Is Nobody There? Good! Globally Measuring Connection Tampering without Responsive Endhosts

Sadia Nourin Erik Rye Kevin Bock Nguyen Phong Hoang Dave Levin

MAX PLANCK INSTITUTE
FOR INFORMATICS

UNIVERSITY OF
MARYLAND

THE UNIVERSITY
OF BRITISH COLUMBIA

Network Interference

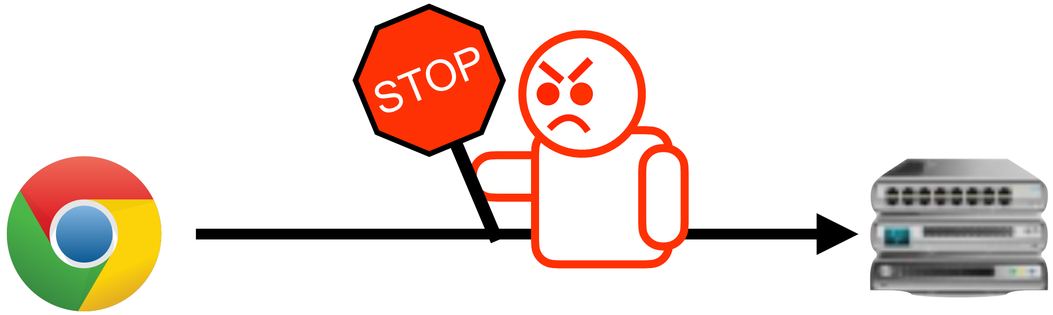

Nation-state censorship

University firewalls

Corporate IDS, IPS

Network Interference

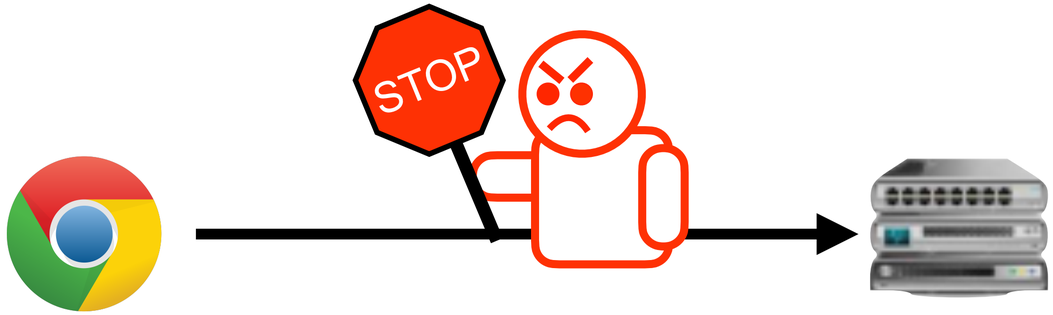

Nation-state censorship

University firewalls

Corporate IDS, IPS

Policies and mechanisms can vary drastically from across networks

Measuring network interference globally

Broad

Should measure as many networks as possible

Deep

Should test as many domains as possible

Policies and mechanisms can vary drastically from across networks

Global measurements rely on participating end-hosts

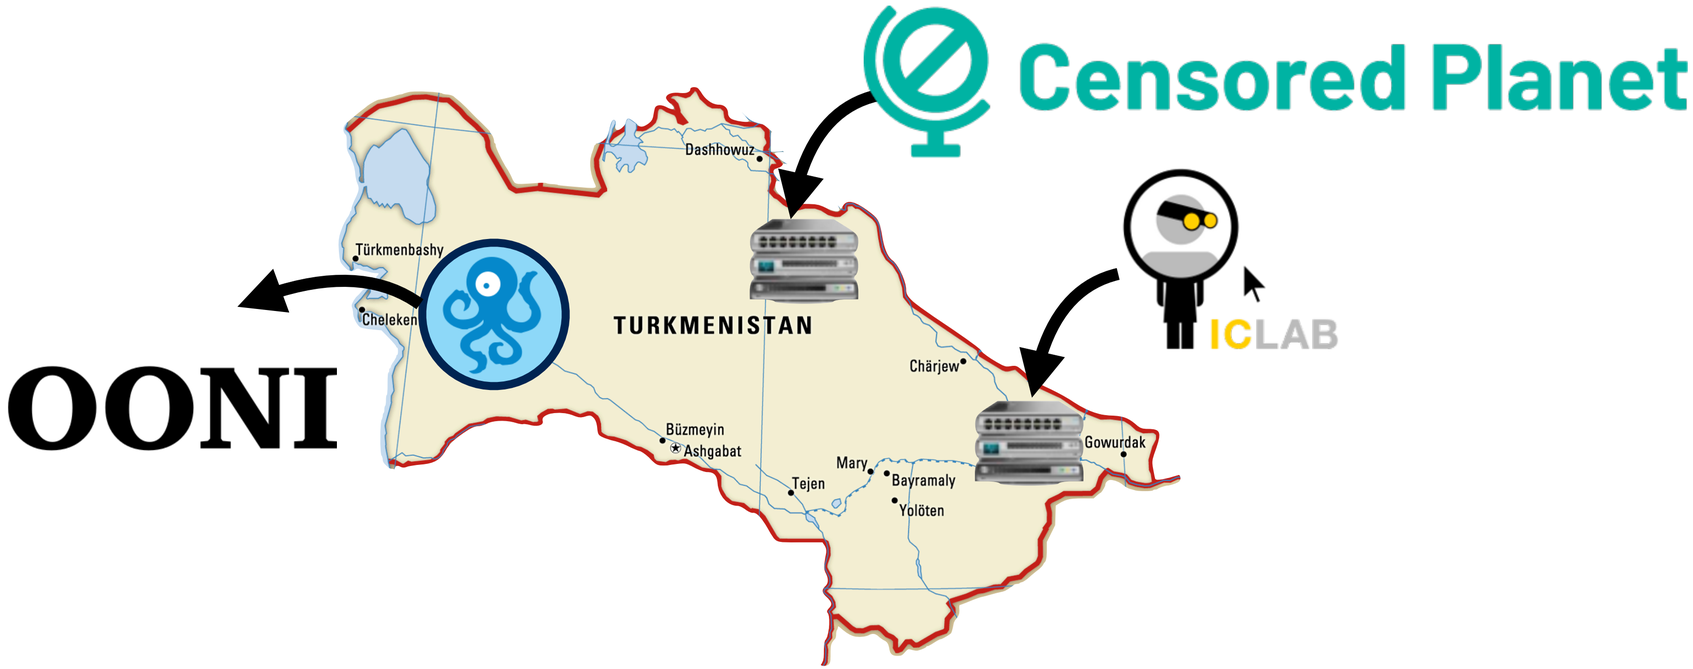

Global measurements rely on participating end-hosts

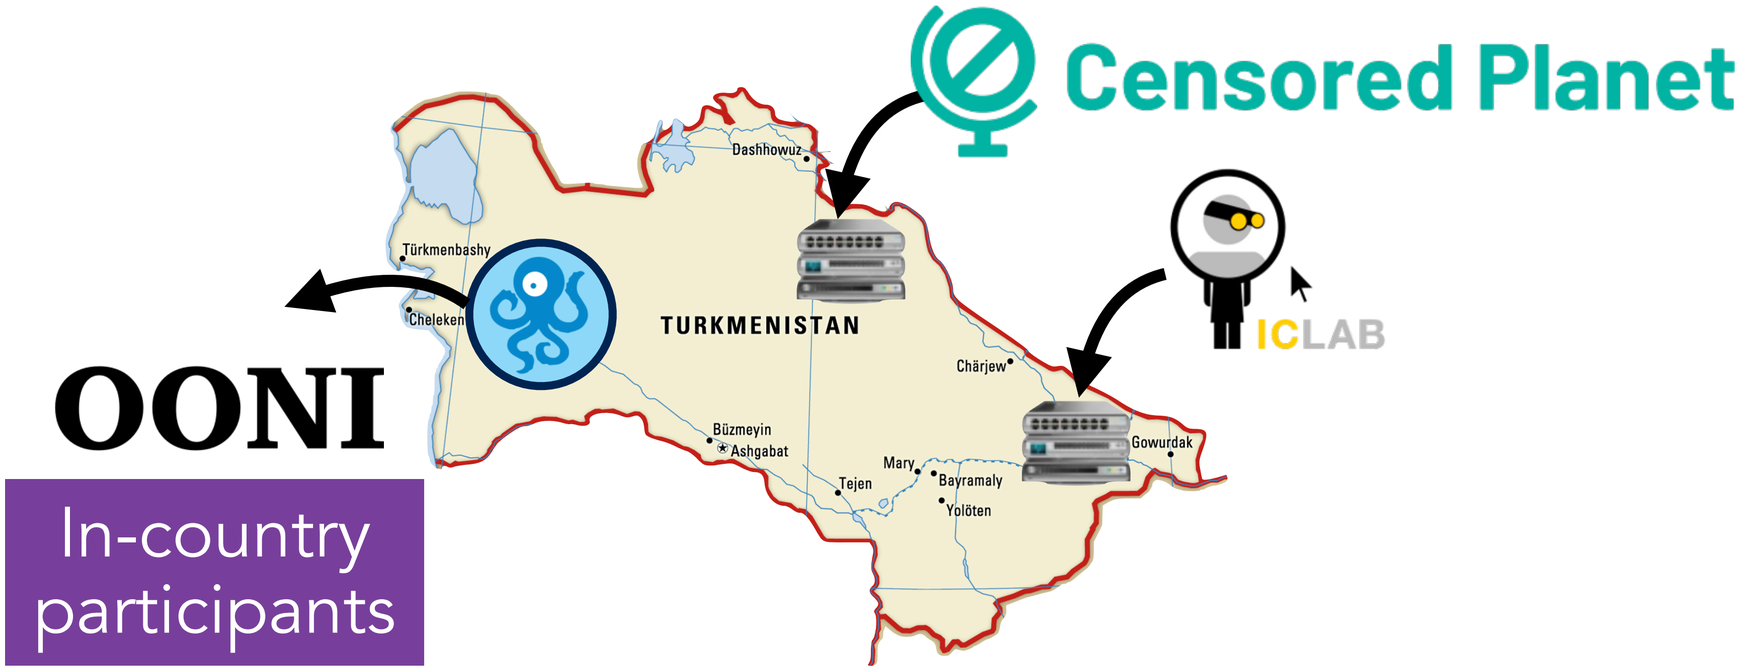

Global measurements rely on participating end-hosts

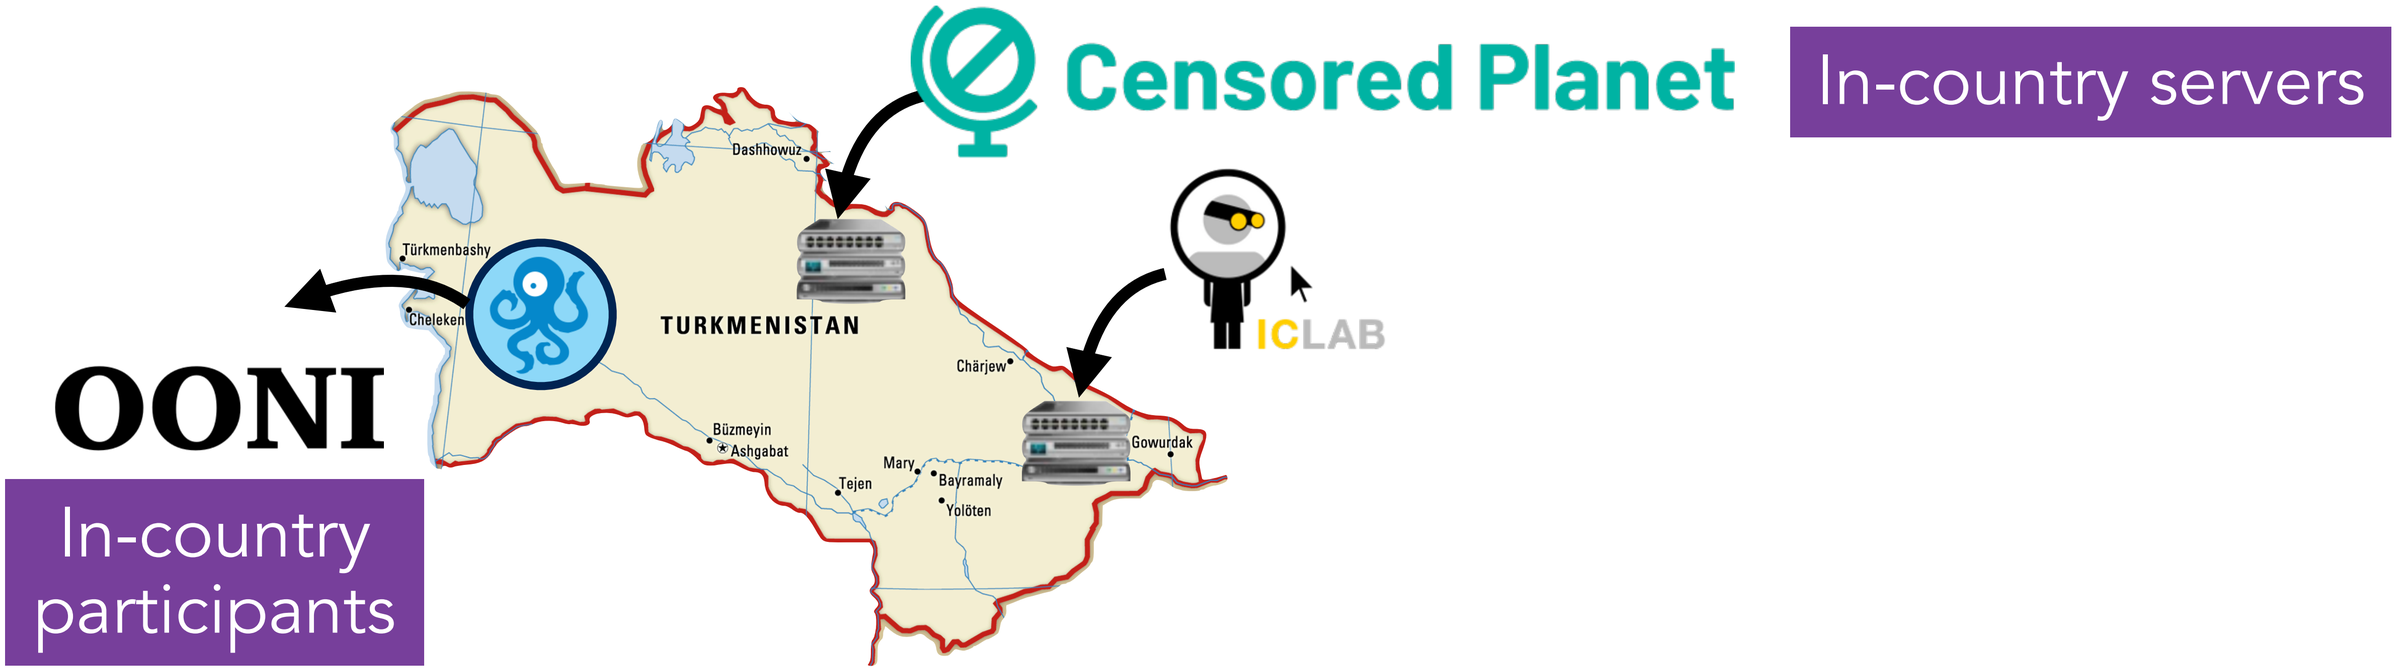

Global measurements rely on participating end-hosts

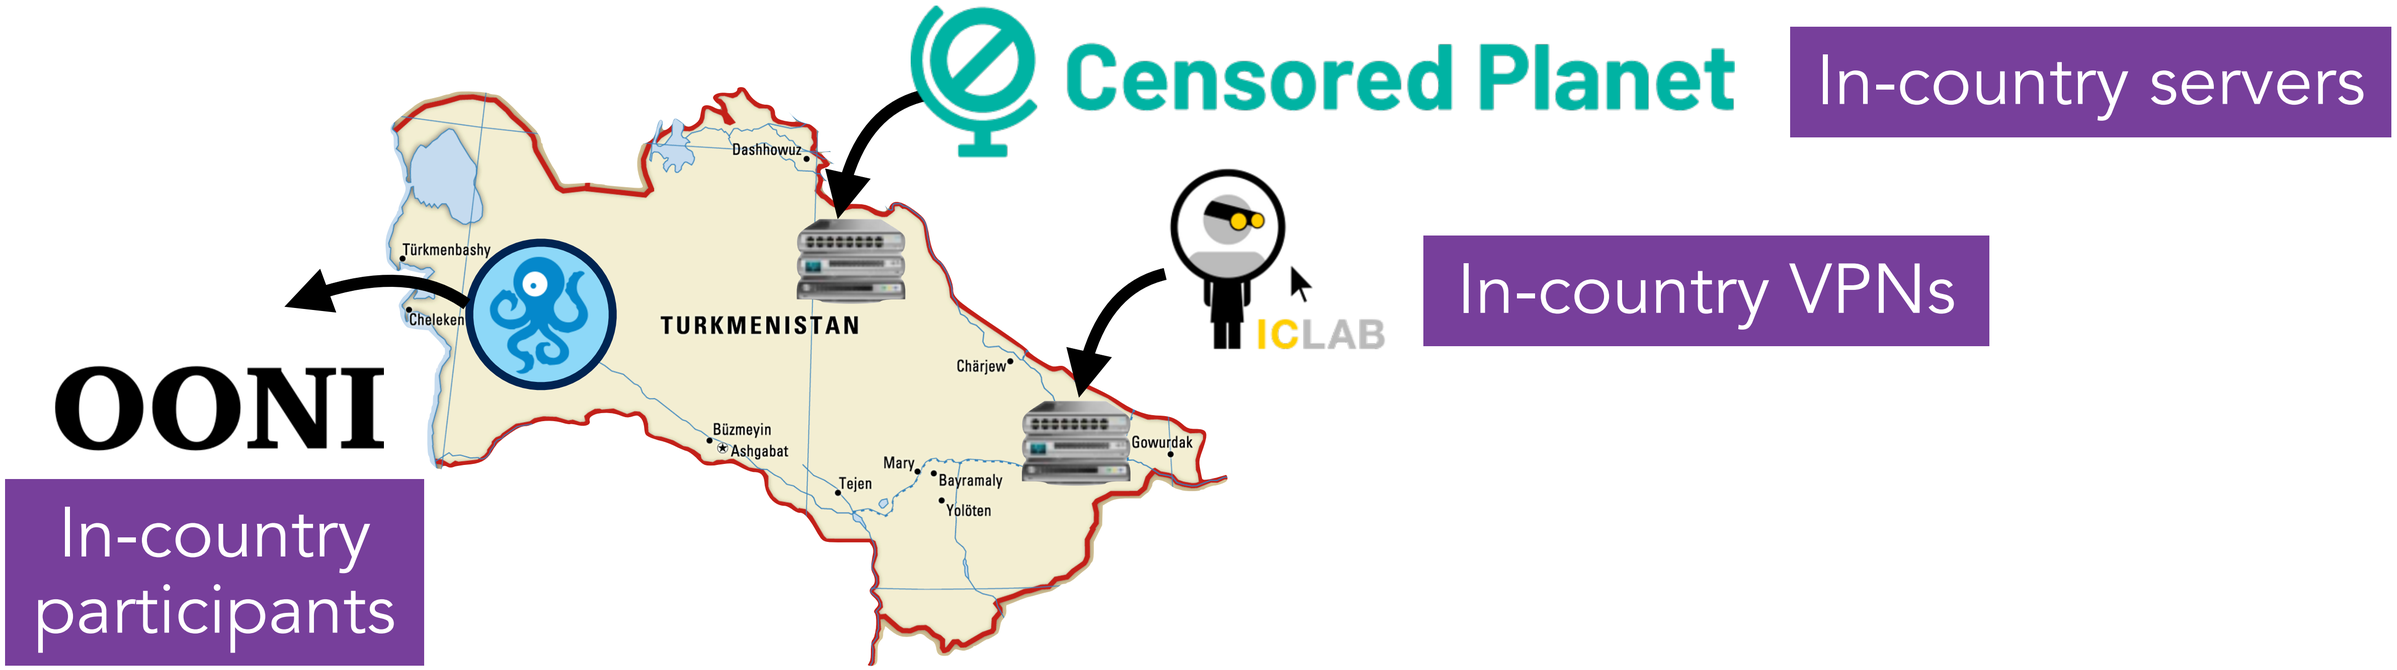

Global measurements rely on participating end-hosts

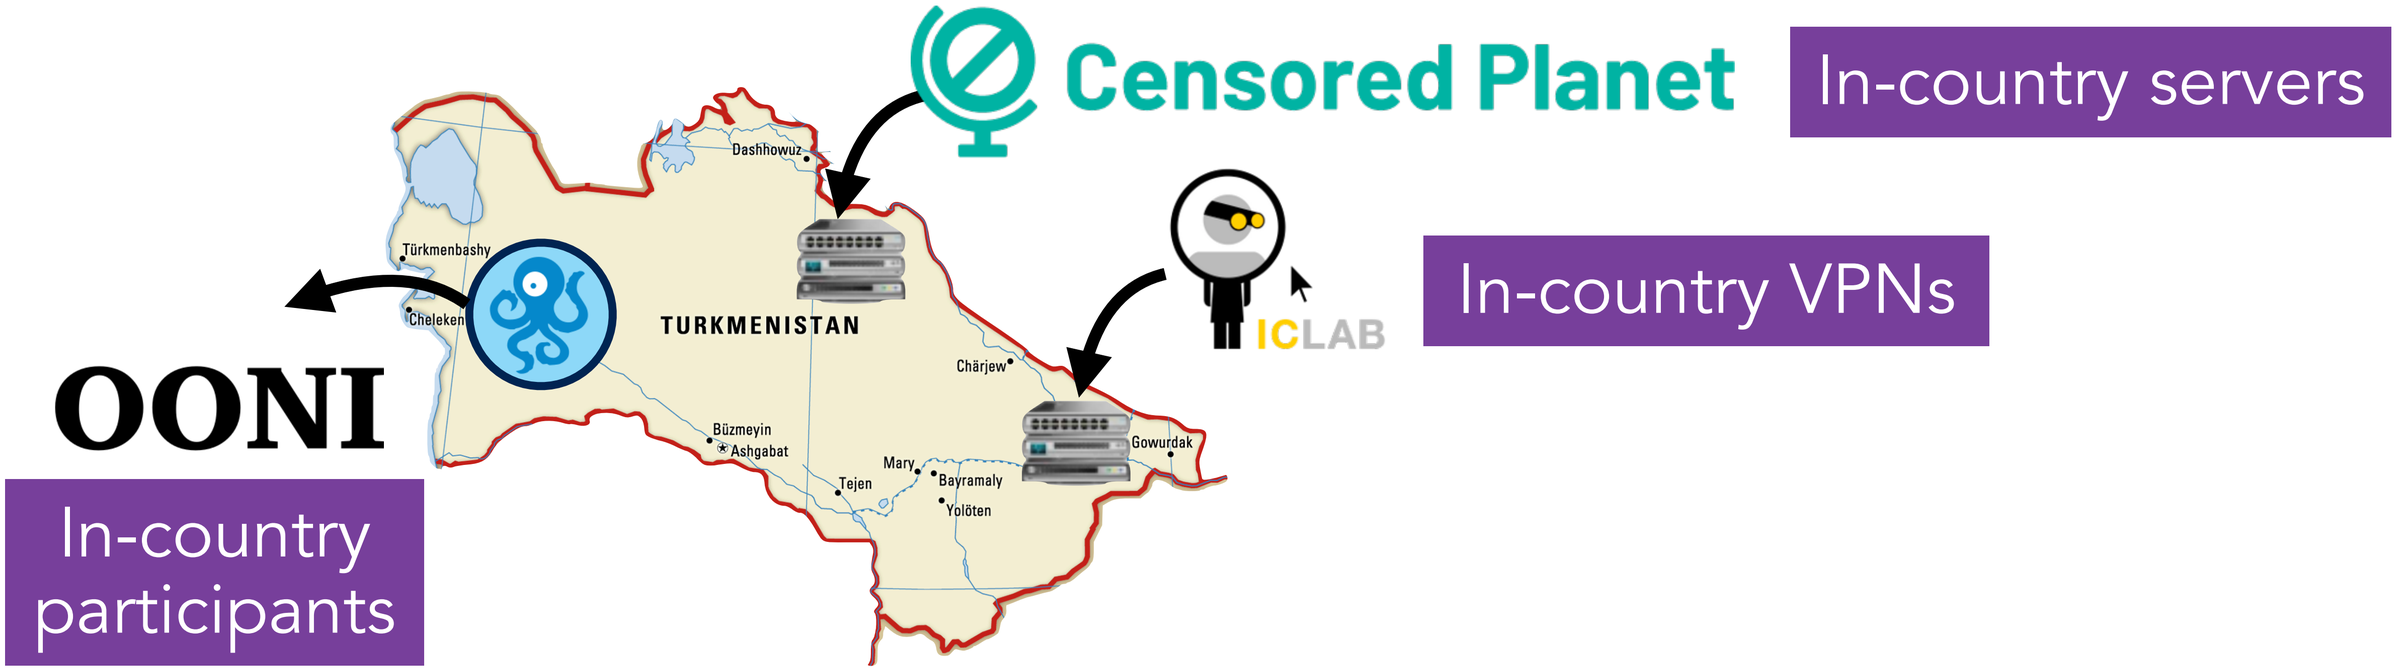

Broad?

Limited to where servers and participants *safely* are

Global measurements rely on participating end-hosts

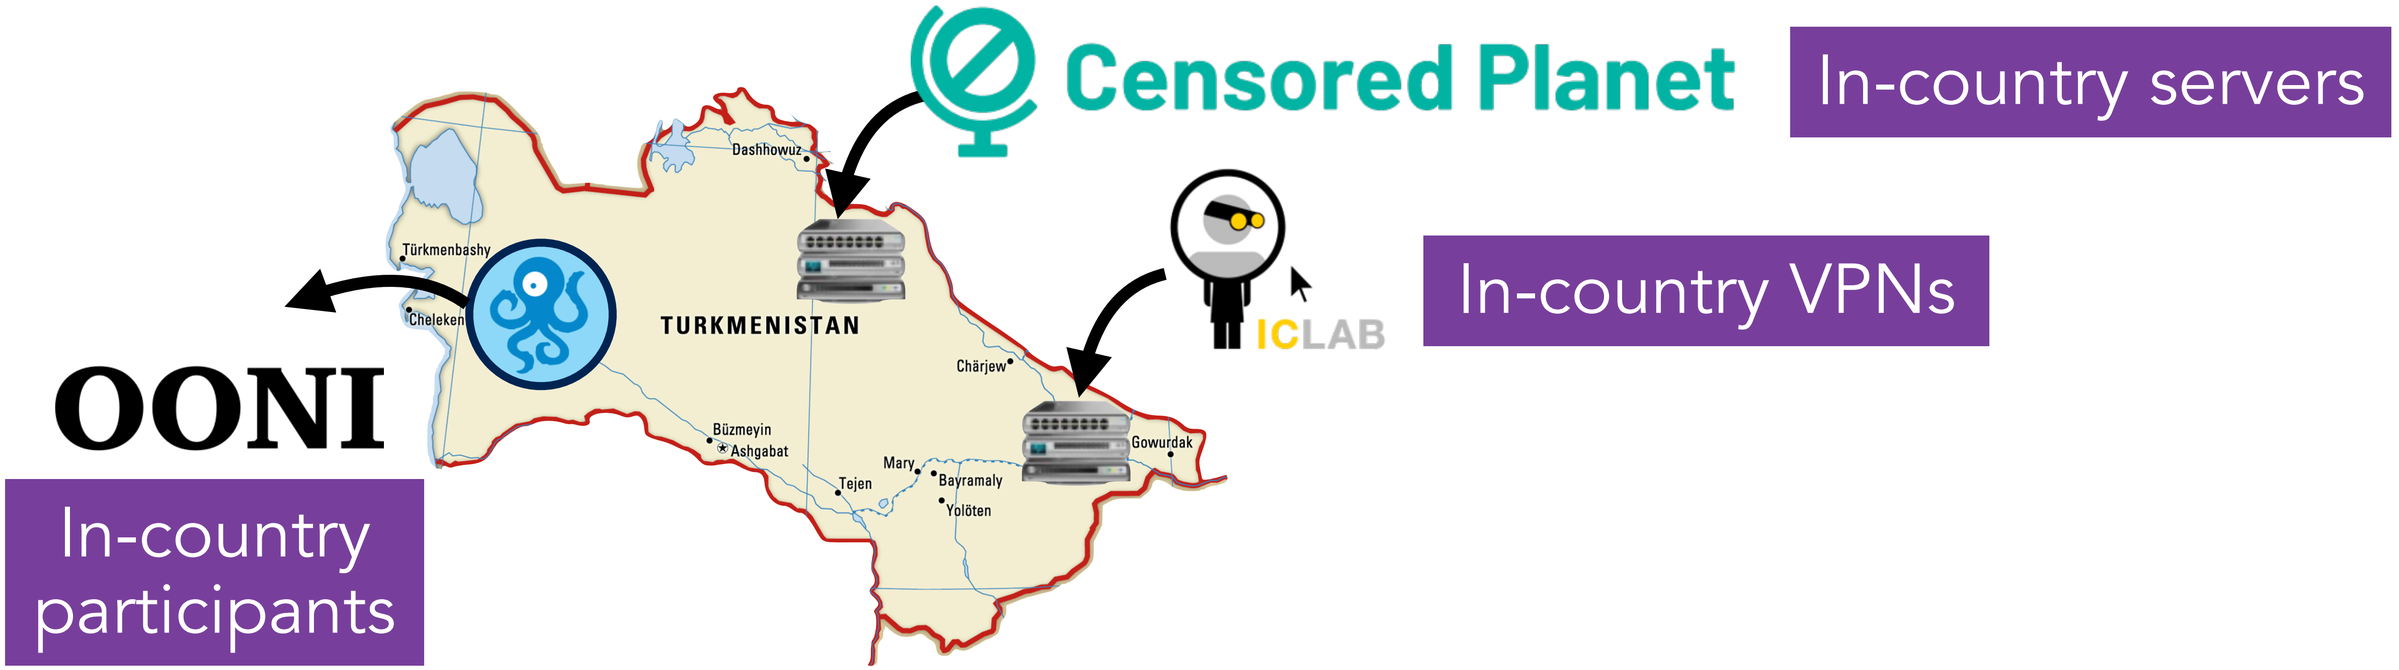

Broad? Limited to where servers and participants *safely* are

Deep? Limited to rates that don't affect live end-hosts

Global measurements rely on participating end-hosts

Censored Planet

In-country servers

ICLAB

In-country VPNs

OONI

In-country participants

Difficult to measure:

Small populations

Low Internet penetrations

IPv6

No global measurements

Highly oppressive regimes

Poor Internet infrastructure

Mint

Measuring **I**nterference with **N**on-responsive **T**argets

Mint sends measurements to ***non-responsive IP addresses***

Mint

Measuring **I**nterference with **N**on-responsive **T**argets

Mint sends measurements to ***non-responsive IP addresses***

Broad — Almost every network has non-responsive hosts; **esp. IPv6**

Mint

Measuring **I**nterference with **N**on-responsive **T**argets

Mint sends measurements to ***non-responsive IP addresses***

Broad	Almost every network has non-responsive hosts; **esp. IPv6**
Deep	Higher rates don't affect end-hosts

Measuring **I**nterference with **N**on-responsive **T**argets

First **global, active** measurement of HTTP(S) network interference **without endhosts**

First network interference measurement **over IPv6 at scale**

Mint builds off of two previous studies

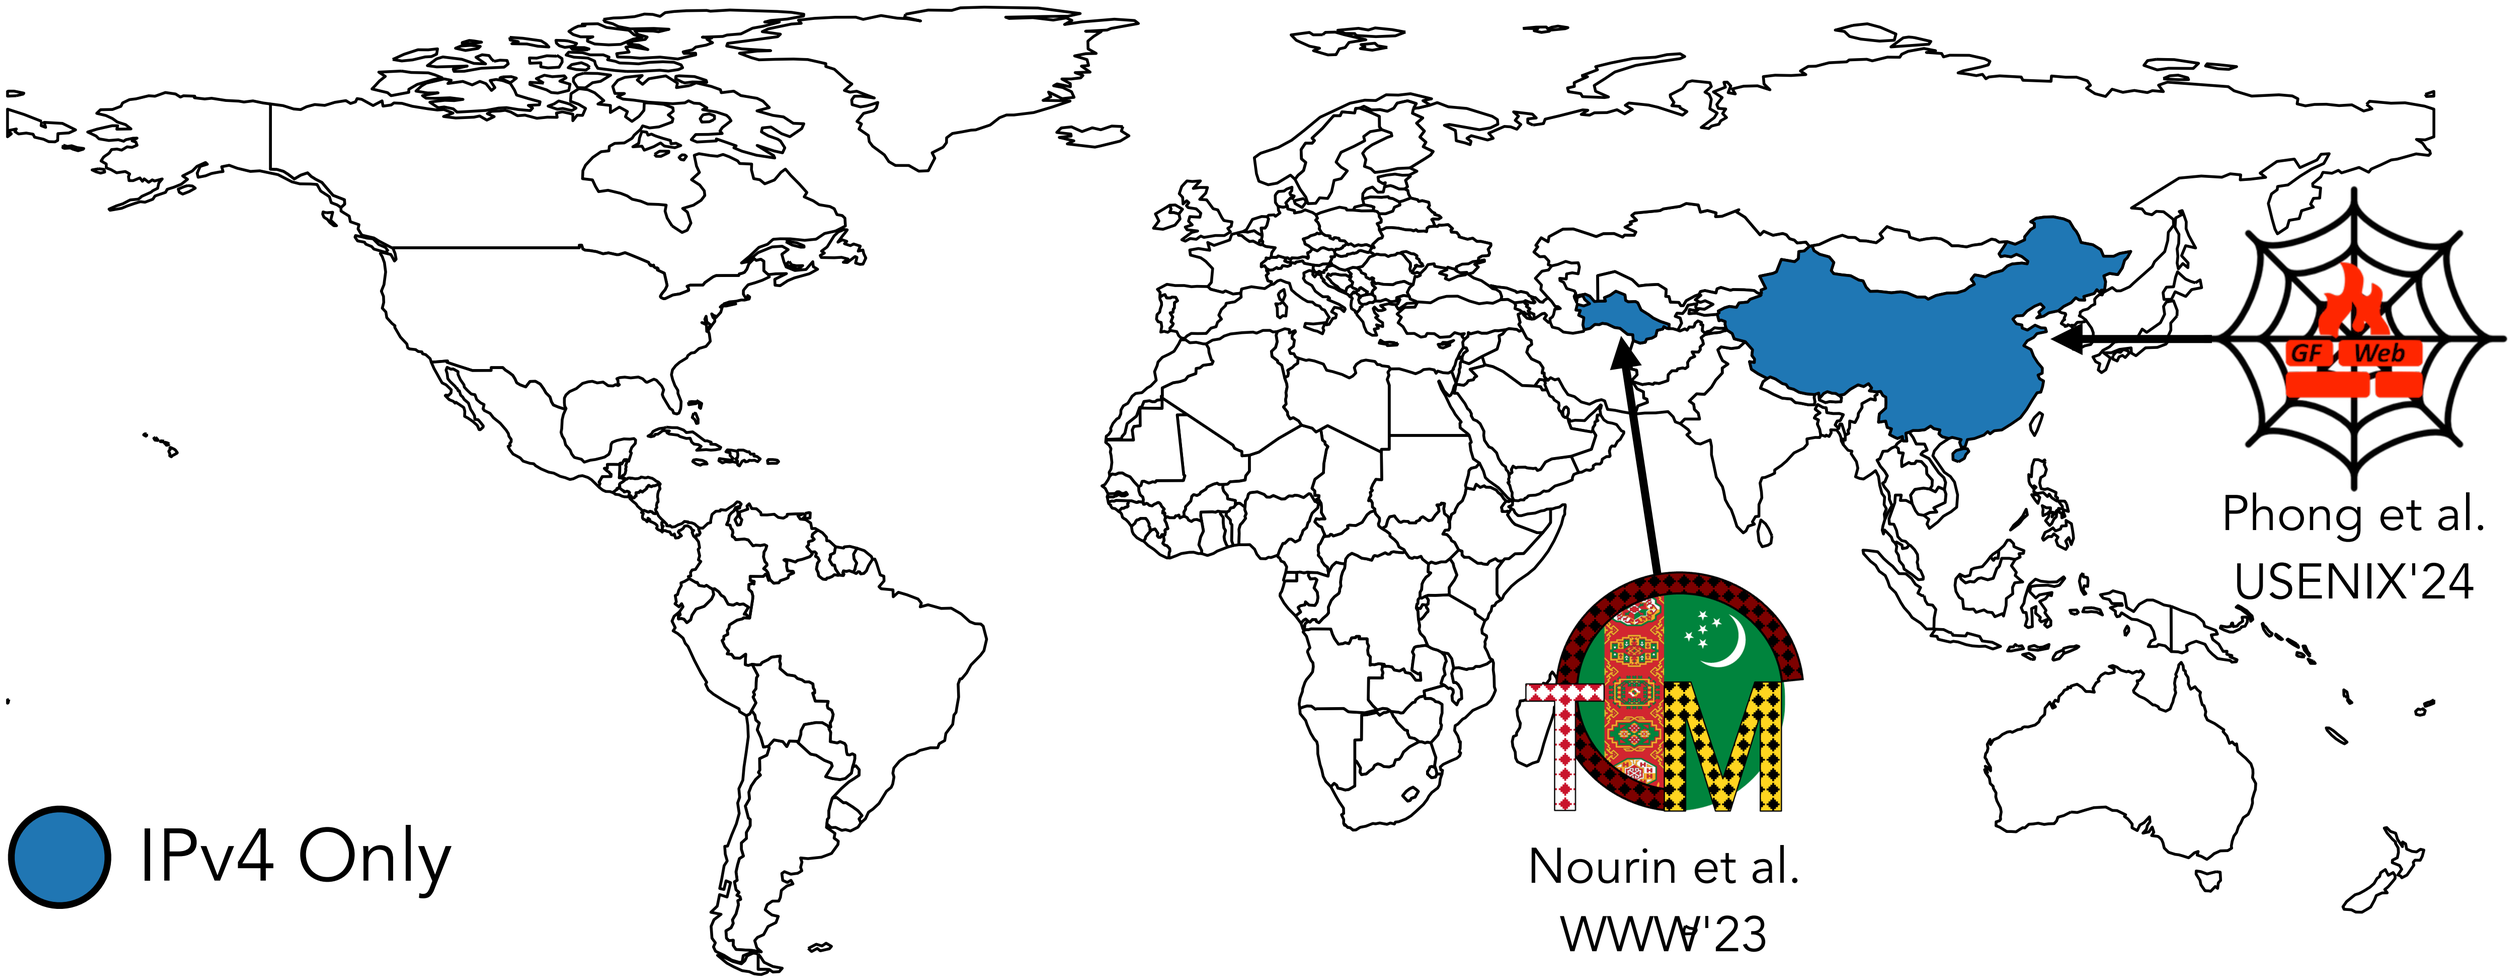

Mint measures IPv4 and IPv6 globally

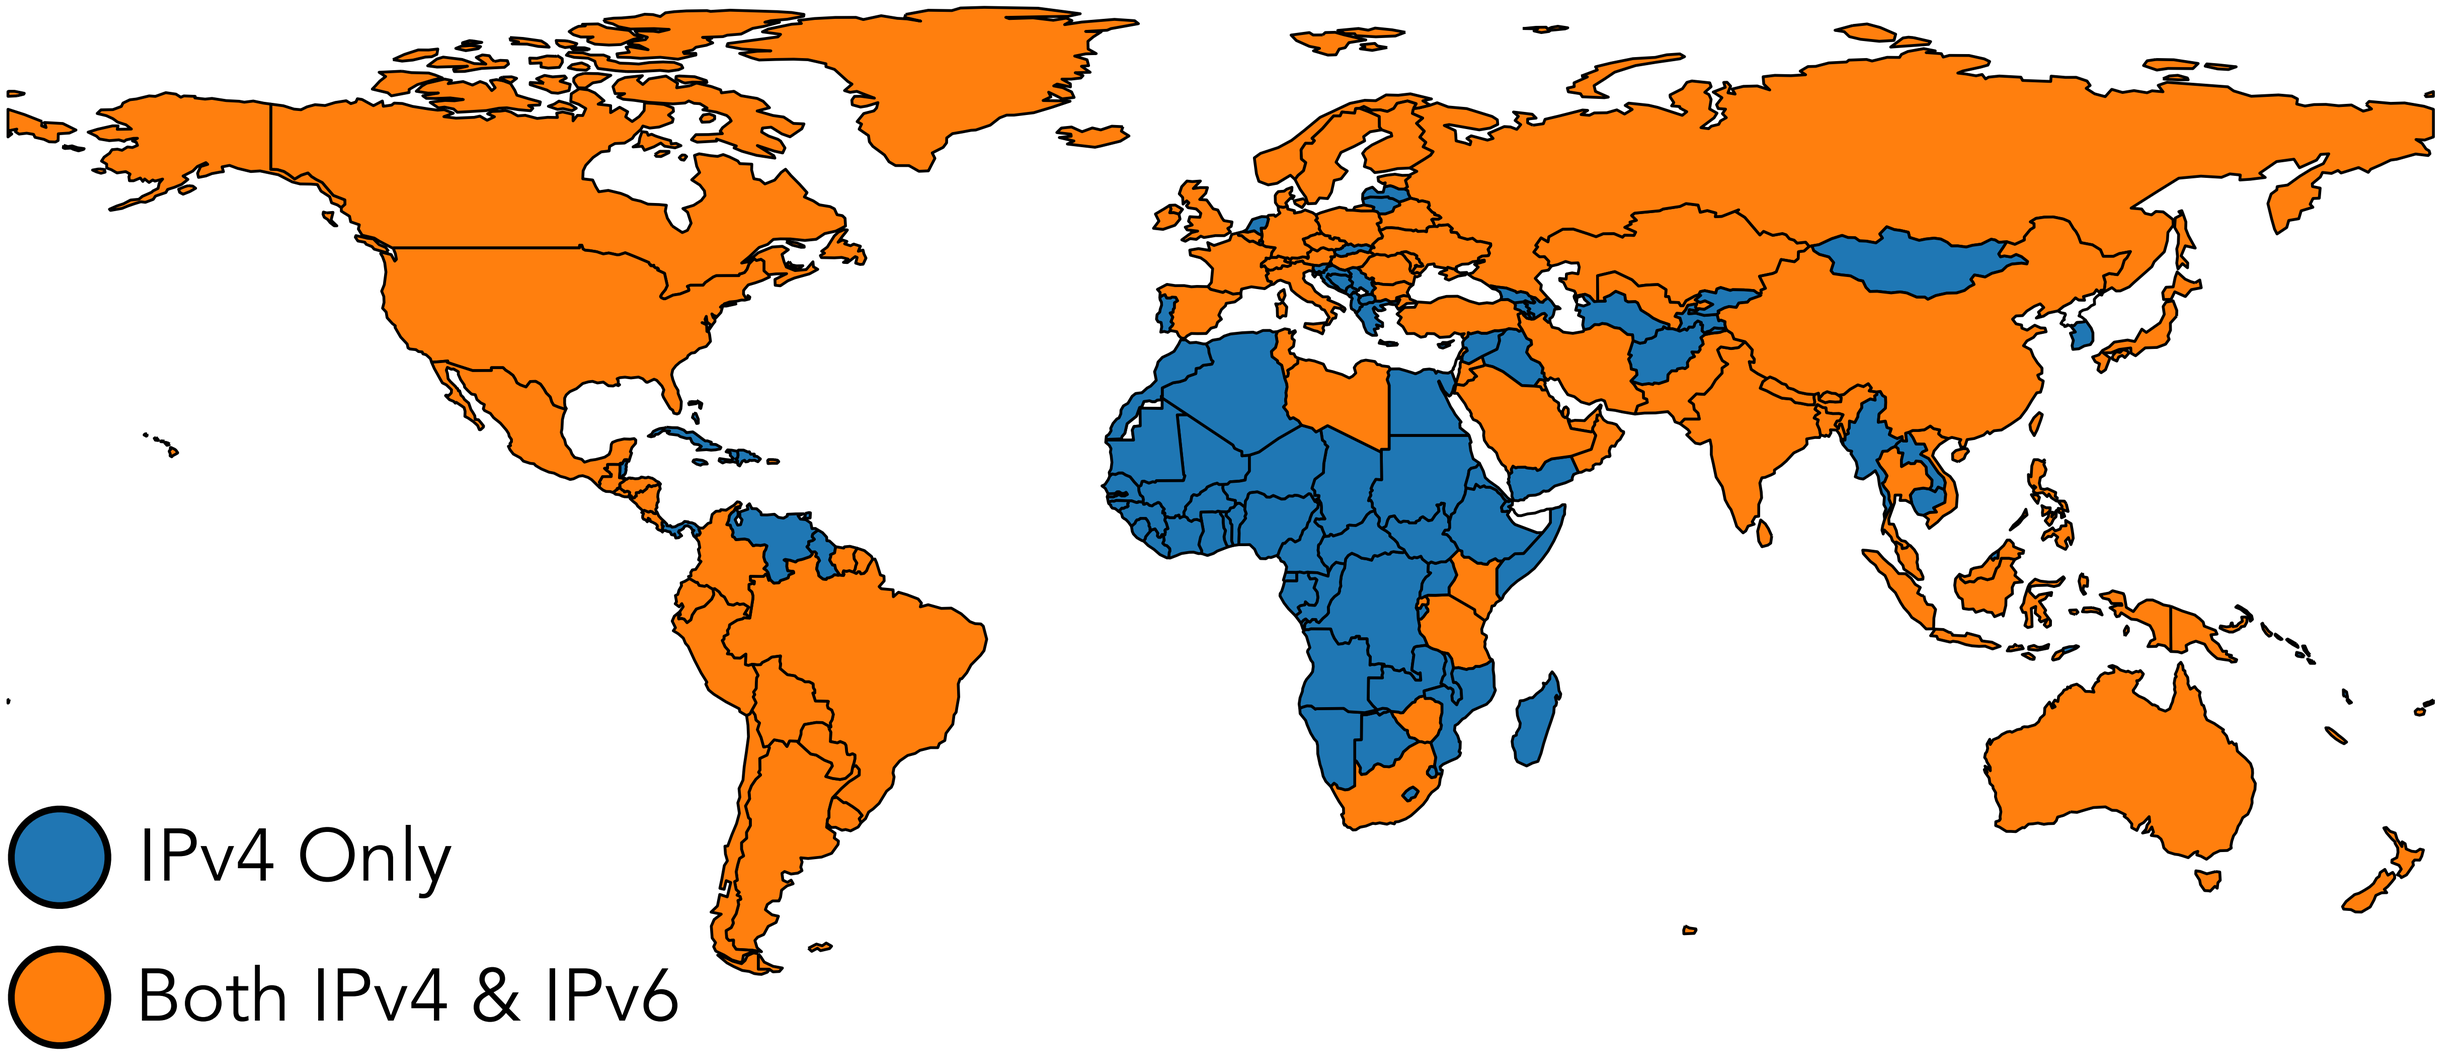

How does Mint work?

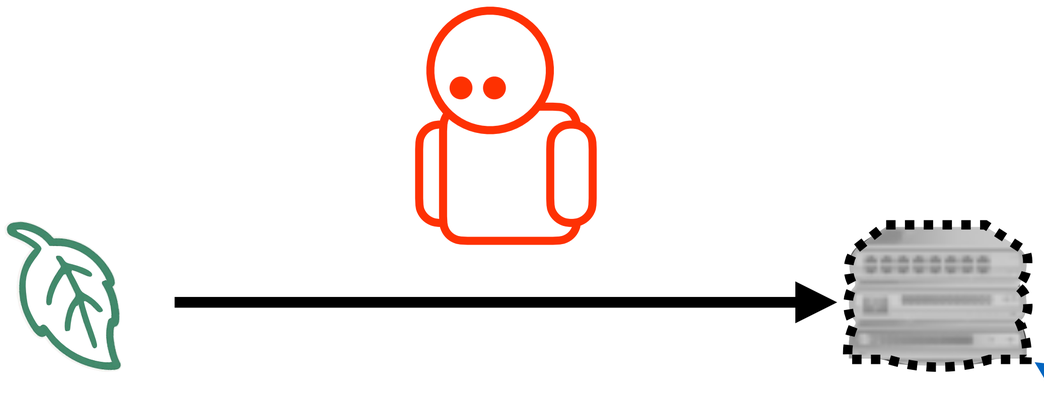

Measures **interference** by probing non-responsive targets

How does Mint work?

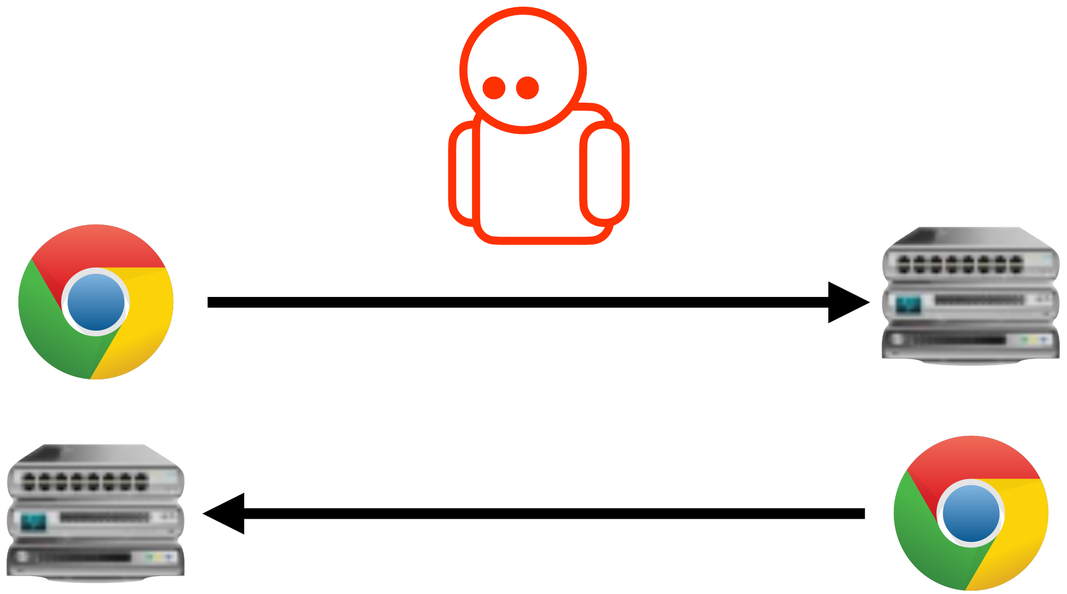

1. Bidirectional interference
   - Agnostic to client/server

How does Mint work?

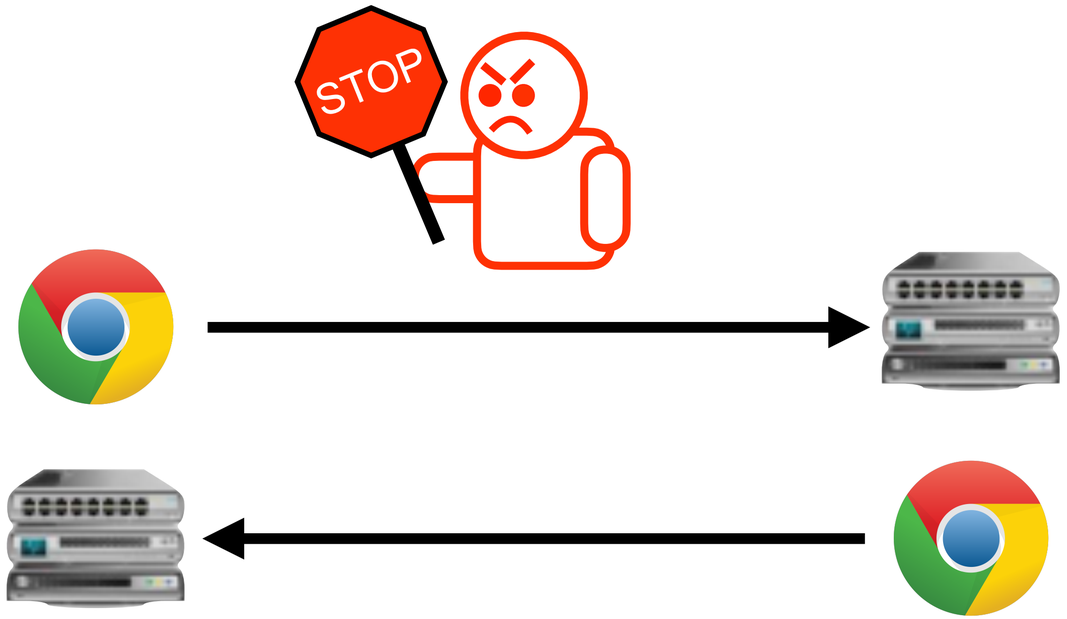

1 Bidirectional interference

IPv4	HTTP	HTTPS
IPv4 addrs	16%	5.4%
/24s	5.8%	3.7%
Countries	100%	100%

Every country has some bidirectional tampering

How does Mint work?

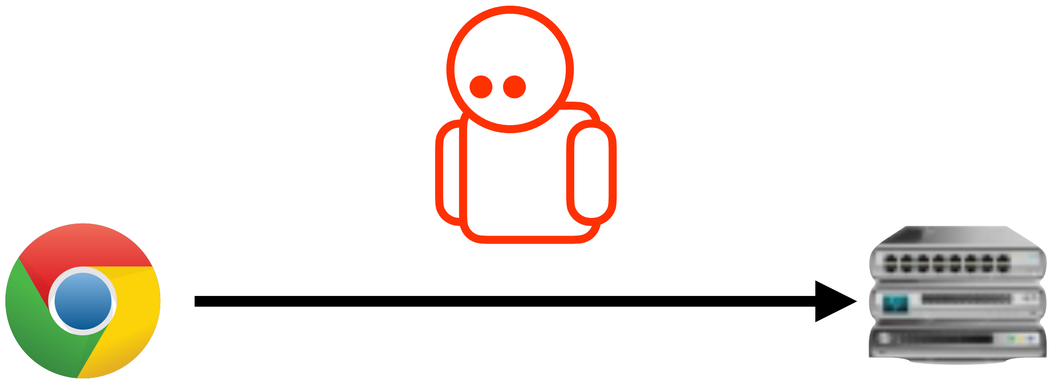

TCP 3-way handshake

Sensitive request

1. Bidirectional interference
2. TCP-noncompliant middleboxes

How does Mint work?

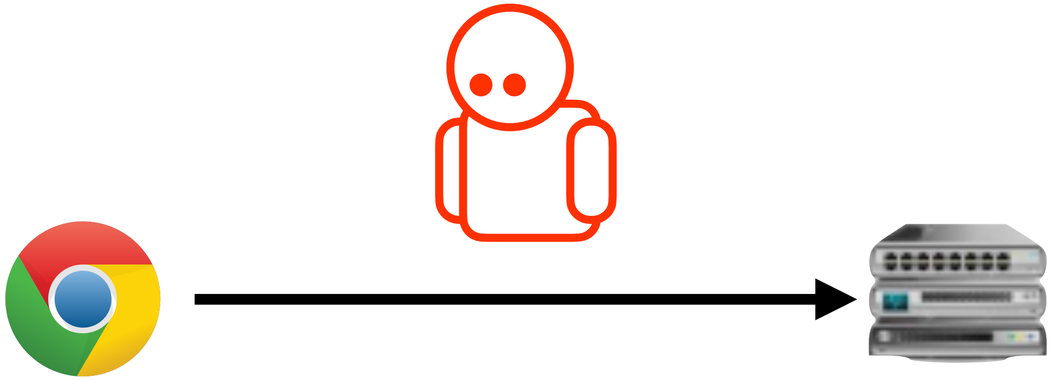

~~TCP 3-way handshake~~

Sensitive request

1. Bidirectional interference
2. TCP-noncompliant middleboxes

   Requires **no packets** from hosts within countries of study

How does Mint work?

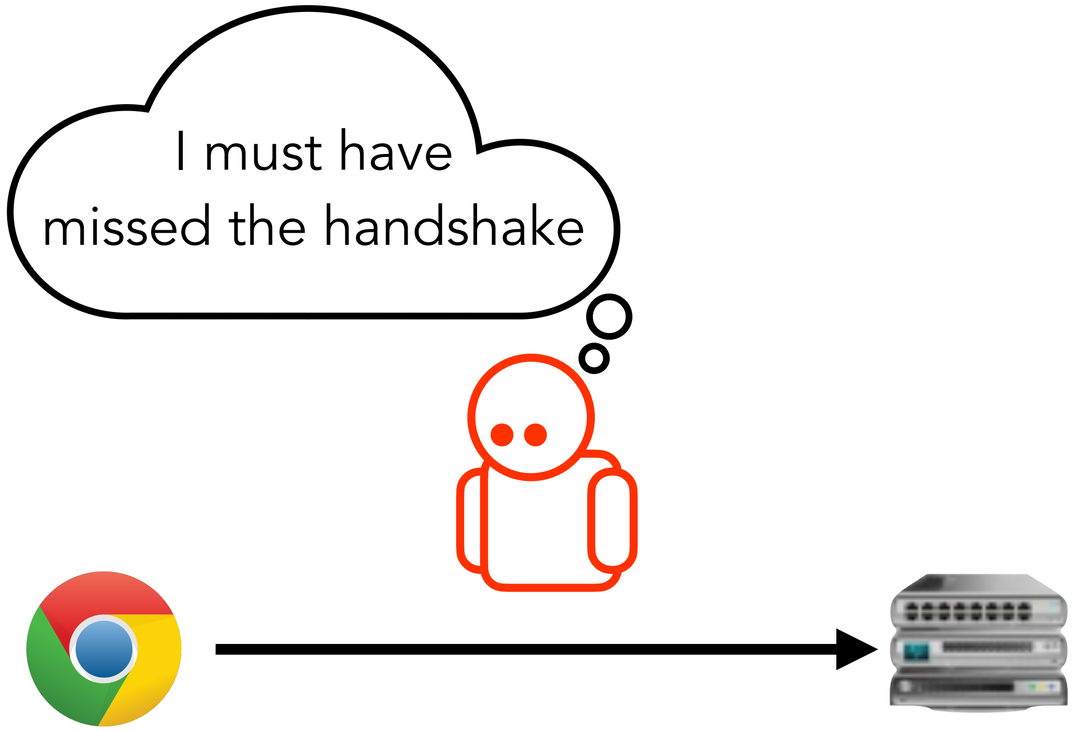

1. Bidirectional interference
2. TCP-noncompliant middleboxes

   Requires **no packets** from hosts within countries of study

How does Mint work?

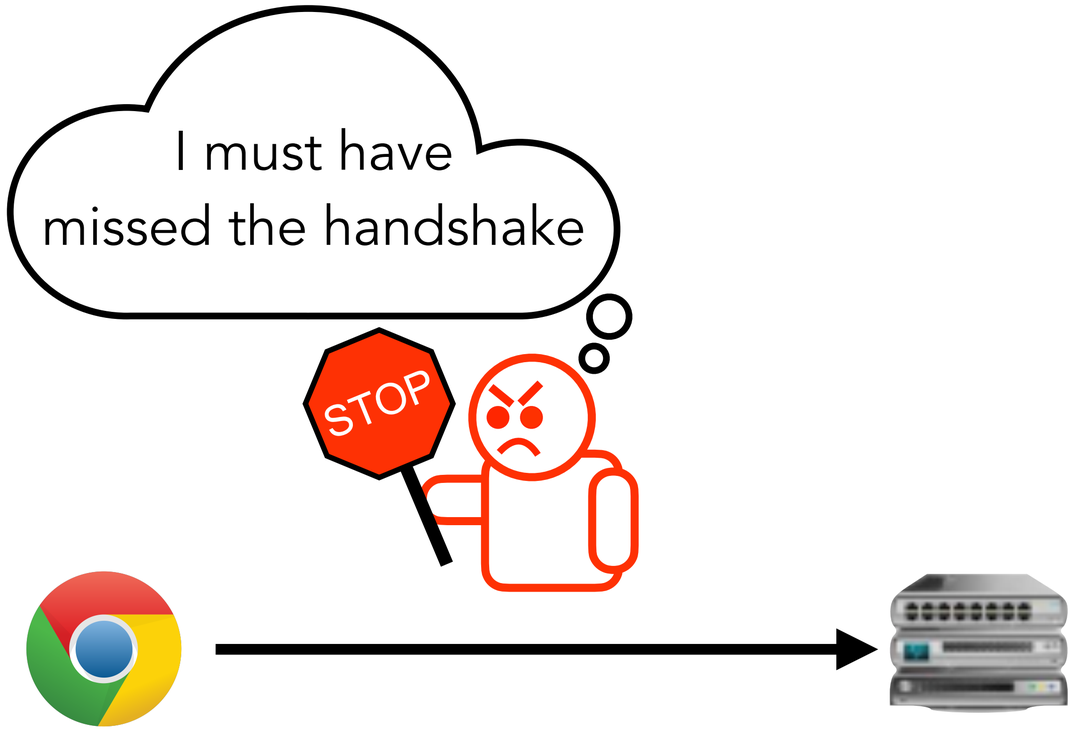

~~TCP 3-way handshake~~

Sensitive request

1. Bidirectional interference
2. TCP-noncompliant middleboxes

   Requires **no packets** from hosts within countries of study

How does Mint work?

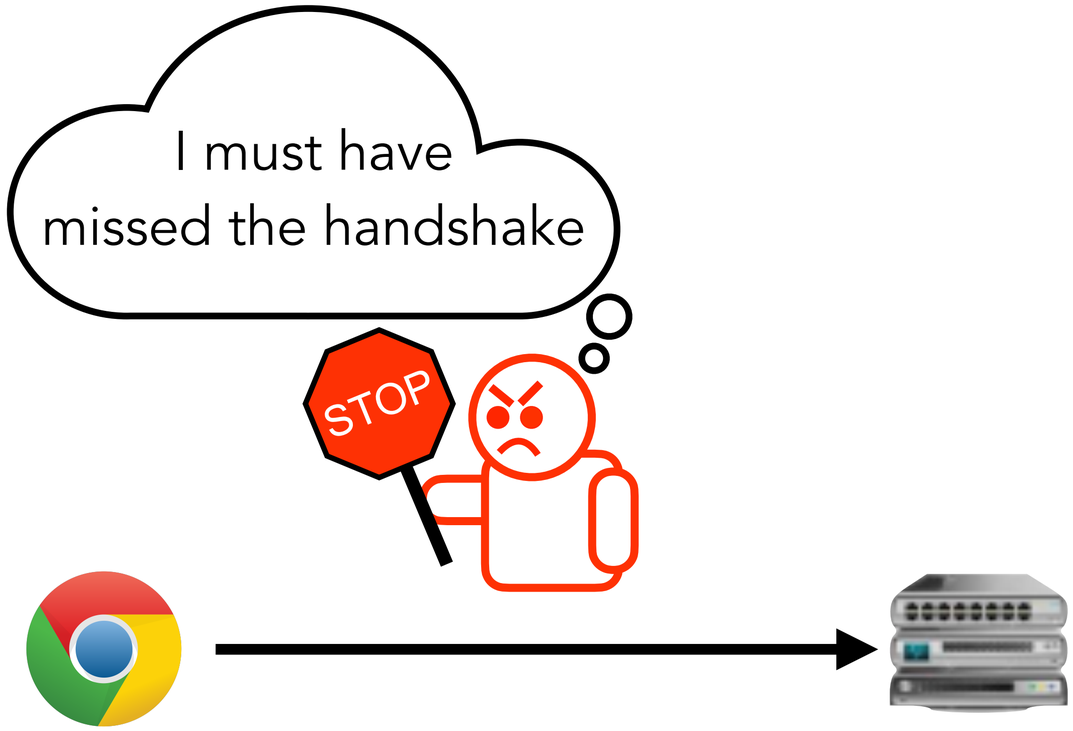

~~TCP 3-way handshake~~

Sensitive request

1. Bidirectional interference
2. TCP-noncompliant middleboxes
   - Requires **no packets** from hosts within countries of study

	HTTP	HTTPS
IPv4 /24s	8.6%	8.4%
IPv6 /48s	**33%**	**33%**

How does Mint work?

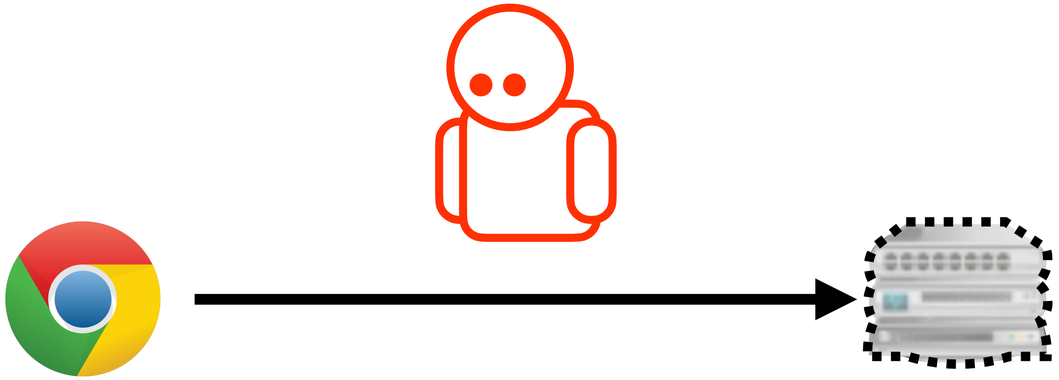

	HTTP	HTTPS
IPv4	91%	90%
IPv6	100%	100%

1. Bidirectional interference
2. TCP-noncompliant middleboxes
3. Non-responsive IP addresses
   - Reduces ethical concerns
   - Applicable in more networks
   - Can test millions of domains

Mint's Scanning Methodology

Find non-responsive hosts in every IPv4/IPv6 network

Probe them for many domains, with many packet sequences

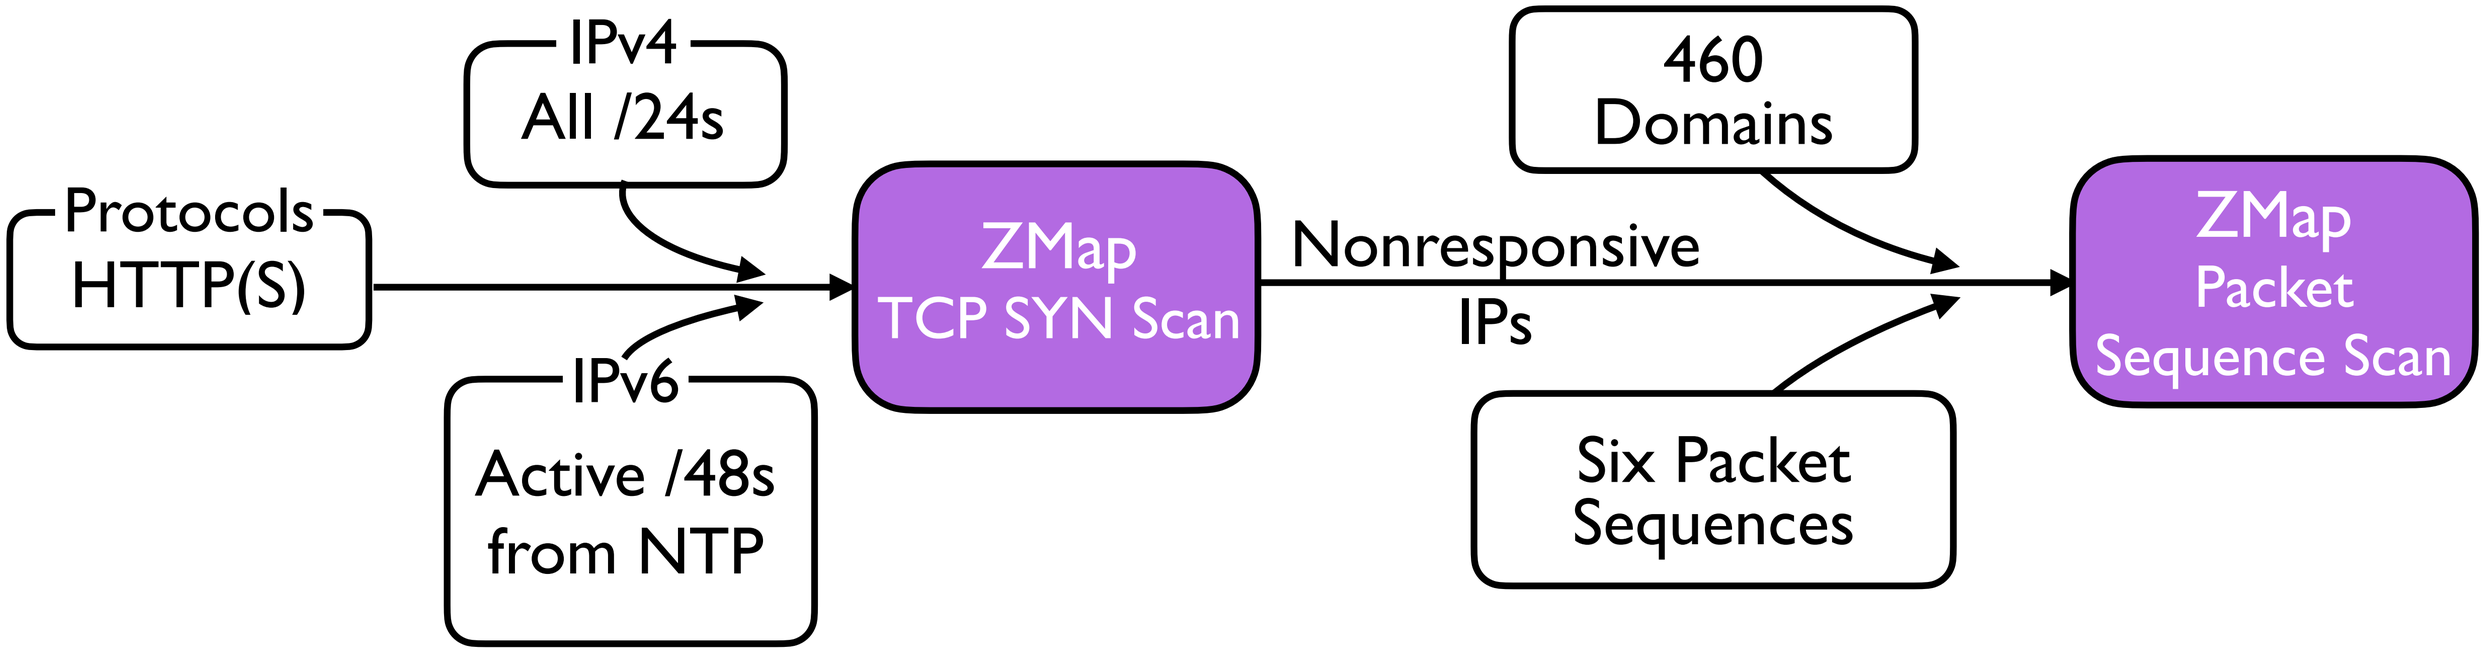

Mint's Scanning Methodology

Find non-responsive hosts in every IPv4/IPv6 network

Probe them for many domains, with many packet sequences

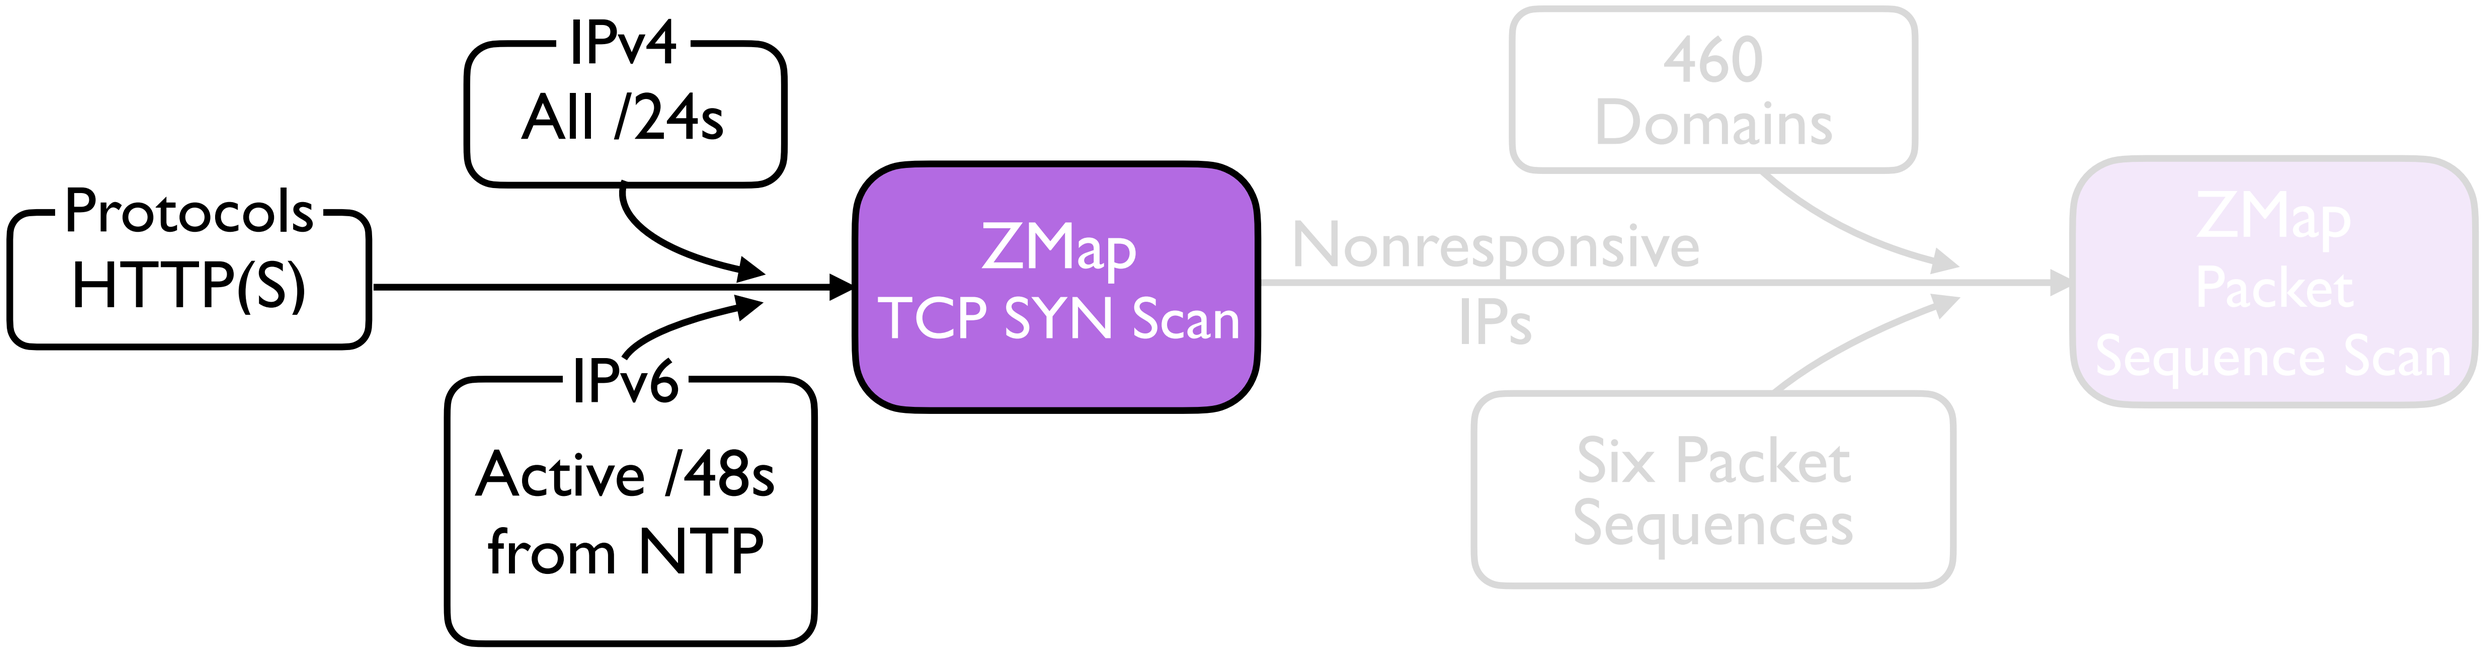

Mint's Scanning Methodology

Find non-responsive hosts in every IPv4/IPv6 network

Probe them for many domains, with many packet sequences

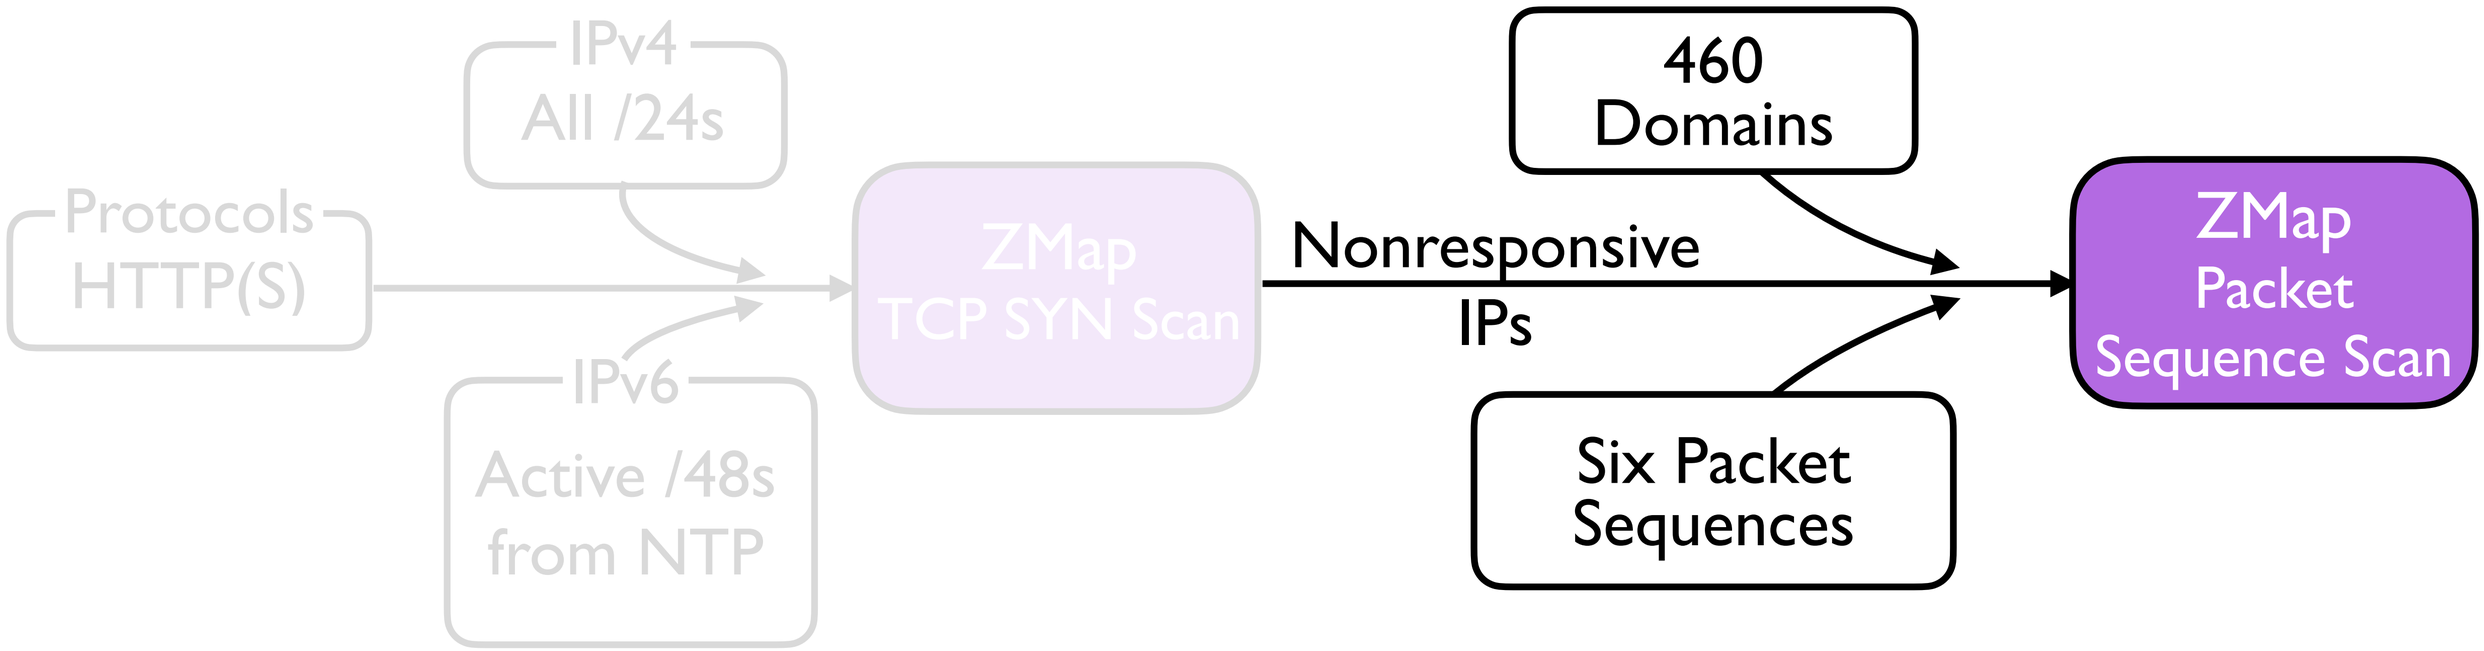

Mint's Scanning Methodology

Find non-responsive hosts in every IPv4/IPv6 network

Probe them for many domains, with many packet sequences

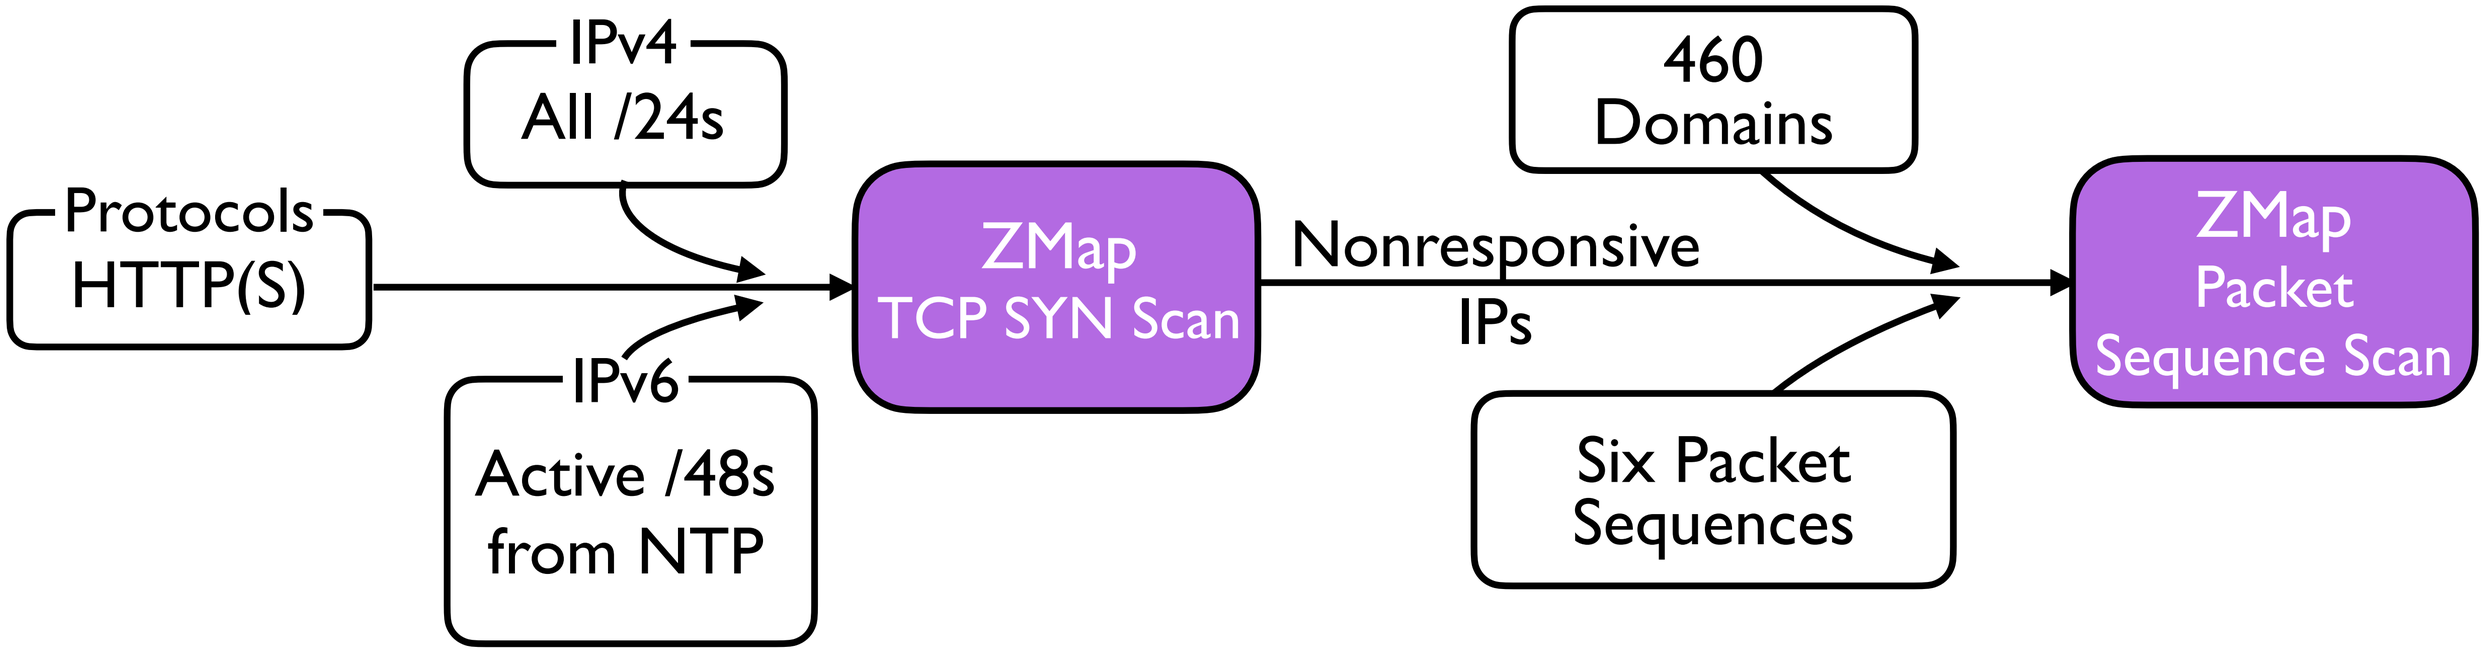

Where does Mint work?

Fraction of /24s Triggered over Total Prefixes Probed

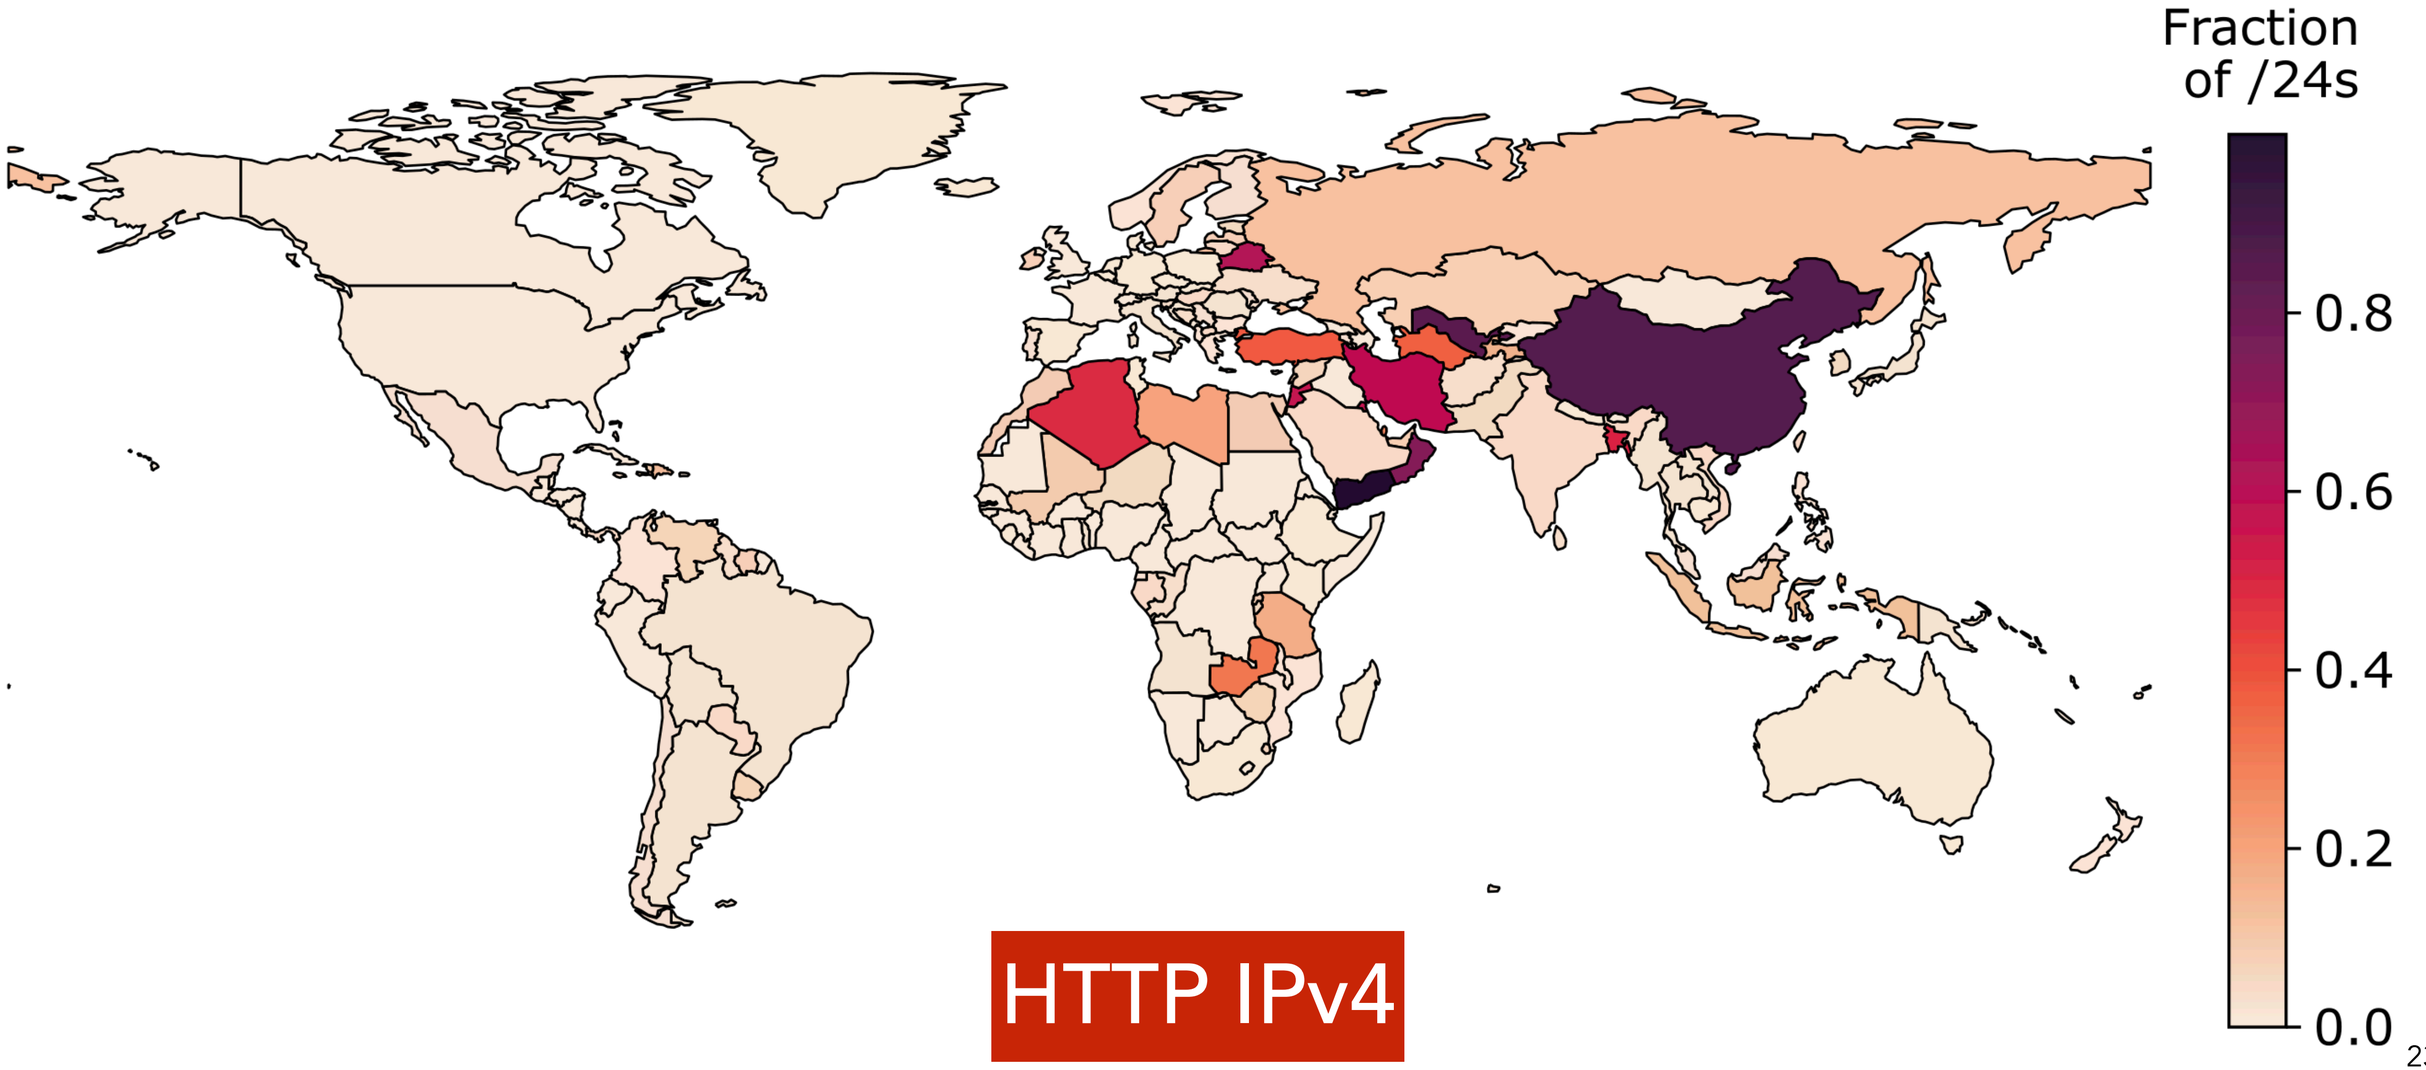

HTTP IPv4

Where does Mint work?

Fraction of /48s Triggered over Total Prefixes Probed

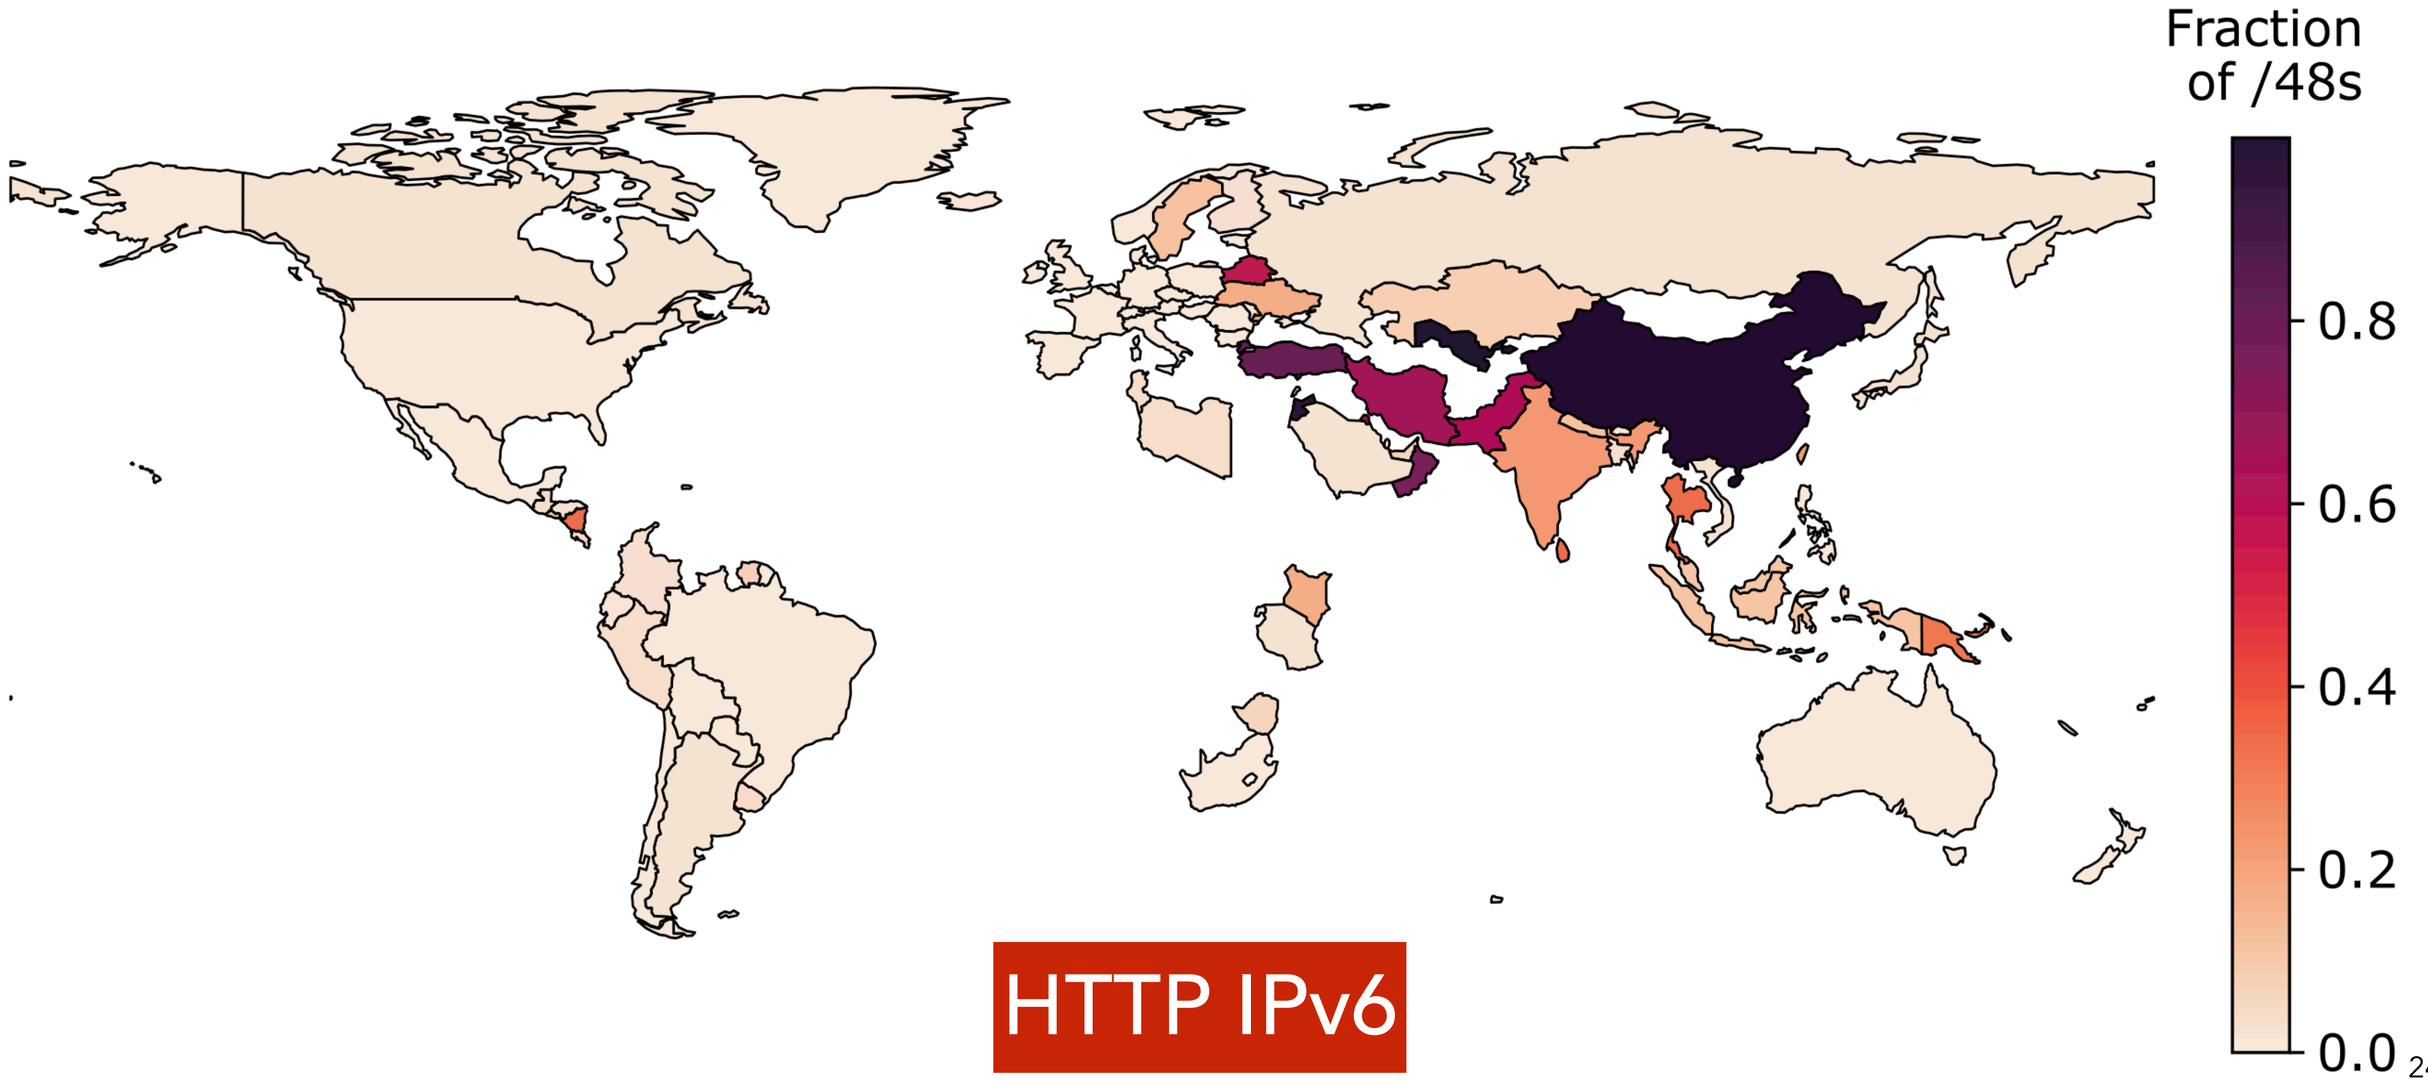

HTTP IPv6

How does Mint compare to previous tools?

Mint triggers **6,348** more ASes than Censored Planet

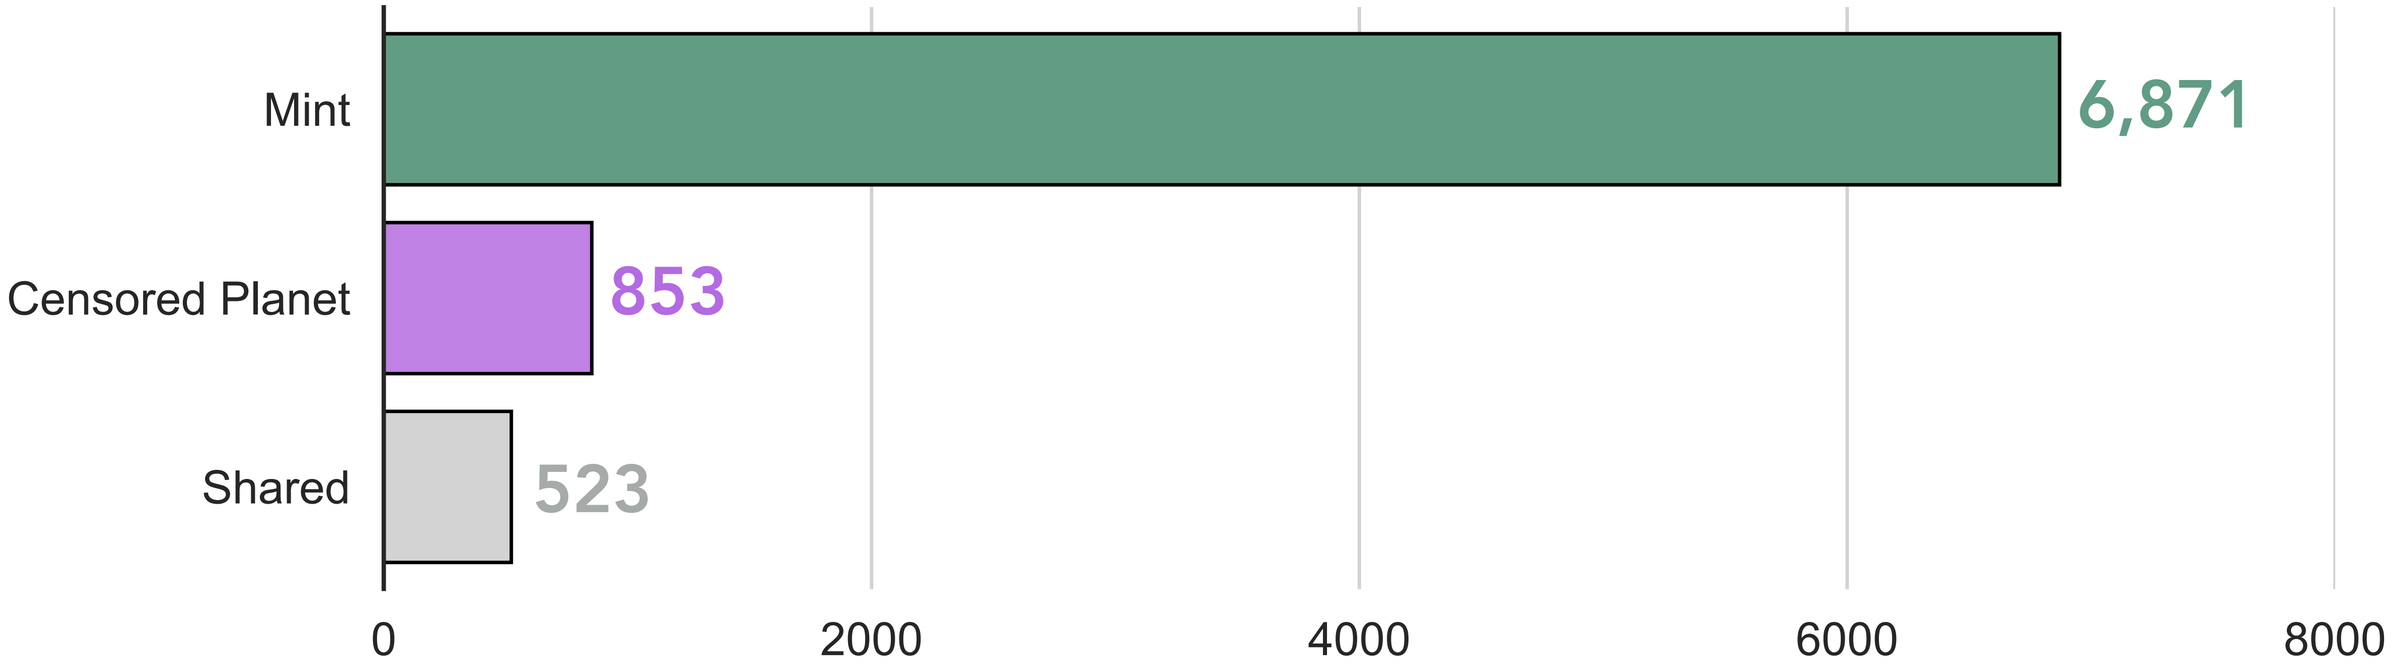

Number of ASes Triggered over **HTTP IPv4**

How does Mint compare to previous tools?

OONI triggers **931** more ASes than Mint

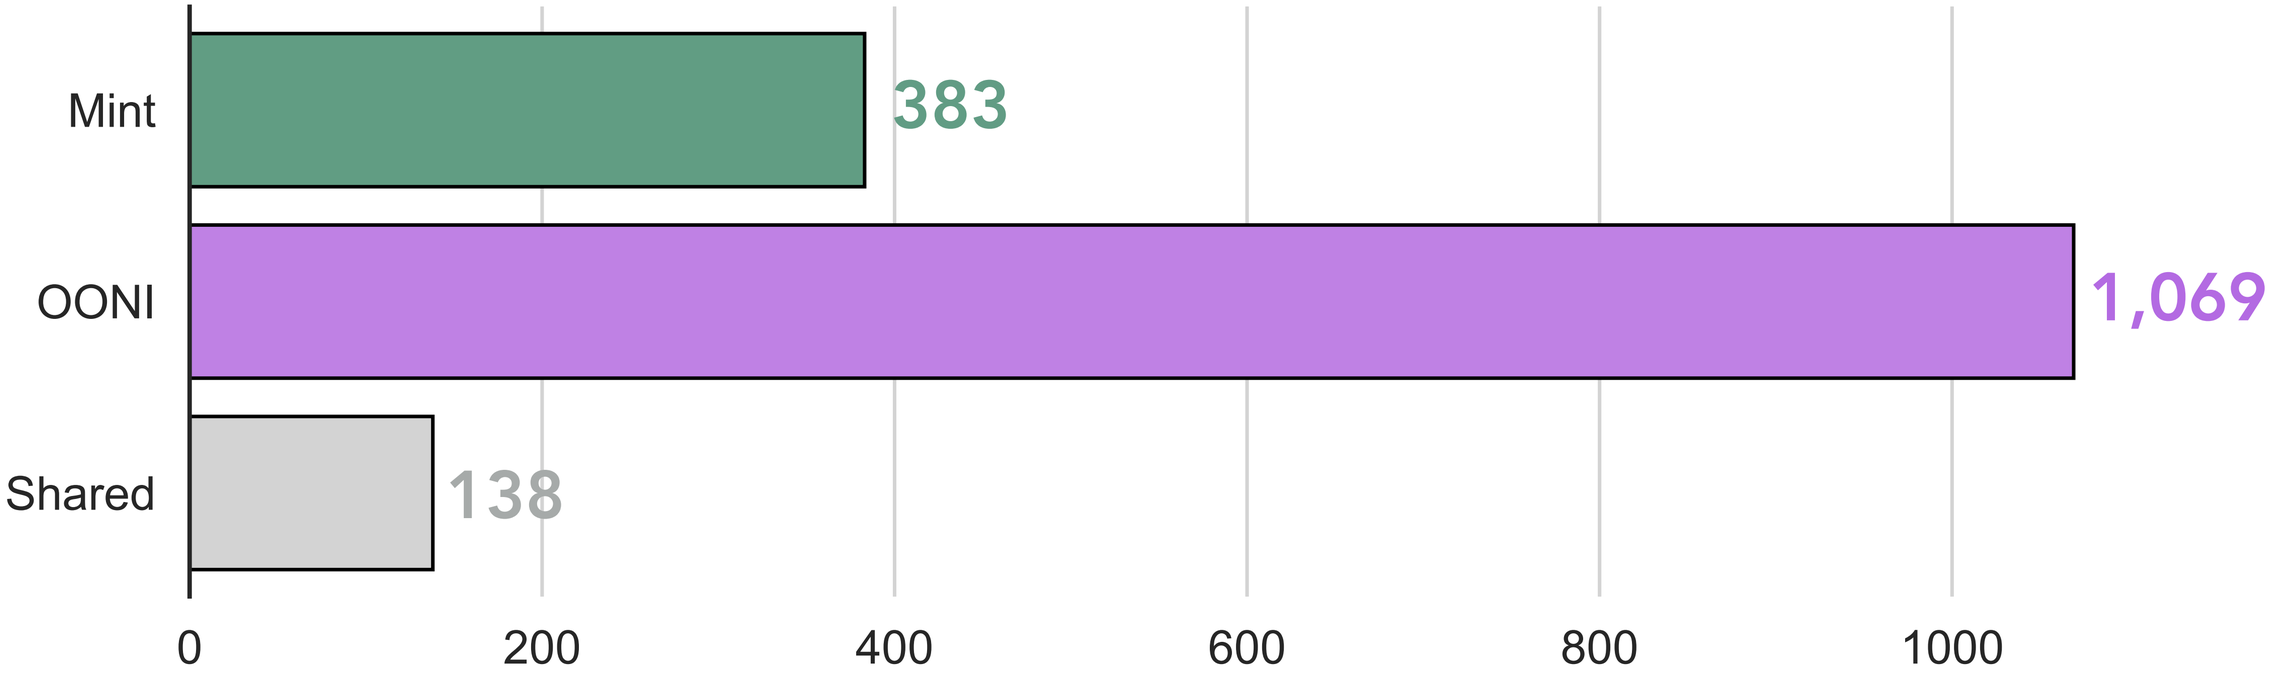

Number of ASes Triggered over **HTTP IPv6**

Broad

Able to measure many networks in each country

Deep

Able to test many domains rapidly

Mint enables studies that were not previously possible

How centralized are interference *mechanisms?*

Centralized

Expect them to interfere with the same **packet sequences**

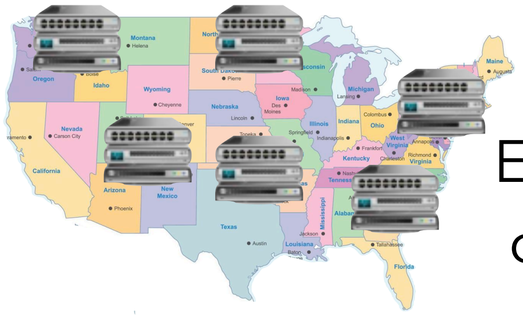

Decentralized

Expect interference with a wide diversity of **packet sequences**

How centralized are interference *mechanisms?*

Centralized

Expect them to interfere with the same **packet sequences**

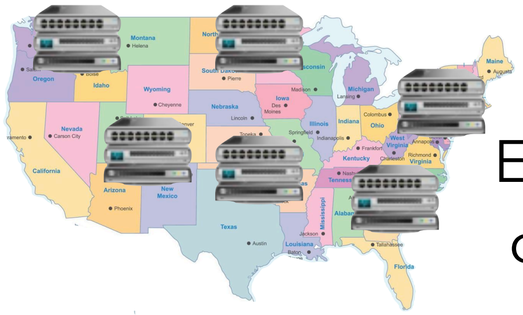

Decentralized

Expect interference with a wide diversity of **packet sequences**

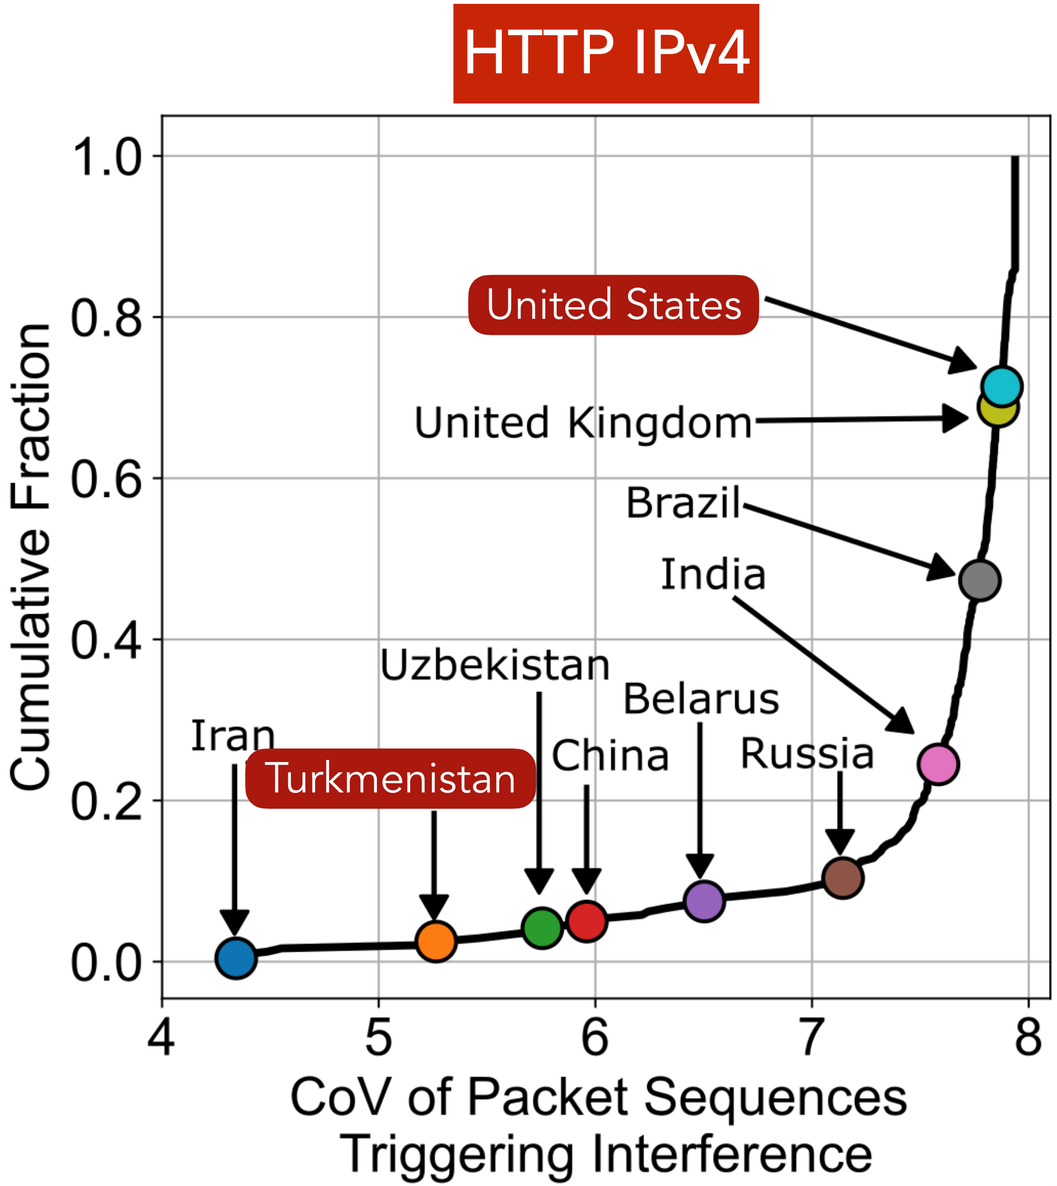

How centralized are interference *policies?*

Centralized

Expect them to interfere with the same **domains**

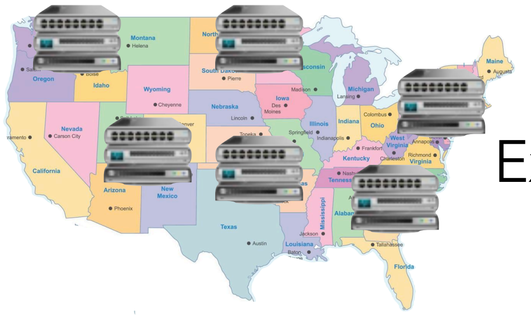

Decentralized

Expect interference with a wide diversity of **domains**

How centralized are interference *policies?*

Centralized

Expect them to interfere with the same **domains**

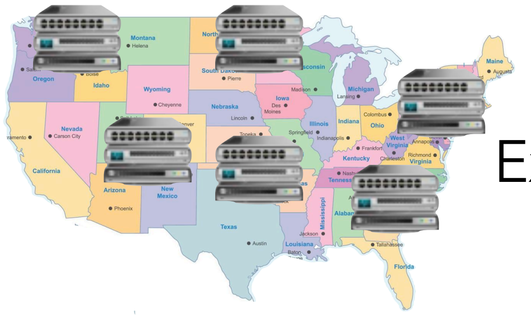

Decentralized

Expect interference with a wide diversity of **domains**

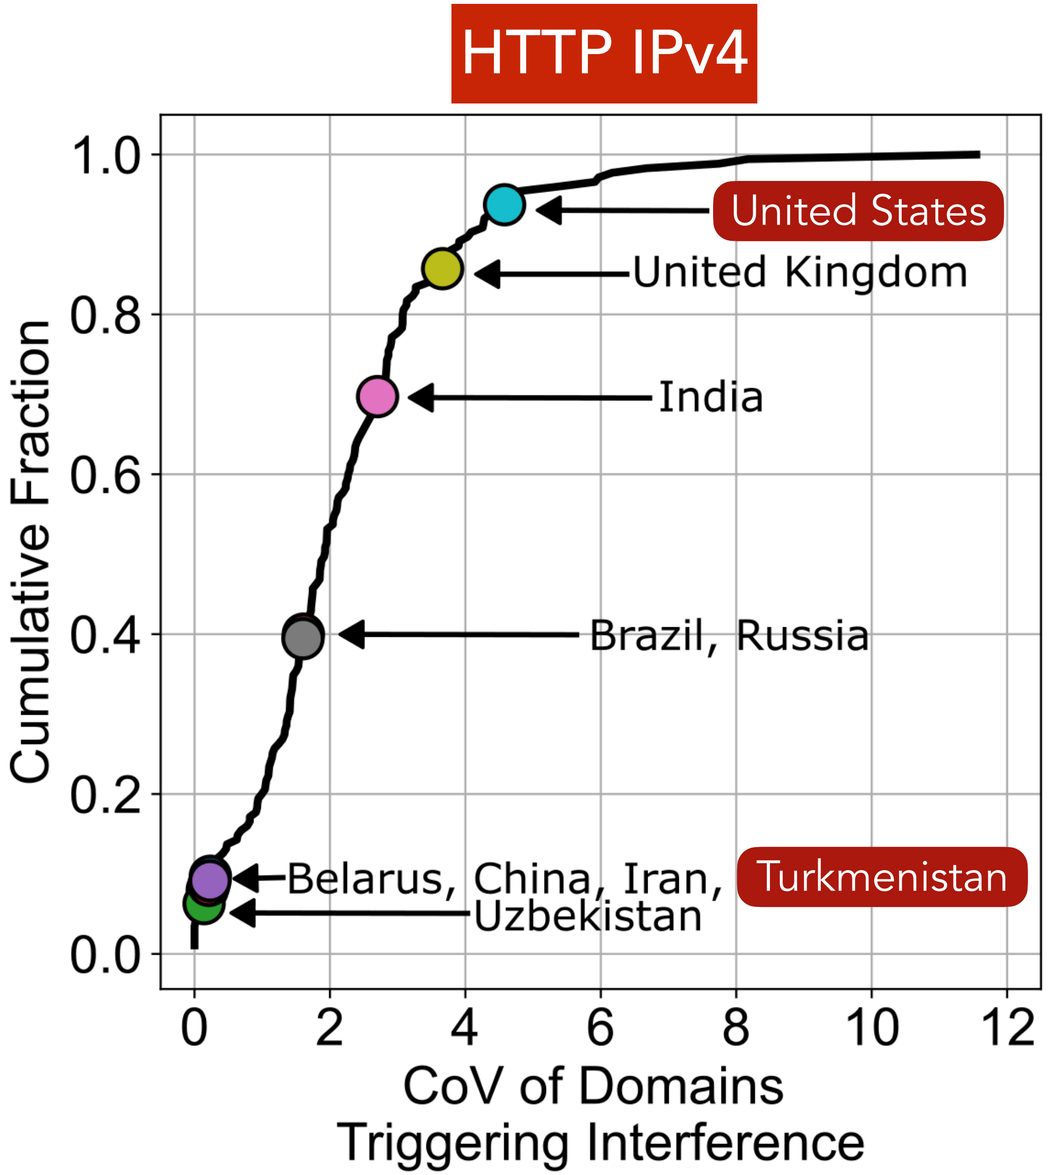

Other results in the paper

Is Nobody There? Good!
Globally Measuring Connection Tampering without Responsive Endhosts

Sadia Nourin*† Erik Rye* Kevin Bock* Nguyen Phong Hoang‡ Dave Levin*

*University of Maryland †Max Planck Institute for Informatics ‡University of British Columbia

***Abstract*—Many techniques have been introduced to measure network interference—tampering performed by nation-state censors or corporate firewalls to block unwanted traffic. However, virtually all prior measurement techniques require some degree of participation from endpoints within each country of study: including VPNs, cloud providers, or volunteers willing to run measurement software on their personal devices at their own risk. However, such endpoints are not always available in all countries that tamper with connections, leaving many networks unmeasurable.**

In this paper, we present the first *global*, active, network interference measurements that require *no participating endpoints within any country of study*. Our techniques extend two recent studies that use packet sequences that trigger network interference from outside the country of study by tricking middleboxes into believing a connection exists. Our system, Mint, generalizes and automates this approach—which had previously only been applied to two countries—to allow it to apply to the global IPv4 and IPv6 Internet. We use Mint to conduct the first global measurements of network interference without using any participating endpoints, and the first comprehensive scans of IPv6 interference. We show that we are able to measure networks, autonomous systems, and even entire countries that previous methods could not. We also present several case studies that highlight how our tool can be used to perform new measurement studies of network interference.

1. Introduction

Network interference occurs when a third-party middlebox (such as a firewall) drops or resets a connection. This is a common technique performed by censoring regimes to restrict access to information, but also in more benign settings like corporate firewalls or schools.

Empirically measuring network interference is critical for informing policy makers and censorship circumvention efforts as to what is being interfered with, who is performing interference, and how they do it. Moreover, because network interference can vary drastically from one network to another [19], [43], [30], it is important that measurement efforts be as *broad* (cover as many networks) and as *deep* (cover as many domains) as possible.

To this end, a diverse set of measurement platforms and techniques have been introduced, but they share common limitations. In particular, nearly all existing measurement techniques require participating endpoints within each country of study, such as: (1) recruiting willing volunteers to run measurement software on their personal devices at their own risk [32], [14], (2) renting vantage points from commercial cloud or VPN providers [28], [19], [20], or (3) making use of unwitting live servers within each country of study [48], [35], [53], [40], [39].

When participating endpoints are available, these techniques are successful, but they face two key limitations:

Participating endpoints are often unavailable. Many autonomous systems (ASes)—especially cellular networks—lack live servers, cloud providers, or users willing to take on the ethical risks of volunteering their personal devices to be used in censorship measurement research. This problem is exacerbated in IPv6, in which it can be extremely difficult to identify live servers [56], [47]; as a result, there are no comprehensive studies of censorship over IPv6 to date.

Even when available, participating endpoints offer limited resources. To avoid saturating endpoints, researchers must limit the number of measurements they perform when using live servers or volunteers.[1] For instance, OONI [14] and Censored Planet [39] commonly test only a small number of URLs on each endpoint per measurement, e.g., less than 100 for OONI at the time of writing.

These limitations show that relying on participating endpoints within a country of study constrains the breadth and depth at which network interference can be measured.

Two recent studies introduced an alternative approach that does not require participation from within a country of study. Instead, they use "packet sequences"—subsets of a standard TCP connection—to essentially trick interfering middleboxes into believing that a connection has been established. Nourin et al. [30] used one packet sequence to study Turkmenistan's censorship infrastructure, and Hoang et al. [18] used a different one to study China's Great Firewall. These two isolated studies demonstrate that packet-sequence-based measurement can, in at least some cases, allow for measurement to take place without the standard limitations. However, we are aware of no work that has

1. Rented VPNs and cloud providers typically do not have this limitation.

Bidirectional Interference — Countries and networks that exhibit bidirectional network interference

Case Studies — Case studies on ASes that deploy IPv4 and **IPv6 censorship** differently

IPv6 Residual Censorship — Identifying and circumventing IPv6 **residual censorship in China**

HTTP, HTTPS, IPv4, IPv6 — Results for all combinations of protocols

Is Nobody There? Good! Globally Measuring Connection Tampering without Responsive Endhosts

Mint **globally** measures network interference **without endhosts**

Mint is the first tool **to measure IPv6** network interference **at scale**

Mint lets us explore new questions about network interference, like quantifying **levels of centralization**

Mint

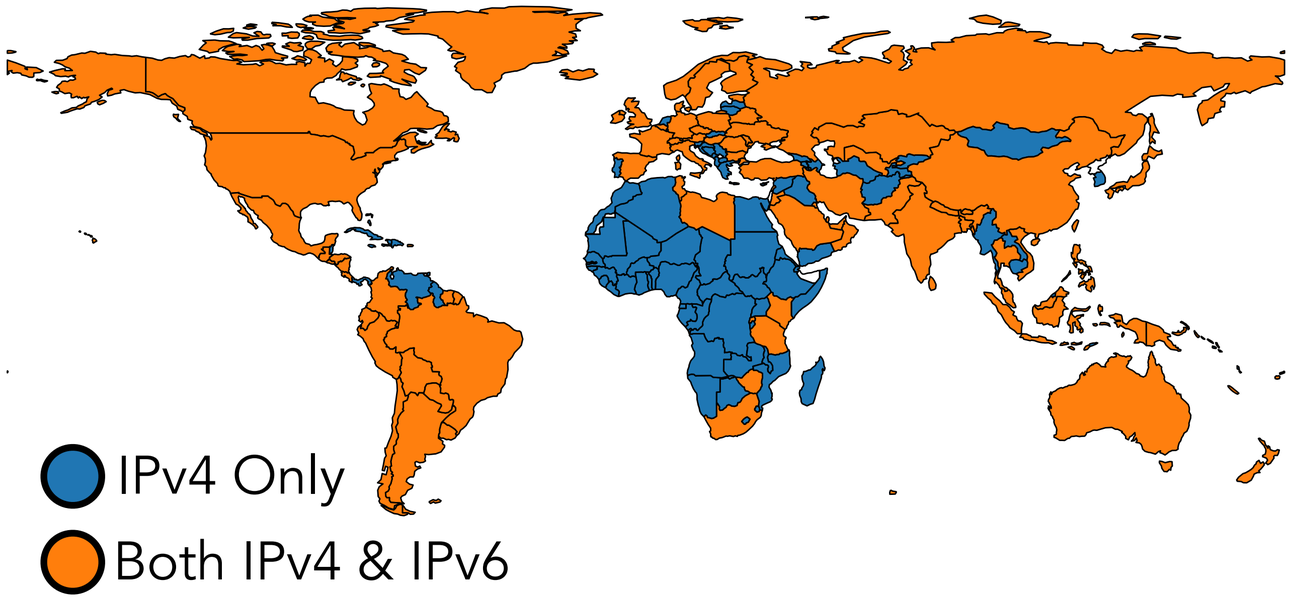

Website censorship.ai