


IP Routing -2

❖ *Dynamic Routing*

❖ *RIP, OSPF*



Introduction

Dynamic routing is when **protocols are used to find networks and update routing tables** on routers. True, this is easier than using static or default routing, but it'll **cost you in terms of router CPU processes and bandwidth on the network links**.

A routing protocol defines the set of rules used by a router when it communicates routing information between neighbor routers. The routing protocol includes **Routing Information Protocol (RIP) versions 1 and 2**, with a bit of **Interior Gateway Routing Protocol (IGRP)** thrown in.

Two types of routing protocols are used in internetworks: **interior gateway protocols (IGPs)** and **exterior gateway protocols (EGPs)**.

IGPs are used to exchange routing information with routers in the same autonomous system (AS). An AS is a collection of networks under a common administrative domain, which basically means that all routers sharing the same routing table information are in the same AS.

EGPs are used to communicate between ASes. An example of an EGP is Border

Gateway Protocol (BGP), which is beyond the scope of our lab.

Administrative Distances

The administrative distance (AD) is used to **rate the trustworthiness of routing information received on a router from a neighbor router**. An administrative distance is an integer from 0 to 255, where 0 is the most trusted and 255 means no traffic will be passed via this route.

If a router receives two updates listing the same remote network, the first thing the router checks is the AD.

If one of the advertised routes has a lower AD than the other, then the route with the lowest AD will be placed in the routing table. If both advertised routes to the same network have the same AD, then routing protocol **metrics** (such as hop count or bandwidth of the lines) will be used to find the best path to the remote network. The advertised route with the lowest metric will be placed in the routing table.

But if both advertised routes have the **same AD as well as the same metrics**, then the routing protocol will **load-balance** to the remote network (which means that it sends packets down each link).

Table below shows default AD

Route Source	Default AD
Connected interface	0
Static route	1
EIGRP	90
IGRP	100
OSPF	110
RIP 1	120
External EIGRP	170
Unknown	266

The smaller the AD is , the more preferable to route is.

Classes of routing protocols

Distance vector

The distance-vector protocols find the best path to a remote network by judging distance. Each time a packet goes through a router, that's called a hop. **The route with the least number of hops to the network is determined to be the best route.** The vector indicates the direction to the remote network. Both RIP and IGRP are distance-vector routing protocols. **They send the entire routing table to directly connected neighbors.**

Link state

In link-state protocols, also called **shortest-path-first protocols**. Link-state routers know more about the internetwork than any distance vector routing protocol. OSPF is an IP routing protocol that is completely link state. Linkstate protocols send updates containing the state of their own links to all other routers on the network.


Hybrid

Hybrid protocols use aspects of both distance vector and link state—for example, EIGRP. There's no set way of configuring routing protocols for use with every business. This is something you really have to do on a case-by-case basis. If you understand how the different routing protocols work, you can make good, solid decisions that truly meet the individual needs of any business.

Distance vector routing protocols


The distance-vector routing algorithm passes complete routing table contents to neighboring routers, which then combine the received routing table entries with their own routing tables to complete the router's routing table. This is called **routing by rumor**, because a router receiving an update from a neighbor router believes the information about remote networks without actually finding out for itself.

Example:



the four routers start off with only their directly connected networks in their routing tables. After a distance-vector routing protocol is started on each router, the routing tables are updated with all route information gathered from neighbor routers.

After convergence, the routing tables will look like this :



Routing loops in DV and how to solve it?

Routing loops can occur because every router isn't updated simultaneously, or even close to it. Here's an example assume network 10.4.0.0 goes down-figure a-, and before R3 advertises that **(by sending routing poisoning message: hope count > 15 = infinity)**, it receives an update form R2 which contains information about network 10.4.0.0 –see figure b-

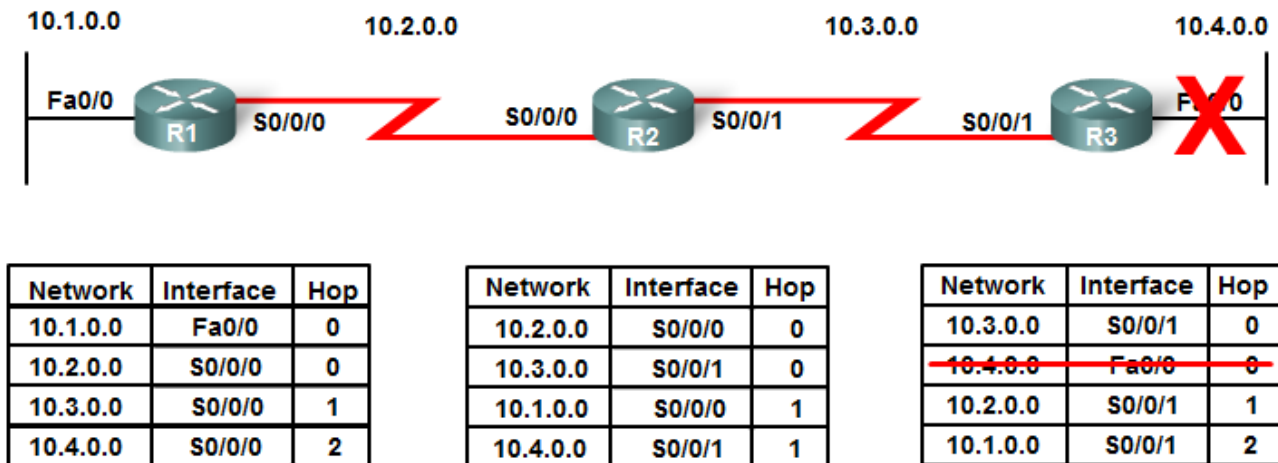


Figure a

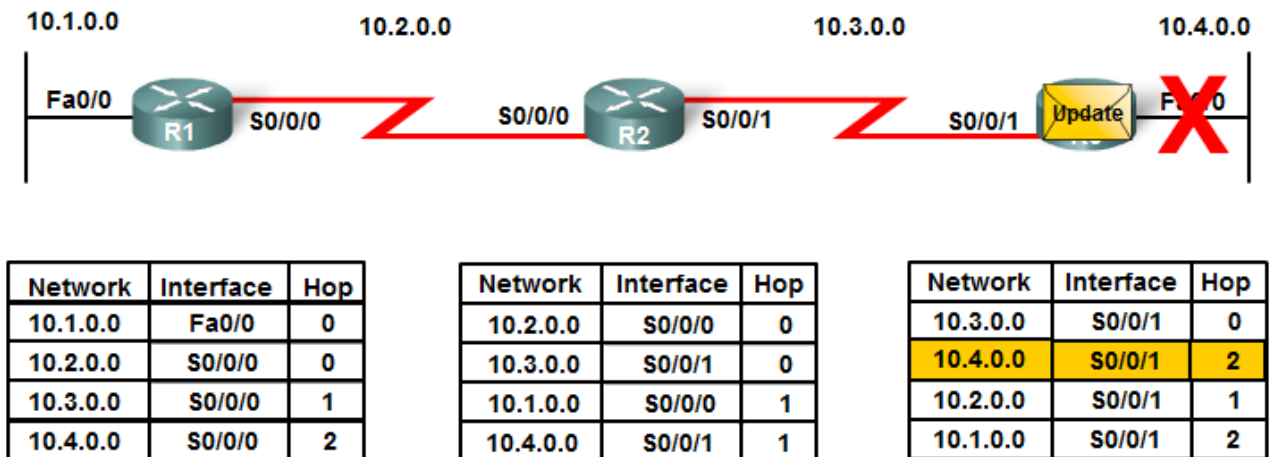
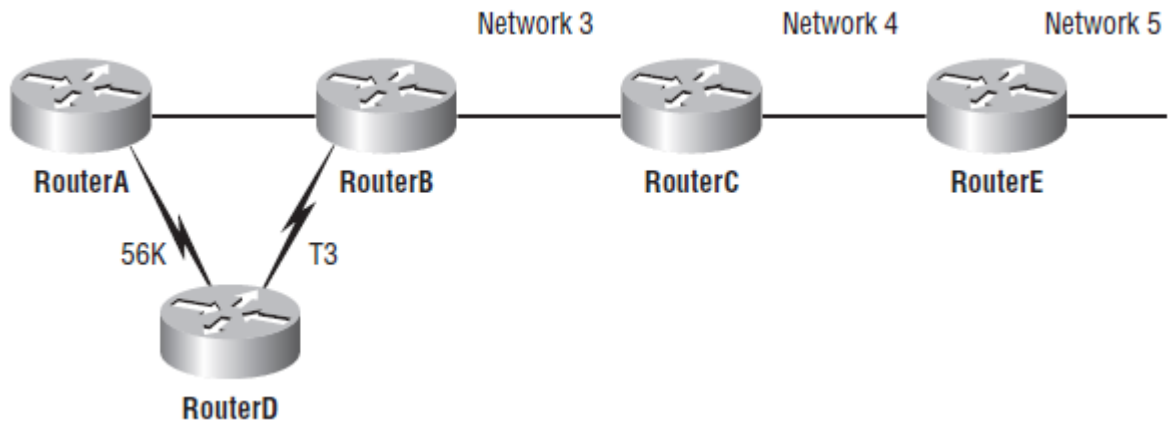


Figure b

Now assume a PC in network 10.1.0.0 tries to send data to a PC in network 10.4.0.0? what will happen?!

Another example —let's say that the interface to Network 5 fails. All routers know about Network 5 from RouterE. RouterA, in its tables, has a path to Network 5 through RouterB.



When Network 5 fails, RouterE tells RouterC **(by sending routing poisoning message: hope count > 15 = infinity)**. This causes RouterC to stop routing to Network 5 through RouterE. But routers A, B, and D don't know about Network 5 yet, so they keep sending out update information. RouterC will eventually send out its update and cause B to stop routing to Network 5, but routers A and D are still not updated. To them, it appears that Network 5 is still available through RouterB with a metric of 3. **The problem occurs** when RouterA sends out its regular 30-second "Hello, I'm still here— these are the links I know about" message, which includes the ability to reach Network 5, and now routers B and D receive the wonderful news that Network 5 can be reached from RouterA.

Loops?

Router A thinks he know how to reach Network 5 form network B by 3 hops and B thinks he know it form A by 4 hops. When B advertise, A will receive that B knows Network 5 by 4 now ! then he alters his table and advertise. Same thing when B receives it. And so on.


Solution :

Maximum hope count ,, Split horizon >>> HOW?

Routing Information Protocol (RIP)

Routing Information Protocol (RIP) is a true distance-vector routing protocol. RIP sends the complete routing table out to all active interfaces every **30 seconds**. **RIP only uses hop count to determine the best way to a remote network**, but it has a maximum allowable **hop count of 15** by default, meaning that 16 is deemed unreachable. RIP works well in small networks, but it's inefficient on large networks with slow WAN links or on networks with a large number of routers installed.

What will happen using RIP?



RIP Timers

RIP uses four different kinds of timers to regulate its performance:

Route update timer

Sets the interval (typically 30 seconds) between periodic routing updates in which the router sends a complete copy of its routing table out to all neighbors.

Route invalid timer

If an update has not been received to refresh an existing route after 180 seconds (the default), the route is marked as invalid by setting the metric to 16. The route is retained in the routing table until the flush timer expires.

Hold down timer

This sets the amount of time during which routing information is suppressed. Routes will enter into the holddown state when an update packet is received that indicated the route is unreachable. This continues either until an update packet is received with a better metric or until the holddown timer expires. The default is 180 seconds.

(see exploration2 4.4.4.1)

Route flush timer

Sets the time between a route becoming invalid and its removal from the routing table (240 seconds). Before it's removed from the table, the router notifies its neighbors of that route's impending demise. The value of the route invalid timer must be less than that of the route flush timer. This gives the router enough time to tell its neighbors about the invalid route before the local routing table is updated.

RIP Versions

	RTP V1	RIP V2
AD	120	120
Metric	Hope count	Hope count
Max hope count	15	15
Sending update every	30 sec	30 sec
Sending updates using	Broadcast	multicast
VLSM\CIDR	Not supported	Supported

RIP practical part

RIPV1

Lab_A#**config t**

Lab_A(config)#**router rip**

Lab_A(config-router)#**network 192.168.10.0 (only net without mask)**

Lab_A(config-router)#**passive-interface serial 0/0** (This command prevents RIP update broadcasts from being sent out a specified interface, yet that same interface can still receive RIP updates.)

RIPV2

RIPv2 is considered classless because subnet information is sent with each route update

```
Lab_C(config)#router rip  
Lab_C(config-router)#network 192.168.40.0  
Lab_C(config-router)#network 192.168.50.0  
Lab_C(config-router)#version 2
```

Interior Gateway Routing Protocol (IGRP)

No longer supported by CISCO so we won't waste our time learning it ☺.

Open Shortest Path First (OSPF)

OSPF works by using the **Dijkstra algorithm**. First, a shortest path tree is constructed, and then the routing table is populated with the resulting best paths. OSPF converges quickly, although perhaps not as quickly as EIGRP, and it supports multiple, equal-cost routes to the same destination. Like EIGRP, it does support both IP and IPv6 routed protocols.

OSPF provides the following features:

- Consists of areas and autonomous systems
- Minimizes routing update traffic
- Allows scalability
- Supports VLSM/CIDR
- Has unlimited hop count
- Allows multi-vendor deployment (open standard)

OSPF Terminology

Router ID

The Router ID (RID) is an IP address used to identify the router. Cisco chooses the Router ID by using the highest IP address of all configured interfaces.

Neighbor

Neighbors are two or more routers that have an interface on a common network, such as two routers connected on a point-to-point serial link.

Adjacency

An adjacency is a relationship between two OSPF routers that permits the direct exchange of route updates. OSPF shares routes only with neighbors that have also established adjacencies.

Hello protocol

The OSPF Hello protocol provides dynamic neighbor discovery and maintains neighbor relationships.

Neighborship database

The neighborship database is a list of all OSPF routers for which Hello packets have been seen.

Link State Advertisement

A Link State Advertisement (LSA) is an OSPF data **packet** containing link-state and routing information that's shared among OSPF routers.(contains directly connected interface, cost, type and between whom? Ex:

R1-R2 Serial network 10.4.0.0 cost 40).

Topological database

The topological database contains information from all of the Link State Advertisement packets that have been received for an area.

Designated router

A Designated Router (DR) is elected whenever OSPF routers are connected to the same **multi-access network** (later). DR is chosen (elected) to disseminate/receive routing information to/from the remaining routers. This ensures that their topology tables are synchronized.

The DR is the one with the highest router ID.

Backup designated router

A Backup Designated Router (BDR) is a hot standby for the DR on BDR receives all routing updates from OSPF adjacent routers but doesn't flood LSA updates.

OSPF areas

An OSPF area is a grouping of contiguous networks and routers. All routers in the same area share a common Area ID. Areas also play a role in establishing a hierarchical network organization—something that really enhances the scalability of OSPF!

Loopback Address

If the OSPF router-id command is not used and loopback interfaces are configured, OSPF will choose highest IP address of any of its loopback interfaces. A loopback address is a virtual interface and is automatically in the up state when configured. You already know the commands to configure a loopback interface:

```
Router(config)#interface loopback 0
```

```
Router(config-if)#ip address 10.3.0.0 255.255.255.255
```

Link costs

Simply by dividing 10^8 on every interface bandwidth ☺ see table below

Interface Type	$10^8/\text{bps} = \text{Cost}$
Fast Ethernet and faster	$10^8/100,000,000 \text{ bps} = 1$
Ethernet	$10^8/10,000,000 \text{ bps} = 10$
E1	$10^8/2,048,000 \text{ bps} = 48$
T1	$10^8/1,544,000 \text{ bps} = 64$
128 kbps	$10^8/128,000 \text{ bps} = 781$
64 kbps	$10^8/64,000 \text{ bps} = 1562$
56 kbps	$10^8/56,000 \text{ bps} = 1785$

Hint : When the serial interface is not actually operating at the default speed, the interface requires manual modification. Both sides of the link should be configured to have the same value.

```
Router(config)#interface s0/0  
Router(config-if)# Bandwidth 64
```

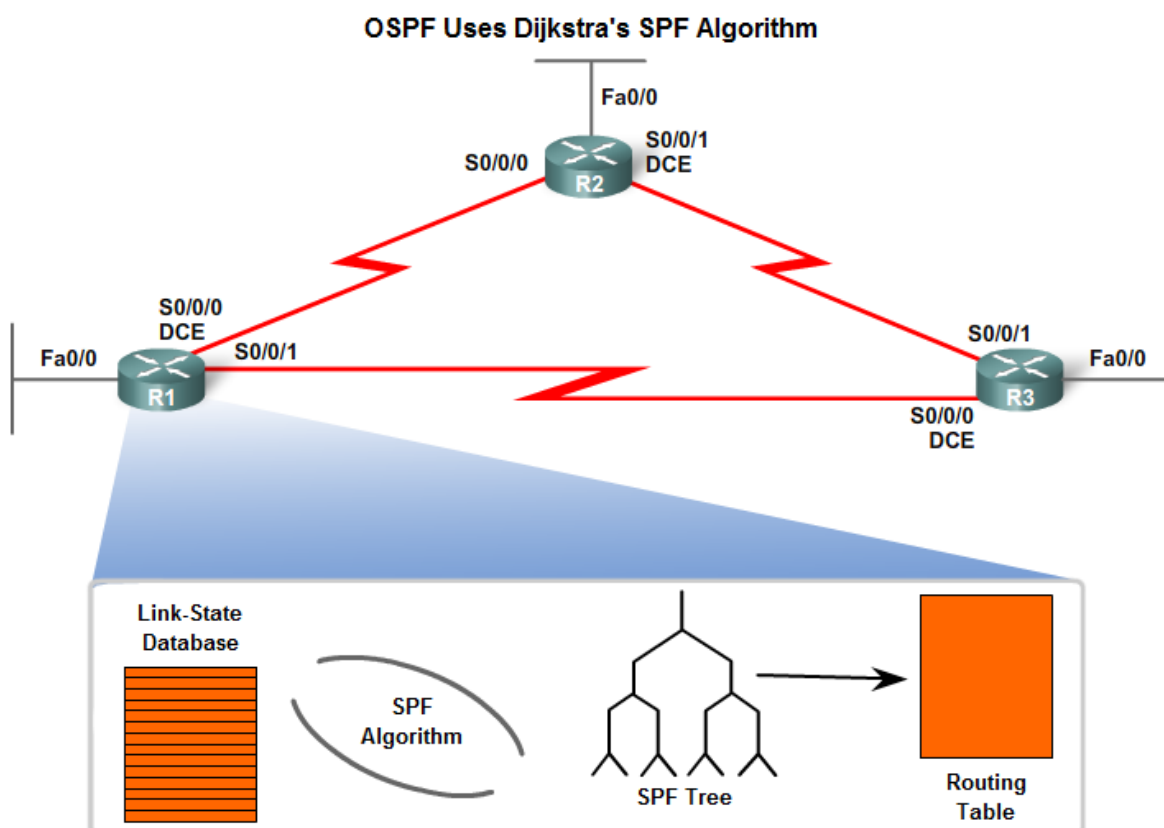
Then the interface cost will be 1562 ☺ Or simply

```
R1(config)#interface serial 0/0/0  
R1(config-if)#ip ospf cost 1562
```

Multi access networks

See exploration2 : 11.4.1.4 and 11.4.1.5

Everything ready? Start *shortest path first (SPF)* algorithm ☺ who?



OSPF practical part

Lab_A# confi g t
Lab_A(config)# router ospf 1
1 : is OSPF process number : out of scope for CCNA, range : <1-65535>
Lab_A(config-router)# network 10.0.0.0 0.255.255.255 area 0
0.255.255.255 : an example of wildcard

Wildcard:

The wildcard mask can be configured as the inverse of a subnet mask. For example, IP 172.16.1.16/28 network. The subnet mask for this interface is /28 or 255.255.255.240. The inverse of the subnet mask results in the wildcard mask .

255.255.255.255	
-255.255.255.252	Subtract the subnet mask

0. 0. 0. 15	Wildcard mask

Next lab

- EIGRP
- Layer 2 Switching and Spanning Tree Protocol (STP)