Equipping Today's Instructors for Tomorrow's Students

Academy Conference 2013

Cisco Networking Academy

What I need to know about IPv6 to teach CCNA2 Routing Protocols

Rick Graziani
CS/CIS Instructor
Cabrillo College

Before we begin….

- Our virtual audience can submit questions via the Q & A window anytime during the presentation
- Simply type your question in the Q & A window and click ‘Submit’ - subject matter experts will be answering questions live during the presentation

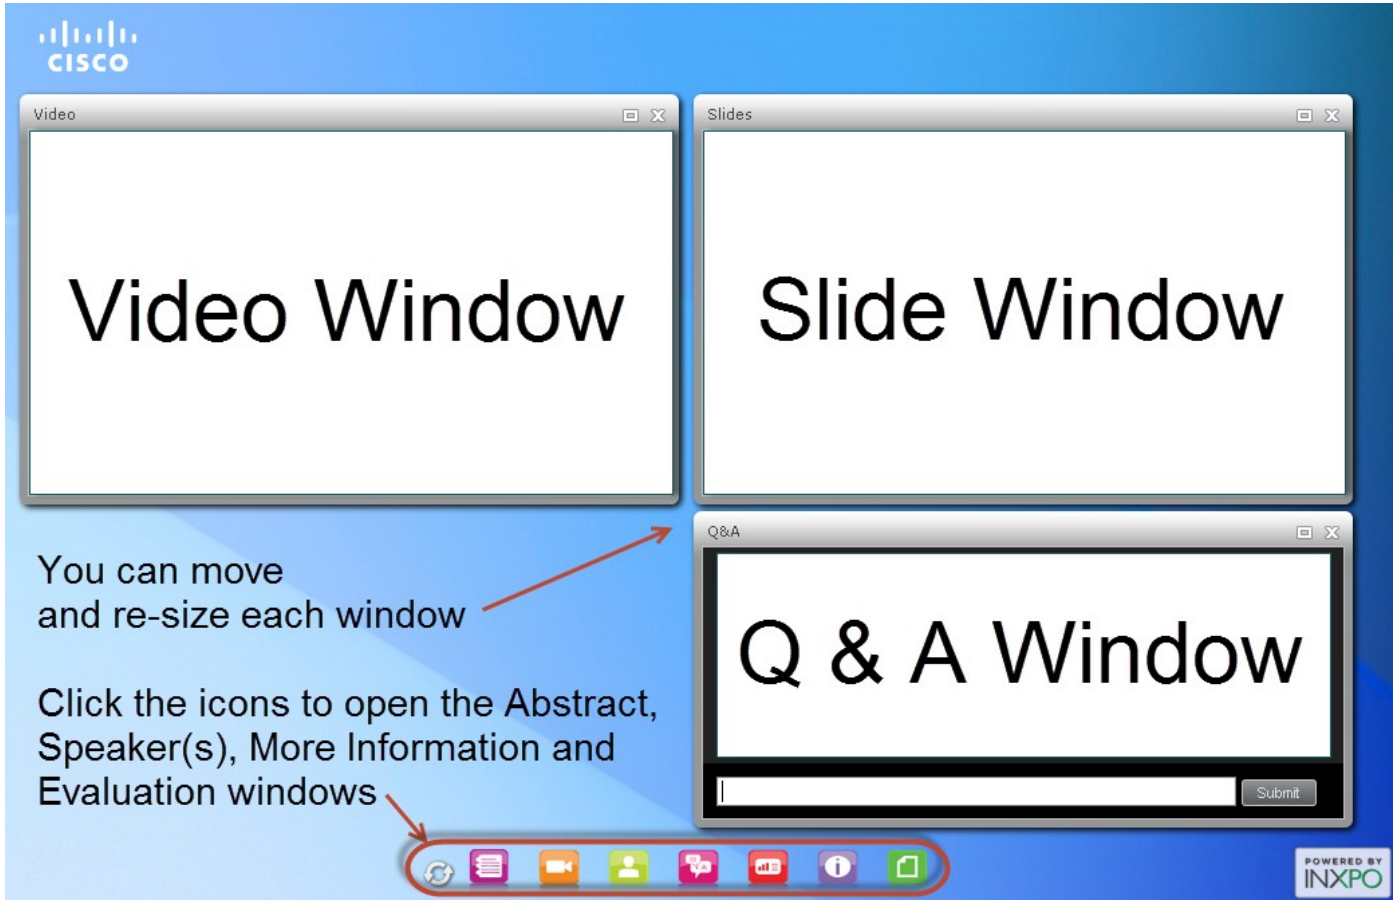

Who am I?

- Rick Graziani - graziani@cabrillo.edu
- CS/CIS instructor at Cabrillo College, Santa Cruz, California
- Working in IT since 1980
- Cisco Networking Academy instructor since 1997
- Practice what I preach…
  - Implementing native IPv6 at Cabrillo College
  - Home: Run native IPv6 (and IPv4) to the Internet
- Curriculum Development Team for Cisco Networking Academy
- When not working, hopefully I'm surfing

Topics

- Quick review of IPv6
  - For more information see my CCNA 1 IPv6 Presentation
- IPv6 Static Routes
- EIGRP for IPv6
- OSPFv3
- Multi-Area OSPFv3
- IPv6 Access Control Lists

So we can finish, please hold questions until the end

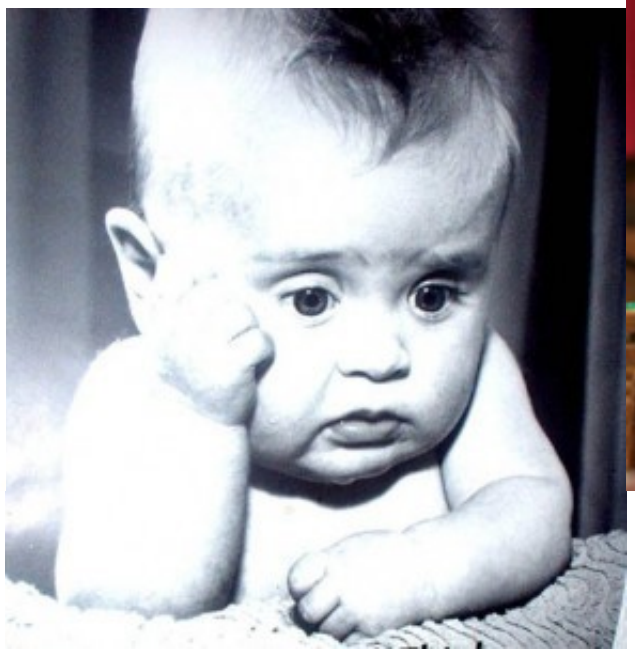

Works for Both Routing and Routing & Switching Courses

Routing Protocols

R&S Part 1

R&S Part 2

CCNA 2 – Routing Protocols

There is a lot of new information in CCNA 2 besides IPv6!

- Chapter 1: Routing Concepts → Basic IPv6 address configuration and verification
- Chapter 2: **Static Routing** → We will review when introduced
- Chapter 3: Routing Dynamically → Specific IPv6 info in each chapter
- Chapter 4: **EIGRP**
- Chapter 5: **Advanced EIGRP**
- Chapter 6: **Single Area OSPF**
- Chapter 7: **Advanced Single Area OSPF**
- Chapter 8: **Multi-Area OSPF**
- Chapter 9: **Access Control Lists**
- Chapter 10: IOS File Management → No IPv6

Quick Review of IPv6 Addresses

IPv6 Address Notation

Dec.	Hex.	Binary	Dec.	Hex.	Binary
0	0	0000	8	8	1000
1	1	0001	9	9	1001
2	2	0010	10	A	1010
3	3	0011	11	B	1011
4	4	0100	12	C	1100
5	5	0101	13	D	1101
6	6	0110	14	E	1110
7	7	0111	15	F	1111

One Hex digit = 4 bits

2001:0DB8:AAAA:1111:0000:0000:0000:0100/64

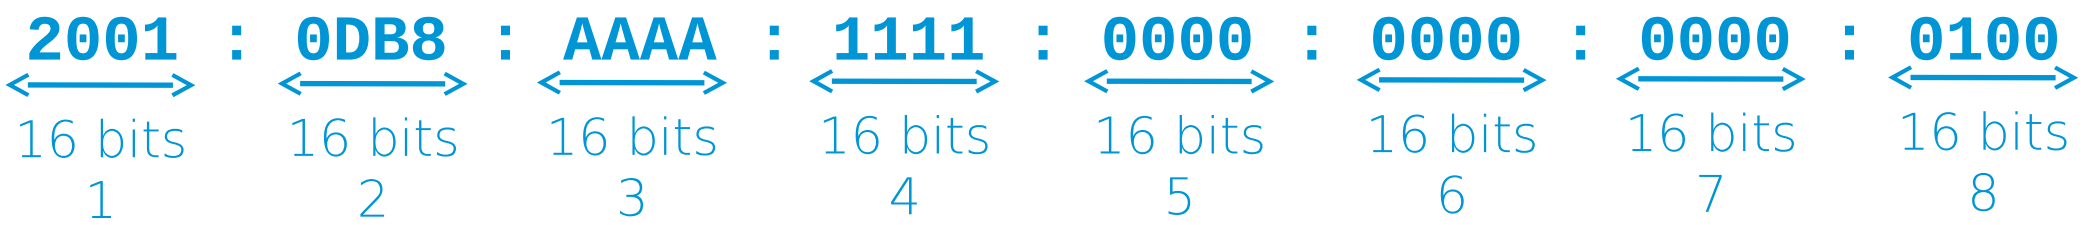

- IPv6 addresses are 128-bit addresses represented in:
  - Eight 16-bit segments or “hextets” (not a formal term)
  - Hexadecimal (non-case sensitive) between 0000 and FFFF
  - Separated by colons

Compressed Notation

- 1st Rule: Leading zeroes in any 16-bit segment do not have to be written
- 2nd Rule: Any single, contiguous string of one or more 16-bit segments consisting of all zeroes can be represented with a double colon

2001 : 0DB8 : AAAA : 1111 : ~~0000 : 0000 : 0000~~ : ~~0~~100

Second Rule First Rule

2001 : 0DB8 : AAAA : 1111 :: 100

Focus on: Global Unicast (GUA) and Link-Local Unicast

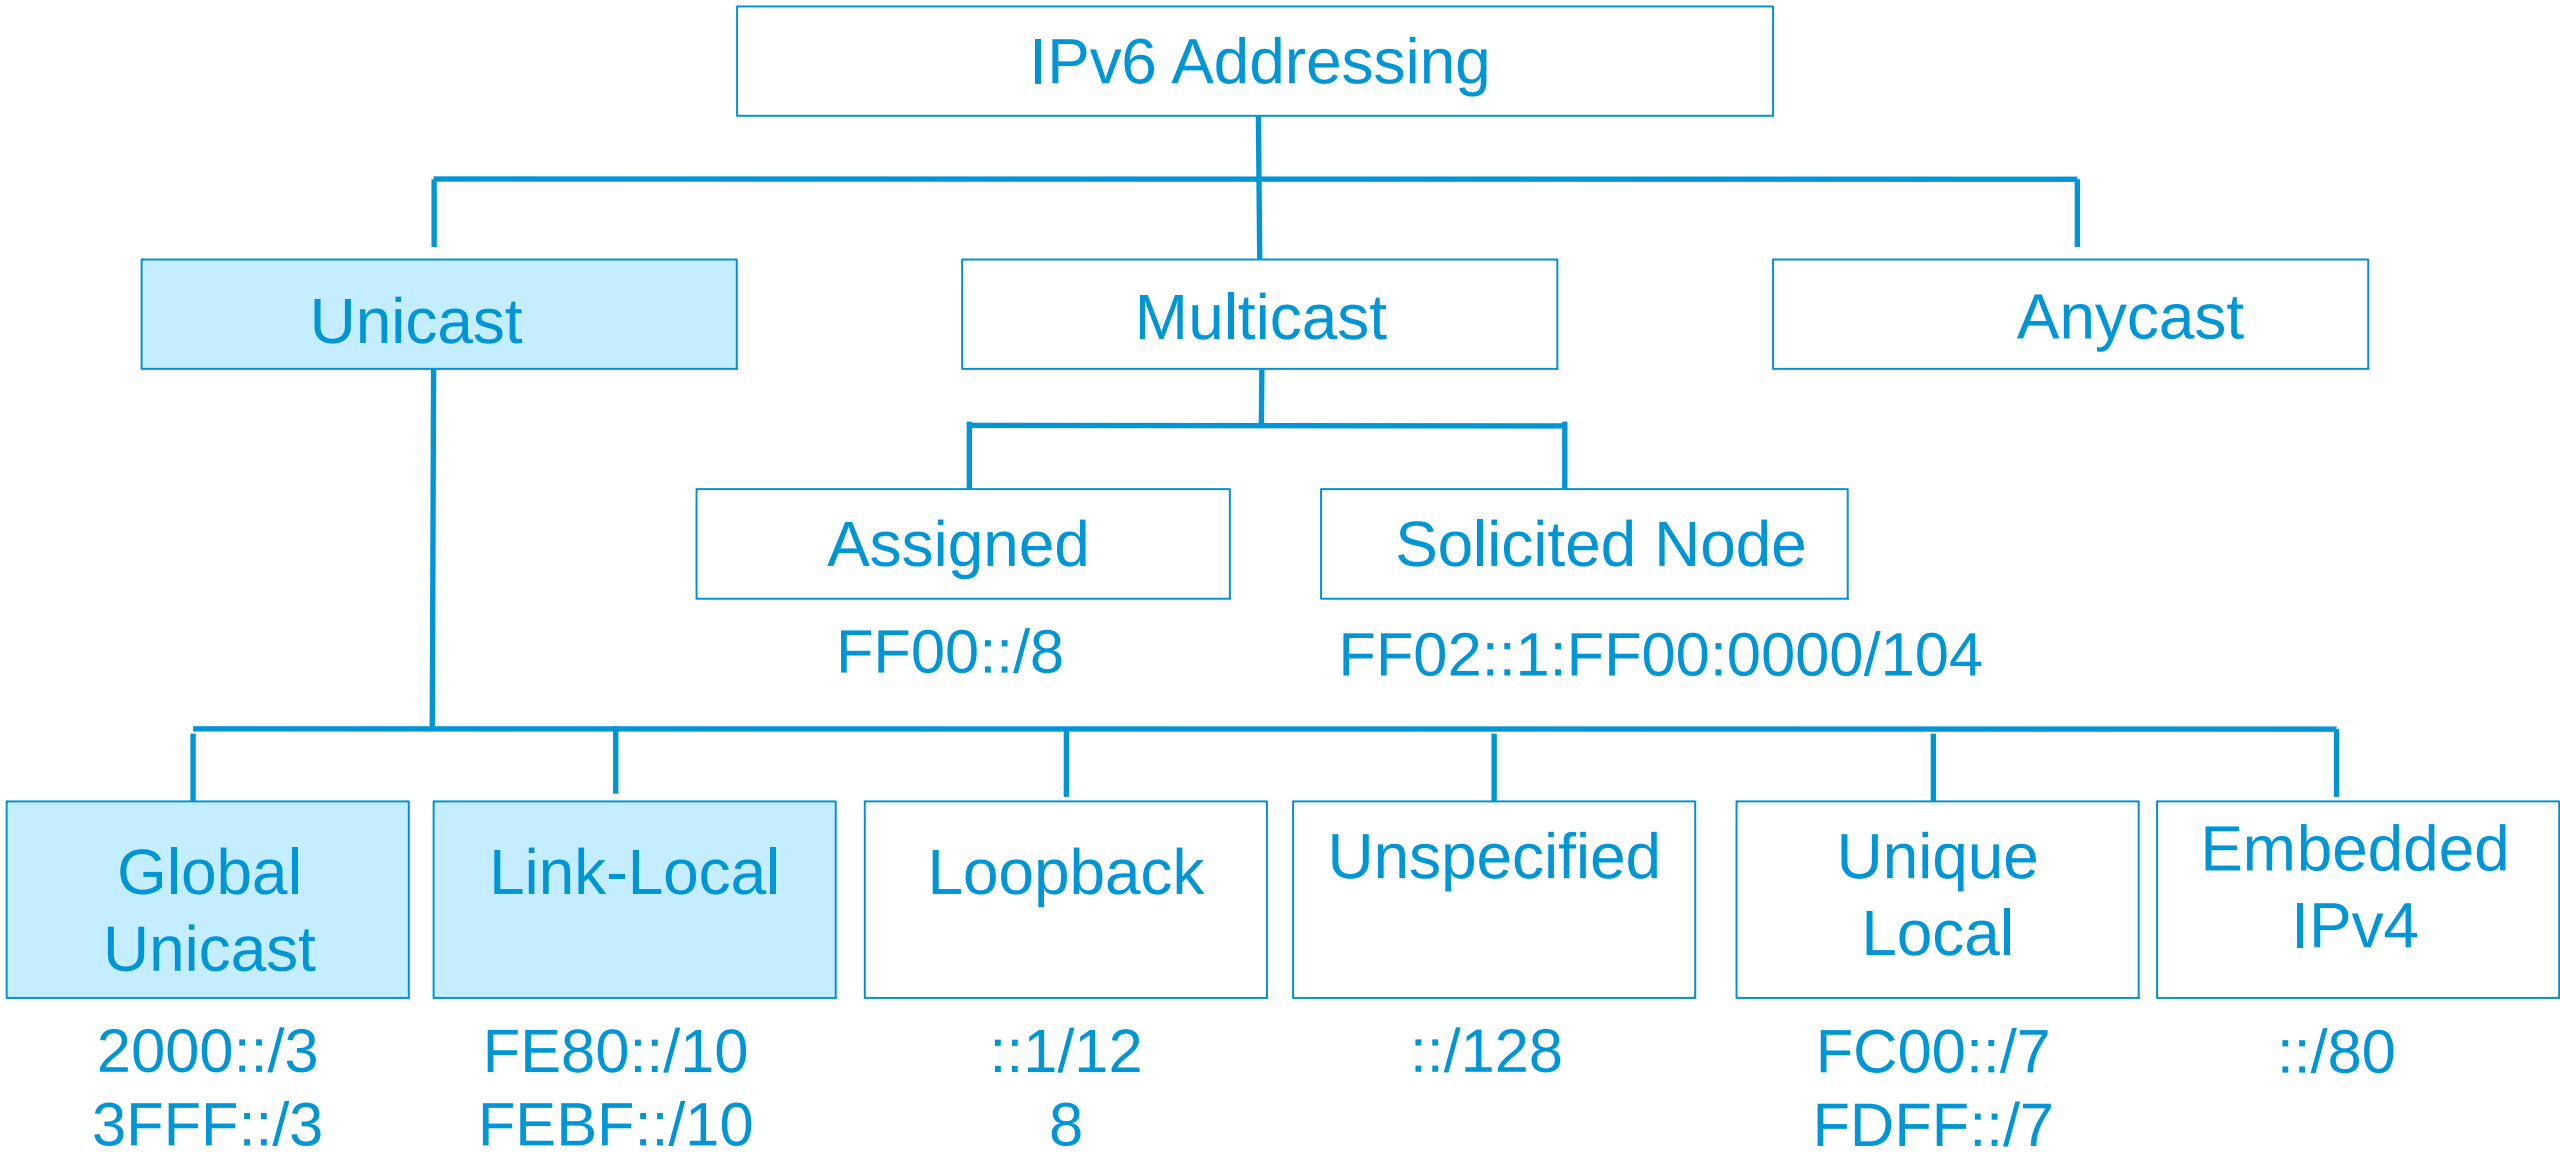

Global Unicast Address (GUA)

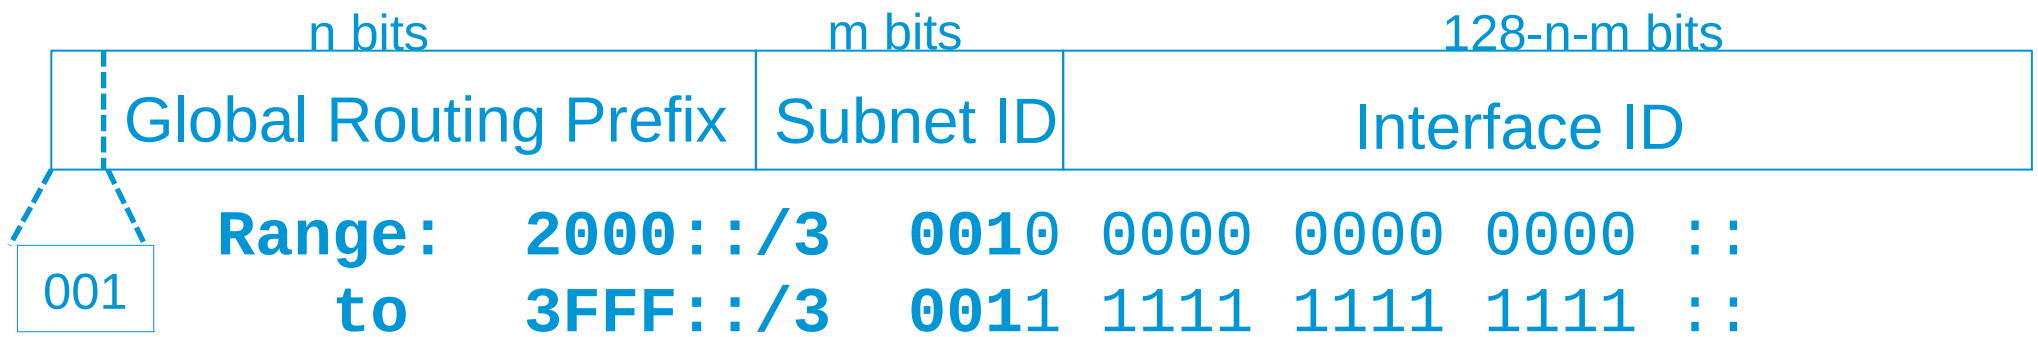

- Global unicast addresses are similar to IPv4 addresses
  - Equivalent to IPv4 public addresses
  - Except under very specific circumstances, all end users will have a global unicast address
  - Routable
  - Unique
  - Terminology:
  - **Prefix** equivalent to *network address*
  - **Prefix length** equivalent to *subnet mask in IPv4*
  - **Interface ID** equivalent to *host portion*

Typical GUA and Why We Love IPv6!

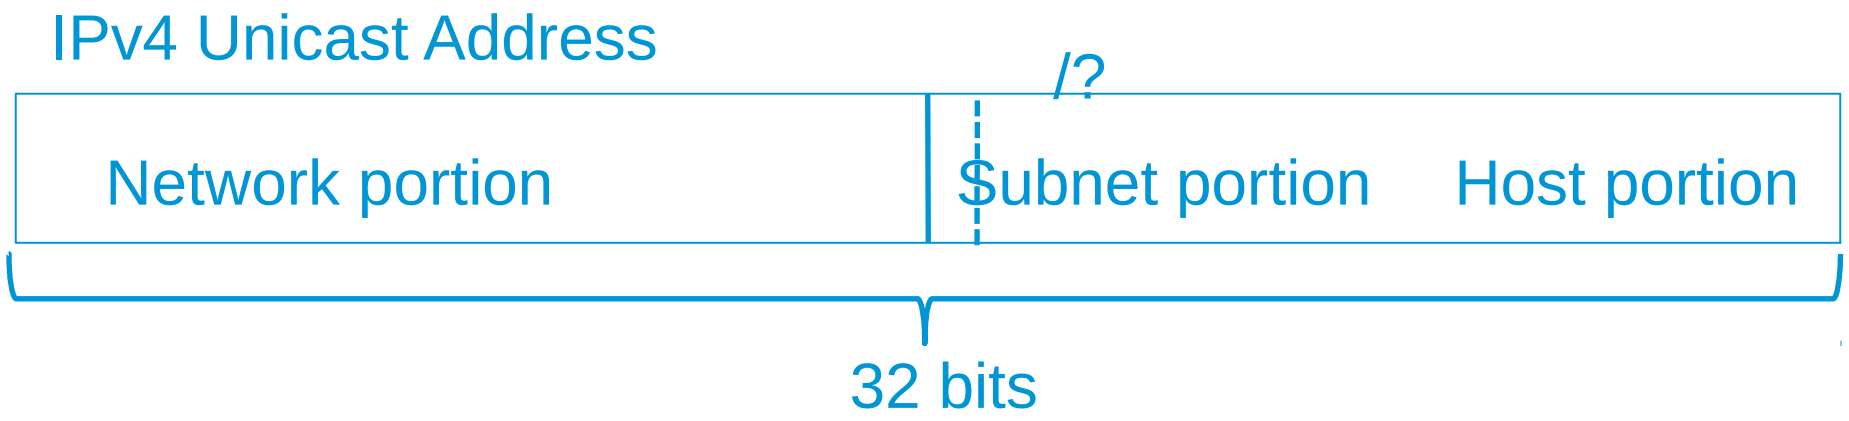

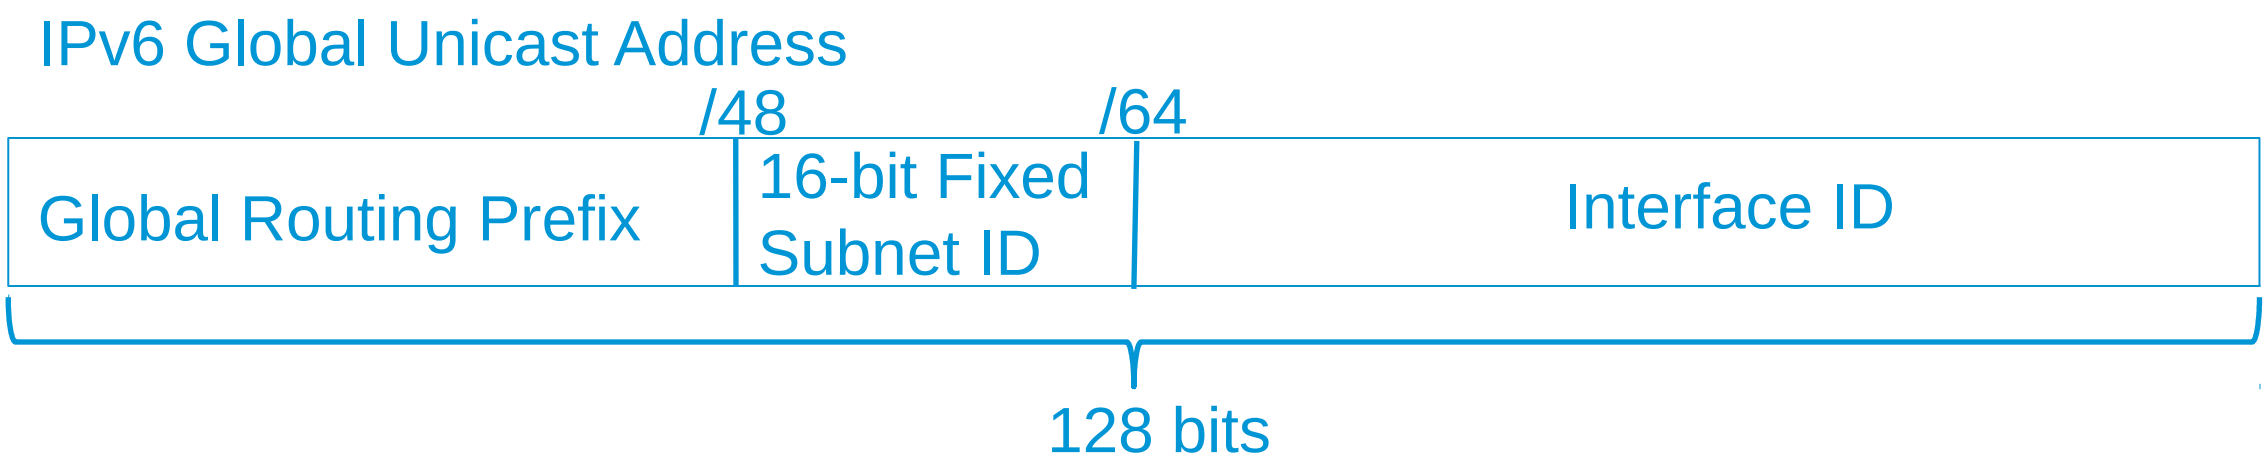

- 16-bit Subnet ID = 65,536 subnets
- 64-bit Interface ID = 18 quintillion (18,446,744,073,709,551,616) devices/subnet

/64 Global Unicast Addresses and the 3-1-4 rule

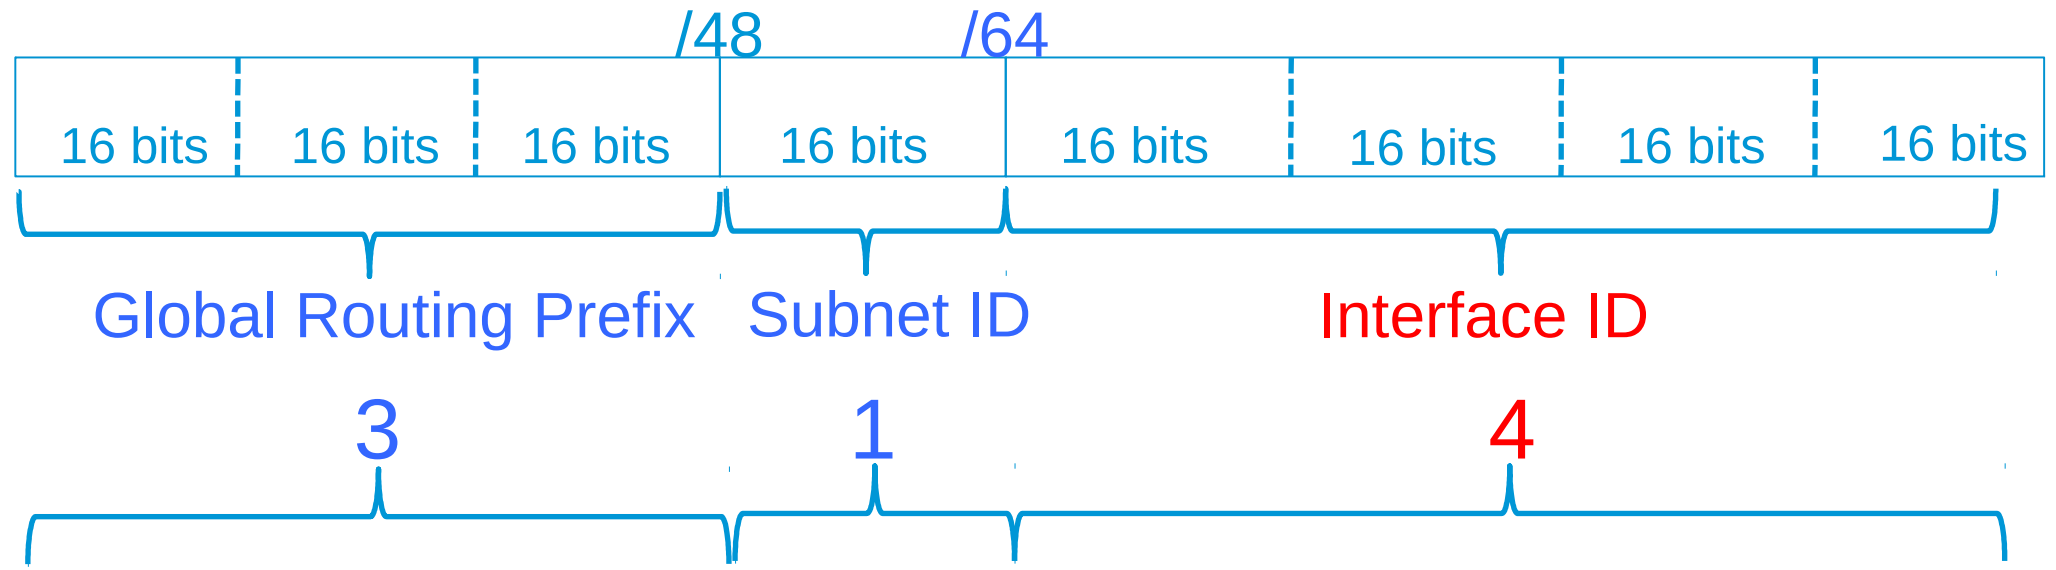

2001 : 0DB8 : AAAA : 1111 : 0000 : 0000 : 0000 : 0100

3 + 1 = 4 (/64) : 4

2001:0DB8:AAAA:1111:0000:0000:0000:0100/64

2001:0DB8:AAAA:1111::100/64

Subnetting IPv6 and Why Our Students Will Love IPv6

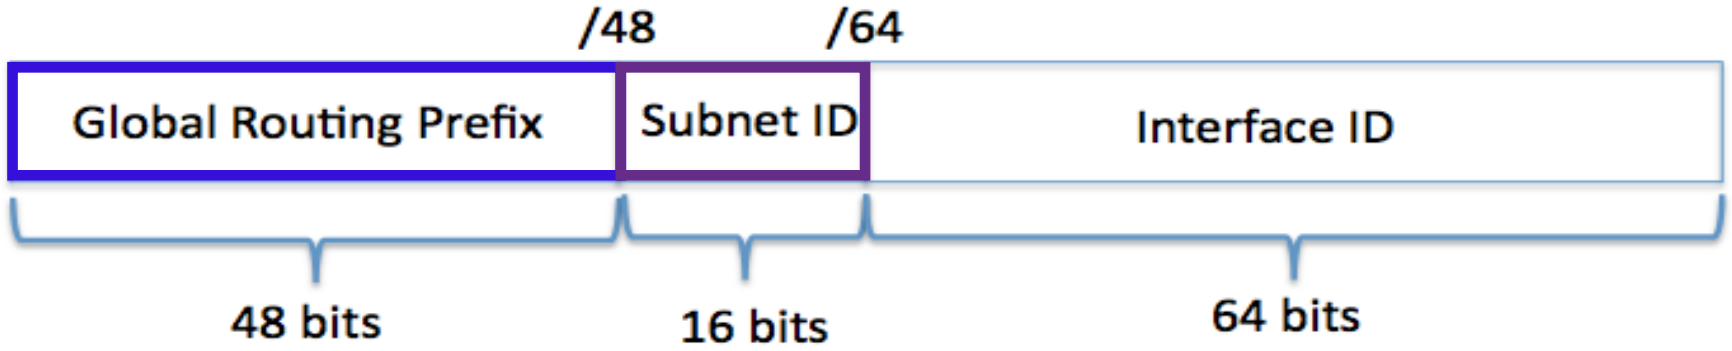

- Just increment by 1 in Hexadecimal:
  - **2001:0DB8:AAAA:0000::/64**
  - **2001:0DB8:AAAA:0001::/64**
  - **2001:0DB8:AAAA:0002::/64**
  - **2001:0DB8:AAAA:000A::/64**
- Valid abbreviation is to remove the 3 leading 0's from the first shown quartet
  - **2001:0DB8:AAAA:1::/64**

Link-Local Unicast

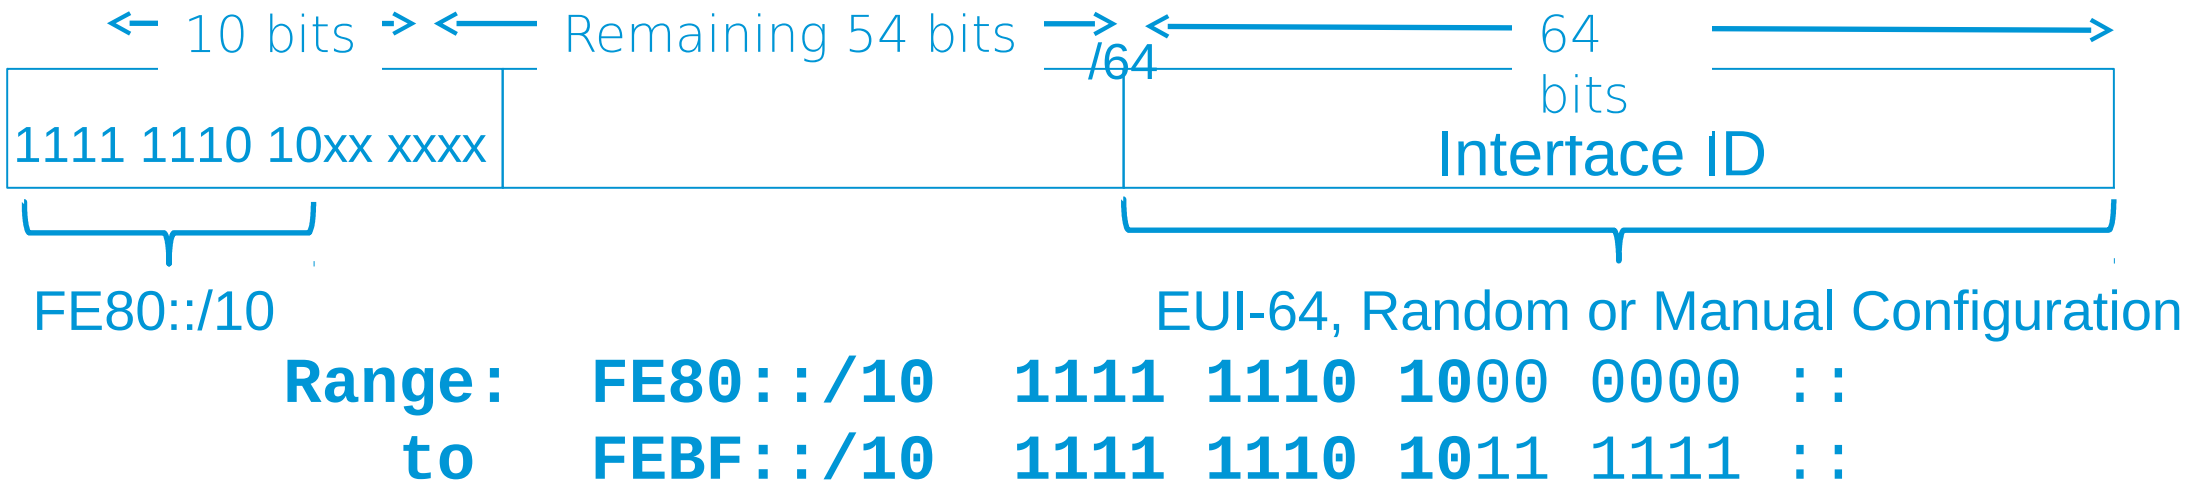

Range: FE80::/10 1111 1110 1000 0000 ::
to FEBF::/10 1111 1110 1011 1111 ::

- Used to communicate with other devices on the link (network)
- NOT routable off the link
- An IPv6 device must have at least a link-local address
- Used by:
  - Hosts to communicate to the IPv6 network before it has a global unicast address
  - Used as the default gateway address by hosts
  - Adjacent routers to exchange routing updates

ipv6 enable command

```
Router(config)# interface fastethernet 0/1
Router(config-if)# ipv6 enable
Router(config-if)# end
Router# show ipv6 interface brief
FastEthernet0/1                [up/up]
    FE80::20C:30FF:FE10:92E1
Router#
```

Link-local unicast address only

- Link-local addresses are automatically created whenever a global unicast address is configured
- The **ipv6 enable** command will:
  - Create a link-local address when there is no global unicast address
  - Maintain the link-local address even when the global unicast address is removed

Ping Link-local Address

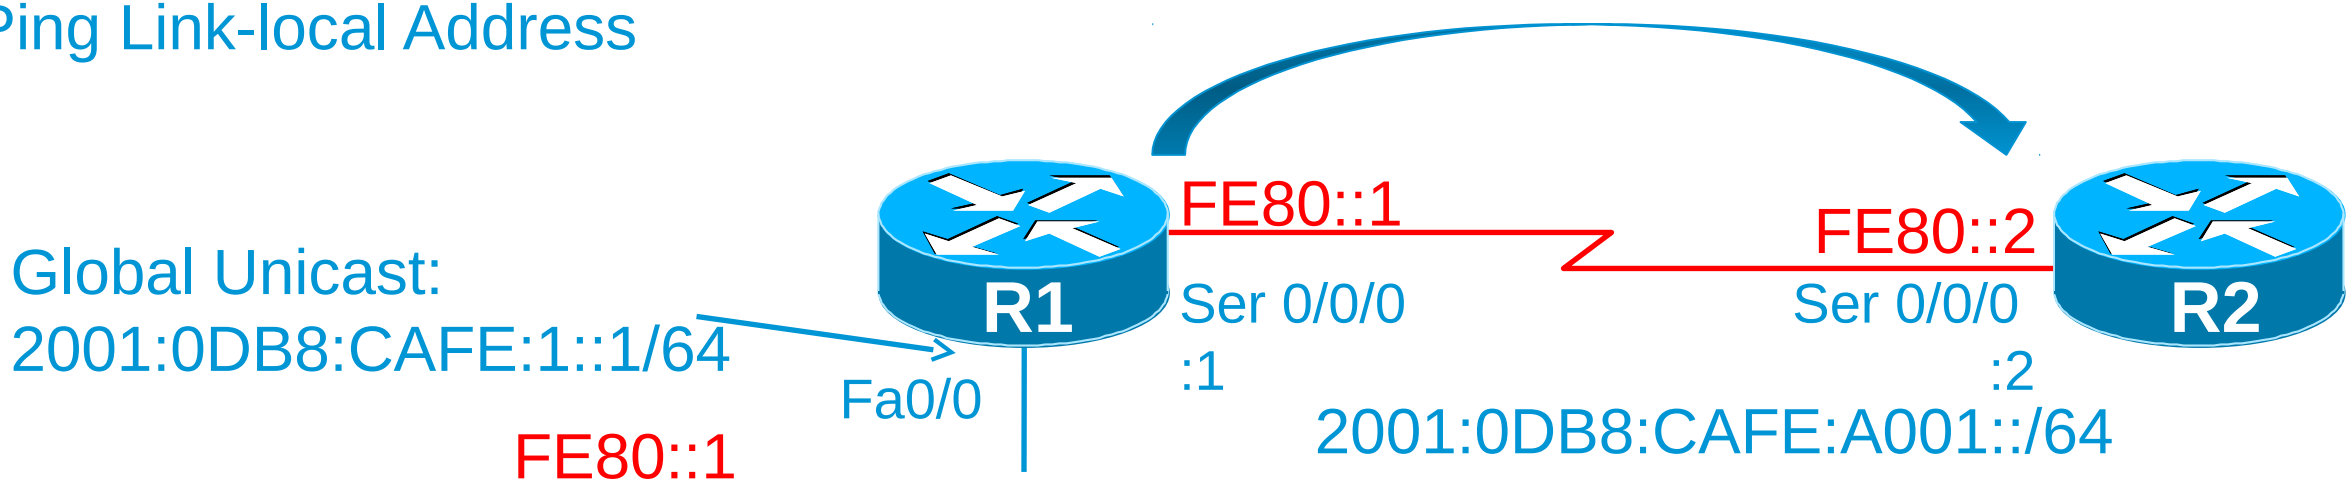


```
R1# ping fe80::2
Output Interface: ser 0/0/0
% Invalid interface. Use full interface name without
   spaces (e.g. Serial0/1)
Output Interface: serial0/0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::2, timeout is 2
   secs:
!!!!!
```

Must include exit-interface

IPv6 Static Routes

Making a Router an IPv6 Router

```
Router(config)# ipv6 unicast-routing
```

- A router's interfaces can be enabled (configured with an IPv6 address) for IPv6 like any other device on the network
- For the router to "act" as an IPv6 router it must be enabled with the **`ipv6-unicast routing`** command
- This enables the router to:
  - Send ICMPv6 Router Advertisement messages
  - Enable the forwarding of IPv6 packets
  - Configure static routing and participate in IPv6 routing protocols (EIGRP for IPv6, OSPFv3)
- Similar to "old" **`ip routing`** command for IPv4 which is enabled by default

IPv6 Static Routes

```
Router(config)#ipv6 route ipv6-prefix/ipv6-prefix-length
                          {ipv6-address | exit-interface}
```

- Static Route with a next hop IPv6 address:

```
Router(config)# ipv6 route 2001:db8:acad:2::/64 2001:db8:feed::1
```

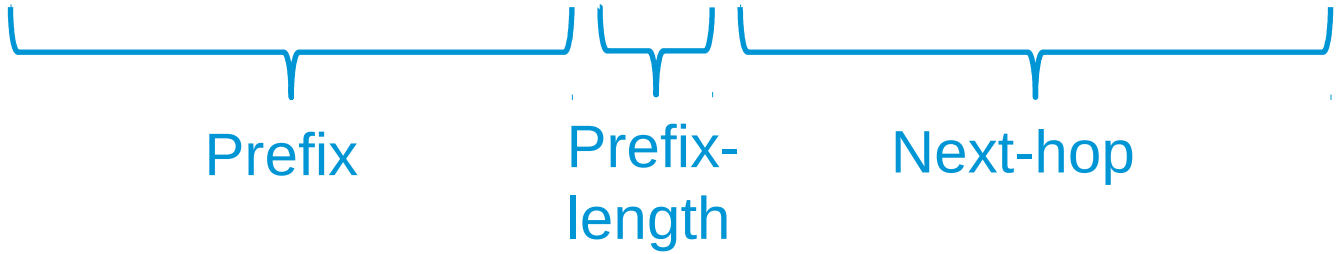

Note: Static routes using only an exit interface on point-to-point networks are common, however the use of the default CEF forwarding mechanism makes this practice unnecessary

For reasons, beyond the scope of this presentation, there are advantages to using a static route with a next-hop address

IPv6 Static Route

- Static Route with an exit interface (unnecessary with CEF enabled):

```
R1(config)# ipv6 route 2001:db8:acad:2::/64 g0/0
```

- A fully specified static route includes an exit interface and the next hop address (unnecessary with CEF enabled except when using a next hop link-local address):

```
R1(config)# ipv6 route 2001:db8:acad:2::/64 g0/0 2001:db8:feed::1
```

- A floating static route, administrative distance greater than primary route:

```
R1(config)# ipv6 route 2001:db8:acad:2::/64 2001:db8:feed::1
R1(config)# ipv6 route 2001:db8:acad:2::/64 2001:db8:feed::2 5
```

- A summary route:

```
R1(config)# ipv6 route 2001:db8:acad::/48 2001:db8:feed::1
```

Our Topology

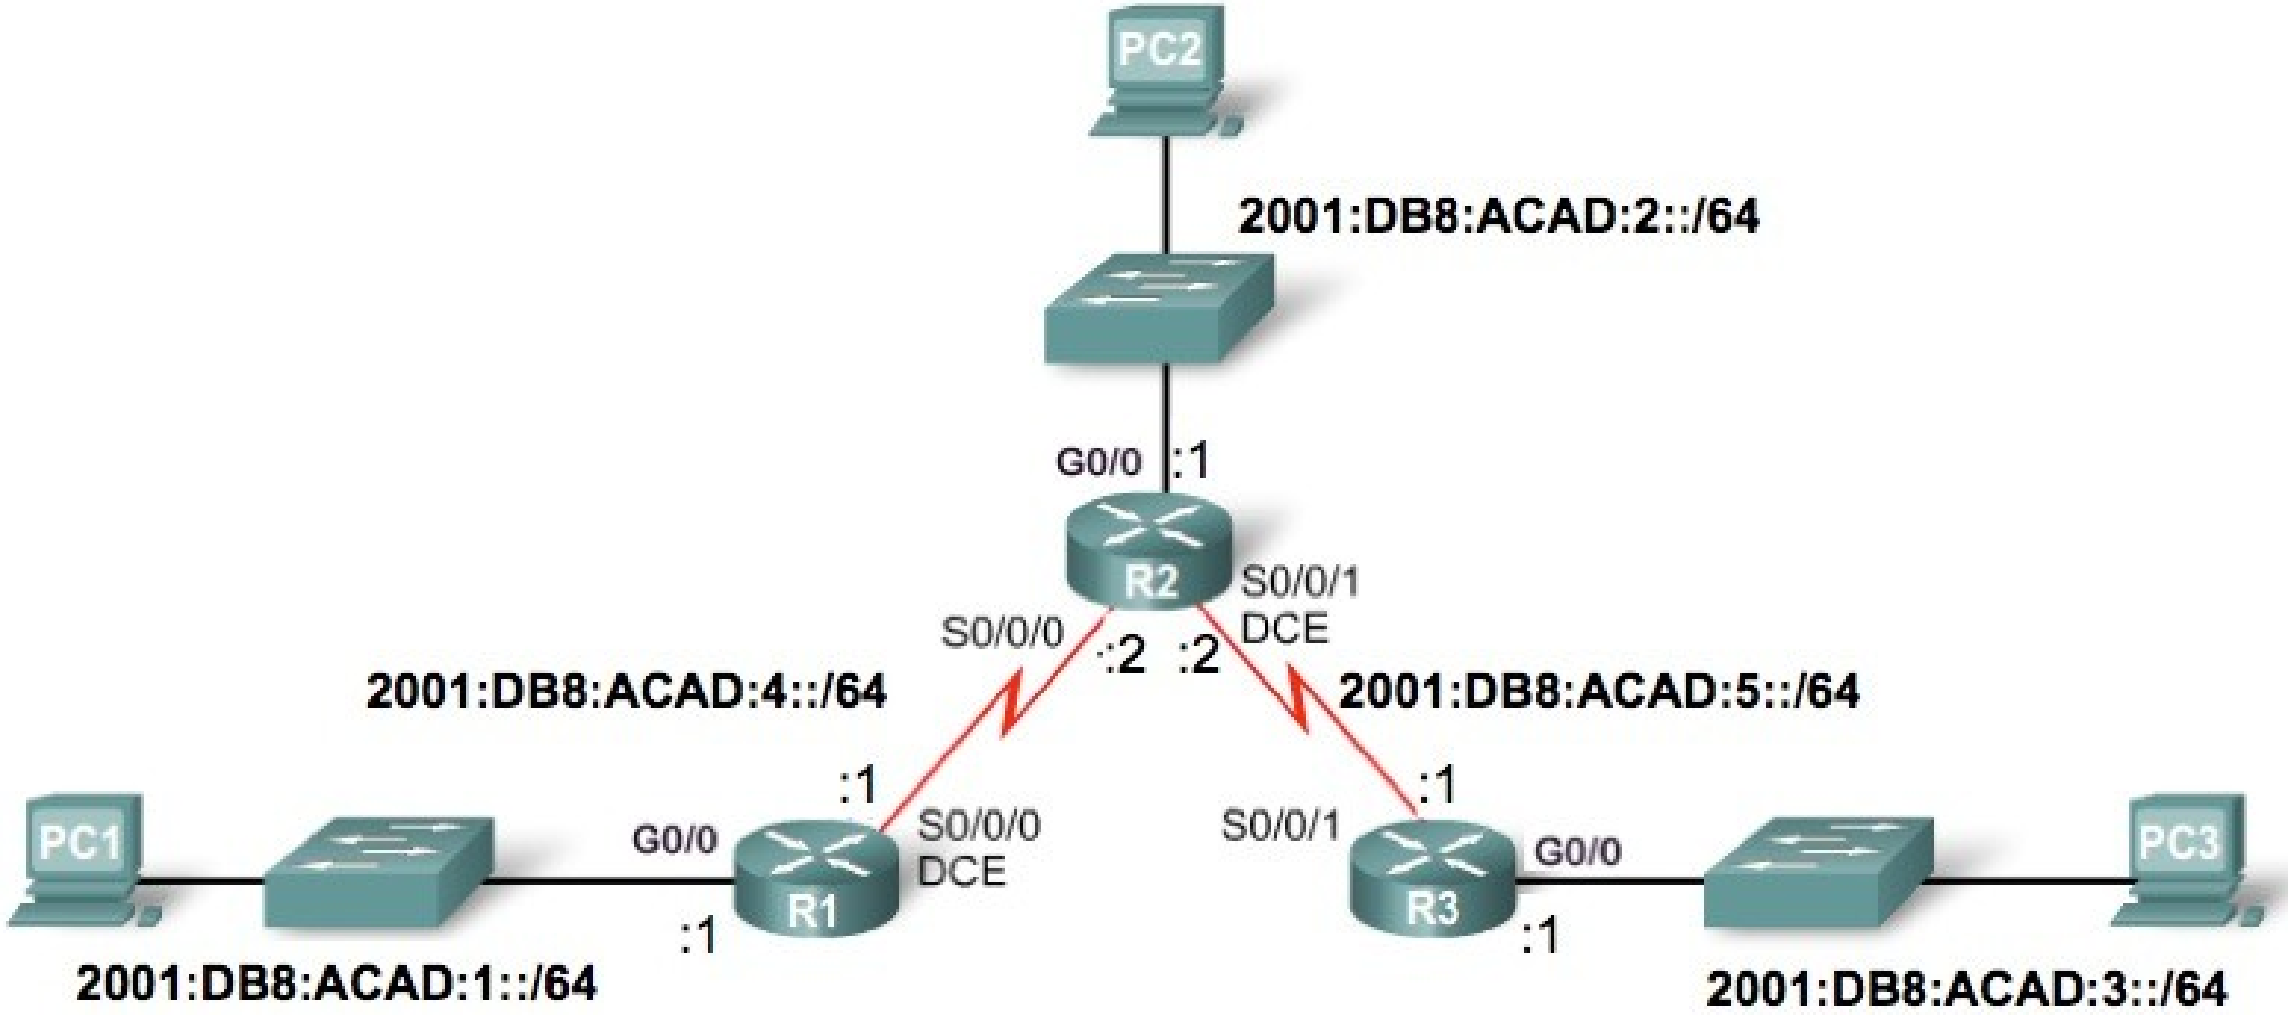

IPv6 Address Configuration

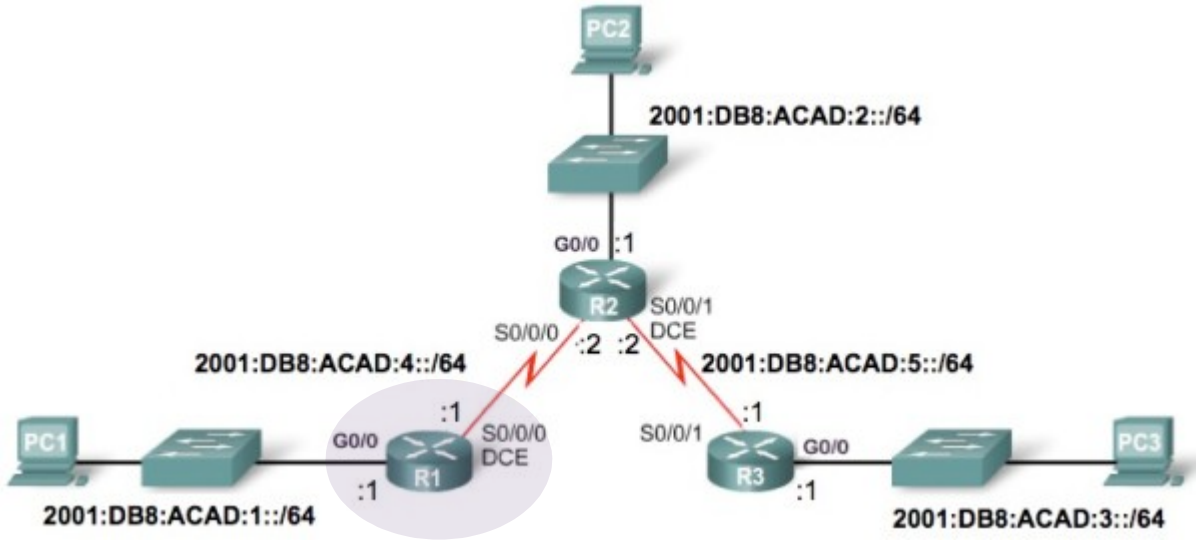


```
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# exit
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 address 2001:db8:acad:4::1/64
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# exit
```

Ugly EUI-64 Interface ID is used by default

Verifying IPv6 Interfaces

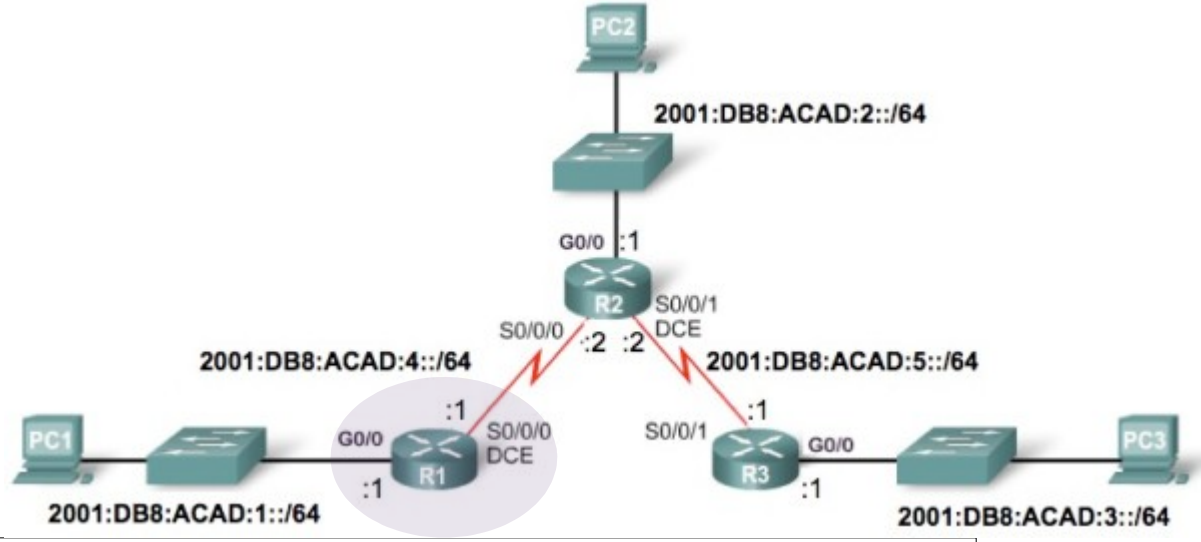


```
R1# show ipv6 interface brief
GigabitEthernet0/0      [up/up]
    FE80::1
    2001:DB8:ACAD:1::1
Serial0/0/0             [up/up]
    FE80::1
    2001:DB8:ACAD:4::1
R1#
```

Same Link-local address on all interfaces

IPv6 Routing Table: Directly Connected Networks

```
R1#show ipv6 route

<output omitted>

C   2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0, receive
C   2001:DB8:ACAD:4::/64 [0/0]
     via Serial0/0/0, directly connected
L   2001:DB8:ACAD:4::1/128 [0/0]
     via Serial0/0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
R1#
```

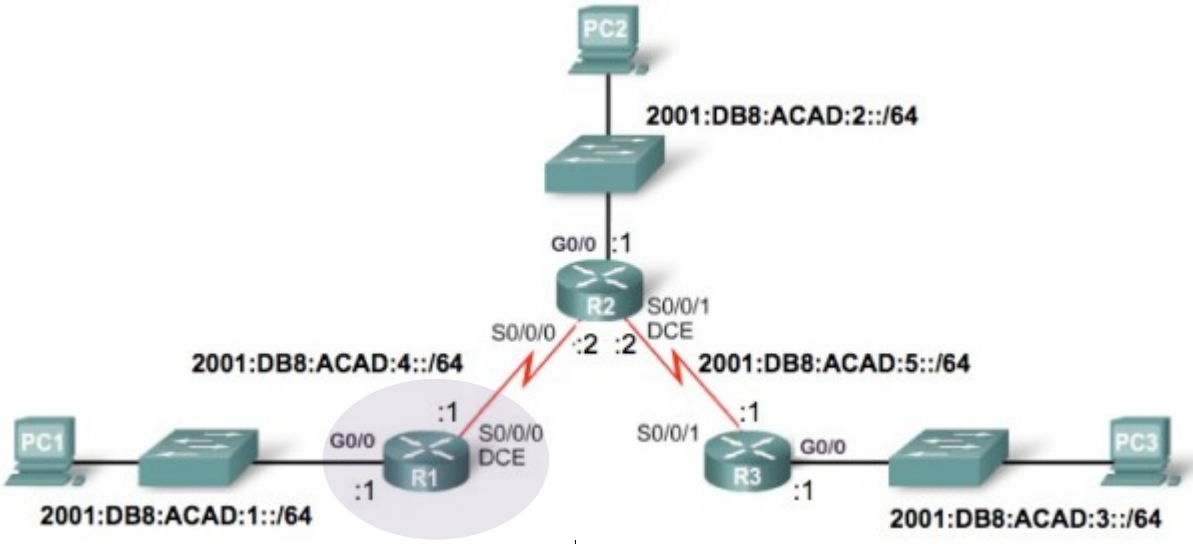

- Connected routes occur for any interface with an IPv6 unicast address that has more than link local scope
- **Link-local** addresses are not included in the routing table because they are not routable off the link

IPv6 Routing Table: Local Addresses

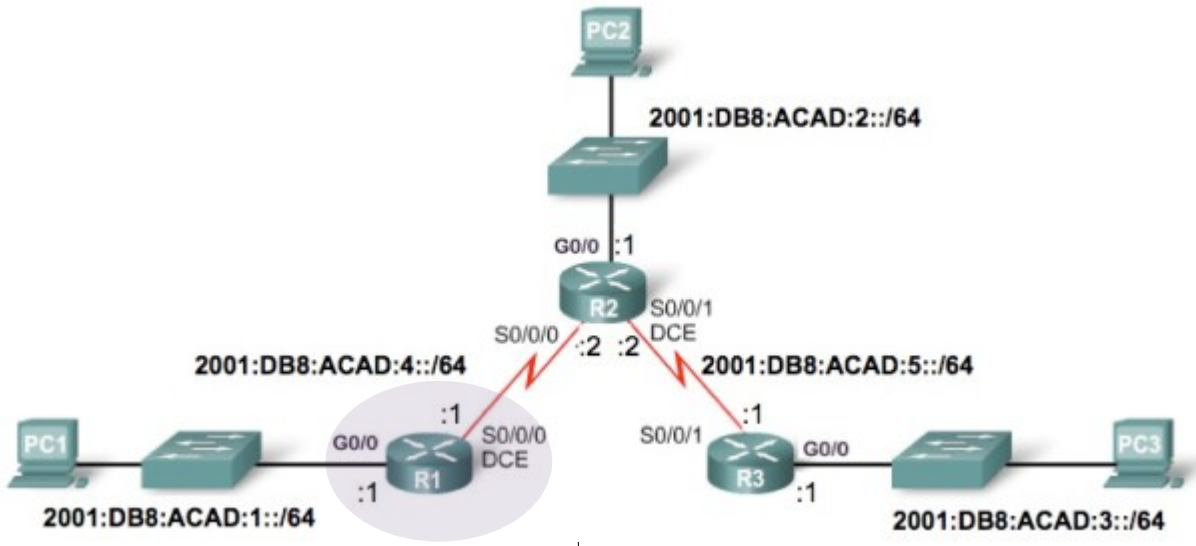


```
R1#show ipv6 route

<output omitted>

C    2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0, directly connected
L    2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0, receive
C    2001:DB8:ACAD:4::/64 [0/0]
     via Serial0/0/0, directly connected
L    2001:DB8:ACAD:4::1/128 [0/0]
     via Serial0/0/0, receive
L    FF00::/8 [0/0]
     via Null0, receive
R1#
```

Multicast packets
Not routed

- The local routes are all /128 routes (host routes) for the router's IPv6 unicast address
- Allow the router to more efficiently process packets directed to the router itself rather than for packet forwarding

Configuring IPv6 Static Routes

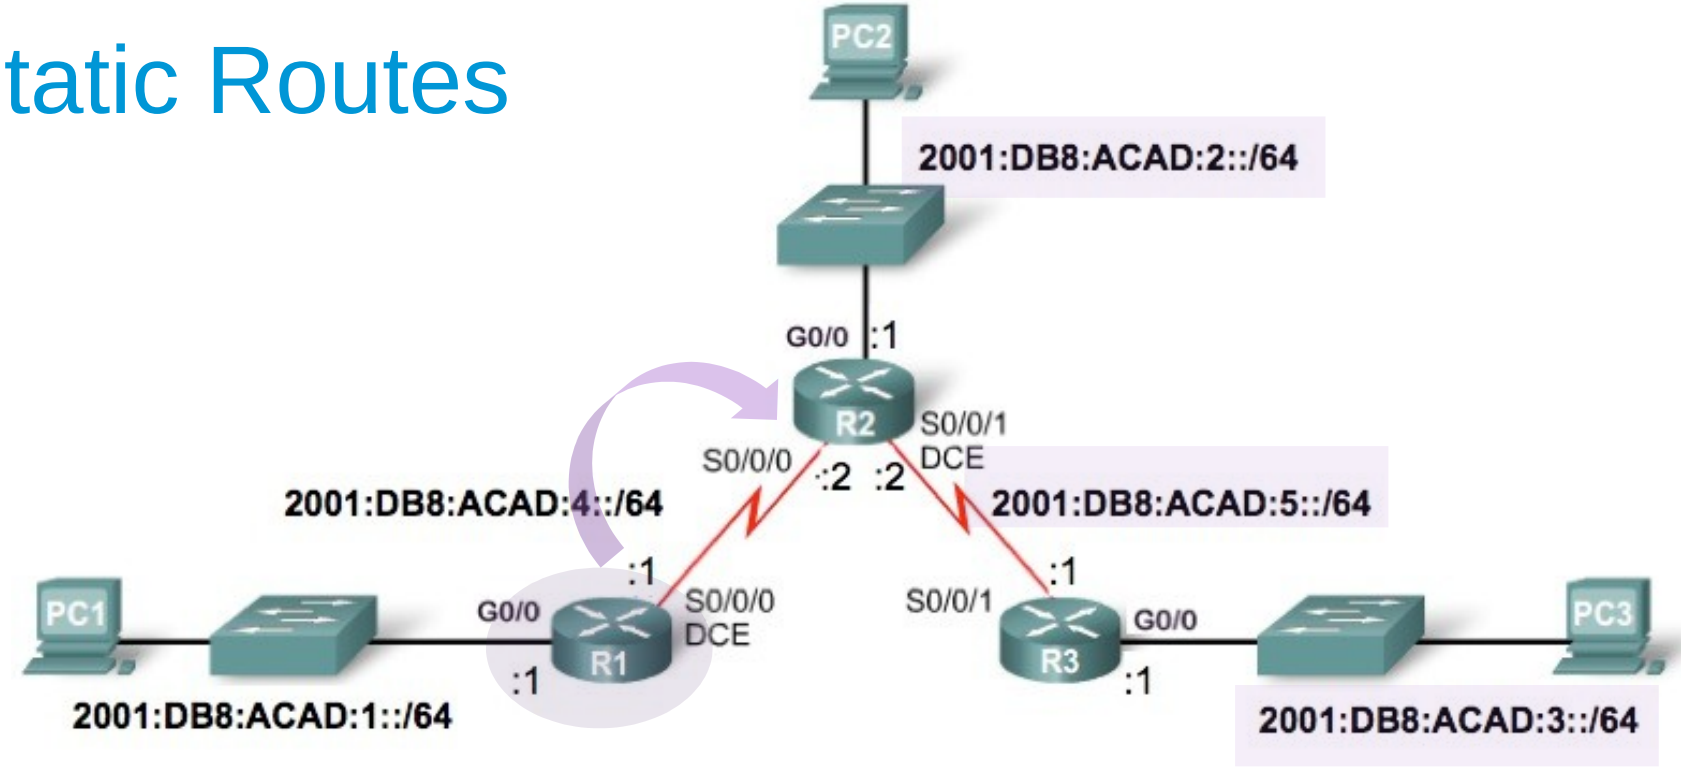


```
R1(config)#ipv6 route 2001:DB8:ACAD:2::/64 2001:DB8:ACAD:4::2
R1(config)#ipv6 route 2001:DB8:ACAD:5::/64 2001:DB8:ACAD:4::2
R1(config)#ipv6 route 2001:DB8:ACAD:3::/64 2001:DB8:ACAD:4::2
R1(config)#
```

IPv6 Routing Table
Static Routes

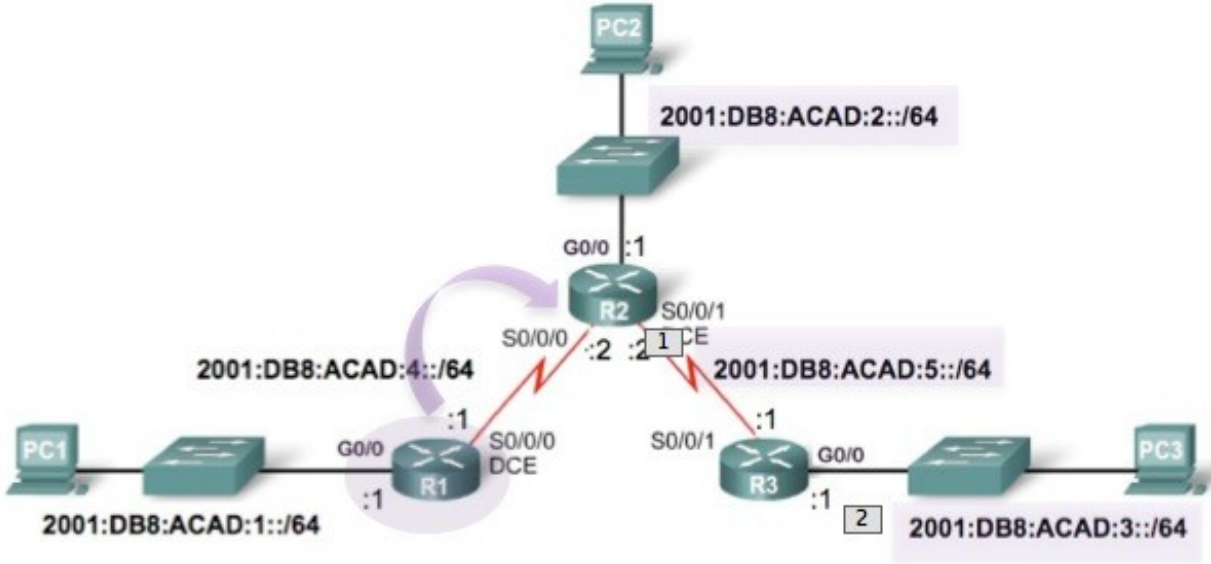


```
R1#show ipv6 route static
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, <Output omitted>
S   2001:DB8:ACAD:2::/64 [1/0]
     via 2001:DB8:ACAD:4::2
S   2001:DB8:ACAD:3::/64 [1/0]
     via 2001:DB8:ACAD:4::2
S   2001:DB8:ACAD:5::/64 [1/0]
     via 2001:DB8:ACAD:4::2
R1#
```

IPv6 Routing Table
Default Static Routes

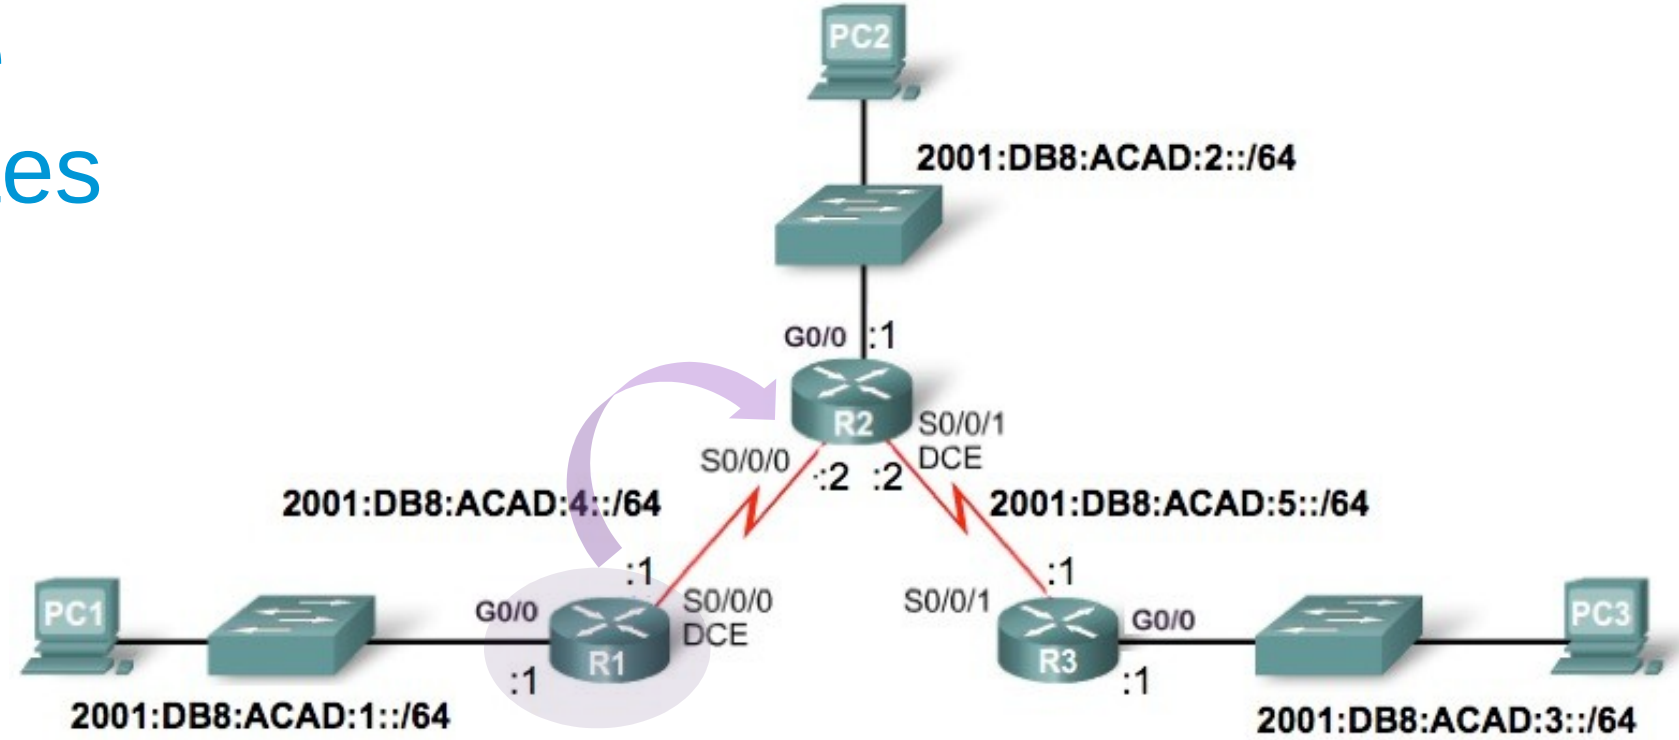


```
R1(config)# ipv6 route ::/0 2001:db8:acad:4::2
R1(config)# end
R1# show ipv6 route
S   ::/0 [1/0]
     via 2001:DB8:ACAD:4::2
```

Static Route using a Link-Local Next-Hop Address

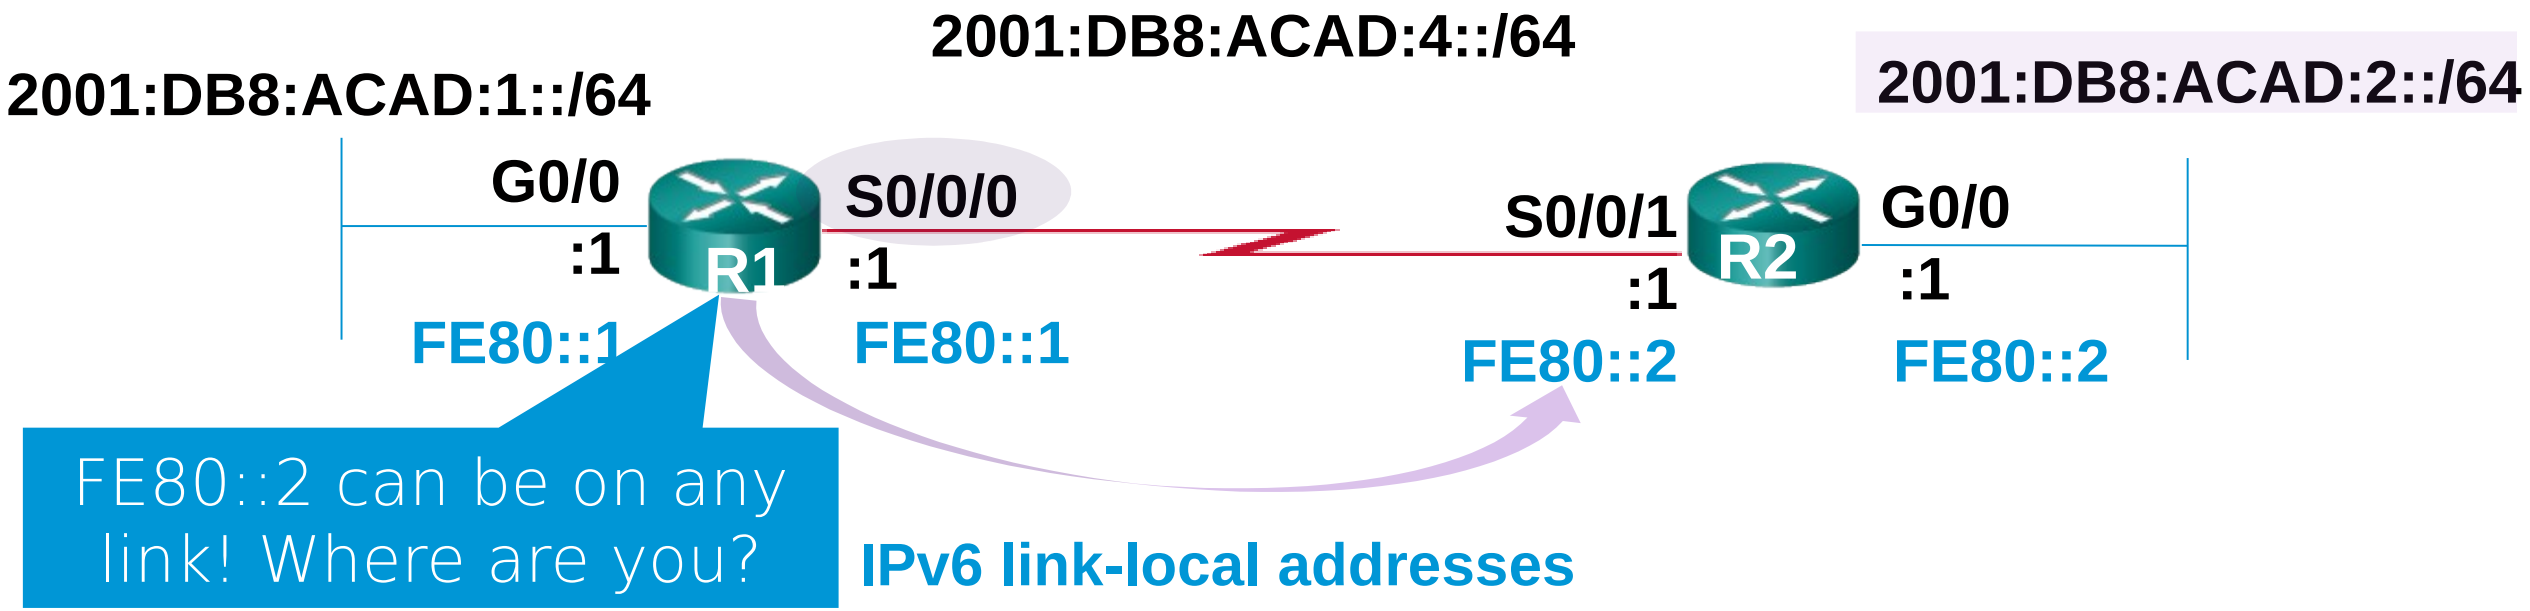


```
R1(config)# ipv6 route 2001:db8:acad:2::/64 fe80::2
% Interface has to be specified for a link-local nexthop
R1(config)# ipv6 route 2001:db8:acad:2::/64 s0/0/0 fe80::2
R1(config)# end
R1# show ipv6 route
S   2001:DB8:ACAD:2::/64 [1/0]
     via FE80::2, Serial0/0/0
```

Dynamic Routing Protocols

IPv6 Routing Protocols

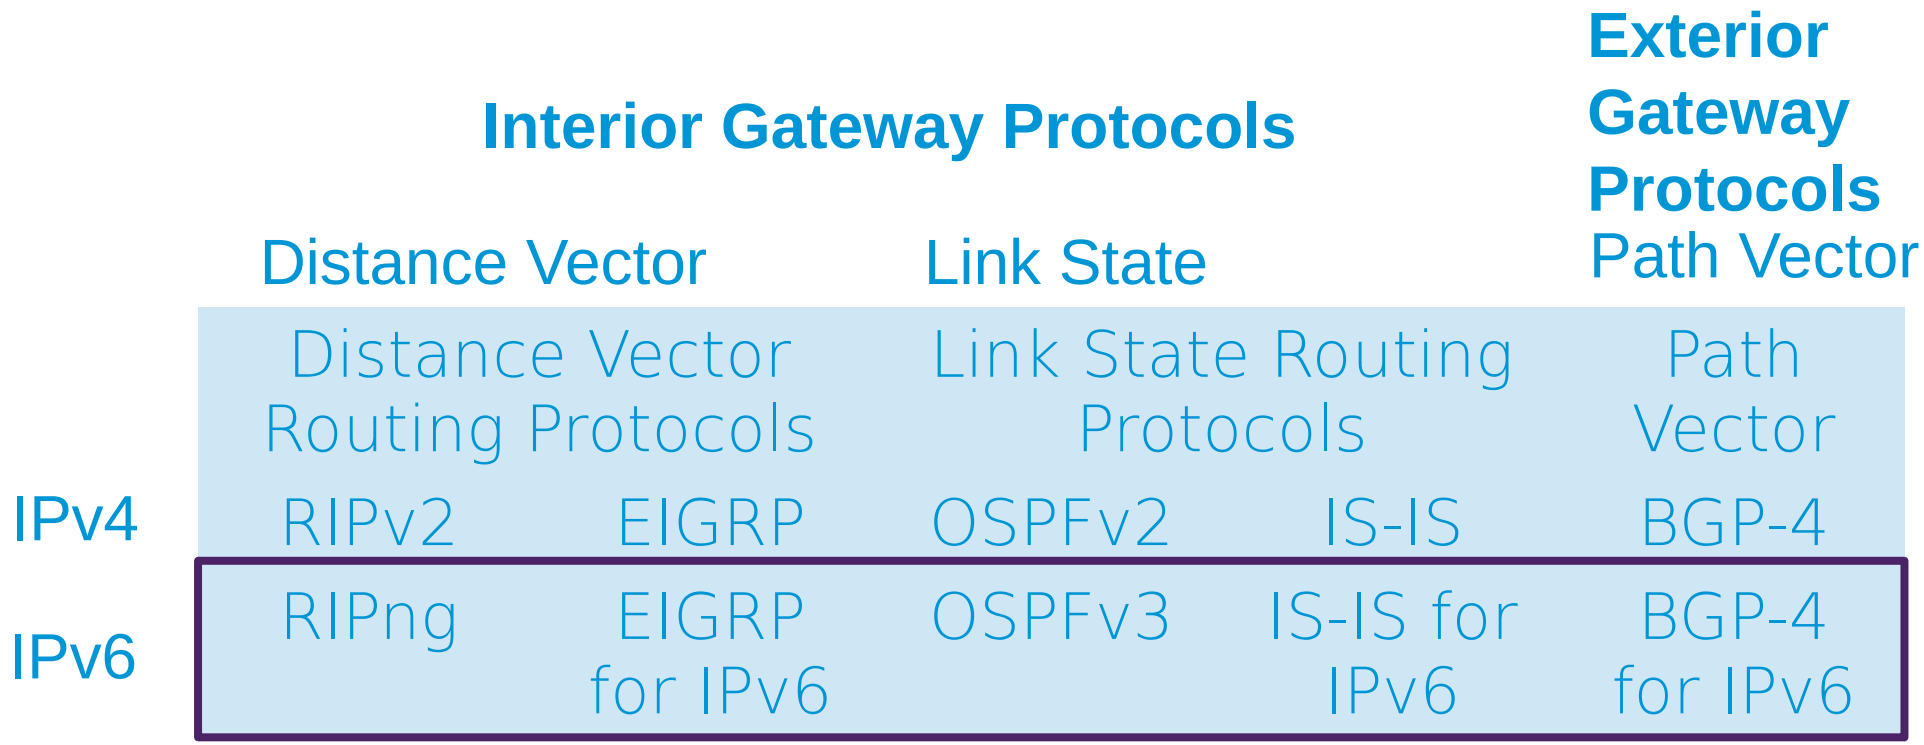

Most IPv6 routing protocol commands are identical to their IPv4 counterpart
Just need to substitute "**`ipv6`**" for "**`ip`**"

Routing Protocols… Same Payload, Different Truck

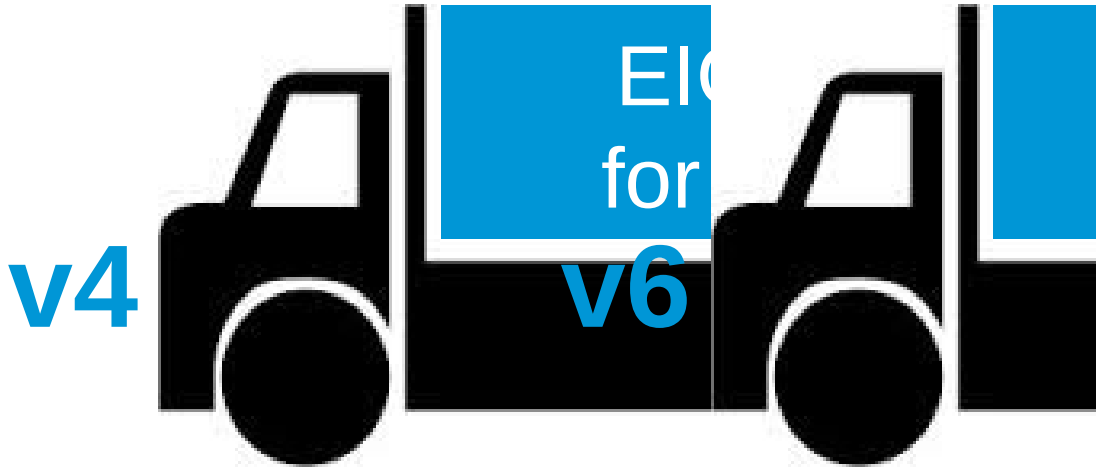

In case you are wondering….

Yes, I do drive a v6!

EIGRP for IPv6

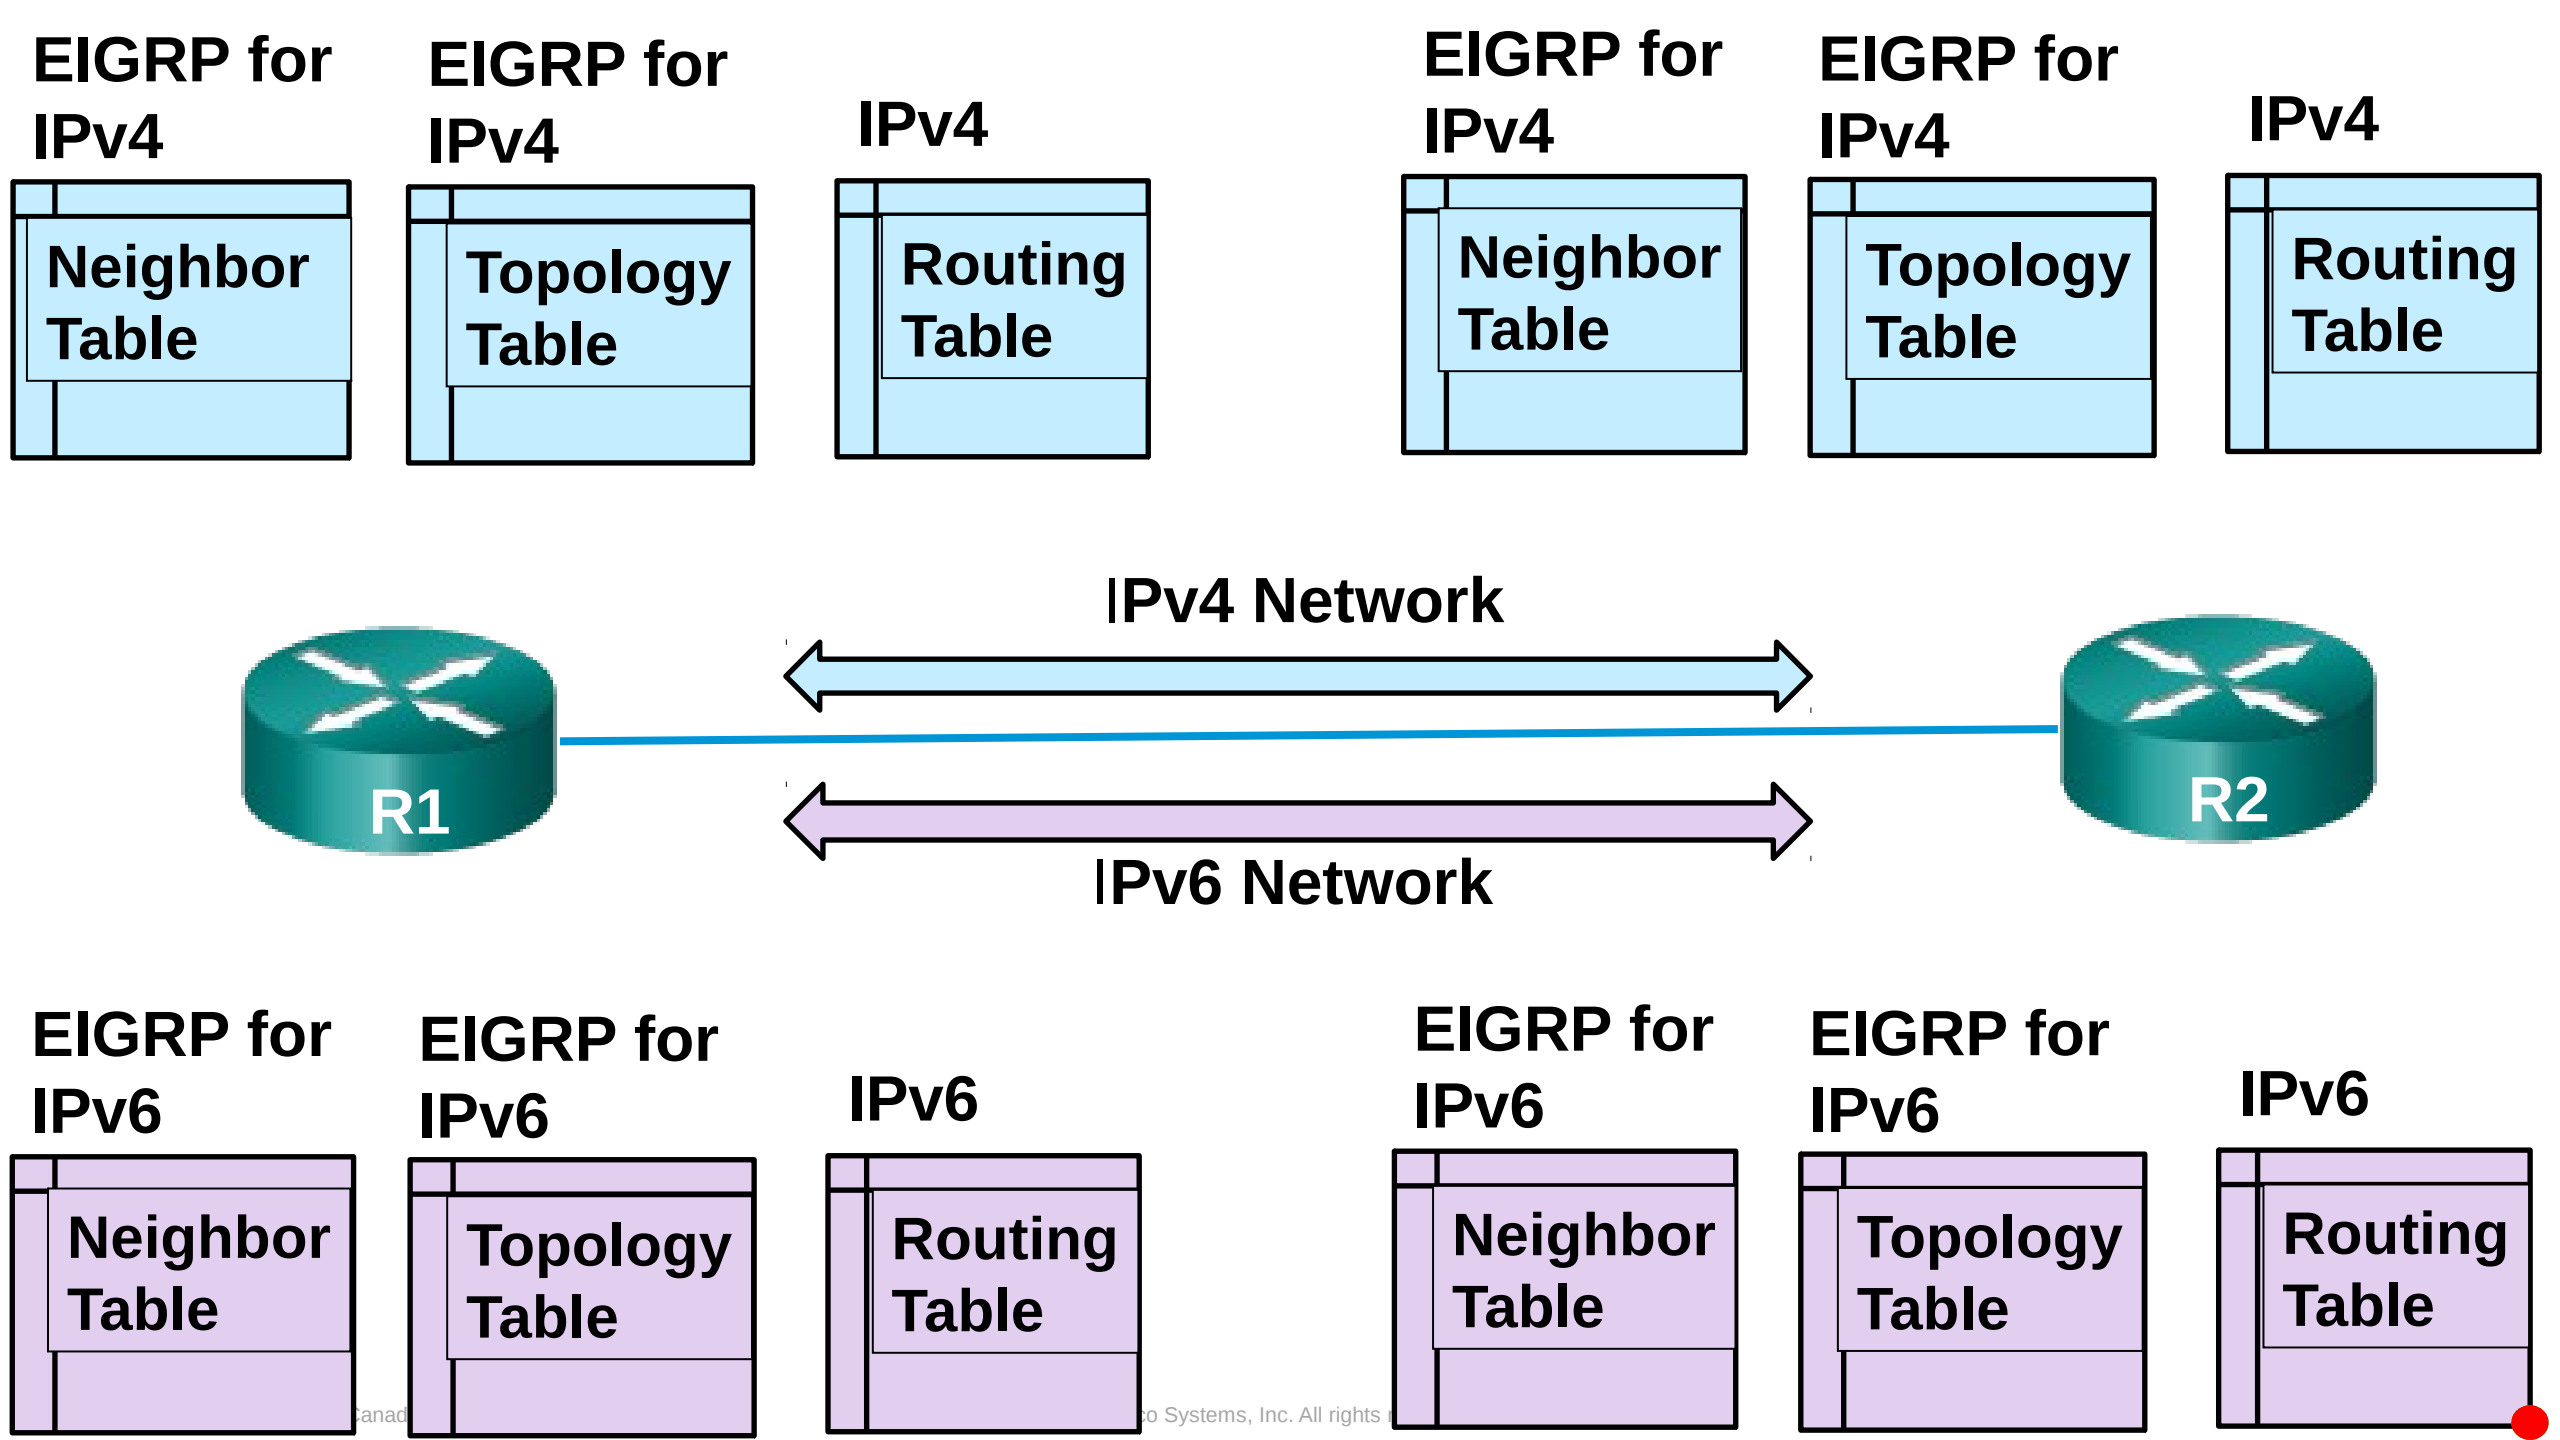
EIGRP for IPv4
Neighbor Table
EIGRP for IPv4
Topology Table
IPv4
Routing Table
EIGRP for IPv4
Neighbor Table
EIGRP for IPv4
Topology Table
IPv4
Routing Table
IPv4 Network
R1
R2
IPv6 Network
EIGRP for IPv6
Neighbor Table
EIGRP for IPv6
Topology Table
IPv6
Routing Table
EIGRP for IPv6
Neighbor Table
EIGRP for IPv6
Topology Table
IPv6
Routing Table

	EIGRP for IPv4	EIGRP for IPv6
Advertised routes		
Distance vector		
Convergence technology		
Metric		
Transport protocol		
Update messages		
Neighbor discovery		
Source address; destination addresses		
Authentication		
Router ID		

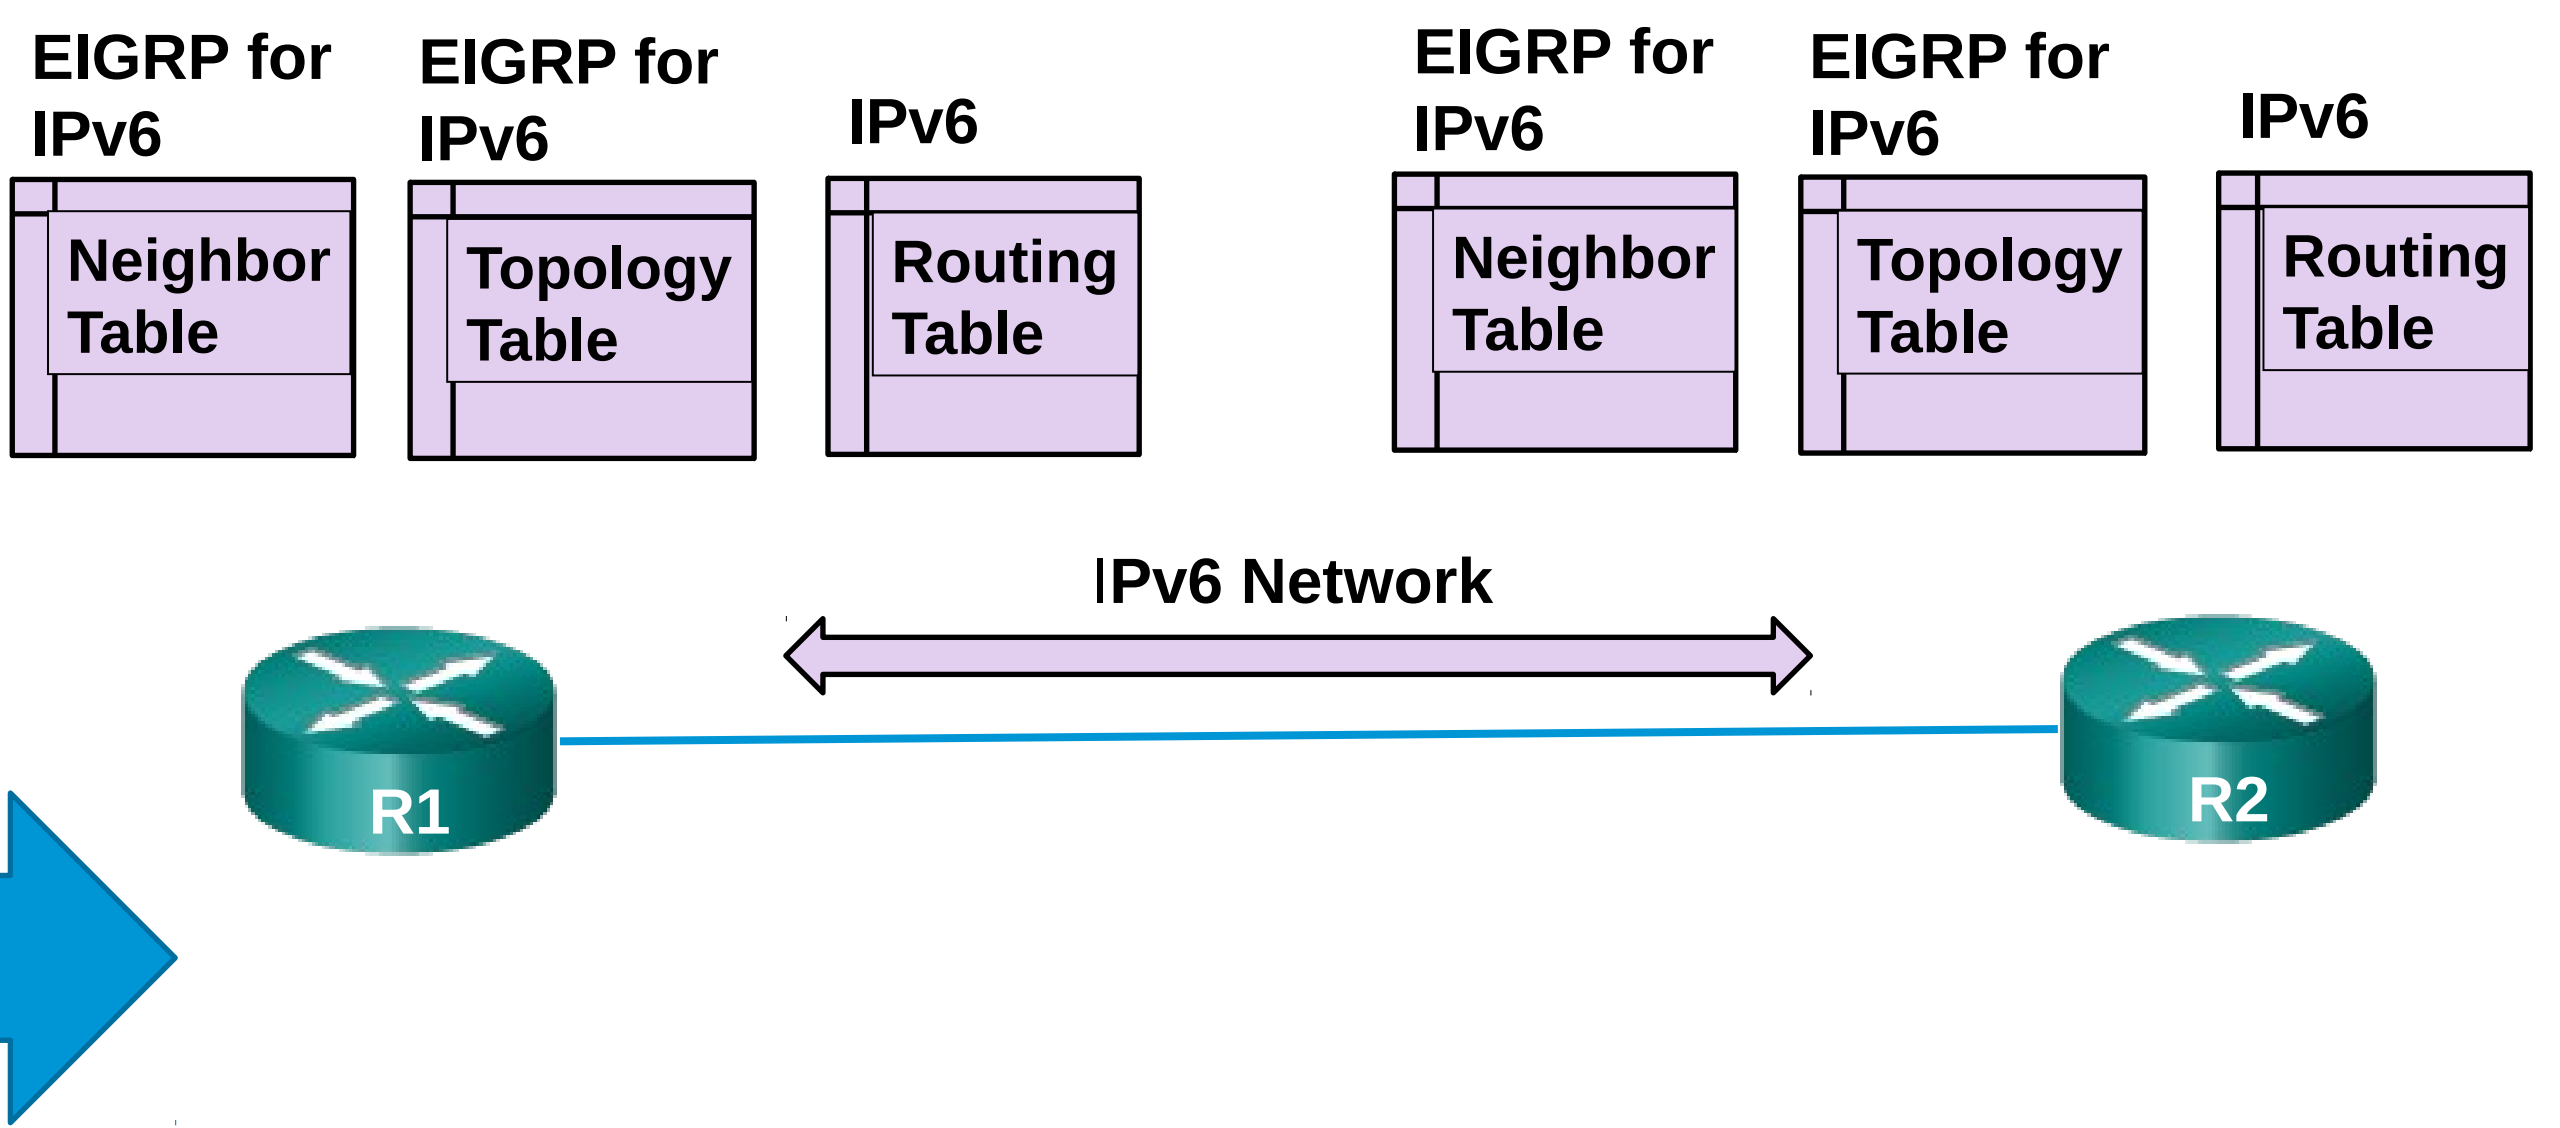

EIGRP for IPv6 Topology

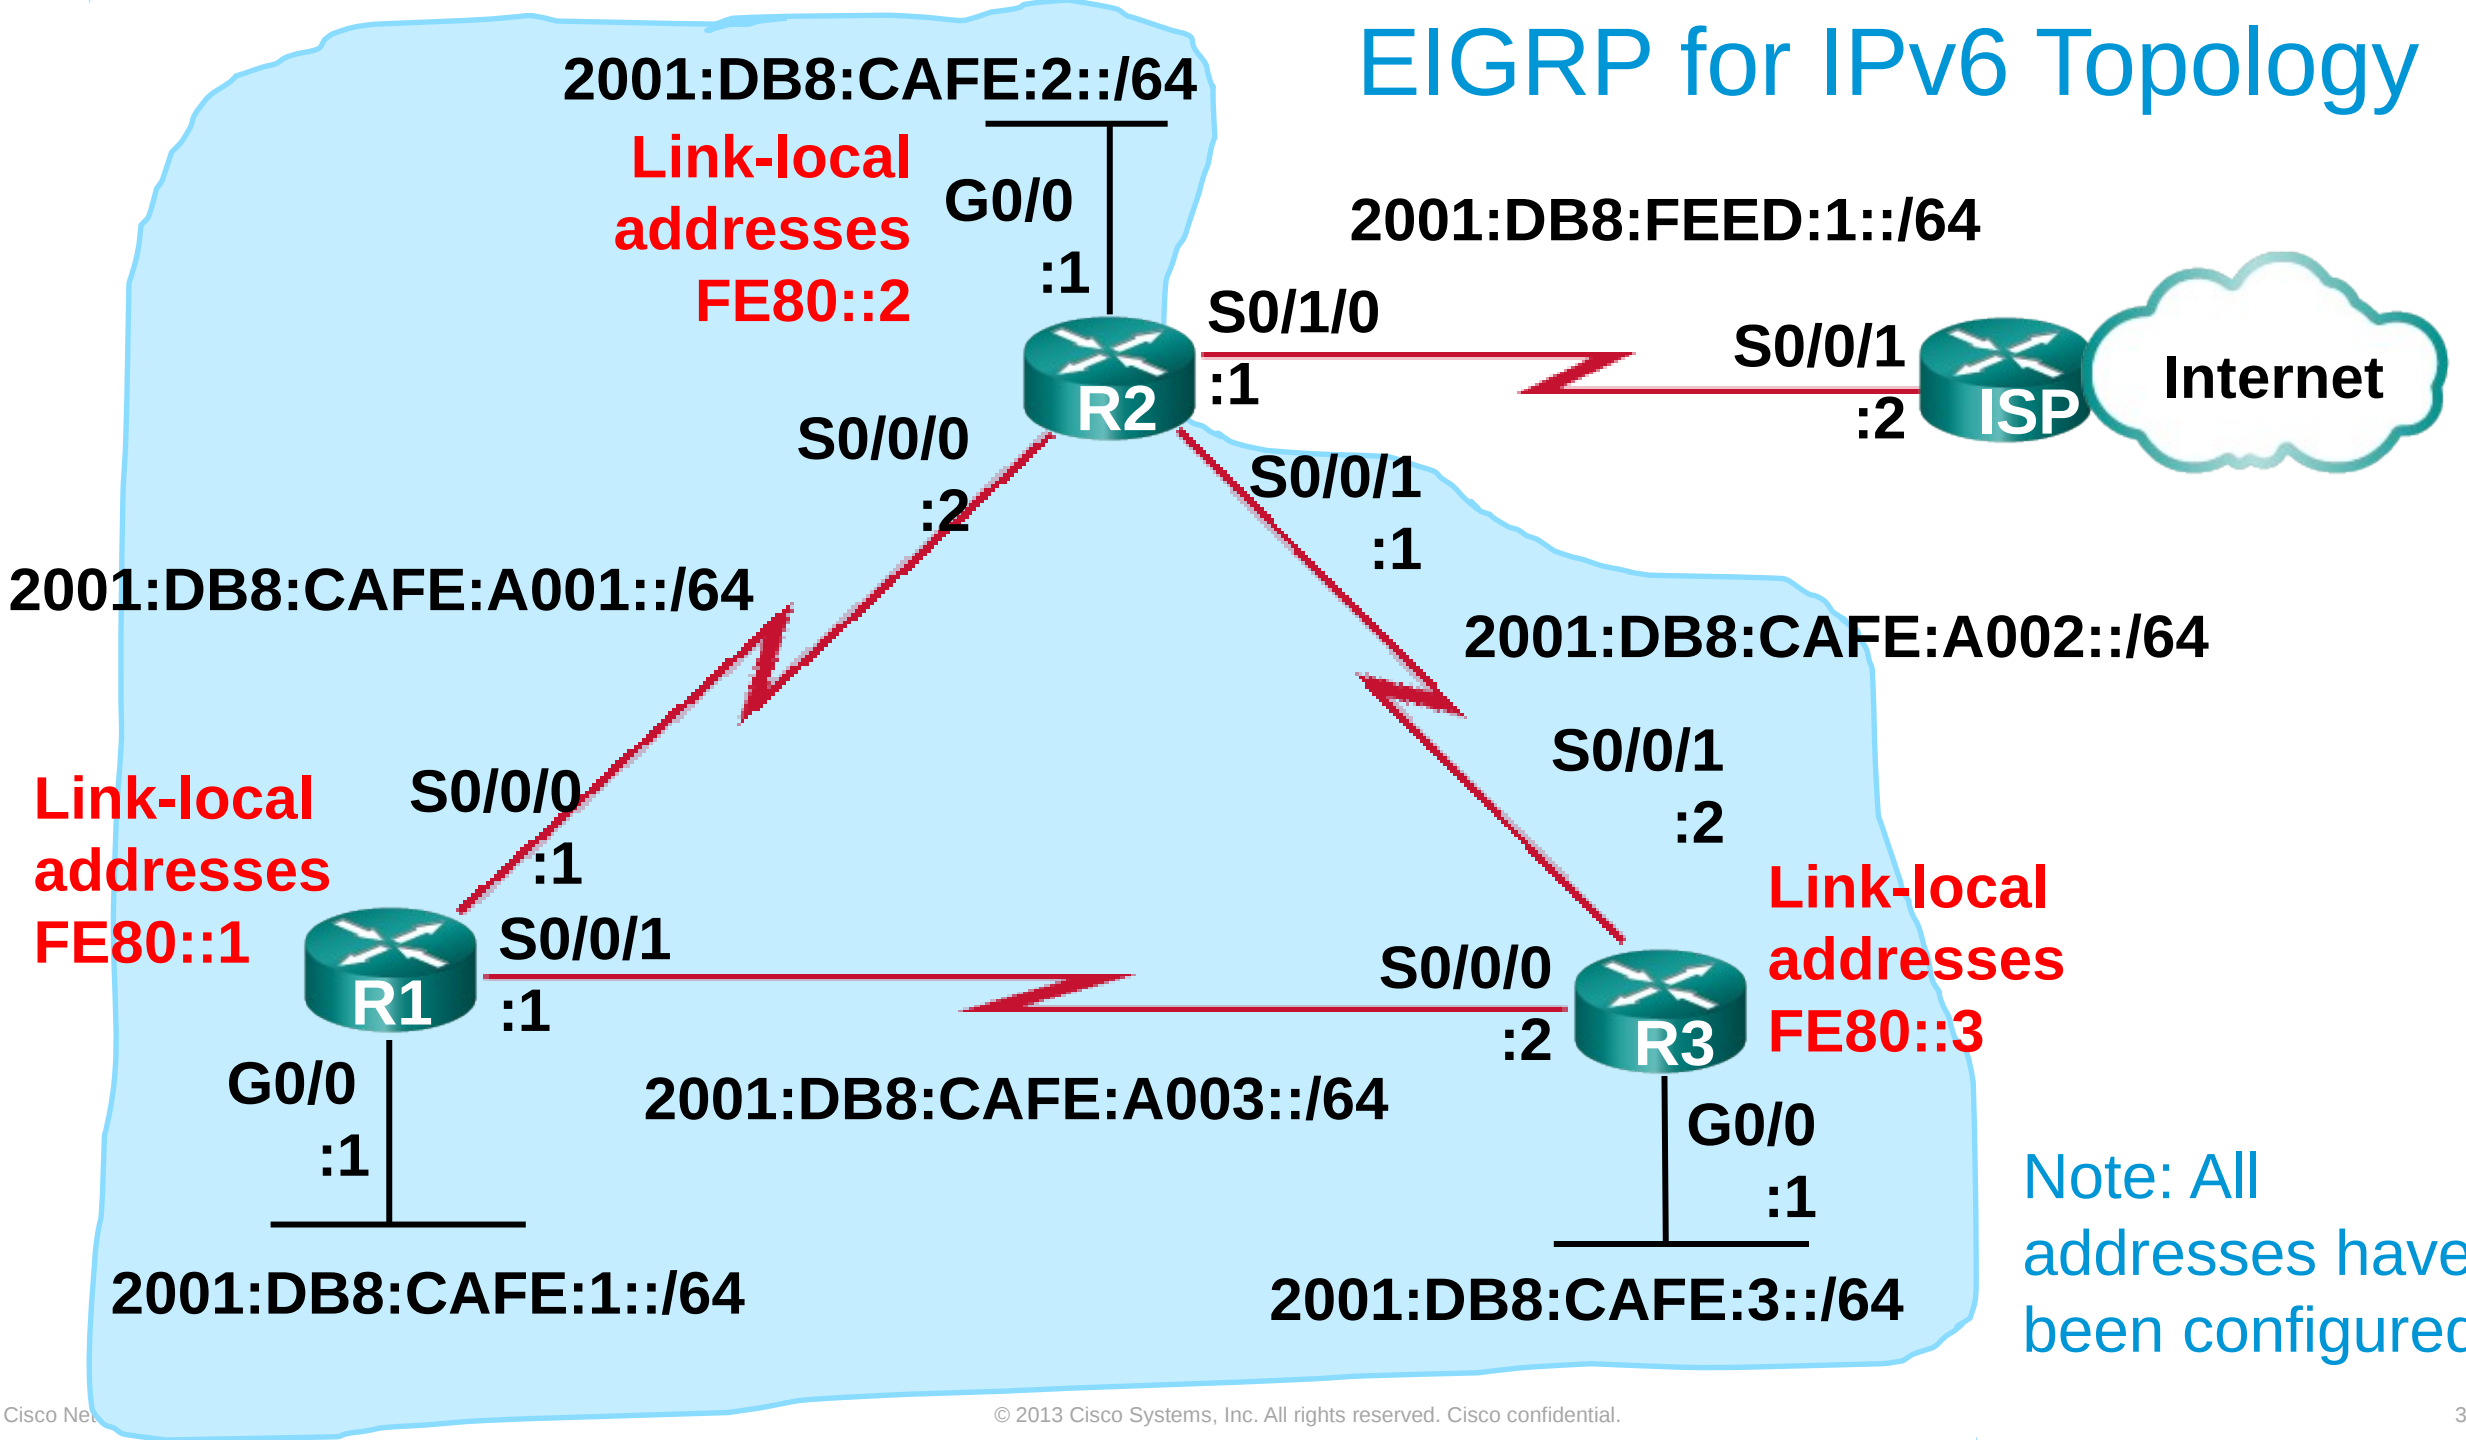

Note: All addresses have been configured

Configuring the EIGRP for IPv6 Routing Process

EIGRP for IPv6 was made available in Cisco IOS, Release 12.4(6)T

```
R1(config)# ipv6 router eigrp 2
% IPv6 routing not enabled
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 router eigrp 2
R1(config-rtr)# eigrp router-id 1.0.0.0
R1(config-rtr)# no shutdown
R1(config-rtr)#
```

Enables IPv6 routing

Must be same on all routers

This is specific to EIGRP for IPv6

- EIGRP uses a 32-bit Router ID for both IPv4 and IPv6 ("**eigrp**" not always required)
- The **eigrp router-id** command takes precedence over any loopback or physical interface IPv4 addresses
- If there are no active IPv4 interfaces, then the **eigrp router-id** command is required
- Router ID should be a unique otherwise, routing inconsistencies can occur

Enabling EIGRP for IPv6 on the Interface

No **network** commands needed!

```
R1(config)#interface g0/0
R1(config-if)#ipv6 eigrp 2
R1(config-if)#exit
R1(config)#interface s 0/0/0
R1(config-if)#ipv6 eigrp 2
R1(config-if)#exit
R1(config)#interface s 0/0/1
R1(config-if)#ipv6 eigrp 2
R1(config-if)#
```

Perform same process on other routers in domain…..

```
R2(config)#ipv6 unicast-routing
R2(config)#ipv6 router eigrp 2
R2(config-rtr)#eigrp router-id 2.0.0.0
R2(config-rtr)#no shutdown
R2(config-rtr)#
R2(config)#interface g 0/0
R2(config-if)#ipv6 eigrp 2
R2(config-if)#exit
R2(config)#interface s 0/0/0
R2(config-if)#ipv6 eigrp 2
R2(config-if)#exit
%DUAL-5-NBRCHANGE: EIGRP-IPv6 2: Neighbor FE80::1
(Serial0/0/0) is up: new adjacency
R2(config)#interface s 0/0/1
R2(config-if)#ipv6 eigrp 2
R2(config-if)#
```

Verifying EIGRP for IPv6 Neighbor Adjacencies

Similar information as EIGRP for IPv6

```
R1#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(2)
H   Address                 Interface Hold   Uptime    SRTT    RTO  Q  Seq
                                      (sec)             (ms)       Cnt Num
1   Link-local address:     Se0/0/1     13   00:37:17    45    270  0  8
    FE80::3
0   Link-local address:     Se0/0/0     14   00:53:16    32   2370  0  8
    FE80::2
R1#
```

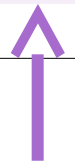

Link-local addresses are used to form adjacencies and source messages

Verifying EIGRP for IPv6 Parameters

```
R1#show ipv6 protocols
<Some output omitted for brevity>
IPv6 Routing Protocol is "eigrp 2"
EIGRP-IPv6 Protocol for AS(2)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  NSF-aware route hold timer is 240
  Router-ID: 1.0.0.0
  Topology : 0 (base)
    Distance: internal 90 external 170
    Maximum path: 16
    Maximum hopcount 100
    Maximum metric variance 1

  Interfaces:
    GigabitEthernet0/0
    Serial0/0/0
    Serial0/0/1
R1#
```

Routing protocol and Process ID (AS Number)

Same K values used in composite metric

EIGRP Router ID

Same EIGRP Administrative Distances

Interfaces enabled for this EIGRP for IPv6

Verifying EIGRP for IPv6 Learned Prefixes

```
R1#show ipv6 route eigrp
<Some output omitted>

D   2001:DB8:CAFE:2::/64 [90/3524096]
     via FE80::3, Serial0/0/1
D   2001:DB8:CAFE:3::/64 [90/2170112]
     via FE80::3, Serial0/0/1
D   2001:DB8:CAFE:A002::/64 [90/3523840]
     via FE80::3, Serial0/0/1
R1#
```

Link-local addresses are used as next hop addresses

EIGRP for IPv6 Manual Summarization

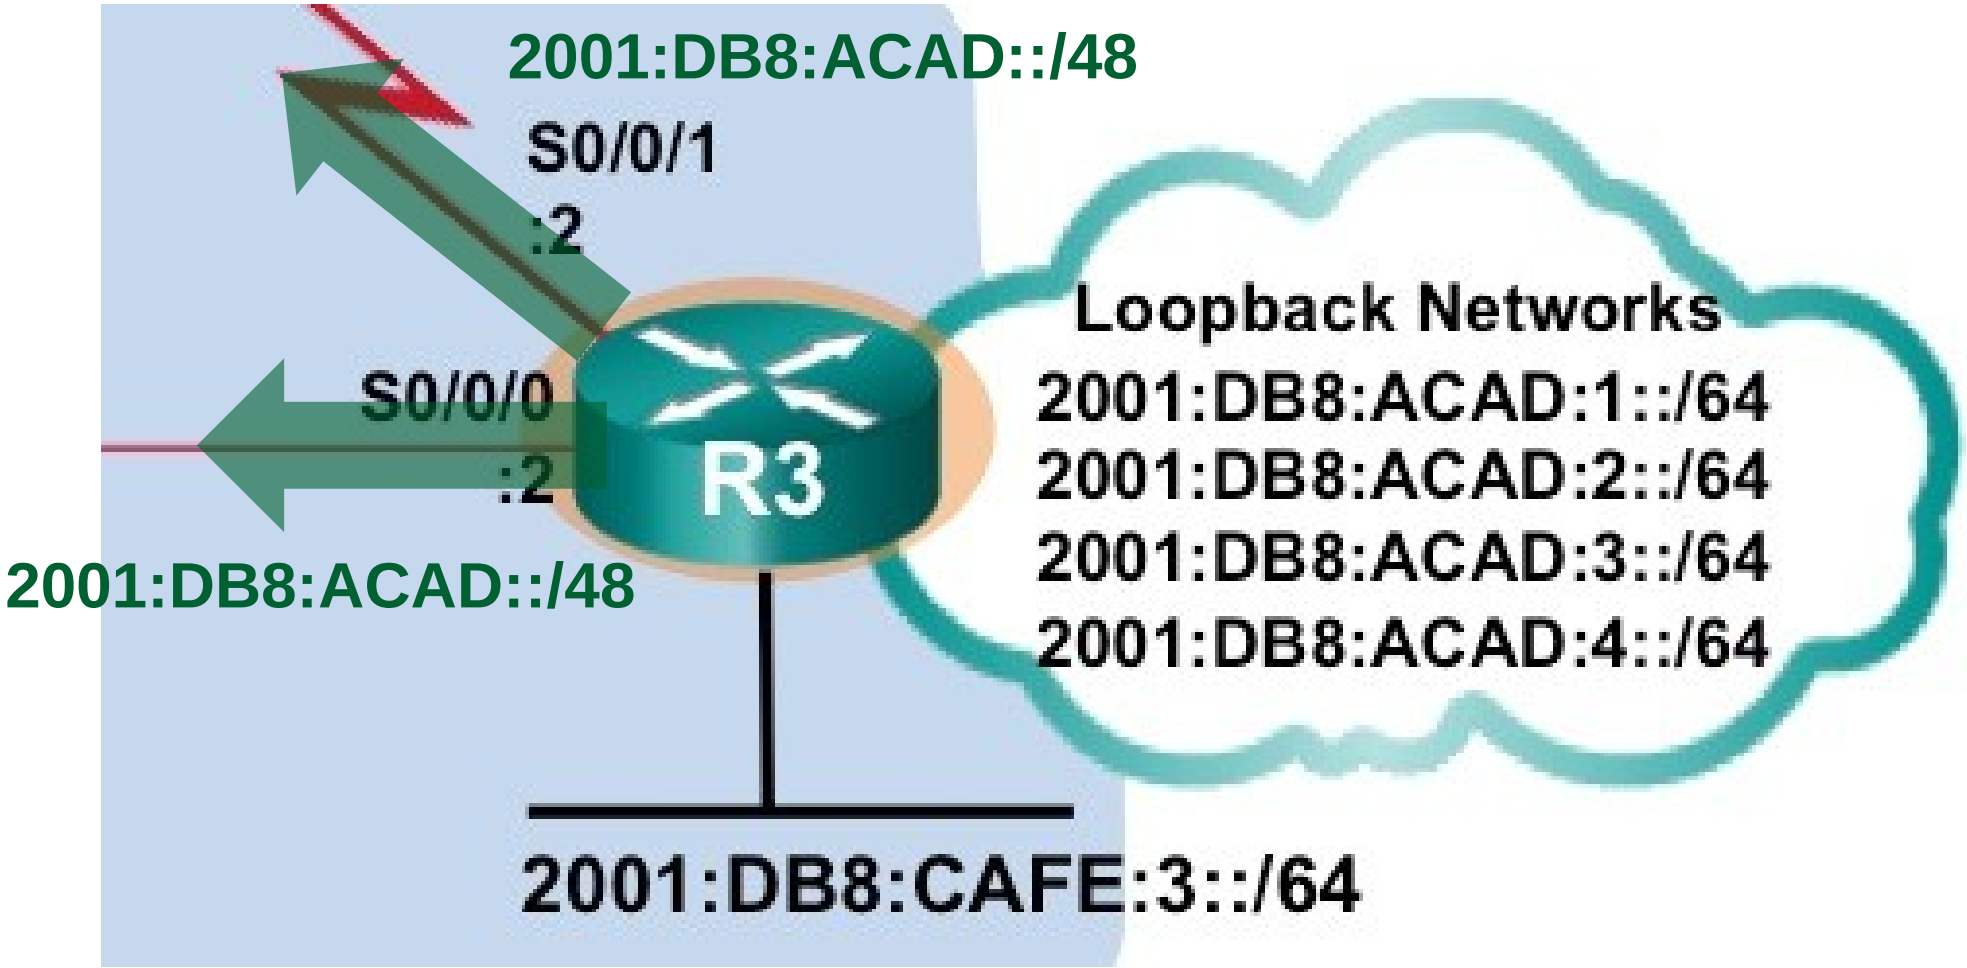

Note 1: There is no automatic summarization in IPv6 (no classful networks)

Note 2: EIGRP for IPv4 automatic summarization is disabled by default beginning with Cisco IOS Release 15.0(1)M and 12.2(33)

EIGRP for IPv6 Manual Summarization

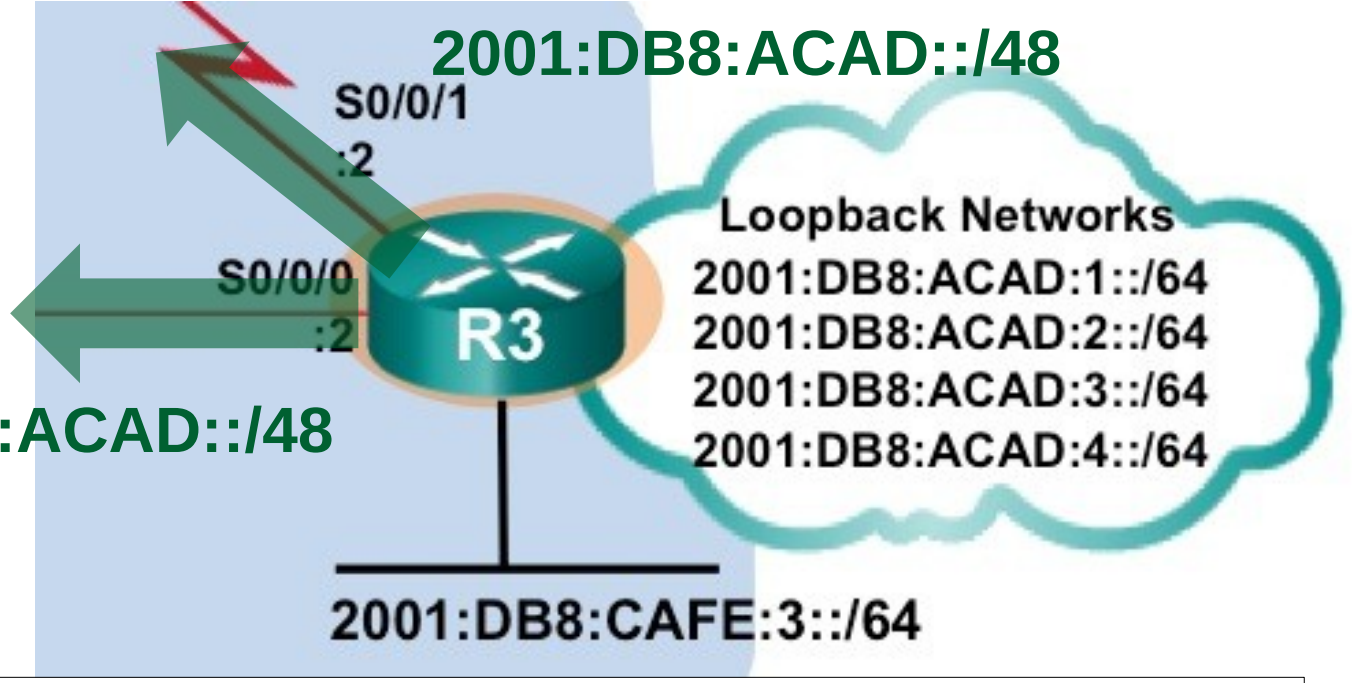

Similar command as EIGRP for IPv4

```
R3(config)# interface serial 0/0/0
R3(config-if)# ipv6 summary-address eigrp 2 2001:db8:acad::/48
R3(config-if)# exit
R3(config)# interface serial 0/0/1
R3(config-if)# ipv6 summary-address eigrp 2 2001:db8:acad::/48
R3(config-if)# end
R3# show ipv6 route

D   2001:DB8:ACAD::/48 [5/128256]
     via Null0, directly connected
```

Similar to EIGRP for IPv4, R3 includes a summary route to null0 as a loop prevention mechanism

```
R1#  show ipv6 route | include 2001:DB8:ACAD:
D   2001:DB8:ACAD::/48 [90/2297856]
R1#
```

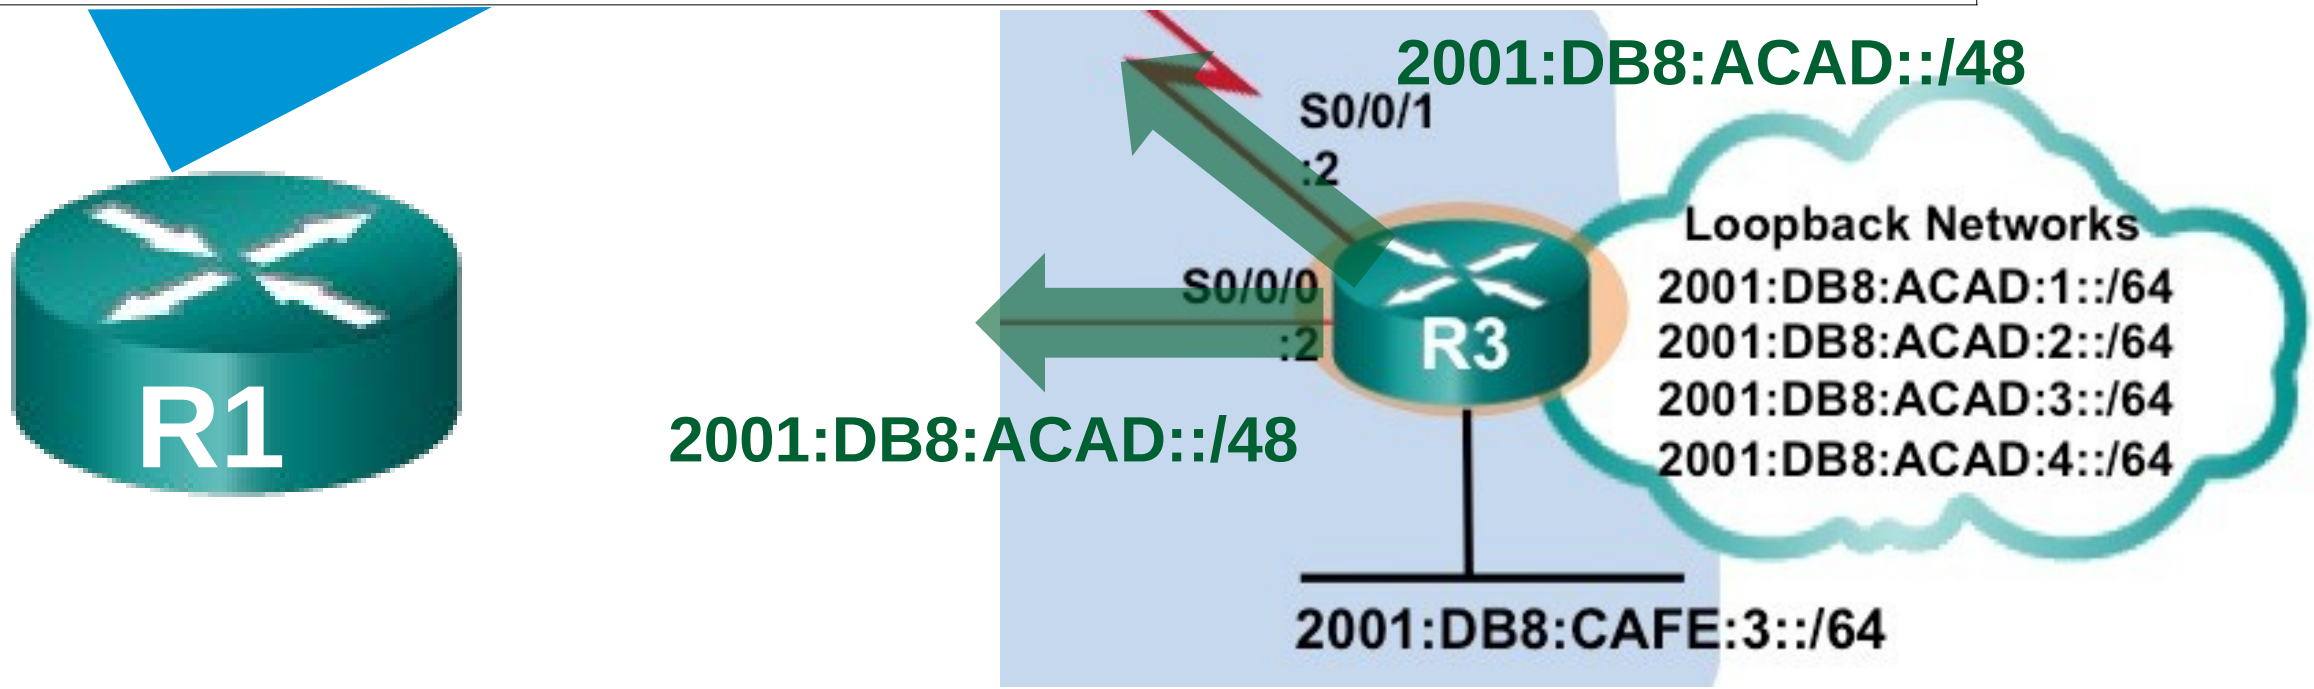

Propagating a Default Static Route in EIGRP for IPv6

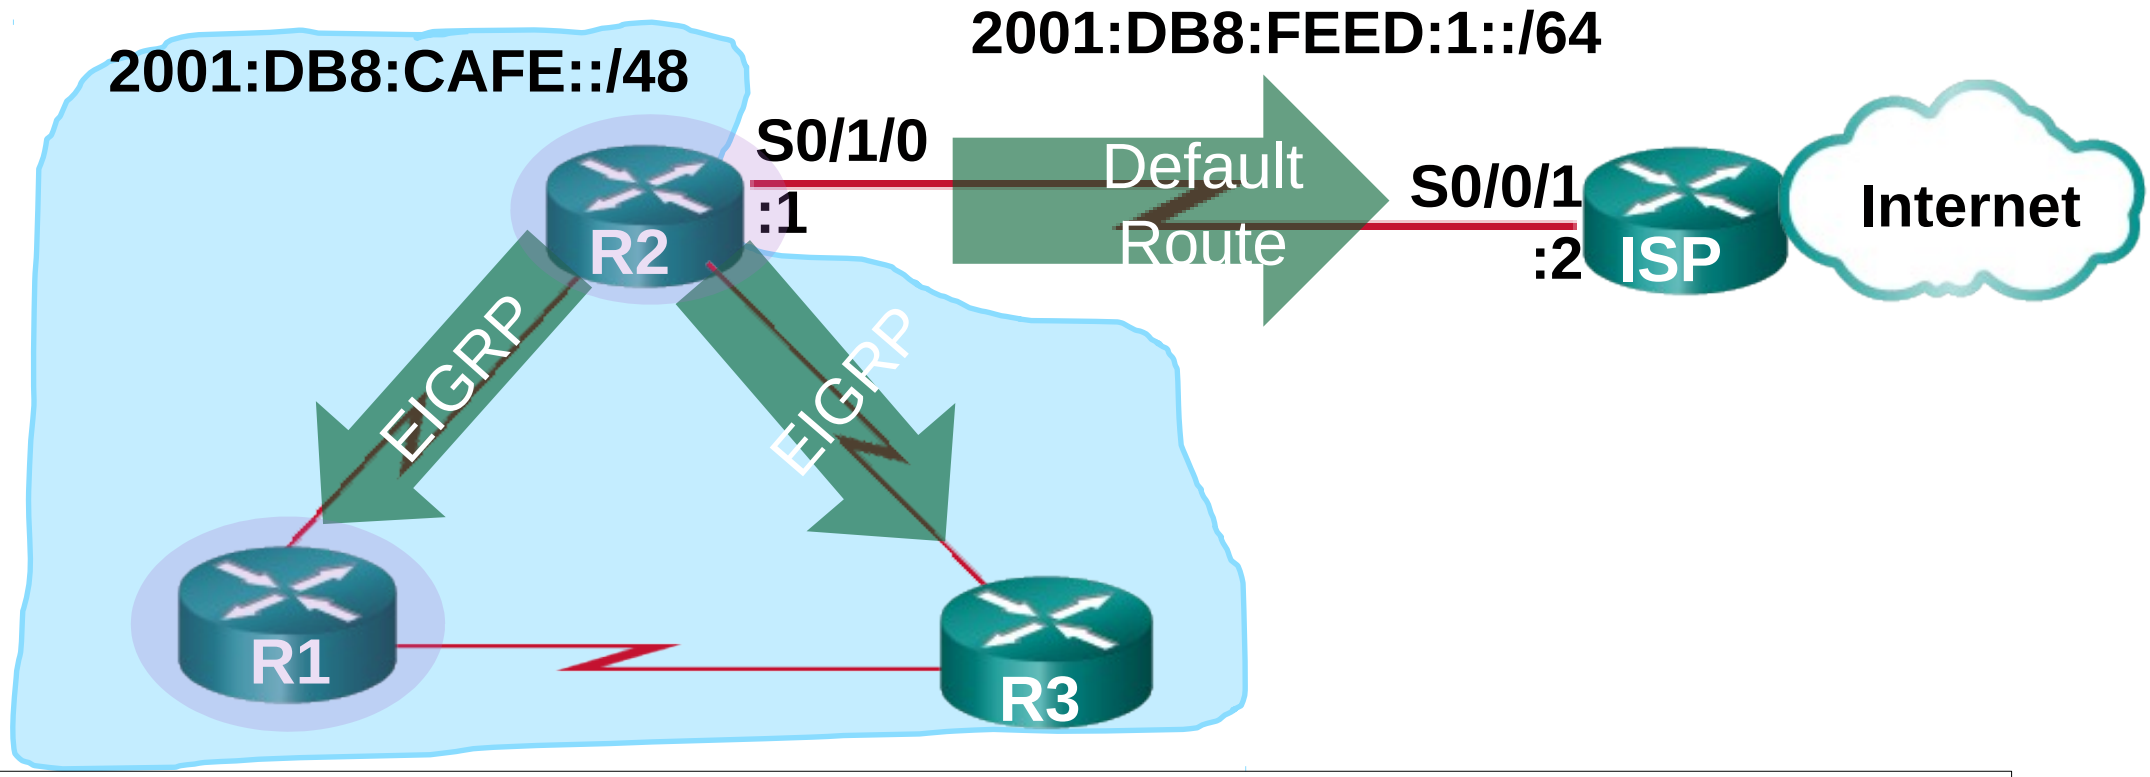


```
R2(config)# ipv6 route ::/0 2001:DB8:FEED:1::2
R2(config)# ipv6 router eigrp 2
R2(config-rtr)# redistribute static
```

```
R1# show ipv6 route
EX   ::/0 [170/3523840]
      via FE80::3, Serial0/0/1
```

EX = EIGRP External

Other commands…. Exactly the same, only different

Bandwidth utilization, Hello and Hold Timers remain the same

```
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 bandwidth-percent eigrp 2 50
R1(config-if)# ipv6 hello-interval eigrp 2 60
R1(config-if)# ipv6 hold-time eigrp 2 180
R1(config-if)#
```

Just add “v6”!

Same with MD5 Authentication (only IPv6 relevant commands are shown)

```
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 authentication mode eigrp 2 md5
R1(config-if)# ipv6 authentication key-chain eigrp 2 EIGRPV6_KEY
```

OSPFv3
Single Area and Multi-Area

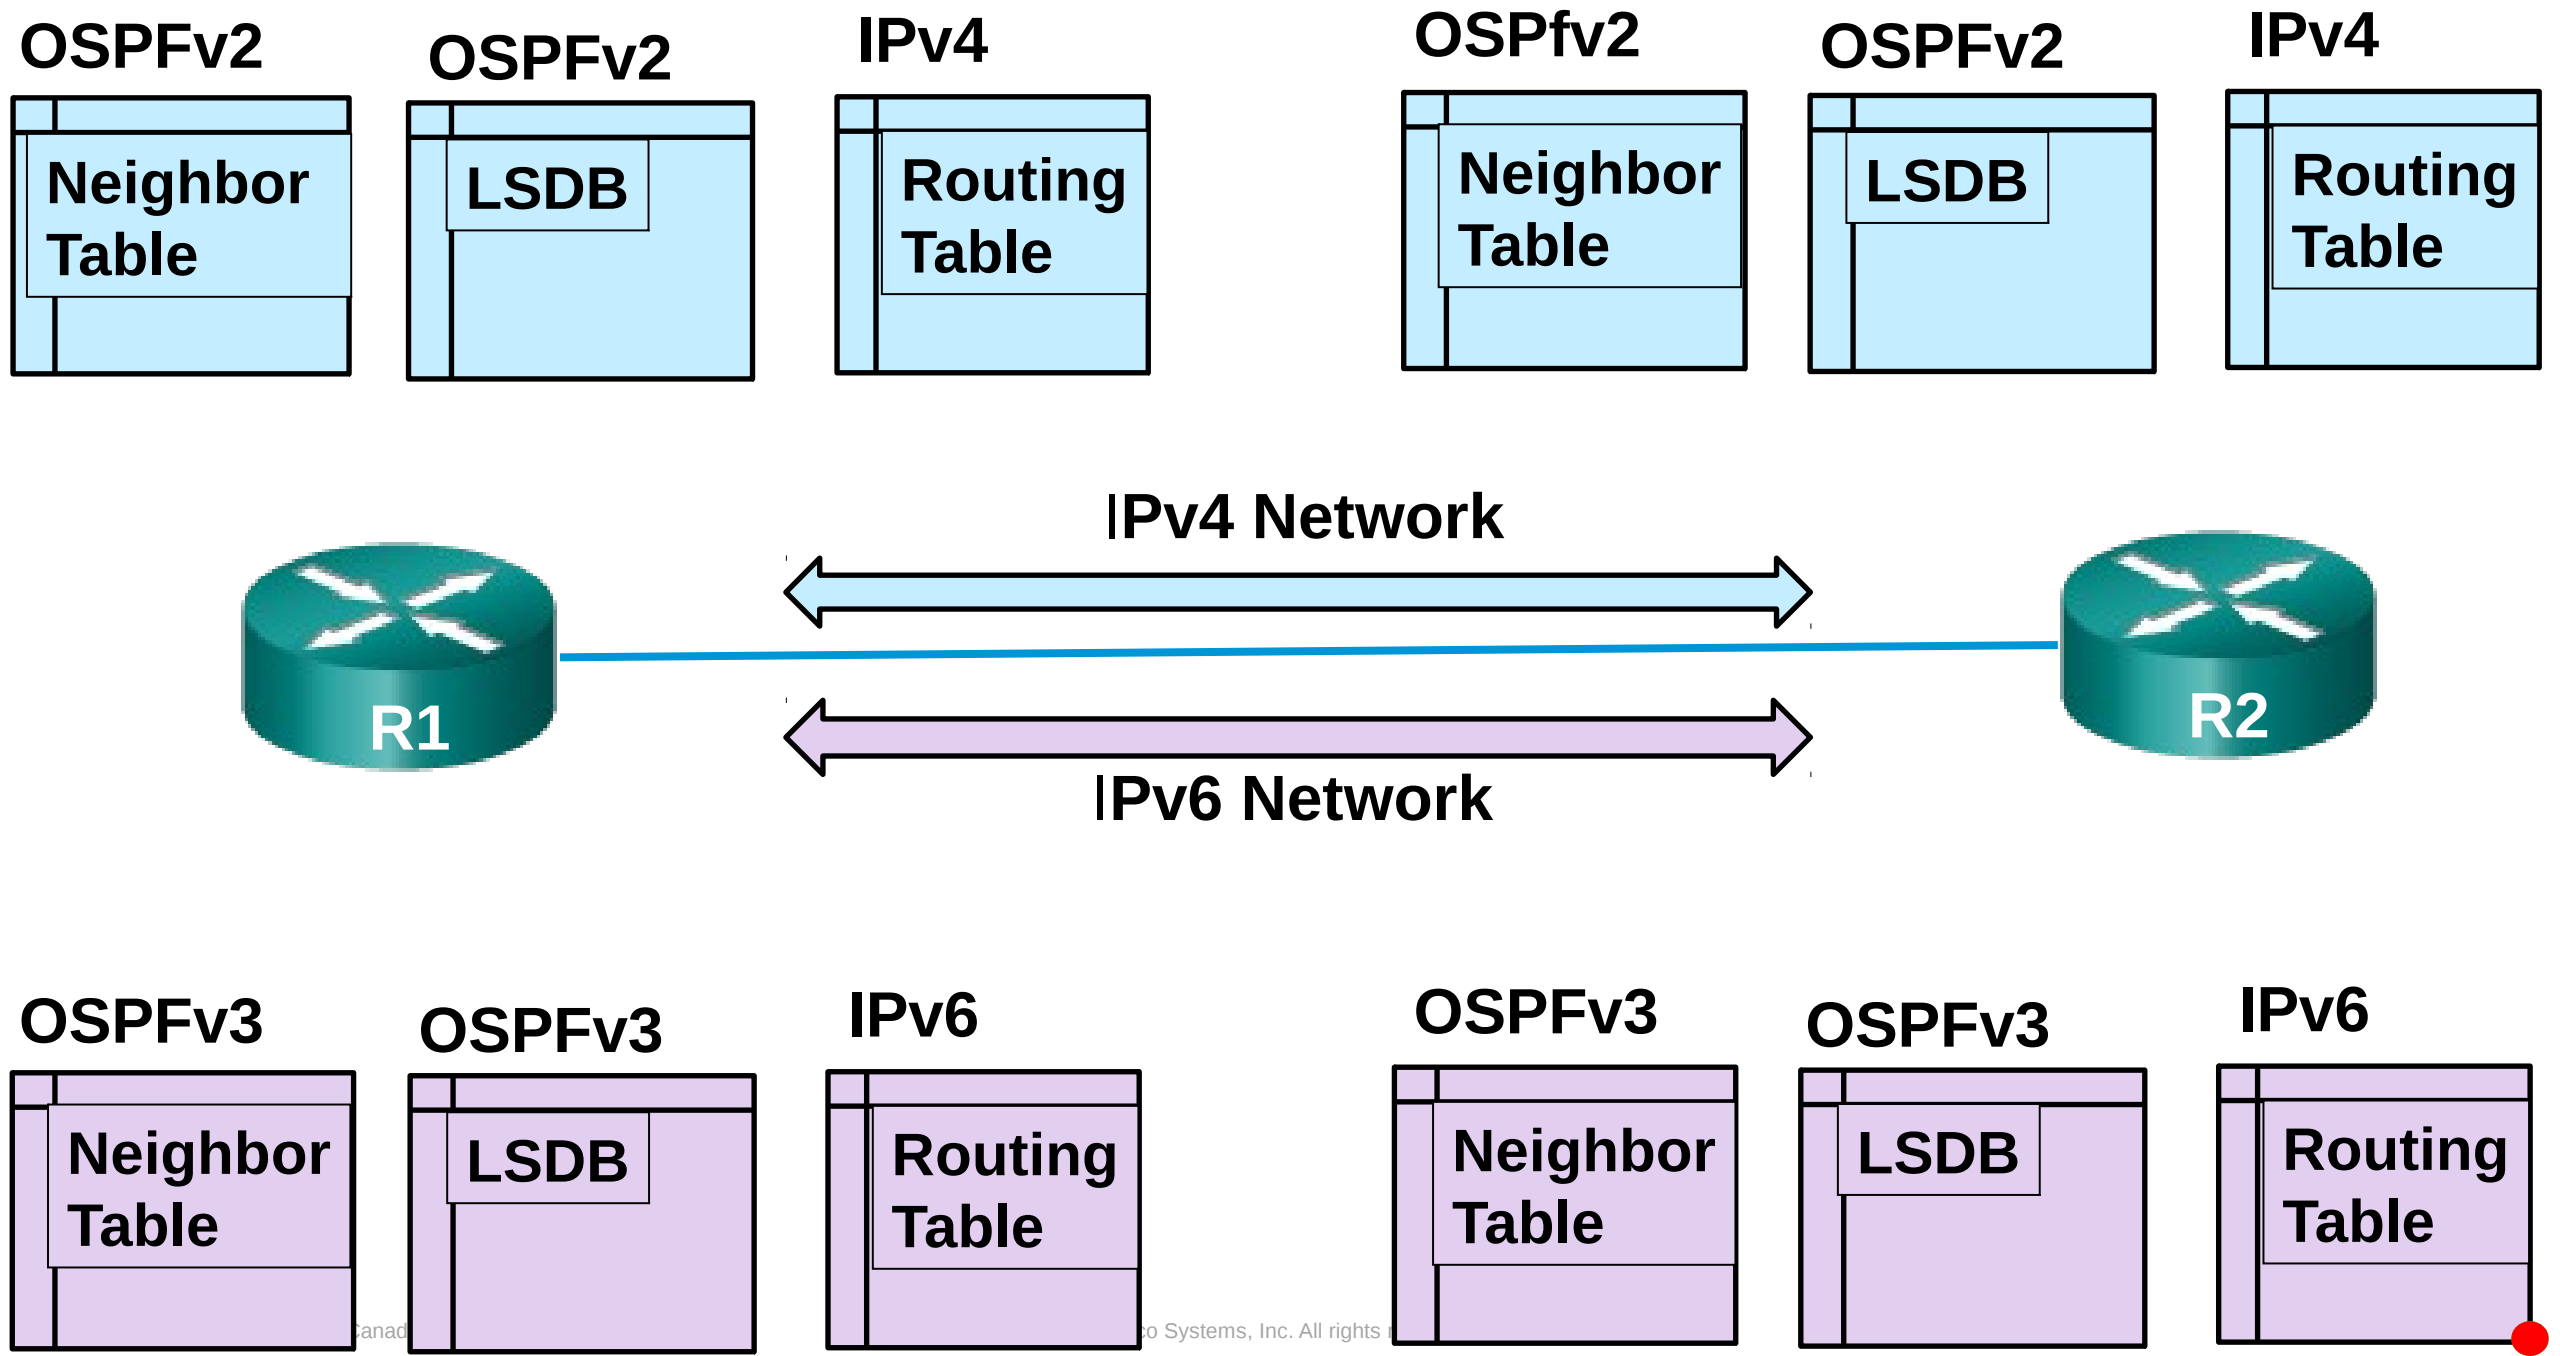

	OSPFv2	OSPFv3
Advertises		
Link-State		
Routing Algorithm		
Metric		
Source address		
Destination address		
Authentication		
Areas		
Packet types		
Neighbor discovery		
DR and BDR		
Router ID		

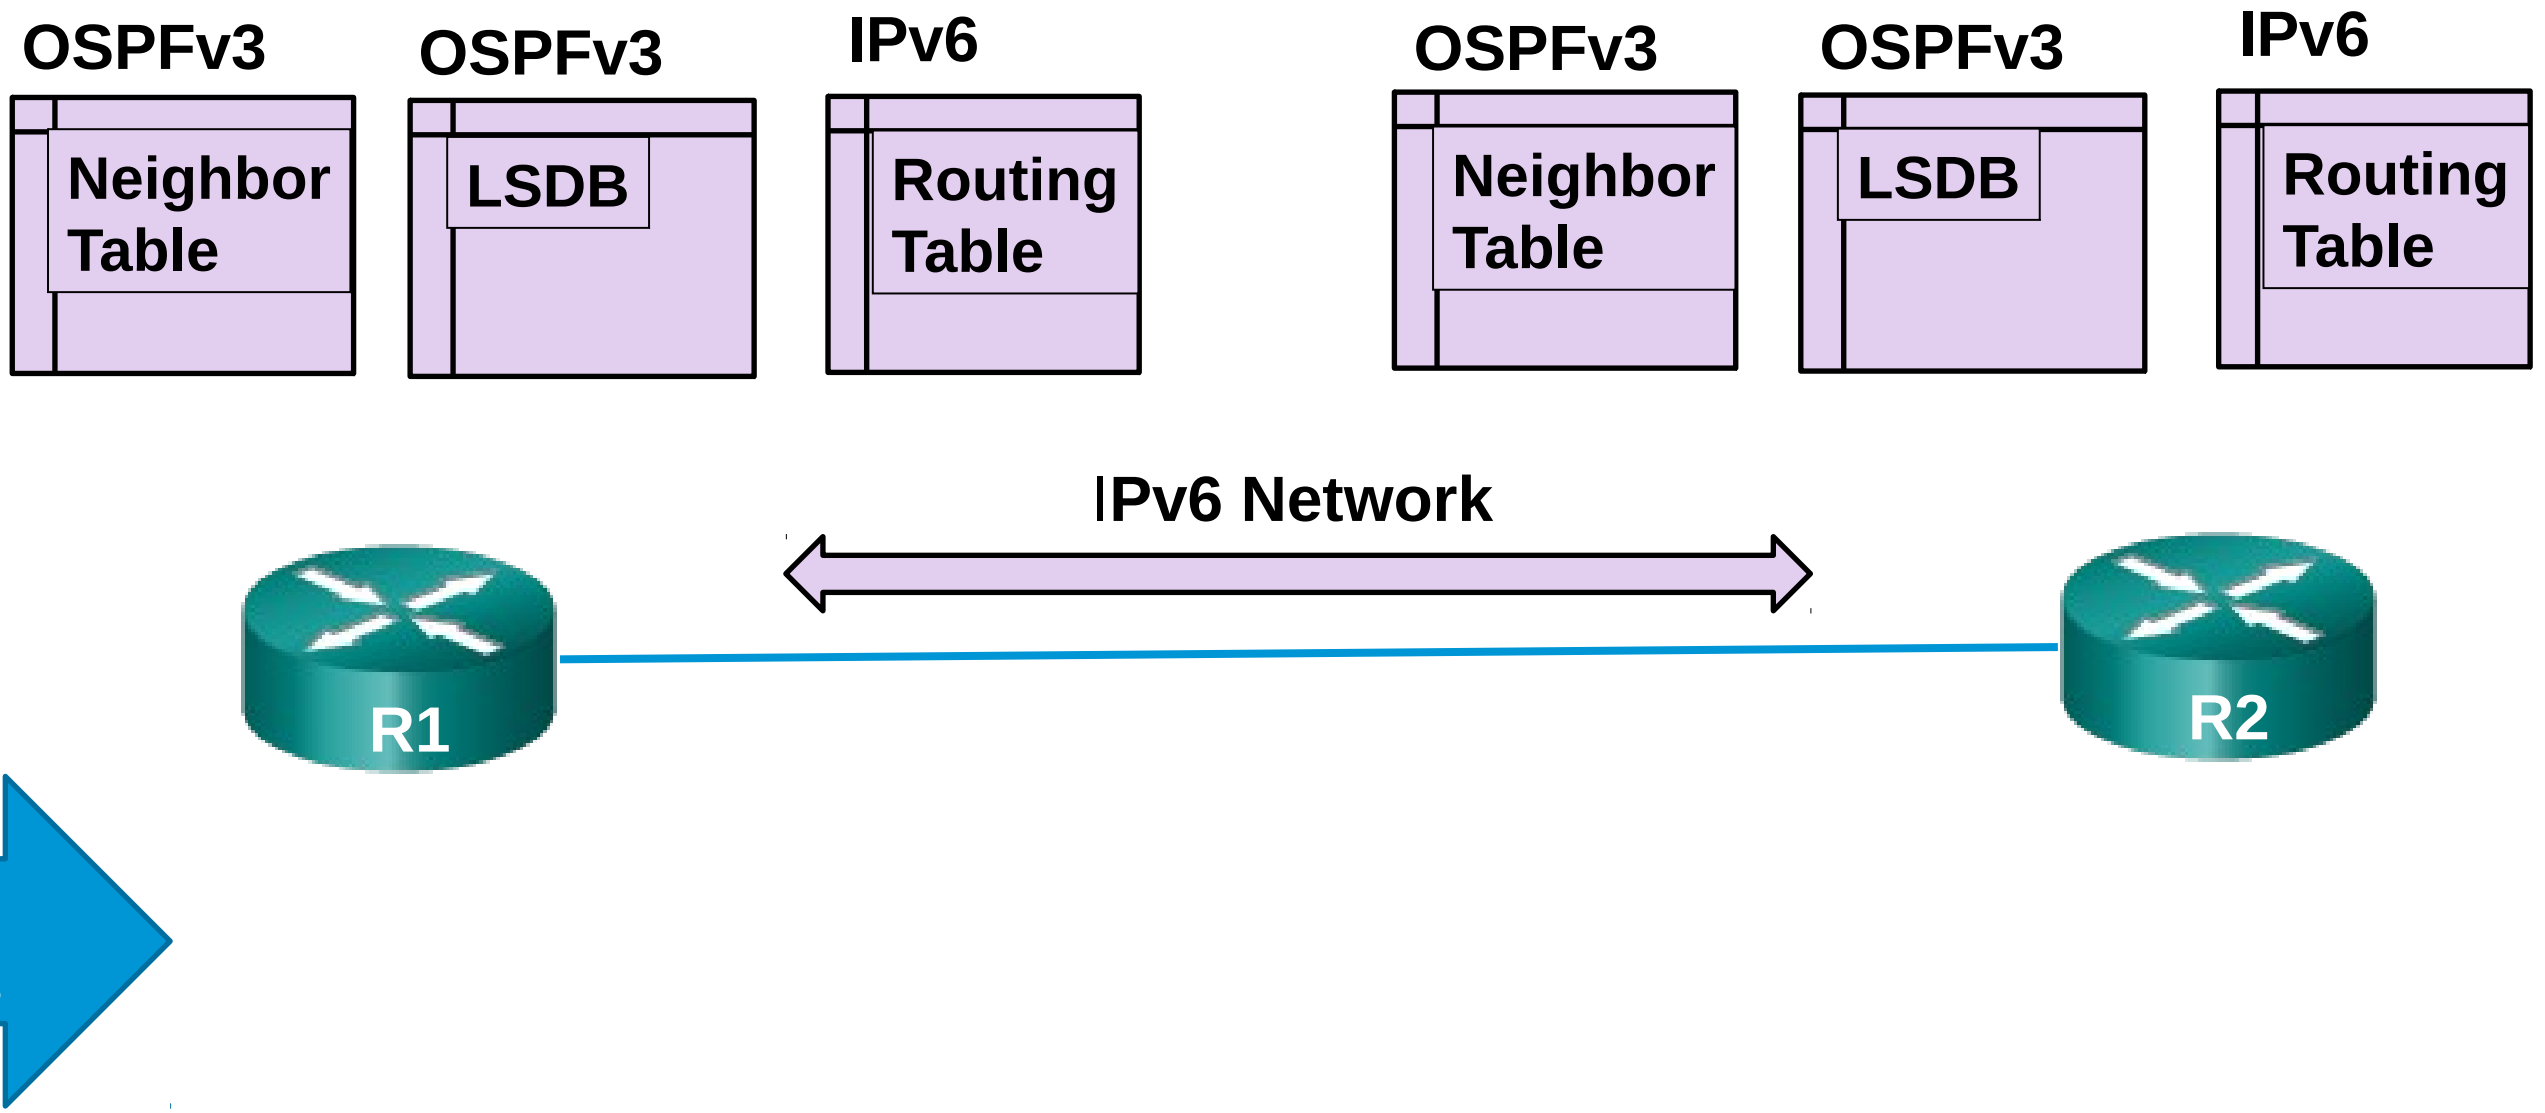

OSPFv3 Topology

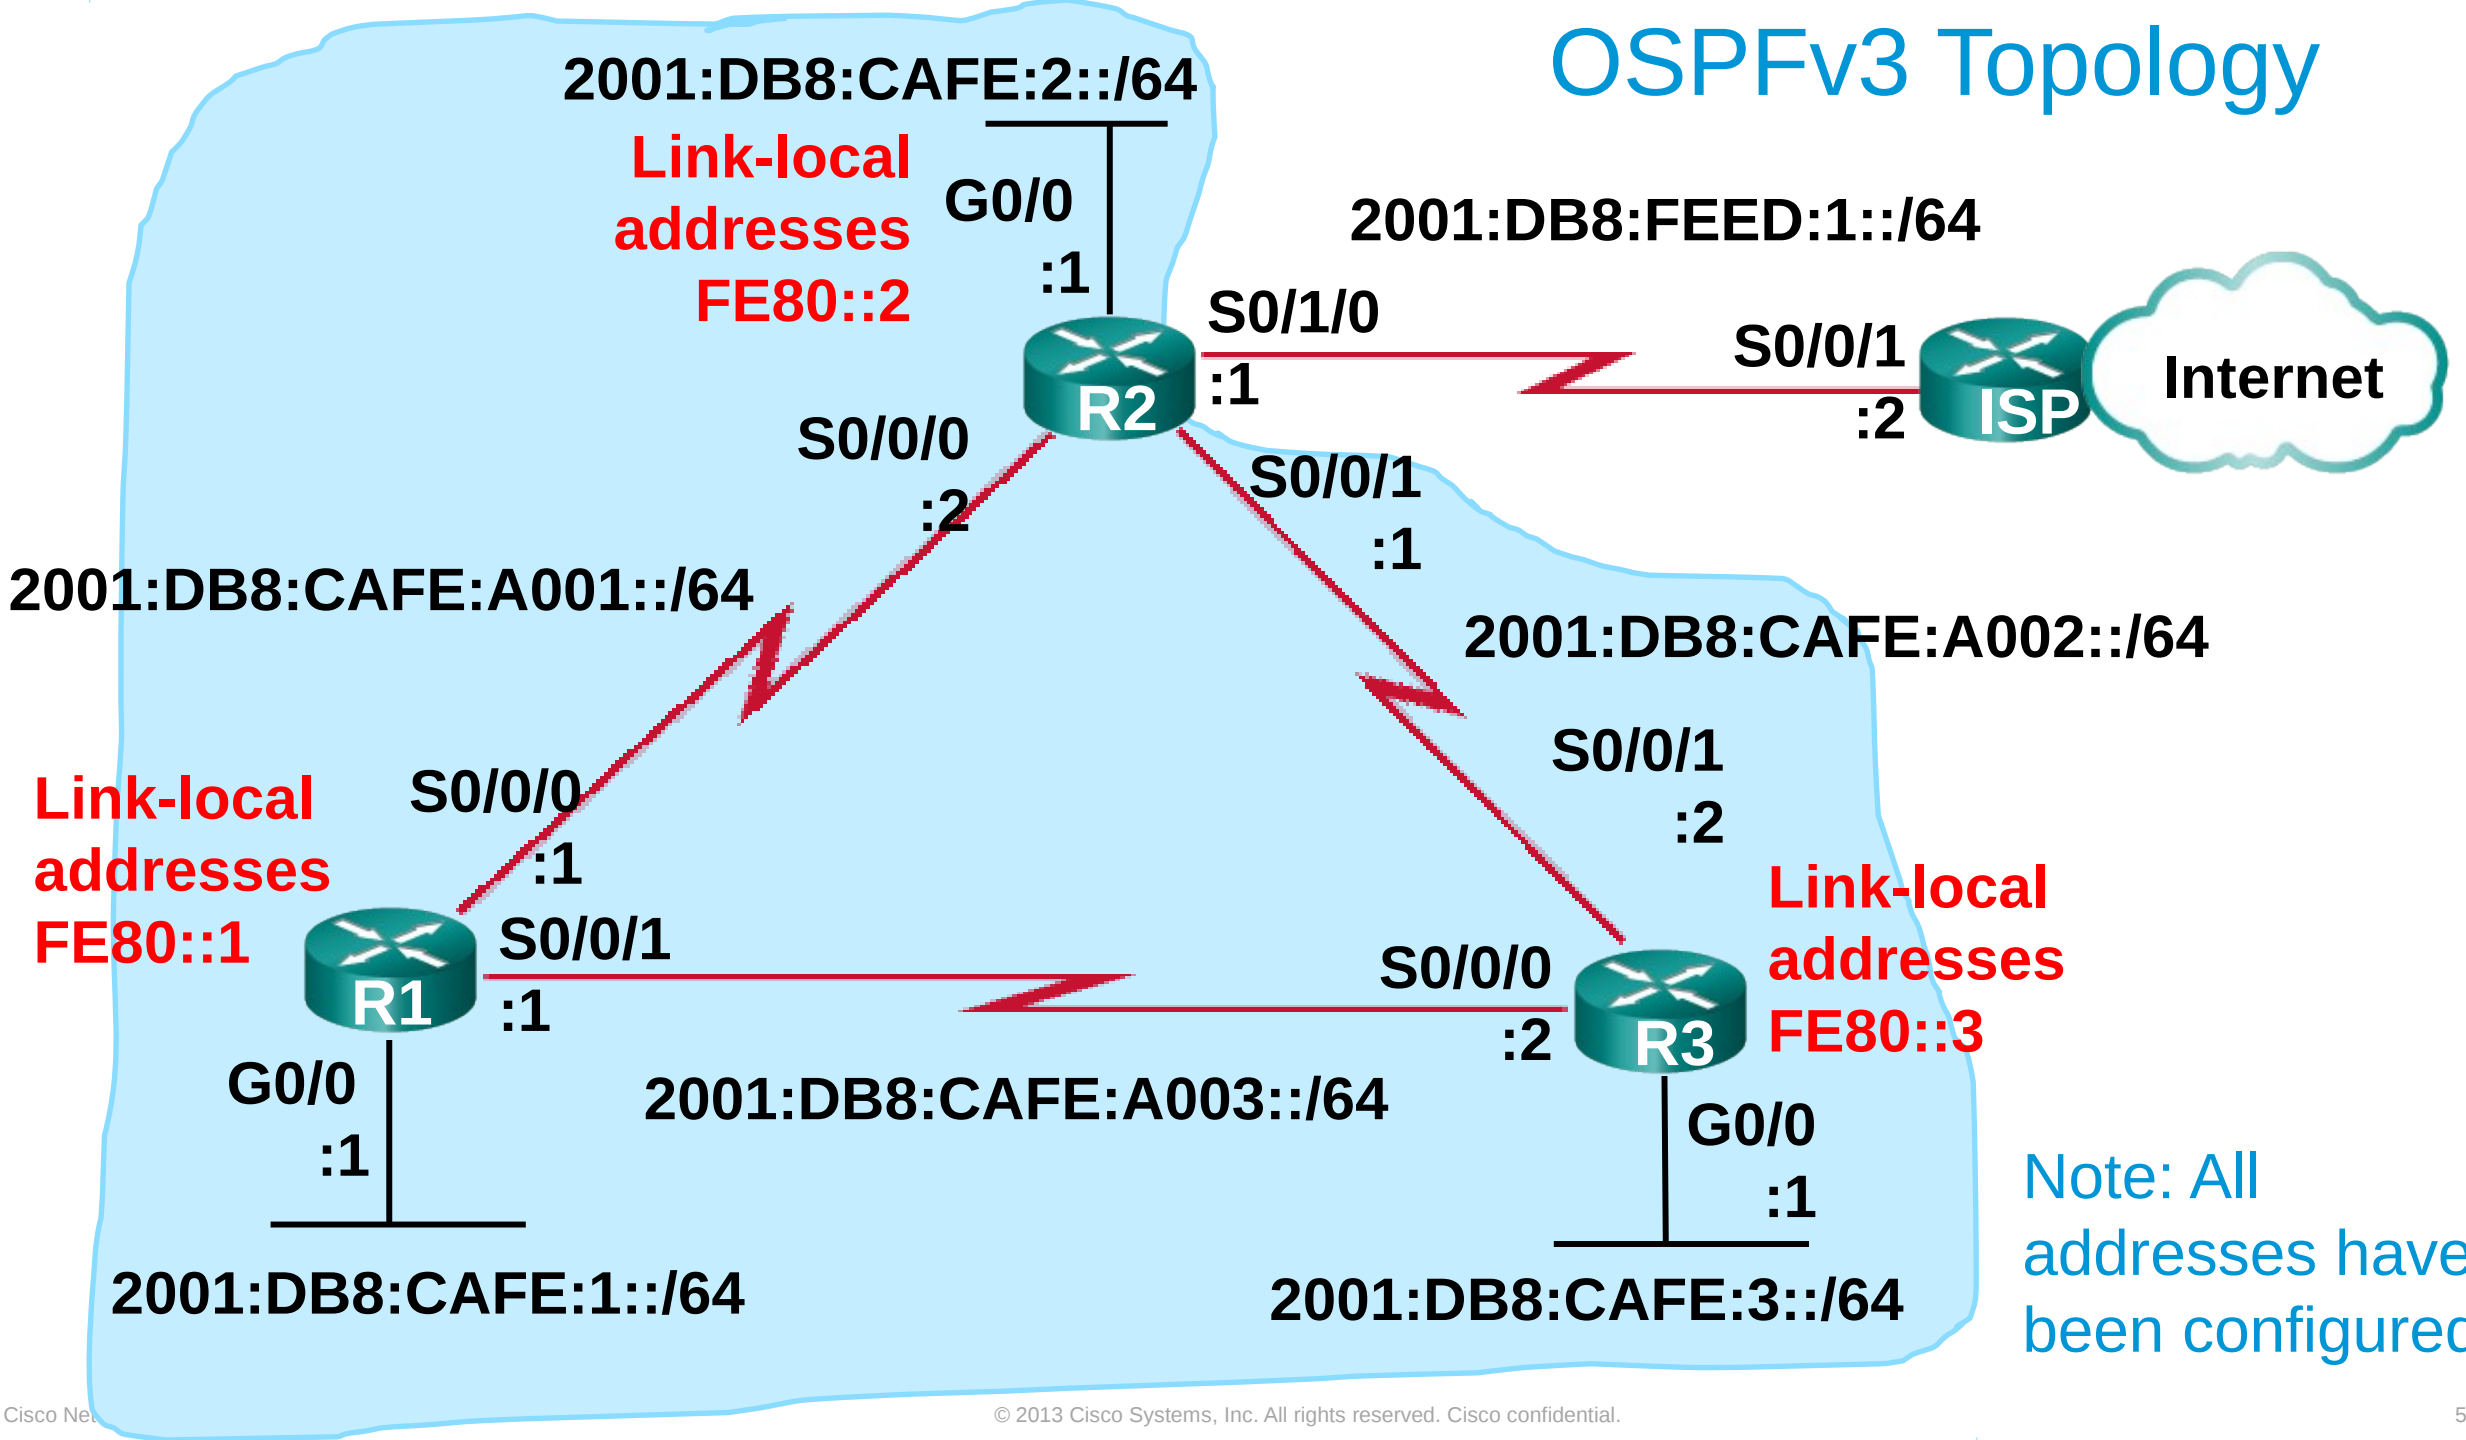

Note: All addresses have been configured

Configuring the OSPFv3 Routing Process

```
R1(config)#ipv6 router ospf 10
R1(config-rtr)#
*Mar 29 11:21:53.739: %OSPFv3-4-NORTRID: Process OSPFv3-1-IPv6
could not pick a router-id, please configure manually
R1(config-rtr)#
R1(config-rtr)#router-id 1.1.1.1
R1(config-rtr)#auto-cost reference-bandwidth 1000
% OSPFv3-1-IPv6:  Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent across
        all routers.
R1(config-rtr)#end
R1#
```

32-bit Router ID similar to OSPFv2

Must modify reference bandwidth because we have gigabit Ethernet links, otherwise Fastethernet and faster would have the same cost

There is no "**no shutdown**" ☺

Auto-cost reference-bandwidth 1000

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
Gigabit Ethernet 1 Gbps	1,000,000,000	÷	1,000,000,000	**1**
Fast Ethernet 100 Mbps	1,000,000,000	÷	100,000,000	**10**
Ethernet 10 Mbps	1,000,000,000	÷	10,000,000	**100**
Serial 1.544 Mbps	1,000,000,000	÷	1,544,000	**647**
Serial 128 kbps	1,000,000,000	÷	128,000	**7812**
Serial 64 kbps	1,000,000,000	÷	64,000	**15625**

Enabling OSPFv3 on the Interface

```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# exit
R1(config)#interface Serial0/0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# exit
R1(config)#interface Serial0/0/1
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# end
R1#
R1#show ipv6 ospf interfaces brief
Interface    PID   Area                Intf ID    Cost  State Nbrs F/C
Se0/0/1      10    0                   7          15625 P2P   0/0
Se0/0/0      10    0                   6          647   P2P   0/0
Gi0/0        10    0                   3          1     WAIT  0/0
R1#
```

No **network** commands needed!

Configure OSPFv3 on the other Routers….

```
R2(config)# router ospf 10
R2(config-rtr)# router-id 2.2.2.2
R1(config-rtr)# auto-cost reference-bandwidth 1000
R2(config-rtr)# exit
R2(config)# interface GigabitEthernet 0/0
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# exit
R2(config)# interface Serial0/0/0
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# exit
R2(config)# interface Serial0/0/1
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)#
```

Verifying OSPFv3 Neighbor Adjacencies

```
R1#show ipv6 ospf neighbor

            OSPFv3 Router with ID (1.1.1.1) (Process ID 10)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
3.3.3.3           0   FULL/  -        00:00:39    6               Serial0/0/1
2.2.2.2           0   FULL/  -        00:00:36    6               Serial0/0/0
R1#
```

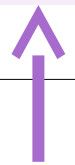

Neighbors' 32-bit Router IDs

Verifying OSPFv3 Parameters

```
R1 #show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Serial0/0/1
    Serial0/0/0
    GigabitEthernet0/0
  Redistribution:
    None
R1#
```

Routing protocol and Process ID

OSPFv3 Router ID

Interfaces enabled for OSPFv3

Verifying OSPFv3 Learned Prefixes

```
R1# show ipv6 route ospf

O   2001:DB8:CAFE:2::/64 [110/657]
     via FE80::2, Serial0/0/0
O   2001:DB8:CAFE:3::/64 [110/1304]
     via FE80::2, Serial0/0/0
O   2001:DB8:CAFE:A002::/64 [110/1294]
     via FE80::2, Serial0/0/0
R1#
```

Link-local addresses are used as next hop addresses

Other commands…. Exactly the same, only different

Hello and Dead Timers remain the same

```
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 ospf hello-interval 5
R1(config-if)# ipv6 ospf dead-interval eigrp 20
R1(config-if)#
*Apr 10 15:03:51.175: %OSPFv3-5-ADJCHG: Process 10, Nbr 2.2.2.2
on Serial0/0/0 from FULL to DOWN, Neighbor Down: Dead timer
expired
R1(config-if)#
```

Just add “v6”!

Changing interface priority (DR/BDR) and cost

```
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 ospf priority 10
R1(config-if)# ipv6 ospf cost 65
```

Propagating a Default Static Route with OSPFv3

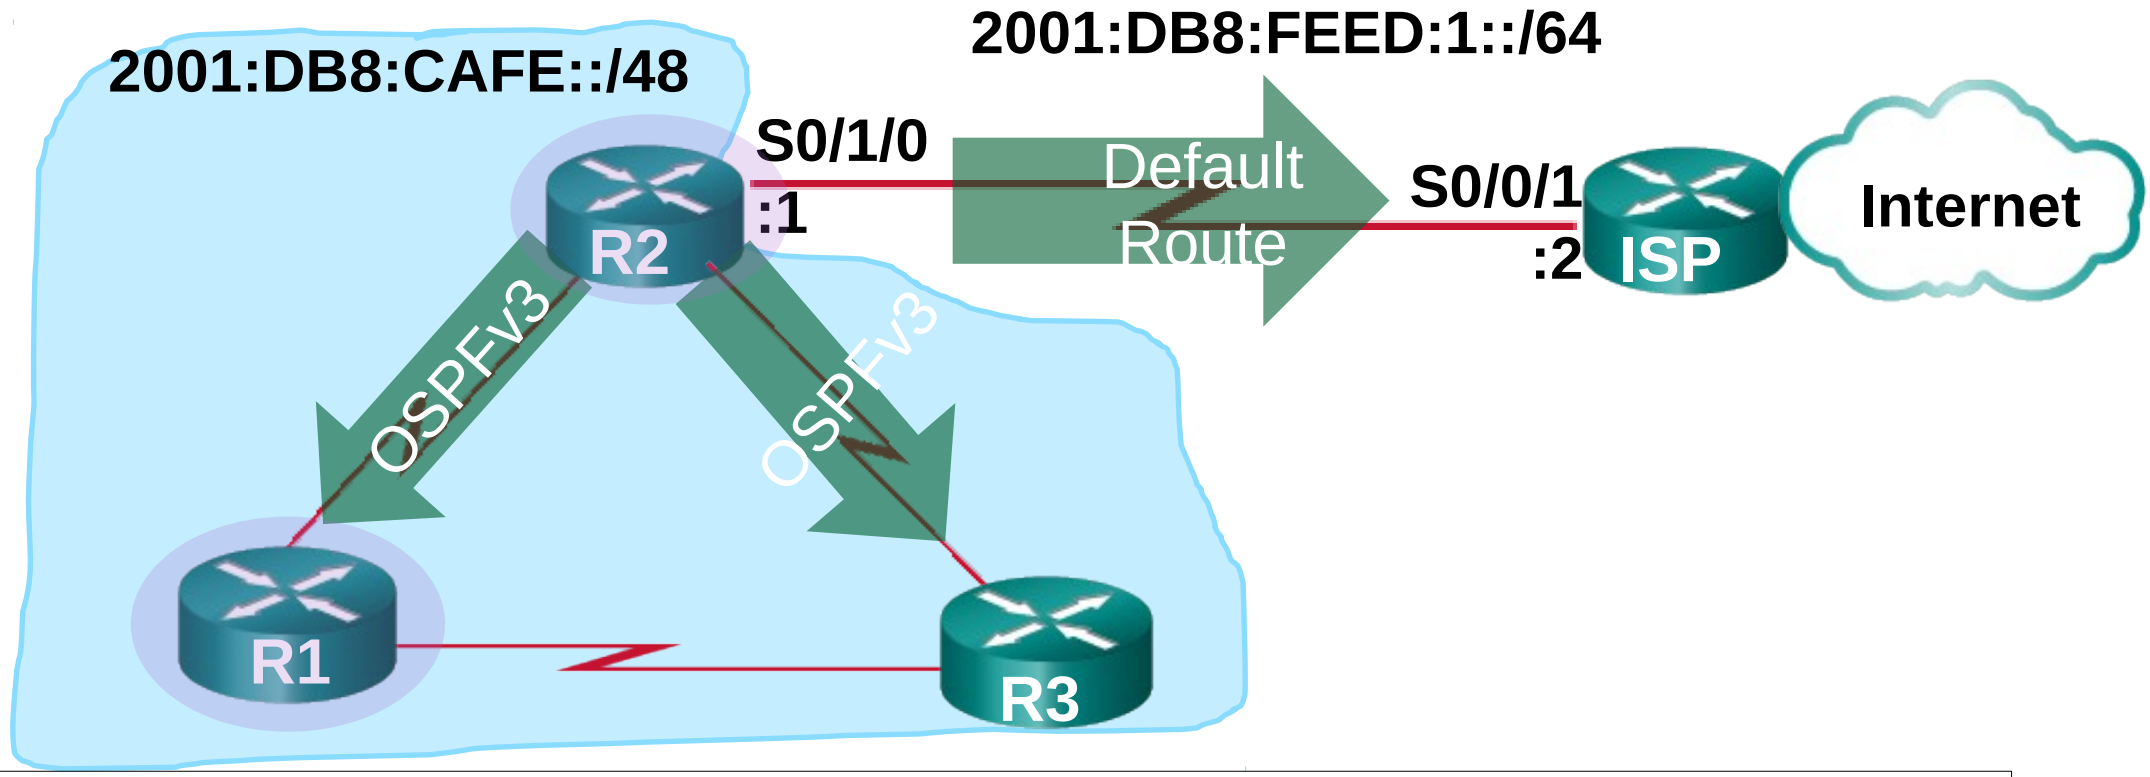


```
R2(config)# ipv6 route ::/0 2001:DB8:FEED:1::2
R2(config)# ipv6 router ospf 10
R2(config-rtr)# default-information originate
```

```
R1# show ipv6 route
OE2 ::/0 [110/1], tag 10
      via FE80::2, Serial0/0/0
```

Multi-Area OSPFv3

Multi-Area OSPF

OSPF is good….
0 SPF is not good!

- Multi-Area OSPF is new (optional) to CCNA
- We don’t have time to cover multi-area OSPF except for the commands associated with OSPFv3
- But here is a quick overview of the concepts…

Issues with a Large OSPF Area

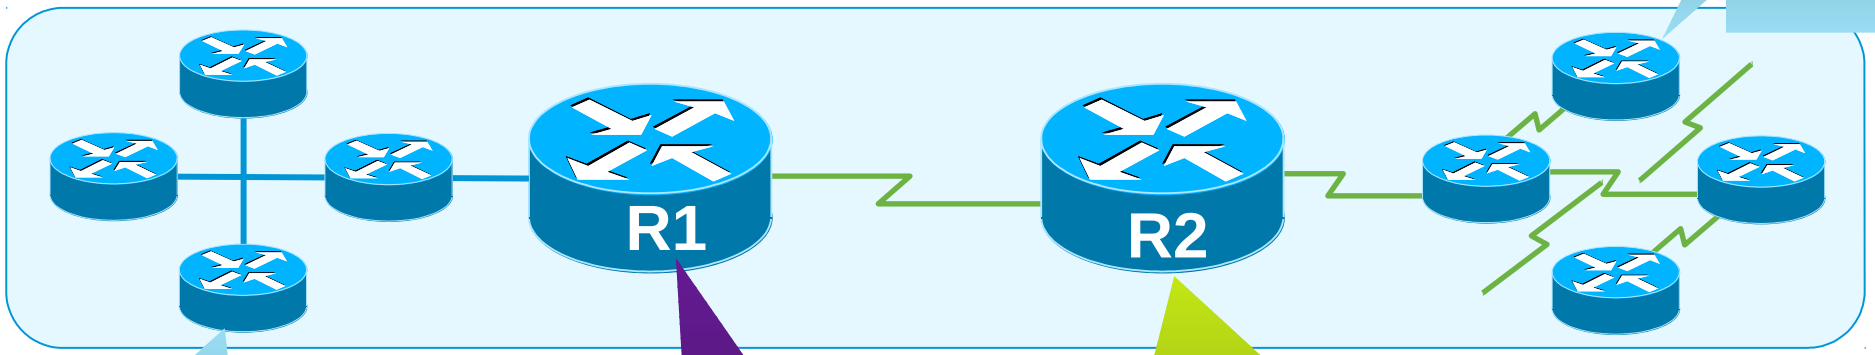

I'm tired of listening to Rick and I need coffee

My routing table is too big and I am running low on memory

I'm receiving too many LSAs

My SPF algorithm is running too often for me to route properly

Advantages of Multi-Area OSPF

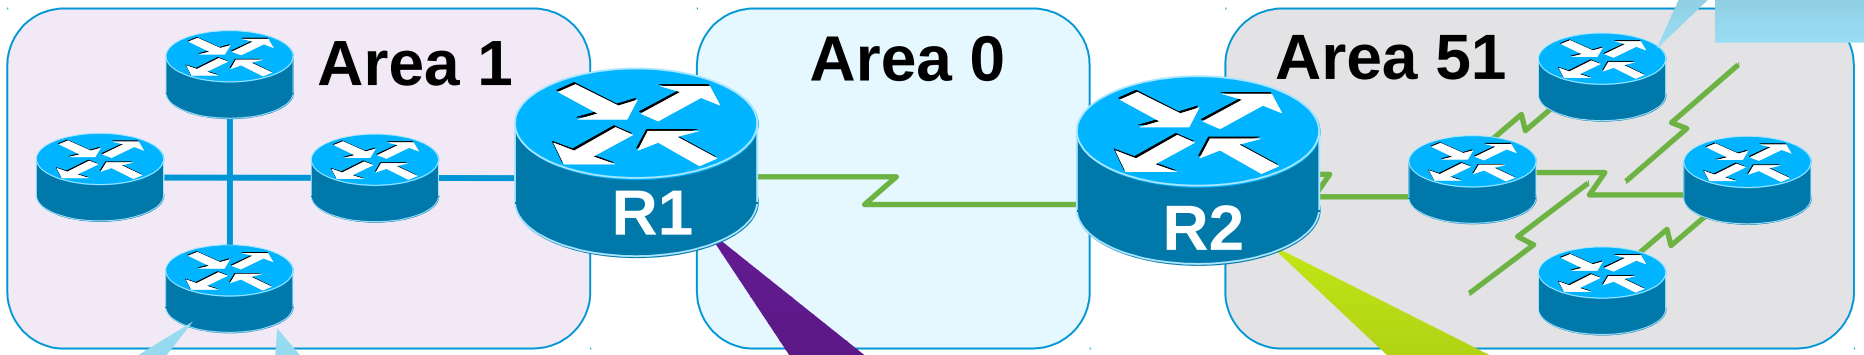

I still need coffee

My routing table is smaller as it does not have external routes (default).

I do not need to rerun my SPF algorithm if there is a change in another area.

I'm now only receiving LSAs from area 0 and area 1

I now only need to run the SPF algorithm when there is a change in area 0 or area 51

Advantages of Multi-Area OSPF

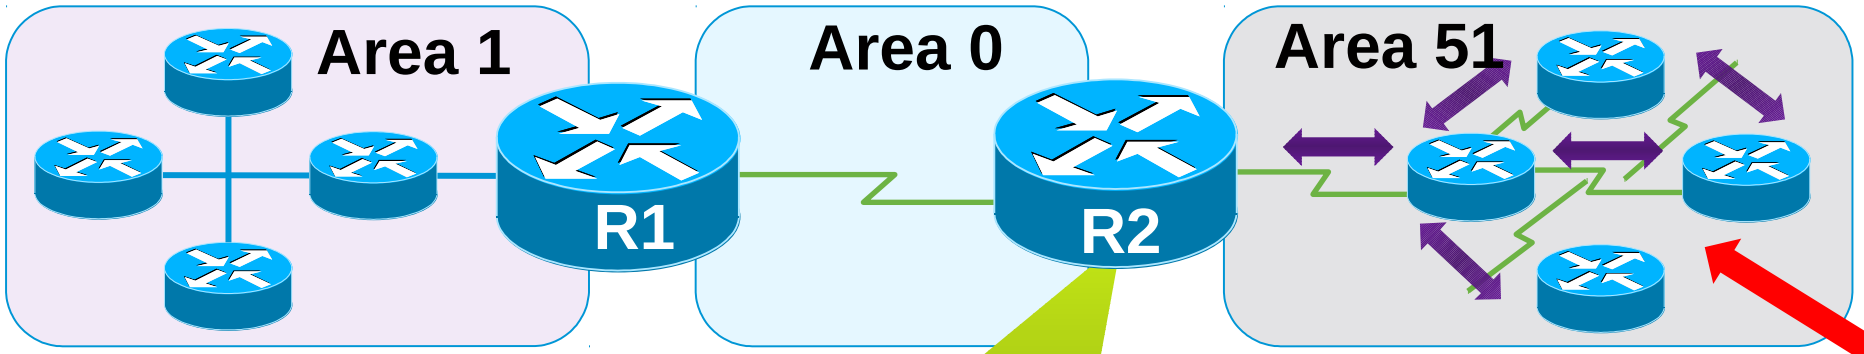

Only R2 and routers in area 51 exchange router LSAs and run the SPF algorithm

Link fails

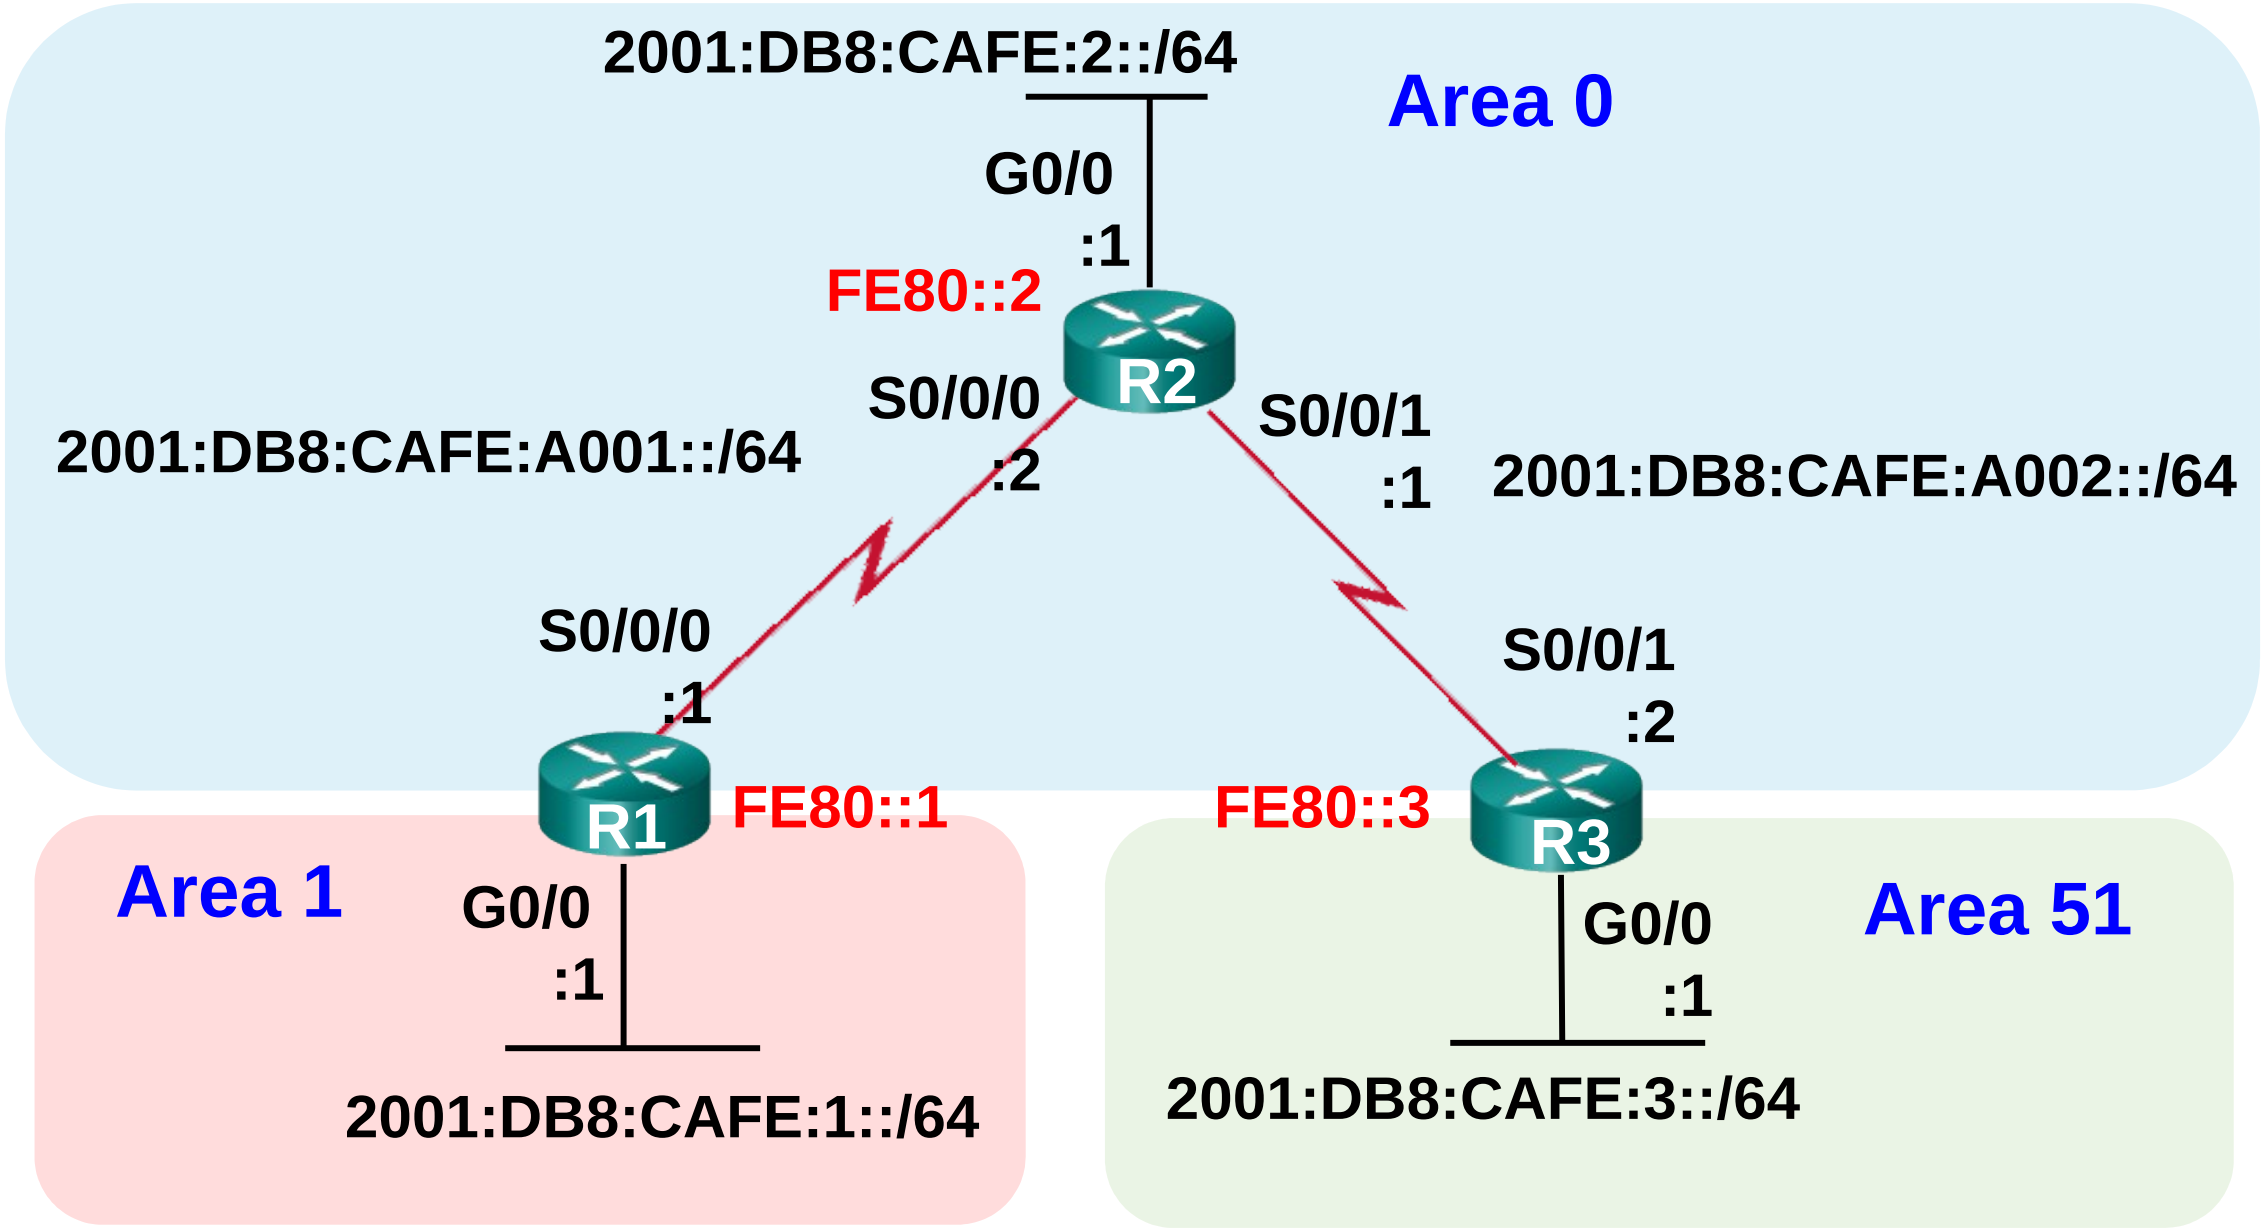

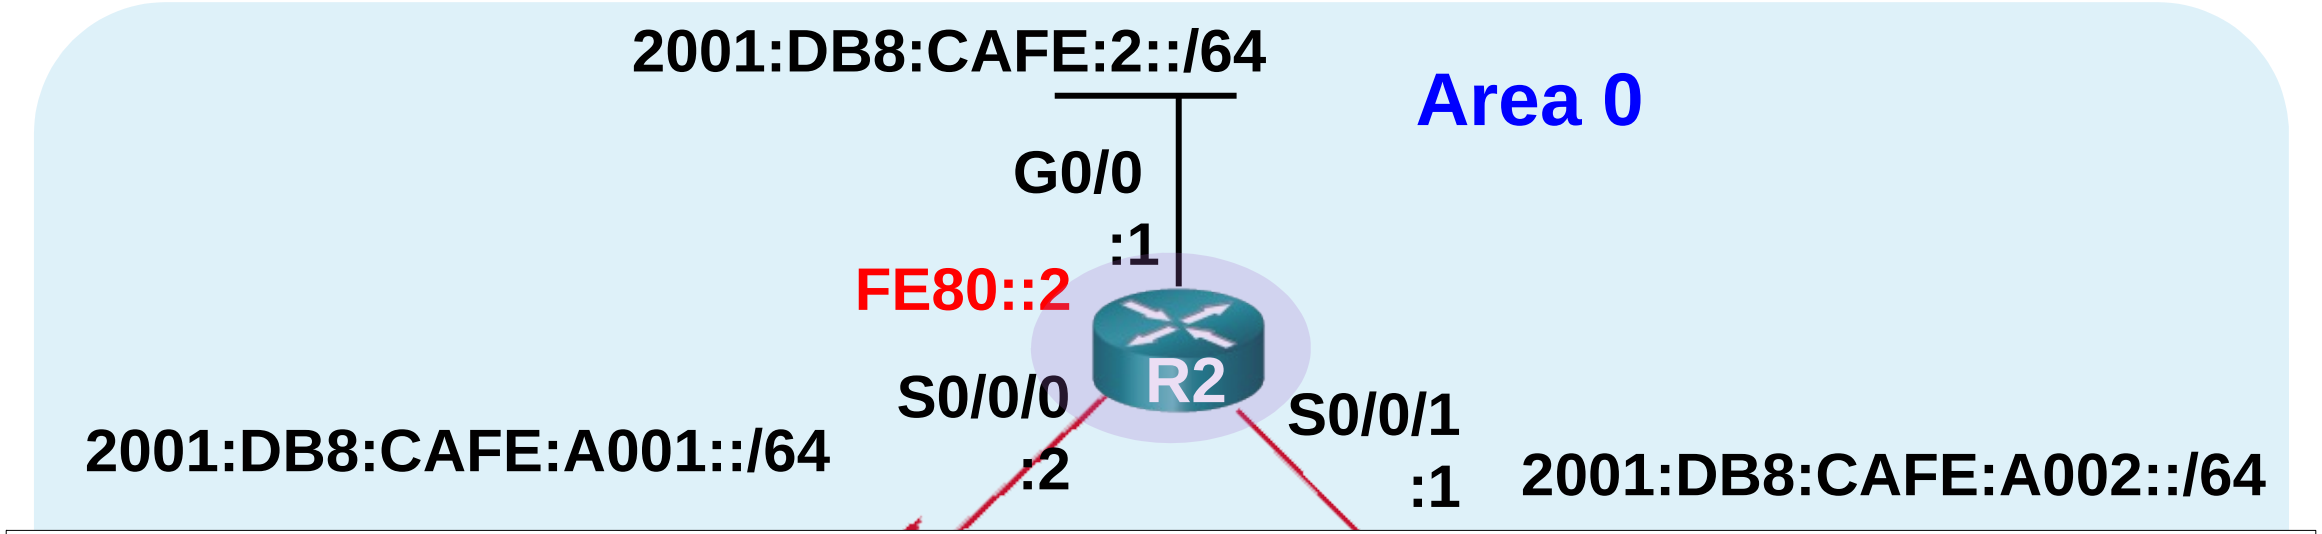


```
R2(config)# router ospf 10
R2(config-rtr)# router-id 2.2.2.2
R2(config-rtr)# exit
R2(config)# interface GigabitEthernet 0/0
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# exit
R2(config)# interface Serial0/0/0
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# exit
R2(config)# interface Serial0/0/1
R2(config-if)# ipv6 ospf 10 area 0
```

No changes to R2
All interfaces in Area 0

```
R1(config)# ipv6 router ospf 10
R1(config-rtr)# router-id 1.1.1.1
R1(config-rtr)# exit
R1(config)# interface Serial 0/0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# exit
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ipv6 ospf 10 area 1
R1(config-if)#
```

An ABR (Area Border Router) has interfaces in more than one area

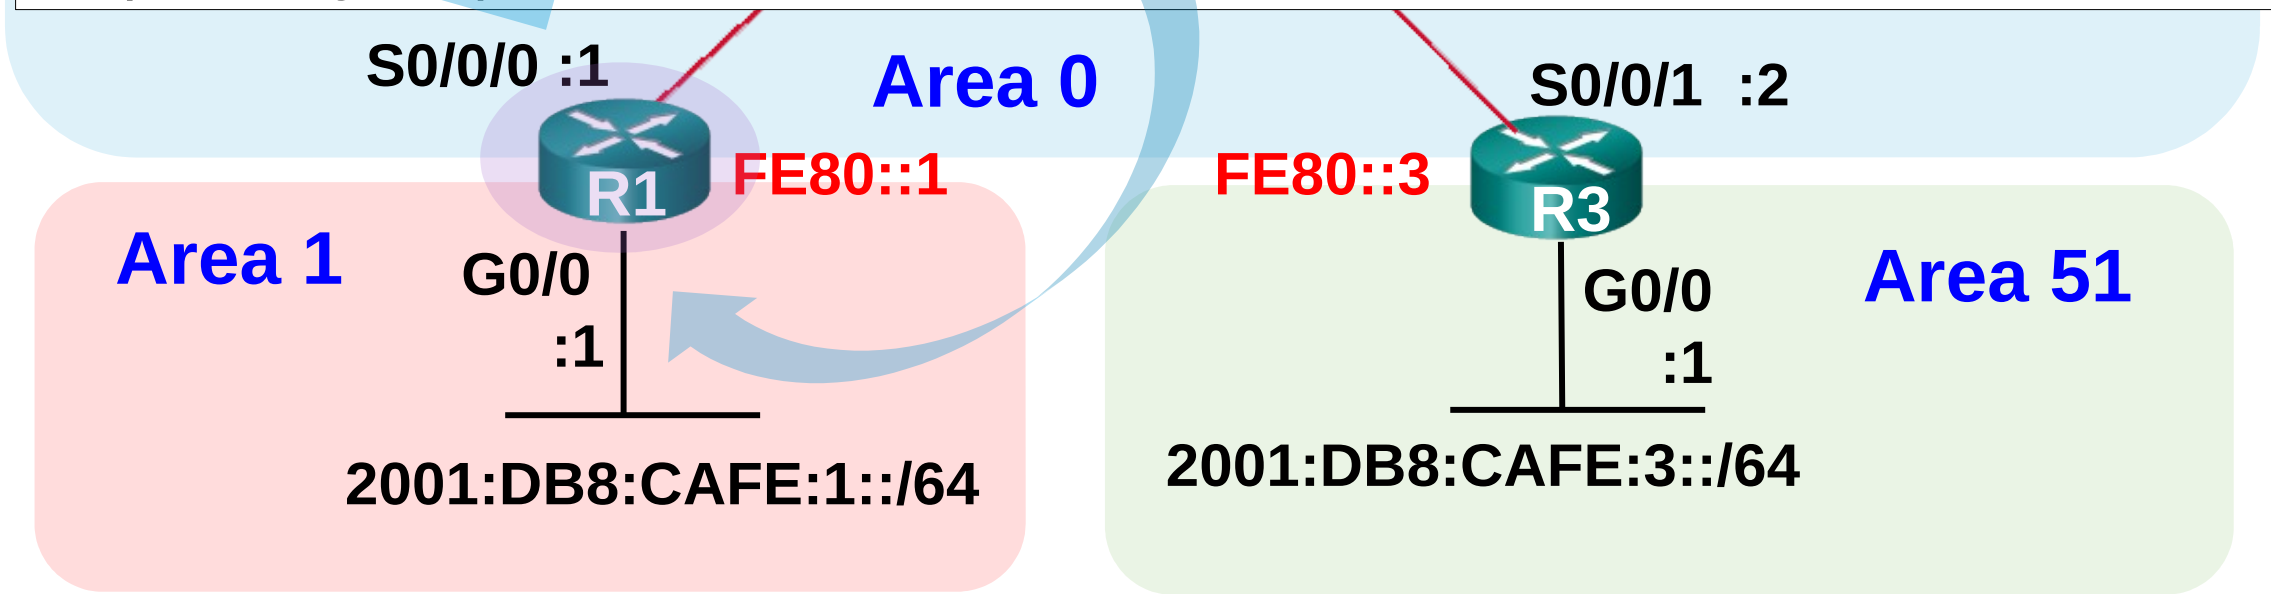

IPv6 Access Control Lists

IPv6 Access Control Lists

IPv4 ACLs

- Standard
  - Numbered
  - Named
- Extended
  - Numbered
  - Named

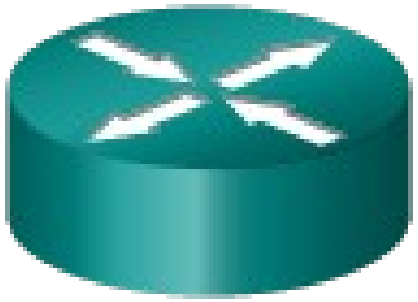

IPv6 ACLs

- Named only
- Similar features to Extended ACLs

IPv6 ACLs are very similar to IPv4 ACLs but with three significant differences

1. Applying an IPv6 ACL to an interface:

- **`ip access-group`** is used in IPv4
- **`ipv6 traffic-filter`** is used to apply an IPv6 ACL to an IPv6 interface

2. No Wildcard Masks

- Prefix-lengths are used instead of wildcard masks

3. Two additional default statements (besides implicit **deny any**)

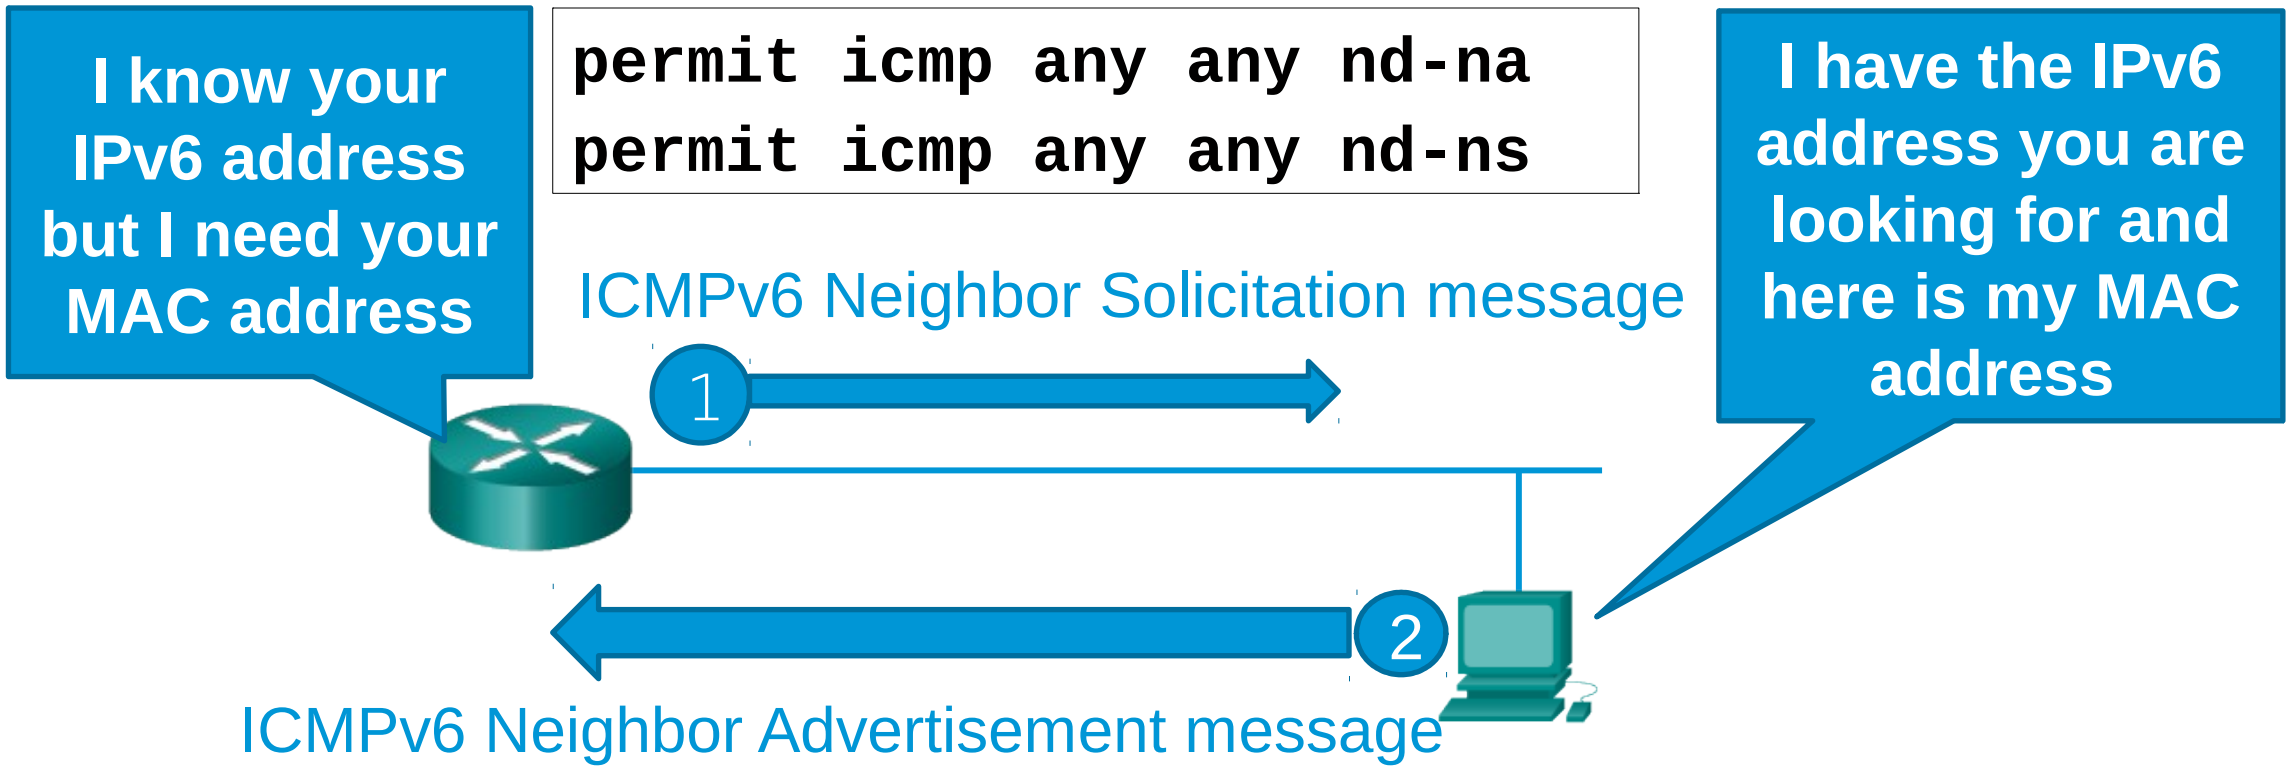

- Two default statements allow the router to participate in the IPv6 equivalent of ARP
- IPv6 uses ICMPv6 Neighbor Discovery (ND) messages to accomplish name address resolution encapsulated in IPv6 packets
- ARP does not use IPv4
- IPv6 ACLs need to implicitly permit ND packets to be sent and received on an interface

Example 1: Deny an IPv6 Prefix

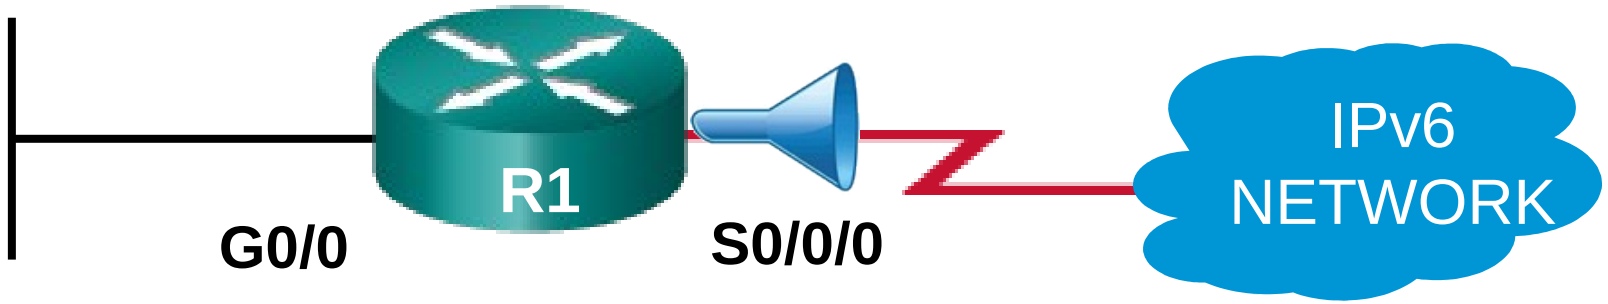


```
R1(config)# ipv6 access-list NO-R3-LAN-ACCESS
R1(config-ipv6-acl)# deny ipv6 2001:db8:cafe:30::/64 any
R1(config-ipv6-acl)# permit ipv6 any any
R1(config-ipv6-acl)# exit
R1(config)# interface s0/0/0
R1(config-if)# ipv6 traffic-filter NO-R3-LAN-ACCESS in
R1(config-if)#
```

- Deny all IPv6 packets from the 2001:DB8:CAFE:30::/64 coming into Serial 0/0/0
- Permit all other IPv6 packets

Example 2: Deny FTP Traffic to an IPv6 Prefix


```
R1(config)#ipv6 access-list NO-FTP-TO-11
R1(config-ipv6-acl)#deny tcp any 2001:db8:cafe:11::/64 eq ftp
R1(config-ipv6-acl)#deny tcp any 2001:db8:cafe:11::/64 eq ftp-data
R1(config-ipv6-acl)#permit ipv6 any any
R1(config-ipv6-acl)#exit
R1(config)#interface g0/0
R1(config-if)#ipv6 traffic-filter NO-FTP-TO-11 in
R1(config-if)#
```

- Deny FTP traffic from Gig0/0 to 2001:DB8:CAFE:11::/64

What we covered…

- Quick review of IPv6
- IPv6 Static Routes
- EIGRP for IPv6
- OSPFv3
- Multi-Area OSPFv3
- IPv6 Access Control Lists

Web Site, Book, Etc.

- Rick Graziani - graziani@cabrillo.edu
- PowerPoints for CCNA, CCNP, IPv6
  - www.cabrillo.edu/~rgraziani
  - Username = cisco
  - Password = perlman

Shameless plug!

Quality time with my two nieces…

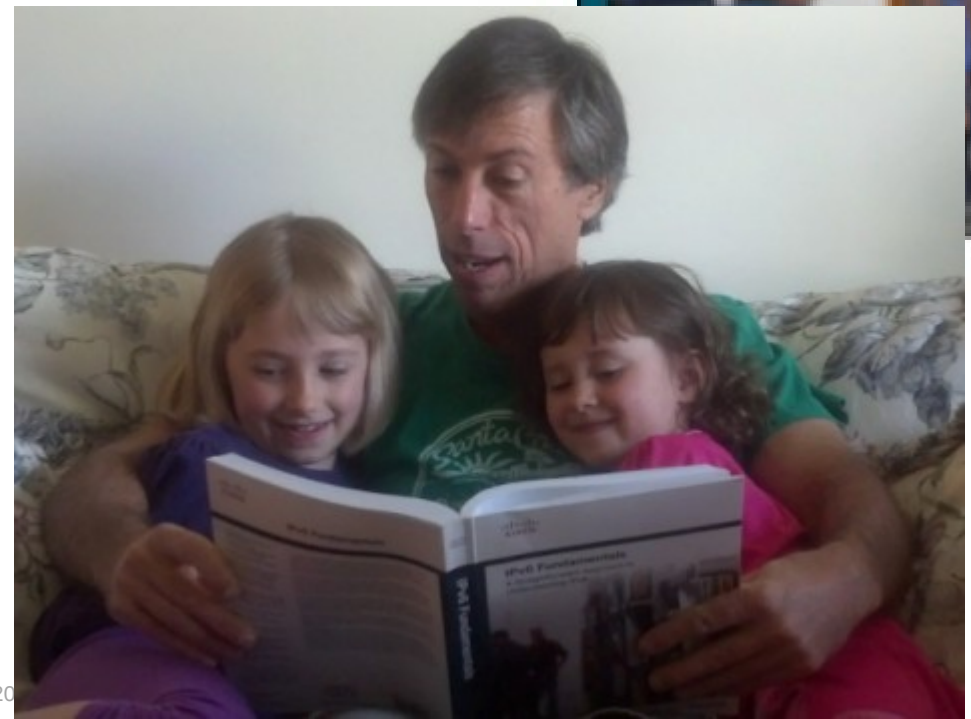

We Appreciate Your Feedback

- Please take a moment to complete the evaluation
- Our virtual audience can click on the Evaluation icon  at the bottom of the screen to access the evaluation form
- Our audience here in San Jose will receive an email after the session with a link to the survey

And…… Thank you very much!

Rick Graziani - graziani@cabrillo.edu

www.cabrillo.edu/~rgraziani

Username = cisco

Password = perlman