

Remote-triggered Blackholing for network operators

White paper

Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks are becoming more and more harmful. Not only the targets of such attacks suffer, but also the whole infrastructure of a network operator can be affected. Remote triggered Blackholing is an instrument that can be used to stop DoS and DDoS attacks to specific targets in your network. You need to set up an infrastructure for this well before an attack happens. In this white paper for network engineers, we give you an example of how to implement a remote-triggered Blackholing.

Definitions

Distributed Denial of Service attack (DDoS) is an attack against a system via the Internet. The attacker uses multiple (sometimes millions of) network sources to send more traffic towards the attacked system than the system can handle. The collateral damage is quite often the network infrastructure to which the attacked system is connected.

Blackholing is a method for discarding unwanted or malicious traffic. Instead of forwarding unwanted packets to their destination, they are discarded as early as possible.

The goal of Blackholing

If DoS or DDoS packets get dropped as early as possible, the target system is no longer reachable for those packets and there is no collateral damage. Or at least the collateral damage is kept to a minimum. The goal is to drop the packets outside the attacked network, or if this is not possible, at the earliest possible stage within the attacked network. In this white paper we share an example of how to use Blackholing to:

- drop malicious packets within the network at all possible places,
- signal to neighboring networks to do the same.

How to set it up

Being under attack is stressful. Once you have determined the target of an attack and know which IP address(es) should be blackholed, you want to be able to start (and stop) Blackholing quickly and easily, and monitor the success of the process.

Example network

The example network in this white paper is a small ISP, with three iBGP-speaking routers and two upstream providers, that peer at DE-CIX Frankfurt. All IP addresses and autonomous system numbers (ASNs) used in the examples are fictional – please do not copy and paste the examples, but instead use your own IP addresses.

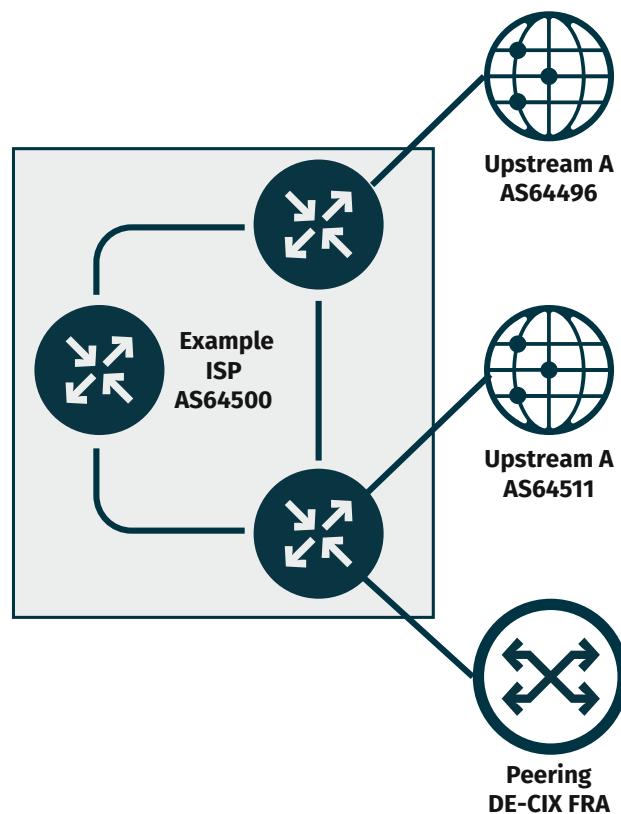
The injector devices

Using a router

You need to have something in place already to import prefixes into BGP. The idea is to use the same mechanism to import prefixes to be blackholed into BGP. In general, there are two ways of importing prefixes into BGP in Cisco IOS:

- 1) using a network statement
- 2) using “redistribute” with a route-map for filtering

In the case that you are using network statements, you must also use a route-map to set the necessary parameters, like BLACKHOLE community, and you need a corresponding route in your router’s routing table. Configuration example for Cisco IOS:



```
ip route 10.1.1.3 255.255.255.255 Null0
!
route-map set-blackholing permit 1000
  set community 65535:666 additive
!
router bgp 64500
  network 10.1.1.3 mask 255.255.255.255 route-map set-blackholing
```

The route-map can be pre-configured at any time, but the network statement in BGP and the static IP route must both be entered to activate Blackholing, and both removed to disable it. (It would be enough to remove one of the statements for deactivation, but after some time your router config would look very messy.)

If you redistribute static routes into BGP using some sort of filtering, you can easily extend this to also accommodate Blackholing. Again, you need a static route in your routing table, and here we also add a tag statement to it:

```
ip route 10.1.1.3 255.255.255.255 Null0 tag 666
!
route-map static-to-bgp permit 1000
  match tag 666
  set community 65535:666 additive
!
route-map static-to-bgp deny 65000
!
router bgp 64500
  redistribute static route-map static-to-bgp
```

Note that the route-map static-to-bgp is not complete – you also need statements to add regular (not to be blackholed) routes to BGP. Using a “tag” statement is very elegant, as you only have to add one line of configuration to start Blackholing.

You can achieve this also without a “tag”, by using an access-list or prefix-list. But in this case, you would have to add/remove two lines of config, and therefore we would not really recommend it.

```
ip route 10.1.1.3 255.255.255.255 Null0
!
ip access-list standard blackhole-list
  permit 10.1.1.3
!
route-map static-to-bgp permit 1000
  match ip address blackhole-list
  set community 65535:666 additive
!
router bgp 64500
  redistribute static route-map static-to-bgp
```

For IPv6 this is very similar. Please see the recommended version below, but note that you can use the same route-map as for IPv4:

```
ipv6 route 2001:DB8:0:1::1/128 Null0 tag 666
!
route-map static-to-bgp permit 1000
  match tag 666
  set community 65535:666 additive
!
route-map static-to-bgp deny 65000
!
router bgp 64500
  address-family ipv6
    redistribute static route-map static-to-bgp
!
```

Using ExaBGP

Instead of a router, you can also use the open-source tool ExaBGP to inject the prefixes. For this, you need to connect the server where you run ExaBGP to one of your routers as a route reflector. Another best practice in BGP is to only allow in or out what is really needed. So we only allow in IPv4 /32s with the Blackholing community already set. Incoming we set the IP next-hop address to an address routed to Null0 on all routers. The config below is for the router that the ExaBGP Blackholing server is directly connected to:

```
interface GigabitEthernet1/0
  description to ExaBGP Server
  ip address 192.168.2.13 255.255.255.252
!
ip route 192.168.66.66 255.255.255.255 Null0
!
route-map send-nothing deny 100
!
ip access-list extended only-32s
  permit ip any host 255.255.255.255
!
ip community-list standard blackhole permit 65535:666
!
route-map only-blackhole permit 100
  match ip address only-32s
  match community blackhole
  set ip next-hop 192.168.66.66
!
router bgp 64500
  neighbor exabgp peer-group
  neighbor exabgp remote-as 64500
  neighbor exabgp route-reflector-client
  neighbor exabgp soft-reconfiguration inbound
  neighbor exabgp route-map only-blackhole in
  neighbor exabgp route-map send-nothing out
  neighbor 192.168.2.14 peer-group exabgp
!
```

On the ExaBGP server, you need to define a static route entry for each prefix to be blackholed, as in the example below. Note that ExaBGP has many more features – this is just a basic configuration example for one prefix to be blackholed:

```
neighbor 192.168.2.13 {
  router-id 192.168.2.14;
  local-address 192.168.2.14;
  local-as 64500;
  peer-as 64500;

  static {
    route 10.1.1.1/32 {
      community [ 65535:666 ];
      next-hop 192.168.66.66;
    }
  }
}
```

Your iBGP-speaking routers

Blackholing means that you want to sink traffic everywhere and as early as possible. Bad packets towards the target under-attack are not necessarily from a source outside of your own network. Compromised customer devices can also be part of the attack, and packets from these must be discarded just like packets from the outside.

Because of this, you will want to implement similar mechanisms to those on your network borders also inside your network. The trick here is to set a next-hop address for all prefixes to be blackholed and route this next-hop address to the Null0 interface. The example below lists both IPv4 and IPv6 configurations. For this next-hop we are using 192.168.66.66/32 in IPv4 and 2001:DB8:666::666/128 in IPv6 as examples. When implementing this, please use IP addresses from within your own range.

We also make sure the Null0 interface does not send back any ICMP destination unreachable packets:

```
ip route 192.168.66.66 255.255.255.255 Null0
ipv6 route 2001:DB8:666::666/128 Null0
!
interface Null0
 no ip unreachable
 no ipv6 unreachable
!
ip community-list standard blackhole permit 65535:666
!
route-map internal-in-ipv4 permit 100
 match community blackhole
 set ip next-hop 192.168.66.66
!
route-map internal-in-ipv4 permit 10000
!
route-map internal-in-ipv6 permit 100
 match community blackhole
 set ipv6 next-hop 2001:DB8:666::666
!
route-map internal-in-ipv6 permit 1000
!
```

We now apply this route-map to our iBGP sessions. Again, the example shows both IPv4 and IPv6:

```
router bgp 64500
 no bgp default ipv4-unicast
 neighbor internal peer-group
 neighbor internal remote-as 64500
 neighbor internal update-source Loopback0
 neighbor internalv6 peer-group
 neighbor internalv6 remote-as 64500
 neighbor internalv6 update-source Loopback0
 address-family ipv4
  neighbor internal send-community both
  neighbor internal next-hop-self
  neighbor internal route-map internal-in in
 exit-address-family
!
 address-family ipv6
  neighbor internalv6 send-community both
  neighbor internalv6 next-hop-self
  neighbor internalv6 route-map internal-in-ipv6 in
 exit-address-family
```

This results in the following:

- All prefixes received where the BLACKHOLE community is set get their next-hop IP address changed;
- This changed next-hop points to a Null0 interface
- The router therefore discards all traffic to this IP address;
- This happens on every iBGP-speaking router in your network.

eBGP sessions to your upstreams

Once your internal devices “know” what prefixes to blackhole, you can also announce these prefixes to your upstreams. The only thing you need to do is to prepare the configuration of your filters in your eBGP-speaking routers. If your upstreams also support the well-known BLACKHOLE BGP community, you only have to set your filters so that all prefixes tagged with BLACKHOLE are also announced externally. Make sure that you announce only your own and your customers’ prefixes (do not re-announce prefixes you have received from a peer for Blackholing).

If you have to set a special BGP community to tell one of your upstreams to blackhole, you have to configure this in your outgoing route-map. Be aware that most of the time the to-be-blackholed prefixes are very short ones, so the filters for Blackholing have to be processed before the filters which block prefixes that are too short.

Below is an example code in Cisco IOS if you only pass along to well-known BLACKHOLE community. The example upstream provider here is AS64496. Here, we allow only the /16 address block out for our regular BGP announcement, but also smaller subnets up to /32 in case BLACKHOLE is set:

```
ip prefix-list my-prefixes seq 5 permit 10.1.0.0/16
ip prefix-list my-blackholes seq 5 permit 10.1.0.0/16 ge 32
!
route-map upstream-out permit 500
  match ip address prefix-list my-blackholes
  match community blackhole
  set community no-export additive
!
route-map upstream-out permit 1000
  match ip address prefix-list my-prefixes
!
route-map upstream-out deny 65000
!
router bgp 64500
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor upstream peer-group
  neighbor 172.16.2.1 remote-as 64496
  neighbor 172.16.2.1 peer-group upstream
  address-family ipv4
    neighbor upstream send-community both
    neighbor upstream route-map upstream-out out
  !
```

Below is example code in Cisco IOS if you have to set a special community (in this example 64511:666). The example upstream provider here is AS64511:

```
ip prefix-list my-prefixes seq 5 permit 10.1.0.0/16
ip prefix-list my-blackholes seq 5 permit 10.1.0.0/16 ge 32
!
route-map upstream-out permit 500
  match ip address prefix-list my-blackholes
  match community blackhole
  set community no-export 64511:666 additive
!
route-map upstream-out permit 1000
  match ip address prefix-list my-prefixes
!
route-map upstream-out deny 65000
!
router bgp 64500
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor upstream peer-group
  neighbor 172.17.2.1 remote-as 64511
  neighbor 172.17.2.1 peer-group upstream
  address-family ipv4
    neighbor upstream send-community both
    neighbor upstream route-map upstream-out out
!
```

eBGP sessions to DE-CIX peering

Similar to your upstreams, you can also signal to your peers to blackhole traffic. DE-CIX supports this by sinking traffic with a **special next-hop IP address** set at its ingress interfaces if your peers do not already block the traffic inside their networks.

As with your upstreams, for this to work it is important that your peers accept the IPv4 /32s and IPv6 /128s you announce to them via BGP. The example below only shows the route-map part which you need for Blackholing at DE-CIX Frankfurt. The rest of the configuration is the same as the eBGP to your upstreams:

```
route-map peering-out permit 500
  match ip address prefix-list my-blackholes
  match community blackhole
  set community no-export additive
  set ip next-hop 80.81.193.66
```


How to use Blackholing

Be prepared

At least once every half year you should schedule an emergency exercise. How you do this depends on your organizational and network structure, but it should involve everyone in your operational departments who would also be involved in case of a real attack. You can have this scheduled and announced to your teams well beforehand (recommended for the first few exercises) or as a “surprise” (not recommended if you do not have a well-trained team).

All documentation on how to start Blackholing and how to monitor it should be kept up to date (a good idea is to check after each emergency exercise whether the documentation still matches reality) and should be easily accessible for your team. This might include keeping a printed version available, or having a PDF document on your teams’ phones. Keep in mind that attacks do also happen during the night and on weekends, when your staff might not be in the office.

Also, you need to make sure that your management network (the network you use to configure your routers) is separate from your production network and is shielded from attacks. You have to be able to initiate the Blackholing after the attack has started.

When under attack

Use your prepared plan to initiate Blackholing. An example plan might read like this:

1. Find out what the target is. Sounds easy, but if multiple attacks happen at the same time, this might be challenging.
2. Initiate Blackholing of the targets IP address(es). This will:
 - a. sink the attack traffic within your network as early as possible,
 - b. signal your upstream provider(s) and peers to blackhole at their side.
3. Notify your customer(s)! Let your customer(s) know that they are under attack and that you have taken steps.
4. Check if the Blackholing is effective. Not all of your upstreams or peers might honor the Blackholing request. Talk to your upstreams and peers who do not. If the attack is extremely severe and still hurting your network, be prepared to shut down connections to parties who still send you attack traffic. This should only be the last option. Talk to your peers. Everybody has experienced attacks and they might be able to help.
5. Monitor the attack. If it subsides, stop Blackholing (but be prepared to re-initiate it if the attack increases again).

DE-CIX Blackholing guide

More information about how Blackholing works at DE-CIX can be found in the DE-CIX Blackholing guide available [on our website](#).

What you can achieve

If you set up your infrastructure accordingly, well before an attack happens, remote-triggered Blackholing is a very good instrument to stop DoS and DDoS attacks to specific targets in your network. Having such an infrastructure will help you to prevent collateral damage.

Questions?

Check out the Blackholing guide **on our website** and find out more **about our Blackholing service**.

About DE-CIX

DE-CIX is the world's largest peering and interconnection provider operating several carrier and data center neutral interconnection platforms in Europe, the Middle East, North America, and Asia. Founded in 1995, DE-CIX's Internet Exchange in Frankfurt is the world's leading interconnection platform, managing more than 9 Terabits per second peak traffic. With DE-CIX Apollon, we maintain the world's largest and most advanced Ethernet-based platform, delivering high-availability peering with full 100G Ethernet capabilities. Find out more at **www.de-cix.net**.

Get in touch

Phone: +49 69 1730902-12
Email: sales@de-cix.net