DDoS Tutorial

Krassimir Tzvetanov
krassimir@a10networks.com

NANOG 65

Introduction and overview

Introduction

- Who am I?
- What is the target audience of this tutorial?
- Let's make it interactive!

Overview

- Discuss what DDoS is, general concepts, adversaries, etc.
- What is currently fashionable?
  - DDoS, NTP, SSDP
  - SYN Flood (Prince quote here)
- Go through a networking technology overview, in particular the OSI layers, sockets and their states
- Look at popular attack types at the different layers
- Discuss reflection and amplification
- Challenges
- Mitigations

What is DoS/DDoS?

What is Denial of Service?

- Resource exhaustion… which leads to lack of availability
- Consider:
  - How is it different from CNN pointing to somebody's web site?
  - How is that different from company's primary Internet connection going down?

What is Denial of Service?

- From security point of view?
  - Decreased availability
- From operations point of view?
  - An outage
- From business point of view?
  - Financial losses

DDoS is an Outage!

- Well, as service providers, we all know how to deal with outages

DoS vs. DDoS?

- One system is sending the traffic vs many systems are sending the traffic
- In the past it _usually_ meant difference in volume
  - Over the past 3 years, due to reflective attacks, this has been changing rapidly

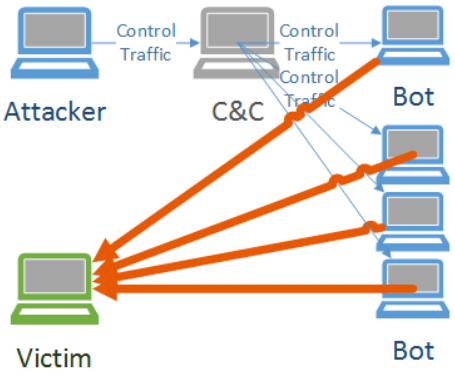

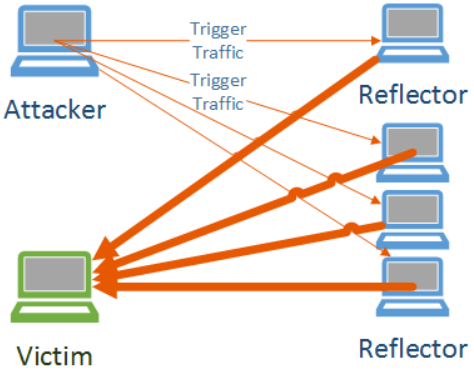

The problem?

Let's look at attack bandwidth

- Bandwidth in 2010 – little over 100 Gbps?
- 2013 – over 300 Gbps
- 2014 – over 400 GBps

Source: Arbor Networks Yearly Report

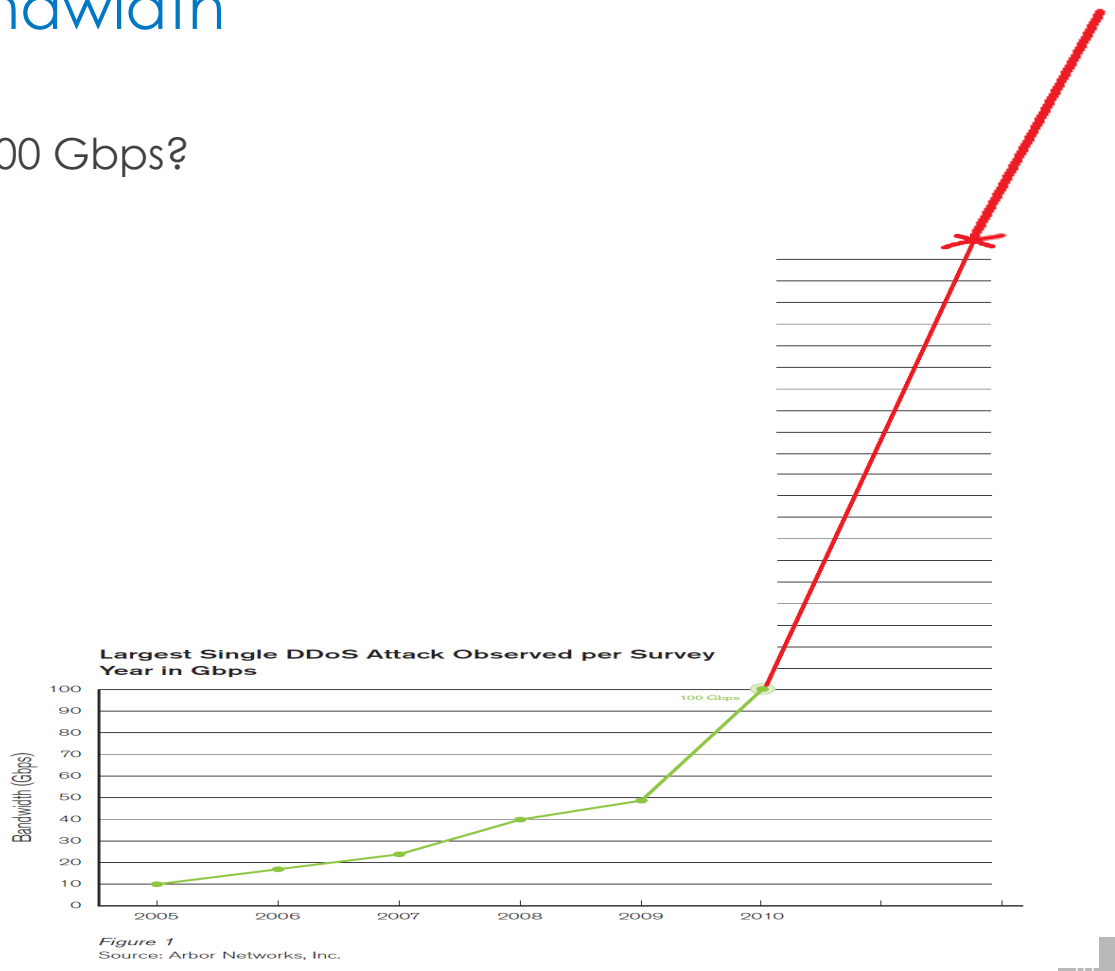

Figure 1
Source: Arbor Networks, Inc.

Contributing factors

- Embedded devices (mostly home routers)
- Available reflectors (DNS, NTP, SSDP)

  …with ability to amplify
- Outdated Content Management Systems (CMSes)

Who is the adversary?

Adversary

- Wide range of attackers
  - Gamers – on the rise!!! ☺
  - Professional DDoS operators and booters/stressors
  - Some of the attacks have been attributed to nation states
  - Hacktivists – not recently

  …and more

Motivation

- Wide range of motivating factors as well
  - Financial gain
    - extortion (DD4BC)
    - taking the competition offline during high-gain events
  - Political statement
  - Divert attention (seen in cases with data exfiltration)
  - Immature behavior

Skill level

- Wide range of skills
  - Depending on the role in the underground community
  - Mostly segmented between operators and tool-smiths
  - Tool-smiths are not that sophisticated (at this point) and there is a large reuse of code and services
  - This leads to clear signatures for some of the tools

- Increasing complexity
  - DirtJumper
  - xnote.1
  - XOR Botnet

What is new(-ish)?

What is new?

- Booters/Stressors
- Embedded home and SOHO devices
- Content management systems – in the past

Booters/Stressors

- Inexpensive
- Tools are sold for cheap on the black market (forums)
- Range 5-10 Gbps and more
- Usually short duration
- Popular among gamers

Booters/Stressors

- What are the booter services?
- A picture is worth a thousand words:
  - Think about the audience they are trying to attract
- Google: "Gwapo's Professional DDOS"

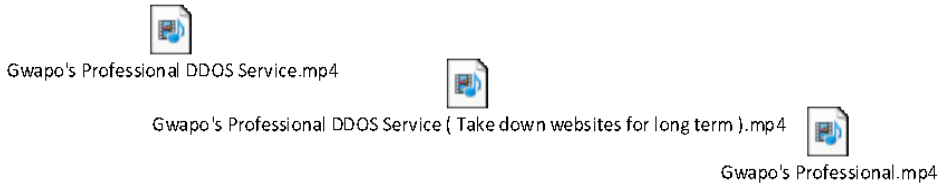

Home routers

- Embedded home and SOHO devices
  - Krebs on security: http://krebsonsecurity.com/2015/01/lizard-stresser-runs-on-hacked-home-routers/
- XBOX and Sony attacks over Christmas
  - Default username password
  - Open DNS recursive resolvers
  - NetUSB bug (from last week)
- Is that intentional? – "follow the money"

Technology and Terminology Overview

Technology Overview

- The purpose of this section is to level set
- Topics we'll cover
  - OSI and Internet models
  - TCP and sockets
  - Look at the operation of tools like netstat, netcat, tcpdump and wireshark
  - DNS operation and terminology
  - NTP, SNMP, SSDP operation
  - Some terminology and metrics
- Let me know if the pace is too slow or too fast

Attack types and terminology

Attack classification classifications (pun intended) ;)

- By volume
  - Volumetric
  - Logic/Application
- Symmetry
  - Asymmetric
  - Symmetric
- Direction
  - Direct
  - Reflected

- Source
  - Single source
  - Distributed
- State change
  - Permanent
  - Recoverable

- Based on network layer

Metrics

- Bandwidth (Kbps, Gbps)
- PPS
- QPS
- Storage
- CPU
- Application specific – usually latency

Attack surface

Network Layers – OSI vs Internet Model

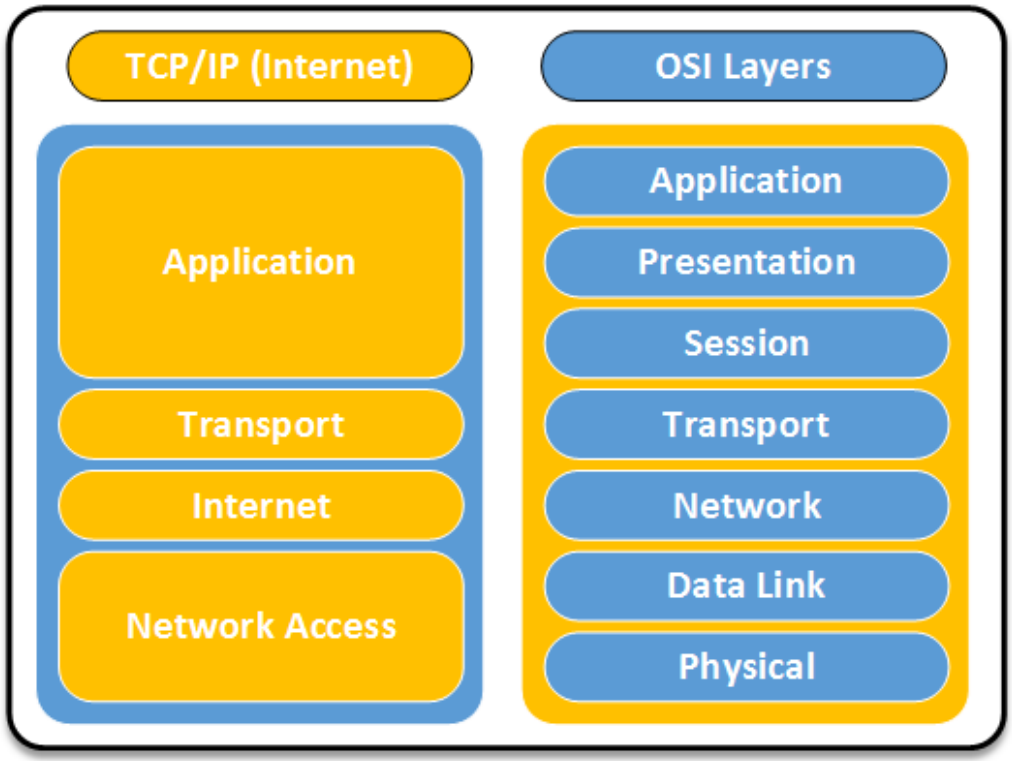

Physical and Data-link Layers

- Cut cables
- Jamming
- Power surge
- EMP

- MAC Spoofing
- MAC flood

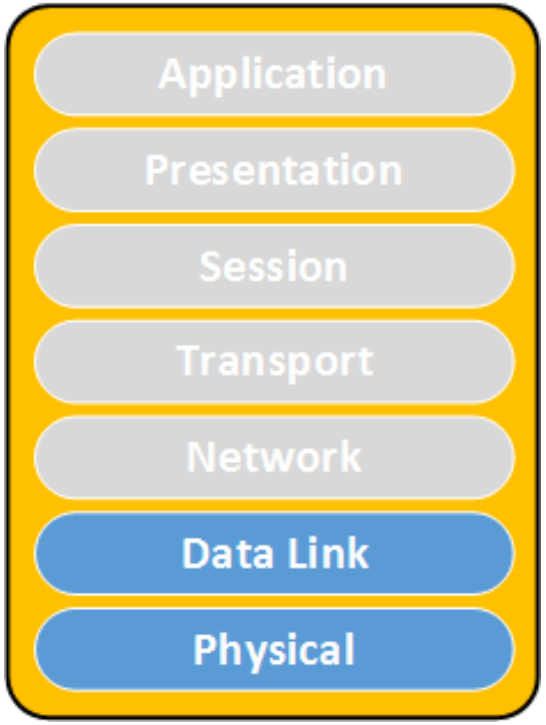

Network Layer

- Floods (ICMP)
- Teardrop
  (overlapping IP segments)

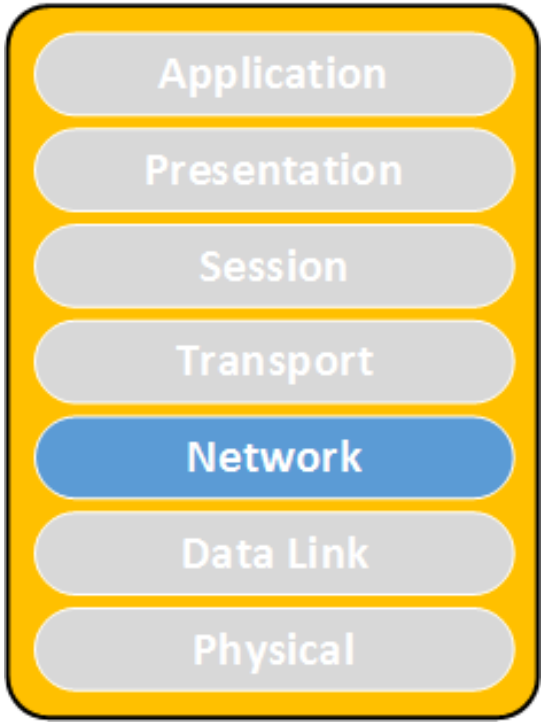

Transport Layer

- SYN Flood
- RST Flood
- FIN Flood
- You name it…

- Window size 0
  (looks like Slowloris)
- Connect attack
- LAND (same IP as src/dst)

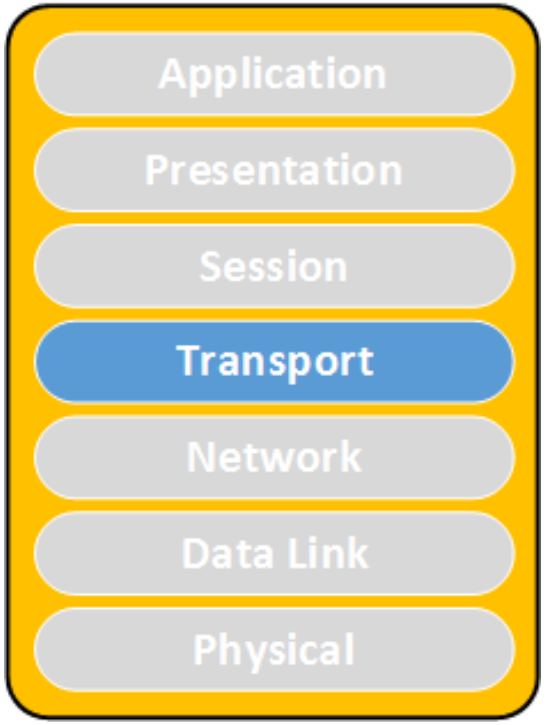

Session Layer

- Slowloris
- Sending data to a port with no NL in it (long headers, long request lines)
- Send data to the server with no CR

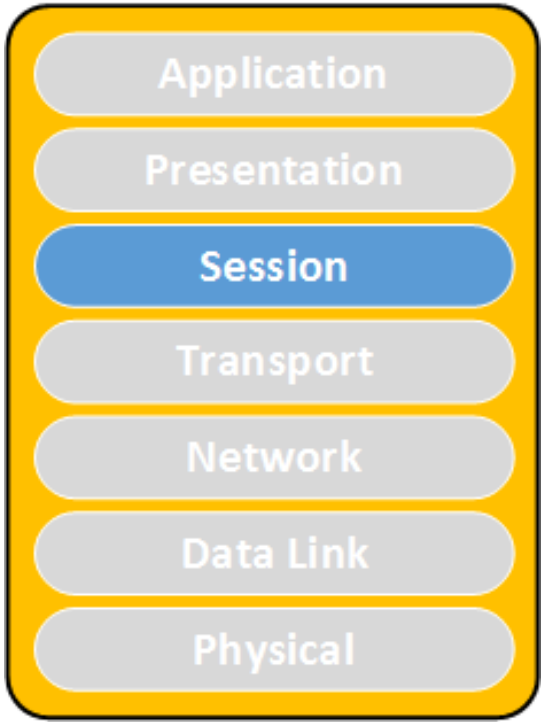

Presentation Layer

- Expensive queries (repeated many times)
- XML Attacks

```
<!DOCTYPE lolz
[
<!ENTITY lol1 "&lol2;">
<!ENTITY lol2 "&lol1;">
 ]>
<lolz>&lol1;</lolz>
```

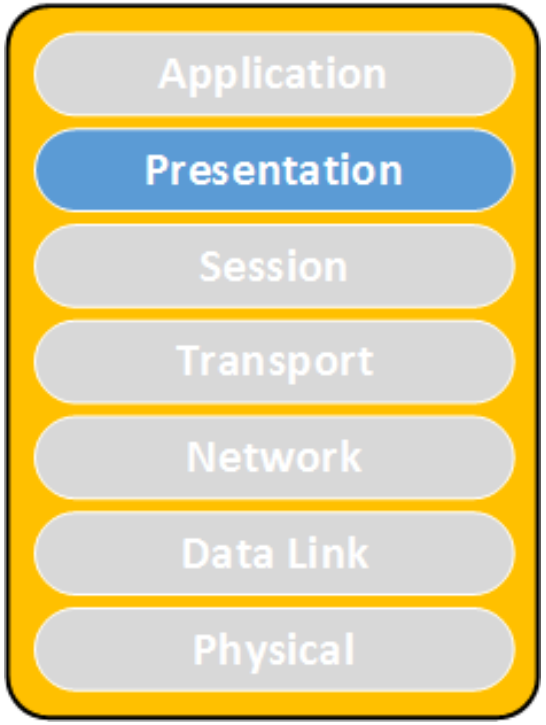

Application Layer

- SPAM?
- DNS queries
- Black fax

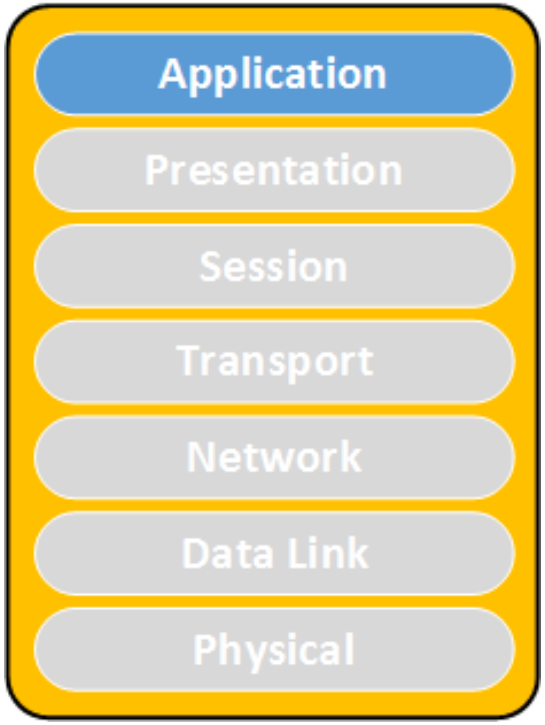

Attack summary by layer

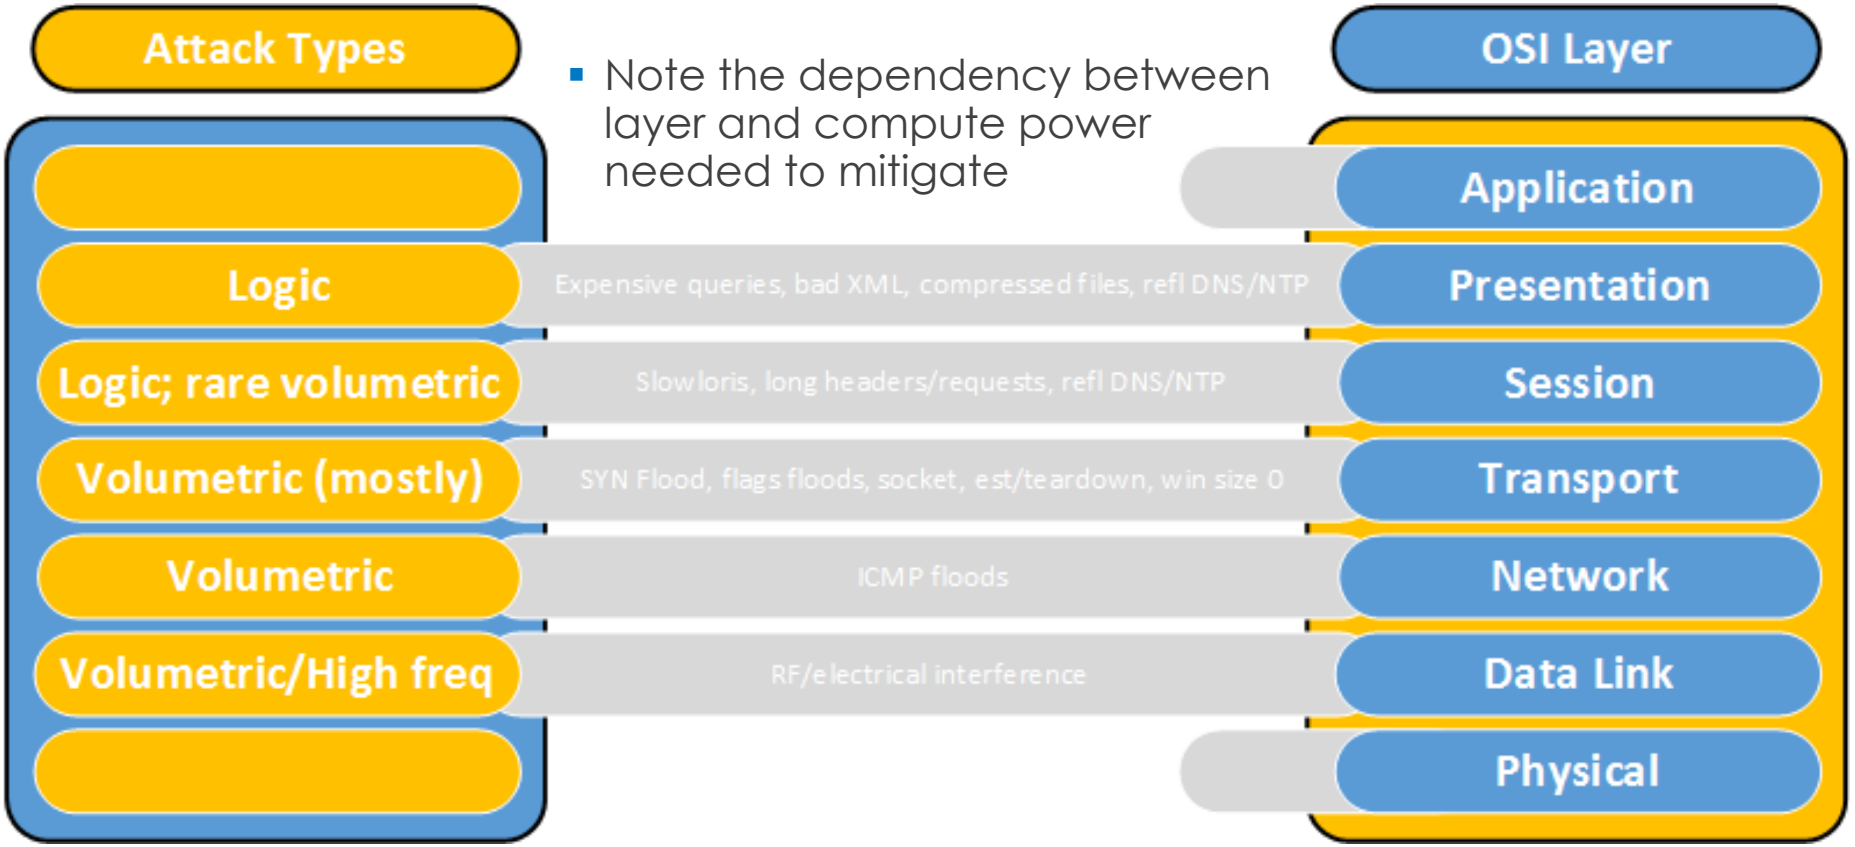

- Note the dependency between layer and compute power needed to mitigate

Attack types and terminology

Transmission Control Protocol (TCP)

Sockets

- Socket is an abstraction allowing an application to bind to a transport layer address (aka network port)
- It is described by a state machine
- Throughout its life time it goes through a number of states

Socket States

- Here are some of the socket states of importance:
  - LISTEN – waiting for a connection request
  - SYN_RECV – received request still negotiating
  - ESTABLISHED – connection working OK
  - FIN-WAIT1/2 – one side closed the connection
  - TIME-WAIT – waiting for a while…
    - What is MSL?

- In most of the states a socket is characterized by:
  - IP address
  - TCP/UDP address

Use of netstat for troubleshooting

```
[root@knight ghost]# netstat -nap | grep 12345
tcp        0      0 0.0.0.0:12345           0.0.0.0:*               LISTEN      2903/nc
[root@knight ghost]# netstat -nap | grep 12345
tcp        0      0 127.0.0.1:12345         127.0.0.1:49188         ESTABLISHED 2903/nc
[root@knight ghost]# netstat -nap | grep 12345
tcp        0      0 127.0.0.1:49188         127.0.0.1:12345         TIME_WAIT   -
[root@knight ghost]# netstat -nap | grep 12345
[root@knight ghost]#
```

SYN Flood

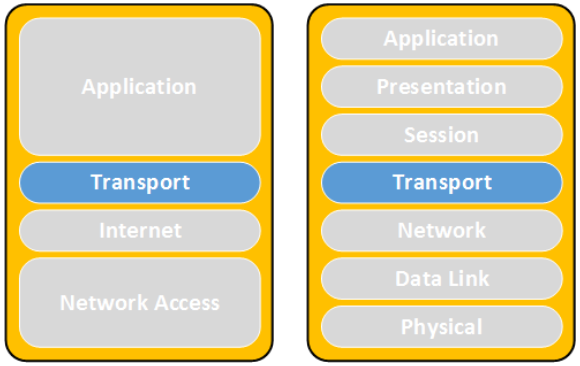
Application
Transport
Internet
Network Access
Application
Presentation
Session
Transport
Network
Data Link
Physical

What is a SYN flood?

- What is a 3-way handshake?

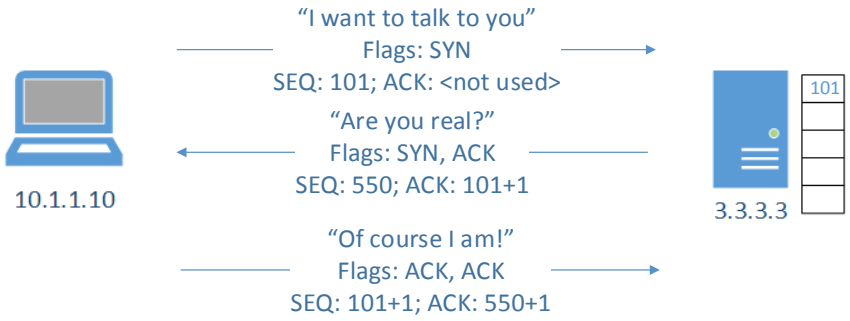

SYN flood

- Exploits the limited slots for pending connections
- Overloads them

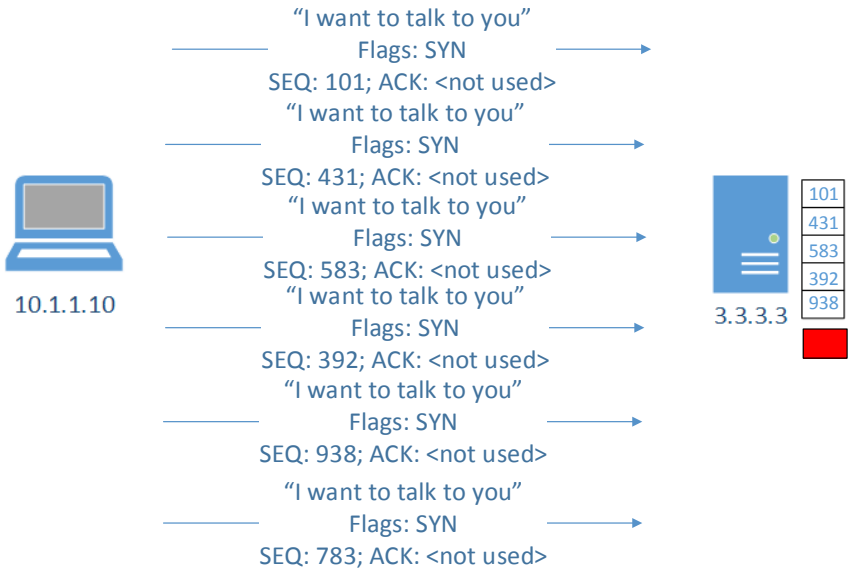

SYN flood through the eyes of netstat

- netstat –anp

```
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address           Foreign Address         State       PID/Program name
tcp        0      0 0.0.0.0:111             0.0.0.0:*               LISTEN      1339/rpcbind
tcp        0      0 0.0.0.0:33586           0.0.0.0:*               LISTEN      1395/rpc.statd
tcp        0      0 192.168.122.1:53        0.0.0.0:*               LISTEN      1962/dnsmasq
tcp        0      0 127.0.0.1:631           0.0.0.0:*               LISTEN      1586/cupsd
tcp        0      0 127.0.0.1:25            0.0.0.0:*               LISTEN      2703/sendmail: acce
tcp        0      0 127.0.0.1:25            127.0.0.1:49718         SYN_RECV    -
tcp        0      0 127.0.0.1:25            127.0.0.1:49717         SYN_RECV    -
tcp        0      0 127.0.0.1:25            127.0.0.1:49722         SYN_RECV    -
tcp        0      0 127.0.0.1:25            127.0.0.1:49720         SYN_RECV    -
tcp        0      0 127.0.0.1:25            127.0.0.1:49719         SYN_RECV    -
tcp        0      0 127.0.0.1:25            127.0.0.1:49721         SYN_RECV    -
tcp        0      0 127.0.0.1:25            127.0.0.1:49716         SYN_RECV    -
```

SYN flood mitigation

- Technology
  - SYN Cookies
  - Whitelists
  - TCP Proxy (TCP Intercept – active mode)
  - TCP Resets (TCP Intercept – passive)
  - Nowadays – volumetric
- Device stack optimization
- Dedicated devices

What is a SYN cookie?

- Hiding information in ISN (initial seq no)
- SYN Cookie:

  Timestamp % 32 + MSS + 24-bit hash
- Components of 24-bit hash:
  - server IP address
  - server port number
  - client IP address
  - client port
  - timestamp >> 6 (64 sec resolution)

Enabling SYN-coockie

- To enable SYN cookies:

  echo 1 > /proc/sys/net/ipv4/tcp_syncookies

- All TCP related settings are located in /proc/sys/net/ipv4/
  - tcp_max_syn_backlog
  - tcp_synack_retries
  - tcp_syn_retries

Socket Exhaustion

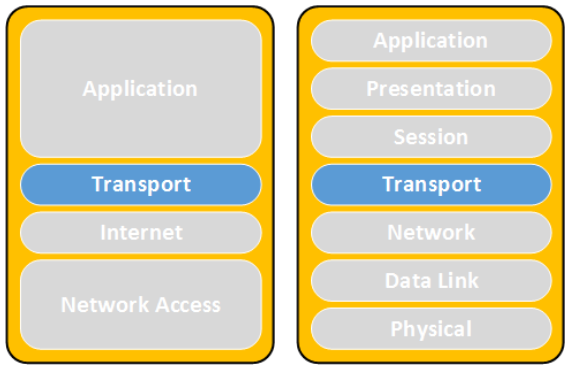

Socket Exhaustion

- What is a socket?
- What is Maximum Segment Lifetime (MSL)?
  - How old is the Internet?
  - What is Time To Live (TTL) measured in?
- What is socket exhaustion?

Socket Exhaustion through the eyes of netstat

- Socket exhaustion would look likethis:

Active Internet connections (servers and established)

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State	PID/Program name
tcp	0	0	0.0.0.0:111	0.0.0.0:*	LISTEN	1339/rpcbind
tcp	0	0	0.0.0.0:33586	0.0.0.0:*	LISTEN	1395/rpc.statd
tcp	0	0	192.168.122.1:53	0.0.0.0:*	LISTEN	1962/dnsmasq
tcp	0	0	127.0.0.1:631	0.0.0.0:*	LISTEN	1586/cupsd
tcp	0	0	127.0.0.1:25	0.0.0.0:*	LISTEN	2703/sendmail: acce
tcp	0	0	0.0.0.0:1241	0.0.0.0:*	LISTEN	1851/nessusd: waiti
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60365**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60240**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60861**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60483**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60265**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60618**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60407**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60423**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60211**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60467**	**TIME_WAIT**	**-**
tcp	**0**	**0**	**127.0.0.1:25**	**127.0.0.1:60213**	**TIME_WAIT**	**-**

How to enable socket reuse

- Enable socket reuse

  echo 1 > /proc/sys/net/ipv4/tcp_tw_recycle

  echo 1 > /proc/sys/net/ipv4/tcp_tw_reuse

Slowloris

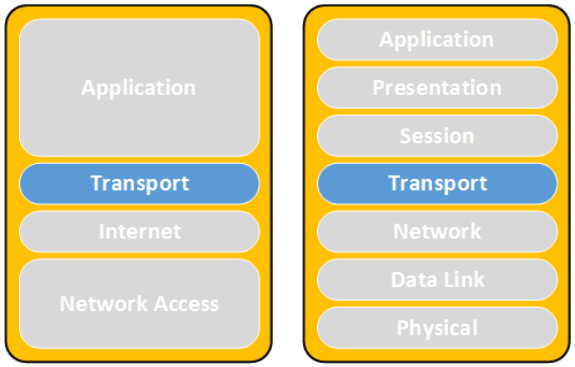
Application
Transport
Internet
Network Access
Application
Presentation
Session
Transport
Network
Data Link
Physical

Connection handling architectures

- Process based connection handling?
  - Think “Apache”

- Event based connection handling?
  - Think “nginx”

Slowloris

- Exploits the process based model but opening a number of concurrent connections and holds them open for as long as possible with the least amount of bandwidth possible

Slowloris mitigation

- Change of the software architecture
- Use of event driven reverse proxy to protect the server (like nginx)
- Dedicated hardware devices

Reflection and amplification attacks

Two different terms

- Reflection – using an intermediary to deliver the attack traffic

- Amplification – ability to deliver larger response than the trigger traffic

Reflection

Reflective attacks

- Attacks where the an unwilling intermediary is used to deliver the attack traffic
- The attacker would normally send a packet with a forged source IP address to the intermediary. The forget address is going to be the one of the target. The intermediary will deliver a response which will go to the target instead of the attacker
- Note to audience: think what protocols we can use for that?

What is reflection(ed) attack

- Attacks where the an unwilling intermediary is used to deliver the attack traffic

- Attacker sends a packet with a spoofed source IP set to the victim's
- Reflectors respond to the victim

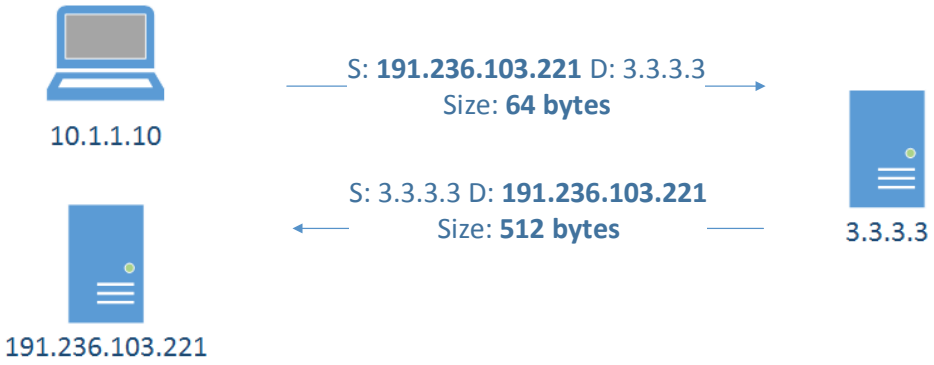

Reflector types

The ones that are of interest are:

- DNS
- NTP
- SSDP
- SNMP
- RPC (reported lately but not really large)

Amplification

What is amplification attack?

- Asymmetric attack where response is much larger than the original query

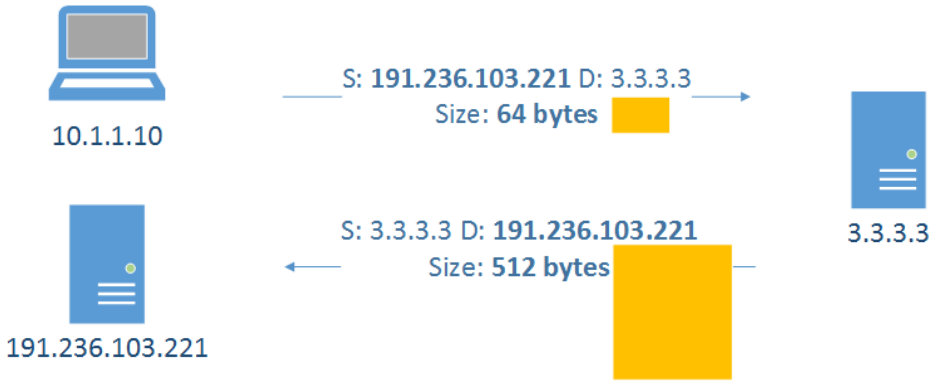

Amplifiers types

- The ones that are of interest and provide amplifications are:
  - DNS
  - SSDP
  - NTP
  - SNMP

- Amplification factors:
  https://www.us-cert.gov/ncas/alerts/TA14-017A

Amplification quotients

Protocol	Bandwidth Amplification Factor	Vulnerable Command
DNS	28 to 54	Multiple
NTP	556.9	Multiple
SNMPv2	6.3	GetBulk request
NetBIOS	3.8	Name resolution
SSDP	30.8	SEARCH request
CharGEN	358.8	Character generation request
QOTD	140.3	Quote request
BitTorrent	3.8	File search
Kad	16.3	Peer list exchange
Quake Network Protocol	63.9	Server info exchange
Steam Protocol	5.5	Server info exchange

- Source: US-CERT: https://www.us-cert.gov/ncas/alerts/TA14-017A

DNS Resolution

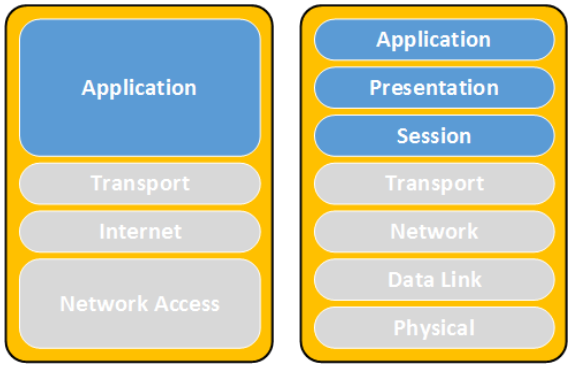

DNS resolution

- Authoritative
- Recursive

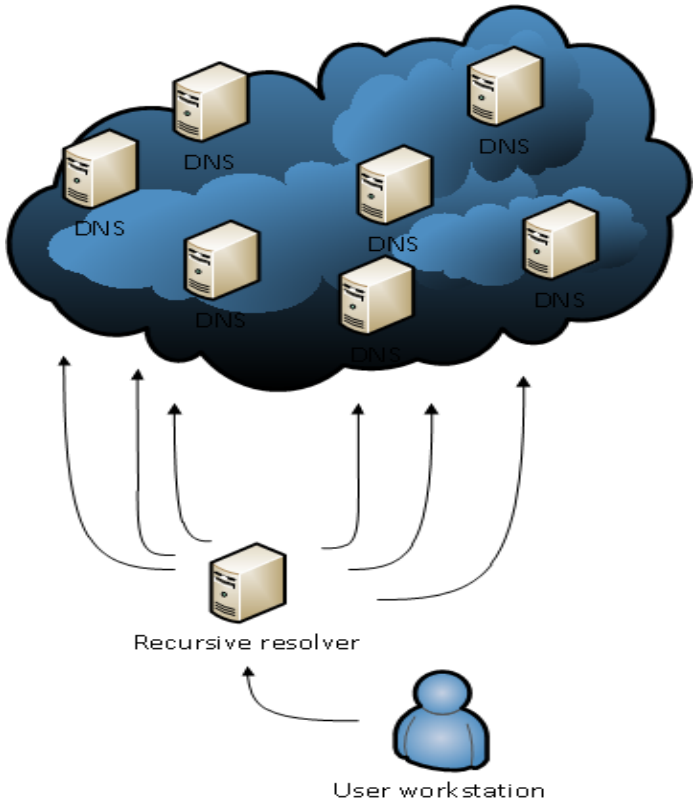

DNS resolution

- How does DNS work?
- User talks to recursive resolver
- The recursive goes on the Internet and talks to the authoritative servers
- When an answer is obtained (or not) it reports back to the user

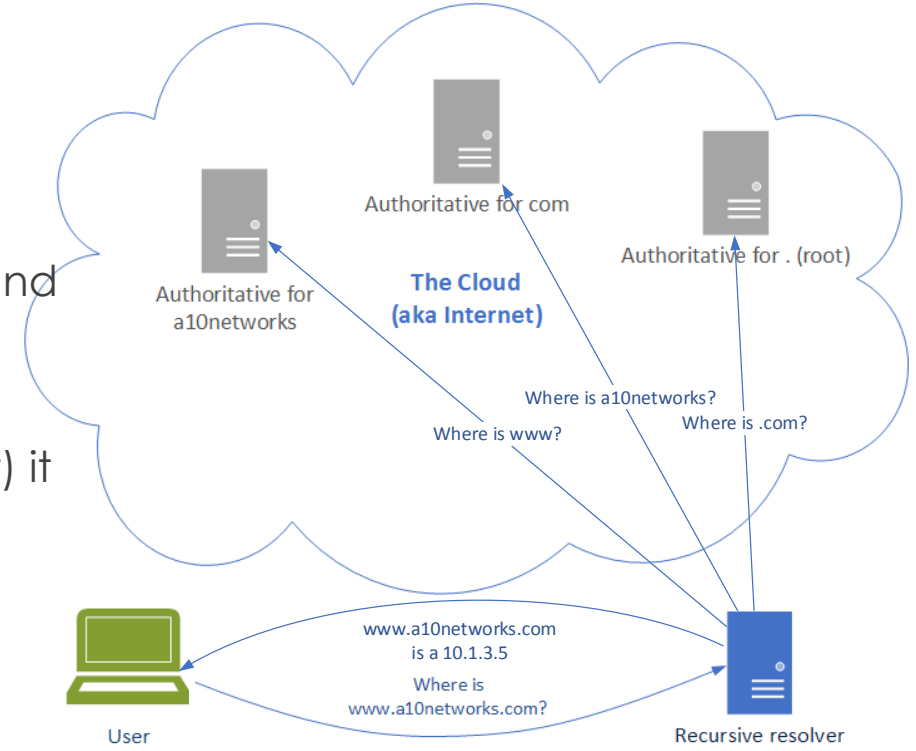

DNS resolution at the packet level

- The process of mapping:
  www.a10networks. com => 191.236.103.221

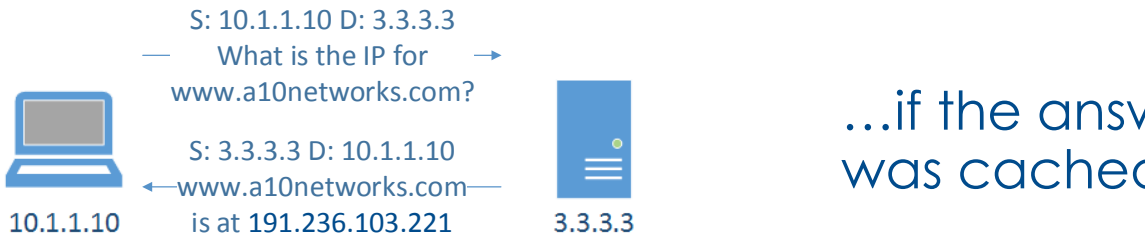

…if the answer was cached

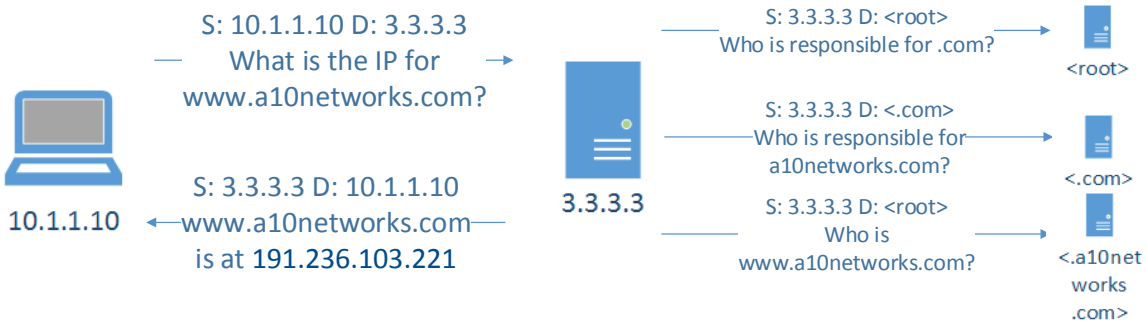

DNS Reflection

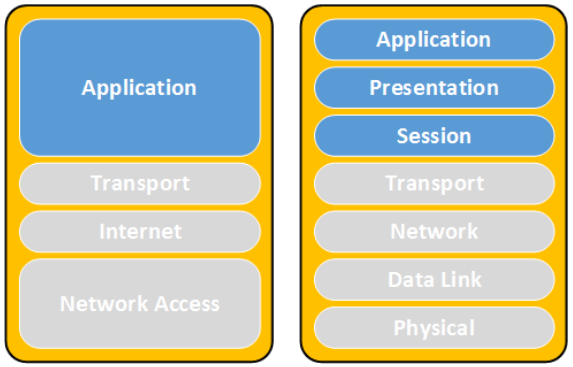

What is DNS reflection attack?

- What happens if an attacker forges the victim address as its source?

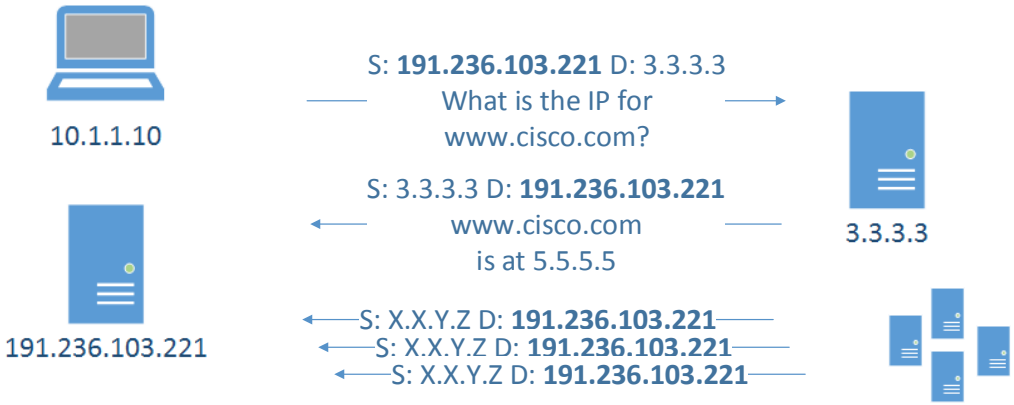

…the reflected traffic goes to the target server

- … and what if hundreds of misconfigured open DNS resolvers are used?

Consider this query

- Triggered by something like:
- dig ANY isc.org @3.3.3.3

- Example:~$ dig ANY isc.org @172.20.1.1 # My home lab
- Flip over for answer

Consider this (cont'd)

ghostwood@sgw:~$ dig ANY isc.org @172.20.1.1

;; ANSWER SECTION:

isc.org. 481 IN RRSIG DS 7 2 86400 20130607155725 20130517145725 42353 org. KHMs09DaFMx416/7xXhaD9By0NrqCiQ4kBnqi6oq2VocZRREAbUHHrAY KydlgKO5vOaw6l1Fy86/oiODkk3yyHspciwdJvjlefu4PktdUnd1IQxW 791q/jWgHBL5iQQigBYv7Z5lfY1ENn+6fPOchAywWqEBYcdqW8pzzOjz zlU=

isc.org. 481 IN DS 12892 5 2 F1E184C0E1D615D20EB3C223ACED3B03C773DD952D5F0EB5C777586D E18DA6B5

isc.org. 481 IN DS 12892 5 1 982113D08B4C6A1D9F6AEE1E2237AEF69F3F9759

isc.org. 5725 IN RRSIG A 5 2 7200 20130620134150 20130521134150 50012 isc.org. iCBy1Jj9P6mXVYjaSc62JClrZW+hvYAUGHo7WwRmxGRaipS8l9+LCvRl 2erglomkBP79m9ahnFOxWEAaueA6TlHClGxOkgrk3hBtMFjUB9rhvklm uxO2D8gc1DJDLl5egfpJCF2flTFhEvWzeMt6QGNwicWMxBsFHCxM7Fms D8l=

isc.org. 5725 IN A 149.20.64.42

isc.org. 5725 IN RRSIG DNSKEY 5 2 7200 20130620130130 20130521130130 12892 isc.org. dfxTGA/f6vdhulqojp+Konkdt8c4y3WiU+Vs5TjznvhdEyH14qPh/cHh +y1vA6+gAwTHl4X+GpzctNxiElwaSwVu3m9Nocniwl/AZQoL/SyDgEsI bJM/X+ZXY5qrgQrV2grOcKAAA91Bus3behYQZTsdaH2TStAKjKINEgvm yQ5xWEo6zE3p0ygtPq4eMNO4fRT9UQDhTRD3v3ztxFlNXKvBsQWZGBH0 5tQcbC6xnGyn1bBptJEEGhCBG01ncJt1MCyEf98VGHKJFeowORiirDQ3 cjJRFPTCCkA8n4j8vnsimIUP/TGl+Mg4ufAZpE96jJnvFBsdcC/iOo6i XkQVIA==

isc.org. 5725 IN RRSIG DNSKEY 5 2 7200 20130620130130 20130521130130 50012 isc.org. o18F3KIFkYedFRw1e5MP4qDo3wSg0XK9l5WCYD75aGhs9Rl5eyc/ 6KEW Se4lZXRhf6d77xXlerMYCrsfh/GHdjPRoE1xL/nzH/hTBJAI9XDbC5I/ EUpFlGVLVdQy43XKtywm0j2nyc5MdGa2VeLKo+hHTmH3St3pGRVJp2IK 5Z0=

isc.org. 5725 IN DNSKEY 257 3 5 BEAAAAOhHQDBrhQbtphgq2wQUpEQ5t4DtUHxoMVFu2hWLDMvoOMRXjGr hhCeFvAZih7yJHf8ZGfW6hd38hXG/ xylYCO6Krpbdojwx8YMXLA5/kA+ u50WIL8ZR1R6KTbsYVMf/Qx5RiNbPClw+vT+U8eXEJmO20jIS1ULgqy3 47cBB1zMnnz/4LJpA0da9CbKj3A254T515sNIMcwsB8/2+2E63/zZrQz Bkj0BrN/9Bexjpiks3jRhZatEsXn3dTy47R09Uix5WcJt+xzqZ7+ysyL KOOedS39Z7SDmsn2eA0FKtQpwA6LXeG2w+jxmw3oA8lVUgEf/rzeC/bB yBNsO70aEFTd

isc.org. 5725 IN DNSKEY 256 3 5 BQEAAAABwuHz9Cem0BJ0JQTO7C/a3McR6hMaufljs1dfG/inaJpYv7vH XTrAOm/MeKp+/ x6eT4QLru0KoZkvZJnqTI8JyaFTw2OM/ltBfh/hL2lm Cft2O7n3MfeqYtvjPnY7dWghYW4sVfH7VVEGm958o9nfi79532Qeklxh x8pXWdeAaRU=

a.root-servers.net. 297269 IN A 198.41.0.4

a.root-servers.net. 415890 IN AAAA 2001:503:ba3e::2:30

b.root-servers.net. 298007 IN A 192.228.79.201

Reflection and Amplification

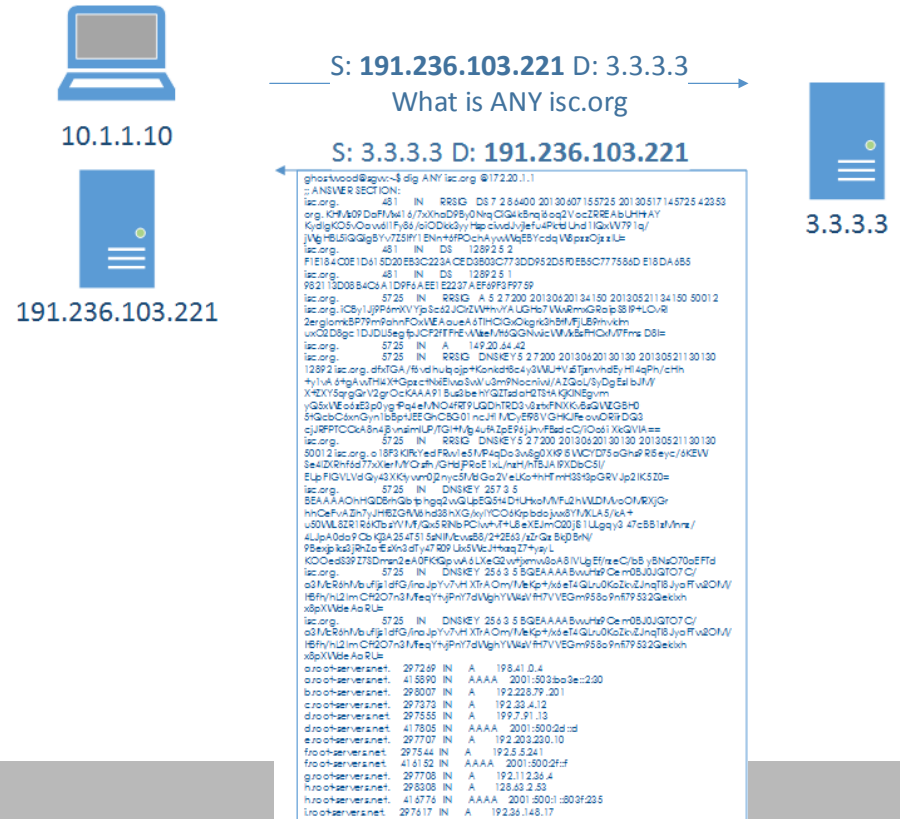

DNS attacks mitigation (victim)

- Validate packet and query structure
- Whitelisting
- Challenges*
- High performance equipment
  - Variety of techniques
  - Vendor dependent
- Drop known reflector traffic: http://openresolverproject.org/

DNS attacks mitigation (victim - DNS challenge)

- What is a DNS challenge?

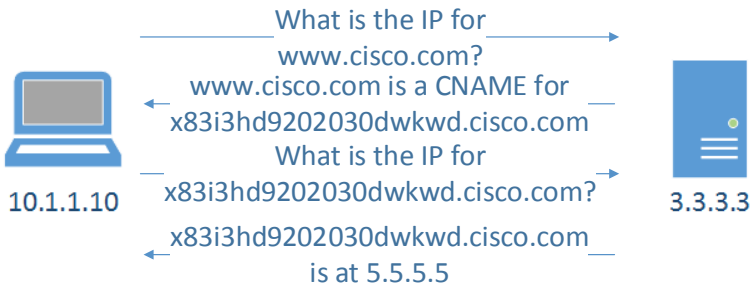

- Challenges with DNS challenge?
  - Two times the amount of traffic
  - Two times the packet rate
  - Computational resources

Large scale mitigation and load distribution: Anycast

- Multiple points of presence advertise the same address space
- Network ensures user is routed to the "closest" instance

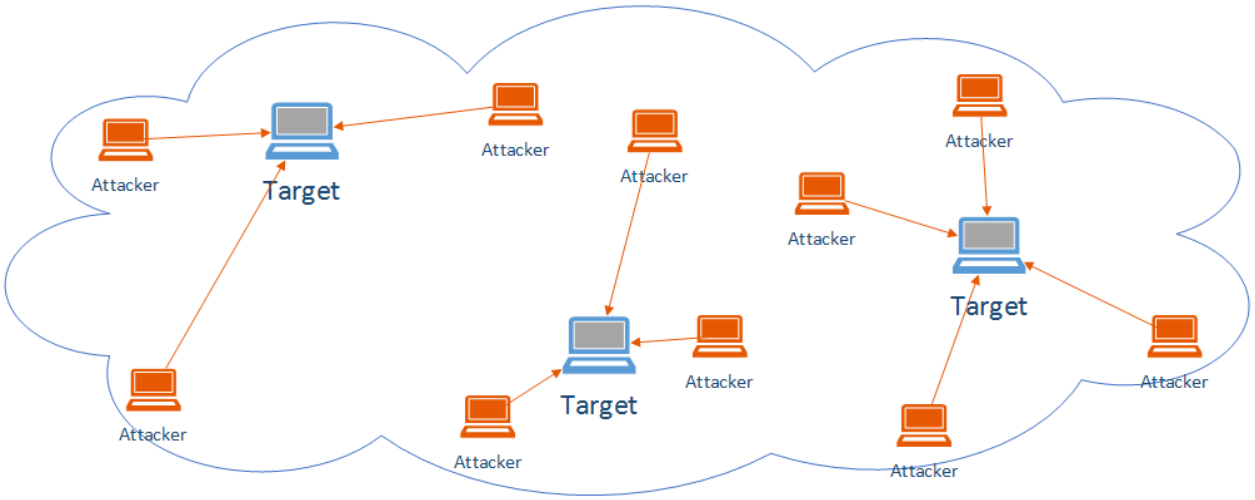

DNS Rate limits (reflector)

- Not specified for recursive but you can still tweak it to something that works for you
- Configuration example:

```
rate-limit {
        responses-per-second 5;
        window 5;
};
```

- Reference:

http://www.redbarn.org/dns/ratelimits

Proper resolver configuration (reflector)

```
acl "trusted" {
        192.168.0.0/16;
        10.153.154.0/24;
        localhost;
        localnets;
};

options {
        ...
        allow-query { trusted; }; // allow-query { any; };
        allow-recursion { trusted; };
        allow-query-cache { trusted; };
        ...
};
```

Large scale mitigation and load distribution: Anycast

- Multiple points of presence advertise the same address space
- Network ensures user is routed to the "closest" instance

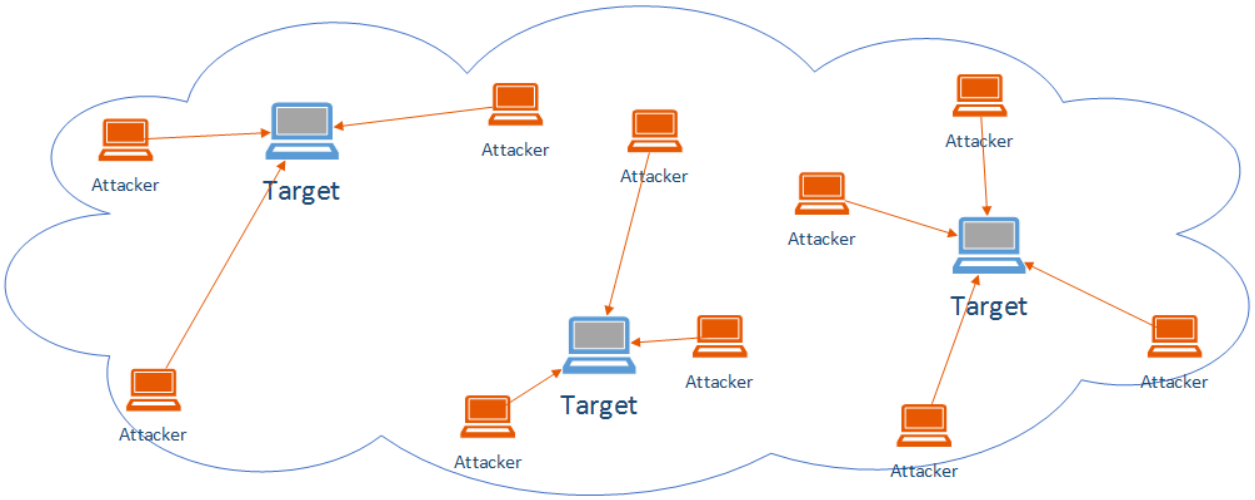

IPS/DDoS mitigation gear

- Depends on vendor
- Different techniques
- Different mitigation rates for different packet types

Network Time Protocol (NTP)

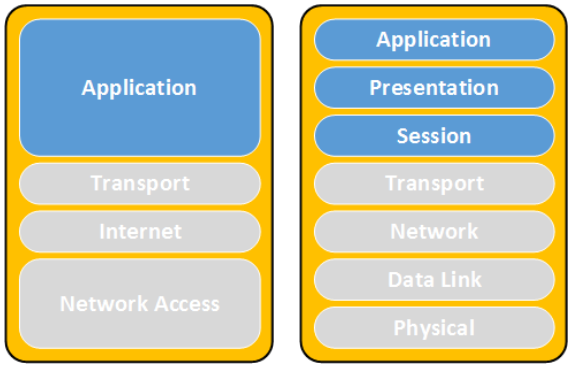

NTP servers

- Stratum servers
- NTP queries

- MONLIST command
  - provides
    a list of clients that have
    time readings

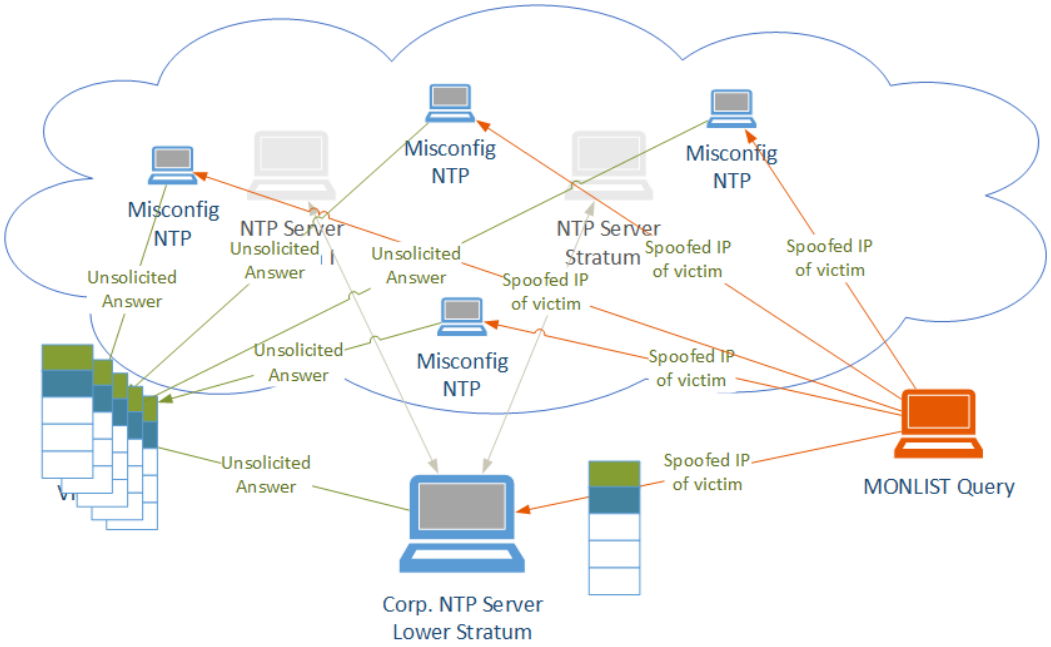

NTP server configuration

- Access lists
- NTP authentication
- Disable the MONLIST command
- Useful hints:
  http://www.team-cymru.org/secure-ntp-template.html
- List of open NTP reflectors:
  http://openntpproject.org/

Cache busting (back to DNS)

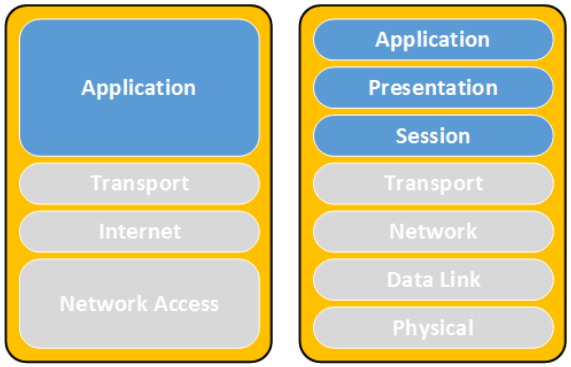

DNS resolution (rehash)

- Let's focus on the number of requests per second
- User talks to recursive resolver, which:
  - Caches answers
  - Answers a large number of requests
- The recursive talks to different level of authoritative servers, which:
  - Do not cache answers (they are auths)
  - Relatively lower number of queries
- Consider caching and authoritative capacity

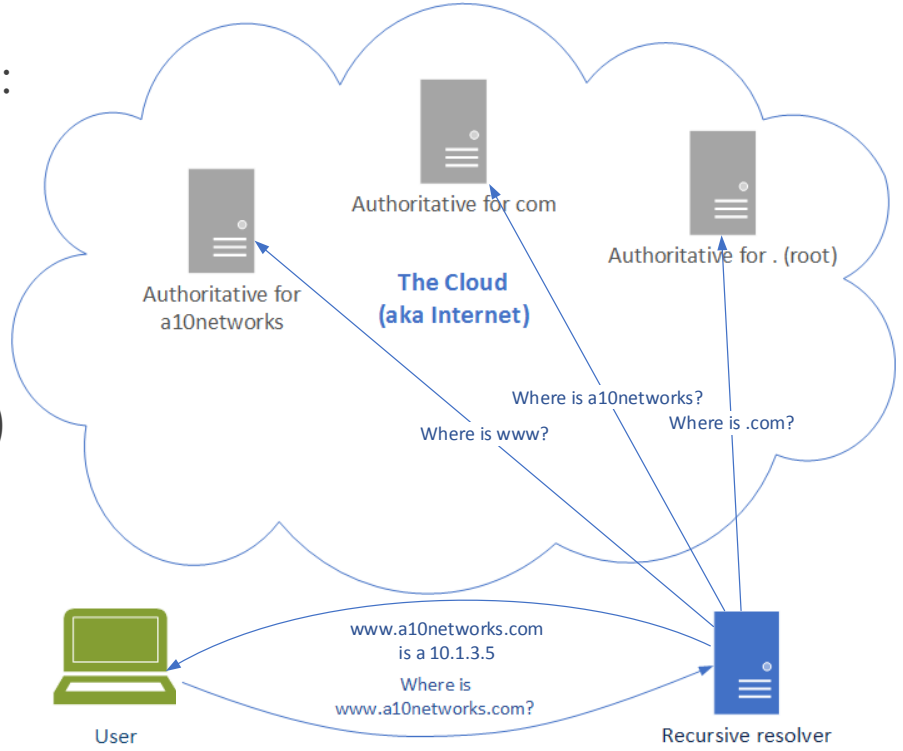

What cache busting?

- Attacker sends a query to recursive/reflector
- Recursive forwards the query
- And so on…
- Imagine one more recursive resolver
- Rinse and repeat…

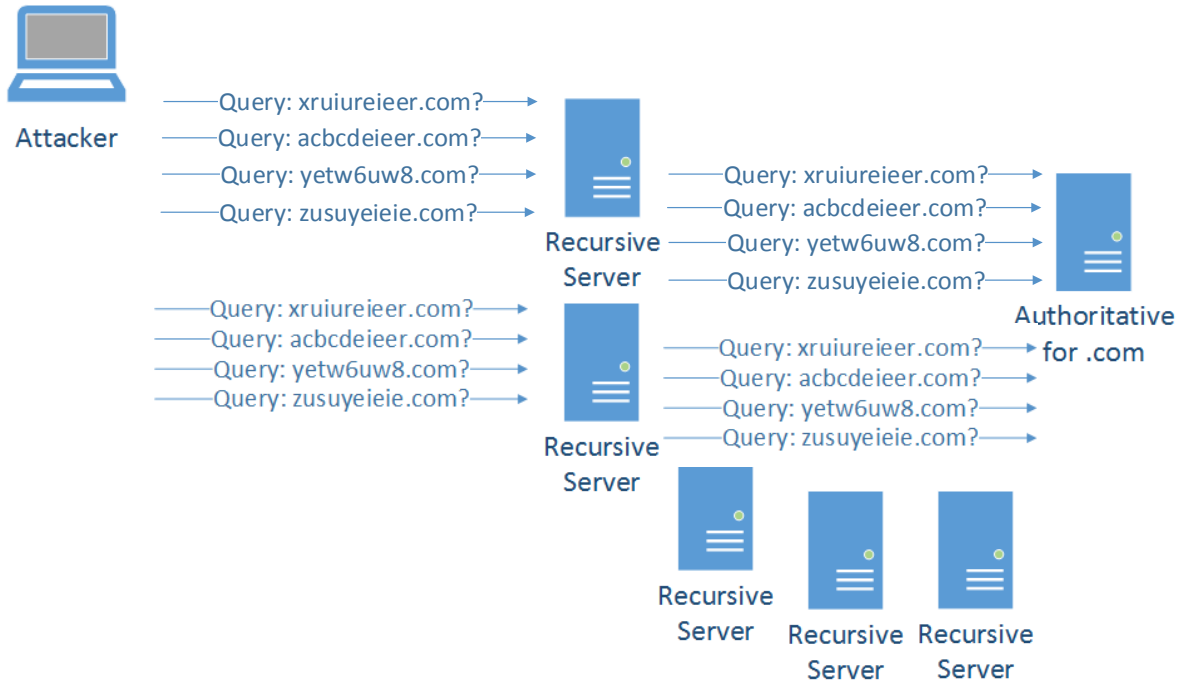

Backscatter

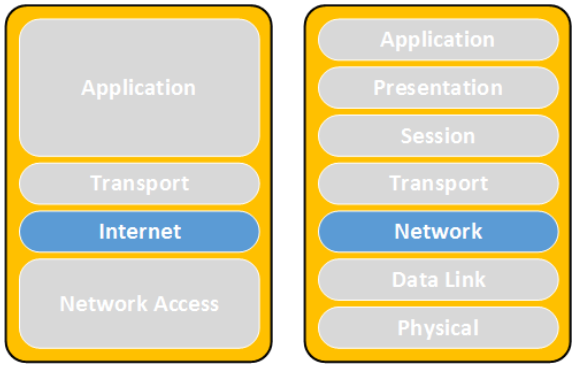

Backscatter

- Traffic that is a byproduct of the attack

- Why is that interesting?
  - It is important to distinguish between the actual attack traffic and unintended traffic sent by the victim
  - Imagine a SYN flood against a "victim" protected by a major scrubbing provider spoofed from IP address X
    - What is the traffic to X going to look like?

SYN Flood Backscatter?

- Cookie flood ☺

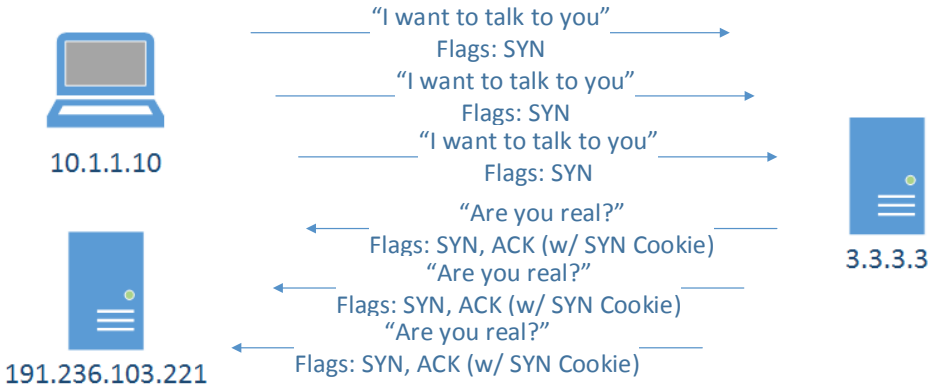

Are you a reflector? (Backscatter)

- In some cases return traffic/backscatter

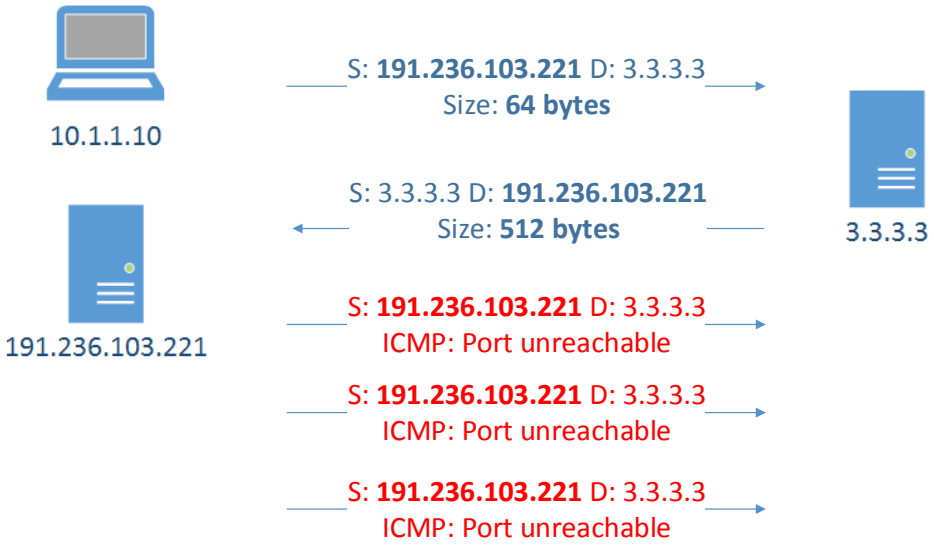

Mitigation

Risk Pyramid

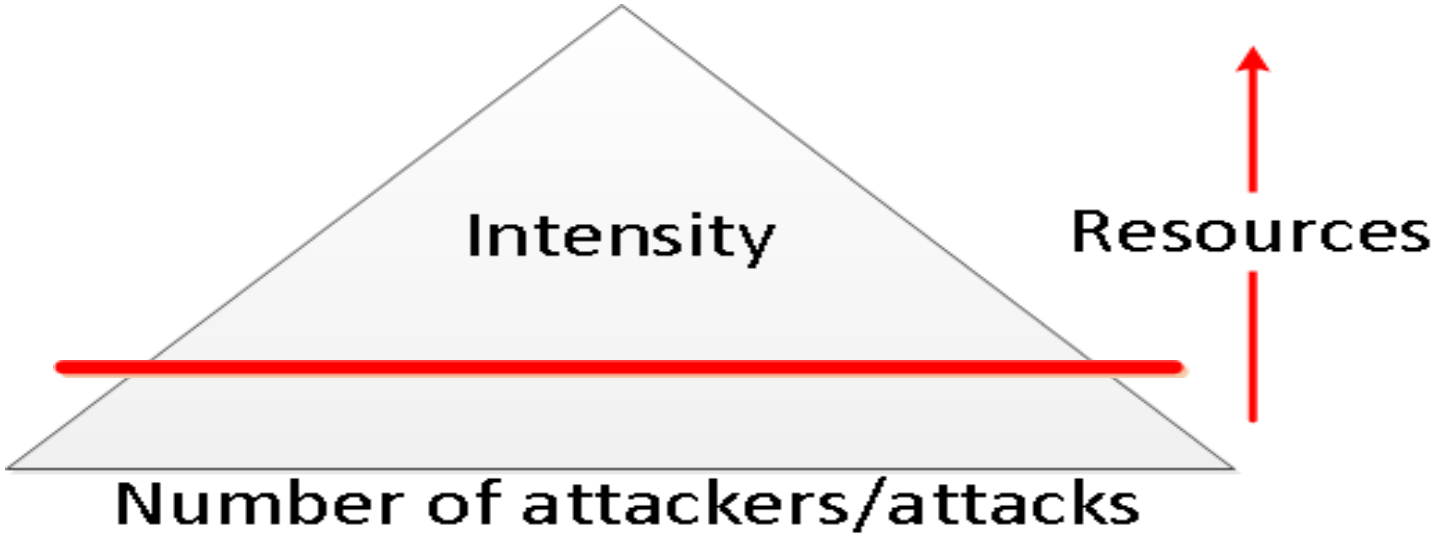

The cost of a minute?

- How much does a minute of outage cost to your business?
- Are there other costs associated with it? Reputation?
- Are you in a risk category?
- How much is executive management willing to spend to stay up?

- Are there reasons you need to mitigate on-site vs offsite? Latency?

On-site / DIY

- Bandwidth
- Equipment
- Qualified personnel
- More expensive overall but cheaper per MB
- Need for a backup plan

Outsource / scrubbing center

- Limited protocol support (usually HTTP/S)
- Added latency
- May loose visibility to source IP of the client
- Pay per MB of clean traffic (usually)
- Fast setup/Lower overhead
- More expensive per MB

Good Internet citizenship

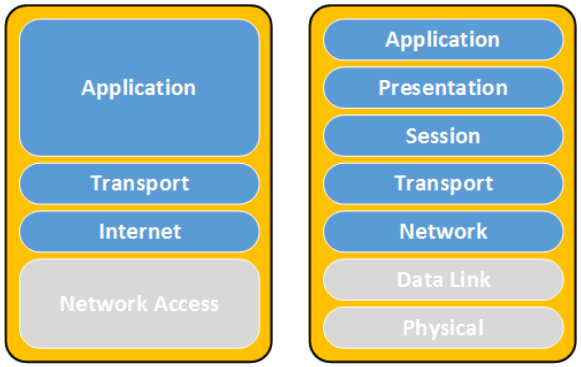

Mitigations

- Defend yourself
  - Anycast
  - Some form of IPS/DDoS mitigation gear
  - Overall network architecture
- Defend the Internet
  - Rate-limiting
  - BCP38/140 (outbound filtering) source address validation
  - Securely configured DNS, NTP and SNMP servers
  - No open resolvers
- Talk to the professionals

Are you noticing the imbalance?

Defend yourself

- Anycast (DNS)
- Some form of IPS/DDoS mitigation gear

- **Lots of money**

Defend the Internet

- Rate-limiting
- BCP38/140 (outbound filtering) source address validation
- Securely configured authoritative DNS servers
- No open resolvers

- **Somewhat cheap**

What's the point I'm trying to make?

- It's not feasible to mitigate those attacks single handedly
- We need cooperation
- Companies need to start including "defending the Internet from themselves" as a part of their budget – not only "defending themselves from the Internet"

What can I do about it?

- RFC 2827/BCP 38 – Paul Ferguson
- If possible filter all outgoing traffic and use proxy
- uRPF

- BCP 140: "Preventing Use of Recursive Nameservers in Reflector Attacks"
- http://tools.ietf.org/html/bcp140
- Aka RFC 5358

Resources

- DNS
- http://openresolverproject.org/

- NTP
- http://openntpproject.org/

- If you see your IP space in the lists provided by those sites – resolve it

Summary

- Discuss what DDoS is, general concepts, adversaries, etc.
- Went through a networking technology overview, in particular the OSI layers, sockets and their states, tools to inquire system state or capture and review network traffic
- Dove into specifics what attack surface the different layers offer
- Discussed different attack types
- Terminology
- Tools

Thank you