NETIX

Carrier Class Ethernet Exchange

What is Blackholing Service ?

- It’s a process of diverting unwanted data flow to a predefined (Blackhole) Next-hop where traffic is discarded
- It protects the services located within the ”blackholed” prefix so that no “bad“ traffic reach them
- This service is an effective way to mitigate the effects of Distributed Denial of Service (DDoS) attacks

How does the NetIX Blackholing Service work ?

- Next-hop change
  - Members advertise normally their prefixes with a next-hop IP address belonging to their ASN
  - When we receive a prefix marked with the Blackhole community (65499:999) we will change the next-hop IP address to 193.218.0.99 (Blackhole Next-hop)
- L2 Filtering
  - Blackhole Next-hop (BN) has a unique MAC address (determined by ARP or ND for the BN IP address)

Important: *Further, same security checks apply as usual (whether the advertised prefix belongs to the member's ASN, etc.)*

EXAMPLE

Normal situation

- AS65001 announce prefixes
  - via direct peering session (here AS65002)
  - via the route server, which propagates prefixes to other peers peering with the route server (here AS65003, AS65004, AS65005)
- The other ASes learn the BGP Next-hop for the announced prefix
  - prefix is received/accepted and chosen as best-path
- The corresponding Next-hop MAC is learned via ARP or ND

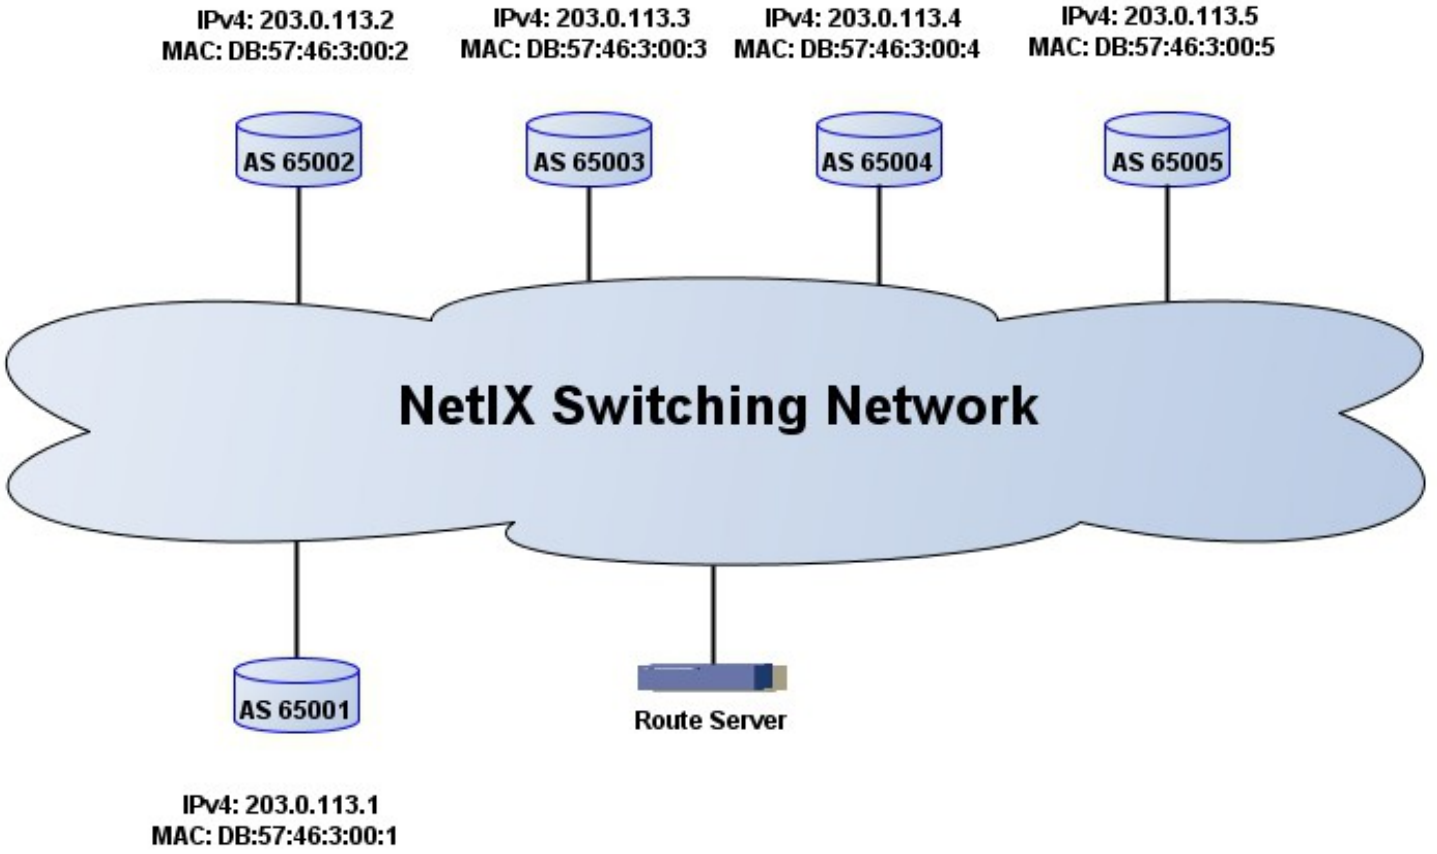
IPv4: 203.0.113.2
MAC: DB:57:46:3:00:2
IPv4: 203.0.113.3
MAC: DB:57:46:3:00:3
IPv4: 203.0.113.4
MAC: DB:57:46:3:00:4
IPv4: 203.0.113.5
MAC: DB:57:46:3:00:5
AS 65002
AS 65003
AS 65004
AS 65005
NetIX Switching Network
AS 65001
Route Server
IPv4: 203.0.113.1
MAC: DB:57:46:3:00:1

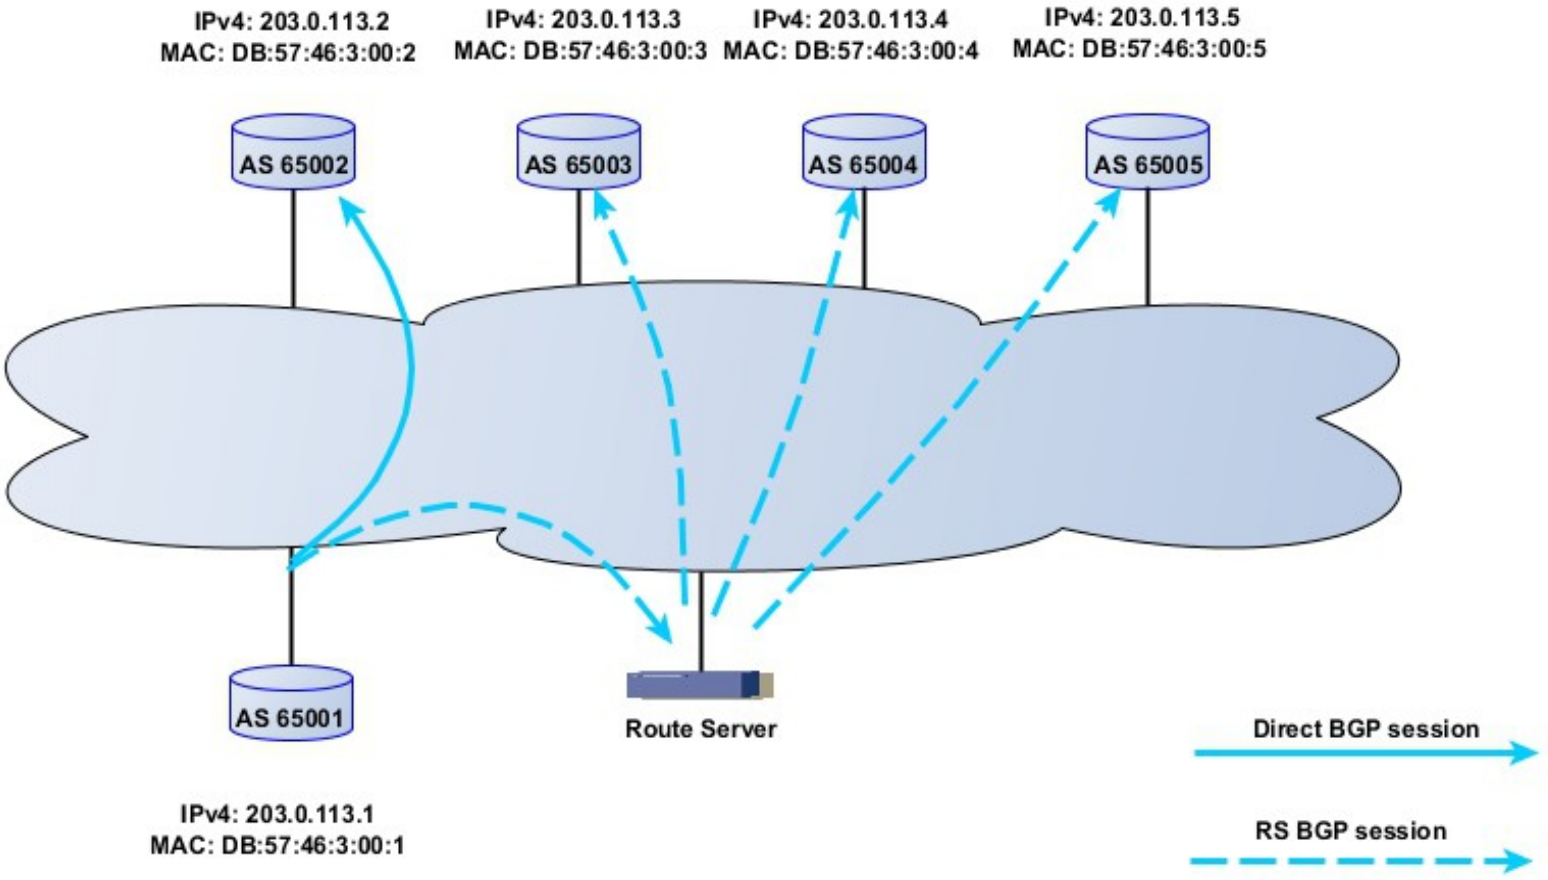
IPv4: 203.0.113.2
MAC: DB:57:46:3:00:2
IPv4: 203.0.113.3
MAC: DB:57:46:3:00:3
IPv4: 203.0.113.4
MAC: DB:57:46:3:00:4
IPv4: 203.0.113.5
MAC: DB:57:46:3:00:5
AS 65002
AS 65003
AS 65004
AS 65005
AS 65001
Route Server
IPv4: 203.0.113.1
MAC: DB:57:46:3:00:1
Direct BGP session
RS BGP session

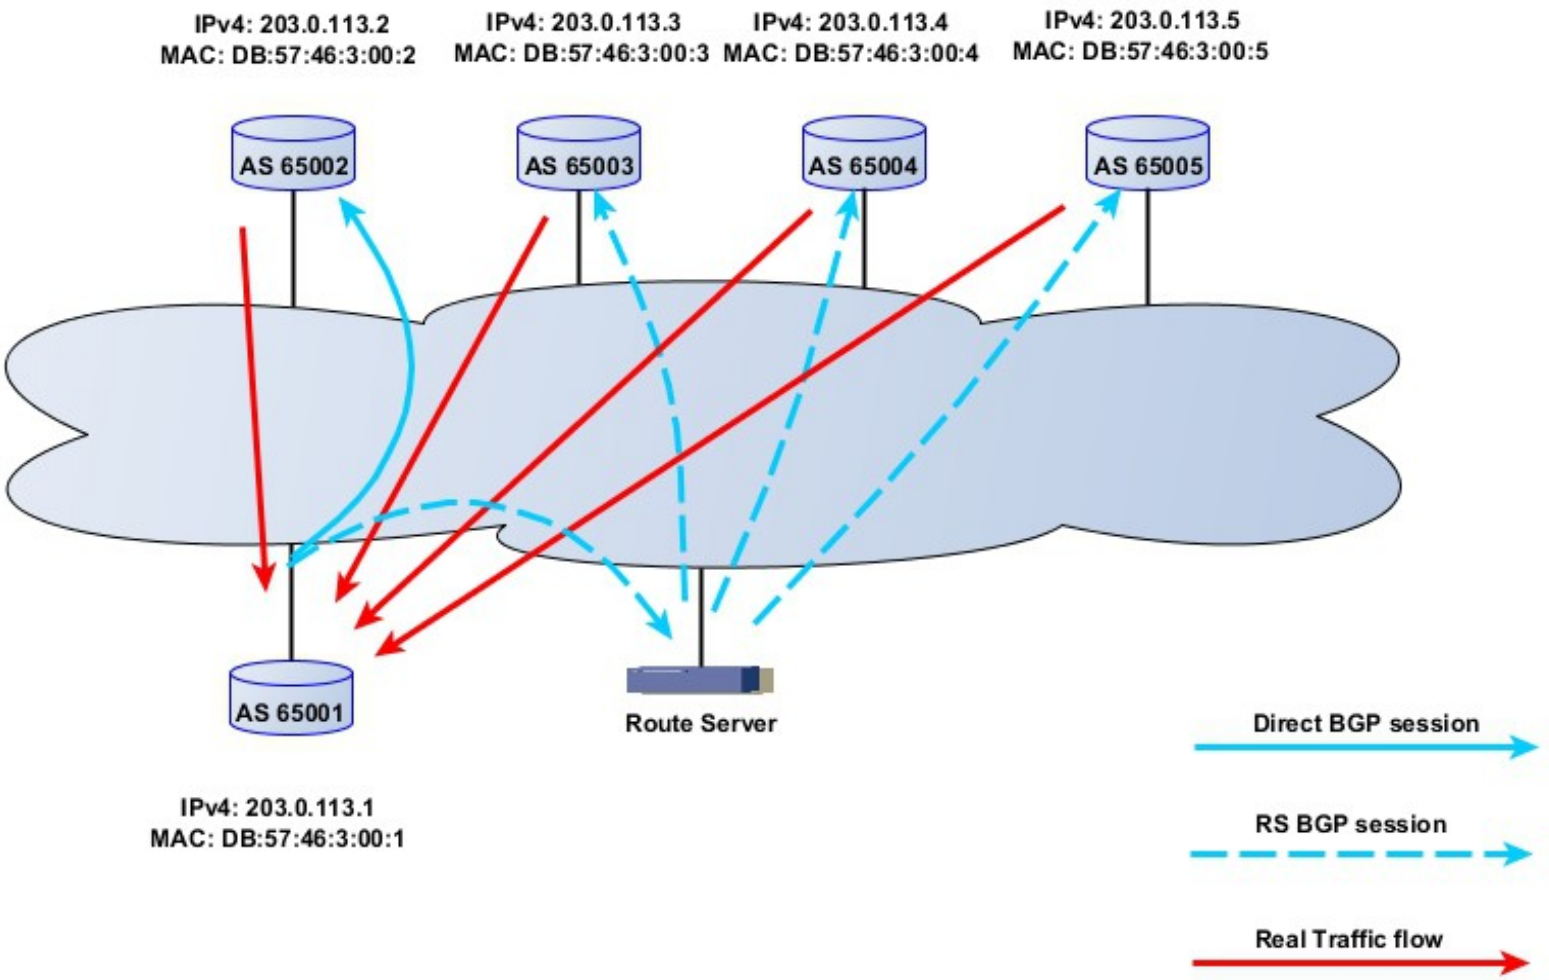
IPv4: 203.0.113.2
MAC: DB:57:46:3:00:2
IPv4: 203.0.113.3
MAC: DB:57:46:3:00:3
IPv4: 203.0.113.4
MAC: DB:57:46:3:00:4
IPv4: 203.0.113.5
MAC: DB:57:46:3:00:5
AS 65002
AS 65003
AS 65004
AS 65005
AS 65001
Route Server
IPv4: 203.0.113.1
MAC: DB:57:46:3:00:1
Direct BGP session
RS BGP session
Real Traffic flow

Example of an attack

- AS65001 has a destination/prefix under attack (198.51.100.1/24)
- AS65001 also announces other prefixes than the attacked one
- AS65002, AS65003, AS65004 originate traffic, which is part of the attack
- AS65005 doesn’t originate malicious traffic
- AS65002 is a direct peering of AS65001
- AS65003, AS65004 and AS65005 see AS65001 via the route server

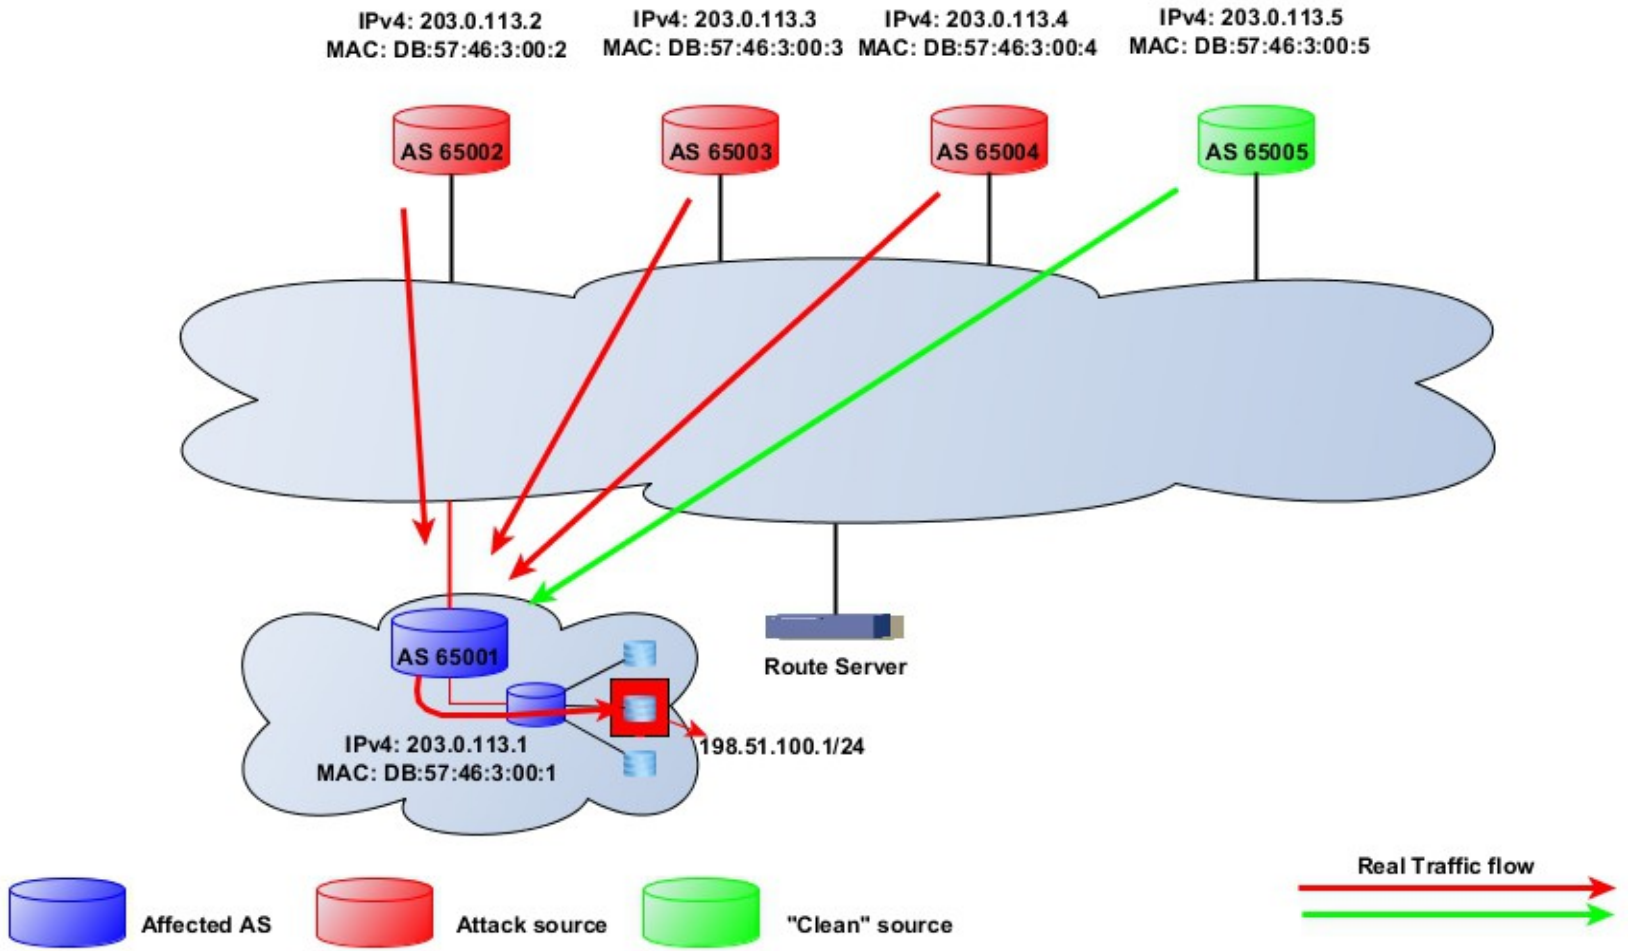
IPv4: 203.0.113.2
MAC: DB:57:46:3:00:2
IPv4: 203.0.113.3
MAC: DB:57:46:3:00:3
IPv4: 203.0.113.4
MAC: DB:57:46:3:00:4
IPv4: 203.0.113.5
MAC: DB:57:46:3:00:5
AS 65002
AS 65003
AS 65004
AS 65005
AS 65001
Route Server
IPv4: 203.0.113.1
MAC: DB:57:46:3:00:1
198.51.100.1/24
Affected AS
Attack source
"Clean" source
Real Traffic flow

NETIX

Considerations

- The attacked prefix (198.51.100.1/24) behind AS65001 is not fully reachable via the peering platform any more
- Other prefixes behind AS65001 might be affected by this attack as well port congestion
- flapping BGP sessions
- high router CPUs, etc ...
- AS65005 has a degraded reachability of 198.51.100.1/24 (even it is not part of the attack)

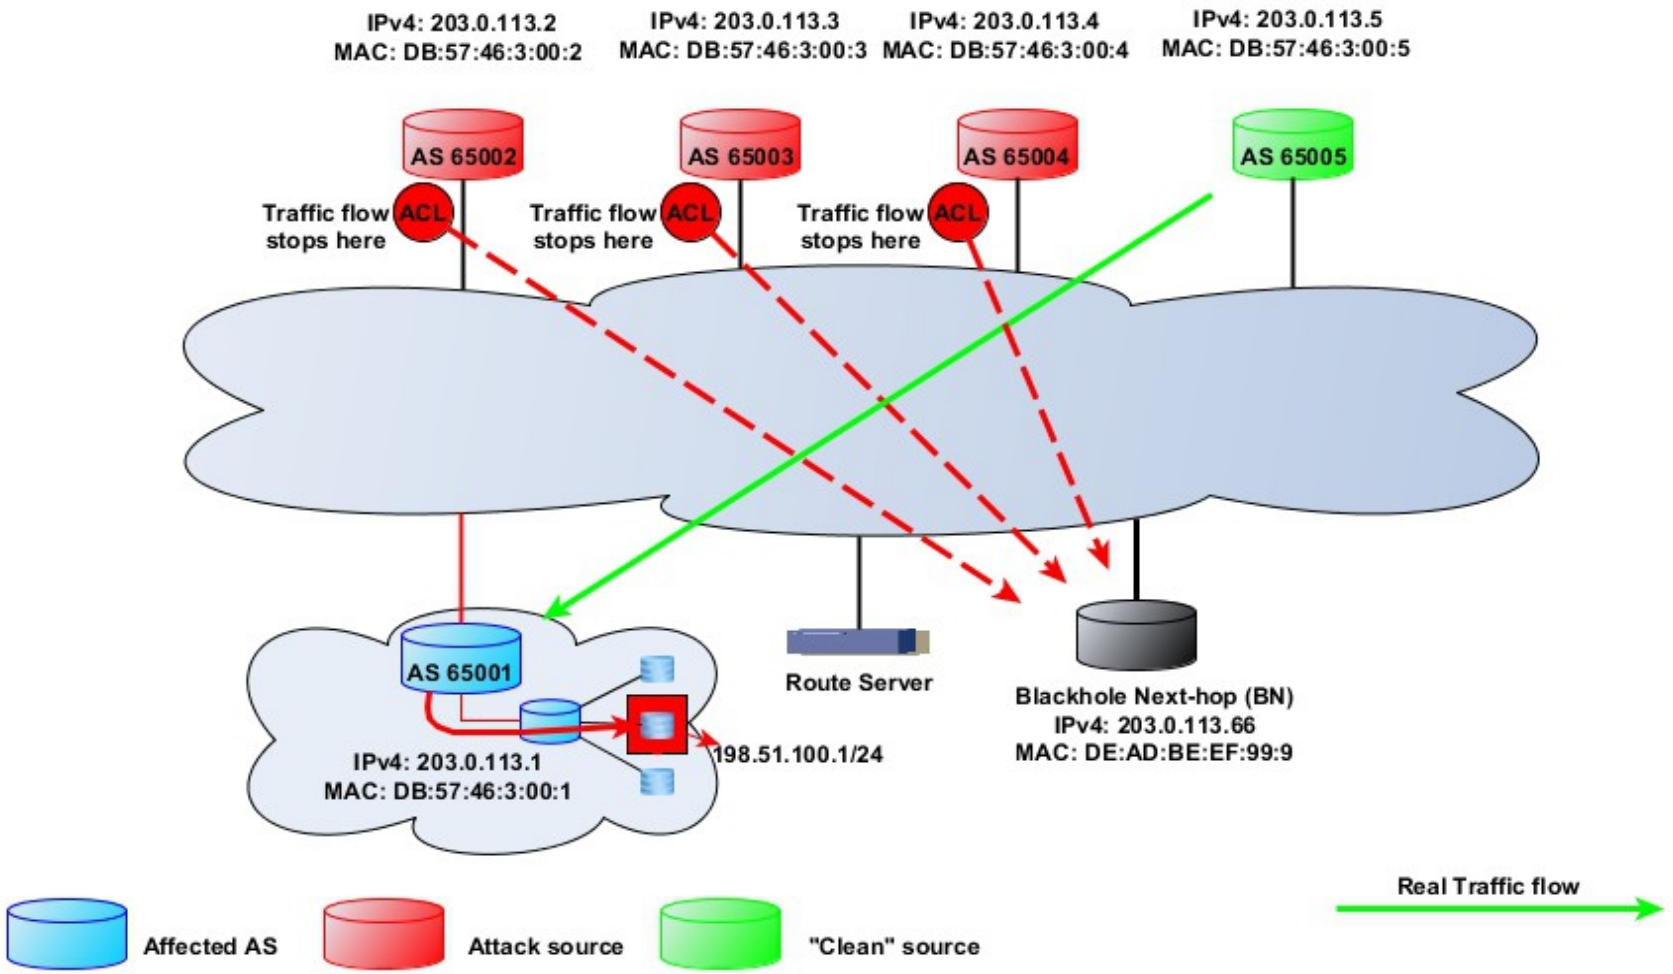
IPv4: 203.0.113.2
MAC: DB:57:46:3:00:2
IPv4: 203.0.113.3
MAC: DB:57:46:3:00:3
IPv4: 203.0.113.4
MAC: DB:57:46:3:00:4
IPv4: 203.0.113.5
MAC: DB:57:46:3:00:5
AS 65002
AS 65003
AS 65004
AS 65005
Traffic flow stops here
ACL
Traffic flow stops here
ACL
Traffic flow stops here
ACL
AS 65001
IPv4: 203.0.113.1
MAC: DB:57:46:3:00:1
198.51.100.1/24
Route Server
Blackhole Next-hop (BN)
IPv4: 203.0.113.66
MAC: DE:AD:BE:EF:99:9
Affected AS
Attack source
"Clean" source
Real Traffic flow

NETIX

Example Summary

- AS65001 selectively announce the attacked prefix with the Blackhole community (BC)
- All peers which select this new prefix as best-path, learn the BN's MAC address via ARP or ND provided by NetIX
- Traffic destined to the BN's MAC is dropped ingress via ACL closer to the source
- AS65001 has a chance to selectively blackhole traffic

Important Notes

- Traffic from all upstream host to the “blackholed” prefix is discarded
- Including the normal/non-malicious traffic
- Solution: if the origin ASN(s) from where the attack is coming is known, blackhole routes with appropriate BGP communities may be announced

Peer configuration example (IPv4, Cisco)

```
!
ip prefix-list blackholing seq 5 permit <blackholed prefix/32>
!
router bgp <your ASN>
no bgp enforce-first-as
neighbor <RS> remote-as <NetIX ASN>
!
address-family ipv4
network <blackholed prefix/32>
neighbor <RS> send-community
neighbor <RS> route-map Send_blackhole out
exit-address-family
!
route-map Send_blackhole permit 10
match ip address prefix-list blackhole
set community 65499:999
route-map Send_to_NetIX permit 20
```

NetIX Blackholing details:

BNv4: 193.218.0.99

BNv6: 2001:67c:29f0::9999

MAC: a0f3.c170.99a8

NETIX

How does the NetIX Blackholing Service work?

- L2 Filtering

  - Blackhole Next-hop (BN) has a unique MAC address (determined by ARP for the BN IP address)

  - All frames with destination MAC address belonging to the BN are filtered ingress by the L2 ACL applied on all member ports on NetIX switches

- In this case, all traffic to the “blackholed” prefix is discarded on the closest to the attacker switch, and hence the attacked member’s resources are protected

Contact us

- Sales team: sales@netix.net
- NOC team: noc@netix.net

NETIX