Routing Basic

Nazrul Islam

nazrul13@gmail.com

Routing

- Forwarding of packets from one network to another network
- Choosing the best path from the routing table
- Best path selection is based on the type of routing we are using (static/Dynamic)

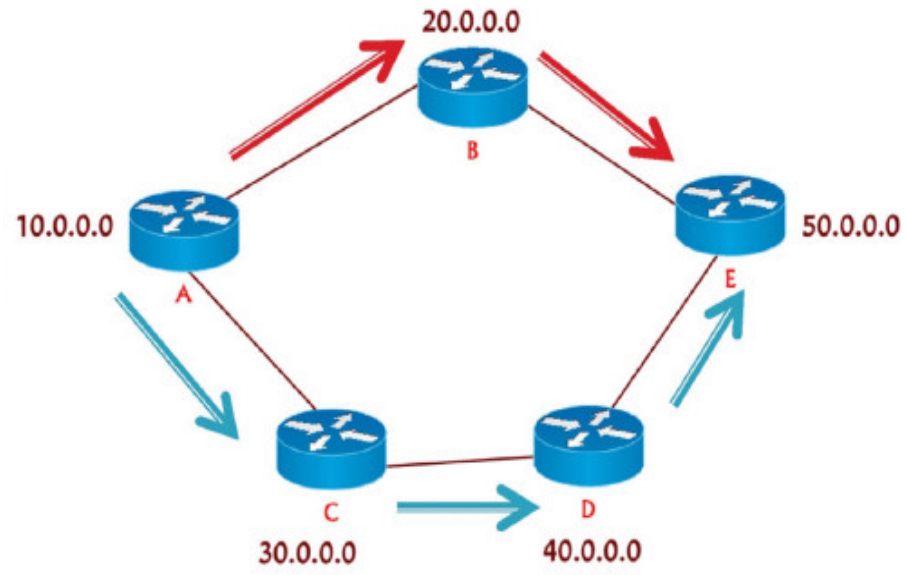

Types of Routing

1. Static Routing
2. Default Routing
3. Dynamic Routing

Static Routing:

- Best path is configured manually by Administrator
- Mandatory need of Destination Network ID
- It is Secure & fast

- Disadvantages
  - Everything to manually
  - Used for small network
  - Network change effect complete network

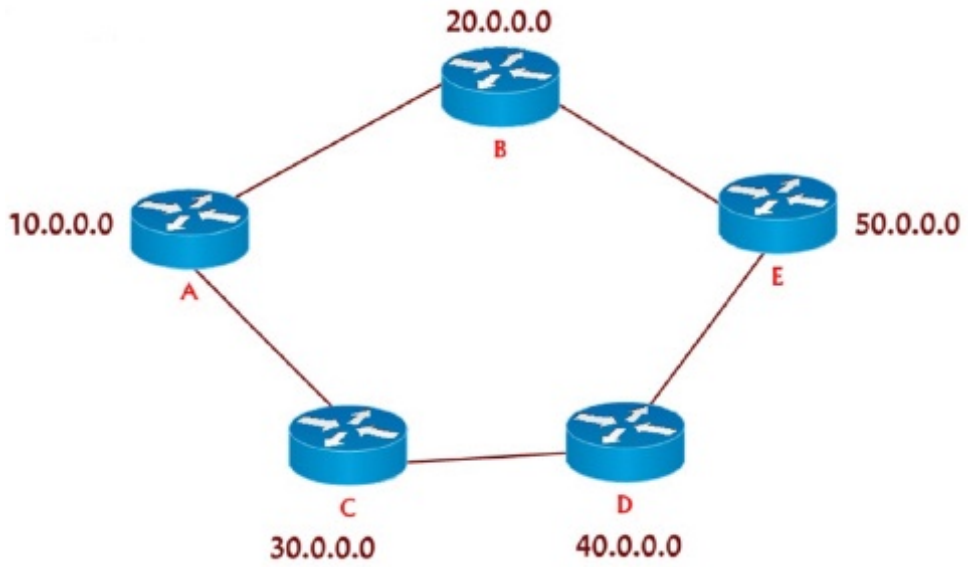

Configuration (Static Route)

- Router(config)#

ip route < Destination Network ID> < Destination Subnet Mask> < Next-hop IP Address>

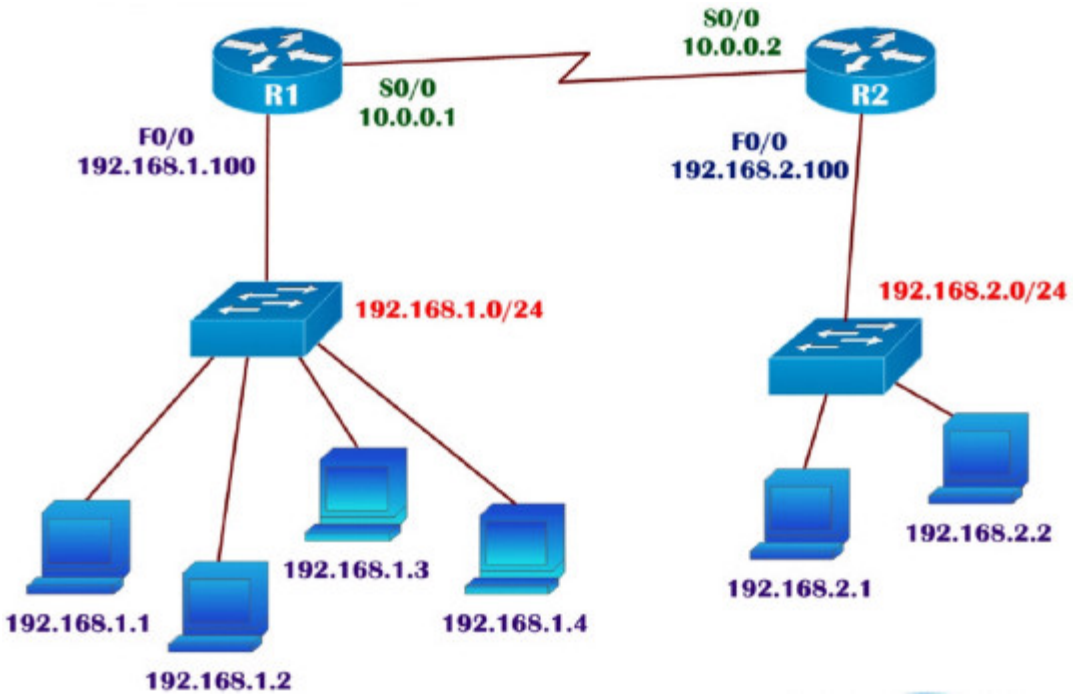

Verification before Static Routing

```
R1# show ip route
C 10.0.0.0/8 is directly connected, Serial0/0
C 192.168.1.0/24 is directly connected, FastEthernet0/0

PC> ping 192.168.2.1
Pinging 192.168.2.1 with 32 bytes of data:

Reply from 192.168.1.100: Destination host unreachable.
Reply from 192.168.1.100: Destination host unreachable.
Reply from 192.168.1.100: Destination host unreachable.
Reply from 192.168.1.100: Destination host unreachable.
```

Static Routing

```
R1(config)# ip route 192.168.2.0 255.255.255.0 10.0.0.2
R2(config)# ip route 192.168.1.0 255.255.255.0 10.0.0.1
```

Verification:

```
R1# show ip route
Gateway of last resort is not set
C 10.0.0.0/8 is directly connected, Serial0/0
C 192.168.1.0/24 is directly connected, FastEthernet0/0
S 192.168.2.0/24[1/0] via 10.0.0.2
```

```
R1# show ip route
Gateway of last resort is not set
C 10.0.0.0/8 is directly connected, Serial0/0
C 192.168.1.0/24 is directly connected, FastEthernet0/0
S 192.168.1.0/24[1/0] via 10.0.0.2

PC> ping 192.168.2.1
Pinging 192.168.2.1 with 32 bytes of data:
Reply from 192.168.2.1: bytes=32 time=20ms TTL=126
Reply from 192.168.2.1: bytes=32 time=20ms TTL=126
Reply from 192.168.2.1: bytes=32 time=21ms TTL=126
Reply from 192.168.2.1: bytes=32 time=21ms TTL=126
```

```
PC> tracert 192.168.2.1

      Tracing route to 192.168.2.1 over a maximum of 30 hops:

      1  44 ms  9 ms  10 ms  192.168.1.100

      2  13 ms  13 ms  12 ms  10.0.0.2

      3  17 ms  22 ms  20 ms  192.168.2.1
```

Static Routing (using 3 router)

```
R1(config)# ip route 192.168.2.0 255.255.255.0 10.0.0.2
R1(config)# ip route 192.168.3.0 255.255.255.0 10.0.0.2
R1(config)# ip route 11.0.0.0 255.0.0.0 10.0.0.2
```

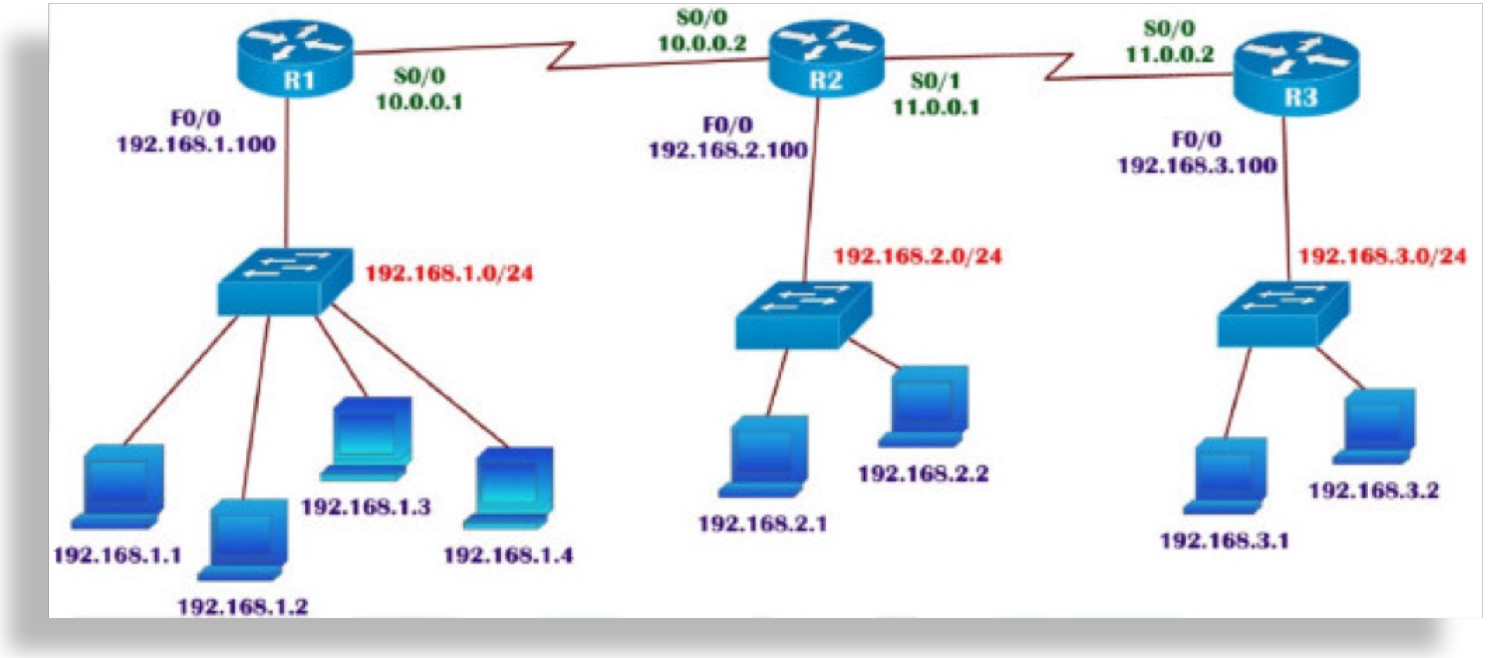

```
R2(config)# ip route 192.168.1.0 255.255.255.0 10.0.0.1
R2(config)# ip route 192.168.3.0 255.255.255.0 11.0.0.2

R3(config)# ip route 192.168.2.0 255.255.255.0 11.0.0.1
R3(config)# ip route 192.168.1.0 255.255.255.0 11.0.0.1
R3(config)# ip route 10.0.0.0 255.0.0.0 11.0.0.1
```

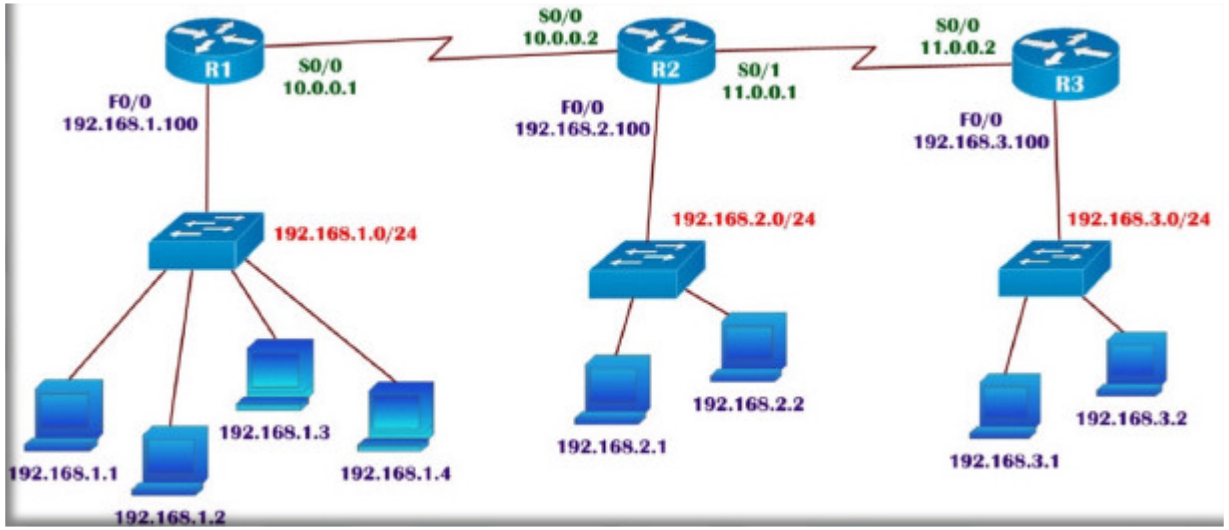

Routing Lookup

```
R1# show ip route
Gateway of last resort is not set
C 10.0.0.0/8 is directly connected, Serial0/0
S 11.0.0.0/8[1/0] via 10.0.0.2
C 192.168.1.0/24 is directly connected, FastEthernet0/0
S 192.168.2.0/24[1/0] via 10.0.0.2
S 192.168.3.0/24[1/0] via 10.0.0.2

R2# show ip route
Gateway of last resort is not set
C 10.0.0.0/8 is directly connected, Serial0/0
C 11.0.0.0/8 is directly connected, Serial0/1
S 192.168.1.0/24[1/0] via 10.0.0.1
C 192.168.2.0/24 is directly connected, FastEthernet0/0
S 192.168.3.0/24[1/0] via 11.0.0.2
```

Static Routing (Redundant Link)

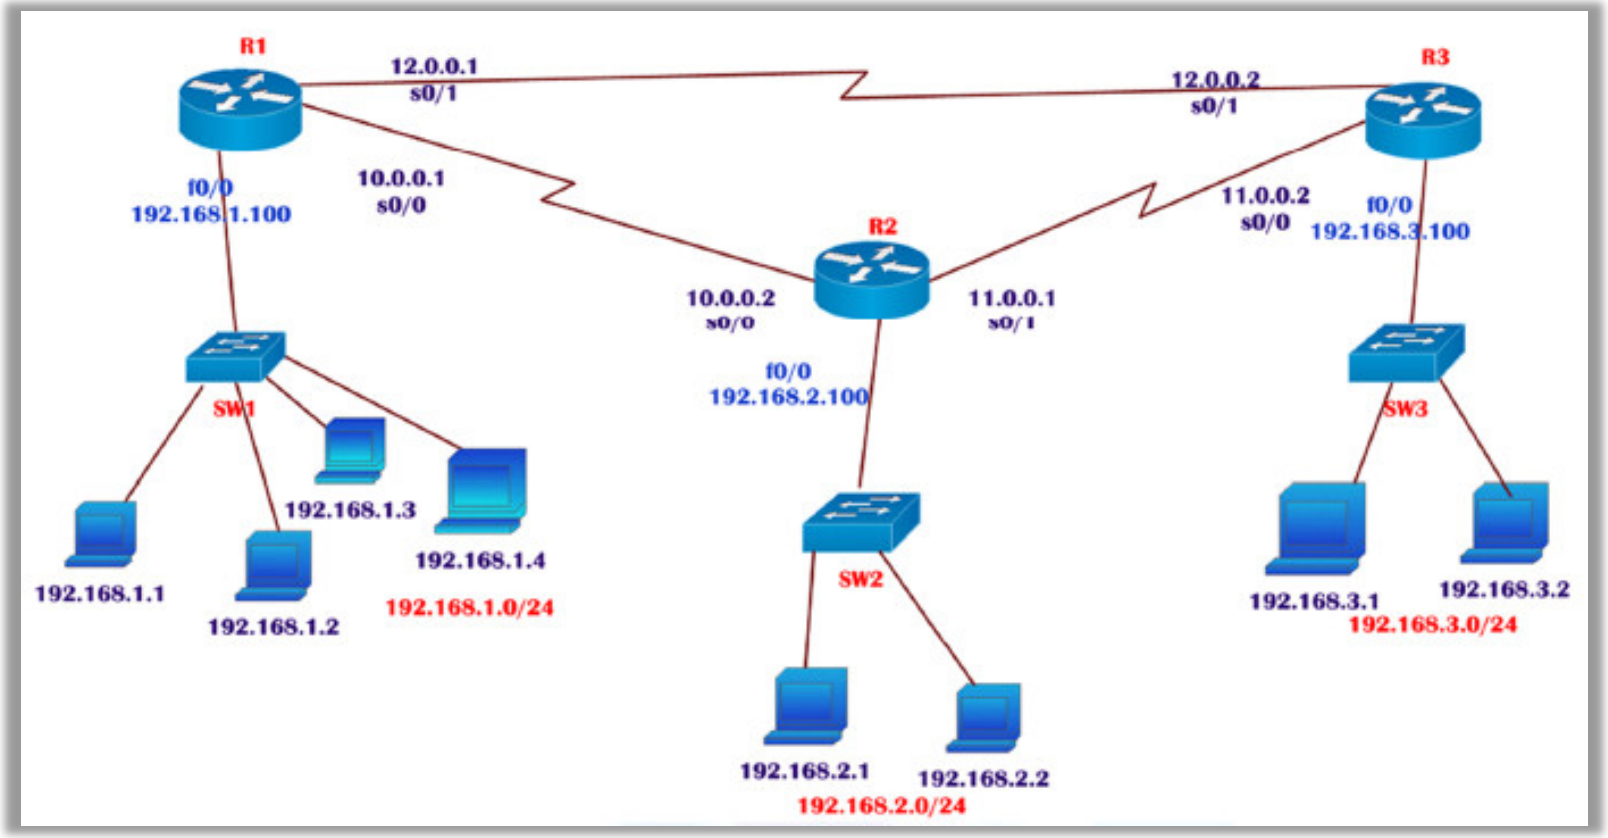

Default Routing

- Used to route traffic for unknown destinations (internet)

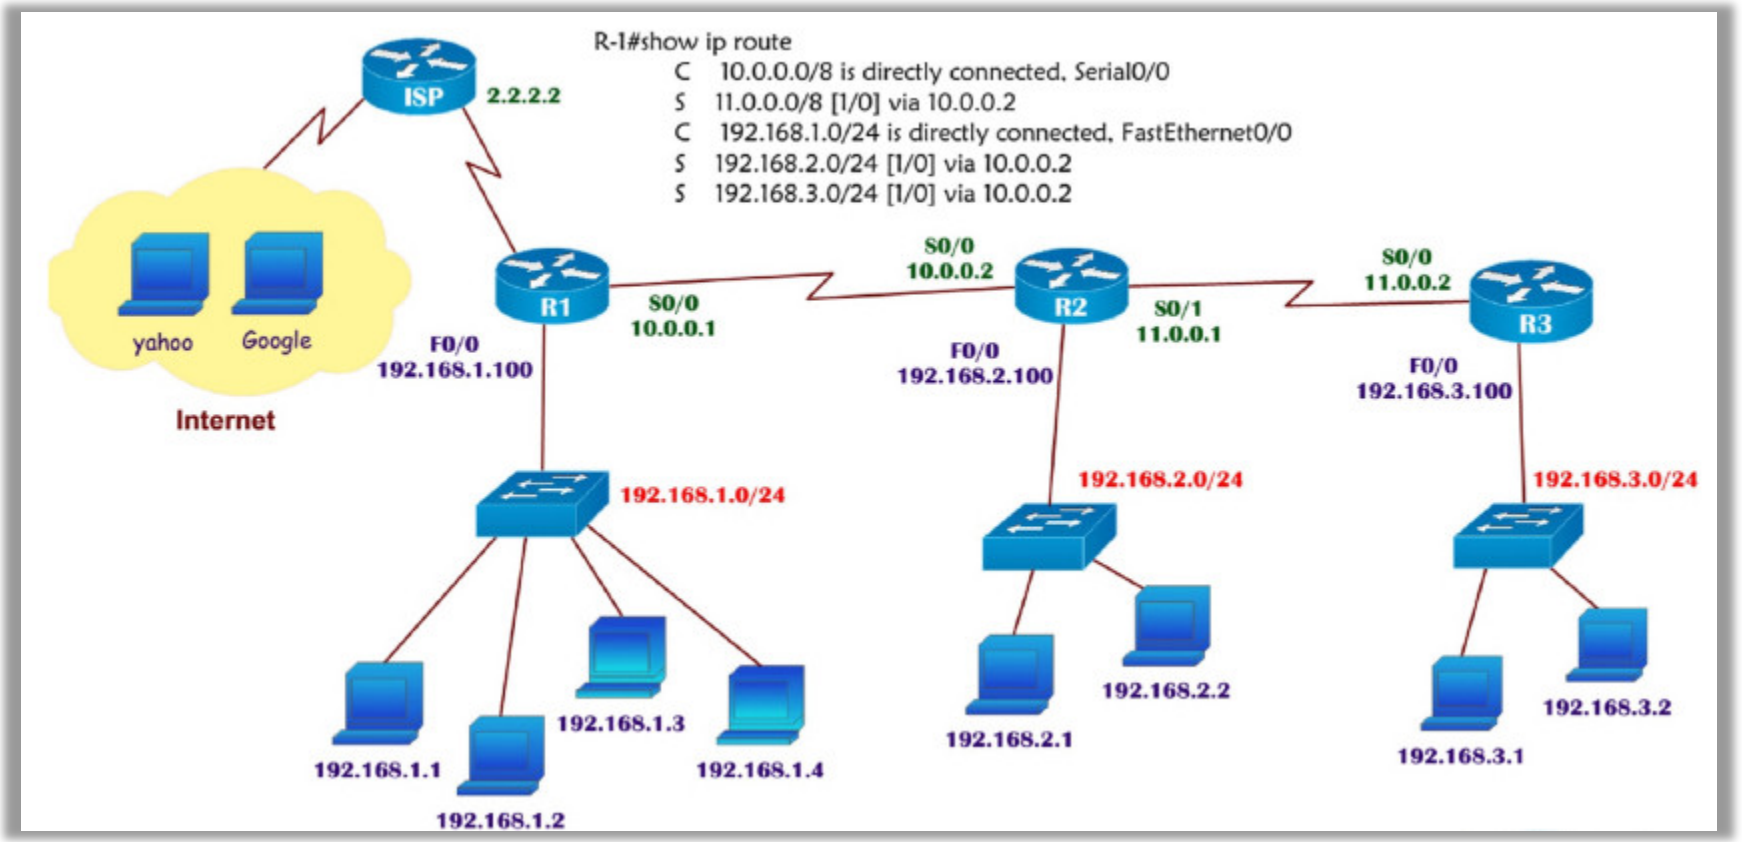

Default Routing

- Also can be used at end locations (optional)

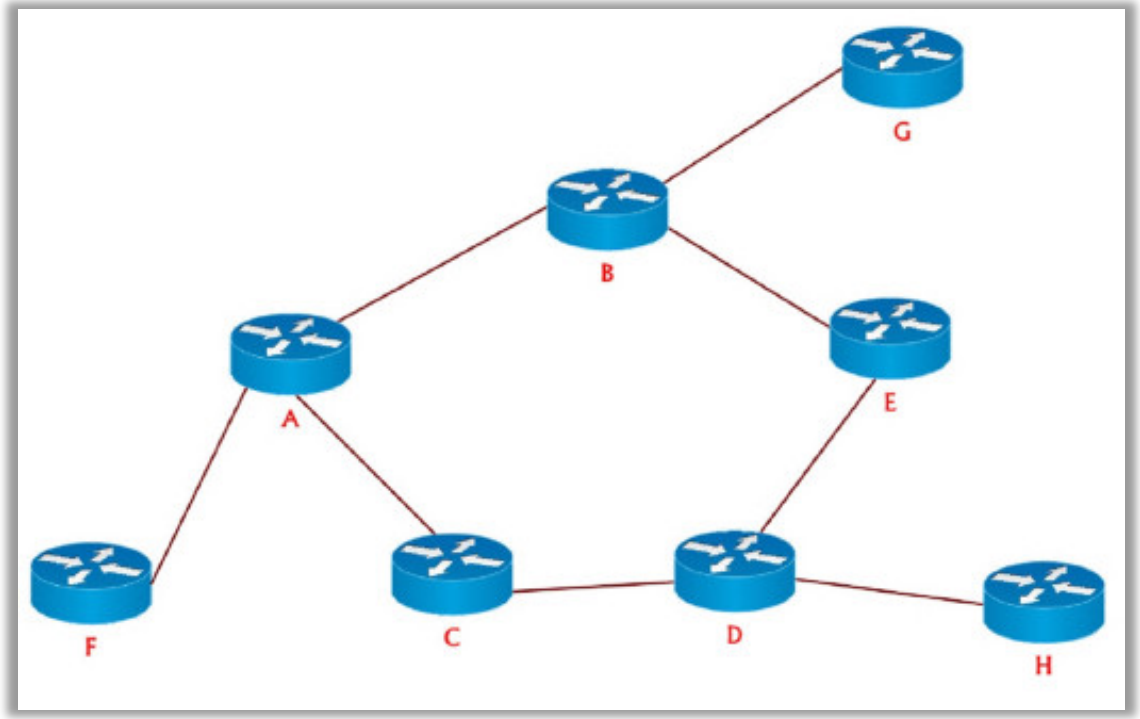

Dynamic Routing

- Advantages of Dynamic over Static:
  - No Need of Manual configuration (unlike static routing)
  - Learns about other networks via advertisements (of directly connected network)
  - Automatically select the best route (builds routing table)
  - Updates the topology changes dynamically
  - No need to know the destination network. (other network)
  - Administrative work is reduced
  - Applicable for large organization

Types of Dynamic Routing Protocol

- Distance Vector Protocol
- Link state Protocol
- Hybrid Protocol

Types of Dynamic Routing Protocol

Distance Vector	Link State	Hybrid (advance Distance vector)
Works with Bellman ford algorithm	Works with Dijkstra algorithm	Works with DUAL algorithm
Periodic Updates	Incremental Update	Incremental updates
Full routing tables are exchanged	Missing routes are exchanged	Missing route are exchanged
Classful routing Protocol	Classless routing protocol	Classless routing protocol
Updates are through broadcast	Updates are through multicast	Updates are through multicast
Example: RIPv1,RIPv2, IGRP	Example: OSPF, IS-IS	Example: EIGRP
Less overhead	More Overhead	Less overhead
Easy to configure	Difficult to configure	Easy to configure

- Classfull Protocols
  - Classfull routing protocol do not carry the subnet mask information along with updates
  - Which means that all devices in the network must use the same subnet mask (FLSM or default same class)
    - Ex: RIPv1, IGRP

- Classless Protocols
  - Classless routing protocol carry the subnet mask information along with updates
  - That’s why they support sub networks (VLSM and FLSM) and default networks also.
    - Ex. RIPv2, EIGRP, OSPF, IS-IS

Routing Information Protocol (RIP)

- Open standard protocol (cisco/no-cisco)
- Classfull routing protocol (no carry subnet mask)
- Updates are broadcasted via 255.255.255.255
- Metric: Hop count
- Load Balancing up to 4 equal paths
- Max Hop Counts =15 /Max router =16
- Applicable for small organizations
- Administrative distance is 120
- Exchange entire routing table for every 30 seconds(periodic update)

Routing Loop

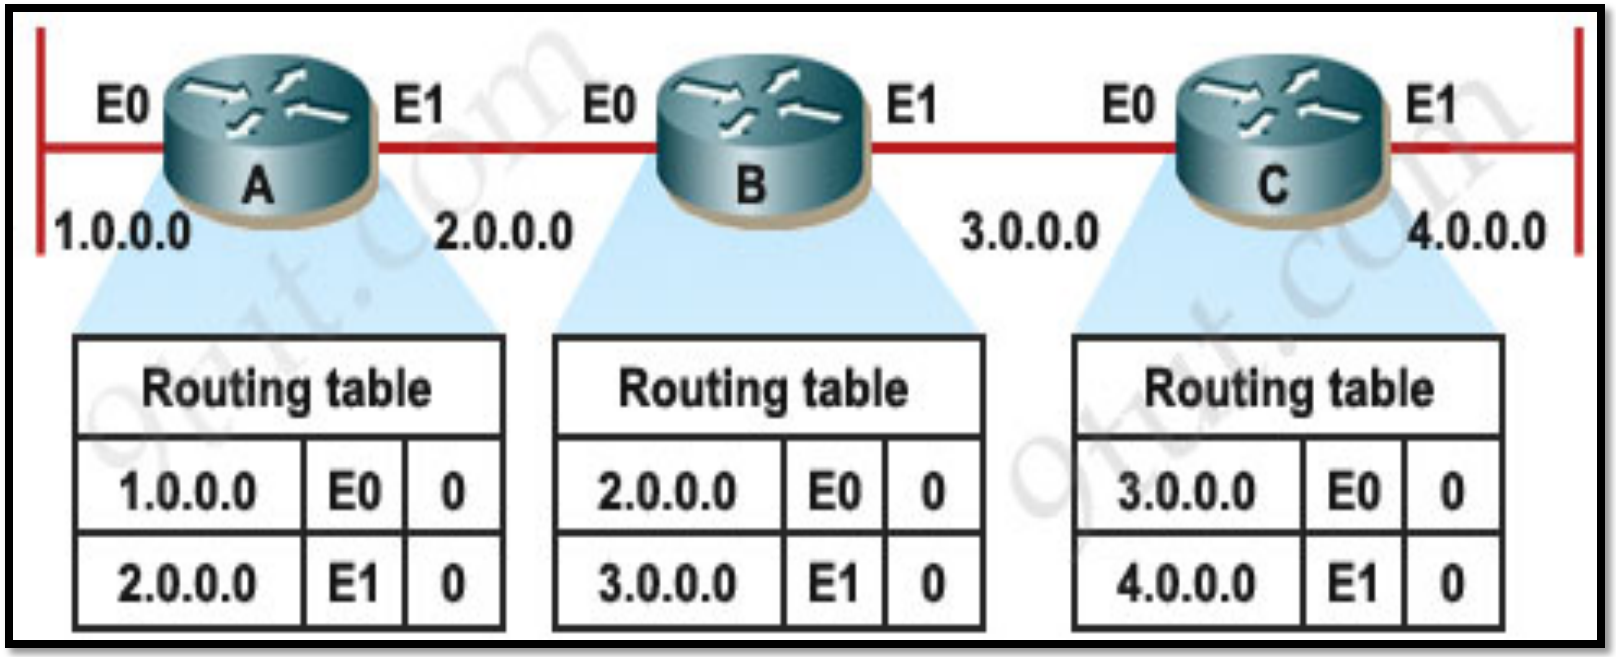

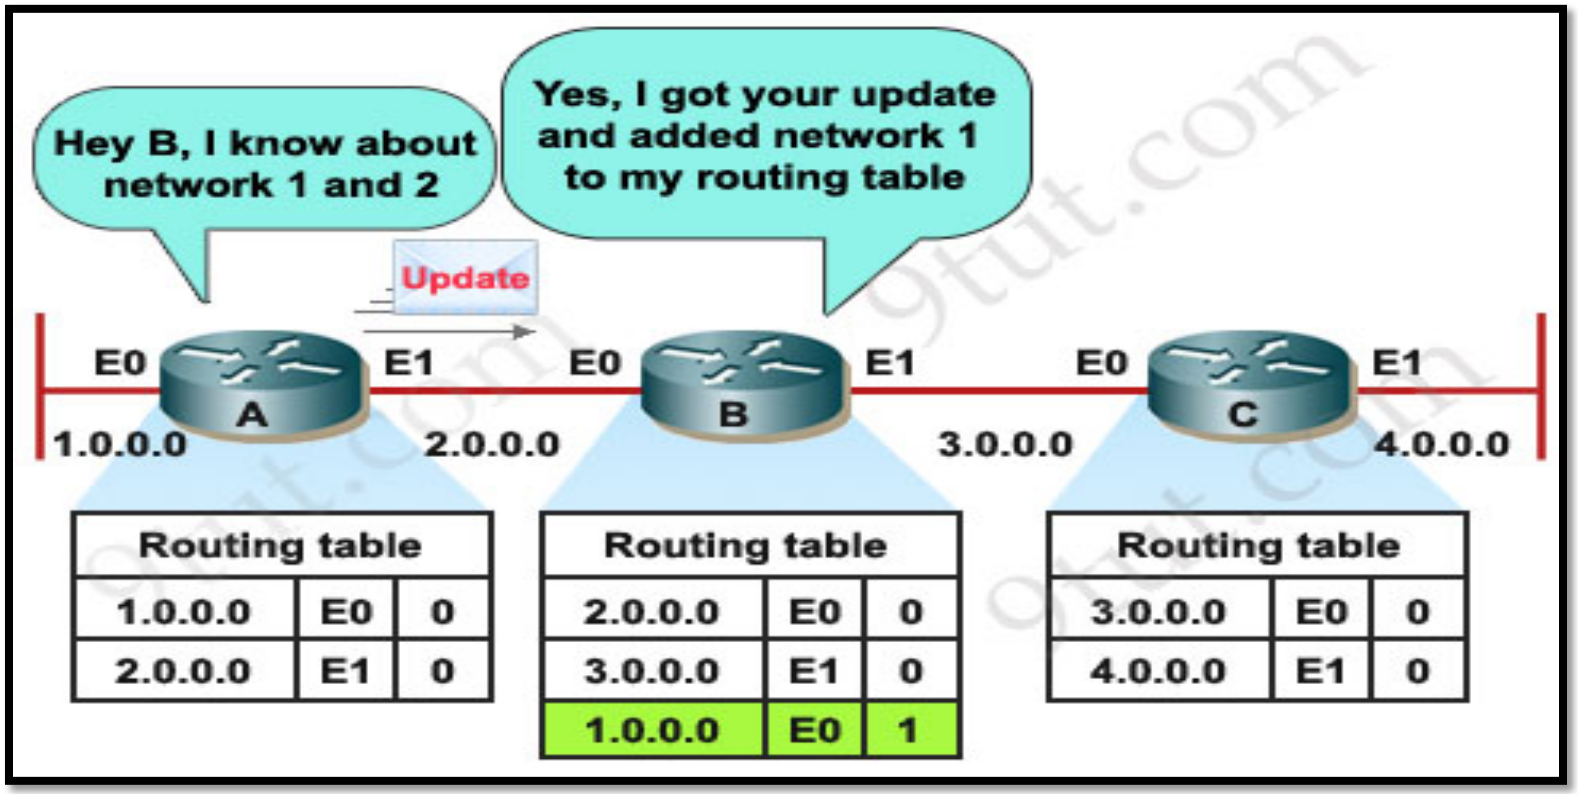
Hey B, I know about network 1 and 2
Yes, I got your update and added network 1 to my routing table
Update
E0
E1
E0
E1
E0
E1
A
B
C
1.0.0.0
2.0.0.0
3.0.0.0
4.0.0.0
Routing table
1.0.0.0 E0 0
2.0.0.0 E1 0
Routing table
2.0.0.0 E0 0
3.0.0.0 E1 0
1.0.0.0 E0 1
Routing table
3.0.0.0 E0 0
4.0.0.0 E1 0

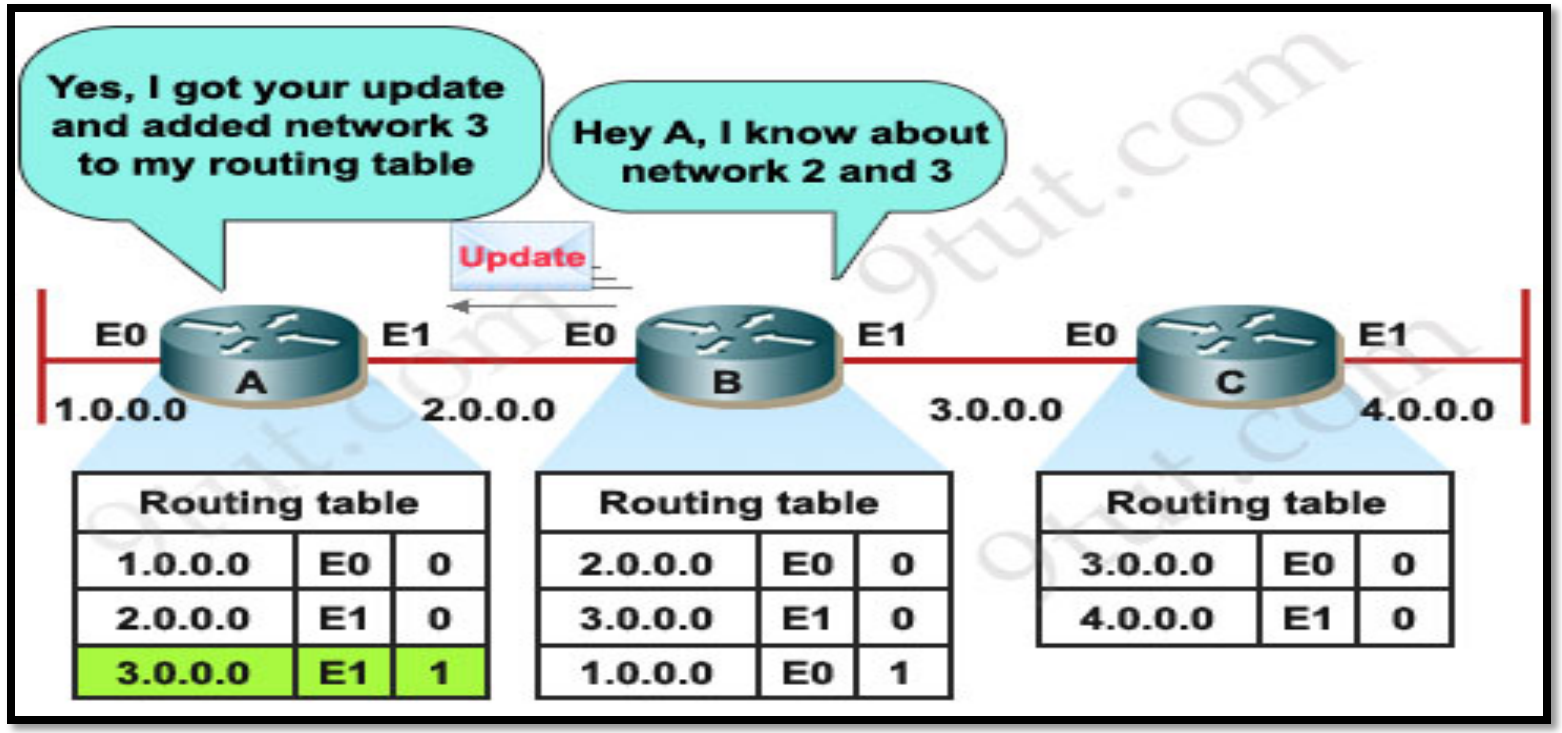
Yes, I got your update and added network 3 to my routing table
Hey A, I know about network 2 and 3
Update
E0
E1
E0
E1
E0
E1
A
B
C
1.0.0.0
2.0.0.0
3.0.0.0
4.0.0.0
Routing table
1.0.0.0 E0 0
2.0.0.0 E1 0
3.0.0.0 E1 1
Routing table
2.0.0.0 E0 0
3.0.0.0 E1 0
1.0.0.0 E0 1
Routing table
3.0.0.0 E0 0
4.0.0.0 E1 0

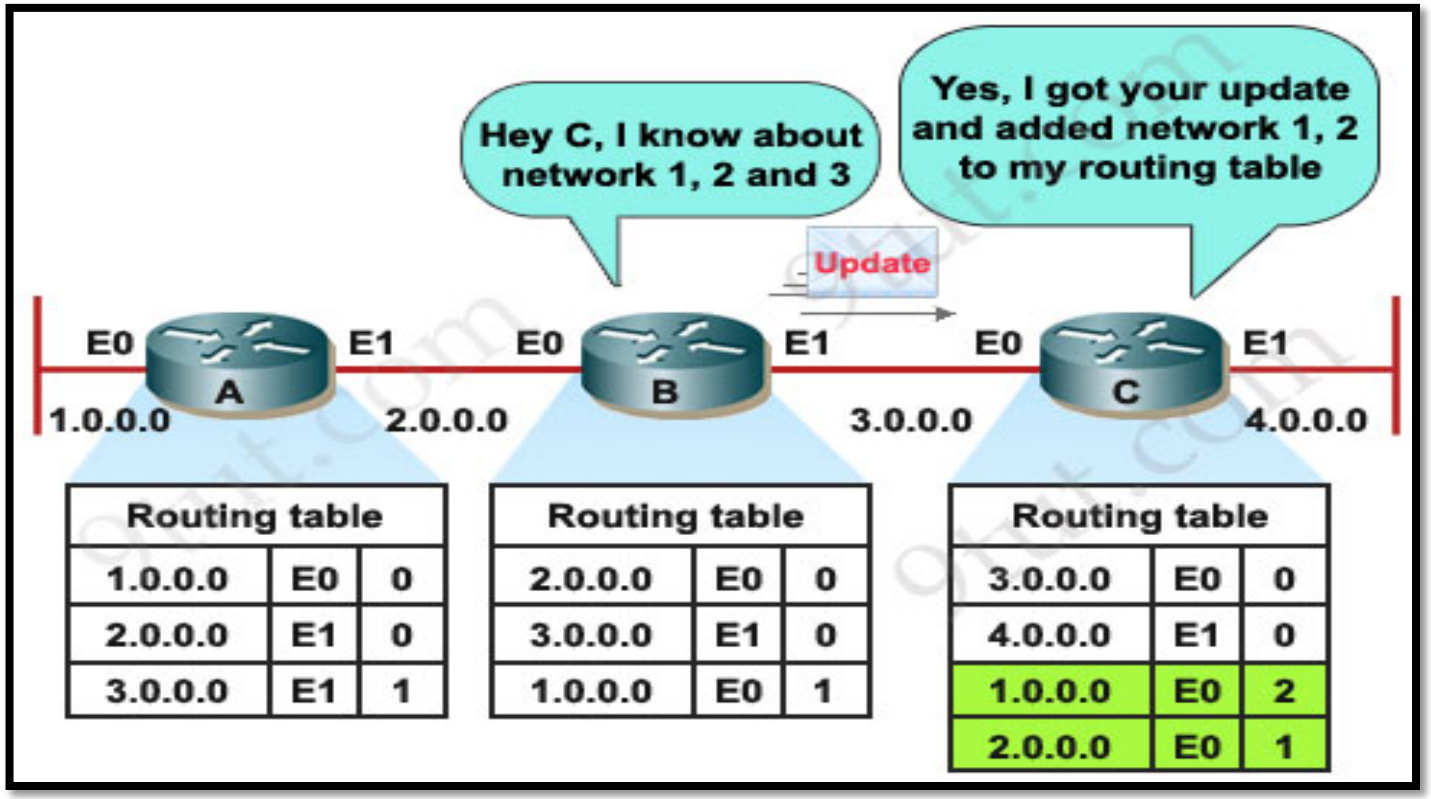
Hey C, I know about network 1, 2 and 3
Yes, I got your update and added network 1, 2 to my routing table
Update
E0
E1
A
B
C
1.0.0.0
2.0.0.0
3.0.0.0
4.0.0.0
Routing table
1.0.0.0 E0 0
2.0.0.0 E1 0
3.0.0.0 E1 1
Routing table
2.0.0.0 E0 0
3.0.0.0 E1 0
1.0.0.0 E0 1
Routing table
3.0.0.0 E0 0
4.0.0.0 E1 0
1.0.0.0 E0 2
2.0.0.0 E0 1

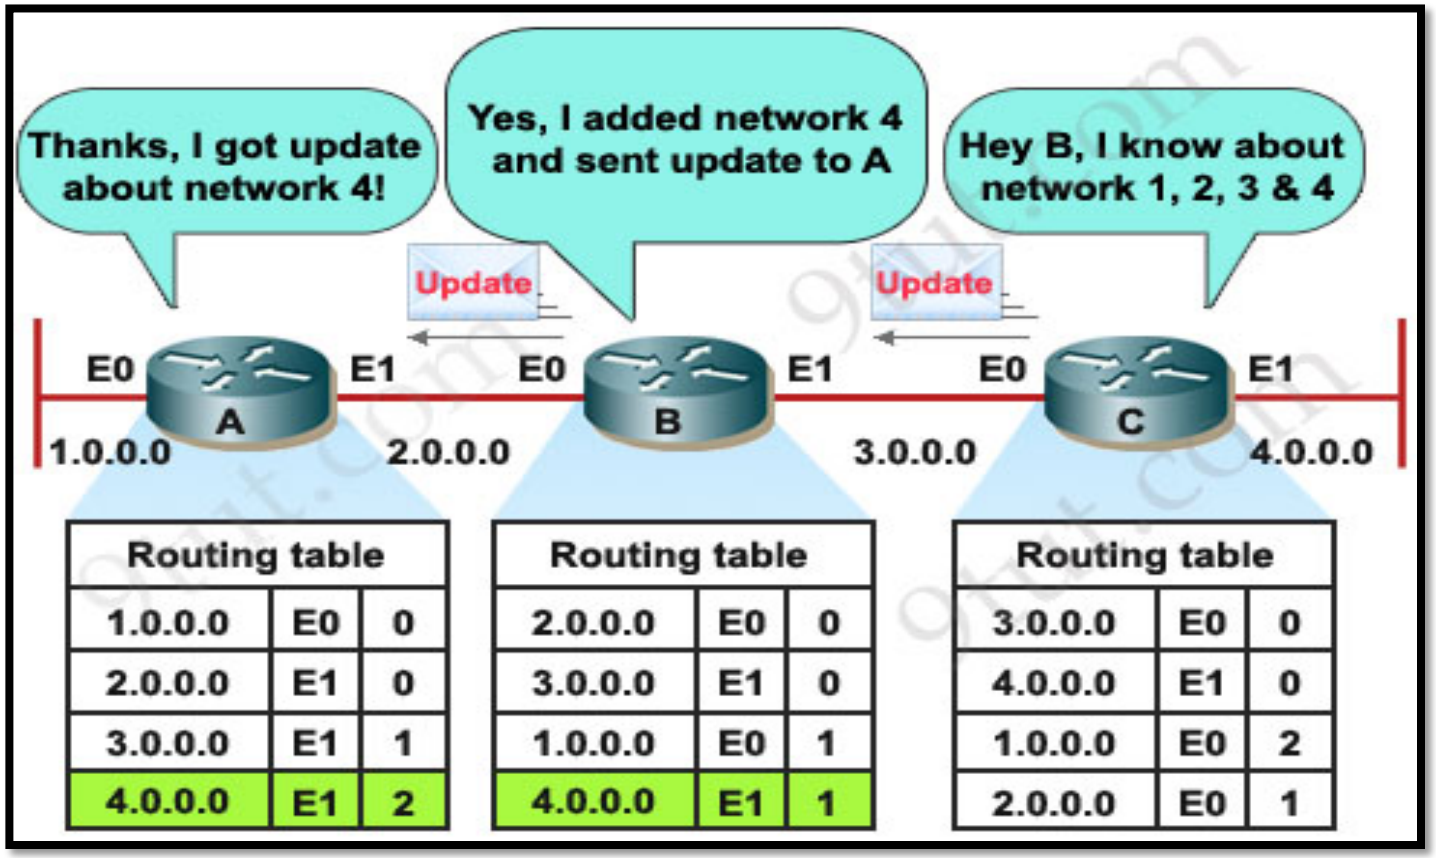
Thanks, I got update about network 4!
Yes, I added network 4 and sent update to A
Hey B, I know about network 1, 2, 3 & 4
Update
Update
E0
E1
E0
E1
E0
E1
A
B
C
1.0.0.0
2.0.0.0
3.0.0.0
4.0.0.0
Routing table
1.0.0.0 E0 0
2.0.0.0 E1 0
3.0.0.0 E1 1
4.0.0.0 E1 2
Routing table
2.0.0.0 E0 0
3.0.0.0 E1 0
1.0.0.0 E0 1
4.0.0.0 E1 1
Routing table
3.0.0.0 E0 0
4.0.0.0 E1 0
1.0.0.0 E0 2
2.0.0.0 E0 1

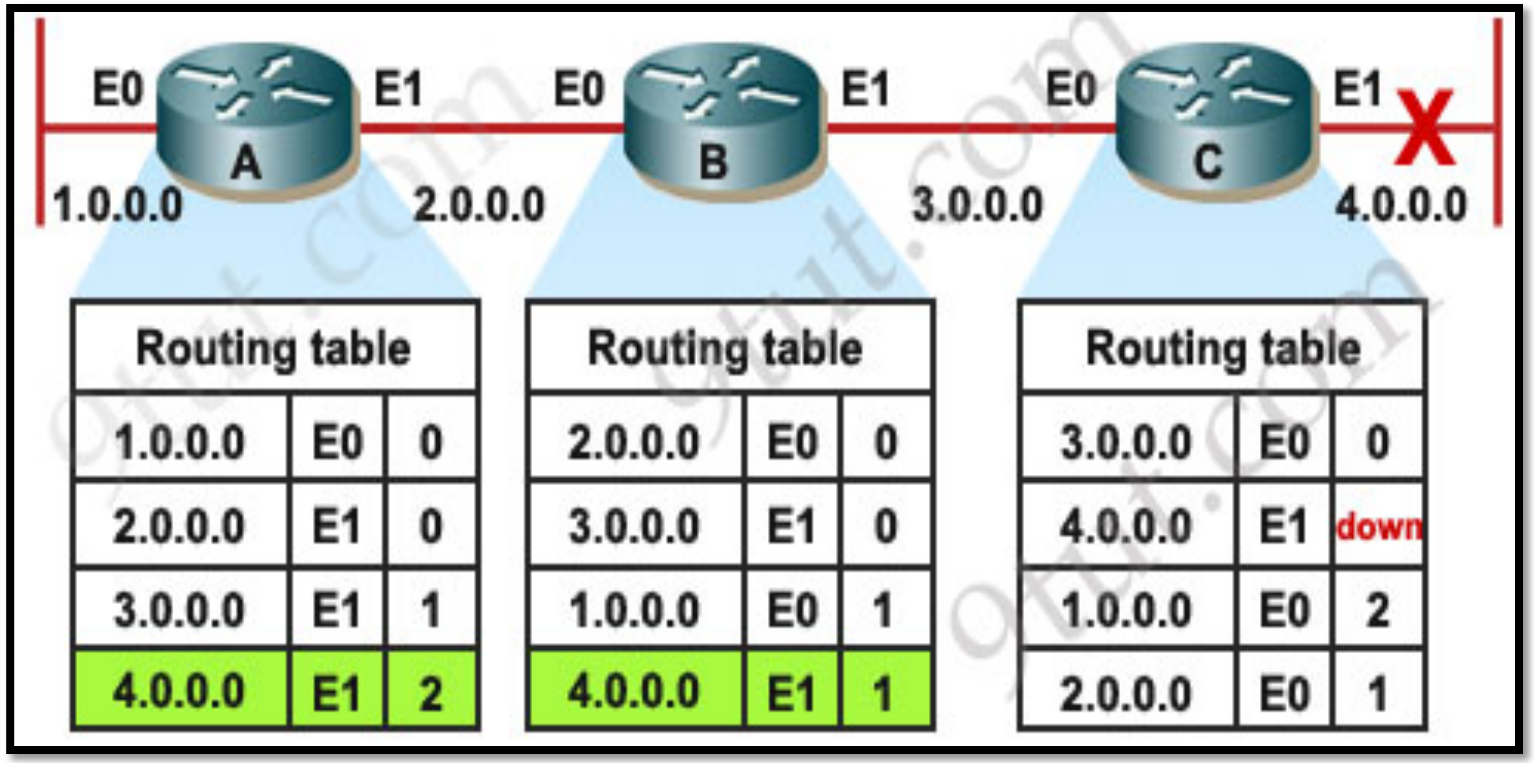
E0
E1
E0
E1
E0
E1
A
B
C
X
1.0.0.0
2.0.0.0
3.0.0.0
4.0.0.0
Routing table
1.0.0.0 E0 0
2.0.0.0 E1 0
3.0.0.0 E1 1
4.0.0.0 E1 2
Routing table
2.0.0.0 E0 0
3.0.0.0 E1 0
1.0.0.0 E0 1
4.0.0.0 E1 1
Routing table
3.0.0.0 E0 0
4.0.0.0 E1 down
1.0.0.0 E0 2
2.0.0.0 E0 1

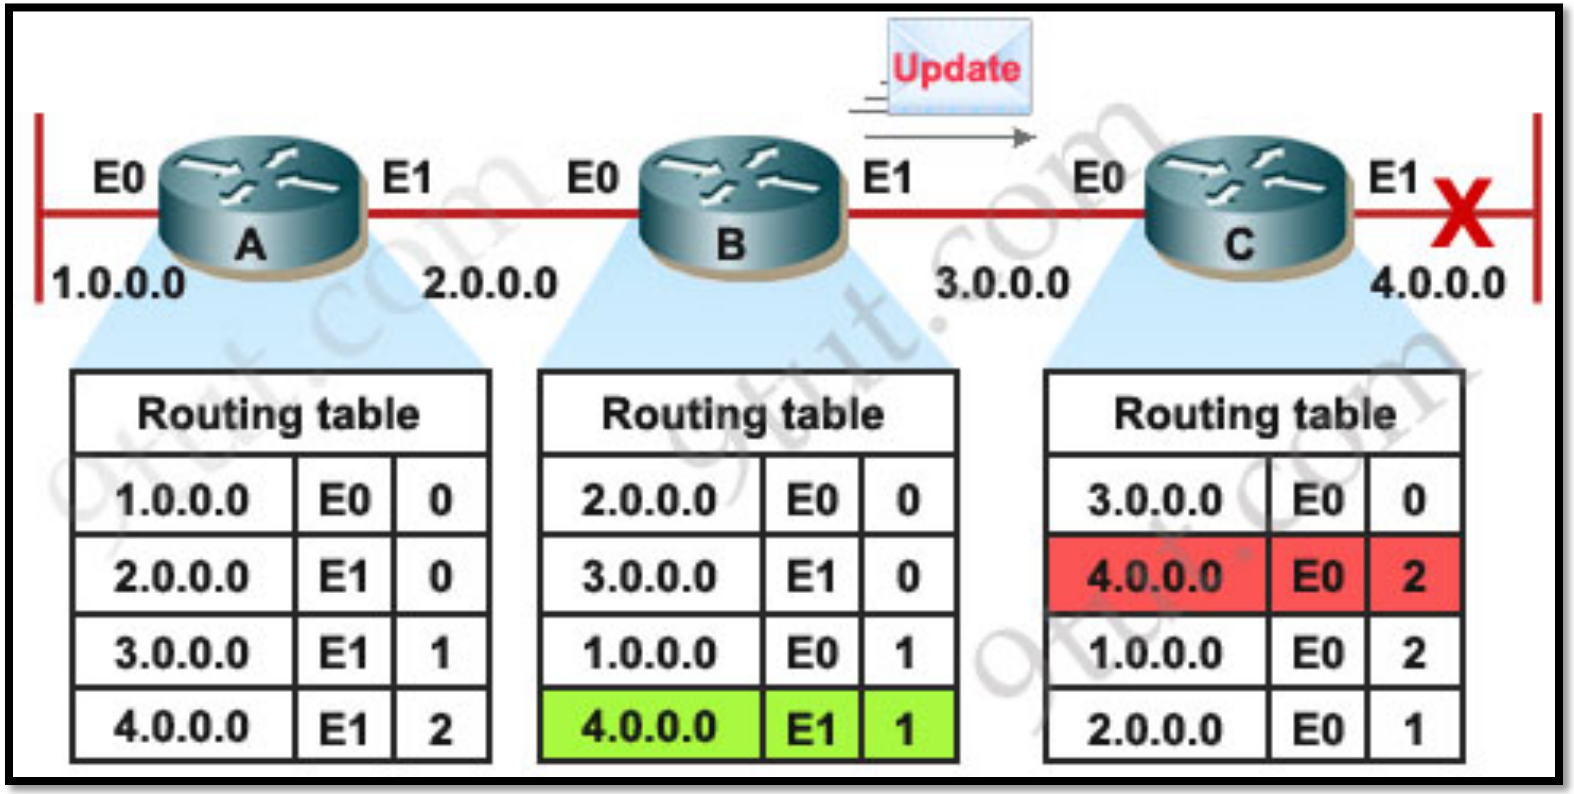
Update
E0
E1
E0
E1
E0
E1
A
B
C
1.0.0.0
2.0.0.0
3.0.0.0
4.0.0.0
Routing table
1.0.0.0 E0 0
2.0.0.0 E1 0
3.0.0.0 E1 1
4.0.0.0 E1 2
Routing table
2.0.0.0 E0 0
3.0.0.0 E1 0
1.0.0.0 E0 1
4.0.0.0 E1 1
Routing table
3.0.0.0 E0 0
4.0.0.0 E0 2
1.0.0.0 E0 2
2.0.0.0 E0 1

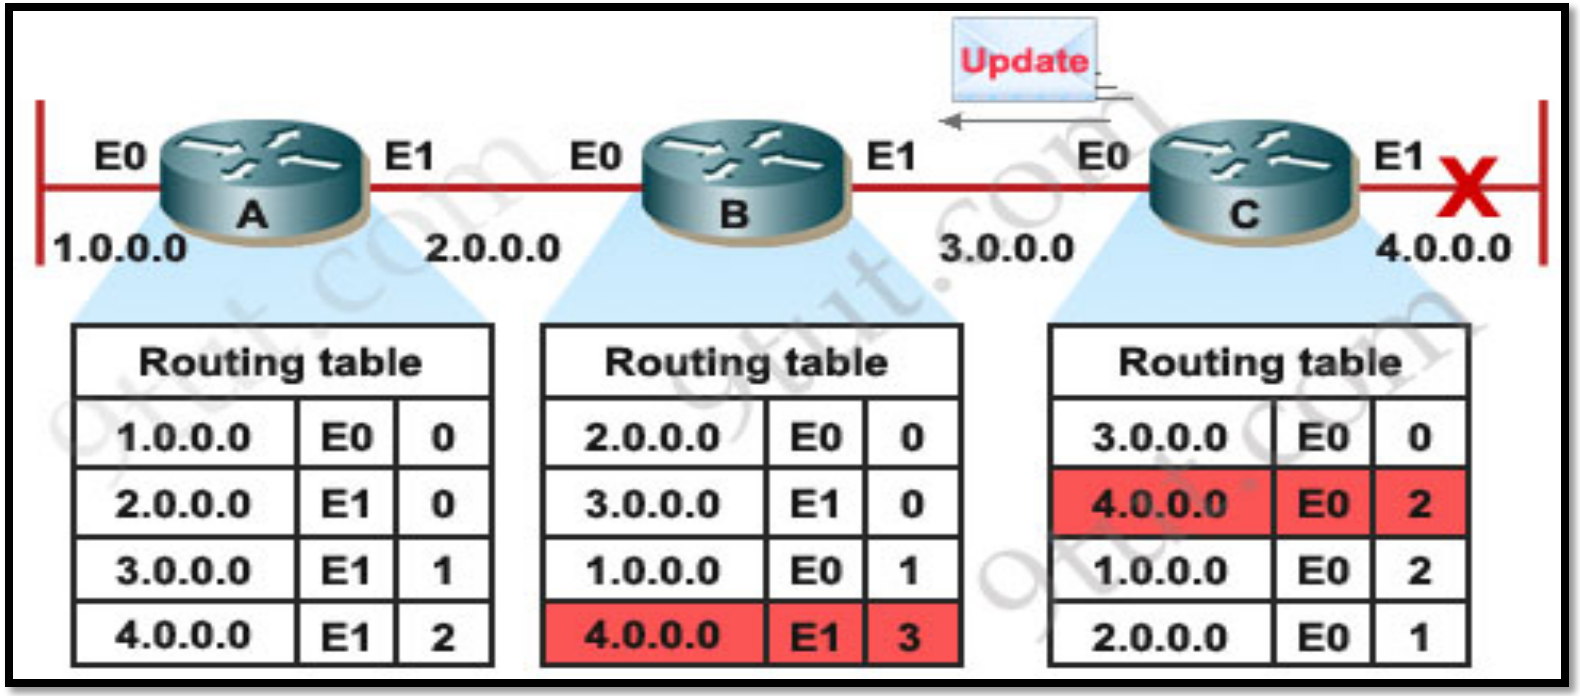
Update
E0
E1
E0
E1
E0
E1
A
B
C
1.0.0.0
2.0.0.0
3.0.0.0
4.0.0.0
Routing table
1.0.0.0 E0 0
2.0.0.0 E1 0
3.0.0.0 E1 1
4.0.0.0 E1 2
Routing table
2.0.0.0 E0 0
3.0.0.0 E1 0
1.0.0.0 E0 1
4.0.0.0 E1 3
Routing table
3.0.0.0 E0 0
4.0.0.0 E0 2
1.0.0.0 E0 2
2.0.0.0 E0 1

Routing Loop Prevention

- **Split Horizon**

- A router never sends information about a route back in same direction which is original information came.

- **Route Poisoning**

-Router consider route advertised with an infinitive metric to have failed (metric=16) instead of marking it down.

- **Poison Reverse**

-The poison reverse rule overwrites split horizon rule.

- **Hold down Timers**
- **Triggered Update**
- **Counting to Infinity**

RIP Timers

- Updates Time:

  Exchanged entire routing table for every 30 seconds (periodic update)

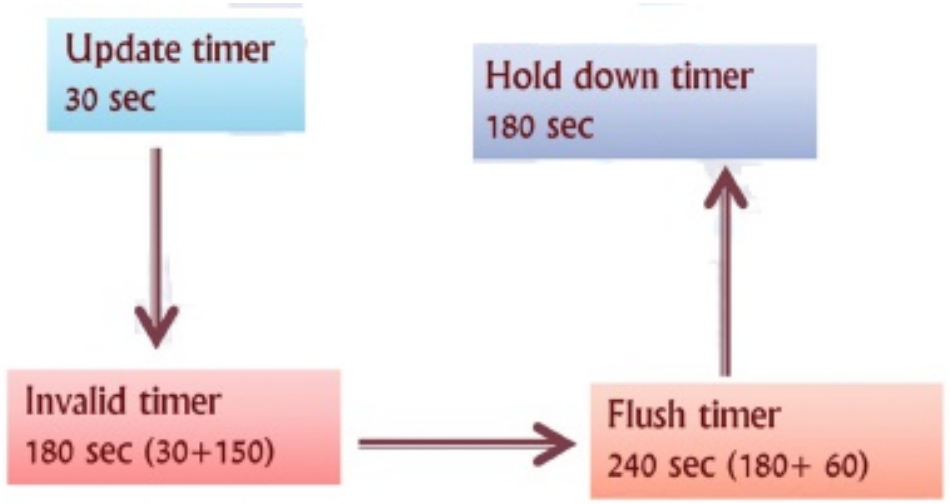

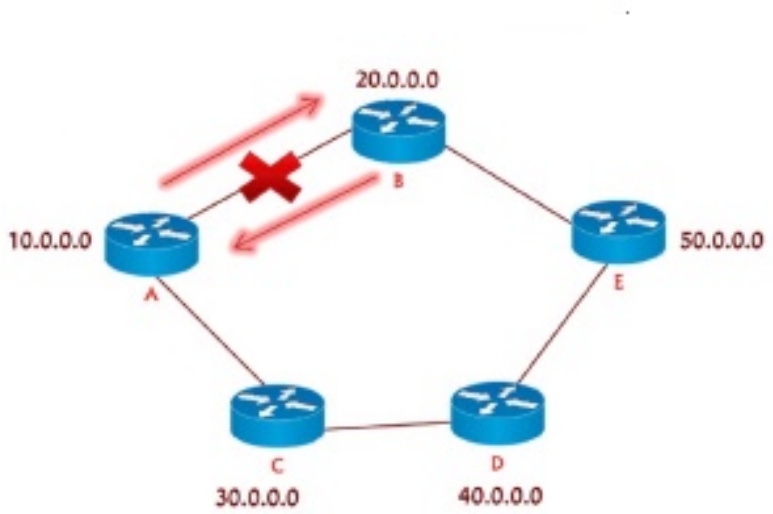

RIP Timers

- Update Time: 30 sec
  -Time Between consecutive updates
- Invalid Timer: 180 sec
  - Time a router waits to hear updates
  - The route is marked unreachable if there is no update during this interval
- Flush Timer: 240 sec
  - Time before the invalid route is removed from the routing table
- Hold Down timer: 180 sec
  - Stabilizes routing information and helps preventing routing loops during period when the topology is converging on new information.

RIPv1	RIPv2
Classfull routing protocol (not carry mask)	Classless routing protocol (carry mak)
Updates via broadcast (255.255.255.255)	Updates via multicast address 224.0.0.9
No support for authentication	Supports authentication

RIP Configuration

- Configure RIPv1:

```
Router(config)# router rip
Router(config-router)# network <Network ID>
```

- Configure RIPv2:

```
Router(config)# router rip
Router(config-router)# network <Network ID>
Router(config-router)# version 2
```

RIP Topology

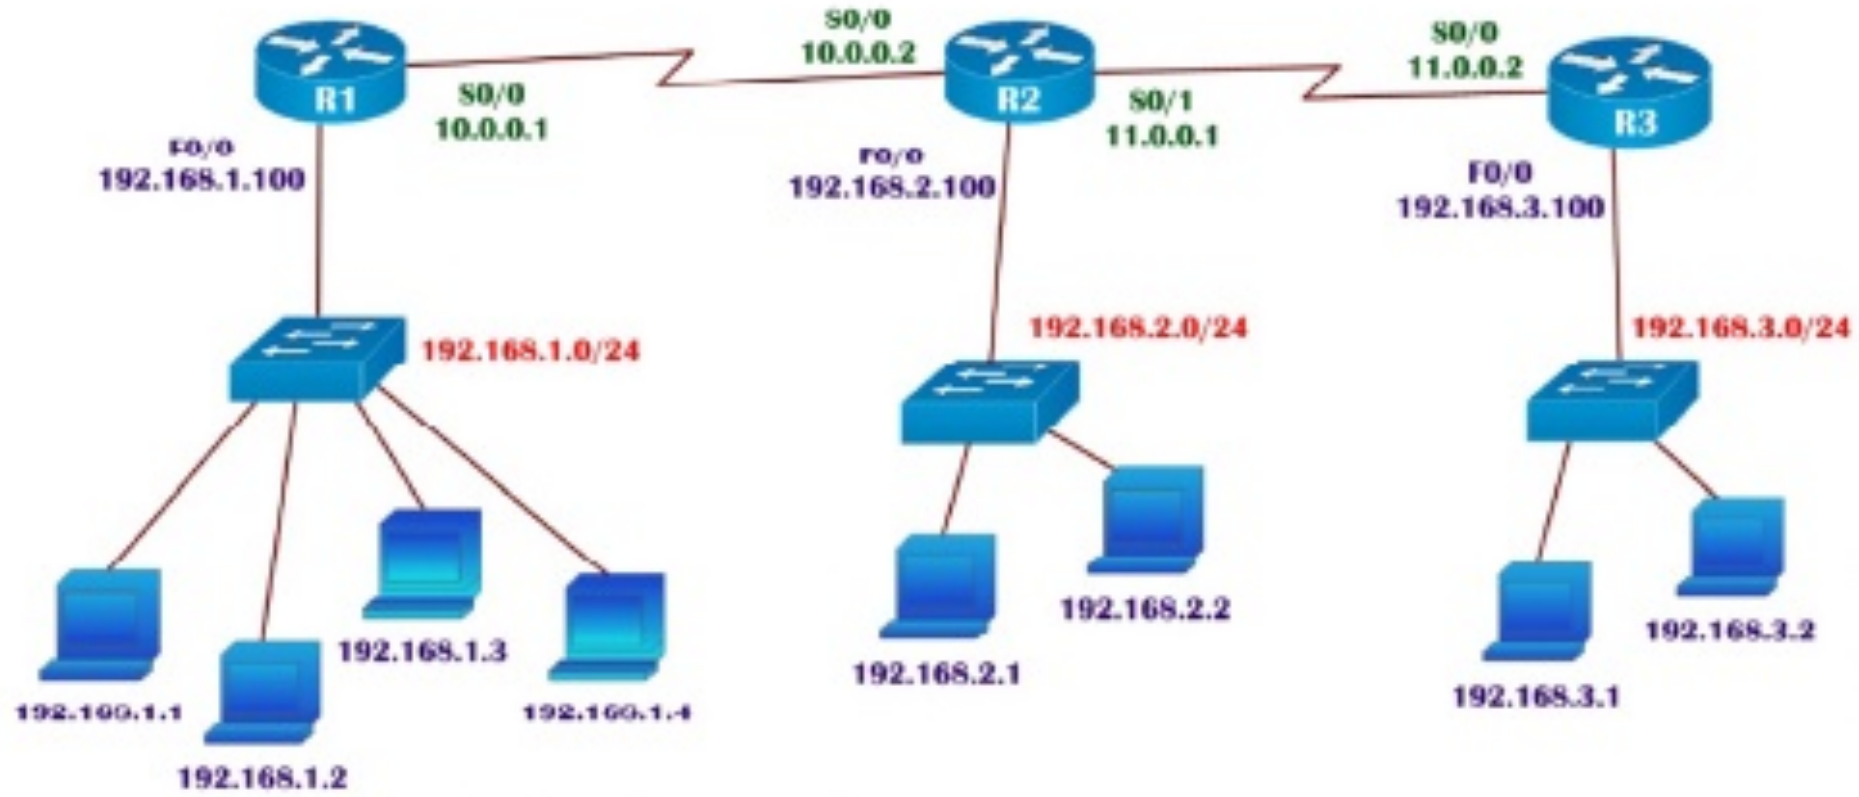

LAB: RIPv2

```
R1(config)#router rip
R1(config-router)# network 192.168.1.0
R1(config-router)# network 10.0.0.0
R1(config-router)#version 2

R2(config)#router rip
R2(config-router)# network 192.168.2.0
R2(config-router)# network 10.0.0.0
R2(config-router)# network 11.0.0.0
R2(config-router)#version 2

R1(config)#router rip
R1(config-router)# network 192.168.3.0
R1(config-router)# network 11.0.0.0
R1(config-router)#version 2
```

Verification

- R1# sh ip route
- PC> ping 192.168.3.1
- PC> traceroute 192.168.3.1

- Advantages of RIP
  - Easy to configure
  - No design constraints (unlike OSPF)
  - Less overhead

- Disadvantages of RIP
  - Bandwidth utilization is very high as broadcast for every 30 seconds
  - Works only on hop count (no consider BW)
  - Applicable for small organizations
  - Slow convergence (240sec)

Administrative Distance

- Trustworthiness of the information received by the router.
- The Number is between 0 and 255
- Less value is more preferred routing
  - Default Administrative Distance:
    - Directly Connected = 0
    - Static route = 1
    - IGRP = 100
    - EIGRP = 90
    - OSPF = 110
    - RIP = 120

Autonomous System Number

- Is a collection of networks under a common administrative domain.
- A unique number identifying the Routing domain of the routers (one organization)
- Ranges from 1- 65535
  - Public AS (in between SP) = 1 - 64512
  - Private AS (Same SP) = 64513 - 65535

Interior vs Exterior Gateway Routing protocols

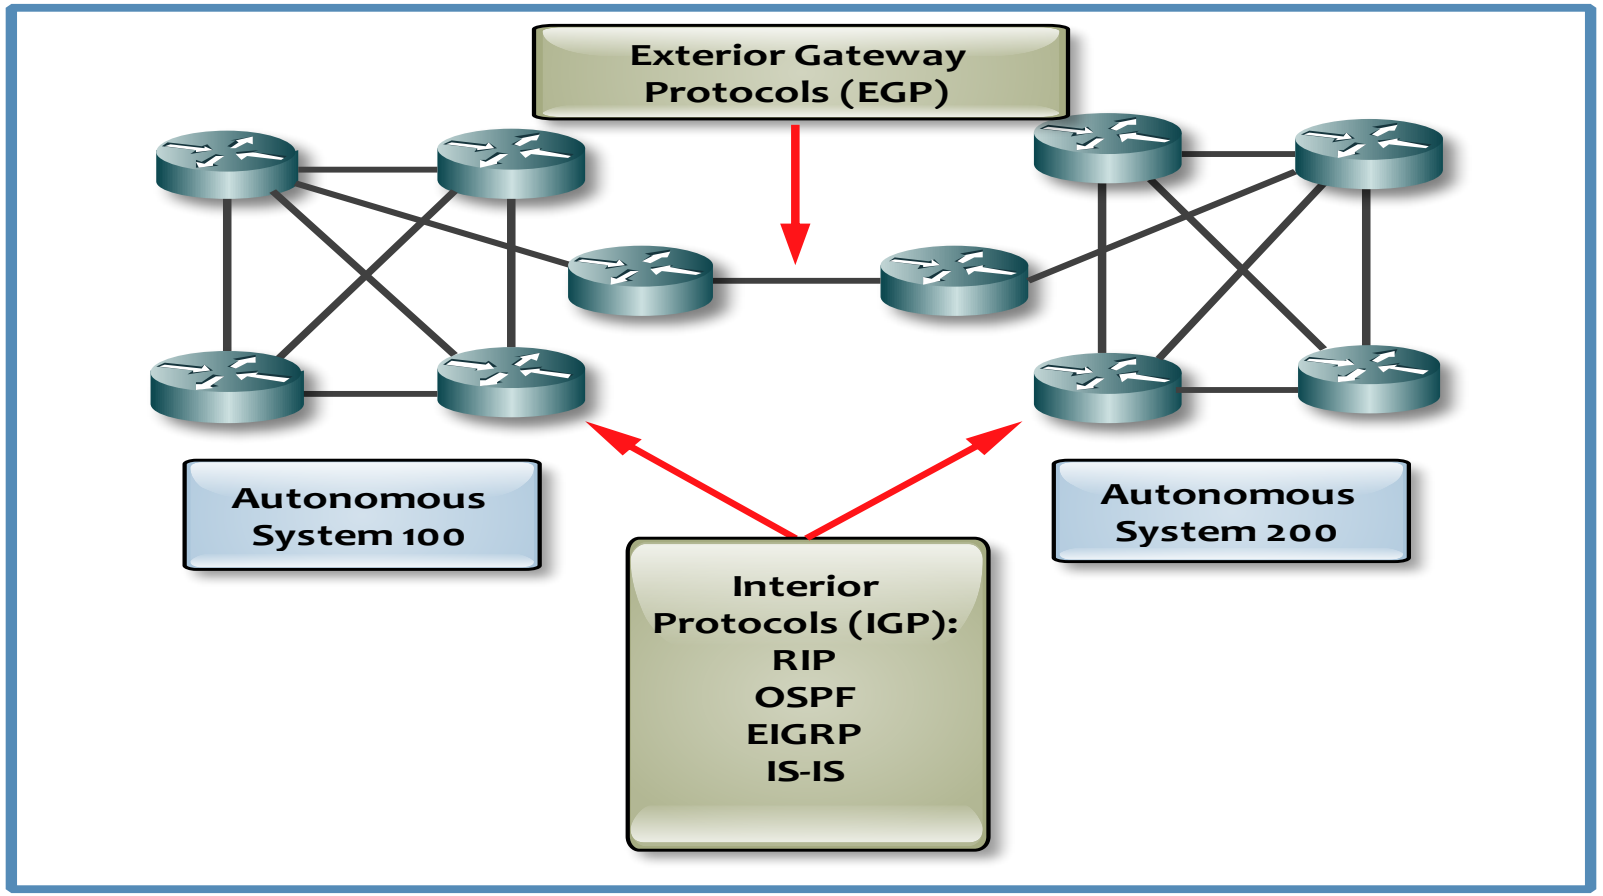

Routing Protocols Metrics

A metric is a value used by routing protocols to assign costs to reach remote networks.

Metrics Parameters:

- **Hop count** - A simple metric that counts the number of routers a packet must traverse
- **Bandwidth** - Influences path selection by preferring the path with the highest bandwidth
- **Load** - Considers the traffic utilization of a certain link
- **Delay** - Considers the time a packet takes to traverse a path
- **MTU –** Minimum path Maximum Transmission Unit (MTU) (never used in the metric calculation)
- **Reliability** - Assesses the probability of a link failure, calculated from the interface error count or previous link failures
- **Cost** - A value determined either by the IOS or by the network administrator to indicate preference for a route.
- Cost can represent a metric, a combination of metrics or a policy.

Convergence

- Convergence is when all routers' routing tables are at a state of consistency.
- The network has converged when all routers have complete and accurate information about the network.
- Convergence time is the time it takes routers to share information, calculate best paths, and update their routing tables.
- A network is not completely operable until the network has converged; therefore, most networks require short convergence times.
- RIP and IGRP are slow to converge
- EIGRP and OSPF are faster to converge.

Enhanced Interior Gateway Routing Protocol (EGRP)

- Cisco proprietary protocol
- First released in 1994 with IOS version 9.21
- Includes all features of IGRP
- Classless routing protocol(same as RIPv2 and OSPF)
- Support for VLSM/CIDR
- Metric (32 bit) : Composite Metric (BW + Delay + load + MTU + reliability
- Administrative distance is 90 for internal EIGRP 170 for External EIGRP.
- Updates are through Multicast (224.0.0.10)
- Max Hop count is 255 (100 by default)
- Hello packets are sent every 5 seconds
- Convergence rate is fast
- Support authentication (MD5)
- Support for IP, IPX, and AppleTalk via protocol-dependent modules
- Support for summaries and discontinuous networks
- Efficient neighbor discovery
- Communication via Reliable Transport Protocol (RTP)

EIGRP Packets

- **Hello:** used to identify neighbors. They are sent as periodic multicasts
- **Update:** used to advertise routes, only sent as multicasts when something is changed
- **Ack:** acknowledges receipt of an update. In fact, Ack is Hello packet without data. It is always unicast and uses UDP.
- **Query:** used to find alternate paths when all paths to a destination have failed
- **Reply:** is sent in response to query packets to instruct the originator not to recompute the route because feasible successors exist. Reply packets are always unicast to the originator of the query

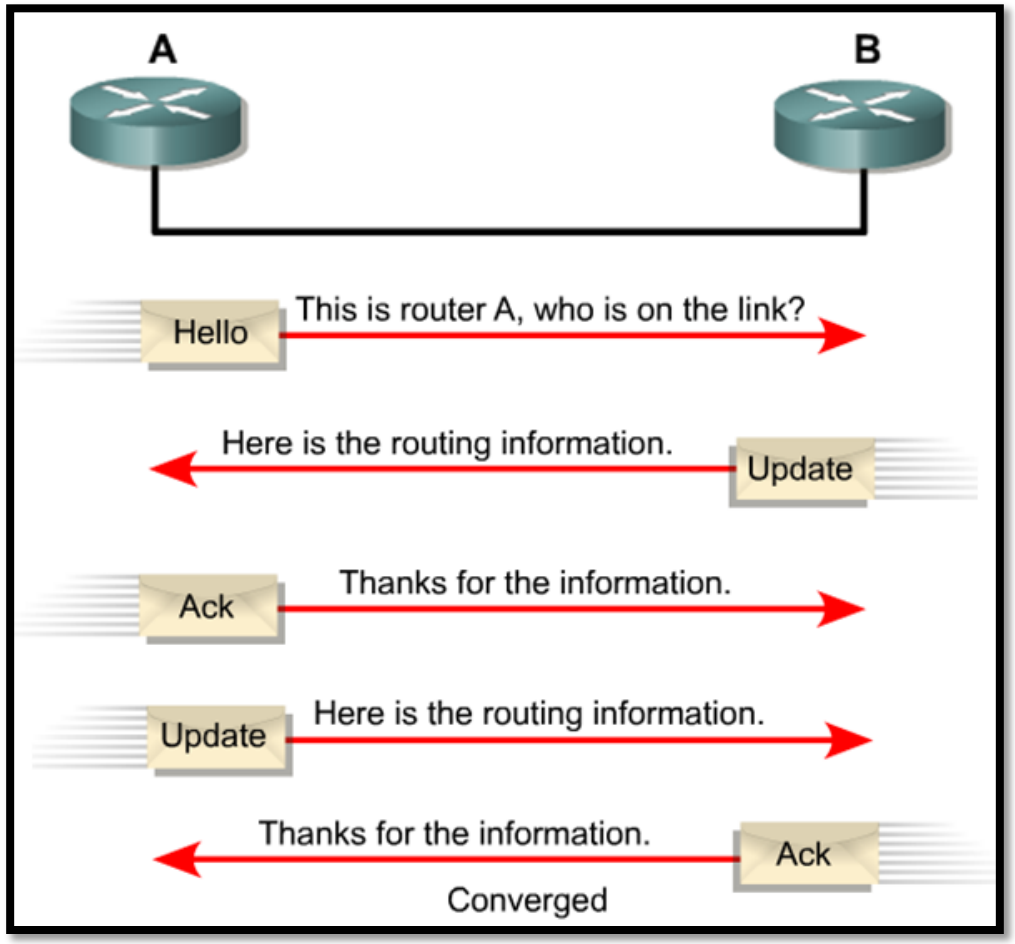
A
B
Hello
This is router A, who is on the link?
Here is the routing information.
Update
Ack
Thanks for the information.
Update
Here is the routing information.
Thanks for the information.
Ack
Converged

EIGRP Table

- **Neighbor Table** - lists directly connected routers running EIGRP with which this router has an adjacency
- **Topology Table** - Lists all the paths to all destinations learned. However, it shows only metric and the next-hop router and does not illustrate the real topology.
- **Routing Table** - The best metric routes are going to be used to send the traffic.

Router C's tables:

IP EIGRP Neighbor Table

Next-Hop Router	Interface
Router A Router B	Ethernet 0 Ethernet 1

IP EIGRP Topology Table

Network	Feasible Distance (EIGRP Metric)	Advertised Distance	EIGRP Neighbor
10.1.1.0 /24 10.1.1.0 /24	2000 2500	1000 1500	Router A (E0) Router B (E1)

The IP Routing Table

Network	Metric (Feasible Distance)	Outbound Interface	Next Hop (EIGRP Neighbor)
10.1.1.0 /24	2000	Ethernet 0	Router A

The EIGRP Topology Table

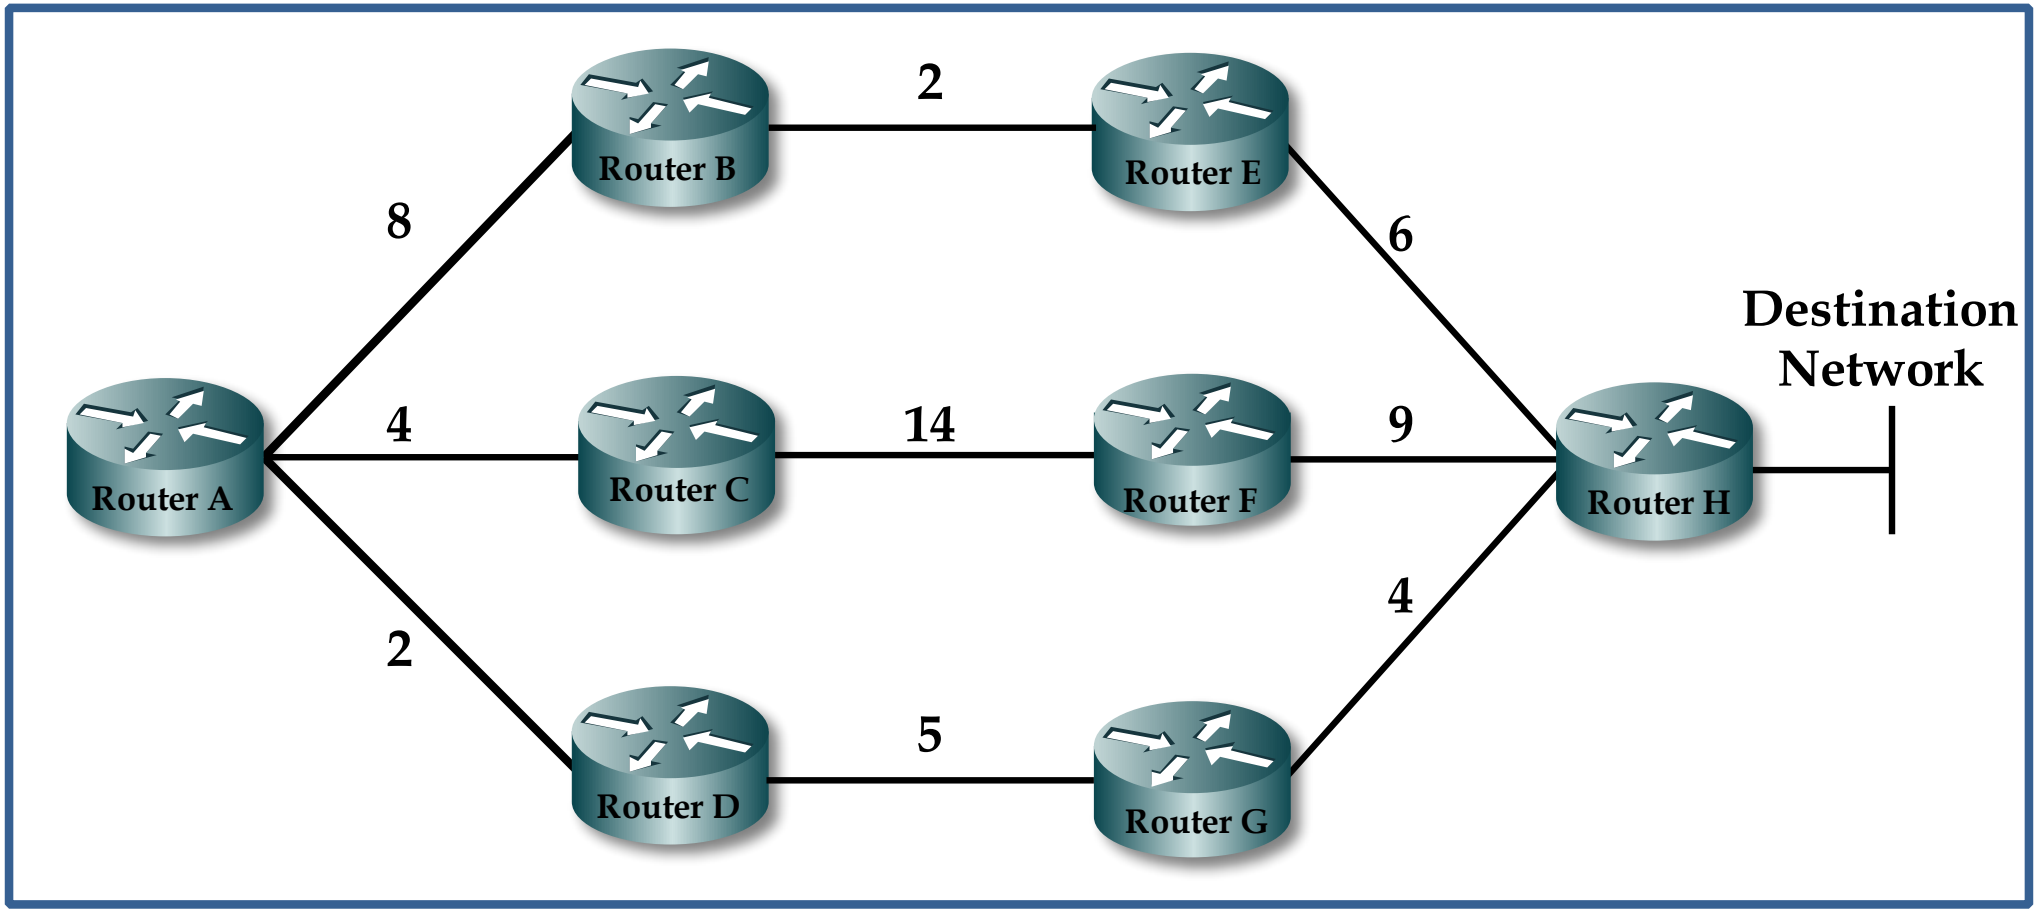

- **Feasible distance**
- **Reported distance**
- **Neighbor table**
- **Topology table**
- **Successor**
- **Feasible successor**

Topology table Example

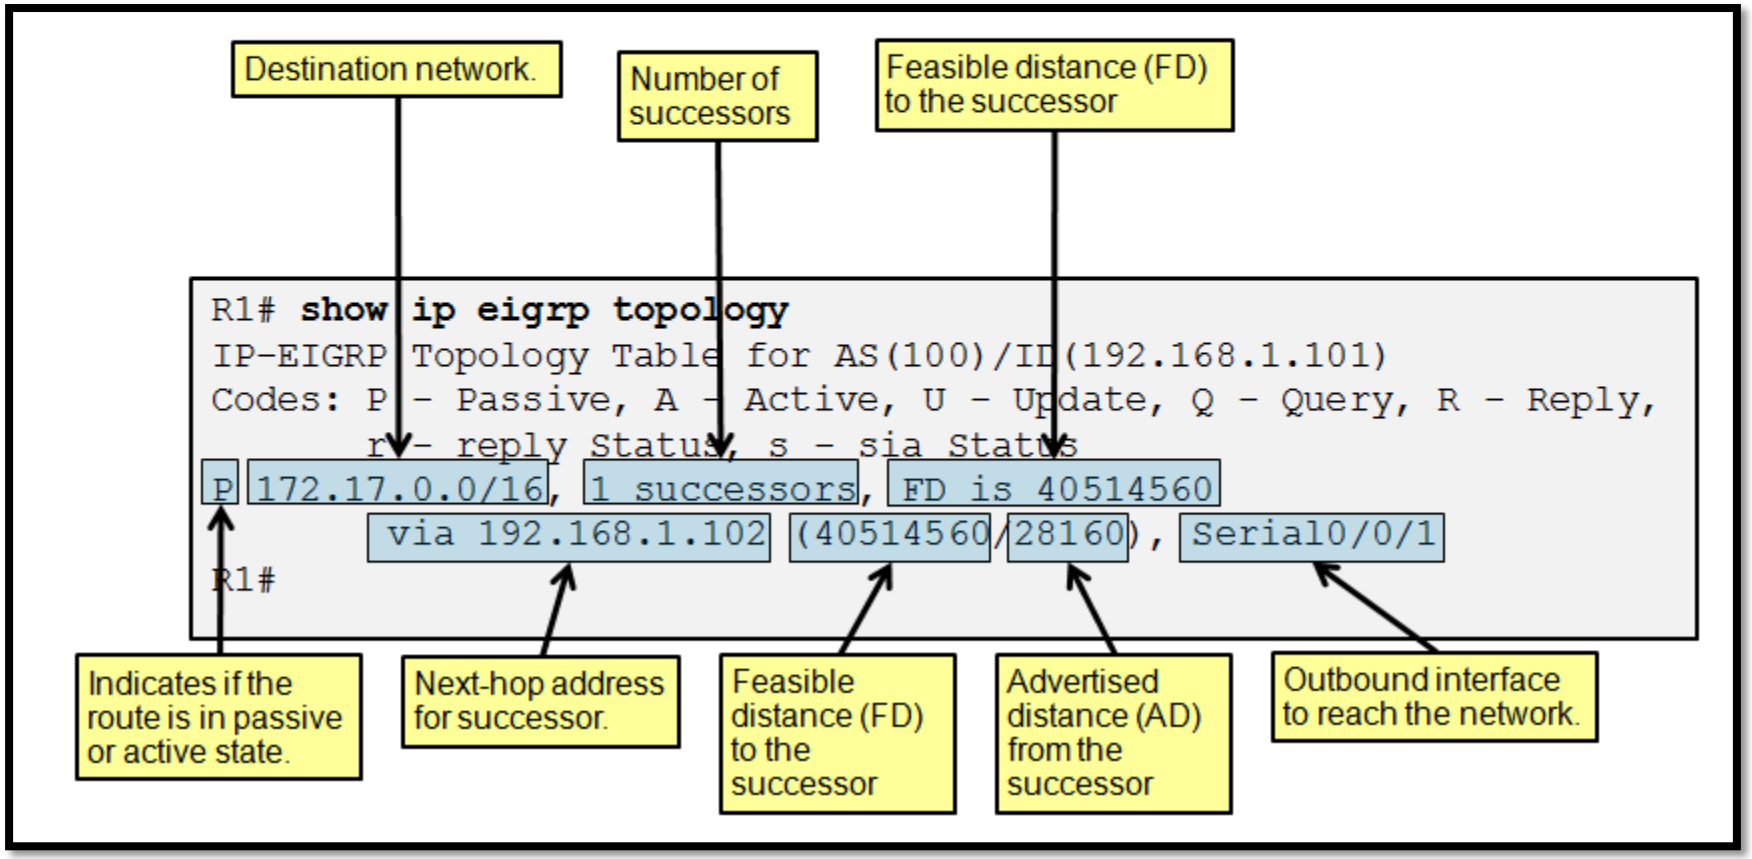

Neighbor table

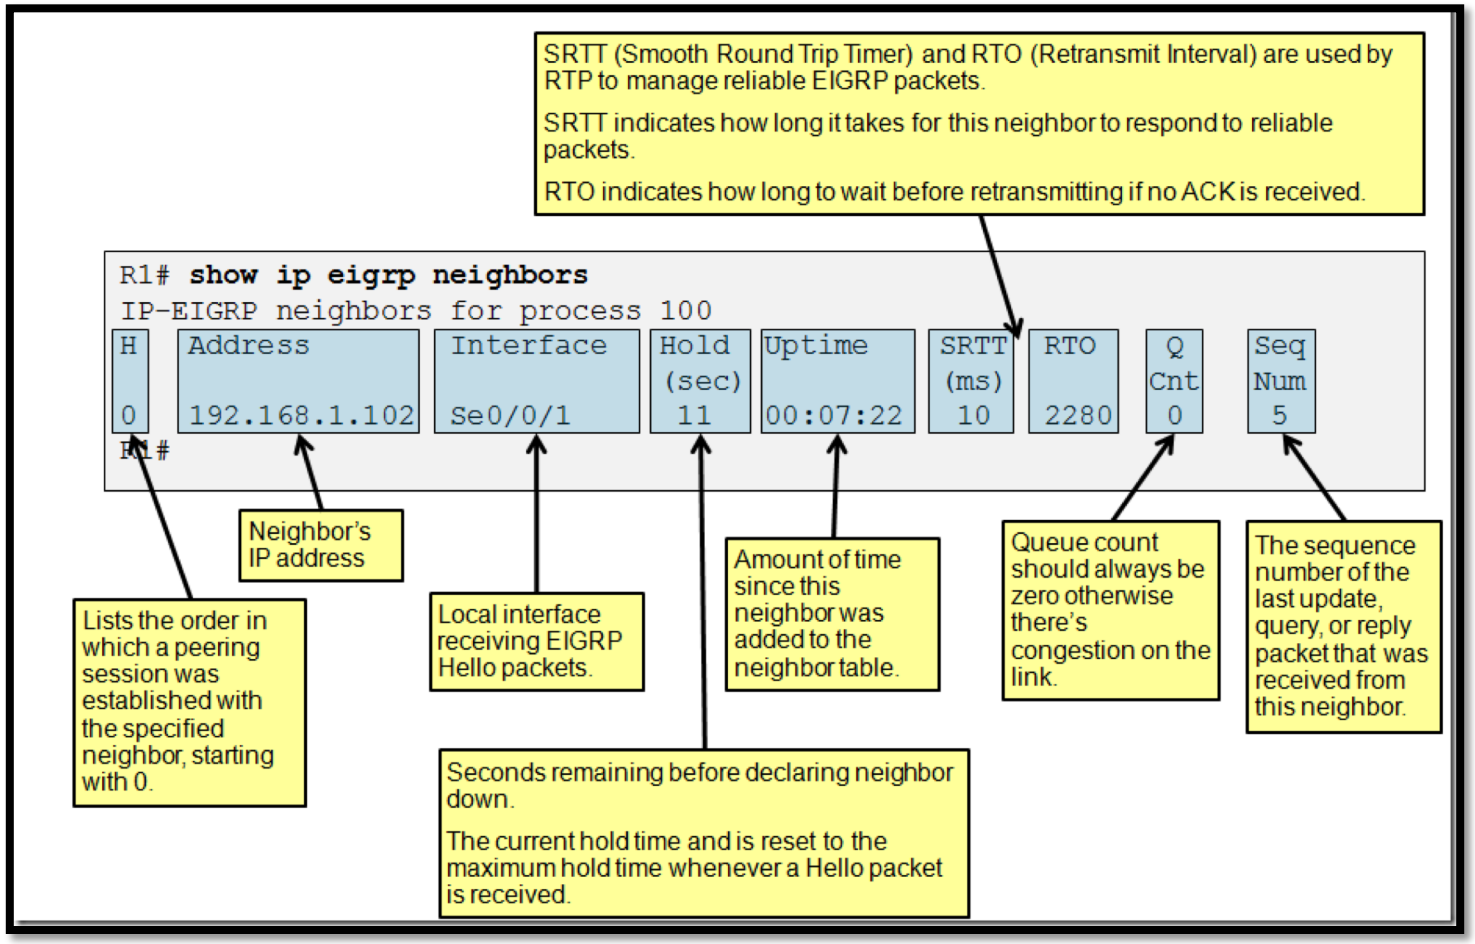

Routing Table

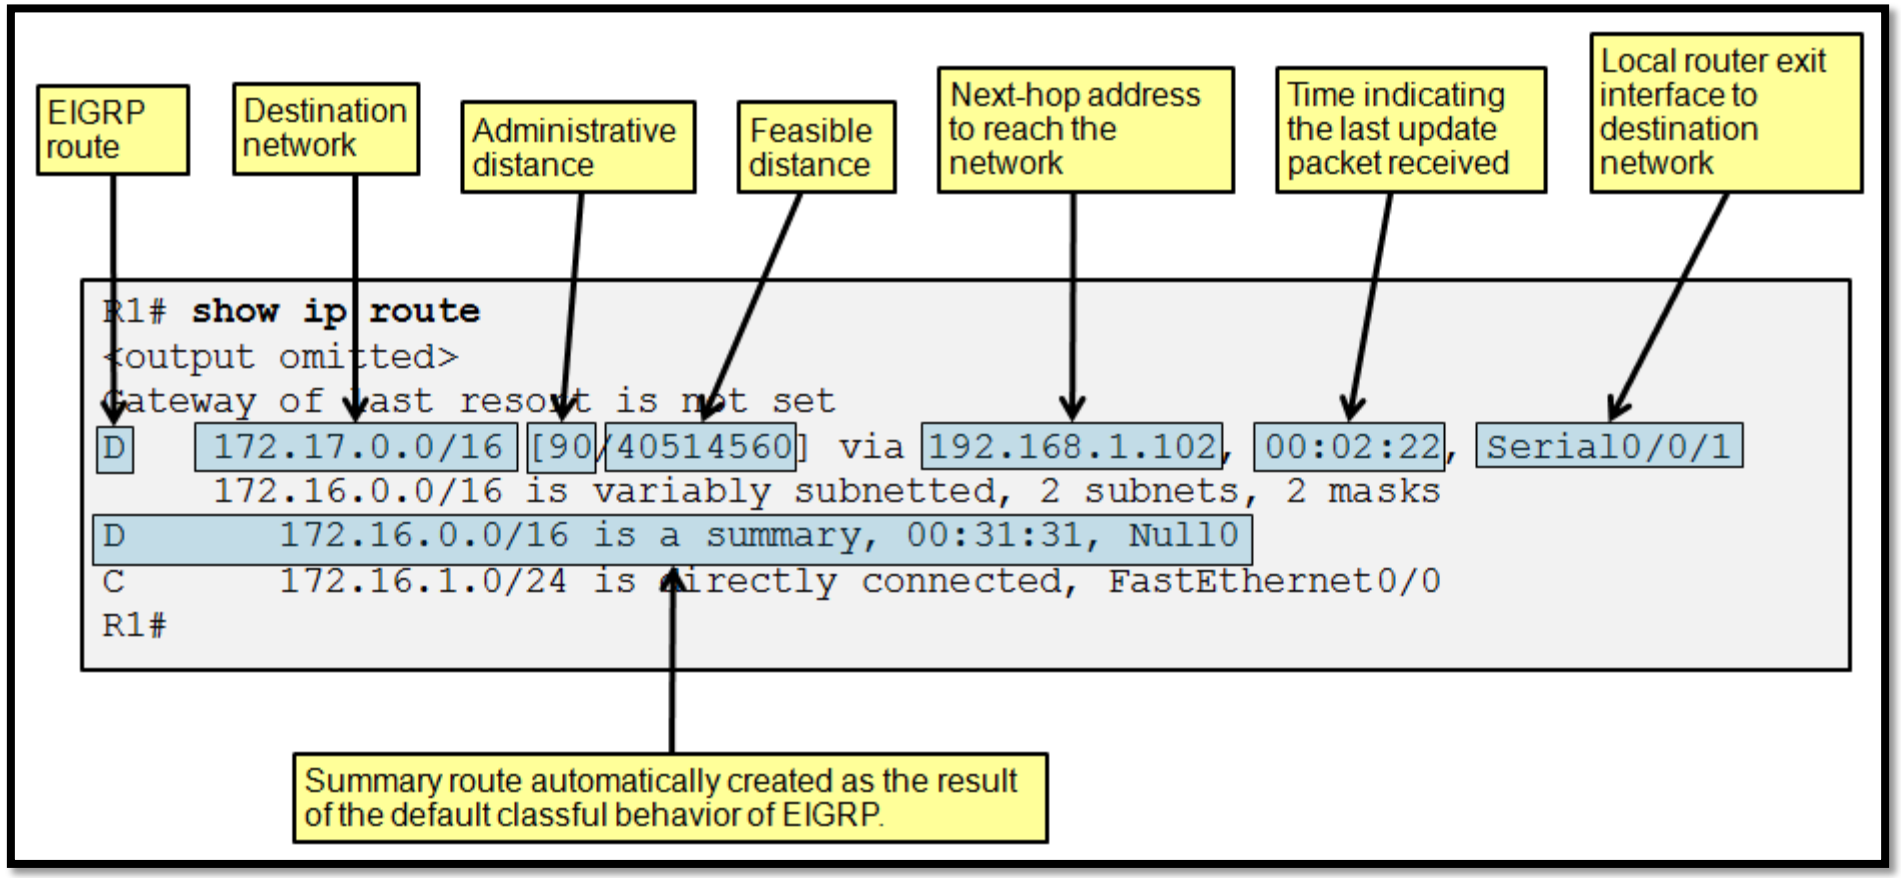

Open Shortest Path Fast (OSPF)

- Routing information is shared through Link-state updates (LSAs)
- It’s a link state protocol
- It uses SPF (shortest path first) or dijkistra algorithm
- Unlimited hop count
- Metric is cost (cost=10 ^8/B.W.)
- Administrative distance is 110
- It is a classless routing protocol. It sends the subnet mask in the routing update.
- It supports VLSM and CIDR
- It supports only equal cost load balancing
- Introduces the concept of Area’s to ease management and control traffic
- Provides hierarchical network design with multiple different areas
- Must have one area called as area 0
- All the areas must connect to area 0
- Scales better than Distance Vector Routing protocols.
- Supports Authentication
- Updates are sent through multicast address 224.0.0.5
- Faster convergence.

Link-State Packets (LSPs)

- There are five types of OSPF Link-State Packets (LSPs).
  - **Hello**: are used to establish and maintain adjacency with other OSPF routers. They are also used to elect the Designated Router (DR) and Backup Designated Router (BDR) on multi-access networks (like Ethernet or Frame Relay).
  - **Database Description** (DBD or DD): contains an abbreviated list of the sending router's link-state database and is used by receiving routers to check against the local link-state database
  - **Link-State Request** (LSR): used by receiving routers to request more information about any entry in the DBD
  - **Link-State Update** (LSU): used to reply to LSRs as well as to announce new information. LSUs contain seven different types of Link-State Advertisements (LSAs)
  - **Link-State Acknowledgement** (LSAck): sent to confirm receipt of an LSU message

OSPF Process

- Seven Tages:
  - 1. Down
  - 2. init
  - 3. 2 way
  - 4. Exstart
  - 5. Exchange
  - 6. Loading
  - 7. Full

Down stage

Hello

I am 192.168.1.100 ,who is on the link – 224.0.0.5

Init stage

(Hello) I am 10.0.0.2 - unicast

Hello

Two-way stage

They become neighbors, built neighbor table

OSPF router ID

- Used to identify the router inside the OSPF database
- OSPF identify using same ID in all directions
- Default uses highest IP address of active physical interface

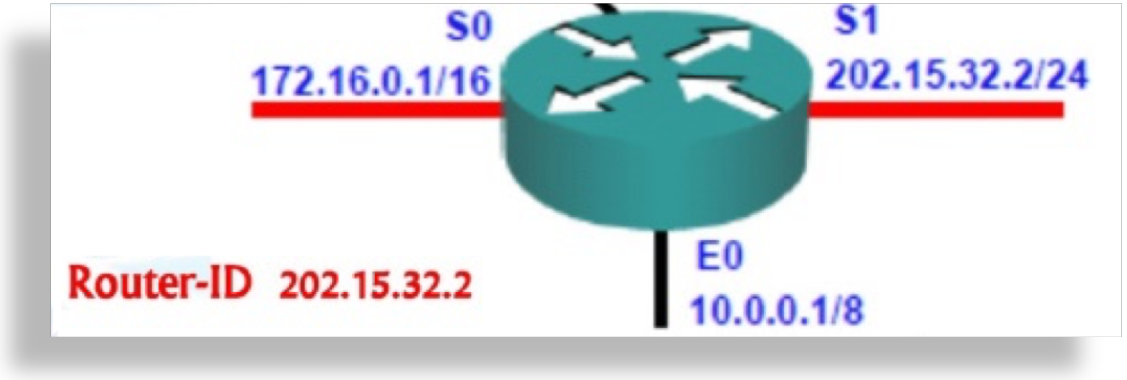

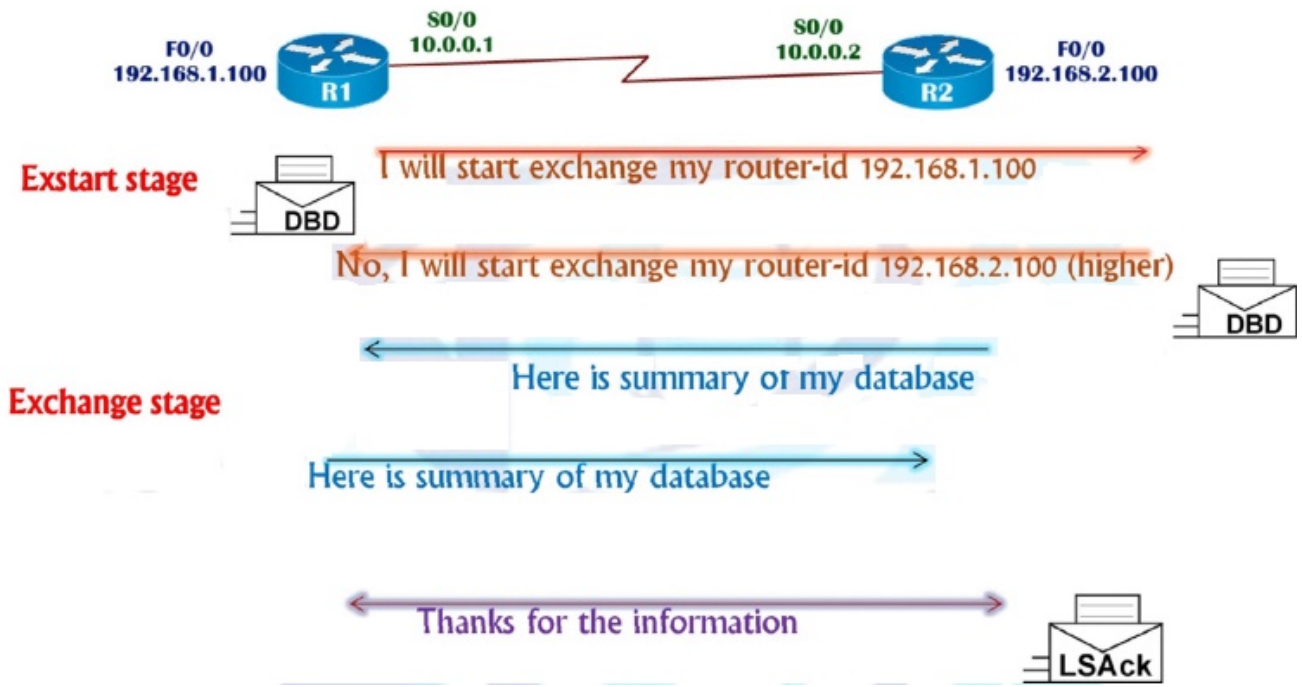
F0/0
192.168.1.100
R1
S0/0
10.0.0.1
S0/0
10.0.0.2
R2
F0/0
192.168.2.100
Exstart stage
DBD
I will start exchange my router-id 192.168.1.100
No, I will start exchange my router-id 192.168.2.100 (higher)
DBD
Here is summary of my database
Exchange stage
Here is summary of my database
Thanks for the information
LSAck

OSPF Tables

- Neighbor Table
  - Contains list of directly connected routes
  - # show ip ospf neighbor
- Routing Table
  - The best route to destination
  - # show ip route
- Database table
  - Referred to as LSDB (link state database)
  - Contains information about all the possible routes to the networks with in the area
  - # show ip ospf database

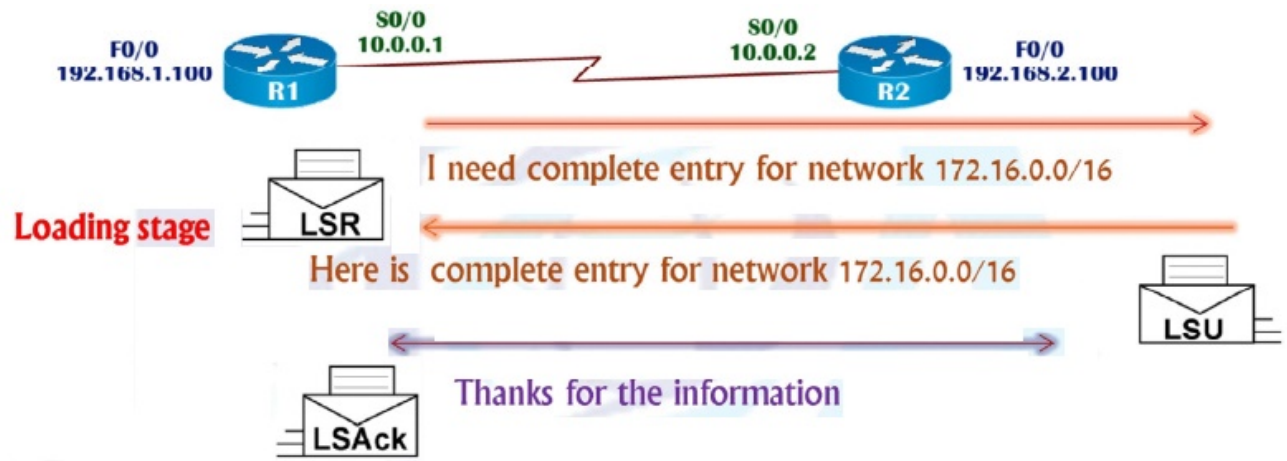
F0/0
192.168.1.100
R1
S0/0
10.0.0.1
S0/0
10.0.0.2
R2
F0/0
192.168.2.100
I need complete entry for network 172.16.0.0/16
LSR
Loading stage
Here is complete entry for network 172.16.0.0/16
LSU
Thanks for the information
LSAck
Full stage
Routers have synchronized database
Calculate the best route (10^8/BW)

OSPF Convergence

- Incremental updates
- Periodically send hello packets are sent every 10 seconds (dead-40 sec)
- Convergence rate is fast (40 sec)

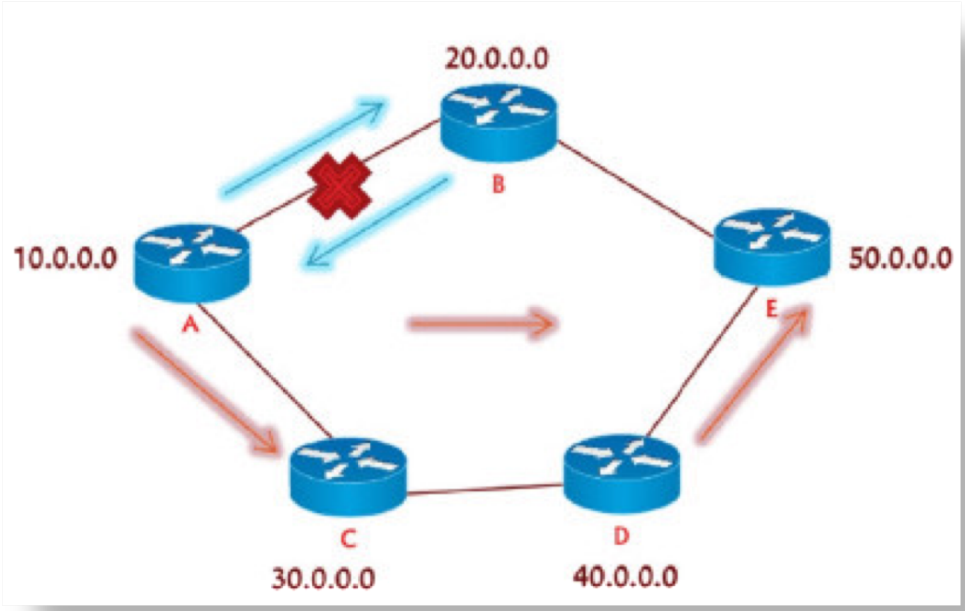

Issue with Maintaining a Large OSPF

- All the routers must have the same database
- Any change in the database advertised to all routers.

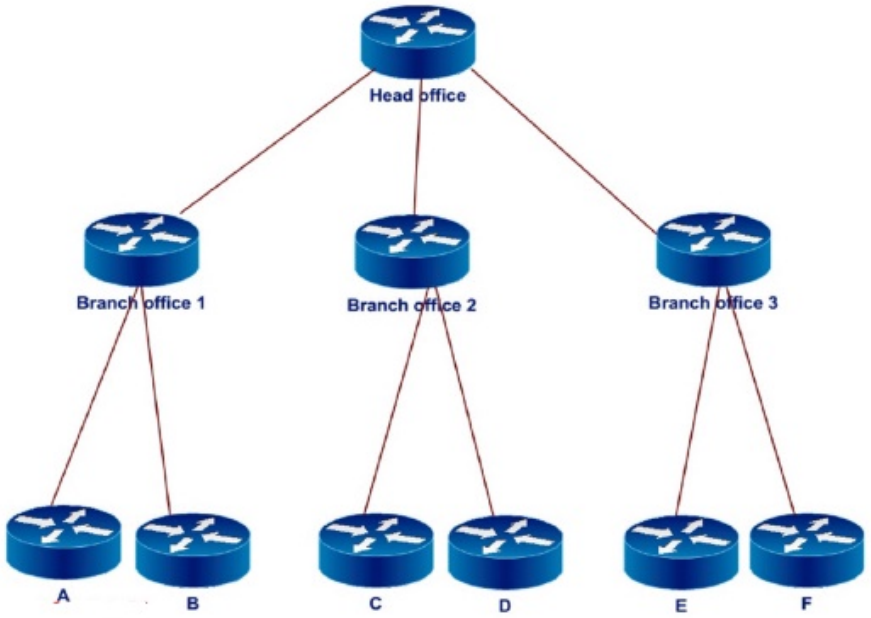

Issue with Maintaining a Large OSPF

- OSPF area: Area is logical grouping of Routers.
  - All the routes must have the save database with in area
  - Any change in the database advertised to all routers with in area.

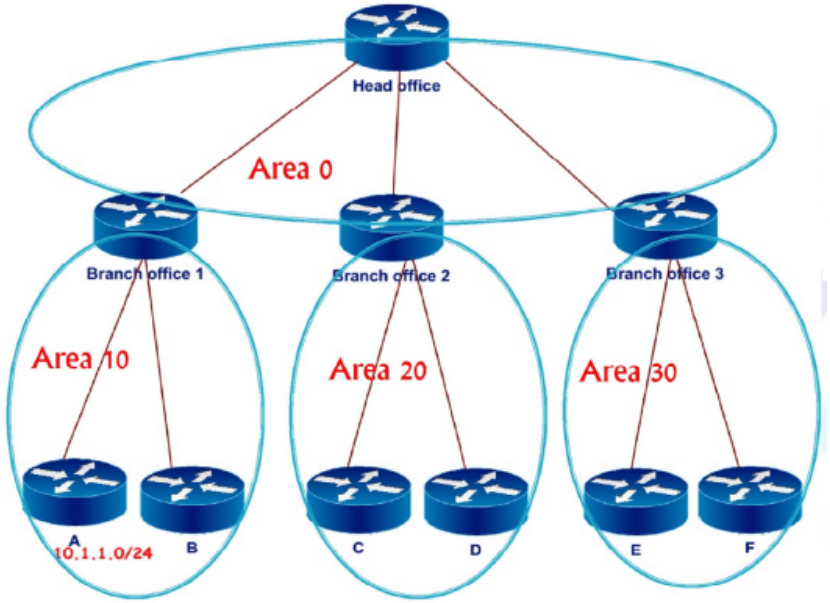

Concept of Area

- Area is logical grouping of Routers
- All the routers maintain same database within the same Area
- Any change impact all the routers with the same area.

- Impact:
  - Minimizes size of database
  - Restrict any changes with in that area. (not flood outside area)
  - Routers with in the same area participate in Algorithm

OSPF area design Rules

- Must have on area called as area 0 (called backbone area)
- All the no-backbone areas must connect to area 0. (Area 0 must be transit area)
- At least one-Area Border Router (connecting two or more area)
- Interface of both routers facing must be in the same area.

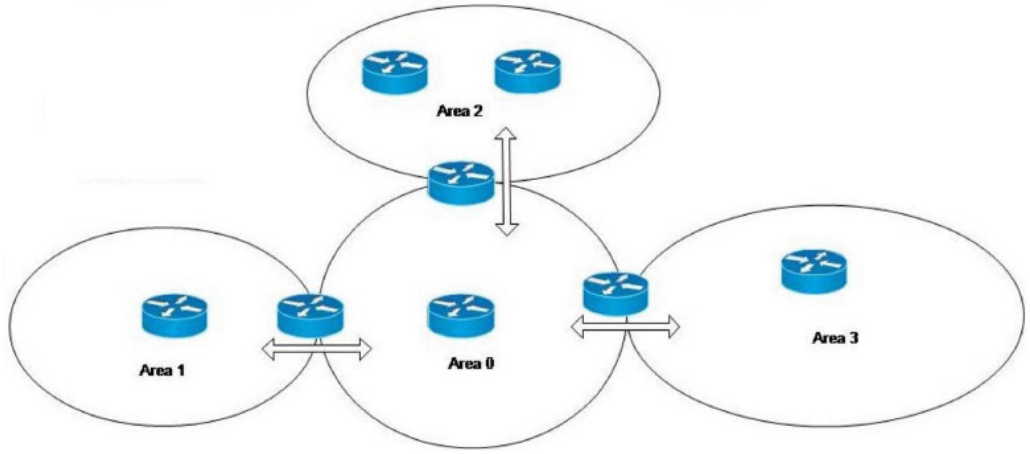

Configuring OSPF

- Router(config)# router ospf <process id>
- Router(config-router)# network <Network Id> < wildcard mask> area <area id>

  - Process ID is a number used to identify an OSPF routing process on the router.
  - Multiple OSPF processes can be started on the same router.
  - The number can be any value between 1 and 65,535