Cisco IOS eBGP (External) Protocol

Network Topology

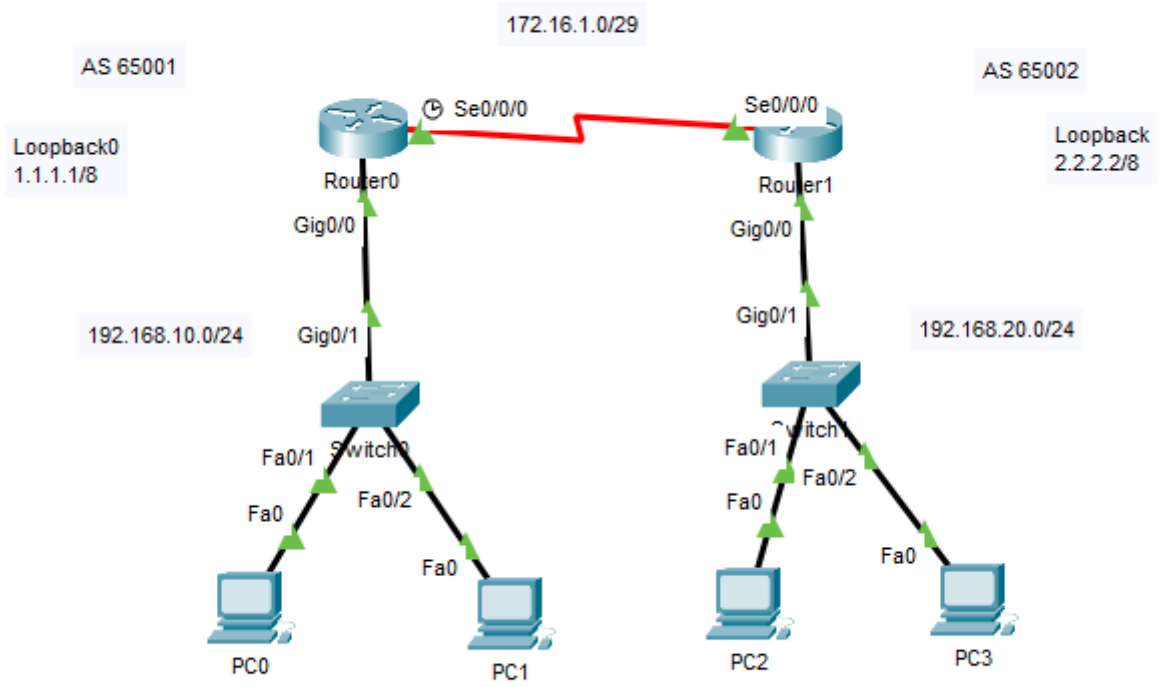

As of Cisco Packet Tracer 8.2 internal BGP (iBGP) is not supported only external BGP (eBGP) is supported. If you try and configure iBGP on Packet Tracer, you will see the following message.

```
%Cisco Packet Tracer does not support internal BGP in this version. Only external
neighbors are supported.
```

So, in this example the focus will be on setting up eBGP between two separate autonomous systems.

Router0 Configuration Commands

Router0>enable
Router0#configure terminal
Router0(config)#interface serial 0/0/0

```
Router0(config-if)#ip address 172.16.1.1 255.255.255.248
Router0(config-if)#no shutdown
Router0(config-if)#interface gigabitEthernet 0/0
Router0(config-if)#ip address 192.168.10.1 255.255.255.0
Router0(config-if)#no shutdown
Router0(config-if)#int Loopback 0
Router0(config-if)#ip address 1.1.1.1 255.0.0.0
Router0(config-if)#no shutdown
Router0(config-if)#exit
Router0#router bgp 65001
Router0(config-router)#neighbor 172.16.1.2 remote-as 65002
Router0(config-router)#network 1.1.1.1 mask 255.0.0.0
Router0(config-router)#network 192.168.10.0 mask 255.255.255.0
Router0(config-router)#bgp router-id 1.1.1.1
Router0(config-router)#no synchronization
```

Router1 Configuration Commands

```
Router1>enable
Router1#configure terminal
Router1(config)#interface serial 0/0/0
Router1(config-if)#ip address 172.16.1.2 255.255.255.248
Router1(config-if)#no shutdown
Router1(config-if)#interface gigabitEthernet 0/0
Router1(config-if)#ip address 192.168.20.1 255.255.255.0
Router1(config-if)#no shutdown
Router1(config-if)#int Loopback 0
Router1(config-if)#ip address 2.2.2.2 255.0.0.0
Router1(config-if)#no shutdown
Router1(config-if)#exit
Router1#router bgp 65001
Router1(config-router)#neighbor 172.16.1.1 remote-as 65001
Router1(config-router)#network 2.2.2.2 mask 255.0.0.0
Router1(config-router)#network 192.168.20.0 mask 255.255.255.0
Router1(config-router)#bgp router-id 2.2.2.2
Router1(config-router)#no synchronization
```

A Special Word about the No Synchronization Command

The site BGPExpert has an excellent explanation of what the *no synchronization* command is and what is does.

Check BGP Configuration

```
Router0#show ip route
```

```
Router0#sho ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       1.0.0.0/8 is directly connected, Loopback0
L       1.1.1.1/32 is directly connected, Loopback0
B    2.0.0.0/8 [20/0] via 172.16.1.2, 00:00:00
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C       172.16.1.0/29 is directly connected, Serial0/0/0
L       172.16.1.1/32 is directly connected, Serial0/0/0
     192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.10.0/24 is directly connected, GigabitEthernet0/0
L       192.168.10.1/32 is directly connected, GigabitEthernet0/0
B    192.168.20.0/24 [20/0] via 172.16.1.2, 00:00:00
```

Router0#show ip bgp summary

```
Router0#show ip bgp summary
BGP router identifier 1.1.1.1, local AS number 65001
BGP table version is 7, main routing table version 6
4 network entries using 528 bytes of memory
4 path entries using 208 bytes of memory
2/2 BGP path/bestpath attribute entries using 368 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 1 (at peak 1) using 32 bytes of memory
BGP using 1184 total bytes of memory
BGP activity 4/0 prefixes, 4/0 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.1.2      4 65002     883     879        7    0    0 08:02:30        4
```

Router0#show ip bgp

```
Router0#show ip bgp
BGP table version is 7, local router ID is 1.1.1.1
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
*> 1.0.0.0/8         0.0.0.0                  0         0 32768 i
*> 2.0.0.0/8         172.16.1.2               0         0     0 65002 i
*> 192.168.10.0/24   0.0.0.0                  0         0 32768 i
*> 192.168.20.0/24   172.16.1.2               0         0     0 65002 i
```

Router0#show bgp ipv4 unicast neighbor 172.16.1.2 advertised-routes

This is another command that is not available in Cisco Packet Tracer as of v8.2. So, here is an example from GNS3.

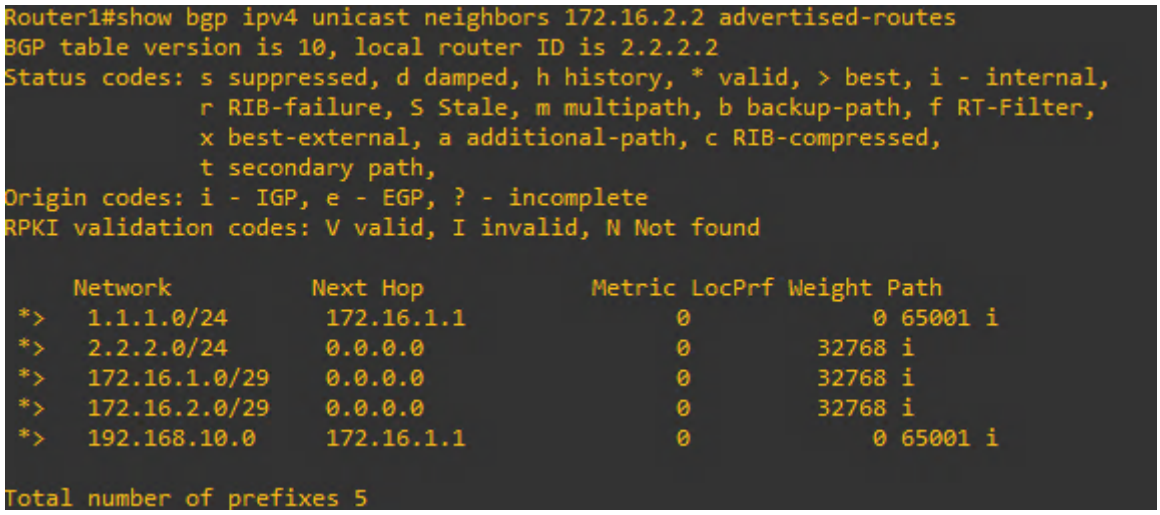

```
Router1#show bgp ipv4 unicast neighbors 172.16.2.2 advertised-routes
BGP table version is 10, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>   1.1.1.0/24       172.16.1.1               0             0 65001 i
 *>   2.2.2.0/24       0.0.0.0                  0         32768 i
 *>   172.16.1.0/29    0.0.0.0                  0         32768 i
 *>   172.16.2.0/29    0.0.0.0                  0         32768 i
 *>   192.168.10.0     172.16.1.1               0             0 65001 i

Total number of prefixes 5
```

Ping and Tracert Connectivity Tests

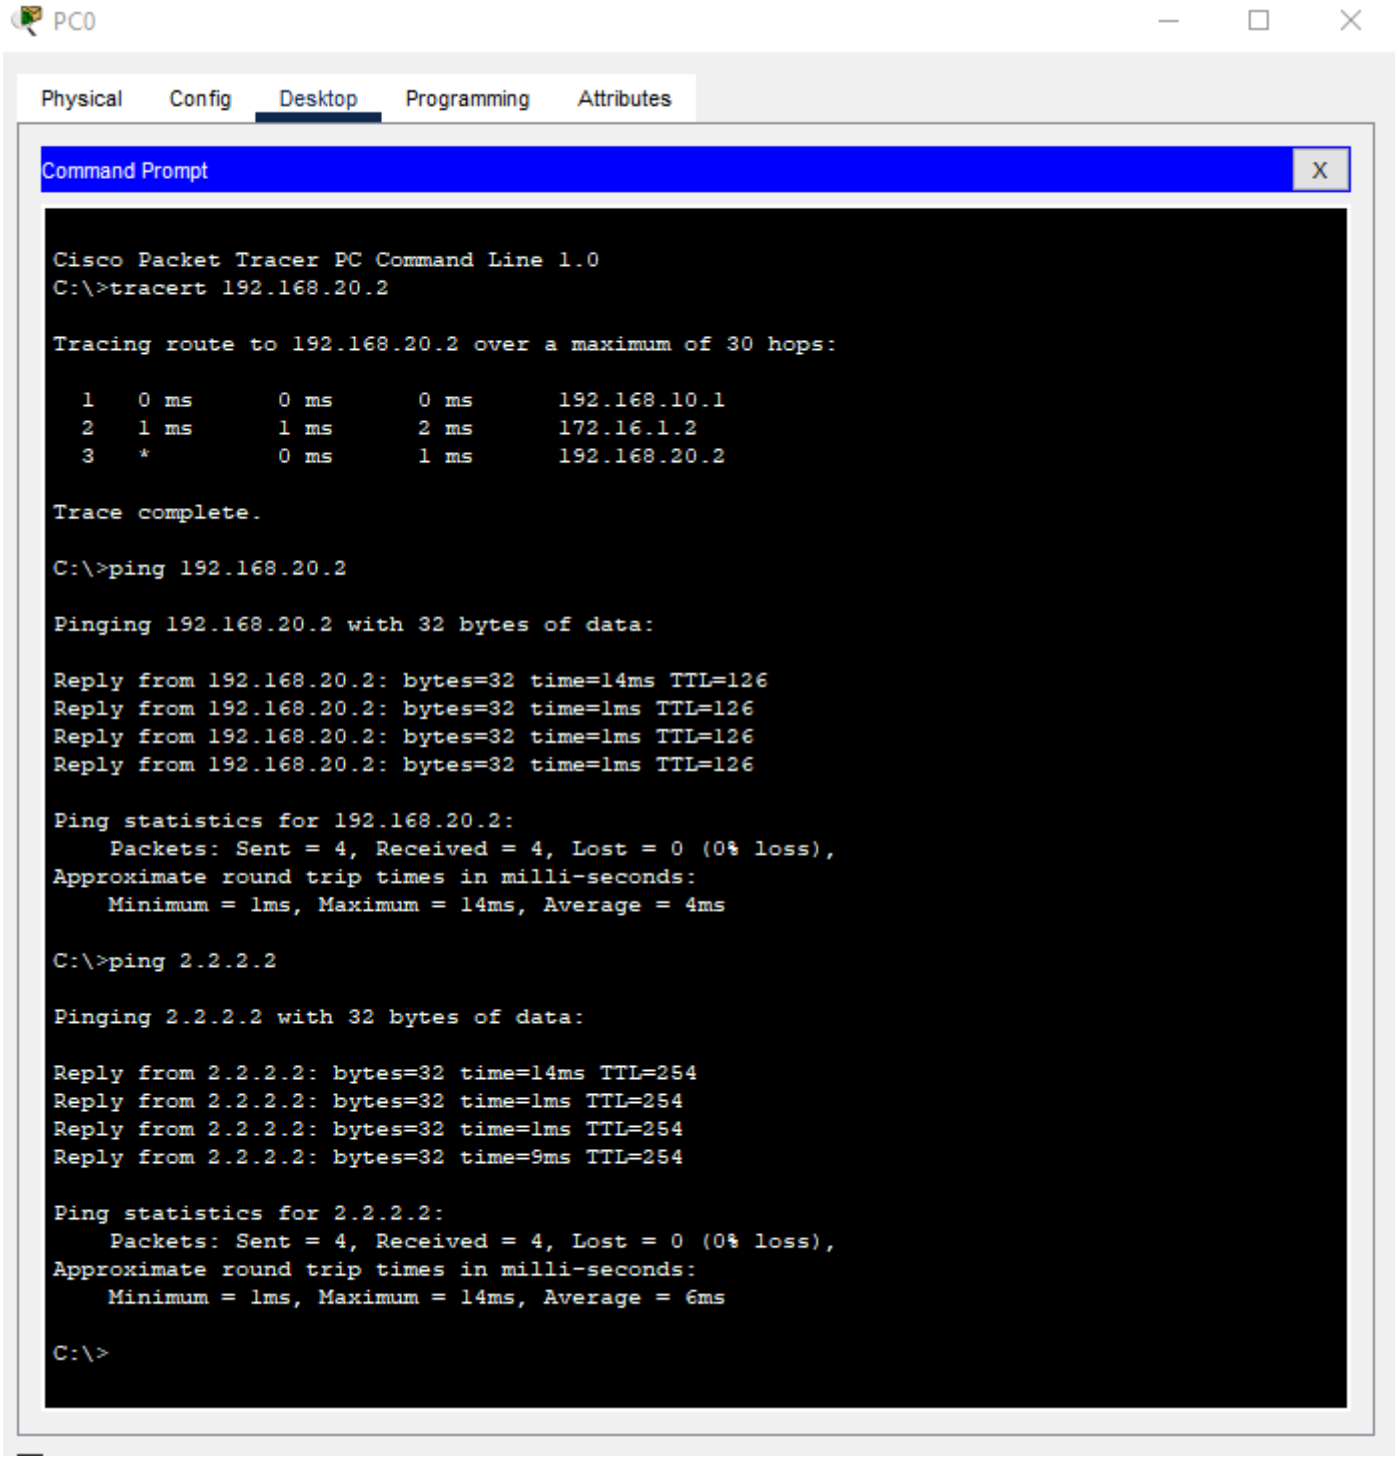
PC0
Physical
Config
Desktop
Programming
Attributes
Command Prompt
X
Cisco Packet Tracer PC Command Line 1.0
C:\>tracert 192.168.20.2
Tracing route to 192.168.20.2 over a maximum of 30 hops:
1 0 ms 0 ms 0 ms 192.168.10.1
2 1 ms 1 ms 2 ms 172.16.1.2
3 * 0 ms 1 ms 192.168.20.2
Trace complete.
C:\>ping 192.168.20.2
Pinging 192.168.20.2 with 32 bytes of data:
Reply from 192.168.20.2: bytes=32 time=14ms TTL=126
Reply from 192.168.20.2: bytes=32 time=1ms TTL=126
Reply from 192.168.20.2: bytes=32 time=1ms TTL=126
Reply from 192.168.20.2: bytes=32 time=1ms TTL=126
Ping statistics for 192.168.20.2:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 1ms, Maximum = 14ms, Average = 4ms
C:\>ping 2.2.2.2
Pinging 2.2.2.2 with 32 bytes of data:
Reply from 2.2.2.2: bytes=32 time=14ms TTL=254
Reply from 2.2.2.2: bytes=32 time=1ms TTL=254
Reply from 2.2.2.2: bytes=32 time=1ms TTL=254
Reply from 2.2.2.2: bytes=32 time=9ms TTL=254
Ping statistics for 2.2.2.2:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 1ms, Maximum = 14ms, Average = 6ms
C:\>

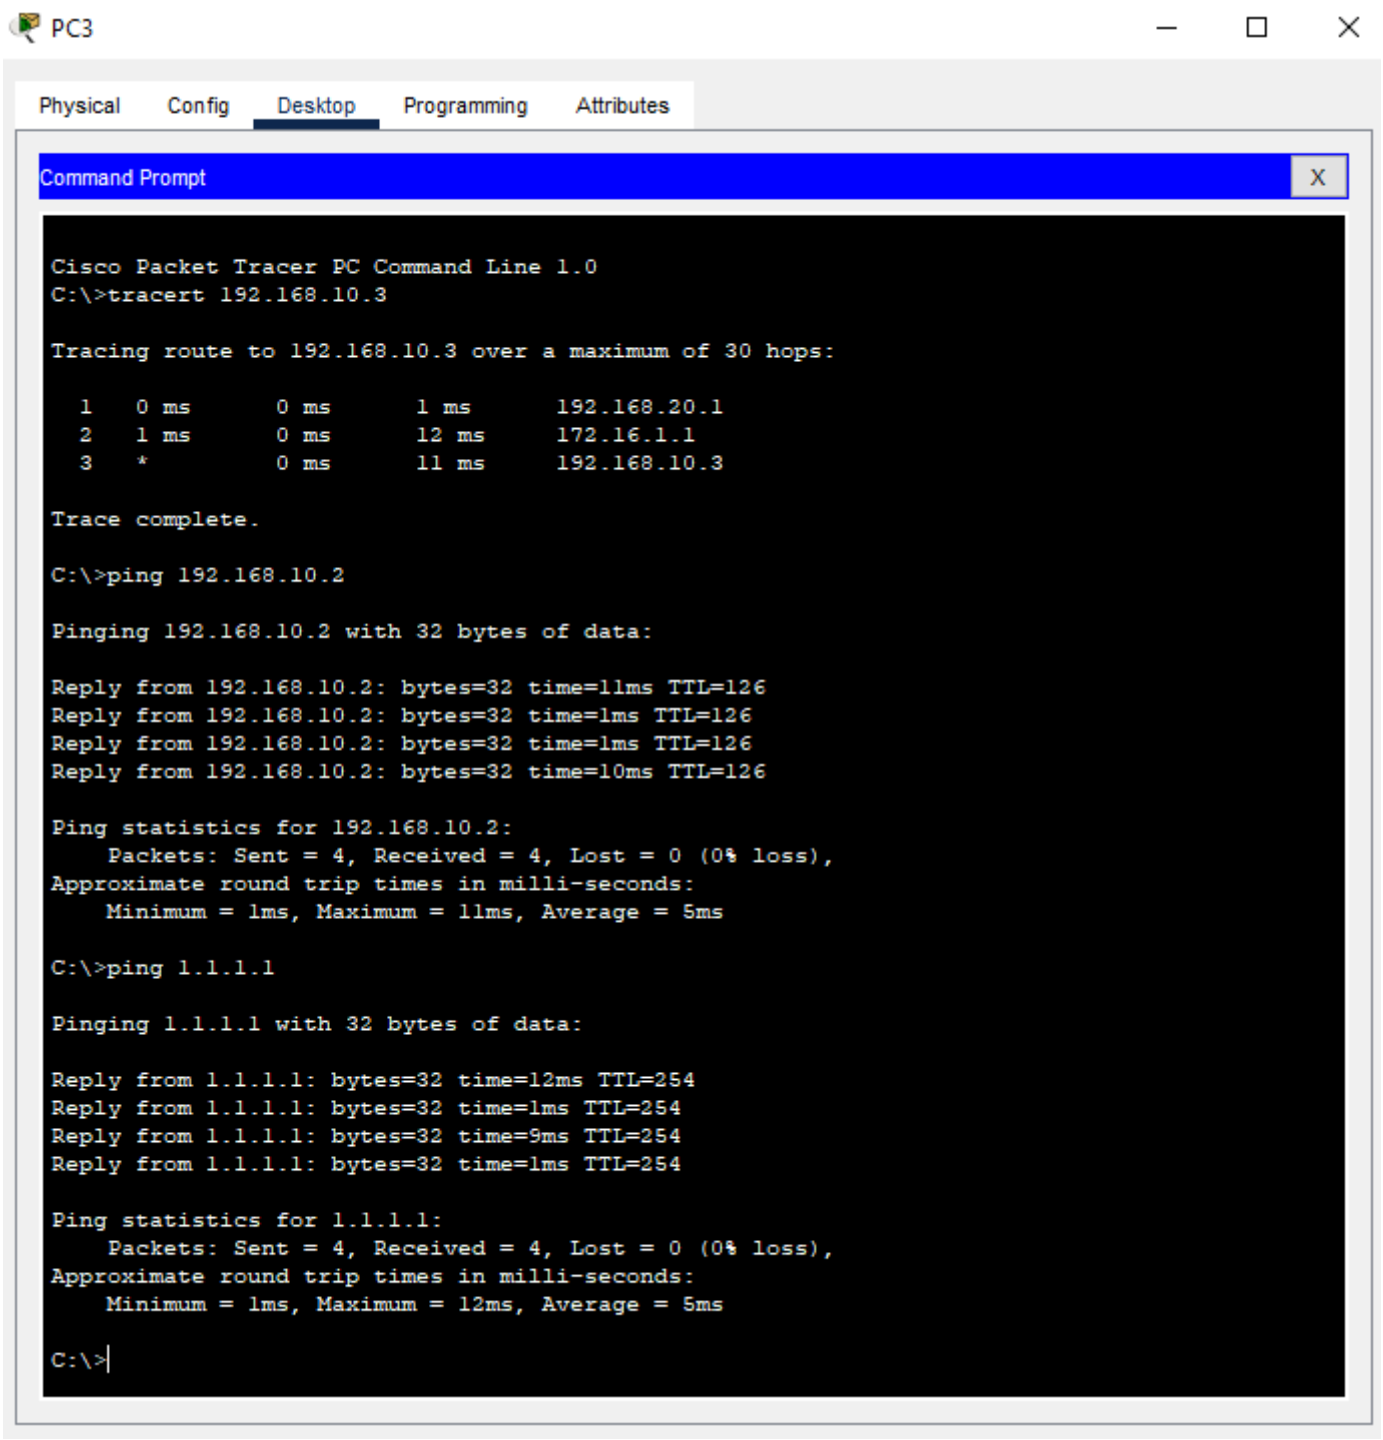

More Notes on BGP

- BGP communicates using TCP port 179
- BGP can form neighbor adjacencies with directly connected routers, which isn't a surprise as other routing protocols do that. But BGP can also form neighbor adjacencies with routers multiple hops away.

Directly Connected BGP Neighbors	Multihop BGP Neighbors
BGP will use the Arp table to locate the Layer 2 address of the peer.	BGP will use routing table information to find the peer's IP address.

- BGP is a path-vector routing protocol meaning it uses path attributes that are associated with each network path when selecting the best route. This also helps ensure that the path taken is loop free.
- BGP path attributes are defined in RFC 4271 (January 2006 release date).
- RFC 1654 defined BGP and termed it an Inter-Autonomous System routing protocol.
- The 'AS' in the network topology is short for Autonomous System. An Autonomous System is. 'the entire routing domain controlled by a company, ISP, or other organization.
- Inter-Autonomous means that BGP is able to route packets across organizations' routing domains. This makes BGP perfect for the routing of the Internet.
- An organization requests an Autonomous System Number (ASN) from the Internet Service Provider (ISP) or more typically from the Internet Assigned Numbers Authority IANA.
- The ASN is a 16-bit or 32-bit number.
- 32-bit ASN length provides for 4,294,967,295 unique ASNs.
- There are private ASNs that any organization can use. These are similar in concept to the private IP ranges that any organization can use internally listed below.
  - Class A 10.0.0.0 - 10.255.255.255
  - Class B 172.16.0.0 - 172.31.255.255
  - Class C 192.168.0.0 - 192.168.255.255

Private 16-bit ASN Range	Private 32-bit ASN Range
64,512 - 65,535	4,200,000,000 - 4,294,967,294

- BGP Peer Communication Message types:
  - Open Message
    - establishes the BGP adjacency
    - Contains - BGP version number, AS number, Hold down timer, other parameters
  - Keep Alive
    - Ensures the neighbors are still active
    - KeepAlive timer set for every 60 seconds by default
    - 1/3 of the HoldDown timer, which is 180 seconds by default.
  - Update
    - These messages are for an update/change to the network.
    - NLRI - Network Layer Reachability Information is included in an update message
  - Notification
    - Error detected
      - Neighbor down
      - Expiration of HoldDown timer
    - BGP Session Reset Requests (clear ip bgp *)

Router0#show ip bgp neighbors

```
Message statistics:
  InQ depth is 0
  OutQ depth is 0

                     Sent       Rcvd
  Opens:                2          2
  Notifications:        0          1
  Updates:              4          4
  Keepalives:         235        235
  Route Refresh:        0          0
  Total:              241        242
```

Another Limitation of Cisco Packet Tracer

As of Packet Tracer 8.2 there isn't a command for debug ip bgp which is very odd because there are debug commands for other protocols

```
Router0#debug ip ?
  eigrp    IP-EIGRP information
  icmp     ICMP transactions
  nat      NAT events
  ospf     OSPF information
  packet   Packet information
  rip      RIP protocol transactions
  routing  Routing table events
```

So, for this next section I will be using GNS3 to show the BGP connection handshake. The BGP establishes a TCP session with a neighboring BGP peer or peers. The connection between the peers **may** report the following states while the connection is established.

- Idle
- Connect
- Active
- OpenSent
- OpenConfirm
- Established

The states can be quite fluid in the ordering as can be seen by and official Cisco diagram showing the states.

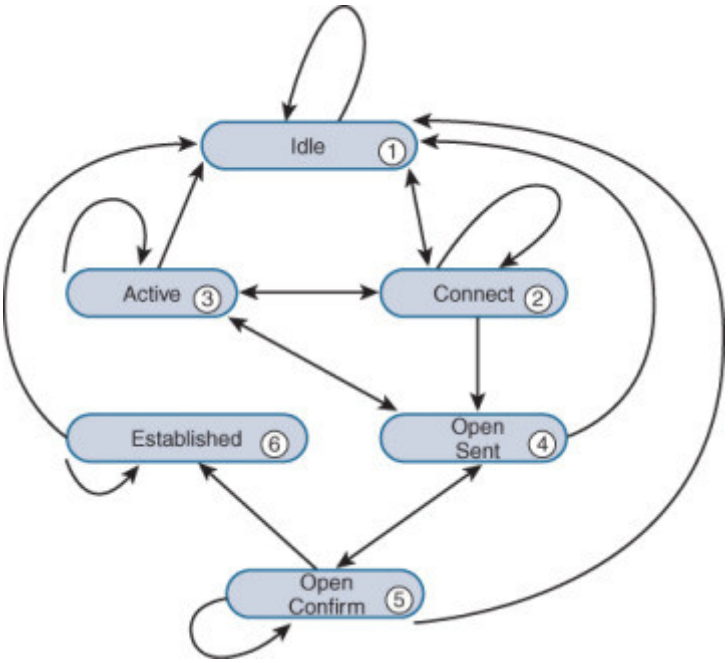

To demo these states in GNS3 I will initiate a debug for the BGP protocol.

Router0#debug ip bgp
Router0#clear ip bgp *

IDLE:

```
*Jan  9 19:56:21.862: BGP: 172.16.1.2 went from Established to Closing
*Jan  9 19:56:21.862: BGP: ses global 172.16.1.2 (0x10EB8290:1) Send NOTIFICATION 6/4 (Administrative Reset) 0 bytes
*Jan  9 19:56:21.863: %BGP-3-NOTIFICATION_MANY: sent to 1 sessions 6/4 (Administrative Reset) for all peers
*Jan  9 19:56:21.863: BGP: system reset due to User reset
*Jan  9 19:56:21.866: BGP: nbr_topo global 172.16.1.2 IPv4 Unicast:base (0x10EB8290:1) NSF delete stale NSF not active
*Jan  9 19:56:21.866: BGP: nbr_topo global 172.16.1.2 IPv4 Unicast:base (0x10EB8290:1) NSF no stale paths state is NSF not a
ctive
*Jan  9 19:56:21.867: BGP: nbr_topo global 172.16.1.2 IPv4 Unicast:base (0x10EB8290:1) Resetting ALL counters.
*Jan  9 19:56:21.867: BGP: nbr global 172.16.1.2 Active open failed - can't get active topologies
*Jan  9 19:56:21.873: BGP: nbr_topo global 172.16.1.2 IPv4 Unicast:base (0x10EB8290:1) NSF delete stale NSF not active
*Jan  9 19:56:21.874: BGP: nbr_topo global 172.16.1.2 IPv4 Unicast:base (0x10EB8290:1) NSF no stale paths state is NSF not a
ctive
*Jan  9 19:56:21.874: BGP: nbr_topo global 172.16.1.2 IPv4 Unicast:base (0x10EB8290:1) Resetting ALL counters.
*Jan  9 19:56:21.874: BGP: 172.16.1.2 closing
*Jan  9 19:56:21.876: BGP: ses global 172.16.1.2 (0x10EB8290:1) Session close and reset neighbor 172.16.1.2 topostate
*Jan  9 19:56:21.877: BGP: nbr topo global 172.16.1.2 IPv4 Unicast:base (0x10EB8290:1) Resetting ALL counters.
*Jan  9 19:56:21.877: BGP: 172.16.1.2 went from Closing to Idle
*Jan  9 19:56:21.878: %BGP-5-ADJCHANGE: neighbor 172.16.1.2 Down User reset
```

CONNECT:

For the connect phase, I never explicitly saw this reported in the debug logs. I tried two different routers with different Cisco IOS versions and in both cases Idle to Connect never logged. This must be what Cisco was referring to when they said, "*The connection between the peers **may** report the following states while the connection is established.*"

ACTIVE:

```
*Jan  9 19:56:21.878: %BGP_SESSION-5-ADJCHANGE: neighbor 172.16.1.2 IPv4 Unicast topology base removed from session  User re
set
Router0#
*Jan  9 19:56:21.878: BGP: ses global 172.16.1.2 (0x10EB8290:1) Removed topology IPv4 Unicast:base
*Jan  9 19:56:21.878: BGP: ses global 172.16.1.2 (0x10EB8290:1) Removed last topology
*Jan  9 19:56:21.879: BGP: nbr global 172.16.1.2 Open active delayed 10240ms (35000ms max, 60% jitter)
*Jan  9 19:56:21.879: BGP: nbr global 172.16.1.2 Active open failed - open timer running
Router0#
*Jan  9 19:56:31.735: BGP: 172.16.1.2 active went from Idle to Active
*Jan  9 19:56:31.735: BGP: 172.16.1.2 open active, local address 172.16.1.1
*Jan  9 19:56:31.741: BGP: ses global 172.16.1.2 (0x10EB7938:0) act Adding topology IPv4 Unicast:base
*Jan  9 19:56:31.742: BGP: ses global 172.16.1.2 (0x10EB7938:0) act Send OPEN
*Jan  9 19:56:31.742: BGP: ses global 172.16.1.2 (0x10EB7938:0) act Building Enhanced Refresh capability
```

OPENSENT:

```
*Jan  9 19:56:31.742: BGP: ses global 172.16.1.2 (0x10EB7938:0) act Building Enhanced Refresh capability
*Jan  9 19:56:31.742: BGP: 172.16.1.2 active went from Active to OpenSent
*Jan  9 19:56:31.743: BGP: 172.16.1.2 active sending OPEN, version 4, my as: 65001, holdtime 180 seconds, ID 1010101
*Jan  9 19:56:31.752: BGP: 172.16.1.2 active rcv message type 1, length (excl. header) 38
*Jan  9 19:56:31.752: BGP: ses global 172.16.1.2 (0x10EB7938:0) act Receive OPEN
*Jan  9 19:56:31.753: BGP: 172.16.1.2 active rcv OPEN, version 4, holdtime 180 seconds
*Jan  9 19:56:31.753: BGP: 172.16.1.2 active rcv OPEN w/ OPTION parameter len: 28
*Jan  9 19:56:31.753: BGP: 172.16.1.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 6
*Jan  9 19:56:31.753: BGP: 172.16.1.2 active OPEN has CAPABILITY code: 1, length 4
*Jan  9 19:56:31.753: BGP: 172.16.1.2 active OPEN has MP_EXT CAP for afi/safi: 1/1
*Jan  9 19:56:31.754: BGP: 172.16.1.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 2
*Jan  9 19:56:31.754: BGP: 172.16.1.2 active OPEN has CAPABILITY code: 128, length 0
*Jan  9 19:56:31.754: BGP: 172.16.1.2 active OPEN has ROUTE-REFRESH capability(old) for all address-families
*Jan  9 19:56:31.754: BGP: 172.16.1.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 2
*Jan  9 19:56:31.754: BGP: 172.16.1.2 active OPEN has CAPABILITY code: 2, length 0
*Jan  9 19:56:31.754: BGP: 172.16.1.2 active OPEN has ROUTE-REFRESH capability(new) for all address-families
*Jan  9 19:56:31.755: BGP: 172.16.1.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 2
*Jan  9 19:56:31.755: BGP: 172.16.1.2 active OPEN
Router0#has CAPABILITY code: 70, length 0
*Jan  9 19:56:31.755: BGP: ses global 172.16.1.2 (0x10EB7938:0) act Enhanced Refresh cap received in open message
*Jan  9 19:56:31.755: BGP: 172.16.1.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 6
*Jan  9 19:56:31.755: BGP: 172.16.1.2 active OPEN has CAPABILITY code: 65, length 4
*Jan  9 19:56:31.756: BGP: 172.16.1.2 active OPEN has 4-byte ASN CAP for: 65002
*Jan  9 19:56:31.756: BGP: 172.16.1.2 active rcvd OPEN w/ remote AS 65002, 4-byte remote AS 65002
```

OPENCONFIRM:

```
*Jan  9 19:56:31.756: BGP: 172.16.1.2 active rcvd OPEN w/ remote AS 65002, 4-byte remote AS 65002
*Jan  9 19:56:31.756: BGP: 172.16.1.2 active went from OpenSent to OpenConfirm
*Jan  9 19:56:31.758: BGP: ses global 172.16.1.2 (0x10EB7938:0) act read request no-op
*Jan  9 19:56:31.759: B
```

ESTABLISHED:

```
Router0#GP: 172.16.1.2 active went from OpenConfirm to Established
*Jan  9 19:56:31.759: BGP: ses global 172.16.1.2 (0x10EB7938:1) act Assigned ID
*Jan  9 19:56:31.760: BGP: ses global 172.16.1.2 (0x10EB7938:1) Up
*Jan  9 19:56:31.761: %BGP-5-ADJCHANGE: neighbor 172.16.1.2 Up
*Jan  9 19:56:32.757: BGP_Router: unhandled major event code 128, minor 0
Router0#
*Jan  9 19:57:02.814: BGP: topo global:IPv4 Unicast:base Scanning routing tables
*Jan  9 19:57:02.815: BGP: Applying map to find origin for 1.0.0.0/8
*Jan  9 19:57:02.815: BGP: Applying map to find origin for 192.168.10.0/24
*Jan  9 19:57:02.815: BGP: topo global:IPv4 Multicast:base Scanning routing tables
*Jan  9 19:57:02.816: BGP: topo global:L2VPN E-VPN:base Scanning routing tables
*Jan  9 19:57:02.816: BGP: topo global:MVPNv4 Unicast:base Scanning routing tables
```

Now we can check the TCP connectivity.

Router0#show tcp brief

```
Router0#show tcp brief
TCB        Local Address               Foreign Address             (state)
10EB7B40  172.16.1.1.57597          172.16.1.2.179              ESTAB
```

The connection is established. Note that the TCP port on the peer (foreign address) is 179. This peer has the higher IP address and thus manages the connect phase and establishes the port 179 connectivity. The peer with the lower IP will get a randomized port for the connectivity. In this example the randomized port is 57597.

Additional Information about BGP from Cloudflare

Additional Information about BGP Neighbor States and Connectivity from CiscoPress

Cisco Packet Tracer File

net14 bgp.pkt

Revision #14
Created 6 January 2023 18:39:36 by Glen Taylor
Updated 24 January 2023 22:45:52 by Glen Taylor