INDICE

1 Introduzione

Il presente documento si propone di definire il comportamento di links IBGP (Internal BGP)

2 IBGP two peering, three node, RIB-BGP table

IBGP rule advertising between IBGP not permitted – rib failure – different advertisement

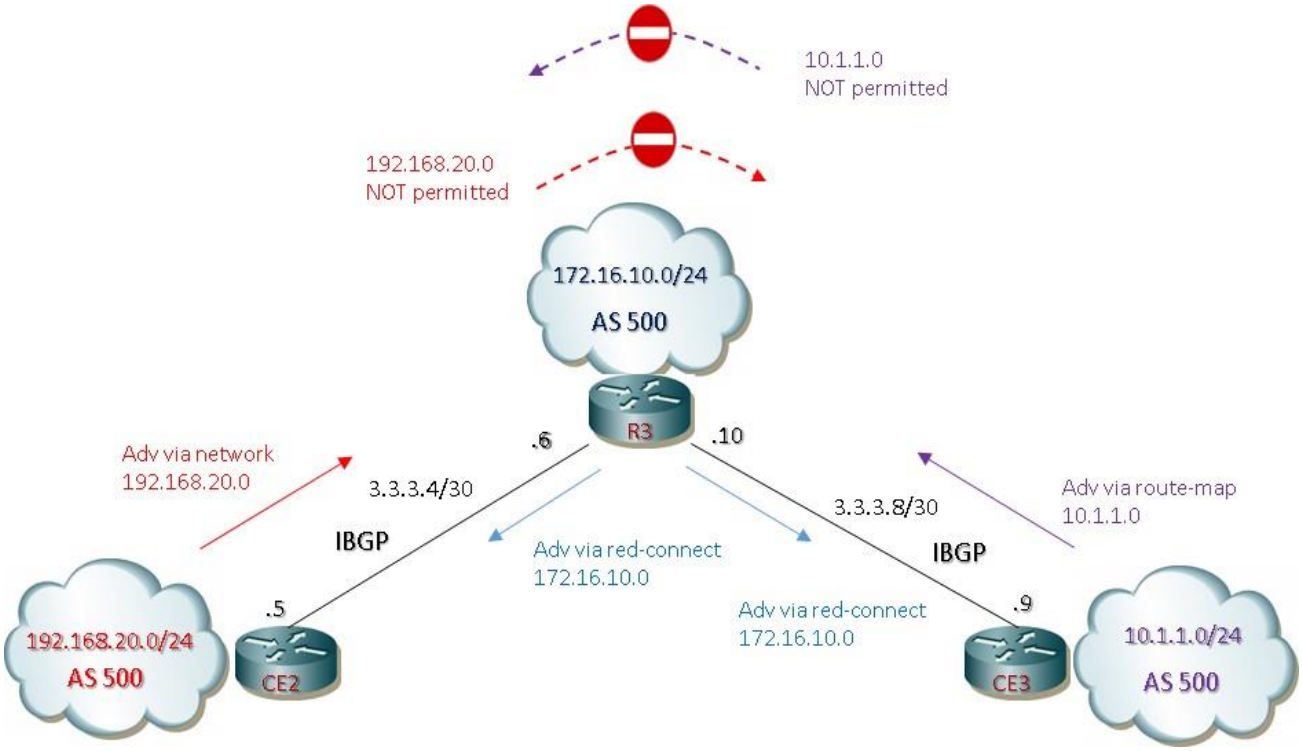

Figura 1: two IBGP , three node

Prima configurazione di base:

R3

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.5 remote-as 500
 neighbor 3.3.3.9 remote-as 500
!
 address-family ipv4
  redistribute connected
  neighbor 3.3.3.5 activate
  neighbor 3.3.3.9 activate
```

CE2

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.6 remote-as 500
!
 address-family ipv4
  network 192.168.20.0
  neighbor 3.3.3.6 activate
```

CE3

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.10 remote-as 500
!
 address-family ipv4
  redistribute connected route-map LAN10
  neighbor 3.3.3.10 activate
!
access-list 10 permit 10.1.1.0  0.0.0.255
!
route-map LAN10 permit 10
match ip address 10
!
route-map LAN10 permit 20
```

RIB and BGP table outputs from **R3**

```
R3#show ip route
Gateway of last resort is not set
   3.0.0.0/8   is variable subnetted, subnets, 2 masks
C  3.3.3.4/30   is directly connected, GigabitEthernet0/2
L  3.3.3.6/32   is directly connected, GigabitEthernet0/2
C  3.3.3.8/30   is directly connected, GigabitEthernet0/1
L  3.3.3.10/32  is directly connected, GigabitEthernet0/1
   10.0.0.0/24  is subnetted, 1 subnet
B  10.1.1.0  [200/0], via 3.3.3.9, 00:28:16
   172.16.10.0  is variable subnetted, 2 subnet, 2 masks
C  172.16.10.0/24 is directly connected, GigabitEthernet0/0
L  172.16.10.1/32 is directly connected, GigabitEthernet0/0
B  192.168.20.0/24  [200/0], via 3.3.3.5, 00:40:26
!
```

```
R3#show ip bgp
     Network          Next-Hop      Metric      LocPrf      Weight      Path
*>   3.3.3.4/30       0.0.0.0            0                   32768      ?
*  i 3.3.3.8/30       3.3.3.9            0         100           0      ?
*>                    0.0.0.0            0                   32768      ?
*> i 10.1.1.0/24      3.3.3.9            0         100           0      ?
*>   172.16.10.0/24   0.0.0.0            0                   32768      ?
*> i 192.168.20.0     3.3.3.5            0        100            0      i

i= origine della prefix è il router stesso (IGP) attraverso il comando network (? Incomplete è
quando la prefix viene annunciata via redistribuzione o route-map)


R3# show ip bgp 3.3.3.4
BGP routing table entry for 3.3.3.4/30 version 14
Paths: (1 available, best#1, table default)
 Advertised to update-group
     1
 Refresh Epoch 1
 Local
   0.0.0.0 from 0.0.0.0 (172.16.10.1)
     Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
     rx pathid: 0, tx pathid: 0x0


R3# show ip bgp 3.3.3.8
BGP routing table entry for 3.3.3.8/30 version 15
Paths: (2 available, best#2, table default)
 Advertised to update-group
     1
 Refresh Epoch 2
 Local
   3.3.3.9 from 3.3.3.9 (3.3.3.9)
     Origin incomplete, metric 0, localpref 100, weight 32768, valid, internal
     rx pathid: 0, tx pathid: 0x0
Refresh Epoch 1
 Local
   0.0.0.0 from 0.0.0.0 (172.16.10.1)
     Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
     rx pathid: 0, tx pathid: 0x0


L'output in Epoch 2 è significativo per la redistribuzione via red-connected route-map da parte
di CE3 (3.3.3.9); pertanto il router R3 avrà nella sua tabella BGP due differenti paths di cui
uno direttamente connesso attraverso la sua interfaccia gi0/1 (visto in Epoch 1 e best-path) ed
uno attraverso la redistribuzione delle connesse con route-map di cui sopra (visto in Epoch 2);
Se ad esempio presso il router CE3, annulliamo la redistribuzione ed annunciamo la prefix via
network statement (annunciando solo quello definito dal comando), avremmo:

CE3#

router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.10 remote-as 500
!
```

```
 address-family ipv4
  no redistribute connected route-map LAN10
  network 10.1.1.0 mask 255.255.255.0
  neighbor 3.3.3.10 activate
```

La nuova tabella BGP di R3 sarà:

```
R3#show ip bgp
       Network          Next-Hop      Metric      LocPrf      Weight      Path
*>     3.3.3.4/30       0.0.0.0            0                   32768      ?
*>     3.3.3.8/30       0.0.0.0            0                   32768      ?
*> i   10.1.1.0/24      3.3.3.9            0         100           0      ?
*>     172.16.10.0/24   0.0.0.0            0                   32768      ?
*> i   192.168.20.0     3.3.3.5            0        100            0      i
```

E quindi attraverso lo show ip bgp della network 3.3.3.8 avremmo un solo path disponibile:

```
R3# show ip bgp 3.3.3.8
BGP routing table entry for 3.3.3.8/30 version 15
Paths: (1 available, best#1, table default)
 Advertised to update-group
     1
 Refresh Epoch 1
 Local
   0.0.0.0 from 0.0.0.0 (172.16.10.1)
     Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
     rx pathid: 0, tx pathid: 0x0
```

Gli annunci e le route ricevute sono indicate dai seguenti outputs:

```
R3# show ip bgp neighbors 3.3.3.5 advertised-routes     # to CE2 (via red-connected)
      Network          Next-Hop      Metric      LocPrf      Weight      Path
*>    3.3.3.4/30        0.0.0.0           0                   32768       ?
*>    3.3.3.8/30        0.0.0.0           0                   32768       ?
*>    172.16.10.0/24    0.0.0.0           0                   32768       ?
!
!
R3# show ip bgp neighbors 3.3.3.5 routes    # from CE2 (via network)
       Network          Next-Hop      Metric      LocPrf      Weight      Path
*> i   192.168.20.0/24   3.3.3.5           0         100           0      i
!
!
R3# show ip bgp neighbors 3.3.3.9 advertised-routes    # to CE3 (via red-connected)
      Network          Next-Hop        Metric      LocPrf      Weight      Path
*>    3.3.3.4/30        0.0.0.0             0                   32768      ?
```

```
*>    3.3.3.8/30          0.0.0.0             0                    32768      ?

*>    172.16.10.0/24      0.0.0.0             0                    32768      ?

!
!
R3# show ip bgp neighbors 3.3.3.9 routes     # from CE3 (via network)

        Network           Next-Hop       Metric       LocPrf      Weight      Path

*> i    10.1.1.0/24       3.3.3.9             0          100           0      ?


OPPURE CON DIVERSA REDISTRIBUZIONE FROM CE3



R3# show ip bgp neighbors 3.3.3.9 routes     # from CE3 (via red-connected route-map)

         Network           Next-Hop       Metric       LocPrf      Weight      Path

*  i     3.3.3.8/30        3.3.3.9             0          100           0      ?

*> i     10.1.1.0/24       3.3.3.9             0          100           0      ?
```

Tabella 1: RIB and BGP Table from R3 quale router of advertisement prefix

RIB and BGP table outputs from **CE2**

```
CE2#show ip route

Gateway of last resort is not set

    3.0.0.0/8  is variable subnetted, 3 subnets, 2 masks

C   3.3.3.4/30  is directly connected, GigabitEthernet0/2

L   3.3.3.5/32  is directly connected, GigabitEthernet0/2

B   3.3.3.8/30  [200/0], via 3.3.3.6, 02:10:36

    172.16.10.0/24 is subnetted, 1 subnet

B   172.16.10.0  [200/0] via 3.3.3.6, 01:40:42

    192.168.20.0 is variable subnetted, 2 subnets, 2 masks

C   192.168.20.0/24 is directly connected, GigabitEthernet0/1

L   192.168.20.1/32 is directly connected, GigabitEthernet0/1

!
```

La prefix 10.1.1.0/24 annunciata dal CE3 (non direttamente connesso al CE2) non viene propagate da R3 (nodo intermedio IBGP tra i due CE) per il corretto comportamento del protocollo I-BGP per il quale le routes/prefix apprese via IBGP non possono essere re-annunciate sempre via IBGP ad altri neighbors (regola dello split-horizon).

```
CE2#show ip bgp
        Network          Next-Hop      Metric       LocPrf        Weight      Path
r> i    3.3.3.4/30       3.3.3.6            0          100             0      ?
*> i    3.3.3.8/30       3.3.3.6            0          100             0      ?
*> i    172.16.10.0/24   3.3.3.6            0          100             0      ?
*>      192.168.20.0     0.0.0.0            0                      32768      i
!
```

RIB-Failure (r>) significa che solo la best-route può essere in RIB-Failure e pertanto la route 3.3.3.4/30 annunciata da R3 attraverso il comando redistribute connected, viene inserita in questo stato perché dal punto di vista di CE2 esiste già una route con questa subnet presente nella tabella di routing con una migliore distanza amministrativa

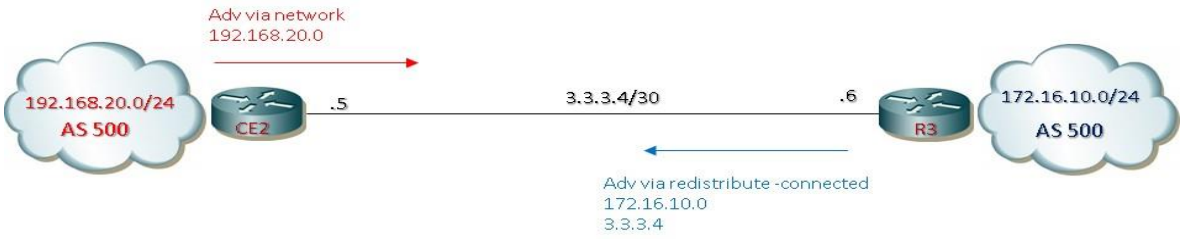


```
CE2#show ip bgp rib-failure
        Network          Next-Hop                RIB-failure       RIB-NH Matches
        3.3.3.4/30       3.3.3.6          High admin distance                n/a
```

Se volessimo annullare la presenza di rib-failure presso la tabella BGP di CE2, basta settare presso il router R3 il comando network statement, e quindi annunciare la sola prefix definita dal comando stesso, oppure settare lato CE2 un redistribute connected, annunciando non solo le prefix internal ma anche le direttamente connesse ma questa volta avendo una pari distanza amministrativa dal punto di vista di entrambi i peers IBGP.

Example

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.6 remote-as 500
!
 address-family ipv4
  no network 192.168.20.0
  redistribute connected
  neighbor 3.3.3.6 activate
```

Ed avremmo:

```
CE2#show ip bgp

        Network          Next-Hop      Metric      LocPrf        Weight      Path
*>      3.3.3.4/30       0.0.0.0            0                     32768      ?
*  i                     3.3.3.6            0         100             0      ?
*> i    3.3.3.8/30       3.3.3.6            0         100             0      ?
*> i    172.16.10.0/24   3.3.3.6            0         100             0      ?
*>      192.168.20.0     0.0.0.0            0                     32768      i
!

Gli annunci e le route ricevute sono indicate dai seguenti outputs:


CE2# show ip bgp neighbor 3.3.3.6 advertised-routes     # to R3 (via red-connected)

         Network          Next-Hop       Metric       LocPrf       Weight      Path
*        3.3.3.4/30       0.0.0.0             0                     32768      ?
*>       192.168.20.0/24  3.3.3.9             0                     32768      ?

CE2# show ip bgp neighbor 3.3.3.6 advertised-routes     # from R3 (via red-connected)

         Network          Next-Hop       Metric       LocPrf       Weight      Path
*  i     3.3.3.4/30       3.3.3.6             0          100             0      ?
*> i     3.3.3.8/30       3.3.3.6             0          100             0      ?
*> i     172.16.10.0/24   3.3.3.6             0          100             0      ?

OPPURE CAMBIANDO LA REDISTRIBUZIONE SU CE2 RIPORTANDOLA COME ALL'INIZIO:

CE2# show ip bgp neighbors 3.3.3.6 advertised-routes     # to R3 (via network statement)

      Network          Next-Hop      Metric       LocPrf      Weight      Path
*>    192.168.20.0/24  0.0.0.0            0                     32768      i


CE2# show ip bgp neighbors 3.3.3.6 routes        # from R3 (via red-connected)

       Network          Next-Hop      Metric       LocPrf      Weight      Path
r> i   3.3.3.4/30       3.3.3.6            0          100             0      ?
*> i   3.3.3.8/30       3.3.3.6            0          100             0      ?
*> i   172.16.10.0/24    3.3.3.6            0          100             0      ?
!
```

Tabella 2: RIB and BGP Table from CE2 with direct IBGP with R3

RIB and BGP table outputs from **CE3**

```
CE3#show ip route
Gateway of last resort is not set
    3.0.0.0/8  is variable subnetted, 3 subnets, 2 masks
B   3.3.3.4/30  [200/0] via 3.3.3.10, 04:43:37
C   3.3.3.8/30  is directly connected, GigabitEthernet0/2
L   3.3.3.9/32  is directly connected, GigabitEthernet0/2
   10.0.0.0/8   is variable subnetted, 2 subnets, 2 masks
C  10.1.1.0/24  is directly connected, GigabitEthernet0/1
L  10.1.1.1/32  is directly connected, GigabitEthernet0/1
   172.16.10.0/24 is subnetted, 1 subnet
B  172.16.10.0  [200/0] via 3.3.3.10, 04:43:37
!
```

La prefix 192.168.20.0/24 annunciata dal CE2 (non direttamente connesso al CE3) non viene propagate da R3 (nodo intermedio IBGP tra i due CE) per il corretto comportamento del protocollo I-BGP per il quale le routes/prefix apprese via IBGP non possono essere re-annunciate sempre via IBGP ad altri neighbors (regola dello split-horizon).

```
CE3#show ip bgp
          Network          Next-Hop      Metric       LocPrf       Weight      Path
*> i      3.3.3.4/30       3.3.3.10           0          100            0      ?
*>        3.3.3.8/30       0.0.0.0            0                     32768      ?
*  i                       3.3.3.10           0          100            0      ?
*>        10.1.1.0/24      0.0.0.0            0                     32768      ?
*> i      172.16.10.0/24   3.3.3.10           0          100            0      ?
!
```

Gli annunci e le route ricevute sono indicate dai seguenti outputs:

```
CE3# show ip bgp neighbors 3.3.3.10 advertised-routes   # to R3 (via red-connected route-map)
         Network          Next-Hop      Metric       LocPrf         Weight      Path
*>       3.3.3.8/30       0.0.0.0            0                     32768       ?
*>       10.1.1.0/24      0.0.0.0            0                     32768       ?
!
```

```
CE3# show ip bgp neighbors 3.3.3.10 routes   # from R3 (via red-connected)

       Network            Next-Hop     Metric       LocPrf      Weight      Path
*> i   3.3.3.4/30         3.3.3.10          0          100           0      ?
*  i   3.3.3.8/30         3.3.3.10          0          100           0      ?
*> i   172.16.10.0/24     3.3.3.10          0          100           0      ?
```

Tabella 3: RIB and BGP Table from CE3 with direct IBGP with R3

3 IBGP three peering full-mesh, three node, RIB-BGP table, IGP OSPF

IBGP with two physical peers, one logical peers, Full-Mesh, IGP OSPF, different advertisement, synchronization rules

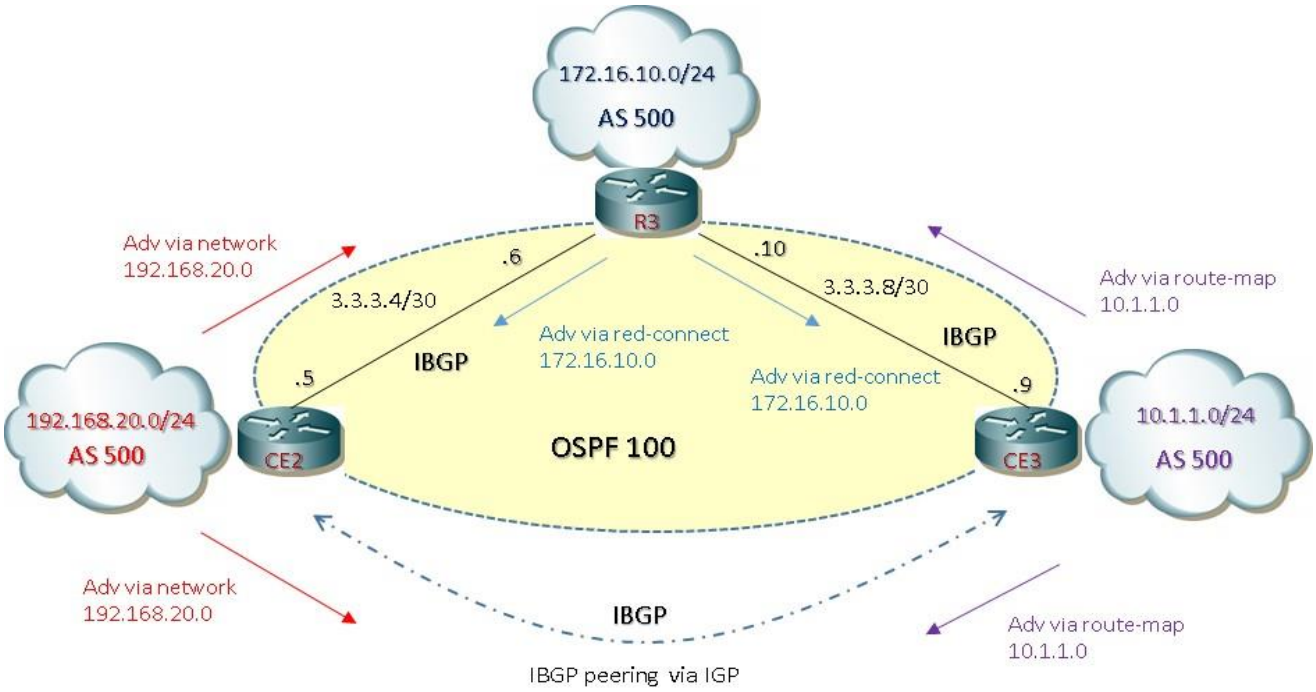

Figura 2: three IBGP Full-Mesh , three node, OSPF-IGP

Con questa architettura possiamo avere una completa redistribuzione delle Prefix sorgenti annunciate rispettivamente attraverso i rispettivi peering, di cui:

- ✓ peering IBGP fisico R3 – CE2
- ✓ peering IBGP fisico R3 – CE3
- ✓ peering IBGP logico via IGP CE2 – CE3

Si crea pertanto una architettura IBGP Full-Mesh dove ogni nodo ha sessioni IBGP con i rispettivi neighbors all'interno del proprio AS e questo permette una redistribuzione delle prefix in modo equivalente avendo tutti tabelle di routing convergenti.

NOTA: in una architettura reale si preferisce stabilire sessioni IBGP Full-Mesh attraverso interfacce di loopback essendo queste logiche e pertanto sempre up a differenza di indirizzi IP configurati su interfacce fisiche.

Seconda configurazione di base:

R3

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.5 remote-as 500
 neighbor 3.3.3.9 remote-as 500
!
 address-family ipv4
  redistribute connected
  neighbor 3.3.3.5 activate
  neighbor 3.3.3.9 activate
!
router ospf 100
network 3.3.3.0  0.0.0.255 area 0
```

CE2

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.6 remote-as 500
 neighbor 3.3.3.9 remote-as 500
!
 address-family ipv4
  network 192.168.20.0
  neighbor 3.3.3.6 activate
  neighbor 3.3.3.9 activate
!
router ospf 100
network 3.3.3.0  0.0.0.255 area 0
```

CE3

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.10 remote-as 500
 neighbor 3.3.3.5 remote-as 500
!
 address-family ipv4
  redistribute connected route-map LAN10
  neighbor 3.3.3.10 activate
  neighbor 3.3.3.5 activate
!
router ospf 100
network 3.3.3.0  0.0.0.255 area 0
!
access-list 10 permit 10 permit 10.1.1.0  0.0.0.255
!
route-map LAN10 permit 10
match ip address 10
!
route-map LAN permit 20
```

RIB and BGP table outputs from **R3** versione full-mesh

```
R3#show ip route
Gateway of last resort is not set
   3.0.0.0./8   is variable subnetted, subnets, 2 masks
C  3.3.3.4/30   is directly connected, GigabitEthernet0/2
L  3.3.3.6/32   is directly connected, GigabitEthernet0/2
C  3.3.3.8/30   is directly connected, GigabitEthernet0/1
L  3.3.3.10/32  is directly connected, GigabitEthernet0/1
   10.0.0.0/24  is subnetted, 1 subnet
B  10.1.1.0  [200/0], via 3.3.3.9, 00:28:16
   172.16.10.0  is variable subnetted, 2 subnet, 2 masks
C  172.16.10.0/24 is directly connected, GigabitEthernet0/0
L  172.16.10.1/32 is directly connected, GigabitEthernet0/0
B  192.168.20.0/24  [200/0], via 3.3.3.5, 00:40:26
!
R3#show ip bgp
       Network          Next-Hop      Metric      LocPrf      Weight      Path
*>     3.3.3.4/30       0.0.0.0            0                   32768      ?
*  i   3.3.3.8/30       3.3.3.9            0         100           0      ?
*>                      0.0.0.0            0                   32768      ?
*> i   10.1.1.0/24       3.3.3.9            0         100           0      ?
*>     172.16.10.0/24    0.0.0.0            0                   32768      ?
*> i   192.168.20.0      3.3.3.5            0        100            0      i
```

i= origine della prefix è il router stesso (IGP) attraverso il comando network (? Incomplete è quando la prefix viene annunciata via redistribuzione o route-map)

La tabella RIB dal punto di vista R3 è identica a prima poiché era l'unico nodo ad avere due sessioni IBGP con i rispettivi neighbors IBGP.
Le tre prefix sorgenti sono regolarmente presenti in tabella.

```
R3# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.5     4    500       123         122         6       0      0   01:49:00            1
3.3.3.9     4    500       123         124         6       0      0   01:48:59            2
```

Con l'ingresso del protocollo OSPF abbiamo:

```
R3# show ip ospf neighbor

Neighbor ID         Pri      State         Dead Time      Address       Interface
192.168.20.1          1      FULL/DR       00:00:35       3.3.3.5       GigabitEthernet0/2
10.1.1.1              1      FULL/DR       00:00:31       3.3.3.9       GigabitEthernet0/1
```

Tabella 4: RIB and BGP Table from R3 with Full-Mesh IBGP

RIB and BGP table outputs from **CE2** versione full-mesh

```
CE2# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.6     4    500       284         286         9       0      0   04:17:08            3
3.3.3.9     4    500       285         287         9       0      0   04:16:17            2
```

Con il protocollo OSPF attivo abbiamo:

```
CE2# show ip ospf neighbor

Neighbor ID         Pri      State         Dead Time      Address       Interface
172.16.10.1           1      FULL/DR       00:00:39       3.3.3.6       GigabitEthernet0/2
```

ed allo stesso tempo attiviamo il commando " synchronization " lato CE2 e CE3 attraverso la configurazione:

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.6 remote-as 500
 neighbor 3.3.3.9 remote-as 500
!
 address-family ipv4
  network 192.168.20.0
  synchronization
  neighbor 3.3.3.6 activate
  neighbor 3.3.3.9 activate
```

abbiamo una regola che ci dice che una route/prefix per essere appresa da un neighbor IBGP deve essere prima presente e quindi conosciuta via IGP ed essere cosi presente nella sua tabella di routing

Con la synchronizaton UP abbiamo:

```
CE2#show ip route

Gateway of last resort is not set

   3.0.0.0./8   is variable subnetted, subnets, 2 masks

C  3.3.3.4/30   is directly connected, GigabitEthernet0/2

L  3.3.3.5/32   is directly connected, GigabitEthernet0/2

O  3.3.3.8/30   [110/2] via 3.3.3.6, 00:25:18, GigabitEthernet0/2

   192.168.20.0  is variably subnetted, 2 subnets, 2 masks

C  192.168.20.0/24 is directly connected, GigabitEthernet0/1

L  192.168.20.1/32 is directly connected, GigabitEthernet0/1
```

Le due prefix sorgenti 10.1.1.0/24 from CE3 e 172.16.10.0/24 from R3 non sono presenti nella tabella di routing

La possibilità di avere una tabella RIB/FIB completa di tutte le prefix sorgenti è possibile disattivando il comando synchronization:

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.6 remote-as 500
 neighbor 3.3.3.9 remote-as 500
!
 address-family ipv4
  network 192.168.20.0
  no synchronization
  neighbor 3.3.3.6 activate
  neighbor 3.3.3.9 activate
```

```
CE2#show ip route

Gateway of last resort is not set

    3.0.0.0/8  is variable subnetted, 3 subnets, 2 masks

C   3.3.3.4/30  is directly connected, GigabitEthernet0/2

L   3.3.3.5/32  is directly connected, GigabitEthernet0/2

O   3.3.3.8/30  [110/2], via 3.3.3.6, 02:10:36

    10.0.0.0/24 is subnetted, 1 subnet

B   10.1.1.0 [200/0] via 3.3.3.9, 00:00:11

    172.16.10.0/24 is subnetted, 1 subnet

 B   172.16.10.0 [200/0] via 3.3.3.6, 01:40:42

    192.168.20.0 is variable subnetted, 2 subnets, 2 masks

C   192.168.20.0/24 is directly connected, GigabitEthernet0/1

L   192.168.20.1/32 is directly connected, GigabitEthernet0/1

!
```

La prefix 10.1.1.0/24 annunciata dal CE3 (non direttamente connesso al CE2) e la prefix 172.16.10.0/24 annunciata da R3 ora sono presenti nella RIB regolarmente.

```
CE2#show ip bgp
        Network          Next-Hop      Metric     LocPrf     Weight     Path
r> i    3.3.3.4/30       3.3.3.6            0        100          0     ?
r  i    3.3.3.8/30       3.3.3.9            0        100          0     ?
r> i                     3.3.3.6            0        100          0     ?
*> i    10.1.1.0/24      3.3.3.9            0        100          0     ?
*> i    172.16.10.0/24   3.3.3.6            0        100          0     ?
*>      192.168.20.0     0.0.0.0            0                 32768     i
!
```

Tabella 5: RIB and BGP Table from CE2 with Full-Mesh IBGP

RIB and BGP table outputs from **CE3** versione full-mesh

```
CE2# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.5     4    500       17          17         7       0      0   00:12:41            1
3.3.3.10    4    500       19          17         7       0      0   00:12:42            3
```

Con il protocollo OSPF attivo abbiamo:

```
CE2# show ip ospf neighbor

Neighbor ID         Pri      State        Dead Time      Address       Interface
172.16.10.1          1      FULL/DR       00:00:31       3.3.3.10       GigabitEthernet0/2
```

```
CE3#show ip route
Gateway of last resort is not set
    3.0.0.0/8  is variable subnetted, 3 subnets, 2 masks
O   3.3.3.4/30  [110/2] via 3.3.3.10, 01:45:50, GigabitEthernet0/2
C   3.3.3.8/30  is directly connected, GigabitEthernet0/2
L   3.3.3.9/32  is directly connected, GigabitEthernet0/2
```

```
    10.0.0.0/24  is variably subnetted, 2 subnet, 2 masks
C   10.1.1.0/24  is directly connected, GigabitEthernet0/1
L   10.1.1.1/32  is directly connected, GigabitEthernet0/1
    172.16.10.0/24 is subnetted, 1 subnet
B   172.16.10.0 [200/0] via 3.3.3.10, 00:14:20
B   192.168.20.0 [200/0] via 3.3.3.5, 00:14:15
!
```

La prefix 192.168.20.0/24 annunciata dal CE2 e la prefix 172.16.10.0/24 annunciata da R3 ora sono presenti nella RIB regolarmente.

```
CE3#show ip bgp
        Network          Next-Hop      Metric      LocPrf        Weight      Path
r> i    3.3.3.4/30       3.3.3.10           0         100             0      ?
*>      3.3.3.8/30       0.0.0.0            0                     32768      ?
*  i                     3.3.3.10           0         100             0      ?
*>      10.1.1.0/24      0.0.0.0            0                     32768      ?
*> i    172.16.10.0/24   3.3.3.10           0         100             0      ?
*> i    192.168.20.0     3.3.3.5            0         100             0      i
!
```

Tabella 6: RIB and BGP Table from CE3 with Full-Mesh IBGP

4 IBGP Confederation

IBGP peering without BGP Confederation

La regola dello split-horzon IBGP può essere risolta attraverso un'architettura di tipo Full-Mesh peering IBGP; nel caso di un numero elevato di peering tra nodi si rende necessaria la soluzione scalabile dei Router Reflector oppure il BGP Confederation.

Consideriamo la seguente architettura:

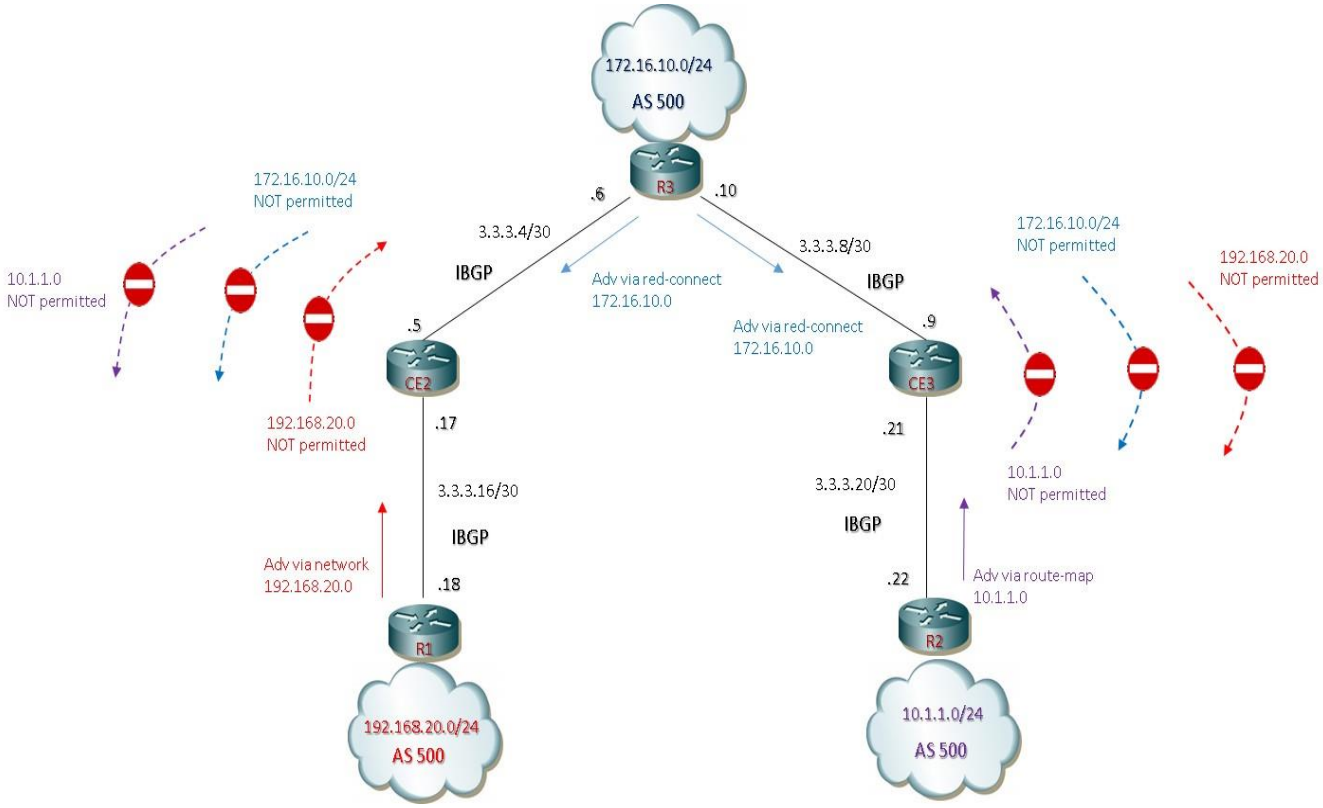

Figura 3: IBGP peering with advertisement NOT-permitted (NO CONFEDERATION)

Terza configurazione di base (SENZA BGP CONFEDERATION):

R3

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.5 remote-as 500
 neighbor 3.3.3.9 remote-as 500
!
 address-family ipv4
  redistribute connected
  neighbor 3.3.3.5 activate
  neighbor 3.3.3.9 activate
!
```

CE2

```
router bgp 500
 bgp router-id 2.2.2.2              # vedi nota sotto
 bgp log-neighbor-changes
 neighbor 3.3.3.6 remote-as 500
 neighbor 3.3.3.18 remote-as 500
!
 address-family ipv4
  neighbor 3.3.3.6 activate
  neighbor 3.3.3.18 activate
```

NOTA: è fondamentale inserire in questo caso il commando bgp router-id <ip-address-loopback> per evitare la non formazione del peering IBGP a causa del seguente massaggio di logging:

CE2# Apr 21 13:57:20:215: %BGP-5-NBR_RESET: Neighbor 3.3.3.18 passive reset (BGP Notification sent)

CE2# Apr 21 13:57:20:217: %BGP-5-ADJCHANGE: neighbor 3.3.3.18 passive Down BGP Notification sent

CE2# Apr 21 13:57:20:222: %BGP-3-NOTIFICATION: received from neighbor 3.3.3.18 active 2/3 (**BGP identifier wrong**) 4 bytes C0A81401

CE2# Apr 21 13:57:20:223: %BGP-5-NBR_RESET: Neighbor 3.3.3.18 active reset (BGP Notification received)

CE2# Apr 21 13:57:20:223: %BGP-5-ADJCHANGE: neighbor 3.3.3.18 active Down BGP Notification received

CE2# Apr 21 13:57:20:223: %BGP_SESSION-5-ADJCHANGE: neighbor 3.3.3.18 IPv4 Unicast topology base removed from session BGP Notification received

CE2# Apr 21 13:57:25:732: %BGP-3-NOTIFICATION: sent to neighbor 3.3.3.18 passive 2/3 (BGP identifier wrong) 4 bytes C0A81401

CE2# Apr 21 13:57:25:732 %BGP-4-MSGDUMP: unsupported or mal-formatted message received from 3.3.3.18

FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0039 0104 01F4 00B4 C0A8 1401 1C02 0601 0400 0100 0102 0280 0002 0202 0002 0246 0002 0641 0400 0001 F4

L'errore consiste in uno sbagliato router-id identificativo scambiato tra i due nuovi peers durante l'instaurazione della sessione.

CE3

```
router bgp 500
 bgp router-id 3.3.3.3
 bgp log-neighbor-changes
 neighbor 3.3.3.10 remote-as 500
 neighbor 3.3.3.22 remote-as 500
!
 address-family ipv4
  neighbor 3.3.3.10 activate
  neighbor 3.3.3.22 activate
```

R1

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.17 remote-as 500
 !
 address-family ipv4
  network 192.168.20.0 mask 255.255.255.0
  neighbor 3.3.3.17 activate
```

R2

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 3.3.3.21 remote-as 500
 !
 address-family ipv4
  redistribute connected route-map LAN10
  neighbor 3.3.3.21 activate
!
access-list 10 permit 10.1.1.0  0.0.0.255
!
route-map LAN10 permit 10
 match ip address 10
!
route-map LAN 10 permit 20
```

Le tabelle seguenti mettono in evidenza la RIB/BGP Table di ciascun nodo peers IBGP, senza il contributo del BGP Confederation:

RIB and BGP table outputs from **R3** senza BGP Confederation

```
R3#show ip route
Gateway of last resort is not set
   3.0.0.0./8   is variable subnetted, subnets, 2 masks
C  3.3.3.4/30   is directly connected, GigabitEthernet0/2
L  3.3.3.6/32   is directly connected, GigabitEthernet0/2
C  3.3.3.8/30   is directly connected, GigabitEthernet0/1
L  3.3.3.10/32  is directly connected, GigabitEthernet0/1
   172.16.10.0  is variable subnetted, 2 subnet, 2 masks
C  172.16.10.0/24 is directly connected, GigabitEthernet0/0
L  172.16.10.1/32 is directly connected, GigabitEthernet0/0
```

```
R3#show ip bgp
     Network          Next-Hop      Metric      LocPrf      Weight      Path
*>   3.3.3.4/30       0.0.0.0            0                   32768      ?
*>   3.3.3.8/30       0.0.0.0            0                   32768      ?
*>   172.16.10.0/24   0.0.0.0            0                   32768      ?

Le due prefix sorgenti dai rispettivi nodi R1 ed R2 NON sono presenti in tabella di routing
regolarmente a causa della regola dello split-horizon imposto da CE2 e CE3.

R3# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.5     4    500       123         137        28       0      0   01:49:00            0
3.3.3.9     4    500       123         136        28       0      0   01:48:59            0
```

Tabella 7: RIB and BGP Table from R3 without BGP Confederation

RIB and BGP table outputs from **R1** senza BGP Confederation

```
R1#show ip route
Gateway of last resort is not set
   3.0.0.0./8   is variable subnetted, subnets, 2 masks
C  3.3.3.16/30   is directly connected, GigabitEthernet0/3
L  3.3.3.18/32   is directly connected, GigabitEthernet0/3
   192.168.20.0  is variable subnetted, 2 subnet, 2 masks
C  192.168.20.0/24 is directly connected, GigabitEthernet0/0
L  192.168.20.1/32 is directly connected, GigabitEthernet0/0

R1#show ip bgp
     Network          Next-Hop      Metric      LocPrf      Weight      Path
*>   192.168.20.0/24    0.0.0.0            0                   32768      ?

Le due prefix sorgenti dai rispettivi nodi R2 ed R3 NON sono presenti in tabella di routing
regolarmente a causa della regola dello split-horizon imposto da CE2 e CE3.

R1# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.17    4    500       145         146         2       0      0   02:10:00            0
```

Tabella 8: RIB and BGP Table from R1 without BGP Confederation

RIB and BGP table outputs from **R2** senza BGP Confederation

```
R2#show ip route
Gateway of last resort is not set
   3.0.0.0./8   is variable subnetted, subnets, 2 masks
C  3.3.3.20/30   is directly connected, GigabitEthernet0/3
L  3.3.3.22/32   is directly connected, GigabitEthernet0/3
   10.0.0.0/8   is variable subnetted, 2 subnet, 2 masks
C  10.1.1.0/24  is directly connected, GigabitEthernet0/0
L  10.1.1.1/32  is directly connected, GigabitEthernet0/0

R2#show ip bgp
       Network            Next-Hop      Metric      LocPrf      Weight      Path
*>     3.3.3.20/30         0.0.0.0           0                   32768      ?
*>     192.168.20.0/24     0.0.0.0           0                   32768      ?

Le due prefix sorgenti dai rispettivi nodi R1 ed R3 NON sono presenti in tabella di routing
regolarmente a causa della regola dello split-horizon imposto da CE2 e CE3.

R2# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.21    4    500        23          23         3       0      0   00:17:33            0
```

Tabella 9: RIB and BGP Table from R2 without BGP Confederation

RIB and BGP table outputs from **CE2** senza BGP Confederation

```
CE2#show ip route
Gateway of last resort is not set
   3.0.0.0./8  is variable subnetted, 5 subnets, 2 masks
C  3.3.3.4/30  is directly connected, GigabitEthernet0/2
L  3.3.3.5/32  is directly connected, GigabitEthernet0/2
B  3.3.3.8/30 [200/0] via 3.3.3.6, 02:14:34
C  3.3.3.16/30  is directly connected, GigabitEthernet0/3
L  3.3.3.17/32   is directly connected, GigabitEthernet0/3
   172.16.0.0/24 is subnetted, 1 subnet
B  172.16.10.0 [200/0] via 3.3.3.6, 02:14:34
B 192.168.20.0/24 [200/0] via 3.3.3.18, 00:21:35
```

```
CE2#show ip bgp
          Network          Next-Hop      Metric      LocPrf      Weight      Path
r*>i      3.3.3.4/30        3.3.3.6           0         100           0      ?
*> i      3.3.3.8/30        3.3.3.6           0         100           0      ?
*> i      172.16.10.0/24    3.3.3.6           0         100           0      ?
*> i      192.168.20.0/24   3.3.3.18          0         100           0      ?
```

Il CE2 per la sua posizione intermedia tra R3 ed R1 avendo rispettivamente sessioni IBGP con entrambi ha nella sua tabella le prefix annunciate dai medesimi peers,senza però avere il ruolo di re-annunciarle verso altri/opposti neighbors IBGP presenti.

```
CE2# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.6     4    500       170         167         5       0      0   02:29:53            3
3.3.3.18    4    500       168         166         5       0      0   02:29:06            1
```

Tabella 10: RIB and BGP Table from CE2 without BGP Confederation

RIB and BGP table outputs from **CE3** senza BGP Confederation

```
CE3#show ip route
Gateway of last resort is not set
   3.0.0.0./8  is variable subnetted, 5 subnets, 2 masks
B  3.3.3.4/30 [200/0] via 3.3.3.10, 02:30:31
C  3.3.3.8/30  is directly connected, GigabitEthernt0/2
L  3.3.3.9/32  is directly connected, GigabitEthernet0/2
C  3.3.3.20/30  is directly connected, GigabitEthernet0/3
L  3.3.3.21/32   is directly connected, GigabitEthernet0/3
   10.0.0.0/24  is subnetted, 1 subnet
B  10.1.1.0/24 [200/0]via 3.3.3.22, 00:34:09
   172.16.0.0/24 is subnetted, 1 sunbet
B  172.16.10.0/24 [200/0] via 3.3.3.10, 02:30:31
```

```
CE3#show ip bgp
          Network          Next-Hop      Metric      LocPrf      Weight      Path
*> i      3.3.3.4/30       3.3.3.10           0         100           0      ?
r> i      3.3.3.8/30       3.3.3.10           0         100           0      ?
r> i      3.3.3.20/30      3.3.3.22           0         100           0      ?
*> i      10.1.1.0/24      3.3.3.22           0         100           0      ?
*> i      172.16.10.0/24   3.3.3.10           0         100           0      ?
```

Il CE3 per la sua posizione intermedia tra R3 ed R2 avendo rispettivamente sessioni IBGP con entrambi ha nella sua tabella le prefix annunciate dai medesimi peers,senza però avere il ruolo di re-annunciarle verso altri/opposti neighbors IBGP presenti.

```
CE3# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
3.3.3.10    4    500       180         176         6       0      0   02:38:53            3
3.3.3.22    4    500        50          51         6       0      0   00:42:20            2
```

Tabella 11: RIB and BGP Table from CE3 without BGP Confederation

IBGP peering with BGP Confederation

Consideriamo la seguente architettura:

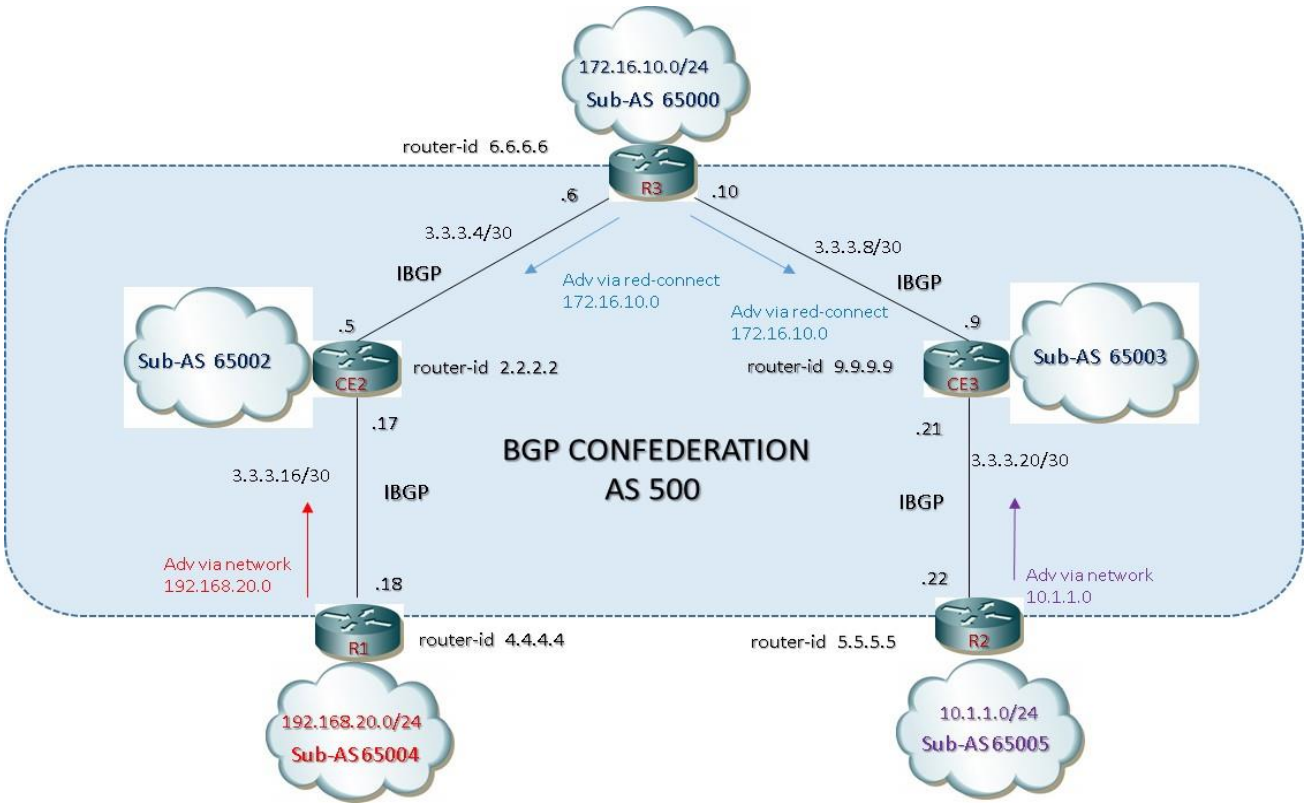

Figura 4: IBGP peering with advertisement permitted (BGP CONFEDERATION)

Quarta configurazione di base (BGP CONFEDERATION):

R3

```
router bgp 65000
 bgp router-id 6.6.6.6
 bgp log-neighbor-changes
 bgp confederation identifier 500
 bgp confederation peer 65002 65003
 neighbor 2.2.2.2 remote-as 65002
 neighbor 2.2.2.2 ttl-security hops 5
 neighbor 2.2.2.2 update-source loopback0
 neighbor 2.2.2.2 version 4
 neighbor 9.9.9.9 remote-as 65003
 neighbor 9.9.9.9 ttl-security hops 5
 neighbor 9.9.9.9 update-source loopback0
 neighbor 9.9.9.9 version 4
!
 address-family ipv4
  redistribute connected
  neighbor 2.2.2.2 activate
  neighbor 9.9.9.9 activate
!
```

```
router ospf 100
network 3.3.3.4 0.0.0.3 area 0
network 3.3.3.8 0.0.0.3 area 0
network 6.6.6.6 0.0.0.0 area 0
```

In alternativa con l'ausilio di peer-session:

```
router bgp 65000
 template peer-session IBGP-65002
 remote-as 65002
 ttl-security hops 5
 update-source loopback0
!
template peer-session IBGP-65003
 remote-as 65003
 ttl-security hops 5
 update-source loopback0
!
bgp router-id 6.6.6.6
bgp log-neighbor-changes
bgp confederation identifier 500
bgp confederation peers 65002 65003
neighbor 2.2.2.2 inherit peer-session 65002
neighbor 9.9.9.9 inherit peer-session 65003
!
address-family ipv4
redistribute connected
neighbor 2.2.2.2 activate
neighbor 9.9.9.9 ctivate
```

CE2

```
router bgp 65002
 bgp router-id 2.2.2.2
 bgp log-neighbor-changes
 bgp confederation identifier 500
 bgp confederation peer 65000 65004
 neighbor 4.4.4.4 remote-as 65004
 neighbor 4.4.4.4 ttl-security hops 5
 neighbor 4.4.4.4 update-source loopback0
 neighbor 4.4.4.4 version 4
 neighbor 6.6.6.6 remote-as 65000
 neighbor 6.6.6.6 ttl-security hops 5
 neighbor 6.6.6.6 update-source loopback0
 neighbor 6.6.6.6 version 4
!
 address-family ipv4
  neighbor 4.4.4.4 activate
  neighbor 6.6.6.6 activate
!
router ospf 100
network 2.2.2.2 0.0.0.0 area 0
network 3.3.3.4 0.0.0.3 area 0
network 3.3.3.16 0.0.0.0 area 0
```

In alternativa con l'ausilio di peer-session:

```
router bgp 65002
 template peer-session IBGP-65000
 remote-as 65000
 ttl-security hops 5
 update-source loopback0
!
template peer-session IBGP-65004
 remote-as 65004
 ttl-security hops 5
 update-source loopback0
!
bgp router-id 2.2.2.2
bgp log-neighbor-changes
bgp confederation identifier 500
bgp confederation peers 65000 65004
neighbor 4.4.4.4 inherit peer-session 65004
neighbor 6.6.6.6 inherit peer-session 65000
!
address-family ipv4
neighbor 4.4.4.4 activate
neighbor 6.6.6.6 activate
```

CE3

```
router bgp 65003
 bgp router-id 9.9.9.9
 bgp log-neighbor-changes
 bgp confederation identifier 500
 bgp confederation peer 65000 65005
 neighbor 5.5.5.5 remote-as 65005
 neighbor 5.5.5.5 ttl-security hops 5
 neighbor 5.5.5.5 update-source loopback0
 neighbor 5.5.5.5 version 4
 neighbor 6.6.6.6 remote-as 65000
 neighbor 6.6.6.6 ttl-security hops 5
 neighbor 6.6.6.6 update-source loopback0
 neighbor 6.6.6.6 version 4
!
 address-family ipv4
  neighbor 5.5.5.5 activate
  neighbor 6.6.6.6 activate
!
router ospf 100
network 3.3.3.8 0.0.0.3 area 0
network 3.3.3.20 0.0.0.0 area 0
network 9.9.9.9 0.0.0.0 area 0
```

In alternativa con l'ausilio di peer-session:

```
router bgp 65003
 template peer-session IBGP-65000
 remote-as 65000
 ttl-security hops 5
 update-source loopback0
!
```

```
template peer-session IBGP-65005
 remote-as 65005
 ttl-security hops 5
 update-source loopback0
!
bgp router-id 9.9.9.9
bgp log-neighbor-changes
bgp confederation identifier 500
bgp confederation peers 65000 65005
neighbor 5.5.5.5 inherit peer-session 65005
neighbor 6.6.6.6 inherit peer-session 65000
!
address-family ipv4
neighbor 5.5.5.5 activate
neighbor 6.6.6.6 activate
```

R1

```
router bgp 65004
 bgp router-id 4.4.4.4
 bgp log-neighbor-changes
 bgp confederation identifier 500
 bgp confederation peer 65002
 neighbor 2.2.2.2 remote-as 65002
 neighbor 2.2.2.2 ttl-security hops 5
 neighbor 2.2.2.2 update-source loopback0
 neighbor 2.2.2.2 version 4
!
 address-family ipv4
  network 192.168.20.0
  neighbor 2.2.2.2 activate
!
router ospf 100
 network 3.3.3.16 0.0.0.0 area 0
 network 4.4.4.4 0.0.0.0 area 0
```

In alternativa con l'ausilio di peer-session:

```
router bgp 65004
 template peer-session IBGP-65002
 remote-as 65002
 ttl-security hops 5
 update-source loopback0
!

bgp router-id 4.4.4.4
bgp log-neighbor-changes
bgp confederation identifier 500
bgp confederation peers 65002
neighbor 2.2.2.2 inherit peer-session 65002
!
address-family ipv4
network 192.168.20.0
neighbor 2.2.2.2 activate
```

R2

```
router bgp 65005
 bgp router-id 5.5.5.5
 bgp log-neighbor-changes
 bgp confederation identifier 500
 bgp confederation peer 65003
 neighbor 9.9.9.9 remote-as 65003
 neighbor 9.9.9.9 ttl-security hops 5
 neighbor 9.9.9.9 update-source loopback0
 neighbor 9.9.9.9 version 4
!
 address-family ipv4
  network 10.1.1.0 mask 255.255.255.0
  neighbor 9.9.9.9 activate
!
router ospf 100
 network 3.3.3.20 0.0.0.0 area 0
 network 5.5.5.5 0.0.0.0 area 0


In alternativa con l'ausilio di peer-session:


router bgp 65005
 template peer-session IBGP-65003
 remote-as 65003
 ttl-security hops 5
 update-source loopback0
!

bgp router-id 5.5.5.5
bgp log-neighbor-changes
bgp confederation identifier 500
bgp confederation peers 65003
neighbor 9.9.9.9 inherit peer-session 65003
!
address-family ipv4
network 10.1.1.0 mask 255.255.255.0
neighbor 9.9.9.9 activate
```

RIB and BGP table outputs from **R3** con BGP Confederation

```
R3#show ip route
Gateway of last resort is not set
   2.0.0.0/32   is subnetted, 1 subnet
O  2.2.2.2/32 [110/2] via 3.3.3.5, 00:30:43, GigabitEthernet0/2
   4.0.0.0/32   is subnetted, 1 subnet
O  4.4.4.4/32 [110/3] via 3.3.3.5, 00:30:44, GigabitEthernet0/2
   5.0.0.0/32   is subnetted, 1 subnet
O  5.5.5.5/32 [110/3] via 3.3.3.9, 00:19:41, GigabitEthernet0/1
   6.0.0.0/32  is subnetted, 1 subnet
C  6.6.6.6/32  is directly connected, Loopback0
   9.0.0.0/32   is subnetted, 1 subnet
O  9.9.9.9/32 [110/2] via 3.3.3.9, 00:25:41, GigabitEthernet0/1
   10.0.0.0/24  is subnetted, 1 subnet
B  10.1.1.0 [200/0] via 5.5.5.5, 00:17:33
   172.16.0.0  is variable subnetted, 2 subnet, 2 masks
C  172.16.10.0/24 is directly connected, GigabitEthernet0/0
L  172.16.10.1/32 is directly connected, GigabitEthernet0/0
B  192.168.20.0/24 [200/0] via 4.4.4.4, 00:21:00


R3#show ip bgp
       Network          Next-Hop      Metric      LocPrf      Weight      Path
*>     3.3.3.4/30       0.0.0.0            0                   32768      ?
*>     3.3.3.8/30       0.0.0.0            0                   32768      ?
*>     6.6.6.6/32       0.0.0.0            0                   32768      ?
*>     10.1.1.0/24      5.5.5.5            0         100           0      (65003 65005) i
*>     172.16.10.0/24   0.0.0.0            0                   32768      ?
*>     192.168.20.0     4.4.4.4            0          100          0      (65002 65004) i


R3# show ip bgp summary

Neighbor    V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
2.2.2.2     4   65002      53          54         9        0      0   00:42:57            1
9.9.9.9     4   65003      47          49         9        0      0   00:37:57            1
```

```
R3# show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 9
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  (65003 65005)
  5.5.5.5 (metric 3) from 9.9.9.9 (9.9.9.9)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0



R3# show ip bgp 192.168.20.0
BGP routing table entry for 192.168.20.0/24, version 7
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  (65002 65004)
  4.4.4.4 (metric 3) from 2.2.2.2 (2.2.2.2)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0



R3# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 3
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
  0.0.0.0 from 0.0.0.0 (6.6.6.6)
    Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
    rx pathid: 0,  tx pathid: 0x0
```

Tabella 12: RIB and BGP Table from R3 with BGP Confederation

RIB and BGP table outputs from **CE2** con BGP Confederation

```
CE2#show ip route
Gateway of last resort is not set
   2.0.0.0/32  is subnetted, 1 subnet
C  2.2.2.2/32  is directly connected, Loopback0
   4.0.0.0/32   is subnetted, 1 subnet
O  4.4.4.4/32 [110/2] via 3.3.3.18, 00:44:59, GigabitEthernet0/3
   5.0.0.0/32   is subnetted, 1 subnet
O  5.5.5.5/32 [110/4] via 3.3.3.6, 00:41:48, GigabitEthernet0/2
   6.0.0.0/32  is subnetted, 1 subnet
C  6.6.6.6/32 [110/2] via 3.3.3.6, 00:52:47, GigabitEthernet0/2
   9.0.0.0/32   is subnetted, 1 subnet
O  9.9.9.9/32 [110/3] via 3.3.3.6, 00:47:38, GigabitEthernet0/2
    10.0.0.0/24  is subnetted, 1 subnet
B   10.1.1.0 [200/0] via 5.5.5.5, 00:39:40
    172.16.0.0  is subnetted, 1 subnet
B   172.16.10.0/24 [200/0] via 6.6.6.6, 00:52:37
B  192.168.20.0/24 [200/0] via 4.4.4.4, 00:43:07
```

```
CE2#show ip bgp
       Network          Next-Hop    Metric      LocPrf      Weight      Path
r>     3.3.3.4/30       6.6.6.6          0         100           0      (65000) ?
r>     3.3.3.8/30       6.6.6.6          0         100           0      (65000) ?
r>     6.6.6.6/32       6.6.6.6          0         100           0      (65000) ?
*>     10.1.1.0/24      5.5.5.5          0         100           0      (65000 65003 65005) i
*>     172.16.10.0/24   6.6.6.6          0         100           0      (65000) ?
*>     192.168.20.0     4.4.4.4          0         100           0      (65004) i
```

```
CE2# show ip bgp summary
Neighbor     V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
4.4.4.4      4   65004      64          68        13        0      0   00:54:07            1
6.6.6.6      4   65000      76          74        13        0      0   01:02:37            5
```

```
CE2# show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 13
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  (65000 65003 65005)
  5.5.5.5 (metric 3) from 6.6.6.6 (6.6.6.6)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0


CE2# show ip bgp 192.168.20.0
BGP routing table entry for 192.168.20.0/24, version 11
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  (65004)
  4.4.4.4 (metric 3) from 4.4.4.4 (4.4.4.4)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0


CE2# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 9
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  (65005)
  6.6.6.6 (metric 3) from 6.6.6.6 (6.6.6.6)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0
```

Tabella 13: RIB and BGP Table from CE2 with BGP Confederation

RIB and BGP table outputs from **CE3** con BGP Confederation

```
CE3#show ip route
Gateway of last resort is not set
   2.0.0.0/32  is subnetted, 1 subnet
O  2.2.2.2/32 [110/3] via 3.3.3.10, 01:03:55, GigabitEthernet0/2
   4.0.0.0/32   is subnetted, 1 subnet
O  4.4.4.4/32 [110/4] via 3.3.3.10, 01:01:11, GigabitEthernet0/2
   5.0.0.0/32   is subnetted, 1 subnet
O  5.5.5.5/32 [110/2] via 3.3.3.22, 00:58:00, GigabitEthernet0/3
   6.0.0.0/32  is subnetted, 1 subnet
C  6.6.6.6/32 [110/2] via 3.3.3.10, 01:03:55, GigabitEthernet0/2
   9.0.0.0/32  is subnetted, 1 subnet
O  9.9.9.9/32  is directly connected, Loopback0
    10.0.0.0/24  is subnetted, 1 subnet
B   10.1.1.0 [200/0] via 5.5.5.5, 00:55:52
    172.16.0.0  is subnetted, 1 subnet
B   172.16.10.0/24 [200/0] via 6.6.6.6, 01:03:49
B  192.168.20.0/24 [200/0] via 4.4.4.4, 00:59:19


CE3#show ip bgp
       Network          Next-Hop    Metric      LocPrf      Weight      Path
r>     3.3.3.4/30       6.6.6.6          0         100           0      (65000) ?
r>     3.3.3.8/30       6.6.6.6          0         100           0      (65000) ?
r>     6.6.6.6/32       6.6.6.6          0         100           0      (65000) ?
*>     10.1.1.0/24      5.5.5.5          0         100           0      (65005) i
*>     172.16.10.0/24   6.6.6.6          0         100           0      (65000) ?
*>     192.168.20.0     4.4.4.4          0         100           0      (65000 65002 65004) i


CE3# show ip bgp summary

Neighbor     V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
5.5.5.5      4   65005      75          79        13       0      0   01:04:21            1
6.6.6.6      4   65000      86          83        13       0      0   01:11:17            5
```

```
CE3# show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 13
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  (65005)
  5.5.5.5 (metric 2) from 5.5.5.5 (5.5.5.5)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0



CE3# show ip bgp 192.168.20.0
BGP routing table entry for 192.168.20.0/24, version 11
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 2
  (65000 65002 65004)
  4.4.4.4 (metric 3) from 6.6.6.6 (6.6.6.6)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0



CE3# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 9
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  (65000)
  6.6.6.6 (metric 3) from 6.6.6.6 (6.6.6.6)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0
```

Tabella 14: RIB and BGP Table from CE3 with BGP Confederation

RIB and BGP table outputs from **R1** con BGP Confederation

```
R1#show ip route
Gateway of last resort is not set
   2.0.0.0/32  is subnetted, 1 subnet
O  2.2.2.2/32 [110/2] via 3.3.3.17, 01:14:06, GigabitEthernet0/3
   4.0.0.0/32  is subnetted, 1 subnet
O  4.4.4.4/32  is directly connected, Loopback0
   5.0.0.0/32   is subnetted, 1 subnet
O  5.5.5.5/32 [110/5] via 3.3.3.17, 01:10:48, GigabitEthernet0/3
   6.0.0.0/32  is subnetted, 1 subnet
C  6.6.6.6/32 [110/3] via 3.3.3.17, 01:14:06, GigabitEthernet0/3
   9.0.0.0/32  is subnetted, 1 subnet
O  9.9.9.9/32 [110/4] via 3.3.3.17, 01:14:06, GigabitEthernet0/3
    10.0.0.0/24  is subnetted, 1 subnet
B   10.1.1.0 [200/0] via 5.5.5.5, 01:08:40
    172.16.0.0  is subnetted, 1 subnet
B   172.16.10.0/24 [200/0] via 6.6.6.6, 01:13:07
    192.168.20.0/24 is variable subnetted, 2 subnet, 2 masks
C   192.168.20.0/24 is directly connected, GigabitEthernet0/0
L   192.168.20.1/32 is directy connected, GigabitEthernet0/0


R1#show ip bgp
       Network          Next-Hop    Metric     LocPrf     Weight     Path
r>     3.3.3.4/30       6.6.6.6          0        100          0     (65002 65000) ?
r>     3.3.3.8/30       6.6.6.6          0        100          0     (65002 65000) ?
r>     6.6.6.6/32       6.6.6.6          0        100          0     (65002 65000) ?
*>     10.1.1.0/24      5.5.5.5          0        100          0     (65002 65000 65003 65005) i
*>     172.16.10.0/24   6.6.6.6          0        100          0     (65002 65000) ?
*>     192.168.20.0     0.0.0.0          0                 32768     (65000 65002 65004) i


R1# show ip bgp summary

Neighbor     V     AS   MsgRcvd     MsgSent     TblVer     InQ   OutQ   Up/Down     State/PfxRcd
2.2.2.2      4   65002      97          94         12         0      0   01:21:10            5
```

```
R1# show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 12
Paths: (1 available, best #1, table default)
  Not advertised to any peer:
  Refresh Epoch 2
  (65002 65000 65003 65005)
  5.5.5.5 (metric 2) from 2.2.2.2 (2.2.2.2)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0



R1# show ip bgp 192.168.20.0
BGP routing table entry for 192.168.20.0/24, version 11
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
      1
  Refresh Epoch 2
  Local
  0.0.0.0 from 0.0.0.0 (4.4.4.4)
    Origin IGP, metric 0, localpref 100, valid, sourced, local, best
    rx pathid: 0,  tx pathid: 0x0



R1# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 9
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  (65002 65000)
  6.6.6.6 (metric 3) from 2.2.2.2 (2.2.2.2)
    Origin incomplete, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0
```

Tabella 15: RIB and BGP Table from R1 with BGP Confederation

RIB and BGP table outputs from **R2** con BGP Confederation

```
R2#show ip route
Gateway of last resort is not set
   2.0.0.0/32  is subnetted, 1 subnet
O  2.2.2.2/32 [110/4] via 3.3.3.21, 01:25:37, GigabitEthernet0/3
   4.0.0.0/32  is subnetted, 1 subnet
O  4.4.4.4/32 [110/5] via 3.3.3.21, 01:25:38, GigabitEthernet0/3
   5.0.0.0/32  is subnetted, 1 subnet
O  5.5.5.5/32  is directly connected, Loopback0
   6.0.0.0/32  is subnetted, 1 subnet
C  6.6.6.6/32 [110/3] via 3.3.3.21, 01:25:38, GigabitEthernet0/3
   9.0.0.0/32  is subnetted, 1 subnet
O  9.9.9.9/32 [110/2] via 3.3.3.21, 01:25:37, GigabitEthernet0/3
    10.0.0.0/8  is variable subnetted, 2 subnets, 2 masks
C   10.1.1.0/24 is directly connected, GigabitEthernet0/0
L   10.1.1.1/32 is directy connected, GigabitEthernet0/0
    172.16.0.0  is subnetted, 1 subnet
B   172.16.10.0/24 [200/0] via 6.6.6.6, 01:24:19
B   192.168.20.0/24 [200/0] via 4.4.4.4, 01:24:19


R2#show ip bgp
       Network          Next-Hop    Metric     LocPrf     Weight     Path
r>     3.3.3.4/30       6.6.6.6          0        100          0    (65003 65000) ?
r>     3.3.3.8/30       6.6.6.6          0        100          0    (65003 65000) ?
r>     6.6.6.6/32       6.6.6.6          0        100          0    (65003 65000) ?
*>     10.1.1.0/24      0.0.0.0          0                 32768     i
*>     172.16.10.0/24   6.6.6.6          0        100          0    (65003 65000) ?
*>     192.168.20.0     4.4.4.4          0        100          0    (65003 65000 65002 65004) i


R2# show ip bgp summary

Neighbor     V     AS   MsgRcvd     MsgSent    TblVer     InQ   OutQ   Up/Down     State/PfxRcd
9.9.9.9      4   65003       109         105        12        0      0   01:31:41            5
```

```
R2# show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 12
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  Local
  0.0.0.0 from 0.0.0.0 (5.5.5.5)
    Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best
    rx pathid: 0,  tx pathid: 0x0



R2# show ip bgp 192.168.20.0
BGP routing table entry for 192.168.20.0/24, version 11
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  (65003 65000 65002 65004)
  4.4.4.4 (metric 5) from 9.9.9.9 (9.9.9.9)
    Origin IGP, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0



R2# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 10
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  (65003 65000)
  6.6.6.6 (metric 3) from 9.9.9.9 (9.9.9.9)
    Origin incomplete, metric 0, localpref 100, valid, confed-external, best
    rx pathid: 0,  tx pathid: 0x0
```

Tabella 16: RIB and BGP Table from R2 with BGP Confederation