DDoS Mitigation

Krassimir Tzvetanov

NANOG 64

Introduction and overview

Introduction

- Who am I?
- What is the target audience of this tutorial?

Overview

- Discuss what DDoS is, general concepts, adversaries, etc.
- Go 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
- Dive into specifics what attack surface the different layers offer
- Discuss reflection, amplification and back scatter
- Terminology
- Tools

What is DoS?

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?

- Conclusion: It is a condition which leads to lack of availability of a resource

What is Denial of Service?

- The main point:

DoS is an Outage!

DoS vs. DDoS?

- One system is sending the traffic vs many systems are sending the traffic
- TODO: Elaborate on the differences

Common misconceptions

- You do not need a botnet?
- It’s not a matter of application or devices weaknesses but rather capacity

The problem?

Let's look at attack bandwidth

- Bandwidth in 2010 – little over 100 Gbps?
- Last year – over 300Gbps
- This year?
- Over 400 Gbps

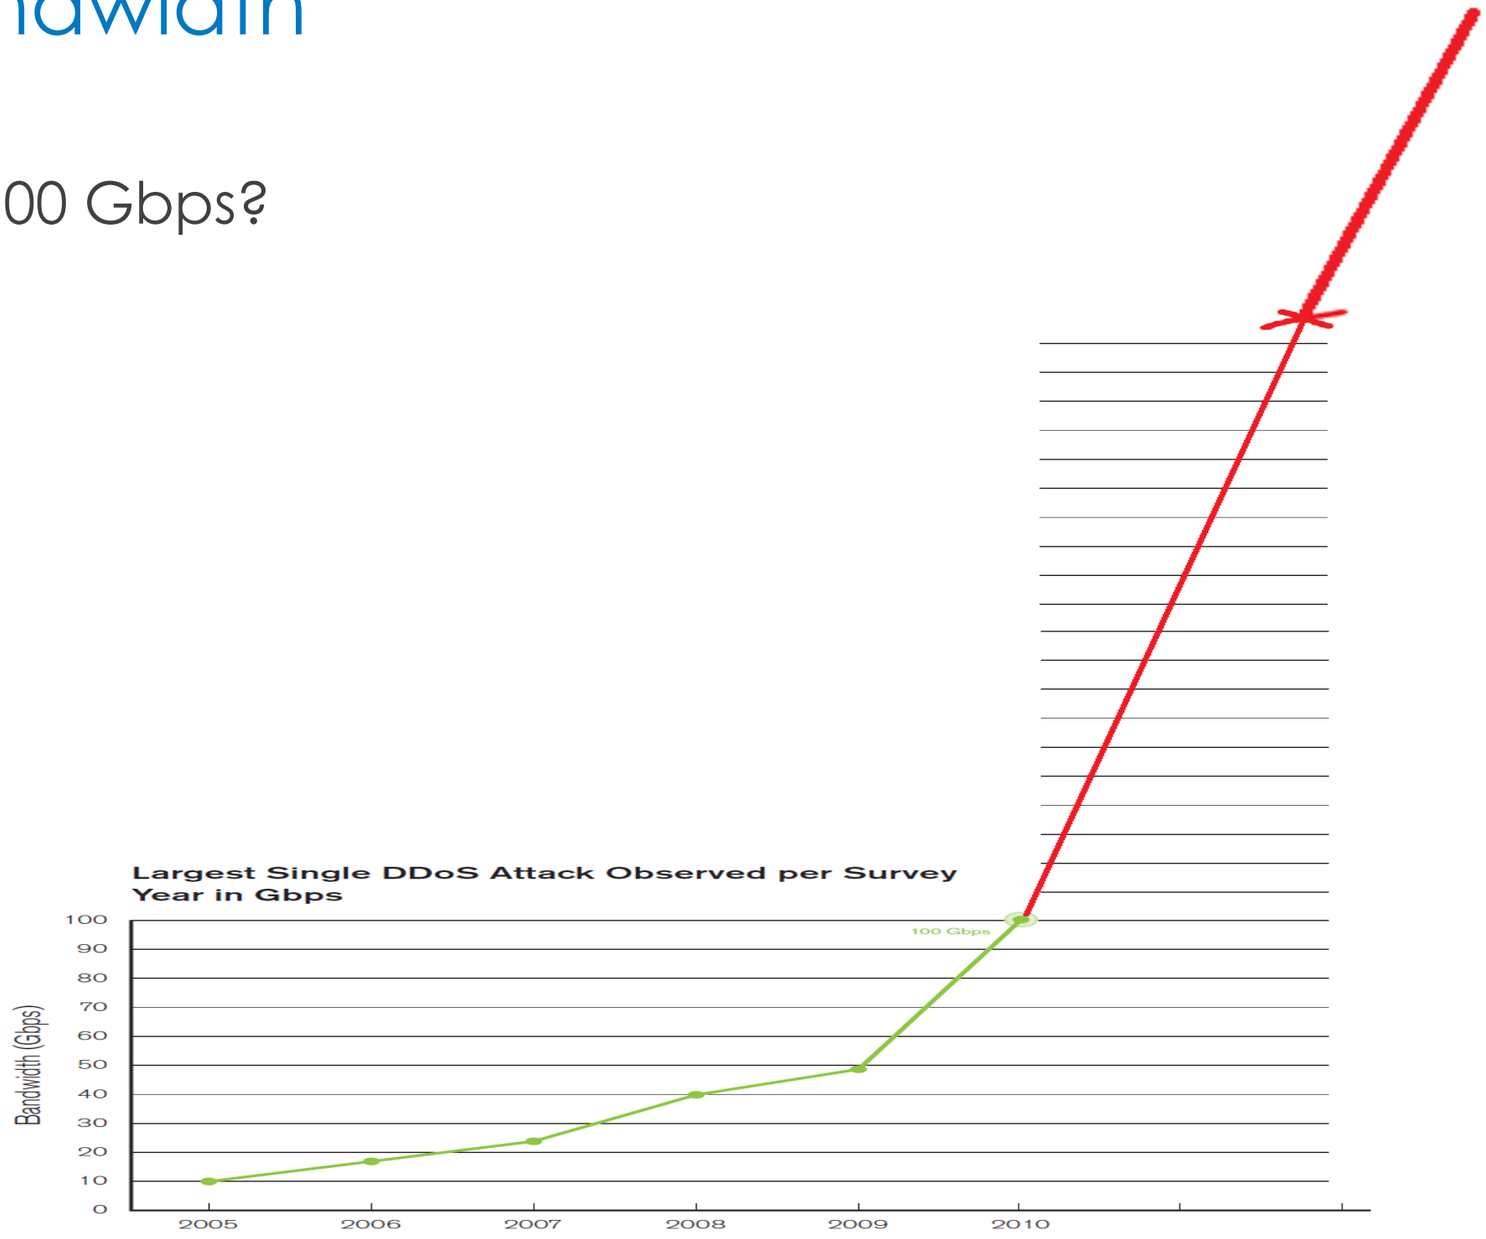

Figure 1
Source: Arbor Networks, Inc.

Contributing factors

- Not patched Content Management Systems (CMSes)
- Available reflectors (DNS, NTP, SSDP)
- …with ability to amplify
- More bandwidth available
- Unpatched embedded devices (mostly home routers)

Who is the adversary?

Adversary

- Wide range of attackers
  - High-school pranks
  - Frustrated “hackers”
  - Professional DDoS operators
  - State sponsored actors
  - Hacktivists
  - Did I miss anybody?

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 and there is a large reuse of code and services
  - This leads to clear signatures for some of the tools

Motivation

- Financial gain
  - extortion
  - taking the competition offline during high-gain events
- Political statement
- Divert attention (seen in cases with data exfiltration)
- Immature behavior
- etc.

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
- Automation
  - Manual (LOIC)
  - Automated
- Backscatter*
- Based on network layer

Metrics

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

Backscatter

- What is backscatter and why do I care?
- 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

Attack surface

Network Layers – OSI vs Internet Model

- OSI – Open Systems Interconnect

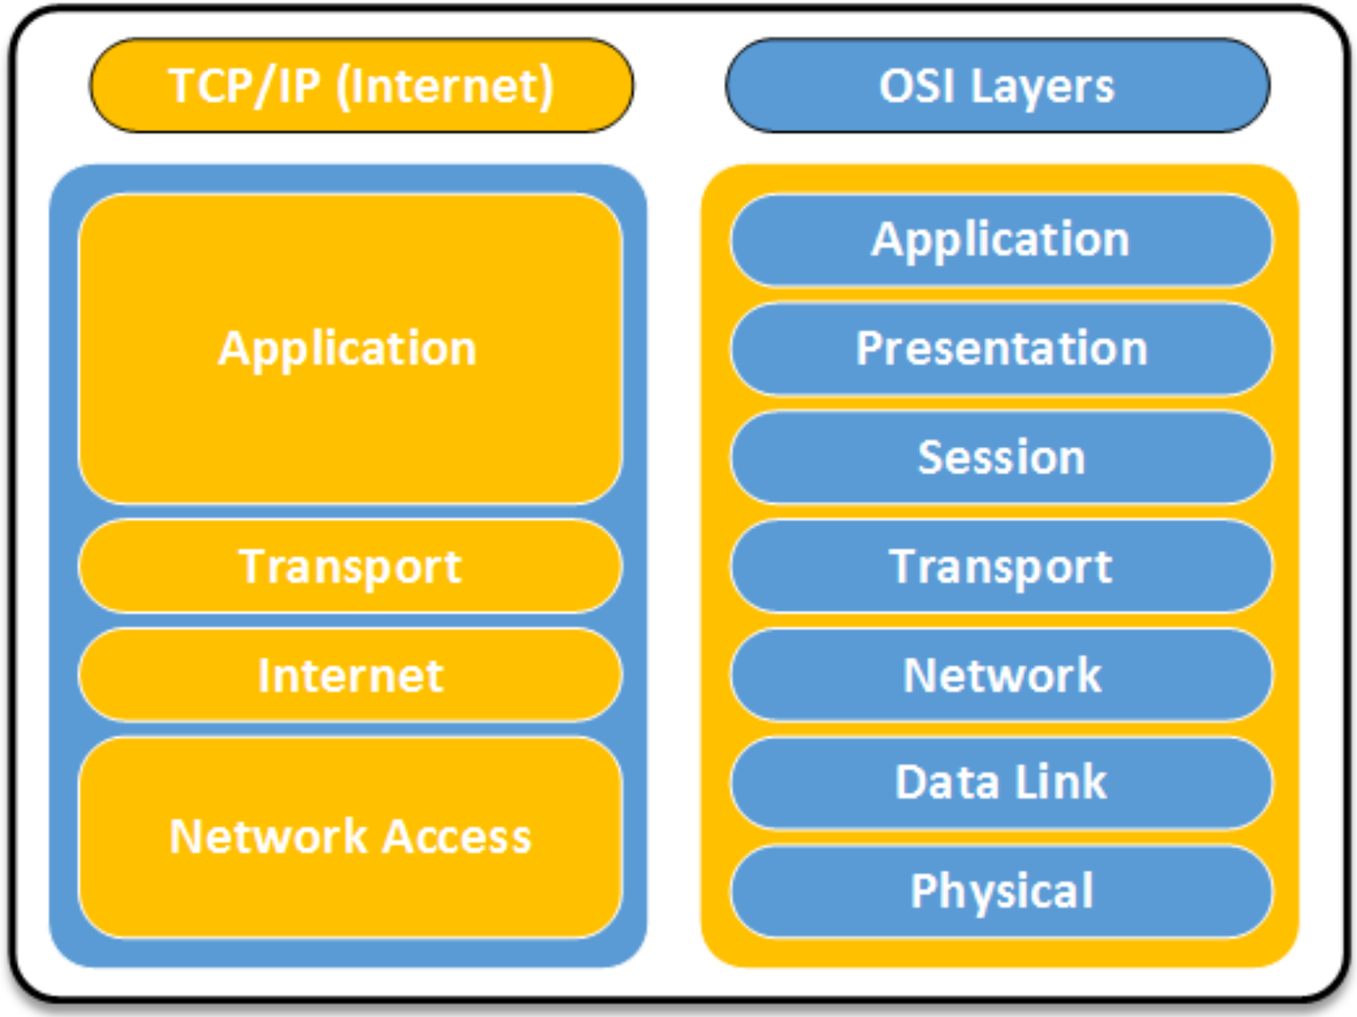

Physical and Data-link Layers

- Cut cables
- Jamming
- Power surge
- EMP

- MAC Spoofing
- MAC flood

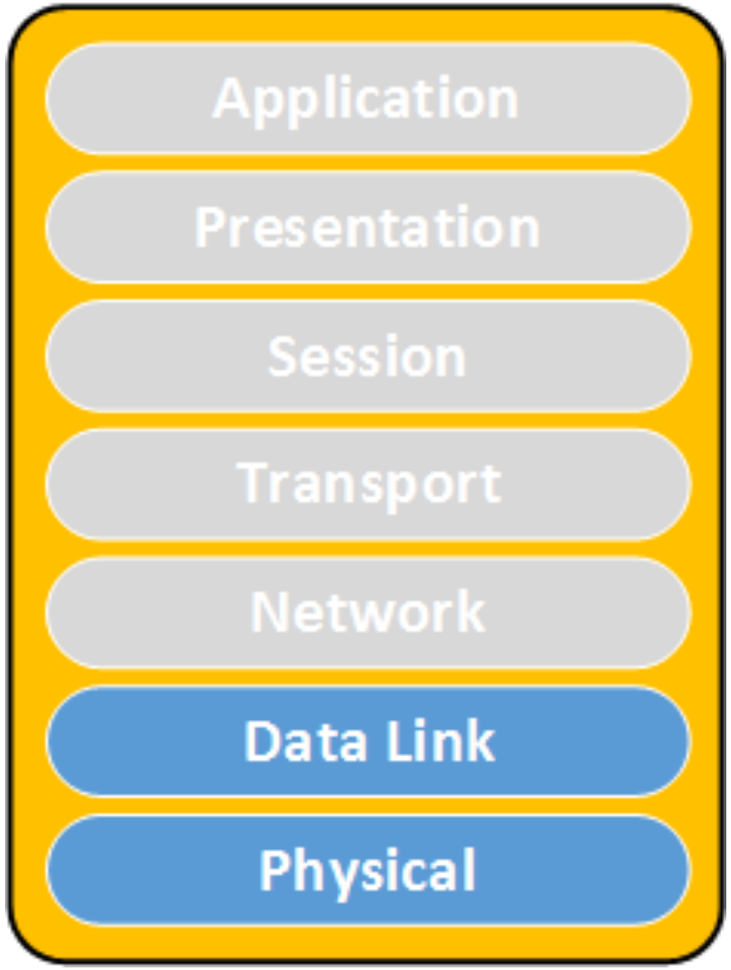

Network Layer

- Floods (ICMP)

- Teardrop
  (overlapping IP segments)

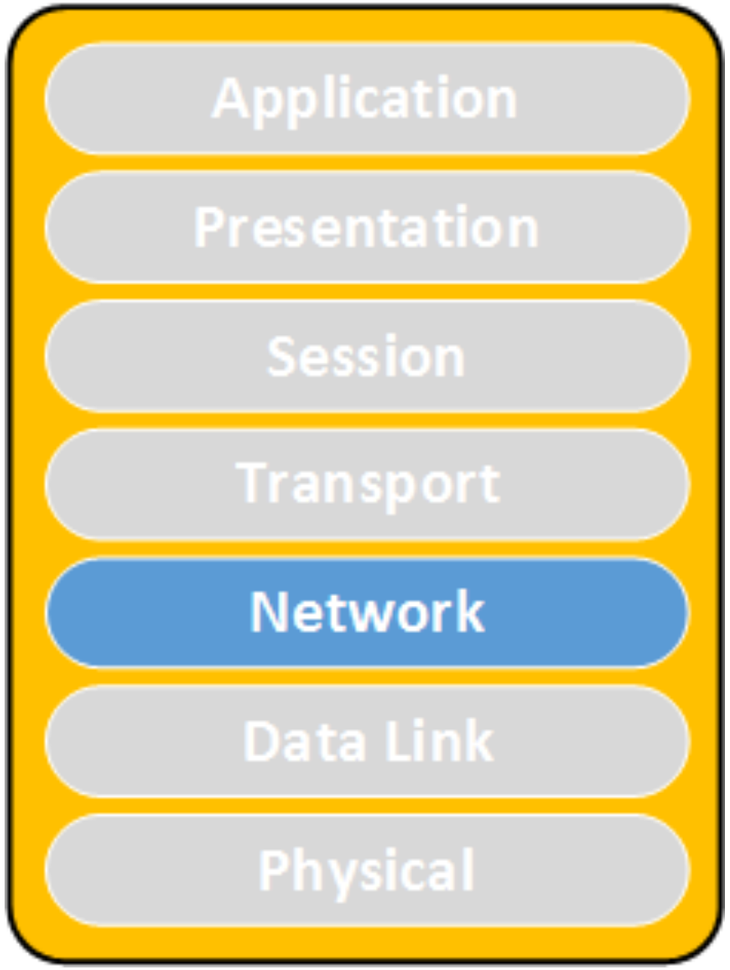

Transport Layer

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

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

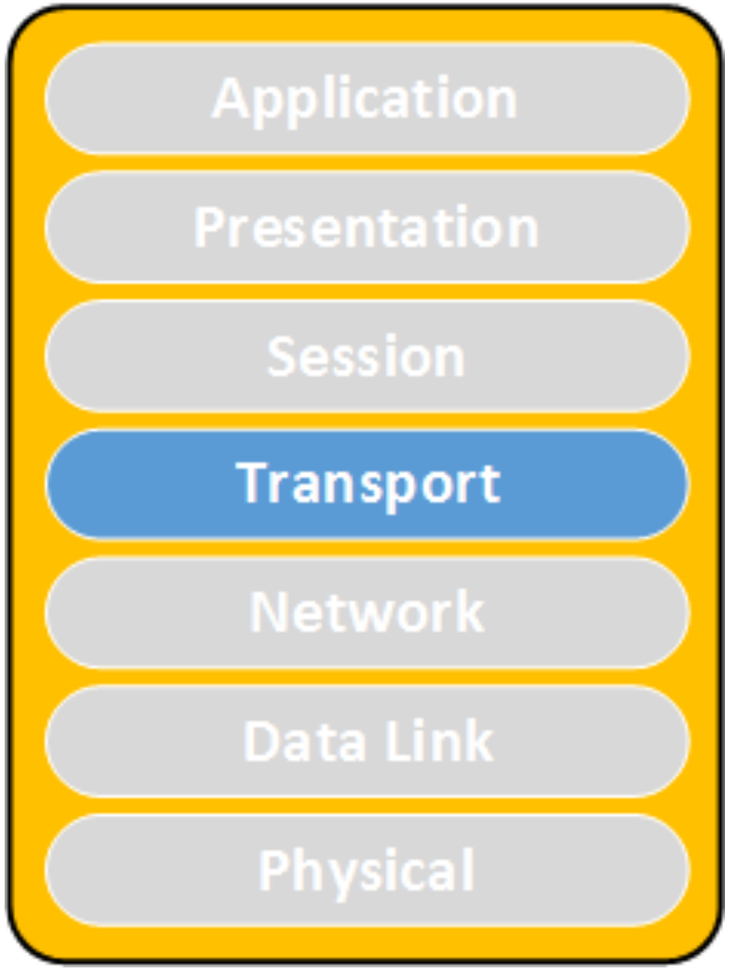

Session Layer

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

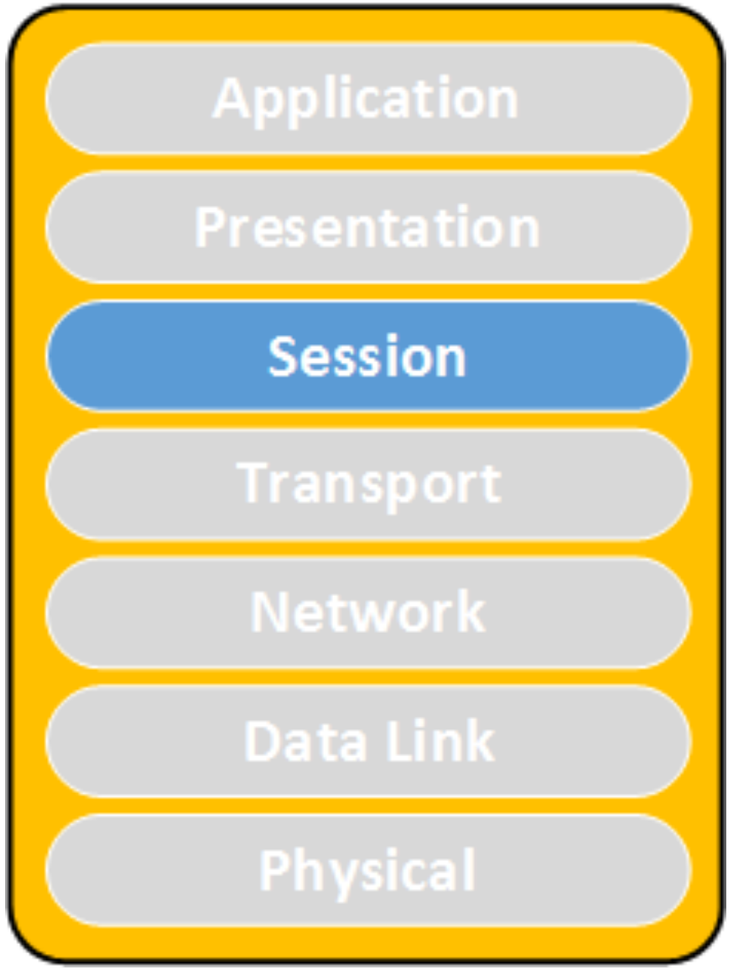

Presentation Layer

- Expensive queries (repeated many times)
- XML Attacks

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

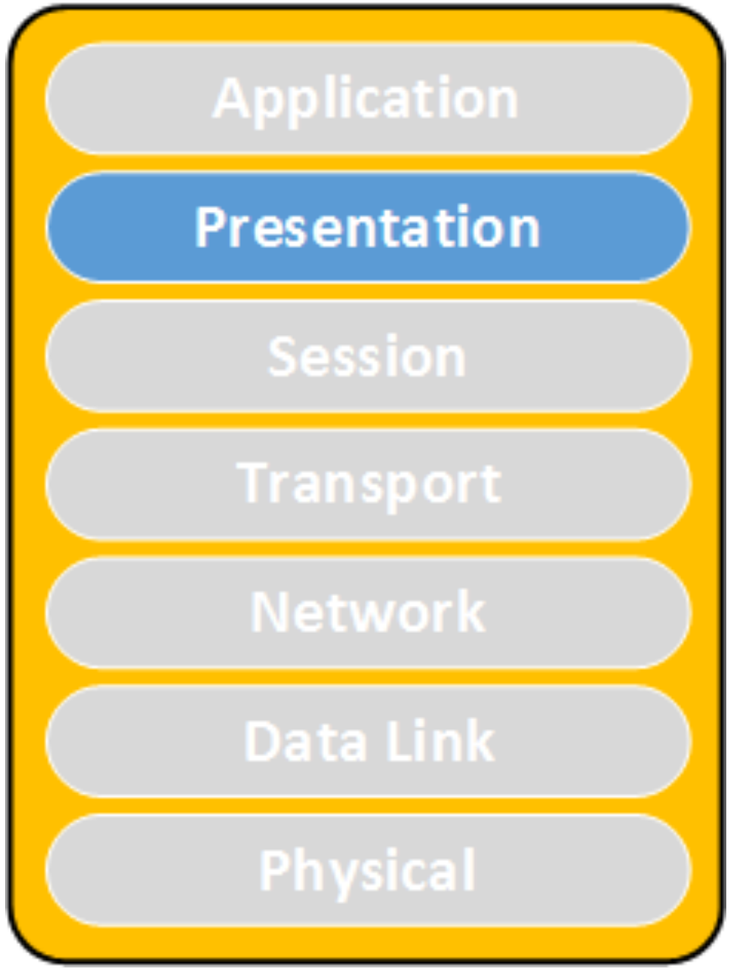

Application Layer

- SPAM?
- DNS queries
- Black fax

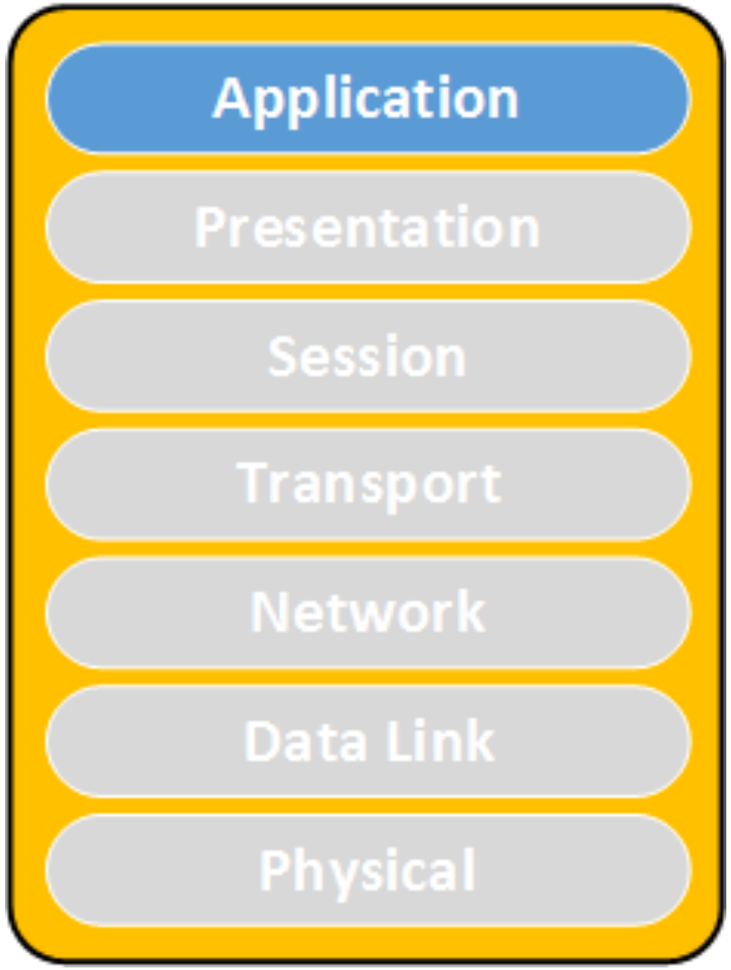

Attack summary by layer

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

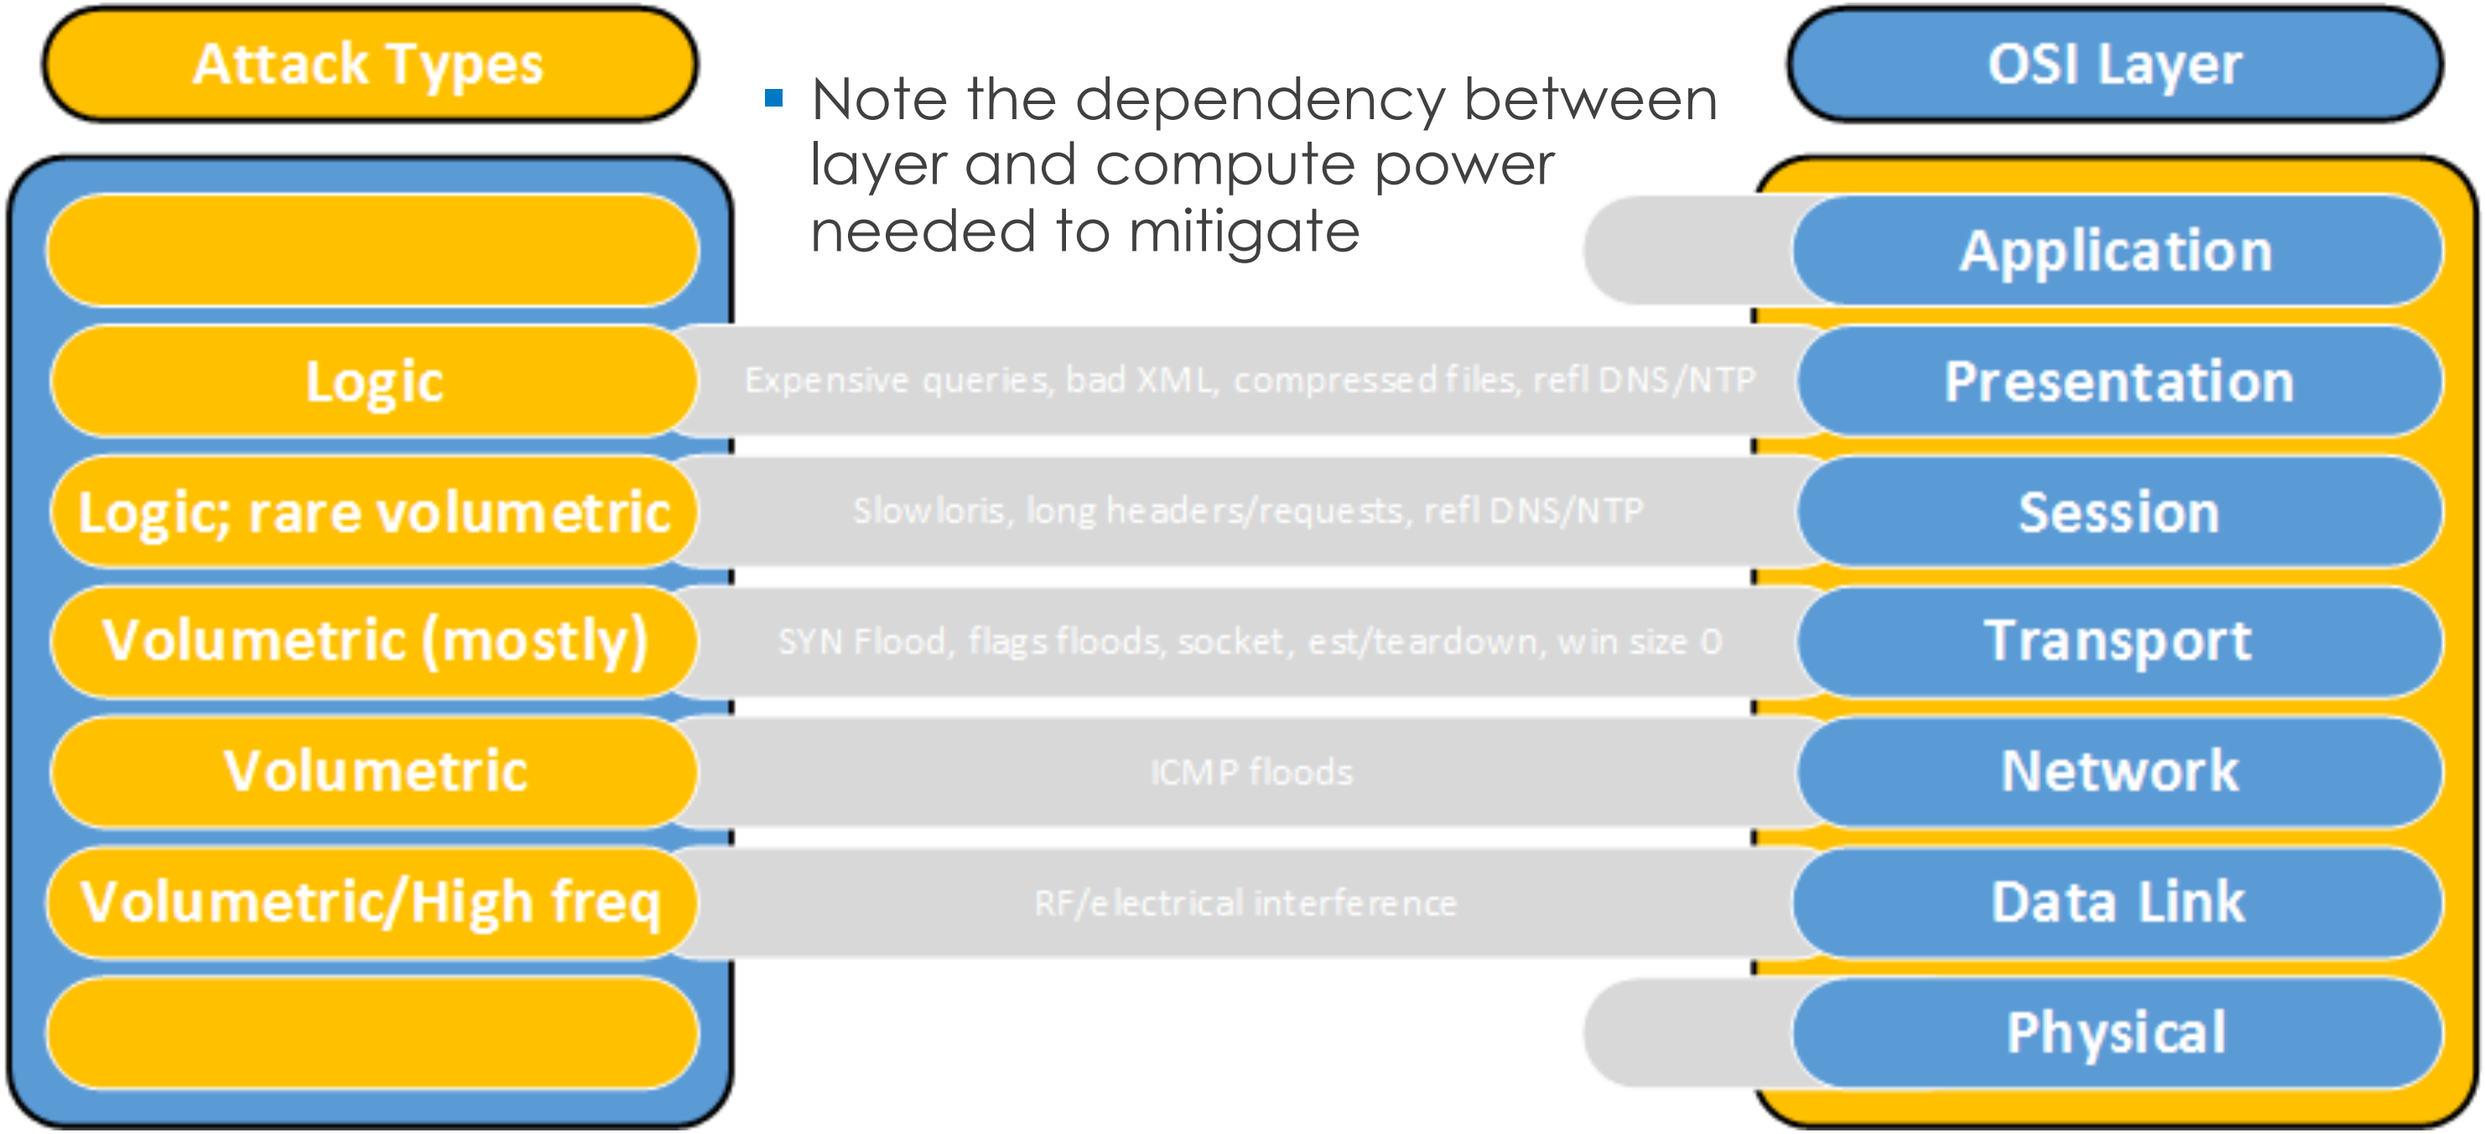

Transmission Control Protocol (TCP)

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

What is a SYN flood?

- What is a 3-way handshake?

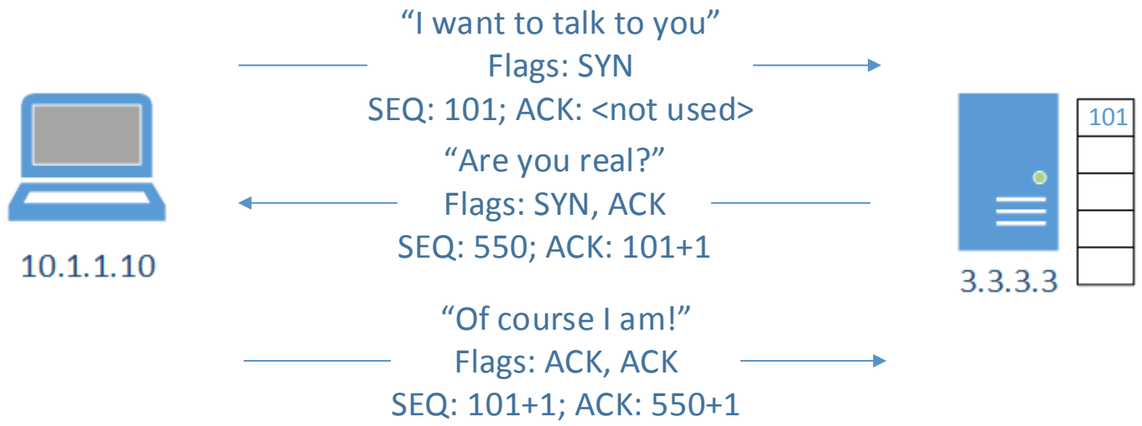

SYN flood

- Exploits the limited slots for pending connections
- Overloads them

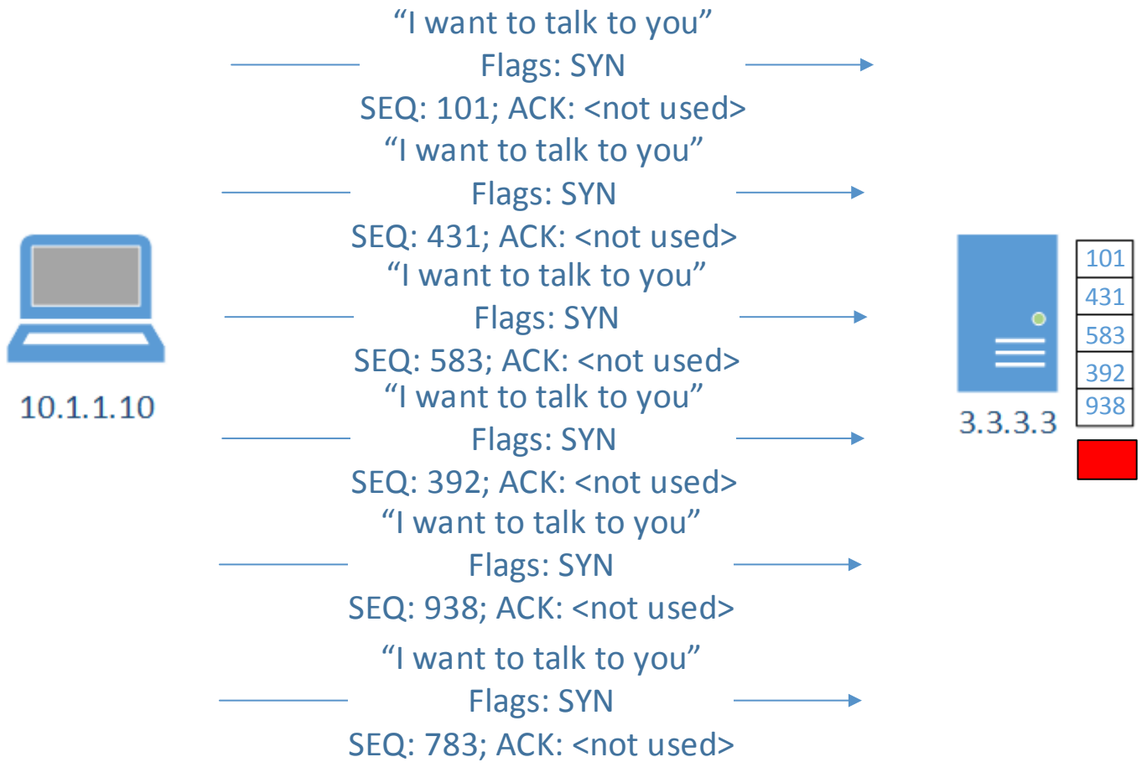

SYN flood through the eyes of netstat

- 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

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)
- Implement challenges using Nginx plugin – Roboo (ECL-LABS.ORG)
- Dedicated hardware devices

Lab: slowloris

- Open a web browser and go to the local web site
  - http://127.0.0.1
- Open a terminal
  - Go to the tools directory
  - Execute: ./slowloris.pl -dns 127.0.0.1
- Refresh the browser a few times to see the effect on it page

Lab: slowloris mitigation

- Reconfigure Apache to switch to port 8088
  - /etc/apache2/ports.conf
    - Hint: Listen 80
  - /etc/apache2/sites-available/000-default.conf
    - Hint: <VirtualHost *:80>
  - Restart Apache
    - root@ubuntu:/etc/apache2# /etc/init.d/apache2 restart
- Reconfigure Nginx to listen on port 80
  - /etc/nginx/sites-available/default
    - Hint: listen 88
  - Restart Nginx
    - /etc/init.d/nginx restart
- Repeat the previous experiment

DNS Resolution

DNS resolution

- Authoritative
- Open recursive

- www.a10networks.com.
- www a10networks com <root>

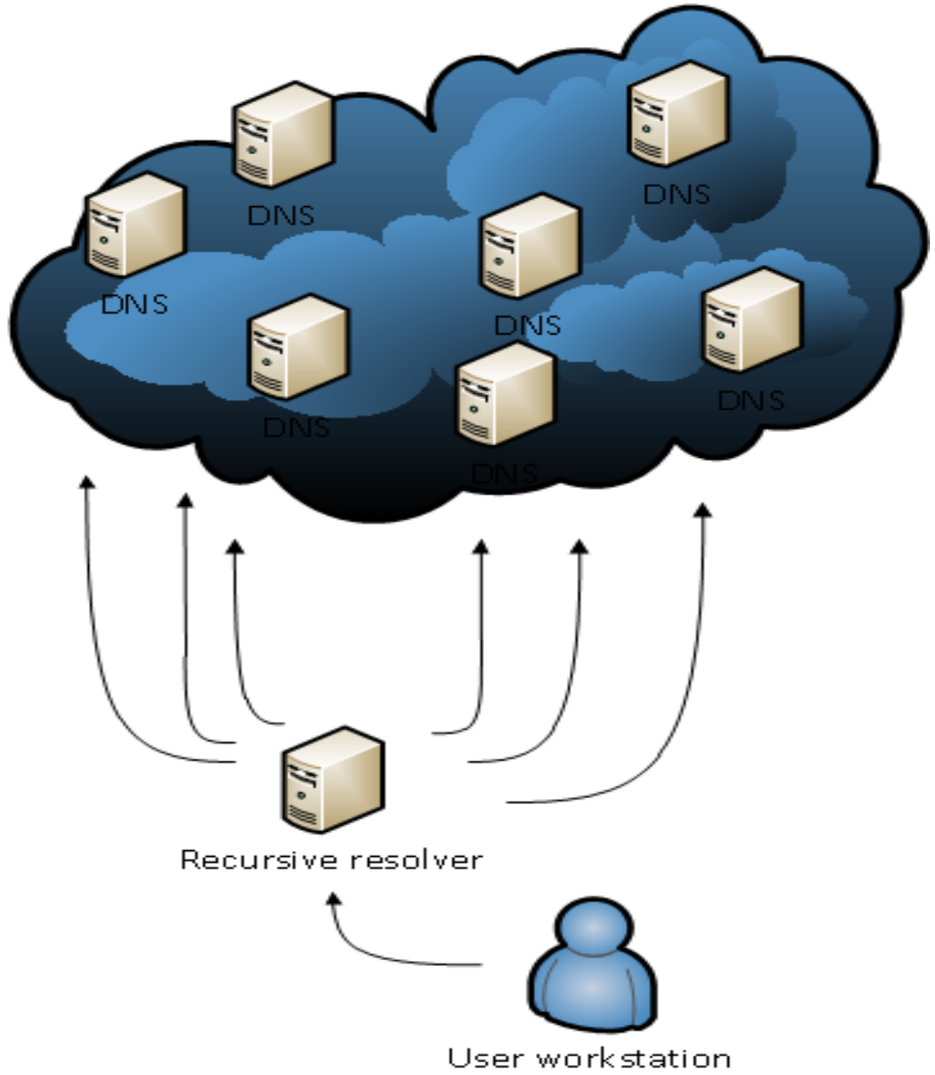

What is DNS resolution?

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

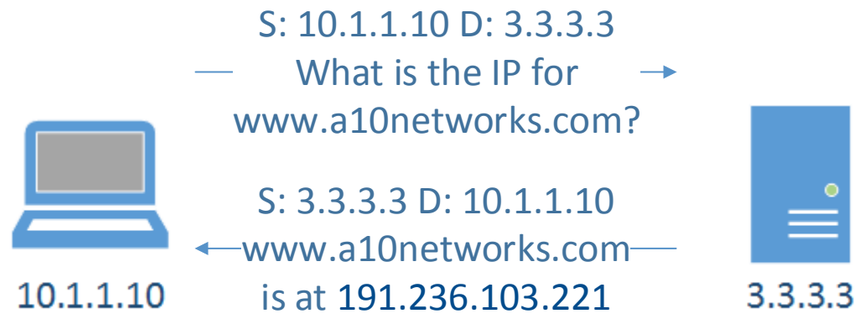

…if the answer was cached

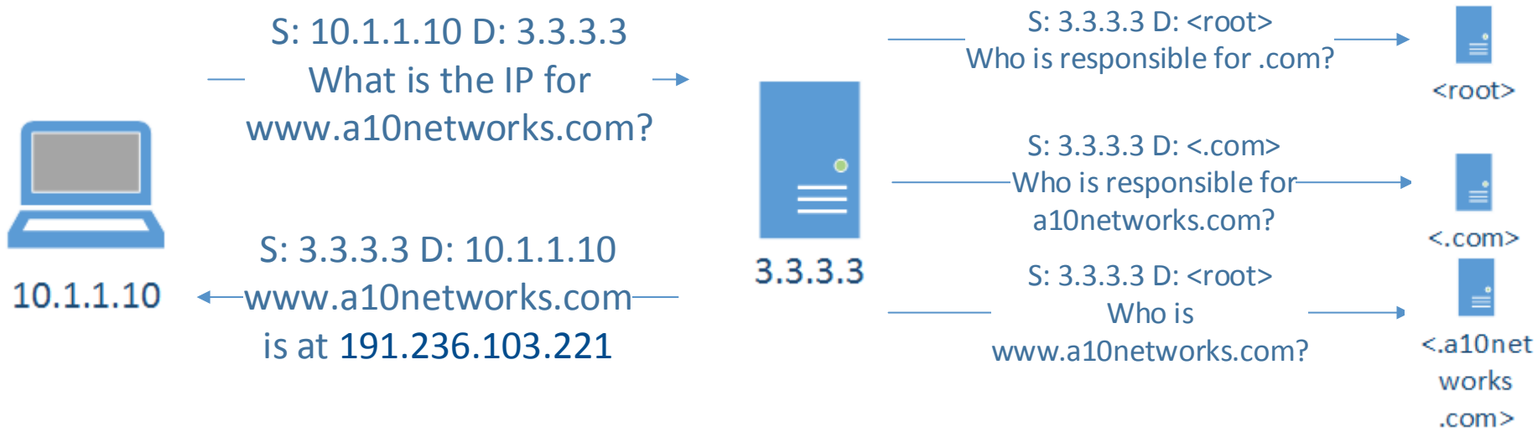

DNS Reflection
TODO: Reorganize

What is DNS reflection?

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

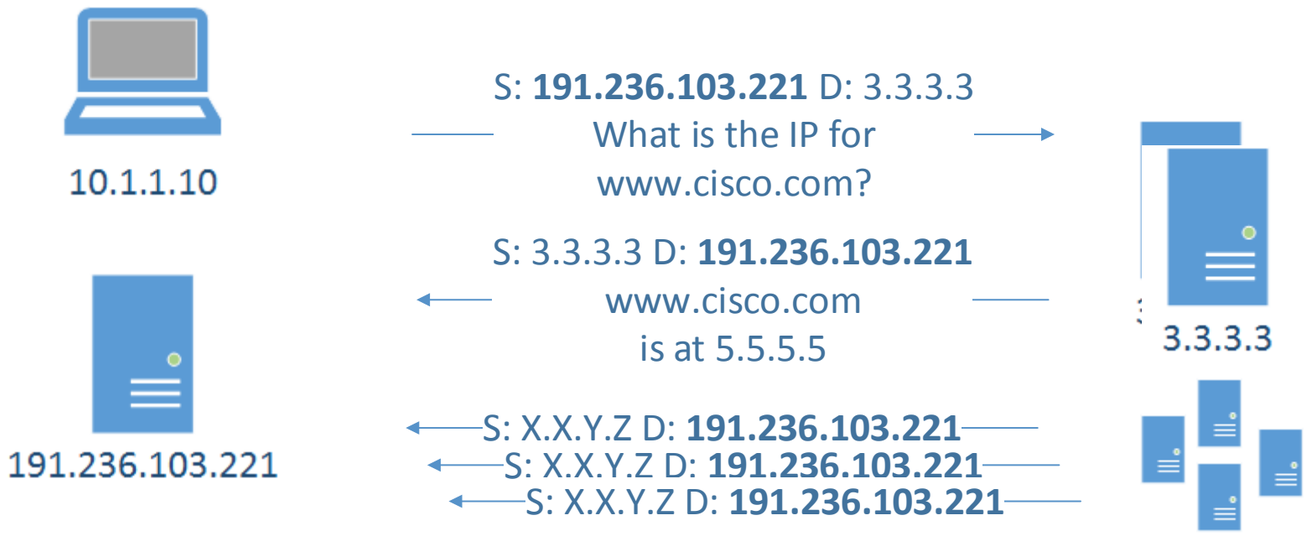

…the reflected traffic goes to the target server

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

Two concepts to remember

- Reflection
- Amplification

What is reflected attack

- Attacker spoofs the source with the IP of the victim
- Reflectors respond to the victim

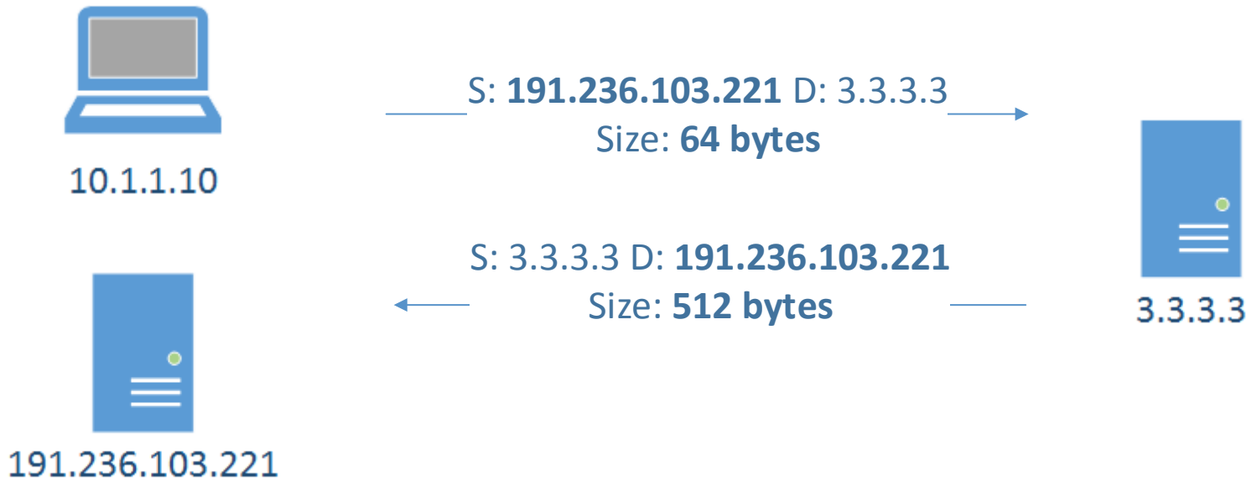

What is DNS reflection?

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

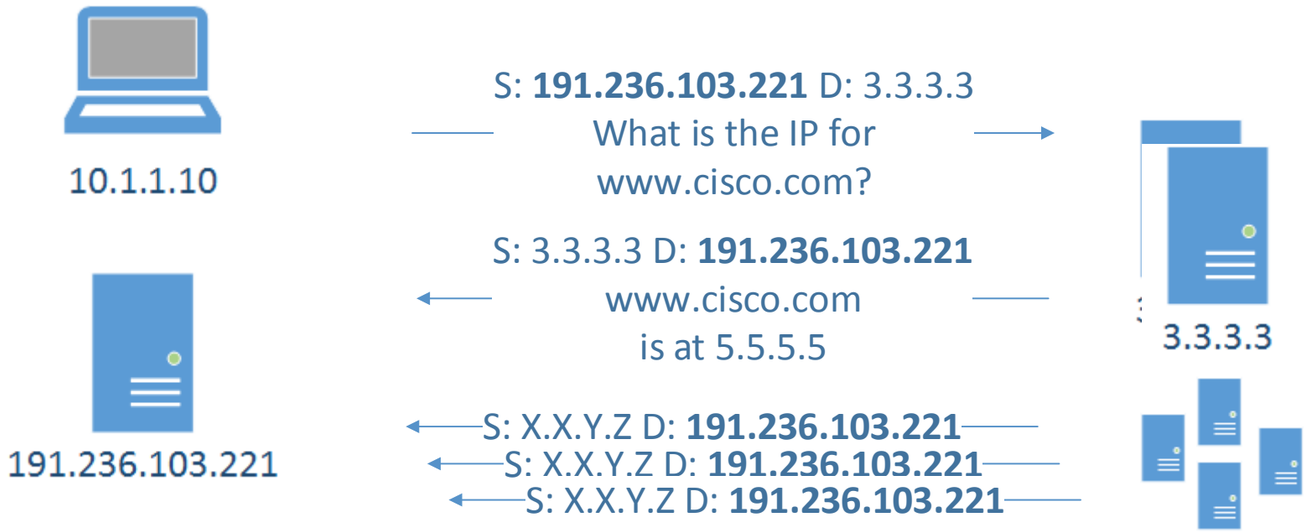

…the reflected traffic goes to the target server

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

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?

Reflector types

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

What is amplification attack?

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

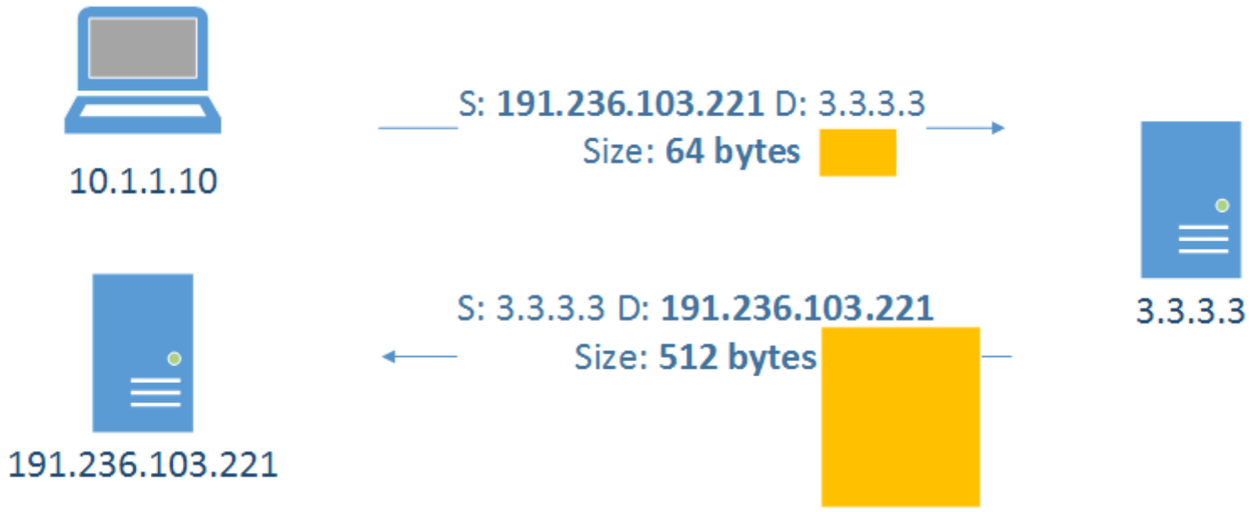

Amplification attacks

- The response to a request exceeds it by a large factor
- Protocols:
- DNS
- NTP
- SNMP
- SSDP
- What else?

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/oiODkk3yyHspciwdJvjIefu4PktdUnd1IQxW 791q/jWgHBL5iQQigBYv7Z5IfY1ENn+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
2erglomkBP79m9ahnFOxWEAaueA6TlHClGxOkgrk3hBtMFjUB9rhvkIm uxO2D8gc1DJDLl5egfpJCF2flTFhEvWzeMt6QGNwicWMxBsFHCxM7Fms D8I=

isc.org.          5725   IN     A       149.20.64.42

isc.org.          5725   IN     RRSIG   DNSKEY 5 2 7200 20130620130130 20130521130130 12892 isc.org. dfxTGA/f6vdhuIqojp+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. o18F3KIFkYedFRw1e5MP4qDo3wSg0XK9I5WCYD75aGhs9RI5eyc/6KEW
Se4lZXRhf6d77xXlerMYCrsfh/GHdjPRoE1xL/nzH/hTBJAI9XDbC5I/ EUpFIGVLVdQy43XKtywm0j2nyc5MdGa2VeLKo+hHTmH3St3pGRVJp2IK 5Z0=

isc.org.          5725   IN     DNSKEY  257 3 5 BEAAAAOhHQDBrhQbtphgq2wQUpEQ5t4DtUHxoMVFu2hWLDMvoOMRXjGr hhCeFvAZih7yJHf8ZGfW6hd38hXG/
xylYCO6Krpbdojwx8YMXLA5/kA+ u50WIL8ZR1R6KTbsYVMf/Qx5RiNbPClw+vT+U8eXEJmO20jlS1ULgqy3 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+/x6eT4QLru0KoZkvZJnqTl8JyaFTw2OM/ItBfh/
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

c.root-servers.net.    297373  IN      A       192.33.4.12
```

Reflection and Amplification

10.1.1.10

191.236.103.221

S: **191.236.103.221** D: 3.3.3.3

What is ANY isc.org

S: 3.3.3.3 D: **191.236.103.221**

3.3.3.3

```
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. KHMs09 DaFMw416/7xXhaD9By0NrqCiQ4kBnqi6oq2VocZRREAbUHHAY
KydlgKO5vOavvd1Fy86/oiODkk3yyHspciwdJvjefu4PktdUnd1IQixW791q/
jWgHBL5iQiQlgBYv7Z5ffY1ENn+6fPOchAyvvNqEBYcdqVNBpzxOjtxiU=
isc.org.            481     IN      DS      12892 5 2
F1E184C0E1D6155D20EB3C223ACED3B03C773DD952D5F0EB5C777586DE18DA6B5
isc.org.            481     IN      DS      12892 5 1
98211D08B4C8A1D9F6AEE1E2237AEF69F3F9759
isc.org.            5725    IN      RRSIG   A 5 2 7200 20130620134150 20130521134150 50012
isc.org. iCBy1JjPP6mXVYjpSc62JCrZWHnVYAUGHb7WxRmxGRoip88P+LOvRl
2erglomkBP79m9ohnFOxNEAoueA6TIHCiGxOkgnk3hBfMfjLB9rhvichn
uxC02D8gc1DJDU5egfpJCF2ftFhEvNkeMvt6QGNwicWWkBsFHCxM7Fms= D8I=
isc.org.            5725    IN      A       149.20.64.42
isc.org.            5725    IN      RRSIG   DNSKEY 5 2 7200 20130620130130 20130521130130
12892 isc.org. dfxTGA/f6vdhuiqojp+Konkd8c4y3WiJHVs5TjnvhdEyHI4qPh/cHh
+y1vA6+gAwTHiX+GpscctNxEIvoSvWu3m9Nocnivu//AZGioL/SyDgEsIbJMY
X+ZXY5qrgGrV2grOcKAAA91Bus8beIhYQiZTsdaH2TS+AKjKINEgvm
yQi5xNEodzE3p0ygfPq4eMNO4FRT9UQDhTRD3v8zfxFNXKvBsQWZGBH0
5+QicbC6xnGyn1bBptJEEGnCBG01ncJ+IMCyEf86VGHKJfeovORIrDQ3
cjJRFPTCOkA8n4jBvnsimUP/TGI+Mg4ufAZpE96jJhvFBadcC/iOo6i XicGVIA==
isc.org.            5725    IN      RRSIG   DNSKEY 5 2 7200 20130620130130 20130521130130
50012 isc.org. o18F3KIRkYedFRwIe5MP4qDo3vi8g0XK95WCYDT5oGhs9R5eyc/6KEW
Se4IZXRhf6d77xXierMYOrafh/GHdJPRoE1xL/nzH/hTBJA9XDbC5I/
EUpFIGVLVdGy43XKtyvm0J2nycSMdGa2VeLKo+hHTmH3S+3pGRVJp2IK5Z0=
isc.org.            5725    IN      DNSKEY  257 3 5
BEAAAAOhHQDBrhGb+phgq2wGUpEQ5t4DtUHkoMMFu2hWLDMvoOMRXjGr
hhCeFvAZh7yJHfBZGRMN6hd38hXG/xyIYCO6KrpbdoJwx6YMKLA5/kA+
u50WL8ZR1R6KTbzYMf/Qx5RNbPCIwtvT+L8eXEJmO20j81ULgqy3 47cBB1zMvnnz/
4LuJpA0do9QbKjBA2547515aNIMctwsB8/2+2E63/zZrGtzBkjDBrN/
9BexjpIks8jRhZo+EaXn3dTy47R09 Ux5WcJ++xaqZ7+yayL
KOOed839Z7SDmsn2eA0FKKqpvvA6LXeG2wtjxmvu9oA8IVUgEf/reCr/bB yBNsO70aEFTd
isc.org.            5725    IN      DNSKEY  256 3 5 BGEAAAABwuHz9Cem0BJ0JGIO7C/
a8MtR6hMbuFja1dfG/inaJpYv7vH XTrAOm/MeKp+/x6eT4GLru0KaZkvZJnqTI8JyaRFv20MV/
HBfh/hL2ImCft2O7n3MfeqY+vjPnY7dWghYW4aVfH7VVEGm958a9nfi79532Qekkxh
x8pXWdeAaRU=
isc.org.            5725    IN      DNSKEY  256 3 5 BGEAAAABwuHz9Cem0BJ0JGIO7C/
a8MtR6hMbuFja1dfG/inaJpYv7vH XTrAOm/MeKp+/x6eT4GLru0KaZkvZJnqTI8JyaRFv20MV/
HBfh/hL2ImCft2O7n3MfeqY+vjPnY7dWghYW4aVfH7VVEGm958a9nfi79532Qekkxh
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
c.root-servers.net.  297373  IN      A       192.33.4.12
d.root-servers.net.  297555  IN      A       199.7.91.13
d.root-servers.net.  417805  IN      AAAA    2001:500:2d::d
e.root-servers.net.  297707  IN      A       192.203.230.10
f.root-servers.net.  297544  IN      A       192.5.5.241
f.root-servers.net.  416152  IN      AAAA    2001:500:2f::f
g.root-servers.net.  297708  IN      A       192.112.36.4
h.root-servers.net.  298308  IN      A       128.63.2.53
h.root-servers.net.  416776  IN      AAAA    2001:500:1::803f:235
i.root-servers.net.  297617  IN      A       192.36.148.17
```

DNS Rate limits

- 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

```
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

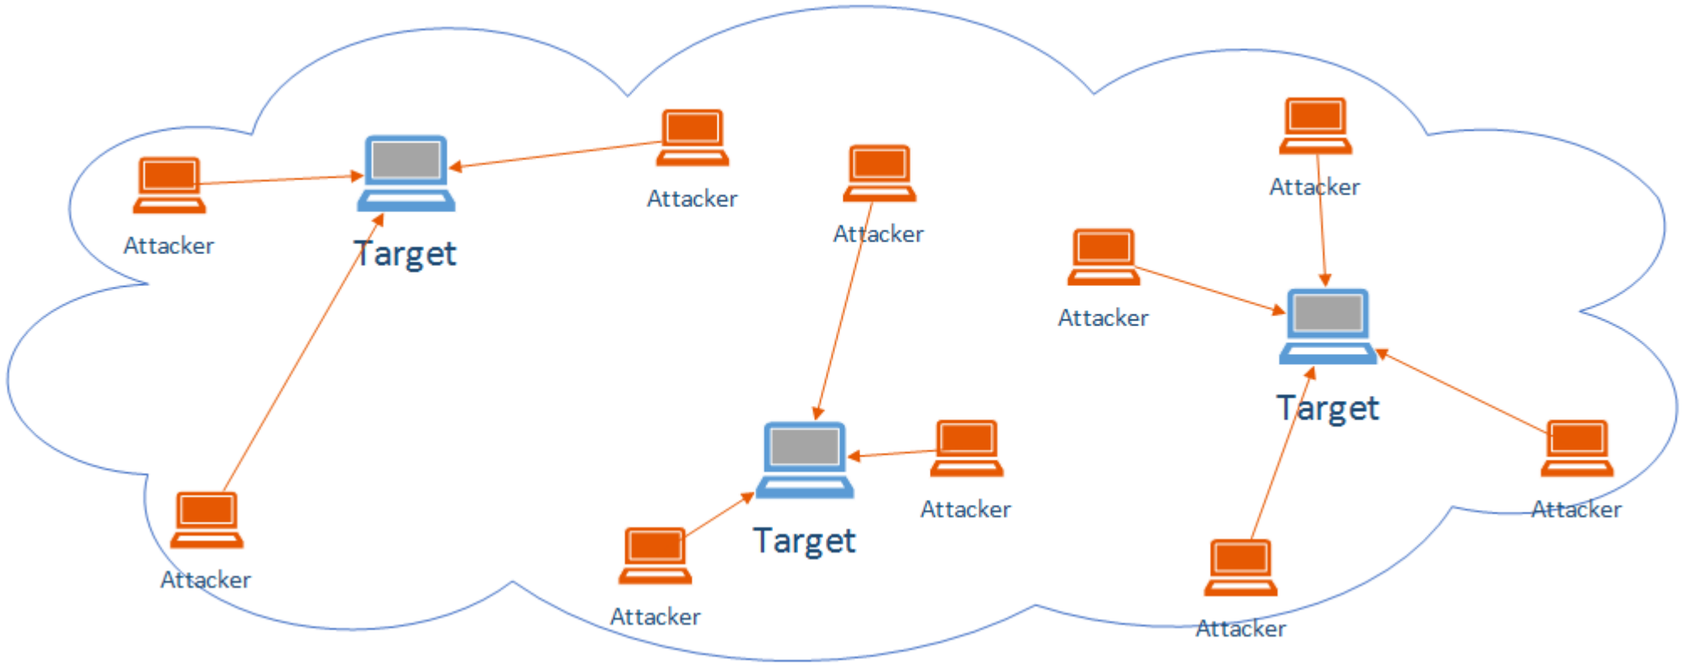

IPS/DDoS mitigation gear

- 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

- What's next?

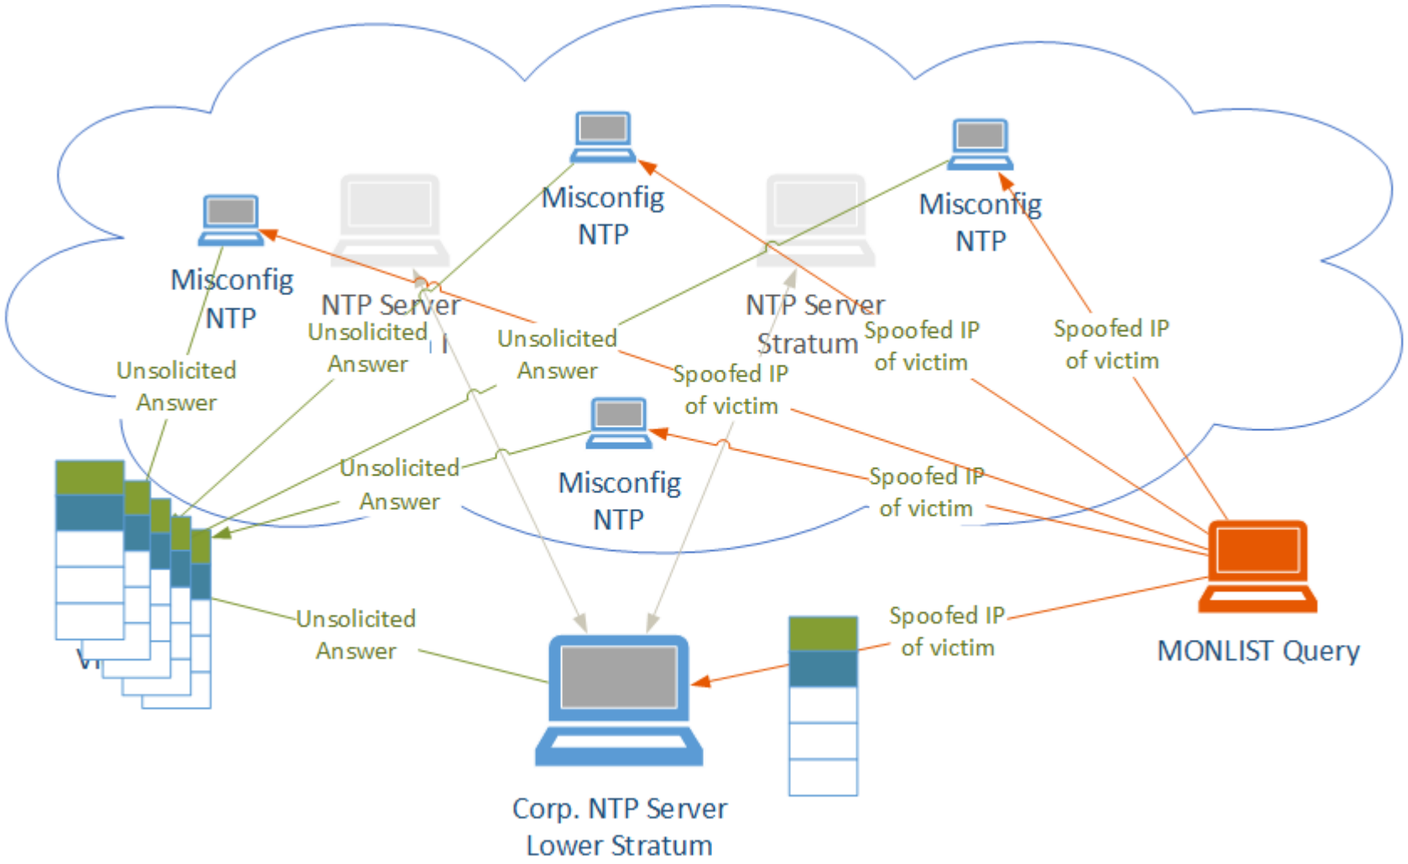

Good Internet citizenship

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

But WHY?!?!?

Why would you ever want to run an open DNS resolver?

- OpenDNS/Google DNS
- Authoritative name servers (non-recursive of course)
- Because you have not read the Cricket book

Are you a reflector?

- In some cases return traffic/backscatter

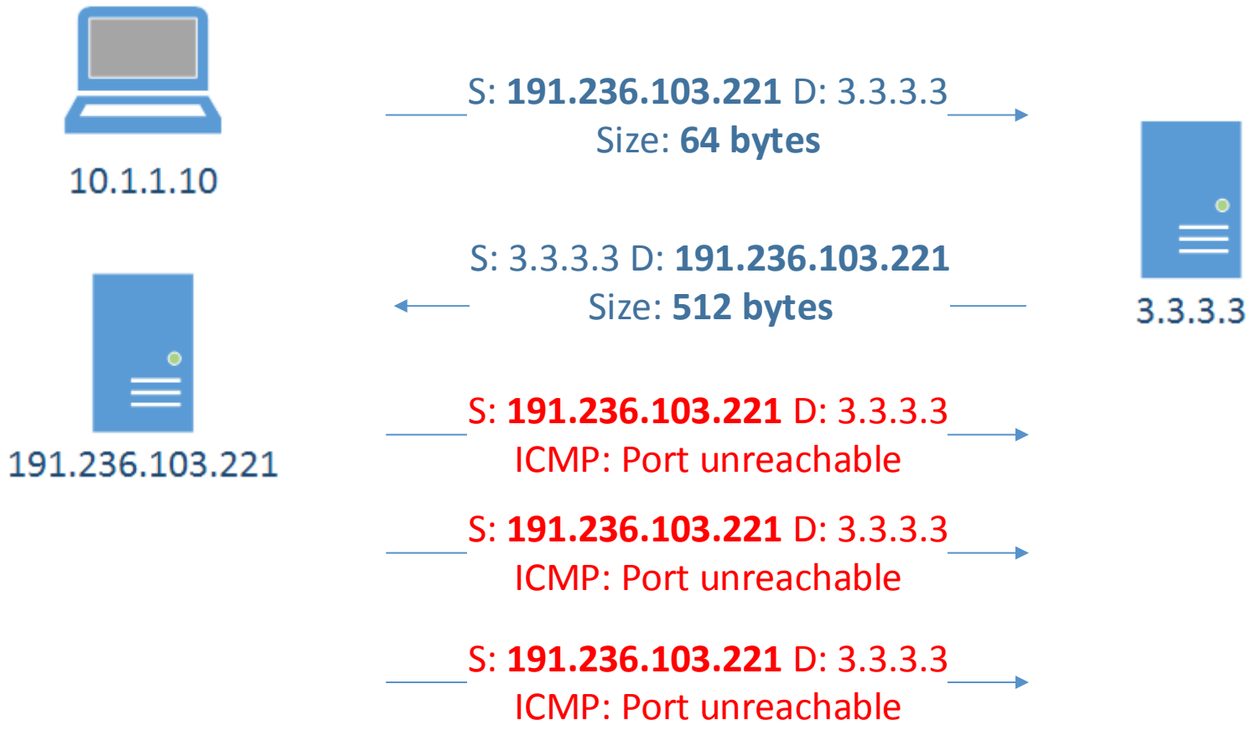

Resources

- 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.
- Go 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
- Dive into specifics what attack surface the different layers offer
- Discuss reflection, amplification and back scatter
- Terminology
- Tools

Thank you