INDICE

1 Introduzione

Il presente documento si propone di definire alcuni parametri del BGP e le sue applicazioni in modo teorico accompagnati da configurazioni e verifiche pratiche attraverso il laboratorio.

2 BGP Convergence Mode

BGP annuncia le sue routes con una metrica pari alla lunghezza del suo AS-PATH attribute (path-vector protocol); come un qualsiasi protocollo distance vector, il processo di BGP routing accetta in ingresso multipli update di routing ed annuncia solo il best-path per ogni determinata prefix ai suoi neighbors.

Anziché utilizzare periodici update, BGP utilizza " explicit withdraw " all'interno di update per segnalare la perdita di un particolare path; oltre a questi " explicit withdraw " BGP utilizza un " implicit signaling " dove poter segnalare aggiornamenti riguardo la stessa prefix e poter sostituire cosi le informazioni precedentemente ricevute.

I principali fattori che possono limitare la convergenza BGP sono la grandezza della tabella BGP, il livello di trasporto settato (Tuning BGP transport), e il valore di advertisement delay

Il processo di selezione del best-path è proporzionale, quindi, alla quantità di contenuti in tabella BGP cosi come al tempo richiesto per l'aggiornamento di questi update.

BGP explicit withdraw

Numero di volte che una Prefix è ritirata (withdraw) causa mancanza di fattibilità o sua raggiungibilità

BGP implicit withdraw

Numero di volte che una Prefix è stata ritirata ritirata e re-annunciata; da chiarire che:

- ✓ Quando un BGP sender annuncia una route (percorso) verso un indirizzo precedentemente irragiungibile significa che questo è un nuovo annuncio;
- ✓ Quando un BGP sender annuncia una route (percorso) verso un indirizzo attualmente raggiungibile e questo percorso è identico al percorso corrente significa che questo annuncio è duplicato oppure sta sostituendo il percorso attuale con uno nuovo e questo è definito " ritiro implicito " (implicit withdraw)

BGP UPDATE-DELAY

Appena due nodi stabiliscono una sessione BGP scambiandosi OPEN messages, uno dei due entra in una modalità detta " BGP Read-Only" e non inizierà nessun best-path process sino a quando non avrà ricevuto tutti gli annunci interessati oppure raggiunge il valore di timeout del BGP Read-Only definito dal comando " bgp update-delay "

Il processo BGP determina la fine degli Update messages in due modi:

- ✓ attraverso la ricezione di Keepalive messages
- ✓ attraverso la ricezione di End-of-RIB messages (anche usato per BGP graceful restart)

BGP ROUTER PROCESS

Appena un nodo cessa la modalità Read-Only", allora inizia il processo di selezione best-path conosciuto come BGP Routing Process ed il suo contenuto compare all'interno della Local-BGP-RIB, con un best-path per ciascuna prefix annunciata.

Appena terminata la computazione del best-path process, il nodo aggiorna tutte le routes nella sua RIB (Routing Information Base) prima di annunciarle ai suoi peer neighbors (questo è un prerequisito di tutti i distance-vector protocol); la RIB quindi aggiornerà la sua FIB (Forwarding Information Base) per la trasmissione attraverso le sue line-card.

BGP DYNAMIC UPDATE GROUP

Utilizzato nel processo di replicazione del best-path, una volta aggiornata la RIB per tutti i peers interessati a ricevere l'annuncio (questo passaggio è molto a carico della CPU memory di un nodo) il nodo deve performare la tabella BGP con i suoi contenuti per ciascun neighbors e costruire una corrispondenza BGP to Adj-RIB-Out

Con l'implementazione " dynamic update group " è possibile dinamicamente trovare i rispettivi neighbors e condividerli all'interno di una output policy, eleggere un peer con l'indirizzo più basso del gruppo come leader del gruppo stesso e generare tutti gli update verso questo group-leader; tutti gli altri nodi appartenenti al gruppo riceveranno gli stessi annunci.

Il BGP update group trova impiego nei Router Reflector dove questi hanno centinaia di peering BGP che condividono le stesse policy; è possibile verificare gli update group con il comando "show ip bgp replication".

BGP ADVERTISEMENT INTERVAL

E' un parametro utilizzato per prevenire instabilità tra peers all'interno della rete a cui fanno parte a causa di una eccessiva trasmissione di messaggi update tra loro per aggiornamenti vari.

Pertanto un nodo prima di iniziare a trasmettere messaggi update controlla se il valore di timer è in corso per un dato peer e qualora fosse attivo, ritarda la trasmissione sino a quando questo valore di timer è finito.

Il comando che definisce questo timer è " neighbor <ip_addr_neighbor> advertisment-interval <times> " ; il valore di default è pari a 30 sec per EBGP (Inter-AS) ed invece pari a 5 sec IBGP (Intra-AS)

Nella implementazione Cisco seguire i seguenti passi:

1) Il BGP ADD-PATH è possibile dalle versioni software IOS 15.3T, IOS XE 10S e IOS XE 4.0.0
2) Specificare se i nodi devono negoziare la funzionalità ADD-PATH CAPABILITY e la modalità:
    a. Send
    b. Reveive
    c. Send / Receive

3) Selezionare l'insieme per percorsi da annunciare altre al best-path
4) Specificare per ciascun nodo, l'insieme dei percorsi da annunciare

https://tools.ietf.org/html/draft-ietf-idr-add-paths-guidelines-07

3 Architettura di riferimento V-LAB

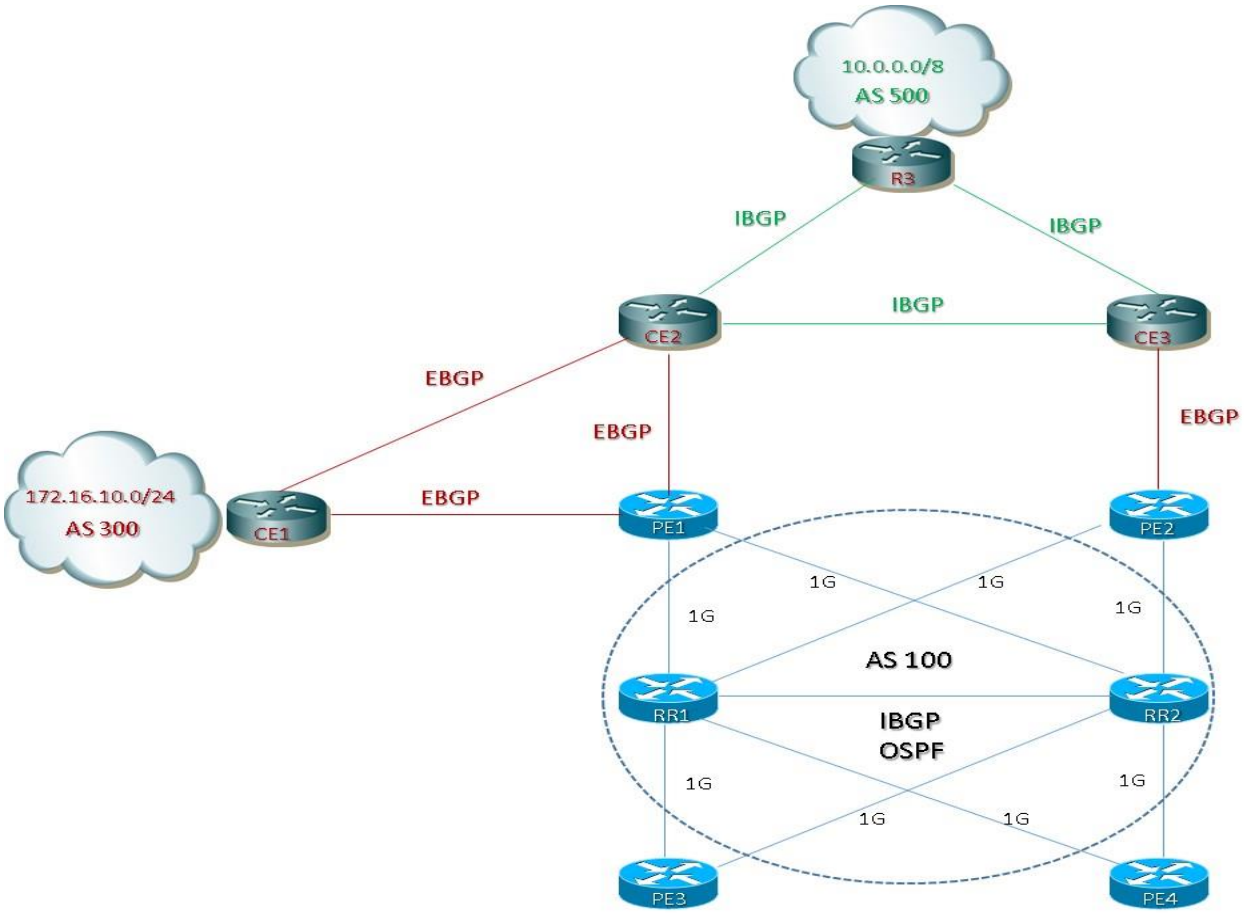

***Figura** 1: architettura di riferimento VLAB BGP*

Prendiamo in considerazione questa prima parte di rete evidenziata in giallo dalla figura seguente:

Obiettivo: redistribuzione della reti external tra gli AS300 ed AS500 attraverso i links EBGP e l'advertisement tra EBGP to IBGP sia in tabella BGP che in tabella di routing presso i routers CE3 e R3.

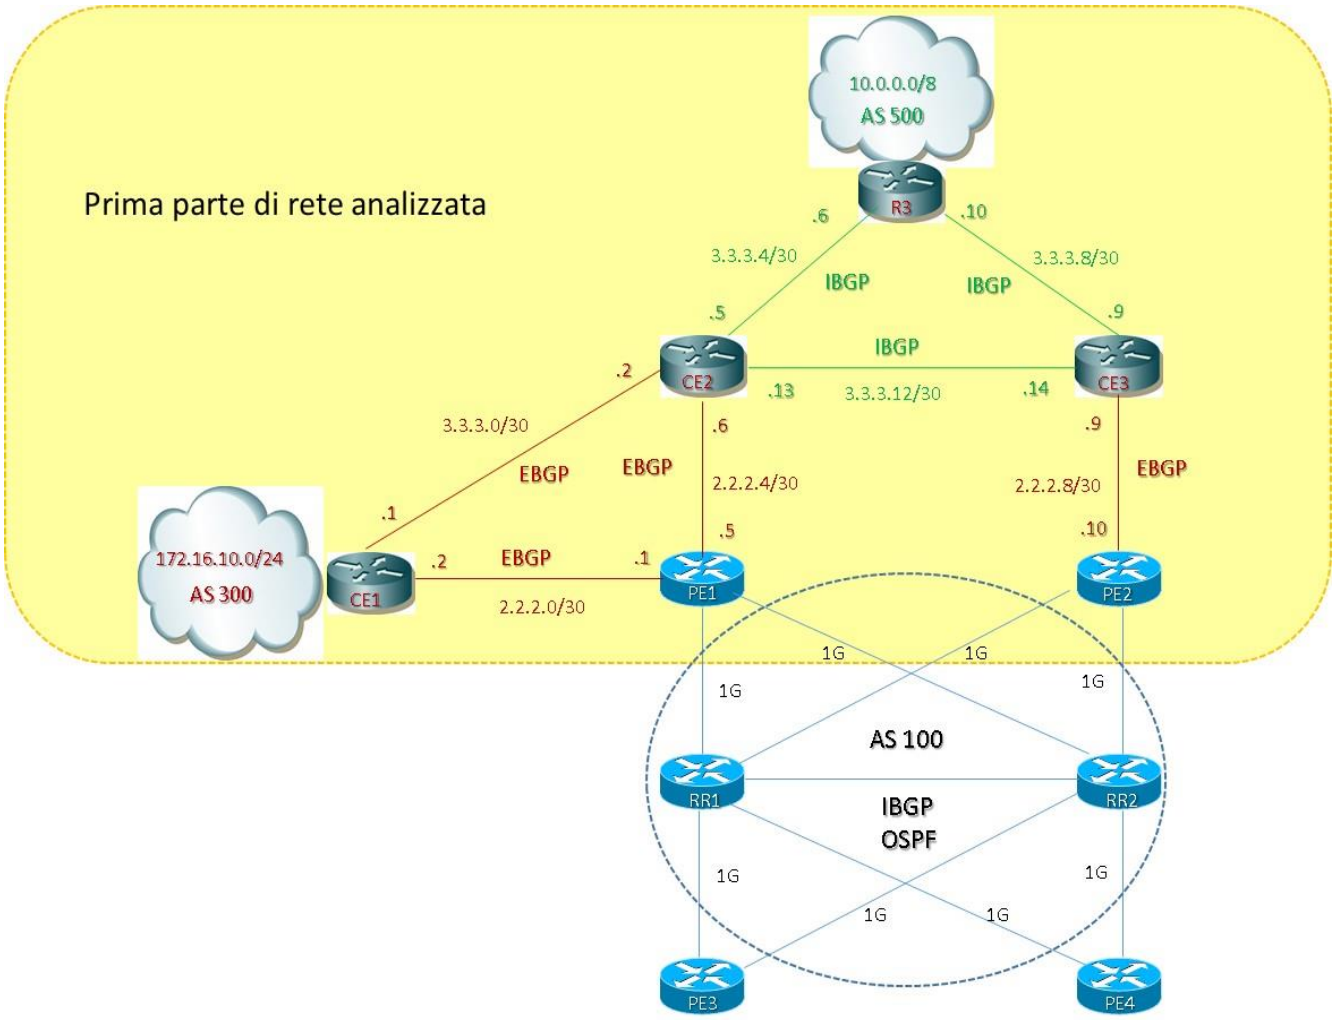

Figura 2: *architettura di riferimento VLAB prima parte analizzata*

Configurazione BGP di base prima parte analizzata

CE1

```
router bgp 300
 bgp log-neighbor-changes
 neighbor 2.2.2.1 remote-as 100
 neighbor 3.3.3.2 remote-as 500
!
 address-family ipv4
  network 172.16.10.0 mask 255.255.255.0
  neighbor 2.2.2.1 activate
  neighbor 3.3.3.2 activate
```

CE2

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 2.2.2.5 remote-as 100
 neighbor 3.3.3.1 remote-as 300
 neighbor 3.3.3.6 remote-as 500
 neighbor 3.3.3.14 remote-as 500
!
 address-family ipv4
  neighbor 2.2.2.5 activate
  neighbor 3.3.3.1 activate
  neighbor 3.3.3.6 activate
  neighbor 3.3.3.14 activate
```

CE3

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 2.2.2.10 remote-as 100
 neighbor 3.3.3.10 remote-as 500
 neighbor 3.3.3.13 remote-as 500
!
 address-family ipv4
  neighbor 2.2.2.10 activate
  neighbor 3.3.3.10 activate
  neighbor 3.3.3.13 activate
```

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
```

PE1

```
router bgp 100
 bgp log-neighbor-changes
 neighbor 2.2.2.2 remote-as 300
 neighbor 2.2.2.6 remote-as 500
!
 address-family ipv4
  neighbor 2.2.2.2 activate
  neighbor 2.2.2.6 activate
```

PE2

```
router bgp 100
 bgp log-neighbor-changes
 neighbor 2.2.2.9 remote-as 500
!
 address-family ipv4
  neighbor 2.2.2.9 activate
```

A seguito di questa configurazione base del BGP abbiamo i seguenti outputs dai rispettivi routers CEs per le tabelle di routing (RIB) e la tabelle BGP ipv4

RIB and BGP table prima parte analizzata

RIB and BGP table outputs from **CE1 (AS300)**

```
CE1#show ip route

Gateway of last resort is not set

B   10.0.0.0/8 [20/0] via 3.3.3.2,  03:39:31

    172.16.10.0 is variable subnetted, 2 subnet, 2 masks

C   172.16.10.0/24 is directly connected, GigabitEthernet0/1

L   172.16.10.1/32 is directly connected, GigabitEthernet0/1

!
!
CE1#show ip bgp

     Network          Next-Hop      Metric      LocPrf      Weight      Path

*    10.0.0.0         2.2.2.1                                    0      100 500 ?

*>                    3.3.3.2                                    0      500 ?

*>   172.16.10.0/24   0.0.0.0           0                    32768      i

!
```

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

Commento:

1) Le due reti 10/8 è 172.16.10.0 sono presenti sia in RIB che in BGP table;
2) La rete 10/8 è vista con due next hop EBGP di cui quello best-path è il 3.3.3.2 perché con un AS-PATH migliore;
3) La rete 172.16.10.0/24 è direttamente collegata per cui è vista con next-hop 0.0.0.0

Regolare comportamento del E-BGP

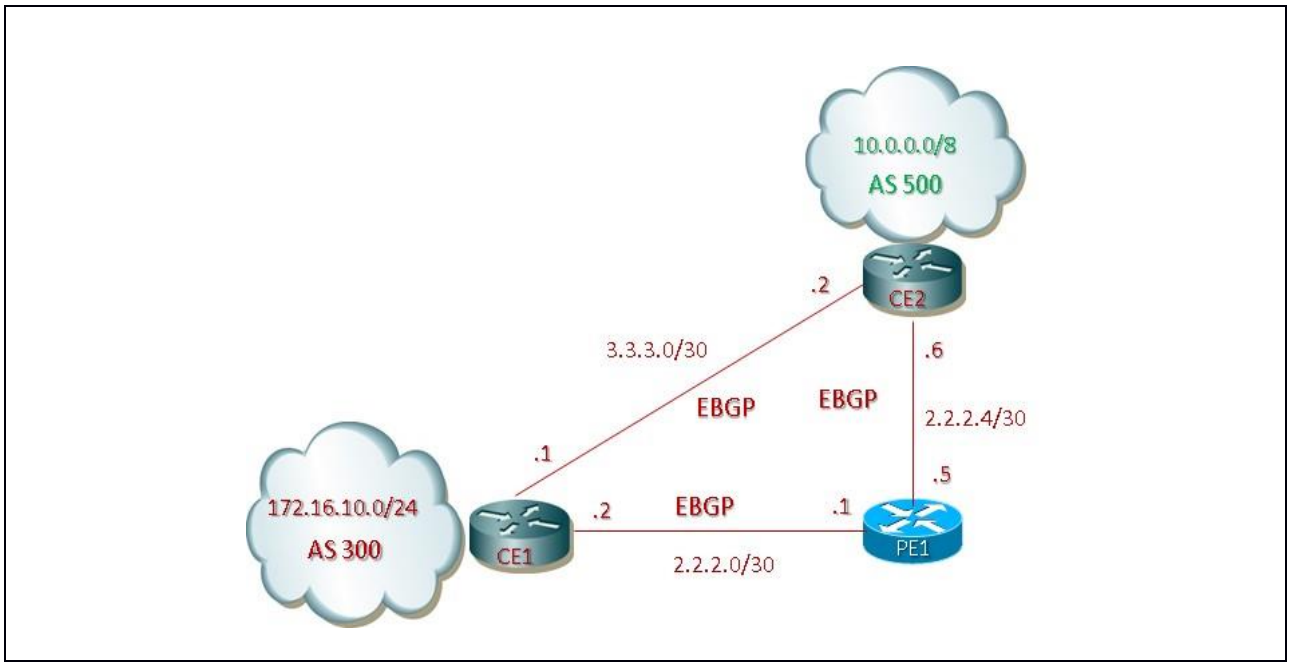

Tabella 1: RIB and BGP Table from CE1 prima parte analizzata

RIB and BGP table outputs from **CE2 (AS500)**

```
CE2#show ip route

Gateway of last resort is not set

B  10.0.0.0/8 [200/0] via 3.3.3.6,  04:02:55

   172.16.10.0/24  is subnetted,  1 subnet

B  172.16.10.0 [20/0] via 3.3.3.1, 04:04:56

!
!
CE2#show ip bgp

      Network          Next-Hop      Metric      LocPrf      Weight      Path

*>i   10.0.0.0         3.3.3.6            0         100           0      ?

*>    172.16.10.0/24   3.3.3.1            0                       0      300 i

*                      2.2.2.5                                    0      100 300 i

!
Commento:
```

1) La ret1 10/8 e 172.16.10.0 sono presenti sia in RIB che in BGP table;
2) La rete 10/8 è vista con next-hop via IBGP collegato al router R3 che annuncia la prefix (l'origin incomplete è perché ho usato il comando redistribute connected in R3);
3) La rete 172.16.10.0/24 è vista con due next-hop EBGP di cui uno con best-path via 3.3.3.1 poiché con AS-Path migliore

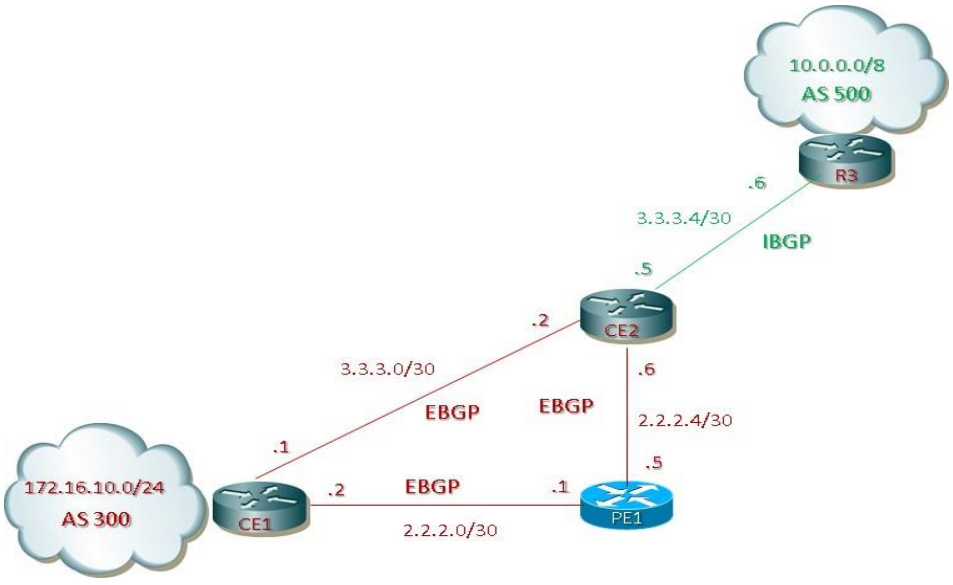

Tabella 2: RIB and BGP Table from CE2 prima parte analizzata

RIB and BGP table outputs from **CE3 (AS500)**

```
CE3#show ip route

Gateway of last resort is not set

B  10.0.0.0/8 [200/0] via 3.3.3.10,  03:44:14

!
CE3#show ip bgp
       Network          Next-Hop     Metric      LocPrf      Weight      Path

*>i    10.0.0.0         3.3.3.10          0         100           0      ?

* i    172.16.10.0/24   3.3.3.1           0         100           0      300 i

!
```

Commento:

Dal punto di vista CE3 dove non abbiamo un collegamento diretto tra i due AS300 ed AS500:

1) Solo la rete 10/8 è presente in RIB;
2) La rete 10/8 è vista il next-hop via IBGP collegato al router R3 che annuncia la prefix;
3) La rete 172.16.10.0/24 è vista solo in tabella BGP (con next-hop che ha originato l'annuncio della Prefix via EBGP (comportamento di default)
4) La rete 172.16.10.0/24 non viene installata nella RIB poiché è assente una route che permetta di conoscere la subnet con la quale viene annunciata la Prefix da CE1

```
CE3# show ip route 3.3.3.1
% subnet not in table
```

5) Questo è il caso di utilizzare il commando " next-hop-self " lato CE2 in quanto in questo modo il CE2 sostituisce il next-hop che ha originato la Prefix (comportamento di default BGP) con il proprio indirizzo di next-hop il quale è presente come routes nella tabella RIB di CE3 (stesso discorso vale per R3).

FROM CE2 CONFIG:

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 2.2.2.5 remote-as 100
 neighbor 3.3.3.1 remote-as 300
 neighbor 3.3.3.6 remote-as 500
 neighbor 3.3.3.14 remote-as 500
!
 address-family ipv4
  neighbor 2.2.2.5 activate
  neighbor 3.3.3.1 activate
  neighbor 3.3.3.6 activate
  neighbor 3.3.3.6 next-hop-self
  neighbor 3.3.3.14 activate
  neighbor 3.3.3.14 next-hop-self
```

Ripetiamo di nuovo gli outputs per la RIB e la BGP Table:

```
CE3#show ip route

Gateway of last resort is not set

B  10.0.0.0/8 [200/0] via 3.3.3.10,  03:44:14

   172.16.0.0/24 is subnetted, 1 subnet

B  172.16.10.0 [200/0] via 3.3.3.13, 00:00:24

!

CE3#show ip bgp

       Network          Next-Hop     Metric      LocPrf      Weight      Path

*>i    10.0.0.0         3.3.3.10          0         100           0      ?

* i    172.16.10.0/24   3.3.3.13          0         100           0      300 i


CE3# show ip route 3.3.3.13
Routing entry for 3.3.3.12/30
 Known via connected, distance 0, metric 0 (connected, via interface)
```

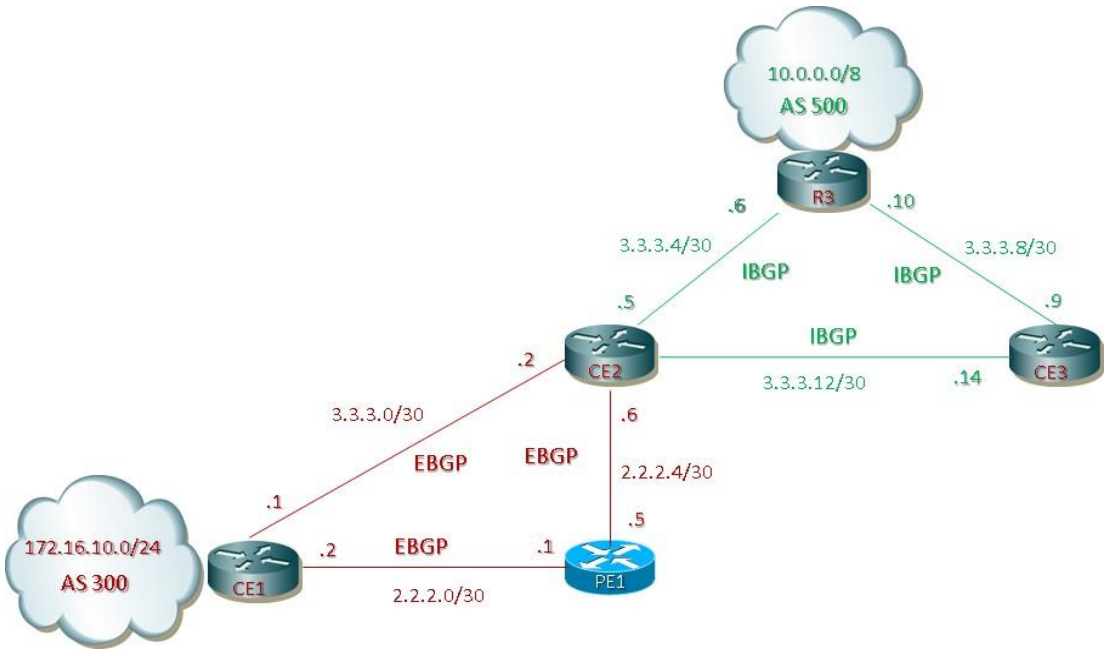

Tabella 3: RIB and BGP Table from CE3 prima parte analizzata

RIB and BGP table outputs from **R3 (AS500) internal roter**

```
R3#show ip route
Gateway of last resort is not set
   10.0.0.0/8  is variably subnetted, 2 subnet,  2 mask
C  10.0.0.0/8  is directly connected, GigabitEthernet0/0
L  10.0.0.1/32 is directly connected, GigabitEthernet0/0
   172.16.0.0/24 is subnetted, 1 subnet
B  172.16.10.0 [200/0] via 3.3.3.5, 00:09:32
!

R3#show ip bgp
       Network          Next-Hop      Metric      LocPrf      Weight      Path
*>     10.0.0.0         0.0.0.0            0                   32768      ?
*>i    172.16.10.0/24   3.3.3.5            0         100           0      300 i
!
```

Commento:

Dal punto di vista R3 vale lo stessa cosa vista nella tabella outputs di CE3 sopra:

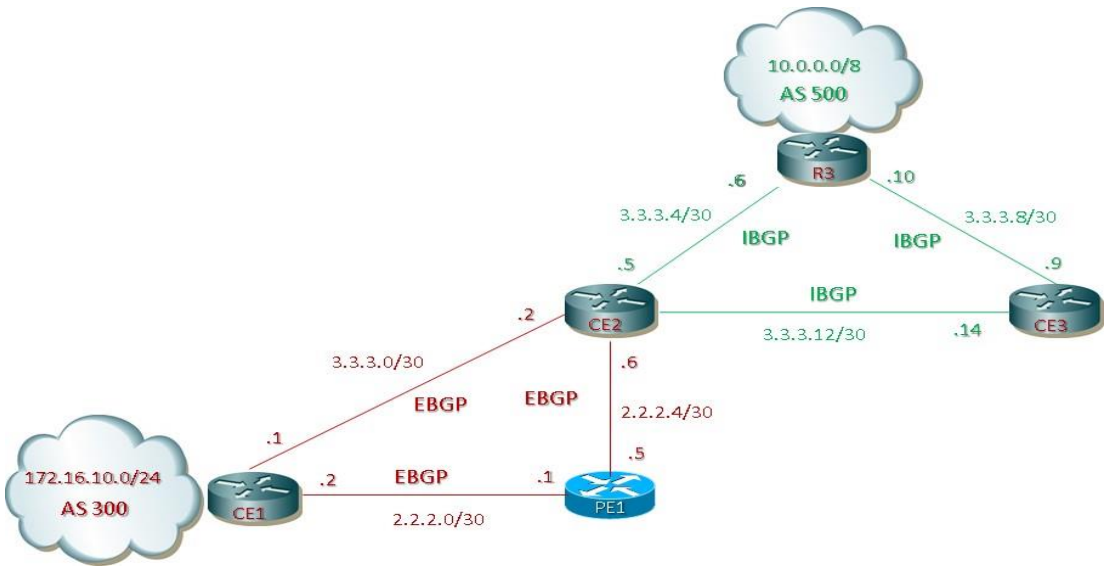

Tabella 4: RIB and BGP Table from R3 prima parte analizzata

No-Export community

Per questo test è stato tolto il collegamento EBGP tra i due AS300 ed AS500 come in figura evidenziato:

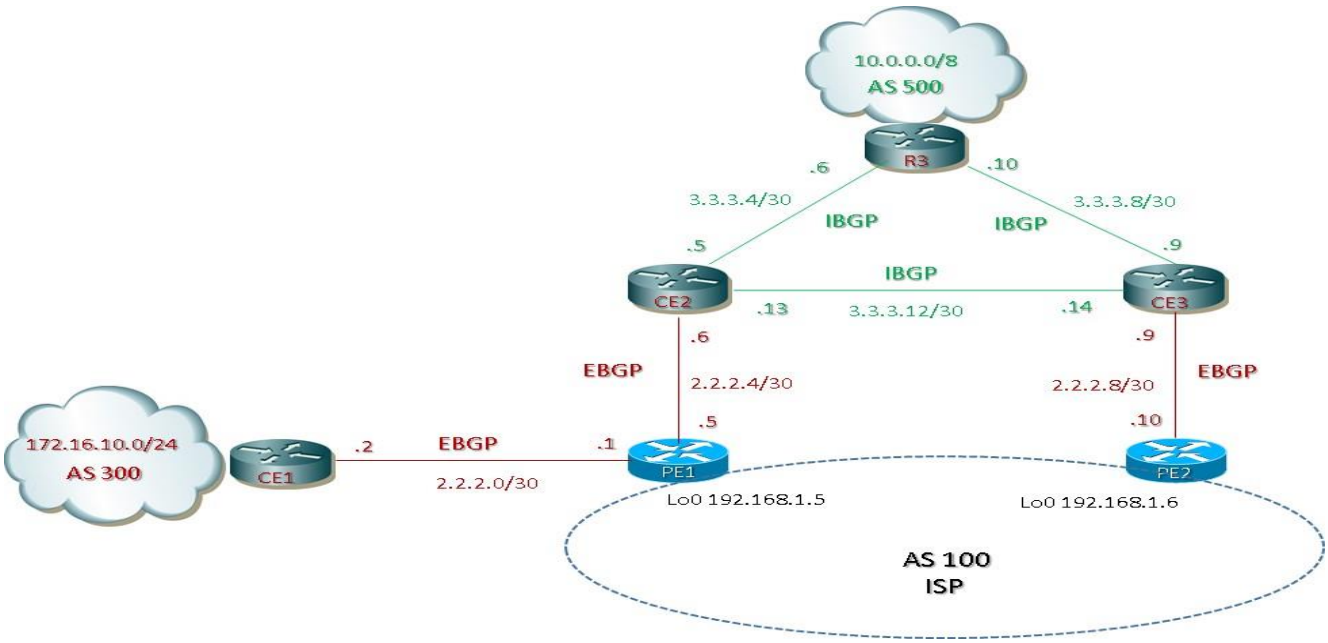

***Figura** 3: uso delle community well-know no-advertisement*

In questo scenario i due AS external al dominio ISP non sono direttamente collegati e pertanto le due Prefix, rispettivamente annunciate dai due AS external, transitano per l'AS 100 (ISP) e sono annunciati dai router PE di bordo; la tabella di outputs mette in evidenza le differenze tra la prima parte di architettura (quella con l'EBGP tra i due AS external e quella senza di figura 4)

RIB and BGP table outputs from **CE1 (AS300)**
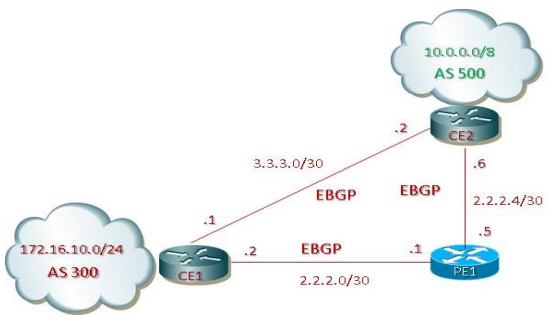

```
CE1#show ip route

Gateway of last resort is not set

B  10.0.0.0/8 [20/0] via 3.3.3.2,  03:39:31

   172.16.10.0 is variable subnetted, 2 subnet, 2 masks

C  172.16.10.0/24 is directly connected, GigabitEthernet0/1

L  172.16.10.1/32 is directly connected, GigabitEthernet0/1

!
!
CE1#show ip bgp

    Network          Next-Hop      Metric      LocPrf      Weight      Path

*   10.0.0.0         2.2.2.1                                0      100 500 ?

*>                   3.3.3.2                                0      500 ?

*>  172.16.10.0/24   0.0.0.0           0                    32768      i

!
```

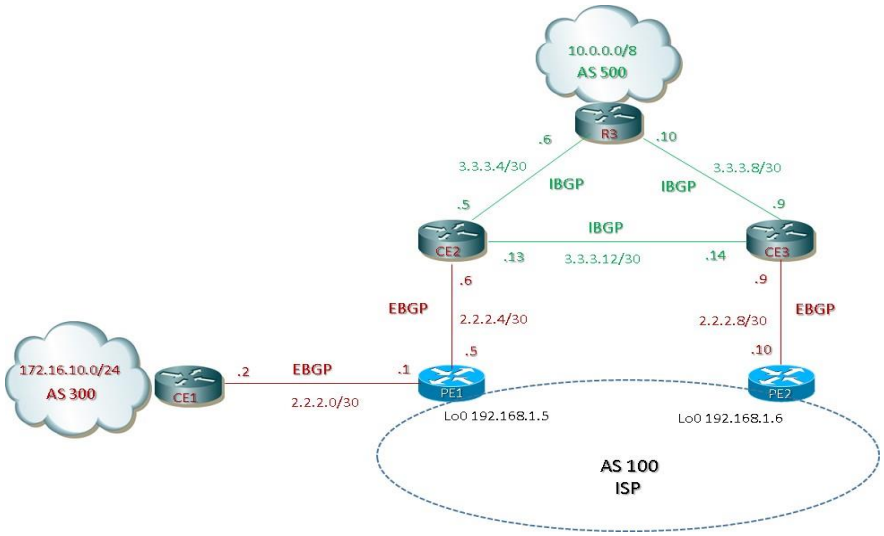


```
CE1#show ip route

Gateway of last resort is not set

B  10.0.0.0/8 [20/0] via 2.2.2.1,  00:16:12

   172.16.10.0 is variable subnetted, 2 subnet, 2 masks

C  172.16.10.0/24 is directly connected, GigabitEthernet0/1

L  172.16.10.1/32 is directly connected, GigabitEthernet0/1

!
!
CE1#show ip bgp

    Network          Next-Hop      Metric      LocPrf      Weight      Path

*>   10.0.0.0         2.2.2.1                                0       100 500 ?

*>  172.16.10.0/24   0.0.0.0           0                    32768      i
```

DAL PUNTO DI VISTA DI PE1 quale router di congiunzione tra i due AS abbiamo i seguenti outputs

```
PE1#show ip route
Gateway of last resort is not set
B  10.0.0.0/8 [20/0] via 2.2.2.6,  1d:3h
   172.16.10.0  is variable subnetted, 1 subnet, 1 masks
B  172.16.10.0/24 [20/0] via 2.2.2.2,  1d:04h

PE1#show ip bgp
    Network          Next-Hop     Metric      LocPrf      Weight      Path
*>   10.0.0.0          2.2.2.6                                0      500 ?
*>   172.16.10.0/24    2.2.2.2           0                    0      300 i
```

OBIETTIVO DEL TEST E' NON PERMETTERE LA REDISTRIBUZIONE DELLA PREFIX 10/8 ALL'INTERNO DELL'AS 300:

Per questo motivo dobbiamo realizzare questa configurazione presso i seguenti routers:

CE2 (nodo che annuncia la Prefix 10/8)

```
access-list 10 permit 10.0.0.0  0.255.255.255
!
route-map community permit 10
 match ip address 10
 set community no-export
!
router bgp 500
 address-family ipv4
 neighbor 2.2.2.5 route-map community out
 neighbor 2.2.2.5 send-community
```

Il comando send-community per aver effetto deve essere configurato anche lato peer (PE1):

PE1

```
router bgp 100
 address-family ipv4
 neighbor 2.2.2.6 send-community
```

A seguito di questa configurazione vediamo di nuovo se la prefix 10/8 è presente in CE1 di AS300.

```
CE1#show ip route 10.0.0.0
% Subnet not in table

CE1#show ip bgp
    Network          Next-Hop      Metric      LocPrf      Weight      Path
*>  172.16.10.0/24   0.0.0.0            0                   32768      i
```

Si ricorda il significato delle community di tipo well-Know quali:

no-export: tutti gli annunci aventi questo tipo di community non devono essere propagati al di fuori di un AS (che riceve l'annuncio) attraverso peering di tipo EBGP;

no-advertse: come sopra ma attraverso peering di tipo EBGP e IBGP

Tabella 5: RIB and BGP Table a seguito applicazione community no-export

4 Architettura completa con RR e PE routers

Architettura:

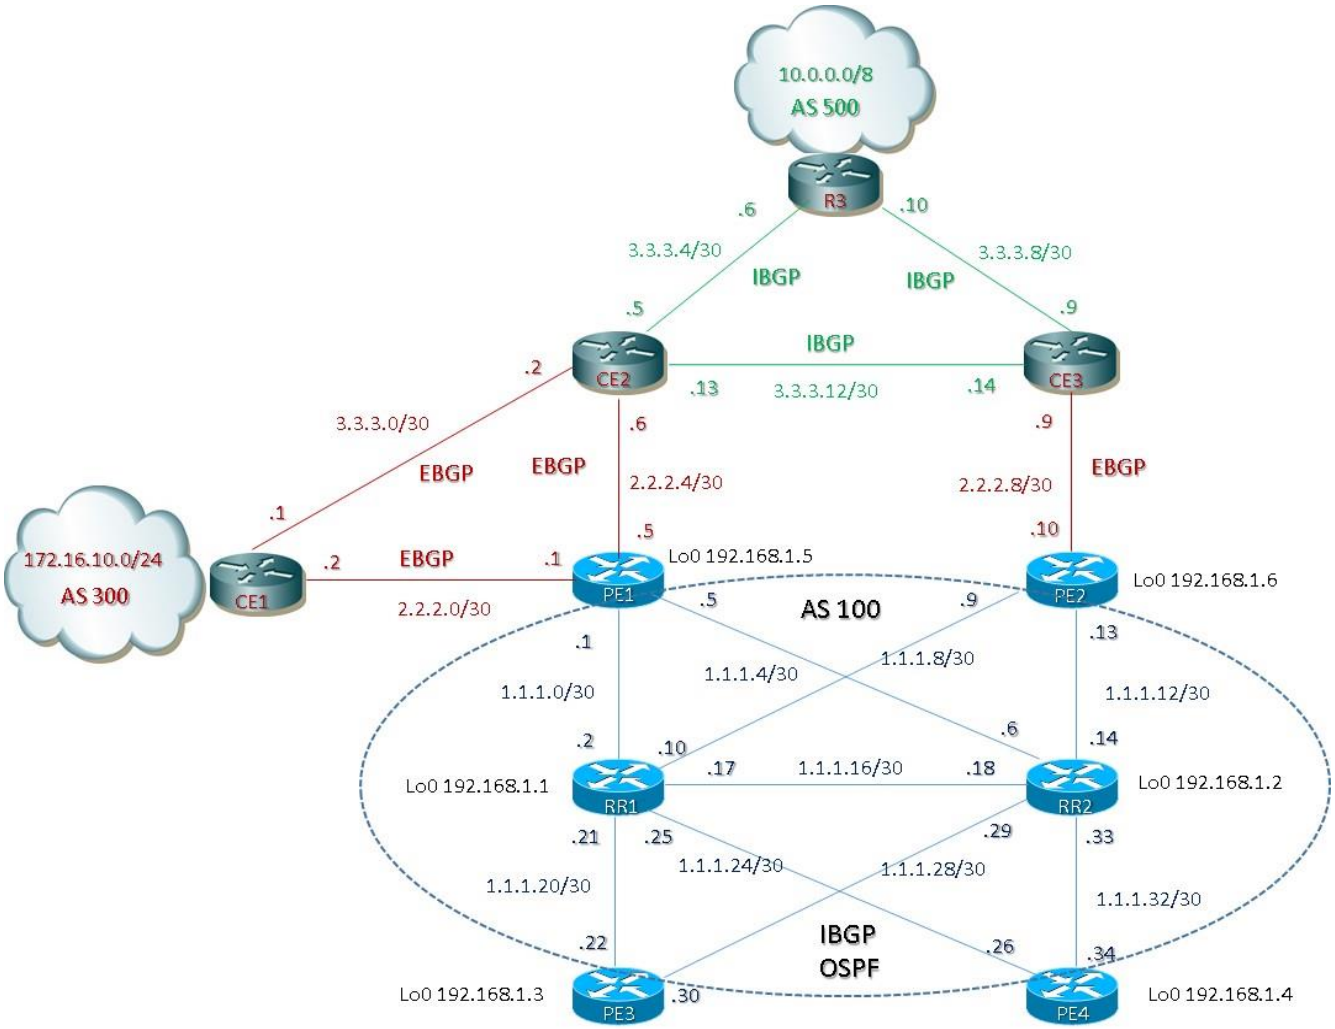

***Figura** 4: architettura completa con RR e PE routers*

Configurazione BGP di base Router Reflector con stesso cluster ID

RR1

```
interface loopback0
description RID
ip address 192.168.1.1 255.255.255.255
!
router ospf 100
 router-id 192.168.1.1
 network 1.1.1.0  0.0.0.255 area 0
 network 192.168.1.0  0.0.0.255 area 0
!
```

```
router bgp 100
 bgp router-id 192.168.1.1
 bgp cluster-id 100
 bgp log-neighbor-changes
 neighbor reflectors peer-group
 neighbor reflectors remote-as 100
 neighbor reflectors update-source loopback0
 neighbor clients peer-group
 neighbor clients remote-as 100
 neighbor clients update-source loopback0
 neighbor 192.168.1.2 peer-group reflectors
 neighbor 192.168.1.3 peer-group clients
 neighbor 192.168.1.4 peer-group clients
 neighbor 192.168.1.5 peer-group clients
 neighbor 192.168.1.6 peer-group clients
!
 address-family ipv4
  neighbor clients route-reflector-client
  neighbor 192.168.1.2 activate
  neighbor 192.168.1.3 activate
  neighbor 192.168.1.4 activate
  neighbor 192.168.1.5 activate
  neighbor 192.168.1.6 activate
```

RR2

```
interface loopback0
description RID
ip address 192.168.1.2 255.255.255.255
!
router ospf 100
 router-id 192.168.1.2
 network 1.1.1.0  0.0.0.255 area 0
 network 192.168.1.0  0.0.0.255 area 0
!
router bgp 100
 bgp router-id 192.168.1.1
 bgp cluster-id 100
 bgp log-neighbor-changes
 neighbor reflectors peer-group
 neighbor reflectors remote-as 100
 neighbor reflectors update-source loopback0
 neighbor clients peer-group
 neighbor clients remote-as 100
 neighbor clients update-source loopback0
 neighbor 192.168.1.1 peer-group reflectors
 neighbor 192.168.1.3 peer-group clients
 neighbor 192.168.1.4 peer-group clients
 neighbor 192.168.1.5 peer-group clients
 neighbor 192.168.1.6 peer-group clients
!
 address-family ipv4
  neighbor clients route-reflector-client
  neighbor 192.168.1.1 activate
  neighbor 192.168.1.3 activate
  neighbor 192.168.1.4 activate
  neighbor 192.168.1.5 activate
  neighbor 192.168.1.6 activate
```

Configurazione BGP di base PE routers

PE1

```
interface loopback0
description RID
ip address 192.168.1.5 255.255.255.255
!
router ospf 100
 router-id 192.168.1.5
 network 1.1.1.0  0.0.0.255 area 0
 network 192.168.1.0  0.0.0.255 area 0
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 2.2.2.2 remote-as 300
 neighbor 2.2.2.6 remote-as 500
 neighbor 192.168.1.1 remote-as 100
 neighbor 192.168.1.1 update-source loopback0
 neighbor 192.168.1.2 remote-as 100
 neighbor 192.168.1.2 update-source loopback0
!
 address-family ipv4
 neighbor 2.2.2.2 activate
 neighbor 2.2.2.6 activate
 neighbor 192.168.1.1 activate
 neighbor 192.168.1.1 next-hop-self
 neighbor 192.168.1.2 activate
 neighbor 192.168.1.2 next-hop-self
```

PE2

```
interface loopback0
description RID
ip address 192.168.1.6 255.255.255.255
!
router ospf 100
 router-id 192.168.1.6
 network 1.1.1.0  0.0.0.255 area 0
 network 192.168.1.0  0.0.0.255 area 0
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 2.2.2.9 remote-as 500
 neighbor 192.168.1.1 remote-as 100
 neighbor 192.168.1.1 update-source loopback0
 neighbor 192.168.1.2 remote-as 100
 neighbor 192.168.1.2 update-source loopback0
!
 address-family ipv4
 neighbor 2.2.2.9 activate
 neighbor 192.168.1.1 activate
 neighbor 192.168.1.1 next-hop-self
 neighbor 192.168.1.2 activate
 neighbor 192.168.1.2 next-hop-self
```

PE3

```
interface loopback0
description RID
ip address 192.168.1.3 255.255.255.255
!
router ospf 100
 router-id 192.168.1.3
 network 1.1.1.0  0.0.0.255 area 0
 network 192.168.1.0  0.0.0.255 area 0
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 192.168.1.1 remote-as 100
 neighbor 192.168.1.1 update-source loopback0
 neighbor 192.168.1.2 remote-as 100
 neighbor 192.168.1.2 update-source loopback0
!
 address-family ipv4
 neighbor 192.168.1.1 activate
 neighbor 192.168.1.1 next-hop-self
 neighbor 192.168.1.2 activate
 neighbor 192.168.1.2 next-hop-self
```

PE4

```
interface loopback0
description RID
ip address 192.168.1.4 255.255.255.255
!
router ospf 100
 router-id 192.168.1.4
 network 1.1.1.0  0.0.0.255 area 0
 network 192.168.1.0  0.0.0.255 area 0
!
router bgp 100
 bgp log-neighbor-changes
 neighbor 192.168.1.1 remote-as 100
 neighbor 192.168.1.1 update-source loopback0
 neighbor 192.168.1.2 remote-as 100
 neighbor 192.168.1.2 update-source loopback0
!
 address-family ipv4
 neighbor 192.168.1.1 activate
 neighbor 192.168.1.1 next-hop-self
 neighbor 192.168.1.2 activate
 neighbor 192.168.1.2 next-hop-self
```

RIB and BGP table per RR e PE routers

RIB and BGP table outputs from **PE1 (AS100)**

```
PE1#show ip route
Gateway of last resort is not set
B   10.0.0.0/8 [20/0] via 2.2.2.6,  1d:01h
    172.16.0.0/24 is subnetted,  1 subnets
B   172.16.10.0 [20/0] via 2.2.2.2, 1d:02h
!

PE1#show ip bgp
       Network          Next-Hop      Metric      LocPrf      Weight      Path
*      10.0.0.0         2.2.2.2                                0      300 500 ?
*>                      2.2.2.6                                0      500 ?
*      172.16.10.0/24   2.2.2.6                                0      500 300 i
*>                      2.2.2.2            0                   0      300 i
!
```

Le due reti 10/8 e 172.16.10.0 sono presenti in RIB e nella BGP table
in grassetto sono evidenziate le best-path e riportate anche attraverso i seguenti outputs:

```
PE1# show ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/8, version 5
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
      1           4
  Refresh Epoch   1
  300 500
     2.2.2.2 from 2.2.2.2 (172.16.10.1)
         Origin incomplete, localpref 100, valid, external
         rx pathid: 0,  tx pathid: 0
   Refresh Epoch  1
   500
     2.2.2.6 from 2.2.2.6 (3.3.3.13)
        Origin incomplete, localpref 100, valid, external, best
        rx pathid: 0,  tx pathid: 0


PE1# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 2
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
      1           4
  Refresh Epoch   1
  500 300
     2.2.2.6 from 2.2.2.6 (3.3.3.13)
         Origin incomplete, localpref 100, valid, external
         rx pathid: 0,  tx pathid: 0
```

```
   Refresh Epoch  1
   300
     2.2.2.2 from 2.2.2.2 (172.16.10.1)
        Origin incomplete, localpref 100, valid, external, best
        rx pathid: 0,  tx pathid: 0



Il router PE advertise le due Prefix ricevute via EBGP dai rispettivi peers ai due RR quali
neighborship IBGP:



PE1# show ip bgp neighbors 192.168.1.1 advertised-routes

       Network            Next-Hop     Metric      LocPrf      Weight      Path

*>     10.0.0.0           2.2.2.6                                  0      500 ?

*>     172.16.10.0/24     2.2.2.2                                  0      300 i

!

PE1# show ip bgp neighbors 192.168.1.2 advertised-routes

       Network            Next-Hop     Metric      LocPrf      Weight      Path

*>     10.0.0.0           2.2.2.6                                  0      500 ?

*>     172.16.10.0/24     2.2.2.2                                  0      300 i

!
```

Tabella 6: RIB and BGP Table from PE1

RIB and BGP table outputs from **PE2 (AS100)**

```
PE2#show ip route
Gateway of last resort is not set
B   10.0.0.0/8 [20/0] via 2.2.2.9,  1d:03h
    172.16.0.0/24 is subnetted,  1 subnets
B   172.16.10.0 [200/0] via 192.168.1.5, 01:15:48
!
PE2#show ip bgp
         Network          Next-Hop         Metric      LocPrf      Weight      Path
*i       10.0.0.0         192.168.1.5           0         100           0      500 ?
*i                        192.168.1.5           0         100           0      500 ?
*>                        2.2.2.9                                       0      500 ?
*       172.16.10.0/24    2.2.2.9                                       0      500 300 i
*i                        192.168.1.5           0         100           0      300 i
*>i                       192.168.1.5           0         100           0      300 i
!
```

Le due reti 10/8 e 172.16.10.0 sono presenti in RIB e nella BGP table;
in grassetto sono evidenziate le best-path e riportate anche attraverso i seguenti outputs:

```
PE2# show ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/8, version 5
Paths: (3 available, best #3, table default)
  Advertised to update-groups:
      4
  Refresh Epoch   1
  500
     192.168.1.5 (metric 3) from 192.168.1.2 (192.168.1.2)
         Origin incomplete, metric 0, localpref 100, valid, internal
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0
Refresh Epoch   1
  500
     192.168.1.5 (metric 3) from 192.168.1.1 (192.168.1.1)
         Origin incomplete, metric 0, localpref 100, valid, internal
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0
   Refresh Epoch  1
   500
     2.2.2.9 from 2.2.2.9 (3.3.3.14)
        Origin incomplete, localpref 100, valid, external, best
        rx pathid: 0,  tx pathid: 0
!
!
```

```
PE2# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 13
Paths: (3 available, best #3, table default)
  Advertised to update-groups:
      1
  Refresh Epoch   1
  500 300
     2.2.2.9 from 2.2.2.9 (3.3.3.14)
         Origin IGP, localpref 100, valid, external
         rx pathid: 0,  tx pathid: 0
Refresh Epoch   1
  300
     192.168.1.5 (metric 3) from 192.168.1.2 (192.168.1.2)
         Origin IGP, metric 0, localpref 100, valid, internal
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0
   Refresh Epoch  1
   300
     192.168.1.5 (metric 3) from 192.168.1.1 (192.168.1.1)
         Origin IGP, metric 0, localpref 100, valid, internal, best
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0



Il router PE2 advertise la Prefix 10/8 ricevute via EBGP dai rispettivo peer ai due RR quali
neighborship IBGP:


PE1# show ip bgp neighbors 192.168.1.1 advertised-routes

       Network            Next-Hop     Metric      LocPrf      Weight      Path

*>     10.0.0.0           2.2.2.9                                  0      500 ?

!

PE1# show ip bgp neighbors 192.168.1.2 advertised-routes

       Network            Next-Hop     Metric      LocPrf      Weight      Path

*>     10.0.0.0           2.2.2.9                                  0      500 ?

!
```

Tabella 7: RIB and BGP Table from PE2

RIB and BGP table outputs from **RR1 (AS100) router reflector**

```
RR1#show ip route
Gateway of last resort is not set
B   10.0.0.0/8 [200/0] via 192.168.1.5,  02:40:41
    172.16.0.0/24 is subnetted,  1 subnets
B   172.16.10.0 [200/0] via 192.168.1.5, 02:40:41
!
RR1#show ip bgp
         Network          Next-Hop         Metric      LocPrf      Weight      Path
*>i      10.0.0.0         192.168.1.5           0         100           0      500 ?
* i                       192.168.1.6           0         100           0      500 ?
*>i      172.16.10.0/24   192.168.1.5           0         100           0      300 i
!
```

Le due reti 10/8 e 172.16.10.0 sono presenti in RIB e nella BGP table;
in grassetto sono evidenziate le best-path e riportate anche attraverso i seguenti outputs:

```
RR1# show ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/8, version 19
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
      1           2
  500 (Received from a RR-client)
     192.168.1.5 (metric 2) from 192.168.1.5 (192.168.1.5)
         Origin incomplete, metric 0, localpref 100, valid, internal, best
  500 (Received from a RR-client)
     192.168.1.6 (metric 2) from 192.168.1.6 (192.168.1.6)
         Origin incomplete, metric 0, localpref 100, valid, internal
!
RR1# 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)
  Advertised to update-groups:
      1           2
  300 (Received from a RR-client)
     192.168.1.5 (metric 2) from 192.168.1.5 (192.168.1.5)
         Origin IGP, metric 0, localpref 100, valid, internal, best
```

Il router RR1 ha il compito di riflettere gli annunci ricevuti dai suoi client a tutti gli altri peers appartenenti al gruppo medesimo:

```
RR1# show ip bgp neighbor 192.168.1.3 advertised-routes
BGP table version is 19, local router ID is 192.168.1.1

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR1# show ip bgp neighbor 192.168.1.4 advertised-routes
BGP table version is 19, local router ID is 192.168.1.1

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR1# show ip bgp neighbor 192.168.1.5 advertised-routes
BGP table version is 19, local router ID is 192.168.1.1

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR1# show ip bgp neighbor 192.168.1.6 advertised-routes
BGP table version is 19, local router ID is 192.168.1.1

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR1# show ip bgp neighbor 192.168.1.2 advertised-routes
BGP table version is 19, local router ID is 192.168.1.1

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i
```

Tabella 8: RIB and BGP Table from RR1 Router Reflector

RIB and BGP table outputs from **RR2 (AS100) router reflector**

```
RR2#show ip route
Gateway of last resort is not set
B   10.0.0.0/8 [200/0] via 192.168.1.5,  02:58:32
    172.16.0.0/24 is subnetted,  1 subnets
B   172.16.10.0 [200/0] via 192.168.1.5, 02:58:32
!
RR2#show ip bgp
         Network          Next-Hop         Metric      LocPrf      Weight      Path
*>i      10.0.0.0         192.168.1.5           0         100           0      500 ?
* i                       192.168.1.6           0         100           0      500 ?
*>i      172.16.10.0/24   192.168.1.5           0         100           0      300 i
!
```

Le due reti 10/8 e 172.16.10.0 sono presenti in RIB e nella BGP table;
in grassetto sono evidenziate le best-path e riportate anche attraverso i seguenti outputs:

```
RR2# show ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/8, version 19
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
      1           2
  500 (Received from a RR-client)
     192.168.1.5 (metric 2) from 192.168.1.5 (192.168.1.5)
         Origin incomplete, metric 0, localpref 100, valid, internal, best
  500 (Received from a RR-client)
     192.168.1.6 (metric 2) from 192.168.1.6 (192.168.1.6)
         Origin incomplete, metric 0, localpref 100, valid, internal
!
RR2# 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)
  Advertised to update-groups:
      1           2
  300 (Received from a RR-client)
     192.168.1.5 (metric 2) from 192.168.1.5 (192.168.1.5)
         Origin IGP, metric 0, localpref 100, valid, internal, best
```

Il router RR2 ha il compito di riflettere gli annunci ricevuti dai suoi client a tutti gli altri peers appartenenti al gruppo medesimo:

```
RR2# show ip bgp neighbor 192.168.1.3 advertised-routes
BGP table version is 19, local router ID is 192.168.1.2

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR2# show ip bgp neighbor 192.168.1.4 advertised-routes
BGP table version is 19, local router ID is 192.168.1.2

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR2# show ip bgp neighbor 192.168.1.5 advertised-routes
BGP table version is 19, local router ID is 192.168.1.2

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR2# show ip bgp neighbor 192.168.1.6 advertised-routes
BGP table version is 19, local router ID is 192.168.1.2

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i


RR2# show ip bgp neighbor 192.168.1.1 advertised-routes
BGP table version is 19, local router ID is 192.168.1.2

        Network             Next-Hop         Metric      LocPrf      Weight      Path
*>i     10.0.0.0            192.168.1.5           0         100           0      500 ?
*>I     172.16.10.0/24      192.168.1.5           0         100           0      300 i
```

Tabella 9: RIB and BGP Table from RR2 Router Reflector

Commento:

I due Router Reflectors eleggono percorsi in modo consistente per entrambi ed annunciano in modo simmetrico ai loro peers clients lo stesso percorso;

Entrambi i PE3 e PE4 ricevono l'annuncio per le due Prefix da entrambi i due RRs ed usano il percorso via il PE1 (192.168.1.5).

RIB and BGP table outputs from **PE3 (AS100)**

```
PE3#show ip route
Gateway of last resort is not set
B  10.0.0.0/8 [200/0] via 192.168.1.5,  03:02:27
   172.16.0.0/24 is subnetted,  1 subnets
B  172.16.10.0 [200/0] via 192.168.1.5, 03:02:27
!
PE3#show ip bgp
        Network          Next-Hop         Metric      LocPrf      Weight      Path
* i     10.0.0.0         192.168.1.5           0         100           0      500 ?
*>i                      192.168.1.5           0         100           0      500 ?
* i    172.16.10.0/24    192.168.1.5           0         100           0      300 i
*>i                      192.168.1.5           0         100           0      300 i
!
Le due reti 10/8 e 172.16.10.0 sono presenti in RIB e nella BGP table;
in grassetto sono evidenziate le best-path e riportate anche attraverso i seguenti outputs:


PE3# show ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/8, version 34
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  Refresh Epoch   1
  500
     192.168.1.5 (metric 3) from 192.168.1.2 (192.168.1.2)
         Origin incomplete, metric 0, localpref 100, valid, internal
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0
  Refresh Epoch  1
  500
     192.168.1.5 (metric 3) from 192.168.1.1 (192.168.1.1)
        Origin incomplete, metric 0, localpref 100, valid, internal, best
        Originator: 192.168.1.5, Cluster list: 0.0.0.100
        rx pathid: 0,  tx pathid: 0
!
!

PE3# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 19
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  Refresh Epoch   1
  300
     192.168.1.5 (metric 3) from 192.168.1.2 (192.168.1.2)
         Origin IGP, metric 0, localpref 100, valid, internal
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
```

```
        rx pathid: 0,  tx pathid: 0

Refresh Epoch  1
300
   192.168.1.5 (metric 3) from 192.168.1.1 (192.168.1.1)
        Origin IGP, metric 0, localpref 100, valid, internal, best
        Originator: 192.168.1.5, Cluster list: 0.0.0.100
        rx pathid: 0,  tx pathid: 0
```

Tabella 10: RIB and BGP Table from PE3

RIB and BGP table outputs from **PE4 (AS100)**

```
PE4#show ip route

Gateway of last resort is not set

B  10.0.0.0/8 [200/0] via 192.168.1.5,  03:35:28

   172.16.0.0/24 is subnetted,  1 subnets

B  172.16.10.0 [200/0] via 192.168.1.5, 03:35:28

!

PE4#show ip bgp

        Network          Next-Hop         Metric      LocPrf      Weight      Path

* i     10.0.0.0         192.168.1.5           0         100           0      500 ?

*>i                      192.168.1.5           0         100           0      500 ?

* i    172.16.10.0/24    192.168.1.5           0         100           0      300 i

*>i                      192.168.1.5           0         100           0      300 i

!
Le due reti 10/8 e 172.16.10.0 sono presenti in RIB e nella BGP table;
in grassetto sono evidenziate le best-path e riportate anche attraverso i seguenti outputs:


PE4# show ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/8, version 34
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  Refresh Epoch   1
  500
     192.168.1.5 (metric 3) from 192.168.1.2 (192.168.1.2)
         Origin incomplete, metric 0, localpref 100, valid, internal
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0
  Refresh Epoch  1
  500
     192.168.1.5 (metric 3) from 192.168.1.1 (192.168.1.1)
        Origin incomplete, metric 0, localpref 100, valid, internal, best
        Originator: 192.168.1.5, Cluster list: 0.0.0.100
        rx pathid: 0,  tx pathid: 0
!
!

PE4# show ip bgp 172.16.10.0
BGP routing table entry for 172.16.10.0/24, version 19
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  Refresh Epoch   1
  300
     192.168.1.5 (metric 3) from 192.168.1.2 (192.168.1.2)
         Origin IGP, metric 0, localpref 100, valid, internal
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0

 Refresh Epoch  1
 300
```

```
    192.168.1.5 (metric 3) from 192.168.1.1 (192.168.1.1)
         Origin IGP, metric 0, localpref 100, valid, internal, best
         Originator: 192.168.1.5, Cluster list: 0.0.0.100
         rx pathid: 0,  tx pathid: 0
```

Tabella 11*: RIB and BGP Table from PE4*

BGP NEIGHBORS TABLE FROM CE, PE and RR

In questo paragrafo vediamo le impostazioni di default che ciascun router ha in relazione ai suoi peers neighbors

BGP NEIGHBORS table from **CE1 (AS300)**

```
CE1# ship ip bgp neighbors
BGP neighbor is 2.2.2.1,  remote AS 100, external link
BGP version 4, remote router ID 192.168.1.5
BGP state = Established, up for 19:15:18
Last read 00:00:29, last write 00:00:00, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not multisession capable (disabled)
Neighbor capabilities:
 Route Refhesh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised and received
 Multisession Capability:
 Stateful switchover support enabled: NO for session  1

Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               3         12
Keepalive:          1274       1266
Route Refresh:         0          0
Total:              1278       1281
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 unicast
 Session: 2.2.2.1
```

```
 BGP table version 16, neighbor version 1/0
 Output queue size: 0
 Index 2, Advertise bit 0
      2          update-group member
 Slow-peer deterction is disabled
 Slow-peer split-update-group dynamic is disabled

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          3 (consumes 240 bytes)
   Prefixes Total:               8         16
   Implicit Withdraw:            4          9
   Explicit Withdraw:            0          4
   Used as bestpath:           n/a          0
   Used as multipath:          n/a          0

                                   Outbound      Inbound
Local Policy Denied Prefixes
    AS_Path Loop:                       n/a            1
    Bestpath from this peer               7          n/a
    Total                                 7            1

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 2
Last Sent Refresh Start-of-rib: never
Last Sent Refresh End-of-rib: never
Last Received Refresh Start-of-rib: 19:15:18
Last Received Refresh End-of-rib: 19:15:18
Refresh-In took 0 seconds

                                   Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:            0         1
   Refresh End-of-RIB:              0         1

 Address tracking is enabled, the RIB does have a route to 2.2.2.1
 Connections established 2; dropped 1
```

```
 Last reset 19:15:27, due to BGP protocol initialization
 External BGP neighbor configured for connected checks (single-hop no-disable-
 connected-check)
 Interface associated; GigabitEthernet0/0 (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 1
Local host: 2.2.2.2, Local port: 38802
Foreign host: 2.2.2.1, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
-------

BGP neighbor is 3.3.3.2,  remote AS 500, external link
BGP version 4, remote router ID 3.3.3.13
BGP state = Established, up for 18:10:56
Last read 00:00:19, last write 00:00:45, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not multisession capable (disabled)
Neighbor capabilities:
 Route Refhesh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised and received
 Multisession Capability:
 Stateful switchover support enabled: NO for session  1


Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               4          4
Keepalive:          1202       1204
```

```
Route Refresh:        0          0
Total:             1209       1211
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 unicast
 Session: 3.3.3.2
 BGP table version 16, neighbor version 1/0
 Output queue size: 0
 Index 2, Advertise bit 0
      1          update-group member
 Slow-peer deterction is disabled
 Slow-peer split-update-group dynamic is disabled

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          3 (consumes 240 bytes)
   Prefixes Total:               8          3
   Implicit Withdraw:            4          0
   Explicit Withdraw:            0          0
   Used as bestpath:           n/a          3
   Used as multipath:          n/a          0

                                   Outbound      Inbound
Local Policy Denied Prefixes
    AS_Path Loop:                       n/a            2
    Bestpath from this peer               7          n/a
    Total                                 7            2

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 2
Last Sent Refresh Start-of-rib: 18:10:36
Last Sent Refresh End-of-rib: 18:10:36
Last Received Refresh Start-of-rib: 18:10:36
Last Received Refresh End-of-rib: 18:10:36
Refresh-In took 0 seconds
                                  Sent      Rcvd
```

```
 Refresh activity:
   Refresh Start-of-RIB:           1         1
   Refresh End-of-RIB:             1         1

 Address tracking is enabled, the RIB does have a route to 3.3.3.1
 Connections established 2; dropped 1
 Last reset 20:03:53, due to Active open failed
 External BGP neighbor configured for connected checks (single-hop no-disable-
 connected-check)
 Interface associated; GigabitEthernet0/3 (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 1
Local host: 3.3.3.1, Local port: 179
Foreign host: 3.3.3.2, Foreign port: 36516
Connection tableid (VRF): 0
Maximum output segment queue size: 50
```

Tabella 12*: BGP Neighbor Table from CE1 customers routers*

BGP NEIGHBORS table from **CE2 (AS500)**

```
CE2# ship ip bgp neighbors
BGP neighbor is 2.2.2.5,  remote AS 100, external link
BGP version 4, remote router ID 192.168.1.5
BGP state = Established, up for 19:58:41
Last read 00:00:45, last write 00:00:29, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not multisession capable (disabled)
Neighbor capabilities:
 Route Refhesh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised and received
 Multisession Capability:
 Stateful switchover support enabled: NO for session  1

Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               6          7
Keepalive:          1323       1319
Route Refresh:         0          0
Total:              1330       1329
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 unicast
 Session: 2.2.2.5
 BGP table version 16, neighbor version 1/0
 Output queue size: 0
 Index 13, Advertise bit 0
      13         update-group member
 Slow-peer deterction is disabled
 Slow-peer split-update-group dynamic is disabled
```

```
                             Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          1 (consumes 80 bytes)
   Prefixes Total:              8          2
   Implicit Withdraw:           4          1
   Explicit Withdraw:           1          0
   Used as bestpath:          n/a          0
   Used as multipath:         n/a          0

                                 Outbound      Inbound
Local Policy Denied Prefixes
    AS_Path Loop:                     n/a            6
    Bestpath from this peer             1          n/a
    Total                               1            6

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 2
Last Sent Refresh Start-of-rib: never
Last Sent Refresh End-of-rib: never
Last Received Refresh Start-of-rib: 18:55:18
Last Received Refresh End-of-rib: 18:55:18
Refresh-In took 0 seconds
                                Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:           0         1
   Refresh End-of-RIB:             0         1

 Address tracking is enabled, the RIB does have a route to 2.2.2.5
 Connections established 3; dropped 2
 Last reset 19:58:47, due to BGP protocol initialization
 External BGP neighbor configured for connected checks (single-hop no-disable-
 connected-check)
 Interface associated; GigabitEthernet0/0 (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled
```

```
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 1
Local host: 2.2.2.6, Local port: 179
Foreign host: 2.2.2.5, Foreign port: 59765
Connection tableid (VRF): 0
Maximum output segment queue size: 50
-------

BGP neighbor is 3.3.3.1,  remote AS 300, external link
BGP version 4, remote router ID 172.16.10.1
BGP state = Established, up for 18:56:31
Last read 00:00:04, last write 00:00:39, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not multisession caplable (disabled)
Neighbor capabilities:
 Route Refhesh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised and received
 Multisession Capability:
 Stateful switchover support enabled: NO for session  1

Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               4          4
Keepalive:          1254       1253
Route Refresh:         0          0
Total:              1261       1260
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 unicast
 Session: 3.3.3.1
 BGP table version 6, neighbor version 6/0
```

```
 Output queue size: 0
 Index 2, Advertise bit 0
      13      update-group member
 Slow-peer deterction is disabled
 Slow-peer split-update-group dynamic is disabled

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          1 (consumes 80 bytes)
   Prefixes Total:              8          1
   Implicit Withdraw:           4          0
   Explicit Withdraw:           1          0
   Used as bestpath:          n/a          1
   Used as multipath:         n/a          0

                                   Outbound      Inbound
Local Policy Denied Prefixes
    AS_Path Loop:                      n/a            6
    Bestpath from this peer              1          n/a
    Total                                1            6

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 2
Last Sent Refresh Start-of-rib: 18:56:31
Last Sent Refresh End-of-rib: 18:56:31
Last Received Refresh Start-of-rib: 18:56:31
Last Received Refresh End-of-rib: 18:56:31
Refresh-In took 0 seconds
                                Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:          1         1
   Refresh End-of-RIB:            1         1

 Address tracking is enabled, the RIB does have a route to 3.3.3.1
 Connections established 2; dropped 1
 Last reset 20:49:48, due to Active open failed
 External BGP neighbor configured for connected checks (single-hop no-disable-
```

```
 connected-check)
 Interface associated; GigabitEthernet0/3 (peering address in same link)
  Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 1
Local host: 3.3.3.2, Local port: 36516
Foreign host: 3.3.3.1, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
```

Tabella 13*: BGP Neighbor Table from CE2 customer router*

BGP NEIGHBORS table from **PE1 (AS100 ISP)**

```
PE1# ship ip bgp neighbors
BGP neighbor is 2.2.2.2,  remote AS 300, external link
BGP version 4, remote router ID 172.16.10.1
BGP state = Established, up for 21:28:31
Last read 00:00:07, last write 00:00:40, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not multisession capable (disabled)
Neighbor capabilities:
 Route Refhesh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised and received
 Multisession Capability:
 Stateful switchover support enabled: NO for session  1

Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:              12          3
Keepalive:          1413       1421
Route Refresh:         0          0
Total:              1430       1426
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 unicast
 Session: 2.2.2.2
 BGP table version 39, neighbor version 39/0
 Output queue size: 0
 Index 1, Advertise bit 0
      14        update-group member
 Slow-peer detection is disabled
```

```
 Slow-peer split-update-group dynamic is disabled

                             Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          4 (consumes 320 bytes)
   Prefixes Total:             28          4
   Implicit Withdraw:          16          0
   Explicit Withdraw:           5          0
   Used as bestpath:          n/a          1
   Used as multipath:         n/a          0

                                  Outbound      Inbound
Local Policy Denied Prefixes
    Well-know community:                6          n/a
    Bestpath from this peer             3          n/a
    Total                               9            0

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 1
Last Sent Refresh Start-of-rib: 21:32:14
Last Sent Refresh End-of-rib: 21:32:14
Last Received Refresh Start-of-rib: never
Last Received Refresh End-of-rib: never
Refresh-In took 0 seconds
                               Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:          1         0
   Refresh End-of-RIB:            1         0

 Address tracking is enabled, the RIB does have a route to 2.2.2.2
 Connections established 2; dropped 1
 Last reset 21:33:43, due to BGP Notification received of session 1, Administ reset
 External BGP neighbor configured for connected checks (single-hop no-disable-
 connected-check)
 Interface associated; GigabitEthernet0/0 (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
```

```
 SSO is disabled


Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 1
Local host: 2.2.2.1, Local port: 179
Foreign host: 2.2.2.2, Foreign port: 38802
Connection tableid (VRF): 0
Maximum output segment queue size: 50
-------


BGP neighbor is 2.2.2.6,  remote AS 500, external link
BGP version 4, remote router ID 3.3.3.13
BGP state = Established, up for 21:34:34
Last read 00:00:31, last write 00:00:41, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not multisession capable (disabled)
Neighbor capabilities:
 Route Refhesh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised and received
 Multisession Capability:
 Stateful switchover support enabled: NO for session  1


Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               7          6
Keepalive:          1425       1429
Route Refresh:         0          0
Total:              1435       1436
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 30 seconds


For address family: IPv4 unicast
 Session: 2.2.2.6
```

```
 BGP table version 39, neighbor version 39/0
 Output queue size: 0
 Index 1, Advertise bit 0
      1     update-group member
 Slow-peer deterction is disabled
 Slow-peer split-update-group dynamic is disabled

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          4 (consumes 320 bytes)
   Prefixes Total:              28          7
   Implicit Withdraw:           16          3
   Explicit Withdraw:            5          0
   Used as bestpath:           n/a          3
   Used as multipath:          n/a          0

                                   Outbound      Inbound
Local Policy Denied Prefixes
    Well-know Communiy:                   6          n/a
    Bestpath from this peer               3          n/a
    Total                                 9            0

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 1
Last Sent Refresh Start-of-rib: 20:33:10
Last Sent Refresh End-of-rib: 20:33:10
Last Received Refresh Start-of-rib: never
Last Received Refresh End-of-rib: never
Refresh-In took 0 seconds
                                 Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:            1         0
   Refresh End-of-RIB:              1         0

 Address tracking is enabled, the RIB does have a route to 2.2.2.6
 Connections established 3; dropped 2
 Last reset 21:38:57, due to BGP Notification received of session 1, Administ Reset
```

```
 External BGP neighbor configured for connected checks (single-hop no-disable-
 connected-check)
 Interface associated; GigabitEthernet0/1 (peering address in same link)
  Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 1
Local host: 2.2.2.5, Local port: 59765
Foreign host: 2.2.2.6, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
------------


BGP Neighbor is 192.168.1.1, remote AS 100, internal link
 BGP version 4, remote router ID 192.168.1.1
 BGP state = Established, up for 1d18h
 Last read 00:00:20, last write 00:00:40, hold time is 180, keepalive interval is 60 seconds
 Neighbor sessions:
    1 active, is multisession capable (disabled)
 Neighbor capabilities:
  Route Refresh: advertised and received (new)
  Four-octets ASN capability: advertised and received
  Address family IPv4 unicast: advertised and received
  Enhanced Refresh Capability: advertised
  Multisession Capability: and received
  Stateful switchover support enabled: NO for session  1

 Message statistics, state Established:
    InQ depth is 0
    OutQ depth is 0
                     Sent       Rcvd
Opens:                  1          1
Notifications:          0          0
Updates:               15          9
Keepalive:           2782       2788
Route Refresh:          0          0
Total:               2798       2798
```

```
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 Unicast
  Session: 192.168.1.1
  BGP table version 39, neighbor version 39/0
  Output queue size: 0
  Index 4, Advertise bit 1
      3          update-group-member
  NEXT_HOP is always this router for eBGP paths
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          0
   Prefixes Total:             13          6
   Implicit Withdraw:           2          0
   Explicit Withdraw:           7          6
   Used as bestpath:          n/a          0
   Used as multipath:         n/a          0

                                   Outbound      Inbound
Local Policy Denied Prefixes
    ORIGINATOR loop:                    n/a           11
    Bestpath from this peer              15          n/a
    Bestpath from iBGP peer               3          n/a
    Total                                18           11

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 1
Last Sent Refresh Start-of-rib: never
Last Sent Refresh End-of-rib: never
Last Received Refresh Start-of-rib: never
Last Received Refresh End-of-rib: never
```

```
                                Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:           0         0
   Refresh End-of-RIB:             0         0

 Address tracking is enabled, the RIB does have a route to 192.168.1.1
 Connections established 3; dropped 2
 Last reset 1d18h, due to Active open failed
 Interface associated; none (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.5, Local port: 42599
Foreign host: 192.168.1.1, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50

----------------------

BGP Neighbor is 192.168.1.2, remote AS 100, internal link
 BGP version 4, remote router ID 192.168.1.2
 BGP state = Established, up for 1d18h
 Last read 00:00:13, last write 00:00:27, hold time is 180, keepalive interval is 60 seconds
 Neighbor sessions:
    1 active, is multisession capable (disabled)
 Neighbor capabilities:
  Route Refhesh: advertised and received (new)
  Four-octets ASN capability: advertised and received
  Address family IPv4 unicast: advertised and received
  Enhanced Refresh Capability: advertised
  Multisession Capability: and received
  Stateful switchover support enabled: NO for session  1

 Message statistics, state Established:
    InQ depth is 0
    OutQ depth is 0
```

```
                        Sent       Rcvd
Opens:                  1          1
Notifications:          0          0
Updates:               17          9
Keepalive:           2805       2800
Route Refresh:          0          0

Total:               2823       2810
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 Unicast
  Session: 192.168.1.2
  BGP table version 39, neighbor version 39/0
  Output queue size: 0
  Index 4, Advertise bit 1
      4           update-group-member
  NEXT_HOP is always this router for eBGP paths
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          0
   Prefixes Total:              13          6
   Implicit Withdraw:            2          0
   Explicit Withdraw:            7          6
   Used as bestpath:           n/a          0
   Used as multipath:          n/a          0

                                     Outbound      Inbound
Local Policy Denied Prefixes
    ORIGINATOR loop:                    n/a           11
    Bestpath from this peer              15          n/a
    Bestpath from iBGP peer               3          n/a
    Total                                18           11

Number of NLRIs in the update sent: max 3, min 0
Last detected as dynamic slow peer: never
```

```
Dynamic slow peer recovered: never
Refresh Epoch: 1
Last Sent Refresh Start-of-rib: never
Last Sent Refresh End-of-rib: never
Last Received Refresh Start-of-rib: never
Last Received Refresh End-of-rib: never

                                 Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:         0         0
   Refresh End-of-RIB:           0         0

 Address tracking is enabled, the RIB does have a route to 192.168.1.2
 Connections established 3; dropped 2
 Last reset 1d18h, due to Active open failed
 Interface associated; none (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.5, Local port: 24491
Foreign host: 192.168.1.1, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
```

Tabella 14: BGP Neighbor Table from PE1 edge router

BGP NEIGHBORS table from **RR1 (AS100 ISP)**

```
RR1# ship ip bgp neighbors
BGP neighbor is 192.168.1.2,  remote AS 100, internal link
Member of peer-group reflectors for session parameters
BGP version 4, remote router ID 192.168.1.2
BGP state = Established, up for 1d:18h
Last read 00:00:47, last write 00:00:11, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 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:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          9
Keepalive:          3073       3059
Route Refresh:         0          0
Total:              3083       3069
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds


For address family: IPv4 unicast
 Session: 192.168.1.2 session 1
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 2
      2           update-group member
      reflector peer-group member
```

```
                             Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          0
   Prefixes Total:             17          0
   Implicit Withdraw:          12          0
   Explicit Withdraw:           1          0
   Used as bestpath:          n/a          1
   Used as multipath:         n/a          0

                                  Outbound      Inbound
Local Policy Denied Prefixes
    CLUSTER_LIST loop:                 n/a           17
    Invalid Path                         1          n/a
    Total                                1           17

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.2
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.1, Local port: 51253
Foreign host: 192.168.1.2, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
-------

BGP neighbor is 192.168.1.3,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.3
BGP state = Established, up for 1d:18h
Last read 00:00:00, last write 00:00:41, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          3
Keepalive:          3077       3074
Route Refresh:         0          0
Total:              3087       3078
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.3 session 1
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 1
 Route-Reflector Client
      1          update-group member
      clients peer-group member

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          0
   Prefixes Total:              17          0
   Implicit Withdraw:           12          0
   Explicit Withdraw:            1          0
   Used as bestpath:           n/a          1
   Used as multipath:          n/a          0

                                   Outbound      Inbound
Local Policy Denied Prefixes
```

```
    Invalid Path                            1          n/a
    Total                                   1           17

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.3
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.1, Local port: 179
Foreign host: 192.168.1.3, Foreign port: 57325
Connection tableid (VRF): 0
Maximum output segment queue size: 50

---------------------

BGP neighbor is 192.168.1.4,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.4
BGP state = Established, up for 1d:18h
Last read 00:00:45, last write 00:00:30, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          3
Keepalive:          3070       3067
```

```
Route Refresh:        0          0
Total:             3082       3073
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 unicast
 Session: 192.168.1.4
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 1
 Route Reflector Client
      1          update-group member
      clients peer-group member

                             Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          0
   Prefixes Total:             17          0
   Implicit Withdraw:          12          0
   Explicit Withdraw:           1          0
   Used as bestpath:          n/a          0
   Used as multipath:         n/a          0

                                  Outbound      Inbound
Local Policy Denied Prefixes
    Invalid Path                         1          n/a
    Total                                1            0

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.4
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.1, Local port: 47604
Foreign host: 192.168.1.4, Foreign port: 179
```

```
Connection tableid (VRF): 0
Maximum output segment queue size: 50
----------------------


BGP neighbor is 192.168.1.5,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.5
BGP state = Established, up for 1d:18h
Last read 00:00:01, last write 00:00:45, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9         15
Keepalive:          3085       3081
Route Refresh:         0          0
Total:              3096       3098
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds


For address family: IPv4 unicast
 Session: 192.168.1.5
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 1
 Route-Reflector Client
      1          update-group member
      clients peer-group member
```

```
                             Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          4 (Consumes 208 bytes)
   Prefixes Total:             17         17
   Implicit Withdraw:          12          6
   Explicit Withdraw:           1          7
   Used as bestpath:          n/a          4
   Used as multipath:         n/a          0

                                  Outbound      Inbound
Local Policy Denied Prefixes
    Invalid Path                         1          n/a
    Total                                1            0

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.5
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.1, Local port: 179
Foreign host: 192.168.1.5, Foreign port: 42599
Connection tableid (VRF): 0
Maximum output segment queue size: 50

-----------------

BGP neighbor is 192.168.1.6,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.6
BGP state = Established, up for 1d:18h
Last read 00:00:33, last write 00:00:33, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          6
Keepalive:          3091       3093
Route Refresh:         0          0
Total:              3102       3101
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.6
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 1
 Route-Reflector Client
      1          update-group member
      clients peer-group member

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          3 (Consumes 156 bytes)
   Prefixes Total:              17          7
   Implicit Withdraw:           12          3
   Explicit Withdraw:            1          1
   Used as bestpath:           n/a          0
   Used as multipath:          n/a          0
                                    Outbound      Inbound
Local Policy Denied Prefixes
    Invalid Path                          1          n/a
    Total                                 1            0
```

```
Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.6
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.1, Local port: 179
Foreign host: 192.168.1.6, Foreign port: 47070
Connection tableid (VRF): 0
Maximum output segment queue size: 50
```

Tabella 15*: BGP Neighbor Table from RR1 Router Reflector*

BGP NEIGHBORS table from **RR2 (AS100 ISP)**

```
RR2# ship ip bgp neighbors
BGP neighbor is 192.168.1.1,  remote AS 100, internal link
Member of peer-group reflectors for session parameters
BGP version 4, remote router ID 192.168.1.1
BGP state = Established, up for 1d:18h
Last read 00:00:06, last write 00:00:35, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          9
Keepalive:          3093       3108
Route Refresh:         0          0
Total:              3103       3118
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.1 session 1
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 1
      1      update-group member
      reflector peer-group member
```

```
                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          0
   Prefixes Total:              17          0
   Implicit Withdraw:           12          0
   Explicit Withdraw:            1          0
   Used as bestpath:           n/a          0
   Used as multipath:          n/a          0
                                    Outbound      Inbound
Local Policy Denied Prefixes
    CLUSTER_LIST loop:                 n/a           17
    Invalid Path                         1          n/a
    Total                                1           17

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.1
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.2, Local port: 179
Foreign host: 192.168.1.1, Foreign port: 51253
Connection tableid (VRF): 0
Maximum output segment queue size: 50
-------

BGP neighbor is 192.168.1.3,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.3
BGP state = Established, up for 1d:18h
Last read 00:00:40, last write 00:00:39, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          3
Keepalive:          3101       3100
Route Refresh:         0          0
Total:              3111       3104
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.3 session 1
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 1
 Route-Reflector Client
      2           update-group member
      clients peer-group member

                               Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          0
   Prefixes Total:              17          0
   Implicit Withdraw:           12          0
   Explicit Withdraw:            1          0
   Used as bestpath:           n/a          1
   Used as multipath:          n/a          0

                                    Outbound      Inbound
Local Policy Denied Prefixes
    Invalid Path                           1          n/a
    Total                                  1           17

Number of NLRIs in the update sent: max 3, min 0
```

```
 Address tracking is enabled, the RIB does have a route to 192.168.1.3
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled


Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.2, Local port: 179
Foreign host: 192.168.1.3, Foreign port: 57325
Connection tableid (VRF): 0
Maximum output segment queue size: 50


---------------------


BGP neighbor is 192.168.1.4,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.4
BGP state = Established, up for 1d:18h
Last read 00:00:28, last write 00:00:42, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          3
Keepalive:          3104       3102
Route Refresh:         0          0
Total:              3114       3106
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds
```

```
For address family: IPv4 unicast
 Session: 192.168.1.4
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 2
 Route Reflector Client
      2          update-group member
      clients peer-group member

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          0
   Prefixes Total:              17          0
   Implicit Withdraw:           12          0
   Explicit Withdraw:            1          0
   Used as bestpath:           n/a          0
   Used as multipath:          n/a          0

                                   Outbound      Inbound
Local Policy Denied Prefixes
    Invalid Path                          1          n/a
    Total                                 1            0

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.4
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.2, Local port: 43153
Foreign host: 192.168.1.4, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
----------------------
```

```
BGP neighbor is 192.168.1.5,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.5
BGP state = Established, up for 1d:18h
Last read 00:00:13, last write 00:00:31, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9         17
Keepalive:          3114       3119
Route Refresh:         0          0
Total:              3124       3137
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.5
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 2
 Route-Reflector Client
      2          update-group member
      clients peer-group member
                                  Sent       Rcvd
Prefix activity:
   Prefixes Current:              4          4 (Consumes 208 bytes)
   Prefixes Total:               17         21
   Implicit Withdraw:            12         10
```

```
   Explicit Withdraw:           1           7
   Used as bestpath:          n/a           4
   Used as multipath:         n/a           0


                                  Outbound      Inbound
Local Policy Denied Prefixes
    Invalid Path                         1          n/a
    Total                                1           0

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.5
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.2, Local port: 179
Foreign host: 192.168.1.5, Foreign port: 24491
Connection tableid (VRF): 0
Maximum output segment queue size: 50

-----------------

BGP neighbor is 192.168.1.6,  remote AS 100, internal link
Member of peer-group clients for session parameters
BGP version 4, remote router ID 192.168.1.6
BGP state = Established, up for 1d:18h
Last read 00:00:49, last write 00:00:33, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is not 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
Message statistics:
```

```
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               9          7
Keepalive:          3125       3118
Route Refresh:         0          0
Total:              3135       3127
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.6
 BGP table version 19, neighbor version 19/0
 Output queue size: 0
 Index 1
 Route-Reflector Client
      2           update-group member
      clients peer-group member

                              Sent       Rcvd
Prefix activity:
   Prefixes Current:             4          3 (Consumes 156 bytes)
   Prefixes Total:              17         10
   Implicit Withdraw:           12          6
   Explicit Withdraw:            1          1
   Used as bestpath:           n/a          0
   Used as multipath:          n/a          0
                                    Outbound      Inbound
Local Policy Denied Prefixes
    Invalid Path                          1          n/a
    Total                                 1           0

Number of NLRIs in the update sent: max 3, min 0
 Address tracking is enabled, the RIB does have a route to 192.168.1.6
 Connections established 1; dropped 0
 Last reset never
 Transport (tcp) path-mtu-discovery is enabled
```

```
 Graceful-Restart is disabled


Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.1.2, Local port: 179
Foreign host: 192.168.1.6, Foreign port: 29409
Connection tableid (VRF): 0
Maximum output segment queue size: 50
```

Tabella 16*: BGP Neighbor Table from RR2 Router Reflector*

BGP NEIGHBORS table from **PE3 (AS100 ISP)**

```
PE3# ship ip bgp neighbors
BGP neighbor is 192.168.1.1,  remote AS 100, internal link
BGP version 4, remote router ID 192.168.1.1
BGP state = Established, up for 1d:23h
Last read 00:00:50, last write 00:00:06, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is multisession capable (disabled)
Neighbor capabilities:
 Route Refresh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised
 Multisession Capability: and received
 Stateful switchover support enabled: NO for session  1
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               3          9
Keepalive:          3144       3147
Route Refresh:         0          0
Total:              3148       3157
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.1
 BGP table version 34, neighbor version 34/0
 Output queue size: 0
 Index 3, Advertise bit 0
      3          update-group member
 NEXT-HOP is always this router for eBGP paths
 Slow-peer detection is disabled
 Slow-peer split-update-group dynamic is disabled
```

```
                             Sent       Rcvd
Prefix activity:
   Prefixes Current:            4          4 (consumes 320 bytes)
   Prefixes Total:              0         17
   Implicit Withdraw:           0         12
   Explicit Withdraw:           4          1
   Used as bestpath:          n/a          4
   Used as multipath:         n/a          0

                                  Outbound      Inbound
Local Policy Denied Prefixes
    Bestpath from this peer:            20          n/a
    Bestpath from iBGP peer:             4          n/a
    Total                               24            0

Number of NLRIs in the update sent: max 2, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 1
Last Sent Refresh Start-of-rib: never
Last Sent Refresh End-of-rib: never
Last Received Refresh Start-of-rib: never
Last Received Refresh End-of-rib: never
                                Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:           0         0
   Refresh End-of-RIB:             0         0

 Address tracking is enabled, the RIB does have a route to 192.168.1.1
 Connections established 3; dropped 2
 Last reset 1d:23h, due to Active open failed
 Interface associated; none (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
```

```
Local host: 192.168.1.3, Local port: 57325
Foreign host: 192.168.1.1, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
-------

BGP neighbor is 192.168.1.2,  remote AS 100, internal link
BGP version 4, remote router ID 192.168.1.2
BGP state = Established, up for 1d:23h
Last read 00:00:50, last write 00:00:49, hold time is 180, keepalive interval is 60
seconds
Neighbor session:
 1 active, is multisession capable (disabled)
Neighbor capabilities:
 Route Refresh: advertised and received (new)
 Four-octets ASN capability: advertised and received
 Address family IPv4 unicast: advertised and received
 Enhanced Refresh Capability: advertised
 Multisession Capability: and received
 Stateful switchover support enabled: NO for session  1
Message statistics:
  InQ depth: 0
  OutQ depth: 0
                    Sent       Rcvd
Opens:                 1          1
Notifications:         0          0
Updates:               3          9
Keepalive:          3141       3142
Route Refresh:         0          0
Total:              3145       3152
Do log neighbor state change (via global configuration)
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 unicast
 Session: 192.168.1.2
 BGP table version 34, neighbor version 34/0
 Output queue size: 0
 Index 3, Advertise bit 0
      3          update-group member
```

```
 NEXT-HOP is always this router for eBGP paths
 Slow-peer detection is disabled
 Slow-peer split-update-group dynamic is disabled

                             Sent       Rcvd
Prefix activity:
   Prefixes Current:            0          4 (consumes 320 bytes)
   Prefixes Total:              0         17
   Implicit Withdraw:           0         12
   Explicit Withdraw:           4          1
   Used as bestpath:          n/a          4
   Used as multipath:         n/a          0

                                  Outbound      Inbound
Local Policy Denied Prefixes
    Bestpath from this peer:            20          n/a
    Bestpath from iBGP peer:             4          n/a
    Total                               24            0

Number of NLRIs in the update sent: max 2, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 1
Last Sent Refresh Start-of-rib: never
Last Sent Refresh End-of-rib: never
Last Received Refresh Start-of-rib: never
Last Received Refresh End-of-rib: never
                                Sent      Rcvd
 Refresh activity:
   Refresh Start-of-RIB:           0         0
   Refresh End-of-RIB:             0         0

 Address tracking is enabled, the RIB does have a route to 192.168.1.2
 Connections established 3; dropped 2
 Last reset 1d:23h, due to Active open failed
 Interface associated; none (peering address in same link)
 Transport (tcp) path-mtu-discovery is enabled
 Graceful-Restart is disabled
```

```
 SSO is disabled

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255

Local host: 192.168.1.3, Local port: 44056
Foreign host: 192.168.1.2, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50
```

Tabella 17*: BGP Neighbor Table from PE3 edge router*

5 Simulazione Fault PE1 node

Architettura di riferimento:

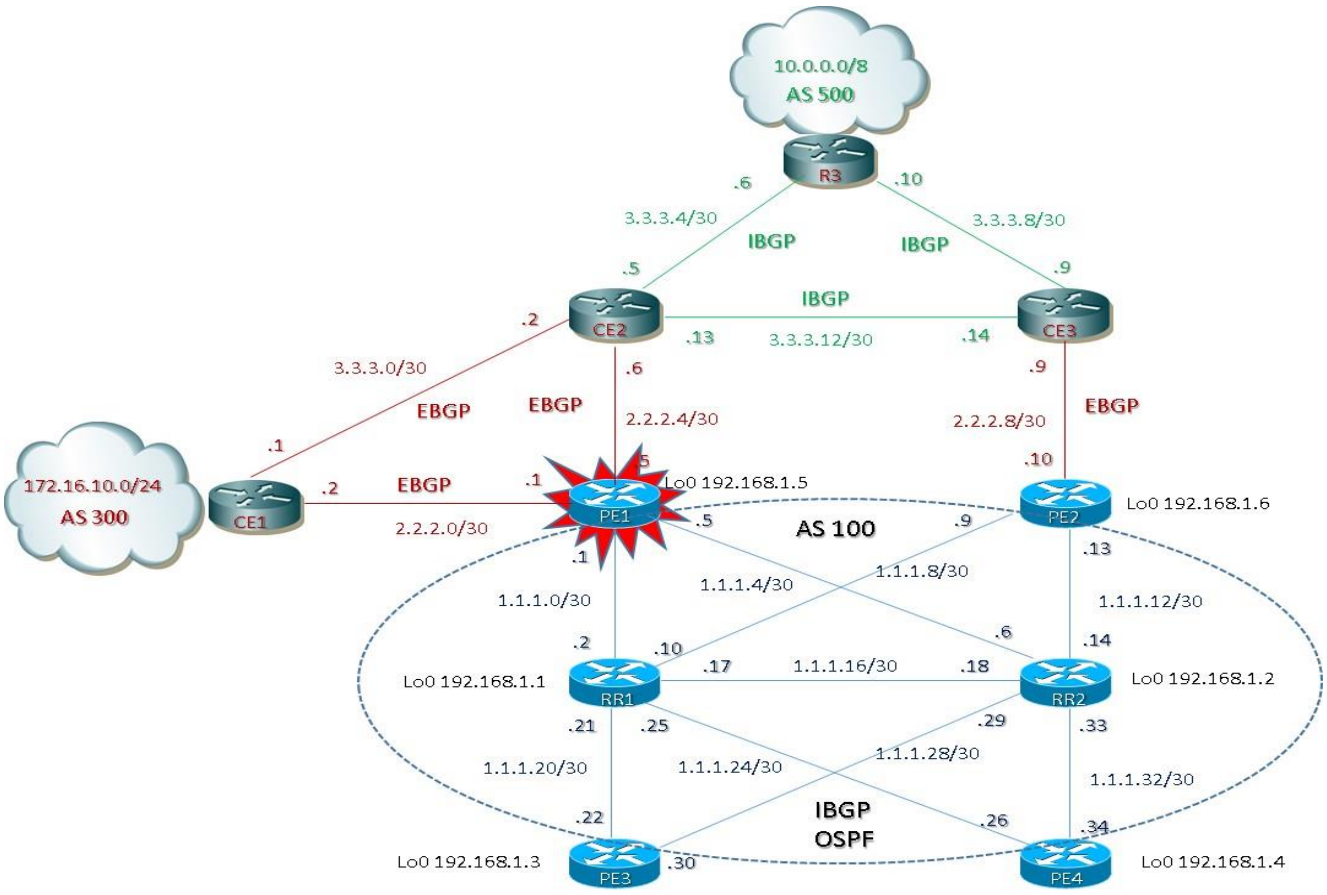

***Figura** 5: architettura con Fault nodo PE1 egress traffic*

E' possibile effettuare un re-inoltro del traffico qualora si verificasse un fault della rete in caso ci siano più percorsi disponibili equivalenti (a parità di costo) e gestita dal piano di forwarding non appena il fault è segnalato.

L'uso di Router Reflector ha significato sulla ridondanza di percorsi nascondendo percorsi alternativi ma forse non sempre percorribile in ambienti reali.

La rilevazione di guasti e la sua propagazione come evento, attraverso BGP può essere lenta e dipende molto dal numero di Prefix gestite.

Soluzione che accelerano la convergenza del BGP a seguito di un fault, quali:

- ✓ Tuning BGP transport
- ✓ BGP fast peering session deactivation
- ✓ BGP and IGP interaction
- ✓ BGP PIC and multiple path propagation (add-path)

BGP TUNING TRANSPORT

E' un meccanismo molto utile ed importante per migliorare le performance di convergenza BGP

TCP è il protocollo utilizzato per stabilire sessioni BGP e scambiare le informazioni update tra peers; settare opportuni parametri TCP è fondamentale, tra questi:

- ✓ abilitare TCP Path MTU discovery per ogni nodo; questo permette un migliore MSS size (Maximum TCP Segment Size); pacchetti che eccedono il valore di MSS, a differenza del MTU, non vengono frammentati ma semplicemente scartati

- ✓ aggiustare lo spazio di accodamento in ingresso di un nodo per permettere una grande quantità di TCP ACK messages; quando un nodo inizia a replicare BGP update, ogni peers inizia a rispondere con TCP ACK messages per ogni secondo segmento trasmesso.

BGP FAST PEERING SESSION DEACTIVATION

Con il solo meccanismo di BGP convergence la rilevazione di un fault è basato su keepalive timers:

valore di default = 60/180 sec

Miglioramento della convergenza BGP con il setting di questi timers a 1/3 sec con il rischio però di provocare inutili flapping tra sessioni BGP

Altre soluzioni possono definirsi attraverso meccanismi BGP come pure il Fast Peering Session Deactivation (quest'ultima di default per sessioni EBGP) tracciando l'interfaccia outgoing associata alla sessione); appena l'interfaccia o il next-hop EBGP viene visto down, la sessione viene deattivata (la tecnica di Event Dampening può mitigare eventuali flapping di interfaccia e quindi per la stessa sessione interessata).

Il comando che disabilità la funzionalità di Fast Peering Session Deactivation è *no bgp fast-external-fallover*;

Questa feature è di default disabilitata per sessioni IBGP, e si suppone che il protocollo IGP sia responsabile di meccanismi per re-routed e restore.

L'uso del BFD è migliore per multipoint interface quale ethernet che non supporta fast link down detection;

BFD fallover è attivato attraverso il comando *neighbor fall-over bfd*

BGP IGP INTERACTION

BGP Prefix si basano su ricorsivi next-hop resolution; di solito il next-hop associato con la Prefix annunciata non sono direttamente connesse ma risolte via un protocollo IGP.

Il processo chiamato BGP Scanner è implementato per una interazione tra il BGP ed IGP; questo processo lavora per performare le tabelle BGP e validare il next-hop; questa validazione consiste principalmente nel riconoscimento del next-hop attraverso la RIB (Routing Information Base) in modo ricorsivo e possibilmente cambiando le informazioni su base forwarding table in risposta a un IGP event.

BGP Scanner di default è pari a 60 sec; possiamo cambiare questo valore attraverso il comando *bgp scan-time* (attenzione ai parametri di timer di questa feature per non sovraccaricare il lavoro CPU di un router qualora fosse presenta una tabella BGP molto larga di informazioni/Prefix)

Il comportamento del BGP Scanner potrebbe essere ancora lento in relazione ad eventuali eventi IGP; IGP può essere settato per reagire ad eventuali cambiamenti della rete all'interno di centinaia di millisecondi.

BGP Next-Hop Tracking (NHT) si basa su una registrazione del next-hop con la RIB chiamata "watcher" ed un richiamo "call back" ogni volta che le informazioni riguardo una Prefix corrispondono ad un cambiamento del next-hop.

BGP NHT si utilizza per rilevare guasti e/o errori nella rete e propagare velocemente questo evento a tutti i nodi BGP in modo da attivare il processo di modifica delle informazioni di forwarding a condizione che sia in uso la FIB gerarchica (questo processo è anche conosciuto come BGP PIC Prefix Independent Convergence poiché non dipende dal numero di BGP Prefix e non è soggetto a best-path re-election)

La FIB gerarchica è un ottimo contributo che migliora un processo di fast convergence e prevede una sua organizzazione in modo appunto gerarchica anziché flat;

Nella FIB flat il piano di forwarding per ogni BGP Prefix è associato ad informazioni quali l'interfaccia di uscita, il MAC rewrite, la label MPLS e se ad esempio dovesse cambiare il valore di next-hop è previsto un ricalcolo ed un aggiornamento per tutte le Prefix associate a quel next-hop, oppure se il next-hop dovesse rimanere lo stesso ma l'interfaccia di uscita cambia, il processo di aggiornamento della FIB è ancora soggetto ad un ricalcolo per tutte le Prefix interessate e riprogrammare il piano di forwarding.

In caso di FIB gerarchica, invece, qualsiasi IGP evento/guasto accada non influisce sulle BGP Prefix e riduce in modo importante i tempi di latenza della tabella di forwarding che è soggetta solo a ricalcoli per le sole informazioni IGP.

La FIB gerarchica non richiede configurazioni ed è di default abilitata nei principali routers BGP.

BGP NHT permette di reagire molto rapidamente all'interno di un dominio AS attraverso un corretto tuning dell'IGP sottostante; il processo quindi di fast convergence copre core links e node come pure edge links e node.

Di default il delay IGP di un evento è stimato di default a 5 sec ed è configurabile attraverso il bgp nextop trigger delay <value>; questo delay è importante in vari scenari e soprattutto quando BGP ha necessità che a seguito di evento dove un guasto della rete può interessare più collegamenti, tutti i nodi IGP siano a conoscenza dell'evento e sono completamente in linea; di solito possiamo settare il NHT delay per essere al di sopra del tempo necessario all'IGP per convergere completamente (possiamo settare questo delay anche a zero secondi qualora volessimo che ogni evento sia segnalato immediatamente; da tenere in considerazione con una accurata messa a punto dei valori IGP per evitare inutili oscillazioni)

Multi-Hop BFD è una feature che mitiga la problematica legata ad IGP route summarization che può nascondere dettagli di cambiamenti che possono avvenire in rete;

BGP PIC

Il BGP PIC (Prefix Independent Convergence) è un meccanismo di fast reroute il quale provvede in termini di sub-seconds a convergenza via IGP a seguito di un failure links e/o node della rete;

La funzionalità PIC richiede la presenza di due annunci BGP della stessa prefix ma con next-hop differenti; pertanto è utile utilizzarlo in scenari dove abbiamo multi-path per una determinata destinazione.

Quando un nodo riceve multiple path (il comportamento del BGP prevede un solo best-path avvertito ed accettato per una determinata prefix per un determinato peer), per la stessa Prefix all'interno della stessa sessione significa che sta utilizzando una estensione chiamata Add-Path che appunto permette di propagare ed accettare multipli path per la stessa destinazione.

BGP ADD-PATH

Il BGP ADD-PATH è una funzionalità che permette di annunciare oltre il best-path anche il best-path alternativo;

Spesso in reti ISP sono presenti cluster di Router Reflector il cui compito è quello di anninciare prefix attraverso il solo best-path (comportamento di default). La possibilità di inviare più annunci è quella del BGP ADD-PATH oppure il BGP DIVERSE-PATH.

L'idea del BGP ADD_PATH si traduce in una estensione del NLRI (Network Layer Reachability Information) associando a ciascun annuncio un identificativo chiamato Path-ID generato automaticamente dal processo BGP.

Di fatto si crea un NLRI costituito da una coppia " Prefix + Path-ID " univoca;

Da notare che il valore Path-ID è fondamentale in questa unicità, altrimenti per via del meccanismo implicit withdraw il secondo annuncio sostituirebbe il primo lasciando nella tabella BGP sempre e solo un annuncio.

Il BGP ADD-PATH permette più annunci verso lo stesso prefisso, evitando il meccanismo di implicit withdraw con la funzionalità conosciuta come Path-Diversity (anziché del path-hiding).

Il Nodo BGP che riceve i due (o più) annunci li installa entrambi nella sua tabella BGP considerandoli diversi; E' infine necessario che i due nodi alla estremità della sessione I-BGP siano entrambi in grado di generare annunci con NLRI modificato (ossia Prefix + Path-ID) e ricevere e quindi interpretare correttamente il nuovo NLRI.

BGP NLRI può essere originato con:

- ✓ Network statement
- ✓ Redistribute statement
- ✓ Aggregate-Address statement
- ✓ Inject-Map statement

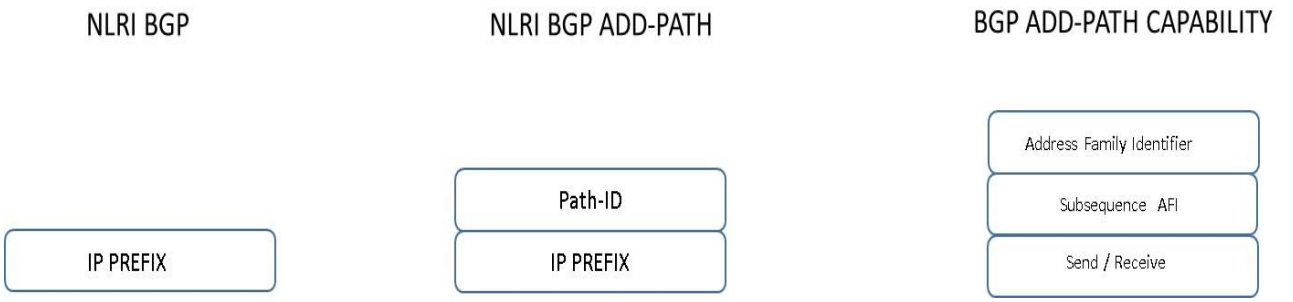

Figura 6: NLRI BGP, NLRI BGP ADD-PATH, ADD-PATH CAPABILITY

Quindi ADD-PATH permette ai peering BGP di negoziare se supportano la capability di advertising / receiving multiple path per-prefix.

Uno speciale path-ID è aggiunto al NLRI proprio per differenziare multipli percorsi per la stessa destinazione all'interno della sessione.

Da notare che BGP ancora considera tutti i percorsi come possibili e confrontabili per il processo di selezione best-path, pertanto sono tutti memorizzati nella tabella di routing RIB e solo uno è scelto come best-path; di fatto gli altri path sono visti come backup paths.

Da notare infine che la funzionalita BGP PIC è possibile anche senza l'ADD-PATH capability in RR scenario, a condizione che i RR propagano i percorsi alternativi ai nodi di bordo (PE); questo potrebbe richiedere una manipolazione delle metriche IGP per gestire che i diversi punti di uscita siano selezionati dagli RRs oppure altre tecniche quali ad esempio l'uso di differenti RD per ambienti multi-homed attachement points.

SIMULAZIONE FAULT PE1 NODE AND VERIFY CONVERGENCE TIME of DEFAULT WITH RRs

DEBUG OSPF BGP from RR1

PE1 SHUTDOWN 17:47

```
RR1#

*Apr 15 17:47:33.763: OSPF: Detect change in topology Base with MTID-0, in LSA type 1, LSID 192.168.1.2 from 192.168.1.2
area 0

*Apr 15 17:47:33.763: OSPF: Schedule SPF in area 0, topology Base with MTID 0

      Change in LS ID 192.168.1.2, LSA type R, spf-type Full

*Apr 15 17:47:33.783: OSPF: Detect MAXAGE in LSA type 2, LS ID 1.1.1.6, from 192.168.1.2

*Apr 15 17:47:33.783: OSPF: Detect generic change in LSA type 2, LSID 1.1.1.6, from 192.168.1.2 area 0

*Apr 15 17:47:33.783: OSPF: Schedule SPF in area 0, topology Base with MTID 0

      Change in LS ID 1.1.1.6, LSA type N, spf-type Full

RR1#

*Apr 15 17:47:35.623: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.1.5 on GigabitEthernet3/0 from FULL to DOWN, Neighbor Down:
Dead timer expired

RR1#

*Apr 15 17:47:36.123: OSPF: Schedule SPF in area 0, topology Base with MTID 0

      Change in LS ID 192.168.1.1, LSA type R, spf-type Full

*Apr 15 17:47:36.123: OSPF: Schedule SPF in area 0, topology Base with MTID 0

      Change in LS ID 1.1.1.2, LSA type N, spf-type Full

RR1#

*Apr 15 17:47:38.763: OSPF: running SPF for area 0, Topology Base with MTID 0 SPF-type Full

*Apr 15 17:47:38.763: OSPF: Initializing to run spf

*Apr 15 17:47:38.763:  OSPF - spf_intra() - rebuilding the tree for topology Base with MTID 0

*Apr 15 17:47:38.763:  It is a router LSA 192.168.1.1. Link Count 6

*Apr 15 17:47:38.763:   Processing link 0, id 192.168.1.1, link data 255.255.255.255, type 3

*Apr 15 17:47:38.763: OSPF:   Add better path to LSA ID 192.168.1.1, gateway 192.168.1.1, dist 1

*Apr 15 17:47:38.763: OSPF:   Add path: next-hop 192.168.1.1, interface Loopback0

*Apr 15 17:47:38.763:   Processing link 1, id 1.1.1.25, link data 1.1.1.25, type 2

*Apr 15 17:47:38.763: OSPF:   Add better path to LSA ID 1.1.1.25, gateway 1.1.1.25, dist 1

*Apr 15 17:47:38.763: OSPF: putting LSA on the clist LSID 1.1.1.25, Type 2, Adv Rtr. 192.168.1.1

*Apr 15 17:47:38.763: OSPF:   Add path: next-hop 1.1.1.25, interface GigabitEthernet5/0

*Apr 15 17:47:38.763:   Processing link 2, id 1.1.1.21, link data 1.1.1.21, type 2

*Apr 15 17:47:38.763: OSPF:   Add better path to LSA ID 1.1.1.21, gateway 1.1.1.21, dist 1

*Apr 15 17:47:38.763: OSPF: putting LSA on the clist LSID 1.1.1.21, Type 2, Adv Rtr. 192.168.1.1

*Apr 15 17:47:38.763: OSPF: upheap LSA on the clist LSID 1.1.1.21, Type 2, Adv Rtr. 192.168.1.1,

                                from index 2 to index 2

*Apr 15 17:47:38.763: OSPF:   Add path: next-hop 1.1.1.21, interface GigabitEthernet4/0
```

```
*Apr 15 17:47:38.763:   Processing link 3, id 1.1.1.18, link data 1.1.1.17, type 2

*Apr 15 17:47:38.763: OSPF:   Add better path to LSA ID 1.1.1.18, gateway 1.1.1.17, dist 1

*Apr 15 17:47:38.763: OSPF: putting LSA on the clist LSID 1.1.1.18, Type 2, Adv Rtr. 192.168.1.2

*Apr 15 17:47:38.763: OSPF: upheap LSA on the clist LSID 1.1.1.18, Type 2, Adv Rtr. 192.168.1.2,

                                from index 3 to index 3

*Apr 15 17:47:38.763: OSPF:   Add path: next-hop 1.1.1.17, interface GigabitEthernet1/0

*Apr 15 17:47:38.763:   Processing link 4, id 1.1.1.10, link data 1.1.1.10, type 2

*Apr 15 17:47:38.763: OSPF:   Add better path to LSA ID 1.1.1.10, gateway 1.1.1.10, dist 1

*Apr 15 17:47:38.763: OSPF: putting LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1

*Apr 15 17:47:38.763: OSPF: upheap LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1,

                                from index 4 to index 4

*Apr 15 17:47:38.763: OSPF:   Add path: next-hop 1.1.1.10, interface GigabitEthernet2/0

*Apr 15 17:47:38.763:   Processing link 5, id 1.1.1.0, link data 255.255.255.252, type 3

*Apr 15 17:47:38.763: OSPF:   Add better path to LSA ID 1.1.1.3, gateway 1.1.1.0, dist 1

*Apr 15 17:47:38.763: OSPF:   Add path: next-hop 1.1.1.2, interface GigabitEthernet3/0

*Apr 15 17:47:38.763: OSPF: downheap LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1,

                                from index 1 to index 2

*Apr 15 17:47:38.763: OSPF: Route update succeeded for 1.1.1.24/255.255.255.252, metric 1, Next Hop:
GigabitEthernet5/0/1.1.1.25 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.763:  It is a network LSA 1.1.1.25. Router Count 2

*Apr 15 17:47:38.763:   Processing router id 192.168.1.1

*Apr 15 17:47:38.763:   New newdist 1 olddist 0

*Apr 15 17:47:38.763:   Processing router id 192.168.1.4

*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.4, gateway 1.1.1.26, dist 1

*Apr 15 17:47:38.767: OSPF: putting LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4

*Apr 15 17:47:38.767: OSPF: upheap LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4,

                                from index 4 to index 4

*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.26, interface GigabitEthernet5/0

*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4,

                                from index 1 to index 2

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.20/255.255.255.252, metric 1, Next Hop:
GigabitEthernet4/0/1.1.1.21 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767:  It is a network LSA 1.1.1.21. Router Count 2

*Apr 15 17:47:38.767:   Processing router id 192.168.1.1

*Apr 15 17:47:38.767:   New newdist 1 olddist 0

*Apr 15 17:47:38.767:   Processing router id 192.168.1.3

*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.3, gateway 1.1.1.22, dist 1

*Apr 15 17:47:38.767: OSPF: putting LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3

*Apr 15 17:47:38.767: OSPF: upheap LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3,

                                from index 4 to index 2

*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.22, interface GigabitEthernet4/0

*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4,
```

```
                                from index 1 to index 3

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.8/255.255.255.252, metric 1, Next Hop:
GigabitEthernet2/0/1.1.1.10 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767:  It is a network LSA 1.1.1.10. Router Count 2

*Apr 15 17:47:38.767:   Processing router id 192.168.1.1

*Apr 15 17:47:38.767:   New newdist 1 olddist 0

*Apr 15 17:47:38.767:   Processing router id 192.168.1.6

*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.6, gateway 1.1.1.9, dist 1

*Apr 15 17:47:38.767: OSPF: putting LSA on the clist LSID 192.168.1.6, Type 1, Adv Rtr. 192.168.1.6

*Apr 15 17:47:38.767: OSPF: upheap LSA on the clist LSID 192.168.1.6, Type 1, Adv Rtr. 192.168.1.6,

                                from index 4 to index 2

*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.9, interface GigabitEthernet2/0

*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3,

                                from index 1 to index 1

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.16/255.255.255.252, metric 1, Next Hop:
GigabitEthernet1/0/1.1.1.17 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767:  It is a network LSA 1.1.1.18. Router Count 2

*Apr 15 17:47:38.767:   Processing router id 192.168.1.2

*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.2, gateway 1.1.1.18, dist 1

*Apr 15 17:47:38.767: OSPF: putting LSA on the clist LSID 192.168.1.2, Type 1, Adv Rtr. 192.168.1.2

*Apr 15 17:47:38.767: OSPF: upheap LSA on the clist LSID 192.168.1.2, Type 1, Adv Rtr. 192.168.1.2,

                                from index 4 to index 1

*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.18, interface GigabitEthernet1/0

*Apr 15 17:47:38.767:   Processing router id 192.168.1.1

*Apr 15 17:47:38.767:   New newdist 1 olddist 0

*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 192.168.1.6, Type 1, Adv Rtr. 192.168.1.6,

                                from index 1 to index 1

*Apr 15 17:47:38.767:  It is a router LSA 192.168.1.2. Link Count 6

*Apr 15 17:47:38.767:   Processing link 0, id 192.168.1.2, link data 255.255.255.255, type 3

*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.2, gateway 192.168.1.2, dist 2

*Apr 15 17:47: 38.767: OSPF:   Add path: next-hop 1.1.1.18, interface GigabitEthernet1/0

*Apr 15 17:47:38.767:   Processing link 1, id 1.1.1.29, link data 1.1.1.29, type 2

*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 1.1.1.29, gateway 1.1.1.29, dist 2

*Apr 15 17:47:38.767: OSPF: putting LSA on the clist LSID 1.1.1.29, Type 2, Adv Rtr. 192.168.1.2

*Apr 15 17:47:38.767: OSPF: upheap LSA on the clist LSID 1.1.1.29, Type 2, Adv Rtr. 192.168.1.2,

                                from index 4 to index 4

*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.18, interface GigabitEthernet1/0

*Apr 15 17:47:38.767:   Processing link 2, id 1.1.1.33, link data 1.1.1.33, type 2

*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 1.1.1.33, gateway 1.1.1.33, dist 2

*Apr 15 17:47:38.767: OSPF: putting LSA on the clist LSID 1.1.1.33, Type 2, Adv Rtr. 192.168.1.2

*Apr 15 17:47:38.767: OSPF: upheap LSA on the clist LSID 1.1.1.33, Type 2, Adv Rtr. 192.168.1.2,

                                from index 5 to index 5
```

```
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.18, interface GigabitEthernet1/0
*Apr 15 17:47:38.767:   Processing link 3, id 1.1.1.14, link data 1.1.1.14, type 2
*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 1.1.1.14, gateway 1.1.1.14, dist 2
*Apr 15 17:47:38.767: OSPF: putting LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2
*Apr 15 17:47:38.767: OSPF: upheap LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2,
                                from index 6 to index 6
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.18, interface GigabitEthernet1/0
*Apr 15 17:47:38.767:   Processing link 4, id 1.1.1.4, link data 255.255.255.252, type 3
*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 1.1.1.7, gateway 1.1.1.4, dist 2
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.18, interface GigabitEthernet1/0
*Apr 15 17:47:38.767:   Processing link 5, id 1.1.1.18, link data 1.1.1.18, type 2
*Apr 15 17:47:38.767:   Ignore newdist 2 olddist 1
*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2,
                                from index 1 to index 3
*Apr 15 17:47:38.767:  It is a router LSA 192.168.1.6. Link Count 3
*Apr 15 17:47:38.767:   Processing link 0, id 192.168.1.6, link data 255.255.255.255, type 3
*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.6, gateway 192.168.1.6, dist 2
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.9, interface GigabitEthernet2/0
*Apr 15 17:47:38.767:   Processing link 1, id 1.1.1.14, link data 1.1.1.13, type 2
*Apr 15 17:47:38.767:   Add equal-length path to 1.1.1.14, dist 2
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.9, interface GigabitEthernet2/0
*Apr 15 17:47:38.767:   Processing link 2, id 1.1.1.10, link data 1.1.1.9, type 2
*Apr 15 17:47:38.767:   Ignore newdist 2 olddist 1
*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 1.1.1.33, Type 2, Adv Rtr. 192.168.1.2,
                                from index 1 to index 4
*Apr 15 17:47:38.767:  It is a router LSA 192.168.1.4. Link Count 3
*Apr 15 17:47:38.767:   Processing link 0, id 192.168.1.4, link data 255.255.255.255, type 3
*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.4, gateway 192.168.1.4, dist 2
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.26, interface GigabitEthernet5/0
*Apr 15 17:47:38.767:   Processing link 1, id 1.1.1.33, link data 1.1.1.34, type 2
*Apr 15 17:47:38.767:   Add equal-length path to 1.1.1.33, dist 2
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.26, interface GigabitEthernet5/0
*Apr 15 17:47:38.767:   Processing link 2, id 1.1.1.25, link data 1.1.1.26, type 2
*Apr 15 17:47:38.767:   Ignore newdist 2 olddist 1
*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 1.1.1.33, Type 2, Adv Rtr. 192.168.1.2,
                                from index 1 to index 2
*Apr 15 17:47:38.767:  It is a router LSA 192.168.1.3. Link Count 3
*Apr 15 17:47:38.767:   Processing link 0, id 192.168.1.3, link data 255.255.255.255, type 3
*Apr 15 17:47:38.767: OSPF:   Add better path to LSA ID 192.168.1.3, gateway 192.168.1.3, dist 2
*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.22, interface GigabitEthernet4/0
*Apr 15 17:47:38.767:   Processing link 1, id 1.1.1.21, link data 1.1.1.22, type 2
```

```
*Apr 15 17:47:38.767:   Ignore newdist 2 olddist 1

*Apr 15 17:47:38.767:   Processing link 2, id 1.1.1.29, link data 1.1.1.30, type 2

*Apr 15 17:47:38.767:   Add equal-length path to 1.1.1.29, dist 2

*Apr 15 17:47:38.767: OSPF:   Add path: next-hop 1.1.1.22, interface GigabitEthernet4/0

*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2,

                                from index 1 to index 2

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.28/255.255.255.252, metric 2, Next Hop:
GigabitEthernet4/0/1.1.1.22 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.28/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.18 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767:  It is a network LSA 1.1.1.29. Router Count 2

*Apr 15 17:47:38.767:   Processing router id 192.168.1.2

*Apr 15 17:47:38.767:   New newdist 2 olddist 1

*Apr 15 17:47:38.767:   Processing router id 192.168.1.3

*Apr 15 17:47:38.767:   New newdist 2 olddist 1

*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2,

                                from index 1 to index 1

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.32/255.255.255.252, metric 2, Next Hop:
GigabitEthernet5/0/1.1.1.26 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.32/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.18 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767:  It is a network LSA 1.1.1.33. Router Count 2

*Apr 15 17:47:38.767:   Processing router id 192.168.1.2

*Apr 15 17:47:38.767:   New newdist 2 olddist 1

*Apr 15 17:47:38.767:   Processing router id 192.168.1.4

*Apr 15 17:47:38.767:   New newdist 2 olddist 1

*Apr 15 17:47:38.767: OSPF: downheap LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2,

                                from index 1 to index 1

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.12/255.255.255.252, metric 2, Next Hop:
GigabitEthernet2/0/1.1.1.9 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.12/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.18 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.767:  It is a network LSA 1.1.1.14. Router Count 2

*Apr 15 17:47:38.767:   Processing router id 192.168.1.2

*Apr 15 17:47:38.767:   New newdist 2 olddist 1

*Apr 15 17:47:38.767:   Processing router id 192.168.1.6

*Apr 15 17:47:38.767:   New newdist 2 olddist 1

*Apr 15 17:47:38.767: OSPF: Adding Stub nets

*Apr 15 17:47:38.767: OSPF: Route update succeeded for 1.1.1.0/255.255.255.252, metric 1, Next Hop:
GigabitEthernet3/0/1.1.1.2 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Route update succeeded for 1.1.1.4/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.18 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Route update succeeded for 192.168.1.1/255.255.255.255, metric 1, Next Hop:
Loopback0/192.168.1.1 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Route update succeeded for 192.168.1.2/255.255.255.255, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.18 area 0, topo/MTID Base/0, process OSPF-100 Router
```

```
*Apr 15 17:47:38.771: OSPF: Route update succeeded for 192.168.1.3/255.255.255.255, metric 2, Next Hop:
GigabitEthernet4/0/1.1.1.22 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Route update succeeded for 192.168.1.4/255.255.255.255, metric 2, Next Hop:
GigabitEthernet5/0/1.1.1.26 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Route update succeeded for 192.168.1.6/255.255.255.255, metric 2, Next Hop:
GigabitEthernet2/0/1.1.1.9 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Entered intra-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Entered intra-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.771: OSPF: Check VLs

*Apr 15 17:47:38.771: OSPF: ospf_gen_asbr_sum_all_areas

*Apr 15 17:47:38.771: OSPF: running spf for summaries area 0, topology Base with MTID 0

*Apr 15 17:47:38.775: OSPF: Entered inter-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.775: OSPF: Entered inter-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.775: OSPF: Started Building Type 5 External Routes

*Apr 15 17:47:38.775: OSPF: Started Building Type 7 External Routes

*Apr 15 17:47:38.775: OSPF: Entered External route sync - area dummy area, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.775: OSPF: Entered External route sync - area dummy area, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.775: OSPF: Entered NSSA route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:38.775: OSPF: Entered NSSA route sync - area 0, topo/MTID Base/0, process OSPF-100 Router



*Apr 15 17:47:43.799: BGP_Router: unhandled major event code 128, minor 0

*Apr 15 17:47:45.951: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:47:45.951: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR1#

*Apr 15 17:48:45.959: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:48:45.959: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR1#

*Apr 15 17:49:45.963: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:49:45.963: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR1#

*Apr 15 17:49:51.623: BGP: 192.168.1.5 connection timed out 180800ms (last update) 180000ms (hold time)

*Apr 15 17:49:51.623: BGP: 192.168.1.5 went from Established to Closing

*Apr 15 17:49:51.623: %BGP-5-ADJCHANGE: neighbor 192.168.1.5 Down BGP Notification sent

*Apr 15 17:49:51.623: %BGP-3-NOTIFICATION: sent to neighbor 192.168.1.5 4/0 (hold time expired) 0 bytes

*Apr 15 17:49:51.623: BGP: ses global 192.168.1.5 (0x6890E914:1) Send NOTIFICATION 4/0 (hold time expired) 0 bytes

*Apr 15 17:49:51.623: BGP: 192.168.1.5 local error close after sending NOTIFICATION

*Apr 15 17:49:51.623: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6890E914:1) NSF delete stale NSF not active

*Apr 15 17:49:51.623: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6890E914:1) NSF no stale paths state is NSF
not active

*Apr 15 17:49:51.623: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6890E914:1) Resetting ALL counters.

*Apr 15 17:49:51.623: BGP: 192.168.1.5 closing

*Apr 15 17:49:51.623: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6890E914:1) Resetting ALL counters.

*Apr 15 17:49:51.623: BGP: 192.168.1.5 went from Closing to Idle
```

```
*Apr 15 17:49:51.627: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.1.5 IPv4 Unicast topology base removed from session  BGP
Notification sent

*Apr 15 17:49:51.627: BGP: ses global 192.168.1.5 (0x6890E914:1) Removed topology IPv4 Unicast:base

*Apr 15 17:49:51.627: BGP: ses global 192.168.1.5 (0x6890E914:1) Removed last topology

*Apr 15 17:49:51.627: BGP: nbr global 192.168.1.5 Active open failed - route to peer is invalid

*Apr 15 17:49:51.627: BGP: nbr global 192.168.1.5 Active open failed - route to peer is invalid

RR1#

*Apr 15 17:50:45.967: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:50:45.967: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR1#

*Apr 15 17:51:45.975: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:51:45.975: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR1#

*Apr 15 17:52:45.979: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:52:45.979: BGP: topo global:IPv4 Multicast:base Scanning routing tables
```

Tabella 18*: Degub OSPF and BGP from RR1*

DEBUG OSPF BGP from RR2

PE1 SHUTDOWN 17:47

```
RR2#
*Apr 15 17:45:42.079: BGP: topo global:IPv4 Unicast:base Scanning routing tables
*Apr 15 17:45:42.079: BGP: topo global:IPv4 Multicast:base Scanning routing tables
*Apr 15 17:46:42.083: BGP: topo global:IPv4 Unicast:base Scanning routing tables
*Apr 15 17:46:42.083: BGP: topo global:IPv4 Multicast:base Scanning routing tables
RR2#
*Apr 15 17:47:32.347: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.1.5 on GigabitEthernet2/0 from FULL to DOWN, Neighbor Down:
Dead timer expired
RR2#
*Apr 15 17:47:32.855: OSPF: Schedule SPF in area 0, topology Base with MTID 0
      Change in LS ID 192.168.1.2, LSA type R, spf-type Full
*Apr 15 17:47:32.855: OSPF: Schedule SPF in area 0, topology Base with MTID 0
      Change in LS ID 1.1.1.6, LSA type N, spf-type Full
*Apr 15 17:47:35.235: OSPF: Detect change in topology Base with MTID-0, in LSA type 1, LSID 192.168.1.1 from 192.168.1.1
area 0
*Apr 15 17:47:35.235: OSPF: Schedule SPF in area 0, topology Base with MTID 0
      Change in LS ID 192.168.1.1, LSA type R, spf-type Full
*Apr 15 17:47:35.263: OSPF: Detect MAXAGE in LSA type 2, LS ID 1.1.1.2, from 192.168.1.1
*Apr 15 17:47:35.263: OSPF: Detect generic change in LSA type 2, LSID 1.1.1.2, from 192.168.1.1 area 0
*Apr 15 17:47:35.263: OSPF: Schedule SPF in area 0, topology Base with MTID 0
      Change in LS ID 1.1.1.2, LSA type N, spf-type Full
*Apr 15 17:47:37.855: OSPF: running SPF for area 0, Topology Base with MTID 0 SPF-type Full
*Apr 15 17:47:37.855: OSPF: Initializing to run spf
*Apr 15 17:47:37.855:  OSPF - spf_intra() - rebuilding the tree for topology Base with MTID 0
*Apr 15 17:47:37.855:  It is a router LSA 192.168.1.2. Link Count 6
*Apr 15 17:47:37.855:   Processing link 0, id 192.168.1.2, link data 255.255.255.255, type 3
*Apr 15 17:47:37.855: OSPF:   Add better path to LSA ID 192.168.1.2, gateway 192.168.1.2, dist 1
*Apr 15 17:47:37.855: OSPF:   Add path: next-hop 192.168.1.2, interface Loopback0
*Apr 15 17:47:37.855:   Processing link 1, id 1.1.1.29, link data 1.1.1.29, type 2
*Apr 15 17:47:37.855: OSPF:   Add better path to LSA ID 1.1.1.29, gateway 1.1.1.29, dist 1
*Apr 15 17:47:37.855: OSPF: putting LSA on the clist LSID 1.1.1.29, Type 2, Adv Rtr. 192.168.1.2
*Apr 15 17:47:37.855: OSPF:   Add path: next-hop 1.1.1.29, interface GigabitEthernet5/0
*Apr 15 17:47:37.855:   Processing link 2, id 1.1.1.33, link data 1.1.1.33, type 2
*Apr 15 17:47:37.855: OSPF:   Add better path to LSA ID 1.1.1.33, gateway 1.1.1.33, dist 1
*Apr 15 17:47:37.855: OSPF: putting LSA on the clist LSID 1.1.1.33, Type 2, Adv Rtr. 192.168.1.2
*Apr 15 17:47:37.855: OSPF: upheap LSA on the clist LSID 1.1.1.33, Type 2, Adv Rtr. 192.168.1.2,
                                from index 2 to index 2
*Apr 15 17:47:37.855: OSPF:   Add path: next-hop 1.1.1.33, interface GigabitEthernet4/0
```

```
*Apr 15 17:47:37.855:   Processing link 3, id 1.1.1.14, link data 1.1.1.14, type 2

*Apr 15 17:47:37.855: OSPF:   Add better path to LSA ID 1.1.1.14, gateway 1.1.1.14, dist 1

*Apr 15 17:47:37.855: OSPF: putting LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2

*Apr 15 17:47:37.855: OSPF: upheap LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2,

                                from index 3 to index 3

*Apr 15 17:47:37.855: OSPF:   Add path: next-hop 1.1.1.14, interface GigabitEthernet3/0

*Apr 15 17:47:37.855:   Processing link 4, id 1.1.1.4, link data 255.255.255.252, type 3

*Apr 15 17:47:37.855: OSPF:   Add better path to LSA ID 1.1.1.7, gateway 1.1.1.4, dist 1

*Apr 15 17:47:37.855: OSPF:   Add path: next-hop 1.1.1.6, interface GigabitEthernet2/0

*Apr 15 17:47:37.855:   Processing link 5, id 1.1.1.18, link data 1.1.1.18, type 2

*Apr 15 17:47:37.855: OSPF:   Add better path to LSA ID 1.1.1.18, gateway 1.1.1.18, dist 1

*Apr 15 17:47:37.855: OSPF: putting LSA on the clist LSID 1.1.1.18, Type 2, Adv Rtr. 192.168.1.2

*Apr 15 17:47:37.855: OSPF: upheap LSA on the clist LSID 1.1.1.18, Type 2, Adv Rtr. 192.168.1.2,

                                from index 4 to index 4

*Apr 15 17:47:37.855: OSPF:   Add path: next-hop 1.1.1.18, interface GigabitEthernet1/0

*Apr 15 17:47:37.855: OSPF: downheap LSA on the clist LSID 1.1.1.18, Type 2, Adv Rtr. 192.168.1.2,

                                from index 1 to index 2

*Apr 15 17:47:37.855: OSPF: Route update succeeded for 1.1.1.28/255.255.255.252, metric 1, Next Hop:
GigabitEthernet5/0/1.1.1.29 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.855:  It is a network LSA 1.1.1.29. Router Count 2

*Apr 15 17:47:37.855:   Processing router id 192.168.1.2

*Apr 15 17:47:37.855:   New newdist 1 olddist 0

*Apr 15 17:47:37.855:   Processing router id 192.168.1.3

*Apr 15 17:47:37.855: OSPF:   Add better path to LSA ID 192.168.1.3, gateway 1.1.1.30, dist 1

*Apr 15 17:47:37.855: OSPF: putting LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3

*Apr 15 17:47:37.859: OSPF: upheap LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3,

                                from index 4 to index 4

*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.30, interface GigabitEthernet5/0

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3,

                                from index 1 to index 2

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.32/255.255.255.252, metric 1, Next Hop:
GigabitEthernet4/0/1.1.1.33 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859:  It is a network LSA 1.1.1.33. Router Count 2

*Apr 15 17:47:37.859:   Processing router id 192.168.1.2

*Apr 15 17:47:37.859:   New newdist 1 olddist 0

*Apr 15 17:47:37.859:   Processing router id 192.168.1.4

*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 192.168.1.4, gateway 1.1.1.34, dist 1

*Apr 15 17:47:37.859: OSPF: putting LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4

*Apr 15 17:47:37.859: OSPF: upheap LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4,

                                from index 4 to index 2

*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.34, interface GigabitEthernet4/0

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3,
```

```
                                from index 1 to index 3

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.16/255.255.255.252, metric 1, Next Hop:
GigabitEthernet1/0/1.1.1.18 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859:  It is a network LSA 1.1.1.18. Router Count 2

*Apr 15 17:47:37.859:   Processing router id 192.168.1.2

*Apr 15 17:47:37.859:   New newdist 1 olddist 0

*Apr 15 17:47:37.859:   Processing router id 192.168.1.1

*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 192.168.1.1, gateway 1.1.1.17, dist 1

*Apr 15 17:47:37.859: OSPF: putting LSA on the clist LSID 192.168.1.1, Type 1, Adv Rtr. 192.168.1.1

*Apr 15 17:47:37.859: OSPF: upheap LSA on the clist LSID 192.168.1.1, Type 1, Adv Rtr. 192.168.1.1,

                                from index 4 to index 2

*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.17, interface GigabitEthernet1/0

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4,

                                from index 1 to index 1

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.12/255.255.255.252, metric 1, Next Hop:
GigabitEthernet3/0/1.1.1.14 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859:  It is a network LSA 1.1.1.14. Router Count 2

*Apr 15 17:47:37.859:   Processing router id 192.168.1.2

*Apr 15 17:47:37.859:   New newdist 1 olddist 0

*Apr 15 17:47:37.859:   Processing router id 192.168.1.6

*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 192.168.1.6, gateway 1.1.1.13, dist 1

*Apr 15 17:47:37.859: OSPF: putting LSA on the clist LSID 192.168.1.6, Type 1, Adv Rtr. 192.168.1.6

*Apr 15 17:47:37.859: OSPF: upheap LSA on the clist LSID 192.168.1.6, Type 1, Adv Rtr. 192.168.1.6,

                                from index 4 to index 1

*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.13, interface GigabitEthernet3/0

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 192.168.1.1, Type 1, Adv Rtr. 192.168.1.1,

                                from index 1 to index 1

*Apr 15 17:47:37.859:  It is a router LSA 192.168.1.6. Link Count 3

*Apr 15 17:47:37.859:   Processing link 0, id 192.168.1.6, link data 255.255.255.255, type 3

*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 192.168.1.6, gateway 192.168.1.6, dist 2

*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.13, interface GigabitEthernet3/0

*Apr 15 17:47:37.859:   Processing link 1, id 1.1.1.14, link data 1.1.1.13, type 2

*Apr 15 17:47:37.859:   Ignore newdist 2 olddist 1

*Apr 15 17:47:37.859:   Processing link 2, id 1.1.1.10, link data 1.1.1.9, type 2

*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 1.1.1.10, gateway 1.1.1.9, dist 2

*Apr 15 17:47:37.859: OSPF: putting LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1

*Apr 15 17:47:37.859: OSPF: upheap LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1,

                                from index 4 to index 4

*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.13, interface GigabitEthernet3/0

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1,

                                from index 1 to index 3

*Apr 15 17:47:37.859:  It is a router LSA 192.168.1.1. Link Count 6
```

```
*Apr 15 17:47:37.859:   Processing link 0, id 192.168.1.1, link data 255.255.255.255, type 3
*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 192.168.1.1, gateway 192.168.1.1, dist 2
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.17, interface GigabitEthernet1/0
*Apr 15 17:47:37.859:   Processing link 1, id 1.1.1.25, link data 1.1.1.25, type 2
*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 1.1.1.25, gateway 1.1.1.25, dist 2
*Apr 15 17:47:37.859: OSPF: putting LSA on the clist LSID 1.1.1.25, Type 2, Adv Rtr. 192.168.1.1
*Apr 15 17:47:37.859: OSPF: upheap LSA on the clist LSID 1.1.1.25, Type 2, Adv Rtr. 192.168.1.1,
                                from index 4 to index 4
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.17, interface GigabitEthernet1/0
*Apr 15 17:47:37.859:   Processing link 2, id 1.1.1.21, link data 1.1.1.21, type 2
*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 1.1.1.21, gateway 1.1.1.21, dist 2
*Apr 15 17:47:37.859: OSPF: putting LSA on the clist LSID 1.1.1.21, Type 2, Adv Rtr. 192.168.1.1
*Apr 15 17:47:37.859: OSPF: upheap LSA on the clist LSID 1.1.1.21, Type 2, Adv Rtr. 192.168.1.1,
                                from index 5 to index 5
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.17, interface GigabitEthernet1/0
*Apr 15 17:47:37.859:   Processing link 3, id 1.1.1.18, link data 1.1.1.17, type 2
*Apr 15 17:47:37.859:   Ignore newdist 2 olddist 1
*Apr 15 17:47:37.859:   Processing link 4, id 1.1.1.10, link data 1.1.1.10, type 2
*Apr 15 17:47:37.859:   Add equal-length path to 1.1.1.10, dist 2
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.17, interface GigabitEthernet1/0
*Apr 15 17:47:37.859:   Processing link 5, id 1.1.1.0, link data 255.255.255.252, type 3
*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 1.1.1.3, gateway 1.1.1.0, dist 2
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.17, interface GigabitEthernet1/0
*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 1.1.1.21, Type 2, Adv Rtr. 192.168.1.1,
                                from index 1 to index 4
*Apr 15 17:47:37.859:  It is a router LSA 192.168.1.3. Link Count 3
*Apr 15 17:47:37.859:   Processing link 0, id 192.168.1.3, link data 255.255.255.255, type 3
*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 192.168.1.3, gateway 192.168.1.3, dist 2
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.30, interface GigabitEthernet5/0
*Apr 15 17:47:37.859:   Processing link 1, id 1.1.1.21, link data 1.1.1.22, type 2
*Apr 15 17:47:37.859:   Add equal-length path to 1.1.1.21, dist 2
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.30, interface GigabitEthernet5/0
*Apr 15 17:47:37.859:   Processing link 2, id 1.1.1.29, link data 1.1.1.30, type 2
*Apr 15 17:47:37.859:   Ignore newdist 2 olddist 1
*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 1.1.1.21, Type 2, Adv Rtr. 192.168.1.1,
                                from index 1 to index 2
*Apr 15 17:47:37.859:  It is a router LSA 192.168.1.4. Link Count 3
*Apr 15 17:47:37.859:   Processing link 0, id 192.168.1.4, link data 255.255.255.255, type 3
*Apr 15 17:47:37.859: OSPF:   Add better path to LSA ID 192.168.1.4, gateway 192.168.1.4, dist 2
*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.34, interface GigabitEthernet4/0
*Apr 15 17:47:37.859:   Processing link 1, id 1.1.1.33, link data 1.1.1.34, type 2
```

```
*Apr 15 17:47:37.859:   Ignore newdist 2 olddist 1

*Apr 15 17:47:37.859:   Processing link 2, id 1.1.1.25, link data 1.1.1.26, type 2

*Apr 15 17:47:37.859:   Add equal-length path to 1.1.1.25, dist 2

*Apr 15 17:47:37.859: OSPF:   Add path: next-hop 1.1.1.34, interface GigabitEthernet4/0

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1,

                                from index 1 to index 2

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.24/255.255.255.252, metric 2, Next Hop:
GigabitEthernet4/0/1.1.1.34 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.24/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.17 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859:  It is a network LSA 1.1.1.25. Router Count 2

*Apr 15 17:47:37.859:   Processing router id 192.168.1.1

*Apr 15 17:47:37.859:   New newdist 2 olddist 1

*Apr 15 17:47:37.859:   Processing router id 192.168.1.4

*Apr 15 17:47:37.859:   New newdist 2 olddist 1

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1,

                                from index 1 to index 1

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.20/255.255.255.252, metric 2, Next Hop:
GigabitEthernet5/0/1.1.1.30 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.20/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.17 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859:  It is a network LSA 1.1.1.21. Router Count 2

*Apr 15 17:47:37.859:   Processing router id 192.168.1.1

*Apr 15 17:47:37.859:   New newdist 2 olddist 1

*Apr 15 17:47:37.859:   Processing router id 192.168.1.3

*Apr 15 17:47:37.859:   New newdist 2 olddist 1

*Apr 15 17:47:37.859: OSPF: downheap LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1,

                                from index 1 to index 1

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.8/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.17 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.8/255.255.255.252, metric 2, Next Hop:
GigabitEthernet3/0/1.1.1.13 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859:  It is a network LSA 1.1.1.10. Router Count 2

*Apr 15 17:47:37.859:   Processing router id 192.168.1.1

*Apr 15 17:47:37.859:   New newdist 2 olddist 1

*Apr 15 17:47:37.859:   Processing router id 192.168.1.6

*Apr 15 17:47:37.859:   New newdist 2 olddist 1

*Apr 15 17:47:37.859: OSPF: Adding Stub nets

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.0/255.255.255.252, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.17 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.859: OSPF: Route update succeeded for 1.1.1.4/255.255.255.252, metric 1, Next Hop:
GigabitEthernet2/0/1.1.1.6 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.863: OSPF: Route update succeeded for 192.168.1.1/255.255.255.255, metric 2, Next Hop:
GigabitEthernet1/0/1.1.1.17 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.863: OSPF: Route update succeeded for 192.168.1.2/255.255.255.255, metric 1, Next Hop:
Loopback0/192.168.1.2 area 0, topo/MTID Base/0, process OSPF-100 Router
```

```
*Apr 15 17:47:37.863: OSPF: Route update succeeded for 192.168.1.3/255.255.255.255, metric 2, Next Hop:
GigabitEthernet5/0/1.1.1.30 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.863: OSPF: Route update succeeded for 192.168.1.4/255.255.255.255, metric 2, Next Hop:
GigabitEthernet4/0/1.1.1.34 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.863: OSPF: Route update succeeded for 192.168.1.6/255.255.255.255, metric 2, Next Hop:
GigabitEthernet3/0/1.1.1.13 area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.863: OSPF: Entered intra-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.863: OSPF: Entered intra-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.867: OSPF: Check VLs

*Apr 15 17:47:37.867: OSPF: ospf_gen_asbr_sum_all_areas

*Apr 15 17:47:37.867: OSPF: running spf for summaries area 0, topology Base with MTID 0

*Apr 15 17:47:37.867: OSPF: Entered inter-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.867: OSPF: Entered inter-area route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.867: OSPF: Started Building Type 5 External Routes

*Apr 15 17:47:37.867: OSPF: Started Building Type 7 External Routes

*Apr 15 17:47:37.867: OSPF: Entered External route sync - area dummy area, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.867: OSPF: Entered External route sync - area dummy area, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.867: OSPF: Entered NSSA route sync - area 0, topo/MTID Base/0, process OSPF-100 Router

*Apr 15 17:47:37.867: OSPF: Entered NSSA route sync - area 0, topo/MTID Base/0, process OSPF-100 Router


*Apr 15 17:47:42.087: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:47:42.087: BGP: topo global:IPv4 Multicast:base Scanning routing tables

*Apr 15 17:47:42.895: BGP_Router: unhandled major event code 128, minor 0

RR2#

*Apr 15 17:48:42.095: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:48:42.095: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR2#

*Apr 15 17:49:42.099: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:49:42.099: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR2#

*Apr 15 17:49:46.971: BGP: 192.168.1.5 connection timed out 180120ms (last update) 180000ms (hold time)

*Apr 15 17:49:46.971: BGP: 192.168.1.5 went from Established to Closing

*Apr 15 17:49:46.971: %BGP-5-ADJCHANGE: neighbor 192.168.1.5 Down BGP Notification sent

*Apr 15 17:49:46.971: %BGP-3-NOTIFICATION: sent to neighbor 192.168.1.5 4/0 (hold time expired) 0 bytes

*Apr 15 17:49:46.971: BGP: ses global 192.168.1.5 (0x6889F134:1) Send NOTIFICATION 4/0 (hold time expired) 0 bytes

*Apr 15 17:49:46.971: BGP: 192.168.1.5 local error close after sending NOTIFICATION

*Apr 15 17:49:46.971: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6889F134:1) NSF delete stale NSF not active

*Apr 15 17:49:46.971: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6889F134:1) NSF no stale paths state is NSF
not active

*Apr 15 17:49:46.971: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6889F134:1) Resetting ALL counters.

*Apr 15 17:49:46.971: BGP: 192.168.1.5 closing

*Apr 15 17:49:46.971: BGP: nbr_topo global 192.168.1.5 IPv4 Unicast:base (0x6889F134:1) Resetting ALL counters.

*Apr 15 17:49:46.971: BGP: 192.168.1.5 went from Closing to Idle
```

```
*Apr 15 17:49:46.971: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.1.5 IPv4 Unicast topology base removed from session  BGP
Notification sent

*Apr 15 17:49:46.975: BGP: ses global 192.168.1.5 (0x6889F134:1) Removed topology IPv4 Unicast:base

*Apr 15 17:49:46.975: BGP: ses global 192.168.1.5 (0x6889F134:1) Removed last topology

*Apr 15 17:49:46.975: BGP: nbr global 192.168.1.5 Active open failed - route to peer is invalid

*Apr 15 17:49:46.975: BGP: nbr global 192.168.1.5 Active open failed - route to peer is invalid

RR2#

*Apr 15 17:50:42.103: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:50:42.103: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR2#

*Apr 15 17:51:42.107: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 17:51:42.107: BGP: topo global:IPv4 Multicast:base Scanning routing tables

RR2#

*Apr 15 18:03:42.155: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 18:03:42.155: BGP: topo global:IPv4 Multicast:base Scanning routing tables
```

Tabella 19: Degub OSPF and BGP from RR2

DEBUG OSPF BGP from PE2

PE1 SHUTDOWN 17:47

```
PE2#
*Apr 15 15:47:31.766: OSPF-100 SPF  : Detect change in LSA type 1, LSID 192.168.1.2 from 192.168.1.2 area 0
*Apr 15 15:47:31.786: OSPF-100 SPF  : Detect MAXAGE in LSA type 2, LS ID 1.1.1.6, from 192.168.1.2
*Apr 15 15:47:31.786: OSPF-100 SPF  : Detect generic change in LSA type 2, LSID 1.1.1.6, from 192.168.1.2 area 0
*Apr 15 15:47:34.131: OSPF-100 SPF  : Detect change in LSA type 1, LSID 192.168.1.1 from 192.168.1.1 area 0
*Apr 15 15:47:34.161: OSPF-100 SPF  : Detect MAXAGE in LSA type 2, LS ID 1.1.1.2, from 192.168.1.1
*Apr 15 15:47:34.161: OSPF-100 SPF  : Detect generic change in LSA type 2, LSID 1.1.1.2, from 192.168.1.1 area 0
*Apr 15 15:47:36.766: OSPF-100 MON  : Begin SPF at 4447.864ms, process time 82ms
*Apr 15 15:47:36.766: OSPF-100 INTRA: Running SPF for area 0, SPF-type Full
*Apr 15 15:47:36.766: OSPF-100 INTRA: Initializing to run spf
*Apr 15 15:47:36.767: OSPF-100 INTRA: spf_intra() - rebuilding the tree
*Apr 15 15:47:36.767: OSPF-100 INTRA:  It is a router LSA 192.168.1.6. Link Count 3
*Apr 15 15:47:36.767: OSPF-100 INTRA:   Processing link 0, id 192.168.1.6, link data 255.255.255.255, type 3
*Apr 15 15:47:36.767: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.6, gateway 192.168.1.6, dist 1
*Apr 15 15:47:36.767: OSPF-100 SPF  :    Add path: next-hop 192.168.1.6, interface Loopback0
*Apr 15 15:47:36.767: OSPF-100 INTRA:   Processing link 1, id 1.1.1.14, link data 1.1.1.13, type 2
*Apr 15 15:47:36.768: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.14, gateway 1.1.1.13, dist 1
*Apr 15 15:47:36.768: OSPF-100 INTRA:    Putting LSA on the clist LSID 1.1.1.14, Type 2, Adv Rtr. 192.168.1.2
*Apr 15 15:47:36.768: OSPF-100 SPF  :    Add path: next-hop 1.1.1.13, interface GigabitEthernet0/3
*Apr 15 15:47:36.768: OSPF-100 INTRA:   Processing link 2, id 1.1.1.10, link data 1.1.1.9, type 2
*Apr 15 15:47:36.768: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.10, gateway 1.1.1.9, dist 1
*Apr 15 15:47:36.768: OSPF-100 INTRA:    Putting LSA on the clist LSID 1.1.1.10, Type 2, Adv Rtr. 192.168.1.1
*Apr 15 15:47:36.768: OSPF-100 INTRA:      Upheap LSA ID 1.1.1.10, Type 2, Adv 192.168.1.1 on clist from index 2 to 2
*Apr 15 15:47:36.768: OSPF-100 SPF  :    Add path: next-hop 1.1.1.9, interface GigabitEthernet0/2
*Apr 15 15:47:36.768: OSPF-100 INTRA:      Downheap LSA ID 1.1.1.10, Type 2, Adv 192.168.1.1 on clist from index 1 to 1
*Apr 15 15:47:36.768: OSPF-100 INTRA: Route update succeeded for 1.1.1.12/255.255.255.252, metric 1, Next Hop:
GigabitEthernet0/3/1.1.1.13 area 0
*Apr 15 15:47:36.769: OSPF-100 INTRA:  It is a network LSA 1.1.1.14. Router Count 2
*Apr 15 15:47:36.769: OSPF-100 INTRA:   Processing router id 192.168.1.2
*Apr 15 15:47:36.769: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.2, gateway 1.1.1.14, dist 1
*Apr 15 15:47:36.769: OSPF-100 INTRA:    Putting LSA on the clist LSID 192.168.1.2, Type 1, Adv Rtr. 192.168.1.2
*Apr 15 15:47:36.769: OSPF-100 INTRA:      Upheap LSA ID 192.168.1.2, Type 1, Adv 192.168.1.2 on clist from index 2 to 2
*Apr 15 15:47:36.770: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3
*Apr 15 15:47:36.770: OSPF-100 INTRA:   Processing router id 192.168.1.6
*Apr 15 15:47:36.770: OSPF-100 INTRA:   New newdist 1 olddist 0
*Apr 15 15:47:36.770: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.2, Type 1, Adv 192.168.1.2 on clist from index 1 to 1
*Apr 15 15:47:36.770: OSPF-100 INTRA: Route update succeeded for 1.1.1.8/255.255.255.252, metric 1, Next Hop:
GigabitEthernet0/2/1.1.1.9 area 0
```

```
*Apr 15 15:47:36.770: OSPF-100 INTRA:  It is a network LSA 1.1.1.10. Router Count 2
*Apr 15 15:47:36.770: OSPF-100 INTRA:   Processing router id 192.168.1.1
*Apr 15 15:47:36.770: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.1, gateway 1.1.1.10, dist 1
*Apr 15 15:47:36.770: OSPF-100 INTRA:    Putting LSA on the clist LSID 192.168.1.1, Type 1, Adv Rtr. 192.168.1.1
*Apr 15 15:47:36.770: OSPF-100 INTRA:      Upheap LSA ID 192.168.1.1, Type 1, Adv 192.168.1.1 on clist from index 2 to 1
*Apr 15 15:47:36.770: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2
*Apr 15 15:47:36.771: OSPF-100 INTRA:   Processing router id 192.168.1.6
*Apr 15 15:47:36.771: OSPF-100 INTRA:   New newdist 1 olddist 0
*Apr 15 15:47:36.771: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.2, Type 1, Adv 192.168.1.2 on clist from index 1 to 1
*Apr 15 15:47:36.771: OSPF-100 INTRA:  It is a router LSA 192.168.1.1. Link Count 6
*Apr 15 15:47:36.771: OSPF-100 INTRA:   Processing link 0, id 192.168.1.1, link data 255.255.255.255, type 3
*Apr 15 15:47:36.771: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.1, gateway 192.168.1.1, dist 2
*Apr 15 15:47:36.771: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2
*Apr 15 15:47:36.771: OSPF-100 INTRA:   Processing link 1, id 1.1.1.25, link data 1.1.1.25, type 2
*Apr 15 15:47:36.771: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.25, gateway 1.1.1.25, dist 2
*Apr 15 15:47:36.771: OSPF-100 INTRA:    Putting LSA on the clist LSID 1.1.1.25, Type 2, Adv Rtr. 192.168.1.1
*Apr 15 15:47:36.771: OSPF-100 INTRA:      Upheap LSA ID 1.1.1.25, Type 2, Adv 192.168.1.1 on clist from index 2 to 2
*Apr 15 15:47:36.772: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2
*Apr 15 15:47:36.772: OSPF-100 INTRA:   Processing link 2, id 1.1.1.21, link data 1.1.1.21, type 2
*Apr 15 15:47:36.772: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.21, gateway 1.1.1.21, dist 2
*Apr 15 15:47:36.772: OSPF-100 INTRA:    Putting LSA on the clist LSID 1.1.1.21, Type 2, Adv Rtr. 192.168.1.1
*Apr 15 15:47:36.772: OSPF-100 INTRA:      Upheap LSA ID 1.1.1.21, Type 2, Adv 192.168.1.1 on clist from index 3 to 3
*Apr 15 15:47:36.772: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2
*Apr 15 15:47:36.772: OSPF-100 INTRA:   Processing link 3, id 1.1.1.18, link data 1.1.1.17, type 2
*Apr 15 15:47:36.772: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.18, gateway 1.1.1.17, dist 2
*Apr 15 15:47:36.772: OSPF-100 INTRA:    Putting LSA on the clist LSID 1.1.1.18, Type 2, Adv Rtr. 192.168.1.2
*Apr 15 15:47:36.772: OSPF-100 INTRA:      Upheap LSA ID 1.1.1.18, Type 2, Adv 192.168.1.2 on clist from index 4 to 4
*Apr 15 15:47:36.772: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2
*Apr 15 15:47:36.772: OSPF-100 INTRA:   Processing link 4, id 1.1.1.10, link data 1.1.1.10, type 2
*Apr 15 15:47:36.773: OSPF-100 INTRA:   Ignore newdist 2 olddist 1
*Apr 15 15:47:36.773: OSPF-100 INTRA:   Processing link 5, id 1.1.1.0, link data 255.255.255.252, type 3
*Apr 15 15:47:36.773: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.3, gateway 1.1.1.0, dist 2
*Apr 15 15:47:36.773: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2
*Apr 15 15:47:36.773: OSPF-100 INTRA:      Downheap LSA ID 1.1.1.18, Type 2, Adv 192.168.1.2 on clist from index 1 to 2
*Apr 15 15:47:36.773: OSPF-100 INTRA:  It is a router LSA 192.168.1.2. Link Count 6
*Apr 15 15:47:36.773: OSPF-100 INTRA:   Processing link 0, id 192.168.1.2, link data 255.255.255.255, type 3
*Apr 15 15:47:36.773: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.2, gateway 192.168.1.2, dist 2
*Apr 15 15:47:36.773: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3
*Apr 15 15:47:36.773: OSPF-100 INTRA:   Processing link 1, id 1.1.1.29, link data 1.1.1.29, type 2
*Apr 15 15:47:36.774: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.29, gateway 1.1.1.29, dist 2
*Apr 15 15:47:36.774: OSPF-100 INTRA:    Putting LSA on the clist LSID 1.1.1.29, Type 2, Adv Rtr. 192.168.1.2
```

```
*Apr 15 15:47:36.774: OSPF-100 INTRA:      Upheap LSA ID 1.1.1.29, Type 2, Adv 192.168.1.2 on clist from index 4 to 4

*Apr 15 15:47:36.774: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.774: OSPF-100 INTRA:   Processing link 2, id 1.1.1.33, link data 1.1.1.33, type 2

*Apr 15 15:47:36.774: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.33, gateway 1.1.1.33, dist 2

*Apr 15 15:47:36.774: OSPF-100 INTRA:    Putting LSA on the clist LSID 1.1.1.33, Type 2, Adv Rtr. 192.168.1.2

*Apr 15 15:47:36.774: OSPF-100 INTRA:      Upheap LSA ID 1.1.1.33, Type 2, Adv 192.168.1.2 on clist from index 5 to 5

*Apr 15 15:47:36.774: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.774: OSPF-100 INTRA:   Processing link 3, id 1.1.1.14, link data 1.1.1.14, type 2

*Apr 15 15:47:36.774: OSPF-100 INTRA:   Ignore newdist 2 olddist 1

*Apr 15 15:47:36.774: OSPF-100 INTRA:   Processing link 4, id 1.1.1.4, link data 255.255.255.252, type 3

*Apr 15 15:47:36.775: OSPF-100 SPF  :    Add better path to LSA ID 1.1.1.7, gateway 1.1.1.4, dist 2

*Apr 15 15:47:36.775: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.775: OSPF-100 INTRA:   Processing link 5, id 1.1.1.18, link data 1.1.1.18, type 2

*Apr 15 15:47:36.775: OSPF-100 INTRA:   Add equal-length path to 1.1.1.18, dist 2

*Apr 15 15:47:36.775: OSPF-100 INTRA:    LSA already on the clist LSID 1.1.1.18, Type 2, Adv Rtr. 192.168.1.2

*Apr 15 15:47:36.775: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.775: OSPF-100 INTRA:      Downheap LSA ID 1.1.1.33, Type 2, Adv 192.168.1.2 on clist from index 1 to 4

*Apr 15 15:47:36.775: OSPF-100 INTRA: Route update succeeded for 1.1.1.24/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/2/1.1.1.10 area 0

*Apr 15 15:47:36.775: OSPF-100 INTRA:  It is a network LSA 1.1.1.25. Router Count 2

*Apr 15 15:47:36.775: OSPF-100 INTRA:   Processing router id 192.168.1.1

*Apr 15 15:47:36.775: OSPF-100 INTRA:   New newdist 2 olddist 1

*Apr 15 15:47:36.776: OSPF-100 INTRA:   Processing router id 192.168.1.4

*Apr 15 15:47:36.776: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.4, gateway 1.1.1.26, dist 2

*Apr 15 15:47:36.776: OSPF-100 INTRA:    Putting LSA on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4

*Apr 15 15:47:36.776: OSPF-100 INTRA:      Upheap LSA ID 192.168.1.4, Type 1, Adv 192.168.1.4 on clist from index 5 to 5

*Apr 15 15:47:36.776: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2

*Apr 15 15:47:36.776: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.4, Type 1, Adv 192.168.1.4 on clist from index 1 to 4

*Apr 15 15:47:36.776: OSPF-100 INTRA: Route update succeeded for 1.1.1.16/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/3/1.1.1.14 area 0

*Apr 15 15:47:36.776: OSPF-100 INTRA: Route update succeeded for 1.1.1.16/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/2/1.1.1.10 area 0

*Apr 15 15:47:36.776: OSPF-100 INTRA:  It is a network LSA 1.1.1.18. Router Count 2

*Apr 15 15:47:36.776: OSPF-100 INTRA:   Processing router id 192.168.1.2

*Apr 15 15:47:36.776: OSPF-100 INTRA:   New newdist 2 olddist 1

*Apr 15 15:47:36.776: OSPF-100 INTRA:   Processing router id 192.168.1.1

*Apr 15 15:47:36.777: OSPF-100 INTRA:   New newdist 2 olddist 1

*Apr 15 15:47:36.777: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.4, Type 1, Adv 192.168.1.4 on clist from index 1 to 2

*Apr 15 15:47:36.777: OSPF-100 INTRA: Route update succeeded for 1.1.1.28/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/3/1.1.1.14 area 0

*Apr 15 15:47:36.777: OSPF-100 INTRA:  It is a network LSA 1.1.1.29. Router Count 2

*Apr 15 15:47:36.777: OSPF-100 INTRA:   Processing router id 192.168.1.2

*Apr 15 15:47:36.777: OSPF-100 INTRA:   New newdist 2 olddist 1
```

```
*Apr 15 15:47:36.777: OSPF-100 INTRA:   Processing router id 192.168.1.3

*Apr 15 15:47:36.777: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.3, gateway 1.1.1.30, dist 2

*Apr 15 15:47:36.777: OSPF-100 INTRA:    Putting LSA on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3

*Apr 15 15:47:36.777: OSPF-100 INTRA:      Upheap LSA ID 192.168.1.3, Type 1, Adv 192.168.1.3 on clist from index 4 to 2

*Apr 15 15:47:36.777: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.777: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.4, Type 1, Adv 192.168.1.4 on clist from index 1 to 3

*Apr 15 15:47:36.777: OSPF-100 INTRA: Route update succeeded for 1.1.1.32/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/3/1.1.1.14 area 0

*Apr 15 15:47:36.777: OSPF-100 INTRA:  It is a network LSA 1.1.1.33. Router Count 2

*Apr 15 15:47:36.777: OSPF-100 INTRA:   Processing router id 192.168.1.2

*Apr 15 15:47:36.777: OSPF-100 INTRA:   New newdist 2 olddist 1

*Apr 15 15:47:36.777: OSPF-100 INTRA:   Processing router id 192.168.1.4

*Apr 15 15:47:36.777: OSPF-100 INTRA:    LSA already on the clist LSID 192.168.1.4, Type 1, Adv Rtr. 192.168.1.4

*Apr 15 15:47:36.777: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.777: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.4, Type 1, Adv 192.168.1.4 on clist from index 1 to 1

*Apr 15 15:47:36.778: OSPF-100 INTRA: Route update succeeded for 1.1.1.20/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/2/1.1.1.10 area 0

*Apr 15 15:47:36.778: OSPF-100 INTRA:  It is a network LSA 1.1.1.21. Router Count 2

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing router id 192.168.1.1

*Apr 15 15:47:36.778: OSPF-100 INTRA:   New newdist 2 olddist 1

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing router id 192.168.1.3

*Apr 15 15:47:36.778: OSPF-100 INTRA:    LSA already on the clist LSID 192.168.1.3, Type 1, Adv Rtr. 192.168.1.3

*Apr 15 15:47:36.778: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2

*Apr 15 15:47:36.778: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.3, Type 1, Adv 192.168.1.3 on clist from index 1 to 1

*Apr 15 15:47:36.778: OSPF-100 INTRA:  It is a router LSA 192.168.1.4. Link Count 3

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing link 0, id 192.168.1.4, link data 255.255.255.255, type 3

*Apr 15 15:47:36.778: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.4, gateway 192.168.1.4, dist 3

*Apr 15 15:47:36.778: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.778: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing link 1, id 1.1.1.33, link data 1.1.1.34, type 2

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Ignore newdist 3 olddist 2

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing link 2, id 1.1.1.25, link data 1.1.1.26, type 2

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Ignore newdist 3 olddist 2

*Apr 15 15:47:36.778: OSPF-100 INTRA:      Downheap LSA ID 192.168.1.3, Type 1, Adv 192.168.1.3 on clist from index 1 to 1

*Apr 15 15:47:36.778: OSPF-100 INTRA:  It is a router LSA 192.168.1.3. Link Count 3

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing link 0, id 192.168.1.3, link data 255.255.255.255, type 3

*Apr 15 15:47:36.778: OSPF-100 SPF  :    Add better path to LSA ID 192.168.1.3, gateway 192.168.1.3, dist 3

*Apr 15 15:47:36.778: OSPF-100 SPF  :    Add path: next-hop 1.1.1.10, interface GigabitEthernet0/2

*Apr 15 15:47:36.778: OSPF-100 SPF  :    Add path: next-hop 1.1.1.14, interface GigabitEthernet0/3

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing link 1, id 1.1.1.21, link data 1.1.1.22, type 2

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Ignore newdist 3 olddist 2

*Apr 15 15:47:36.778: OSPF-100 INTRA:   Processing link 2, id 1.1.1.29, link data 1.1.1.30, type 2
```

```
*Apr 15 15:47:36.778: OSPF-100 INTRA:   Ignore newdist 3 olddist 2

*Apr 15 15:47:36.779: OSPF-100 INTRA: Adding Stub nets

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 1.1.1.0/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/2/1.1.1.10 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 1.1.1.4/255.255.255.252, metric 2, Next Hop:
GigabitEthernet0/3/1.1.1.14 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 192.168.1.1/255.255.255.255, metric 2, Next Hop:
GigabitEthernet0/2/1.1.1.10 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 192.168.1.2/255.255.255.255, metric 2, Next Hop:
GigabitEthernet0/3/1.1.1.14 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 192.168.1.3/255.255.255.255, metric 3, Next Hop:
GigabitEthernet0/3/1.1.1.14 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 192.168.1.3/255.255.255.255, metric 3, Next Hop:
GigabitEthernet0/2/1.1.1.10 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 192.168.1.4/255.255.255.255, metric 3, Next Hop:
GigabitEthernet0/2/1.1.1.10 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 192.168.1.4/255.255.255.255, metric 3, Next Hop:
GigabitEthernet0/3/1.1.1.14 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Route update succeeded for 192.168.1.6/255.255.255.255, metric 1, Next Hop:
Loopback0/192.168.1.6 area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Entered intra-area route sync for area 0

*Apr 15 15:47:36.779: OSPF-100 INTRA: Entered intra-area route sync for area 0

*Apr 15 15:47:36.779: OSPF-100 INTER: Check and generate summary LSA into all areas

*Apr 15 15:47:36.779: OSPF-100 INTER: Running spf for summaries area 0

*Apr 15 15:47:36.779: OSPF-100 INTER: Entered inter-area route sync for area 0

*Apr 15 15:47:36.779: OSPF-100 INTER: Entered inter-area route sync for area 0

*Apr 15 15:47:36.779: OSPF-100 EXTER: Started Building Type 5 External Routes

*Apr 15 15:47:36.779: OSPF-100 EXTER: Started Building Type 7 External Routes

*Apr 15 15:47:36.779: OSPF-100 EXTER: Entered External route sync for area dummy area

*Apr 15 15:47:36.779: OSPF-100 EXTER: Entered External route sync for area dummy area

*Apr 15 15:47:36.779: OSPF-100 EXTER: Entered NSSA route sync for area 0

*Apr 15 15:47:36.779: OSPF-100 EXTER: Entered NSSA route sync for area 0

*Apr 15 15:47:36.779: OSPF-100 MON  : End SPF at 4447.877ms, Total elapsed time 13ms

*Apr 15 15:47:40.735: BGP: topo global:IPv4 Unicast:base Scanning routing tables

*Apr 15 15:47:40.735: BGP: topo global:IPv4 Multicast:base Scanning routing tables

*Apr 15 15:47:40.735: BGP: topo global:L2VPN E-VPN:base Scanning routing tables

*Apr 15 15:47:40.735: BGP: topo global:MVPNv4 Unicast:base Scanning routing tables

*Apr 15 15:47:41.780: BGP_Router: unhandled major event code 128, minor 0
```

Tabella 20: Degub OSPF and BGP from PE2

RIB and BGP TABLE DOPO IL FAULT PE1

PE2 table RIB and BGP

Prima del Fault	Dopo il Fault
PE2#show ip route Gateway of last resort is not set B 10.0.0.0/8 [20/0] via 2.2.2.9, 1d:03h 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [200/0] via 192.168.1.5, 01:15:48 ! PE2#show ip bgp Network Next-Hop Metric LocPrf Weight Path *i 10.0.0.0 192.168.1.5 0 100 0 500 ? *i 192.168.1.5 0 100 0 500 ? *> 2.2.2.9 0 500 ? * 172.16.10.0/24 2.2.2.9 0 500 300 i *i 192.168.1.5 0 100 0 300 i *>i 192.168.1.5 0 100 0 300 i !	PE2#show ip route Gateway of last resort is not set B 10.0.0.0/8 [20/0] via 2.2.2.9, 1d:03h 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [20/0] via 2.2.2.9, 01:15:48 ! PE2#show ip bgp Network Next-Hop Metric LocPrf Weight Path *> 10.0.0.0 2.2.2.9 0 100 0 500 ? *> 172.16.10.0/24 2.2.2.9 0 100 0 500 300 i

Tabella 21: RIB and BGP table from PE2 after Fault

RR1 table RIB and BGP

Prima del Fault	Dopo il Fault
RR1#show ip route Gateway of last resort is not set B 10.0.0.0/8 [200/0] via 192.168.1.5, 02:40:41 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [200/0] via 192.168.1.5, 02:40:41 ! RR1#show ip bgp Network Next-Hop Metric LocPrf Weight Path *>i 10.0.0.0 192.168.1.5 0 100 0 500 ? * i 192.168.1.6 0 100 0 500 ? *>i 172.16.10.0/24 192.168.1.5 0 100 0 300 i	RR1#show ip route Gateway of last resort is not set B 10.0.0.0/8 [200/0] via 192.168.1.6, 02:40:41 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [200/0] via 192.168.1.6, 02:40:41 ! RR1#show ip bgp Network Next-Hop Metric LocPrf Weight Path *>i 10.0.0.0 192.168.1.6 0 100 0 500 ? *>i 172.16.10.0/24 192.168.1.6 0 100 0 500 300 i

Tabella 22: RIB and BGP table from RR1 after Fault

RR2 table RIB and BGP	
Prima del Fault	**Dopo il Fault**
RR1#show ip route Gateway of last resort is not set B 10.0.0.0/8 [200/0] via 192.168.1.5, 02:40:41 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [200/0] via 192.168.1.5, 02:40:41 ! RR1#show ip bgp Network Next-Hop Metric LocPrf Weight Path *>i 10.0.0.0 192.168.1.5 0 100 0 500 ? * i 192.168.1.6 0 100 0 500 ? *>i 172.16.10.0/24 192.168.1.5 0 100 0 300 i	RR1#show ip route Gateway of last resort is not set B 10.0.0.0/8 [200/0] via 192.168.1.6, 02:40:41 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [200/0] via 192.168.1.6, 02:40:41 ! RR1#show ip bgp Network Next-Hop Metric LocPrf Weight Path *>i 10.0.0.0 192.168.1.6 0 100 0 500 ? *>i 172.16.10.0/24 192.168.1.6 0 100 0 500 300 i

Tabella 23: RIB and BGP table from RR3 after Fault

PE3 table RIB and BGP	
Prima del Fault	**Dopo il Fault**
PE3#show ip route Gateway of last resort is not set B 10.0.0.0/8 [200/0] via 192.168.1.5, 03:02:27 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [200/0] via 192.168.1.5, 03:02:27 ! **PE3#show ip bgp** Network Next-Hop Metric LocPrf Weight Path * i 10.0.0.0 192.168.1.5 0 100 0 500 ? *>i 192.168.1.5 0 100 0 500 ? * i 172.16.10.0/24 192.168.1.5 0 100 0 300 i *>i 192.168.1.5 0 100 0 300 i	RR1#show ip route Gateway of last resort is not set B 10.0.0.0/8 [200/0] via 192.168.1.6, 02:40:41 172.16.0.0/24 is subnetted, 1 subnets B 172.16.10.0 [200/0] via 192.168.1.6, 02:40:41 ! RR1#show ip bgp Network Next-Hop Metric LocPrf Weight Path *>i 10.0.0.0 192.168.1.6 0 100 0 500 ? * i 192.168.1.6 0 100 0 500 ? *>i 172.16.10.0/24 192.168.1.6 0 100 0 500 300 i * i 192.168.1.6 0 100 0 500 300 i

Tabella 24: RIB and BGP table from PE3 after Fault

TUNING TIMERS IGP OSPF

OSPF mantiene traccia dei suoi links (collegamenti) per costruire la sua topologia di rete all'interno di un database; i nodi OSPF inviano LSA conosciuti come link-state-advertisement per considerare consistente il suo database costruito con ognuno degli altri nodi appartenenti all'area di competenza OSPF.

I tipi di LSA sono:

LSA type 1: contiene una lista dei collegamenti locali (quelli direttamente connessi) al router e quindi il rispettivo costo (status); questo tipo di LSA è generato da tutti i router della rete IGP e trasmessi ai neighbors all'interno dell'area di competenza.

LSA type 2: contiene una lista dei routers collegati al Designated Router all'interno di un segmento di rete di tipo broadcast come ad esempio una rete ethernet LAN; questo tipo di LSA è generato dai Designated Routers aventi questo ruolo.

LSA type 3: contiene una lista delle network/prefix di destinazione all'interno dell'area di competenza ed è permesso una trasmissione di queste LSA tra inter-area (tra aree diverse); questo tipo di LSA sono generati da router con ruolo di ABR (Area Border Router).

LSA type 4: contiene una route (quindi un percorso) per ogni router che abbia il ruolo di ASBR (Autonomous System Border Router); questo tipo di LSA è generato dai router ABR all'interno della area "locale" in modo che i routers interni a questa area abbiamo conoscenza di un punto di uscita dal proprio AS di competenza.

LSA type 5: contiene tutte le network/prefix esterne all'AS OSPF; questo tipo di LSA sono generati da router con ruolo di ASBR e sono trasmesse in ogni area del sistema OSPF.

LSA type 7: contiene anch'essa le network/prefix esterne all'AS e trasmesse all'interno di una particolare chiamata NSSA no-so-stubby-area generate da un router con ruolo ASBR appartenente a questo tipo di area; questo tipo di LSA non sono trasmesse all'interno di altre aree ma sono convertite in LSA type 5 per essere poi trasmesse in altre aree del sistema OSPF.

Ogni LSA viene trasmessa e propagata a seguito:

- ✓ quando una nuova adiacenza è formata;
- ✓ quando un evento/cambiamento della topologia di rete occorre;
- ✓ quando un LSA raggiunge il suo massimo tempo di vita (max age refresh = 30 minuti di default)

I timers throttle OSPF prevedono tre parametri SPF (Shortest Path First) misurati in millisecondi e rappresentano il ricalcolo dei tempi in caso di evento/cambiamento topologia della rete IGP.

- ✓ Il primo considera un tempo di attesa per dare inizio ad un ricalcolo SPF appena ricevuto un evento IGP conosciuto come start SPF;

- ✓ Il secondo considera un tempo di attesa conosciuto come hold-time SPF tra due differenti processi di ricalcolo SPF qualora il protocollo IGP OSPF riceve un altro evento/cambiamento di topologia durante questo primo tempo di attesa configurato e pertanto inizia un secondo ricalcolo

considerando questo tempo di attesa doppio sino al raggiungimento del massimo tempo di attesa possibile configurato come terzo parametro.

- ✓ Il terzo considera appunto il tempo massimo di attesa conosciuto come hold-time SPF max

DEFAULT TIMERS PARAMETERS	SETTING TIMERS PARAMETERS
SPF Start = 5000 msec SPF Hold = 10000 msec SPF Hold Max = 10000 msec ! LSA Start = 0 msec LSA Hold = 5000 msec LSA Hold Max = 5000 msec ! When using OSPF throttling it is also required set the minumum interval for accepting the same LSA; if the same LSA arrives sooner than the interval setting, the LSA is dropped. (Cisco recommends the arrival interval be less than or equal to the LSA hold timer interval of the timers throttle lsa command) ! ! Hello Time = 10 sec Dead Time = 40 sec	SPF Start = 10 msec SPF Hold = 100 msec SPF Hold Max = 1000 msec ! LSA Start = 10 msec LSA Hold = 100 msec LSA Hold Max = 1000 msec ! LSA-arrival = 50 msec Hello Time = 1 sec Dead Time = 3 sec

Tabella 25*: IGP OSPF Timers*

A livello interfaccia abbiamo due parametri di timers quali:

- ✓ Hello timer: controllo quanto spesso il router trasmette messaggi di routine al proprio neighbor per indicare la sua presenza viva; se il neighbor non sentisse questo tipo di messaggio per un tempo definito dal valore indicato nell'intervallo conosciuto come dead-interval, allora considera il router non più raggiungibile ed annulla la relativa adiancenza con esso;

- ✓ Dead timer: il tempo per il quale il neighbor dichiara down il router.

TUNING COMANDI PER BGP ADD-PATH

Il comando bgp additional-path install viene utilizzato in ambienti non-BGP multipath permettendo di installare path di backup oltre il best-path di regola eletto dal comportamento naturale del BGP.

Il comando no bgp recursion host viene utilizzato quando ci sono contesti di rete dove additonal-path repair è presente, allora la funzionalità di default del BGP di quando perde una rotta di next-hop per una determinata Prefix, il processo CEF proverà a cercare (look-up) una prossima next-longest prefix per quel next-hop non è necessaria, abbassando cosi i tempi di convergenza.

```
router ospf 100
timers throttle spf 10 100 1000
timers throttle lsa all 10 100 1000
timers lsa arrival 50
!
router bgp 100
 bgp additional-paths install
 no bgp recursion host
```

Tabella 26: comandi add-path e tuning OSPF per i router AS100