PROTECT NETWORK EDGE WITH BGP, URPF AND S/RTBH

by John Brown, CityLink Telecommunications, LLC

About Me

- Based in Albuquerque, NM US
- Will travel for packet$, food, and good Scotch.!
- MikroTik Trainer
- CityLink does Fiber/Wireless to Business and Home
- Ran IANA's L-Root DNS for 3 years
- Built several global ATM networks while in SV
- CityLink has around 800 BGP peers
- LAX, SJC, ASH, AMS-IX(soon)
- We offer 1Gig for $595 and 10Gig for $1995

What can we protect?

- We can protect the Edge
- We can protect the Core
- We can protect based destination address
- We can protect based on source address
- We can prevent source spoofed IP addresses!

Protecting Other Networks

- First, lets help protect other people's network
- Source Spoofed Packets
- BCP 38 https://tools.ietf.org/html/bcp38
- http://www.bcp38.info/index.php/Main_Page
- Some people don't do it
  - Some think its to hard
  - Some don't know
- JUST DO IT ! ☺

Where to Prevent Source Spoofed

- Where do you have control over this ?
  - Access Edge ?
  - Distribution Network?
  - Inter-Provider Edge?
  - Peers ?
  - Transit ?

How do we prevent bad packets?

- ACL's
  - You can write ACL's (Filters) and apply them to each interface in your network.
  - Place them on the access edge
  - Easy to maintain even when you have 100's of routers (NOT)
- Is there a better way ??

Are there easier, better ways ?

- BGP
  - Blackhole based on destination
  - Remote Trigger with BGP Peers

- uRPF (https://www.ietf.org/rfc/rfc3704.txt)
  - Typically done in hardware
  - Very low impact to CPU
  - Has some interesting side benefits

BGP Blackhole, how it works

- You and your BGP peers (transit, peers) agree.
- Special community tag will signal routes to BH
- Say Community 666:666
- You send your BGP peer a route with 666:666
- They then adjust next-hop for that route to be BLACKHOLE, or NULL, etc (vendor dependent)
- Traffic won't goto that prefix anymore.
- You've pushed the problem upstream.
- But your victim can't get ANY traffic!! Lesser of two

BGP Blackhole

- Use BGP to tell others (peers, transit) to drop traffic towards a particular prefix.
- So your customer is getting DDOS'd and its impacting the rest of your network.
- You can use blackhole to tell peers to drop towards the victim.
- POOF. Victim is now really a victim, but rest of network is now happier.

How does uRPF Work

- Routers typically make decisions based on Dest IP
- With uRPF, router now also looks at Source IP.
- Two Modes (Strict and Loose)
- Router looks at Source IP and then the routing table.
- If source is reachable via the input interface, then good, else drop (Strict)
- If source reachable via any route in routing table, then good, else drop. (Loose), except null.

Filtering Spoofed from your customers

- Enabling uRPF Strict on your single homed customers
  - Will prevent spoofed packets from entering your net
  - Will prevent your customers from participating in DDOS
  - Will keep The “NET” cleaner
  - Will SAVE YOU MONEY!! Less wasted bandwidth
  - Will help your wireless network, Less wasted bandwidth
- MAKE SURE YOU TEST THIS IN A LAB FIRST
- MAKE SURE YOU UNDERSTAND IT!
- You CAN break your network, if not careful.

But I can't drop everything !

- If I drop all traffic my special customer will be dead.
- Can't I just write an ACL ??
- Sure, can you write one fast enough for 3000 random spoofed source addresses ??
- And can you apply it to 4 core routers fast enough ?
- And what will you do if hacker changes SRC_IP ?

Use uRPF LOOSE Feature

- What you can do is use a part of uRPF/Loose
- Inject route into your FIB (Forward Info Base)
- Have the next hop o that route be BLACKHOLE
- uRPF/Loose will see NxtHop is Blackhole and DROP.
- So you inject the Src_IP's of the bad traffic with NxtHop as Blackhole and Poof, traffic FROM (Source) goes away at your edge.
- You can use tools to real-time create these injects

Use BGP to Inject

- You can use BGP to inject sources you want to drop
- Have a "Injector Machine"
  - Put routes you DO NOT want to receive from (Sources)
  - Tag those routes with two communities
  - NO-EXPORT and Say 65000:666
- Your BGP production routers will peer with Injector
  - The have a a bgp-in filter from your injector that sets next hop to BLACKHOLE / NULL, etc
- Route is updated at Injector, all of your BGP edge drops from that source.

CAUTION CAUTION CAUTION

- You *MUST* (https://www.ietf.org/rfc/rfc2119.txt)
  - Make sure you do NOT redistribute these SOURCE prefixes to other Peers / Transit.
  - DANGER DANGER DANGER
  - Make sure your BGP-OUT filters DROP all prefixes with your 65000:666 community tag.

What can I use as Injector

- My favorite is exabgp
  - https://github.com/Exa-Networks/exabgp/wiki
- Its FREE
- It WORKS
- It supports cool things like JSON
- It has an API
- You can control it via many methods
- You can run it on something simple / low cost

Data sources for injectors

- NetFlow (nfsen, etc)
  - Lorenzo Busatti
  - http://mum.mikrotik.com/presentations/EU16/presentation_3049_1456752471.pdf
- Bro
  - https://www.bro.org/index.html
- Suricata
  - https://suricata-ids.org/

THANK YOU

- Congratulations to all the new Mikrotik Trainers
- Thank you MikroTik for an awesome product and awesome people!
- Shout Out to Tom Smyth, my inspiration for presentation!
- Shout Out to Lorenzo Busatti, inspiration for Tik Flows☺
- I can be reached at:
  - mum2016@citylinkfiber.com
  - +1.505.938.6309
  - PGP FINGERPRINT: 5A6126CF