BGP实施使用32位AS号码配置示例

目录

简介

使用32位AS编号，本文描述如何配置边界网关协议(BGP)。在BGP中，每路由域是单个管理域并且有一个唯一AS编号分配到它和在统一内操作设置路由策略。它也维护域间路由。

在本文中， BGP并列配置在16位和32位发言的BGP路由器之间。新的32位AS模式是与the16-bit AS模式兼容。在32位模式能经营的BGP对等体在新的模式确实地回应给新的功能和该会话经营。另一方面， 32位BGP对等体，当通信用16位BGP扬声器， 16位发言路由器时忽略此新建的功能并且操作他们的16位模式的BGP会话。

先决条件

要求

思科建议您有BGP基础知识。

硬件与软件版本

本文中的配置基于装有 Cisco IOS® 软件版本 15.0(1) 的 Cisco 7200 系列路由器。

规则

有关文档规则的详细信息，请参阅 Cisco 技术提示规则。

配置

使用16位AS模式，在本例中，路由器R1和R3配置是AS 100形成的iBGP关系。使用32位AS模式，路由器R2和R4是配置的AS 10.1，并且形成IBGP同位体。路由器R1和R2运行和IGP协议，在此示例OSPF在彼此之间并且形成eBGP相邻在他们之间。

注意： 使用命令查找工具（仅限注册用户）查找有关本文档所使用命令的详细信息。

网络图

本文档使用以下网络设置：

配置

本文档使用以下配置：

- 路由器 R1
- 路由器 R2
- 路由器 R3
- 路由器R4

路由器 R1

```
R1#show run
Building configuration...
!
version 15.0
!
hostname R1
!
ip cef
!
interface Loopback0
ip address 1.1.1.1 255.255.255.255
!
interface Loopback10
ip address 192.168.100.1 255.255.255.0
!
interface Loopback20
ip address 192.168.200.1 255.255.255.0
!
interface FastEthernet1/0
ip address 192.168.10.1 255.255.255.0
duplex auto
 speed auto
!
interface Serial2/0
 ip address 10.10.100.1 255.255.255.0
 serial restart-delay 0
!
router ospf 1
 log-adjacency-changes
 network 1.1.1.1 0.0.0.0 area 0
 network 10.10.100.0 0.0.0.255 area 0
!
router bgp 100 !--- BGP is configured using 16-bit AS
number no synchronization bgp router-id 10.10.10.10 bgp
asnotation dot !--- This command change the default
asplain notation to dot notation. !--- Note that without
```

```
this command the AS number will treated as asplain
notation i.e. 10.1 will be displayed as 655361 bgp log-
neighbor-changes network 192.168.100.0 network
192.168.200.0 neighbor 2.2.2.2 remote-as 10.1 !--- The
AS number of the eBGP peer in 32-bit neighbor 2.2.2.2
ebgp-multihop 255 neighbor 2.2.2.2 update-source
Loopback0 neighbor 192.168.10.2 remote-as 100 neighbor
192.168.10.2 next-hop-self no auto-summary ! end
```

路由器 R2

```
R2#show run
!
version 15.0
!
hostname R2
!
ip cef
!
interface Loopback0
ip address 2.2.2.2 255.255.255.0
!
interface Loopback10
ip address 10.1.1.1 255.255.255.255
!
interface Loopback20
ip address 20.1.1.1 255.255.255.255
!
interface FastEthernet1/0
ip address 172.16.10.1 255.255.255.0
duplex auto
speed auto
!
interface Serial2/0
ip address 10.10.100.2 255.255.255.0
serial restart-delay 0
!
!
router ospf 1
 log-adjacency-changes
 network 2.2.2.2 0.0.0.0 area 0
 network 10.10.100.0 0.0.0.255 area 0
!
router bgp 10.1 !--- BGP is configured using 32-bit AS
number no synchronization bgp router-id 20.20.20.20 bgp
asnotation dot bgp log-neighbor-changes network 10.1.1.1
mask 255.255.255.255 network 20.1.1.1 mask
255.255.255.255 neighbor 1.1.1.1 remote-as 100 neighbor
1.1.1.1 ebgp-multihop 255 neighbor 1.1.1.1 update-source
Loopback0 neighbor 172.16.10.2 remote-as 10.1 neighbor
172.16.10.2 next-hop-self no auto-summary ! end
```

路由器 R3

```
R3#show run
Building configuration...
!
version 15.0
ip cef
!
interface Loopback0
 ip address 30.30.30.30 255.255.255.255
!
interface FastEthernet1/0
```

```
 ip address 192.168.10.2 255.255.255.0
 duplex auto
 speed auto
!
router bgp 100 no synchronization bgp router-id 3.3.3.3
bgp log-neighbor-changes network 30.30.30.30 mask
255.255.255.255 neighbor 192.168.10.1 remote-as 100
neighbor 192.168.10.1 next-hop-self no auto-summary !---
iBGP peering is formed between routers R1 and R3 using
16-bit AS number. ! end
```

路由器R4

```
R4#show run
Building configuration...
!
version 15.0
ip cef
!
interface Loopback0
 ip address 40.40.40.40 255.255.255.255
!
interface FastEthernet1/0
 ip address 172.16.10.2 255.255.255.0
 duplex auto
 speed auto
!
router bgp 10.1 no synchronization bgp router-id 4.4.4.4
bgp asnotation dot bgp log-neighbor-changes network
40.40.40.40 mask 255.255.255.255 neighbor 172.16.10.1
remote-as 10.1 no auto-summary ! end !--- iBGP peering
is formed between routers R2 and R4 using 32-bit AS
number.
```

验证

使用本部分可确认配置能否正常运行。

命令输出解释程序（仅限注册用户）(OIT) 支持某些 **show** 命令。使用 OIT 可查看对 show 命令输出的分析。

显示命令

为了验证BGP可以支持32位ASN，请使用show ip bgp neighbor命令。

show ip bgp neighbor

在路由器 R1 中

```
R1#show ip bgp neighbor 2.2.2.2 BGP neighbor is 2.2.2.2,
remote AS 10.1, external link BGP version 4, remote
router ID 20.20.20.20 BGP state = Established, up for
03:28:22 Last read 00:00:41, last write 00:00:29, hold
time is 180, keepalive interval is 60 seconds Neighbor
sessions: 1 active, is multisession capable Neighbor
capabilities: Route refresh: advertised and
received(new) Four-octets ASN Capability: advertised and
received Address family IPv4 Unicast: advertised and
received Multisession Capability: advertised and
received Message statistics, state Established: InQ
depth is 0 OutQ depth is 0 Sent Rcvd Opens: 1 1
Notifications: 0 0 Updates: 3 3 Keepalives: 229 230
```

```
Route Refresh: 0 0 Total: 233 234 !--- Output omitted---
!
```

要显示条目在BGP路由表里，请使用show ip bgp命令。

show ip bgp

在路由器 R1 中

```
R1#sh ip bgp
BGP table version is 13, local router ID is 10.10.10.10
Status codes: s suppressed, d damped, h history, *
valid, > best, I - internal,
              r RIB-failure, S Stale
Origin codes: I - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf
Weight Path
*> 10.1.1.1/32 2.2.2.2 0 0 10.1 I *> 20.1.1.1/32 2.2.2.2
0 0 10.1 I *>i30.30.30.30/32 192.168.10.2 0 100 0 I *>
40.40.40.40/32 2.2.2.2 0 10.1 I *> 192.168.100.0 0.0.0.0
0 32768 I *> 192.168.200.0 0.0.0.0 0 32768 I !--- Note
that the routes highlighted are received from the eBGP
peer router R2 which is in 32-bit AS 10.1. In router R3
R3#sh ip bgp
BGP table version is 11, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, *
valid, > best, I - internal,
              r RIB-failure, S Stale
Origin codes: I - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf
Weight Path
*>i10.1.1.1/32 192.168.10.1 0 100 0 655361 I
*>i20.1.1.1/32 192.168.10.1 0 100 0 655361 I *>
30.30.30.30/32 0.0.0.0 0 32768 I *>i40.40.40.40/32
192.168.10.1 0 100 0 655361 I *>i192.168.100.0
192.168.10.1 0 100 0 I *>i192.168.200.0 192.168.10.1 0
100 0 I !--- The router R3 does not have bgp asnotation
dot configured in it. Therefore, the route received from
the router in 32-bit AS AS 10.1 is displayed as 655361.
In router R4 R4#sh ip bgp
BGP table version is 7, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, *
valid, > best, I - internal,
              r RIB-failure, S Stale
Origin codes: I - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf
Weight Path
*>i10.1.1.1/32      172.16.10.1              0    100
0 I
*>i20.1.1.1/32      172.16.10.1              0    100
0 I
*>i30.30.30.30/32   172.16.10.1              0    100
0 100 I
*> 40.40.40.40/32   0.0.0.0                  0
32768 I
*>i192.168.100.0    172.16.10.1              0    100
0 100 I
*>i192.168.200.0    172.16.10.1              0    100
0 100 I

 !--- The above output shows the entries in BGP routing
```

```
table of router R4.
```

为了验证在路由器之间的可接通性，请使用ping命令。

ping

从路由器R3

```
R3#ping 40.40.40.40 Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 40.40.40.40, timeout
is 2 seconds: !!!!! Success rate is 100 percent (5/5),
round-trip min/avg/max = 68/101/148 ms
```

从路由器R4

```
R4#ping 30.30.30.30 Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.30.30.30, timeout
is 2 seconds: !!!!! Success rate is 100 percent (5/5),
round-trip min/avg/max = 56/89/112 ms !--- The above
output shows that End to End connectivity is established
between R3 and R4, where R3 is AS 100(16-bit AS) and
router R4 is in AS 10.1(32-bit AS).
```

相关信息

- Cisco IOS BGP 4字节ASN支持
- BGP 支持页
- BGP 案例分析
- 测试的自治系统编号
- 技术支持和文档 - Cisco Systems