Our Use of RIPE Atlas in Our Work on “The Effect of DNS on Tor’s Anonymity”

Benjamin Greschbach
KTH Royal Institute of Technology

Tobias Pulls
Karlstad University

Laura M. Roberts
Princeton University

Philipp Winter
Princeton University

Nick Feamster
Princeton University

October 5, 2017

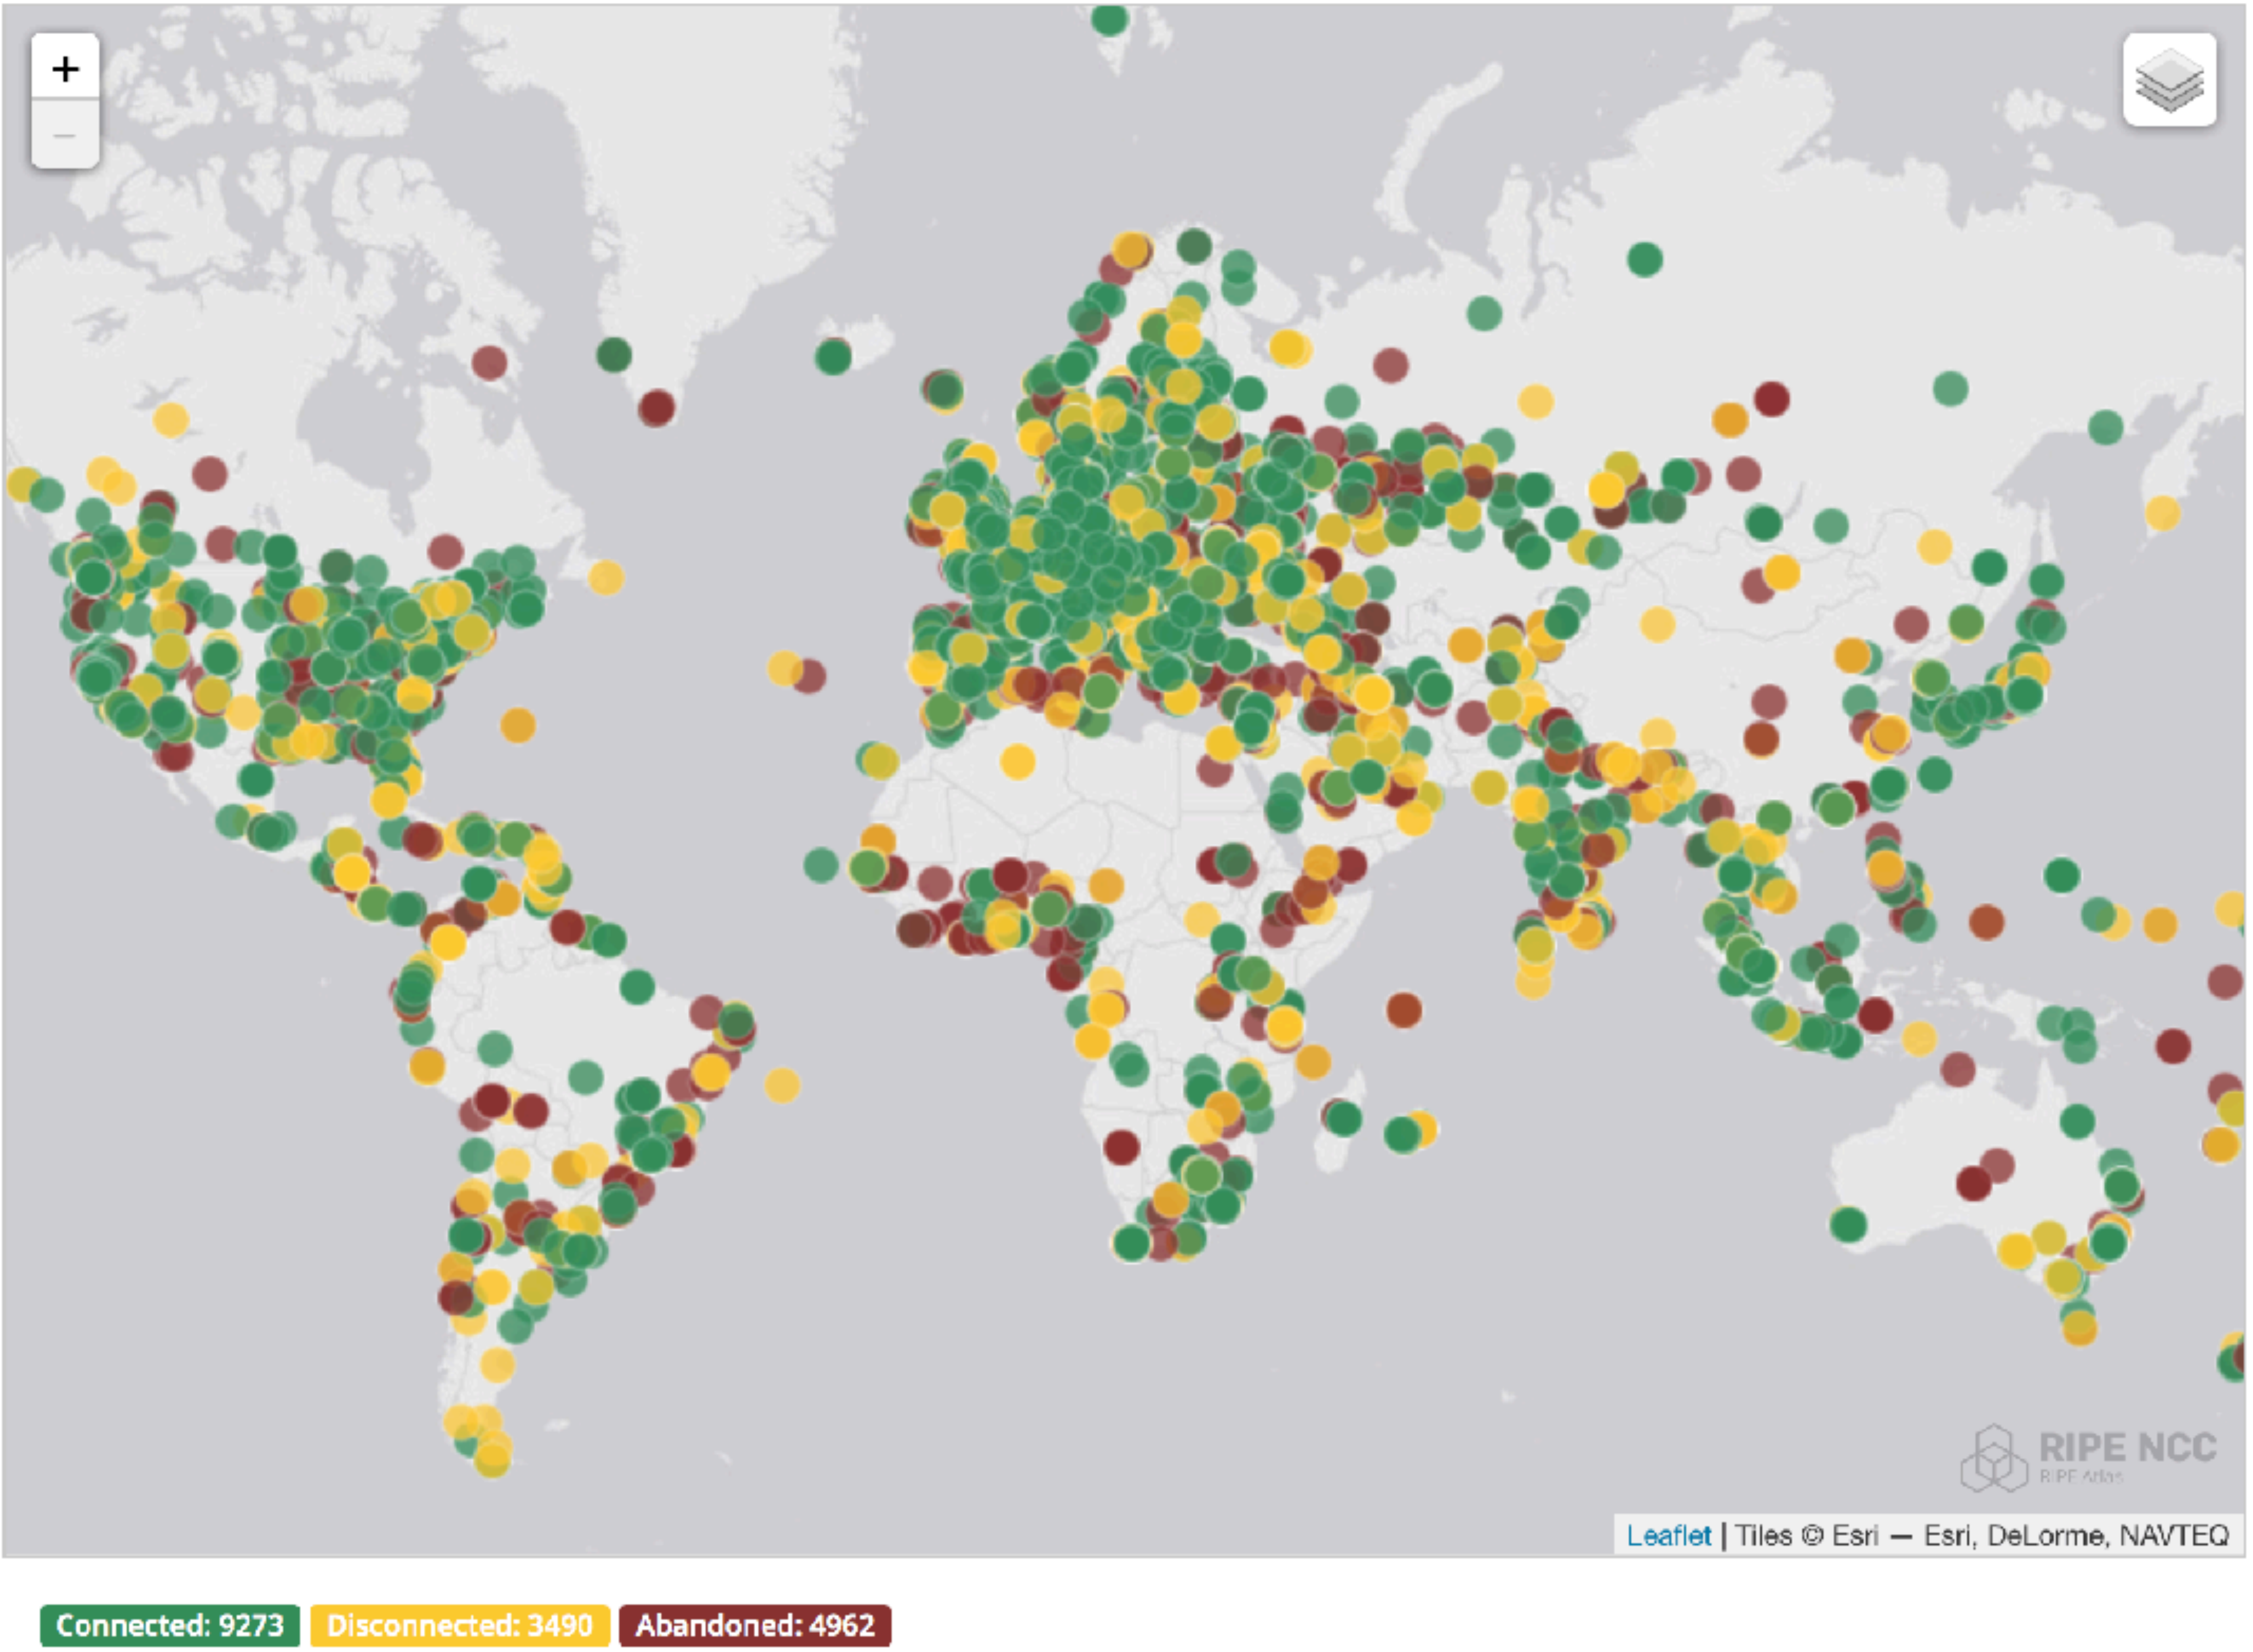

This talk will guide you through the details of how we used RIPE Atlas to perform our Internet-scale simulations.

This talk will guide you through the details of how we used RIPE Atlas to perform our Internet-scale simulations.

Tor®

TorProject.org

Background

This talk will guide you through the details of how we used RIPE Atlas to perform our Internet-scale simulations.

Background

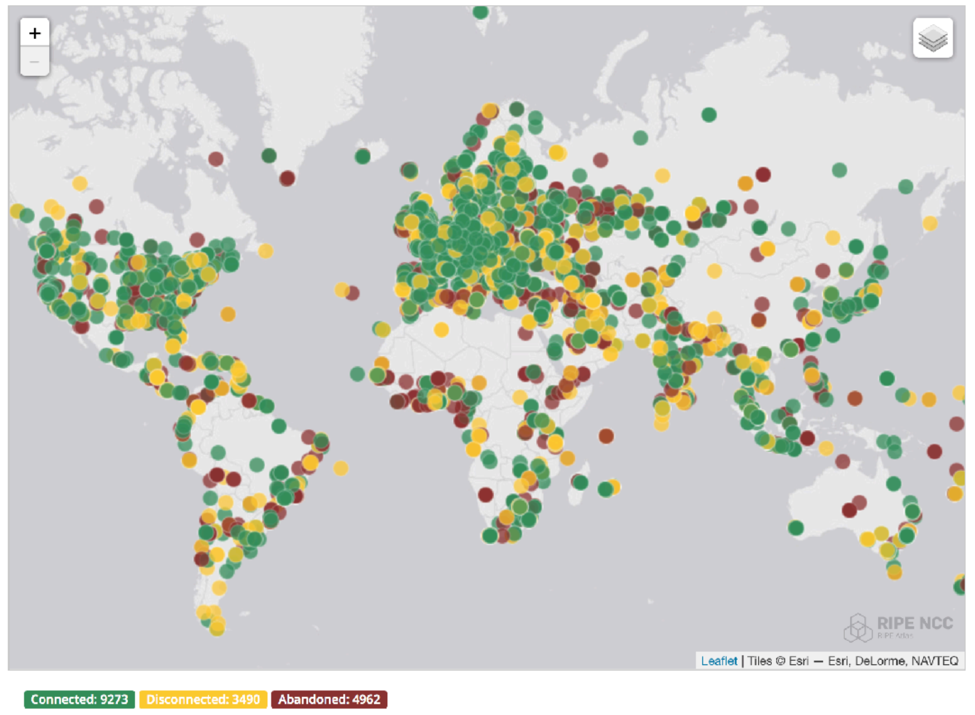

Our use of RIPE Atlas

This talk will guide you through the details of how we used RIPE Atlas to perform our Internet-scale simulations.

TorProject.org

Background

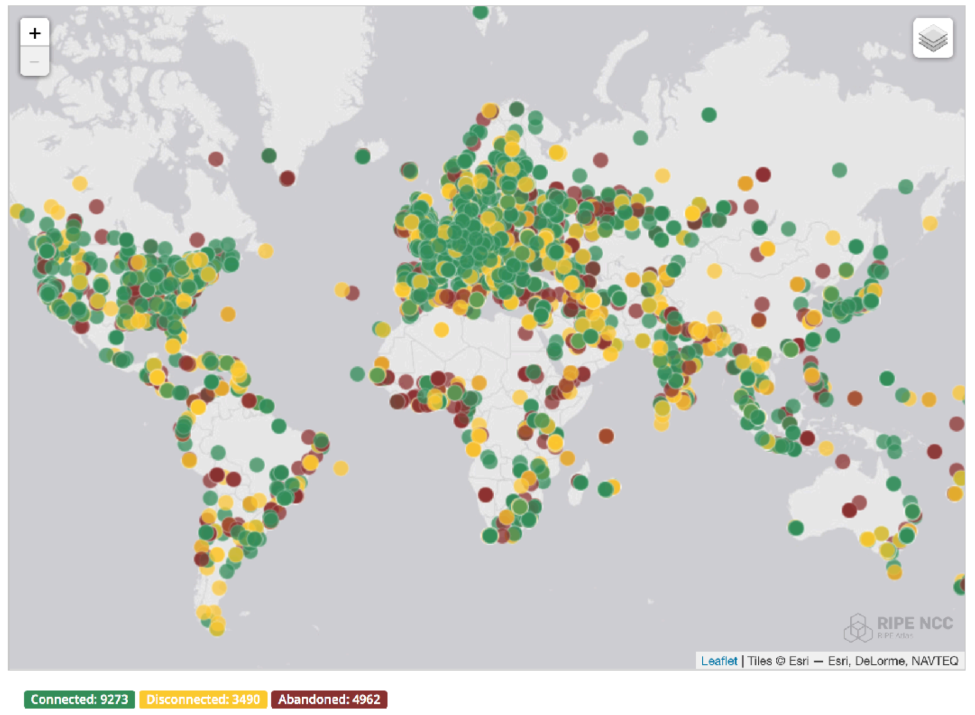

Our use of RIPE Atlas

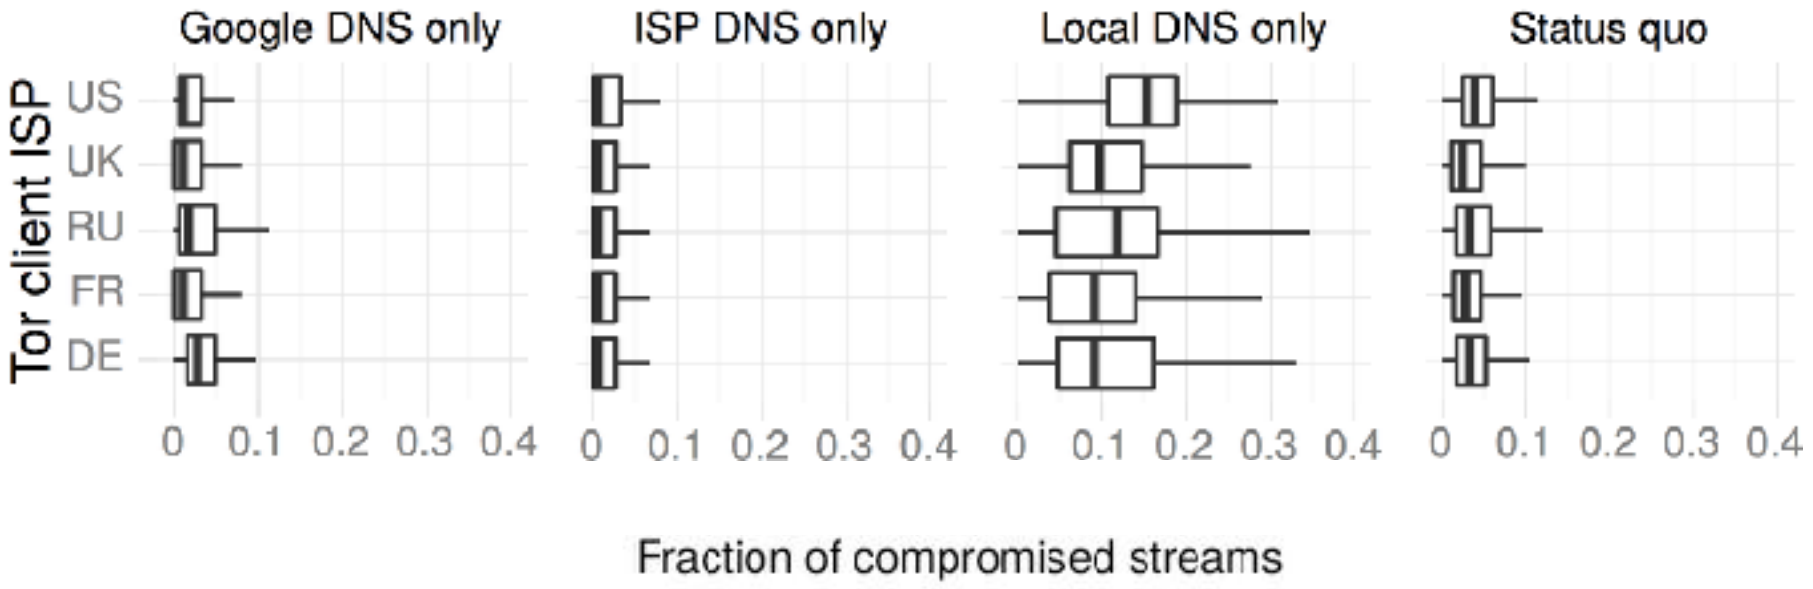

(a) The fraction of compromised streams of simulated Tor clients.

Results

TorProject.org

Background

This work appeared at NDSS 2017.

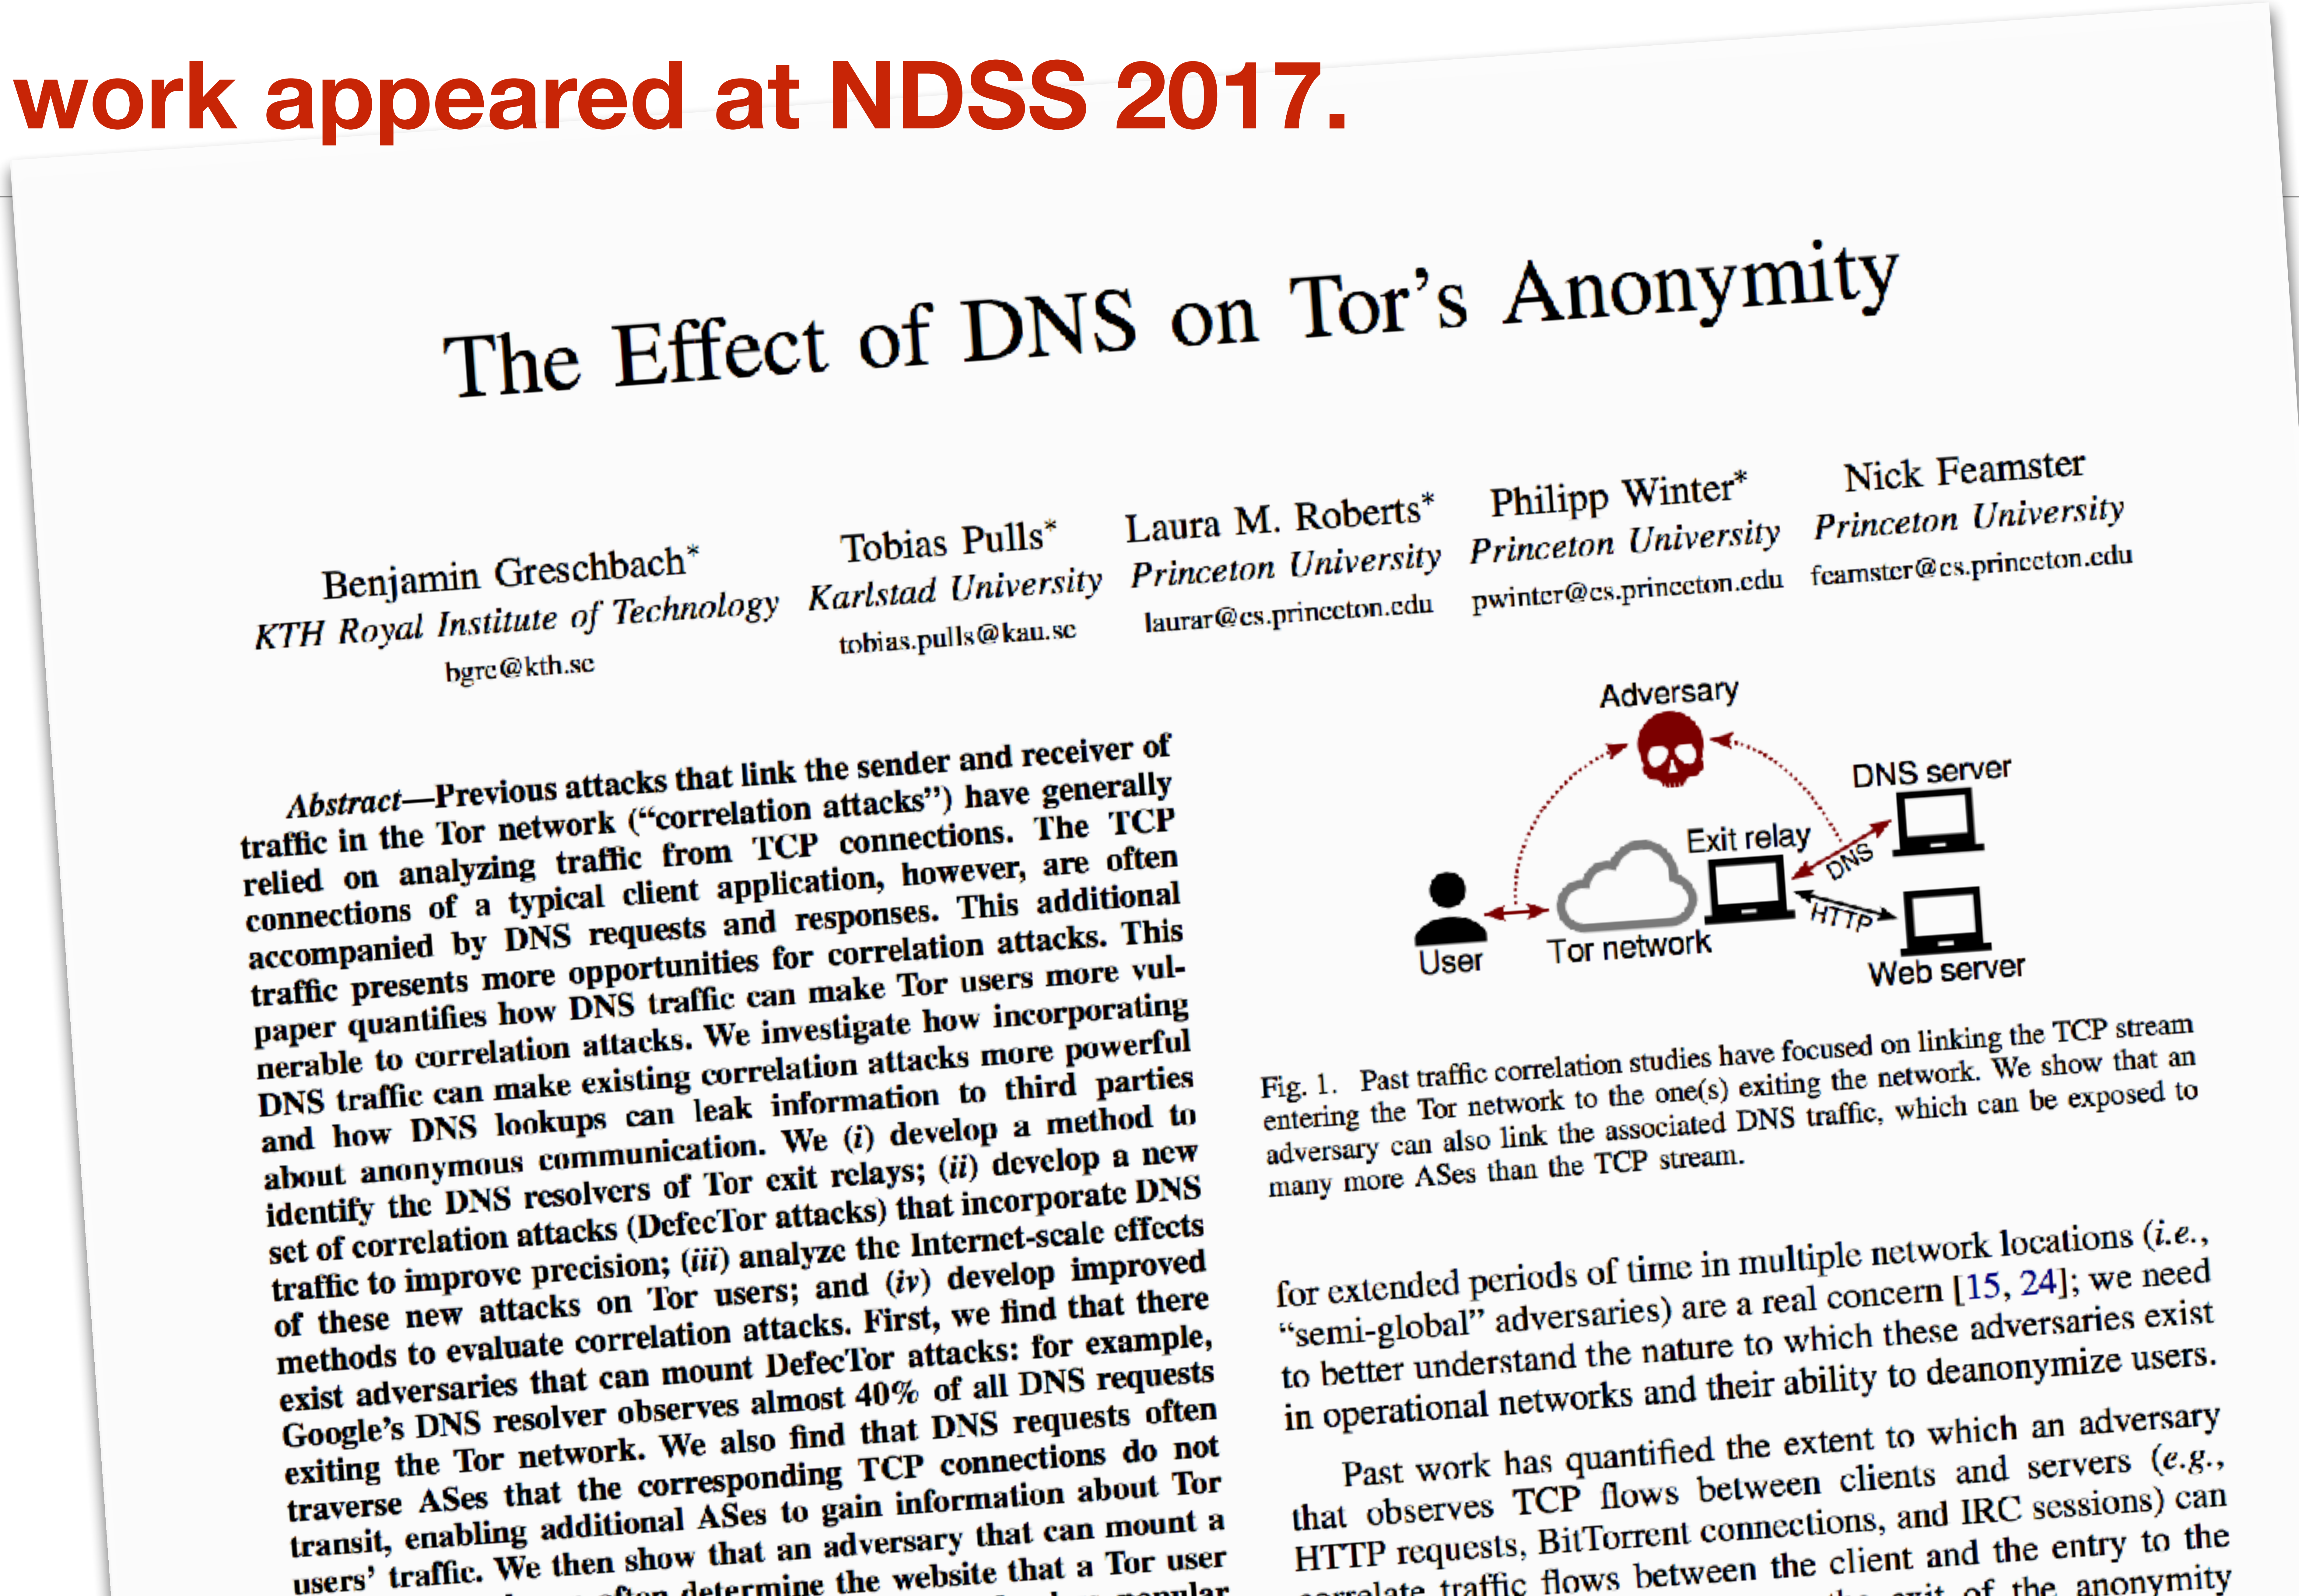

The Effect of DNS on Tor's Anonymity

Benjamin Greschbach*
KTH Royal Institute of Technology
bgre@kth.se

Tobias Pulls*
Karlstad University
tobias.pulls@kau.se

Laura M. Roberts*
Princeton University
laurar@cs.princeton.edu

Philipp Winter*
Princeton University
pwinter@cs.princeton.edu

Nick Feamster
Princeton University
feamster@cs.princeton.edu

***Abstract*—Previous attacks that link the sender and receiver of traffic in the Tor network ("correlation attacks") have generally relied on analyzing traffic from TCP connections. The TCP connections of a typical client application, however, are often accompanied by DNS requests and responses. This additional traffic presents more opportunities for correlation attacks. This paper quantifies how DNS traffic can make Tor users more vulnerable to correlation attacks. We investigate how incorporating DNS traffic can make existing correlation attacks more powerful and how DNS lookups can leak information to third parties about anonymous communication. We (*i*) develop a method to identify the DNS resolvers of Tor exit relays; (*ii*) develop a new set of correlation attacks (DefecTor attacks) that incorporate DNS traffic to improve precision; (*iii*) analyze the Internet-scale effects of these new attacks on Tor users; and (*iv*) develop improved methods to evaluate correlation attacks. First, we find that there exist adversaries that can mount DefecTor attacks: for example, Google's DNS resolver observes almost 40% of all DNS requests exiting the Tor network. We also find that DNS requests often traverse ASes that the corresponding TCP connections do not transit, enabling additional ASes to gain information about Tor users' traffic. We then show that an adversary that can mount a often determine the website that a Tor user**

Fig. 1. Past traffic correlation studies have focused on linking the TCP stream entering the Tor network to the one(s) exiting the network. We show that an adversary can also link the associated DNS traffic, which can be exposed to many more ASes than the TCP stream.

for extended periods of time in multiple network locations (*i.e.*, "semi-global" adversaries) are a real concern [15, 24]; we need to better understand the nature to which these adversaries exist in operational networks and their ability to deanonymize users.

Past work has quantified the extent to which an adversary that observes TCP flows between clients and servers (*e.g.*, HTTP requests, BitTorrent connections, and IRC sessions) can correlate traffic flows between the client and the entry to the

I spoke about our work at RIPE 74.

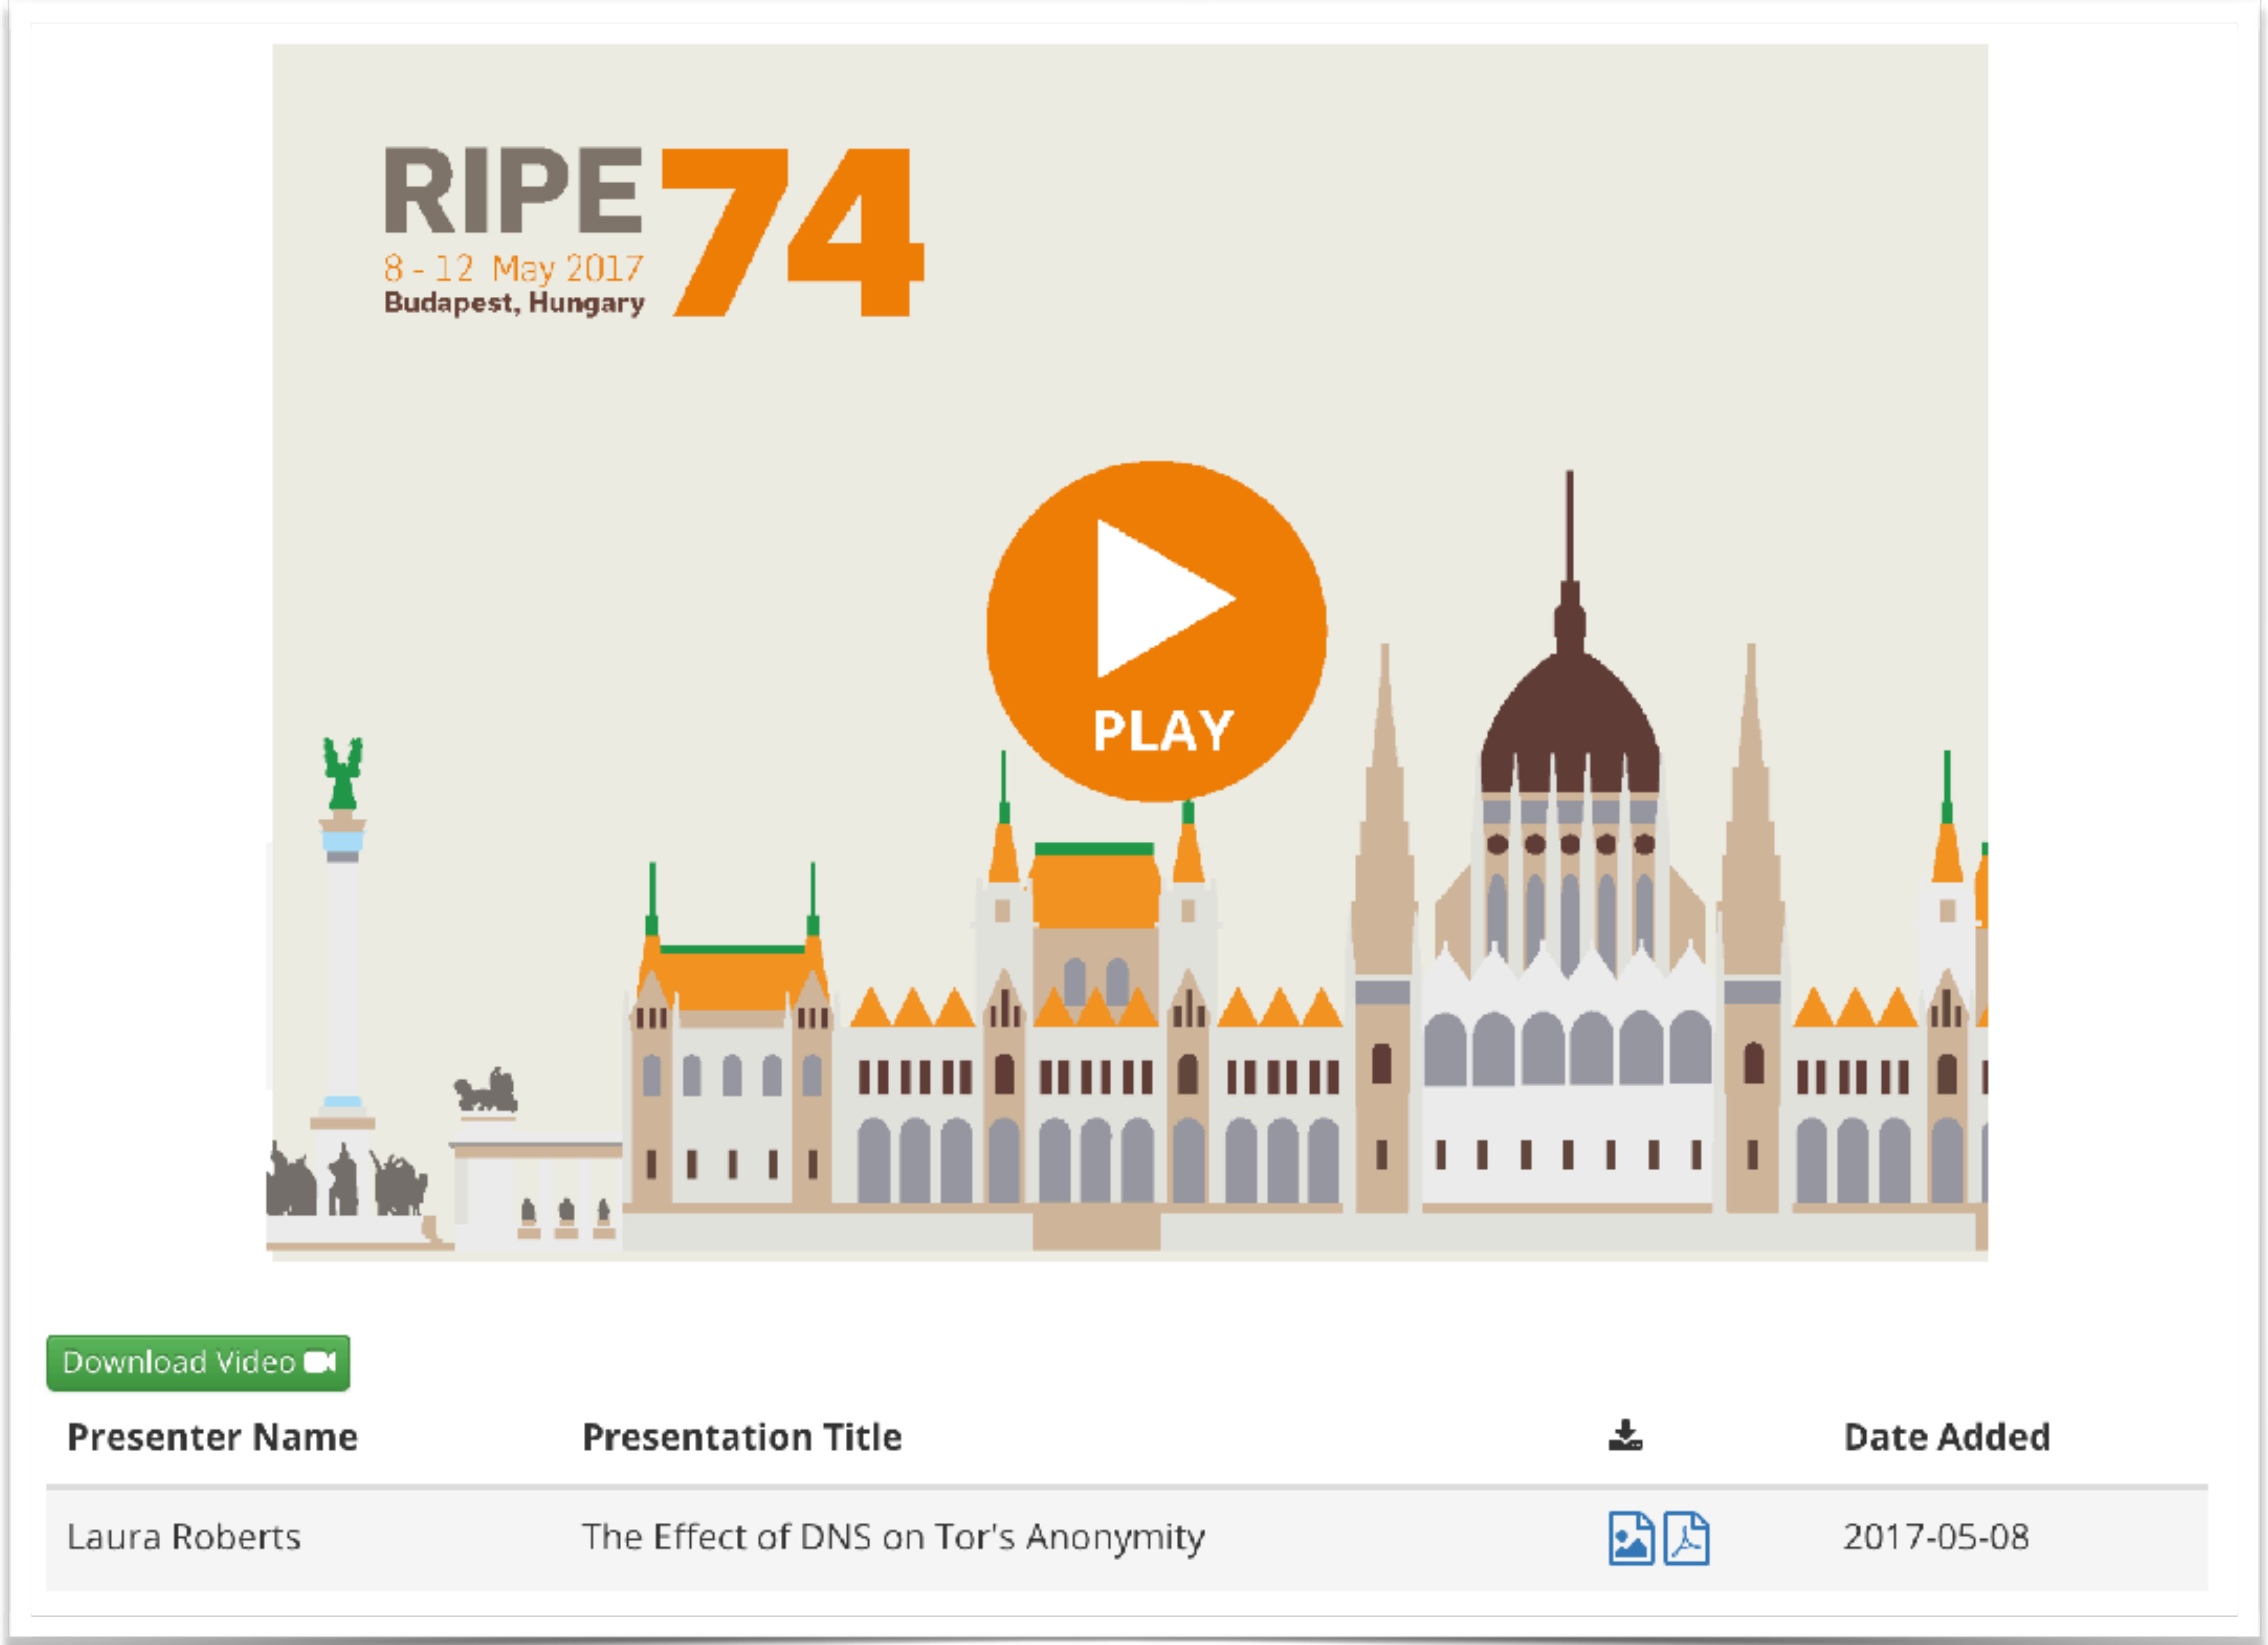

Presenter Name	Presentation Title	⬇	Date Added
Laura Roberts	The Effect of DNS on Tor's Anonymity	🖼 📄	2017-05-08

Tor is an anonymity network.

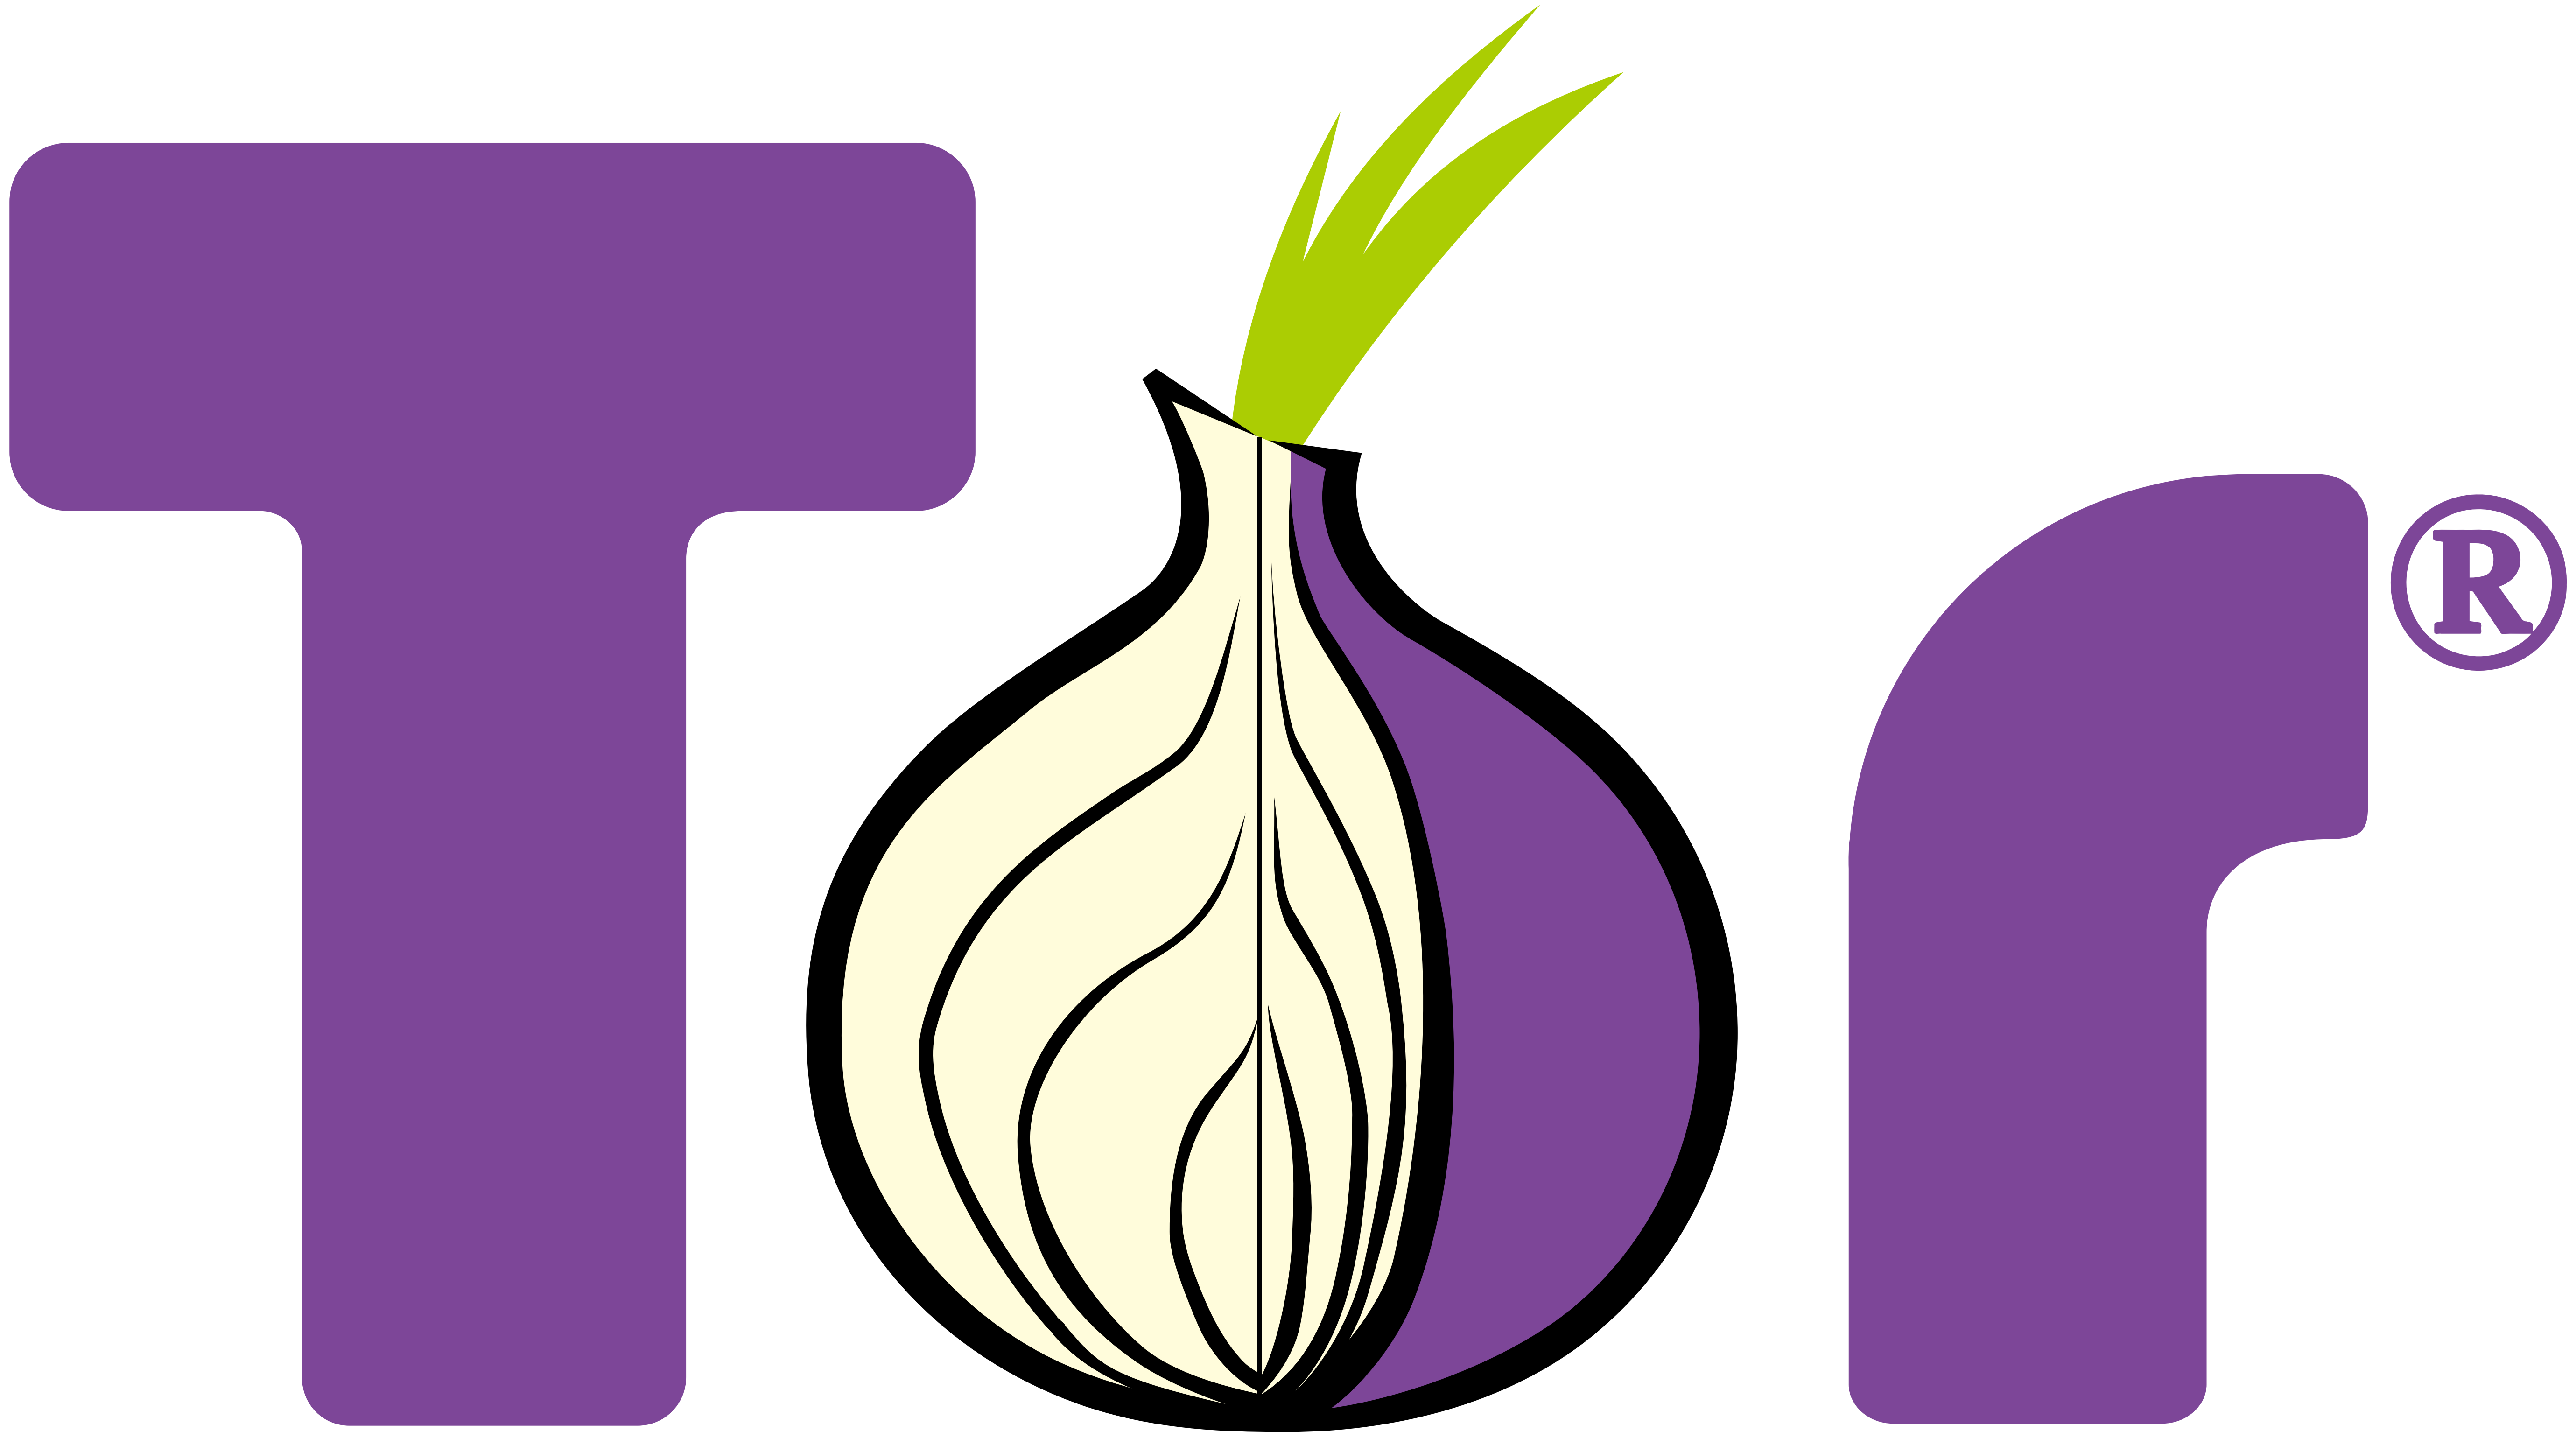

TorProject.org

Relay operators and developers make Tor work.

Relay operators who donate their computers' resources

Relay operators and developers make Tor work.

Relay operators who donate their computers' resources

Developers who maintain the code

Relay operators and developers make Tor work.

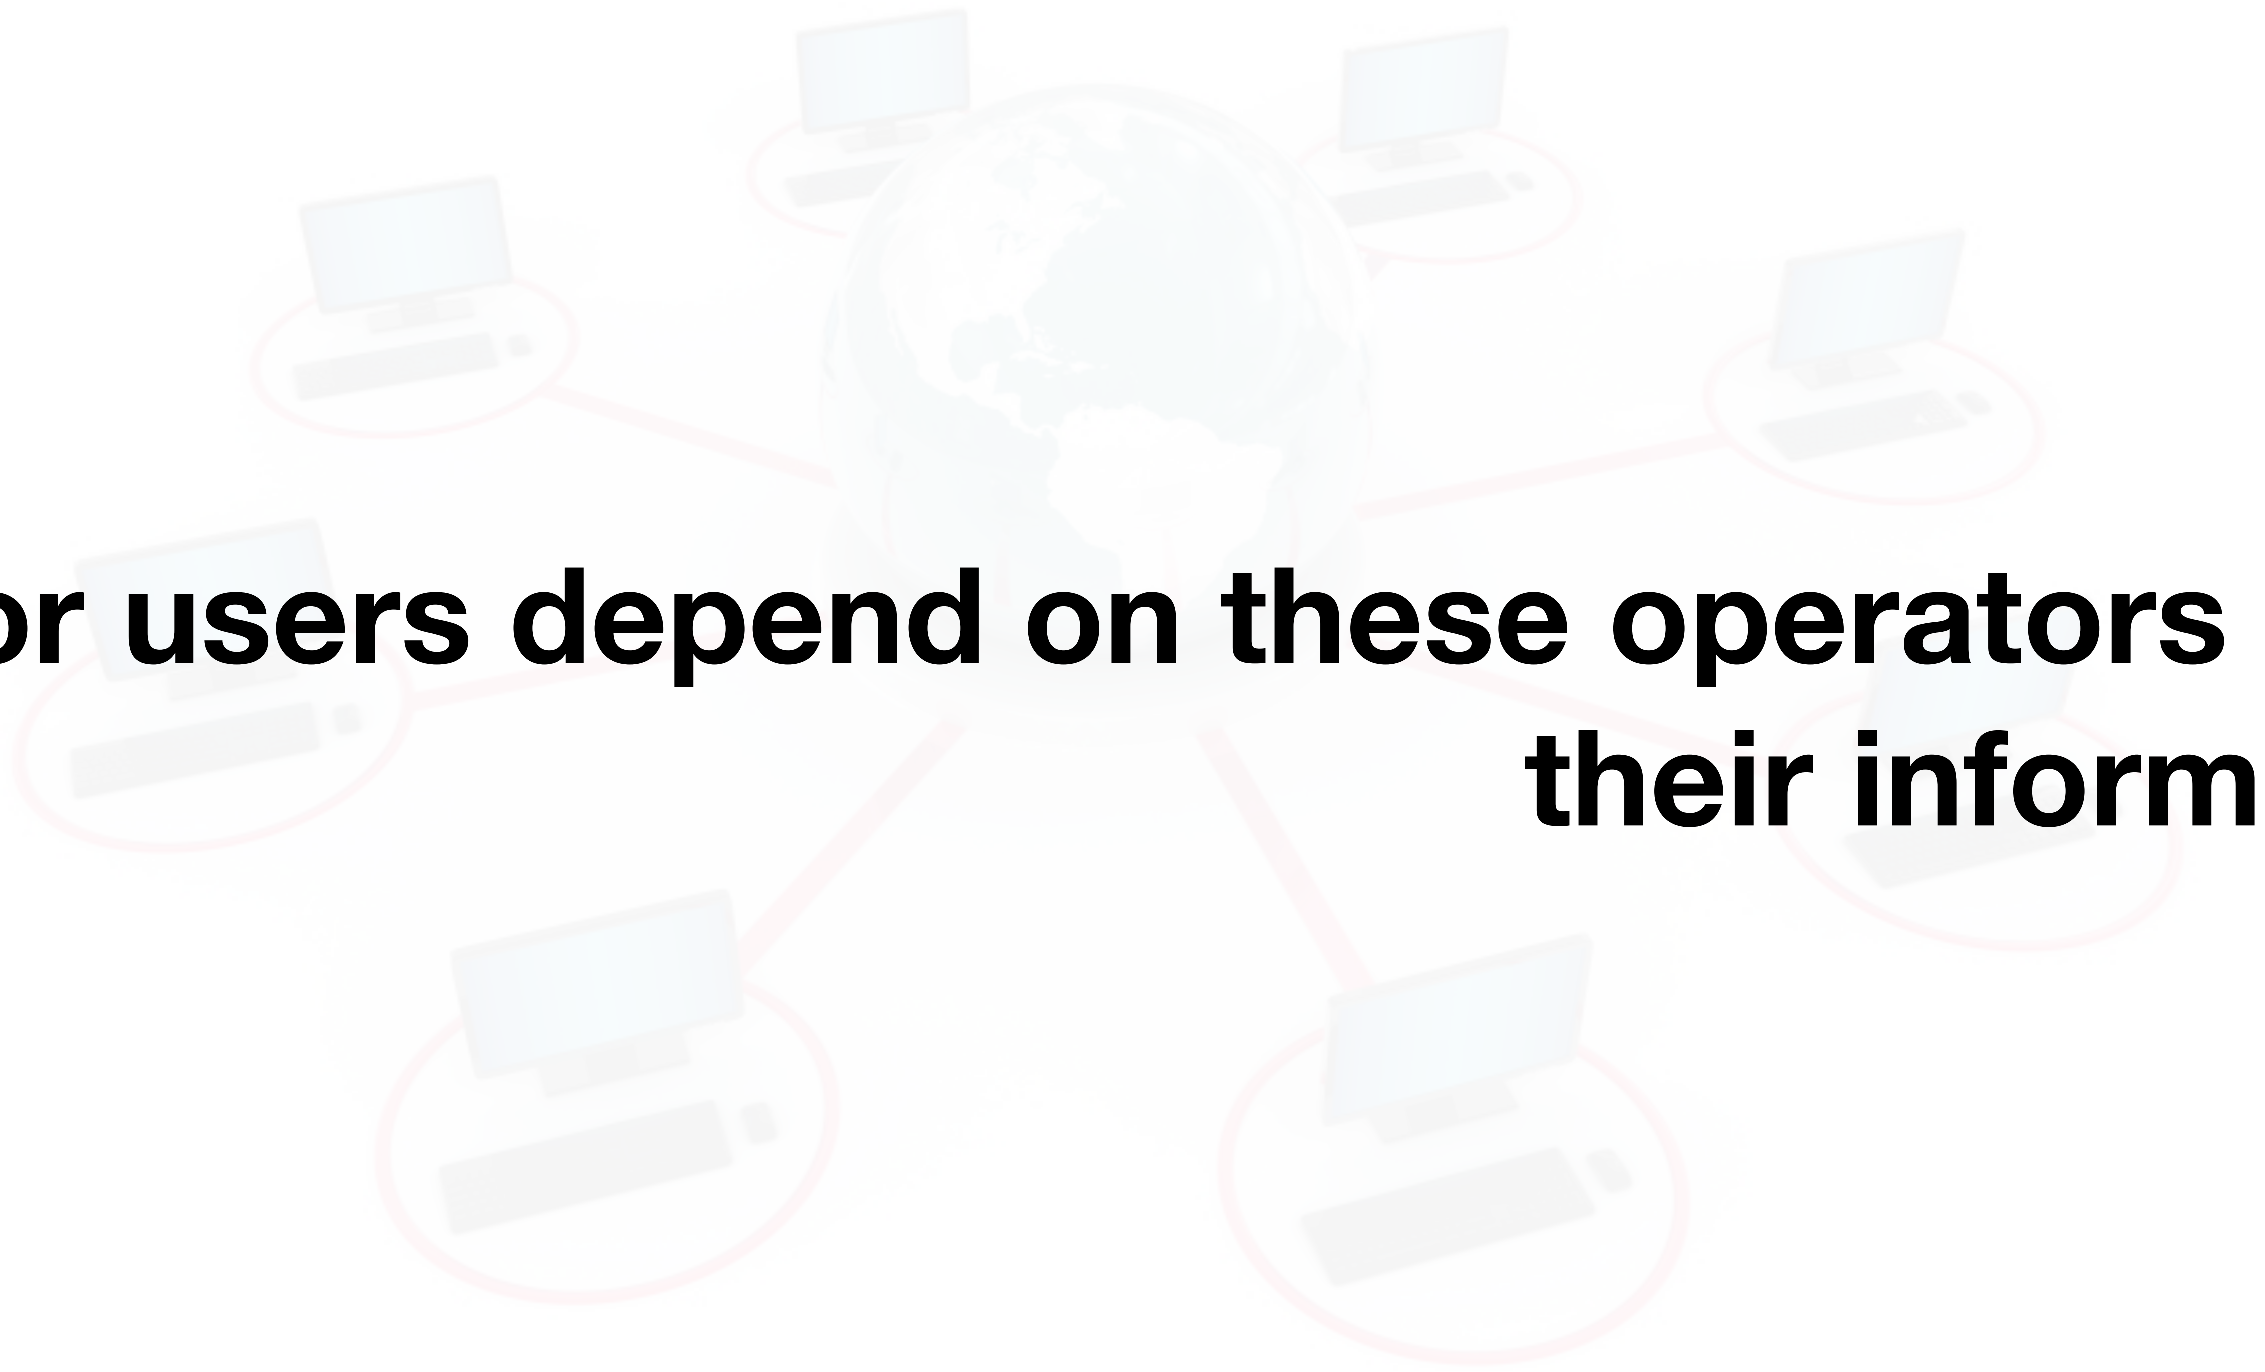

Relay operators who donate their computers' resources

Developers who maintain the code

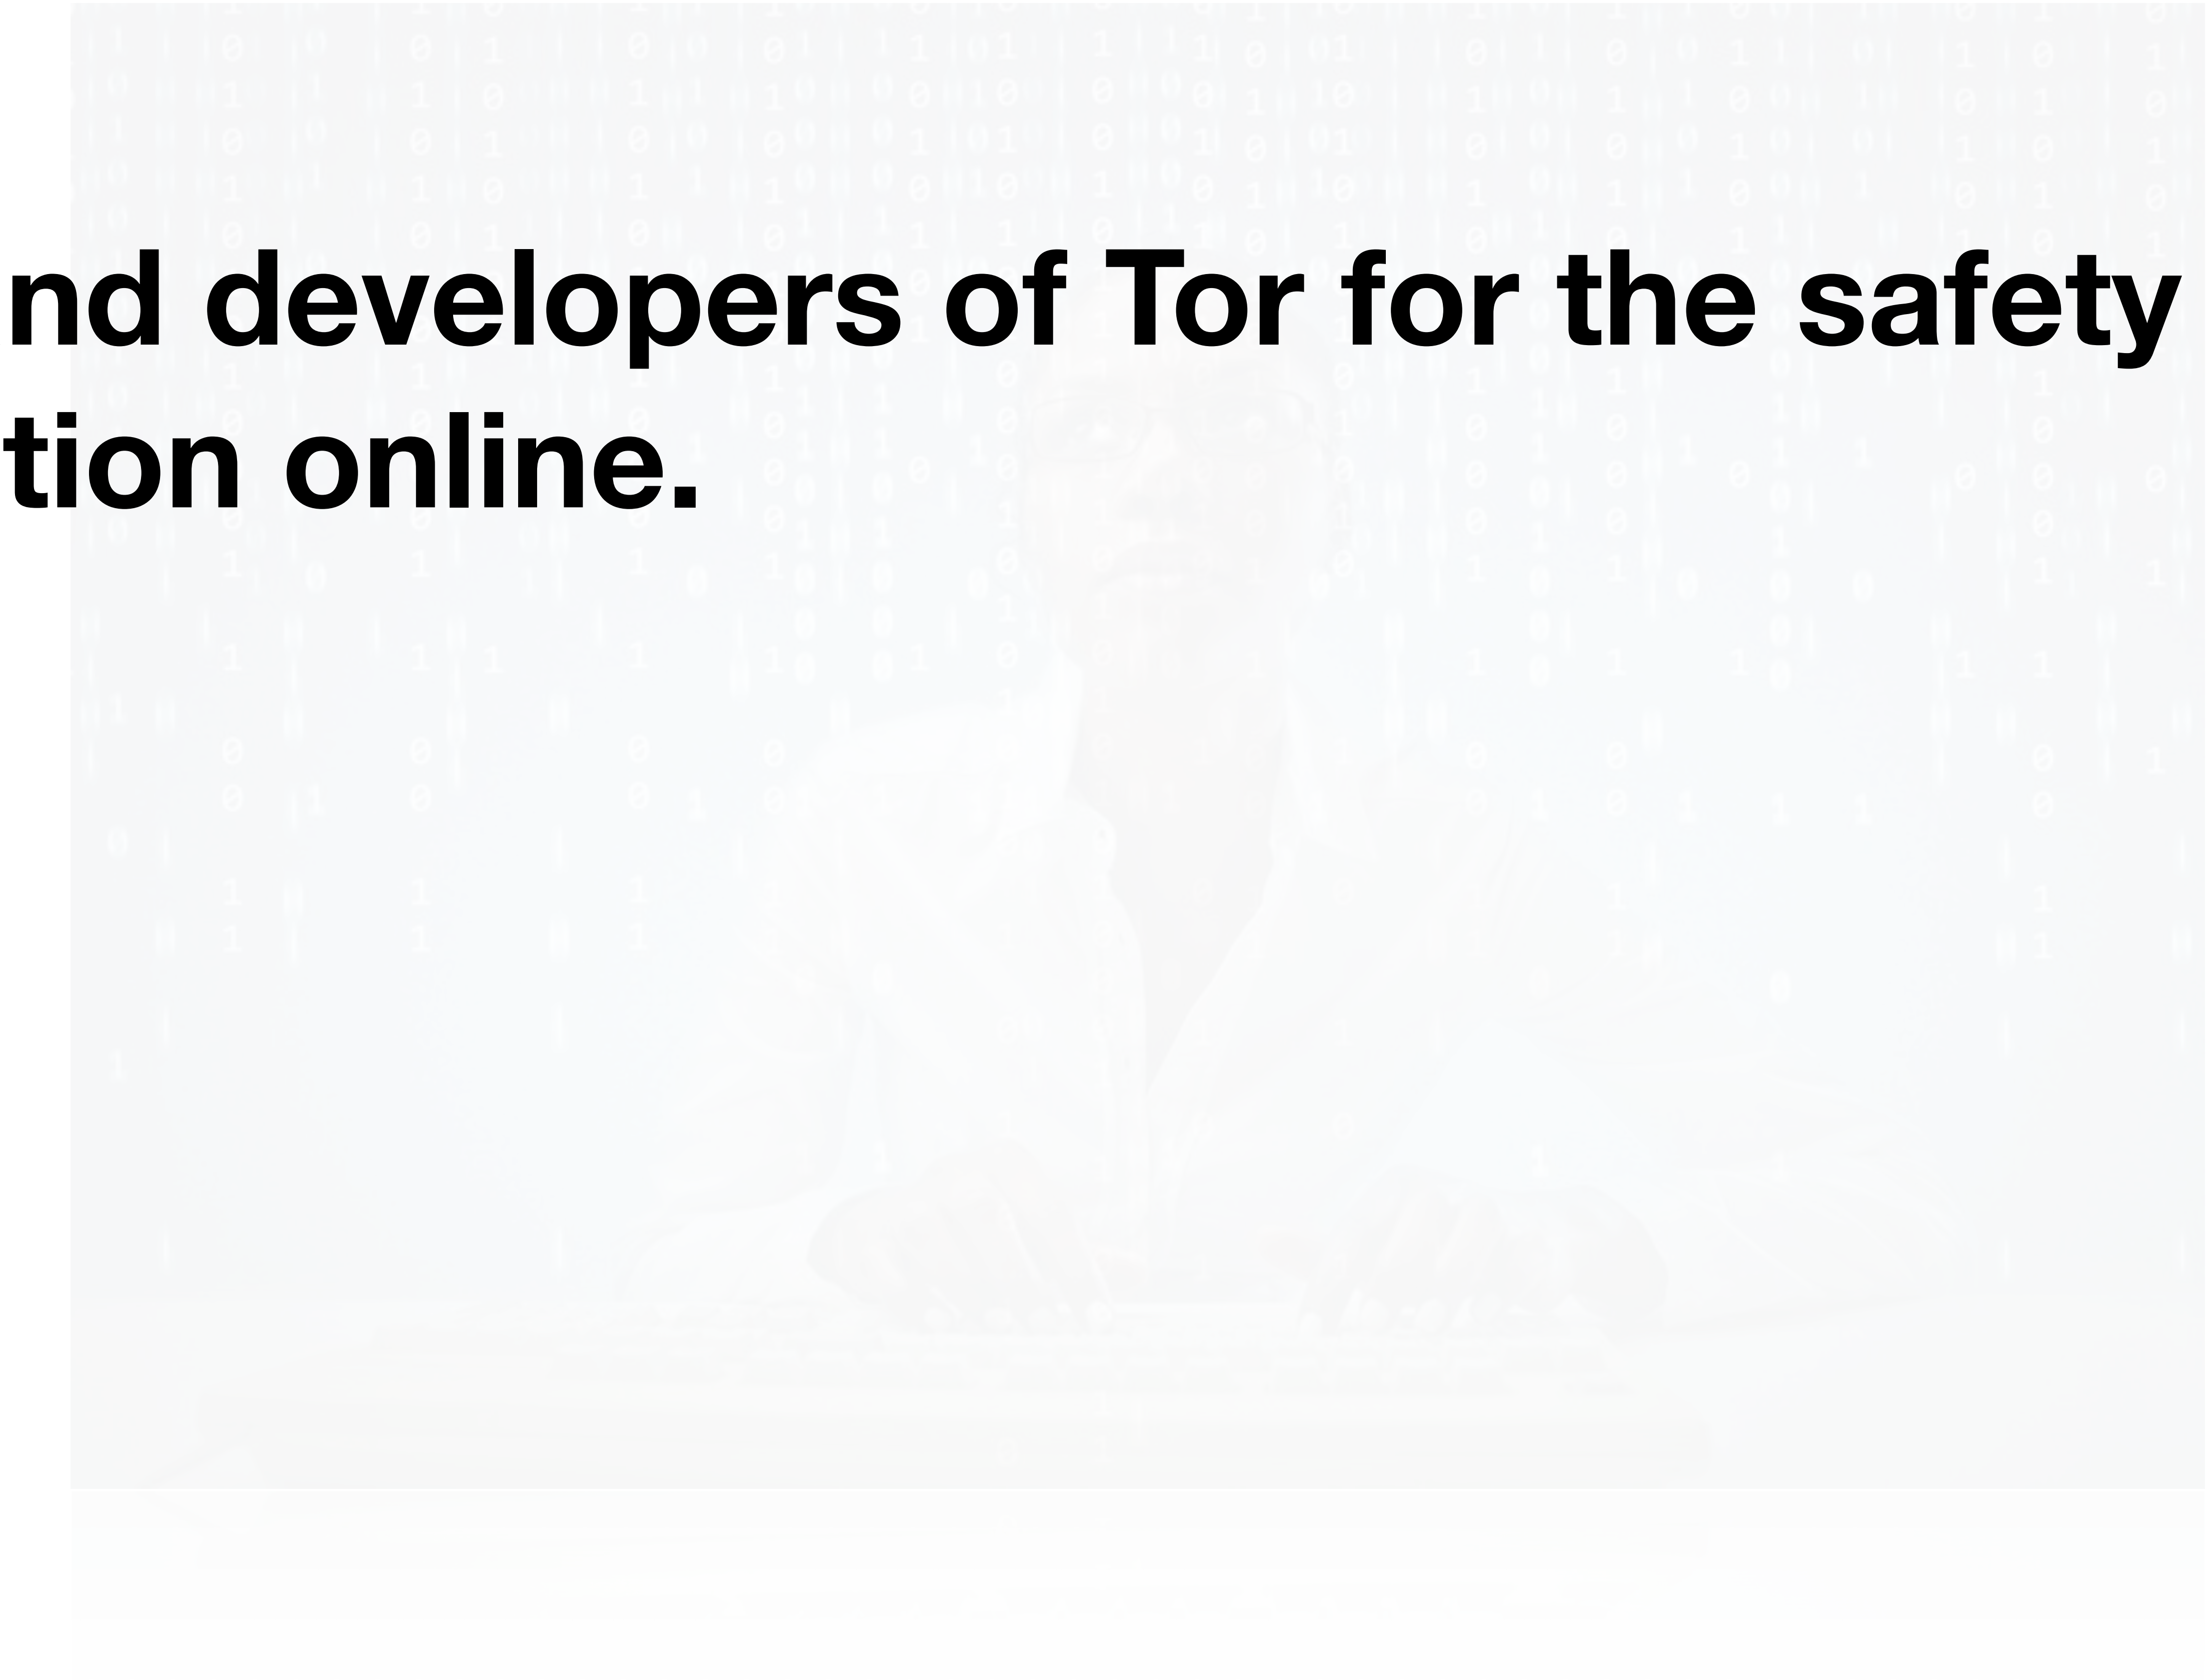

Tor users depend on these operators and developers of Tor for the safety of their information online.

Tor has adversaries.

Our goal was to show that the *domain name system (DNS)* can be used to compromise Tor and to provide recommendations.

Contributions

We discovered that DNS exposes Tor users’ behavior to more adversaries than previously thought. (We focus on autonomous systems as adversaries in our work.)

We discovered that Google gets to learn a lot about what websites Tor users are visiting via DNS.

We created proof-of-concept deanonymization attacks (“DefecTor” attacks) that take advantage of DNS.

We performed simulations at Internet-scale in order to understand how our attacks could affect real people

Contributions

We discovered that DNS exposes Tor users' behavior to more adversaries than previously thought. (We focus on autonomous systems as adversaries in our work.)

We discovered that Google gets to learn a lot about what websites Tor users are visiting via DNS.

We created proof-of-concept deanonymization attacks ("DefecTor" attacks) that take advantage of DNS.

We performed simulations at Internet-scale in order to understand how our attacks could affect real people.

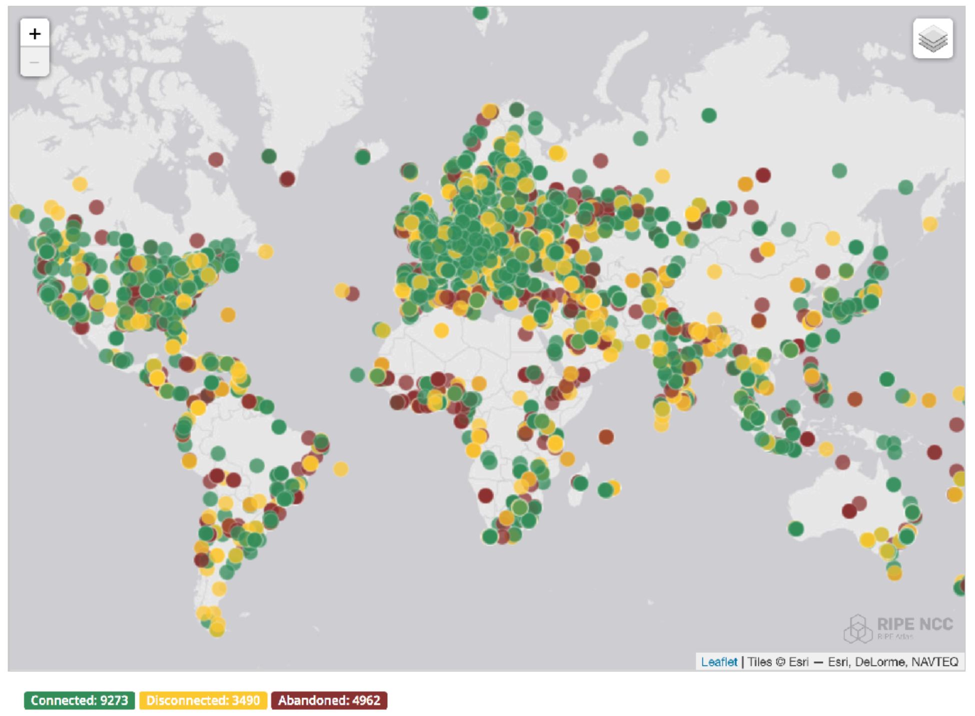

Our use of RIPE Atlas

We quantify the likelihood that any AS is in position to mount DefecTor attacks.

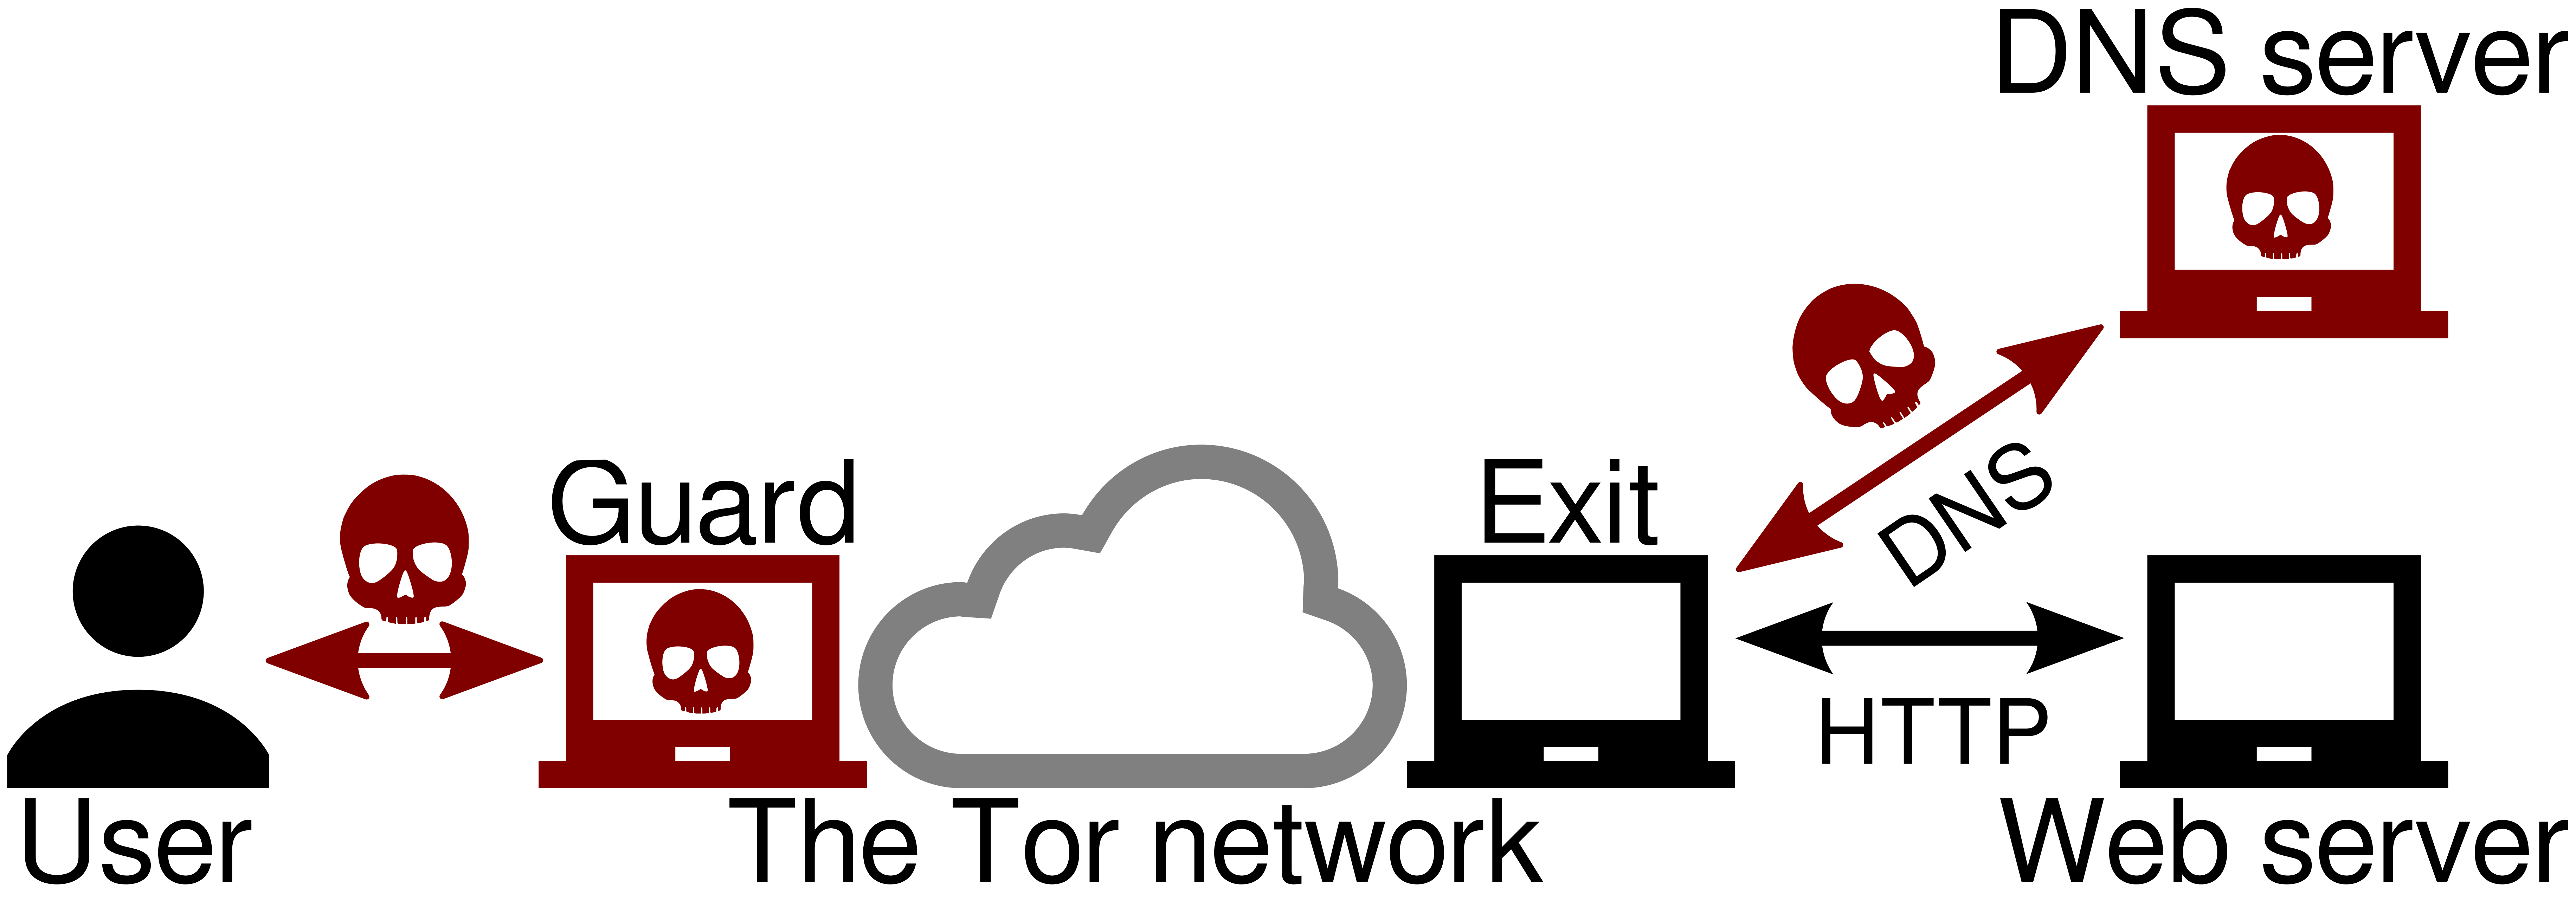

An AS needs to be on both ends of the Tor network in order to mount DefecTor attacks.

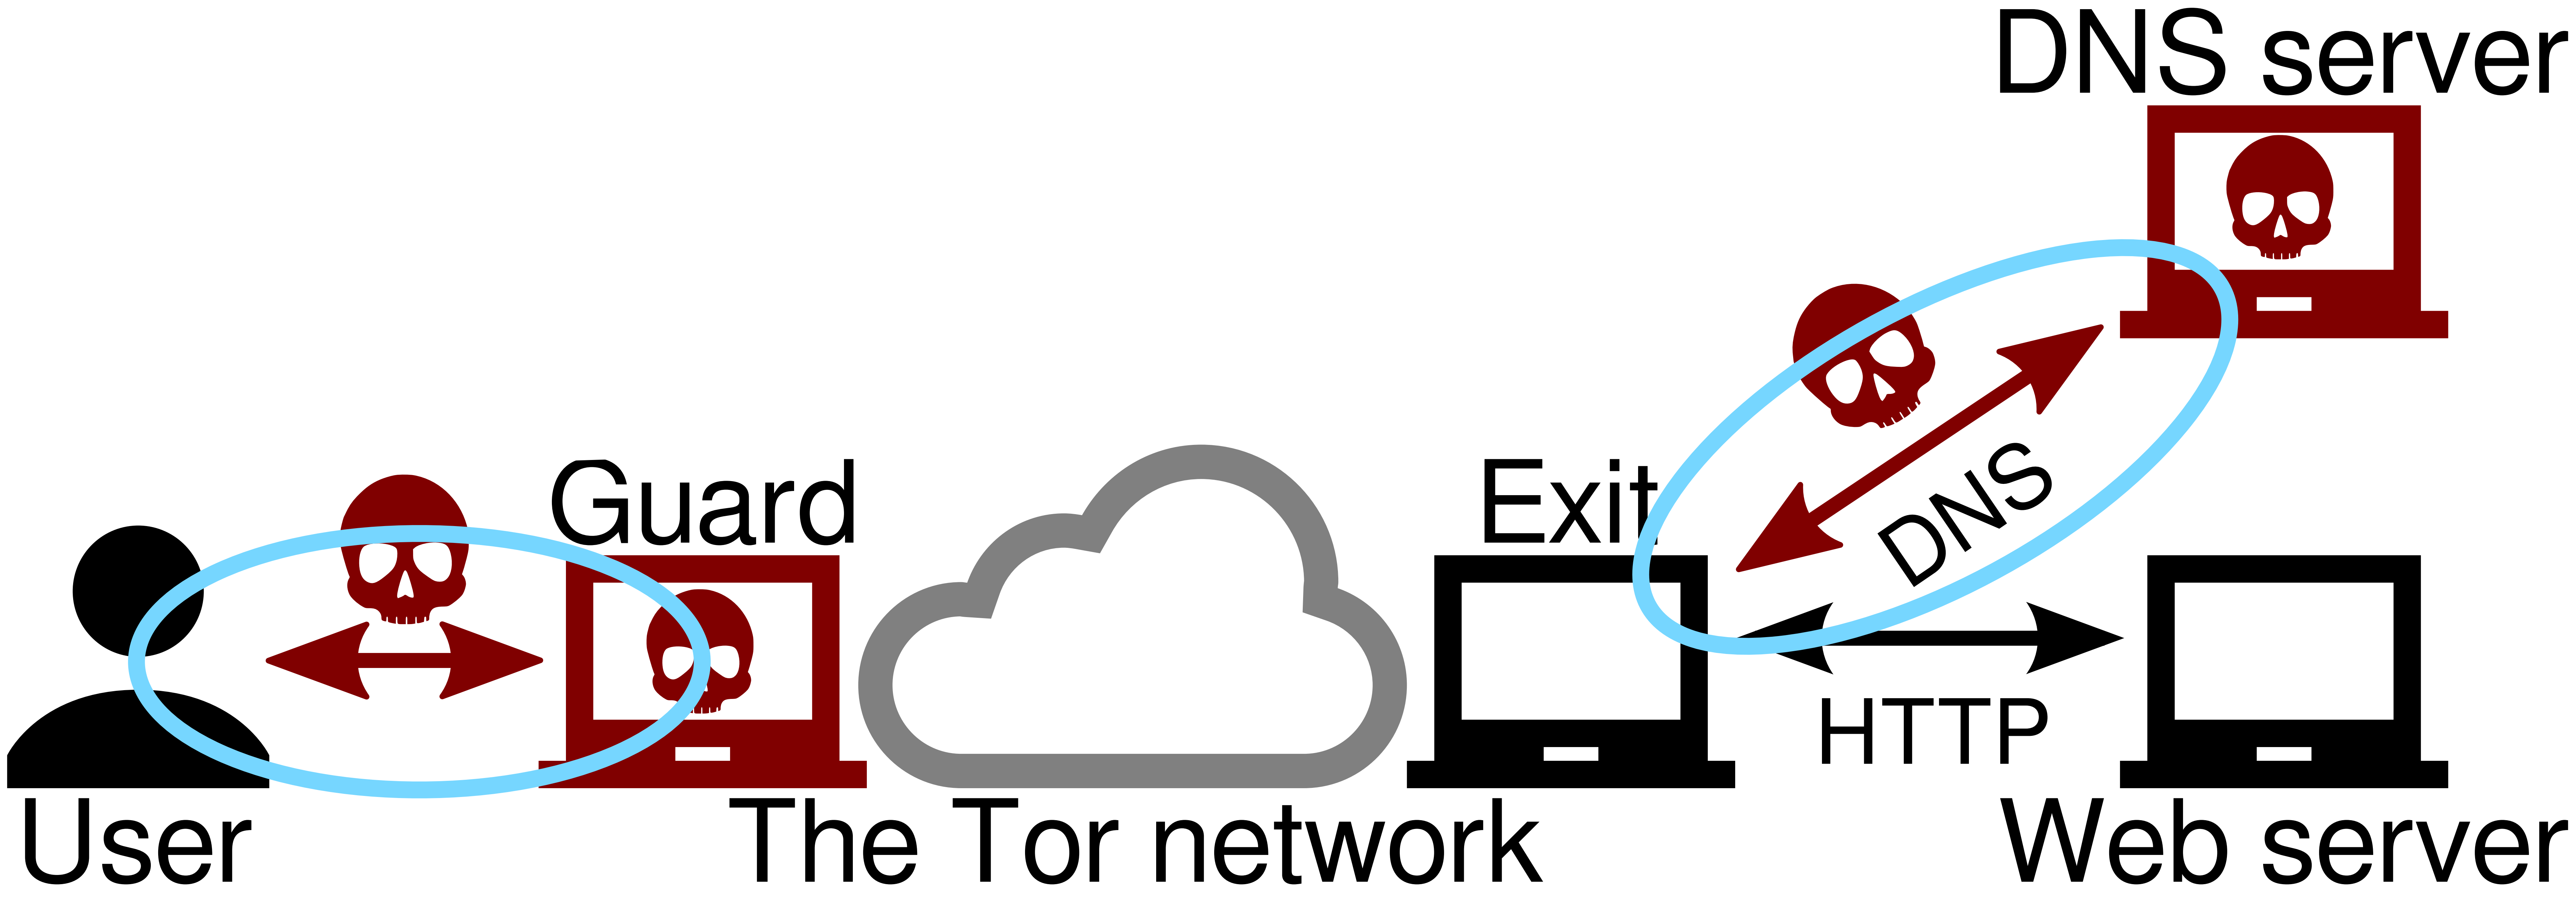

An AS needs to be on both ends of the Tor network in order to mount DefecTor attacks.

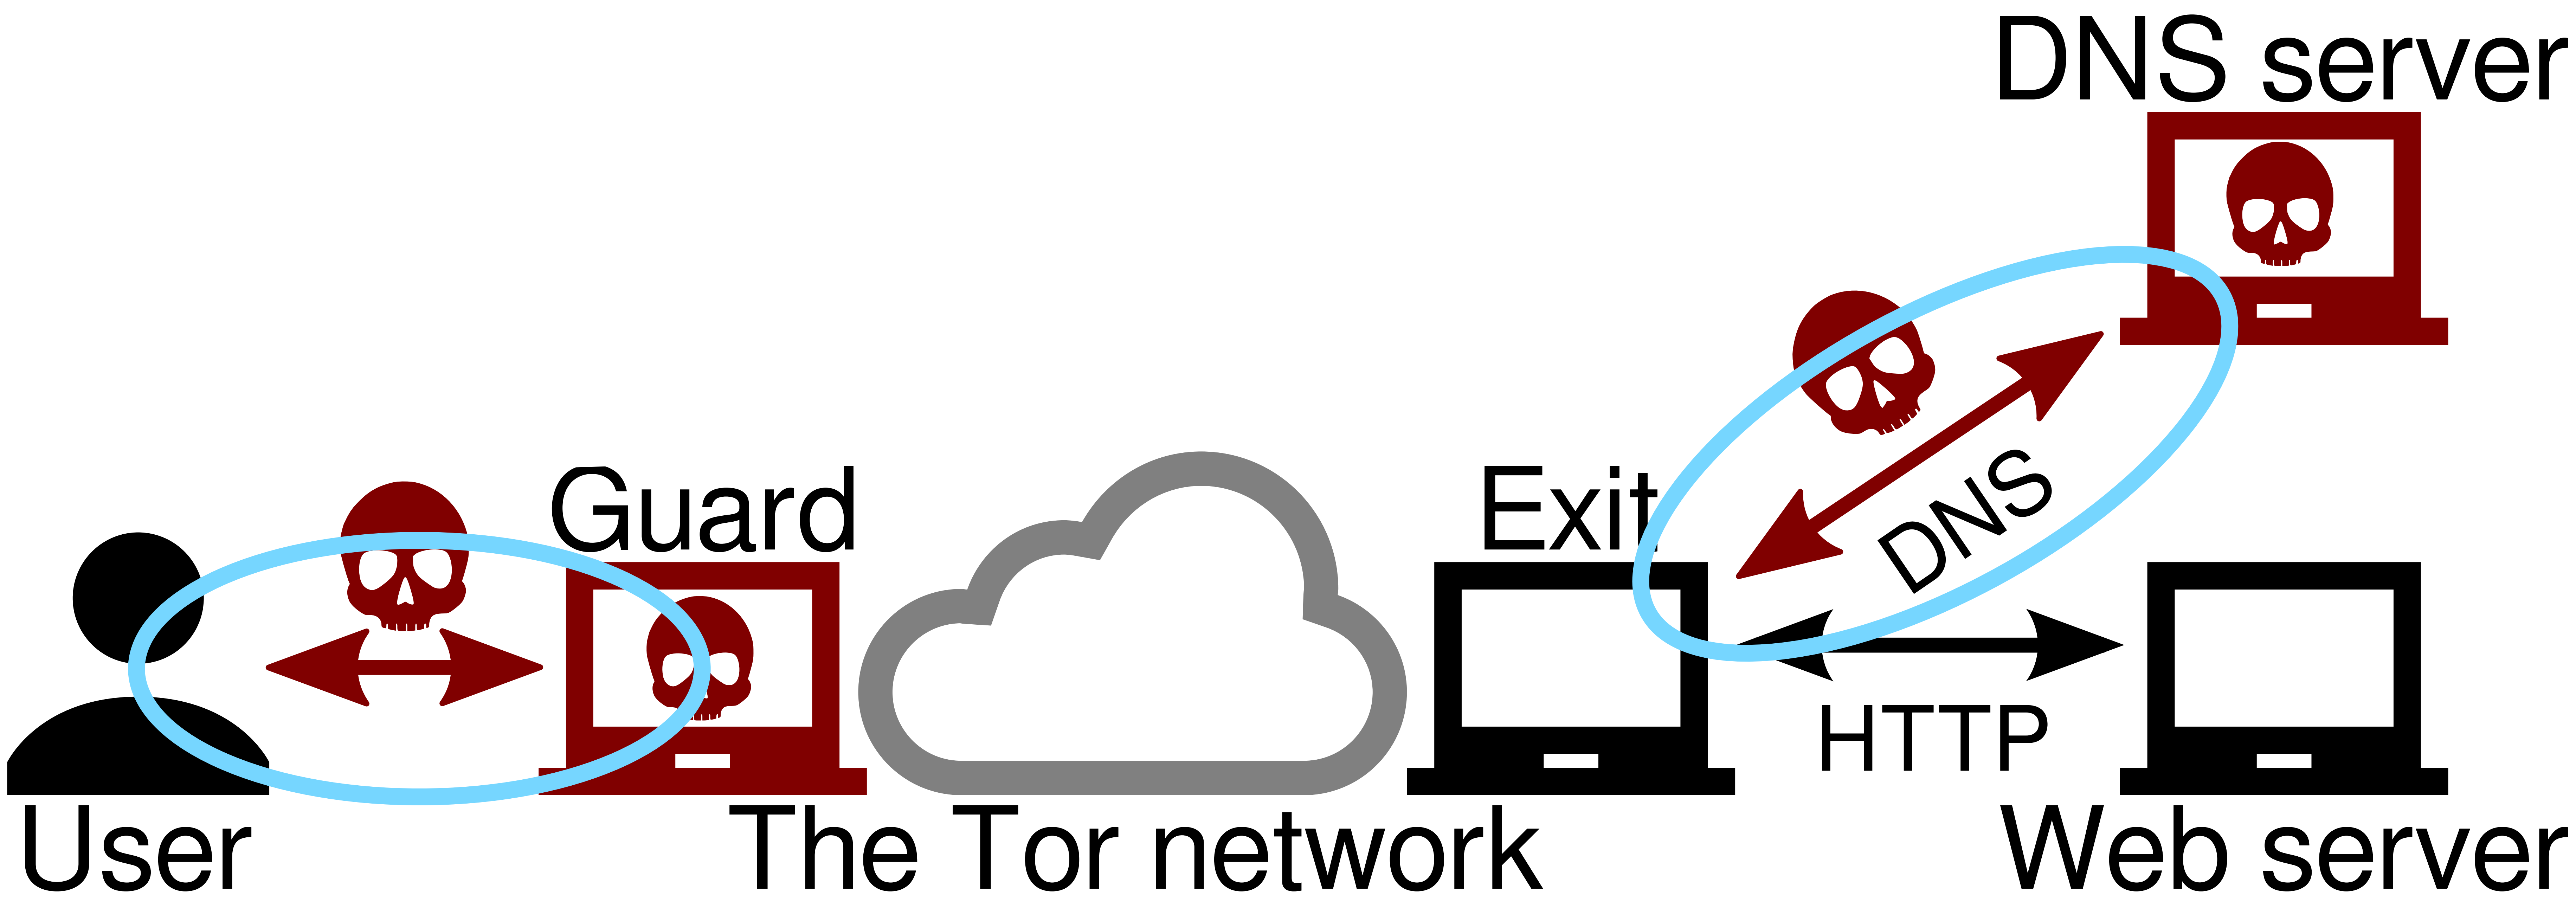

An AS needs to be on both ends of the Tor network in order to mount DefecTor attacks.

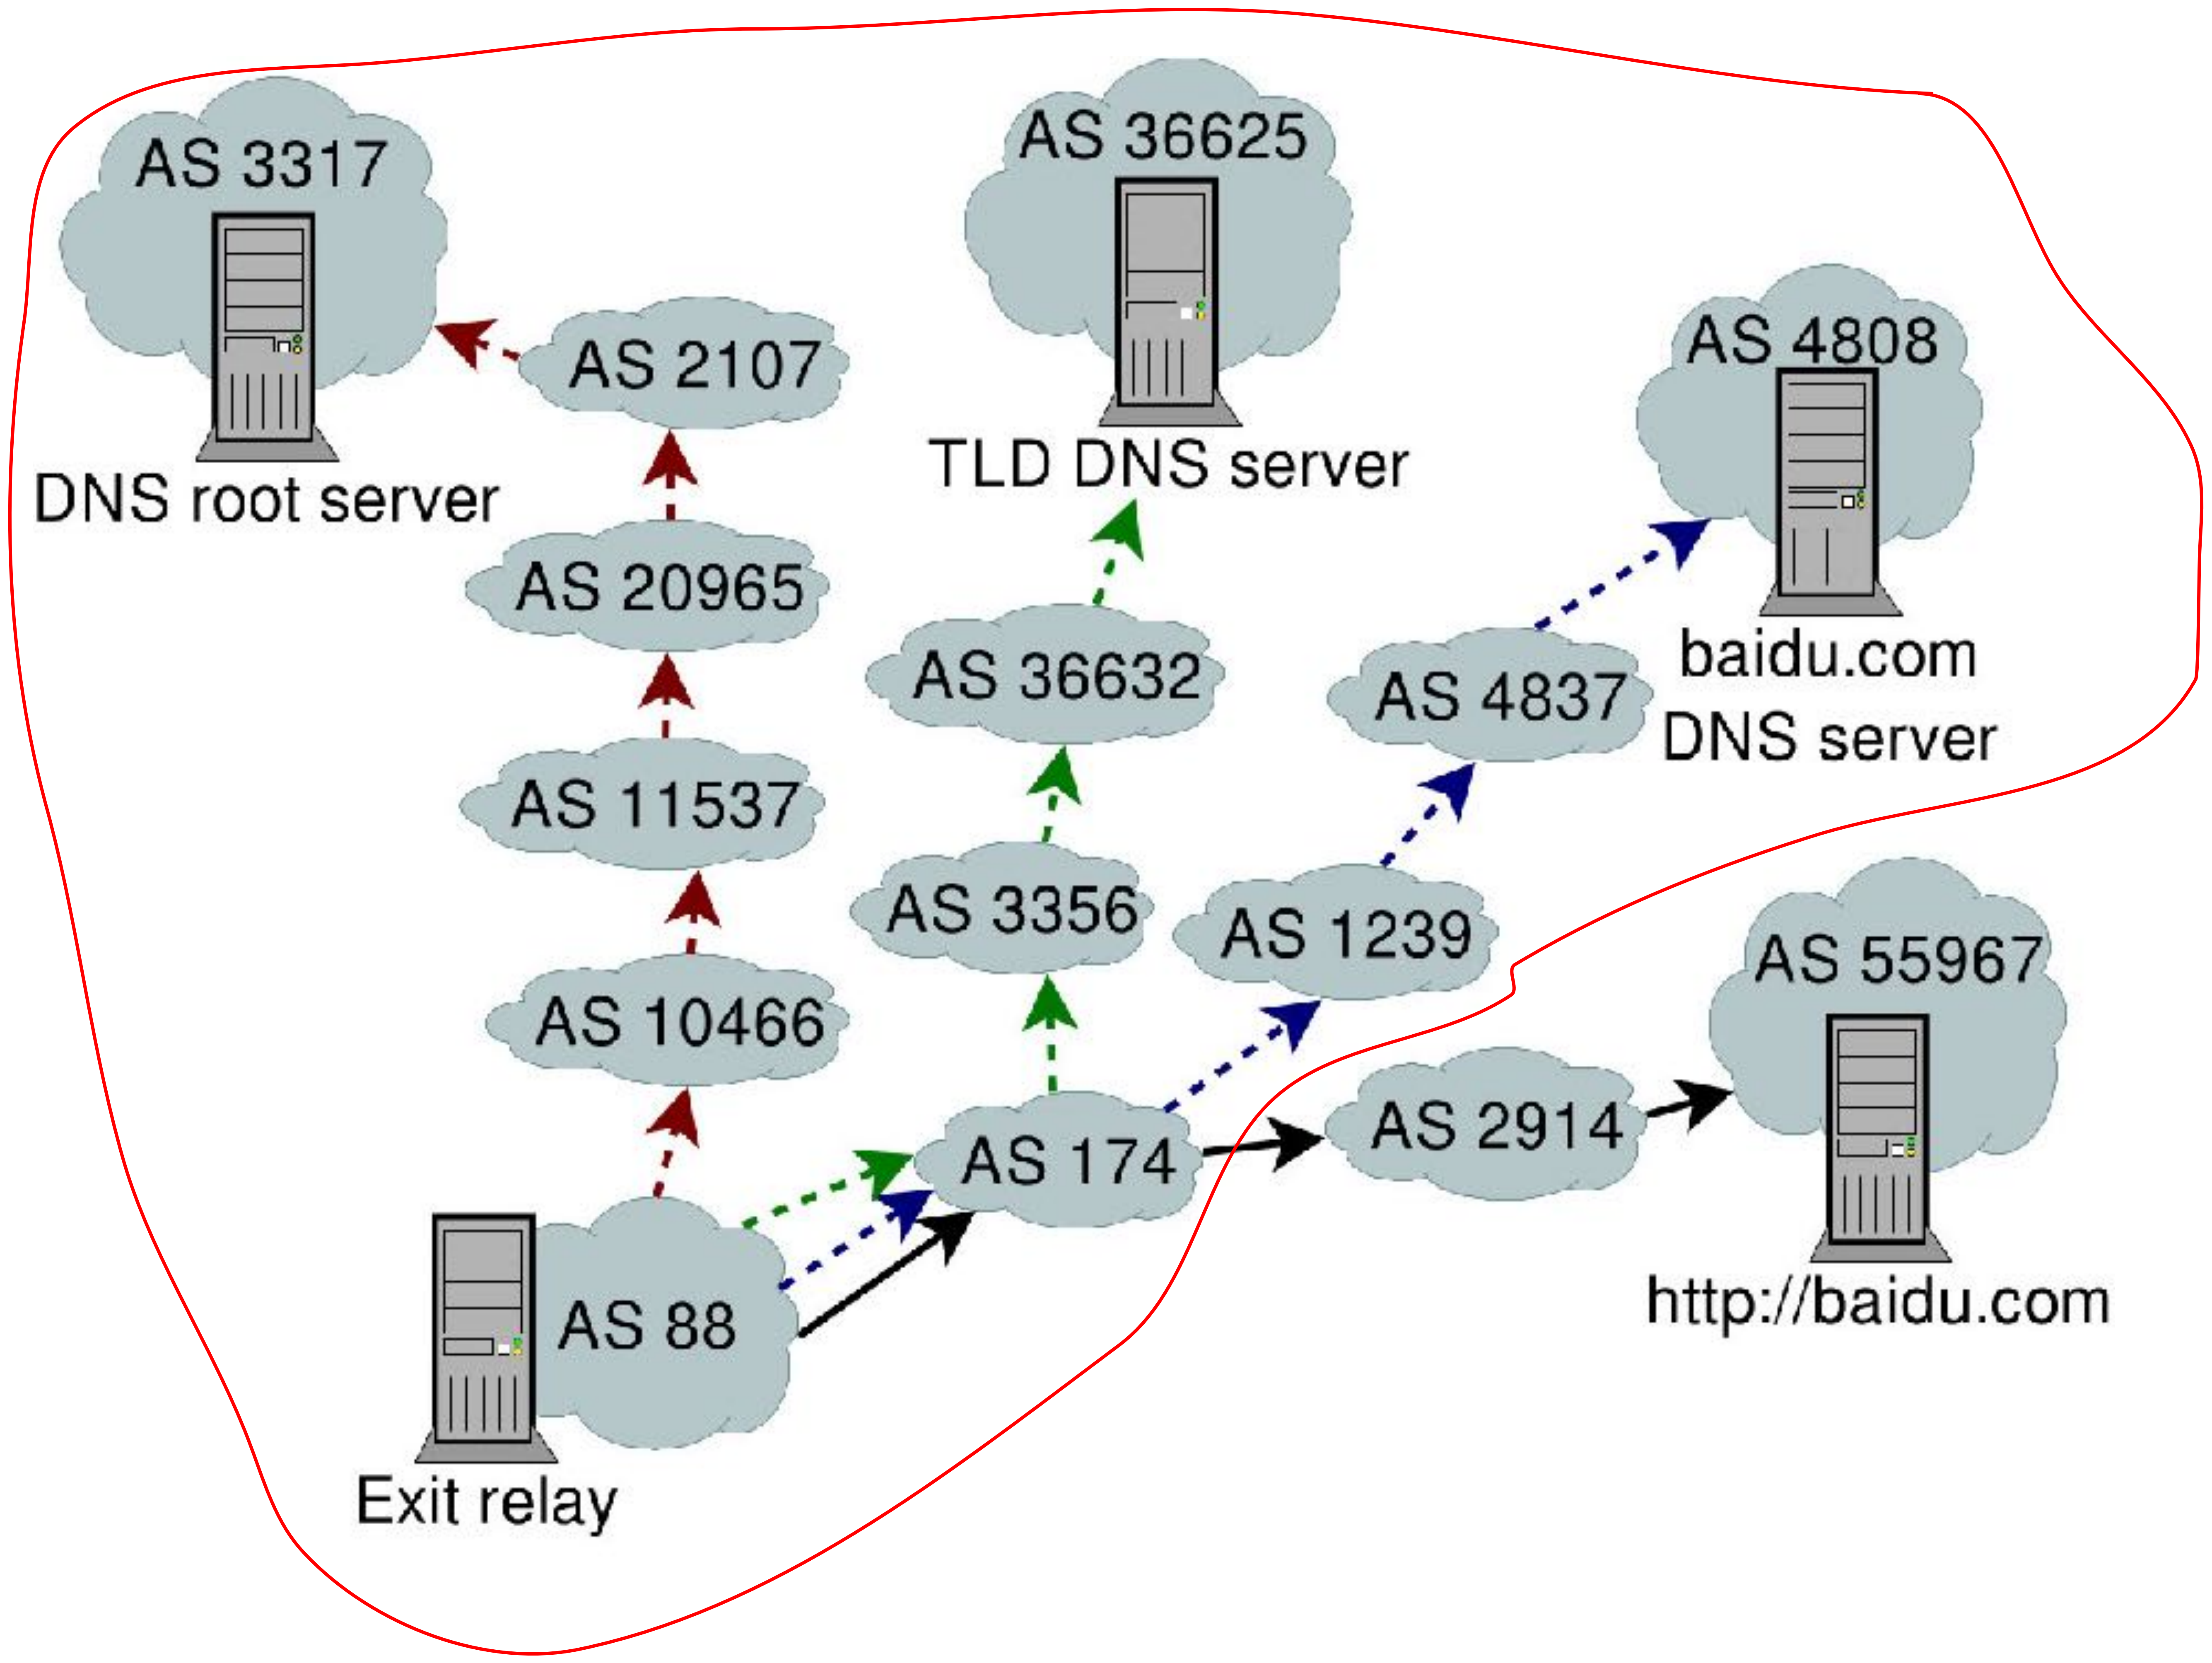

An AS needs to be on both ends of the Tor network in order to mount DefecTor attacks.

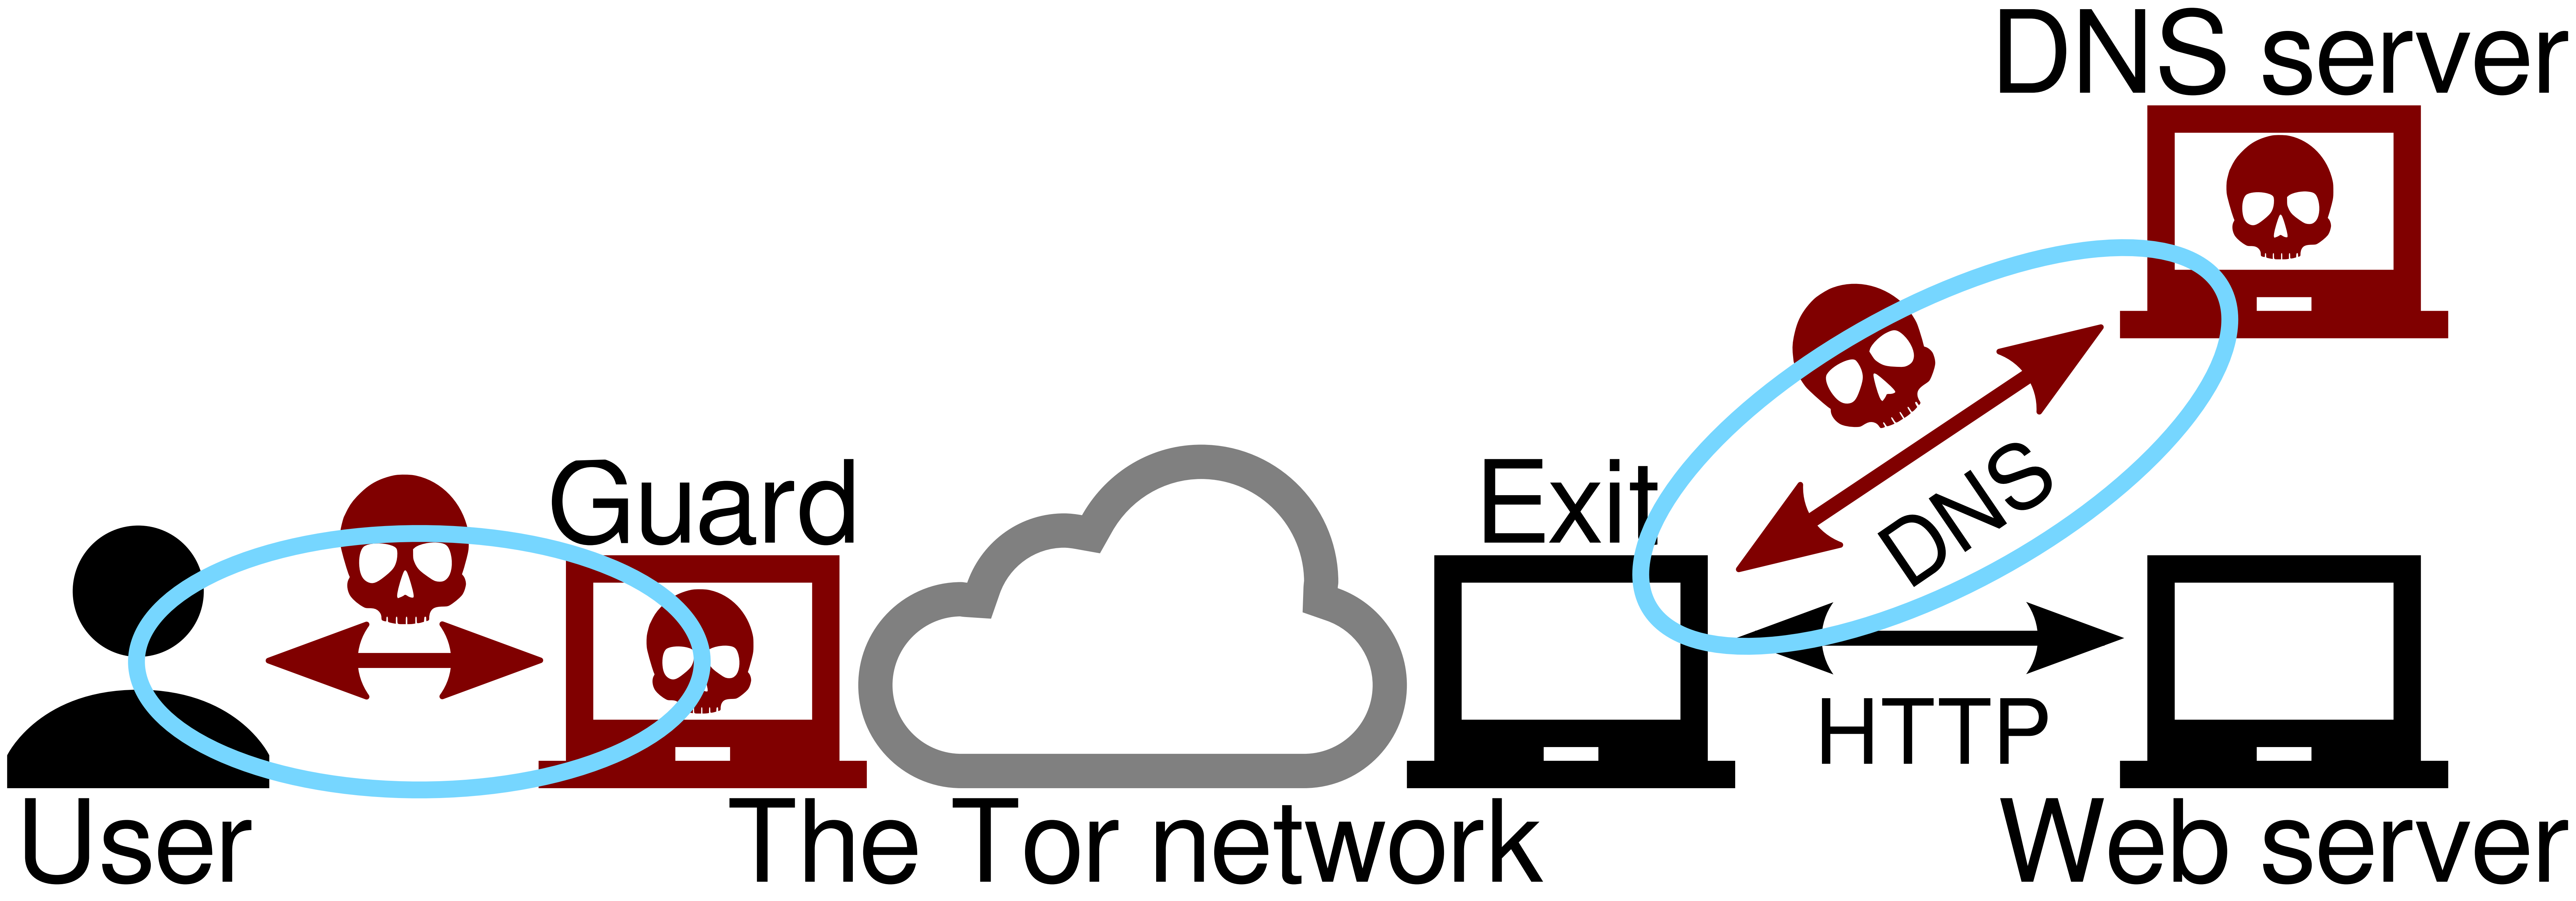

We simulate Tor user activity with Tor Path Simulator (TorPS).

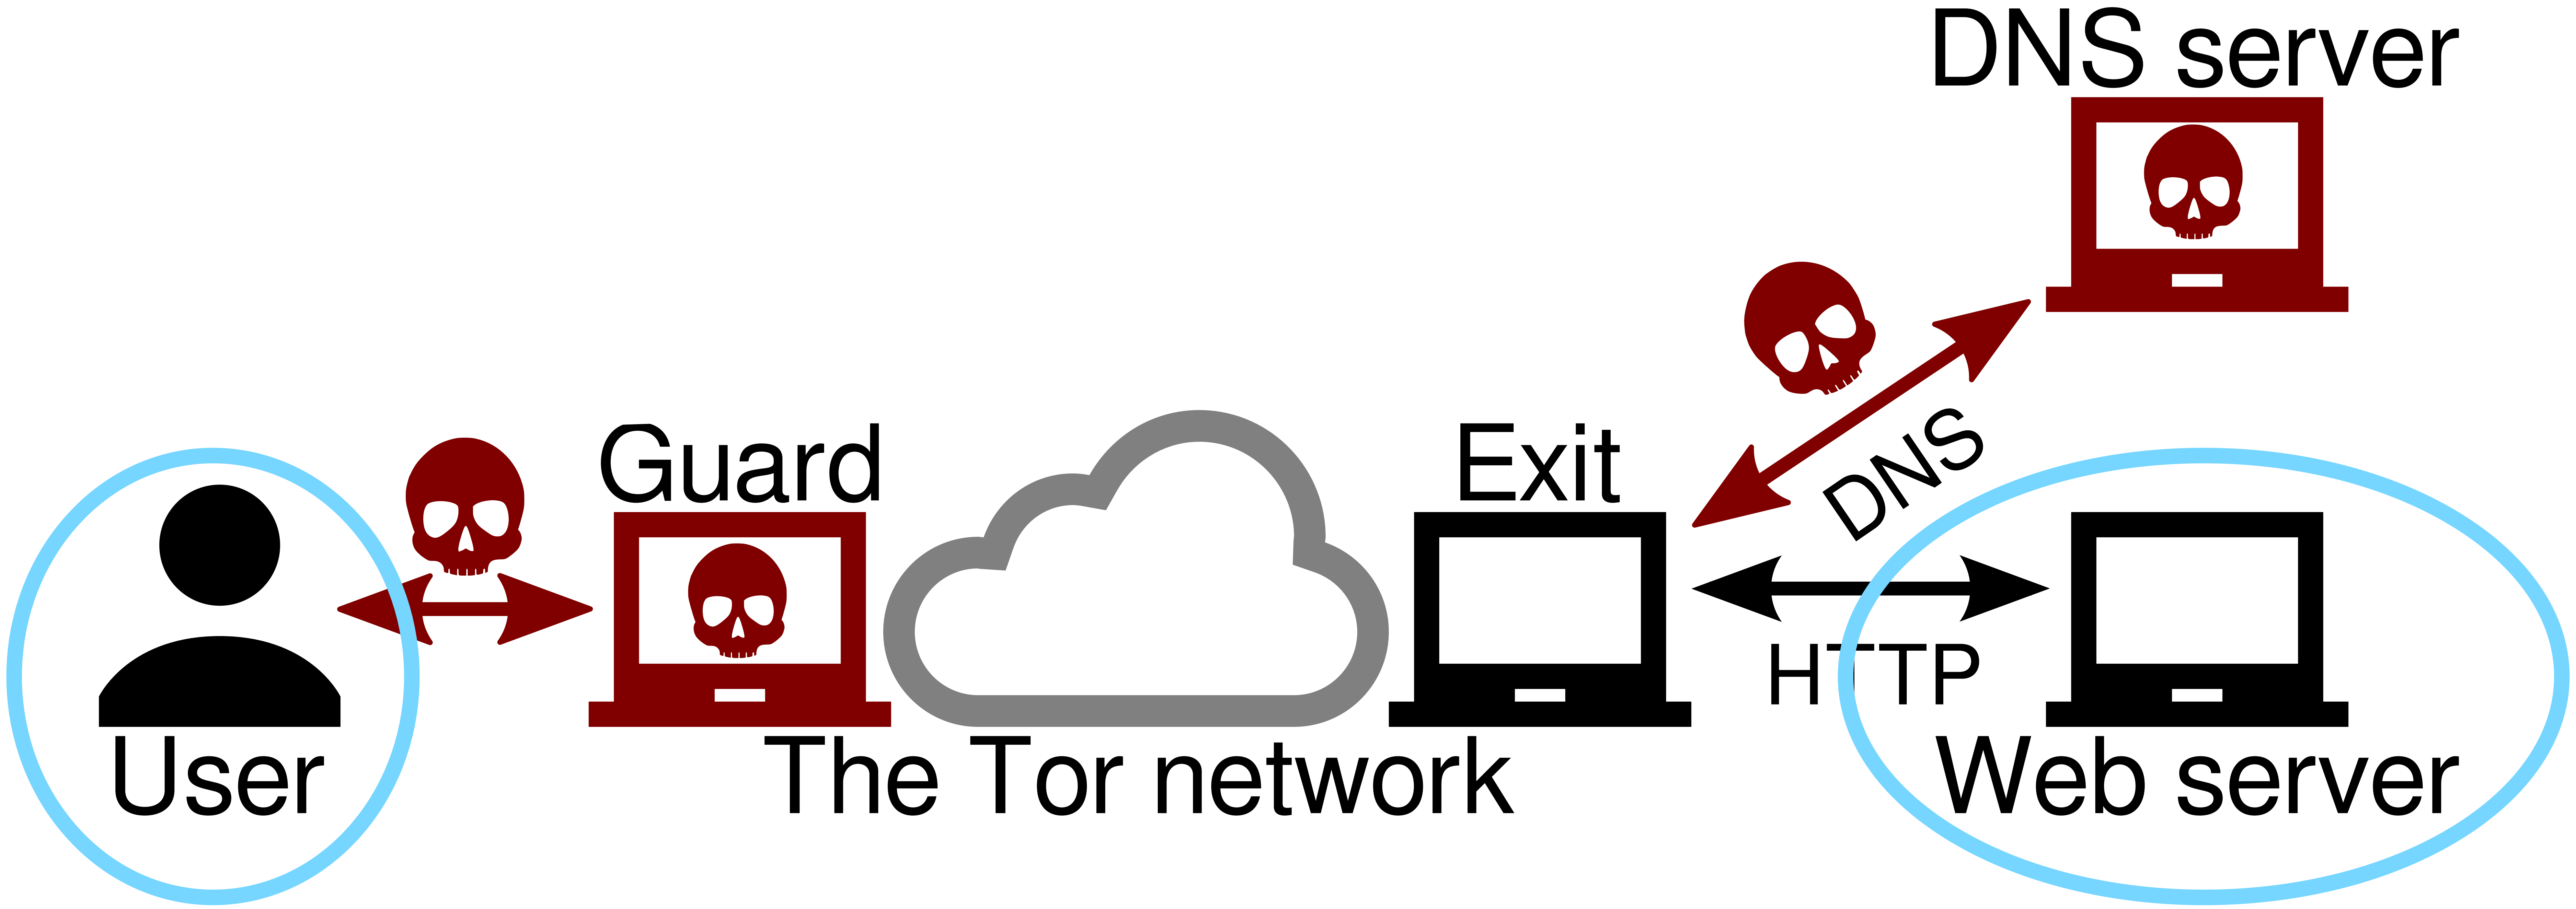

We simulate Tor user activity with Tor Path Simulator (TorPS).

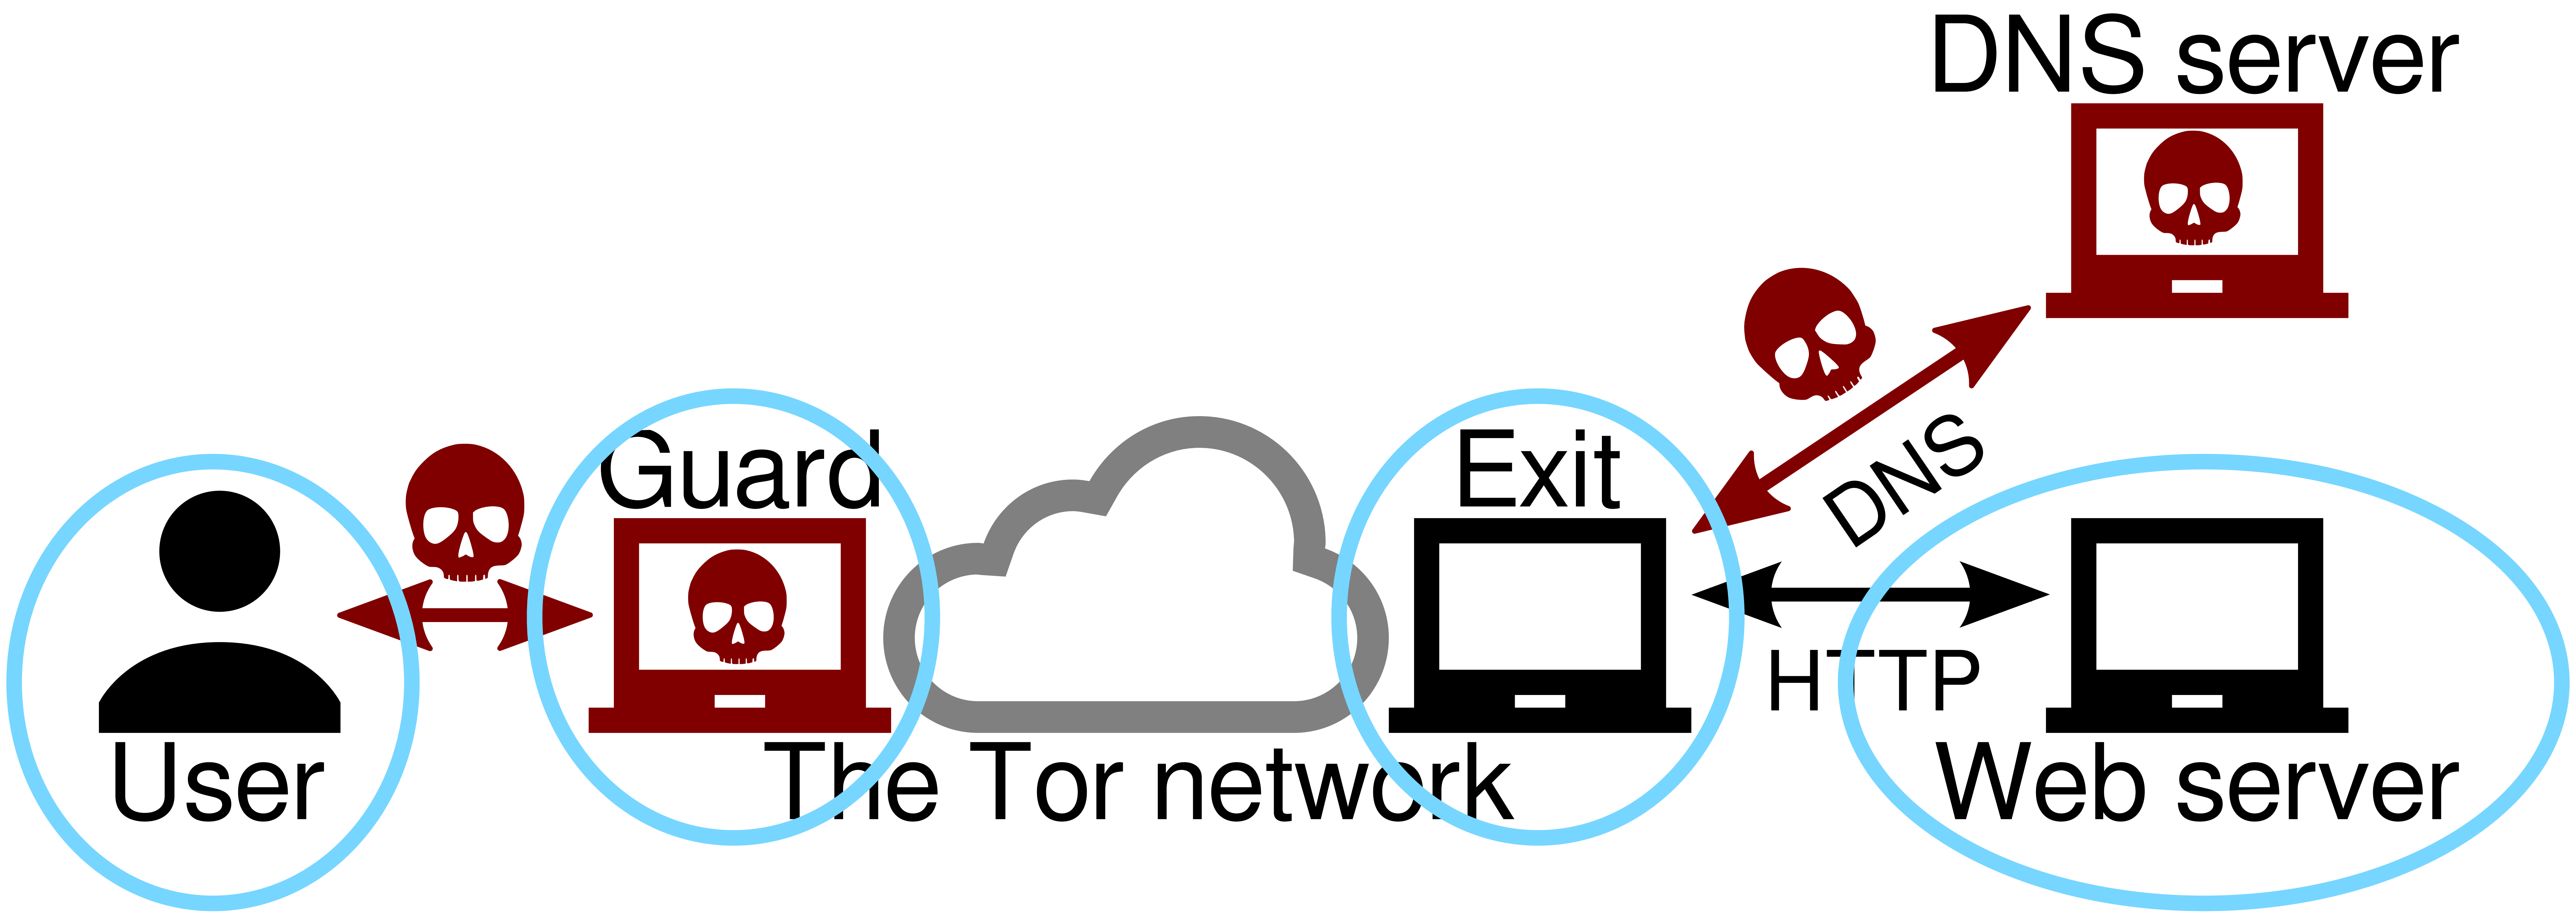

We simulate Tor user activity with Tor Path Simulator (TorPS).

March 2016

9am EST: Gmail and Twitter

12pm: Google Calendar and Google Docs

3pm: Facebook and Instagram

6pm: google.com, startpage.com, and ixquick.com

6:20pm: google.com, startpage.com, and ixquick.com

User

Web server

We simulate Tor user activity with Tor Path Simulator (TorPS).

Our user model made 12 website visits per day for 31 days

Had TorPS create a 100,000-sample simulation for us

Each sample had 372 opportunities to be compromised by an AS because 12 website visits * 31 days = 372 circuits, or “streams”

We placed our simulated user in five popular Tor ASes.

United States - Comcast

Russia - Rostelecom

Germany - Deutsche Telekom

France - Orange

U.K. - British Telecom

We analyzed four Tor exit relay DNS set-up scenarios.

First, what if all Tor exit relays were set up to use their ISPs’ resolvers?

We analyzed four Tor exit relay DNS set-up scenarios.

First, what if all Tor exit relays were set up to use their ISPs’ resolvers?

Second, what if all Tor exit relays were set up to use Google’s 8.8.8.8 public resolver?

We analyzed four Tor exit relay DNS set-up scenarios.

First, what if all Tor exit relays were set up to use their ISPs' resolvers?

Second, what if all Tor exit relays were set up to use Google's 8.8.8.8 public resolver?

Third, what if all Tor exit relays were set up to do their own DNS resolution

We analyzed four Tor exit relay DNS set-up scenarios.

First, what if all Tor exit relays were set up to use their ISPs' resolvers?

Second, what if all Tor exit relays were set up to use Google's 8.8.8.8 public resolver?

Third, what if all Tor exit relays were set up to do their own DNS resolution

Fourth, what if all Tor exit relays were set up as they currently are?

We did not use AS path inference to obtain AS-level paths.

We preferred to use traceroutes to obtain AS-level paths.

We obtained AS-level paths using traceroutes via RIPE Atlas.

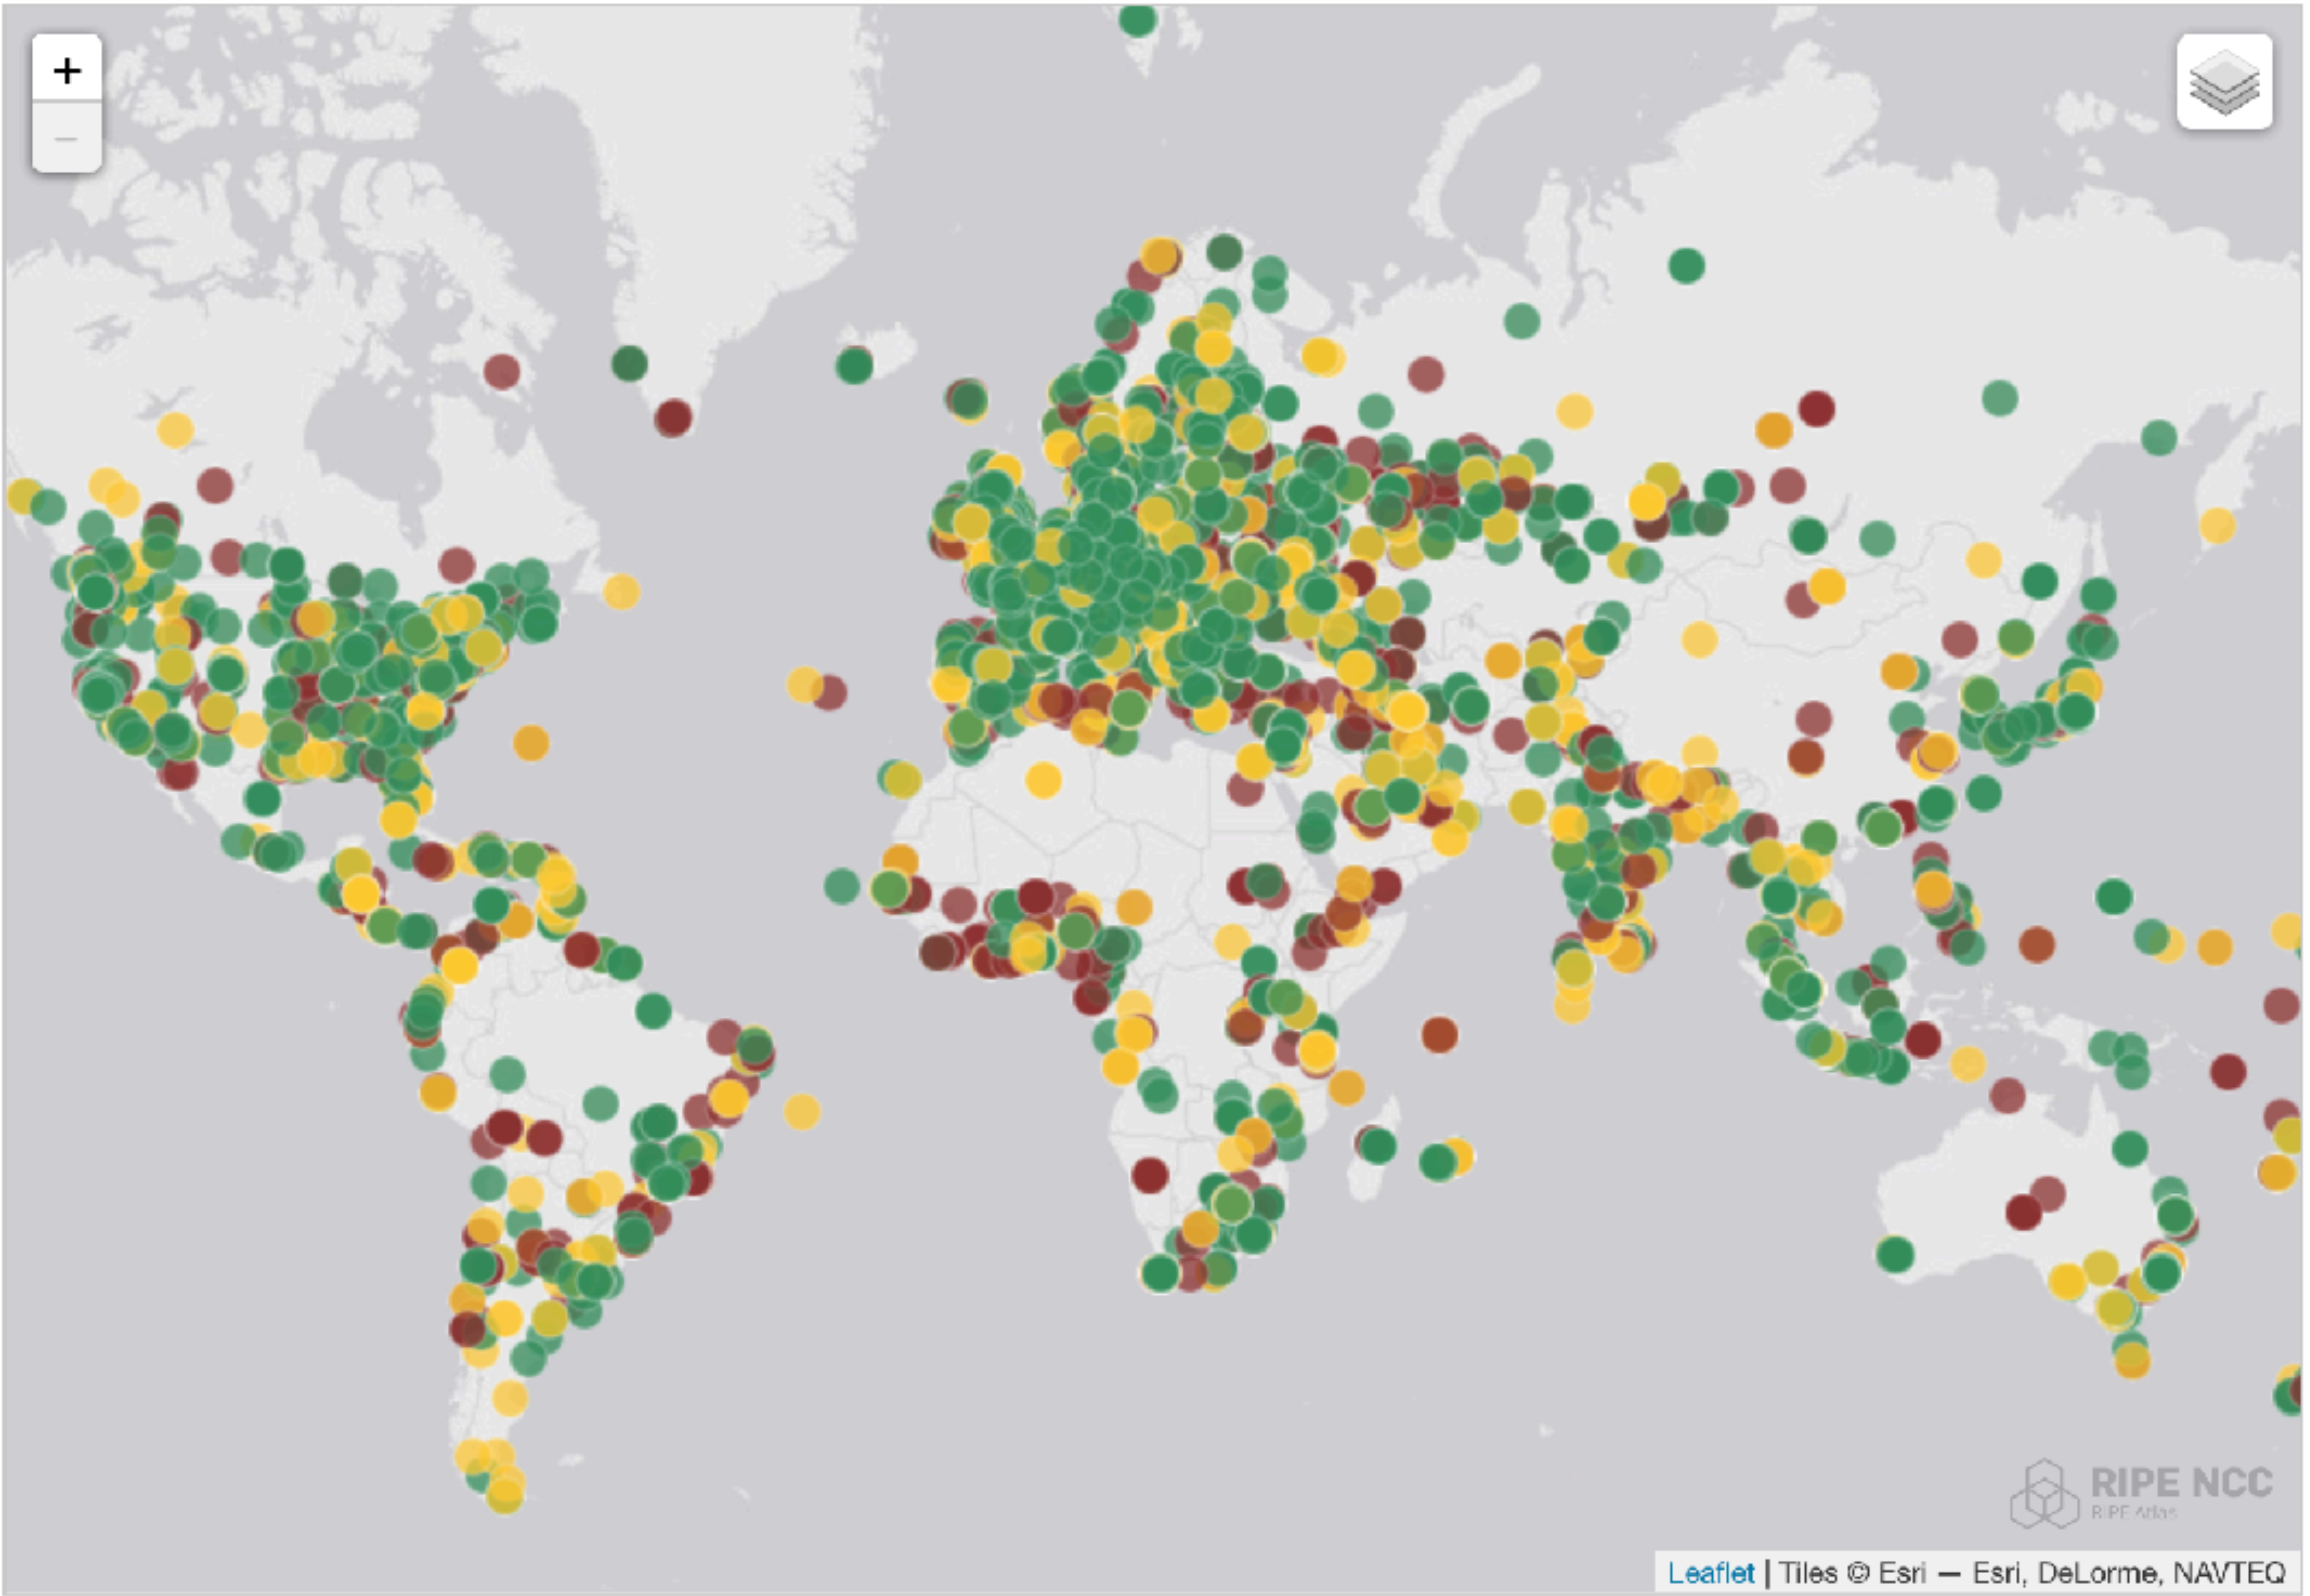

Connected: 9273 Disconnected: 3490 Abandoned: 4962

RIPE Atlas probe coverage of Tor guard and exit relay ASes in May 2016.

Atlas probe coverage	Tor guard ASes (%)	Tor exit ASes (%)
By number	50.69	52.25
By bandwidth	73.59	57.53

Here are the RIPE Atlas traceroutes we ran for the ingress side.

Ingress side traceroutes

- From RIPE Atlas probes in the five ISPs to all guard relay IP addresses used in March 2016

Here are the traceroutes we ran for our four Tor exit relay DNS resolution scenarios.

Egress side traceroutes

ISP DNS only: no traceroutes

Google DNS only: traceroutes from co-located RIPE Atlas probes to 8.8.8.8

Local DNS only: traceroutes from co-located RIPE Atlas probes to all IP addresses in our websites' delegation paths

Status quo: traceroutes from co-located RIPE Atlas probes to IP addresses of the resolvers that exit relays actually use

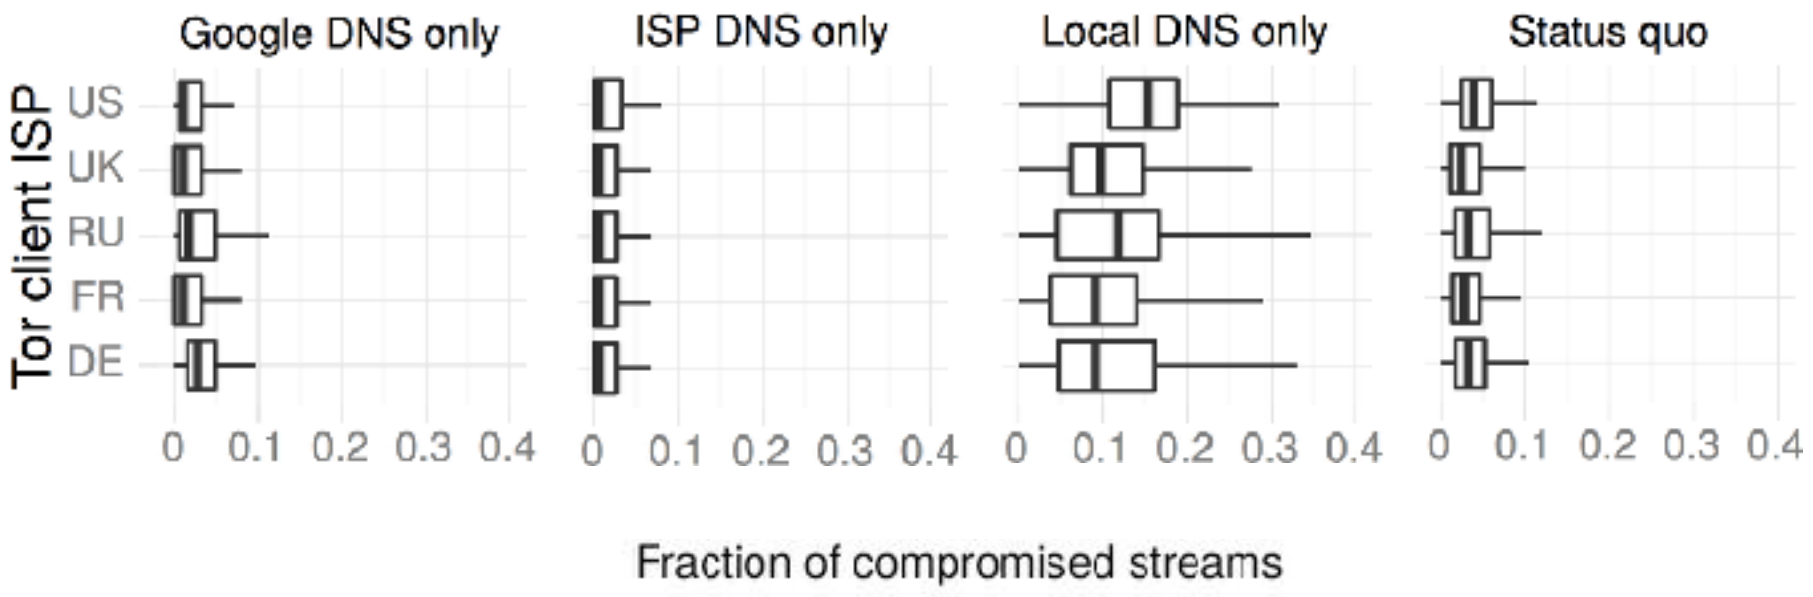

(a) The fraction of compromised streams of simulated Tor clients.

Results

RIPE Atlas allowed us to study the fraction of compromised streams for our four DNS scenarios.

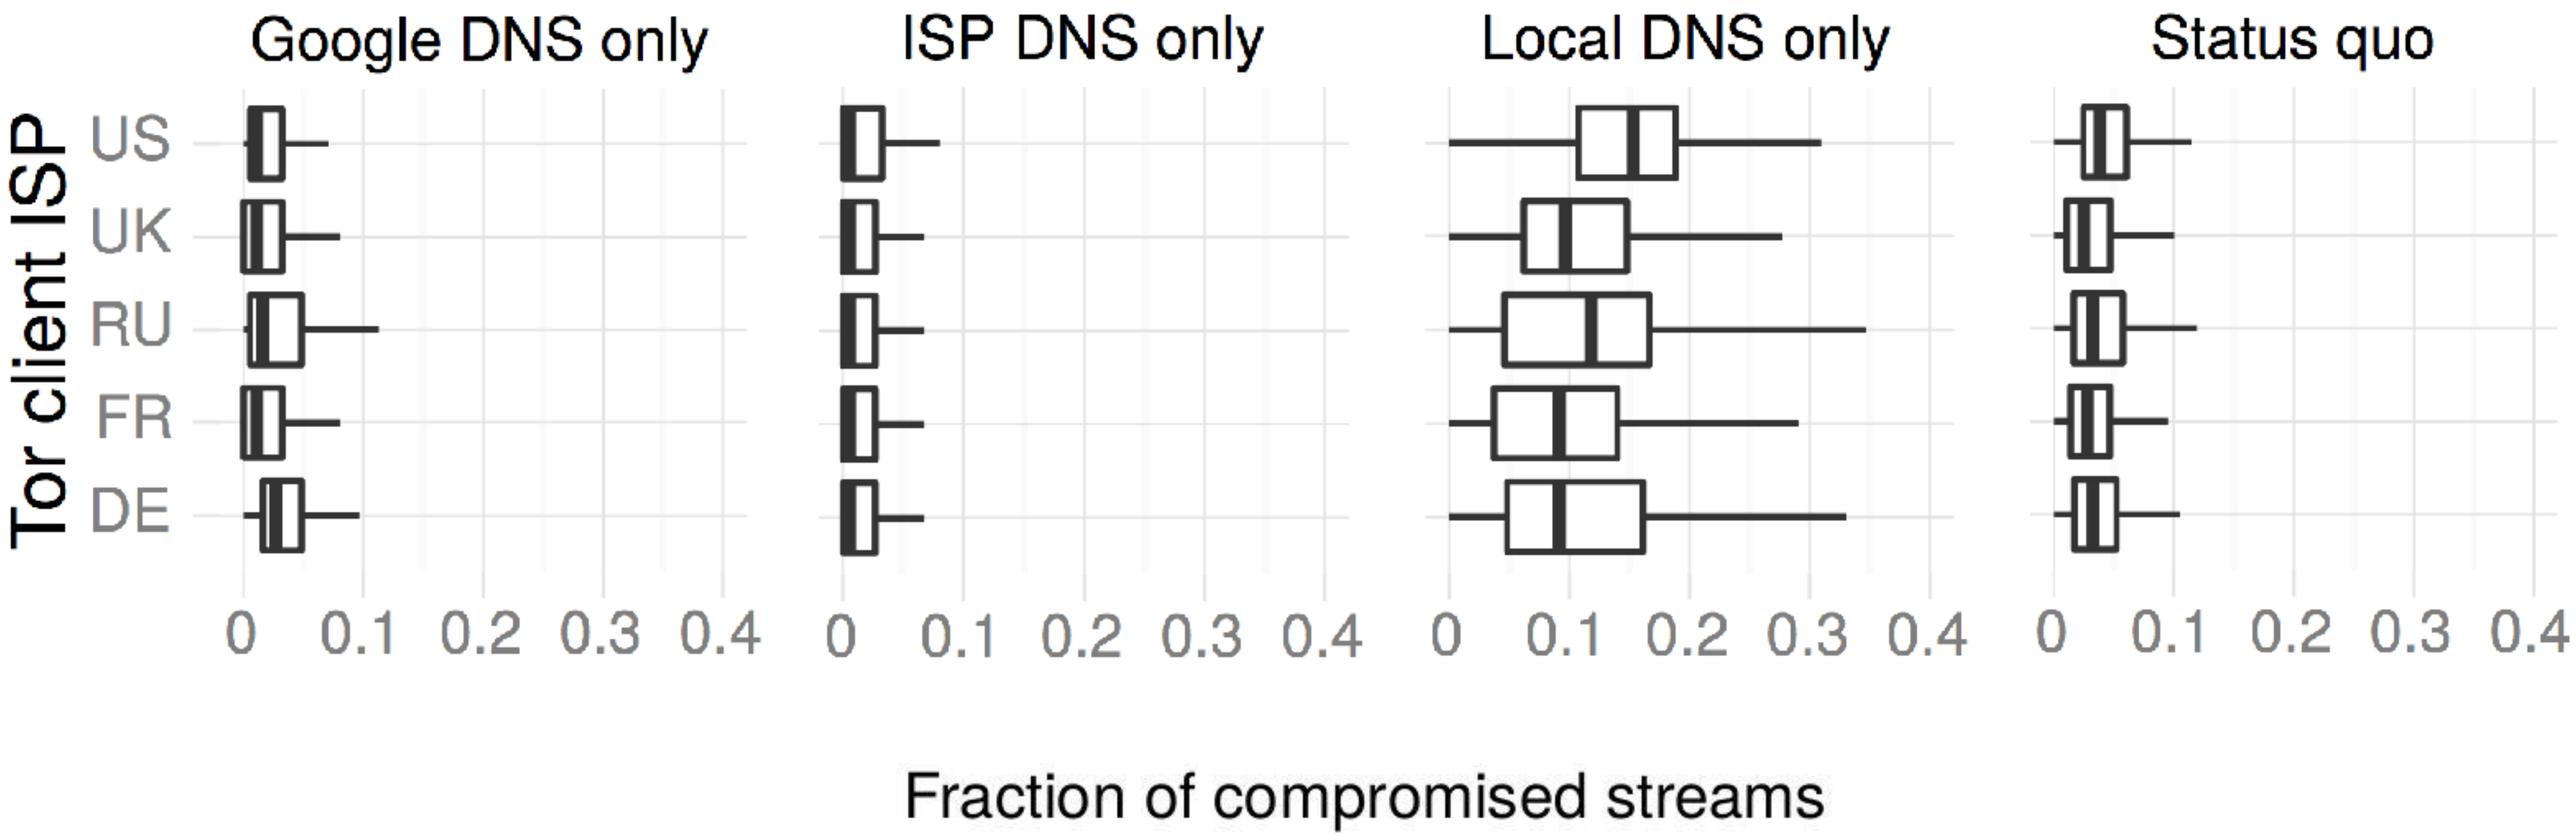

(a) The fraction of compromised streams of simulated Tor clients.

RIPE Atlas allowed us to study the fraction of compromised streams for our four DNS scenarios.

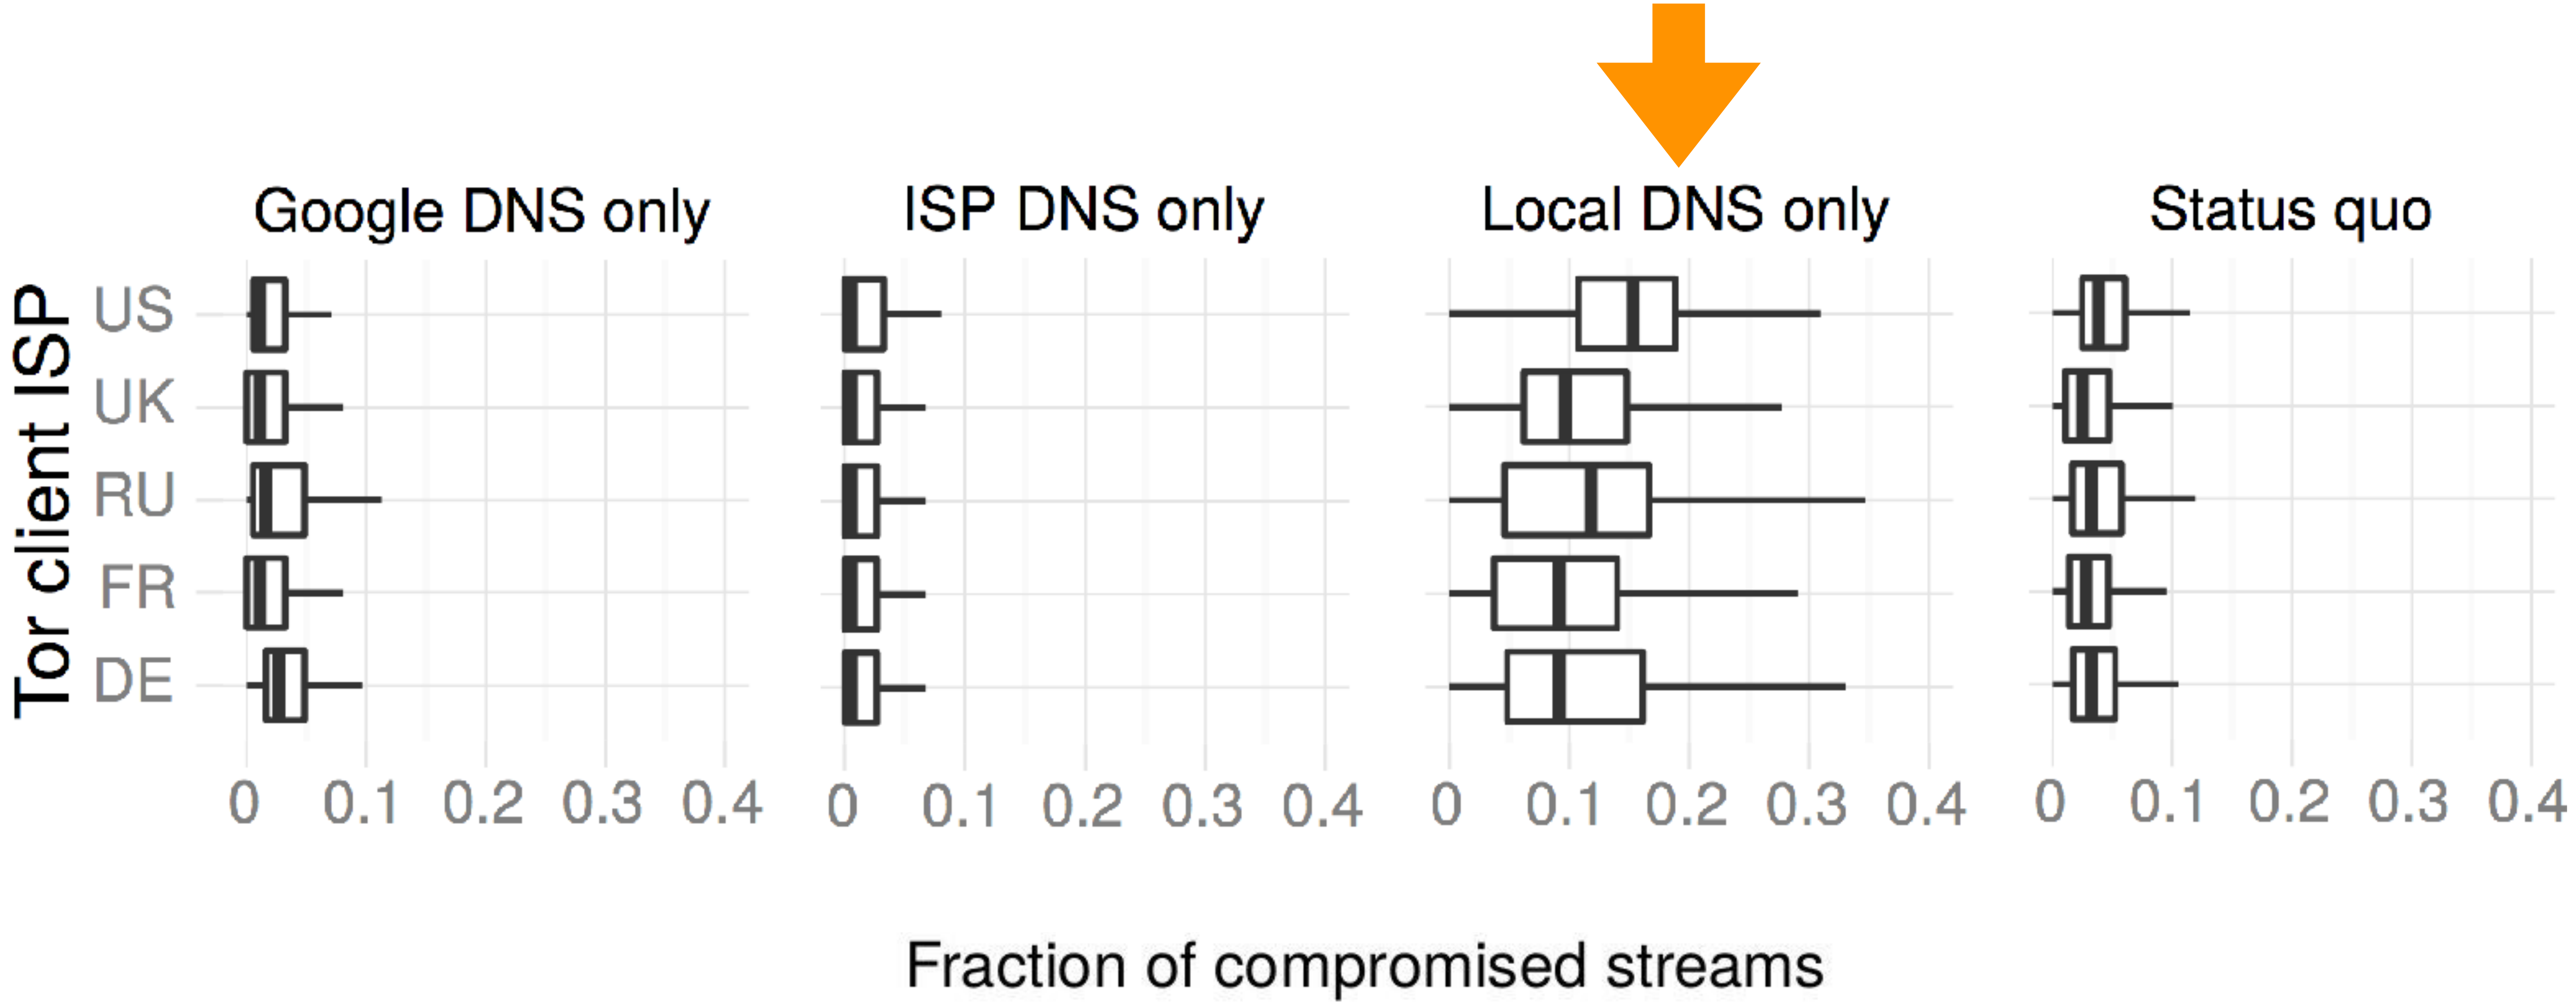

(a) The fraction of compromised streams of simulated Tor clients.

RIPE Atlas allowed us to study the fraction of compromised streams for our four DNS scenarios.

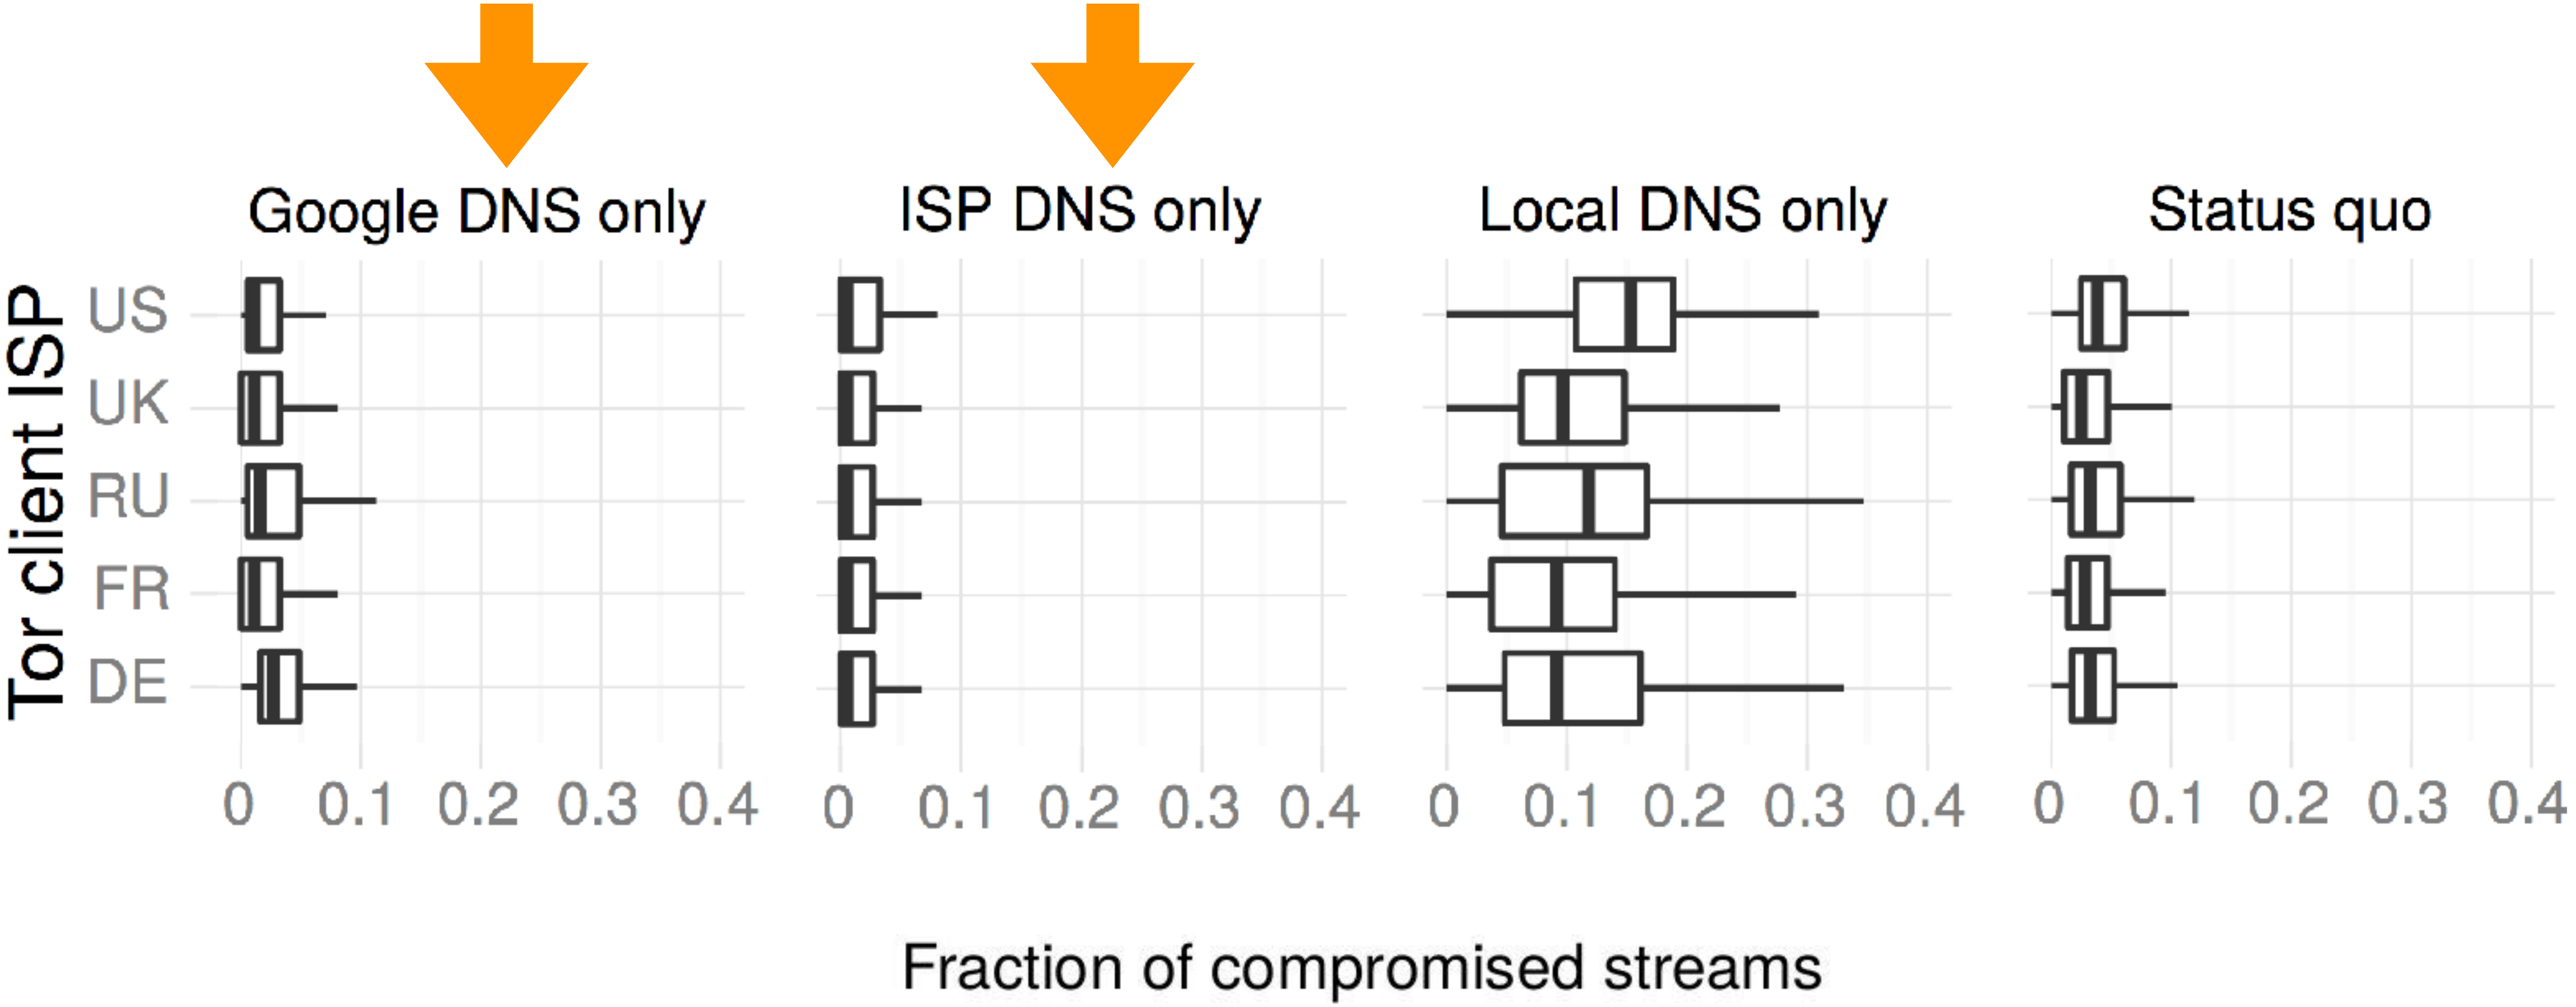

(a) The fraction of compromised streams of simulated Tor clients.

In summary, RIPE Atlas played a key role in our security paper.

In summary, RIPE Atlas played a key role in our security paper.

Our paper, data, code, and replication instructions can be found at https://nymity.ch/tor-dns/

Contact me at laurar@cs.princeton.edu.

Instructions on replicating our experiments

Benjamin Greschbach — bgre@kth.se
Tobias Pulls — tobias.pulls@kau.se
Laura M. Roberts — laurar@princeton.edu
Philipp Winter — pwinter@cs.princeton.edu
Nick Feamster — feamster@cs.princeton.edu

September 28, 2016

In this document, we provide instructions on how to replicate the results from our research paper "The Effect of DNS on Tor's Anonymity".[1] Each section discusses the replication of a specific experiment, providing both code and data necessary to replicate. Our project page is available online at https://nymity.ch/tor-dns/.

Contents