CSE227 – Graduate Computer Security

Censorship

UC San Diego

Housekeeping

General course things to know

- Midpoint check-in is done! Final stretch towards the projects
  - Early feedback – good progress, many pivots, make sure you're building towards something by the end of the quarter
  - Grades should be released by Friday
- Optional meetings with me available on Thursday the 20th during normal class time
  - https://calendly.com/kumarde/cse-227-project-meeting
  - Only 3 groups have signed up so far

Today's lecture

Learning Objectives

- Learn about network censorship, motivations for censorship, how it's done in practice
- Discuss methods, techniques, and general ethos around *measuring* censorship
- Talk about the Censored Planet paper
- Talk about the Throttling Twitter paper

Preliminaries

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

- What is an Internet shutdown?

How Internet Censors Work

Major mechanisms for running Internet Censorship

- What is an Internet shutdown?
  - Removing Internet service altogether – much easier in some countries than others

How Internet Censors Work

Major mechanisms for running Internet Censorship

- What is an Internet shutdown?
  - Removing Internet service altogether – much easier in some countries than others
- What is throttling?

How Internet Censors Work

Major mechanisms for running Internet Censorship

- What is an Internet shutdown?
  - Removing Internet service altogether – much easier in some countries than others
- What is throttling?
  - Making certain services slower in country boundaries

How Internet Censors Work

Major mechanisms for running Internet Censorship

- What is an Internet shutdown?
  - Removing Internet service altogether – much easier in some countries than others
- What is throttling?
  - Making certain services slower in country boundaries
- What are content takedowns?

How Internet Censors Work

Major mechanisms for running Internet Censorship

- What is an Internet shutdown?
  - Removing Internet service altogether – much easier in some countries than others
- What is throttling?
  - Making certain services slower in country boundaries
- What are content takedowns?
  - Removal of “offensive” content from online services

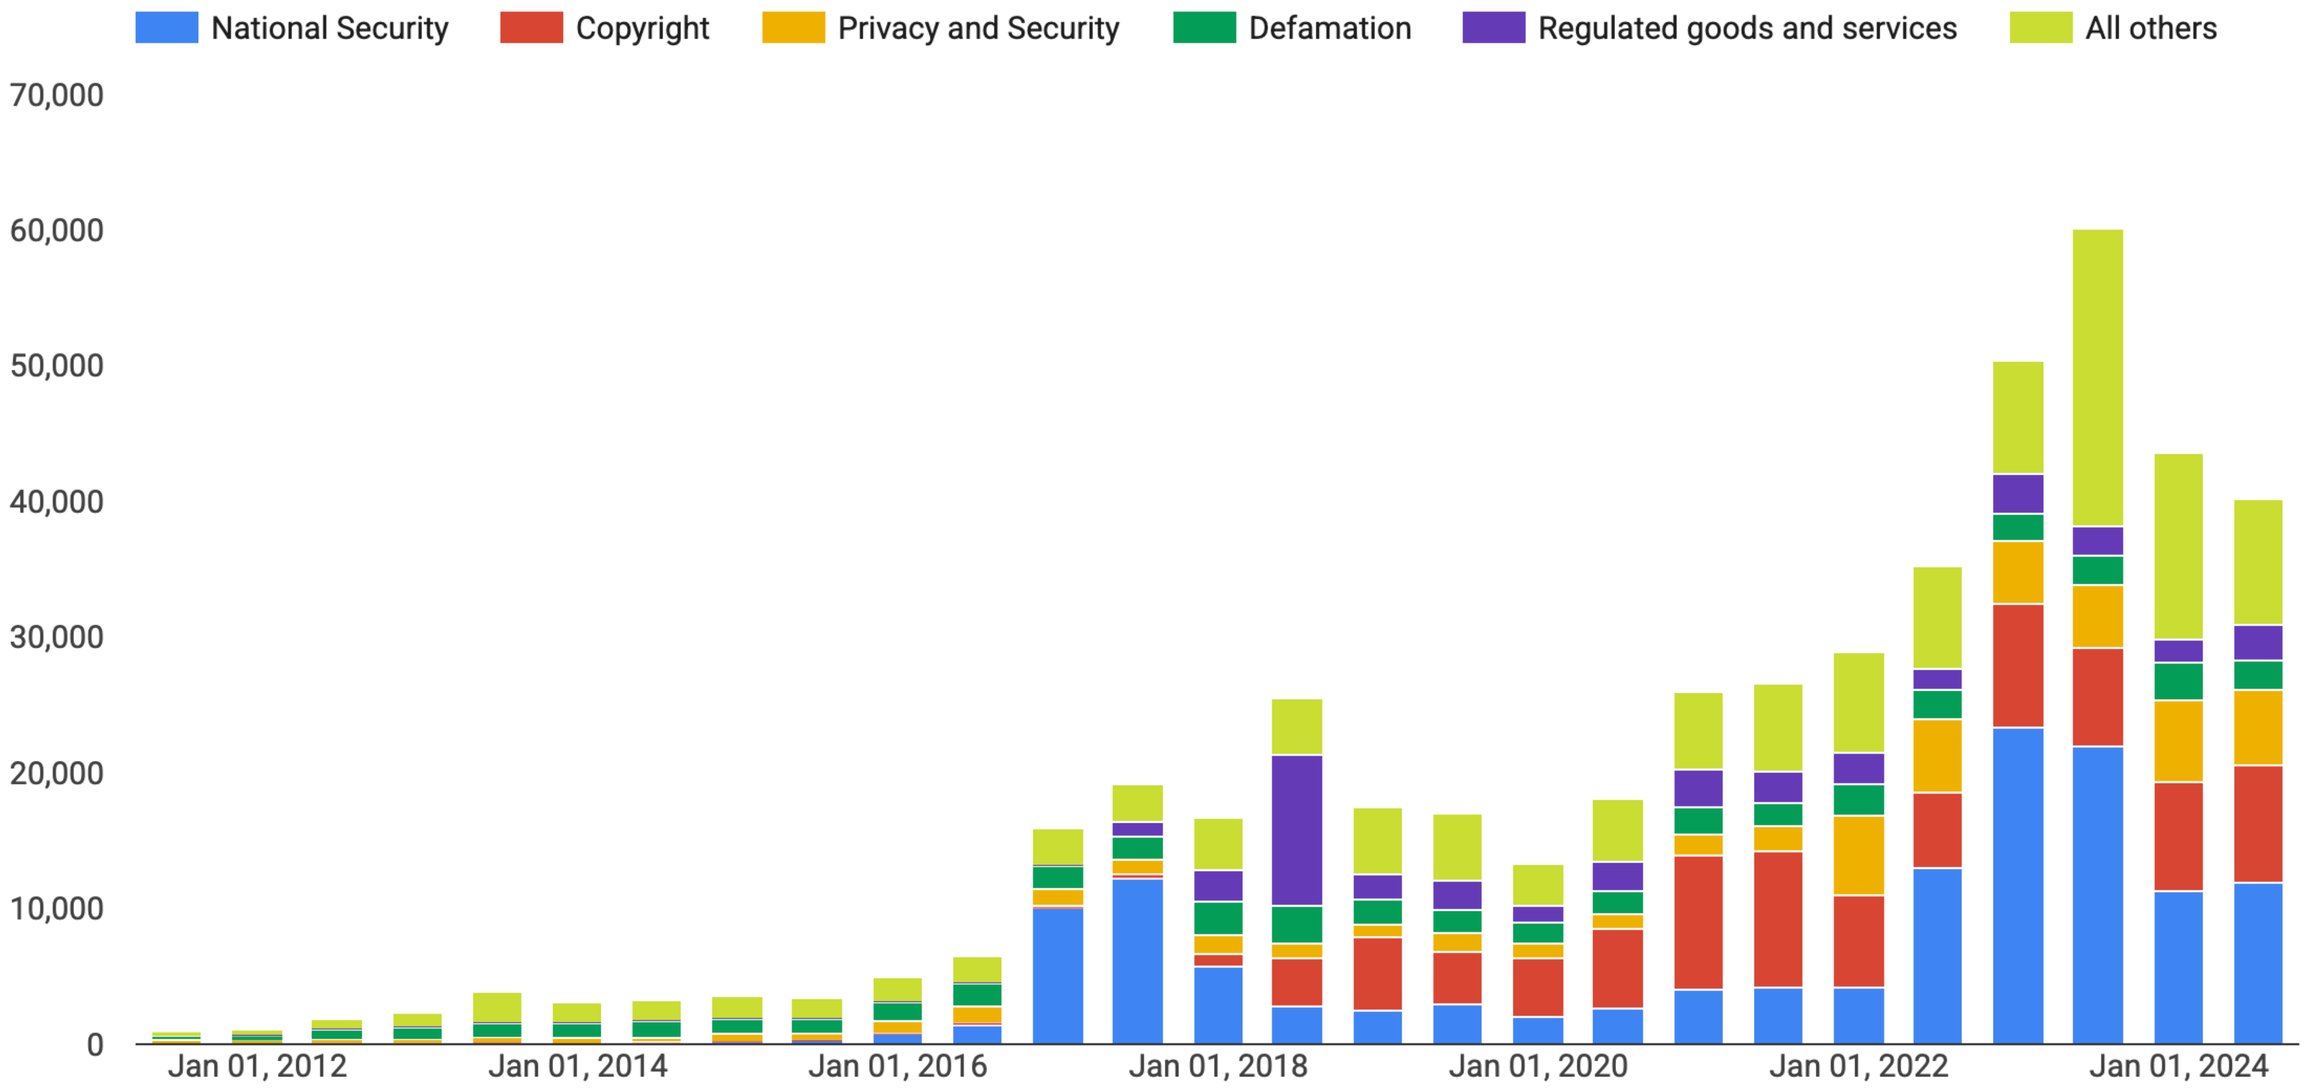
National Security
Copyright
Privacy and Security
Defamation
Regulated goods and services
All others
70,000
60,000
50,000
40,000
30,000
20,000
10,000
0
Jan 01, 2012
Jan 01, 2014
Jan 01, 2016
Jan 01, 2018
Jan 01, 2020
Jan 01, 2022
Jan 01, 2024

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
  - What is DNS manipulation?

Censorship during an Internet connection

Modes of website blocking

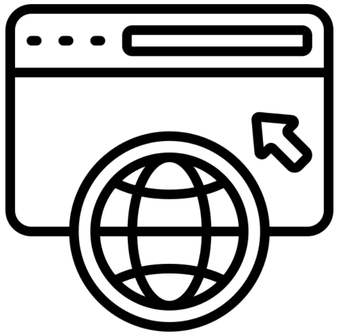

Client

CSE227 Server

Censorship during an Internet connection

DNS Manipulation

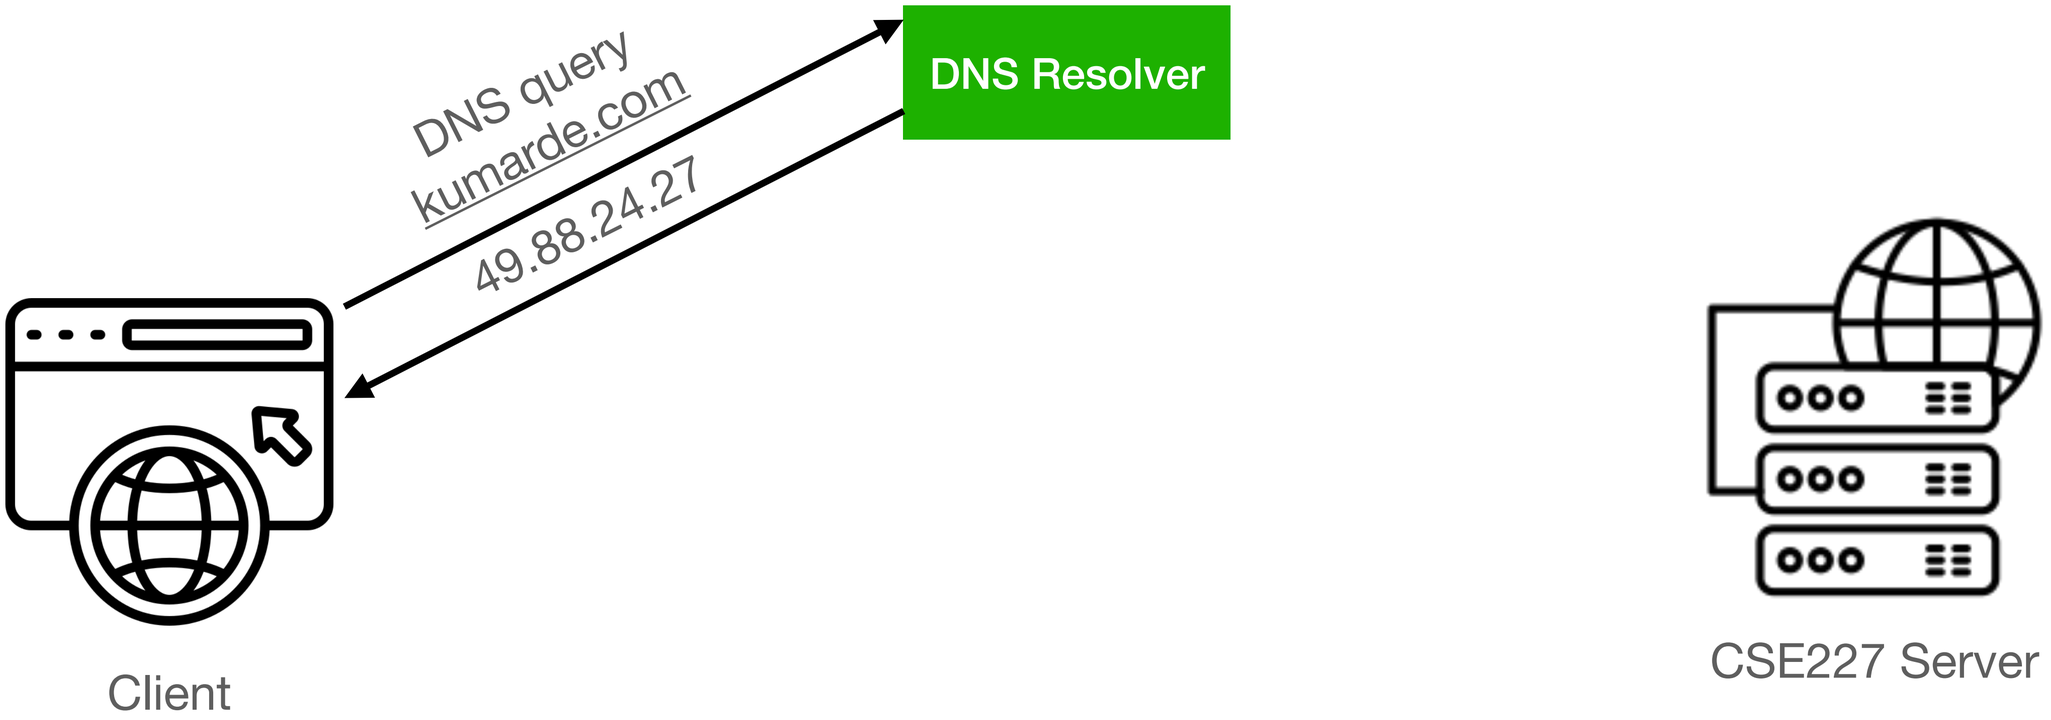

Censorship during an Internet connection

DNS Manipulation

DNS Resolver

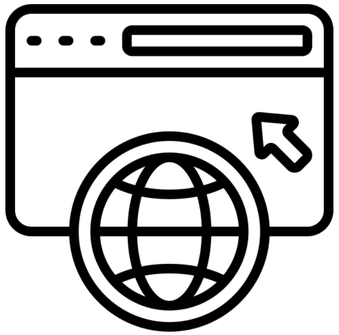

Client

CSE227 Server

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
  - What is DNS manipulation?
  - What is IP blocking?

Censorship during an Internet connection

IP Blocking

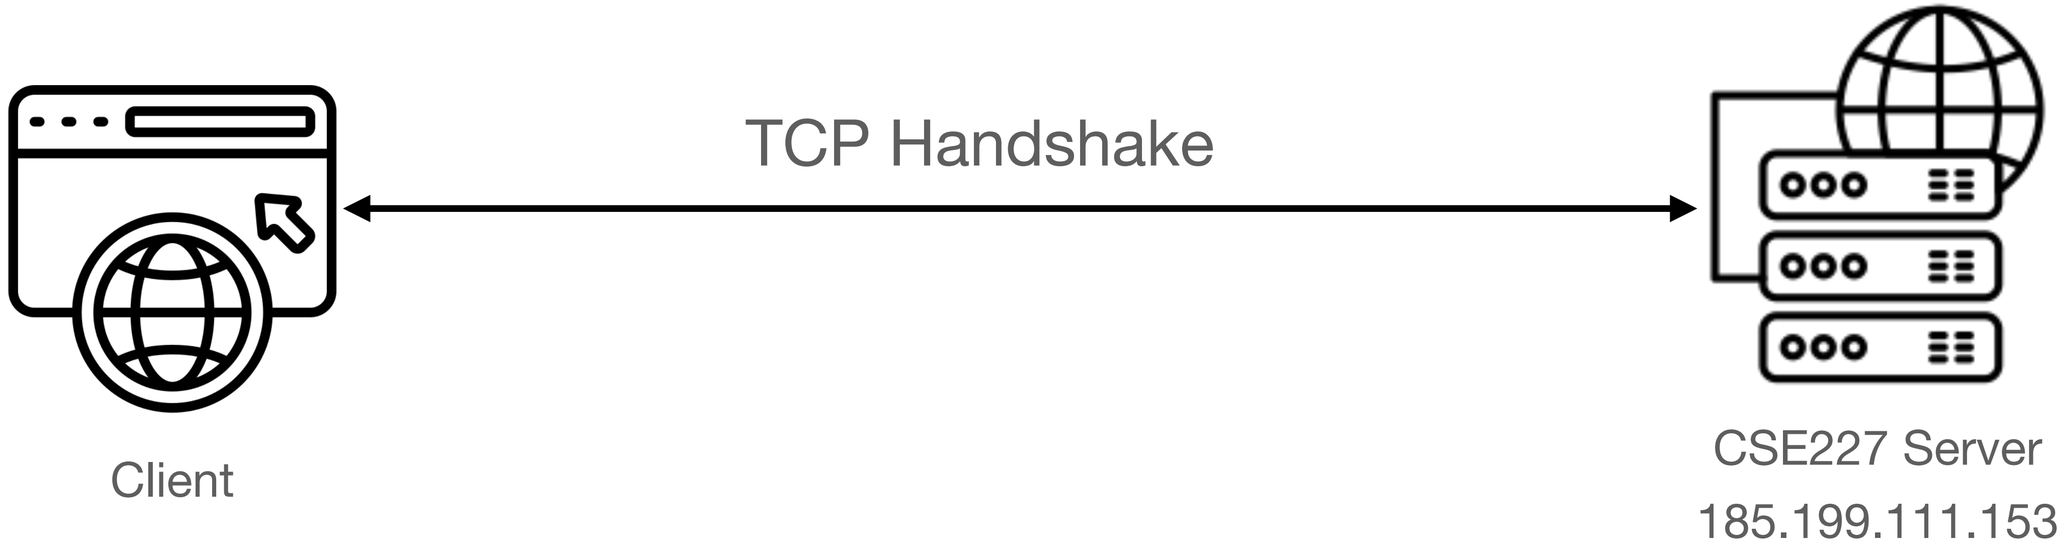

Censorship during an Internet connection

IP Blocking

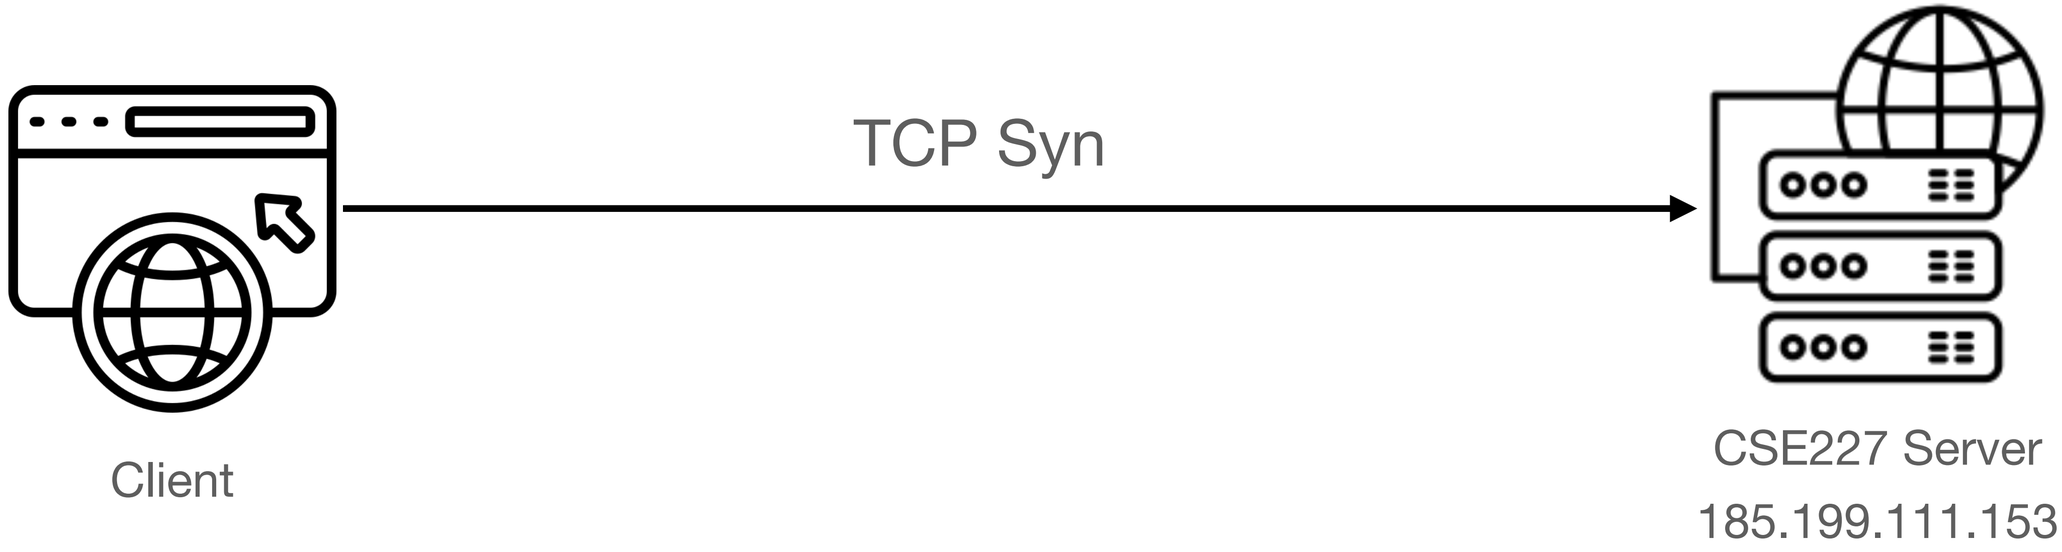

Censorship during an Internet connection

IP Blocking

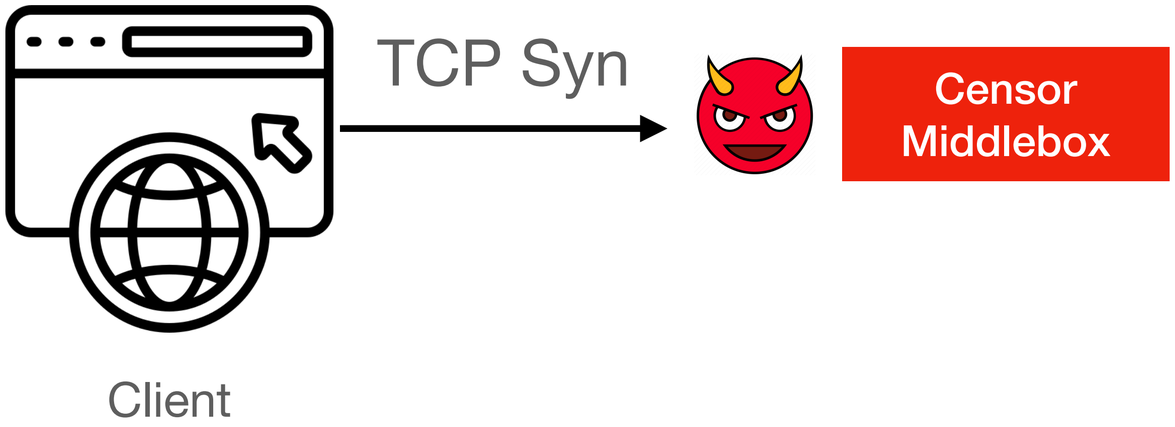

Censorship during an Internet connection

IP Blocking

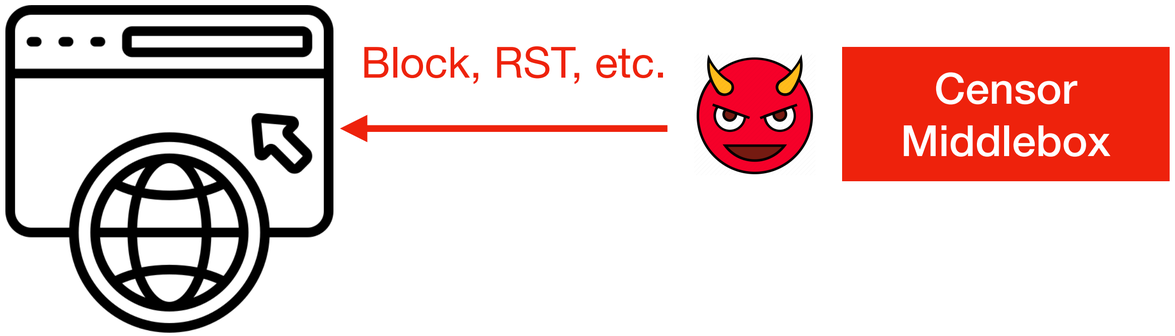

Censorship during an Internet connection

IP Blocking

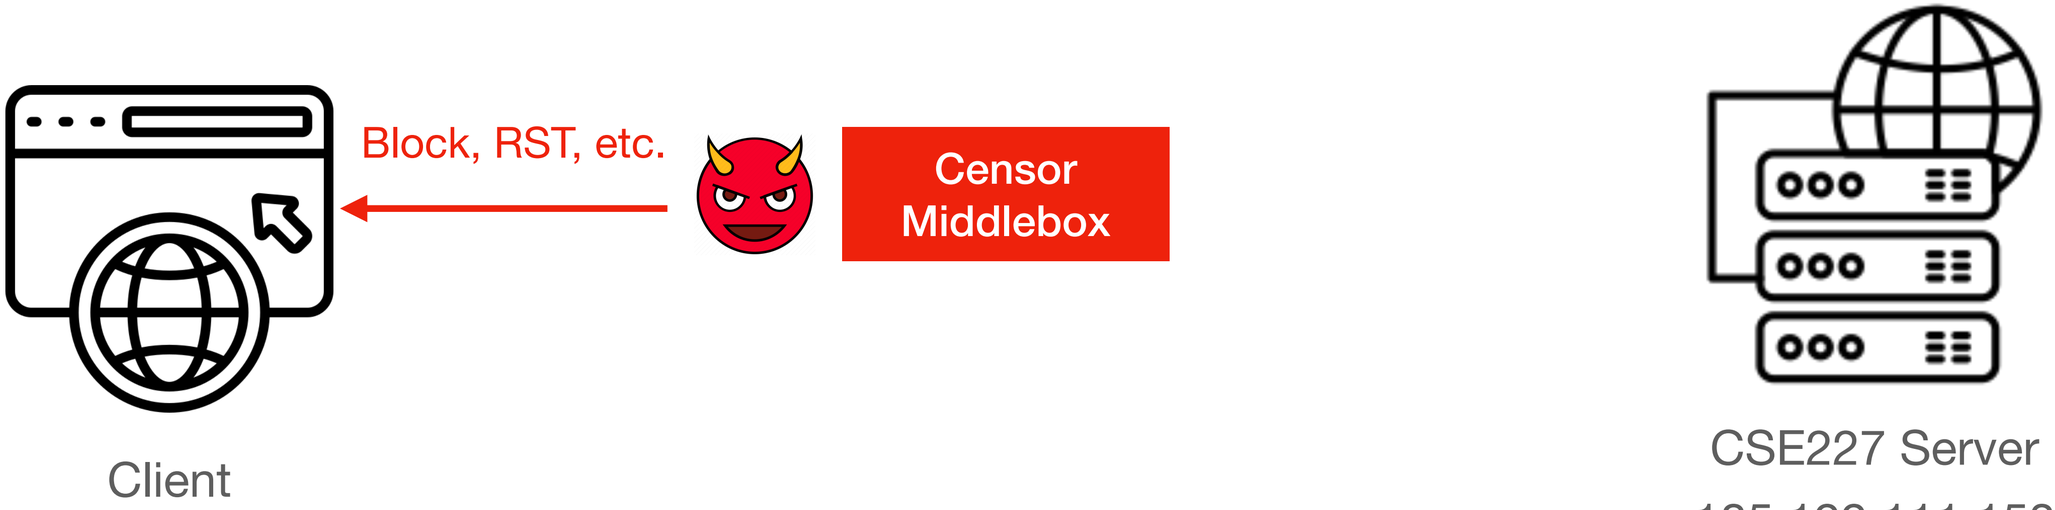

How easy do we think this is to implement?

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
  - What is DNS manipulation?
  - What is IP blocking?
  - What is application layer blocking?

Censorship during an Internet connection

Application Layer Blocking

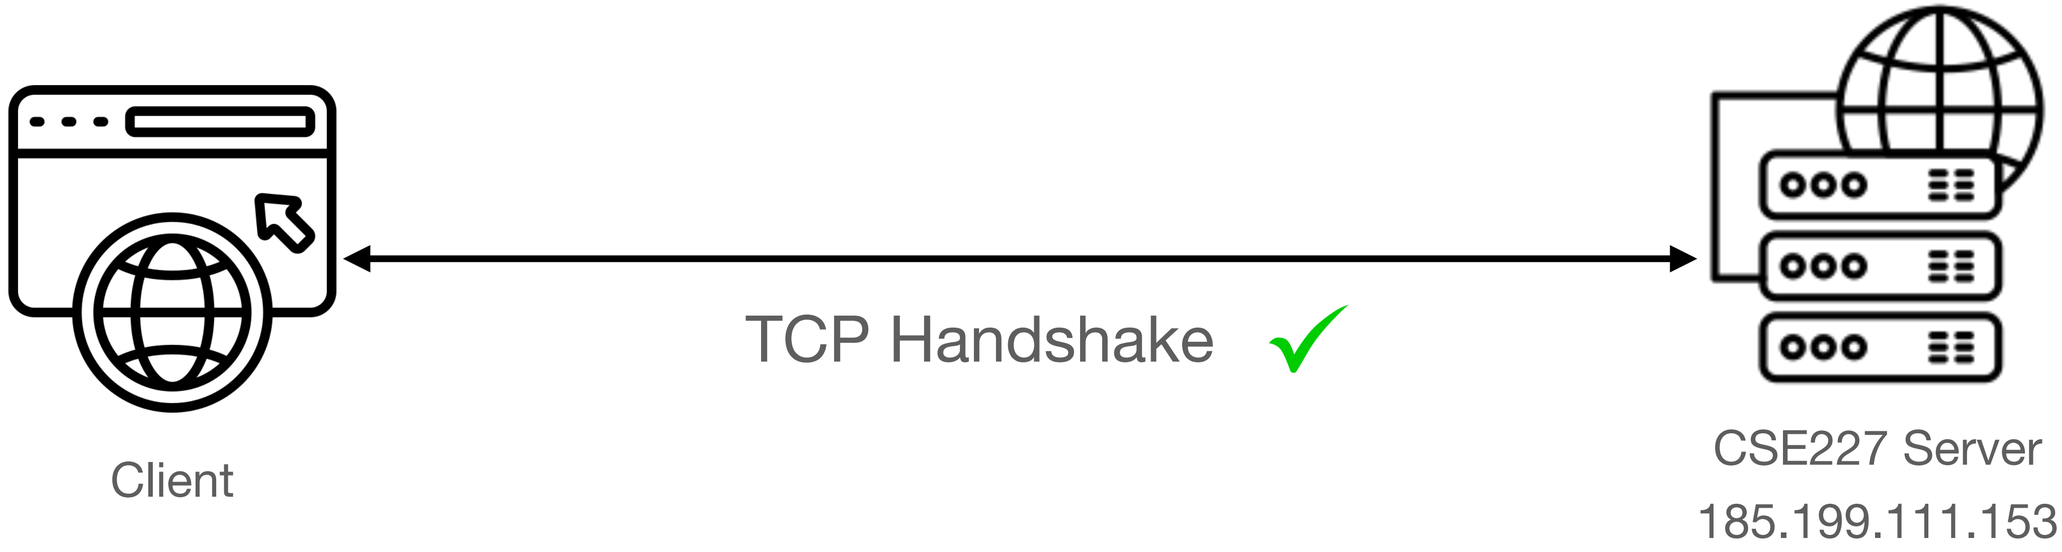

Censorship during an Internet connection

Application Layer Blocking

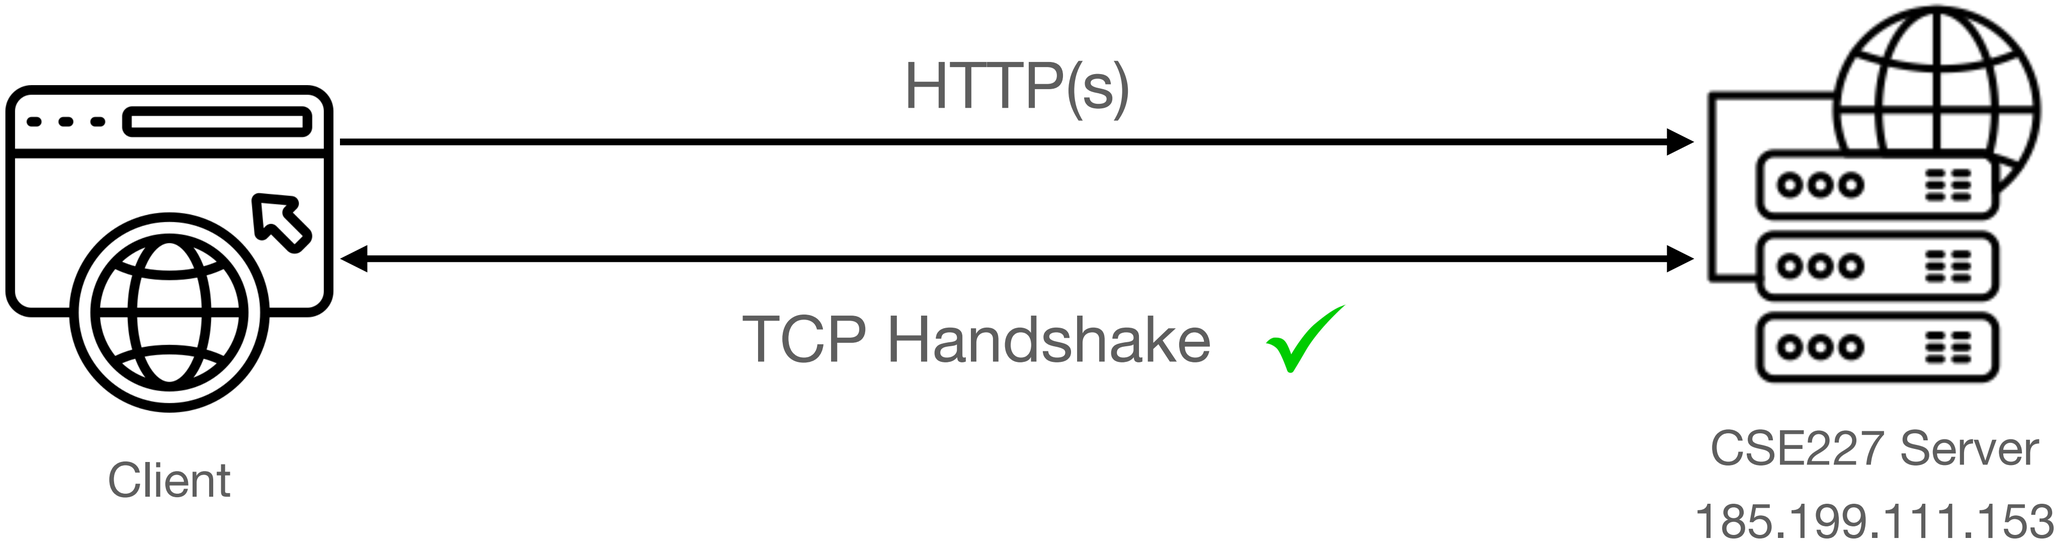

Censorship during an Internet connection

Application Layer Blocking

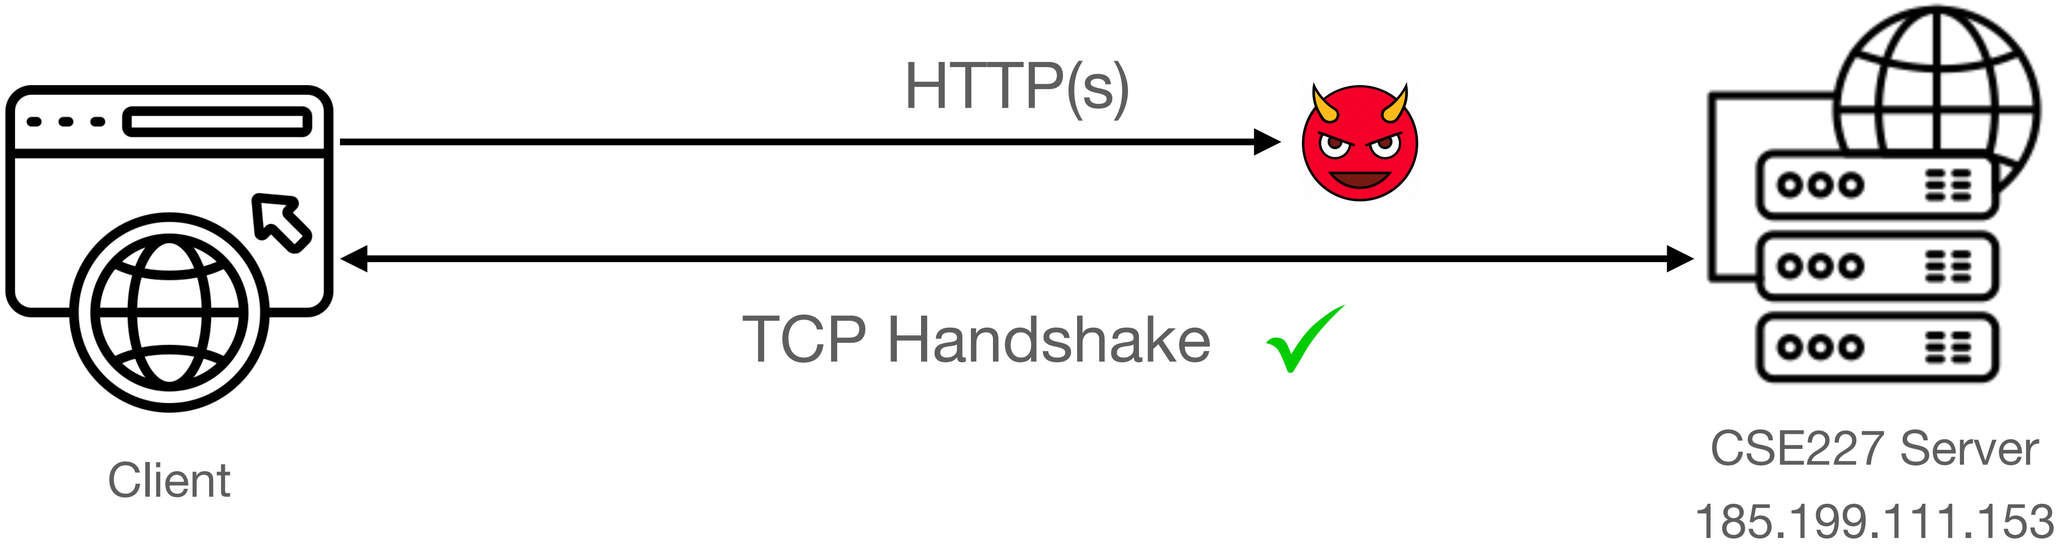

Censorship during an Internet connection

Application Layer Blocking

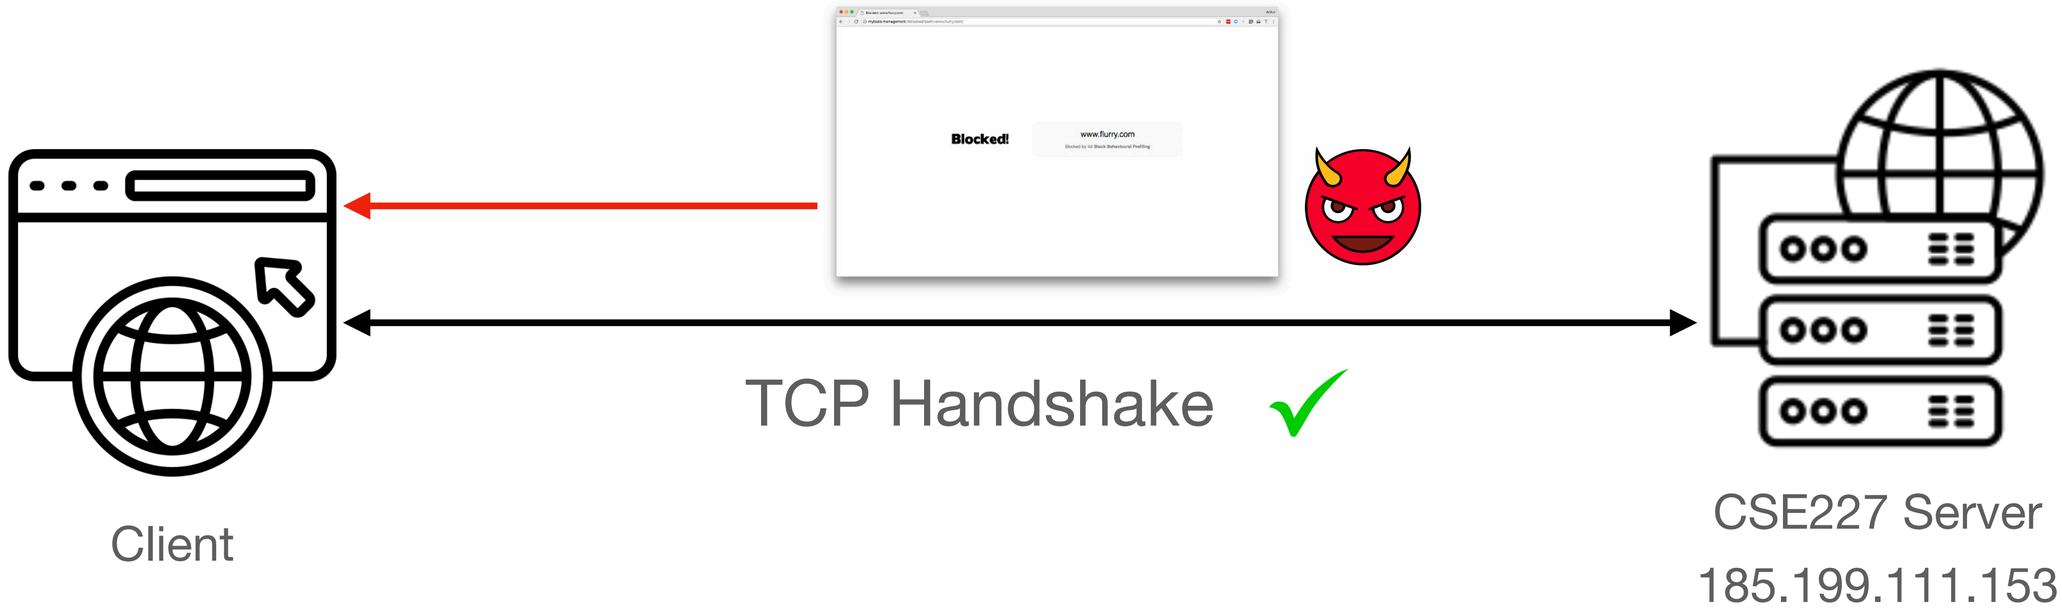

How easy do we think this is to implement?

Censorship during an Internet connection

Application Layer Blocking

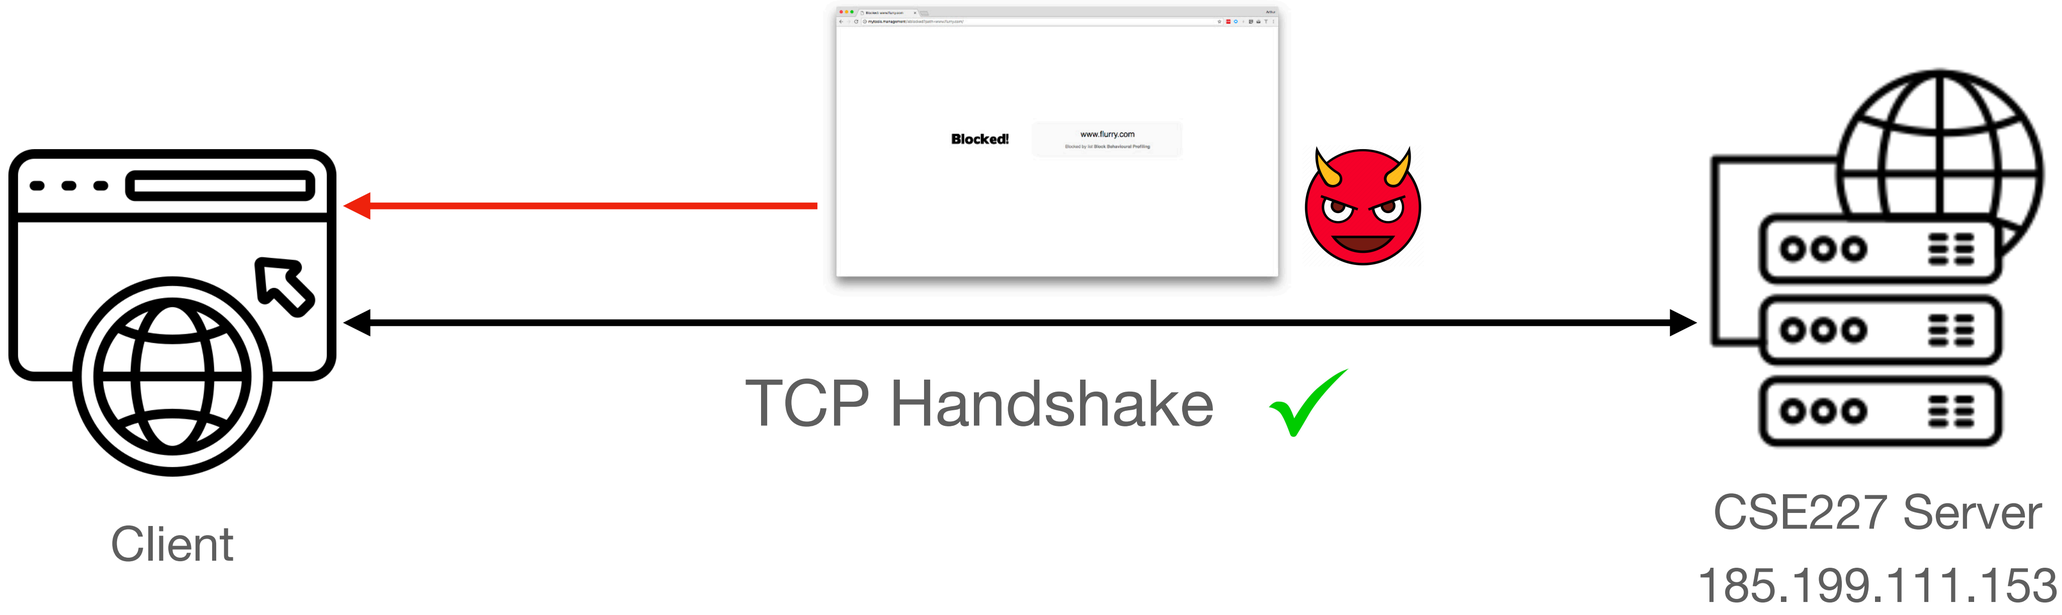

Censored Planet

Why measure censorship?

Why measure censorship?

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

Building a censorship observatory

- The goal of Censored Planet, or a system like it, is to build and provide a censorship observatory
  - Why do we need a censorship observatory?

Building a censorship observatory

- The goal of Censored Planet, or a system like it, is to build and provide a censorship observatory
  - Why do we need a censorship observatory?
- Censored planet claims to “use network side-channels to efficiently and remotely detect network anomalies… without relying on dedicated probing infrastructure in the field”
  - They draw a distinction between what they do, and volunteer censorship measurement
  - What is volunteer censorship measurement?

Volunteer Censorship Measurement in Action

OONI

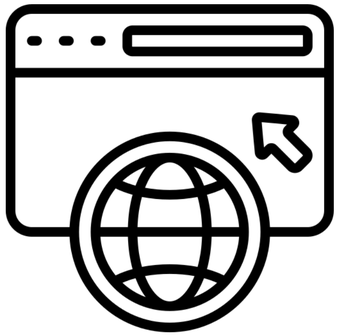

Volunteer client
in-country

Censor

Server of Interest

Control Relay

Volunteer Censorship Measurement in Action

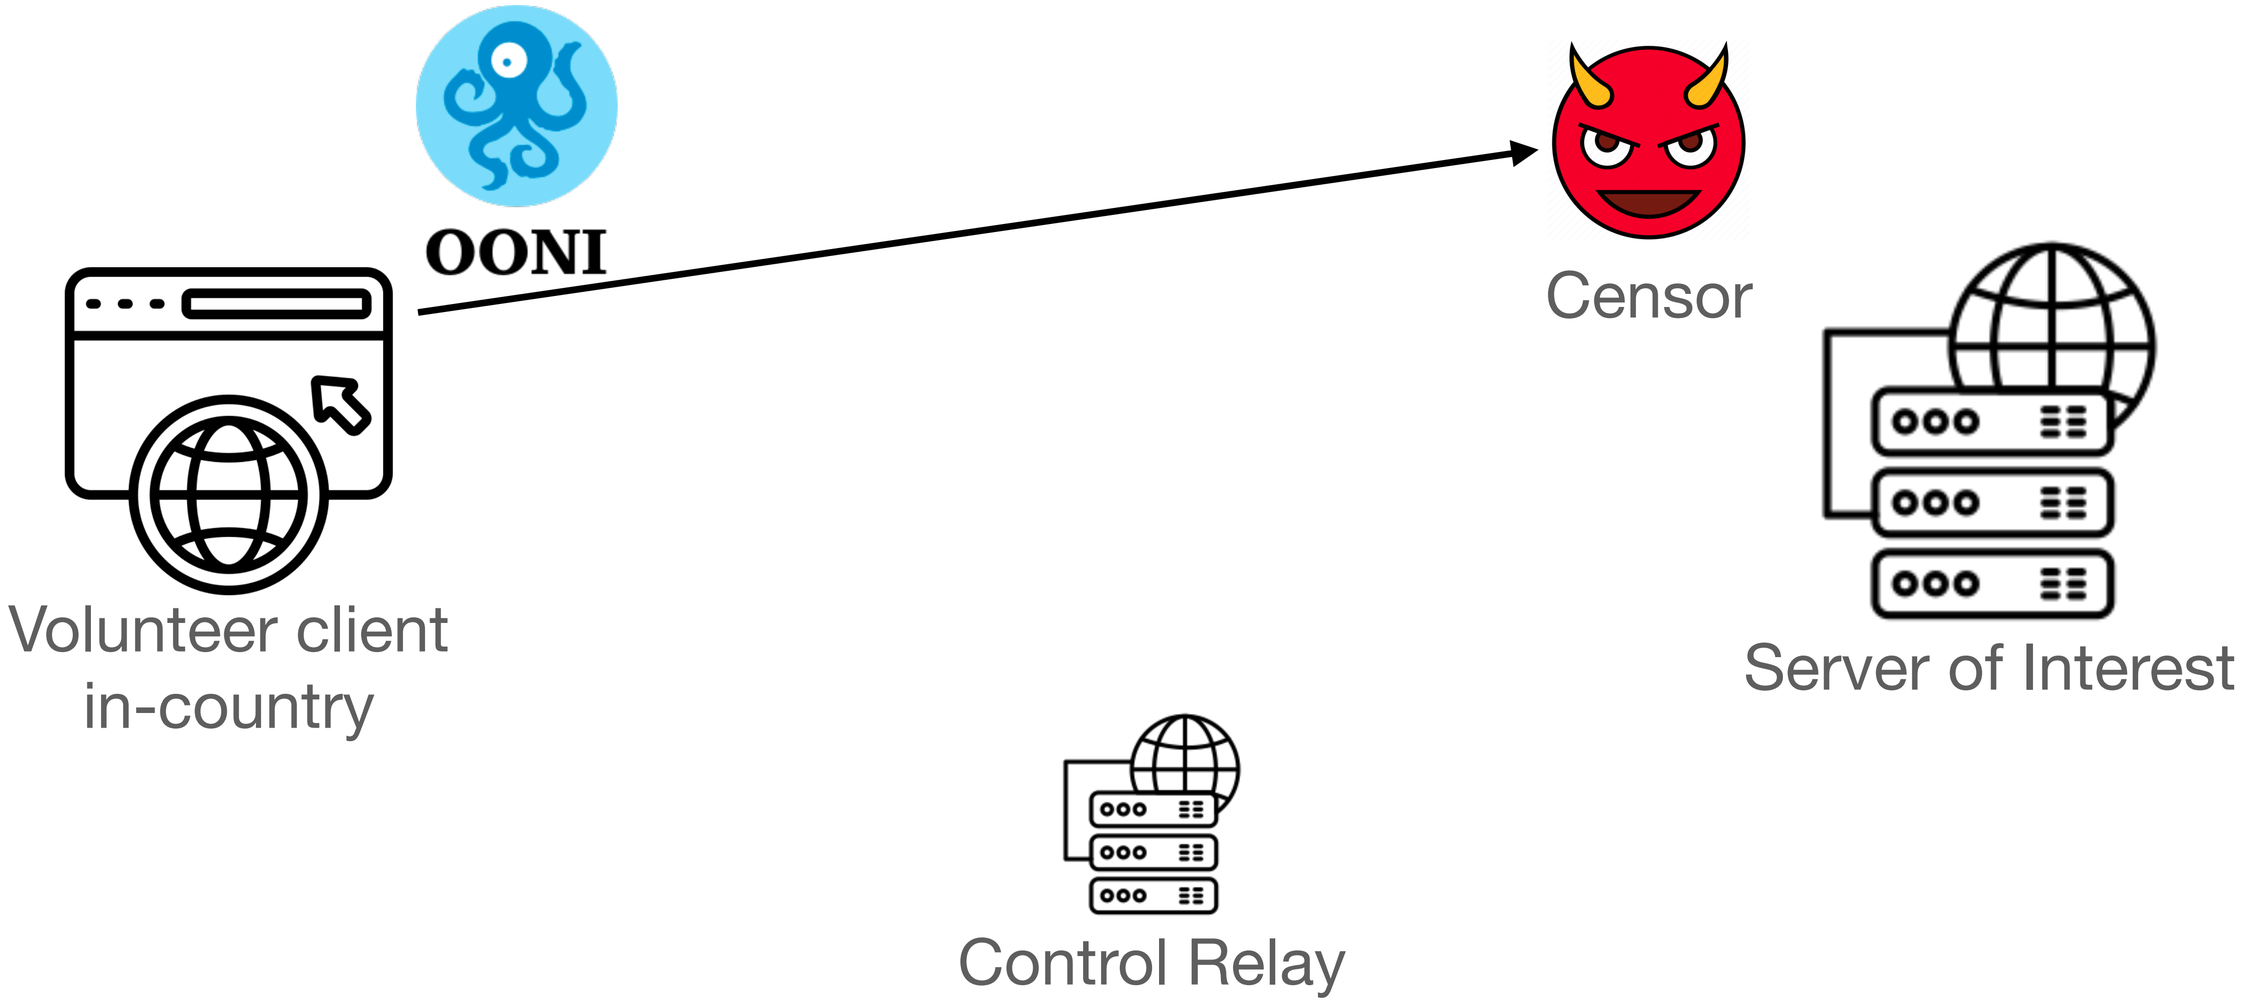

Volunteer Censorship Measurement in Action

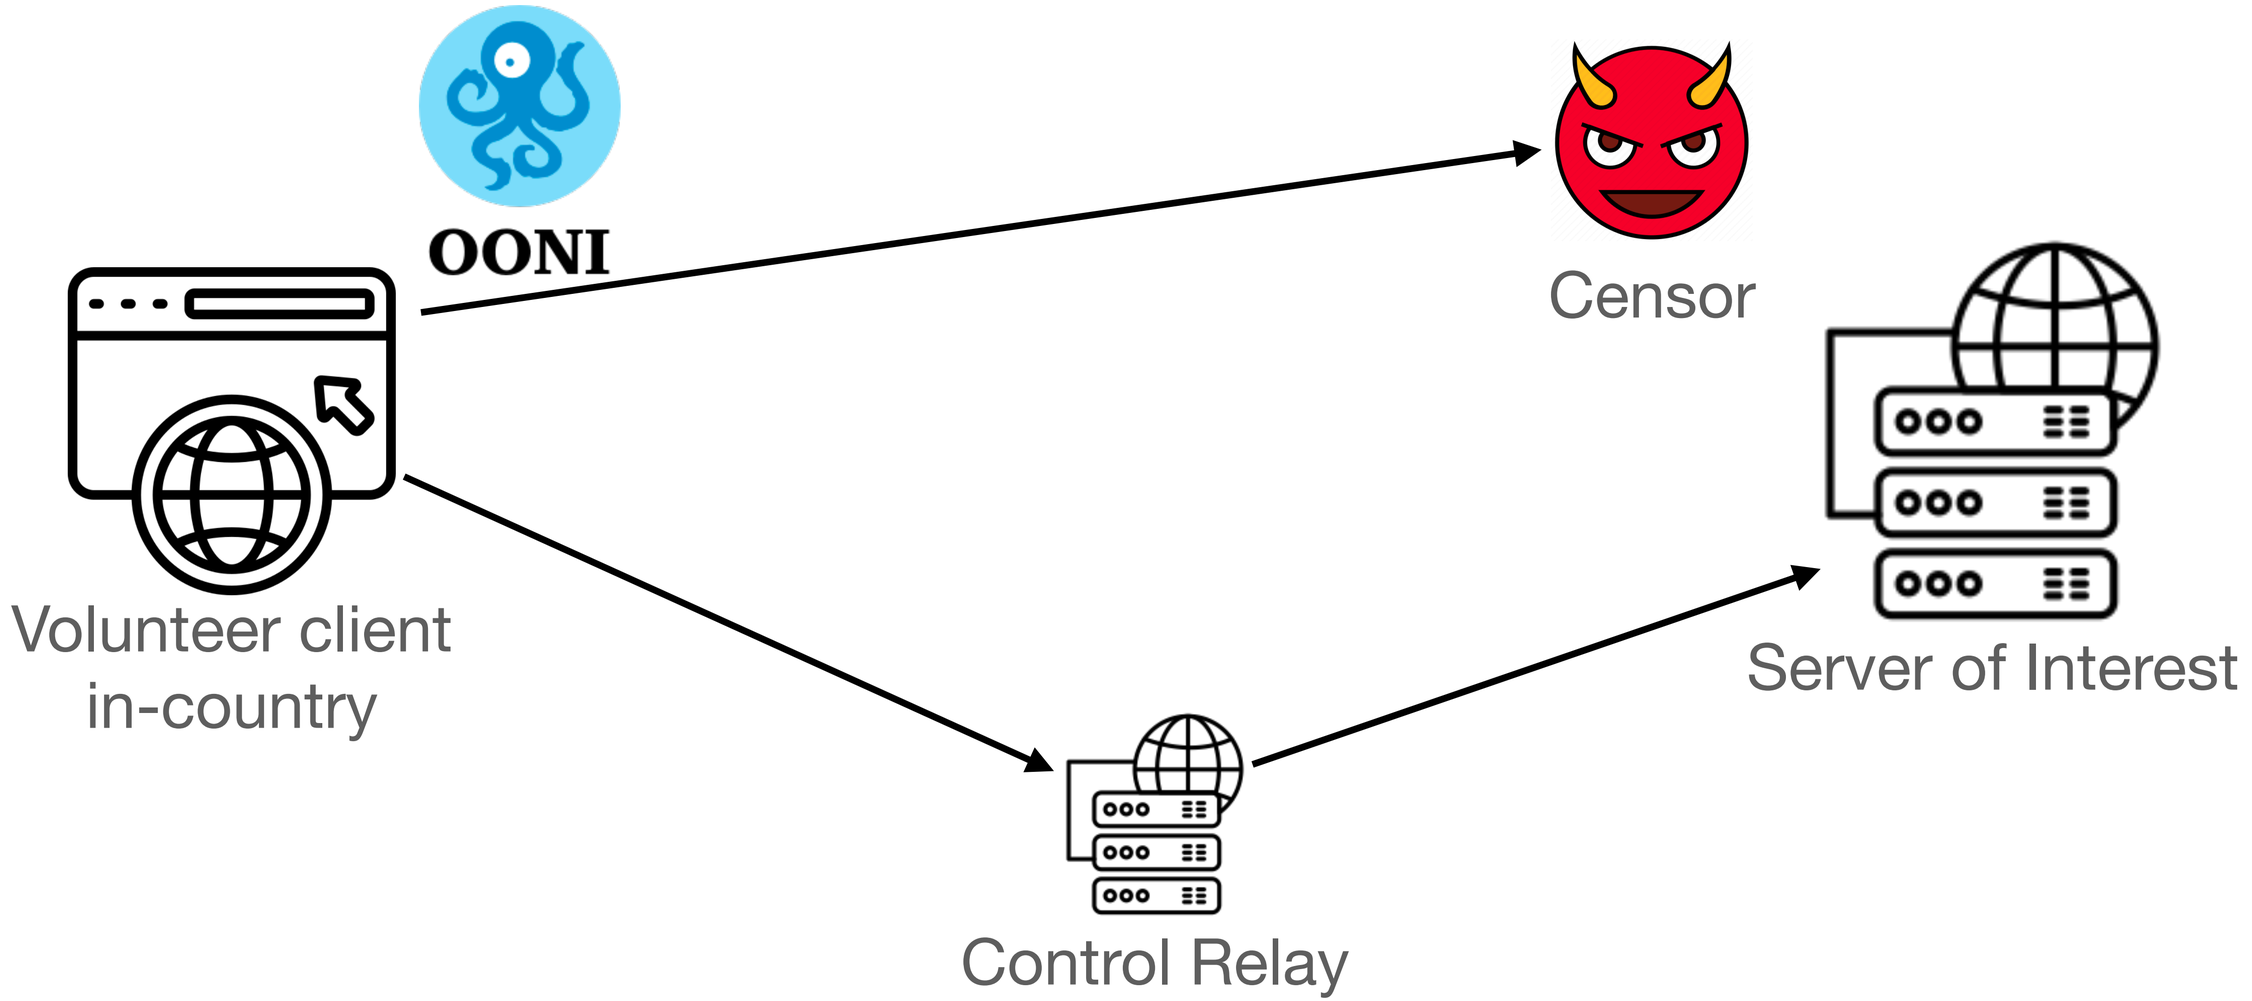

Limitations of Volunteer Measurement

- What are the scale limitations of volunteer measurements?
- What are the coverage limitations of volunteer measurements?
- What are the continuity limitations of volunteer measurements?
- What are the *ethical* or *safety* limitations of volunteer measurements?

OONI

Enter: Censored Planet

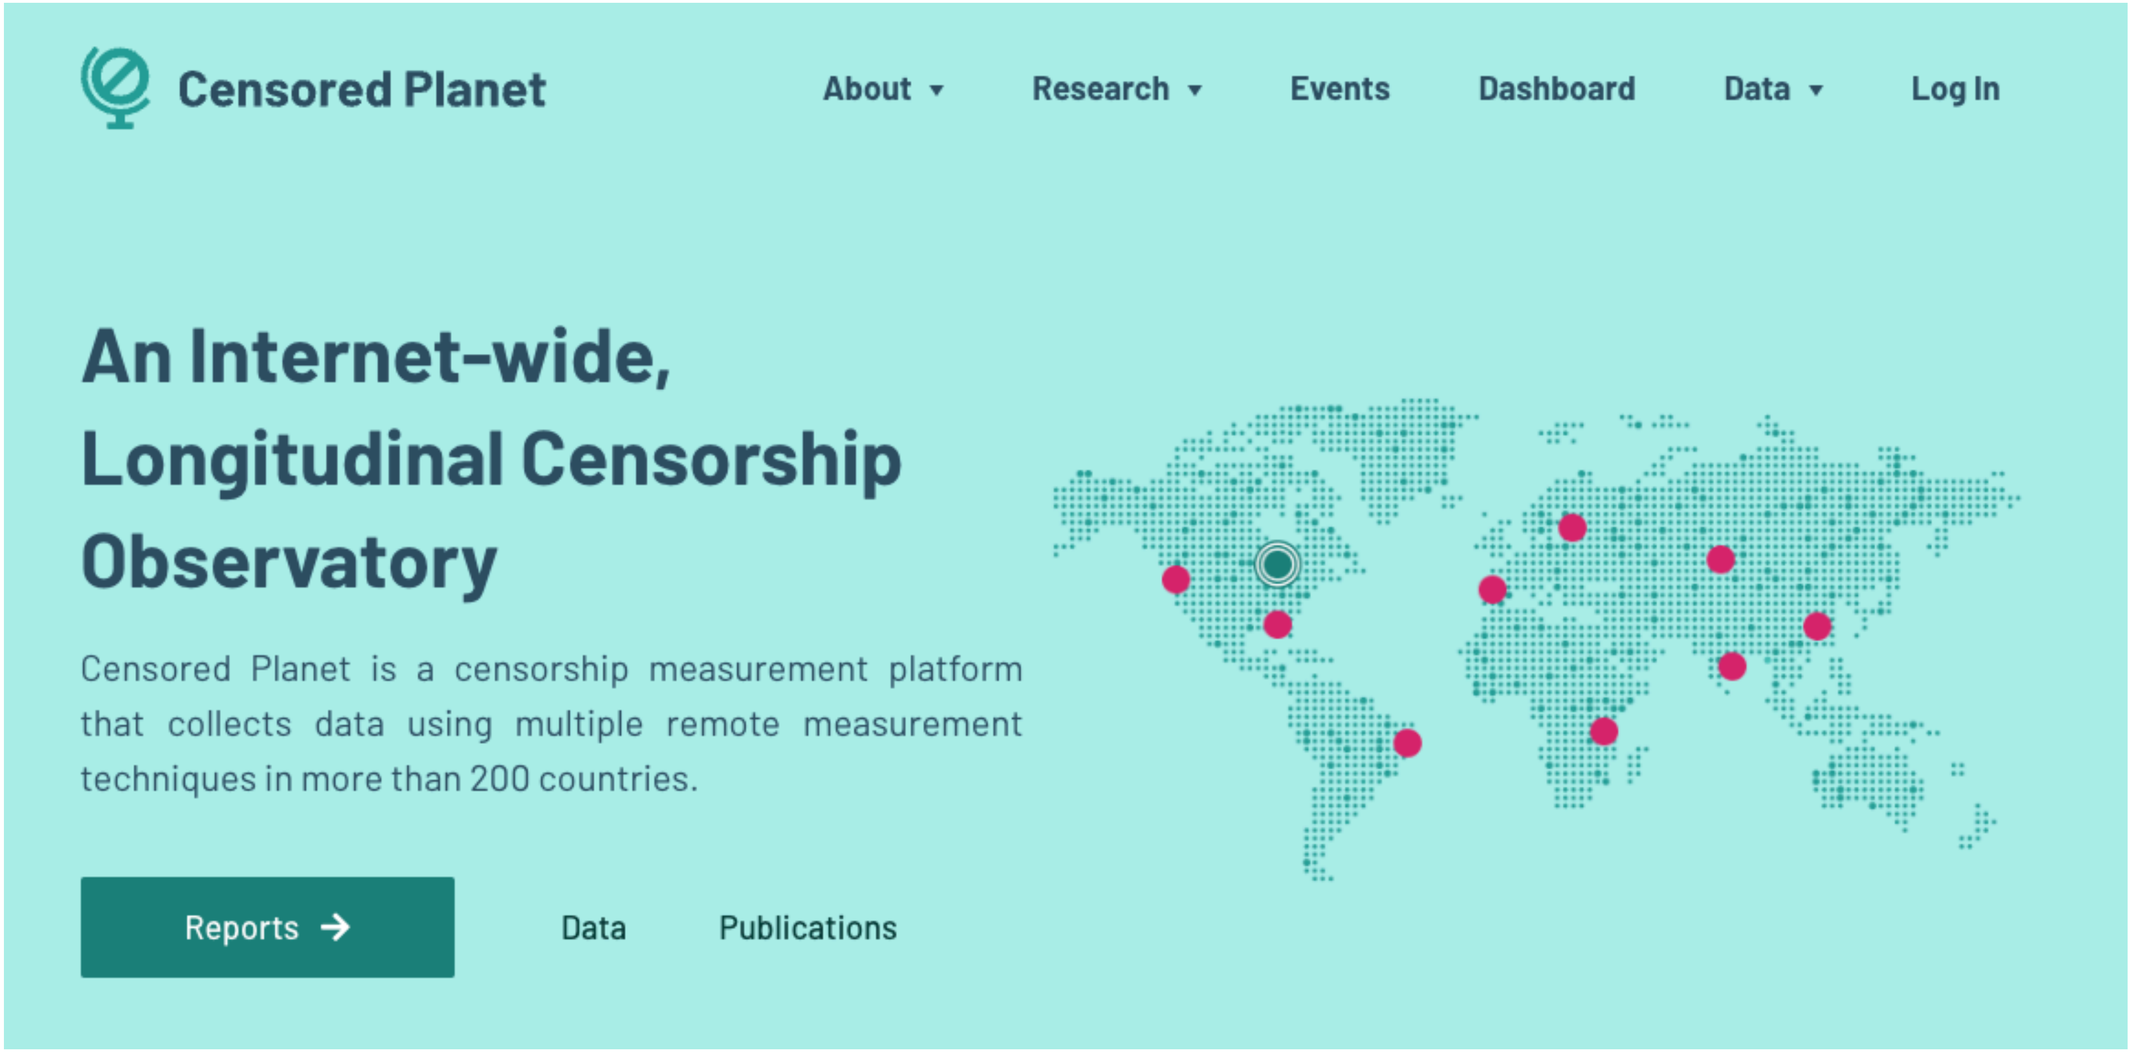

Censored Planet Observatory

Collects data using remote measurement techniques on **6 Internet protocols** (TCP/IP, DNS, HTTP, HTTPS, Echo, Discard)

Satellite

Hyperquack

Spooky Scan

Continuous baseline of reachability data for **2000 websites each week**

More than **95,000 vantage points** in **221 countries and territories**

45 billion

Measurements over 36 Months

221 countries

42%-360% increase compared to other platforms

8 ASes (median)/country

Median increase of 4-7 ASes per country

Censored Planet Broad Strategy

- Distinguishes between responses between “control” servers and “test” servers
    - What are control servers and test servers? How do they help us determine network censorship?

Censored Planet Broad Strategy

- Distinguishes between responses between “control” servers and “test” servers
  - What are control servers and test servers? How do they help us determine network censorship?
- What are some challenges or limitations in using this strategy broadly?

Using Echo (Quack) to Measure App-Layer Blocking

Network Working Group
Request for Comments: 862

J. Postel
ISI
May 1983

Echo Protocol

This RFC specifies a standard for the ARPA Internet community. Hosts on the ARPA Internet that choose to implement an Echo Protocol are expected to adopt and implement this standard.

A very useful debugging and measurement tool is an echo service. An echo service simply sends back to the originating source any data it receives.

TCP Based Echo Service

One echo service is defined as a connection based application on TCP. A server listens for TCP connections on TCP port 7. Once a connection is established any data received is sent back. This continues until the calling user terminates the connection.

UDP Based Echo Service

Another echo service is defined as a datagram based application on UDP. A server listens for UDP datagrams on UDP port 7. When a datagram is received, the data from it is sent back in an answering datagram.

Using Echo (Quack) to Measure App-Layer Blocking

Client

Server of Interest

Using Echo (Quack) to Measure App-Layer Blocking

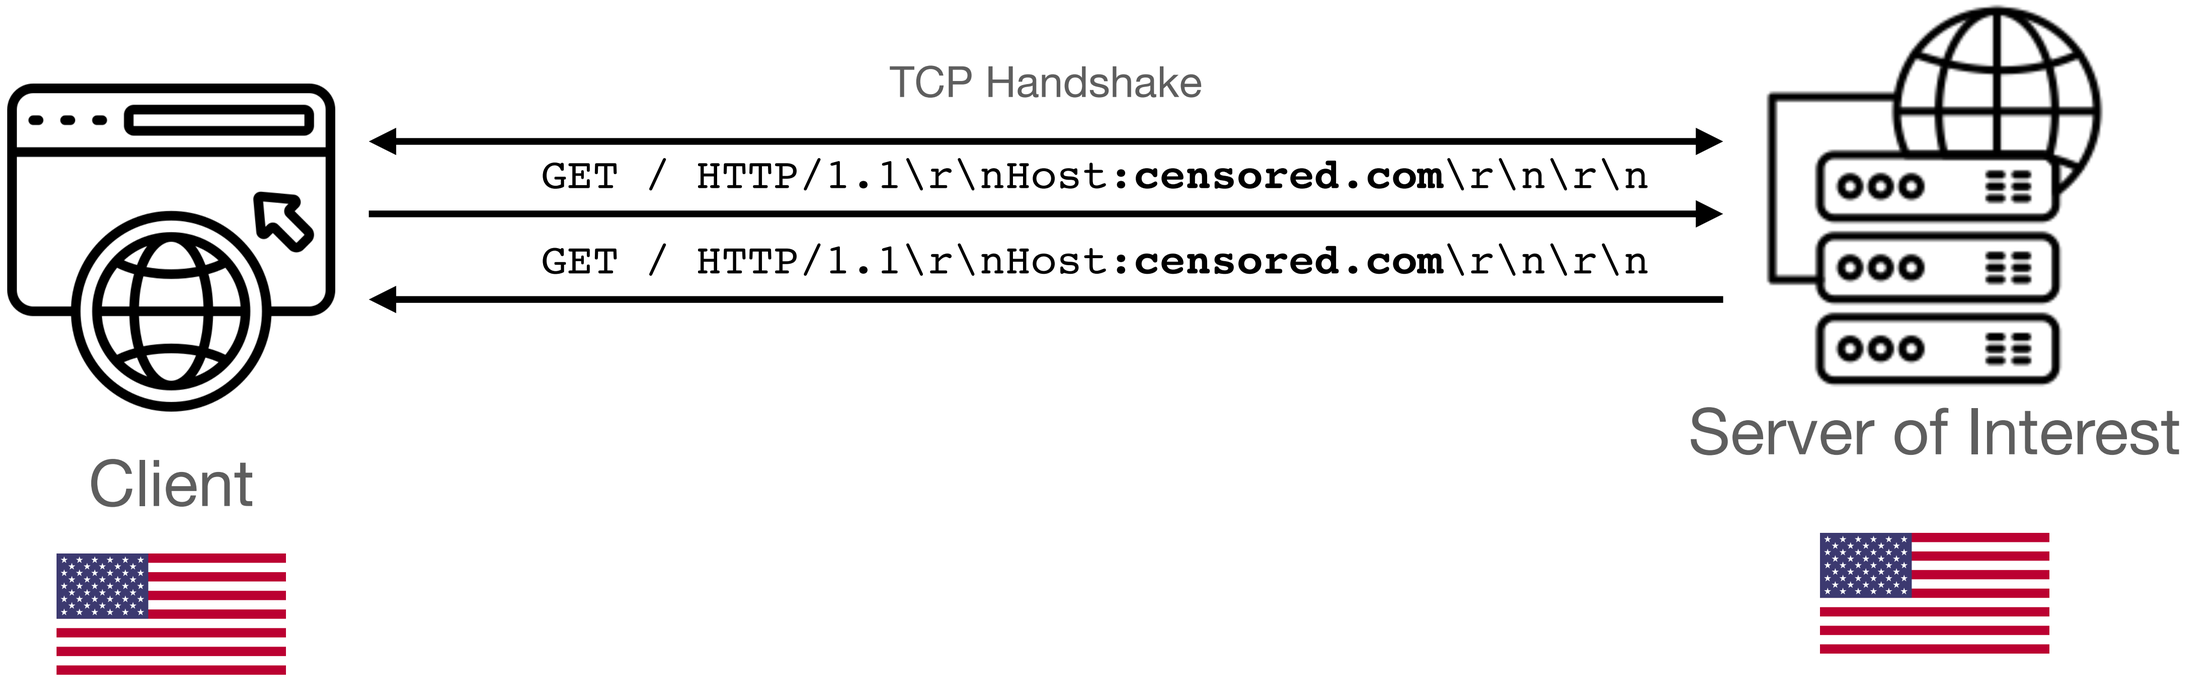

Using Echo (Quack) to Measure App-Layer Blocking

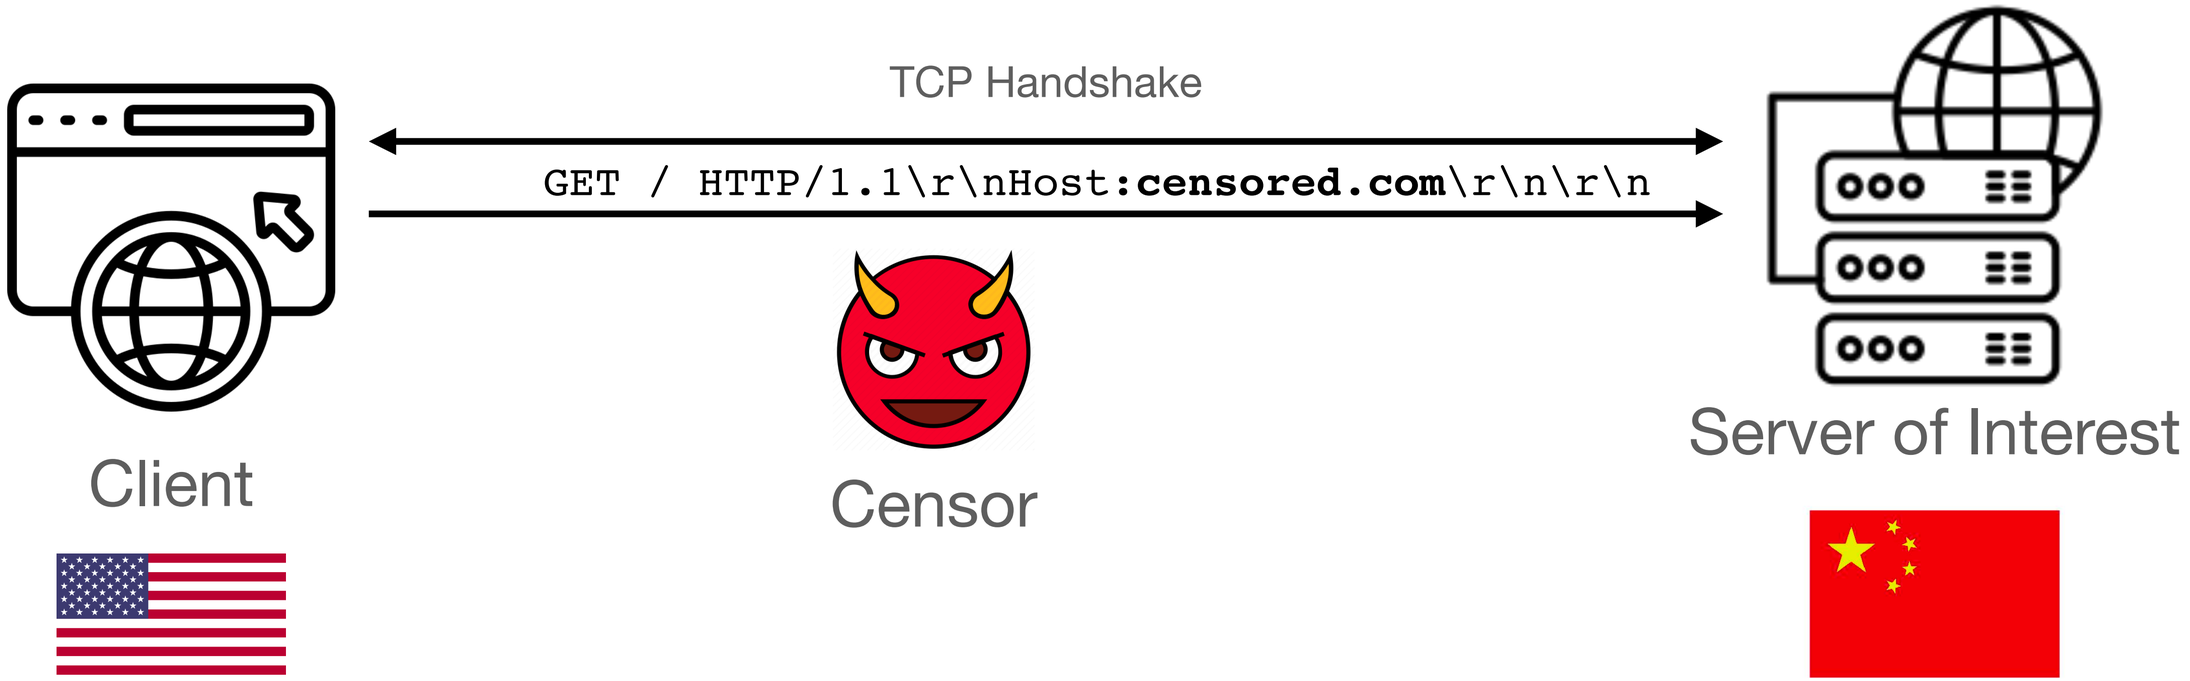

What will happen in the presence of censorship?

Using Echo (Quack) to Measure App-Layer Blocking

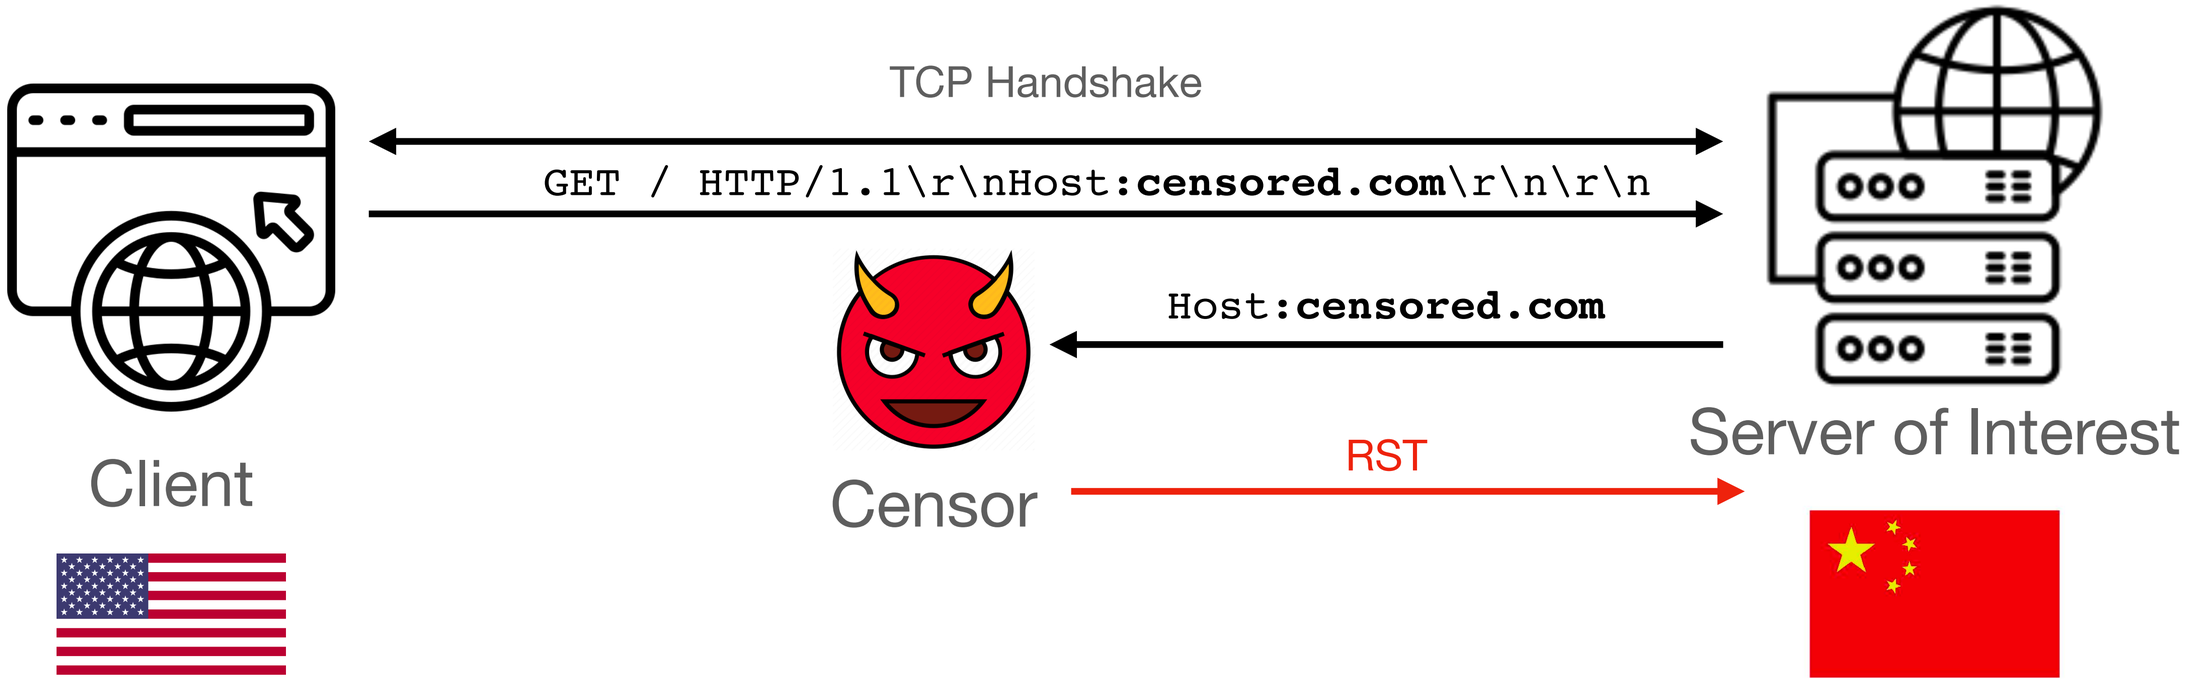

Meta Thoughts on this Paper

- What did we think about this paper?
- Do we think censorship measurement is important? Why or why not?
- *Who is this research for?*

Break Time + Attendance

https://tinyurl.com/cse227-attend

Codeword:
This-Code-Is-Censored

Throttling Twitter

The Circumstances

- This paper is about Russia's throttling of Twitter's service in 2021 to a **country-wide** rate-limit
  - What was the primary reason Russia did this?

The Circumstances

- This paper is about Russia's throttling of Twitter's service in 2021 to a **country-wide** rate-limit
  - What was the primary reason Russia did this?

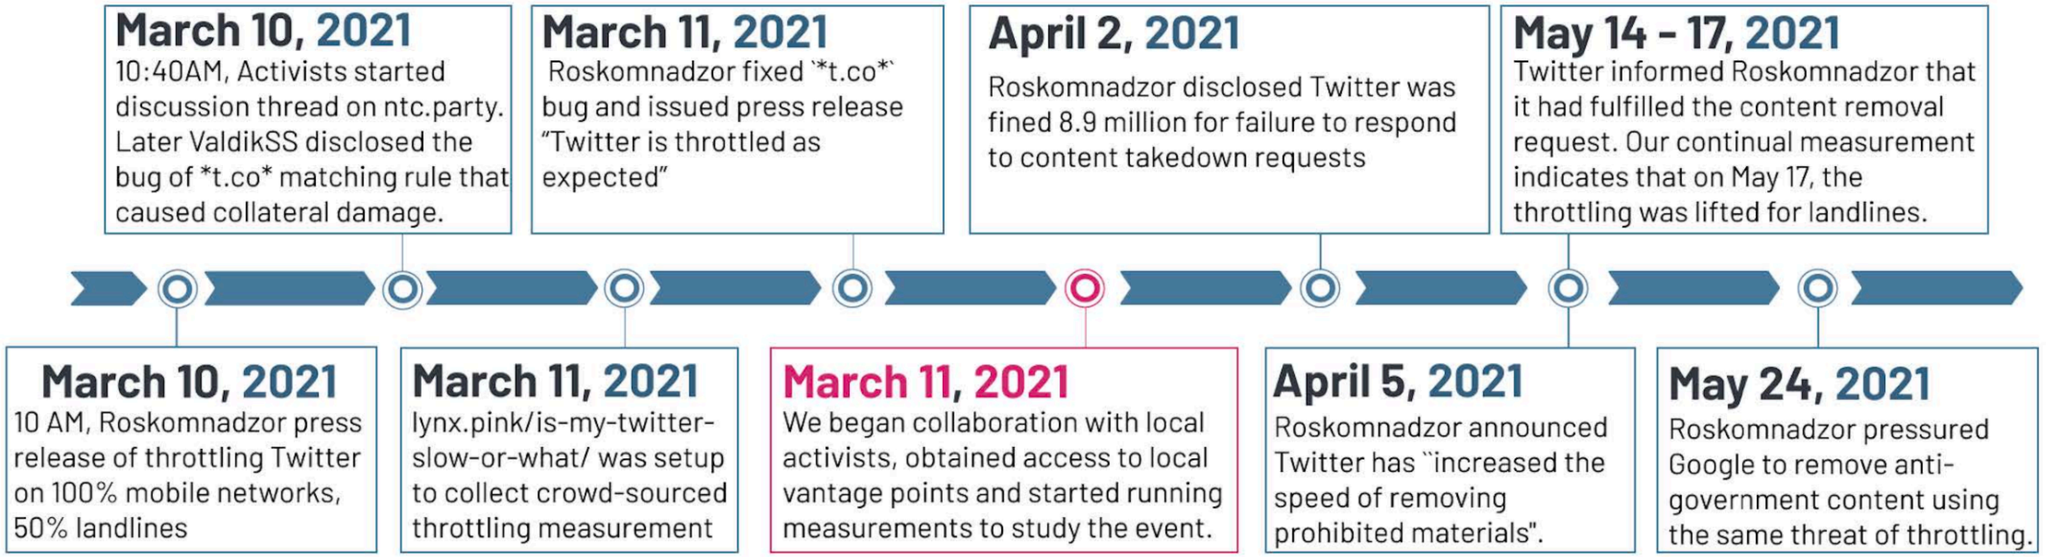

Measuring the Problem

- How did the authors primarily measure the problem?

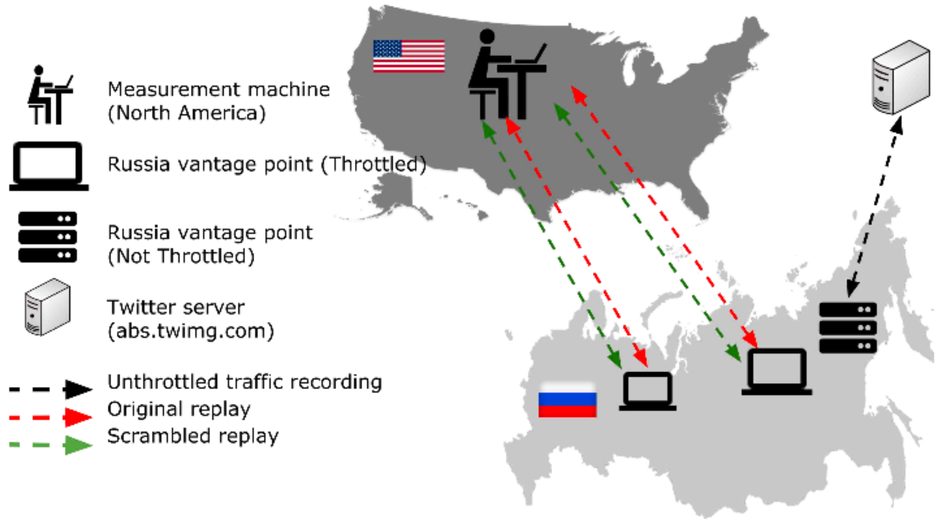

Measuring the Problem

- How did the authors primarily measure the problem?
  - Authors secured 8 local vantage-points, 7 of which experienced throttling, one of which served as a control (luck)

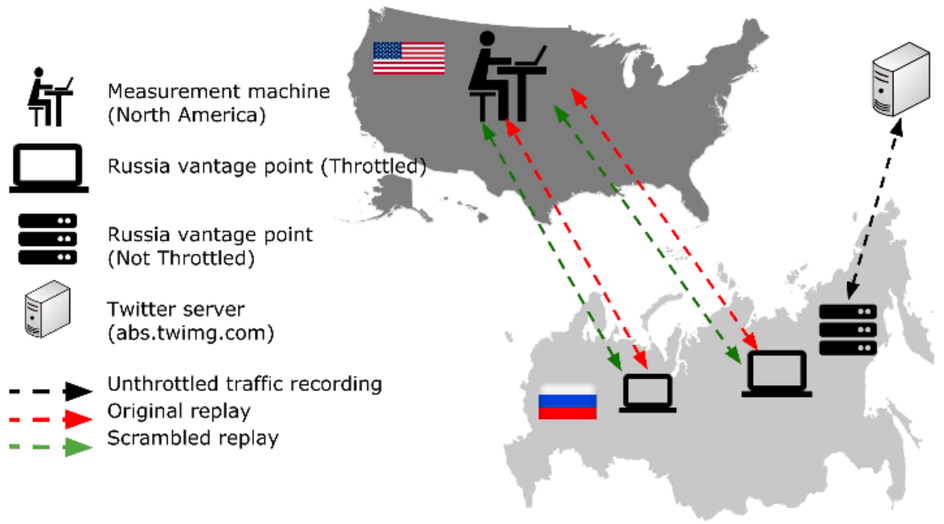

Measuring the Problem

- How did the authors primarily measure the problem?
  - Authors secured 8 local vantage-points, 7 of which experienced throttling, one of which served as a control (luck)
- What is a “Record and Replay” measurement system?

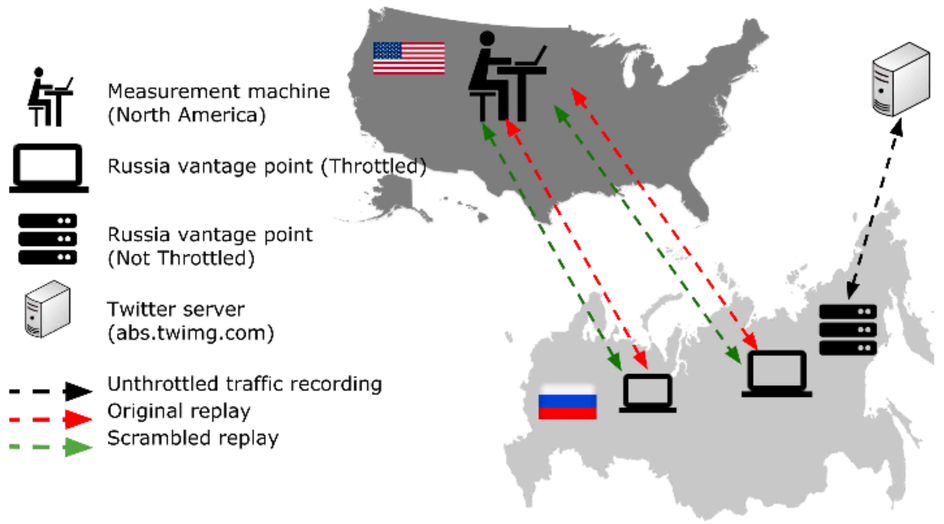

Measuring the Problem

- How did the authors primarily measure the problem?
  - Authors secured 8 local vantage-points, 7 of which experienced throttling, one of which served as a control (luck)
- What is a “Record and Replay” measurement system?
  - Authors record an unthrottled session, and replay it on their server to identify app-level blocking

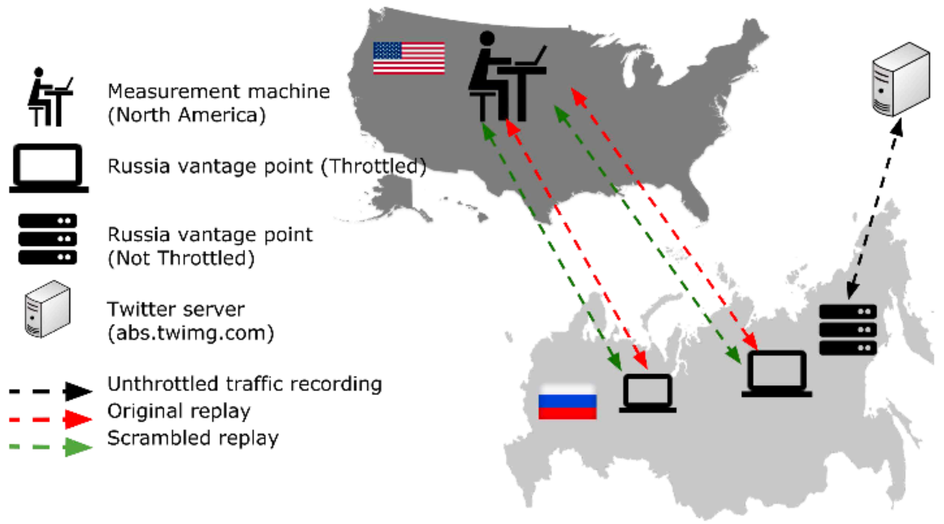

What they found

- Throttler uses traffic policing: drop packets that are above a particular rate limit
- Throttler parses network packets from both directions and throttled based upon observing sensitive SNI in the TLS Client Hello
  - How can Russia do this? Isn't TLS encrypted?
- Throttling was initially implemented poorly
  - They wanted to block t.co, but they also blocked microsof**t.co**m, and reddi**t.co**m
- Throttling happened close to the user, suggesting **nationwide** control near the **edge** (scary!)

Discussion

- What surprised you in this paper?
- What can we do as a result of this paper?
- For more: see **censorship circumvention** (e.g., Telex, Refractive Networking, etc.)

Next time…

- No lecture! Sign up for a meeting with me if you wish
- See you next Tuesday! Work on your projects!