

## Chapter 4: Manipulating Routing Updates



# CCNP ROUTE: Implementing IP Routing

Cisco | Networking Academy®  
Mind Wide Open™



# Chapter 4 Objectives

- Network performance issues and ways to control routing updates and traffic.
- The purpose of and considerations for using multiple routing protocols in a network.
- Configure and verify route redistribution of multiple protocols.
- Various methods for controlling routing update traffic.



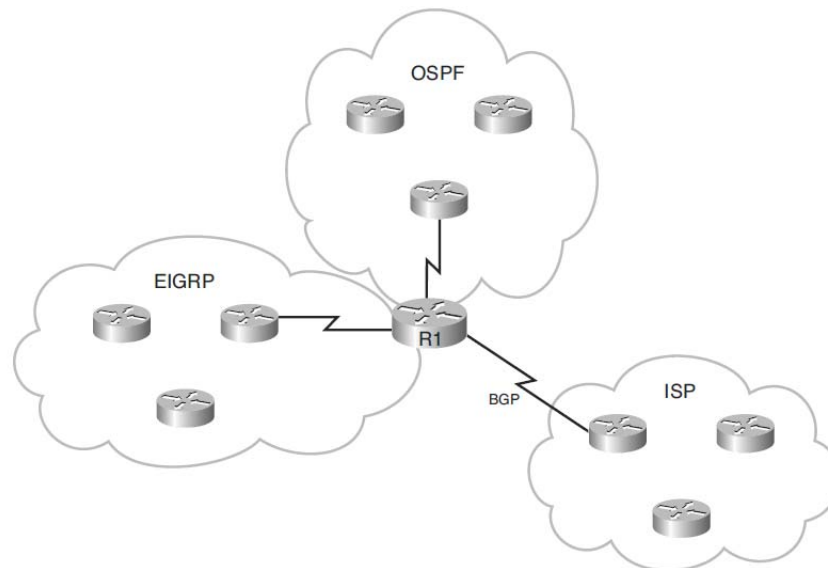
# Common Routing Performance Issues

- Excessive routing updates
  - CPU utilization can easily spike during this processing depending on:
    - The size of the routing update
    - The frequency of the updates
    - The Layer 3 network design
- The presence of any incorrectly configured route maps or filters
- The number of routing protocols running in the same autonomous system



# Running Multiple Protocols

- Different routing protocols were not designed to interoperate with one another.
  - Each protocol collects different types of information and reacts to topology changes in its own way.
  - Redistribution is the capability of a boundary router to exchange and advertise routing information between domains.
- Running multiple routing protocols increases CPU utilization and requires more memory resources to maintain all the topology, database and routing table





# Routing Protocol Performance Solutions

- Design changes, such as limiting the number of routing protocols used.
- Using passive interfaces to prevent routing protocol updates from being advertised out an interface.
- Route filtering techniques to block specific routes from being advertised:
  - Access control lists (ACLs)
  - Route maps
  - Distribute lists
  - Prefix lists



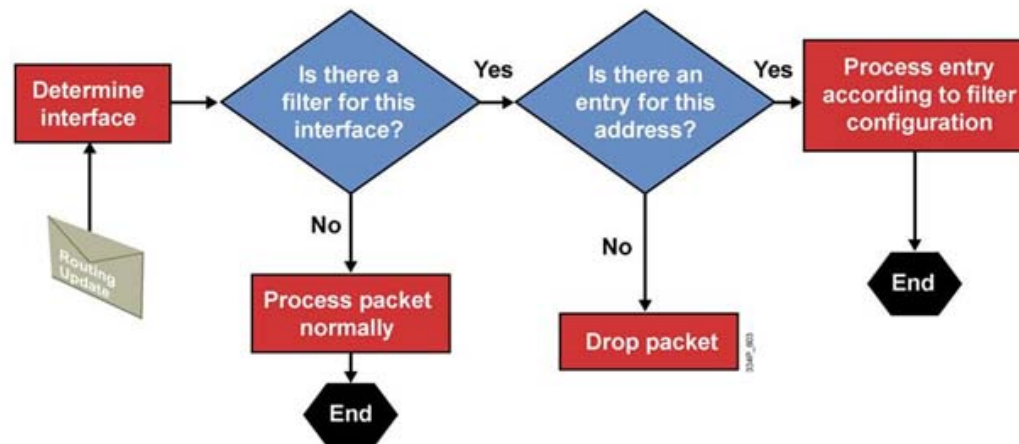
# Route Filtering

- Using route maps, distribute lists, or prefix lists instead of access lists provides greater route filtering flexibility.
- Filters can be configured to:
  - Prevent updates through router interfaces.
  - Control the advertising of routes in routing updates.
  - Control the processing of routing updates.
- If filters are not configured correctly or if filters are applied to wrong interfaces, network performance issues may occur.



# Route Filtering Process

1. A router stores the incoming routing update in the buffer and triggers a decision.
2. Is there an incoming filter applied to this interface?
  - If no, then the routing update packet is processed normally.
3. Otherwise, is there an entry in the filter matching the routing update packet?
  - If no, then the routing update packet is dropped.
4. Otherwise, the router processes the routing update according to the filter.





# Multiple routing protocols

- Simple routing protocols work well for simple networks.
  - Typically only require one routing protocol.
- Running a single routing protocol throughout your entire IP internetwork is desirable.
- However, as networks grow they become more complex and large internetworks may have to support several routing protocols.
  - Proper inter-routing protocol exchange is vital.





# Why have multiple routing protocols?

- Interim during conversion
  - Migrating from an older IGP to a new IGP.
- Application-specific protocols
  - One size does not always fit all.
- Political boundaries
  - Multiple departments managed by different network administrators
  - Groups that do not work well with others
- Mismatch between devices
  - Multivendor interoperability
- Company mergers



# Complex Networks

- Complex networks require careful routing protocol design and traffic optimization solutions, including the following:
  - Redistribution between routing protocols
  - Route filtering (covered in the next chapter)
  - Summarization (covered in EIGRP and OSPF)



# Redistribution

- Cisco routers allow different routing protocols to exchange routing information through a feature called *route redistribution*.
  - Route redistribution is defined as the capability of boundary routers connecting different routing domains to exchange and advertise routing information between those routing domains (autonomous systems).



# Redistributed Routes

- Redistribution is always performed outbound;
  - the router doing redistribution does not change its routing table.
- The boundary router's neighbors see the redistributed routes as external routes.
- Routes must be in the routing table for them to be redistributed.



# Redistribution Considerations

- The key issues that arise when using redistribution:
  - **Routing loops**
    - If more than one boundary router is performing route redistribution, then the routers might send routing information received from one autonomous system back into that same autonomous system.
  - **Incompatible routing information**
    - Each routing protocol uses different metrics to determine the best path, therefore path selection using the redistributed route information might not be optimal.
  - **Inconsistent convergence times**
    - Different routing protocols converge at different rates.
- Good planning should solve the majority of issues but additional configuration might be required.
  - Some issues might be solved by changing the administrative distance, manipulating the metrics, and filtering using route maps, distribute lists, and prefix lists.



# Selecting the Best Route

Routers use the following two parameters to select the best path:

- **Administrative distance:**

- Used to rate a routing protocol's believability (also called its trustworthiness).
- This criterion is the first thing a router uses to determine which routing protocol to believe if more than one protocol provides route information for the same destination.


- **Routing metric:**

- The routing metric is a value representing the path between the local router and the destination network, according to the routing protocol being used.
- The metric is used to determine the routing protocol's "best" path to the destination.



# Cisco IOS Administrative Distance

Routing Protocol	Default Administrative Distance Value
Connected interface	0
Static route out an interface	1
Static route to a next-hop address	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIPv1 and RIP v2	120
Exterior Gateway Protocol (EGP)	140
On-Demand Routing (ODR)	160
External EIGRP	170
Internal BGP	200
Unknown	255





# Routing Metric

- A boundary router must be capable of translating the metric of the received route into the receiving routing protocol.
  - Redistributed route must have a metric appropriate for the receiving protocol.
- The Cisco IOS assigns the following default metrics when a protocol is redistributed into the specified routing protocol:

Protocol That Route Is Redistributed Into ...	Default Seed Metric
RIP	0 (interpreted as infinity)
IGRP / EIGRP	0 (interpreted as infinity)
OSPF	20 for all except BGP routes (BGP routes have a default seed metric of 1)
IS-IS	0
BGP	BGP metric is set to IGP metric value





# Defining a Seed Metric

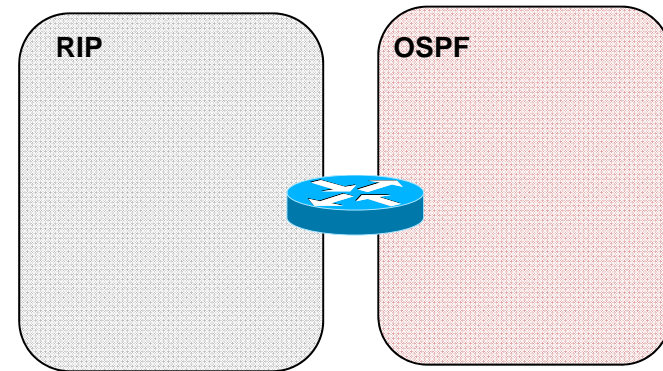
- A seed metric, different than the default metric, can be defined during the redistribution configuration.
  - After the seed metric for a redistributed route is established, the metric increments normally within the autonomous system.
    - The exception to this rule is OSPF E2 routes.
- Seed metrics can be defined in two ways:
  - The `default-metric` router configuration command establishes the seed metric for all redistributed routes.
  - The `redistribute` can also be used to define the seed metric for a specific protocol.




# Redistribution Methods

- Redistribution can be done through:
  - **One-point redistribution**
    - Only one router is redistributing one-way or two-way (both ways).
      - There could still be other boundary routers but they are not configured to redistribute.
  - **Multipoint redistribution**
    - Multiple routers are used to redistribute either one-way or two-way (both ways).
    - More prone to routing loop problems.

## One-Point Redistribution



## Multipoint Redistribution

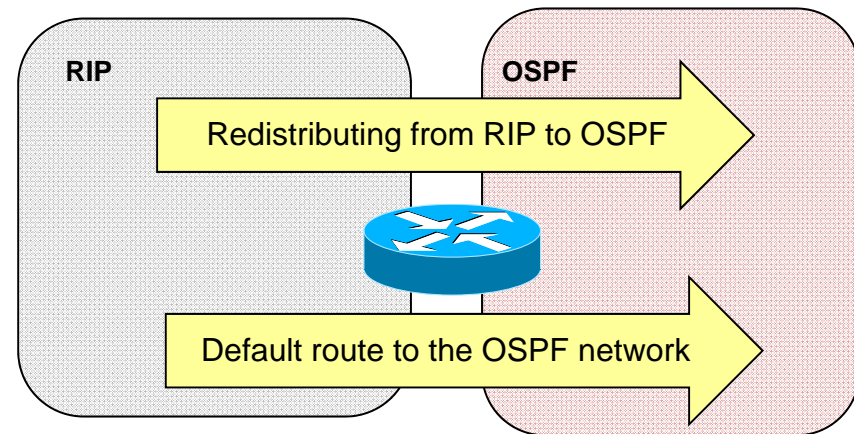




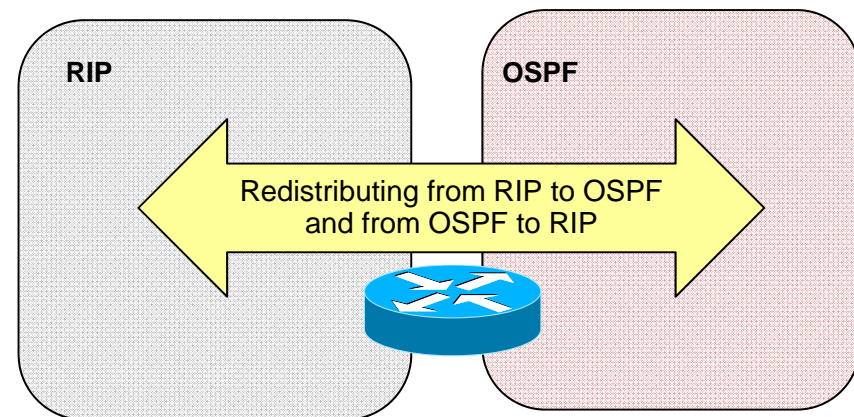
# One-Point Redistribution

- One-point redistribution can be configured in either:
  - **One-point One-way**
    - Redistributes networks from one routing protocol into the other routing protocol.
    - Typically uses a default or static route so that devices in that other part of the network can reach the first part of the network.
  - **One-point Two-way**
    - Redistributes routes between the two routing processes, in both directions.

**One-Point One-Way Redistribution**



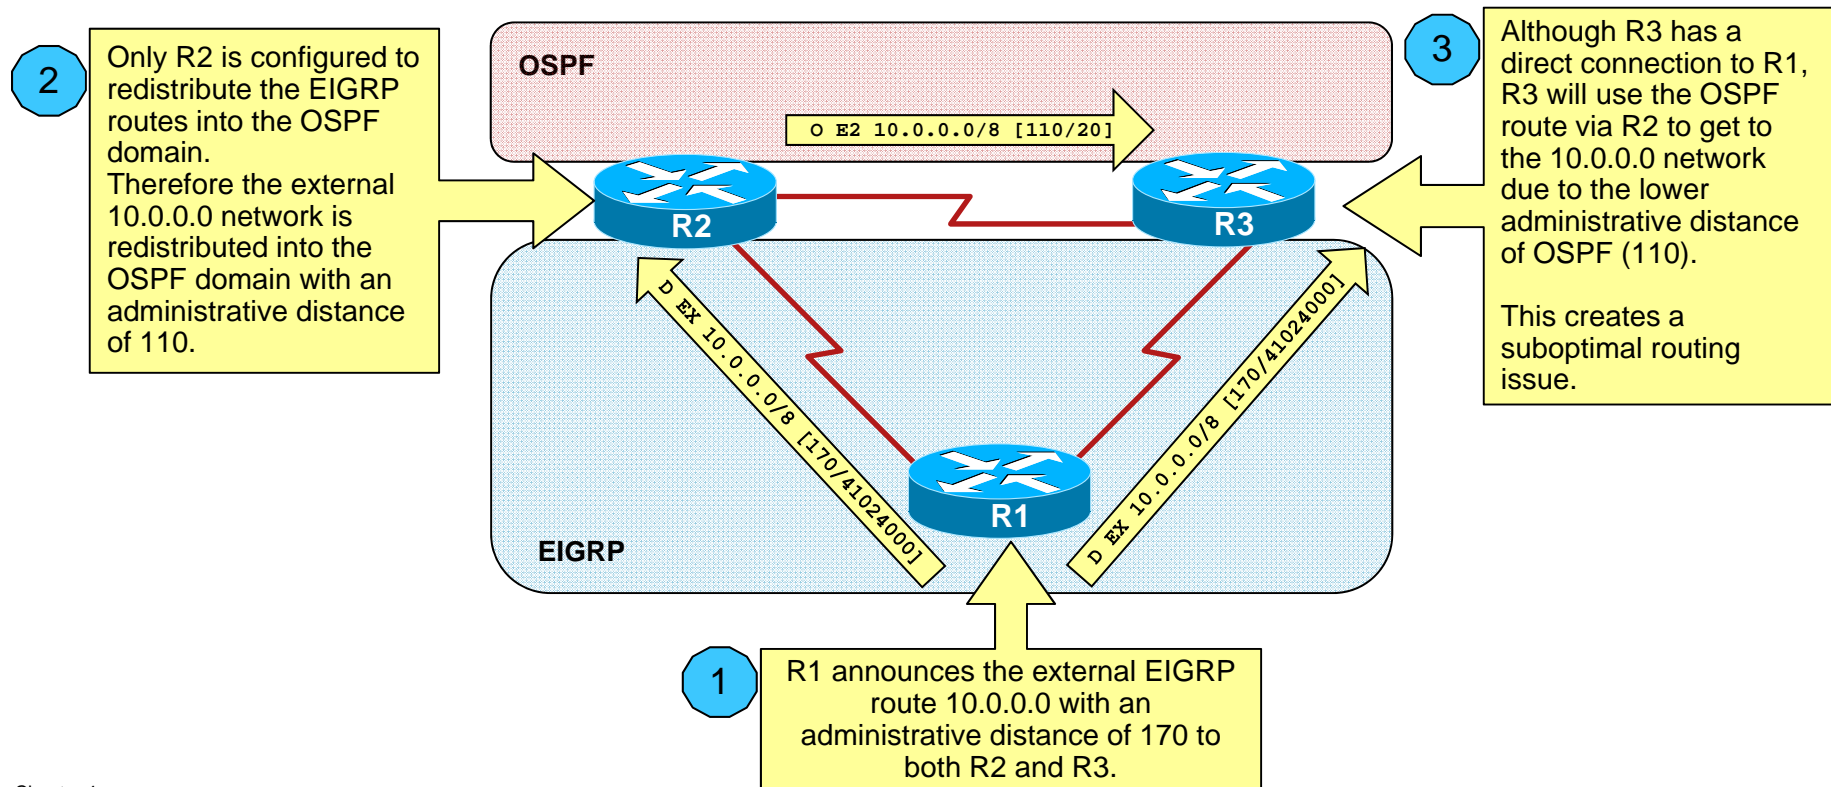
**One-Point Two-Way Redistribution**





# One-Point One-Way Redistribution Issue

- Although one-point one-way or two-way redistribution is usually safe from routing loops, issues can still occur if multiple boundary routers exist and only one router is performing one-point one-way redistribution.
  - In this example, R2 is redistributing an external EIGRP route into the OSPF domain.



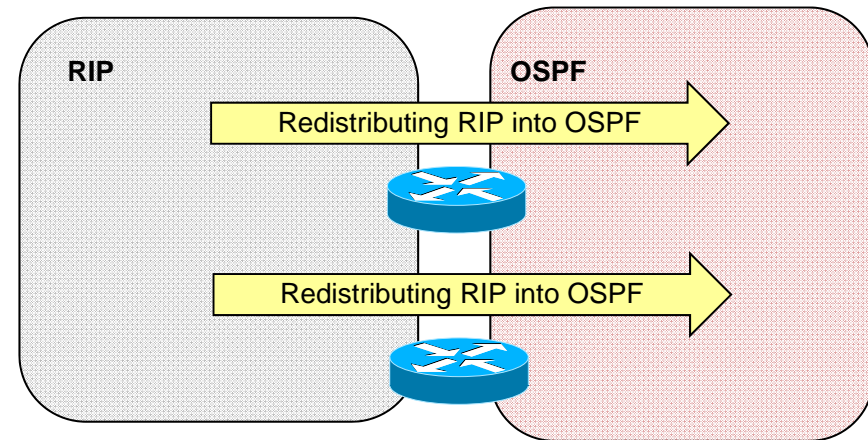




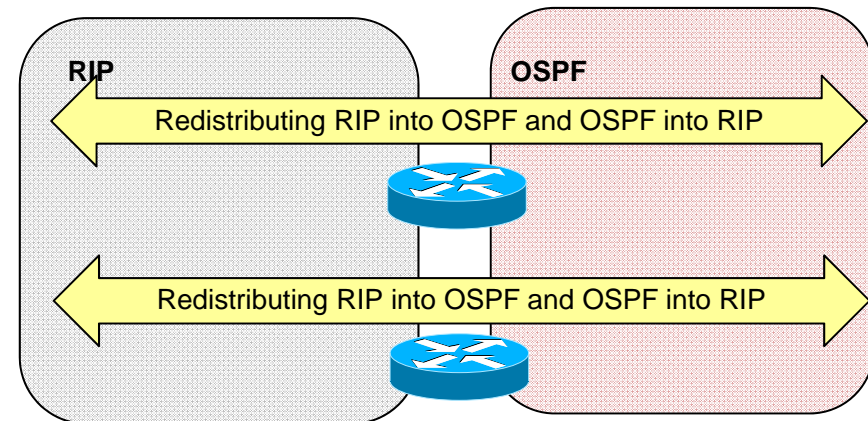
# Multipoint Redistribution

- Multipoint redistribution has two (or more) separate routers running both routing protocols.
- Redistribution can be configured as:
  - Multipoint one-way redistribution
  - Multipoint two-way redistribution
- Although multipoint two-way redistribution is especially problematic, either method is likely to introduce potential routing feedback loops.

## Multipoint One-Way Redistribution



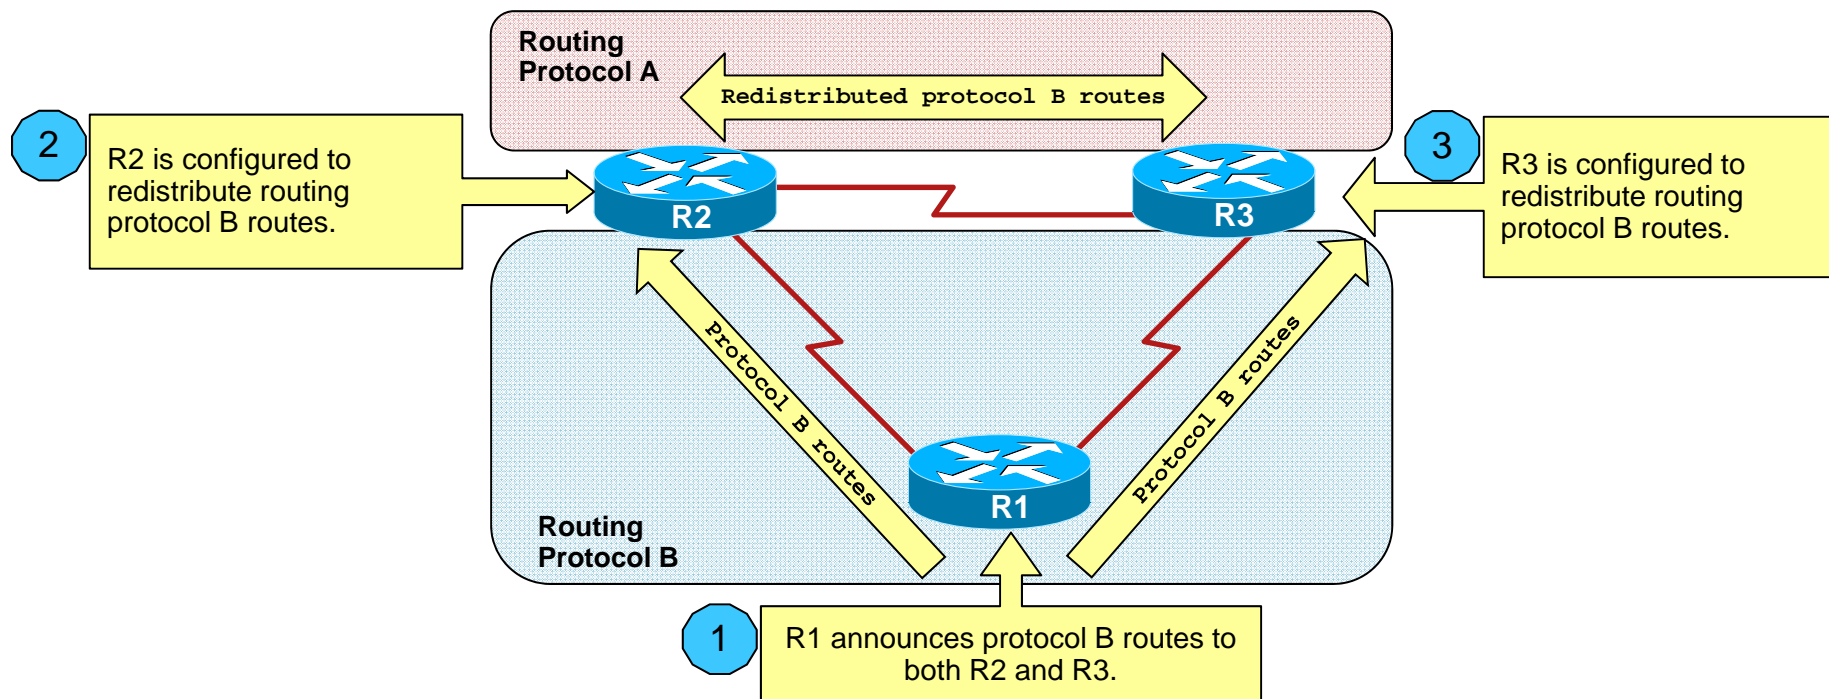
## Multipoint Two-Way Redistribution





# Multipoint Redistribution

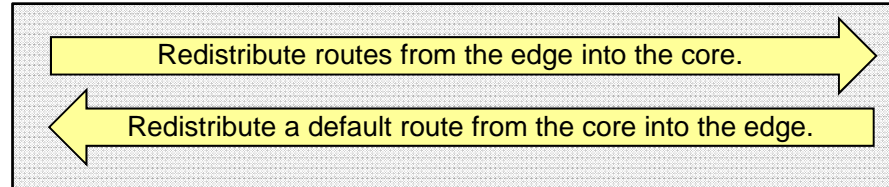
- Multipoint one-way redistribution only works well if:
  - The receiving routing protocol is either EIGRP, BGP and OSPF because they support different administrative distances for internal and external routes.
  - The administrative distance of protocol B's external routes is higher than the administrative distance of protocol A's routes, so that R2 and R3 will use the appropriate routes to destinations in the protocol A side of the network.



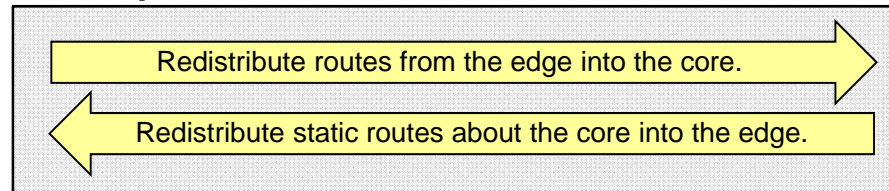


# Redistribution Techniques

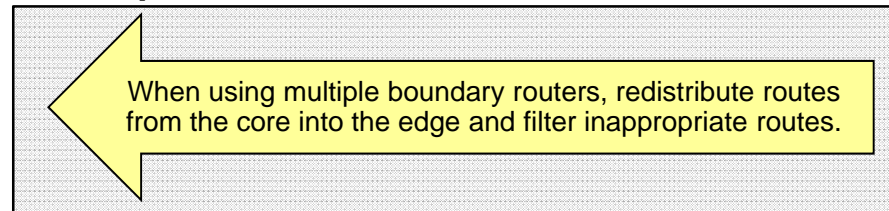
## Technique #1





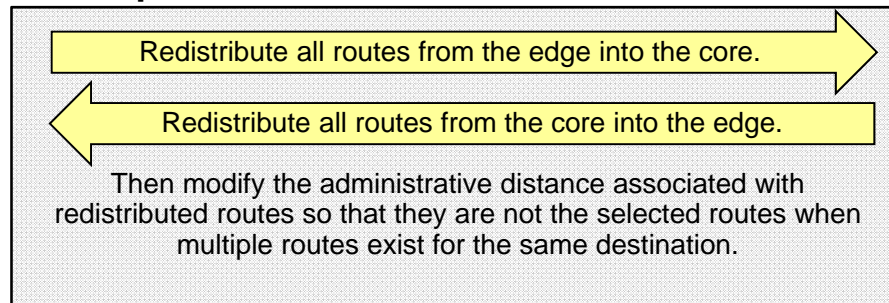
## Technique #2



## Technique #3



## Technique #4





# Preventing Routing Loops

- The safest way to perform redistribution is to redistribute routes in only one direction, on only one boundary router within the network.
  - However, this results in a single point of failure in the network.
- If redistribution must be done in both directions or on multiple boundary routers, the redistribution should be tuned to avoid problems such as suboptimal routing and routing loops.





# Redistribution Guidelines

- Do not overlap routing protocols.
  - Do not run two different protocols in the same Internetwork.
  - Instead, have distinct boundaries between networks that use different routing protocols.
- Be familiar with your network.
  - Knowing the network will result in the best decision being made.




# Generic Redistribution Steps

1. Identify the boundary router(s) that will perform redistribution.
2. Determine which routing protocol is the core protocol.
3. Determine which routing protocol is the edge protocol.
  - Determine whether all routes from the edge protocol need to be propagated into the core and consider methods that reduce the number of routes.
4. Select a method for injecting the required routes into the core.
  - Summarized routes at network boundaries minimizes the number of new entries in the routing table of the core routers.
5. Consider how to inject the core routing information into the edge protocol.



# Which Path From R1 to 10.0.0.0 /8?

- RIP, OSPF, and EIGRP are all configured on the routers.
- Which path would R1 choose if:
  - RIP made the decision? R1 ⇒ R4 ⇒ R6
  - OSPF made the decision? R1 ⇒ R2 ⇒ R3 ⇒ R5 ⇒ R6
  - EIGRP made the decision? R1 ⇒ R2 ⇒ R3 ⇒ R5 ⇒ R6
- Because EIGRP has the lowest administrative distance of the three protocols, only the EIGRP path to 10.0.0.0/8 is put into the routing table.





# Quiz Question

- Assume that a router has three routing processes running simultaneously on it, and each process has received these routes:
  - EIGRP (internal):            192.168.32.0/26
  - RIP:                            192.168.32.0/24
  - OSPF:                         192.168.32.0/19
- Which of these routes will be installed in the routing table?
  - All of them!
    - Although EIGRP has the best administrative distance, each of these routes has a different prefix length (subnet mask).
    - They are therefore considered different destinations and are all installed in the routing table.



# Modifying the Administrative Distance

- When routes are redistributed between two different routing protocols, some information may be lost making route selection more confusing.
- One approach to correct this is to control the administrative distance to indicate route selection preference and ensure that route selection is unambiguous.
  - Although, this approach does not always guarantee the best route is selected, only that route selection will be consistent.
- For all protocols use the **distance** *administrative-distance* router configuration command.
  - Alternatively for OSPF, use the **distance ospf** command.
  - Alternatively for EIGRP, use the **distance eigrp** command.



# Controlling Routing Updates

- Propagating routing information can be controlled by using:
  - Passive interface
  - Static routes
  - Default route
  - Route maps
  - Distribute lists
  - Prefix lists
- **NOTE:**
  - There is not one type of route filter that is appropriate for every situation.
  - A variety of techniques may be used to make the network run smoothly.



# Passive Interfaces

- Passive interfaces prevent routing updates from being sent and/or received for a specified protocol.
  - RIP interfaces listen but will not send updates.
  - OSPF and EIGRP interfaces do not listen for or send updates and therefore no neighbor adjacencies can be established.

Routing protocol	Suppresses outgoing routing updates	Suppresses incoming routing updates	Stops neighbor adjacency
RIP	✓		
EIGRP	✓	✓	✓
OSPF	✓	✓	✓
IS-IS	✓	✓	✓



# Static and Default Routes

- Static routes are manually configured routes that are used to:
  - Define specific routes to use when two autonomous systems must exchange routing information.
  - Define routes to destinations over a WAN link to eliminate the need for a dynamic routing protocol.
- Static route configuration considerations:
  - If you want a router to advertise a static route in a routing protocol, it might need to be redistributed.
  - To reduce the number of static route entries, define a default static route.





# Understanding Route Maps

- Route maps are similar in function to ACLs, but provide far more control.
- Route maps are more similar to a scripting language.
  - They can be named rather than numbered for easier documentation.
  - Lines are sequence-numbered for easier editing.
  - Match and set criteria can be used, similar to the “if, then” logic.
    - They allow conditions to be tested using match commands and if the conditions match, actions specified by set commands can be taken to modify attributes of the packet or routes.
- Just as ACLs are used by a variety of Cisco IOS features, route maps can also be used for various applications.
  - The actual route map implementation will vary based on how its applied.



# Route Map Applications

- **Route filtering during redistribution**
  - All IP routing protocols can use route maps for redistribution filtering.
  - Applied using the `redistribute protocol route-map` router configuration command.
- **Policy-based routing (PBR)**
  - PBR allows the operator to define routing policy other than basic destination-based routing using the routing table.
  - Applied using the `ip policy route-map` interface configuration command.
- **NAT**
  - Route maps provide more control over which private addresses are translated to public addresses.
- **BGP**
  - Route maps are the primary tools for implementing BGP policy.



# Route Map Operation Logic

- A route map consists of a list of statements.
  - The list is processed top-down like an access list.
  - Sequence numbers are used for inserting or deleting specific statements.
- Route map permit or deny determines if the candidate will be redistributed.
  - At least one reference must permit the route for it to be a candidate for redistribution.
- The first match found for a route is applied.
  - The match statement may contain multiple references.
    - Multiple match criteria in the same line use a logical OR.
    - Multiple match criteria in multiple separate lines use a logical AND.
  - Once there is a match, set the action (if defined) and leave the route map.
    - Other route-map statements are not processed.



# Configuring Route Maps for PBR

- PBR allows the operator to define a routing policy other than basic destination-based routing using the routing table.
  - For example to make packets to take a route other than the obvious shortest path.
- Sample implementation plan:
  - Define and name the route map with the `route-map` command.
    - Define the conditions to match (the `match` statements).
    - Define the action to be taken when there is a match (the `set` statements).
  - Define which interface the route map will be attached to using the `ip policy route-map` interface configuration command.
    - PBR is applied to incoming packets.



# Configuring Route Maps for PBR Example

```

R1(config)# access-list 1 permit 172.21.16.18 0.0.0.0
R1(config)#
R1(config)# route-map MY-ROUTE-MAP permit 10
R1(config-route-map)# match ip address 1
R1(config-route-map)# set ip next-hop 172.30.3.20
R1(config-route-map)#
R1(config-route-map)# interface S0/0/0
R1(config-if)# ip policy route-map MY-ROUTE-MAP
  
```

- The route map has only one **permit** statement.
  - Any packets that match the IP address specified by ACL 1 (172.21.16.18) should be sent to the next hop IP address 172.30.3.20.
- This route map applies to incoming packets on the S0/0/0 interface.



# Configuring Route Maps for Redistribution

- Use route maps when you want detailed control over how routes are redistributed between routing protocols.
- Sample implementation plan:
  - Define and name the route map with the `route-map` command.
    - Define the conditions to match (the `match` statements).
    - Define the action to be taken when there is a match (the `set` statements).
  - Specify the route map to use when redistributing.
    - Use the `redistribute protocol route-map map-tag` router configuration command.



# Configuring Route Maps for Redistribution


```

R1(config)# access-list 23 permit 10.1.0.0 0.0.255.255
R1(config)# access-list 29 permit 172.16.1.0 0.0.0.255
R1(config)# access-list 37 permit 10.0.0.0 0.255.255.255
R1(config)#
R1(config)# route-map REDIS-RIP permit 10
R1(config-route-map)# match ip address 23 29
R1(config-route-map)# set metric 500
R1(config-route-map)# set metric-type type-1
R1(config-route-map)#
R1(config-route-map)# route-map REDIS-RIP deny 20
R1(config-route-map)# match ip address 37
R1(config-route-map)#
R1(config-route-map)# route-map REDIS-RIP permit 30
R1(config-route-map)# set metric 5000
R1(config-route-map)# set metric-type type-2
R1(config-route-map)#
R1(config-route-map)# router ospf 10
R1(config-router)# redistribute rip route-map REDIS-RIP subnets
R1(config-router)#
  
```

- The route map REDIS-RIP tests the following;
  - In sequence 10, any routes matching ACLs 23 or 29 will have their metric changed accordingly.
  - In sequence 20, any routes matching ACLs 37 will not be redistributed.
  - In sequence 30, all other routes will have their metric changed accordingly.
- Finally, all RIP routes and subnets will be redistributed into OSPF according to the REDIS-RIP route map statements.



# Route Feedback



- There is a possibility that routing feedback might cause suboptimal routing when routes are redistributed by more than one router such as in the two-way multipoint redistribution configuration on R1 and R2.
- The following explains the routing feedback loop for this scenario:
  - RIPv2 on R3 advertises network 192.168.1.0.
  - R1 redistributes the 192.168.1.0 network into OSPF.
  - OSPF then propagates this route through the OSPF domain.
  - An OSPF router eventually advertises the 192.168.1.0 network to R2.
  - R2 then redistributes 192.168.1.0 from OSPF back into the original RIPv2 network creating a routing feedback loop.



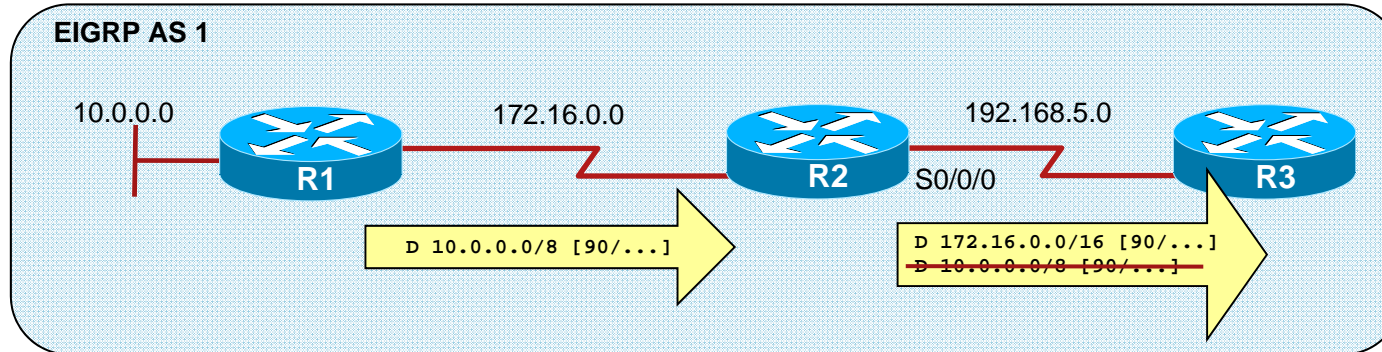


# Using Distribute Lists

- Another way to control routing updates is to use a distribute list which allows an ACL to be applied to routing updates for filtering purposes.
  - Administrators control which routes get distributed.
  - This control is for security, overhead, and management reasons.
- It's important to understanding that the distribution lists are used to control (filter) routing updates while ACLs filter user traffic.
- Sample implementation plan:
  - Identify network traffic to be filtered using an ACL or route map.
  - Associate the distribute list with the ACL or route-map using the `distribute-list` router configuration command.



# Filter Outgoing Routing Updates Example 1a

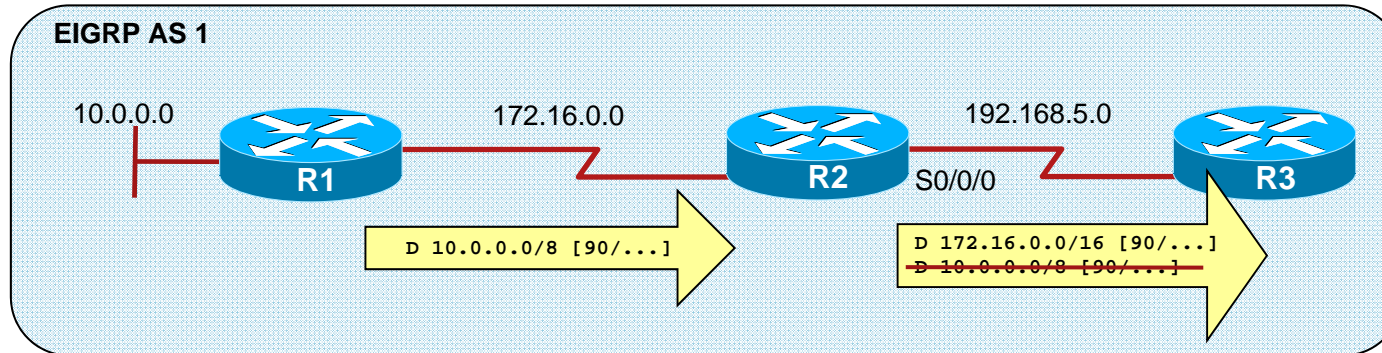


```
R2(config)# access-list 7 permit 172.16.0.0 0.0.255.255
R2(config)#
R2(config)# router eigrp 1
R2(config-router)# network 172.16.0.0
R2(config-router)# network 192.168.5.0
R2(config-router)# distribute-list 7 out Serial0/0/0
R2(config-router)#
```

- In this example, the network 10.0.0.0 must be hidden from the devices in network 192.168.5.0.
  - The **distribute-list out** command on R2 applies ACL 7 to packets going out S0/0/0 which only permits 172.16.0.0 routing information to be distributed out.
  - The implicit **deny any** at the end of the ACL prevents updates about any other networks from being advertised and as a result, network 10.0.0.0 is hidden.



# Filter Outgoing Routing Updates Example 1b



```
R2(config)# access-list 7 deny 10.0.0.0 0.255.255.255
R2(config)# access-list 7 permit any
R2(config)#
R2(config)# router eigrp 1
R2(config-router)# network 172.16.0.0
R2(config-router)# network 192.168.5.0
R2(config-router)# distribute-list 7 out Serial0/0/0
R2(config-router)#
```

- As an alternative, network 10.0.0.0 can be explicitly denied and all other routes are valid.
  - The **distribute-list out** command on R2 applies ACL 7 to packets going out S0/0/0 which denies the 10.0.0.0/8 network but permits all other routes.



# Drawback of Distribute Lists

- Using distribute lists as route filters has several drawbacks, including:
  - A subnet mask cannot be easily matched.
  - ACLs are evaluated sequentially for every IP prefix in the routing update.
  - An extended ACL can be cumbersome to configure.
  - A distribute list hides network information, which could be considered a drawback in some circumstances.
    - For example, in a network with redundant paths, a distribute list might permit routing updates for only specific paths, to avoid routing loops.
    - In this case, if the primary path goes down, the backup paths are not used because the rest of the network does not know they exist.
    - When redundant paths exist, use other techniques.



# Using Prefix Lists

- Prefix lists can be used as an alternative to access lists in many route filtering commands.
- Prefix list characteristics include:
  - A significant performance improvement over ACLs in loading and route lookup of large lists.
  - Support for incremental modifications.
  - An improved user-friendly command-line interface.
  - Greater flexibility in specifying subnet mask ranges.



# Similarities Between Prefix Lists and ACLs

- A prefix list can consist of any number of lines, each of which indicates a test and a result.
- When a router evaluates a route against the prefix list, the first line that matches results in either a permit or deny.
- If none of the lines in the list match, the result is “implicitly deny,” just as it is in an access list.




# Prefix List Filtering rules

- An empty prefix list permits all prefixes.
- If a prefix is permitted, the route is used. If a prefix is denied, the route is not used.
- Prefix lists consist of statements with sequence numbers. The router begins the search for a match at the top of the prefix list, which is the statement with the lowest sequence number.
- When a match occurs, the router does not need to go through the rest of the prefix list. For efficiency, you might want to put the most common matches (permits or denies) near the top of the list by specifying a lower sequence number.
- An implicit deny is assumed if a given prefix does not match any entries in a prefix list.





# Prefix-list Scenario #1



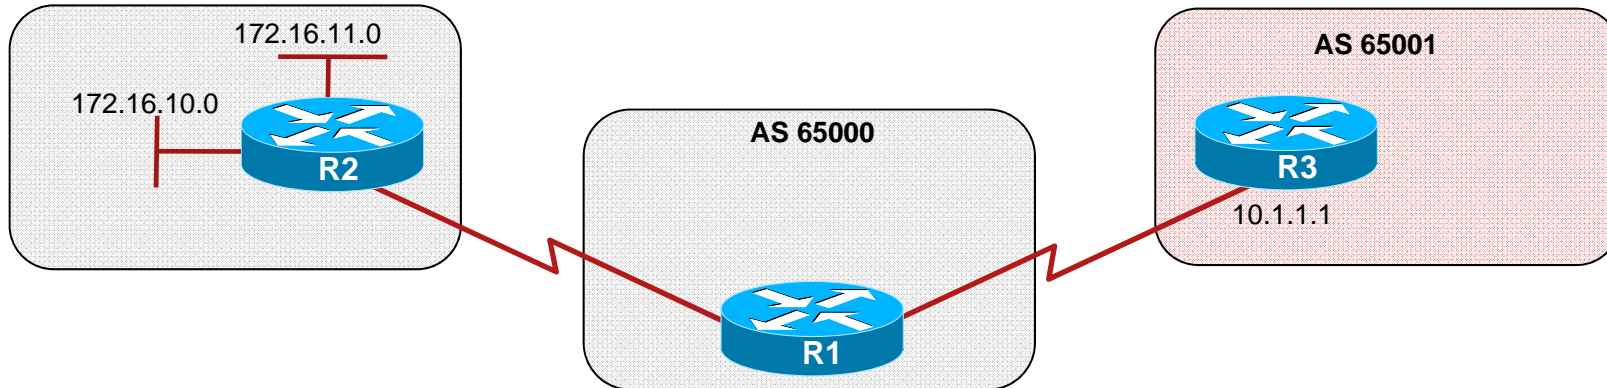
```
R1(config)# ip prefix-list TEN-ONLY permit 172.16.10.0/8 le 24
R1(config)# router bgp 65000
R1(config-router)# aggregate-address 172.16.0.0 255.255.0.0
R1(config-router)# neighbor 10.1.1.1 remote-as 65001
R1(config-router)# neighbor 10.1.1.1 prefix-list TEN-ONLY out
R1(config-router)# exit
R1(config)# do show running-config | include ip prefix-list
ip prefix-list TEN-ONLY seq 5 permit 172.0.0.0/8 le 24
R1(config)#
```

- Notice that the last line of this configuration changed to **ip prefix-list TEN-ONLY permit 172.0.0.0/8 le 24**
  - This is because only the first 8 bits in the address are considered significant when a prefix length of /8 is used.
- In this case, neighbor R3 learns about 172.16.0.0/16, 172.16.10.0/24, and 172.16.11.0/24.
  - These are the routes that match the first 8 bits of 172.0.0.0 and have a prefix length between 8 and 24.





# Prefix-list Scenario #2

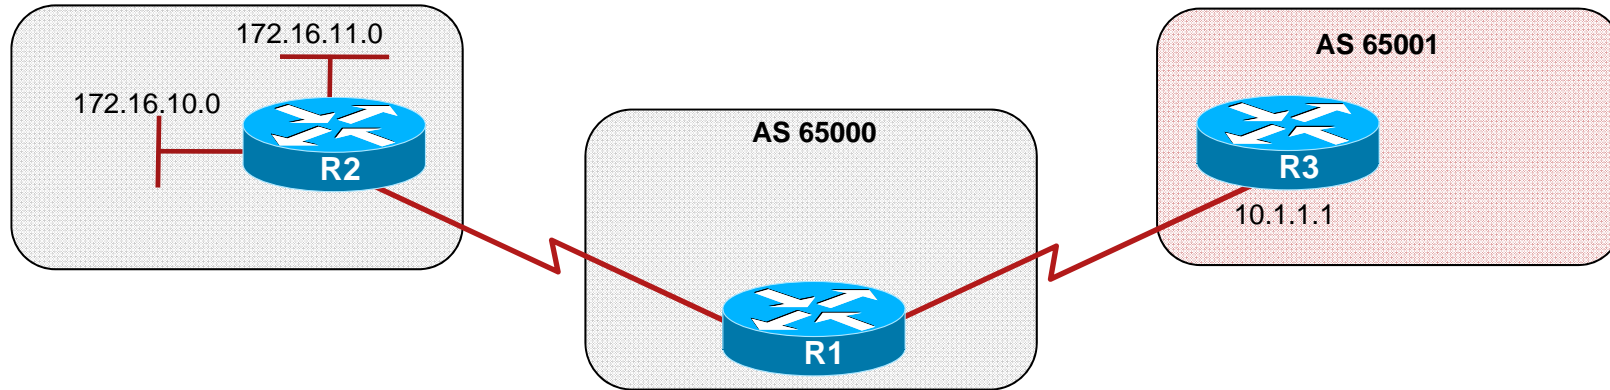


```
R1(config)# ip prefix-list TEN-ONLY permit 172.16.10.0/8 le 16
R1(config)# router bgp 65000
R1(config-router)# aggregate-address 172.16.0.0 255.255.0.0
R1(config-router)# neighbor 10.1.1.1 remote-as 65001
R1(config-router)# neighbor 10.1.1.1 prefix-list TEN-ONLY out
R1(config-router)# exit
R1(config)#
```

- Now neighbor R3 learns only about 172.16.0.0/16.
  - This is the only route that matches the first 8 bits of 172.0.0.0 and has a prefix length between 8 and 16.



# Prefix-list Scenario #3

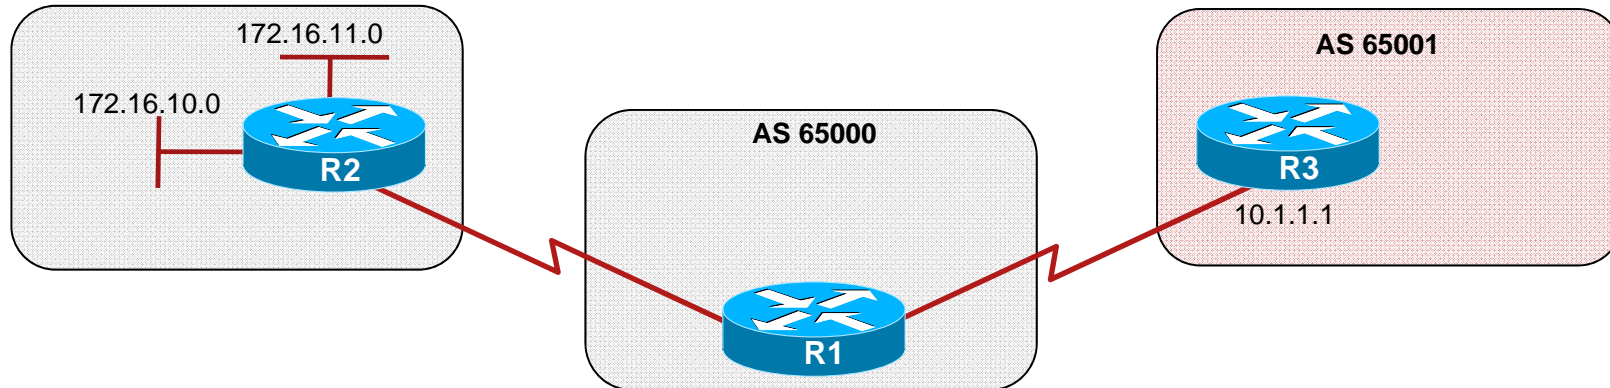


```
R1(config)# ip prefix-list TEN-ONLY permit 172.16.10.0/8 ge 17
R1(config)# router bgp 65000
R1(config-router)# aggregate-address 172.16.0.0 255.255.0.0
R1(config-router)# neighbor 10.1.1.1 remote-as 65001
R1(config-router)# neighbor 10.1.1.1 prefix-list TEN-ONLY out
R1(config-router)# exit
R1(config)#
```

- Now neighbor R3 learns only about 172.16.10.0/24 and 172.16.11.0/24.
  - R1 ignores the /8 parameter and treats the command as if it had the parameters `ge 17 le 32`.



# Prefix-list Scenario #4

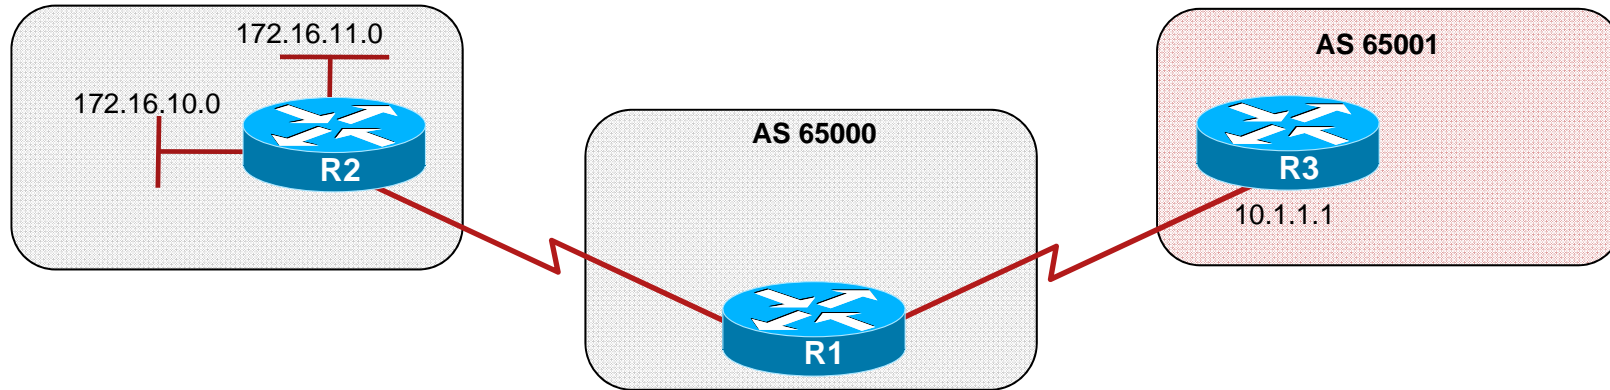


```
R1(config)# ip prefix-list TEN-ONLY permit 172.16.10.0/8 ge 16 le 24
R1(config)# router bgp 65000
R1(config-router)# aggregate-address 172.16.0.0 255.255.0.0
R1(config-router)# neighbor 10.1.1.1 remote-as 65001
R1(config-router)# neighbor 10.1.1.1 prefix-list TEN-ONLY out
R1(config-router)# exit
R1(config)#
```

- Now neighbor 10.1.1.1 learns about 172.16.0.0/16, 172.16.10.0/24, and 172.16.11.0/24.
  - R1 ignores the /8 parameter and treats the command as if it had the parameters `ge 16 le 24`.



# Prefix-list Scenario #5

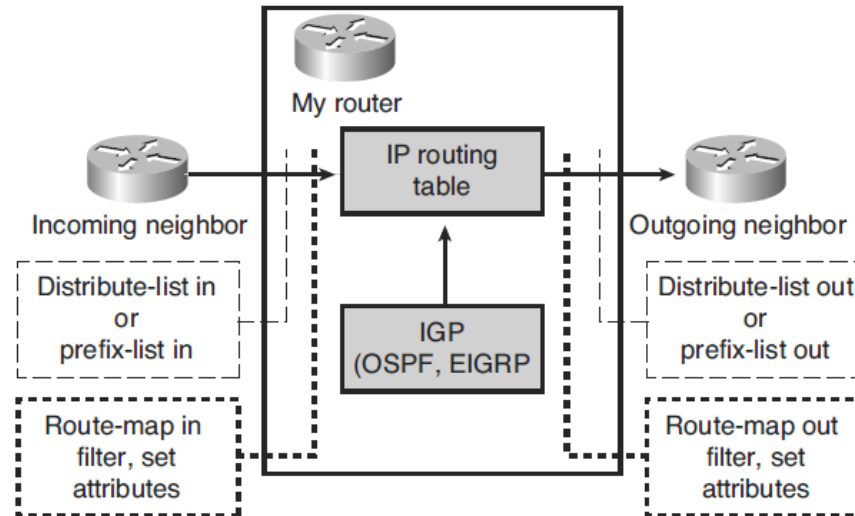


```
R1(config)# ip prefix-list TEN-ONLY permit 172.16.10.0/8 ge 17 le 24
R1(config)# router bgp 65000
R1(config-router)# aggregate-address 172.16.0.0 255.255.0.0
R1(config-router)# neighbor 10.1.1.1 remote-as 65001
R1(config-router)# neighbor 10.1.1.1 prefix-list TEN-ONLY out
R1(config-router)# exit
R1(config)#
```

- Now neighbor 10.1.1.1 learns about 172.16.10.0/24 and 172.16.11.0/24.
  - R1 ignores the /8 parameter and treats the command as if it had the parameters `ge 17 le 24`.



# Multiple Methods to Control Routing Updates



- The example displays how a combination of prefix lists, distribute lists, and route maps can be applied to incoming or outgoing information.
  - All must permit the routes that are received from a neighbor before they will be accepted into the IP routing table.
  - Outgoing routes must pass the outgoing distribute list, the outgoing prefix list, and the outgoing route map before being forwarded to the neighbor.

Cisco | Networking Academy<sup>®</sup>  
Mind Wide Open<sup>™</sup>