LINK TECHNOLOGIES INC

www.linktechs.net
(314)-735-0270

Advanced BGP using Route Reflectors

By Dennis Burgess

About your speaker

- By Dennis Burgess
- CTO Link Technologies, Inc.
- Advanced WISP Support/Engineering / Consulting
  - BGP – OSPF – VPLS – MPLS - Routing
- www.linktechs.net
- 314-735-0270

Some Basics

eBGP and iBGP

- External BGP
  - Any BGP Peer that peers with a different AS than yours
  - Your router will have differences for routes learns though eBGP vs iBGP.

eBGP and iBGP

- Internal BGP -- iBGP
  - Any BGP peer that peers with the SAME AS number.
  - Typically used for cross connects as well as other BGP routers in your network.

Route Reflector

Route Reflector

- Common to have a Route Reflector inside a complicated BGP structure
- Allows for simplified communications and routes.
  - Typically you need two RRs to ensure route availability in the unlikely event that one goes down.

Route Reflector – Command/Control

- Typically small board or CHR
  - Can be as small as a 2011
  - This is a command and control router
  - This router is iBGP connected to all of your BGP routers that talk eBGP

Central Control Route Reflector

- Configure routers and filters accordingly
  - CogentCo – Uses Black Hole Routers to back hole traffic
  - ATT – Uses Community String to black hole
  - Your Provider – Ask them for their BGP guild, this should give you all of the information, such as community strings etc. that you could need.

Central Control Ro

- This is from CogentCo
  - There BGP Guide.
  - Gives information such as where you will form a eBGP peer to connect with a black hole router.

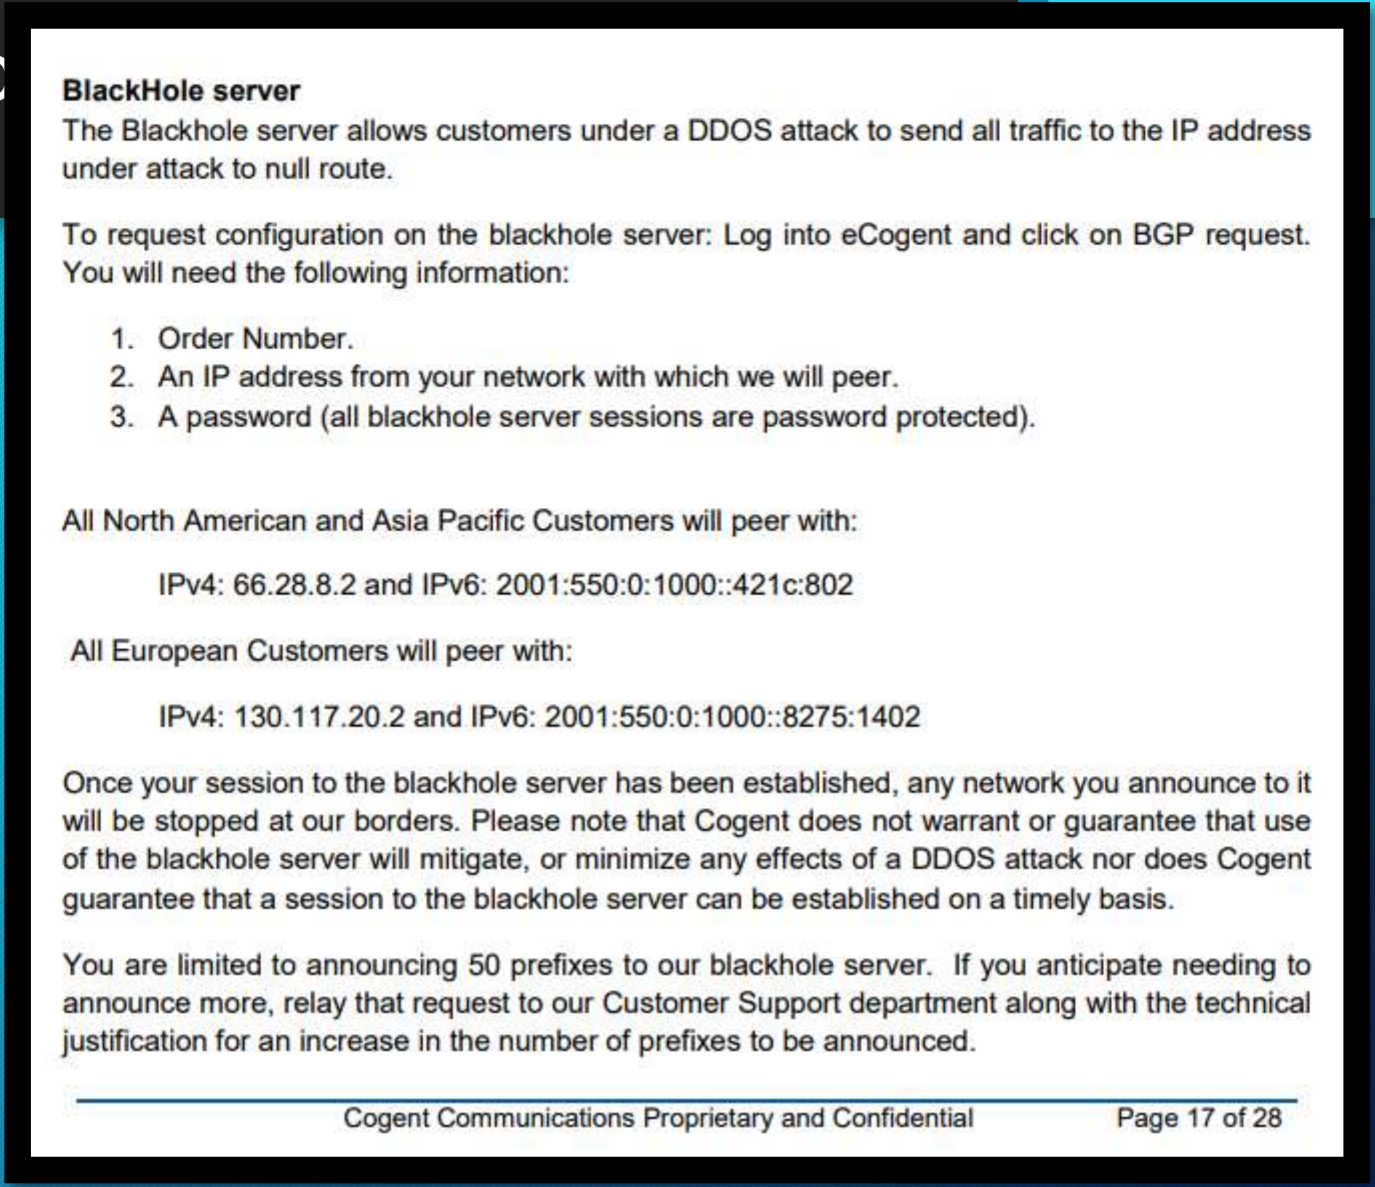

BlackHole server

The Blackhole server allows customers under a DDOS attack to send all traffic to the IP address under attack to null route.

To request configuration on the blackhole server: Log into eCogent and click on BGP request. You will need the following information:

1. Order Number.
2. An IP address from your network with which we will peer.
3. A password (all blackhole server sessions are password protected).

All North American and Asia Pacific Customers will peer with:

IPv4: 66.28.8.2 and IPv6: 2001:550:0:1000::421c:802

All European Customers will peer with:

IPv4: 130.117.20.2 and IPv6: 2001:550:0:1000::8275:1402

Once your session to the blackhole server has been established, any network you announce to it will be stopped at our borders. Please note that Cogent does not warrant or guarantee that use of the blackhole server will mitigate, or minimize any effects of a DDOS attack nor does Cogent guarantee that a session to the blackhole server can be established on a timely basis.

You are limited to announcing 50 prefixes to our blackhole server. If you anticipate needing to announce more, relay that request to our Customer Support department along with the technical justification for an increase in the number of prefixes to be announced.

Cogent Communications Proprietary and Confidential Page 17 of 28

Central Control Route Reflector

- I have some customers that use this control RR to advertise their prefixes out, and use it to control how that is done.
  - Communities can be used to control how your prefix is announced.
  - Typically its located at their eBGP peering point. Not all of the time

Central Control Route Reflector

- BOGON Lists and other BGP Feeds
  - In some cases, we have a bit more CPU behind this.
  - We will setup CYMRU BOGON Peering with this router, then distribute the BOGON List with edge routers.
    - BOGON are prefixes that have not be officially assigned to anyone by a IANA or other Regional Authority.
    - They are high jacked and are being used to send spam.
    - FREE SERVICE – You SHOULD have a BOGON list and BLACKHOLE any routes on that list.

Central Control Route Reflector

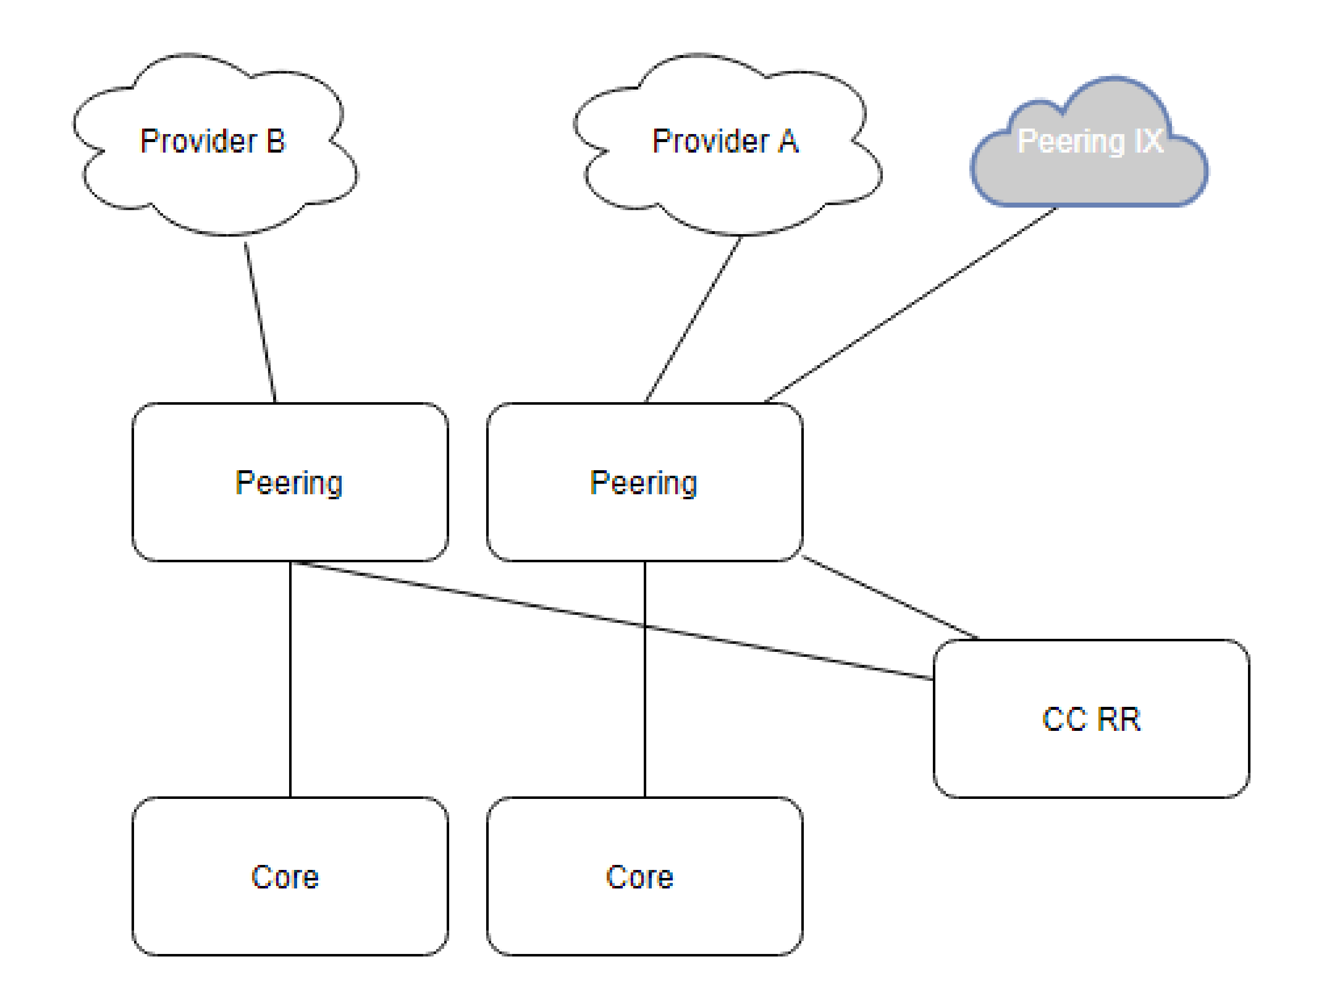

Central Control Route Reflector

- In this case, we can do our announcements from the CC RR router. This is our route reflector.
  - By adjusting the BGP communities that we have on this router, we can affect how and what is advertised out what provider.
  - Furthermore, we can have our BGP customers peer with our RRs, vs directly with our peering routers.
  - Then our set of rules, they can get a copy of (BGP Guide) and then they can affect how their prefixes are announced as well!

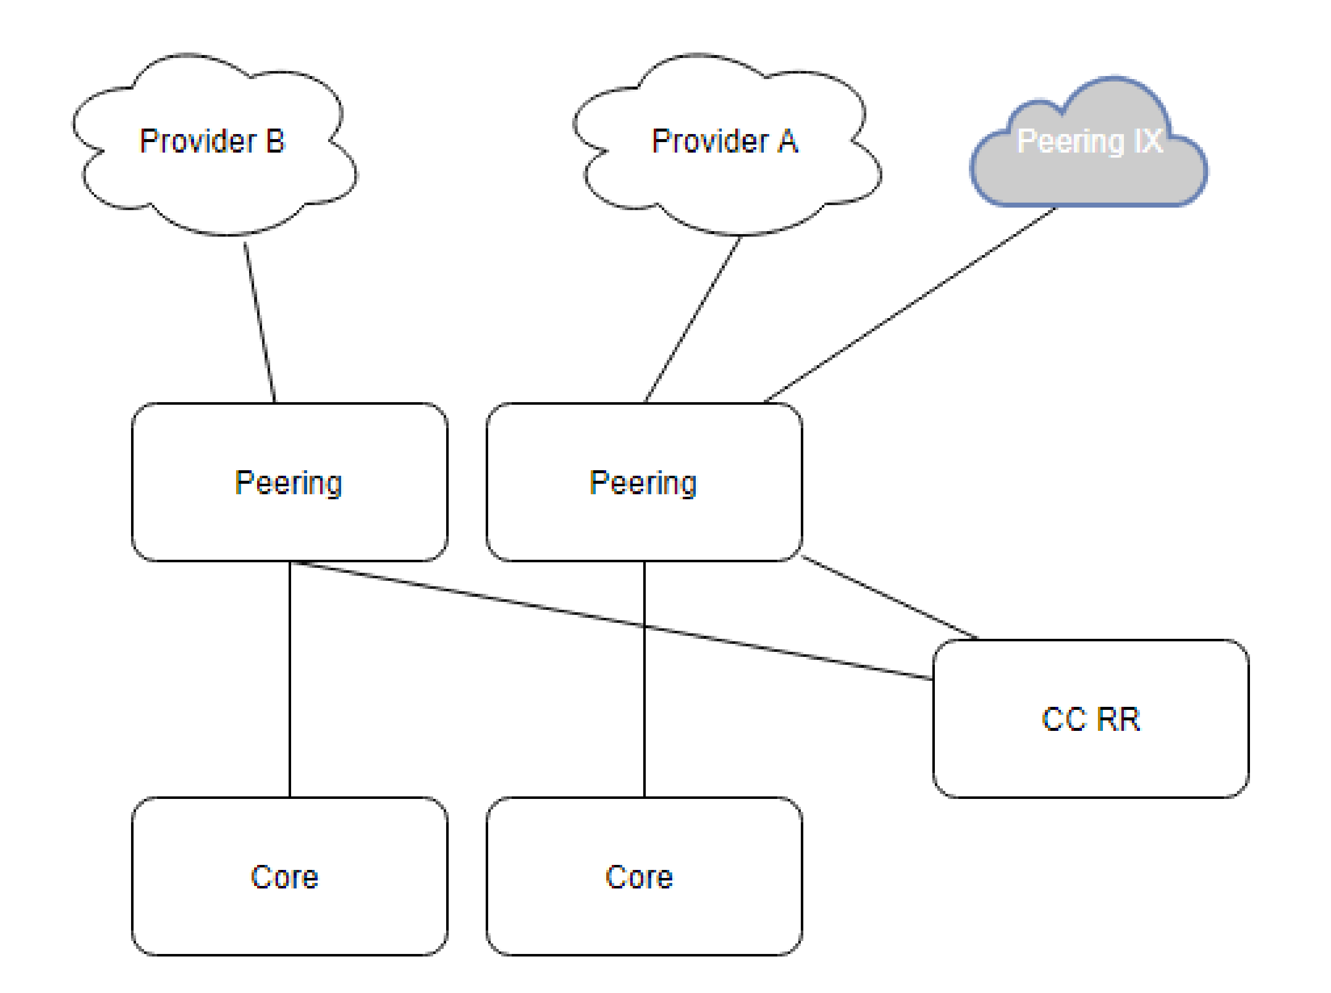

Survive DDOS via Centralized Route Reflector

Surviving a DDOS attack

- Two Stages and one sub-stage

Surviving a DDOS attack

Enough Bandwidth to sustain the attack

Not Enough bandwidth

Surviving a DDOS attack

Enough Bandwidth to sustain the attack

Not Enough bandwidth

BLACKHOLE

Surviving a DDOS attack

- How to survive a DDOS
  - **HAVE ENOUGH BANDWIDTH TO SUSTAIN THE ATTACK**

Surviving a DDOS attack

- How to survive a DDOS
  - **HAVE ENOUGH BANDWIDTH TO SUSTAIN THE ATTACK**

  - Place rules on edge routers to protect against attackers.
  - These edge routers they should be secure but should have minimal rule set
  - Turn off services not being used
  - Typically 10 gig interfaces

Surviving a DDOS attack

- How to survive a DDOS
  - HAVE ENOUGH BANDWIDTH TO SUSTAIN THE ATTACK
    - If you have enough bandwidth, then you can rely on your router rules to block.
    - IF you have 4 eBGP peers, all 10gig, then an attack from 20,000 attackers comes in.
      - Some attackers will come in on each peer, block them for 5 min. If they continue to attack it will block them again
      - 99% of the time, once an attacker starts the attack they will check to see if its effective. If not, they stop, or grow the attack.

Surviving a DDOS attack – Blackhole

Surviving a DDOS attack – Blackhole

- Backhole
    - You advertise a /32 or /128 IP address to your upstreams
    - You do this via a predefined method (BGP Guild)
    - Once done, your up streams block traffic to that IP
        - Note this takes your single IP off-line! But stops the packets from even making it to your interface
        - Thus you can sustain the attack

Surviving a DDOS attack

- Once you are out of bandwidth
  - You have a choice
    - **Have one customer go off-line and be mad at you or have everyone be mad at you?**
      - I would choose the single customer

Surviving a DDOS attack

- How to survive a DDOS
  - You place the single IP or IPs being attacked that you wish to black hole on this RR .
  - You don't even have to LOGIN to your edges!
  - This pushes that /32 to your edges, and since it's a /32 your edges know where and how to advertise it.
    - They will advertise with the correct community string to black hole
    - And/or submit it to black hole servers

Surviving a DDOS attack

- How to survive a DDOS
  - Now that /32 is blocked
  - It is blocked at your up streams before it gets to you
  - They, the /32 is OFFLINE but your other customers are happy

Stages of DDOS Attack

Surviving a DDOS attack

- Attack coming in..

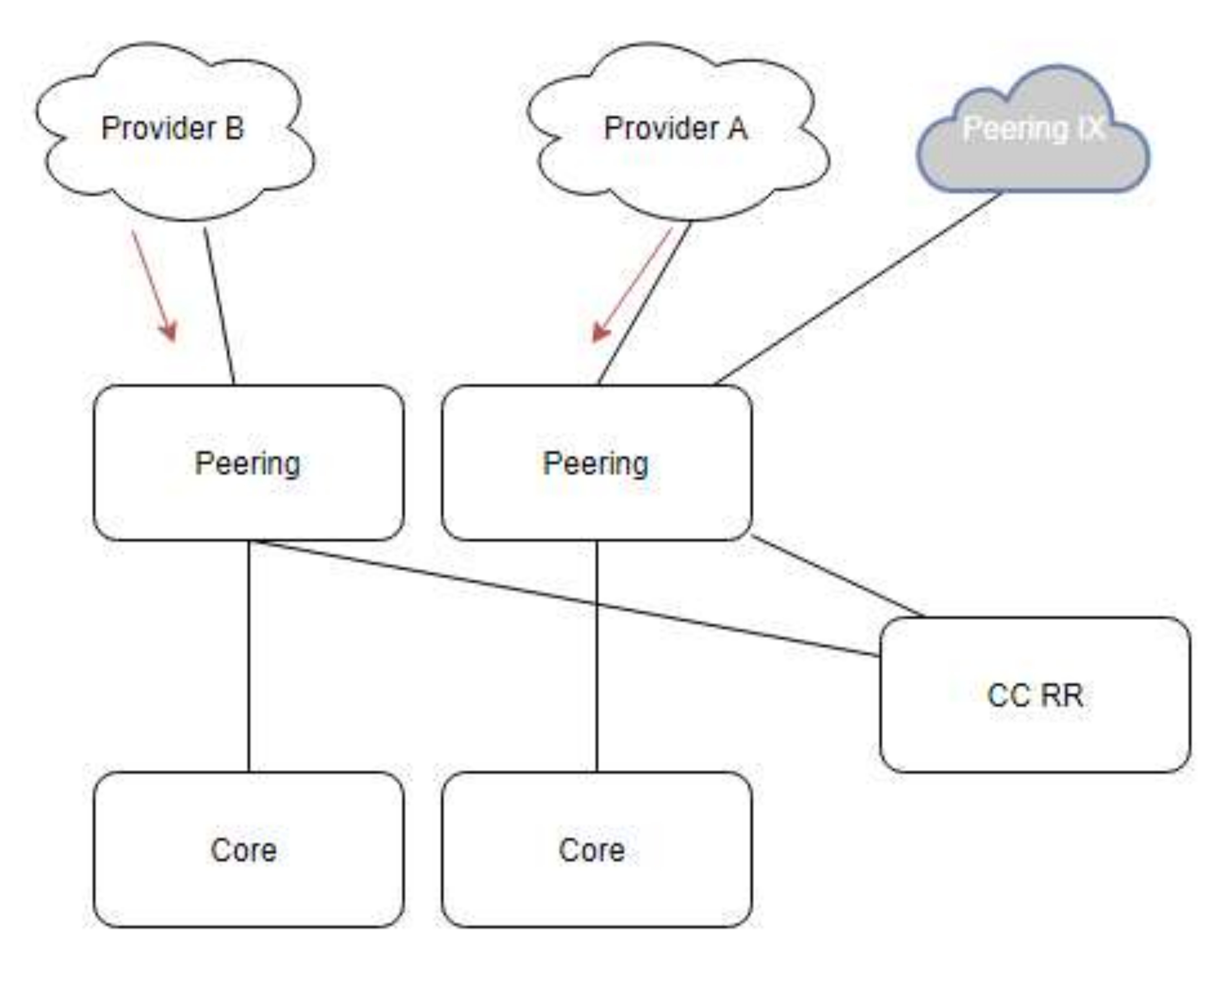

Surviving a DDOS attack

- Attack coming in..
  - Goes on though to your network.
  - Until its identified!

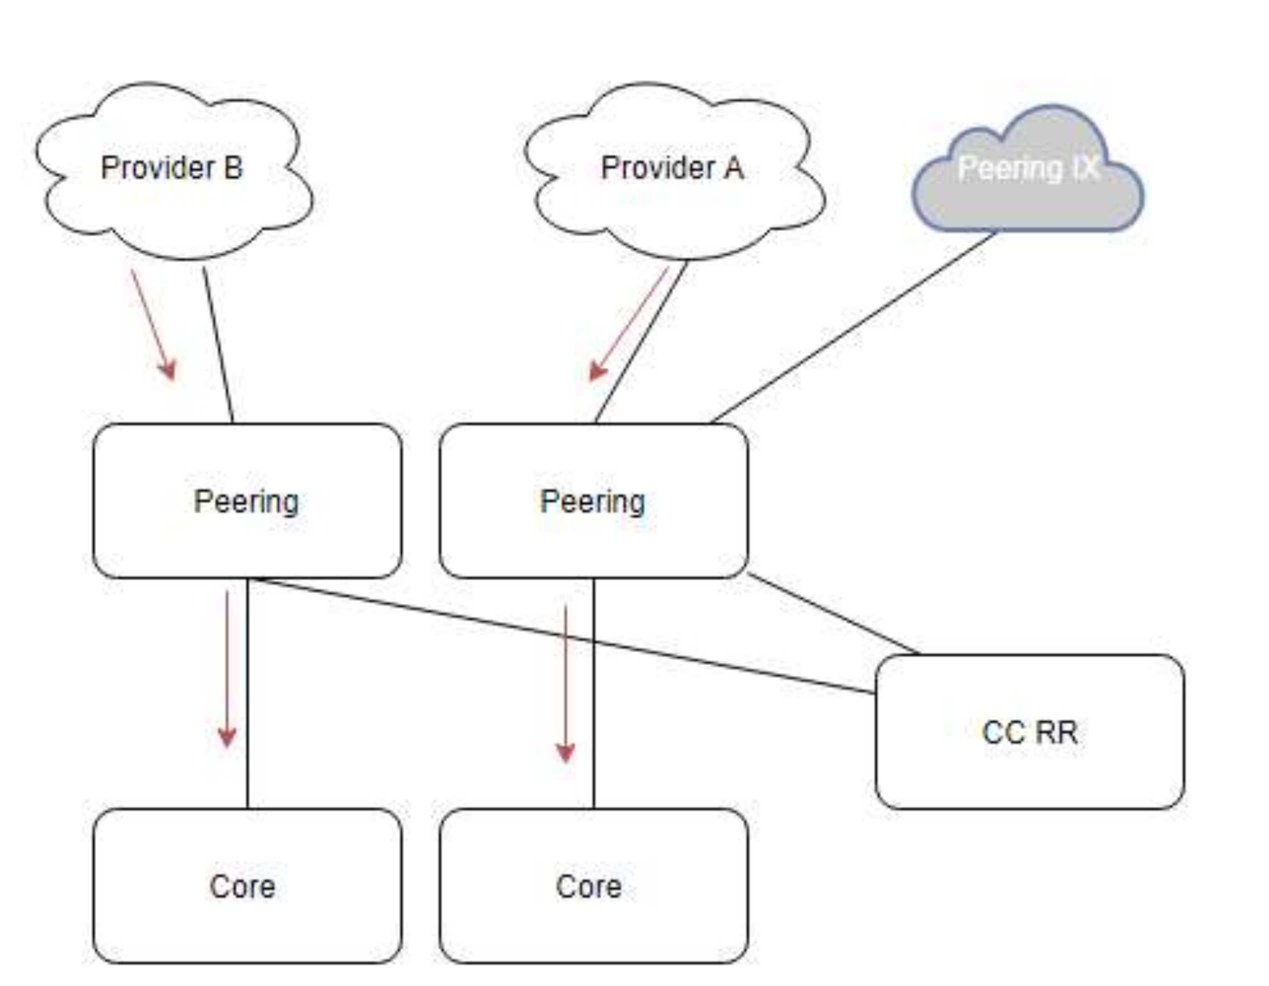

Surviving a DDOS attack

- Attack coming in..
  - Your rules on your peering router identify who is being attacked.
  - As well as who is attacking.
  - We add addresses to the address list, then we start to block data!

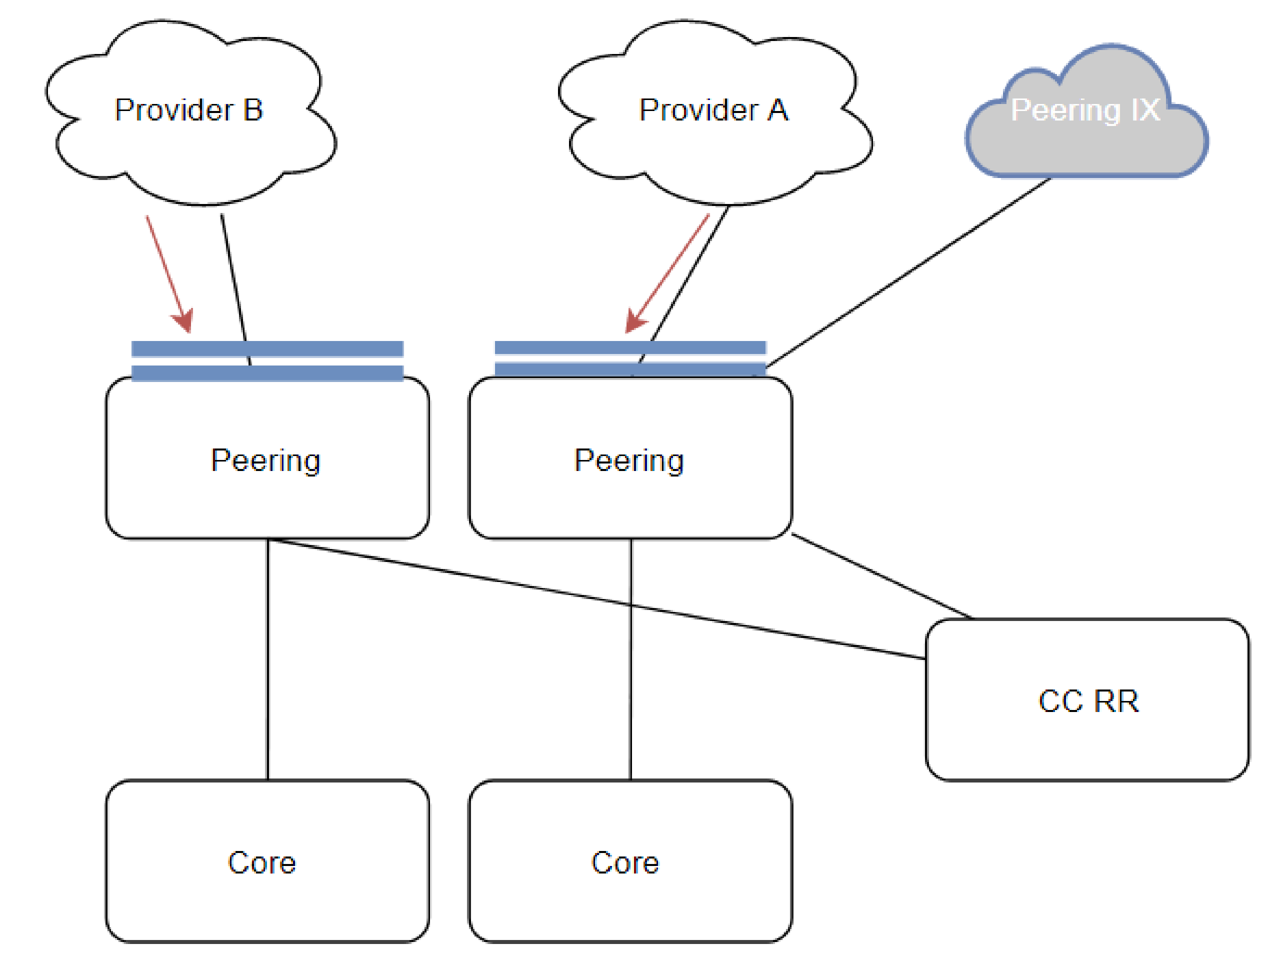

Surviving a DDOS attack

This peer connection is only 1Gig, and is being maxed!

But traffic is being dropped at your edge, still affecting your customers!

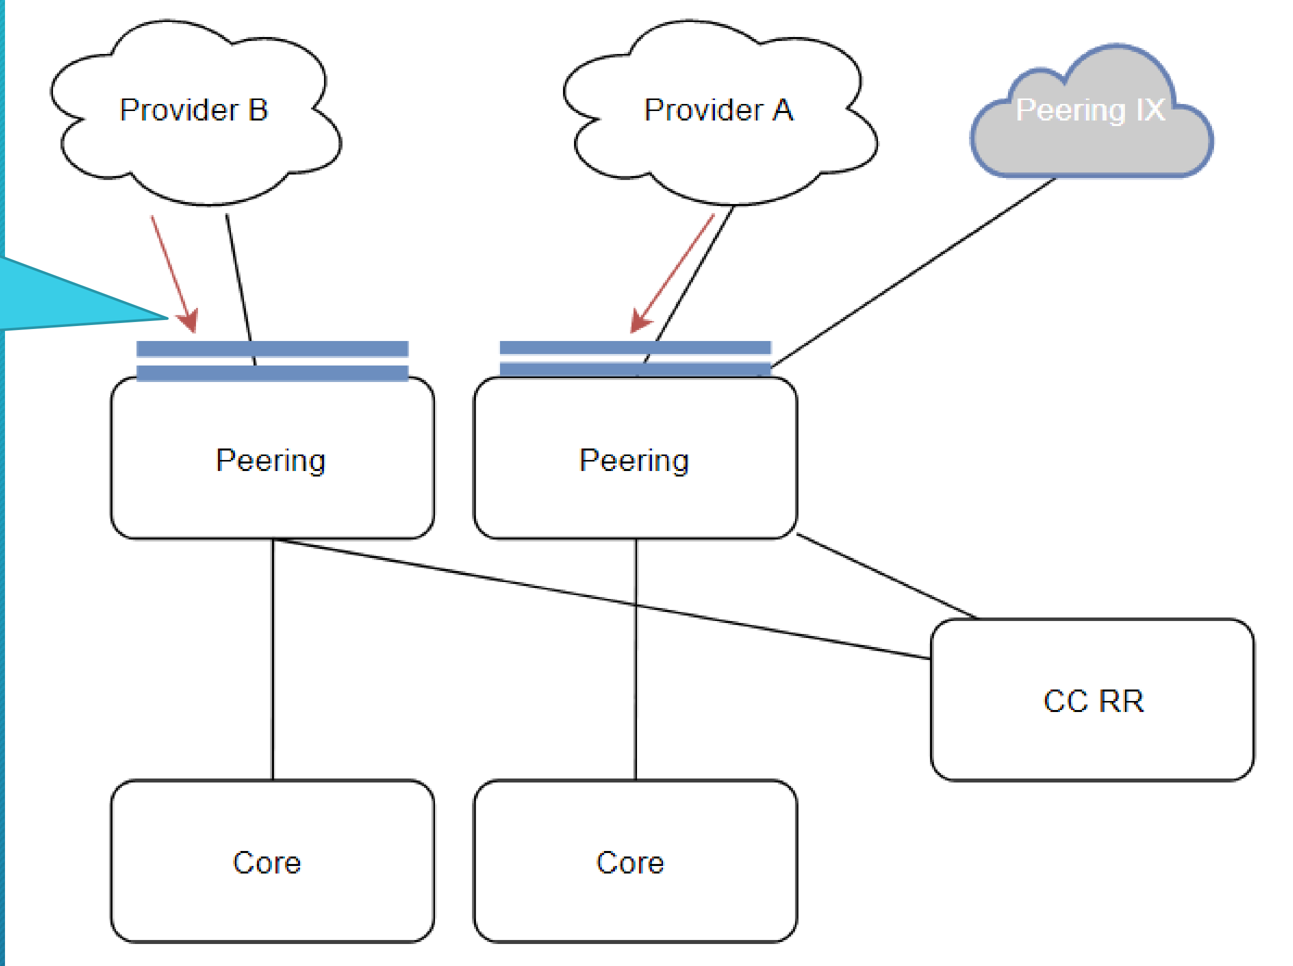

Surviving a DDOS attack

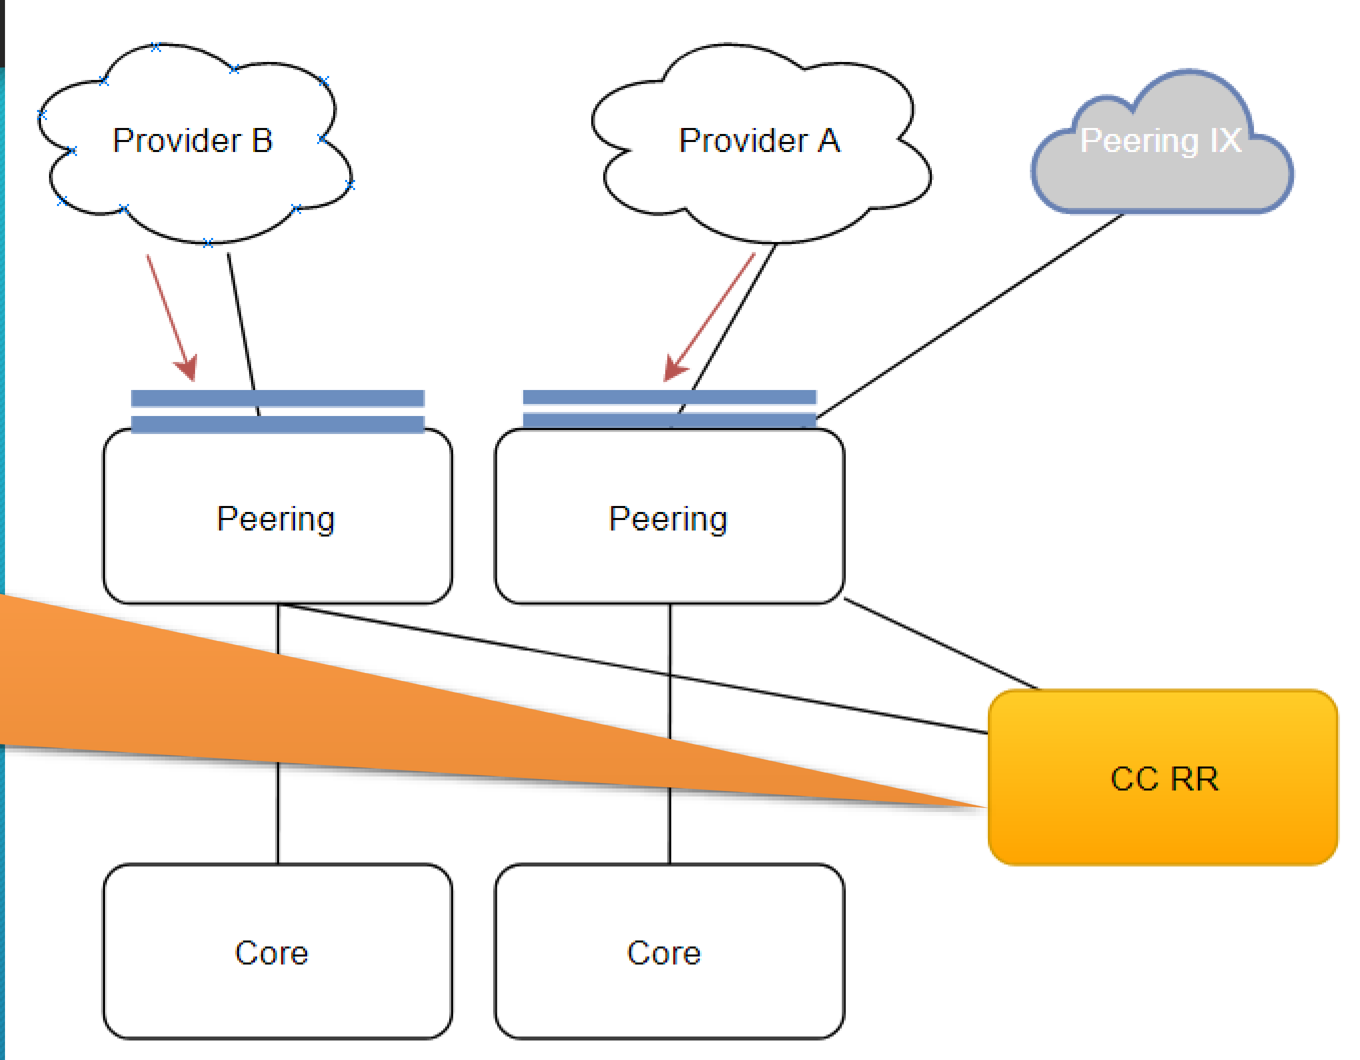

You look and both of your peering routers have identified the destination IP. The Attacked!

We place a /32 in the CCRR Router!

Surviving a DDOS attack

The /32 is sent to both your peering routers!

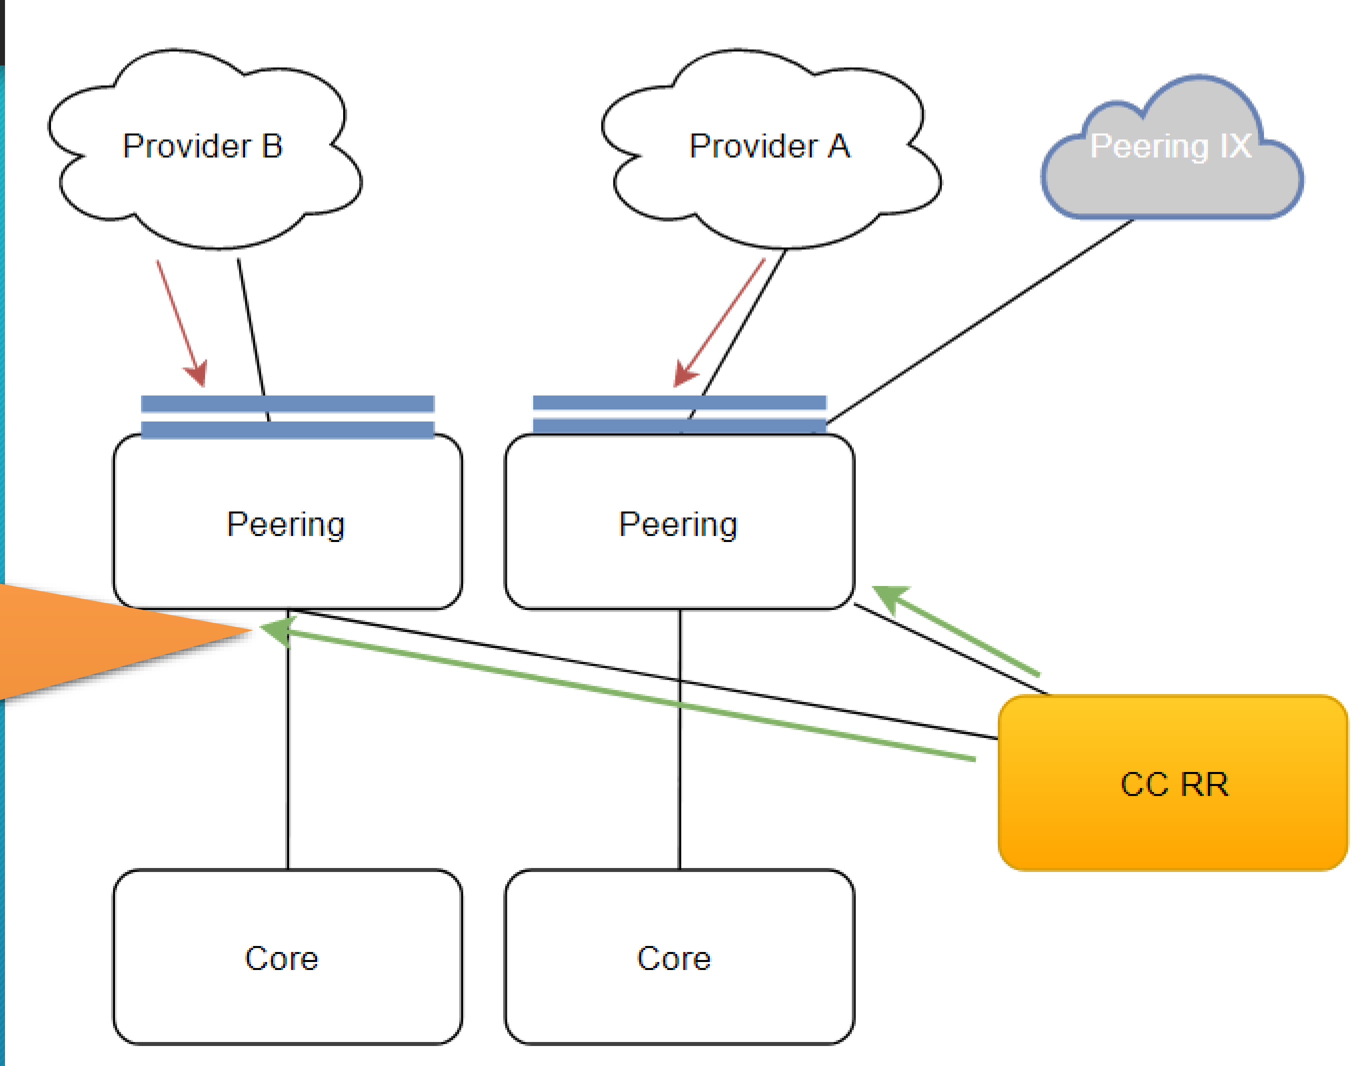

Surviving a DDOS attack

Both peering routers add the correct community string to send to the upstreams and BOOM, that traffic is not coming in anymore!

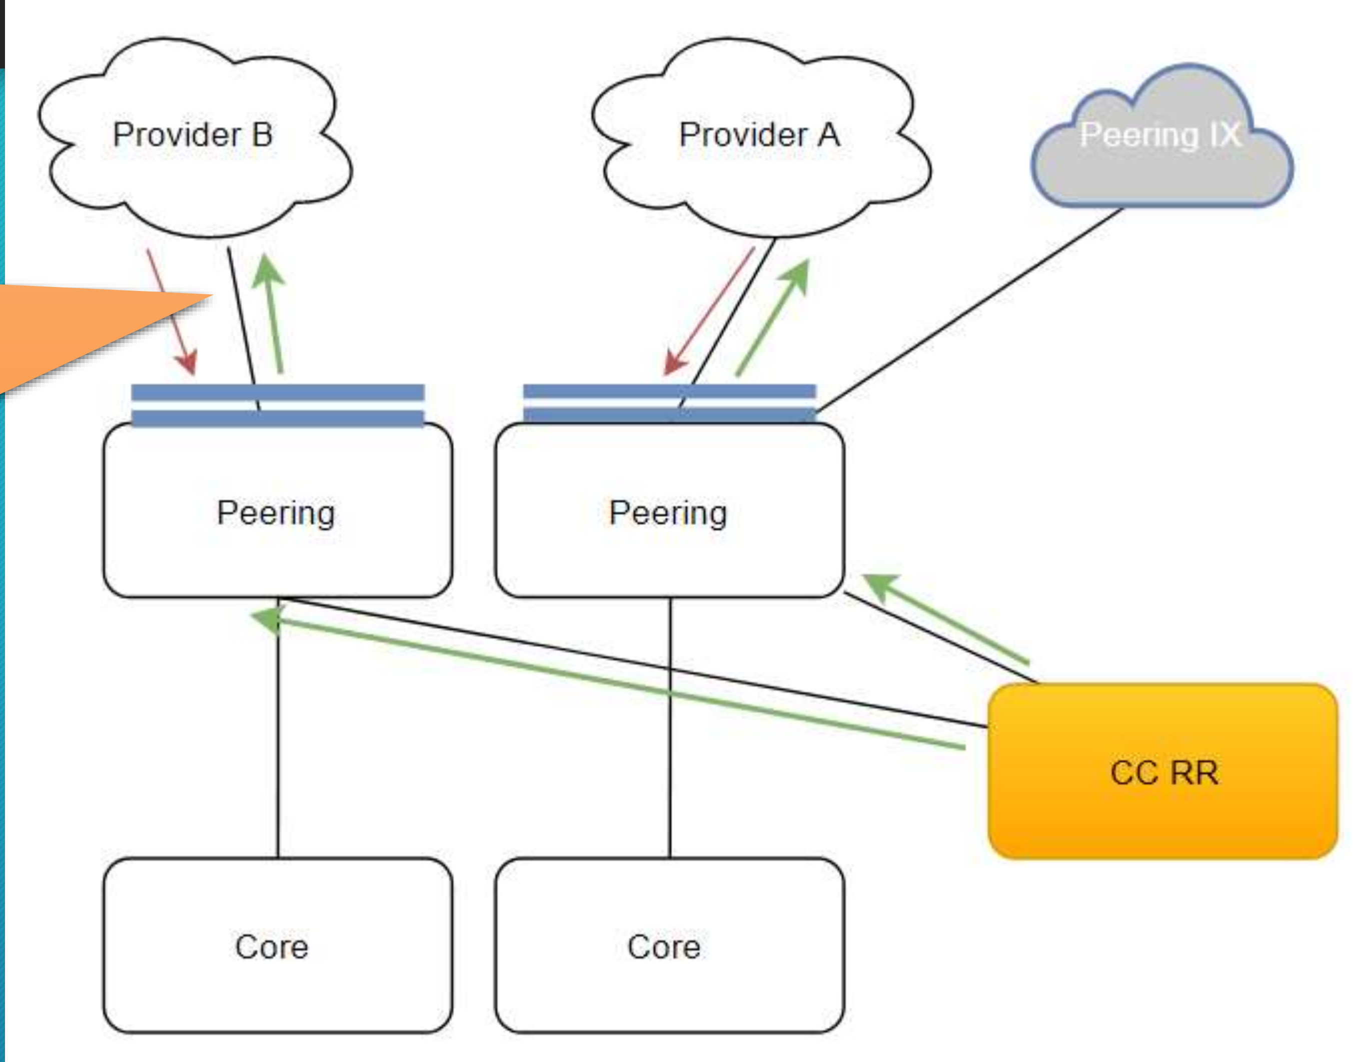

Other DDOS Services/Protection

BGP DDOS Protection

- Other methods of using BGP to survive an attack
  - DDOS Scrubbing Service
  - NetFlow BGP Monitor

BGP DDOS Protection

- Other methods of using BGP to survive an attack
  - DDOS Scrubbing Service
    - This is used in conjunction to net flow monitoring. You send all of your data to a collector, then said collector monitors for DDOS activity. If found, they announce more specific prefixes, they get your inbound traffic and then deliver them to you via tunnel (GRE or IPIP) .
    - This increases your latency but the “Service” has lots of bandwidth on-hand to handle the DDOS attack.
    - If you have many prefixes, they typically will only announce a /24 or the minimum needed to take care of the attack. Once the attack slows, they will denounce and your prefix will come directly to you as well.
    - Most of the time this is fairly costly.

BGP DDOS Protection

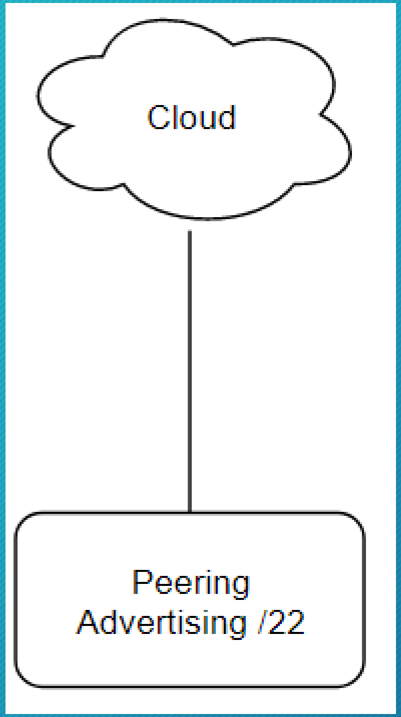

BGP DDOS Protection

- Tunnel Service
  - You send netflow data to scrubbing service
  - They are looking for a DDOS attack on any one of your 4 /24s.

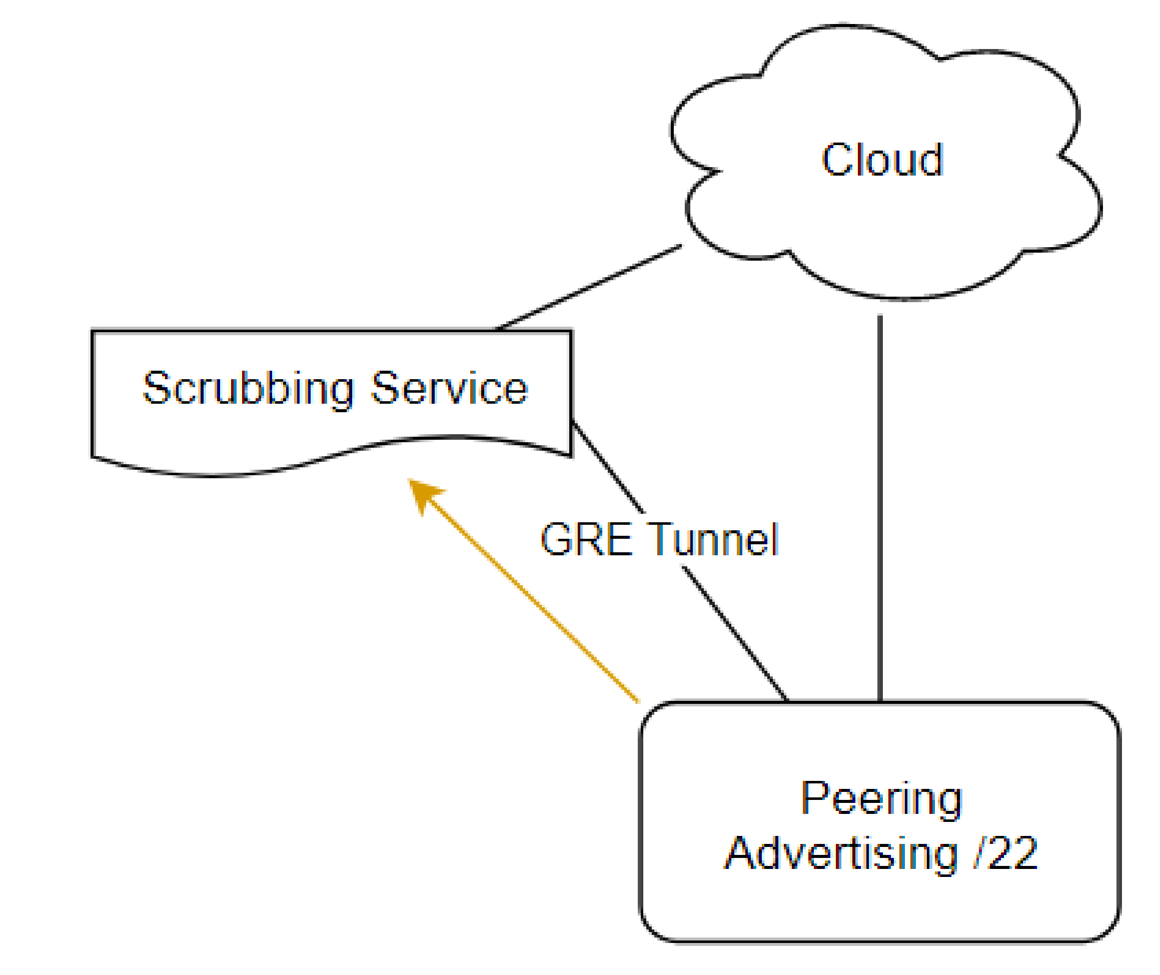

BGP DDOS Protection

- Tunnel Service
  - When they find a DDOS, they send out a BGP advertisement announcing the /24 being attacked
  - BGP prefers longer prefixes, so all traffic going to your /24 that was involved in being attacked is changed over to going to the scrubbing service.
  - They then block the DDOS and send filtered good data on the /24 though to you via tunnel.

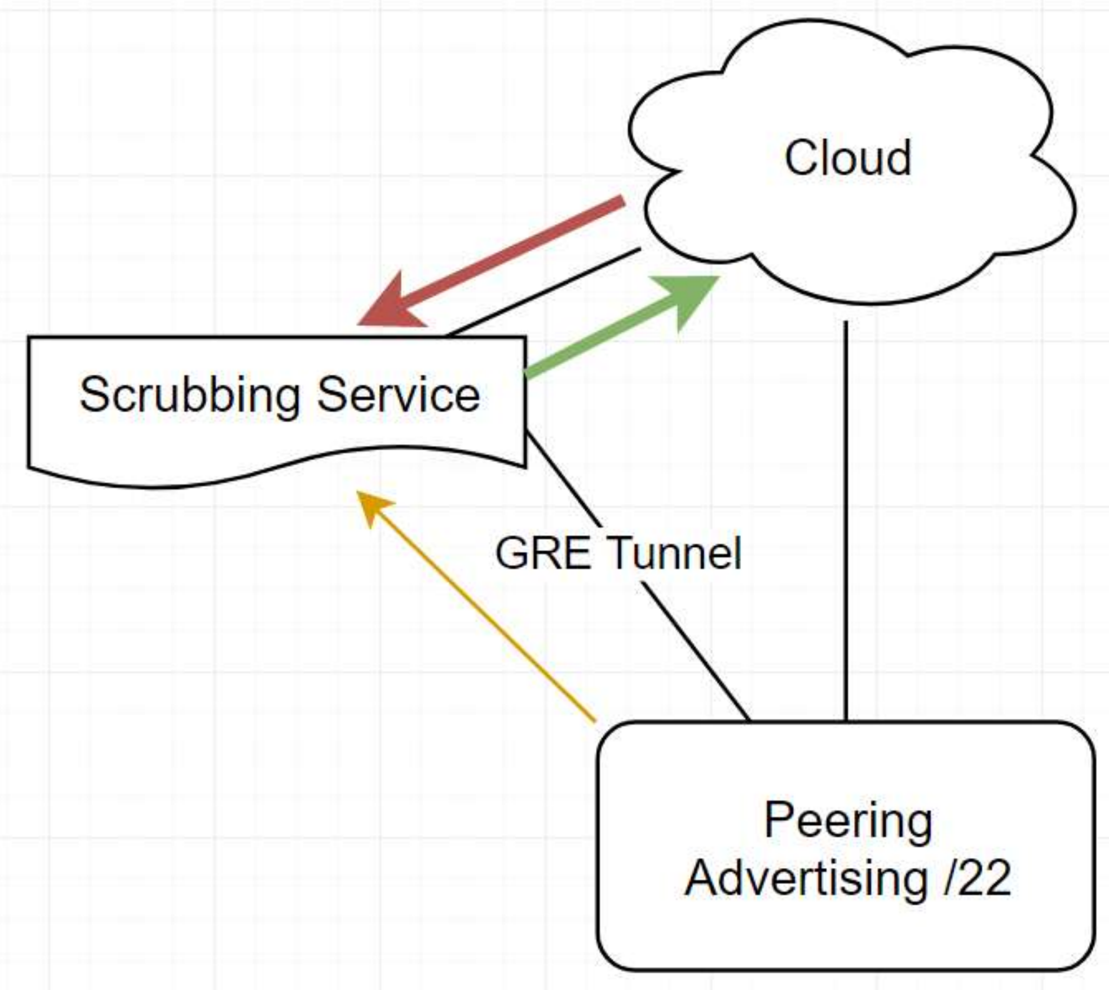

BGP DDOS Protection

- Tunnel Service
  - When the traffic ceases, they will withdraw the /24 and then your traffic will normally traverse your internet connection.

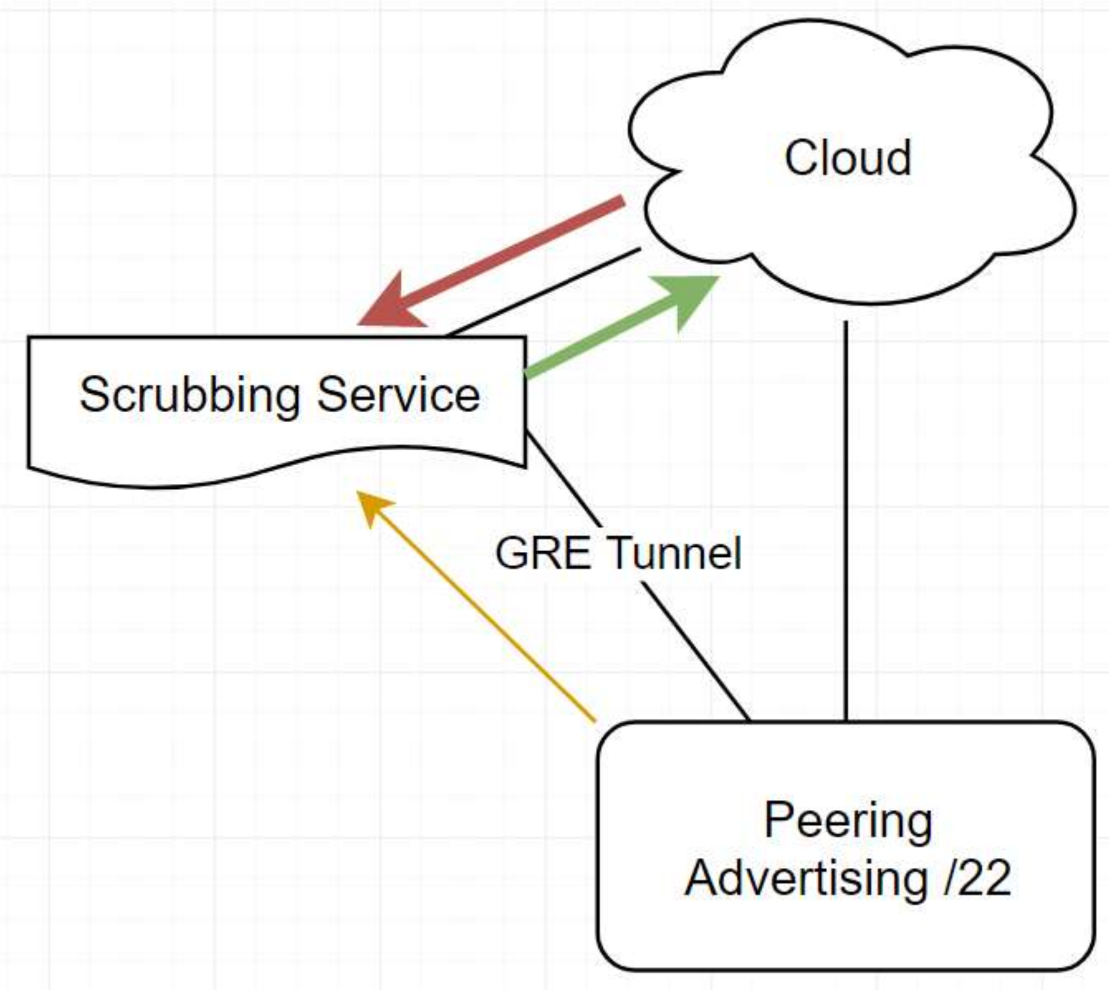

BGP DDOS Protection

- CONS
  - Typically fairly expensive to do this
  - Need technical expertise to assist in setting this up.
  - Added Latency for traffic going though subbing service.

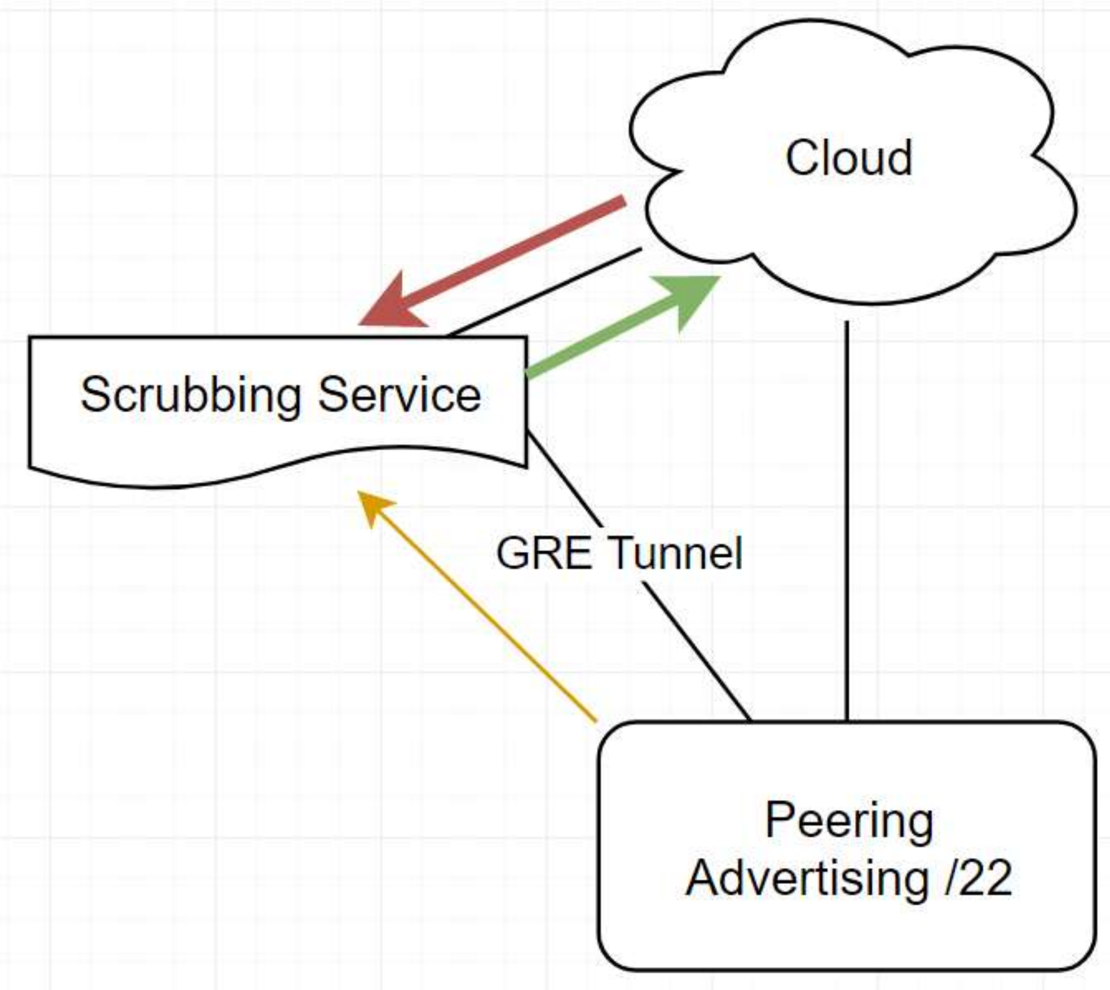

BGP DDOS Protection

- PROS
  - Does monitor your service and can filter out DDOS attacks.
  - Typically this is for more sustained attacks.
  - Works well with incoming bandwidth under 10gig.

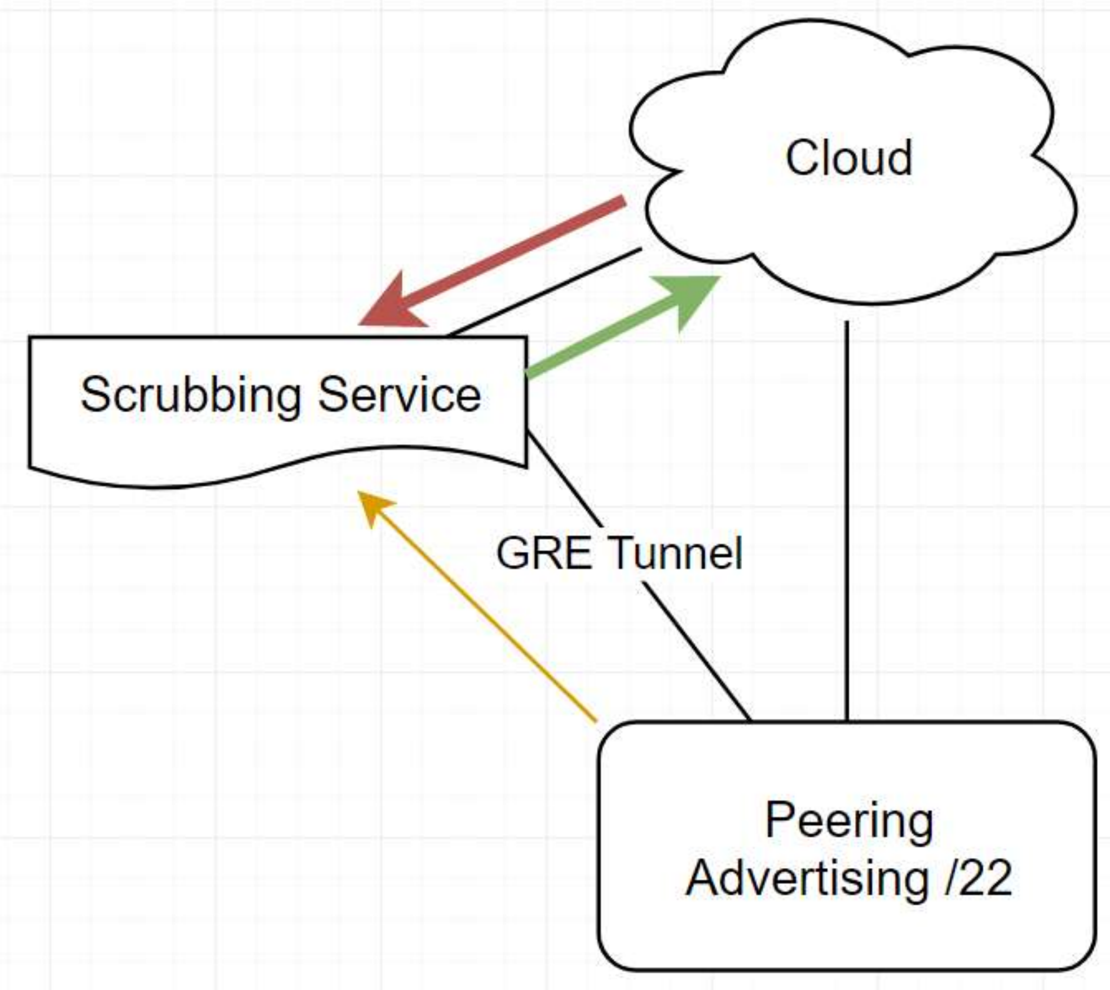

BGP DDOS Protection

- Other methods of using BGP to survive an attack
  - NetFlow BGP Monitor
    - Fairly similar to DDOS Subbing, but you have monitors that take net flow data in, and upon identifying an attack, it simply injects BGP black holes for you.
    - You still must have enough bandwidth to sustain the attack.
    - I have also used software that monitors your BGP sessions, as well as latency to common websites, and if you have a lower latency connection using the non-preferred BGP Peer, it will inject BGP adjustments to cause the non-preferred peer to become active, in effect lowering overall latencies. The software I used ran around $2500 a month.

BGP DDOS Protection

- Your server sits on-site with you
  - It receives netflow data from your peers
  - Upon it recognizing that you have a DDOS, it advertised to your peering router what to block
- Based on this, you can have it go up to your upstream if you wish.

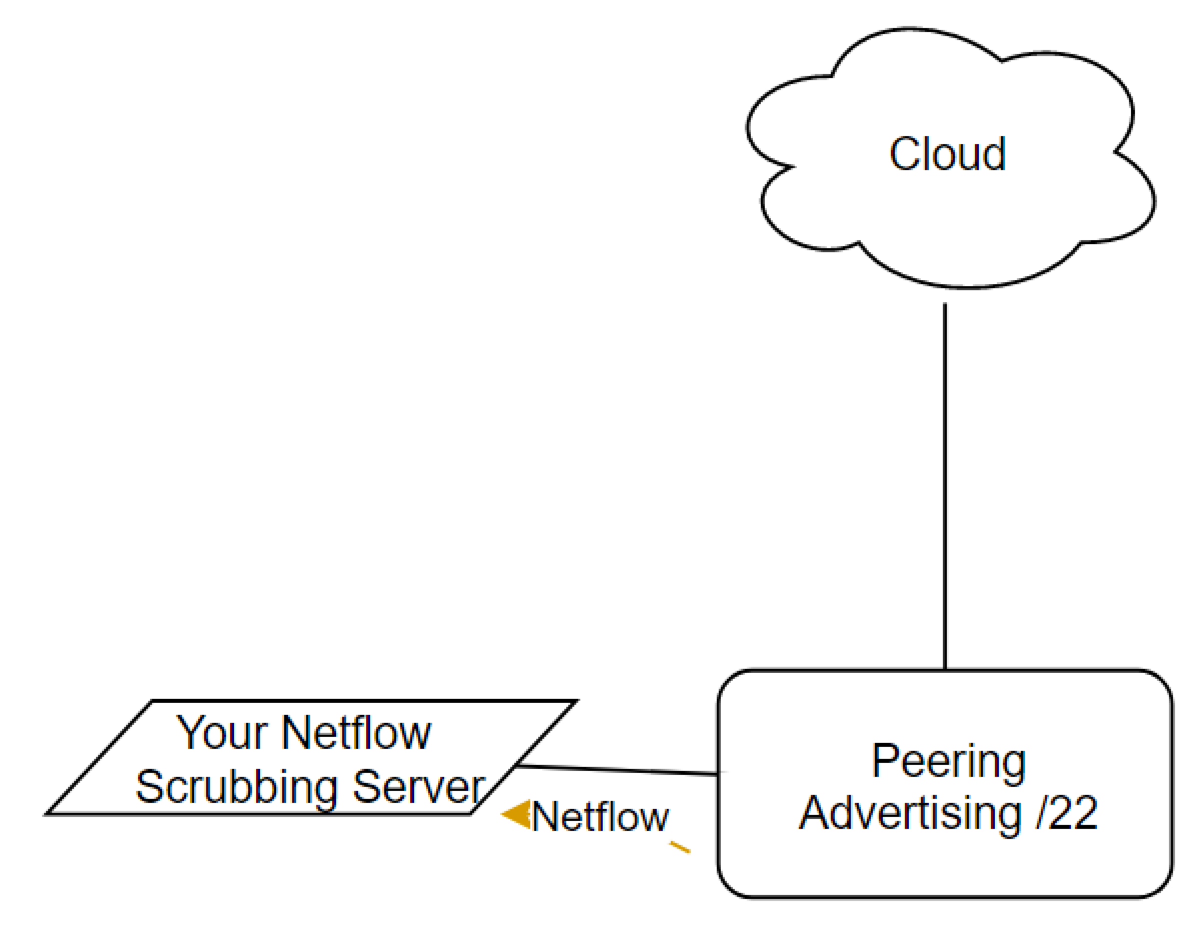

Questions!

- By Dennis Burgess
- **Link Technologies, Inc**.
  - Senior Technical Officer
- Advanced WISP Support/Engineering / Consulting
- BGP – OSPF – VPLS – MPLS - Routing
- www.linktechs.net
- Office:314-735-0270
- Email: dmburgess@linktechs.net