DE-CIX Blackholing Service

How to mitigate effects
of Distributed Denial of Service
(DDoS) attacks

DE-CIX

What is Blackholing?

- Blackholing effectively means diverting the flow of data to a specific IP next-hop (Blackhole), where the traffic is then discarded
- As a result no traffic reaches the original destination, meaning peering links, networks, and hosts located within the blackholed prefix are protected
- Therefore Blackholing is an effective way of mitigating the effects of Distributed Denial of Service (DDoS) attacks

How does the Blackholing service work?

- ***Default case – Blackholing is not used***
  - Customers advertise their IP prefix(es) with the next-hop IP of their advertising router. DE-CIX Route Servers accept the following prefix lengths:
    - IPv4: /8 ≤ prefix length ≤ /24
    - IPv6: /19 ≤ prefix length ≤ /48
- ***Blackholing case: To protect against a massive DDoS attack***
  - Customers advertise their IP prefix(es) tagged with the BGP BLACKHOLE Community (65535:666). Accepted prefix lengths are:
    - IPv4: /8 ≤ prefix length ≤ /32 (if and only if BLACKHOLE is set)
    - IPv6: /19 ≤ prefix length ≤ /128 (if and only if BLACKHOLE is set)
  - Prefix validation (RIR filtering) is applied as usual, to prevent unauthorized Blackholing

How does the Blackholing service work?

- ➔ L2 filtering
  - If the BGP BLACKHOLE Community is set, the DE-CIX Route Servers rewrite the next-hop of the advertised IP prefix(es) to the address of the Blackhole next-hop (BN)
  - BNs has a unique MAC address (determined by ARP/NDP)
  - All frames with destination MAC address belonging to the BN are ingress filtered by a L2 ACL applied on all customer ports of the switching platform
- ➔ As a result, all traffic to the blackholed IP prefix(es) is discarded on the switching infrastructure already, hence the victim's resources are protected

DE-CIX‘s Blackholing service is available at:

DE CIX

Where networks meet

www.de-cix.net

Example

DE CIX

Where networks meet

www.de-cix.net

Default case – Blackholing is not used

- → AS 64511 announces IP prefixes
  - directly to other peers (here AS 64501)
  - via the Route Servers, which re-distribute the prefixes to other peers
  - Other ASs also peering with the Route Servers:

    AS 64502, AS 64503, AS 64504
- → The other ASs learn the BGP next-hop for the announced IP prefixes
  - IP prefix is chosen as best-path
- → The corresponding next-hop MAC is learned via ARP/NDP

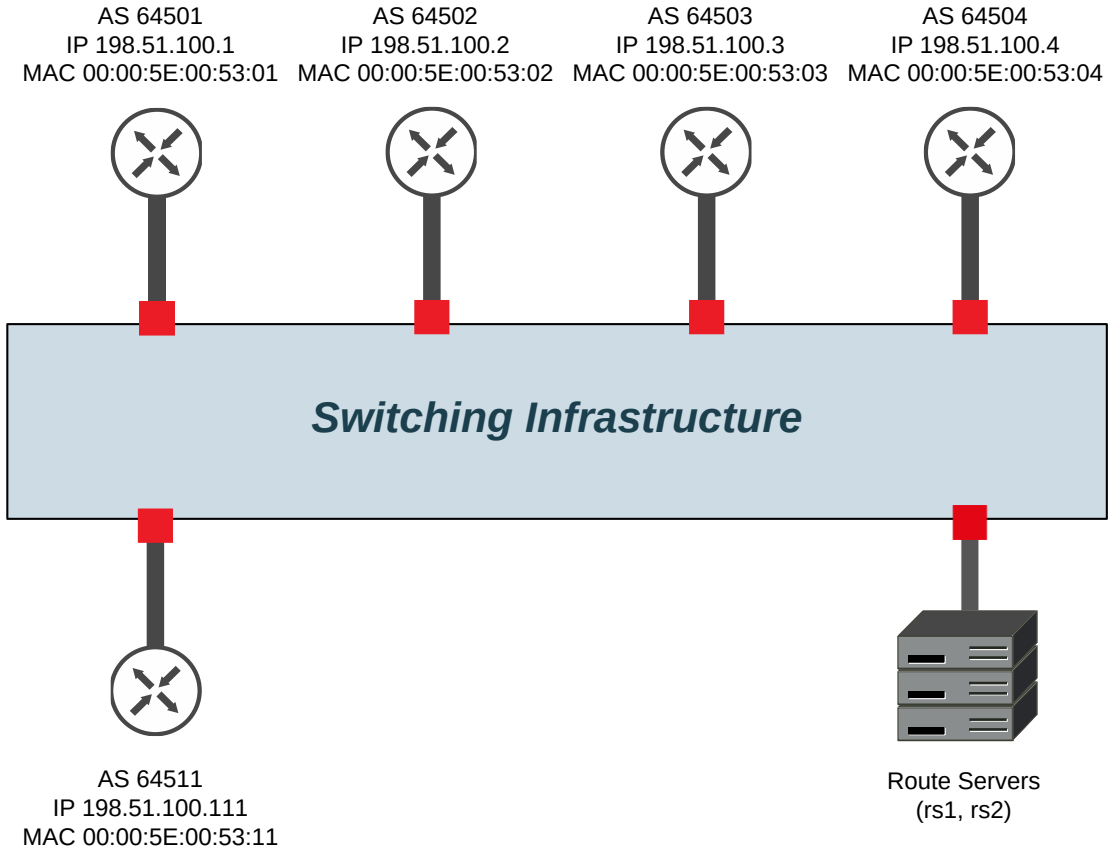
AS 64501
IP 198.51.100.1
MAC 00:00:5E:00:53:01
AS 64502
IP 198.51.100.2
MAC 00:00:5E:00:53:02
AS 64503
IP 198.51.100.3
MAC 00:00:5E:00:53:03
AS 64504
IP 198.51.100.4
MAC 00:00:5E:00:53:04
Switching Infrastructure
AS 64511
IP 198.51.100.111
MAC 00:00:5E:00:53:11
Route Servers
(rs1, rs2)

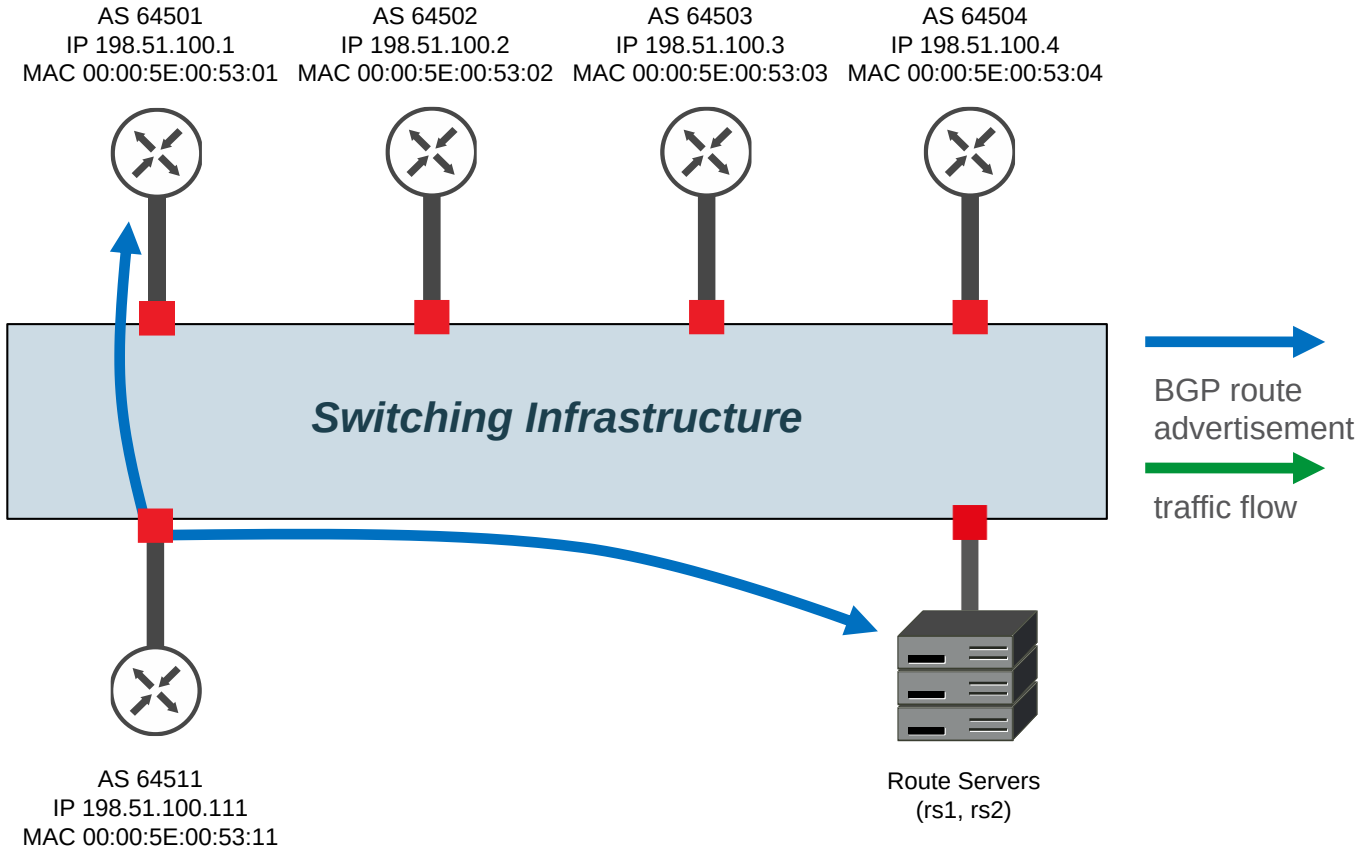
AS 64501
IP 198.51.100.1
MAC 00:00:5E:00:53:01
AS 64502
IP 198.51.100.2
MAC 00:00:5E:00:53:02
AS 64503
IP 198.51.100.3
MAC 00:00:5E:00:53:03
AS 64504
IP 198.51.100.4
MAC 00:00:5E:00:53:04
Switching Infrastructure
BGP route advertisement
traffic flow
AS 64511
IP 198.51.100.111
MAC 00:00:5E:00:53:11
Route Servers
(rs1, rs2)

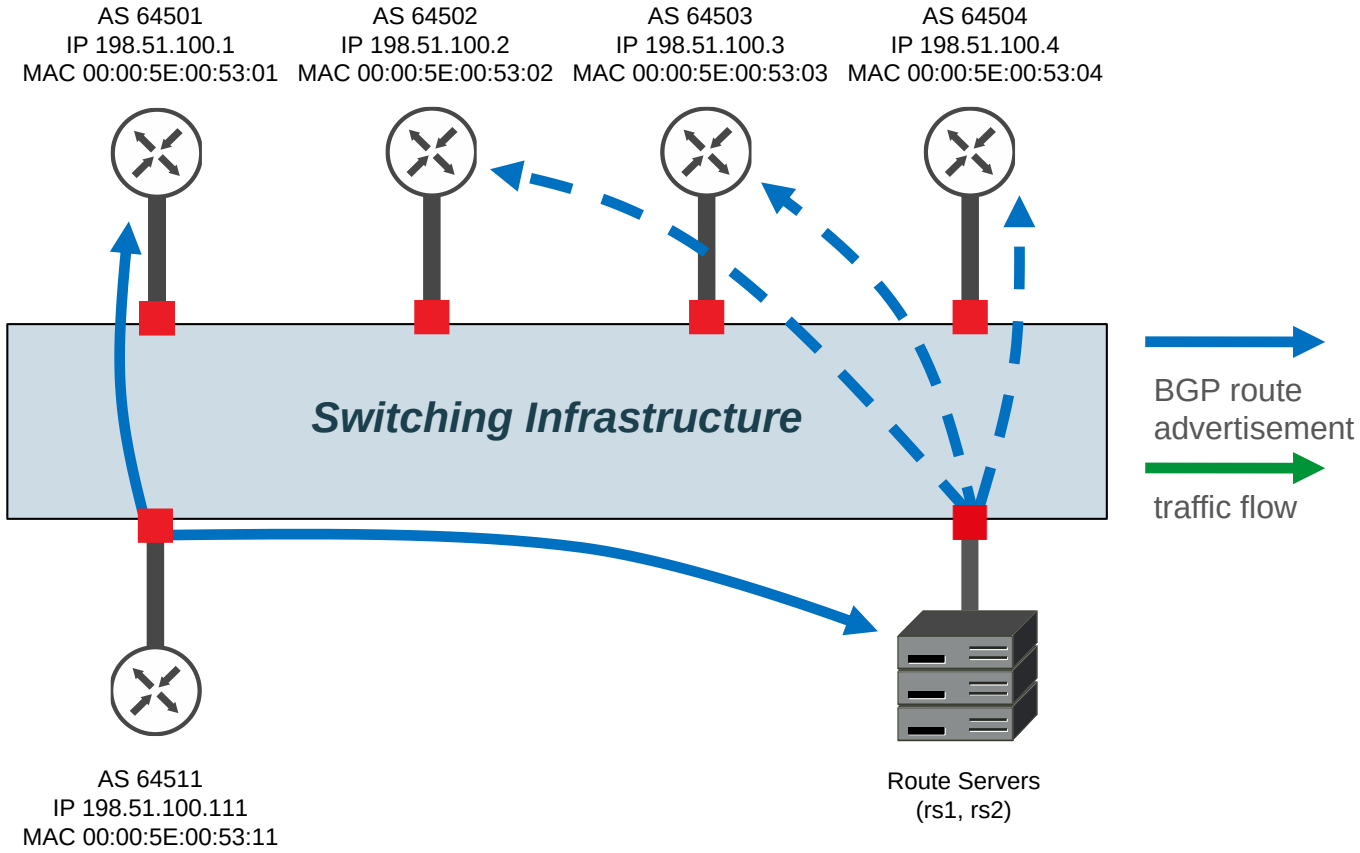
AS 64501
IP 198.51.100.1
MAC 00:00:5E:00:53:01
AS 64502
IP 198.51.100.2
MAC 00:00:5E:00:53:02
AS 64503
IP 198.51.100.3
MAC 00:00:5E:00:53:03
AS 64504
IP 198.51.100.4
MAC 00:00:5E:00:53:04
Switching Infrastructure
BGP route advertisement
traffic flow
AS 64511
IP 198.51.100.111
MAC 00:00:5E:00:53:11
Route Servers
(rs1, rs2)

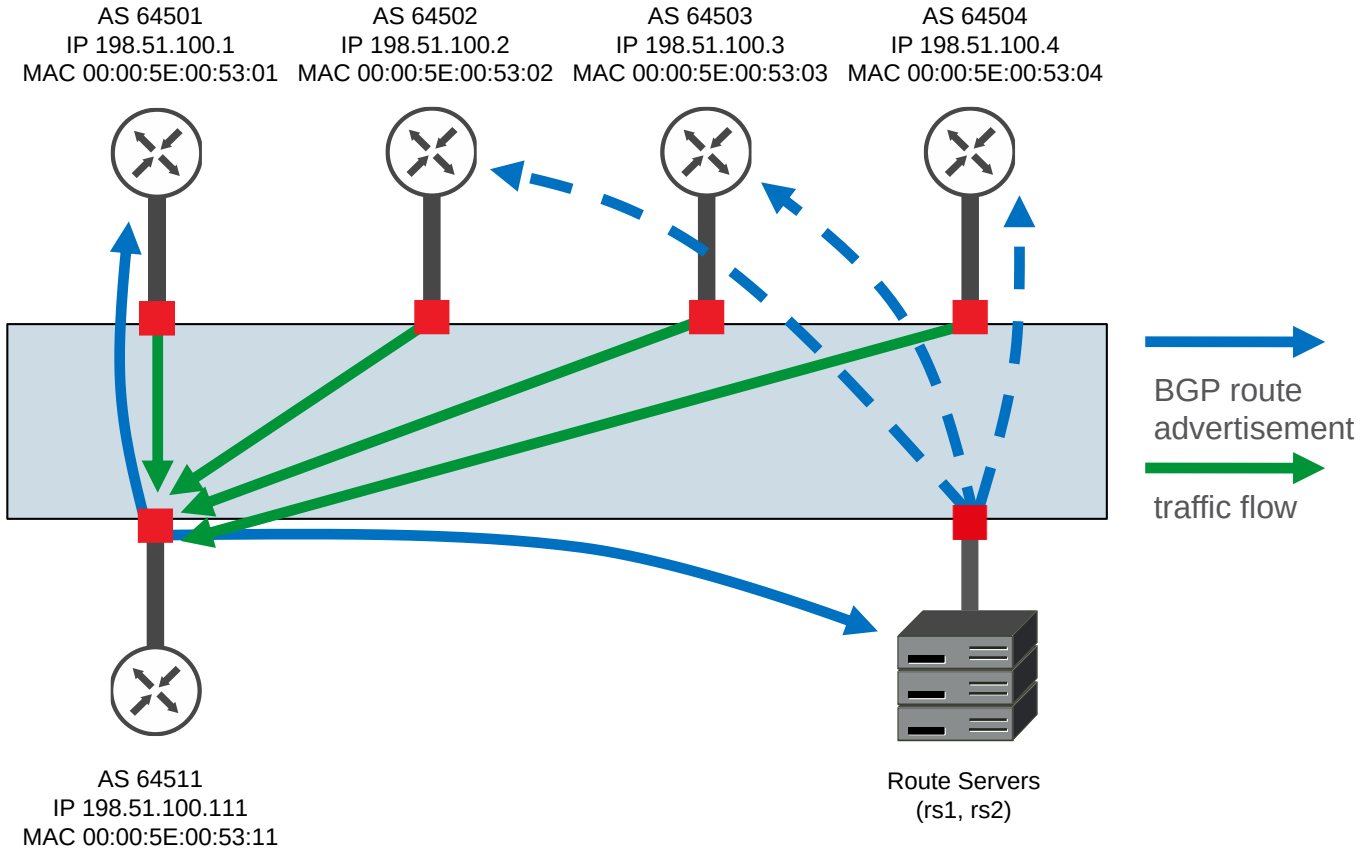
AS 64501
IP 198.51.100.1
MAC 00:00:5E:00:53:01
AS 64502
IP 198.51.100.2
MAC 00:00:5E:00:53:02
AS 64503
IP 198.51.100.3
MAC 00:00:5E:00:53:03
AS 64504
IP 198.51.100.4
MAC 00:00:5E:00:53:04
BGP route advertisement
traffic flow
AS 64511
IP 198.51.100.111
MAC 00:00:5E:00:53:11
Route Servers
(rs1, rs2)

Blackholing case: To protect against a massive DDoS attack

- A destination within the IP prefix 203.0.113.0/24 of AS 64511 is a target of a massive DDoS attack
- AS 64511 also announces other IP prefixes than the attacked one
- AS 64501, AS 64502 and AS 64503 originate traffic, which is part of the DDoS attack
- AS 64504 originates legitimate traffic
- AS 64501 directly peers with AS 64511
- AS 64502, AS 65403 and AS 64504 only see AS 64511's IP prefixes via the Route Servers

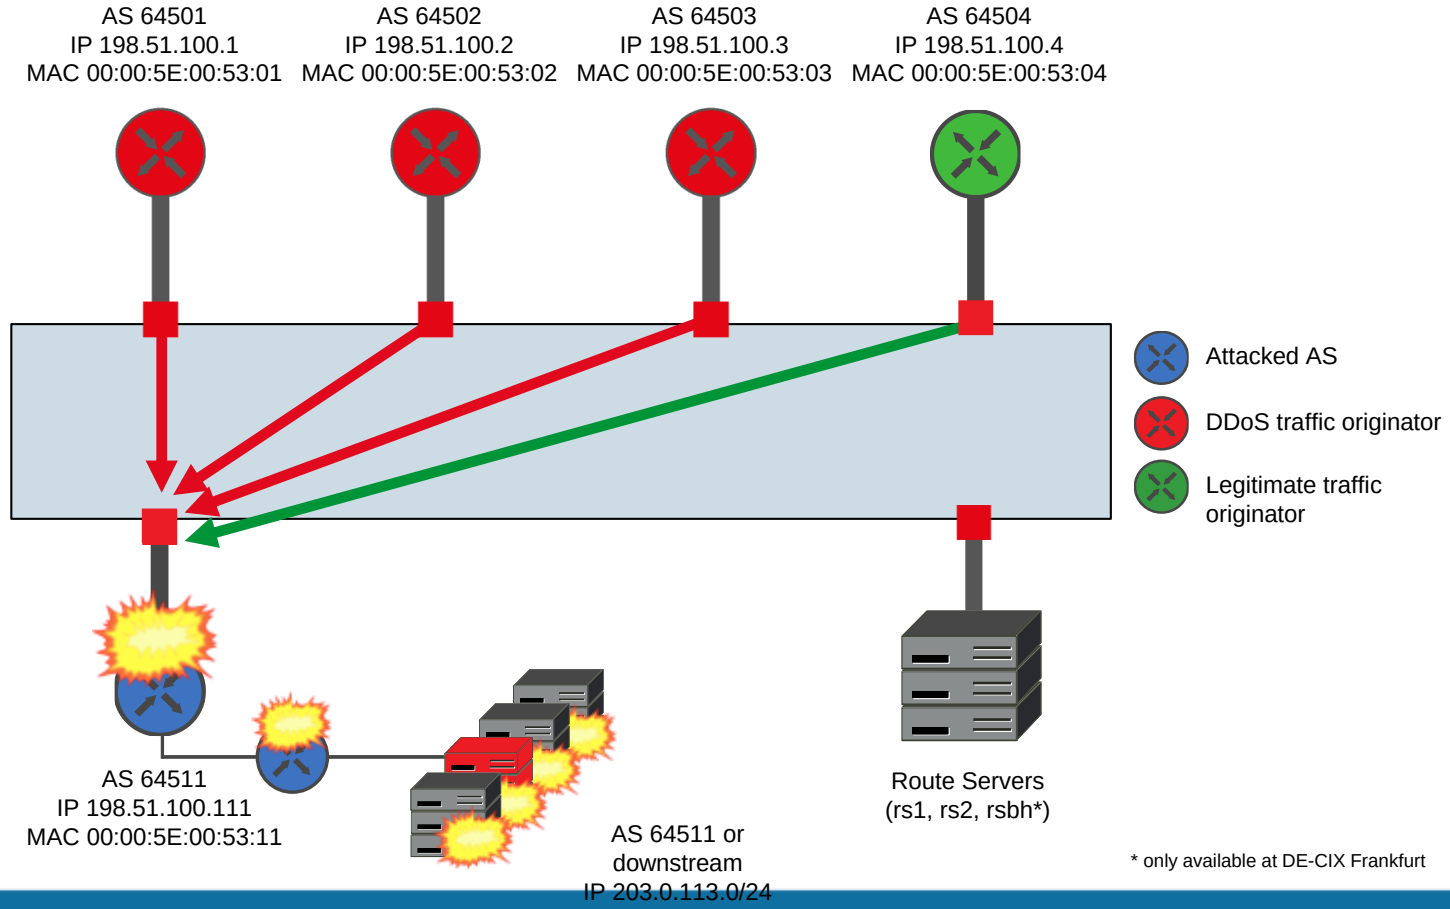

* only available at DE-CIX Frankfurt

Considerations

- ➔ The reachability of the attacked IP prefix (203.0.113.0/24) behind AS 64511 is limited as the peering link, the router and the network is congested
- ➔ Collateral damage on other resources (e.g. reachability of IP prefixes) might occur
- ➔ AS 64504 has a degraded reachability of 203.0.113.0/24, even it is not attacked directly

Solution: Blackholing

- AS 64511 announces the attacked IP prefix(es) to be blackholed by using the BGP BLACKHOLE Community (65535:666)
  - DE-CIX Route Server allow control over the re-distribution process of blackholed IP prefixes by utilizing BGP communities
    - Example: To order the Route Servers to advertise blackholed prefixes to all peers except AS 64504 the following BGP communities must be set: (6695:6695) (0:64504)
      6695 is the ASN for DE-CIX Frankfurt and must be adjusted to the appropriate value for other DE-CIX locations
- DE-CIX provides ARP reply or NDP for BN's MAC
- All Frames with destination MAC address belonging to the BN are filtered within each local Apollon switch

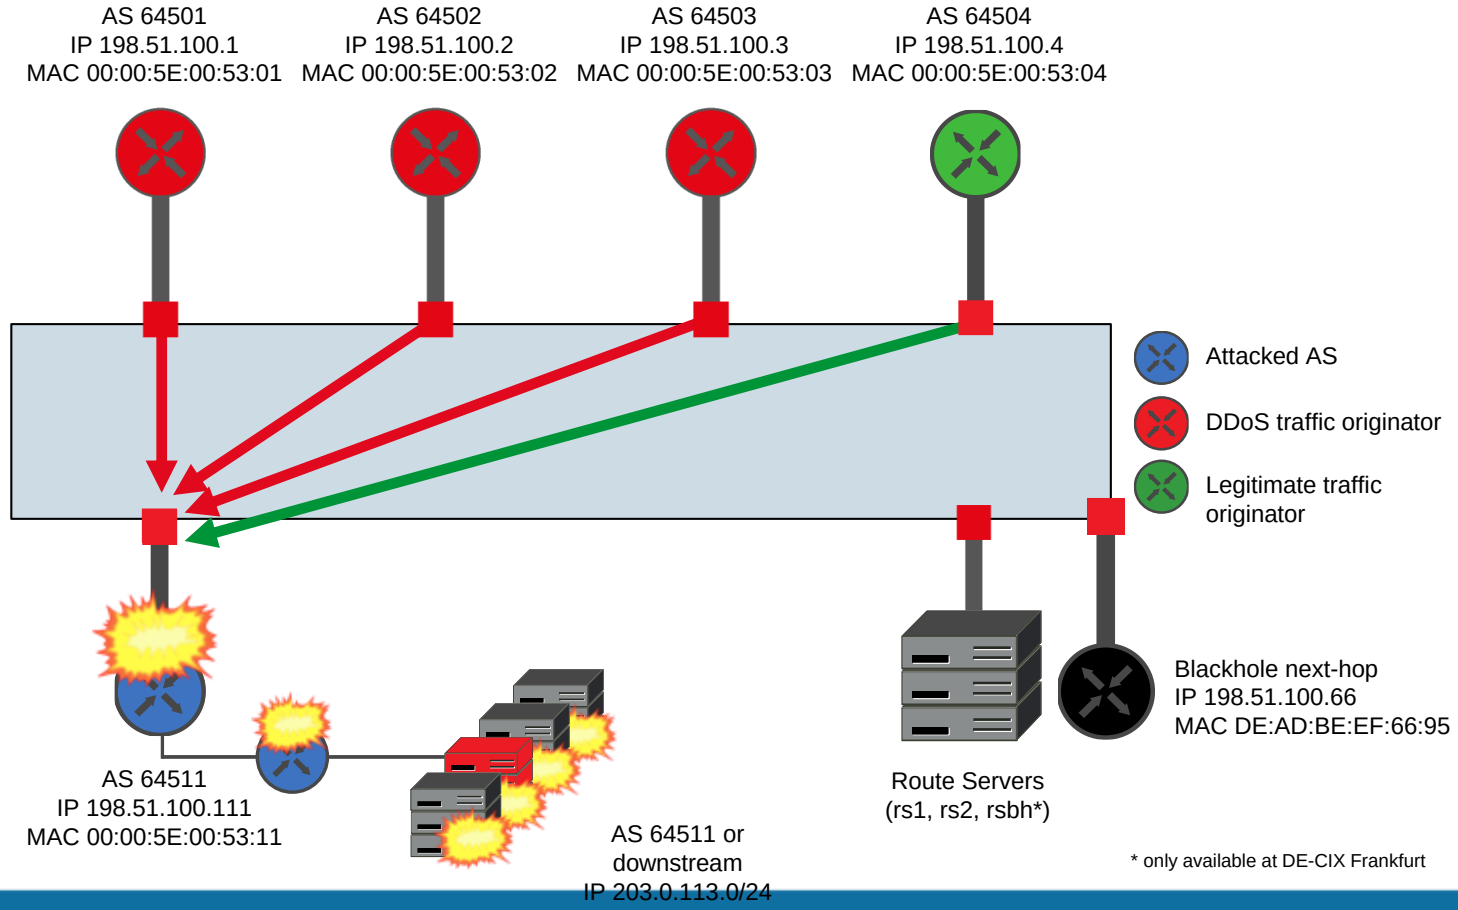

* only available at DE-CIX Frankfurt

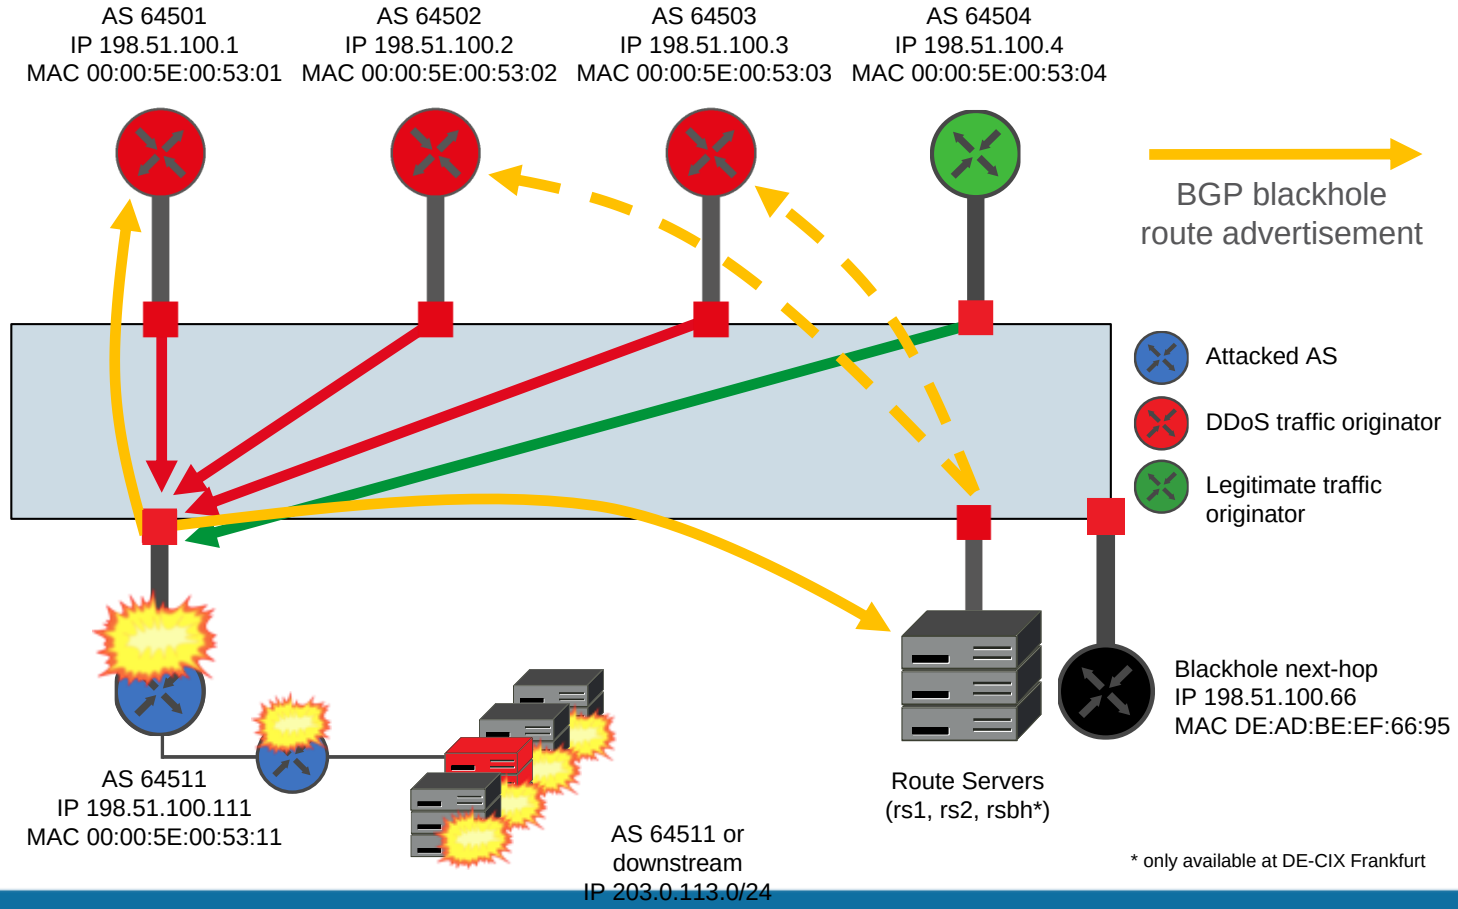

* only available at DE-CIX Frankfurt

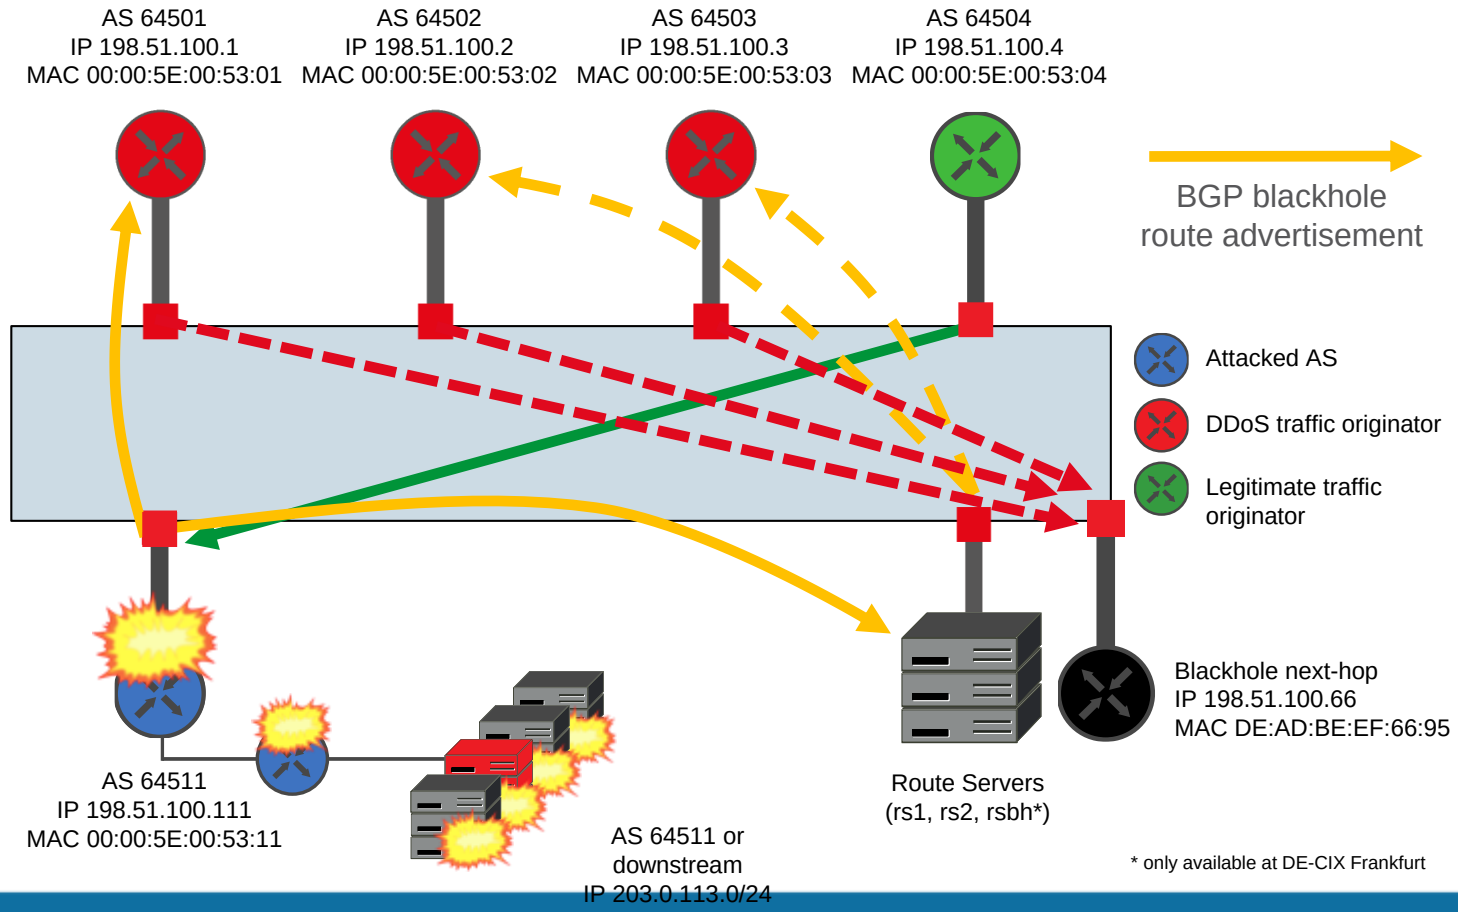

* only available at DE-CIX Frankfurt

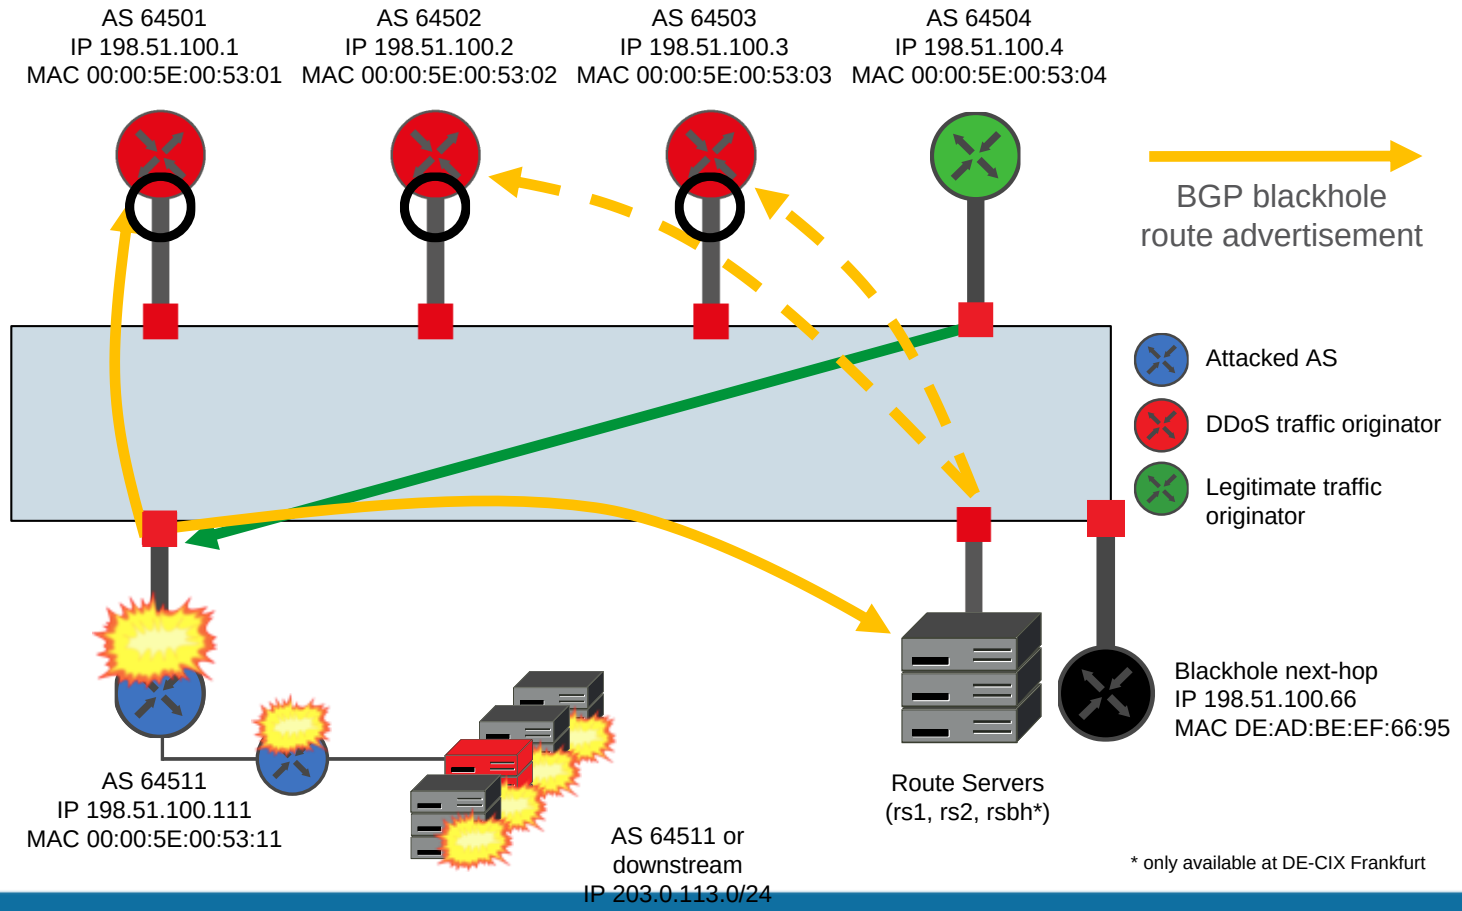

* only available at DE-CIX Frankfurt

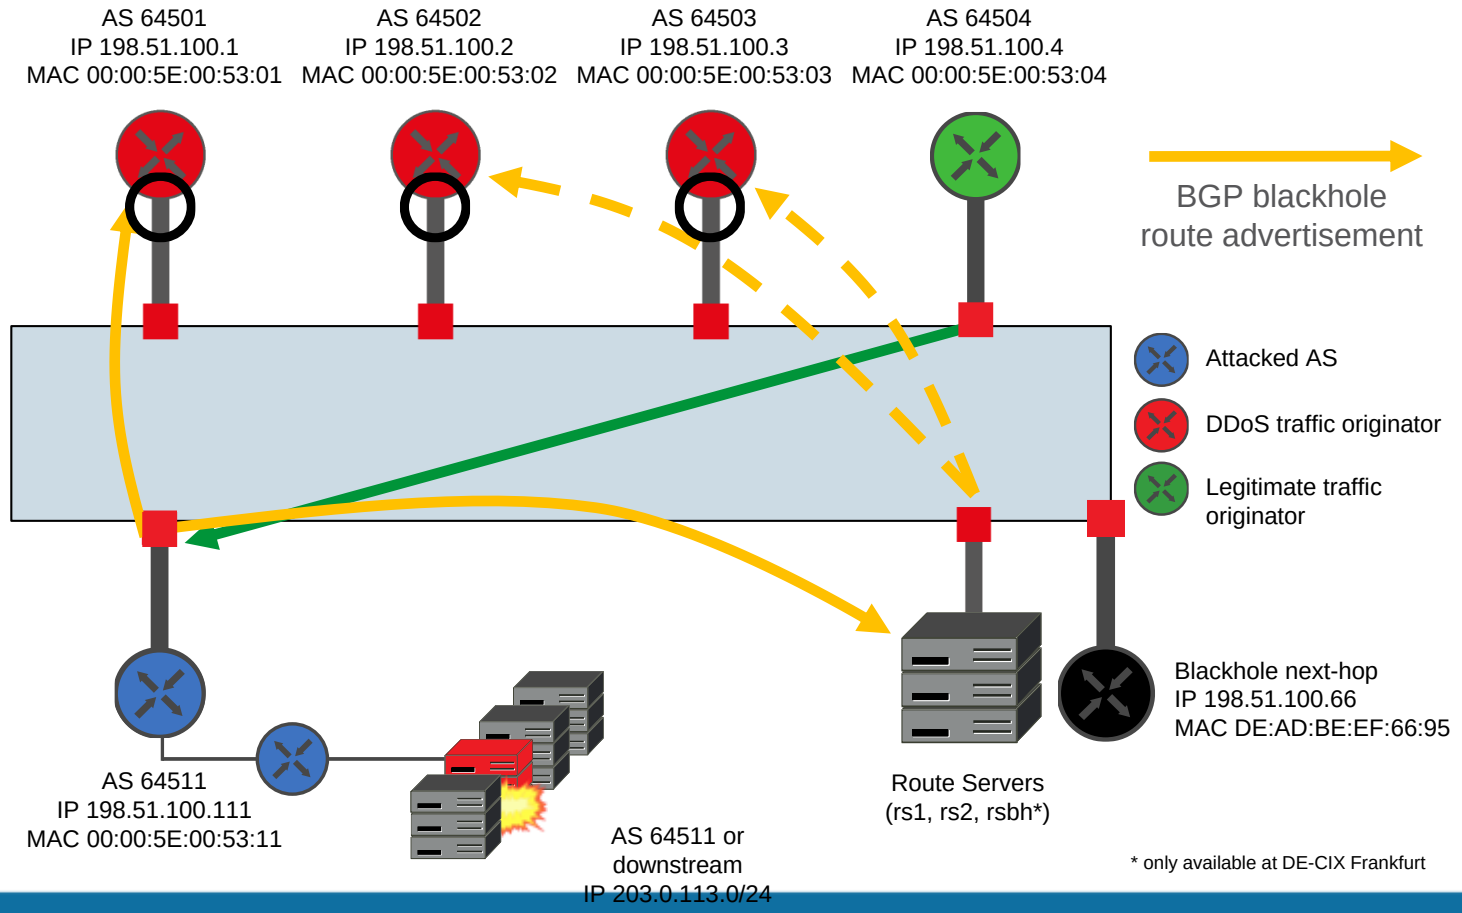

* only available at DE-CIX Frankfurt

Example summary

- AS 64511 selectively announces the attacked IP prefix with the BGP BLACKHOLE Community → (6695:6695) (0:64504)
- The Route Servers rewrite the BGP next-hop to the pre-defined IP of the Blackhole next-hop
- All peers which select this new IP prefix as best-path, learn the BN's MAC address via ARP/ND provided by DE-CIX
- Traffic destined to the BN's MAC is dropped ingress via L2 ACL
- AS 64511 has a chance to selectively blackhole traffic

DE-CIX FRA: Advertise to which route server?

- In addition to rs1 and rs2, there's a Route Server dedicated to Blackholing announcements (rsbh) available at DE-CIX Frankfurt
- All Route Servers (rs1, rs2, rsbh) support
  - Blackholing of IP prefixes with BGP BLACKHOLE Community set
  - Exporting those IP prefixes with Blackhole next-hop and the well-known BGP community NO-EXPORT and BLACKHOLE
- Benefits of using the Blackholing Route Server
  - All your blackholed prefixes on one session; no need to modify your sessions to rs1/rs2 (alter prefix lists, route maps, etc)
  - No need to accept > /24 or > /48 on your existing sessions to rs1/rs2 (just on the one to the Blackholing Route Server)
- Recommended usage: Advertise prefixes to be blackholed only to rsbh
  - We require you to set the BGP BLACKHOLE Community on rsbh as well, otherwise the IP prefixes are not accepted by rsbh
- You'll receive blackholed prefixes advertised to rs1/rs2 only from the blackholing Route Server as well (and vice versa)

Important notes

- ➔ Traffic from all of your peers to the blackholed IP prefix(es) is discarded
  - Including the legitimate traffic
  - Solution: Advertise the prefix(es) to be blackholed only to certain ASNs (which are originating DDoS traffic) by using the appropriate DE-CIX Route Server control BGP communities

- ➔ Traffic towards all hosts within the blackholed IP prefix is discarded
  - Including any hosts not under DDoS attack
  - Solution: You can blackhole prefixes as specific as /32 (IPv4) or /128 (IPv6)

DE CIX

Thank you!

Any questions? Contact us!

DE-CIX Management GmbH | Lindleystr. 12 | 60314 Frankfurt | Germany
Phone + 49 69 1730 902 0 | sales@de-cix.net | www.de-cix.net