DE·CIX

Where networks meet

DDoS Mitigation:
Customer-Triggered Blackholing @ DE-CIX

Internet Security Days 2012

09/2012

Petr Marciniak
petr.marciniak@de-cix.net

Where networks meet DE·CIX

Who we are?

- *DE-CIX is the „connectivity cloud“ and the service called „Peering“. Peering is the short cut for IP packets between origin and destination.*
- *Benefits of peering at DE-CIX are*
  - *Routing around congested Internet paths*
  - *Reducing latency*
  - *Reducing transit costs*
  - *Control over IP routing*
  - *Better end-user experience*
  - *Enjoying marketing benefits*
- *DE-CIX is located in Frankfurt, Germany and is the worlds largest Internet Exchange by peak traffic (1935 Gbit/s)*

DE·CIX Where networks meet

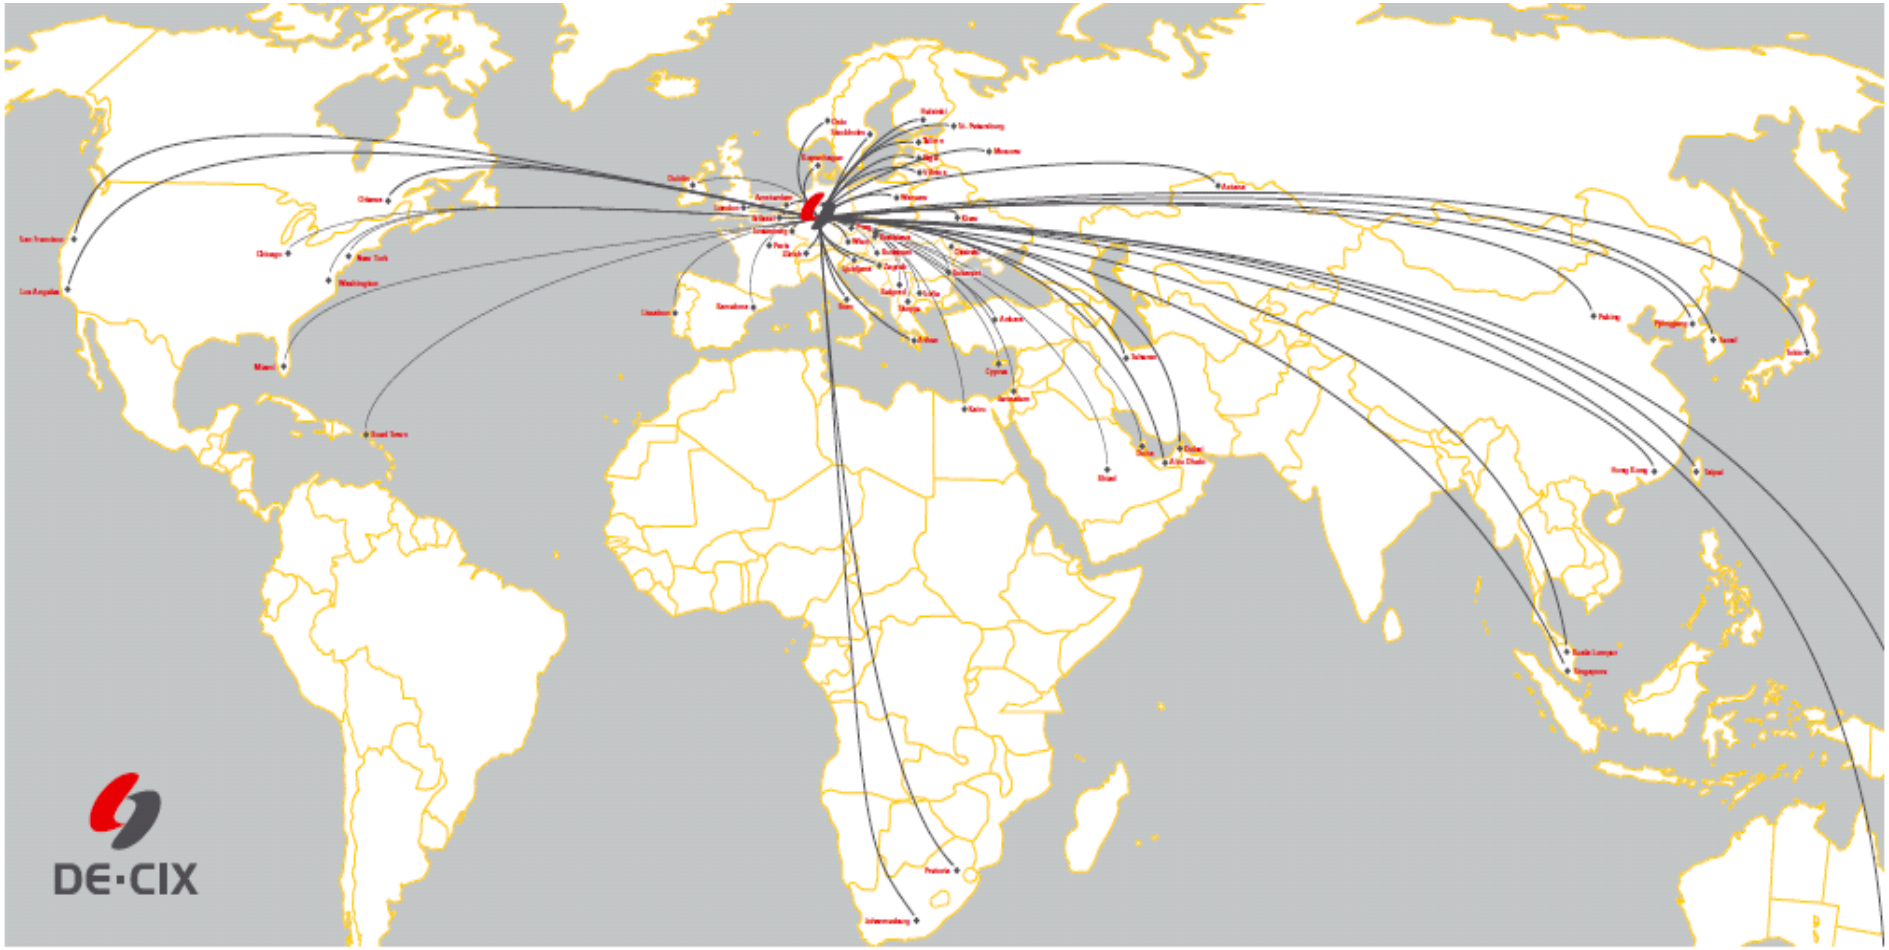

customers

452!

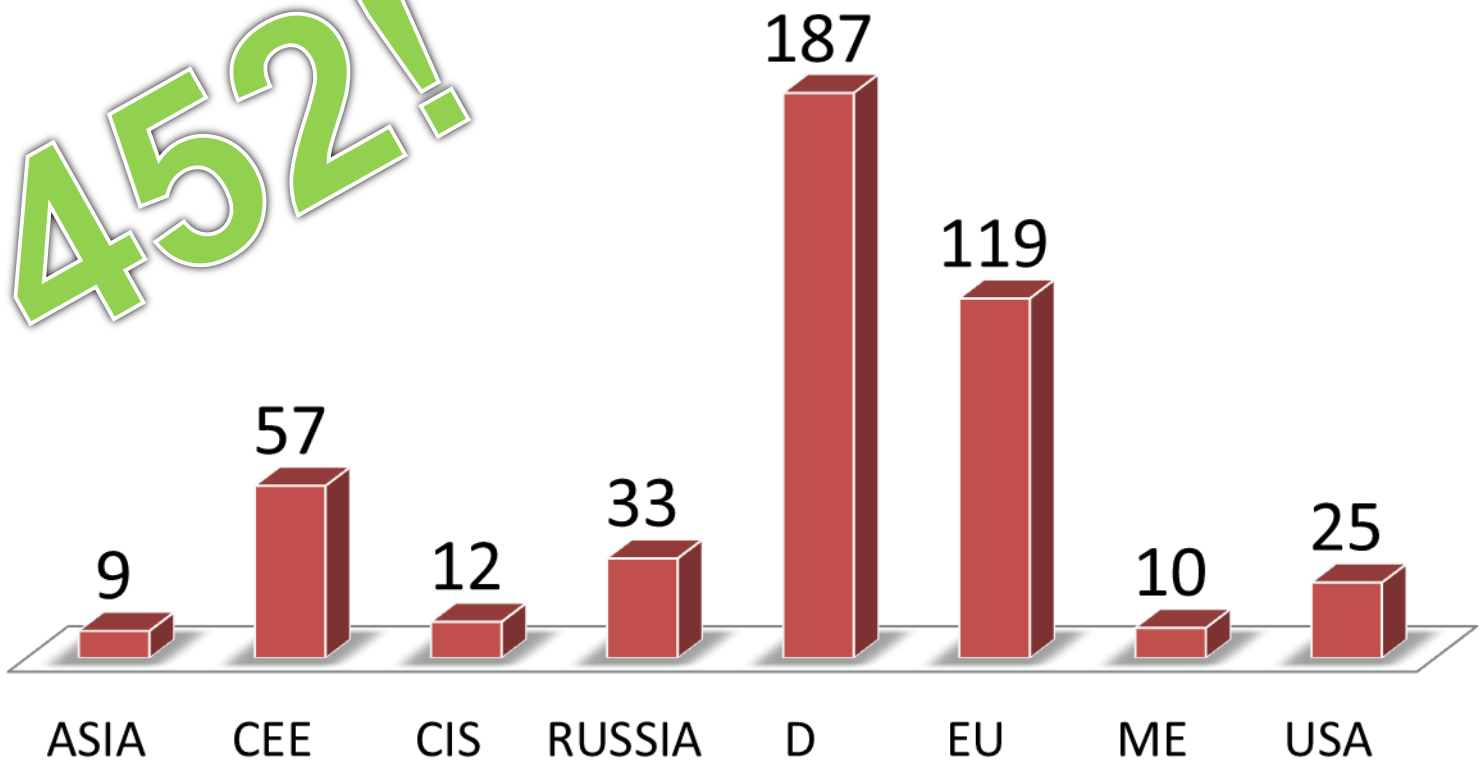

DE·CIX Where networks meet

capacity in Gbps

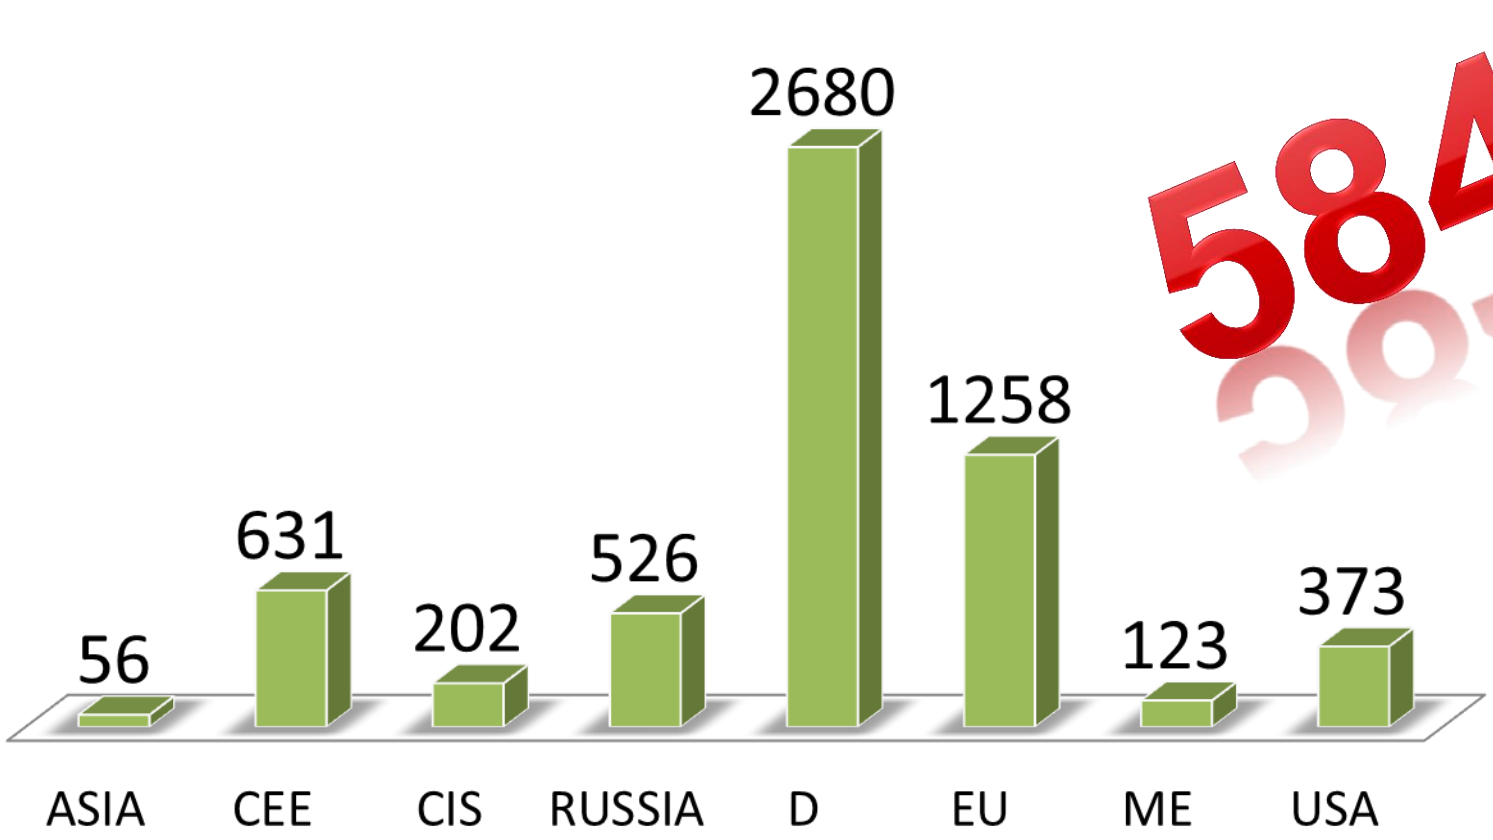

5848

Motivation

- *Customers saw attacks going over our platform*
- *Questions for help arose:*
  - *Can you filter traffic from X going to Y?*
  - *Can you mitigate?*
  - *Can you **blackhole**?*

Motivation

- *The answer now is:*

  –YES!

What is blackholing?

- *Blackholing effectively means diverting the flow of data to a different (Blackhole) Next-hop, where the traffic is discarded*
- *The result is that no traffic is reaching the original destination and hence hosts located within the „blackholed“ prefix are protected*
- *Thus blackholing is an effective way of mitigating the effects of Distributed Denial of Service (DDoS) attacks, etc.*

Blackholing @DE-CIX

DE-CIX Blackholing Service – basic principle (I)

- *In standard conditions*
  - *Customers advertise their prefixes with a next-hop IP address belonging to their AS*
    - *IPv4: /8 <= and <= /24*
    - *IPv6: /19 <= and <= /48*
- *In case of attack*
  - *Customers advertise their prefixes with a unique DE-CIX-provided Blackhole Next-hop IP address (BN)*
    - *IPv4: /8 <= **up to = /32** (if and only if the BN is set)*
    - *IPv6: /19 <= **up to = /128** (if and only if the BN is set)*
  - *Further, same security checks apply as usual (whether the advertised prefix belongs to customer's ASN, etc.)*

DE-CIX Blackholing Service – basic principle (2)

- *L2 filtering*
  - *Blackhole Next-hop (BN) has a unique MAC address (determined by ARP for the BN IP address)*
  - *New „deny“ rule was introduced in L2 ACLs on all customer ports*
  - *All traffic with BN‘s MAC address as destination is denied ingress*
- *As a result, all traffic to the attacked and „blackholed“ prefix is discarded already on the switch, and hence victim‘s resources are protected*

Example

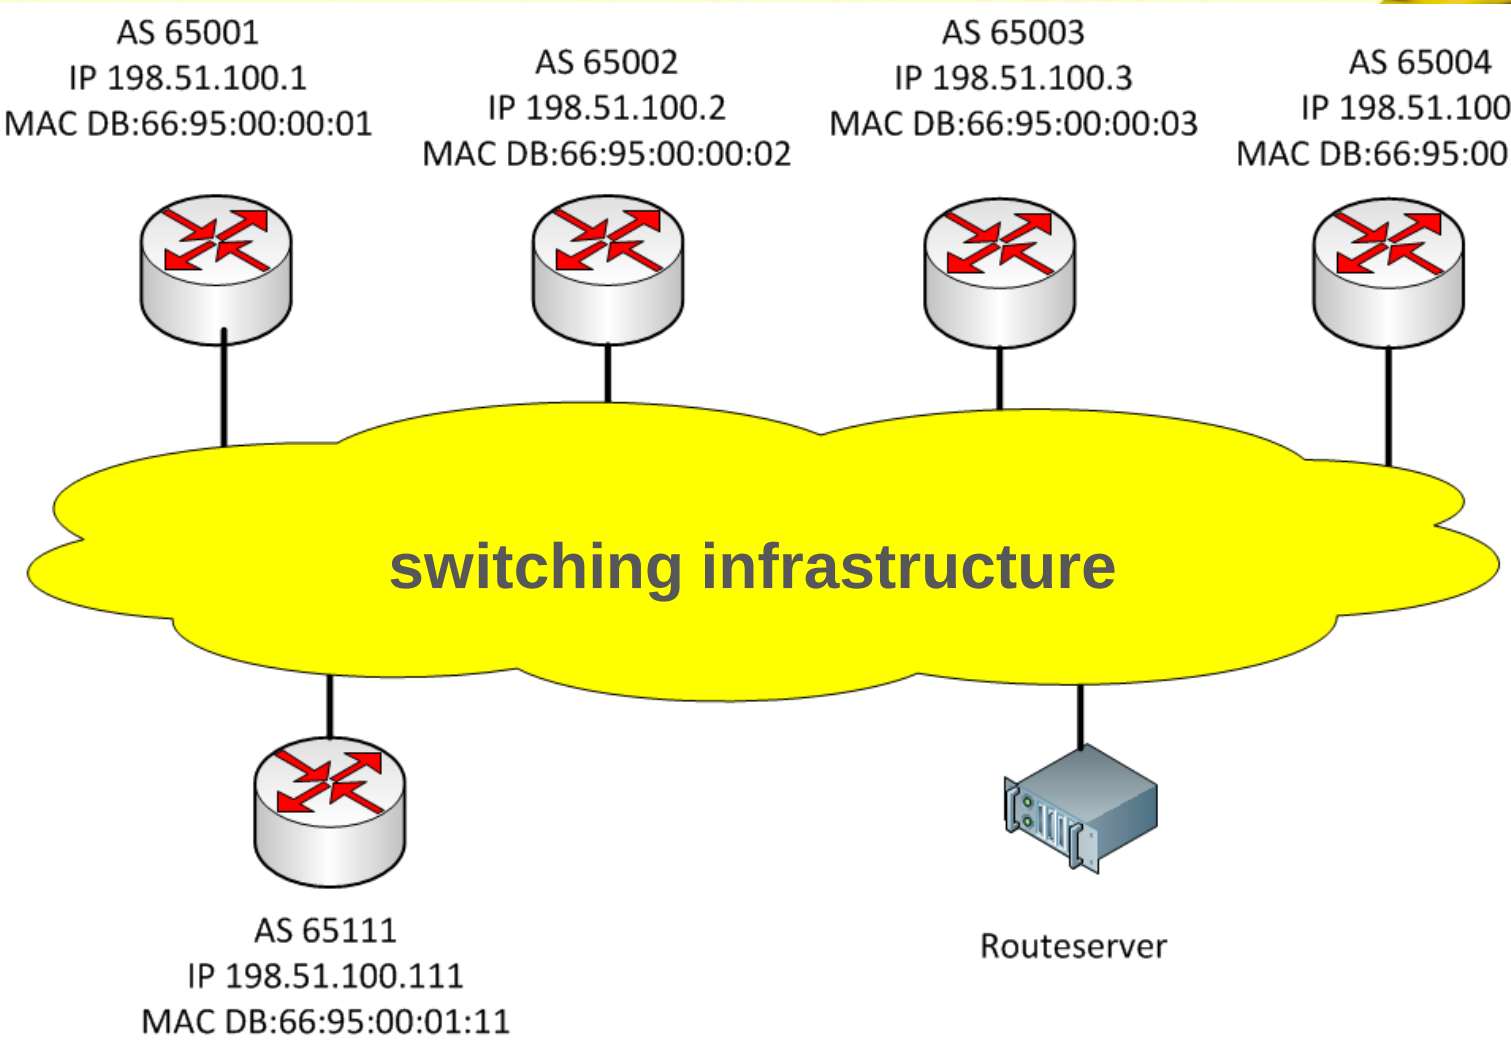
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
switching infrastructure
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver

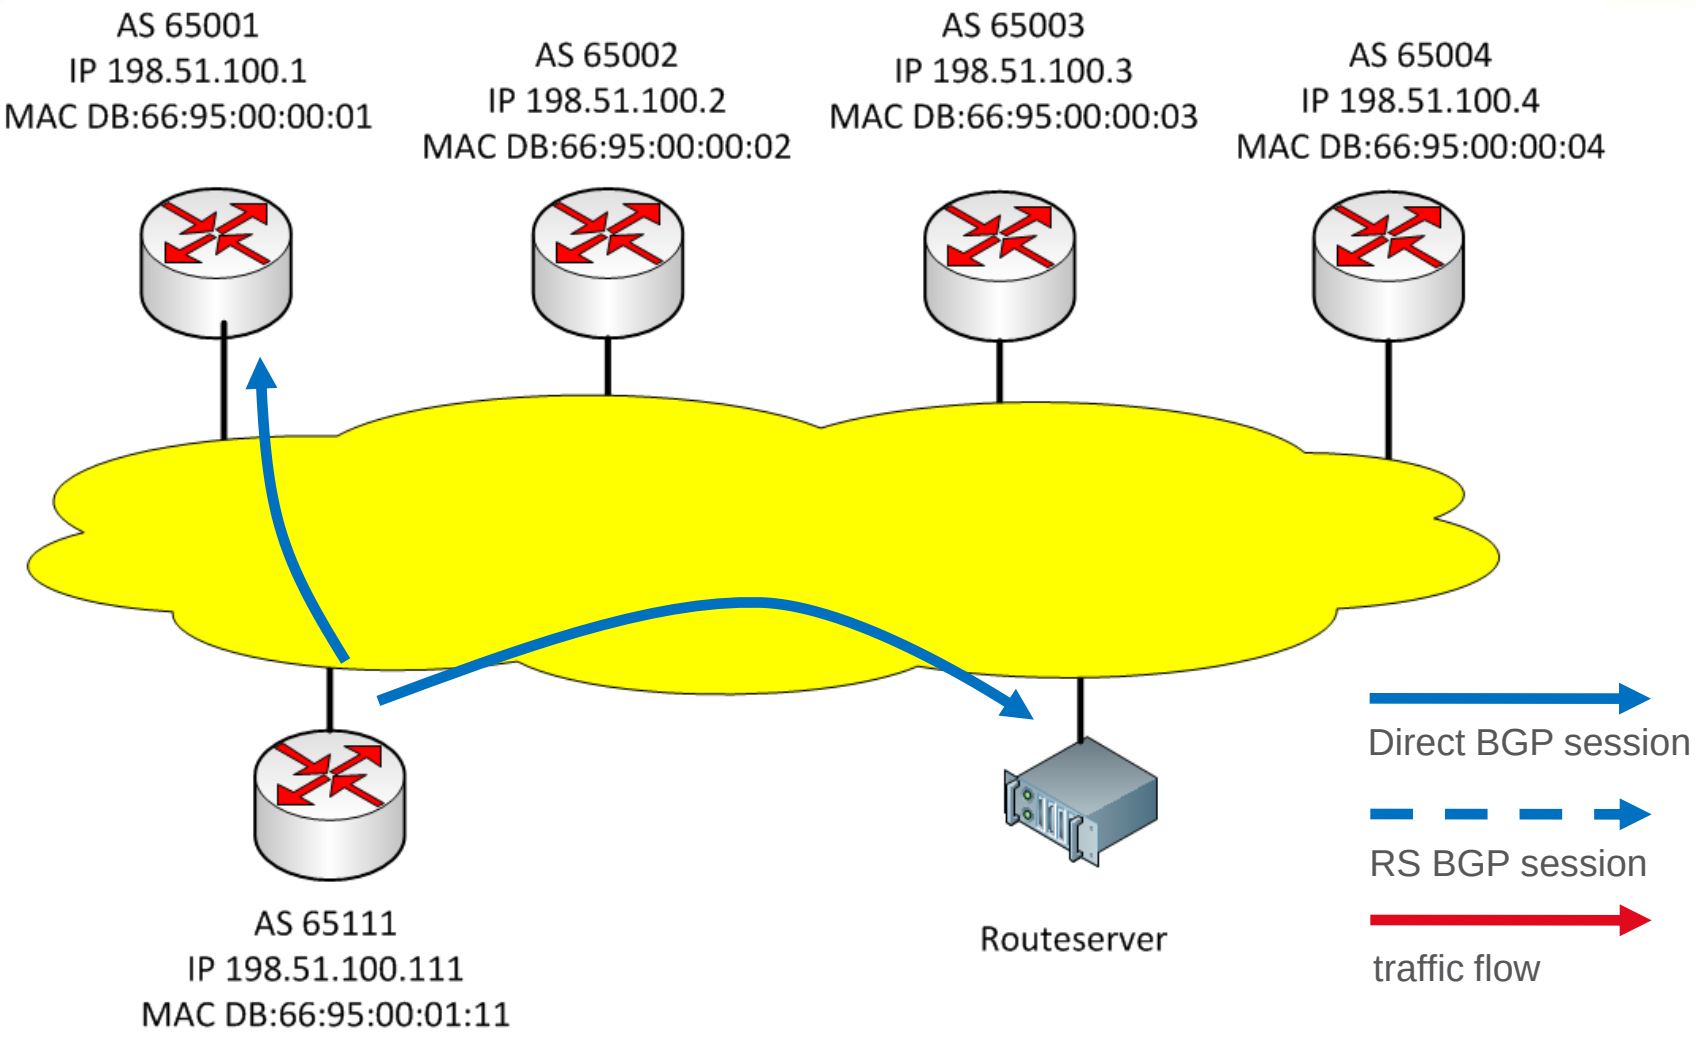
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Direct BGP session
RS BGP session
traffic flow

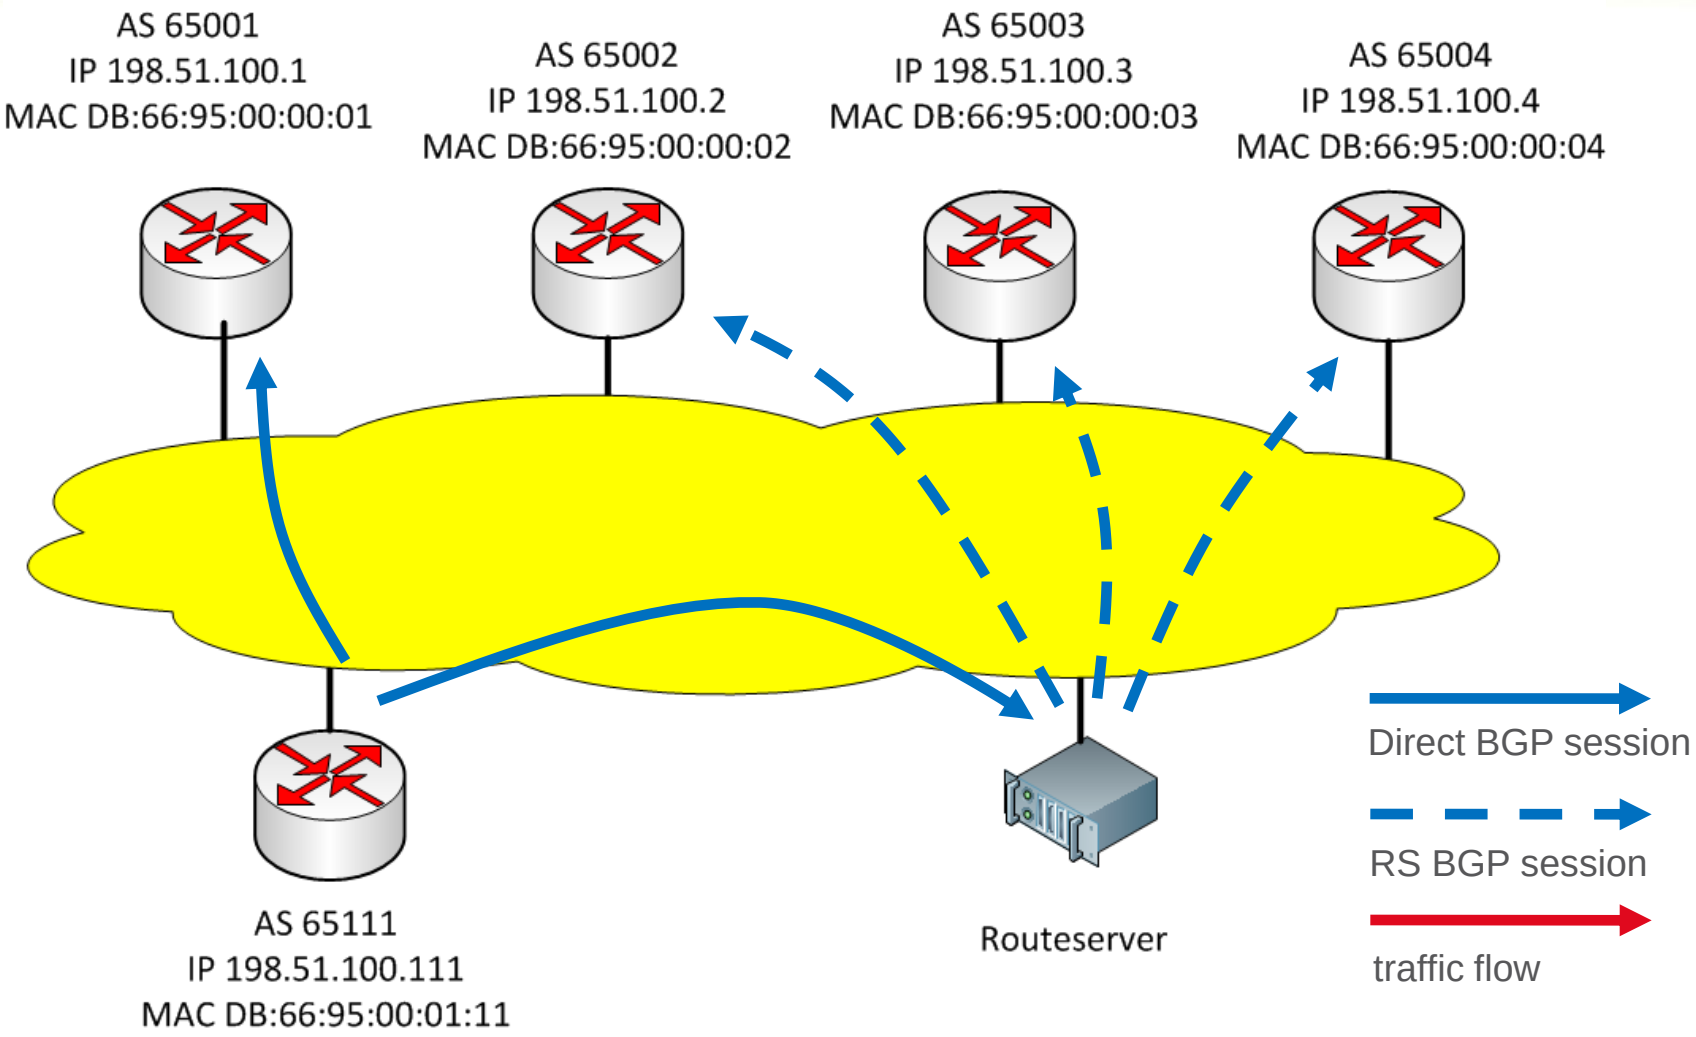
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Direct BGP session
RS BGP session
traffic flow

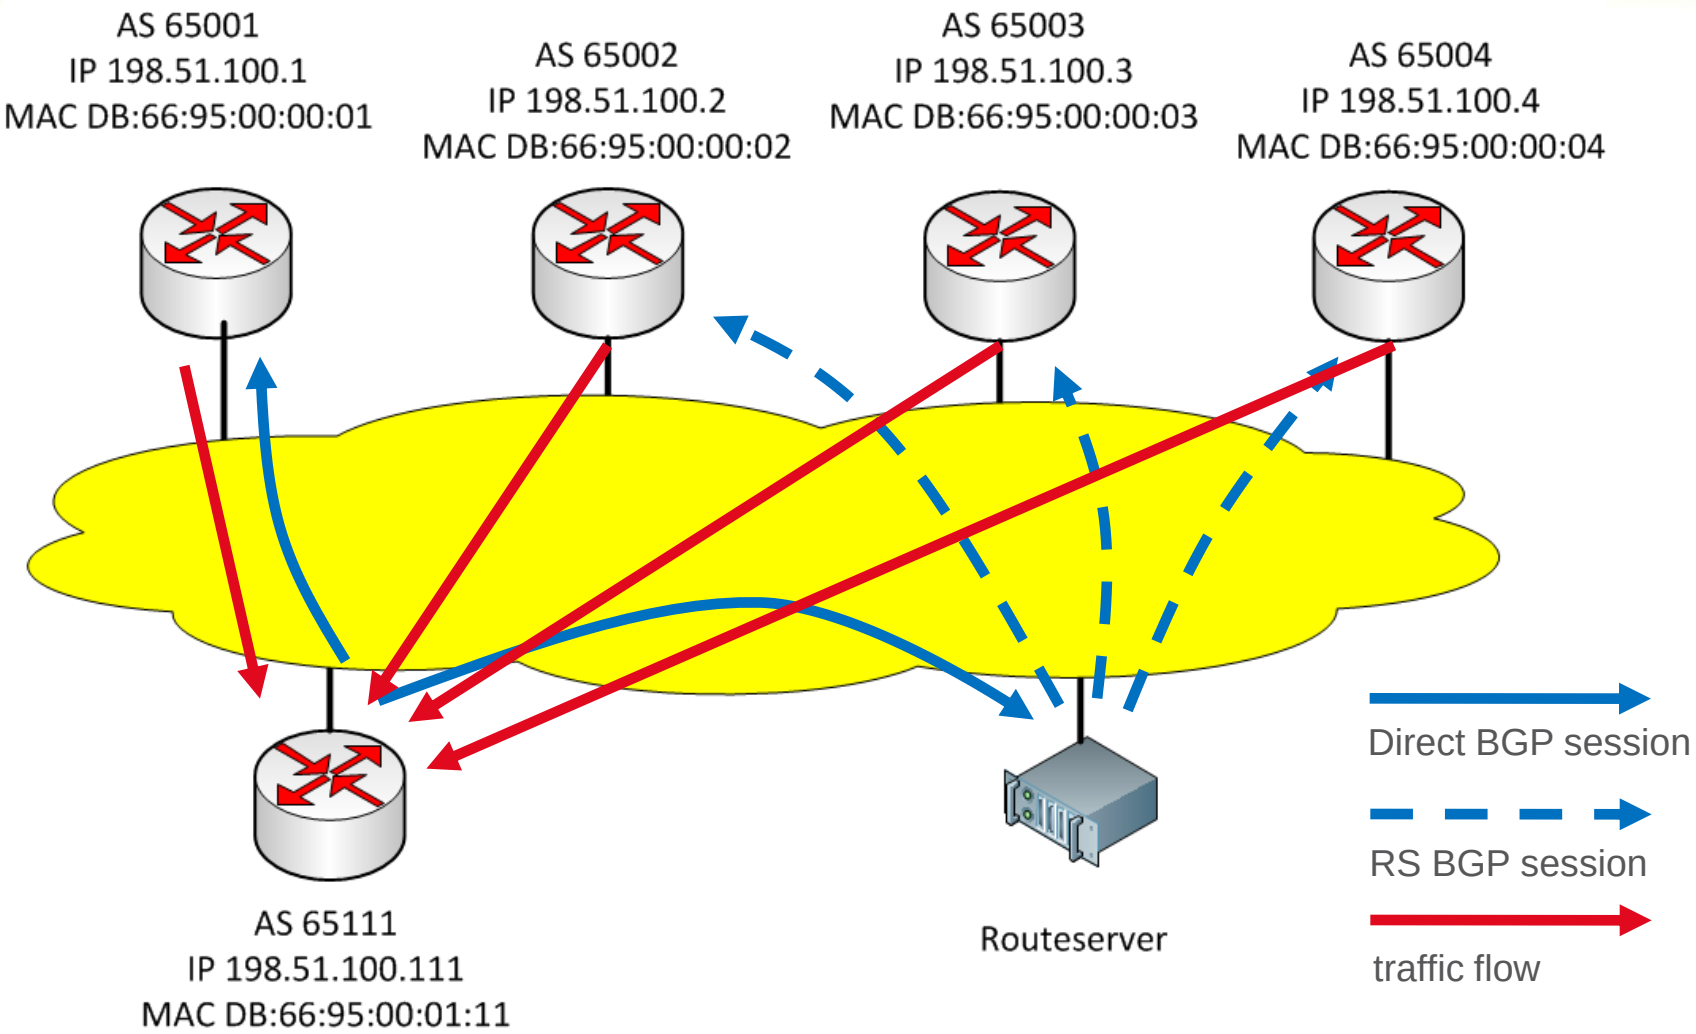
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Direct BGP session
RS BGP session
traffic flow

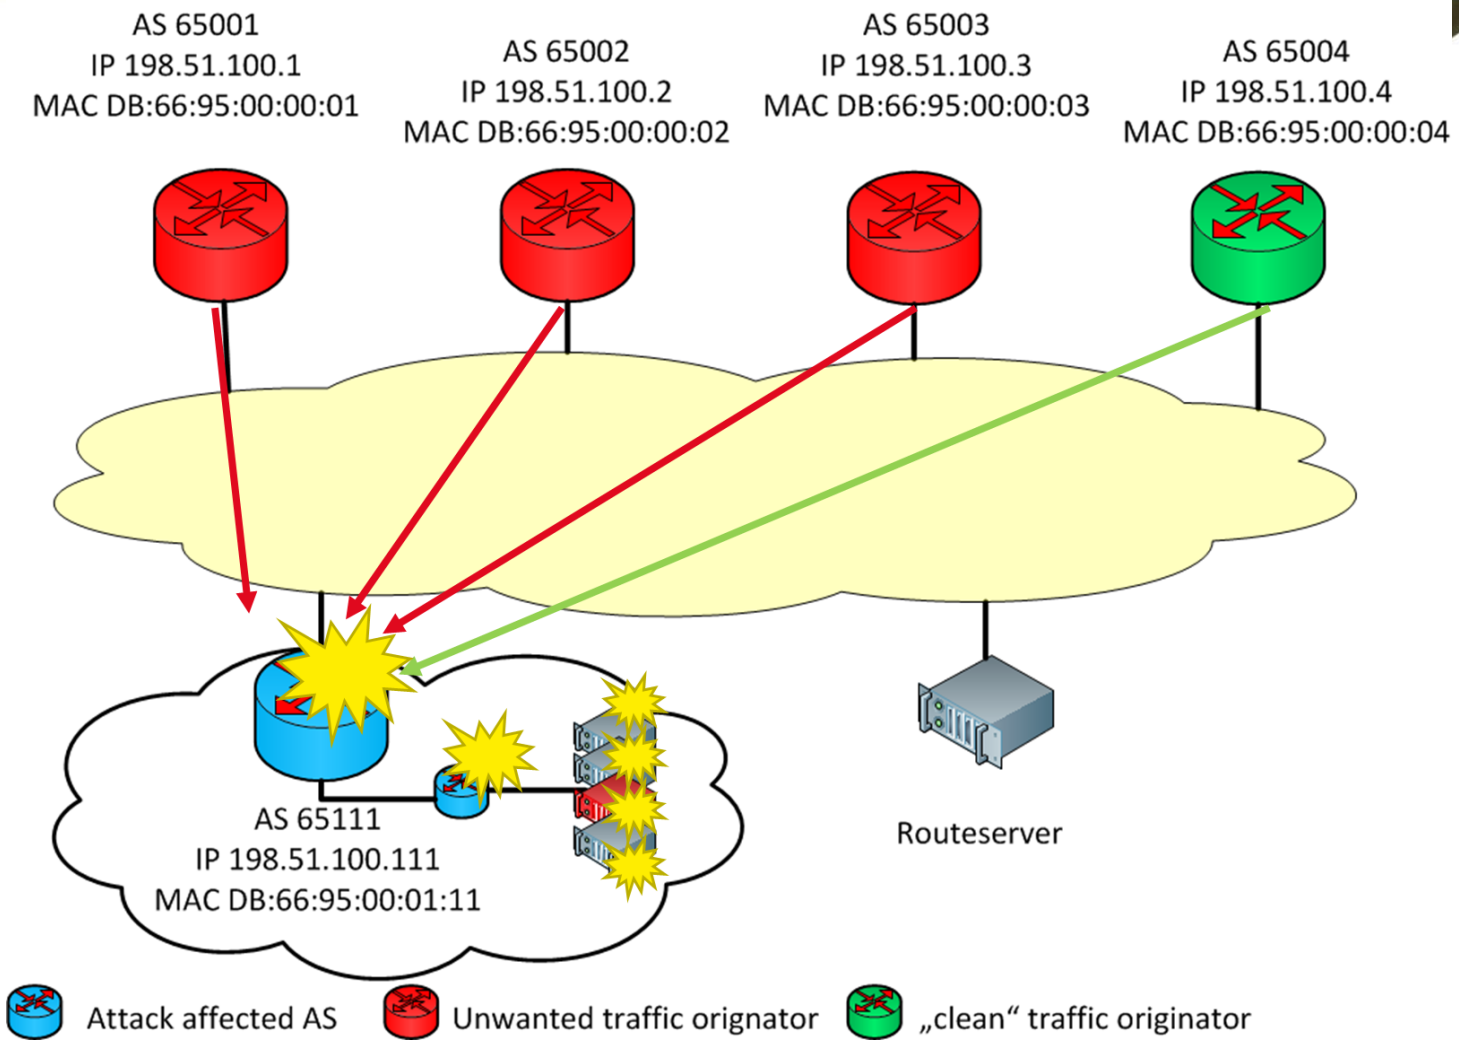
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Attack affected AS
Unwanted traffic orignator
„clean“ traffic originator

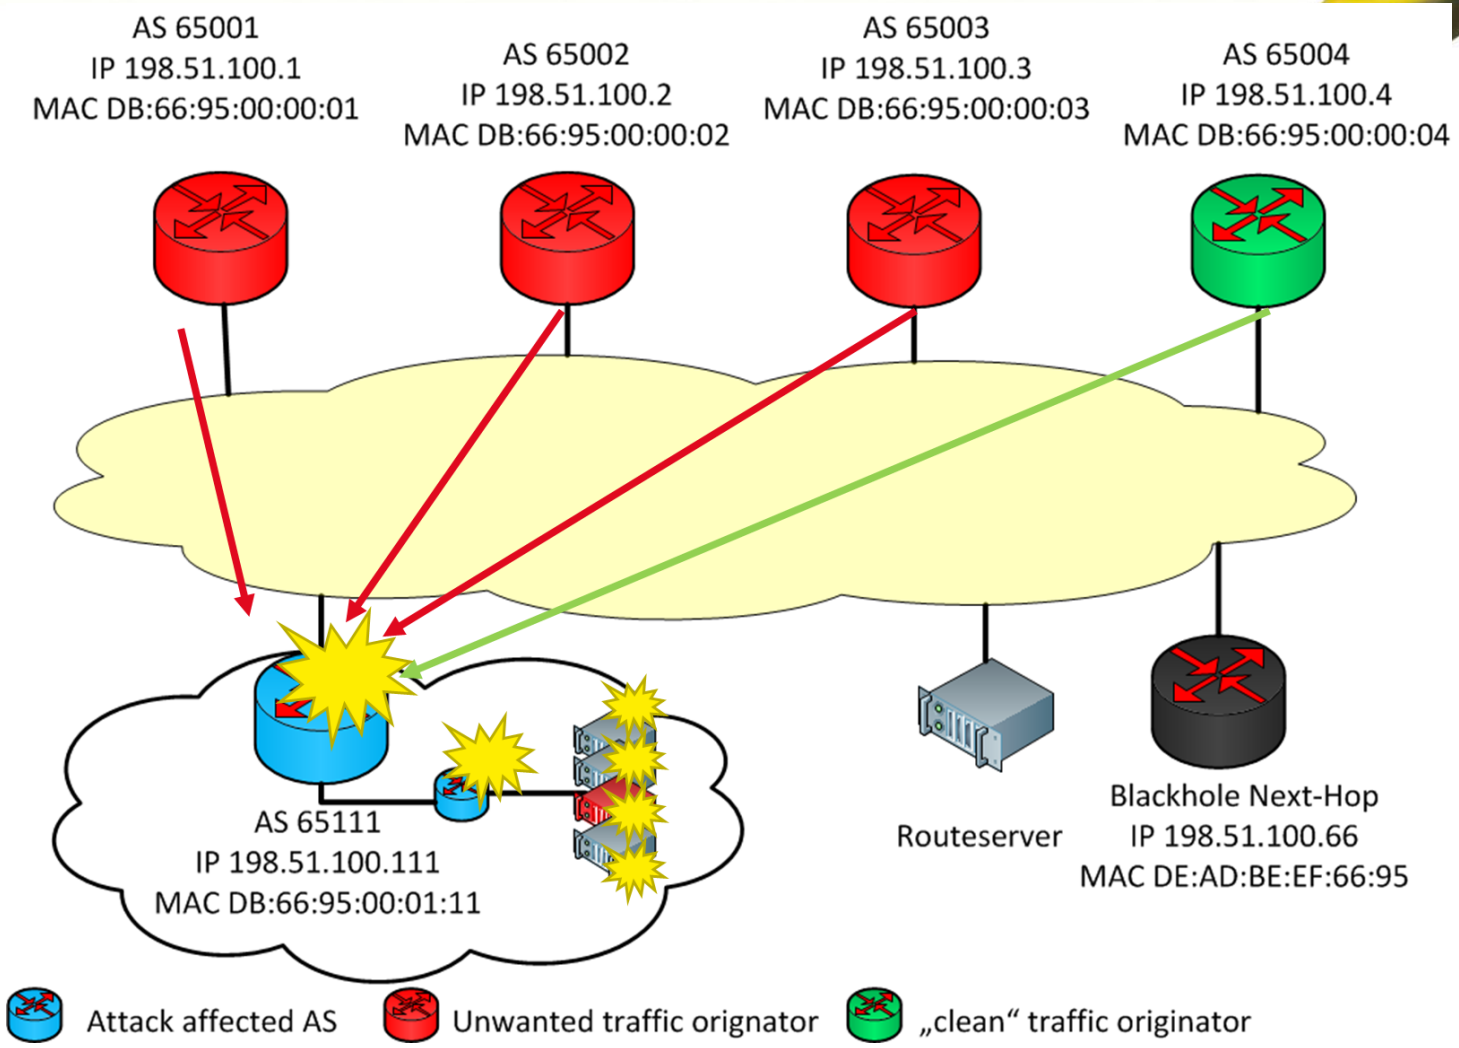
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Blackhole Next-Hop
IP 198.51.100.66
MAC DE:AD:BE:EF:66:95
Attack affected AS
Unwanted traffic orignator
„clean“ traffic originator

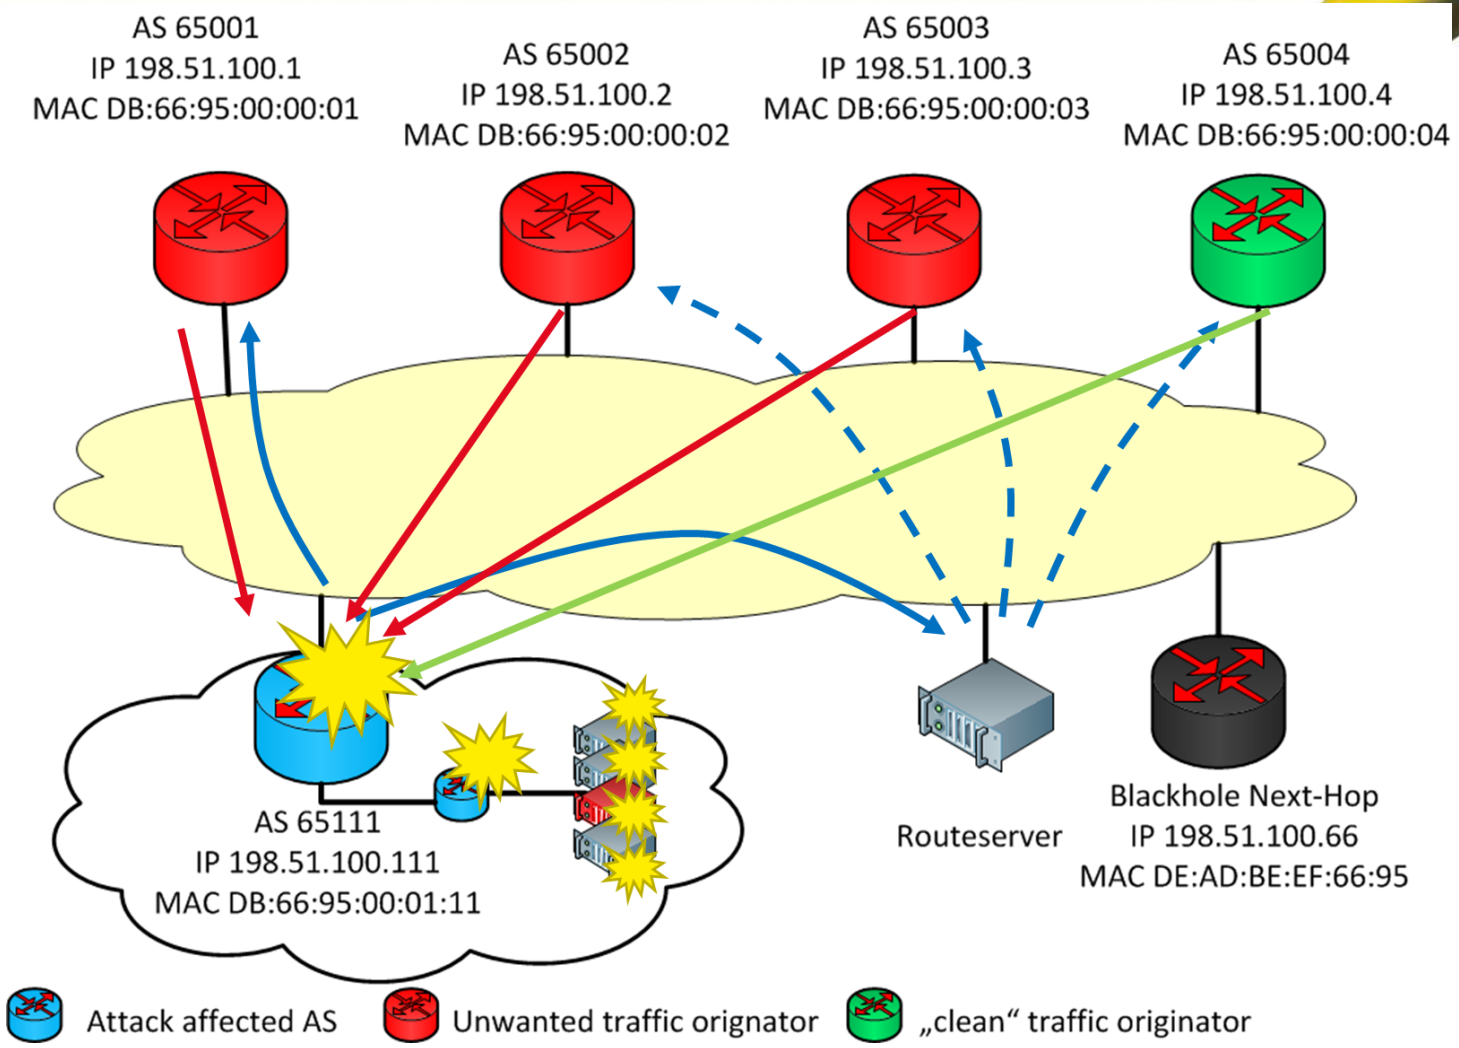
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Blackhole Next-Hop
IP 198.51.100.66
MAC DE:AD:BE:EF:66:95
Attack affected AS
Unwanted traffic orignator
„clean“ traffic originator

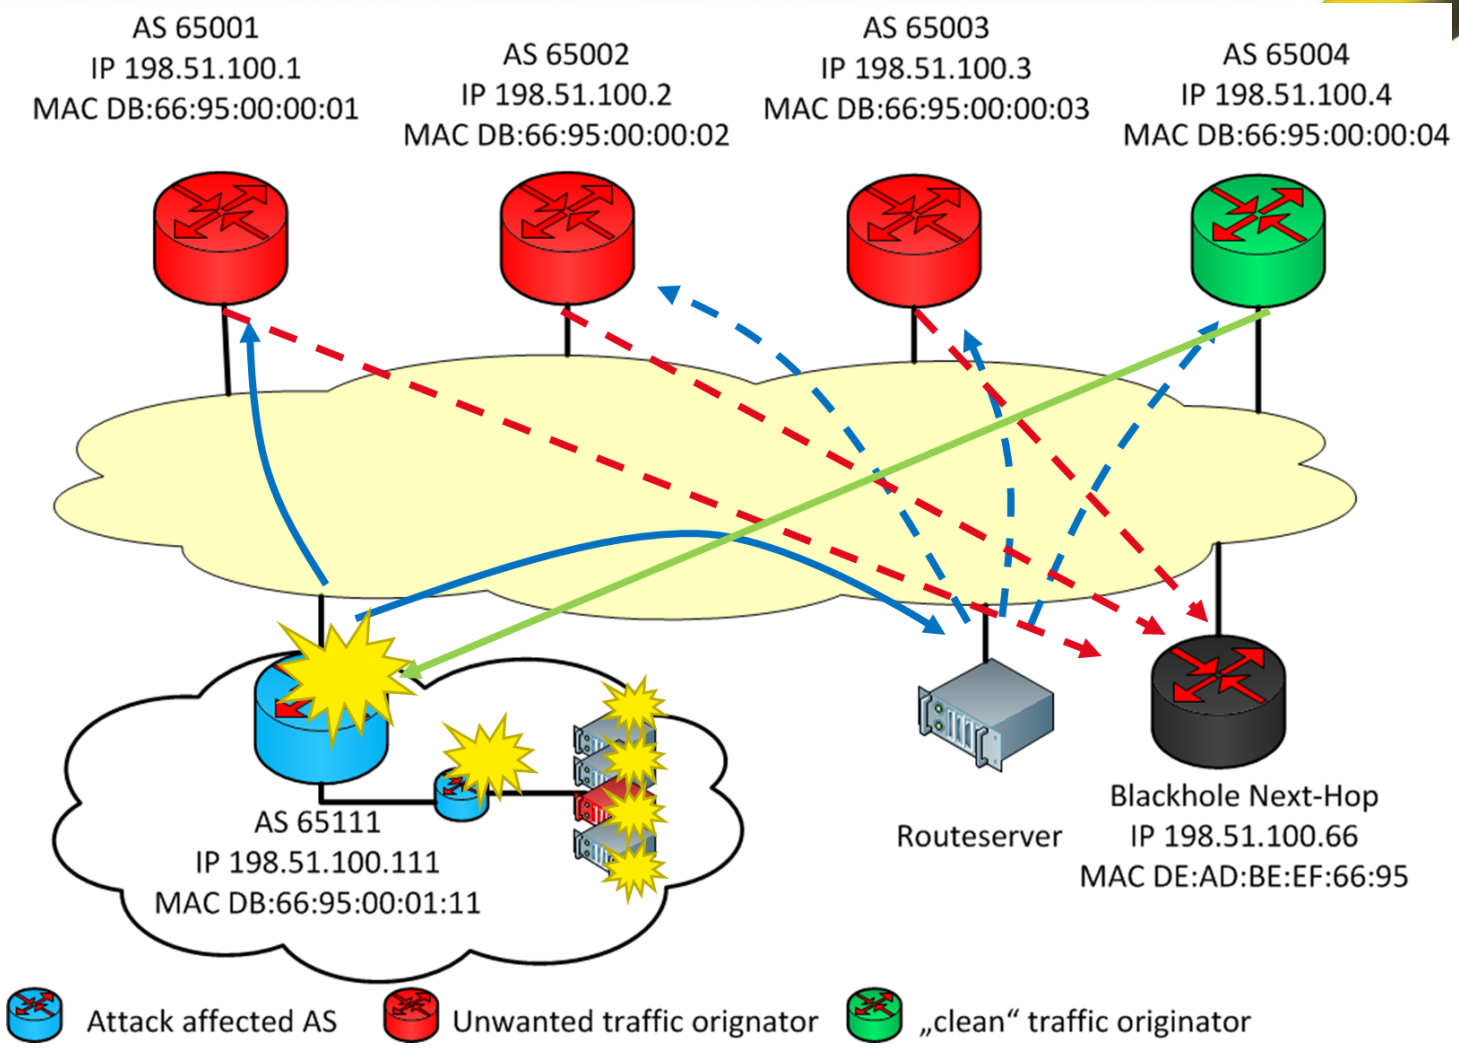
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Blackhole Next-Hop
IP 198.51.100.66
MAC DE:AD:BE:EF:66:95
Attack affected AS
Unwanted traffic orignator
„clean“ traffic originator

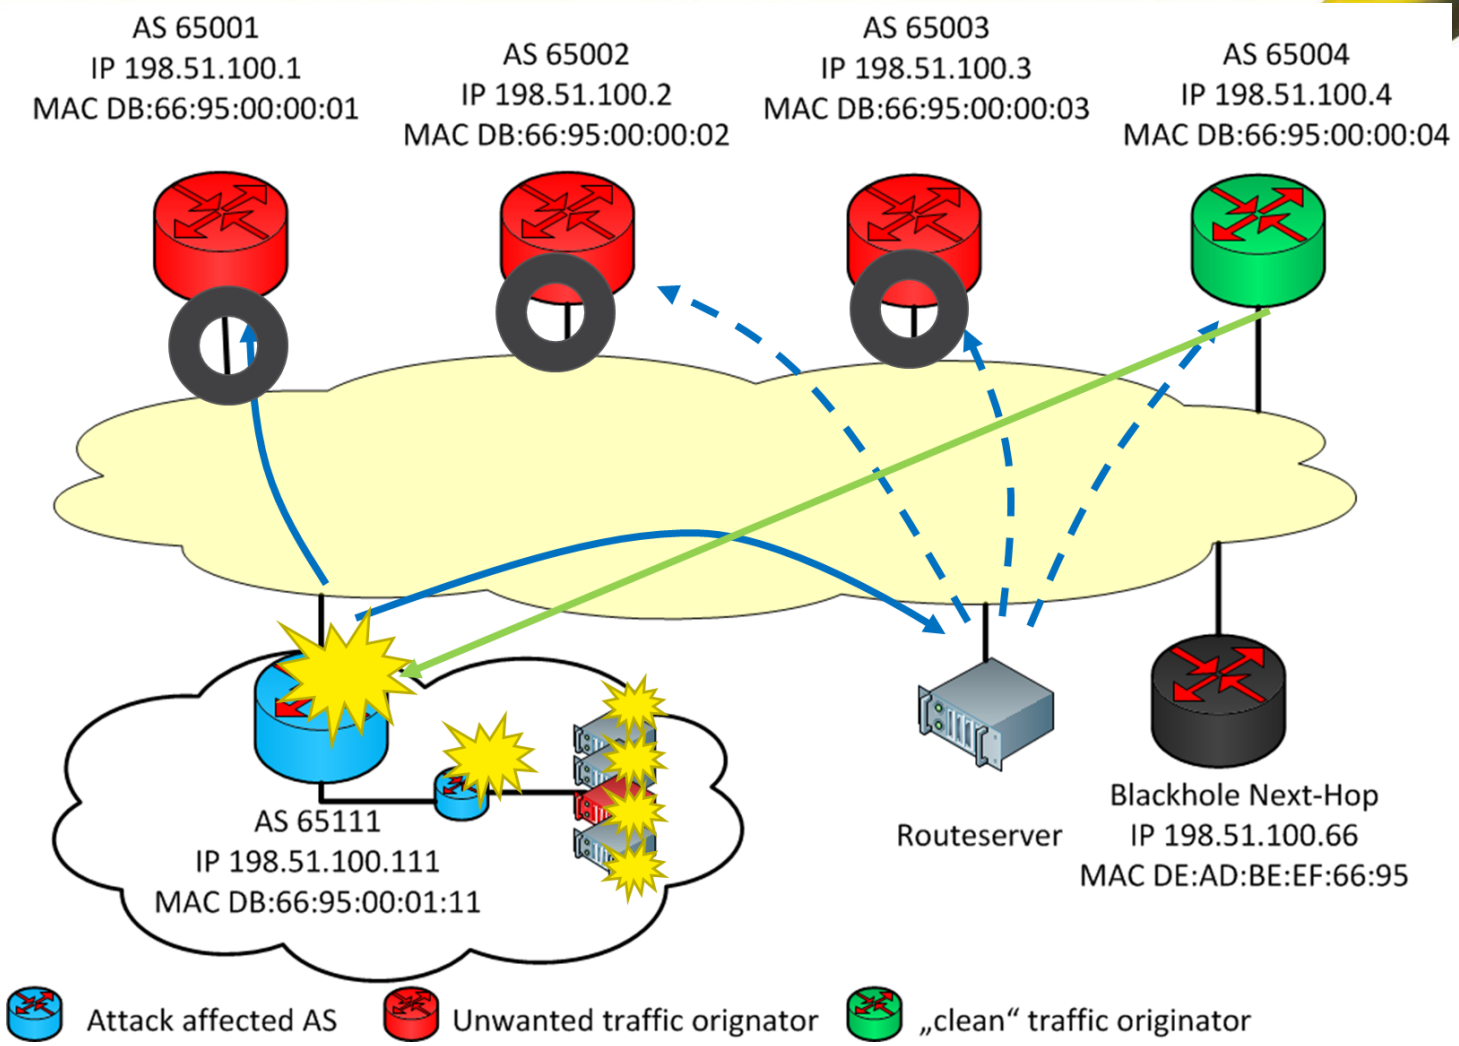
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Blackhole Next-Hop
IP 198.51.100.66
MAC DE:AD:BE:EF:66:95
Attack affected AS
Unwanted traffic orignator
„clean“ traffic originator

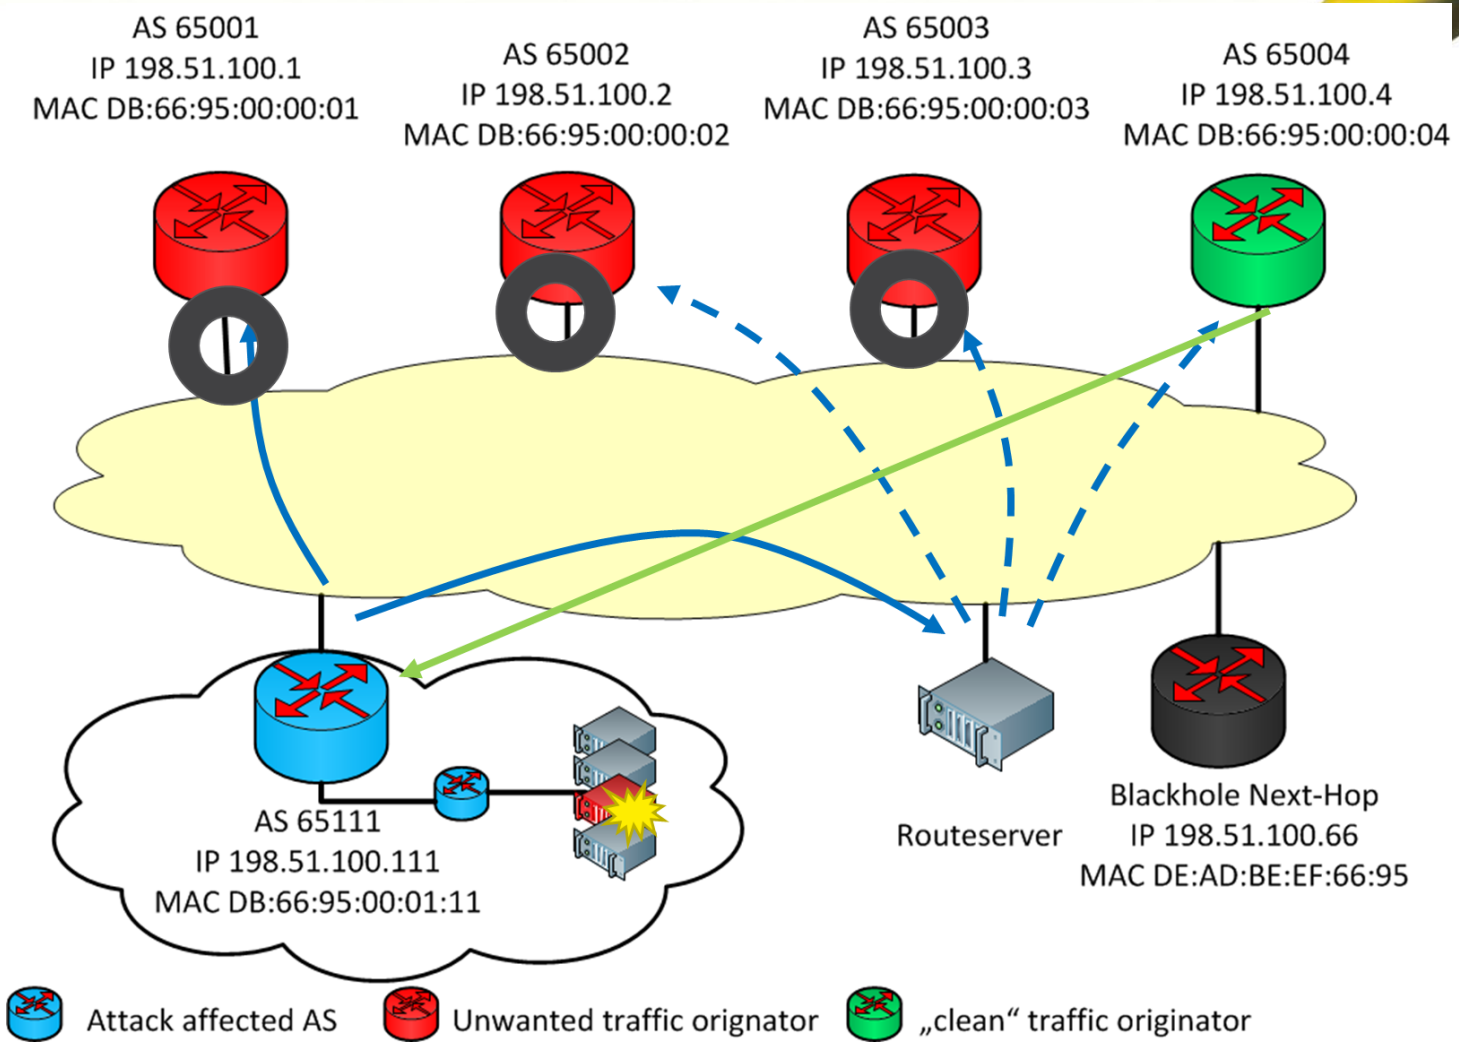
AS 65001
IP 198.51.100.1
MAC DB:66:95:00:00:01
AS 65002
IP 198.51.100.2
MAC DB:66:95:00:00:02
AS 65003
IP 198.51.100.3
MAC DB:66:95:00:00:03
AS 65004
IP 198.51.100.4
MAC DB:66:95:00:00:04
AS 65111
IP 198.51.100.111
MAC DB:66:95:00:01:11
Routeserver
Blackhole Next-Hop
IP 198.51.100.66
MAC DE:AD:BE:EF:66:95
Attack affected AS
Unwanted traffic orignator
„clean“ traffic originator

Example - Summary

- *AS65111 selectively announced the attacked prefix with the Blackhole Next-hop IP address*
- *All peers who had recieved this new route, learned the BN's MAC address via ARP/ND*
- *Traffic destined to the BN's MAC is dropped ingress via the L2 ACL*
- *AS65111's resources are preserved*

Blackholing – Important notes

- *Traffic from all upstream hosts to the „blackholed“ prefix is discarded*
  - *Including the normal/non-malicious traffic*
  - *Solution: If you know the origin ASN(s) from where the attack is coming, announce the blackhole routes with appropriate BGP communities (behavior similar to source-based blackholing)*
- *Traffic to all hosts in the „blackholed“ prefix is discarded*
  - *Including the hosts not under attack*
  - *Solution: You can advertise blackhole routes for prefixes as specific as /32 (IPv4) or /128 (IPv6)* *)

*) Please note, that according to CBP some of your peers might be filtering out the „more specific“ routes

Benefits of DE-CIX's solution

- *„KISS" approach*
- *Works for both direct and route server peering*
- *Easy configuration*
  - *On both customer and IX side*
  - *No need to support new types of community, etc.*
- *Robust solution*
  - *Dedicated unique BN IP and MAC addresses*
- *Customer-triggered*
  - *Customers can announce blackhole routes without having to ask for DE-CIX's approval*

Peer configuration example (IPv4, Cisco IOS 12.4(24)T)

```
!
router bgp <your ASN>
 no bgp enforce-first-as
 bgp log-neighbor-changes
 neighbor <RS> remote-as 6695
 !
 address-family ipv4
  neighbor <RS> activate
  neighbor <RS> route-map blackhole_out out
  network <your prefix> mask <mask>
 exit-address-family
!
ip prefix-list blackholing seq 5 permit <blackholed prefix>
!
 route-map blackhole_out permit 5
 match ip address prefix-list blackholing
 set ip next-hop 80.81.193.66
!
route-map blackhole_out permit 10
 set ip next-hop <your IP>
!
```

DE-CIX Blackholing Service – FAQs

- *How many blackhole routes can I advertise?*
  - *Blackhole routes are included in the maximum number of advertised prefixes, hence number of your normal + blackhole routes should not exceed the allowed maximum*
- *How specific can the „blackholed" prefix be?*
  - *The prefix can be as specific as /32 (IPv4) or /128 (IPv6)* *)
- *Do I have to pay for using the DE-CIX Blackholing Service?*
  - *No – once a DE-CIX customer, use of blackholing is free of charge*
- *At which locations is the DE-CIX Blackholing Service available?*
  - *The service is currently available only at DE-CIX Frankfurt*

http://go.de-cix.net/blackholing

*) Please note, that according to CBP some of your peers might be filtering out the „more specific" routes

Challenges and future work

- *Development of a monitoring solution for customers*
  - *Current implementation*
    - *No data about the blackholed traffic available*
  - *Next step*
    - *Provide customers with blackholed traffic statistics*
    - *Thus help them decide, whether to reannounce the routes with the correct next-hop IP address again*
- *Deployment of the DE-CIX Blackholig Service at new locations*
- *Share your ideas with us – join DE-CIX -> Competence Group Security*

Questions/Discussion

http://go.de-cix.net/blackholing/

http://blip.tv/web-montag-frankfurt-am-main/wmfra37_3-6093481

Thank you

Join DE-CIX now!

DE-CIX Competence Center
Lindleystrasse 12
60314 Frankfurt/Germany

Phone +49 69 1730 902 - 0
info@de-cix.net

DE-CIX Competence Center @ Kontorhaus Building

Frankfurt Osthafen (Docklands)