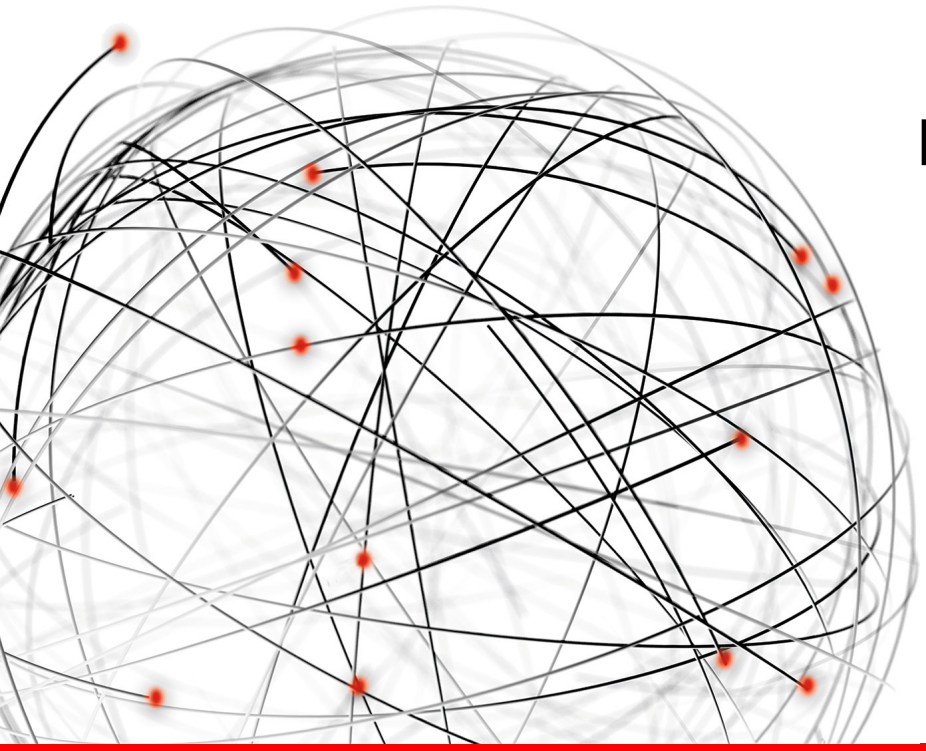

EQUINIX

Remote Triggered Black Hole

Equinix Internet Exchange Platform

Presented at SANOG 28

Date: 02.08.2016

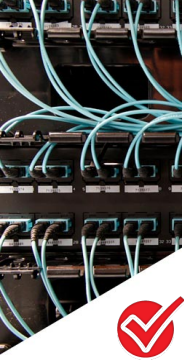

Remote Triggered Black Hole

Equinix Internet Exchange

- **Agenda**
  - DDoS Overview
  - RTBH Overview
  - How it works
  - Verification in Test Lab
  - Operations & Limitation
  - Configuration Guidelines
  - Questions & Answers

DDoS Overview

Equinix Internet Exchange

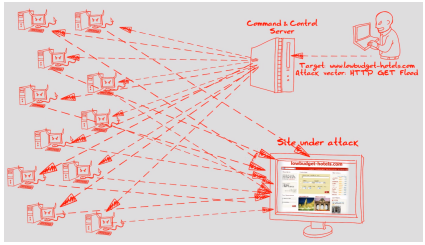

Overview

- DDoS is a type of DOS attack where multiple compromised systems, which are often infected with a Trojan, are used to target a single system causing a Denial of Service (DoS) attack.

Attack Class: Four common categories of attacks

- TCP Connection Attacks - *Occupying connections*
- Volumetric Attacks - *Using up bandwidth*
- Fragmentation Attacks - *Pieces of packets*
- Application Attacks - *Targeting applications*

Amplification: Two ways attacks can multiply traffic they can send.

- DNS Reflection – Small request, big reply
- Chargen Reflection – Steady Streams of text

Remote Triggered Black Hole

Equinix Internet Exchange

- **Overview**
  - Purpose
    - To provide self-managed black hole solution to EIE customers.
    - Enable fast responsive action to DDoS attack.
  - Deployment Status
    - RTBH has already been rolled out in APAC Region.

Remote Triggered Black Hole

Equinix Internet Exchange

- **How it works**

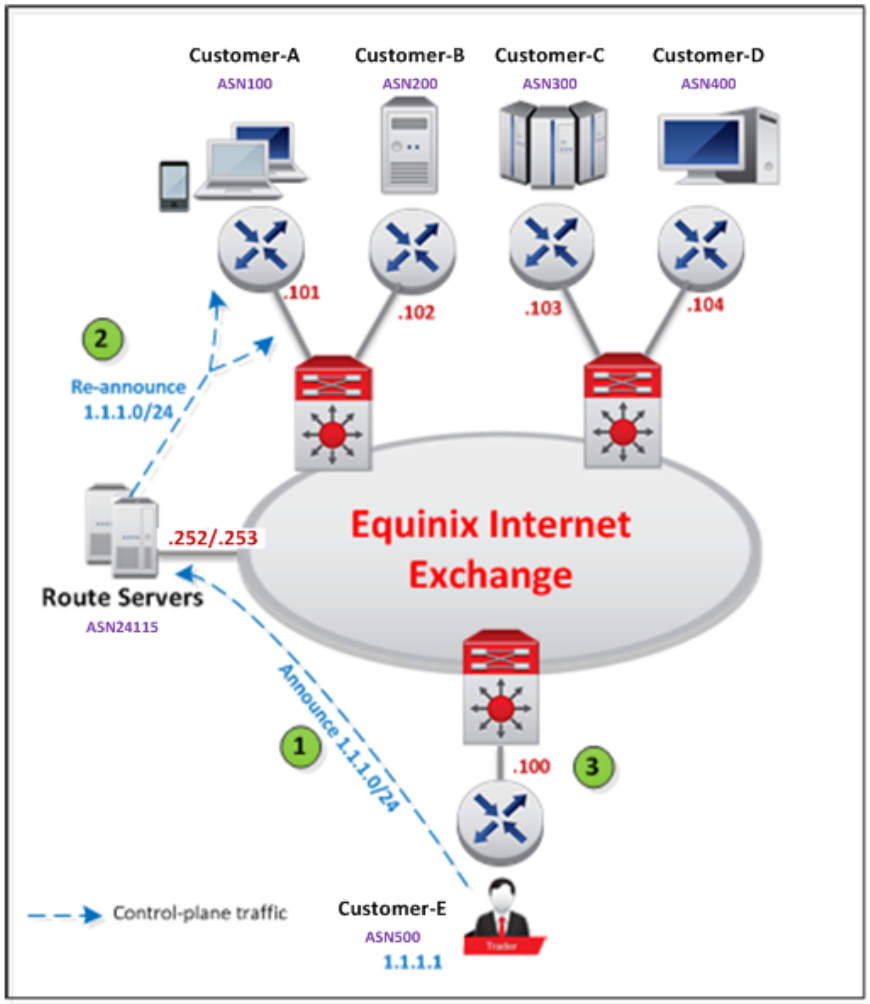

Initially:

1. All customers established BGP peering to route servers via MLPE IX peering subnet. The victim customer announced their prefix **1.1.1.0/24** to route servers.
2. Route servers re-announce it back to other peering participants.
3. The next-hop to reach 1.1.1.0/24 prefix is **.100** which is the peering IP address of victim customer.

Remote Triggered Black Hole

Equinix Internet Exchange

- **How it works**

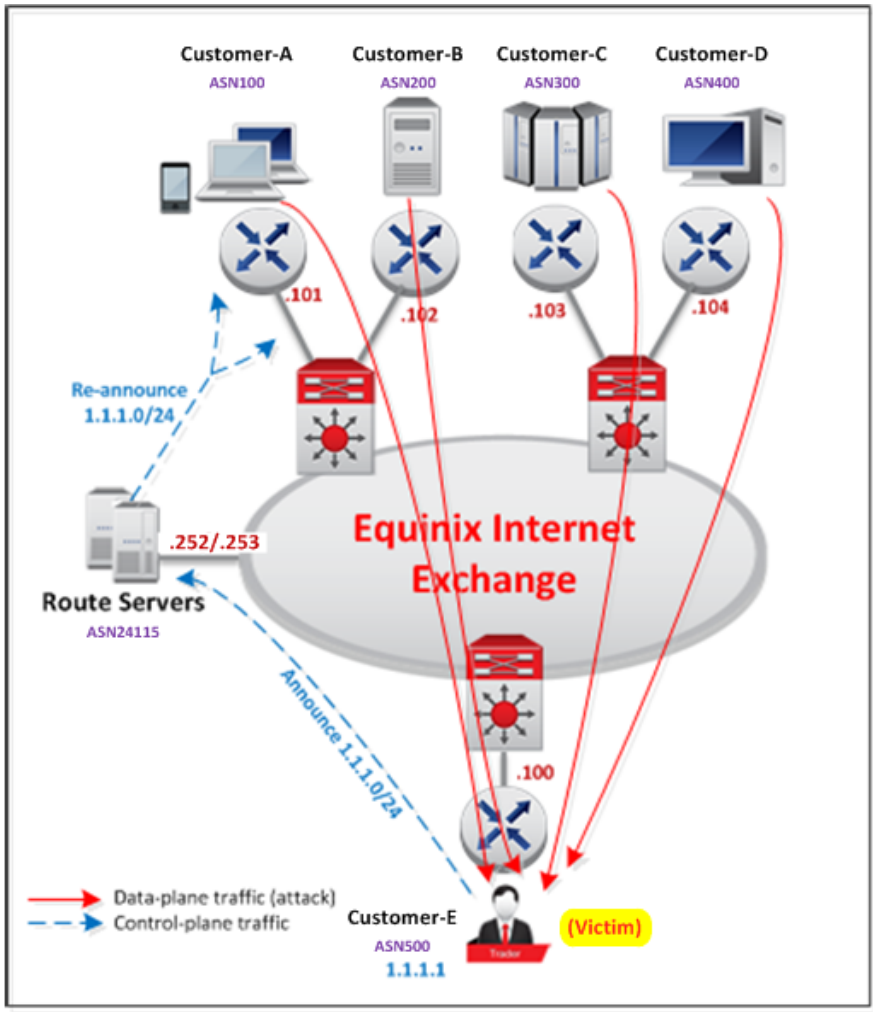

Under DDoS Attack:

- There is an DDoS attack traffic towards 1.1.1.1 which is one of the web servers.
- The victim's EIE port has been flooded and inbound (towards customer) utilization hits to 100%.
- This flooding caused service disruption to all production services.
- The victim needs to free up the port utilization by stopping traffic to 1.1.1.1.
  - » by RTBH.

Remote Triggered Black Hole

Equinix Internet Exchange

- **How it works**

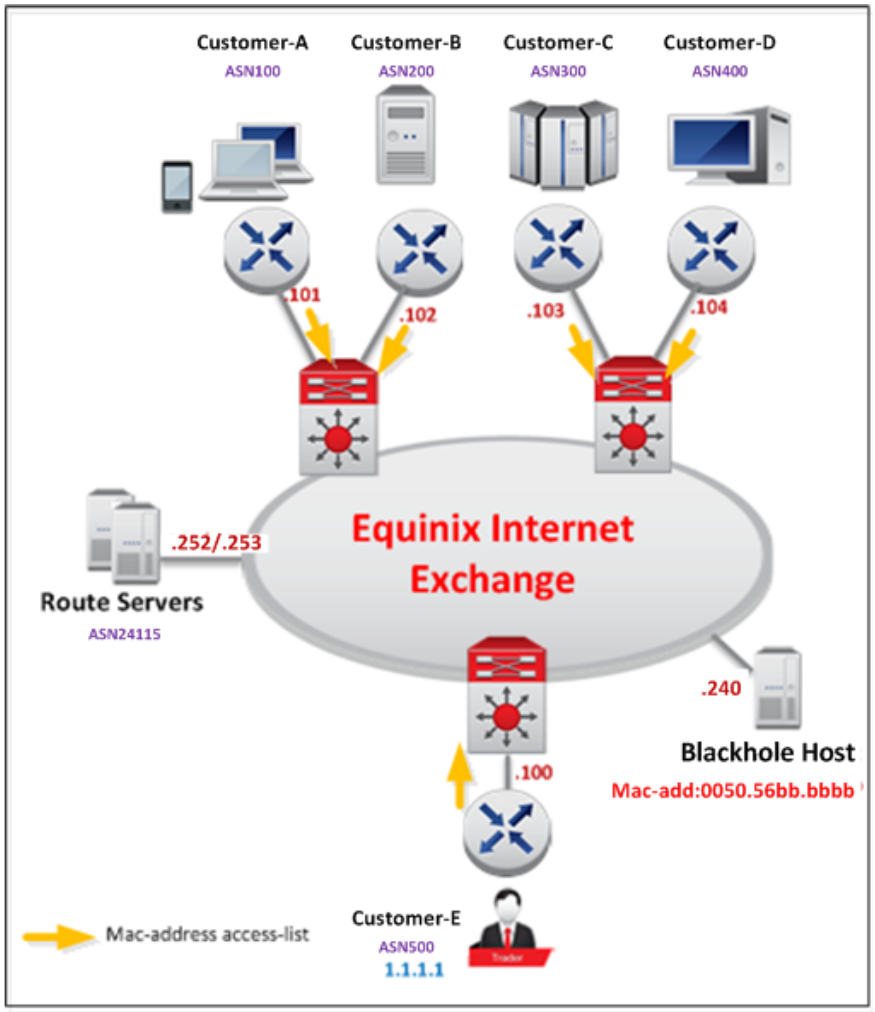

RTBH Setup:

- Equinix provides **Blackhole Host** with IP address **.240** whose mac address is **0050.56bb.bbbb.**
- All unicast traffic towards the Blackhole Host will be denied at customer facing ports (by mac-address ACL).

Remote Triggered Black Hole

Equinix Internet Exchange

- **How it works**

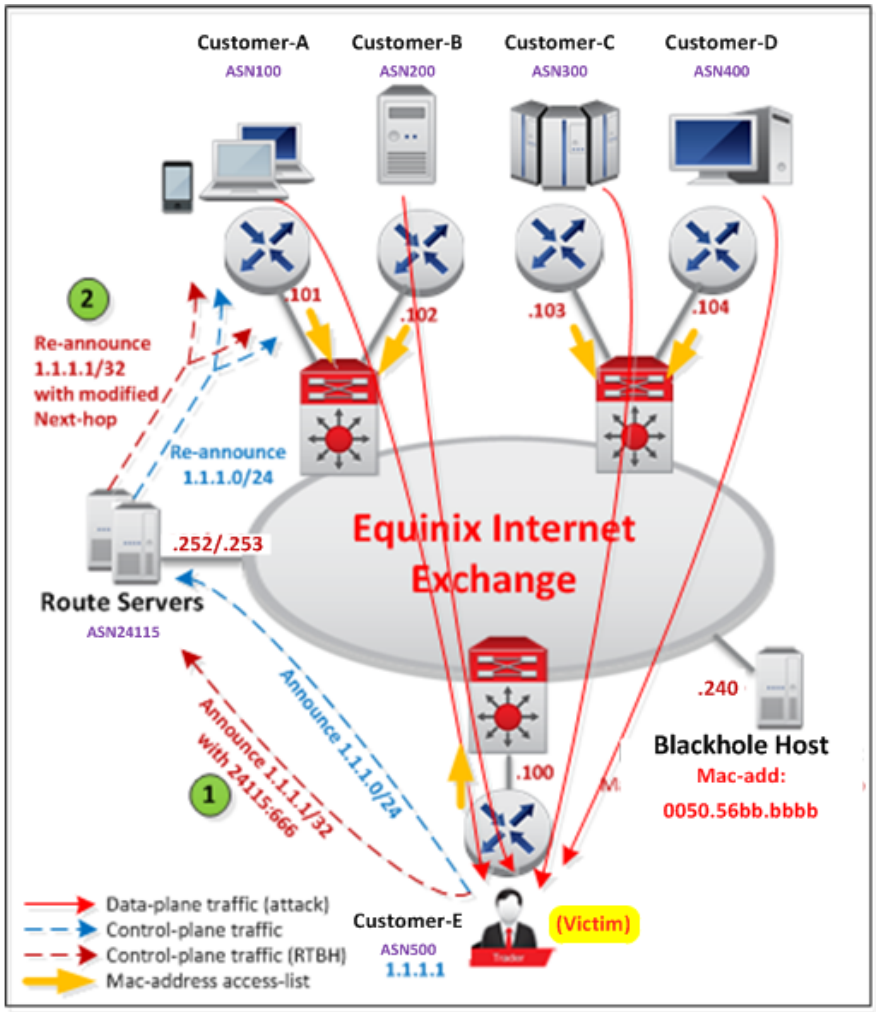

Mitigation:

1. The victim starts announce 1.1.1.1/32 with Blackhole BGP community **65535:666**.
2. Route servers will modified these prefixes (tagged with 65535:666) with it's next-hop to **.240** and re-announced back to other peering participants.

Remote Triggered Black Hole

Equinix Internet Exchange

- **How it works**

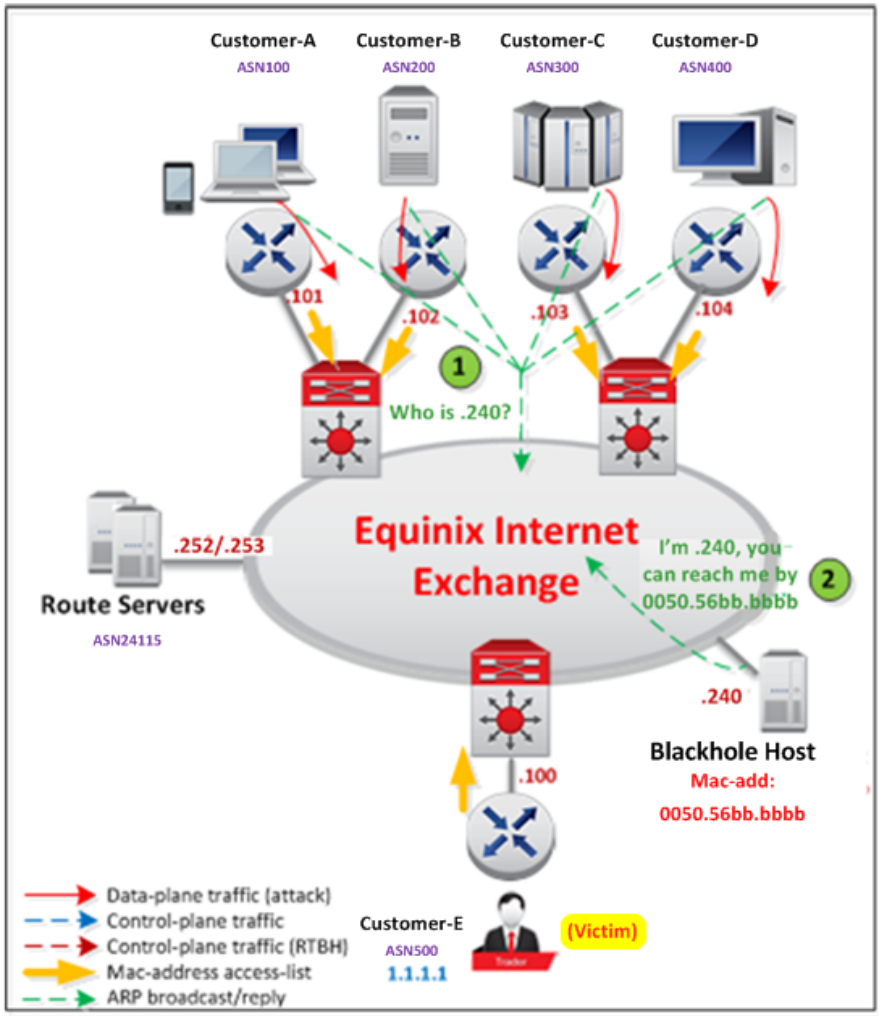

Mitigation (continued..):

1. Peering partners start to resolve next-hop IP address **.240** in order to reach 1.1.1.1.
2. Blackhole Host replied ARP with mac-address **0050.56bb.bbbb.**

Remote Triggered Black Hole

Equinix Internet Exchange

- **How it works**

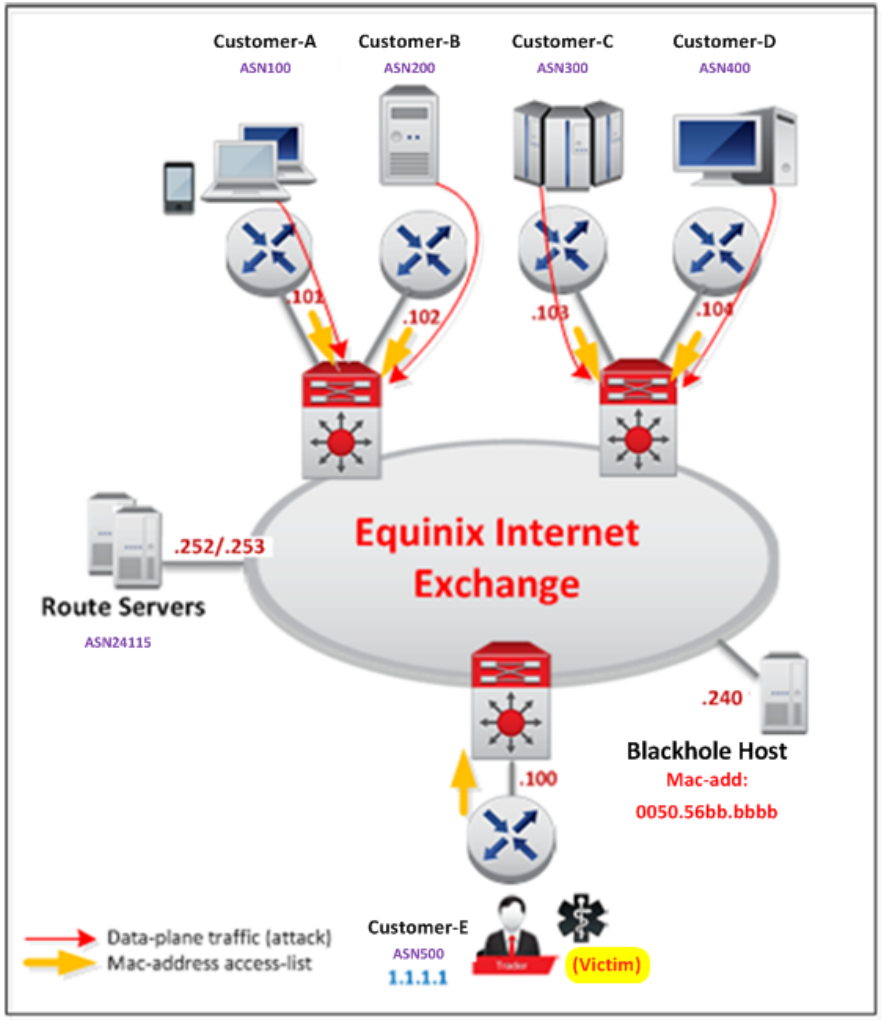

Mitigation (completed):

- The attack traffic with next-hop **.240** has been stopped by EIE switch **inbound access-list**.
- The victim's switch port has been mitigated and no more congestion.

Remote Triggered Black Hole

Equinix Internet Exchange

- **Verification**
  in the Lab

- Initially, Customer (Victim) is announcing their prefix 103.10.0.0/24 to route server via MLPE.
- Various Sources initiate DDoS attack to the destination IP, 103.0.10.1.
- The attack traffic has been forwarded to the destination port
- The victim port has been flooded with DDoS attack traffic.

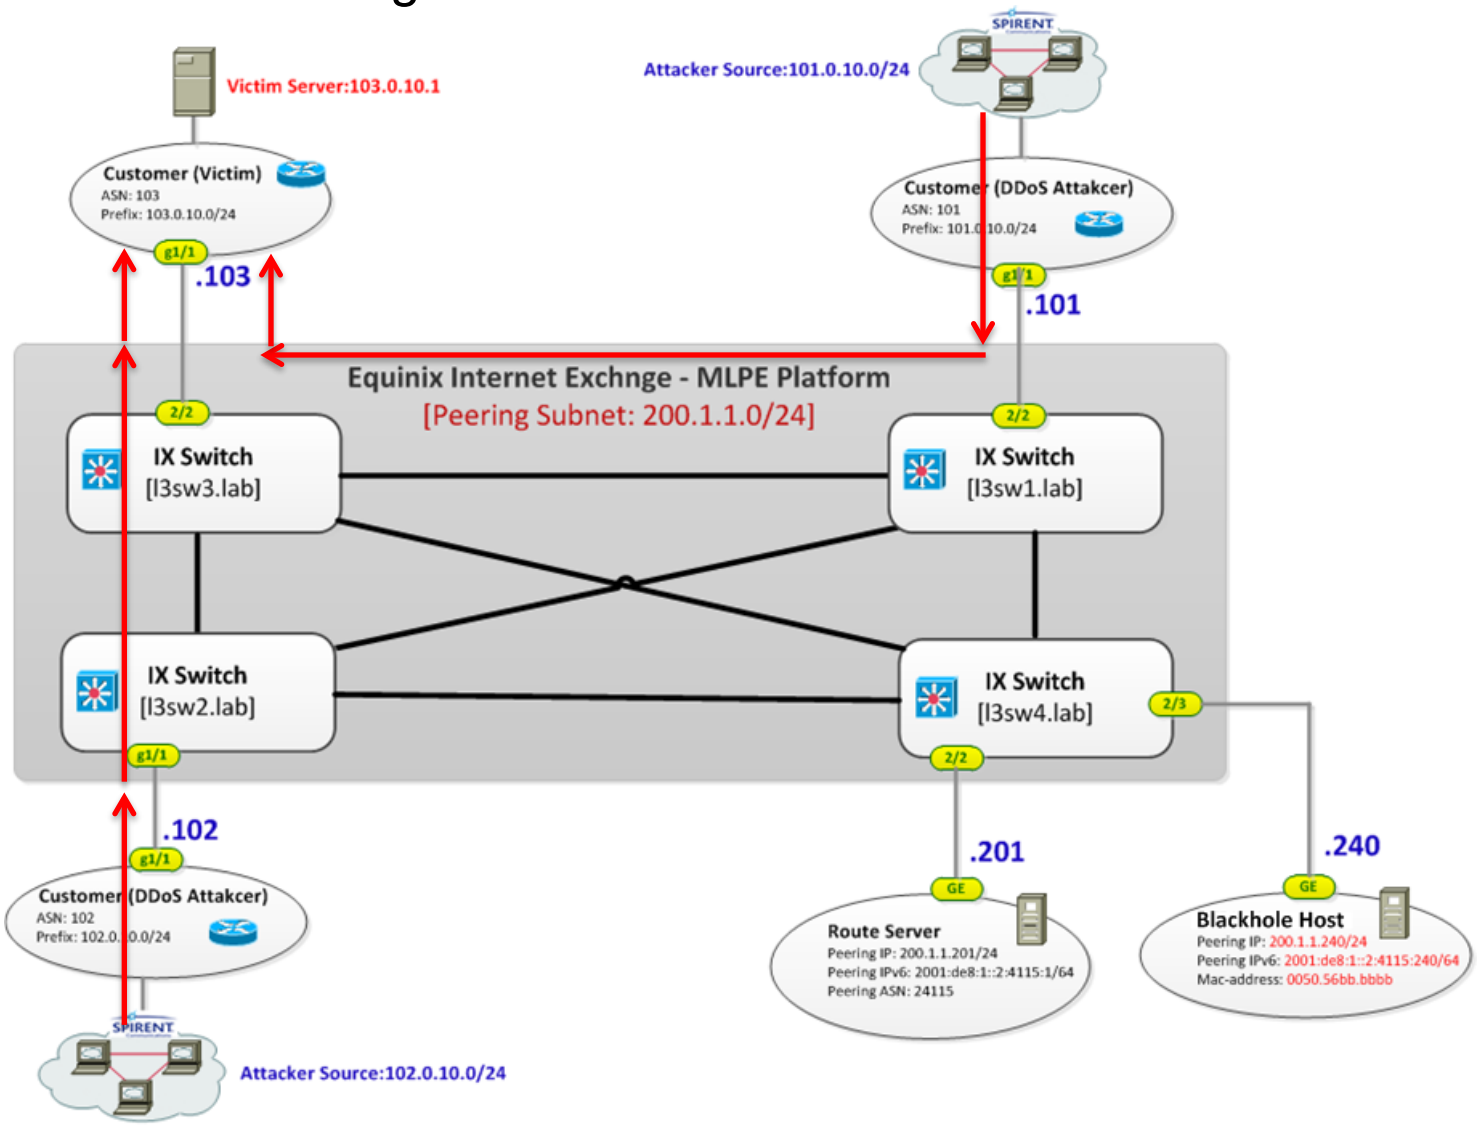

Remote Triggered Black Hole

Equinix Internet Exchange

- **Verification**
  in the Lab

- Initially, Customer (Victim) is announcing their prefix **103.0.10.0/24** to route server via MLPE.
- Other Peering Customers see their announcement with next-hop **200.1.1.103**
- Other Peering Customers is able to reach the victim IP address **103.0.10.1**

INITIAL STATUS

```
! CUSTOMER VICTIM
r4.lab#sh bgp vpnv4 unicast vrf EIE-TEST neighbors 200.1.1.201 advertised-routes | be Route
Route Distinguisher: 103:103 (default for vrf EIE-TEST)
*> 103.0.10.0/24    0.0.0.0                  0         32768 i
*> 103.0.20.0/24    0.0.0.0                  0         32768 i
*> 103.0.30.0/24    0.0.0.0                  0         32768 i
*> 103.0.40.0/24    0.0.0.0                  0         32768 i
*> 103.0.50.0/24    0.0.0.0                  0         32768 i

Total number of prefixes 5


! OTHER CUSTOMERS (ATTACKER)
r2.lab#sh ip route vrf EIE-TEST 103.0.10.1
Routing Table: EIE-TEST
Routing entry for 103.0.10.0/24
  Known via "bgp 102", distance 20, metric 0
  Tag 103, type external
  Last update from 200.1.1.103 00:10:28 ago
  Routing Descriptor Blocks:
  * 200.1.1.103, from 200.1.1.201, 00:10:28 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 103
      MPLS label: none


r2.lab#ping vrf EIE-TEST 103.0.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103.0.10.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 m
```

Remote Triggered Black Hole

Equinix Internet Exchange

- **Verification**
  in the Lab

- The victim customer starts announce 103.0.10.1/32 prefix with community 65535:666
- Route server received /32 prefix advertisement.
- The route server changed the next-hop to 200.1.1.240 and re-announce back to other peering participants.

MITIGATION STATUS

```
! VICTIM CUSTOMER
r4.lab#sh bgp vpnv4 unicast vrf EIE-TEST neighbors 200.1.1.201 advertised-routes | be Route
Route Distinguisher: 103:103 (default for vrf EIE-TEST)
*> 103.0.10.0/24    0.0.0.0                  0         32768 i
*> 103.0.10.1/32    0.0.0.0                  0         32768 i
*> 103.0.20.0/24    0.0.0.0                  0         32768 i
*> 103.0.30.0/24    0.0.0.0                  0         32768 i
*> 103.0.40.0/24    0.0.0.0                  0         32768 i
*> 103.0.50.0/24    0.0.0.0                  0         32768 i

Total number of prefixes 6

! ROUTE SERVER
[root@ixrs1 ~]# birdc show route protocol A200_1_1_103
BIRD 1.3.9 ready.
103.0.40.0/24      via 200.1.1.103 on eth1 [A200_1_1_103 13:54] * (100) [AS103i]
103.0.50.0/24      via 200.1.1.103 on eth1 [A200_1_1_103 13:54] * (100) [AS103i]
103.0.10.1/32      via 200.1.1.103 on eth1 [A200_1_1_103 13:54] * (100) [AS103i]
103.0.10.0/24      via 200.1.1.103 on eth1 [A200_1_1_103 13:54] * (100) [AS103i]
103.0.20.0/24      via 200.1.1.103 on eth1 [A200_1_1_103 13:54] * (100) [AS103i]
103.0.30.0/24      via 200.1.1.103 on eth1 [A200_1_1_103 13:54] * (100) [AS103i]

[root@ixrs1 ~]# birdc show route 103.0.10.1/32 all
BIRD 1.3.9 ready.
103.0.10.1/32      via 200.1.1.103 on eth1 [A200_1_1_103 13:54] * (100) [AS103i]
        Type: BGP unicast univ
        BGP.origin: IGP
        BGP.as_path: 103
        BGP.next_hop: 200.1.1.240
        BGP.med: 0
        BGP.local_pref: 100
        BGP.community: (24115,103) (65535,666)
        BGP.ext_community: (rt, 103, 103)
```

Remote Triggered Black Hole

Equinix Internet Exchange

- **Verification**
  in the Lab

- Attacker receive /32 prefix advertisement from route server and installed in the routing table.
- The next-hop address is pointing to **Blackhole host, 200.1.1.240.**
- Since the traffic towards Blackhole host has been filtered, the attack traffic has been stopped at IX entry points.

AFTER MITIGATION

```
! ATTACKER-1
r2.lab#sh ip route vrf EIE-TEST 103.0.10.1
Routing Table: EIE-TEST
Routing entry for 103.0.10.1/32
  Known via "bgp 102", distance 20, metric 0
  Tag 103, type external
  Last update from 200.1.1.240 00:12:33 ago
  Routing Descriptor Blocks:
  * 200.1.1.240, from 200.1.1.201, 00:12:33 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 103
      MPLS label: none
r2.lab#ping vrf EIE-TEST 103.0.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103.0.10.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

! ATTACKER-2
r3.lab#sh ip route vrf EIE-TEST 103.0.10.1

Routing Table: EIE-TEST
Routing entry for 103.0.10.1/32
  Known via "bgp 104", distance 20, metric 0
  Tag 103, type external
  Last update from 200.1.1.240 00:13:11 ago
  Routing Descriptor Blocks:
  * 200.1.1.240, from 200.1.1.201, 00:13:11 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 103
      MPLS label: none
r3.lab#ping vrf EIE-TEST 103.0.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103.0.10.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Remote Triggered Black Hole

Equinix Internet Exchange

Operations & Limitation

- **RTBH Users must Announce its Own Prefixes**
  - If the customer wishes to announce RTBH 1.1.1.1/32
    - 1.1.1.1 prefix must be subset of it's own announced prefixes.
- **RTBH Prefix Announcement must be Accepted by Other Peering Participants**
  - In order RTBH to take it effect, the black hole announcement must be accepted by other peering partners.
  - Need to advise all peering participants to accept the prefixes with (prefix length > /24 and BGP community **65535:666**).
  - It's totally voluntary setup.
- **MLPA & BLPA**
  - In MLPA environment, EIE route servers manipulate Blackhole Host next-hop IP address.
  - In BLPA environment, the two peering partners need to engage for defining BGP community and next-hop IP address.
  - BLPA partners still can use Blackhole Host next-hop mac-address (0050.56bb.bbbb) in their implementation.
- **Support for IPv6**
  - It supports IPv6.
  - Route Server will change next-hop IPv6 address to Blackhole Host (example 2001::2:4115:240) if the customer prefixes are tagged with BGP community **65535:666**.
  - Blackhole Host will response to IPv6 ND.

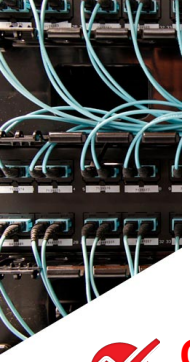

Remote Triggered Black Hole

Equinix Internet Exchange

Configuration Guidelines

– **Customers who Announce RTBH Route**

- The route has to be tagged with Blackhole BGP community 65535:666 and send to route servers.
- Make sure to send the community to peers (some router doesn't send by default).
- Clear BGP outbound (if it's necessary).
- Sample Configuration

To accept RTBH announcements by other peers

Sending RTBH announcement for 1.1.1.1/32

```
router bgp <asn>
 network 1.1.1.1 mask 255.255.255.255
 neighbor <route-server> remote-as <REMOTE-AS-Number>
 neighbor <route-server> route-map RM-RTBH-IN in
 neighbor <route-server> route-map RM-RTBH-OUT out
 neighbor <route-server> send-community
!
route-map RM-RTBH-IN permit 10
 match ip address prefix-list PL-RTBH-IN
 match community 65535:666
!
route-map RM-RTBH-IN deny 20
 match ip address prefix-list PL-RTBH-IN
!
route-map RM-RTBH-IN permit 30
!
ip prefix-list PL-RTBH-IN seq 10 permit 0.0.0.0/24 le 32
!
route-map RM-AS103-V4-OUT permit 10
 match ip address prefix-list PL-RTBH-OUT
 set community 65535:666
!
ip prefix-list PL-RTBH-OUT seq 10 permit 1.1.1.1/32
ip route 1.1.1.1 255.255.255.255 Null0
```

➢ The sample configuration is just for guideline, not to be used as it is.

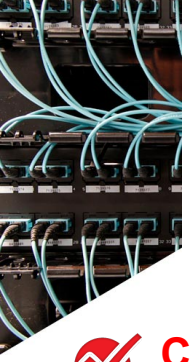

Remote Triggered Black Hole

Equinix Internet Exchange

Configuration Guidelines (cont..)

– **Customers who Receive RTBH Routes**

- To accept > /24 Prefixes which are tagged with BGP community 65536:666.
  - Reference: https://tools.ietf.org/html/draft-ymbk-grow-blackholing-01
- To tag 'No-Export', 'No-Advertise' communities to RTBH prefix once it's received.
- Suggested not to re-advertise back to any upstream, downstream or peers.

- Sample Configuration as below:

```
router bgp <asn>
 neighbor <route-server> remote-as <REMOTE-AS-Number>
 neighbor <route-server> route-map RM-IXRX-IN in
 neighbor <route-server> send-community
 neighbor <ibgp-neighbor> remote-as <asn>
!
route-map RM-IXRS-IN permit 10
 match community 65535:666
 set community 65535:65281 additive   <- set no-export community
 set ip next-hop 192.168.0.1          <- set to next-hop with RFC1918 address
!
ip route 192.168.0.1 255.255.255.255 Null0   <- Point next-hop to local Null Interface
```

➢ The sample configuration is just for guideline, not to be used as it is.

Remote Triggered Black Hole

Equinix Internet Exchange

Bi-Lateral Peering Note

- Note
  - Some peering partners may not accept IX Well-known Blackhole Community (65535:666).
  - So BLPA peers must bilaterally agree on their own defined Blackhole Community + the router configuration to accept /32 announcement for that community.
  - Also customers are recommended to implement a route-map in which prefixes with blackhole community are changed its next-hop to local Null0 interface and re-advertise back to their peers.

The Remote Triggered Black Hole (DDOS blocking) is based on the standard described in these evolving feature documents (use the last one of these, as it is most current):

- https://tools.ietf.org/id/draft-ymbk-grow-blackholing-00.txt
- https://tools.ietf.org/id/draft-ymbk-grow-blackholing-01.txt
- https://tools.ietf.org/id/draft-ietf-grow-blackholing-00.txt
- https://tools.ietf.org/id/draft-ietf-grow-blackholing-01.txt

Remote Triggered Black Hole

Equinix Internet Exchange

Questions?

mabdulrasheed@ap.equinix.com