



- A router may receive multiple possible paths to get to a destination network
- Only the best path will make it into the routing table and be used
- The different IGPs use different methods to calculate the best path to a destination network

Metric



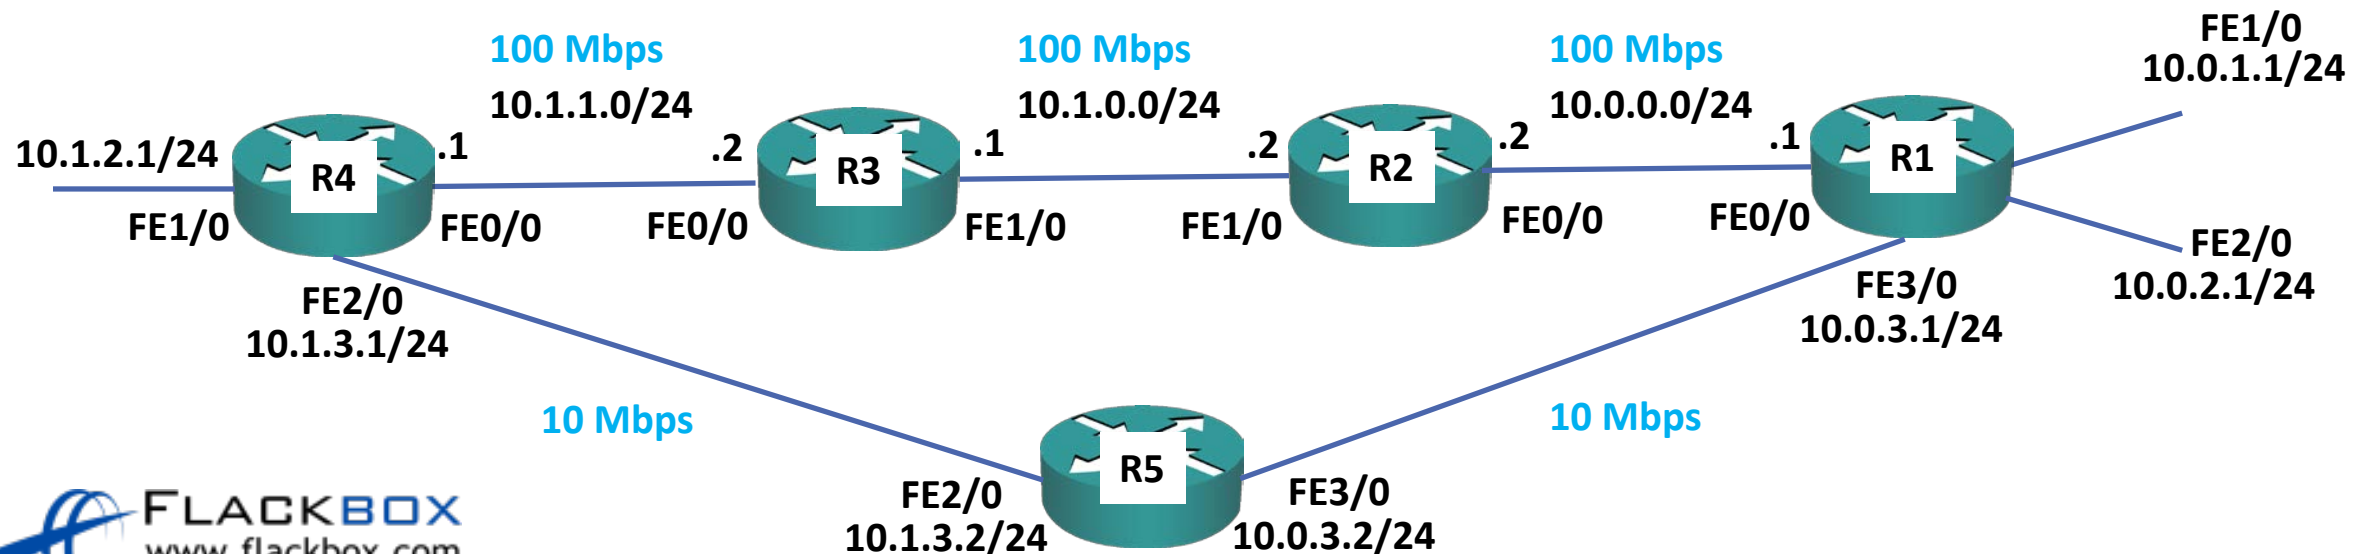
- Each possible path will be assigned a 'metric' value by the routing protocol which indicates how preferred the path is
- The lowest metric value is preferred
- Distance Vector routers advertise to each other the networks they know about, and their metric to get to each of them
- Link State routers advertise all the links in their area of the network to each other
- Each router will take this information and then make an independent calculation of its own best path to get to each destination



- If the best path to a destination is lost (for example because a link went down) it will be removed from the routing table and replaced with the next best route

RIP Metric – Hop Count

- RIP uses Hop Count as the metric
- The maximum hop count by default is 15. Paths which are more than 15 hops away are marked as unreachable
- Path R4>5>1 will be preferred for 10.0.1.0/24 in the example below
- RIP is typically used only in small or test environments



RIP Metric – Hop Count

- RIP uses Hop Count as the metric

R1: "You can get to these networks via me":

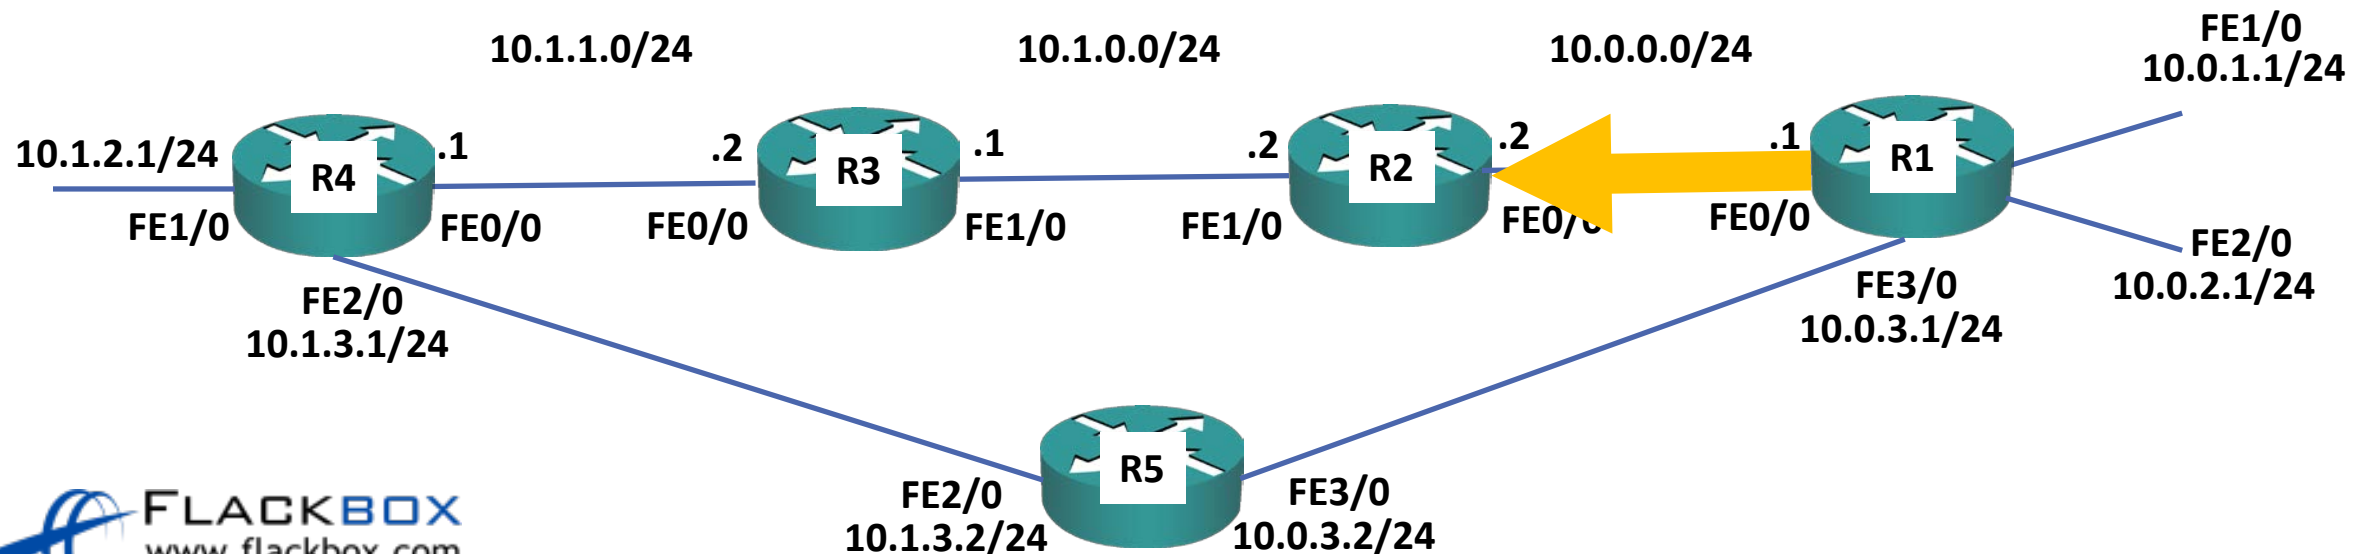
10.0.1.0/24 – 1 hop

10.0.2.0/24 – 1 hop

10.0.3.0/24 – 1 hop

10.1.3.0/24 – 2 hops

10.1.2.0/24 – 3 hops



RIP Metric – Hop Count



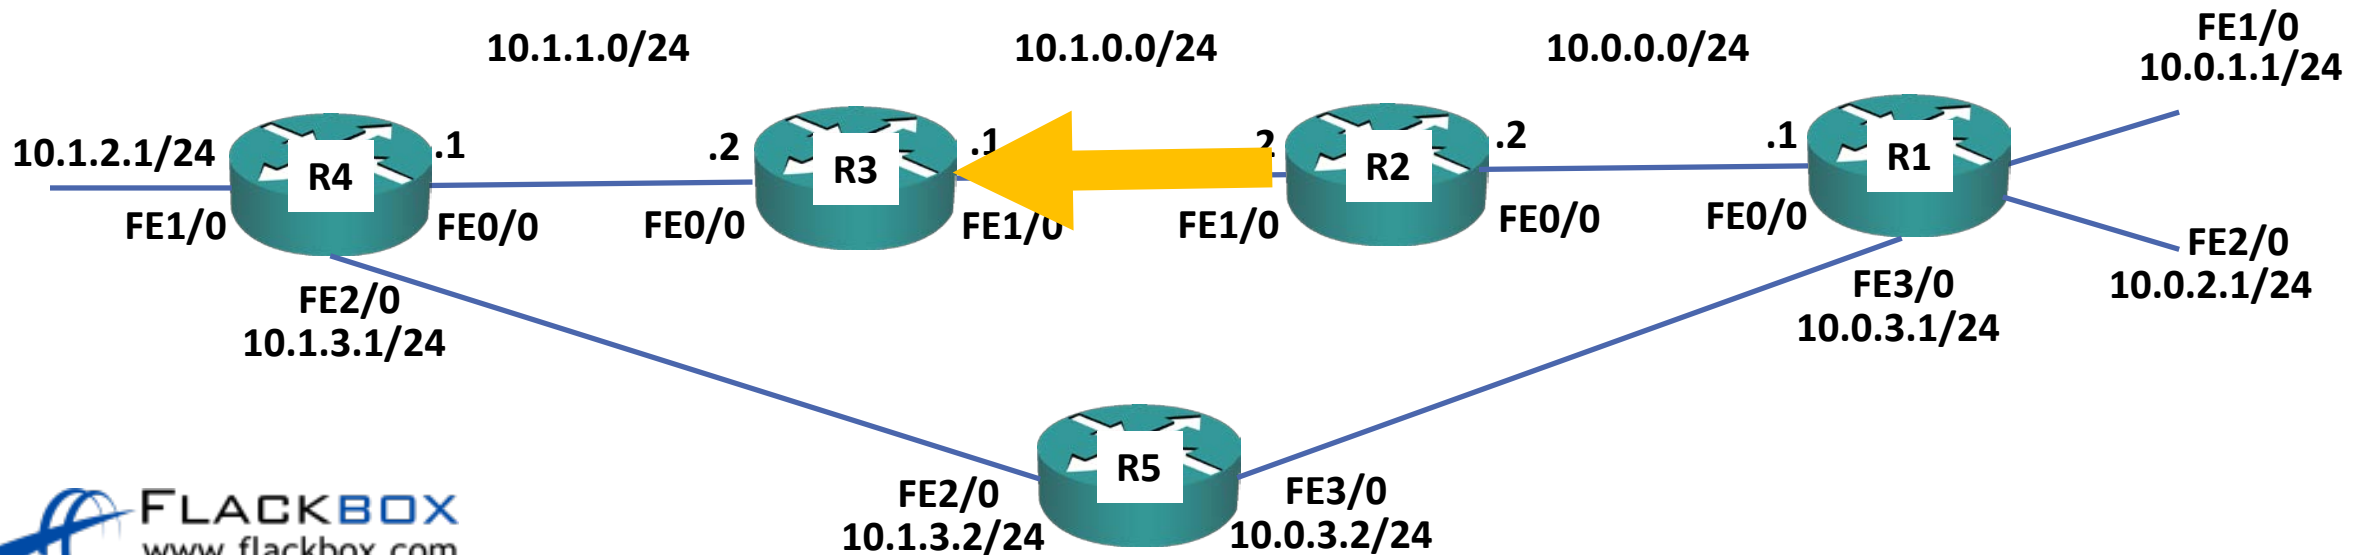
R2: "You can get to these networks via me":

10.0.0.0/24 – 1 hop

10.0.1.0/24 – 2 hops

10.0.2.0/24 – 2 hops

10.0.3.0/24 – 2 hops



RIP Metric – Hop Count



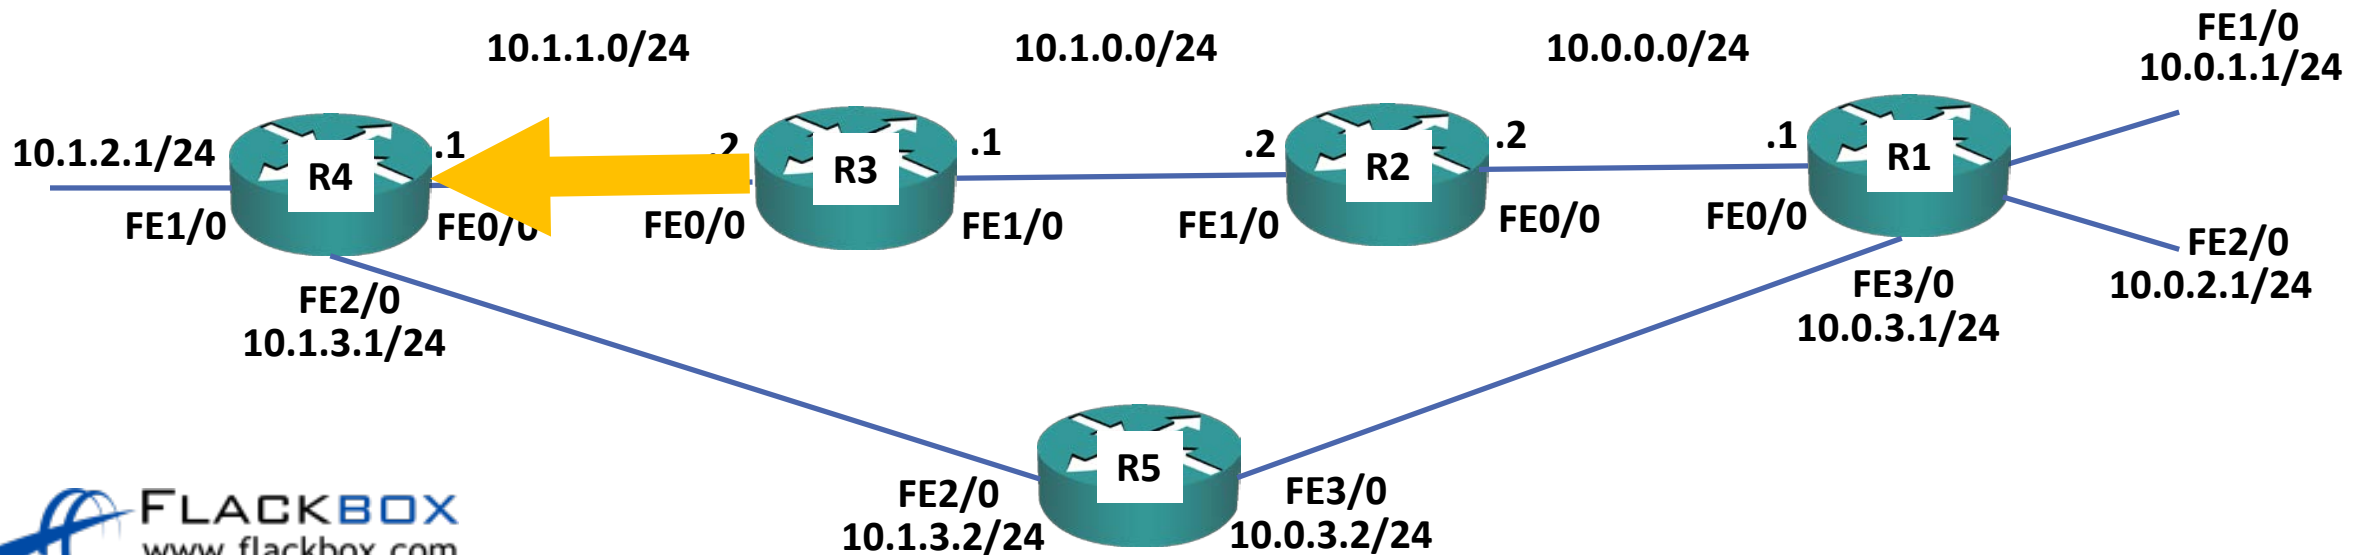
R3: “You can get to these networks via me”:

10.0.0.0/24 – 2 hops

10.0.1.0/24 – 3 hops

10.0.2.0/24 – 3 hops

10.1.0.0/24 – 1 hop



RIP Metric – Hop Count




R3: “You can get to these networks via me”:

10.0.0.0/24 – 2 hops

10.0.1.0/24 – 3 hops

10.0.2.0/24 – 3 hops

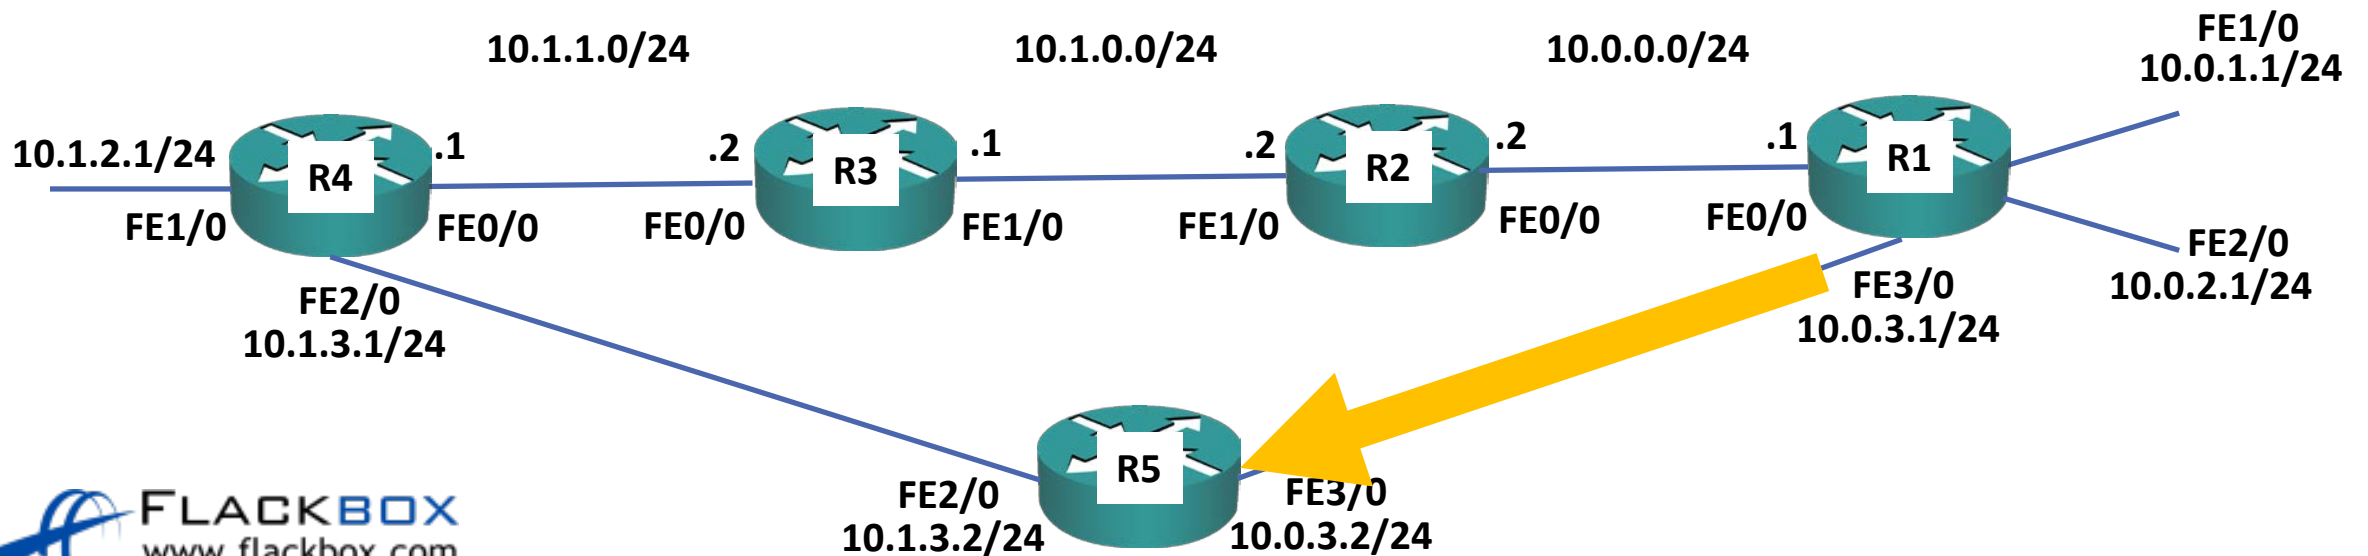
10.1.0.0/24 – 1 hop



RIP Metric – Hop Count



R1: "You can get to these networks via me":
10.0.0.0/24 – 1 hop
10.0.1.0/24 – 1 hop
10.0.2.0/24 – 1 hop
10.1.0.0/24 – 2 hops



RIP Metric – Hop Count



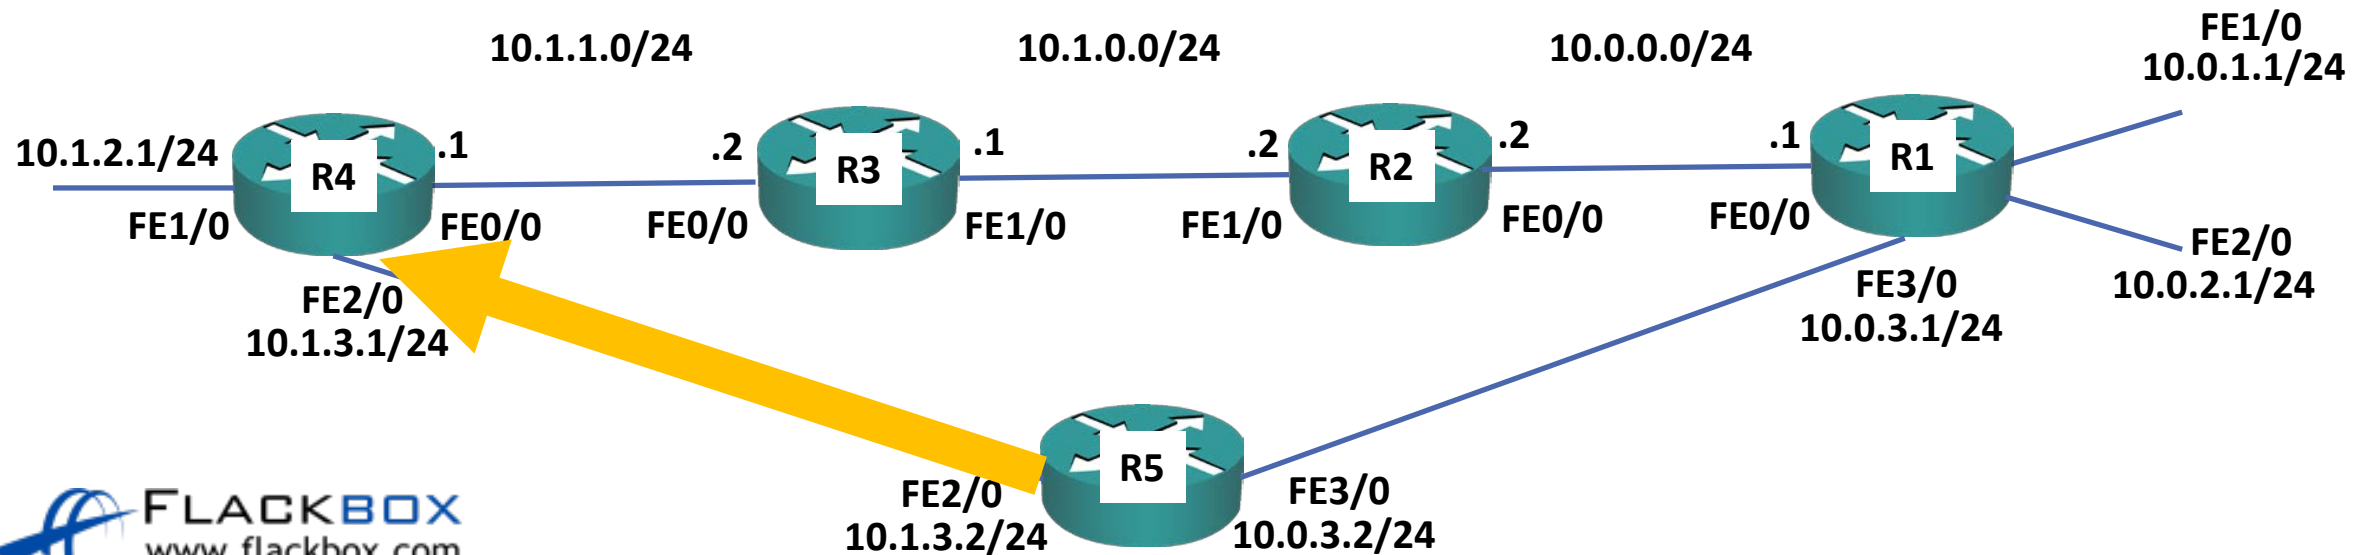
R5: "You can get to these networks via me":

10.0.0.0/24 – 2 hops

10.0.1.0/24 – 2 hops

10.0.2.0/24 – 2 hops

10.0.3.0/24 – 1 hops



RIP Metric – Hop Count




R5: "You can get to these networks via me":

10.0.0.0/24 – 2 hops

10.0.1.0/24 – 2 hops


10.0.2.0/24 – 2 hops

10.0.3.0/24 – 1 hops



RIP Metric – Hop Count

R4: I learned 2 possible routes to get to the 10.0.1.0/24 network:
3 hops via 10.1.1.2 out FE0/0
2 hops via 10.1.3.2 out F2/0
I'll put the best one in my routing table



RIP Metric – Hop Count




R4#sh ip route

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 11 subnets, 2 masks
R   10.0.0.0/24 [120/2] via 10.1.3.2, 00:00:06, FastEthernet2/0
    [120/2] via 10.1.1.2, 00:00:16, FastEthernet0/0
R   10.0.1.0/24 [120/2] via 10.1.3.2, 00:00:06, FastEthernet2/0
R   10.0.2.0/24 [120/2] via 10.1.3.2, 00:00:06, FastEthernet2/0
R   10.0.3.0/24 [120/1] via 10.1.3.2, 00:00:06, FastEthernet2/0
R   10.1.0.0/24 [120/1] via 10.1.1.2, 00:00:16, FastEthernet0/0
C   10.1.1.0/24 is directly connected, FastEthernet0/0
L   10.1.1.1/32 is directly connected, FastEthernet0/0
C   10.1.2.0/24 is directly connected, FastEthernet1/0
L   10.1.2.1/32 is directly connected, FastEthernet1/0
C   10.1.3.0/24 is directly connected, FastEthernet2/0
L   10.1.3.1/32 is directly connected, FastEthernet2/0
203.0.113.0/24 is variably subnetted, 2 subnets, 2 masks
C   203.0.113.0/24 is directly connected, FastEthernet3/0
L   203.0.113.1/32 is directly connected, FastEthernet3/0
```

OSPF Metric – Cost

- OSPF uses 'Cost' as the metric, which is automatically derived from interface bandwidth by default
- You can manually configure the cost of links if you want to manipulate the path
- Path R4>3>2>1 will be preferred for 10.0.1.0/24 in the example below



OSPF Metric – Cost



R4#sh ip route


Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 11 subnets, 2 masks

- O 10.0.0.0/24 [110/3] via 10.1.1.2, 00:03:57, FastEthernet0/0
- O 10.0.1.0/24 [110/4] via 10.1.1.2, 00:03:57, FastEthernet0/0**
- O 10.0.2.0/24 [110/4] via 10.1.1.2, 00:03:57, FastEthernet0/0
- O 10.0.3.0/24 [110/13] via 10.1.1.2, 00:03:57, FastEthernet0/0
- O 10.1.0.0/24 [110/2] via 10.1.1.2, 00:03:57, FastEthernet0/0
- C 10.1.1.0/24 is directly connected, FastEthernet0/0
- L 10.1.1.1/32 is directly connected, FastEthernet0/0
- C 10.1.2.0/24 is directly connected, FastEthernet1/0
- L 10.1.2.1/32 is directly connected, FastEthernet1/0
- C 10.1.3.0/24 is directly connected, FastEthernet2/0
- L 10.1.3.1/32 is directly connected, FastEthernet2/0
- 203.0.113.0/24 is variably subnetted, 2 subnets, 2 masks
- C 203.0.113.0/24 is directly connected, FastEthernet3/0
- L 203.0.113.1/32 is directly connected, FastEthernet3/0

IS-IS Metric – Cost


- IS-IS also uses 'Cost' as the metric, but it is not automatically derived from interface bandwidth. All links have an equal cost by default
- You can manually configure the cost of links if you want to manipulate the path
- If you do not manually set the link costs then the path with the lowest hop count will be used
- Path R4>5>1 will be preferred for 10.0.1.0/24 in the example below



EIGRP Metric



- EIGRP uses the bandwidth and delay of links to calculate the metric
- (Load and reliability can also be considered but are ignored by default)
- A fixed delay value is used based on the interface bandwidth, the protocol does not dynamically measure current delay
- You can manually configure the delay on links if you want to manipulate the path
- Path R4>3>2>1 will be preferred for 10.0.1.0/24 in the example below



Choosing a Routing Protocol



- RIP uses hop count and has a default maximum metric of 15. It is not usually used in production networks because of its scalability limitations.
- EIGRP is very simple to maintain, calculates changes very quickly and its metric calculation will normally choose the best path by default. It is typically only supported on Cisco routers however.
- OSPF's metric calculation will typically choose the best path by default. It is an open standard which is supported by all vendor's routers and is the most commonly deployed IGP today. It is however more complicated to maintain than EIGRP.
- IS-IS links need to be manually configured or it will use hop count to determine the best path. It is typically only used in Service Provider networks or large organisations with their own MPLS network who choose it because of its scalability.

Lab

