Privacy Enhancing Technologies

Marc Juarez

COSIC KU Leuven and iMinds

marc.juarez {at} kuleuven.be

SecAppDev, March 2016

Outline

1. Introduction to traffic analysis
2. The traffic analysis threat model
3. PETs to protect against traffic analysis.
4. The Onion Router (Tor)
5. Traffic analysis attacks and defences in Tor

Outline

1. **Introduction to traffic analysis**
2. The traffic analysis threat model
3. PETs to protect against traffic analysis.
4. The Onion Router (Tor)
5. Traffic analysis attacks and defences in Tor

Traffic Analysis (TA)

"Making use of the traffic data (**metadata**) of a communication to extract information".

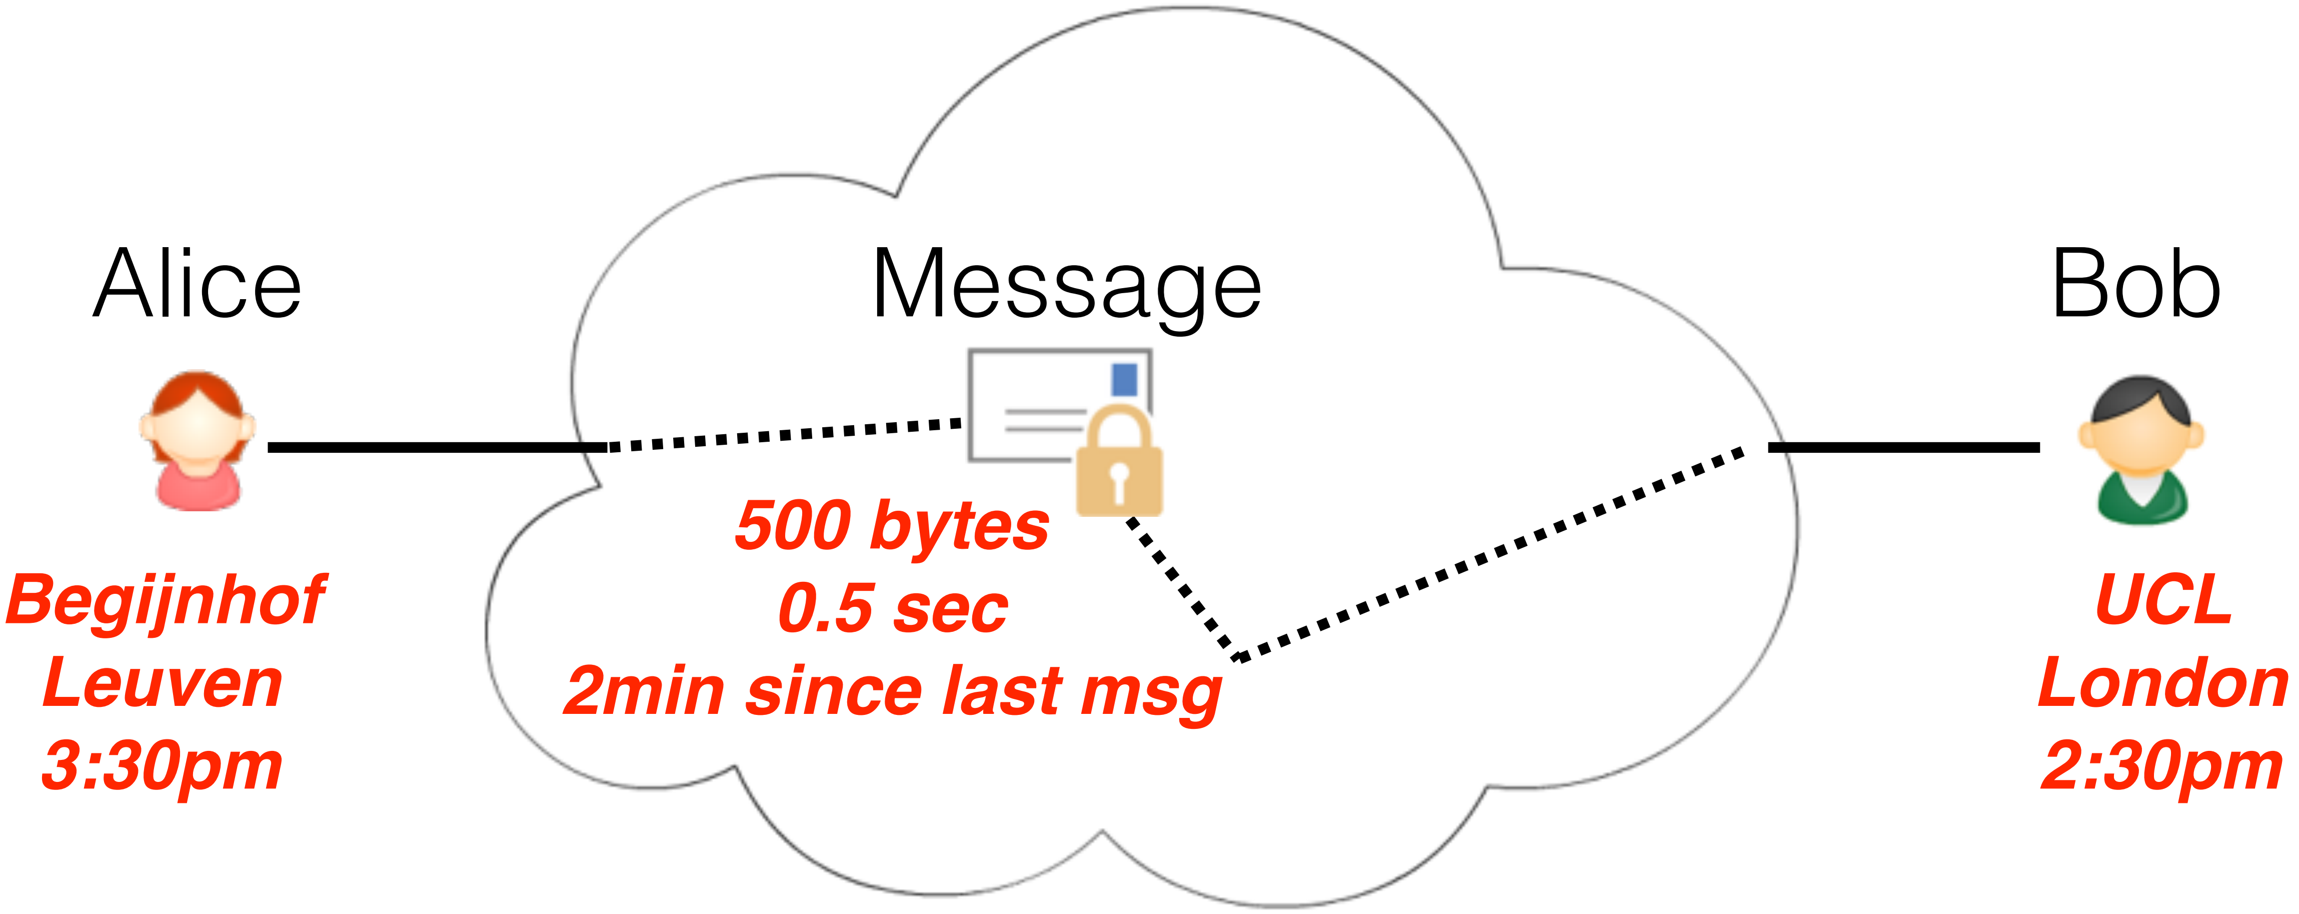

“Traffic analysis, not cryptanalysis, is the backbone of communications intelligence.”

–B. W. Diffie and S. Landau,
’Privacy on the Line': The Politics of Wiretapping and Encryption, 1999.

Source: http://generalhaydenisred.blogspot.be

“We kill people based on metadata.”

–General Michael Hayden, former director of the NSA and the CIA, during the Johns Hopkins Foreign Affairs Symposium, 2014.

Why is traffic analysis so valuable?

- Compared to cryptanalysis it is both easier and cheaper to extract and process.
- Often used to tackle the ‘needle in a haystack’ problem and perform ‘target selection’.

Credit: G. Danezis

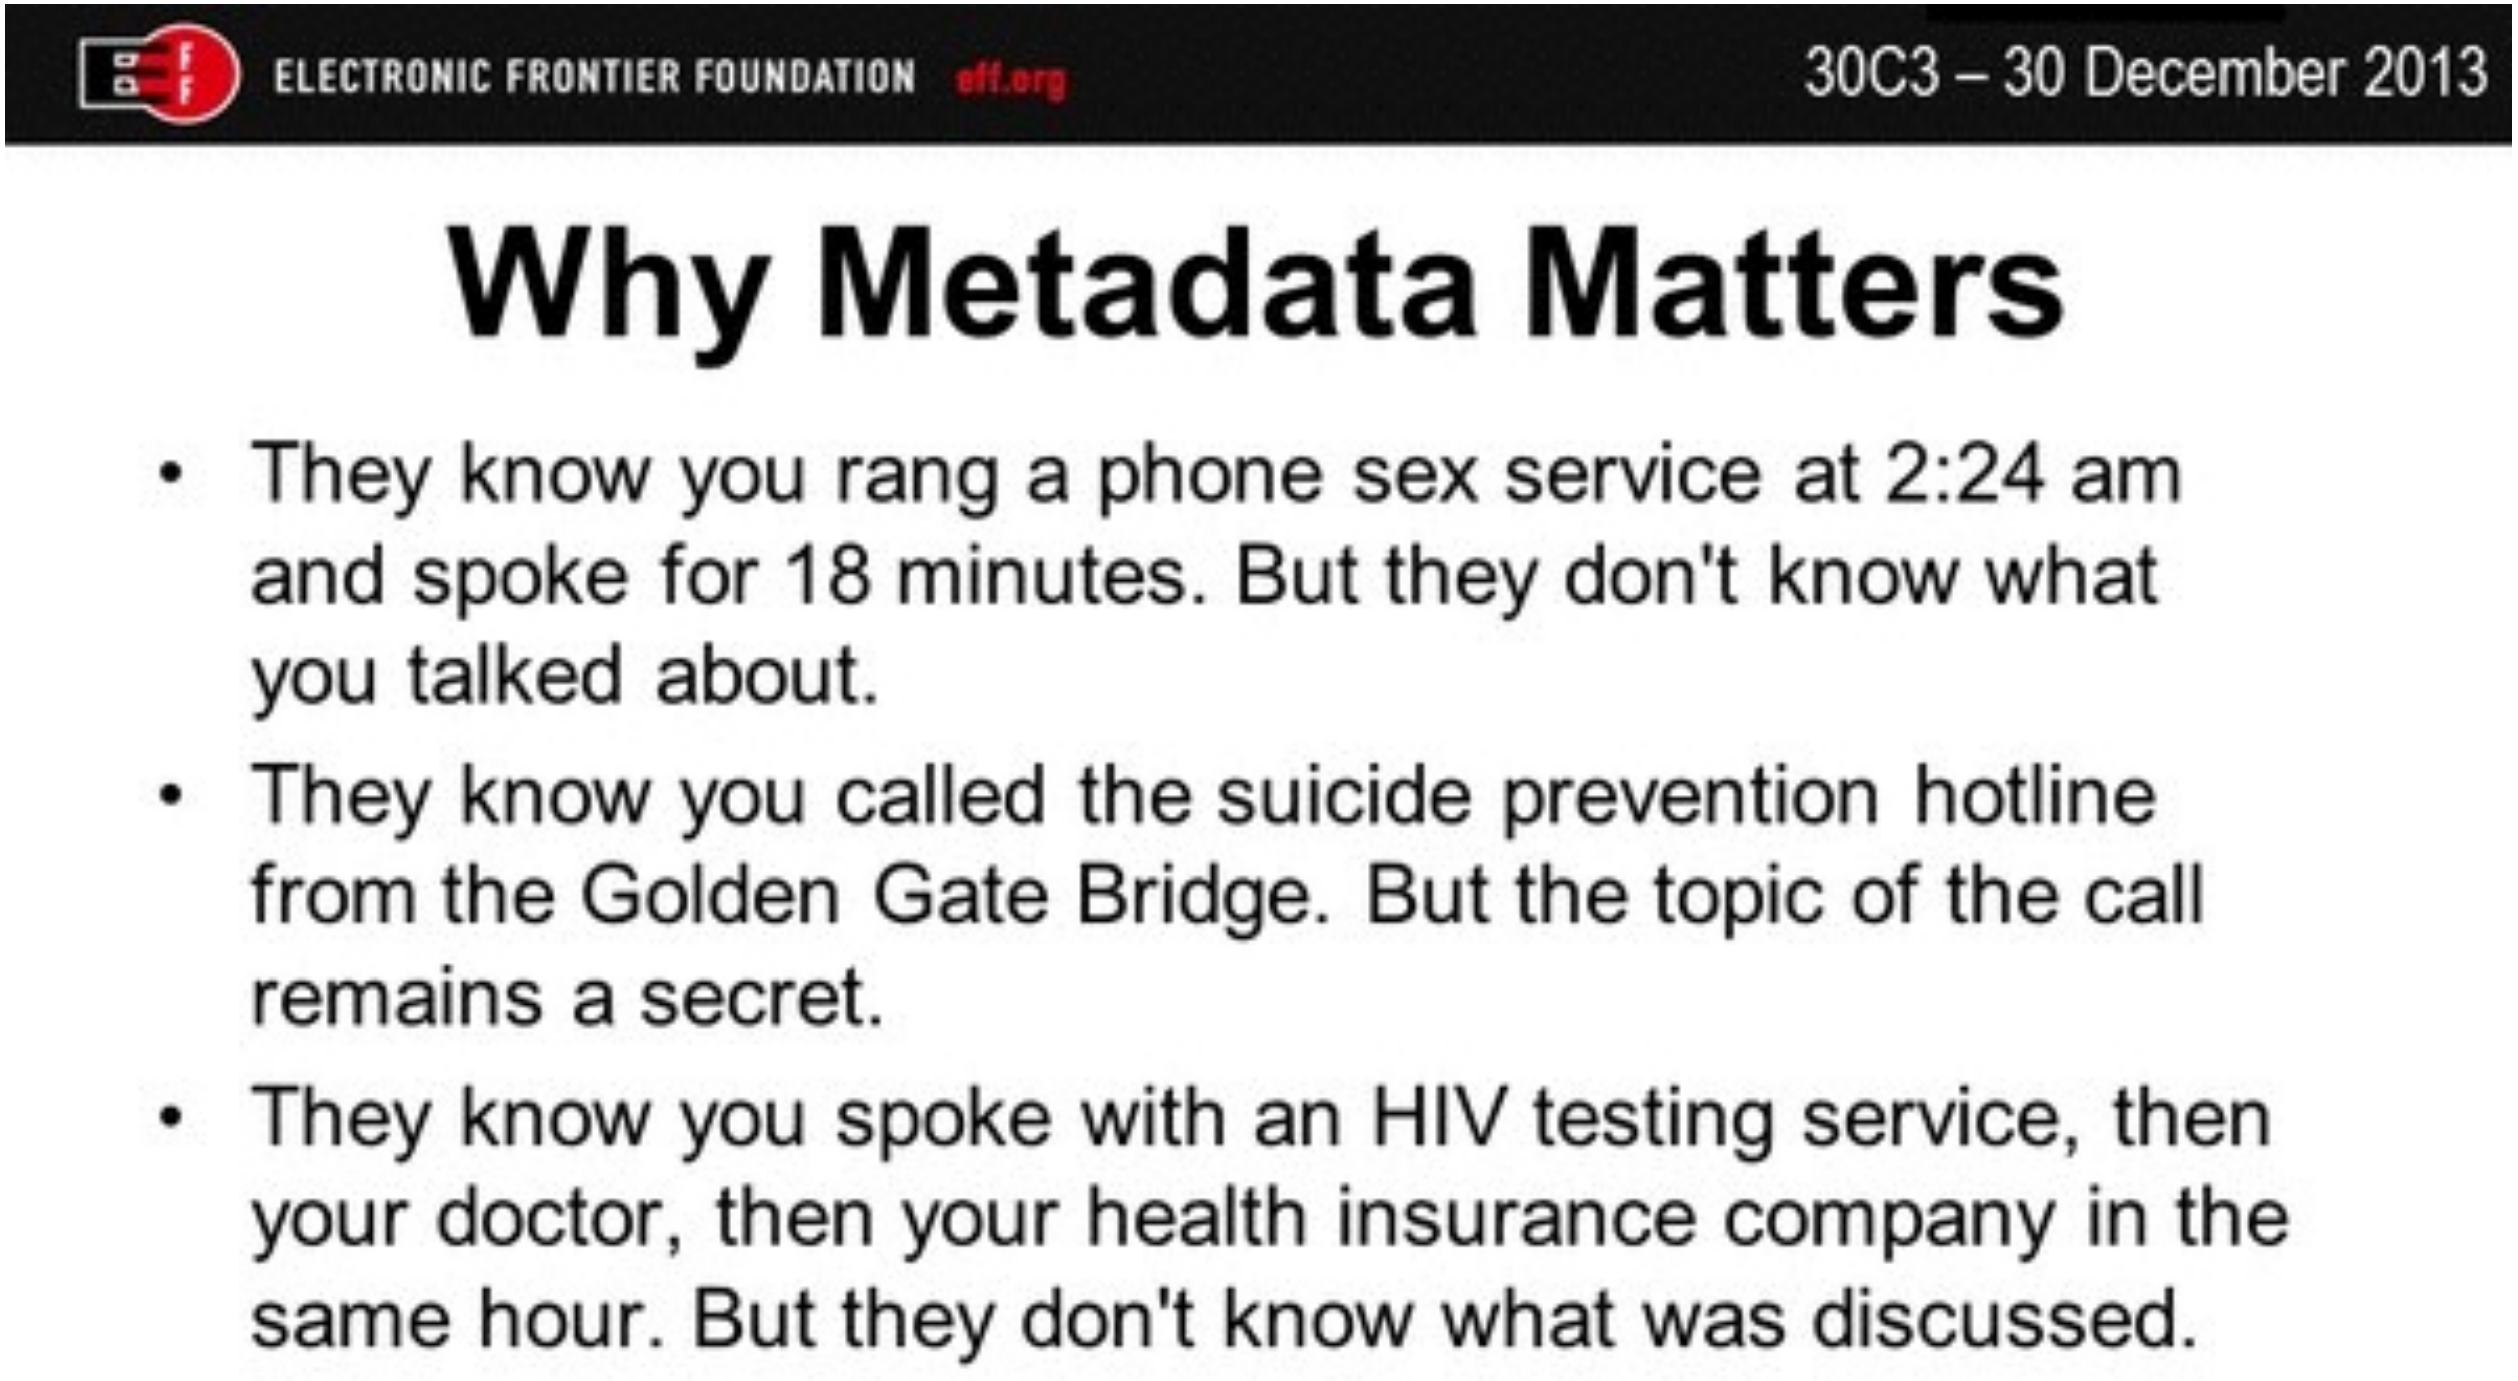

ELECTRONIC FRONTIER FOUNDATION eff.org

30C3 – 30 December 2013

Why Metadata Matters

- They know you rang a phone sex service at 2:24 am and spoke for 18 minutes. But they don't know what you talked about.
- They know you called the suicide prevention hotline from the Golden Gate Bridge. But the topic of the call remains a secret.
- They know you spoke with an HIV testing service, then your doctor, then your health insurance company in the same hour. But they don't know what was discussed.

Source: tweet by Kurt Opsahl, EFF

Traffic Analysis on the Internet

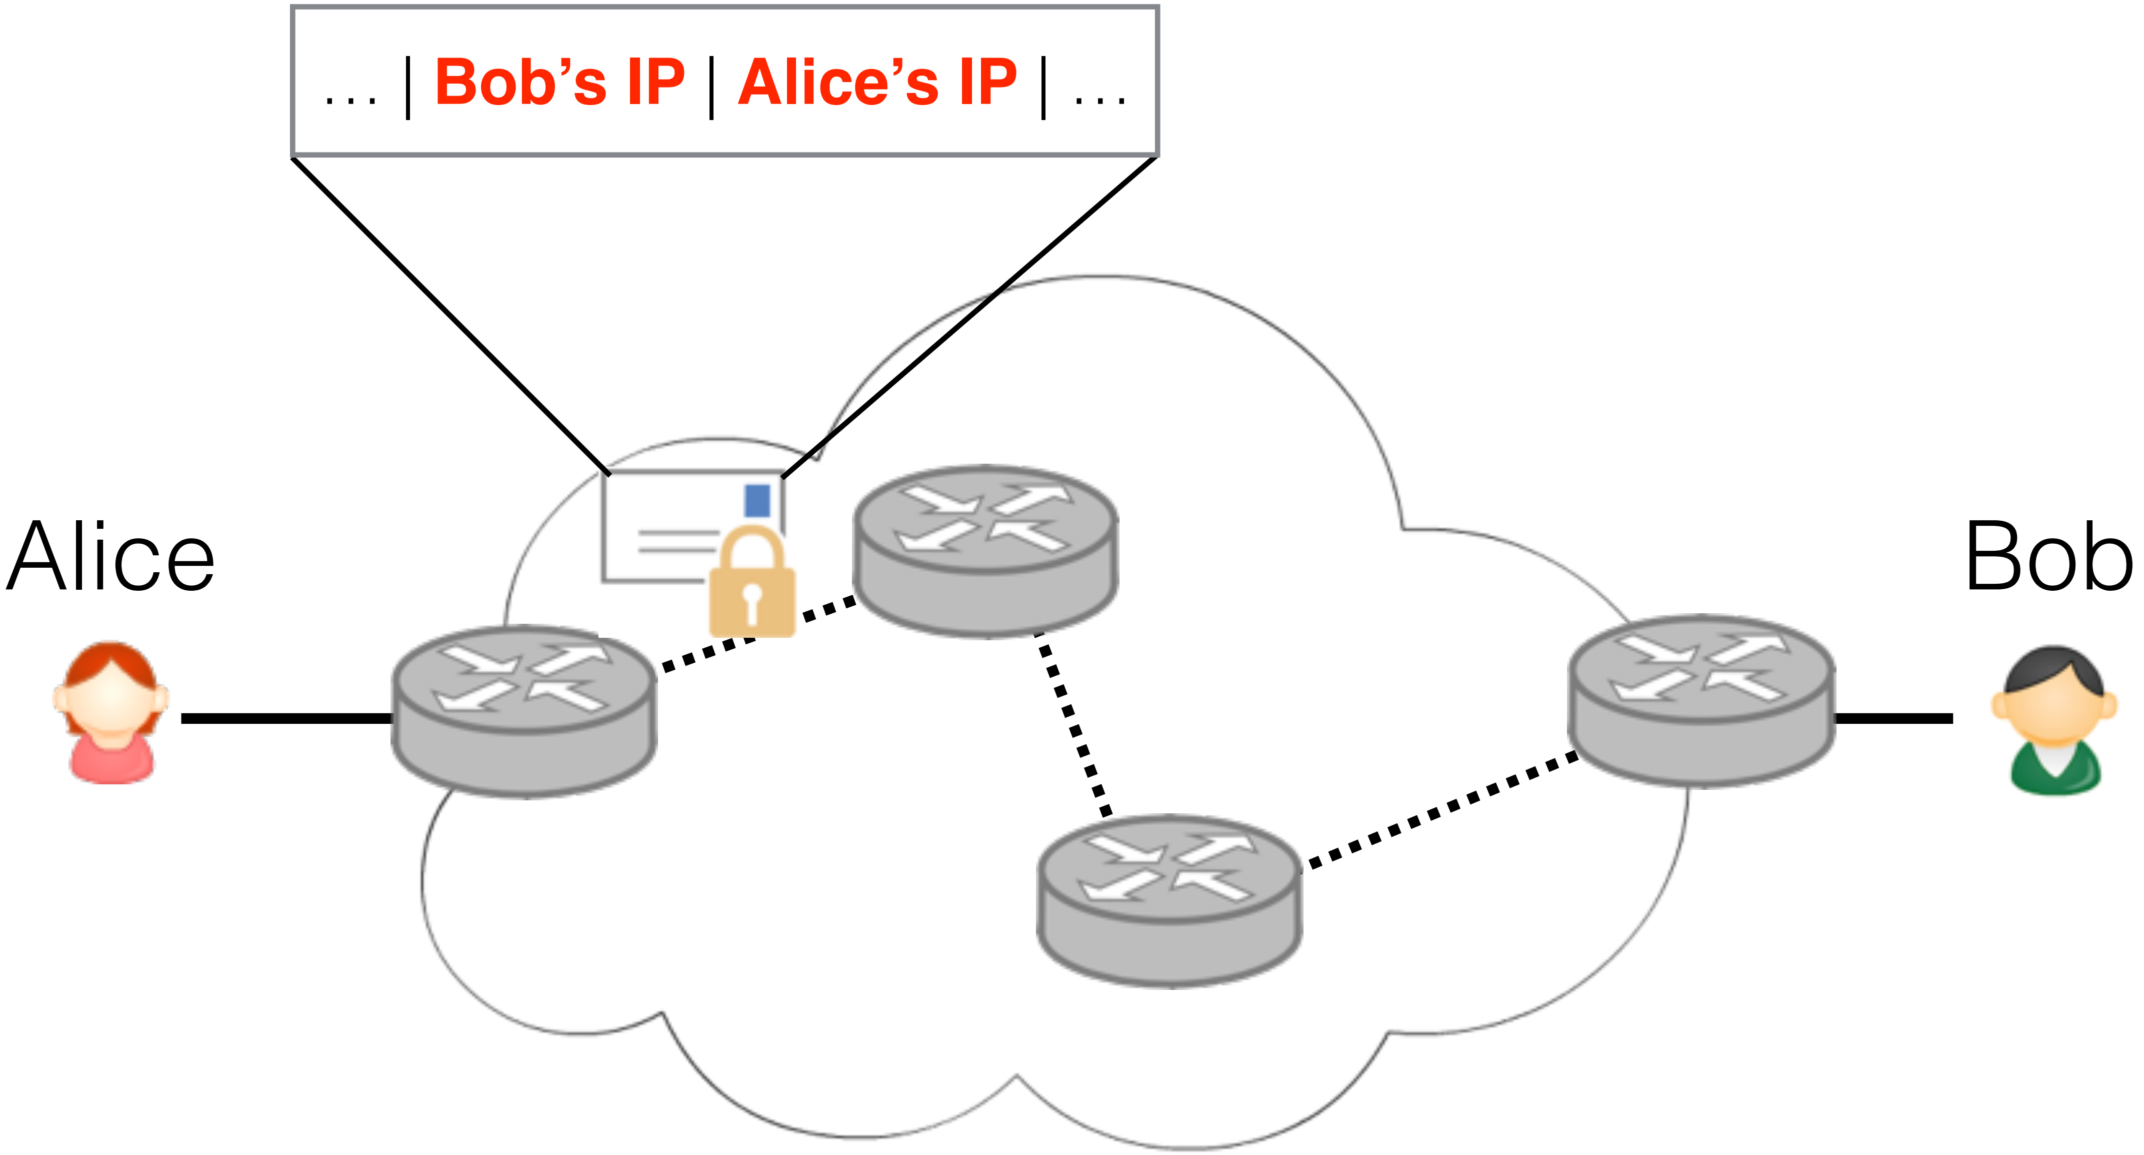

Does SSH protect your privacy?

- SSH is used for secure remote access and file transfer.
- What can we extract about a password typed in an SSH session?
- Observation: each key press is transmitted separately. Length of password is observable.
- Depending on the position of the key on the keyboard, different inter-key timings.
- Song et al.: reduce the entropy of password – fewer guesses required.

Credit: G. Danezis

Outline

1. Introduction to traffic analysis

2. **The traffic analysis threat model**

3. PETs to protect against traffic analysis.

4. The Onion Router (Tor)

5. Traffic analysis attacks and defences in Tor

Traffic analysis resistance properties

- **Anonymity**: one of the communicating parties cannot trace the identity of the others, or both can communicate without knowing each other's identities.

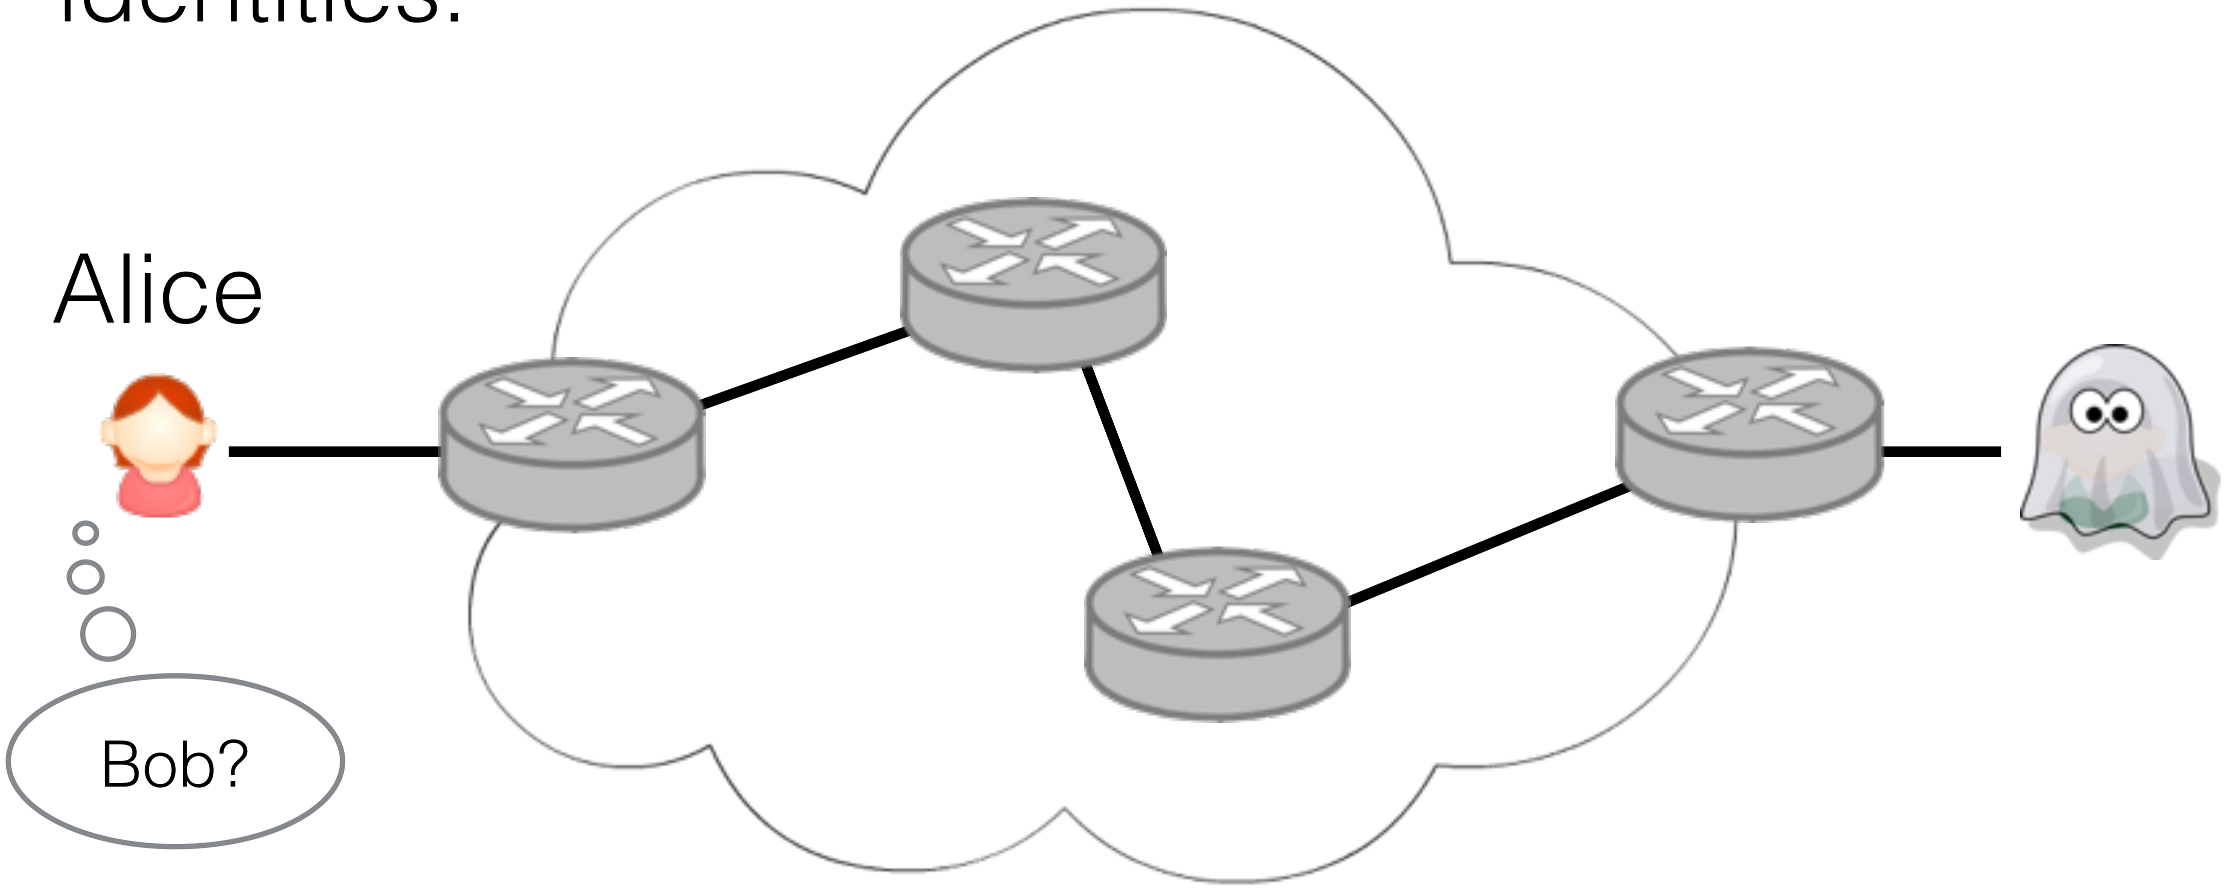

Traffic analysis resistance properties

- **3rd-party anonymity**: communicating parties trust each-other but do not want other parties to learn they are communicating.

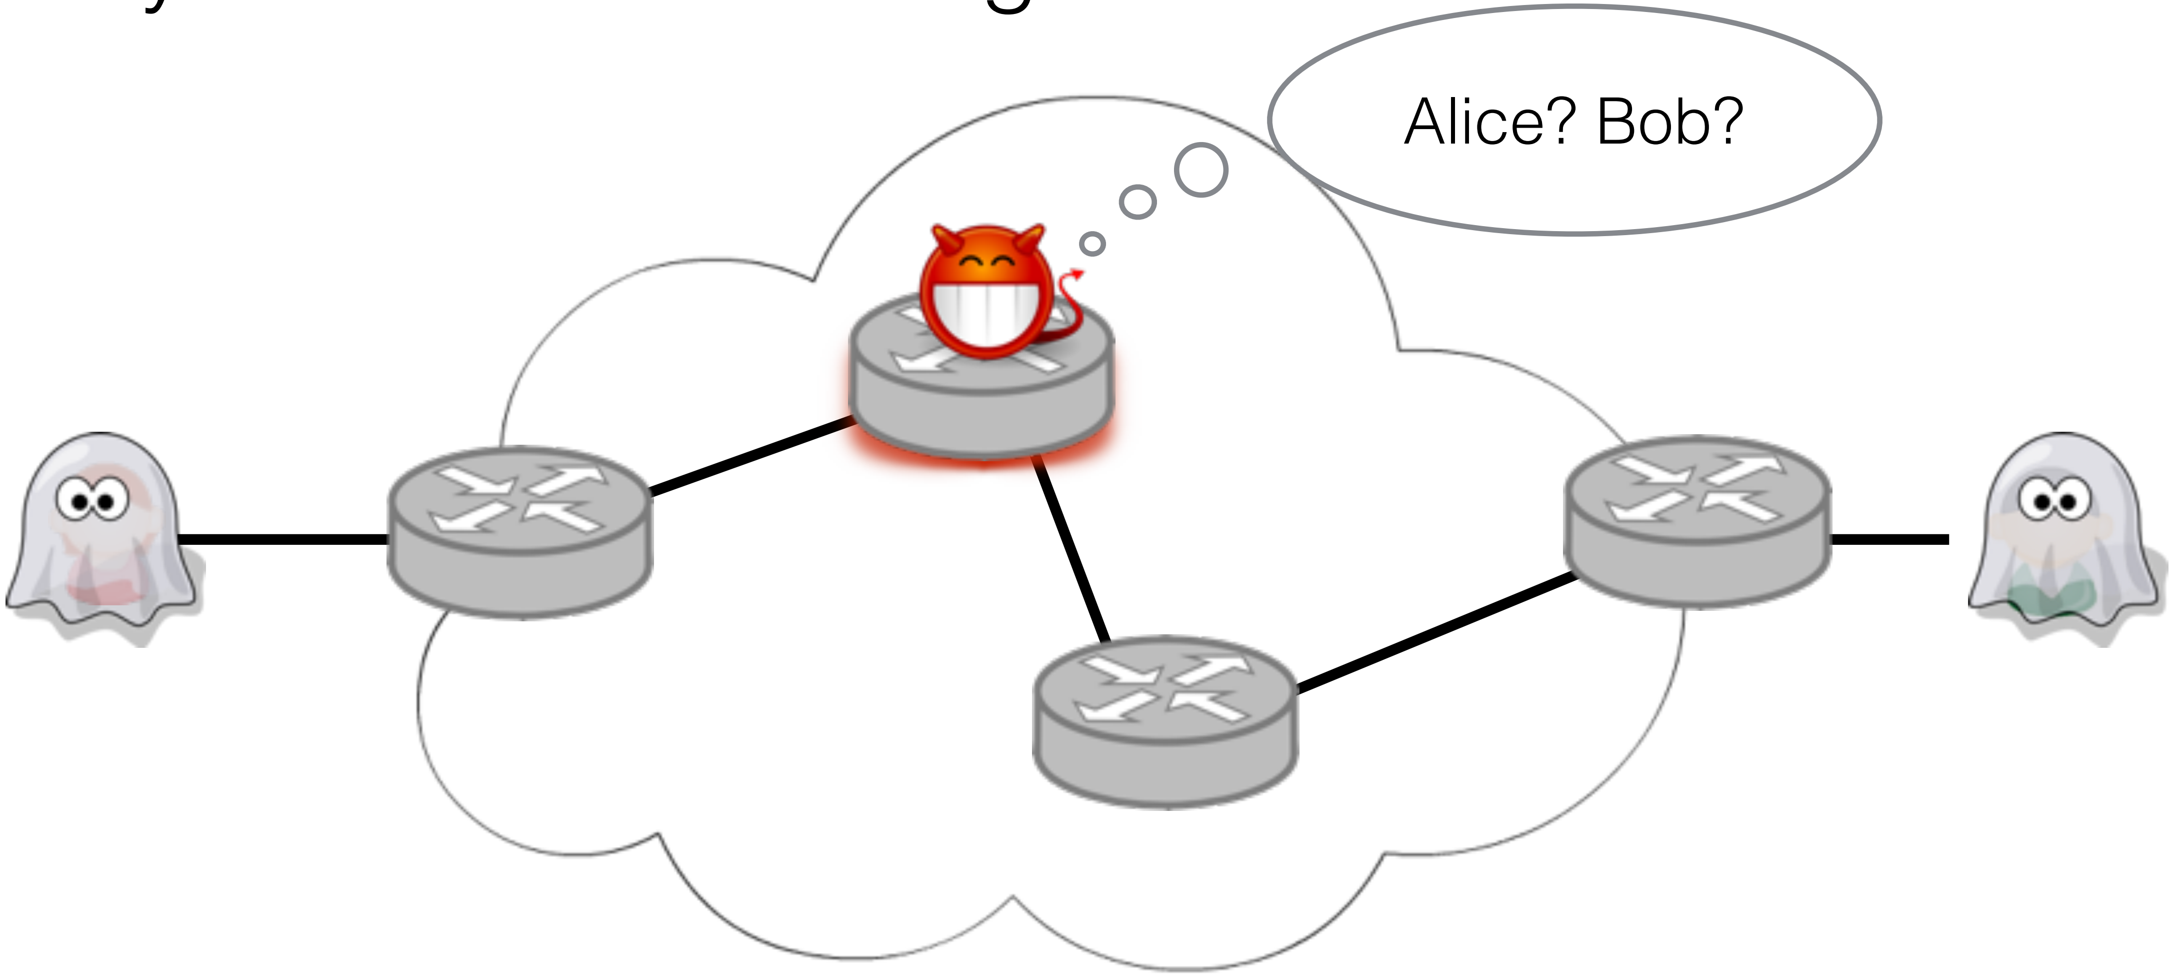

Traffic analysis resistance properties

- **Unlinkability**: the attacker cannot link two events (e.g., parties, messages, etc.).

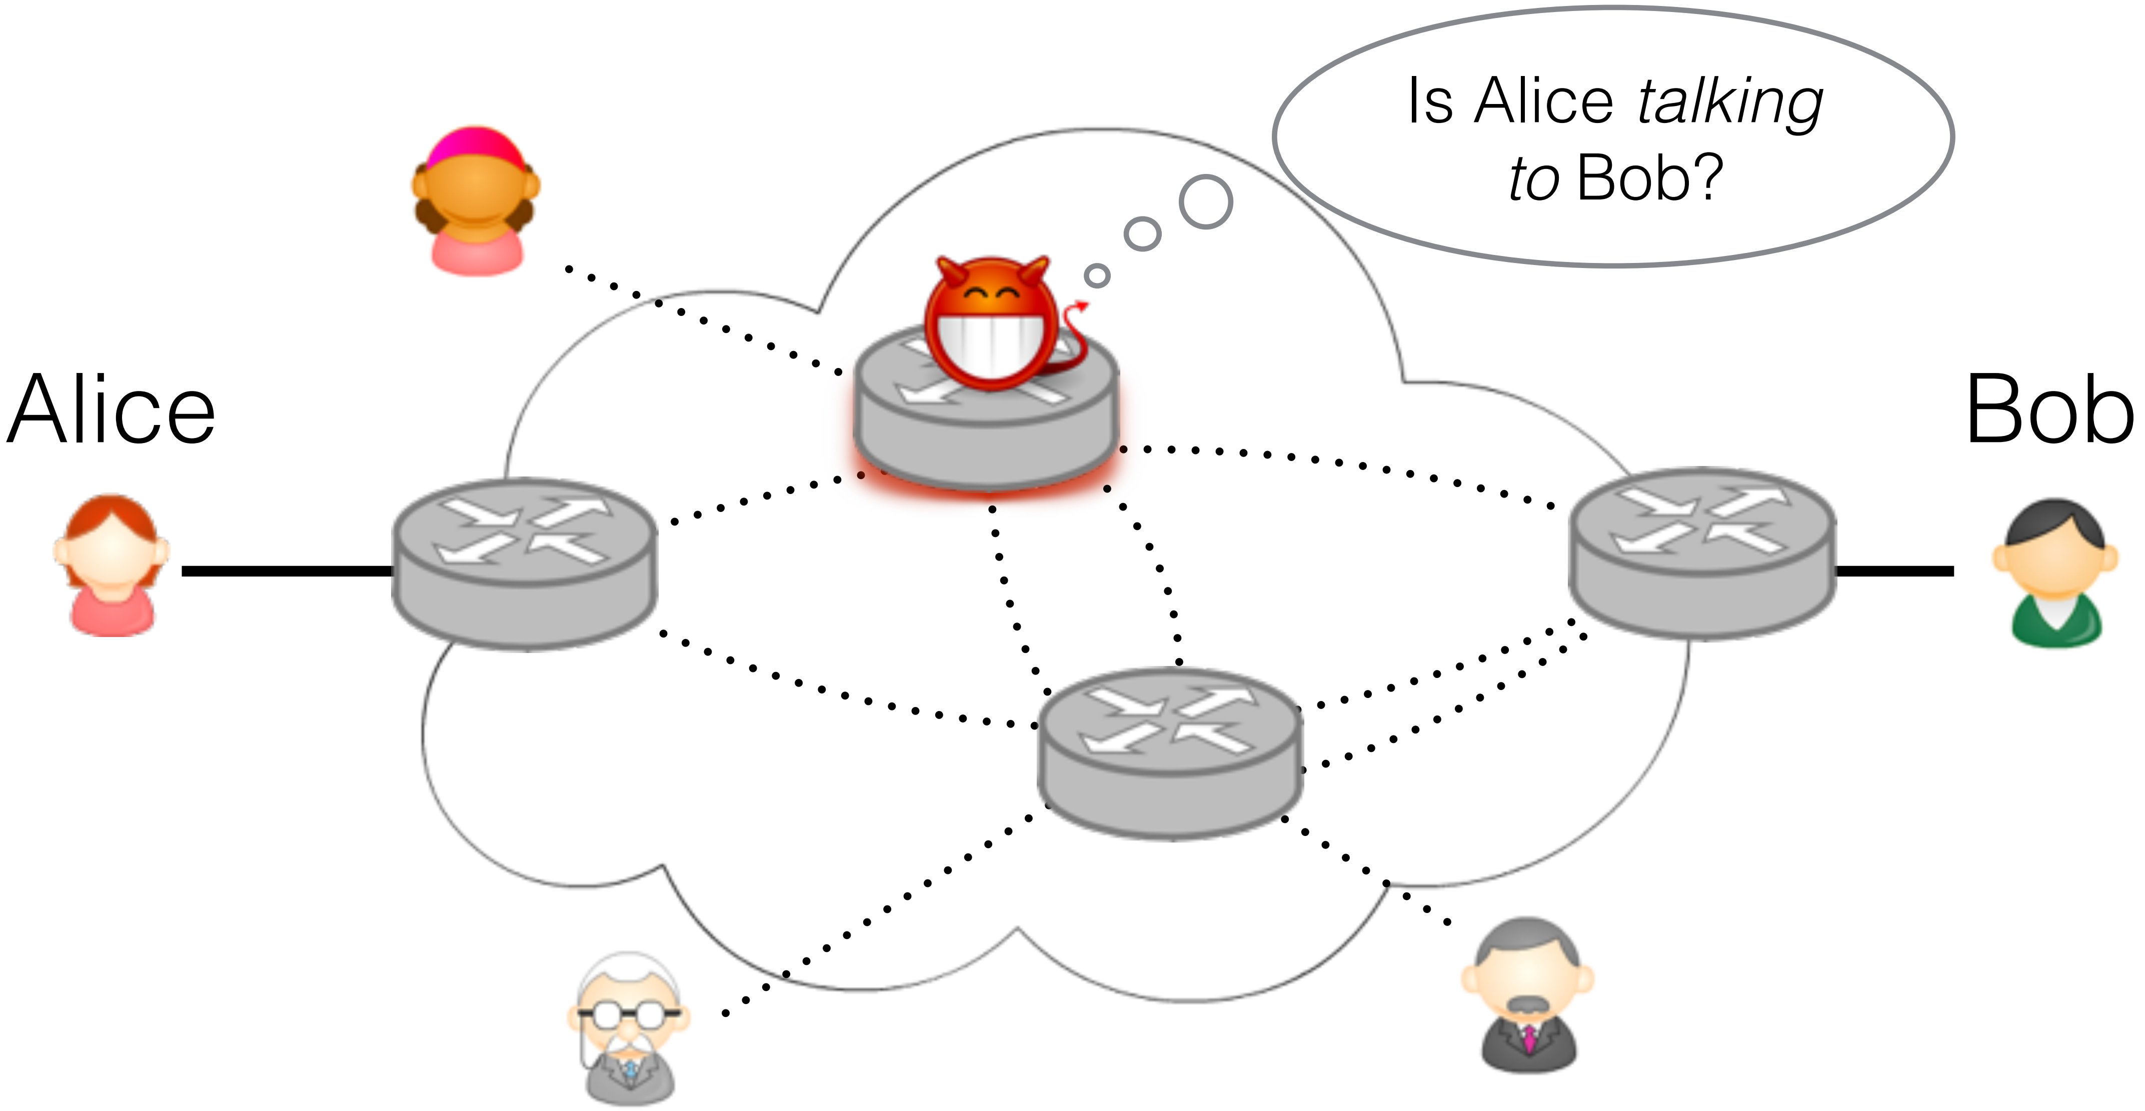

Adversary’s capabilities

Generally **passive** but it can also be **active**:

- *Add* messages at any point of the network.
- *Delete* or *delay* messages.
- *Modify* messages to help with tracing.

Limitations:

- Cannot break cryptographic primitives.
- Cannot see inside nodes he does not control.

Credit: C. Diaz

Adversary’s view of the network

- **Global**: can observe all communication links and look at all traffic.

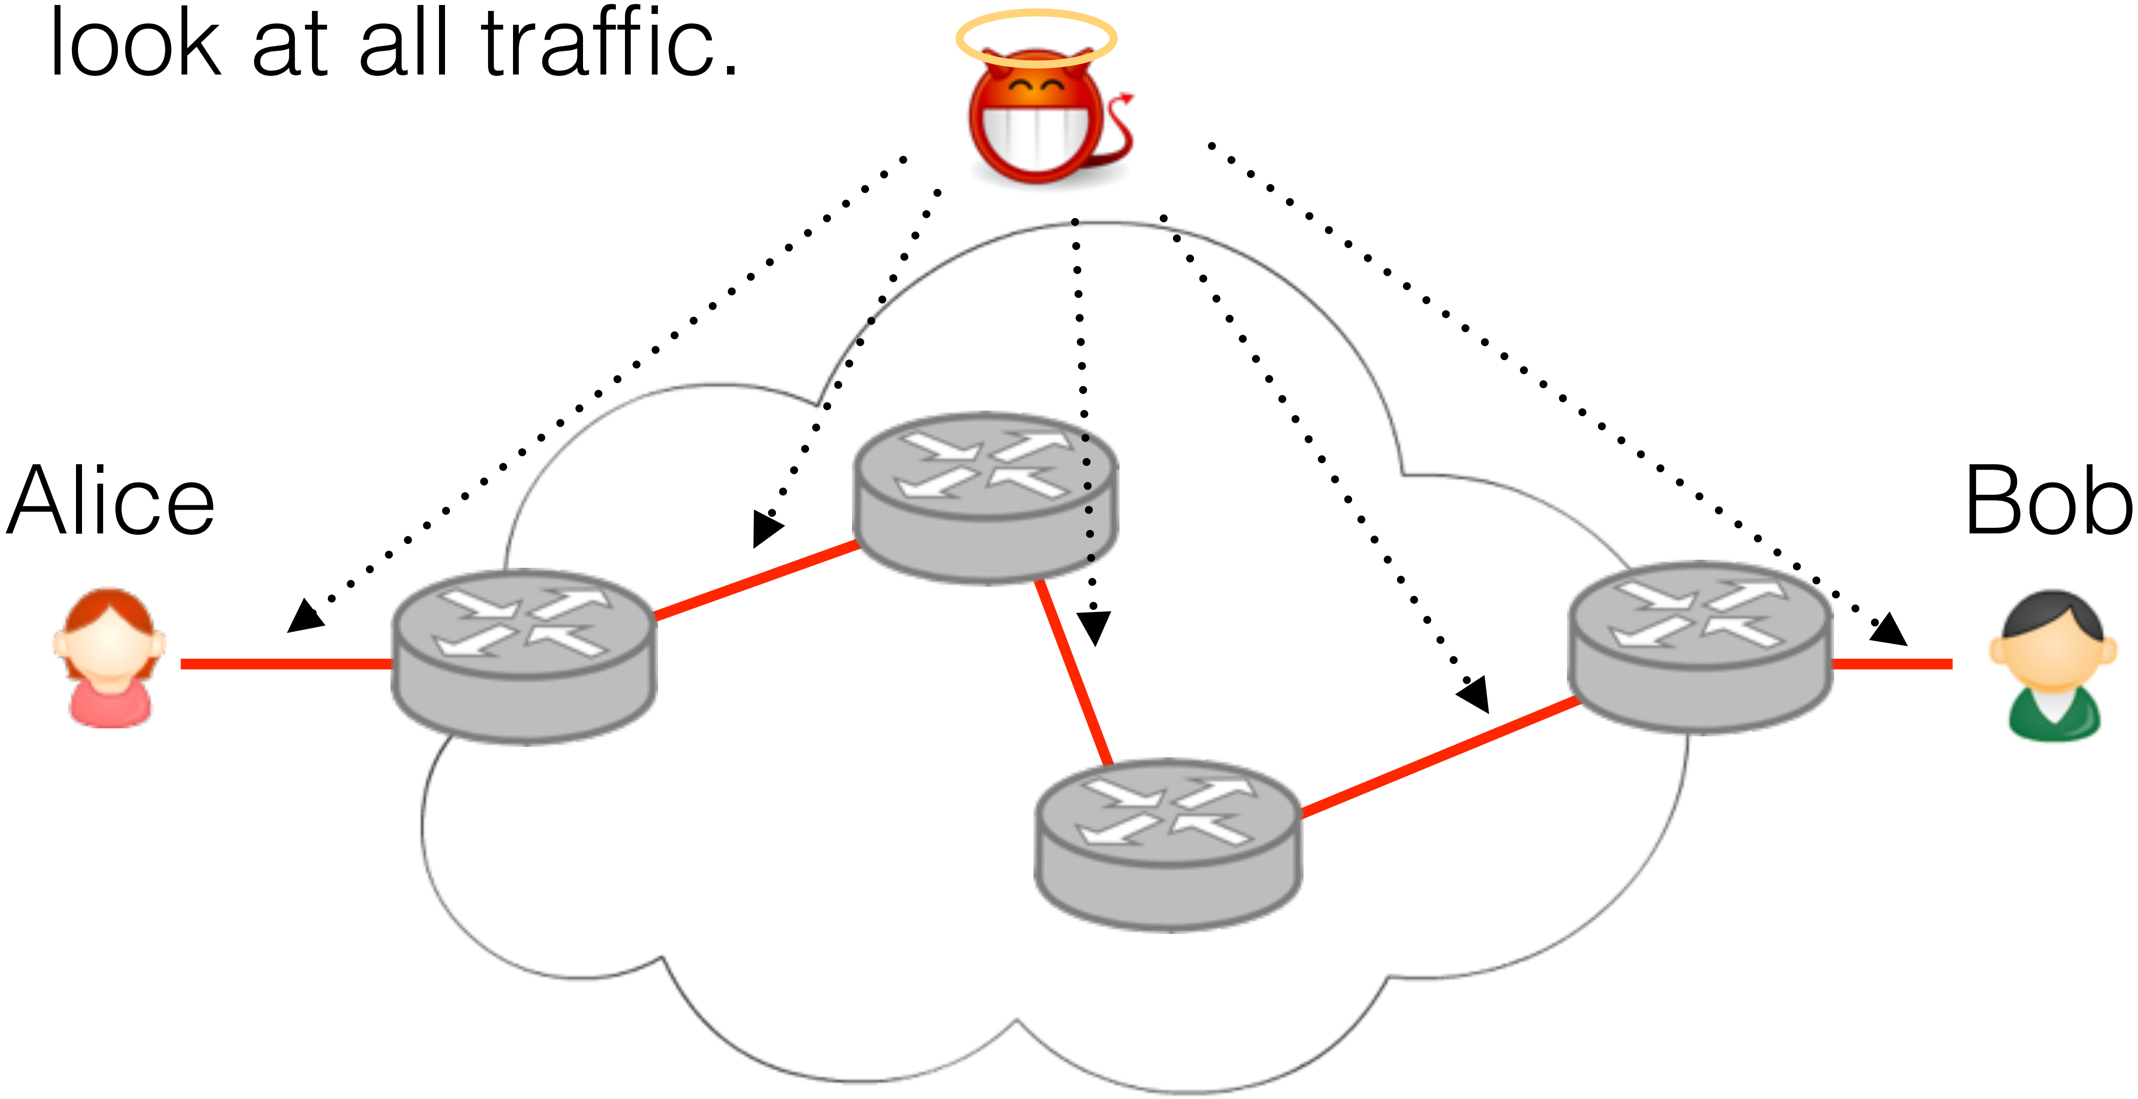

Adversary’s view of the network

- **Local**: controls an edge of the network (ISP, employer, malicious sender/receiver).

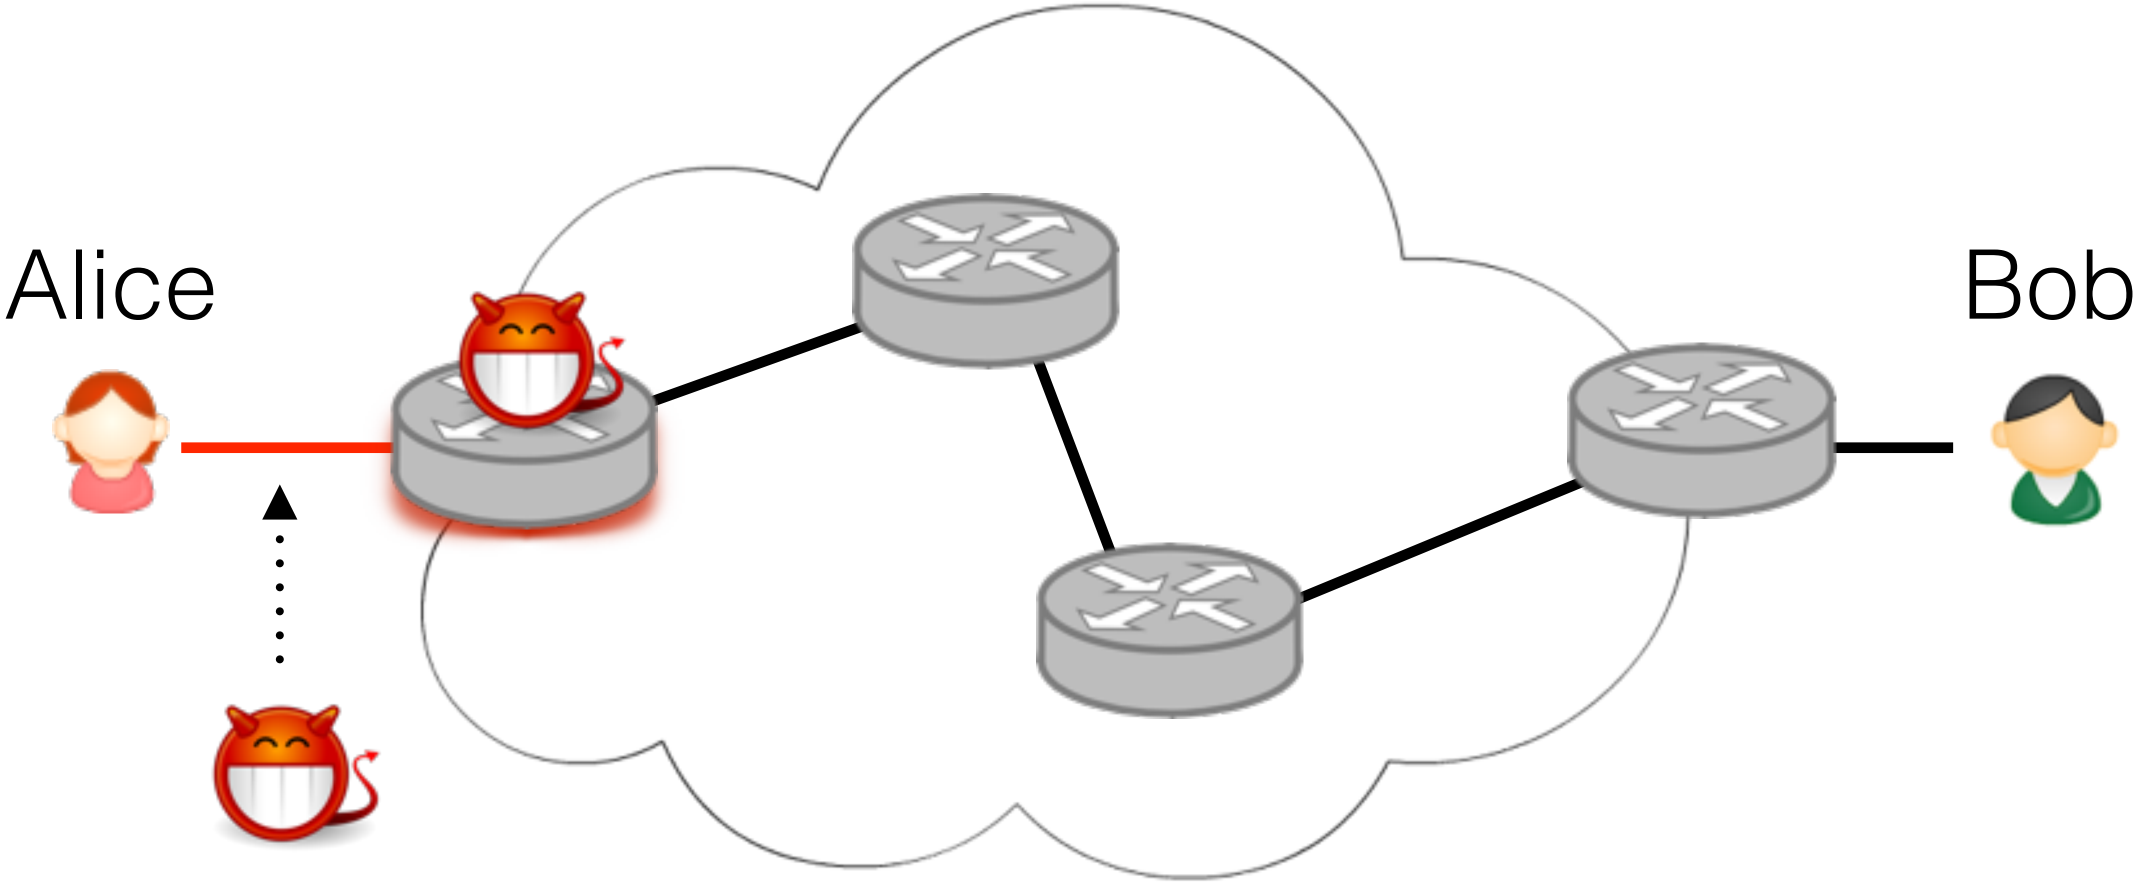

Corrupt participants

- Corrupt insiders: “anonymity loves company”.
- Corrupt nodes.
- Sybil attacks.

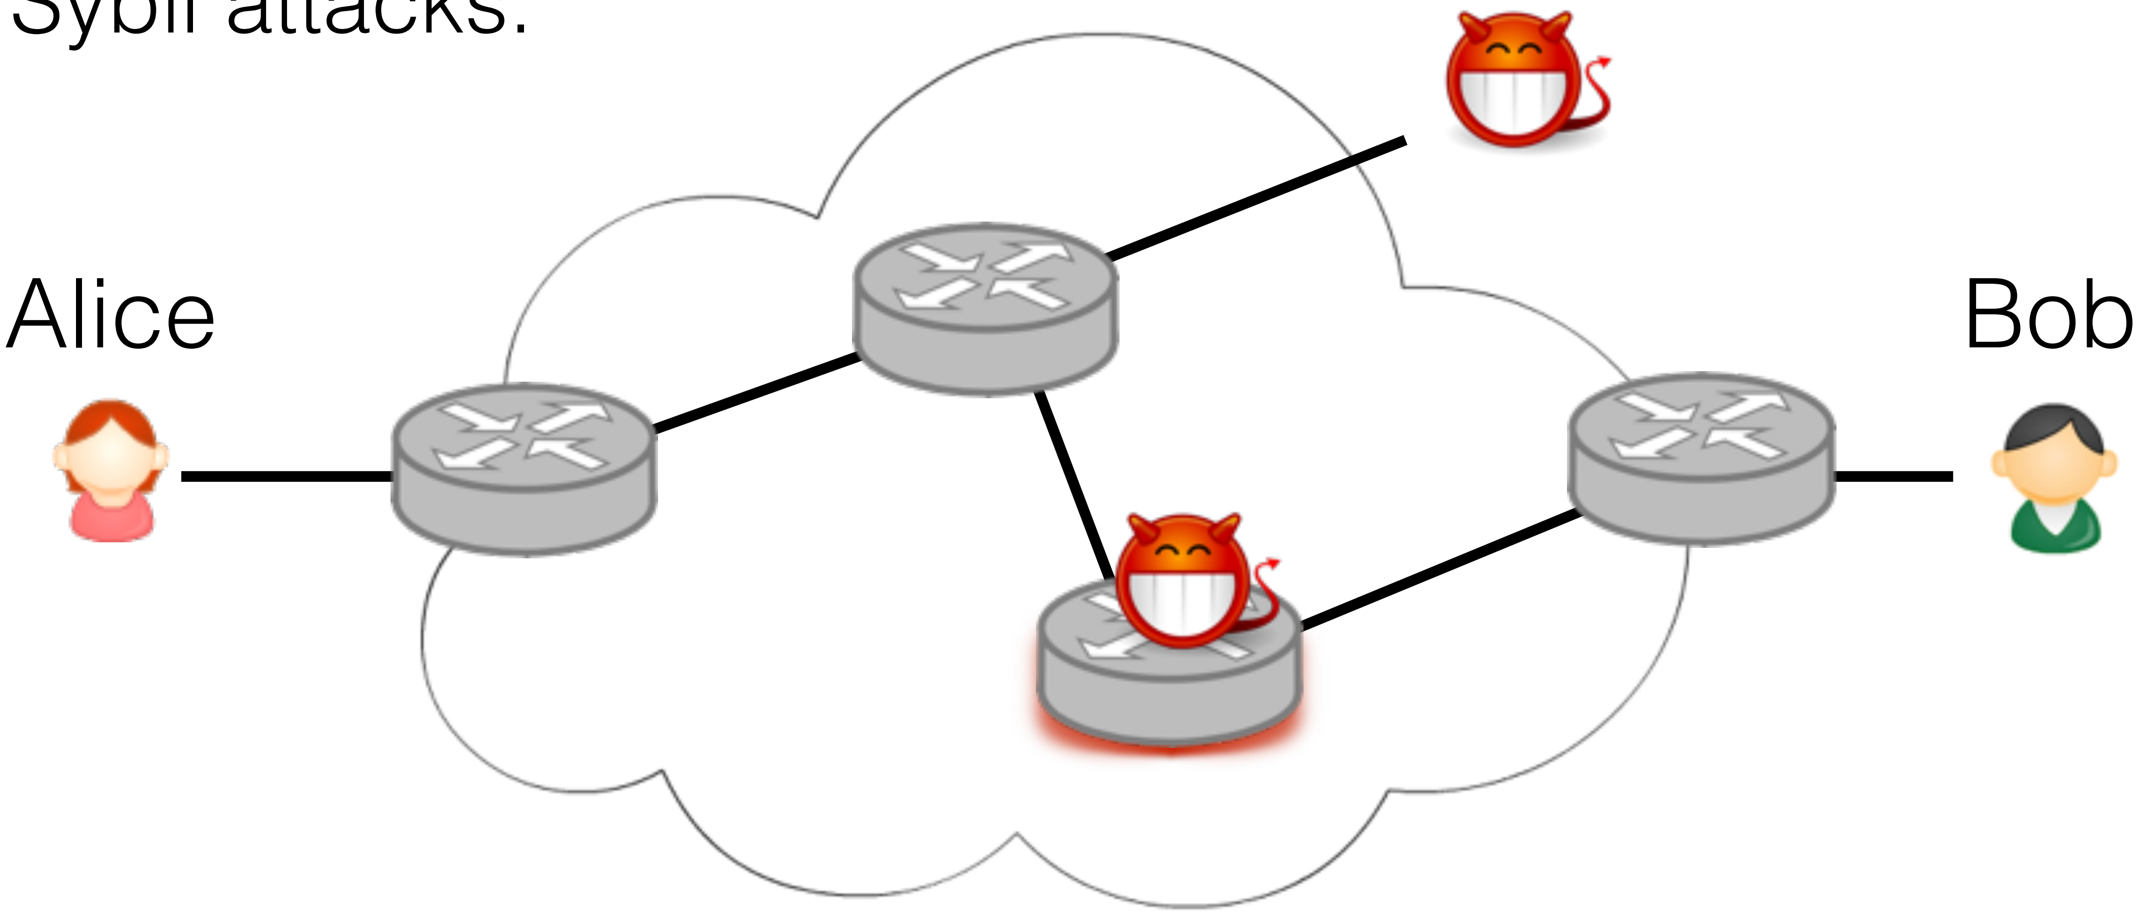

Credit: C. Diaz

Coercion

- Honest nodes may be forced to cooperate with the adversary. Blackmail, bribery, legal or physical threats.
  - Nodes should:
    - Know as few secrets as possible.
    - Be given the opportunity to (plausibly) lie.

Credit: C. Diaz

Outline

1. Introduction to traffic analysis
2. The traffic analysis threat model
3. **PETs to protect against traffic analysis.**
4. The Onion Router (Tor)
5. Traffic analysis attacks and defences in Tor

Chaumian Mix

- D. Chaum, 1982: fundamental building block.
- Goal: an adversary observing input and output of the mix should not be able to relate input messages with output messages.
- Messages of fixed size (large messages divided into blocks if needed).

Credit: C. Diaz

Mix example

- Phase 1: collect inputs.
- Parameter T (threshold): T = 4 in the example.

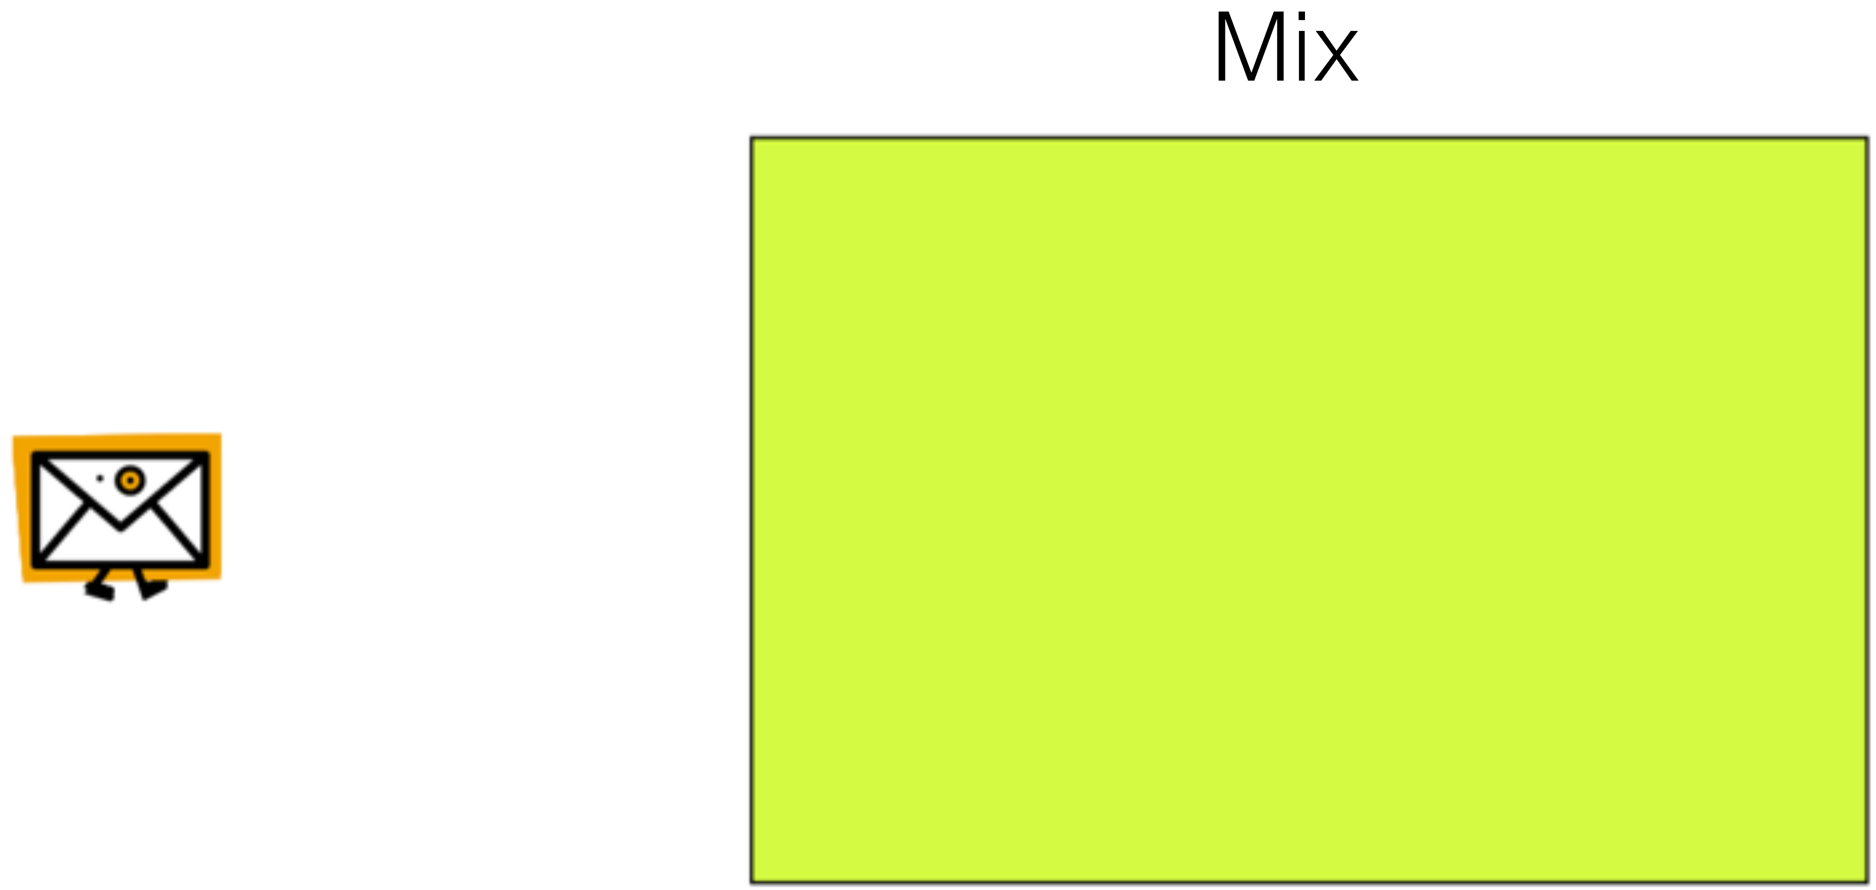

Mix

Credit: C. Diaz

Mix example

- Phase 2: mix and flush.

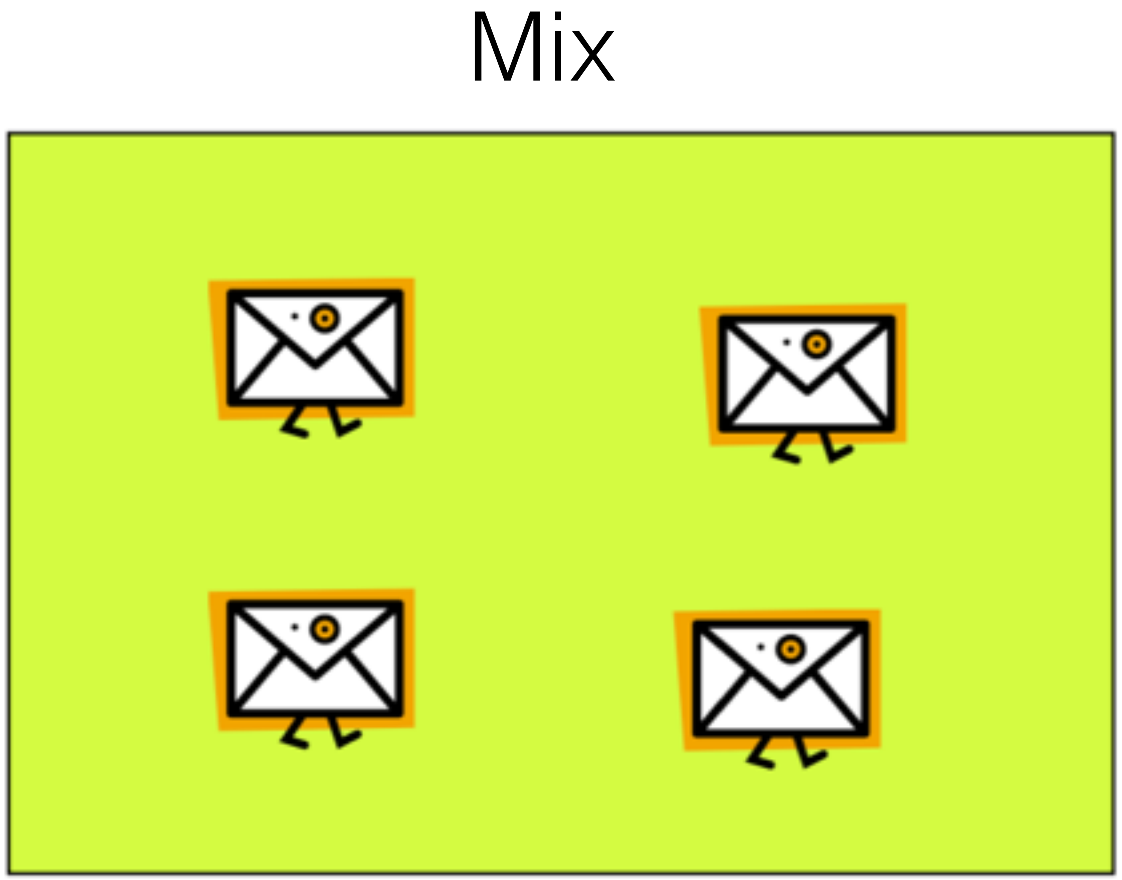

Credit: C. Diaz

Attacking mixes

- "*n-1*" attacks.
- Long-term intersection attacks:
  - Danezis and Serjantov: Statistical Disclosure.
  - Troncoso: Perfect matching Disclosure.

Credit: C. Diaz

Onion Routing

- Designed for **low-latency** applications.
  - Web browsing, secure shell, instant messaging.
  - No mixing, padding or traffic shaping (just forward).
  - Connection-based (instead of message-based).
- Started at the US Navy Research Lab (1996).
  - Need to “mix” with civilians!

Credit: C. Diaz

Onion Routing

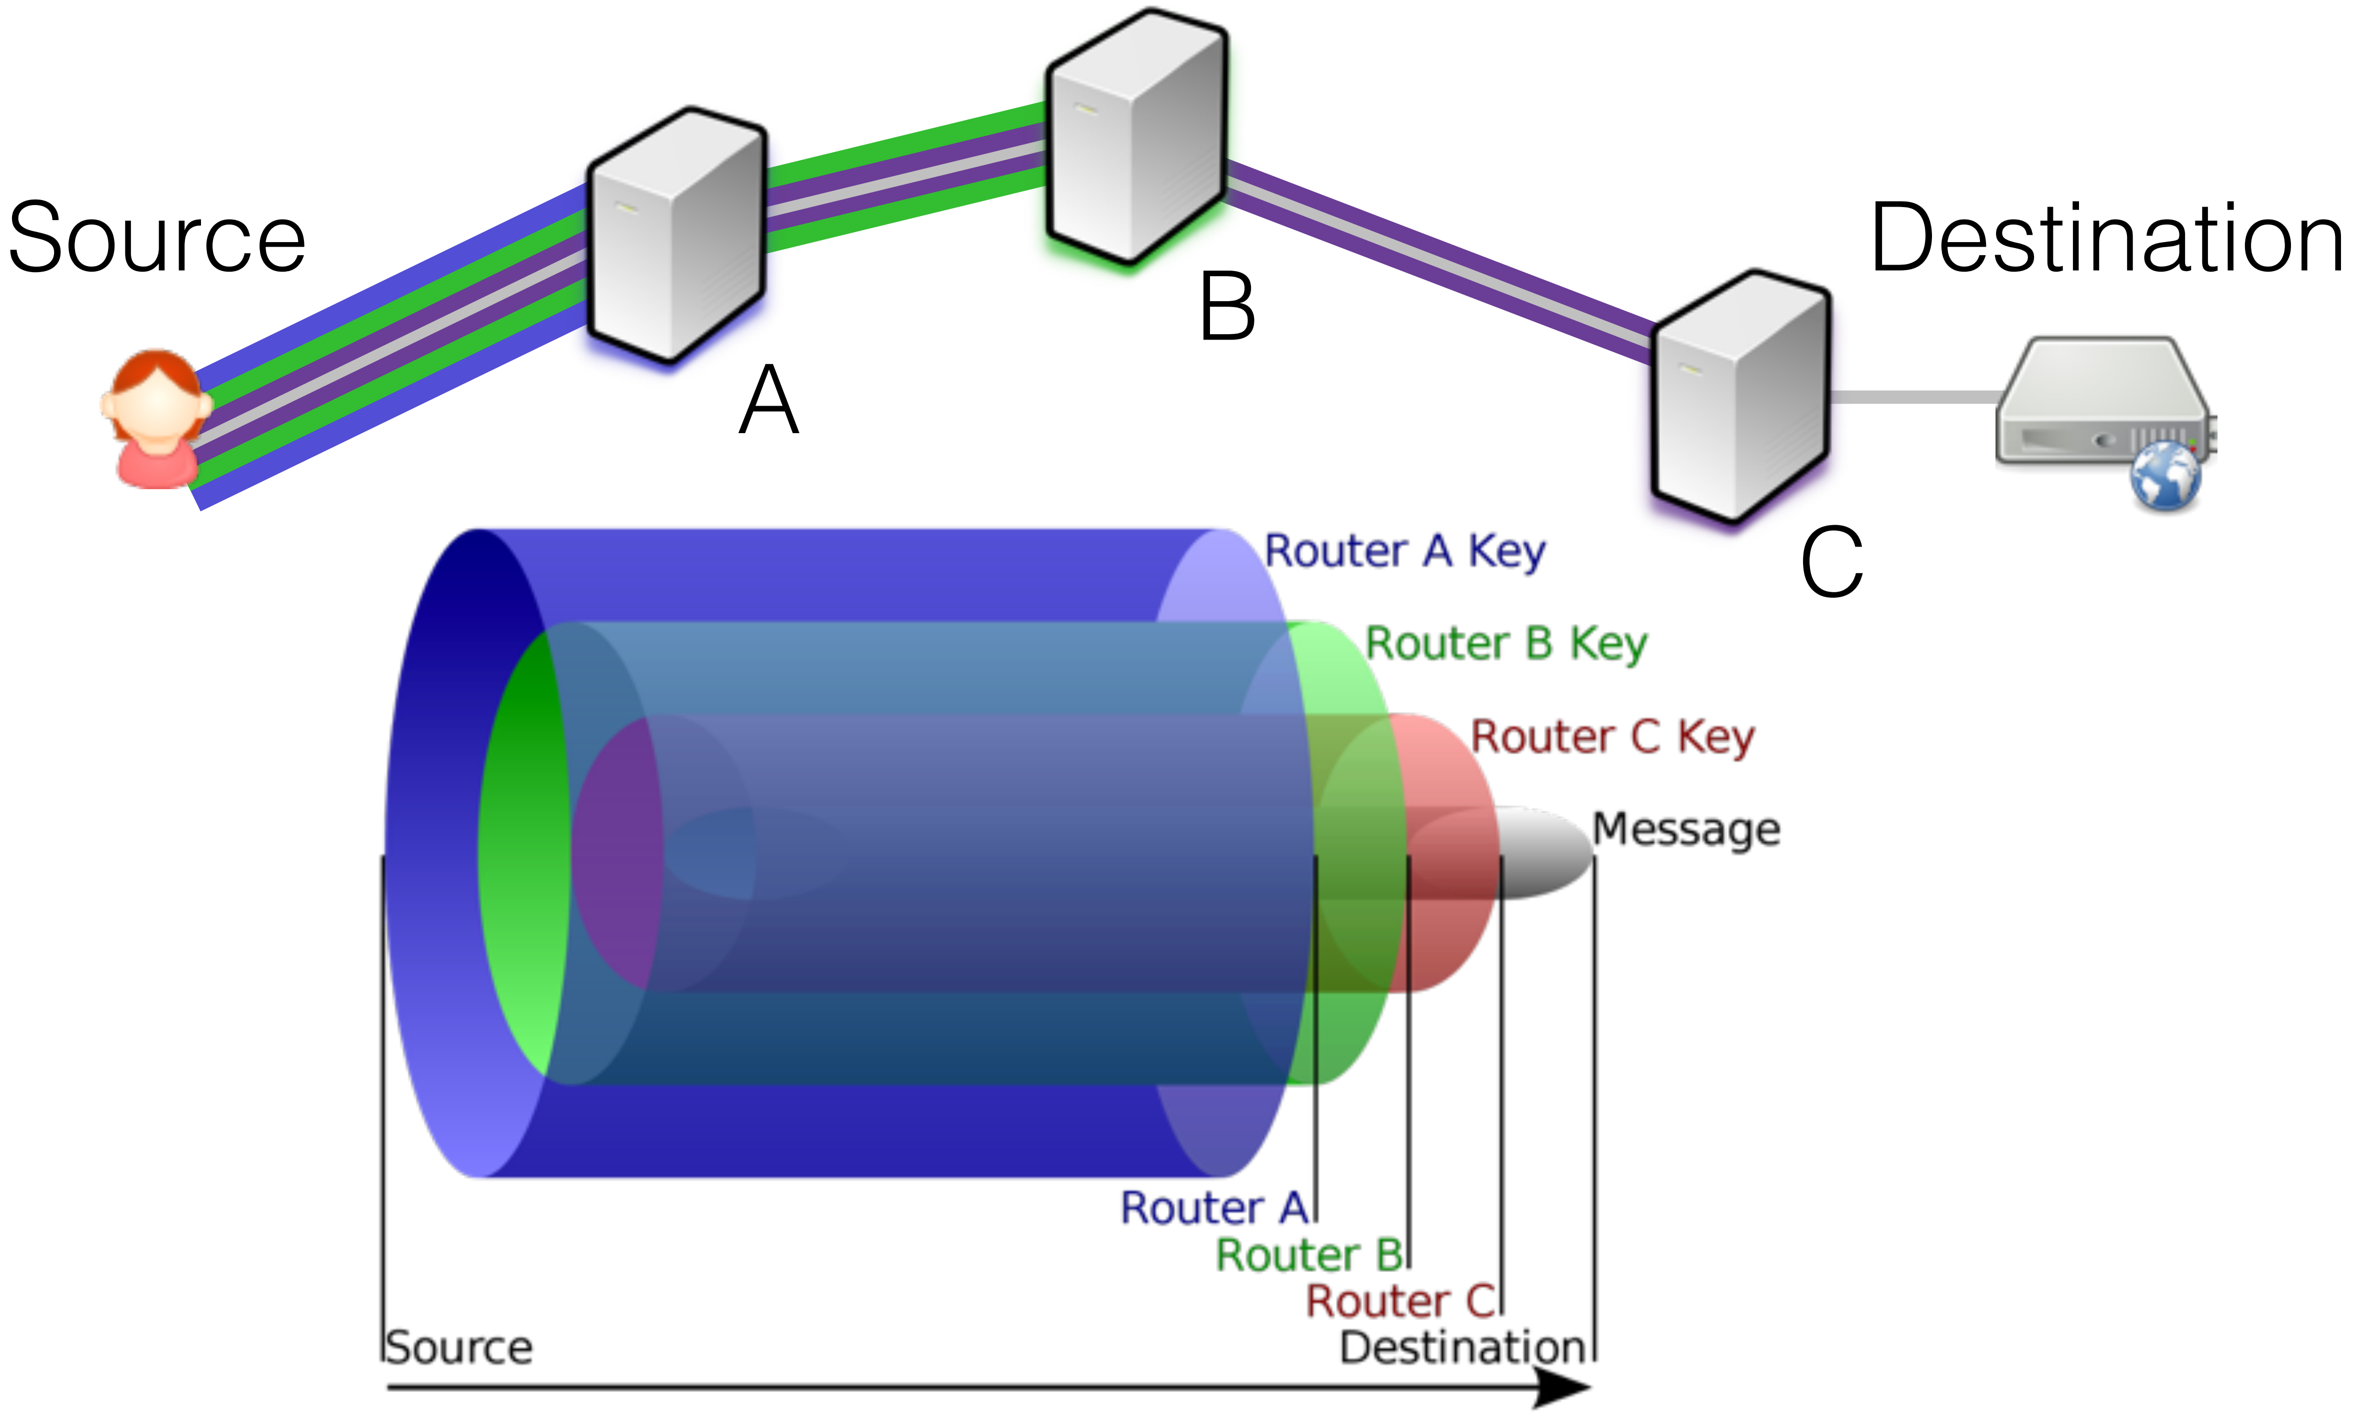

Source: wikimedia.org

Extending the route

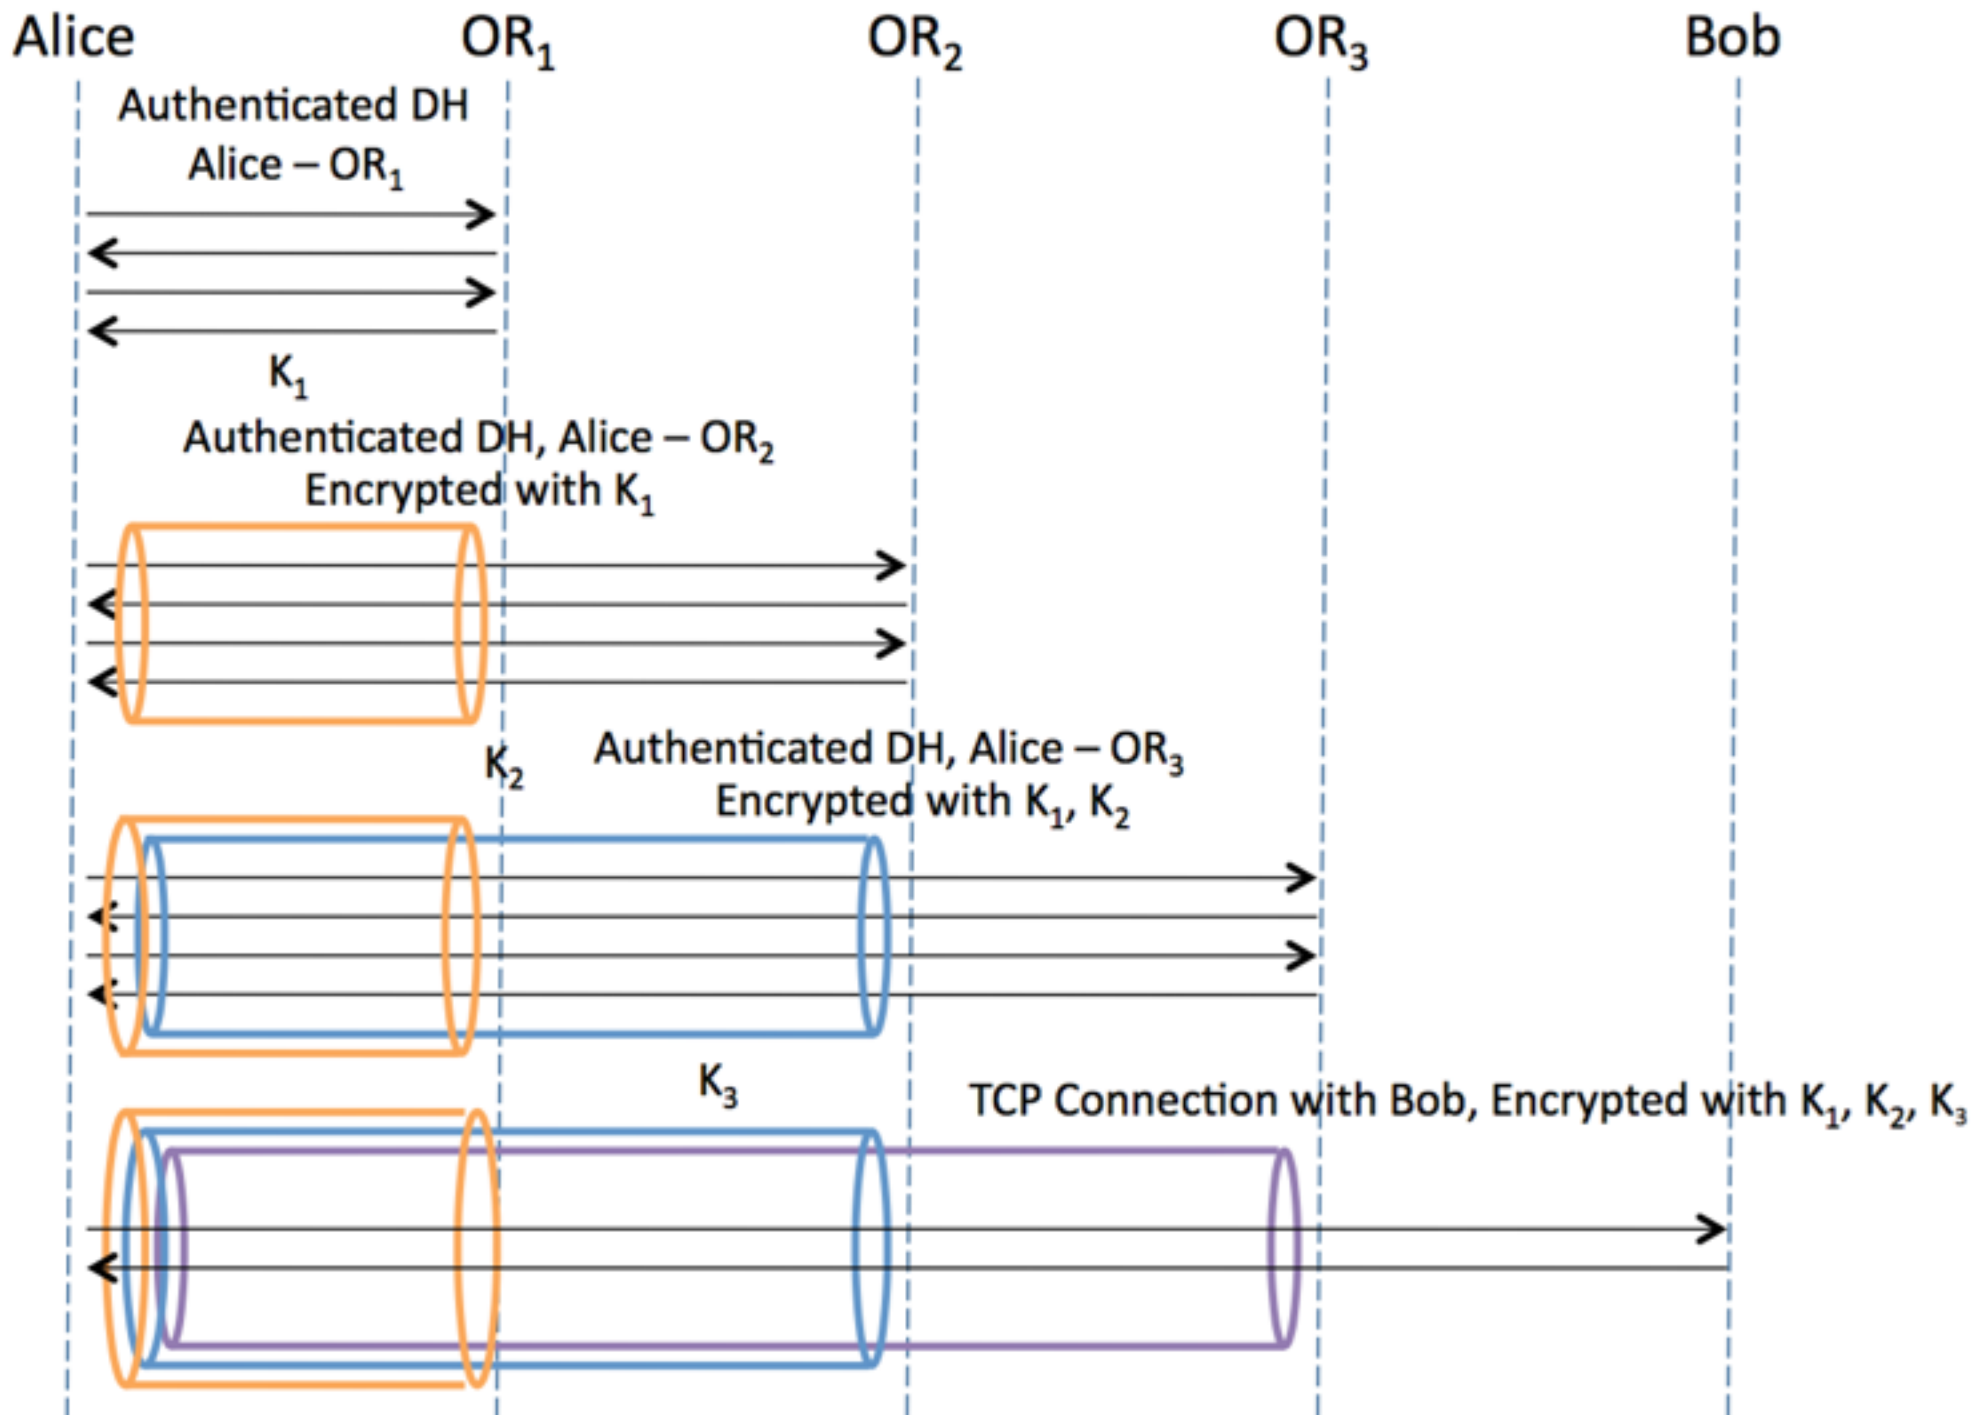
Alice
OR$_1$
OR$_2$
OR$_3$
Bob
Authenticated DH
Alice – OR$_1$
K$_1$
Authenticated DH, Alice – OR$_2$
Encrypted with K$_1$
K$_2$
Authenticated DH, Alice – OR$_3$
Encrypted with K$_1$, K$_2$
K$_3$
TCP Connection with Bob, Encrypted with K$_1$, K$_2$, K$_3$

Outline

1. Introduction to traffic analysis

2. The traffic analysis threat model

3. PETs to protect against traffic analysis.

4. **The Onion Router (Tor)**

5. Traffic analysis attacks and defences in Tor

- Users: currently 2M estimated
- Usability:
  - Easy to install and use.
  - CloudFlare is showing CAPTCHAs.

Directly connecting users

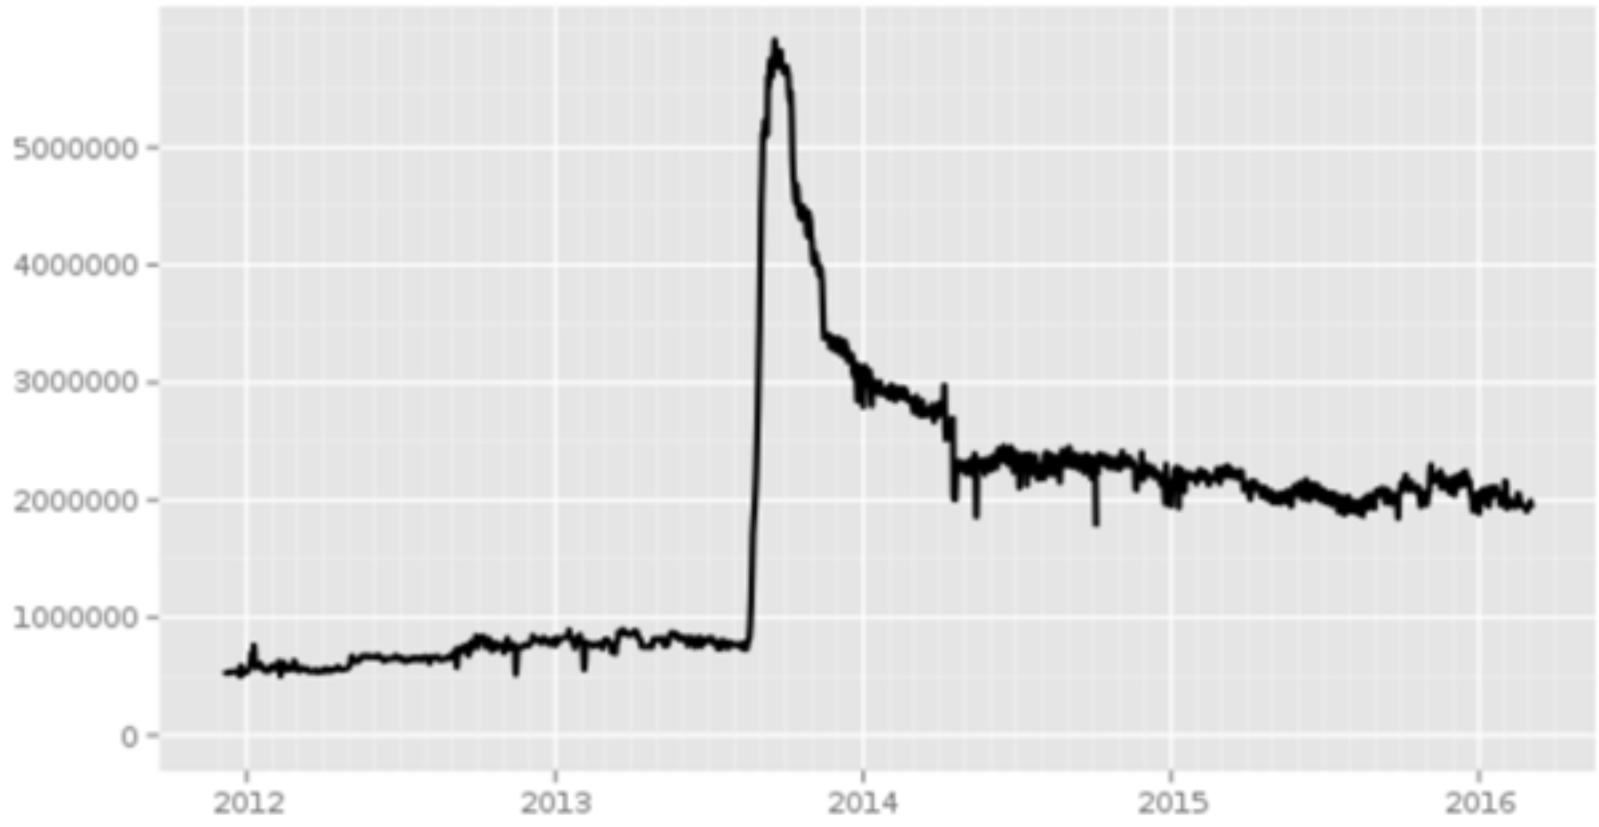

The Tor Project - https://metrics.torproject.org/

Tor volunteer nodes: currently ~7000

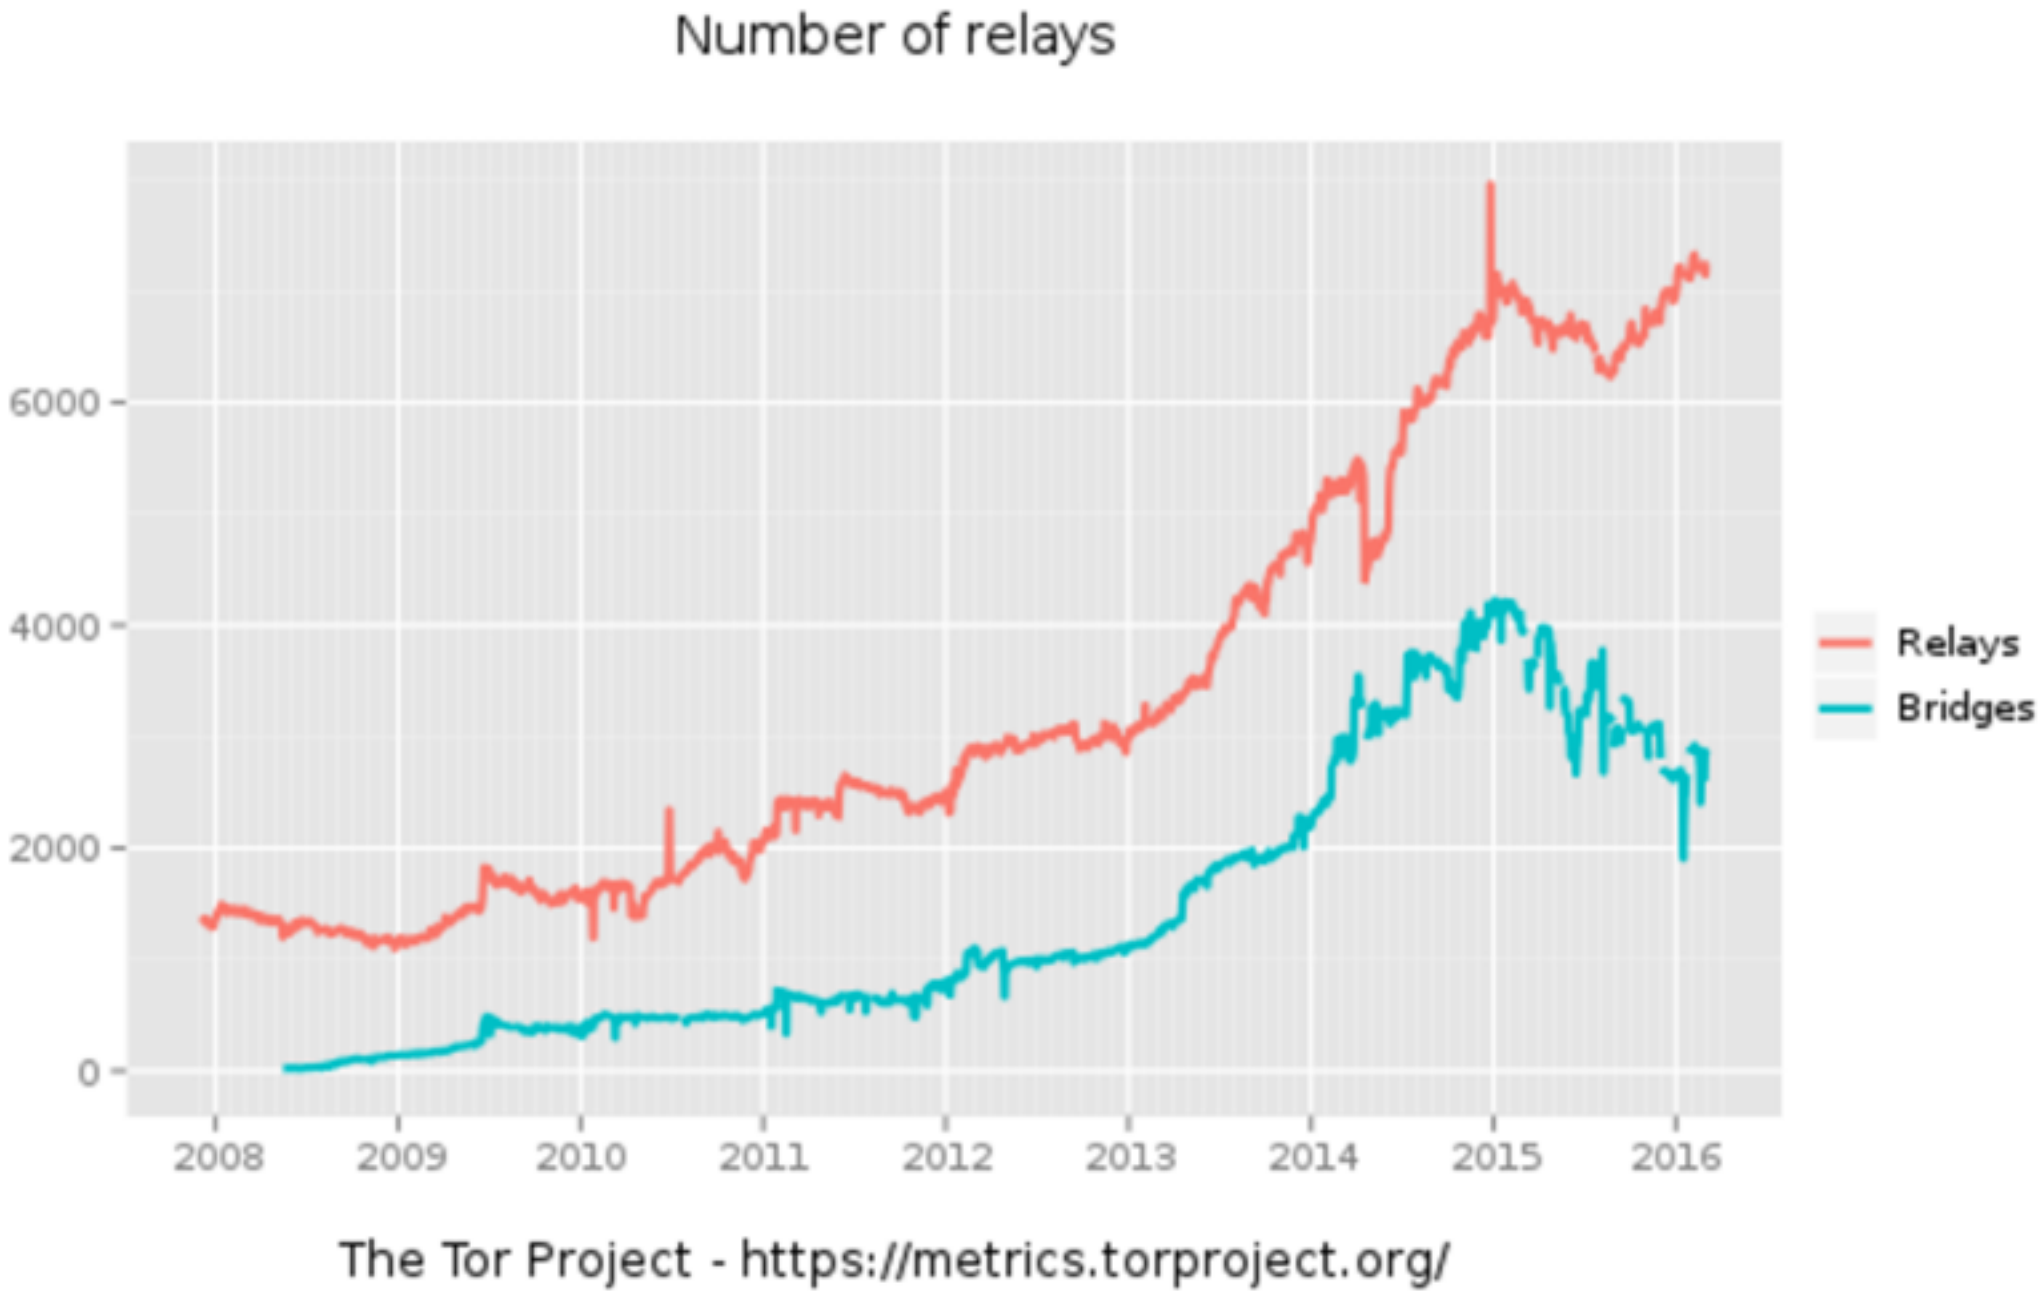

Credit: C. Diaz

Importance of network diversity

- What happens if the first and the last nodes are own by the same entity?
  - The Tor routing policy does not allow choosing more than one node in the same IP subnet,
  - but does not solve other diversity problems.

- Not that slow anymore.

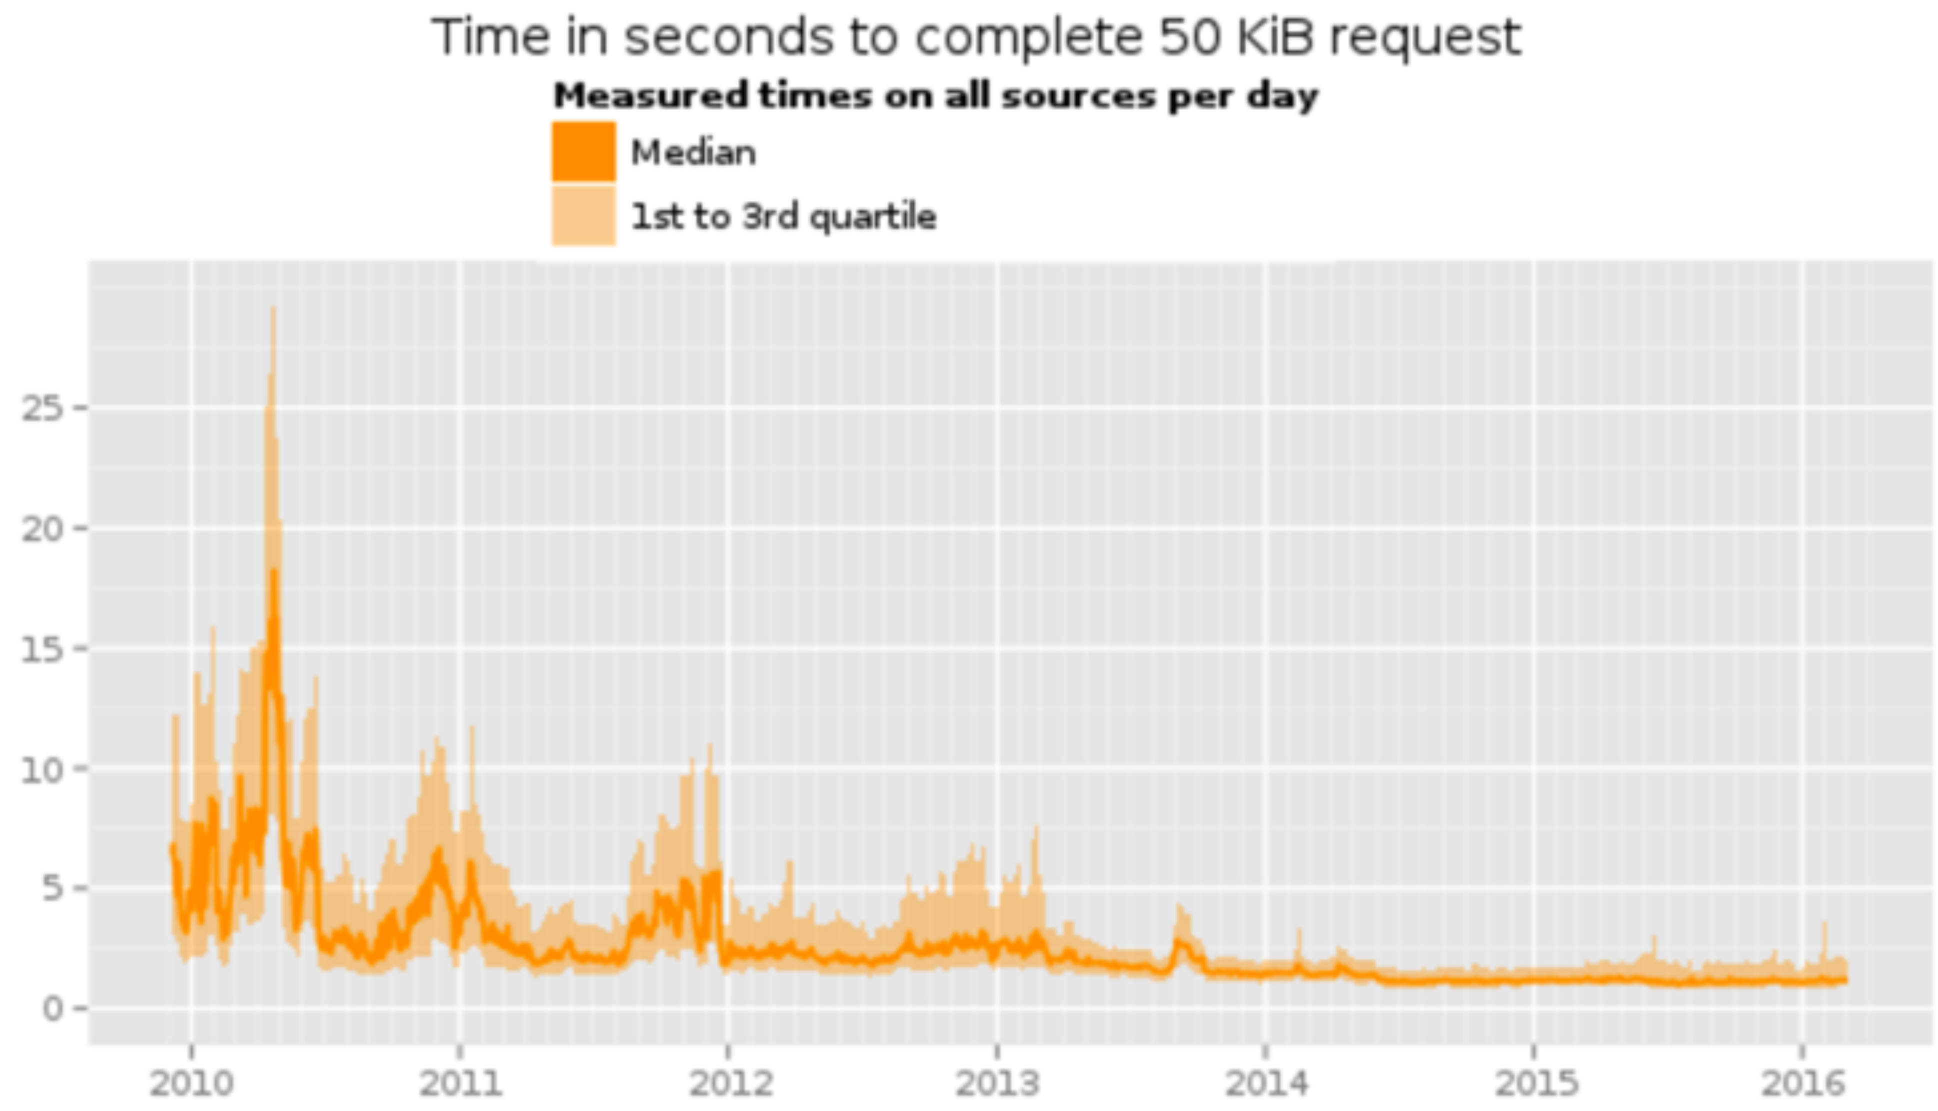

The Tor Project - https://metrics.torproject.org/

Credit: C. Diaz

Basic characteristics (I)

- **Overlay** network, over TCP.
- Fixed-size message: **cells**.
- Clients select a set of routers that constitute the anonymous channel (aka **circuit**).
- Special nodes to enter the network: **entry guards**.
- Client fetches **Directory Authorities** to obtain a list of relays.

Basic characteristics (II)

- Also anonymity for servers: Onion Services (aka Hidden Services). Some examples:
  - Silk Road (taken down by FBI).
  - New Yorker Strongbox
  - Facebook
- Censorship circumvention: Bridges and Pluggable Transports.
- Ecosystem: torsocks, arm, stem, tails, exitmap, exonerator, gettor, torflow, tor2web, orbot, and so on.

Some attacks against Tor

- Selective DoS.
- Identifying users: e.g., browser fingerprinting.
- Linking users with pages: e.g., website fingerprinting.
- Identifying Hidden Services: e.g., guard discovery.
- Others…

Selective Denial of Service (DoS) to break anonymity

- DoS good nodes or paths.
  - Force reconstruction of paths until you control all nodes in the circuit path (or entry+exit).
- Paths are either fully honest or corrupted.
- Probabilities of compromise?

Outline

1. Introduction to traffic analysis

2. The traffic analysis threat model

3. PETs to protect against traffic analysis.

4. The Onion Router (Tor)

5. **Traffic analysis attacks and defences in Tor**

Browser Fingerprinting in Tor

Browser Fingerprinting

- Browser plugins, fonts, clock skew, canvas, JS engine performance and so on.
- Panopticlick showed the feasibility of it.
- Alternative to cookies: not subject to deletion or expiration, hard to spoof and notice.
- "it's just about devices, not personally identifying information (PII)"

Credit: G. Acar

JS-based Fingerprinting

- Fallback when Flash is not available.
- Featuring font probing.

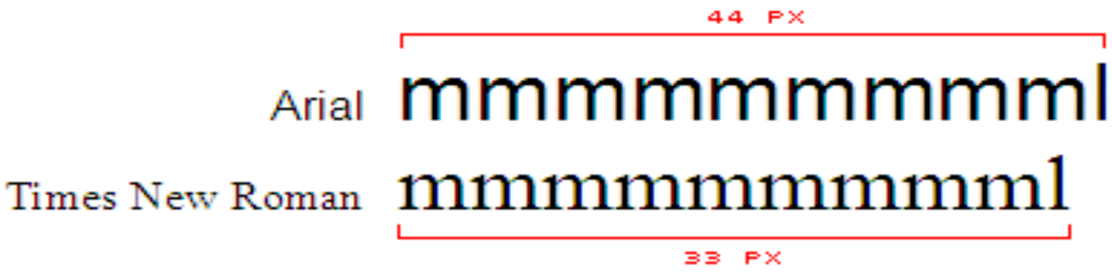

  - Measure dimensions of invisible string with different fonts.
  - Compare the size to fallback font.
  - DEMO.

Credit: G. Acar

Canvas Fingerprinting

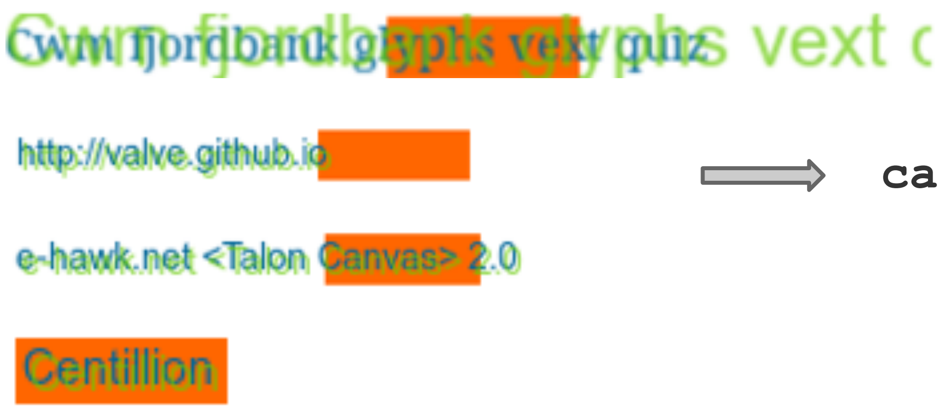

⟹ `canvas.toDataURL()` ⟹

data:image/png;base64,iVBORw0KGgoAAAANSUhEUgAAAAUAAAAFCAYAAACNbyblAAAAHElEQVQI12P4//8/w38GIAXDIBKE0DHxgljNBAAO9TXL0Y4OHwAAAABJRU5ErkJggg==

- Acar et al.: We found 8 providers including a very popular 3rd party script: AddThis.

K. Mowery and H. Shacham. **Pixel perfect: Fingerprinting canvas in HTML5**. In Proceedings of W2SP 2012, IEEE Computer Society, 2012.

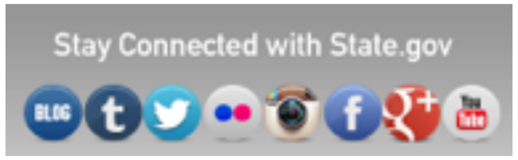

- DEMO.

Credit: G. Acar

Countermeasures in Tor Browser

- Anonymity set fingerprint.
- Limits the number of fonts a site can load.
- Alert resize window in Tor Browser.
- Alert when accessing to Canvas and return blank.
- New attack vectors.

Credit: G. Acar

Website Fingerprinting in Tor

Website Fingerprinting

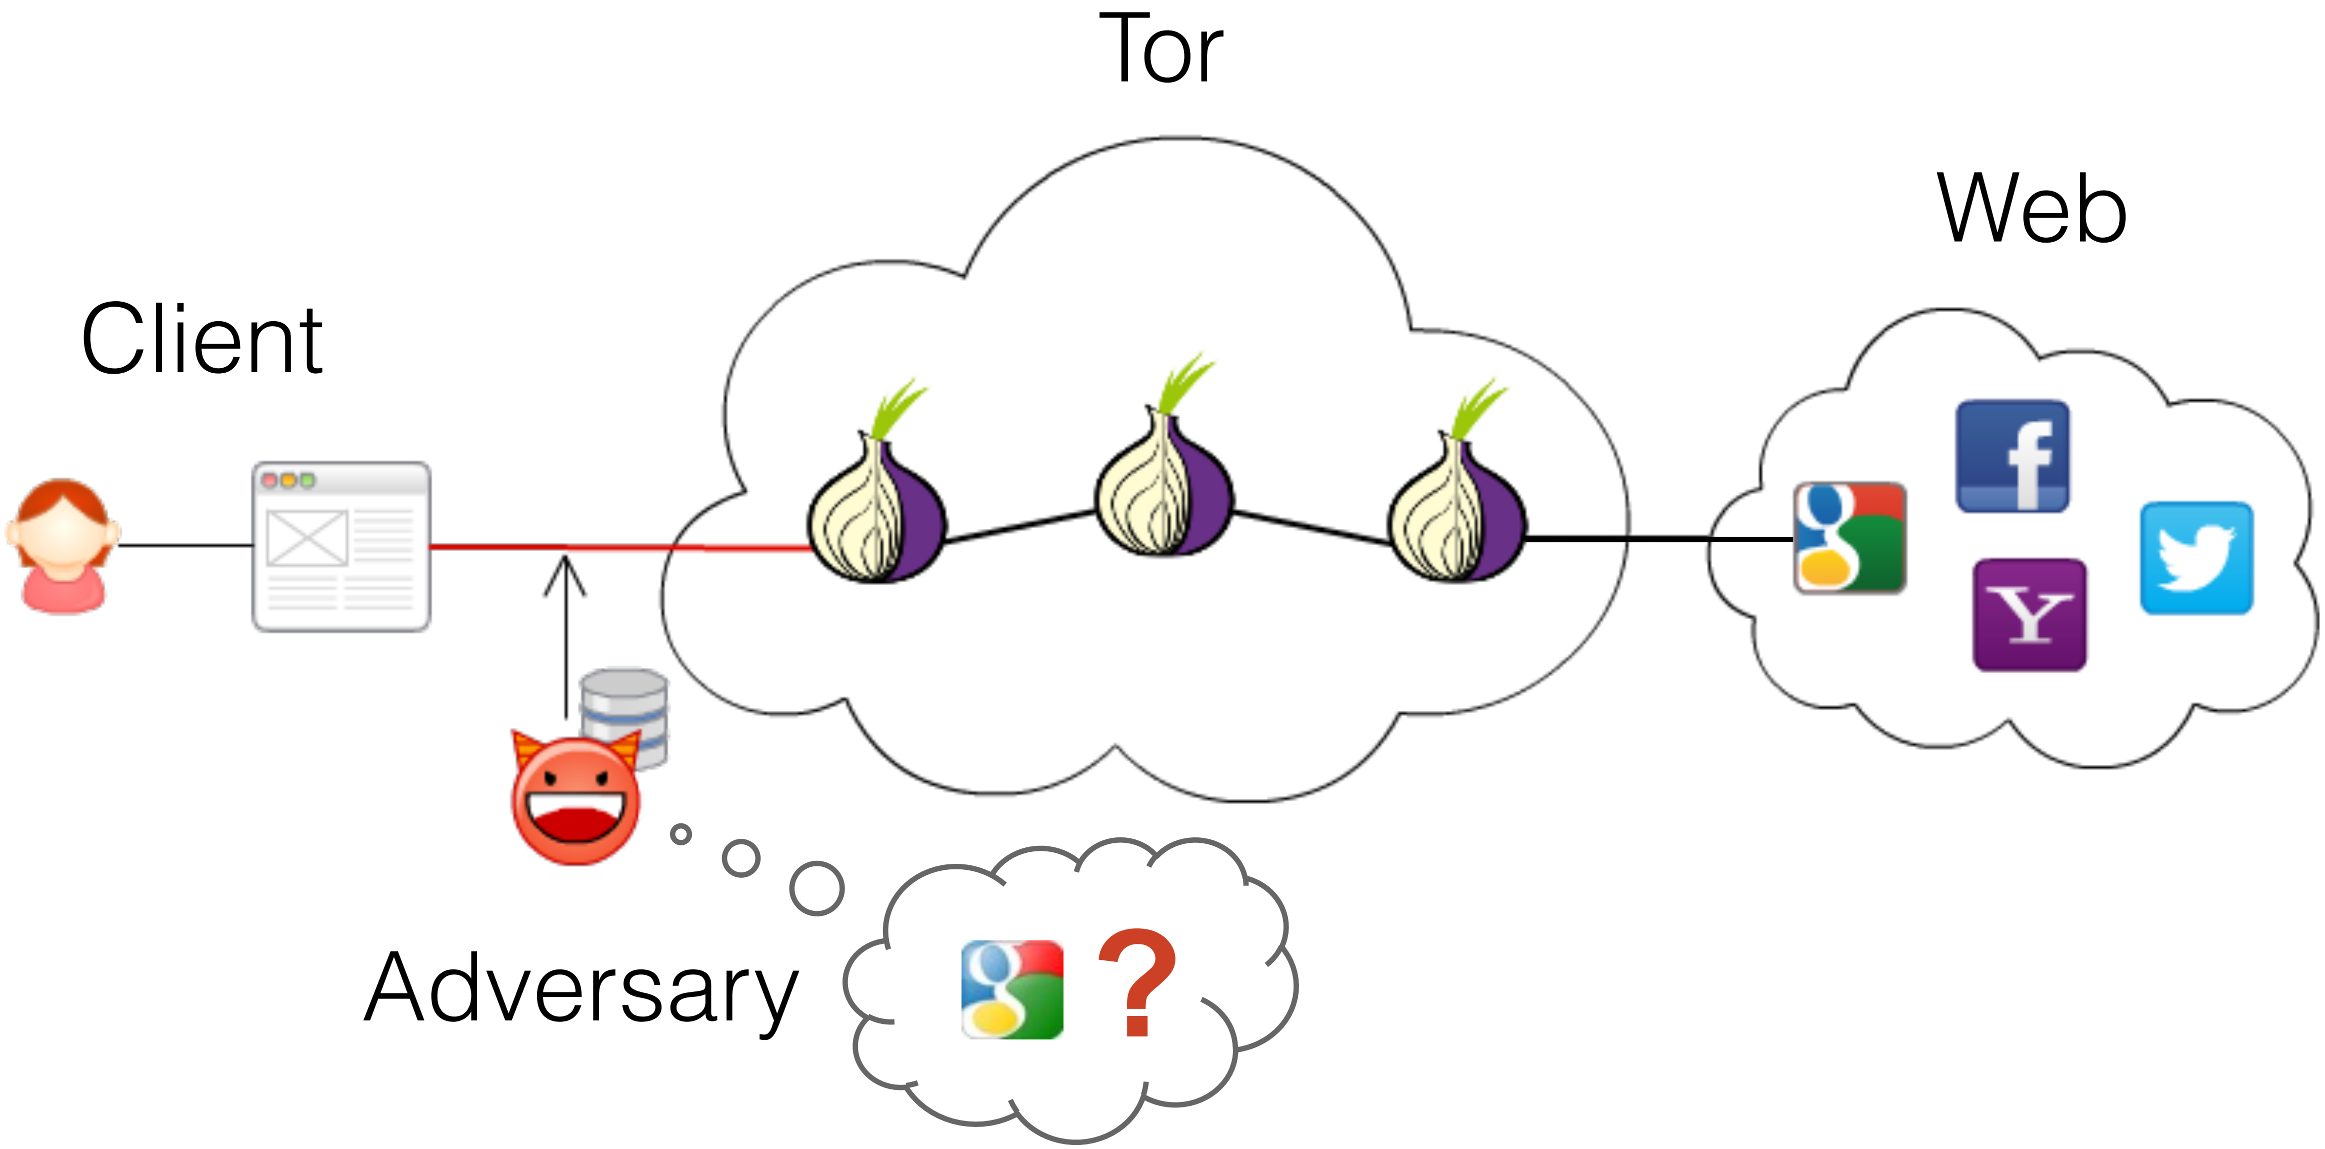

Closed vs. Open world

- Early prior WF works considered closed world of pages: client can only visit *monitored* pages.
- In practice: extremely large universe of web pages.
- How likely is the user (a priori) to visit a monitored web page?
- If the prior is not a good estimate: *base rate fallacy*.

The base rate fallacy

- A new ‘terrorist detection’ system in a city:
  - 0.88 detects true terrorists (true positive rate).
  - 0.05 false positive rate.
- Eve is detected by the system as a terrorist:
  - What is the probability that she is a terrorist?
  - Is it 0.88? 0.95? Something in between?

The base rate fallacy

- A new ‘terrorist detection’ system in a city:
  - 0.88 detects true terrorists (true positive rate).
  - 0.05 false positive rate.
- Eve is detected by a terrorist:
  - What is the proba is a terrorist?
  - Is it 0.88? 0.95? Something in between?

Only 0.1!!

The Base Rate Fallacy

- The circumference represents the population of the city.

- Each dot represents an inhabitant.

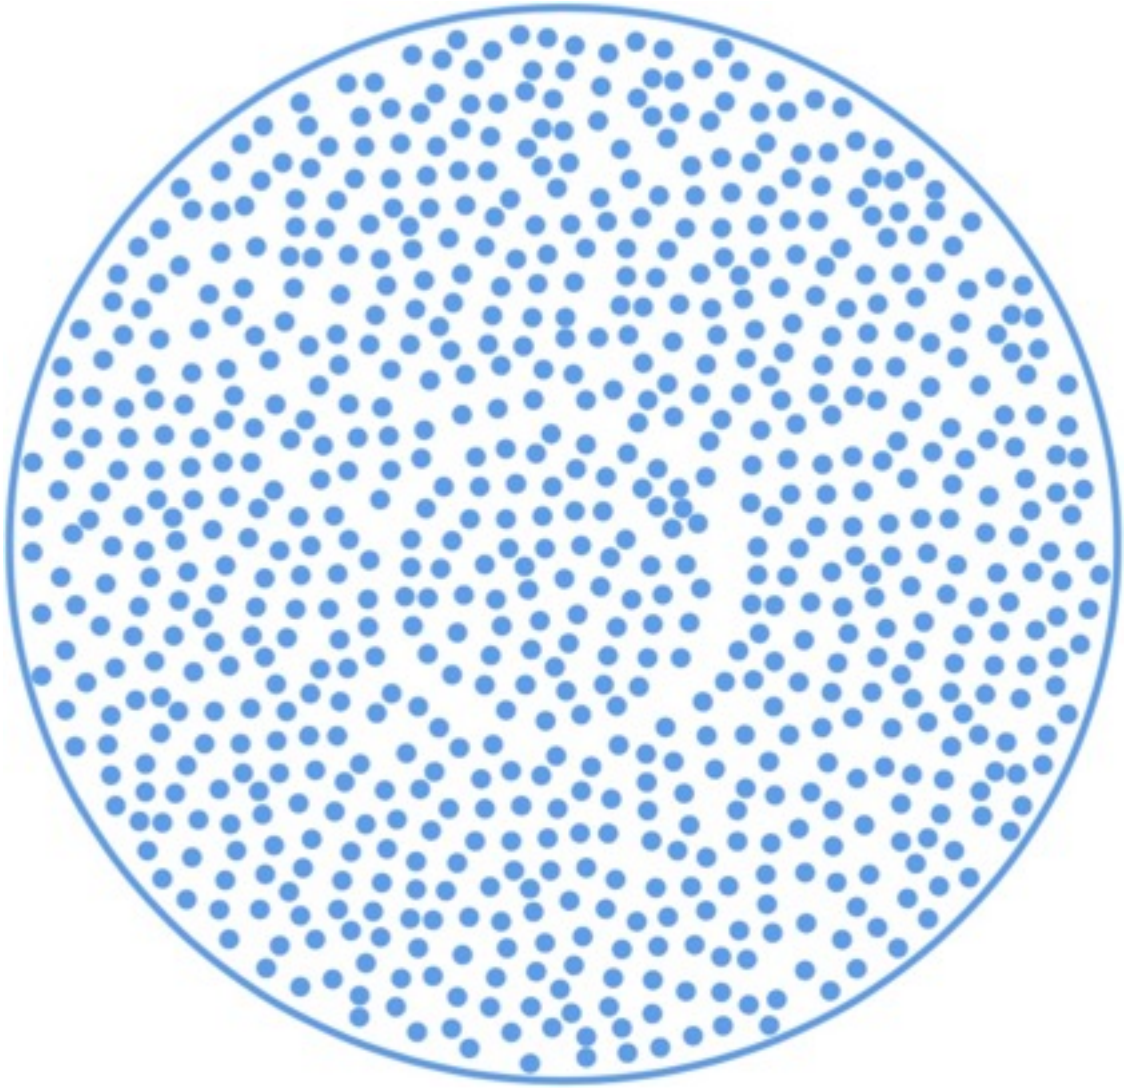

The Base Rate Fallacy

- 1% of the people are terrorists (base rate or prior).

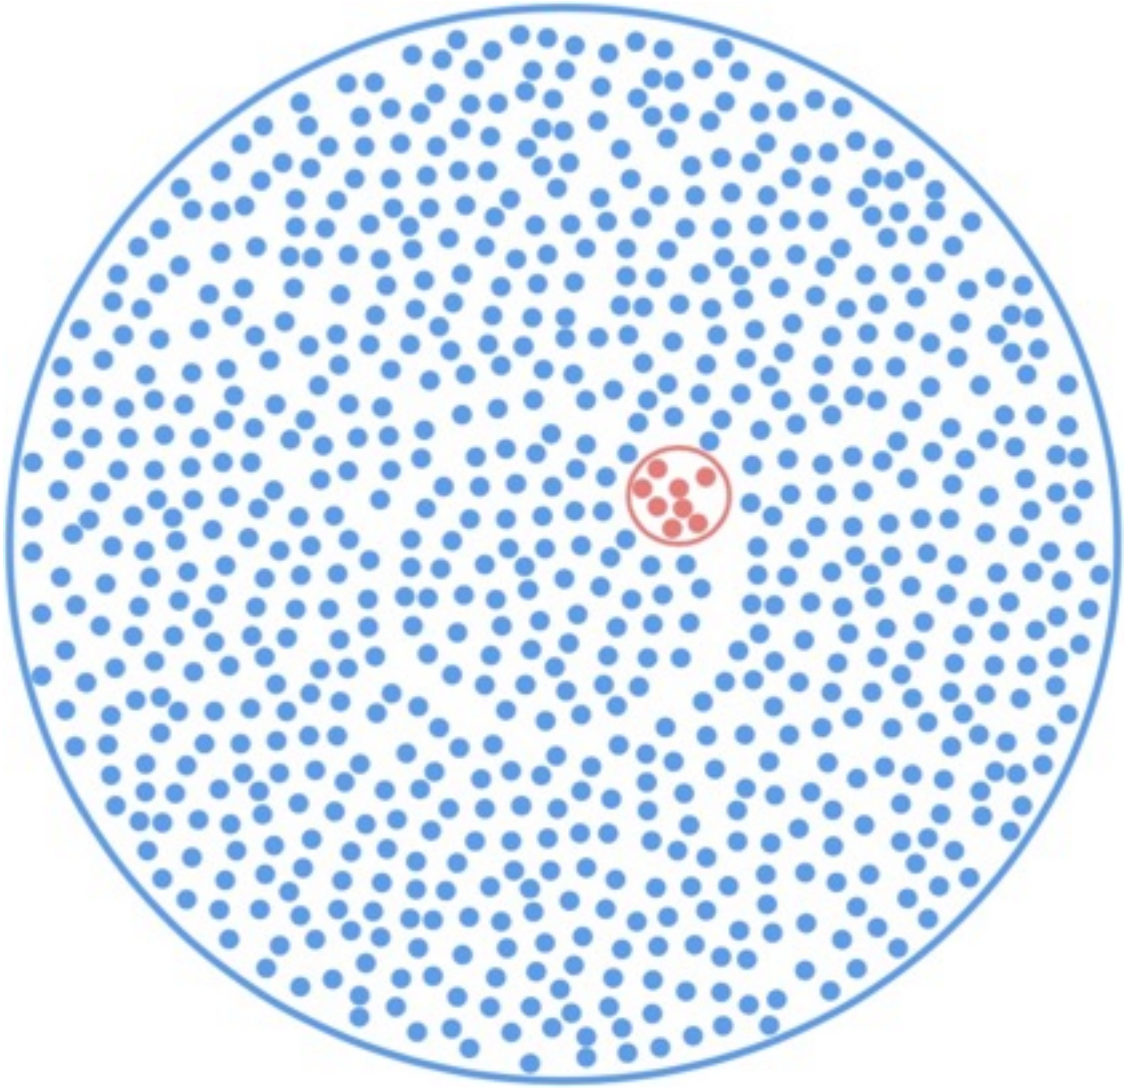

The Base Rate Fallacy

- From the terrorists, 88% are identified as terrorists by the detection system.

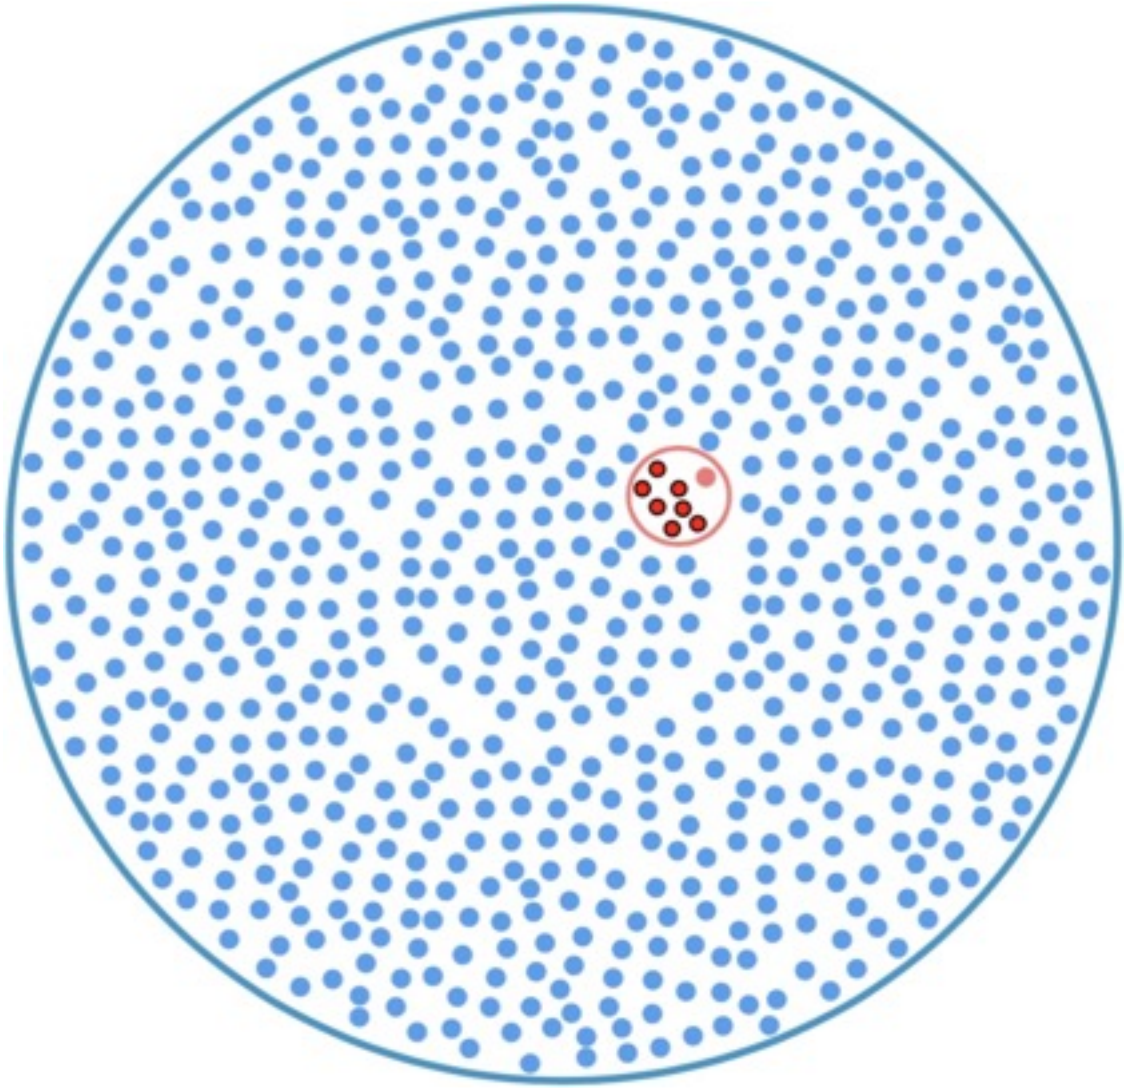

The Base Rate Fallacy

- From the non-terrorists, 5% are erroneously identified as terrorists.

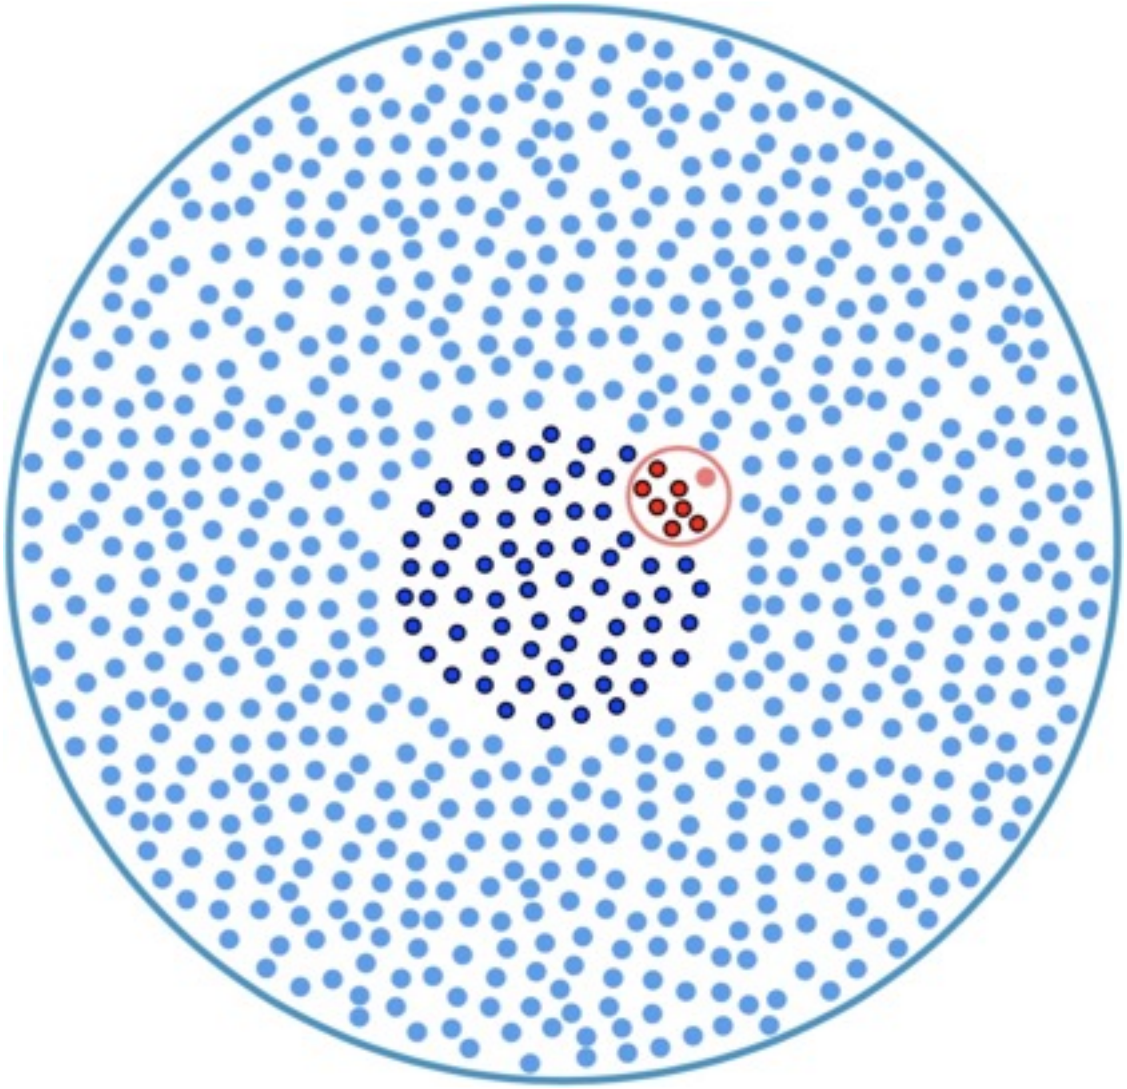

The Base Rate Fallacy

- Eve must be within the black circumference.
- Ratio of red dots within the black circumference:

Red / Total = 7 / 70 = 0.1 !

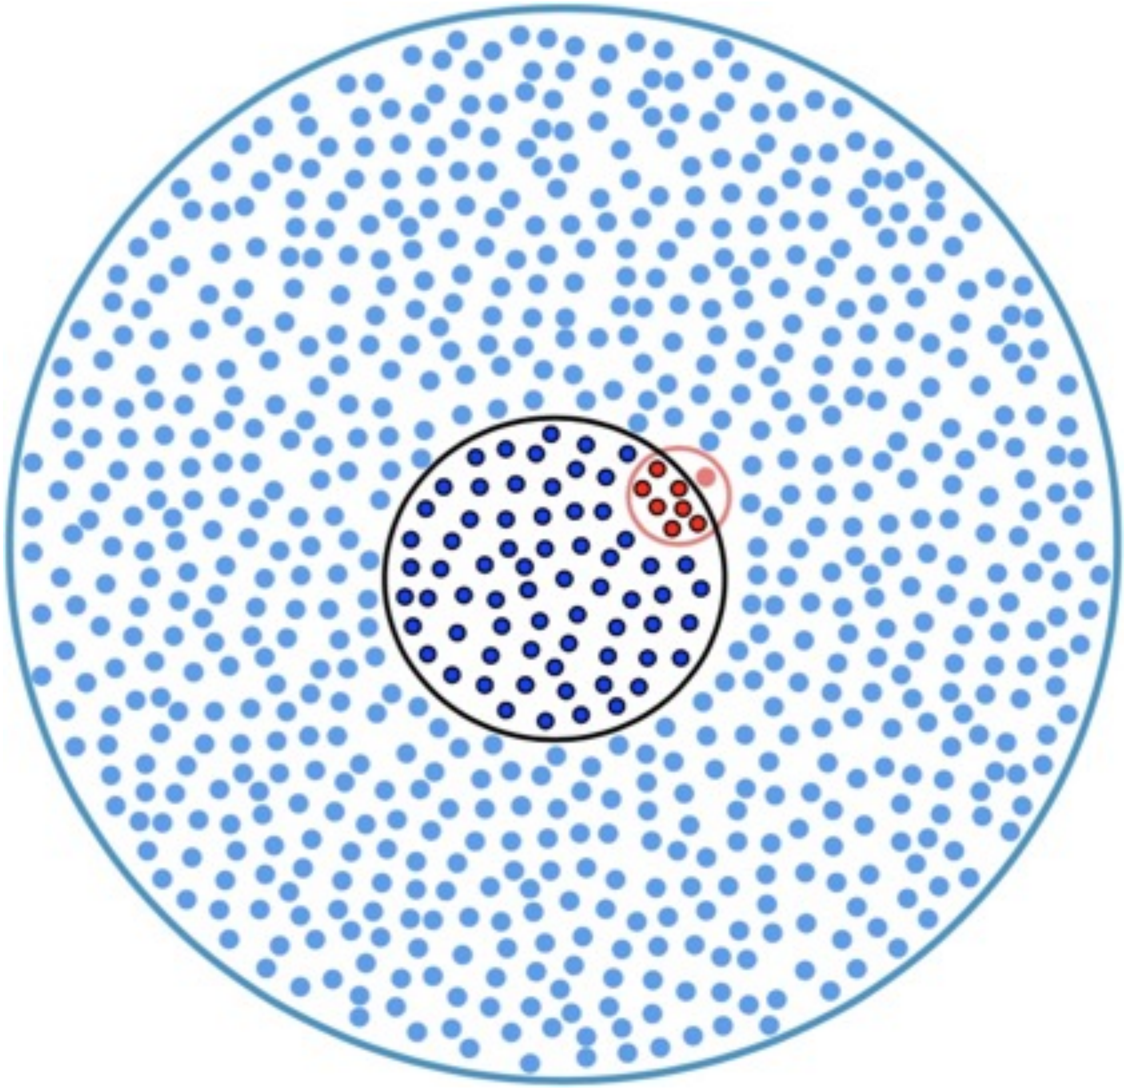

Base Rate Fallacy in WF

- What is the probability of visiting a *monitored* site?
- For non-popular pages, WF is ineffective.
- If the attacker has some background knowledge (targeted attack), then the attack might be very effective.

Attacks on Onion Services

Tor Hidden Services (HS)

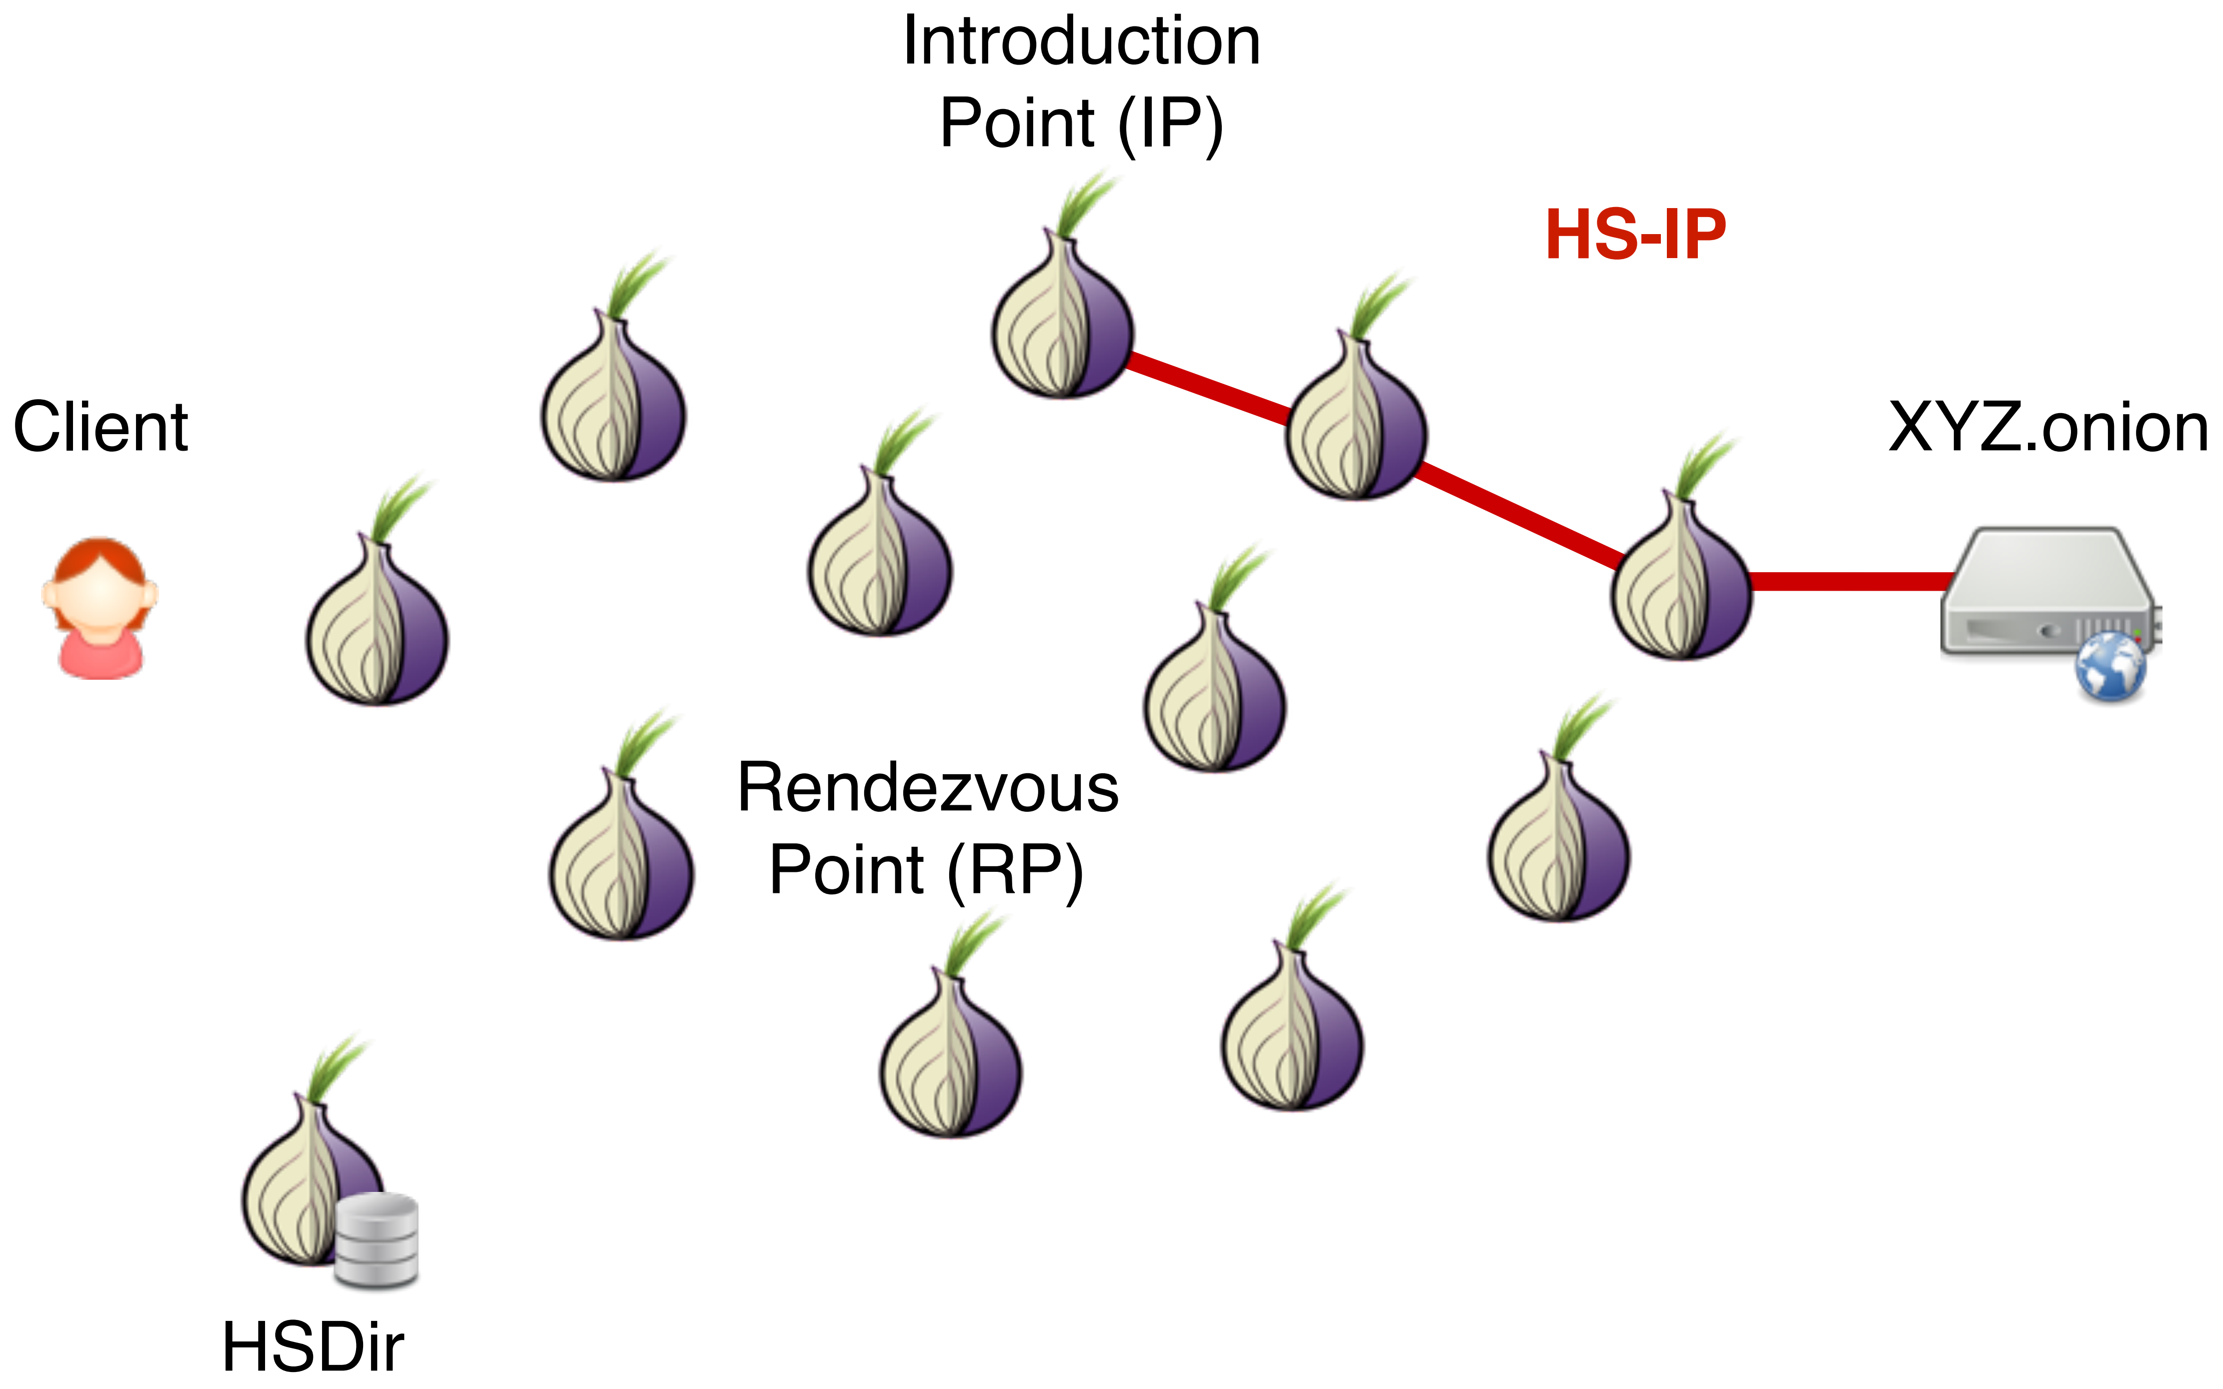

Tor Hidden Services (HS)

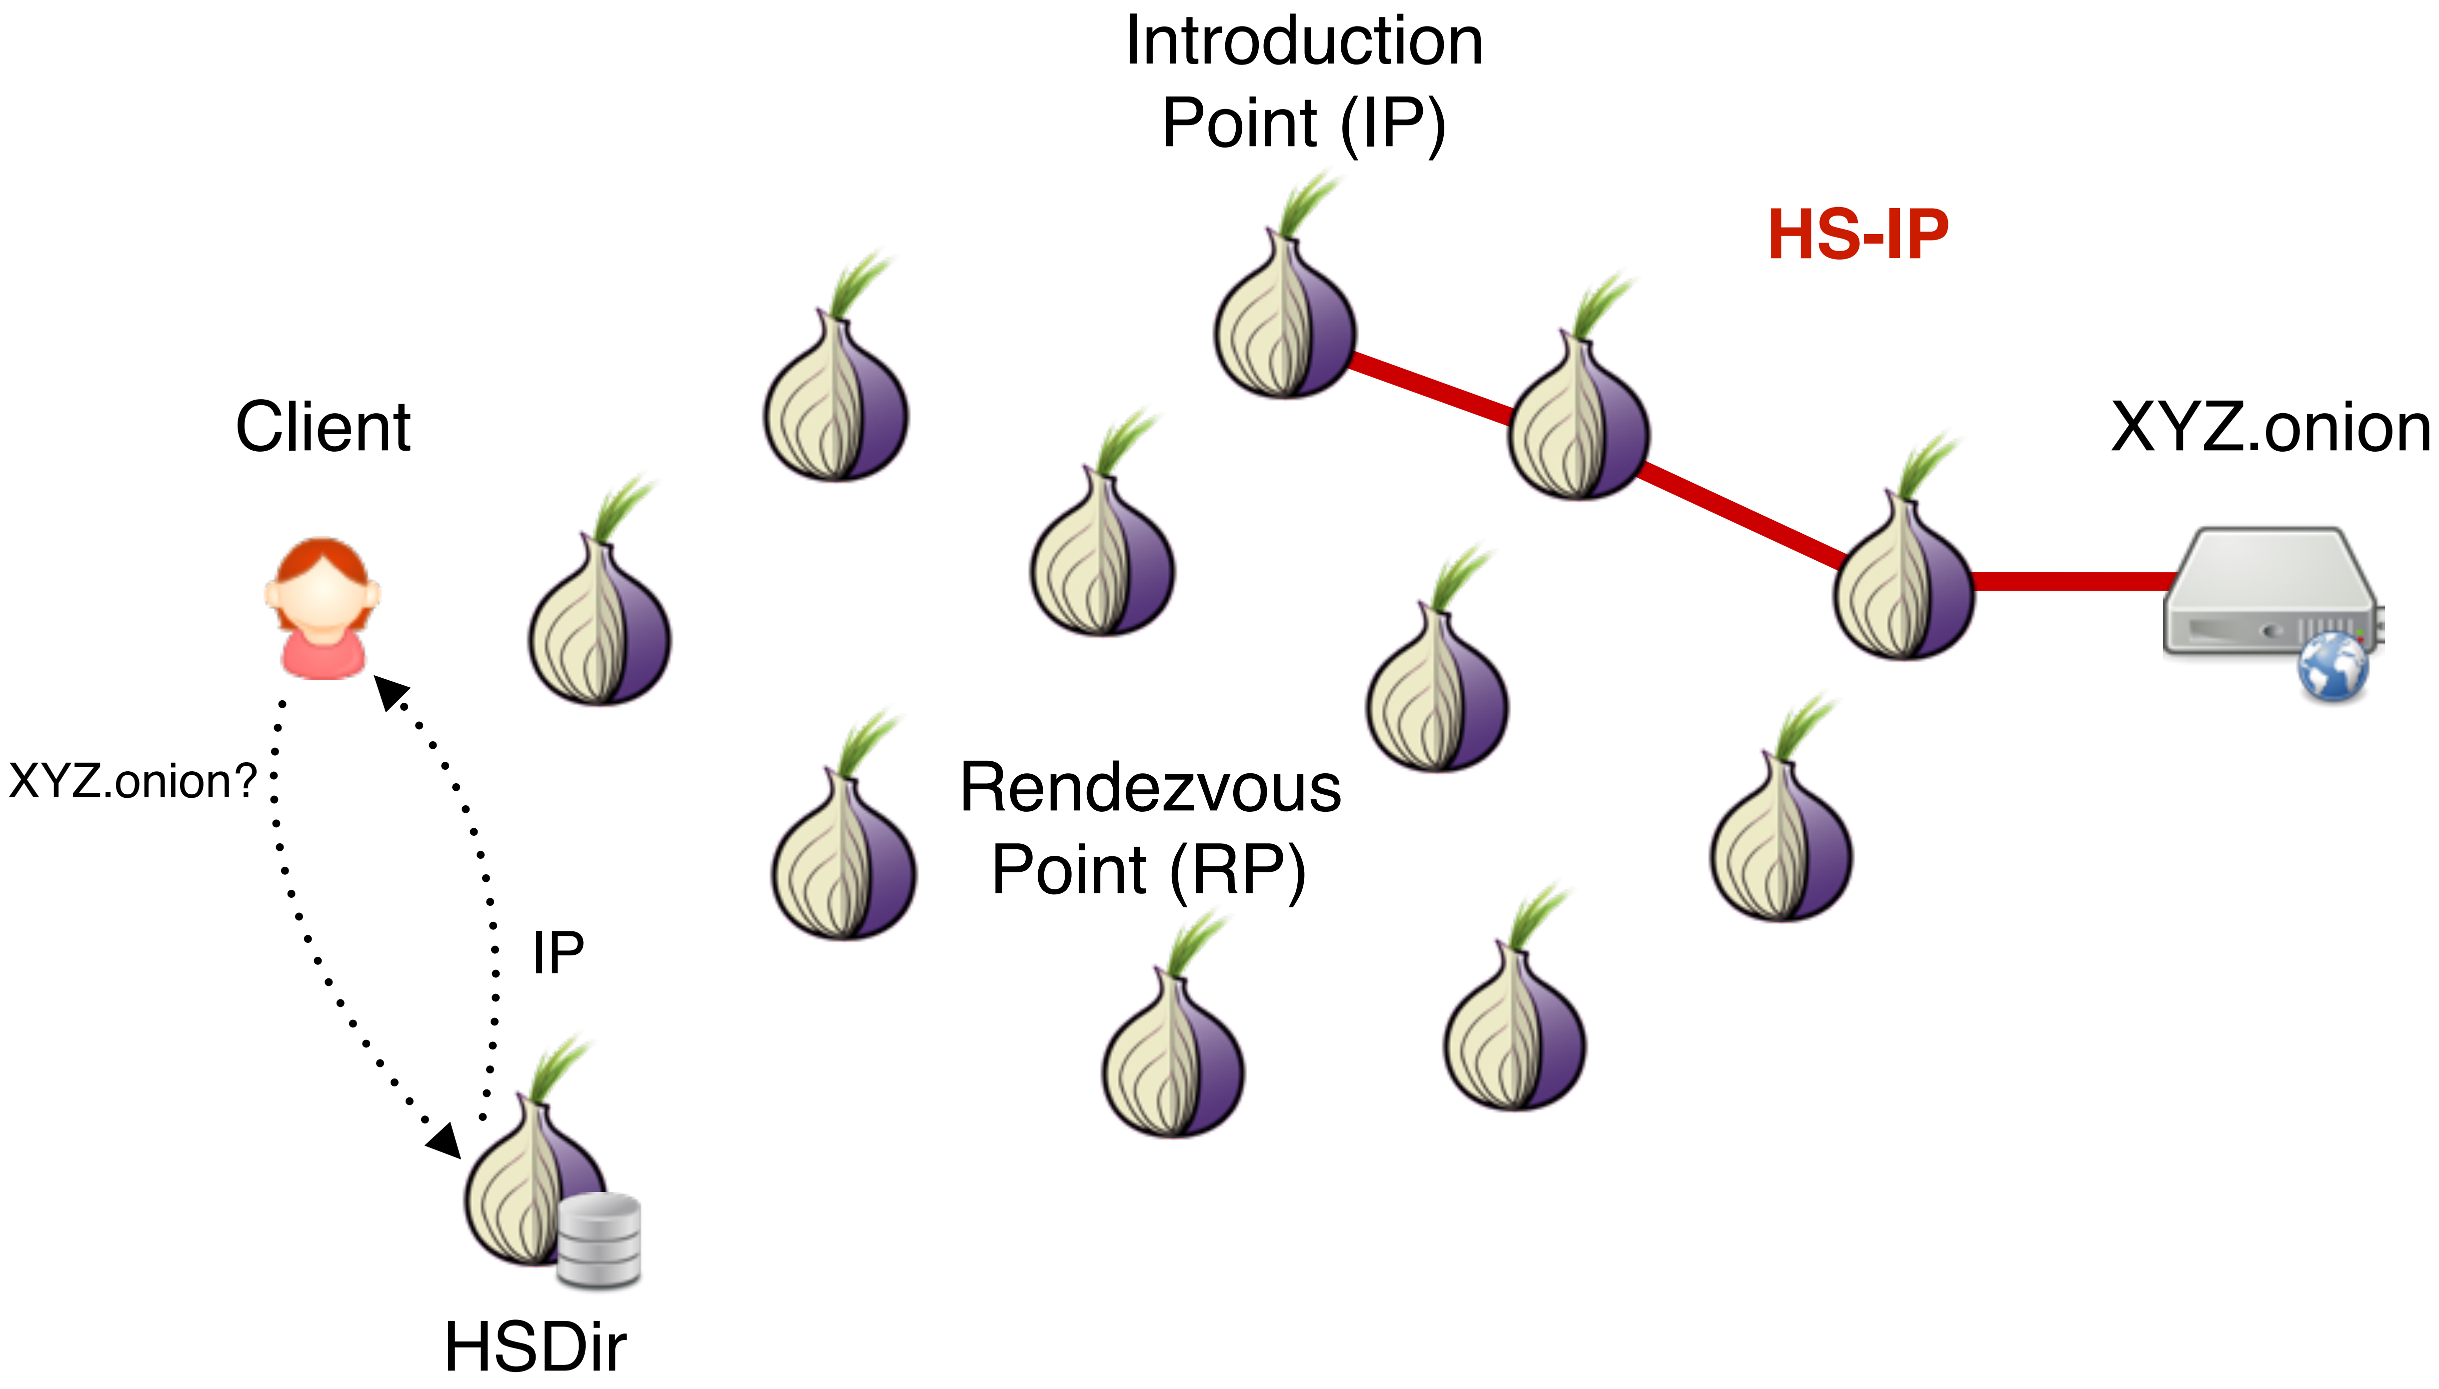

Tor Hidden Services (HS)

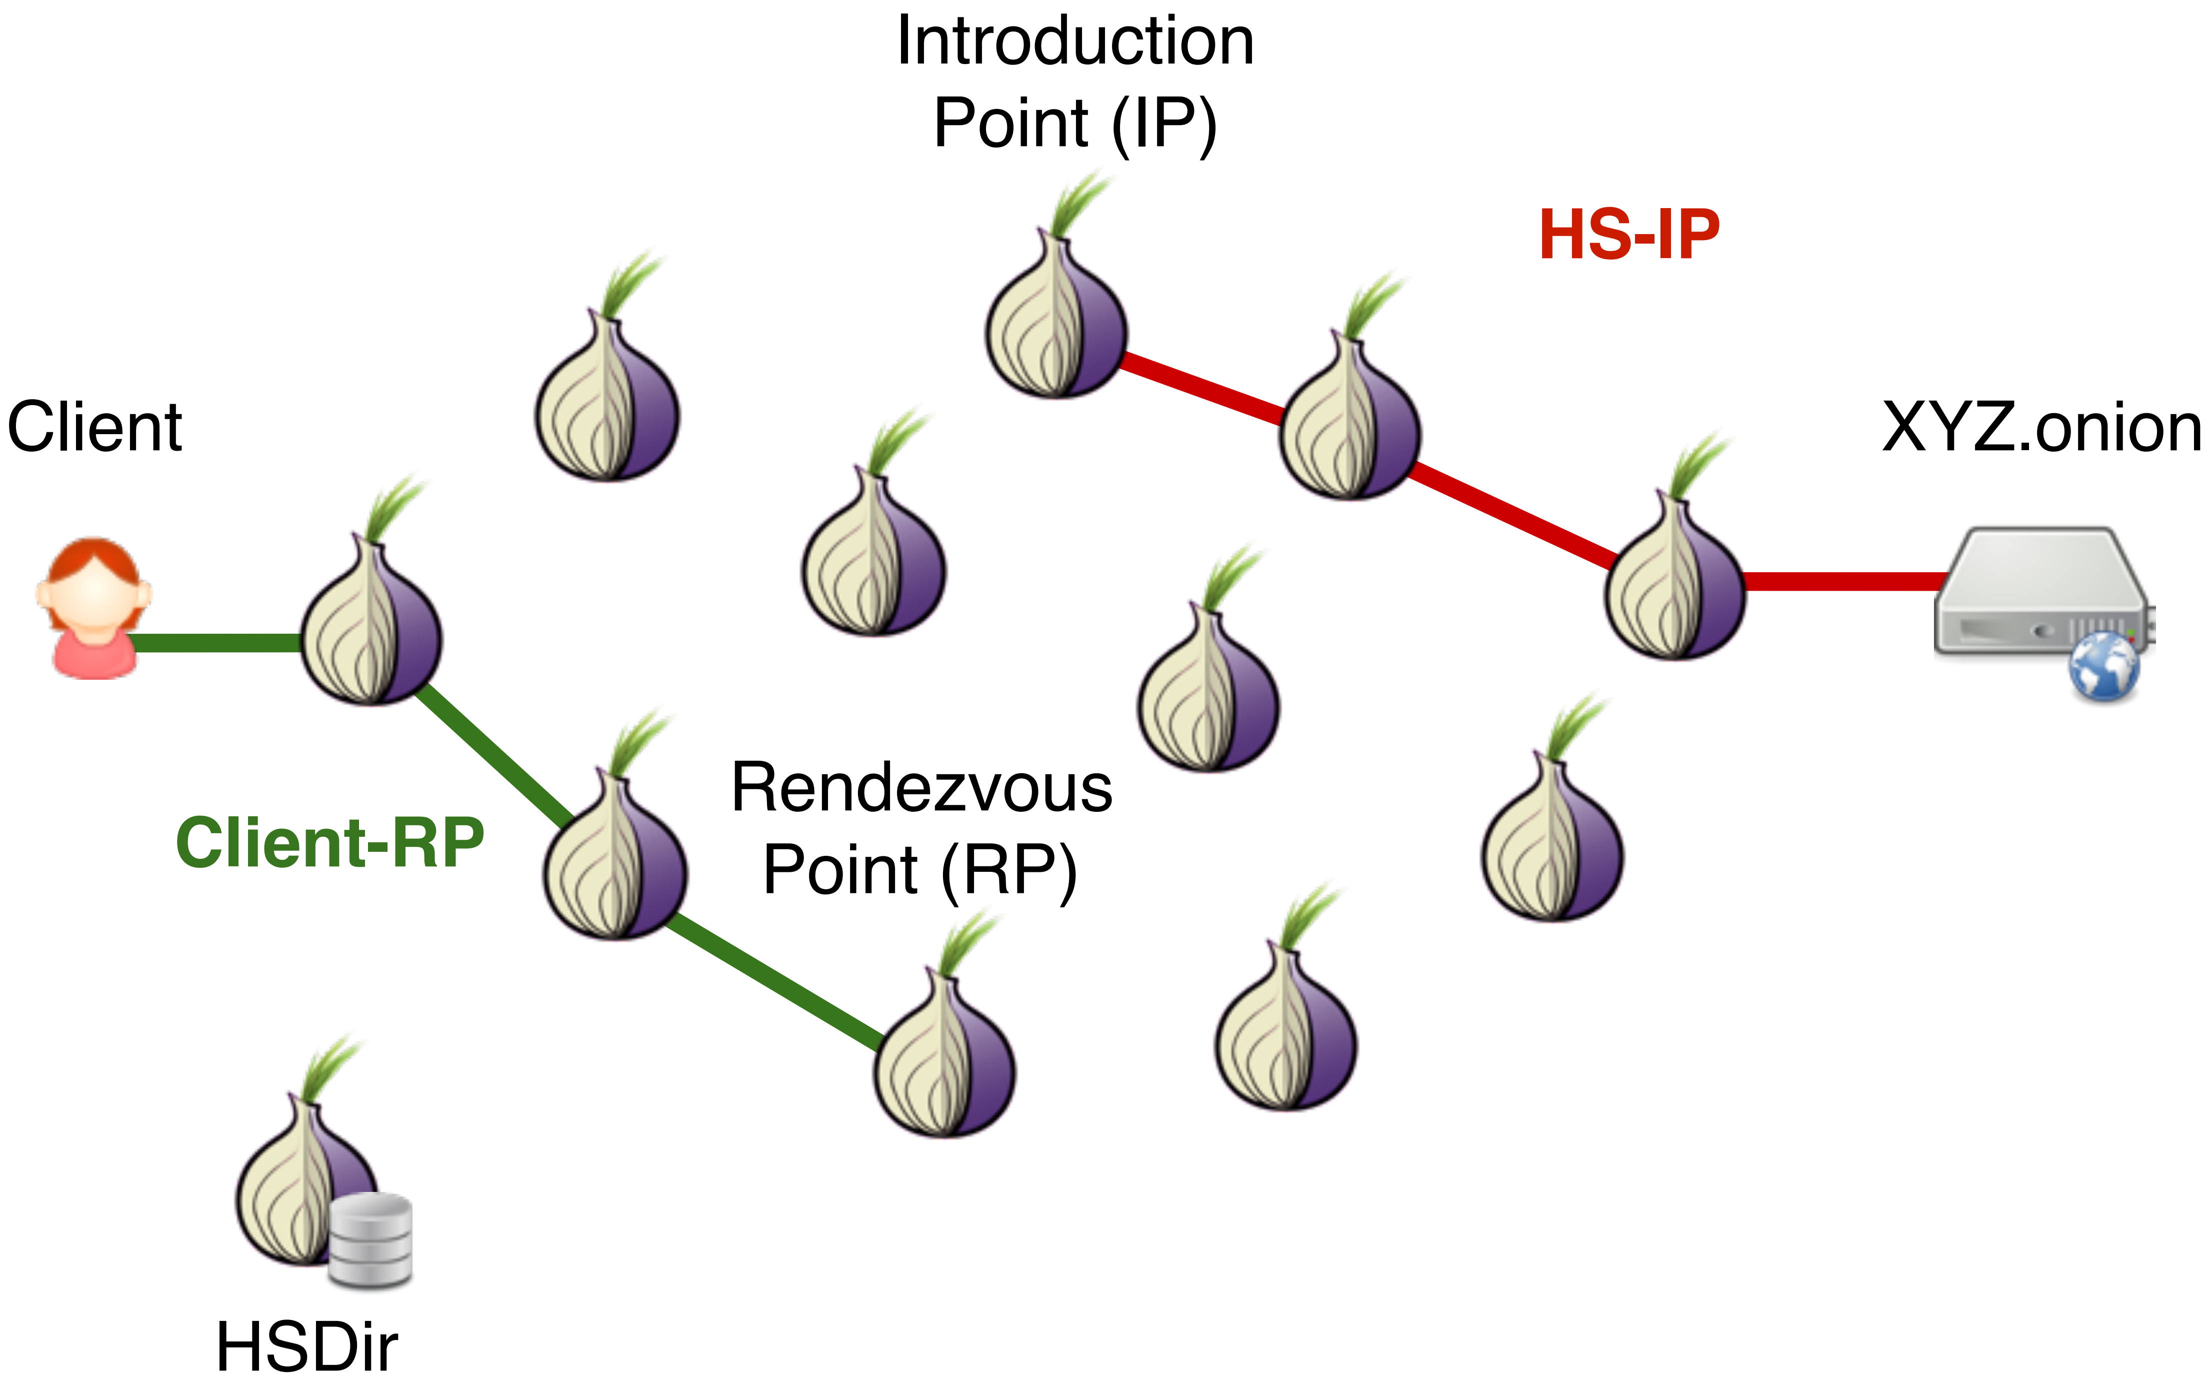

Tor Hidden Services (HS)

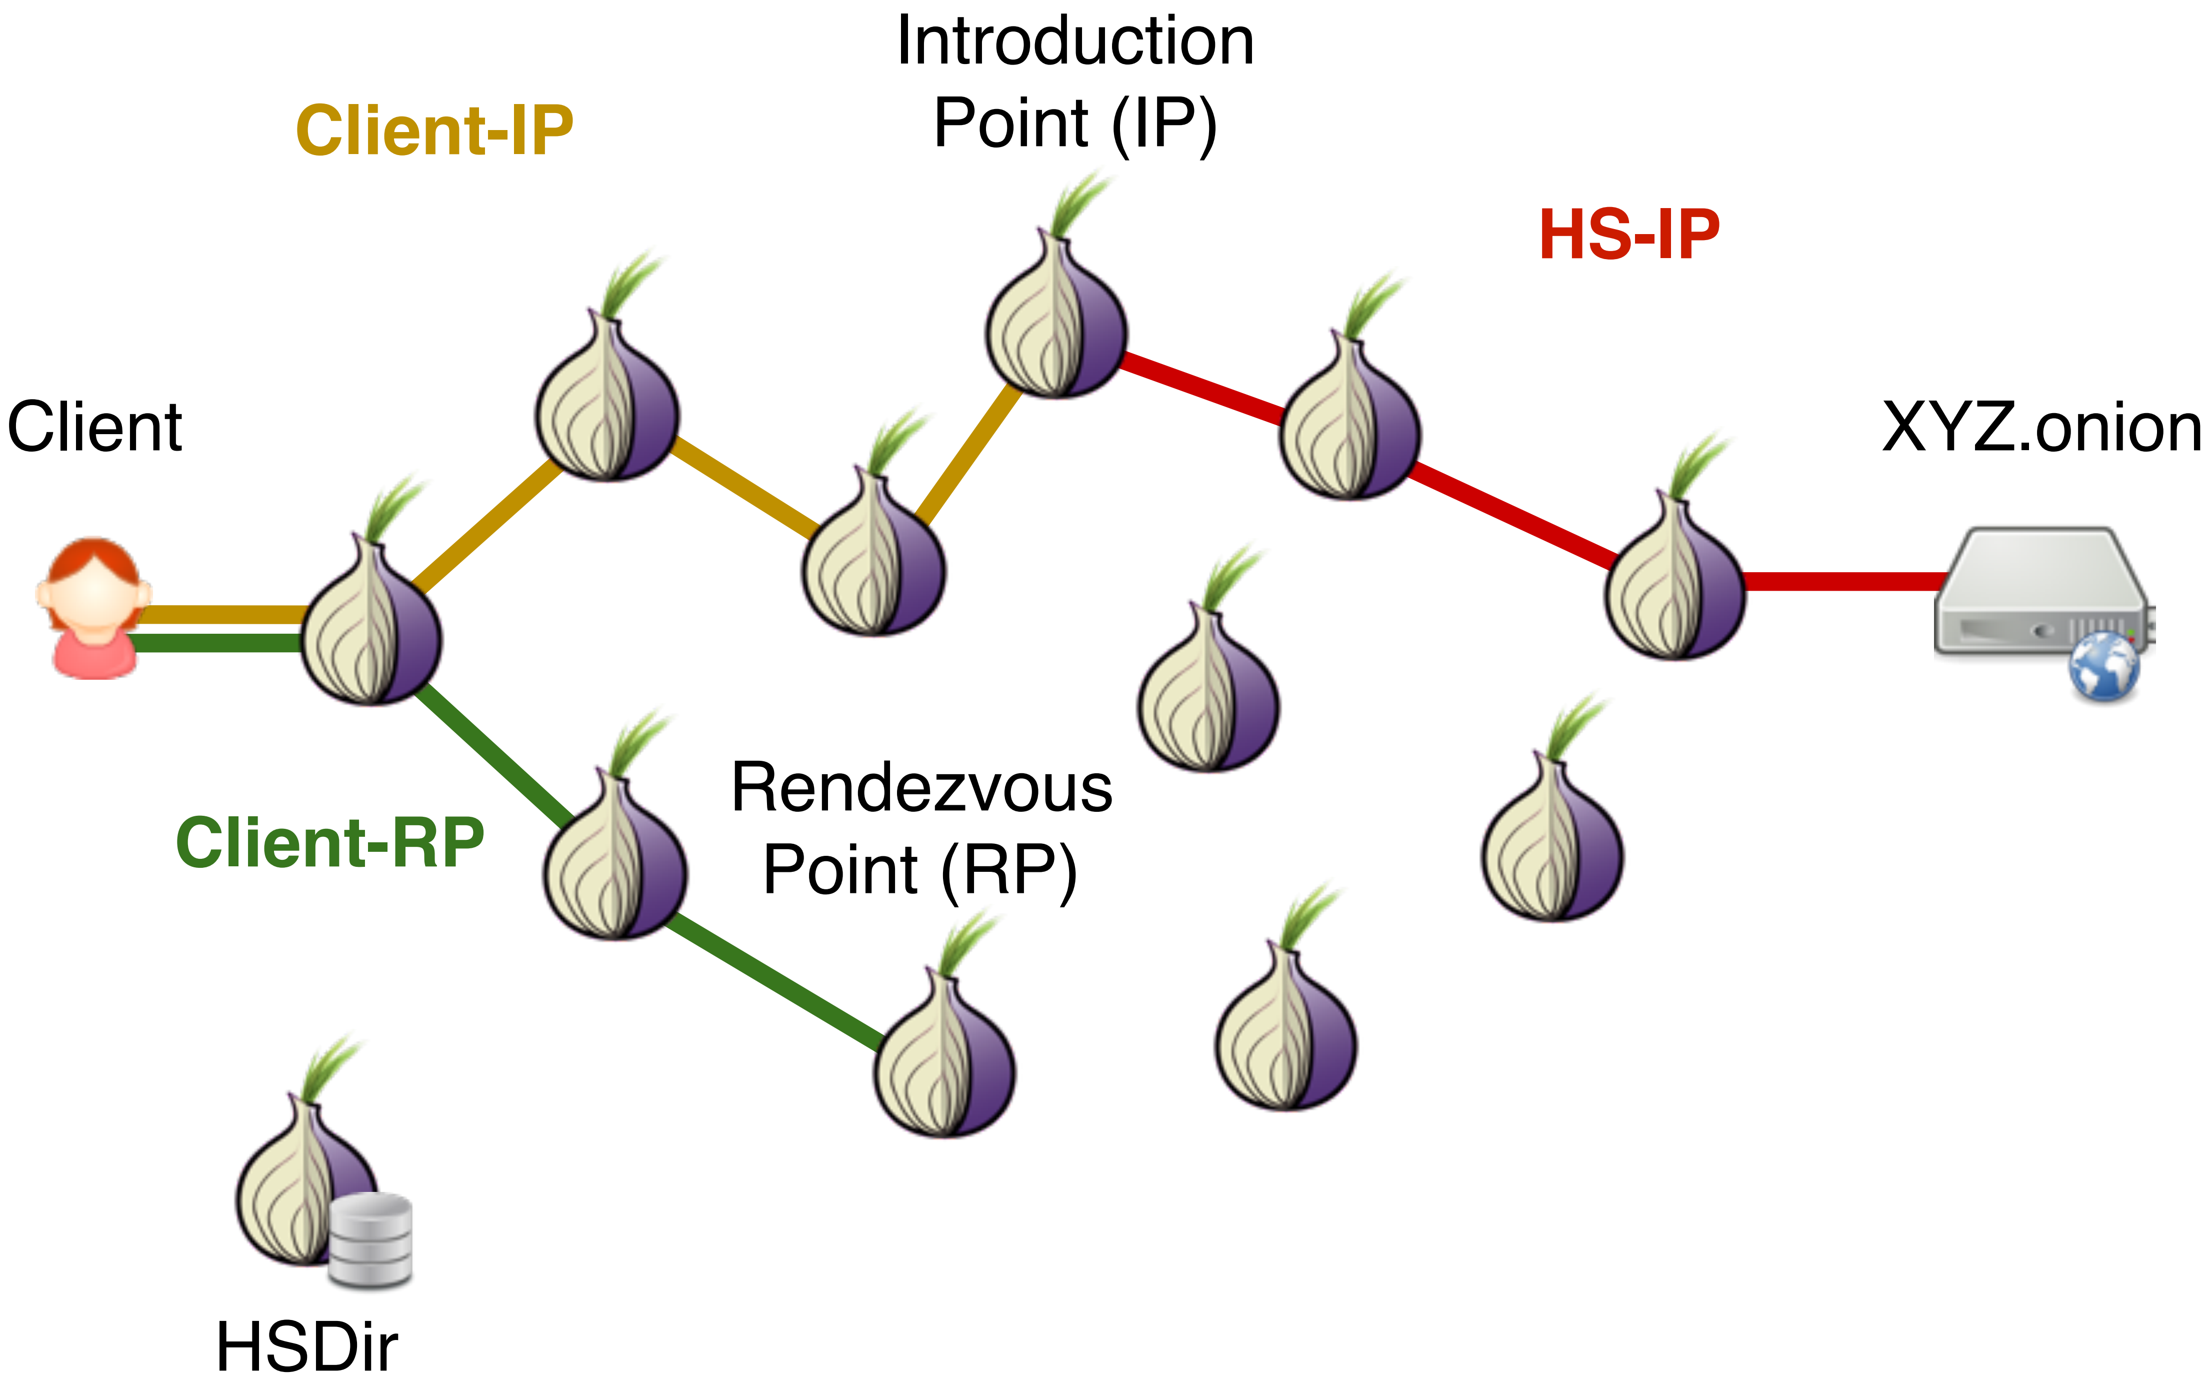

Tor Hidden Services (HS)

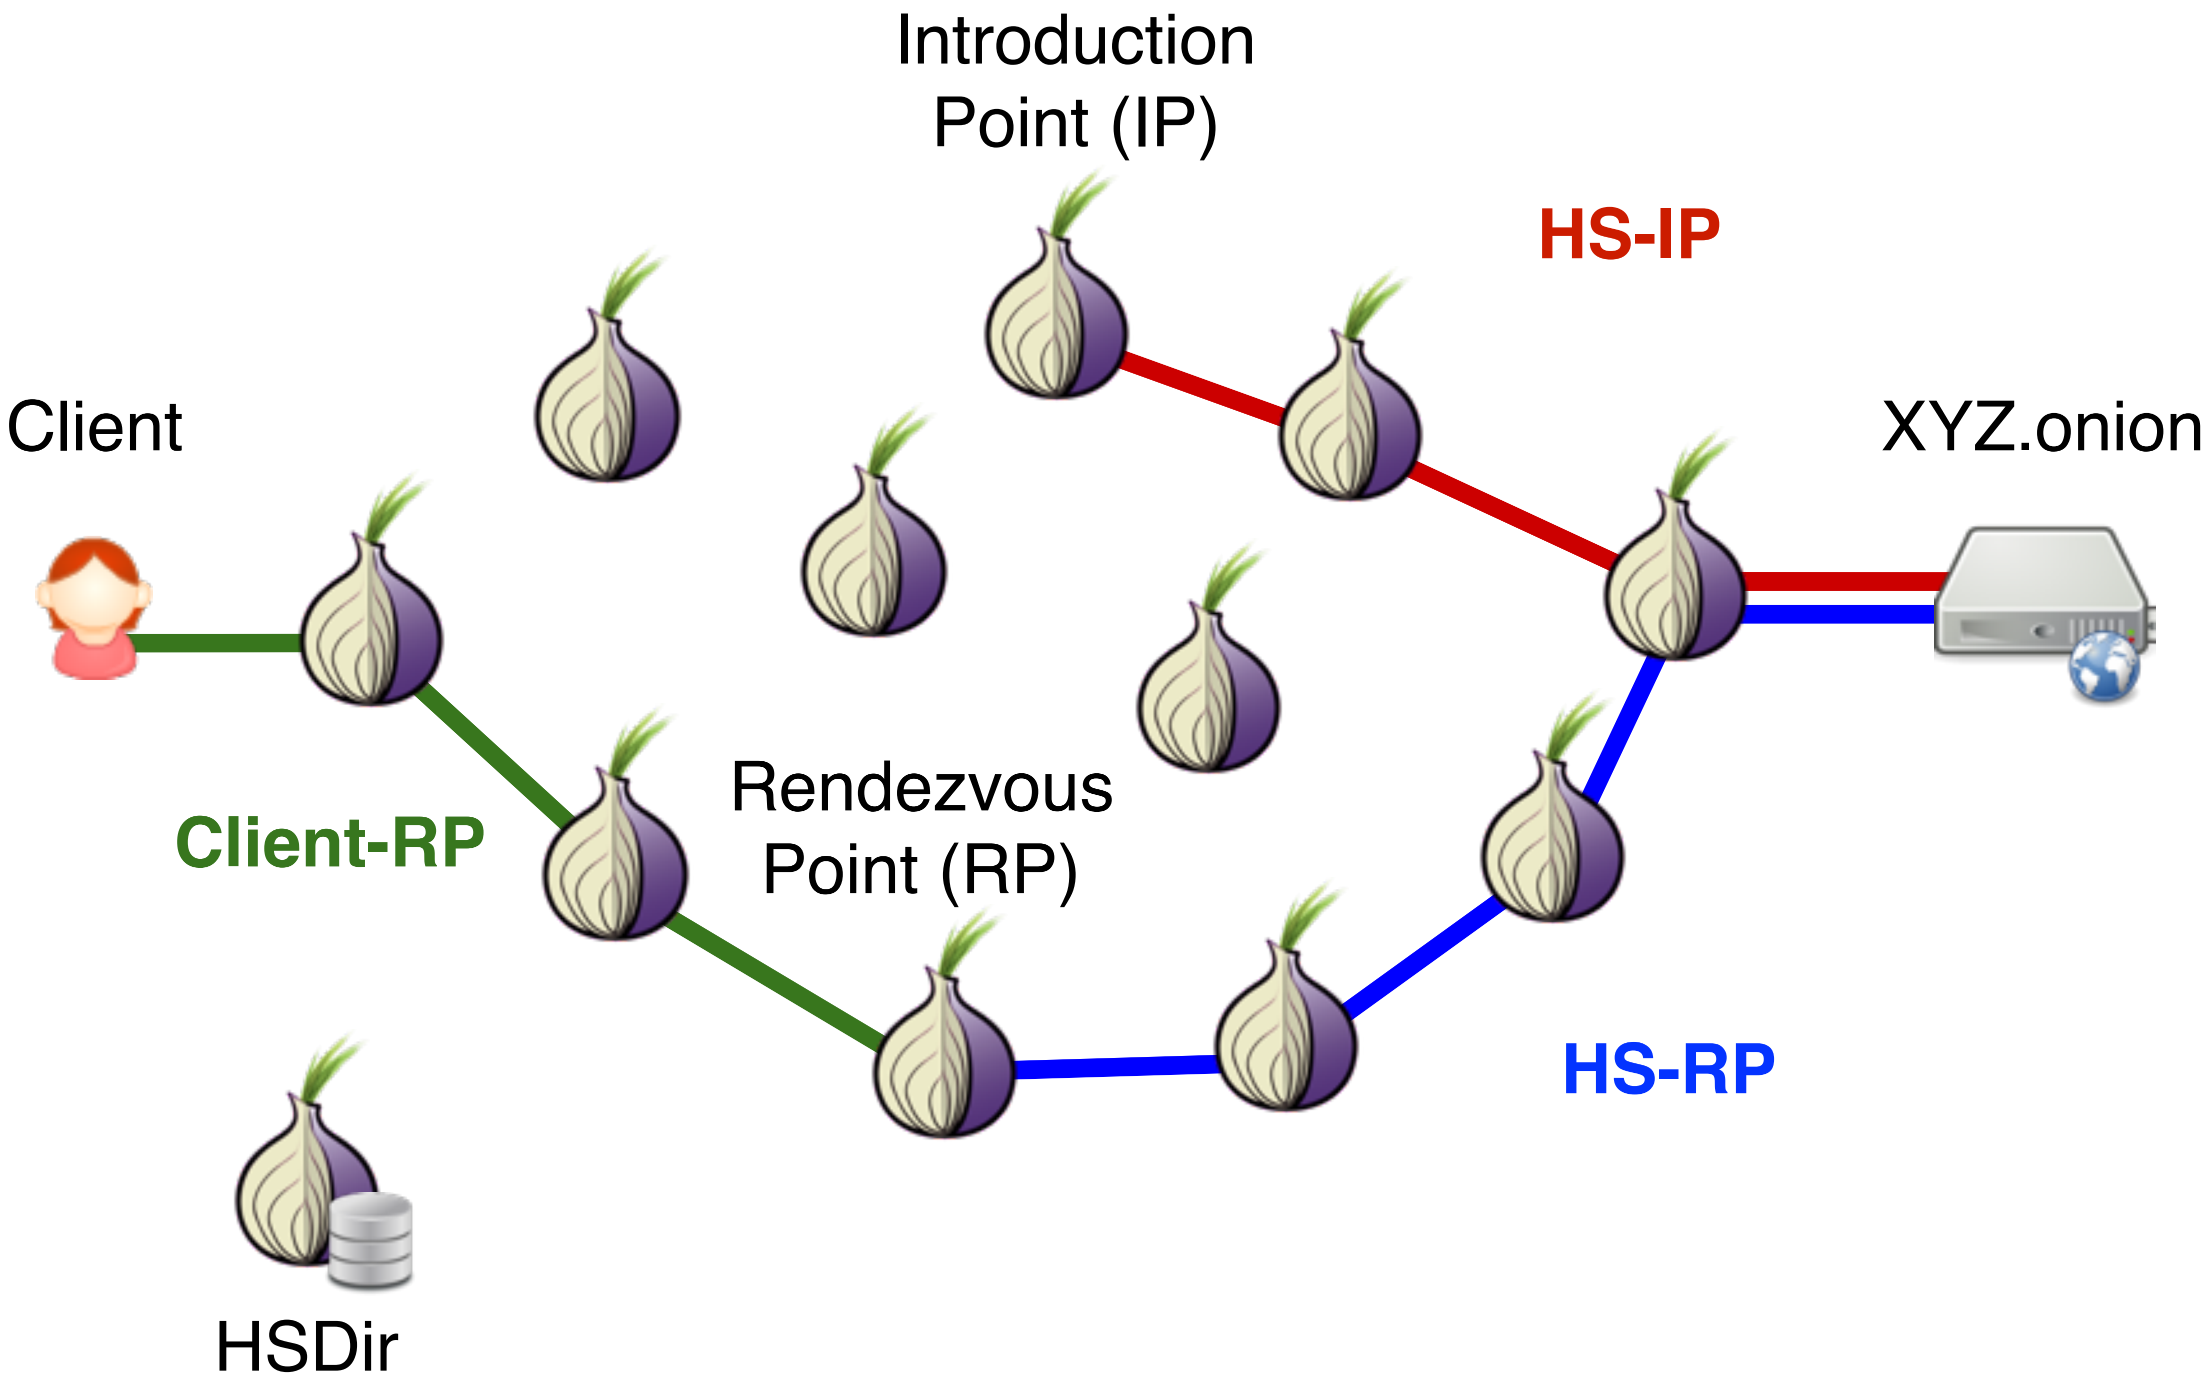

Fingerprinting HS activity

Kwon et al.: distinguish HS traffic from normal Tor traffic.

Distinguishers:

- Duration of Activity.
- Number of Incoming/Outgoing cells.
- Sequence of N first cells.

Observations: HS traffic

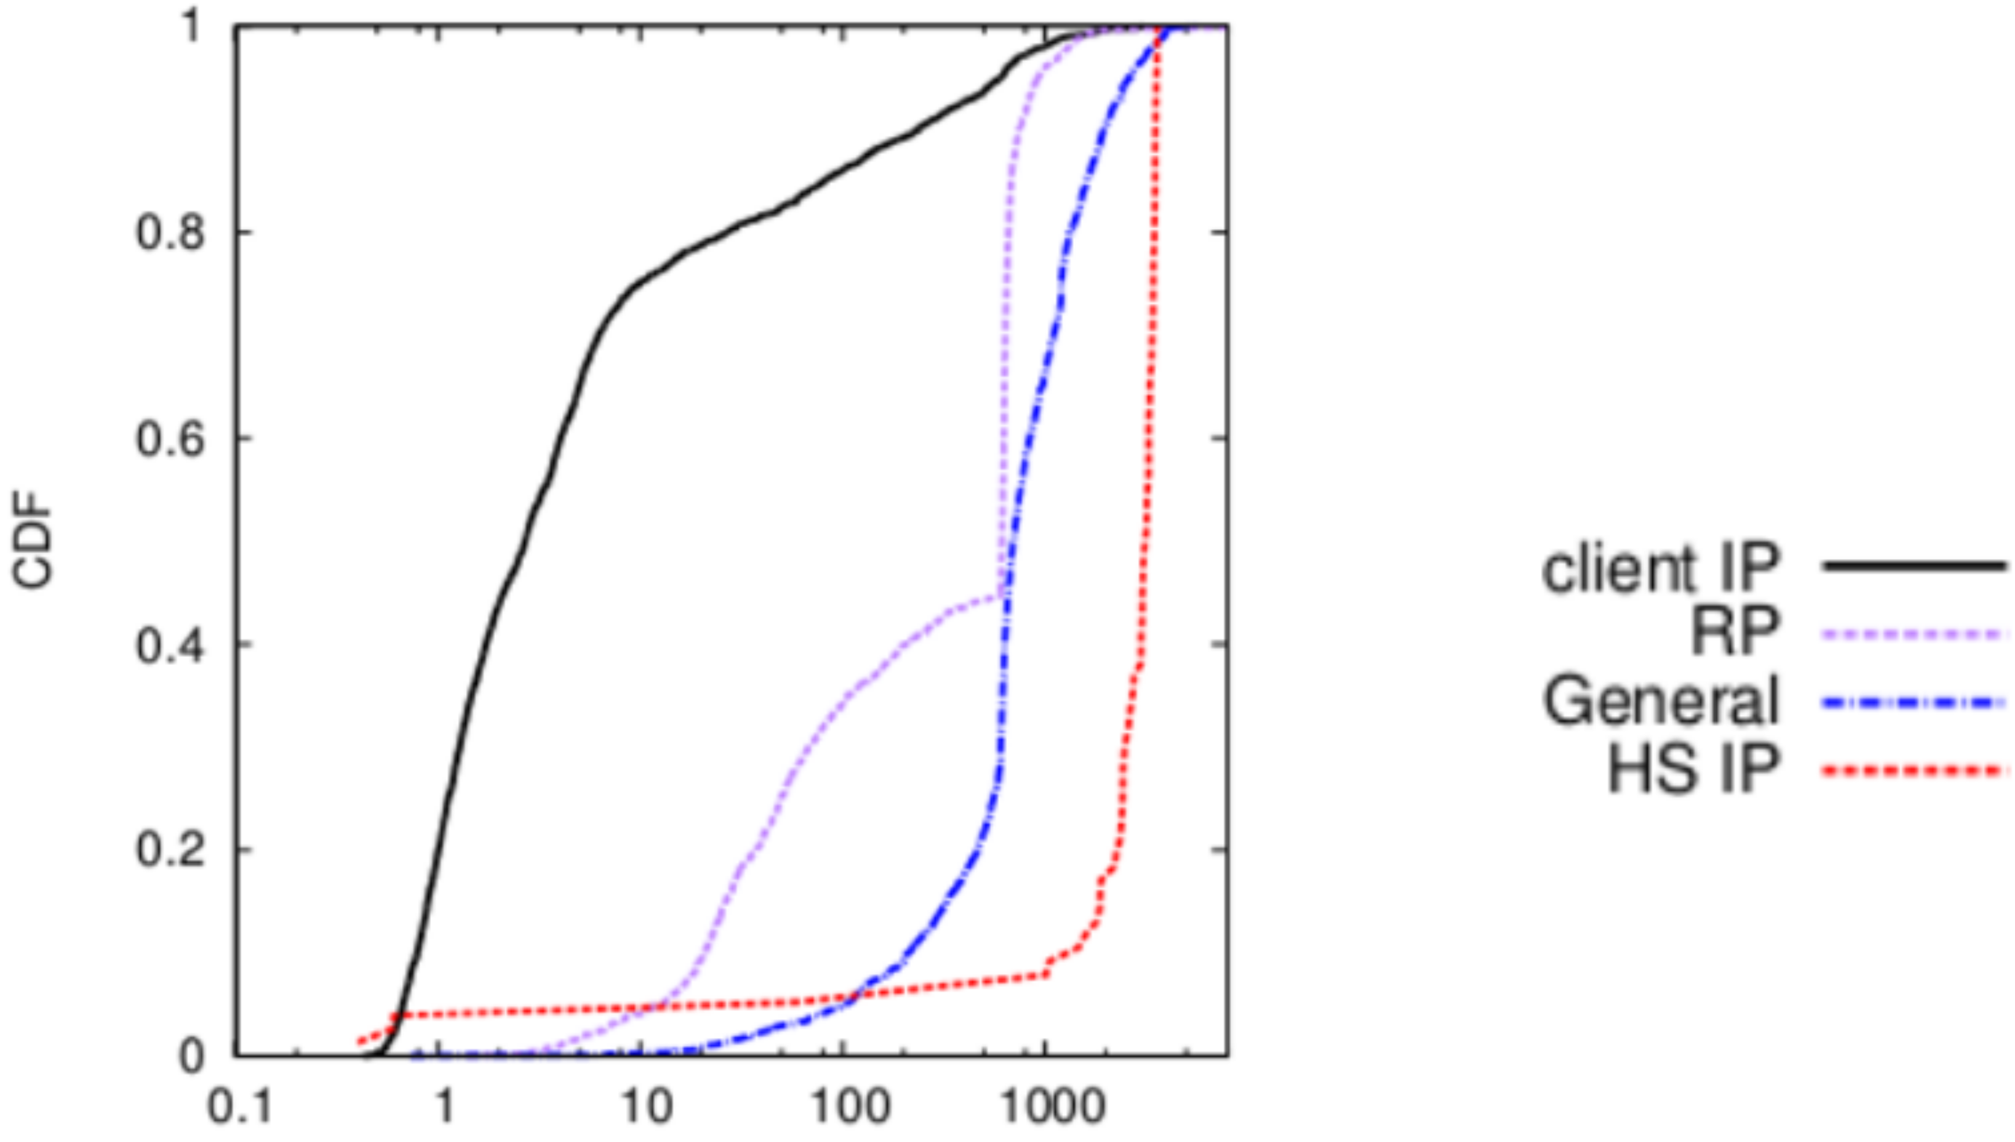

Duration of Activity (seconds)

HS Fingerprinting Takeaways

Website Fingerprinting is more effective for HS.

- Less background traffic.
- Less variability.
- **Smaller world.**

HS guard discovery attack

- Biryukov et al.: anybody can force HS to build a circuit to a RP.
- Attacker's RP introduces a signal. Attacker waits until one of his middle nodes observes the signal.
- Tor proposal #247 change path selection to prevent:
  - Guard discovery: two layers of guards (reduce rotation time).
  - Compromise: pin 3rd node with respect to the second guard.

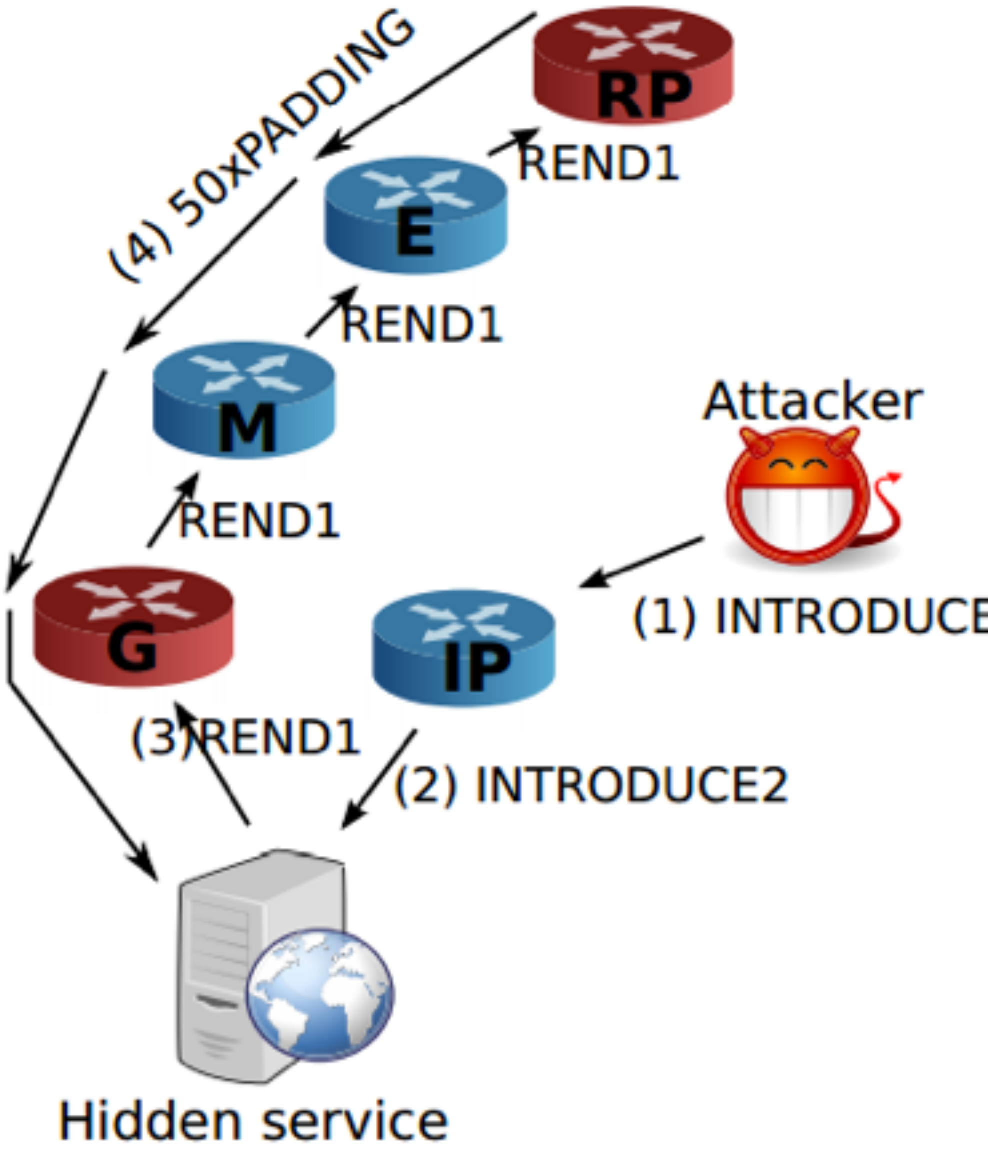

Source: A. Biryukov et al.

"Relay early" traffic confirmation attack (I)

- Feb 2016, Dept. of Defense paid CMU to indiscriminately target all users of hidden services.

- In 2014 BlackHat conference researchers withdraw a talk about a new attack against Tor.

- Signed up around 115 fast non-exit relays, all running on 50.7.0.0/16 or 204.45.0.0/16.

- Together these relays summed to about 6.4% of the Guard capacity in the network.

“Relay early” traffic confirmation attack (II)

- The attacker encoded the name of the hidden service in the injected signal using headers of cells (RELAY vs RELAY EARLY).

- The HSDirs add signal and Guards read it (during HS descriptor lookup).

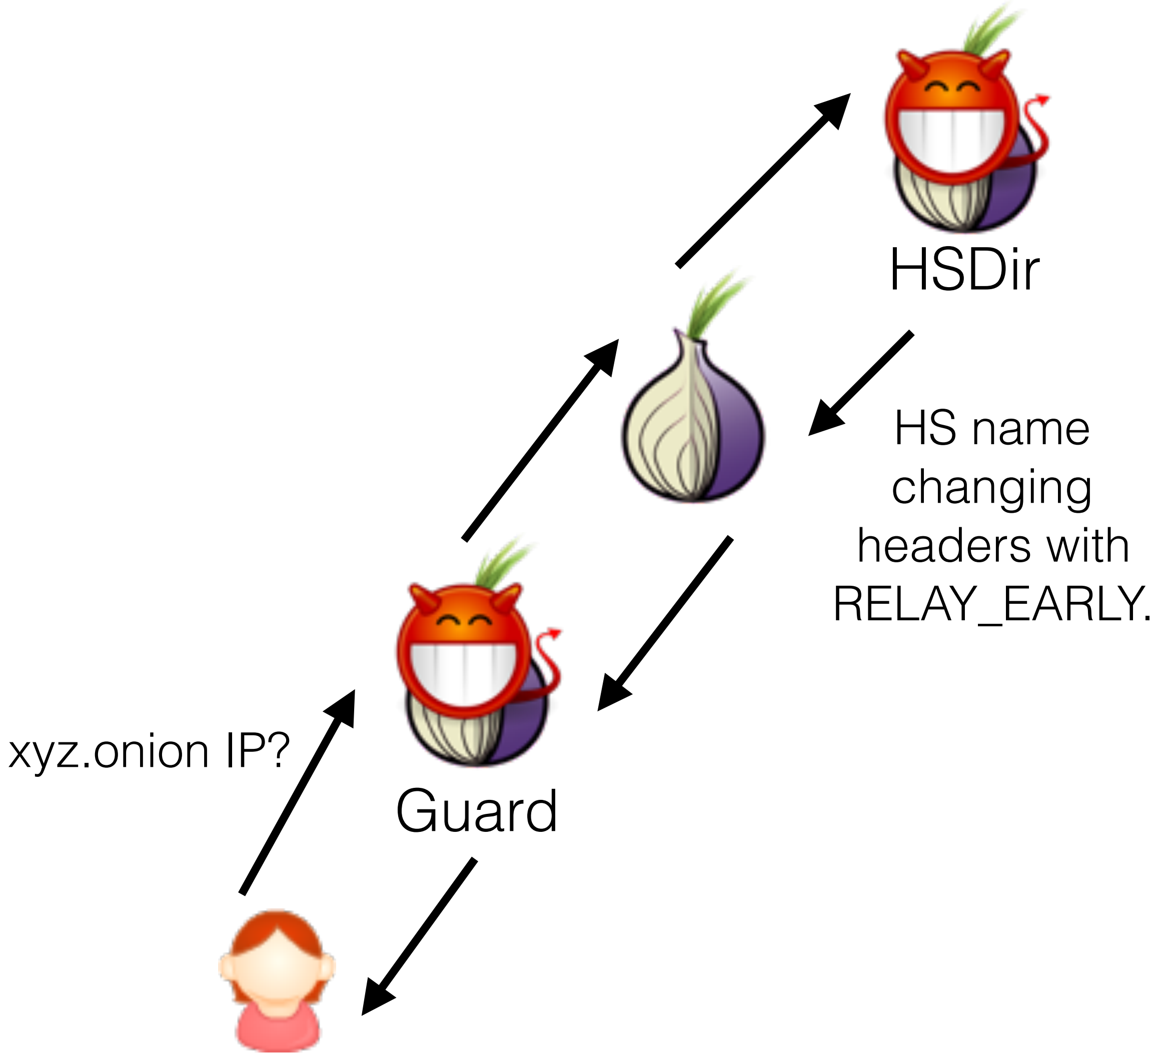

Defences against Traffic Analysis in Tor

Pluggable Transport (PT) Architecture

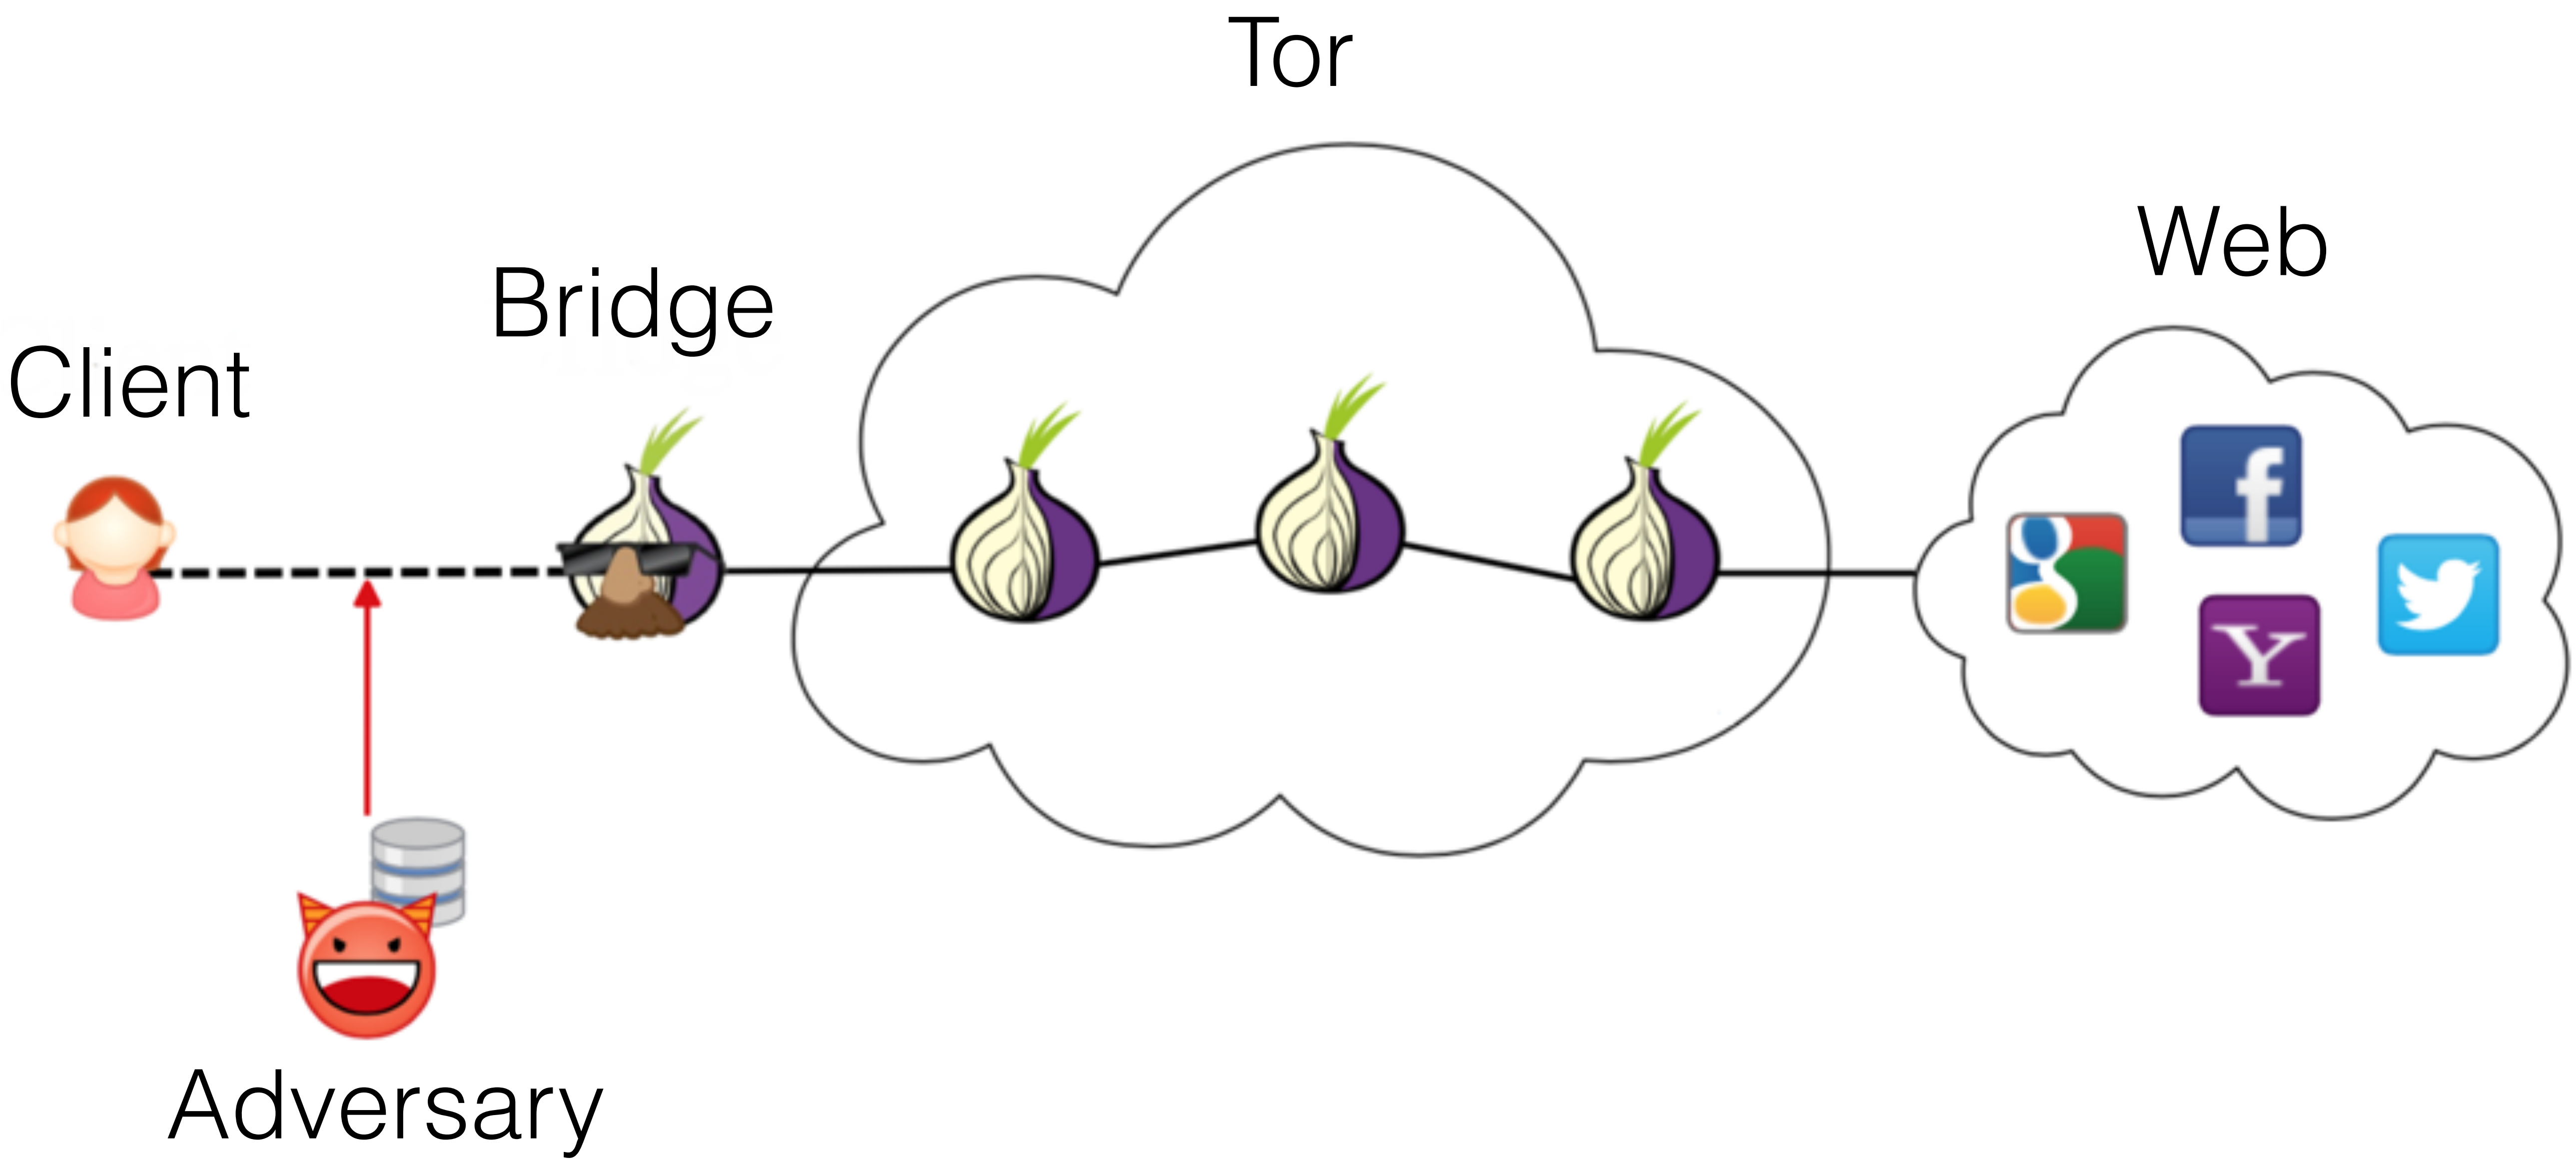

Popular PTs

- Obfsproxy.
- Format-Transforming-Encryption.
- Meek
- Flashproxy and Snowflake.

Obfs4

- Based on Obfsproxy.
- Maintained by Yawning Angel.
- Similarly to *Scramblesuit,* it's an obfuscation layer for TCP protocols.
- Authenticated key exchange (prevent MITM): *nTor.*
- Public key obfuscation: *Elligator2* (Bernstein et al.)

Format-Transforming Encryption (FTE)

- Maintained by: Kevin P. Dyer.

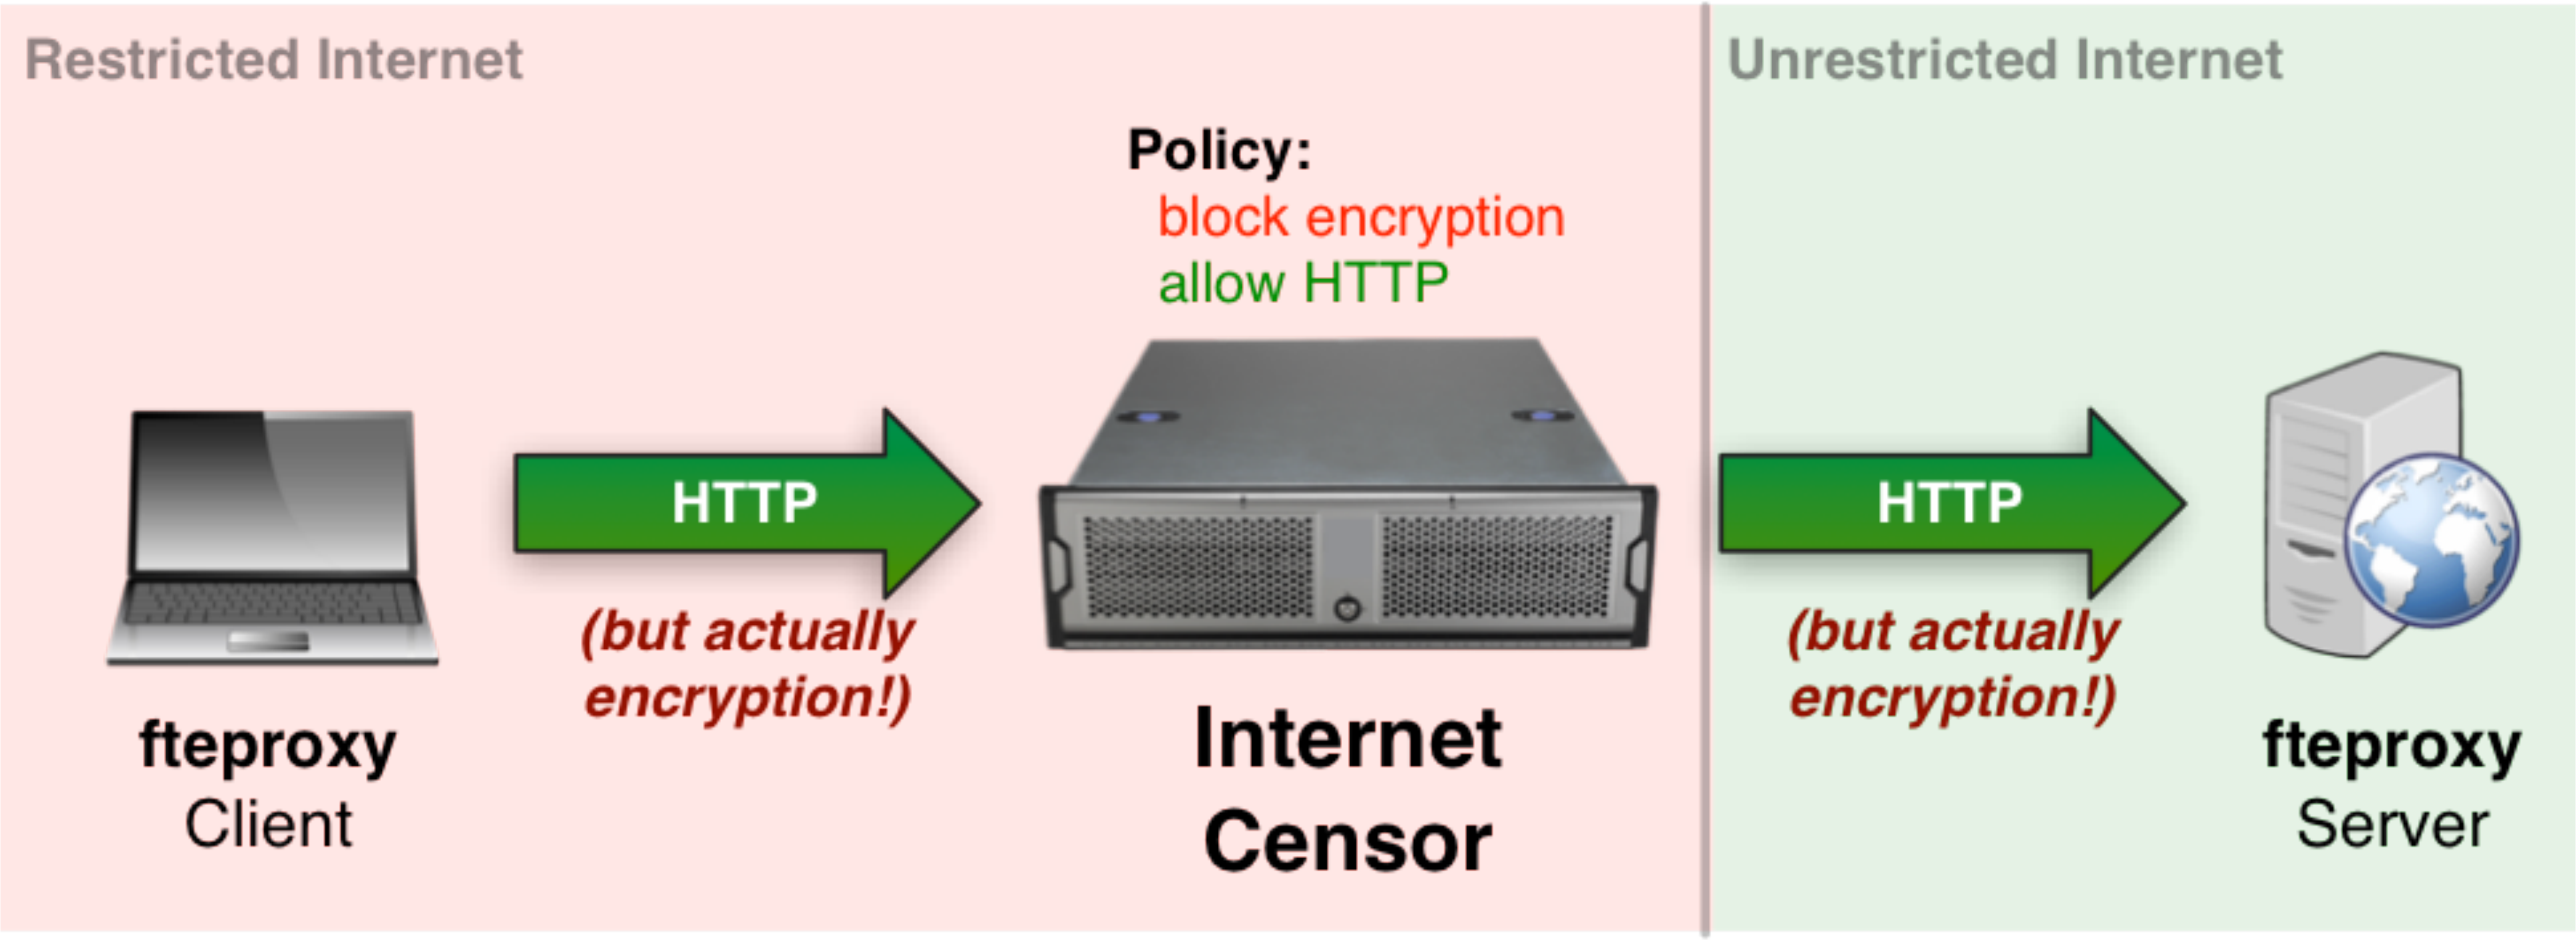

Source: fteproxy.org

Meek

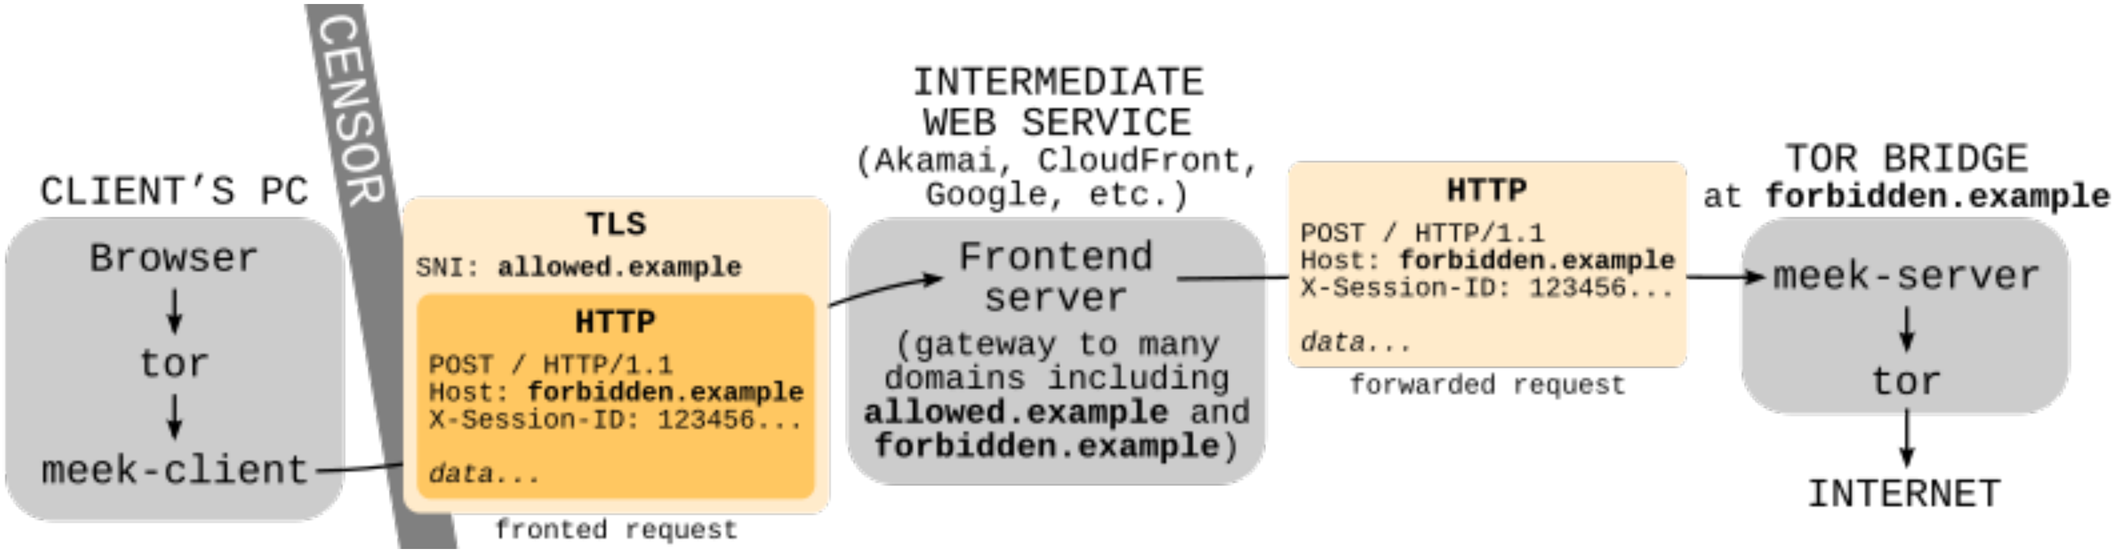

Source: trac.torproject.org

- Maintained by: David Fifield
- Based on *Domain Fronting:* collateral damage of censorship.
- Demo in Amazon “Cloudfront” CDN:

wget -q -O - https://a0.awsstatic.com/ --header 'Host: d2zfqthxsdq309.cloudfront.net'

Concluding remarks

- The Internet was not designed to protect against traffic analysis.
- Widespread use of traffic analysis techniques that threats the privacy of Internet users.
- The need for better performance uncovers new side-channels.
- Tor, the best anonymity tool we have, still has a lot of attacks that need to be solved.

Resources

- Danezis and Clayton. “Introducing Traffic Analysis”.
- Danezis-Diaz-Syverson. “Systems for Anonymous Communication”.
- Chaum original mix and dinning cryptographers papers.
- Statistical disclosure attacks (Danezis).
- Browser fingerprinting: https://securehomes.esat.kuleuven.be/~gacar/persistent/
- Juarez et al. “A Critical Evaluation of Website Fingerprinting Attacks”, CCS’13.
- Kown et al. “Circuit Fingerprinting Attacks: Passive Deanonymization of Tor Hidden Services”, USENIX’15.
- Tor: https://blog.torproject.org
- More: http://freehaven.net/anonbib/

Thank you!

Marc Juarez

marcjuarez {at} kuleuven.be

http://homes.esat.kuleuven.be/~mjuarezm