# 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 |

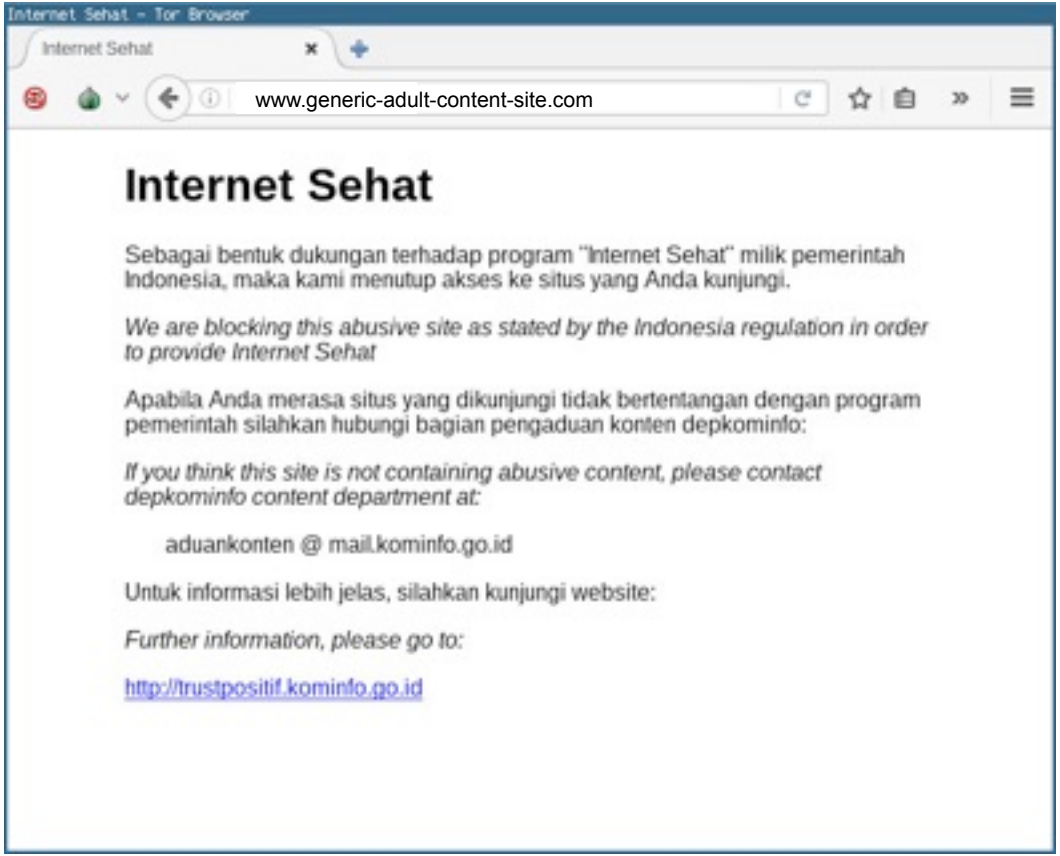
Internet Sehat - Tor Browser
Internet Sehat
www.generic-adult-content-site.com
Internet Sehat
Sebagai bentuk dukungan terhadap program "Internet Sehat" milik pemerintah Indonesia, maka kami menutup akses ke situs yang Anda kunjungi.
We are blocking this abusive site as stated by the Indonesia regulation in order to provide Internet Sehat
Apabila Anda merasa situs yang dikunjungi tidak bertentangan dengan program pemerintah silahkan hubungi bagian pengaduan konten depkominfo:
If you think this site is not containing abusive content, please contact depkominfo content department at:
aduankonten @ mail.kominfo.go.id
Untuk informasi lebih jelas, silahkan kunjungi website:
Further information, please go to:
http://trustpositif.kominfo.go.id

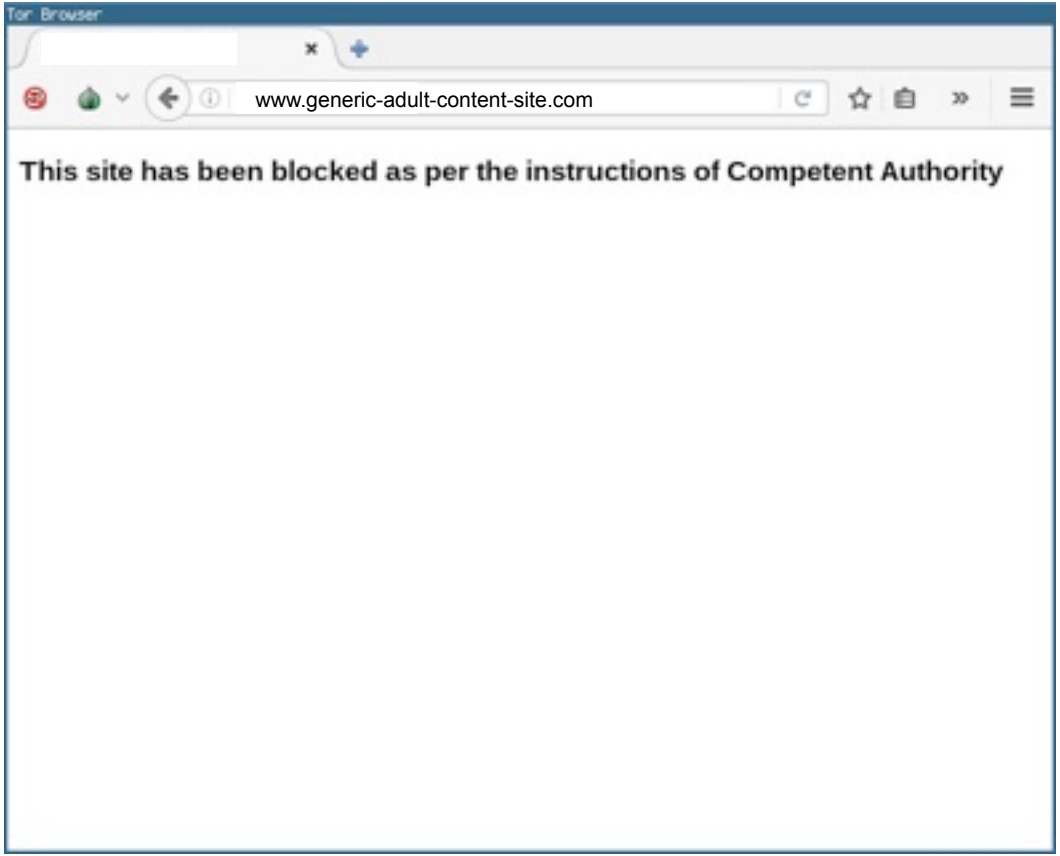
Tor Browser
www.generic-adult-content-site.com
This site has been blocked as per the instructions of Competent Authority

# How is DNS handled in Tor?

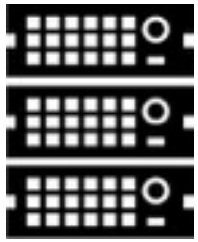

DNS resolver

Tor client

Guard

Middle

Exit

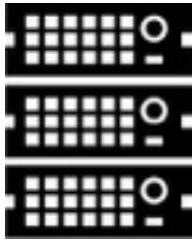

example.com

# How is DNS handled in Tor?

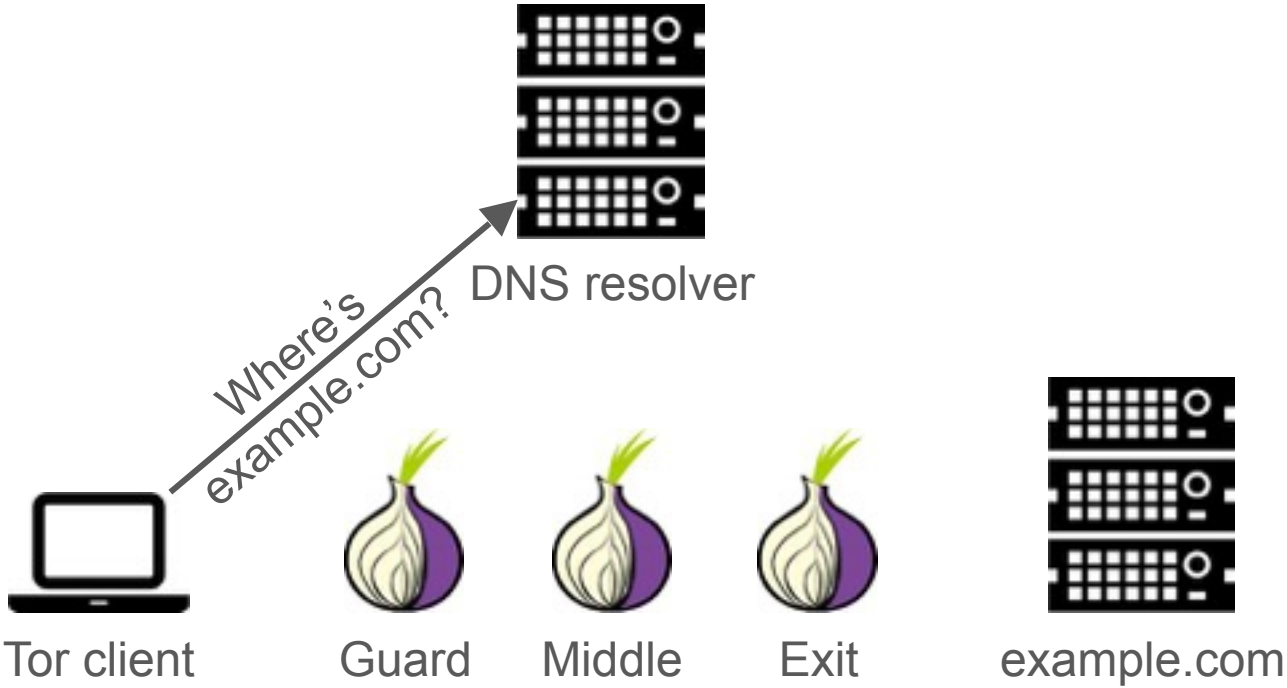

# How is DNS handled in Tor?

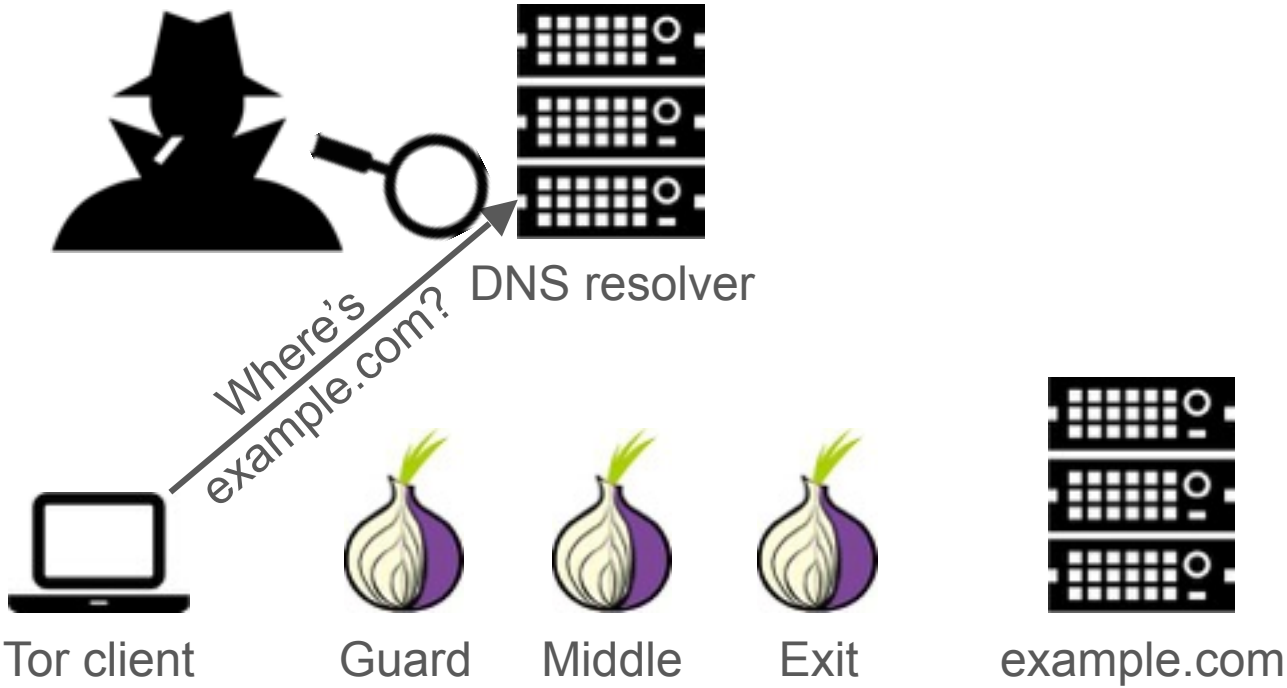

# Exit relays perform DNS resolution.

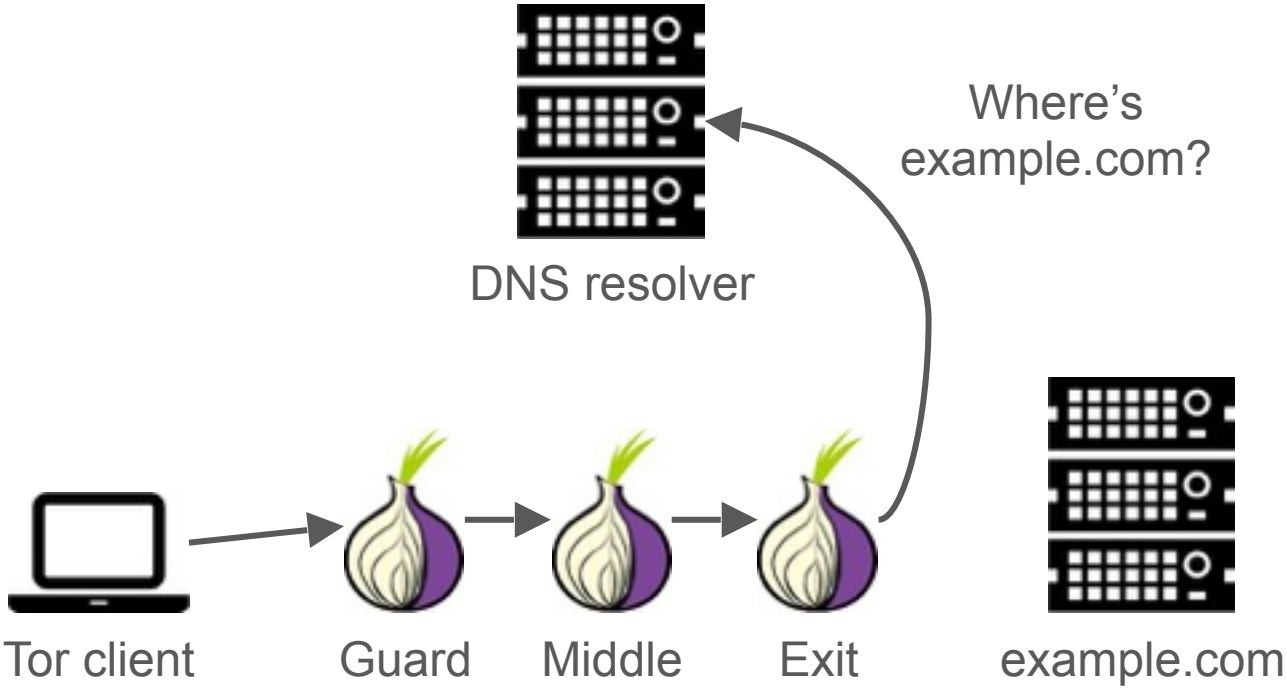

# Research Questions

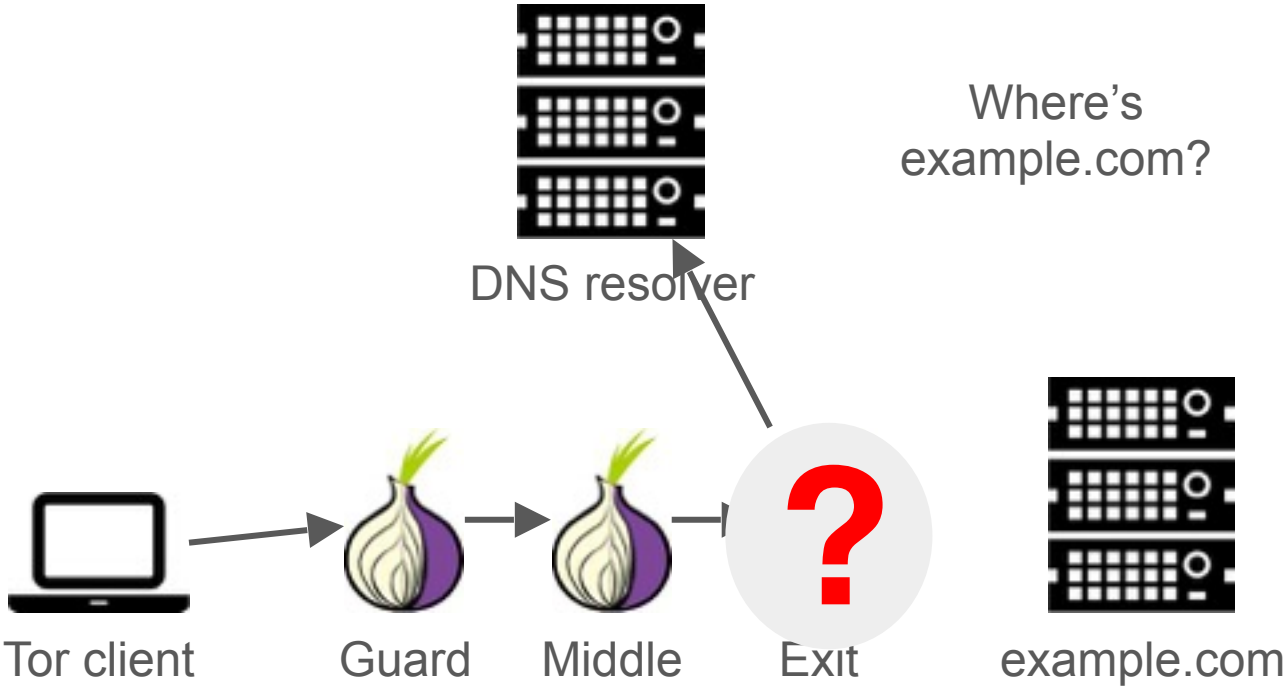

# Research Questions

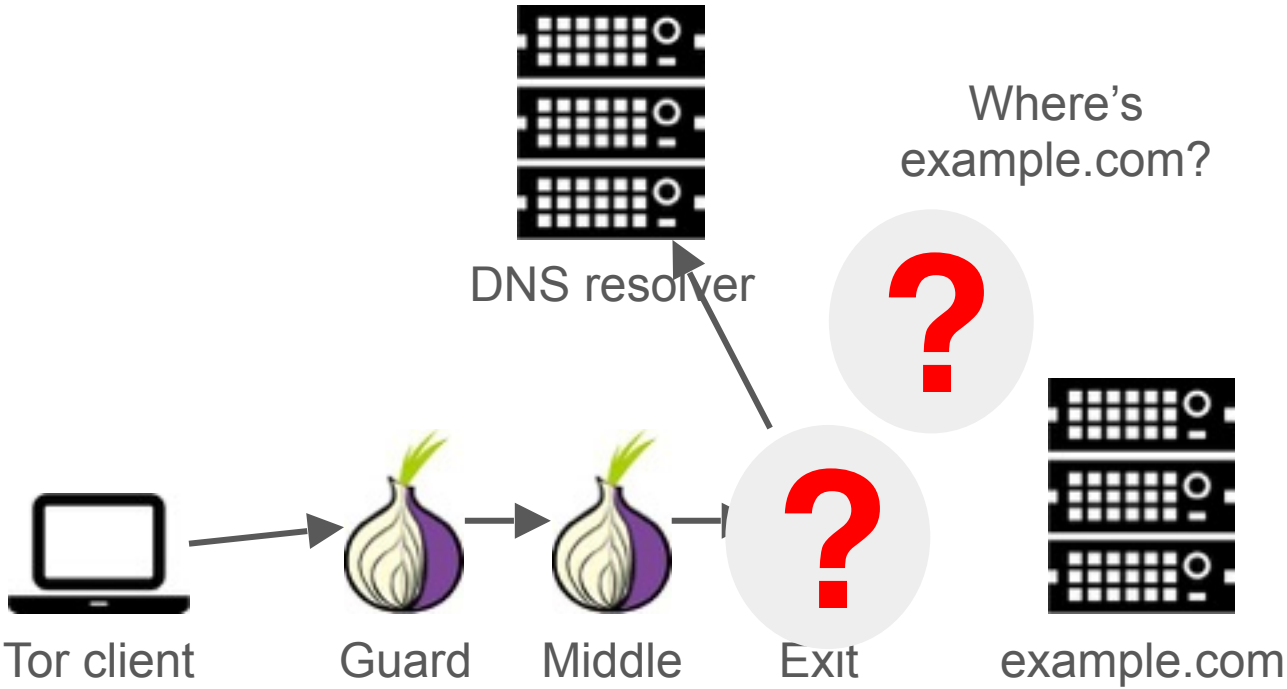

# How DNS can be used to compromise Tor.

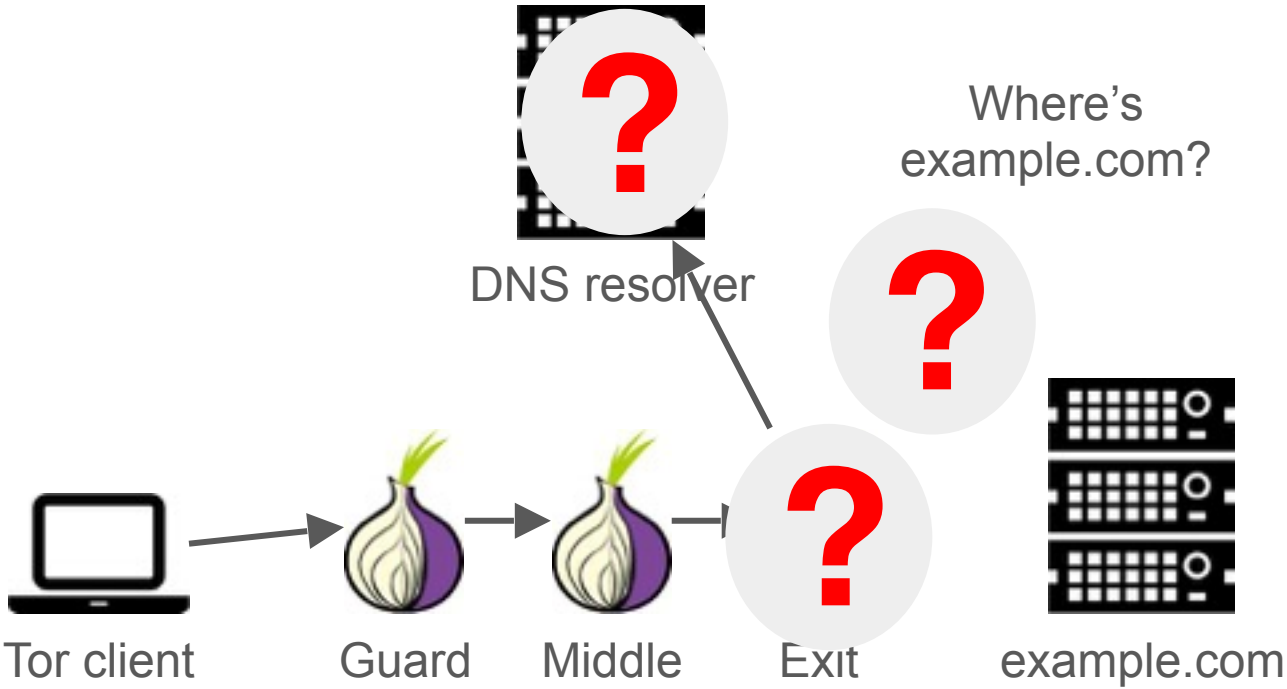

# How exposed are DNS queries?

# How exposed are DNS queries?

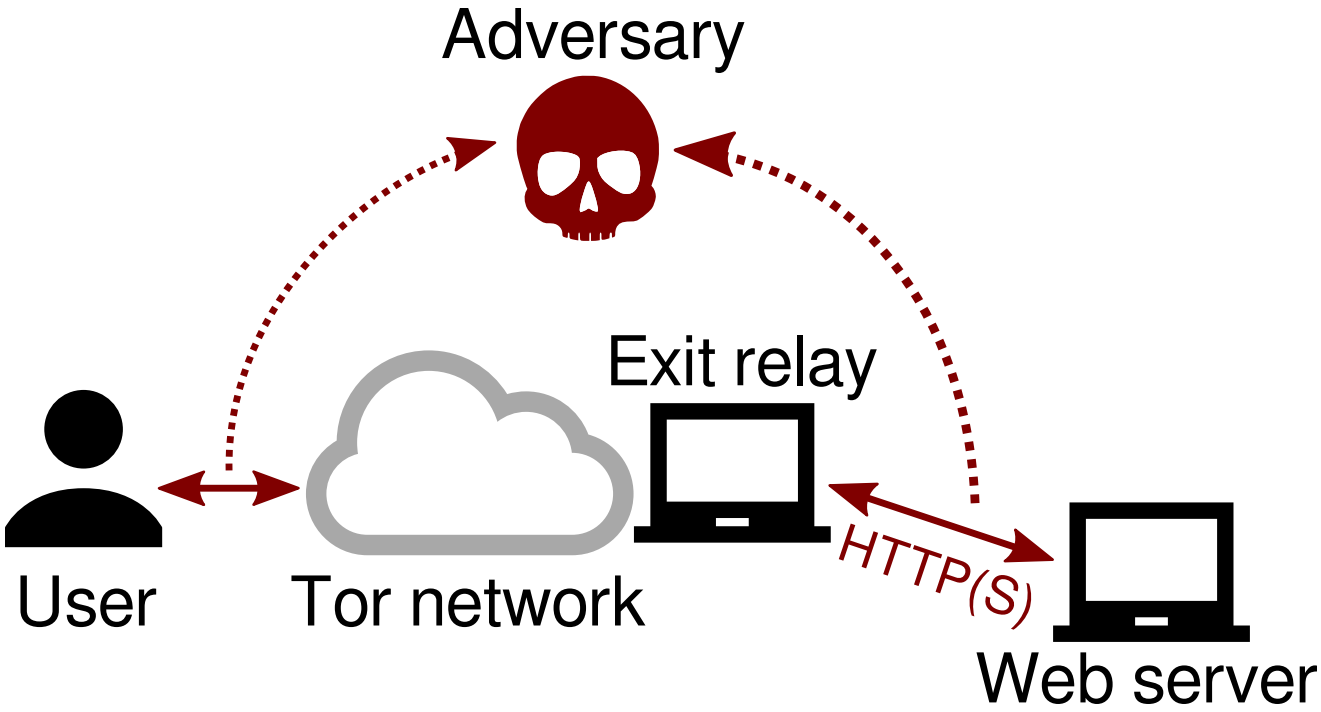

# How exposed are DNS queries?

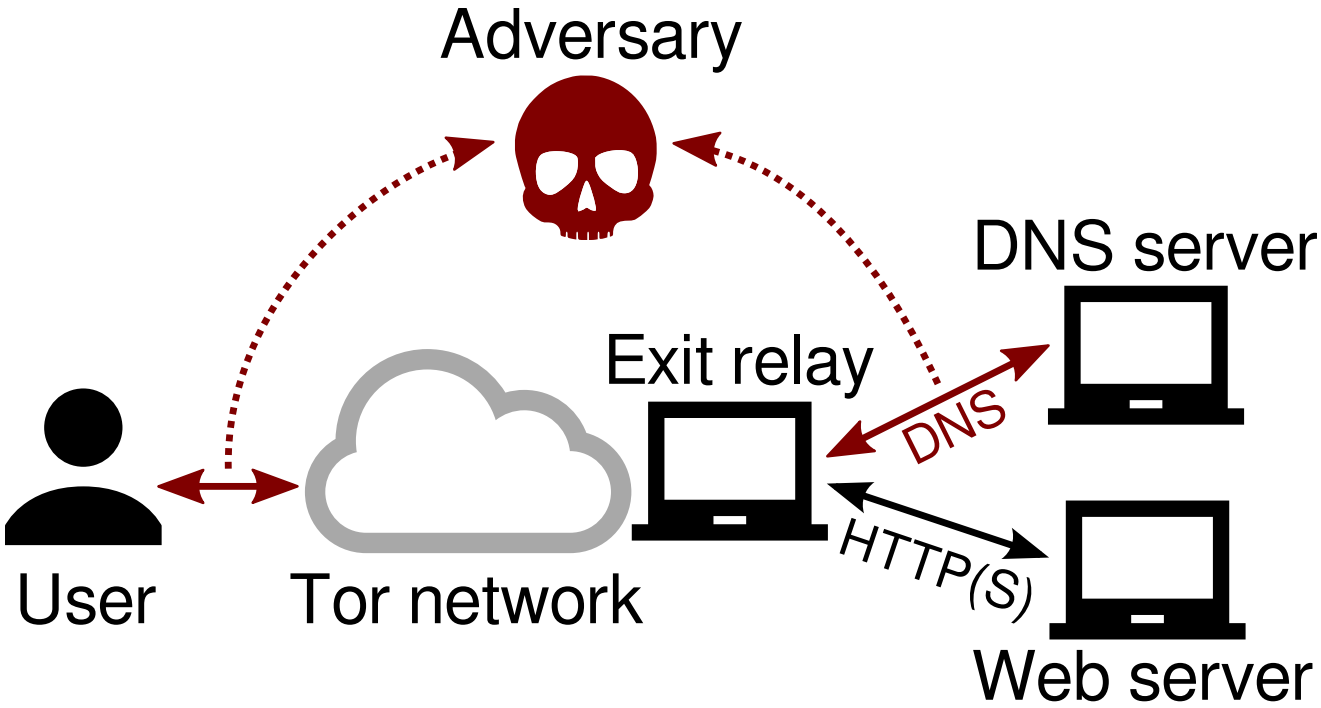

# How exposed are DNS queries?

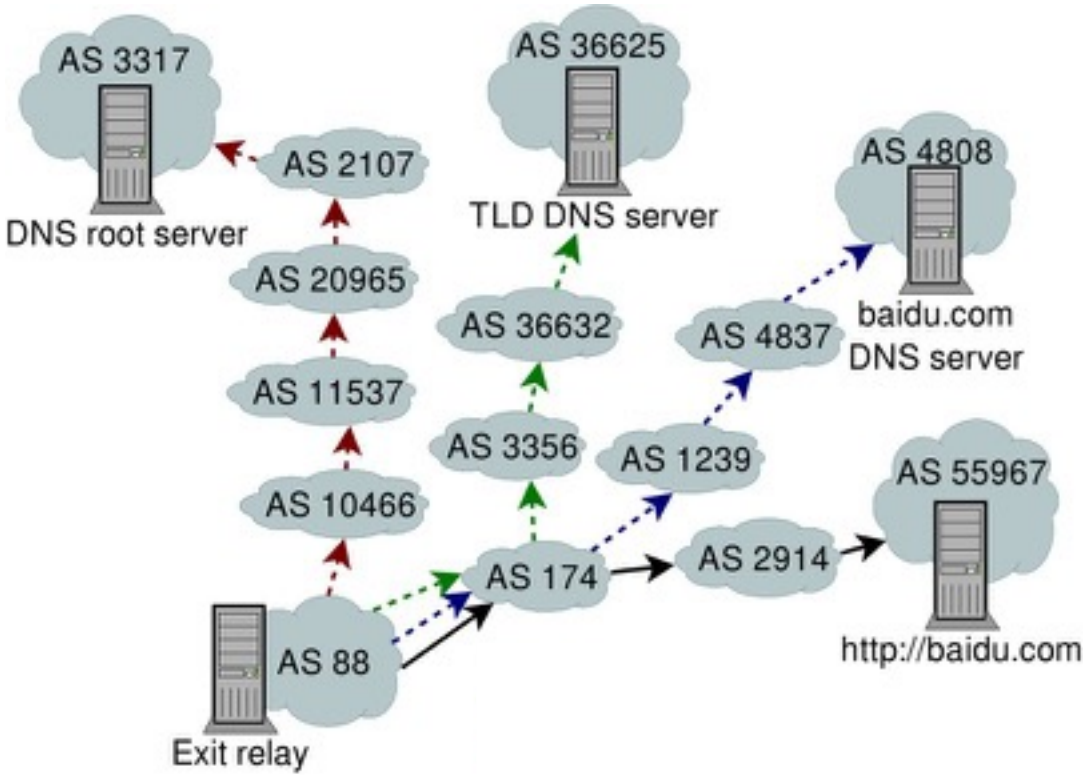

# How exposed are DNS queries?

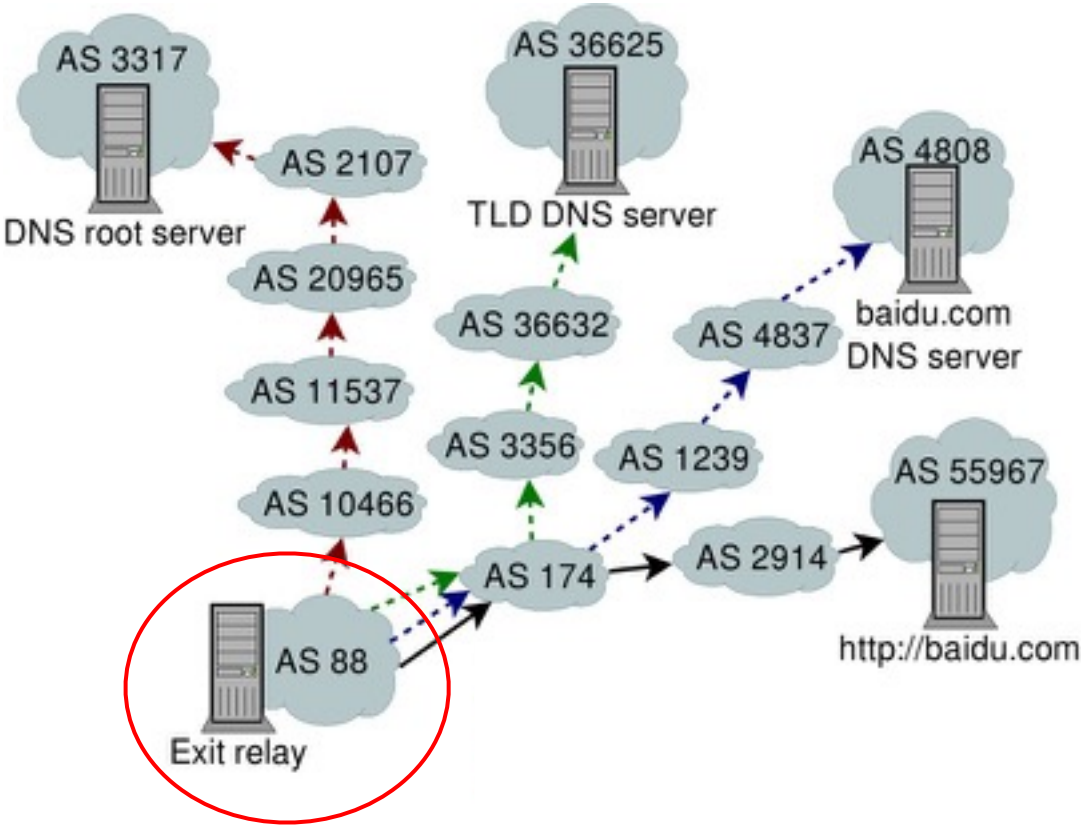

# How exposed are DNS queries?

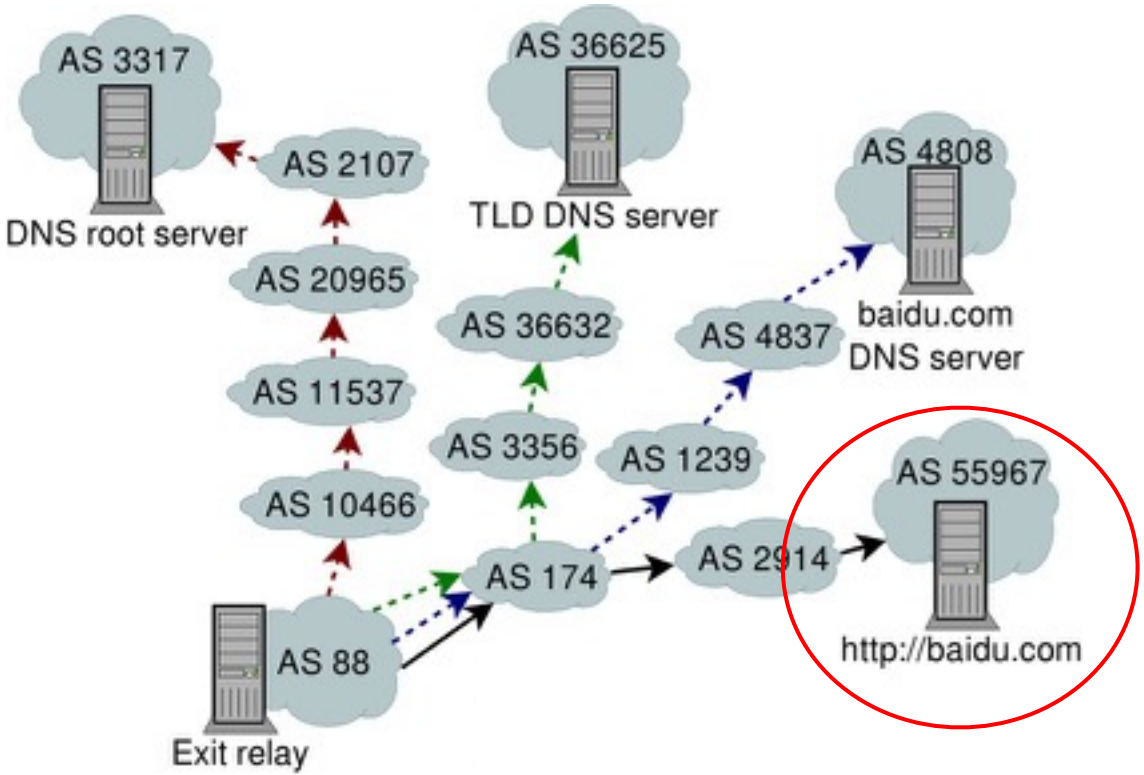

# How exposed are DNS queries?

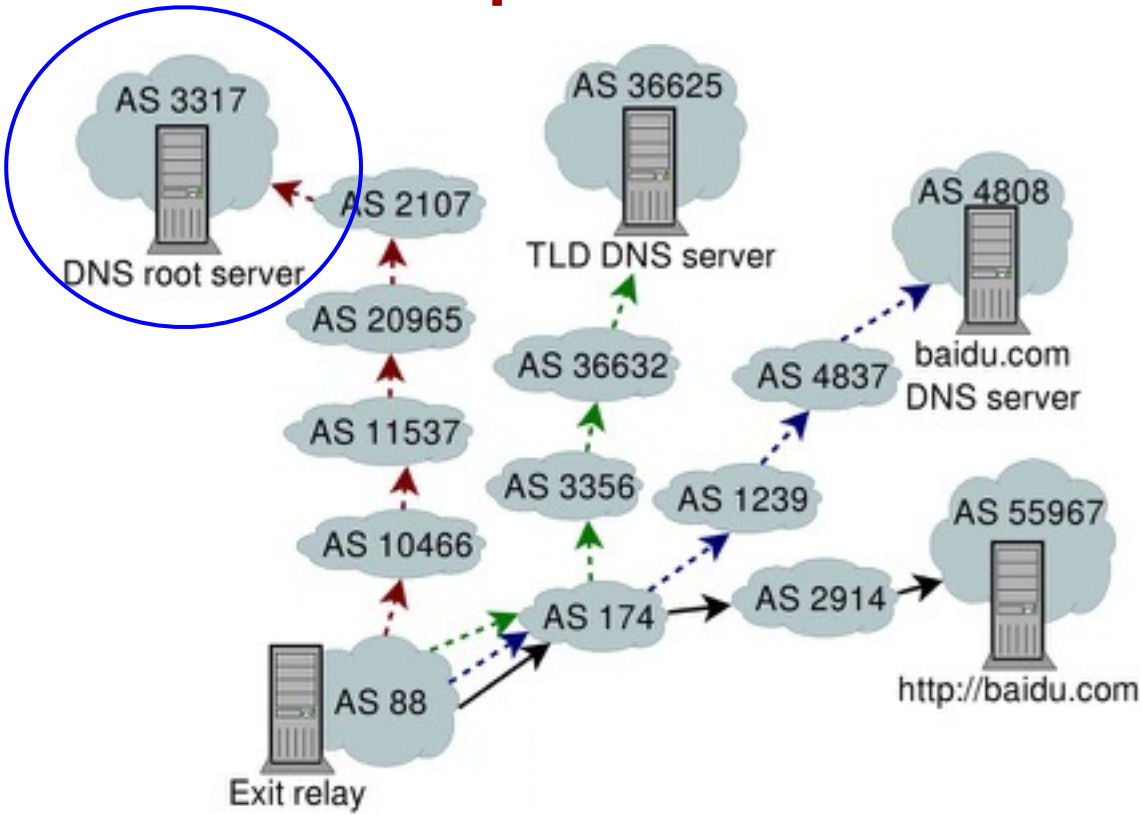

# How exposed are DNS queries?

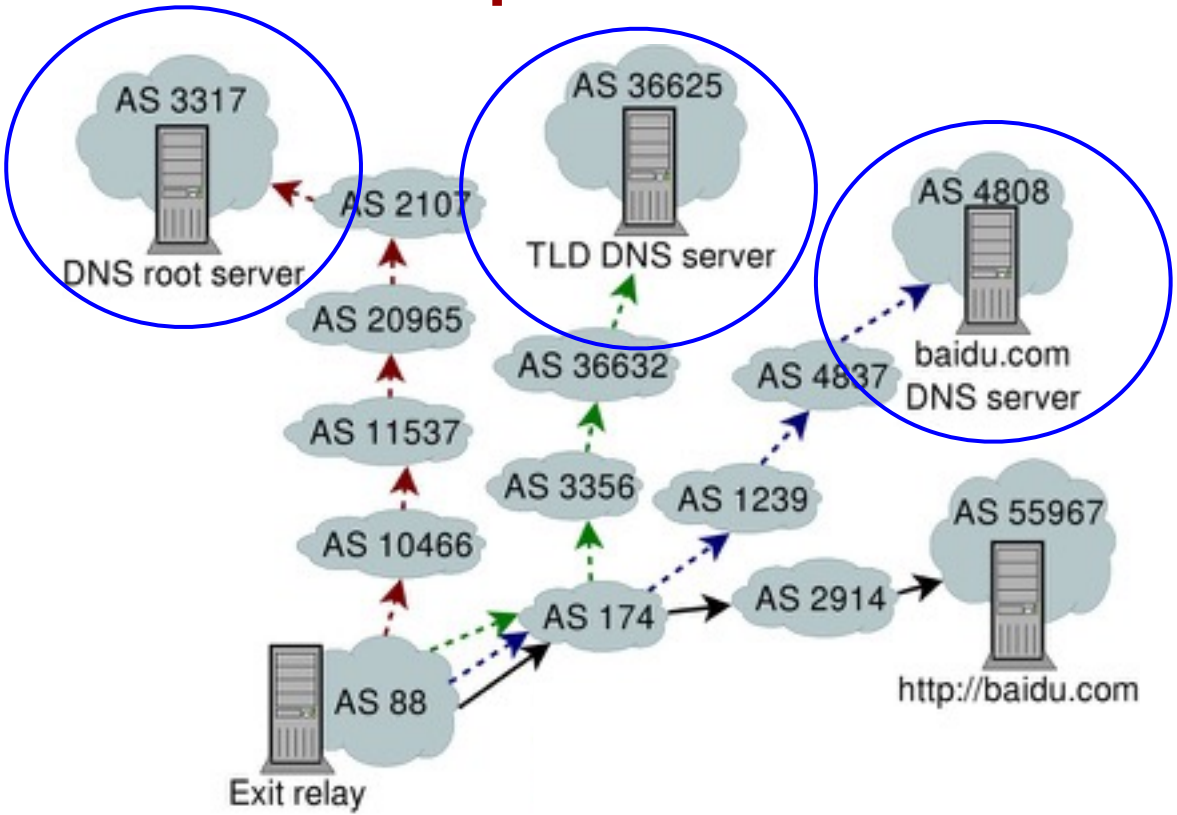

# How exposed are DNS queries?

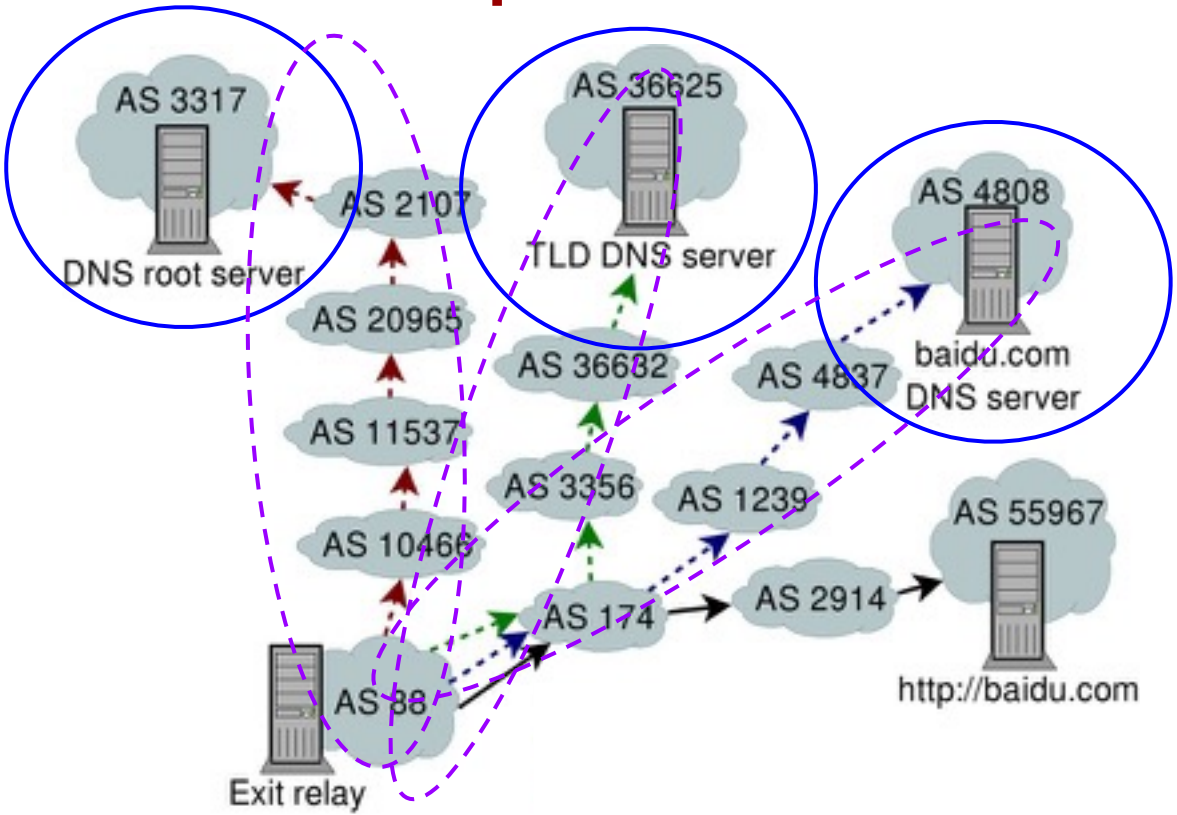

# How exposed are DNS queries?

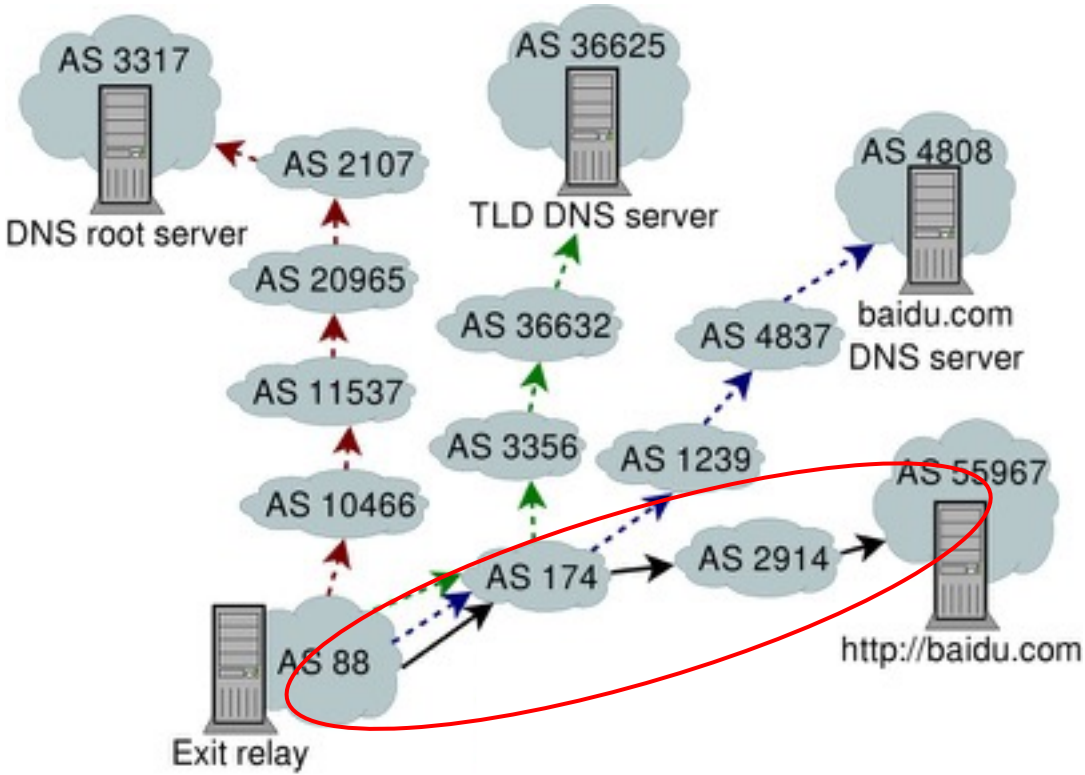

# How exposed are DNS queries?

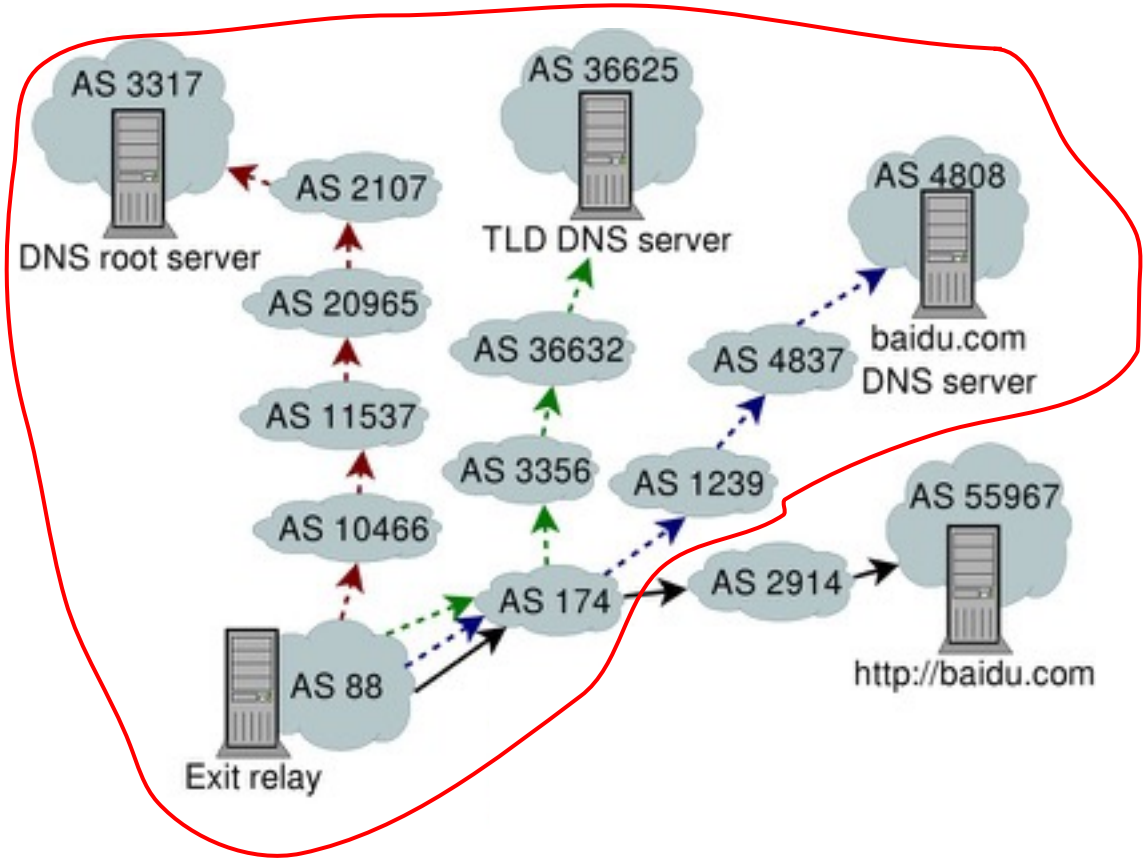

# How exposed are DNS queries?

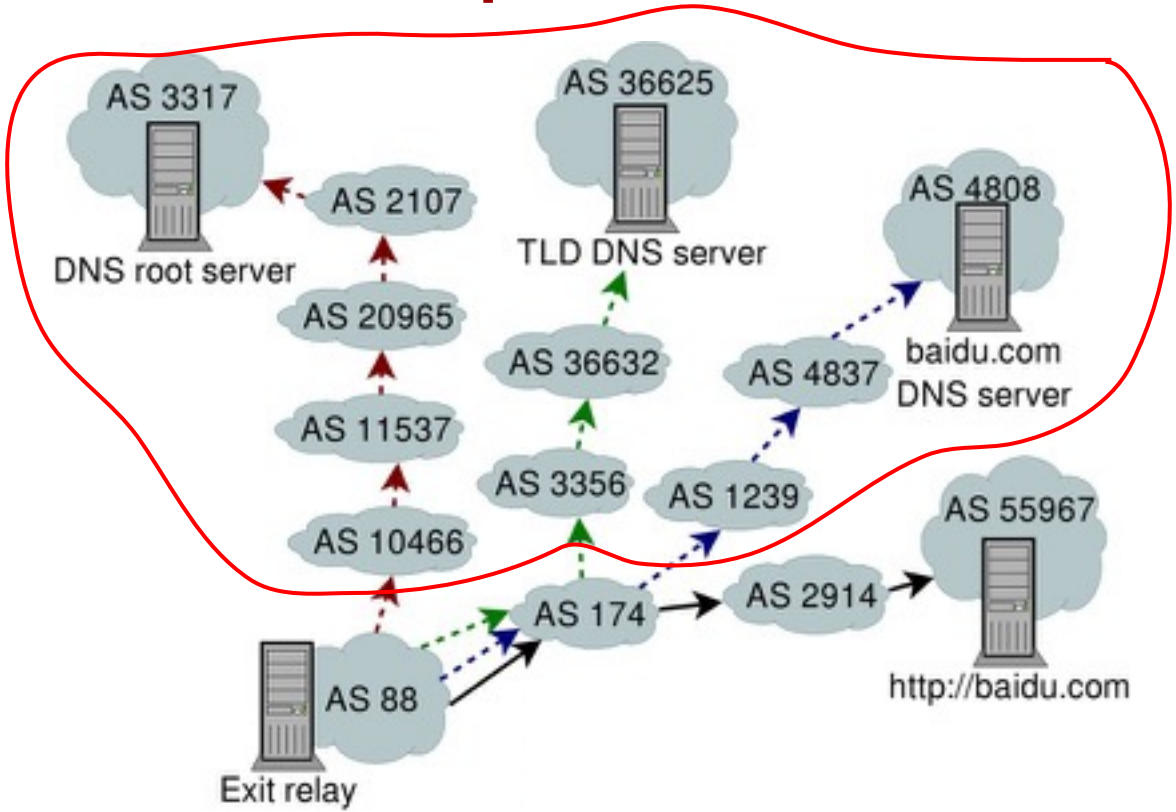

# DNS traffic traverses ASes that are not otherwise traversed by TCP traffic.

For half of all of the Alexa Top 1,000 websites, DNS-only ASes account for 57% or more of all traversed ASes

# What resolvers do exit relays use?

# What resolvers do exit relays use?

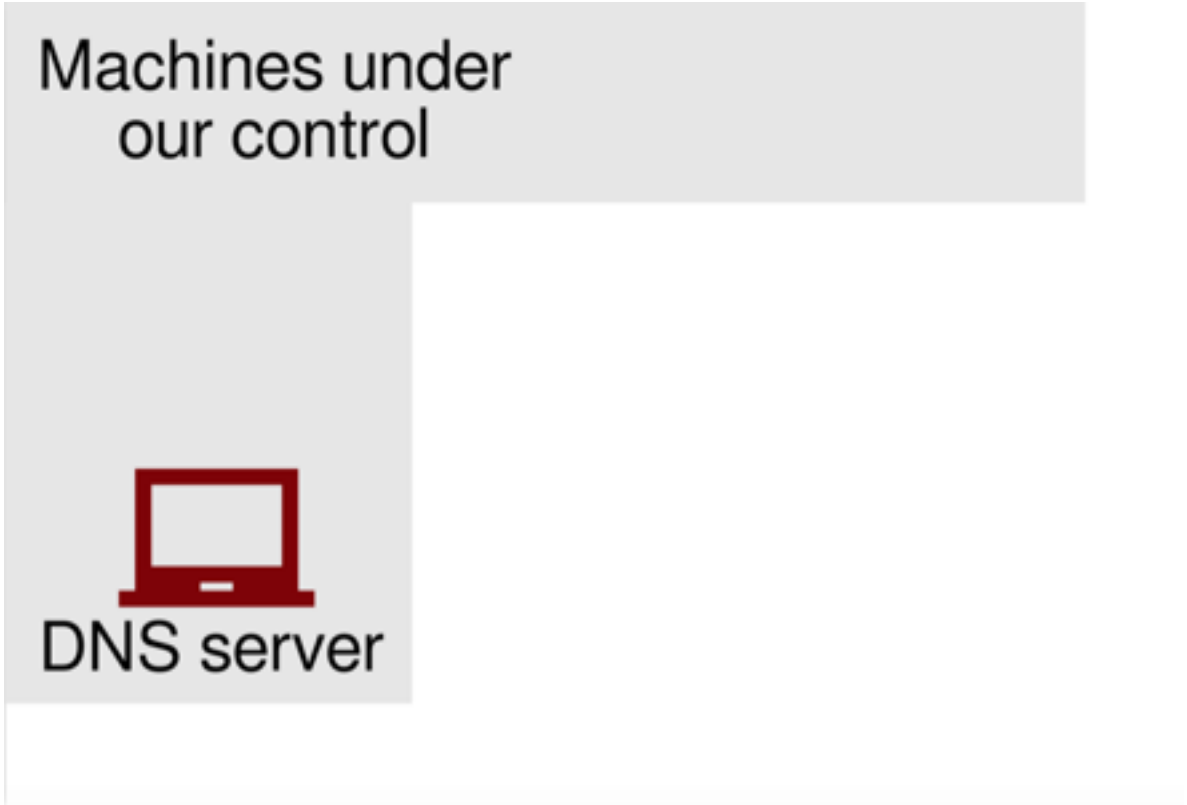

# What resolvers do exit relays use?

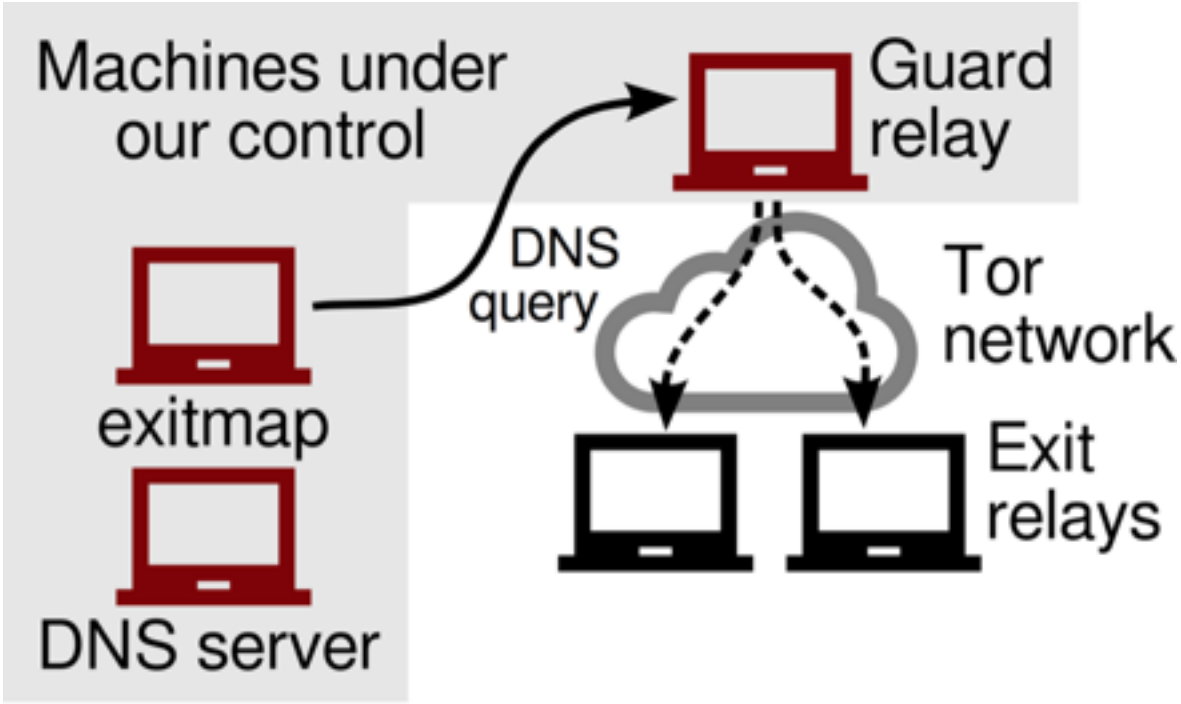

# What resolvers do exit relays use?

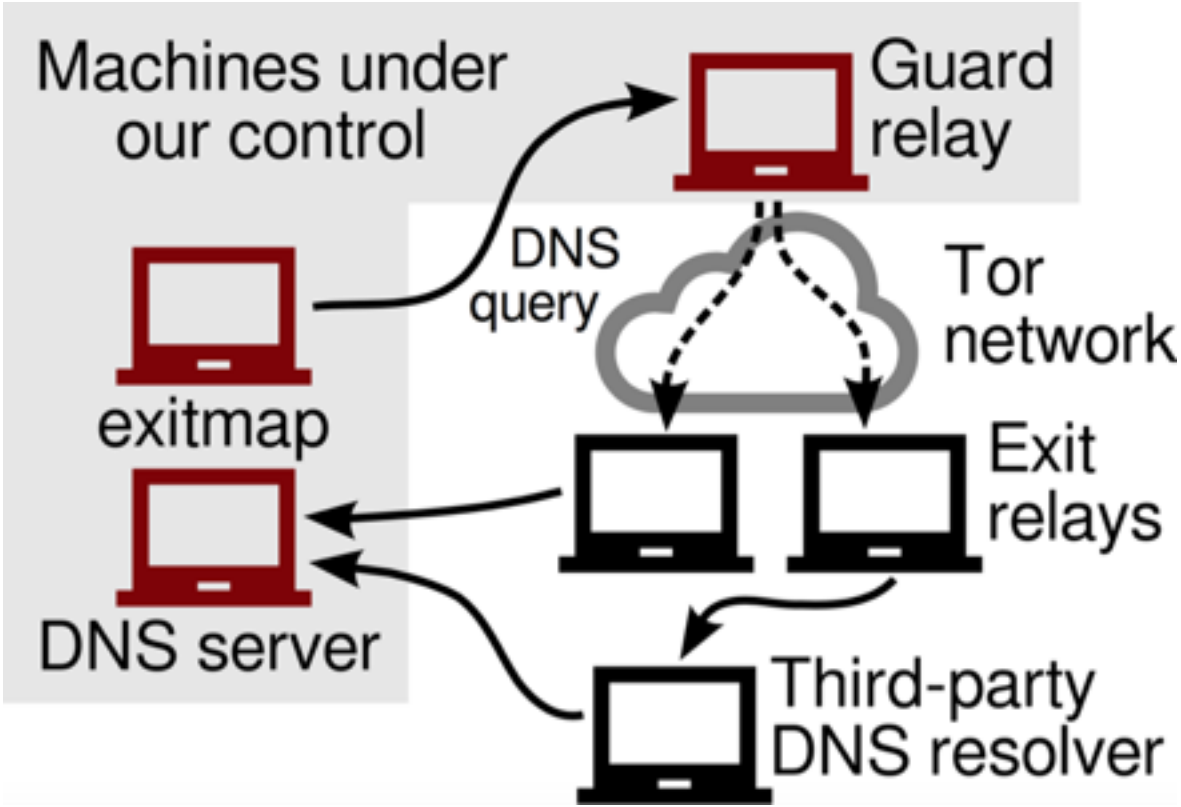

# What resolvers do exit relays use?

| Resolver | Min (%) | Max (%) | Median (%) |
|---|---|---|---|
| Google | 23.57 | 42.33 | 32.84 |
| Local | 7.71 | 15.95 | 11.56 |
| OVH | 1.96 | 14.13 | 6.57 |
| OpenDNS | 0.05 | 5.62 | 0.76 |

Percentage of observed DNS queries

# What resolvers do exit relays use?

| Resolver | Min (%) | Max (%) | Median (%) |
|---|---|---|---|
| Google | 23.57 | 42.33 | 32.84 |
| Local | 7.71 | 15.95 | 11.56 |
| OVH | 1.96 | 14.13 | 6.57 |
| OpenDNS | 0.05 | 5.62 | 0.76 |

Percentage of observed DNS queries

# What resolvers do exit relays use?

| Resolver | Min (%) | Max (%) | Median (%) |
|---|---|---|---|
| Google | 23.57 | 42.33 | 32.84 |
| Local | 7.71 | 15.95 | 11.56 |
| OVH | 1.96 | 14.13 | 6.57 |
| OpenDNS | 0.05 | 5.62 | 0.76 |

Percentage of observed DNS queries

# How can an attacker leverage DNS?

# How can an attacker leverage DNS?

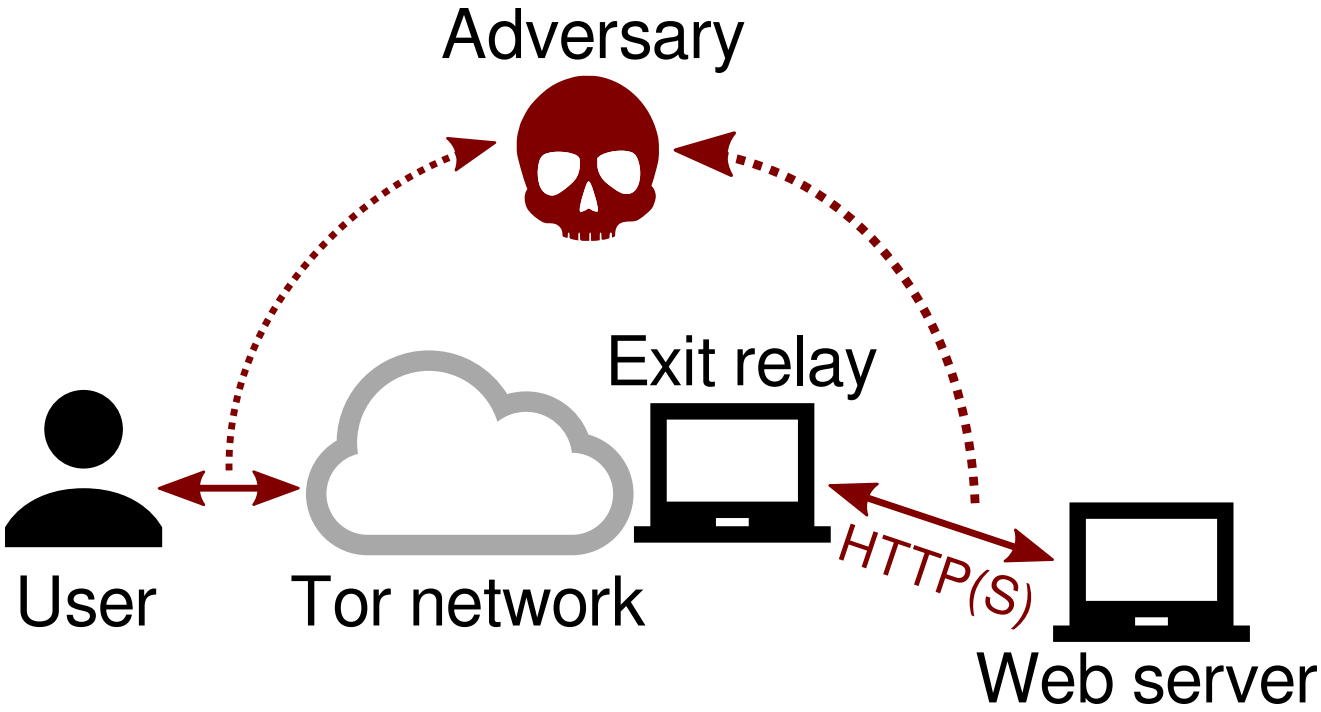

# How can an attacker leverage DNS?

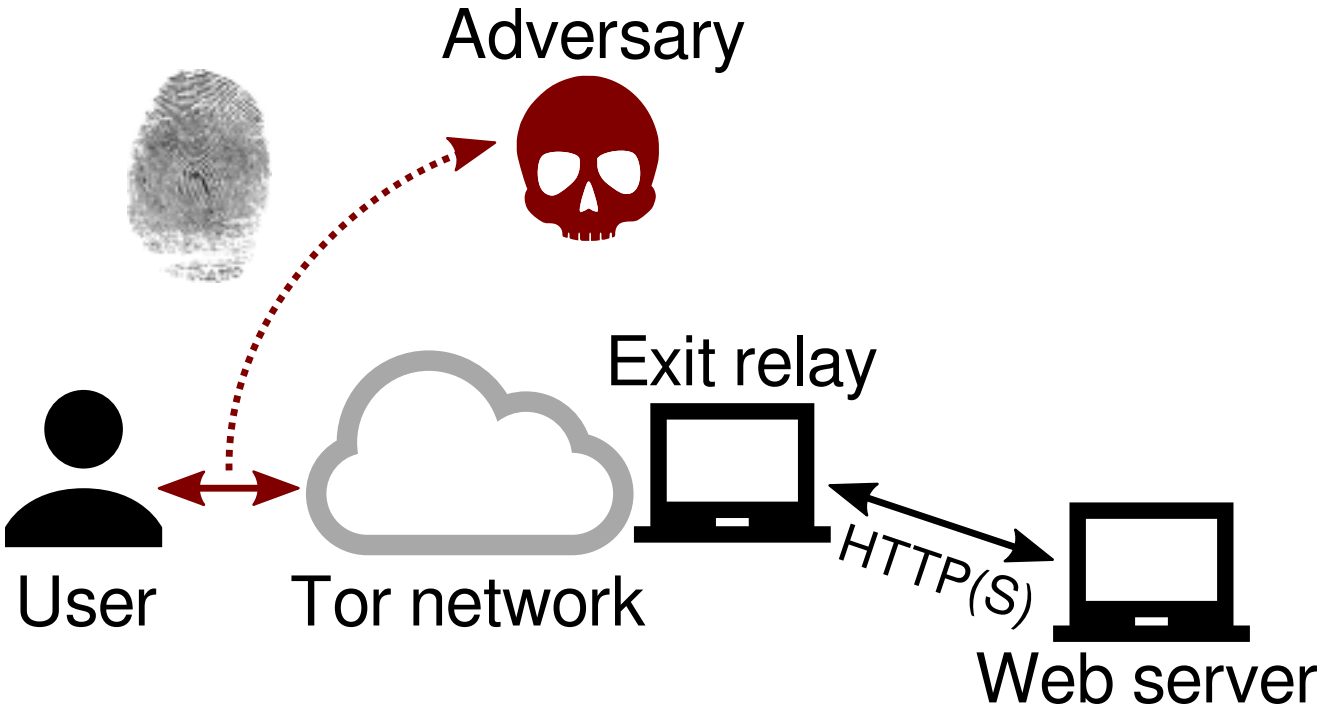

# How can an attacker leverage DNS?

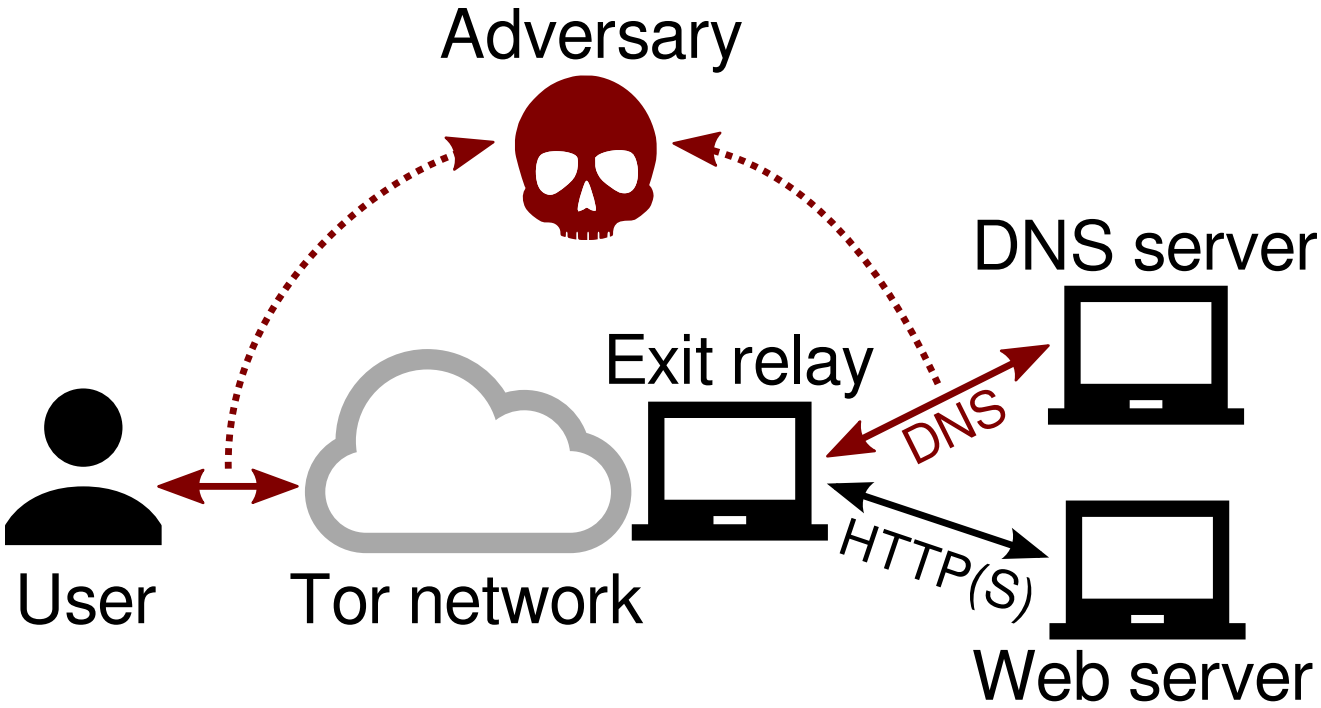

# How can an attacker leverage DNS?

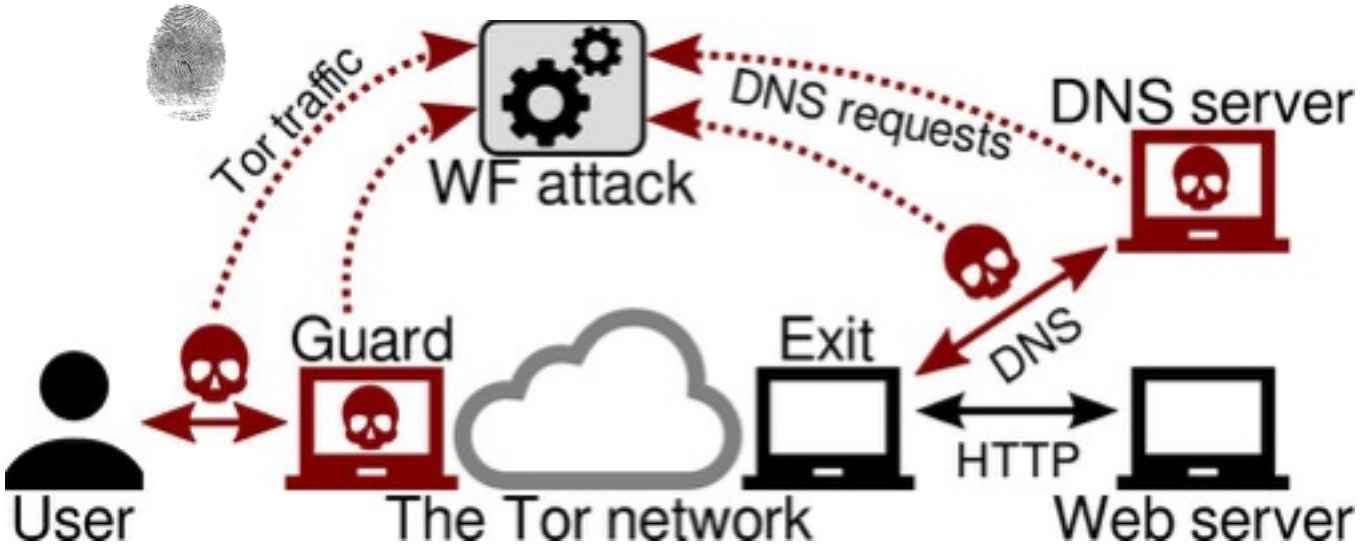

# Attacker augments website fingerprinting attack with DNS data

- We extended Wang et al.'s Wa-kNN classifier (USENIX Security'14)
- Close-the-world attack
- High precision attack
    - Accepts Wa-kNN's website classification only if that website was observed in DNS traffic
- Our attacks are very precise for unpopular websites

# Our attacks at Internet-scale

- Place Tor clients in top five Tor usage countries
- Simulate clients' online behavior
  - Cf. Johnson et al. CCS'13
- Simulate Tor clients' path selection
  - TorPS (github.com/torps/torps)
- Run traceroutes client →guard and exit → destination
  - Use RIPE Atlas!
- Check for overlapping autonomous systems
  - Set intersection

# RIPE Atlas probes

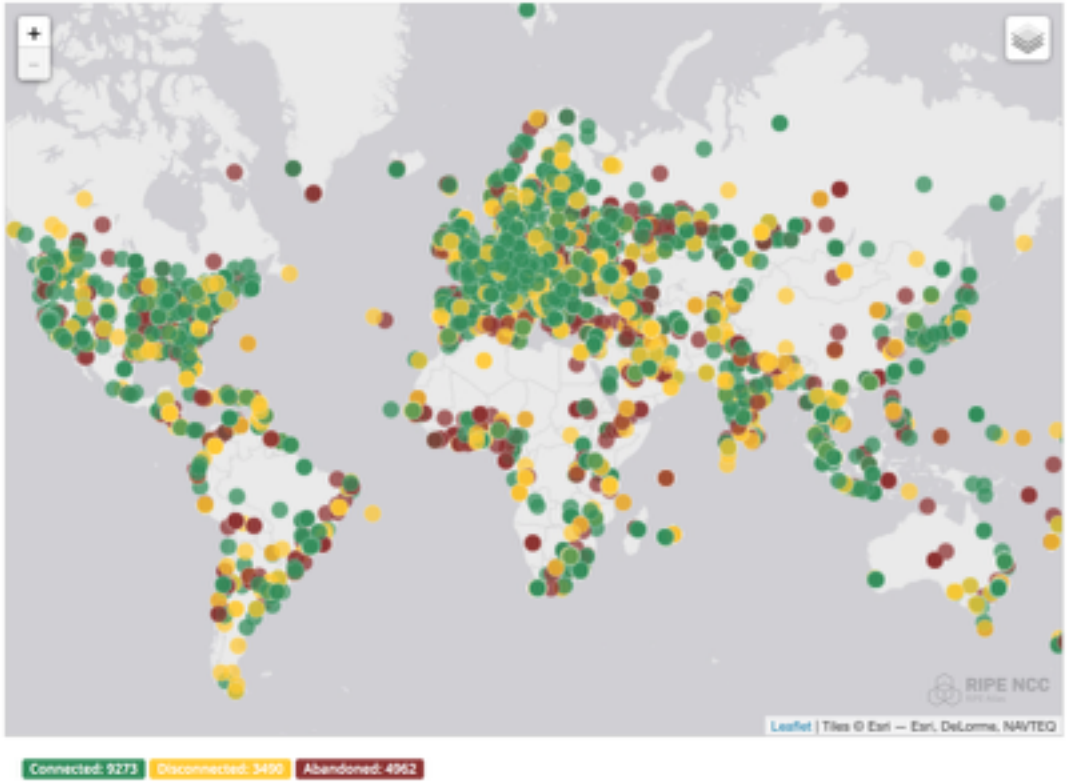

# Analyzed four Tor exit relay DNS set-up scenarios

What if all Tor exit relays were set up to use their ISPs' resolvers?

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

What if all Tor exit relays were set up to do their own DNS resolution?

What if all Tor exit relays were set up as they currently are (status quo)?

# Fraction of compromised streams

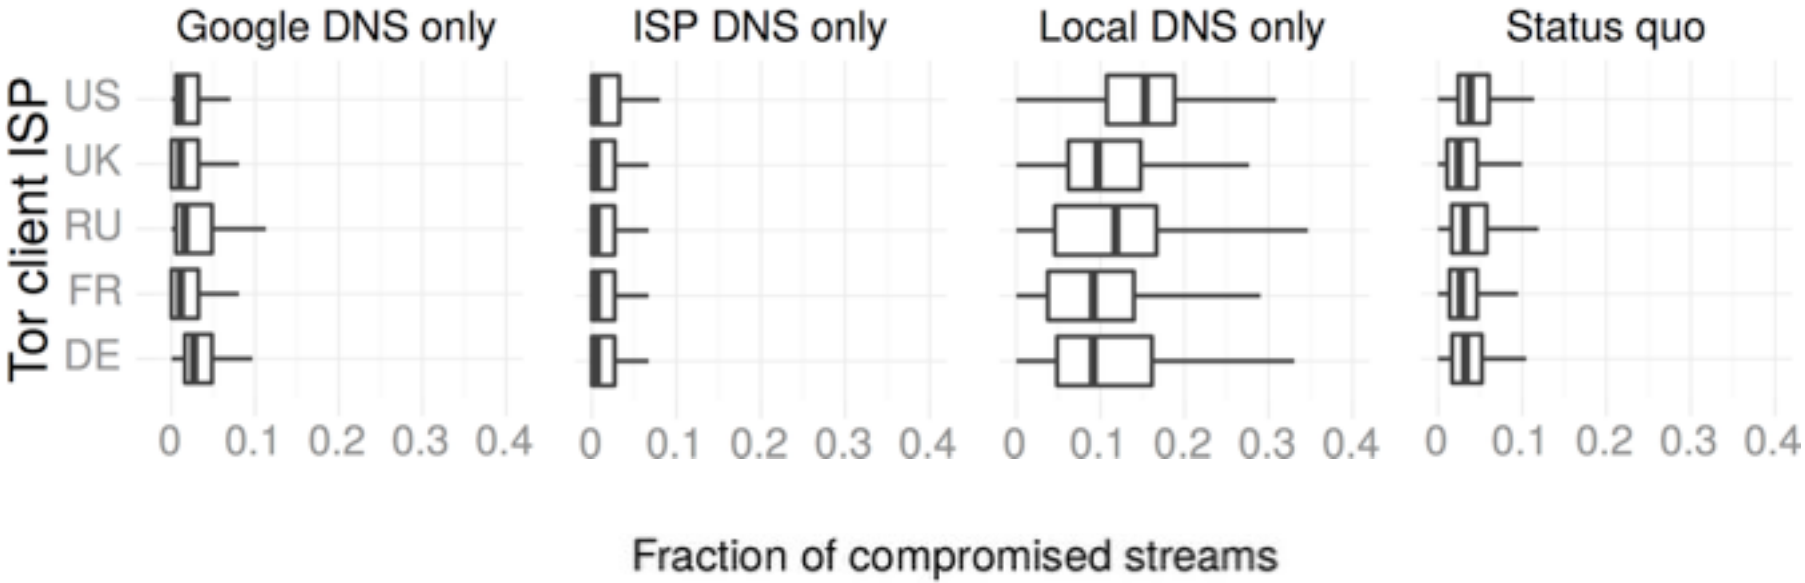


(a) The fraction of compromised streams of simulated Tor clients.

# Immediate Countermeasures

- Recommendations for exit relay operators
  - Don’t use Google’s 8.8.8.8
  - Use ISP’s resolver
  - Run their own resolver with QNAME minimization
- But it depends! Further study is required

# Long-term Solutions

- Add confidentiality to DNS
  - T-DNS (Zhu et al. Oakland'15)
- Improve website fingerprinting defenses

# Contributions

- Discovered that DNS exposes Tor users' behavior to more adversaries than previously thought

# Contributions

Discovered that DNS exposes Tor users' behavior to more adversaries than previously thought

Discovered that Google gets to learn a lot about Tor users' online activity

# Contributions

Discovered that DNS exposes Tor users' behavior to more adversaries than previously thought

Discovered that Google gets to learn a lot about Tor users' online activity

Created proof-of-concept deanonymization attacks that demonstrate how DNS can make website fingerprinting attacks more precise

# Contributions

Discovered that DNS exposes Tor users' behavior to more adversaries than previously thought

Discovered that Google gets to learn a lot about Tor users' online activity

Created proof-of-concept deanonymization attacks that demonstrate how DNS can make website fingerprinting attacks more precise

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

# Contributions

Our work compels researchers to continue exploring how to make DNS more secure

# Fin

- Paper, data, code, and replication instructions: https://nymity.ch/tor-dns/
- Contact: laurar@cs.princeton.edu

Nick

Tobias

Laura

Benjamin

Philipp