CLOUDFLARE®

Spoofing and Denial of Service: A risk to the decentralized Internet

DDoS: The real story with BCP38

Tom Paseka

APRICOT 2017

Global Network

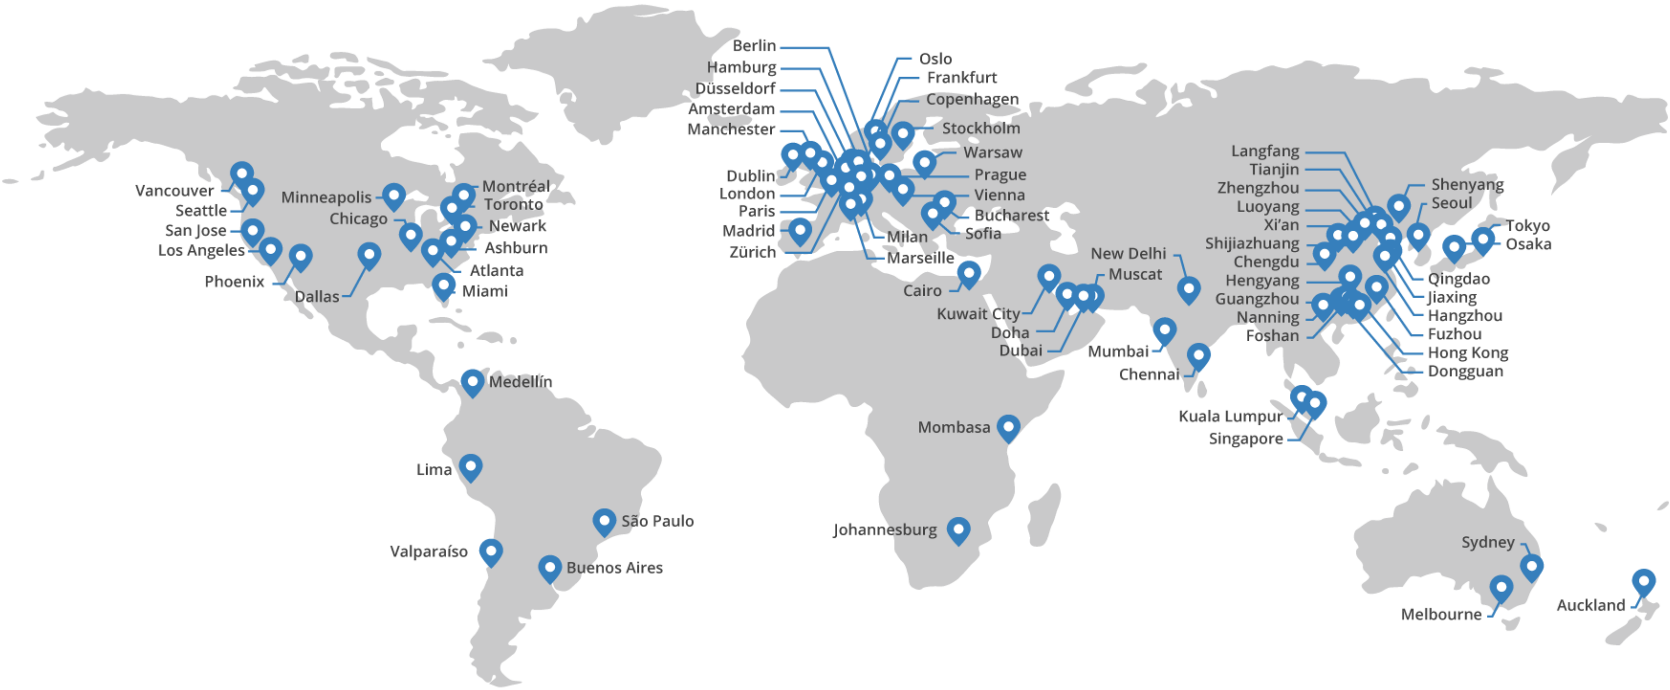

CLOUDFLARE

Content Neutral

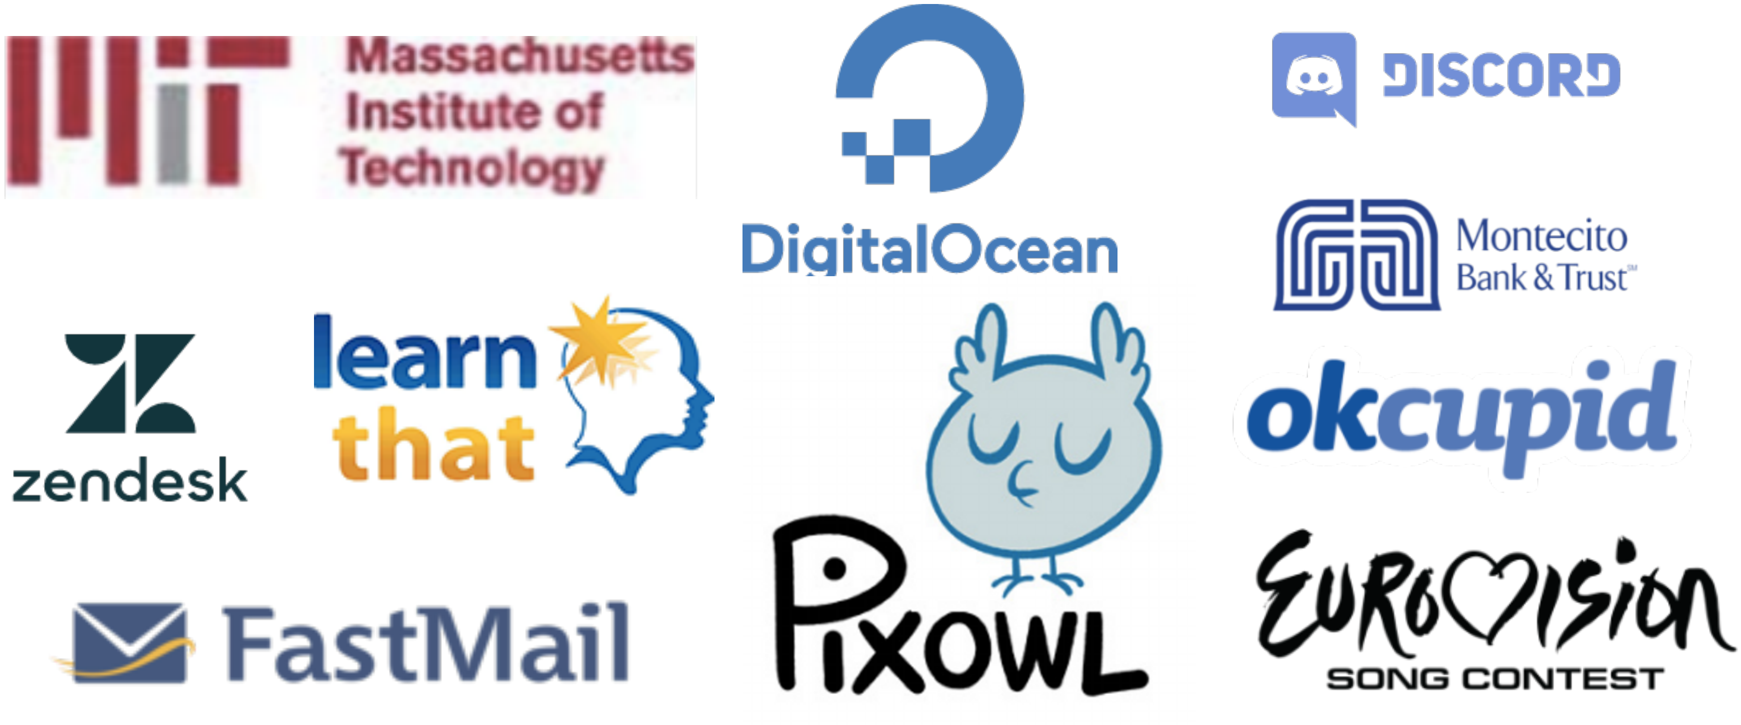

Daily Attacks

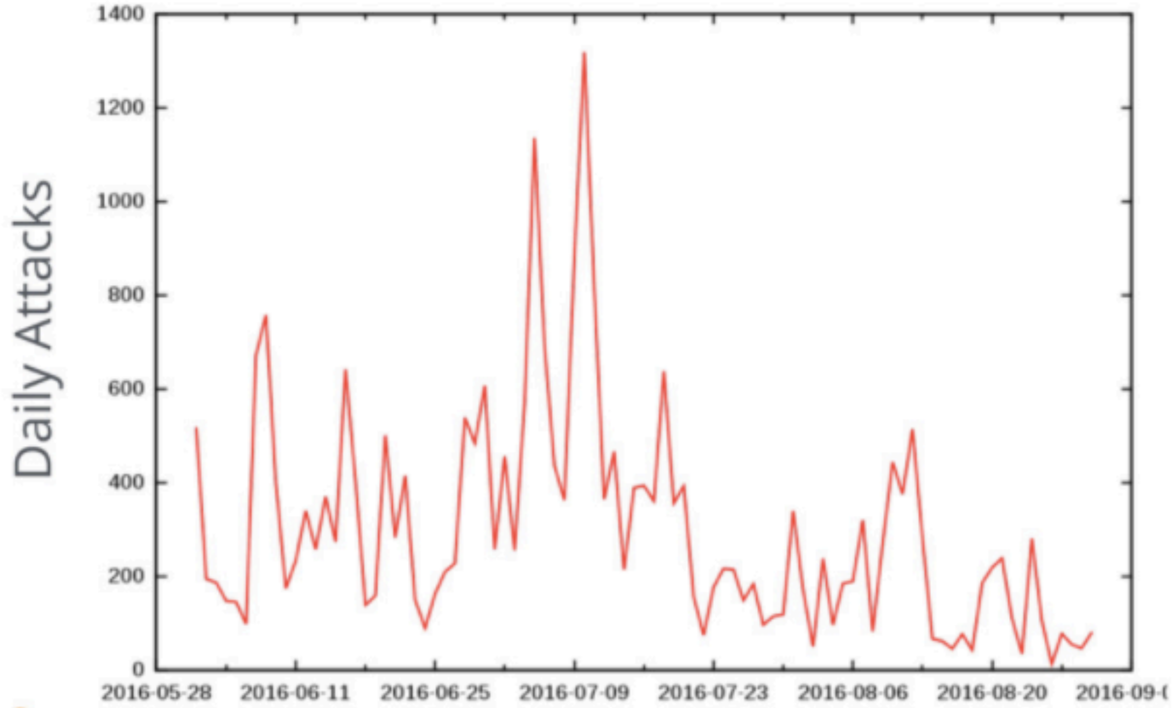

CLOUDFLARE

Daily Attacks

- Because we have such a broad view of the internet, we see a lot of attacks
- This graph is showing count of different attacks
- Sometimes, seeing more than 1,400 unique attacks daily

We have to solve attacks

Record Breaking Attacks

Nickname	Type	Volume
SNMP Amp	SNMP Amplification/Reflection	80Gbps
Spamhaus	DNS Amplification/Reflection	300Gbps
"Winter of Attacks"	Direct	400Gbps
IoT	Direct	500Gbps+

Record Breaking Attacks

- Around 5 years ago we saw some SNMP reflection attacks
- Cable modems from a very large Cable ISP in North America were reflecting SNMP walks towards us
- We then saw the infamous “Spamhaus” attacks. Attacks which were directed at us and internet infrastructure, resulting in impact to hundreds of thousands of internet users
- From September 2016, the “IoT” attacks, most famously the Mirai (未来) botnet with attacks breaking 500Gbps

Most big attacks have a few things in common

Flood of IP Packets

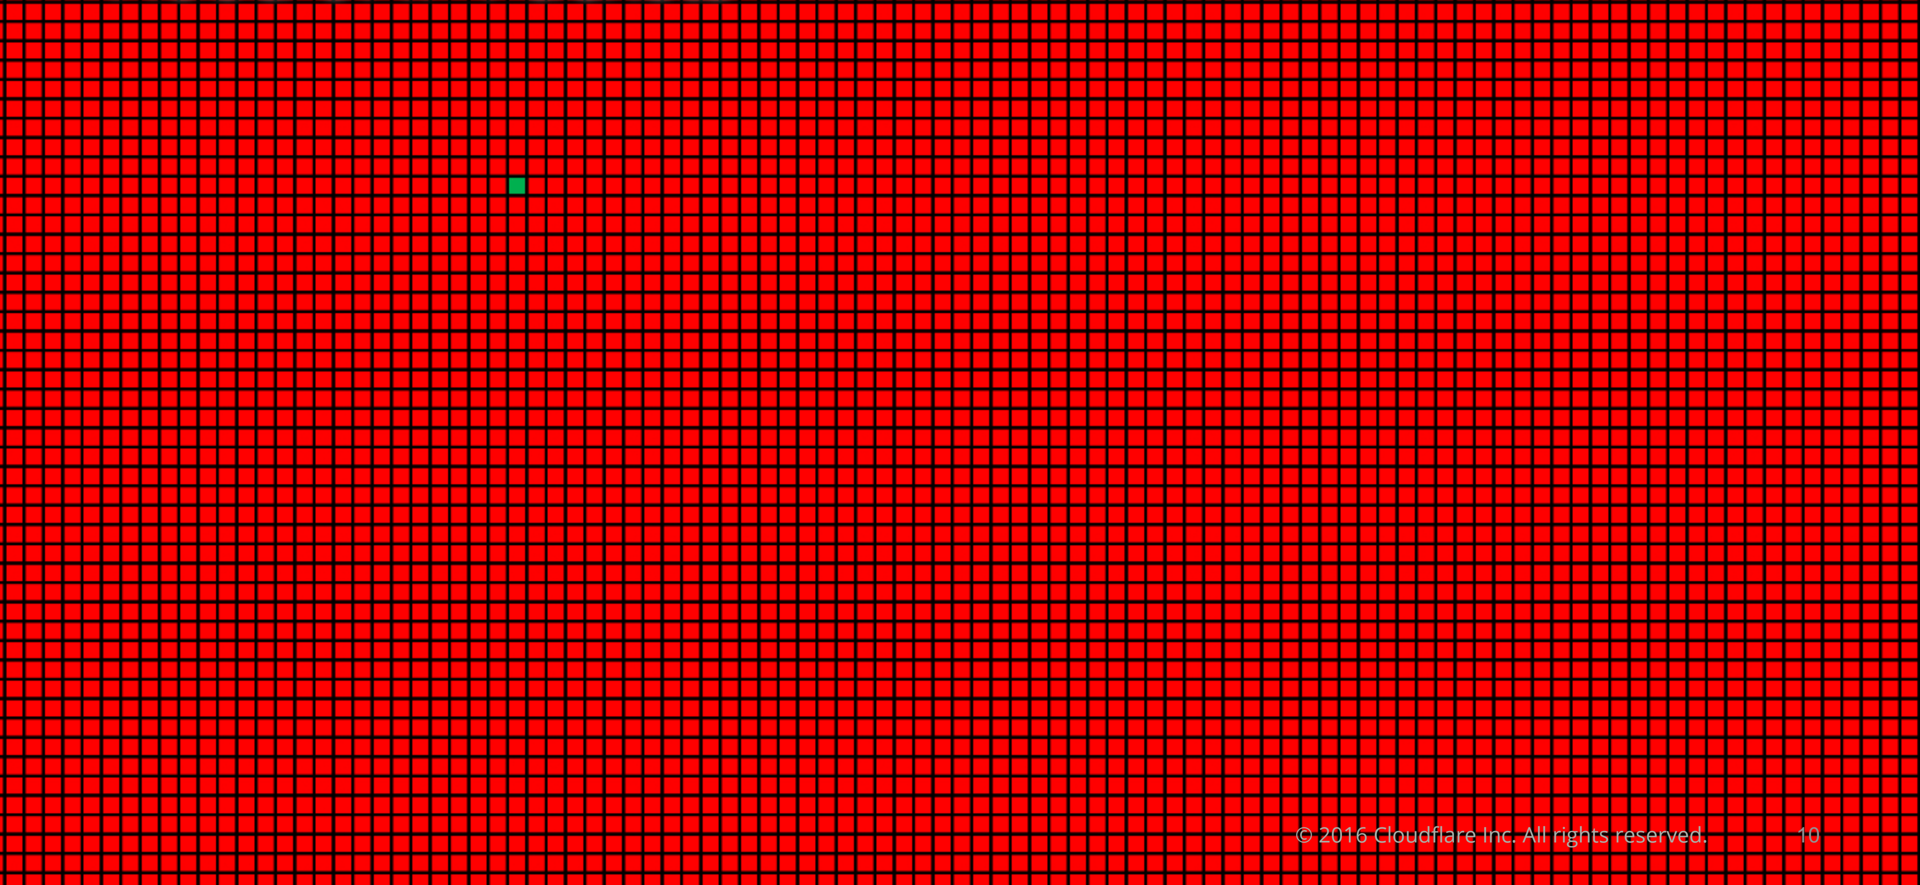

IP Spoofing

Spoofing Enables Impersonation

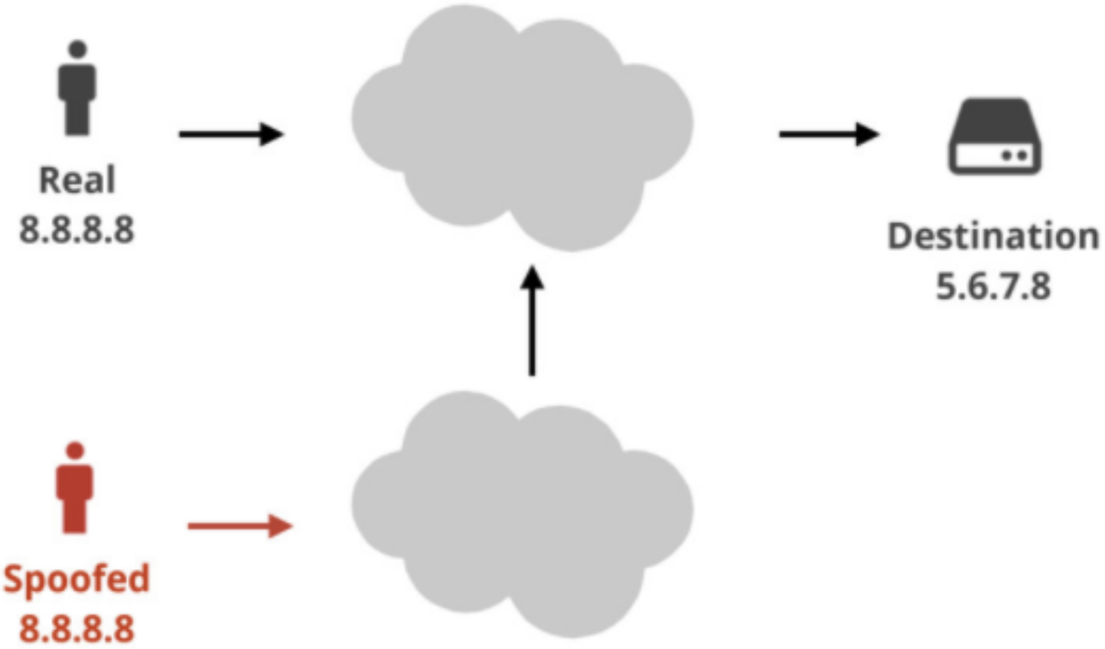

Spoofing?

- Why is spoofing an issue?
- This is my good friend Walt Wollny
- Let’s say, he was assaulted, but it was by masked assailant
- Without removing the mask, there can’t be legal retribution

May 2000: BCP38

BCP38

- BCP, Best Common Practice #38 was published in May 2000
- It gave guidance on how to configure your network to prefer spoofing
- This document is nearly 17 years old, why it isn't engrained yet?
- Vendors Faults? Operators Fault?
- Regardless, IT'S. JUST. NOT. THERE.

Caida Spoofer Stats

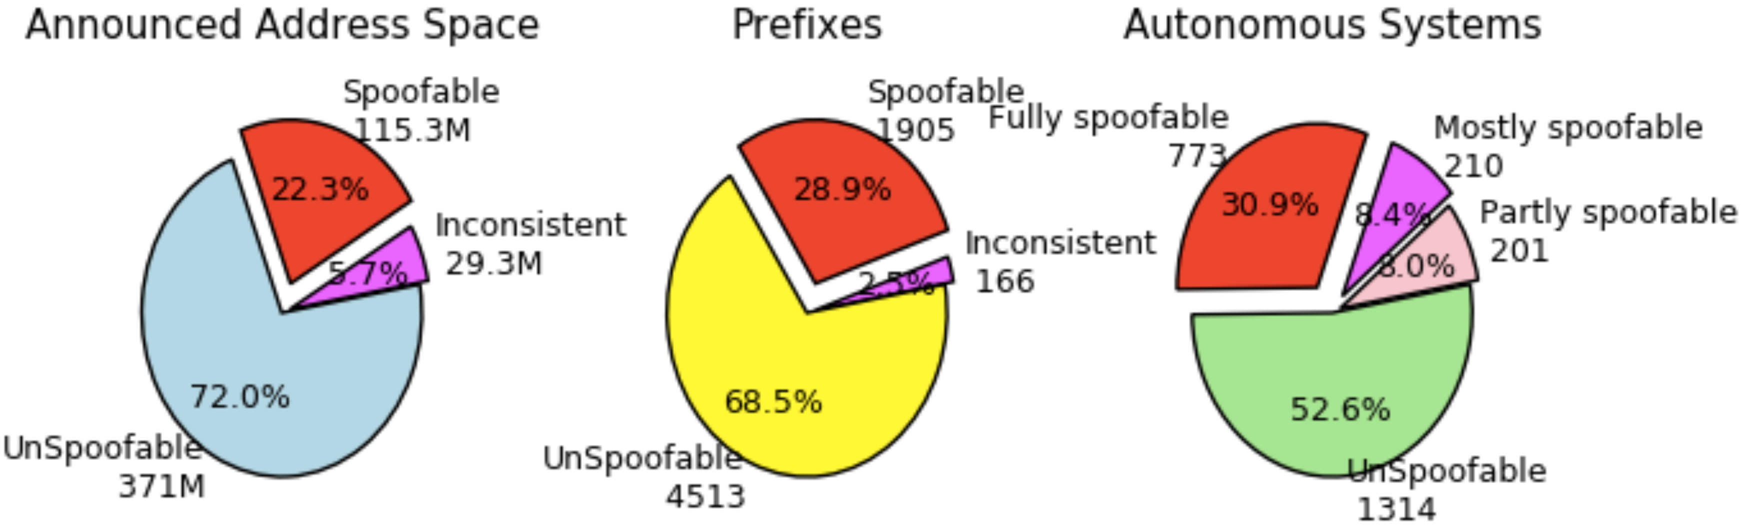

Updated: Feb 2017. Source: https://spoofer.caida.org

Filter close to the source

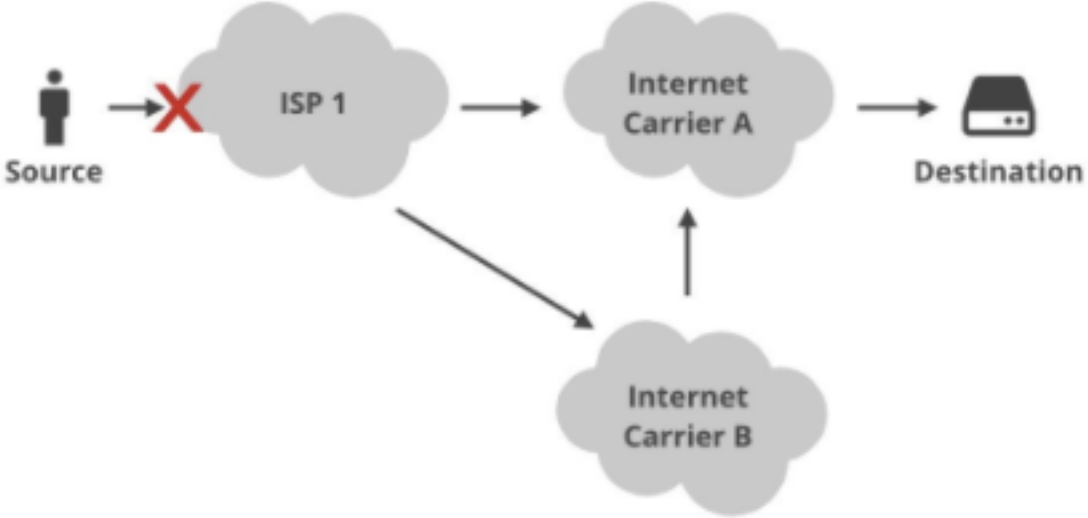

Filter close to the source

- Filtering at the ingress from your customer is really how to stop filtering
- You should also be filtering at the egress if your network for multiple layers, incase of some misconfiguration
- Unicast Reverse Path Forwarding doesn’t scale well
- What about simple ACLs?
- Yet this still isn’t there!

IP Spoofing:

- Enables Impersonation
- Isn’t solved

IP Spoofing

1. Tracing back is impossible
2. Allows sophisticated attacks

IP Spoofing

1. Tracing back is ~~impossible~~ very hard!!!
2. Allows sophisticated attacks

Where did the attack come from?

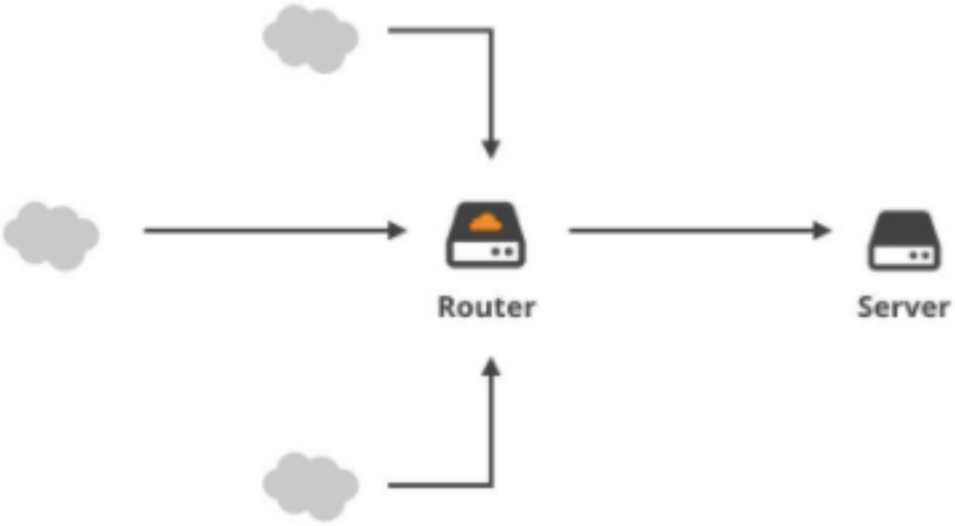

Where did the attack come from?

- The “Server” in this slide, gets attack traffic
- It has one link out, to its router, so we know it came from the ‘router’
- But from there, where did it come from?
- There are multiple input interfaces, which one could be sending the traffic? Which network?
- We can trace this down a bad way, by looking at graphs

Identifying interfaces

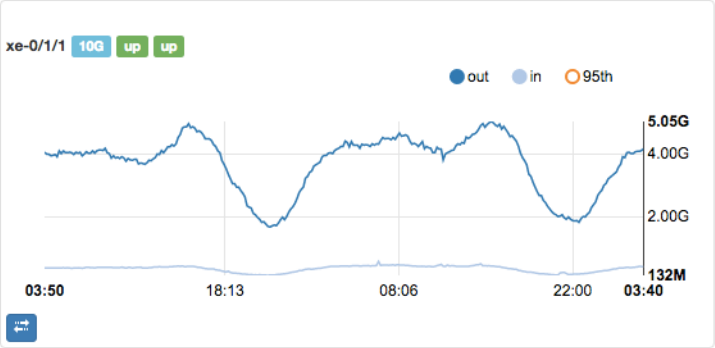

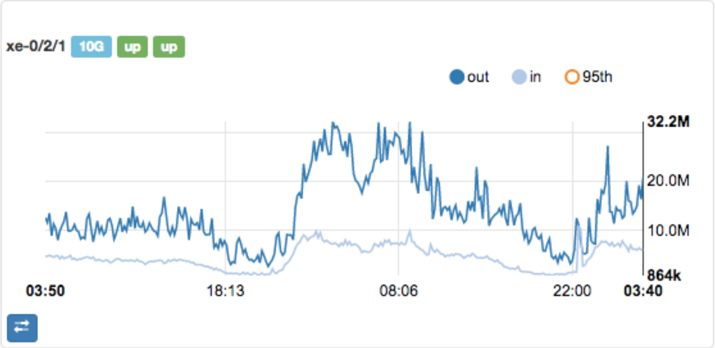

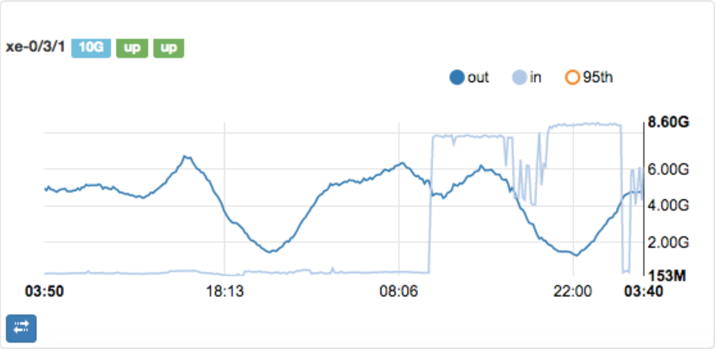

Identifying interfaces

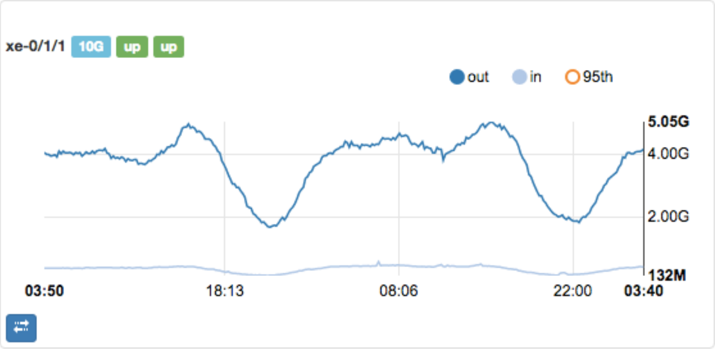

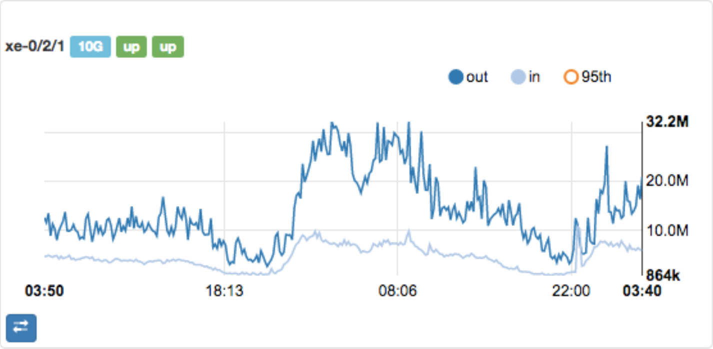

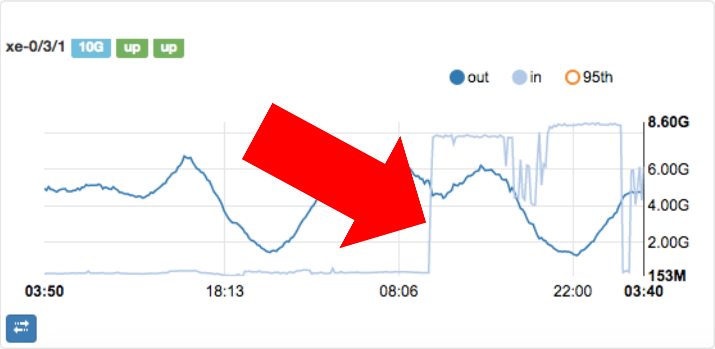

What’s on the other side of the Cable?

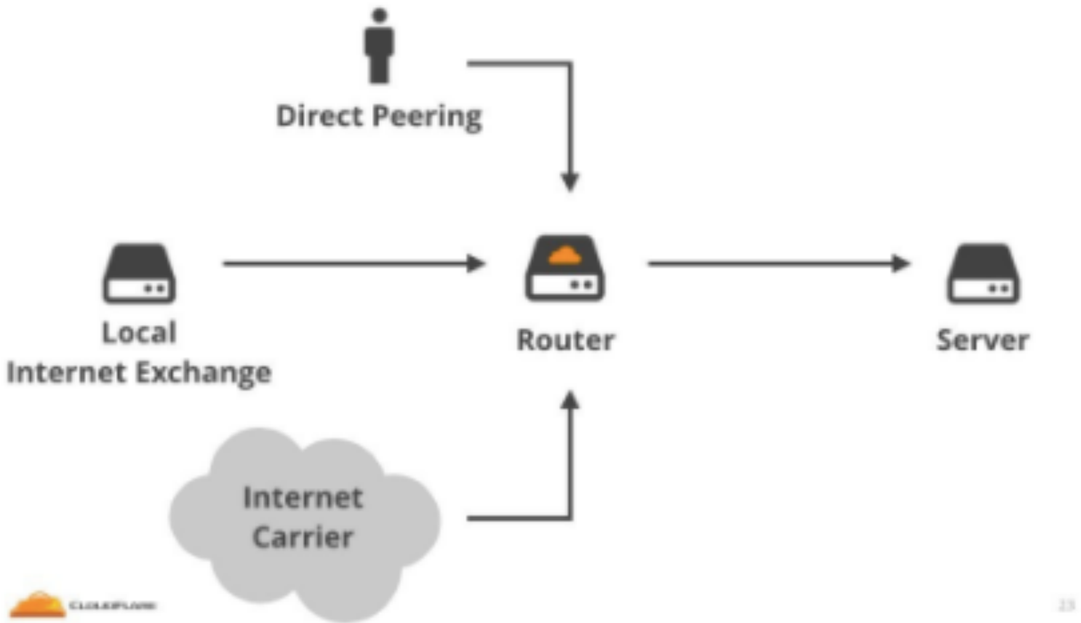

Direct Peering

Local Internet Exchange

Router

Server

Internet Carrier

What’s on the other side of the Cable?

- For most internet networks, there are several types of input sources:
    - Direct Peering: Where you have a single network and their customer cone on that interfaces
    - Internet Exchange: many networks connected to a single fabric. Possible hundreds of direct networks and thousands of in-direct networks
    - Internet Carrier / Transit Provider: The whole Internet

1. Direct Peering

Direct Peering
Router

1. Direct Peering

- Where we have direct peering with another network, you have a pretty good idea of what’s on the other side
- This is going to be limited to that network and their customers
- In a case like this, it’s pretty easy to identify at least the ISP responsible for traffic

2. IXP / Internet Exchange Point

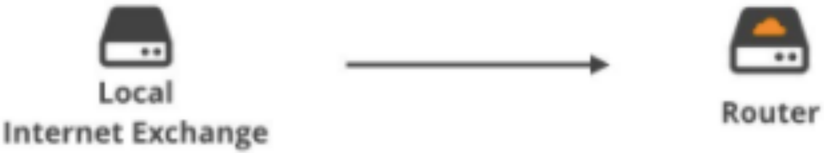

3. Transit Provider

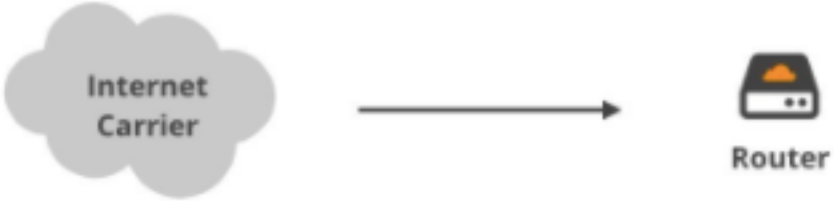

IXPs and Transit Providers

- Both of these represent an issue
- There is any number of networks where traffic could be coming from
- No easy way to identify the source over either of these
- Let’s explore a little but more about IXPs

2. IXP / Internet Exchange Point

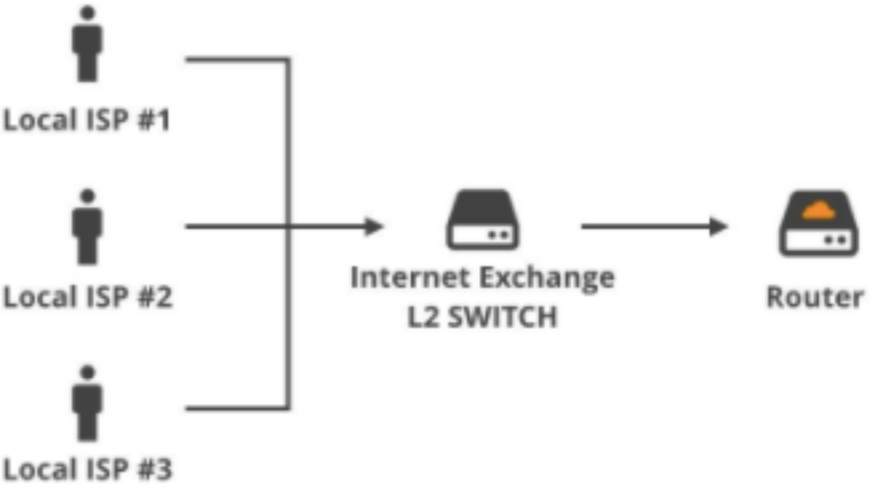

2. IXP / Internet Exchange Point

?.?.?.?

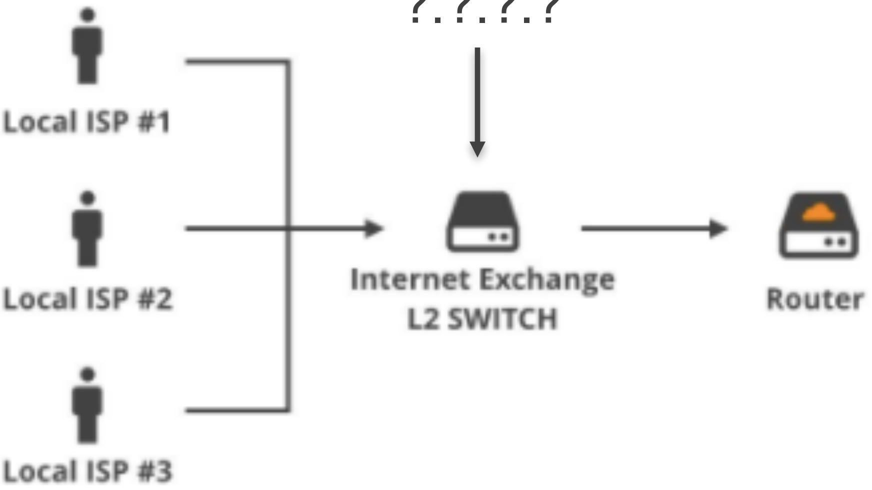

2. IXP / Internet Exchange Point

- When traffic enters the IXP, we have no idea where the source came from
- Since you’re on one big fabric, anyone can inject it
- Very hard to track back
- Some ways to trace, but poorly implemented. I’ll touch on this later.

3. Transit Provider

Src ip = 8.8.8.8

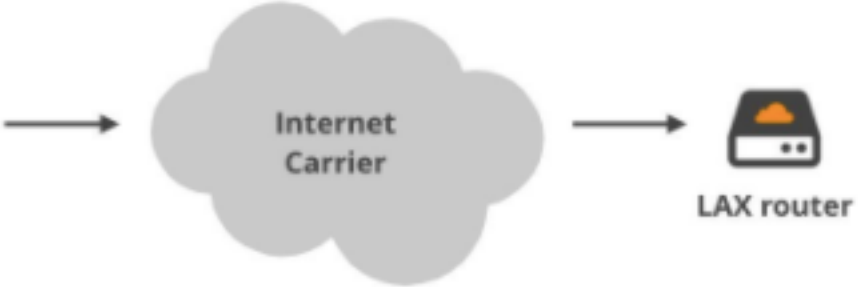

3. Transit Provider

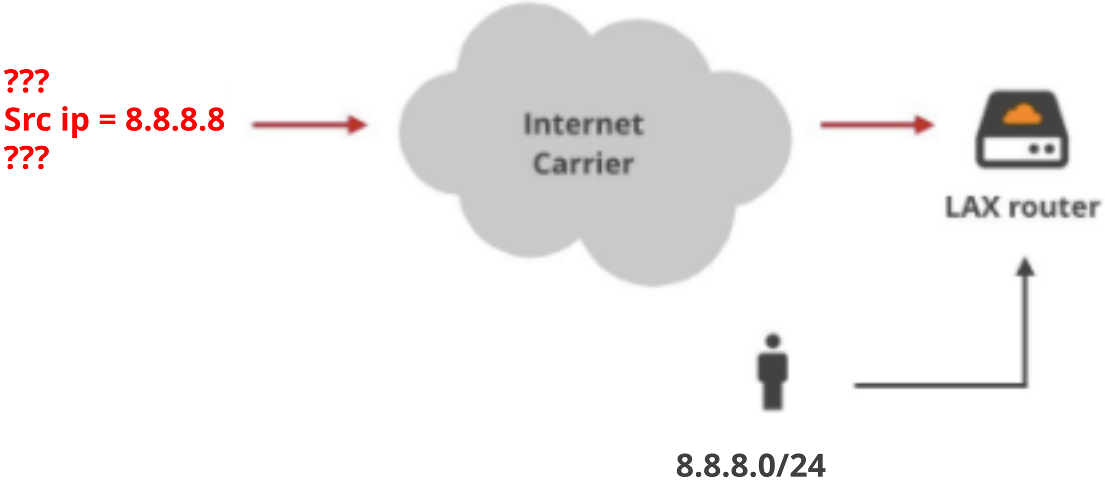

???
Src ip = 8.8.8.8
???

8.8.8.0/24

3. Transit Provider

- So, we see an attack coming from 8.8.8.8
- This is coming in over a transit provider
- But we have direct peering with the network that represents this traffic
- Why isn’t this traffic coming over the peering?
- ….Because it’s spoofed.

Lack of Attribution

IP Spoofing

Very hard!!!

1. Tracing back is ~~impossible~~
2. Allows sophisticated attacks

CLOUDFLARE

Amplification

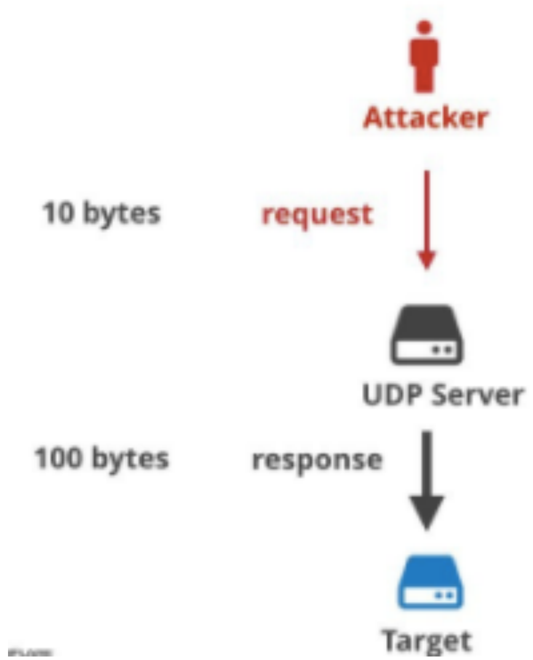

Amplification

- We know about amplification attacks, so I’m not going to go into technical detail
- The premise: Send a small request and get a big response directed at your target
- Amplification means you can knock off a service, much larger than you are, without using all your resources.

March 2013: Spamhaus

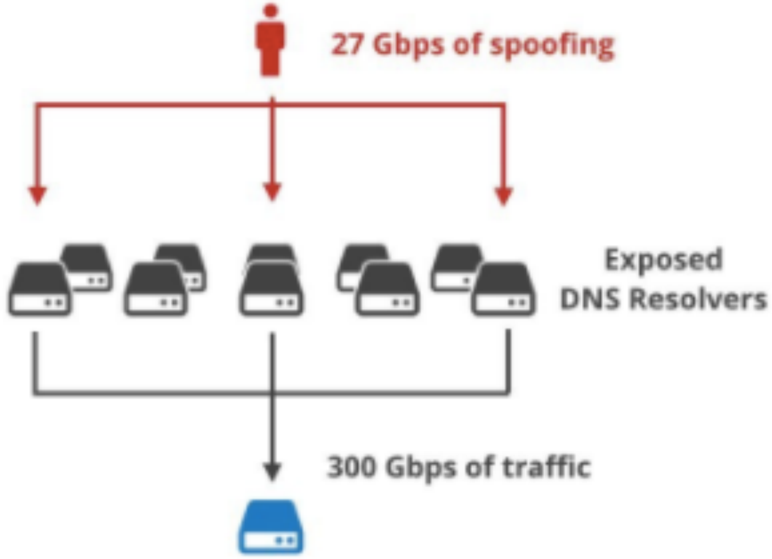

March 2013: Spamhaus

- During the Spamhaus attacks, DNS amplification was used
- Large DNS replies (eg. ANY isc.org ~4,000 byte reply to a very small query)
- 37Gbps of attack traffic was able to be amplified to 300Gbps of attack traffic

Amplification is relatively easy to block....

- ...If you have the bandwidth. (few networks can absorb hundreds of Gbps)

- Block on firewall:
  - src UDP/53 > deny

- Internet is fighting amplification sources:
  - openresolverproject.org
  - openntpproject.org

Source IP Addresses

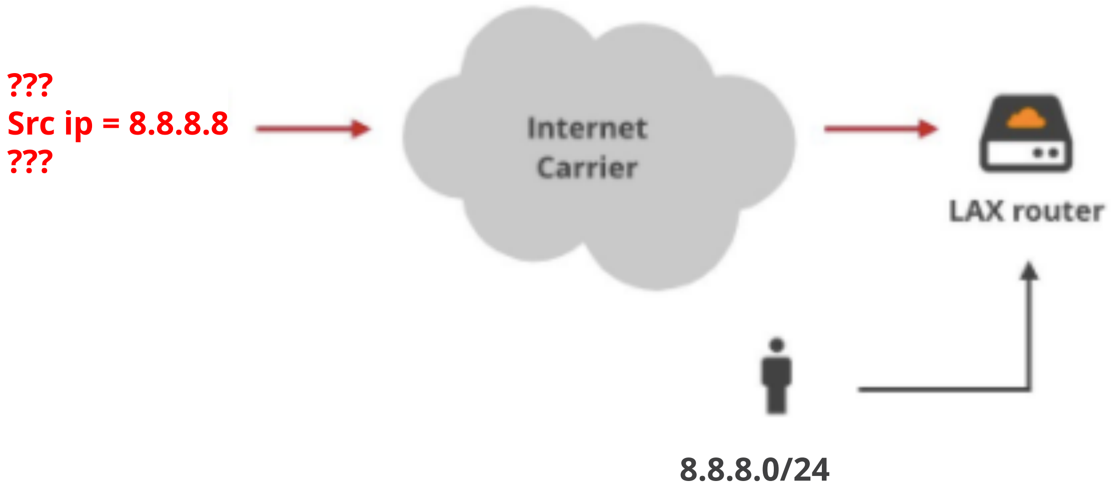

???
Src ip = 8.8.8.8
???

8.8.8.0/24

Source IP Addresses

- So, what happens when we trace the source IP address in attacks.
- Taking this lovely picture from xkcd, we see a map of what the internet is

Source IP Addresses

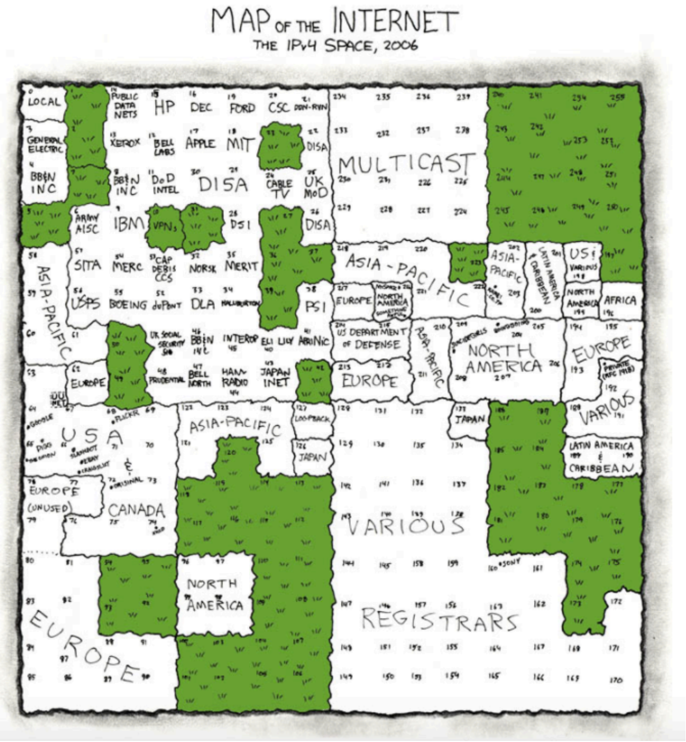

https://xkcd.com/195/

Source IP Addresses

- What does this same map look like, when we see a large scale attack?

Source IP Addresses

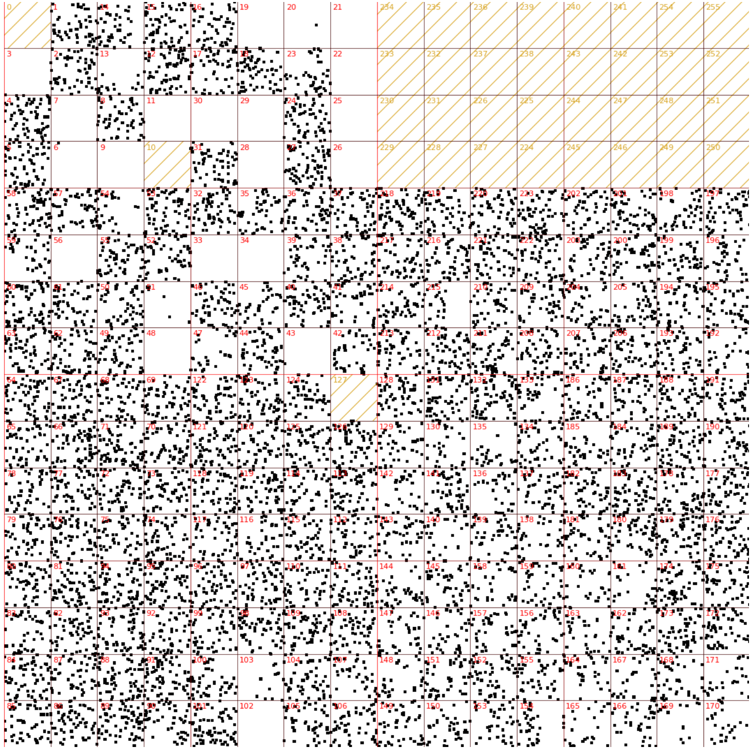

Source IP Addresses

- What about a different type of attack?

- This attack is coming from a single network, the graph on the left is the view of what is routed by that network
- The graph on the right is attack sources from that network
- Is this network doing egress filtering? Is it spoofed or all direct from that network?

Source IP Addresses

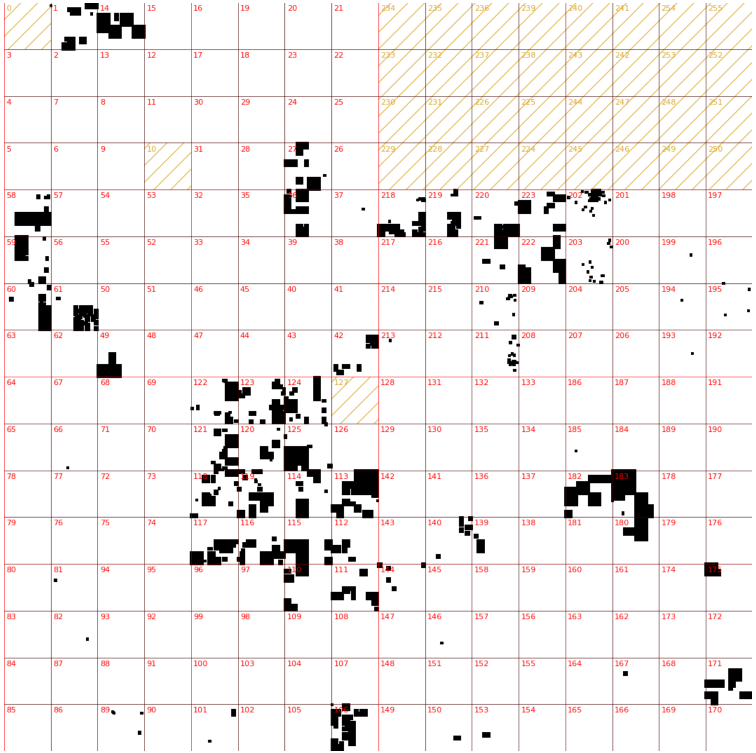

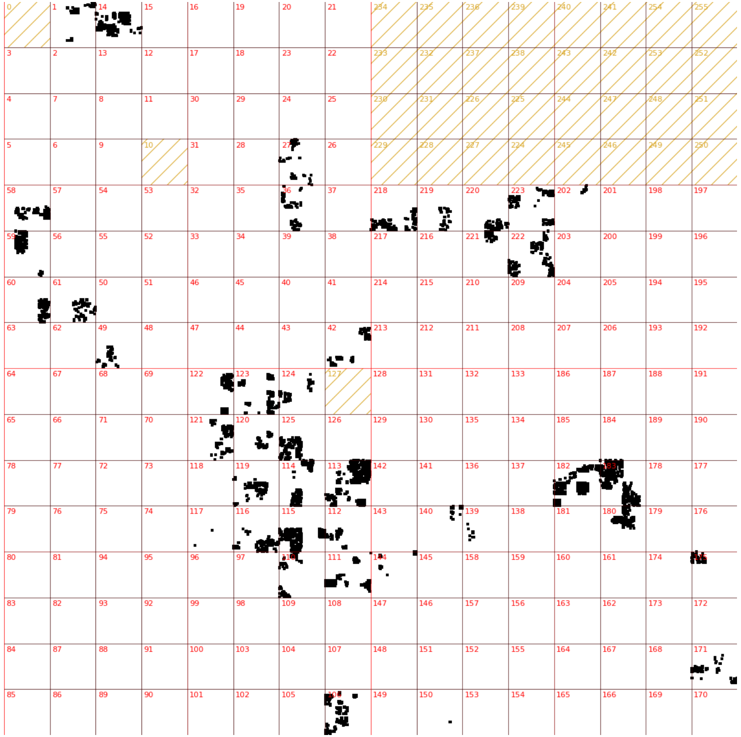

Dealing with Attacks

Null Routing

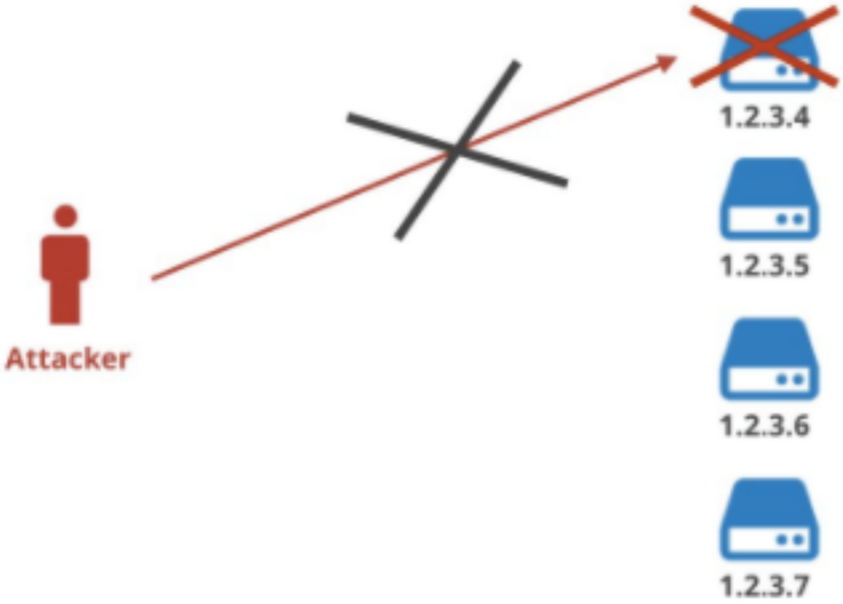

Null Routing

- Probably the simplest way to deal with an attack
- You instruct your ISP not to route traffic for a single host, or a series of hosts in your network
- Except, you’ve **just let the attacker win**
- If you null route your service, you’ve taken it offline. Perhaps you have an advanced system and can quickly renumber, but the attacker can update their attack too

The only way to stay online is to absorb the attack

Receive and Process

Receive and Process

- To absorb the attack you need to receive and process it
- This means you need to scale up infrastructure or develop advanced techniques to deal with attacks
- Both of these need huge amounts of capacity, both physical and logical
- Few networks are ready for it, so you outsource
- But this breaks the model of de-centralization

Centralization

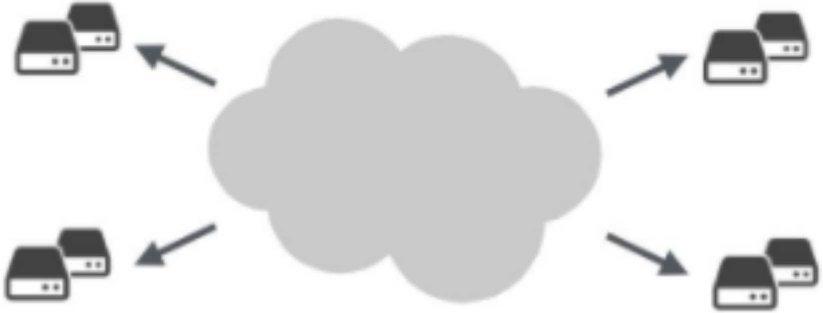

Solution?

Technical solutions to IP Spoofing have failed

Don’t just solve the IP Spoofing

Don’t just solve the IP Spoofing…

…solve the attribution!

NETFLOW
ALL THE THINGS

Netflow

- Opensource Toolsets are great
- Scales very well
- Privacy Concerns?
  - This is very very simple data
  - Rotate (delete) logs every few days
  - Use a high sampling rate. 1/16,000

Netflow

- H/W vendors must get better
- Netflow v9 supports src/dst MAC
    - Which vendor supports it?

Photo: The Simpsons/FOX

NetFlow

- It is **EMBARRASING** that a transit provider doesn't know where packets ingress their networks
- It's even more embarrassing that service providers who have NetFlow equipment, be it open sourced / in house or provided by a vendor don't know how to use it
- It's also **EMBARRASING** that hardware vendors don't support full NetFlow v9

- This needs to be resolved now

This is the first step

Attribution allows informed discussion

DDoS Causes centralization

CLOUDFLARE

To fix DDoS we need attribution

To make the internet better for everyone

CLOUDFLARE