

DE-CIX Academy: BGP - Multihoming

Links and Examples

*Where
networks
meet*



Notice of Liability

Despite careful checking of content, we accept no liability for the content of external links. Content on linked sites is exclusively the responsibility of the respective website operator.

Links visited during the webinar

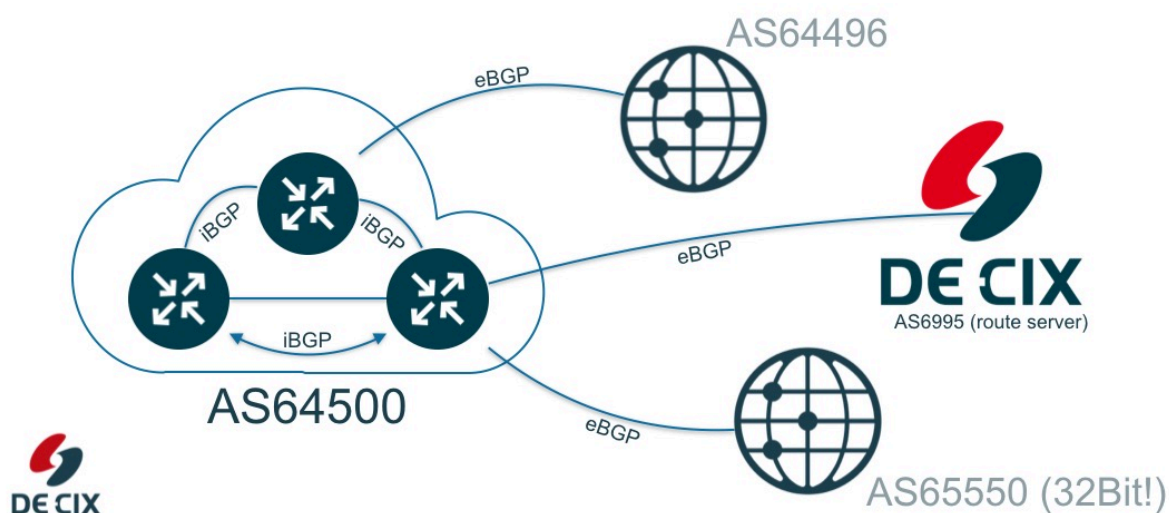
- BGP 01 - Introduction: [recorded webinar](#)
- RFCs about BGP:
 - [RFC4271](#)
 - see [5.1.5](#) for a definition of Local Preference
 - see [9.1](#) for the BGP routing algorithm
- BGP Route Selection Algorithm by vendor:
 - [Cisco](#)
 - [Juniper](#)
 - [Mikrotik](#)
 - [Nokia](#)
 - [BIRD](#)
 - [Quagga](#)

DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet


Example of a Network



→ We are using AS64500 as an example

- AS64500 has two upstreams
 - AS64496
 - AS65550 - note that this is a 32bit AS
- Also AS64500 has peering at DE-CIX Frankfurt

→ Prefixes are announced by AS517 via upstream and peering



DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

BGP Routing Algorithm

Bold lines were covered in this webinar.

1	NextHop reachable?	Continue if "yes"
2	Local Preference	higher wins
3	AS Path	shorter wins
4	Origin Type	IGP over EGP over Incomplete
5	MED	lower wins
6	eBGP, iBGP	eBGP wins
7	Exit	nearest wins
8	Age of route	older wins
9	Router ID	lower wins
10	Neighbor IP	lower wins

Local Preference is...

- a 32bit integer value (0-4294967295)
- Propagated via iBGP inside an Autonomous System
- Usually set using rules when receiving prefixes
 - According to your routing policy
- Typical values
 - 10000 (high value) for customer prefixes
 - 1000 (medium value) for prefixes received via peering
 - 100 (low value) for prefixes received via upstream
- Rules to adjust local preference can be as complex as your router software allows it to be.

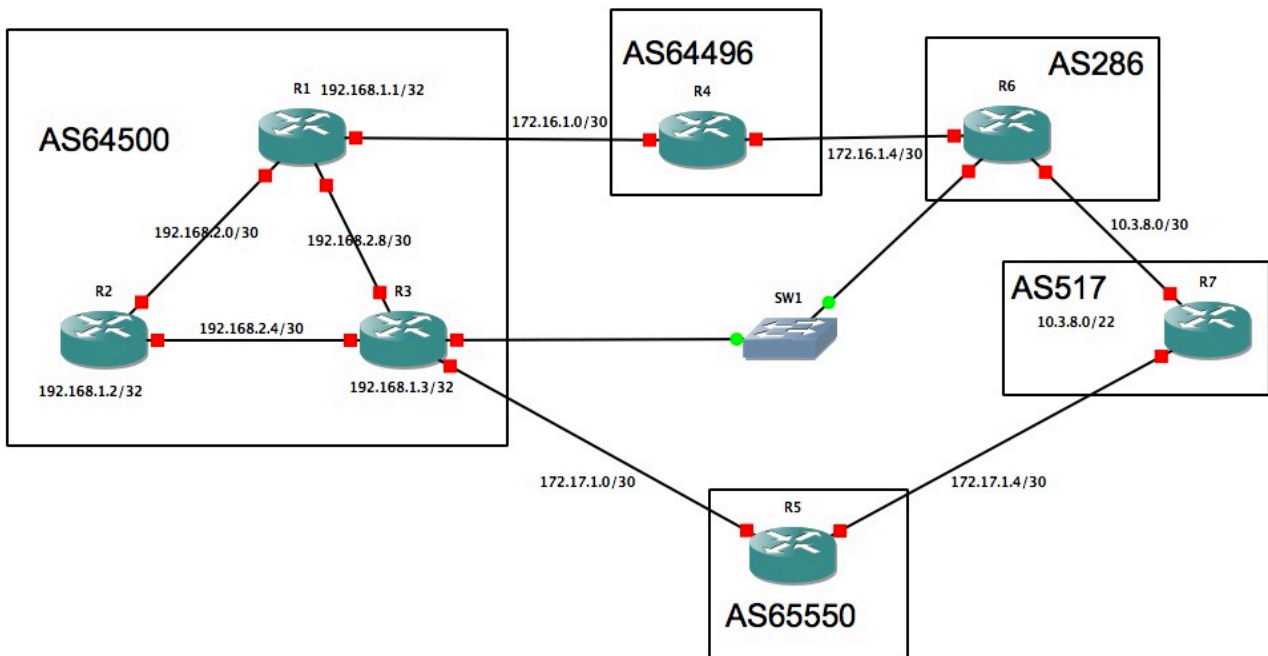
DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Appendix - Example Network Topology

Screenshot from GNS3



DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Appendix - Router configurations

Find below initial router configurations for GNS3 network emulator

In this example, all routers are Cisco 72xx style routers with 1-3 (Gig)Ethernet interfaces. Only relevant parts of the configurations are listed.

Router R2

```
!  
hostname R2  
ip bgp-community new-format  
!  
interface Loopback0  
 ip address 192.168.1.2 255.255.255.255  
!  
interface GigabitEthernet0/0  
 ip address 192.168.2.1 255.255.255.252  
!  
interface GigabitEthernet1/0  
 ip address 192.168.2.5 255.255.255.252  
!  
router ospf 64500  
 redistribute connected subnets route-map internal-only  
 network 192.168.2.0 0.0.0.3 area 0  
 network 192.168.2.4 0.0.0.3 area 0  
!  
router bgp 64500  
 neighbor internal peer-group  
 neighbor internal remote-as 64500  
 neighbor internal update-source Loopback0  
 neighbor internal next-hop-self  
 neighbor internal send-community both  
 neighbor 192.168.1.1 peer-group internal  
 neighbor 192.168.1.3 peer-group internal  
!  
ip prefix-list internal seq 5 permit 192.168.0.0/16 le 32  
!  
route-map internal-only permit 10  
 match ip address prefix-list internal  
!
```

DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Router R1

```
!  
hostname R1  
ip bgp-community new-format  
!  
interface Loopback0  
 ip address 192.168.1.1 255.255.255.255  
!  
interface GigabitEthernet0/0  
 ip address 192.168.2.2 255.255.255.252  
!  
interface GigabitEthernet1/0  
 ip address 172.16.1.2 255.255.255.252  
!  
interface GigabitEthernet2/0  
 ip address 192.168.2.9 255.255.255.252  
!  
router ospf 64500  
 redistribute connected subnets route-map internal-only  
 network 192.168.2.0 0.0.0.3 area 0  
 network 192.168.2.8 0.0.0.3 area 0  
!  
router bgp 64500  
 neighbor internal peer-group  
 neighbor internal remote-as 64500  
 neighbor internal update-source Loopback0  
 neighbor internal next-hop-self  
 neighbor upstream peer-group  
 neighbor upstream route-map upstream-in in  
 neighbor upstream route-map upstream-out out  
 neighbor 172.16.1.1 remote-as 64496  
 neighbor 172.16.1.1 peer-group upstream  
 neighbor 192.168.1.2 peer-group internal  
 neighbor 192.168.1.3 peer-group internal  
!  
ip prefix-list internal seq 5 permit 192.168.0.0/16 le 32  
route-map internal-only permit 10  
 match ip address prefix-list internal  
!  
route-map upstream-in permit 100  
 set local-preference 100  
!  
route-map upstream-out permit 100  
!  
!
```

DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Router R3

```
!  
hostname R3  
ip bgp-community new-format  
!  
interface Loopback0  
 ip address 192.168.1.3 255.255.255.255  
!  
interface GigabitEthernet0/0  
 ip address 172.17.1.2 255.255.255.252  
!  
interface GigabitEthernet1/0  
 ip address 192.168.2.6 255.255.255.252  
!  
interface GigabitEthernet2/0  
 ip address 192.168.2.10 255.255.255.252  
!  
interface GigabitEthernet3/0  
 description to emulated DE-CIX  
 ip address 80.81.194.110 255.255.252.0  
!  
router ospf 64500  
 redistribute connected subnets route-map internal-only  
 network 192.168.2.4 0.0.0.3 area 0  
 network 192.168.2.8 0.0.0.3 area 0  
!  
router bgp 64500  
 neighbor internal peer-group  
 neighbor internal remote-as 64500  
 neighbor internal update-source Loopback0  
 neighbor internal next-hop-self  
 neighbor internal send-community both  
 neighbor upstream peer-group  
 neighbor upstream send-community both  
 neighbor upstream soft-reconfiguration inbound  
 neighbor upstream route-map upstream-in in  
 neighbor upstream route-map upstream-out out  
 neighbor peering peer-group  
 neighbor peering send-community both  
 neighbor peering soft-reconfiguration inbound  
 neighbor peering route-map peering-out out  
 neighbor peering route-map peering-in in  
 neighbor 80.81.192.86 remote-as 286  
 neighbor 80.81.192.86 peer-group peering  
 neighbor 172.17.1.1 remote-as 65550  
 neighbor 172.17.1.1 peer-group upstream  
 neighbor 192.168.1.1 peer-group internal  
 neighbor 192.168.1.2 peer-group internal  
!
```

DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Router R3 (continued from last page)

```
!  
ip prefix-list internal seq 5 permit 192.168.0.0/16 le 32  
route-map internal-only permit 10  
  match ip address prefix-list internal  
!  
route-map upstream-in permit 100  
  set local-preference 100  
!  
route-map upstream-out permit 100  
!  
route-map peering-in permit 100  
  set local-preference 1000  
!  
route-map peering-out permit 100
```

Router R4 (AS64496)

```
!  
hostname R4  
ip bgp-community new-format  
!  
interface GigabitEthernet0/0  
  ip address 172.16.1.5 255.255.255.252  
!  
interface GigabitEthernet1/0  
  ip address 172.16.1.1 255.255.255.252  
!  
router bgp 64496  
  neighbor upstream peer-group  
  neighbor upstream send-community both  
  neighbor upstream soft-reconfiguration inbound  
  neighbor upstream route-map upstream-in in  
  neighbor upstream route-map upstream-out out  
  neighbor 172.16.1.2 remote-as 64500  
  neighbor 172.16.1.6 remote-as 286  
  neighbor 172.16.1.6 peer-group upstream  
!  
route-map upstream-in permit 100  
  set local-preference 100  
!  
route-map upstream-out permit 100  
!
```


DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Router R5 (AS5550)

```
!  
hostname R5  
ip bgp-community new-format  
!  
interface GigabitEthernet0/0  
 ip address 172.17.1.1 255.255.255.252  
!  
interface GigabitEthernet1/0  
 ip address 172.17.1.5 255.255.255.252  
!  
router bgp 65550  
neighbor customer peer-group  
 neighbor customer send-community both  
 neighbor customer soft-reconfiguration inbound  
 neighbor customer route-map customer-in in  
 neighbor customer route-map customer-out out  
 neighbor 172.17.1.2 remote-as 64500  
 neighbor 172.17.1.2 peer-group customer  
 neighbor 172.17.1.6 remote-as 517  
 neighbor 172.17.1.6 peer-group customer  
!  
route-map customer-in permit 100  
 set local-preference 10000  
!  
route-map customer-out permit 100  
!  
!
```

DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Router R6 (AS286)

```
!  
hostname R6  
ip bgp-community new-format  
!  
interface GigabitEthernet0/0  
 ip address 172.16.1.6 255.255.255.252  
!  
interface GigabitEthernet1/0  
 ip address 10.3.8.2 255.255.255.252  
!  
interface GigabitEthernet2/0  
 description emulated DE-CIX  
 ip address 80.81.192.86 255.255.252.0  
!  
router bgp 286  
 neighbor customer peer-group  
 neighbor customer send-community both  
 neighbor customer soft-reconfiguration inbound  
 neighbor customer route-map customer-in in  
 neighbor customer route-map customer-out out  
 neighbor peering peer-group  
 neighbor peering send-community both  
 neighbor peering soft-reconfiguration inbound  
 neighbor peering route-map peering-in in  
 neighbor peering route-map peering-out out  
 neighbor 10.3.8.1 remote-as 517  
 neighbor 10.3.8.1 peer-group customer  
 neighbor 80.81.194.110 remote-as 64500  
 neighbor 80.81.194.110 peer-group peering  
 neighbor 172.16.1.5 remote-as 64496  
 neighbor 172.16.1.5 peer-group customer  
!  
route-map customer-in permit 100  
 set local-preference 10000  
!  
route-map customer-out permit 100  
!  
route-map peering-out permit 100  
!  
route-map peering-in permit 100  
 set local-preference 1000  
!
```

DE-CIX Academy: BGP - Multihoming

Links and Examples

Where
networks
meet

Router R7 (AS517)

```
!  
hostname R7  
ip bgp-community new-format  
!  
interface GigabitEthernet0/0  
 ip address 10.3.8.1 255.255.255.252  
!  
interface GigabitEthernet1/0  
 ip address 172.17.1.6 255.255.255.252  
!  
router bgp 517  
 network 10.3.8.0 mask 255.255.252.0  
 neighbor upstream peer-group  
 neighbor upstream send-community both  
 neighbor upstream soft-reconfiguration inbound  
 neighbor upstream route-map upstream-in in  
 neighbor upstream route-map upstream-out out  
 neighbor 10.3.8.2 remote-as 286  
 neighbor 10.3.8.2 peer-group upstream  
 neighbor 172.17.1.5 remote-as 65550  
 neighbor 172.17.1.5 peer-group upstream  
!  
ip as-path access-list 1 permit ^$  
!  
ip route 10.3.8.0 255.255.252.0 Null0  
!  
route-map upstream-in permit 100  
 set local-preference 100  
!  
route-map upstream-out permit 100  
 match as-path 1  
!
```