

# BGP AS Migration Process

Moving from one ASN to another

## Introduction

BGP Autonomous System (AS) migration is a process of migration from one AS to another AS. It is typically a scenario when a company, which we will refer to as company B, which has been assigned its AS number (ASN) by the Regional Internet Registry (RIR), is purchased by another company, which we will refer to as company A and this company has a different ASN. As a result of this, company B needs to reconfigure the ASN on its BGP enabled routers to the ASN assigned to company A. Sounds quite simple, however, in real life, there are certain challenges associated with AS migration. We are going to discuss them in this eBook.

First of all, imagine that companies A and B are Internet Service Providers (ISPs). As we have mentioned, company B is going to migrate its ASN to the ASN assigned to company A. After a change of the ASN on ISP B routers, all the customers of ISP B are forced to reconfigure their Customer Edge (CE) routers. Their CE routers have established external BGP (eBGP) peer session with the Provider Edge (PE) routers of the ISP B. Basically they have to change the ASN in a neighbor command, under BGP configuration to reflect a new ASN assigned to the ISP B. It might be a tricky process since the ISP B can have a lot of customers. Moreover, it would be hard to coordinate reconfiguration on a multitude of CE routers at the same time in order to shorten network outages. For this reason, common AS migration mechanisms must be implemented. The mechanisms require no interaction with custom-

ers during and after ASN migration.

In fact, BGP AS migration mechanisms are not included in RFC 4271 - Border Gateway Protocol 4 (BGP-4). Instead, a separate RFC 7705 is created to describe them. The mechanisms allow ISPs to migrate the old (legacy) ASN to a new ASN while BGP configuration of CE routers is kept untouched. These additional features of the BGP protocol ensure that a change of the ASN on a PE routers is not visible for CE routers. Thanks to this, customer's routers can keep the legacy ASN for their BGP peer - a PE router for the period of the migration process. However, during a next maintenance window, customers can freely reconfigure the ASN value under BGP configuration on their CE router to reflect a new ASN assigned to their ISP. In other words, a change of the ASN on a provider side does not impact the customer's side. Customers are not forced to cooperate during an AS migration.

## Migration Mechanisms

The paragraph above introduces a key factor to a successful AS migration process for service providers - independent from customer intervention. It is achieved by configuring **local-as** keyword followed by the old legacy ASN on PE router of an ISP. Let's say that an ISP B has assigned ASN 64510 and it is going to migrate AS from a number 64510 to 64500. The IP address of customer's

## AS Migration Process

---

CE BGP peer is 172.19.1.1, with an assigned ASN 64496. In this case, we can probably find the following statements on PE router configuration.

```
router bgp 64510
neighbor 172.19.1.1 remote-as 64496
```

During migration, a network administrator makes changes to the configuration as described below

```
no router bgp 64510
router bgp 64500
neighbor 172.19.1.1 remote-as 64496
neighbor 172.19.1.1 local-as 64510
```

A globally configured ASN is now 64500. A locally configured ASN is 64510. The entry "local-as 64510" enables a service provider's PE router to use the legacy ASN 64510 in the My AS field inside a BGP open message instead of the globally configured ASN 64500. The PE and CE routers can establish an eBGP session with each other, enabling customers' router to believe that a PE router still uses the legacy ASN 64510. This is what makes the AS migration on a provider's side invisible to its customers.

The second requirement of a successful AS number migration involves manipulation of the AS\_PATH attribute on a provider's PE router. As we know, BGP path selection algorithm selects a route with the shortest path AS\_PATH to a remote prefix.

For this reason, a route with less AS hops is preferred to the route with more AS hops if parameters such as weight or local preference do not break a tie. However, a one significant issue is connected with configuration of local-as statement. If a local-as entry is configured without **no-prepend** option on the ISP PE router for CE neighbor, a locally configured (legacy) ASN is automatically prepended to AS\_PATH to advertised or installed BGP routes received from customer CE router. The ISP's router announces these routes to the iBGP peers within the same AS. The peers announce the routes to their eBGP neighbors that can notice an increase of AS-PATH length. As a result, BGP enabled routers might select a different route to a customer network which leads to unexpected changes in traffic patterns. Below is the same line from previous configuration on a PE router of an ISP with added no-prepend entry.

```
neighbor 172.19.1.1 local-as 64510 no-prepend
```

In order to prevent the ISP's router prepending a locally configured ASN in the routing updates sent to its iBGP neighbors, thus artificially increasing the AS\_PATH length, the no-prepend option must be added to local-as configuration on the ISP's PE router. However, we need to configure another BGP AS migration mechanism to the ISP's PE router configuration to prevent increasing the AS\_PATH length in the opposite direction. The keyword replace-old prevents adding globally configured ASN 64500 to AS\_PATH attribute inside the BGP update messages sent by the ISP's PE router towards the CE router. Below is a particular line in the BGP configuration of the PE router of an ISP.

## AS Migration Process

```
neighbor 172.19.1.1 local-as 64510 no-prepend replace-as
```

The last AS migration mechanism, which we did not mention yet, is a dual-as option. We configure this option under a neighbor statement along with local-as, no-prepend and replaces-as options on the PE router of an ISP. The dual-as option allows the CE router to establish an eBGP peer session with the PE router of said ISP, either using the legacy ASN 64510 or the globally configured ASN 64500.

```
neighbor 172.19.1.1 local-as 64510 no-prepend replace-as dual-as
```

### Configuration of BGP AS Migration Mechanisms

In this part, we are going to setup a simple network lab for the purpose of showing the configuration of BGP AS migration mechanisms. We will also explain their impact on AS\_PATH length attribute and show how to fix existing issues.



**NOTE:** The lab is not intended to be a step-by-step guide for an AS migration. It serves the purpose of demonstrating the impact of BGP AS migration mechanisms on AS\_PATH length attribute.

A network topology consists of four routers with configured BGP routing protocol. The ISPs A and B provide connection to the Internet for their customers - C and D. The Customer Edge (CE) router CE-D belongs to a network of the customer D and it has an established eBGP session with a Provider Edge (PE) router of the ISP A. The CE-C router is a part of a network of the customer C and it has an established eBGP session with Provider Edge (PE) router of ISP B. The routers of ISP A and B are eBGP peers.

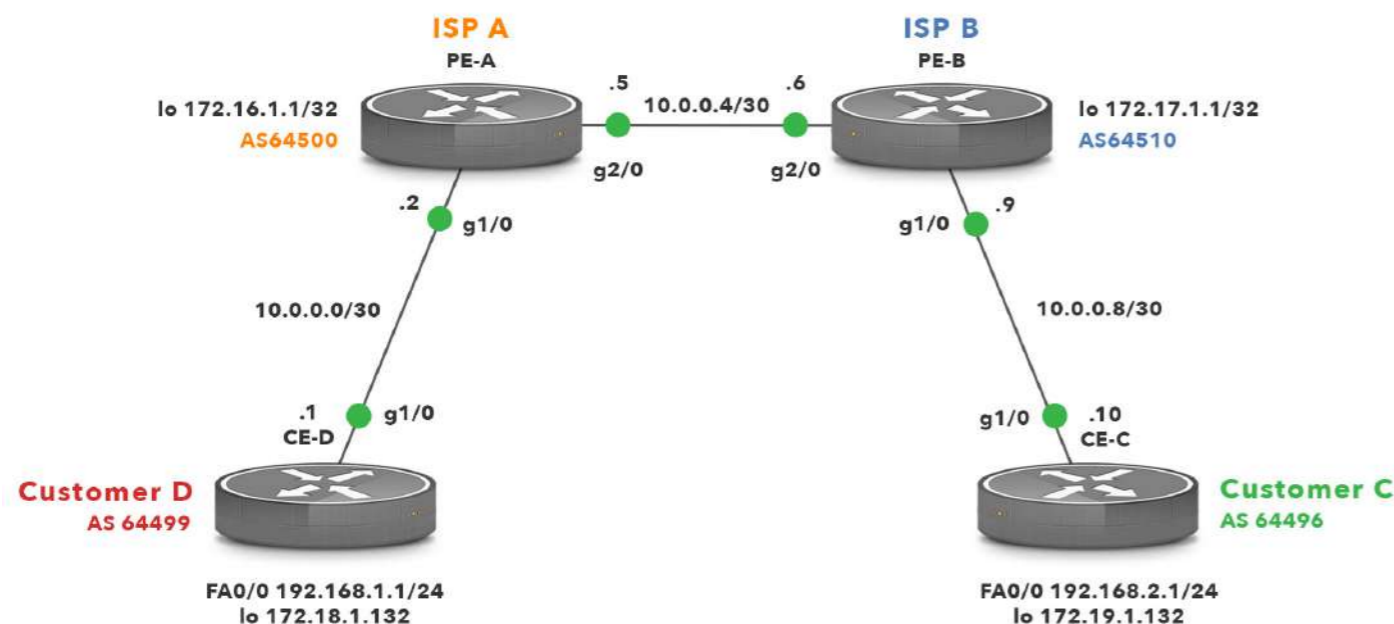
Let's say that the ISP B is acquired by ISP A, thus we must reconfigure a router PE-B to change its globally configured ASN from 64510 to 64500. First, we will check the BGP configuration on all routers before changing anything.



**NOTE:** Routers have configured AS numbers that are reserved for use in documentation. Also, they use IP addresses from private IP ranges. In real life, public IP addresses would be configured on routers.



# BGP AS Migration Process



Picture 1- Network Topology BEFORE Migration of PE-B Router to New ASN

## BGP Configuration on CE-D

```
CE-D# show running-config | section bgp
```

```
router bgp 64499
  bgp log-neighbor-changes
  network 10.0.0.0 mask 255.255.255.252
  network 172.18.1.1 mask 255.255.255.255
  network 192.168.1.0
  neighbor 172.16.1.1 remote-as 64500
  neighbor 172.16.1.1 disable-connected-check
  neighbor 172.16.1.1 update-source Loopback0
```

## BGP Configuration on PE-A

```
PE-A# show running-config | section bgp
```

```
router bgp 64500
  bgp log-neighbor-changes
  network 10.0.0.0 mask 255.255.255.252
  network 10.0.0.4 mask 255.255.255.252
  network 172.16.1.1 mask 255.255.255.255
  neighbor 172.17.1.1 remote-as 64510
  neighbor 172.17.1.1 disable-connected-check
  neighbor 172.17.1.1 update-source Loopback0
  neighbor 172.18.1.1 remote-as 64499
  neighbor 172.18.1.1 disable-connected-check
  neighbor 172.18.1.1 update-source Loopback0
```

## BGP Configuration on PE-B

```
PE-B# show running-config | section bgp
```

```
router bgp 64510
  bgp log-neighbor-changes
  network 10.0.0.4 mask 255.255.255.252
  network 10.0.0.8 mask 255.255.255.252
  network 172.17.1.1 mask 255.255.255.255
  neighbor 172.16.1.1 remote-as 64500
  neighbor 172.16.1.1 disable-connected-check
  neighbor 172.16.1.1 update-source Loopback0
  neighbor 172.19.1.1 remote-as 64496
  neighbor 172.19.1.1 disable-connected-check
  neighbor 172.19.1.1 update-source Loopback0
```

# BGP AS Migration Process

## BGP Configuration on CE-C

```

router bgp 64496
  bgp log-neighbor-changes
  network 10.0.0.8 mask 255.255.255.252
  network 172.19.1.1 mask 255.255.255.255
  network 192.168.2.0
  neighbor 172.17.1.1 remote-as 64510
  neighbor 172.17.1.1 disable-connected-check
  neighbor 172.17.1.1 update-source Loopback0
    
```

To check that our routers are properly configured, we will inspect the output of the BGP table on router CE-D. The route to prefix 192.168.2.0/24 is over AS 64500, 64510 and 64496.

```

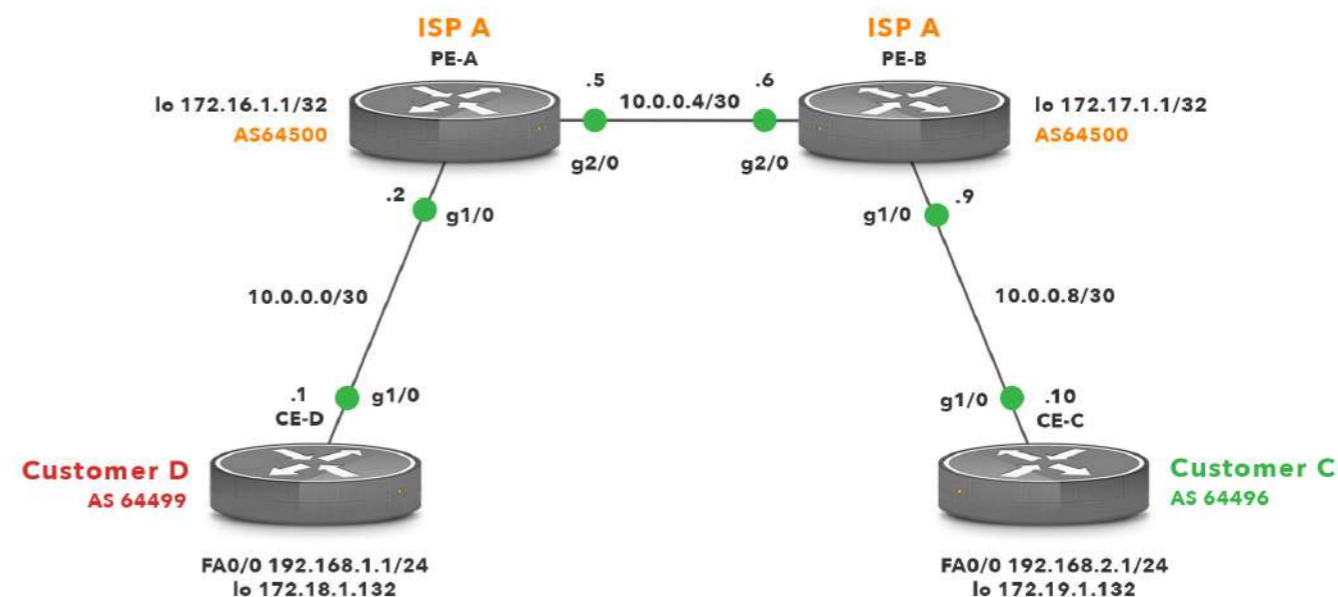
CE-D#show ip bgp | begin Network
Network      Next Hop    Metric LocPrf  Weight  Path
* 10.0.0.0/30 172.16.1.1  0       0       0       64500 i
*>           0.0.0.0    0       0       32768   i
*> 10.0.0.4/30 172.16.1.1  0       0       0       64500 i
*> 10.0.0.8/30 172.16.1.1  0       0       0       64500 64510 i
r> 172.16.1.1/32 172.16.1.1  0       0       0       64500 i
*> 172.17.1.1/32 172.16.1.1  0       0       0       64500 64510 i
*> 172.18.1.1/32 0.0.0.0     0       0       32768   i
*> 172.19.1.1/32 172.16.1.1  0       0       0       64500 64510 64496 i
*> 192.168.1.0 0.0.0.0     0       0       32768   i
*> 192.168.2.0 172.16.1.1  0       0       0       64500 64510 64496 i
    
```

We will do the same for router CE-C. The route to prefix 192.168.1.0/24 is over AS 64510, 64500 and 64499.

```

CE-D#show ip bgp | begin Network
Network      Next Hop    Metric LocPrf  Weight  Path
*> 10.0.0.0/30 172.17.1.1  0       0       0       64510 64500 i
*> 10.0.0.4/30 172.17.1.1  0       0       0       64510 i
* 10.0.0.8/30 172.17.1.1  0       0       0       64510 i
*>           0.0.0.0    0       0       32768   i
*> 172.16.1.1/32 172.17.1.1  0       0       0       64510 64500 i
r> 172.17.1.1/32 172.17.1.1  0       0       0       64510 i
*> 172.18.1.1/32 172.17.1.1  0       0       0       64510 64500 64499 i
*> 172.19.1.1/32 0.0.0.0     0       0       32768   i
*> 192.168.1.0 172.17.1.1  0       0       0       64510 64500 64499 i
*> 192.168.2.0 0.0.0.0     0       0       32768   i
    
```

Now let's start with reconfiguration of the router PE-B. According to the scheme below, we will change a globally configured ASN 64510 to 64500.



Picture 2 - Network Topology AFTER Migration of PE-B Router to New AS

## BGP AS Migration Process

---

Below are the configuration steps that we will copy to router PE-B of ISP A.

### BGP Configuration on PE-B

```
no router bgp 64510
router bgp 64500
  network 10.0.0.4 mask 255.255.255.252
  network 10.0.0.8 mask 255.255.255.252
  network 172.17.1.1 mask 255.255.255.255
  neighbor 172.16.1.1 remote-as 64500
  neighbor 172.16.1.1 update-source Loopback0
  neighbor 172.19.1.1 remote-as 64496
  neighbor 172.19.1.1 disable-connected-check
  neighbor 172.19.1.1 update-source Loopback0
  neighbor 172.19.1.1 local-as 64510
```



**NOTE:** To tell iBGP neighbor PE-A to use a next-hop IP address 172.17.1.1 (PE-B) instead of the IP 172.19.1.1 (CE-C) for a route 192.168.2.0/24 (customer C), we need to add the following line into the BGP configuration of PE-B.

```
neighbor 172.16.1.1 next-hop-self
```

Similarly, here are the configuration steps for BGP configuration on router PE-A of ISP A.

### BGP Configuration on PE-A

```
router bgp 64500
  neighbor 172.17.1.1 remote-as 64500
  neighbor 172.17.1.1 next-hop-self
```

Likewise, we have configured a router PE-A as a BGP next-hop-self (IP address 172.16.1.1) for router PE-B instead of 172.18.1.1.

### Checking BGP Table of CE-D and CE-C Routers

At this point we should be able to ping the router CE-C with the IP address 192.168.2.1 from the router CE-D and vice versa.

```
E-D# ping 192.168.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2
seconds:
!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
76/78/84 ms
```

## BGP AS Migration Process

---

However let's check the BGP table for route 192.168.2.0 on router CE-D. Notice that AS\_PATH length is 64500, 54510 and 64496. It is because the router PE-B has prepended a legacy ASN 64510 to the AS\_PATH attribute in the updates announced to its iBGP peer, the router PE-A.

```
CE-D# show ip bgp | include 192.168.2.0
*> 192.168.2.0      172.16.1.1    0 64500 64510 64496 i
```

Also a check of the BGP table on router CE-C reveals that router PE-B has added a globally configured ASN 64500 to an outgoing routing update for CE-C.

```
CE-C# show ip bgp | include 192.168.1.0
*> 192.168.1.0      172.17.1.1      0 64510 64500 64499 i
```

To fix these issues we will modify a statement for a neighbor 172.19.1.1 in the BGP configuration on router PE-B as follows.

```
neighbor 172.19.1.1 local-as 64510 no-prepend replace-as  
dual-as
```



**NOTE:** The dual-as statement allows customer C to use either ASN 64510 or 64500 for BGP peer 172.17.1.1.

Let's inspect the BGP table for route 192.168.2.0/24 on CE-D router again. The AS\_PATH length is now reduced to ASN 64500 and 64496.

```
CE-D# show ip bgp | include 192.168.2.0
*> 192.168.2.0      172.16.1.1      0 64500 64496 i
```

Similarly, we will check the BGP table for route 192.168.1.0/24 on the router CE-C. The AS\_PATH length is reduced to ASN 64510 and 6449.

```
CE-C# show ip bgp | include 192.168.1.0
*> 192.168.1.0      172.17.1.1      0 64510 64499 i
```

Thanks to the configuration of BGP AS migration mechanisms in our network topology, we have decreased the length of AS\_PATH attribute in both directions. Below is a complete BGP configuration on router PE-A of ISP A.



## BGP AS Migration Process

---

```
PE-A# show running-config | section bgp
router bgp 64500
  bgp log-neighbor-changes
  network 10.0.0.0 mask 255.255.255.252
  network 10.0.0.4 mask 255.255.255.252
  network 172.16.1.1 mask 255.255.255.255
  neighbor 172.17.1.1 remote-as 64500
  neighbor 172.17.1.1 update-source Loopback0
  neighbor 172.17.1.1 next-hop-self
  neighbor 172.18.1.1 remote-as 64499
  neighbor 172.18.1.1 disable-connected-check
  neighbor 172.18.1.1 update-source Loopback0
```

BGP configuration on a router PE-B of ISP A follows.

```
PE-B# show running-config | section bgp
router bgp 64500
  bgp log-neighbor-changes
  network 10.0.0.4 mask 255.255.255.252
  network 10.0.0.8 mask 255.255.255.252
  network 172.17.1.1 mask 255.255.255.255
  neighbor 172.16.1.1 remote-as 64500
  neighbor 172.16.1.1 update-source Loopback0
  neighbor 172.16.1.1 next-hop-self
  neighbor 172.19.1.1 remote-as 64496
  neighbor 172.19.1.1 local-as 64510 no-prepend replace-as
  dual-as
  neighbor 172.19.1.1 disable-connected-check
  neighbor 172.19.1.1 update-source Loopback0
```

## ASN Retirement and Recycling

Now we ask the question of what is going to happen to the old ASN 64510? The no longer used AS Numbers must be returned to the issuing authority, e.g. ARIN. Therefore, it is very likely that the ASN 64510 will be voluntarily returned to the unallocated pool of ASNs, held there for a period of 60 days prior to a subsequent reallocation.

## Conclusion

There are common BGP AS migration mechanisms created and widely used to help with AS migration. The process of migrating ASNs described above allows multihomed network operators to move clients to new configurations during their normal hours of service without compromising the established peering arrangements.



This ebook was brought to you by [Noction](#).

Noction Intelligent Routing Platform enables enterprises and service providers to maximize end-to-end network performance and safely reduce infrastructure costs. The platform evaluates critical network performance metrics in real-time and responds quickly by automatically rerouting traffic through a better path to avoid outages and congestion.

Request a free trial today and see how IRP can boost your network performance.

[Start a Free Trial](#)