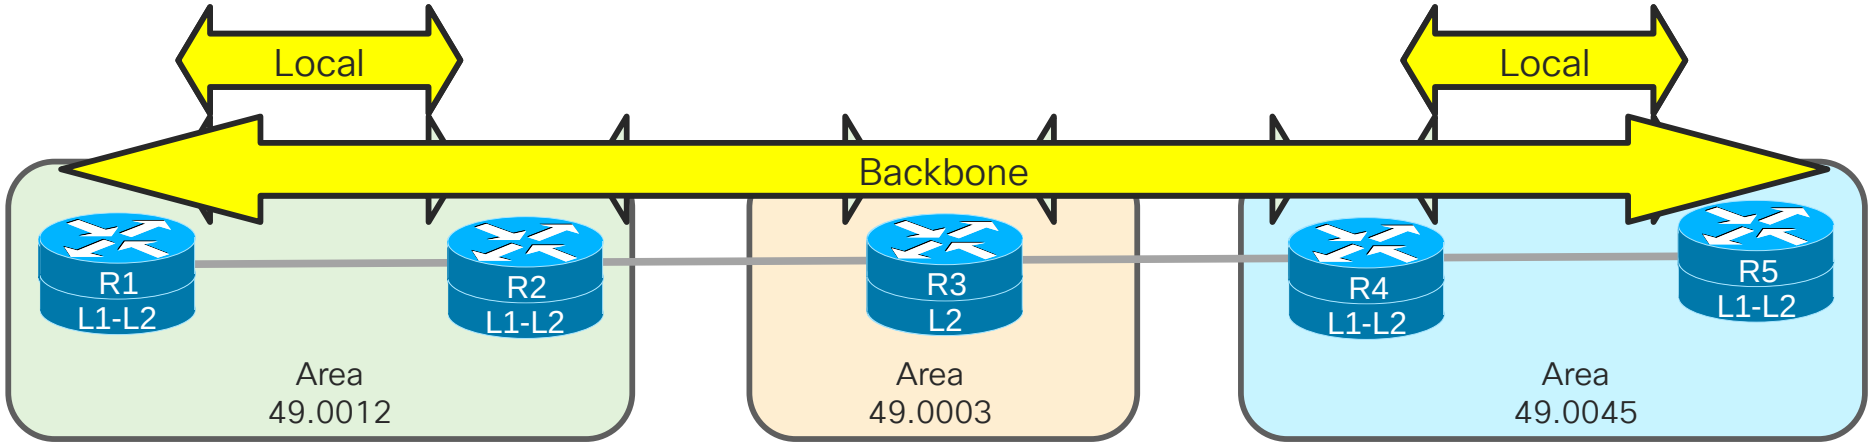

Understanding Route Advertisement

- IS-IS builds the topology using TLV#2.
- Routes are built using TLV#128 and/or TLV#135
- All routers in the same L1/L2 area have the same LSPDBs for that Level

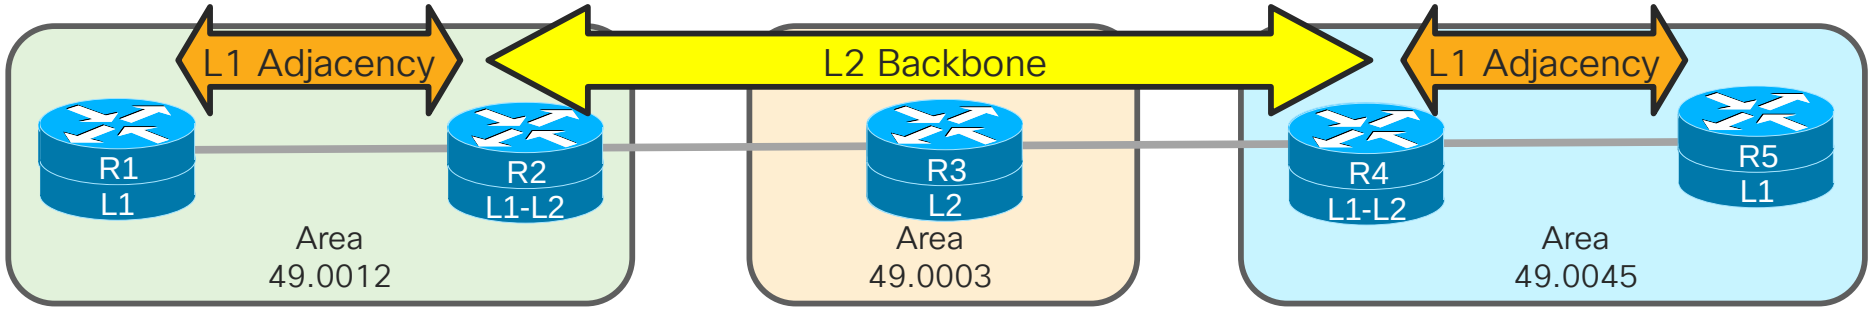

Viewing an IS-IS Topology (IOS)

```
R1#show isis topology
Tag CISCOLIVE:
IS-IS TID 0 paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2                   Gi0/2       fa16.3ed4.04f5
```

```
R2#show isis topology
Tag CISCOLIVE:
IS-IS TID 0 paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   10         R1                   Gi0/1       fa16.3e5c.91c1
R2                   --
IS-IS TID 0 paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R2                   --
R3                   10         R3                   Gi0/3       fa16.3e94.673d
R4                   20         R3                   Gi0/3       fa16.3e94.673d
R5                   30         R3                   Gi0/3       fa16.3e94.673d
```

R1 is no longer present in L2 Topology/Database

Building of an IS-IS Topology

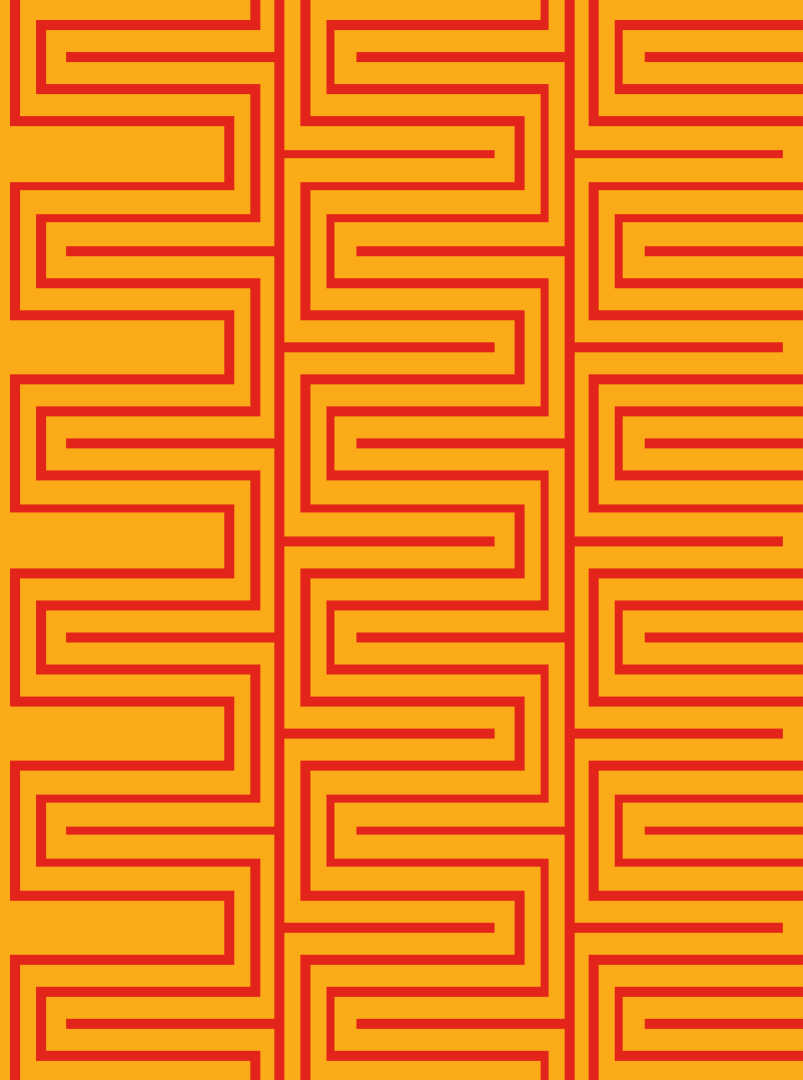

Understanding How the Topology is Built

- Topology is built off of TLV#2 (IS-Neighbors) and the LSP-ID
- LSP-IDs that end with 00 are those of routers themselves
  - Remember the SEL being set to 00 back from the NET addressing?
- LSP-IDs that DO NOT end with 00 are those of DIS (pseudonode)

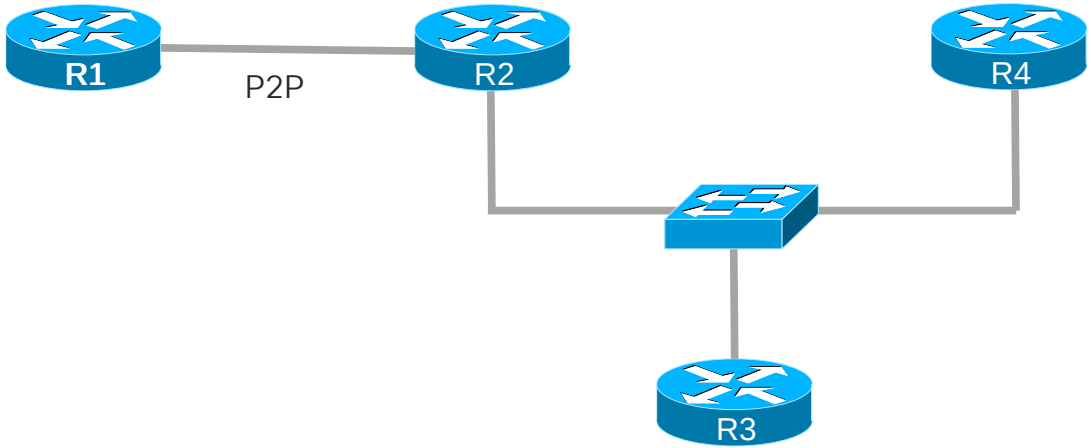

Understanding How the Topology is Built

`show isis database` [*LSP-ID*] `[level-1|level-2] [detail]`

- Displays all the LSPs from a specific router (or DIS), Sequence Number, Holdtime, Attribute fields (Partition, Attached/Overload/Router Type)

```
R1#show isis database
Tag CISCOLIVE:
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
R1.00-00            * 0x00000007   0x3E7A                757/*         0/0/0
R2.00-00              0x0000000A   0x40A6                576/1199      0/0/0
```

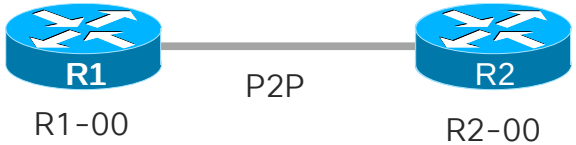

Understanding How the Topology is Built (P2P)

`show isis database` [*LSP-ID*] `[level-1|level-2] [detail]`

```
R1#show isis database detail
Tag CISCOLIVE:
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
R1.00-00            * 0x00000007   0x3E7A                335/*         0/0/0
  Area Address: 49.1234
  NLPID:        0xCC 0x8E
  Hostname: R1
  Metric: 10          IS R2.00
  IP Address:   192.168.1.1
  Metric: 10          IP 10.12.1.0 255.255.255.0
  Metric: 10          IP 10.1.1.0 255.255.255.0
R2.00-00              0x0000000C   0xF0BB                939/1199      0/0/0
..
```

R1-00 R2-00

Understanding How the Topology is Built (P2P)

```
R1#show isis database detail | exclude IP|PID|Area

Tag CISCOLIVE:
IS-IS Level-1 Link State Database:
R1.00-00              * 0x0000000A   0x7BEC                1038/*          0/0/0
  Hostname: R1
  Metric: 10          IS R2.00
R2.00-00                0x0000000C   0xF0BB                 791/1199       0/0/0
  Hostname: R2
  Metric: 10          IS R1.00
```

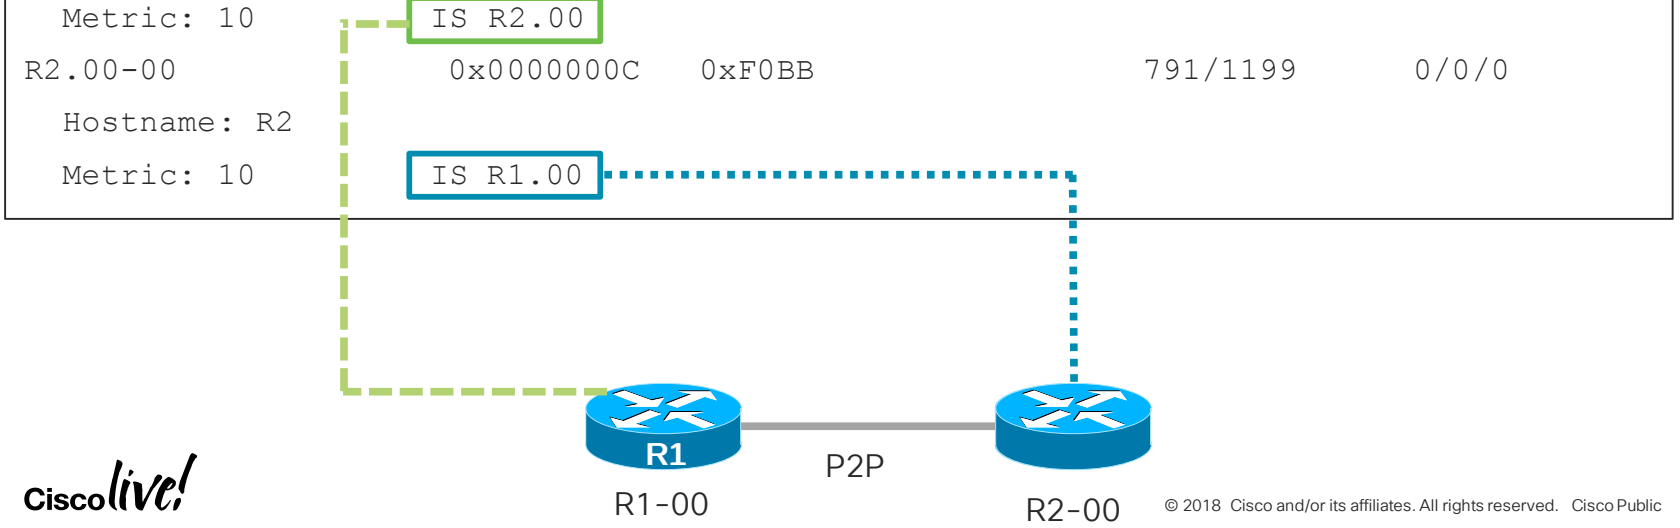

Understanding How the Topology is Built (Multi)

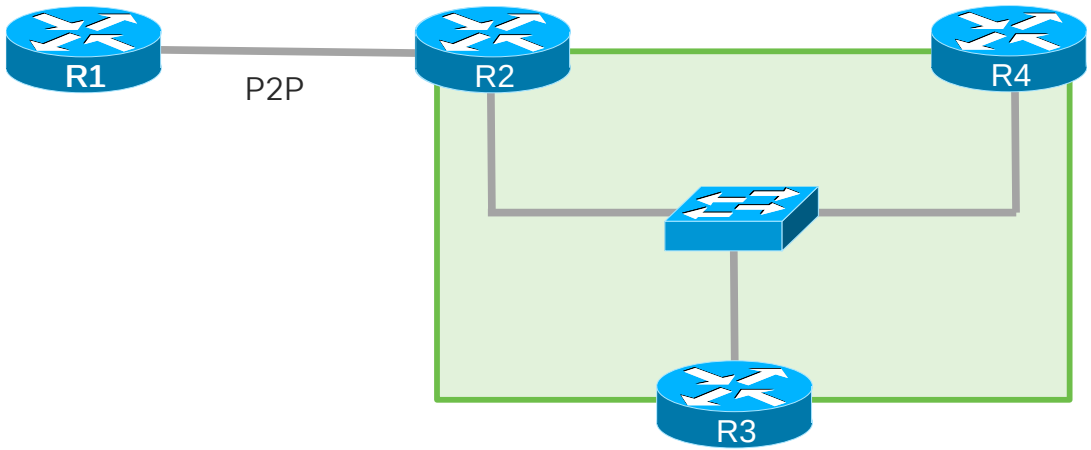

Understanding How the Topology is Built (Multi)

```
R2#show isis database detail | exclude IP|PID|Area
..
R2.00-00              0x0000017E   0x52E6              1130/1199      0/0/0
  Hostname: R2
  Metric: 10          IS R4.01
R3.00-00              0x00000173   0xE6AF              1121/1198      0/0/0
  Hostname: R3
  Metric: 10          IS R4.01
R4.00-00              0x0000017D   0x823F              1120/1198      0/0/0
  Hostname: R4
  Metric: 10          IS R4.01
R4.01-00              0x00000172   0xB040              1130/1198      0/0/0
  Metric: 0           IS R4.00
  Metric: 0           IS R2.00
  Metric: 0           IS R3.00
```

Understanding How the Topology is Built (Multi)

```
R2#show isis database detail | exclude IP|PID|Area
..
R2.00-00                0x0000017E
  Hostname: R2
  Metric: 10            IS R4.01
R3.00-00                0x00000173
  Hostname: R3
  Metric: 10            IS R4.01
R4.00-00                0x0000017D
  Hostname: R4
  Metric: 10            IS R4.01
R4.01-00                0x00000172
  Metric: 0             IS R4.00
  Metric: 0             IS R2.00
  Metric: 0             IS R3.00
```

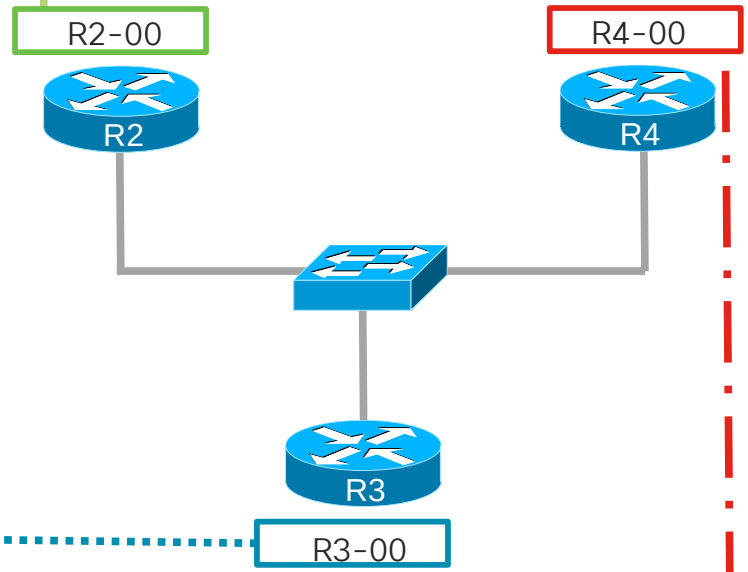

Understanding How the Topology is Built (Multi)

```
R2#show isis database detail | exclude IP|PID|Area
..
R2.00-00                0x0000017E
  Hostname: R2
  Metric: 10
R3.00-00
  Hostname: R3
  Metric: 10            IS R4.01
R4.00-00                0x0000017D
  Hostname:
  Metric: 10            IS R4.01
R4.01-00                0x00000172
  Metric: 0             IS R4.00
  Metric: 0             IS R2.00
  Metric: 0             IS R3.00
```

R4.01-00 must be the DIS

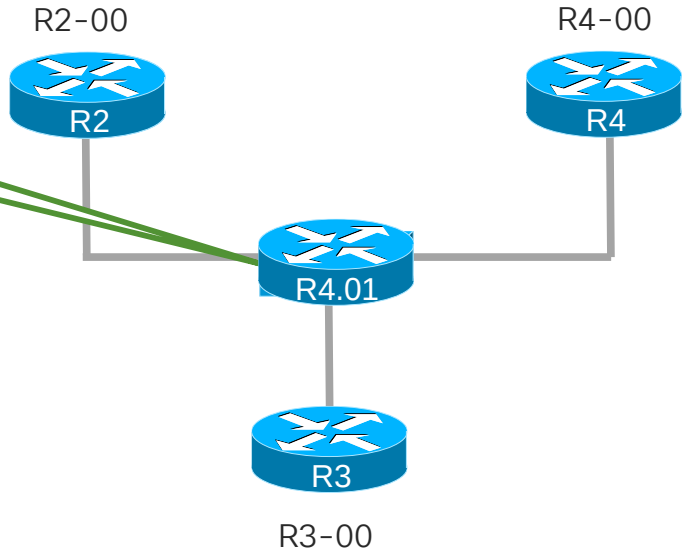

Understanding How the Topology is Built (Multi)

```
R2#show isis database detail | exclude IP|PID|Area
..
R2.00-00                0x0000017E
  Hostname: R2
  Metric: 10            IS R4.01
R3.00-00                0x00000173
  Hostname: R3
  Metric: 10            IS R4.01
R4.00-00                0x0000017D
  Hostname: R4
  Metric: 10            IS R4.01
R4.01-00                0x00000172
  Metric: 0             IS R4.00
  Metric: 0             IS R2.00
  Metric: 0             IS R3.00
```

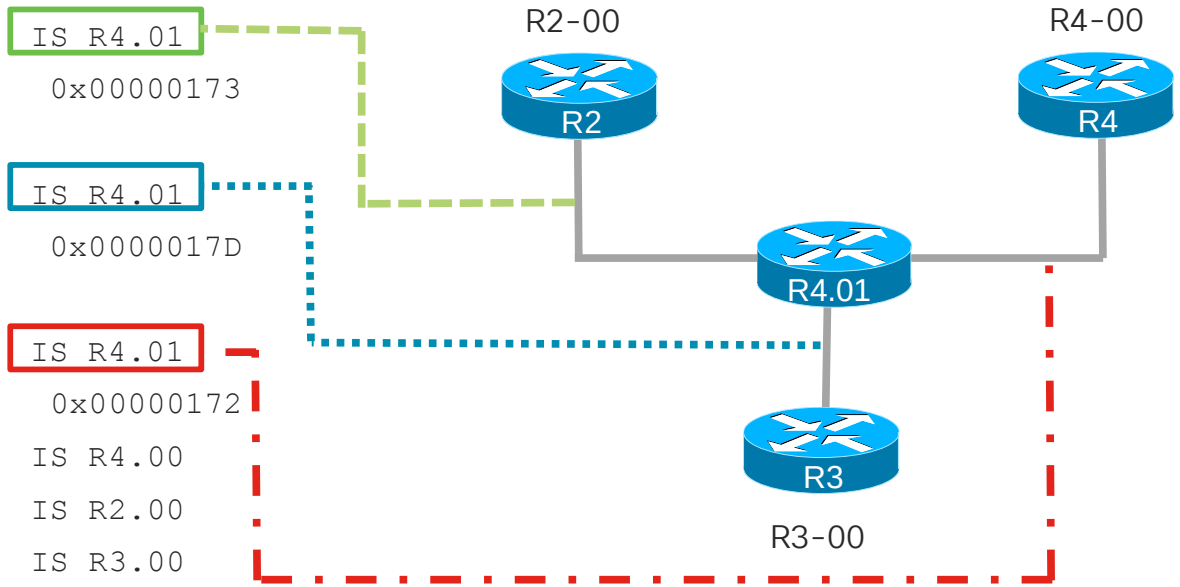

Understanding How the Topology is Built (Multi)

```
R2#show isis database detail | exclude IP|PID|Area
..
R2.00-00                0x0000017E
  Hostname: R2
  Metric: 10          IS R4.01
R3.00-00                0x00000173
  Hostname: R3
  Metric: 10          IS R4.01
R4.00-00                0x0000017D
  Hostname: R4
  Metric: 10          IS R4.01
R4.01-00                0x00000172
  Metric: 0           IS R4.00
  Metric: 0           IS R2.00
  Metric: 0           IS R3.00
```

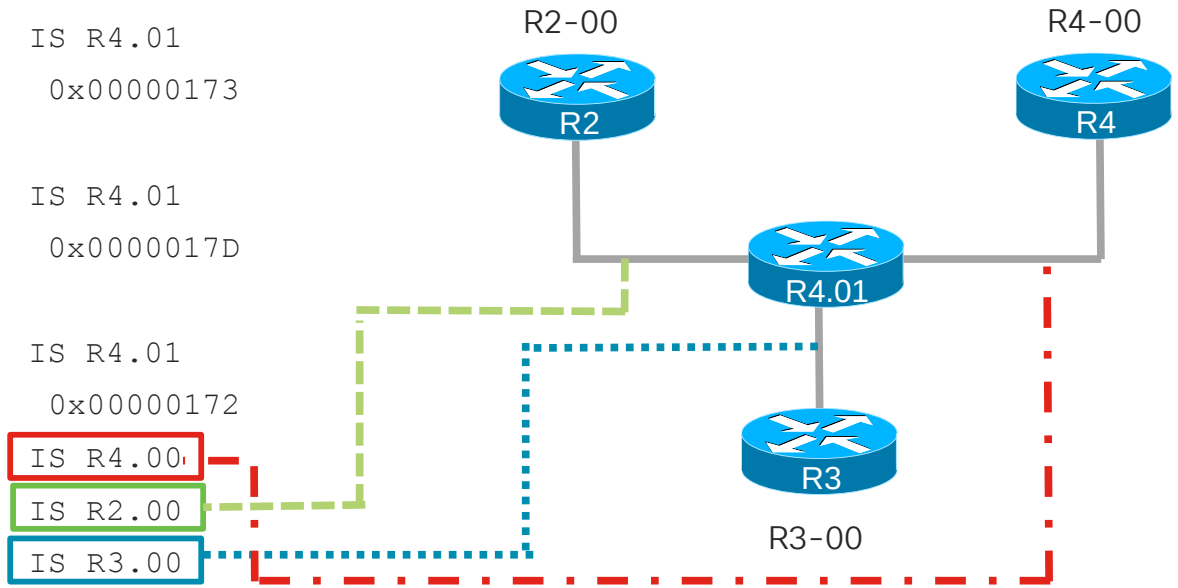

Metrics

Path Computation

After a router has built a topology of routers and their connecting interfaces, it runs a Shortest Path First Computation

The local router is the top of SPF Tree. All other routers are a branch. Calculations are made based off of interface cost

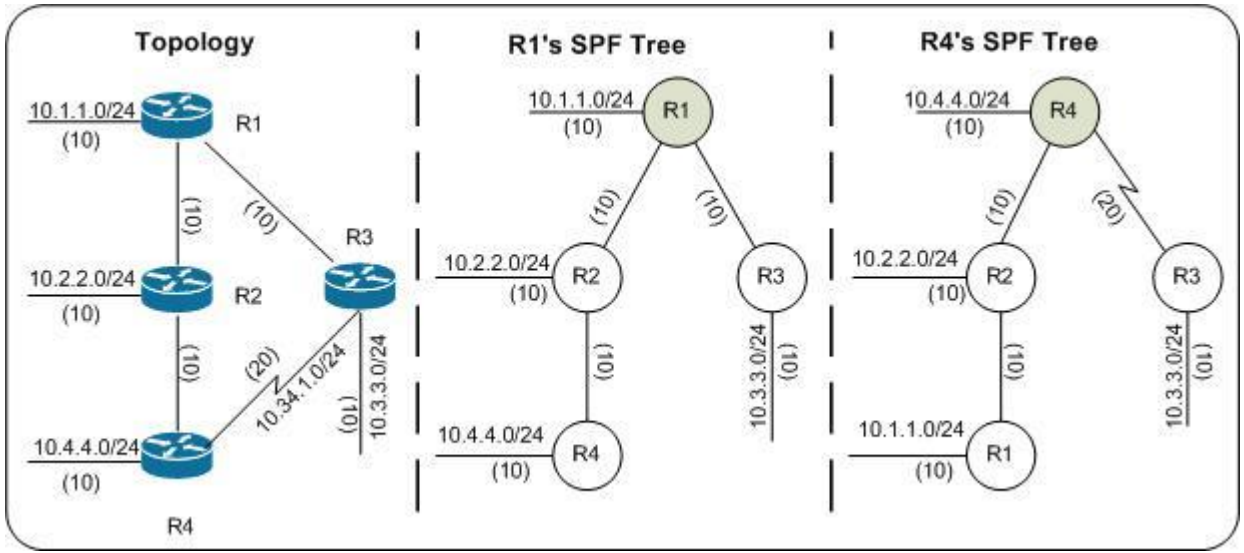

Interface Metrics

RFC 1195 provides a 6-bit field supporting values 1-63

- Stored in TLV# 128
- Industry standard is that all interfaces are statically set to 10 by default

```
R1#show isis database detail
..
R1.00-00              * 0x00000007   0x3E7A                 335/*         0/0/0
  Area Address: 49.1234
  NLPID:        0xCC 0x8E
  Hostname: R1
  Metric: 10         IS R2.00
```

Interface Metrics

RFC 1195 provides a 6-bit field supporting values 1-63

- Stored in TLV# 128
- Industry standard is that all interfaces are statically set to 10 by default
  - Value are changed as needed statically as needed
  - Except Nexus uses a reference bandwidth of 40 Gbps by default
    - 10-Gigabit Interfaces are set to 4
    - Gigabit Interfaces are set to 40

Interface Wide Metrics

Some network engineers thought that 6-bits is not enough to tune a network

- RFC 5305 introduced a new TLV# 135 that supported 32-bit values
- Allows for wide scale of metrics to reflect values from T1 interfaces to 100Gb interfaces
- Does not impact the way a topology is built, using TLV #2

What's Wrong?

```
R1#show isis topology level-2

IS-IS TID 0 paths to level-2 routers

System Id               Metric     Next-Hop              Interface    SNPA

R1                      --
R2                      10         R2                    Gi0/2        fa16.3ed4.04f5
R3                      **
R4                      **
R5                      **
```

IS-IS builds the topology on TLV #2.

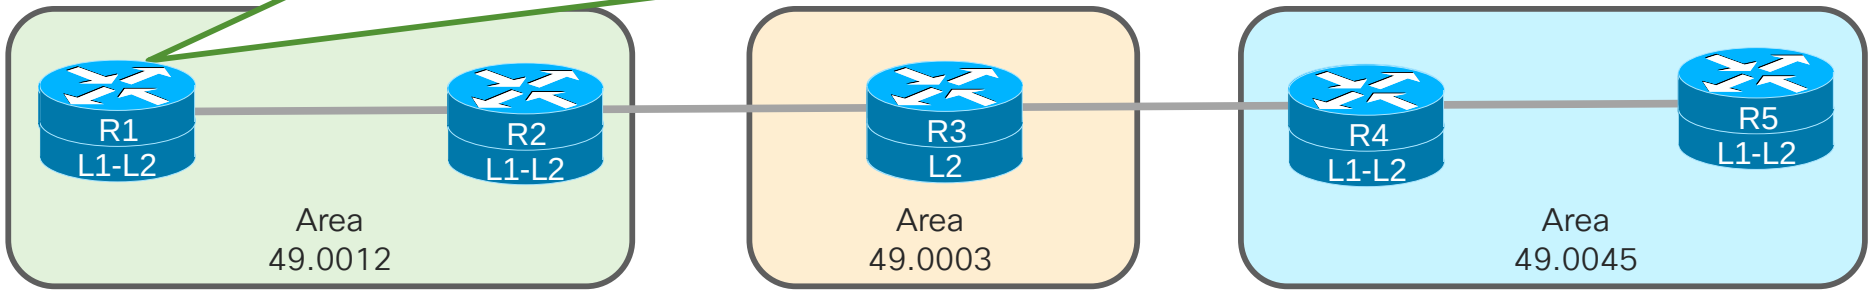

Mismatch Metric Types

- TLVs are transmitted as long as they are recognized
- When a router does not recognize a TLV it drops it.

IOS and IOS XR use Narrow Metrics by Default
NX-OS uses Wide Metrics by default

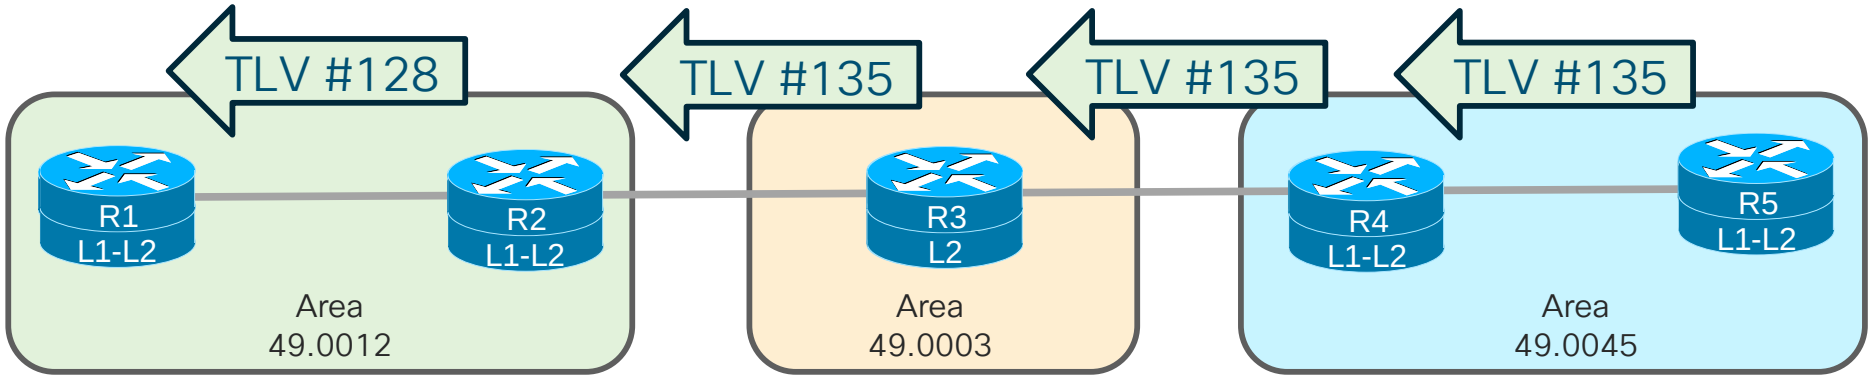

Checking Metric Style

IOS

```
R1#show isis protocol | i narrow|wide
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
```

Checking Metric Style

IOS XR

```
RP/0/0/CPU0:XR1#show isis protocol | i "Level-|style"
Level-1
        Metric style (generate/accept): Narrow/Narrow
      Level-2
        Metric style (generate/accept): Narrow/Narrow
```

NX-OS

```
NX-1# show isis protocol | i Metric
  Metric-style : advertise(wide), accept(narrow, wide)
```

Narrow vs. Wide Metrics

A router can use Narrow, Wide, or Transition Metrics (Both)

IOS

```
router isis CISCOLIVE
metric-style {narrow | transition | wide}
```

IOS XR

```
router isis CISCOLIVE
 address-family ipv4 unicast
   metric-style {narrow | transition | wide}
```

NX-OS

```
router isis CISCOLIVE
  metric-style transition
```

Needs to be consistently Narrow or Wide (Exception is Transition)

Trivia Question Hint

- How do you make the traffic between R1 and R4, take R5?
- How do you make the traffic between R2 and R3, take the direct link?

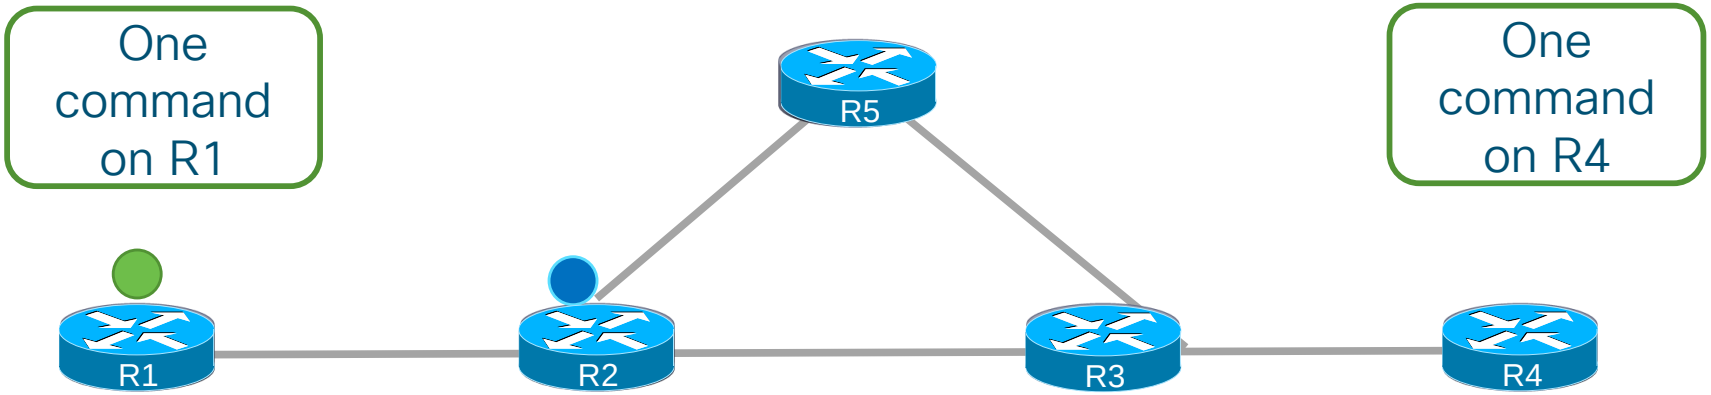

Route Advertisements

Understanding Route Advertisement

- IS-IS builds the topology using TLV#2.
- Routes are built using TLV#128 and/or TLV#135
- All routers in the same L1/L2 area have the same LSPDBs for that Level

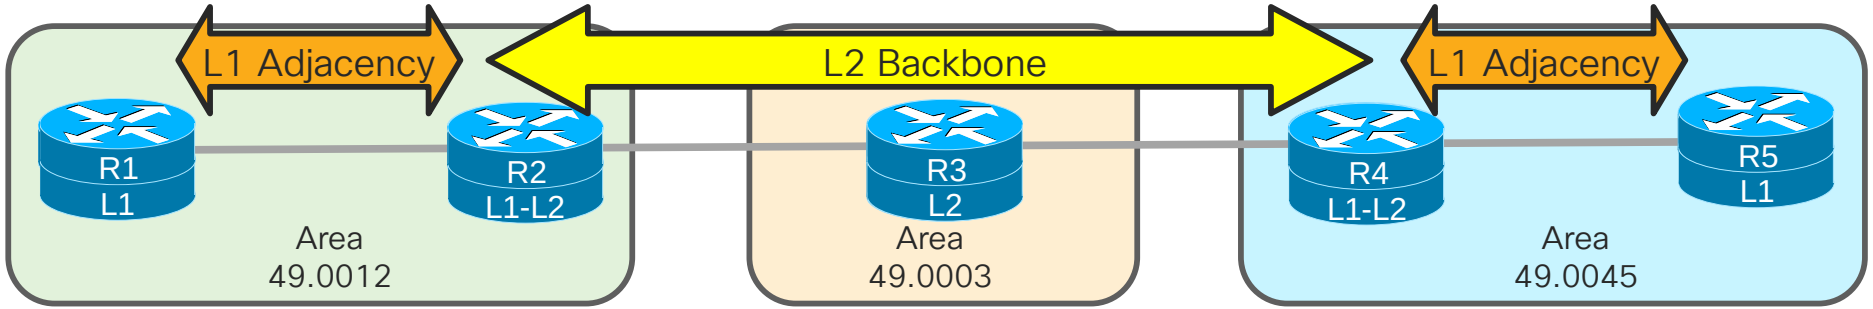

Viewing Routes in the LSPDB

```
R1#show isis database R2.00-00 level-1 detail
Tag
IS-1                    0-00
LSPI                    LSP Seq Num  LSP C          he/Rcvd        ATT/P/OL
R2.00-00                0x0000000E   0x2BD          2/1199         1/0/0
  Area Ad ess: 49.0012
Hostname: R2
  Metric: 10         IS R1.00
  IP Address:   192.168.2.2
  Metric: 10         IP 10.12.1.0 255.255.255.0
  Metric: 10         IP 10.23.1.0 255.255.255.0
  Metric: 10         IP 192.168.2.2 255.255.255.255
  IPv6 Address: 2001:2:2::
  Metric: 10         IPv6 2000:12::/64
  Metric: 10         IPv6 2000:23::/64
  Metric: 10         IPv6 2001:2:2::/128
```

Metric for the topology

Connected IPv4 Networks

Connected IPv6 Networks

Understanding Route Advertisement

- L1 advertises the router 192.168.1.1 within Area 49.0012 via L1
- R2 takes the L1 route and places it into the L2 router as its own
- R2's L2 LSP is Forwarded to R3 and R4
- Does R4 advertise 192.168.1.1 into Area 49.0045?
- No, it does not!

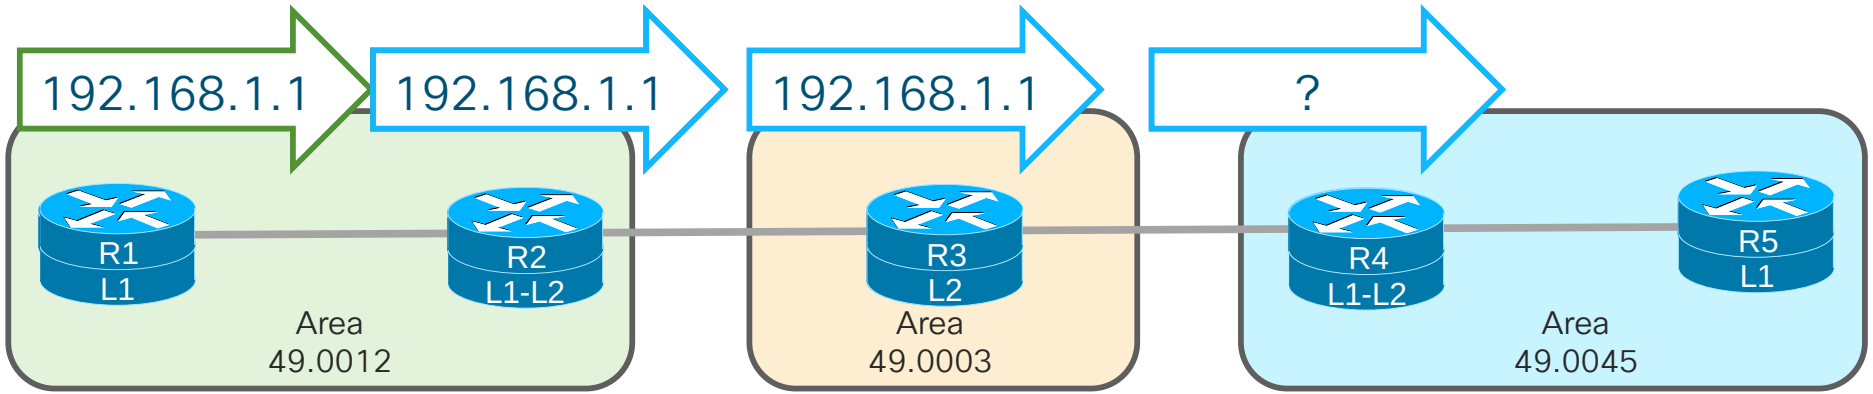

Understanding Route Advertisement

- L1 advertises the router 192.168.1.1 within Area 49.0012 via L1
- R2 takes the L1 route and places it into the L2 router as its own
- R2's L2 LSP is Forwarded to R3 and R4
- Does R4 advertise 192.168.1.1 into Area 49.0045?
- No, it does not!

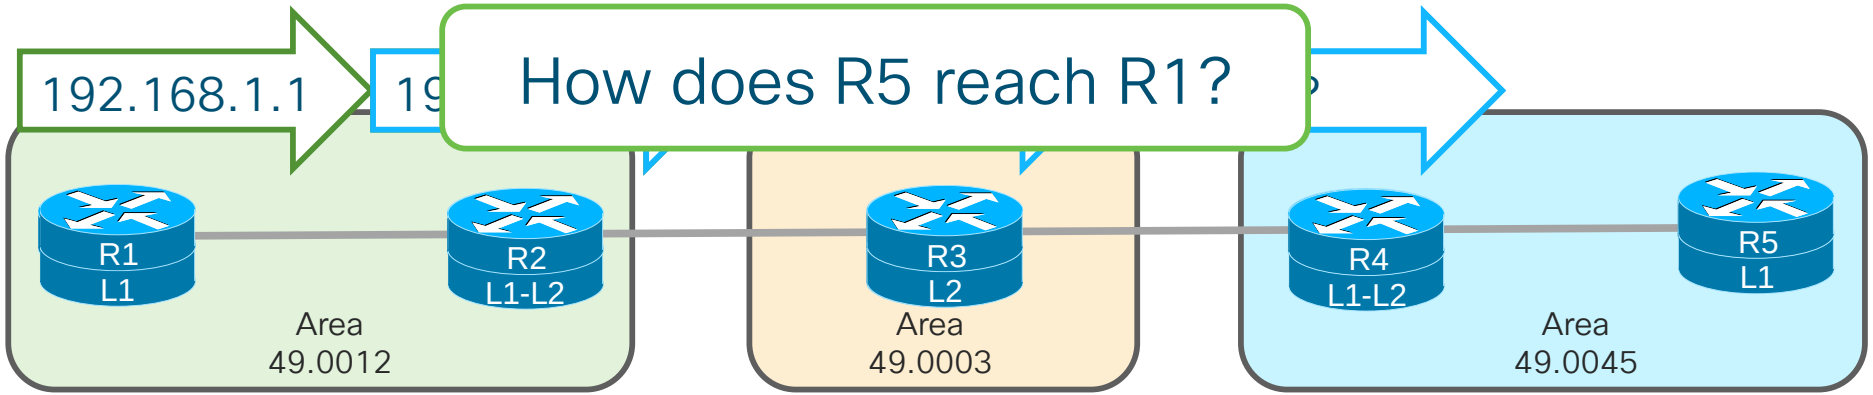

The Attached Bit

- L1 routers use the attach bit to locate their nearest L1-L2 router
  - That L1-L2 router must contain LSPs from a different area.
- The L1-L2 router acts as a gateway
- L1 routers translate the Attach bit as the default gateway

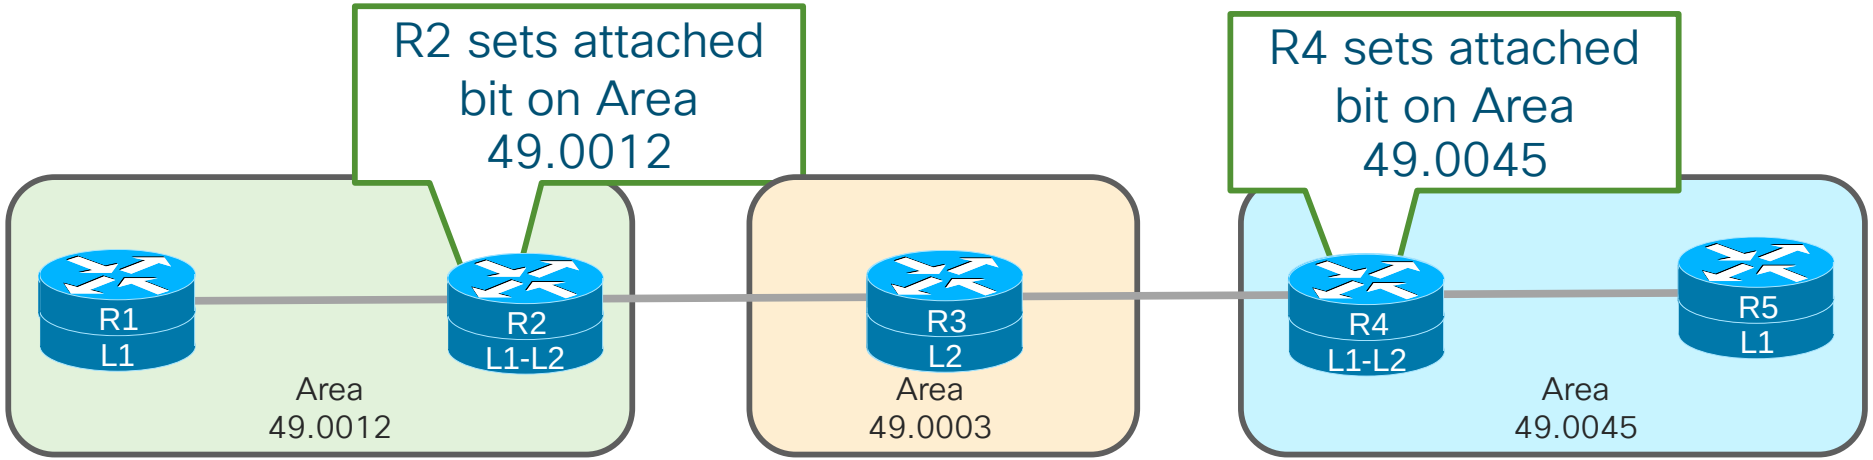

The Attached Bit

- Can be viewed by looking examining Attribute fields

```
R1#show isis database
Tag CISCOLIVE:
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
R1.00-00            * 0x0000002E   0xB3D2              1174/*           0/0/0
R2.00-00              0x0000002C   0xF6DE               932/1199        1/0/0
```

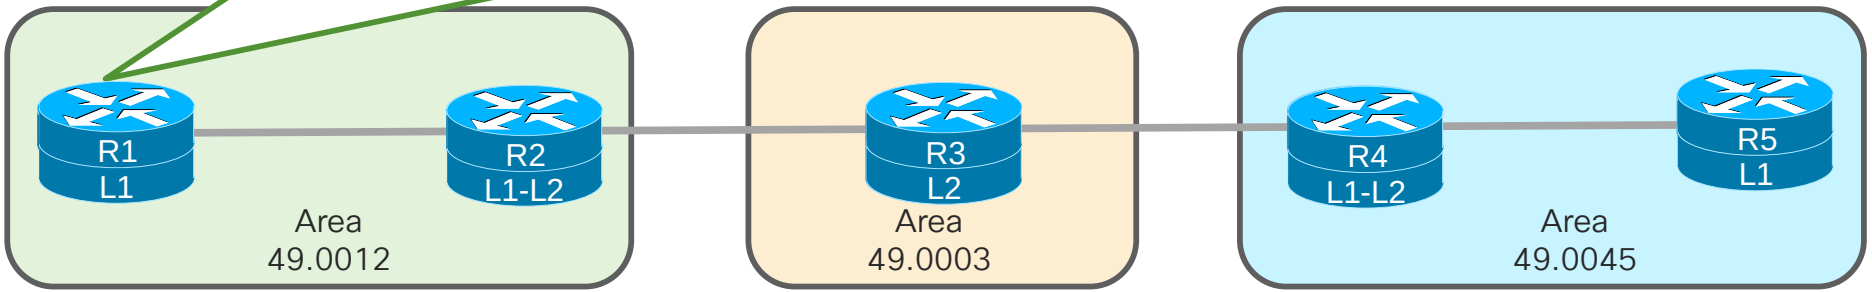

The Attached Bit

- Can be seen by viewing an explicit router's LSP too

```
R1#show isis database detail R2.00-00
IS-IS Level-1 LSP R2.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
R2.00-00              0x0000002E   0xF2E0                725/1199      1/0/0
  Area Address: 49.0012
  NLPID:        0xCC 0x8E
  Hostname: R2
  Metric: 10         IS R1.00
  IP Address:   192.168.2.2
  Metric: 10         IP 10.12.1.0 255.255.255.0
  Metric: 10         IP 10.23.1.0 255.255.255.0
  Metric: 10         IP 192.168.2.2 255.255.255.255
```

The Attached Bit

- Translating it to the Routing Table

```
R1#show ip route isis
i*L1   0.0.0.0/0 [115/10] via 10.12.1.2, 00:03:18, GigabitEthernet0/2
       10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
i L1      10.23.1.0/24 [115/20] via 10.12.1.2, 00:03:18, GigabitEthernet0/2
       192.168.2.0/32 is subnetted, 1 subnets
i L1      192.168.2.2 [115/20] via 10.12.1.2, 00:03:18, GigabitEthernet0/2
```

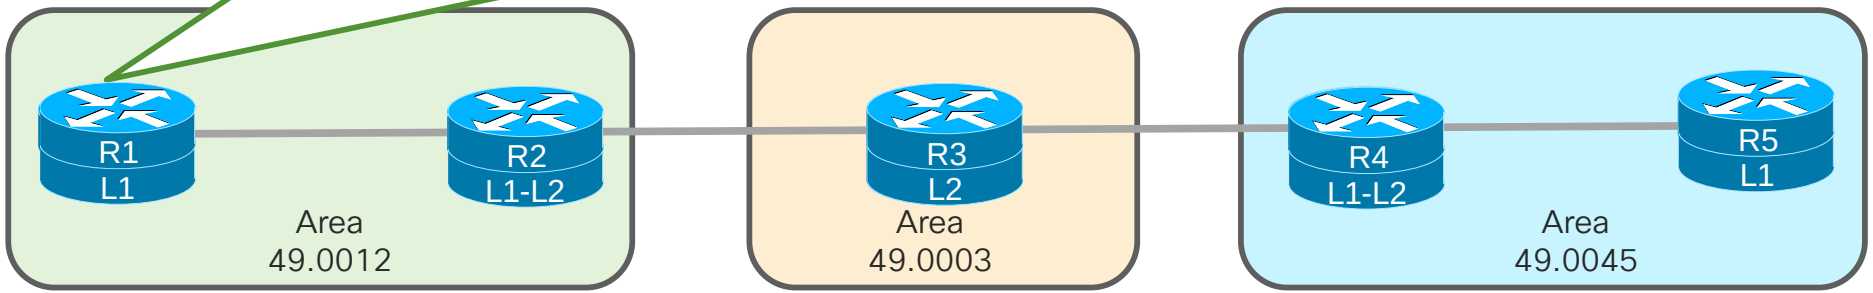

Viewing the Backbone Routing Table

```
R2#show ip route isis | ex subnet
i L2     10.34.1.0/24 [115/20] via 10.23.1.3, 00:18:41, GigabitEthernet0/3
i L2     10.45.1.0/24 [115/30] via 10.23.1.3, 00:18:41, GigabitEthernet0/3
i L1     192.168.1.1 [115/20] via 10.12.1.1, 07:54:35, GigabitEthernet0/1
i L2     192.168.3.3 [115/20] via 10.23.1.3, 00:18:41, GigabitEthernet0/3
i L2     192.168.4.4 [115/30] via 10.23.1.3, 00:18:41, GigabitEthernet0/3
i L2     192.168.5.5 [115/40] via 10.23.1.3, 00:18:41, GigabitEthernet0/3
```

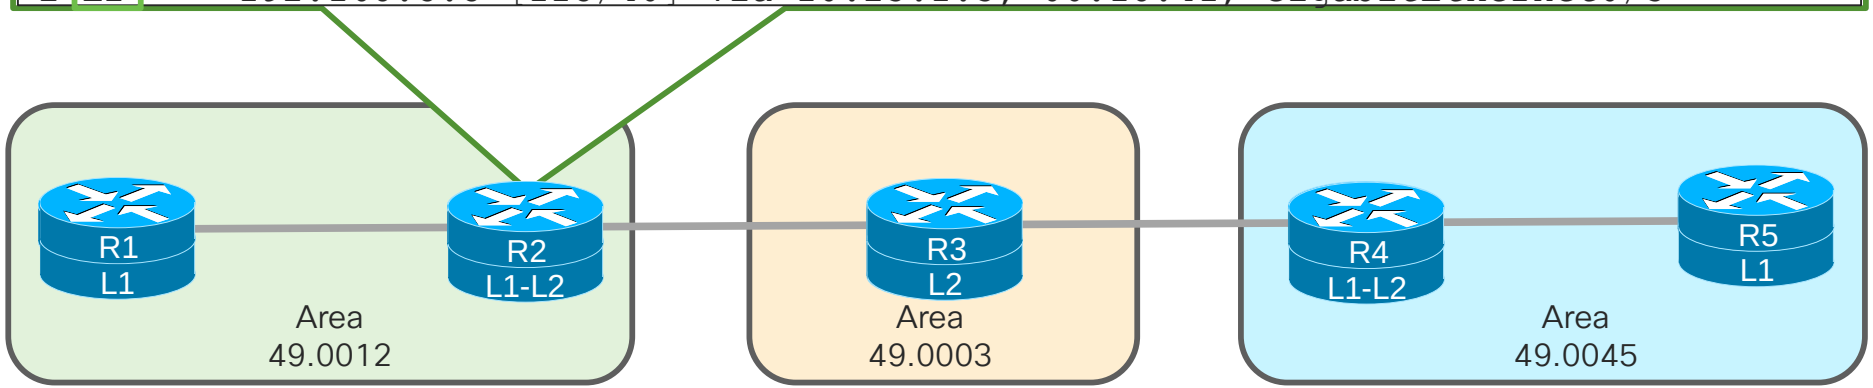

Pop Quiz

Troubleshooting a connectivity problem

```
R1#ping 192.168.5.5 source loopback 0
Sending 5, 100-byte ICMP Echos to 192.168.5.5, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
U.U.U
Success rate is 0 percent (0/5)
R1#trace 192.168.5.5 source loopback 0
Tracing the route to 192.168.5.5
VRF info: (vrf in name/id, vrf out name/id)
  1 10.12.1.2 2 msec 2 msec 2 msec
  2 10.12.1.2 !H  *  !H
```

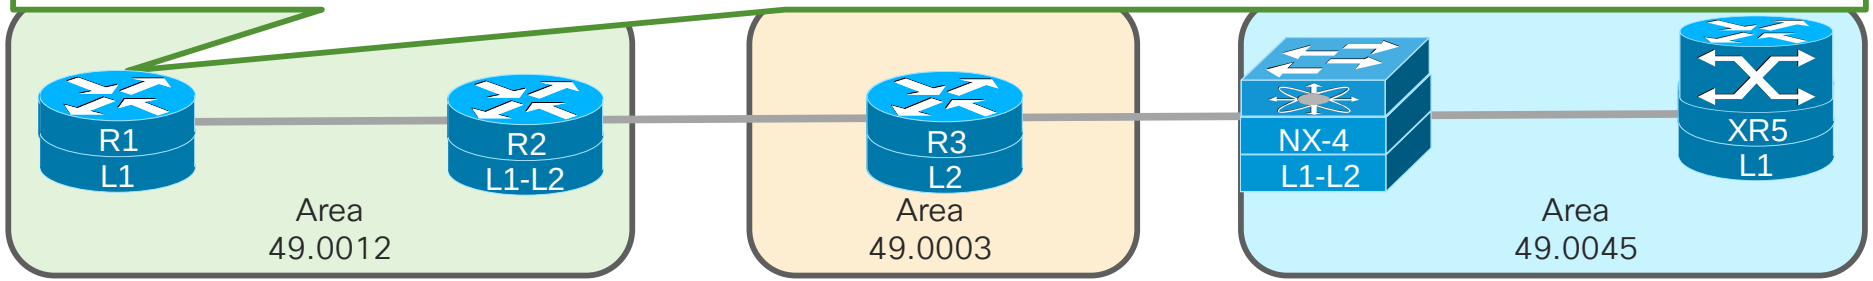

Troubleshooting a connectivity problem

```
R2#show ip route isis | ex subnett
i L2      10.34.1.0/24 [115/20] via 10.23.1.3, 00:46:08, GigabitEthernet0/3
i L2      10.45.1.0/24 [115/60] via 10.23.1.3, 00:04:40, GigabitEthernet0/3
i L1      192.168.1.1 [115/20] via 10.12.1.1, 00:10:43, GigabitEthernet0/1
i L2      192.168.3.3 [115/20] via 10.23.1.3, 00:46:08, GigabitEthernet0/3
i L2      192.168.4.4 [115/21] via 10.23.1.3, 00:04:40, GigabitEthernet0/3
```

What happened to 192.168.5.5?

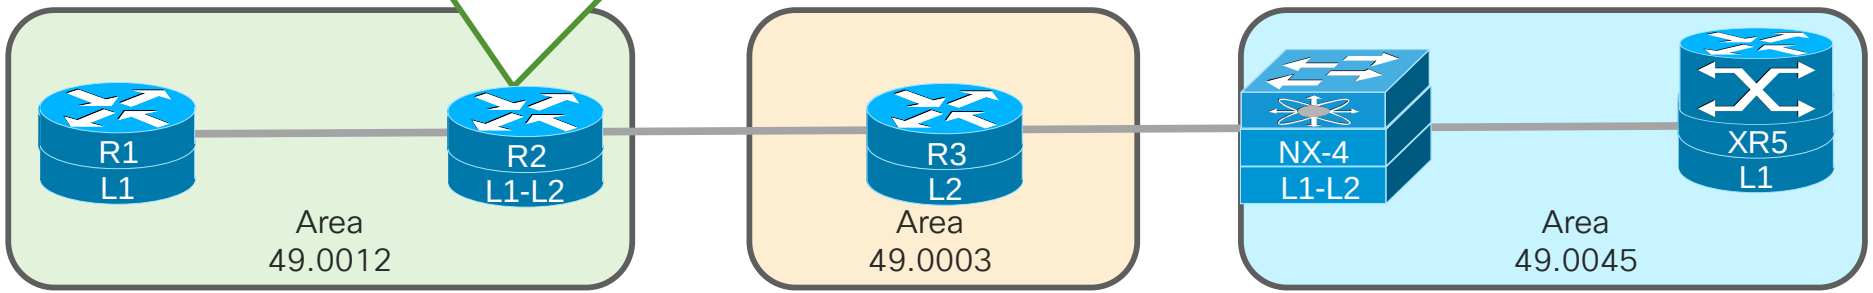

Troubleshooting a connectivity problem

```
RP/0/0/CPU0:XR5#trace 192.168.1.1 source 192.168.5.5
Type escape sequence to abort.
Tracing the route to 192.168.1.1

 1  10.45.1.4 9 msec  0 msec  0 msec
 2   *  *  *
```

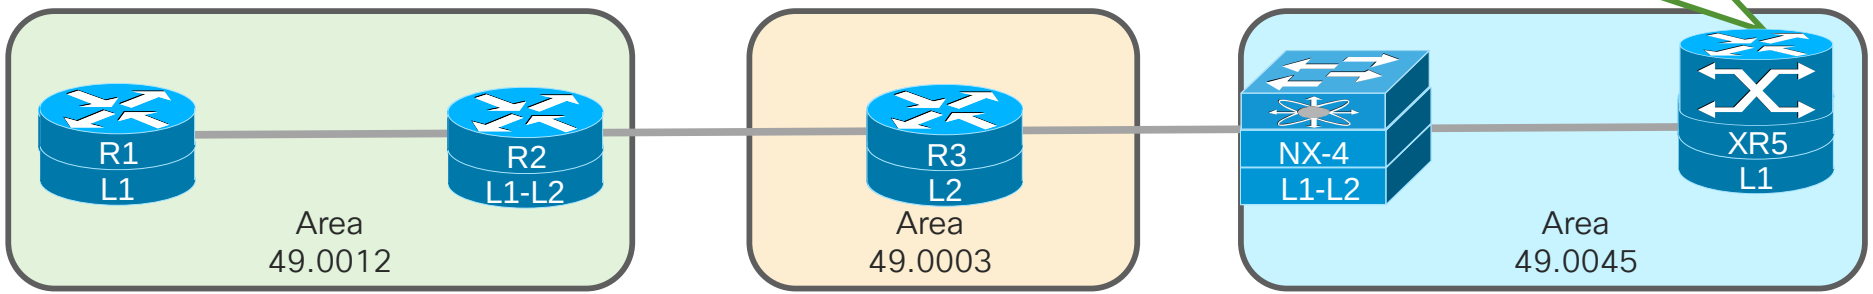

Troubleshooting a connectivity problem

```
NX-4# show ip route isis
{SNIP}
192.168.1.1/32, ubest/mbest: 1/0
    *via 10.34.1.3, Eth2/3, [115/70], 00:15:10, isis-CISCOLIVE, L2
192.168.2.2/32, ubest/mbest: 1/0
    *via 10.34.1.3, Eth2/3, [115/60], 00:17:06, isis-CISCOLIVE, L2
192.168.3.3/32, ubest/mbest: 1/0
    *via 10.34.1.3, Eth2/3, [115/50], 00:17:14, isis-CISCOLIVE, L2
192.168.5.5/32, ubest/mbest: 1/0
    *via 10.45.1.5, Eth2/5, [115/50], 00:17:51, isis-CISCOLIVE, L1
```

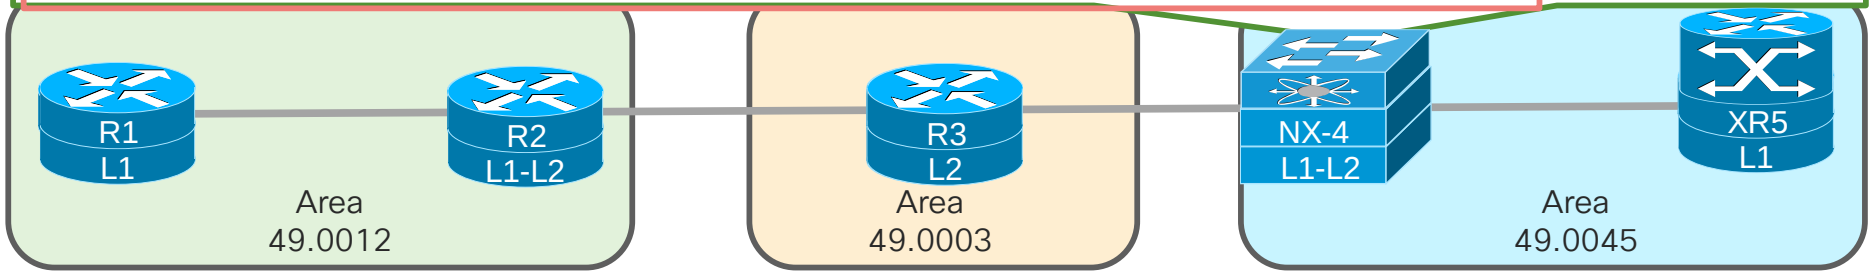

Troubleshooting a connectivity problem

```
NNX-4# show isis database level-2 detail NX-4.00-00
IS-IS Level-2 Link State Database
Extended IS    :  R3.01                Metric : 40
IP Internal    :     192.168.4.4/32    Metric : 1    (I,U)
    IP Internal    :       10.45.1.0/24   Metric : 40   (I,U)
    IP Internal    :       10.34.1.0/24   Metric : 40   (I,U)
    Extended IP    :       10.34.1.0/24   Metric : 40          (U)
    Extended IP    :       10.45.1.0/24   Metric : 40          (U)
    Extended IP    :     192.168.4.4/32   Metric : 1           (U)
```

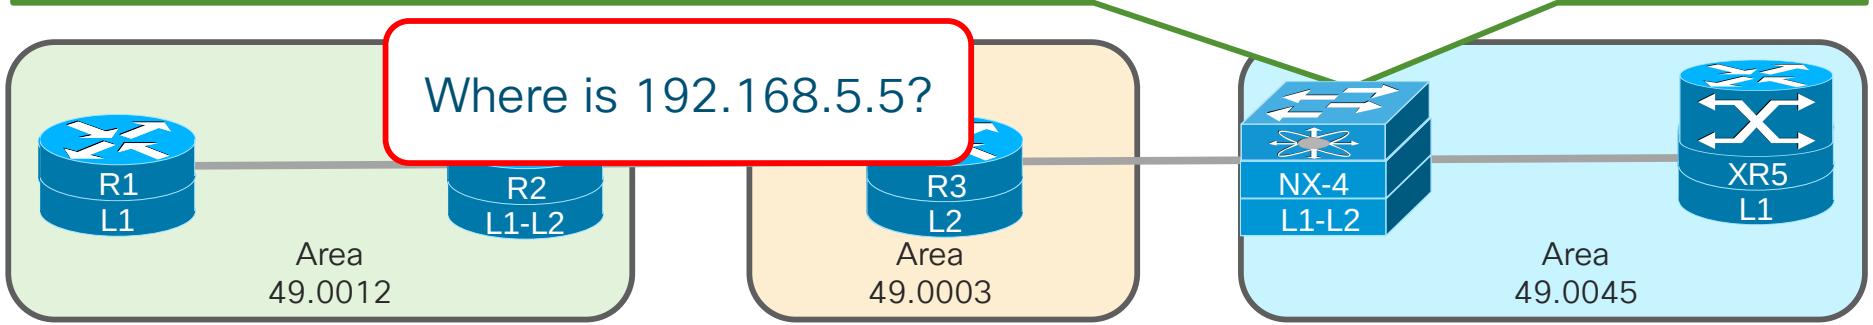

L1-L2 Route Propogation on NX-OS

- A difference in operational behavior between NX-OS and other Cisco operating systems (IOS, IOS XR, etc.).
- Nexus switches require explicit configuration with the command **`distribute level-1 into level-2`** `{`**`all`** `|` *`route-map route-map-name`*`}` on L1-L2 routers to insert L1 routes into the L2 topology.

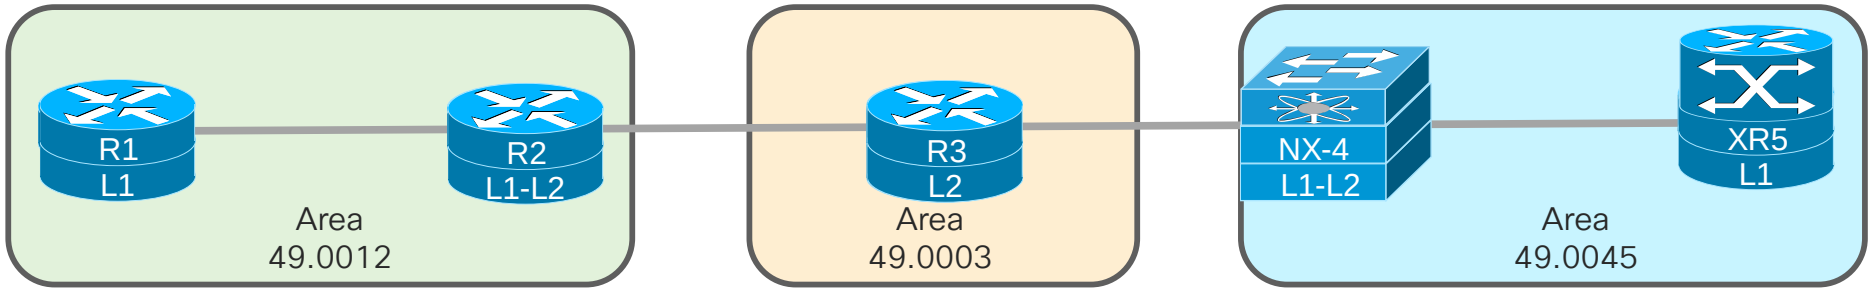

L1-L2 Route Propagation on NX-OS

```
NNX-4# show isis database level-2 detail NX-4.00-00
IS-IS Level-2 Link State Database
IP Internal    :      192.168.4.4/32  Metric : 1    (I,U)
    IP Internal    :        10.45.1.0/24  Metric : 40   (I,U)
    IP Internal    :        10.34.1.0/24  Metric : 40   (I,U)
    Extended IP    :      192.168.5.5/32  Metric : 50           (U)
    Extended IP    :        10.34.1.0/24  Metric : 40           (U)
    Extended IP    :        10.45.1.0/24  Metric : 40           (U)
    Extended IP    :      192.168.4.4/32  Metric : 1            (U)
```

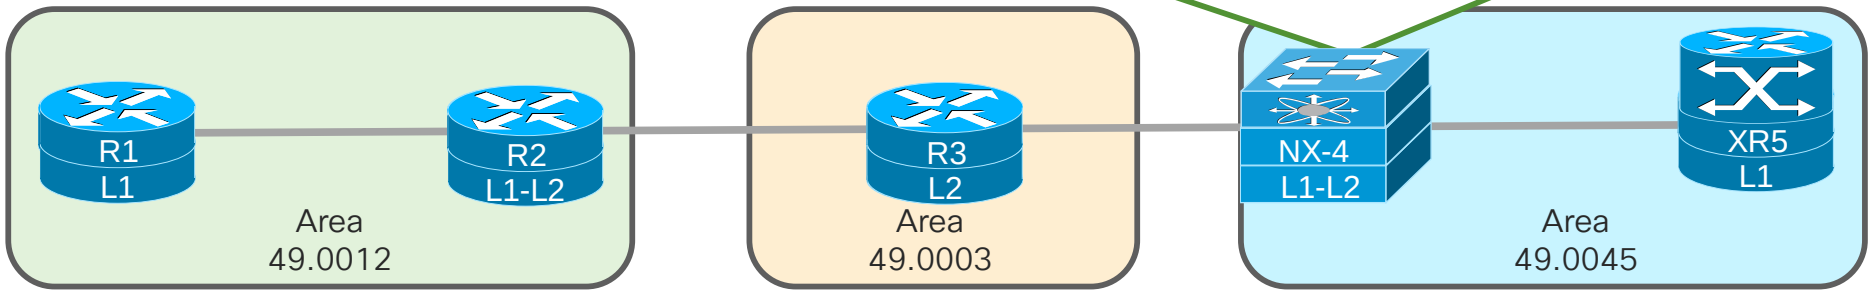

Area Design

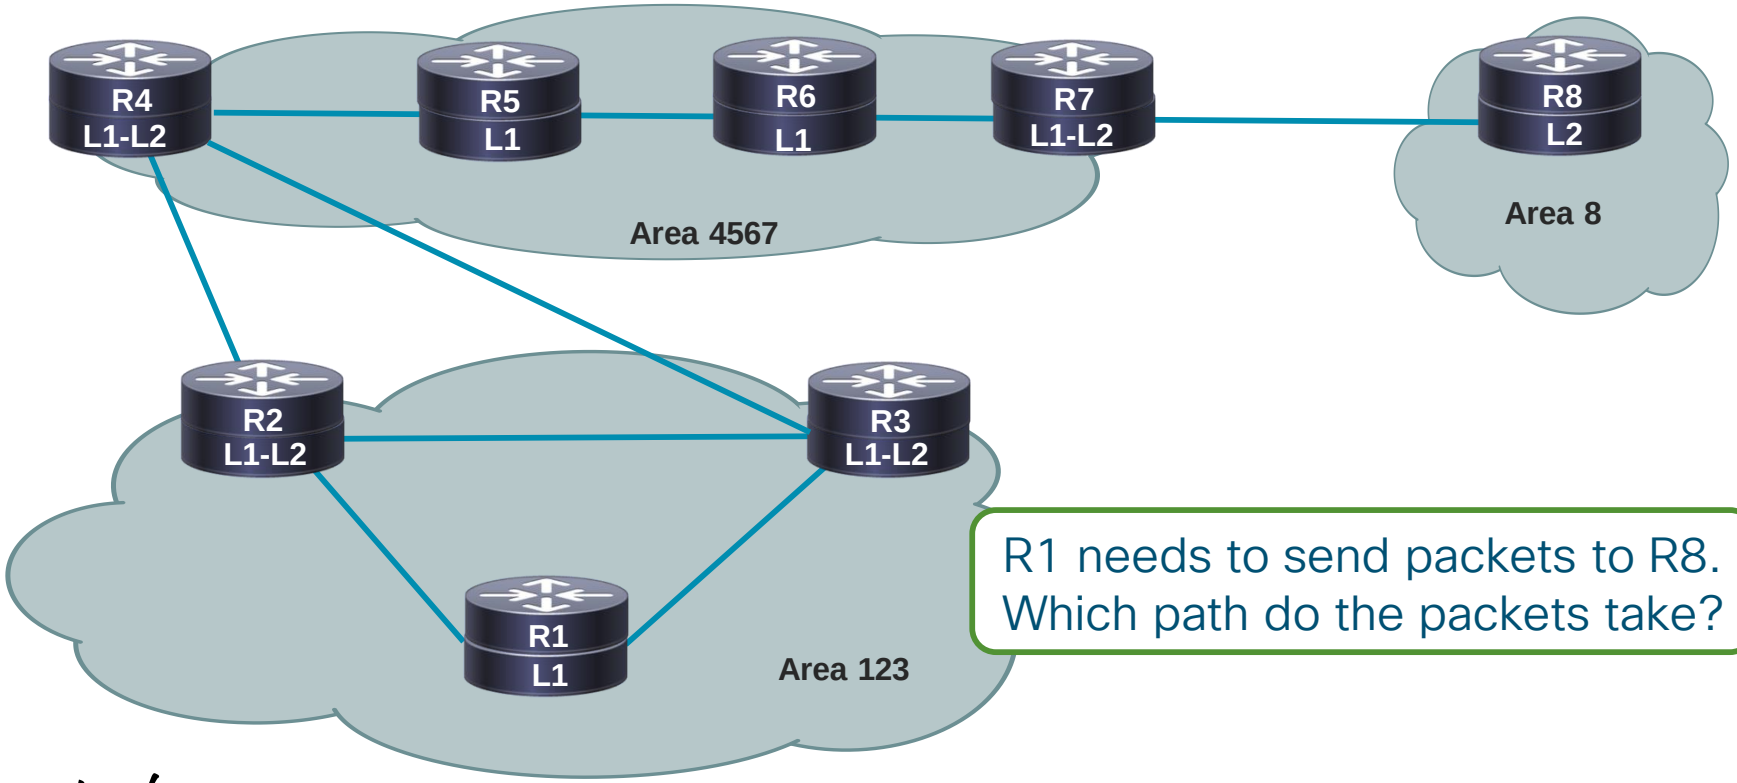

Area Design

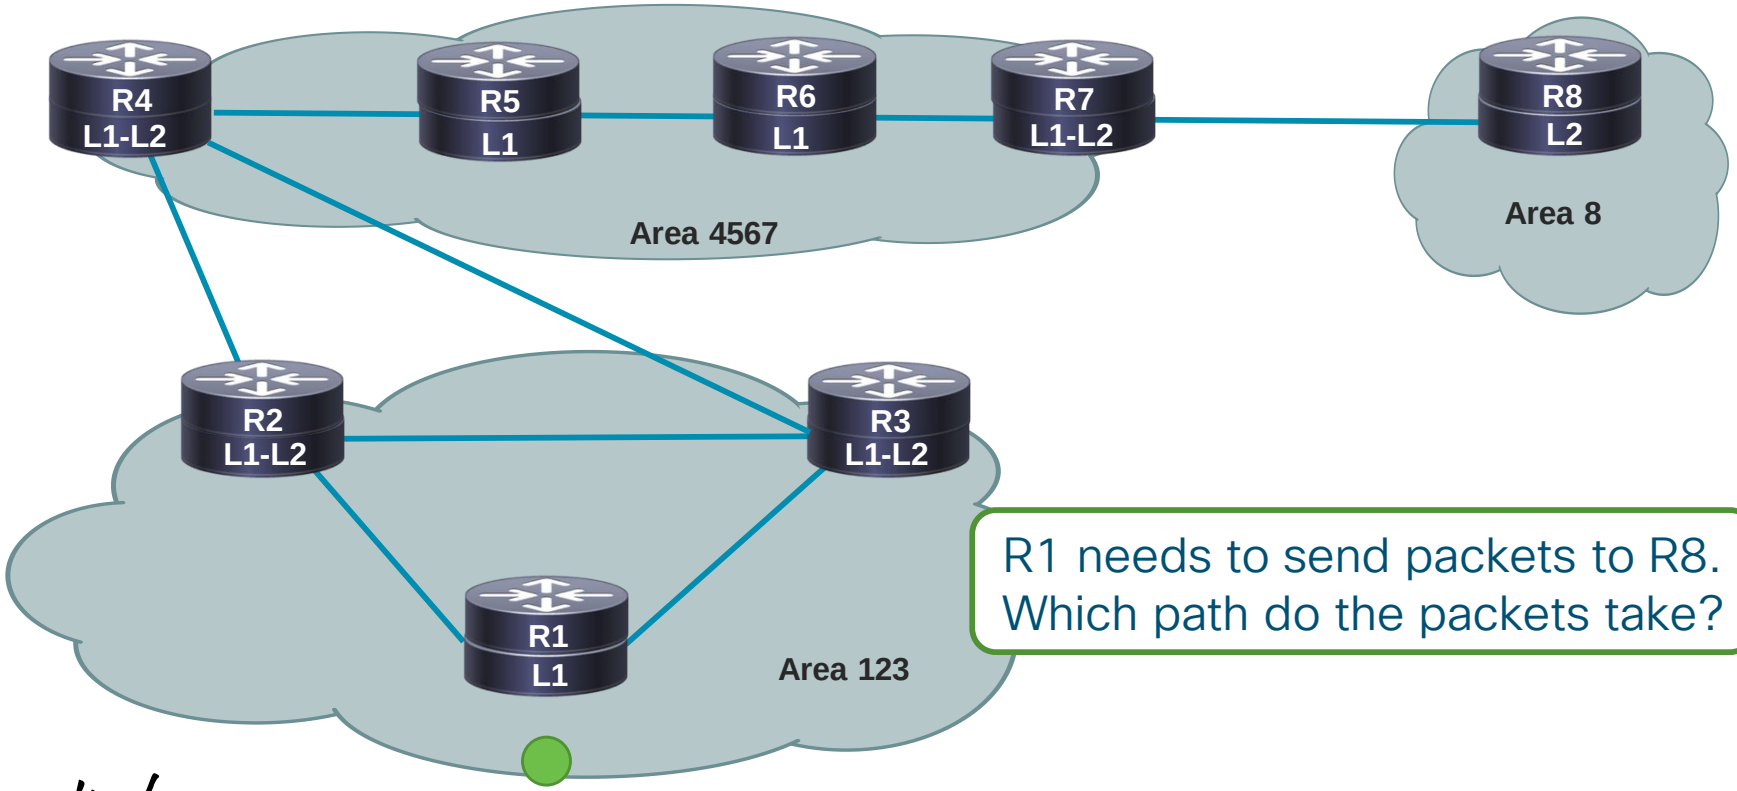

R1 needs to send packets to R8.
Which path do the packets take?

Area Design

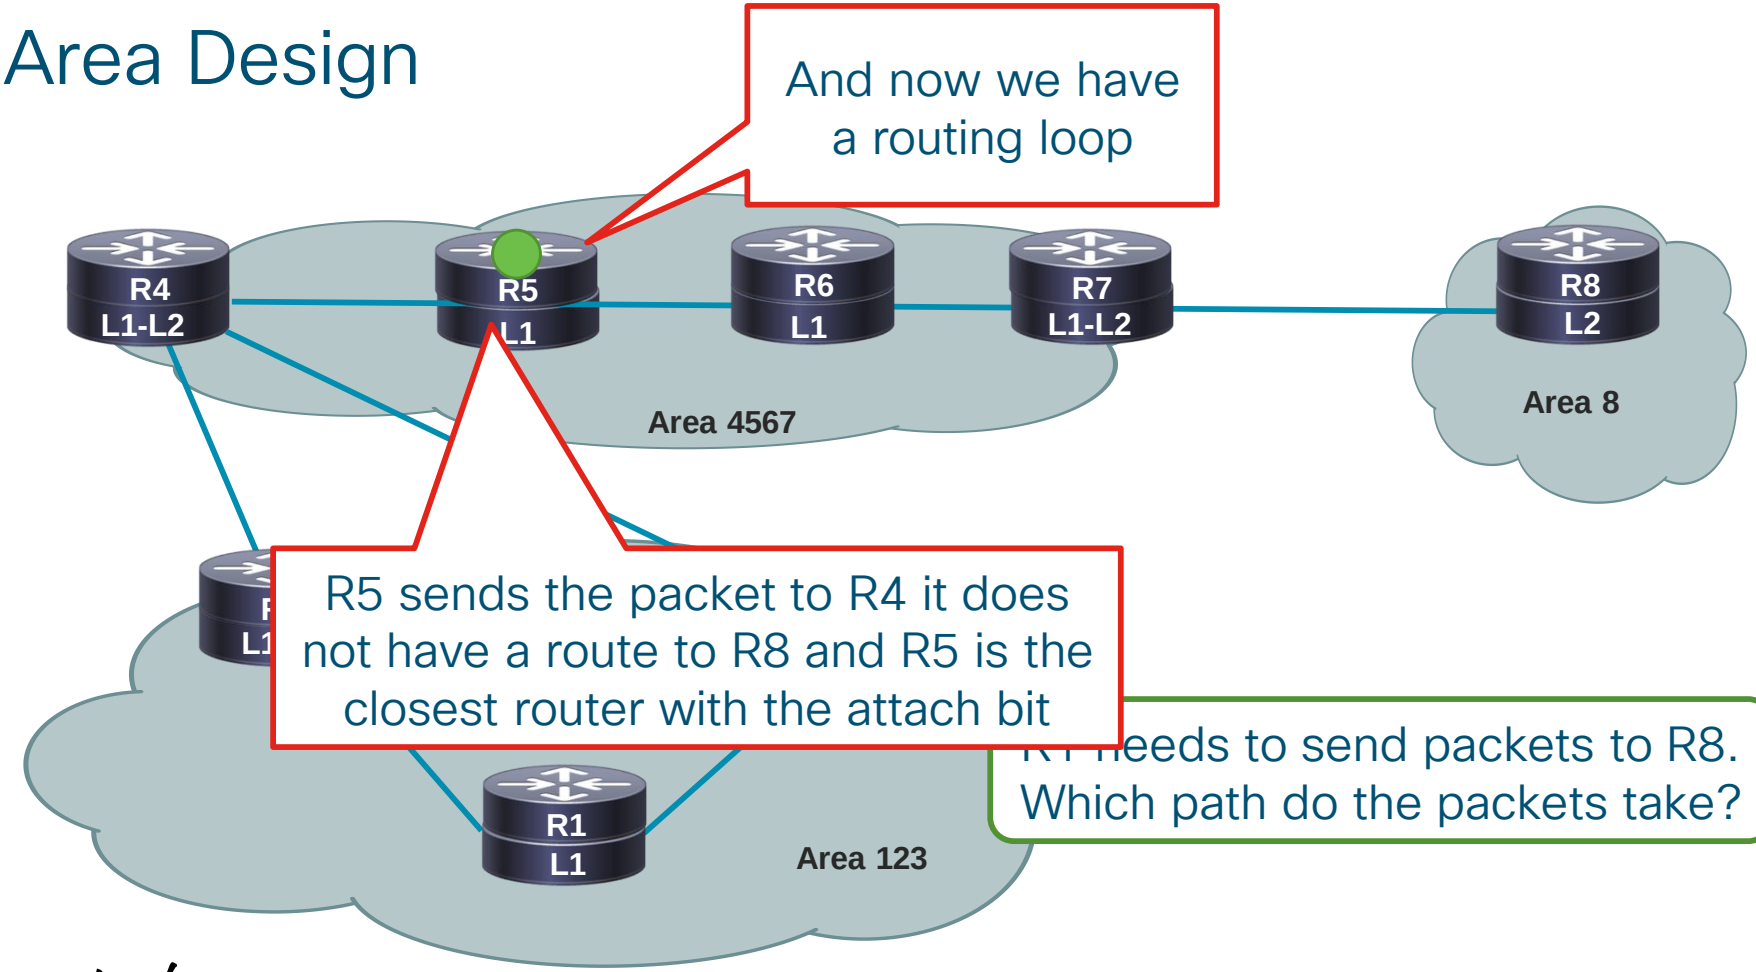

Pop Quiz?

- XR2 cannot see R1's 10.1.1.0/24 network.
- Help me troubleshoot it..

Troubleshooting Missing Routes:

```
R1#show isis topology
Tag CISCOLIVE:
IS-IS TID 0 paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
XR2                  10         XR2                  Gi0/2       fa16.3e7f.311b
```

```
RP/0/0/CPU0:XR2#show isis topology

IS-IS CISCOLIVE paths to IPv4 Unicast (Level-1) routers
System Id       Metric  Next-Hop        Interface       SNPA
R1              10      R1              Gi0/0/0/1       fa16.3e61.7cf8
XR2             --
```

Troubleshooting Missing Routes:

```
RP/0/0/CPU0:XR2#show isis database level 1 R1.00-00 detail

IS-IS CISCOLIVE (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R1.00-00              0x00000005   0x44ed        1189              0/0/0
  Area Address:   49.0123
  NLPID:          0xcc
  Hostname:       R1
  Metric: 10         IS XR2.01
  IP Address:     10.12.1.1
  Metric: 10         IP 10.12.1.0/24
```

Troubleshooting Missing Routes:

```
R1#show protocol
GigabitEthernet0/2 is up, line protocol is up
  Internet address is 10.12.1.1/24
  CLNS enabled
GigabitEthernet0/3 is up, line protocol is up
  Internet address is 10.1.1.1/24
  CLNS enabled
```

Troubleshooting Missing Routes:

```
R1#show run | s router|interface
interface GigabitEthernet0/2
 description to XR2
 ip address 10.12.1.1 255.255.255.0
 ip router isis CISCOLIVE
interface GigabitEthernet0/3
 description to NX-3
 ip address 10.1.1.1 255.255.255.0
 ip router isis
router isis
router isis CISCOLIVE
 net 49.0123.0001.0001.0001.00
 is-type level-1
 log-adjacency-changes
```

```
RP/0/0/CPU0:XR2#show run router isis
router isis CISCOLIVE
 is-type level-1
 net 49.0123.0002.0002.0002.00
 log adjacency changes
 interface GigabitEthernet0/0/0/1
  address-family ipv4 unicast
  !
 !
 interface GigabitEthernet0/0/0/3
  address-family ipv4 unicast
  !
 !
!
```

Recommend using a process-id on IOS boxes to keep consistent with NX-OS and IOS XR

Quiz Afterthoughts

- XR2 would not have seen R10 from a topology perspective
- Drill down to R1 and R10 Adjacency
- Once that was resolved, then it would have been straightforward on the misconfiguration on R1

Sup-Optimal IS-IS Routing

- Area design

Areas and Suboptimal Routing

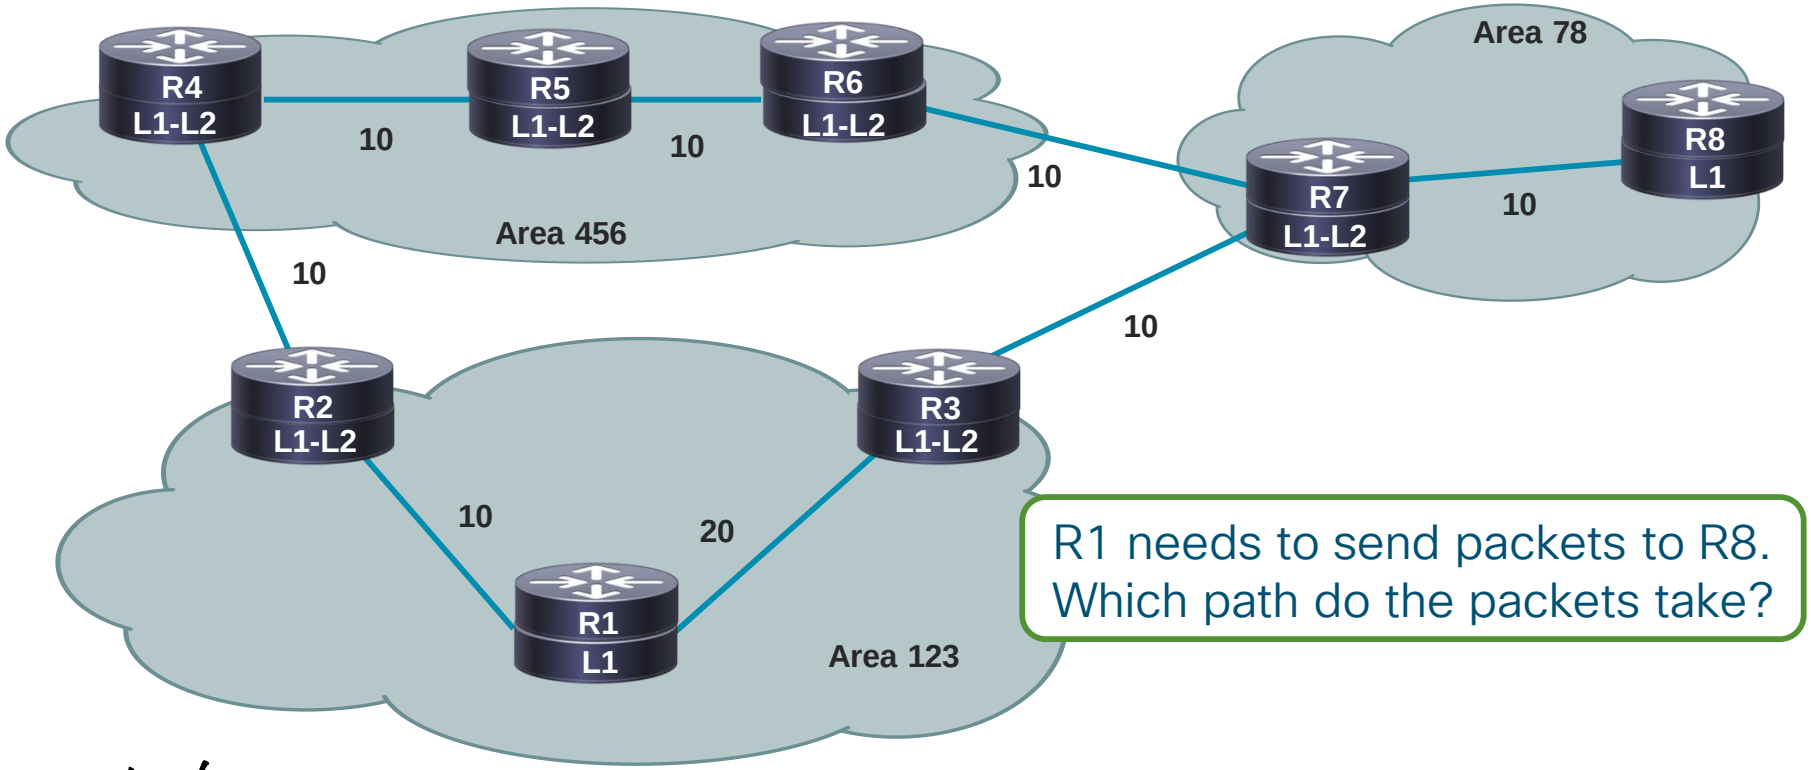

R1 needs to send packets to R8.
Which path do the packets take?

Areas and Suboptimal Routing

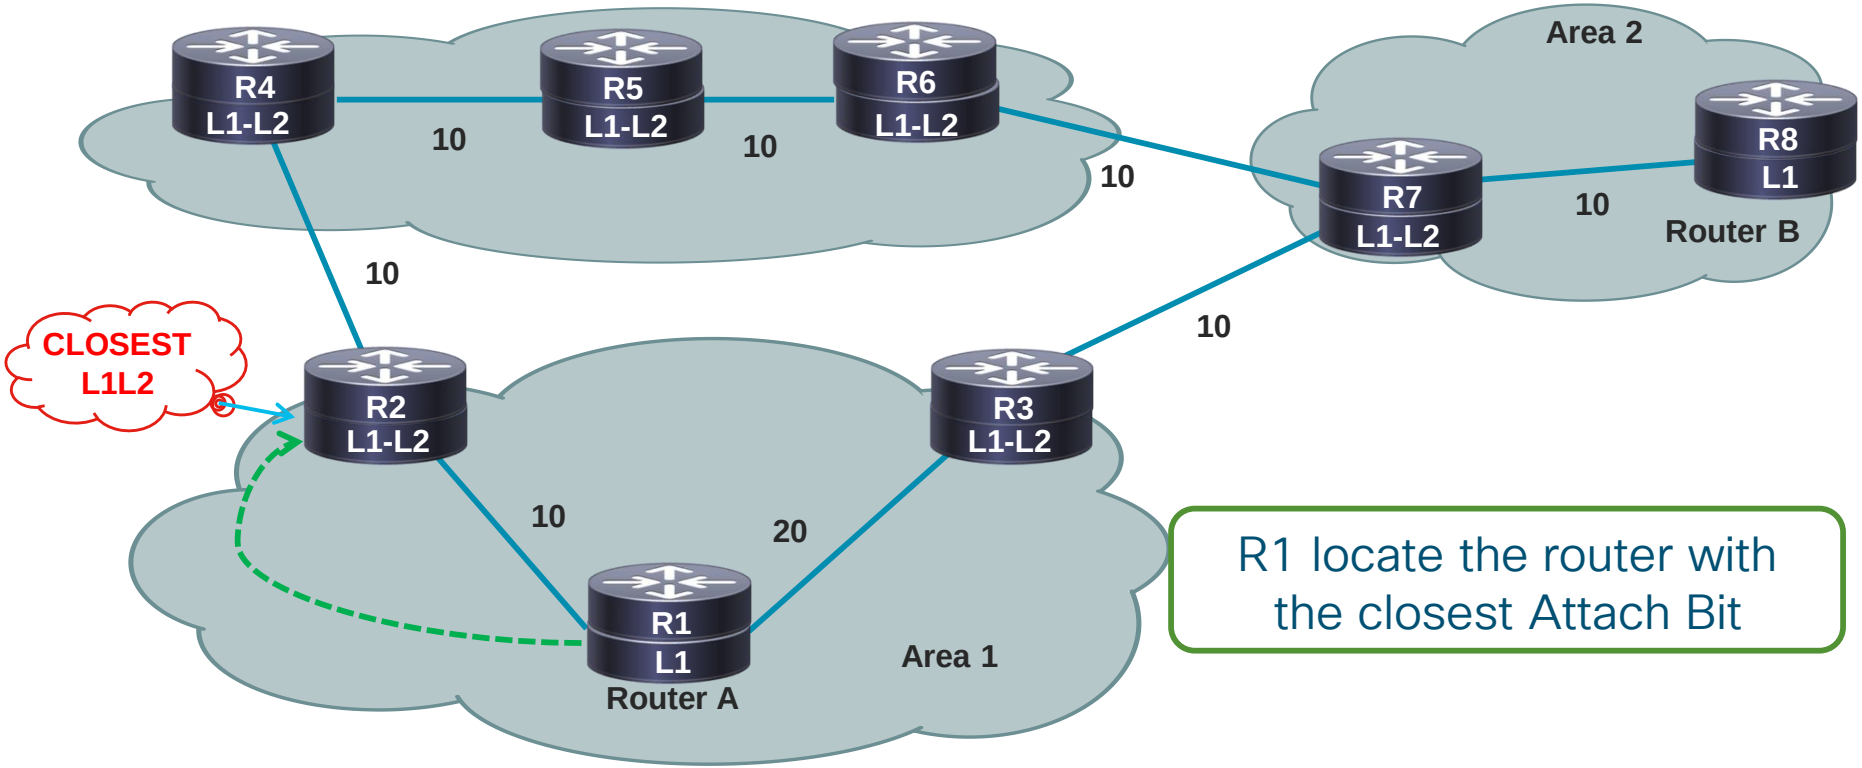

R1 locate the router with the closest Attach Bit

Areas and Suboptimal Routing

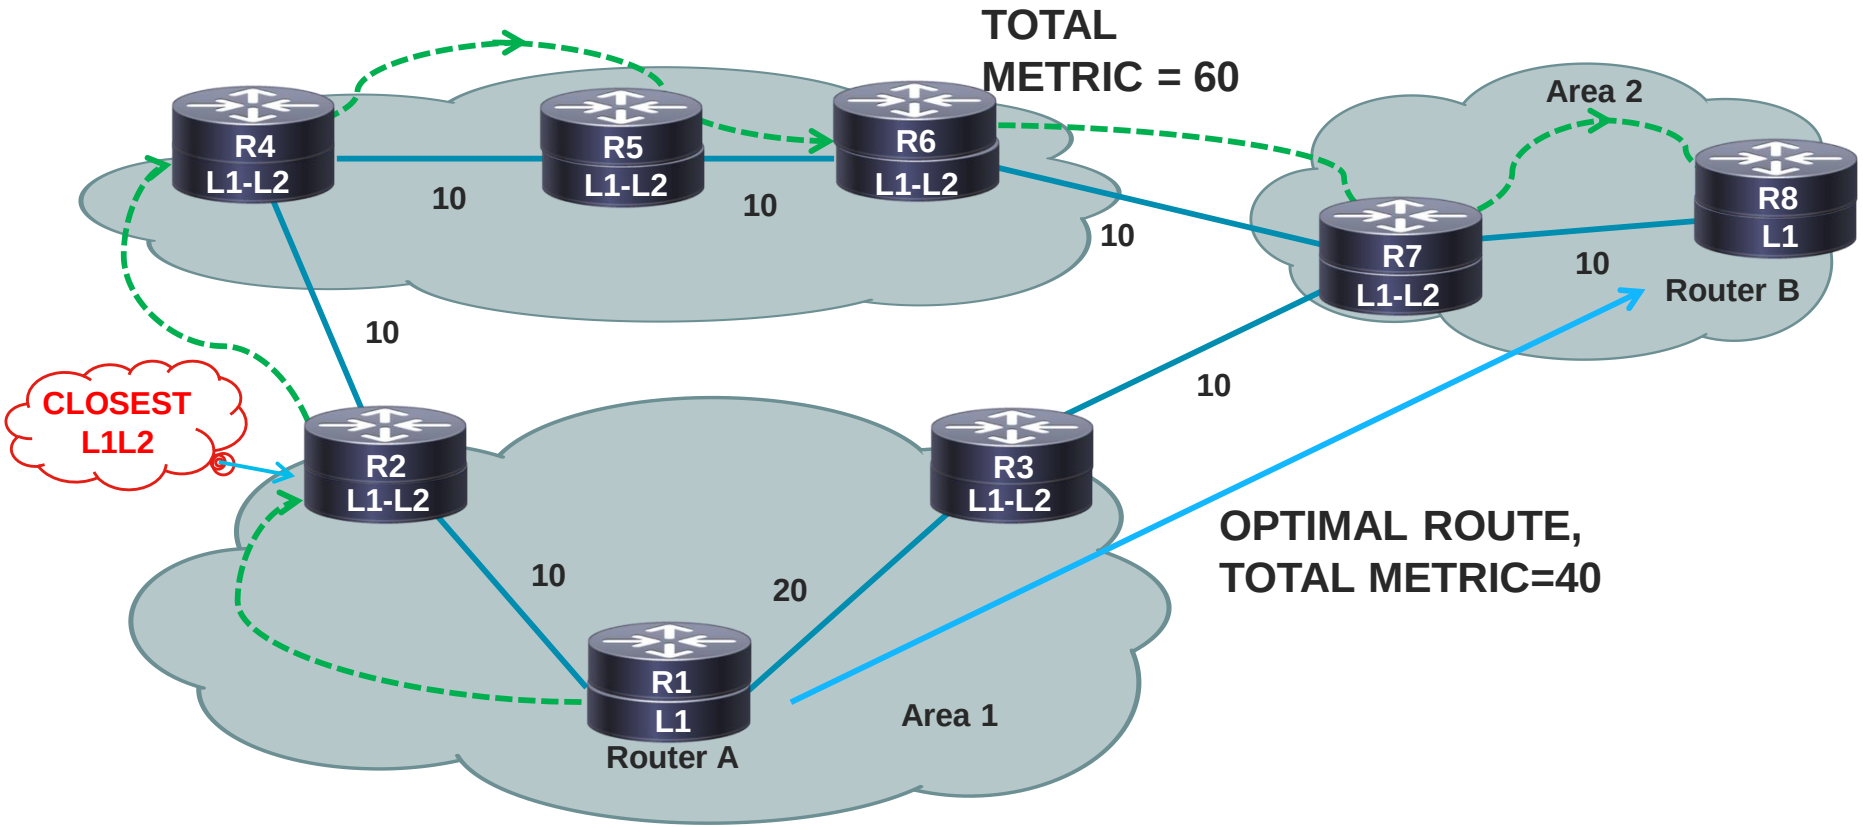

Overcoming Areas and Suboptimal Routing

- All the L1-routers in a given pop will receive the ATT bit set by the L1L2 router at the edge of the POP
  - L1 routers install a default route based on the ATT bit
  - This will cause sub-optimal routing in reaching the prefixes outside the POP by the local routers
- This can be overcome by Leaking more explicit L2 routes into the L1 area

L2 → L1 Leaking Configuration

IOS

```
R1#conf t
R1(config)#router isis CISCOLIVE
R1(config)#redistribute isis ip level-2 into level-1
```

IOS XR

```
RP/0/0/CPU0:XR2#conf t
RP/0/0/CPU0:XR2(config)#router isis CISCOLIVE
RP/0/0/CPU0:XR2(config-isis)#address-family ipv4 unicast
RP/0/0/CPU0:XR2(config-isis-af)# propagate level 2 into level 1
```

NX-OS

```
R1#conf t
NX-3(config)# router isis CISCOLIVE
NX-3(config-router)# distribute level-2 into level-1 all
```

L2 → L1 Leaking Configuration (Conditions)

IOS

```
R1#conf t
R1(config)#router isis CISCOLIVE
R1(config)#redistribute isis ip level-2 into level-1 route-map CONDITIONAL
```

IOS XR

```
RP/0/0/CPU0:XR2#conf t
RP/0/0/CPU0:XR2(config)#router isis CISCOLIVE
RP/0/0/CPU0:XR2(config-isis)#address-family ipv4 unicast
RP/0/0/CPU0:XR2(config-isis-af)# propagate level 2 into level 1 route-policy CONDITIONAL
```

NX-OS

```
R1#conf t
NX-3(config)# router isis CISCOLIVE
NX-3(config-router)# distribute level-2 into level-1 route-map CONDITIONAL
```

Area and Scaling

Areas vs. single area

- ISIS supports a large number of routers in a single area
  - More than 400 routers in the backbone is possible
- Starting with L2-only everywhere is a good choice
  - Backbone continuity is ensured from the start
  - Future implementation of level-1 areas will be easier
- Use areas in places where sub-optimal routing is acceptable
  - areas with a single exit point is a better choice from an optimal routing standpoint

Area Design
L1-Only POP

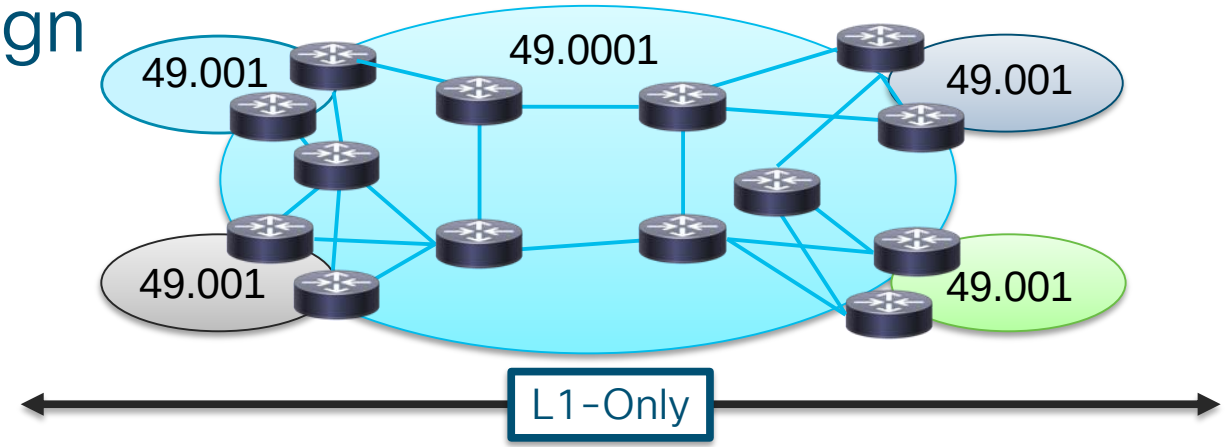

- In this design, all the routers will be running in one area and are all doing L1-only routing
- This design is flat with a single L1-only database running on all the routers
- If you have a change in the topology, the SPF computation will be done in all the routers as they are in the L1-only domain
- SPs picked L1-only to avoid sub-optimal routing problems

Area Design
L2-Only POP

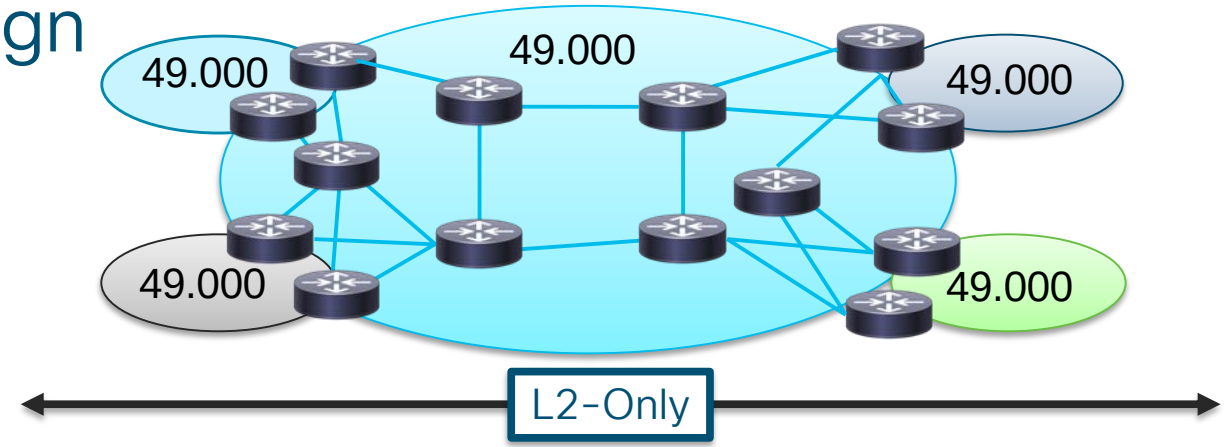

- In this design, all the routers will be running L2-Only in the network
  - With the same Area in all the POPs
- Optimal routing with L2-only database
- Traffic-engineering support with no restrictions, just like L1-only

Area Design
L2-Only POP

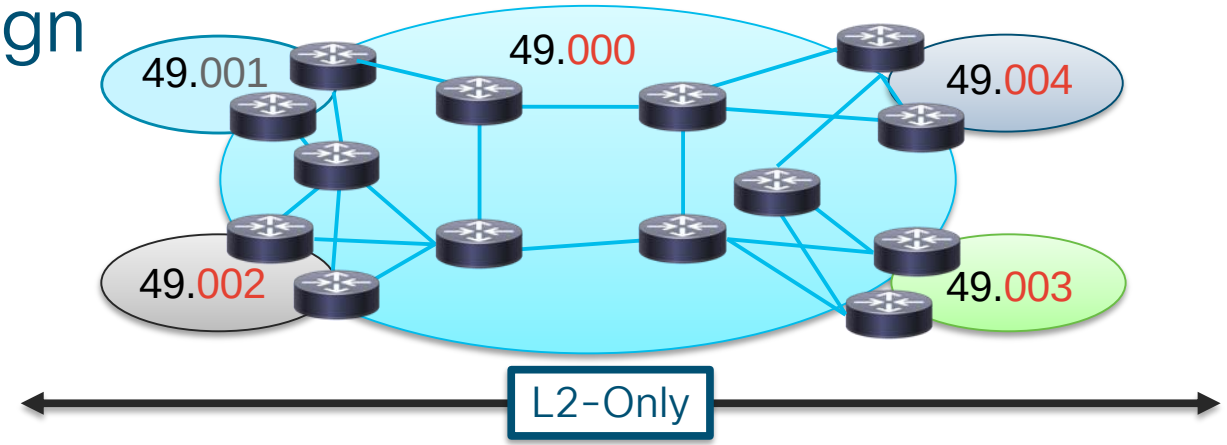

- In this design, all the routers will be running L2-Only in the network
  - With the different Area in all the POPs
  - No summarization and No route-leaking
- All the routers in L2 will share all the LSPs and provides optimal routing (similar to L1-Only POPs)
- As the network grows, easy to bring the L1-only POPs/sub-networks for easy migration

Area Design

L1 in the POP and L2 in the Core

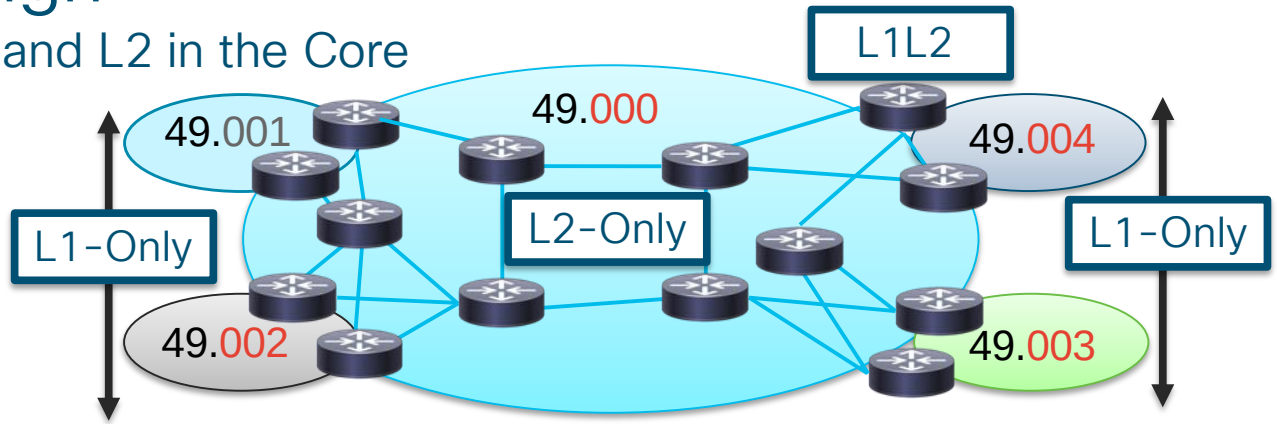

- Within a given local pop—all the routers will be in a separate area
- The L1-L2 routers at the edge of the POPs will be running
  - L1-adj going into the POP
  - L2-adj into the core with the rest of the L1-L2 routers
- The SPF computations will be limited to the respective L1-areas only

Area Design

L1 in the POP and L2 in the Core

- All the L1-routers in a given pop will receive the ATT bit set by the L1L2 router at the edge of the POP
  - L1 routers install a default route based on the ATT bit
- This will cause sub-optimal routing in reaching the prefixes outside the POP by the local routers
- Summarization at the L1L2 boundary
  - potential sub-optimal inter-area routing in certain failure conditions
  - potential black-holing of traffic
  - potential breaking of MPLS LSP among PEs

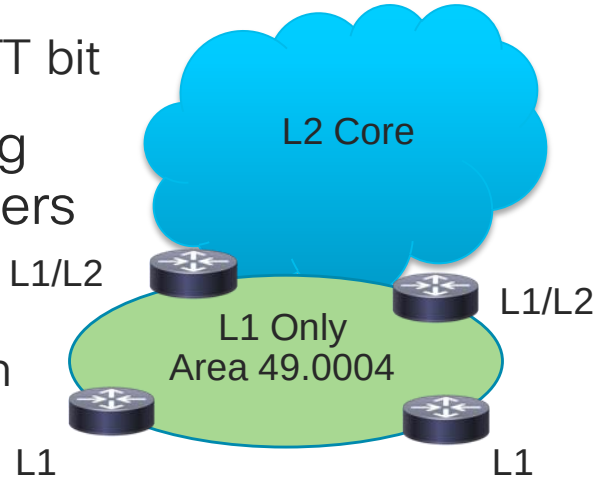

L1-L2 Router at Edge of POP

Route-Leaking

- It is recommended to configure the L1-L2 routers at the edge of the pop with route-leaking capabilities
- Leak BGP next-hops and summarize physical link
- Hence the L1 routers will be able to take the right exit/entry router based on the metric of the leaked IP-prefix
  - Optimal Inter-Area Routing
- Ensure ‘metric-style wide’ is configured when leaking routes e.g. MPLS-VPN (PEs Loopback Reachability and LSP binding)

Sup-Optimal IS-IS Routing - Overload Bit

Overload Bit

Normal traffic flow between XR1 and R4 would be between XR2 and R5 based on metric calculations

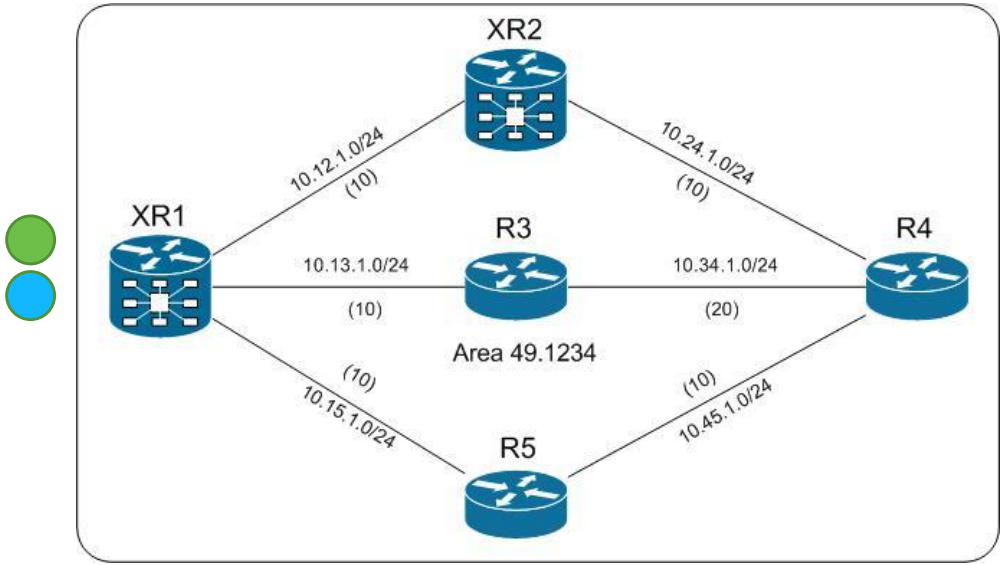

Overload Bit

Traffic flow taken across links that have higher metric are not normal.

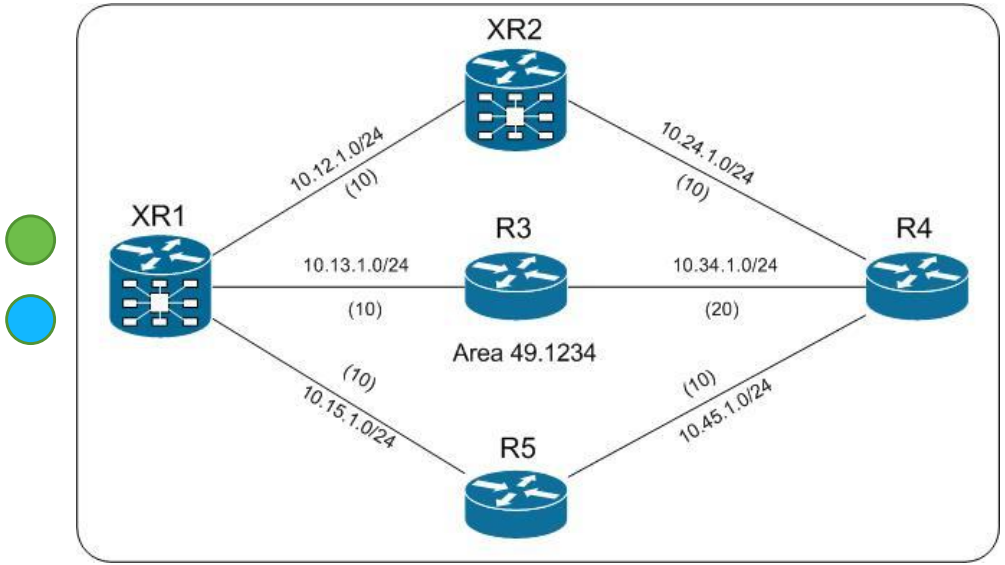

Overload Bit

- The overload bit indicates a router in an overloaded condition.
- Routers avoid sending traffic through routers that set the overload bit.
- Upon recovery, the router advertises a new LSP without the overload bit, and the SPF calculation occurs normally without avoiding routes through the previously overloaded node.

Overload Bit

```
RP/0/0/CPU0:XR1#show isis database
IS-IS ISIS (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
XR1.00-00           * 0x00000007   0x71d6        1046          0/0/0
XR2.00-00             0x0000000c   0x2557        1124          0/0/1
R3.00-00              0x00000009   0x5564        1031          0/0/0
R4.00-00              0x0000000c   0x8baa        1065          0/0/0
R5.00-00              0x00000009   0xa406        1155          0/0/1
R5.03-00              0x00000003   0x7ccc        1124          0/0/0
```

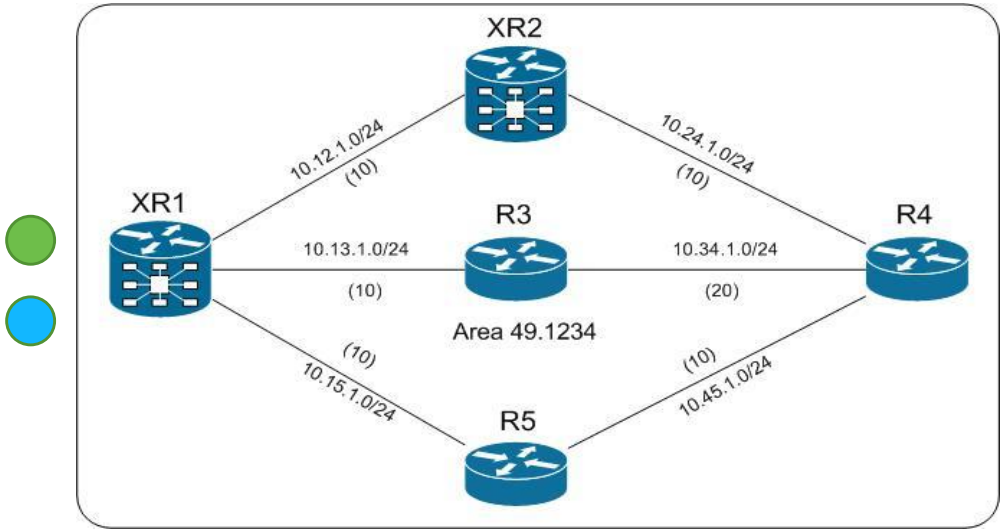

Overload Bit

- Originally, the overload bit signified memory exhaustion, but current routers have a significant amount of memory making those situations very rare.
- Setting the overload bit on a router during maintenance windows is a common technique to route traffic around the nodes being worked on.
- Newer IS-IS functionality allows a router to set the overload bit when it first starts up for a specific amount of time, or until BGP sessions have stabilized.

Sup-Optimal IS-IS Routing

– Other settings

Sub-Optimal IS-IS Routing

- Check for NX-OS devices.
- They auto-cost metric settings which are set to 40 Gbps
- Statically set the interface metrics
- Be aware of paths with wide metrics and narrow metrics (i.e. transition).

IPv6 Topologies

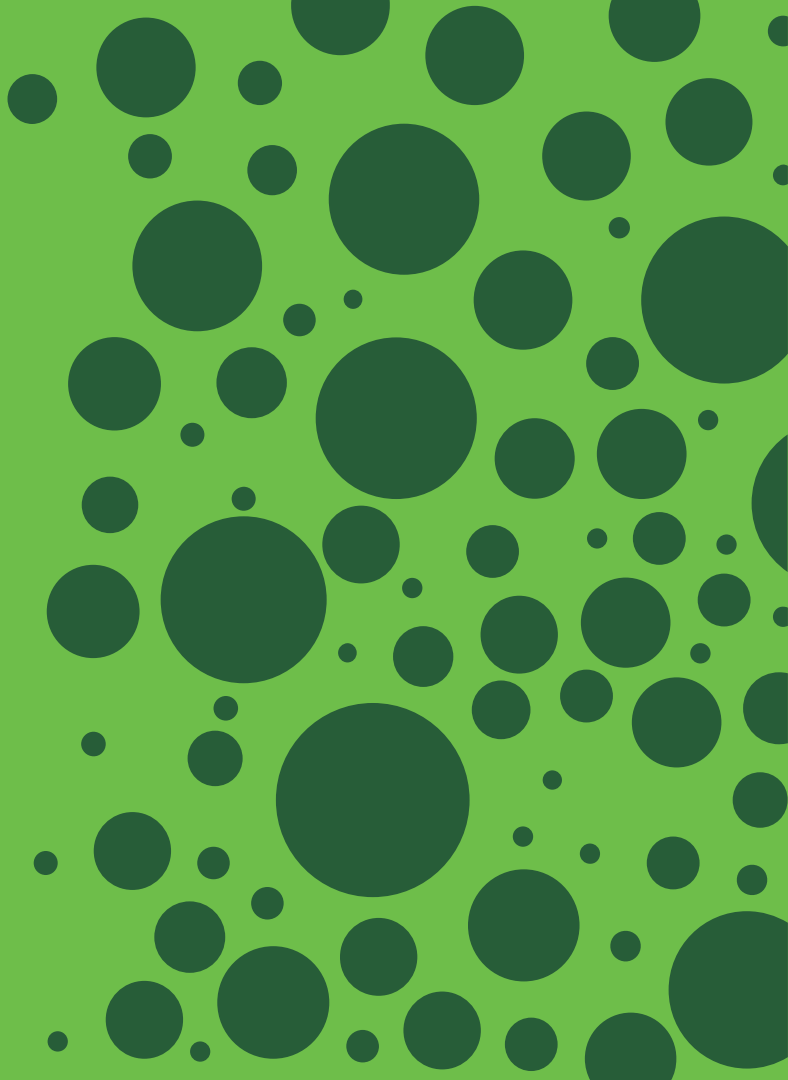

IS-IS for IPv6

- IPv6 Address Family support (RFC 2308)
- 2 new Tag/Length/Values added to introduce IPv6 routing
  - IPv6 Reachability TLV#236:
    - Equivalent to IP Internal/External Reachability TLV's
  - IPv6 Interface Address TLV #232
    - For Hello PDUs, must contain the link-local address
    - For LSP, must contain the non-link local address
- IPv6 NLPID (Network Layer Protocol Identifier) TLV#232 is advertised by IPv6 enabled routers

IS-IS for IPv6

Restrictions with Single Topology

- In Single topology IS-IS for IPv6 uses the same SPF for both IPv4 and IPv6.
  - Remember that the protocol must match for an adjacency to form? IPv4 and IPv6 topologies MUST match exactly
  - Cannot run IS-IS IPv6 on some interfaces, IS-IS IPv4 on others.
  - An IS-IS IPv6-only router will not form an adjacency with an IS-IS IPv4/IPv6 router (Exception is over L2-only interface)

IS-IS for IPv6

Multi-Topology IS-IS extensions

- Multi-Topology IS-IS solves the restrictions of Single topology
  - Two independent topology databases maintained
  - IPv4 uses Multi-Topology ID (MTID) zero(0)
  - New Multi-Topology ID (MTID #2) for IPv6
- Multi-Topology IS-IS has updated packets
  - Hello packets marked with MTID #0 or MTID #2
  - New TLV attributes introduced
  - Each LSP is marked with the corresponding MTID
- Miss-Matched MTID values
  - No effect on broadcast segments, adjacency will form
  - Point-to-point segments, adjacency will not form

IS-IS for IPv6

Choosing Single or Multi-Topology IS-IS

Use Single-Topology for:

- No planned differences in topology between IPv4 and IPv6
- Each interface has the same IPv4 and IPv6 router Level

Use Multi-Topology for:

- Incremental roll-out of IPv6 on an IPv4 topology
- If you plan for differences in topology between IPv4 and IPv6

The optional keyword transition may be used for transitioning existing IS-IS IPv6 single Topology mode to Multi-Topology IS-IS

IS-IS for IPv6

Transition to Multi-Topology IS-IS – Wide Metrics

- Ensure “Wide metric” is enabled
  - Mandatory for Multi-Topology to work
  - When migrating from narrow to wide metrics, care is required
  - Narrow and wide metrics are NOT compatible with each other

- Migration is a two stage process
  - Step 1: make use of the transition keyword

```
router isis
  metric-style transition
```

→

```
router isis
  metric-style wide
```

  - Step 2: Once the whole network is changed to transition support, the metric style can be changed to wide

IPv6 Multi-Topology IS-IS Configuration

IOS

```
R1#conf t
R1(config)#router isis CISCOLIVE
R1(config-router)#metric-style wide
R1(config-router)#address-family ipv6 unicast
R1(config-router-af)#multi-topology
```

IOS XR

```
RP/0/0/CPU0:XR2#conf t
RP/0/0/CPU0:XR2(config)#router isis CISCOLIVE
RP/0/0/CPU0:XR2(config-isis)# metric-style wide
RP/0/0/CPU0:XR2(config-isis)# address-family ipv6 unicast
RP/0/0/CPU0:XR2(config-isis-af)# multi-topology
```

IPv6 Multi-Topology IS-IS Configuration

NX-OS

```
R1#conf t
NX-3(config)# router isis CISCOLIVE
NX-3(config-router)# metric-style wide
NX-3(config-router)# address-family ipv6 unicast
NX-3(config-router)# multi-topology
```

IS-IS with Disjunct IPv6 Networks

- Cannot join two IPv6 areas via an IPv4-only area
  - L2 adjacencies will form OK
  - IPv6 traffic will black-hole in the IPv4 area.

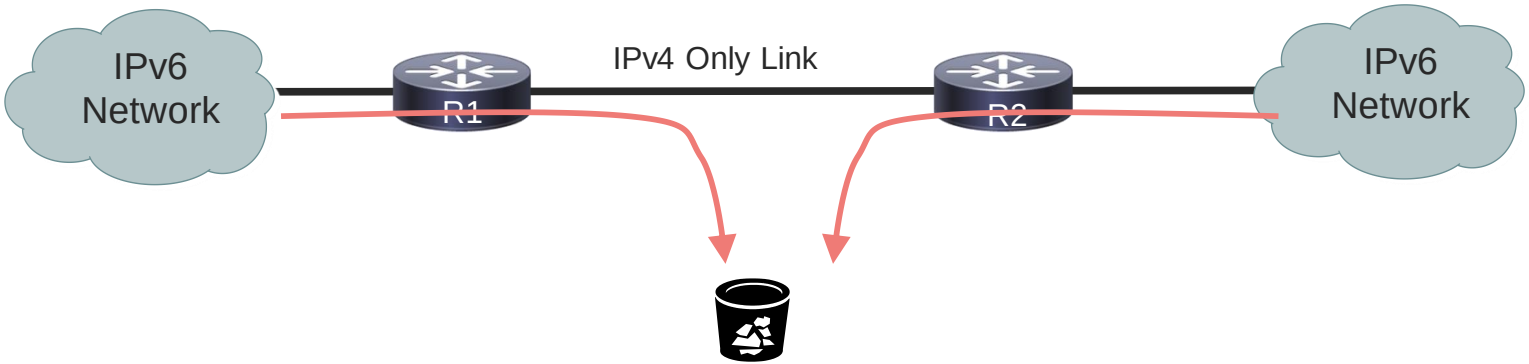

IS-IS with Disjunct IPv6 Networks

- Cannot join two IPv6 areas via an IPv4-only area
  - L2 adjacencies will form OK
  - IPv6 traffic will black-hole in the IPv4 area.

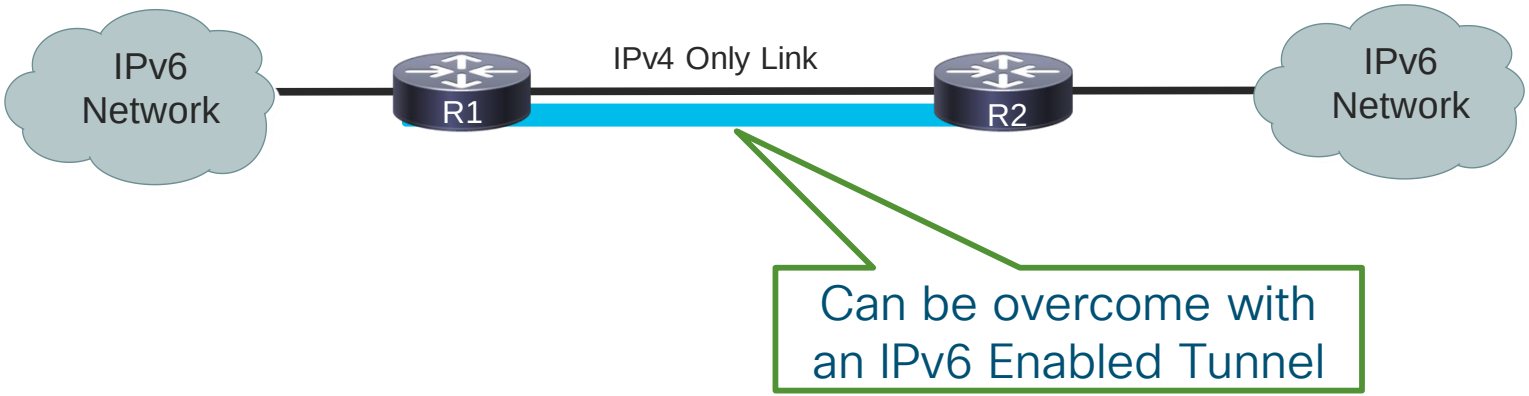

Summary

Troubleshooting Summary

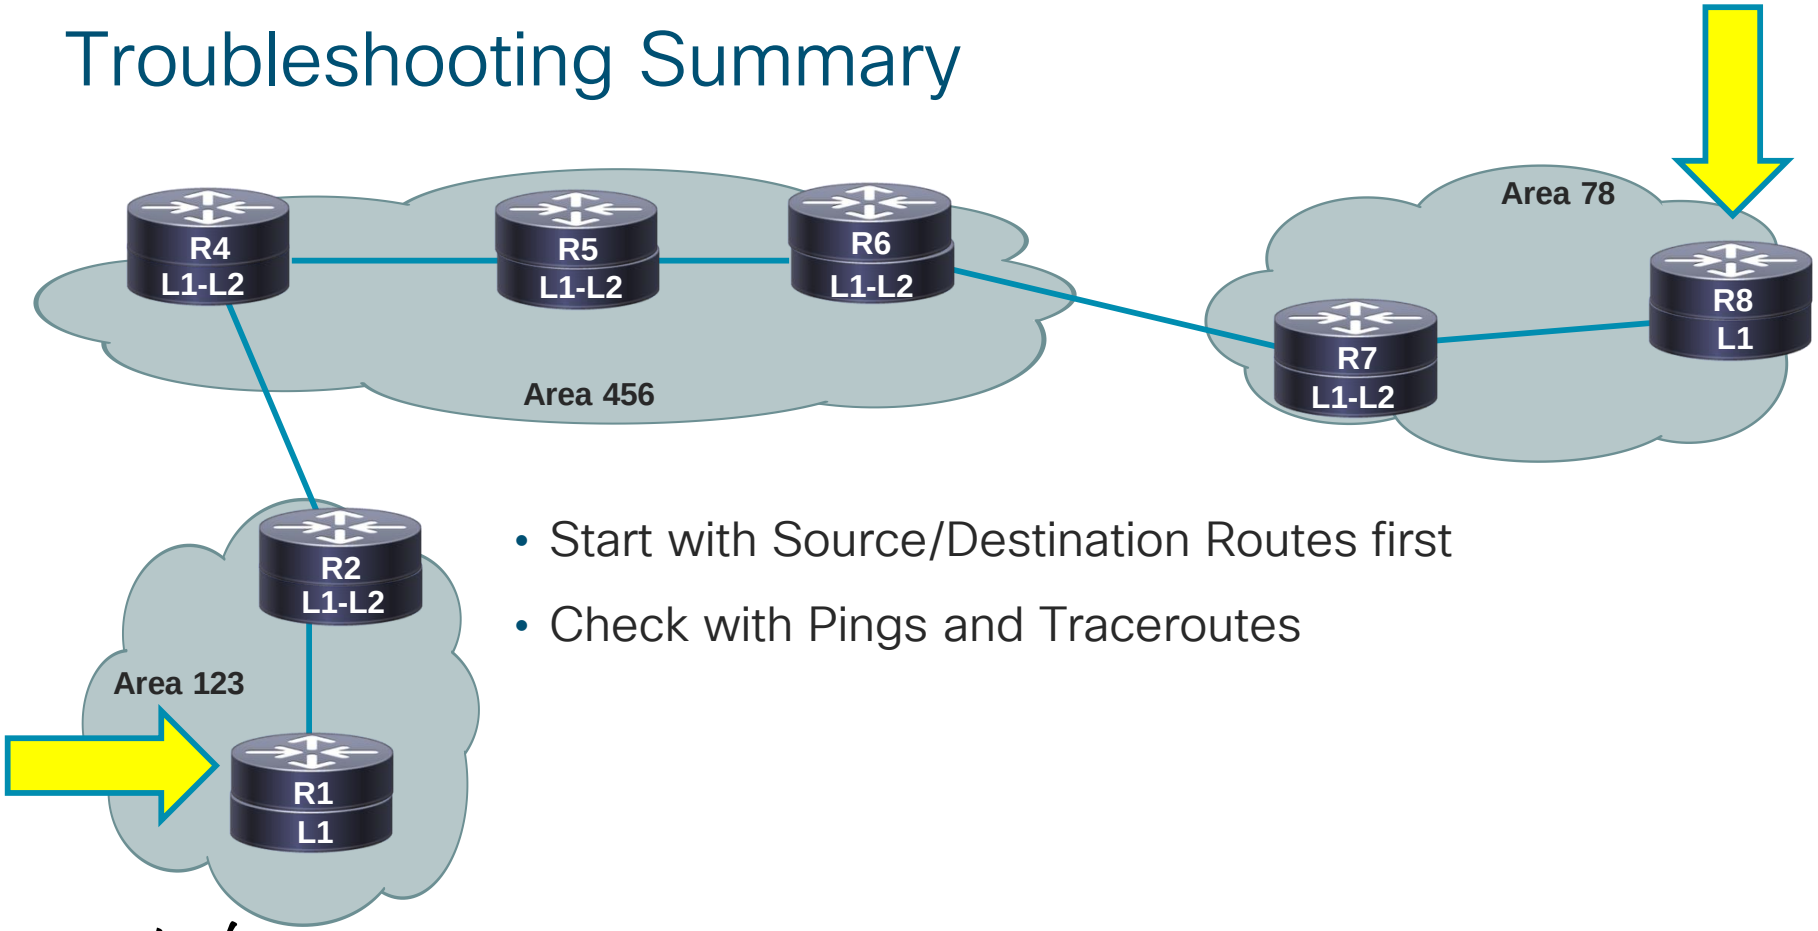

- Start with Source/Destination Routes first
- Check with Pings and Traceroutes

Troubleshooting Summary

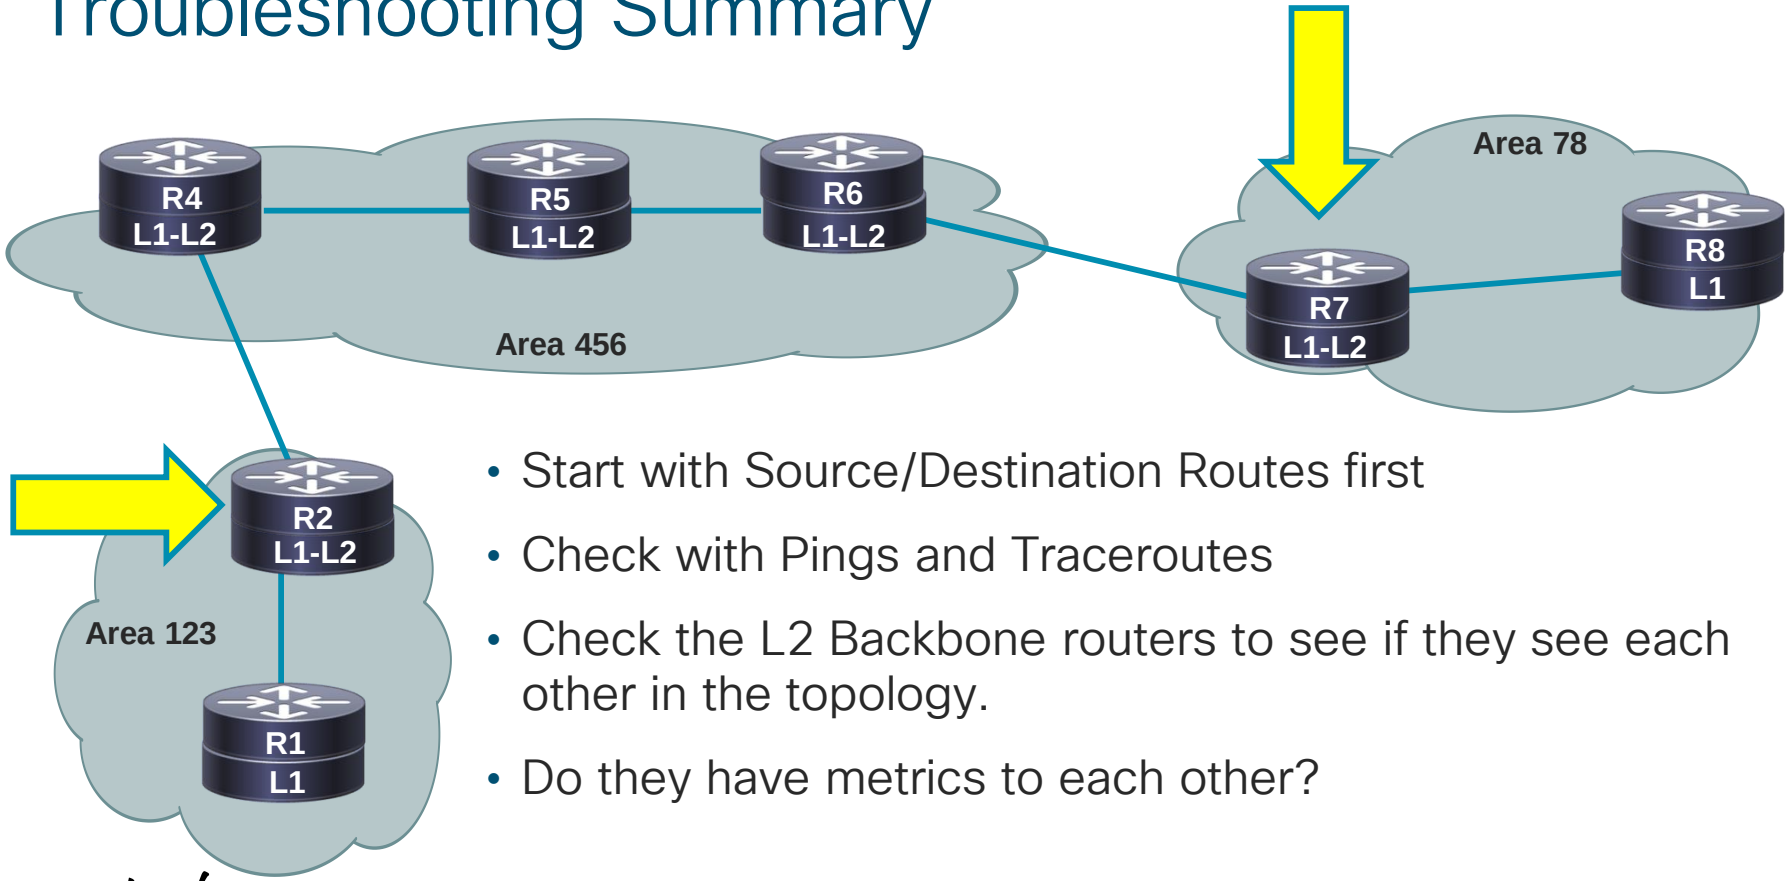

- Start with Source/Destination Routes first
- Check with Pings and Traceroutes
- Check the L2 Backbone routers to see if they see each other in the topology.
- Do they have metrics to each other?

Troubleshooting Summary

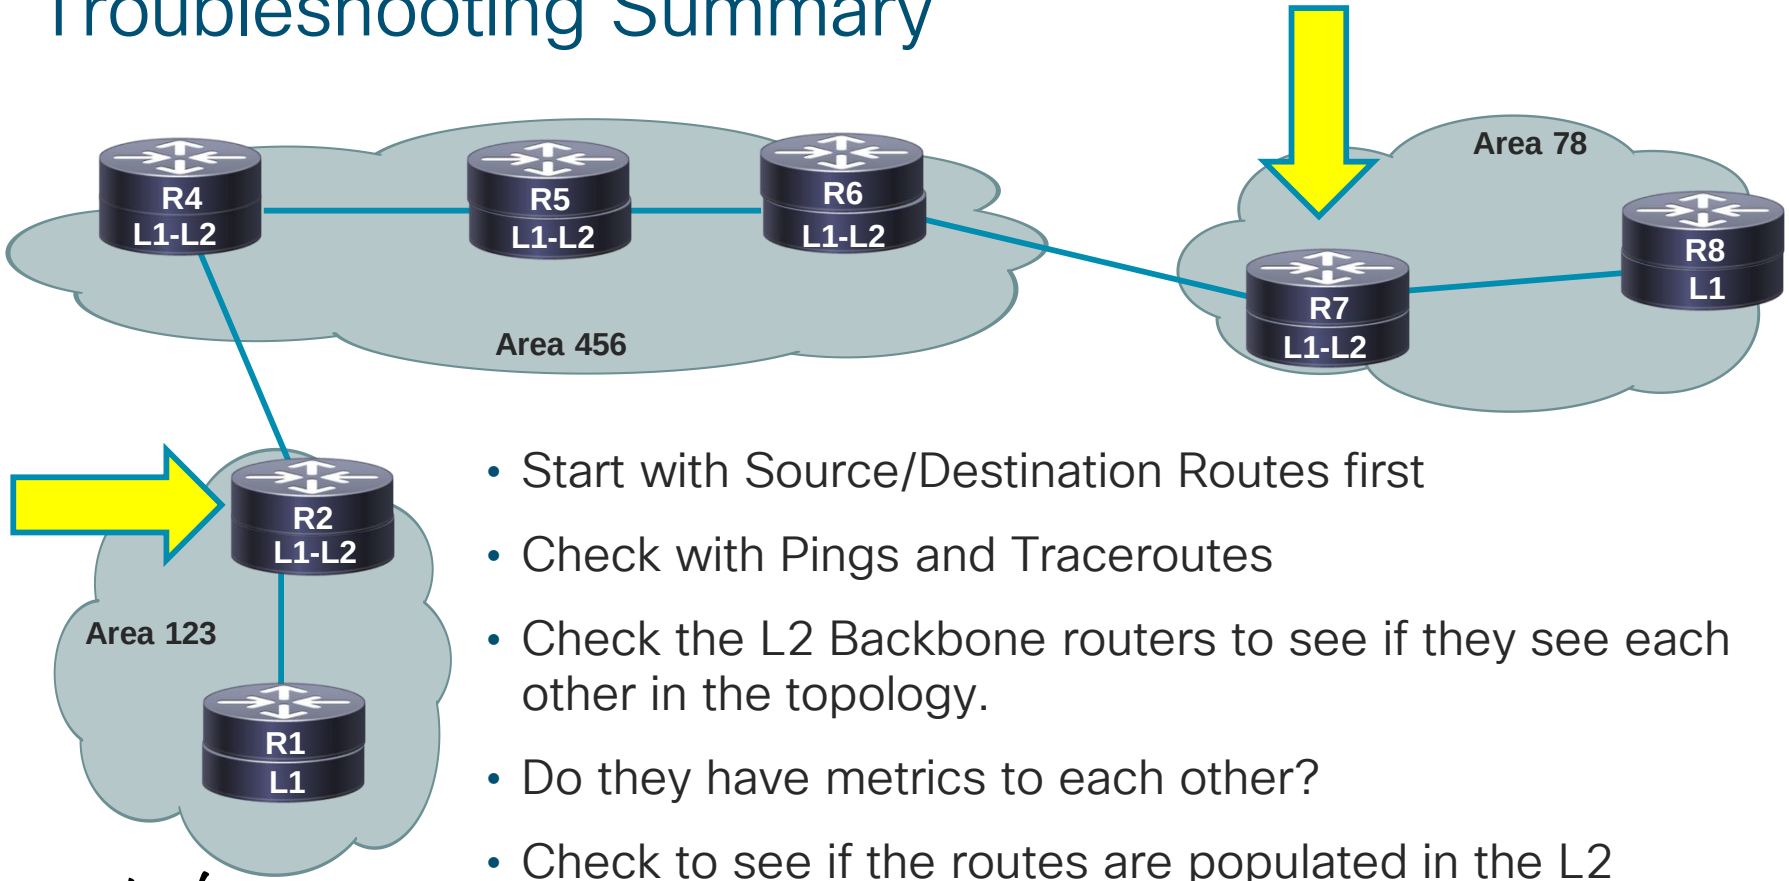

- Start with Source/Destination Routes first
- Check with Pings and Traceroutes
- Check the L2 Backbone routers to see if they see each other in the topology.
- Do they have metrics to each other?
- Check to see if the routes are populated in the L2 LSPDB

Troubleshooting IS-IS Commands (IOS & IOS XR)

IOS

```
show isis neighbors
Show clns interface
Show run | sect router isis
Debug isis adj-packets
Show isis protocol
Show isis topology
show isis database [LSP-ID] [level-1|level-2] [detail]
```

IOS XR

```
show isis neighbors
Show isis interface
Show run router isis
Show isis trace all reverse
Show isis protocol
Show isis topology
show isis database [LSP-ID] [level-1|level-2] [detail]
```

Troubleshooting IS-IS Commands (NX-OS)

NX-OS

```
show isis adjacency
Show isis interface
Show run isis
ethanalyzer local interface inband capture-filter "ether host 01:80:c2:00:00:15"
Show isis event-history
Debug isis iih
Show isis protocol
Show isis topology
show isis database [LSP-ID] [level-1|level-2] [detail]
```

Books for IS-IS

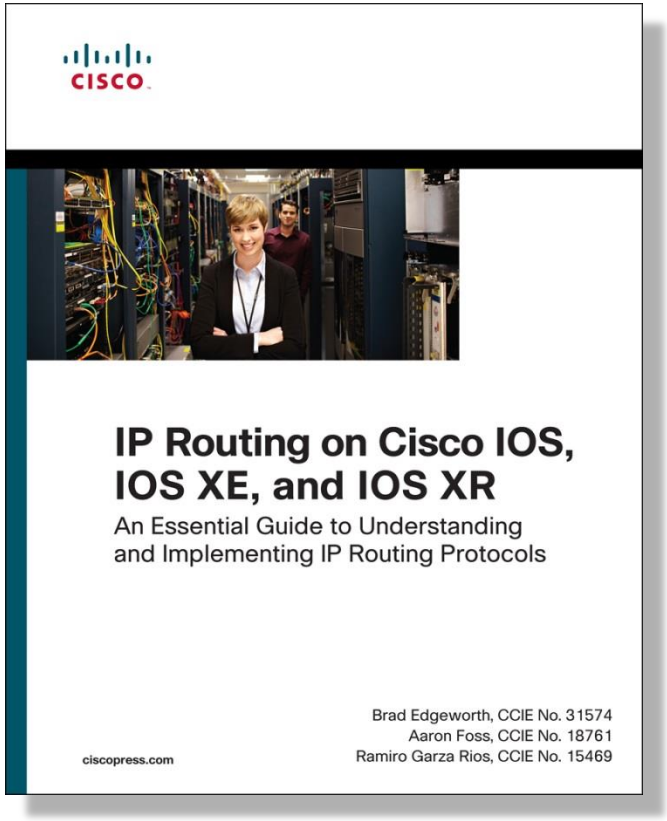

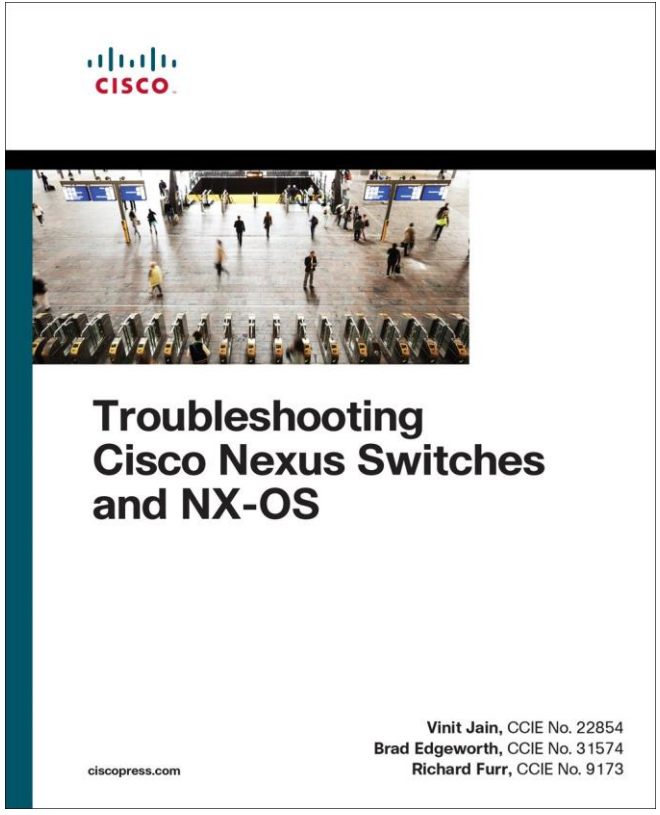

Complete your online session evaluation

Give us your feedback to be entered into a Daily Survey Drawing.

Complete your session surveys through the Cisco Live mobile app or on www.CiscoLive.com/us.

Don't forget: Cisco Live sessions will be available for viewing on demand after the event at www.CiscoLive.com/Online.

Continue your education

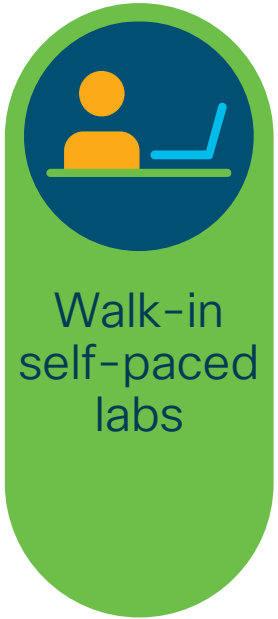

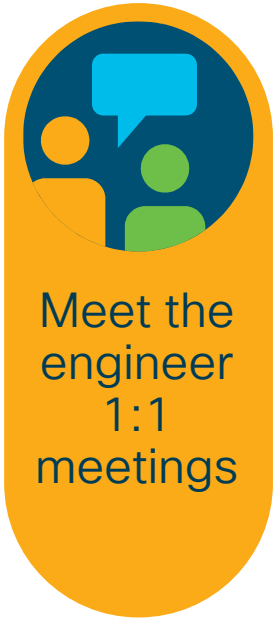

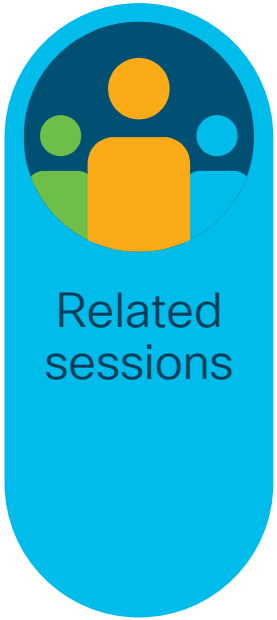

Thank you

INTUITIVE