CSE127, Computer Security

Network Security III

UC San Diego

Housekeeping

General course things to know

- PA4 is out, due on **Thursday**
  - CTF-style assignment, very little instruction, it's fun!
- PA5 released on **Friday**
  - Focuses on Cryptography… next unit in class
  - Last PA… due last class on week 10
- Note, due to travel class is cancelled on **3/10**

Previously on CSE 127…

Recap

- We've been talking a lot about networking attacks
  - MiTM attacker, Eavesdropper, Off-Path attacker
- We've discussed several types of attack classes…
  - Spoofing (IP, ARP, BGP hijacking…) — you might have to do some spoofing on your PA…
  - Poisoning (DNS cache poisoning)

Today's lecture — More attacks, some defenses

Learning Objectives

- Learn about denial-of-service — a classic attack that attacks *availability* of network resources (and is very common on the Internet), and the ways in which DDoS can be implemented
- Learn about network censorship**,** and how malicious ISPs / bad faith actors can manipulate the Internet for end-users
- Discuss ways to mitigate some of these problems through a bevy of network defenses (e.g., firewalls, port forwarding, intrusion detection, VPNs, NATs, etc.)

DoS + DDoS

What is denial of service (DoS)?

What is denial of service (DoS)?

- **Definition:** Any attack that prevents legitimate access to a network service, disrupting *availability*
  - Super broad! Can encompass many different threat models + motivations, and can take place at different positions / layers in the OSI stack

Hyper-Volumetric DDoS Attacks Reach Record 7.3 Tbps, Targeting Key Global Sectors

Ravie Lakshmanan Jul 15, 2025 Botnet / Network Security

New HTTP/2 'MadeYouReset' Vulnerability Enables Large-Scale DoS Attacks

Ravie Lakshmanan Aug 14, 2025 Server Security / Vulnerability

DoS motivations…

- Take three minutes and brainstorm: Who might want to launch DDoS attacks? Why would they want to do this?

DoS motivations…

- Take three minutes and brainstorm: Who might want to launch DDoS attacks? Why would they want to do this?
  - Law enforcement — can legally request service providers to take stuff down
    - Domain seizure, warrants to hosting providers, etc.
  - Governments — that operate networks can restrict access to “unwanted” domains, or even target other governments…
    - E.g., network censorship, advanced persistent threats (APTs)
  - Competitors — businesses can try to knock each other offline or damage adversary infrastructure
    - E.g., finance, gaming, e-commerce… yes, it happens, and it’s crazy

Russia accused of unleashing cyberwar to disable Estonia

- Parliament, ministries, banks, media targeted
- Nato experts sent in to strengthen defences

Ian Traynor in Brussels
The Guardian, Thursday 17 May 2007
Article history

A three-week wave of massive cyber-attacks on the small Baltic
of Estonia, the first known incidence of such an assault on a state
causing alarm across the western alliance, with Nato urgently exa
the offensive and its implications.

August 11th, 2008

Coordinated Russia vs Georgia cyber attack in progress

Posted by Dancho Danchev @ 4:23 pm

Categories: Black Hat, Botnets, Denial of Service (DoS), Governments, Hackers...
Tags: Security, Cyber Warfare, DDoS, Georgia, South Osetia...

62 TalkBacks ADD YOUR OPINION | SHARE | PRINT | E-MAIL | WORTHWHILE? +18 24 VOTES

In the wake of the Russian-Georgian conflict, a week worth of speculations around Russian Internet forums have finally materialized into a coordinated cyber attack against Georgia's Internet infrastructure. The attacks have already managed to compromise several government web sites, with continuing DDoS attacks against numerous other Georgian government sites, prompting the government to switch to hosting locations to the U.S, with Georgia's Ministry of Foreign Affairs undertaking a desperate step in order to disseminate real-time

Extortion via DDoS on the rise

By Denise Pappalardo and Ellen Messmer, Network World, 05/16/05

Criminals are increasingly targeting corporations with distributed denial-of-service attacks designed not to disrupt business networks but to extort thousands of dollars from the companies.

Ivan Maksakov, Alexander Petrov and Denis Stepanov were accused of receiving $4 million from firms that they threatened with cyberattacks.

The trio concentrated on U.K. Internet gambling sites, according to the prosecution. One bookmaker, which refused to pay a demand for $10,000, was attacked and brought offline--which reportedly cost it more than $200,000 a day in lost business.

November 17th, 2008

Anti fraud site hit by a DDoS attack

Posted by Dancho Danchev @ 4:01 pm

Categories: Botnets, Denial of S

Tags: Security, Cybercrime, DDoS

9 TalkBacks

ADD YOUR OPINION

SHAR

continuing to hit the site with
of malware infected hosts mo

U.S. Charges 37 Alleged Mules and Others in Online Bank Fraud Scheme

By Kim Zetter | September 30, 2010 | 3:07 pm | Categories: Crime, Cybersecurity, Hacks and Cracks

Follow @KimZetter

120 Tweet

0 +1

Share

Beyrouti, Babbo and Vitello worked with hackers who breached brokerage accounts at E-Trade and TD Ameritrade. The hackers then executed fraudulent sales of securities and transferred the proceeds from the sale to the mules' accounts. The receiving accounts were set up in the names of shell companies and linked to the hacked accounts.

Meanwhile, the victims' phones received a barrage of calls to prevent the brokerage firms from contacting them to confirm the legitimacy of the transactions. When the victims answered their phone, they would hear silence or a recorded message. About $1.2 million was transferred to shell accounts opened by the suspects, who then transferred the money to other accounts in Asia or withdraw the money from ATMs in the New York area.

Booter / Stresser Services

- Attackers with DoS capabilities sometimes run out of their own victims…. so sometimes, they will rent out services to the “public”

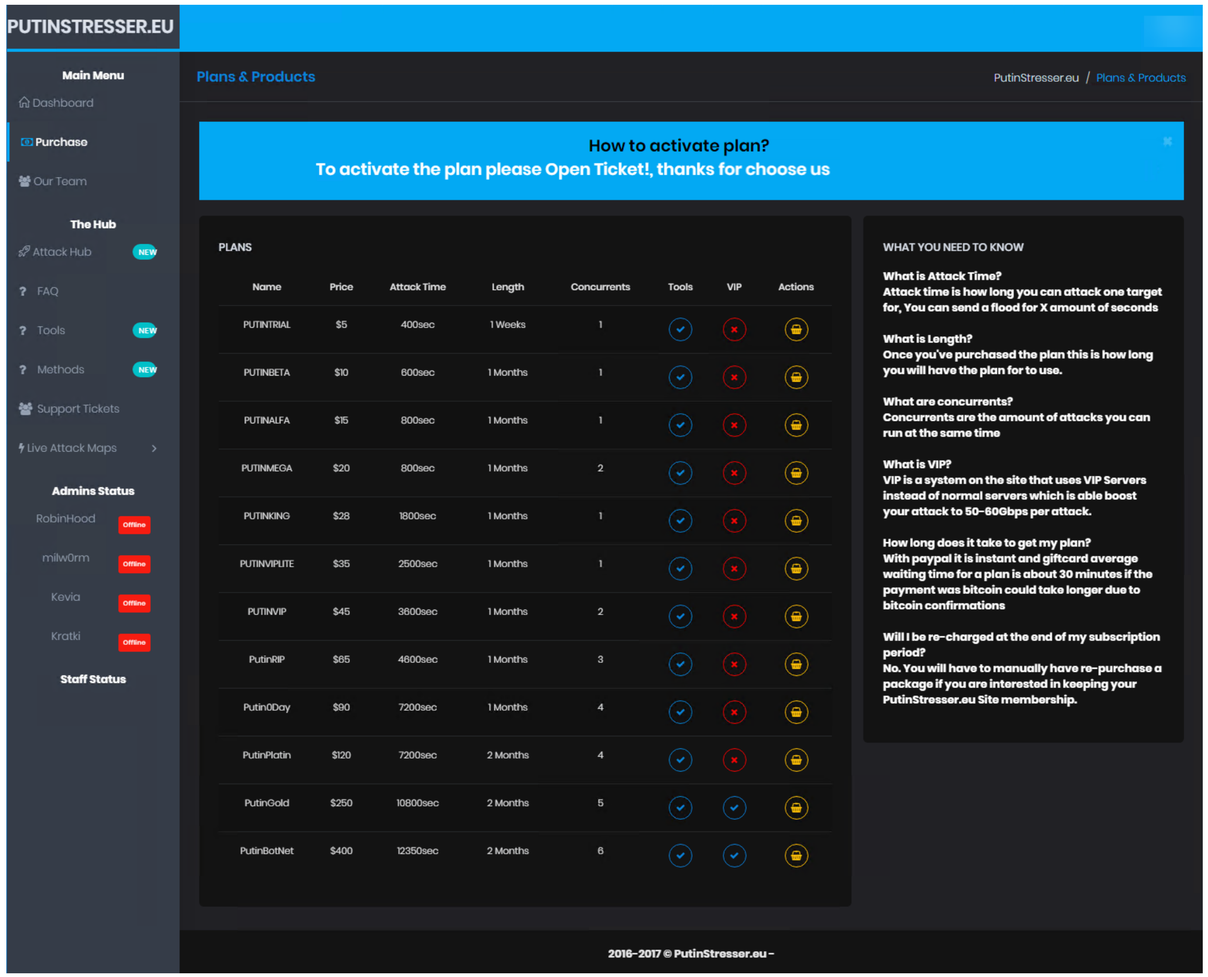

Two kinds of attacks

- Two classes of DoS:
  - **Logic-based**
    - Exploits some fundamental problem in the software that renders the server useless
    - “Ping of death” — sending a malformed ICMP ping used to crash a system
  - **Flooding-based**
    - Overwhelm resources by sending lots of packets
    - SYN flood, HTTP flood, etc.

Resource consumption of service…

- What resources might a DoS attack try to limit?

Resource consumption of service…

- What resources might a DoS attack try to limit?
- Server CPU / memory resources
  - Consume connection state (e.g., SYN flood, HTTP flood)
  - Forces new connections to be dropped and existing connections to time-out
- Network resources
  - Attack the *router* instead of the server, find a bottleneck router that you might be able to slow down
    - If attack is >>> forwarding capacity, good data will be dropped (send big pkts)
    - If router is packet-per second limited, send lots of packets (send small + lots pkts)

Defending against network DoS

- Suppose an attacker has access to a beefy system with high-speed Internet access
  - They pump out packets towards the DoS target at a very high rate
- How might the target defend themselves?

Defending against network DoS

- Suppose an attacker has access to a beefy system with high-speed Internet access
  - They pump out packets towards the DoS target at a very high rate
- How might the target defend themselves?
  - Install a network filter to discard any packets that come from the attacker IP address!
    - E.g., `drop * 66.31.1.37:* -> *:*`

Defending against network DoS

- Suppose an attacker has access to a beefy system with high-speed Internet access
  - They pump out packets towards the DoS target at a very high rate
- How might the target defend themselves?
  - Install a network filter to discard any packets that come from the attacker IP address!
    - E.g., `drop * 66.31.1.37:* -> *:*`
- What can an attacker do to bypass the filter?

Evading naive filtering

- Make traffic appear as though it's coming from many hosts
  - E.g., IP spoofing; pick a random 32-bit source IP address for each packet sent
  - Very hard to defend against.
- Use **many** hosts to send traffic rather than just one
  - Distributed Denial-of-Service = **DDoS** ("dee-doss")
  - Requires defender to install complex filters that only kind of work
  - Today, hosts are cheap to compromise…. *botnets*
    - Underlying issue: broad insecurity (CIA) on the Internet

What are the resources of the attacker?

- Attackers need *bandwidth* to successfully launch attacks
  - At least as much sending capacity as the **bottleneck link** of the target's Internet connection
- How do you get more bandwidth…?
  - Compromise devices into a *botnet* (DDoS)
  - Partner with someone else who has bandwidth
  - …or, you could *reflect and amplify* your attacks through other computers

Reflection & Amplification

- Rather than send your attack packets directly… you can have open servers *reflect* your attack via IP spoofing

Reflection & Amplification

- Rather than send your attack packets directly… you can have open servers *reflect* your attack via IP spoofing

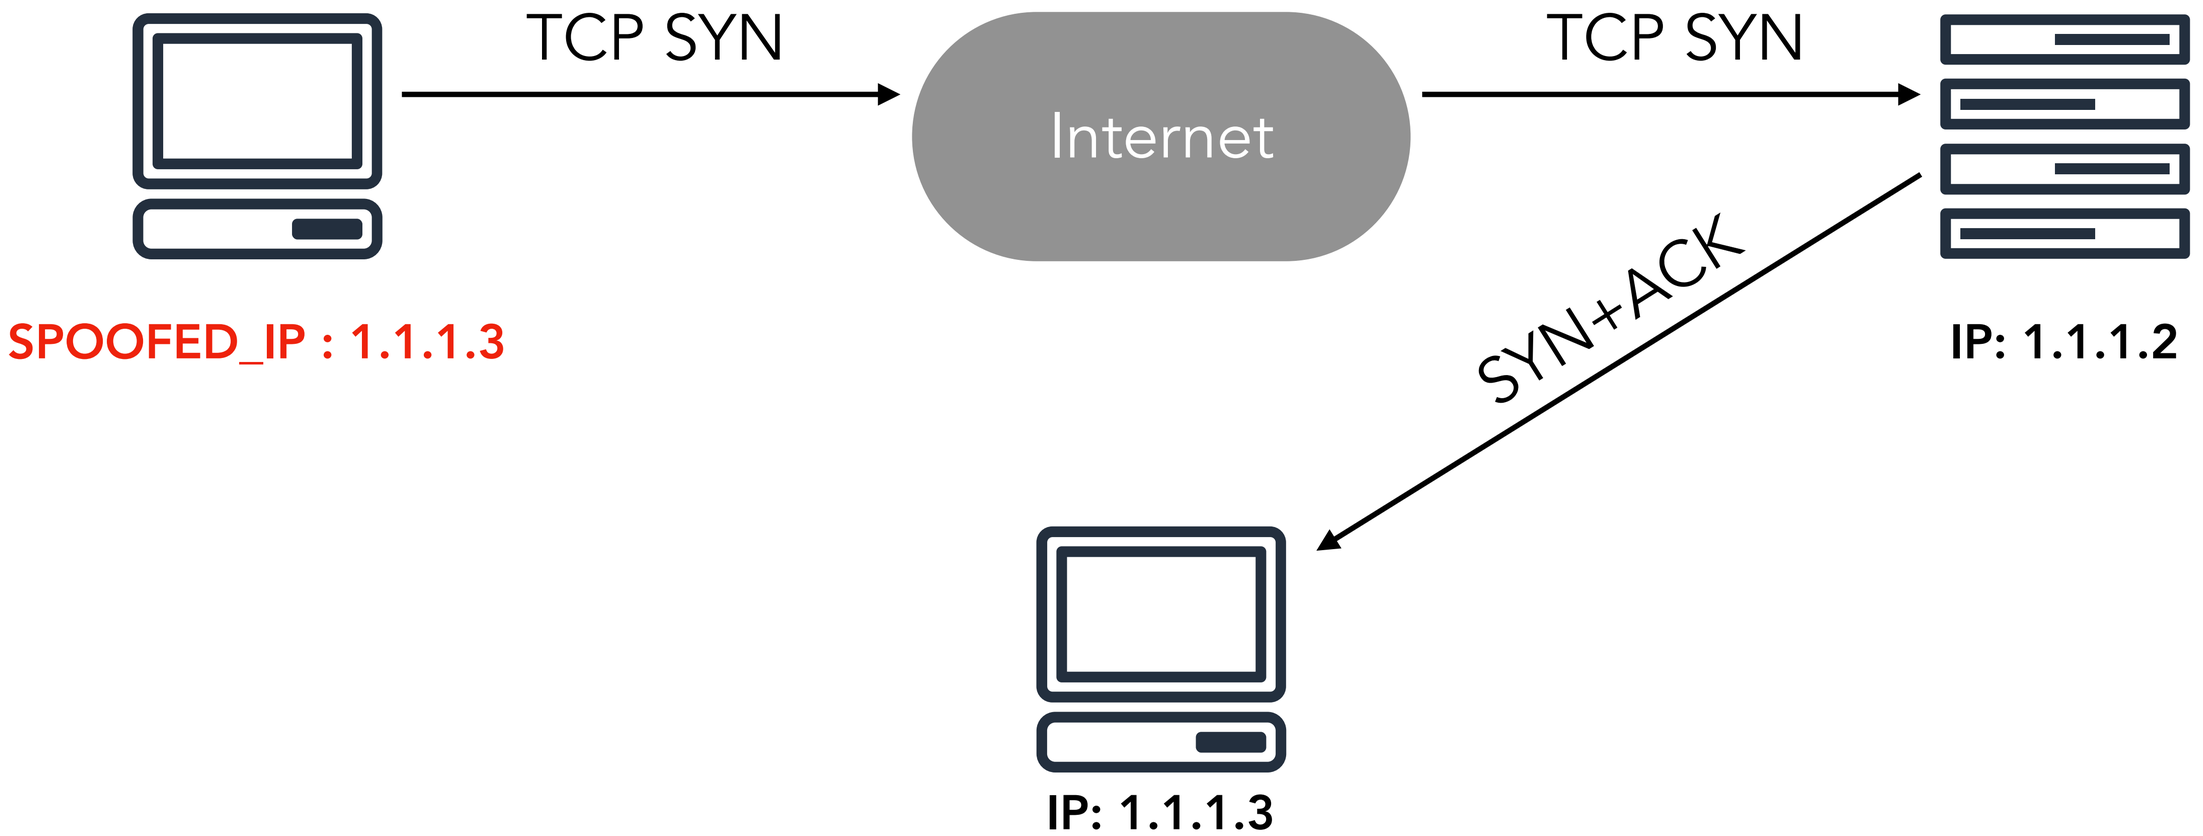

Amplification

Attacker leverages IP spoofing + asymmetric request/response protocol to amplify the load they induce on a resource.

Amplification

Attacker leverages IP spoofing + asymmetric request/response protocol to amplify the load they induce on a resource.

SPOOFED_IP : 1.1.1.3

Victim IP: 1.1.1.3

Amplification

Attacker leverages IP spoofing + asymmetric request/response protocol to amplify the load they induce on a resource.

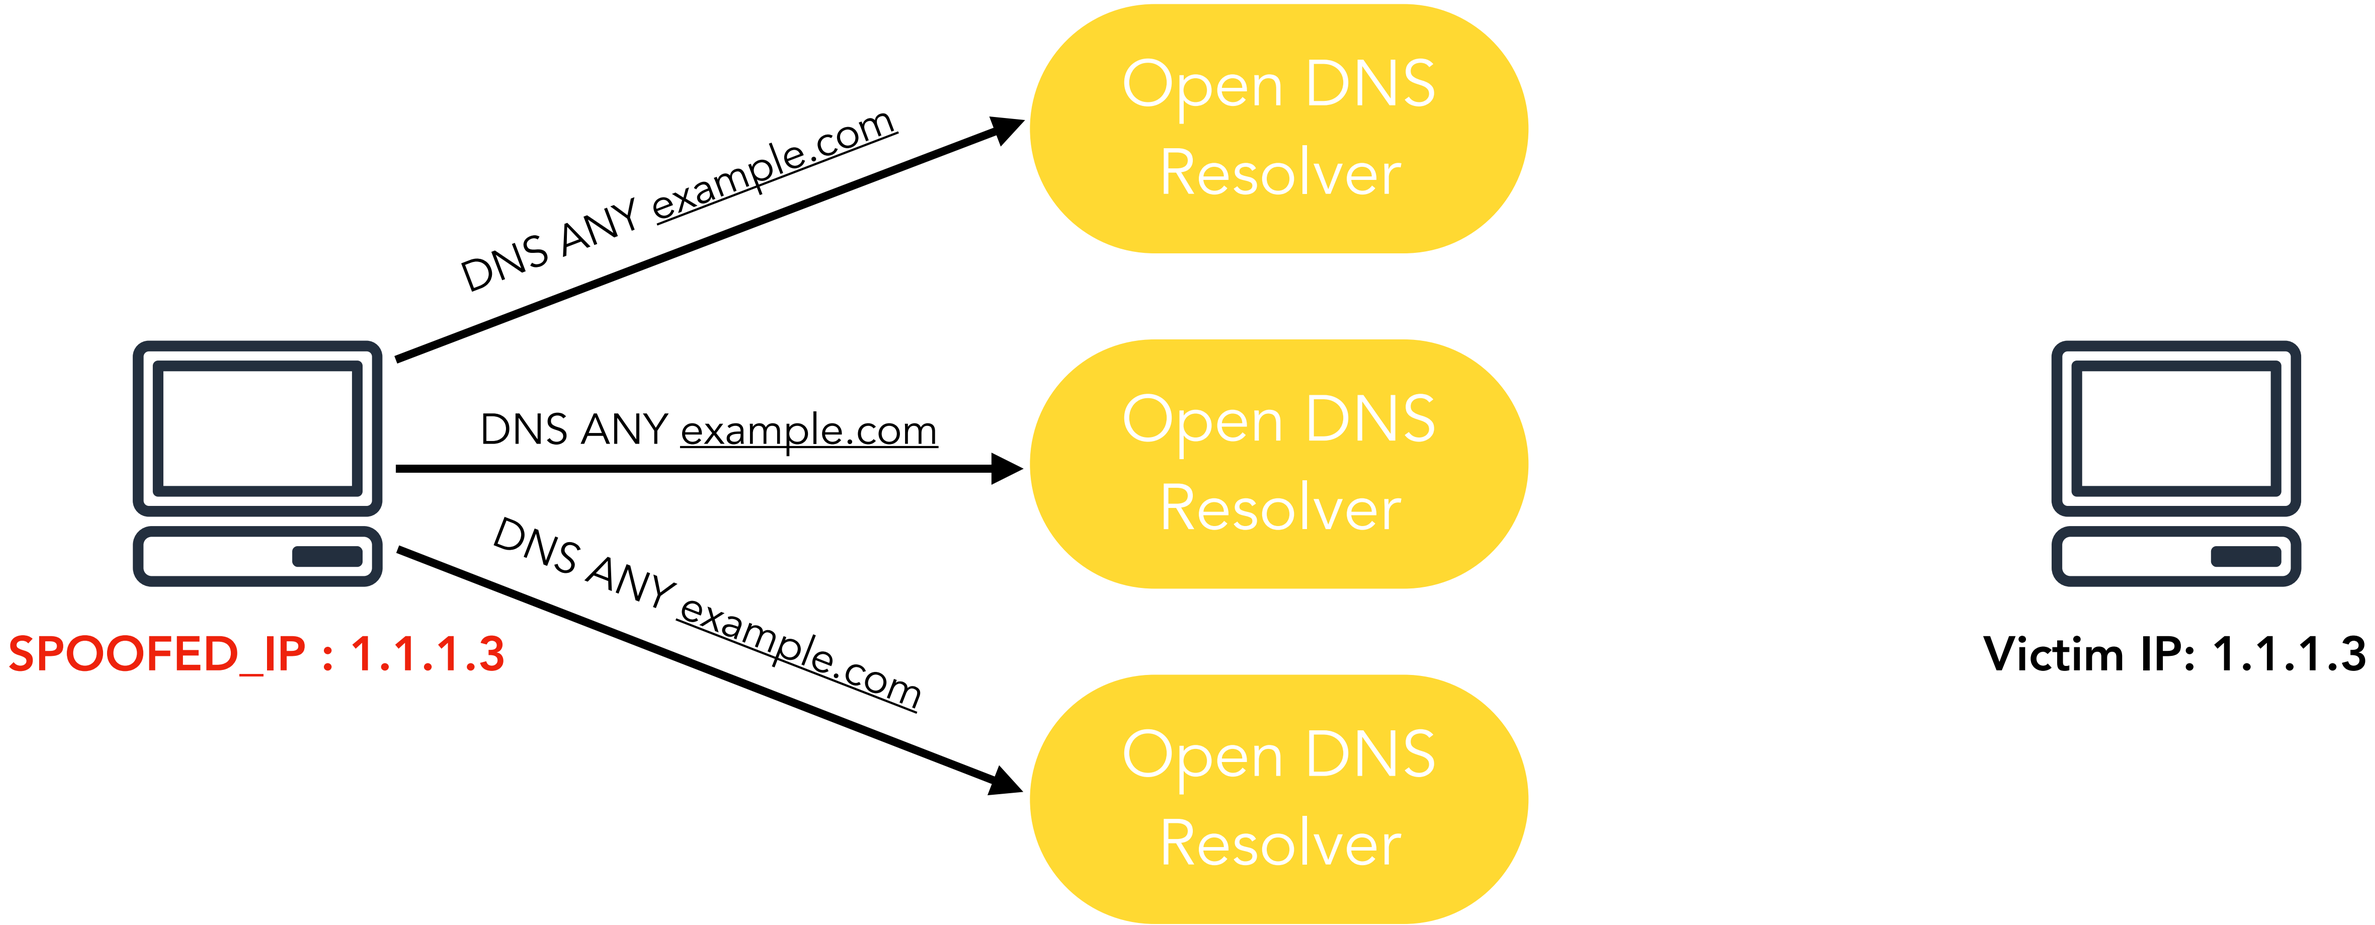

Amplification

Attacker leverages IP spoofing + asymmetric request/response protocol to amplify the load they induce on a resource.

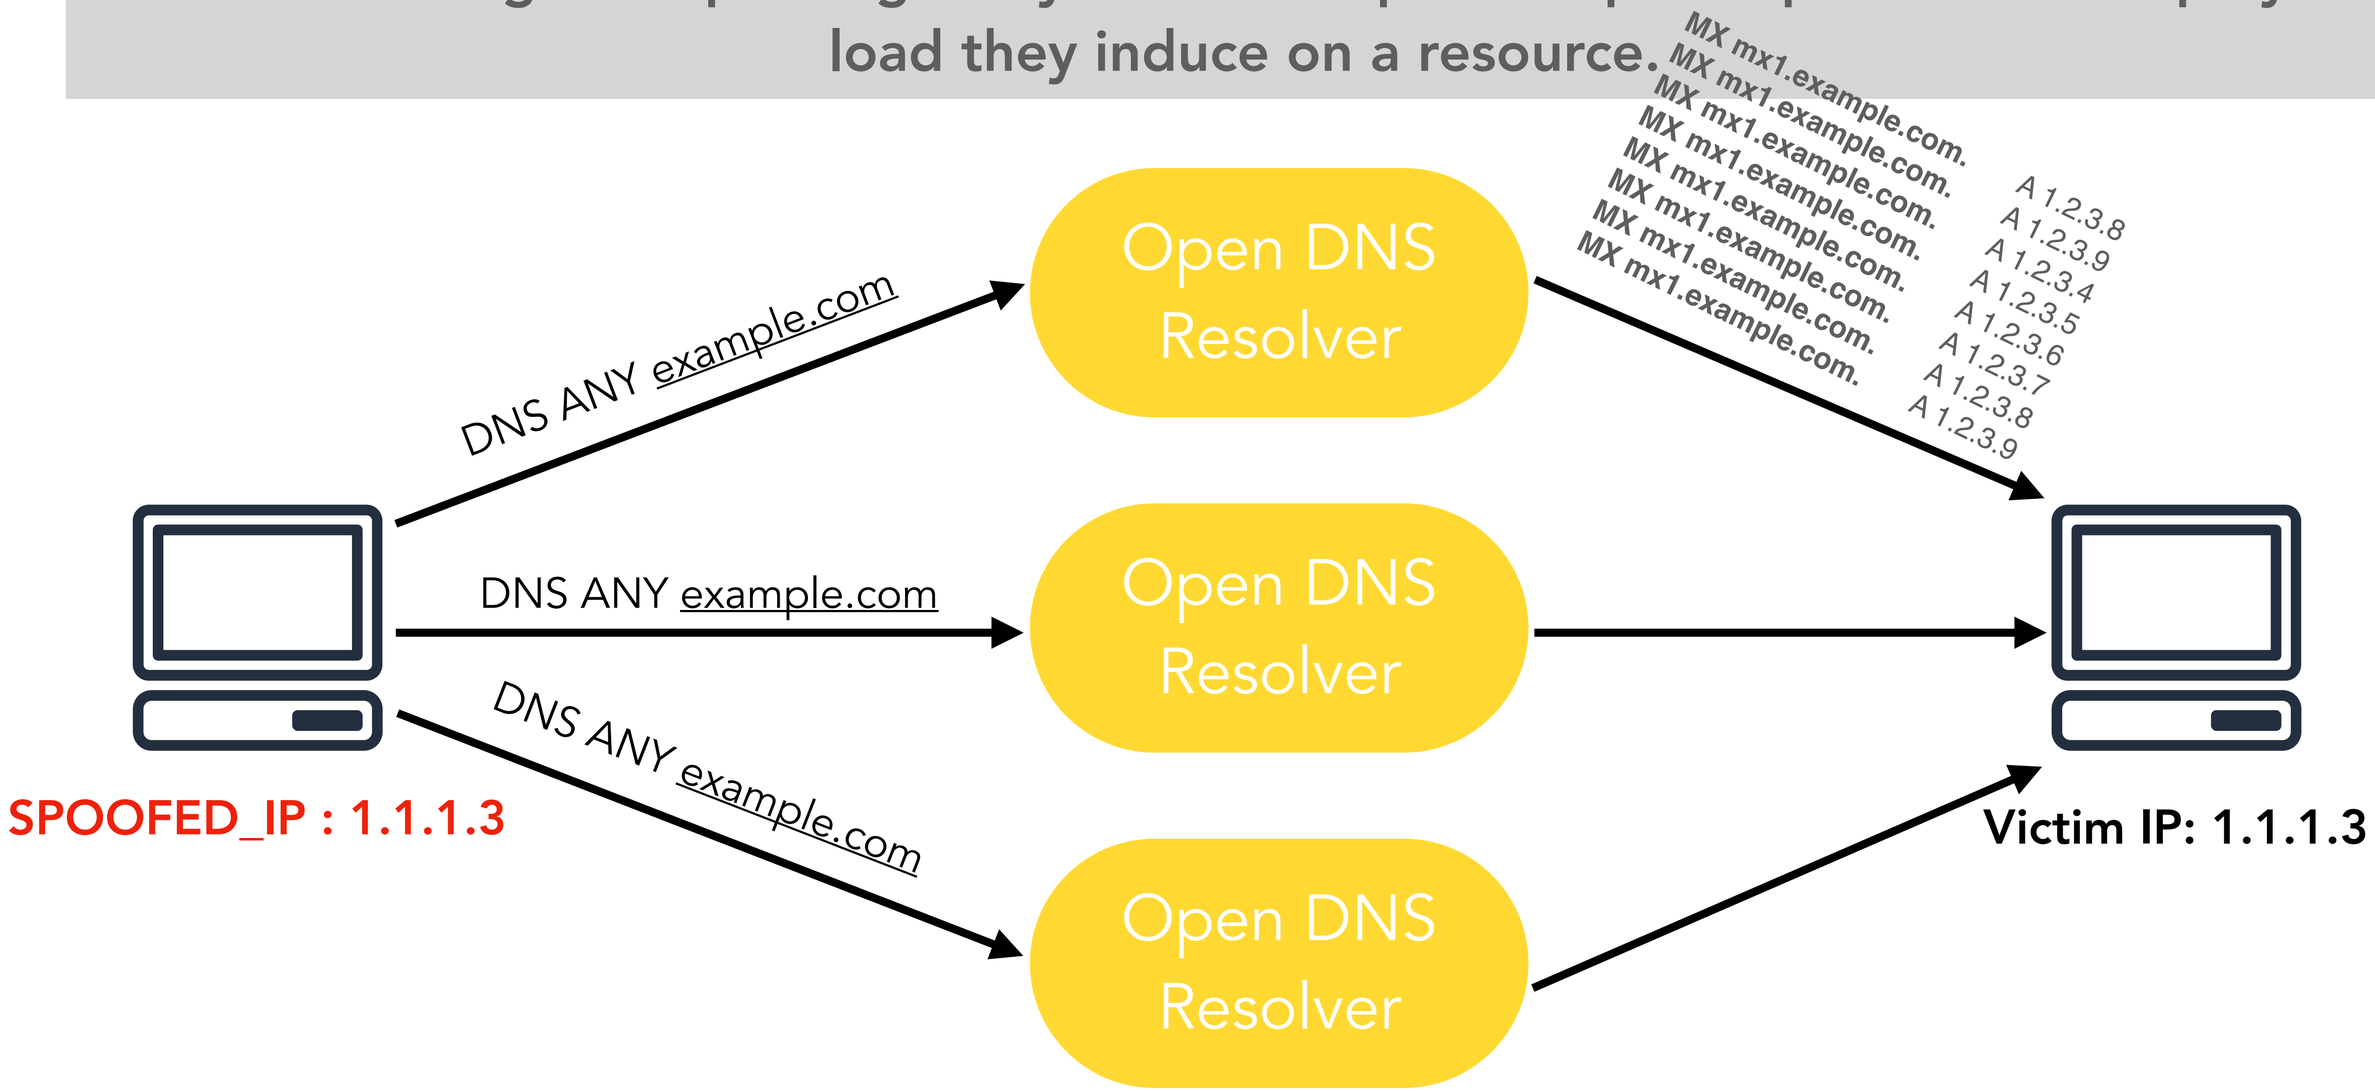

Amplification

Attacker leverages IP spoofing + asymmetric request/response protocol to amplify the load they induce on a resource.

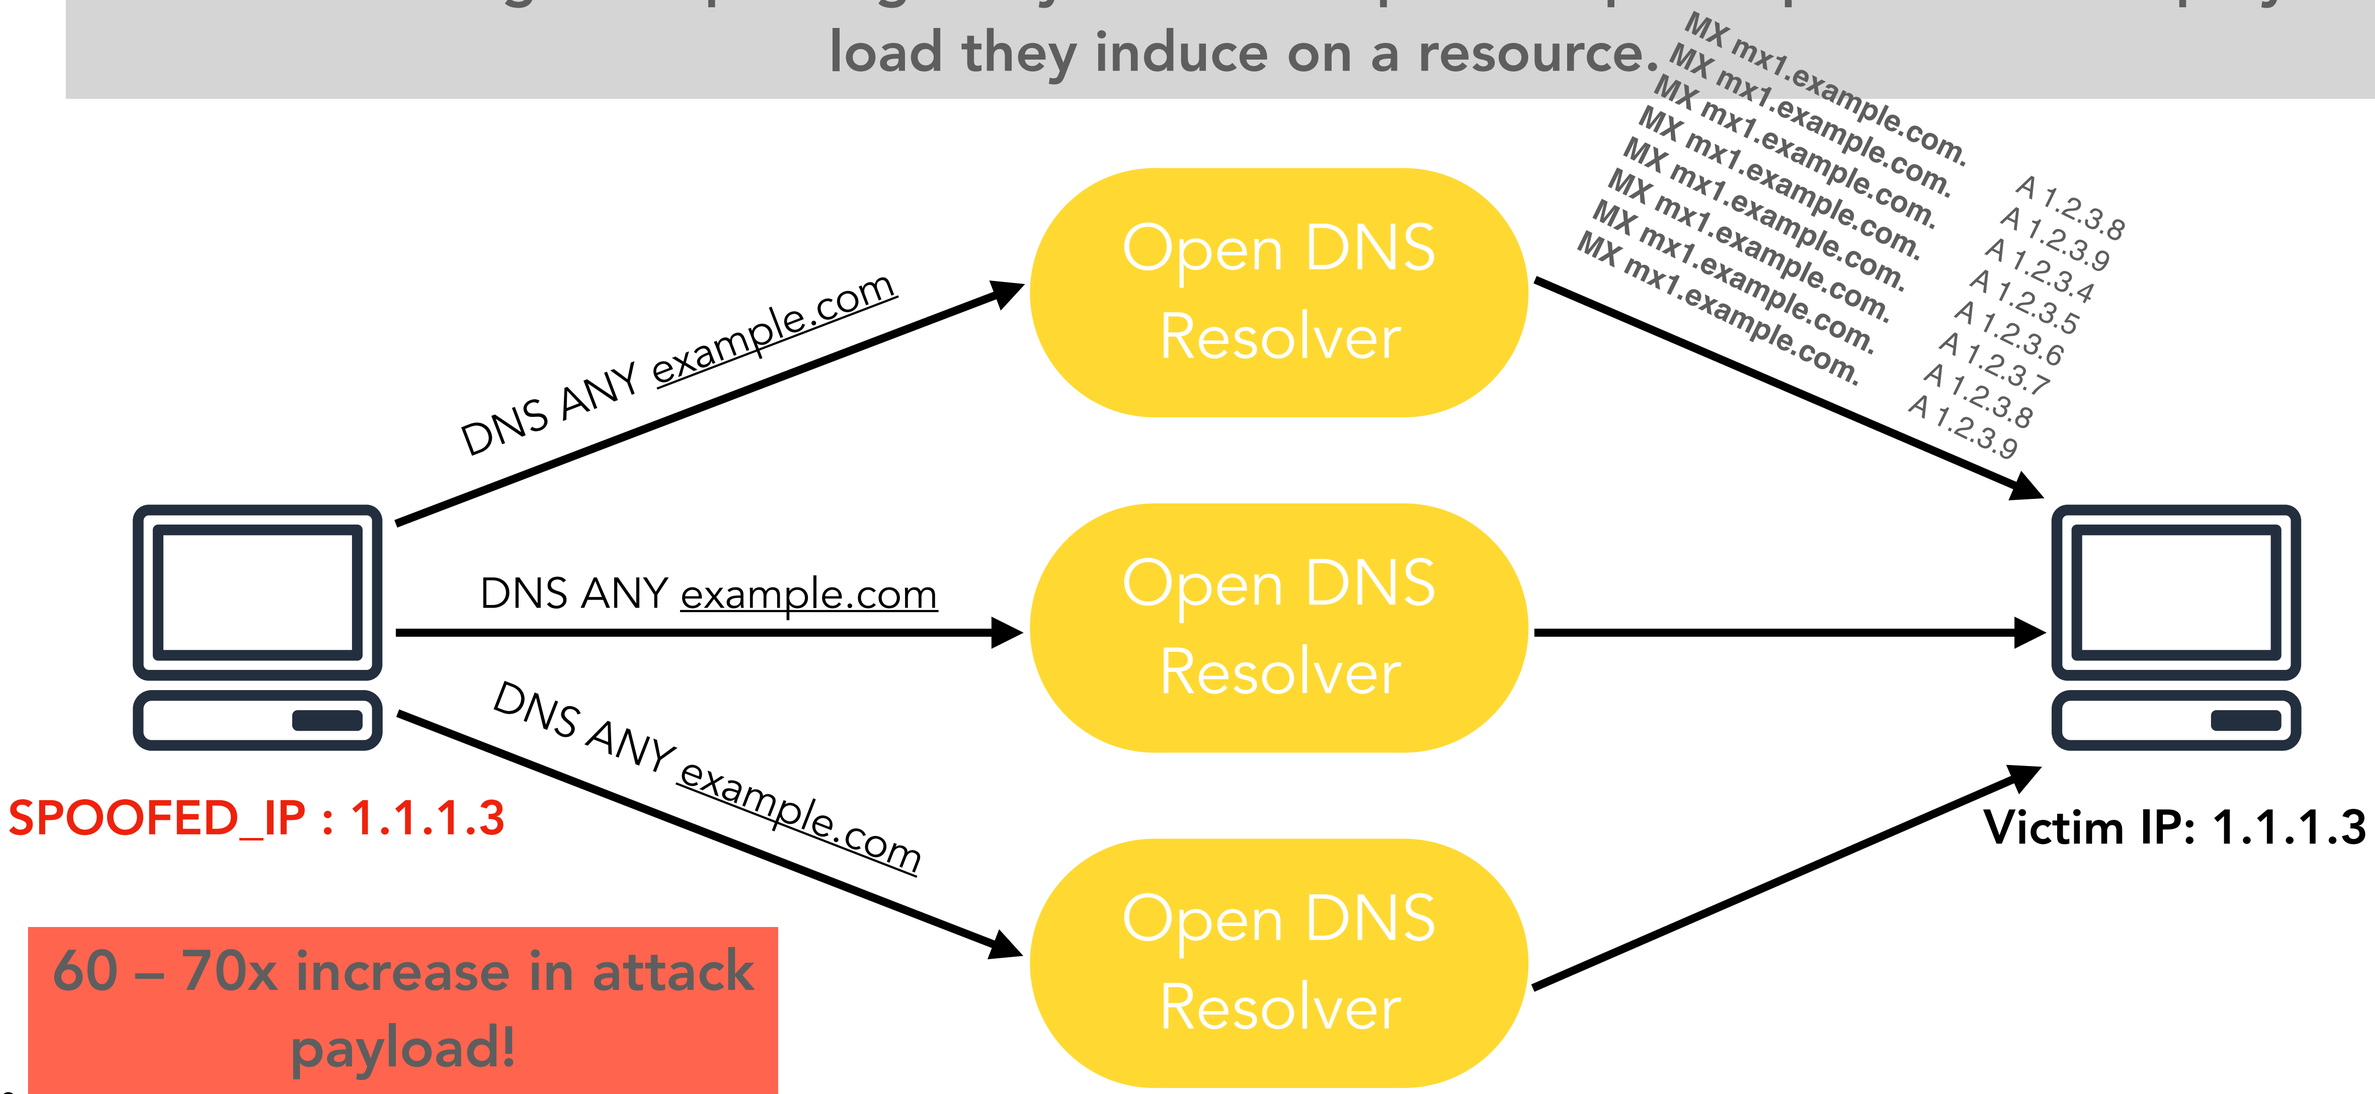

60 – 70x increase in attack payload!

Amplifiers are everywhere…

- Amplifier = any server on the internet that responds to arbitrary requests, where response size / packet count > request size / packet count

Protocols Analyzed

Cat	Protocol	Port(s)	Description
Network Svc	SNMP v2	161	Monitoring network-attached devices
	NTP	123	Time synchronization
	DNS	53	(Primarily) Domain name resolution
	NetBios	137	Name service protocol of NetBios API
	SSDP	1900	Discovery of UPnP-enabled hosts
Leg.	CharGen	19	Legacy character generation protocol
	QOTD	17	Legacy "Quote-of-the-day" protocol
P2P	BitTorrent	any	BitTorrent's Kademlia DHT impl.
	Kad	any	eMule's Kademlia DHT impl.
Gam	Quake 3	27960	Games using the Quake 3 engine
	Steam	27015	Games using the Steam protocol
Bots	ZAv2	164XY	P2P-based rootkit
	Sality	any	P2P-based malware dropper
	Gameover	any	P2P-based banking trojan

Amplifiers Found

Protocol	Amplifiers	Tech.	t_{1k}	t_{100k}
SNMP v2	4,832,000	Scan	1.5s	148.9s
NTP	1,451,000	Scan	2.0s	195.1s
DNS_{NS}	255,819	Crawl	35.3s	3530.0s
DNS_{OR}	7,782,000	Scan	0.9s	92.5s
NetBios	2,108,000	Scan	3.4s	341.5s
SSDP	3,704,000	Scan	1.9s	193.5s
CharGen	89,000	Scan	80.6s	n/a
QOTD	32,000	Scan	228.2s	n/a
BitTorrent	5,066,635	Crawl	0.9s	63.6s
Kad	232,012	Crawl	0.9s	108.0s
Quake 3	1,059	Master	0.6s	n/a
Steam	167,886	Master	1.3s	137.1s
ZAv2	27,939	Crawl	1.5s	n/a
Sality	12,714	Crawl	4.7s	n/a
Gameover	2,023	Crawl	168.5s	n/a

Christian Rossow. "Amplification Hell: Revisiting Network Protocols for DDoS Abuse." NDSS 2014

Case study: The Mirai Botnet

- I was a first year PhD student sitting in a special topics (e.g., CSE 291) security class in my first semester, when I hear Brian Krebs' website is down due to DDoS
- Krebs writes out that it's a huge attack, 620 Gbps (at the time was very large)
- Claim is that it's powered by weak IoT devices (unverified)

Story time…

- Folks in the lab think it's interesting, we want to investigate it
- Then, 9 days later, **the Mirai code is released online,** Internet havoc ensues

[FREE] World's Largest Net:Mirai Botnet, Client, Echo Loader, CNC source code release

Yesterday, 12:50 PM (This post was last modified: Yesterday 04:29 PM by **Anna-senpai**.)

Anna-senpai
L33t Member
L33T

Preface

Greetz everybody,

When I first go in DDoS industry, I wasn't planning on staying in it long. I made my money, there's lots of eyes looking at IOT now, so it
However, I know every skid and their mama, it's their wet dream to have something besides qbot.

So today, I have an amazing release for you. With Mirai, I usually pull max 380k bots from telnet alone. However, after the Kreb DDoS,
shutting down and cleaning up their act. Today, max pull is about 300k bots, and dropping.

So, I am your senpai, and I will treat you real nice, my hf-chan.

Story time…

- A few weeks later, much of the entire Internet was **down** for at least a few hours due to a Mirai DDoS attack on Dyn (dynamic DNS provider)

We got curious!

- What is going on with the Mirai botnet?
  - How do we measure the growth, size, spread, and impact of Mirai?
- What devices enabled Mirai's power, and what was their security posture like?
  - Mirai was enabled by **weak passwords —** IoT devices on the public Internet with default credentials on services that allow shell access: telnet, SSH, CWMP, etc.

How does Mirai work?

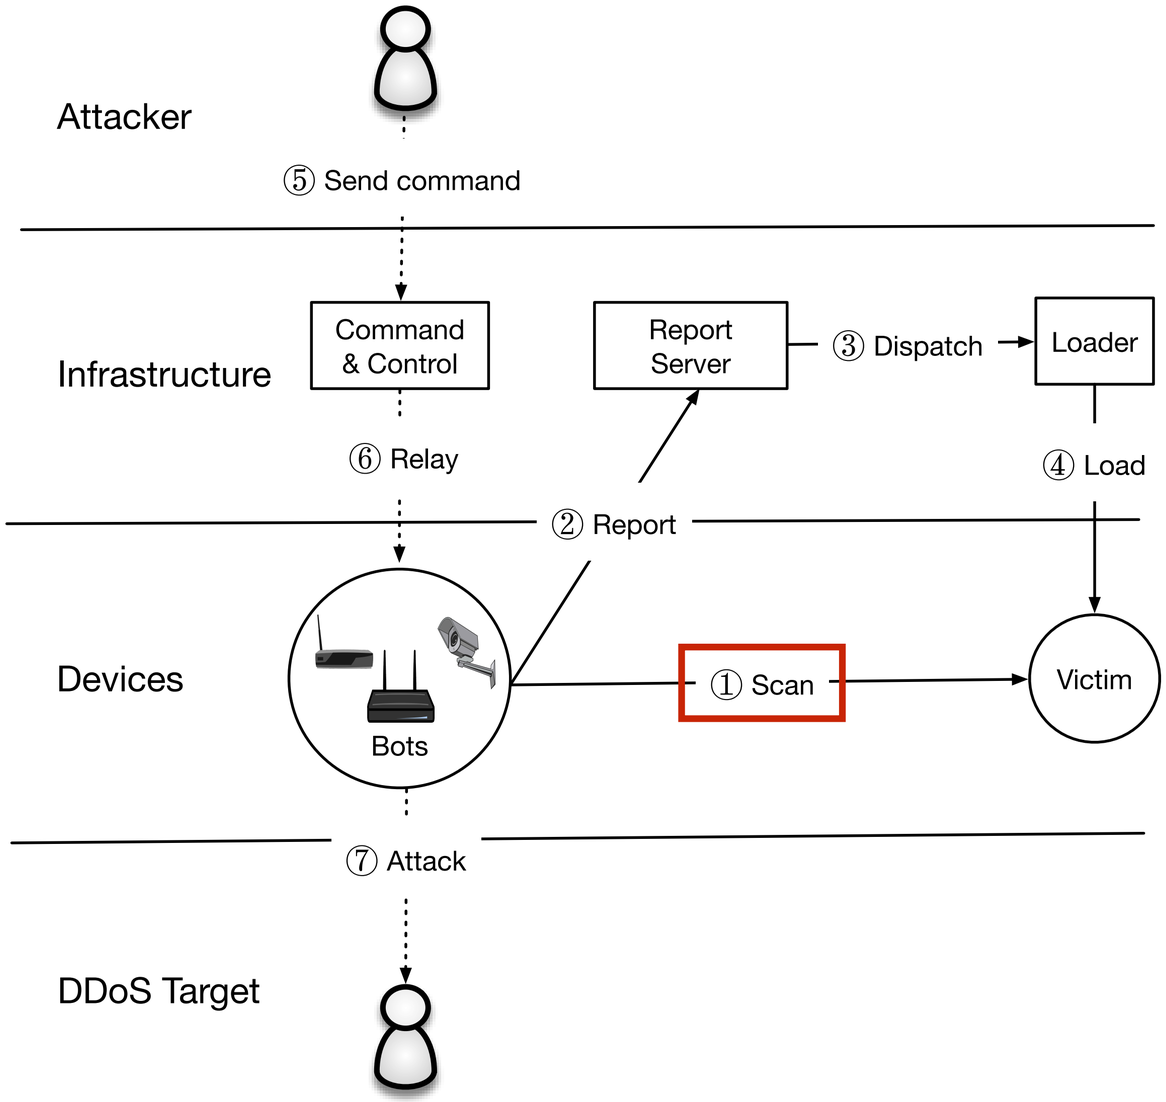

1. SYN Scan to find IPs with open ports
2. Vulnerability scan to test default passwords

How does Mirai work?

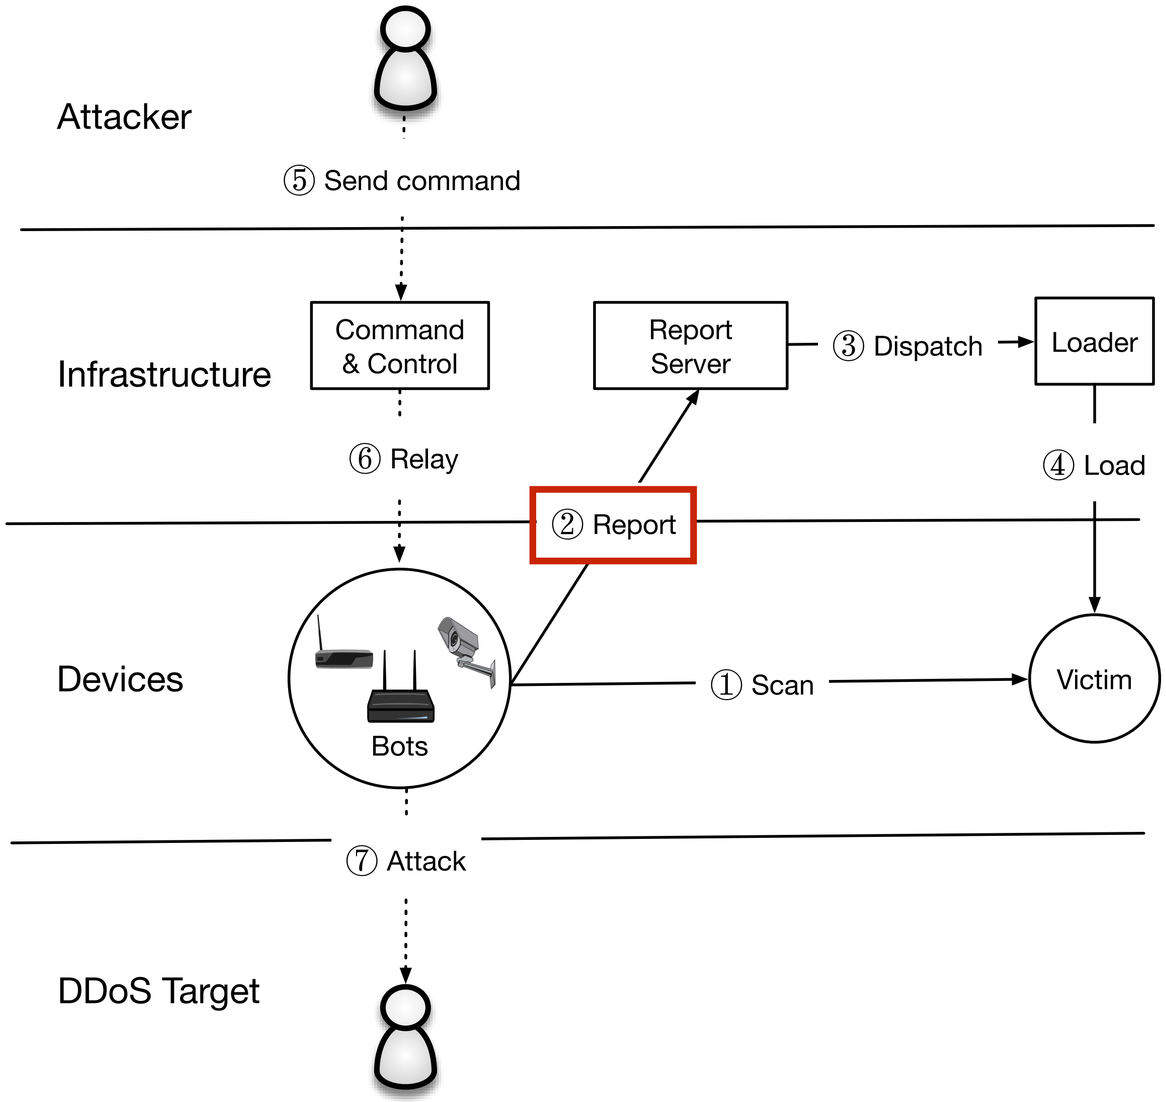

Reports to a server telling it which IPs have vulnerable devices

How does Mirai work?

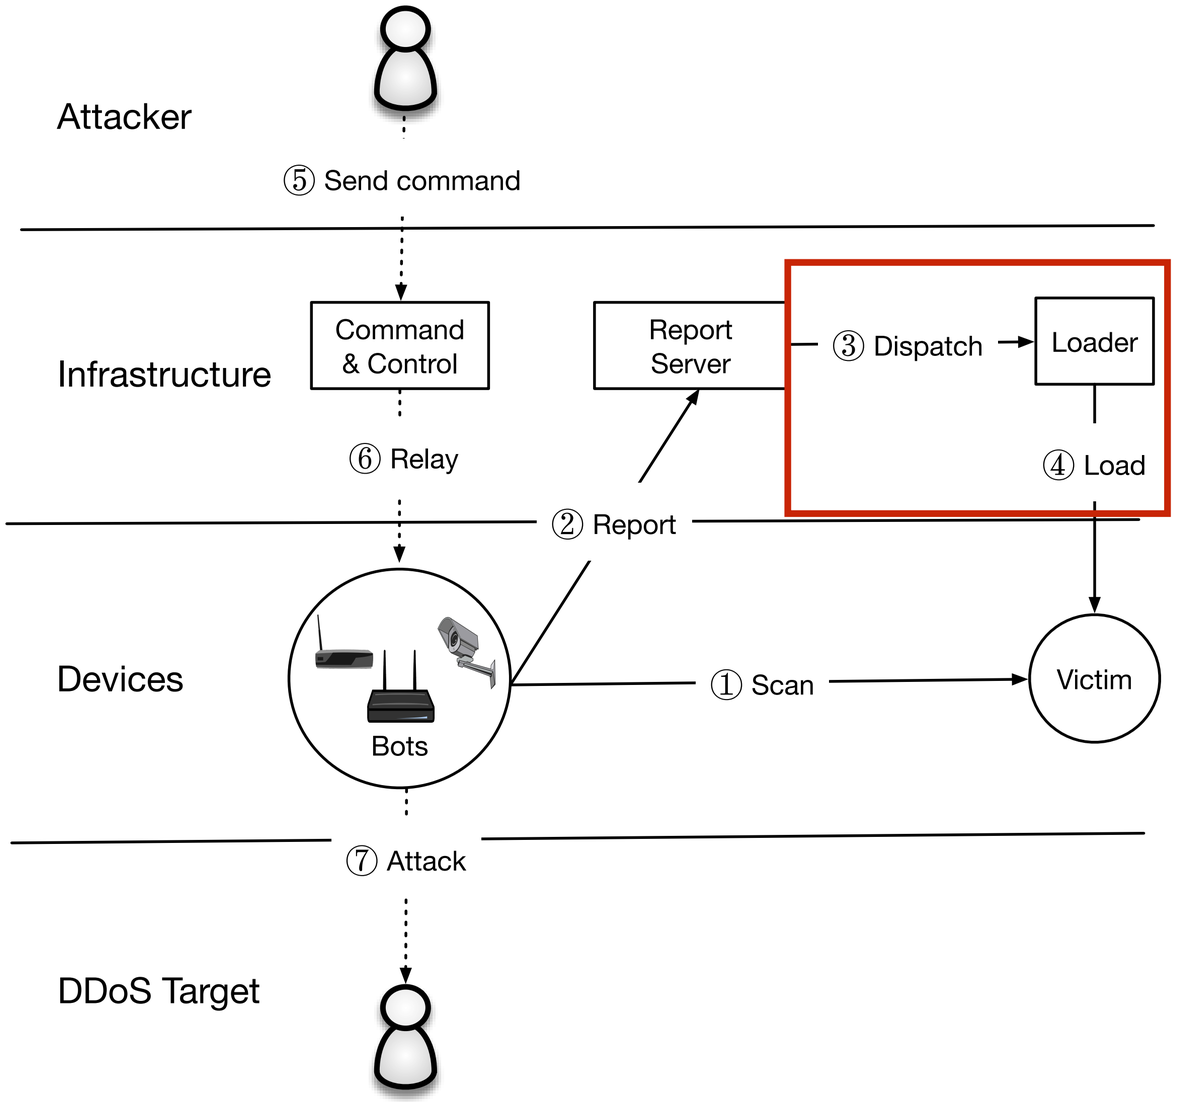

Server then dispatches another machine (usually a compromised bot) to go and download and run malware

How does Mirai work?

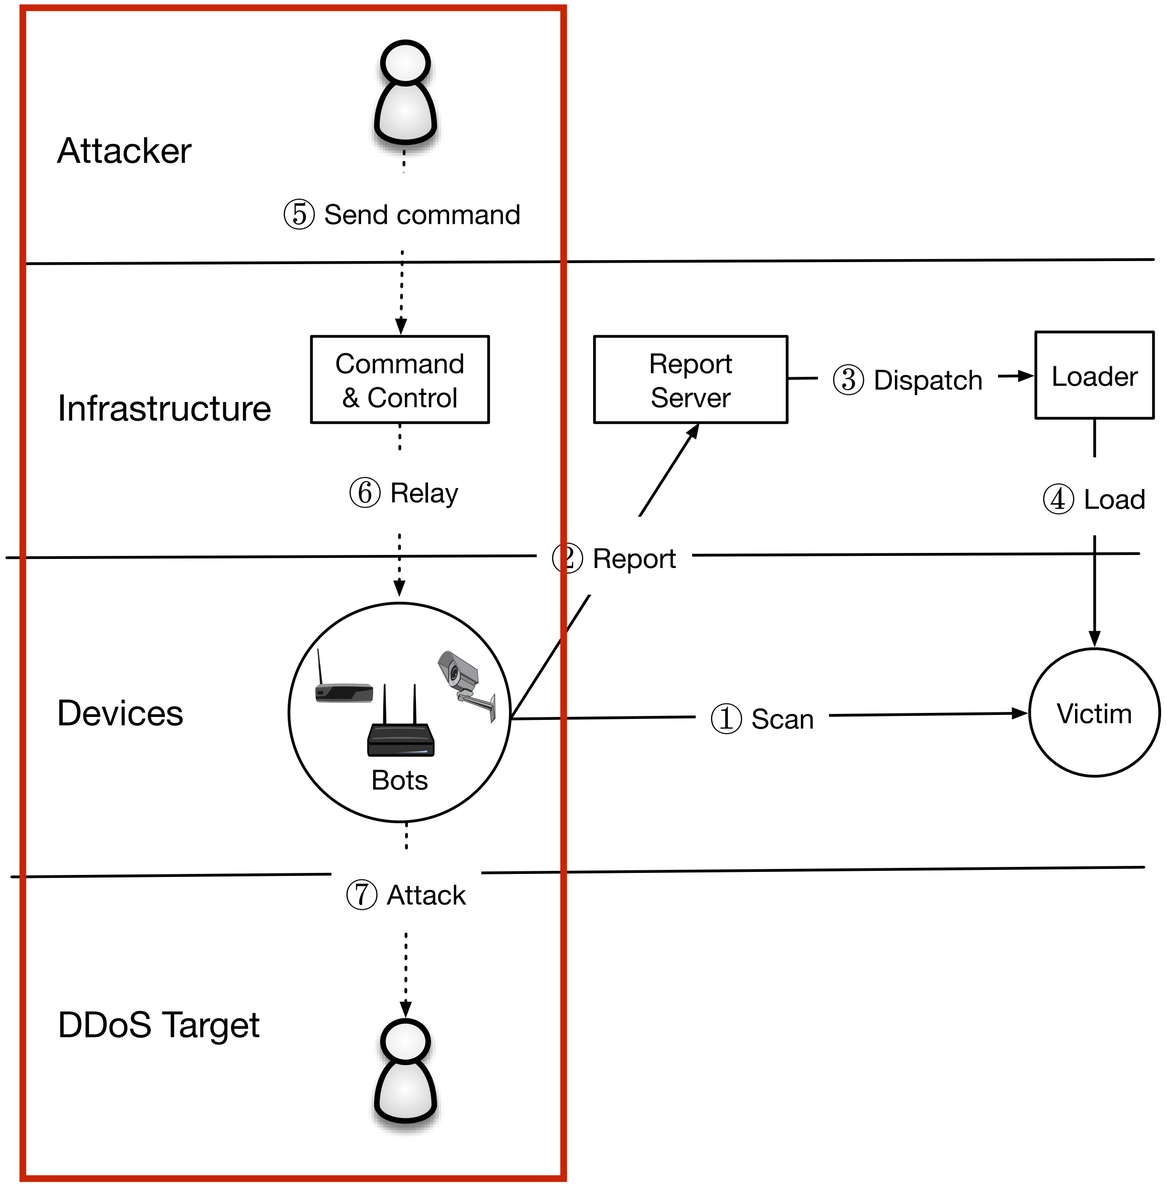

In attack phase, attacker sends commands through a C2 server, which relays to compromised machines, which all target a victim at once!

How did we study Mirai?

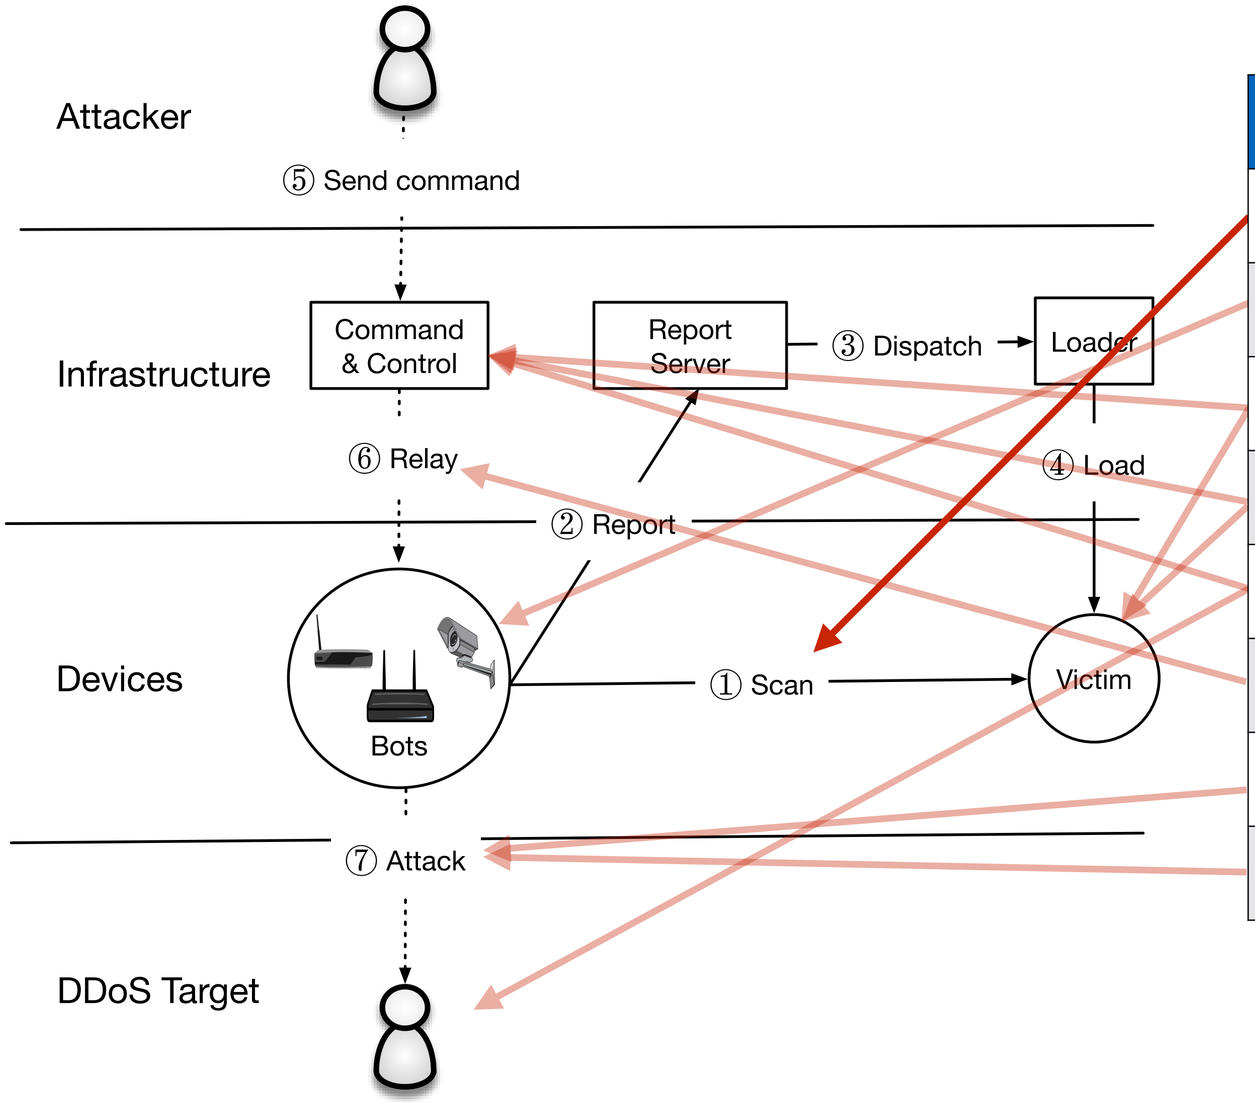

Data Source	Size
Network Telescope	4.7M unused IPs
Active Scanning	136 IPv4 scans
Telnet Honeypots	434 binaries
Malware Repository	594 binaries
Active/Passive DNS	499M daily RRs
C2 Milkers	64K issued attacks
Krebs DDoS Attack	170K attacker IPs
Dyn DDoS Attack	108K attacker IPS

July 2016 - February 2017

Understanding the Dyn attack

The New York Times

"It is possible, investigators say, that the attack on Dyn was conducted by a criminal group that wanted to extort the company. Or it could have been done by "hacktivists." Or a foreign power that wanted to remind the United States of its vulnerability."

Understanding the Dyn attack

The New York Times

> "It is possible, investigators say, that the attack on Dyn was conducted by a criminal group that wanted to extort the company. Or it could have been done by "hacktivists." Or a foreign power that wanted to remind the United States of its vulnerability."

Targeted IP	rDNS	Passive DNS
208.78.70.5	ns1.p05.dynect.net	ns00.playstation.net
204.13.250.5	ns2.p05.dynect.net	ns01.playstation.net
208.78.71.5	ns3.p05.dynect.net	ns02.playstation.net
204.13.251.5	ns4.p05.dynect.net	ns03.playstation.net

- Top targets are linked to Sony PlayStation
- Dyn just happened to be in the same IP block as PSN, **collateral damage**

Myriad of Targets

- **Games:** Minecraft, Runescape, etc.
- **Politics:** Chinese political dissidents, regional Italian politician
- **Anti-DDoS:** DDoS protection services
- **Misc:** Russian cooking blog….?

What happened next?

Who is Anna-Senpai, the Mirai Worm Author?

January 18, 2017 248 Comments

On September 22, 2016, this site was forced offline for nearly four days after it was hit with "**Mirai**," a malware strain that enslaves poorly secured Internet of Things (IoT) devices like wireless routers and security cameras into a botnet for use in large cyberattacks. Roughly a week after that assault, the individual(s) who launched that attack — using the name "**Anna-Senpai**" — released the source code for Mirai, spawning dozens of copycat attack armies online.

After months of digging, KrebsOnSecurity is now confident to have uncovered Anna-Senpai's real-life identity, and the identity of at least one co-conspirator who helped to write and modify the malware.

Mirai co-author Anna-Senpai leaked the source code for Mirai on Sept. 30, 2016.

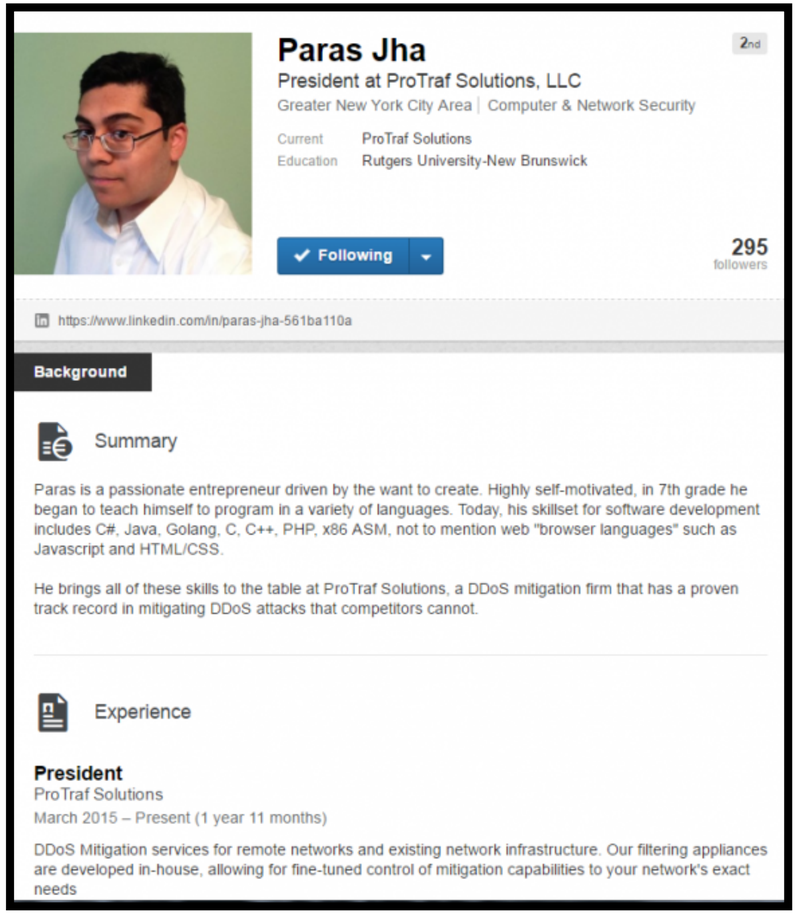

https://krebsonsecurity.com/2017/01/who-is-anna-senpai-the-mirai-worm-author/

What happened next?

The Mirai Confessions: Three Young Hackers Who Built a Web-Killing Monster Finally Tell Their Story

Netflix, Spotify, Twitter, PayPal, Slack. All down for millions of people. How a group of teen friends plunged into an underworld of cybercrime and broke the internet—then went to work for the FBI.

https://www.wired.com/story/mirai-untold-story-three-young-hackers-web-killing-monster/

What to do about DDoS…

- **Very hard.** Remains an open area of study and research.
  - Network service model allows unsolicited requests
  - Bad guys can leverage large # of resources
  - Hard to attribute network actions
- Cat and mouse struggle…
  - You could patch vulnerable devices (not easy)
  - You can try to remove bad amplifiers (e.g., ANY is deprecated… but others will always exist, see GitHub attack of 2018)

BCP 38 – Ingress Filtering

- bcp38.info
- Basic idea: **Ingress filtering**
  - ISPs know which IPs they assigned to each device… so why not block them before the enter the network?
  - AKA: Stop IP spoofing at the source network
- Sounds easy. Why don’t we do it?

BCP 38 – Ingress Filtering

- bcp38.info
- Basic idea: **Ingress filtering**
  - ISPs know which IPs they assigned to each device… so why not block them before the enter the network?
  - AKA: Stop IP spoofing at the source network
- Sounds easy. Why don't we do it?
  - Collective action problem —> everyone will benefit from the change, but doing so in the short term might be hard, so no one does this
  - **All ISPs need to do this –** requiring global coordination

Client Puzzles

- One issue with DDoS: **asymmetry**
  - **Asymmetry:** Attacker can consume victim resources with less effort than it takes to consume their own resources (e.g., amplification)
- Client Puzzle: What if we forced every client (e.g., request initiator) to do a moderate amount of work for every packet sent?
  - Example: Client sends puzzle C
  - Client must solve C; usually takes up some CPU resources
- Sadly, not used a lot… requires significant coordination between servers, clients, built in protocols…. and we usually don’t want to trade off performance for security :(

DDoS Protection Services

- E.g., Google Project Shield
  - https://projectshield.withgoogle.com

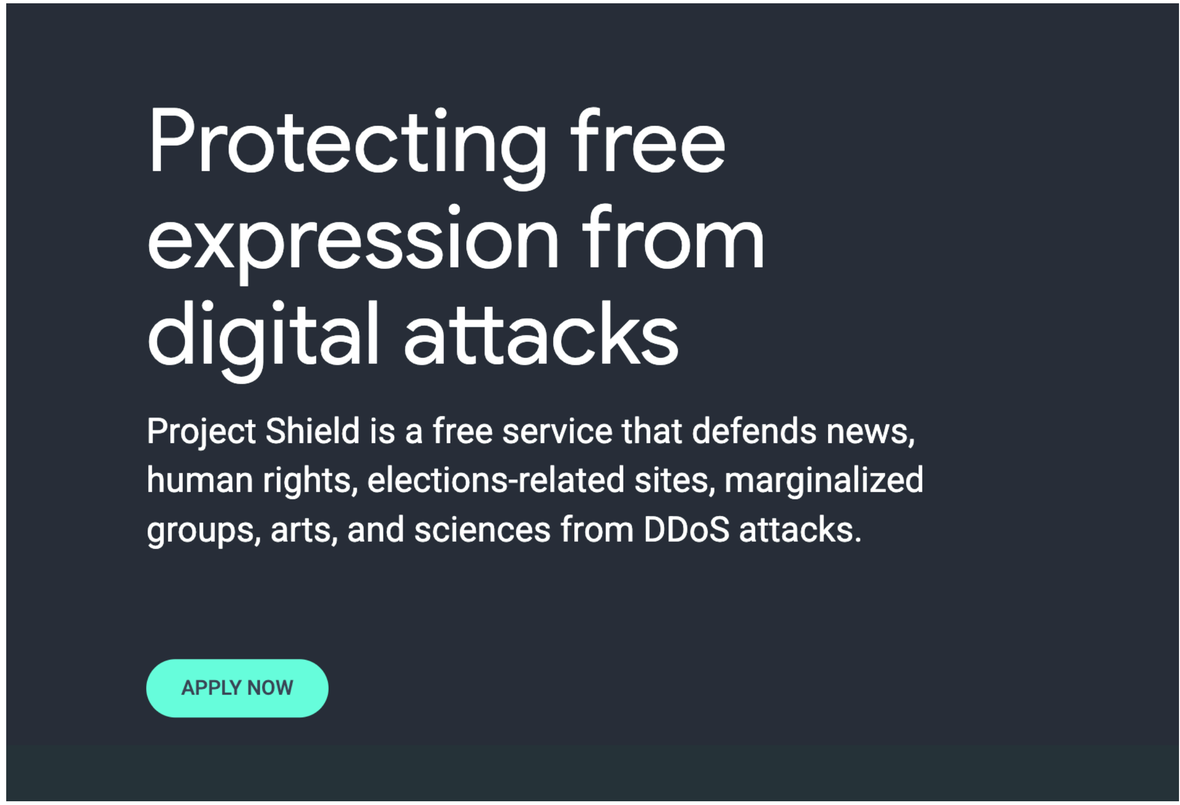

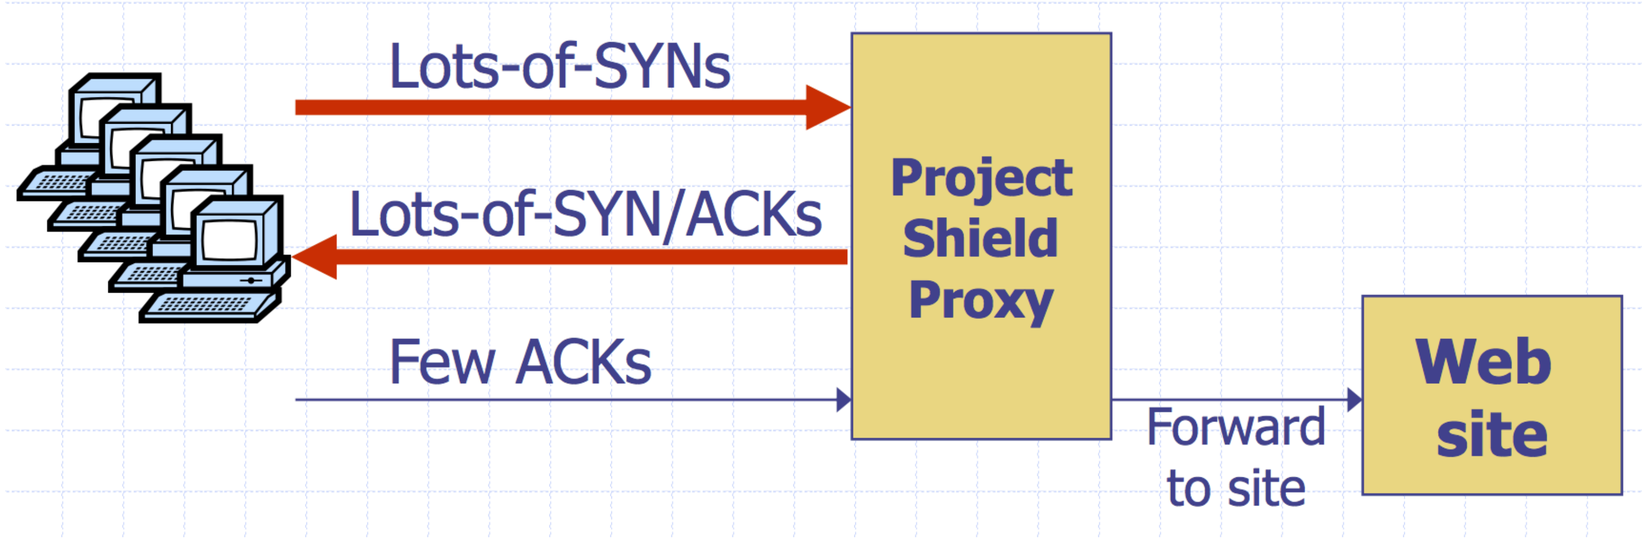

CDNs

- Today, most people get protection from DDoS via content delivery networks (CDNs), like Cloudflare or Akamai
  - These companies “serve” your content by DNS proxy + caching
  - Cloudflare has… infinite bandwidth
- Basic idea: “Scale up” resources when an attack is present
  - Goal is to outmuscle the attacker
- Cons:
  - You have to trust Cloudflare to serve and manage your content :)

Internet Censorship

What is censorship?

What is censorship?

Censorship: The suppression of words, images, ideas that are “offensive,” typically an arm for political or personal control or coercion

What is Internet censorship?

What is Internet censorship?

Internet censorship: Censorship on the Internet – and typically enacted via technical, network-level means

How Internet Censors Work

Major mechanisms for running Internet Censorship

- Internet shutdown
  - Removing Internet service altogether – much easier in some countries than others

155 political shutdowns in 29 countries in 2020

https://slate.com/technology/2020/04/pandemic-internet-shutdown-danger.htm

How Internet Censors Work

Major mechanisms for running Internet Censorship

- Internet shutdown
  - Removing Internet service altogether – much easier in some countries than others
- Throttling
  - Making certain services slower in country boundaries

Russia used throttling to slow down Twitter in March 2021, 2022

How Internet Censors Work

Major mechanisms for running Internet Censorship

- Internet shutdowns
  - Removing Internet service altogether – much easier in some countries than others
- Throttling
  - Making certain services slower in country boundaries
- Content takedowns
  - Removal of “offensive” content from online services

Google received ~550K government takedown requests since 2011

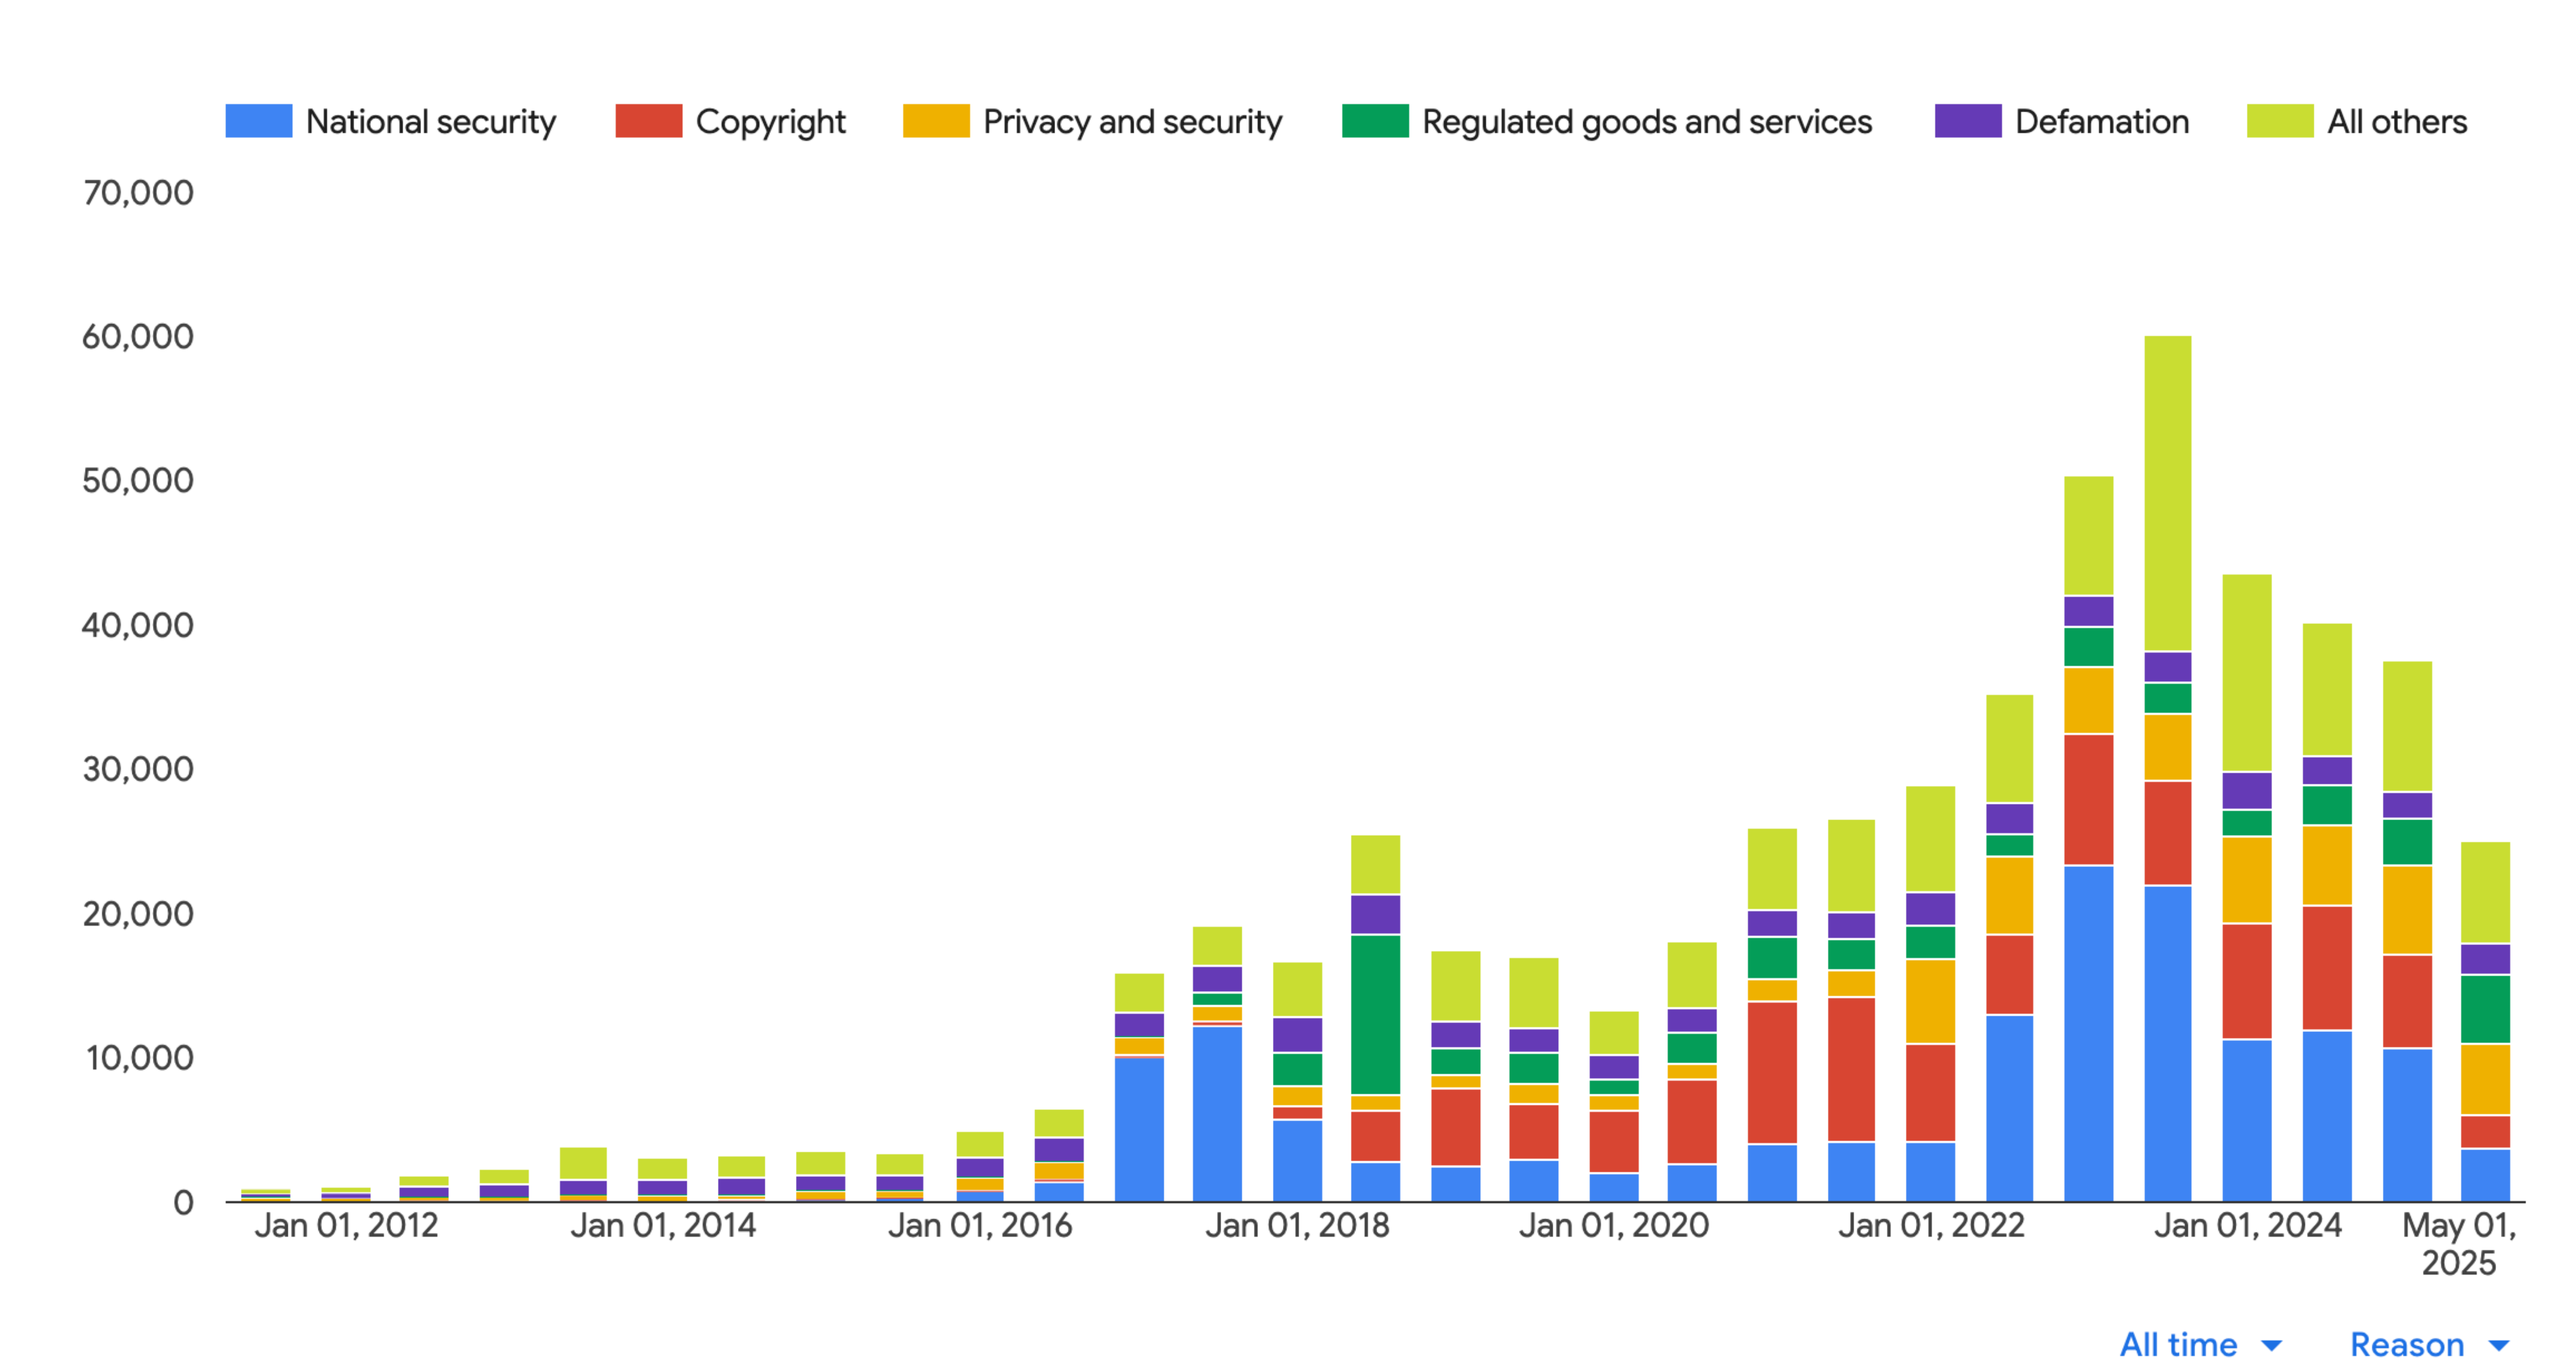

https://transparencyreport.google.com/government-removals/government-requests?hl=en

How Internet Censors Work

Major mechanisms for running Internet Censorship

- Primary form of Internet censorship: network-level blocking

- Three main ways that network-level blocking happens in practice

How Internet Censors Work

Major mechanisms for running Internet Censorship

- Primary form of Internet censorship: network-level blocking
- Three main ways that network-level blocking happens in practice
  - **DNS manipulation**

Censorship during an Internet connection

Modes of website blocking

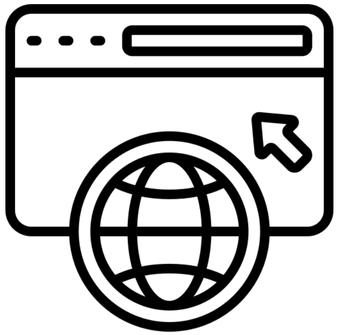

Client

kumarde.com

Censorship during an Internet connection

DNS Manipulation

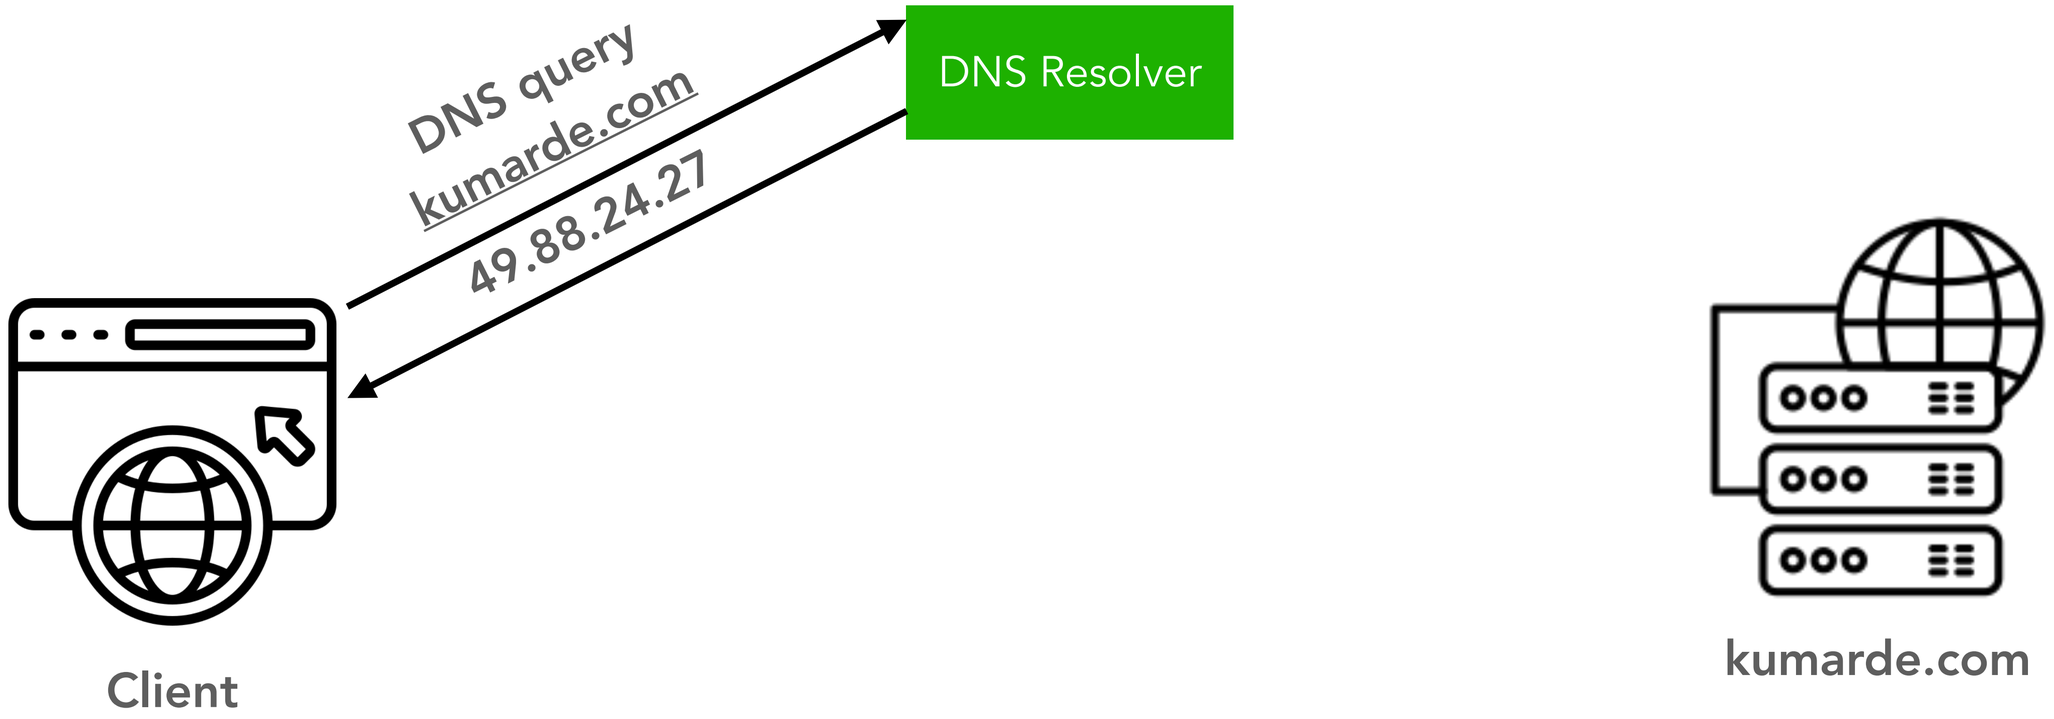

Censorship during an Internet connection

DNS Manipulation

DNS Resolver

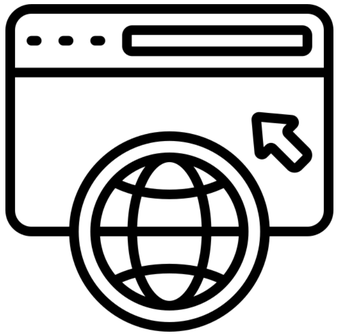

Client

kumarde.com

Censorship during an Internet connection

DNS Manipulation

Censorship during an Internet connection

DNS Manipulation

How easy do we think this is to implement?

How Internet Censors Work

Major mechanisms for running Internet Censorship

- Primary form of Internet censorship: network-level blocking
- Three main ways that network-level blocking happens in practice
  - DNS Manipulation
  - IP Blocking

Censorship during an Internet connection

IP Blocking

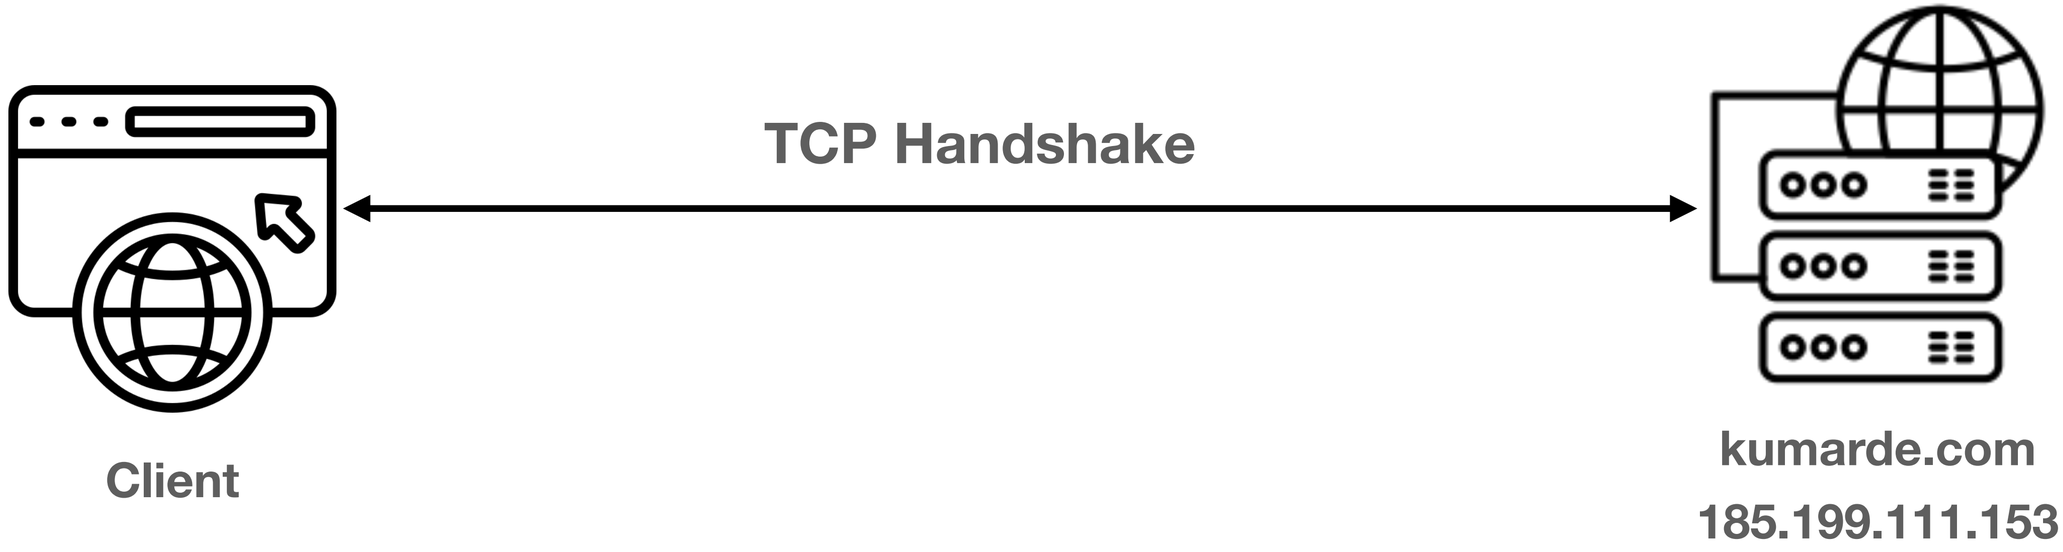

Censorship during an Internet connection

IP Blocking

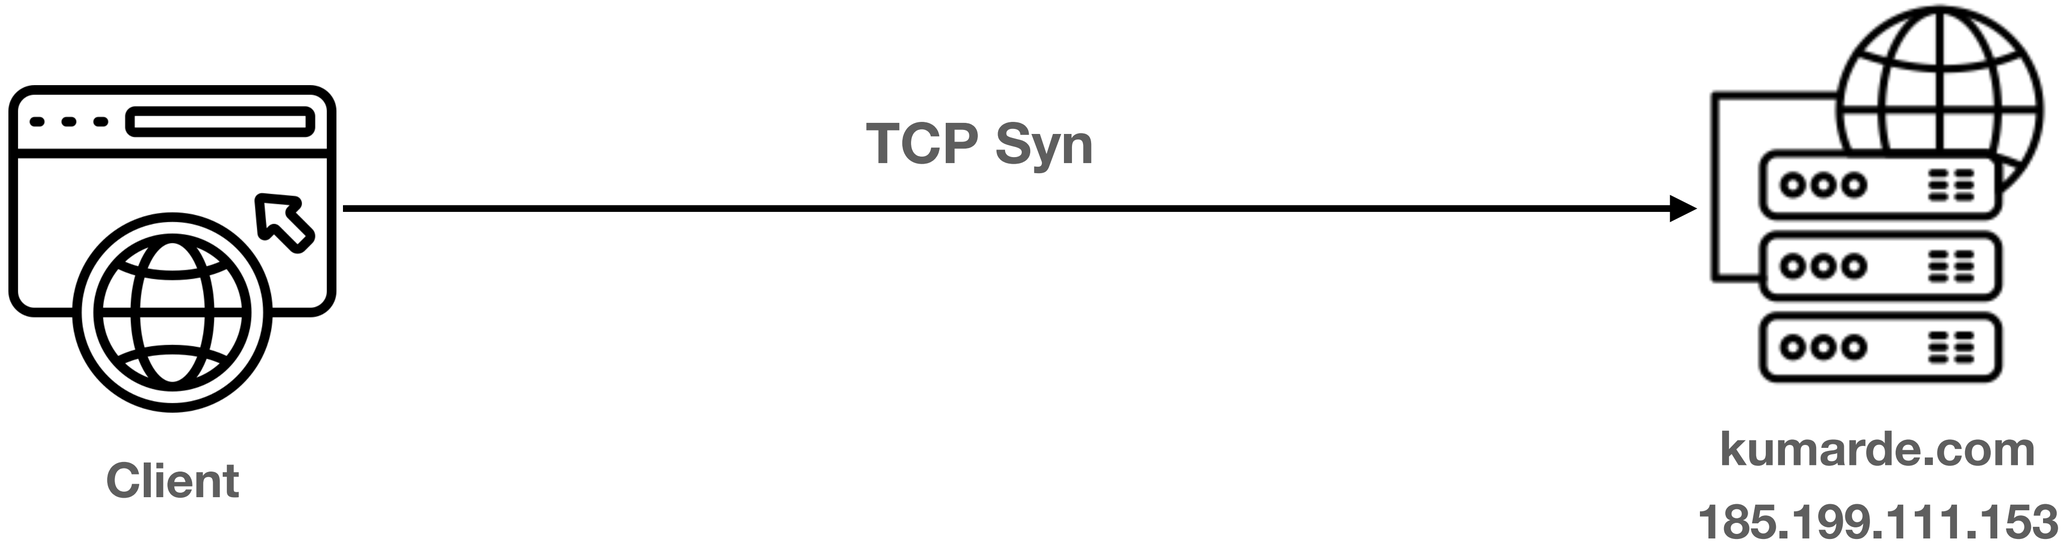

Censorship during an Internet connection

IP Blocking

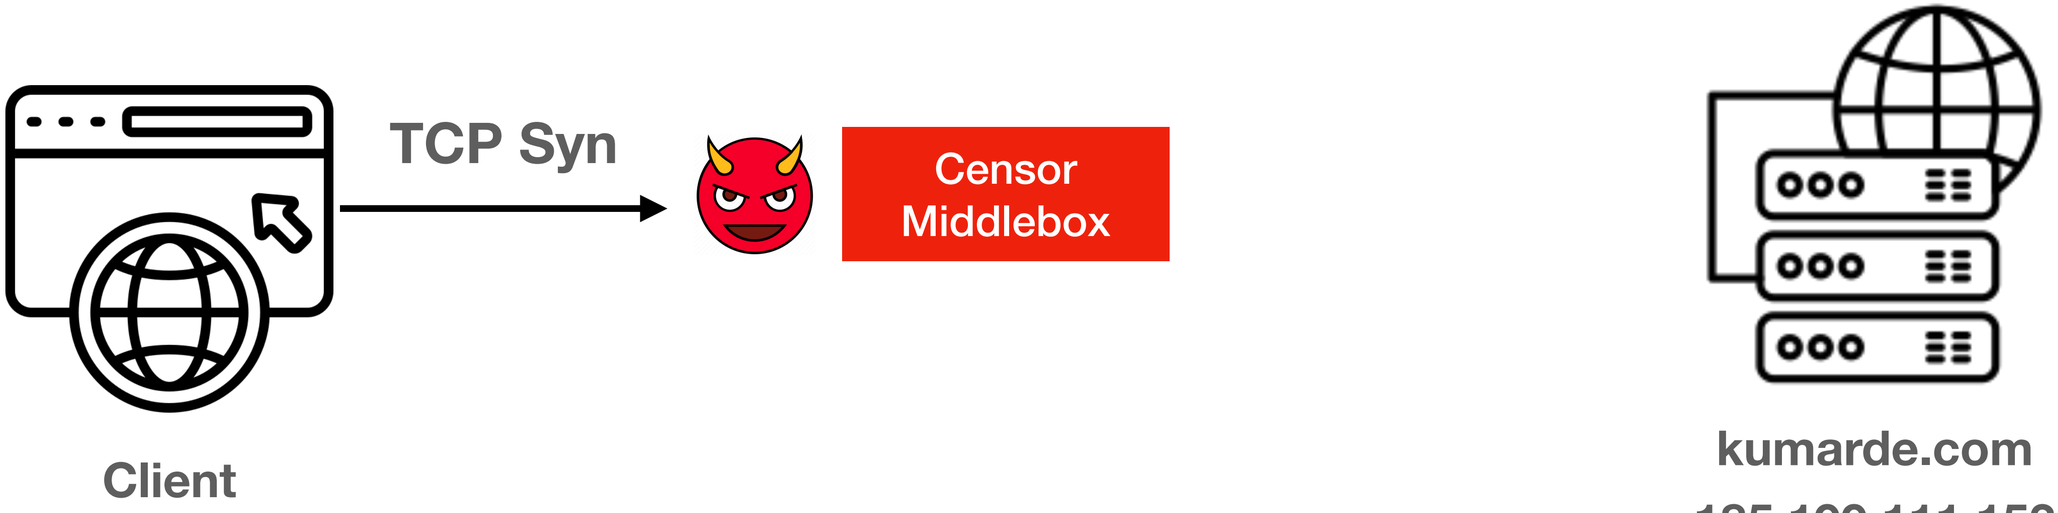

Censorship during an Internet connection

IP Blocking

kumarde.com

185.199.111.153

Censorship during an Internet connection

IP Blocking

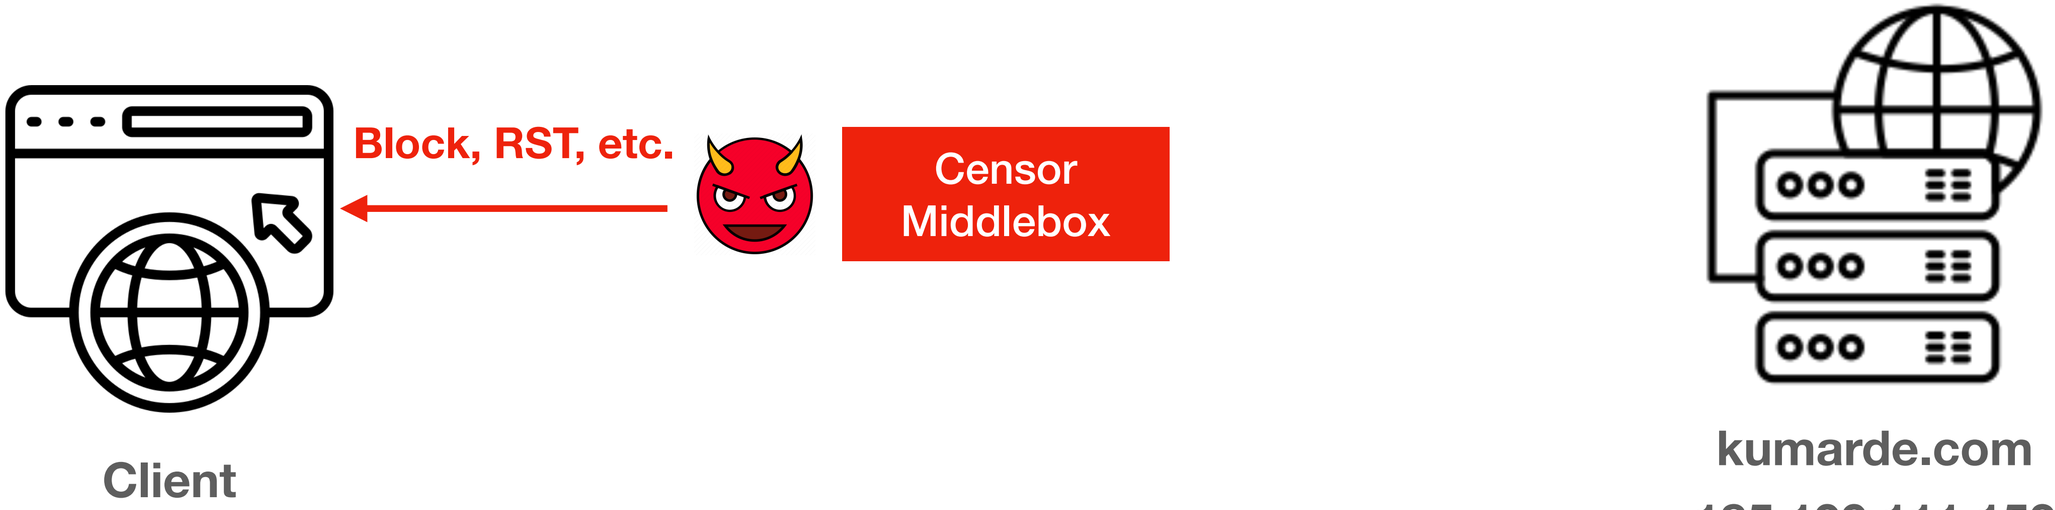

How easy do we think this is to implement?

How Internet Censors Work

Major mechanisms for running Internet Censorship

- Primary form of Internet censorship: network-level blocking
- Three main ways that network-level blocking happens in practice
  - DNS Manipulation
  - IP blocking
  - Application layer blocking

Censorship during an Internet connection

Application Layer Blocking

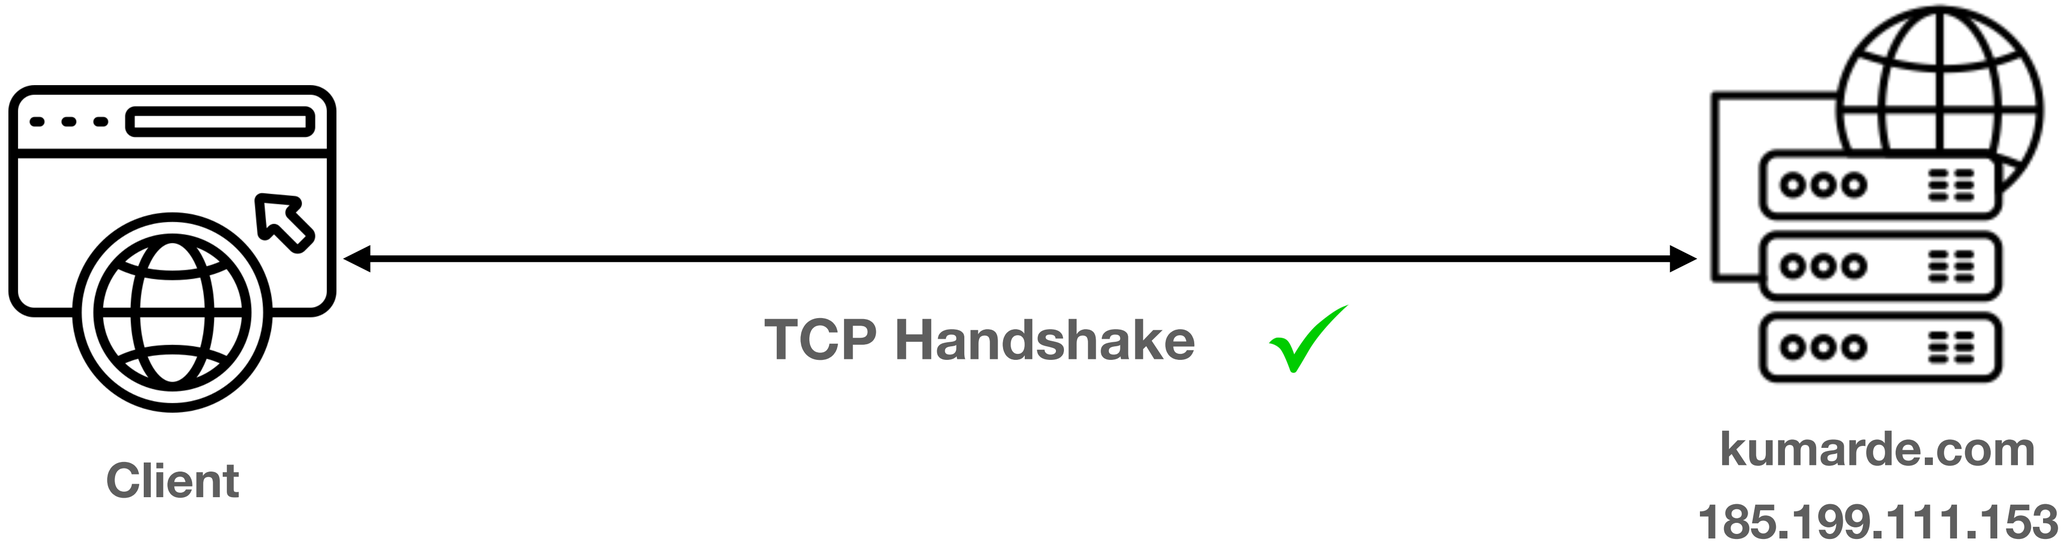

Censorship during an Internet connection

Application Layer Blocking

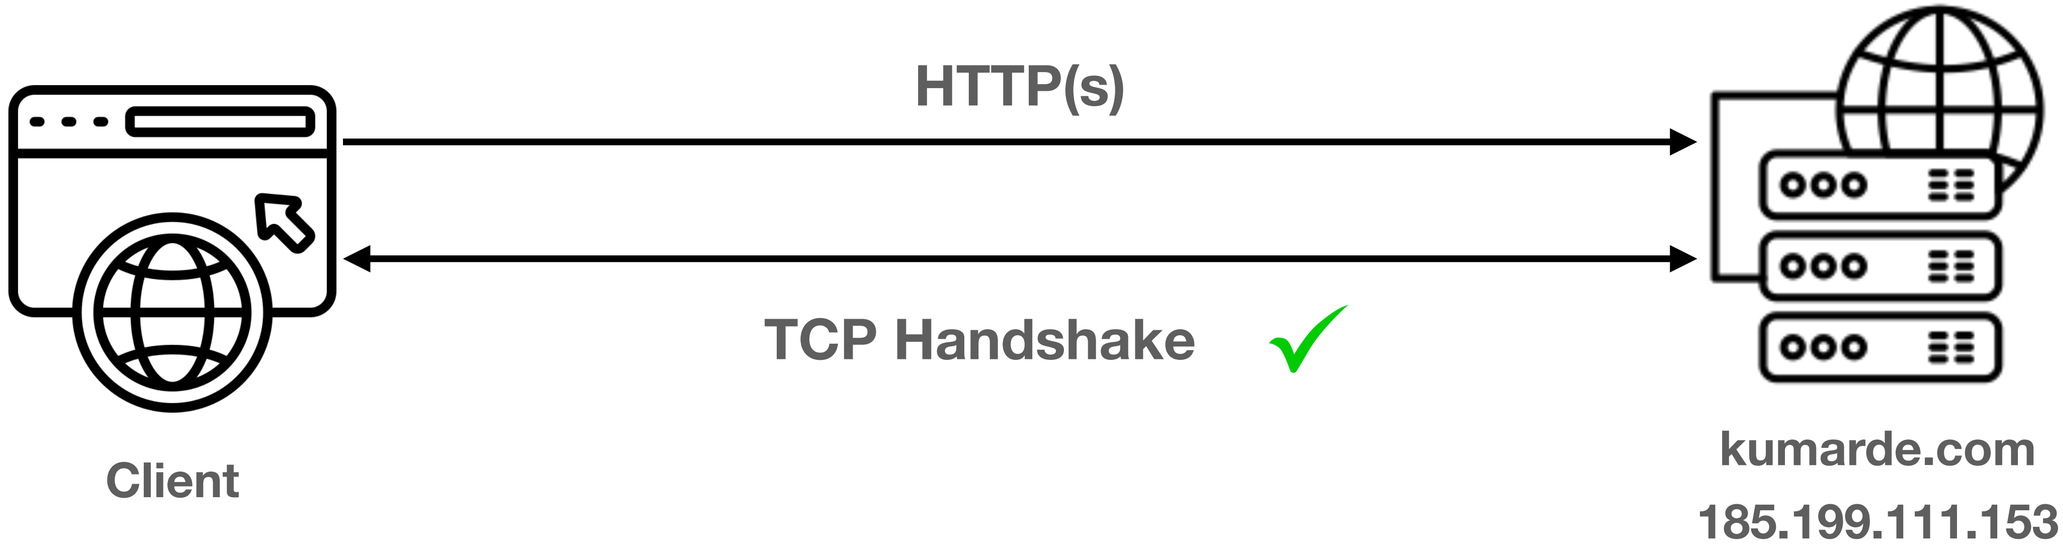

Censorship during an Internet connection

Application Layer Blocking

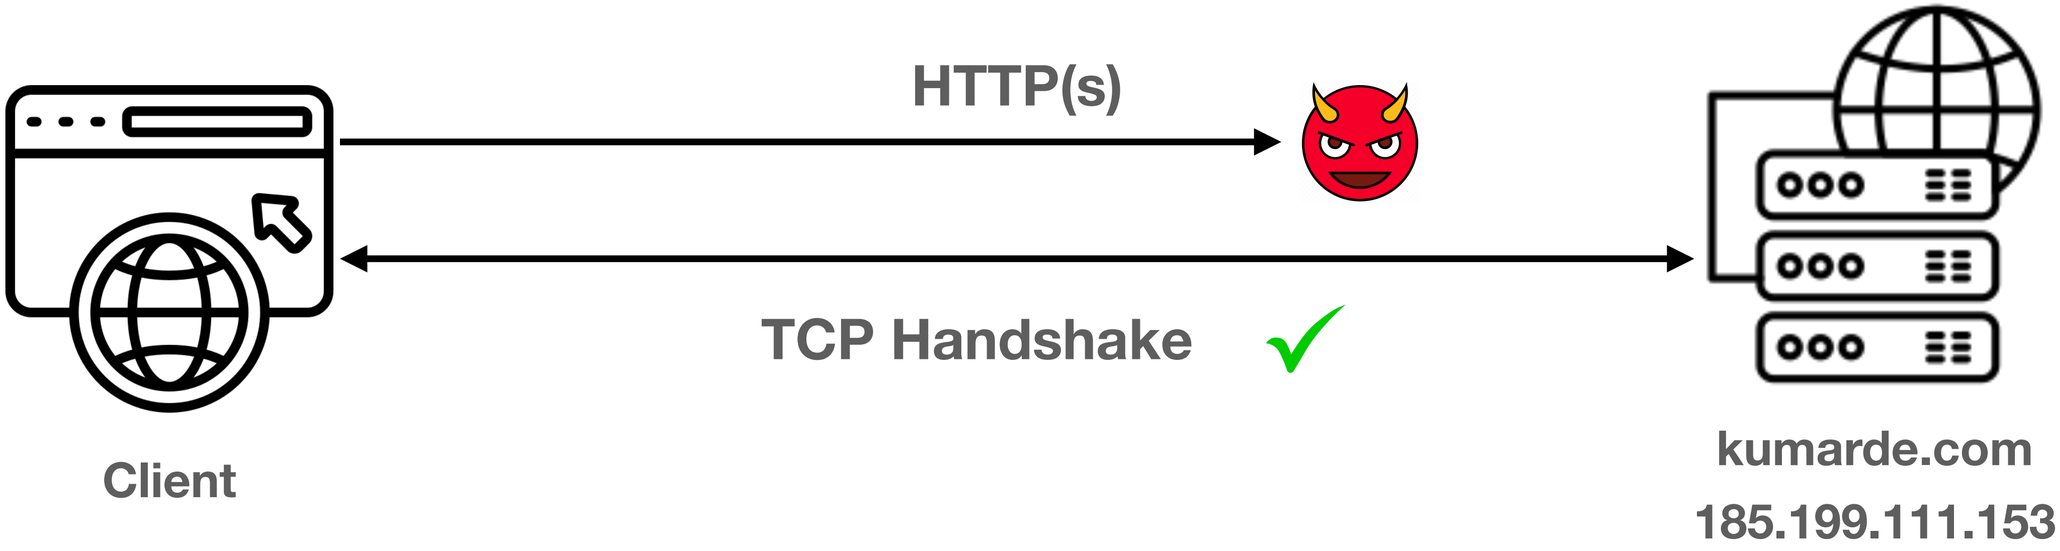

Censorship during an Internet connection

Application Layer Blocking

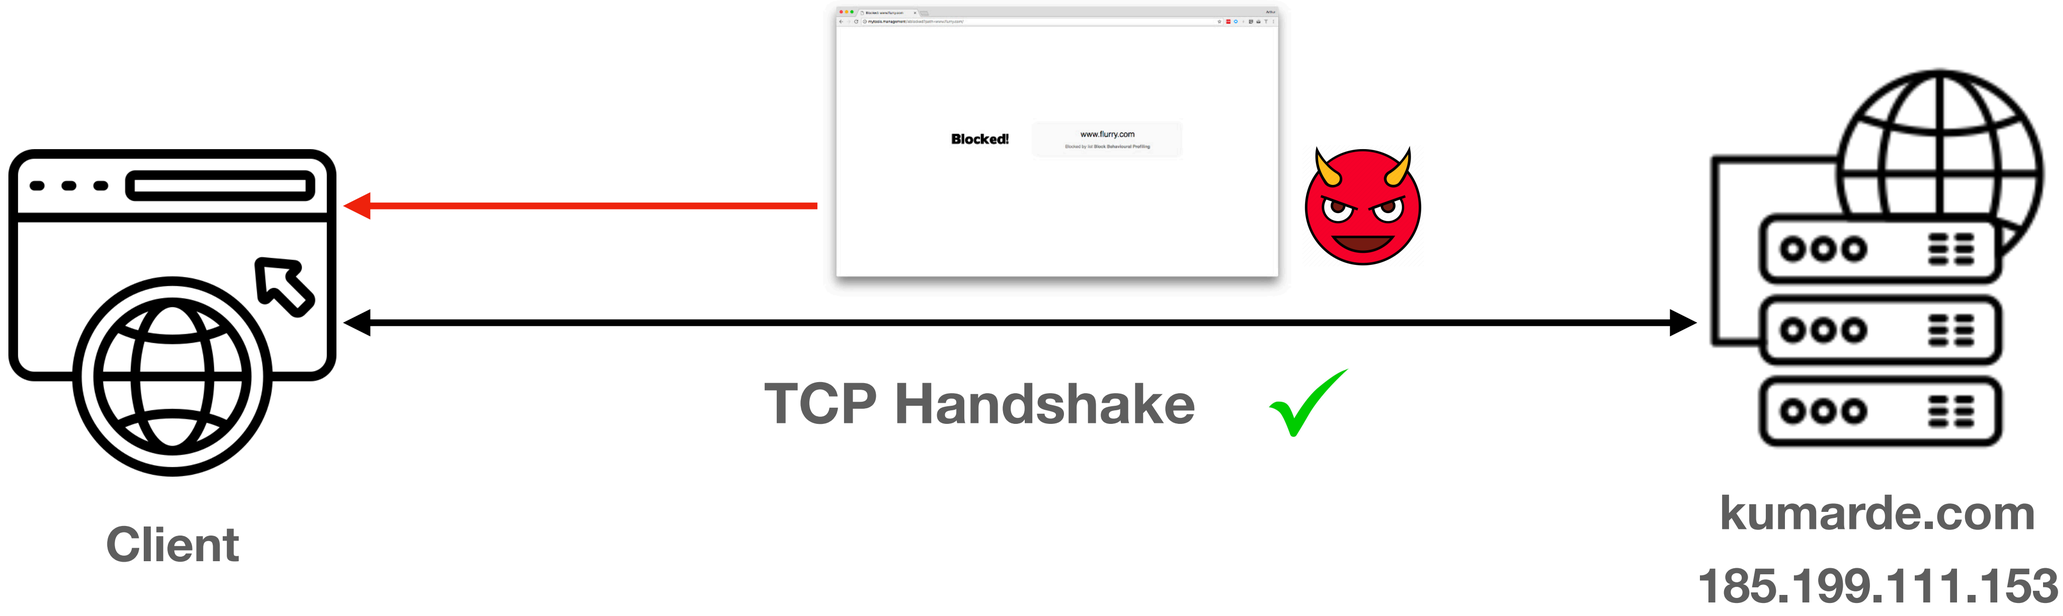

Censorship during an Internet connection

Application Layer Blocking

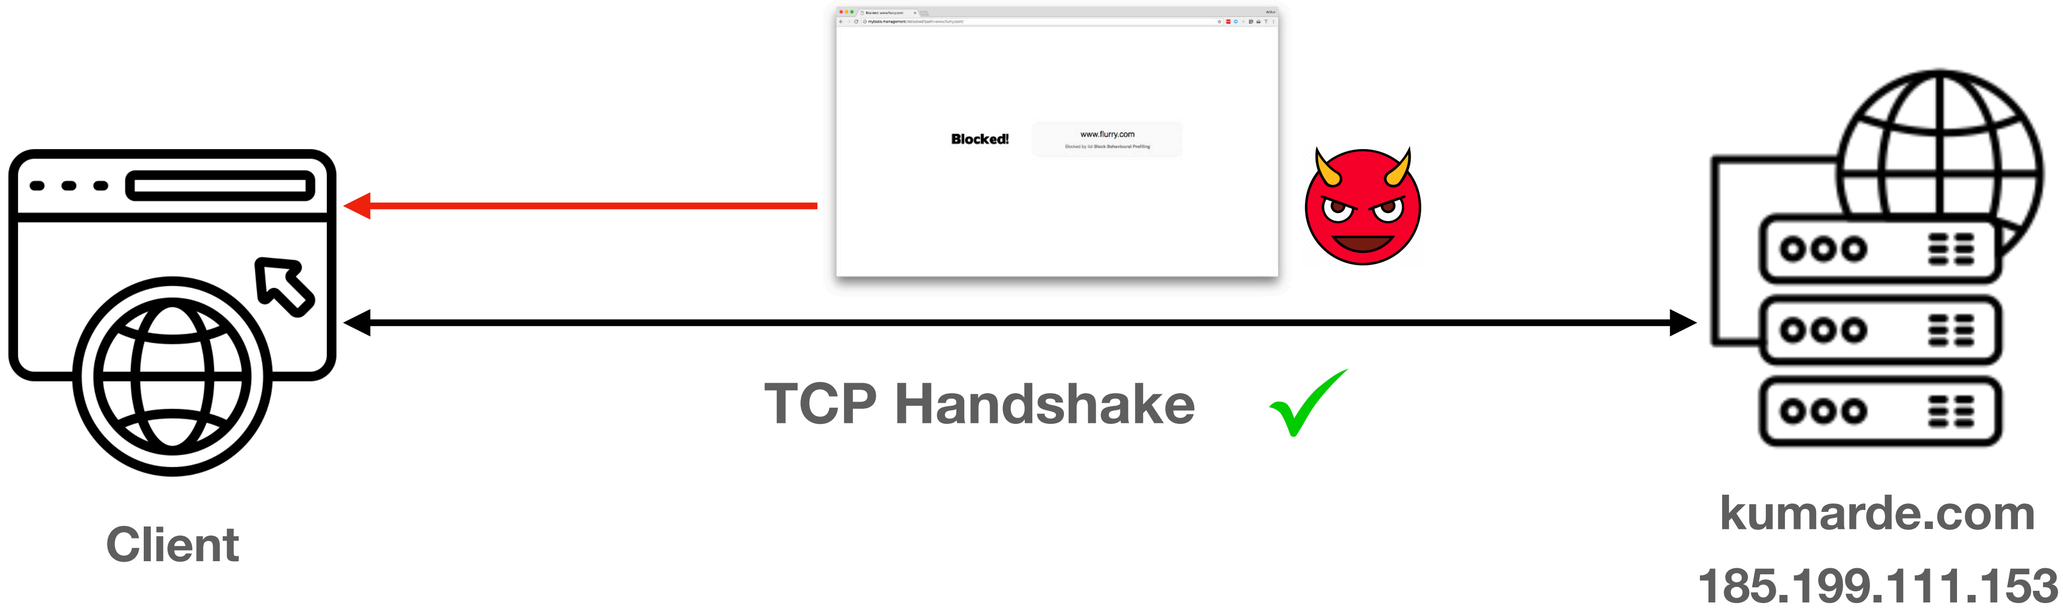

How easy do we think this is to implement?

Why measure censorship?

Censorship harms + how data can help

Network Censorship is on the rise 😞

- Information controls harm citizens
- Spreading beyond just large countries
- Frequently opaque in topic + technique

Measurements help us to:

- Support transparency + accountability
- Improve technical defenses
- Inform users + public policy

Anti-censorship in Turkey in 2014

"…When users become more aware of censorship, they often take actions that enhance Internet freedom and protect fellow users." – Freedom House

Vision for Censorship Measurement Research

Building a “weather map” of censorship

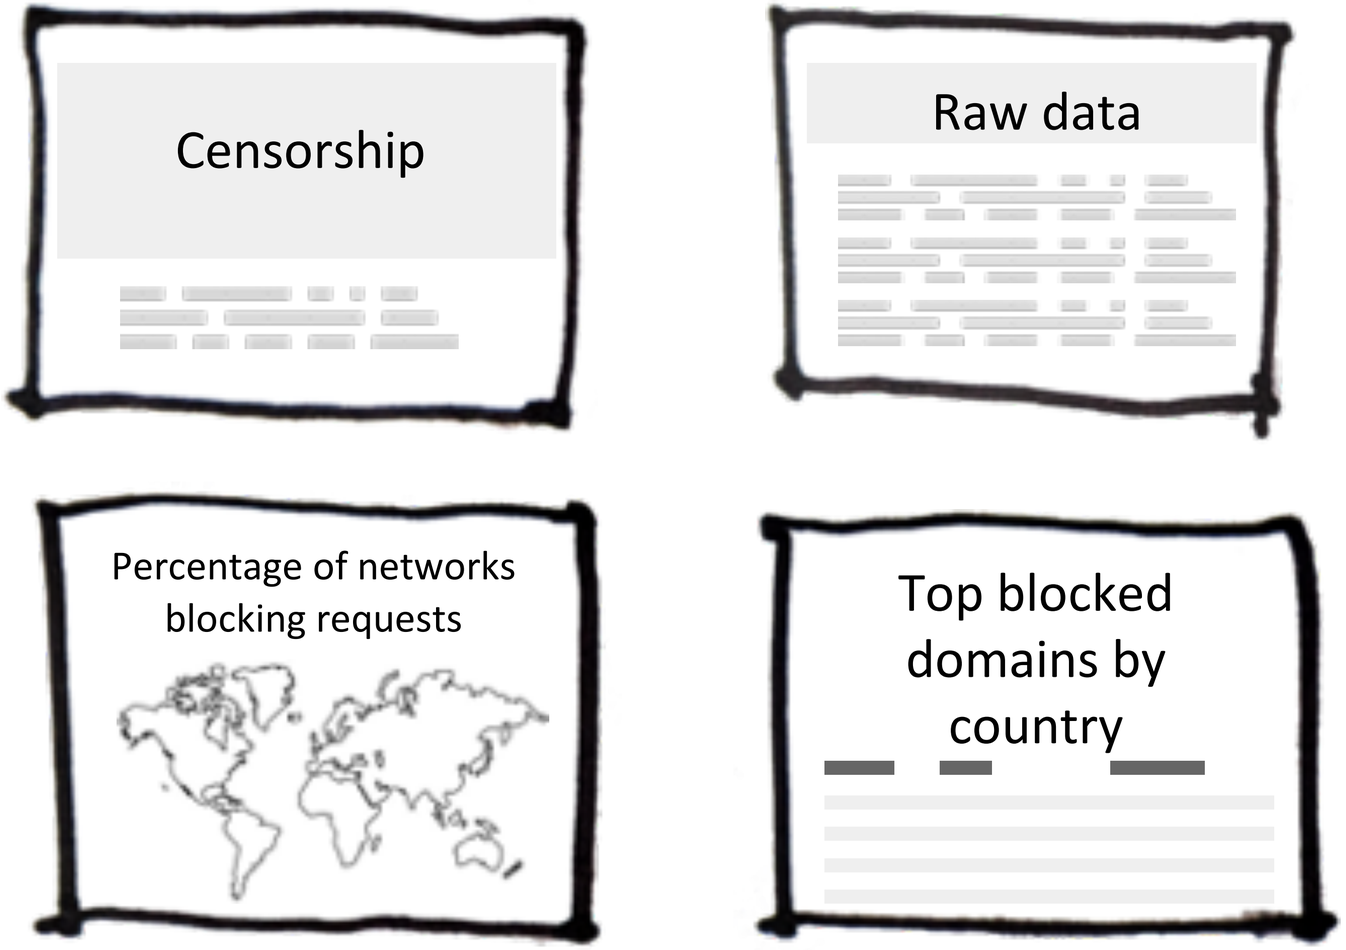

Measuring Internet censorship is hard!

Three challenges for conducting sound measurements

Measuring Internet censorship is hard!

Three challenges for conducting sound measurements

Censorship methods are varied

DNS Manipulation

TCP/IP blocking

Application layer blocking

Measuring Internet censorship is hard!

Three challenges for conducting sound measurements

Censorship methods are varied

DNS Manipulation

TCP/IP blocking

Application layer blocking

Censorship varies around the world

Geographical variance

Network variance

Measuring Internet censorship is hard!

Three challenges for conducting sound measurements

Censorship methods are varied

DNS Manipulation

TCP/IP blocking

Application layer blocking

Censorship varies around the world

Geographical variance

Network variance

Censorship varies over time

Cat + mouse game

First studies into censorship

Few countries, limited snapshots

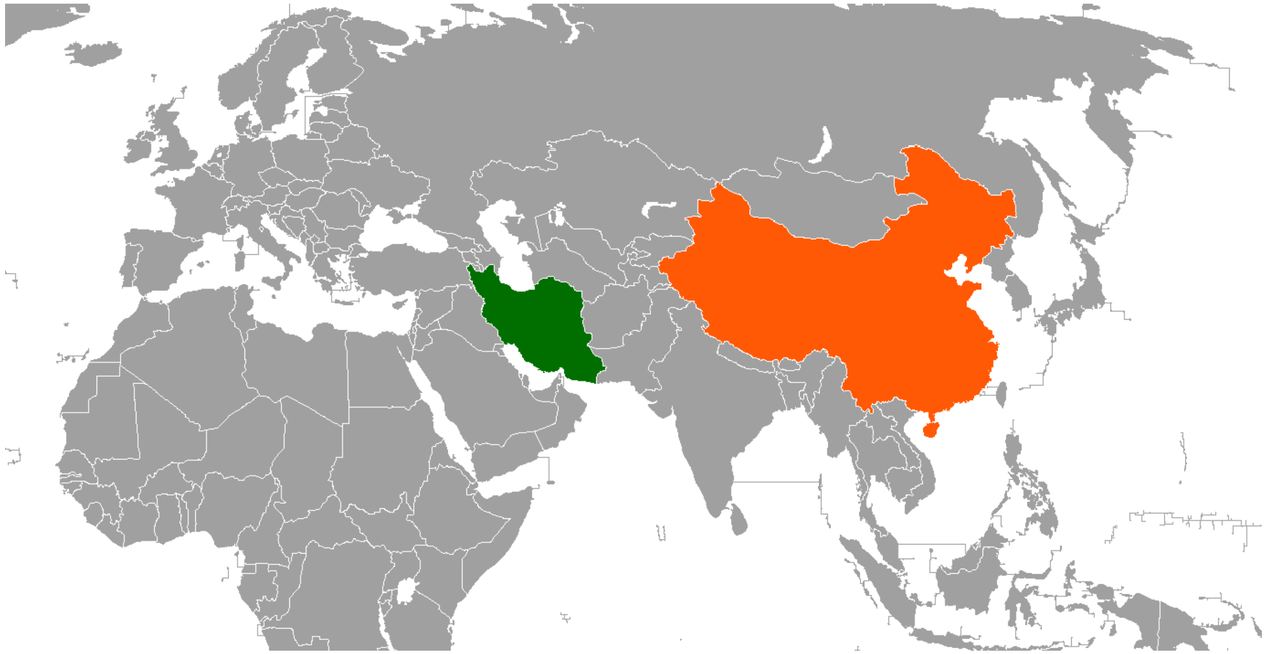

Triplet Censors: Demystifying Great Firewall's DNS Censorship Behavior

Anonymous

Arian Akhavan Niaki
University of Massachusetts Amherst

Nguyen Phong Hoang
Stony Brook University

Phillipa Gill
University of Massachusetts Amherst

Amir Houmansadr
University of Massachusetts Amherst

Internet Censorship in Iran: A First Look

Simurgh Aryan *
Aryan Censorship Project
aryan.censorship.project@gmail.com

Homa Aryan *
Aryan Censorship Project
aryan.censorship.project@gmail.com

J. Alex Halderman
University of Michigan
jhalderm@umich.edu

OONI | Probe

Measure internet censorship

Contribute to the world's largest open dataset on internet censorship

Download OONI Probe for macOS

Other Platforms »

User Guide »

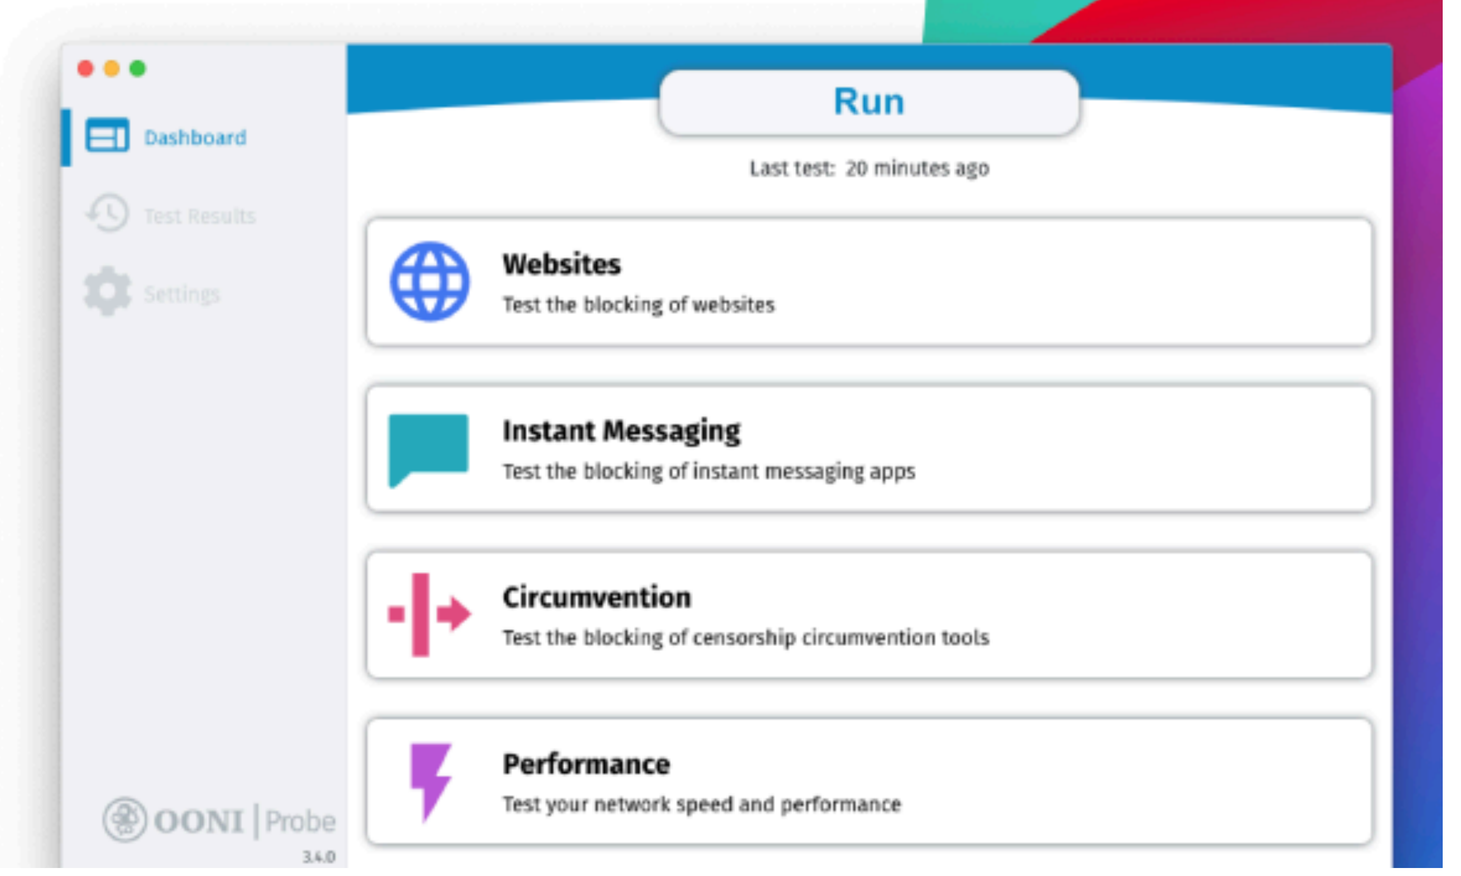

How OONI Works

Volunteer-based direct measurements of censorship

OONI

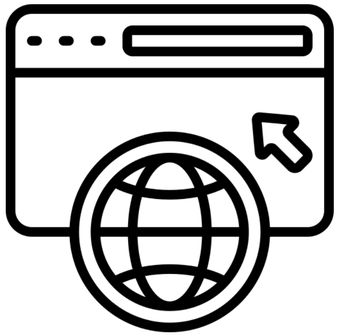

Volunteer client
in-country

Censor

Server of Interest

Control Relay

How OONI Works

Volunteer-based direct measurements of censorship

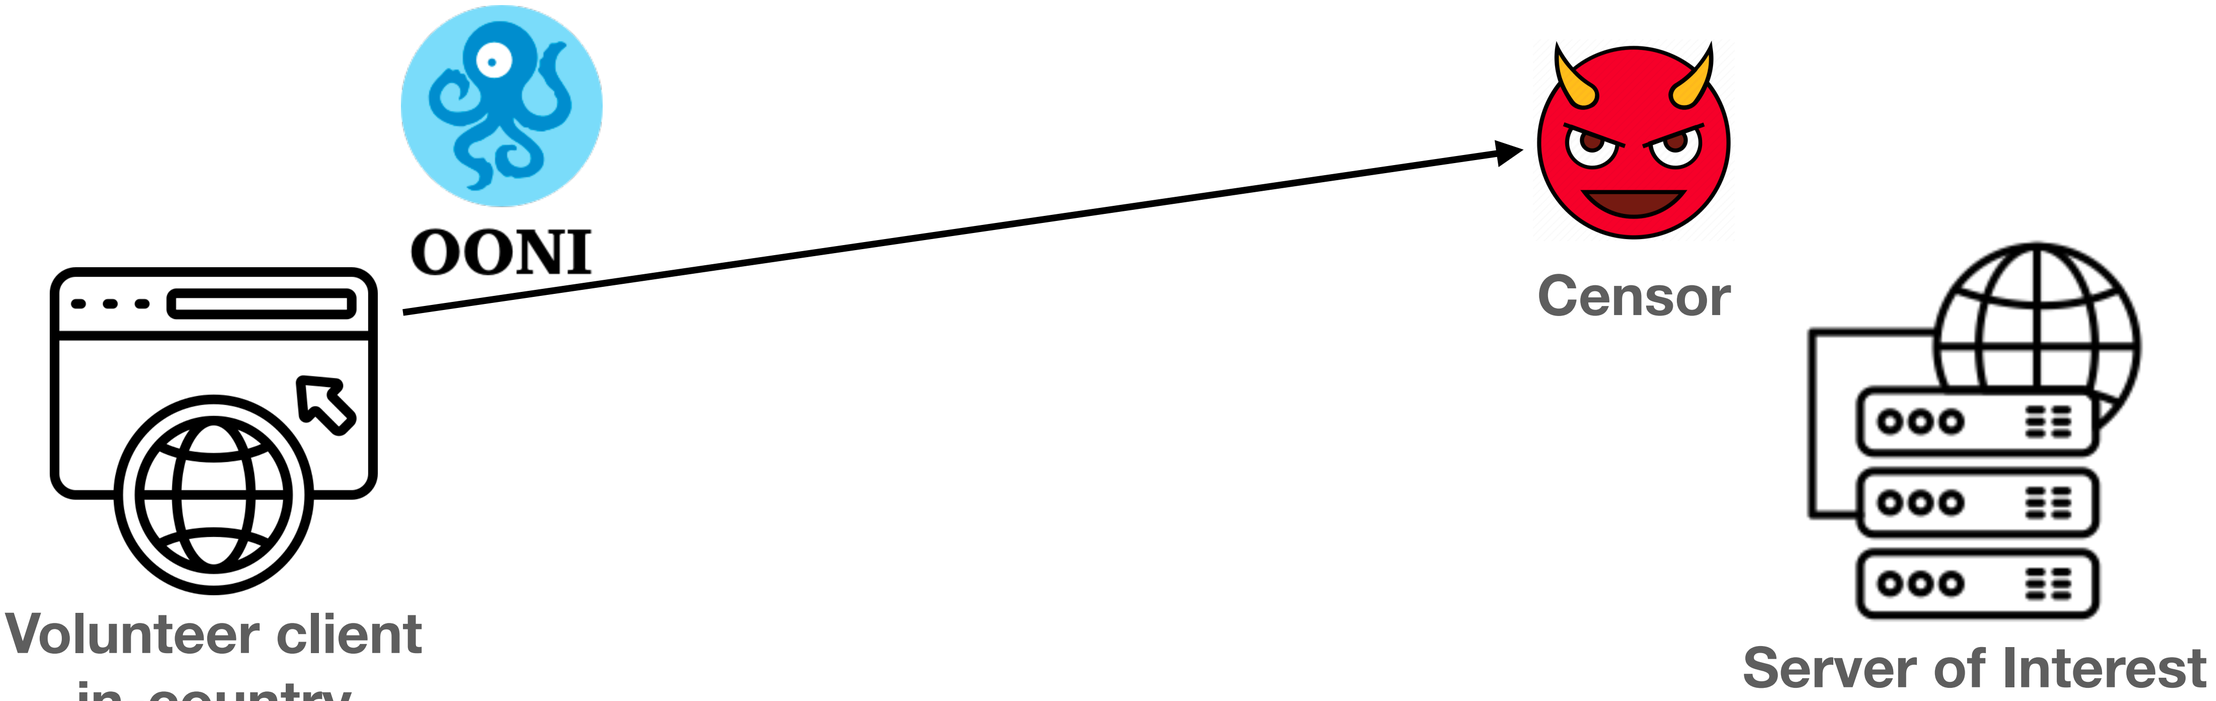

How OONI Works

Volunteer-based direct measurements of censorship

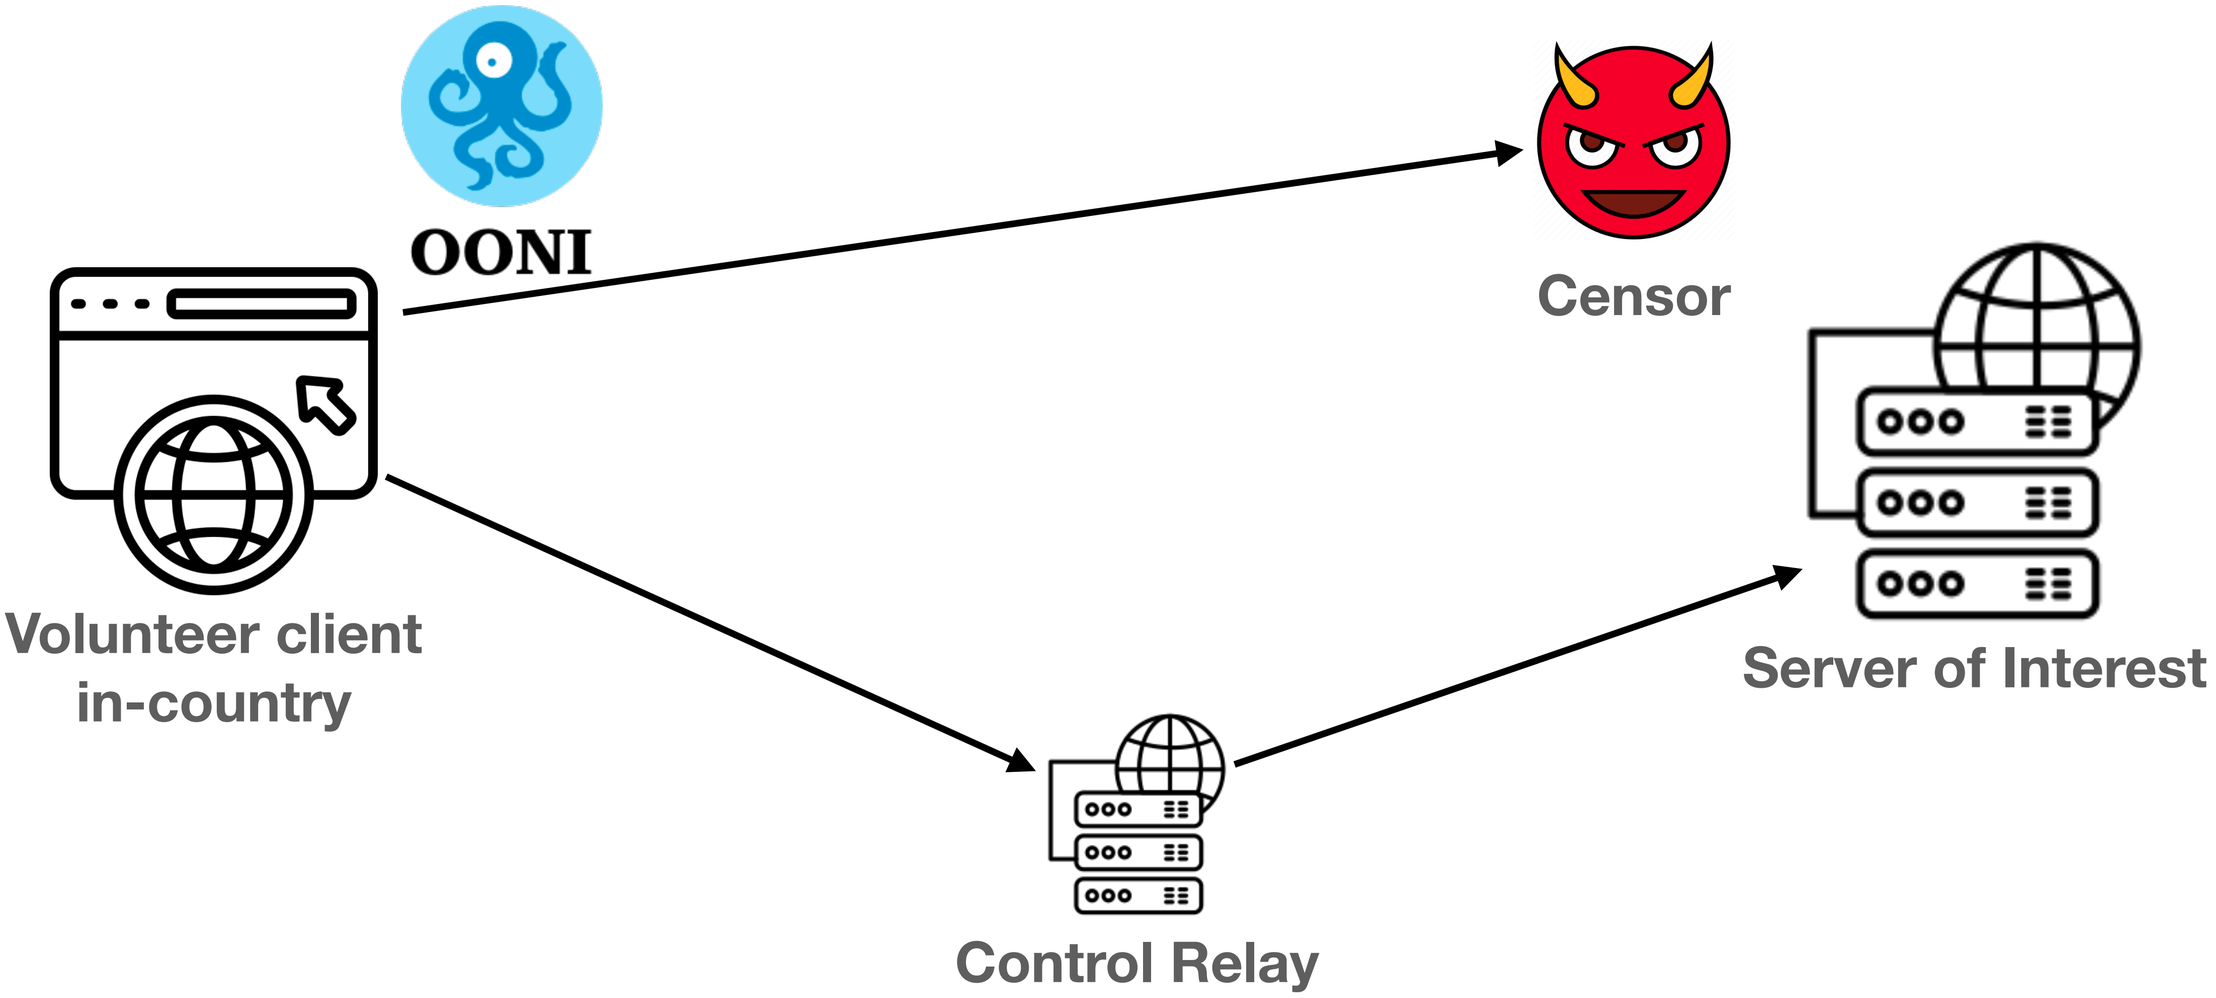

How OONI Works

Volunteer-based direct measurements of censorship

https://explorer.ooni.org

Limitations of volunteer measurements

5 key problems

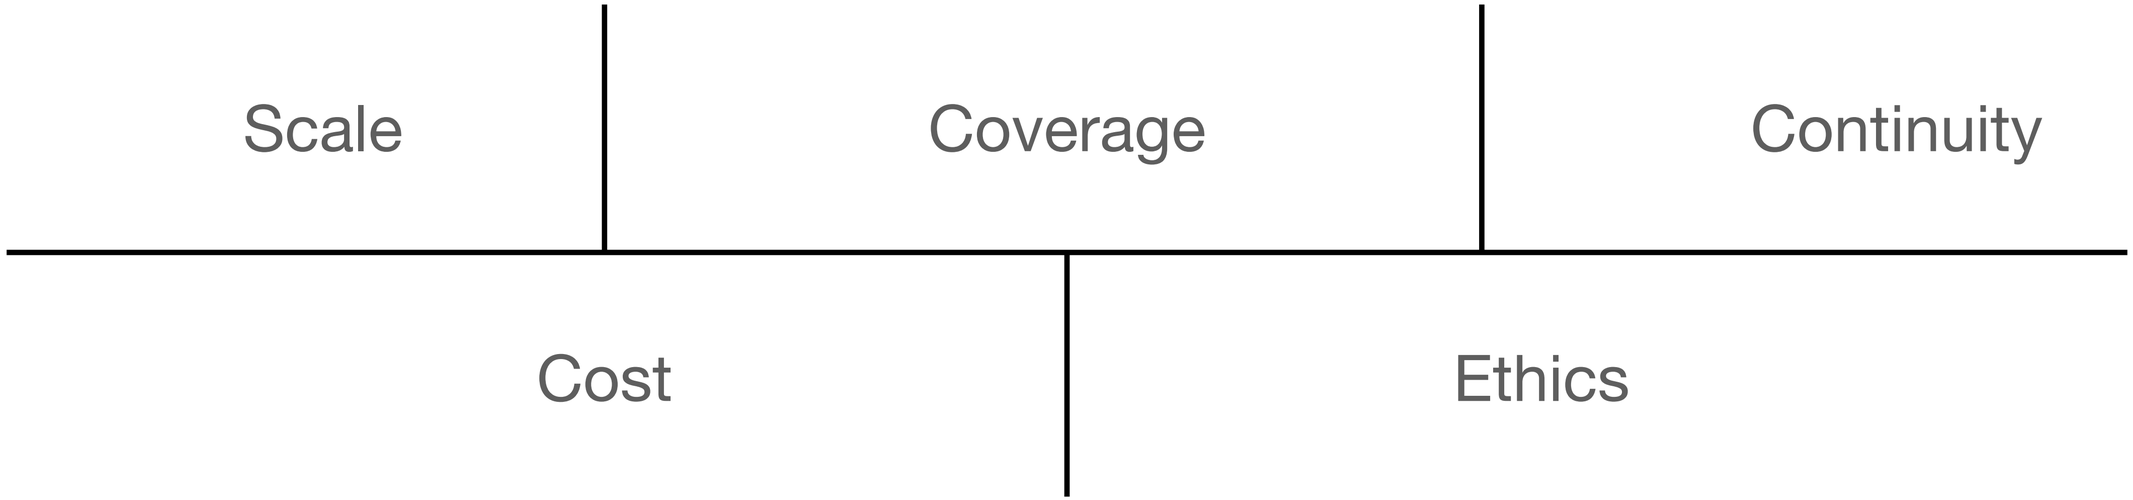

A word on ethics…

Ethical Concerns for Censorship Measurement

Ben Jones, Roya Ensafi, Nick Feamster, Vern Paxson, Nick Weaver

Princeton University, UC Berkeley, International Computer Science Institute

Under what conditions is it safe to use volunteers devices?

What populations of users are affected?

Do people incur no more than minimal risk?

Do the benefits to the population balance the risks?

Circumventing Censors

Proxying requests through “safe” servers — e.g., VPN

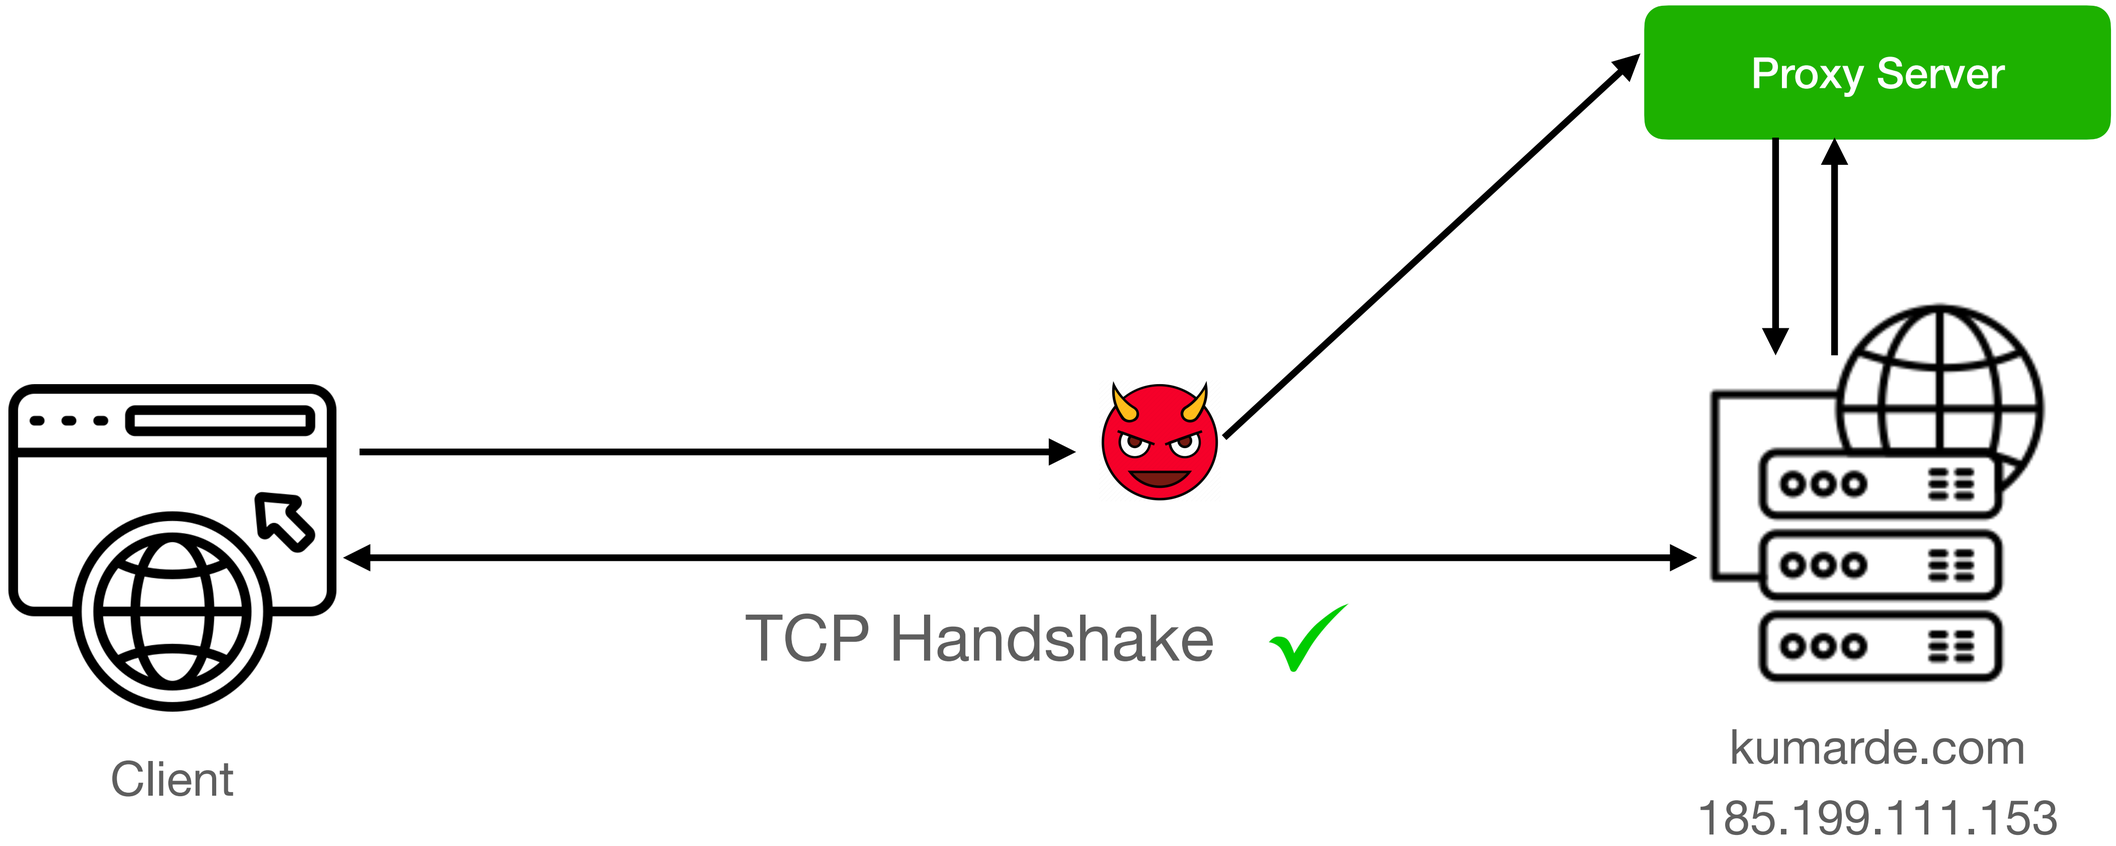

Circumventing Censors

Proxying requests through “safe” servers is easy to detect

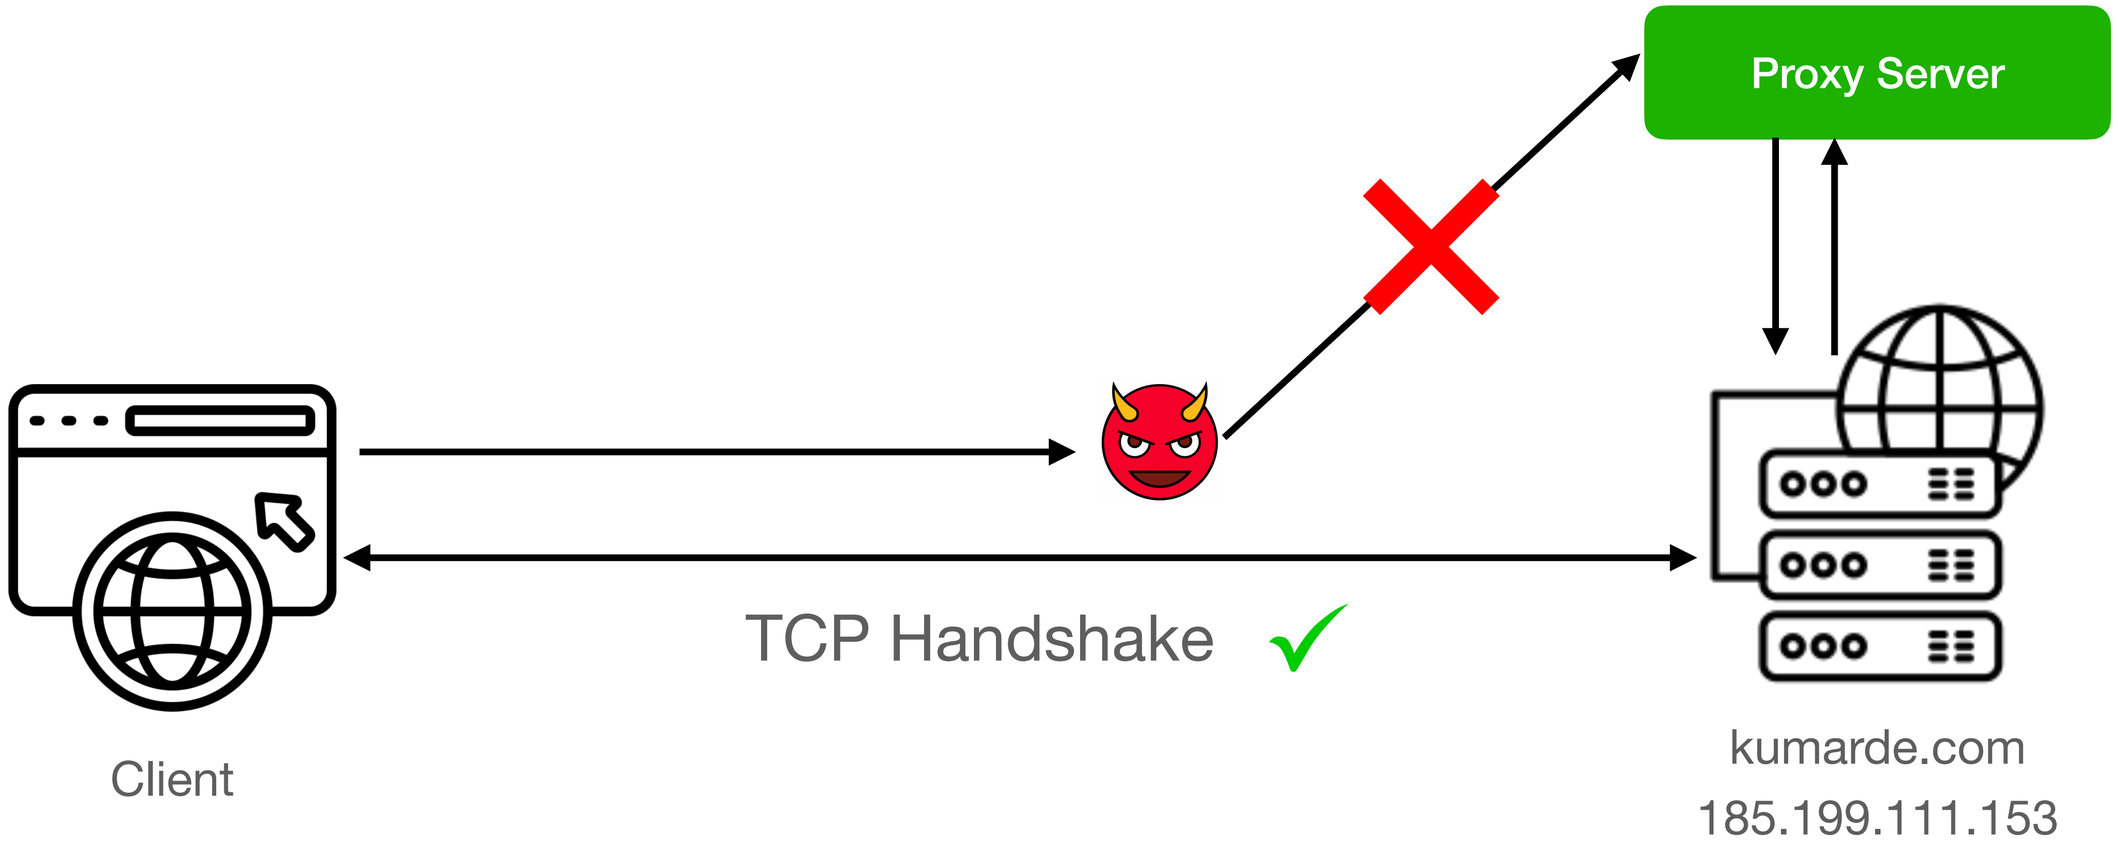

Circumventing Censors

Imitating non-censored protocols

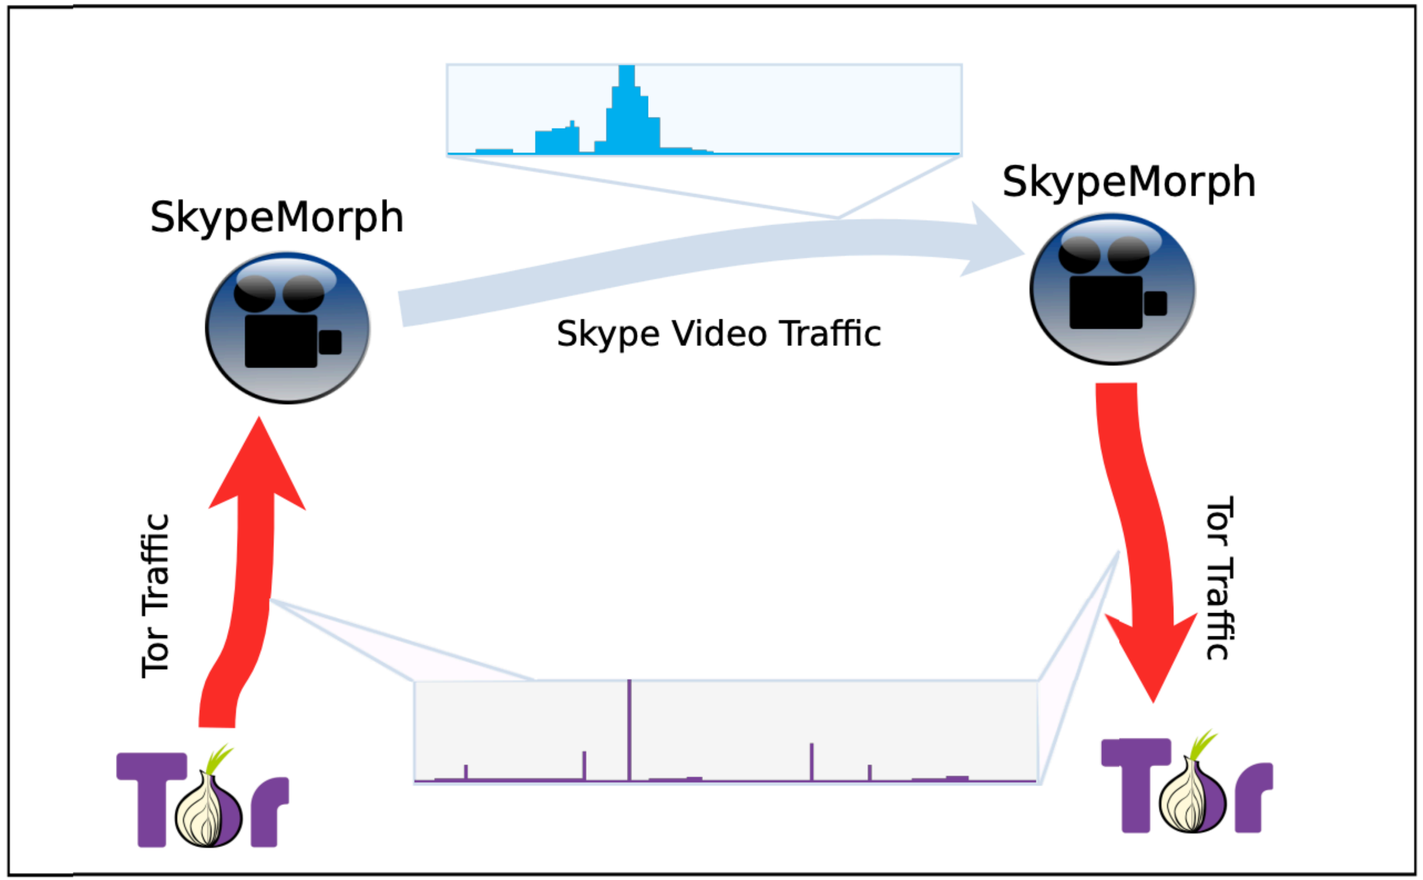

Circumventing Censors

Imitating non-censored protocols has problems

The Parrot is Dead:
Observing Unobservable Network Communications

Amir Houmansadr Chad Brubaker Vitaly Shmatikov
The University of Texas at Austin

Circumventing Censors

Refraction Networking

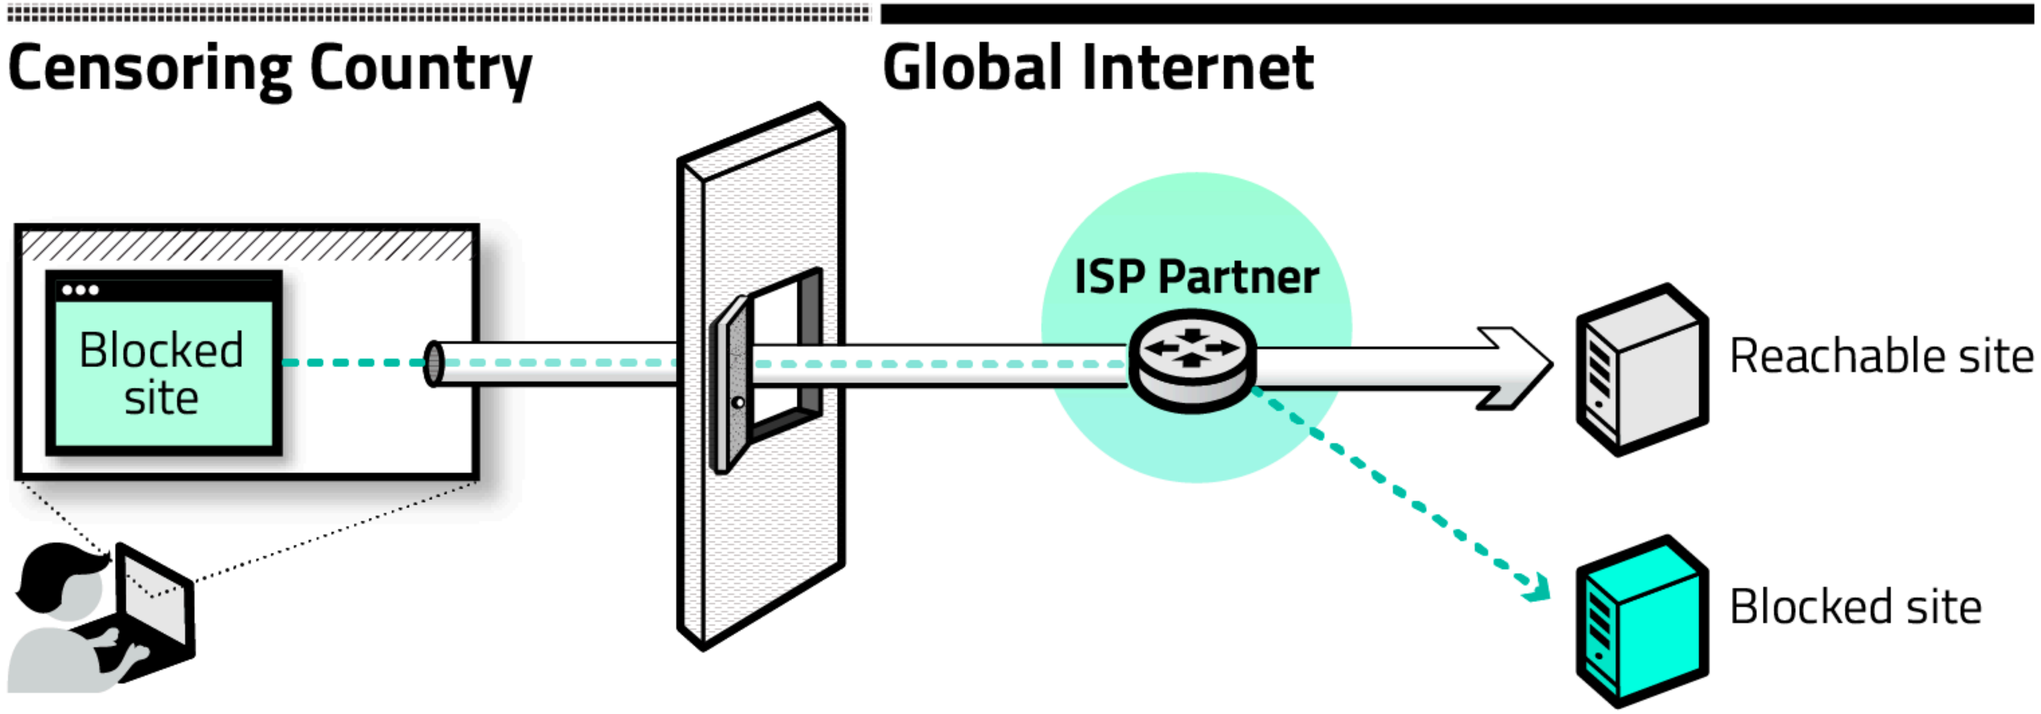

https://refraction.network/

Circumventing Censors

Refraction Networking

Benjamin VanderSloot*, Sergey Frolov, Jack Wampler, Sze Chuen Tan, Irv Simpson, Michalis Kallitsis, J. Alex Halderman, Nikita Borisov, and Eric Wustrow

Running Refraction Networking for Real

An ISP-Scale Deployment of TapDance

Sergey Frolov[1], Fred Douglas[3], Will Scott[5], Allison McDonald[5], Benjamin VanderSloot[5], Rod Hynes[6], Adam Kruger[6], Michalis Kallitsis[4], David G. Robinson[7], Steve Schultze[2], Nikita Borisov[3], J. Alex Halderman[5], and Eric Wustrow[1]

[1]University of Colorado Boulder [2]Georgetown University Law Center [3]University of Illinois Urbana-Champaign
[4]Merit Network [5]University of Michigan [6]Psiphon [7]Upturn

https://refraction.network/

Network Defense

Network Perimeter Defense

- Attacker model: internal network is “privileged”, let’s prevent people from getting access to our internal network
- Idea: network “defenses” on the “outside” of an organization: stop the bad person before they get inside the system
- Typical elements
  - Firewalls and proxies
  - Network address translation
  - Network content analysis / Network Intrusion Detection System (NIDS)

Firewalls

- Problem: You'd like to protect or isolate one part of the network from other parts
  - E.g., protect your network from the global Internet
- So… how? With a **firewall.**
  - Basic idea: filter or otherwise limit network traffic on a number of different fields / protocols
- Key questions:
  - What information do you use to filter?
  - Where do you put the firewall?

Types of firewalls

- What can you filter on?
  - Packet data!
    - Ports, IPs, protocols… you name it
- Where do you do the filtering?
  - Personal firewalls
    - Run at the end hosts; pros are there are more specific application information, cons are that end-users have to set them up
  - Network firewalls
    - Intercept and evaluate communication from many hosts; requires global coordination

Firewall Deployment

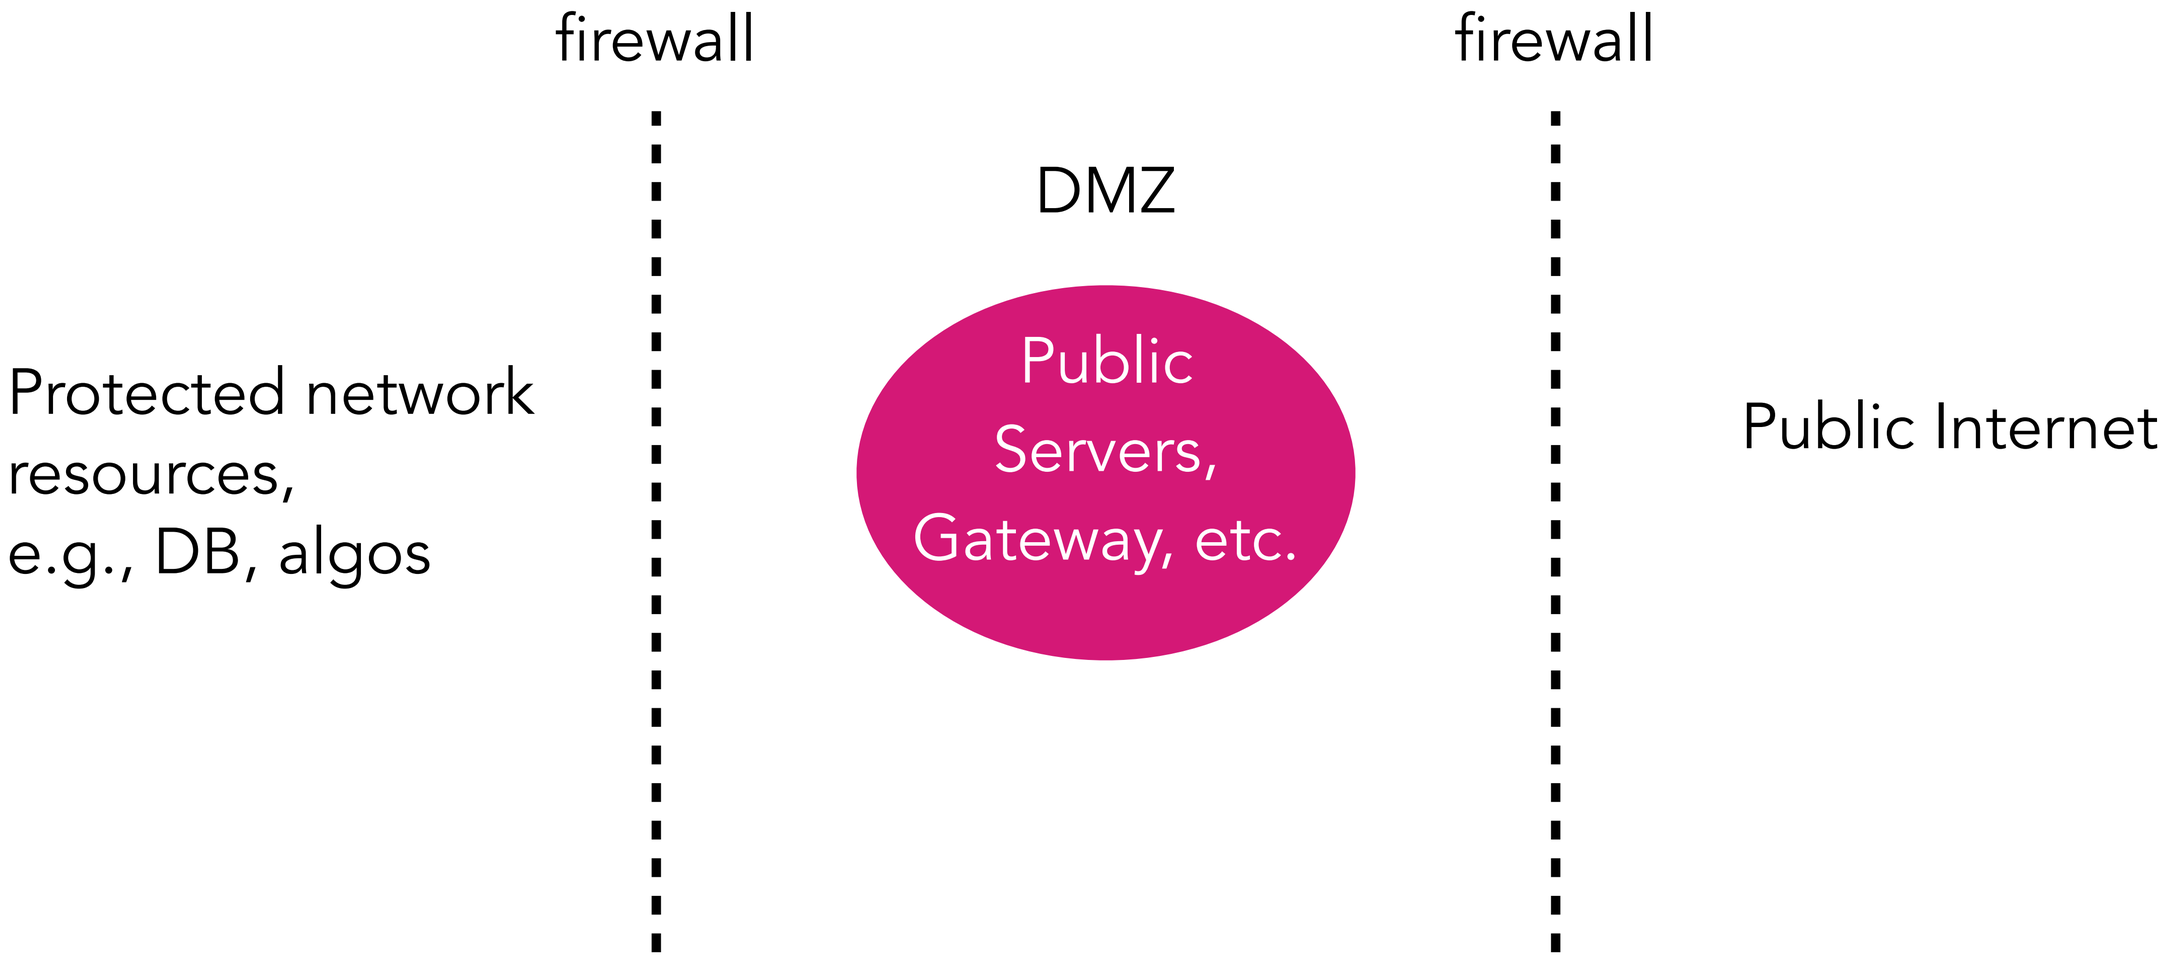

Firewall serves as an attack surface; interactions are limited through it

Packet Filtering Firewalls

- Defines list of access-control rules… checks every packets against those rules and forward or drop
  - See nftables in Linux https://www.netfilter.org/projects/nftables/index.html
- Packet filtering firewalls can take advantage of anything in the packet, e.g.,
  - Source IP
  - Destination IP
  - Source port
  - Destination port
  - Flags (e.g., ACK, SYN, etc.)

Proxy-Based Firewalls

- Basic idea is to *proxy* requests through a DMZ server that does one thing and one thing only
  - This is how my *actual* research network is set up
- Cons: can be painful to maintain and set up; requires careful thought

Network Address Translation

- Basic idea: Multiplex one IP to many with “private IP addresses”
  - Make it so that devices internal to a network are “unseeable” from the outside!
- Pros
  - Only allows connections to the outside that are established from the *inside*
    - Default assumption is inside is protected
- Cons
  - Breaks some protocols (e.g., some streaming protocols), router needs to handle thorny logic…. despite this, NAT is ubiquitous

Network content analysis

- Idea: Devices want to look at network traffic **content** for security
  - Network intrusion detection systems (NIDS)
  - Spam filter
  - Traffic differentiation (e.g., slow down connections at edge…)
  - Deep packet inspection (basically, look beyond just header)
  - …enabled by *middleboxes*
- **Middlebox**: A hop in the path imposed by the network that does *something* to the network flow other than simply forwarding the packet

Network content analysis

- Network vantage point is appealing to implement policy because **every packet has to flow through it**
  - But… can be challenging to infer the semantics of unknown communication
  - These days, end-to-end encryption makes this a challenge… more on this next week
- In a typical network, you might expect to see any combination of all three of these: firewalls, NATing, NIDS… these are our best lines of defense against network attackers

So, in sum..

- The Internet is an incredible resources…. and can be very easily manipulated because we built it with no security in mind
  - We keep stapling stuff on after the fact
  - Bites us in the butt in many contexts, spoofing, DoS, censorship, privilege…. etc.
- But, despite that, we've stapled a lot of stuff these days, and most networks aren't under constant threat of failure…. for the most part :)

Next time…

- Move away from networks and to ***cryptography***; which is our last big unit in the class!
  - Gets a little math-y, FYI, but not too much
- PA4 due Thursday. Good luck!