

## 第九章 BGP

### 9.1. Configuring BGP

提问 在网络中启用 BGPPAN>

回答

Route1 在 AS 65500 中

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#interface Serial0
```

```
Router1(config-if)#ip address 192.168.55.6 255.255.255.252
```

```
Router1(config-if)#exit
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#network 192.168.1.0
```

```
Router1(config-router)#neighbor 192.168.55.5 remote-as 65501
```

```
Router1(config-router)#no synchronization
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

Router2 在 AS 65501 中

```
Router2#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router2(config)#interface Serial0
```

```
Router2(config-if)#ip address 192.168.55.5 255.255.255.252
```

```
Router2(config-if)#exit
```

```
Router2(config)#router bgp 65501
```

```

Router2(config-router)#network 172.25.17.0 mask 255.255.255.0

Router2(config-router)#neighbor 192.168.55.6 remote-as 65500

Router2(config-router)#no synchronization

Router2(config-router)#exit

Router2(config)#end

Router2#

```

注释 在对 BGP 验证的时候比较有用命令是

```

Router1#show ip bgp summary

BGP router identifier 192.168.99.5, local AS number 65500

BGP table version is 7, main routing table version 7

4 network entries and 4 paths using 484 bytes of memory

2 BGP path attribute entries using 196 bytes of memory

BGP activity 11/7 prefixes, 11/7 paths

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.55.5	4	65501	17	18	7	0	0	00:12:38	2
172.25.2.2	4	65531	527	526	0	0	0	21:05:23	Active

Router1#

需要注意的是理想状态是 State 里面是数字，尽管是 Active 也不代表是配置正常，反而是配置出现错误。通过 neighbor 172.20.1.2 update-source Loopback0 命令来限制 BGP 数据包源地址为回环地址，但要确保此地址的连通性

## 9.2. 使用 eBGP Multihop

提问 配置外部 BGP，但是不是直连的路由器

回答

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#ip route 172.20.1.2 255.255.255.255 192.168.1.5 2
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#neighbor 172.20.1.2 remote-as 65530
```

```
Router1(config-router)#neighbor 172.20.1.2 update-source Loopback0
```

```
Router1(config-router)#neighbor 172.20.1.2 ebgp-multihop 3
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

注释 缺省情况下 eBGP 的路由器必须是直连的，如果不是直连的就需要使用此命令。一种说法是此跳数越小越好，但是 RFC 3682 说为了安全还是越大越好，思科在 12.3(7)T 后也采用了这个建议，使用了 neighbor 192.168.55.5 ttl-security hops 1 命令，此命令会丢弃所有 TTL 小于  $255 - 1 = 254$  的 BGP 数据包，这时候如果对端 eBGP 邻居不支持此特性就必须使用下面的命令来配置 neighbor 192.168.55.6 ebgp-multihop 255

### 9.3. 调整 Next-Hop 属性值

提问 在 iBGP 之间宣告路由时候修改下一跳属性值，使其指向内部 AS 的地址

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#neighbor 192.168.1.6 remote-as 65500
```

```
Router1(config-router)#neighbor 192.168.1.6 next-hop-self
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

**注释** 正常情况下 iBGP 之间下一跳属性值是不会修改的，只会在 eBGP 时会进行修改，而此地址会指向 eBGP 邻居的地址，而往往内部 AS 的路由器没有到达此地址的路由。

#### 9.4. 连接两个 ISPs

**提问** 一台路由器连接两个 ISP，保证网络冗余

**回答**

```
Router1#configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

Router1(config)#interface Serial0

Router1(config-if)#description connection to ISP #1, ASN 65510

Router1(config-if)#ip address 192.168.1.6 255.255.255.252

Router1(config-if)#exit

Router1(config)#interface Serial1

Router1(config-if)#description connection to ISP #2, ASN 65520

Router1(config-if)#ip address 192.168.2.6 255.255.255.252

Router1(config-if)#exit

Router1(config)#interface Ethernet0

Router1(config-if)#description connection to internal network, ASN 65500

Router1(config-if)#ip address 172.18.5.2 255.255.255.0

Router1(config-if)#exit

Router1(config)#router bgp 65500

Router1(config-router)#network 172.18.5.0 mask 255.255.255.0

Router1(config-router)#neighbor 192.168.1.5 remote-as 65510

Router1(config-router)#neighbor 192.168.2.5 remote-as 65520

Router1(config-router)#no synchronization
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

注释 注意此配置不是最佳配置，可能导致内部 AS 称为两个 ISP 的 transit AS，同时导致自己路由器接收过多路由

### 9.5. 两台路由器分别连接两个 ISP

提问 内部 AS 有两台路由器，分别连两个 ISP 保证网络冗余

回答

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#interface Serial0
```

```
Router1(config-if)#description connection to ISP #1, ASN 65510
```

```
Router1(config-if)#ip address 192.168.1.6 255.255.255.252
```

```
Router1(config-if)#exit
```

```
Router1(config)#interface Ethernet0
```

```
Router1(config-if)#description connection to internal network, ASN 65500
```

```
Router1(config-if)#ip address 172.18.5.2 255.255.255.0
```

```
Router1(config-if)#exit
```

```
Router1(config)#ip as-path access-list 15 permit ^$
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#network 172.18.5.0 mask 255.255.255.0
```

```
Router1(config-router)#neighbor 172.18.5.3 remote-as 65500
```

```
Router1(config-router)#neighbor 172.18.5.3 next-hop-self
```

```
Router1(config-router)#neighbor 192.168.1.5 remote-as 65510
```

```
Router1(config-router)#neighbor 192.168.1.5 filter-list 15 out

Router1(config-router)#no synchronization

Router1(config-router)#exit

Router1(config)#end

Router1#  
  
Router2#configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

Router2(config)#interface Serial1

Router2(config-if)#description connection to ISP #2, ASN 65520

Router2(config-if)#ip address 192.168.2.6 255.255.255.252

Router2(config-if)#exit

Router2(config)#interface Ethernet0

Router2(config-if)#description connection to internal network, ASN 65500

Router2(config-if)#ip address 172.18.5.3 255.255.255.0

Router2(config-if)#exit

Router2(config)#ip as-path access-list

Router2(config)#router bgp 65500

Router2(config-router)#network 172.18.5.0 mask 255.255.255.0

Router2(config-router)#neighbor 192.168.2.5 remote-as 65520

Router2(config-router)#neighbor 192.168.2.5 filter-list 15 out

Router2(config-router)#neighbor 172.18.5.2 remote-as 65500

Router2(config-router)#neighbor 172.18.5.2 next-hop-self

Router2(config-router)#no synchronization

Router2(config-router)#exit
```

```
Router2(config)#end
```

```
Router2#
```

注释

#### 9.6. 限制向 BGP 对端的网络宣告

提问 限制特定的路由公告给对端的 AS

回答

有三种方法，第一种是扩展 ACL

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#access-list 105 deny ip host 172.25.0.0 host 255.255.0.0
```

```
Router1(config)#access-list 105 permit ip any any
```

```
Router1(config)#route-map ACL-RT-FILTER permit 10
```

```
Router1(config-route-map)#match ip address 105
```

```
Router1(config-route-map)#exit
```

```
Router1(config)#route-map ACL-RT-FILTER deny 20
```

```
Router1(config-route-map)#exit
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#neighbor 192.168.1.5 remote-as 65510
```

```
Router1(config-router)#neighbor 192.168.1.5 route-map ACL-RT-FILTER in
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

第二种是使用 distribute-list:

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#access-list 106 deny ip host 172.25.0.0 host 255.255.0.0
```

```
Router1(config)#access-list 106 permit ip any any
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#neighbor 192.168.1.5 remote-as 65510
```

```
Router1(config-router)#neighbor 192.168.1.5 distribute-list 106 in
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

第三种也是最常用的是使用 prefix lists

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#ip prefix-list PREFIX-FILTER seq 10 deny 172.25.0.0/16
```

```
Router1(config)#ip prefix-list PREFIX-FILTER seq 20 permit 0.0.0.0/0 le 32
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#neighbor 192.168.1.5 remote-as 65510
```

```
Router1(config-router)#neighbor 192.168.1.5 prefix-list PREFIX-FILTER in
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

注释 前两种使用的扩展 ACL 比较奇特，第一个 host 是子网，第二个 host 是子网掩码，而不是传统目的地址，所以 host 172.25.0.0 host 255.255.0.0 就代表网络 172.25.0.0/16，如果用正常的 ACL 就实现不了对无类网络的控制。所以推荐使用第三种方式 prefixlist，此列表支持序列号，可以帮助你修改和插入新的条目 ge 是大于，le 是小于，控制子网掩码 permit 0.0.0.0/0 le 32 就是变相的 permit any

## 9.7. 调整 Local Preference 属性值

提问 调整 Local Preference 属性值来控制路由选择

回答

第一种全局

```
Router1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#bgp default local-preference 200
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

第二种使用 route map 控制

```
Router1#configure terminal
```

```
Enter c
```

```
onfiguration commands, one per line. End with CNTL/Z.
```

```
Router1(config)#ip prefix-list LOW_LP_PREFIXES seq 10 permit 172.22.0.0/16
```

```
Router1(config)#route-map LOCALPREF permit 10
```

```
Router1(config-route-map)#match ip address prefix-list LOW_LP_PREFIXES
```

```
Router1(config-route-map)#set local-preference 50
```

```
Router1(config-route-map)#exit
```

```
Router1(config)#route-map LOCALPREF permit 20
```

```
Router1(config-route-map)#exit
```

```
Router1(config)#router bgp 65500
```

```
Router1(config-router)#neighbor 192.168.1.5 remote-as 65510
```

```
Router1(config-router)#neighbor 192.168.1.5 route-map LOCALPREF in  
Router1(config-router)#exit  
Router1(config)#end  
Router1#
```

注释 此 local preference 属性值只在内部 AS 有用，选路级别高于 AS Path。此值越大优先级越高，缺省值为 100。Show ip bgp 命令可以看到各个路由的 local preference 属性值

#### 9.8. 负载均衡

提问 在 BGP 邻居之间的多链路上负载均衡流量

回答

```
Router1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Router1(config)#router bgp 65500  
Router1(config-router)#maximum-paths 4
```

```
Router1(config-router)#exit
```

```
Router1(config)#end
```

```
Router1#
```

注释 正常情况下 BGP 选路策略会保证只有一条路径，通过此命令可以增加到 4 条，不过要确保所有属性值相同，包括 MED 属性。同时注意此负载均衡只针对出流量而不适合入流量

#### 9.9. 在 AS Path 属性值中清除私有 ASNs

提问 避免内网中的私有 ASN 传播到互联网

回答

```
Router1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Router1(config)#interface Serial0  
Router1(config-if)#description connection to ISP #1, ASN 1
```

```
Router1(config-if)#ip address 192.168.1.6 255.255.255.252

Router1(config-if)#exit

Router1(config)#interface Serial1

Router1(config-if)#description connection to private network, ASN 65500

Router1(config-if)#ip address 192.168.5.1 255.255.255.252

Router1(config-if)#exit

Router1(config)#router bgp 2

Router1(config-router)#neighbor 192.168.5.2 remote-as 65500

Router1(config-router)#neighbor 192.168.1.5 remote-as 1

Router1(config-router)#neighbor 192.168.1.5 remove-private-AS

Router1(config-router)#no synchronization

Router1(config-router)#exit

Router1(config)#end

Router1#
```

注释 注意此命令是不能删除那些在公共 ASN 之间的私有 ASN

#### 9.18. 使用 BGP Route Reflectors

提问 通过路由反射器的方式来简化 iBGP 邻居关系

回答

只要针对三种不同角色路由器的配置

Router1 是 Client Peer:

```
Router1#configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
Router1(config)#interface Ethernet0/0
```

```
Router1(config-if)#ip address 172.18.5.2 255.255.255.0
```

```
Router1(config-if)#exit

Router1(config)#interface Serial0/0

Router1(config-if)#ip address 192.168.1.6 255.255.255.252

Router1(config-if)#exit

Router1(config)#interface Loopback0

Router1(config-if)#ip address 172.18.6.1 255.255.255.255

Router1(config-if)#exit

Router1(config)#router bgp 65500

Router1(config-router)#no synchronization

Router1(config-router)#neighbor 172.18.6.2 remote-as 65500

Router1(config-router)#neighbor 172.18.6.2 next-hop-self

Router1(config-router)#neighbor 172.18.6.2 update-source Loopback0
Router1(config-router)#neighbor 192.168.1.5 remote-as 65510

Router1(config-router)#exit

Router1(config)#ip route 172.18.6.2 255.255.255.255 172.18.5.3

Router1(config)#ip route 172.18.6.3 255.255.255.255 172.18.5.4

Router1(config)#ip route 172.18.6.4 255.255.255.255 172.18.5.10

Router1(config)#end

Router1#  
  
Router4 是 Nonclient Peer:  
  
Router4#configure terminal  
  
Enter configuration commands, one per line. End with CNTL/Z.  
  
Router4(config)#interface Ethernet0  
  
Router4(config-if)#ip address 172.18.5.10 255.255.255.0
```

```
Router4(config-if)#exit

Router4(config)#interface Loopback0

Router4(config-if)#ip address 172.18.6.4 255.255.255.255

Router4(config-if)#exit

Router4(config)#router bgp 65500

Router4(config-router)#no synchronization

Router4(config-router)#neighbor 172.18.6.2 remote-as 65500

Router4(config-router)#neighbor 172.18.6.2 update-source Loopback0

Router4(config-router)#exit

Router4(config)#ip route 172.18.6.1 255.255.255.255 172.18.5.2

Router4(config)#ip route 172.18.6.2 255.255.255.255 172.18.5.3

Router4(config)#ip route 172.18.6.3 255.255.255.255 172.18.5.4

Router4(config)#end

Router4#  
  
R2 是 Route Reflector  
  
Router2#configure terminal  
  
Enter configuration commands, one per line. End with CNTL/Z.  
  
Router2(config)#interface FastEthernet0/0  
  
Router2(config-if)#ip address 172.18.5.3 255.255.255.0  
  
Router2(config-if)#exit  
  
Router2(config)#interface Loopback0  
  
Router2(config-if)#ip address 172.18.6.2 255.255.255.255  
  
Router2(config-if)#exit  
  
Router2(config)#router bgp 65500
```

```
Router2(config-router)#no synchronization

Router2(config-router)#neighbor 172.18.6.1 remote-as 65500

Router2(config-router)#neighbor 172.18.6.1 route-reflector-client

Router2(config-router)#neighbor 172.18.6.1 update-source Loopback0

Router2(config-router)#neighbor 172.18.6.3 remote-as 65500

Router2(config-router)#neighbor 172.18.6.3 route-reflector-client

Router2(config-router)#neighbor 172.18.6.3 update-source Loopback0

Router2(config-router)#neighbor 172.18.6.4 remote-as 65500

Router2(config-router)#neighbor 172.18.6.4 update-source Loopback0

Router2(config-router)#no auto-summary

Router2(config-router)#exit

Router2(config)#ip route 172.18.6.1 255.255.255.255 172.18.5.2

Router2(config)#ip route 172.18.6.3 255.255.255.255 172.18.5.4

Router2(config)#ip route 172.18.6.4 255.255.255.255 172.18.5.10

Router2(config)#end

Router2#
```

注释 路由反射器是解决要求 iBGP 全互联的问题。不过为了保证冗余性还是要配置多个路由反射器，使用 bgp cluster-id 1234 命令来定义 cluster

<!--[if !supportLists]-->9.19. <!--[endif]-->汇总实验

提问 结合前面的方法，重新配置一台路由器两个冗余链路的情况

回答

```
Router1#configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
Router1(config)#interface Serial0

Router1(config-if)#description connection to ISP #1, ASN 65510

Router1(config-if)#ip address 192.168.1.6 255.255.255.252

Router1(config-if)#exit

Router1(config)#interface Serial1

Router1(config-if)#description connection to ISP #2, ASN 65520

Router1(config-if)#ip address 192.168.2.6 255.255.255.252

Router1(config-if)#exit

Router1(config)#interface Ethernet0

Router1(config-if)#description connection to internal network, ASN 65500

Router1(config-if)#ip address 172.18.5.2 255.255.255.0

Router1(config-if)#exit

Router1(config)#ip as-path access-list 15 permit ^$

Router1(config)#ip route 0.0.0.0 0.0.0.0 192.168.101.0 1

Router1(config)#ip route 0.0.0.0 0.0.0.0 192.168.102.0 2

Router1(config)#ip prefix-list CREATE-DEFAULT seq 10 permit 192.168.101.0/24

Router1(config)#ip prefix-list CREATE-DEFAULT seq 20 permit 192.168.102.0/24

Router1(config)#ip prefix-list BLOCK-DEFAULT seq 10 permit 0.0.0.0/0 ge 1

Router1(config)#route-map PREPEND permit 10

Router1(config-route-map)#set as-path prepend 65500 65500

Router1(config-route-map)#exit

Router1(config)#route-map LOCALPREF permit 10

Router1(config-route-map)#set local-preference 75

Router1(config-route-map)#exit
```

```
Router1(config)#route-map DEFAULT-ROUTE permit 10  
  
Router1(config-route-map)#match ip address prefix-list CREATE-DEFAULT  
  
Router1(config-route-map)#exit  
  
Router1(config)#router bgp 65500  
  
Router1(config-router)#network 172.18.5.0 mask 255.255.255.0  
  
Router1(config-router)#neighbor 172.18.5.3 remote-as 65500  
  
Router1(config-router)#neighbor 172.18.5.3 password password_number1  
  
Router1(config-router)#neighbor 172.18.5.3 default-originate route-map DEFAULT-ROUTE  
  
Router1(config-router)#neighbor 192.168.1.5 remote-as 65510  
  
Router1(config-router)#neighbor 192.168.1.5 password password_number2  
  
Router1(config-router)#neighbor 192.168.1.5 filter-list 15 out  
  
Router1(config-router)#neighbor 192.168.1.5 prefix-list CREATE-DEFAULT in  
  
Router1(config-router)#neighbor 192.168.1.5 prefix-list BLOCK-DEFAULT out  
  
Router1(config-router)#neighbor 192.168.2.5 remote-as 65520  
  
Router1(config-router)#neighbor 192.168.2.5 password password_number3  
  
Router1(config-router)#neighbor 192.168.2.5 filter-list 15 out  
  
Router1(config-router)#neighbor 192.168.2.5 prefix-list CREATE-DEFAULT in  
  
Router1(config-router)#neighbor 192.168.2.5 prefix-list BLOCK-DEFAULT out  
  
Router1(config-router)#neighbor 192.168.2.5 route-map PREPEND out  
  
Router1(config-router)#neighbor 192.168.2.5 route-map LOCALPREF in  
  
Router1(config-router)#no synchronization  
  
Router1(config-router)#exit  
  
Router1(config)#end  
  
Router1#
```

