Module 5 Manipulating Routing Updates

Lesson 1 Operating a Network Using Multiple IP Routing Protocols

Redistributing Routing Information

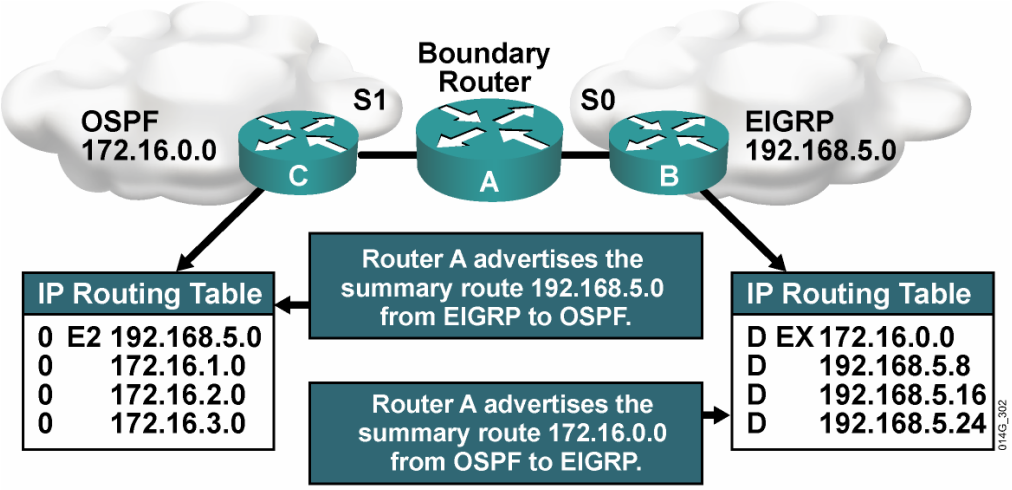

Routes are learned from another routing protocol when a router redistributes the information between the protocols.

Using Seed Metrics

- Use the **`default-metric`** command to establish the seed metric for the route or specify the metric when redistributing.
- Once a compatible metric is established, the metric will increase in increments just like any other route.

  - Exception – OSPF E2, which hold their initial metric regardless of how far they are propagated across an AS
- The initial seed metric should be set to a value larger than the largest metric within the receiving AS to help prevent suboptimal routing and routing loop.

Redistribution with Seed Metric

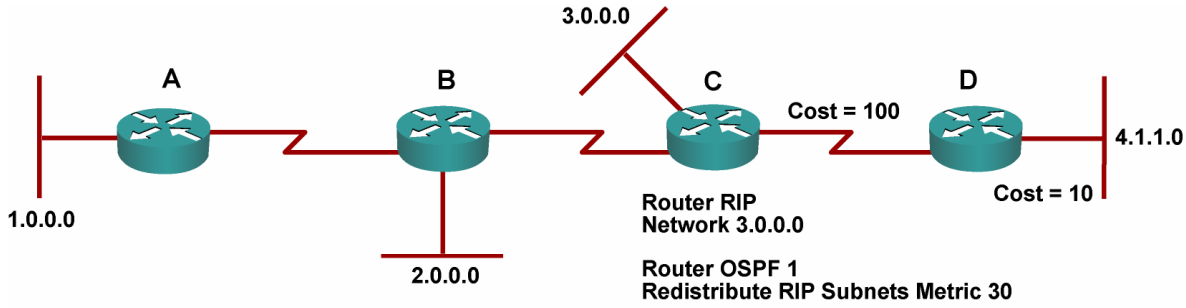

Table A

```
C 1.0.0.0
R 120/1 2.0.0.0
R 120/2 3.0.0.0
```

Table B

```
C 2.0.0.0
R 120/1 1.0.0.0
R 120/1 3.0.0.0
```

Table C

```
C 3.0.0.0
R 120/1   2.0.0.0
R 120/2   1.0.0.0
O 110/110 4.1.1.0
```

Table D

```
C  4.1.1.0/24
E2 110/130 1.0.0.0
E2 110/130 2.0.0.0
E2 110/130 3.0.0.0
```

314P_155

- Default seem metric:
  - Rip – infinity
  - IGRP or EIGRP – infinity
  - OSPF – 20 for all except BGP, which is 1
  - IS-IS – 0
  - BGP – BGP metric is set to IGP metric value

Lesson 2 Confirguring and Verifying Routing Redistribution

Redistribution into RIP

- **R(config-router)#redistributie** *protocol* **[***process-id***] [match** *route-type***] [metric** *metric-value***] [route-map** *map-tag***]**
- Default metric is infinity

Parameter	Description
protocol	Source protocol from which routes are being redistributed. It can be one of the following keywords: **connected**, **bgp**, **eigrp**, **egp**, **igrp**, **isis**, **iso-igrp**, **mobile**, **odr**, **ospf**, **static**, or **rip**.
process-id	This value is an AS number, used for Border Gateway Protocol (BGP), Exterior Gateway Protocol (EGP), EIGRP, or IGRP. For OSPF, this value is an OSPF process ID.
match *route-type*	(Optional) Command parameter used for redistributing OSPF routes into another routing protocol. For OSPF, the criterion by which OSPF routes are redistributed into other routing domains. It can be any of the following: ■ **internal:** Redistributes routes that are internal to a specific AS. ■ **external 1:** Redistributes routes that are external to the AS, but are imported into OSPF as a type 1 external route. ■ **external 2:** Redistributes routes that are external to the AS, but are imported into OSPF as a type 2 external route.
metric *metric-value*	(Optional) Parameter used to specify the RIP seed metric for the redistributed route. When you are redistributing into RIP, this value is not specified and no value is specified using the **default-metric** router configuration command, then the default metric is 0, which is interpreted as infinity, and routes will not be redistributed. The metric for RIP is the hop count.
route-map *map-tag*	(Optional) Identifier of a configured route map to be interrogated to filter the importation of routes from this source routing protocol to the current routing protocol.

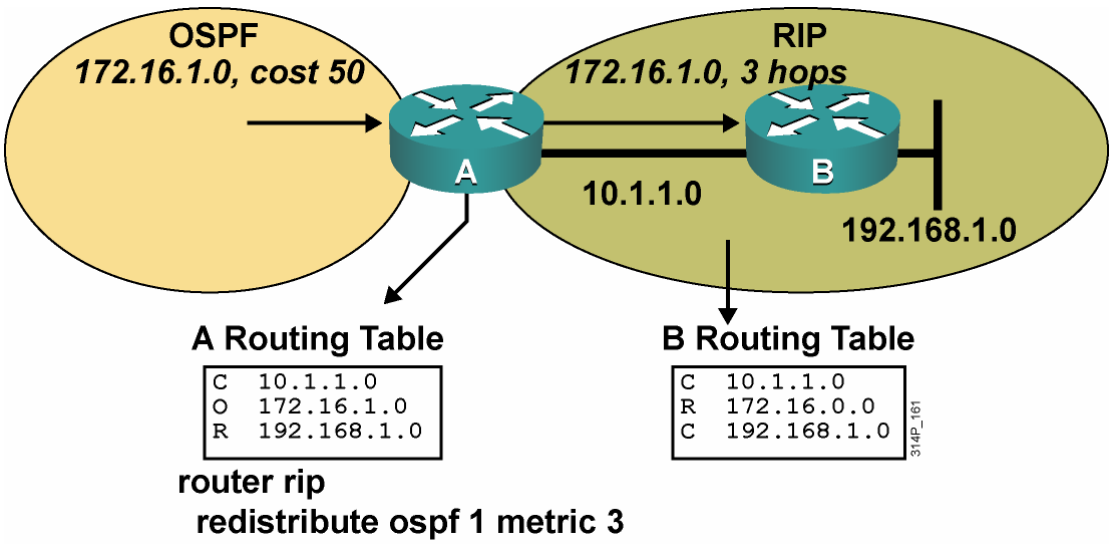

Redistribution into OSPF

- **R(config-router)#redistributie** *protocol* **[***process-id***] [metric** *metric-value***] [metric-type** *type-value***] [route-map** *map-tag***] [subnets] [tag** *tag-value***]**
- Default metric is 20
- Default metric type is 2
- Subnets do not redistribute by default

Parameter	Description
protocol	Source protocol from which routes are being redistributed. It can be one of the following keywords: **connected**, **bgp**, **eigrp**, **egp**, **igrp**, **isis**, **iso-igrp**, **mobile**, **odr**, **ospf**, **static**, or **rip**.
process-id	This value is an AS number, used for BGP, EGP, EIGRP, or IGRP. For OSPF, this value is an OSPF process ID.
metric *metric-value*	(Optional) Parameter that specifies the OSPF seed metric that is used for the redistributed route. When you are redistributing into OSPF, the default metric is 20 (except for BGP, which is 1). Use a value consistent with the destination protocol, in this case, the OSPF cost.
metric-type *type-value*	(Optional) OSPF parameter that specifies the external link type that is associated with the external route that is advertised into the OSPF routing domain. This value can be 1 for type 1 external routes or 2 for type 2 external routes. The default is 2.
route-map *map-tag*	(Optional) Identifier of a configured route map to be interrogated to filter the importation of routes from this source routing protocol to the current routing protocol.
subnets	(Optional) OSPF parameter that specifies that subnetted routes should be redistributed also. Only routes that are not subnetted are redistributed if the **subnets** keyword is not specified.
tag *tag-value*	(Optional) 32-bit decimal value that is attached to each external route. The OSPF protocol does not use this parameter. It may be used to communicate information between AS boundary routers (ASBRs).

- Redistribution into OSPF can also be limited to a defined number of prefixes by the **redistribute maximum-prefix** *maximum* **[***threshold***] [warnning-only]** route configuration command

  - *Threshold* – default to logging a warning at 75 percent of the defined maximum value configured.
  - **warnning-only** – no limitation is placed on redistribution.

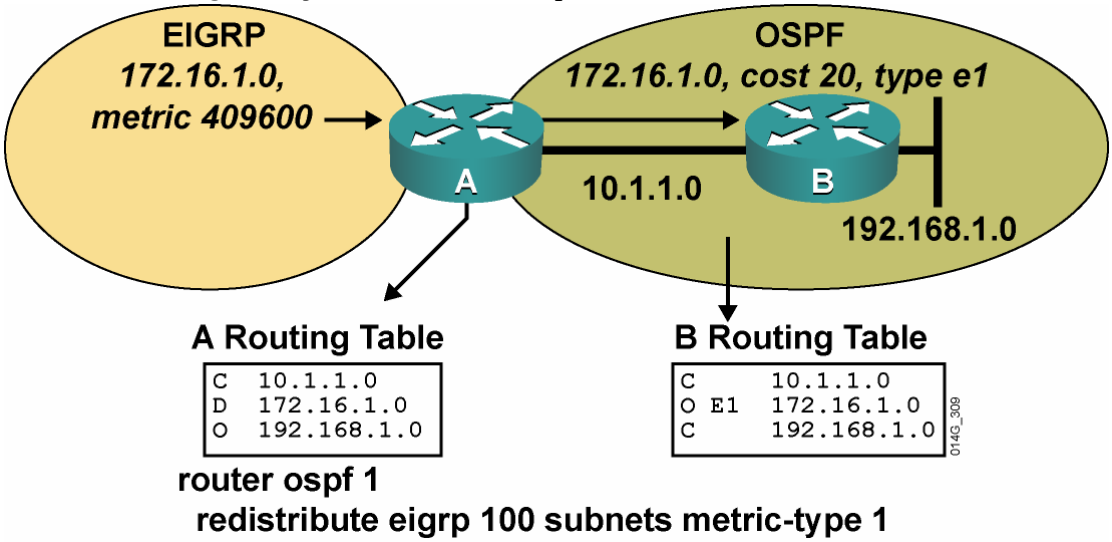

router ospf 1
redistribute eigrp 100 subnets metric-type 1

Redistribution into EIGRP

- **`R(config-router)#redistributie` *`protocol`* `[`*`process-id`*`] [match {internal | external 1 | external 2}] [metric` *`metric-value`*`] [route-map` *`map-tag`*`]`**
- Default metric is infinity
- When erdistributing a static or connected route into EIGRP, the default metric is equal to the metric of the associated interface

Parameter	Description
protocol	Source protocol from which routes are being redistributed. It can be one of the following keywords: **connected**, **bgp**, **eigrp**, **egp**, **igrp**, **isis**, **iso-igrp**, **mobile**, **odr**, **ospf**, **static**, or **rip**.
process-id	This value is an AS number, used for BGP, EGP, EIGRP, or IGRP. For OSPF, this value is an OSPF process ID.
match *route-type*	(Optional) For OSPF, the criterion by which OSPF routes are redistributed into other routing domains. It can be one of the following: ■ **internal:** Redistributes routes that are internal to a specific AS. ■ **external 1:** Redistributes routes that are external to the AS but are imported into OSPF as a type 1 external route. ■ **external 2:** Redistributes routes that are external to the AS but are imported into OSPF as a type 1 external route.
metric *metric-value*	(Optional) Parameter that specifies the EIGRP seed metric, in the order of bandwidth, delay, reliability, load, and maximum transmission unit (MTU), for the redistributed route. When you are redistributing into protocols other than OSPF (including EIGRP), if this value is not specified and no value is specified using the **default-metric** router configuration command, the default metric is 0, zero is interpreted as infinity, and routes are not redistributed. Use a value consistent with the destination protocol. The metric for EIGRP is calculated based only on bandwidth and delay by default.
route-map *map-tag*	(Optional) Identifier of a configured route map that is interrogated to filter the importation of routes from this source routing protocol to the current routing protocol.

- Bandwidth in kilobytes = 10000
- Delay in tens of microseconds = 100
- Reliability = 255 (maximum)
- Load = 1 (minimum)
- MTU = 1500 bytes
- External EIGRP routes hava a higher AD than internal EIGRP (D) routes, so internal EIGRP routes are preferred over external EIGRP routes.

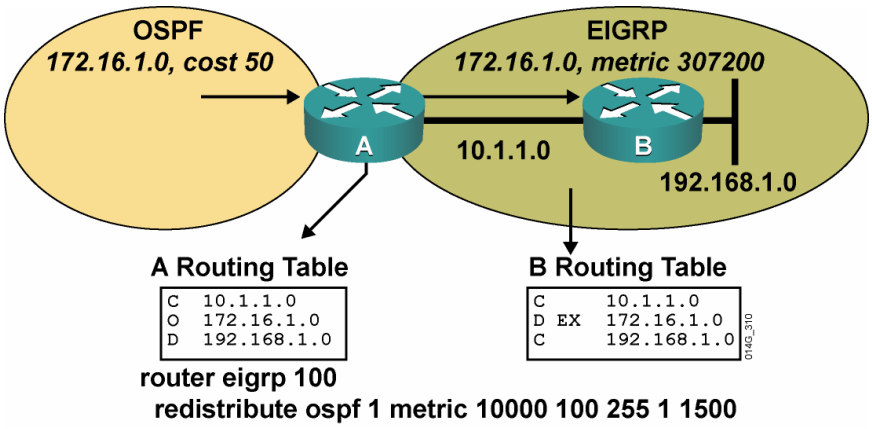

Redistribution into IS-IS

- **R(config-router)#redistributie** *protocol* **[***process-id***] [level** *level-value***] [metric** *metric-value***] [metric-type** *type-value***][route-map** *map-tag***]**
- Routes are introduced as Level 2 with a metric of 0 by default

Parameter	Description
protocol	Specifies the source protocol from which routes are being redistributed. It can be one of the following keywords: **connected**, **bgp**, **eigrp**, **egp**, **igrp**, **isis**, **iso-igrp**, **mobile**, **odr**, **ospf**, **static**, or **rip**.
process-id	Specifies an AS number, used for BGP, EGP, EIGRP, or IGRP. For OSPF, this value is an OSPF process ID.
level *level-value*	Redistributes external routes as Level 1 (*level-1*), Level 1 and Level 2 (*level-1-2*), or Level 2 (*level-2*) routes. The default is Level 2.
metric *metric-value*	Specifies the IS-IS seed metric that is used for the redistributed route. IS-IS uses a default metric of 0. Unlike RIP, IGRP, and EIGRP, a default metric of 0 is not treated as unreachable and is redistributed. The metric is increased in increments as the route is propagated into the IS-IS domain. Use a value consistent with the destination protocol, in this case, the IS-IS cost.
metric-type *type-value*	Specifies the IS-IS metric type as external or internal. The default is internal.
route-map *map-tag*	(Optional) Specifies the identifier of a configured route map to be interrogated to filter the importation of routes from this source routing protocol to the current routing protocol.

- Redistribution into IS-IS can also be limited to a defined number of prefixes by the **redistribute maximum-prefix** *maximum* **[***threshold***] [warnning-only | withdraw]** route configuration command

  - *Threshold* – default to logging a warning at 75 percent of the defined maximum value configured.
  - **warnning-only** – no limitation is placed on redistribution.
  - **withdraw** – cause IS-IS to rebuild link-state protocol data units (PDUs) (link-state packets (LSPs)) without the external (redistributed) IP prefixes.

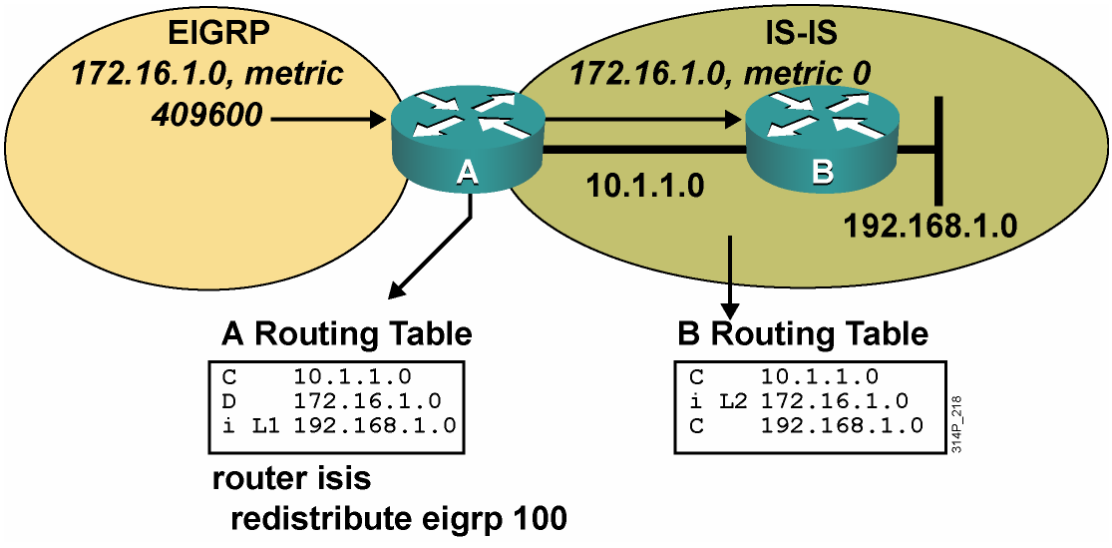

Example: Before Redistribution

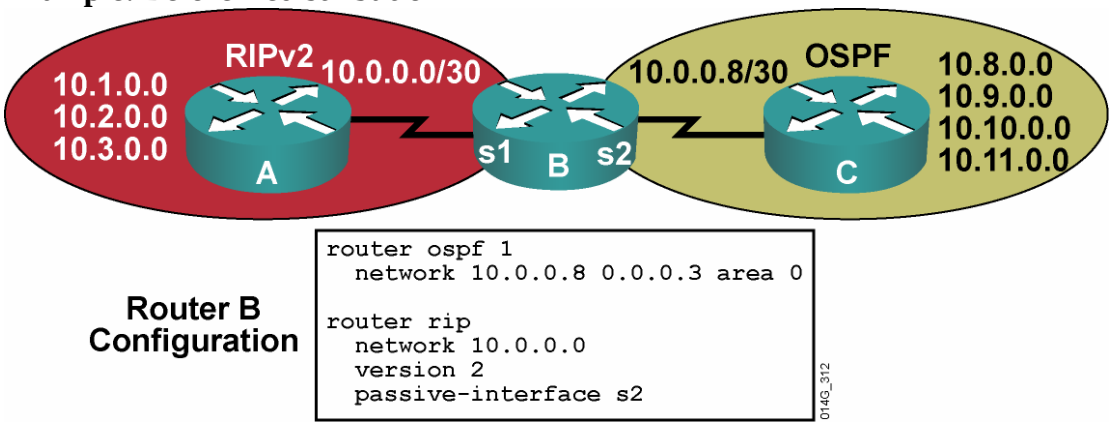

- RIP is required to run on the serial 1 interface only, so **`passive-interface`** command is given for interface serial 2.

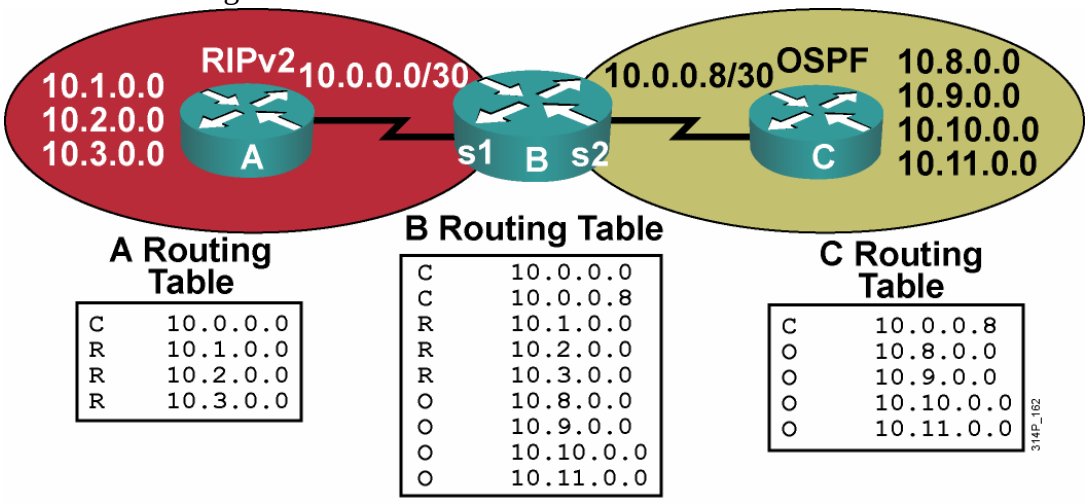

Example: Configuring Redistribution at Router B

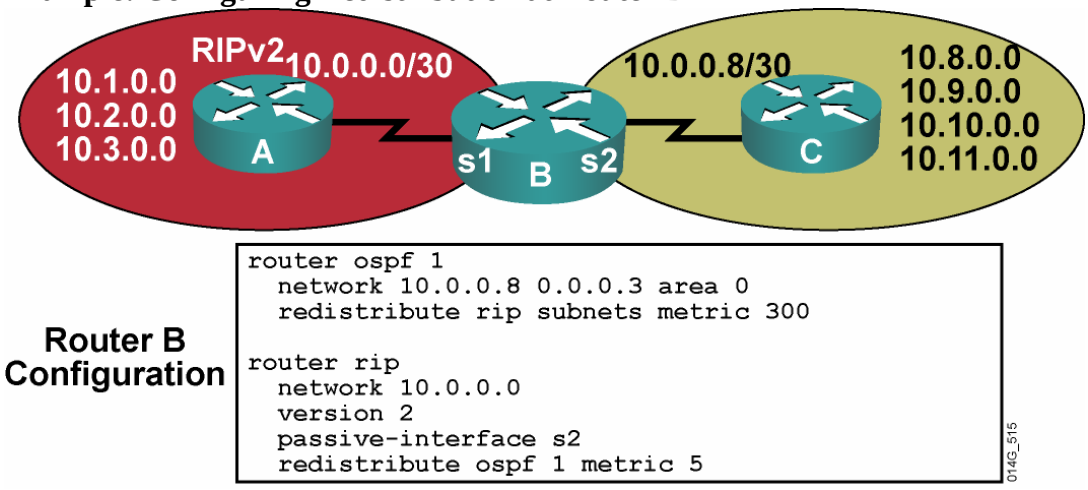

- Under the OSPF process, a value of 300 is selected because it is a worse metric than any of the native OSPF routes
- Under the RIP process, a value of 5 is chosen because it is higher than any metric in the RIP network.

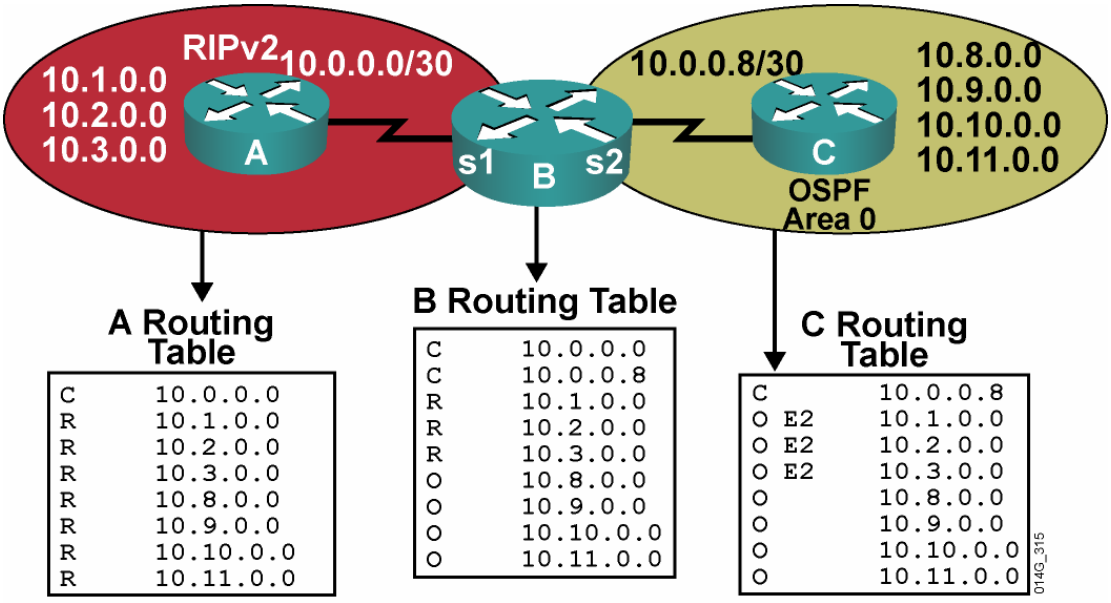

- If router in the other AS are required to track topology changes within the network, then route summarization should not be performed, because it hides information that the routers need.
- A more typical case is that the router need to recognize topology changes only within their own routing domains. In this case, performing route summarization is appropriate.

Example: Routing Tables After Summarizing Routes and Redistribution

```
RouterA(config)#interface s0
RouterA(config-if)#ip summary-address rip 10.0.0.0 255.252.0.0
```

```
RouterC(config)#router ospf 1
RouterC(config-router)#area 1 range 10.8.0.0 255.252.0.0
```

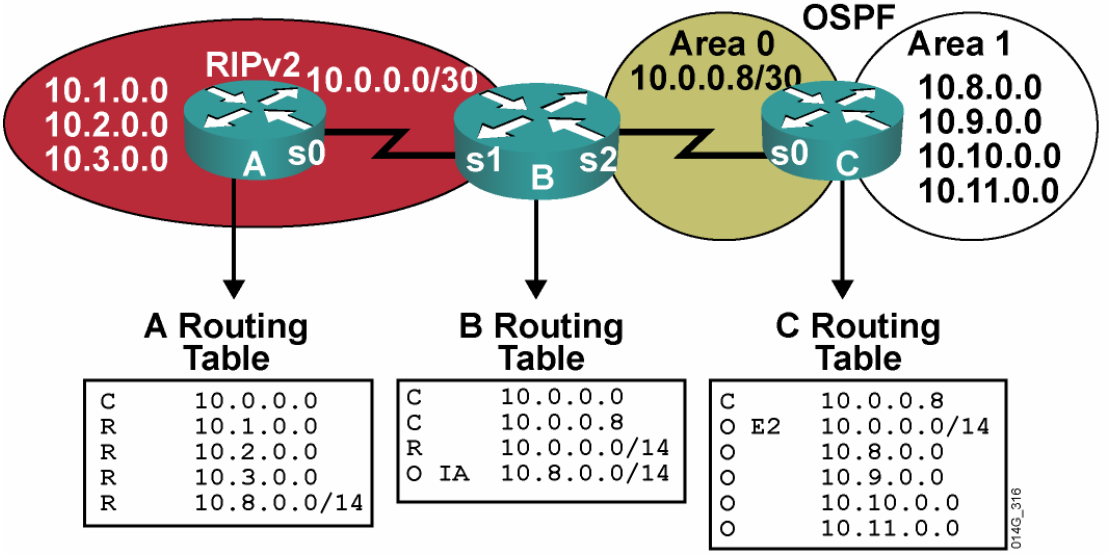

Lesson 3 Controlling Routing Update Traffic

Using the `passive-interface` Command

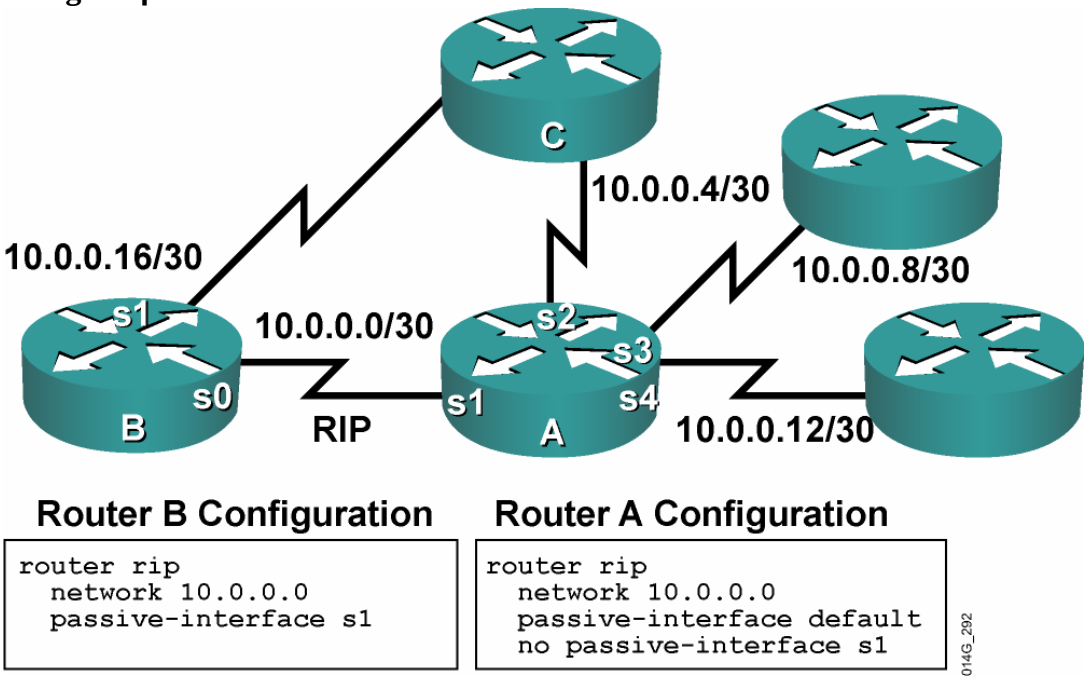

Controlling Routing Update Traffic

- Passive interface: prevent all routing updates from being sent through an interface. For EIGRP, OSPF, IS-IS, this method includes Hello protocol packets.
- Default routes: if no route for a given destination. Send the packet to the default route. Therefore, no dynamic routing updates about the remote destinations are necessary.
- Static routes: routes to remote destinations to be manually configured in the router. Therefor, no dynamic routing updates about the remote destinations are necessary.
- Distribute list: ACL. Options in the **`distribute-list`** command:
  - Incoming interface
  - outgoing interface
  - redistribution from another routing protocol

Configuring `distribute-list`

- For outbound updates:

`R(config-router)#distribute-list {`*`access-list-number | name`***`} out [`***`interface-name | routing-process [routing-process parameter]`***`]`**

- For inbound updates:

`R(config-router)#distribute-list [`*`access-list-number | name`***`] | [route-map `***`map-tag`***`] in [`***`interface-type interface-number`***`]]`**

- Use an access list (for outbound) or route map (for inbound) to permit or deny routes.

- Can be applied to transmitted, received, or redistributed routing updates.
- Configure a distribute list using an ACL:

  - Identify the network address that you want to filter and creat an ACL

  - Determine whether you want to filter traffic on an incoming interface, an outgoing interface, or routes being redistributed from another routing source.

  - use the **`distribute-list out`** to assign the ACL to filter outgoing routing updates or to assign it to route being redistributed into the protocol. **`distribute-list out`** canot be used with link-state routing protocols for blocking outbound link-state advertisements (LSAs) on an interface.

  - use the **`distribute-list in`** to assign the ACL to filter incoming routing updates coming in through an interface. This command prevents most routing protocols from placing the filtered routes in their database. When this command is used with OSPF, the routes are placed in the database but not the routing table.

- **`distribute-list out`**

Parameter	Description
access-list-number | *name*	Specifies standard ACL number or name
out	Applies the ACL to outgoing routing updates
interface-name	(Optional) Specifies name of interface out of which updates are filtered
routing-process	(Optional) Specifies name of the routing process, or the keyword **static** or **connected**, that is being redistributed and from which updates are filtered
routing-process parameter	(Optional) Specifies routing process parameter, such as the autonomous system (AS) number of the routing process

- **`distribute-list in`**

Parameter	Description
access-list-number | *name*	Specifies standard ACL number or name.
map-tag	(Optional) Name of the route map that defines which networks are to be installed in the routing table and which are to be filtered from the routing table. This argument is supported by OSPF only.
in	Applies the ACL to incoming routing updates.
interface-type interface-number	(Optional) Specifies interface type and number from which updates are filtered.

Filtering Routing Updates with a Distribute List

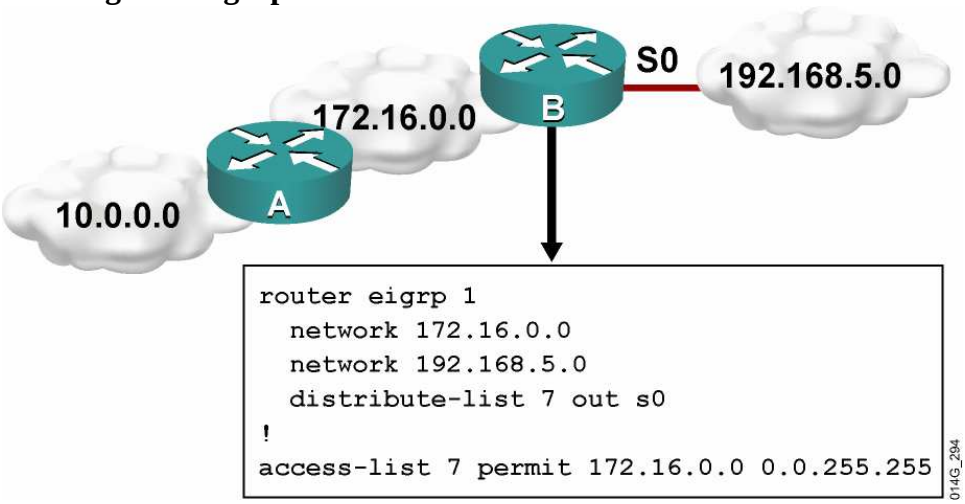

- Hides network 10.0.0.0 using interface filtering

Controlling Redistribution with Distribute Lists

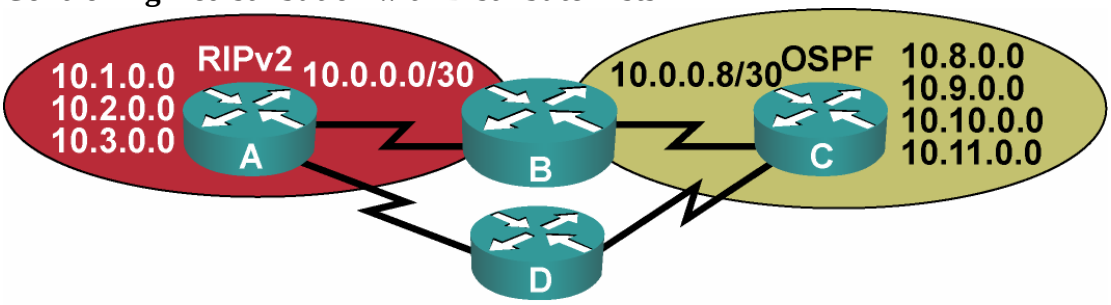

Router B Configuration

```
router ospf 1
  network 10.0.0.8 0.0.0.3 area 0
  redistribute rip subnets
  distribute-list 2 out rip

router rip
  network 10.0.0.0
  version 2
  passive-interface s3
  redistribute ospf 1 metric 5
  distribute-list 3 out ospf 1

access-list 2 permit 10.0.0.0 0.3.255.255

access-list 3 permit 10.8.0.0 0.3.255.255
```

314P_163

Route Maps

- They work like a more sophisticated access list.
  - They offer top-down processing.
  - Once there is a match, leave the route map.
- Lines are sequence-numbered for easier editing.
  - Insertion of lines
  - Deletion of lines
- Route maps are named rather than numbered for easier documentation.
- Match criteria and set criteria can be used, similar to the “if, then” logic in a

scripting language.

Route Map Applications

- Redistribution route filtering: a more sophisticated alternative to distribute lists. (use **`set`** command)
- Policy-based routing (PBR): the ability to determine routing policy based on criteria other than the destination network
- BGP policy implementation: the primary tool for defining BGP routing policies

Route Map Operation

- A list of statements constitutes a route map.
- The list is processed top-down like an access list.
- The first match found for a route is applied.
- The sequence number is used for inserting or deleting specific route map statements.
- An implicit deny at the end of a route map

```
route-map my_bgp permit 10
     { match statements }
     { match statements }
     { set statements }
     { set statements }
route-map my_bgp deny 20
     ::      ::   ::
     ::      ::   ::
route-map my_bgp permit 30
     ::      ::   ::
     ::      ::   ::
```

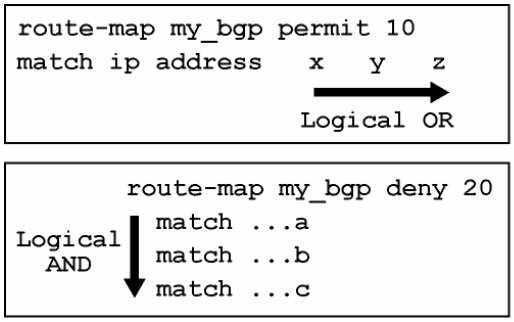

- The match statement may contain multiple references.
- Multiple match criteria in the same line use a logical OR.
- At least one reference must permit the route for it to be a candidate for redistribution.
- Each vertical match uses a logical AND.
- All match statements must permit the route for it to remain a candidate for redistribution.
- Route map permit or deny determines if the candidate will be redistributed.

`route-map` commands

- **`R(config)#route-map`** *`map-tag`* **`[permit | deny]`** **`[`***`sequence-number`***`]`**
  - defines the route map conditions
- **`R(config-route-map)#match {`***`conditions`***`}`**
  - defines the conditions to match
- **`R(config-route-map)#set {`***`actions`***`}`**
  - defines the action to be taken on a match

- **R(config-route)#redistribute** *protocol* **[***process id***] route-map** *map-tag*

  - allows for detailed control of routes being redistributed into a routing protocol.

Parameter	Description
map-tag	Specifies the name of the route map
permit | deny	Specifies the action to be taken if the route map match conditions are met ■ permit = permit the matched route to be redistributed ■ deny = deny the matched route from being redistributed
sequence-number	Specifies the sequence number that indicates the position that a new route map statement will have in the list of route map statements already configured with the same route map name

The `match` Command

Command	Description
match community	Matches a BGP community
match interface	Matches any routes that have the next hop out of one of the interfaces specified
match ip address	Matches any routes that have a destination network number address that is permitted by a standard or extended ACL
match ip next-hop	Matches any routes that have a next-hop router address that is passed by one of the ACLs specified
match ip route-source	Matches routes that have been advertised by routers and access servers at the address that is specified by the ACLs
match length	Matches based on the Layer 3 length of a packet
match metric	Matches routes with the metric specified
match route-type	Matches routes of the specified type
match tag	Matches tag of a route

The `set` Command

Command	Description
`set as-path`	Modifies an AS path for BGP routes
`set automatic-tag`	Computes automatically the tag value
`set community`	Sets the BGP communities attribute
`set default interface`	Indicates where to output packets that pass a match clause of a route map for policy routing and have no explicit route to the destination
`set interface`	Indicates where to output packets that pass a match clause of a route map for policy routing
`set ip default next-hop`	Indicates where to output packets that pass a match clause of a route map for policy routing and for which Cisco IOS software has no explicit route to a destination
`set ip next-hop`	Indicates where to output packets that pass a match clause of a route map for policy routing
`set level`	Indicates where to import routes for IS-IS and OSPF
`set local-preference`	Specifies a BGP local preference value
`set metric`	Sets the metric value for a routing protocol
`set metric-type`	Sets the metric type for the destination routing protocol
`set tag`	Sets tag value for destination routing protocol
`set weight`	Specifies the BGP weight value

Route Maps and Redistribution Commands

```
Router(config)# router ospf 10
Router(config-router)# redistribute rip route-map redis-rip
```

```
Router(config)#
route-map redis-rip permit 10
match ip address 23  29
set metric 500
set metric-type type-1

route-map redis-rip deny 20
match ip address 37

route-map redis-rip permit 30
set metric 5000
set metric-type type-2
```

```
Router(config)#
access-list 23 permit 10.1.0.0 0.0.255.255
access-list 29 permit 172.16.1.0 0.0.0.255
access-list 37 permit 10.0.0.0 0.255.255.255
```

- Routes matching either access list 23 or 29 are redistributed with an OSPF cost of 500, external type 1.
- Routes permitted by access list 37 are not redistributed.
- All other routes are redistributed with an OSPF cost metric of 5000, external type 2.

Administrative Distance

Route Source	Default Distance
Connected interface	0
Static route	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIPv1, RIPv2	120
External EIGRP	170
Internal BGP	200
Unknown	255

Modifying Administrative Distance

- **`R(config-route)#distance`** *`administrative distance`* **`[`***`address wildcard-mask`***`] [`***`access-list-number`*** **`|`** ***`name`***`]`

  - Used for all protocols except EIGRP and BGP redistribution

Parameter	Description
`administrative distance`	Sets the administrative distance. An integer from 1 to 255. Routes with a distance of 255 are not installed in the routing table. A value of 0 is reserved for directly connected networks.
`address`	(Optional) Specifies the IP address and filters networks according to the IP address of the router supplying the routing information.
`wildcard-mask`	(Optional) Specifies the wildcard mask for an IP address. A bit set to 1 in the mask argument instructs the software to ignore the corresponding bit in the address value. Use an address/mask of 0.0.0.0 255.255.255.255 to match any IP address (any source router supplying the routing information).
`access-list-number` **`|`** *`name`*	(Optional) Specifies the number or name of the standard ACL to apply to incoming routing updates. It allows filtering of the advertised networks.

- **`R(config-route)#distance eigrp`** *`internal-distance external-distance`*

  - Used for EIGRP

Parameter	Description
`external-distance`	Sets the administrative distance for EIGRP external routes. External routes are routes for which the best path is learned from a neighbor external to the AS.
`internal-distance`	Specifies the administrative distance for EIGRP internal routes. Internal routes are routes that are learned from another entity within the AS.

- **`R(config-route)#distance bgp`**
  - Used for BGP

Example: Redistribution Using Administrative Distance

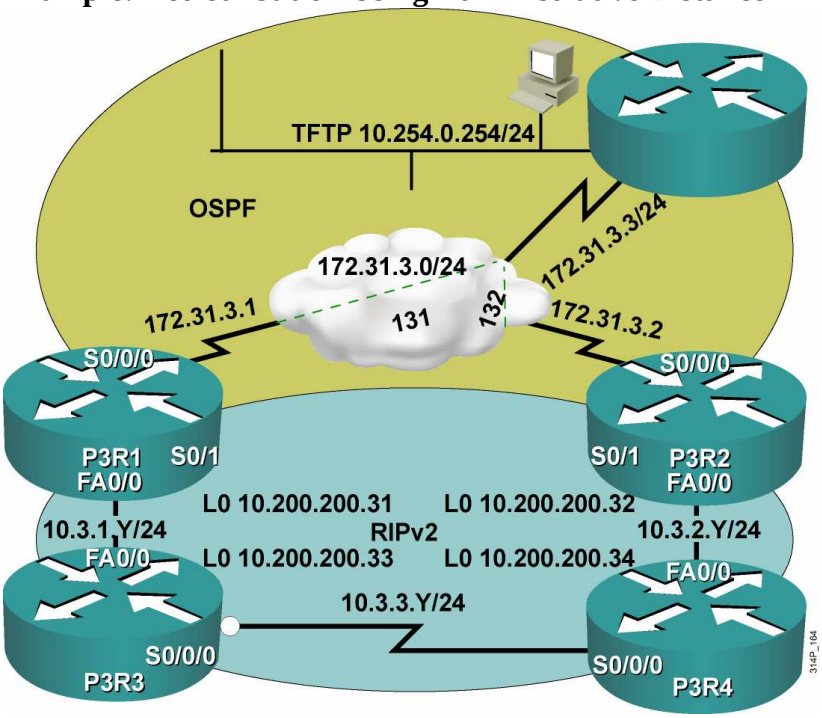

- P3R1 or P3R2 learned 10.3.3.0 via RIPv2 and OSPF. But use OSPF because lower AD even though the path via OSPF might be the longer (worse) path.
- Configuration for P3R1

```
router ospf 1
 redistribute rip metric 10000 metric-type 1 subnets
 network 172.31.0.0 0.0.255.255 area 0
!
router rip
 version 2
 redistribute ospf 1 metric 5
 network 10.0.0.0
 no auto-summary
```

- Configuration for P3R2

```
router ospf 1
 redistribute rip metric 10000 metric-type 1 subnets
 network 172.31.3.2 0.0.0.0 area 0
!
router rip
 version 2
 redistribute ospf 1 metric 5
 network 10.0.0.0
 no auto-summary
```

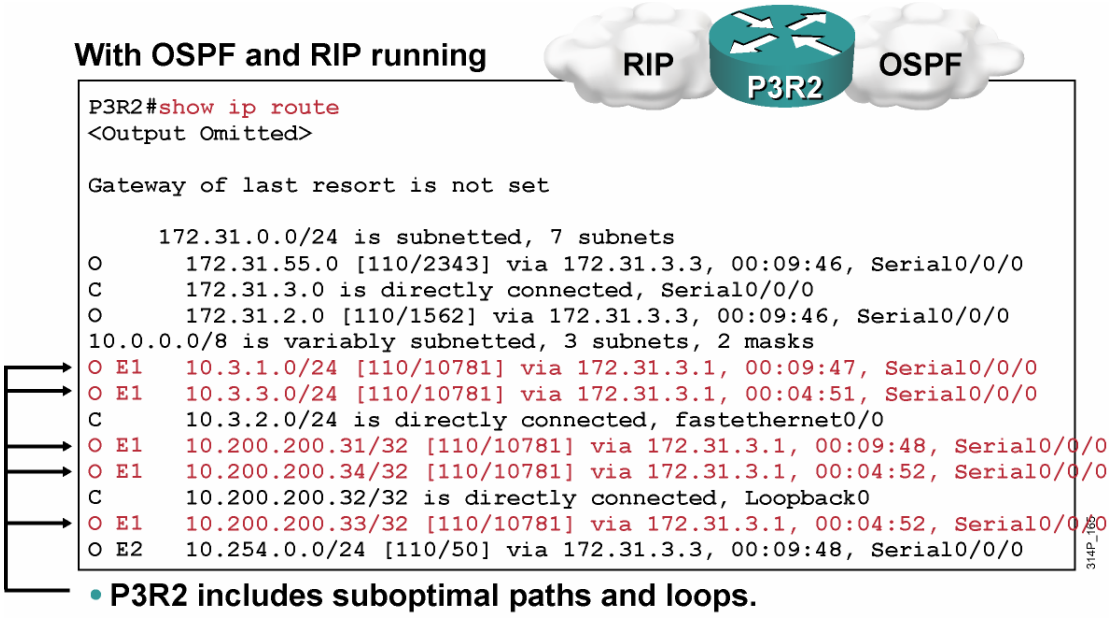

- Something needed to be done to prevent suboptimal path.
- Change the AD of the redistributed RIP route to ensure that the boundary routers select the native RIP routes.

- P3R1

```
hostname P3R1
!
router ospf 1
 redistribute rip metric 10000 metric-type 1 subnets
 network 172.31.0.0 0.0.255.255 area 0
 distance 125 0.0.0.0 255.255.255.255 64
!
router rip
 version 2
 redistribute ospf 1 metric 5
 network 10.0.0.0
 no auto-summary
!
access-list 64 permit 10.3.1.0 0.0.0.255
access-list 64 permit 10.3.3.0 0.0.0.255
access-list 64 permit 10.3.2.0 0.0.0.255
access-list 64 permit 10.200.200.31
access-list 64 permit 10.200.200.34
access-list 64 permit 10.200.200.32
access-list 64 permit 10.200.200.33
```

- P3R2

```
hostname P3R2
!
router ospf 1
redistribute rip metric 10000 metric-type 1 subnets
 network 172.31.3.2 0.0.0.0 area 0
 distance 125 0.0.0.0 255.255.255.255 64
!
router rip
 version 2
 redistribute ospf 1 metric 5
 network 10.0.0.0
 no auto-summary
!
access-list 64 permit 10.3.1.0 0.0.0.255
access-list 64 permit 10.3.3.0 0.0.0.255
access-list 64 permit 10.3.2.0 0.0.0.255
access-list 64 permit 10.200.200.31
access-list 64 permit 10.200.200.34
access-list 64 permit 10.200.200.32
access-list 64 permit 10.200.200.33
```

- The **`distance`** Command Parameters

Parameter	Description
`125`	Defines the administrative distance that specified routes will be assigned
`0.0.0.0 255.255.255.255`	Defines the source address of the router supplying the routing information—in this case, any router
`64`	Defines the ACL to be used to filter incoming routing updates to determine which will have their administrative distance changed

- The **`access-list`** Command Parameters

Parameter	Description
`64`	Displays the ACL number
`permit`	Allows all networks that match the address to be permitted, in this case to have their administrative distance changed
`10.3.1.0`	Displays a network to be permitted, in this case, to have its administrative distance changed

With OSPF changing administrative distance


```
Gateway of last resort is not set

     172.31.0.0/16 is variably subnetted, 8 subnets, 2 masks
O       172.31.55.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
C       172.31.33.0/24 is directly connected, Serial0/0/0
O       172.31.33.1/32 [110/1562] via 172.31.33.4, 00:00:01, Serial0/0/0
O       172.31.33.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
O       172.31.44.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
O       172.31.22.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
O       172.31.11.4/32 [110/781] via 172.31.33.4, 00:00:03, Serial0/0/0
O       172.31.66.4/32 [110/781] via 172.31.33.4, 00:00:03, Serial0/0/0
     10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
R       10.3.1.0/24 [120/2] via 10.3.2.4, 00:00:03, FastEthernet0/0
R       10.3.3.0/24 [120/1] via 10.3.2.4, 00:00:03, FastEthernet0/0
C       10.3.2.0/24 is directly connected, FastEthernet0/0
R       10.200.200.31/32 [120/3] via 10.3.2.4, 00:00:04, FastEthernet0/0
R       10.200.200.34/32 [120/1] via 10.3.2.4, 00:00:04, FastEthernet0/0
C       10.200.200.32/32 is directly connected, Loopback0
R       10.200.200.33/32 [120/2] via 10.3.2.4, 00:00:04, FastEthernet0/0
O E2    10.254.0.0/24 [110/50] via 172.31.33.4, 00:00:04, Serial0/0/0
```

314P_166

- **Router P3R2 prefers RIP routes.**

Lesson 4 Implementing Advanced Cisco IOS Features: Configuring DHCP

DHCP in an Enterprise Network

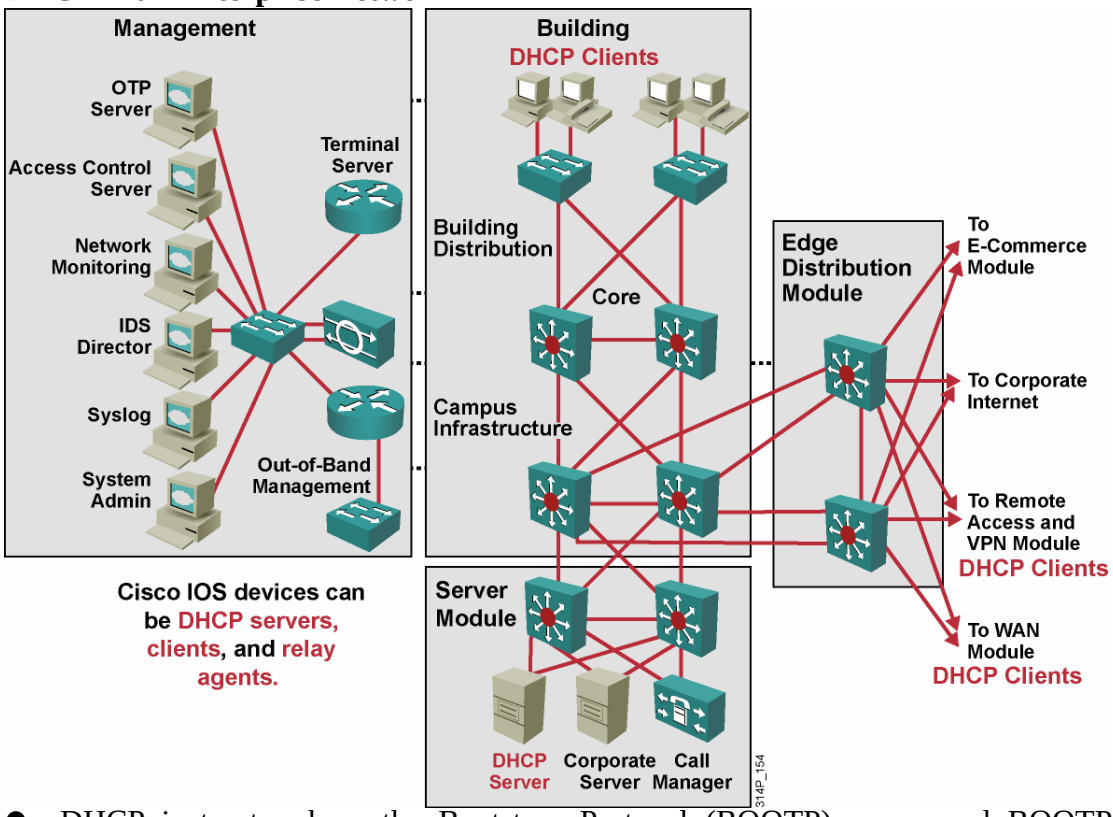

- DHCP is tructured on the Bootstrap Protocol (BOOTP) server and BOOTP well-know ports in User Datagram Protocol (UDP). Previous to DHCP, IP addresses were manually administered to IP hosts, which was a tedious, error-prone, and labor-intensive process.

Dynamic Host Configuration Protocol

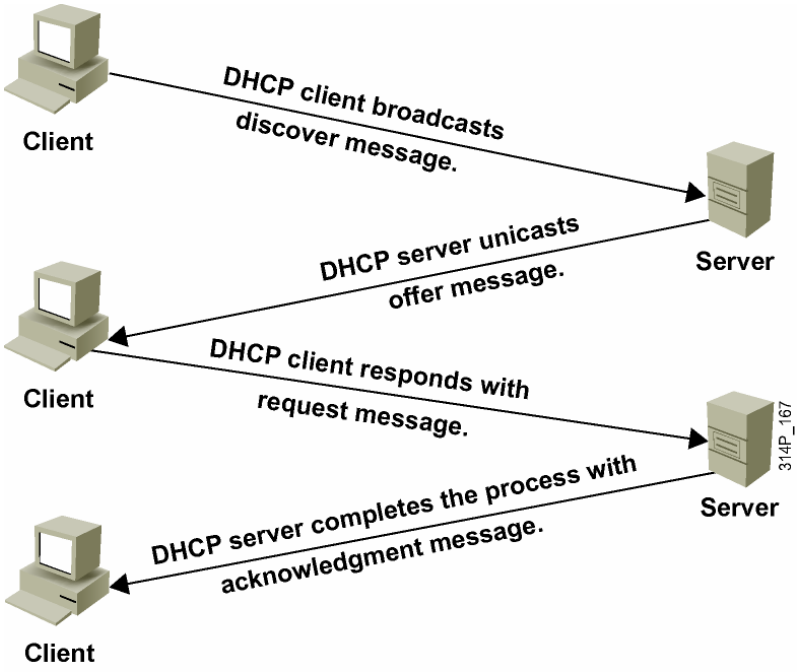

- DHCPDISCOVE broadcast message
- DHCPOFFER unicast message
- DHCPREQUEST broadcast message
- DHCPACK unist message

- DHCP supports three possible address allocation mechanisms:

- Manual: the network administrator assigns the IP address to a specific MAC address. DHCP is used to dispatch the assigned address to the host.
- Automatic: the IP address is permanently assigned to a host.
- Dynamic: the IP address is assigned to a host for a limited time or until the host explicitly releases the address. This mechanism supports automatic address reuse when the host to which the adderss has been assigned no longer needs the address.

Configuring a DHCP Server

- **`R(config)#ip dhcp pool [pool name]`**
  - Enables a DHCP pool for use by hosts
- **`R(config-dhcp)#import all`**
  - imports DNS and WINS information for IPCP
- **`R(config-dhcp)#network [network address] [subnet mask]`**
  - specifies the network and subnet mask of the pool
- **`R(config-dhcp)#default-router [host address]`**
  - specifies the default router for the pool to use

- Cisco IOS DHCP Server Command and Parameters

Command and Parameter	Description
`Router(config)#service dhcp`	Enables DHCP features on router; it is on by default.
`Router(config)#ip dhcp database` *url* [`timeout` *seconds* | *write-delay seconds*]	Configures the database agent and the interval between database updates and database transfers.
`Router(config)#no ip dhcp conflict logging`	Disables DHCP conflict logging. (Used if a DHCP database agent is not configured.)
`Router(config)#ip dhcp excluded-address` *low-address* [*high address*]	Specifies the IP address that the DHCP server should not assign to DHCP clients.
`Router(config)#ip dhcp pool` *name*	Creates a name for the DCHP server address pool and places you in DHCP pool configuration mode.
`Router(config-dhcp)#network` *network-number* [*mask* | *prefix-length*]	Specifies the subnet network number and mask of the DHCP address pool.
`Router(config-dhcp)#domain-name` *domain*	Specifies the domain name for the client.
`Router(config-dhcp)#dns-server` *address* [*address2...address8*]	Specifies the IP address of a Domain Name System (DNS) server that is available to a DHCP client. One is required, but up to eight can be specified.
`Router(config-dhcp)#netbios-name-server` *address* [*address2...address8*]	Same as DNS, but for Windows Internet Name Service (WINS).
`Router(config-dhcp)#default-router` *address* [*address2... address8*]	Specifies the IP address of the default router for a DHCP client.
`Router(config-dhcp)#lease` {*days* [*hours*] [*minutes*] | *infinite*}	Specifies the duration of the lease. The default is a one-day lease.
`Router(config-dhcp)#import all`	Used to import DHCP option parameters into the DHCP server database. Used for remote DHCP pools.

DHCP Server Configuration Example

```
ipdhcp database ftp://user:passwords@172.16.4.253/router-dhcp write-delay 120
ip dhcp excluded-address 172.16.1.100 172.16.1.103
ip dhcp excluded-address 172.16.2.100 172.16.2.103
ip dhcp pool 0
  network 172.16.0.0/16
  domain-name global.com
  dns-server 172.16.1.102 172.16.2.102
  netbios-name-server 172.16.2.103 172.16.2.103
  default-router 172.16.1.100
```

- Before configuring the Cisco IOS DHCP Server featuer:
  - Identify an external FTP, TFTP, or Remote Copy Protocol (RCP) server that will be used to store the DHCP binding database.
  - Identify the IP address range to be assigned by the DHCP server and the IP addresses to be excluded (default routers and other statically assigned addresses within a dynamically assigned range).
  - Identify DHCP options for devices where necessary, including these:
    - Default boot image name
    - Default routers
    - DNS servers
    - Network BISO (NetBIOS) name server
    - IP telephony options, such as option 150
  - Decide on a DNS domain name
- The following example, two DHCP address pools are created, one for subnet 172.16.1.0 and on for subnet 172.16.2.0. the host will receive its address from the DHCP server on the closest device. The router determines how to provide the addresses based on the IP address on the router interface.

```
!
ip dhcp pool 1
  network 172.16.1.0/24
  domain-name global.com
  dns-server 172.16.1.102
  netbios-name-server 172.16.2.103
  default-router 172.16.1.100 172.16.1.101
  lease 30
!
ip dhcp pool 2
  network 172.16.2.0/24
  domain-name global.com
  dns-server 172.16.2.102
  netbios-name-server 172.16.2.103
  default-router 172.16.2.100 172.16.2.101
  lease 30
```

Importing and Autoconfiguration

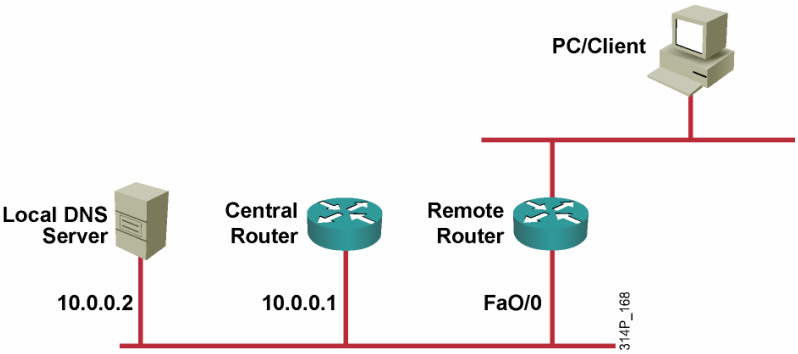


```
Central Router
  ip dhcp-excluded address 10.0.0.1 10.0.0.5
  ip dhcp pool central
  network 10.0.0.0 255.255.255.0
  default-router 10.0.0.1 10.0.0.5
  domain-name central.com
  dns-server 10.0.0.2
  interface fastethernet0/0
    ip address 10.0.0.1 255.255.255.0
Remote Router
  ip dhcp pool client
  network 20.0.0.0 255.255.255.0
  ip dhcp-excluded address 20.0.0.2
  default-router 20.0.0.2
  import all
  interface fastethernet0/0
    ip address dhcp
```

DHCP Client

- **`R(config-if)#ip address dhcp`**
  - Enables a Cisco IOS device to obtain an IP address dynamically from a DHCP server

Helper Addressing Overview

- Routers do not forward broadcasts, by default.
- Helper address provides selective connectivity.
- The Cisco IOS DHCP relay agent is enabled on an interface only when the **`ip helper-address`** is configured.
- The DHCP relay agent is any host that forwards DHCP packets between clients

and servers. Relay agents are used to forward requests and replies between clients and servers when they are not on the same physical subnet.

- Relay agents receive DHCP messages and then generate a new DHCP message to send out on another interface.

Why Use a Helper Address?

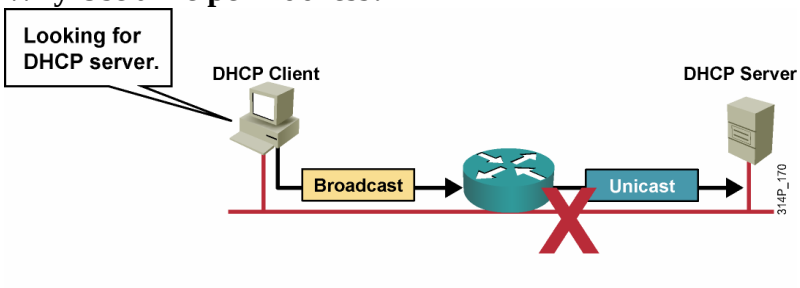

- Sometimes clients do not know the server address.
- Helpers change broadcast to unicast to reach server.
- DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER message, because they do not have information about the network to which they are attached.
- If the client is on a network that does not include a server, UDP broadcasts are normally not forwarded by the attached router.
- The **`ip helper-address`** command causes the UDP broadcast to be changed to a unicast and forwarded out another interface to a unicast IP address specified by the command.
- The relay agent sets the gateway address (giaddr field of the DHCP packet) and, if configured, adds the relay agent information option (option 82) in the packet and forwards it to the DHCP server. The reply from the server is forwarded back to the client after removing option 82.

IP Helper Address Commands

- **`R(config)#ip helper-address`** *`address`*

  - Enables forwarding and specifies destination address for main UDP broadcast packets

  - Changes destination address from broadcast to unicast or directed broadcast address

- **`R(config)#ip forward-protocol {udp [`***`port`***`]}`**

  - Specifies which protocols will be forwarded

Multiple Servers: Remote Networks

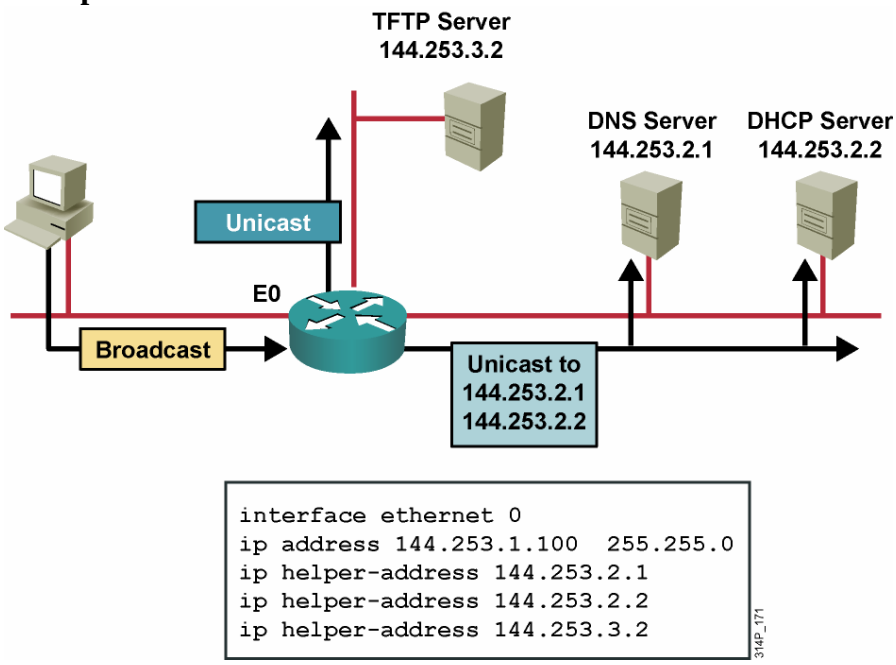

- The UDP well-known ports identified by default forwarding UDP services are:
  - Time: 37
  - TACACS: 49
  - DNS: 53
  - BOOTP/DHCP server: 67
  - BOOTP/DHCP client: 68
  - TFTP: 69
  - NetBIOS name service: 137
  - NetBIOS datagram service: 138

Ports can be eliminated from the forwarding service with the **`no ip forward-protocol`** command, and ports can be added to the forwarding service with the **`ip forward-protocol`** command. The following is an example:

```
!
interface fastethernet0/0
ip address 10.3.3.3 255.255.255.0
ip helper-address 10.1.1.1
no ip forward-protocol upd 137
no ip forward-protocol upd 138
no ip forward-protocol upd 37
ip forward-protocol upd 8000
```

This configuration would cause time and NetBIOS ports to not be forwarded, and UDP port 8000 would be added to the forwarded list.

Relay Agent Option Support

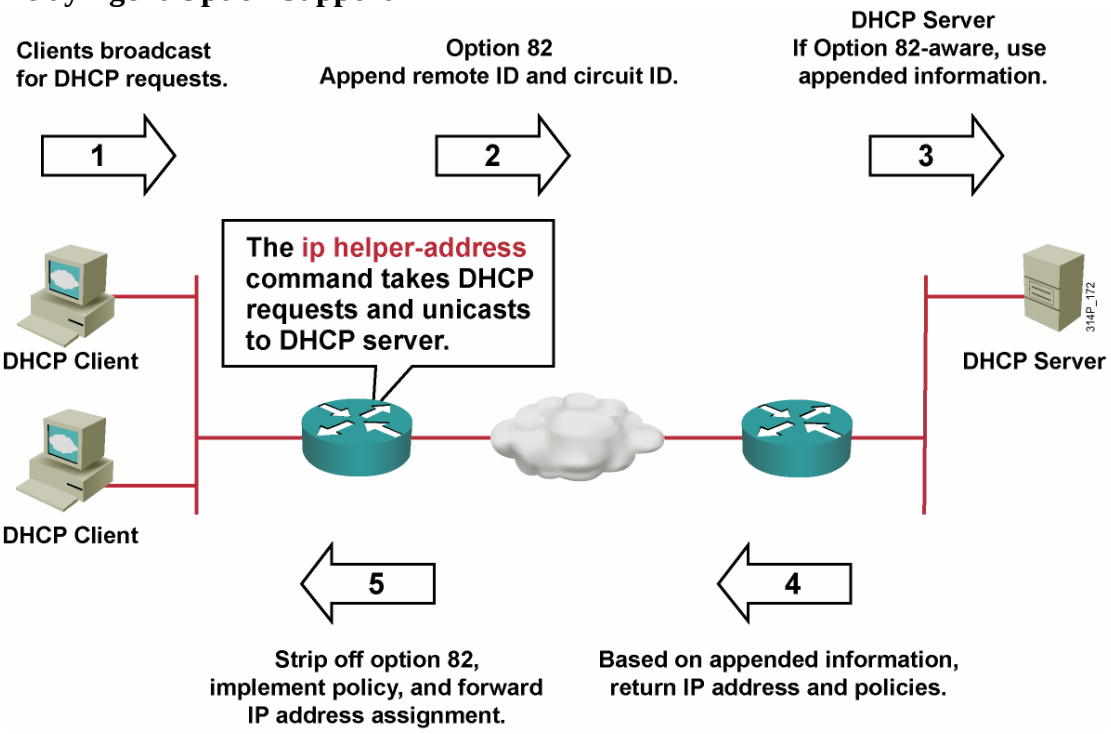

- When you use the **`ip dhcp relay information option`** command, the relay agent adds the circuit identifier suboption and the remote ID suboption to the relay agent information and forwards them to a DHCP server. The following explains the DHCP relay services process:

  1. The DHCP client generates a DHCP request and broadcasts it on the network.
  2. The DHCP relay agent intercepts the broadcast DHCP request packet and inserts the relay agent information option (82) in the packet. The relay agent option contains the related suboptions.
  3. The DHCP relay agent unicasts the DHCP packet to the DHCP server.
  4. The DHCP server receives the packet and uses the suboptions to assign IP addresses and other configuration parameters and forwards them back to the client.
  5. The suboption fields are stripped off of the packet by the relay agent while forwarding to the client.

- Description of Option Support Commands and Parameters

Command and Parameter	Description
`Router(config)#ip dhcp information option`	Enables the system to insert the DHCP relay agent information option (82) in forwarded BOOTREQUEST messages to a DHCP server. Disabled by default.
`Router(config)#ip dhcp information check`	Configures DHCP to check that the relay agent information option in forwarded BOOTREPLY messages is valid. This command is used only to re-enable this function if it is disabled; it is on by default.
`Router(config)#ip dhcp information policy {drop | keep | replace}`	Configures the reforwarding policy for the relay agent (specifies what a relay agent should do if a message already contains relay information).
`Router(config)#ip dhcp relay information trust-all`	Configures all information on a router as trusted sources of the DHCP relay information option. Useful when Ethernet switches are involved in delivery of packet. For an individual interface, the **ip dhcp relay information trusted** command can be used.
`Router#show ip dhcp-relay information trusted-sources`	Displays all interfaces configured to be a trusted source for the DHCP relay information option.

DHCP Verification Commands

- **`R#show ip dhcp database`**
  - Displays recent activity on the DHCP database
- **`R#show ip dhcp server statistics`**
  - Show count information about statistics and messages sent and received
- **`R#show ip route dhcp`**
  - Displays routes added to the routing table by DHCP
- **`R#show ip dhcp server {events | packets | linkage}`**
  - Enables debugging on the DHCP server

Description of DHCP Verification Commands and Parameters

Command and Parameter	Description
`show ip dhcp database`	Displays recent activity in the DHCP database
`show ip dhcp server statistics`	Displays count information about server statistics and messages sent and received
`show ip route dhcp` {*ip-address*}	Displays the routes added to the routing table by the Cisco IOS DHCP server and relay agent
`show ip dhcp binding` {*address*}	Displays a list of all bindings created on a specific DHCP server
`show ip dhcp conflict`	Displays a list of all address conflicts recorded by a specific DHCP server
`show ip dhcp import`	Displays the option parameters that were imported into the DHCP server database
`clear ip dhcp binding` {*address* | *}*	Deletes an automatic address binding from the DHCP database
`clear ip dhcp conflict` *{address | *}*	Clears an address conflict from the DHCP database
`clear ip dhcp server statistics`	Resets all DHCP server counters to 0
`clear ip route dhcp` {*ip-address*}	Removes routes from the routing table added by the Cisco IOS DHCP server and relay agent