# Multihoming with BGP and NAT

## Eliminating ISP as a single point of failure

www.noction.com

## Table of Contents

## Multihoming with BGP and NAT - Introduction

This tutorial discusses the configuration of a multihomed enterprise network where routers CE-1 and CE-2 in AS 64501 are connected to routers ISP-A in AS64500 and ISP-B in AS64502 for redundancy. The connection via ISP-A is used as a primary connection for both outbound and inbound traffic. The connection via ISP-B is a backup connection. Identical inside hosts behind NAT are translated to different addresses assigned by the respective ISPs depending on whether traffic is forwarded via ISP-A or ISP-B to the Internet.

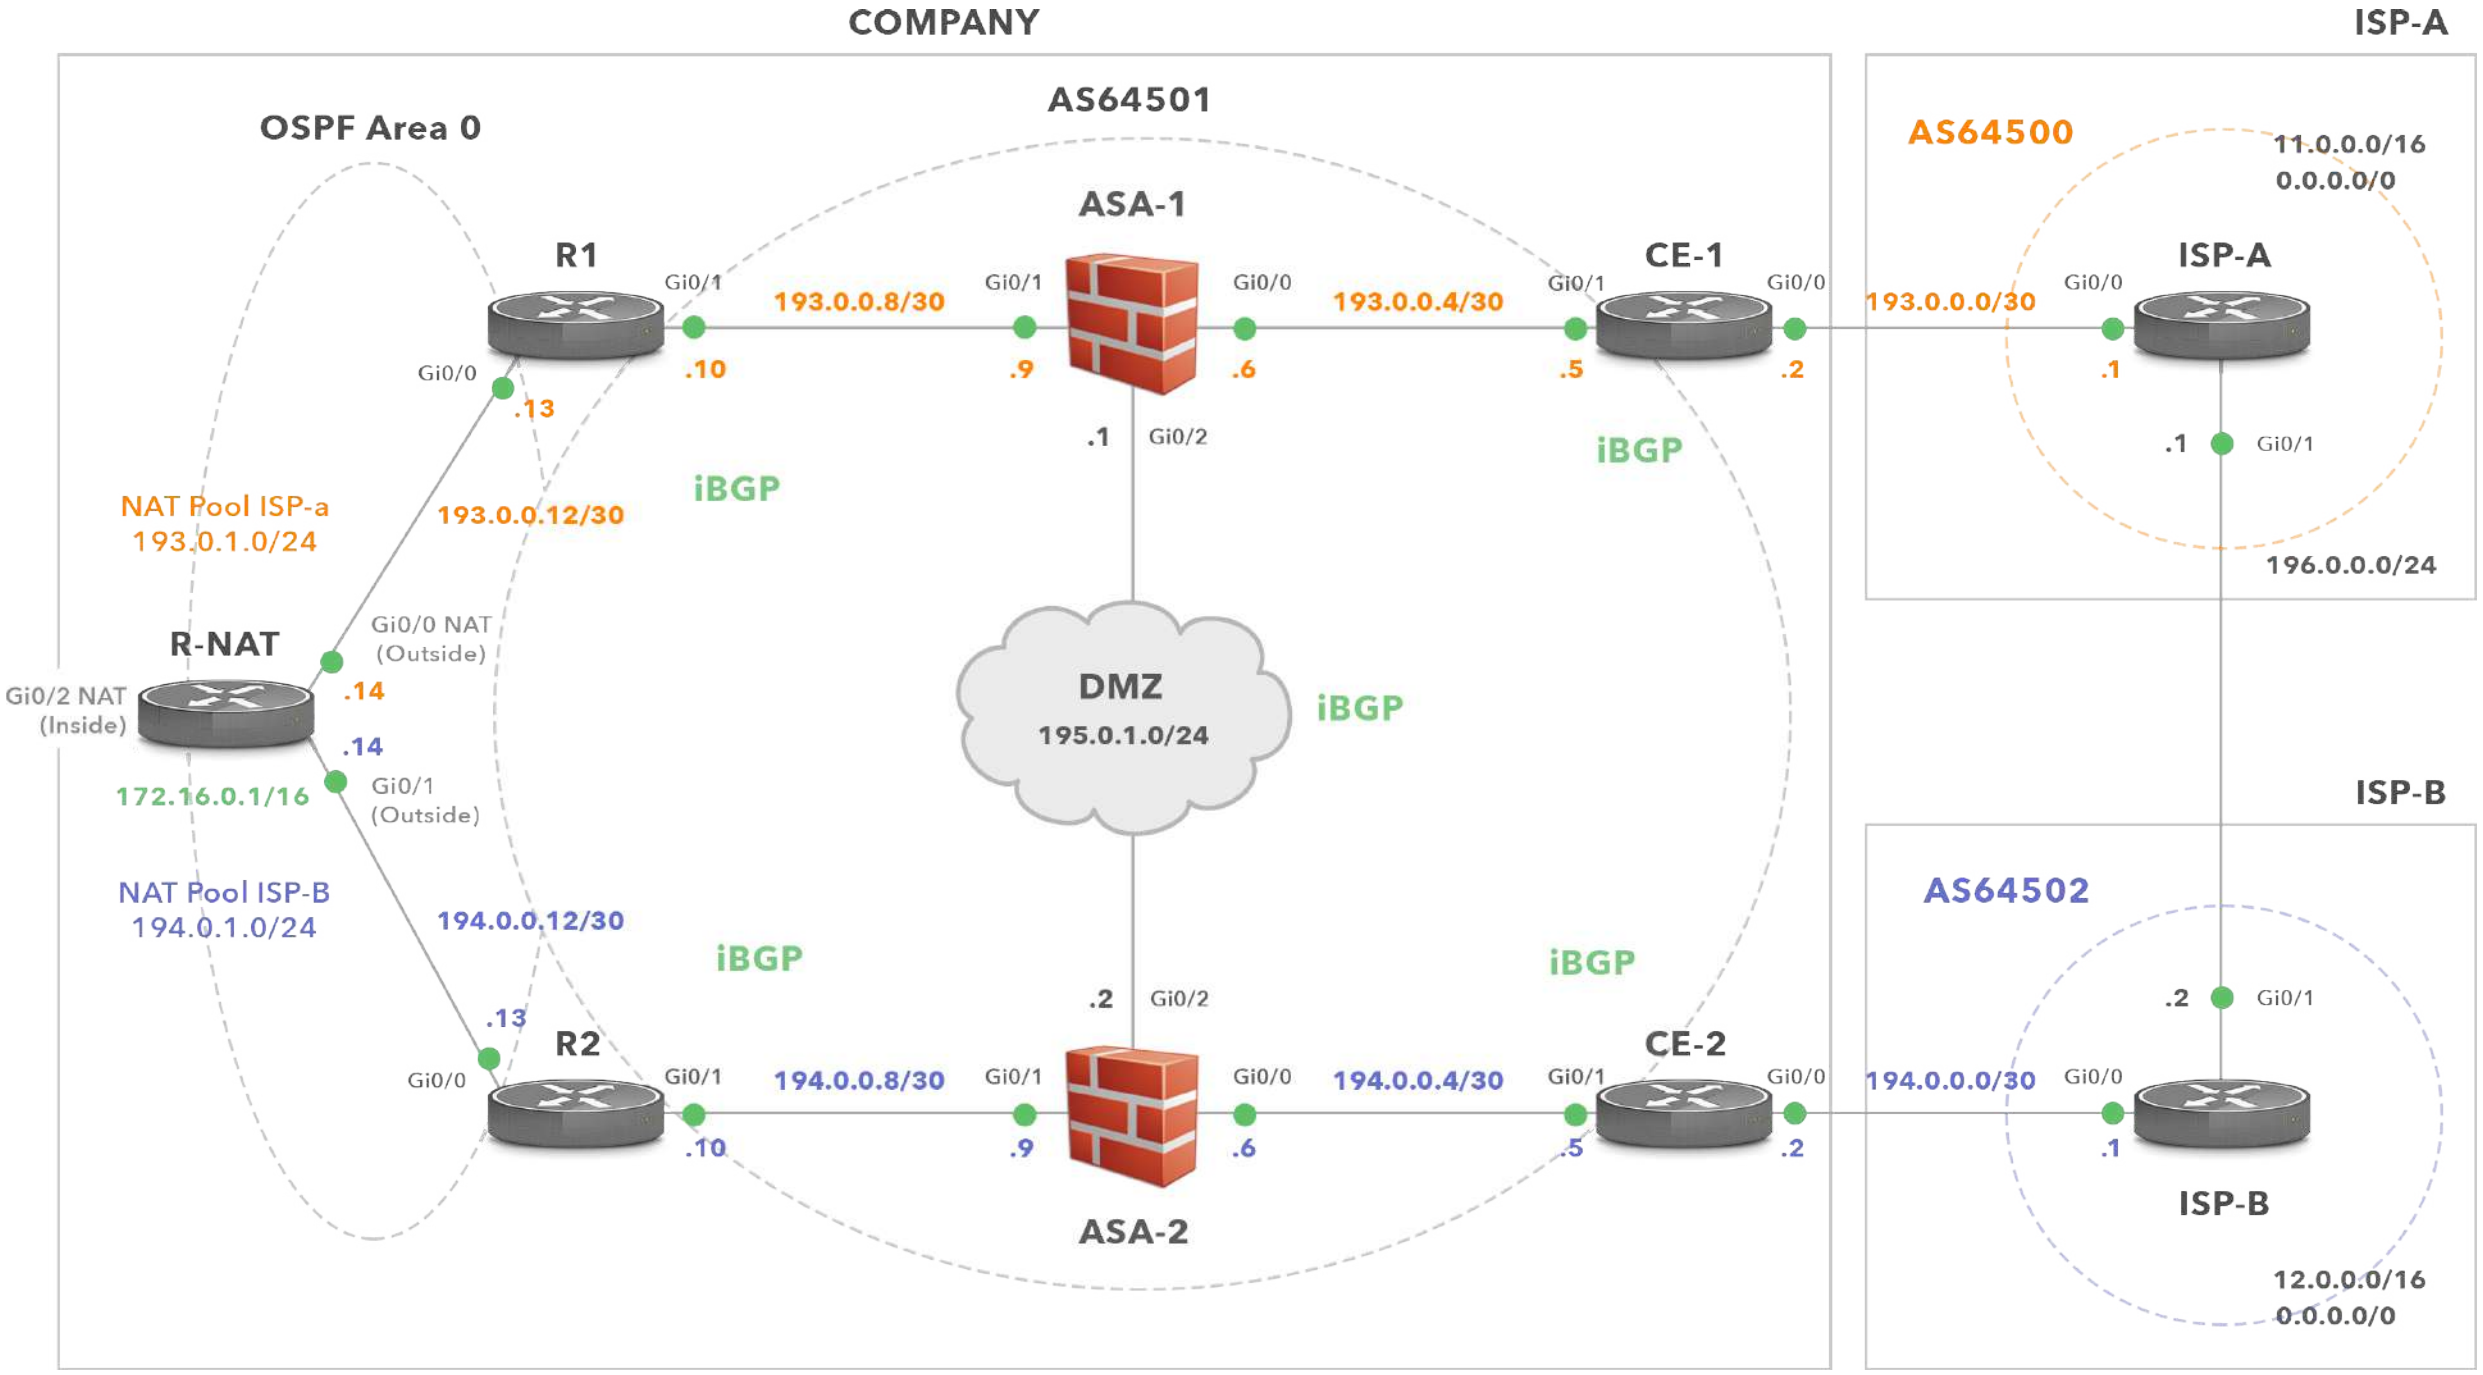

**Diagram 1** - *Enterprise Network (AS64501) is Multi-homed to ISP-A and ISP-B*

Let's go through every line of our configuration to explain its purpose. Below are the rules for outbound traffic from the enterprise to the Internet.

- All outbound traffic from 172.16.0.0/16 forwarded from R-NAT to R1 router has an inside global IP addresses assigned from the 193.0.1.0/24.
- All outbound traffic from 172.16.0.0/16 forwarded to R2 router has an inside global IP address assigned from the 194.0.1.0/24.
- AS64501 prefers the path from CE-1 to ISP-A for outbound traffic. Hence, a preferred outbound path for traffic sent from hosts behind NAT into the Internet is R-NAT-->R1-->ASA-1-->CE-1-->ISP-A. We give this path the name nat-north-internet. NAT pool 193.0.1.0/24 is used for mapping the inside local addressees (172.16.0.0/16) to the inside global addresses when the nat-north-internet path is taken.
- DMZ prefers a path via ASA-1, CE-1 and ISP-A for outbound traffic to the Internet. The path's name is dmz-north-internet path.
- If connectivity to ISP-A fails, all outbound traffic from hosts behind NAT is forwarded via a link from CE-2 to ISP-B. Therefore, it takes the path R-NAT-->R2-->ASA-2-->CE-2-->ISP-B, aka nat-south-internet path. The NAT pool 194.0.1.0/24 is used for mapping the inside local addresses to the inside global addresses in this case.
- If connectivity to ISP-A fails, all outbound traffic from hosts in DMZ to the Internet is routed via a link from CE-2 to ISP-B. In this case, traffic is forwarded to ASA-2-->CE-2-->ISP-B, via the dmz-south-internet path.

Here are the rules for inbound traffic.

- All inbound traffic that comes from the Internet to the enterprise uses the link from ISP-A to CE-1. This traffic is destined either for hosts in DMZ or for NAT pool 193.0.1.0/24 (translated local inside IP addresses 172.16.0.0/16).
- If a link from ISP-A to CE-1 fails, all inbound traffic from the Internet is routed via the link from ISP-B to CE-2. This traffic is destined either for DMZ or for NAT pool 194.0.1.0/24.

**IP Addresses Assignment:**
The enterprise has received a prefix from each ISP. These are used for NAT and interfaces configuration (Diagram 1). In addition, the enterprise has also assigned the prefix 195.0.1.0/24 to be used for DMZ configuration. The prefix 193.0.0.0/23 is assigned from ISP-A. This prefix consists of two /24 subnets - 193.0.0.0 and 193.0.1.0. The company uses the subnet 193.0.0.0/24 for IP address configuration of devices located on the north path (R1, ASA-1 and CE-1). The prefix 193.0.1.0/24 is reserved for NAT pool ISP-A. The enterprise has assigned the prefix 194.0.0.0/23 from ISP-B. The 194.0.0.0/24 sub-

net has been allocated for the south path configuration (R2, ASA-2 and CE-2). The prefix 194.0.1.0/24 is reserved for NAT pool ISP-B.

### Prefix Advertisement and Path Selection

The IP address range 195.0.1.0/24 is used for IP address configuration of devices in DMZ. The prefix is advertised by routers CE-1 and CE-2 via eBGP to ISP-A and ISP-B, respectively. However, since the dmz-north-internet path is preferred over the dmz-south-internet for the outbound traffic from DMZ to the Internet, we set a local preference to 150 for a default route 0.0.0.0/0 via CE-1. It effectively makes the path via CE-1 preferred as the default local preference is 100 for a default route installed in the routing table of CE-2.

The prefix for NAT 193.0.1.0/24 is announced solely by CE-1 in an eBGP update to the ISP-A and from ISP-A into the Internet. As CE-1 is the only router advertising this prefix, the inbound traffic sent from the Internet to the NAT prefix 193.0.1.0/24 takes the internet-north-nat path. The inbound traffic destined for the DMZ 195.0.1.0/24 is also routed via ISP-A to CE-1 and ASA-1 (internet-north-dmz path). The path via ISP-A gets selected by BGP routers located in the other ASs because CE-2 is configured to prepend as-path three times with its own AS 64501 for the DMZ route 195.0.1.0/24 advertised to CE-2. Therefore, the shorter AS_PATH via ISP-A is preferred.
If a link between CE-1 and ISP-A fails, the 193.0.1.0/24 prefix is not advertised by AS64501 at all. It might seem to be a design mistake at first, however when that link fails, the path nat-north-internet is not being used anymore. Instead, outbound traffic from R-NAT to ISP-B is routed via nat-south-internet path with the source IP addresses translated to the pool NAT 194.0.1.0/24. As the prefix 194.0.1.0/24 is advertised by CE-2, the devices behind NAT can communicate with devices in the Internet.

**NOTE:** Routers CE-1 and CE-2 contain the full Internet routing table. The Internet routes are simulated by prefixes 11.0.0.0/16 and 12.0.0./16 advertised by ISPs via an eBGP update to CE routers.

### Default Route Distribution

We have already mentioned that if a link between ISP-A and CE-1 fails, outbound traffic from devices behind NAT into the Internet is routed via a backup path nat-south-internet. But how does this magic work? Both CE-1 and CE-2 advertise a default route to R1 and R2 in an iBGP update message, respectively. However, they do it conditionally as they advertise a default route only if there is an appropriate route (11.0.0.0/16 for ISP-A and 12.0.0.0/16 for ISP-B) along with the ISP's IP address as a next-hop installed in their routing table. If not, CE routers do not advertise a default route to R1, R2 and DMZ.

Routers R1 and R2 advertise a default route to R-NAT conditionally, based on the links between CE and ISP routers being active. For instance, if R1 receives a default route via iBGP from CE-1 it installs it into its routing table. If the route-map CHECK-DEFAULT matches a default route and the next-hop IP address (CE-1), R1 advertises it via OSPF to R1, with the metric 5. R2, however advertises a default route (if the link between CE-2 and ISP-B is active) via OSPF to R-NAT with the metric 30. As a result, R-NAT installs a default route received from R1 with the metric 5 since it is lower than metric 30 of the default route advertised by R2.

### Adjusting Administrative Distance (AD) for iBGP learned Default Route on R1

For the iBGP routes with a default AD value 200 to be prefered over OSPF routes with a default AD value 110, we need to change the AD of iBGP routes bellow 110. If we use the command distance 20 105 200 under the BGP configuration of R1, a default route with AD 105 received from CE-1 in an iBGP update message has an AD lower than 110. The AD of a default route advertised by R-NAT from R2 is 110. Therefore, R1 installs a route 0.0.0.0 via CE-1 into its routing table.

If a link between CE-1 and ISP-1 goes down, R-NAT installs a default route with AD 110 and metric 30 advertised by R2 via OSPF. R1 also installs a default route received from R-NAT into its routing table. Outbound traffic is then sent to ISP-B. If the link between CE-1 and ISP-A goes up, a default route via CE-1 will be reinstalled into the routing table of R1.. The AD of this route is 105 thus it will be preferred to a default route with AD 110 advertised by R-NAT. Outbound traffic will be routed via ISP-A again.

**NOTE:** The default values of the command distance are bgp 20 200 200. The eBGP-learned routes have an administrative distance of 20, iBGP-learned routes have an administrative distance of 200, and local BGP routes have an administrative distance of 200

# 1. R-NAT Configuration

### 1.1 NAT Configuration

AS64501 is multi-homed, connected to ISP-A and ISP-B for redundancy. We use subnet 193.0.1.0/24  for mapping t he inside local addresses to the inside global addresses when outbound traffic from inside hosts is routed via R1. Similarly, the subnet 194.0.1.0/24 is used for translation when outbound traffic from the inside host is routed via R2. Therefore, we have created two NAT pools - 193.0.1.0/24 and 194.0.1.0/24, one for each ISP respectively. The same local inside address from the subnet 172.16.0.0/16 will be translated to different inside global address available in ISP-A and ISP-B NAT pools.

We also implement the pools overload. It means that the first address from the pool and port is to be used. Once all ports are exhausted, the second address from the pool will be used as an inside global address, etc. In theory, up to 65535 inside local addresses could be mapped to a single inside global address (based on the 16-bit port number). This type of NAT is called Dynamic NAT.

Router R-NAT has been configured with the IP address 172.16.0.1/16 on the GigabitEthernet0/2 interface. It is the subnet from a private IPv4 address space (RFC1918) that is going to be translated to NAT pools. Hosts located behind NAT have their IP addresses assigned from this address range. Therefore, we need to tell R-NAT that the interface Gi0/2 is located within the inside network. We will do it with the command ip nat inside.

```
interface GigabitEthernet0/2
 description LAN interface
 ip address 172.16.0.1 255.255.0.0
 ip nat inside
```

The interfaces GigabitEthernet0/0 and 0/1 are located in the outside network  thus we configure them as the outside interfaces with the command ip nat outside.

```
interface GigabitEthernet0/0
 description Link to R1
 ip address 193.0.0.14 255.255.255.252
 ip nat outside

interface GigabitEthernet0/1
 description Link to R2
 ip address 194.0.0.14 255.255.255.252
 ip nat outside
```

Let's configure NAT pools. The NAT pool ISP-A contains addresses assigned by ISP-A.  The ISP-B pool contains IP addresses assigned by ISP-B.

```
ip nat pool ISP-A 193.0.1.1 193.0.1.254 netmask
255.255.255.0
ip nat pool ISP-B 194.0.1.1 194.0.1.254 netmask
255.255.255.0
```

The line below configures a Dynamic NAT mapping of the inside network 172.16.0.0/16 to a global address from the pool ISP-A. Inside addresses are matched by the route-map ISP-A.

```
ip nat inside source route-map ISP-A pool ISP-A overload
```

We will do the same for ISP-B pool.

```
ip nat inside source route-map ISP-B pool ISP-B overload
```

Let's  create a route-map ISP-A. It matches all traffic matched by the access-list (ACL) NAT and going out of the interface Gi0/0.

```
route-map ISP-A permit 10
 match ip address NAT
 match interface GigabitEthernet0/0
```

The route-map ISP-B matches all traffic matched by ACL NAT and going out of interface Gi0/1.

```
route-map ISP-B permit 10
 match ip address NAT
 match interface GigabitEthernet0/1
```

Finally, we create the ACL NAT which is the named source ACL matching traffic from all hosts in the inside network.

```
ip access-list standard NAT
 permit 172.16.0.0 0.0.255.255
```

### 1.2 OSPF Configuration

R-NAT router is running OSPF and forms OSPF adjacency with routers R1 and R2. Therefore, we must advertise the subnets 193.0.0.12/30 and 194.0.0.12/30. The OSPF point-to-point network type must be configured for both loopback interfaces, otherwise the NAT networks will be advertised as host networks with prefix length /32.

```
interface Loopback0
 ip address 193.0.1.1 255.255.255.0
 ip ospf network point-to-point

interface Loopback1
 ip address 194.0.1.1 255.255.255.0
 ip ospf network point-to-point
```

Let's also implement the inter-area filtering on R-NAT to prevent routes from other areas being advertised into the area 0. 0.0.0.0/0 le 32 matches any prefix with a length between 0 and 32 bits (inclusive). This matches all possible IPv4 prefixes.

```
router ospf 1
 area 0 filter-list prefix ADV-TO-R in
 network 193.0.0.12 0.0.0.3 area 0
 network 193.0.1.0 0.0.0.255 area 0
 network 194.0.0.12 0.0.0.3 area 0
 network 194.0.1.0 0.0.0.255 area 0

ip prefix-list ADV-TO-R seq 999 deny 0.0.0.0/0 le 32
```

## 2. R1 and R2 Configuration

### 2.1 R1 Configuration

```
interface GigabitEthernet0/0
 description Link to R-NAT
 ip address 193.0.0.13 255.255.255.252

interface GigabitEthernet0/1
 description Link to ASA-1
 ip address 193.0.0.10 255.255.255.252
```

Now let’s configure OSPF to advertise a default route conditionally, based on whether the link between CE-1 and ISP-1 is active, with the metric 5. Hence, R-NAT will prefer the route via R1 for outgoing traffic to the default route learned via OSPF from R2, with metric 30.

```
router ospf 1
 network 193.0.0.12 0.0.0.3 area 0
 default-information originate metric 5 route-map
CHECK-DEFAULT
```

The R1 router is an iBGP peer with the routers CE-1 and DMZ. The bgp redistribute-internal command can be used as a workaround here. Without this command, even though R1 has a default route with the next-hop 193.0.0.5 learned via iBGP installed in the routing table, a default route is not advertised to R-NAT.

The redistribution of OSPF routes into iBGP is done with the help of the redistribute ospf 1 command. The command distance bgp 20 105 200 changes the default AD from 200 to 105 for routes learned via iBGP.

```
router bgp 64501
 bgp log-neighbor-changes
 bgp redistribute-internal
 redistribute ospf 1
 neighbor 193.0.0.5 remote-as 64501
 neighbor 193.0.0.5 next-hop-self
 neighbor 195.0.1.3 remote-as 64501
 neighbor 195.0.1.3 next-hop-self
 distance bgp 20 105 200
```

Now we need static routes to iBGP peers as they are not directly connected.

```
ip route 193.0.0.4 255.255.255.252 193.0.0.9
ip route 195.0.1.3 255.255.255.255 193.0.0.9

route-map CHECK-DEFAULT permit 10
 match ip address 30
 match ip next-hop 31

access-list 30 permit 0.0.0.0
access-list 31 permit 193.0.0.5
```

### 2.2 R2 Configuration

```
interface GigabitEthernet0/0
description Link to R-NAT
ip address 194.0.0.13 255.255.255.252

interface GigabitEthernet0/1
description Link to ASA-2
ip address 194.0.0.10 255.255.255.252
```

We will configure OSPF to advertise a default route conditionally, based on whether the link between CE-1 and ISP-1 is active, with the metric 30.

```
router ospf 1
 network 194.0.0.12 0.0.0.3 area 0
 default-information originate metric 30 route-map
CHECK-DEFAULT
```

The R2 router is an iBGP peer with the routers CE-2 and DMZ. The command redistribute ospf 1 redistributes OSPF routes into iBGP.

```
router bgp 64501
 bgp log-neighbor-changes
 bgp redistribute-internal
 redistribute ospf 1
 neighbor 194.0.0.5 remote-as 64501
 neighbor 194.0.0.5 next-hop-self
 neighbor 195.0.1.3 remote-as 64501
 neighbor 195.0.1.3 next-hop-self
```

We need static routes to iBGP peers as they are not directly connected.

```
ip route 194.0.0.4 255.255.255.252 194.0.0.9
ip route 195.0.1.3 255.255.255.255 194.0.0.9

route-map CHECK-DEFAULT permit 10
 match ip address 30
 match ip next-hop 31

access-list 30 permit 0.0.0.0
access-list 31 permit 194.0.0.5
```

So far, we have finished the configuration of R-NAT, R1 and R2. Let's continue and complete the configuration of the remaining devices in our topology (Diagram 1). We'll start with the ASA configuration.

## 3. ASA Configuration

The Cisco Adaptive Security Appliance (ASA) protects the inside network and DMZ. As our guide focuses on a multi-homing configuration using BGP, we only cover the basic ASA configuration. It includes the access-lists configuration to allow BGP in all direction. In order to protect the enterprise network from advanced threats, application layer protocol inspection should be configured, in addition to the access-lists configuration.

### 3.1 ASA-1 Configuration

```
interface GigabitEthernet0/0
 description Link to CE-1
 nameif OUTSIDE
 security-level 0
 ip address 193.0.0.6 255.255.255.252

interface GigabitEthernet0/1
 description Link to R1
 nameif INSIDE
 security-level 100
 ip address 193.0.0.9 255.255.255.252

interface GigabitEthernet0/2
 description Link to DMZ
 nameif DMZ
 security-level 50
 ip address 195.0.1.1 255.255.255.0
```

Router R1 can initiate a TCP connection to CE-1 (193.0.0.5), destination TCP port 179 since R1 is connected to the interface Gi0/1 of ASA, configured with security level 100. Therefore, R1 can establish an iBGP adjacency with CE-1. However, we need to configure the access-list 1 (ACL1) that allows to initiate a TCP connection from CE-1 (outside) to R1 (inside), with the destination IP address 193.0.0.10 and TCP port 179. As the interface Gi0/0 is configured with a security level 0, we need to add the rule that permits traffic from CE-1 to DMZ router (195.0.1.3), with the destination TCP port 179. The statement permits traffic from the interface Gi0/0 with the security level 0 to the interface Gi0/2 with higher security level - 50. Therefore, CE-1 can initiate a TCP connection to the DMZ router.

```
access-list ACL1 extended permit tcp host 193.0.0.5
host 193.0.0.10 eq bgp
access-list ACL1 extended permit tcp host 193.0.0.5
host 195.0.1.3 eq bgp
```

The ACL1 is applied on the outside interface (Gi0/0) in the inbound direction.

```
access-group ACL1 in interface OUTSIDE
```

The ACL2 contains a rule that permits TCP traffic from interface Gi0/2 connected to the DMZ and configured with security level 50, to the interface Gi0/1 with a level 100, destination IP 193.0.0.10 and TCP port 179 (BGP).

```
access-list ACL2 extended permit tcp host 195.0.1.3
host 193.0.0.10 eq bgp
```

The ACL2 is applied on the DMZ interface (Gi0/2) in the inbound direction.

```
access-group ACL2 in interface DMZ
```

ASA-1 is not participating in OSPF, so we need static routes in order to forward traffic to subnets that are outside the Gi0/1 interface. The subnets are NAT pools 193.0.1.0/24 (ISP-A), 194.0.1.0/24 (ISP-B), and 193.0.0.12/30, all routed via the next-hop IP address 193.0.0.10 (R1). The default route for forwarding outbound traffic to the Internet is configured with the next-hop 193.0.0.5 (CE-1).

```
route INSIDE 193.0.1.0 255.255.255.0 193.0.0.10
route INSIDE 194.0.1.0 255.255.255.0 193.0.0.10
route INSIDE 193.0.0.12 255.255.255.252 193.0.0.10
route OUTSIDE 0.0.0.0 0.0.0.0 193.0.0.5
```

### 3.2 ASA-2 Configuration

```
interface GigabitEthernet0/0
 description Link to CE-2
 nameif OUTSIDE
 security-level 0
 ip address 194.0.0.6 255.255.255.252

interface GigabitEthernet0/1
 description Link to R2
 nameif INSIDE
 security-level 100
 ip address 194.0.0.9 255.255.255.252

interface GigabitEthernet0/2
 description Link to DMZ
 nameif DMZ
 security-level 50
 ip address 195.0.1.2 255.255.255.0
```

Access-lists configuration is similar to ASA-1. ACL1 permits BGP traffic from the outside BGP peer 194.0.0.5 (CE-1) to the peer 194.0.0.10 (inside) (R2) and to the peer 195.0.1.3 (DMZ) (router DMZ). ACL2 permits BGP traffic from a peer in DMZ interface to the inside R2.

```
access-list ACL1 extended permit tcp host 194.0.0.5
host 194.0.0.10 eq bgp
access-list ACL1 extended permit tcp host 194.0.0.5
host 195.0.1.3 eq bgp
access-list ACL2 extended permit tcp host 195.0.1.3
host 194.0.0.10 eq bgp
```

ACLs are applied in the inbound direction to the outside and DMZ interfaces.

```
access-group ACL1 in interface OUTSIDE
access-group ACL2 in interface DMZ
```

We need to configure static routes to reach subnets behind the interface Gi0/1 and a default static route for outgoing traffic to the Internet.

```
route INSIDE 193.0.1.0 255.255.255.0 194.0.0.10
route INSIDE 194.0.0.12 255.255.255.252 194.0.0.10
route INSIDE 194.0.1.0 255.255.255.0 194.0.0.10
route OUTSIDE 0.0.0.0 0.0.0.0 194.0.0.5
```

## 4. Customer Edge Routers Configuration

### 4.1 CE-1 Configuration

```
interface GigabitEthernet0/0
 description Link to ISP-A
 ip address 193.0.0.2 255.255.255.252

interface GigabitEthernet0/1
 description Link to ASA-1
 ip address 193.0.0.5 255.255.255.252
```

CE-1 router is configured as an eBGP peer with the router ISP-A and as an iBGP with routers R1 and DMZ. Only the prefixes 193.0.1.0/24 (NAT pool ISP-A) and 195.0.1.0/24 (DMZ) matched by route-map ADV-TO-ISPA are advertised to ISP-A. The route-map CHECK-ISPA-ROUTE checks if the route 11.0.0.0/16 with the next-hop 193.0.0.1 is available in the routing table of CE-1. If yes, a default route is advertised to the BGP neighbors: R1 (193.0.0.10) and DMZ (195.0.1.3). The default route is then advertised conditionally, based on whether a link between CE-1 and ISP-A is active.

CE-1 receives a full Internet table from ISP-A. The routes are simulated by the prefix 11.0.0.0/16 advertised by ISP-A. However, only a default route is conditionally advertised to R1 and DMZ routers. Therefore, we permit only advertisement of the default route in outgoing direction to R1 and DMZ with distribute-list that refers to access-list 1.

```
router bgp 64501
 neighbor 193.0.0.1 remote-as 64500
 neighbor 193.0.0.1 route-map ADV-TO-ISPA out
 neighbor 193.0.0.10 remote-as 64501
 neighbor 193.0.0.10 next-hop-self
  neighbor  193.0.0.10  default-originate  route-map
CHECK-ISPA-ROUTE
 neighbor 193.0.0.10 distribute-list 1 out
 neighbor 195.0.1.3 remote-as 64501
 neighbor 195.0.1.3 next-hop-self
  neighbor  195.0.1.3  default-originate  route-map
CHECK-ISPA-ROUTE
 neighbor 195.0.1.3 distribute-list 1 out
```

Static routes to iBGP peers are required because peers are not directly connected.

```
ip route 193.0.0.8 255.255.255.252 193.0.0.6
ip route 195.0.1.3 255.255.255.255 193.0.0.6

route-map CHECK-ISPA-ROUTE permit 10
 match ip address 20
 match ip next-hop 21

route-map ADV-TO-ISPA permit 10
 match ip address 10 11
```

```
access-list 1 permit 0.0.0.0
access-list 10 permit 193.0.1.0 0.0.0.255
access-list 11 permit 195.0.1.0 0.0.0.255
access-list 20 permit 11.0.0.0 0.0.255.255
access-list 21 permit 193.0.0.1
```

### 4.2 CE-2 Configuration

```
interface GigabitEthernet0/0
 description Link to ISP-B
 ip address 194.0.0.2 255.255.255.252

interface GigabitEthernet0/1
 description Link to ASA-2
 ip address 194.0.0.5 255.255.255.252
```

CE-2 is configured as an eBGP peer with ISP-B and as an iBGP peer with routers R2 and DMZ. The prefixes 194.0.1.0/24 (NAT pool ISP-B) and 195.0.1.0/24 (DMZ) are matched by route-map ADV-TO-ISPB and advertised to ISP-B. Default route is advertised conditionally to R2 and DMZ routers based on whether a link between CE-2 and ISP-B is active.

```
router bgp 64501
 neighbor 194.0.0.1 remote-as 64502
 neighbor 194.0.0.1 route-map ADV-TO-ISPB out
 neighbor 194.0.0.10 remote-as 64501
 neighbor 194.0.0.10 next-hop-self
  neighbor  194.0.0.10  default-originate  route-map
```

```
CHECK-ISPB-ROUTE
 neighbor 194.0.0.10 distribute-list 1 out
 neighbor 195.0.1.3 remote-as 64501
 neighbor 195.0.1.3 next-hop-self
  neighbor  195.0.1.3  default-originate  route-map
CHECK-ISPB-ROUTE
 neighbor 195.0.1.3 distribute-list 1 out
```

Static routes to iBGP peers are required because they are not directly connected.

```
ip route 194.0.0.8 255.255.255.252 194.0.0.6
ip route 195.0.1.3 255.255.255.255 194.0.0.6
```

```
route-map CHECK-ISPB-ROUTE permit 10
 match ip address 20
 match ip next-hop 21
```

Both CE-1 and CE-2 routers advertise DMZ route 195.0.1.0/24 to their respective ISPs. However, incoming traffic to DMZ is forwarded via ISP-A because CE-1 is configured to prepend AS_PATH 64501 three times for the route 195.0.1.0/24 advertised to ISP-B. Therefore, BGP routers select a shorter path via ISP-A for traffic to 195.0.1.0/24.

```
route-map ADV-TO-ISPB permit 10
 match ip address 10
```

```
route-map ADV-TO-ISPB permit 20
 match ip address 11
  set as-path prepend 64501 64501 64501

access-list 1 permit 0.0.0.0
access-list 10 permit 194.0.1.0 0.0.0.255
access-list 11 permit 195.0.1.0 0.0.0.255
access-list 20 permit 12.0.0.0 0.0.255.255
access-list 21 permit 194.0.0.1
```

## 5. ISPs Routers Configuration

ISP-A and ISP-B are configured as eBGP peers with CE-1 and CE-2, respectively. They also peer between themselves. Both ISPs advertise full Internet routing table, simulated by prefixes 11.0.0.0/16 and 12.0.0.0/16.

### 5.1 ISP-A Configuration

```
interface GigabitEthernet0/0
 description Link to CE-1
 ip address 193.0.0.1 255.255.255.252

interface GigabitEthernet0/1
 description Link to ISP-B
 ip address 196.0.0.1 255.255.255.252
 router bgp 64500
 network 11.0.0.0 mask 255.255.0.0
 neighbor 193.0.0.2 remote-as 64501
 neighbor 196.0.0.2 remote-as 64502
```

BGP will always advertise the network 11.0.0.0/16 because a null route is installed in the routing table of ISP-A.

```
ip route 11.0.0.0 255.255.0.0 Null0
```

### 5.2 ISP-B Configuration

```
interface GigabitEthernet0/0
 description Link to CE-2
 ip address 194.0.0.1 255.255.255.252

interface GigabitEthernet0/1
 description Link to ISP-A
 ip address 196.0.0.2 255.255.255.252

router bgp 64502
 network 12.0.0.0 mask 255.255.0.0
 neighbor 194.0.0.2 remote-as 64501
 neighbor 196.0.0.1 remote-as 64500
```

BGP will always advertise the network 12.0.0.0/16 because a null route is installed in a routing table of ISP-B.

```
ip route 12.0.0.0 255.255.0.0 Null0
```

In the last part of our tutorial we will verify that our configuration works properly. Let's start with the R-NAT router.

## 6. Verification

As a first step, we are going to check whether NAT is working as expected. The source IP address of traffic going out of the interface Gi0/0 is being translated to an IP address from the NAT pool ISP-A (193.0.1.0/24). Multiple inside local addresses must be translated to a single inside global address as port address translations (PAT) or NAT overload is configured.  Issue the ping command multiple times and check the NAT translation table on R-NAT.

```
R-NAT# ping 193.0.0.13 source GigabitEthernet 0/2
R-NAT# ping 193.0.0.13 source GigabitEthernet 0/2
```

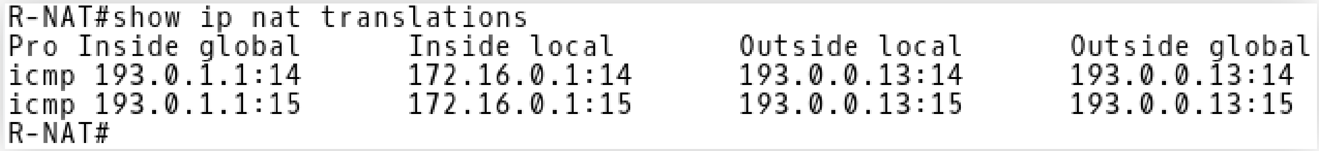

```
R-NAT#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 193.0.1.1:14      172.16.0.1:14      193.0.0.13:14      193.0.0.13:14
icmp 193.0.1.1:15      172.16.0.1:15      193.0.0.13:15      193.0.0.13:15
R-NAT#
```

**Picture 1** - *NAT Translation Table when Traffic is Sent out of Gi0/0*

The protocol is ICMP and an inside local address 172.16.0.1 is translated to the inside global address 193.0.1.1 (NAT pool ISP-A). Notice the numbers 14 and 15 mapped to the IP address. Those are ICMP identifiers included in the ICMP header of the IP packet (Picture 2).  Their map associated ICMP echo requests to echo replies.

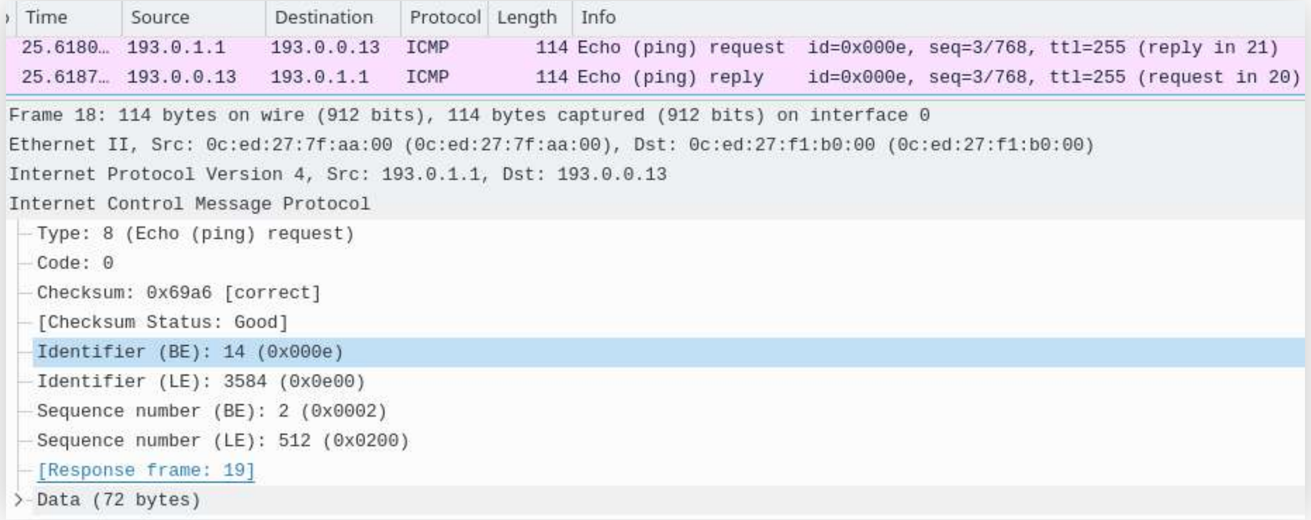

```
 Time       Source      Destination  Protocol Length Info
 25.6180…   193.0.1.1   193.0.0.13   ICMP         114 Echo (ping) request  id=0x000e, seq=3/768, ttl=255 (reply in 21)
 25.6187…   193.0.0.13  193.0.1.1    ICMP         114 Echo (ping) reply    id=0x000e, seq=3/768, ttl=255 (request in 20)
Frame 18: 114 bytes on wire (912 bits), 114 bytes captured (912 bits) on interface 0
Ethernet II, Src: 0c:ed:27:7f:aa:00 (0c:ed:27:7f:aa:00), Dst: 0c:ed:27:f1:b0:00 (0c:ed:27:f1:b0:00)
Internet Protocol Version 4, Src: 193.0.1.1, Dst: 193.0.0.13
Internet Control Message Protocol
  Type: 8 (Echo (ping) request)
  Code: 0
  Checksum: 0x69a6 [correct]
  [Checksum Status: Good]
  Identifier (BE): 14 (0x000e)
  Identifier (LE): 3584 (0x0e00)
  Sequence number (BE): 2 (0x0002)
  Sequence number (LE): 512 (0x0200)
  [Response frame: 19]
> Data (72 bytes)
```

**Picture 2** - *ICMP header of translated IP Packet*

Now we will check whether PAT is in use issuing the ping command to 193.0.0.13 from two hosts (172.16.0.100 and 172.16.0.101), connected to the interface Gi0/2 of the R-NAT router. First, let's delete the NAT translation table of R-NAT.

```
R-NAT# clear ip nat translation *
```

```
Host1$ ping 193.0.0.13
Host2$ ping 193.0.0.13
```

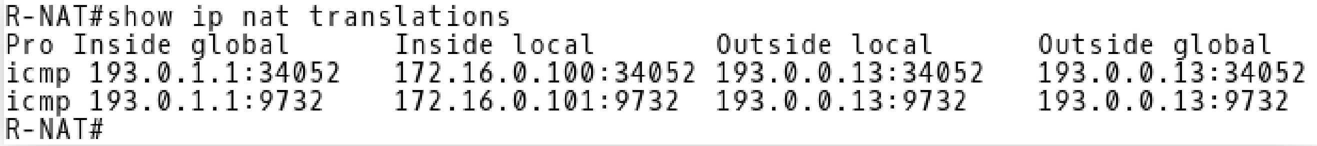

```
R-NAT#show ip nat translations
Pro Inside global        Inside local        Outside local       Outside global
icmp 193.0.1.1:34052     172.16.0.100:34052  193.0.0.13:34052    193.0.0.13:34052
icmp 193.0.1.1:9732      172.16.0.101:9732   193.0.0.13:9732     193.0.0.13:9732
R-NAT#
```

**Picture 3** - *NAT Translation Table when Traffic is Sent out of Gi0/0 from Two Hosts*

PAT is in use because the router R-NAT has translated the inside local addresses 172.16.0.100 and 172.16.0.101 to a single inside global address 193.0.1.1 (Picture 3).

It is time to check the NAT translation table when ICMP echo request is sent out of the interface Gi0/1. The inside local address is dynamically mapped to the inside global address 194.0.1.2 when traffic is forwarded out of the Gi0/1. ICMP identifier is 16 inside the ICMP header of the IP packets related to the first ping command (Picture 4). ICMP identifier is 17 for the IP packets related to the second ping command.

```
R-NAT# clear ip nat translation *
R-NAT# ping 194.0.0.13 source GigabitEthernet 0/2
R-NAT# ping 194.0.0.13 source GigabitEthernet 0/2
```

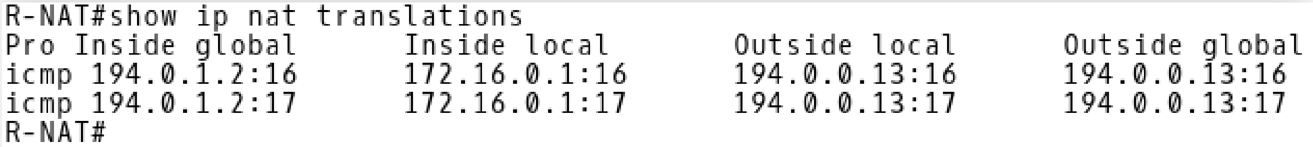

```
R-NAT#show ip nat translations
Pro Inside global      Inside local      Outside local     Outside global
icmp 194.0.1.2:16      172.16.0.1:16     194.0.0.13:16     194.0.0.13:16
icmp 194.0.1.2:17      172.16.0.1:17     194.0.0.13:17     194.0.0.13:17
R-NAT#
```

**Picture 4** - *NAT Translation Table when Traffic is Sent out of Gi0/1*

When the link between CE-1 and ISP-A is active, router R-NAT installs OSPF E2 external type route with a route metric 5 received from 193.0.0.13 (R1) into its routing table (Picture 5).

```
R-NAT# show ip route 0.0.0.0
```

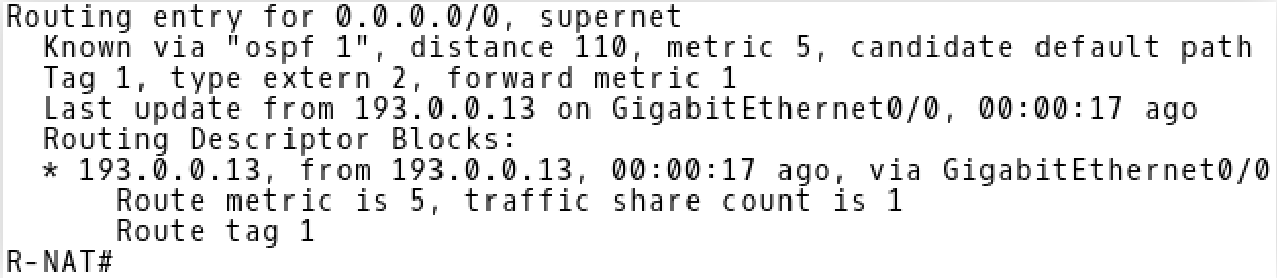

```
Routing entry for 0.0.0.0/0, supernet
  Known via "ospf 1", distance 110, metric 5, candidate default path
  Tag 1, type extern 2, forward metric 1
  Last update from 193.0.0.13 on GigabitEthernet0/0, 00:00:17 ago
  Routing Descriptor Blocks:
  * 193.0.0.13, from 193.0.0.13, 00:00:17 ago, via GigabitEthernet0/0
      Route metric is 5, traffic share count is 1
      Route tag 1
R-NAT#
```

**Picture 5** - OSPF E2 Route Installed in Routing Table of R-NAT when Link Between CE-1 and ISP-A is Active

Now, check the BGP table on R1. The R1 router learns the default route via iBGP from 193.0.0.5 (CE-1) (Picture 6).

```
R1# show ip bgp | begin Network
```

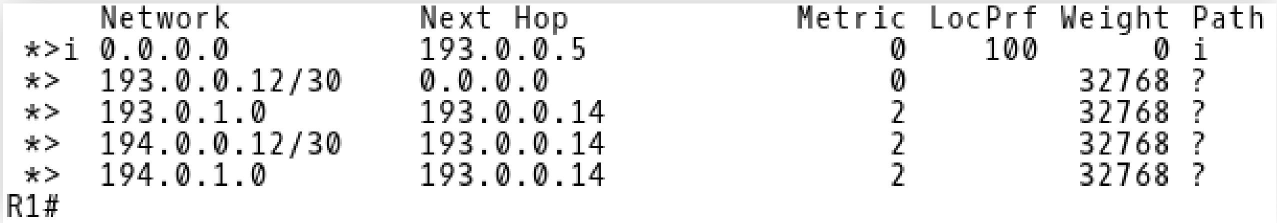

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 0.0.0.0          193.0.0.5                0    100      0 i
 *>  193.0.0.12/30    0.0.0.0                  0         32768 ?
 *>  193.0.1.0        193.0.0.14               2         32768 ?
 *>  194.0.0.12/30    193.0.0.14               2         32768 ?
 *>  194.0.1.0        193.0.0.14               2         32768 ?
R1#
```

**Picture 6** - Default Route in BGP Table of R1 Learned via iBGP from CE-1

R2 learns the default route via iBGP from 194.0.0.5 (CE-2) (Picture 7). However, the route is not installed into the routing table of R2 because it is rejected. There is a default route with a better Administrative Distance (AD) 110 received from another source - 194.0.0.14 (R-NAT) learned via OSPF (Picture 8). The AD of iBGP route is 200. Therefore, a default route advertised by R-NAT with a lower AD 110 is installed into the routing table of R2.

```
R2# show ip bgp | begin Network
```

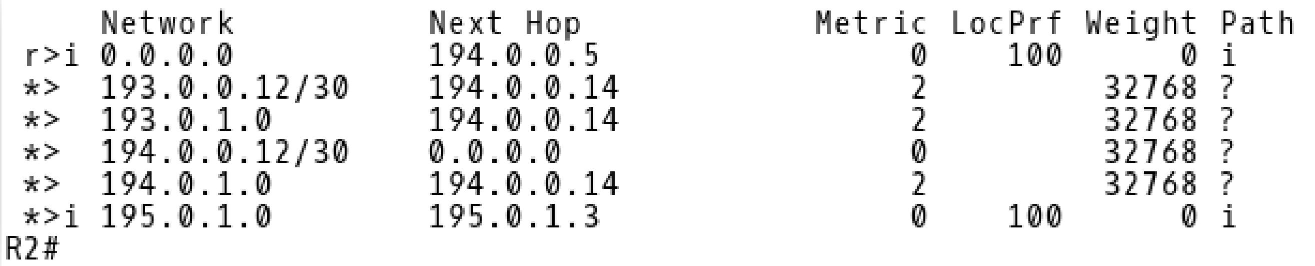

```
     Network          Next Hop            Metric LocPrf Weight Path
 r>i 0.0.0.0          194.0.0.5                0    100      0 i
 *>  193.0.0.12/30    194.0.0.14               2         32768 ?
 *>  193.0.1.0        194.0.0.14               2         32768 ?
 *>  194.0.0.12/30    0.0.0.0                  0         32768 ?
 *>  194.0.1.0        194.0.0.14               2         32768 ?
 *>i 195.0.1.0        195.0.1.3                0    100      0 i
R2#
```

**Picture 7** - *Default Route in BGP Table of R2 Learned via iBGP from CE-2*

```
R2# show ip route ospf
```

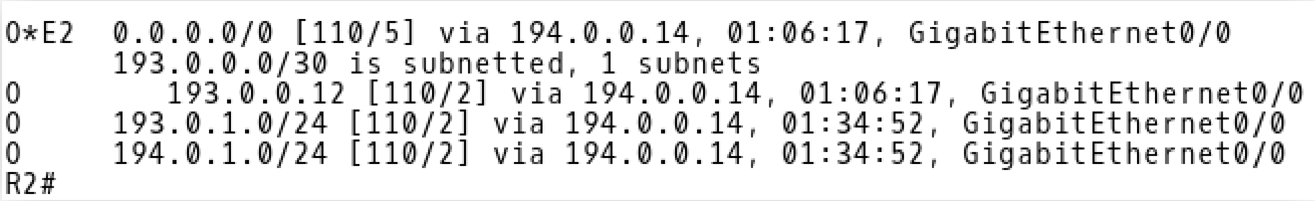

```
O*E2  0.0.0.0/0 [110/5] via 194.0.0.14, 01:06:17, GigabitEthernet0/0
      193.0.0.0/30 is subnetted, 1 subnets
O        193.0.0.12 [110/2] via 194.0.0.14, 01:06:17, GigabitEthernet0/0
O     193.0.1.0/24 [110/2] via 194.0.0.14, 01:34:52, GigabitEthernet0/0
O     194.0.1.0/24 [110/2] via 194.0.0.14, 01:34:52, GigabitEthernet0/0
R2#
```

**Picture 8** - *Default Route in Routing Table of R2 Learned from OSPF*

R1 is the originator of the default route advertised to R-NAT via OSPF and received by R2. Notice the OSPF router ID 193.0.0.13 (Picture 9).

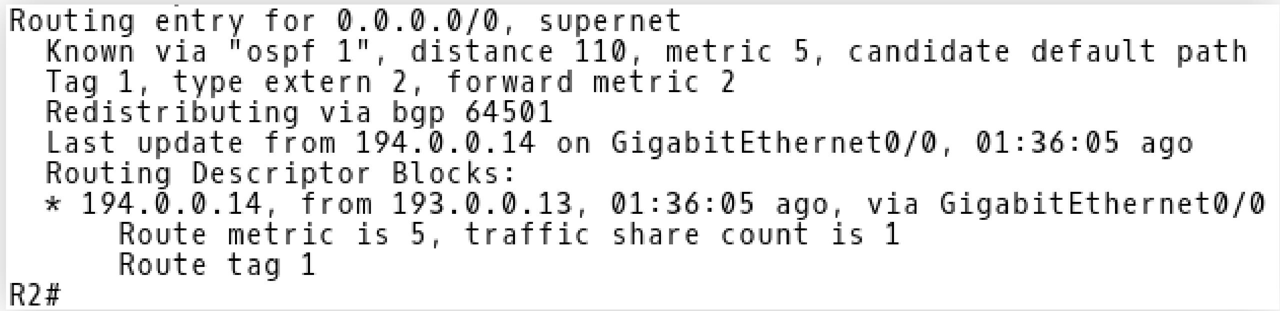

```
Routing entry for 0.0.0.0/0, supernet
  Known via "ospf 1", distance 110, metric 5, candidate default path
  Tag 1, type extern 2, forward metric 2
  Redistributing via bgp 64501
  Last update from 194.0.0.14 on GigabitEthernet0/0, 01:36:05 ago
  Routing Descriptor Blocks:
  * 194.0.0.14, from 193.0.0.13, 01:36:05 ago, via GigabitEthernet0/0
      Route metric is 5, traffic share count is 1
      Route tag 1
R2#
```

**Picture 9** - *OSPF Router ID 193.0.0.13*

The R2 router learns about both NAT prefixes 193.0.1.0/24 and 194.0.1.0/24 from R-NAT via OSPF (Picture 8). OSPF routes are then redistributed to iBGP (Picture 7). The DMZ route 195.0.1.0 is learned via iBGP from DMZ router (195.0.1.3).

Let's check the BGP table of the DMZ router located in DMZ (Picture 10). The router learns the default route from two sources. The local preference for default route 0.0.0.0 via 193.0.0.5 (CE-1) on DMZ router is set to 150. The local preference for the route 0.0.0.0 and the 194.0.0.5 (CE-2) neighbor is 100 by default. The path via iBGP neighbor 193.0.0.5 (CE-1) is preferred because the preference value 150 is higher than 100.

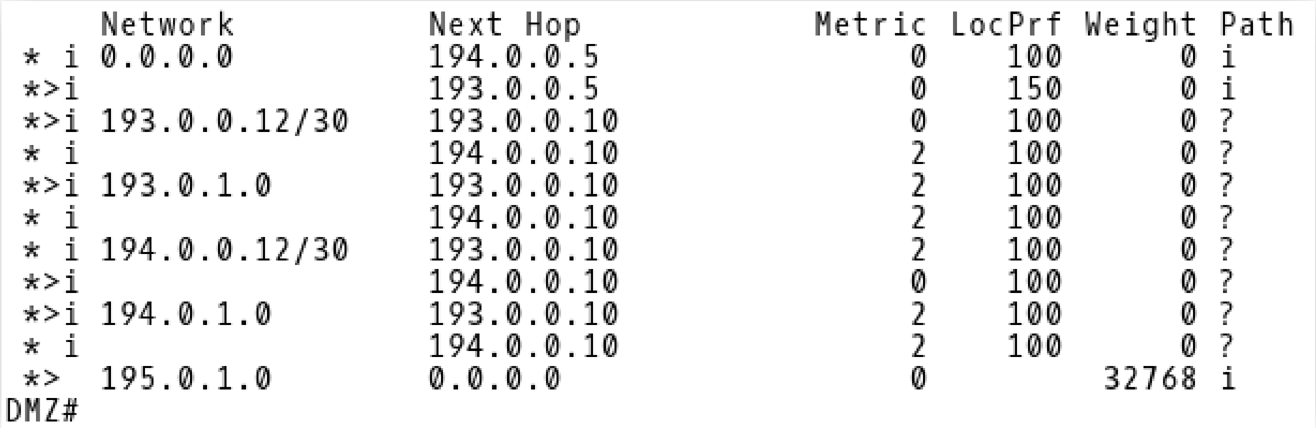

```
     Network          Next Hop            Metric LocPrf Weight Path
 * i 0.0.0.0          194.0.0.5                0    100      0 i
 *>i                  193.0.0.5                0    150      0 i
 *>i 193.0.0.12/30    193.0.0.10               0    100      0 ?
 * i                  194.0.0.10               2    100      0 ?
 *>i 193.0.1.0        193.0.0.10               2    100      0 ?
 * i                  194.0.0.10               2    100      0 ?
 * i 194.0.0.12/30    193.0.0.10               2    100      0 ?
 *>i                  194.0.0.10               0    100      0 ?
 *>i 194.0.1.0        193.0.0.10               2    100      0 ?
 * i                  194.0.0.10               2    100      0 ?
 *>  195.0.1.0        0.0.0.0                  0         32768 i
DMZ#
```

**Picture 10** - *BGP Table of Router DMZ*

The path via 193.0.0.10 (R1) to prefixes 193.0.1.0 and 194.0.1.0 is preferred to the path via 194.0.0.10 because the Router-ID 193.0.0.10 (R1) is lower than Router-ID 194.0.0.10 (R2). The route 195.0.1.0 is originated locally. The default weight for locally originated routes is 32768.

Let's inspect the BGP table of CE-1. CE-1 contains a full Internet routing table size, simulated by the 11.0.0.0/16 and 12.0.0.0/16 networks and received from the 193.0.0.1 (ISP-A) peer (Picture 11).

```
CE-1# show ip bgp
```

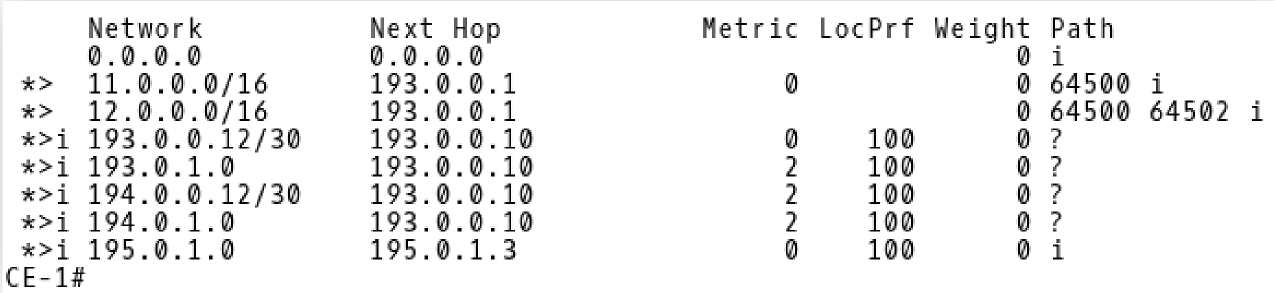

```
     Network          Next Hop            Metric LocPrf Weight Path
     0.0.0.0          0.0.0.0                                 0 i
 *>  11.0.0.0/16      193.0.0.1                0              0 64500 i
 *>  12.0.0.0/16      193.0.0.1                               0 64500 64502 i
 *>i 193.0.0.12/30    193.0.0.10               0    100       0 ?
 *>i 193.0.1.0        193.0.0.10               2    100       0 ?
 *>i 194.0.0.12/30    193.0.0.10               2    100       0 ?
 *>i 194.0.1.0        193.0.0.10               2    100       0 ?
 *>i 195.0.1.0        195.0.1.3                0    100       0 i
CE-1#
```

**Picture 11** - *BGP Table of CE-1*

Similarly, the router CE-2 receives the 12.0.0.0/16 and 11.0.0.0/16 routes from the 194.0.0.1(ISP-B) peer (Picture 12).

```
CE-2# show ip bgp
```

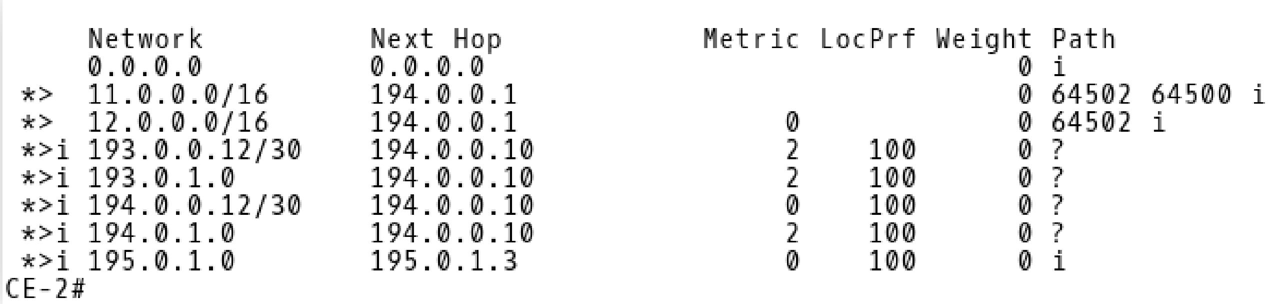

```
     Network          Next Hop            Metric LocPrf Weight Path
     0.0.0.0          0.0.0.0                                 0 i
 *>  11.0.0.0/16      194.0.0.1                               0 64502 64500 i
 *>  12.0.0.0/16      194.0.0.1                0              0 64502 i
 *>i 193.0.0.12/30    194.0.0.10               2    100       0 ?
 *>i 193.0.1.0        194.0.0.10               2    100       0 ?
 *>i 194.0.0.12/30    194.0.0.10               0    100       0 ?
 *>i 194.0.1.0        194.0.0.10               2    100       0 ?
 *>i 195.0.1.0        195.0.1.3                0    100       0 i
CE-2#
```

**Picture 12** - *BGP Table of CE-2*

As the last step of our verification, we will examine the BGP tables of both ISPs. (Picture 13 and 14).

```
ISP-A# show ip bgp
```

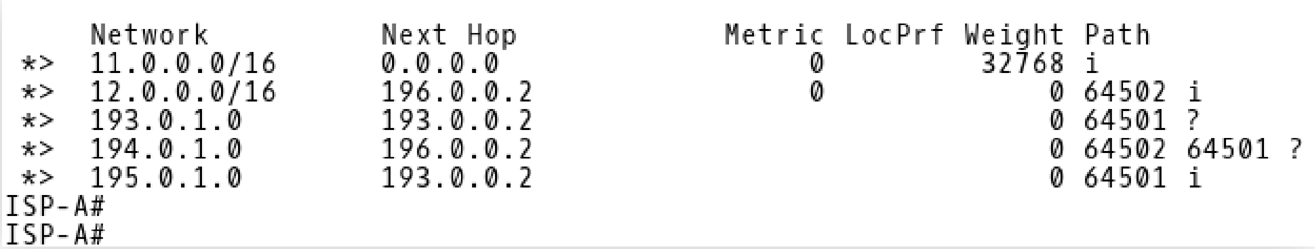

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>  11.0.0.0/16      0.0.0.0                  0          32768 i
 *>  12.0.0.0/16      196.0.0.2                0              0 64502 i
 *>  193.0.1.0        193.0.0.2                               0 64501 ?
 *>  194.0.1.0        196.0.0.2                               0 64502 64501 ?
 *>  195.0.1.0        193.0.0.2                               0 64501 i
ISP-A#
ISP-A#
```

**Picture 13** - *BGP Table of ISP-A*

**ISP-A** learns the DMZ route 195.0.1.0 from the 193.0.0.2 (CE-1) peer (Picture 13). **ISP-B**, however receives the same route from two sources – 196.0.0.1 (ISP-A) and 194.0.0.2 (CE-2). (Picture 14). The path via ISP-A is being installed into the routing table of ISP-B because a path with the shortest AS_PATH is preferred. The requirement is to forward inbound traffic from the Internet to DMZ via CE-1 when a link between CE-1 and ISP-A is active and via CE-2 when the link is down. Therefore, we have configured CE-2 to prepend as-path three times with its own AS 64501 for the route 195.0.1.0/24, advertised to ISP-B.

```
ISP-B# show ip bgp
```

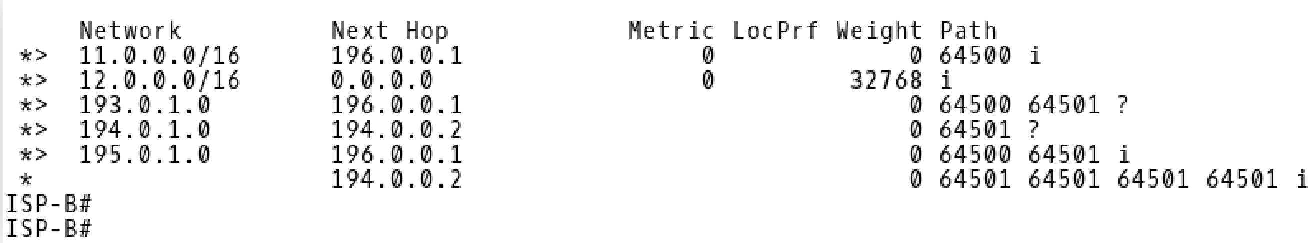

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>  11.0.0.0/16      196.0.0.1                0              0 64500 i
 *>  12.0.0.0/16      0.0.0.0                  0          32768 i
 *>  193.0.1.0        196.0.0.1                               0 64500 64501 ?
 *>  194.0.1.0        194.0.0.2                               0 64501 ?
 *>  195.0.1.0        196.0.0.1                               0 64500 64501 i
 *                    194.0.0.2                               0 64501 64501 64501 64501 i
ISP-B#
ISP-B#
```

**Picture 14** - *BGP Table of ISP-B*

## 7. Troubleshooting

We will bring down the session between CE-1 and ISP-A and shutdown the interface Gi0/0 on CE-1.

```
CE-1(config)# interface gigabitEthernet 0/0
CE-1(config-if)# shutdown
```

Let's check a routing table of the R-NAT router (Picture 15). R-NAT has installed OSPF E2 external route with a route metric 30 and the next-hop 194.0.0.13 (R1) into its routing table. It is the expected behavior as the backup path nat-south-internet is used for outgoing traffic when CE-1 (AS64501) loses eBGP session with ISP-A (AS64500).

```
R-NAT# show ip route ospf
```

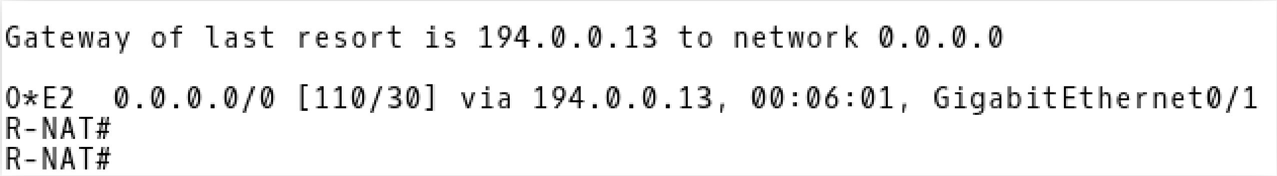

```
Gateway of last resort is 194.0.0.13 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/30] via 194.0.0.13, 00:06:01, GigabitEthernet0/1
R-NAT#
R-NAT#
```

**Picture 15** - *Default Route via R1 Installed in RT of R-NAT*

Issue the traceroute command from R-NAT to ISP-B (Picture 16) and check the NAT translation table (Picture 17).

```
R-NAT# traceroute 194.0.0.1 source gigabitEthernet 0/2
```

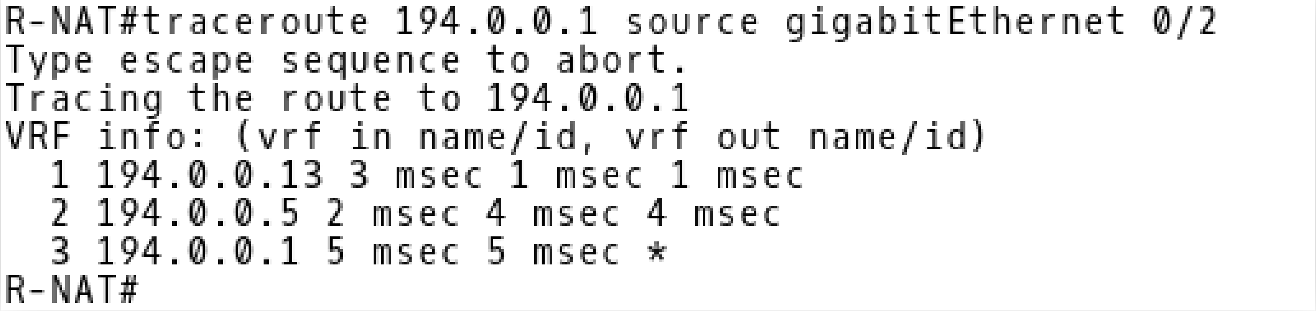

```
R-NAT#traceroute 194.0.0.1 source gigabitEthernet 0/2
Type escape sequence to abort.
Tracing the route to 194.0.0.1
VRF info: (vrf in name/id, vrf out name/id)
  1 194.0.0.13 3 msec 1 msec 1 msec
  2 194.0.0.5 2 msec 4 msec 4 msec
  3 194.0.0.1 5 msec 5 msec *
R-NAT#
```

**Picture 16** - *Traceroute from R-NAT to ISP-B*

The inside local 172.16.0.1 is being translated to the same inside global address 194.0.1.3 from the NAT pool ISP-B (194.0.1.0/24) (Picture 17). The PAT (NAT overload) is in use.

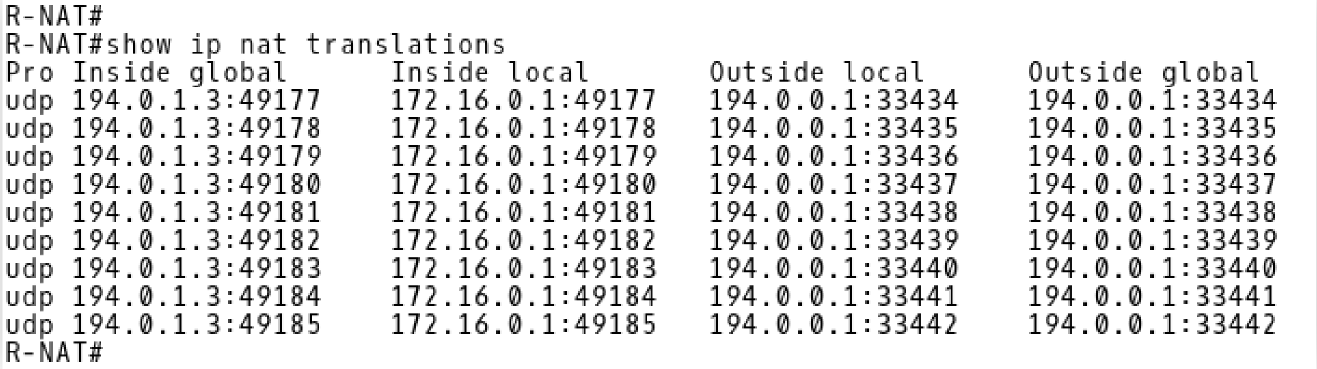

```
R-NAT#
R-NAT#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
udp 194.0.1.3:49177    172.16.0.1:49177   194.0.0.1:33434    194.0.0.1:33434
udp 194.0.1.3:49178    172.16.0.1:49178   194.0.0.1:33435    194.0.0.1:33435
udp 194.0.1.3:49179    172.16.0.1:49179   194.0.0.1:33436    194.0.0.1:33436
udp 194.0.1.3:49180    172.16.0.1:49180   194.0.0.1:33437    194.0.0.1:33437
udp 194.0.1.3:49181    172.16.0.1:49181   194.0.0.1:33438    194.0.0.1:33438
udp 194.0.1.3:49182    172.16.0.1:49182   194.0.0.1:33439    194.0.0.1:33439
udp 194.0.1.3:49183    172.16.0.1:49183   194.0.0.1:33440    194.0.0.1:33440
udp 194.0.1.3:49184    172.16.0.1:49184   194.0.0.1:33441    194.0.0.1:33441
udp 194.0.1.3:49185    172.16.0.1:49185   194.0.0.1:33442    194.0.0.1:33442
R-NAT#
```

**Picture 17** - *NAT Translation Table of R-NAT*

To see how the inbound traffic is routed from the Internet to AS64501, we will check the BGP table of ISP-A (Picture 18).

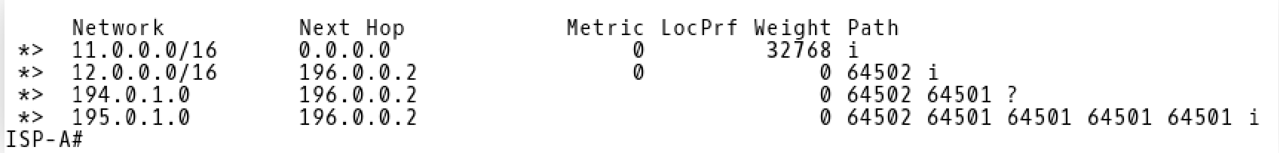

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>  11.0.0.0/16      0.0.0.0                  0         32768 i
 *>  12.0.0.0/16      196.0.0.2                0             0 64502 i
 *>  194.0.1.0        196.0.0.2                              0 64502 64501 ?
 *>  195.0.1.0        196.0.0.2                              0 64502 64501 64501 64501 64501 i
ISP-A#
```

**Picture 18** - *BGP Table of ISP-A*

There are routes 194.0.1.0/24 (NAT pool ISP-B) and 195.0.1.0/24 (DMZ) installed into the BGP table of ISP-A with the next-hop 196.0.0.2 (ISP-B). The route 193.0.1.0/24 (NAT pool ISP-A) is not installed but we do not need it. The inside local addresses are translated to the ISP-B NAT pool.

R1 router redistributes a default route via OSPF conditionally, when a default route 0.0.0.0 via the next-hop 193.0.0.5 (CE-1) is installed into its routing table. But how does the R1 learn about a default route when the route is not advertised from CE-1 to R1? Remember, a link between CE-1 and ISP-A is still down. R1 installed the E2 OSPF default route received from its OSPF neighbor 193.0.0.14 (R-NAT), with a metric 30 (Picture 19).

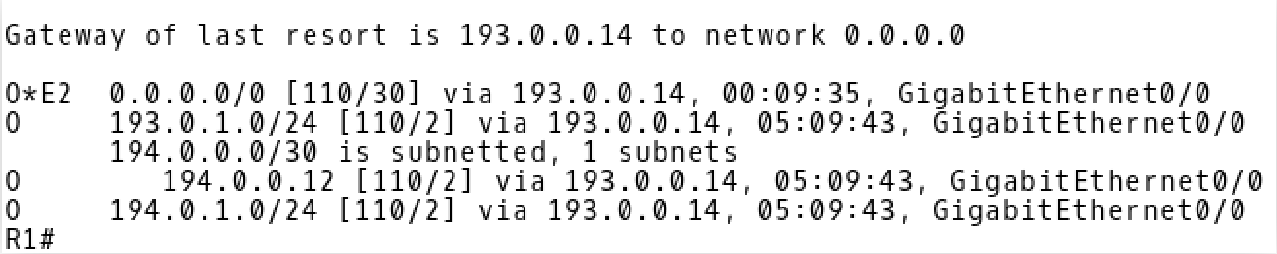

```
Gateway of last resort is 193.0.0.14 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/30] via 193.0.0.14, 00:09:35, GigabitEthernet0/0
O     193.0.1.0/24 [110/2] via 193.0.0.14, 05:09:43, GigabitEthernet0/0
      194.0.0.0/30 is subnetted, 1 subnets
O        194.0.0.12 [110/2] via 193.0.0.14, 05:09:43, GigabitEthernet0/0
O     194.0.1.0/24 [110/2] via 193.0.0.14, 05:09:43, GigabitEthernet0/0
R1#
```

**Picture 19** - *Default Route on R1 Received from R-NAT*

Let's bring the interface Gi0/0 on CE-1 up. The default route 0.0.0.0 via the the next-hop 193.0.0.13 (R1) is reinstalled into the routing table of R-NAT (Picture 20). The source IP addresses of outgoing traffic from the inside network to the Internet are translated to the ISP-A NAT pool again.

```
CE-1(config-if)# no shutdown
```

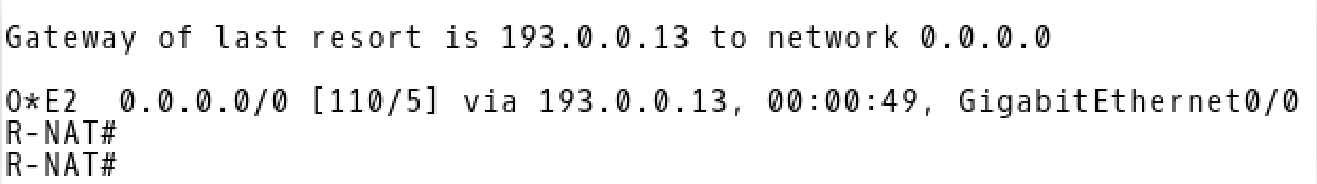

```
Gateway of last resort is 193.0.0.13 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/5] via 193.0.0.13, 00:00:49, GigabitEthernet0/0
R-NAT#
R-NAT#
```

**Picture 20** - *Default Route via R1 Installed in RT of R-NAT*

As soon as R1 receives a default route via iBGP from CE-1, it installs it into its routing table and the OSPF default route via R-NAT is purged (Picture 21). But why are the iBGP routes with the Administrative Distance (AD) 200 preferred over the OSPF route with the distance 110? Remember, we have changed the default AD 200 for iBGP learned routes to 105 in the configuration of R1? Therefore, iBGP routes are preferred over OSPF routes as iBGP AD is set to 105. As a result, outgoing traffic is routed via ISP-A when a link between CE-A and ISP-A is active.

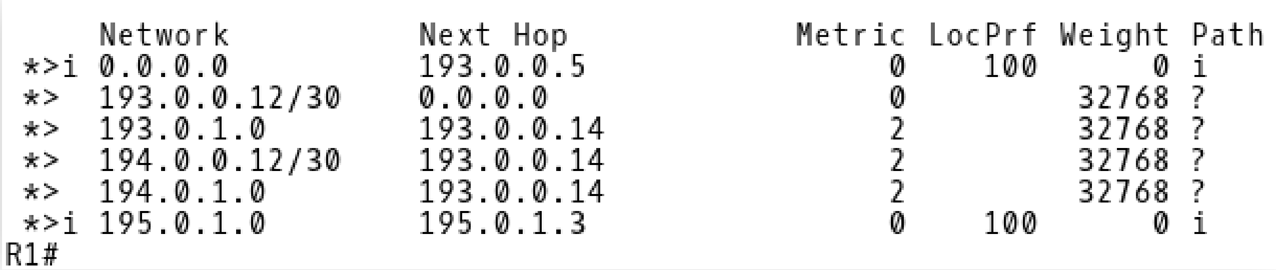

```
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 0.0.0.0          193.0.0.5                0    100      0 i
 *>  193.0.0.12/30    0.0.0.0                  0         32768 ?
 *>  193.0.1.0        193.0.0.14               2         32768 ?
 *>  194.0.0.12/30    193.0.0.14               2         32768 ?
 *>  194.0.1.0        193.0.0.14               2         32768 ?
 *>i 195.0.1.0        195.0.1.3                0    100      0 i
R1#
```

**Picture 21** - *BGP Table of R1*

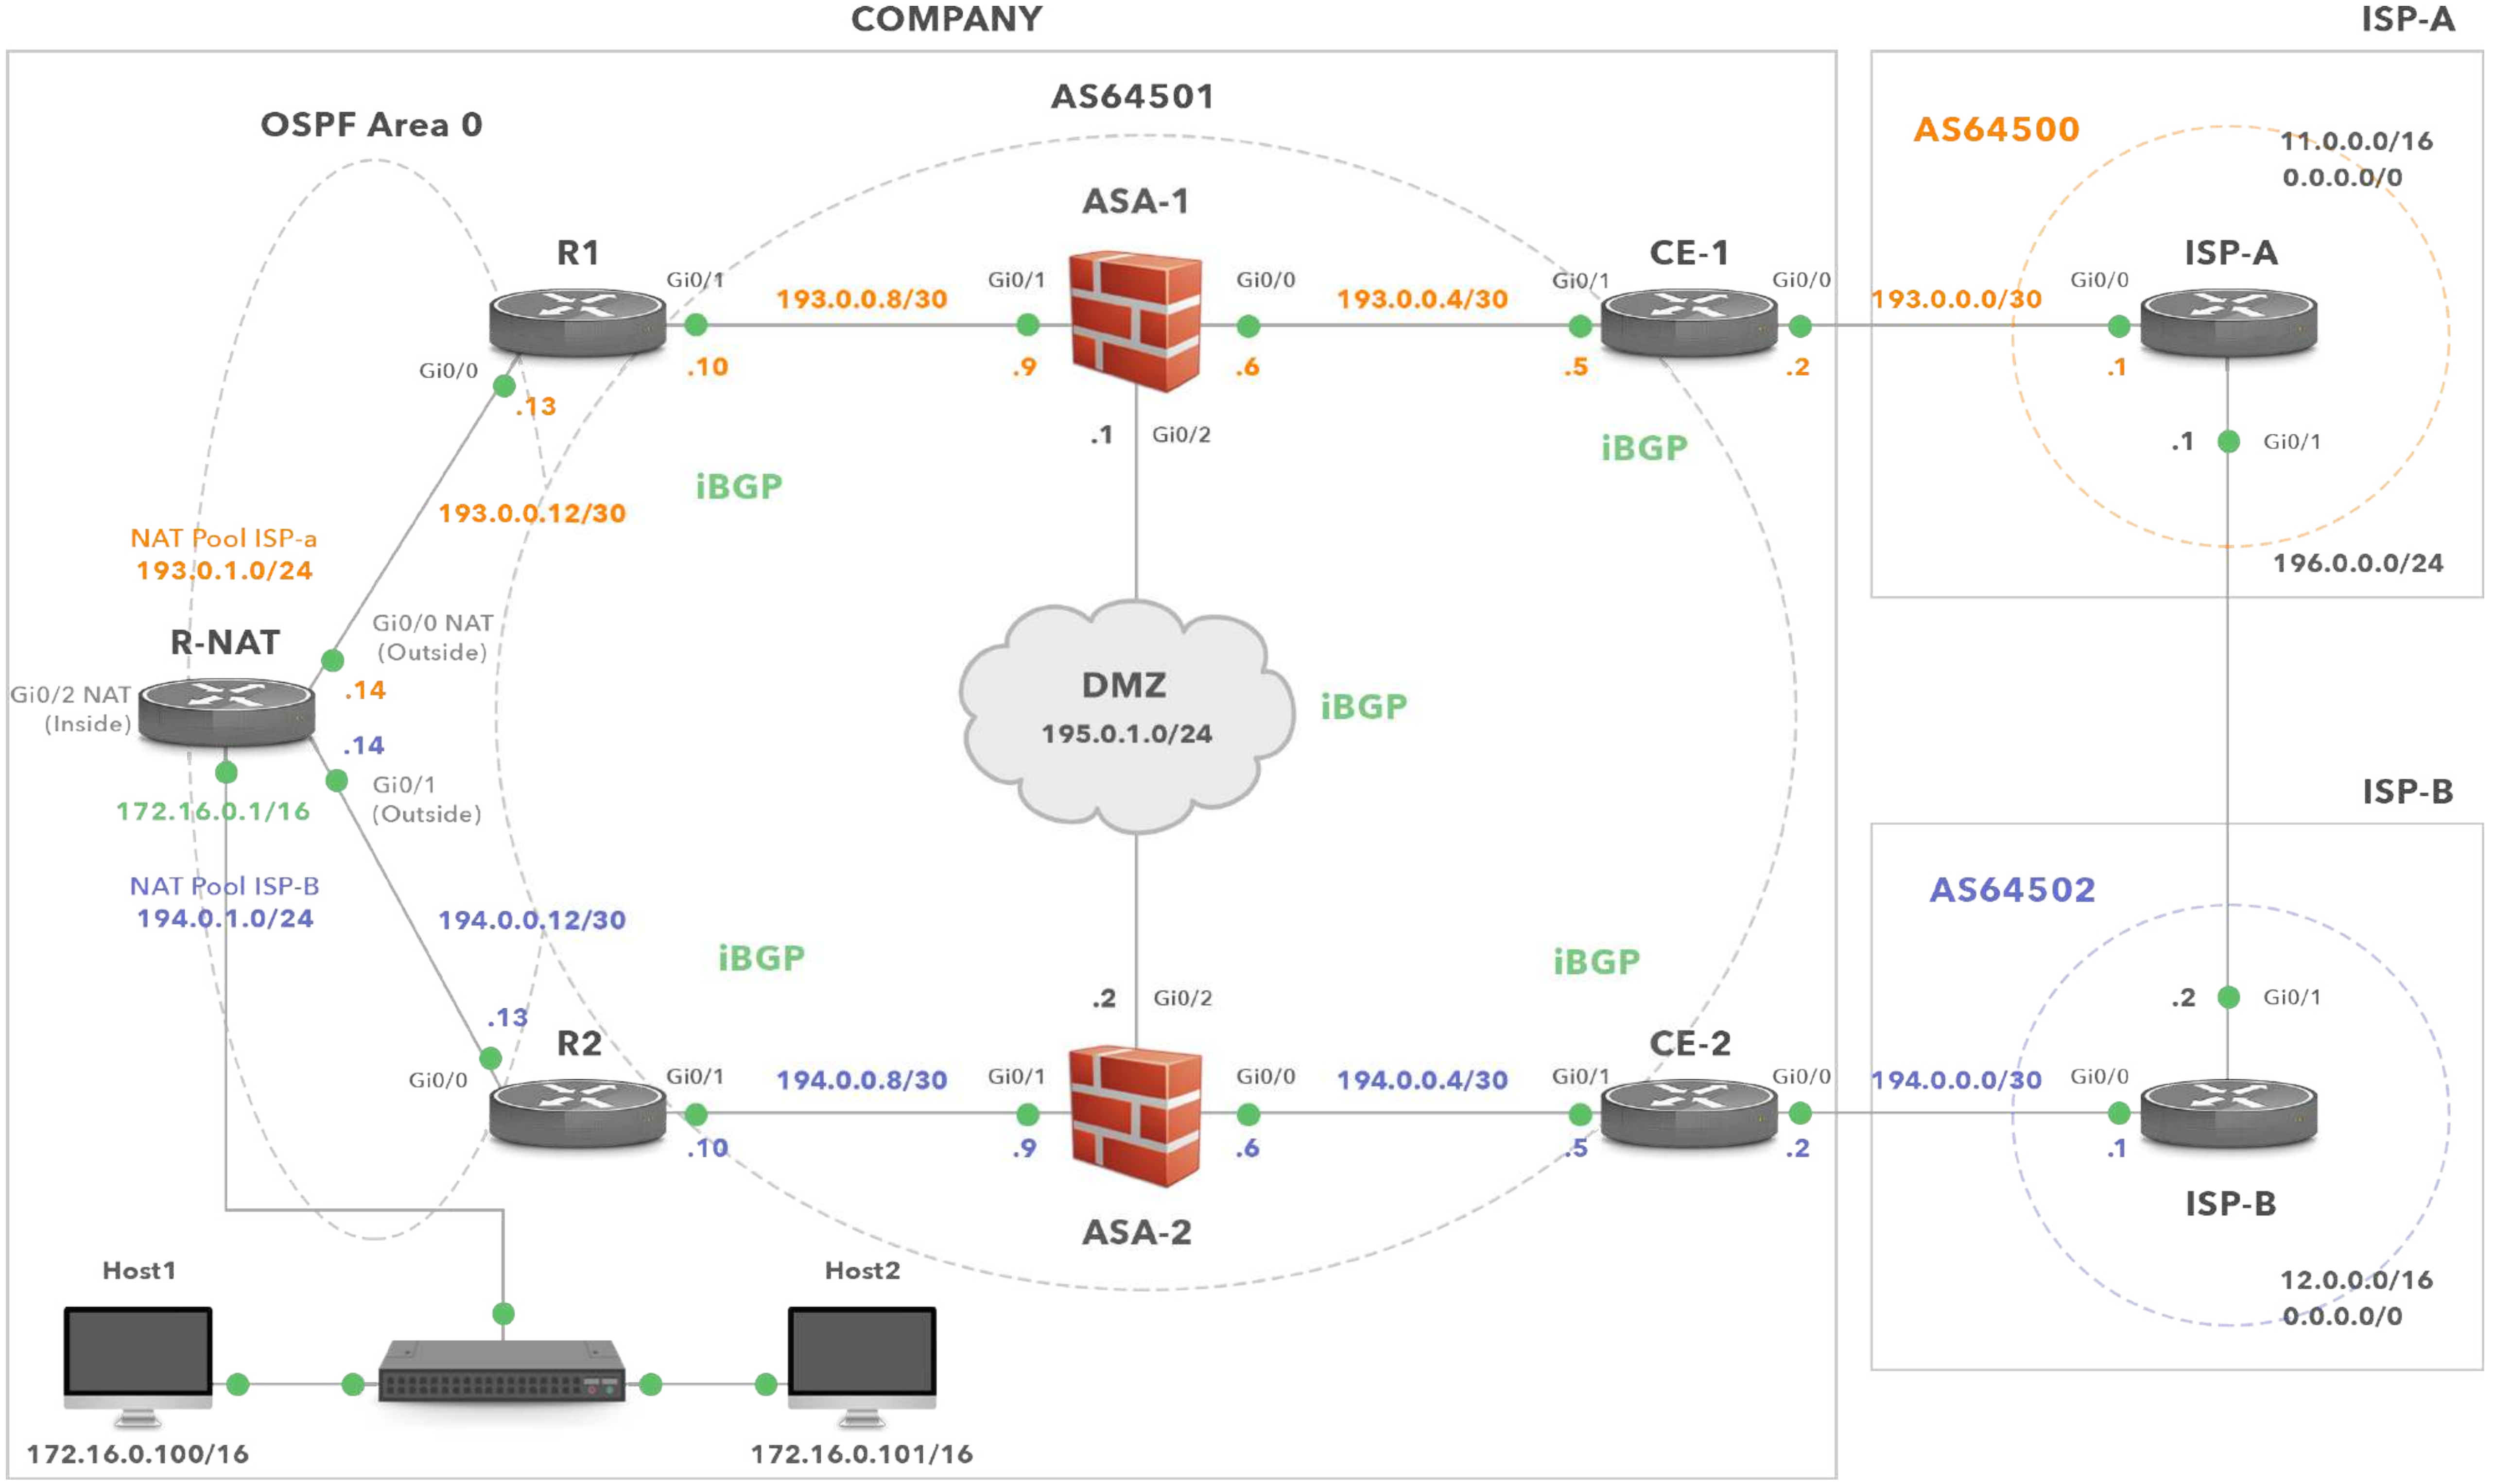

**Diagram 2** - *Enterprise Network (AS64501) is Multi-homed to ISP-A and ISP-B*

## Conclustion

In this tutorial we have covered the configuration of a multi-homed network design where a customer is connected to two different ISPs. The design ensures continuous connectivity, eliminating ISP as a single point of failure. When a primary connection to ISP-A is lost, customer traffic is routed via a backup link to ISP-B. We have also implemented the mechanism where source addresses of the internal users are translated to a different NAT address pool based on the selected path to the ISP.