Security by BGP 101

Building distributed, BGP-based security system

Łukasz Bromirski
lukasz@bromirski.net

May 2017, CERT EE meeting

Roadmap for the session

BGP as security mechanism

BGP blackholing project

Call to action

Q&As and discussion

Flexibility of BGP

Use cases over the last years

Use case	Protocol in use in ~2000	Today and tomorrow
Internet (Peering)	BGP IPv4	BGP IPv4/v6
Private IP services (L3VPN)	BGP IPv4	BGP IPv4/v6 + HA + scalability
Private multicast (Mc VPN)	PIM	BGP Multicast VPN
L2 Services (L2VPN)	LDP VPWS/VPLS	BGP VPLS/VPWS, EVPN
DDoS attacks	CLI, ACL, PBR	BGP Blackholing/FlowSpec/QPPB
Network monitoring	SNMP	BGP monitoring protocol, BGP UPDATEs
Security	Filters	RPKI, BGP FlowSpec
Proximity and application routing		BGP Link State
Scaling DC	IGP (ISIS, OSPF) or L2 (Trill, FP, Vlan)	BGP, BGP SR
MPLS transport	LDP	BGP + Label Unicast (Unified MPLS)
SDN	PBR, OpenFlow (2013), Yang (future)	BGP FlowSpec, BGP Link State, BMP, BGP route controller, BGP Label Unicast, BGP Segment Routing
Overlay routing		VxLAN with BGP, Softwire

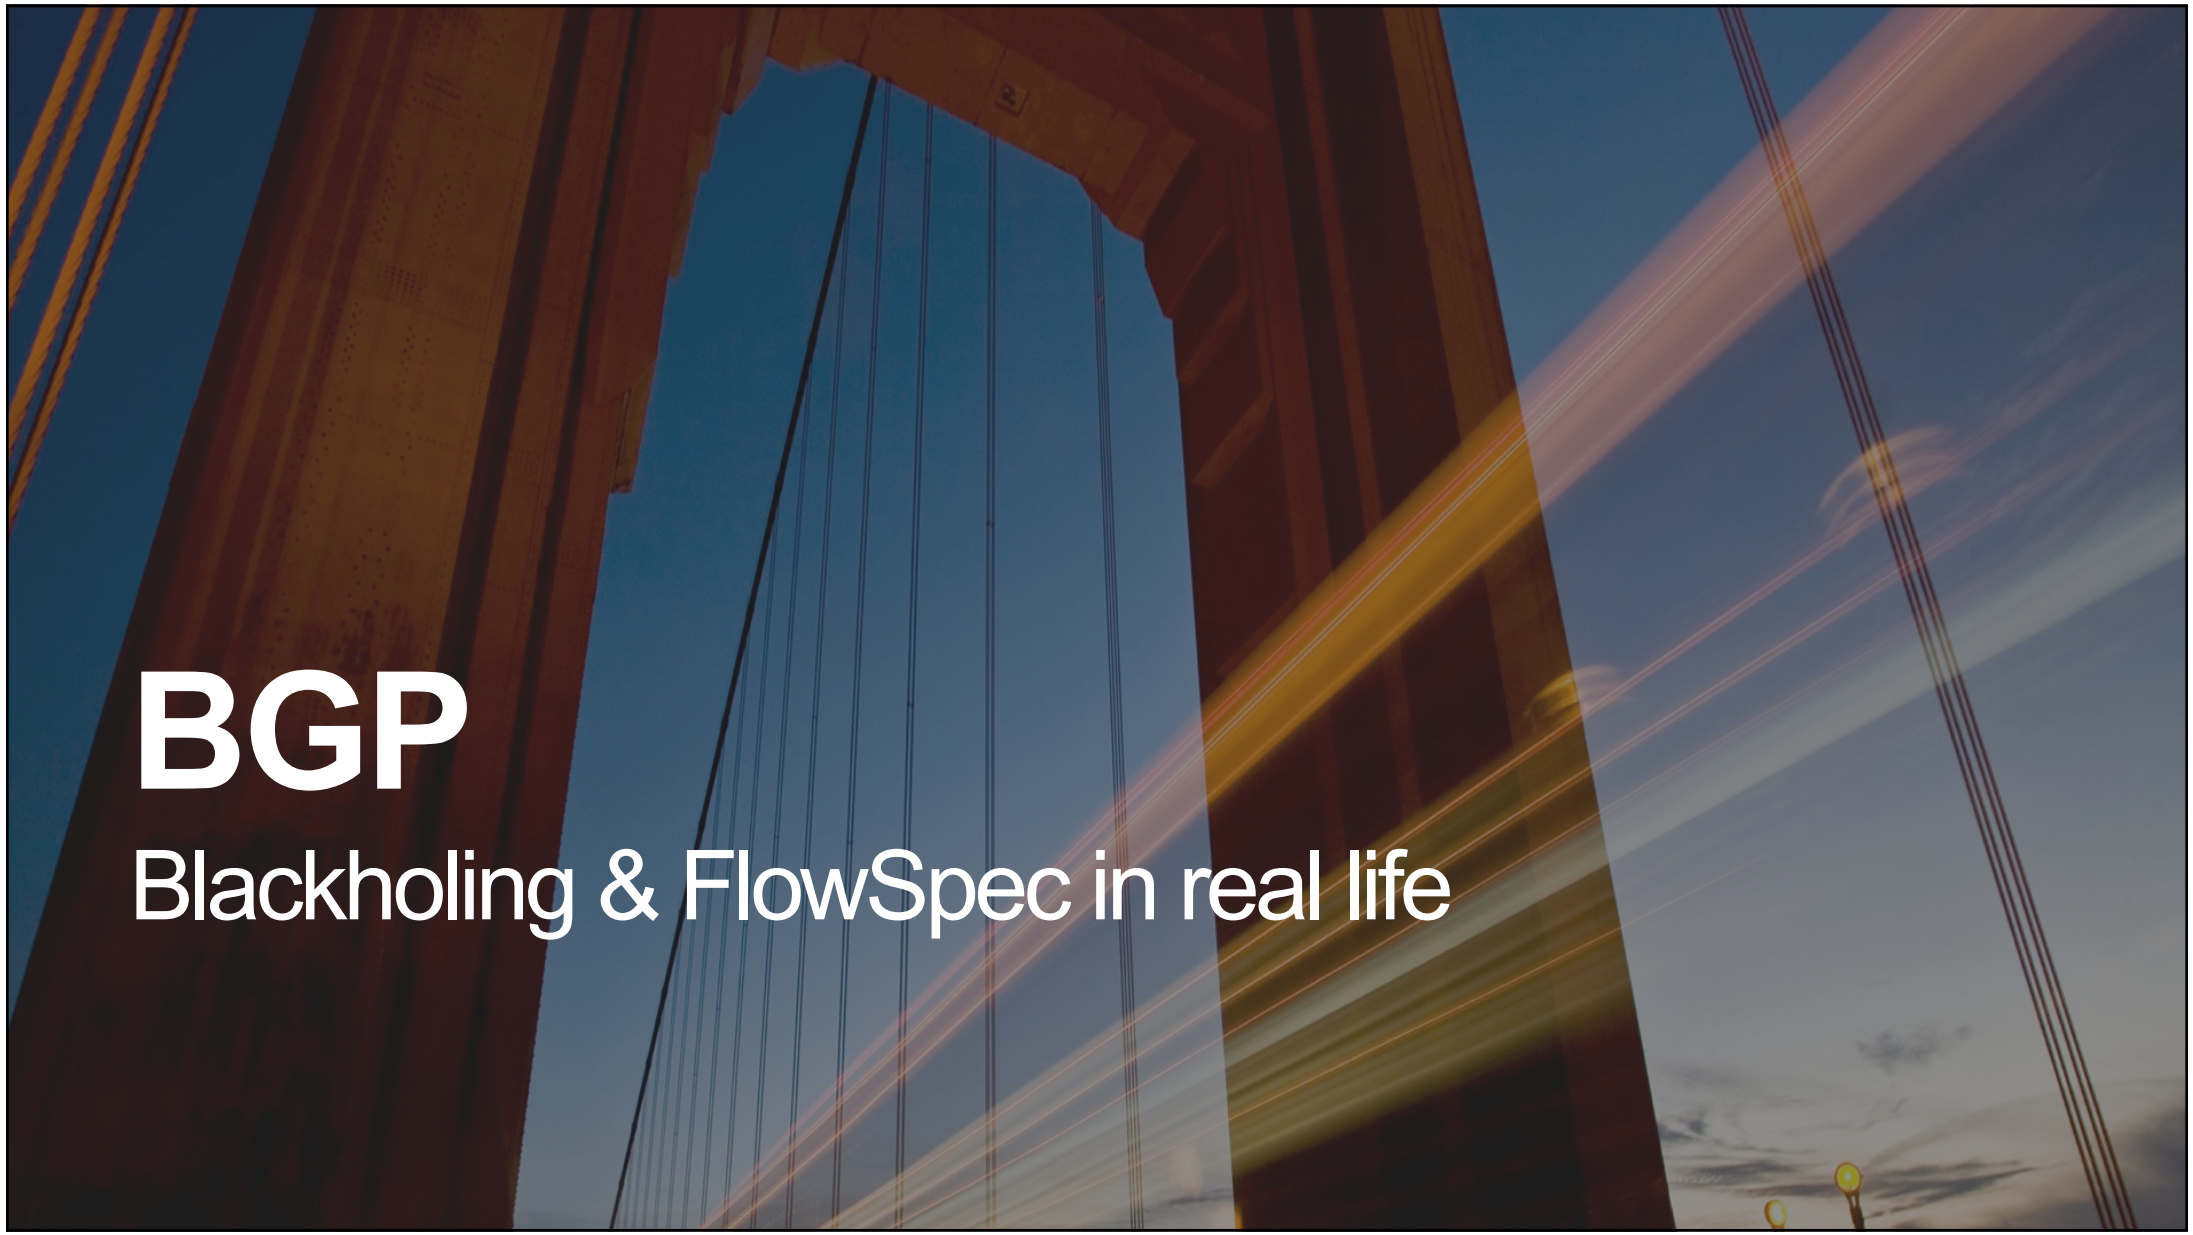
BGP
Blackholing & FlowSpec in real life

BGP Blackholing

I'm a ASN 451!

As soon as You announce Your address space, you can receive traffic

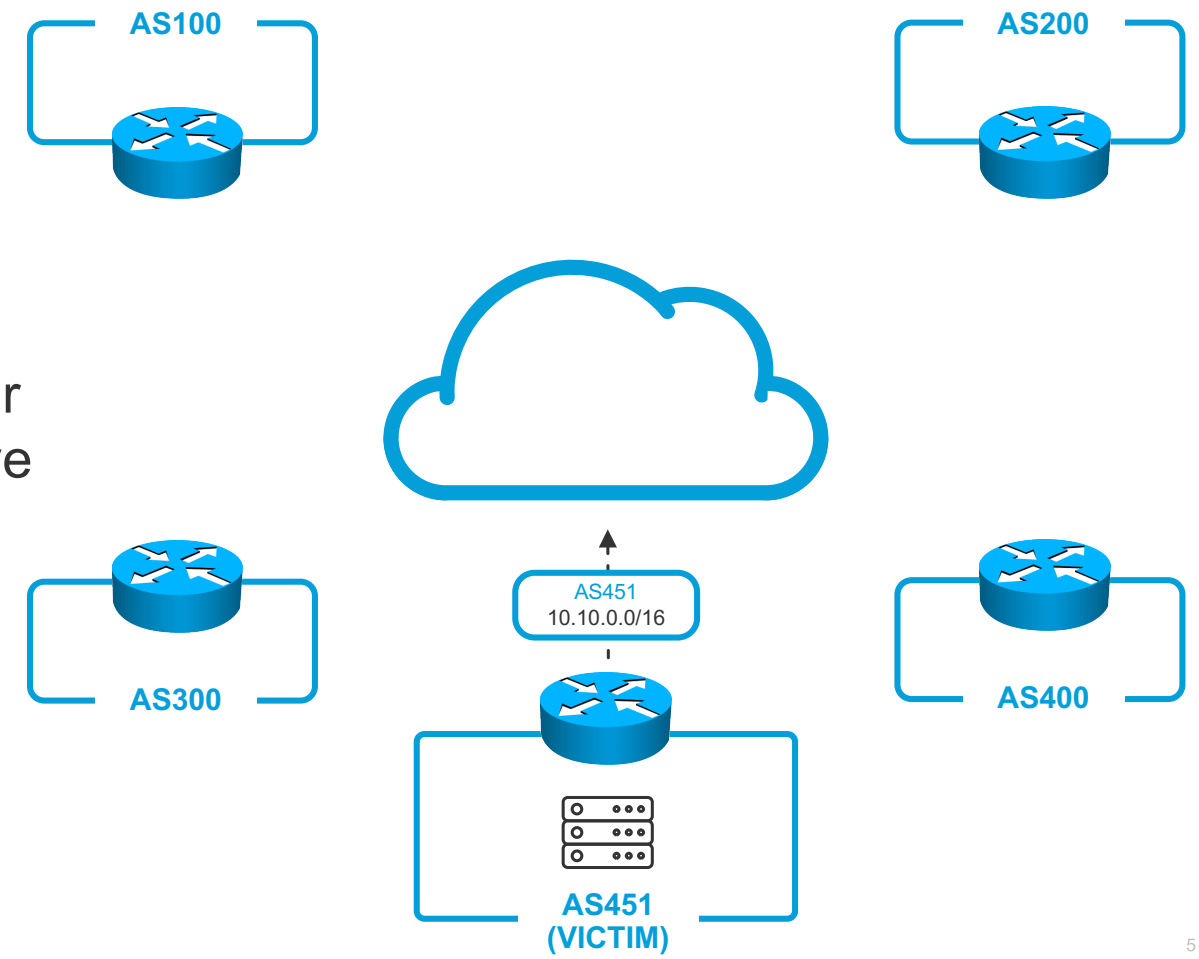

BGP Blackholing

Exposed in Internet

Volumetric and other types of attack simply deliver traffic to you – but unwanted traffic

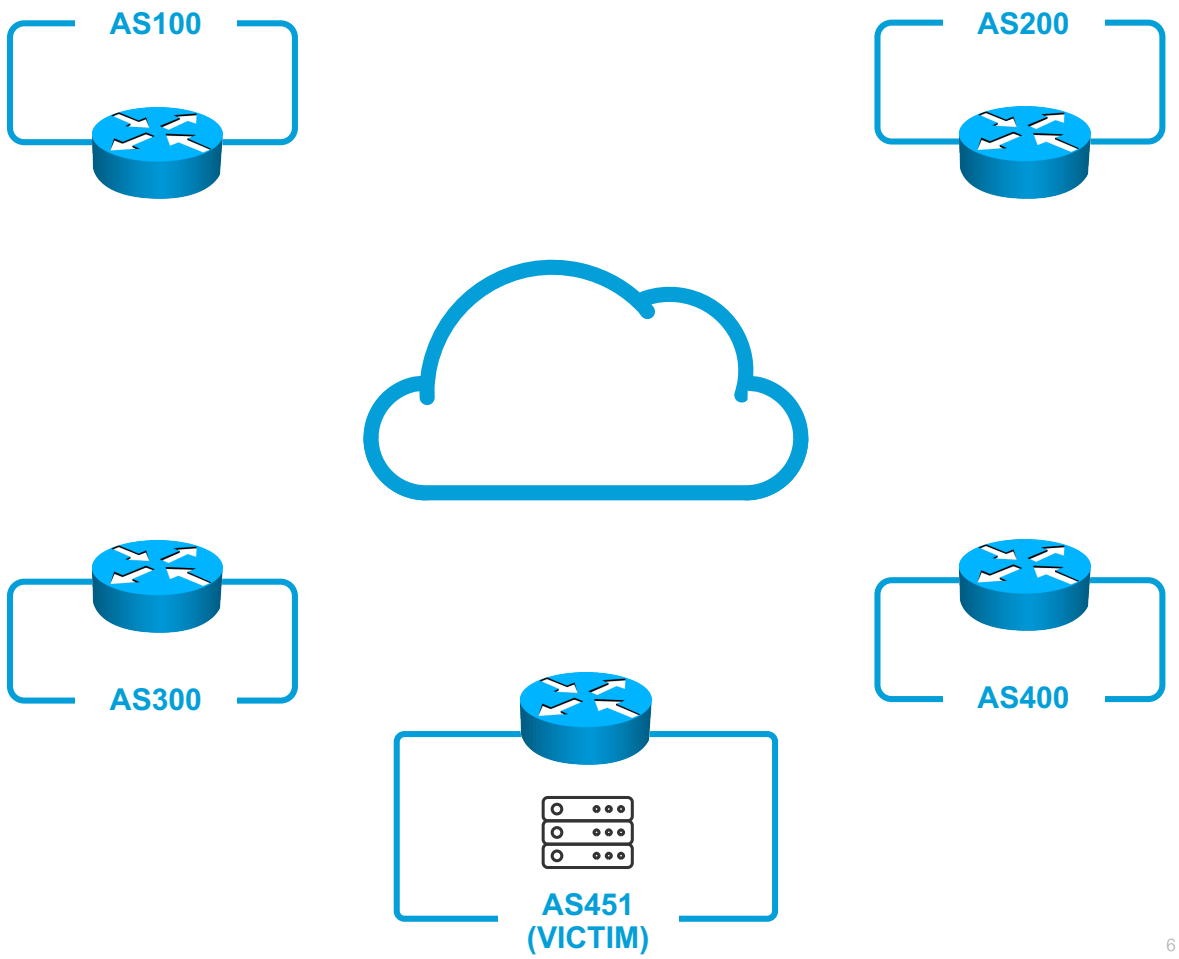

BGP Blackholing

DDoSing the Victim

DDoSes are happening daily – for some of us

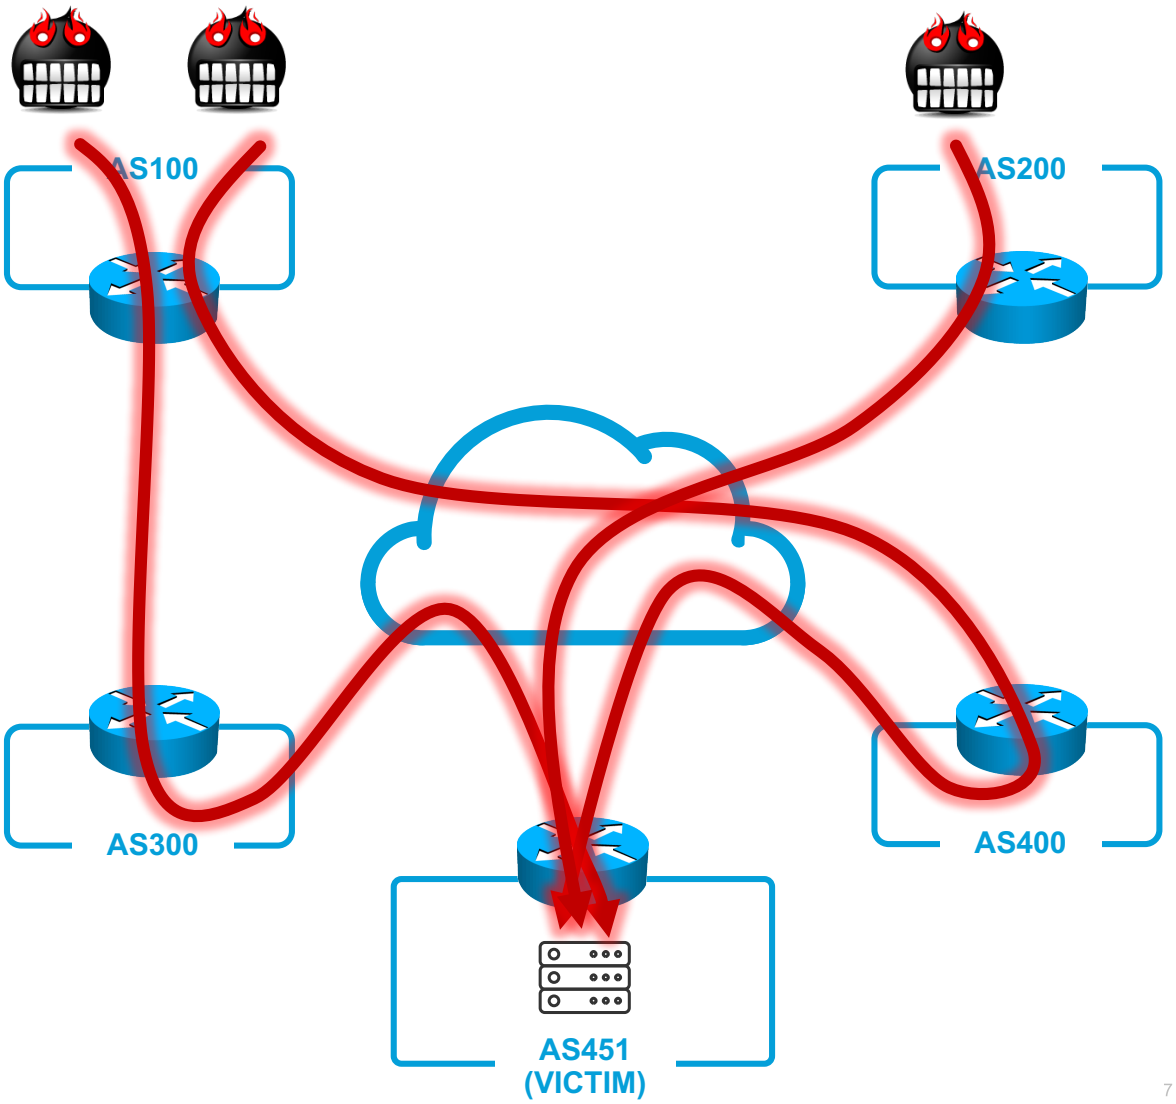

BGP Blackholing

DDoS “solutions”? Plenty

CLOUDFLARE

DDOS radware

Akamai

Incapsula

f5

ARBOR NETWORKS

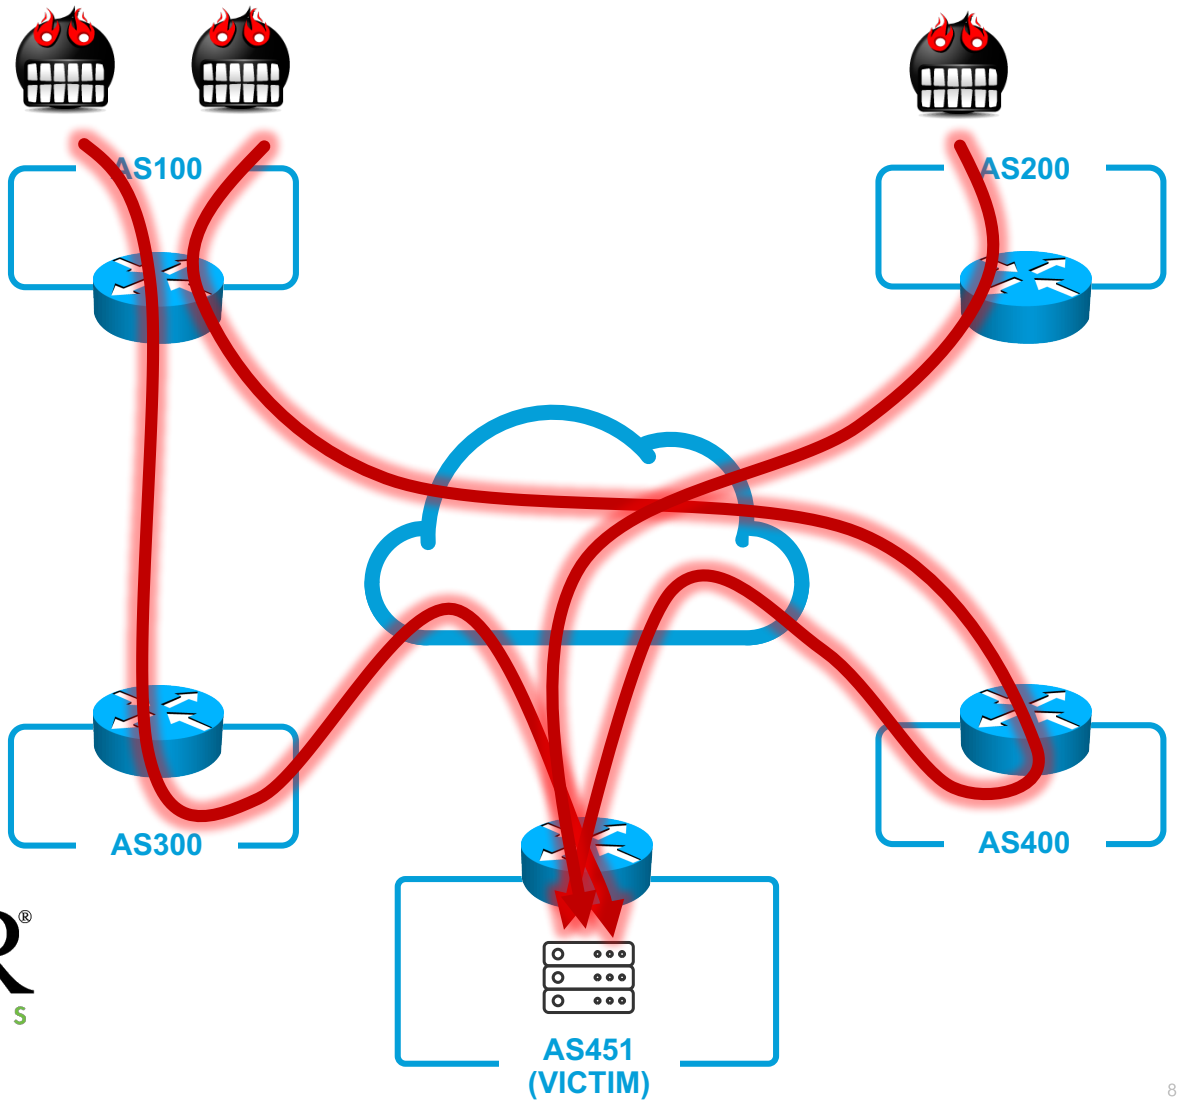

What if…

…you can’t afford the protection?

…you’d like to develop your skills and skills of your team and not pay someone to defend You?

…you anway want some control over what’s going on?

BGP Blackholing

Very effective filtering distribution – within and between ASes

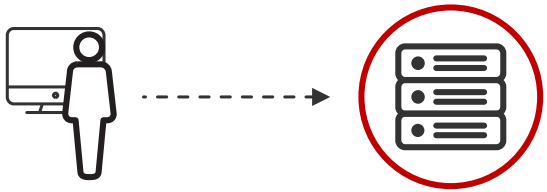

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

Trigger router

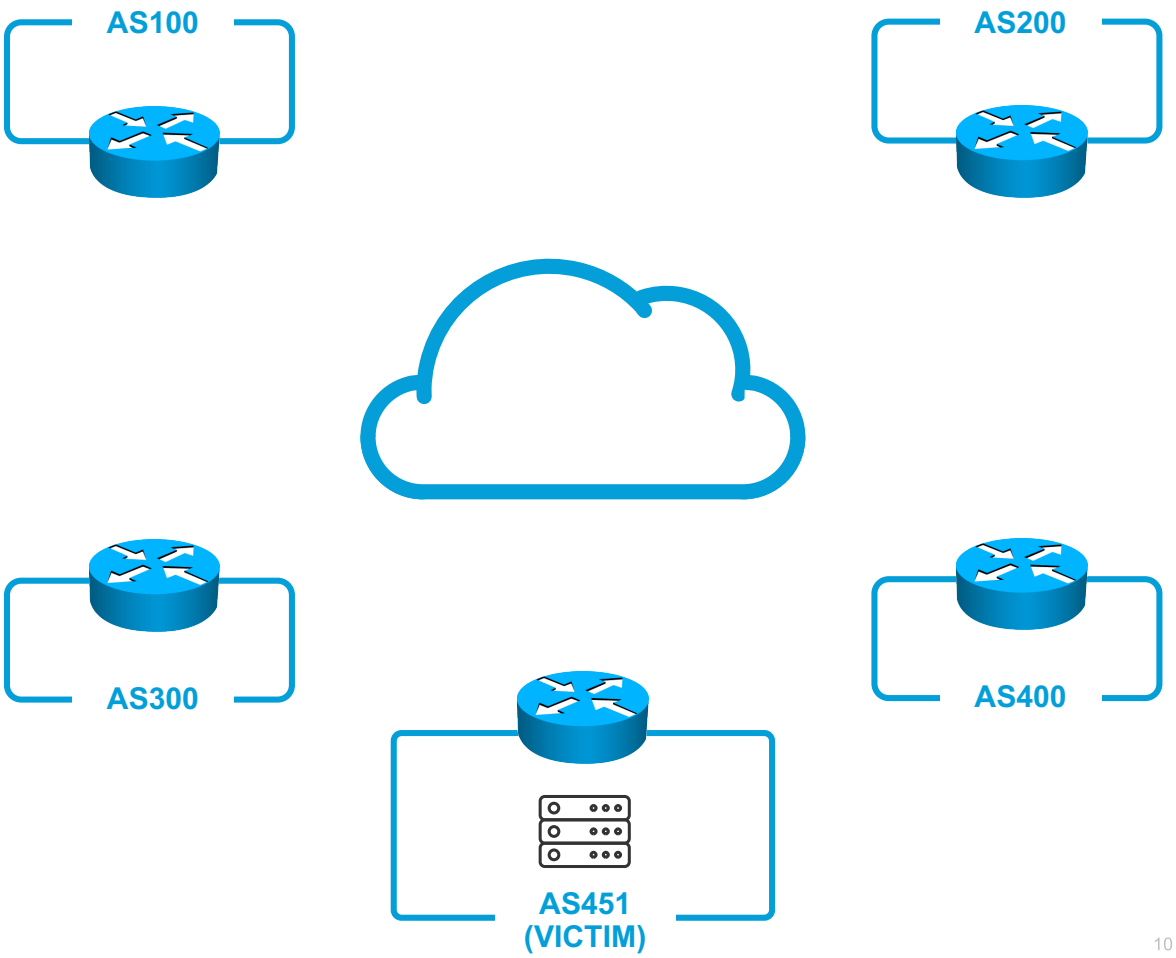

BGP Blackholing

We're announcing victim IP space to ASes

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

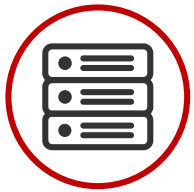

Trigger router

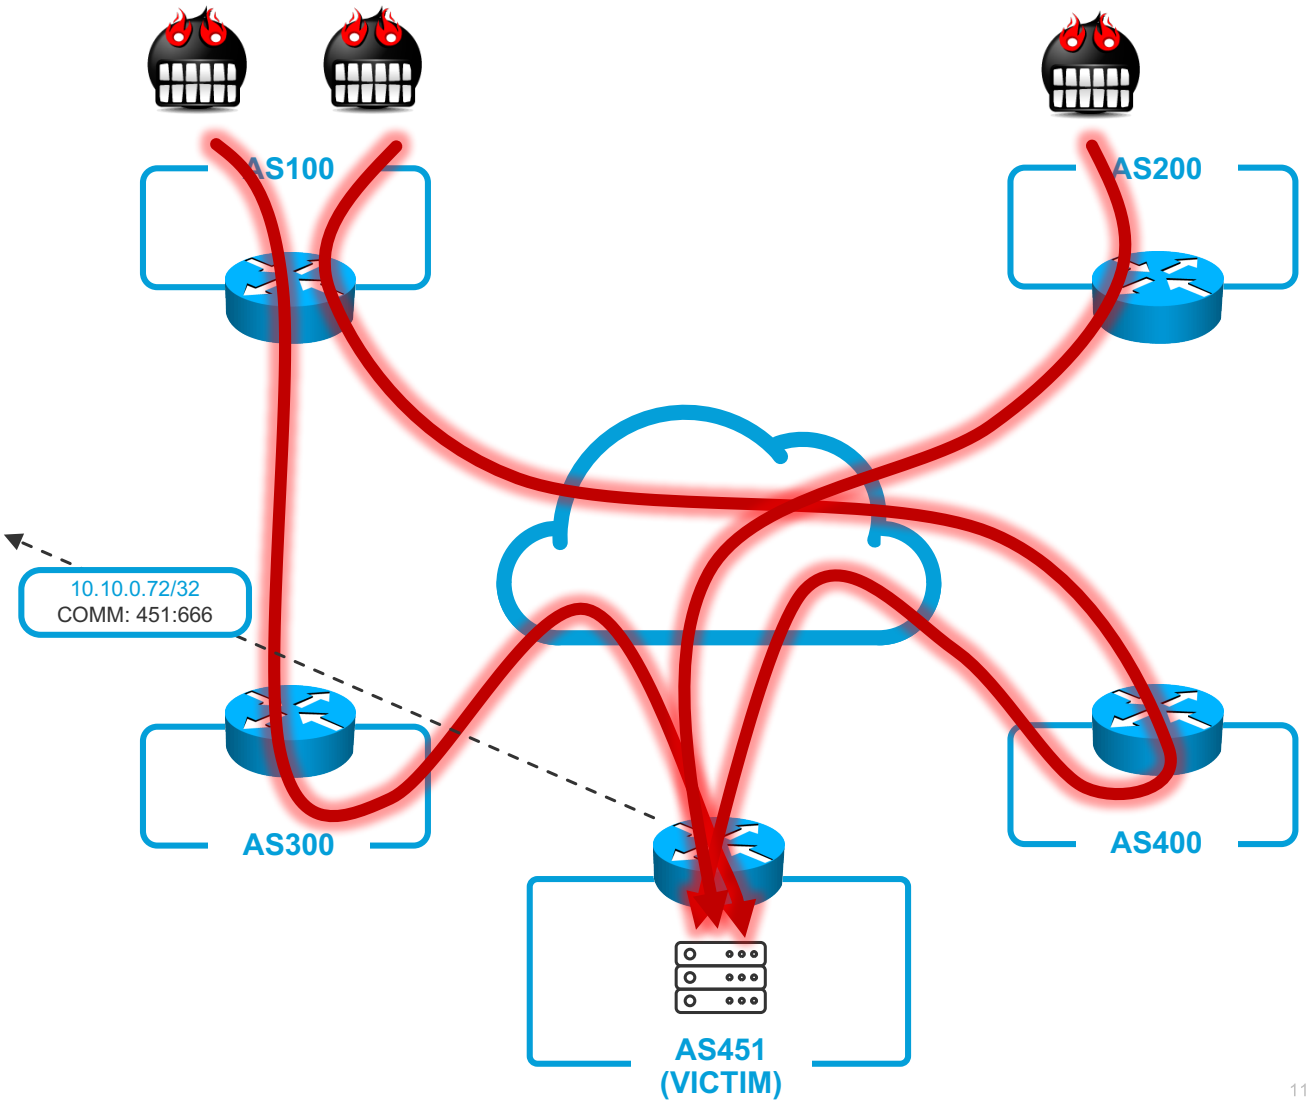

BGP Blackholing

We're announcing victim IP space to members

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

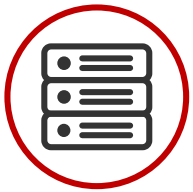

Trigger router

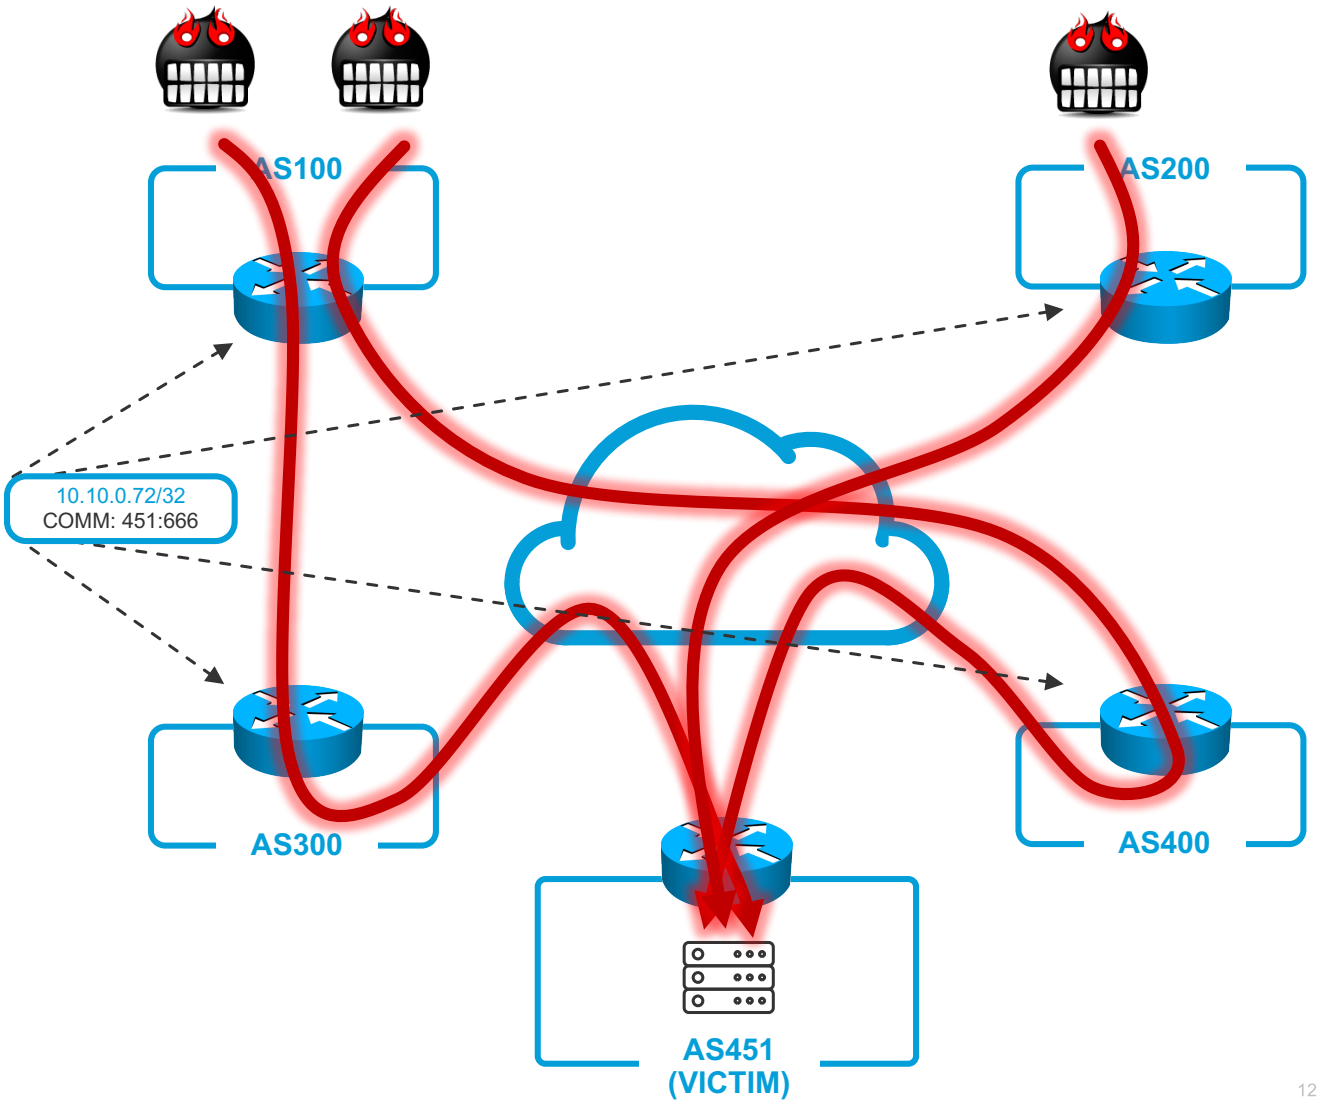

BGP Blackholing

We're announcing victim IP space to ASes

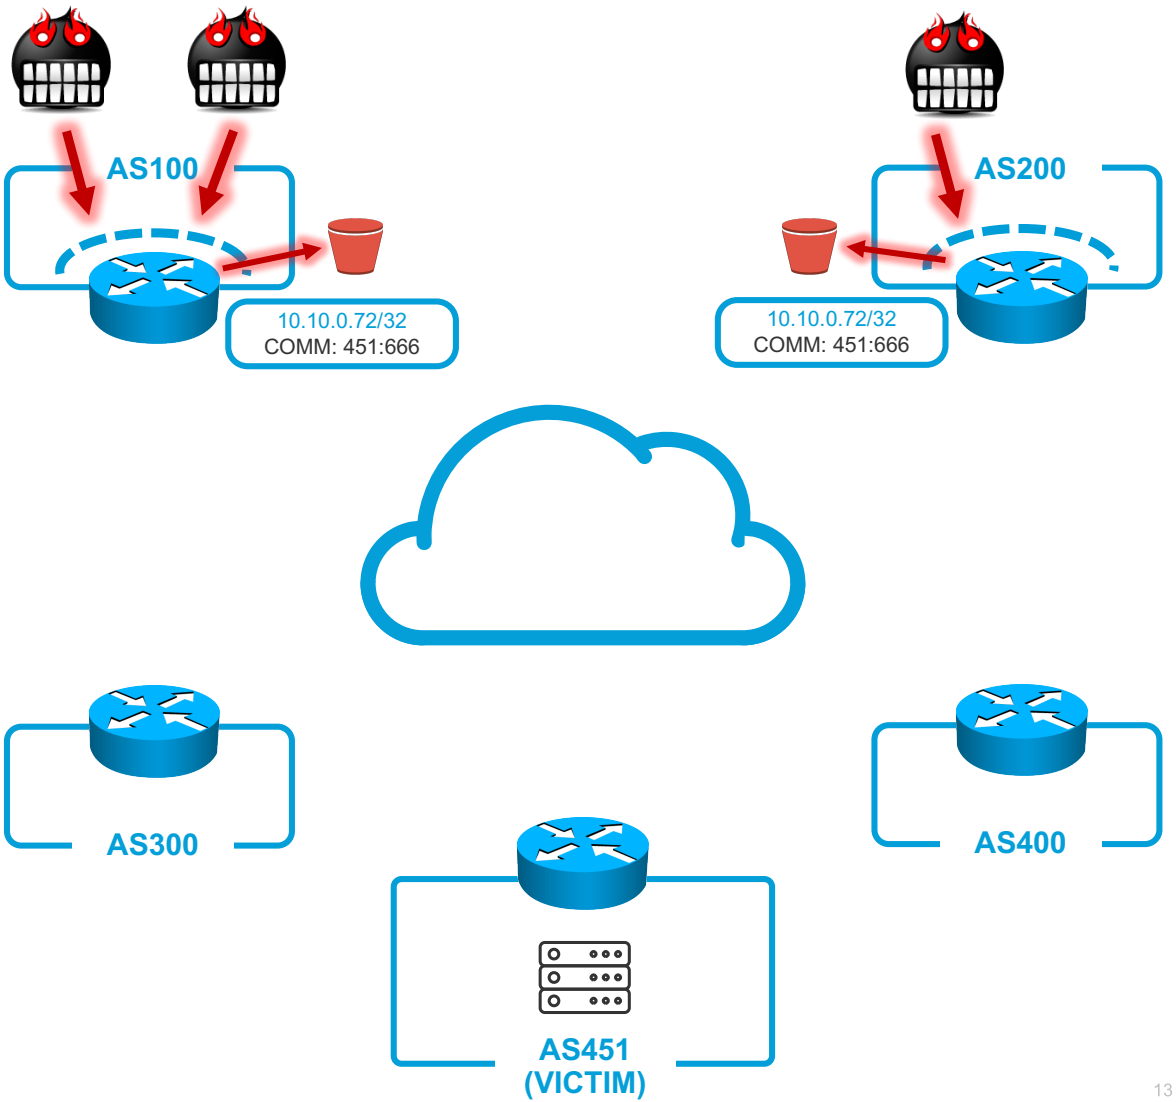

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

Trigger router

BGP Blackholing

"Please dear AS100, block your DDoS coming my way!"

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

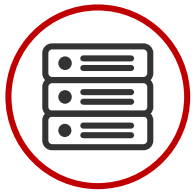

Trigger router

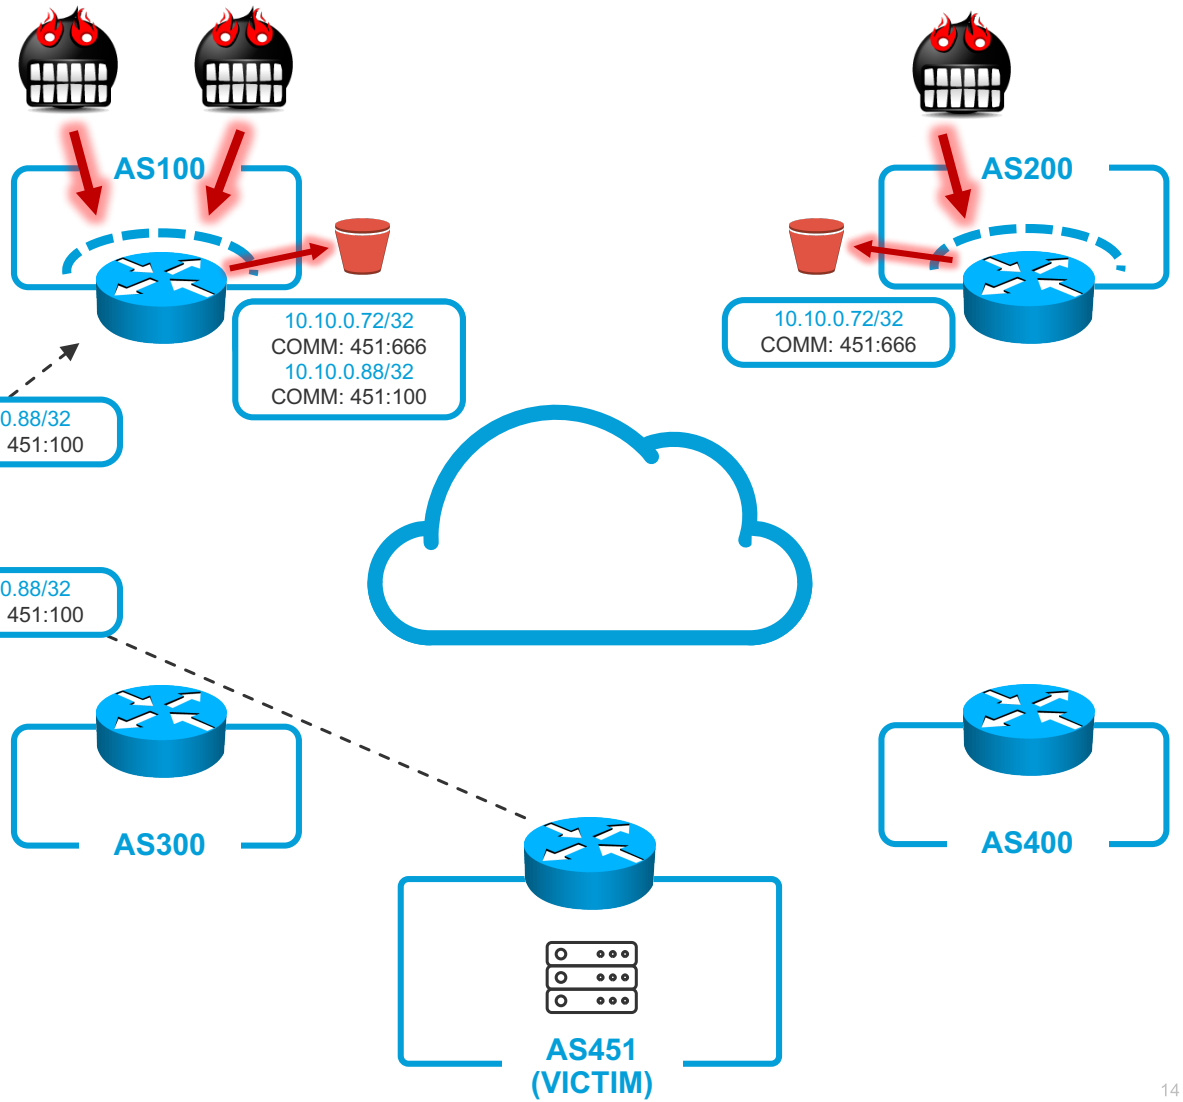

BGP Blackholing

"Please ALL, rate-limit the traffic to 256kbps" (FlowSpec)

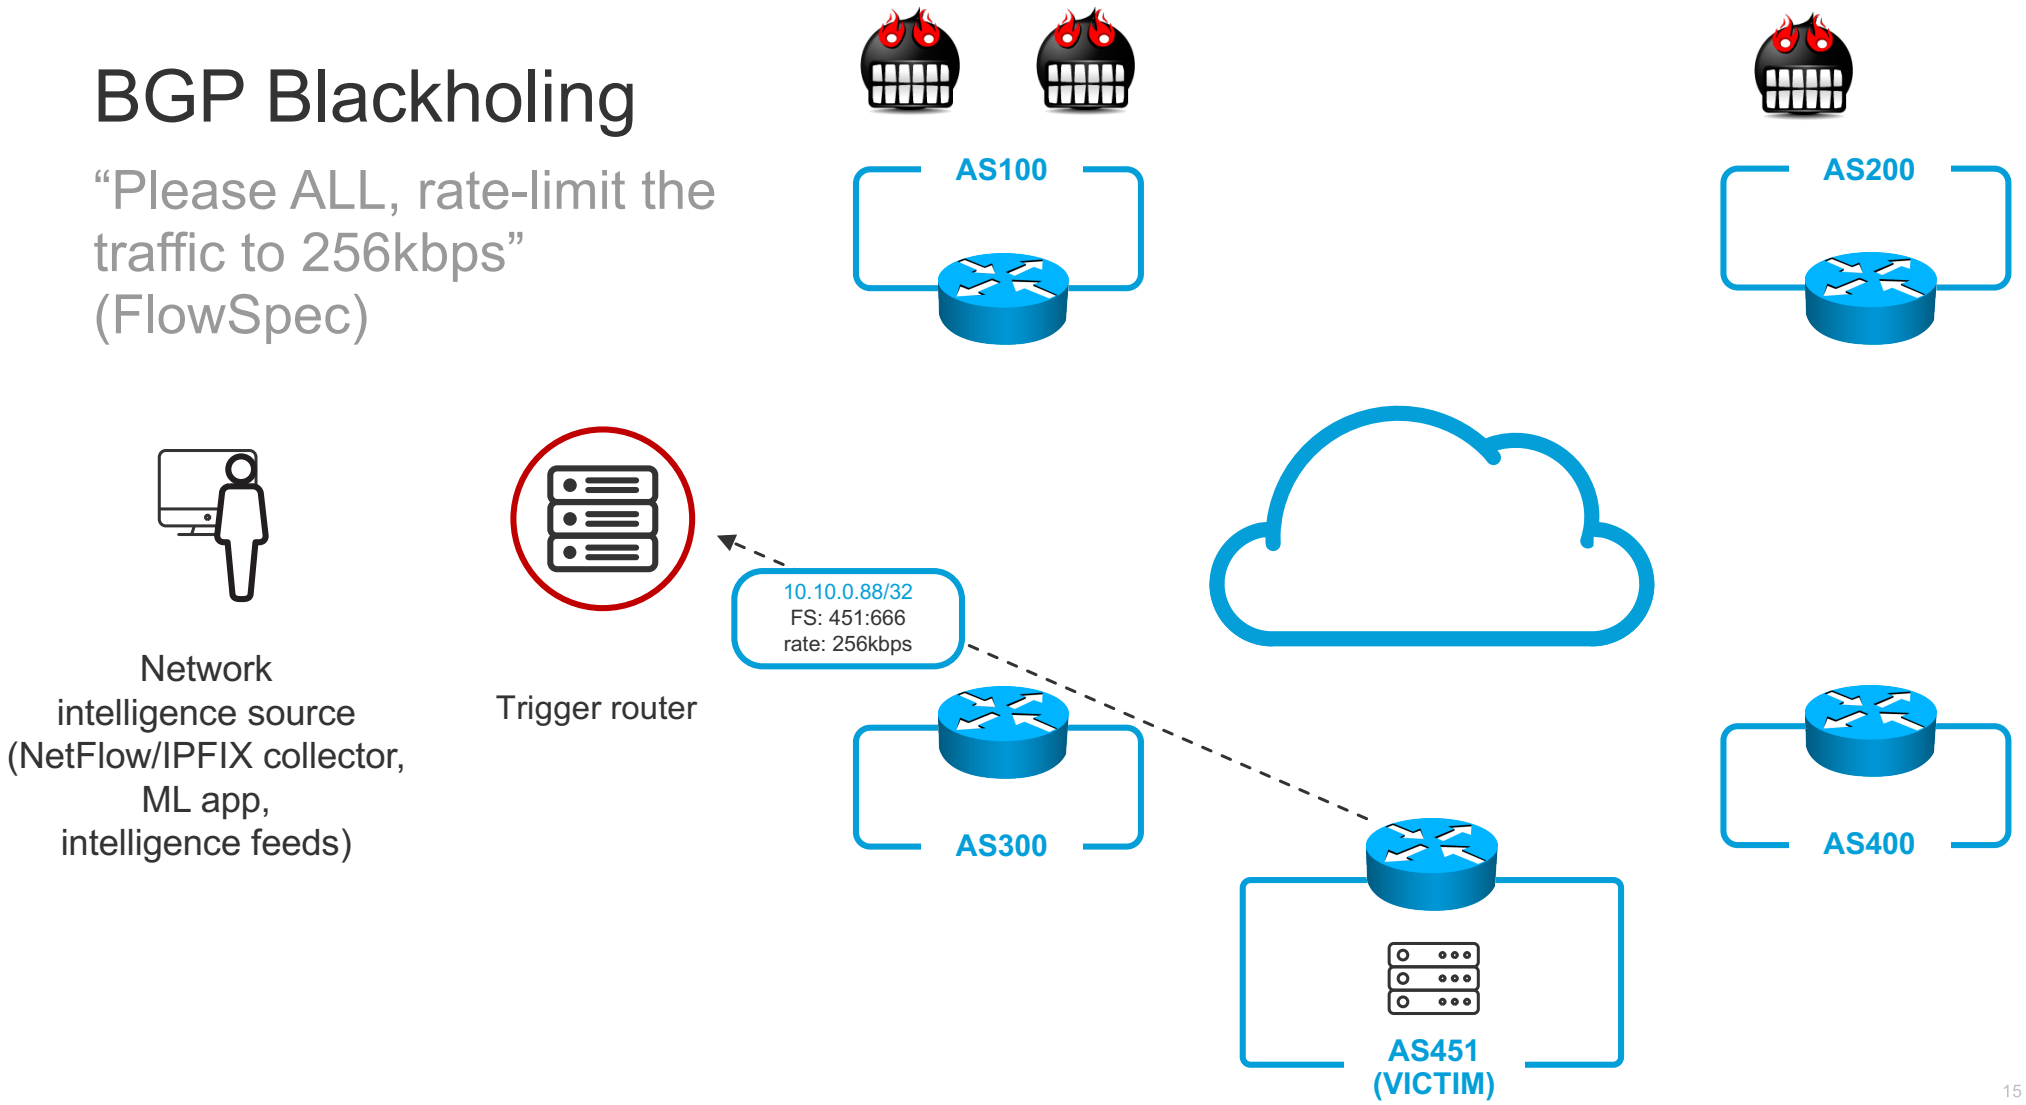

BGP Blackholing

"Please ALL, rate-limit the traffic to 256kbps" (FlowSpec)

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

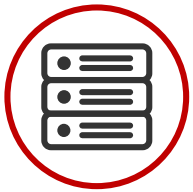

Trigger router

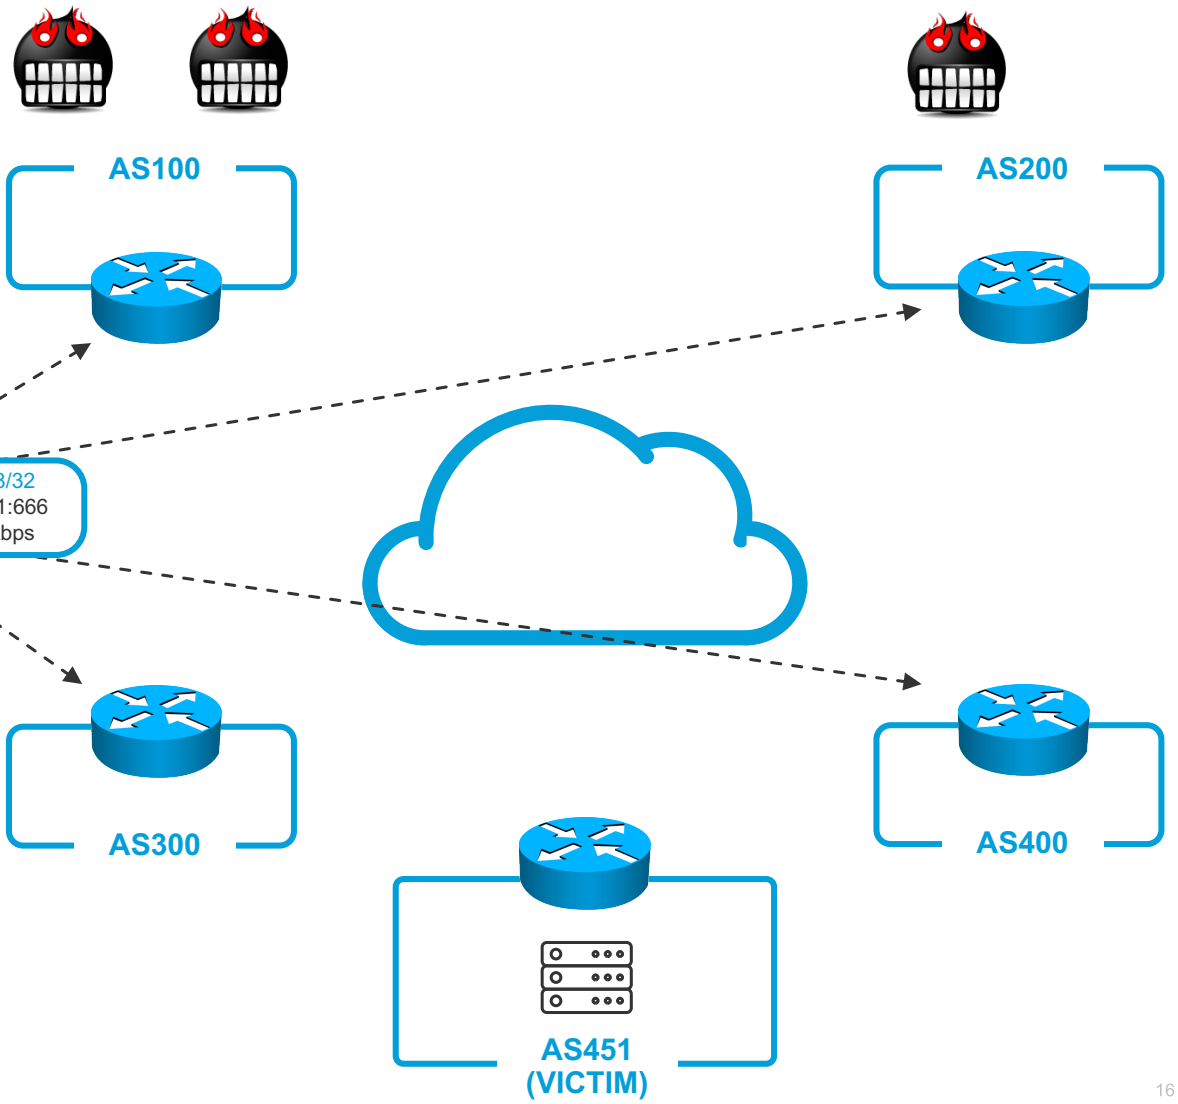

BGP Blackholing

"Please ALL, rate-limit the traffic to 256kbps" (FlowSpec)

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

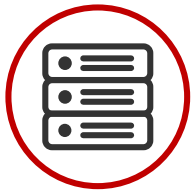

Trigger router

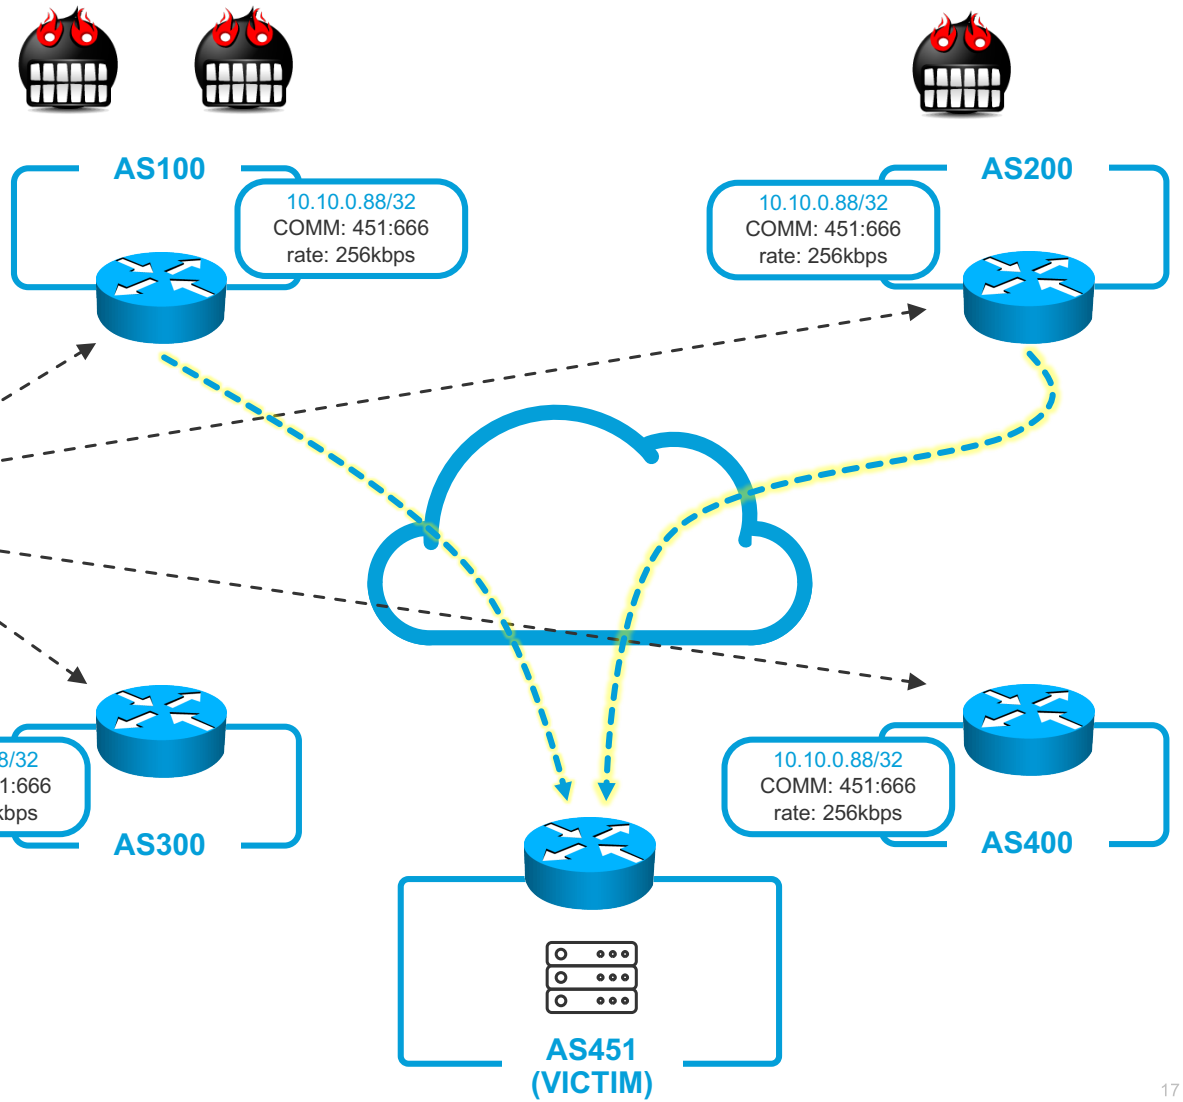

BGP Blackholing

"Please ALL, rate-limit the traffic to 80/tcp to 1Mbps" (FlowSpec)

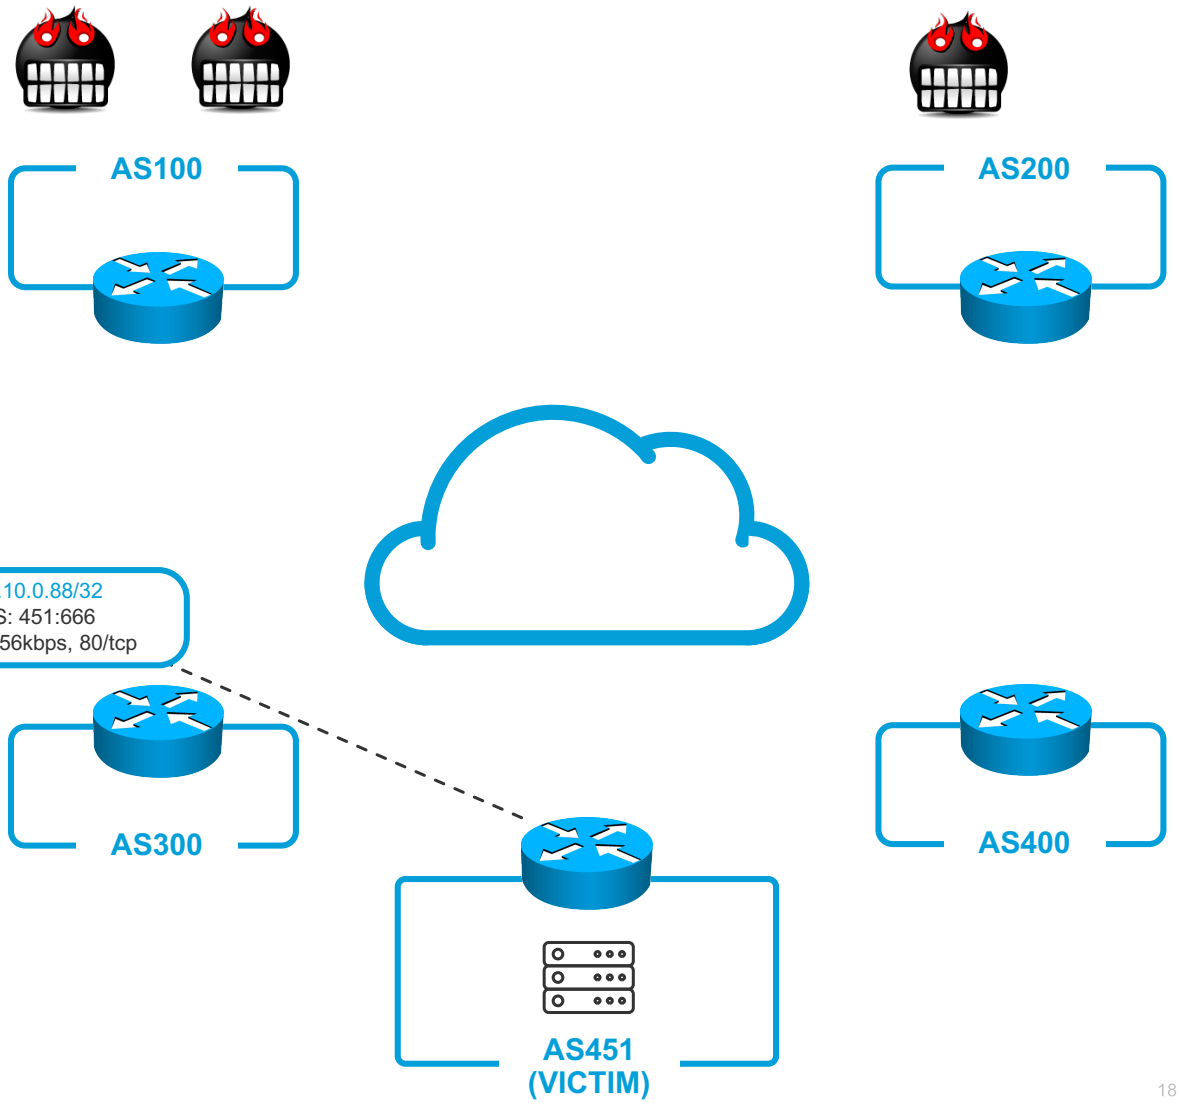

Network intelligence source (NetFlow/IPFIX collector, ML app, intelligence feeds)

Trigger router

Protecting the network using BGP

Options available right now – at your fingers reach

Traditional BGP blackholing	Traditional BGP sinkholing	BGP FlowSpec
Drops ALL traffic to prefix	Redirects ALL traffic to prefix	Drops, redirects, rate-limits traffic to specific L3/L4 combination
Operational best practice – announce and accept only /32s and /128s	Needs underlying infra for redirection: tunnels (GRE/IP/MPLS) and sniffing (receiving) device	Requires devices to understand FlowSpec address family, significantly less scaleable (‘000s vs million entries)

Protecting the network using BGP - configuration choices

Why we're using communities?

BGP announcing next-hop (->192.0.2.1->null0)

BGP propagates next-hop without changes in iBGP

Good for protection/triggering within ASN

BGP announcing community (if XXX:666->next-hop 192.0.2.1->null0)

BGP overwrites next-hop field at AS border (over eBGP sessions)

Scales between ASes and within ASes

“It’s C&C for you to block me from accessing my sites!”

BGP has very flexible policy language, you control EVERYTHING on your end!

Things we advise to block anyway on our peerings:

- Root DNS IPs (list is provided by the project)
- Your choosen public DNS services (like Google 8.8.8.8 or OpenDNS 208.67.222.222, etc)
- Your own AS space (there can be exceptions)
- Important NTP servers (country, European)
- Some other specific networks (will vary)

BGP Blackholing

Configuration for BGP Blackholing: IOS / IOS-XE

Trigger router

```
ip cef
ipv6 cef distributed
!
interface Null0
 no ip unreachables
 no ipv6 unreachables
!
route-map RED-RTBH permit 10
 match tag 666
 set origin igp
 set local-preferences 6666
 set ip next-hop 192.0.2.1
 set community 64999:666
!
ip route 192.0.2.1 255.255.255.255 null0 tag 666
!
router bgp 100
 address-family ipv4 unicast
  redistribute static route-map RED-RTBH
!
```

Edge node

```
ip cef
ipv6 cef distributed
!
interface Null0
 no ip unreachables
 no ipv6 unreachables
!
route-map GET-RTBH permit 10
 match community 64999:666
 set origin igp
 set local-preferences 6666
 set ip next-hop 192.0.2.1
!
ip route 192.0.2.1 255.255.255.255 null0 tag 666
!
router bgp 100
 address-family ipv4 unicast
  neighbor X.X.X.X route-map GET-RTBH in
!
```

BGP Blackholing

Configuration for BGP Blackholing with FlowSpec: IOS-XR

Trigger router

```
class-map type traffic match-all BAD-FLOW06
 match destination-address ipv4 172.16.6.6/32
 match destination-port range 135 139
!
policy-map type pbr FS-BH-GLOBAL
 class type traffic BAD-FLOW06
  drop
!
flowspe
 addres
  servi
!
router
 addres
 neighb
  addre
```

Edge node

```
flowspec
 address-family ipv4
  local-install interface-all
   ! you may choose to select only edge interfaces
 address-family ipv6
  local-install interface-all
```

anyway

```
EDGE.R7#show bgp ipv4 flowspec detail
BGP routing table entry for Source:172.16.6.6/32,DPort:=135:139, version 2
  Paths: (1 available, best #1, table IPv4-Flowspec-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  3356
    0.0.0.0 from 192.168.254.254 (192.168.254.254)
      Origin IGP, localpref 100, valid, external, best
      Extended Community: FLOWSPEC Traffic-rate:64999,0
      rx pathid: 0, tx pathid: 0x0
```


BGP Blackholing
Call to action – let's build community!

BGP blackholing

Community based effort to make the internet more secure place

- We’re not solving **all** the DDoS problems with BGP blackholing, we’re not aiming at that
- We’re helping to educate people, deploy best practices and learn how to operate in increasingly more hostile environment
- This is not paid service, there are no SLAs, but there’s team of people committed to doing “the right things”

BGP Blackholing

International edition – call to action

- Equipment in IXPs across Europe
- 4U with power and Internet connectivity needed if You'd like to colocate us (please do!)
- Let's build and share best practices and intelligence
- Every project member adds value and protection

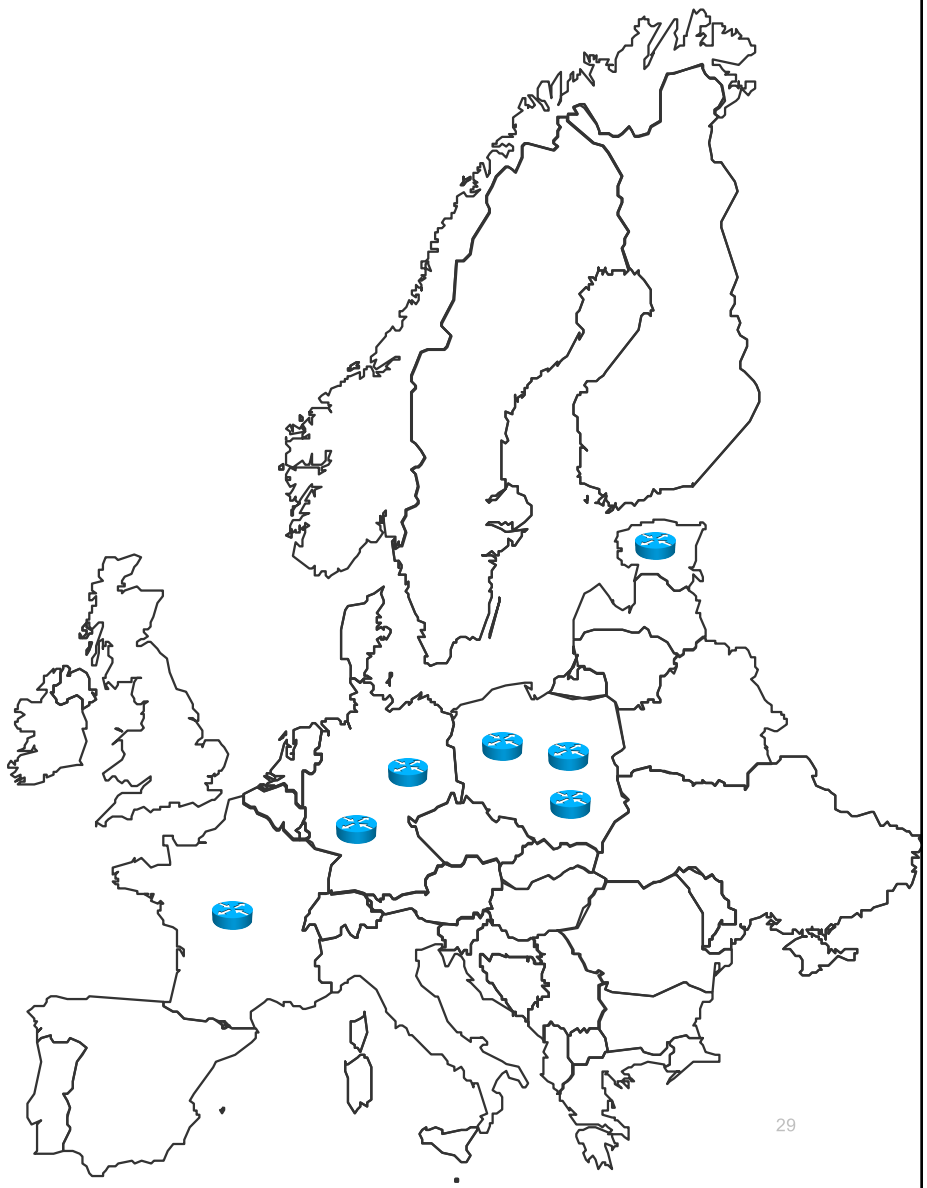

Call to Action

There’s a lot more to do with best practices

- MANRS initiative – antispoofing, RPKI adoption
  https://www.routingmanifesto.org/manrs/
- Engage with *nog sec teams – your knowledge and passion is needed to push NSPs forward!
- Join us – BGP Blackholing PL is going international
  https://null0.pl

Q&A

Łukasz Bromirski
lukasz@bromirski.net