Course Introduction

Deploying Cisco Service Provider Network Routing (SPROUTE) v1.2

Learner Skills and Knowledge

- Students considered for this training will have attended the following classes or obtained equivalent level training:
  - *Building Cisco Service Provider Next-Generation Networks, Part 1* (SPNGN1) v1.2
  - *Building Cisco Service Provider Next-Generation Networks, Part 2* (SPNGN2) v1.2

Course Goal

Upon completing this course, you will be able to:

- Train service provider network professionals on the techniques to plan, implement, and monitor scalable IP routing

Course Flow

	Day 1	Day 2	Day 3	Day 4	Day 5
AM	Course Introduction Service Provider Routing	Implement OSPF in the Service Provider Network (Cont.)	Implement Integrated IS-IS in the Service Provider Network (Cont.)	Routing Protocol Tools and Route Manipulation	Routing Protocol Tools and Route Manipulation (Cont.)
			Lunch		
PM	Implement OSPF in the Service Provider Network	Implement OSPF in the Service Provider Network (Cont.) Implement Integrated IS-IS in the Service Provider Network	Implement BGP in the Service Provider Network	Routing Protocol Tools and Route Manipulation (Cont.)	Routing Protocol Tools and Route Manipulation (Cont.)

Cisco Career Certifications

Cisco Career Certifications

Expand Your Professional Options, Advance Your Career

Cisco CCNP Service Provider

- *Deploying Cisco Service Provider Network Routing* (SPROUTE) v1.2
- *Deploying Cisco Service Provider Advanced Network Routing (SPADVROUTE) v1.2*
- *Implementing Cisco Service Provider Next-Generation Core Network Services* (SPCORE) v1.2
- *Implementing Cisco Service Provider Next-Generation Edge Network Services* (SPEDGE) v1.2

CCNA Prep Center

HOME
LEARNING AND EVENTS
CAREER CERTIFICATIONS AND PATHS
CERTIFICATION RESOURCES
CCNA Prep Center

Career Certifications & Paths

CCNA Prep Center

Get Ready for Your CCNA Certification

CCNA Prep Center provides certification candidates with resources including practice questions, labs, simulations, instructional videos, tips, advice, success stories and peer-discussion forums. The CCNA Prep Center also allows learners easy access to information and formal training from Cisco Learning Partners.

Search CCNA Prep Center

CCNA Prep Center

GO

Advanced Search

Print | Feedback | Help

Overview | CCNA Paths | Exam Study | CCNA TV | Discussions | Additional Information

Planning Your CCNA Preparation

CCNA Paths
CCNA Certification Paths, My Certification History, Pre-Assessment Exam and Exam Registration.

Exam Study
Study Tips, Practice Questions, Remote Labs, Simulations and more.

CCNA TV
Regular broadcasts with experts discussing CCNA topics and answering your questions.

Discussions
Engage in discussions with Cisco experts on technical questions or program issues.

Additional Information
Articles and information about certifications and the job market.

Featured Content

Cisco Subnet Game Beta
Sharpen your subnetting skills as you set up subnetworks in the mysterious Area 51 Military Air Base.
Play the Game
About the Game

Related Tools

Certifications Online Support
Certifications Tracking System Tool
Global Learning Partner Locator
Cisco Learning Connection
Products & Services Tool Index

Related Links

CCNP Prep Center
Certifications Community
Networking Academy

Learning and Events
Career Certification and Paths
About Learning Partners

About Cisco
Cisco Press

Your Opinion Counts

Please take a moment to tell us your thoughts about the CCNA Prep Center.

Learner Introductions

- Your name
- Your company
- Job responsibilities
- Skills and knowledge
- Brief history
- Objective

Cisco Icons and Symbols

Router

Wireless Router

Secure Router

Firewall

Home Office

Workgroup Switch

Access Point

Cisco IOS XR Router

Cisco IOS XE Router

Wireless Connectivity

Line: Serial

Line: Ethernet

Understanding Service Provider Routing Protocols

Service Provider Routing

Cisco IP NGN Architecture

- The Cisco IP NGN is a next-generation service provider infrastructure for video, mobile, and cloud or managed services.
- It provides an all-IP network for services and applications, regardless of access type.

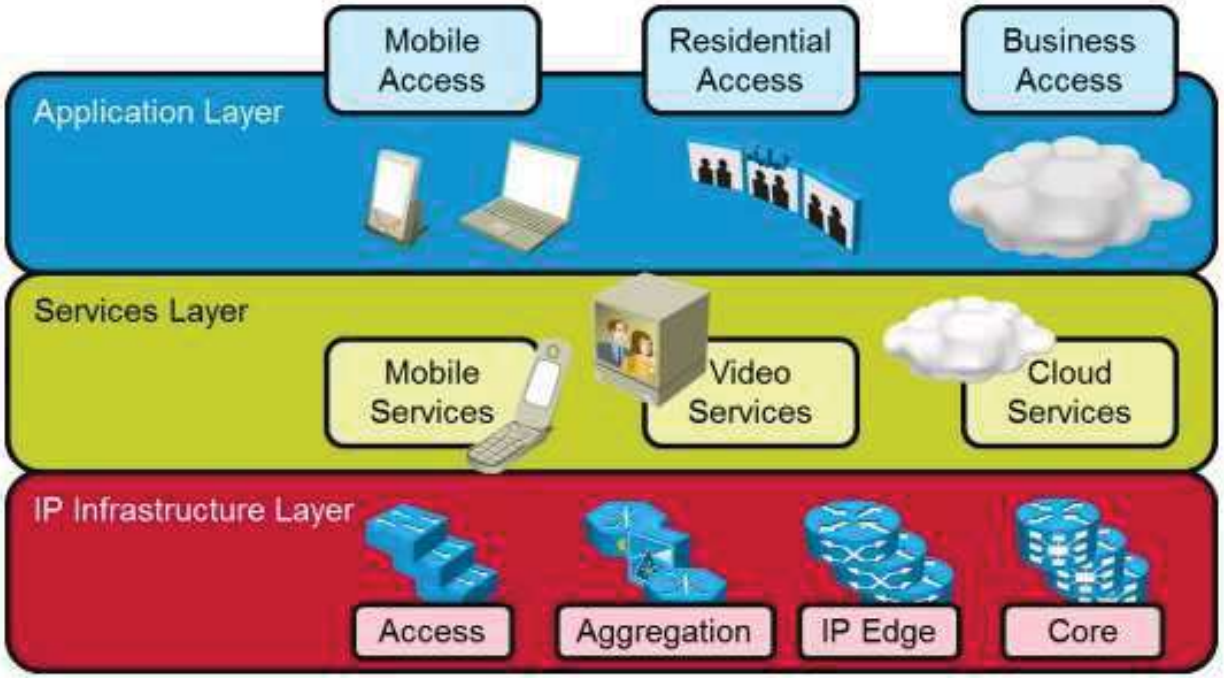

Cisco IP NGN Architecture (Cont.)

Routing protocols that are used in service provider environments focus on the following:

- *IP infrastructure layer* of the Cisco IP NGN
- Service provider *core* and *edge* devices and customer devices

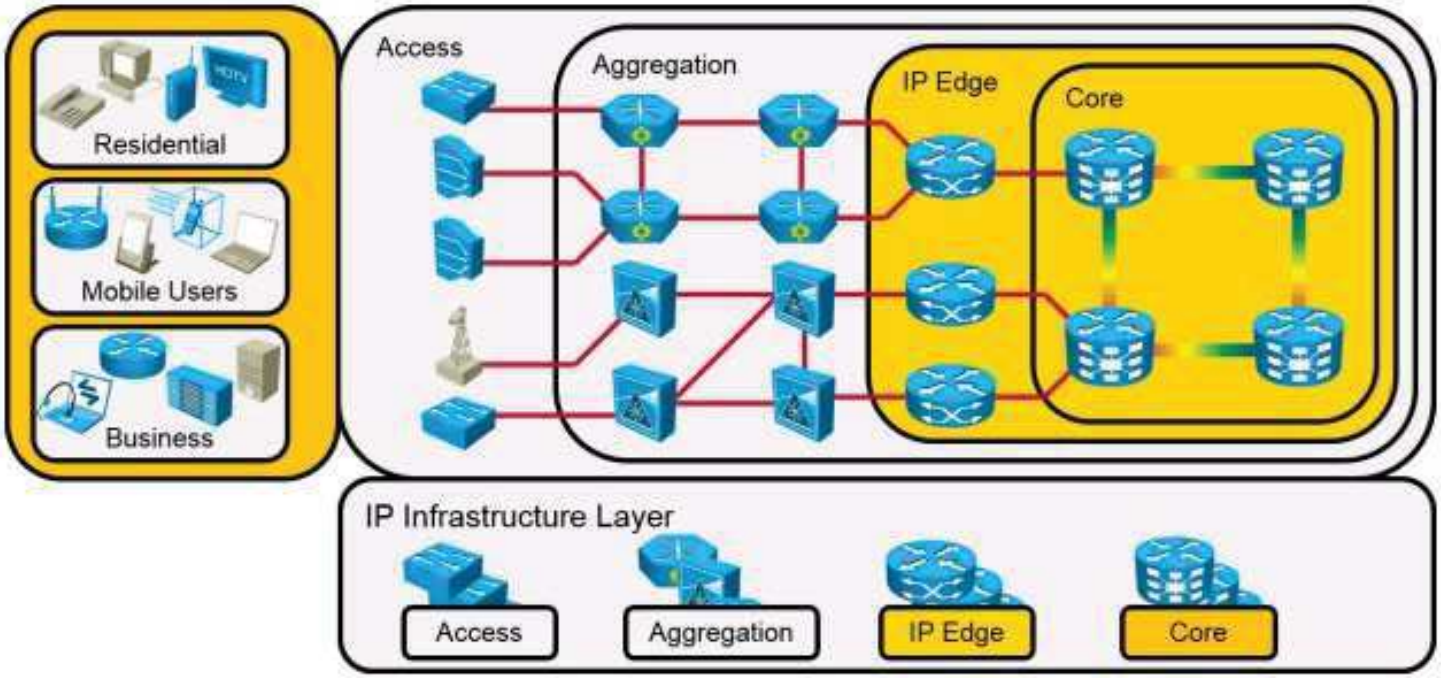

Overview of Routing Protocols

Requirements can be summarized as follows:

- Provide connectivity to the Internet for end customers and subordinate ISPs
- Optionally, provide transit connectivity between service providers (that are Tier 1 ISPs)

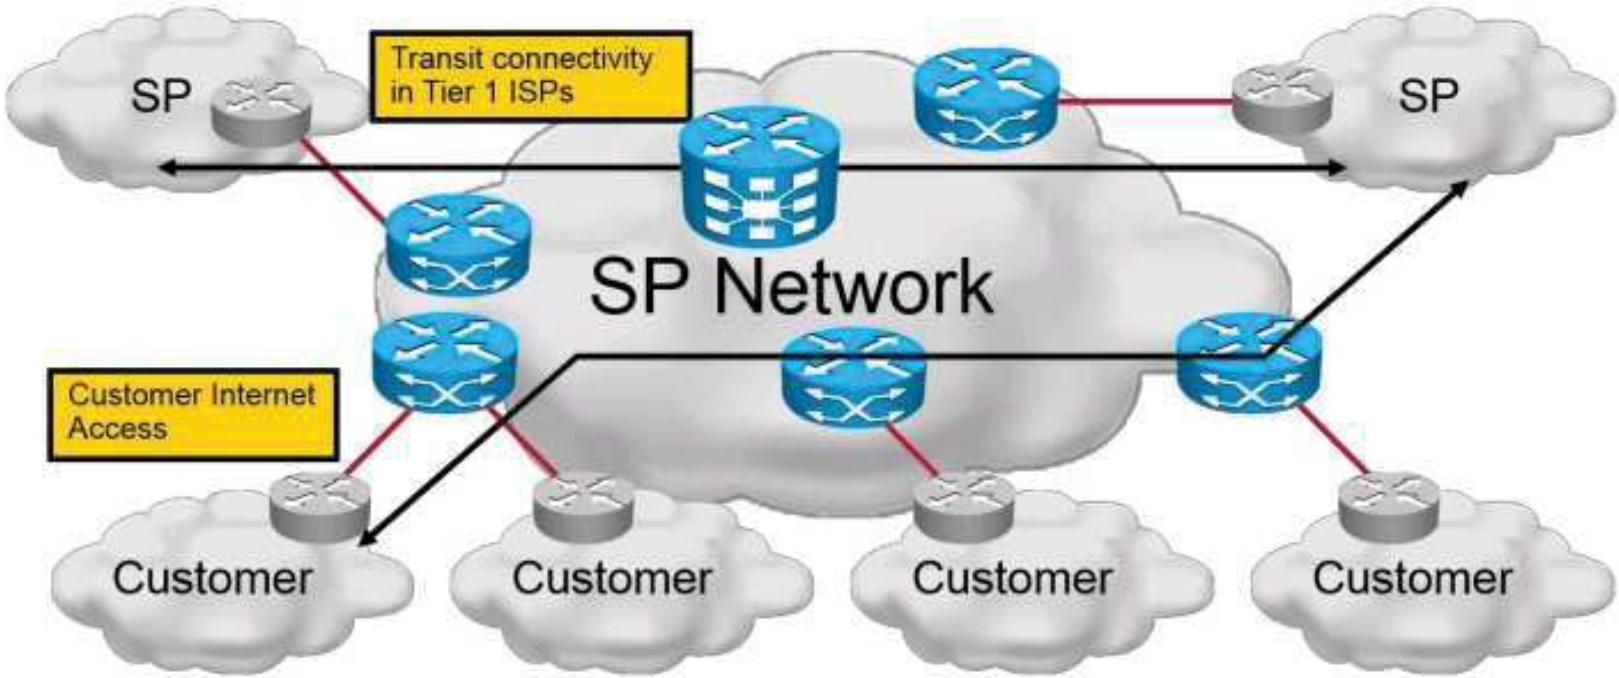

Overview of Routing Protocols (Cont.)

IGP and BGP routing:

- **IGP:** exchange local routing information
- **BGP:** exchange external routing information

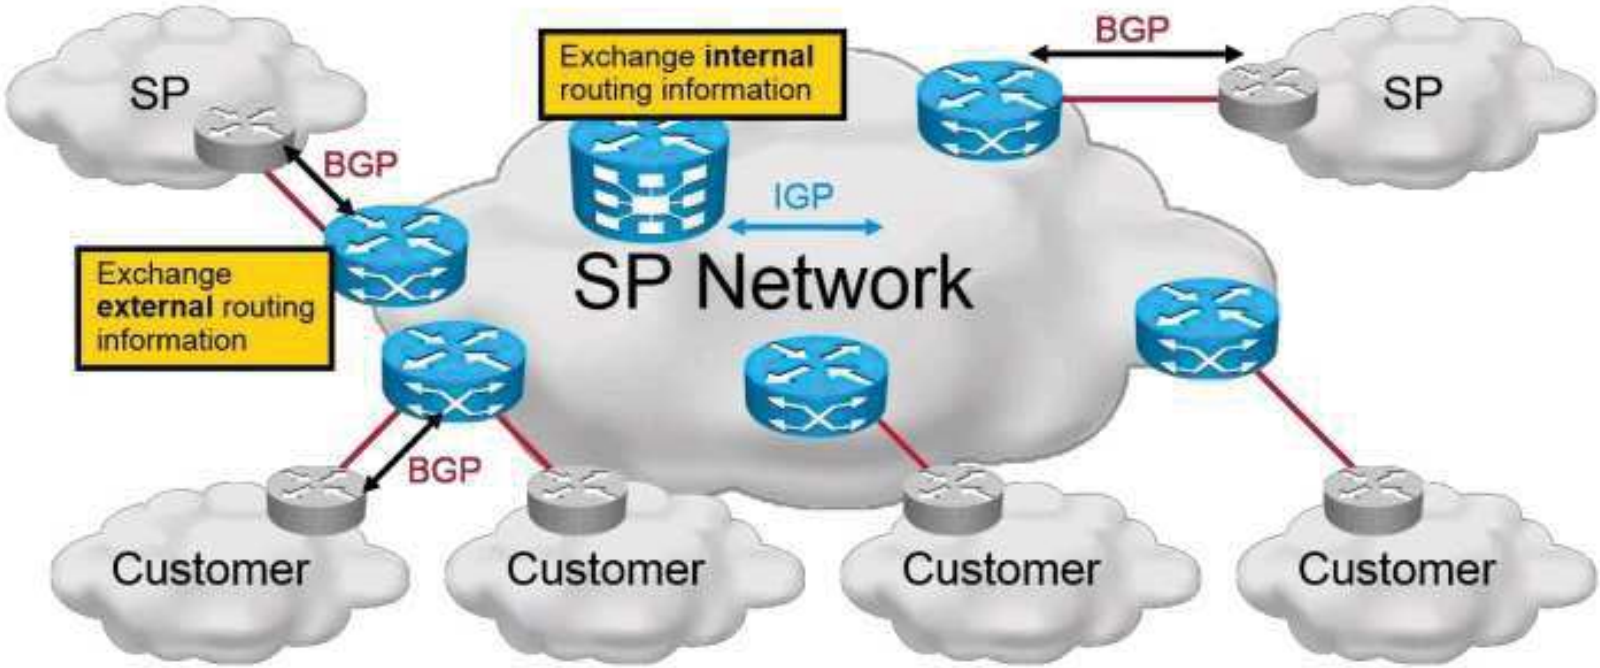

Overview of Routing Protocols (Cont.)

Routing tasks:

- IGP provides reachability for the following:
  - BGP next-hop addresses (typically directly connected edge subnets)
  - BGP neighbors
- BGP provides reachability to remote destinations through next-hop addresses:
  - External BGP sessions with customers and other ISPs
  - An internal BGP session within an AS (administrative domain).

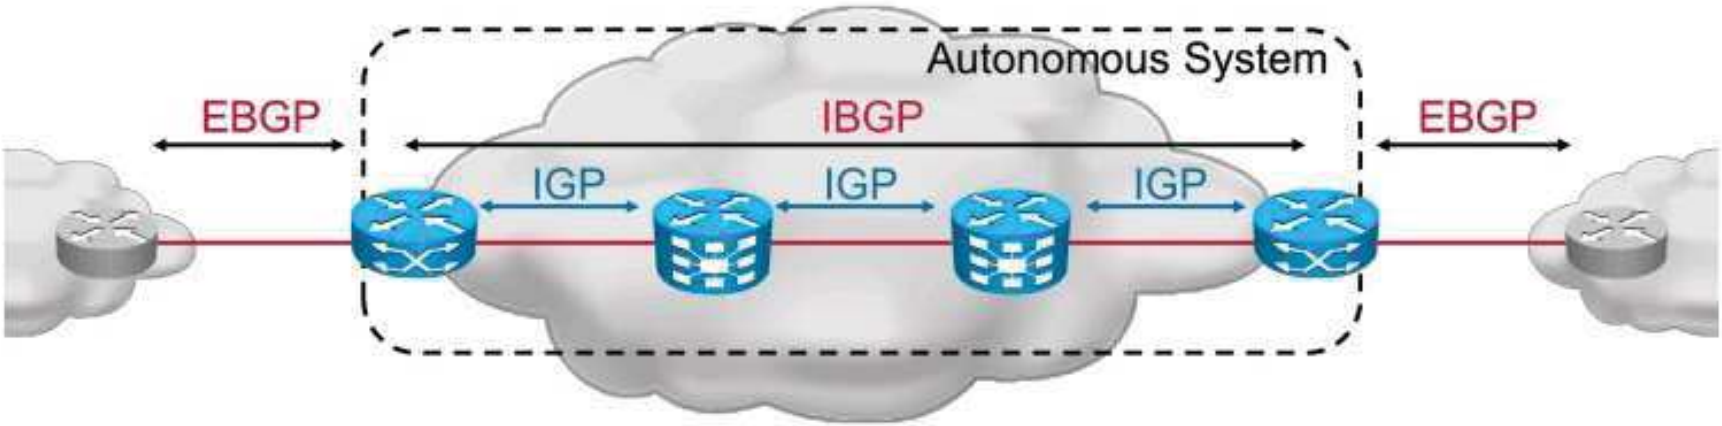

Interior Gateway Protocols

There are three scalable routing protocols for ISP backbones that satisfy the main service provider scalability and performance requirements for an IGP:

- OSPF
- IS-IS
- EIGRP

OSPF and IS-IS are the recommended choices:

- Standard protocols
- Support additional features required in MPLS-enabled networks

Routing Example

Part 1: BGP

1. R1 receives an external BGP update: 209.165.201.0/24; next hop is 192.168.200.2.
2. R4 receives an internal BGP update:
   - By default, the next-hop address does not change.
   - Optionally, BGP on R1 can be configured to change the next-hop address to its own address (typically a loopback address).
3. R4 forwards the update and changes the next-hop address to 192.168.11.1.

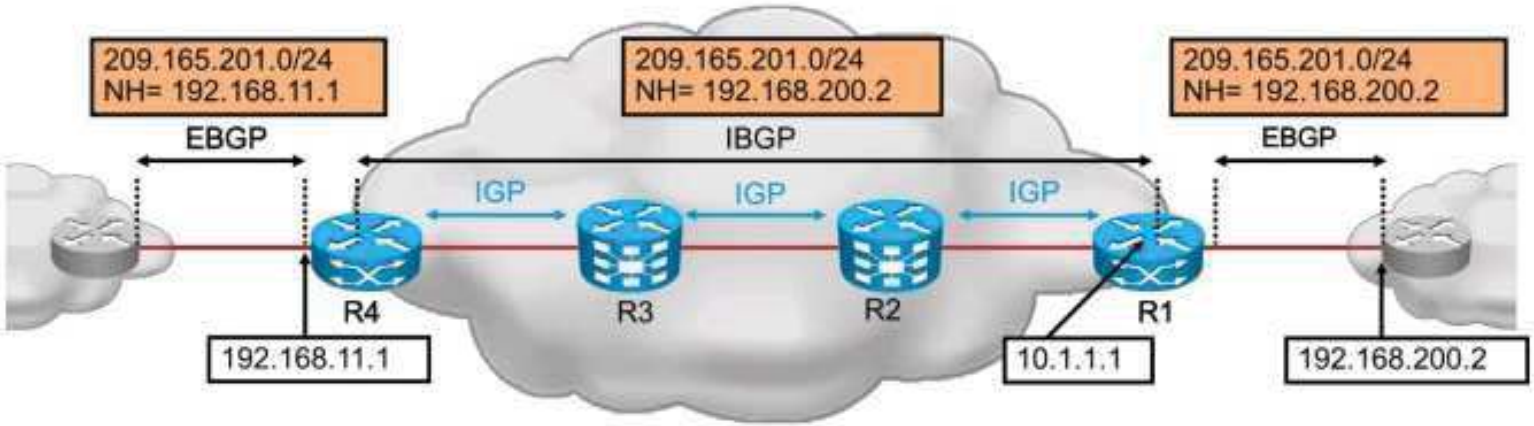

Routing Example (Cont.)

Part 2: IGP

- R1 propagates the BGP next-hop address to all routers in the domain:
  - Edge subnet (192.168.200.0/30) is for the reachability of external BGP next-hop addresses.
  - Loopback address (10.1.1.1/32) is for the reachability of internal BGP neighbors.
- R2 and R3 forward the information:
  - Unchanged (required if the network also uses MPLS-based services such as MPLS VPNs and Cisco MPLS TE).
  - Optionally, summarization can be used within IGP for optimization.

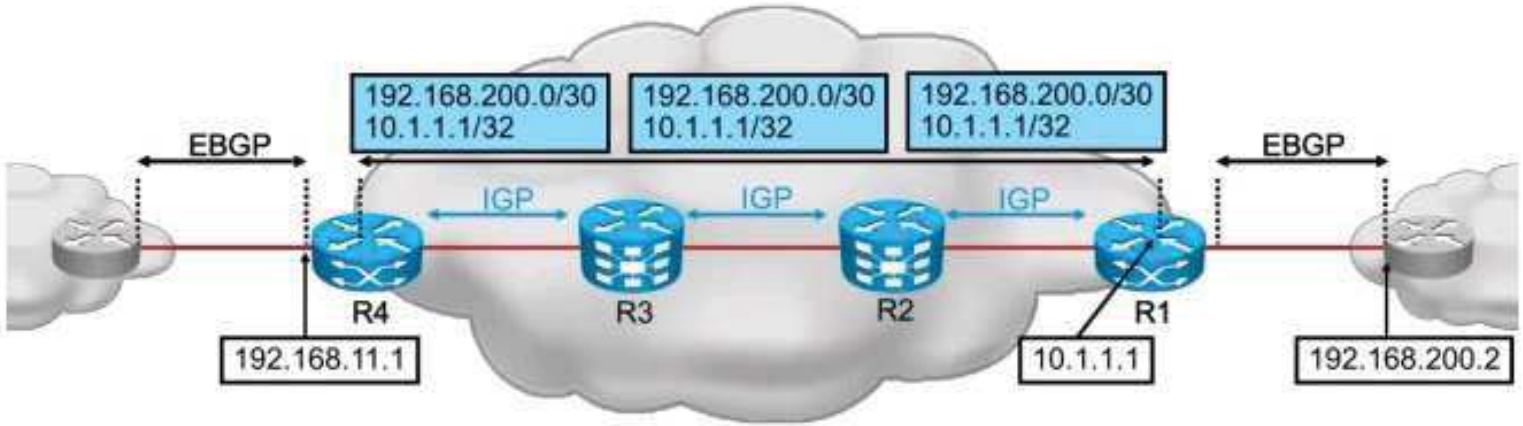

Routing Example (Cont.)

Part 3: Routing Table

- End-to-end connectivity is provided through recursive routing table lookups (optimized by Cisco Express Forwarding):
  - BGP for end prefixes
  - IGP for BGP next-hop reachability

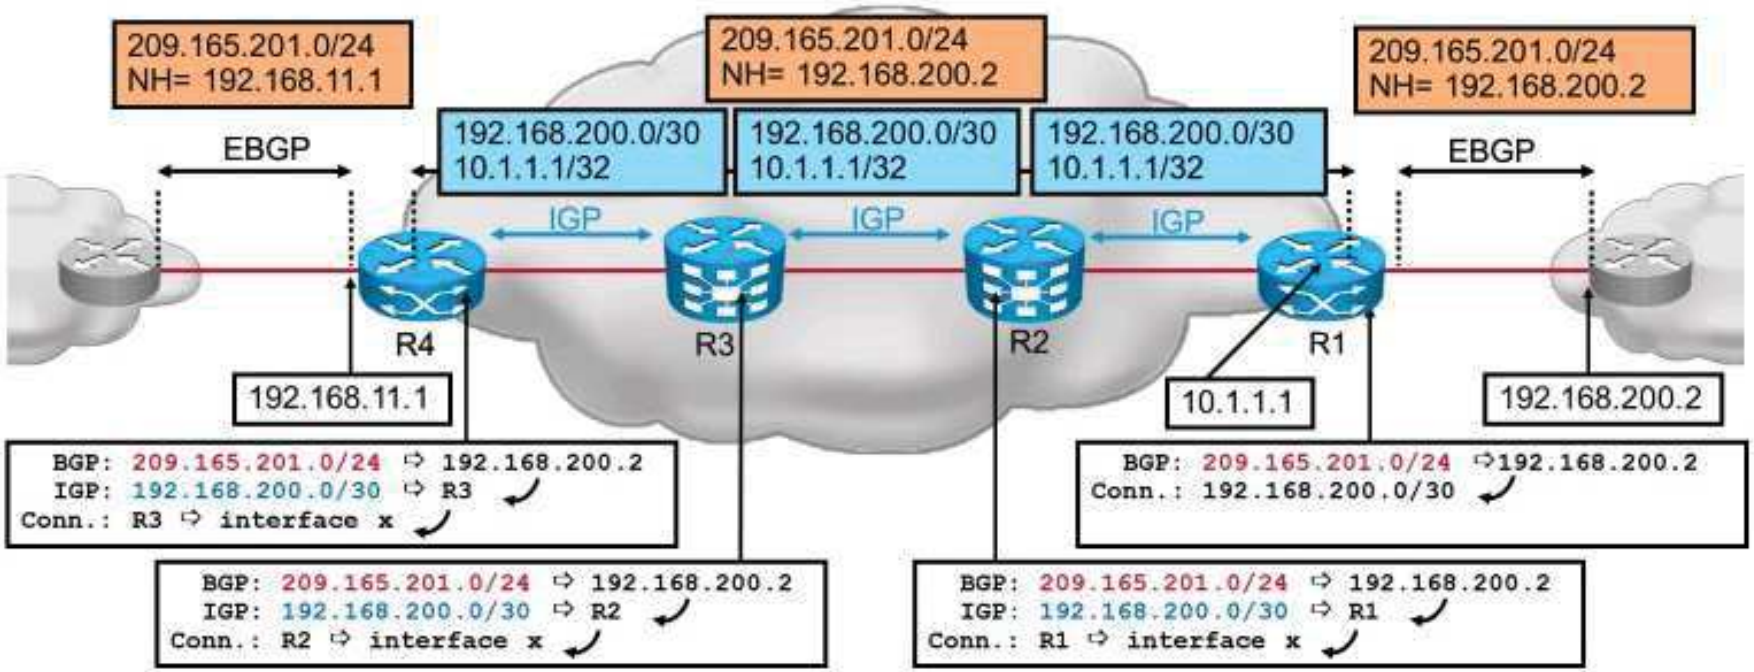

Overview of OSPF

OSPF overview:

- Link-state protocol: routers aware of network topology
- Hierarchical: dual-layer architecture:
  - Backbone area: Area 0
  - Nonbackbone areas interconnected through Area 0

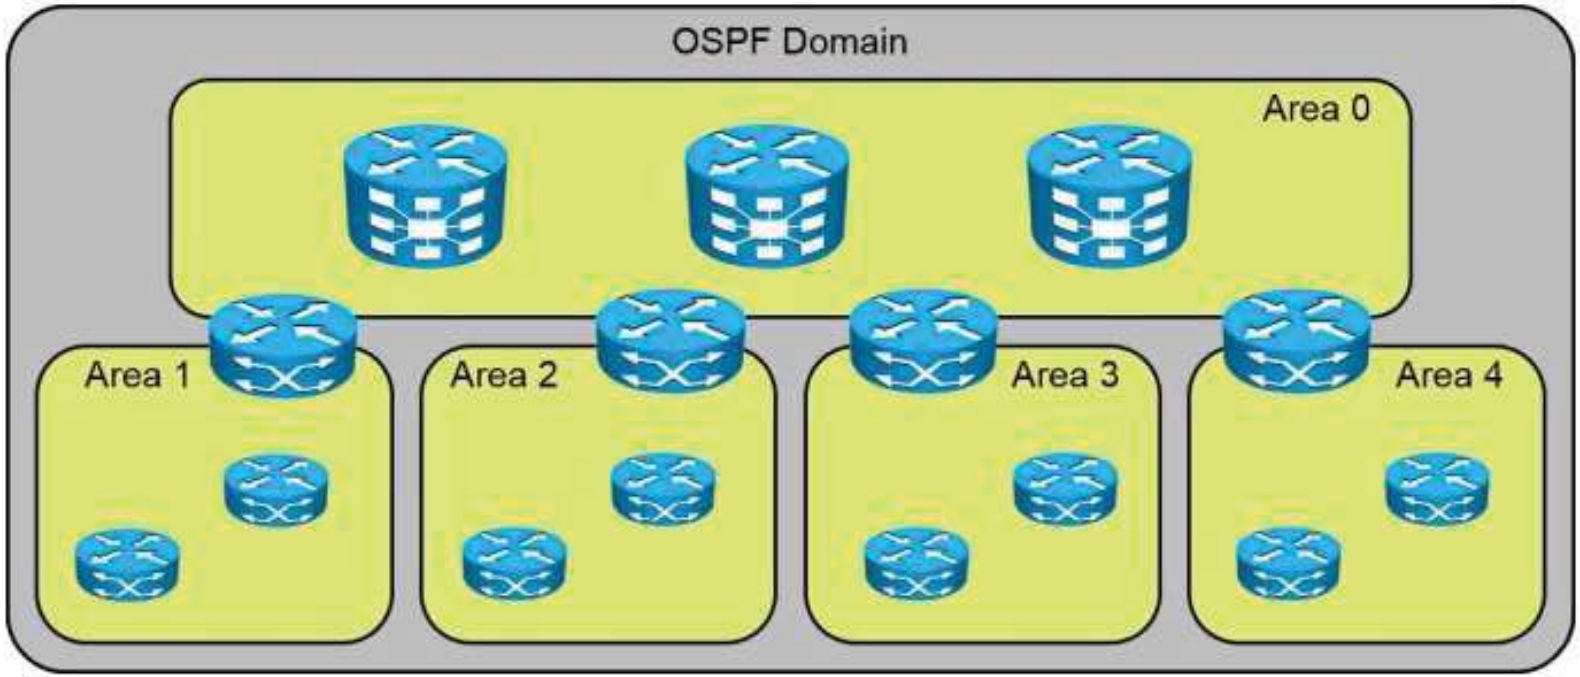

Overview of OSPF (Cont.)

OSPF characteristics:

- Creates a neighbor relationship by exchanging hello packets
- Propagates LSAs rather than routing table updates
  - Link: Router interface
  - State: Description of an interface and its relationship to neighboring routers
- Floods LSAs to all OSPF routers in the area, not just directly connected routers
- Pieces together all of the LSAs that are generated by the OSPF routers to create the OSPF link-state database
- Uses the SPF algorithm to calculate the shortest path to each destination and places it in the routing table

Link-State Data Structures

- Neighbor table:
  - Also known as the adjacency database
  - Contains a list of recognized neighbors
- Topology table:
  - Typically referred to as LSDB
  - Contains all routers and their attached links in the area or network
  - Identical LSDB for all routers within an area
- Routing table:
  - Commonly named a forwarding database
  - Contains a list of best paths to destinations

OSPF Area Terminology and Router Types

OSPF router types:

- ABR
- Internal router
- ASBR
- Backbone router

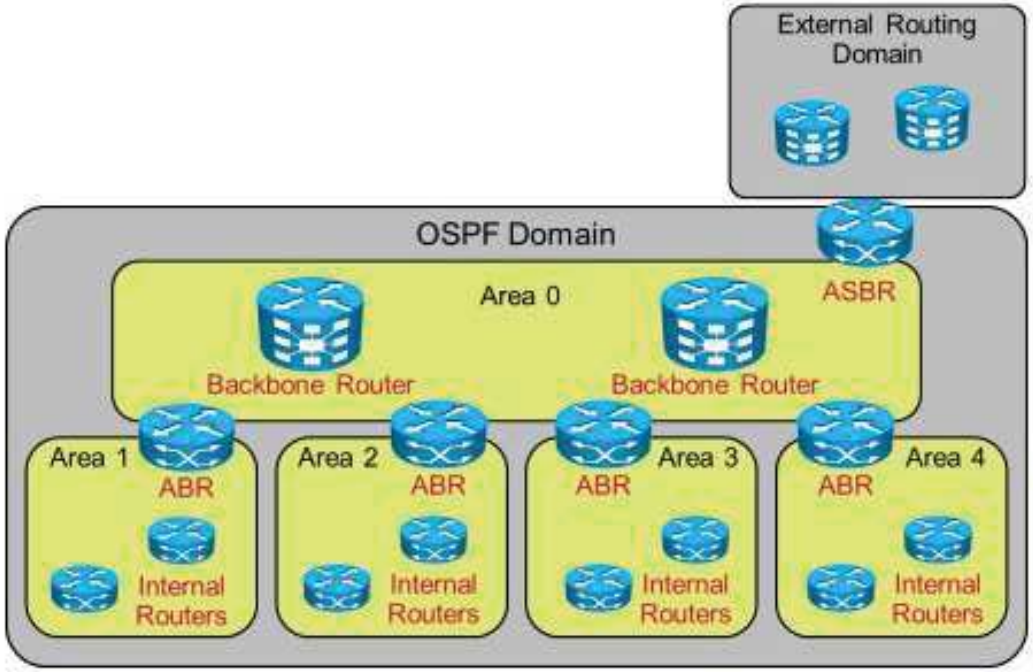

OSPF Areas

OSPF area types follow:

- Backbone area—Area 0
- Regular nonbackbone area
- Stubby area or totally stubby area
- NSSA or totally NSSA

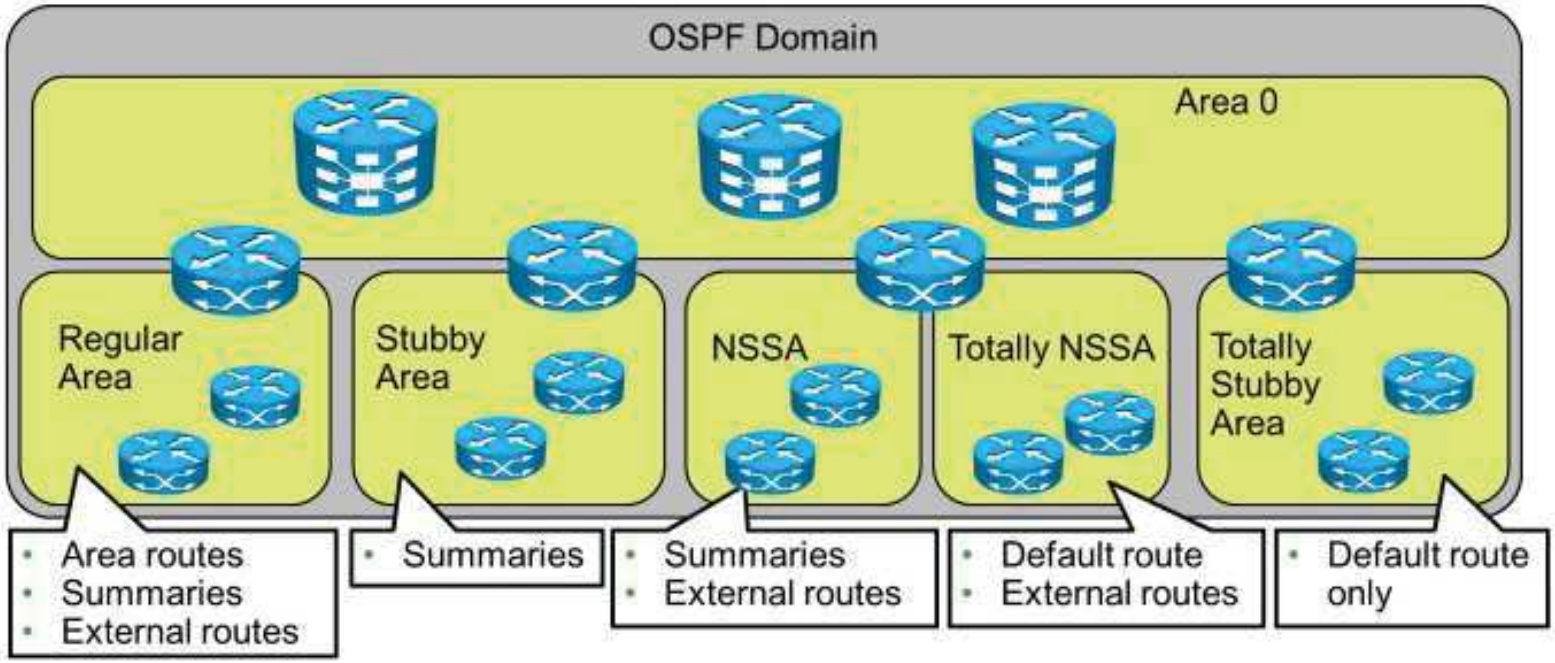

OSPF Metric

- Each link is assigned a cost:
  - The default cost is calculated from interface bandwidth.
  - The default reference bandwidth is 100 Mbps.
  - The modify reference bandwidth is in 1-Gbps networks.
  - The cost can be statically configured for an interface.
- Ensure consistent configuration of costs:
  - The same cost is on both sides of a link when manually configuring the cost.
  - The same reference bandwidth is on all routers in an OSPF domain.

$$\text{Cost} = \frac{\text{Reference Bandwidth}}{\text{Interface Bandwidth}}$$

Typical OSPF Designs

- Single-area design:
  - All routers in Area 0
  - Simple routing design
  - Mostly point-to-point adjacencies
  - Optimal routing decisions
  - Scalability limited to a few hundred routers in the network
- Multiarea design:
  - Regular areas or NSSA are typically used
  - Scales to thousands of routers in the network
  - Mostly point-to-point adjacencies
  - More complex routing design
  - May result in suboptimal routing (for example, dual-attached areas)
  - Less practical in MPLS-enabled networks

Overview of IS-IS

The characteristics of IS-IS in service provider environments follow:

- IS-IS is a stable protocol.
- IS-IS is originally deployed by ISPs because of U.S. government-mandated Internet support of OSI and IP.
- IS = router.
- IS-IS was originally designed as the IGP for the CLNS, which is part of the OSI protocol suite.
- The OSI protocol suite Layer 3 protocol is the CLNP.
- IS-IS uses CLNS addresses to identify routers and build the LSDB.

IS-IS Hierarchical Design

Hierarchical: dual-layer architecture:

- Level 1 is used within local areas.
- Level 2 interconnects areas.

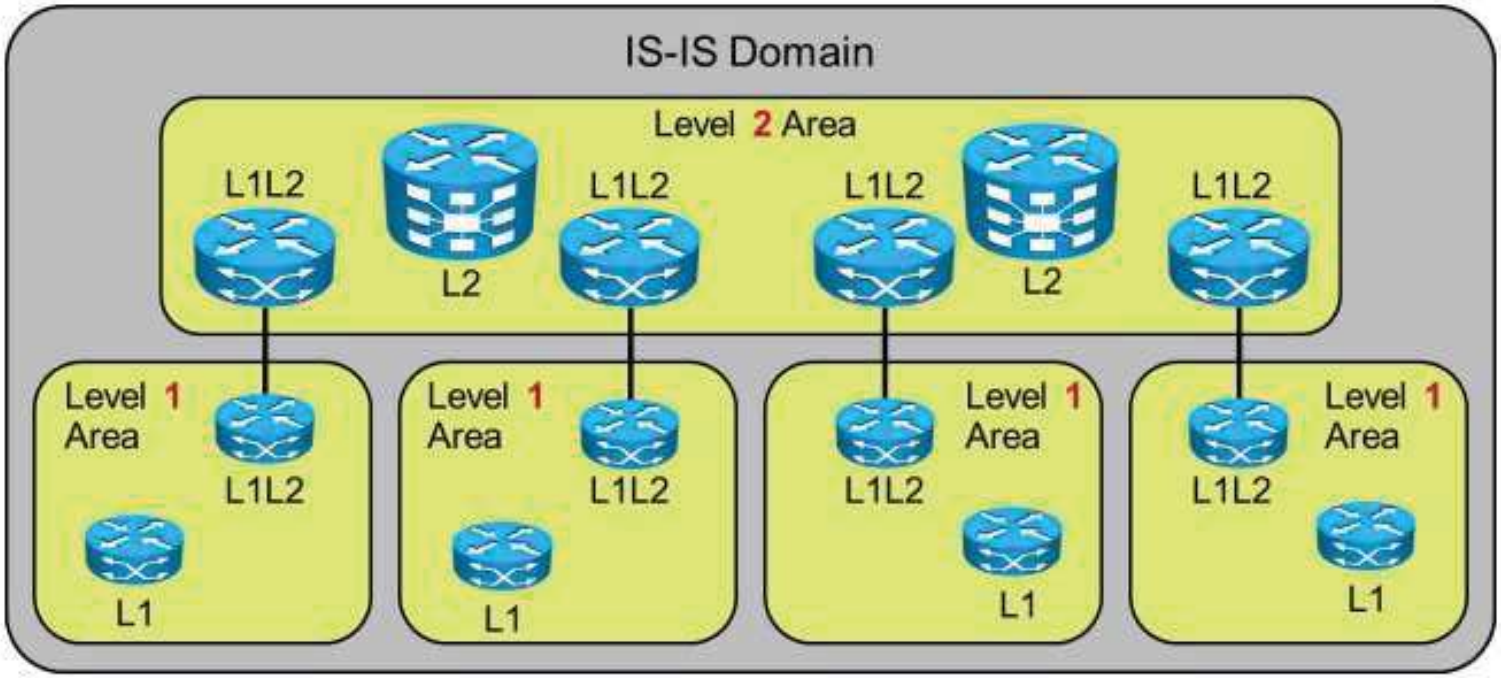

IS-IS Characteristics

An overview of IS-IS characteristics follows:

- It is a link-state routing protocol (routers aware of network topology).
- It supports VLSMs.
- It uses the Dijkstra SPF algorithm, which has fast convergence.
- It uses hellos to establish adjacencies and LSPs to exchange link-state information.
- It has efficient use of bandwidth, memory, and processor.
- It supports two routing levels:
  - **Level 1:** This builds a common topology of system IDs in the local area and routes within an area using the lowest-cost path.
  - **Level 2:** This exchanges prefix information (area addresses) between areas and routes traffic to an area using the lowest-cost path.

IS-IS Characteristics (Cont.)

More IS-IS characteristics follow:

- Each router has topology information for its area.
- IS-IS is part of OSI and was originally used with CLNS only.
- IS-IS still uses CLNS to maintain adjacencies and build an SPF tree.
- Integrated IS-IS can also carry IP routing information in its updates.
- The wide-style metric should be used for large, high-speed service provider networks (24-bit link metric and 32-bit path metric).
- The link cost defaults to 10.
- Each router is identified using a unique NSAP address.

IS-IS Router and Link Types

- Router types follow:
  - Level 1 routers only peer with other Level 1 routers.
  - Level 2 routers only peer with other Level 2 routers.
  - Levels 1 and 2 routers can peer with any router.
- Link types:
  - **Level 1:** Only for Level 1 adjacencies within the same area
  - **Level 2:** Only for Level 2 adjacencies
  - **Level 1 and 2:** For Level 1 adjacencies within the same area and Level 2 adjacencies

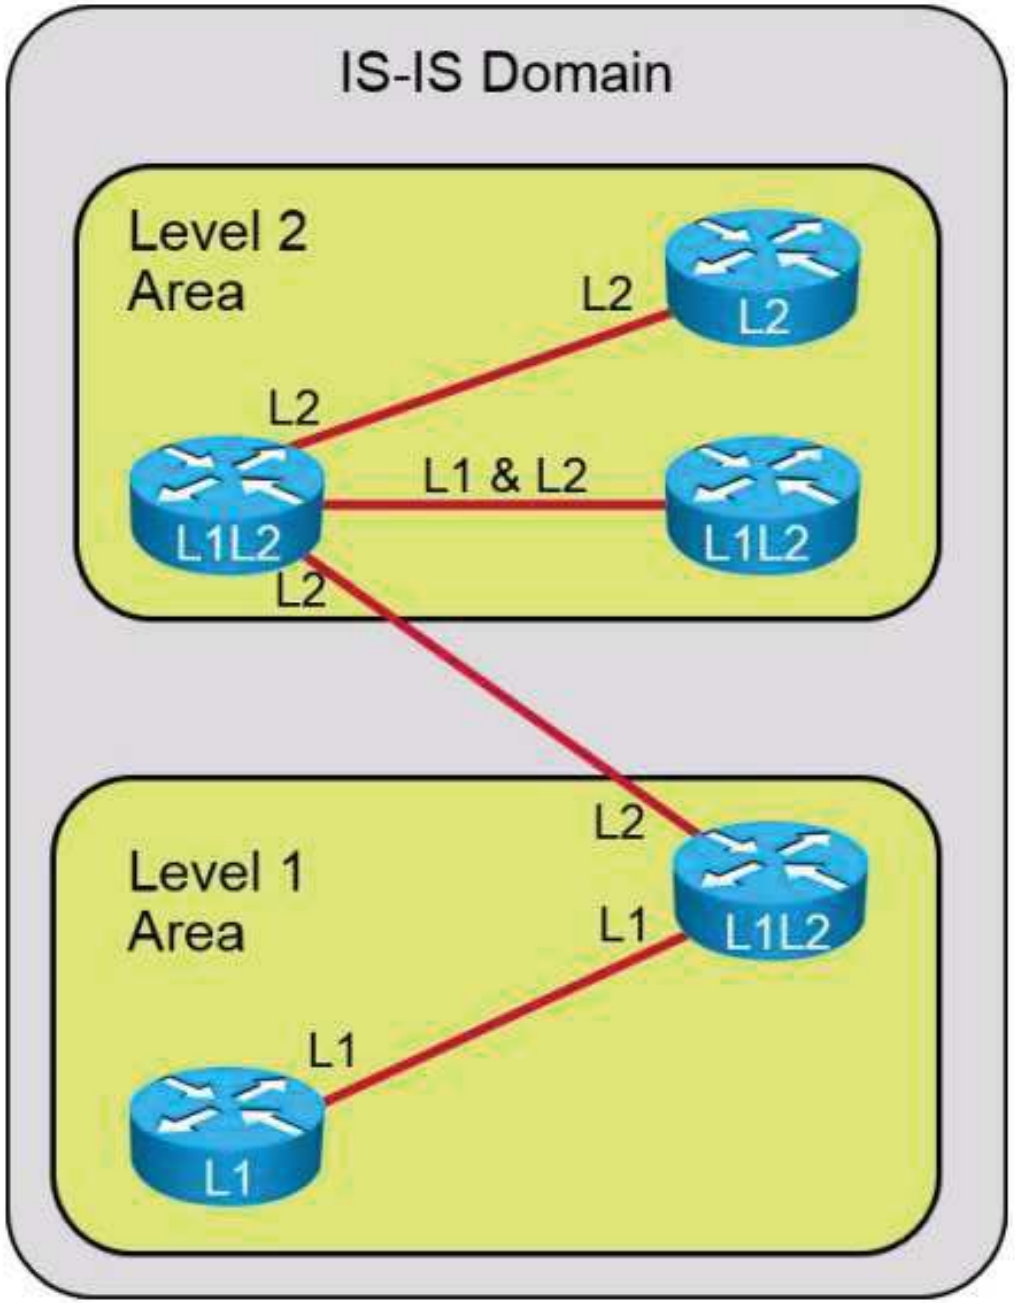

BGP Overview

An overview of BGP follows:

- BGP is designed for routing information exchange between different administrative domains (ASs).
- Each AS is identified using a unique AS number.
- BGP is designed with the following major characteristics:
  - Scalability
  - Stability
  - Security
  - Flexibility

BGP Architecture

There are two types of BGP sessions:

- EBGP sessions exchange routing information between different ASs.
- IBGP sessions exchange routing information between routers within the same AS.

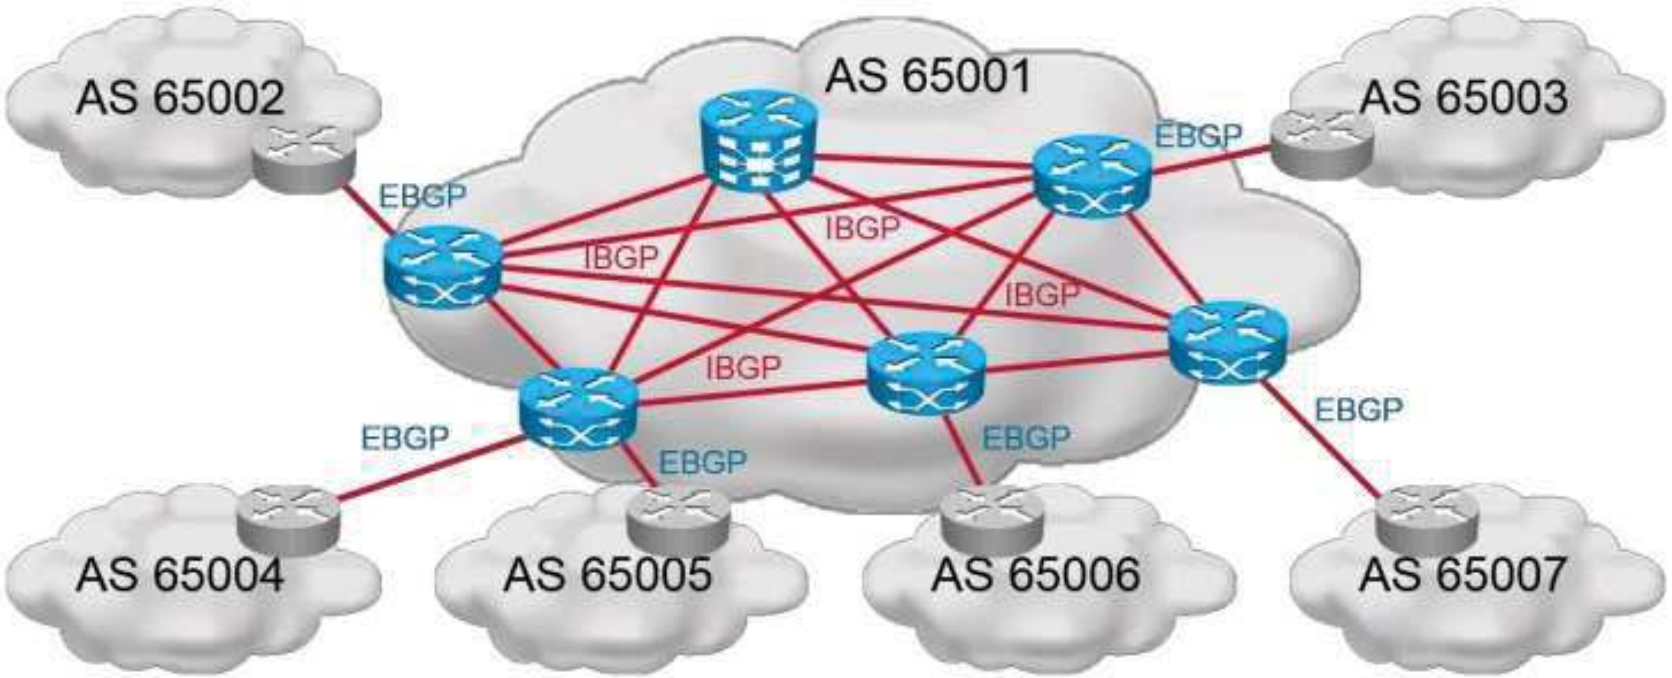

BGP Characteristics

BGP is a path-vector protocol with enhancements:

- Reliable updates
- Triggered updates only
- Rich metrics (called path attributes)
- Designed to scale to huge internetworks

Reliable updates:

- TCP that is used as a transport protocol
- No periodic updates
- Periodic keepalives to verify TCP connectivity
- Triggered updates that are batched and rate-limited
  - Every 5 seconds for the internal peer
  - Every 30 seconds for the external peer

BGP Characteristics (Cont.)

BGP was designed to perform well in these areas:

- Interdomain routing applications
- Huge internetworks with large routing tables
- Environments that require complex routing policies

Common BGP uses follow:

- Customers connected to more than one service provider
- Service provider networks (transit ASs)
- Service providers exchanging traffic at an exchange point (CIX, GIX, NAP, and so on)
- Network cores of large-enterprise customers

BGP AS Number

16-bit AS number:

- Notation: X (for example, “65001”)
- Public range from 1 to 64511 for use on the Internet
- Private range from 64512 to 65535 can be used in isolated environments
- Depleted

32-bit AS number:

- Notation: X.Y (for example, “65100.65200”)
- Carried in a new attribute
- Compatible with old systems:
  - AS 23456 that is used in the old AS path to represent ASs using the new AS number format
  - AS 0.X that is used to encode old AS numbers in the new AS path attribute

BGP Sessions

BGP sessions are established as follows:

- BGP uses TCP on port 179 to establish adjacencies.
- OPEN messages are used at session setup to negotiate fundamental session parameters and capabilities:
  - AS numbers must match configuration and determine session type (EBGP versus IBGP).
  - EBGP peers must be reachable through a directly connected link (default).
  - IBGPs are typically established between loopbacks. (IGP ensures reachability of loopback addresses.)
  - IP addresses must match the configuration.
  - Hold time (default is 180 seconds).

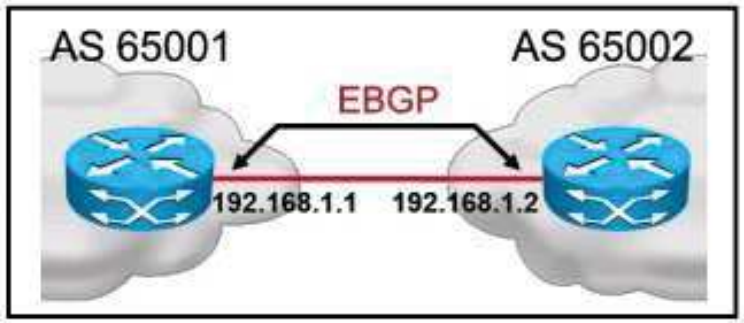

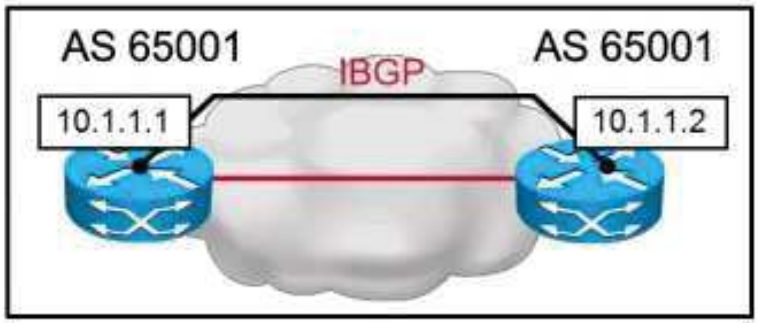

BGP Sessions (Cont.)

EBGP sessions:

- EBGP sessions can form any topology, which is subject to agreements between ASs.
- Received EBGP updates are sent to all other neighbors.
- By default, EBGP neighbors must be directly connected.

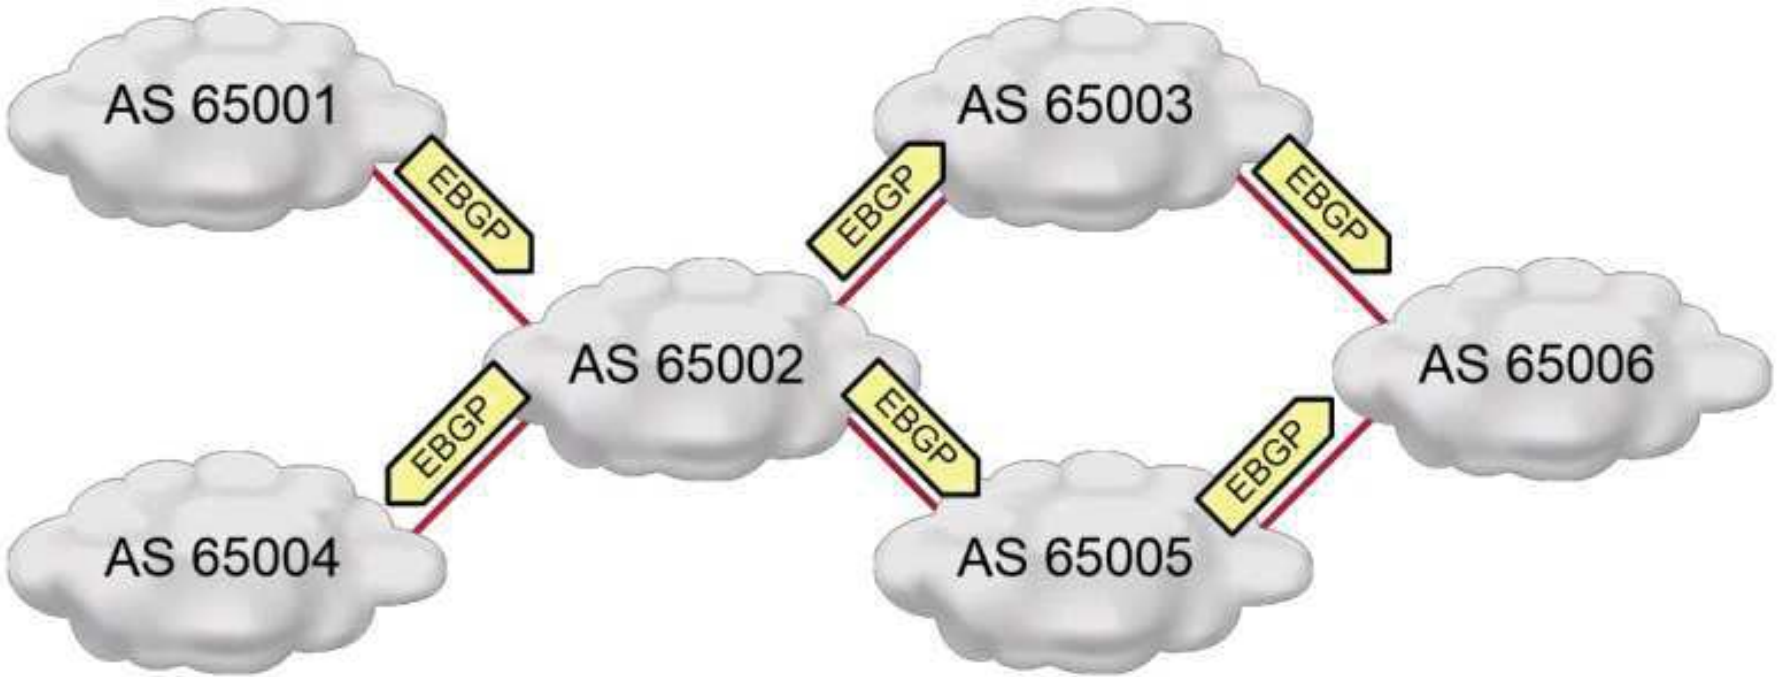

BGP Sessions (Cont.)

IBGP sessions:

- By default, IBGP sessions require a full mesh between all routers within an AS:
  - By default, IBGP updates that are received are not forwarded to other IBGP neighbors.
  - IBGP updates do not scale in large ASs.
- IBGP neighbors can be multiple hops away.

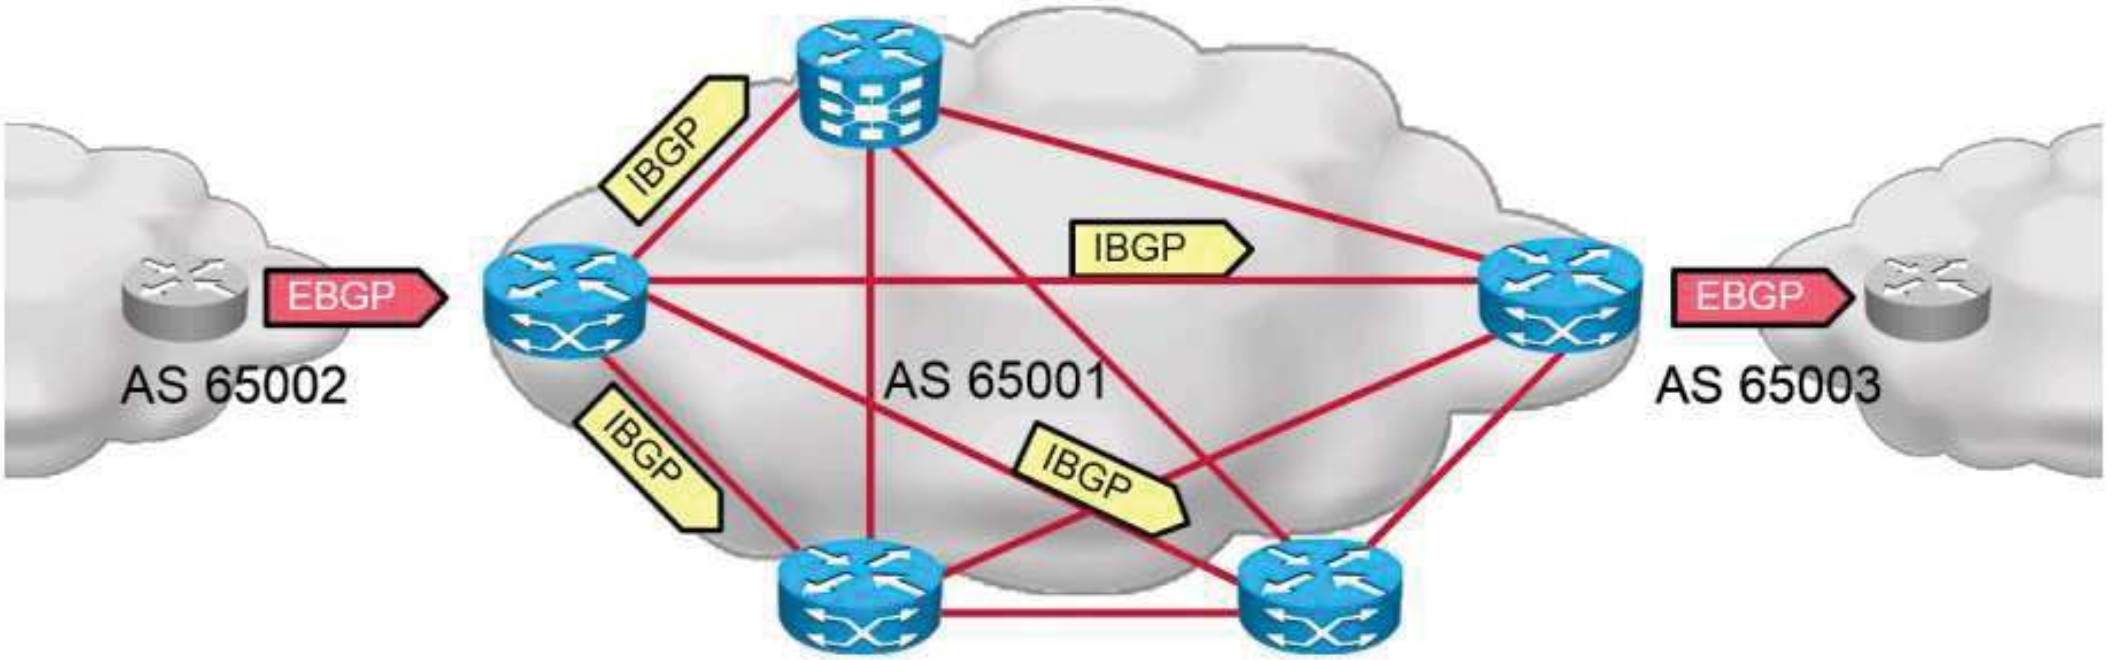

BGP in Customer Connections: Single-Homed Customers

- Single-homed customers, typically, do not require BGP:
  - A static route for the customer ISP-assigned address space is on the edge router.
  - A static default route is on the customer router.
- BGP can be used to detect link failures and trigger dial backup:
  - The ISP originates only the default route.
  - The customer originates address space.
  - Private AS numbers can be assigned to customers by the ISP.

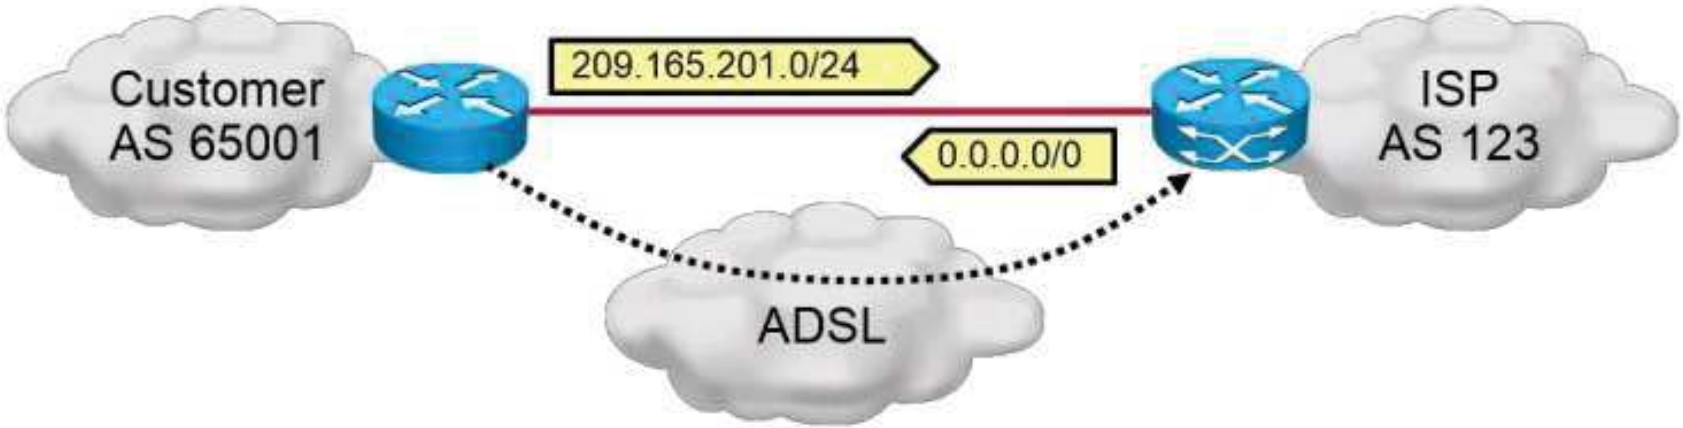

BGP in Customer Connections: Dual-Attached Customers

For dual-attached customers, BGP provides the following services:

- Mitigates link and device failures
- Two design options:
  - Primary and backup routing
  - Load balancing

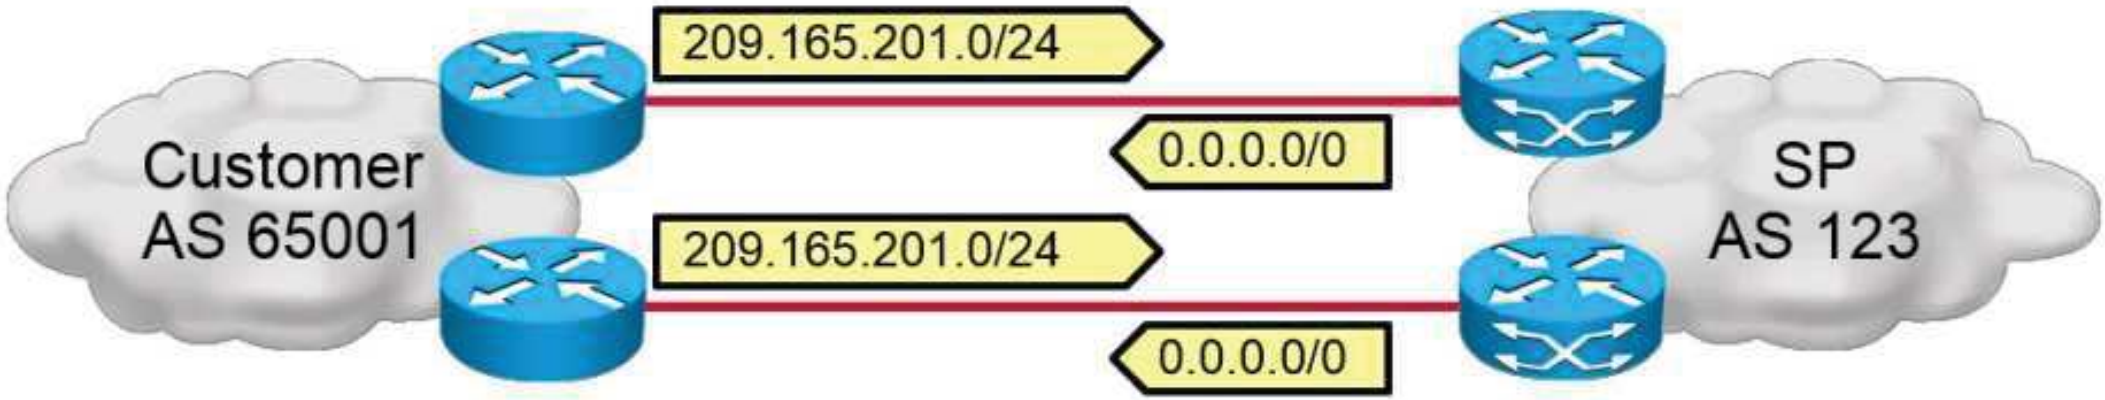

BGP in Customer Connections: Multihomed Customers

For multihomed customers, BGP provides the following services:

- Mitigates link, device, and path failures
- Should connect to independent service providers
- Two design options:
  - Primary and backup routing
  - Load balancing

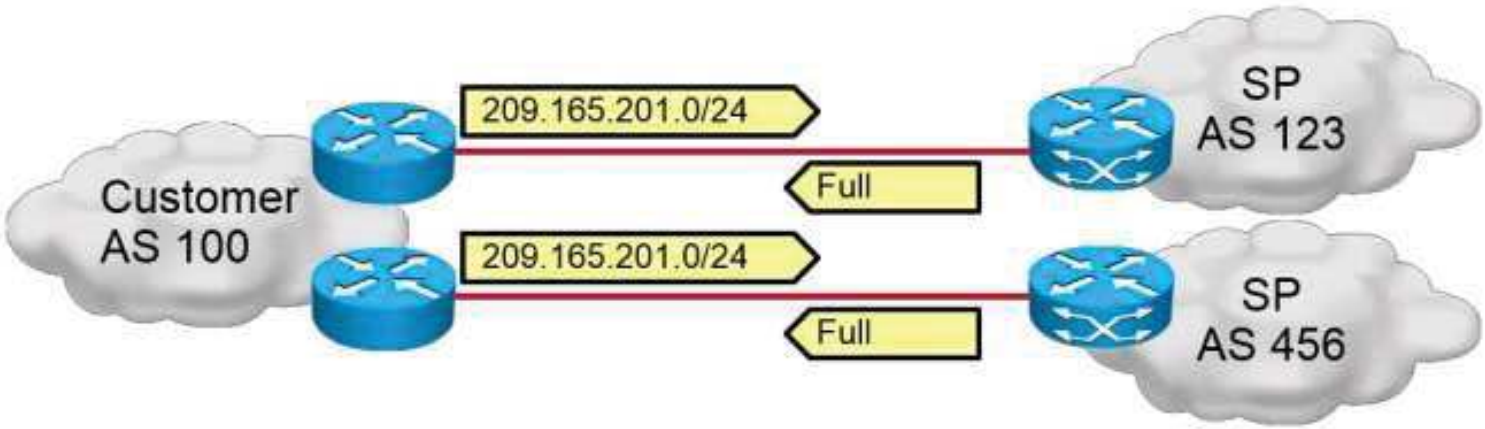

BGP in Customer Connections: Upstream ISP

In an upstream ISP customer environment, BGP provides the following services:

- Mitigates link, device, and path failures
- Should connect to independent upstream ISPs
- Two design options:
  - Primary and backup routing
  - Load balancing
- ISP receives the full Internet routing table
- ISP forwards the following:
  - Summaries for owned address space
  - Prefixes from BGP customers using independent address space

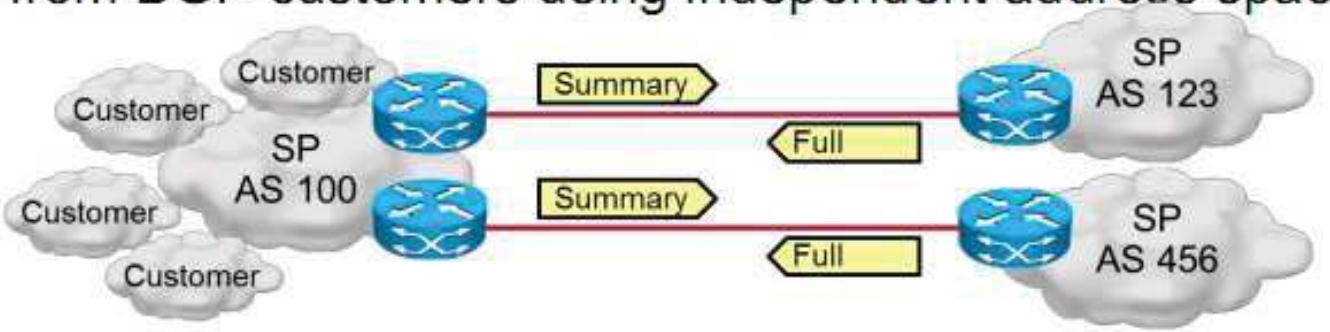

BGP in Customer Connections: Transit ISP

In a transit ISP customer environment, BGP provides the following services:

- Mitigates link, device, and path failures
- Routing policy depends on agreements with other ISPs
- Tier 1 ISP forwards the full Internet routing table

Summary

- Routing protocols are used on the IP infrastructure layer of the Cisco IP NGN.
- BGP is used to exchange external routing information, while IGP is used to exchange local routing information.
- Most service providers use OSPF or IS-IS as IGP.
- IGP propagates BGP next-hop addresses and loopback IP addresses that are used for BGP sessions.
- OPSF uses three data structures: neighbor table, topology table, and routing table.
- The ABR router separates the backbone area from other areas.
- OSPF supports six different types of areas. From a design and implementation perspective, it is preferred to implement OSPF using single-area design.
- OSPF uses link cost as a metric. Link cost is calculated from interface bandwidth.

Summary (Cont.)

- IS-IS is a stable, link-state routing protocol that is often used in service provider environments.
- IS-IS uses hierachical design with Level 1 and Level 2 areas.
- IS-IS is a link-state routing protocol and is similar to OSPF, but it uses CLNS addresses to identify routers and build the LSDB.
- BGP was designed to exchange routing information between autonomous systems.
- BGP uses external and internal BGP sessions to exchange routing information.
- BGP session parameters and capabilities are negotiated during session setup by exchanging OPEN messages.
- BGP is used between service providers and multihomed customers.

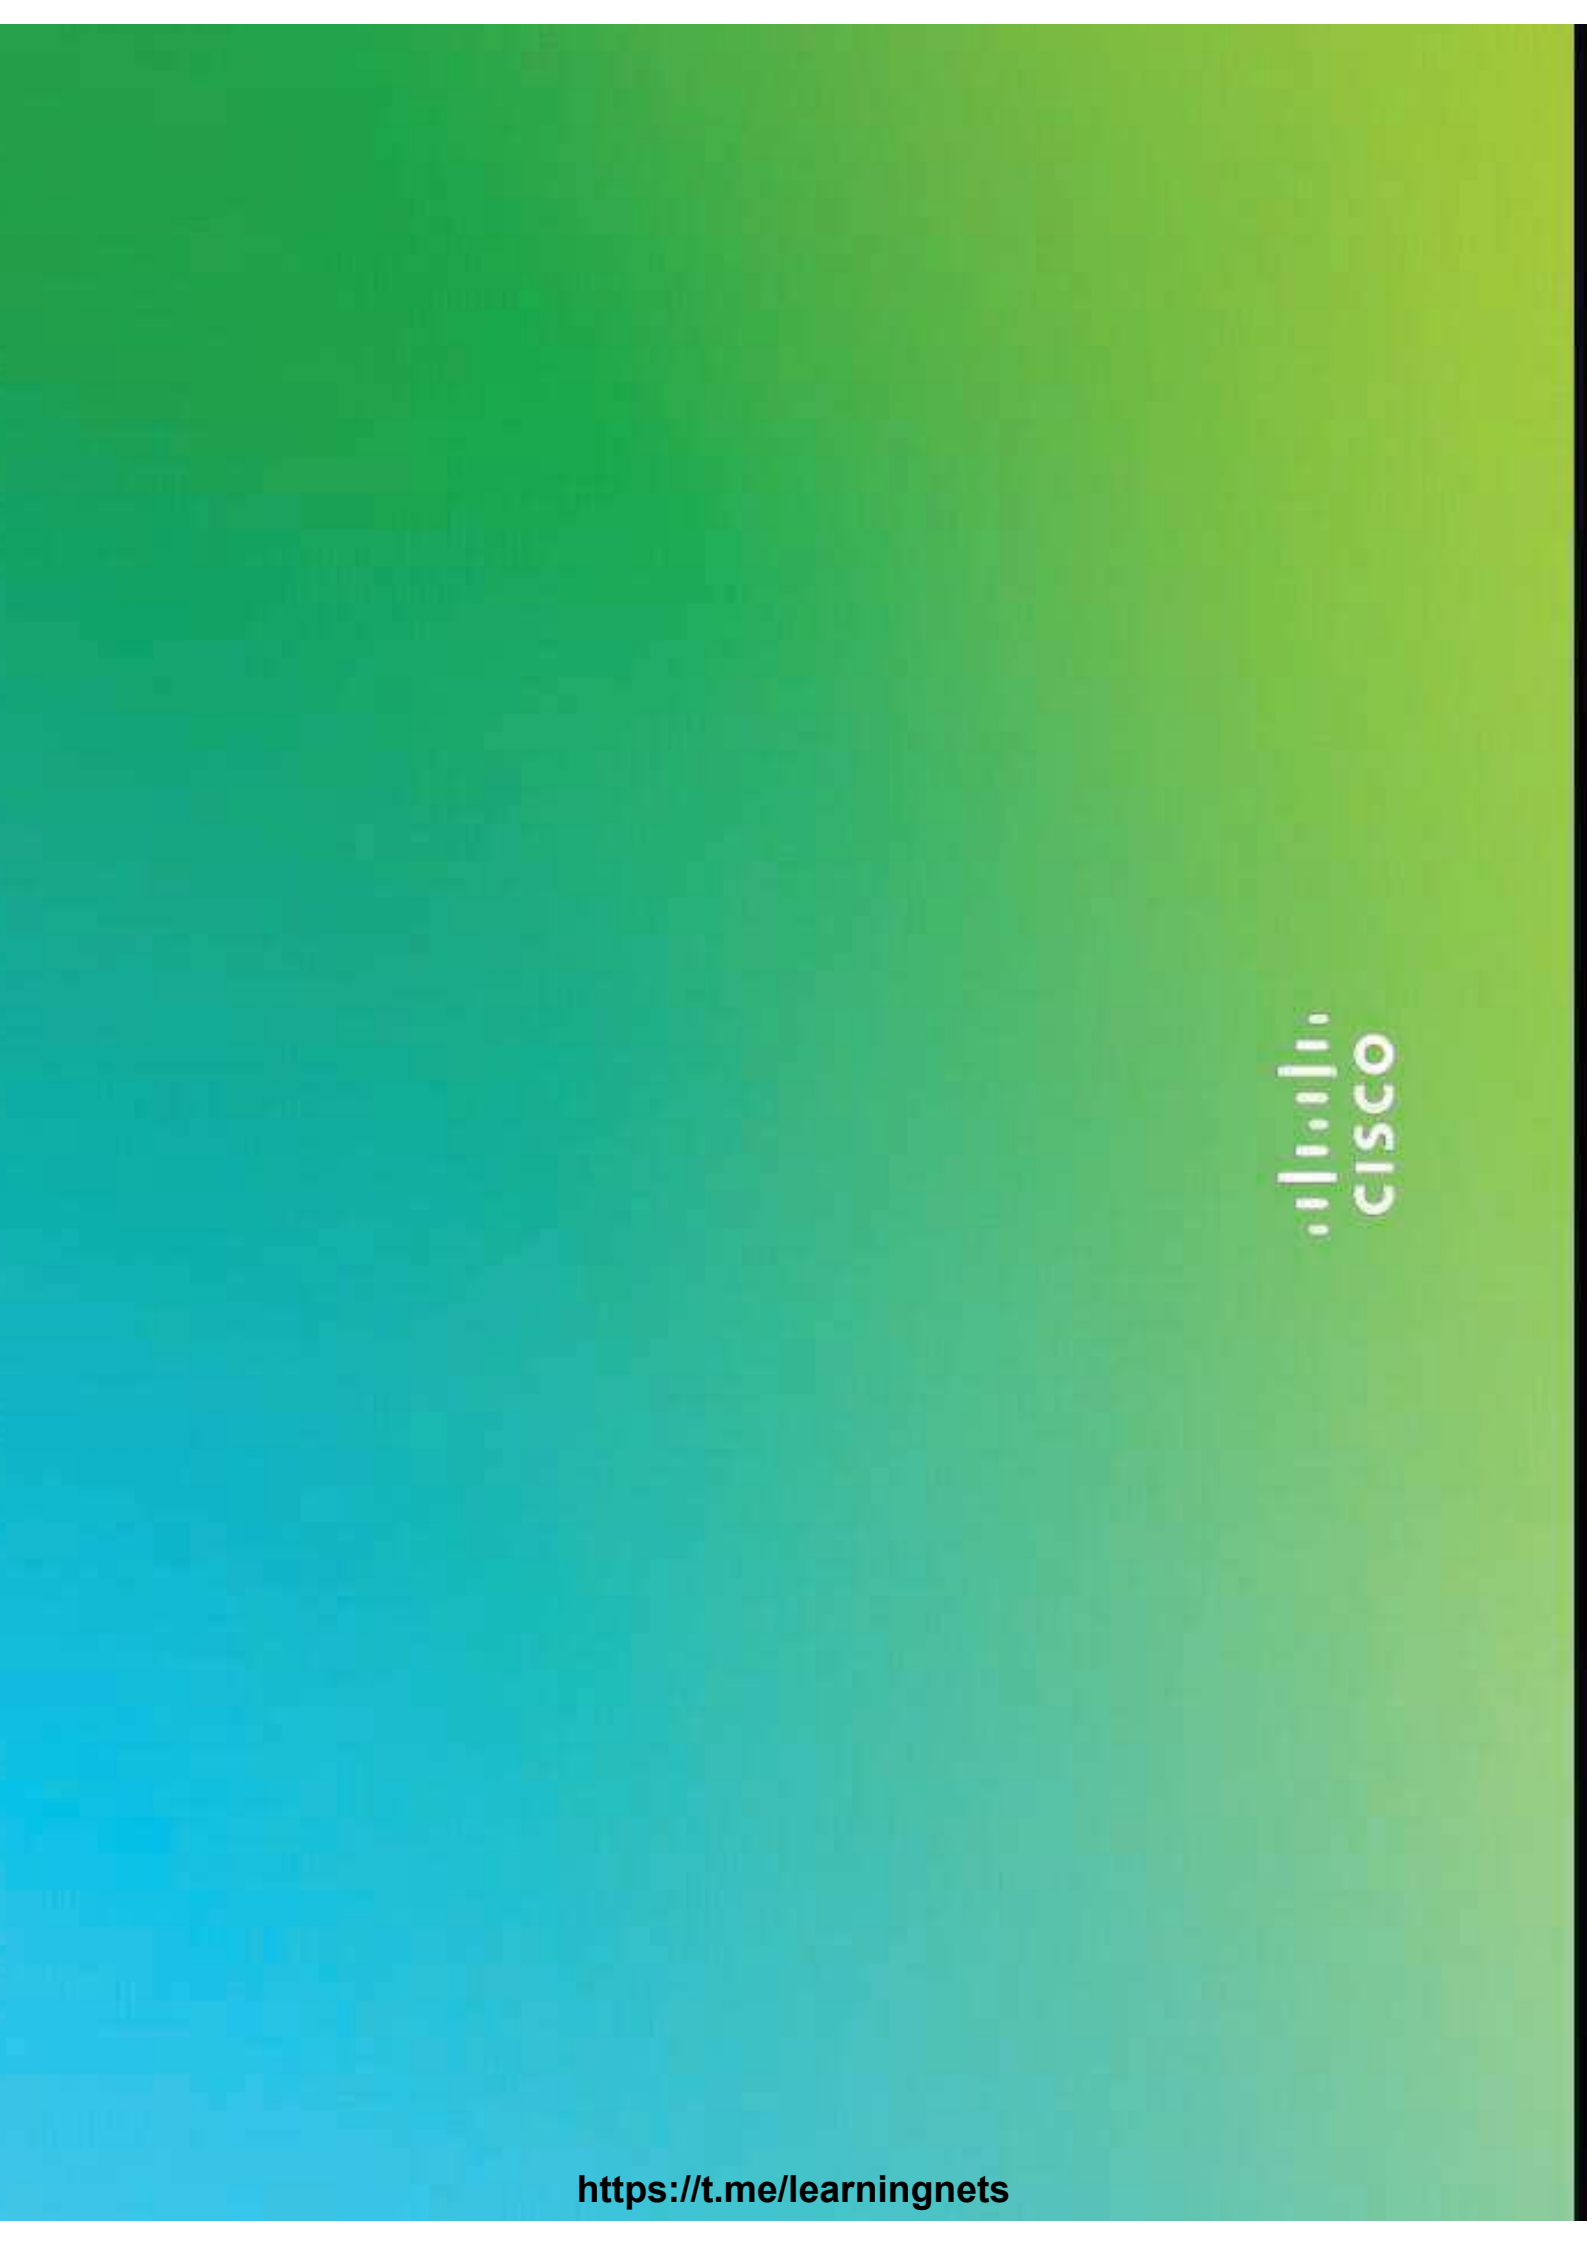

Module Summary

- ISPs use OSPF or IS-IS as IGPs and BGP to provide IP connectivity within the Internet. Routing protocols in the ISP environment should follow these characteristics: scalability, performance, high availability, and security.

Implement OSPF in the Service Provider Network

Deploying Cisco Service Provider Network Routing (SPROUTE) v1.2

Introducing OSPF Routing

Implement OSPF in the Service Provider Network

OSPF in the Cisco IP NGN Architecture

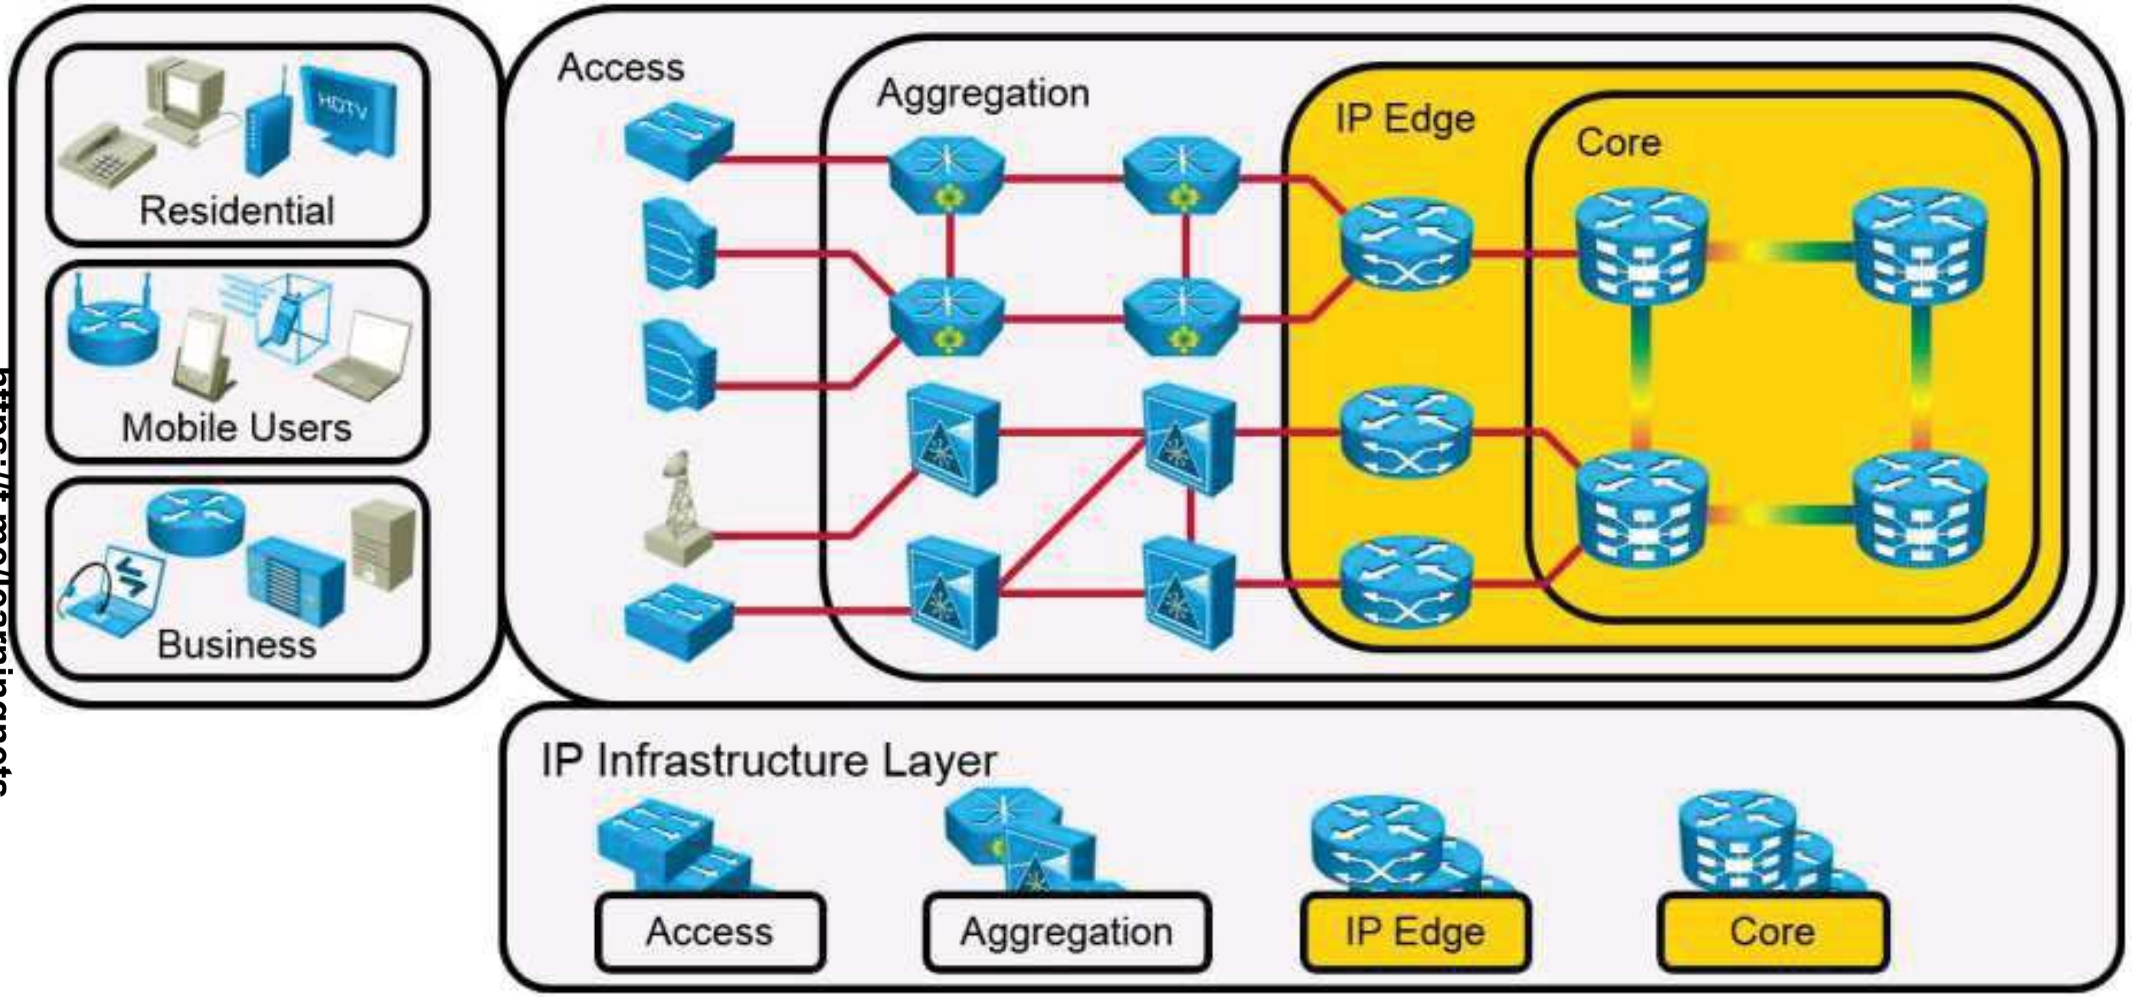

- The OSPF routing protocol that is used in service provider environments focuses on the IP infrastructure layer of the Cisco IP NGN and service provider IP edge and core devices.

OSPF and OSPFv3 Key Characteristics

The key characteristics of OSPF and OSPFv3 follow:

- OSPFv3 is an implementation of the OSPF routing protocol for IPv6.
- OSPFv2 (for IPv4 networks) and OSPFv3 run independently on a network device.
- OSPFv3 has the same key capabilities as OSPFv2:
  - Multiarea network design with ABRs that segment the network
  - SPF algorithm for optimum path calculation
  - Special area types and sophisticated handling of external routes
  - Summarization on area borders simplifies network designs (stub areas)

OSPF Route Entry Creation

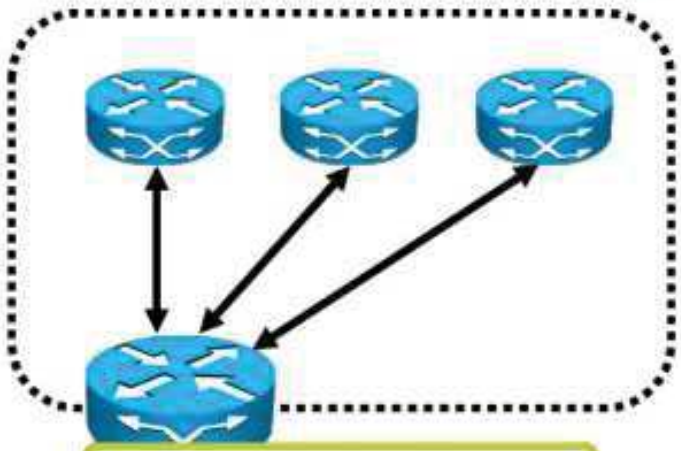

The shortest path to the destination is the lowest total sum of metrics.

Link-State Database → SPF Tree → IP Routing Table

Received LSA updates

Run Dijkstra's algorithm

Best paths are placed in the IP routing table

```
 10.0.0.0/32 is subnetted, 13 subnets
O       10.3.1.1 [110/5] via 192.168.22.2, 3d02h, GigabitEthernet0/0/3
O       10.4.1.1 [110/3] via 192.168.22.2, 4d04h, GigabitEthernet0/0/3
```

```
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.3.1.1        10.3.1.1        255         0x80000393 0x00CE60 2
10.4.1.1        10.4.1.1        674         0x800000B7 0x009A4F 2
```

OSPF Data Structures

- The OSPF data structures foll:
  - Neighbor table (the adjacency database).
  - Topology table (the LSDB).
  - Routing table (the forwarding database).
- Each router has a full picture of the topology.
- Link-state routers tend to make more accurate decisions.

Structure of OSPF Network

Area terminology and router types:

- ABR
- Internal router
- ASBR
- Backbone router

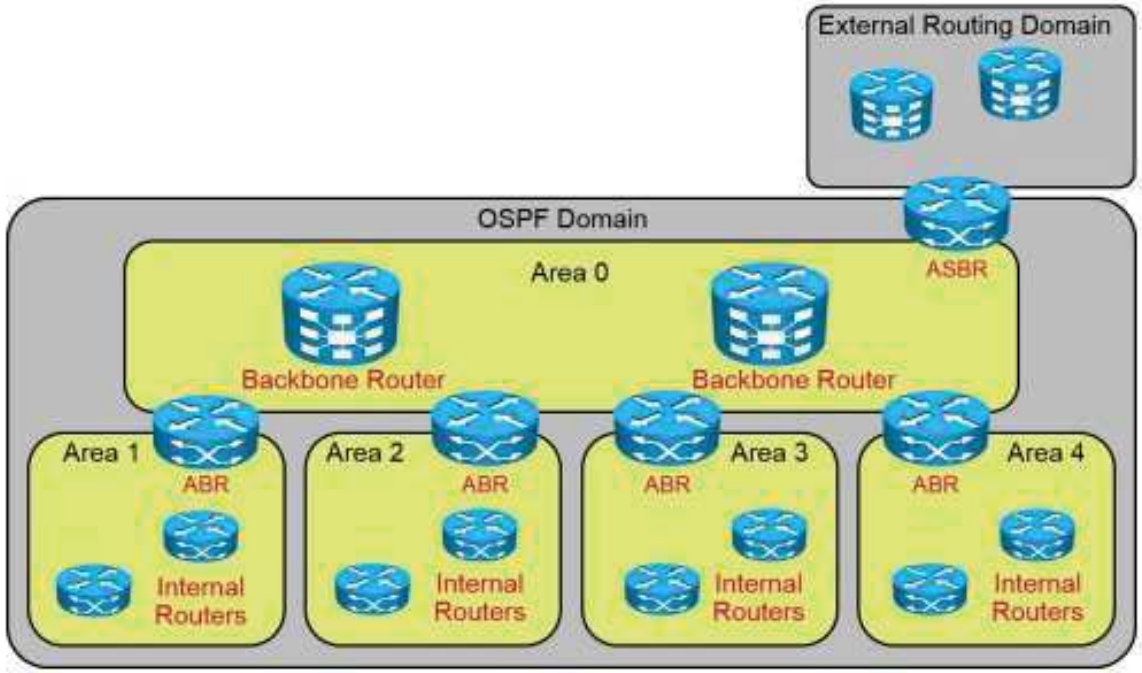

Structure of OSPF Network (Cont.)

OSPF Areas:

- Backbone area—Area 0
- Regular nonbackbone area
- Stubby area or totally stubby area
- NSSA or totally NSSA

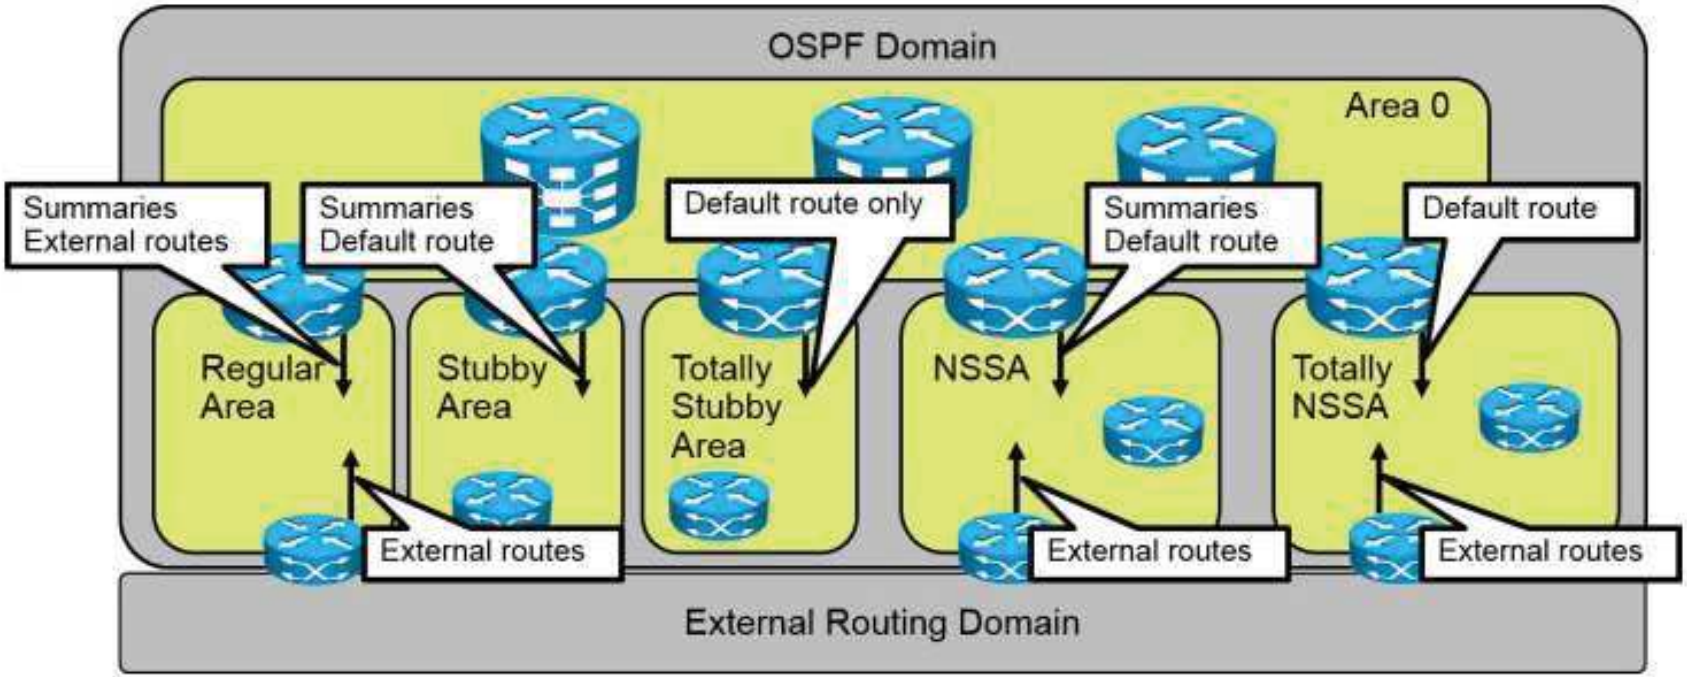

Hierarchical Structure of OSPF in Service Provider Environment

- Link-state routing requires a hierarchical network structure.
- OSPF area characteristics:
  - Minimizes routing table entries
  - Localizes impact of a topology change (link flapping) within an area
  - Detailed LSA flooding stops at area boundary

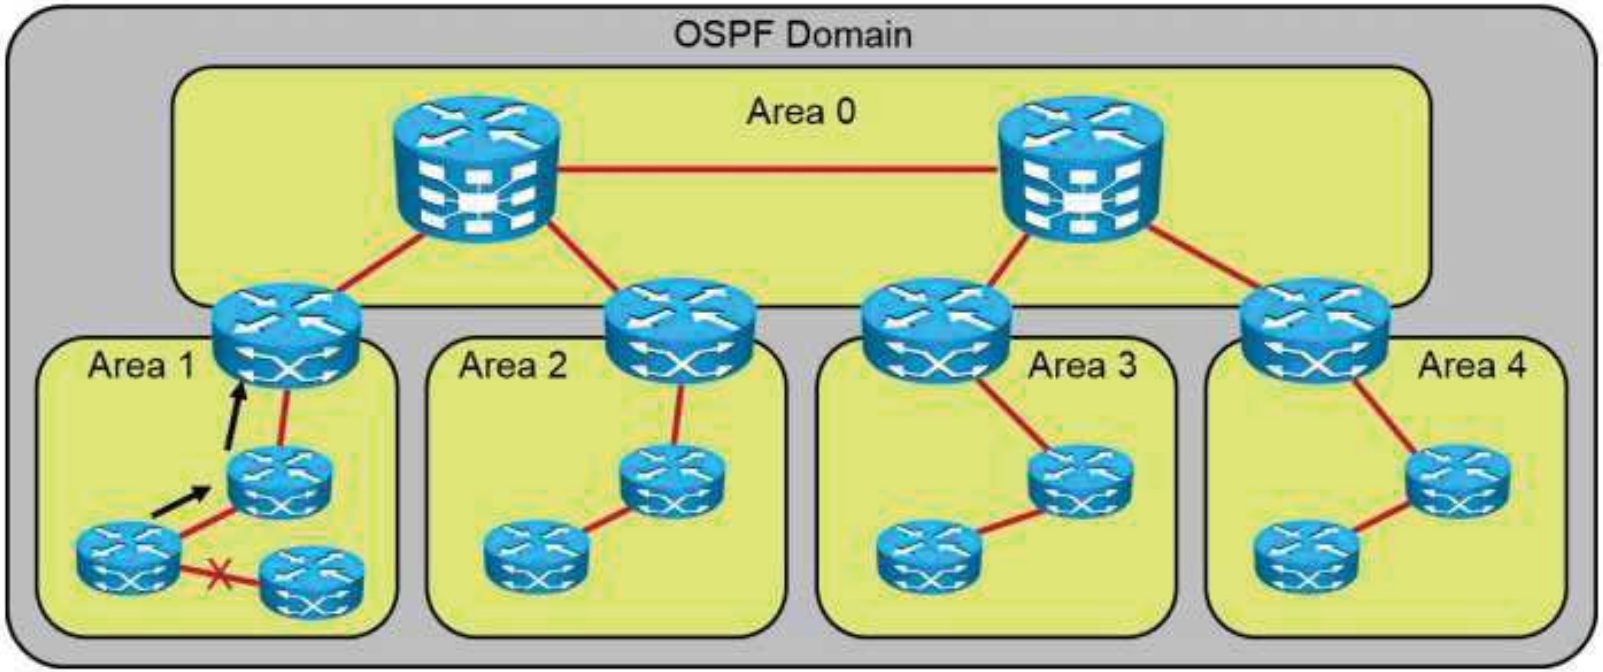

OSPF LSA Types

LSA Type	OSPFv2	OSPFv3
1	Router LSAs	Router LSAs
2	Network LSAs	Network LSAs
3	Summary LSAs	Interarea-prefix LSAs for ABRs
4	Summary LSAs	Interarea-router LSAs for ASBRs
5	External LSAs	AS-external LSAs
6	Multicast OSPF LSAs	Group membership LSAs
7	LSAs defined for NSSA	Type-7 LSAs
8	LSAs defined for NSSA	Link LSAs
9	Opaque LSAs	Intra-area-prefix LSAs
10,11	Opaque LSAs	Opaque LSAs

OSPF Stub Areas

- Stub Area:
  - No external routes
  - Interarea routes present
  - Intra-area routes present
  - Default route generated

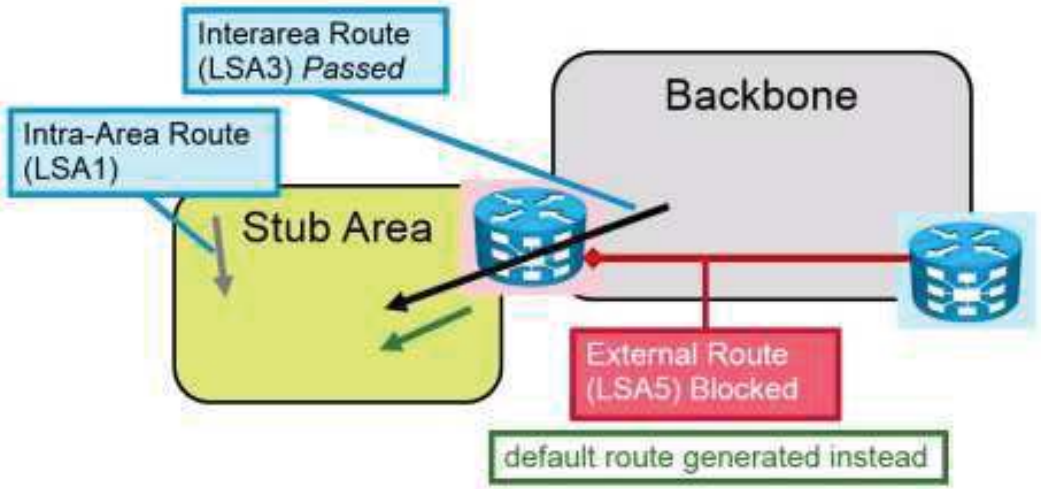

- Totally Stubby Area (stub no-summary):
  - No external routes
  - No interarea routes
  - Intra-area routes present
  - Default route generated
  - Cisco proprietary feature

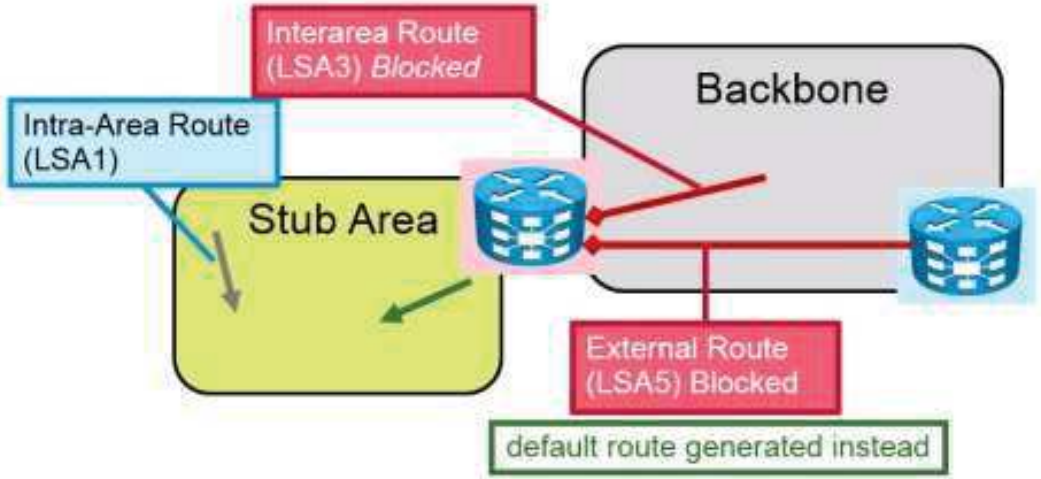

OSPF Not-So-Stubby Areas

- NSSA:
  - Behaves like stub area. It may introduce external routes locally in the area.
- Totally NSSA no-summary:
  - Behaves like totally stubby area. It may introduce external routes locally in the area.
  - Cisco proprietary feature

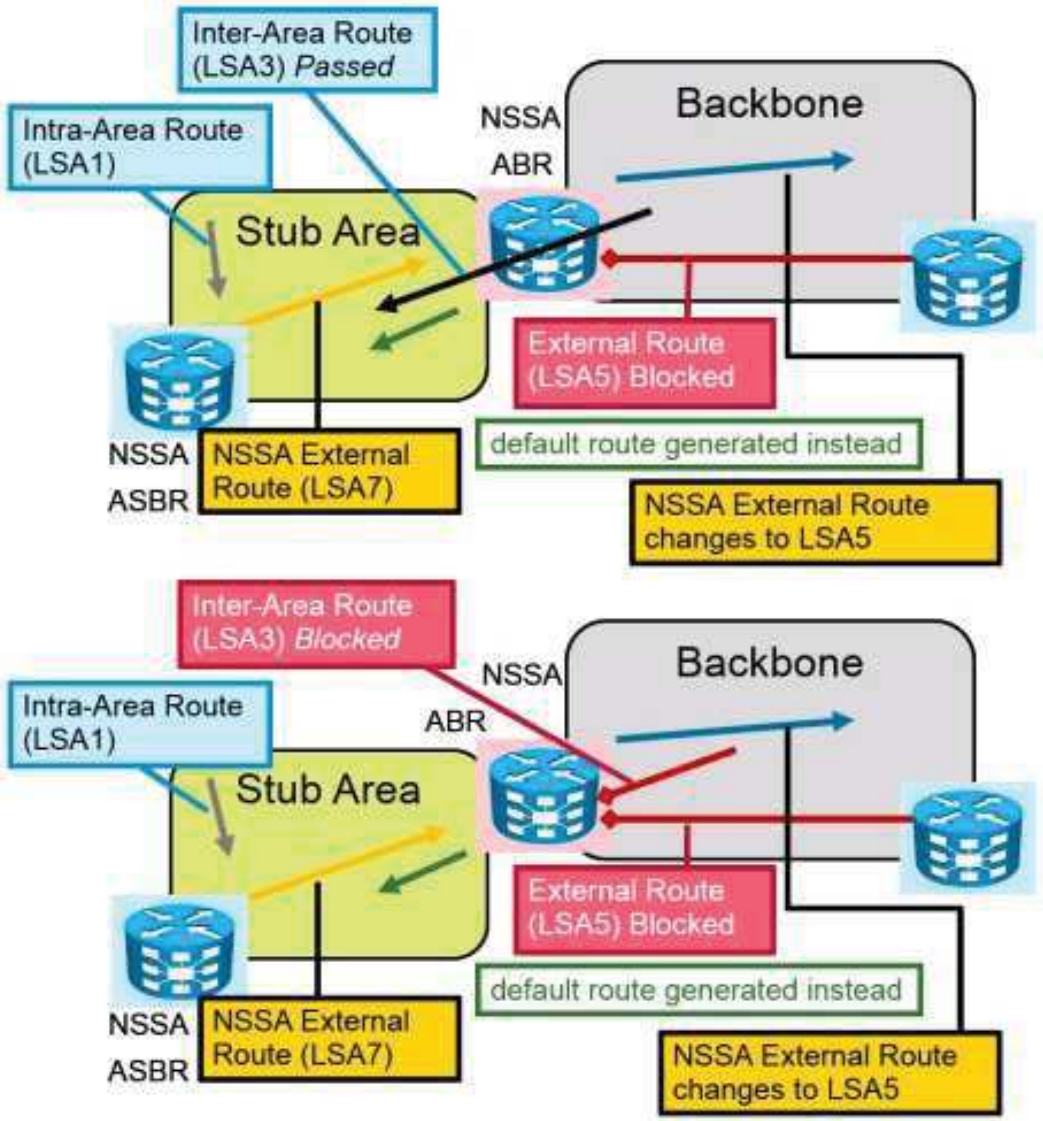

OSPF Operation

OSPF Adjacencies on Point-to-Point Link:

- Routing updates and topology information are passed only between adjacent routers.
- OSPF adjacencies are formed on point-to-point links.
  - Sends OSPF packets using multicast 224.0.0.5 (IPv4) or FF02::5 (IPv6)

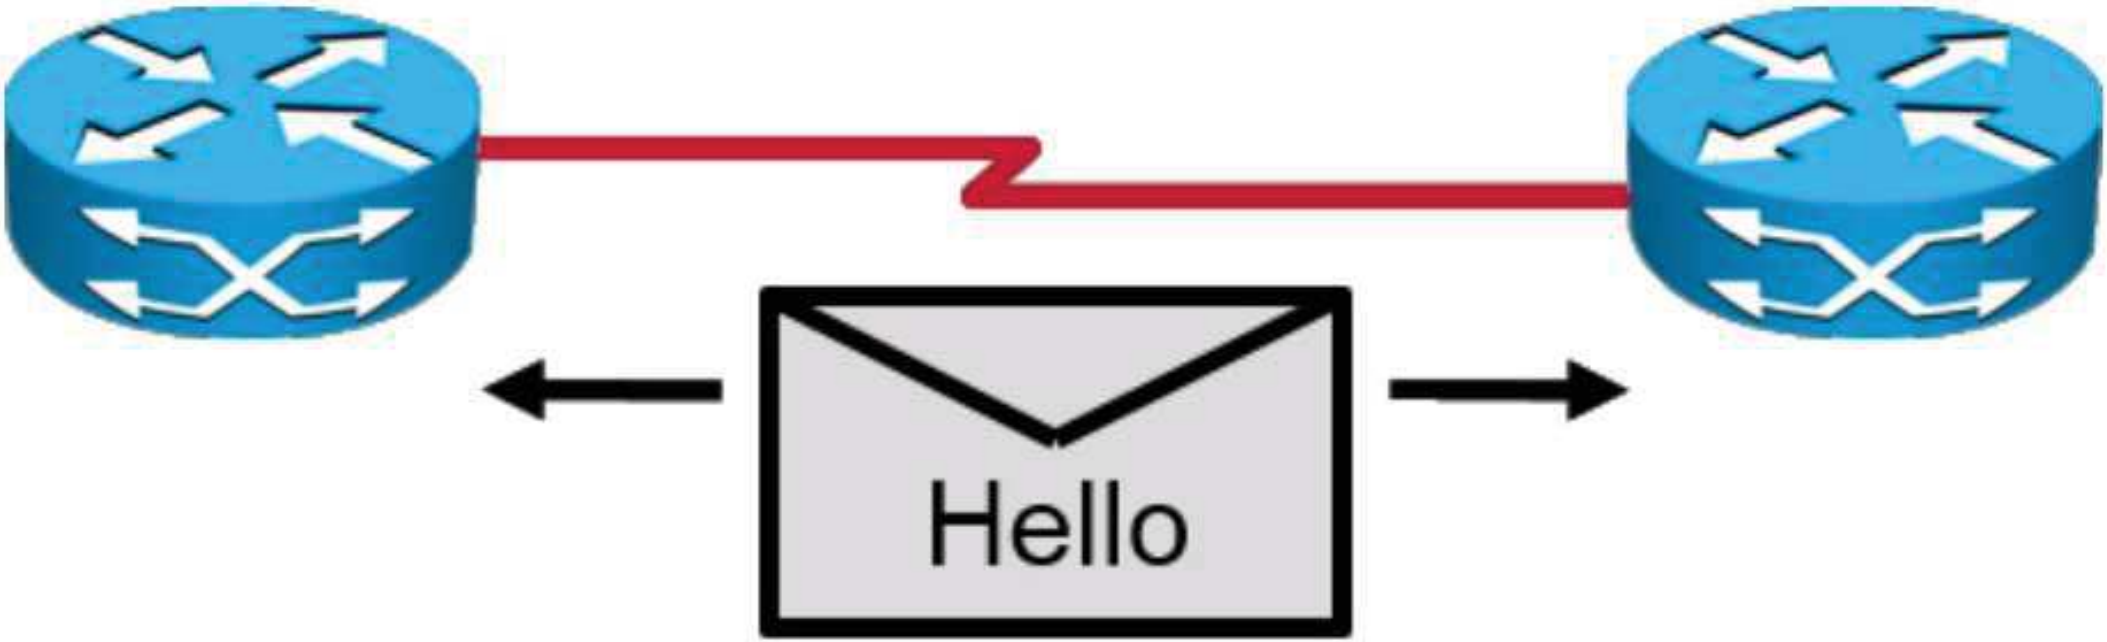

OSPF Operation (Cont.)

OSPF Adjacencies on the LAN Link:

- Forming OSPF adjacencies on LAN links is different than forming them on point-to-point links.
  - Requires DR or BDR election
  - DR and BDR reduce routing update traffic and manage link-state synchronization
  - Sends OSPF packets using multicast 224.0.0.5, 224.0.0.6 (IPv4) or FF02::5, FF02::6 (IPv6)

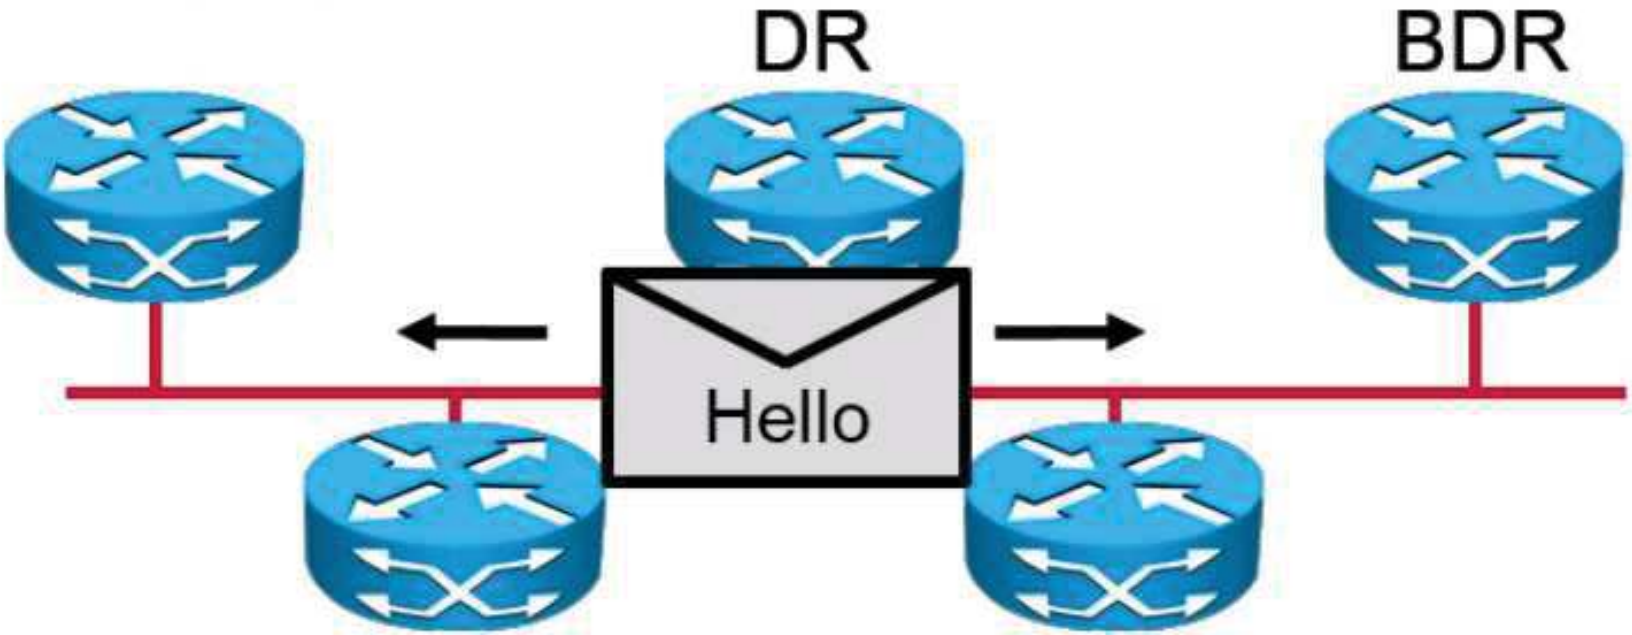

OSPF Best Path Calculation

This is how OSPF calculates the best path to a network:

- Routers find the best paths to destinations by applying the Dijkstra SPF algorithm to the LSDB.
- The best path is calculated based on the lowest total cost and sent to the routing table.

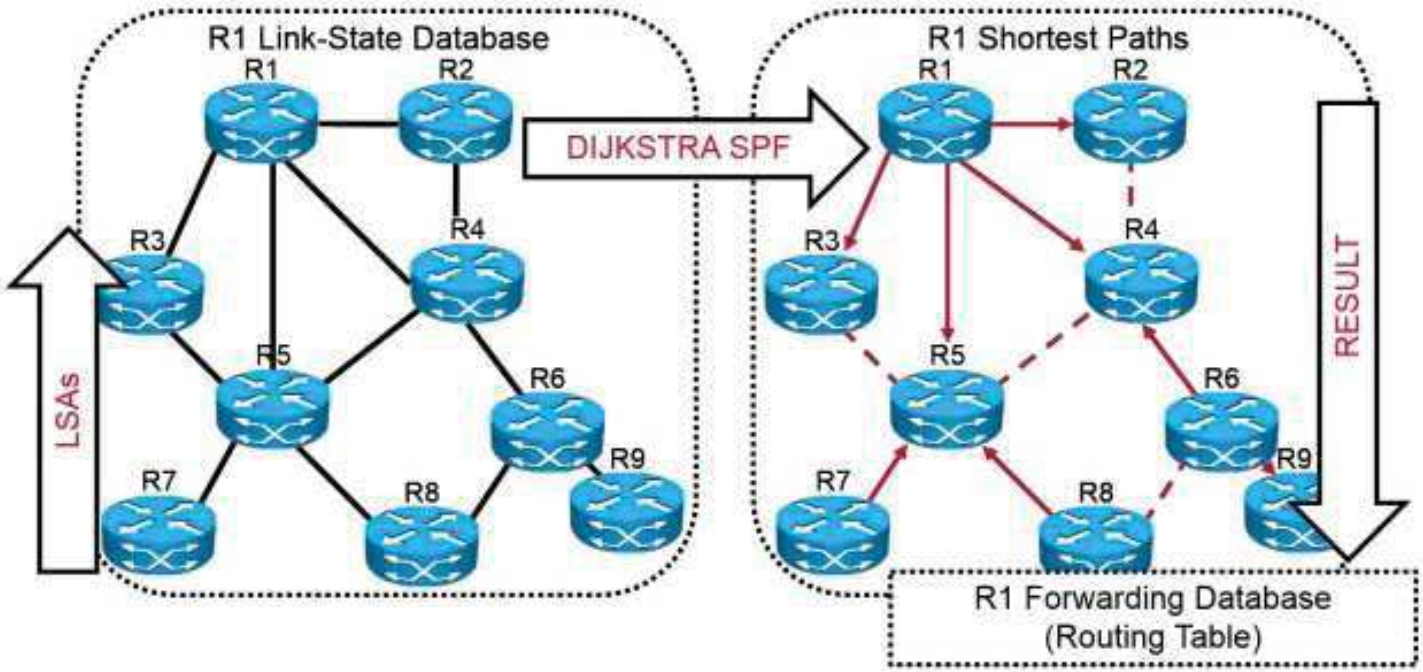

OSPF Metric

These are the main characteristics of the OSPF metric:

- Also called “cost”; COST = 100,000,000 / bandwidth [bps]
- Defined per interface, but may be altered
- Inversely proportional to the bandwidth of that interface

Link Type	Default Cost	Cost with reference 1010 bps
64 kbps serial link	1562	156250
T1 (1.544 Mbps serial link)	64	6476
E1 (2.048 Mbps serial link)	48	4882
Ethernet	10	1000
Fast Ethernet	1	100
Gigabit Ethernet	1	10
10 Gigabit Ethernet	1	1

Building the LSDB

This is how the LSDB is constructed:

- The Hello protocol is used to define neighbors.
- Adjacency is established.
  - Area ID
  - Authentication
  - Hello and dead intervals
  - Stub area flag
- Adjacent routers exchange LSAs.
- Each router builds an LSDB using LSAs.

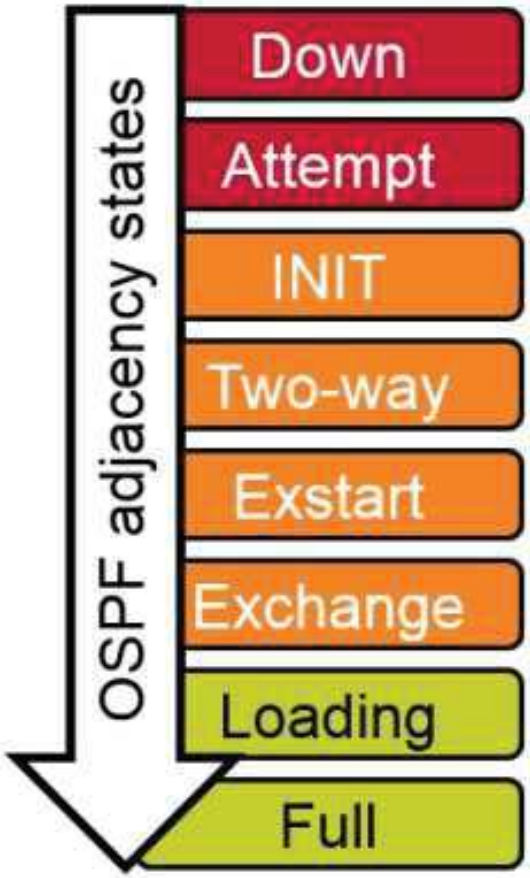

LSA Operation

An LSA is more recent if it has the following:

- A higher sequence number.
- A higher checksum number.
- An age equal to the maximum age (poisoning).
- A significantly smaller link-state age (the LSA is significantly younger).

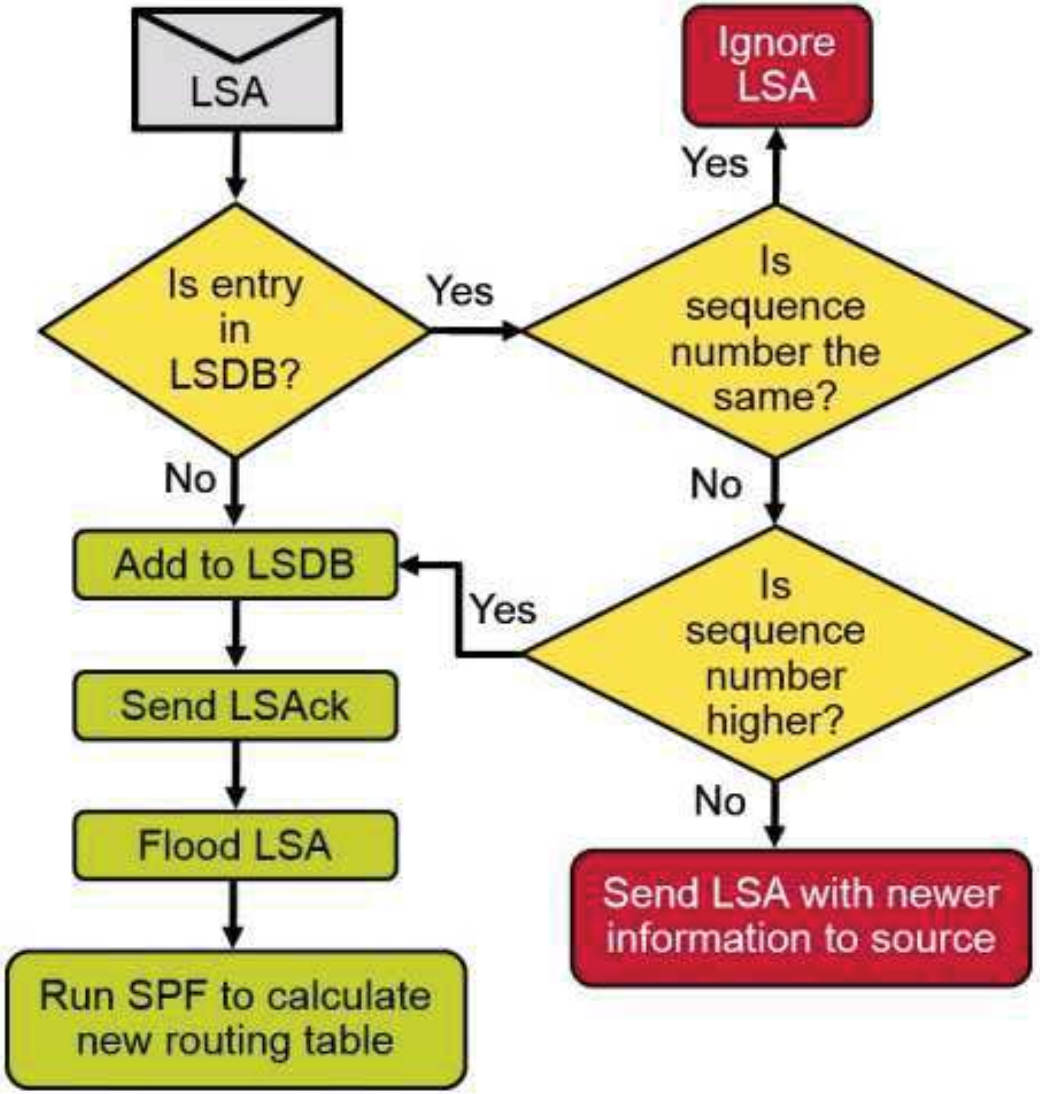

LSA Operation (Cont.)

- Each LSA in the LSDB maintains a sequence number:
  - 4-byte number
  - Begins with 0x80000001; ends with 0x7FFFFFFF
- OSPF floods each LSA every 30 minutes.
  - Each time, the sequence number is incremented by one.
  - The LSA with the higher (newer) sequence number is more recent.
  - Ultimately, a sequence number will wrap around to 0x80000001.
  - The existing LSA was prematurely aged to the maximum age (one hour) and flushed:

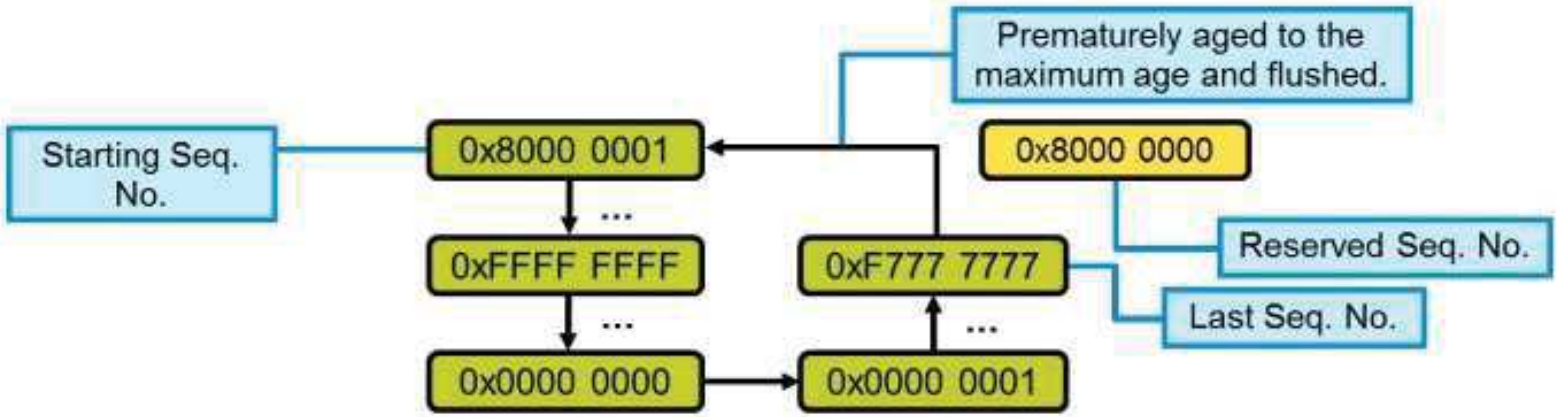

LSA Operation (Cont.)

LSA sequence numbers and maximum age:

- Every OSPF router announces a router LSA for those interfaces that it owns in that area.
- Router with link ID 10.1.10.1 has been updated three times; the last update was 115 seconds ago.


```
RP/0/RSP0/CPU0:PE1# show ospf database
            OSPF Router with ID (10.1.1.1) (Process ID 1)

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        114         0x80000001 0x0049c1 1

                Router Link States (Area 11)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        114         0x80000002 0x004322 1
10.1.10.1       10.1.10.1       115         0x80000003 0x004ddc 2
```

Router and Network LSAs

- One router LSA for every router in an area:
  - Includes a list of directly attached links.
  - Links identified by the IP prefix and link type.
- LSA identified by the router ID of the originating router.
- Floods within its area only; does not cross an ABR.
- Advertised by the DR.
- Floods within its area only; does not cross an ABR.

- One network LSA for each transit broadcast or NBMA network:
  - Includes a list of attached routers on the transit link.
  - Includes a subnet mask of the link.
- Advertised by the DR.
- Floods within its area only; does not cross an ABR.

Router and Network LSAs (Cont.)

Router LSAs

Link Type	Description	Link ID
1	Point-to-point connection to another router	Neighboring router ID
2	Connection to a transit network	IP address of DR
3	Connection to a stub network	IP network or subnet number
4	Virtual link	Neighboring router ID

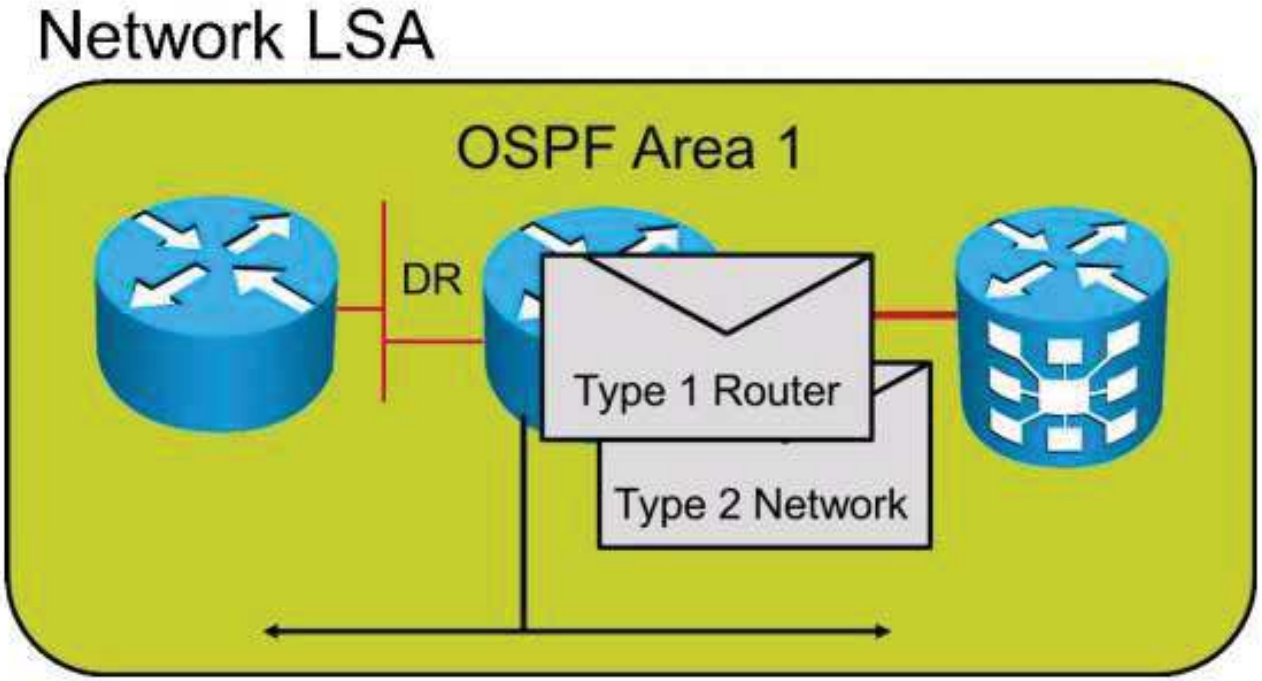

Summary LSAs

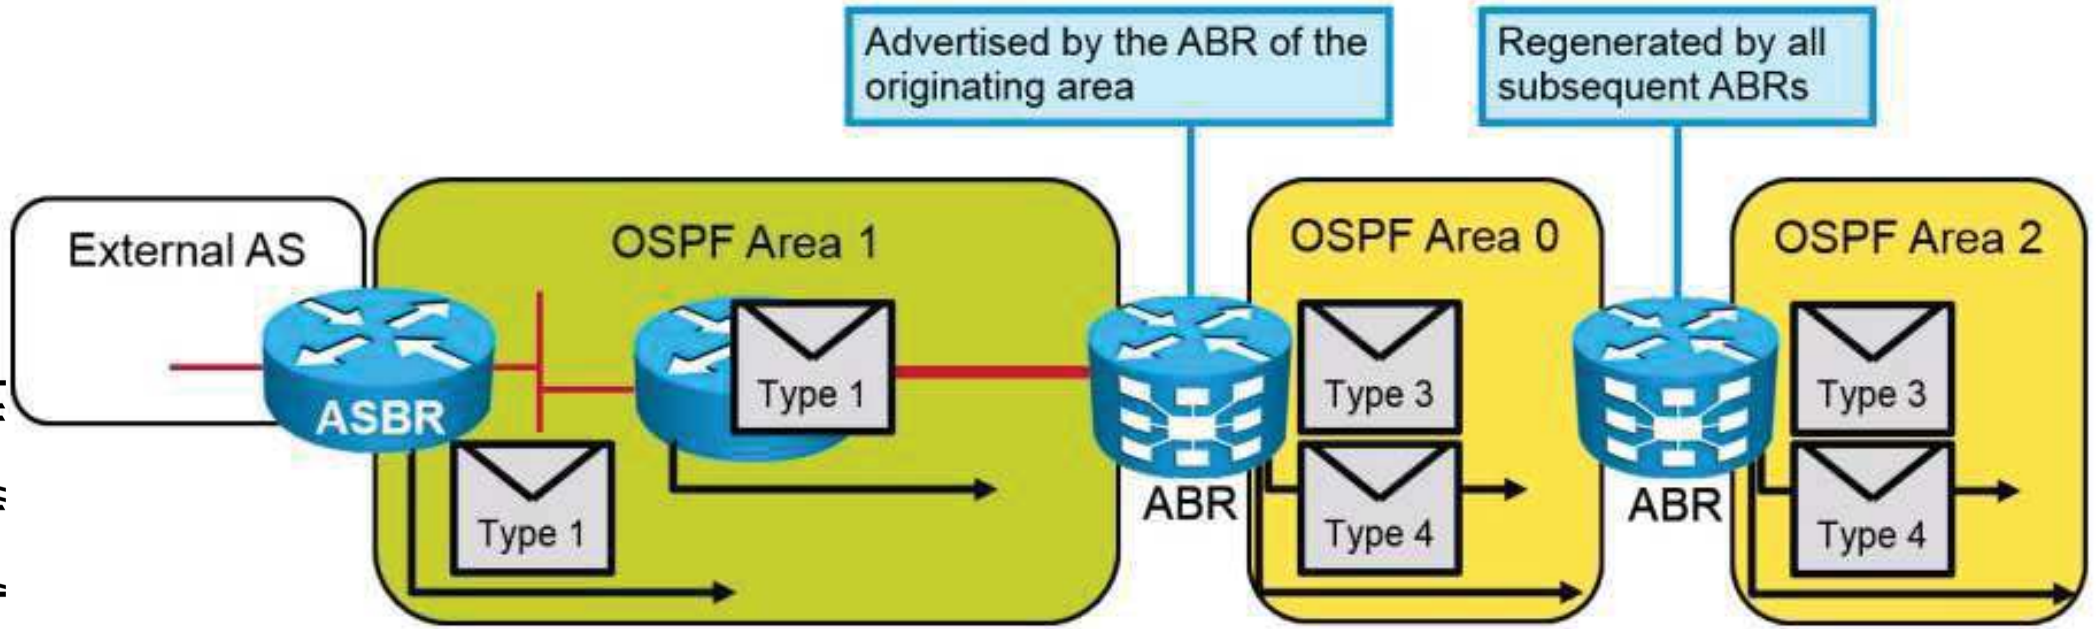

- **LSA type 3:** Used to flood network information to areas outside the originating area
- Describes the network number and mask of the link
- Advertised for every subnet and not summarized, by default

- **LSA type 4:** Used to advertise a metric to the ASBR, which is used for path selection
- Contains the router ID of the ASBR

External LSAs

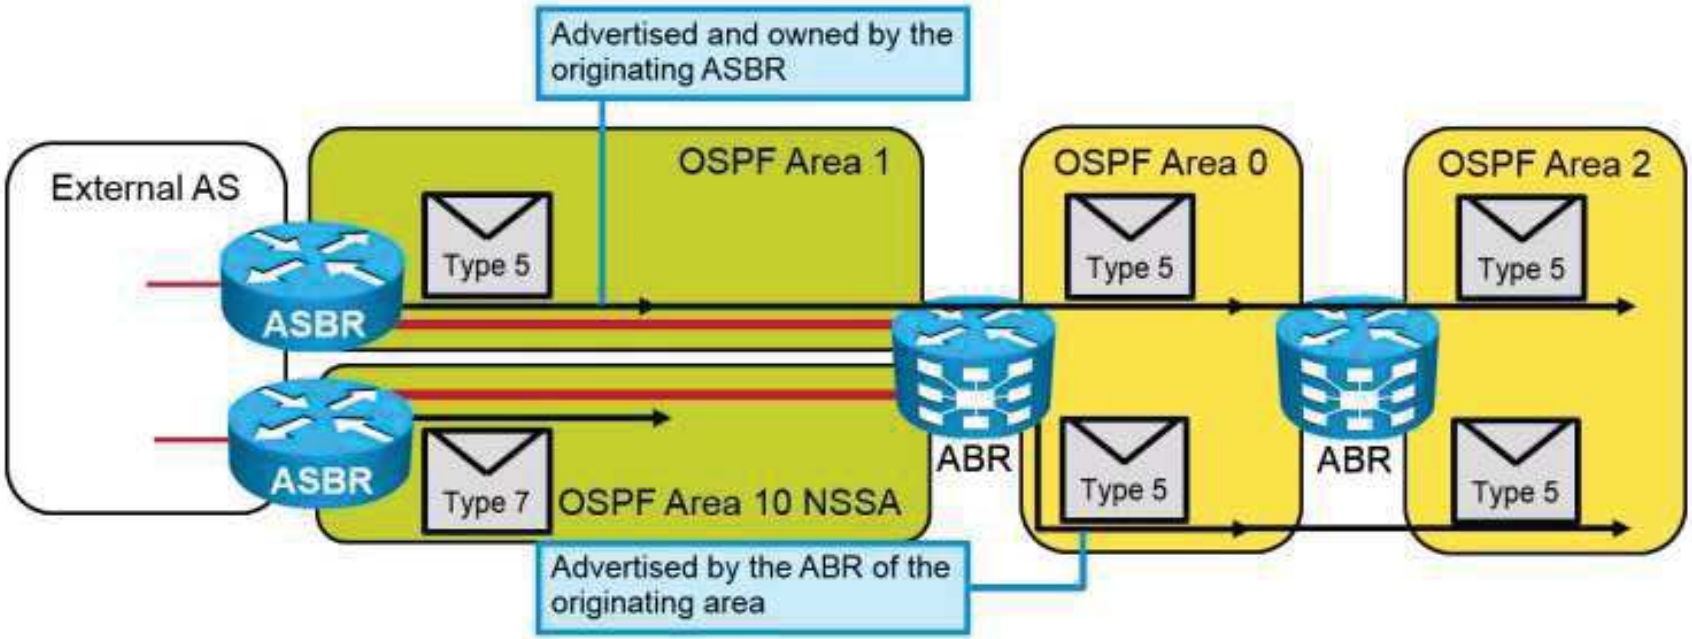

- **LSA type 5:** Used to advertise networks from other autonomous systems
- Flooded throughout entire AS
- Advertising router ID unchanged throughout the AS
- Type 4 LSA is needed to find the ASBR

- **LSA type 7:** Used to advertise networks from other ASs injected into the NSSA
- Same format as a type 5 external LSA
- Translated to LSA type 5 on the NSSA ABR, then propagated as LSA type 5 by subsequent ABR

OSPF Intra-Area Routing

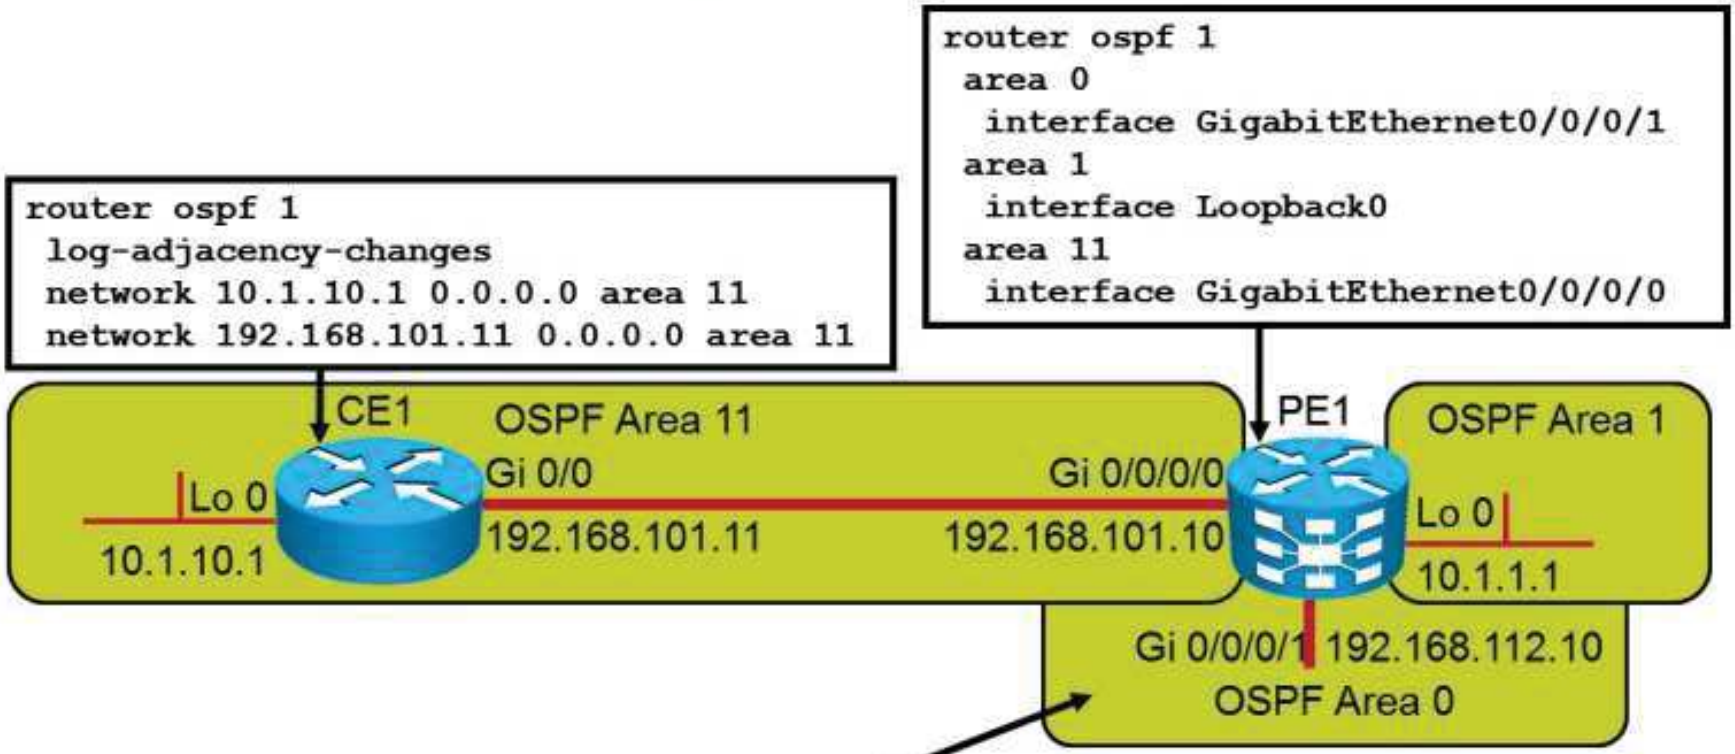


```
RP/0/RSP0/CPU0:PE1# show ospf database
            OSPF Router with ID (10.1.1.1) (Process ID 1)

                Router Link States (Area 0)          <-- LSA Type 1 from Area 0

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        223         0x80000001 0x0069d3 1

                Net Link States (Area 11)            <-- LSA Type 2 from Area 11

Link ID         ADV Router      Age         Seq#       Checksum
192.168.101.11  10.1.10.1       56          0x8000002c 0x006794
```

OSPF Inter-Area Routing

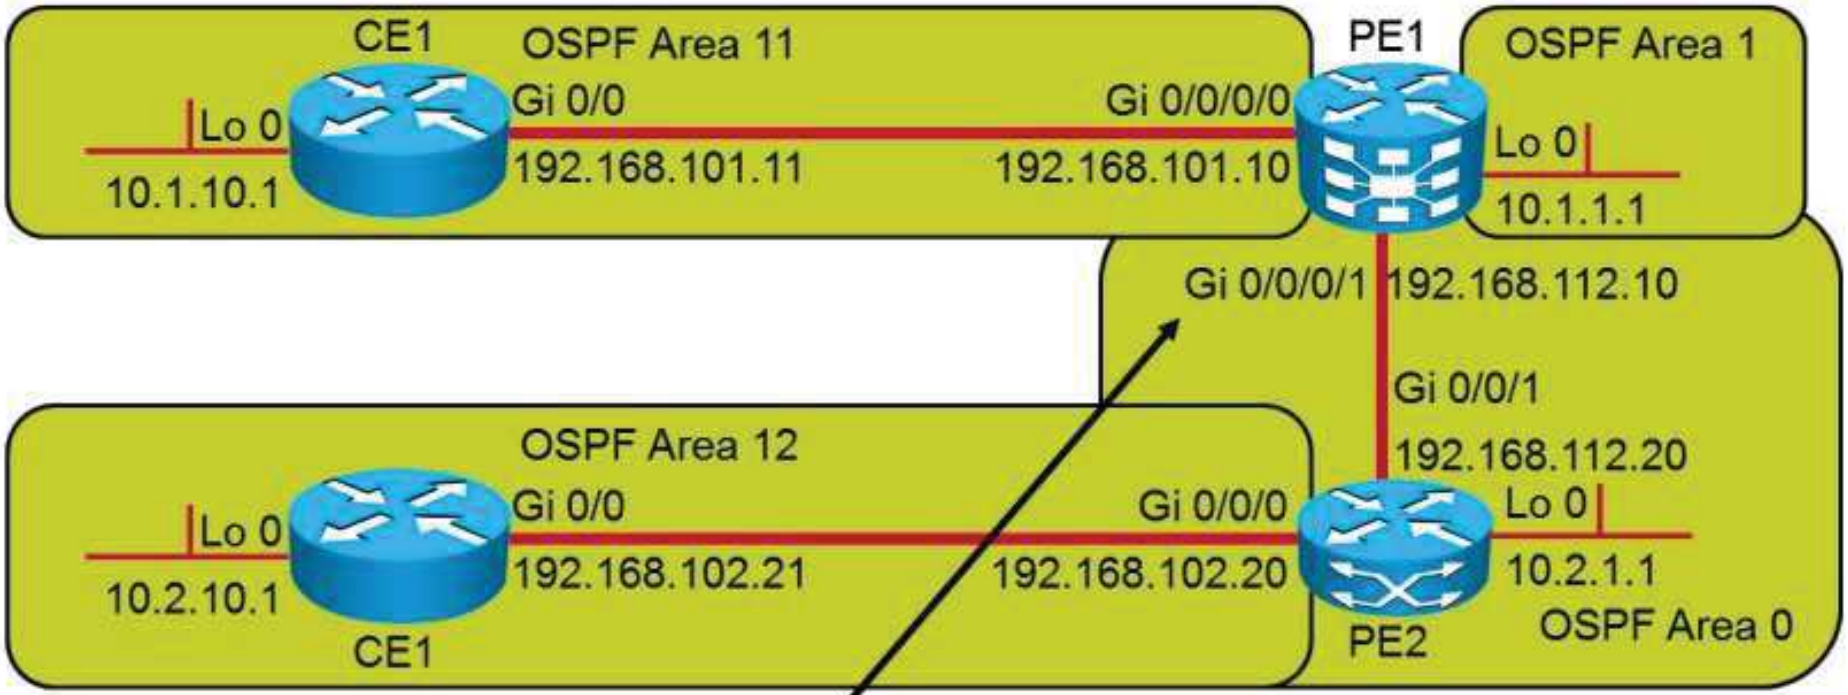


```
RP/0/RSP0/CPU0:PE1# show ospf database
            OSPF Router with ID (10.1.1.1) (Process ID 1)

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        1614        0x80000003 0x007ca3
10.1.10.1       10.1.1.1        1614        0x80000003 0x0023f2
10.2.10.1       10.2.1.1        205         0x80000001 0x001303
192.168.101.0   10.1.1.1        1614        0x80000003 0x00134b
192.168.102.0   10.2.1.1        205         0x80000001 0x00045a
```

LSA Type 3 from Area 0

OSPF External Routes

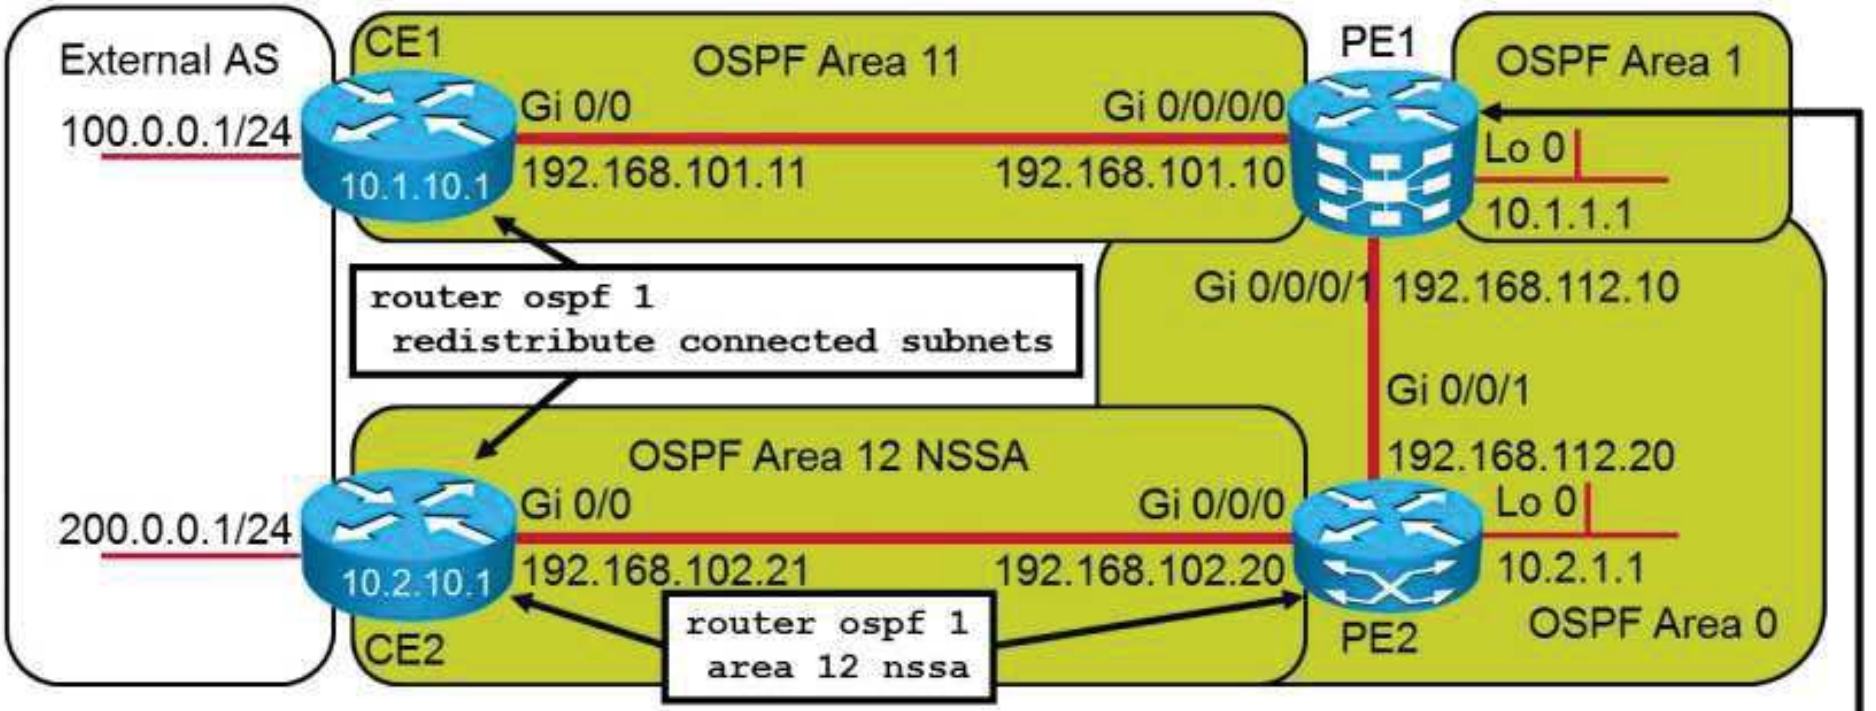


```
RP/0/RSP0/CPU0:PE1# show ospf database
            OSPF Router with ID (10.1.1.1) (Process ID 1)
                Summary ASB Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.10.1       10.1.1.1        1477        0x80000001 0x000f09

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
100.0.0.0       10.1.10.1       1454        0x80000001 0x00ed3d 0
200.0.0.0       10.2.1.1        1267        0x80000001 0x002394 0
```

LSA Type 4 from Area 0

LSA Type 5

Re-originated by ABR (PE2)

OSPF Virtual Link

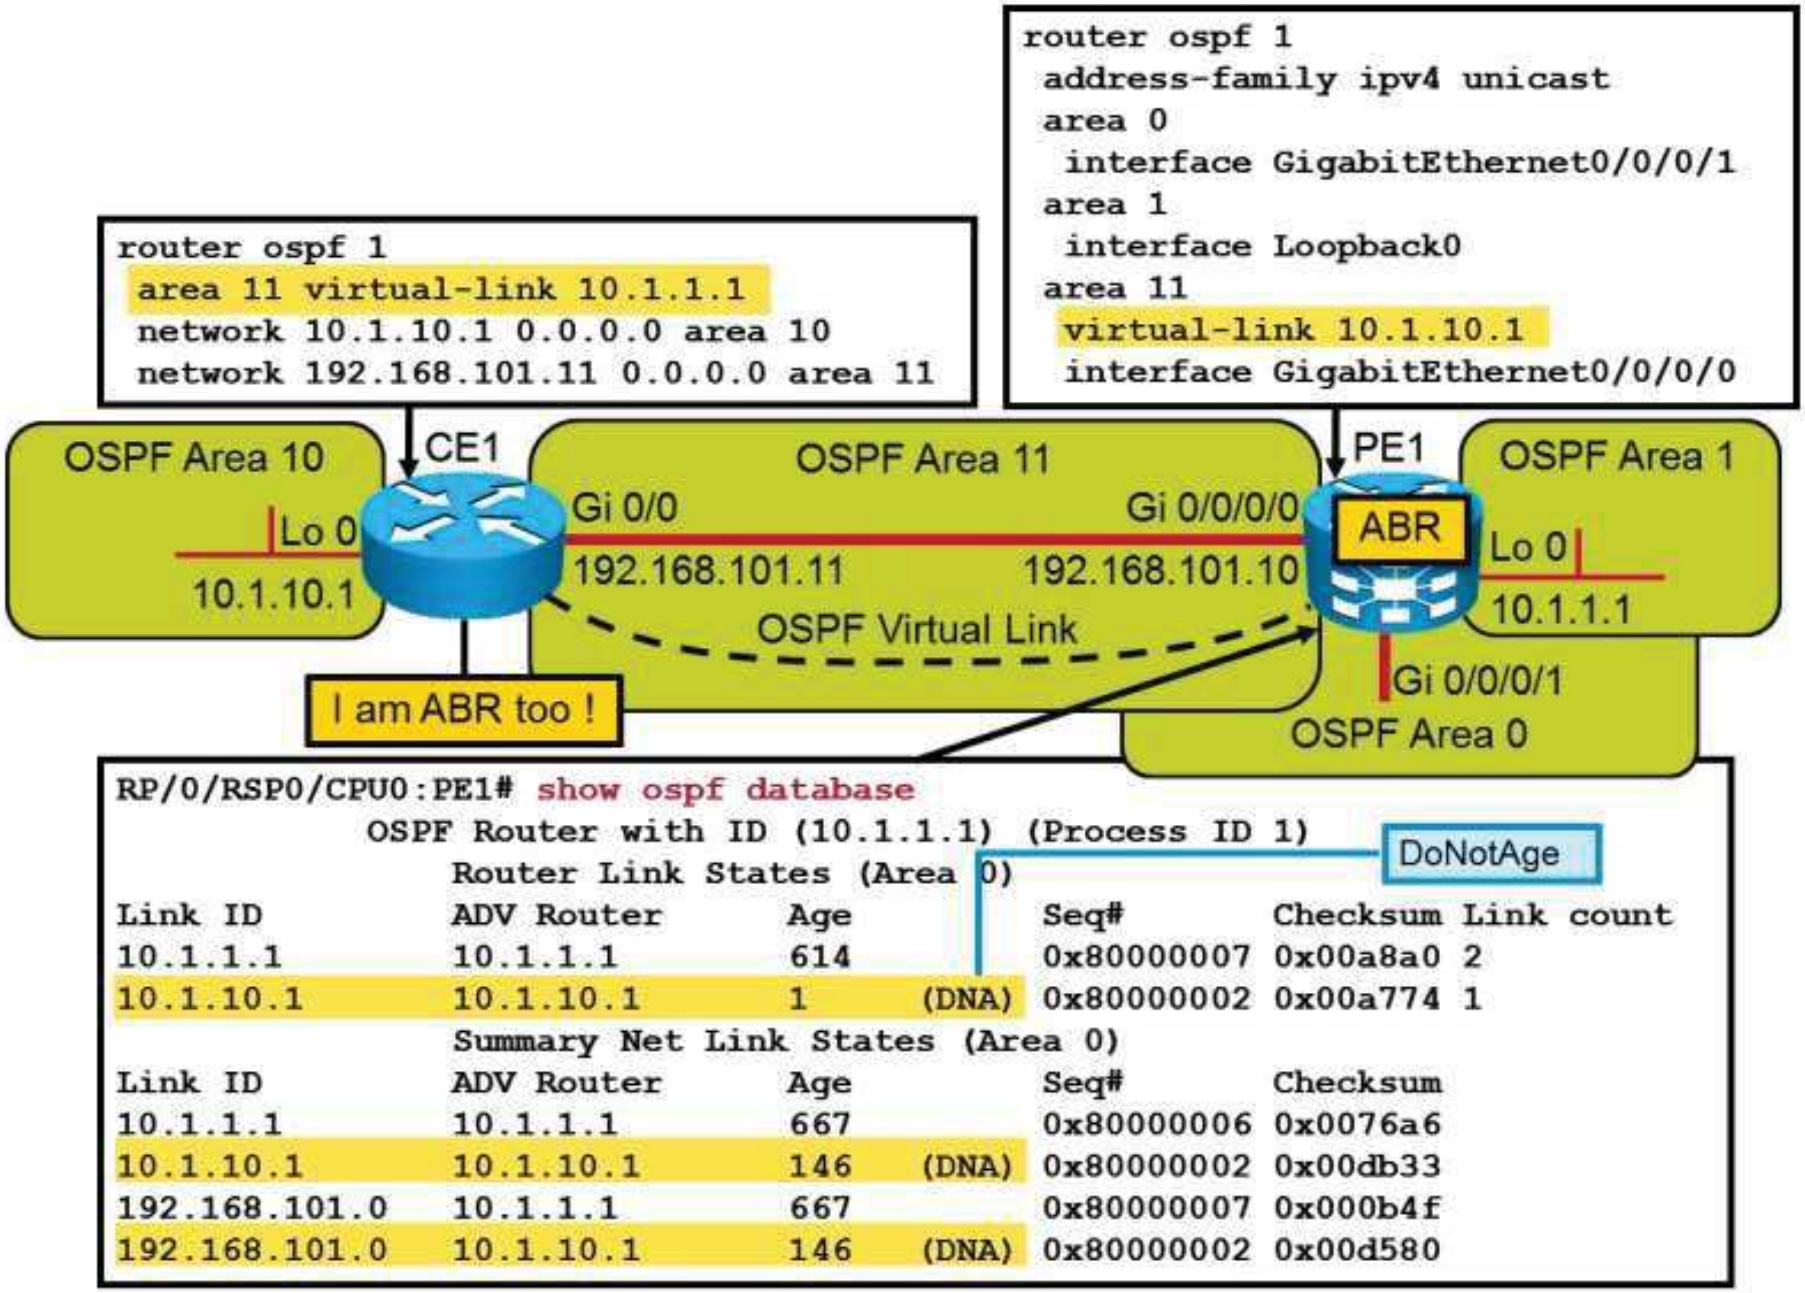

Interpreting OSPF Routes in the Routing Table

OSPF interarea (summary LSA):
- Networks from outside the area of the router, but within the OSPF autonomous system.
- Advertised by means of summary LSAs.

```
RP/0/RSP0/CPU0:PE1# show route ospf
O IA 10.1.10.1/32 [110/2] via 192.168.101.11, 00:31:01, GigabitEthernet0/0/0/0
O IA 10.2.10.1/32 [110/3] via 192.168.112.20, 01:48:41, GigabitEthernet0/0/0/1
O IA 192.168.102.0/24 [110/2] via 192.168.112.20, 01:48:41, GigabitEthernet0/0/0/1
O E2 200.0.0.0/24 [110/15] via 192.168.112.20, 01:26:31, GigabitEthernet0/0/0/1
```

Type 1 or type 2 external routes:
- Networks outside of the autonomous system of the router.
- Advertised by means of external LSAs.


```
CE1# show ip route ospf
<... output omitted ...>
      10.0.0.0/32 is subnetted, 3 subnets
O IA     10.1.1.1 [110/2] via 192.168.101.
O IA     10.2.10.1 [110/4] via 192.168.101
O IA  192.168.102.0/24
          [110/3] via 192.168.101.10, 00:31:56, GigabitEthernet0/0
O     192.168.112.0/24
          [110/2] via 192.168.101.10, 00:31:56, GigabitEthernet0/0
O E2  200.0.0.0/24 [110/15] via 192.168.101.10, 00:34:37, GigabitEthernet0/0
```

OSPF intra-area router LSA and network LSA:
- Networks from within the area of the router.
- Advertised by means of router LSA and network LSAs.

Calculating Costs for E1 and E2 Routes

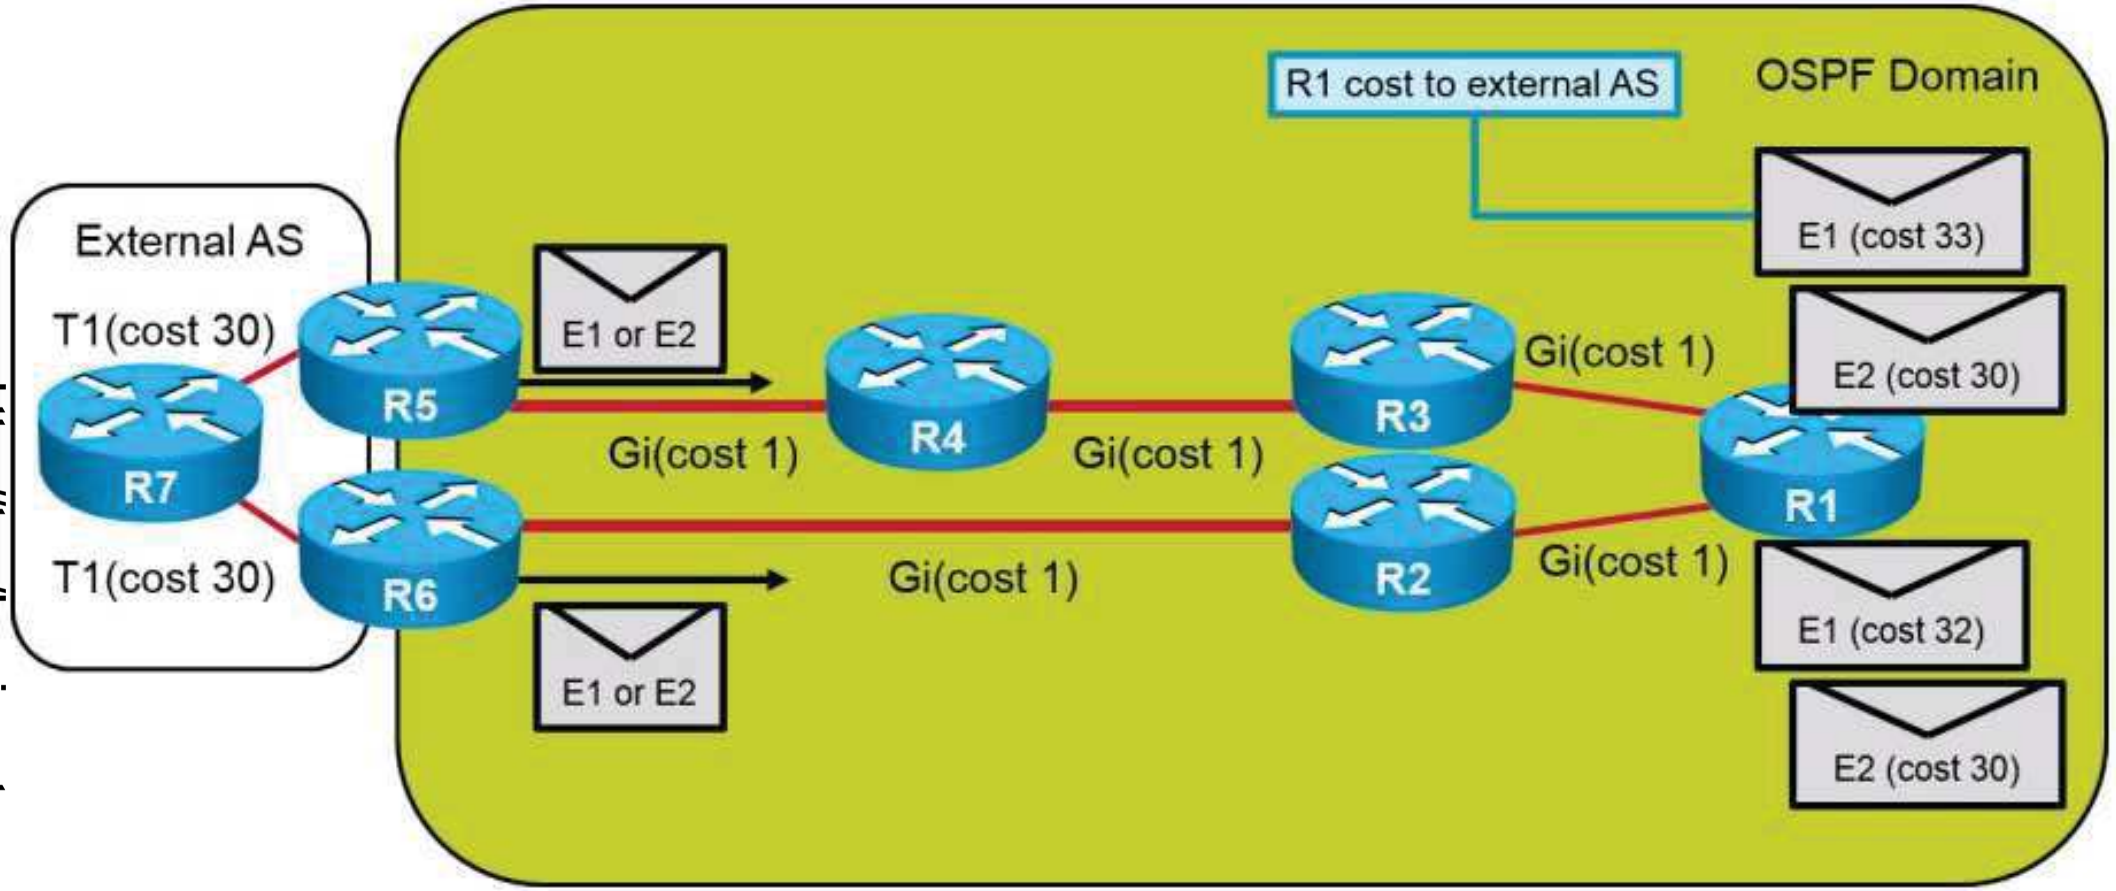

OSPF LSDB Overload Protection

The need for OSPF LSDB Overload Protection feature:

- Excessive LSAs generated by other routers can drain local router resources.
- This feature can limit the processing of nonself-generated LSAs for a defined OSPF process.
- Only a warning message can be sent or neighbors can be ignored.

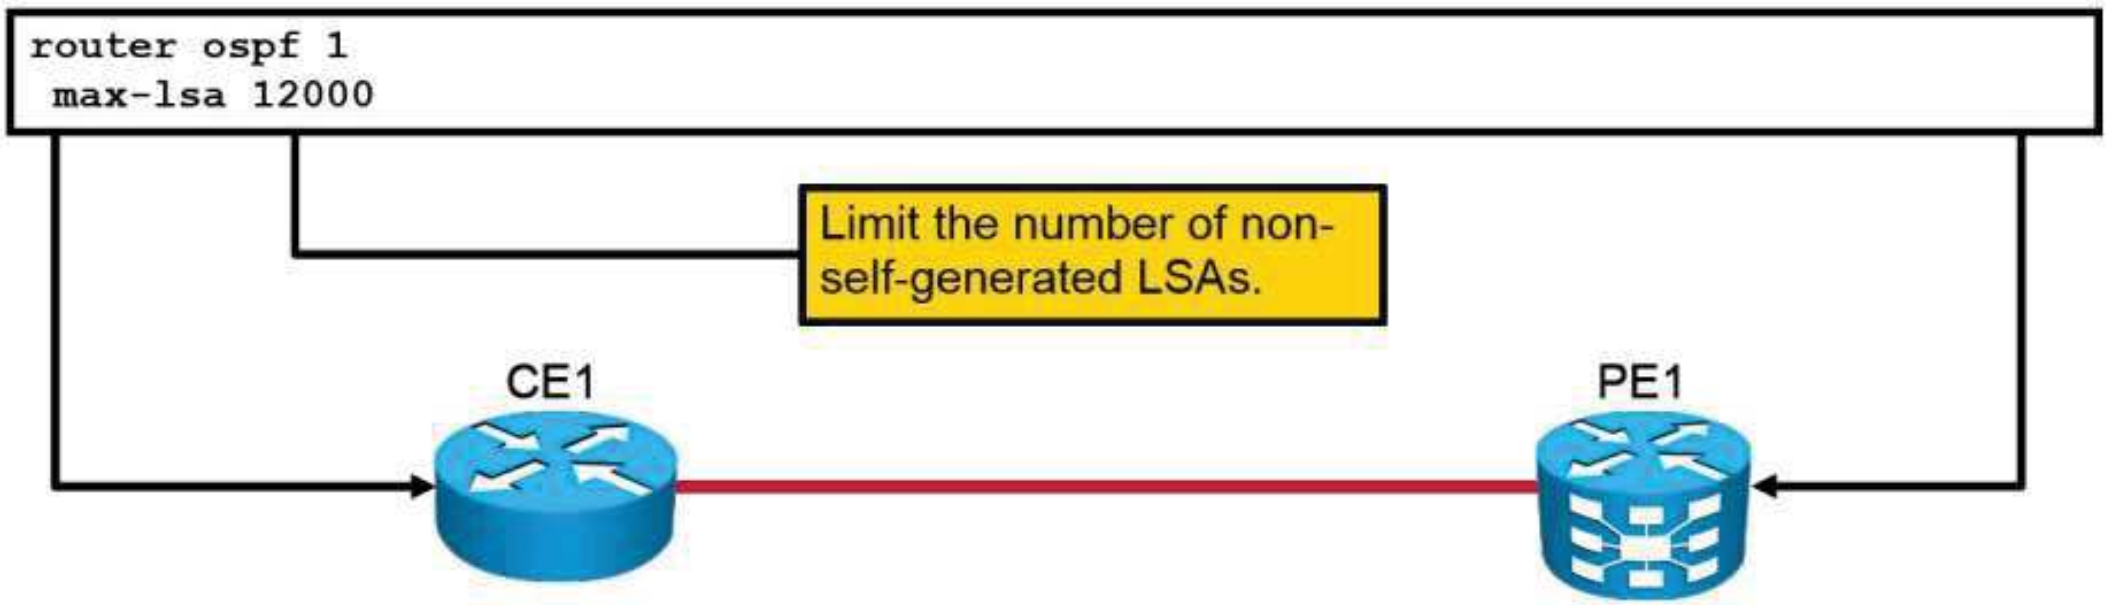

Summary

- OSPF uses link-state advertisements to build a topology database, all OSPF routers have the same picture of network topology.
- ASBR routers connects OSPF area to external routing domain.
- OSPF routing protocol reduces the size of the SPF calculations by partitioning the network into multiple areas.
- Hierarchical structure of OSPF minimizes routing table entries, localizes impact of a topology change and stops LSA flooding.
- There are 11 OSPF LSA types. No external routes are propagated into OSPF stub areas. No external and intra-area routes are propagated into OSPF totally stubby areas.
- OSPF routers find the best path to a destination by applying the Dijkstra SPF algorithm to the topology database.
- OSPF cost of an interface is calculated based on the bandwidth and can be changed by using the OSPF configuration command.

Summary (Cont.)

- OSPF uses hello packets to discover neighbors.
- The following are the most commonly used OSPF LSA types: type 1 router, type 2 network, type 3 and 4 summary, type 5 external, and type 7 external.
- OSPF intra-area LSAs are used to describe routes within an area. OSPF inter-area LSAs are used to describe routes from other areas. OSPF external LSAs are used to describe routes from external autonomous systems.
- Intra-area routes are presented in the routing table with „O“, while inter-area routes are presented with „O IA“.
- External routes can be configured to use either E1 or E2 cost type.
- OSPFv2 and OSPFv3 have run independently on a network device.

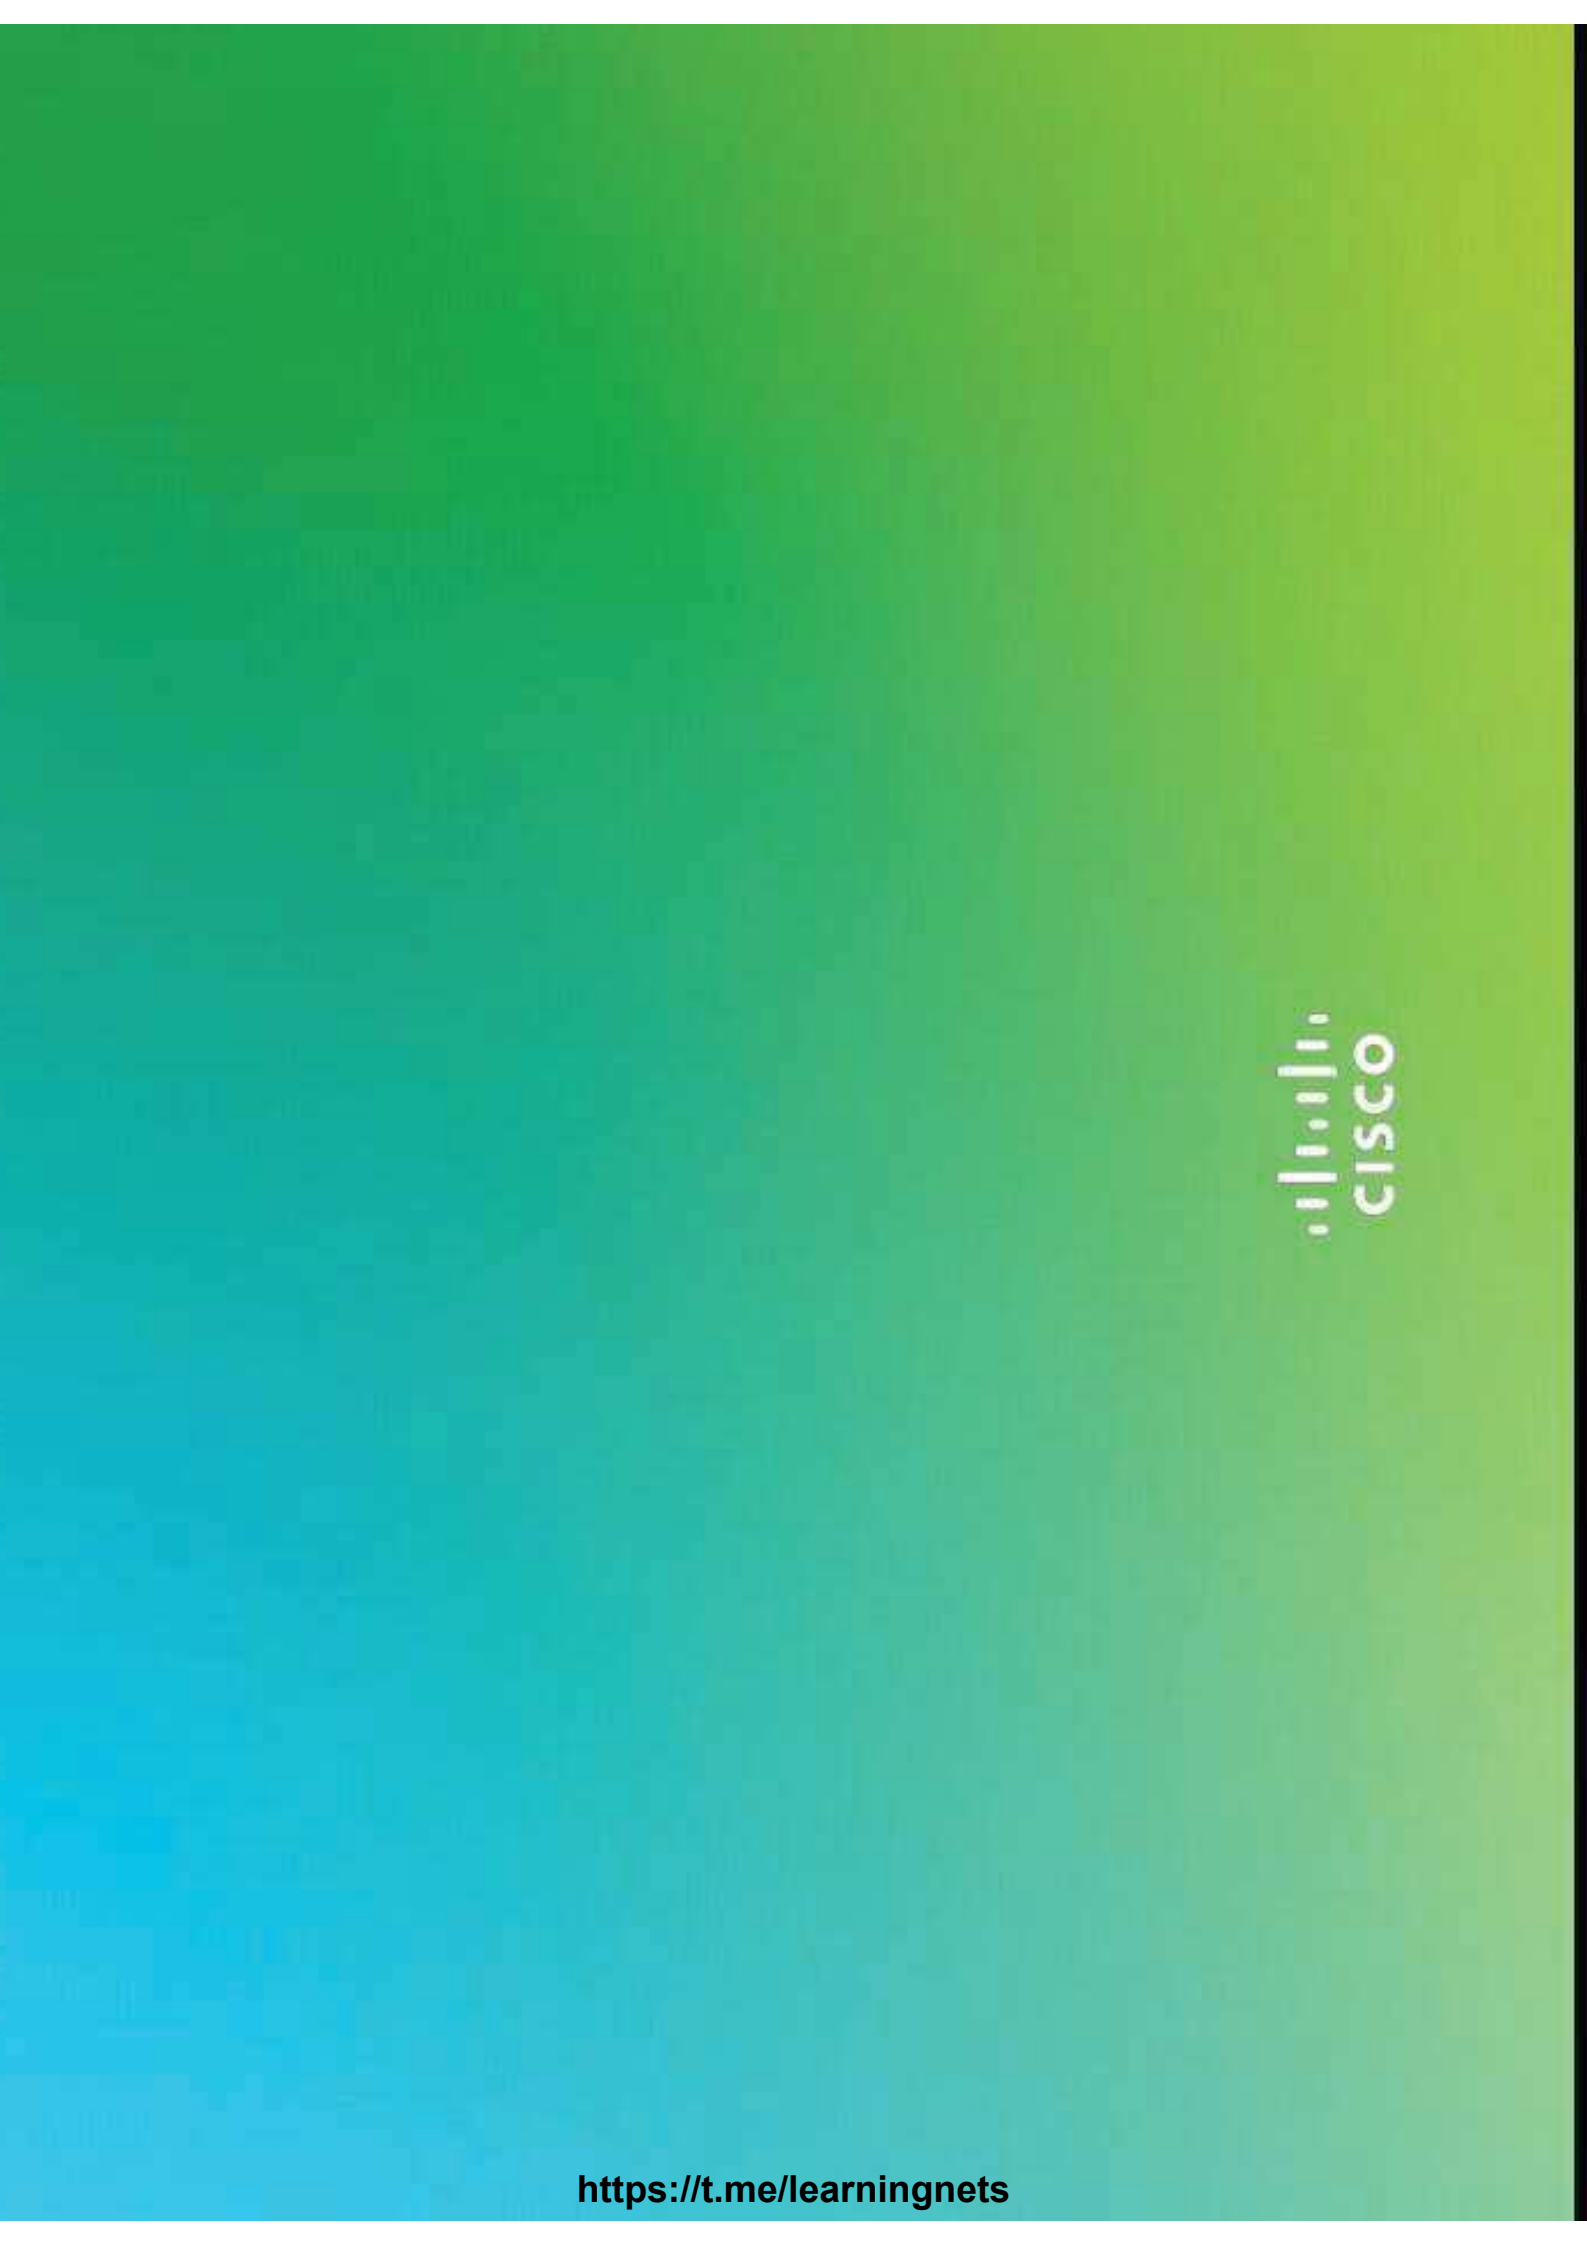

Understanding OSPF Operation

Implement OSPF in the Service Provider Network

OSPF Functions

High-level functions of OSPF include the following:

- Discover neighbors and form adjacencies.
- Flood LSDB information.
- Compute the shortest path.
- Install routes in the route-forwarding table.

Additional functions of OSPF include the following:

- Detect changes in the link state.
- Propagate changes to maintain link-state database synchronization.

OSPF Packet Format

```
CE1# debug ip ospf packet
OSPF packet debugging is on
*Aug 16 09:18:53.974: OSPF: rcv. v:2 t:1 l:48 rid:10.1.1.1
      aid:0.0.0.11 chk:8224 aut:0 auk: from GigabitEthernet0/0
```

Type	Packet	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	DBD	Checks for database synchronization between routers
3	LSR	Requests specific link-state records from another router
4	LSU	Sends specifically requested link-state records
5	LSAck	Acknowledges the other packet types

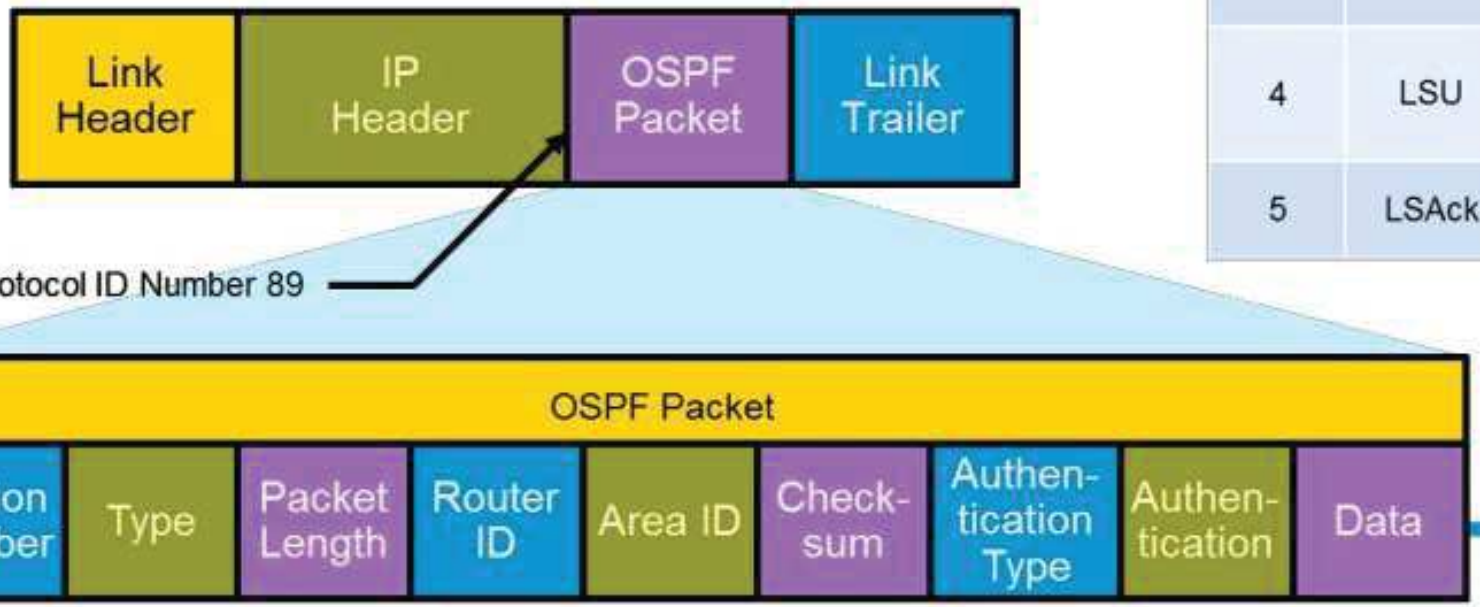

OSPF Packets

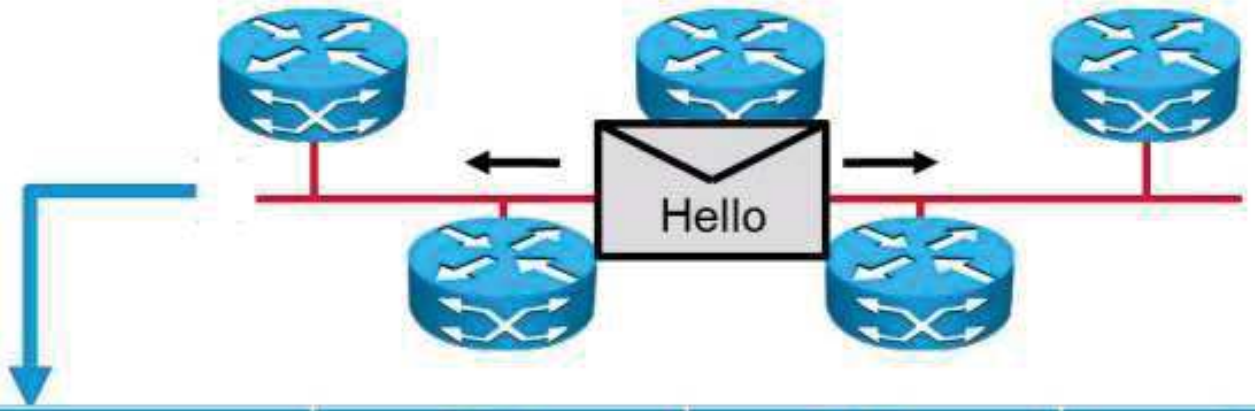

Hello	DBD	LSR	LSU	LSAck
Router ID				
Hello and dead intervals*				
Neighbors				
Area ID*				
Router priority				
DR IP address				
BDR IP address				
Authentication password*				
Stub area flag*				

* Entries must match on neighboring routers

- Use of multicast and unicast IP address.
- Four types of update packets.
- LSDB synchronization process:
  - Discover neighbor.
  - Establish bidirectional communication.
  - Elect a designated router, if desired.
  - Form an adjacency.
  - Discover the network routes.
  - Update and synchronize link-state databases.

OSPF Neighbor States

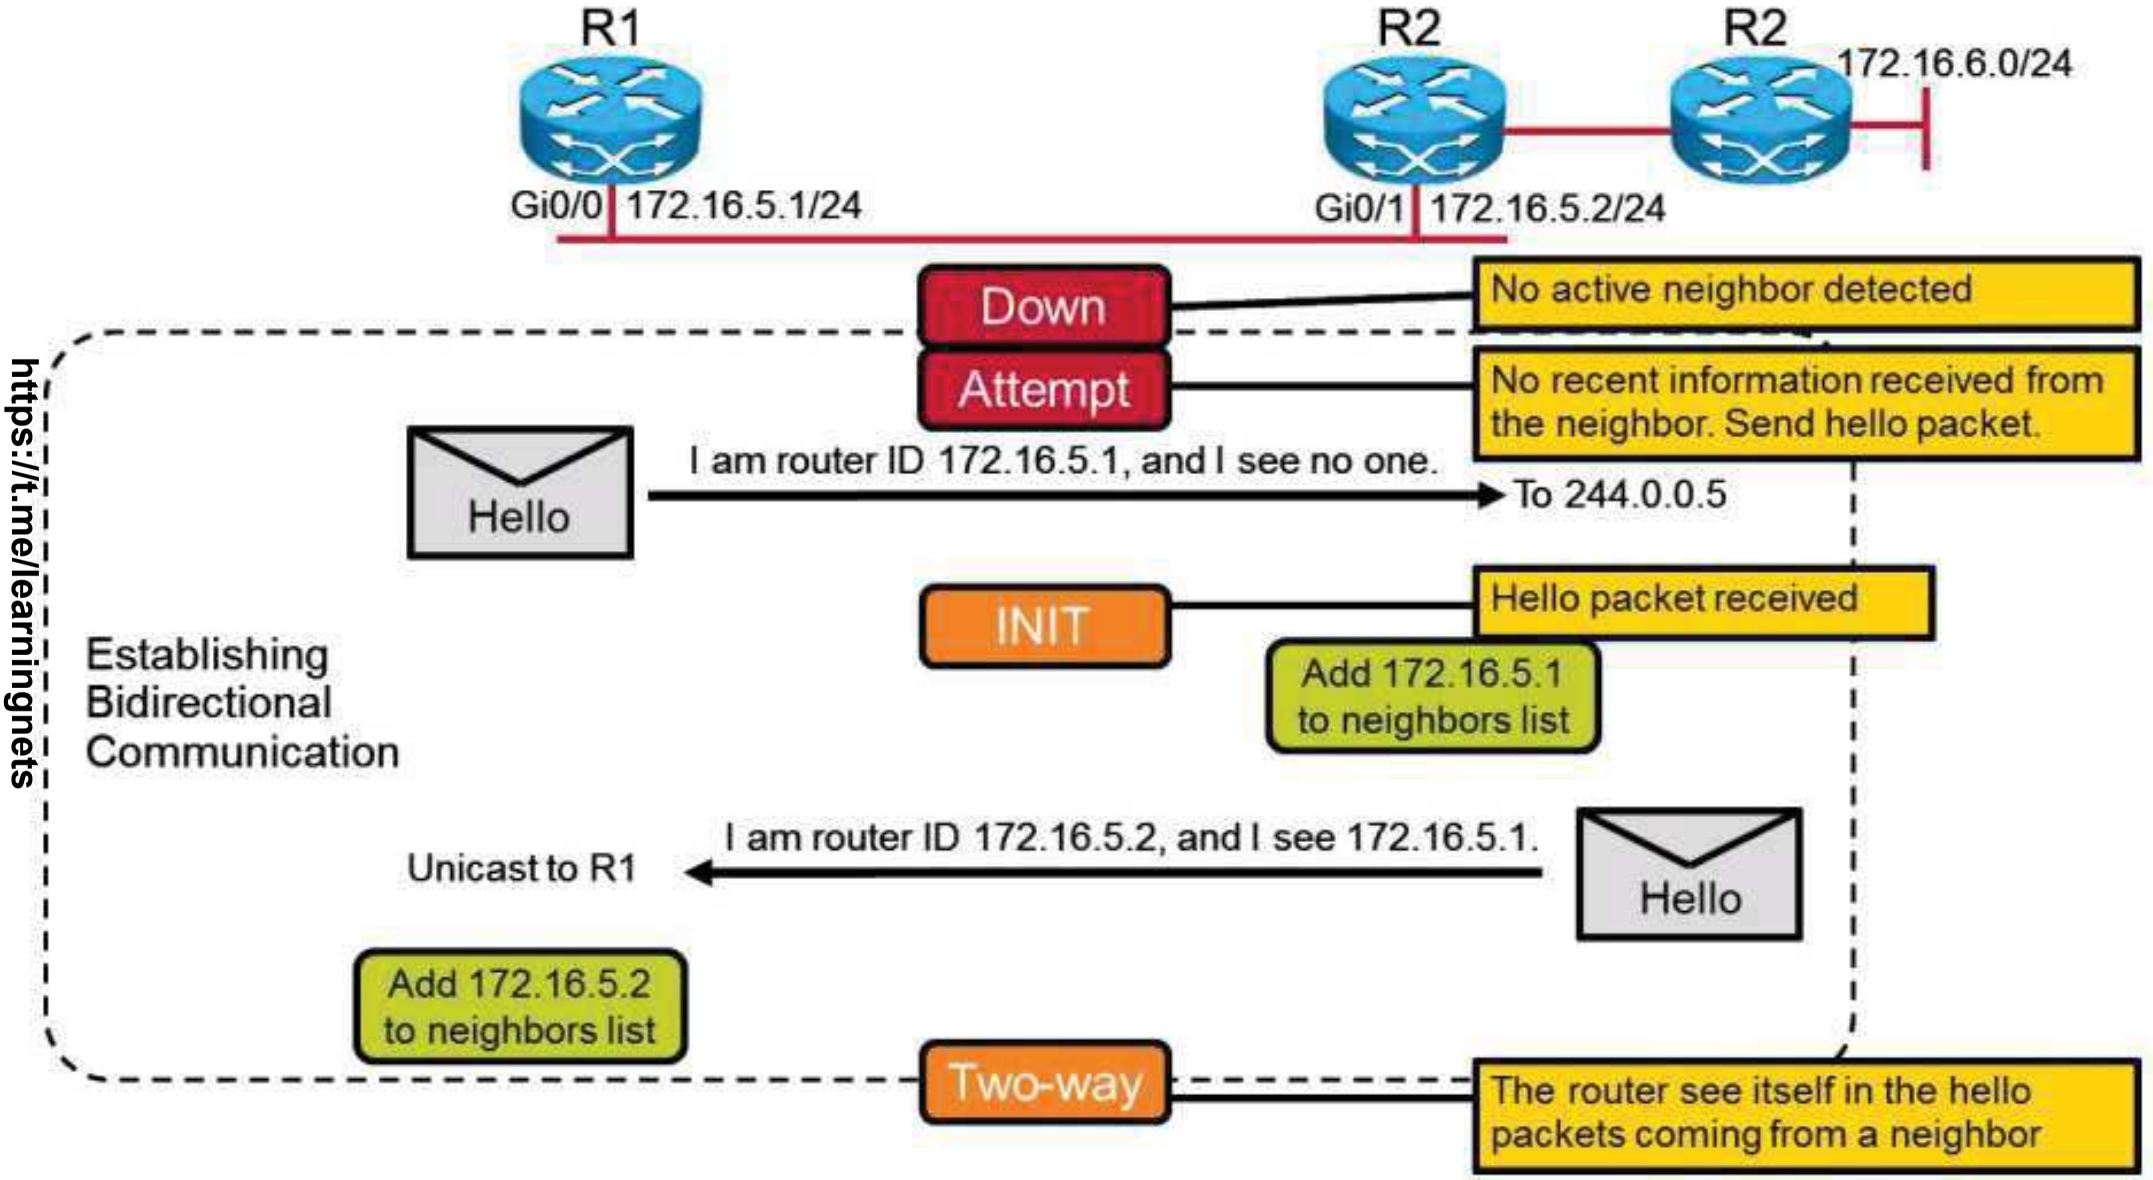

OSPF Neighbor States (Cont.)

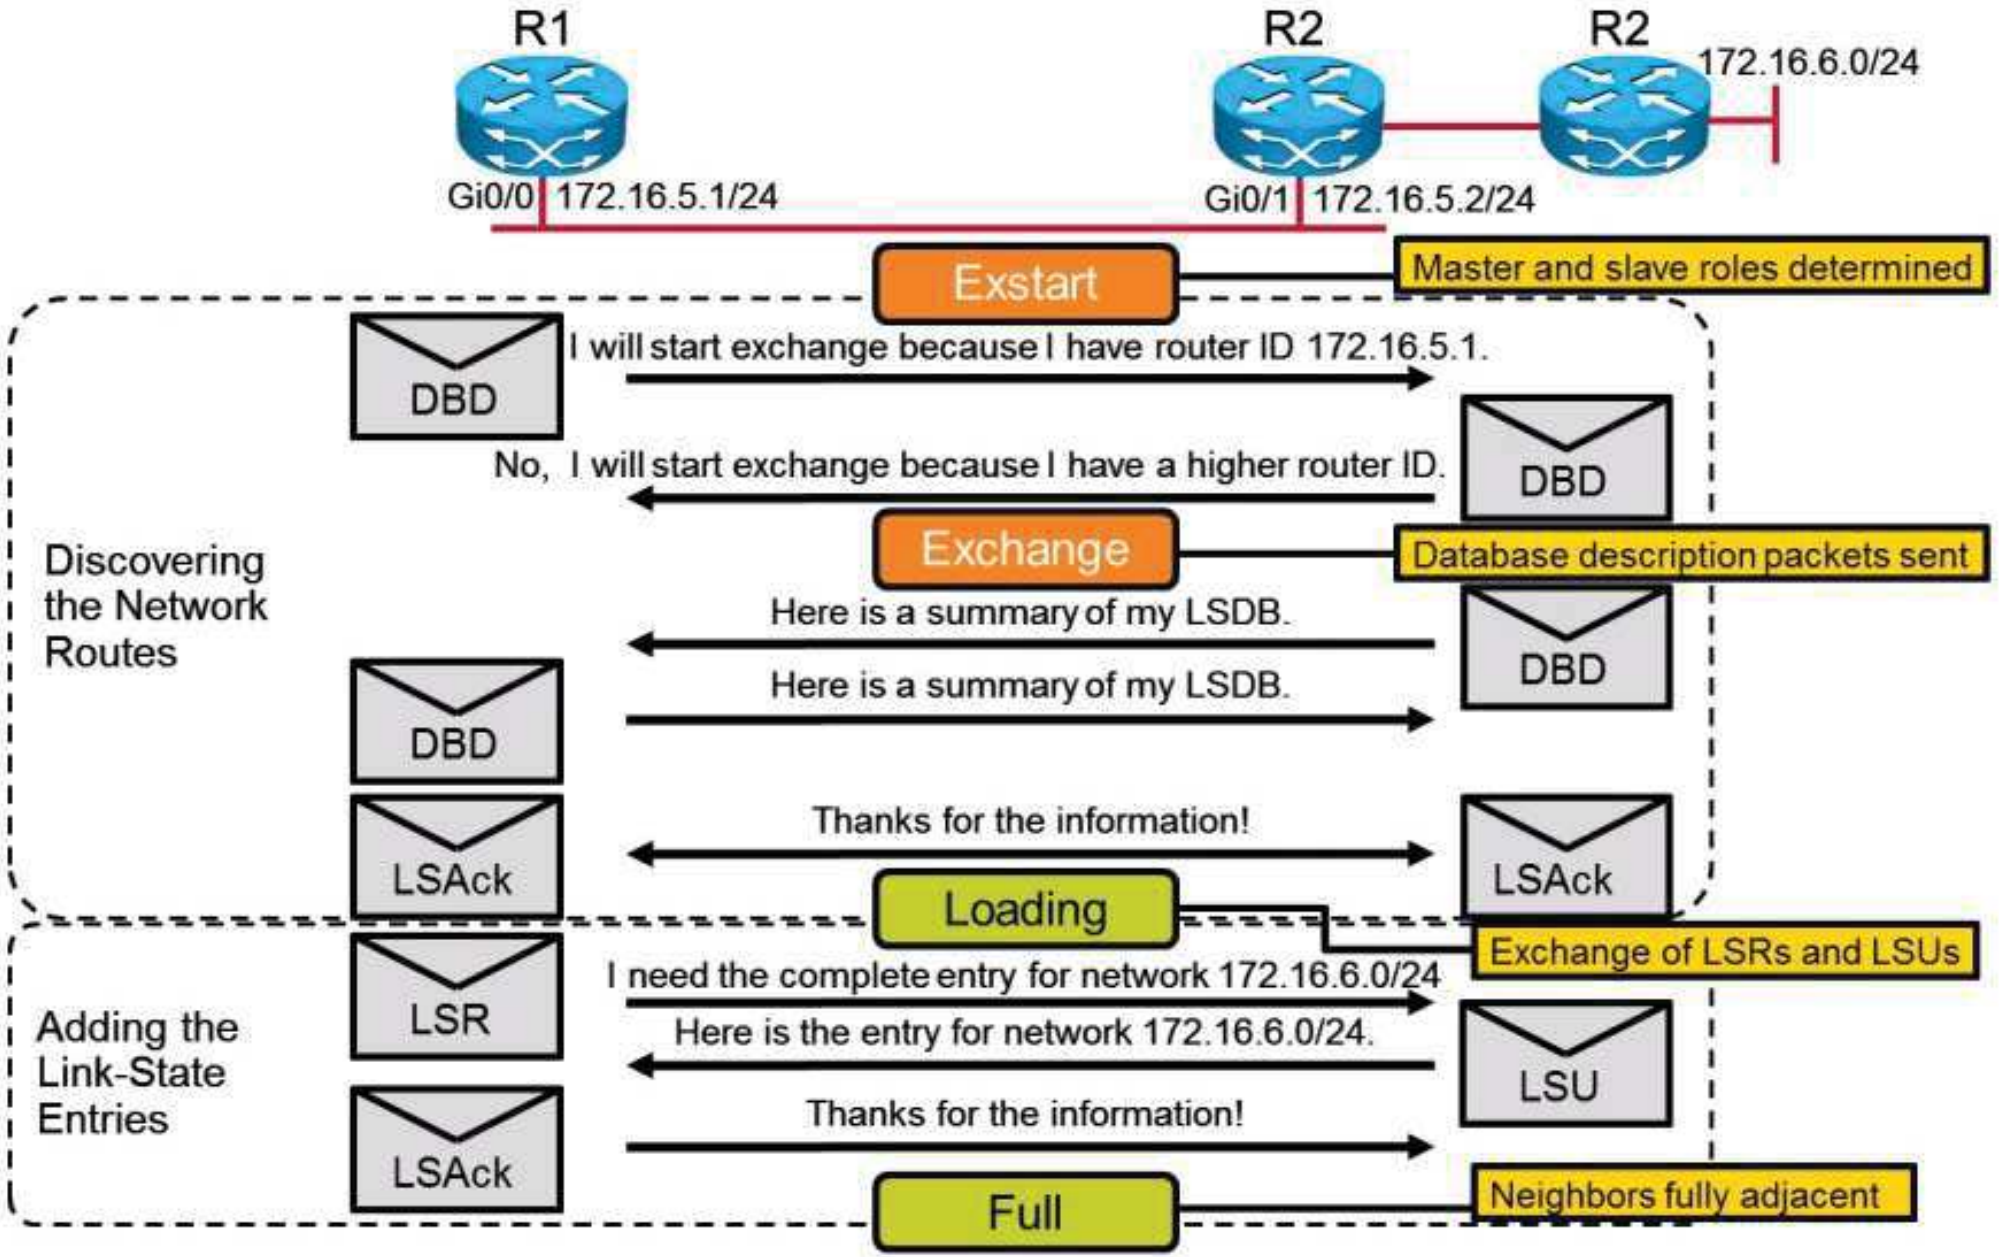

OSPF Link-State Flooding

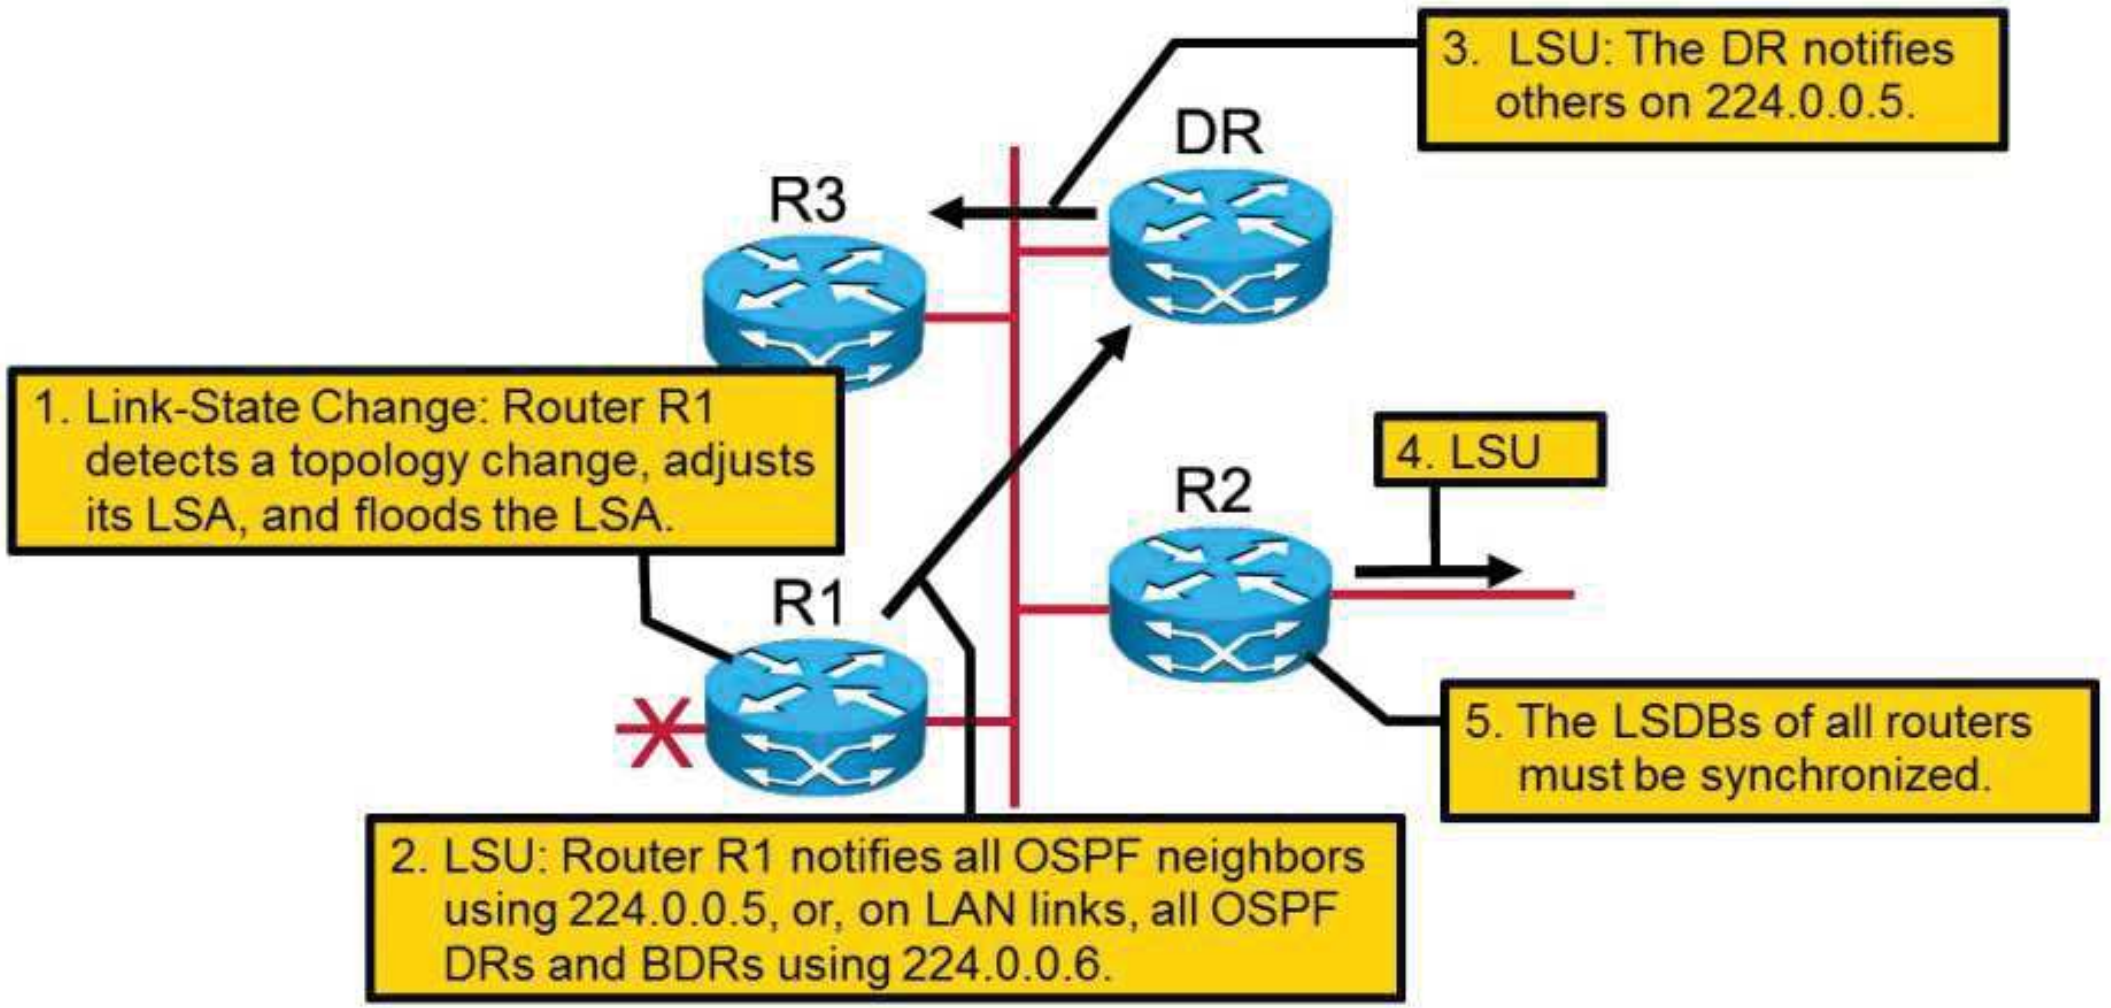

Debug OSPF Packets

Cisco IOS XR: **debug ospf 1 packet**

```
ospf[1010]: Recv: HLO l:48 rid:10.2.1.1 aut:0 auk: from 192.168.112.20 to 224.0.0.5 on
GigabitEthernet0/0/0/1, vrf default vrfid 0x60000000
```

```
OSPF: rcv. v:2 t:1 l:48 rid:10.1.1.1 aid:0.0.0.11 chk:8224 aut:0 auk: from GigabitEthernet0/0
```

Cisco IOS/IOS XE: **debug ip ospf packet**

Debug OSPF Packets (Cont.)

Field	Description
OSPF:	OSPF packet
Recv: / rcv.	was received
v:	Provides the version of OSPF
HLO / t:	Specifies the OSPF packet type: 1: hello; 2: DBD; 3: LSR; 4: LSU; 5: LSAck
l:	Specifies the OSPF packet length in bytes
rid:	Provides the OSPF router ID
aid:	Shows the OSPF area ID
chk:	Displays the OSPF checksum
Aut:	Provides the OSPF authentication type: 0: No authentication; 1: Simple password; 2: MD5
auk:	Specifies the OSPF authentication key, if used
keyid	Displays the MD5 key ID; only used for MD5 authentication
seq	Provides the sequence number; only used for MD5 authentication

OSPF Network Types

Point-to-point:
- A network that joins a single pair of routers
- Does not require DR or BDR election
- Neighbor is automatically detected
- Sends OSPF packets using multicast 224.0.0.5
- Default OSPF hello interval is 10 seconds
- Default OSPF dead interval is 40 seconds

Broadcast:
- DR and BDR selection is required
- Other routers form adjacencies with the DR and BDR only
- Packets to the DR and BDR use 224.0.0.6
- Packets from DR to all other routers use 224.0.0.5
- Default OSPF hello interval is 10 seconds
- Default OSPF dead interval is 40 seconds

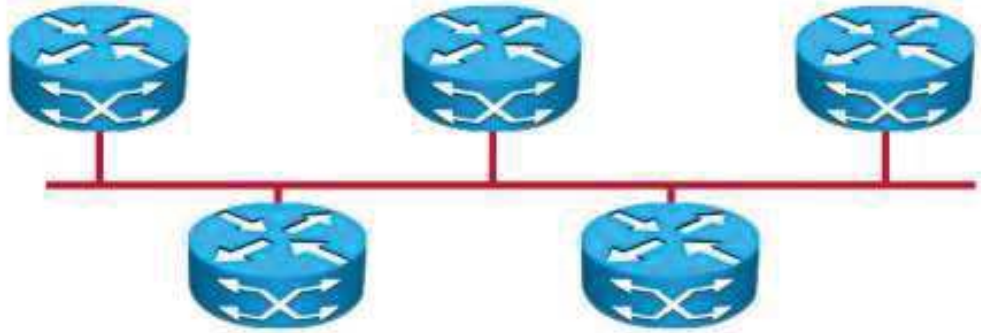

Nonbroadcast multiaccess (NBMA):
- Single interface interconnects multiple sites
- Without broadcasting capabilities
- Default OSPF hello interval is 30 seconds
- Default OSPF dead interval is 120 seconds

Point-to-multipoint:
- No DR or BDR election
- Automatic neighbor discovery
- Each adjacency treated as point-to-point

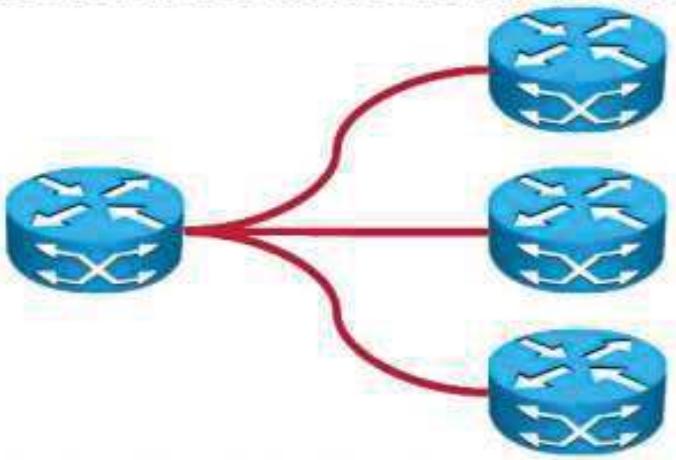

OSPF Network Types (Cont.)

Cisco IP NGN Infrastructure Layer:

- *Broadcast* OSPF network type is used in service provider environments and focuses on *IP edge* devices.
- *Point-to-point* OSPF network type is used in service provider environments and focuses on *core* devices.

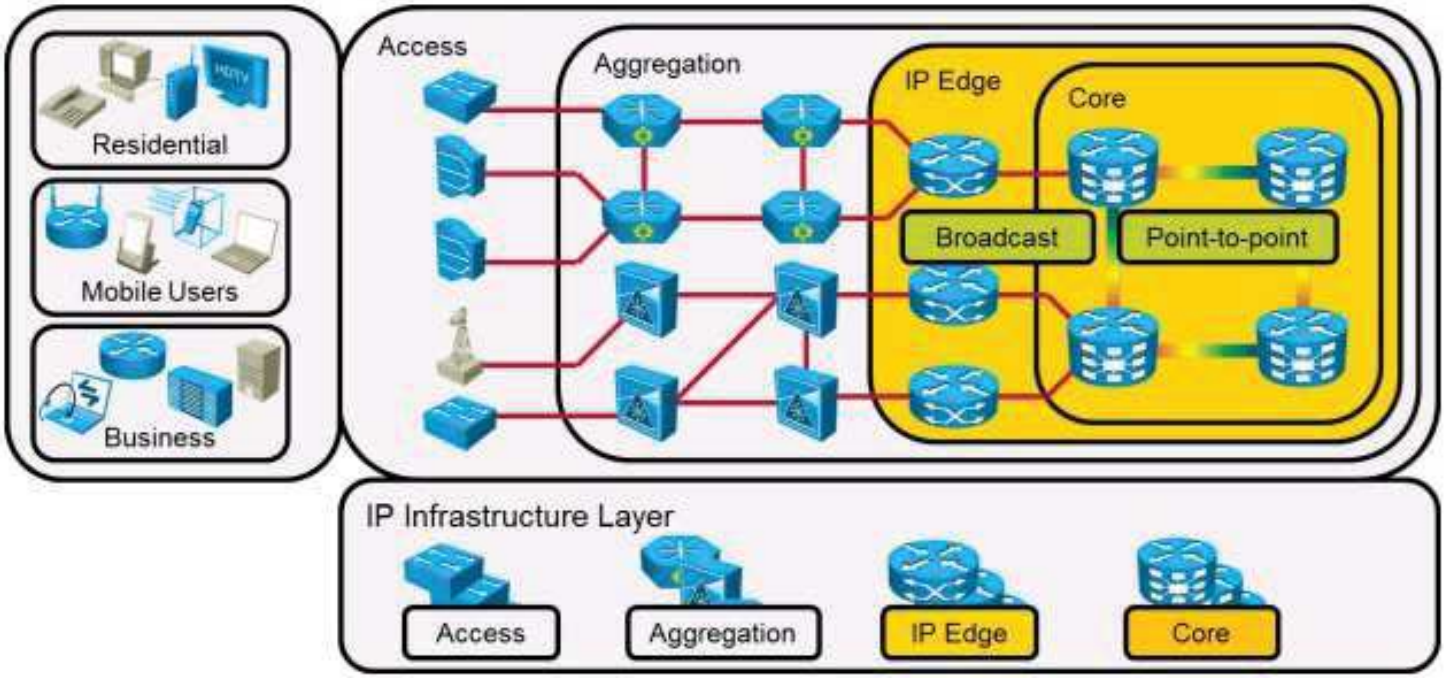

Electing the OSPF DR and BDR

DR and BDR Characteristics:

- Hello packets are exchanged via IP multicast.
- **DR:** The router with the highest OSPF priority
- **BDR:** The router with the second-highest priority value
- OSPF router ID is used as the tie-breaker.
- DR election is nonpre-emptive.

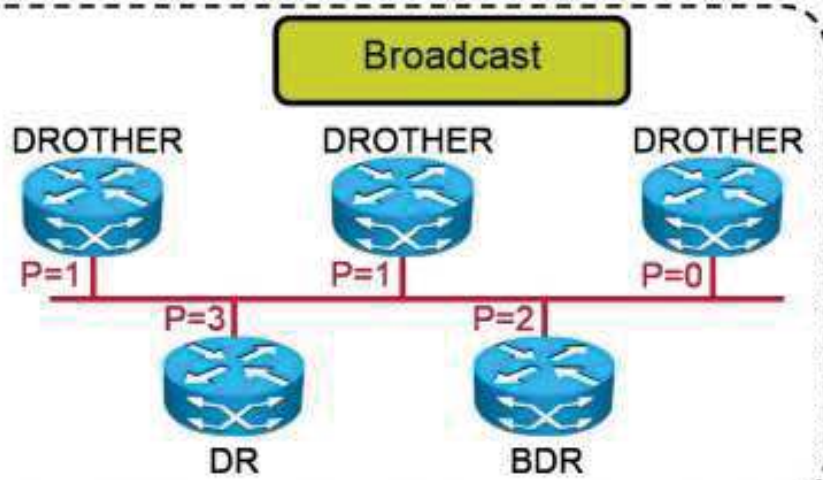

- OSPF considers NBMA to be like other broadcast media.
- DR and BDR need to have fully meshed connectivity with all other routers, but NBMA networks are not always fully meshed.
- DR and BDR each need a list of neighbors.
- OSPF neighbors are not automatically discovered by the router.

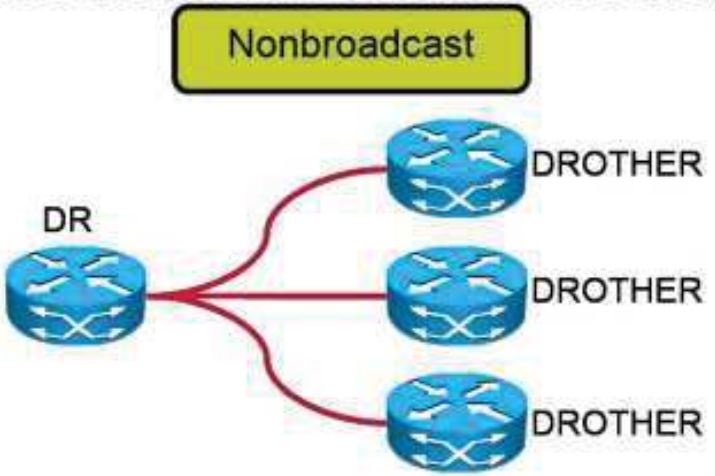

Electing the OSPF DR and BDR (Cont.)

Setting the priority:

- Different interfaces on a router may be assigned different values.
- The default priority is 1. The range is from 0 to 255.
- "0" means the router cannot be the DR or BDR.
- A router that is not the DR or BDR is DROTHER.
- The priority takes effect when the existing DR goes down.

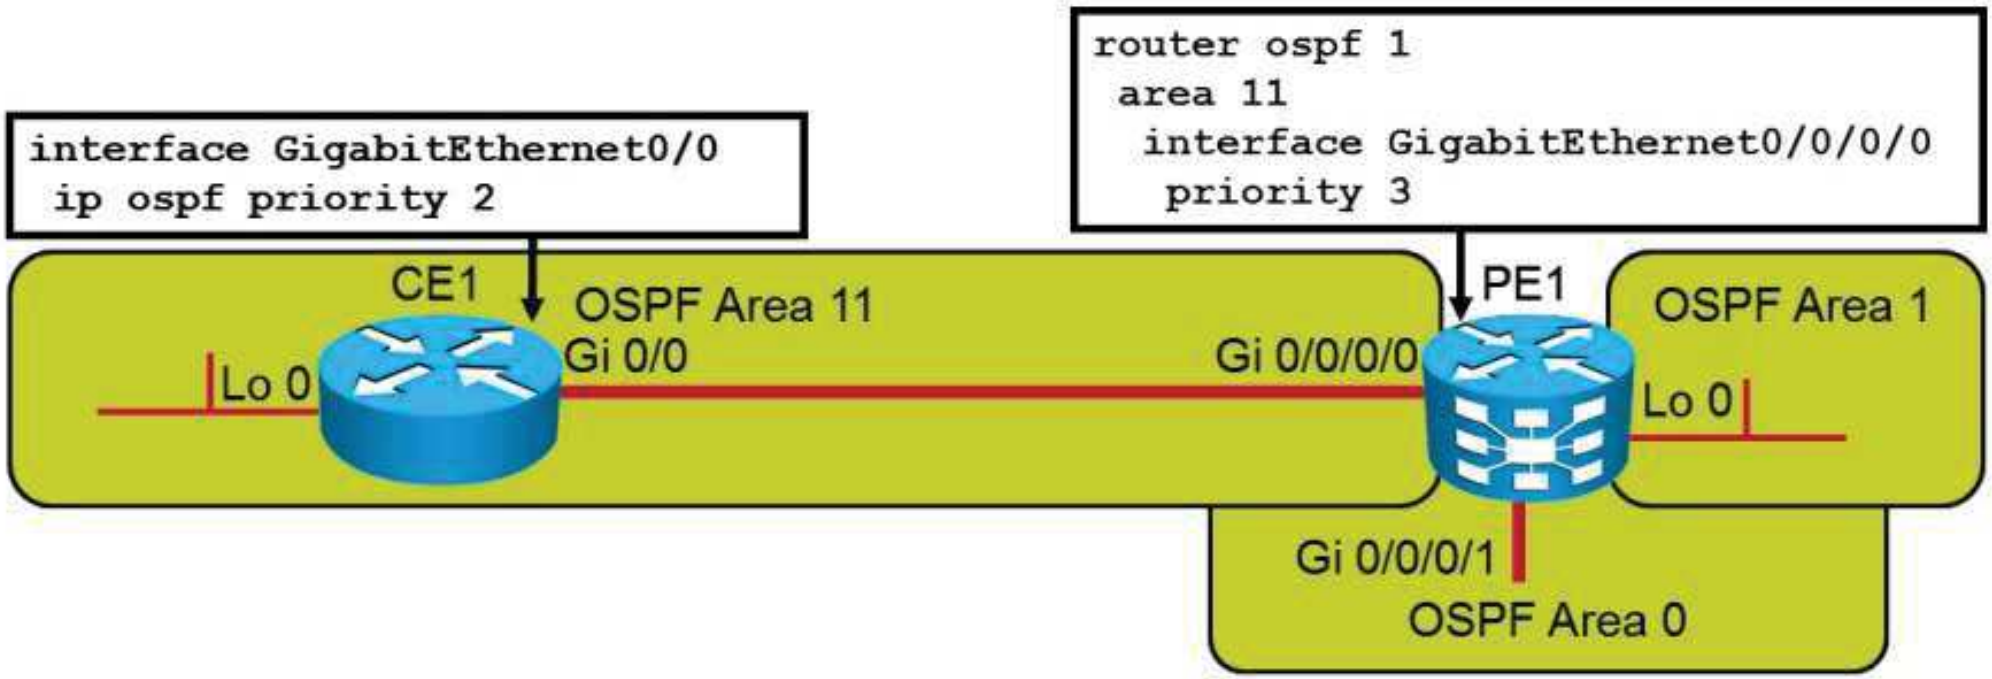

OSPF over NBMA Network Types

These are the possible network types for OSPF over NBMA:

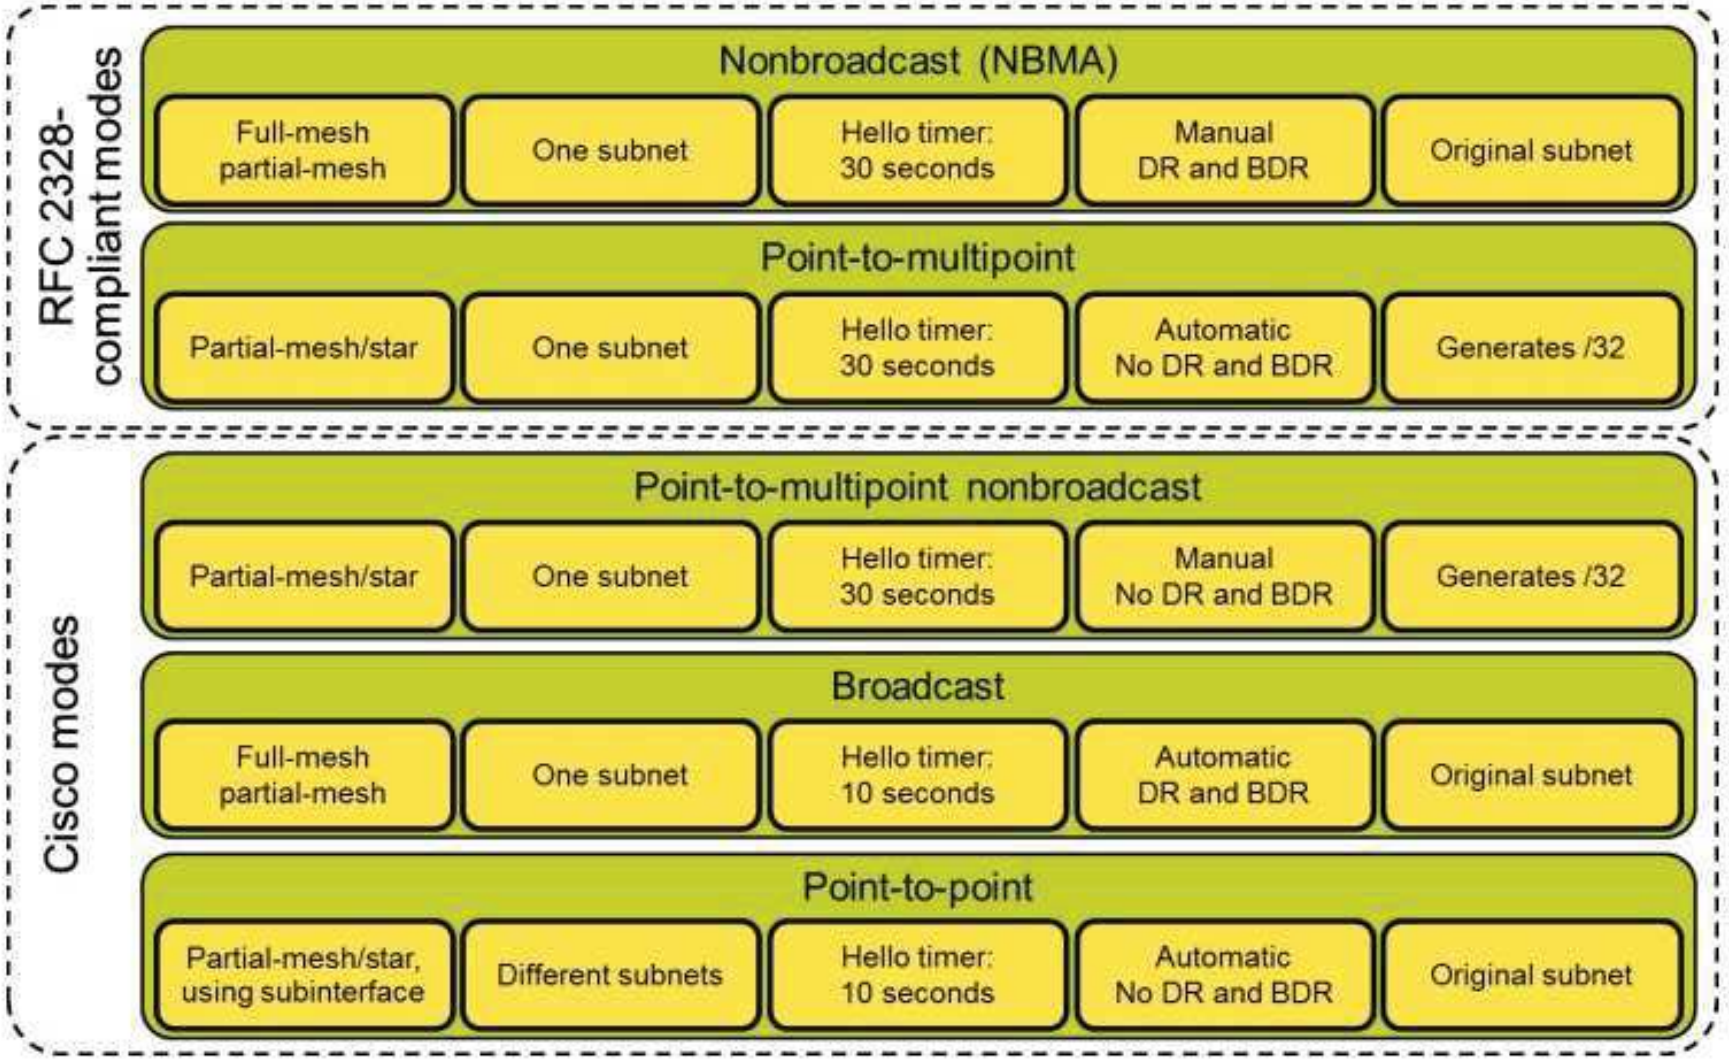

OSPF Adjacency over Metro Ethernet and EoMPLS

These are the characteristics of establishing OSPF adjacency over Metro Ethernet and EoMPLS:

- EoMPLS and Metro Ethernet service does not participate in STP, nor does it learn MAC addresses.
- Customer routers R1 and R2 exchange Ethernet frames via an interface or VLAN subinterfaces.
- OSPF behaves the same as on Ethernet.
  - OSPF network type = multiaccess broadcast network
  - DR and BDR are elected
  - Routers form full adjacencies with DR and BDR only

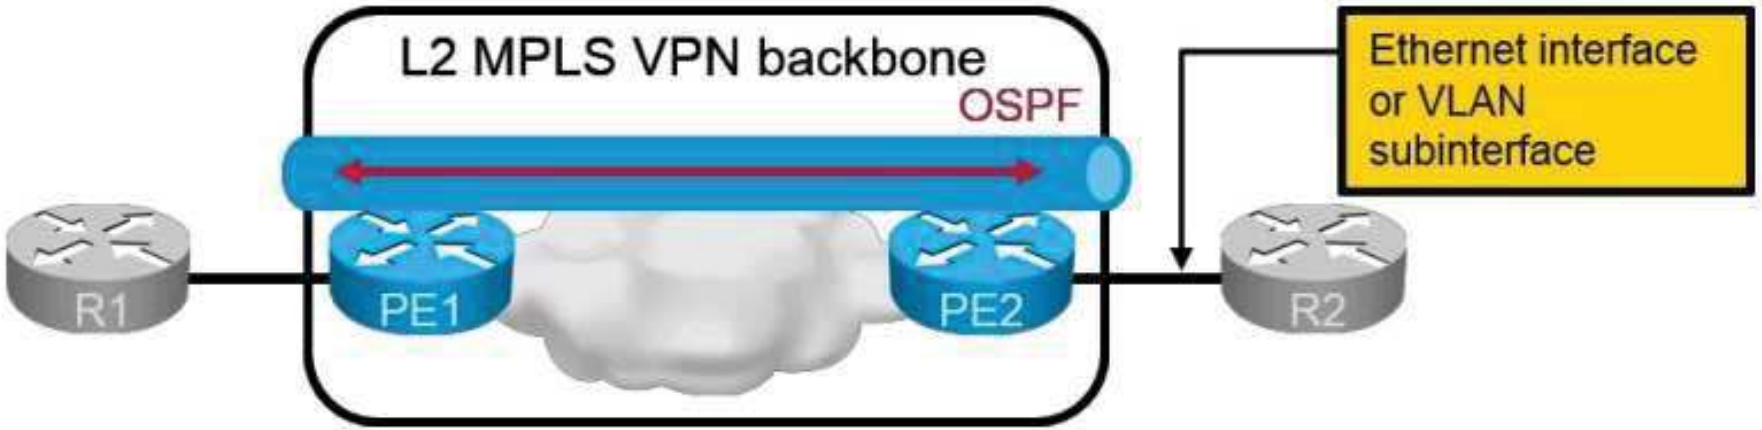

OSPF Adjacency over Layer 3 MPLS VPN

These are the characteristics of establishing OSPF adjacency over Layer 3 MPLS VPN:

- Customer routers run OSPF and exchange routing updates with the PE routers.
  - PE routers appear as another router in the customer's network.
  - Service provider routers are hidden from the customer.
  - Customer routers are unaware of MPLS VPN.
  - Customer and service provider must agree on OSPF parameters.
- Customer routers-to-PE connection can be of any type.
  - OSPF behaves per the connection type (point-to-point, broadcast, NBMA)

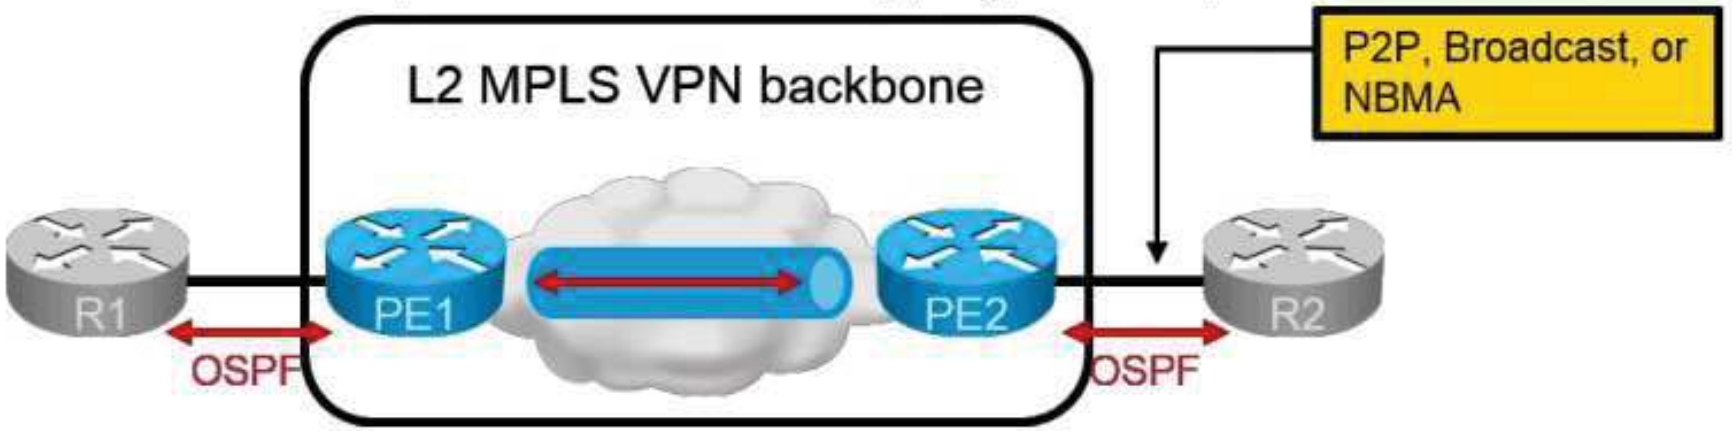

Enabling OSPF on a Link with NBMA Mode

```
interface type number
 ip ospf priority 0
 ip ospf network non-broadcast
!
router ospf 1
 neighbor 192.168.101.10
```

Define the OSPF network type.

Statically define a neighbor relationship.

```
router ospf 1
 area 11
  interface type number
   network non-broadcast
   neighbor 192.168.101.11 priority 0
   neighbor 192.168.101.21 priority 0
```

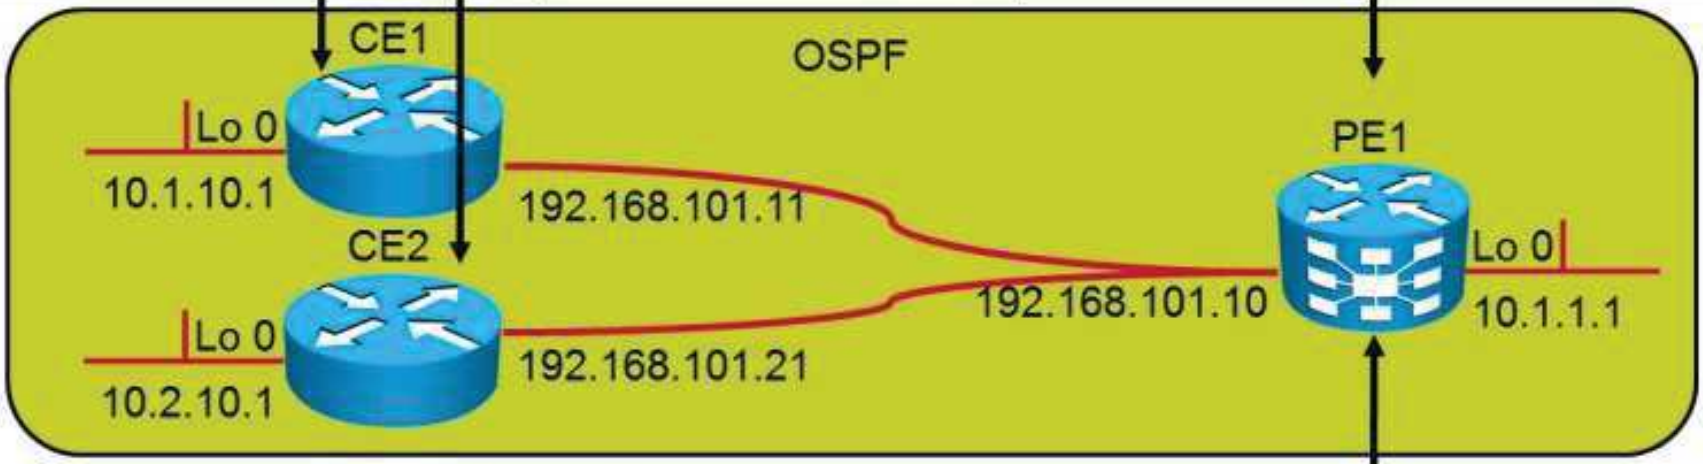


```
RP/0/RSP0/CPU0:PE1# show ospf neighbor
* Indicates MADJ interface
Neighbors for OSPF 1
Neighbor ID     Pri   State           Dead Time   Address         Interface
10.2.10.1       0     FULL/DROTHER    00:01:45    192.168.101.21  type number
    Neighbor is up for 00:15:41
10.1.10.1       0     FULL/DROTHER    00:01:56    192.168.101.11  type number
    Neighbor is up for 00:15:41

Total neighbor count: 2
```

Subinterfaces over Multiaccess WAN Networks

These are the main characteristics of subinterfaces over multiaccess WAN networks:

- Several logical subinterfaces can be created over all multiaccess WAN networks:

- **Point-to-point**

```
interface Serial0.1 point-to-point
```

- **Multipoint**

```
interface Serial0.2 multipoint
```

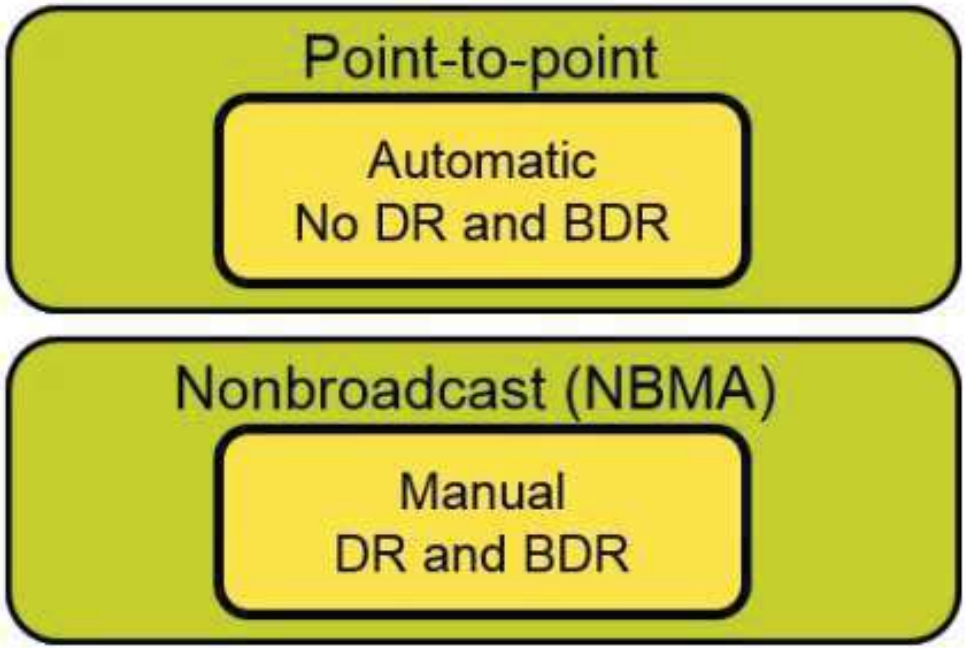

- Each subinterface requires an IP subnet.
- In routing, logical interfaces behave the same way as physical ones.
- Statistics and traffic-shaping behavior differs between interfaces and subinterfaces.

Enabling OSPF on a Link with Point-to-Point Mode

Leased-line emulation

Single subnet

Automatic
No DR and BDR

```
router ospf 1
 area 11
  interface GigabitEthernet0/0/0/0
   network point-to-point
```

```
interface GigabitEthernet0/0
 ip ospf network point-to-point
```

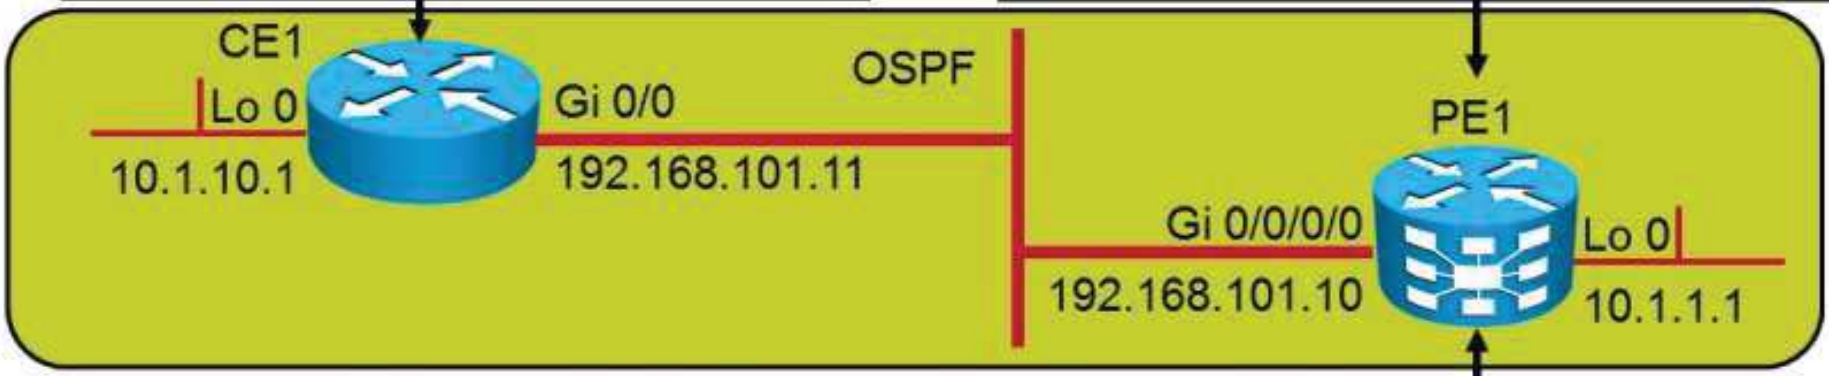


```
RP/0/RSP0/CPU0:PE1# show ospf interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
  Internet Address 192.168.101.10/24, Area 11
  Process ID 1, Router ID 10.1.1.1, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT, MTU 1500, MaxPktSz 1500
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:03
  Index 1/2, flood queue length 0
  <... output omitted ...>
```

Enabling OSPF on a Link with Point-to-Multipoint Mode

Partial-mesh/star

Single subnet

Automatic
No DR and BDR

```
router ospf 1
 area 11
  interface type number
   network point-to-multipoint
```

```
interface type number
 ip ospf network point-to-multipoint
```

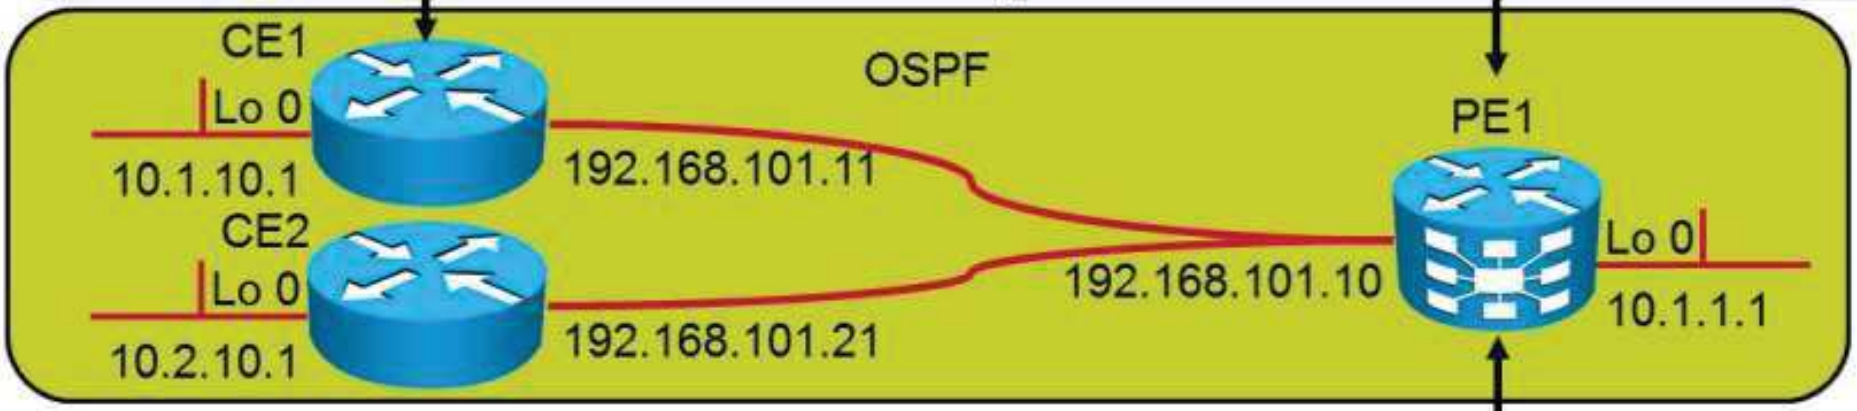


```
RP/0/RSP0/CPU0:PE1# show ospf interface type number
type number is up, line protocol is up
  Internet Address 192.168.101.10/24, Area 11
  Process ID 1, Router ID 10.1.1.1, Network Type POINT_TO_MULTIPOINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT, MTU 1500, MaxPktSz 1500
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
  Index 1/2, flood queue length 0
<... output omitted ...>
```

Enabling OSPF on a Link with Point-to-Multipoint Nonbroadcast Mode

Partial-mesh/star

Single subnet

Manual
No DR and BDR

```
interface type number
 ip ospf network point-to-multipoint non-broadcast
```

```
router ospf 1
 area 11
  interface type number
   network point-to-multipoint non-broadcast
   neighbor 192.168.101.11
   neighbor 192.168.101.21
```


```
RP/0/RSP0/CPU0:PE1# show ospf neighbor
Neighbors for OSPF 1
Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.10.1       0     FULL/  -        00:01:49    192.168.101.11 type number
    Neighbor is up for 00:06:18
10.2.10.1       1     FULL/  -        00:01:41    192.168.101.21 type number
    Neighbor is up for 00:06:32
Total neighbor count: 2
```

IPv6 Support for OSPF Modes

```
CE1(config-if)# ipv6 ospf network ?
  broadcast              Specify OSPF broadcast multi-access network
  non-broadcast          Specify OSPF NBMA network
  point-to-multipoint    Specify OSPF point-to-multipoint network
  point-to-point         Specify OSPF point-to-point network
```

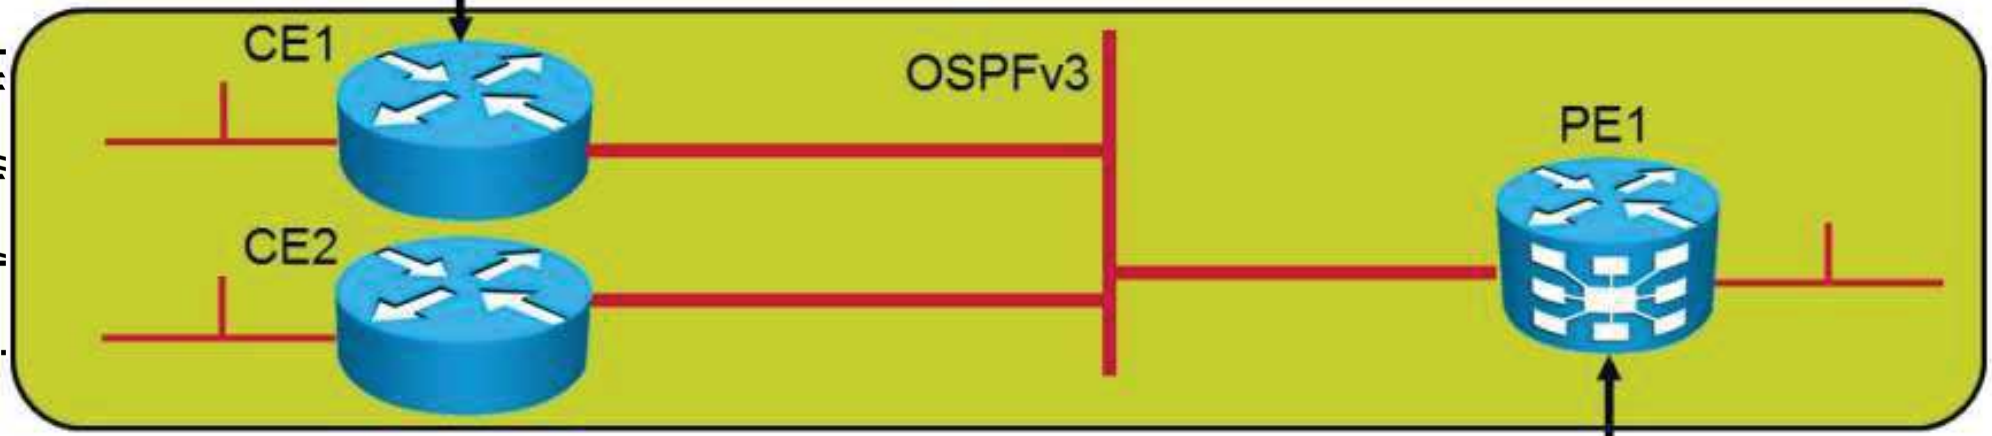


```
RP/0/RSP0/CPU0:PE1(config)# router ospfv3 1
RP/0/RSP0/CPU0:PE1(config-ospfv3)# area 11
RP/0/RSP0/CPU0:PE1(config-ospfv3-ar)# interface type number
RP/0/RSP0/CPU0:PE1(config-ospfv3-ar-if)# network ?
  broadcast              Specify OSPFv3 broadcast multi-access network
  non-broadcast          Specify OSPFv3 NBMA network
  point-to-multipoint    Specify OSPFv3 point-to-multipoint network
  point-to-point         Specify OSPFv3 point-to-point network
```

Summary

- OSPF uses several types of packets to operate properly.
- OSPF packets are encapsulated directly into an IP payload.
- There are five OSPF packet types: hello, DBD, LSU, LSR, and LSAck.
- The Hello protocol forms logical neighbor adjacency relationships.
- A topology change triggers link-state flooding.
- You can use the **debug ospf packets** Cisco IOS XR command to debug OSPF packets.
- Cisco routers support two RFC-compliant OSPF network types: NBMA and point-to-multipoint; and three Cisco defined OSPF network types: point-to-multipoint nonbroadcast, broadcast, and point-to-point.
- A router with the highest priority becomes a DR. In case of a tie, a router with the highest router ID becomes a DR.

Summary (Cont.)

- It is important that adjacent routers agree on the media type. OSPF network type selection depends on the topology, available subnets, and multicast or broadcast supported.
- OSPF over Metro Ethernet and EoMPLS behaves the same as on the Ethernet.
- OSPF over MPLS VPN behaves as per the connection type between a CE and PE router.
- Both adjacent routers must agree on a the media type. Otherwise, adjacencies may not form due to different hello intervals.

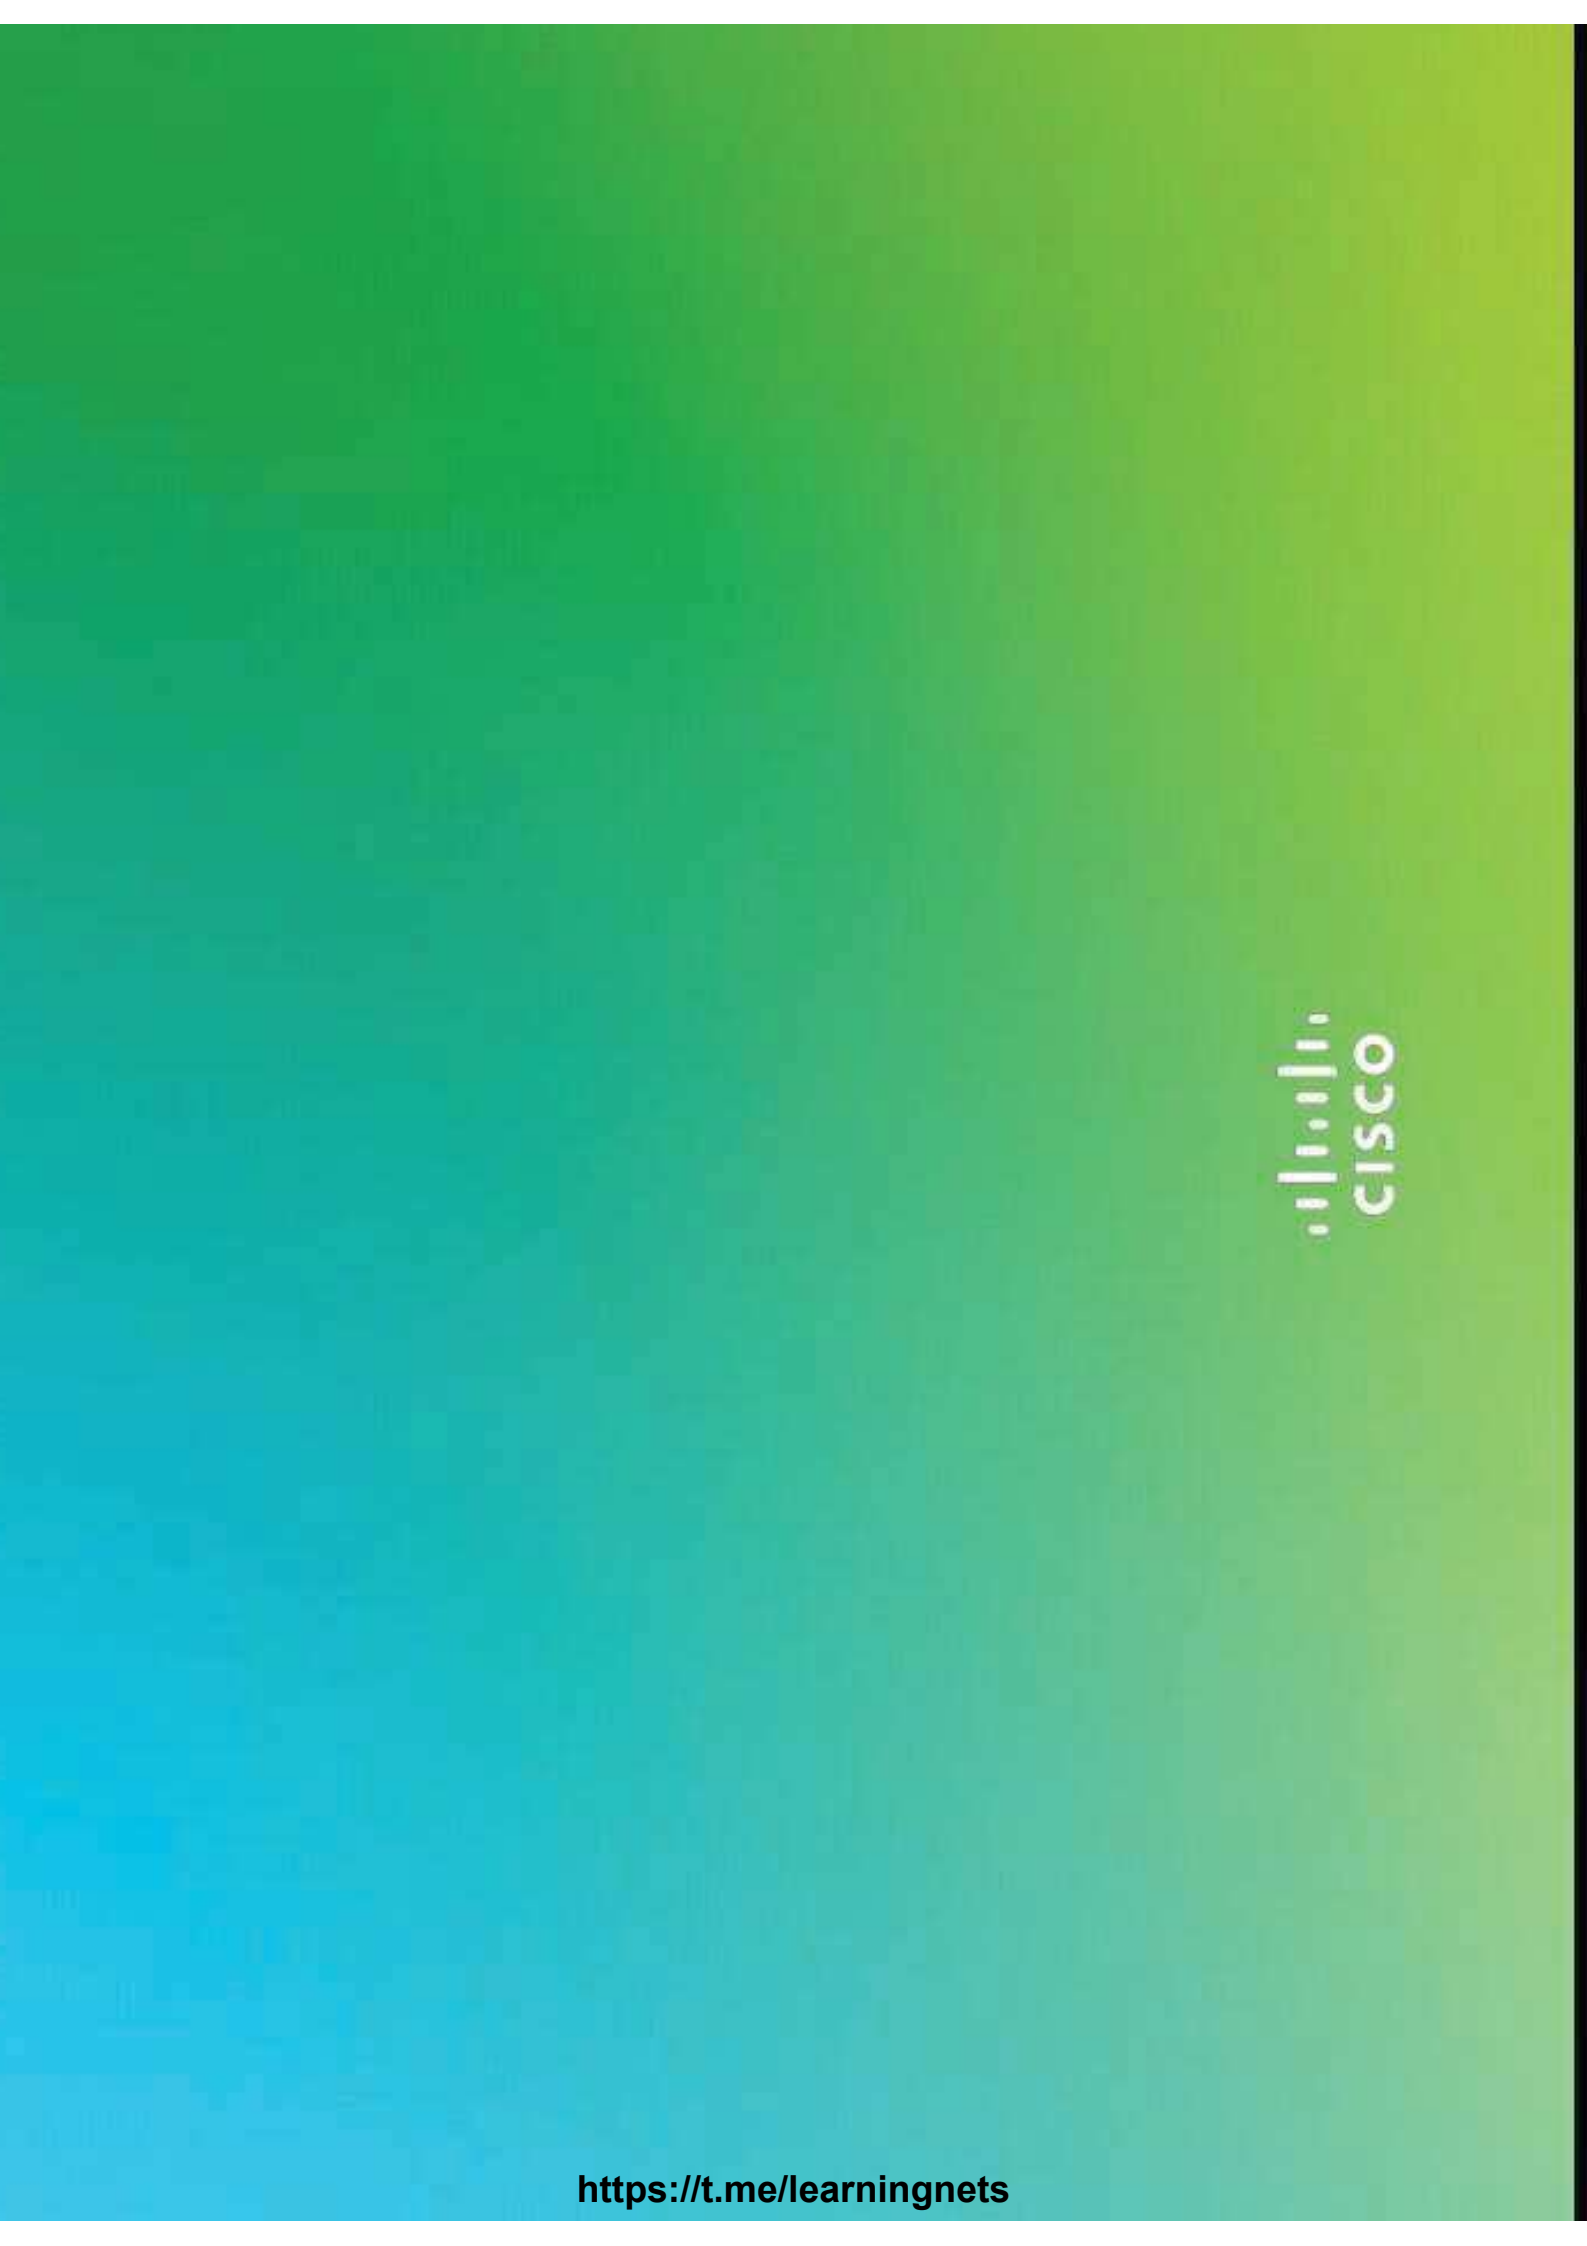

Implementing OSPF Routing

Implement OSPF in the Service Provider Network

Implement OSPF

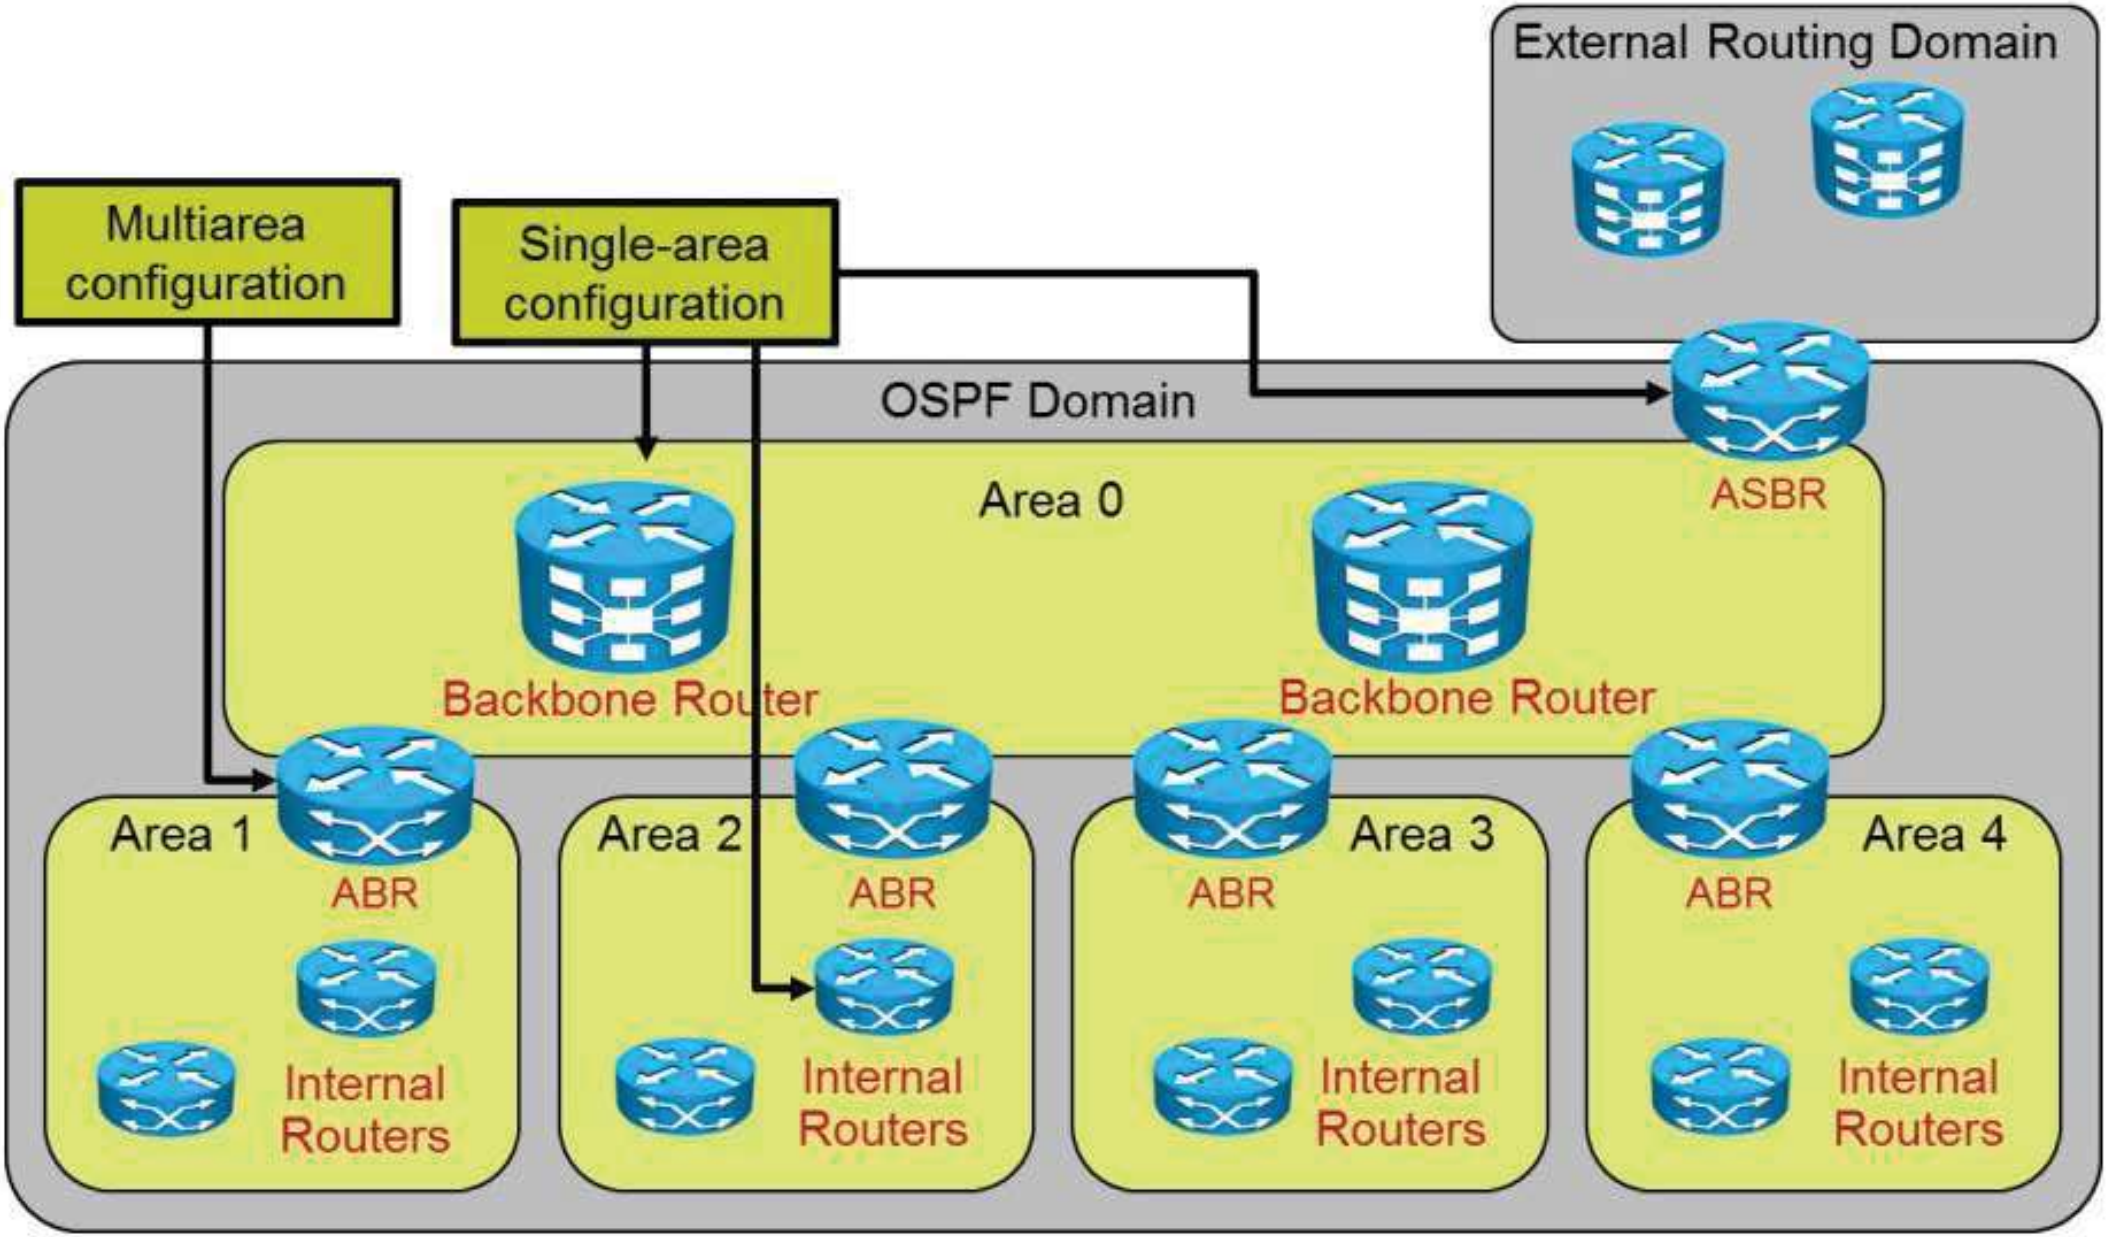

Implement OSPF (Cont.)

Steps to enable OSPF:

- Create an implementation plan:
  - IP addressing
  - Areas and area types
  - ABRs and ASBRs
- Define summarization and redistribution points.
- Configure OSPF routing processes on every OSPF router.

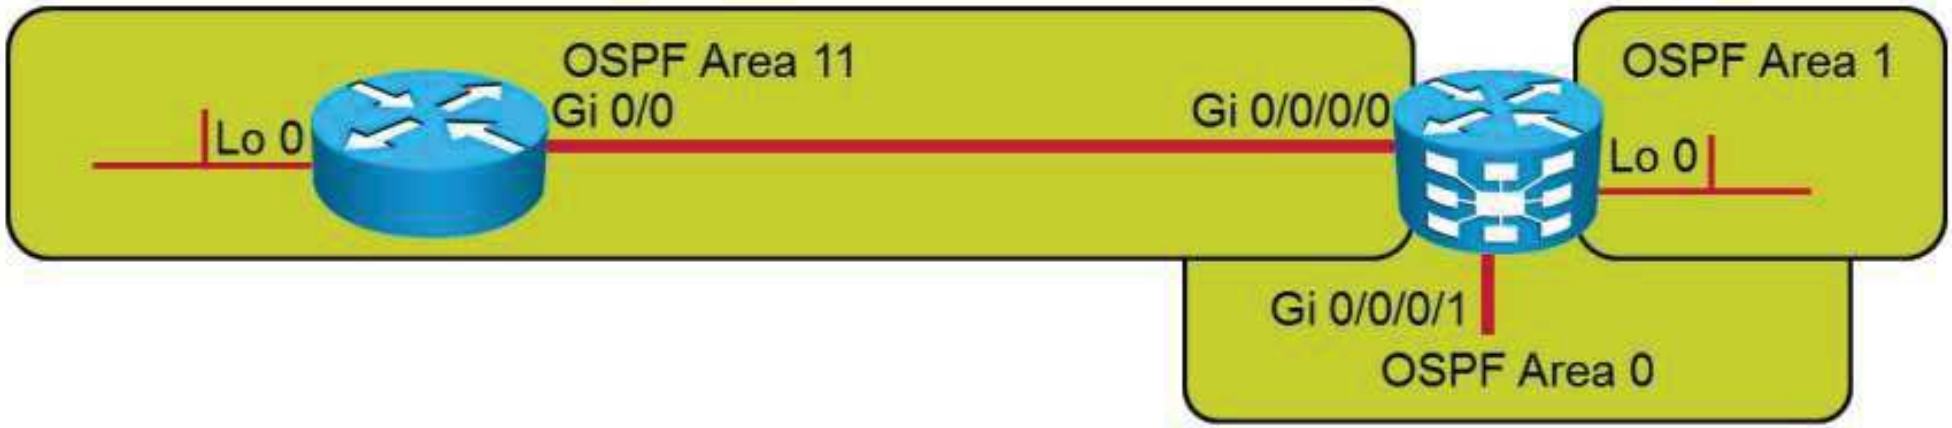

OSPF for Multiple Areas

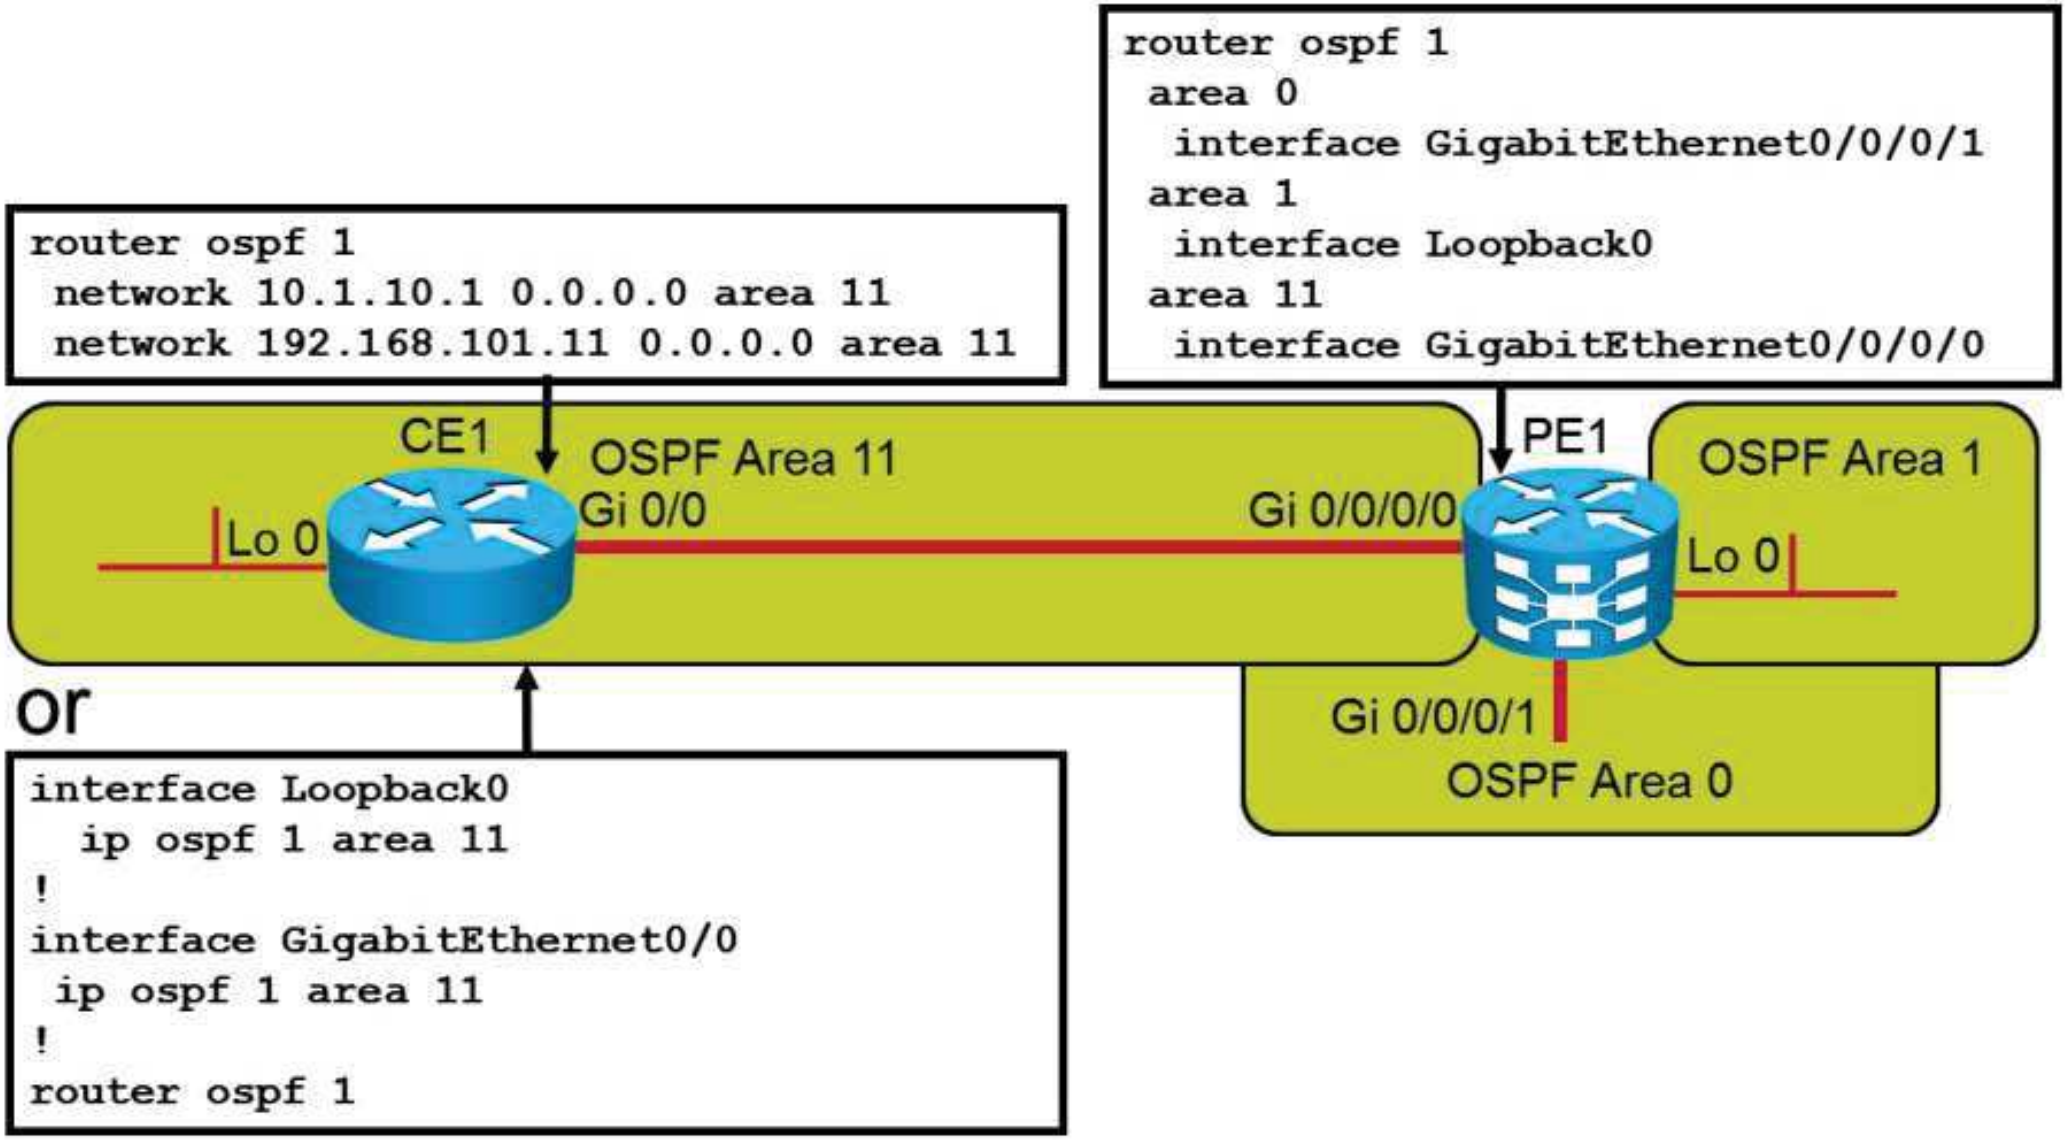

OSPF Router ID

OSPF router ID overview:

- The router is known to OSPF by the router ID number.
- This router ID is used in LSDBs to differentiate routers.
- OSPF requires at least one active interface with an IP address.
- By default, the router ID is:
  - The highest IP address on an active interface at the moment of OSPF process startup.
  - If a loopback interface exists, the router ID is the highest IP address on any active loopback interface. A loopback interface overrides the OSPF router ID.
- The OSPF **router-id** command can be used to override the default OSPF router ID selection process.
- Using a loopback interface or a **router-id** command is recommended for stability.
- OSPFv3 still uses a 32-bit number, written in four octets.

OSPF Router ID (Cont.)

Setting OSPF router ID:

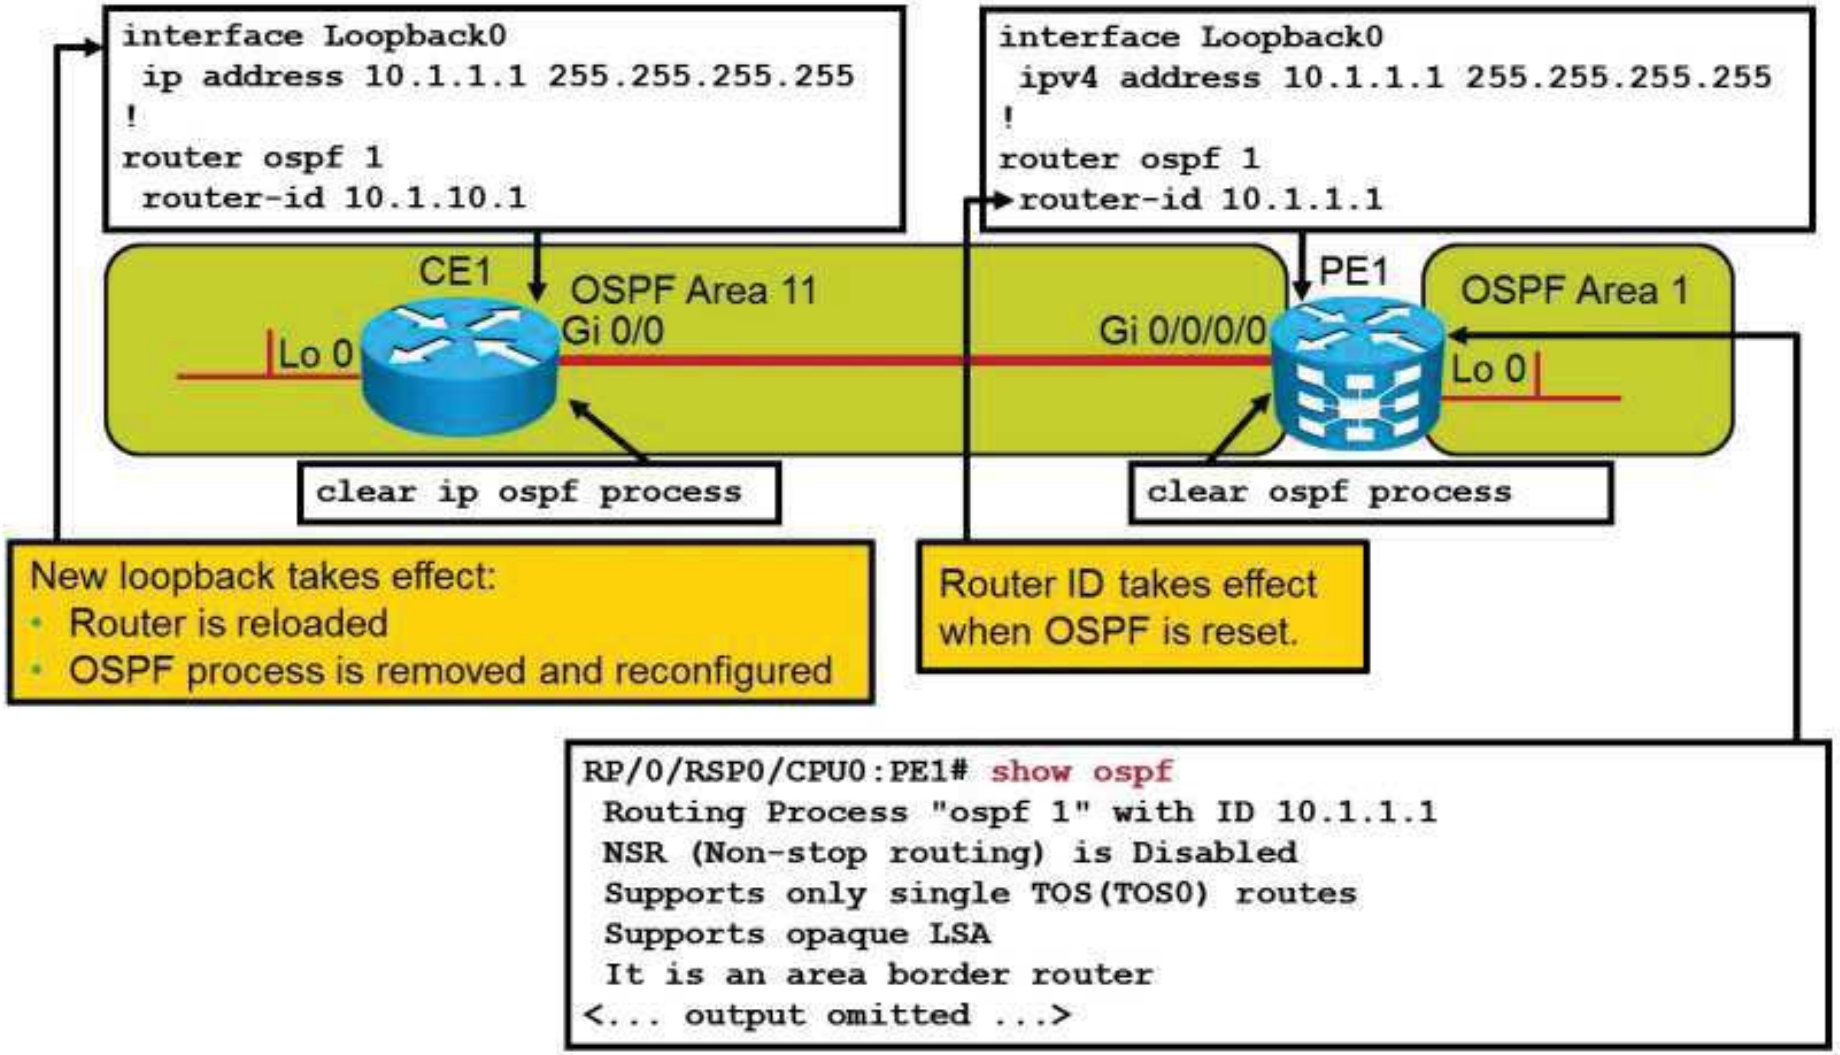

OSPF Passive Interface

These are the characteristics of an OSPF passive interface:

- The sending and receiving of routing updates is disabled.
- The specified interface address appears as a stub network in the OSPF domain.

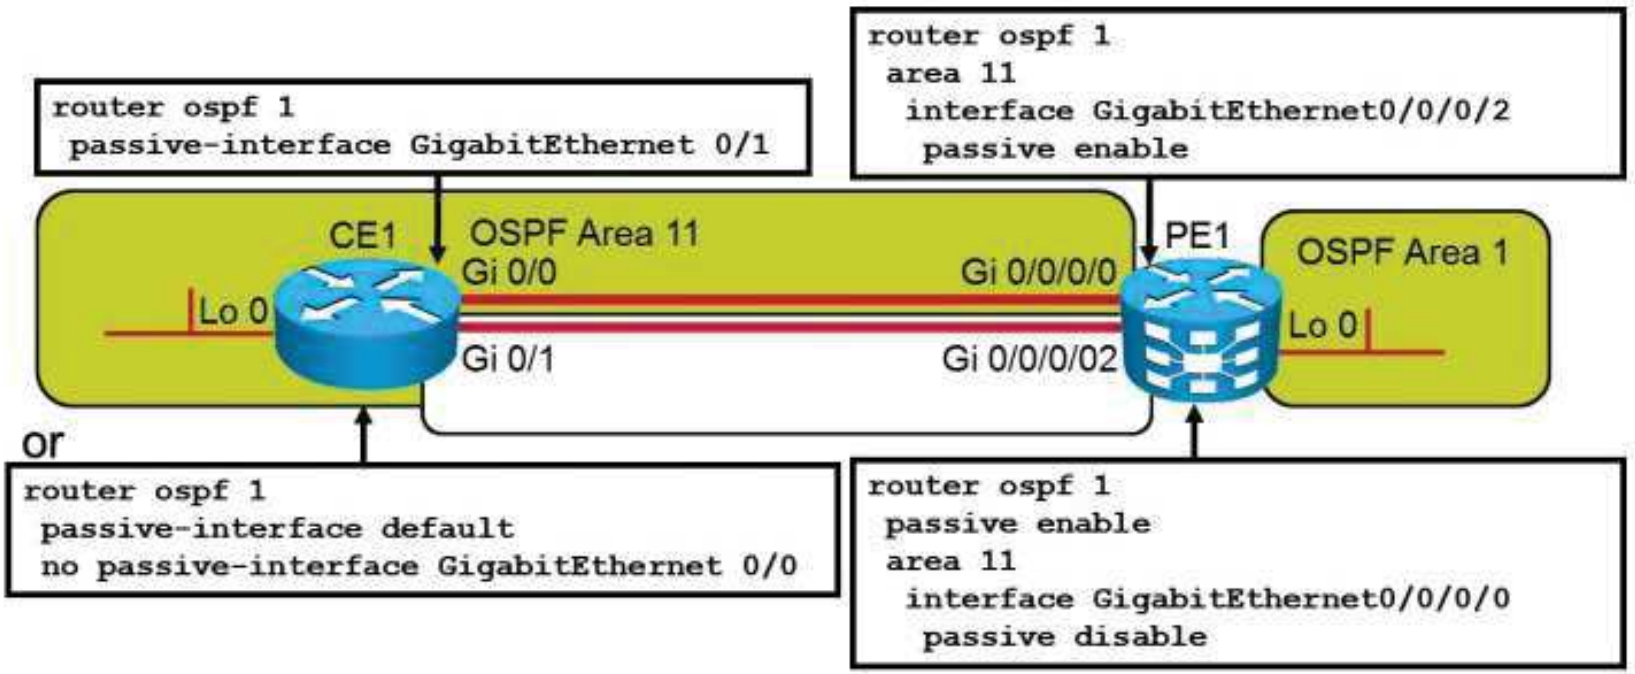

Verify Basic OSPF

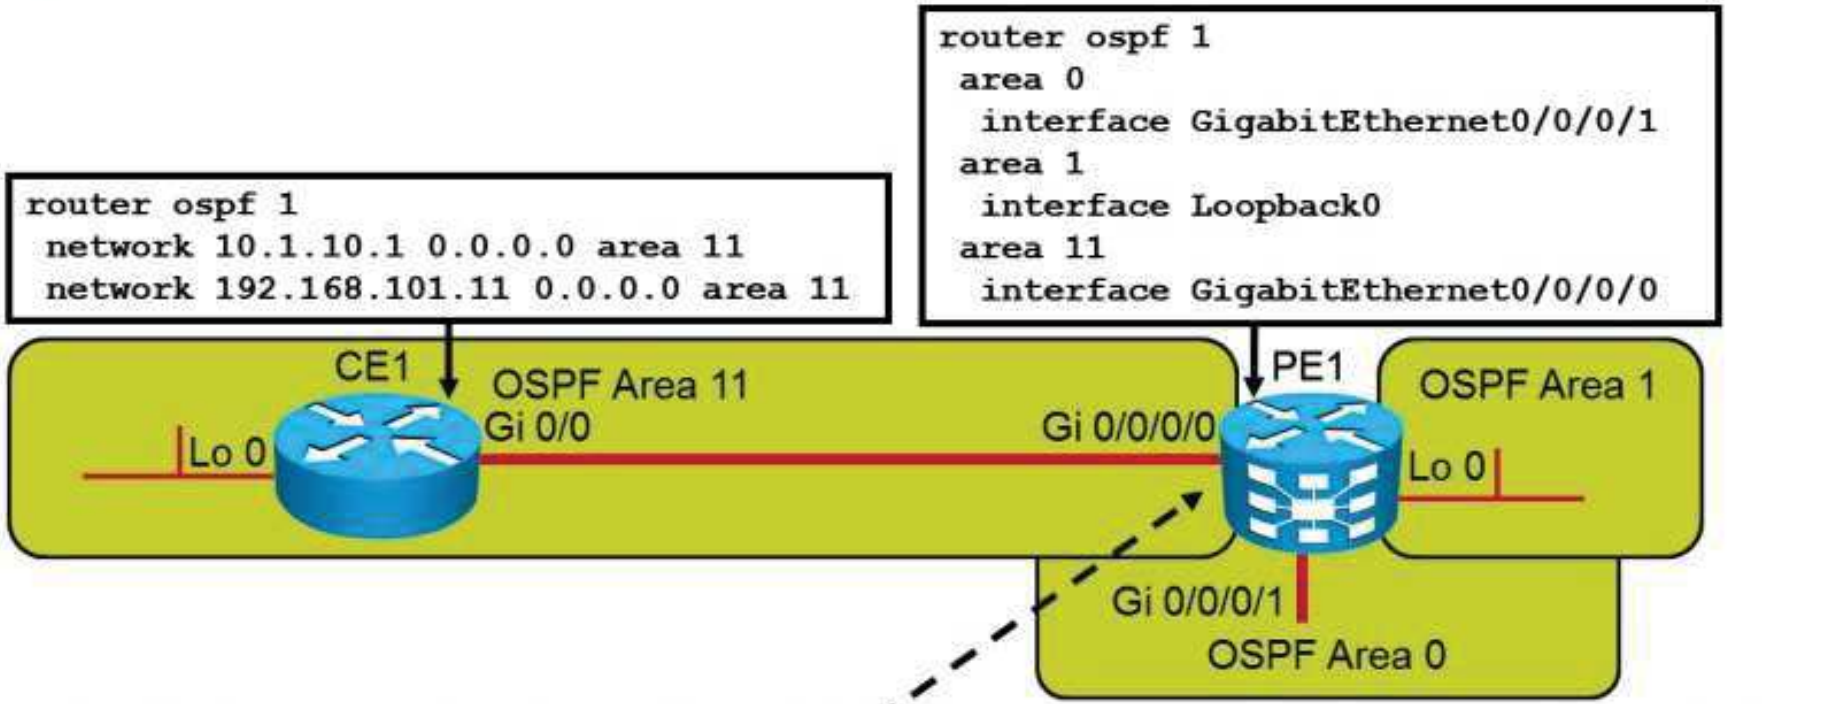

`show ospf`	Verify OSPF routing protocol.
`show ospf interface`	Verify OSPF interface information.
`show ospf neighbor`	Verify OSPF neighbors.
`show route ospf`	Verify OSPF routes learned by the router in the IP routing table.
`show protocols`	Verify configured IP routing protocol processes.
`show ospf database`	Verify OSPF LSDB.

Verify OSPF Routing Protocol

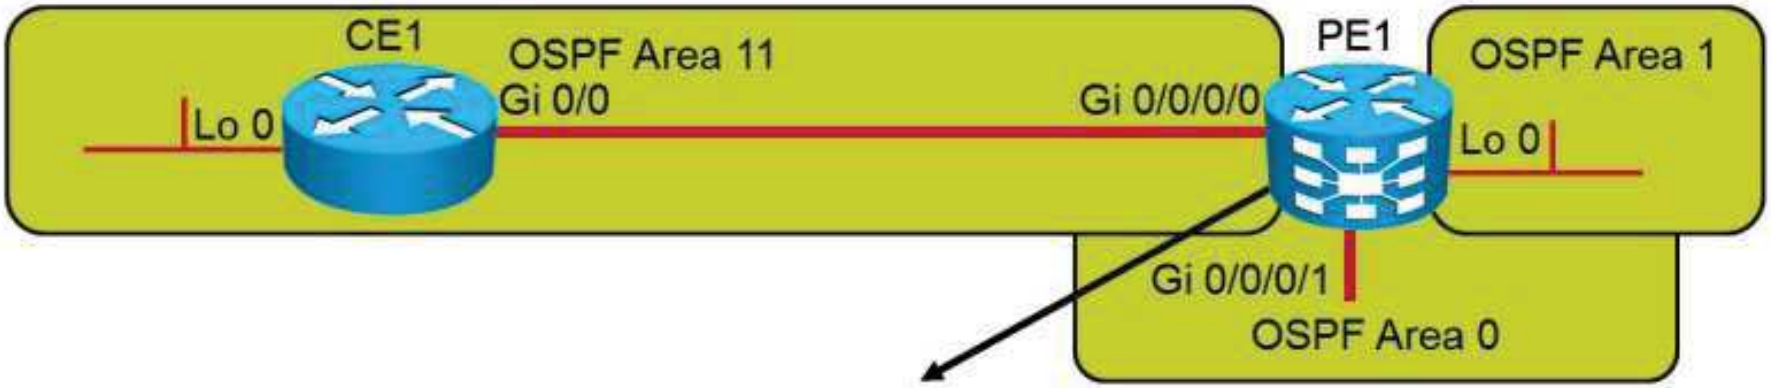


```
RP/0/RSP0/CPU0:PE1# show ospf
 Routing Process "ospf 1" with ID 10.1.1.1
 NSR (Non-stop routing) is Disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 Router is not originating router-LSAs with maximum metric
 Initial SPF schedule delay 50 msecs
 Minimum hold time between two consecutive SPFs 200 msecs
 Maximum wait time between two consecutive SPFs 5000 msecs
 Initial LSA throttle delay 50 msecs
 Minimum hold time for LSA throttle 200 msecs
 Maximum wait time for LSA throttle 5000 msecs
 Minimum LSA interval 200 msecs. Minimum LSA arrival 100 msecs
 LSA refresh interval 1800 seconds
 Flood pacing interval 33 msecs. Retransmission pacing interval 66 msecs
 Adjacency stagger enabled; initial (per area): 2, maximum: 64
    Number of neighbors forming: 0, 2 full
<... output omitted ...>
```

OSPF router ID

OSPF timers

OSPF statistics

Verify OSPF Interface

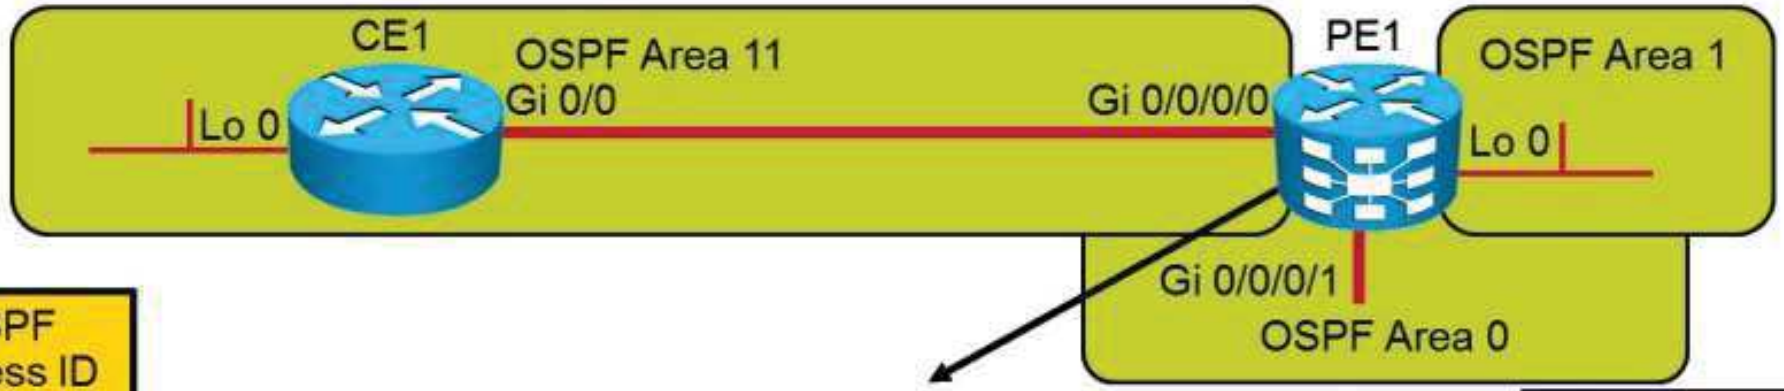

OSPF process ID

OSPF router ID

```
RP/0/RSP0/CPU0:PE1# show ospf interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
  Internet Address 192.168.101.10/24, Area 11
  Process ID 1, Router ID 10.1.1.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1, MTU 1500, MaxPktSz 1500
  Designated Router (ID) 10.1.10.1, Interface address 192.168.101.11
  Backup Designated router (ID) 10.1.1.1, Interface address 192.168.101.10
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Index 1/2, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 1, maximum is 6
  Last flood scan time is 0 msec, maximum is 0 msec
  LS Ack List: current length 0, high water mark 18
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.10.1  (Designated Router)
  Suppress hello for 0 neighbor(s)
  Multi-area interface Count is 0
```

OSPF area

OSPF network type

OSPF timers

OSPF neighbors

Verify OSPF Neighbors

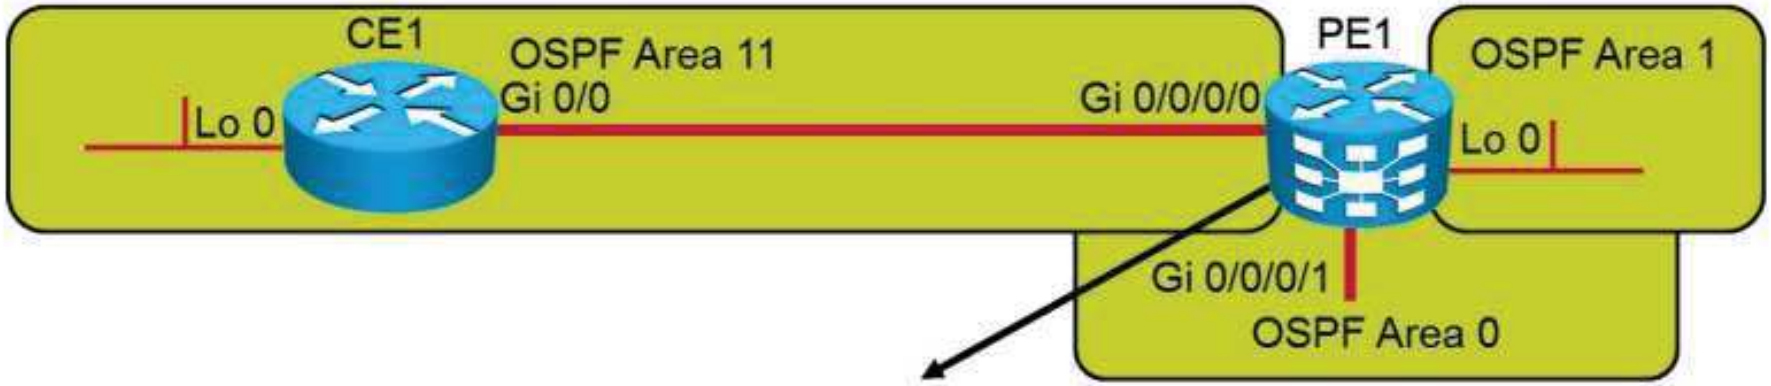


```
RP/0/RSP0/CPU0:PE1# show ospf neighbor
* Indicates MADJ interface
Neighbors for OSPF 1
Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.10.1       1     FULL/DR         00:00:32    192.168.101.11  GigabitEthernet0/0/0/0
    Neighbor is up for 00:41:42
Total neighbor count: 1
```

```
RP/0/RSP0/CPU0:PE1# show ospf neighbor detail
<... output omitted ...>
 Neighbor 10.1.10.1, interface address 192.168.101.11
    In the area 11 via interface GigabitEthernet0/0/0/0
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 192.168.101.11 BDR is 192.168.101.10
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:33
    Neighbor is up for 00:42:09
<... output omitted ...>
```

Verify Routes and Protocols

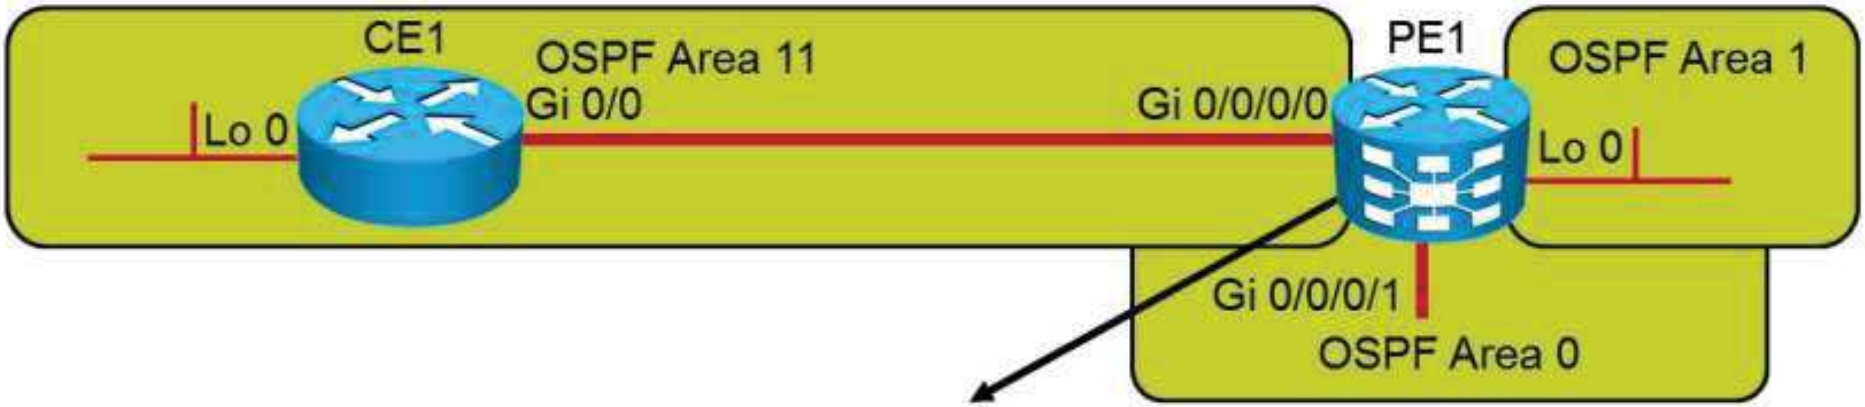


```
RP/0/RSP0/CPU0:PE1# show route ospf
O    10.1.10.1/32 [110/2] via 192.168.101.11, 01:23:03, GigabitEthernet0/0/0/0
O IA 192.168.102.0/24 [110/2] via 192.168.112.20, 00:57:43, GigabitEthernet0/0/0/1
```

```
RP/0/RSP0/CPU0:PE1# show protocols
Routing Protocol OSPF 1
  Router Id: 10.1.1.1
  Distance: 110
  Non-Stop Forwarding: Disabled
  Redistribution:
    None
  Area 0
    GigabitEthernet0/0/0/1
  Area 1
    Loopback0
  Area 11
    GigabitEthernet0/0/0/0
```

OSPF Virtual Links

Design limitations of OSPF:

- If more than one area is configured, one of the areas must be Area 0, the backbone area.
- All areas must be connected to Area 0; Area 0 must be contiguous.

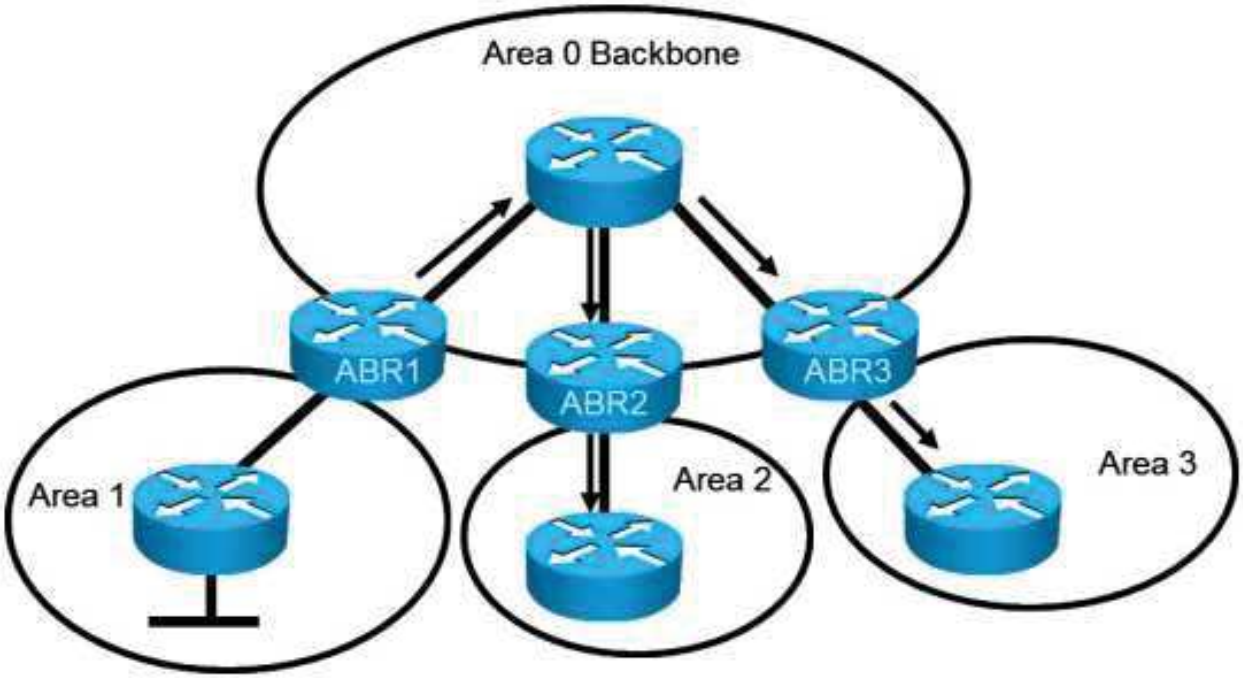

OSPF Virtual Links (Cont.)

Virtual links as a solution:

- An extension to the backbone.
- Carried by a nonbackbone area.
- Cannot be created across a stub or NSSA area, or over unnumbered links.
- Virtual links are used for these purposes:
  - Allow areas to connect to the backbone through a nonbackbone area.
  - Repair a discontiguous Area 0 (for example, if two companies merge and have separate backbone areas).

OSPF Virtual Links (Cont.)

Example: The role of a virtual link when there is no direct physical connection to Area 0:

- Area 20 is added with no physical access to Area 0.
- A virtual link provides a logical path to the backbone area.
- The OSPF database treats the link between routers ABR1 and ABR2 as a direct link.

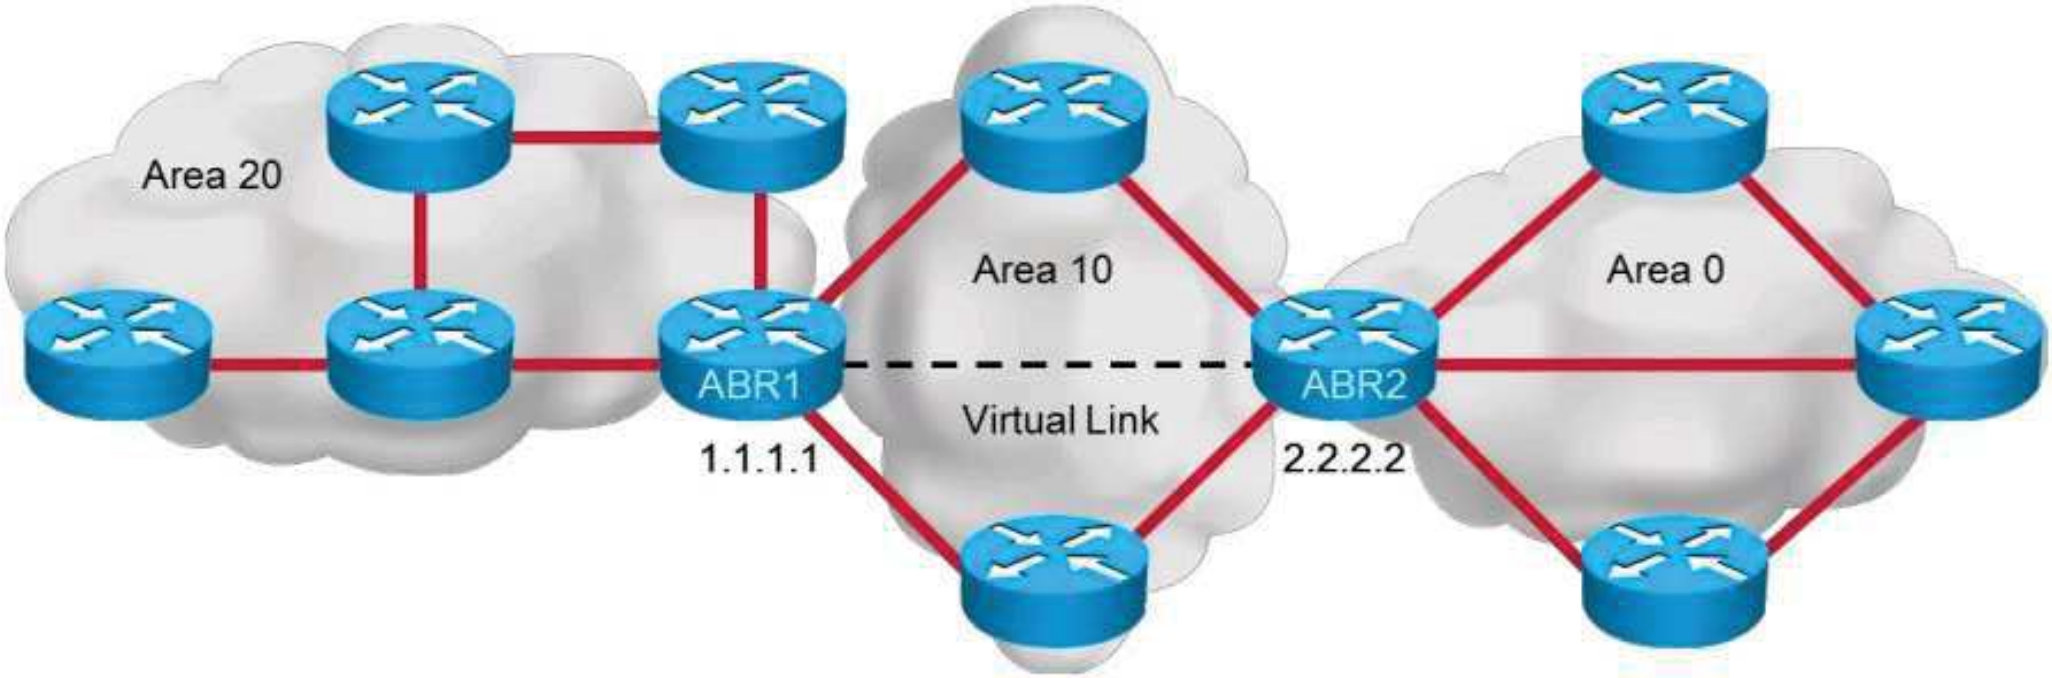

OSPF Virtual Links (Cont.)

Example: The role of a virtual link when having a discontiguous Area 0

- Two companies merge without a direct link between them.
- Virtual links are used to connect the discontiguous Area 0s.
- A logical link is built between routers ABR1 and ABR2.
- Virtual links are also recommended for backup or temporary connections.

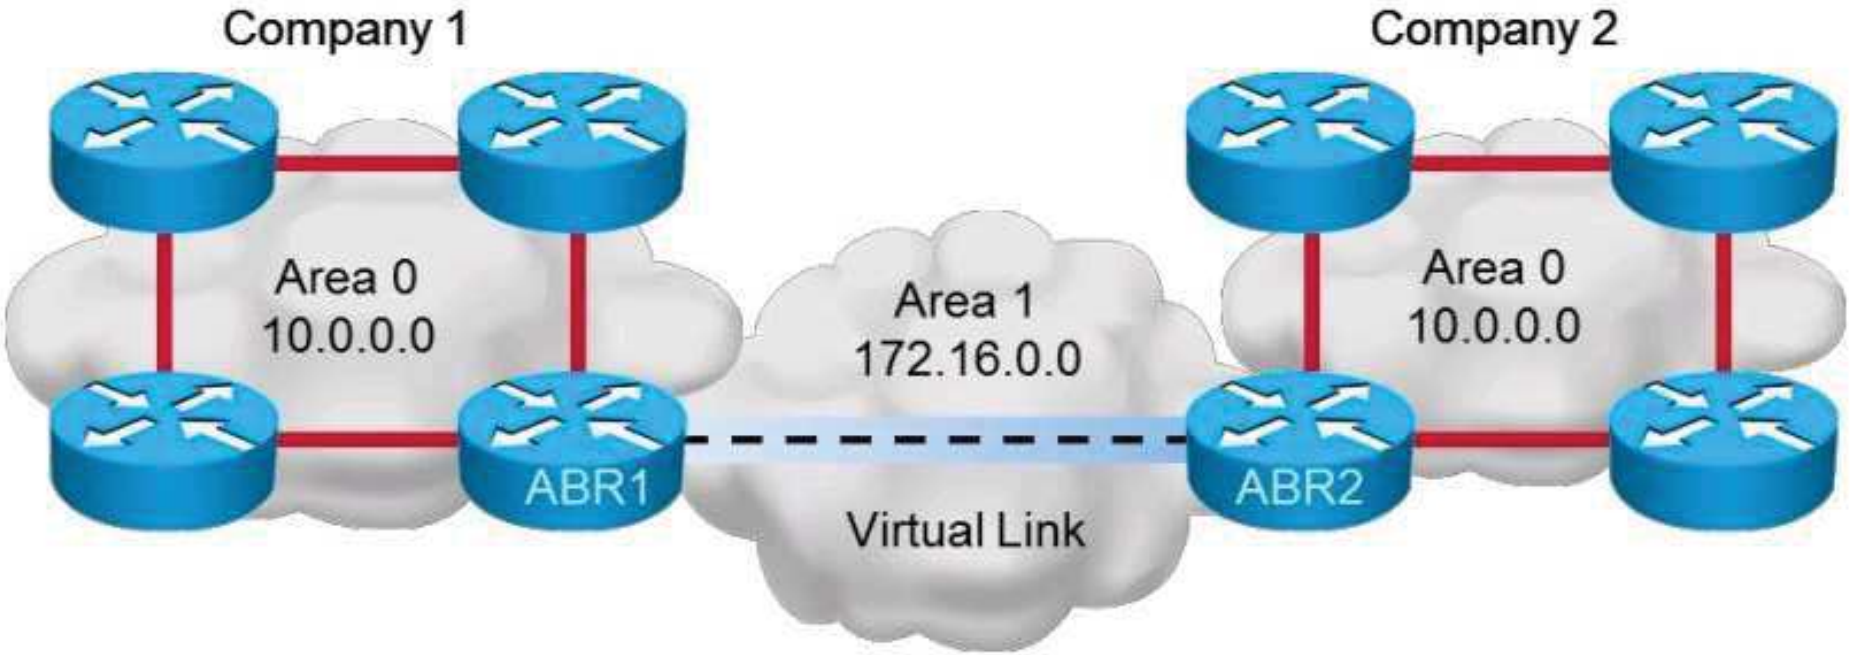

Configuring OSPF Virtual Links

```
router ospf 1
 area 11 virtual-link 10.1.1.1
 network 10.1.10.1 0.0.0.0 area 10
 network 192.168.101.11 0.0.0.0 area 11
```

```
router ospf 1
 address-family ipv4 unicast
 area 0
  interface GigabitEthernet0/0/0/1
 area 1
  interface Loopback0
 area 11
  virtual-link 10.1.10.1
  interface GigabitEthernet0/0/0/0
```

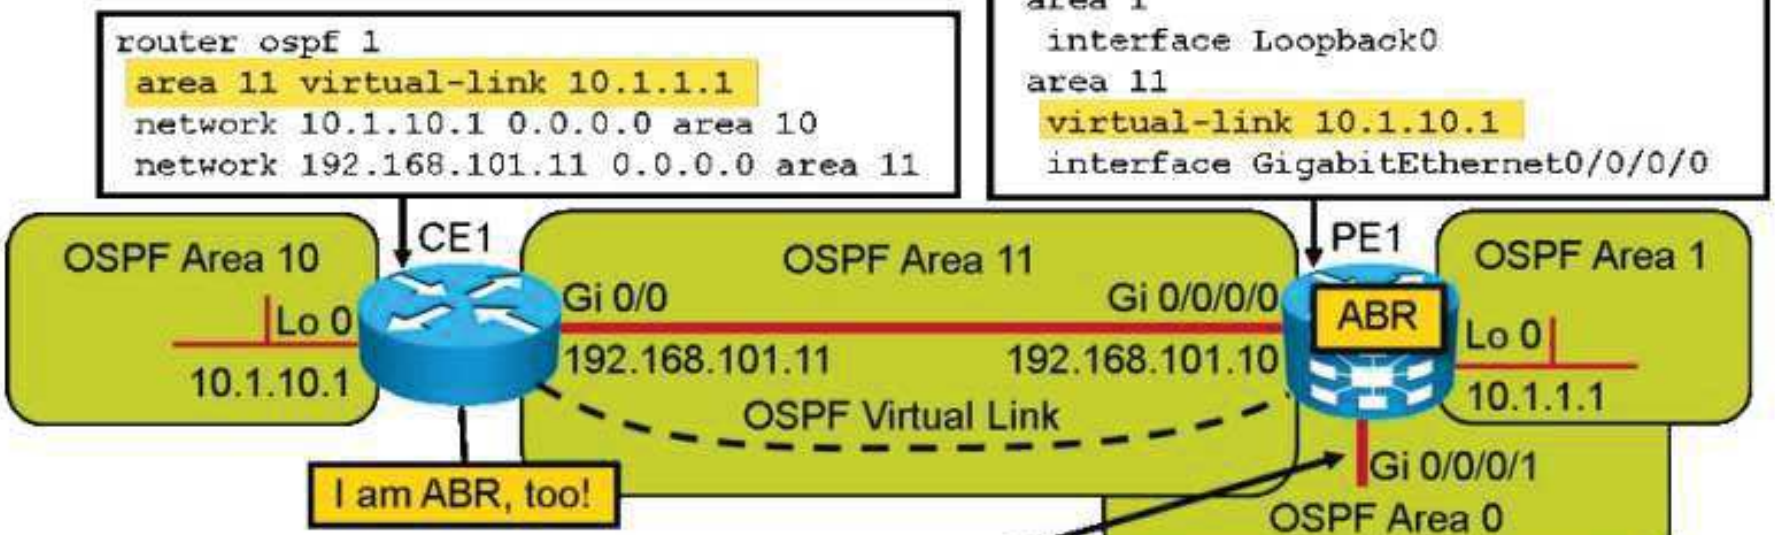


```
RP/0/RSP0/CPU0:PE1# show ospf virtual-links
Virtual Links for OSPF 1
Virtual Link OSPF_VL0 to router 10.1.10.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 11, via interface GigabitEthernet0/0/0/0, Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
    Adjacency State FULL (Hello suppressed)
    Number of DBD retrans during last exchange 0
<... output omitted ...>
```

OSPF Cost

OSPF cost overview:

- The cost, or metric, is an indication of the overhead to send packets over an interface.
- OSPF cost is used as the route selection criteria.
- Dijkstra's algorithm determines the best path by adding all link costs along a path.
- OSPF cost is computed automatically:
  - Cost = 10^8 / Bandwidth (in bps)
  - Bandwidth is specified on the interface with the **bandwidth** command.
- OSPF cost is recomputed after every bandwidth change.

OSPF Cost (Cont.)

Changing OSPF Cost:

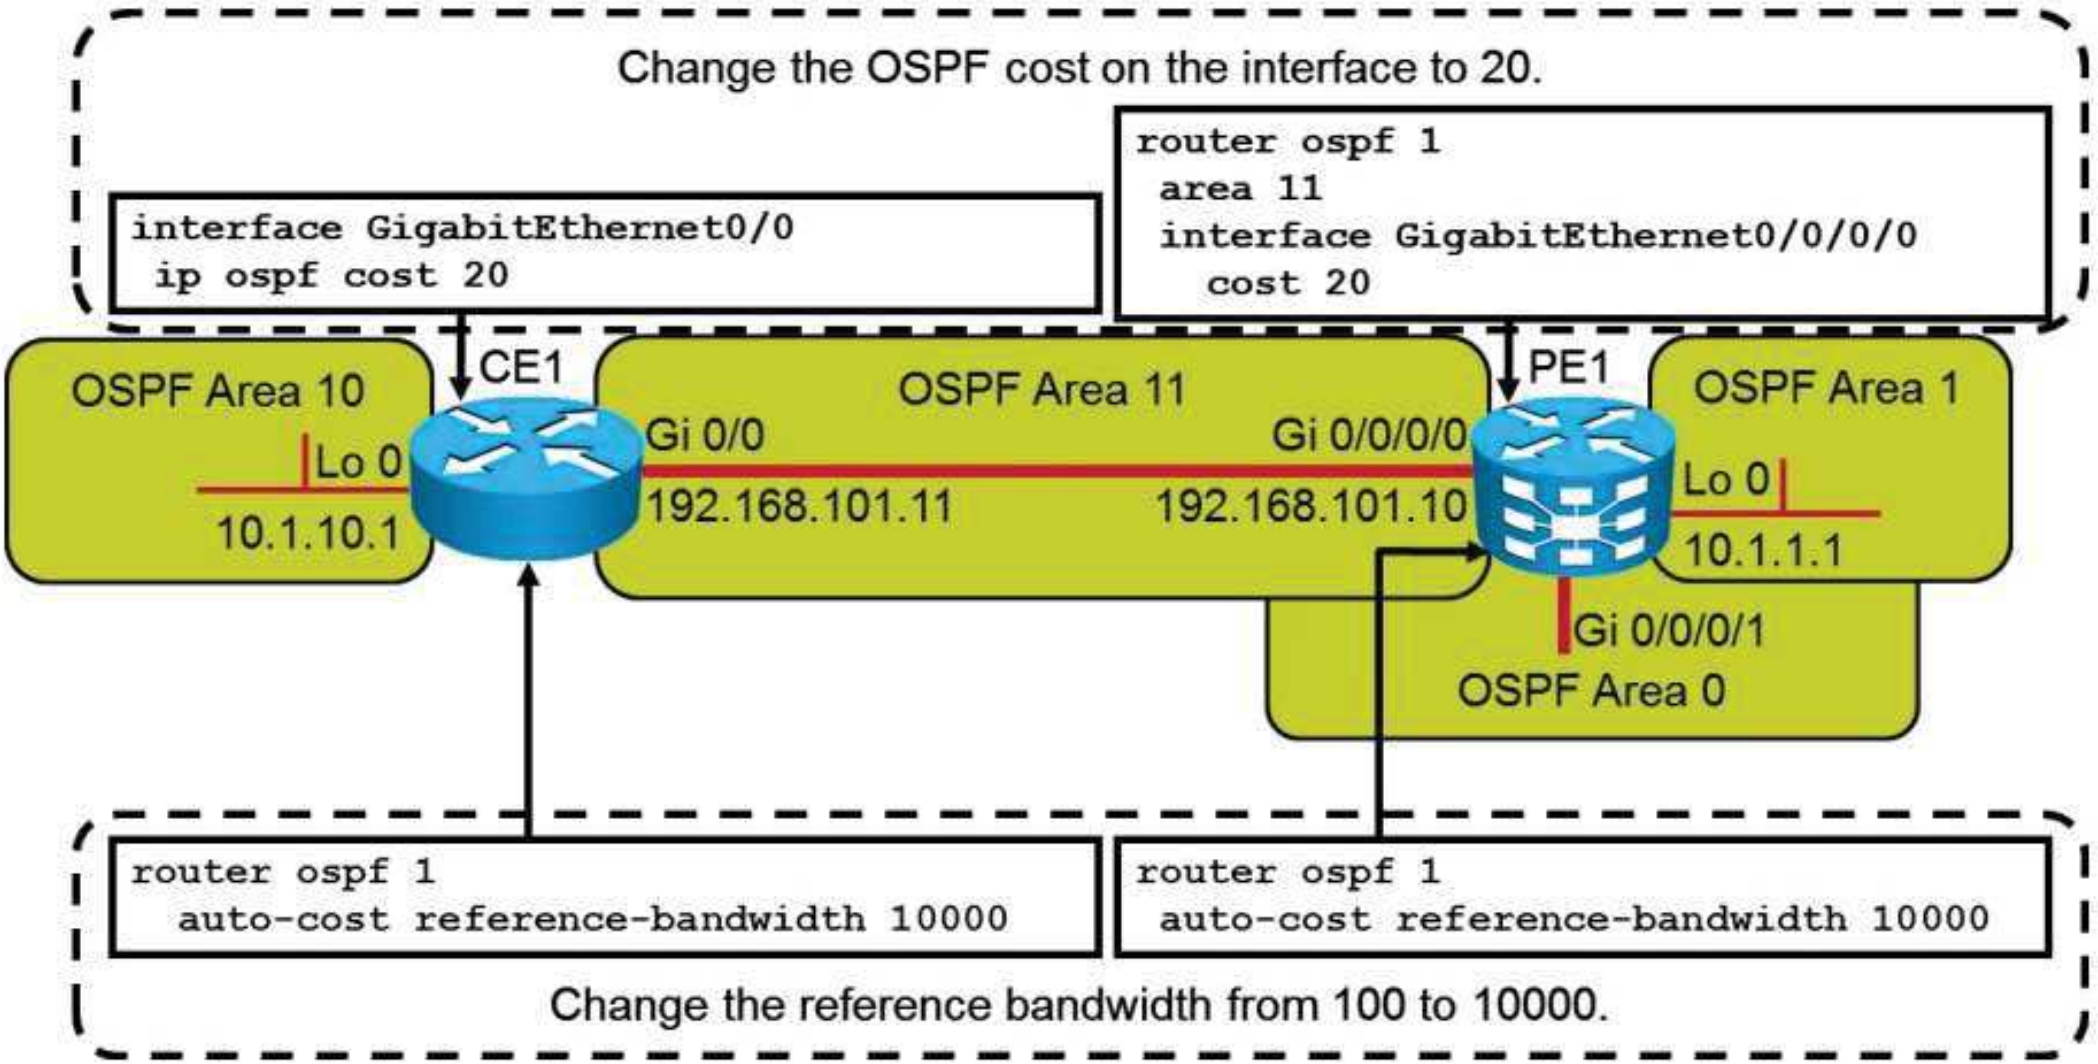

Cisco Nonstop Forwarding

These are the main characteristics of Cisco NSF:

- Cisco NSF is applicable in platforms with dual RPs and works together with SSO.
- Cisco NSF allows the following:
  - Neighbor routers remain established during SSO.
  - Routes on neighboring routers remain valid.
  - Forwarding of data packets continues while routing process on new RP converges.
- Cisco NSF is supported by the following:
  - Routing protocols (OSPF, IS-IS, EIGRP, and BGP).
  - Forwarding operation (Cisco Express Forwarding).
- Device must be NSF-capable.
- Neighboring device must be NSF-aware.

Cisco Nonstop Routing

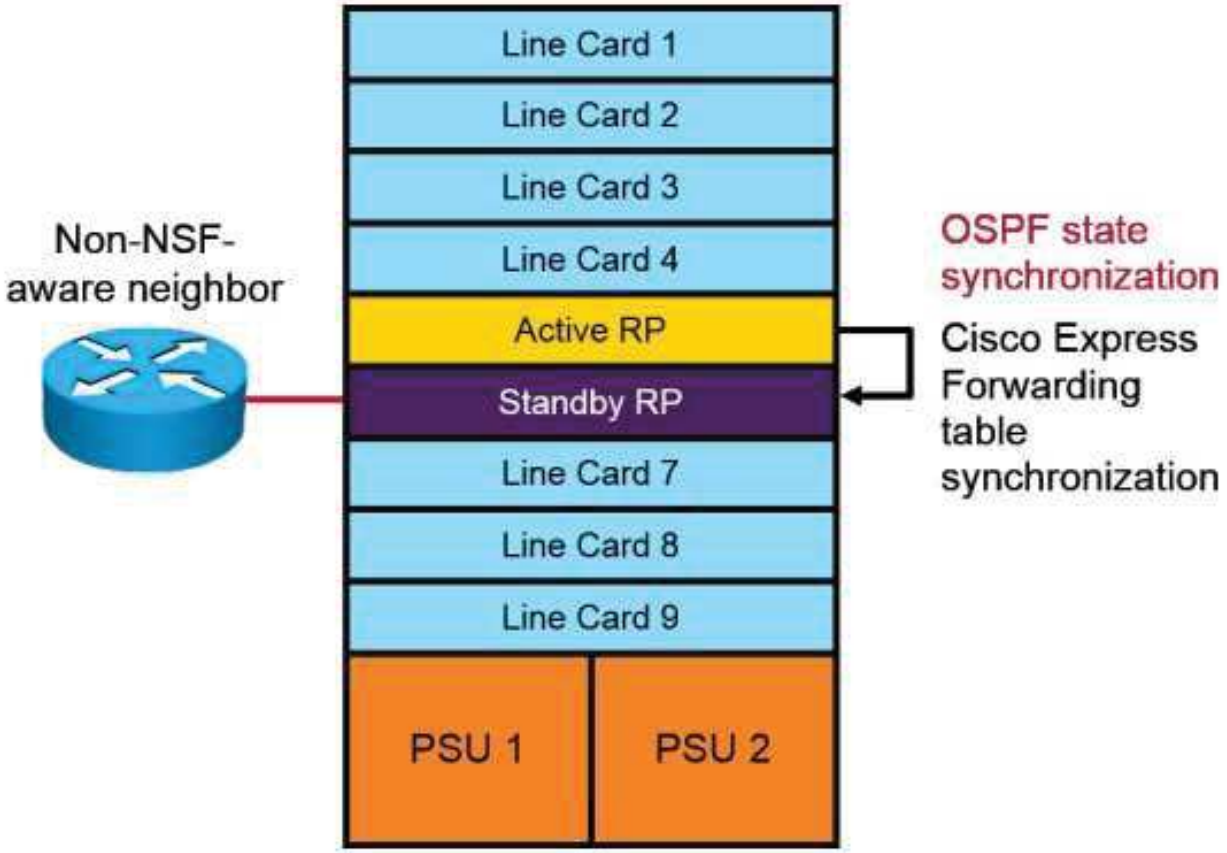

- Used between routers running OSPF.
- Cisco NSR brings Cisco NSF operations to routers that are not NSF-aware.

Cisco NSF and NSR for OSPF

The role of Cisco NSF and NSR for OSPF:

- Minimizes the amount of time that a network is unavailable.
- Continues forwarding IP packets and performs a graceful restart.
- Suppresses routing flaps in Cisco NSF-aware devices.
- Reduces network instability.

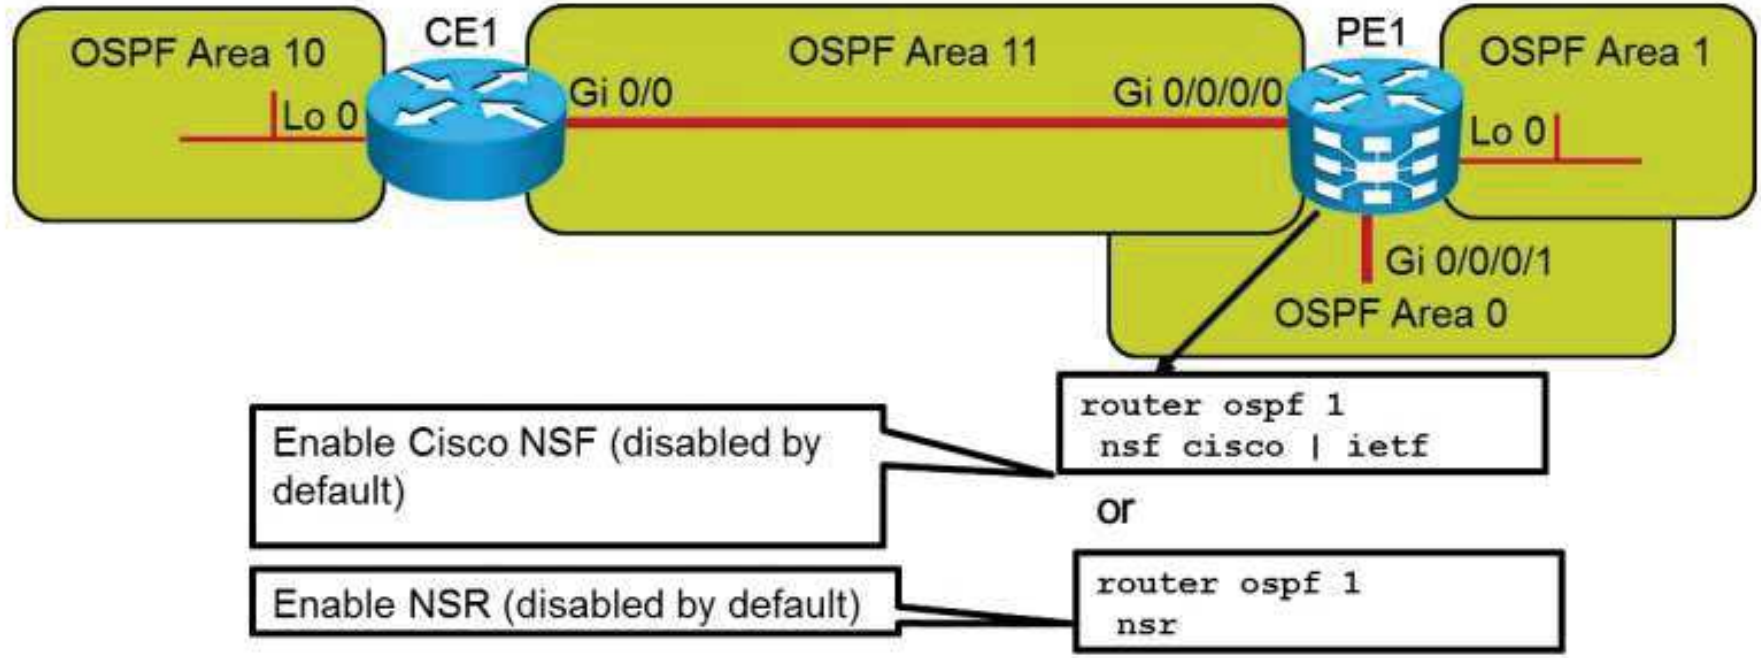

Graceful Restart for OSPFv3

The OSPFv3 graceful restart feature preserves the data plane capability in the following circumstances:

- RP failure, resulting in a switchover to the backup processor.
- Planned OSPFv3 process restart, such as software upgrade or downgrade.
- Unplanned OSPFv3 process restart, such as a process crash.

Bidirectional Forwarding Detection

BFD overview:

- Extremely lightweight Hello protocol that uses UDP to test bidirectional communication
- Used to detect failures in the forwarding path between two adjacent routers
- Millisecond resolution of forwarding plane failure
- Relies on routing protocols to detect neighbors

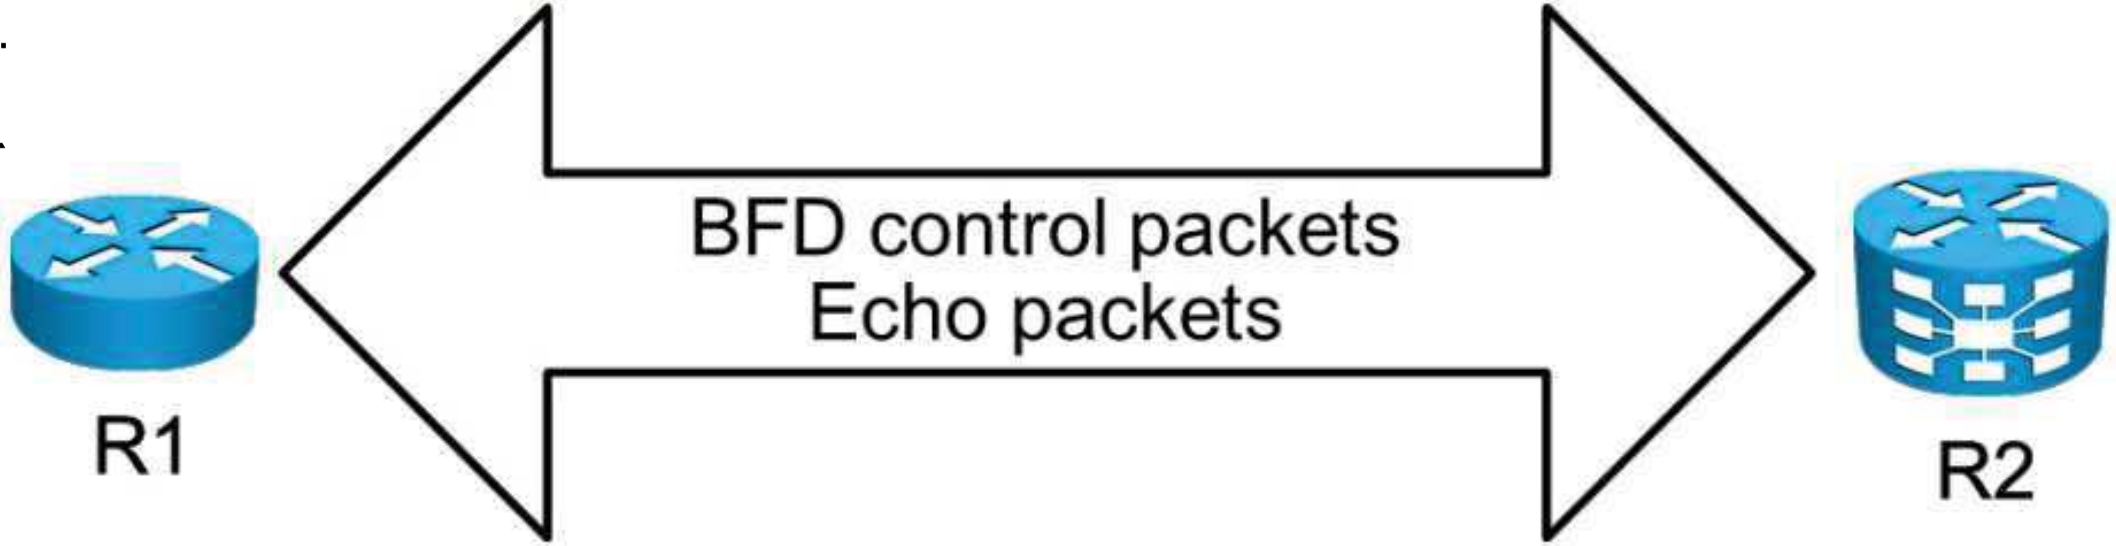

Bidirectional Forwarding Detection (Cont.)

BFD operation:

- Routing protocol (BFD client) bootstraps BFD to create BFD session to a neighbor:
  - BFD client receives link status change notification.
  - Receive and transmit intervals are negotiated and configurable.
- Two systems agree on a method to detect failure.
- In case of failure, BFD notifies BFD client.
  - BFD client independently decides on action.

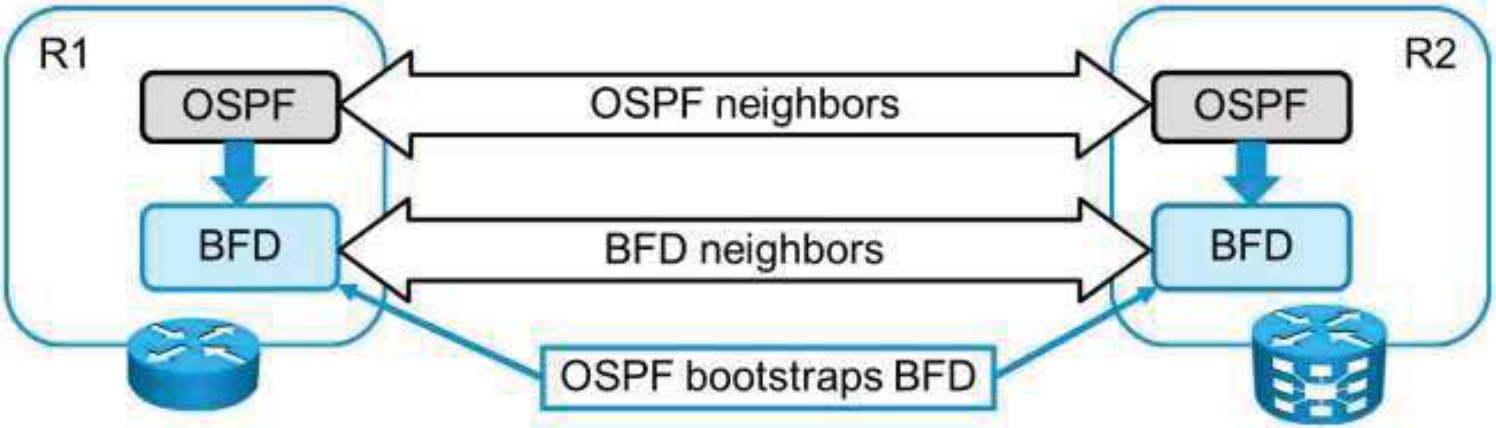

Bidirectional Forwarding Detection for OSPF

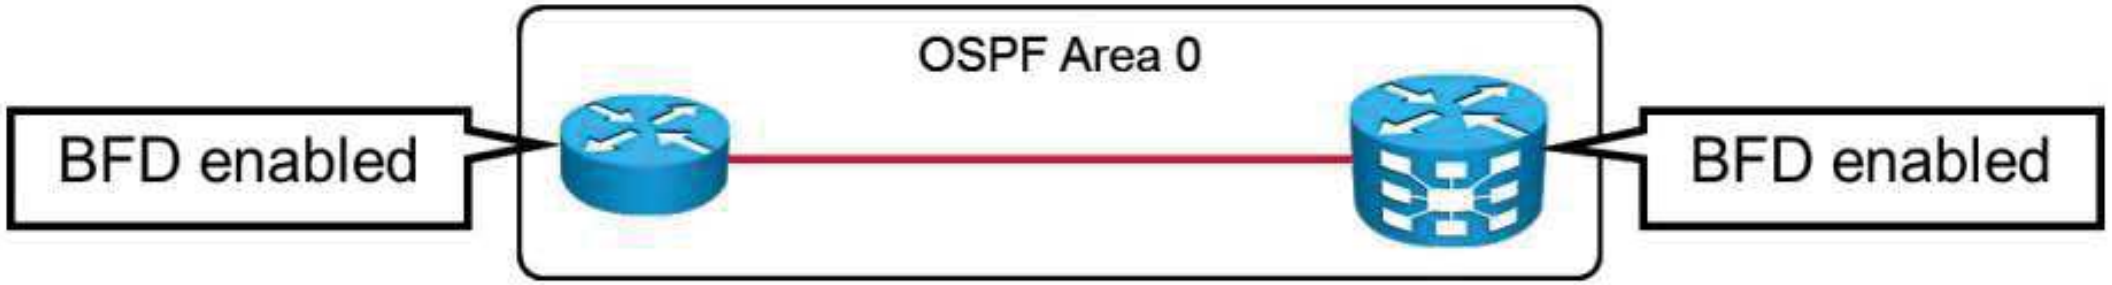

- Enable OSPF to use BFD on all interfaces:
  - Use the following Cisco IOS router command:

```
router ospf 1
bfd all-interfaces
!
```

Bidirectional Forwarding Detection for OSPF (Cont.)

- Enable OSPF to use BFD on a single interface:
  - Use the following Cisco IOS interface configuration command:

```
interface TenGigabitEthernet3/0/1
ip ospf bfd
bfd interval 100 min_rx 100 multiplier 3
!
```

- To enable BFD and adjust parameters, use the following Cisco IOS XR router commands:

```
router ospf 1
area 0
  interface TenGigE0/1/4/0
    bfd fast-detect
    bfd minimum-interval 100
    bfd multiplier 3
```

OSPF Authentication Overview

- OSPF authentication is used to prevent the following:
  - Undesired adjacencies and rogue routes to be inserted into OSPF.
  - Changes in routing information.
- OSPFv2:
  - Plaintext authentication—avoid at all times!
  - MD5 authentication.
  - Authentication material is inserted into OSPF header of every OSPF packet and checked by other router.
- OSPFv3 does not have an authentication mechanism; it relies on IPsec built into IPv6.

OSPFv2 Authentication

OSPFv2 authentication type and key can be configured at different levels:

- Routing process
- Area
- Interface

If authentication is not configured on a lower level, authentication settings are inherited from a higher level.

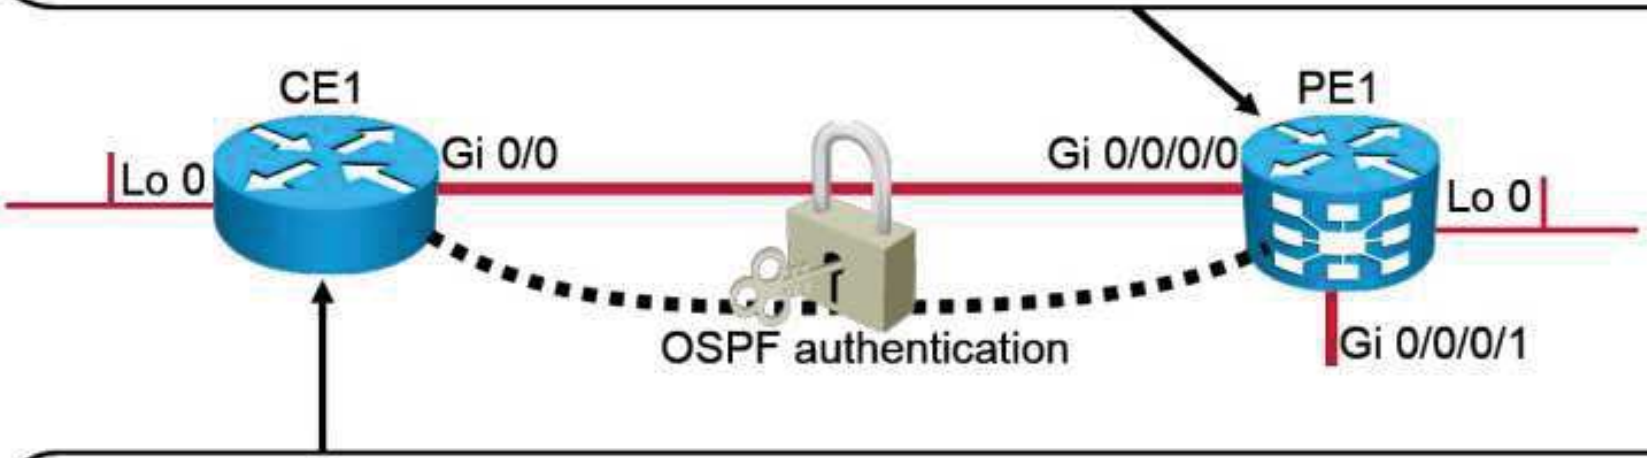

Authentication type can be configured per area in router configuration mode or per interface.

If authentication is not configured per interface, authentication type is inherited from area configuration.

Authentication key can be configured only per interface.

OSPFv3 Authentication

- OSPFv3 uses native functionality offered by IPv6:
  - IPsec AH for authentication and integrity check
  - IPsec ESP for encryption of payload
- Security policy definition on the router is mandatory:
  - SPI value
  - Hashing and encryption algorithms
  - Keys for authentication and encryption

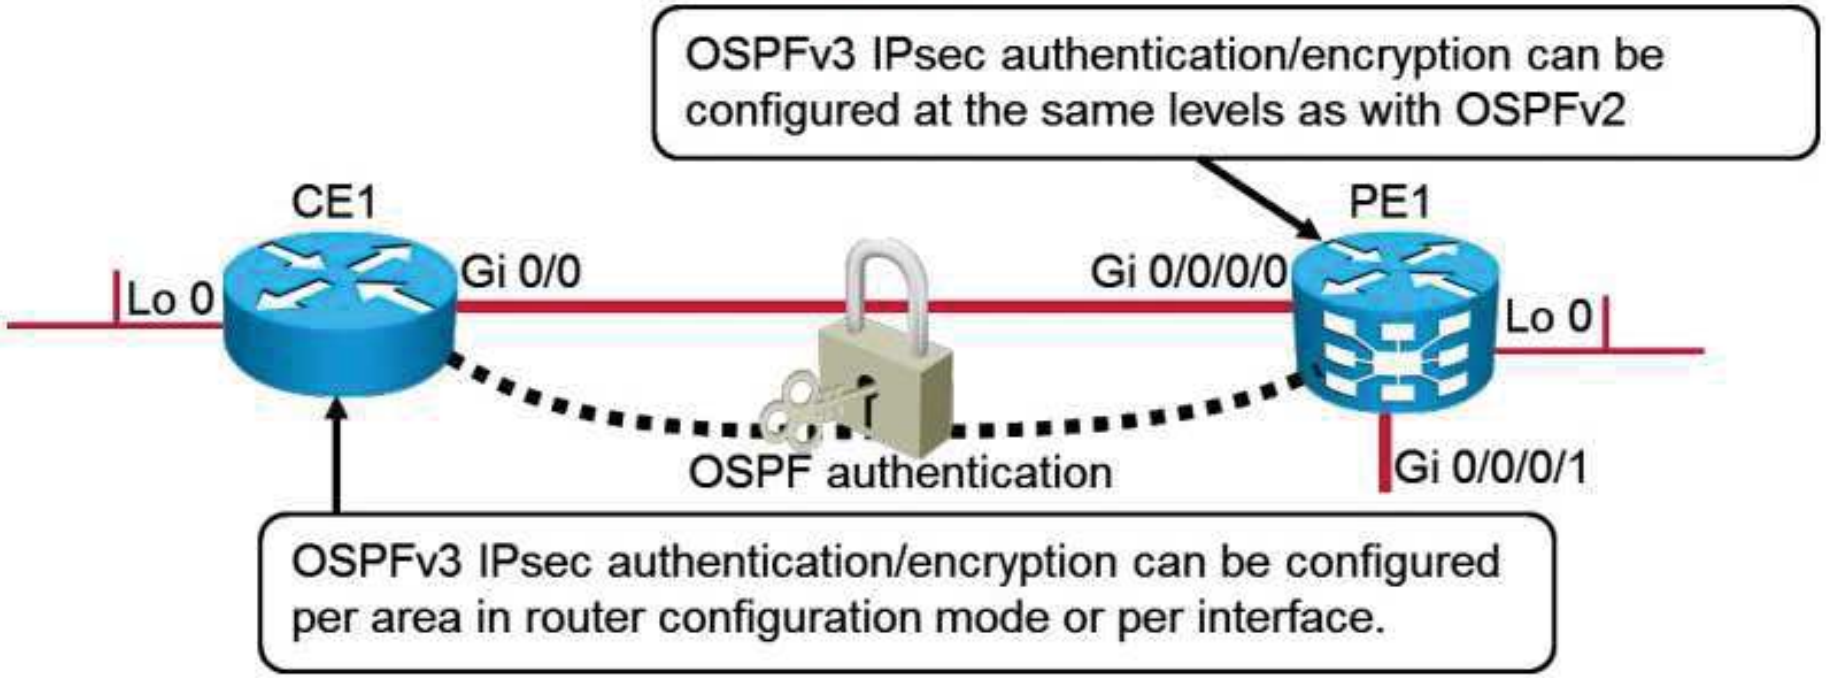

Configuring OSPFv2 Authentication

```
interface GigabitEthernet0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 cisco
```

```
router ospf 1
 area 11
  interface GigabitEthernet0/0/0/0
   authentication message-digest
   message-digest-key 1 md5 encrypted cisco
```


```
RP/0/RSP0/CPU0:PE1# show ospf interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
  Internet Address 192.168.101.10/24, Area 11
  Process ID 1, Router ID 10.1.1.1, Network Type BROADCAST, Cost: 20
  Transmit Delay is 1 sec, State BDR, Priority 1, MTU 1500, MaxPktSz 1500
  Designated Router (ID) 10.1.10.1, Interface address 192.168.101.11
  Backup Designated router (ID) 10.1.1.1, Interface address 192.168.101.10
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06
<... output omitted ...>
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.10.1  (Designated Router)
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
  Multi-area interface Count is 0
```

Configuring OSPFv3 Authentication

```
router ospfv3 1
 area 11
  interface GigabitEthernet0/0/0/0
   authentication ipsec spi 256 sha1 password 12…
```

```
interface GigabitEthernet0/0
 ipv6 ospf 1 area 11
 ipv6 ospf authentication ipsec spi 256 sha1 12…
```

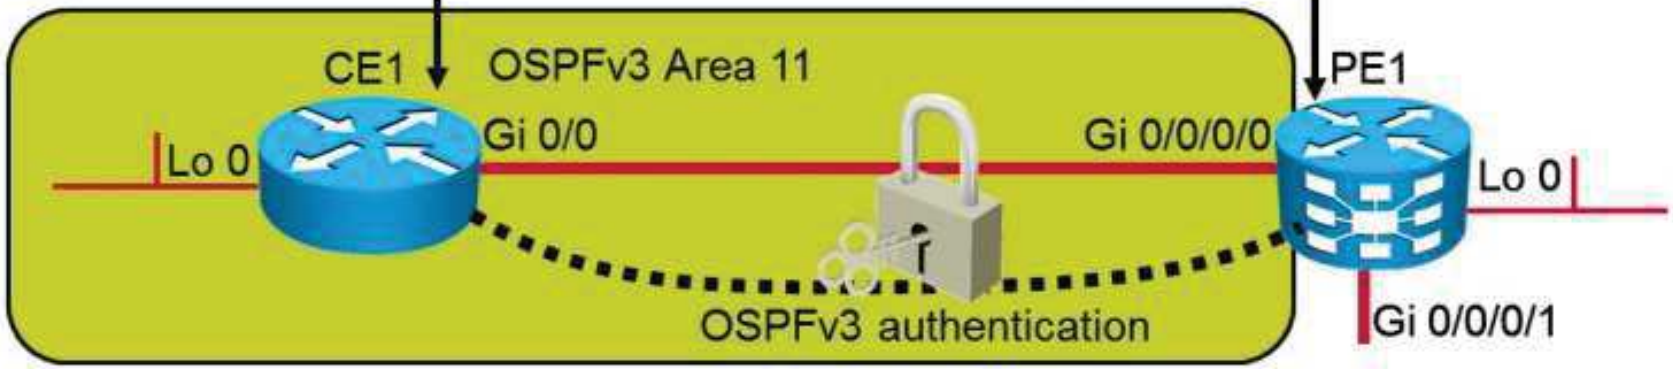

Summary

- Before OSPF is enabled on the service provider network, there must be detailed planning of IP addresses, areas, positioning of ABRs and ASBRs, and redistribution and summarization.
- OSPF requires router ID to uniquely describe each router in a network.
- Designating OSPF interface as passive disables sending and receiving of OSPF packets.
- You can use various **show** commands to verify OSPF.
- A virtual link allows discontiguous Area 0s to be connected, or a disconnected area to be connected to Area 0 via a transit area.

Summary (Cont.)

- You can use the **virtual-link** Cisco IOS XR command to configure OSPF virtual link.
- OSPF cost can be changed by changing bandwidth or cost on the link or by changing reference bandwidth.
- Cisco NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover.
- After an RP failover, the Cisco NSF-capable router sends an OSPF Cisco NSF signal to neighboring Cisco NSF-aware devices to inform the neighbor not to reset the neighbor relationship.
- OSPFv3 graceful restart feature preserves the data plane capability in RP failure or OSPF process restart.

Summary (Cont.)

- BFD provides fast peer failure detection times independently of media types, encapsulations, topologies, and the routing protocols.
- You would not use BFD only if an interface is flaky and you want to retain routing adjacency across short failures.
- OSPF supports authentication to prevent undesired adjacencies and changes in routing information.
- To enable OSPFv2 authentication on the Cisco IOS XR, use the **authentication message-digest** and **message-digest-key md5** router OSPF commands.

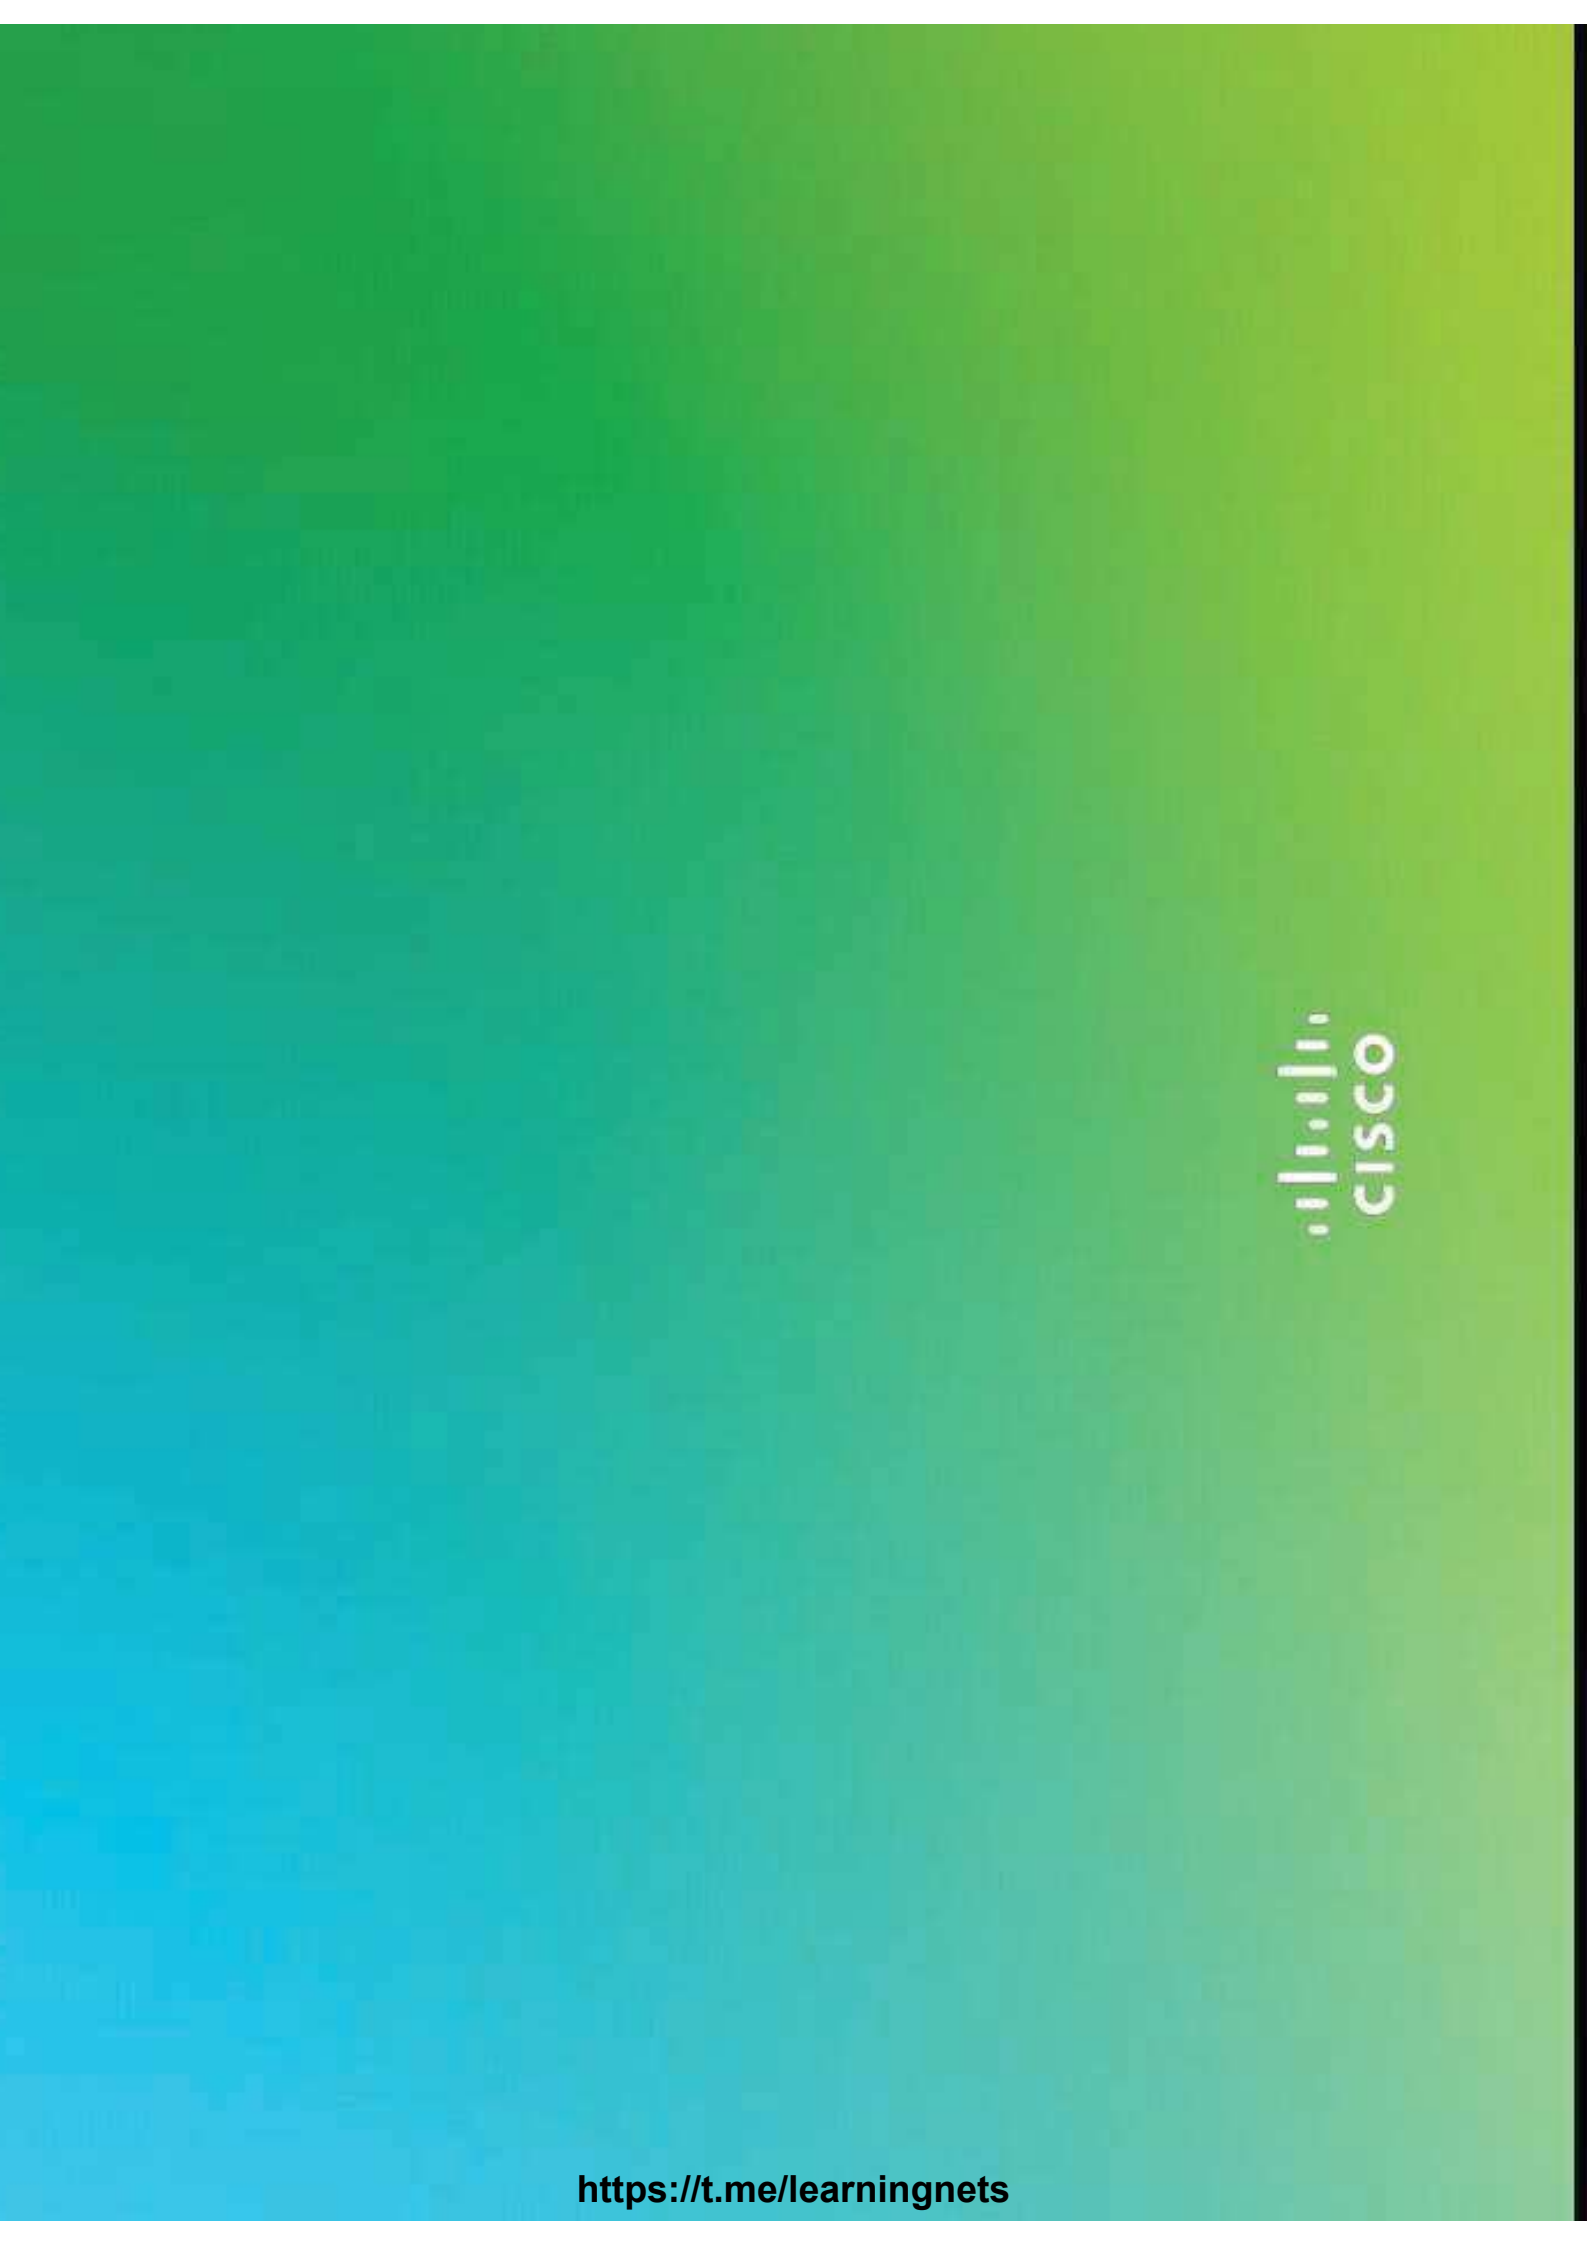

Implementing OSPF Special Area Types

Implement OSPF in the Service Provider Network

OSPF Summarization

OSPF summarization overview:

- Networks are normally translated into type 3 LSAs in other areas.
- Route summarization is the consolidation of advertised addresses.
  - On ABR, summarize type 3 LSAs.
  - On ASBR, summarize type 5 LSAs.
- A good addressing plan is required.
- A drawback is the possibility of suboptimal routing.

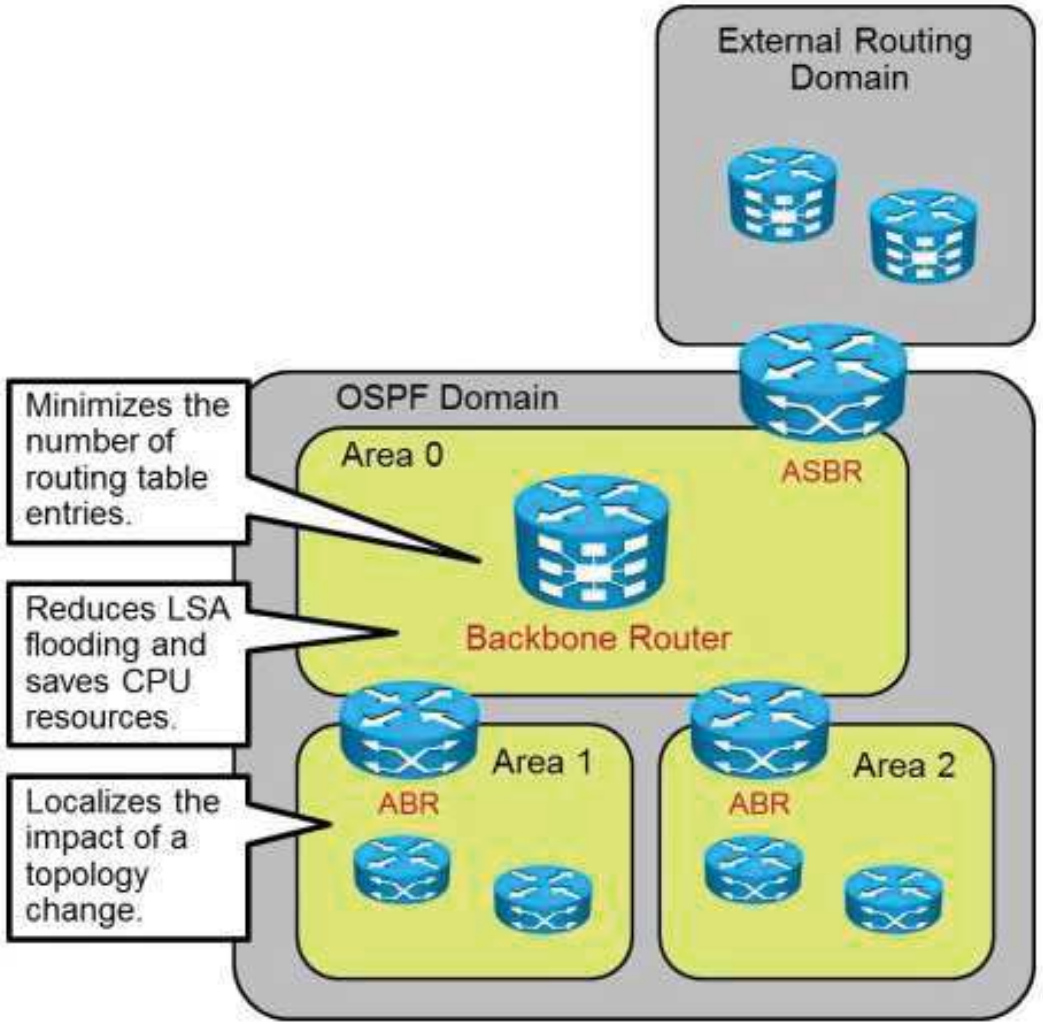

OSPF Interarea Route Summarization

Main characteristics of OSPF interarea route summarization:

- A summary route is generated if at least one subnet within the area falls in the summary address range.
- A summarized route metric will be equal to the lowest cost of all subnets within the summary address range.
- The ABR creates a route to Null0 to avoid loops, only for the summary routes of connected areas.

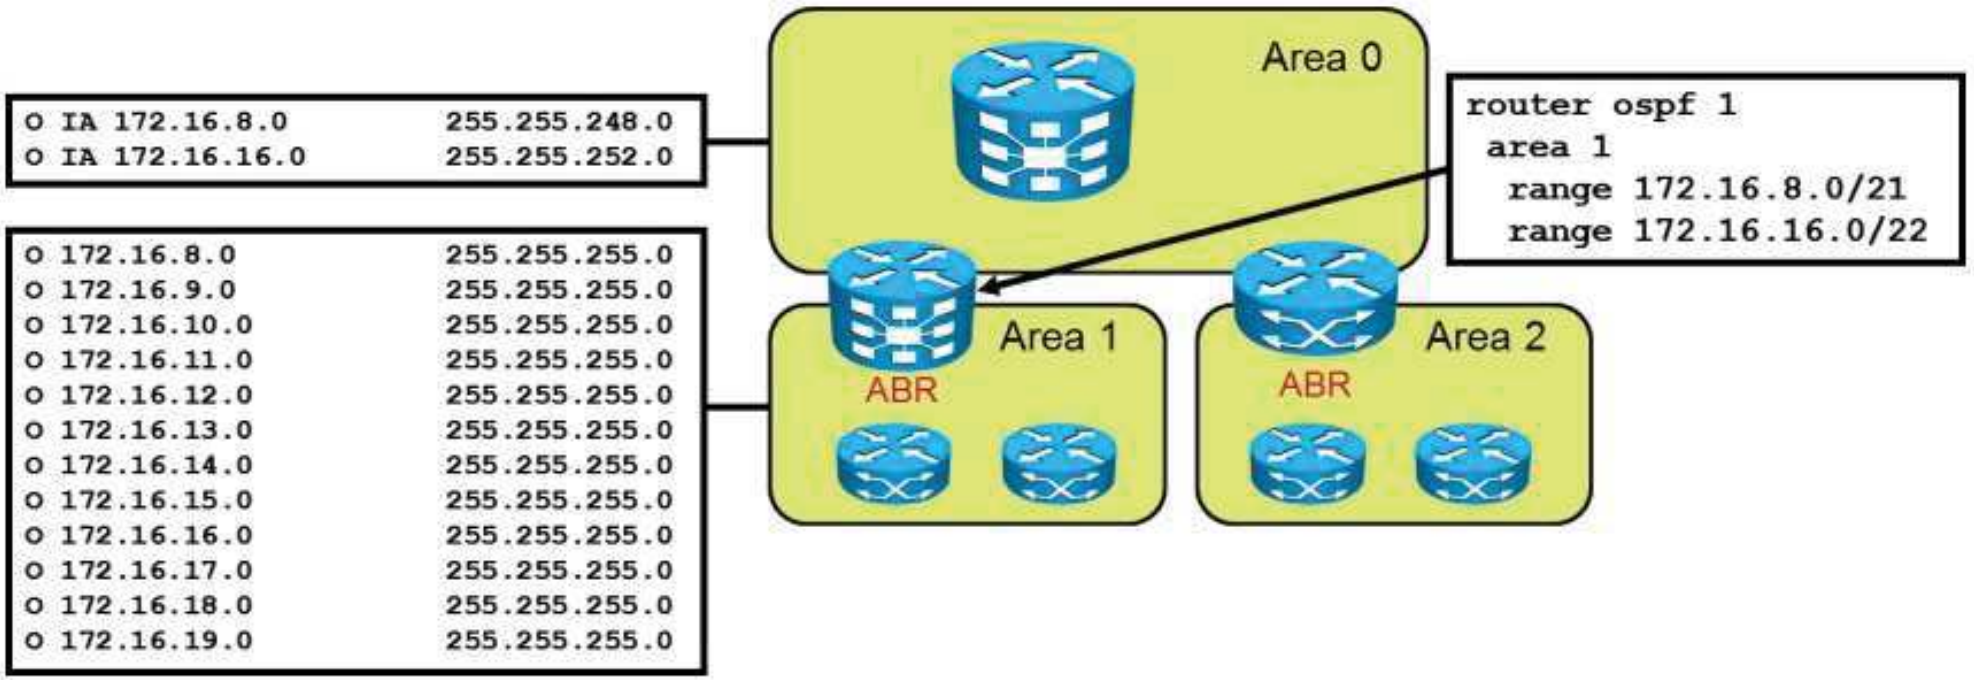

OSPF External Route Summarization

- Summarization can be used for external routes:
  - On an AS boundary for type 5 LSAs (redistributed routes).
  - On an NSSA ABR for type 5 LSAs translated from type 7 LSAs.
- A summary route to Null0 will be created for each summary range.

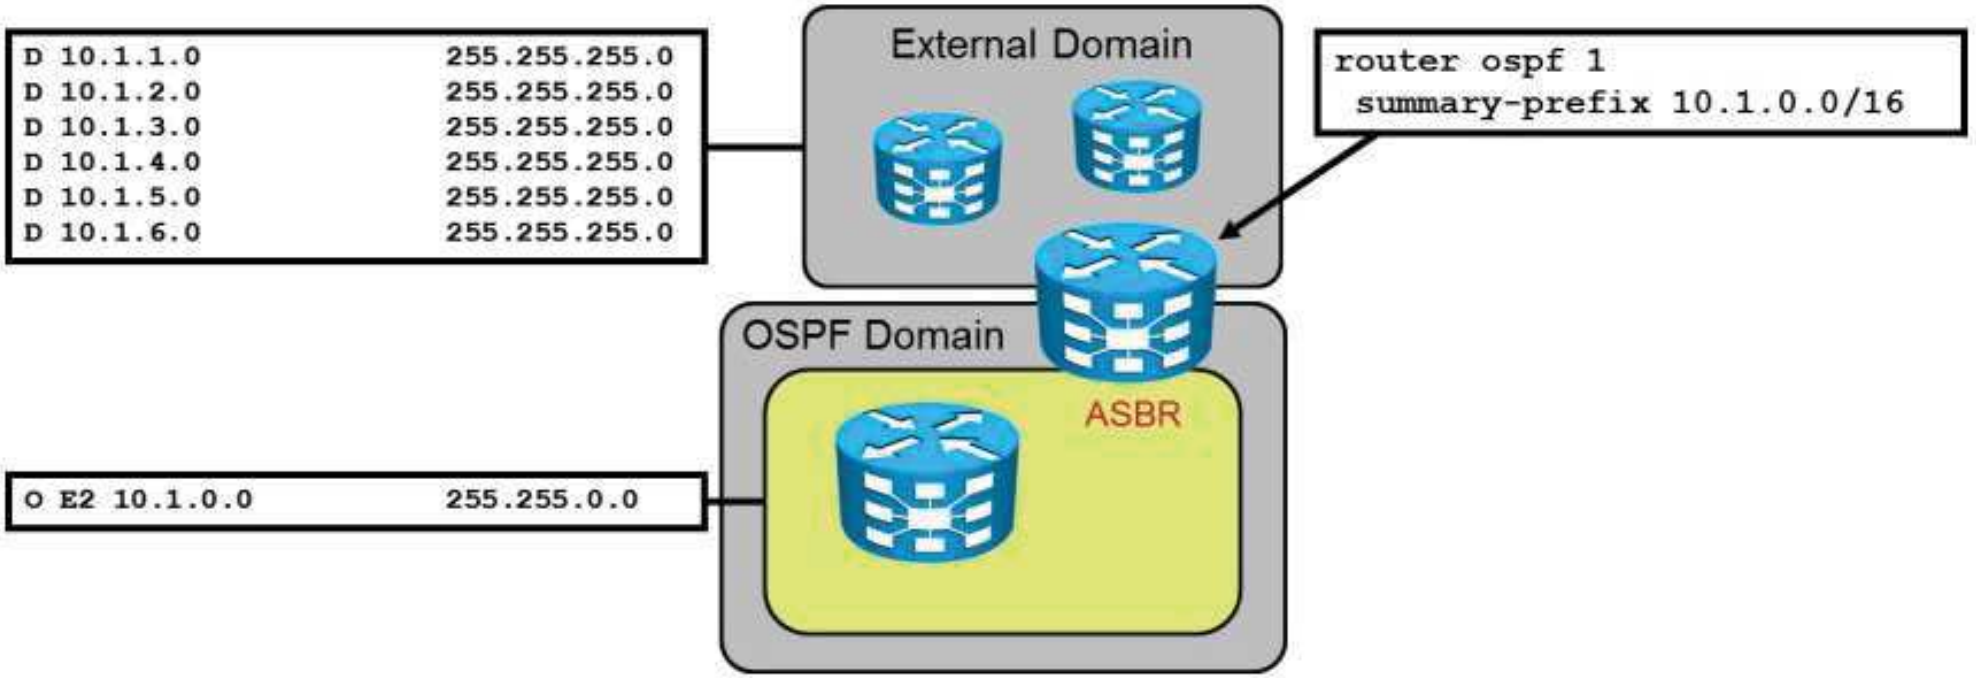

Default Routes in OSPF

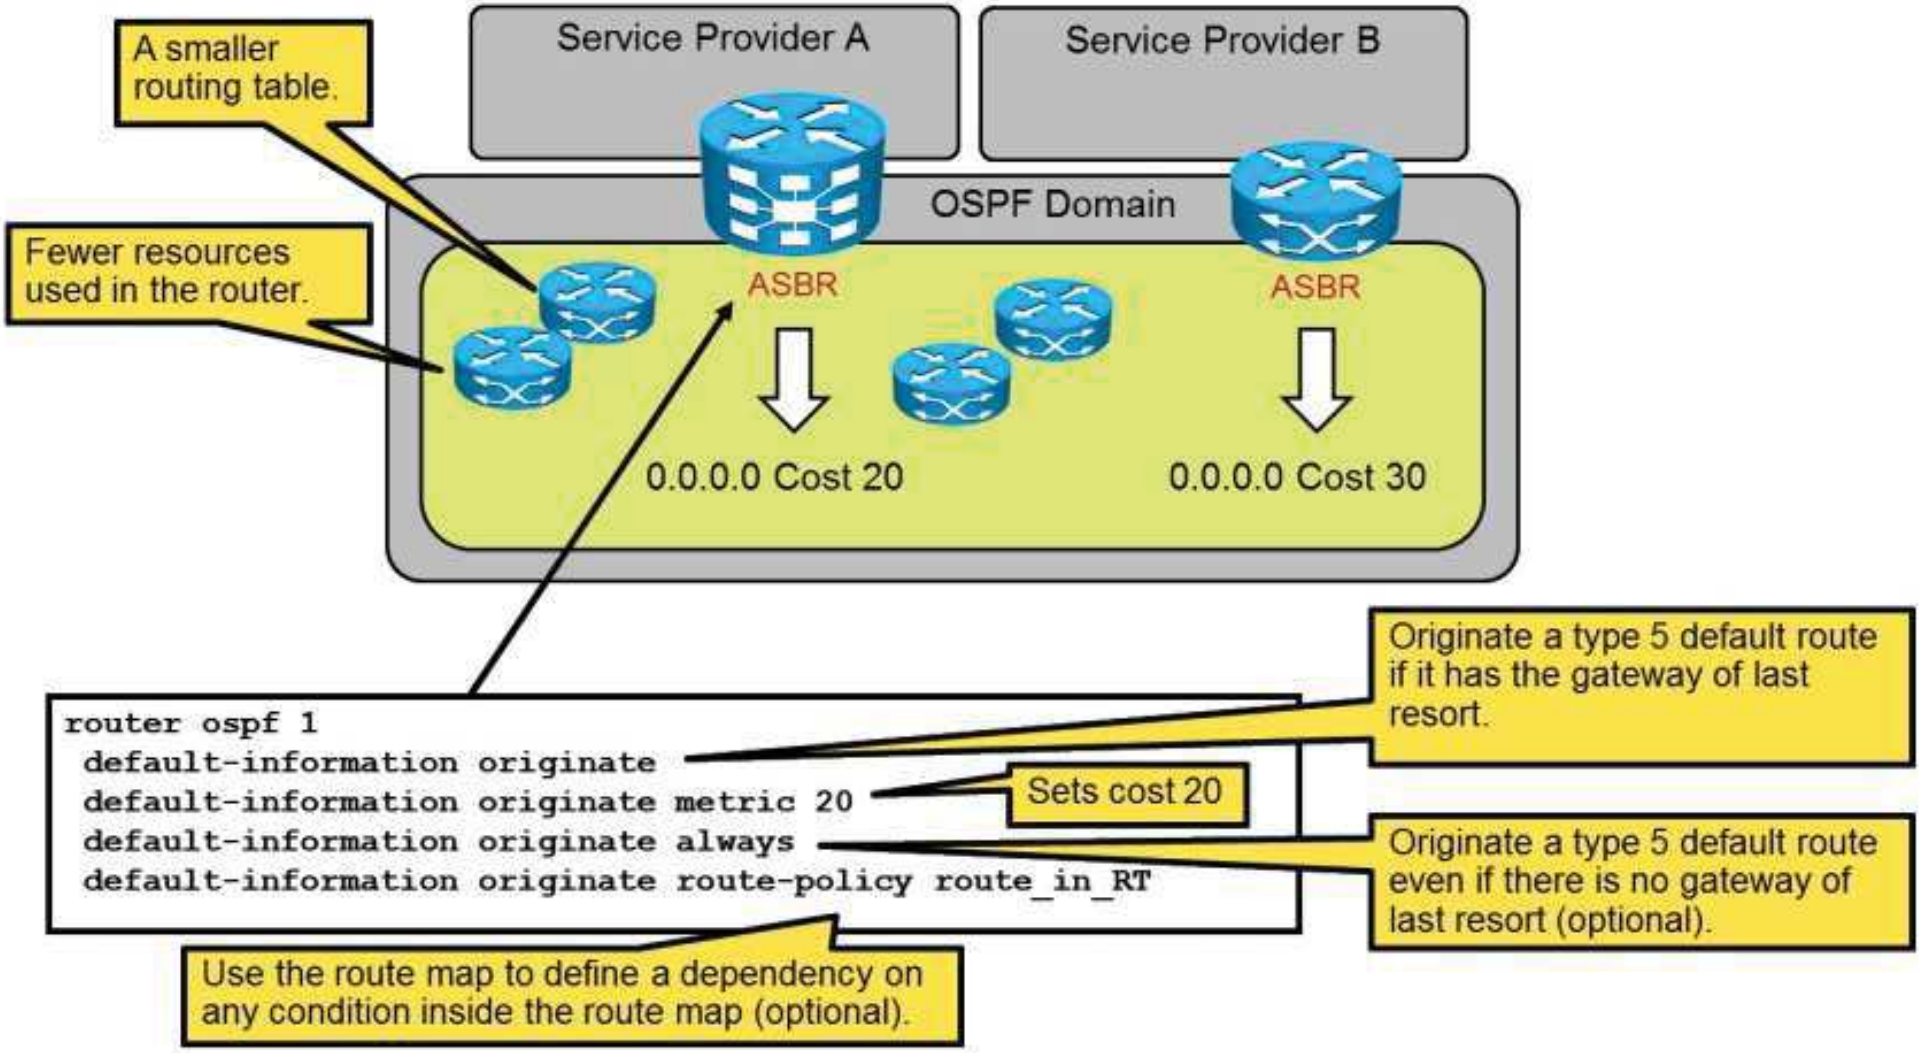

- A default route is injected into OSPF as an external type 5 LSA.
- Default route distribution is not on by default.

OSPF Area Types

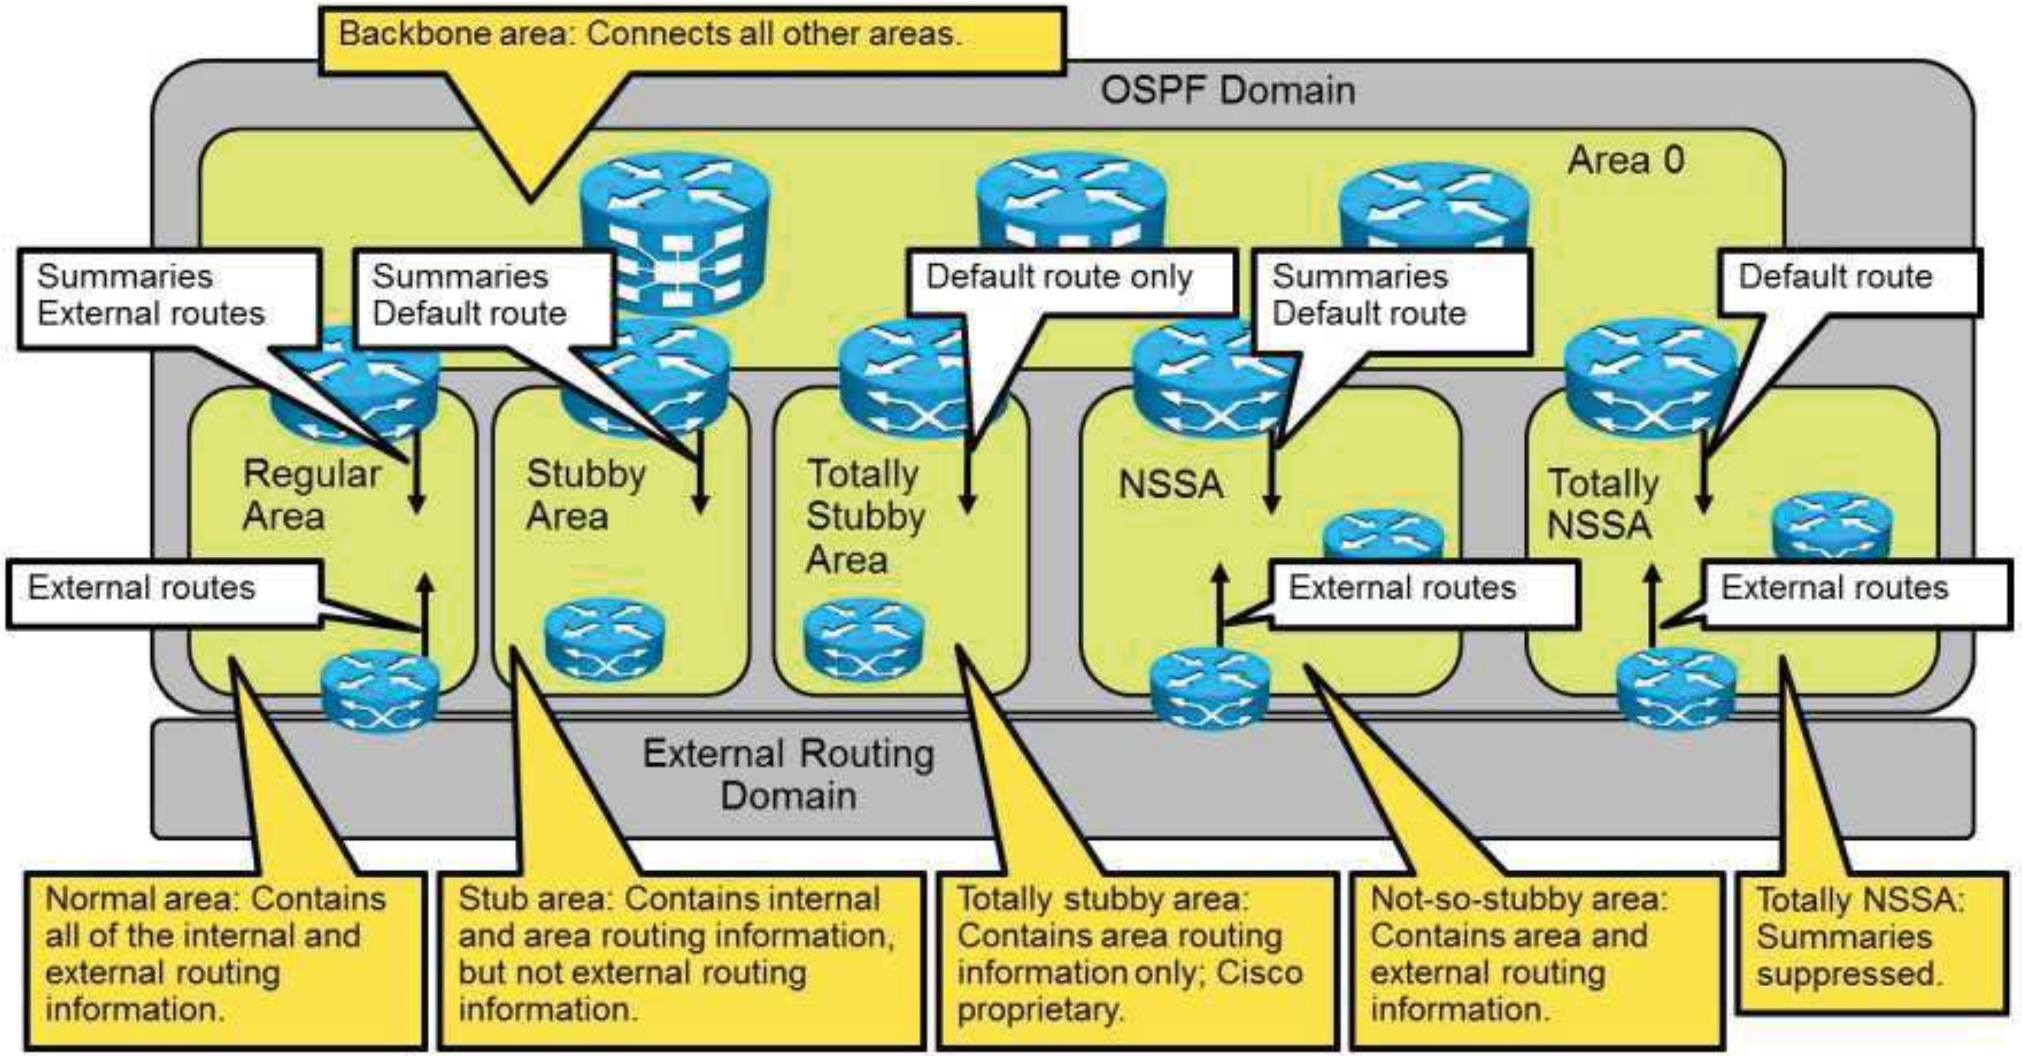

- OSPF is based on a two-level hierarchical area structure.
- Each area has its own topology database.

Configuring OSPF Area Types

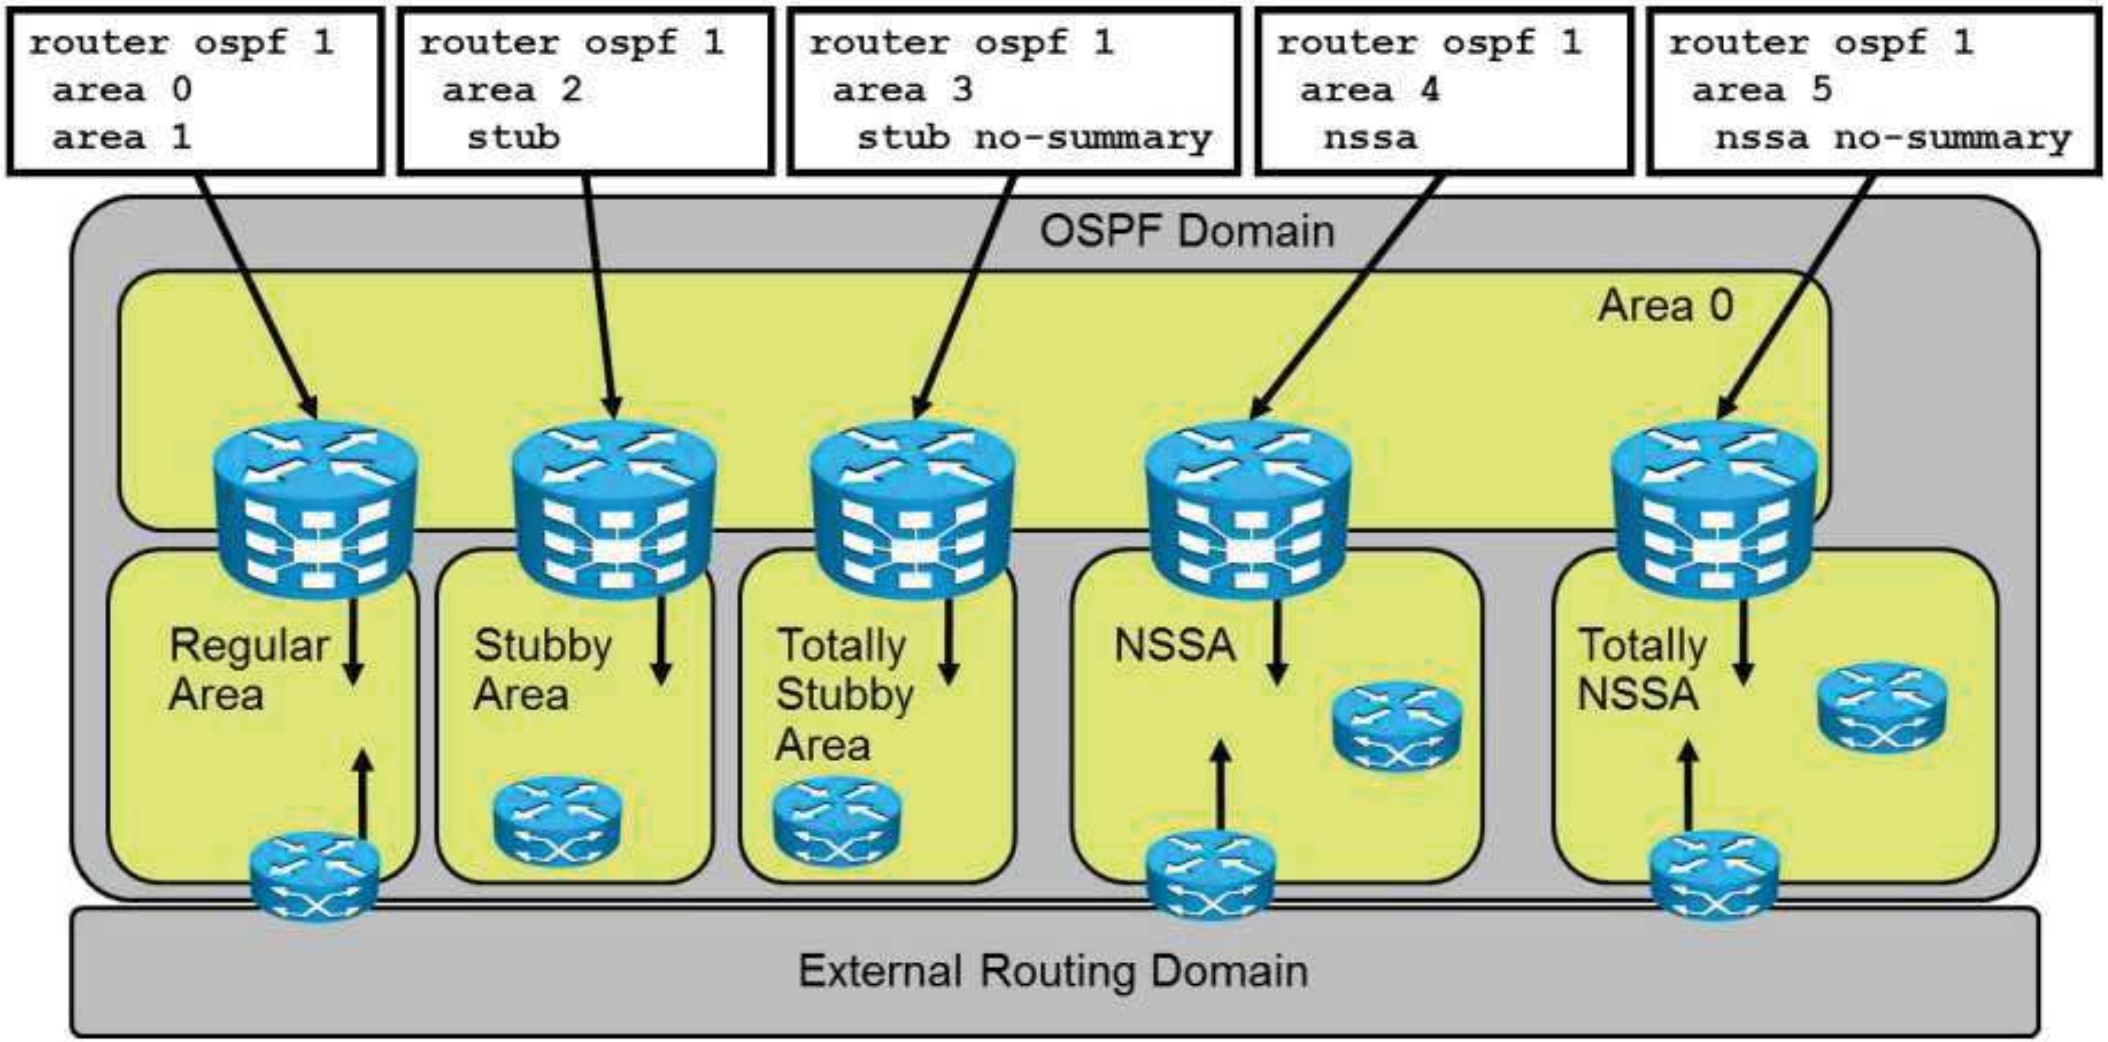

OSPF Router and LSA Types

OSPF border router characteristics:

- ABR generates summary LSAs.
- ASBR generates external LSAs.
- Summary and external LSAs can be blocked and default route sent instead.

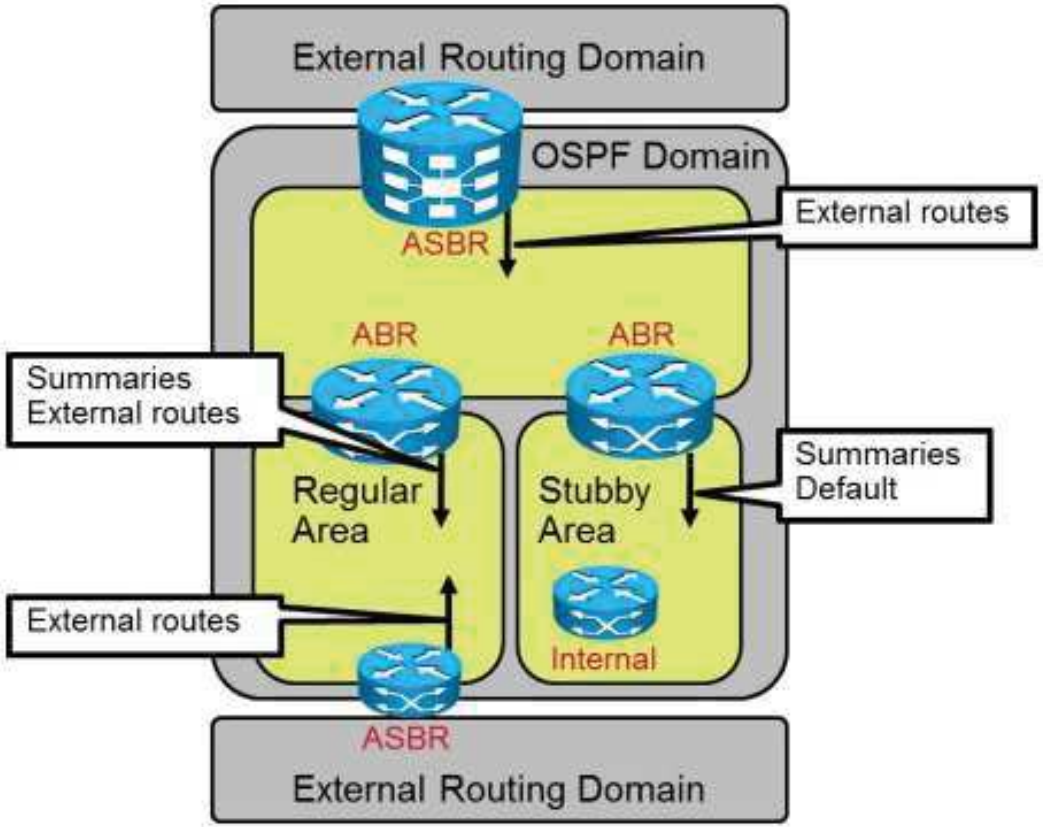

OSPF Stub Areas

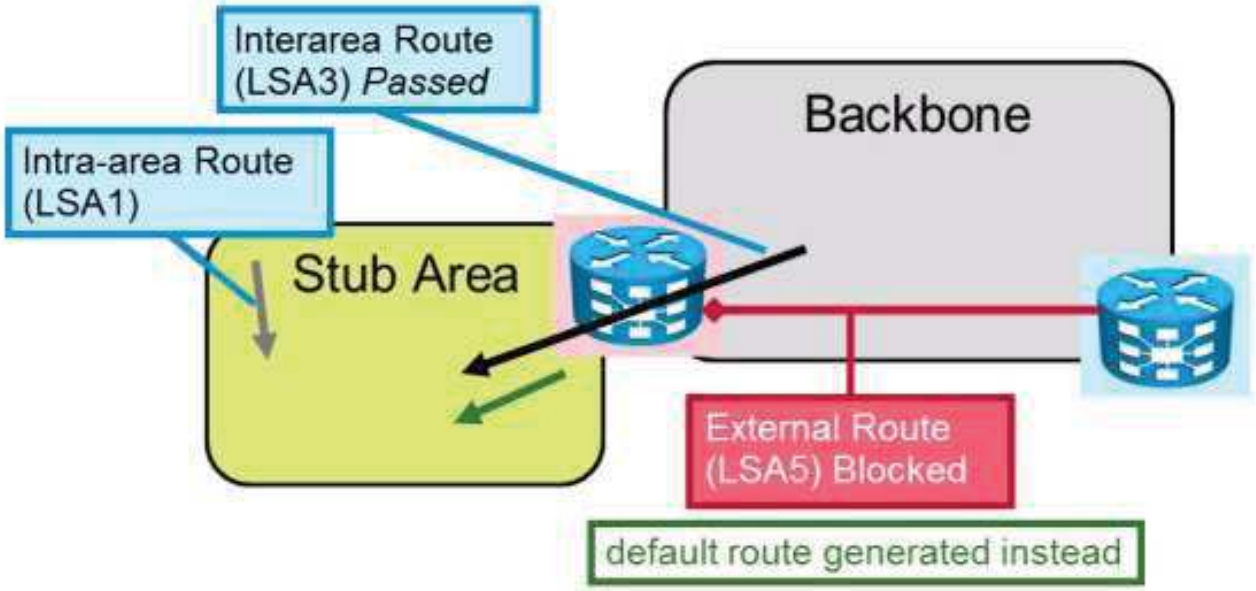

- These are the characteristics of an OSPF Stub Area:
  - No external routes
  - Interarea routes present
  - Intra-area routes present
  - Default route generated

OSPF Totally Stubby Area

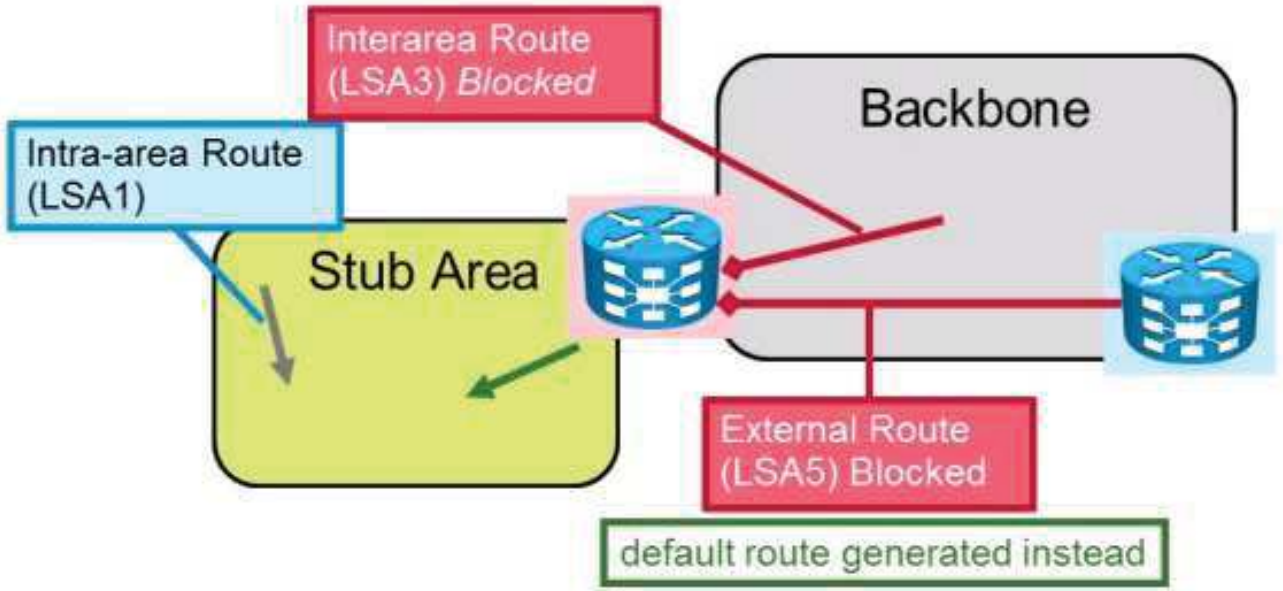

- These are the characteristics of an OSPF Totally Stubby Area (**stub no-summary**):
  - No external routes
  - No interarea routes
  - Intra-area routes present
  - Default route generated
  - Cisco proprietary feature

Configuring OSPF Stub Areas

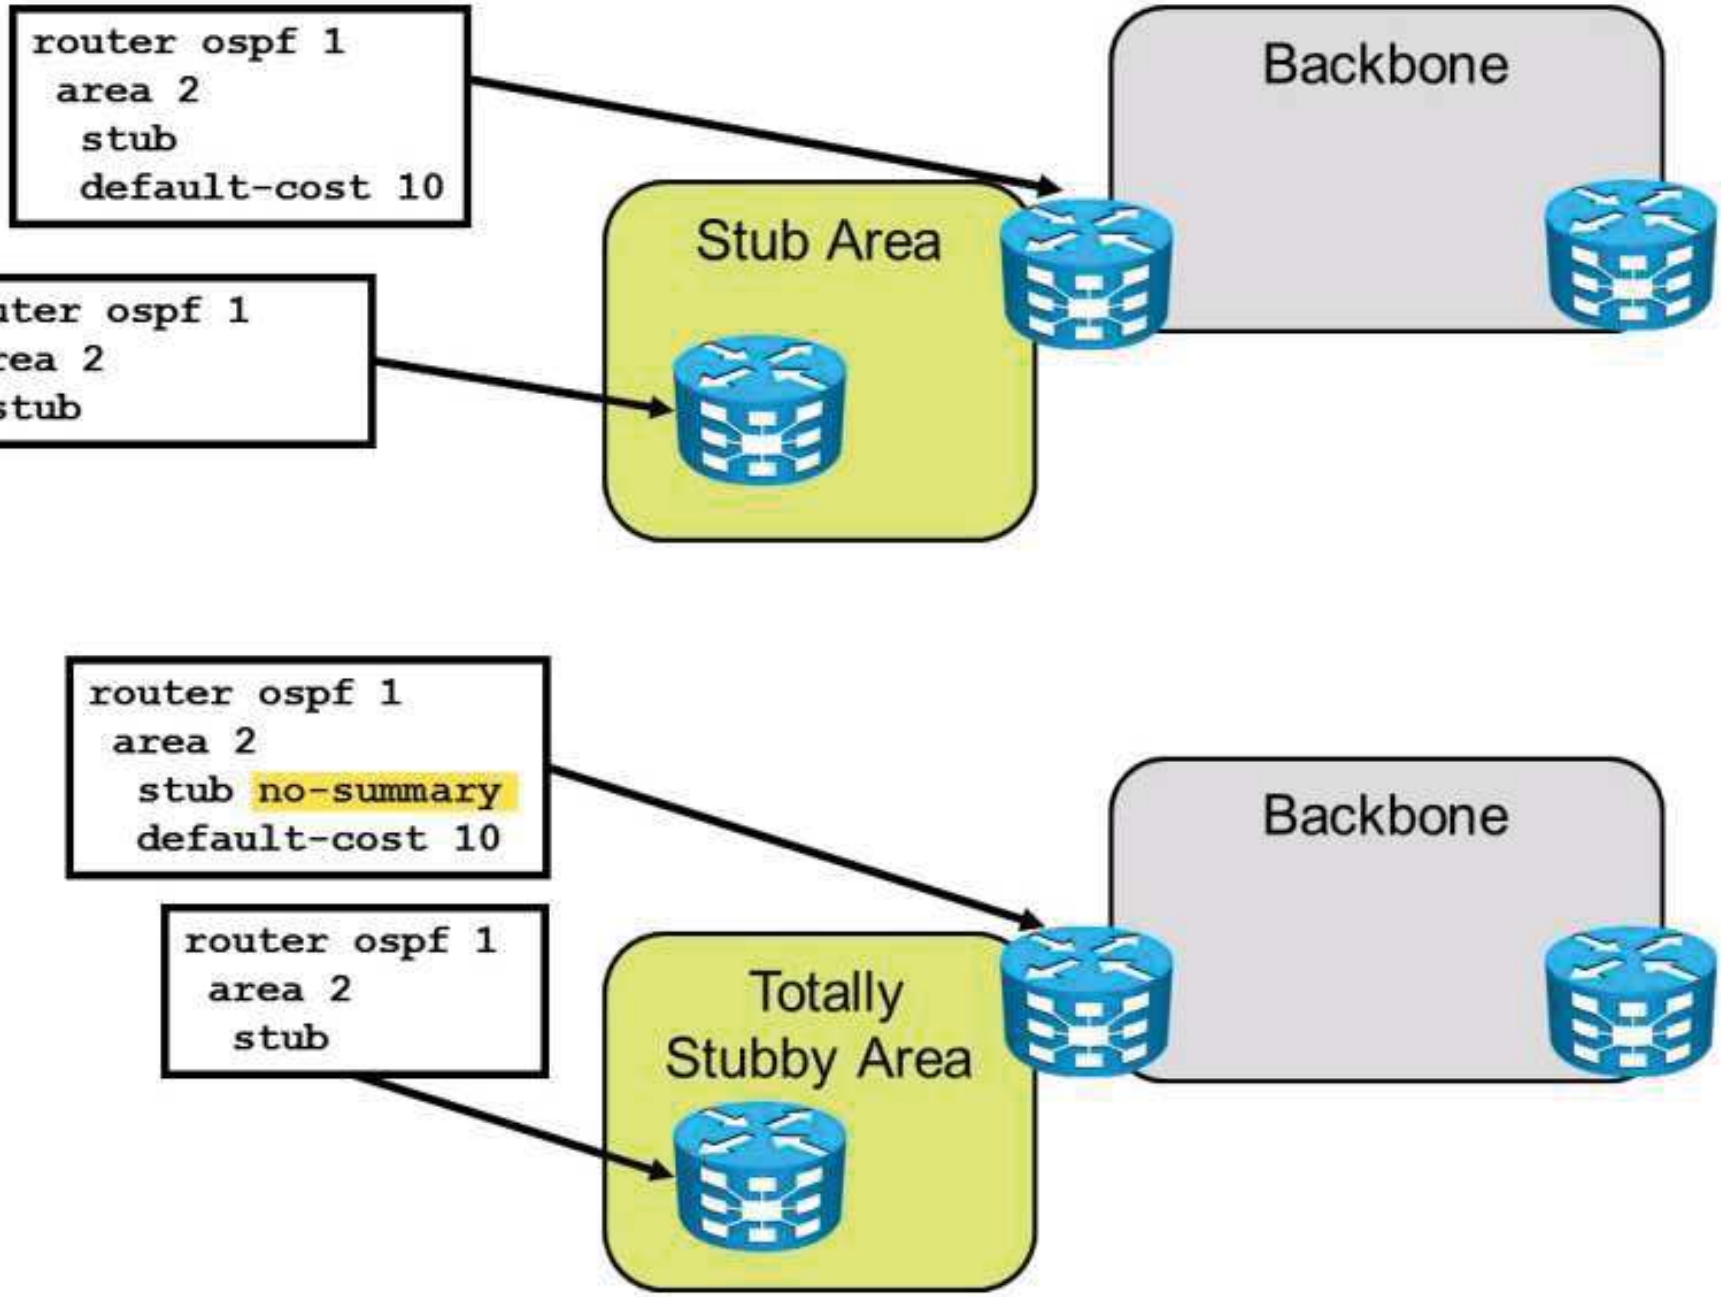

OSPF Not-So-Stubby Area and Totally Not-So-Stubby Area

An NSSA has the following characteristics:

- NSSA breaks stub area rules.
- ASBR is allowed inside.
- LSA type 7 sent by ASBR.
- ABR converts LSA type 7 to LSA type 5.
- ABR sends the default route into NSSA instead of external (LSA type 5) routes.

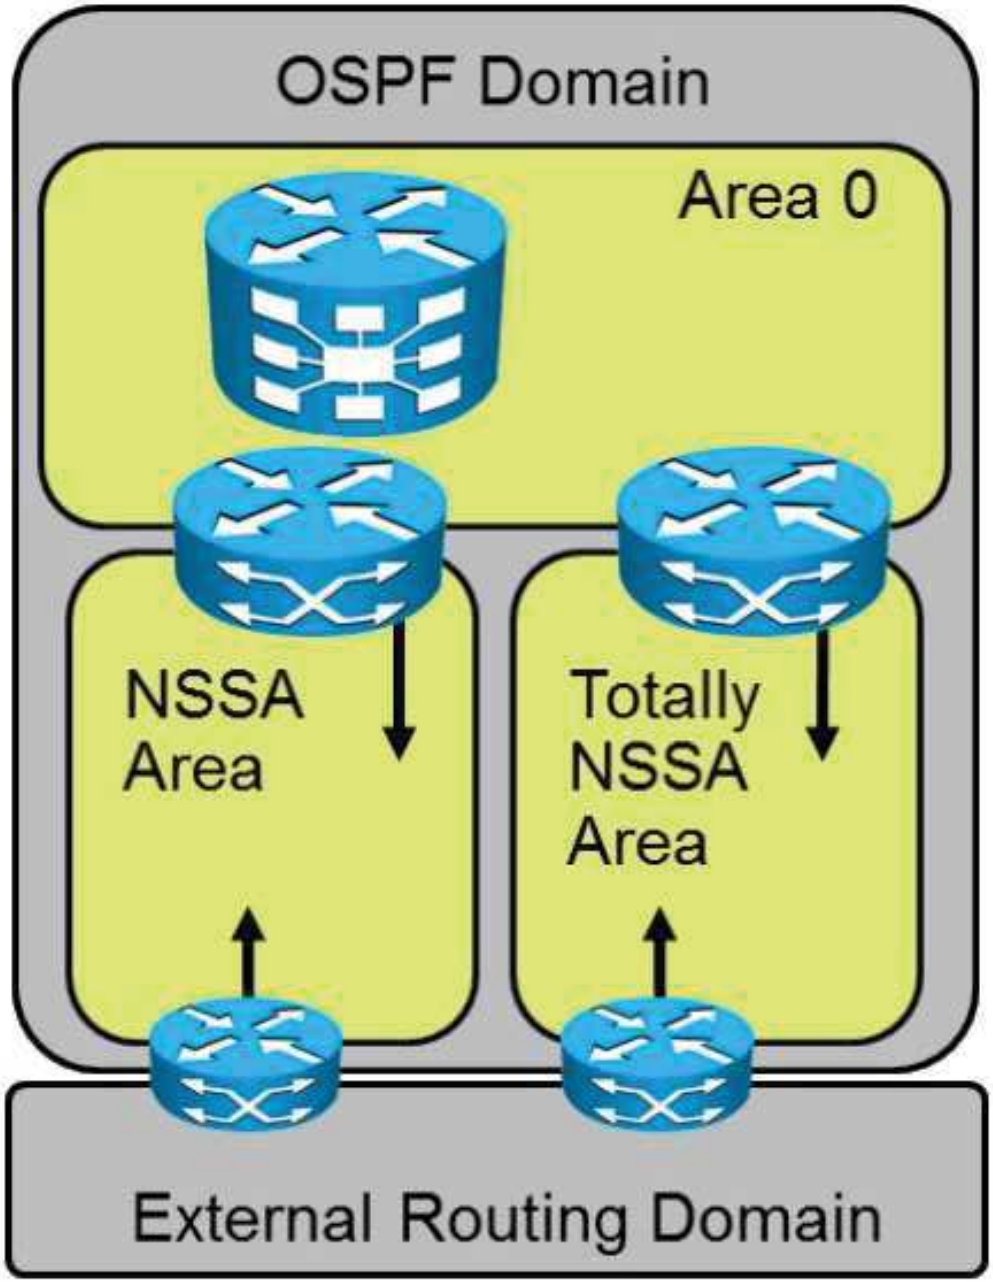

OSPF Not-So-Stubby Area and Totally Not-So-Stubby Area (Cont.)

- NSSA:
  - Behaves like stub area.
  - May introduce external routes locally in the area.
- Totally NSSA (**no-summary**):
  - Behaves like totally stubby area.
  - May introduce external routes locally in the area.
  - Cisco proprietary feature.

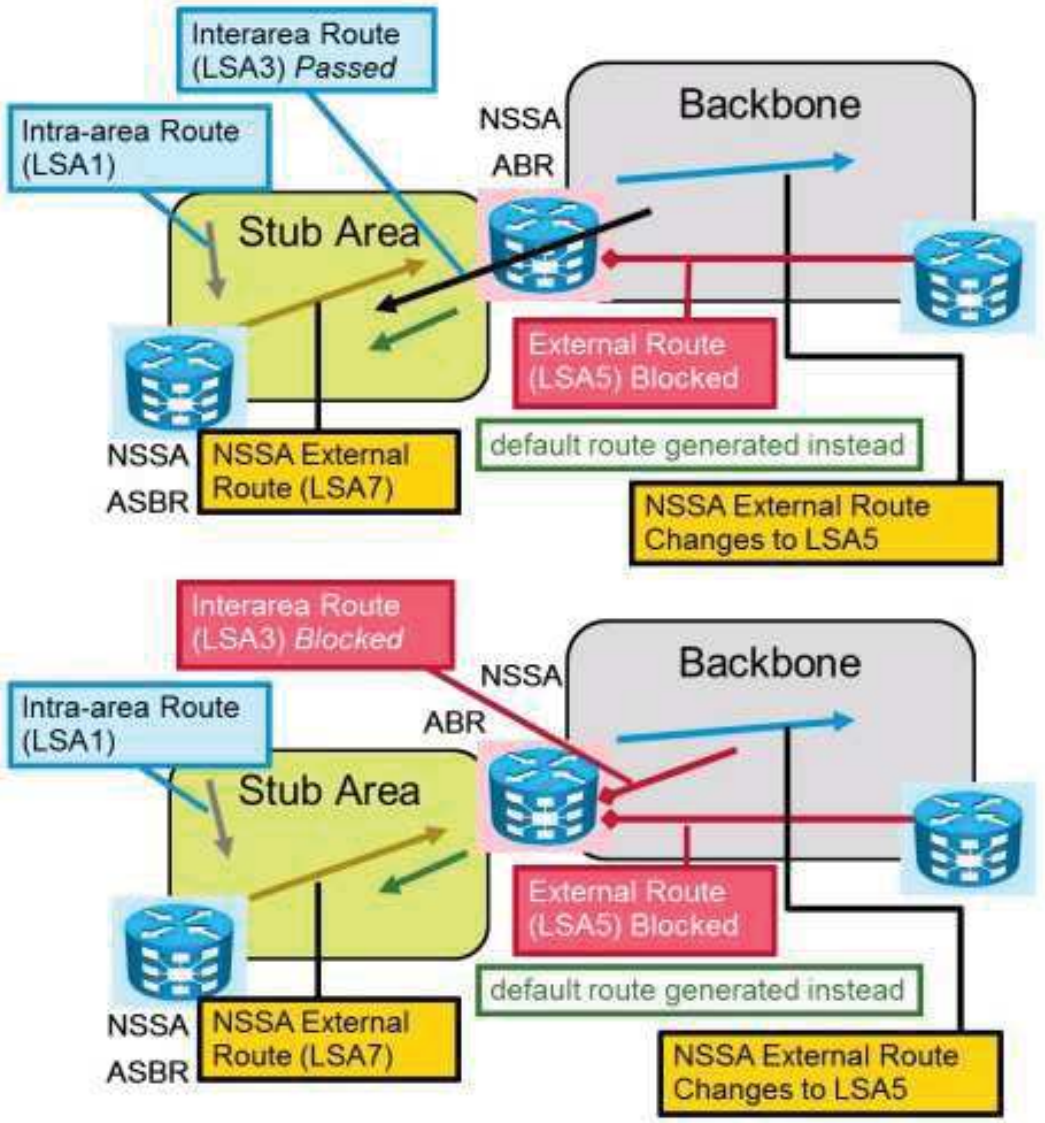

Configuring OSPF NSSA and Totally NSSA

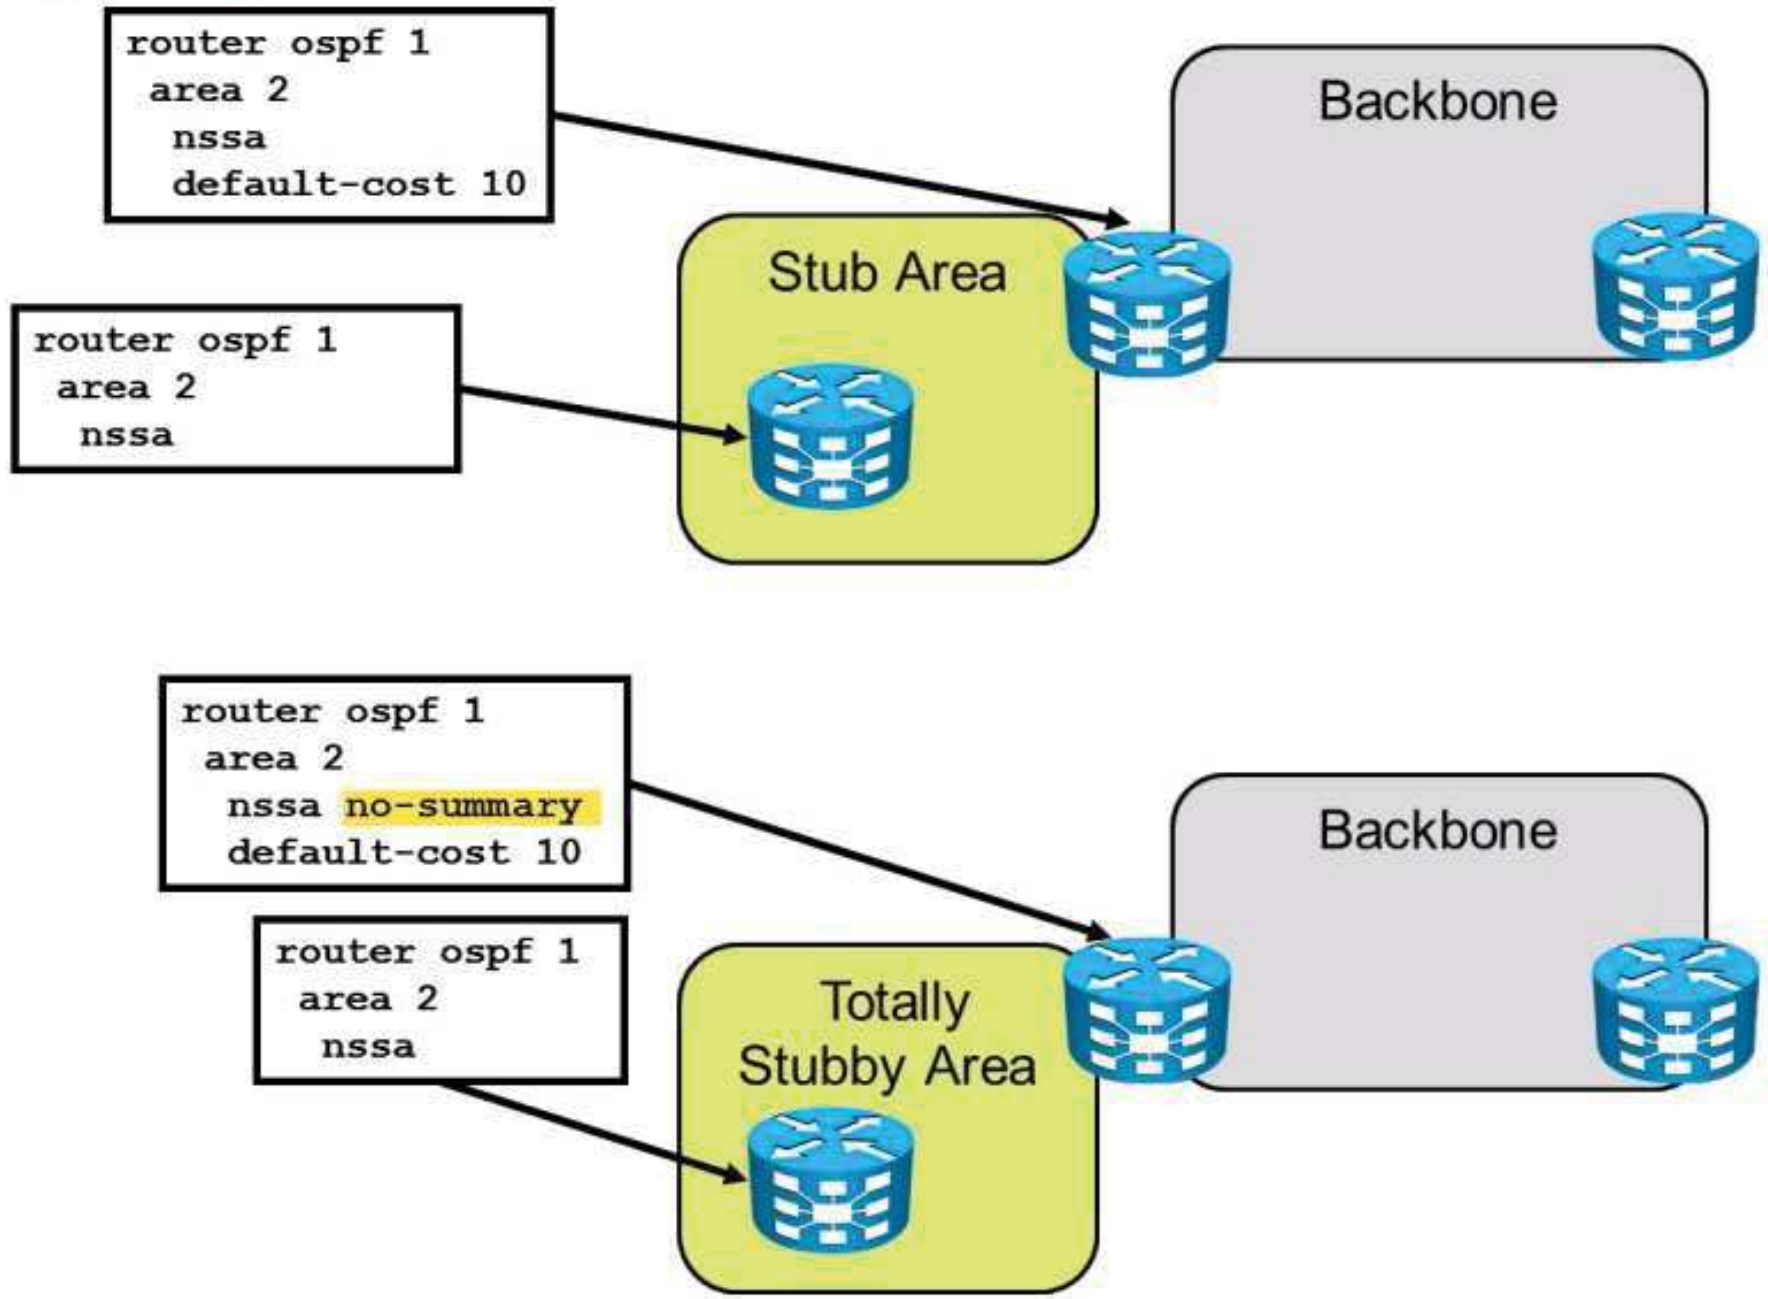

Summary

- Route summarization improves CPU utilization, reduces LSA flooding, and reduces routing table sizes.
- OSPF summarization of internal routes can be done only by ABRs.
- OSPF summarization of external routes can be done by ASBR or NSSA ABR.
- To generate a default external route into an OSPF routing domain, use the **default-information originate** command.
- There are several OSPF area types: normal, backbone, stub, totally stubby, NSSA, and totally NSSA.
- Use the **stub** Cisco IOS XR command to define an area as stubby and add the **no-summary** keyword on the ABR only to define an area as totally stubby.
- Use the **nssa** Cisco IOS XR command to define an area as stubby and add the **no-summary** keyword on the ABR only to define an area as totally NSSA.

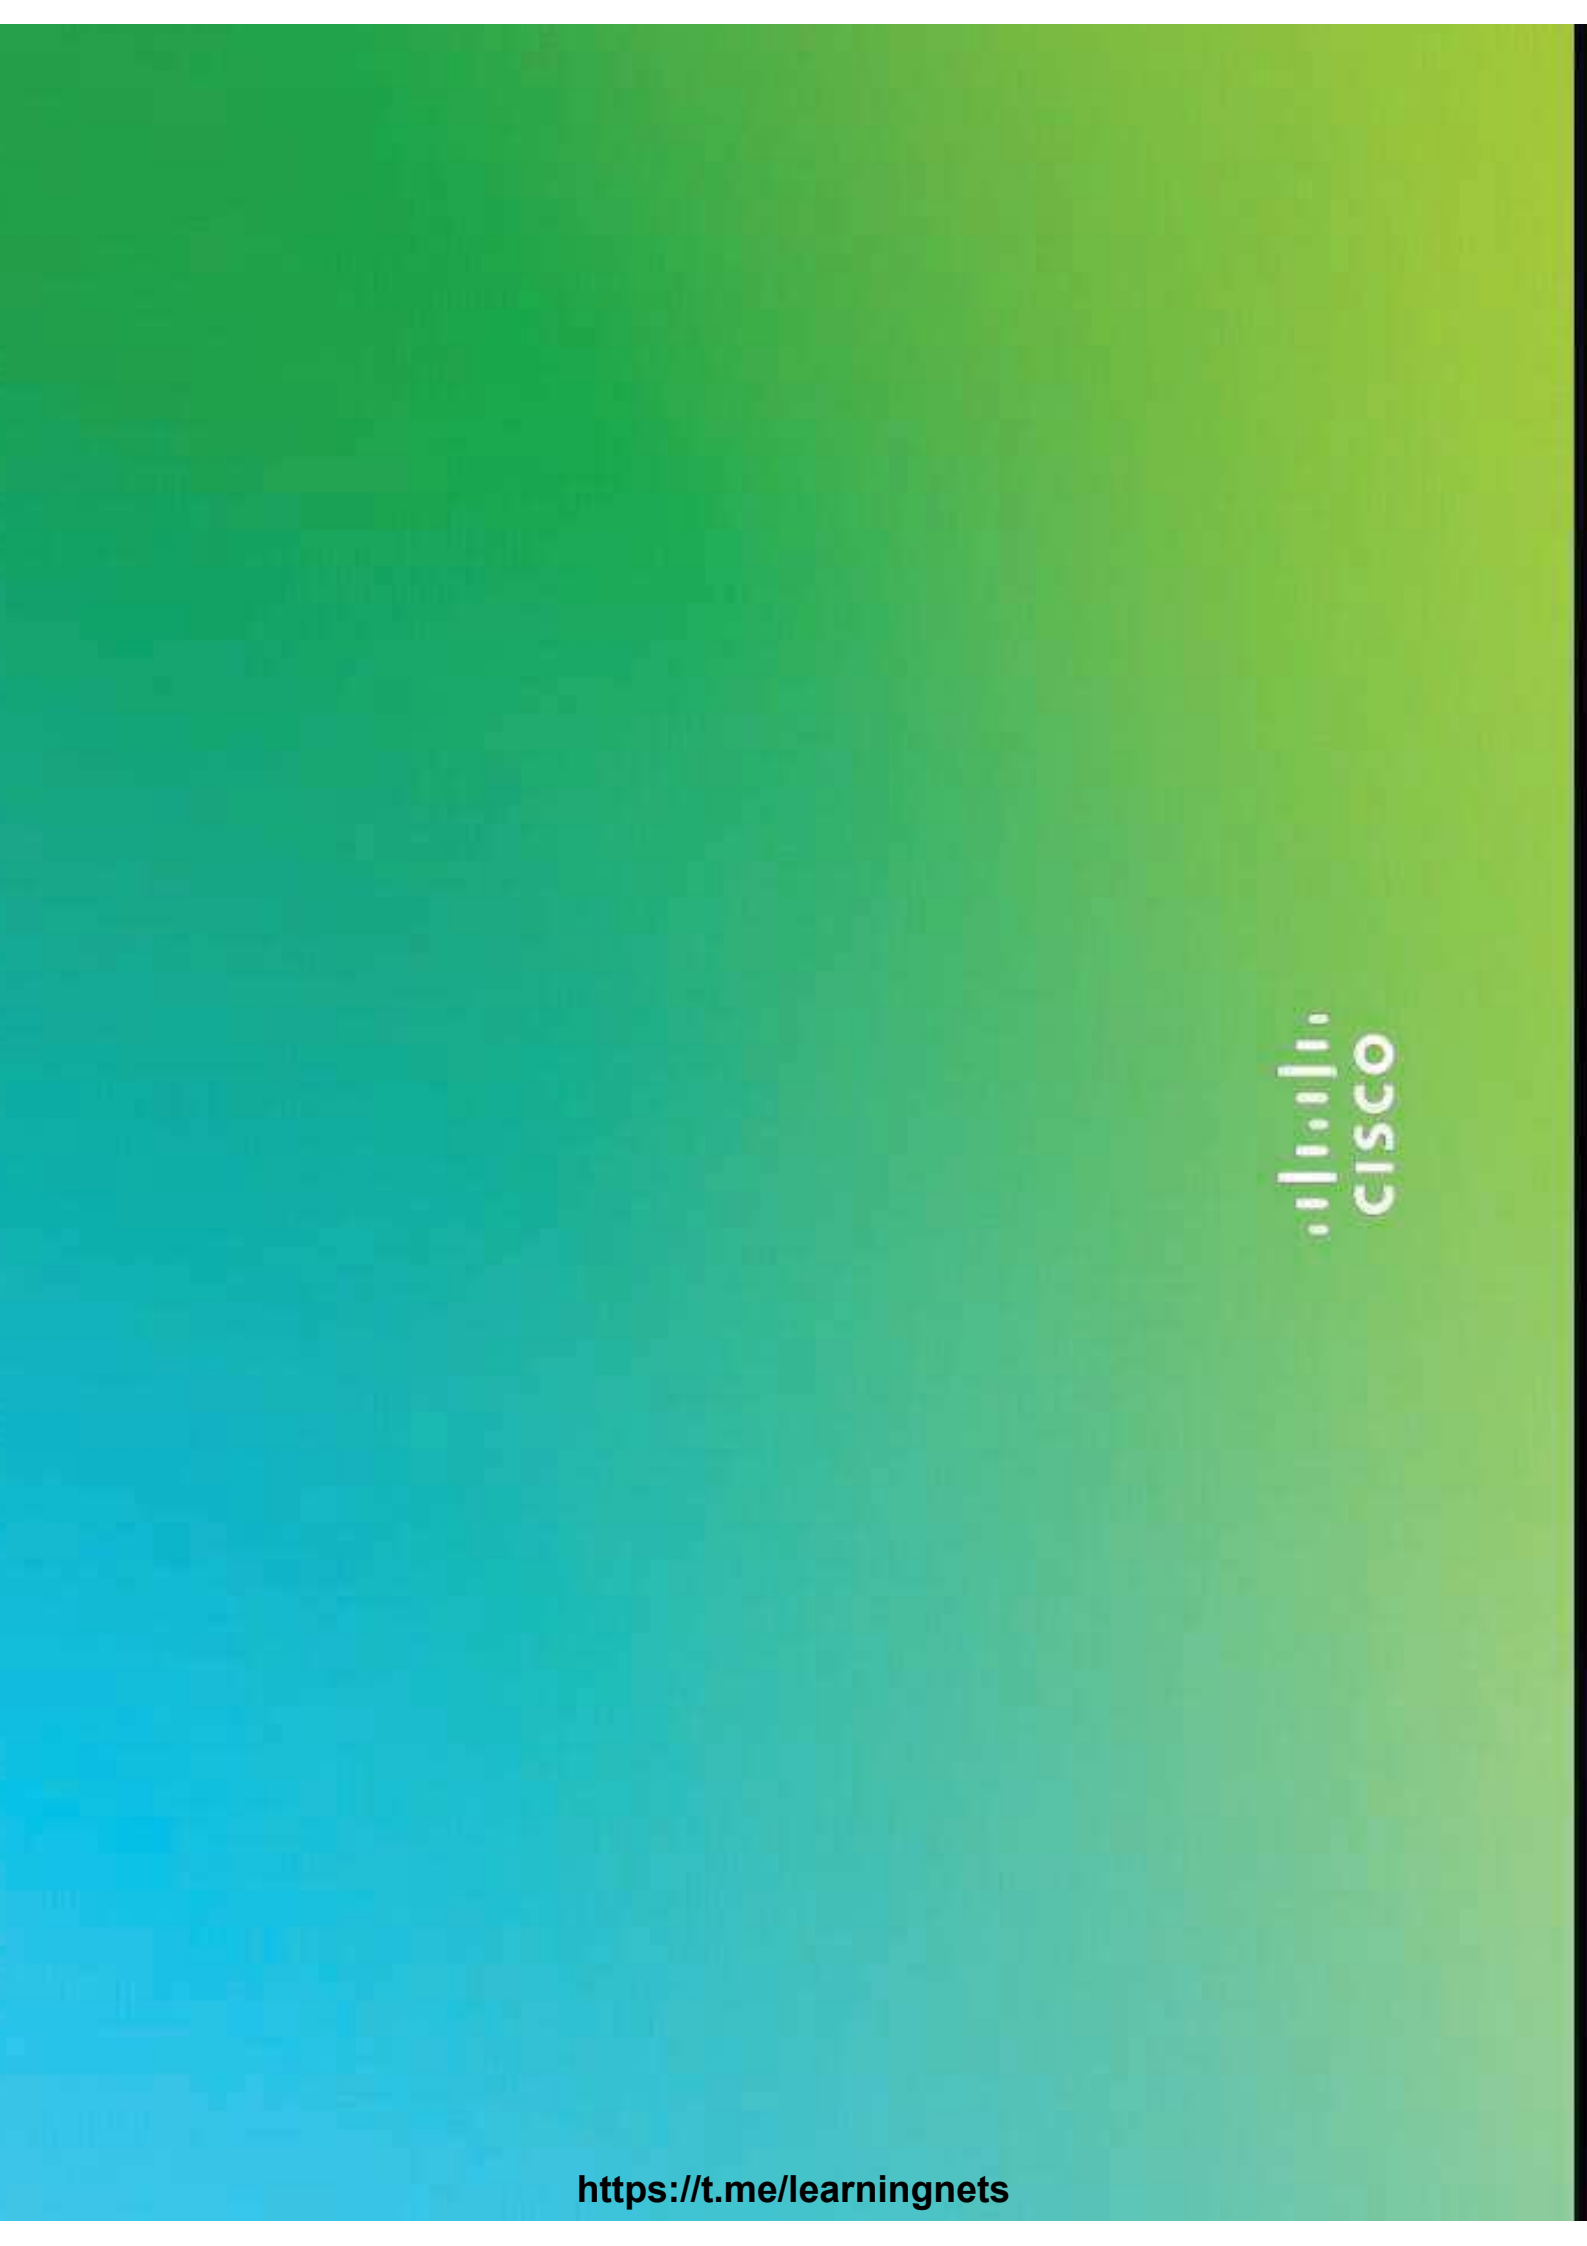

Module Summary

- The OSPF protocol is one of the most commonly used link-state IP routing protocols in service provider networks, and it is an open standard that offers quick convergence and the ability to scale large networks.
- OSPF uses five types of routing protocol packets and six common LSAs.
- The configuration of OSPF is a two-step process: enter the OSPF configuration and start OSPF on the interface.
- Route summarization reduces OSPF LSA flooding and the routing table size, which reduces memory and CPU utilization on routers. Stub area techniques improve OSPF performance by reducing the amount of LSA flooding.