Form BGP Neighbors

Two BGP routers become neighbors after the routers establish a TCP connection between each other. The TCP connection is essential in order for the two peer routers to start the exchange of

The

match interface

match ip address

match ip next-hop

match ip route-source

match metric

match route-type

The related commands for1 12/F1 12Td 22.75 -1set1 12Td12 Td ( )Tj -2(are:j -20atch tag)p4 Td /I

set origin
in
in

neighbor 2.2.2.34.4remote-as 300

neIn this example, if a route matches the IP address 1716 0.0.0, the route has a metric of.2. Then,

neNote: neAlways ask.the question ₩hat happens to routes that do not match any of.the match
ne
Example 2
neSuppose that, inExample 1, you do not want AS100 to accept updates about 1716 0.0.0. ₩u
, mus-luupprouly e inbothe route apsRTC:c 5.

neighbor 2.2.2.34.4remote1as 300

Examp6.852

, feel more comfortabxampply howuto /Frt BGe mnbohowuto byiress 1h.Tt96d /F/F1 12f (ne)Tj0 1.owuto b

neNote:

The format of the **network**

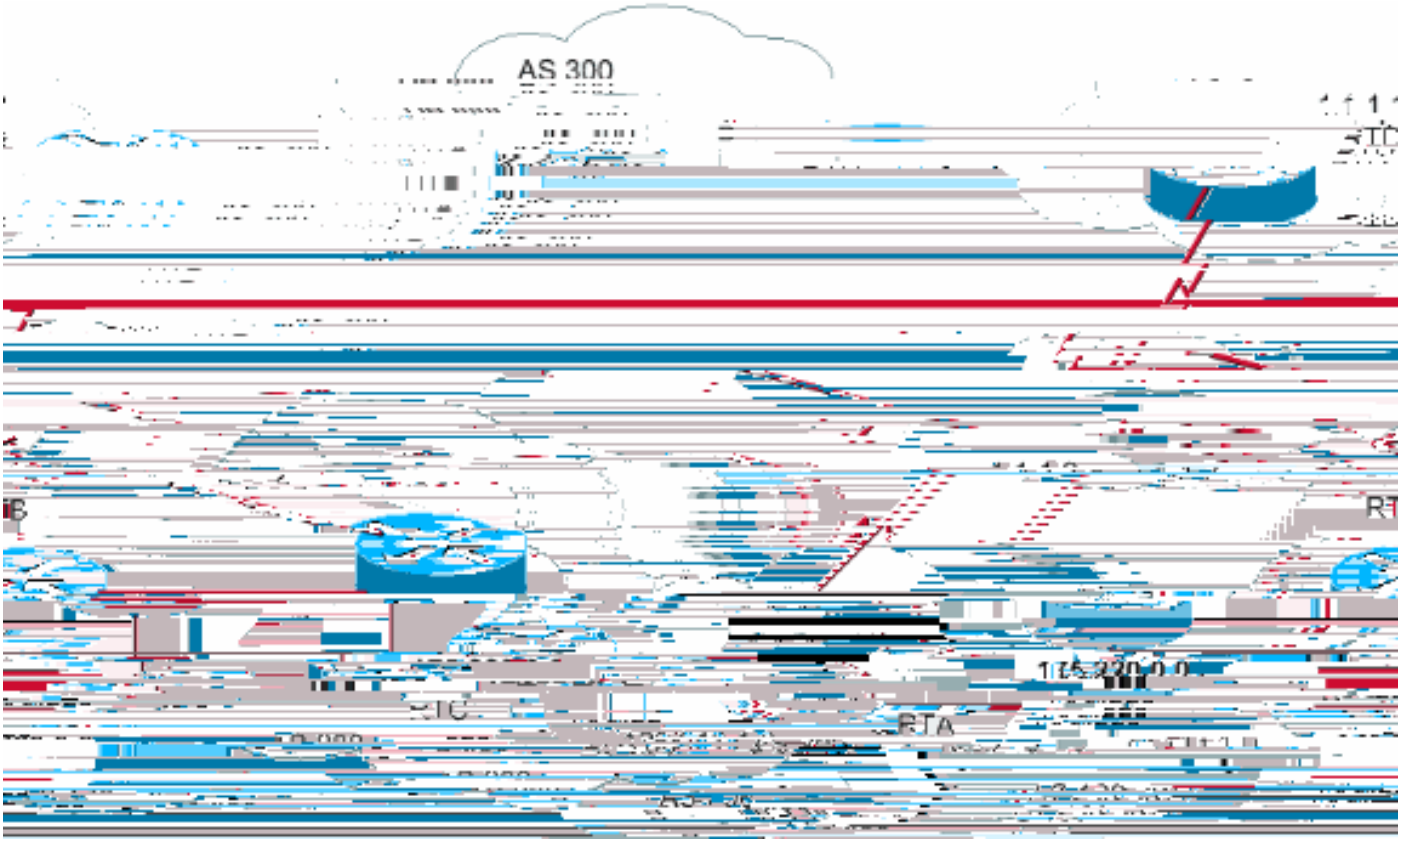
AS 300
RTA

```
network 175.220.0.0
```

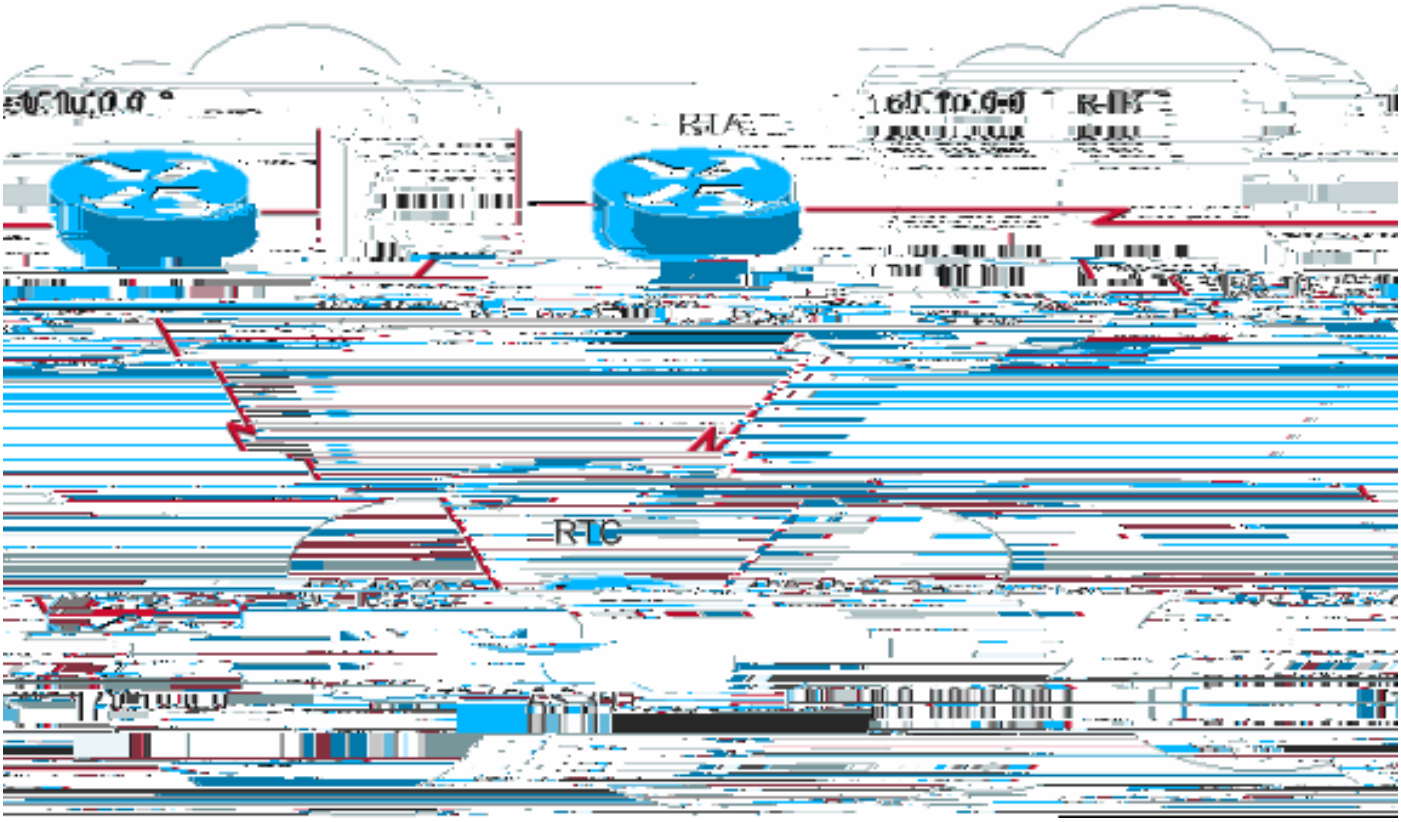

You use iBGP if an AS wants to act as a transit system to other ASs. Is it true that you can do the same thing by learning via eBGP, redistributing into IGP, and then redistributing again into another AS? Yesy r n t you can do the
sab5rT1tore flexibilityo anotore efficientto yher Aexchangtranformation wiribun t you can do the

```
RTA#
router bgp 100
neighbor 190.10.50.1 remote-as 100
neighbor 170.10.20.2 remote-as 300
network 150.10.0.0
redistribute static

ip route 190.10.0.0 255.255.0.0 null0

RTB#
router bgp 100
neighbor 150.10.30.1 remote-as 100
network 190.10.50.0
RTE#
router bgp 300
neighbor 170.10.20.1 remote-as 100
network 170.10.0.0
```

next AS is 100 and that the origin is incomplete and comes from a static route.

BGP Next Hop Attribute

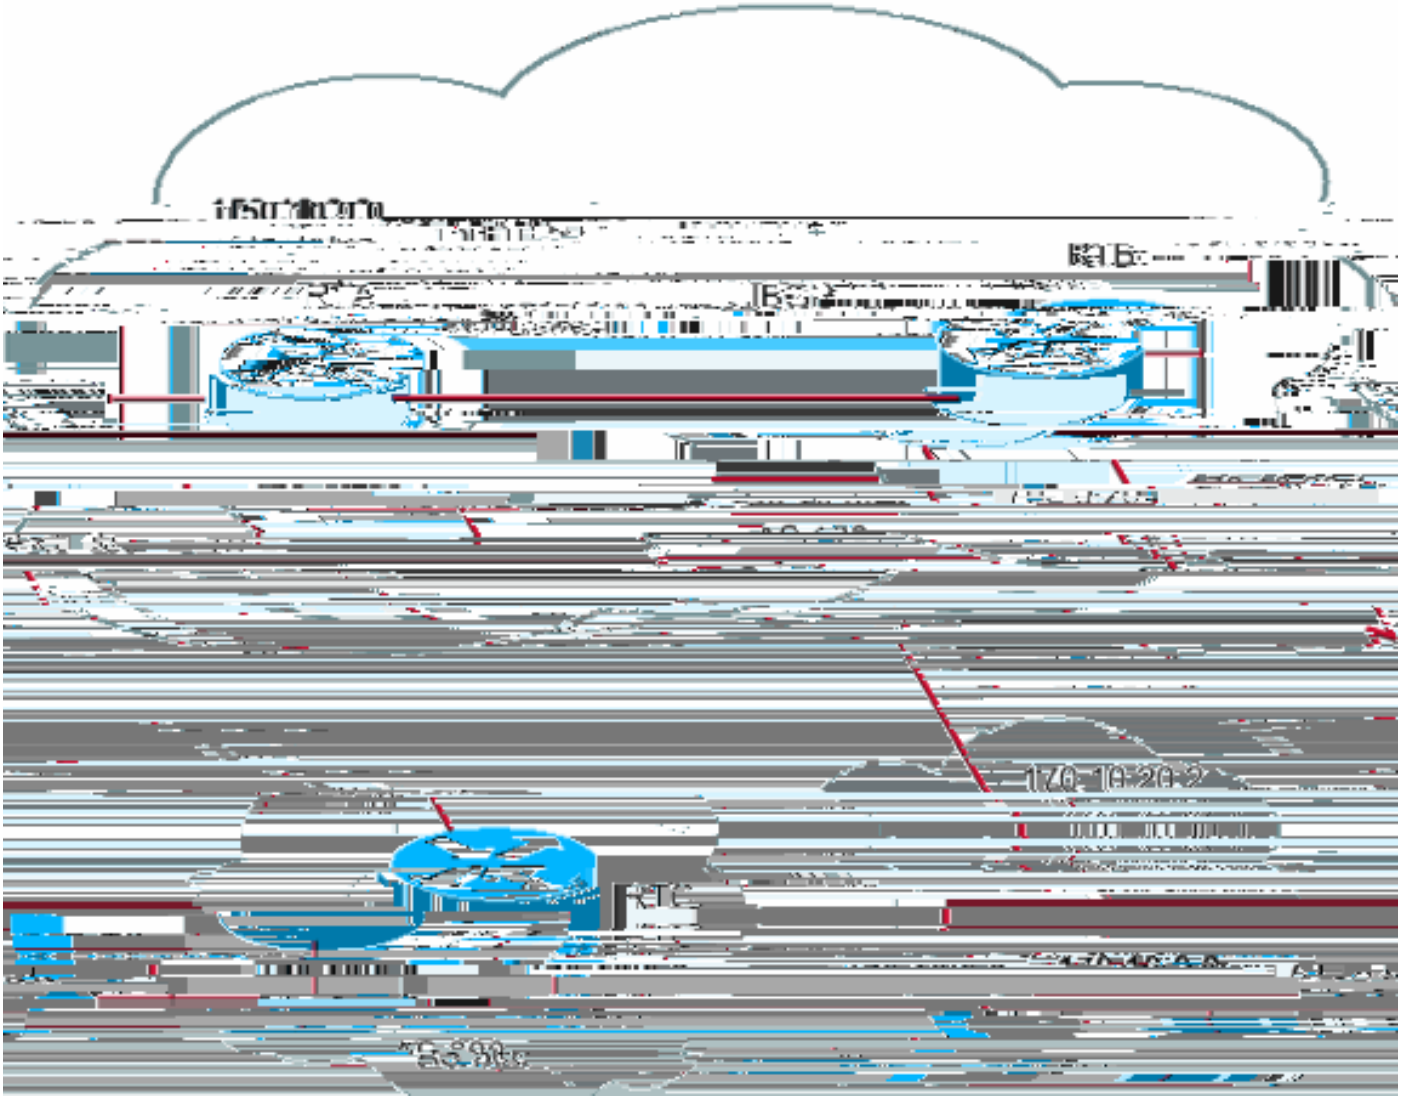

The BGP next hop attribute is the next hop IP address to use in order to reach a certain destination.

For eBGP, the next hop is always the IP address of the neighbor that the **neighbor** command specih0 0 Td 0 -14.4 Td (For eBthat thefn)Tj /exaic r

```
network 150.10.0.0
RTB#
router bgp 100
neighbor 150.10.30.1 remote-as 100
RTC#
router bgp 300
neighbor 170.10.20.1 remote-as 100
network 170.10.0.0
```

Note: RTC advertises 170.10.0.0 to RTA with a next hop equal to 170.10.20.2.

Note: RTA advertises 170.10.0.0 to RTB with a next hop equal to 170.10.20.2. The eBGP next Tus carried in iBGP.

network 180.20.0.0 via 170.10.20.3. When RTC sends a BGP update to RTA with regard to

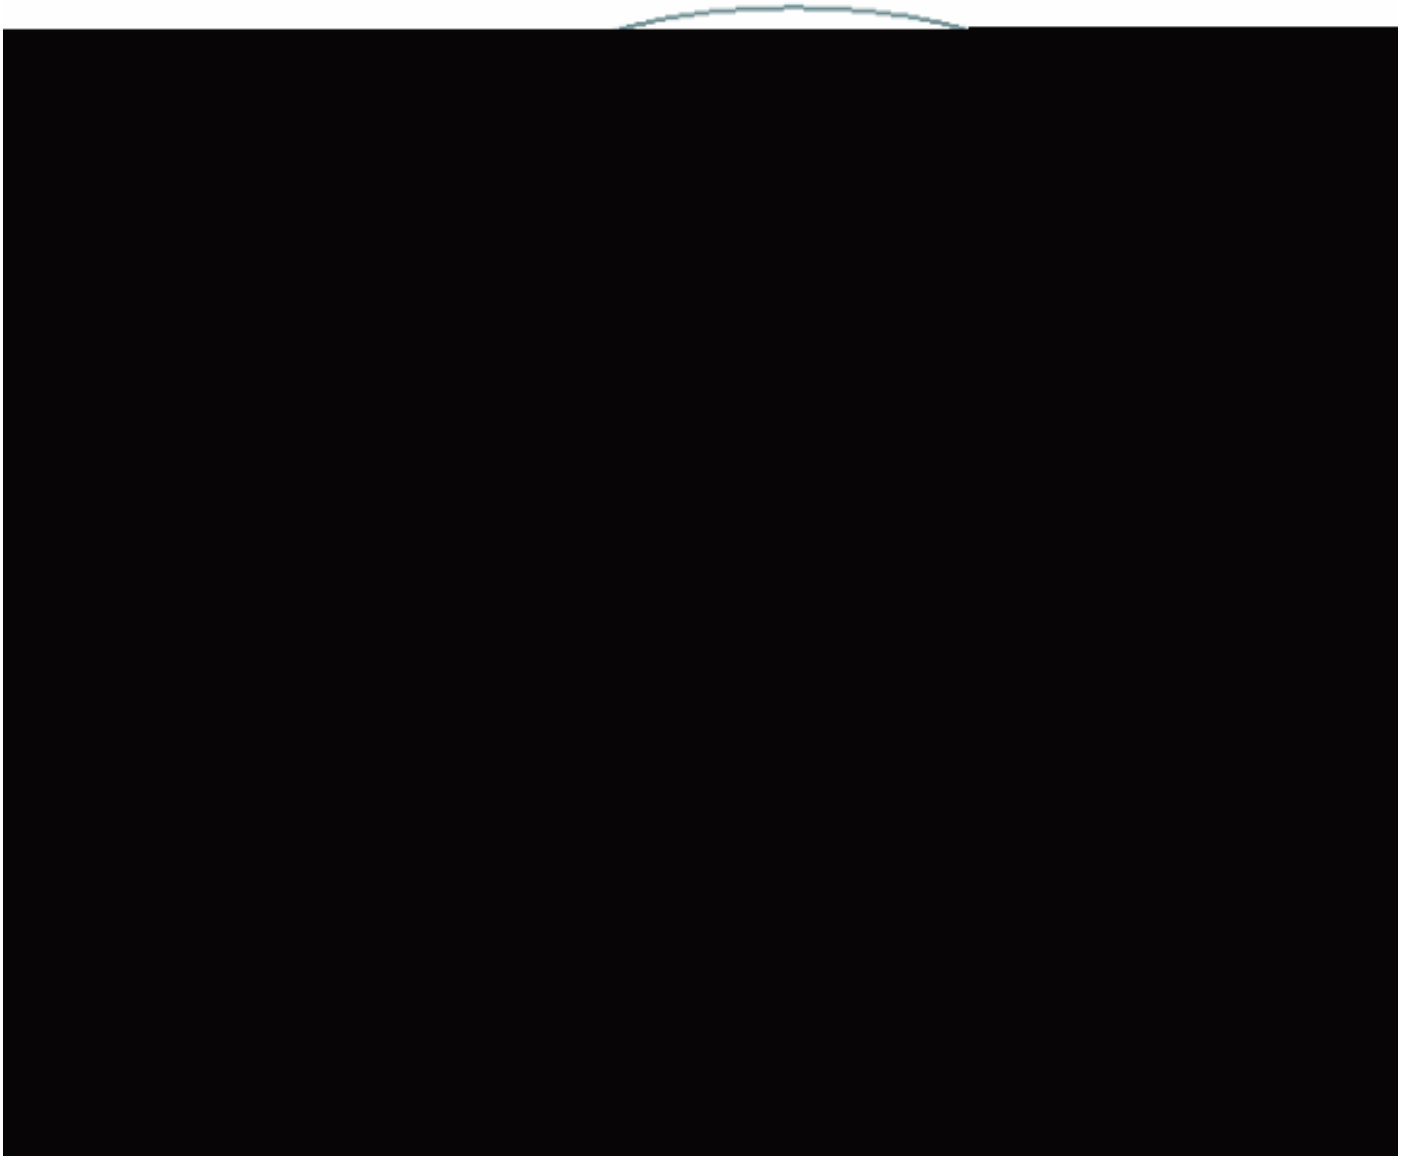

next-hs a I.q20 1 RG 0.75 w 257.45 7mmand

90 for EIGRP

110 for OSPF

```
router bgp 100
neighbor 2.2.2.1 remote-as 300
```

waits until IGP has propagated the route within the AS. Then, BGP advertises the route to external peers.

```
RTB#
router bgp 100
network 150.10.0.0
neighbor 1.1.1.2 remote-as 400
neighbor 3.3.3.3 remote-as 100
no synchronization
```

Use AS_PATH access lists.

> **ip as-path access-list** *access-list-number* **{permit | deny}** *as-regular-expression*
> **neighbor** *ip-address* **filter-list** *access-list-number* **weight** *weight*

Use route maps.

```
RTC#
router bgp 300
neighbor 1.1.1.1 renmit |deny}
```

```
neighbor 128.213.11.2 remote-as 256
bgp default local-preference 150

RTD#
router bgp 256
neighbor 3.3.3.4 remote-as 300
neighbor 128.213.11.1 remote-as 256
bgp default local-preference 200
```

```
route-map setmetricout permit 10
   set metric 50
```

With these configurations, RTA picks RTC as next hop, with consideration of the fact that all other

AS.

Here are two examples of route maps that set the community:

- `route-map communitymap`
  `match ip address 1`
  `set community no-advertise`

  or

- `route-map setcommunity`
  `match as-path 1`
  `set community 200 additive`

If you do not set the **additive** keyword, 200 replaces any old community that already exits. If you use the keyword **additive**, an addition of 200 to the community occurs. Even if you set the community attribute, this attribute does not transmit to neighbors by default. In order to send the attribute to a neighbor, you must use this command:

neighbor {*ip-address* **|** *peer-group-name***} send-community**
Here is an example:

```
RTA#
router bgp 100
neighbor 3.3.3.3 remote-as 300
neighbor 3.3.3.3 send-community
neighbor 3.3.3.3 route-map setcommunity out
```

In Cisco IOS Software Release 12.0 and later, you can configure communities in three different formats: decimal, hexadecimal, and AA:NN. By default, Cisco IOS Software uses the older decimal format. In order to configure and display in AA:NN, issue the **ip bgp-community new-format** global configuration command. The first part of AA:NN represents the AS number, and the second part represents a 2-byte number.

Here is an example:

Without the **ip bgp-community new-format** command in global configuration, an issue of the **show ip bgp 6.0.0.0** command displays the community attribute value in decimal format. In this example, the community attribute value appears as `6553620`.

```
Router# show ip bgp 6.0.0.0 BGP routing table entry for 6.0.0.0/8, version 7 Paths: (1
available, best #1, table Default-IP-Routing-Table) Not advertised to any peer 1 10.10.10.1 from
10.10.10.1 (200.200.200.1) Origin IGP, metric 0, localpref 100, valid, external, best
Community: 6553620
```

Now, issue the **ip bgp-community new-format** command globally on this router.

```
Router# configure terminal Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip bgp-community new-format Router(config)# exit
```

With the **ip bgp-community new-format** global configuration command, the community value displays in AA:NN format. The value appears as `100:20` inre is nity vtputhe

the transmission of all other updates.

end of the input string, or a space.

Piece

A piece is one of these symbols, yseiich follows an atom:j -22.75 0 Td 22.75 -12 Td /F

^$

This expression indicates origination from this AS.

Refer to Using Regulaneiaoression isndi BGPj 0 0 g

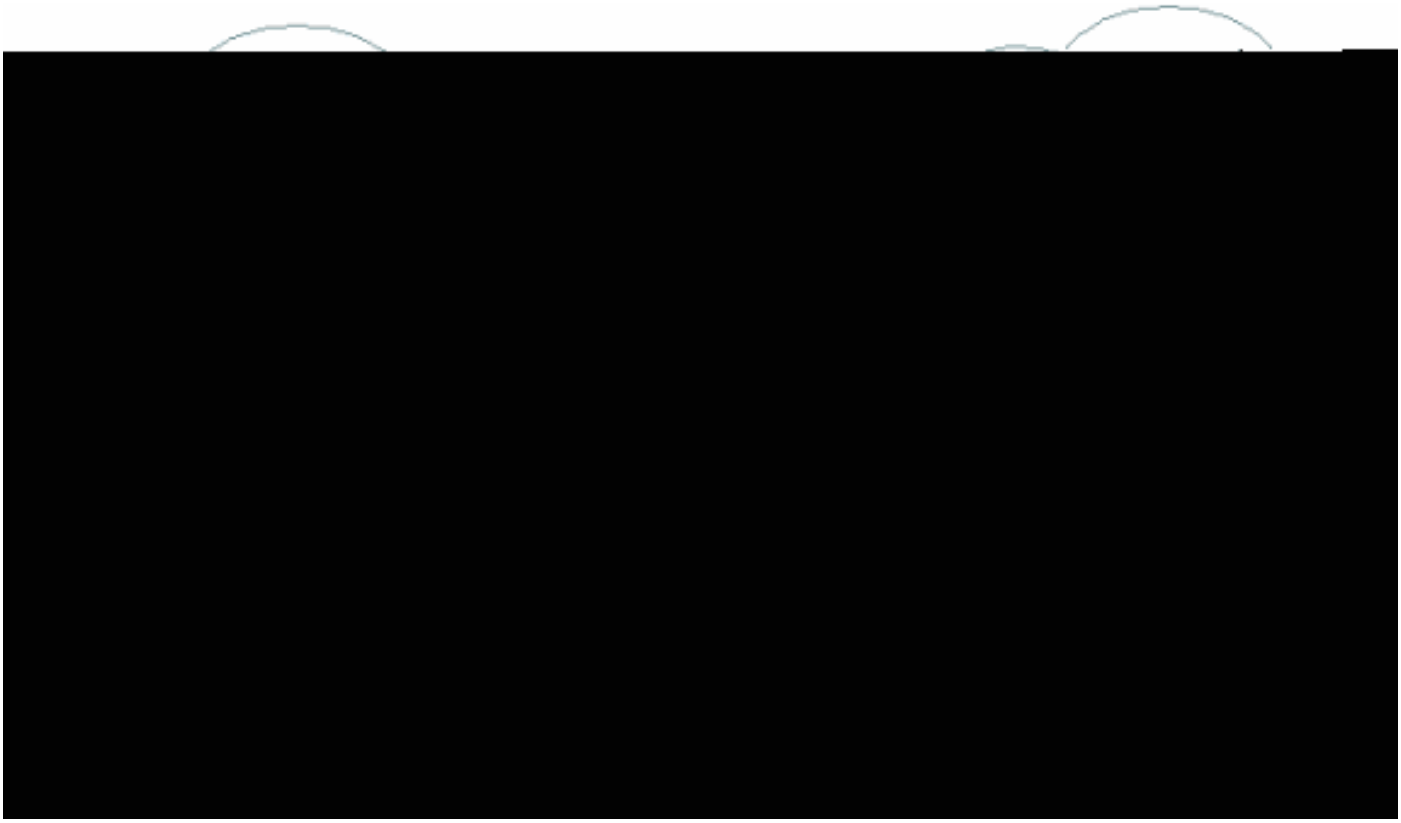

Route maps associated with the **neighbor** statement have no effect on incoming updates when you match based on the IP address:

neighbor *ip-address* **route-map** *route-map-name*

Assume that, in the diagram in this section, you want RTC to learn from AS200 about networks that are local to AS200 and nothing else. Also, you want to set the weight on the accepted routes to 20. Use a combination of **neighbor** and **as-path** access lists:

```
RTC#
   router bgp 300
   network 170.10.0.0
```

aggregate-address

this example, RTC gets updates about 160.20.0.0 from RTA and updates about 160.10.0.0 from RTB. Suppose that RTC wants to aggregate network 160.0.0.0/8 and send the network to RTD. RTD does not know the origin of that route. If you add the **aggregate as-set** statement, you force

```
router bgp 300
neighbor 3.3.3.3 remote-as 200
neighbor 2.2.2.2 remote-as 100
neighbor 4.4.4.4 remote-as 400
aggregate 160.0.0.0 255.0.0.0 summary-only
aggregate 160.0.0.0 255.0.0.0 as-set
!--- This command causes RTC to send RTD updates about 160.0.0.0/8 !--- with an indication that
160.0.0.0 belongs to a set {100 200}.
```

The next two subjects, BGP Confederation and Route Reflectors, are for Internet service providers (ISPs) that want further control of the explosion of iBGP peering inside their ASs.ce providersg5 wpploET

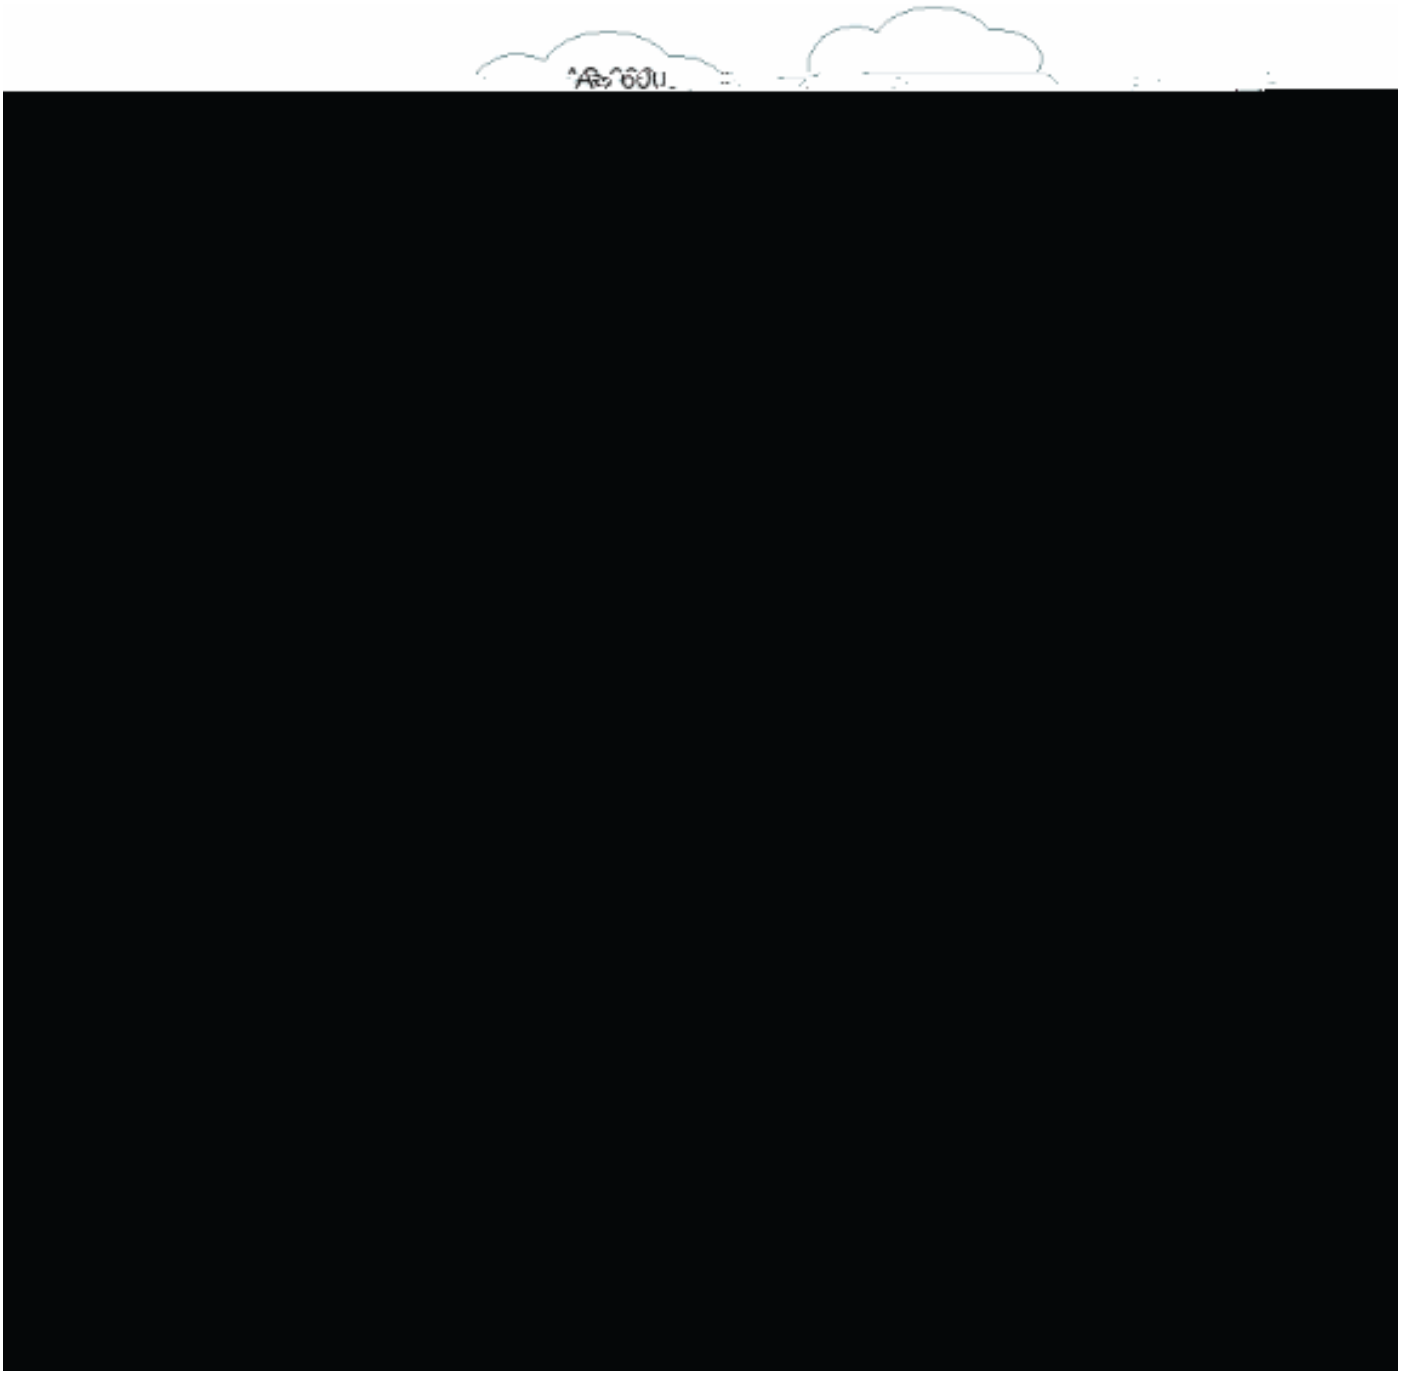

Assume that you have an AS500 that consists of nine BGP speakers. Other non-BGP speakers

In normal cases, maintain a full iBGP mesh between RTA, RTB, and RTC within AS100. If you utilize the RR concept, RTC can be elected as an RR. In this way, RTC has a partial iBGP peering with RTA and RTB. Peering between RTA and RTB is not necessary because RTC is an RR for the updates that come from RTA and RTB.

ugug

exchange updates through the RR. If you configure peer groups, a potential withdrawal to the

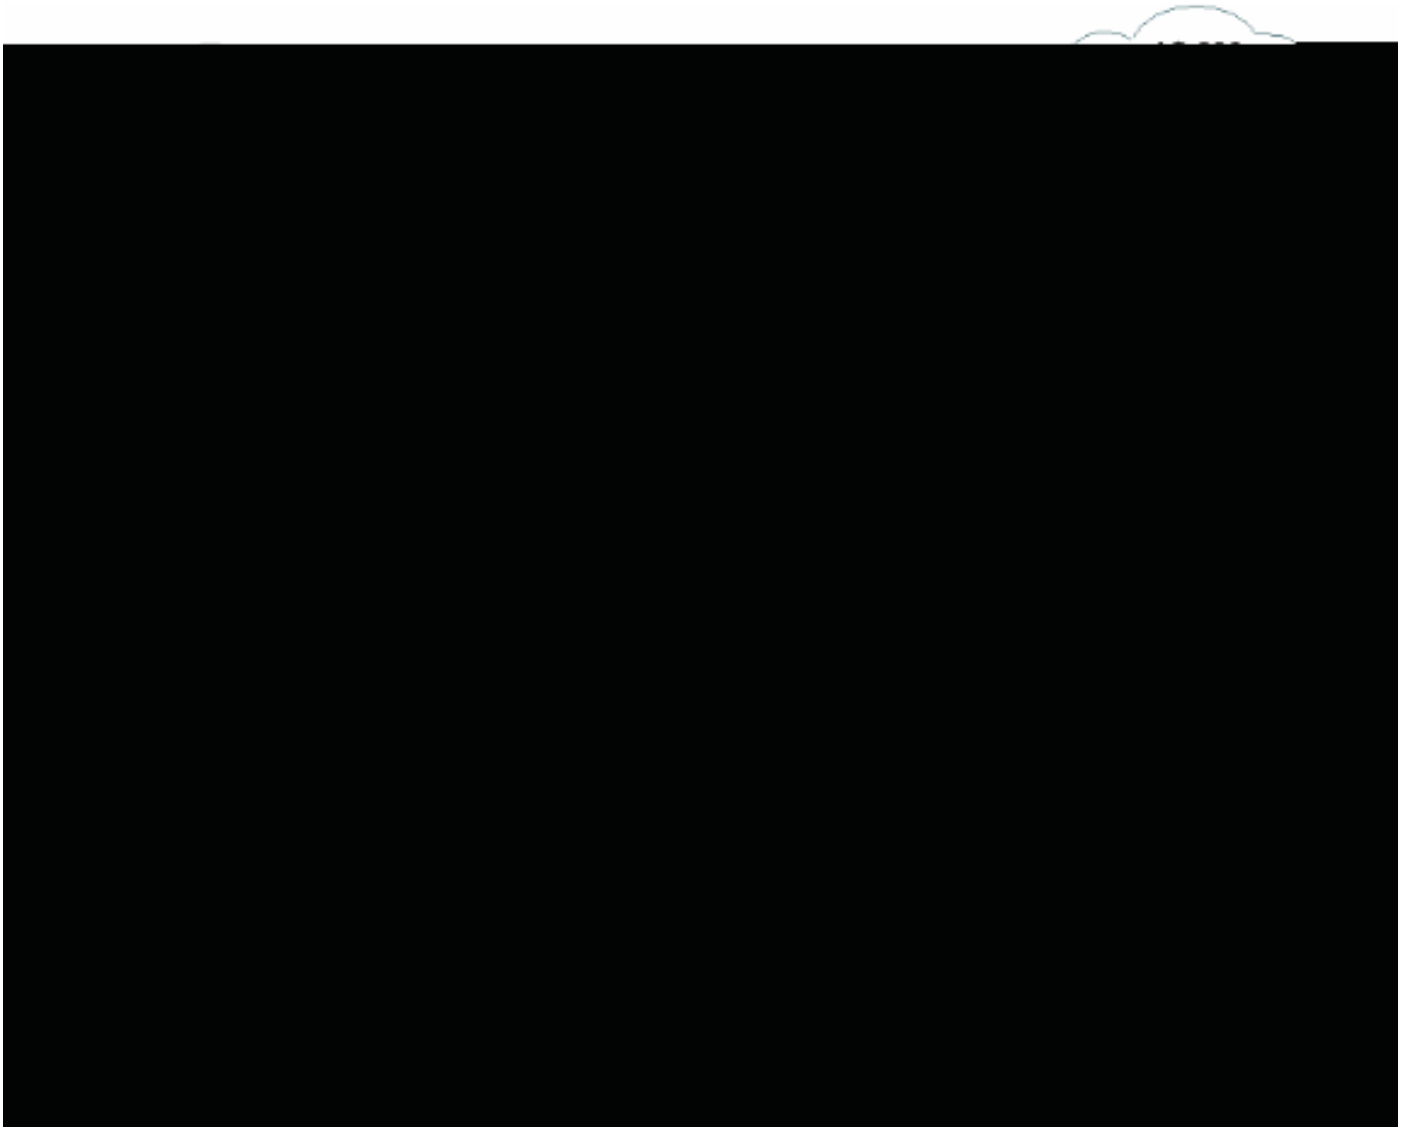

scheme; only the RRs require the upgrade.

Here is the configuration of RTD and RTC:

```
RTD#
```

10bgpampeninBou
10bgpambeninBou 10—C.angethe dahalf-life time.
T is list detailthe dasyntax:
A ndmmas.he at sets all paramer ks athe dasame timetis:
10 seconds. Thhe router keeps the dampeninBoutinformation untilhe daninalty bondmes lesthe an half ofhe da"reuse limit". Athe at point he dauter kepurgethe dainformation.
10bgpampeninBou10—Turns onampeninBou.
Initially,ampeninBoutis off by default. Ifhe dretis a need,he is featuretmay bo given default inablementainhe dafutureThe rse ndmmas. ndntrolauter ampeninBou:

```
interface Serial0
 ip address 203.250.15.2 255.255.255.252

interface Serial1
 ip address 192.208.10.6 255.255.255.252

router bgp 100
 bgp dampening
 network 203.250.15.0
 geao0
```

192.208.10.0 255.255.255.0, version 32 Paths: (1 available, no best path) 300, (suppressed due to dampening) 192.208.10.5 from 192.208.10.5 (192.208.10.174) Origin IGP, metric 0, valid, external Dampinfo: penalty 2615, flapped 3 times in 0:05:18 , reuse in 0:27:00

Practical Design Example

```
RTA#
hostname RTA

ip subnet-zero

interface Loopback0
 ip address 203.250.13.41 255.255.255.0

interface Ethernet0
ip address 203.250.14.1 255.255.255.0

interface Serial0
 ip address 128.213.63.1 255.255.255.252

router ospf 10
 network 203.250.0.0 0.0.255.255 area 0

router bgp 100
 network 203.250.13.0
 network 203.250.14.0
 neighbor 128.213.63.2 remote-as 200
 neighbor 203.250.15.2 remote-as 100
 neighbor 203.250.15.2 update-source Loopback0

RTF#
hostname RTF

ip subnet-zero

interface Ethernet0
 ip address 203.250.14.2 255.255.255.0

interface Serial1
 ip address 203.250.15.1 255.255.255.252

router ospf 10
 network 203.250.0.0 0.0.255.255 area 0

RTB#
hostname RTB

ip subnet-zero

interface Serial0
 ip address 203.250.15.2 255.255.255.252

interface Serial1
 ip address 192.208.10.6 255.255.255.252

router ospf 10
 network 203.250.0.0 0.0.255.255 area 0
```

```
router bgp 100
network 203.250.15.0
 neighbor 192.208.10.5 remote-as 300
 neighbor 203.250.13.41 remote-as 100

RTC#
hostname RTC

ip subnet-zero

interface Loopback0
 ip address 128.213.63.130 255.255.255.192

interface Serial2/0
 ip address 128.213.63.5 255.255.255.252
!
interface Serial2/1
 ip address 128.213.63.2 255.255.255.252

router bgp 200
 network 128.213.0.0
 neighbor 128.213.63.1 remote-as 100
 neighbor 128.213.63.6 remote-as 400

RTD#
hostname RTD

ip subnet-zero

interface Loopback0
ip address 1 0 0 Td 0 d5rg455.255.255.192

interface Serial2/0
 ip address 128.208.10.5 re5.255.255.252
!
interface Serial2/0
 ip address 128.208.10.5 255.255.255.252

router bgp 200
```

```
 ip address 195.211.10.2 255.255.255.252

interface Serial1
 ip address 128.213.63.6 255.255.255.252
 clockrate 1000000

router bgp 400
 network 200.200.10.0
 neighbor 128.213.63.5 remote-as 200
 neighbor 195.211.10.1 remote-as 500

RTG#
hostname RTG

ip subnet-zero

interface Loopback0
 ip address 195.211.10.174 255.255.255.192

interface Serial0
 ip address 192.208.10.1 255.255.255.252

interface Serial1
 ip address 195.211.10.1 255.255.255.252

router bgp 500
 network 195.211.10.0
 neighbor 192.208.10.2 remote-as 300
 neighbor 195.211.10.2 remote-as 400
```

Always use the **network** command or redistribute static entries into BGP to advertise networks. This method is better than a redistribution of IGP into BGP. This example uses the **network** command to inject networks into BGP.

Here, you start with the s1 interface on RTB shutdown, as if the link between RTB and RTD does not exist. This is the RTB BGP table:

```
RTB# show ip bgp BGP table version is 4, local router ID is 203.250.15.2 Status codes: s
suppressed, d damped, h history, * valid, > best, i - internal Origin codes: i - IGP, e - EGP, ?
- incomplete Network Next Hop Metric LocPrf Weight Path *i128.213.0.0 128.213.63.2 0 100 0 200 i
*i192.208.10.0 128.213.63.2 100 0 200 400 500 300 i *i195.211.10.0 128.213.63.2 100 0 200 400
500 i *i200.200.10.0 128.213.63.2 100 0 200 400 i *>i203.250.13.0 203.250.13.41 0 100 0 i
*>i203.250.14.0 203.250.13.41 0 100 0 i *>203.250.15.0 0.0.0.0 0 32768 i
```

In this table, these notations appear:

Note:

because RTA does not synchronize with OSPF because of the difference in masks. Keep

```
 network 192.208.0.0 0.0.255.255 area 0

router bgp 100
 no synchronization
 network 203.250.15.0
 neighbor 192.208.10.5 remote-as 300
 neighbor 203.250.13.41 remote-as 100
```

The BGP tables look like this:

```
RTA# show ip bgp BGP table version is 117, local router ID is 203.250.13.41 Status codes: s
```

```
interface Ethernet0
 ip address 203.250.14.1 255.255.255.0
```

```
 network 203.250.0.0 0.0.255.255 area 0

ip classless

RTB#
hostname RTB

ip subnet-zero

interface Loopback1
 ip address 203.250.15.10 255.255.255.252

interface Serial0
 ip address 203.250.15.2 255.255.255.252
!
interface Serial1

 ip redistriblhd bgp 10 metric 1000bnet-zej0 0 Td 0 -14.4 Td ( )Tj0ilesive-erface Serial1
 ip work 203.250.0.0 0.0.255.255 area 0


!
```

```
access-list 1 permit 195.211.0.0 0.0.255.255
access-list 2 permit any
route-map setcommunity permit 20
 match ip address 2
!
route-map setcommunity permit 10
 match ip address 1
 set community no-export
```

A demonstration of the use of community filtering is on RTG. You add a **no-export** community to 195.211.0.0 updates toward RTD. In this way, RTD does not export that route to RTB. However, in this case, RTB does not accept these routes anyway.

```
RTE#
hostname RTE

ip subnet-zero

interface Loopback0
 ip address 200.200.10.1 255.255.255.0

interface Serial0
 ip address 195.211.10.2 255.255.255.252

interface Serial1
 ip address 128.213.63.6 255.255.255.252

router bgp 400
 network 200.200.10.0
 aggregate-address 200.200.0.0 255.255.0.0 summary-only
 neighbor 128.213.63.5 remote-as 200
 neighbor 195.211.10.1 remote-as 500

ip classless
```

RTE aggregates 200.200.0.0/16. Here are the final BGP and routing tables for RTA, RTF, and RTB:

```
RTA# show ip bgp
```