CCNA 200-301 Volume I

Chapter 19
Understanding OSPF Concepts

Objectives

- Comparing Dynamic Routing Protocol Features
- OSPF Concepts and Operation
- OSPF Areas and LSAs

Comparing Dynamic Routing Protocol Features

- Routing protocol: A set of messages, rules, and algorithms used by routers for the overall purpose of learning routes. This process includes the exchange and analysis of routing information. Each router chooses the best route to each subnet (path selection) and finally places those best routes in its IP routing table. Examples include RIP, EIGRP, OSPF, and BGP.

- Routed protocol and routable protocol: Both terms refer to a protocol that defines a packet structure and logical addressing, allowing routers to forward or route the packets. Routers forward packets defined by routed and routable protocols. Examples include IP Version 4 (IPv4) and IP Version 6 (IPv6).

Routing Protocol Functions

1. Learn routing information about IP subnets from neighboring routers.
2. Advertise routing information about IP subnets to neighboring routers.
3. If more than one possible route exists to reach one subnet, pick the best route based on a metric.
4. If the network topology changes—for example, a link fails—react by advertising that some routes have failed and pick a new currently best route. (This process is called convergence.)

Routing Protocol Functions

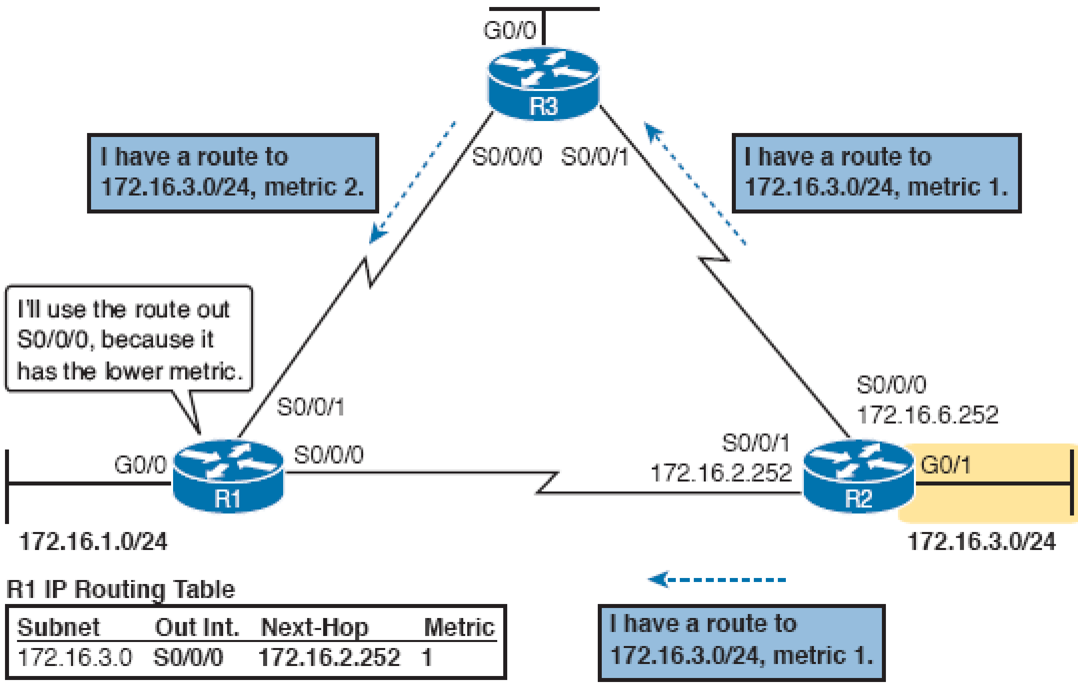

Interior and Exterior Routing Protocols

- There are two major categories of IP routing protocols:
  - IGP: A routing protocol that was designed and intended for use inside a single autonomous system (AS)
  - EGP: A routing protocol that was designed and intended for use between different autonomous systems
- An Autonomous System (AS) is a network under the administrative control of a single organization.
- Some routing protocols work best inside a single AS by design, these are called IGPs
- Other routing protocols work best between ASs by design, these are called EGPs

Comparing Locations for Using IGPs and EGPs

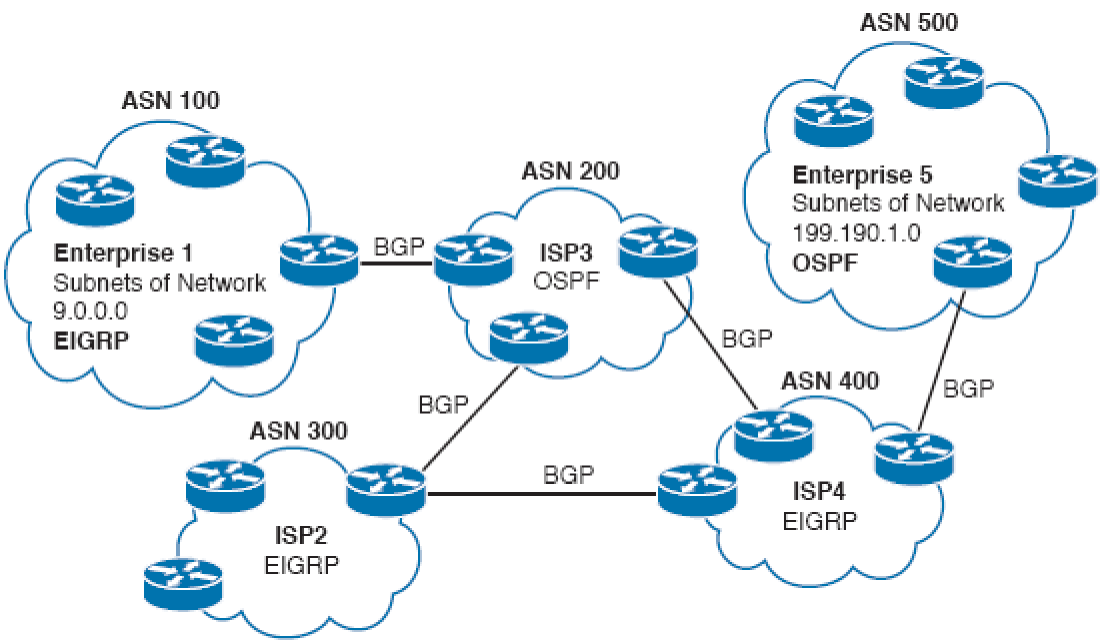

IGP Routing Protocol Algorithms

- There are three main branches of routing protocol algorithms:
  - Distance vector (sometimes called Bellman-Ford after its creators)
  - Advanced distance vector (sometimes called “balanced hybrid”)
  - Link-state
- Distance vector protocols were invented first (RIP and IGRP).
- Link state protocols (OSPF and IS-IS) and advanced distance vector protocols (EIGRP) came later.

IGP Routing Protocol Metrics

- Routing protocols choose the best route to reach a subnet by choosing the route with the lowest metric.

IGP	Metric	Description
RIPv2	Hop count	The number of routers (hops) between a router and the destination subnet
OSPF	Cost	The sum of all interface cost settings for all links in a route, with the cost defaulting to be based on interface bandwidth
EIGRP	Calculation based on bandwidth and delay	Calculated based on the route's slowest link and the cumulative delay associated with each interface in the route

IGP Routing Protocol Algorithms

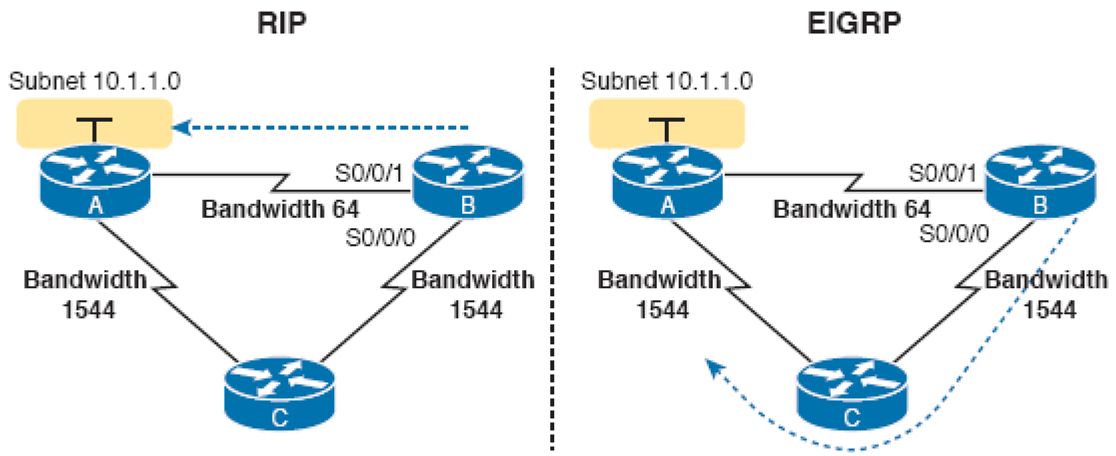

- This example shows why OSPF/EIGRP surpassed RIP; it shows Router B with two possible routes to subnet 10.1.1.0:
  - A shorter route over a very slow 64-Kbps link
  - A longer route over two higher-speed (T1) links
- With EIGRP selecting the better choice even though it goes through multiple routers.

Other IGP Comparisons

Feature	RIPv1	RIPv2	EIGRP	OSPF	IS-IS
Classless/sends mask in updates/supports VLSM	No	Yes	Yes	Yes	Yes
Algorithm (DV, advanced DV, Link State)	DV	DV	adv. DV	LS	LS
Supports manual summarization	No	Yes	Yes	Yes	Yes
Cisco-proprietary	No	No	Yes	No	No
Routing updates are sent to a multicast IP address	No	Yes	Yes	Yes	--
Convergence	Slow	Slow	Fast	Fast	Fast

Administrative Distance

- When two different routing protocols learn routes to the same subnet, IOS cannot compare the metrics.
- When IOS must choose between routes learned using different routing protocols it uses a concept called *administrative distance.*
- Administrative distance is a number that denotes how believable an entire routing protocol is, on a single router.

Administrative Distance

Route Type	Administrative Distance
Connected	0
Static	1
BGP (External routes)	20
EIGRP (Internal routes)	90
IGRP	100
OSPF	110
IS-IS	110
RIP	120
EIGRP (External routes)	170
BGP (Internal routes)	200
DHCP default route	254
Unusable	255

OSPF Overview

- Link State protocols build IP routes with a couple of major steps:
  - The routers learn a lot of information about the network (Routers, links, IP addresses, status information, etc.)
  - The routers then flood this information so that all routers know the same information.
  - Each router can then calculate routes to all subnets from their own perspectives.

Topology Information and LSAs

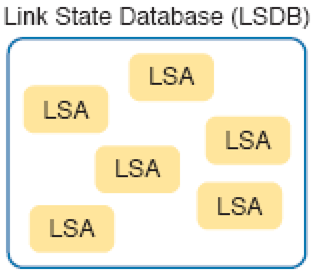

- OSPF organizes topology information using LSAs and the link-state database (LSDB).
- Each LSA is a data structure with some specific information about the network topology.
- The LSDB is a collection of all the LSAs known to a router.

Topology Information and LSAs

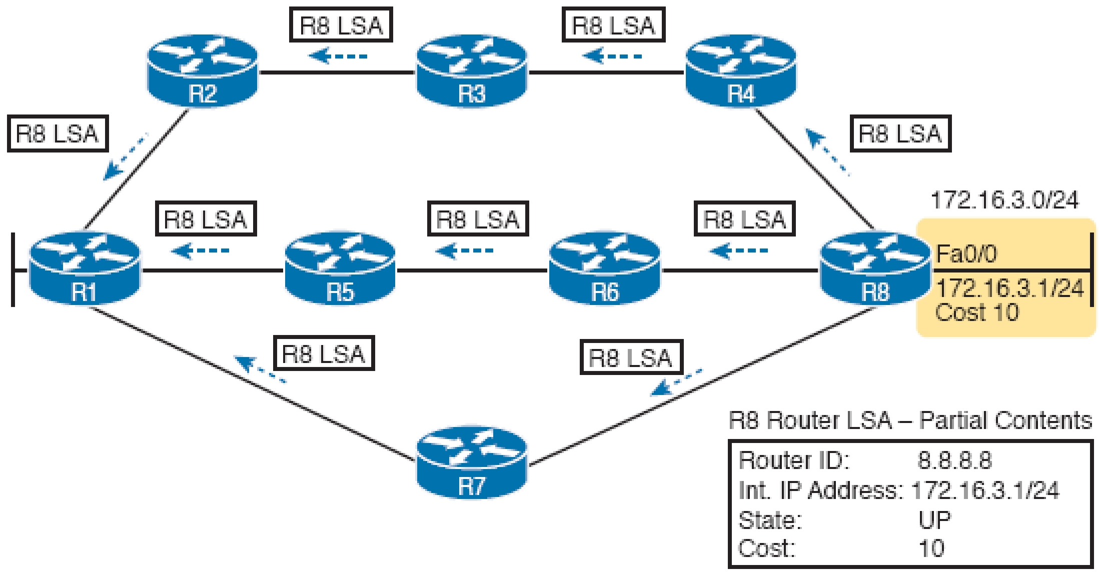

- The flooding process has a way to prevent loops so that the LSAs do not get flooded around in circles:
  - Before sending an LSA to yet another neighbor, routers communicate which LSAs they already have.
  - Only LSAs that are not known are flooded.
- Routers reflood LSA information based on the LSAs separate aging timer (default 30 minutes).

Applying Dijkstra Math to Find the Best Routes

- To build routes, link state routes have to do some math based on the topological information found in the LSDB.
- All link state protocols use a math algorithm called Dijkstra Shortest Path First (SPF) to process the LSDB.
- The algorithm analyzes the LSDB and builds the routes that the local route will add to the IP routing table.

Becoming OSPF Neighbors

- OSPF neighbors are routers that both use OSPF and both sit on the same data link.
- OSPF routers introduce themselves by sending Hello messages.
- Assuming the two neighbors have compatible OSPF parameters, the two will form a neighbor relationship.
- The OSPF neighbor relationship lets OSPF know when a neighbor might not be a good option for routing packets.
- The OSPF neighbor model allows new routers to be dynamically discovered.

Meeting Neighbors and Learning Their Router-ID

- When OSPF exchanges Hello messages they list each router's *router-id* (RID).
- This RID serves as each router's unique name or identifier for OSPF.
- OSPF RIDs are 32-bit numbers.
- IOS chooses its OSPF RID based on an active interface IPv4 address.
- The OSPF RID can also be directly configured.
- As soon as a router has chosen its OSPF RIS and some interfaces come up, the router is ready to meet its OSPF neighbors.

Meeting Neighbors and Learning Their Router-ID

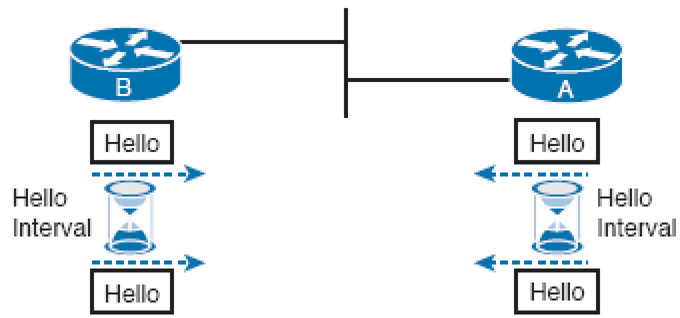

- To discover other OSPF-speaking routers, a router sends multicast OSPF Hello packets to each interface and hopes to receive OSPF hello packets from other routers connected to those interfaces.

Meeting Neighbors and Learning Their Router-ID

- The Hello message follows the IP packet header, with IP protocol type 89.
- Hello packets are sent to multicast IP address 224.0.0.5, a multicast IP address intended for all OSPF-speaking routers.
- OSPF routers listen for packets sent to IP multicast address 224.0.0.5, in part hoping to receive Hello packets and learn about new neighbors.

Early Neighbor States

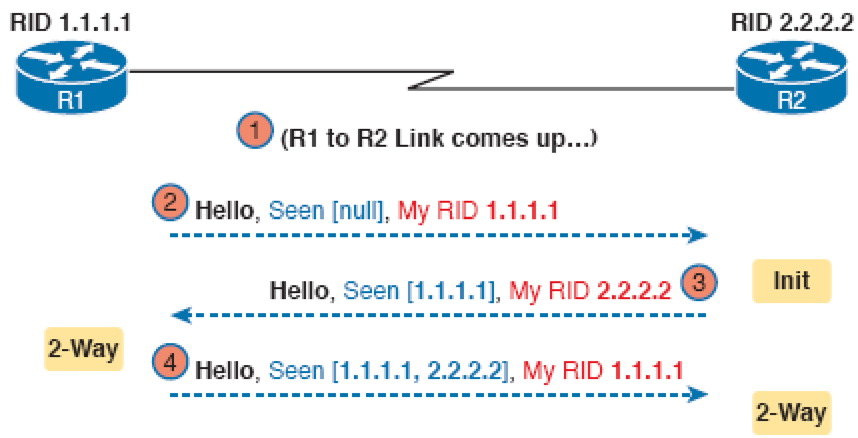

- The 2-way state is a particularly important OSPF state, when in this state the following statements are true:
  - The router received a Hello from the neighbor, with that router's own RID listed as being seen by the neighbor.
  - The router has checked all the parameters in the Hello received from the neighbor, with no problems. The router is willing to become a neighbor.
  - If both routers reach a 2-way state with each other, it means that both routers meet all OSPF configuration requirements to become neighbors. Effectively, at that point, they are neighbors, and ready to exchange their LSDB with each other.

Fully Exchanging LSAs with Neighbors

- After two routers decide to exchange databases, they do not simply send the contents of their entire database:
    - They tell each other a list of LSAs that are in their respective databases,
    - Each router checks to see which LSAs it already has,
    - Each router then requests only those LSAs that it does not know about yet,
- OSPF messages that actually send the LSAs between neighbors are called Link-State Update (LSU) packets.
- Each LSU packet holds data structures called Link State Advertisements (LSA).

OSPF Database Exchange

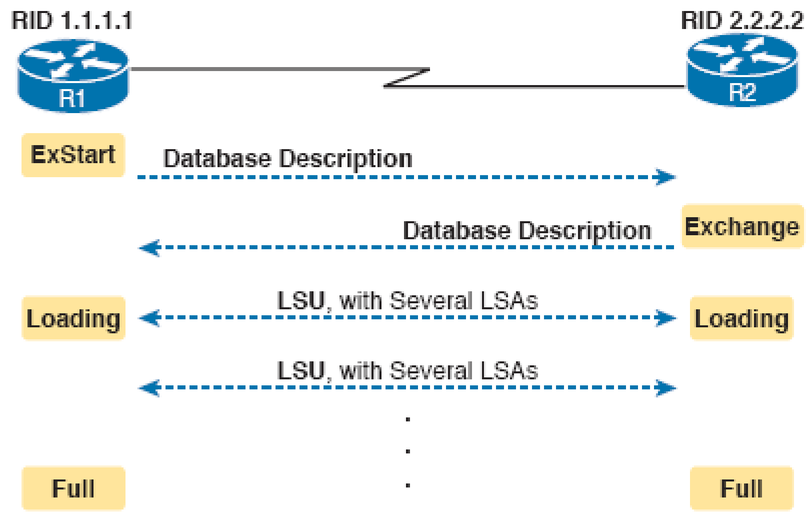

- When finished, the routers reach a full state, meaning they have fully exchanged the contents of their LSDBs

Maintaining Neighbors and the LSDB

- Once two neighbors reach a full state, they have done all the initial work to exchange OSPF information.

- Neighbors still have to do some small ongoing tasks to maintain their neighbor relationship:
  - Routers monitor each neighbor relationship using Hello messages and two related timers: Hello Interval and Dead Interval.
  - Routers send and expect to receive a Hello message from each neighbor based on the Hello interval.
  - If a neighbor is silent for the length of the Dead Interval, the router assumes the neighbor has failed.
  - Routers must also be able to react when the topology changes.

OSPF Maintenance Task Summary

- Maintain neighbor state by sending Hello messages based on the Hello Interval, and listening for Hellos before the Dead Interval expires.
- Flood any changed LSAs to each neighbor.
- Reflood unchanged LSAs as their lifetime expires (default 30 minutes).

Designated Routers on Ethernet Links

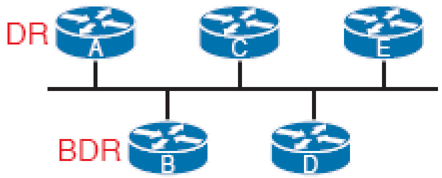

- On Ethernet links, OSPF elects one of the routers on the same subnet to act as the *designated router* (DR).
- The DR plays a key role in how the database exchange process works.

Designated Routers on Ethernet Links

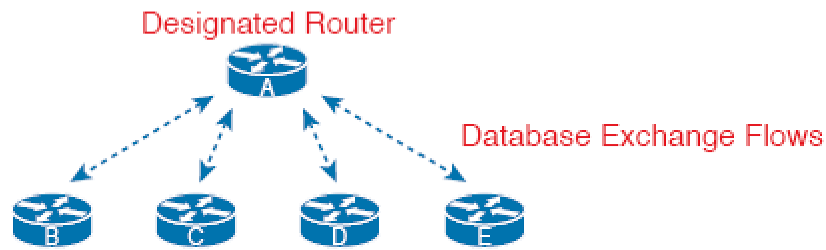

- The database exchange process on an Ethernet link does not happen between every pair of routers on the same VLAN/subnet.
- The database exchange happens between the DR and each of the other routers.
- The *backup designated* router (BDR) watches the status of the DR and takes over for the DR if it fails.
- The DR and BDR both do full database exchanges with all other routers on the LAN, they both reach a full state with all neighbors.

Stable OSPF Neighbor States and Their Meanings

Neighbor State	Term for Neighbor	Term for Relationship
2-way	Neighbor	Neighbor Relationship
Full	Adjacent Neighbor Fully Adjacent Neighbor	Adjacency

Calculating the Best Routes with SPF

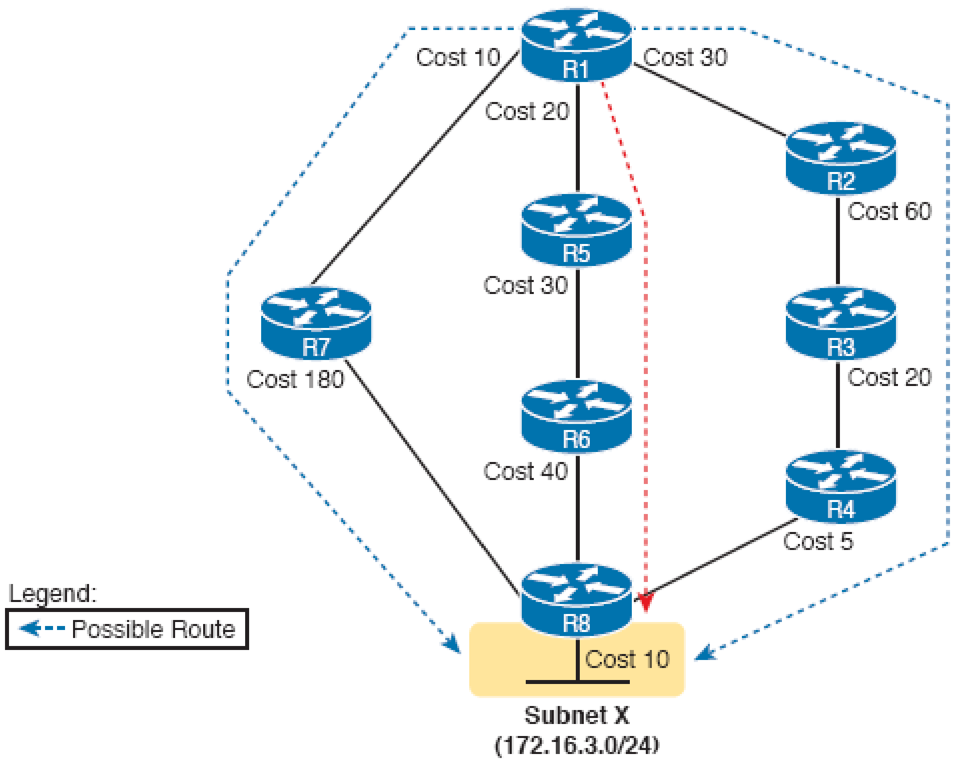

- Once SPF has identified a route, OSPF calculates the metric for a route as follows:
  - The sum of the OSPF interface costs for all outgoing interfaces in the route.

OSPF Area Design

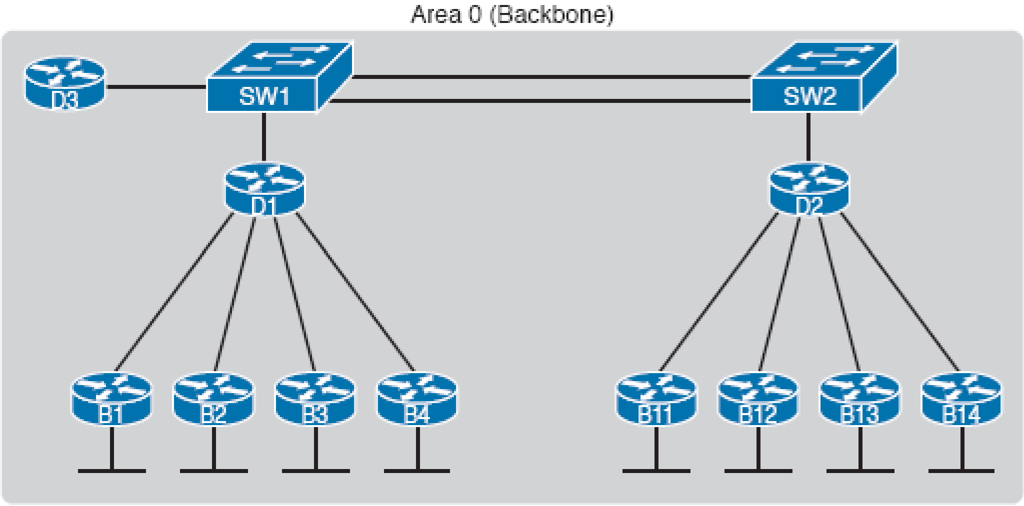

- Larger OSPFv2 networks suffer when using a single area design:
  - A larger topology database requires more memory on each router.
  - Processing the larger topology database with the SPF algorithm requires processing power that grows exponentially with the size of the topology database.
  - A single interface status change, anywhere in the internetwork (up to down, or down to up), forces every router to run SPF again!

OSPF Area Design

- The solution is to take one large LSDB and break it into several smaller LSDBs by using OSPF areas.
- With areas each link is placed into one area.
- SPF does its complicated match on the topology inside the areas and that area’s topology alone.
- Generally, networks larger then a few dozen routers benefit from areas.

OSPF Areas

- OSPF area design follows a couple of basic rules:
  - Put all interfaces connected to the same subnet inside the same area.
  - An area should be contiguous.
  - Some routers may be internal to an area, with all interfaces assigned to that single area.
  - Some routers may be Area Border Routers (ABR), because some interfaces connect to the backbone area, and some connect to nonbackbone areas.
  - All nonbackbone areas must connect to the backbone area (area 0) by having at least one ABR connected to both the backbone area and the nonbackbone area.

OSPF Areas

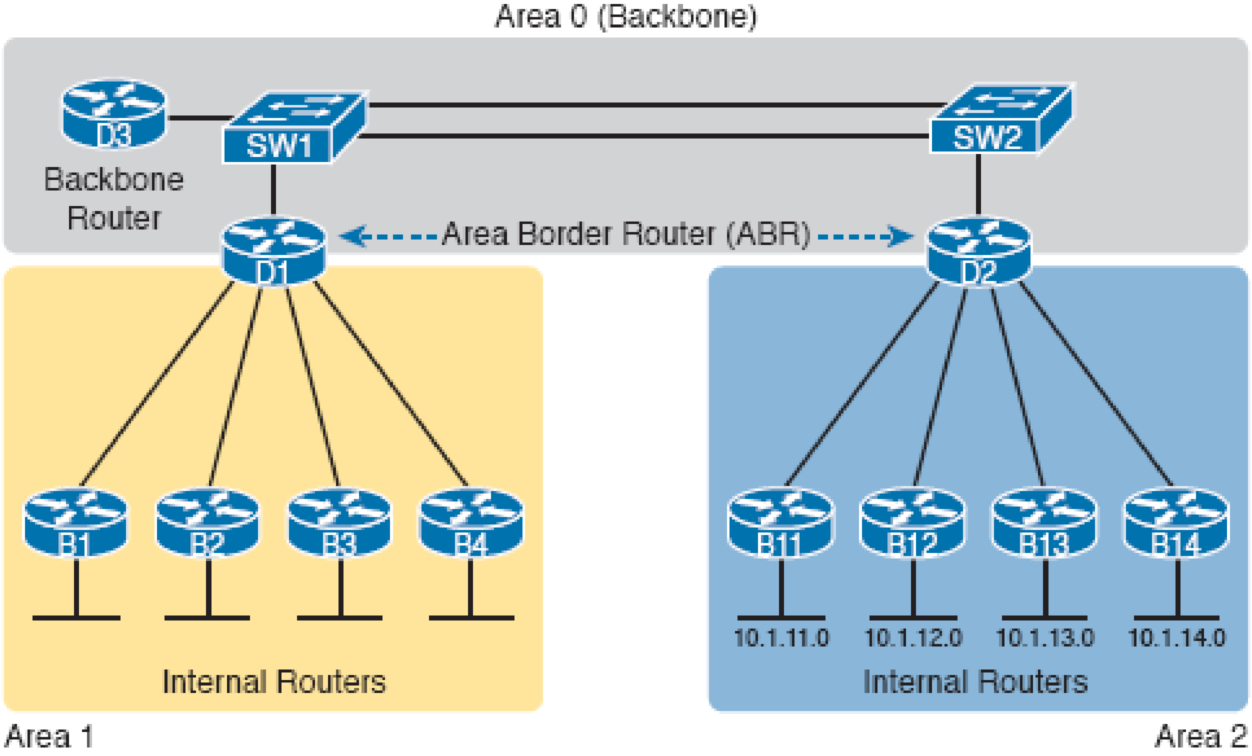

Area 0 (Backbone)
D3
Backbone Router
SW1
SW2
D1
Area Border Router (ABR)
D2
B1
B2
B3
B4
Internal Routers
B11
B12
B13
B14
10.1.11.0
10.1.12.0
10.1.13.0
10.1.14.0
Internal Routers
Area 1
Area 2

OSPF Design Terminology

Term	Description
Area Border Router (ABR)	An OSPF router with interfaces connected to the backbone area and to at least one other area
Backbone Router	A router connected to the backbone area (includes ABRs)
Internal Router	A router in one area (not the backbone area)
Area	A set of routers and links that share the same detailed LSDB information, but not with routers in other areas, for better efficiency
Backbone Area	A special OSPF area to which all other areas must connect—area 0
Intra-area route	A route to a subnet inside the same area as the router
Inter-area route	A route to a subnet in an area of which the router is not a part

How Areas Reduce SPF Calculation Time

- SPF spends most of the its processing time working through all the topology details.
- Areas reduce this workload because the LSDB only lists those routers and links inside that area.
- While the LSDB has less topology information, it still has to have information about all subnets in all areas.
- OSPF uses very brief summary information about the subnets to other areas.
- These LSAs do not include topology information about the other areas.

How Areas Reduce SPF Calculation Time

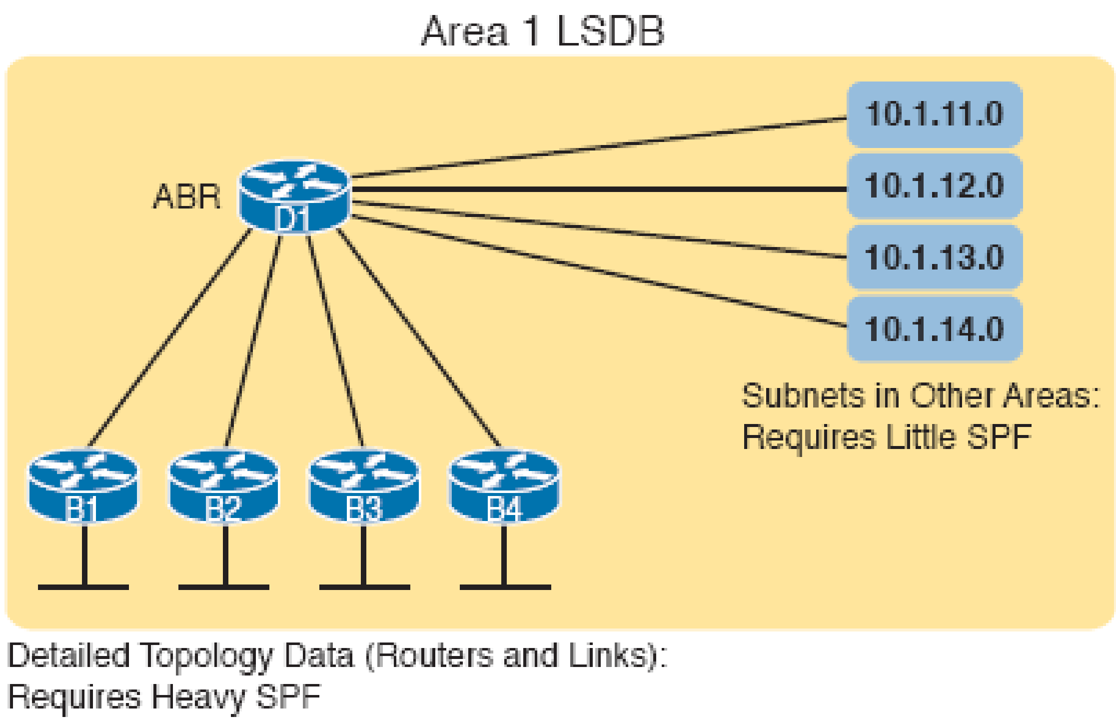

OSPF Area Design Advantages

- The smaller per-area LSDB requires less memory.
- Routers require fewer CPU cycles to process the smaller per-area LSDB with the SPF algorithm, reducing CPU overhead and improving convergence time.
- Changes in the network (for example, links failing and recovering) require SPF calculations only on routers connected to the area where the link changed state, reducing the number of routers that must rerun SPF.
- Less information must be advertised between areas, reducing the bandwidth required to send LSAs.

Enterprise Network with Seven IPv4 Subnets

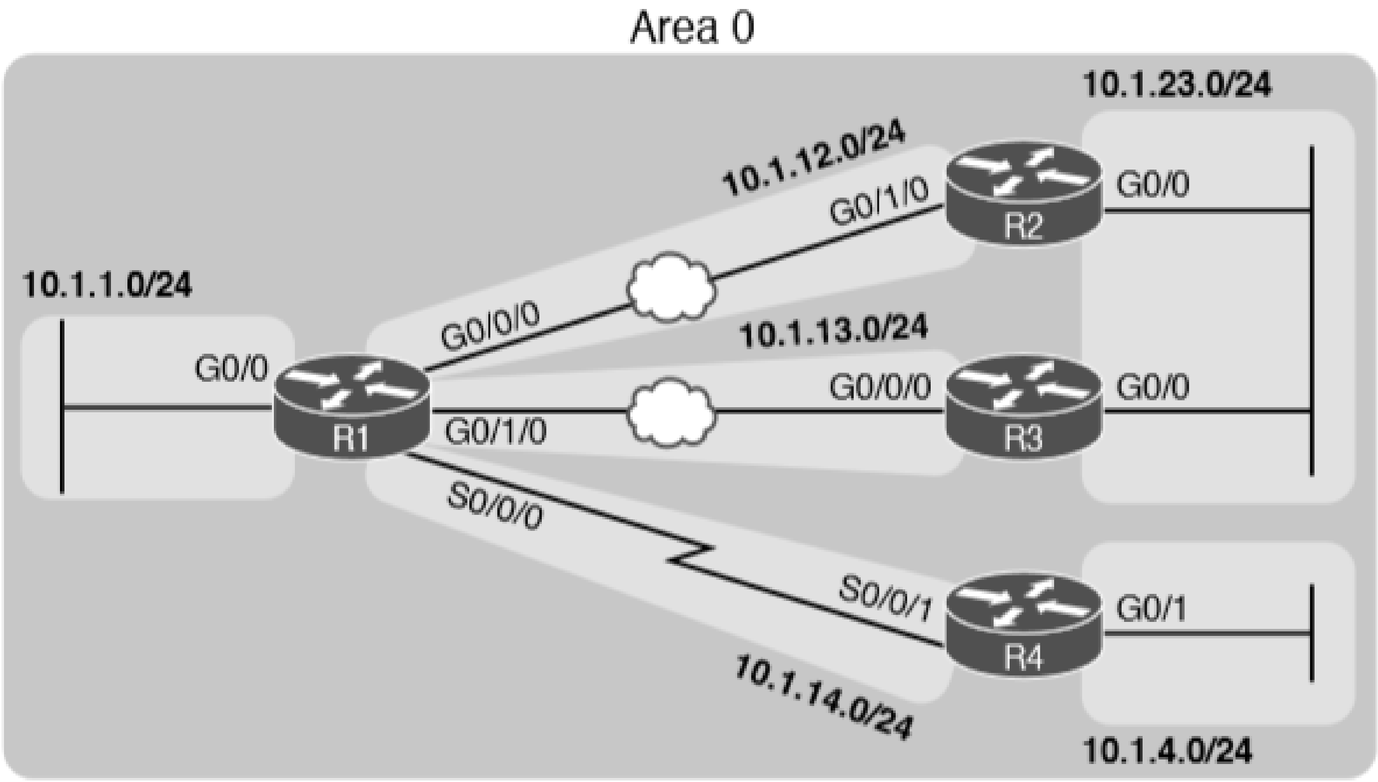

Type 1 LSAs, Assuming a Single-Area Design

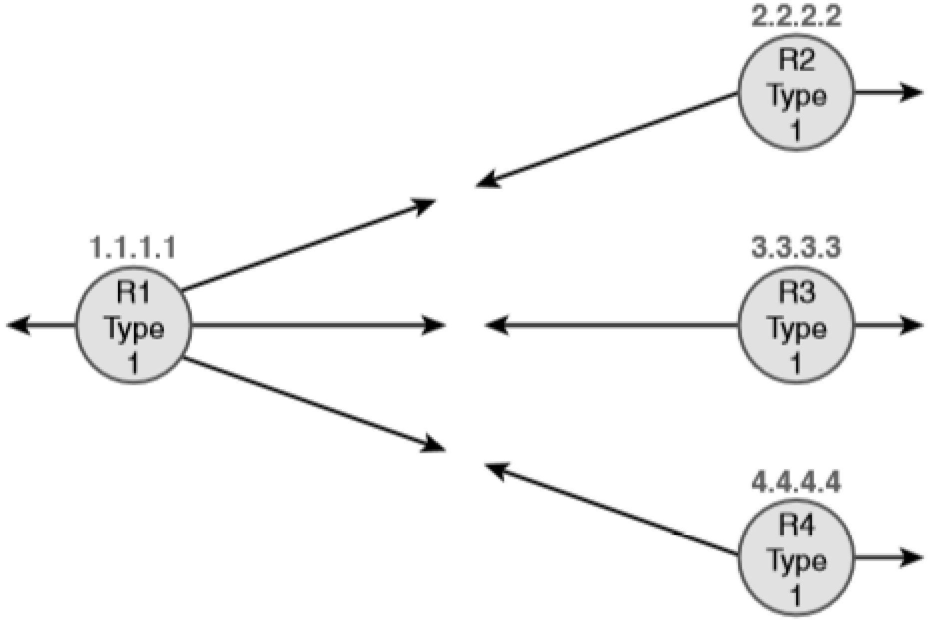

Type 1 and Type 2 LSAs in Area 0, Assuming a Single-Area Design