BGP Blackholing Reconsidered

Job Snijders
NTT Communications
job@ntt.net

What is blackholing? (high level overview)

Blackholing is signaling to adjacent networks that you don't wish to receive traffic for a given destination.

The common use case is that you take one IP address offline, in order to alleviate (potential for) congestion on a circuit, so that other destinations continue to be reachable.

How is blackholing implemented?

Blackholing usually is implemented through BGP signaling:

- **In-band**
  - (through the EBGP session that accompanies the IP Transit circuit)
- **out-of-band**
  - (to a blackhole server via a multi-hop EBGP session).

An operator announces a host route with a specific BGP community to the adjacent network; the adjacent network rewrites the NEXT_HOP to an IP address that is "*null routed*".

Downsides of RFC 3882 blackhole implementation

- The *request* for blackholing also becomes the best path: a /32 or /128 host route usually is the most specific route to a destination.
- On most implementations you need two prefix-list filters per customer: one for normal BGP announcements, another one "le 32" for blackhole routes. (Today's largest router config in NTT's network is 57 megabyte)
- NTT has observed both malicious and accidental blackhole requests

Mock-up routing policy

```
prefix-set AS15562-blackholes
  192.147.168.0/24 ge 32 le 32
end-set

prefix-set AS15562-routes
  192.147.168.0/24
end-set

if (destination in AS15562-blackholes && community 65535:666)
  set next-hop blackhole
  exit-policy
if (destination in AS15562-routes)
  accept
  exit-policy
end-policy
```

Reconsidering blackholing:

Only blackhole traffic...

iff the next less specific route is the best and active path

Blackhole validation procedure

Compare the left most SEQUENCE of the AS_PATH between the second less specific best path and the blackhole request

Or….

Compare whether the NEXT_HOP attribute of the covering route has the same value as the blackhole request

Or…. [insert other ideas]

Implications of only blackholing on active paths

When an IP Transit customer signals “Please don’t send me traffic for this IP address”, and they ain’t the best path for that destination *anyway*, you did your job!

Honoring blackholes will depend on:

- are you the shortest AS_PATH?
- Do you have the highest LOCAL_PREF?
- Did you register the destination in IRR?
- Does the covering less-specific route pass through an RPKI Origin Validation “invalid == reject” policy?
- ARE YOU THE BEST PATH?!

Blackholing Reconsidered “Listener channels”

Listen for blackhole requests via:

- **API**
- **In-band BGP signaling**
  - BMP Adj-RIB-In Pre-Policy
  - BGP no-fib-install, only export to a special “*blackhole validator*”
- **Out-of-band BGP signalling (central collectors)**

Mock-up routing policy – new world order

```
prefix-set AS15562-routes
  192.147.168.0/24
end-set

if (community 65535:666 && prefix-length == 32)
  advertise-only-to-blackhole-server
  do-not-install-in-fib
  exit-policy
if (destination in AS15562-routes)
  accept
  exit-policy
end-policy
```

Implementation plans

As part of the pmacct project, we will create a *blackhole validation server* dubbed “pmblackhole”

This *open-source* software will be able to:

- Ingest via API, BGP, or BMP
- Trigger conditionally on properties such as prefix length & BGP community
- Emit “actions” into a stateful data store such as Redis

Validated Blackholing will work on any platform

Collectors hang off the edges as IBGP Route Reflector Clients

Receive requests via BMP, or BGP:

https://tools.ietf.org/html/draft-ietf-idr-best-external-05

Injection via BGP with ExaBGP, GoBGP, OpenBGPD, or BIRD.

Or inject with statically configured ACLs, or use Flowspec

Many ways to skin the avocado!

RPKI and Blackholing Reconsidered

Honoring blackhole requests (*when they are on the active path)* goes very well together with RPKI Origin Validation: we don't need to do OV on the blackhole requests themselves, we only need to do OV on the covering next less specific route.

Community effort

- NTT Communications
- Telia Carrier
- pmacct

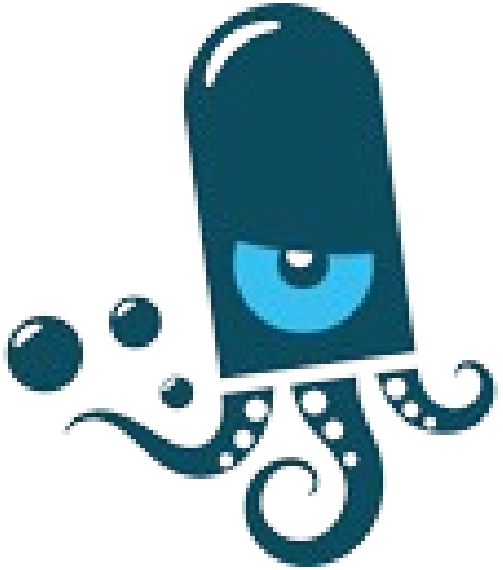

Summary

Only honor requests to blackhole traffic, if you'd send the traffic in that direction anyway.

Blackhole requests must be on the active path!