Internet Censorship

cs249i: The Modern Internet

Stanford
University

whoami

Internet Censorship practices are on the rise

Internet Censorship

How do censors restrict access?

1. Internet shutdowns

155 political shutdowns in 29 countries in 2020

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

Internet Censorship

How do censors restrict access?

1. Internet shutdowns
2. Throttling

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

Internet Censorship

How do censors restrict access?

1. Internet shutdowns
2. Throttling
3. Content Takedowns

Google received ~300K government takedown requests since 2011

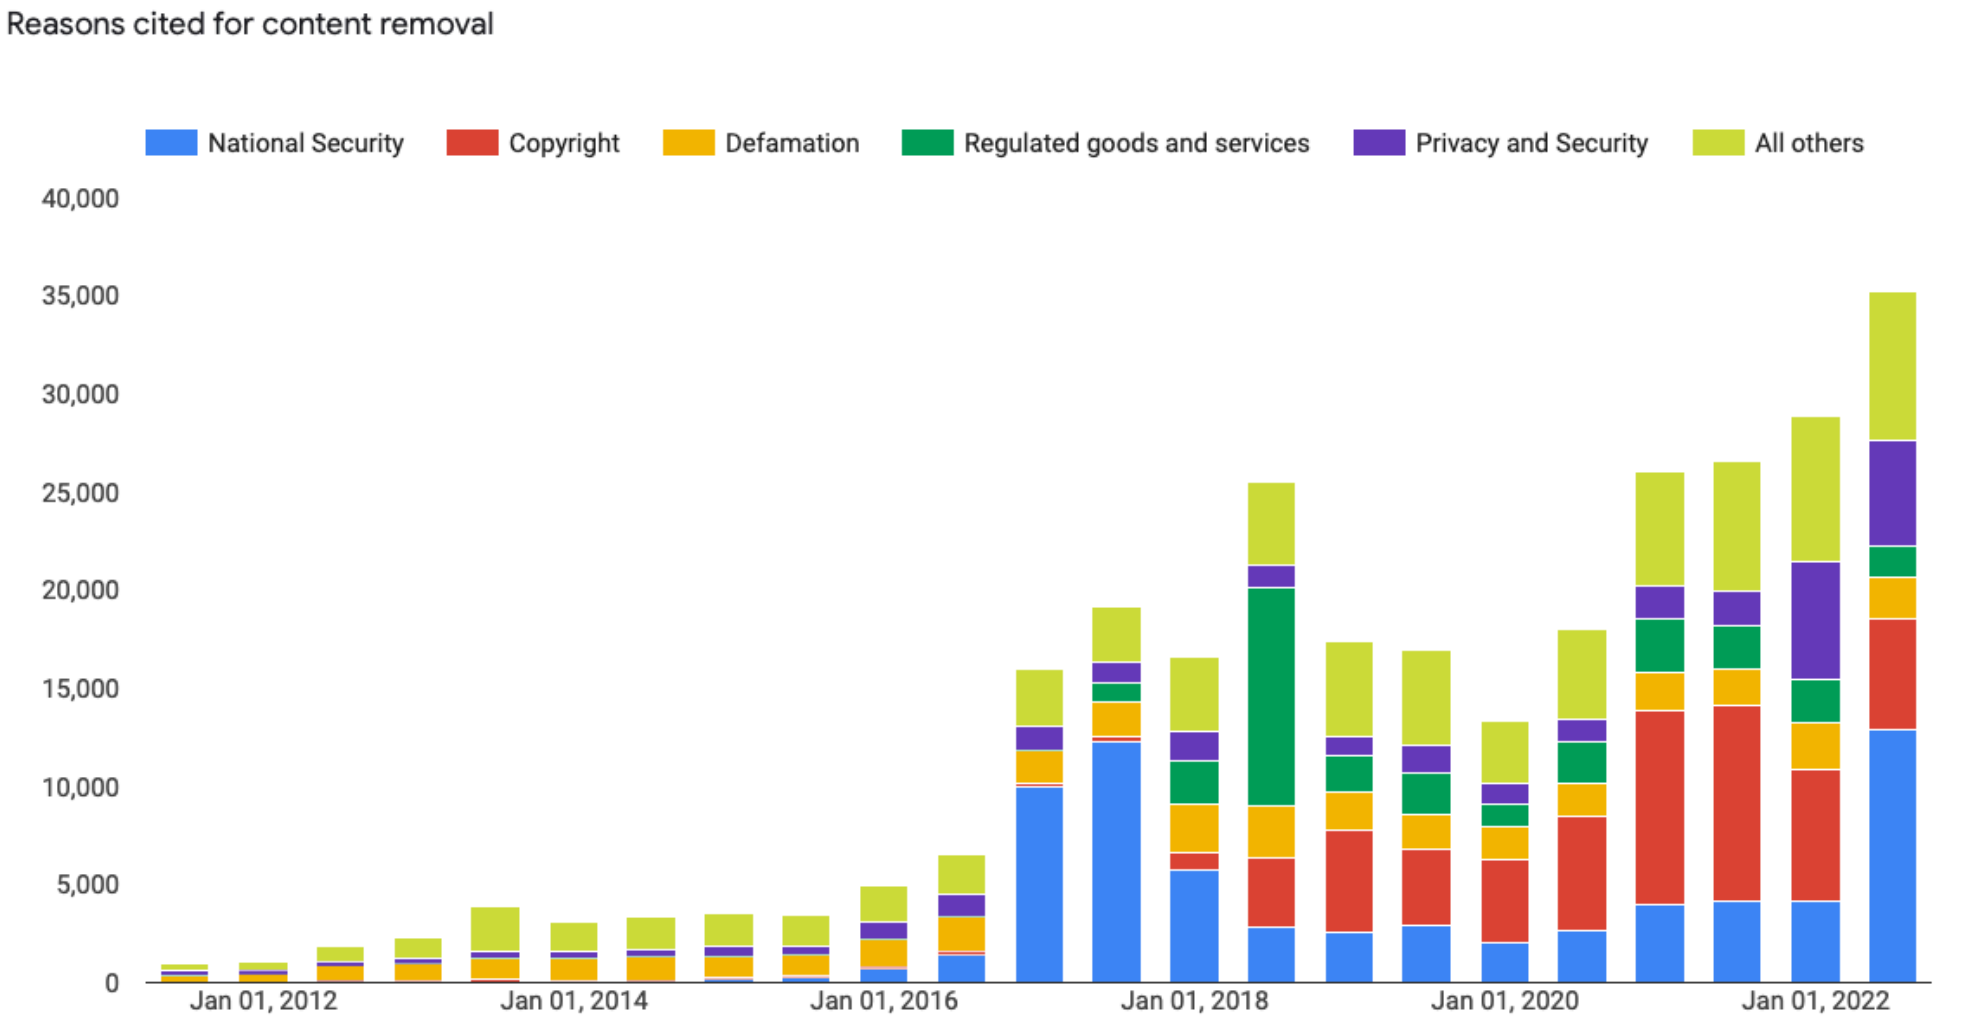

https://transparencyreport.google.com/government-removals/government-requests?hl=en_GB

Internet Censorship

How do censors restrict access?

1. Internet shutdowns
2. Throttling
3. Content Takedowns
4. Website blocking

https://dataprot.net/articles/what-is-internet-censorship/

Censorship during an Internet connection

Modes of website blocking

Censorship during an Internet connection

Modes of website blocking

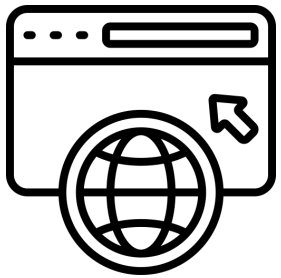

Client

CS249i Server

Censorship during an Internet connection

DNS manipulation

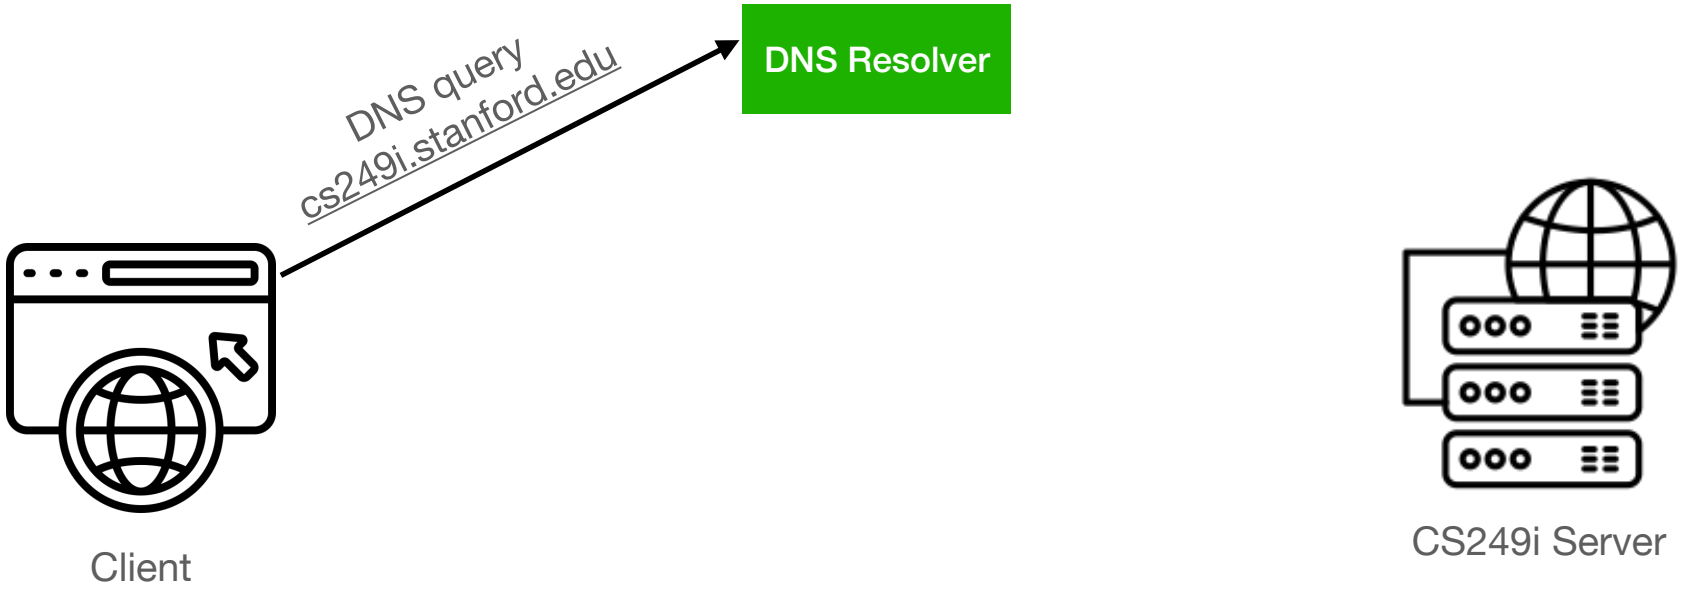

Censorship during an Internet connection

DNS manipulation

DNS Resolver

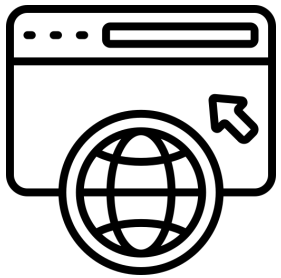

Client

CS249i Server

Censorship during an Internet connection

DNS manipulation

Censorship during an Internet connection

IP blocking

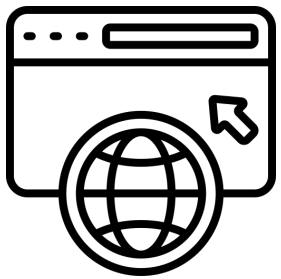

Client

CS249i Server

185.199.111.153

Censorship during an Internet connection

IP blocking

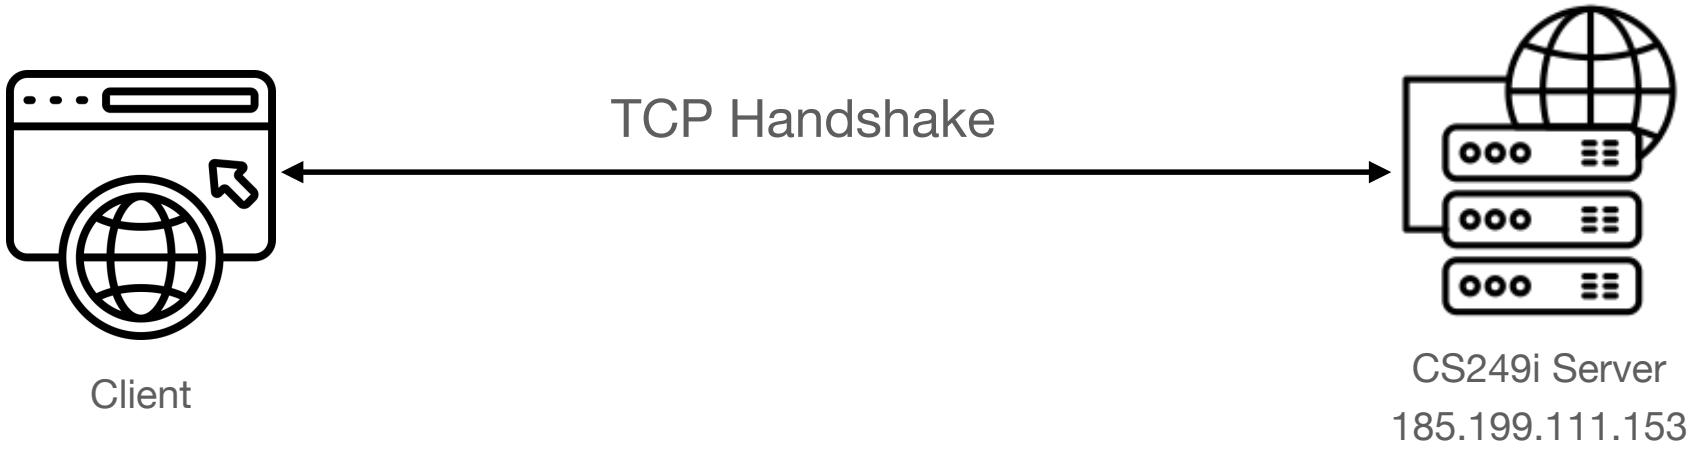

Censorship during an Internet connection

IP blocking

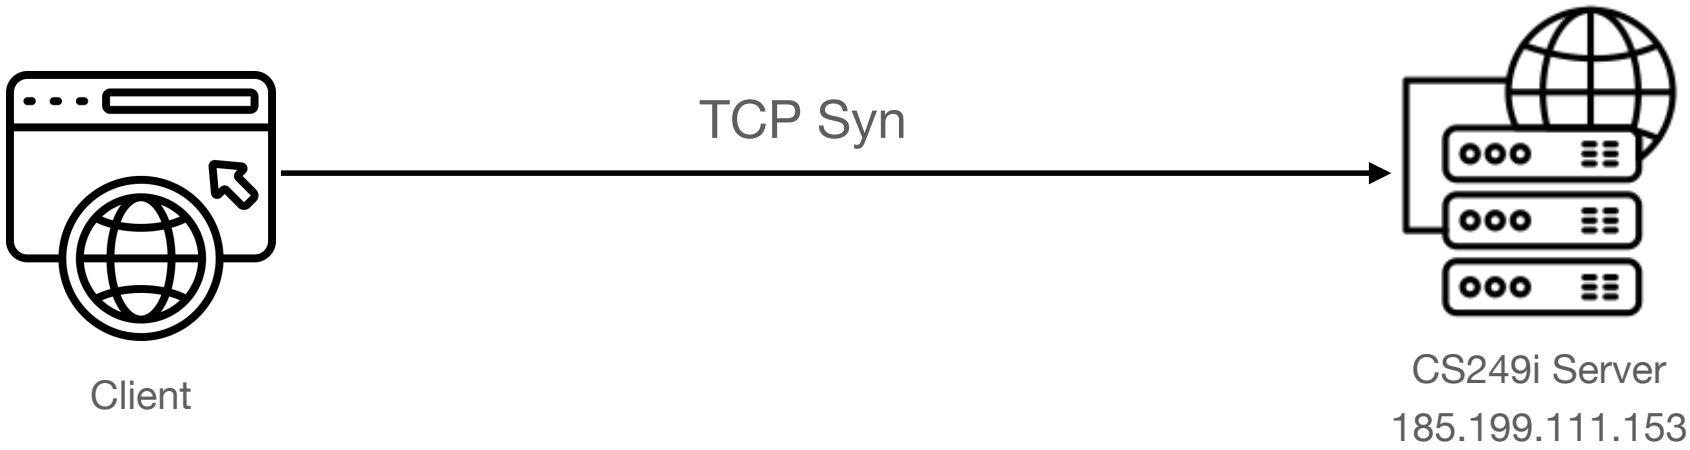

Censorship during an Internet connection

IP blocking

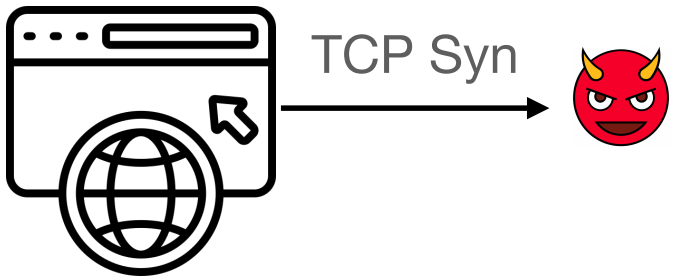

Client

CS249i Server
185.199.111.153

Censorship during an Internet connection

IP blocking

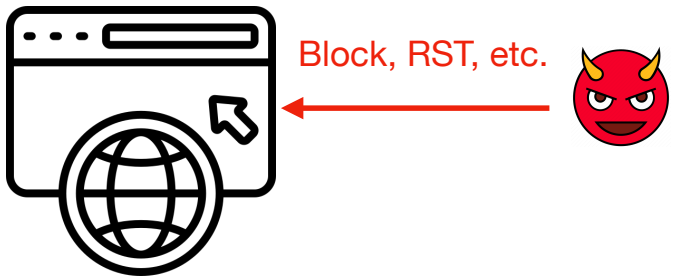

Client

CS249i Server

185.199.111.153

Censorship during an Internet connection

App-layer blocking

Censorship during an Internet connection

App-layer blocking

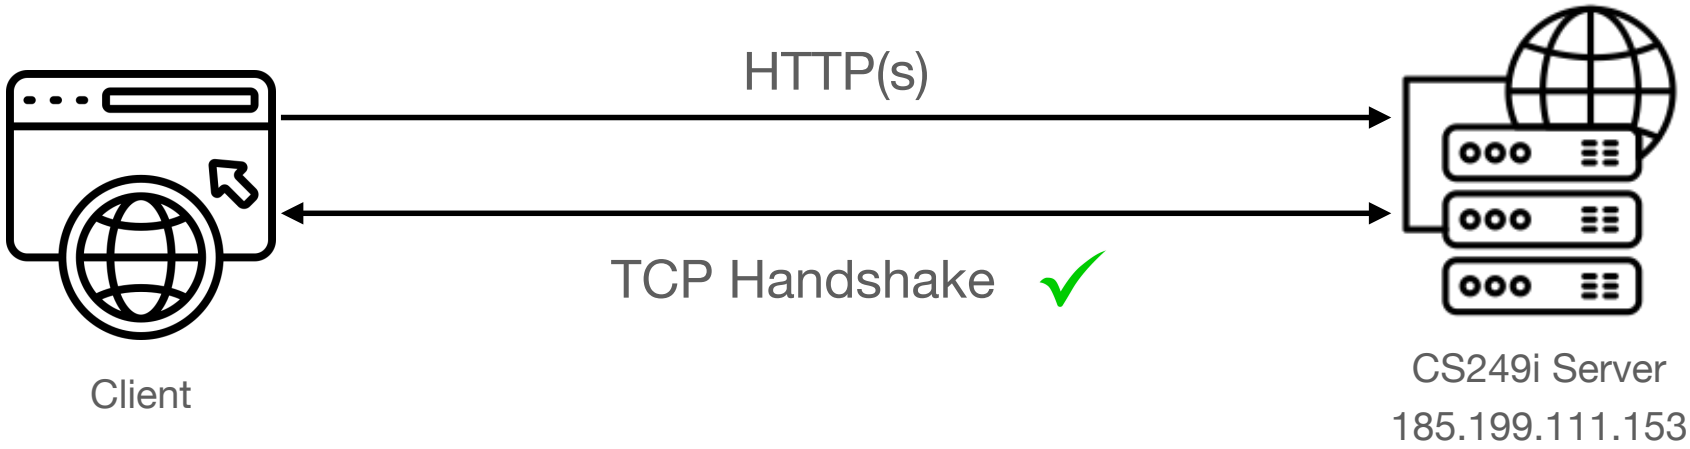

Censorship during an Internet connection

App-layer blocking

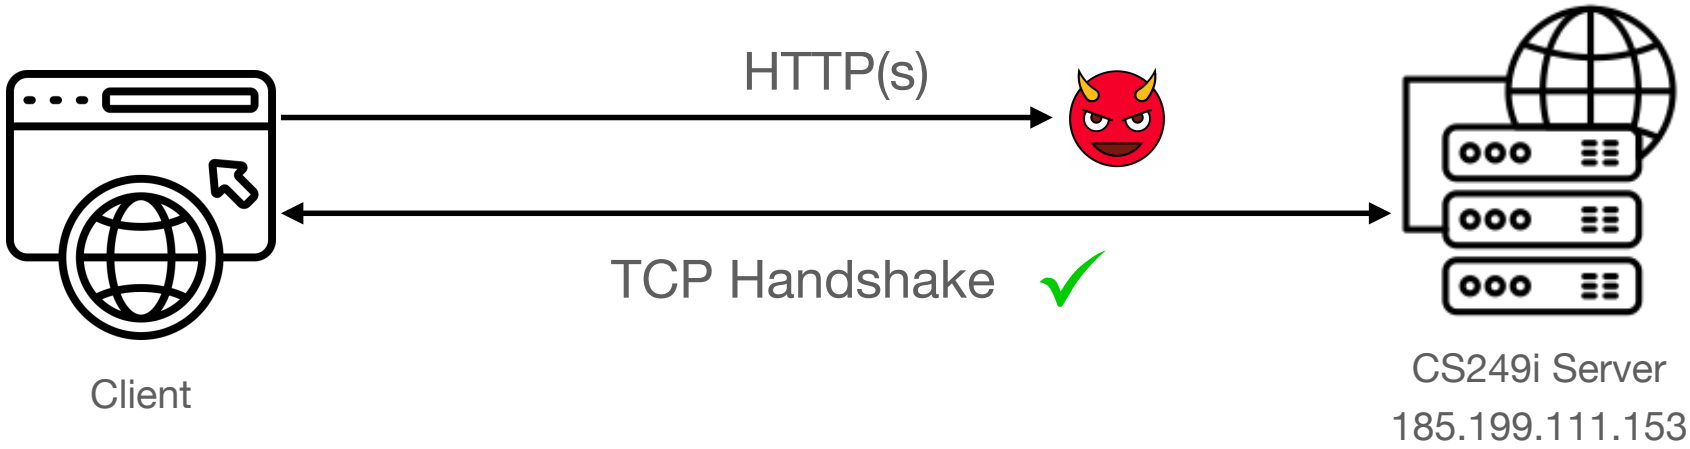

Censorship during an Internet connection

App-layer blocking

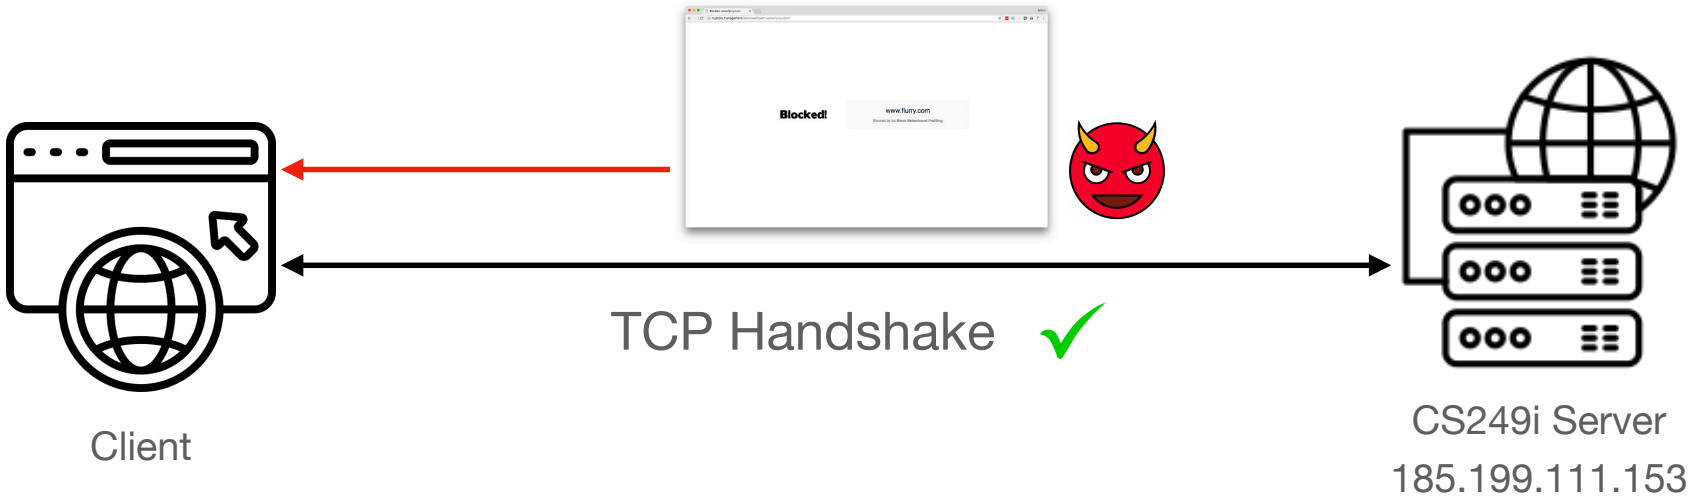

Measuring Internet Censorship

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

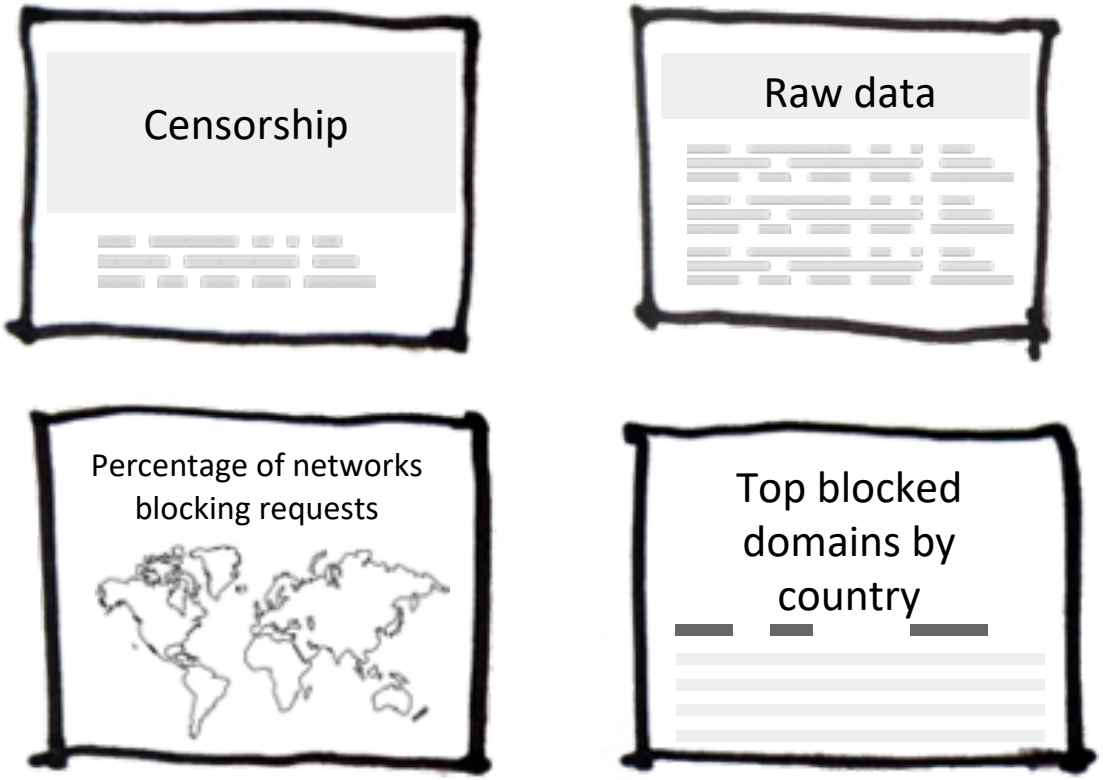

Measuring Internet censorship is hard!

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

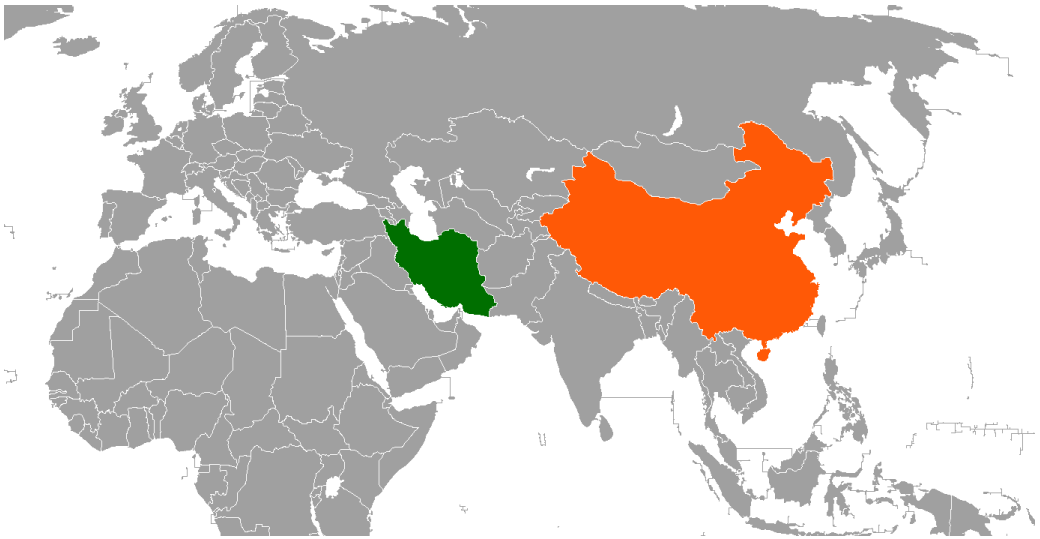

Triplet Censors: Demystifying Great Firewall's DNS Censorship Behavior

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 »

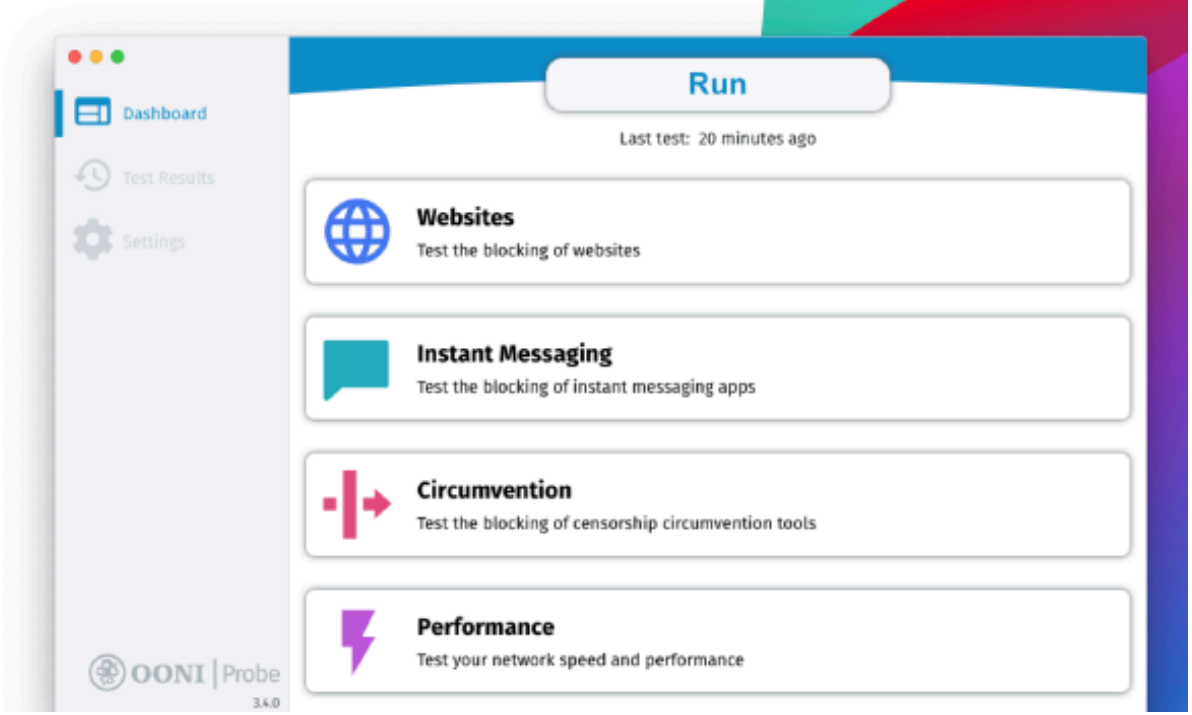

How OONI Works

Volunteer-based direct measurements of censorship

OONI

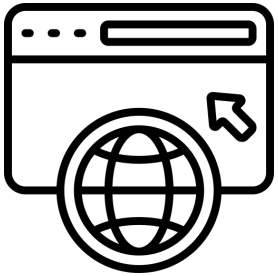

Volunteer client
in-country

Censor

Server of Interest

Control Relay

How OONI Works

Volunteer-based direct measurements of censorship

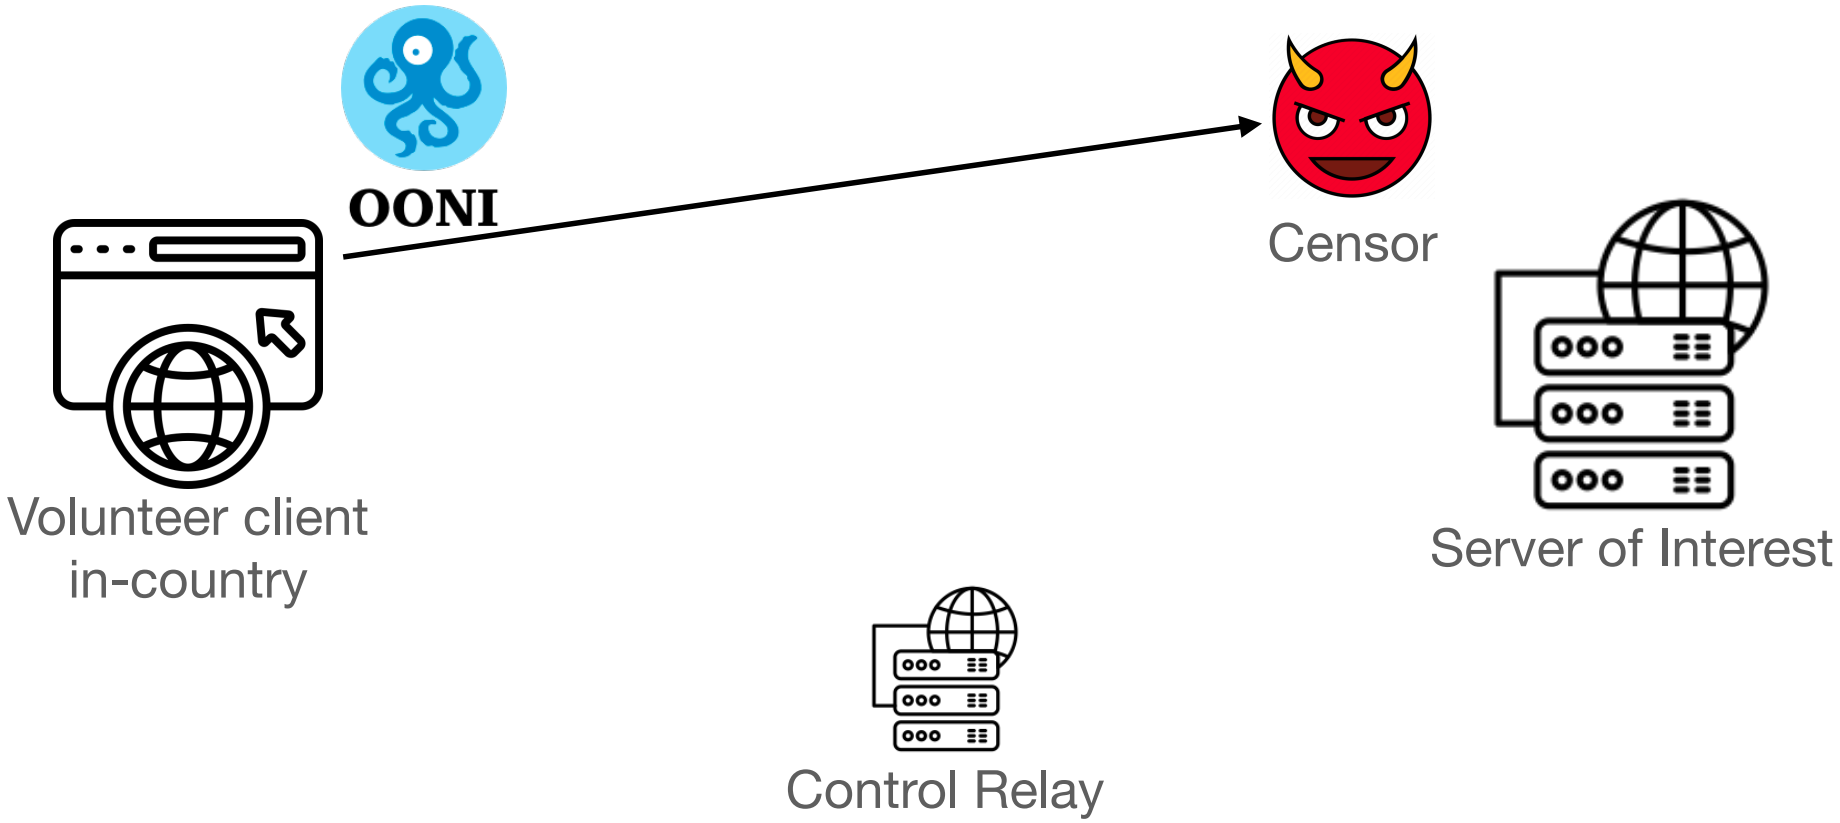

How OONI Works

Volunteer-based direct measurements of censorship

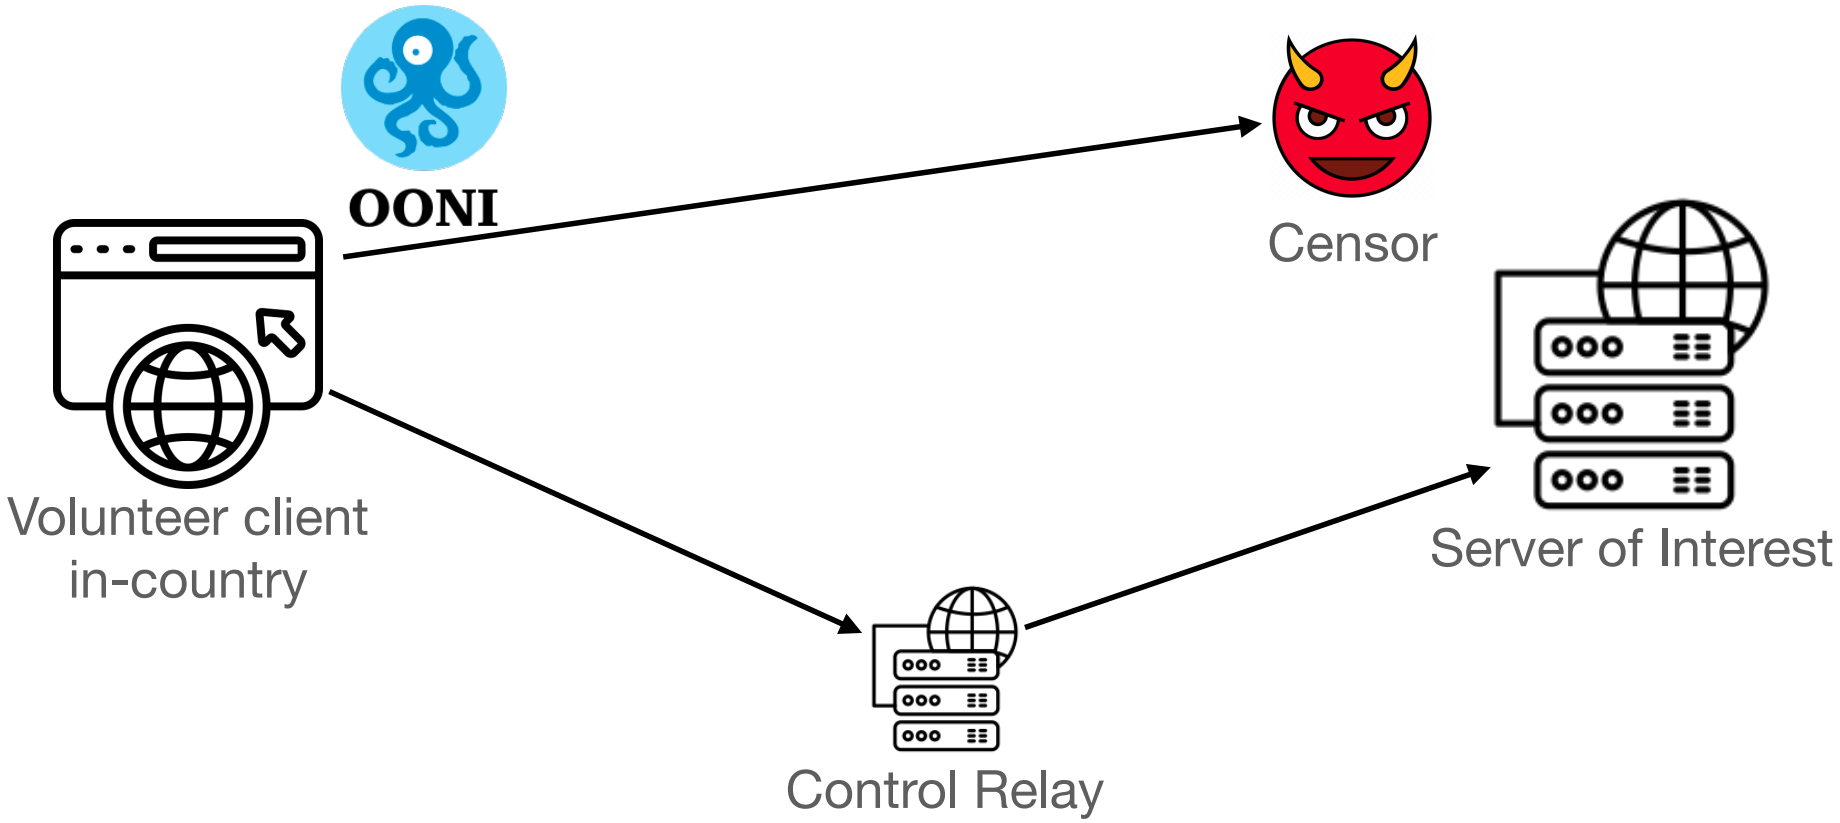

How OONI Works

Volunteer-based direct measurements of censorship

https://explorer.ooni.org

Limitations of volunteer measurements

5 key problems

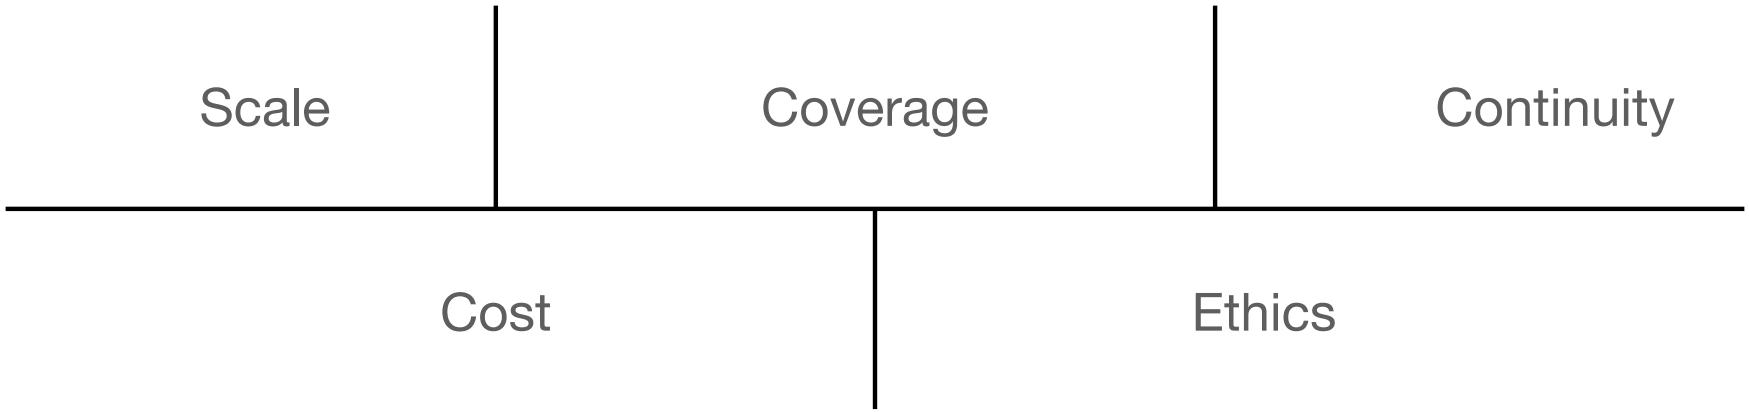

A word on ethics…

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?

Remote Censorship Measurements

A new approach to examining censorship

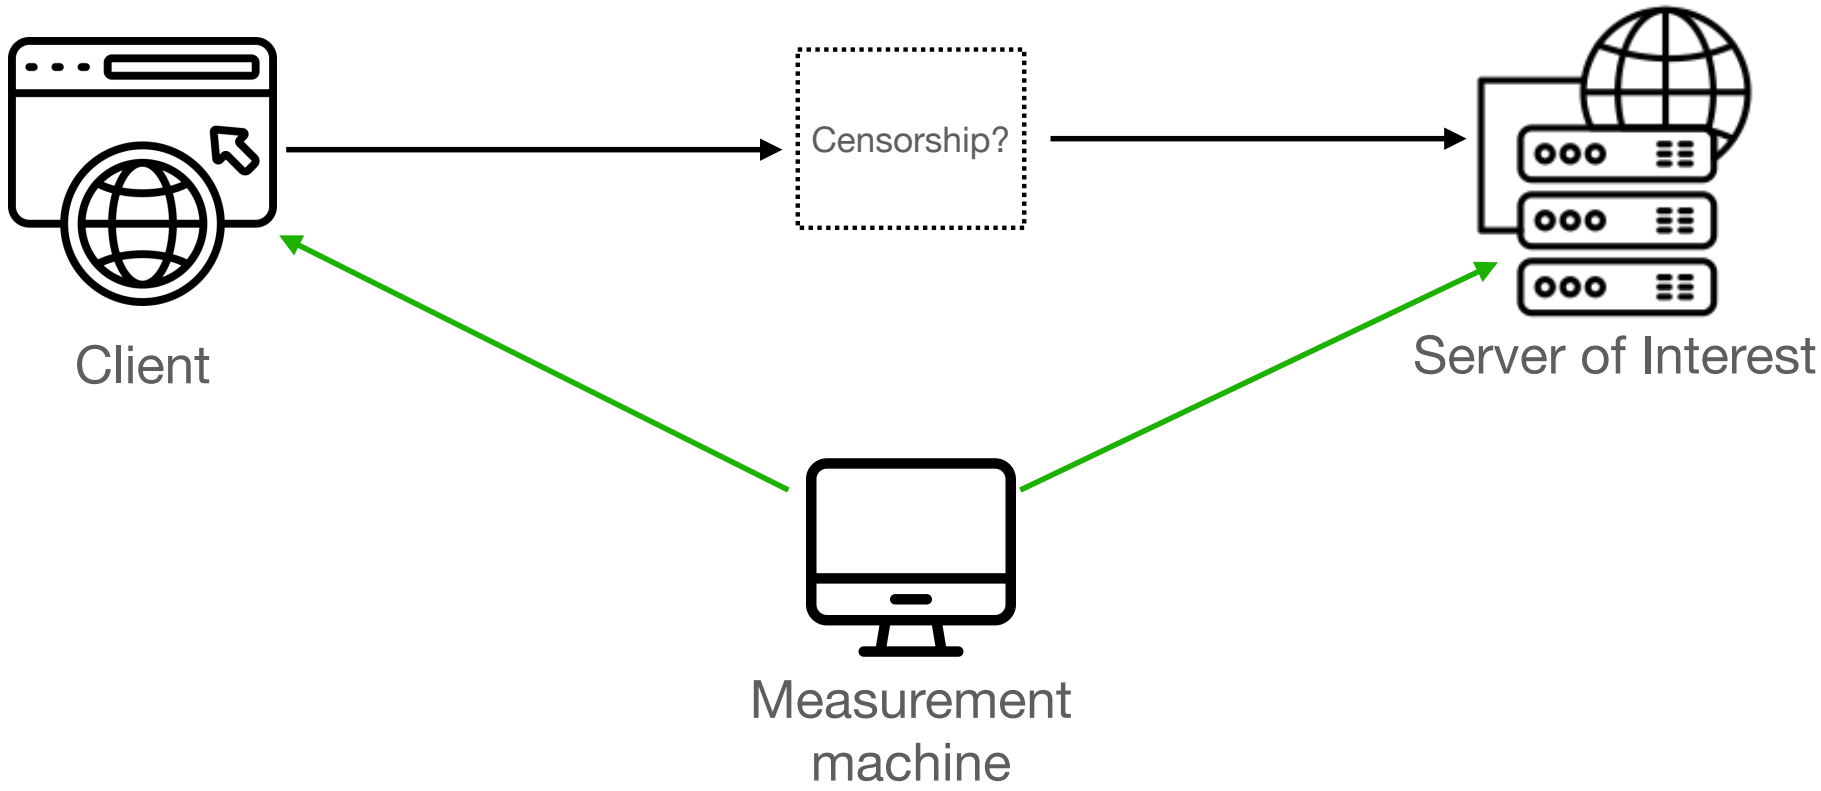

Spooky Scans

Using TCP/IP side-channels to detect if clients + servers can communicate

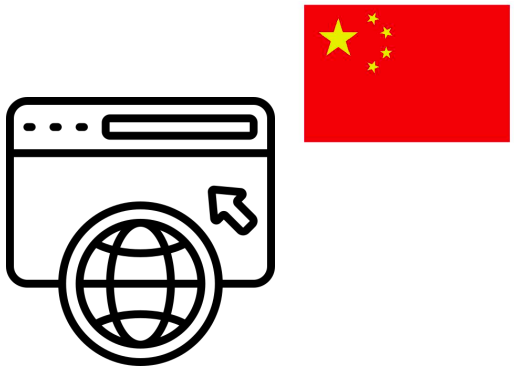

Client

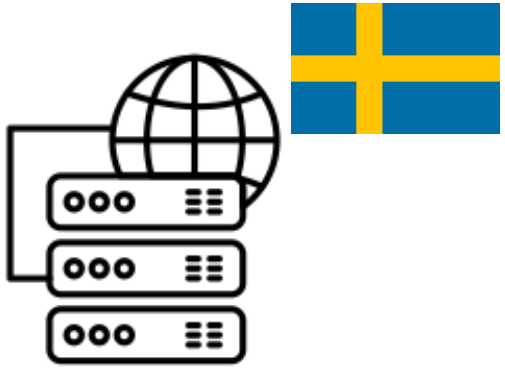

Server of Interest

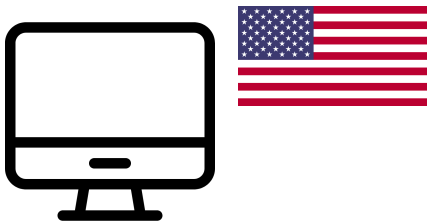

Measurement machine

Detecting Intentional Packet Drops on the Internet via TCP/IP Side Channels
Roya Ensafi, Knockel, Alexander, and Crandall (PAM '14)

TCP / IP Background

IP-IDs

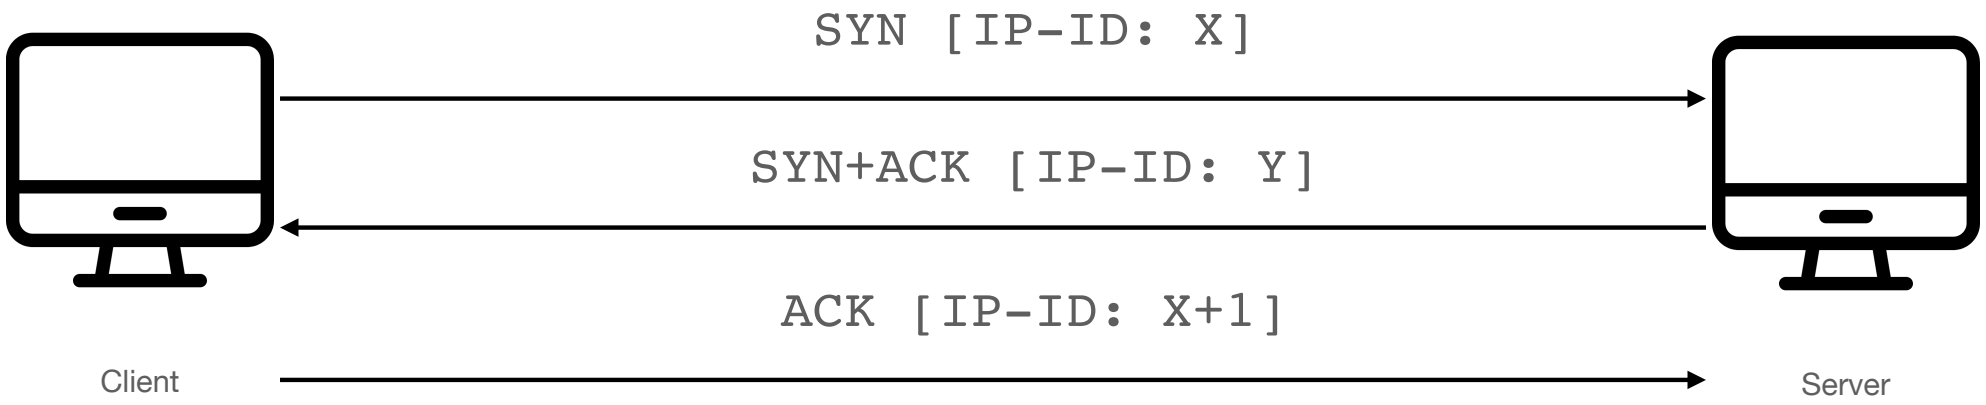

TCP / IP Background

RSTs

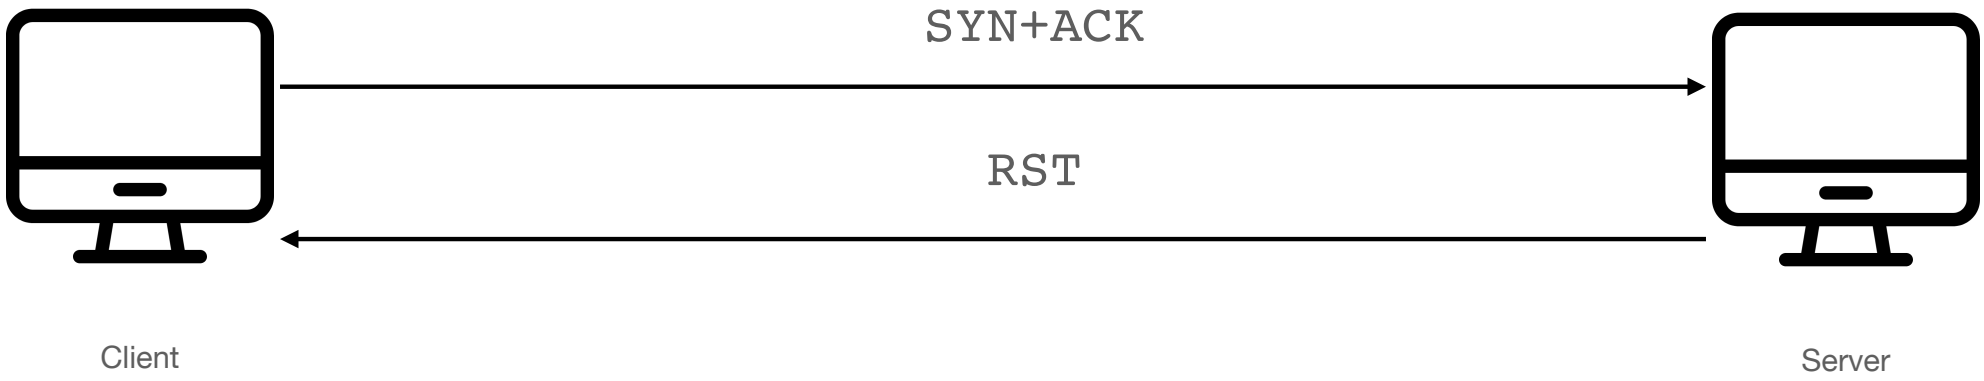

TCP / IP Background

Retries

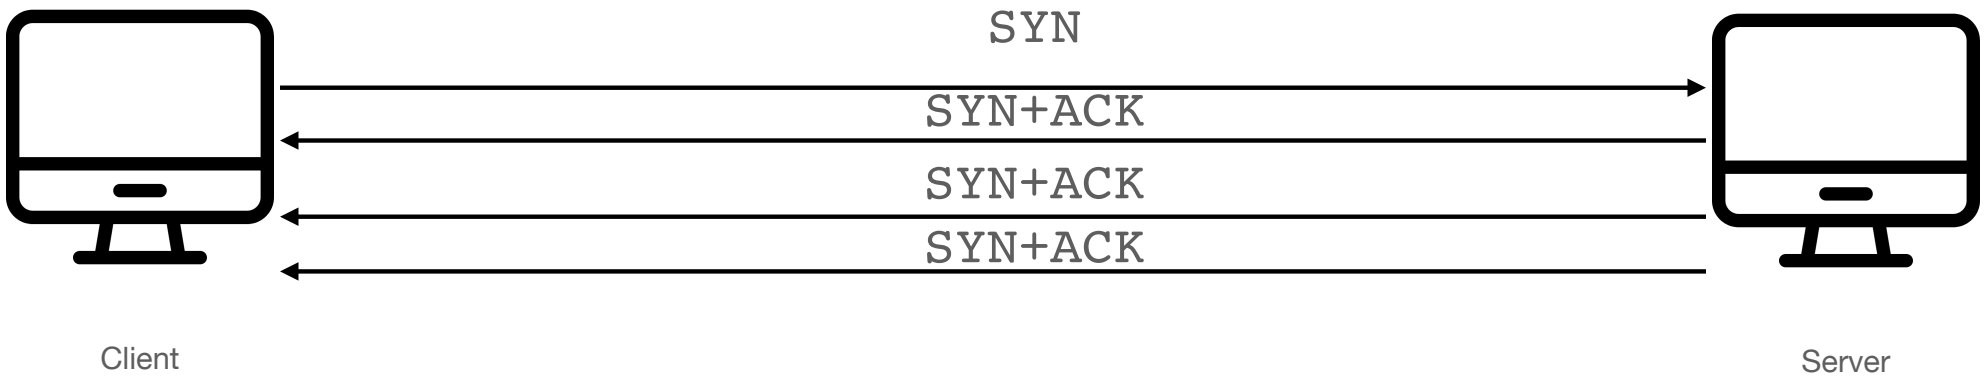

Spooky Scan Requirements

Clients, Server, Spoofing Packets

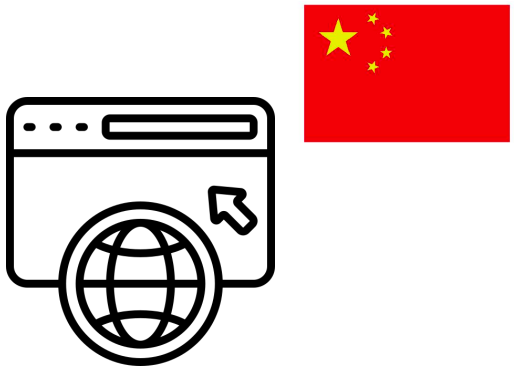

Client

Must maintain a global value for IP_ID

Must be able to spoof packets

Measurement machine

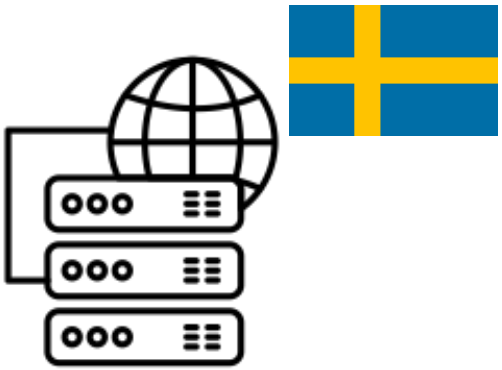

Server of Interest

Open port + retransmitting SYN+ACKs

Spooky Scans

No blocking

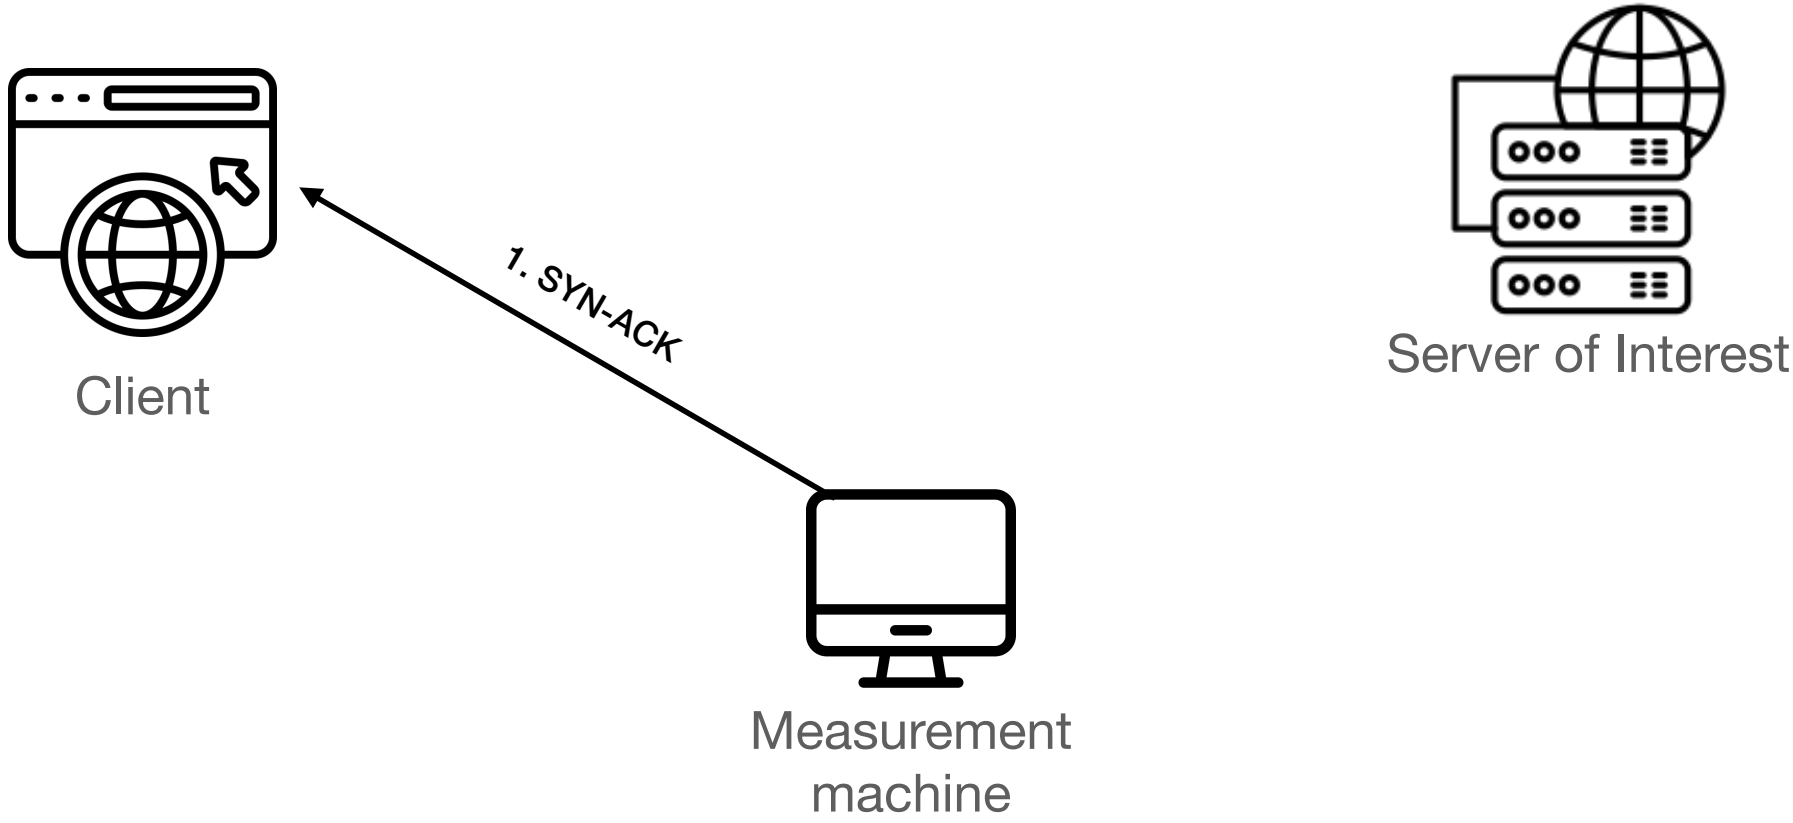

Spooky Scans

No blocking

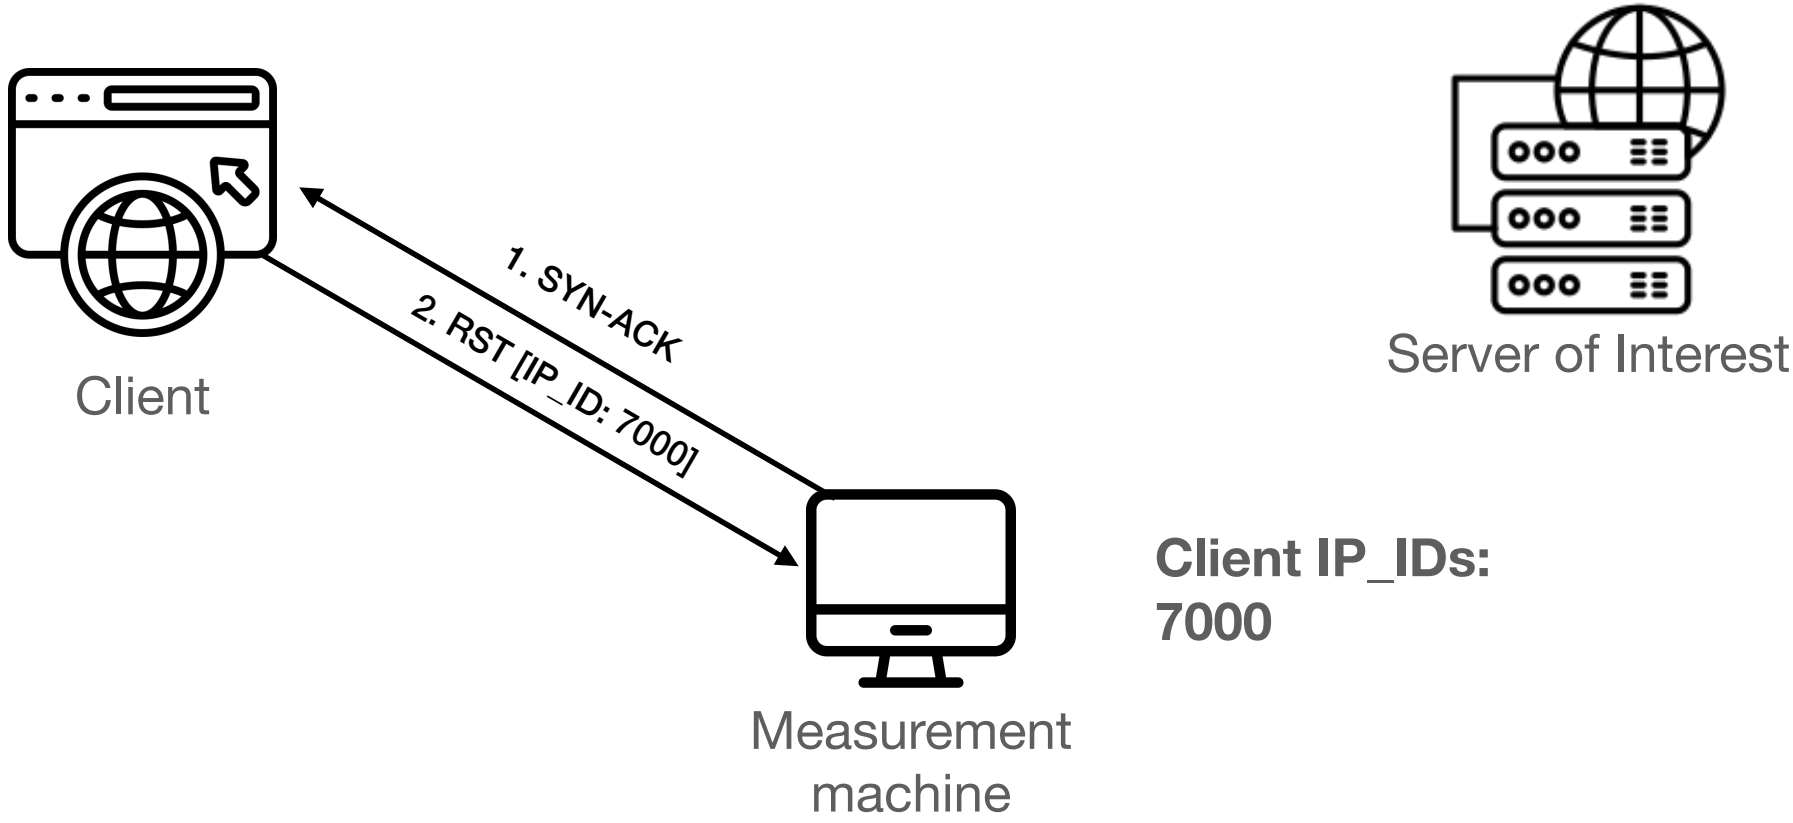

Client IP_IDs:
7000

Spooky Scans

No blocking

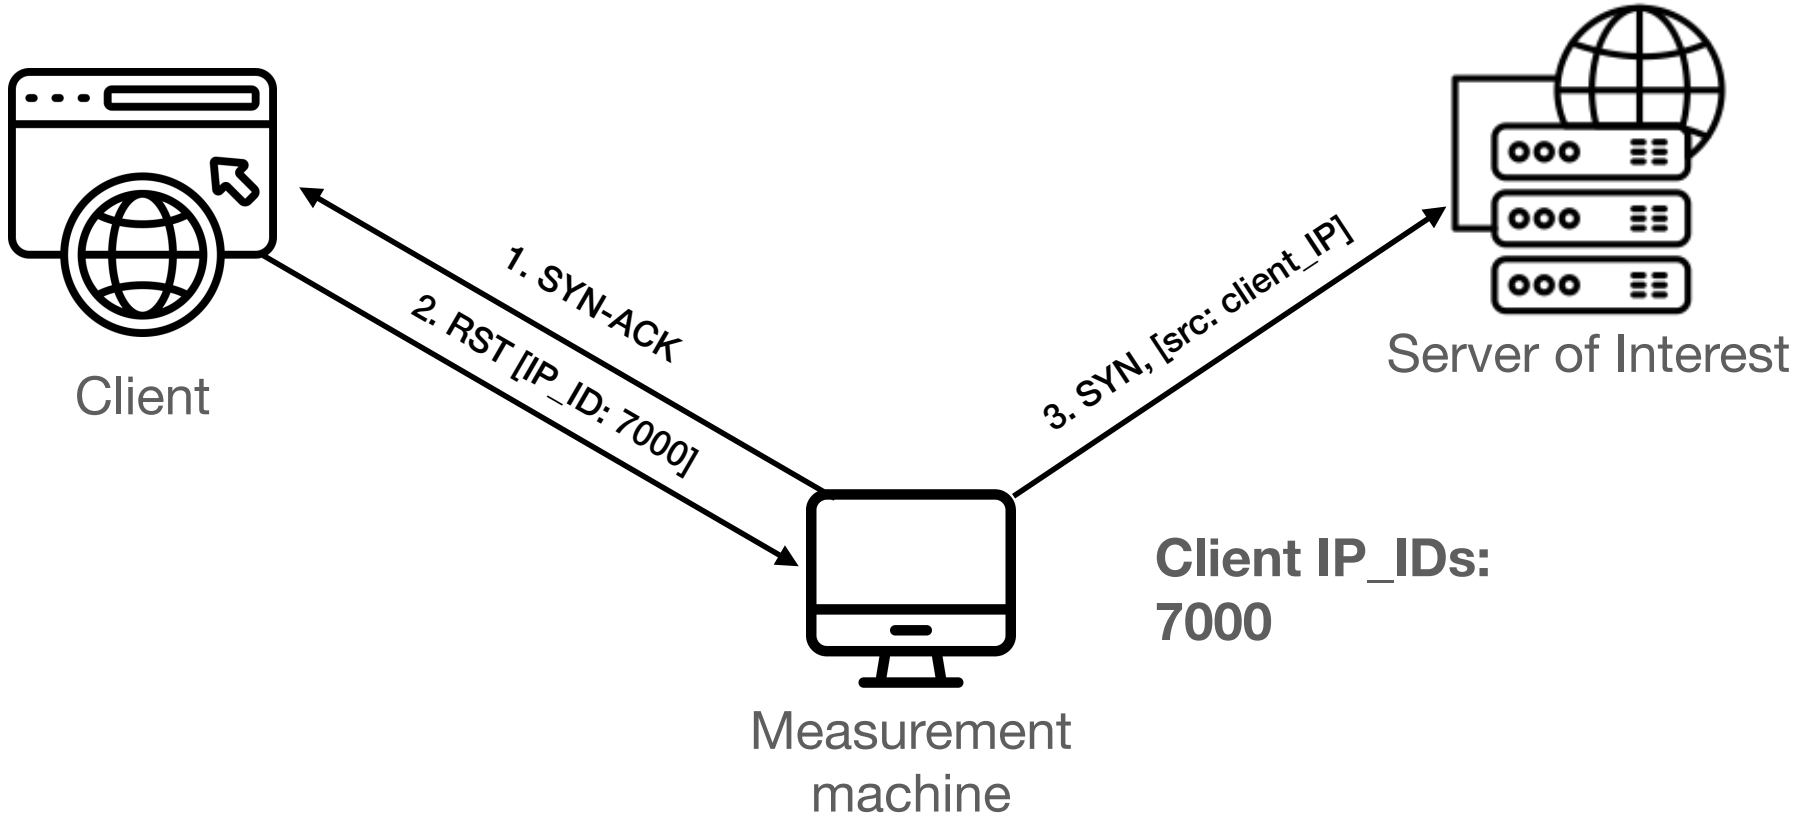

Spooky Scans

No blocking

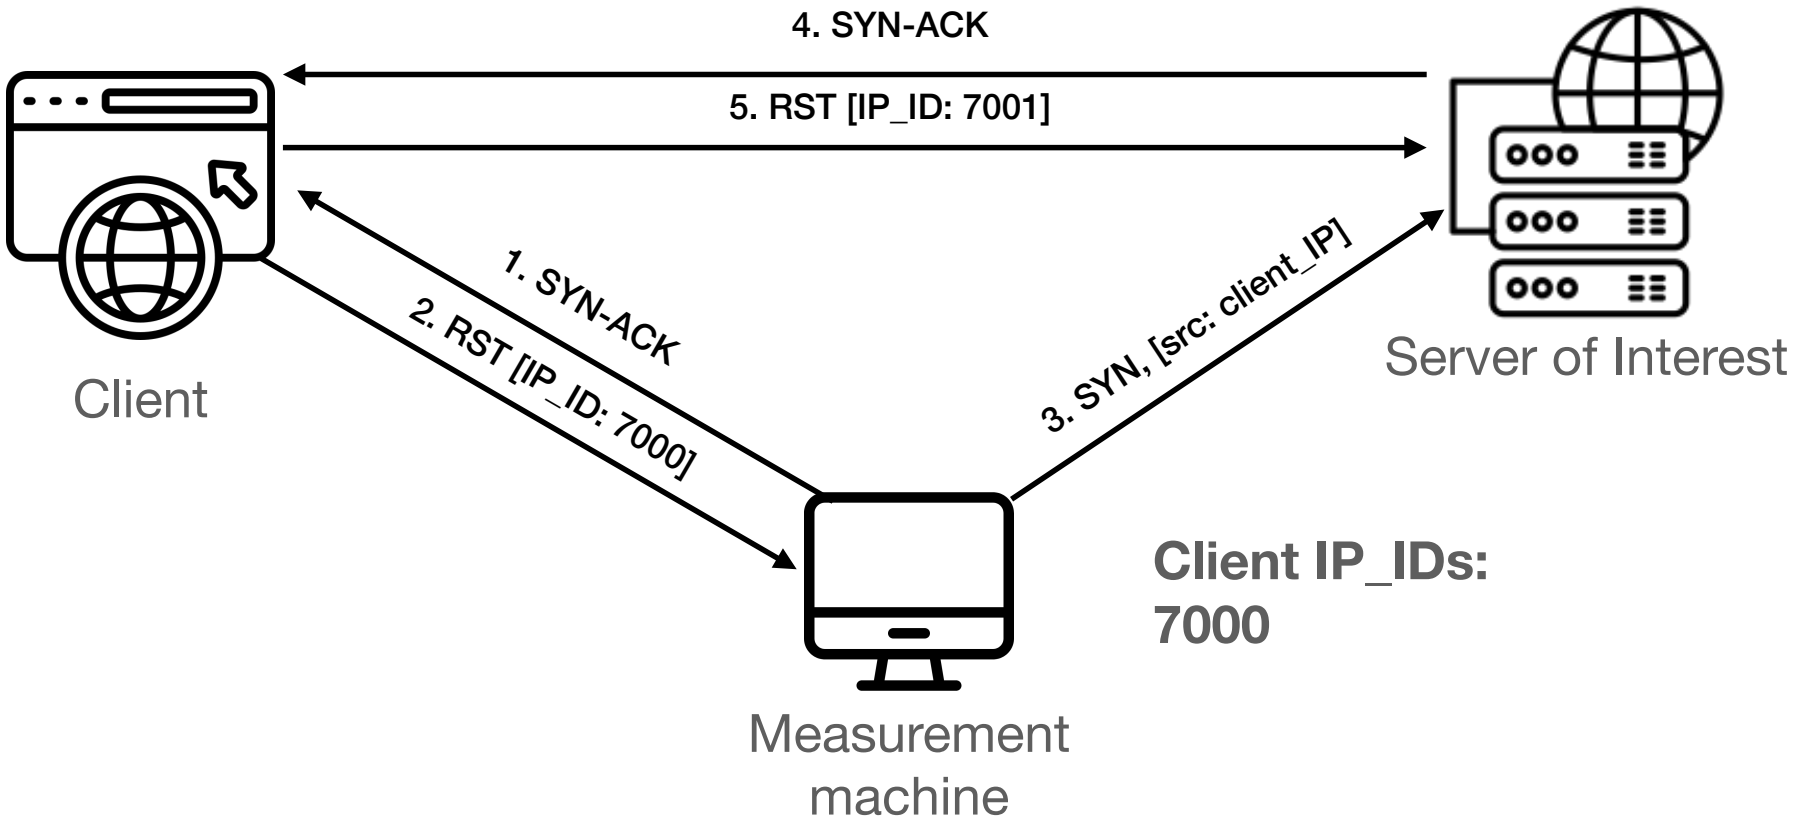

Spooky Scans

No blocking

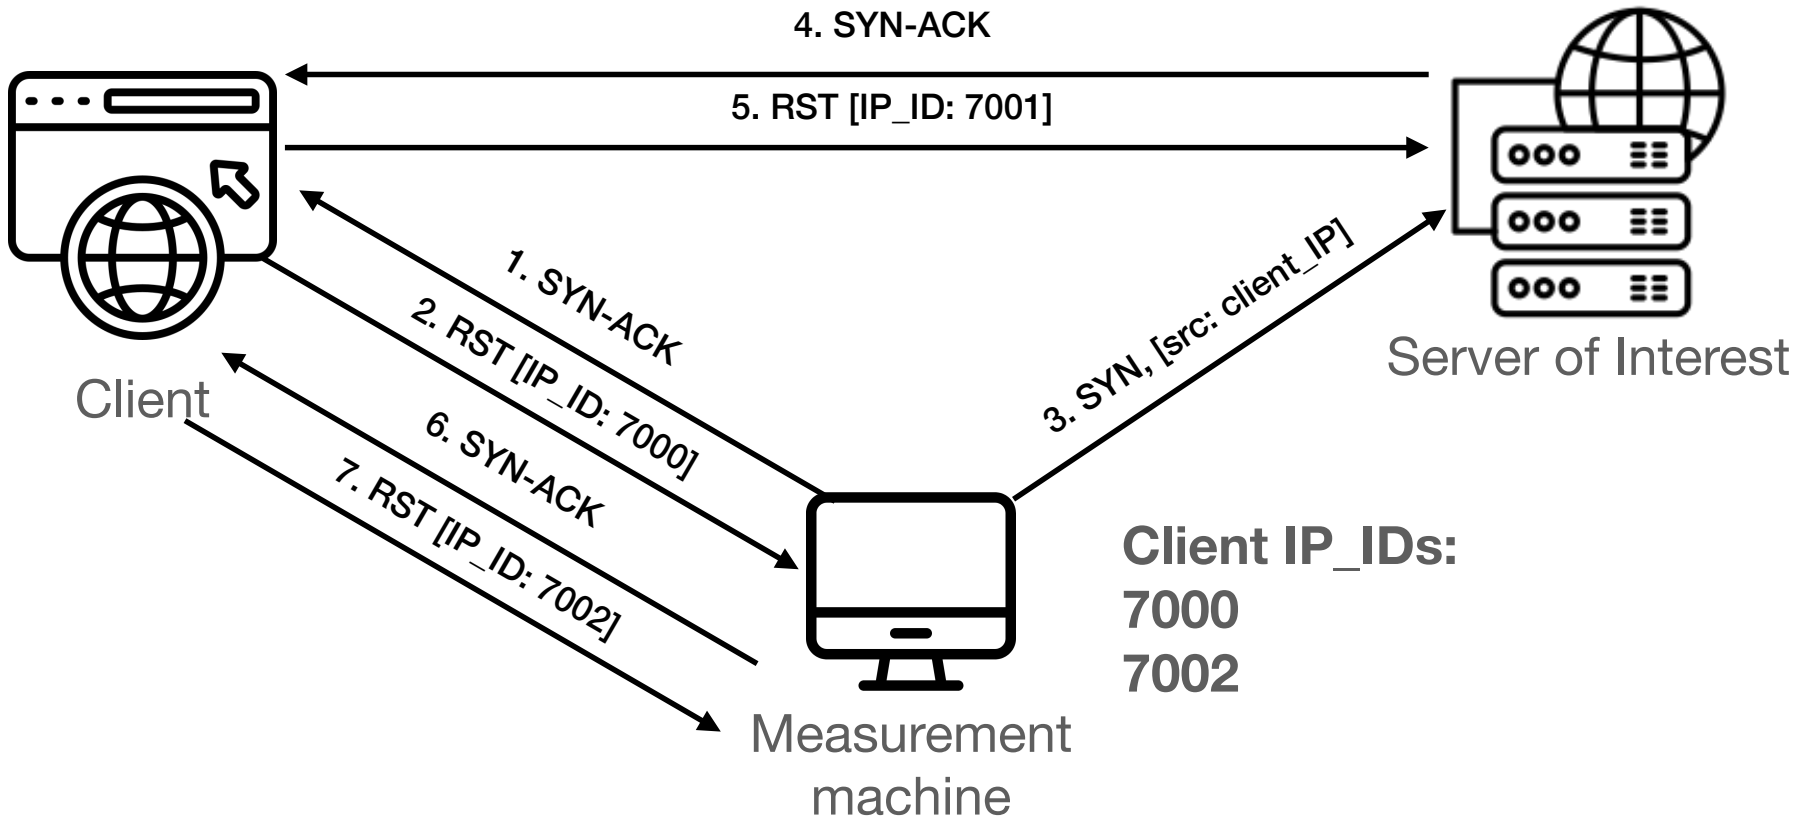

Spooky Scans

No blocking

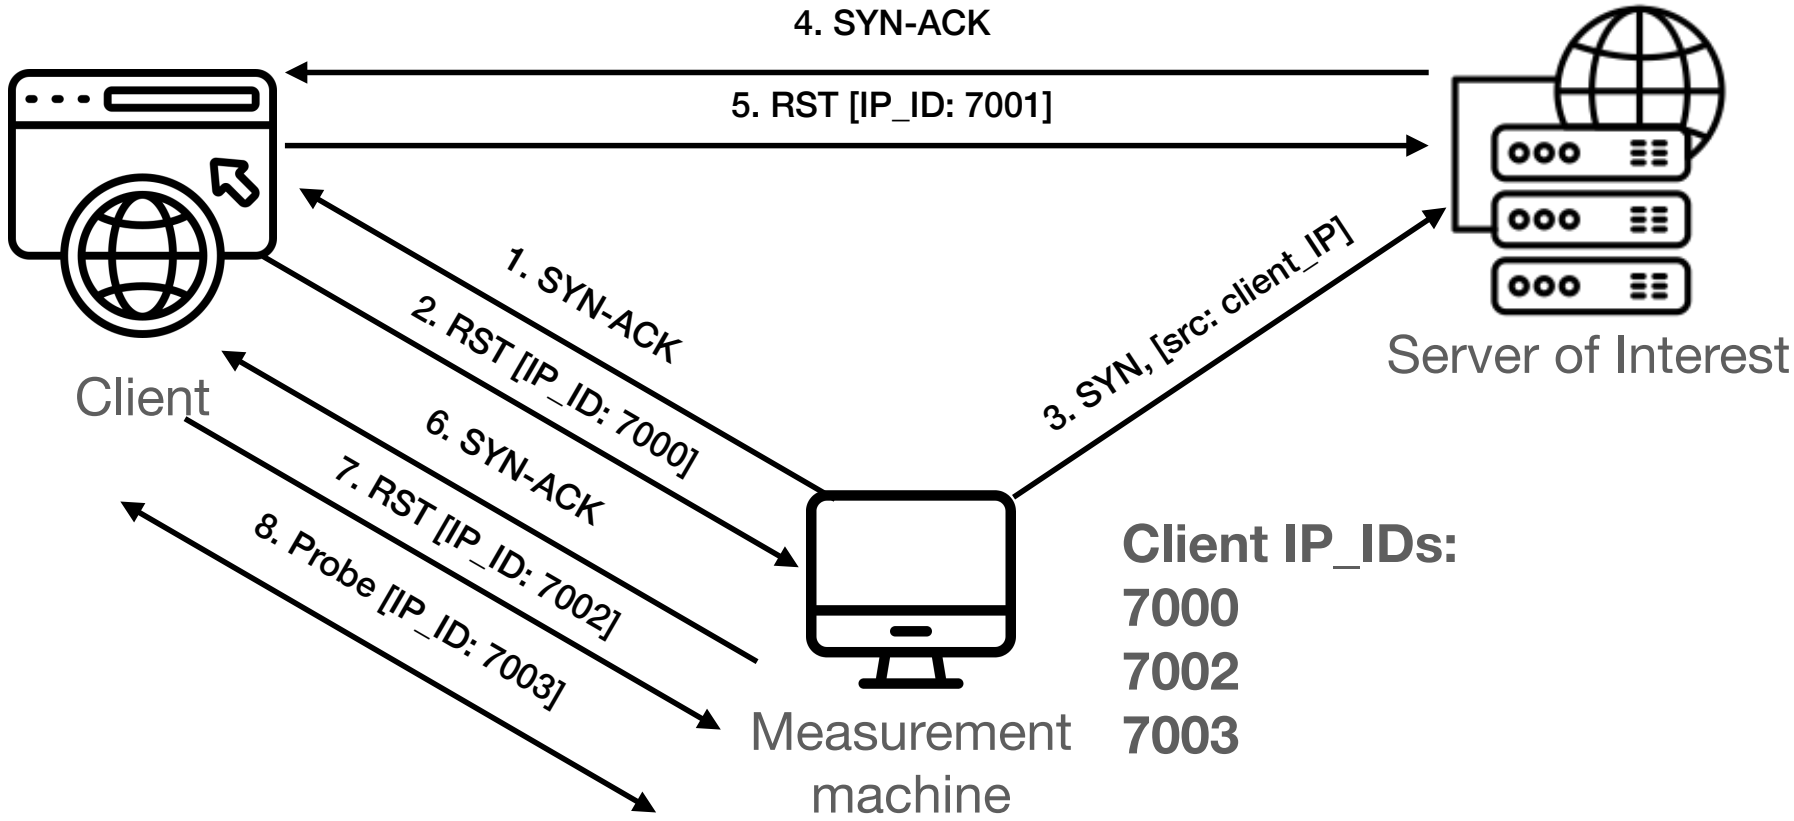

Spooky Scans

Server-to-client is blocked

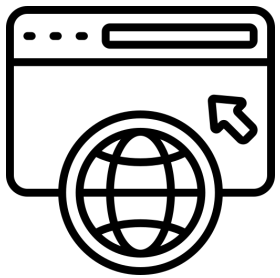

Client

Censor

Server of Interest

Measurement machine

Spooky Scans

Server-to-client is blocked

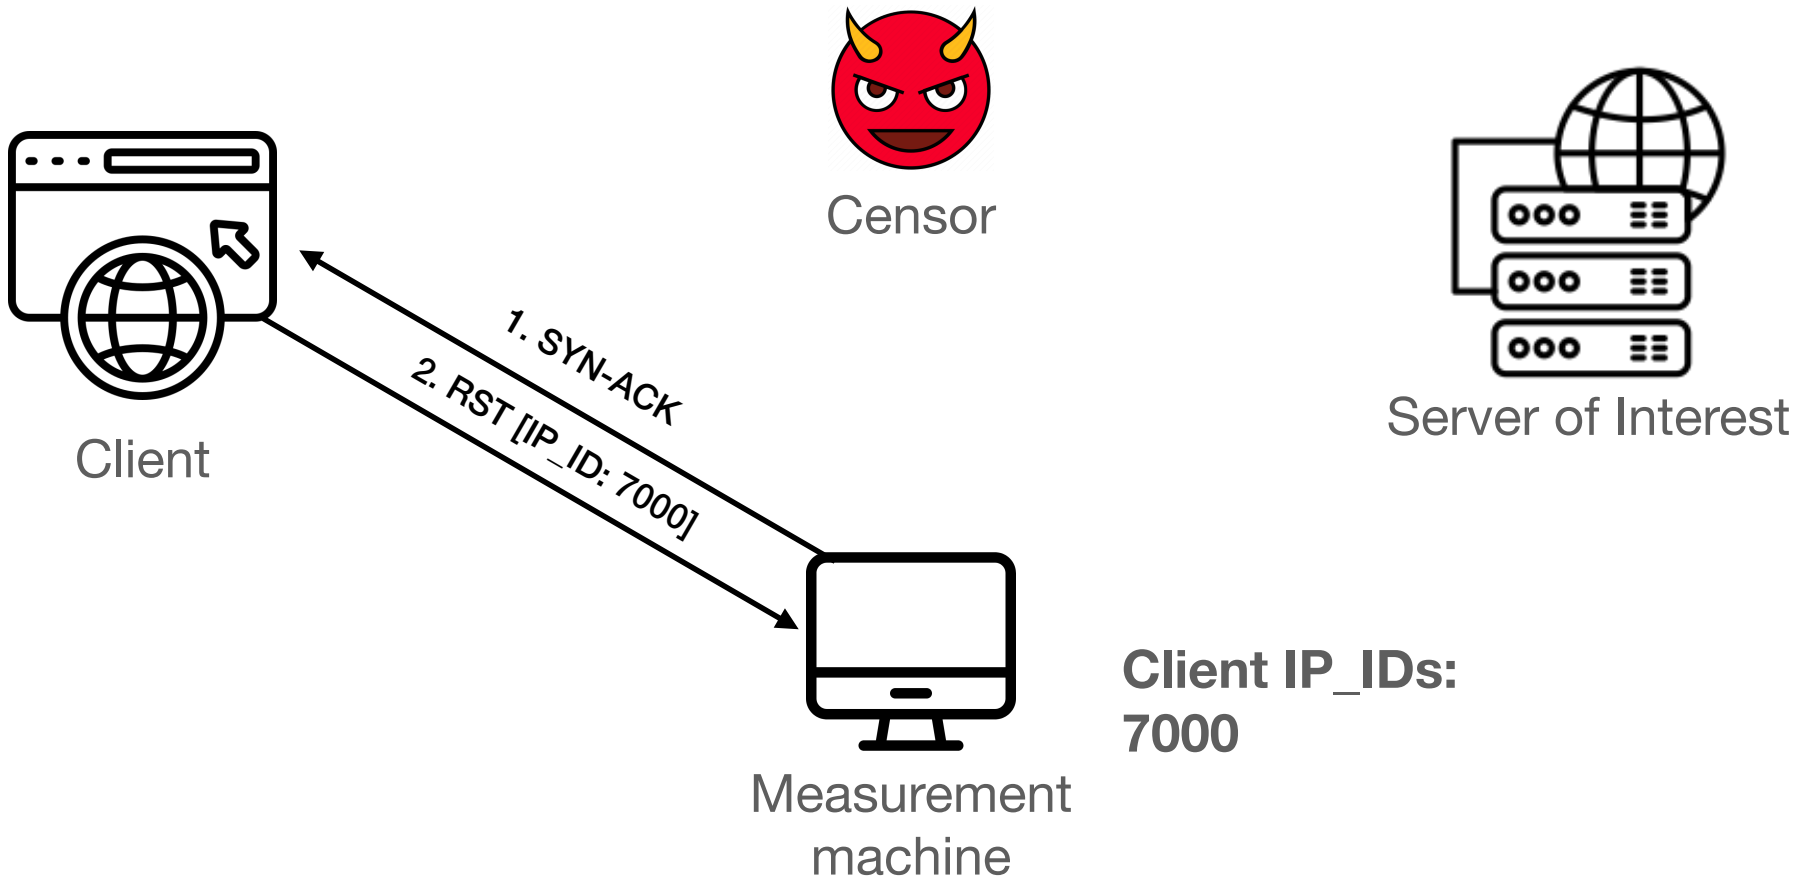

Spooky Scans

Server-to-client is blocked

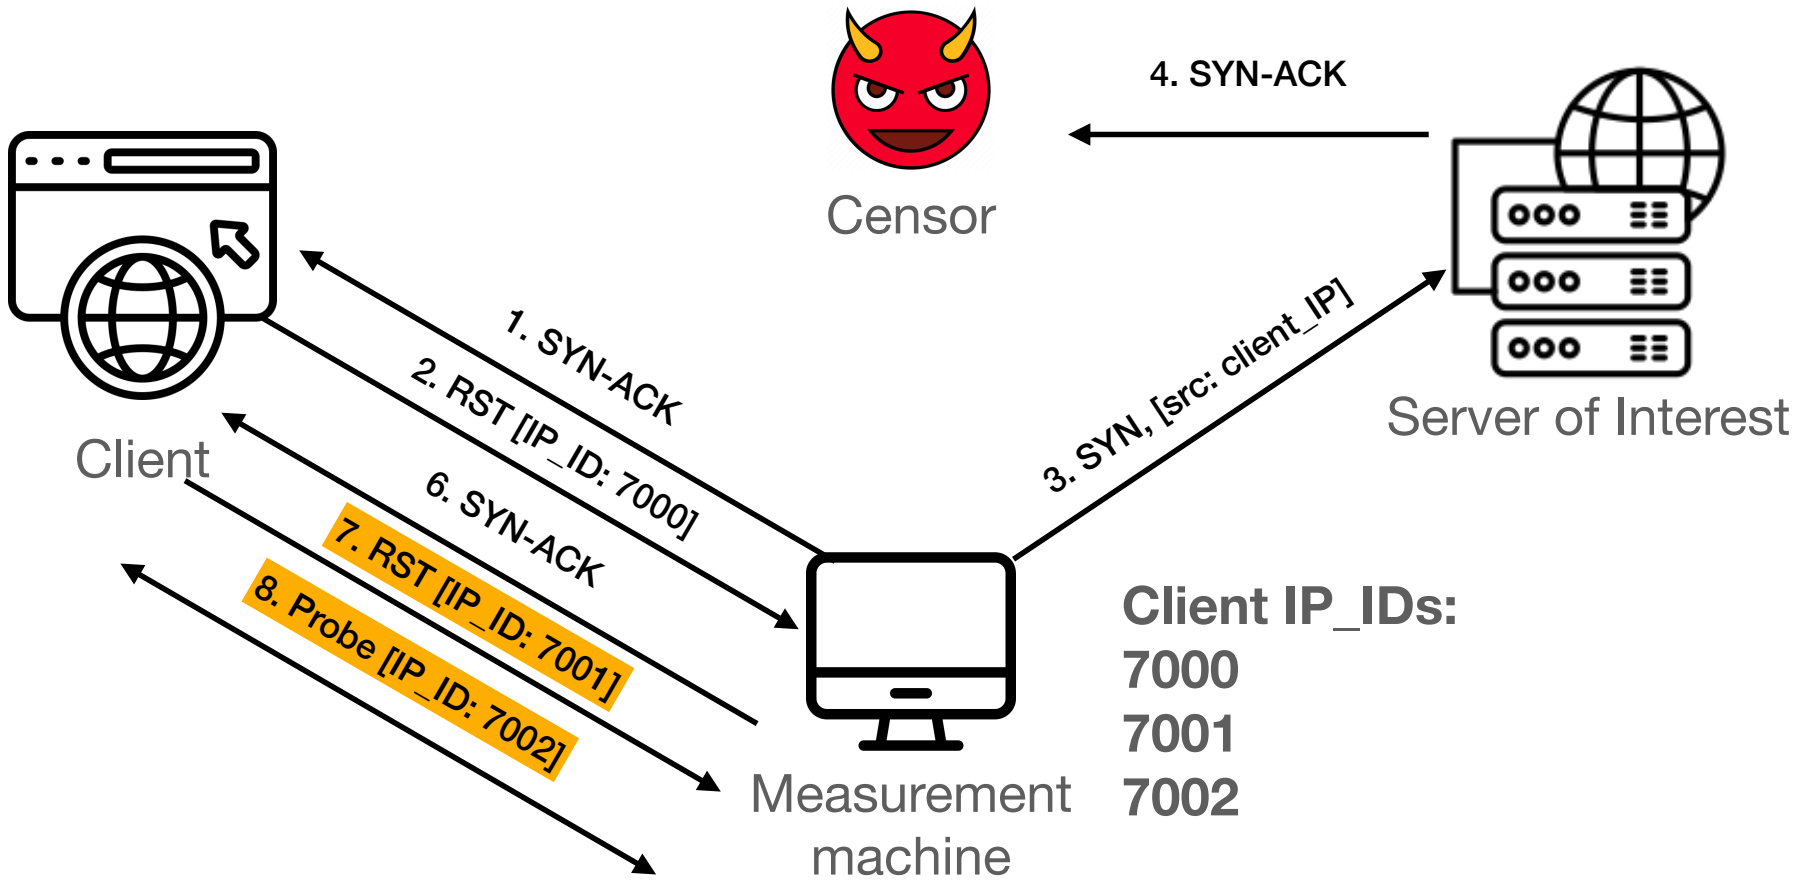

Spooky Scans

Client-to-server is blocked

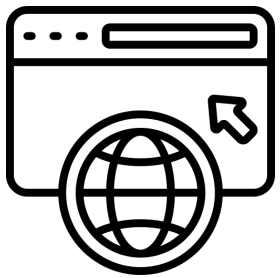

Client

Censor

Server of Interest

Measurement machine

Spooky Scans

Client-to-server is blocked

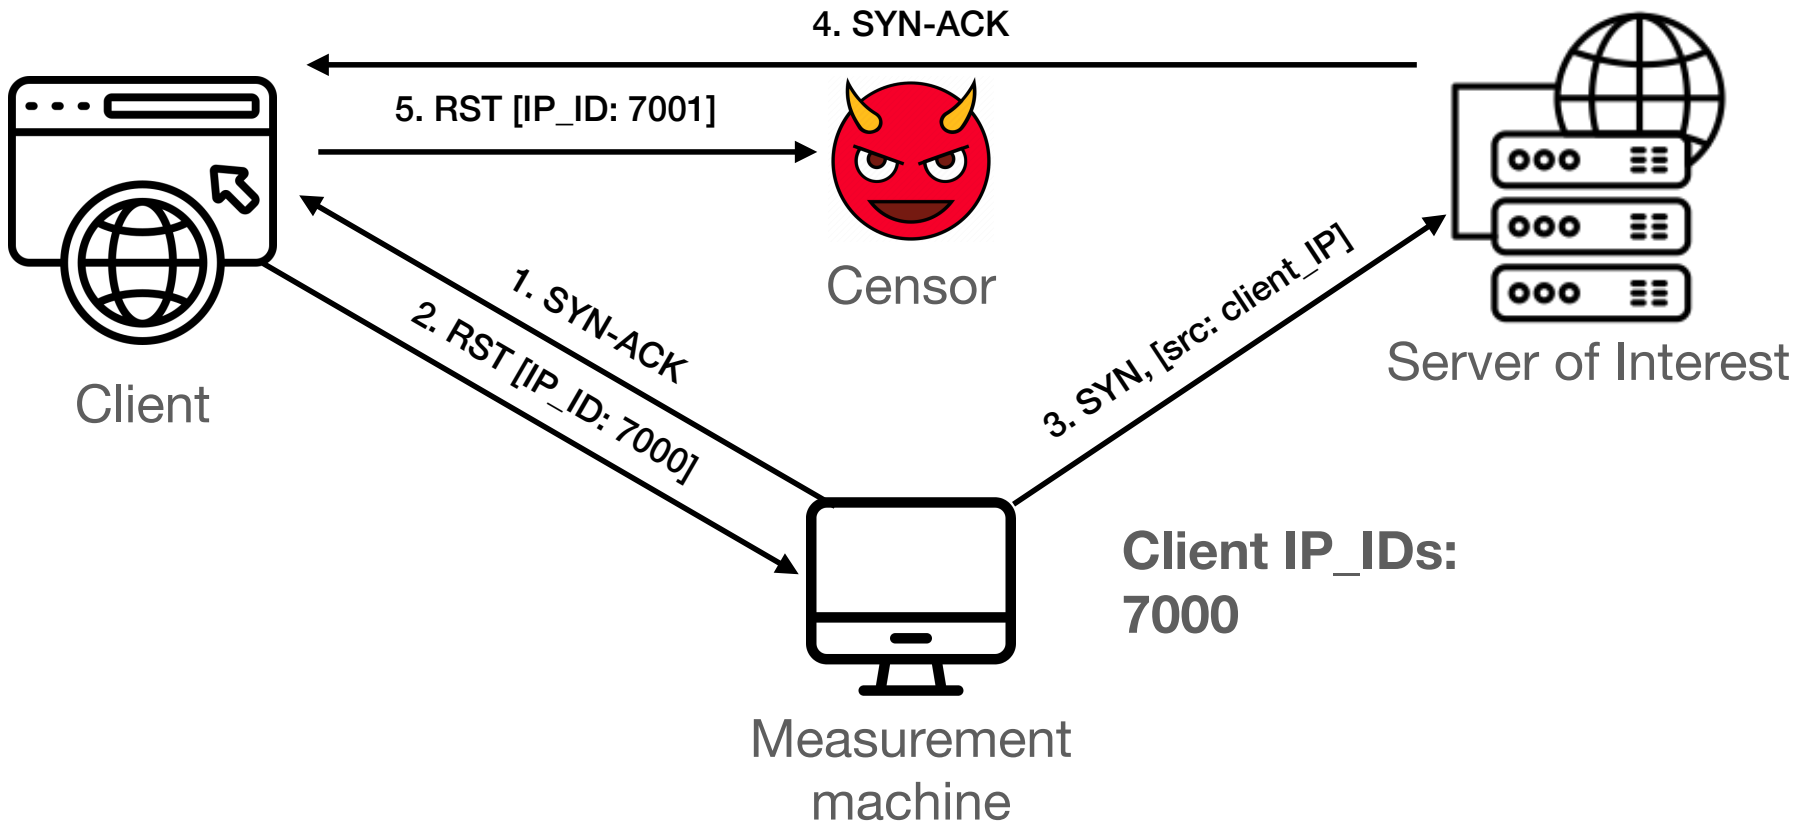

Spooky Scans

Client-to-server is blocked

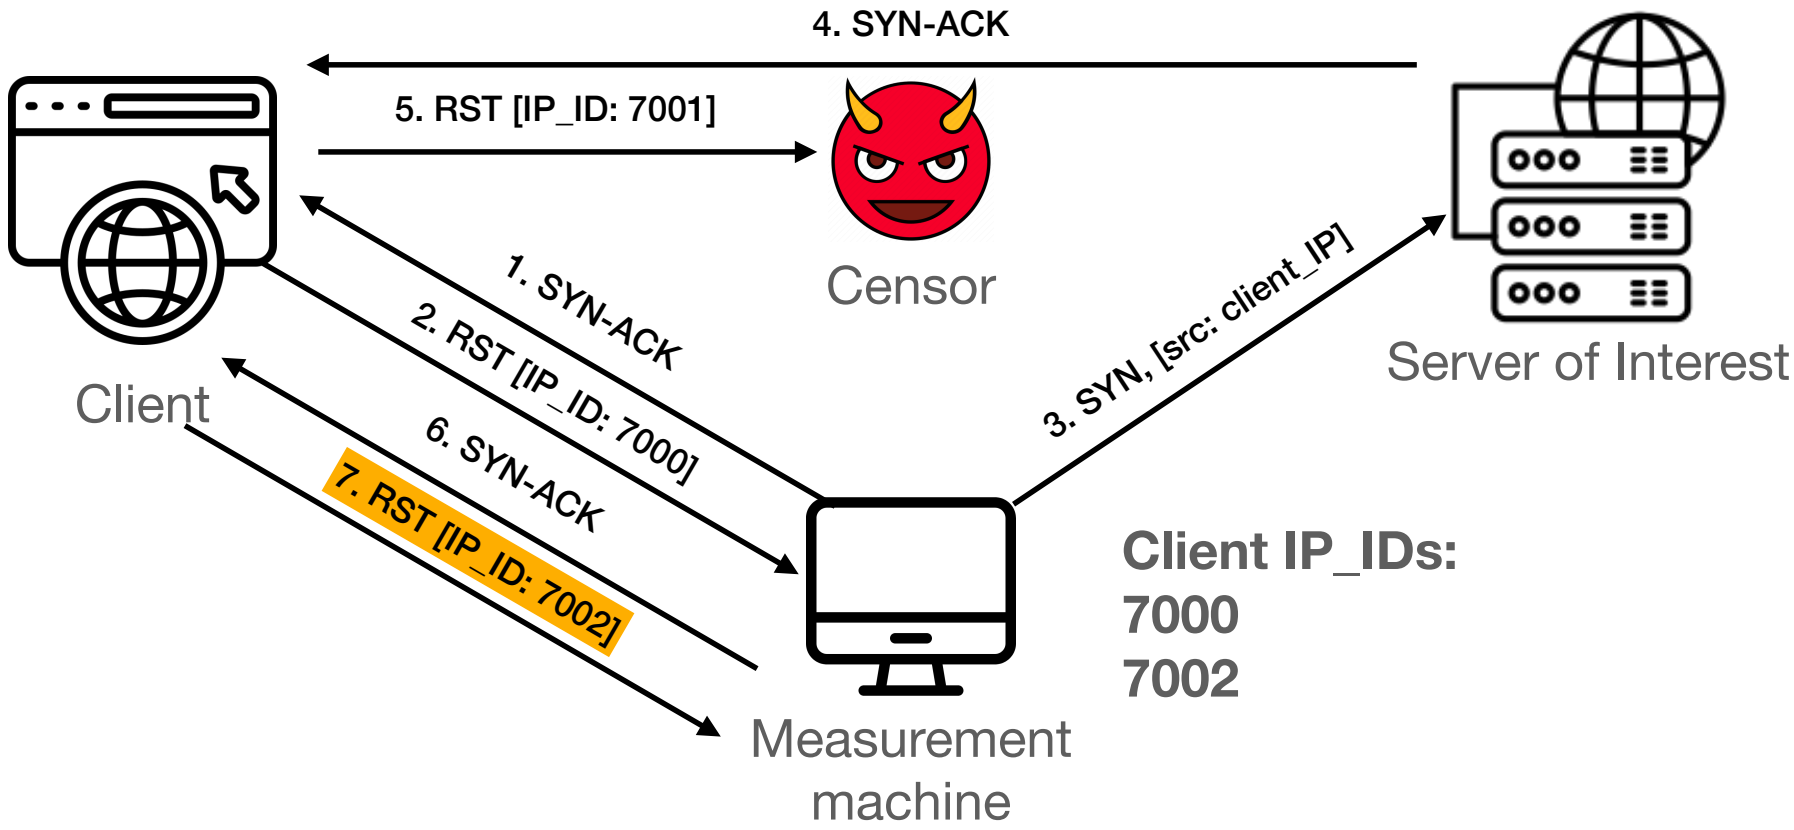

Spooky Scans

Client-to-server is blocked

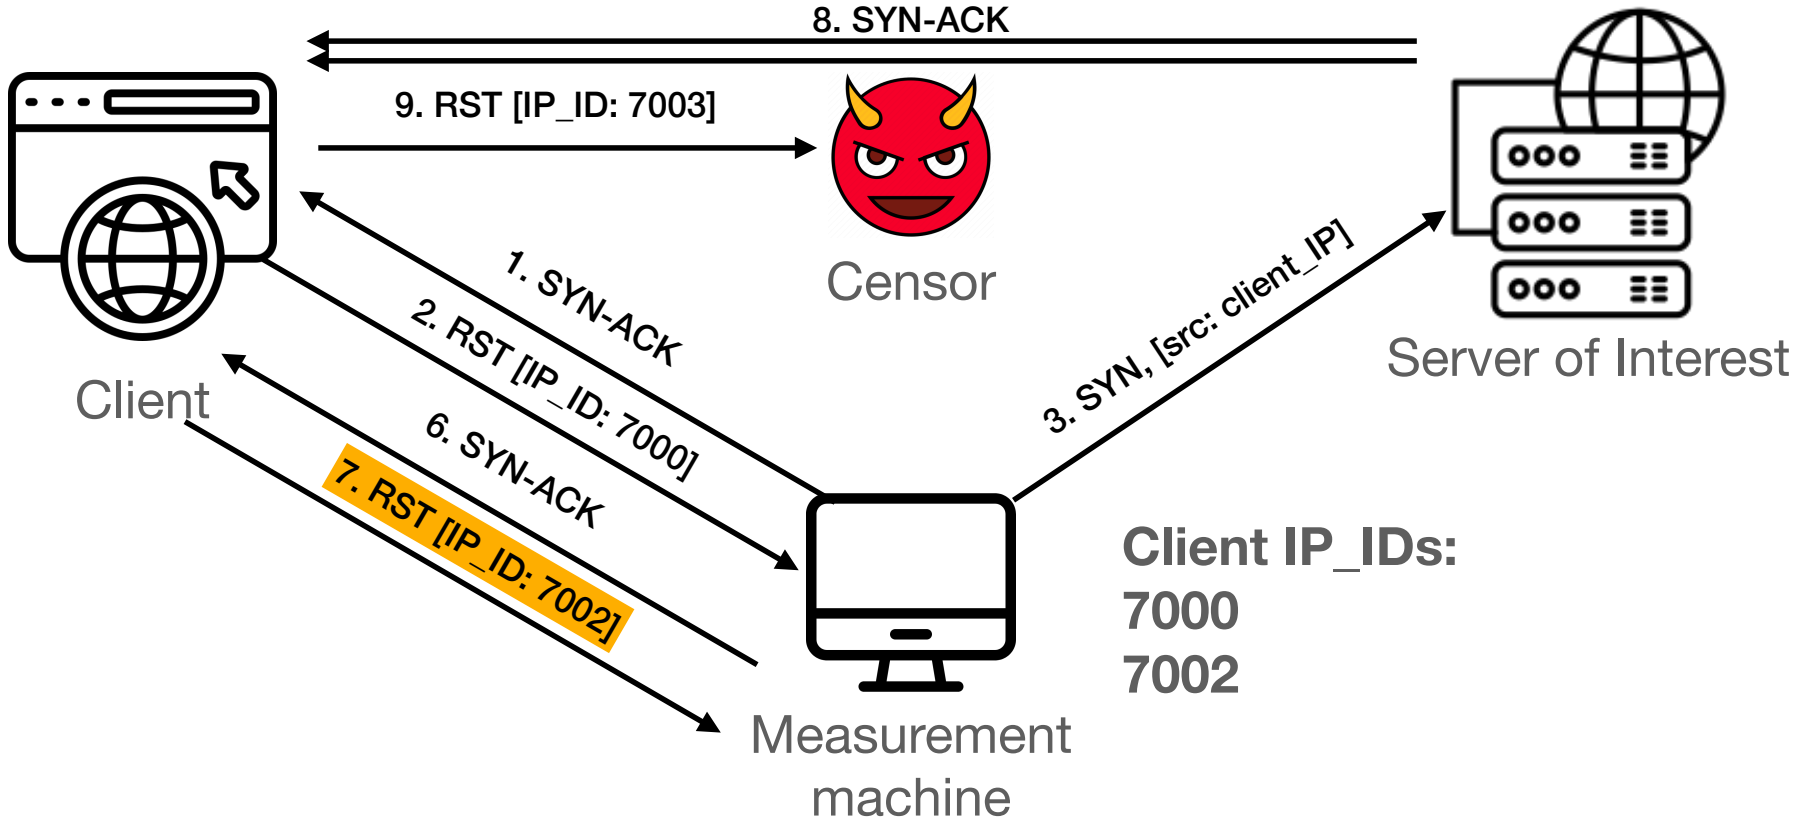

Spooky Scans

Client-to-server is blocked

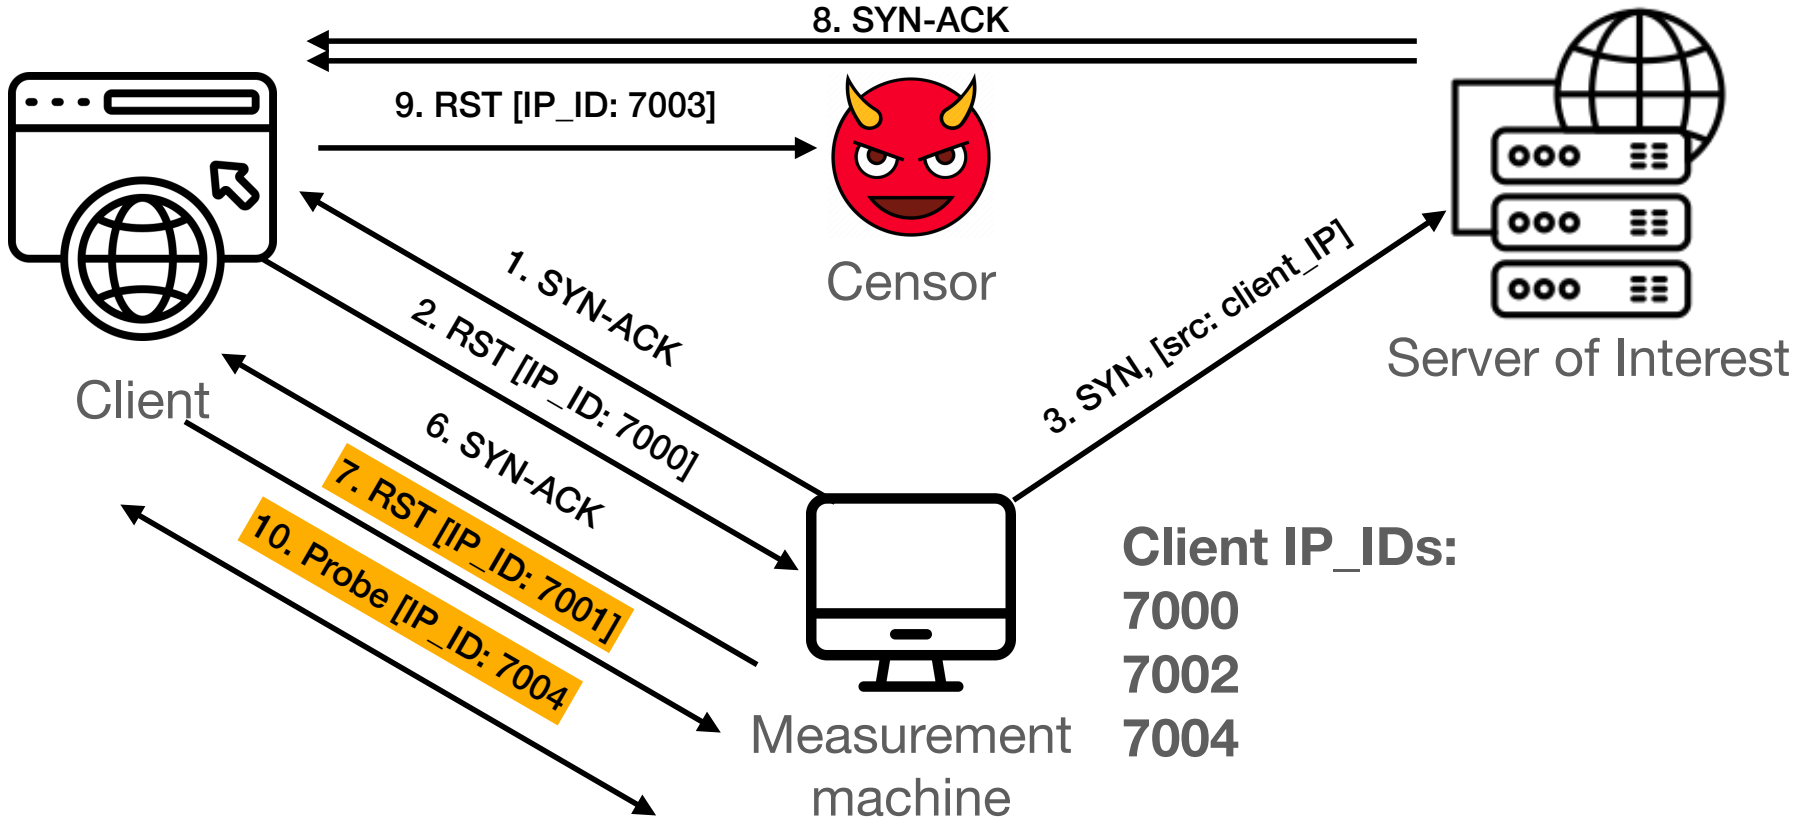

Spooky Scan Outcomes

Three distinct cases

No Blocking	Server-to-client blocking	Client-to-server blocking
IP_ID_1: 7000	IP_ID_1: 7000	IP_ID_1: 7000
IP_ID_2: 7002	IP_ID_2: 7001	IP_ID_2: 7002
IP_ID_3: 7003	IP_ID_3: 7002	IP_ID_3: 7004

Spooky scans can uncover IP blocking!

Augur: Internet-Wide Detection of Connectivity Disruptions

Paul Pearce[†*], Roya Ensafi[§*], Frank Li[†], Nick Feamster[§], Vern Paxson[†]
[†]University of California, Berkeley [§]Princeton University
{pearce, frankli, vern}@berkeley.edu {rensafi, feamster}@cs.princeton.edu

https://www.censoredplanet.org/projects/augur

Measuring Keyword Blocking

Using the Echo Protocol for Fun + Profit

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.

Measuring Keyword Blocking

Using the Echo Protocol for Fun + Profit

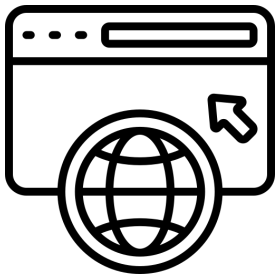

Client

Server of Interest

Measuring Keyword Blocking

Using the Echo Protocol for Fun + Profit

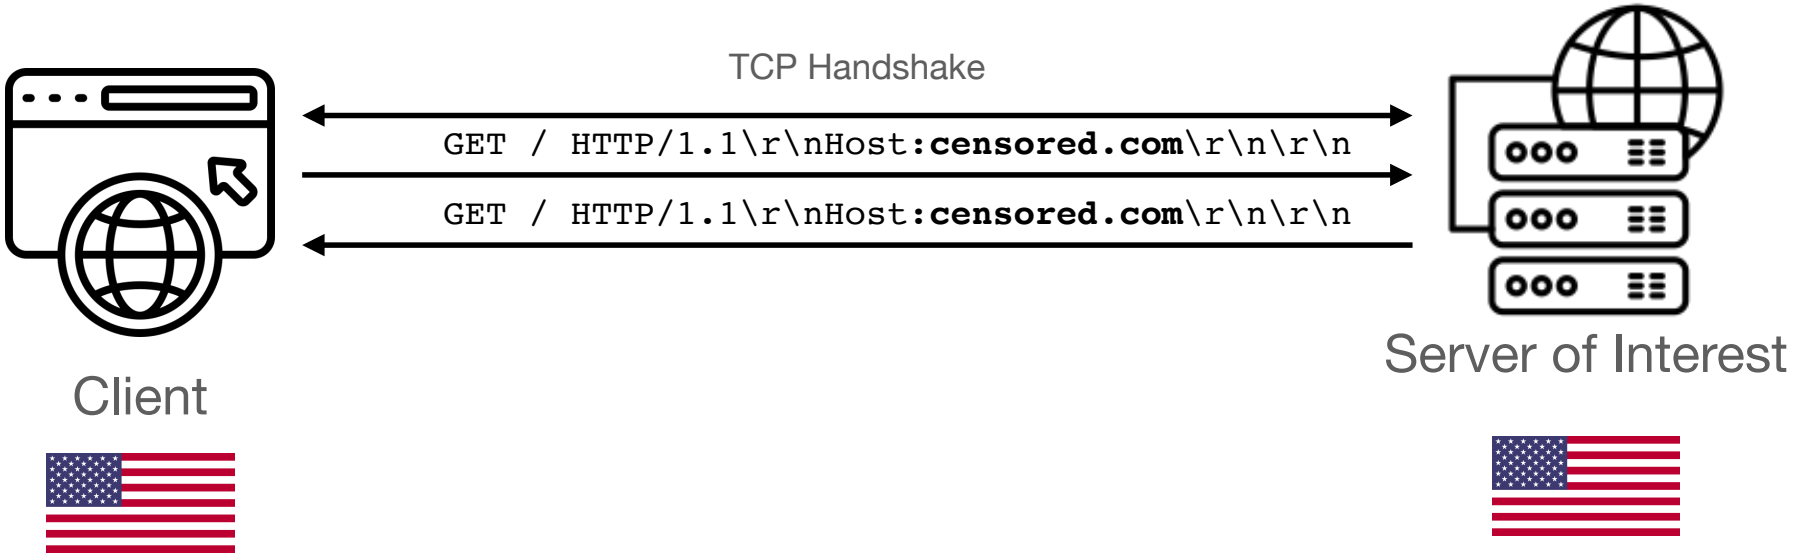

Measuring Keyword Blocking

Using the Echo Protocol for Fun + Profit

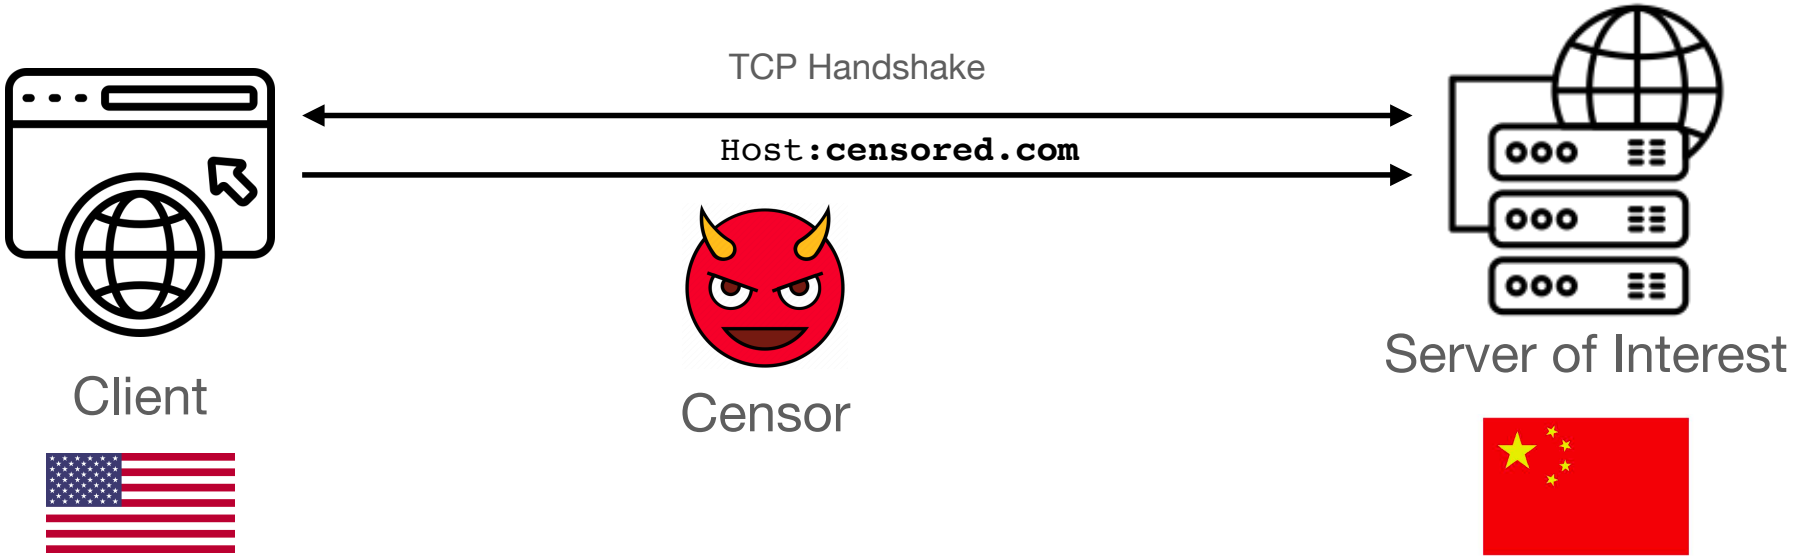

Measuring Keyword Blocking

Using the Echo Protocol for Fun + Profit

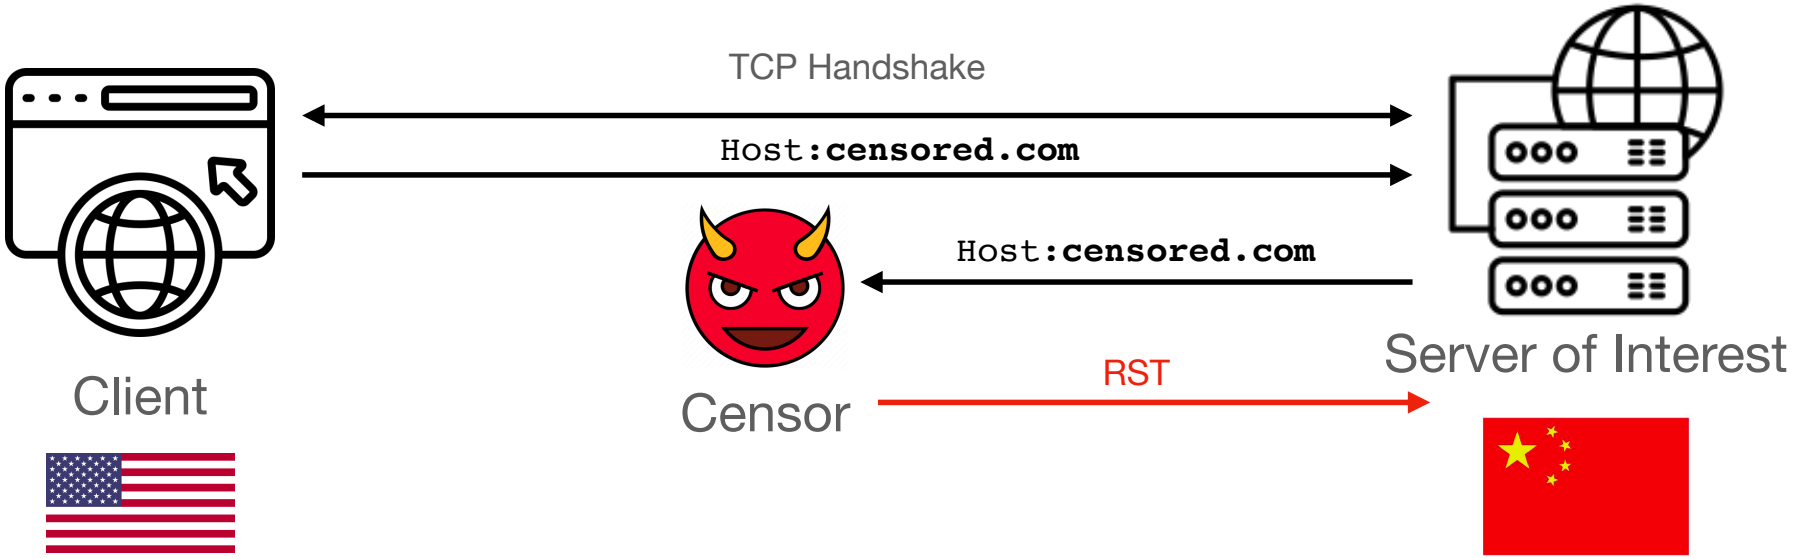

Something's up…

Quack, Hyperquack

Duck-themed censorship research

Quack: Scalable Remote Measurement of Application-Layer Censorship

Benjamin VanderSloot, Allison McDonald, Will Scott, J. Alex Halderman, and Roya Ensafi
University of Michigan
{benvds, amcdon, willscott, jhalderm, ensafi}@umich.edu

Measuring the Deployment of Network Censorship Filters at Global Scale

Ram Sundara Raman*, Adrian Stoll*, Jakub Dalek†, Reethika Ramesh*, Will Scott‡, Roya Ensafi*
*University of Michigan, {ramaks, adrs, reethika, ensafi}@umich.edu
†The Citizen Lab, University of Toronto, jakub.dalek@utoronto.ca
‡Independent, willscott@gmail.com

Censored Planet Observatory

Remote measurement platform

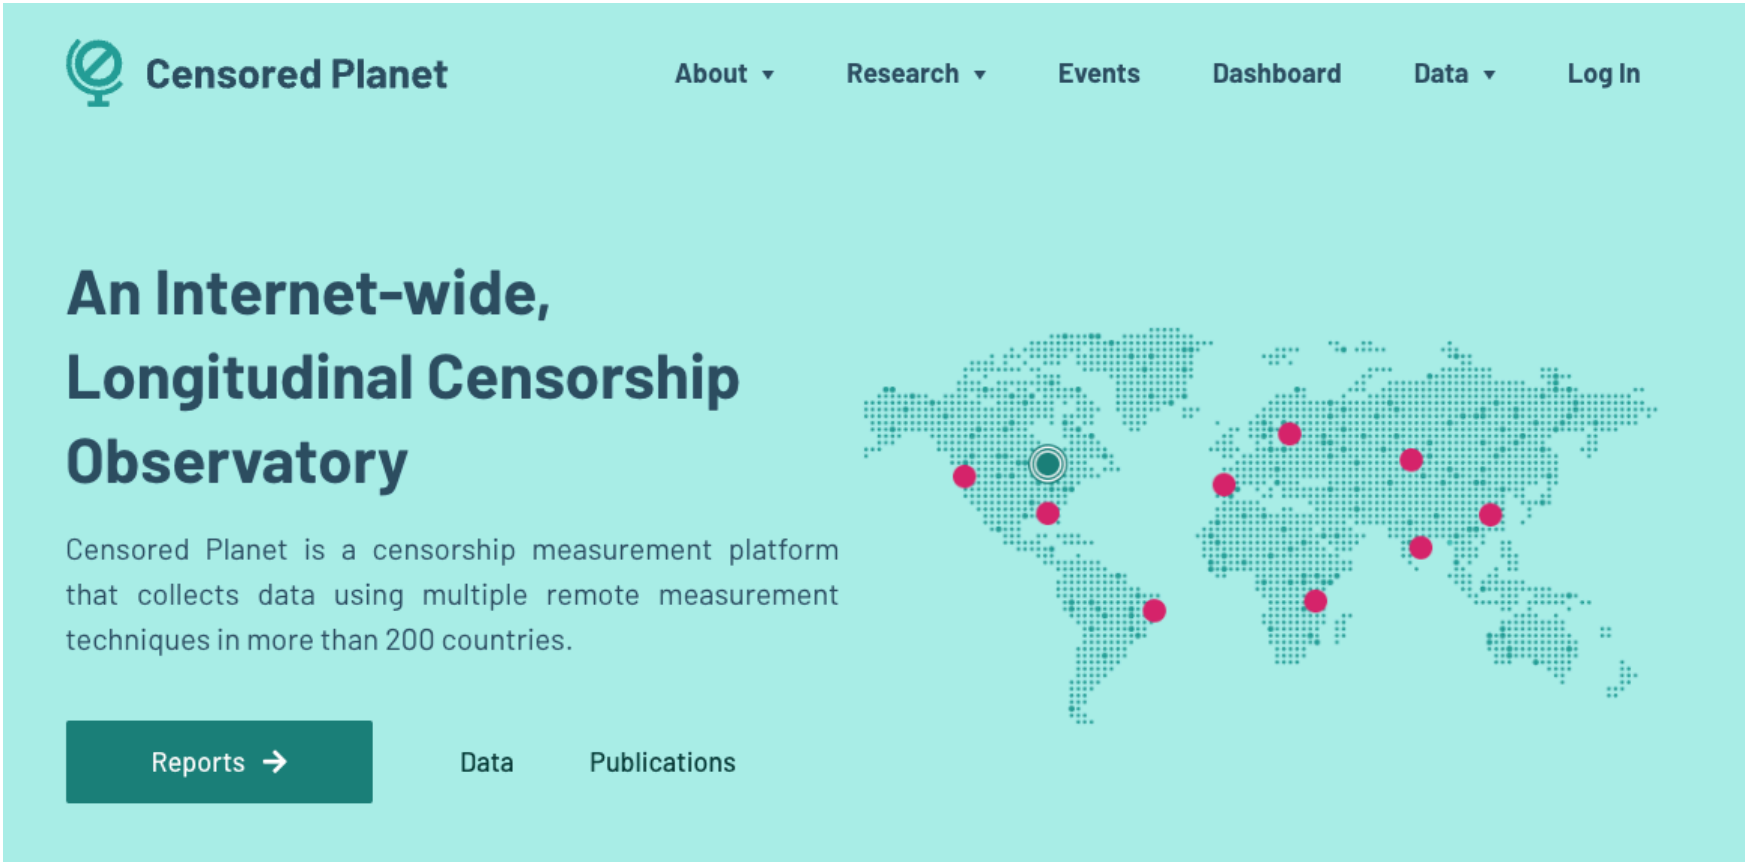

Censored Planet Observatory

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

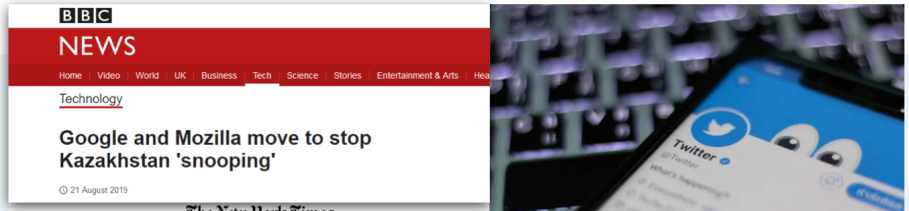

The New York Times

Study: Russia's Web-Censoring Tool Sets Pace for Imitators

By The Associated Press

Nov. 6, 2019

WASHINGTON — Russia is succeeding in imposing a highly effective internet censorship regime across thousands of disparate, privately owned providers in an effort also aimed at making government snooping pervasive, according to a study released Wednesday.

STORIES

Roskomnadzor successfully slows down Twitter. American researchers explained how he did it. They even found a small loophole for users - it's a pity that it's unlikely to help them

01:36, April 8, 2021

Source: Meduza

Real-time monitor tracks the growing use of network filters for censorship

February 21, 2020

The team says their framework can scalably and semi-automatically monitor the use of filtering technologies for censorship at global scale.

Censored Planet's Impact

Rapid response studies

Kazakhstan's National TLS Interception

- **July 17, 2019**: Government started intercepting large fraction of HTTPS traffic within its borders.

Publication -Investigating Large Scale HTTPS Interception in Kazakhstan; R. Sundara Raman, L. Evdokimov, E. Wurstrow, J. A. Halderman, and R. Ensafi; ACM Internet Measurement Conference (IMC), 2020

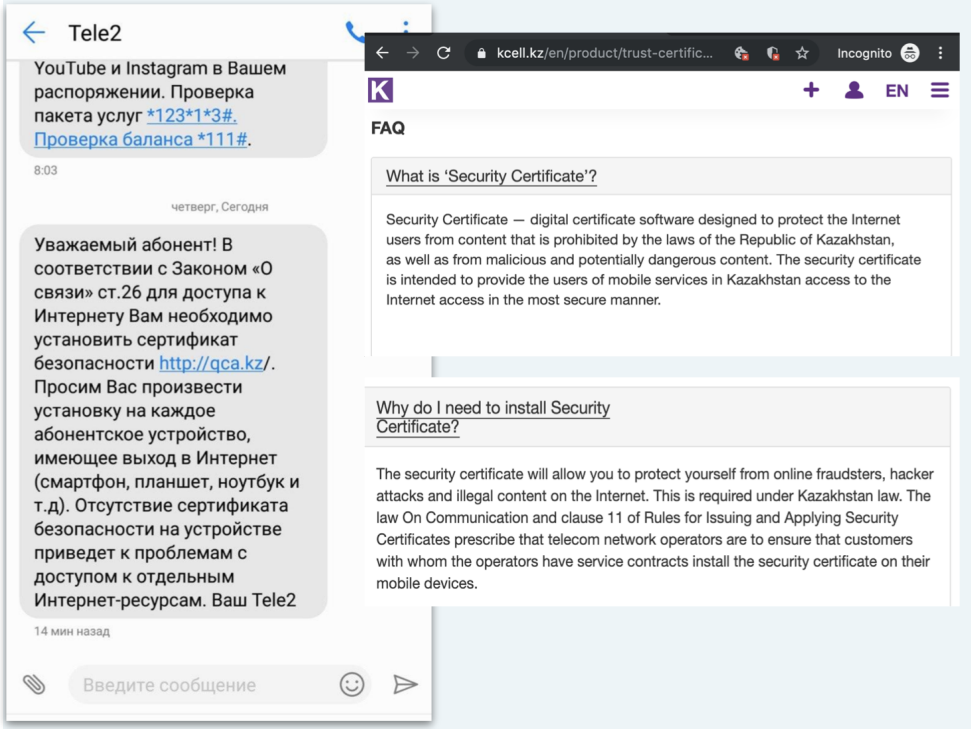

Kazakhstan's National TLS Interception

- Censored Planet detected and studied the interception
- Pilot test
- Interception occurring at large state-owned network and targeting social media websites

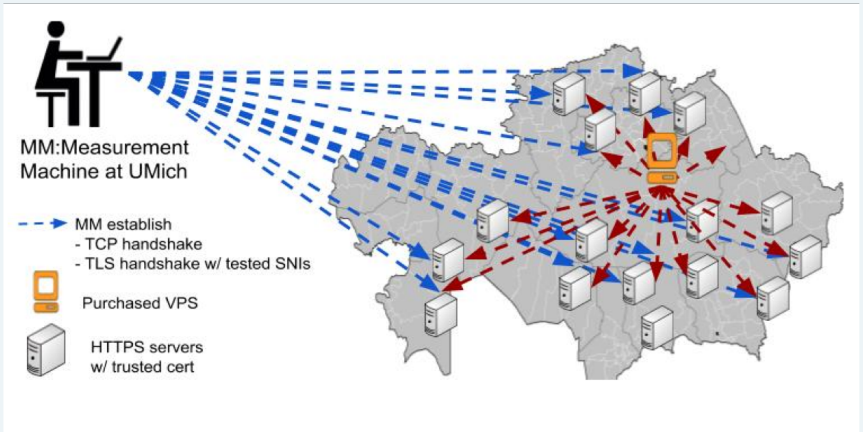

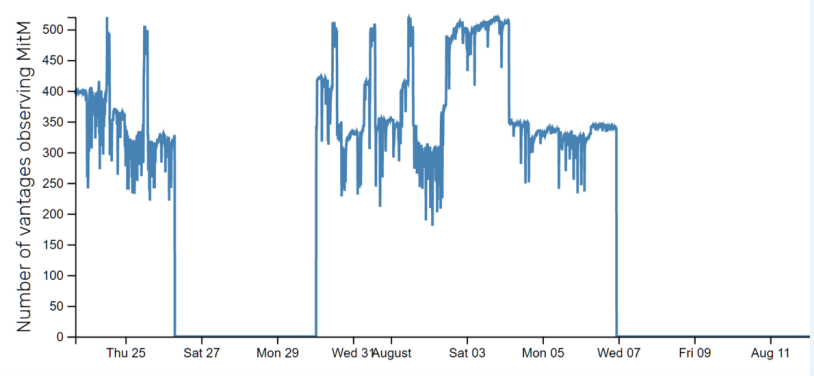

Browsers Take a Stand Against Interception!

The use of 'Qaznet Trust Network' root CA certificate in Chrome, Firefox, and Safari is now prevented

Measurement systems are still being improved upon!

CERTainty: Detecting DNS Manipulation at Scale using TLS Certificates

Elisa Tsai* Deepak Kumar† Ram Sundara Raman* Gavin Li* Yael Eiger* Roya Ensafi*
*Censored Plaent, University of Michigan †Stanford University

Many Roads Lead To Rome: How Packet Headers Influence DNS Censorship Measurement

Abhishek Bhaskar Paul Pearce

Georgia Institute of Technology
{abhaskar, pearce}@gatech.edu

Circumventing Internet Censorship

Circumventing Censors

Proxying requests through “safe” servers

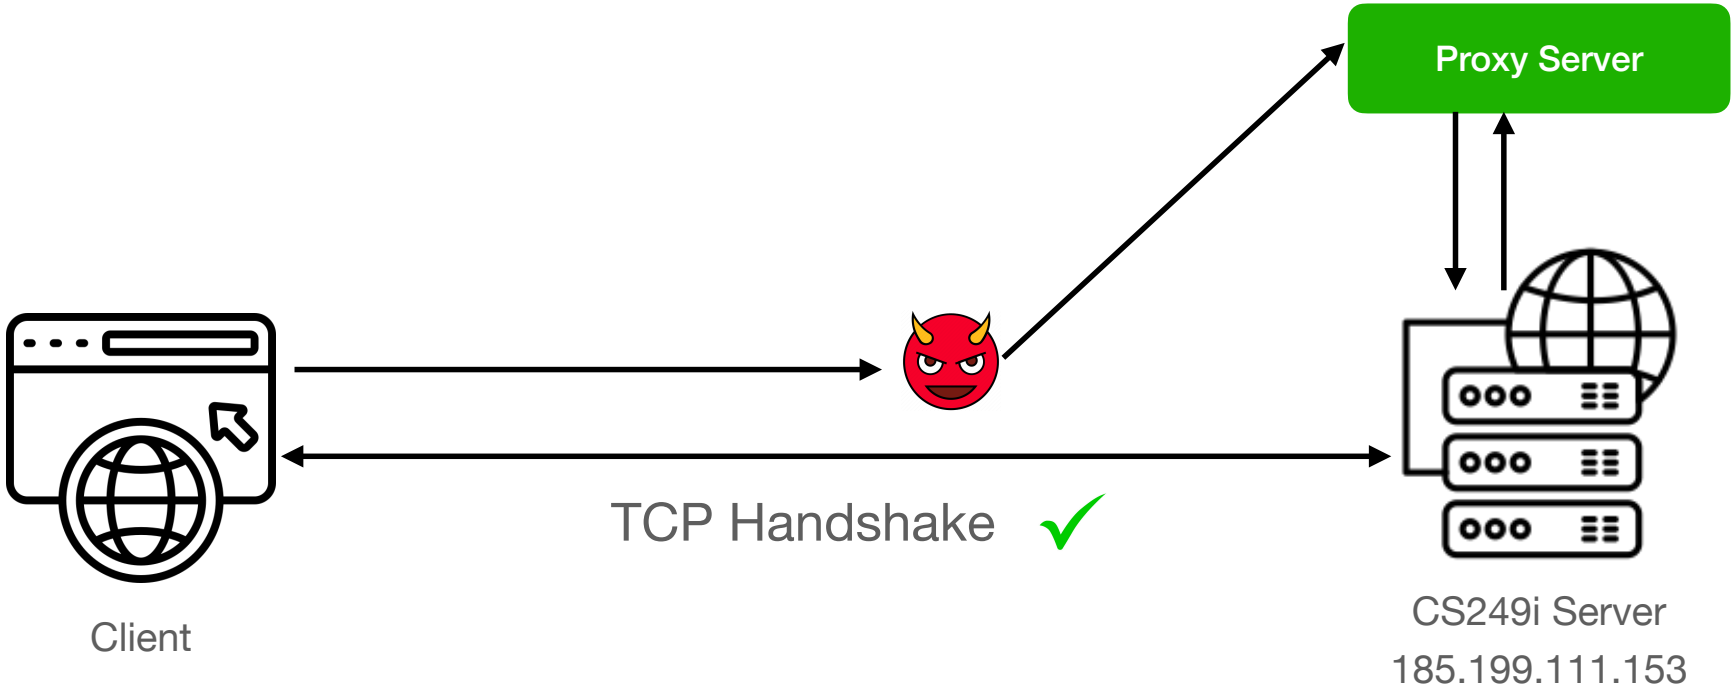

Circumventing Censors

Proxying requests through “safe” servers is easy to detect

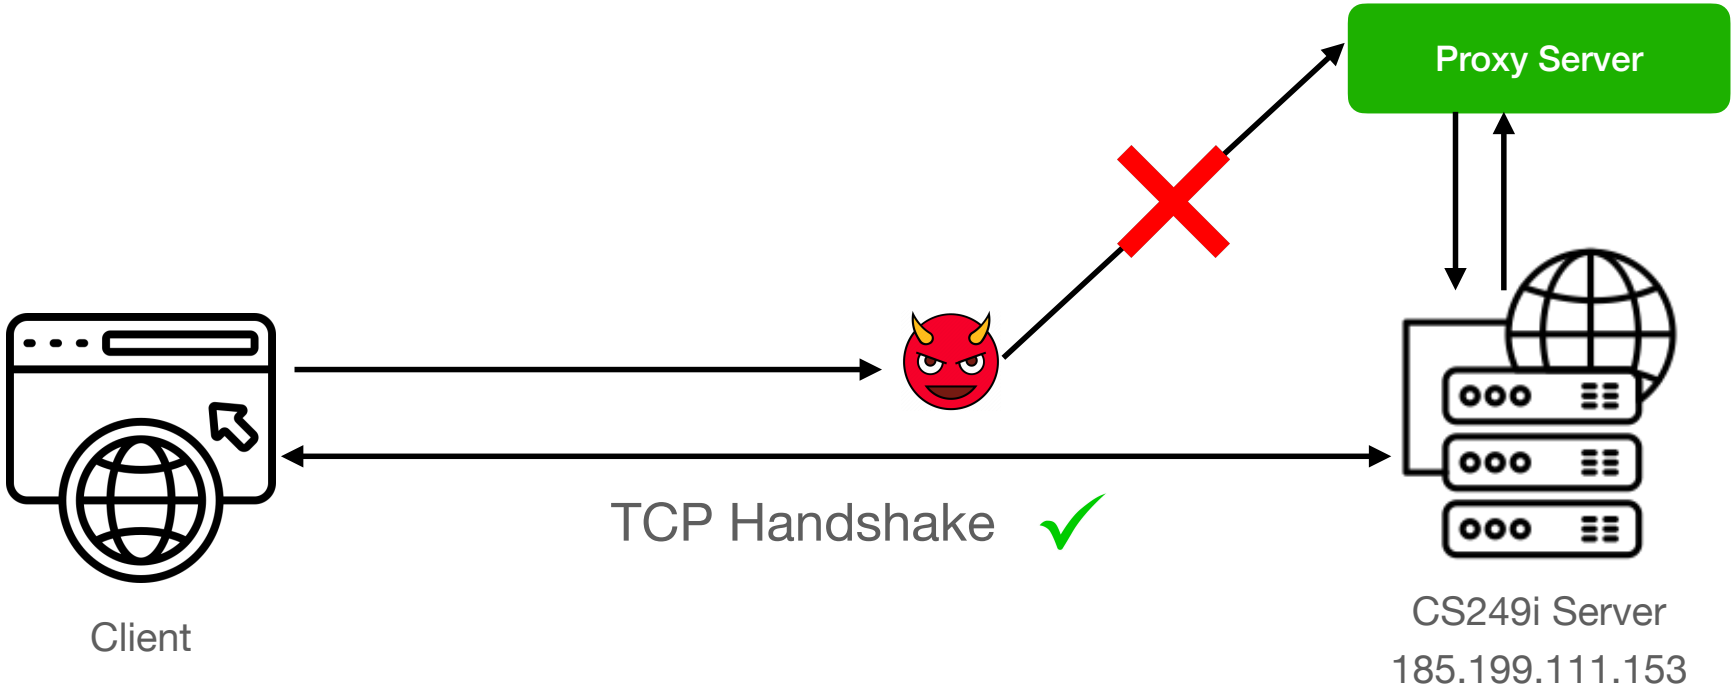

Circumventing Censors

Imitating non-censored protocols

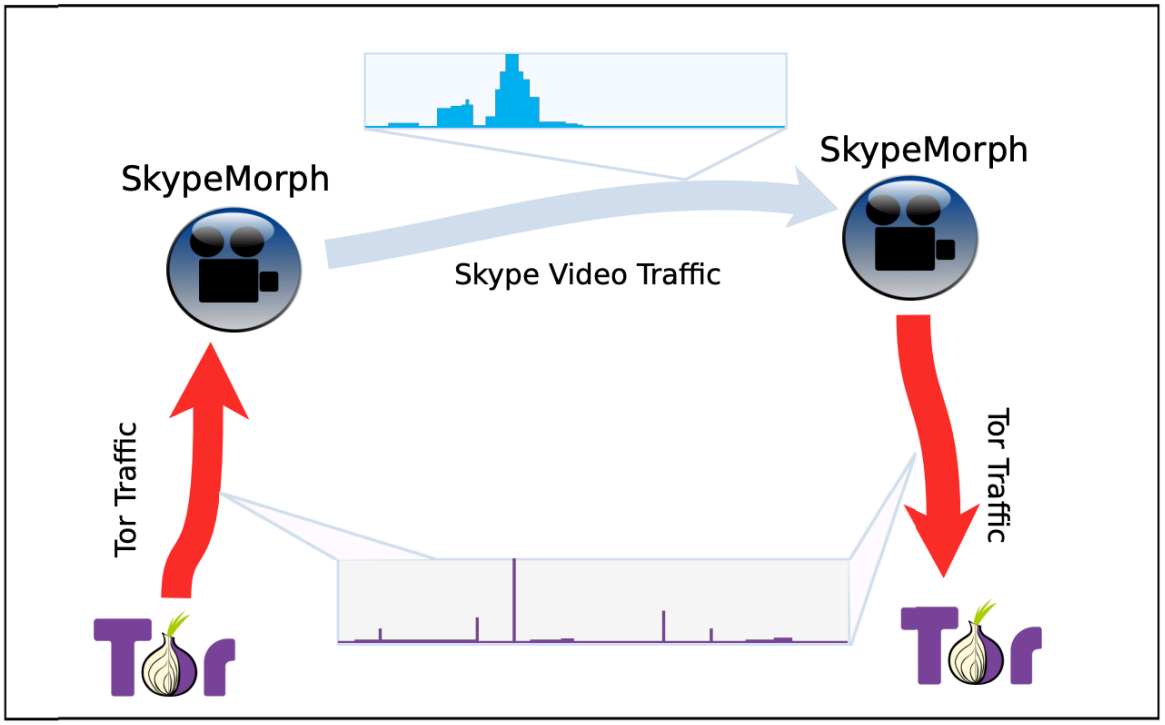

Circumventing Censors

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

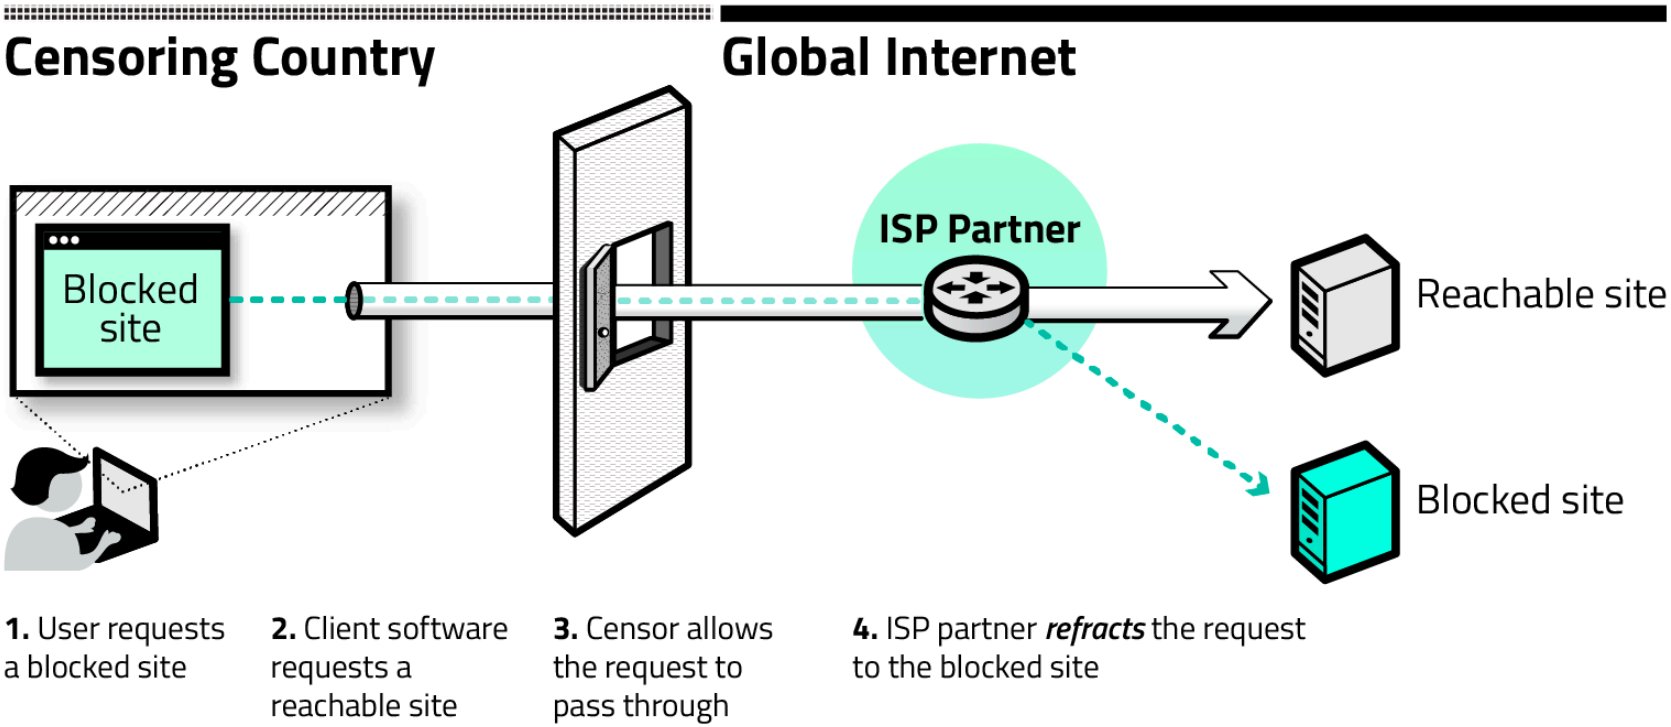

https://refraction.network/

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/

Recap

Censorship techniques, measurements, circumvention

- Internet censorship is on the rise, we're seeing even small countries with deep capabilities
- Many techniques: DNS manipulation, IP blocking, App-layer blocking
- Censorship can be measured through **volunteers** or through carefully designed **remote-measurements**
  - Orgs like OONI, CensoredPlanet are measuring censorship longitudinally!
- Circumvention techniques are moving beyond *simple proxies* towards ISP-mediated refraction, but it's a cat + mouse game

Recap

Censorship techniques, measurements, circumvention

- Internet censorship is on the rise, we're seeing even small countries with deep capabilities
- Many techniques: DNS manipulation, IP blocking, App-layer blocking
- Censorship can be measured through **volunteers** or through carefully designed **remote-measurements**
  - Orgs like OONI, CensoredPlanet are measuring censorship longitudinally!
- Circumvention techniques are moving beyond *simple proxies* towards ISP-mediated refraction, but it's a cat + mouse game

Questions?

Deepak Kumar

kumarde@ucsd.edu

@_kumarde