NIX.CZ – PEERING LAN RTBH

RFC 7999 compliant

NIX.CZ / nmc@nix.cz

20171113 by pj@nix.cz
Valid since 20171120

NIX·CZ

What is RTBH?

- RTBH = Remotely Triggered Black Hole filtering
- RTBH effectively means diverting the flow of data to a different (black hole) next-hop, where the traffic is discarded.
- The result is that no traffic is reaching the original destination and hence hosts located within the „black holed“ prefix are protected.
- Thus black holing is an effective way of mitigating the effects of Distributed Denial of Service (DDoS) attacks, etc.

RTBH @ PEERING – How it works?

- This service is provided by NIX.CZ route servers (RS) rs1.nix.cz and rs2.nix.cz for IPv4 and IPv6 prefixes.
- NIX.CZ route servers (RS) are accepting prefix granularity /(<=32) for IPv4 and /(<=128) for IPv6 only with announced BLACKHOLE community 65535:666.
- In case of an attack, peers can advertise their prefixes with a well-known BLACKHOLE BGP community 65535:666 to RS.
- When RS receives this well-known BGP community, it automatically changes next-hop IP address to the black hole IP address for this prefixes. RS is transparent for 65535:666 community and advertise it to all peers. Peers has then a possibility to build an incoming route-map to mach this community to accept such prefix length and are able to do some other actions described in RFC 7999.
- Black hole next-hop (BN) has a unique MAC address.
- All frames with BN MAC address destination are filtered with L2 ACL on ingress ports to NIX.CZ infrastructure.
- In this case, all traffic to the “black holed” prefixes is dropped before it reaches attacked ISP resources.

RTBH @ PEERING - Conditions

- rs1.nix.cz = 91.210.16.1 / 2001:7F8:14::11
- rs2.nix.cz = 91.210.16.2 / 2001:7F8:14::12
- Standard security checks for received prefixes are applied to all RS peers
- RS peers has to accept /(<=32) IPv4 and /(<=128) IPv6 prefix length or apply incoming route-map to accept such prefix length!
- Black hole MAC address = DE:AD:FA:CE:02:50
- Black hole IPv4 address = 91.210.16.250
- Black hole IPv6 address = 2001:7F8:14::250
- BLACKHOLE community (RFC 7999) = 65535:666

RTBH @ PEERING – RTBH communities

Example	Include community(ies)
Change next-hop to all RS peers	65535:666
Change next-hop to RS peers A, B, C ASNs only, do not send this prefix to other RS peers	47200:A 47200:B 47200:C 65535:666 0:47200
Do not advertise to RS peers ASNs A, B, C and change next-hop to all other RS peers	0:A 0:B 0:C 65535:666

Example	Include BLACKHOLE community and ext. community(ies)
Change next-hop to all RS peers	65535:666
Change next-hop to RS peers A, B, C ASNs only, do not send this prefix to other RS peers	65535:666 + rt:47200:A rt:47200:B rt:47200:C rt:0:47200
Do not advertise to RS peers ASNs A, B, C and change next-hop to all other RS peers	65535:666 + rt:0:A rt:0:B rt:0:C

RTBH @ PEERING – DDoS example

1. Standard situation

→ AS6881 advertise pfx. 195.47.235.0/24 to RS with no BGP community.

→ RS advertise this pfx. to all his peers → prefix is received/accepted and chosen as best-path.

→ The corresponding next-hop IP (91.210.16.245) and MAC is learned via ARP.

→ Peers traffic is flowing through NIX.CZ infrastructure to AS6881.

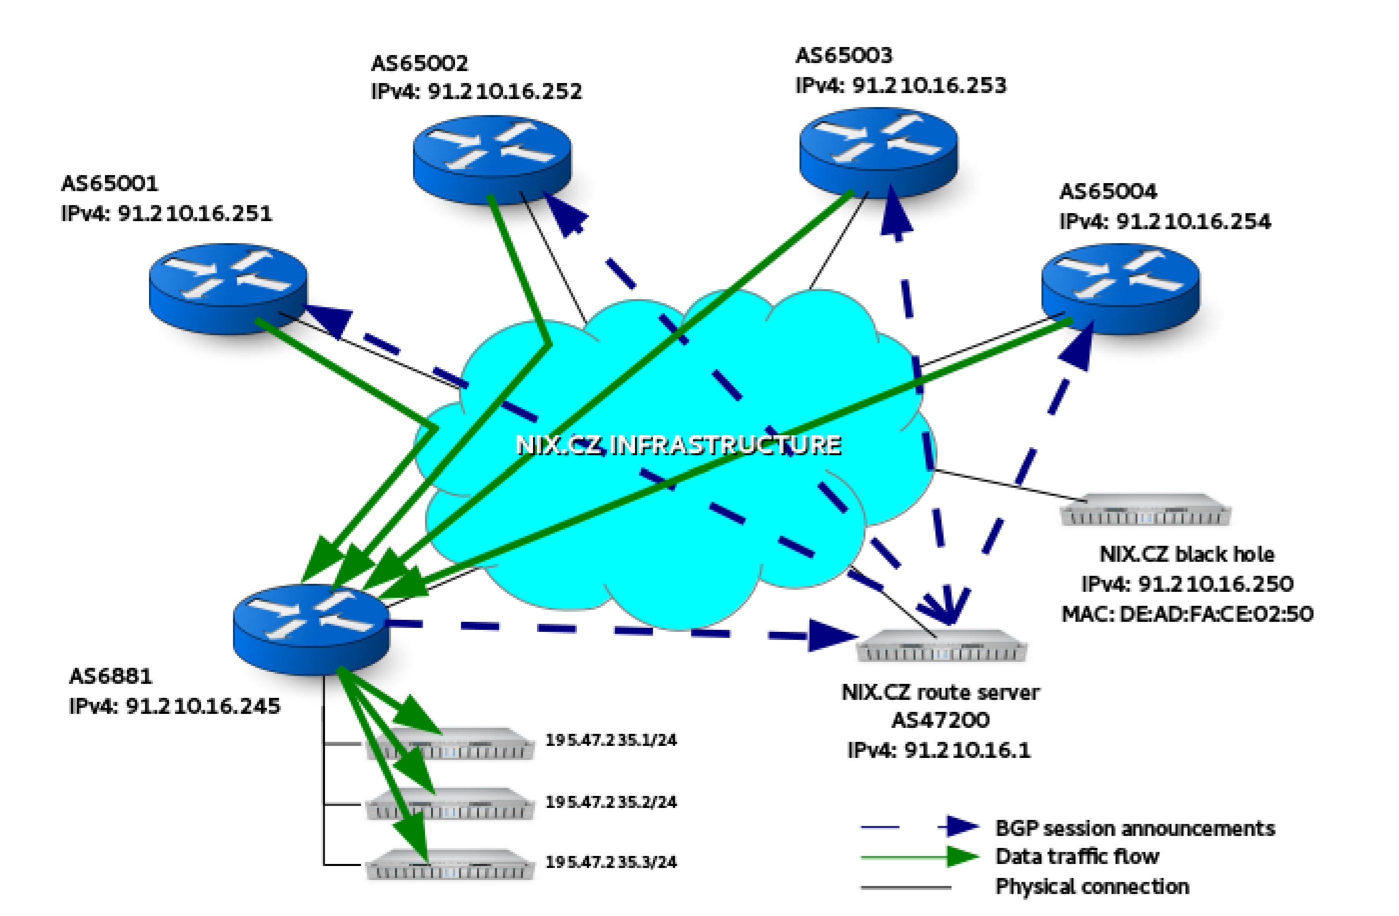

RTBH @ PEERING – DDoS example

2. DDoS attack

→ AS65001-3 are originators of malicious traffic attacking server 195.47.235.1 in AS6881.

→ AS65004 is originator of normal (“clean”) traffic flowing to 195.47.235.1 in AS6881.

→ Server 195.47.235.1 is overloaded → services are unreachable for all peers.

→ Other AS6881 IPs might be affected by this attack as well → port congestion, BGP sessions flapping, overloaded router CPUs, etc ...

RTBH @ PEERING – DDoS example

2. DDoS attack

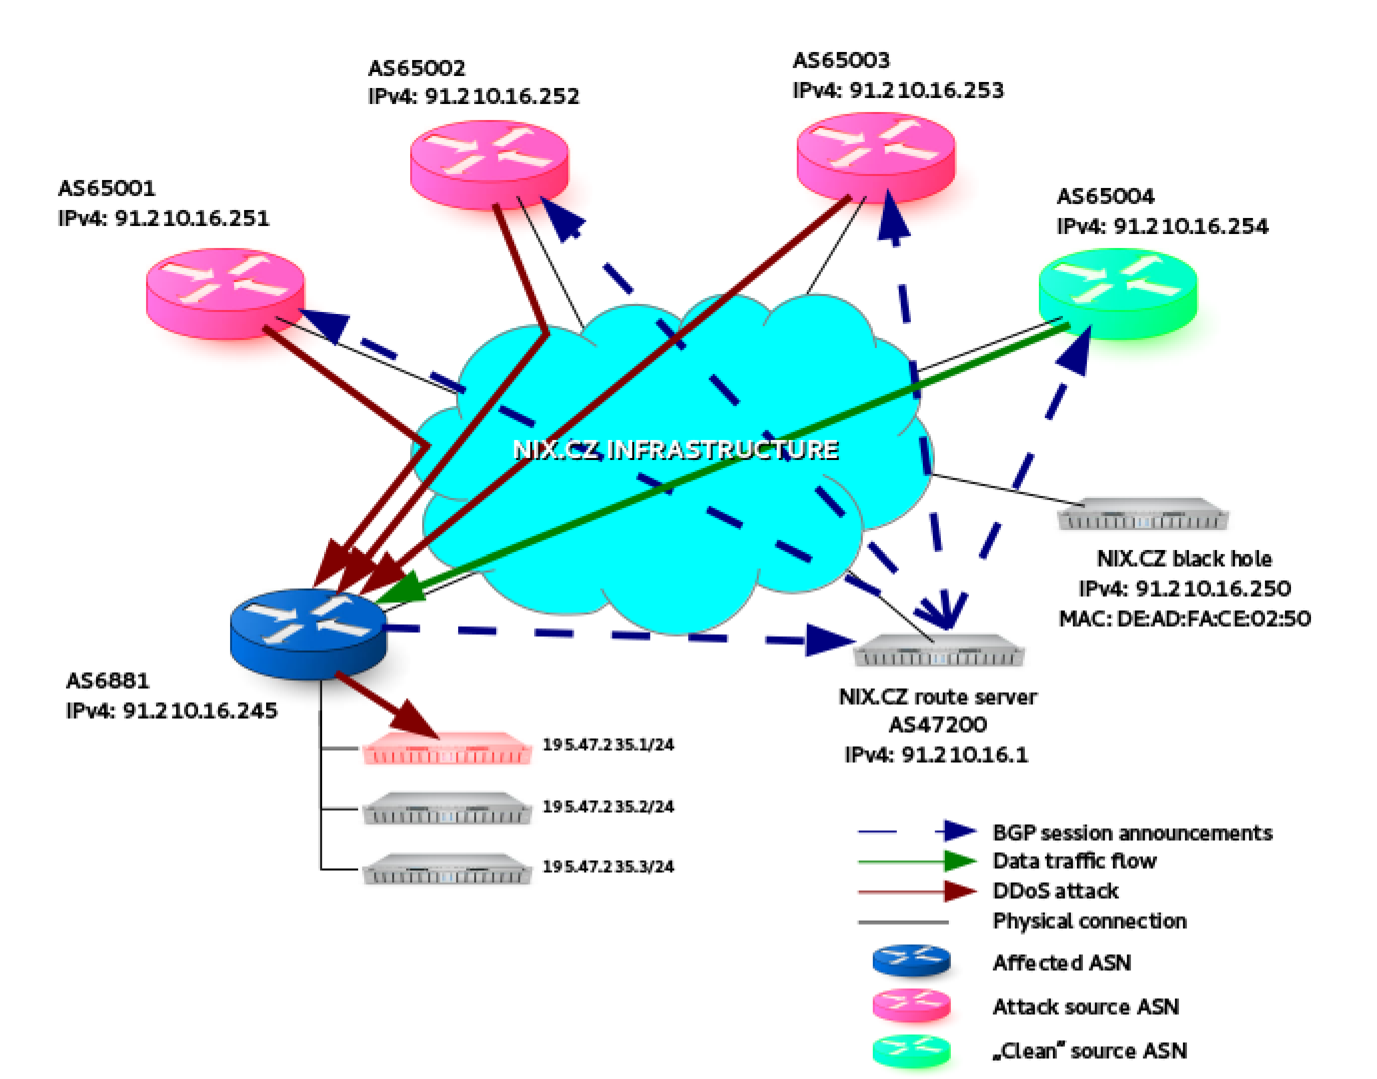

RTBH @ PEERING – DDoS example

3. DDoS attack mitigation

→ AS6881 starts to announce pfx. 195.47.235.1/32 with BGP community 47200:65001 47200:65002 47200:65003 65535:666 0:47200 to the RS.

→ RS receives this community and changes next-hop for the pfx. 195.47.235.1/32 to black hole IP (91.210.16.250) for peers AS65001-3 only.

→ AS65001-3 receives/accepts and choses prefix 195.47.235.1/32 as best-path.

→ AS65001-3 learns corresponding black hole next-hop IP and MAC via ARP.

→ AS65001-3 sends the traffic to black hole IP (91.210.16.250).

→ Traffic destined to the black hole MAC is dropped by ingress L2 ACL on NIX.CZ infrastructure.

→ AS65004 sends “clean” traffic to 195.47.235.1 with no problem, because RS does not change the next-hop for this client.

→ All traffic + DDoS from AS65001-3 to IP 195.47.235.1 is dropped before it reaches AS6881.

RTBH @ PEERING – DDoS example

3. DDoS attack mitigation

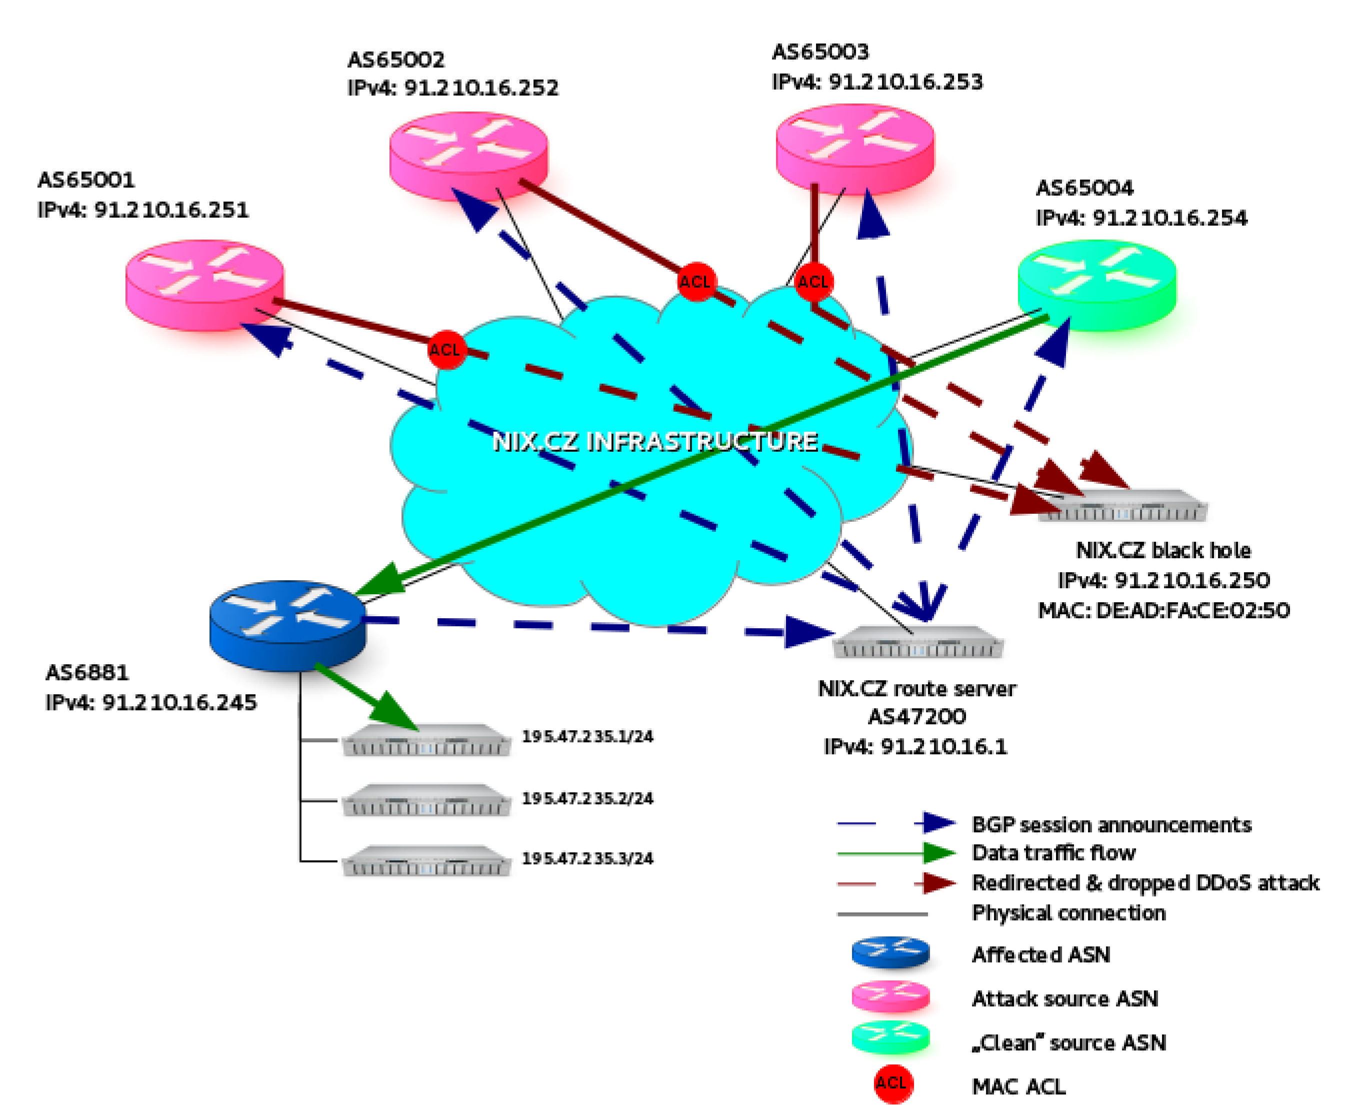

RTBH @ PEERING – DDoS example

3. DDoS attack mitigation

– Example router configuration for RTBH outbound announcements to black hole /32 pfx.:

(Cisco – IPv4)

```
ip prefix-list RTBH seq 5 permit <blackholed prefix/32>
!
router bgp <your ASN>
no bgp enforce-first-as
neighbor <RS> remote-as <NIX.CZ RS ASN>
!
address-family ipv4
network <blackholed prefix/32>
neighbor <RS> route-map RTBH-MAP out
exit-address-family
!
route-map RTBH-MAP permit 10
match ip address prefix-list RTBH
set community 47200:65001 47200:65002 47200:65003 65535:666 0:47200
#
```

RTBH @ PEERING – DDoS example

3. DDoS attack mitigation

– Example router configuration for RTBH inbound route-map to accept /32 pfx. with BLACKHOLE community:

(Cisco – IPv4)

```
ip community-list standard BLACKHOLE permit 65535:666
ip prefix-list IPv4-/24 seq 5 permit 0.0.0.0/0 le 24
ip prefix-list IPv4-/32 seq 5 permit 0.0.0.0/0 le 32
!
route-map AS6881-RS-IPv4-IN permit 10
match ip address prefix-list IPv4-/32
match community BLACKHOLE
set local-preference 666
set community no-export
!
route-map AS6881-RS-IPv4-IN permit 20
match ip address prefix-list IPv4-/24
set local-preference 10
#
```

¿ Questions ?

www.nix.cz
nmc@nix.cz