Network topology:

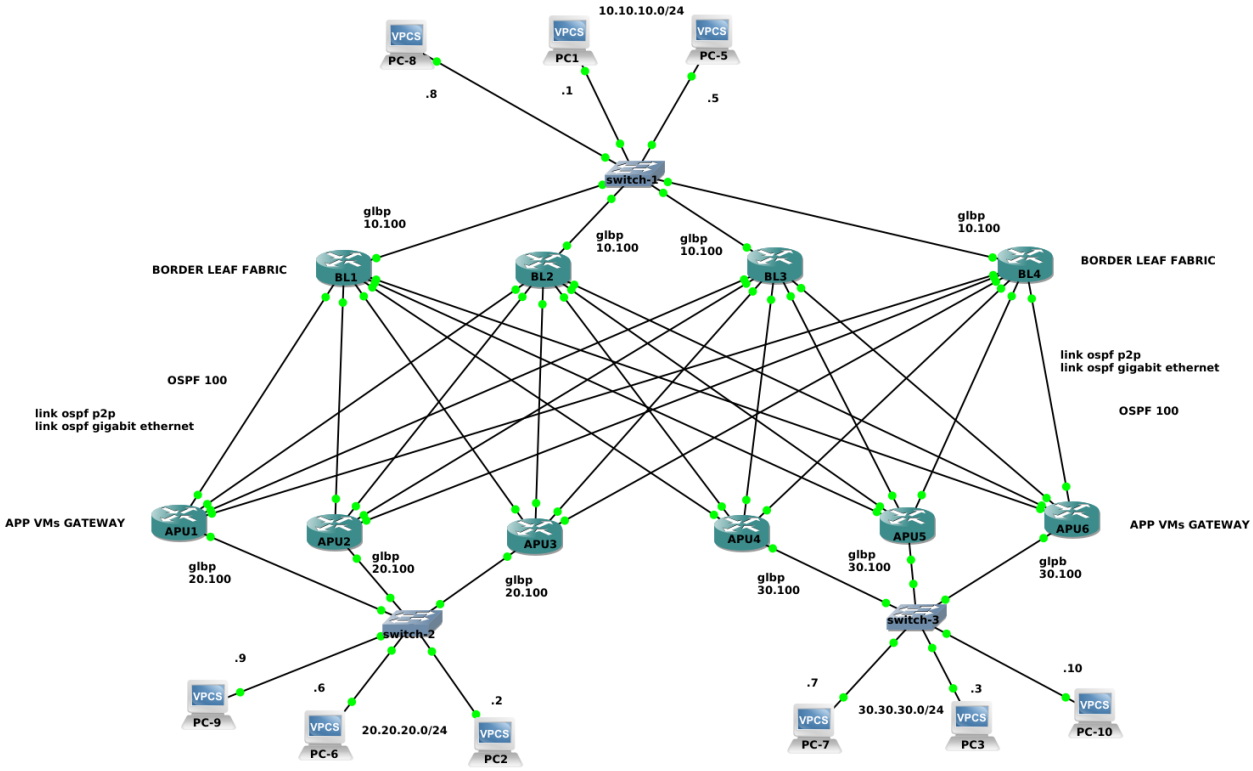

Considerations:

- links ospf are p2p /30 with 1G bandwidth;
- ospf with area 0.0.0.0 intra-area subnets;
- ospf distance amministrative = 110;
- ospf metric = cost of each routers to destination;
- each BL have the same default gateway 10.10.10.100 with GLBP for its VMs;
- APU1, APU2, APU3 have the same default gateway 20.20.20.100 GLBP for its VMs;
- APU4, APU5, APU6 have the same default gateway 30.30.30.100 GLBP for its Vms;
- CEF enable;
- devices GNS3: Cisco 7206VXR (NPE400); 7200 Software (C7200-ADVENTERPRISEK9-M), Version 15.0(1)M3, RELEASE SOFTWARE (fc2)

Target:

- verify the load-balancing per-packet via ospf

- with CEF enable we use the command “ ip load-sharing per-packet “

NOTE

- In per-destination mode all packets for a given destination are forwarded along the same path. This preserves packet order, with potential unequal usage of the links. If one host receives the majority of the traffic all packets will use one link, leaving bandwidth on other links unused.
- Per-packet load balancing guarantees equal load across all links, however potentially the packets may arrive out-of-order at the destination as differential delay may exist within the network.

Before to verify the load-balancing per-packet, we need to report the below output at time-zero:

1) below we see the correct behaviour ospf with a fair load-balancing via its equal cost paths and the active route to the destination for traffic (corresponds to a single packet or an entire flow to a destination, depending on the type of switching configured):

- For process-switching—load balancing is on a per-packet basis and the asterisk (*) points to the interface over which the next packet is sent.
- For fast-switching—load balancing is on a per-destination basis and the asterisk (*) points to the interface over which the next destination-based flow is sent (by default, on Cisco routers, fast switching is enabled under interfaces)
-

CEF allow you to do per-packet and per-destination load-balancing more quickly.
CEF uses a Forwarding Information Base (FIB) to make IP destination prefix-based switching decisions

CEF performs the load-balancing once the routing protocol table is calculated.

From BL: normal ospf operation with load-balancing (4 equal cost path)

```
BL1#sh ip route 20.20.20.0

Routing entry for 20.20.20.0/24

  Known via "ospf 100", distance 110, metric 2,
type intra area

  Last update from 1.1.1.10 on
GigabitEthernet3/0, 01:39:37 ago

  Routing Descriptor Blocks:

    1.1.1.10, from 172.16.1.3, 01:39:37 ago, via
GigabitEthernet3/0

      Route metric is 2, traffic share count is 1

    1.1.1.6, from 172.16.1.2, 01:40:26 ago, via
GigabitEthernet2/0

      Route metric is 2, traffic share count is 1

  * 1.1.1.2, from 172.16.1.1, 01:41:23 ago, via
GigabitEthernet1/0

      Route metric is 2, traffic share count is 1
```

```
Counters interface (with only ospf signaling
traffic)

BL1#sh int gi1/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

36 packets input, 3603 bytes, 0 no buffer

37 packets output, 3711 bytes, 0 underruns

!

BL1#sh int gi2/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

58 packets input, 5591 bytes, 0 no buffer

58 packets output, 5604 bytes, 0 underruns

!

BL1#sh int gi3/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

69 packets input, 6715 bytes, 0 no buffer

69 packets output, 6727 bytes, 0 underruns
```

```
BL1#sh ip route 30.30.30.0

Routing entry for 30.30.30.0/24

  Known via "ospf 100", distance 110, metric 2,
type intra area

  Last update from 1.1.1.22 on
GigabitEthernet6/0, 01:40:41 ago

  Routing Descriptor Blocks:

    1.1.1.22, from 172.16.1.6, 01:40:41 ago, via
GigabitEthernet6/0

      Route metric is 2, traffic share count is 1

    1.1.1.18, from 172.16.1.5, 01:41:24 ago, via
GigabitEthernet5/0

      Route metric is 2, traffic share count is 1

  * 1.1.1.14, from 172.16.1.4, 01:42:04 ago, via
GigabitEthernet4/0

      Route metric is 2, traffic share count is 1
```

```
Counters interface (with only ospf signaling
traffic)

BL1#sh int gi4/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

103 packets input, 10198 bytes, 0 no buffer

103 packets output, 10190 bytes, 0 underruns

!

BL1#sh int gi5/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

138 packets input, 13741 bytes, 0 no buffer

138 packets output, 13437 bytes, 0 underruns

!
```

	BL1#sh int gi6/0 30 sec input rate 0 bits/sec, 0 packets/sec 30 sec output rate 0 bits/sec, 0 packets/sec 160 packets input, 15712 bytes, 0 no buffer 159 packets output, 15623 bytes, 0 underruns

BL2#sh ip route 20.20.20.20
 	
Routing entry for 20.20.20.0/24
 	
  Known via "ospf 100", distance 110, metric 2, type intra area
 	
  Last update from 1.1.1.34 on GigabitEthernet3/0, 01:49:57 ago
 	
  Routing Descriptor Blocks:
 	
  1.1.1.34, from 172.16.1.3, 01:49:57 ago, via GigabitEthernet3/0
 	
  Route metric is 2, traffic share count is 1
 	
  1.1.1.30, from 172.16.1.2, 01:50:46 ago, via GigabitEthernet2/0
 	
  Route metric is 2, traffic share count is 1
 	
  * 1.1.1.26, from 172.16.1.1, 01:51:42 ago, via GigabitEthernet1/0
 	
  Route metric is 2, traffic share count is 1
	Counters interface (with only ospf signaling traffic)
	
	BL2#sh int gi1/0
	
	30 sec input rate 0 bits/sec, 0 packets/sec
	
	30 sec output rate 0 bits/sec, 0 packets/sec
	
	2 packets input, 154 bytes, 0 no buffer
	
	2 packets output, 154 bytes, 0 underruns
	
	!
	
	BL2#sh int gi2/0
	
	30 sec input rate 0 bits/sec, 0 packets/sec
	
	30 sec output rate 0 bits/sec, 0 packets/sec
	
	26 packets input, 2750 bytes, 0 no buffer
	
	25 packets output, 2395 bytes, 0 underruns
	
	!
	
	BL2#sh int gi3/0
	
	30 sec input rate 0 bits/sec, 0 packets/sec
	
	30 sec output rate 0 bits/sec, 0 packets/sec
	
	49 packets input, 4992 bytes, 0 no buffer
	
	49 packets output, 4955 bytes, 0 underruns

BL2#sh ip route 30.30.30.0
 	
Routing entry for 30.30.30.0/24
 	
  Known via "ospf 100", distance 110, metric 2, type intra area
 	
  Last update from 1.1.1.46 on GigabitEthernet6/0, 01:47:48 ago
 	
  Routing Descriptor Blocks:
 	
  1.1.1.46, from 172.16.1.6, 01:47:48 ago, via GigabitEthernet6/0
 	
  Route metric is 2, traffic share count is 1
 	
  1.1.1.42, from 172.16.1.5, 01:48:32 ago, via GigabitEthernet5/0
	Counters interface (with only ospf signaling traffic)
	
	BL2#sh int gi4/0
	
	30 sec input rate 0 bits/sec, 0 packets/sec
	
	30 sec output rate 0 bits/sec, 0 packets/sec
	
	79 packets input, 8050 bytes, 0 no buffer
	
	79 packets output, 8304 bytes, 0 underruns
	
	!
	
	BL2#sh int gi5/0
	
	30 sec input rate 0 bits/sec, 0 packets/sec
	
	30 sec output rate 0 bits/sec, 0 packets/sec

```
      Route metric is 2, traffic share count is 1

  * 1.1.1.38, from 172.16.1.4, 01:49:12 ago, via
GigabitEthernet4/0

      Route metric is 2, traffic share count is 1
```

```
105 packets input, 11223 bytes, 0 no buffer

102 packets output, 11045 bytes, 0 underruns

!

BL2#sh int gi6/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

122 packets input, 12955 bytes, 0 no buffer

121 packets output, 12918 bytes, 0 underruns
```

```
BL3#sh ip route 20.20.20.0

Routing entry for 20.20.20.0/24

  Known via "ospf 100", distance 110, metric 2,
type intra area

  Last update from 1.1.1.58 on
GigabitEthernet3/0, 02:15:39 ago

  Routing Descriptor Blocks:

    1.1.1.58, from 172.16.1.3, 02:15:39 ago, via
GigabitEthernet3/0

      Route metric is 2, traffic share count is 1

    1.1.1.54, from 172.16.1.2, 02:16:28 ago, via
GigabitEthernet2/0

      Route metric is 2, traffic share count is 1

  * 1.1.1.50, from 172.16.1.1, 02:17:25 ago, via
GigabitEthernet1/0

      Route metric is 2, traffic share count is 1
```

```
Counters interface (with only ospf signaling
traffic)

BL3#sh int gi1/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

2 packets input, 154 bytes, 0 no buffer

2 packets output, 154 bytes, 0 underruns

!

BL3#sh int gi2/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

24 packets input, 2402 bytes, 0 no buffer

24 packets output, 2141 bytes, 0 underruns

!

BL3#sh int gi3/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

38 packets input, 3774 bytes, 0 no buffer

38 packets output, 3771 bytes, 0 underruns
```

```
BL3#sh ip route 30.30.30.0

Routing entry for 30.30.30.0/24

  Known via "ospf 100", distance 110, metric 2,
type intra area

  Last update from 1.1.1.70 on
GigabitEthernet6/0, 02:18:00 ago

  Routing Descriptor Blocks:

    1.1.1.70, from 172.16.1.6, 02:18:00 ago, via
GigabitEthernet6/0
```

```
Counters interface (with only ospf signaling
traffic)

BL3#sh int gi4/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

62 packets input, 6176 bytes, 0 no buffer

62 packets output, 6171 bytes, 0 underruns

!
```

```
      Route metric is 2, traffic share count is 1

    1.1.1.66, from 172.16.1.5, 02:18:44 ago, via
GigabitEthernet5/0

      Route metric is 2, traffic share count is 1

  * 1.1.1.62, from 172.16.1.4, 02:19:24 ago, via
GigabitEthernet4/0

      Route metric is 2, traffic share count is 1
```

```
BL3#sh int gi5/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

80 packets input, 7839 bytes, 0 no buffer

79 packets output, 7756 bytes, 0 underruns

!

BL3#sh int gi6/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

95 packets input, 9305 bytes, 0 no buffer

93 packets output, 9127 bytes, 0 underruns
```

```
BL4#sh ip route 20.20.20.20

Routing entry for 20.20.20.0/24

  Known via "ospf 100", distance 110, metric 2,
type intra area

  Last update from 1.1.1.82 on
GigabitEthernet3/0, 02:30:38 ago

  Routing Descriptor Blocks:

    1.1.1.82, from 172.16.1.3, 02:30:38 ago, via
GigabitEthernet3/0

      Route metric is 2, traffic share count is 1

    1.1.1.78, from 172.16.1.2, 02:31:27 ago, via
GigabitEthernet2/0

      Route metric is 2, traffic share count is 1

  * 1.1.1.74, from 172.16.1.1, 02:32:23 ago, via
GigabitEthernet1/0

      Route metric is 2, traffic share count is 1
```

```
Counters interface (with only ospf signaling
traffic)

BL4#sh int gi1/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

1 packets input, 354 bytes, 0 no buffer

1 packets output, 60 bytes, 0 underruns

!

BL4#sh int gi2/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

17 packets input, 1586 bytes, 0 no buffer

18 packets output, 1938 bytes, 0 underruns

!

BL4#sh int gi3/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec

32 packets input, 3035 bytes, 0 no buffer

34 packets output, 3463 bytes, 0 underruns
```

```
BL4#sh ip route 30.30.30.0

Routing entry for 30.30.30.0/24

  Known via "ospf 100", distance 110, metric 2,
type intra area

  Last update from 1.1.1.94 on
GigabitEthernet6/0, 02:31:40 ago
```

```
Counters interface (with only ospf signaling
traffic)

BL4#sh int gi4/0

30 sec input rate 0 bits/sec, 0 packets/sec

30 sec output rate 0 bits/sec, 0 packets/sec
```

Routing Descriptor Blocks: 1.1.1.94, from 172.16.1.6, 02:31:40 ago, via GigabitEthernet6/0 Route metric is 2, traffic share count is 1 1.1.1.90, from 172.16.1.5, 02:32:24 ago, via GigabitEthernet5/0 Route metric is 2, traffic share count is 1 * 1.1.1.86, from 172.16.1.4, 02:33:04 ago, via GigabitEthernet4/0 Route metric is 2, traffic share count is 1	56 packets input, 5731 bytes, 0 no buffer 55 packets output, 5356 bytes, 0 underruns ! BL4#sh int gi5/0 30 sec input rate 0 bits/sec, 0 packets/sec 30 sec output rate 0 bits/sec, 0 packets/sec 70 packets input, 7086 bytes, 0 no buffer 70 packets output, 6787 bytes, 0 underruns ! BL4#sh int gi6/0 30 sec input rate 0 bits/sec, 0 packets/sec 30 sec output rate 0 bits/sec, 0 packets/sec 82 packets input, 8010 bytes, 0 no buffer 82 packets output, 8004 bytes, 0 underruns

Table 1: from BL output show ip route <subnets> and counter interfaces at time-zero

From APU: normal ospf operation with load-balancing (4 equal cost path)	
APU1#sh ip route 10.10.10.0 Routing entry for 10.10.10.0/24 Known via "ospf 100", distance 110, metric 2, type intra area Last update from 1.1.1.25 on GigabitEthernet2/0, 00:23:58 ago Routing Descriptor Blocks: 1.1.1.73, from 192.168.1.1, 00:23:58 ago, via GigabitEthernet4/0 Route metric is 2, traffic share count is 1 1.1.1.49, from 192.168.1.1, 00:23:58 ago, via GigabitEthernet3/0 Route metric is 2, traffic share count is 1 1.1.1.25, from 192.168.1.1, 00:23:58 ago, via GigabitEthernet2/0 Route metric is 2, traffic share count is 1	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table) NOTE: Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route

* 1.1.1.1, from 192.168.1.1, 00:23:58 ago, via GigabitEthernet1/0 Route metric is 2, traffic share count is 1	
APU1#sh ip route 30.30.30.0 Routing entry for 30.30.30.0/24 Known via "ospf 100", distance 110, metric 3, type intra area Last update from 1.1.1.49 on GigabitEthernet3/0, 00:19:57 ago Routing Descriptor Blocks: * 1.1.1.73, from 172.16.1.4, 00:19:57 ago, via GigabitEthernet4/0 Route metric is 3, traffic share count is 1 1.1.1.49, from 172.16.1.4, 00:19:57 ago, via GigabitEthernet3/0 Route metric is 3, traffic share count is 1 1.1.1.25, from 172.16.1.4, 00:19:57 ago, via GigabitEthernet2/0 Route metric is 3, traffic share count is 1 1.1.1.1, from 172.16.1.4, 00:19:57 ago, via GigabitEthernet1/0 Route metric is 3, traffic share count is 1	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table) NOTE: Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route
APU2#sh ip route 10.10.10.0 Routing entry for 10.10.10.0/24 Known via "ospf 100", distance 110, metric 2, type intra area Last update from 1.1.1.5 on GigabitEthernet1/0, 08:18:59 ago Routing Descriptor Blocks: 1.1.1.77, from 192.168.1.4, 08:18:59 ago, via GigabitEthernet4/0 Route metric is 2, traffic share count is 1 1.1.1.53, from 192.168.1.3, 08:18:59 ago, via GigabitEthernet3/0 Route metric is 2, traffic share count is 1 1.1.1.29, from 192.168.1.2, 08:18:59 ago, via GigabitEthernet2/0 Route metric is 2, traffic share count is 1 * 1.1.1.5, from 192.168.1.1, 08:18:59 ago, via GigabitEthernet1/0	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table) NOTE: Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route

Route metric is 2, traffic share count is 1	
APU2#sh ip route 30.30.30.30 Routing entry for 30.30.30.0/24 Known via "ospf 100", distance 110, metric 3, type intra area Last update from 1.1.1.53 on GigabitEthernet3/0, 08:21:11 ago Routing Descriptor Blocks: * 1.1.1.77, from 172.16.1.4, 08:21:11 ago, via GigabitEthernet4/0 Route metric is 3, traffic share count is 1 1.1.1.53, from 172.16.1.4, 08:21:11 ago, via GigabitEthernet3/0 Route metric is 3, traffic share count is 1 1.1.1.29, from 172.16.1.4, 08:21:11 ago, via GigabitEthernet2/0 Route metric is 3, traffic share count is 1 1.1.1.5, from 172.16.1.4, 08:21:11 ago, via GigabitEthernet1/0 Route metric is 3, traffic share count is 1	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table) NOTE: Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route
APU3#sh ip route 10.10.10.0 Routing entry for 10.10.10.0/24 Known via "ospf 100", distance 110, metric 2, type intra area Last update from 1.1.1.9 on GigabitEthernet1/0, 08:24:39 ago Routing Descriptor Blocks: 1.1.1.81, from 192.168.1.4, 08:24:39 ago, via GigabitEthernet4/0 Route metric is 2, traffic share count is 1 1.1.1.57, from 192.168.1.3, 08:24:39 ago, via GigabitEthernet3/0 Route metric is 2, traffic share count is 1 1.1.1.33, from 192.168.1.2, 08:24:39 ago, via GigabitEthernet2/0 Route metric is 2, traffic share count is 1 * 1.1.1.9, from 192.168.1.1, 08:24:39 ago, via GigabitEthernet1/0 Route metric is 2, traffic share count is 1	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table) NOTE: Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route
APU3#sh ip route 30.30.30.30	Counters interface with only ospf signaling

Routing entry for 30.30.30.0/24 Known via "ospf 100", distance 110, metric 3, type intra area Last update from 1.1.1.57 on GigabitEthernet3/0, 08:25:30 ago Routing Descriptor Blocks: *** 1**.1.1.81, from 172.16.1.4, 08:25:30 ago, via GigabitEthernet4/0 Route metric is 3, traffic share count is 1 1.1.1.57, from 172.16.1.4, 08:25:30 ago, via GigabitEthernet3/0 Route metric is 3, traffic share count is 1 1.1.1.33, from 172.16.1.4, 08:25:30 ago, via GigabitEthernet2/0 Route metric is 3, traffic share count is 1 1.1.1.9, from 172.16.1.4, 08:25:30 ago, via GigabitEthernet1/0 Route metric is 3, traffic share count is 1	traffic (the value is similar to value indicated above on the other table) NOTE: Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route
APU4#sh ip route 10.10.10.10 Routing entry for 10.10.10.0/24 Known via "ospf 100", distance 110, metric 2, type intra area Last update from 1.1.1.13 on GigabitEthernet1/0, 08:30:14 ago Routing Descriptor Blocks: 1.1.1.85, from 192.168.1.4, 08:30:14 ago, via GigabitEthernet4/0 Route metric is 2, traffic share count is 1 1.1.1.61, from 192.168.1.3, 08:30:14 ago, via GigabitEthernet3/0 Route metric is 2, traffic share count is 1 1.1.1.37, from 192.168.1.2, 08:30:14 ago, via GigabitEthernet2/0 Route metric is 2, traffic share count is 1 * 1.1.1.13, from 192.168.1.1, 08:30:14 ago, via GigabitEthernet1/0 Route metric is 2, traffic share count is 1	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table) NOTE: Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route
APU4#sh ip route 20.20.20.0	Counters interface with only ospf signaling traffic (the value is similar to value indicated

Routing entry for 20.20.20.0/24
 	
  Known via "ospf 100", distance 110, metric 3, type intra area
 	
  Last update from 1.1.1.13 on GigabitEthernet1/0, 08:28:17 ago
 	
  Routing Descriptor Blocks:
 	
    1.1.1.85, from 172.16.1.1, 08:28:17 ago, via GigabitEthernet4/0
 	
      Route metric is 3, traffic share count is 1
 	
    1.1.1.61, from 172.16.1.1, 08:28:17 ago, via GigabitEthernet3/0
 	
      Route metric is 3, traffic share count is 1
 	
  * 1.1.1.37, from 172.16.1.1, 08:28:17 ago, via GigabitEthernet2/0
 	
      Route metric is 3, traffic share count is 1
 	
    1.1.1.13, from 172.16.1.1, 08:28:17 ago, via GigabitEthernet1/0
 	
      Route metric is 3, traffic share count is 1
	above on the other table)
	
	NOTE:
	
	Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route

APU5#sh ip route 10.10.10.0
 	
Routing entry for 10.10.10.0/24
 	
  Known via "ospf 100", distance 110, metric 2, type intra area
 	
  Last update from 1.1.1.17 on GigabitEthernet1/0, 08:31:28 ago
 	
  Routing Descriptor Blocks:
 	
    1.1.1.89, from 192.168.1.4, 08:31:28 ago, via GigabitEthernet4/0
 	
      Route metric is 2, traffic share count is 1
 	
    1.1.1.65, from 192.168.1.3, 08:31:28 ago, via GigabitEthernet3/0
 	
      Route metric is 2, traffic share count is 1
 	
    1.1.1.41, from 192.168.1.2, 08:31:28 ago, via GigabitEthernet2/0
 	
      Route metric is 2, traffic share count is 1
 	
  * 1.1.1.17, from 192.168.1.1, 08:31:28 ago, via GigabitEthernet1/0
 	
      Route metric is 2, traffic share count is 1
	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table)
	
	NOTE:
	
	Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route

APU5#sh ip route 20.20.20.0 Routing entry for 20.20.20.0/24 Known via "ospf 100", distance 110, metric 3,	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table)

type intra area
 	
  Last update from 1.1.1.17 on GigabitEthernet1/0, 08:29:57 ago
 	
  Routing Descriptor Blocks:
 	
  1.1.1.89, from 172.16.1.1, 08:29:57 ago, via GigabitEthernet4/0
 	
  Route metric is 3, traffic share count is 1
 	
  1.1.1.65, from 172.16.1.1, 08:29:57 ago, via GigabitEthernet3/0
 	
  Route metric is 3, traffic share count is 1
 	
  * 1.1.1.41, from 172.16.1.1, 08:29:57 ago, via GigabitEthernet2/0
 	
  Route metric is 3, traffic share count is 1
 	
  1.1.1.17, from 172.16.1.1, 08:29:57 ago, via GigabitEthernet1/0
 	
  Route metric is 3, traffic share count is 1
	NOTE:
	
	Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route

APU6#sh ip route 10.10.10.0
 	
Routing entry for 10.10.10.0/24
 	
  Known via "ospf 100", distance 110, metric 2, type intra area
 	
  Last update from 1.1.1.21 on GigabitEthernet1/0, 08:32:41 ago
 	
  Routing Descriptor Blocks:
 	
  1.1.1.93, from 192.168.1.4, 08:32:41 ago, via GigabitEthernet4/0
 	
  Route metric is 2, traffic share count is 1
 	
  1.1.1.69, from 192.168.1.3, 08:32:41 ago, via GigabitEthernet3/0
 	
  Route metric is 2, traffic share count is 1
 	
  1.1.1.45, from 192.168.1.2, 08:32:41 ago, via GigabitEthernet2/0
 	
  Route metric is 2, traffic share count is 1
 	
  * 1.1.1.21, from 192.168.1.1, 08:32:41 ago, via GigabitEthernet1/0
 	
  Route metric is 2, traffic share count is 1
	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table)
	
	NOTE:
	
	Correctly ospf load-balance on 4 equal cost path and choose (*) only one way as active route

APU6#sh ip route 20.20.20.0
 	
Routing entry for 20.20.20.0/24
 	
  Known via "ospf 100", distance 110, metric 3, type intra area
 	
  Last update from 1.1.1.21 on
	Counters interface with only ospf signaling traffic (the value is similar to value indicated above on the other table)
	
	NOTE:
	
	Correctly ospf load-balance on 4 equal cost path

GigabitEthernet1/0, 08:31:31 ago
 	
  Routing Descriptor Blocks:
 	
    1.1.1.93, from 172.16.1.1, 08:31:31 ago, via GigabitEthernet4/0
 	
      Route metric is 3, traffic share count is 1
 	
    1.1.1.69, from 172.16.1.1, 08:31:31 ago, via GigabitEthernet3/0
 	
      Route metric is 3, traffic share count is 1
 	
  * 1.1.1.45, from 172.16.1.1, 08:31:31 ago, via GigabitEthernet2/0
 	
      Route metric is 3, traffic share count is 1
 	
    1.1.1.21, from 172.16.1.1, 08:31:31 ago, via GigabitEthernet1/0
 	
      Route metric is 3, traffic share count is 1
	and choose (*) only one way as active route

Table 2: from APU output show ip route <subnets> and counter interfaces time-zero

2) GLBP CONSIDERATION

GLBP (Global Load Balancing Protocol) work correctly balancing each sourced PCs to its default gateway through the three APU gateway.

Enviroment Data Centers: we should have Distributed Anycast Protocol to overcome the limit to have only two routers gateway (which are hsrp or vrrp), building a VXLAN EVPN VTEP Fabric; one advantage of anycast network is load-balancing where dynamic layer 3 routing of Anycast IP addresses nicely load balances traffic over different nodes based on geography. If equal cost route paths are visible from one geography, all nodes can be used.

First NOTE: GLBP is the first level of traffic balancing

See the below example:

GLBP work from PC2, PC6, PC9 (source 20.20.20.0/24 to destinations 10.10.10.1 and 30.30.30.3)	
PC2> ping 10.10.10.1 84 bytes from 10.10.10.1 icmp_seq=1 ttl=62 time=39.122 ms 84 bytes from 10.10.10.1 icmp_seq=2 ttl=62 time=14.786 ms 84 bytes from 10.10.10.1 icmp_seq=3 ttl=62 time=24.343 ms 84 bytes from 10.10.10.1 icmp_seq=4 ttl=62 time=24.436 ms 84 bytes from 10.10.10.1 icmp_seq=5 ttl=62 time=24.439 ms ! PC2> ping 30.30.30.3 84 bytes from 30.30.30.3 icmp_seq=1 ttl=61 time=44.508 ms 84 bytes from 30.30.30.3 icmp_seq=2 ttl=61 time=35.446 ms 84 bytes from 30.30.30.3 icmp_seq=3 ttl=61 time=35.667 ms 84 bytes from 30.30.30.3 icmp_seq=4 ttl=61 time=25.657 ms 84 bytes from 30.30.30.3 icmp_seq=5 ttl=61 time=26.110 ms	PC2> trace 10.10.10.1 trace to 10.10.10.1, 8 hops max, press Ctrl+C to stop 1 20.20.20.102 9.824 ms 8.714 ms 12.471 ms → APU2 2 1.1.1.29 18.673 ms 19.824 ms 20.039 ms 3 *10.10.10.1 28.134 ms ! PC2> trace 30.30.30.3 trace to 30.30.30.3, 8 hops max, press Ctrl+C to stop 1 20.20.20.102 2.256 ms 9.106 ms 9.633 ms → APU2 2 1.1.1.53 22.785 ms 19.628 ms 19.486 ms 3 1.1.1.66 36.031 ms 39.324 ms 39.482 ms 4 *30.30.30.3 39.571 ms
PC-6> ping 10.10.10.1 84 bytes from 10.10.10.1 icmp_seq=1 ttl=62 time=37.439 ms 84 bytes from 10.10.10.1 icmp_seq=2 ttl=62 time=21.052 ms 84 bytes from 10.10.10.1 icmp_seq=3 ttl=62 time=25.919 ms 84 bytes from 10.10.10.1 icmp_seq=4 ttl=62 time=21.895 ms 84 bytes from 10.10.10.1 icmp_seq=5 ttl=62 time=27.355 ms ! PC-6> ping 30.30.30.3 84 bytes from 30.30.30.3 icmp_seq=1 ttl=61 time=44.179 ms 84 bytes from 30.30.30.3 icmp_seq=2 ttl=61 time=48.152 ms 84 bytes from 30.30.30.3 icmp_seq=3 ttl=61 time=47.594 ms 84 bytes from 30.30.30.3 icmp_seq=4 ttl=61 time=46.123 ms 84 bytes from 30.30.30.3 icmp_seq=5 ttl=61 time=37.160 ms	PC-6> trace 10.10.10.1 trace to 10.10.10.1, 8 hops max, press Ctrl+C to stop 1 20.20.20.103 6.575 ms 9.247 ms 9.334 ms → APU3 2 1.1.1.33 30.015 ms 29.874 ms 29.461 ms 3 *10.10.10.1 29.990 ms ! PC-6> trace 30.30.30.3 trace to 30.30.30.3, 8 hops max, press Ctrl+C to stop 1 20.20.20.103 7.398 ms 8.996 ms 15.996 ms → APU3 2 1.1.1.33 22.724 ms 29.759 ms 29.987 ms 3 1.1.1.38 49.680 ms 39.191 ms 39.347 ms 4 *30.30.30.3 39.572 ms
PC-9> ping 10.10.10.1 84 bytes from 10.10.10.1 icmp_seq=1 ttl=62 time=17.680 ms	PC-9> trace 10.10.10.1 trace to 10.10.10.1, 8 hops max, press Ctrl+C to stop 1 20.20.20.101 9.590 ms 9.109 ms 10.068 ms → APU1

84 bytes from 10.10.10.1 icmp_seq=2 ttl=62 time=14.544 ms 84 bytes from 10.10.10.1 icmp_seq=3 ttl=62 time=15.115 ms 84 bytes from 10.10.10.1 icmp_seq=4 ttl=62 time=15.152 ms 84 bytes from 10.10.10.1 icmp_seq=5 ttl=62 time=24.721 ms ! PC-9> ping 30.30.30.3 84 bytes from 30.30.30.3 icmp_seq=1 ttl=61 time=41.544 ms 84 bytes from 30.30.30.3 icmp_seq=2 ttl=61 time=37.737 ms 84 bytes from 30.30.30.3 icmp_seq=3 ttl=61 time=38.113 ms 84 bytes from 30.30.30.3 icmp_seq=4 ttl=61 time=26.412 ms 84 bytes from 30.30.30.3 icmp_seq=5 ttl=61 time=28.867 ms	2 1.1.1.73 19.288 ms 19.564 ms 19.634 ms 3 *10.10.10.1 23.332 ms ! PC-9> trace 30.30.30.3 trace to 30.30.30.3, 8 hops max, press Ctrl+C to stop 1 20.20.20.101 7.530 ms 9.132 ms 9.666 ms → APU1 2 1.1.1.25 20.024 ms 19.500 ms 20.055 ms 3 1.1.1.42 29.079 ms 29.660 ms 29.227 ms 4 *30.30.30.3 46.447 ms

Table 3: GLBP working

3) CEF CONSIDERATION

Cisco Express Forwarding is enabled on APU and BL routers

CEF polarization is the effect when a hash algorithm chooses a particular path and the redundant paths remain completely unused.

CEF switches the packets based on the routing table that is populated by the routing protocols which is OSPF.

CEF performs the load-balancing once the routing protocol table is calculated.

Each traffic from sourced PC is CEF switched; because we have 4 equal cost path (for each APU routers) the source and destination addresses in the packet go through an hash algorithm.

The result is a specific path used to reach the destinations

How IP CEF work with per-destination switching	
PC2> ping 10.10.10.1 **-c 70** PC2> trace 10.10.10.1 trace to 10.10.10.1, 8 hops max, press Ctrl+C to stop 1 20.20.20.102 4.465 ms 9.198 ms 9.744 ms → APU2 2 1.1.1.29 22.902 ms 19.637 ms 19.660 ms → gi2/0	APU2#sh ip cef 10.10.10.1 detail 10.10.10.0/24, epoch 0, per-destination sharing **76 packets, 6432 bytes switched to the prefix** nexthop 1.1.1.5 GigabitEthernet1/0 nexthop 1.1.1.29 GigabitEthernet2/0

3 *10.10.10.1 25.914 ms → BL2	nexthop 1.1.1.53 GigabitEthernet3/0 nexthop 1.1.1.77 GigabitEthernet4/0 ! BL1#sh ip cef 20.20.20.0 det 20.20.20.0/24, epoch 1, per-destination sharing **73 packets, 6096 bytes switched to the prefix** nexthop 1.1.1.2 GigabitEthernet1/0 nexthop 1.1.1.6 GigabitEthernet2/0 nexthop 1.1.1.10 GigabitEthernet3/0
PC2> ping 30.30.30.3 **-c 100** PC2> trace 30.30.30.3 trace to 30.30.30.3, 8 hops max, press Ctrl+C to stop 1 20.20.20.103 4.303 ms 9.068 ms 9.958 ms → APU3 2 1.1.1.9 29.031 ms 23.938 ms 19.385 ms 3 1.1.1.22 34.097 ms 29.897 ms 29.419 ms 4 *30.30.30.3 45.042 ms	APU3#sh ip cef 30.30.30.3 detail 30.30.30.0/24, epoch 0, per-destination sharing **109 packets, 9228 bytes switched to the prefix** nexthop 1.1.1.9 GigabitEthernet1/0 → to BL-1 nexthop 1.1.1.33 GigabitEthernet2/0 nexthop 1.1.1.57 GigabitEthernet3/0 nexthop 1.1.1.81 GigabitEthernet4/0 ! BL1#sh ip cef 30.30.30.3 det 30.30.30.0/24, epoch 1, per-destination sharing **106 packets, 8952 bytes switched to the prefix** nexthop 1.1.1.14 GigabitEthernet4/0 nexthop 1.1.1.18 GigabitEthernet5/0 nexthop 1.1.1.22 GigabitEthernet6/0 → to APU-6 ! TO 20.20.20.0 return traffic APU4#sh ip cef 20.20.20.0 det 20.20.20.0/24, epoch 0, per-destination sharing **103 packets, 8616 bytes switched to the prefix** nexthop 1.1.1.13 GigabitEthernet1/0 nexthop 1.1.1.37 GigabitEthernet2/0 nexthop 1.1.1.61 GigabitEthernet3/0 nexthop 1.1.1.85 GigabitEthernet4/0

PC2> ping 30.30.30.3 **-c 800**

PC2> trace 30.30.30.3

trace to 30.30.30.3, 8 hops max, press Ctrl+C to stop

1 20.20.20.103 10.000 ms 9.417 ms 9.886 ms → APU-3

2 1.1.1.9 26.095 ms 19.185 ms 19.574 ms

3 1.1.1.22 31.716 ms 30.986 ms 40.107 ms

4 *30.30.30.3 50.261 ms

APU3#sh ip cef 30.30.30.3 det

30.30.30.0/24, epoch 0, per-destination sharing

9 packets, 828 bytes switched to the prefix

nexthop 1.1.1.9 GigabitEthernet1/0

nexthop 1.1.1.33 GigabitEthernet2/0

nexthop 1.1.1.57 GigabitEthernet3/0

nexthop 1.1.1.81 GigabitEthernet4/0

!

APU2#sh ip cef 30.30.30.3 det

30.30.30.0/24, epoch 0, per-destination sharing

800 packets, 67200 bytes switched to the prefix

nexthop 1.1.1.5 GigabitEthernet1/0

nexthop 1.1.1.29 GigabitEthernet2/0

nexthop 1.1.1.53 GigabitEthernet3/0

nexthop 1.1.1.77 GigabitEthernet4/0

!

BL1#sh ip cef 30.30.30.0 det

30.30.30.0/24, epoch 1, per-destination sharing

6 packets, 552 bytes switched to the prefix

nexthop 1.1.1.14 GigabitEthernet4/0

nexthop 1.1.1.18 GigabitEthernet5/0

nexthop 1.1.1.22 GigabitEthernet6/0

!

BL3#sh ip cef 30.30.30.3 det

30.30.30.0/24, epoch 0, per-destination sharing

800 packets, 67200 bytes switched to the prefix

nexthop 1.1.1.62 GigabitEthernet4/0

nexthop 1.1.1.66 GigabitEthernet5/0

nexthop 1.1.1.70 GigabitEthernet6/0

!

TO 20.20.20.0 return traffic

APU4#sh ip cef 20.20.20.0 det

20.20.20.0/24, epoch 0, per-destination sharing

335 packets, 28104 bytes switched to the prefix

nexthop 1.1.1.13 GigabitEthernet1/0

	nexthop 1.1.1.37 GigabitEthernet2/0 nexthop 1.1.1.61 GigabitEthernet3/0 nexthop 1.1.1.85 GigabitEthernet4/0 ! APU5#sh ip cef 20.20.20.0 det 20.20.20.0/24, epoch 0, per-destination sharing **234 packets, 19656 bytes switched to the prefix** nexthop 1.1.1.17 GigabitEthernet1/0 nexthop 1.1.1.41 GigabitEthernet2/0 nexthop 1.1.1.65 GigabitEthernet3/0 nexthop 1.1.1.89 GigabitEthernet4/0 ! APU6#sh ip cef 20.20.20.0 det 20.20.20.0/24, epoch 0, per-destination sharing **234 packets, 19656 bytes switched to the prefix** nexthop 1.1.1.21 GigabitEthernet1/0 nexthop 1.1.1.45 GigabitEthernet2/0 nexthop 1.1.1.69 GigabitEthernet3/0 nexthop 1.1.1.93 GigabitEthernet4/0

Table 4: CISCO IP CEF per-destination switching

4) We enable load-balancing per packet with the command " ip load-sharing per-packet " under the egress interface ospf

How IP CEF work with per-packet switching	
PC2> ping 10.10.10.1 -c 50 PC2> trace 10.10.10.1 trace to 10.10.10.1, 8 hops max, press Ctrl+C to stop 1 20.20.20.101 14.544 ms 9.793 ms 9.322 ms → APU-1 2 1.1.1.49 19.221 ms 19.372 ms 19.801 ms 3 *10.10.10.1 36.354 ms	APU1#sh ip cef 10.10.10.1 detail 10.10.10.0/24, epoch 0, per-packet sharing **56 packets, 4752 bytes switched to the prefix** nexthop 1.1.1.1 GigabitEthernet1/0 nexthop 1.1.1.25 GigabitEthernet2/0 nexthop 1.1.1.49 GigabitEthernet3/0 nexthop 1.1.1.73 GigabitEthernet4/0 ! TO 20.20.20.0 return traffic BL2#sh ip cef 20.20.20.0 detail 20.20.20.0/24, epoch 0, per-packet sharing 53 packets, 4416 bytes switched to the prefix nexthop 1.1.1.26 GigabitEthernet1/0 nexthop 1.1.1.30 GigabitEthernet2/0 nexthop 1.1.1.34 GigabitEthernet3/0
PC2> ping 10.10.10.1 -**c 150** PC2> trace 10.10.10.1 trace to 10.10.10.1, 8 hops max, press Ctrl+C to stop 1 20.20.20.103 8.888 ms 9.842 ms 7.632 ms → APU-3 2 1.1.1.9 15.726 ms 24.311 ms 15.147 ms 3 *10.10.10.1 19.400 ms	APU3#sh ip cef 10.10.10.0 det 10.10.10.0/24, epoch 0, per-packet sharing **6 packets, 552 bytes switched to the prefix** nexthop 1.1.1.9 GigabitEthernet1/0 nexthop 1.1.1.33 GigabitEthernet2/0 nexthop 1.1.1.57 GigabitEthernet3/0 nexthop 1.1.1.81 GigabitEthernet4/0 ! APU2#sh ip cef 10.10.10.1 detail 10.10.10.0/24, epoch 0, per-packet sharing **150 packets, 12600 bytes switched to the prefix** nexthop 1.1.1.5 GigabitEthernet1/0 nexthop 1.1.1.29 GigabitEthernet2/0

	nexthop 1.1.1.53 GigabitEthernet3/0 nexthop 1.1.1.77 GigabitEthernet4/0 ! TO 20.20.20.0 return traffic BL3#sh ip cef 20.20.20.0 det 20.20.20.0/24, epoch 0, per-packet sharing **119 packets, 9996 bytes switched to the prefix** nexthop 1.1.1.50 GigabitEthernet1/0 nexthop 1.1.1.54 GigabitEthernet2/0 nexthop 1.1.1.58 GigabitEthernet3/0 ! BL4#sh ip cef 20.20.20.0 detail 20.20.20.0/24, epoch 0, per-packet sharing **34 packets, 2820 bytes switched to the prefix** nexthop 1.1.1.74 GigabitEthernet1/0 nexthop 1.1.1.78 GigabitEthernet2/0 nexthop 1.1.1.82 GigabitEthernet3/0
PC2> ping 10.10.10.1 -**c 300** PC2> trace 10.10.10.1 trace to 10.10.10.1, 8 hops max, press Ctrl+C to stop 1 20.20.20.102 3.405 ms 9.208 ms 9.758 ms → APU-2 2 1.1.1.53 15.376 ms 13.418 ms 23.580 ms 3 *10.10.10.1 8.979 ms	APU2#sh ip cef 10.10.10.1 detail 10.10.10.0/24, epoch 0, per-packet sharing **6 packets, 552 bytes switched to the prefix** nexthop 1.1.1.5 GigabitEthernet1/0 nexthop 1.1.1.29 GigabitEthernet2/0 nexthop 1.1.1.53 GigabitEthernet3/0 nexthop 1.1.1.77 GigabitEthernet4/0 ! APU1#sh ip cef 10.10.10.1 detail 10.10.10.0/24, epoch 0, per-packet sharing **300 packets, 25200 bytes switched to the prefix** nexthop 1.1.1.1 GigabitEthernet1/0 nexthop 1.1.1.25 GigabitEthernet2/0 nexthop 1.1.1.49 GigabitEthernet3/0 nexthop 1.1.1.73 GigabitEthernet4/0 ! BL1#sh ip cef 20.20.20.0 det

	20.20.20.0/24, epoch 1, per-packet sharing **118 packets, 9912 bytes switched to the prefix** nexthop 1.1.1.2 GigabitEthernet1/0 nexthop 1.1.1.6 GigabitEthernet2/0 nexthop 1.1.1.10 GigabitEthernet3/0 ! TO 20.20.20.0 return traffic BL2#sh ip cef 20.20.20.0 detail 20.20.20.0/24, epoch 0, per-packet sharing **118 packets, 9912 bytes switched to the prefix** nexthop 1.1.1.26 GigabitEthernet1/0 nexthop 1.1.1.30 GigabitEthernet2/0 nexthop 1.1.1.34 GigabitEthernet3/0 ! BL3#sh ip cef 20.20.20.0 det 20.20.20.0/24, epoch 0, per-packet sharing **67 packets, 5592 bytes switched to the prefix** nexthop 1.1.1.50 GigabitEthernet1/0 nexthop 1.1.1.54 GigabitEthernet2/0 nexthop 1.1.1.58 GigabitEthernet3/0 !
PC2> ping 30.30.30.3 **-c 800** PC2> trace 30.30.30.3 trace to 30.30.30.3, 8 hops max, press Ctrl+C to stop 1 20.20.20.101 9.395 ms 10.055 ms 9.426 ms –> APU-1 2 1.1.1.1 19.994 ms 22.698 ms 16.355 ms 3 1.1.1.94 43.197 ms 36.617 ms 43.374 ms 4 *30.30.30.3 37.218 ms	APU1#sh ip cef 30.30.30.3 detail 30.30.30.0/24, epoch 0, per-packet sharing **9 packets, 828 bytes switched to the prefix** nexthop 1.1.1.1 GigabitEthernet1/0 nexthop 1.1.1.25 GigabitEthernet2/0 nexthop 1.1.1.49 GigabitEthernet3/0 nexthop 1.1.1.73 GigabitEthernet4/0 ! APU3#sh ip cef 30.30.30.3 detail 30.30.30.0/24, epoch 0, per-packet sharing **800 packets, 67200 bytes switched to the prefix** nexthop 1.1.1.9 GigabitEthernet1/0 nexthop 1.1.1.33 GigabitEthernet2/0

	nexthop 1.1.1.57 GigabitEthernet3/0 nexthop 1.1.1.81 GigabitEthernet4/0 ! BL1#sh ip cef 30.30.30.3 detail 30.30.30.0/24, epoch 1, per-packet sharing **202 packets, 16984 bytes switched to the prefix** nexthop 1.1.1.14 GigabitEthernet4/0 nexthop 1.1.1.18 GigabitEthernet5/0 nexthop 1.1.1.22 GigabitEthernet6/0 ! BL2#sh ip cef 30.30.30.3 det 30.30.30.0/24, epoch 0, per-packet sharing **201 packets, 16892 bytes switched to the prefix** nexthop 1.1.1.38 GigabitEthernet4/0 nexthop 1.1.1.42 GigabitEthernet5/0 nexthop 1.1.1.46 GigabitEthernet6/0 ! BL3#sh ip cef 30.30.30.3 detail 30.30.30.0/24, epoch 0, per-packet sharing **201 packets, 16892 bytes switched to the prefix** nexthop 1.1.1.62 GigabitEthernet4/0 nexthop 1.1.1.66 GigabitEthernet5/0 nexthop 1.1.1.70 GigabitEthernet6/0 ! BL4#sh ip cef 30.30.30.3 detail 30.30.30.0/24, epoch 0, per-packet sharing **202 packets, 16984 bytes switched to the prefix** nexthop 1.1.1.86 GigabitEthernet4/0 nexthop 1.1.1.90 GigabitEthernet5/0 nexthop 1.1.1.94 GigabitEthernet6/0 ! TO 20.20.20.0/24 return traffic APU4#sh ip cef 20.20.20.0 detail 20.20.20.0/24, epoch 0, per-packet sharing

	232 packets, 19488 bytes switched to the prefix nexthop 1.1.1.13 GigabitEthernet1/0 nexthop 1.1.1.37 GigabitEthernet2/0 nexthop 1.1.1.61 GigabitEthernet3/0 nexthop 1.1.1.85 GigabitEthernet4/0 ! APU5#sh ip cef 20.20.20.0 detail 20.20.20.0/24, epoch 0, per-packet sharing **334 packets, 28056 bytes switched to the prefix** nexthop 1.1.1.17 GigabitEthernet1/0 nexthop 1.1.1.41 GigabitEthernet2/0 nexthop 1.1.1.65 GigabitEthernet3/0 nexthop 1.1.1.89 GigabitEthernet4/0 ! APU6#sh ip cef 20.20.20.0 detail 20.20.20.0/24, epoch 0, per-packet sharing **237 packets, 19872 bytes switched to the prefix** nexthop 1.1.1.21 GigabitEthernet1/0 nexthop 1.1.1.45 GigabitEthernet2/0 nexthop 1.1.1.69 GigabitEthernet3/0 nexthop 1.1.1.93 GigabitEthernet4/0

Table 5: CISCO IP CEF per-packet switching