Load Balancing in BGP Networks

By Ivan Pepelnjak

1. 11. 2007

A few years ago, the traditional wisdom was that you could not do load balancing in networks using Border Gateway Protocol (BGP) as their core routing protocol. The technology (actually, its implementation in Cisco IOS) has evolved since then, resulting in a number of load balancing options for BGP-based networks. However, even though it's possible to load-balance in BGP networks, it's still not as easy as Interior Gateway Protocol (IGP; for example OSPF or EIGRP) based load balancing, which happens automatically.

Definitions: Load sharing is the ability to distribute outgoing traffic (or influence the flow of incoming traffic) over multiple paths. Load balancing is the ability to split the load toward the same destination (host or IP prefix) over multiple paths.

Setting the Stage

Before we start discussing load-balancing options available with BGP, let's skim through a few generic load-balancing facts.

Fact #1: Load-balancing is always unidirectional

The load-balancing mechanisms are part of IP packet forwarding performed in the data (forwarding) plane of the layer-3 switches (sometimes known as routers). As IP transport is not connection oriented, IP packets flowing in one direction are not necessarily taking the same path as IP packets flowing in the reverse direction. Load balancing thus has to be tested in each direction separately. For example, in the network shown in Figure 1, the load balancing from the clients to the server is automatic, while the load balancing in the reverse direction is very hard to achieve unless the server itself acts like a very good router; a few alternative approaches are also described in the IP Corner article *Perfect Load-Balancing: How Close Can You Get?*

Figure 1: Simple network design with load balancing challenges

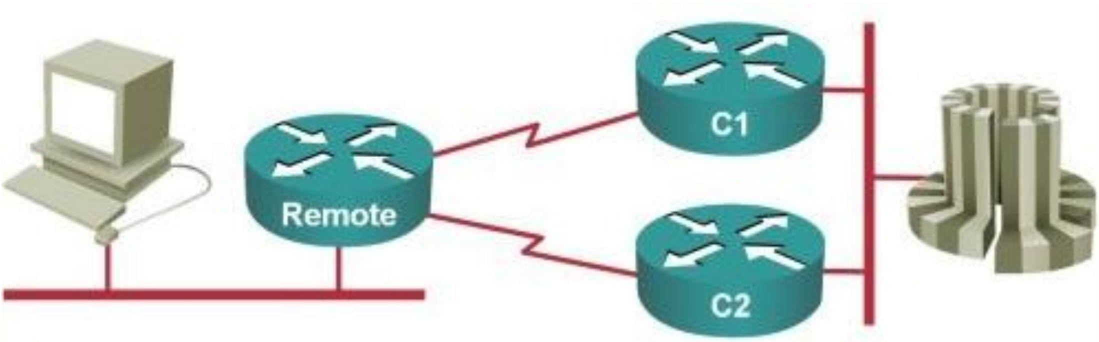

Fact #2: Not all Load-Balancing Mechanisms are Equal

Load balancing toward a destination IP prefix can happen only if the forwarding device has two equal-cost paths to the destination (EIGRP also supports unequal-cost load balancing). Even then, although multiple paths are installed in the IP routing table, only one of them might actually be used to forward the traffic, depending on the forwarding (sometimes called *switching*) mechanism used in the device. For example, if you use Cisco IOS *fast switching* in combination with default routing, only one of the default routes will be used, as fast switching uses per-prefix (not per-destination-host) load balancing.

Rule: If at all possible, use Cisco Express Forwarding (CEF) to achieve better load-balancing results.

CEF can use three load distribution methods:

- *Per-host load-balancing*, where the algorithm selects the outgoing path based on the combination of the source and destination IP address.
- *Per-port load balancing*, where even distinct layer-4 (TCP or UDP) sessions between the same pair of hosts can flow over different paths.
- *Per-packet load balancing*, where each successive packet is sent onto a different outbound path.

Per-packet load balancing can easily result in out-of-order packets that significantly reduce TCP session throughput or result in loss of data in some UDP-based protocols, for example SNA or NetBIOS Fast Sequenced Transport (FST) or Voice-over-IP (VoIP) transport. Even worse, out-of-order packets might be interpreted as attacks by some firewalls.

Rule: Avoid per-packet load balancing at all cost.

Last but not least, sometimes you should replace IP-based load balancing with layer-2 mechanisms, for example link bundling with multilink PPP (for serial links) or EtherChannel (for Ethernet-based point-to-point links).

IGP-based Load Balancing

The only automatic load balancing you get with BGP is the distribution of load over equal-cost paths between ingress and egress points within the same autonomous system (*Edge-1* and *Edge-2* in Figure 2 assuming both paths between them have equal cost) due to the fact that the forwarding of an IP packet sent toward a BGP destination always involves a two-step lookup:

- A BGP route in the IP routing table never points to an actual interface, but toward a BGP next-hop (unless you've messed up your BGP design with route-reflectors using next-hop-self, this should be an IP address directly connected to the egress BGP router).
- The second (recursive) lookup is performed in the IP routing table to find the actual outgoing interface to use in packet forwarding.

Figure 2: IGP-based load balancing between the edge routers

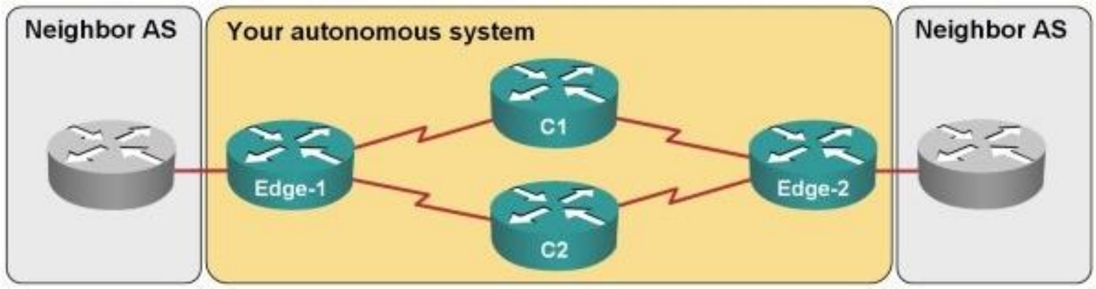

Conclusion: The load balancing between an ingress BGP router and an egress BGP router is automatic. No BGP-specific configuration is needed to fully utilize internal network resources, as long as they offer multiple equal-cost paths between the edge routers.

Note: The edge-to-edge load balancing works even better if you deploy MPLS in your core network, as the core routers don't have to perform IP lookup for BGP destinations. MPLS Traffic Engineering enables you to deploy unequal-cost load balancing.

Obviously, a two-step recursive lookup would be too expensive to perform for every switched IP packet. CEF switching thus pre-computes the equal-cost paths toward a BGP prefix and stores them in the Forwarding Information Base (FIB).

For example, the *Edge-2* router in Figure 2 announces a BGP prefix 10.2.1.0/24 to the *Edge-1* router. When you inspect the prefix in the IP routing table on *Edge-1*, you'll find a single route toward the destination (Listing 1). However there are two equal-cost OSPF paths toward the BGP next hop (Listing 2) resulting in load balancing entries for the BGP prefix in the CEF switching table (Listing 3).

Listing 1: IP routing table entry for BGP route 10.1.2.0/24

```
Edge-1#show ip route 10.1.2.0
Routing entry for 10.1.2.0/24
  Known via "bgp 65001", distance 200, metric 0
  Tag 65300, type internal
  Last update from 172.16.0.22 00:03:31 ago
  Routing Descriptor Blocks:
```

```
  * 172.16.0.22, from 172.16.0.22, 00:03:31 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 65300
```

Listing 2: Equal-cost routes to BGP next-hop 172.16.0.22

```
Edge-1#show ip route 172.16.0.22
Routing entry for 172.16.0.22/32
  Known via "ospf 1", distance 110, metric 101, type intra area
  Last update from 172.16.1.10 on Serial0/1/0, 00:00:00 ago
  Routing Descriptor Blocks:
    172.16.1.10, from 172.16.0.22, 00:00:00 ago, via Serial0/1/0
      Route metric is 101, traffic share count is 1
  * 172.16.1.2, from 172.16.0.22, 00:00:00 ago, via Serial0/0/0.100
      Route metric is 101, traffic share count is 1
```

Listing 3: CEF load balancing information for BGP route 10.1.2.0/24

```
Edge-1#show ip cef 10.1.2.0 internal
10.1.2.0/24, version 40, epoch 0, per-destination sharing
0 packets, 0 bytes
  via 172.16.0.22, 0 dependencies, recursive
    next hop 172.16.1.2, Serial0/0/0.100 via 172.16.0.22/32
    valid adjacency
  Recursive load sharing using 172.16.0.22/32
  Load distribution: 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1 (refcount 2)
  Hash  OK  Interface                 Address           Packets
  1     Y   Serial0/1/0               point2point             0
  2     Y   Serial0/0/0.100           point2point             0
  3     Y   Serial0/1/0               point2point             0
  4     Y   Serial0/0/0.100           point2point             0
… rest deleted …
```

External BGP Load Balancing

The IGP-based load balancing commonly used within an autonomous system is applicable to external BGP (EBGP) sessions as well and allows you to split the traffic load on parallel lines between a pair of EBGP-speaking routers (Figure 3).

Figure 3: Load balancing with multihop EBGP session

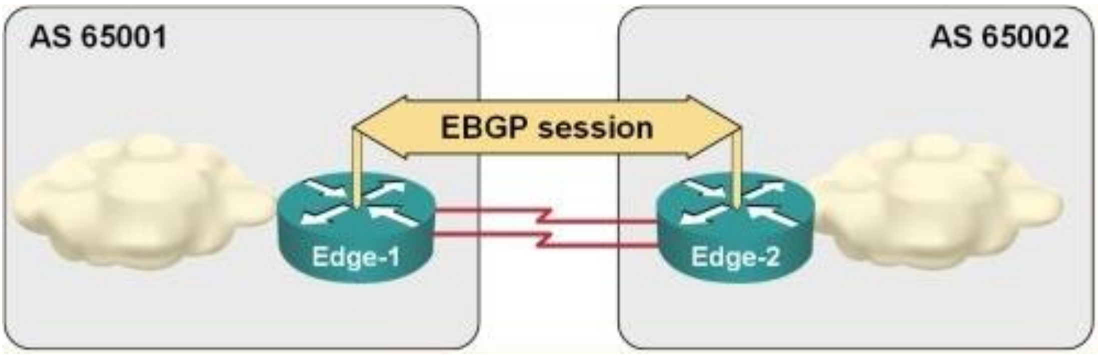

To enable this type of load balancing, you have to configure multihop EBGP session between loopback interfaces and use static host routes (or an IGP) to make the loopback address of the EBGP neighbor reachable. As there are multiple paths to the neighbor's loopback address (which is the BGP next hop), the load balancing works as expected. A detailed configuration is available on Cisco's web site.

Note: An alternate design, available only on serial links, is to use the same IP address on all parallel links.

The true BGP-only EBGP load balancing is configured with the **maximum-paths** router configuration command and allows a BGP router with multiple EBGP sessions to load balance the traffic toward destinations received in EBGP updates. However, as the router would only load balance across "equal-cost" EBGP routes, severe limitations apply:

- All EBGP sessions must be established with routers in the same autonomous system. Routes received from different autonomous systems are never considered to have equal cost.
- All the path selection attributes of the received routes (MED, AS Path) have to be the same. Obviously, the routes also have to have equal BGP local preference and weight (if you set those with inbound **route-map**).

Note: Although the router only considers the autonomous system (AS) path length when selecting the best BGP path, the actual value of the AS path attribute has to match for two paths to be considered equal-cost.

The only applicability of EBGP load balancing is thus the design where a single BGP router has two EBGP sessions into the same autonomous system (Figure 4).

Figure 4: EBGP load balancing

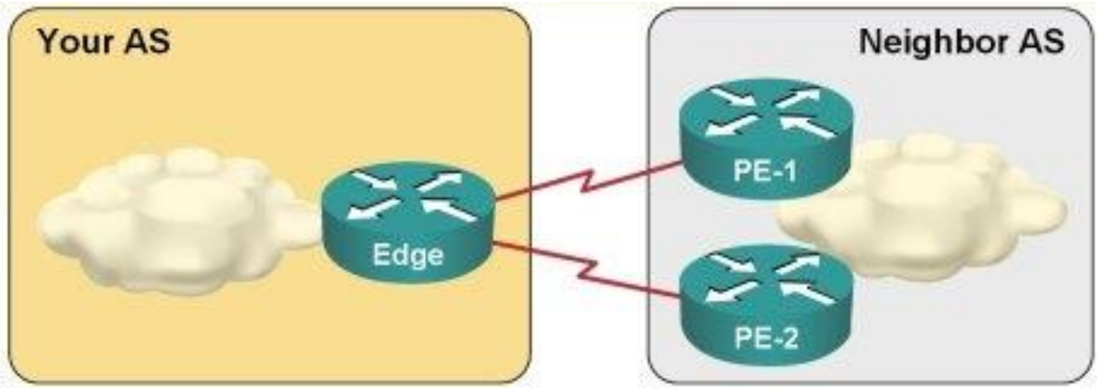

Even then, unless the EBGP sessions terminate on the same router (which would not make sense, you could use the scenario above), you would have to configure IBGP-based load balancing (described in the next section) in the adjacent autonomous system to achieve balanced traffic in reverse direction. In all other designs, the best you could hope for is load sharing, usually achieved with sophisticated design or clever configuration tricks (some examples are described on Cisco's web site).

EBGP Load Balancing Example

The sample network displayed in Figure 5 is a perfect match for the EBGP load balancing feature; the only configuration command to add to the BGP router configuration is the **maximum-paths 2** (Listing 4)

Figure 5: Sample EBGP load balancing scenario

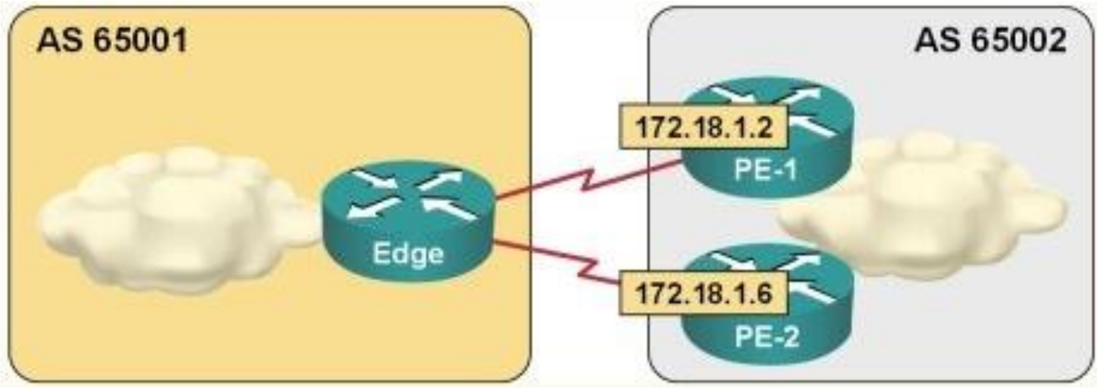

Listing 4: BGP configuration on the Edge router

```
router bgp 65001
 no synchronization
 bgp log-neighbor-changes
 neighbor 172.18.1.2 remote-as 65002
 neighbor 172.18.1.6 remote-as 65002
 maximum-paths 2
```

When an EBGP route is received from both routers in autonomous system (AS) 65002, the *Edge* router in AS 65001 would use both routes to load-balance traffic (Listing 5).

Listing 5: Load balancing across equal-cost EBGP routes

```
edge#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 12
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  65002 65003 65200 65300
    172.18.1.2 from 172.18.1.2 (172.16.0.21)
      Origin IGP, localpref 100, valid, external, multipath, best
  65002 65003 65200 65300
    172.18.1.6 from 172.18.1.6 (172.16.0.12)
      Origin IGP, localpref 100, valid, external, multipath
edge#show ip route 10.1.2.0
Routing entry for 10.1.2.0 255.255.255.0
  Known via "bgp 65001", distance 20, metric 0
  Tag 65002, type external
  Last update from 172.18.1.6 00:00:18 ago
```

```
  Routing Descriptor Blocks:
    172.18.1.6, from 172.18.1.6, 00:00:18 ago
      Route metric is 0, traffic share count is 1
      AS Hops 4
      Route tag 65002
  * 172.18.1.2, from 172.18.1.2, 00:00:18 ago
      Route metric is 0, traffic share count is 1
      AS Hops 4
      Route tag 65002
```

Note: Even though the two BGP routes have "equal cost" and are thus used in packet forwarding, only one of them is selected as the best route and advertised to other BGP peers.

However, if the content of the AS paths of the received BGP routes does not match (even though they have equal length), the load balancing does not work; only the route selected as the BGP best path is inserted into the IP routing table (Listing 6).

Listing 6: Routes with different AS paths are not equal-cost routes

```
edge#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 13
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  65002 65003 65200 65300
    172.18.1.2 from 172.18.1.2 (172.16.0.21)
      Origin IGP, localpref 100, valid, external, best
  65002 65003 65201 65300
    172.18.1.6 from 172.18.1.6 (172.16.0.12)
      Origin IGP, localpref 100, valid, external
edge#show ip route 10.1.2.0
Routing entry for 10.1.2.0 255.255.255.0
  Known via "bgp 65001", distance 20, metric 0
  Tag 65002, type external
  Last update from 172.18.1.2 00:00:08 ago
  Routing Descriptor Blocks:
  * 172.18.1.2, from 172.18.1.2, 00:00:08 ago
```

```
Route metric is 0, traffic share count is 1
AS Hops 4
Route tag 65002
```

IBGP Load Balancing

IBGP load balancing, introduced in IOS release 12.2T and 12.3, addresses the next common design scenario: for redundancy purposes, our network has two egress routers toward a neighboring AS. Although it's clear that these two routers cannot load-balance traffic toward the neighboring AS, at least the other routers in our AS should be able to do so (from *POP* into *neighbor AS* in Figure 6).

Figure 6: IBGP Load Balancing

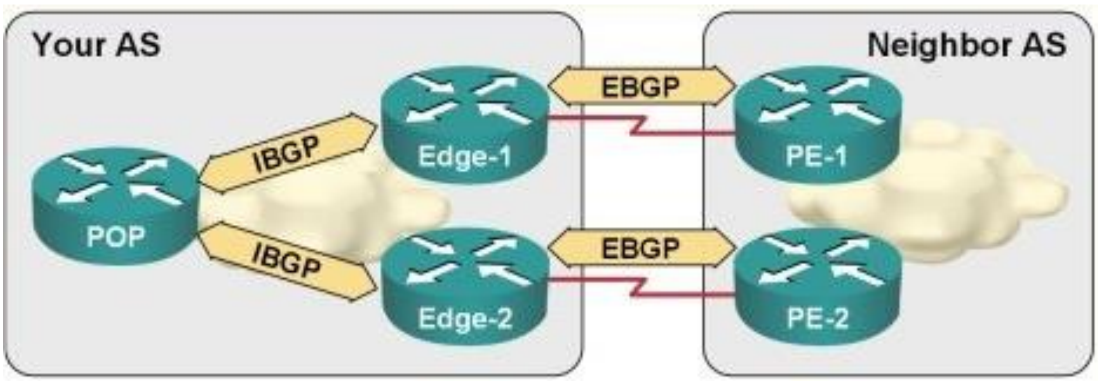

The IBGP load balancing (configured with **maximum-paths ibgp** ***number*** router configuration command) is even more restrictive than the EBGP load balancing. You need a perfect match of BGP attributes (MED, Local preference, complete AS path) and weights, but also equal-cost IGP paths toward the BGP next hops.

The equal-cost requirement might present tough design problems. It's obvious that the IGP cost toward the edge routers has to be the same to enable load balancing, otherwise there's potential for routing loops. Quite often, you'd like to achieve IBGP load balancing even when the interfaces between the neighboring autonomous systems don't have equal bandwidth. In such a case, the paths toward the neighbor AS would not have equal cost if you simply include these interfaces into your IGP (as passive interfaces for security reasons), but there are a few design options to achieve the equal-cost IGP path toward BGP next hop even in these cases.

Option #1: Use next-hop-self

As discussed above, if you have equal-cost path to the edge routers, but the inter-AS interfaces do not have equal bandwidth, the total cost to the BGP next-hop (which is usually the IP address of the edge router in the neighboring AS) will not be the same if you include the inter-AS links as passive interfaces in your IGP (Figure 7).

Figure 7: Inter-AS links are passive interfaces in OSPF

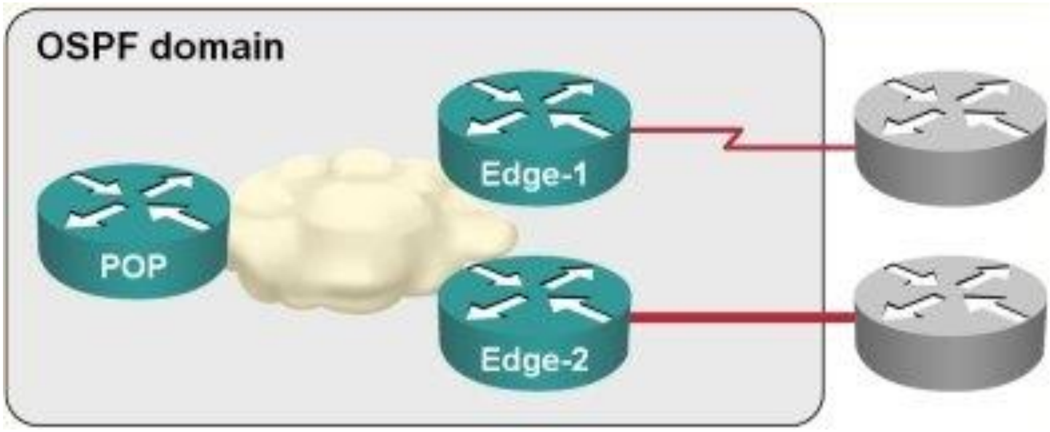

In this case, you could set the BGP next hop in internal BGP updates to be the loopback interface of the edge BGP router with the **neighbor** ***ip-address*** **next-hop-self** router configuration command, to ensure that the cost to all BGP next-hops is the same.

Option #2: External OSPF routes

If you use OSPF as the internal routing protocol in scenario in Figure 7, you could benefit from the fact that the internal OSPF cost is not added to cost of the external type-2 OSPF routes. Redistributing inter-AS links as external type-2 OSPF routes with fixed cost will thus ensure that the cost toward the BGP next-hops will be the same, resulting in IBGP load balancing (Figure 8). Obviously, the intra-AS OSPF cost to the AS boundary router (ASBR; *Edge-1* and *Edge-2* in Figure 8) still has to be equal, as the intra-AS cost is used as a tie-breaker between routes with the same external cost.

Note: You could even fix unequal-cost IGP paths by using external type-1 routes with well considered costs (which you can specify manually in the ***redistribute*** *command).*

Figure 8: Inter-AS subnets are redistributed as type-2 external routes into OSPF

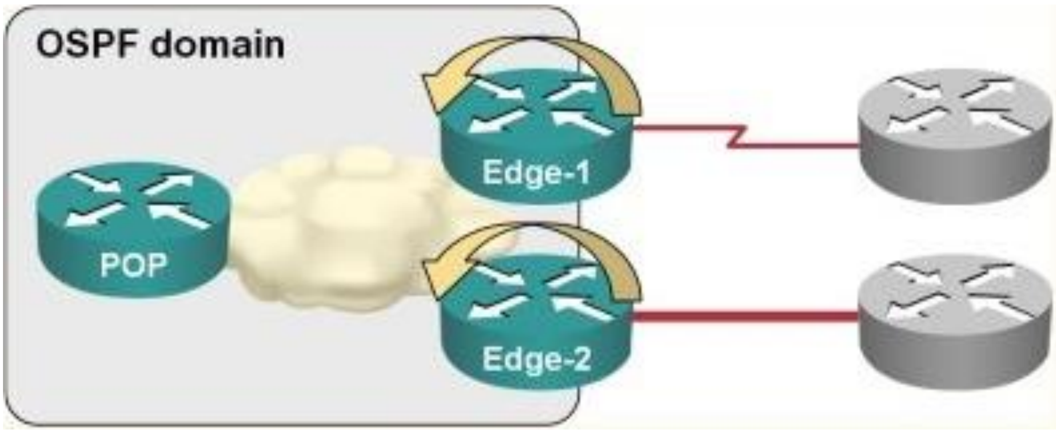

Route Reflectors and IBGP Load Balancing

Before going into sample IBGP load sharing configurations (and associated design challenges), you should be aware of the impact of BGP route reflectors on IBGP load sharing. As discussed in the EBGP load balancing section, each BGP router selects a single best route (even though it might use more than one for packet forwarding) and propagates it to its BGP neighbors. A route reflector might thus use multiple IBGP routes and load balance between them, but its clients would only receive a single route (Figure 9). The impact of this behavior can only be alleviated with a careful placement of route reflectors or a sophisticated BGP design (contact our professional services team if you need help in designing your BGP network).

Figure 9: Route reflector kills load balancing

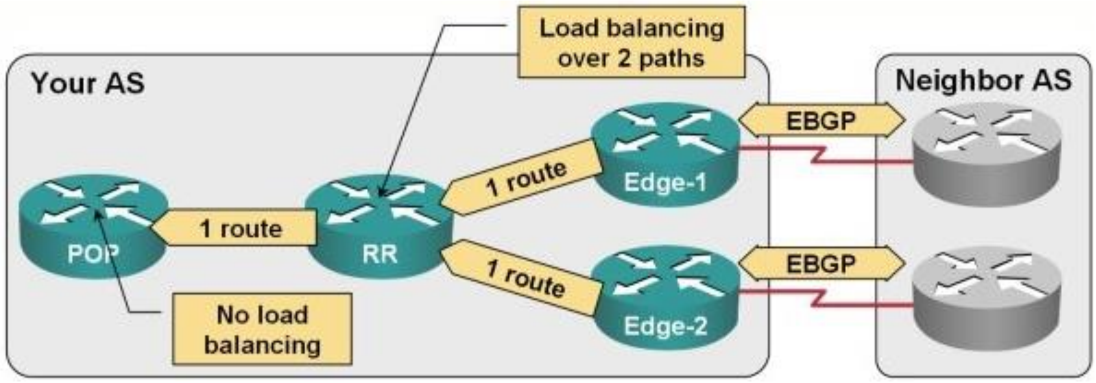

IBGP Load Balancing Example

The IBGP load balancing concepts will be illustrated in the sample three-router network shown in Figure 10. Although the POP router receives two paths for IP prefix 10.1.2.0/24 (one from Edge-1, the other from Edge-2, as shown in Listing 7), the IGP cost toward the BGP next hops is not equal (113 versus 51), so there can be no IBGP load balancing and only one path is inserted in the IP routing table (Listing 8).

Figure 10: Unequal-cost BGP network

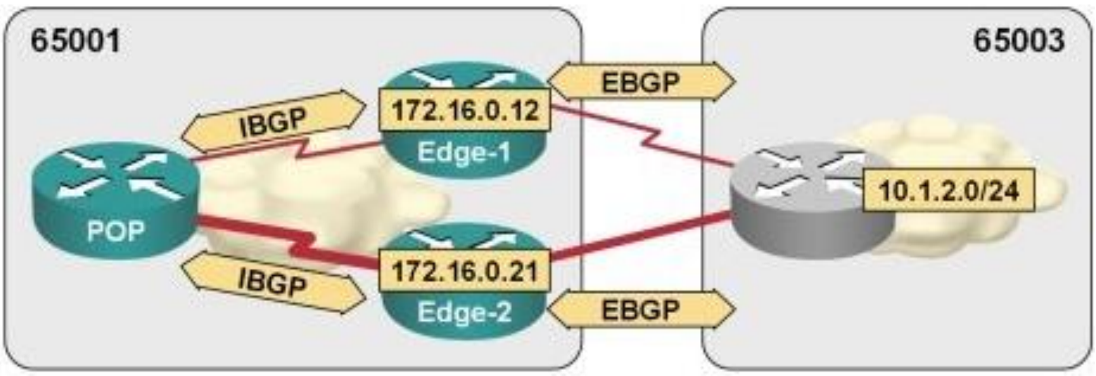

Listing 7: POP router is receiving two IBGP paths toward 10.1.2.0/24

```
POP#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 2
Paths: (2 available, best #2, table Default-IP-Routing-Table)
Multipath: iBGP
Flag: 0x820
  Not advertised to any peer
  65003 65200 65300
    172.16.1.6 (metric 113) from 172.16.0.12 (172.16.0.12)
```

```
      Origin IGP, metric 0, localpref 100, valid, internal
  65003 65200 65300
    192.168.0.6 (metric 51) from 172.16.0.21 (172.16.0.21)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Listing 8: A single BGP route is inserted into IP routing table

```
POP#show ip route 10.1.2.0
Routing entry for 10.1.2.0 255.255.255.0
  Known via "bgp 65001", distance 200, metric 0
  Tag 65003, type internal
  Last update from 192.168.0.6 00:02:39 ago
  Routing Descriptor Blocks:
  * 192.168.0.6, from 172.16.0.21, 00:02:39 ago
      Route metric is 0, traffic share count is 1
      AS Hops 3
      Route tag 65003
```

When the links to the edge of the autonomous system are made equal (their OSPF cost is the same), the two BGP routes have the same cost toward the IGP next-hop (Listing 9) and are both installed in the IP routing table (Listing 10).

Listing 9: BGP routes have equal cost to the BGP next-hop

```
a1#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 5
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Multipath: eBGP iBGP
Flag: 0x900
  Not advertised to any peer
  65003 65200 65300
    172.16.1.6 (metric 100) from 172.16.0.12 (172.16.0.12)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
  65003 65200 65300
    192.168.0.6 (metric 100) from 172.16.0.21 (172.16.0.21)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
```

Listing 10: Both BGP routes are installed into the IP routing table

```
a1#show ip route 10.1.2.0
Routing entry for 10.1.2.0 255.255.255.0
  Known via "bgp 65001", distance 200, metric 0
  Tag 65003, type internal
  Last update from 192.168.0.6 00:00:32 ago
  Routing Descriptor Blocks:
  * 192.168.0.6, from 172.16.0.21, 00:00:32 ago
      Route metric is 0, traffic share count is 1
      AS Hops 3
      Route tag 65003
    172.16.1.6, from 172.16.0.12, 00:00:32 ago
      Route metric is 0, traffic share count is 1
      AS Hops 3
      Route tag 65003
```

BGP Link Bandwidth

In the sample network we have been discussing in the previous sections, the inter-AS links between AS 65001 and 65002 have different speeds (Figure 10), but the *POP* router load balances the traffic in 50:50 ratio between the two *Edge* routers. To achieve unequal-bandwidth load balancing (with traffic sharing proportional to the inter-AS link bandwidth), a new *Link Bandwidth* extended community was defined to influence the load balancing ratio and implemented in IOS releases 12.2T and 12.3.

Note: This extended community is defined in the Section 7 of the Extended Communities Internet draft, but obviously never made it to the final RFC.

To use the BGP Link Bandwidth feature, you have to:

- Configure propagation of extended BGP communities on IBGP sessions with the **neighbor send-community both** router configuration command.
- Configure the attachment of the new extended community to incoming EBGP routes with the **neighbor dmzlink-bw** router

configuration command. The value attached to the received EBGP routes is the bandwidth configured on the interface; you cannot change the *link bandwidth* community with a route-map.

Note: The link bandwidth extended community is attached to incoming BGP updates. To enforce the proper value of this community after the ***neighbor dmzlink-bw*** *command has been entered or after the interface* ***bandwidth*** *has been changed, use the* ***clear ip bgp neighbor soft in*** *command.*

- Configure the link bandwidth-based load sharing on all BGP routers in your AS with the **bgp dmzlink-bw** router configuration command.

In the network in Figure 11 (the example continues from the previous section), you would have to configure community propagation on all IBGP sessions on edge routers as well as marking of incoming routes with *link bandwidth* external community on EBGP sessions. The BGP configuration on the *Edge-2* router is shown in Listing 11.

Figure 11: Unequal bandwidth IBGP load balancing

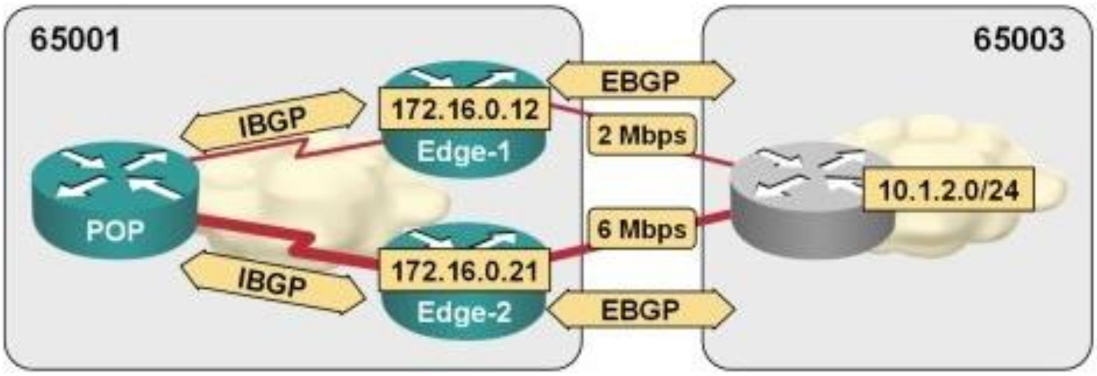

Listing 11: BGP configuration on Edge-2

```
router bgp 65001
 no synchronization
 bgp log-neighbor-changes
 neighbor 172.16.0.11 remote-as 65001
 neighbor 172.16.0.11 update-source Loopback0
 neighbor 172.16.0.11 send-community both
 neighbor 172.16.0.12 remote-as 65001
 neighbor 172.16.0.12 update-source Loopback0
 neighbor 172.16.0.12 send-community both
```



```
neighbor 192.168.0.6 remote-as 65003
neighbor 192.168.0.6 dmzlink-bw
```

When the routes are refreshed across the BGP session between *Edge-2* and the external peer (by clearing the BGP session with the **clear ip bgp** command or requesting route refresh with the **clear ip bgp *neighbor* soft in** command), the external BGP paths are marked with the inter-AS link bandwidth (as configured on the *Edge-2*'s interface with the **bandwidth** command) and propagated to IBGP neighbors in AS 65001.

Note: The extended BGP communities draft specifies that the link bandwidth should not be propagated outside of the AS (similar to the local preference BGP attribute).

After the same configuration changes are applied to *Edge-1*, you can inspect the IBGP routes on the *POP* router and verify that the *link bandwidth* has been attached to all routes received from AS 65003 (Listing 12).

Listing 12: BGP routes to IP prefix 10.1.2.0/24 on the POP router

```
POP#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 12
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Multipath: eBGP iBGP
Flag: 0x800
  Not advertised to any peer
  65003 65200 65300
    172.16.1.6 (metric 20) from 172.16.0.12 (172.16.0.12)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      DMZ-Link Bw 250 kbytes
  65003 65200 65300
    192.168.0.6 (metric 20) from 172.16.0.21 (172.16.0.21)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      DMZ-Link Bw 750 kbytes
```

When the BGP routes are copied into the IP routing table on the *POP* router, the traffic share count reflects the ratio of the DMZ bandwidths (Listing 13).

Listing 13: Unequal-bandwidth load balancing on the POP router

```
POP#show ip route 10.1.2.0
Routing entry for 10.1.2.0 255.255.255.0
  Known via "bgp 65001", distance 200, metric 0
  Tag 65003, type internal
  Last update from 192.168.0.6 00:00:08 ago
  Routing Descriptor Blocks:
    192.168.0.6, from 172.16.0.21, 00:00:08 ago
      Route metric is 0, traffic share count is 3
      AS Hops 3
      Route tag 65003
  * 172.16.1.6, from 172.16.0.12, 00:00:08 ago
      Route metric is 0, traffic share count is 1
      AS Hops 3
      Route tag 65003
```

Summary

In well designed networks, you can achieve a surprising amount of load balancing even when you use BGP as one of your core routing protocols.

It's always possible to load-balance between a pair of edge routers; if you don't have two equal-cost paths between them across your network, you can engineer them with MPLS Traffic Engineering. Similarly, you can always load-balance traffic across parallel links between two edge routers in adjacent autonomous systems.

Furthermore, you can configure EBGP load balancing on edge routers that have multiple links connected to neighbors in the same adjacent autonomous system, as long as routes received from all neighbors have the same AS path and MED (and they have the same local preference and weight if you've set these attributes with an inbound **route-map**). Load balancing between EBGP and IBGP routes is not possible, as it might lead to forwarding loops.

You can also configure IBGP load-balancing, allowing you to split traffic between a number of edge routers (up to 16 with IOS release 12.3T or 12.4); as long as all the BGP paths received from them have identical path-selection attributes and the IGP cost to the BGP next-hop is the same. With the support of *link bandwidth* extended community introduced in IOS release 12.3, you can split the traffic proportionally to the interface bandwidth of the inter-AS links.

NIL – More Than Just a Training Company

NIL Learning delivers the leading-edge Cisco training to IT professionals and companies around the globe. Through field-proven experts — each both active engineer and instructor — NIL Learning enhances the standard learning curriculum with real-life experience and helps clients to maximize their training investment.

NIL Learning is part of NIL, a leading global IT solutions provider. Since 1992, NIL has been at the forefront of advanced contributors to strategic partner Cisco's technologies, learning curriculum and value-added solutions deployed to clients around the globe. Today, NIL has earned the highest certifications offered by Cisco, VMware, EMC, HP, IBM, Microsoft, F5, Jive, MobileIron, RSA, VCE and others. Their portfolio of solutions consists of managed services, professional services and learning services.

NIL is headquartered in Slovenia, with regional offices in Croatia, Serbia, Saudi Arabia, the U.S., Turkey, South Africa, Morocco, Nigeria, Kenya and Botswana.

Why learn at NIL LEARNING?

- All NIL LEARNING instructors are **field-proven experts** - each both active engineer and instructor.
- **75% of NIL LEARNING engineers hold CCSI** certifica- tions, and **18 have already achieved the respected CCIE rank**.
- NIL LEARNING **enhances the standard learning curriculum** with real-life experience and helps clients to maximize their training investment.

- NIL has been a Cisco Training Partner for many years; it became a **Cisco Learning Partner in 1993**, and has been a **Cisco Gold Partner since 1995**.
- NIL was awarded the **Cisco Most Business Relevant Learning Partner in MEA in 2010 and the most innova- tive learning partner in MEA**.
- NIL received the **Innovation Award for its Technology Led Training and its extensive contribution to Cisco learning solutions** at the Cisco EMEAR Learning Partner Summit in 2012.
- NIL received the **Innovation Award for its Technology Led Training and Advanced Engineer Program** at the Cisco Global Learning Partner Summit in 2013.
- NIL LEARNING runs a **centralized training schedule across the whole EMEAR region**.

More Info

Slovenia

T: +386 1 4746 500

E: sales-support@nil.com

Botswana

T: +267 318 1684

E: training@it-iq.bw

Croatia

T: +385 (0)51 583 255

E: info-nilcroatia@nil.com

Kenya

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Morocco

T: +212(0) 660 808 394

E: info-nilmorocco@nil.com

Nigeria

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Saudi Arabia

T: +966 1 465 4641

E: info.nilme@nil.com

Serbia

T: +381 11 2282 818

E: info-nilserbia@nil.co.rs

South Africa

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Turkey

T: +902 123 81 8639

E: info-nilturkey@nil.com

USA

T: +1 612 886 3900

E: info-nilusa@nil.com

www.learning.nil.com