Trawling for Tor Hidden Services: Detection, Measurement, Deanonymization

A. Biryukov, I. Pustogarov, R.P. Weinmann
University of Luxembourg
Ivan.pustogarov@uni.lu

May 20, 2013

Overview

- Background
- Measuring the popularity of hidden services
- DoSing hidden services.
- Harvesting onion addresses.
- Revealing the guards.
- Opportunistic deanonymisation.

Overview

- **Background**
- Measuring the popularity of hidden services
- DoSing hidden services.
- Harvesting onion addresses.
- Revealing the guards.
- Opportunistic deanonymisation.

Tor anonymity network

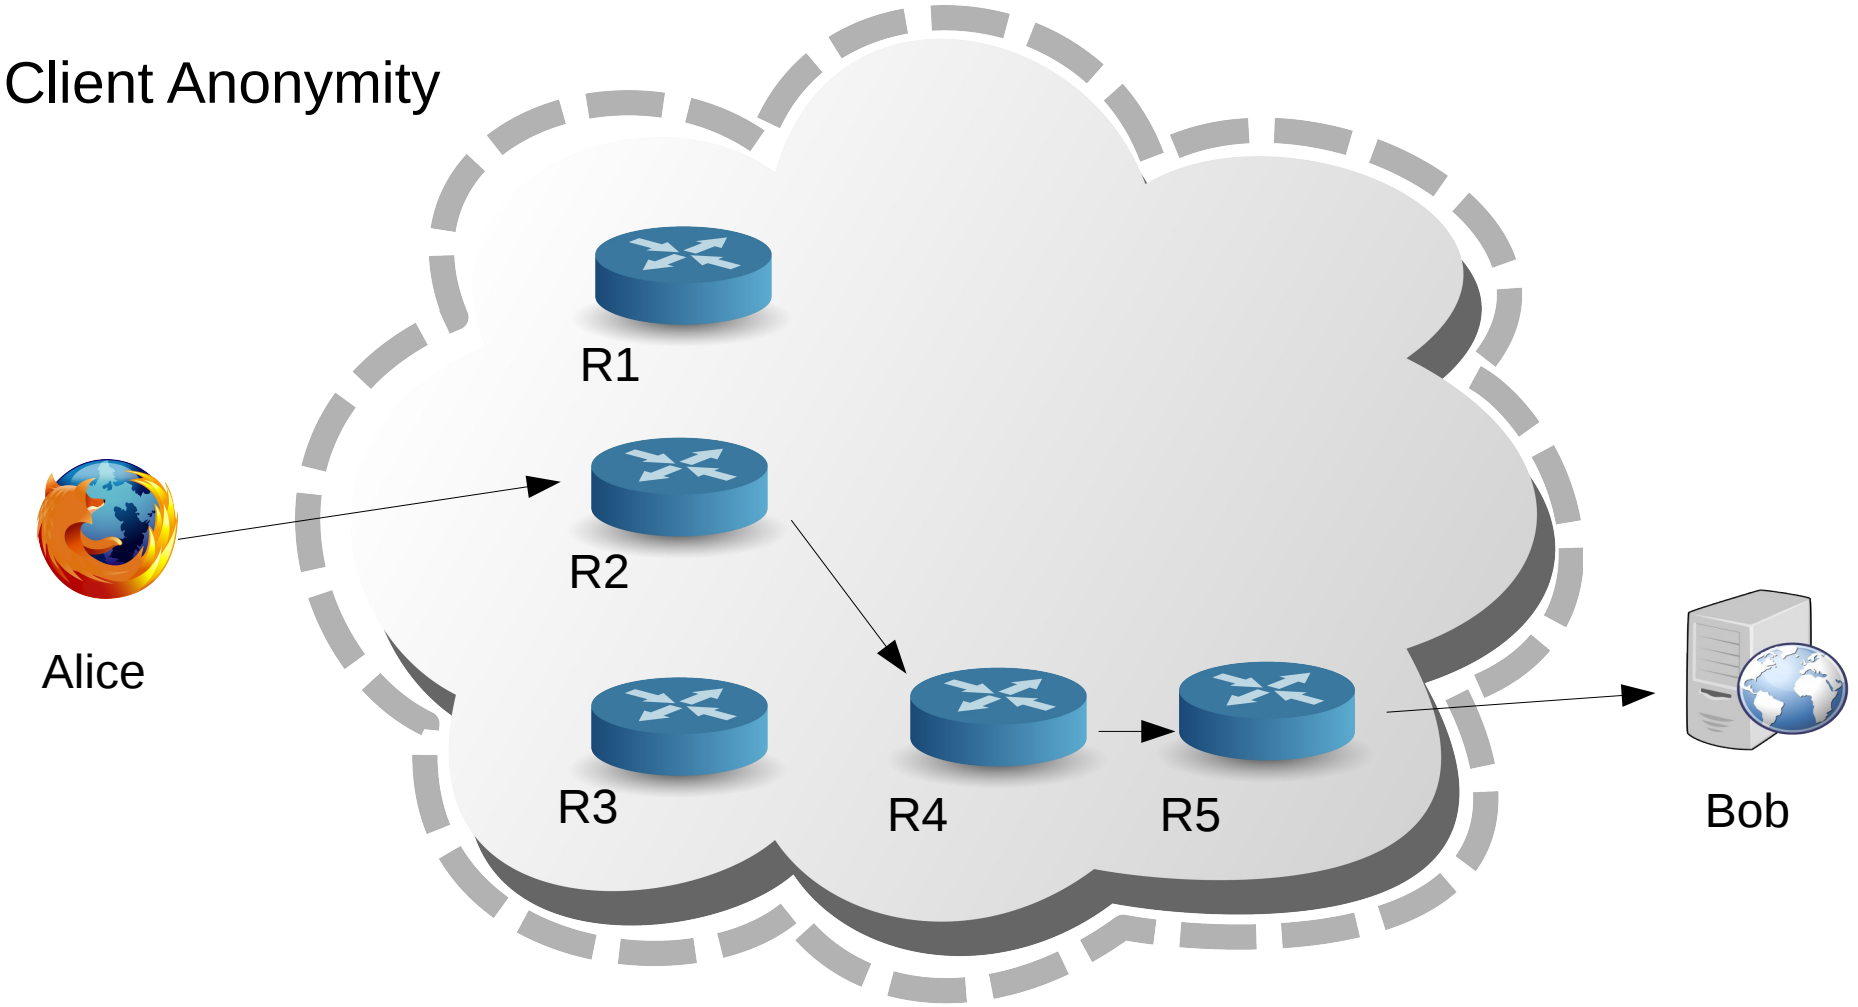

Tor anonymity network

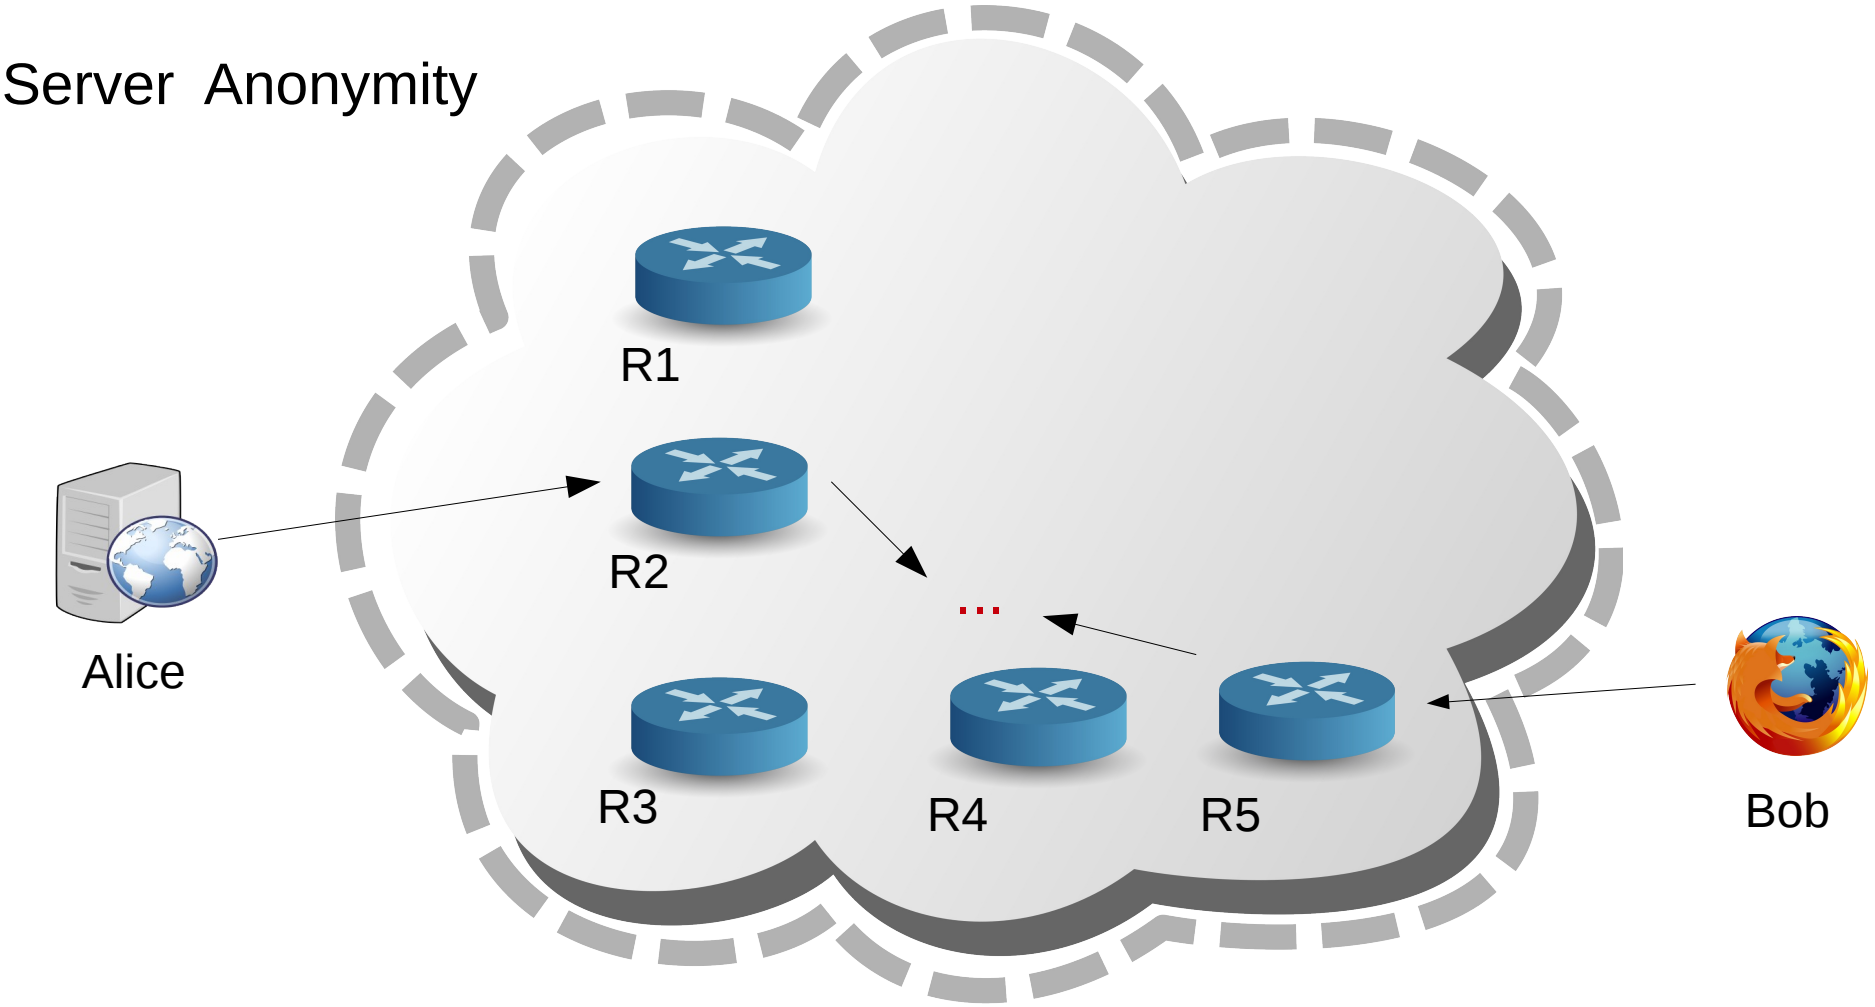

Tor anonymity network

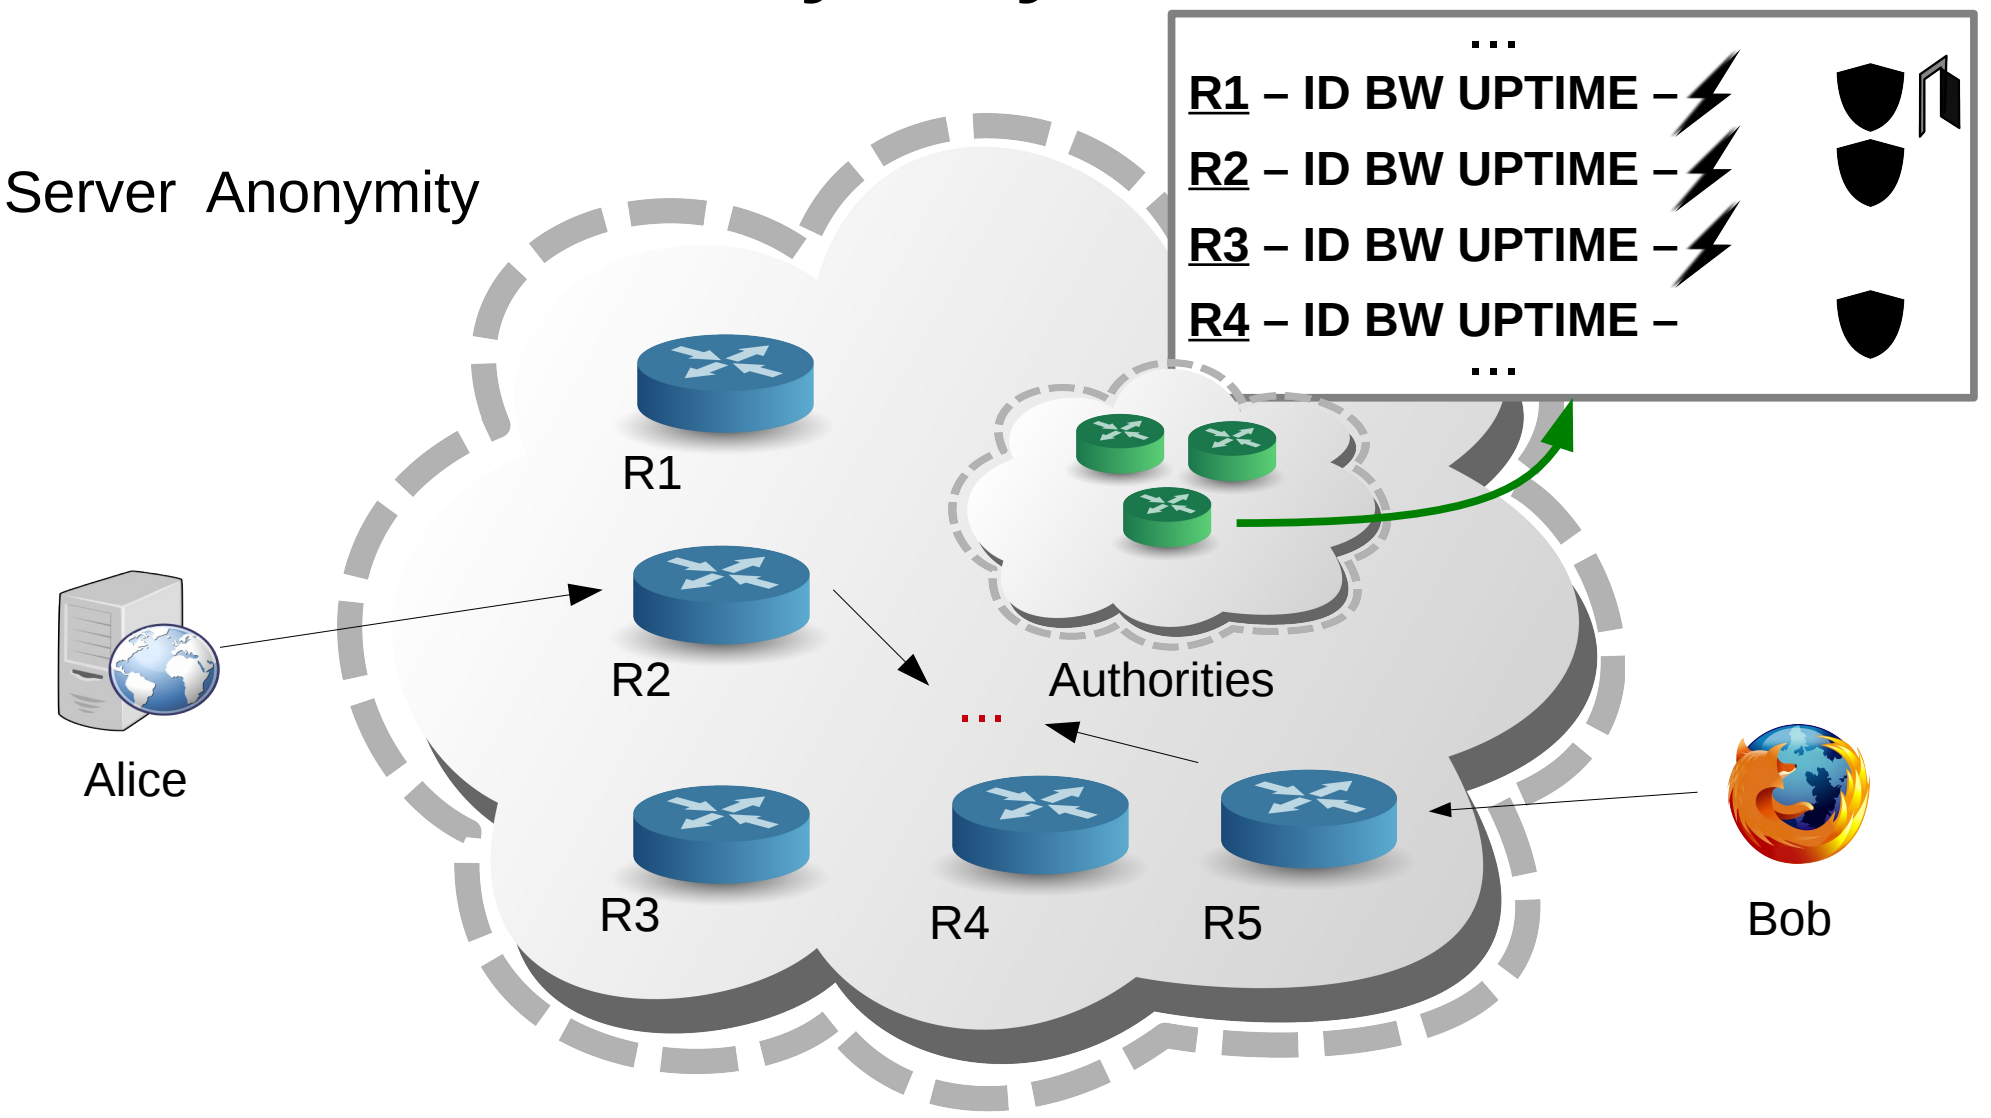

Consensus

menTor	1737	**67 d**	55863896.cust.multi.fi [85.134.56.150]
microshaft	2820	**66 d**	tor-exit.microshaft.org [208.201.249.3]
minisausage	3348	**35 d**	50.7.184.58 [50.7.184.58]
morphium	298	**51 d**	this.is.a.Tor.server.please.see.tor.morphium.info [91.143.90.25]
NetromAc	2115	**47 d**	1385160742.business.dbnet.dk [82.143.224.38]
Nitr0x	175	**78 d**	50.97.1.36-static.reverse.softlayer.com [50.97.1.36]
OhCanada	419	**51 d**	van1.zworg.com [209.17.191.117]
onconnex80	392	**213 d**	tor01.onconnex.com [184.105.231.11]
PasToutAFaitNet1	261	**176 d**	91.229.20.159 [91.229.20.159]
PasToutAFaitNet2	763	**196 d**	tor2.pastoutafait.net [95.130.11.247]
plebia	3599	**79 d**	tor-exit.plebia.org [37.59.162.218]
pps	9	**59 d**	184-22-164-107.static.hostnoc.net [184.22.164.107]
PrivaTOReu	4229	**100 d**	torexit.privator.eu [88.208.90.1]
programmercpp	149	**36 d**	proxy [213.171.220.40]
PsyNetNP	155	**52 d**	broadband-95-84-148-164.nationalcablenetworks.ru [95.84.148.164]
Qwerty	91	**157 d**	93.167.245.178 [93.167.245.178]

http://torstatus.blutmagie.de/

Guards

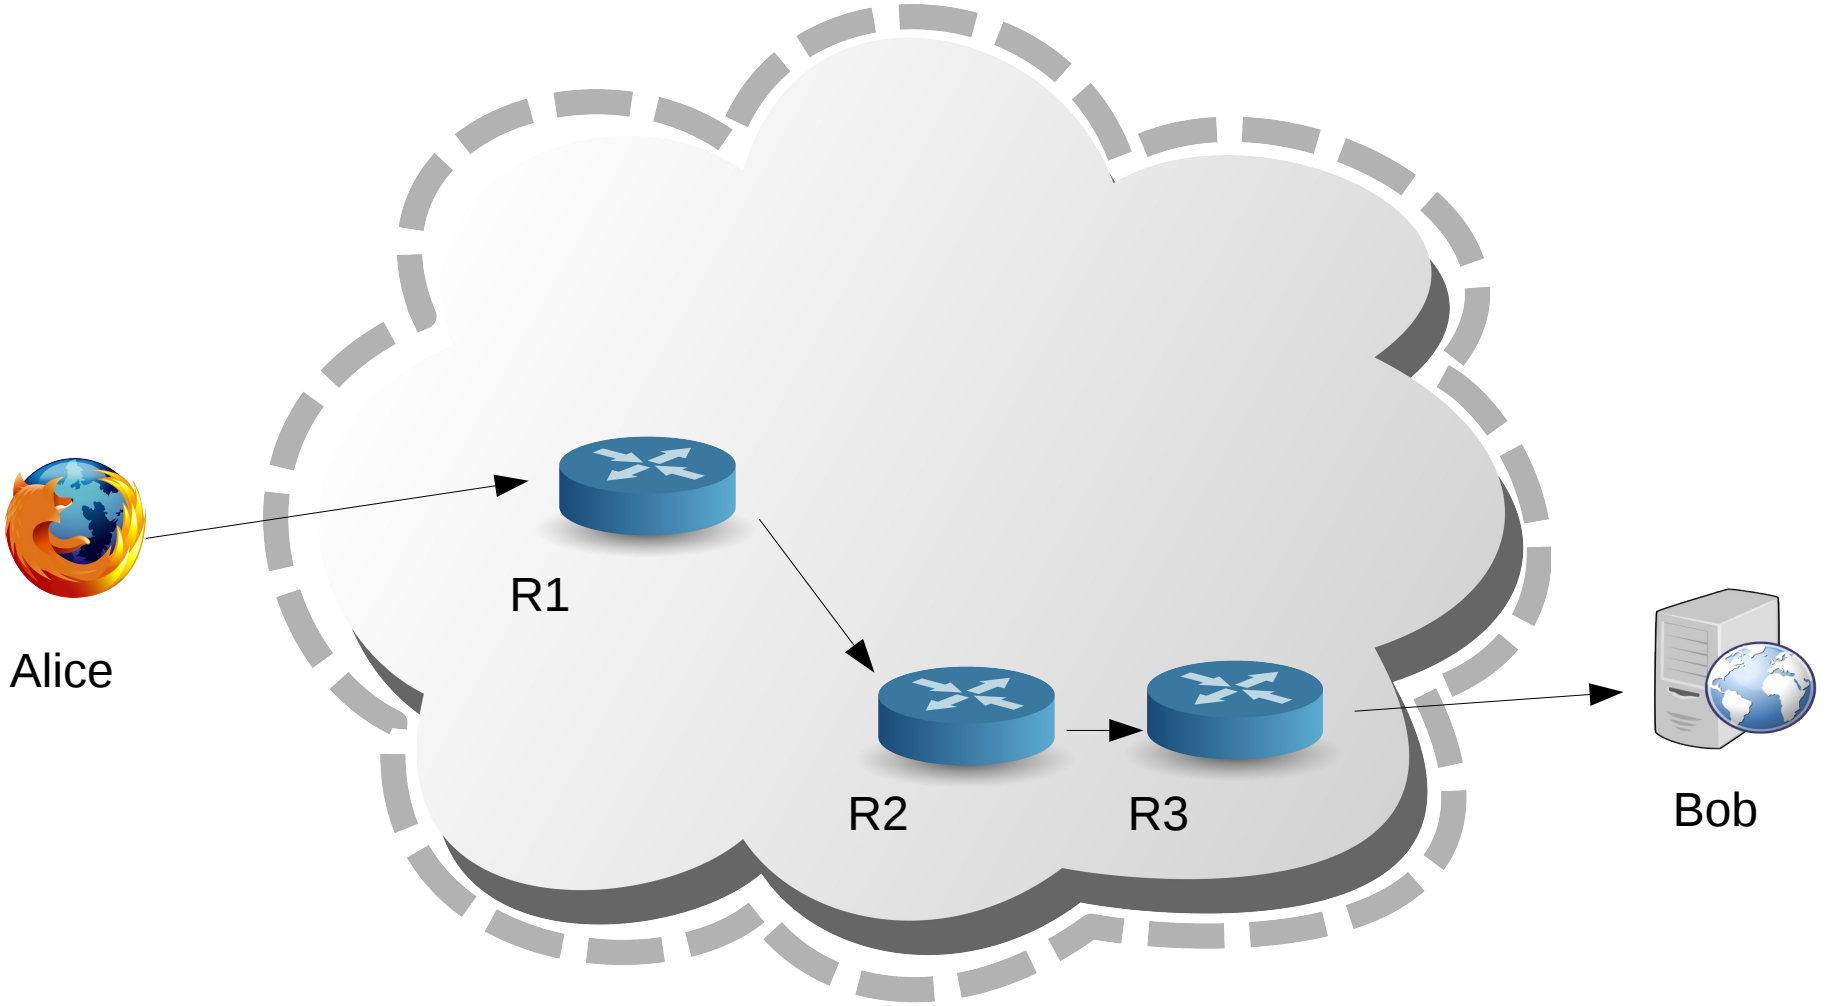

Guard = high uptime + high bandwidth

Every client has 3 Guard nodes

Guards

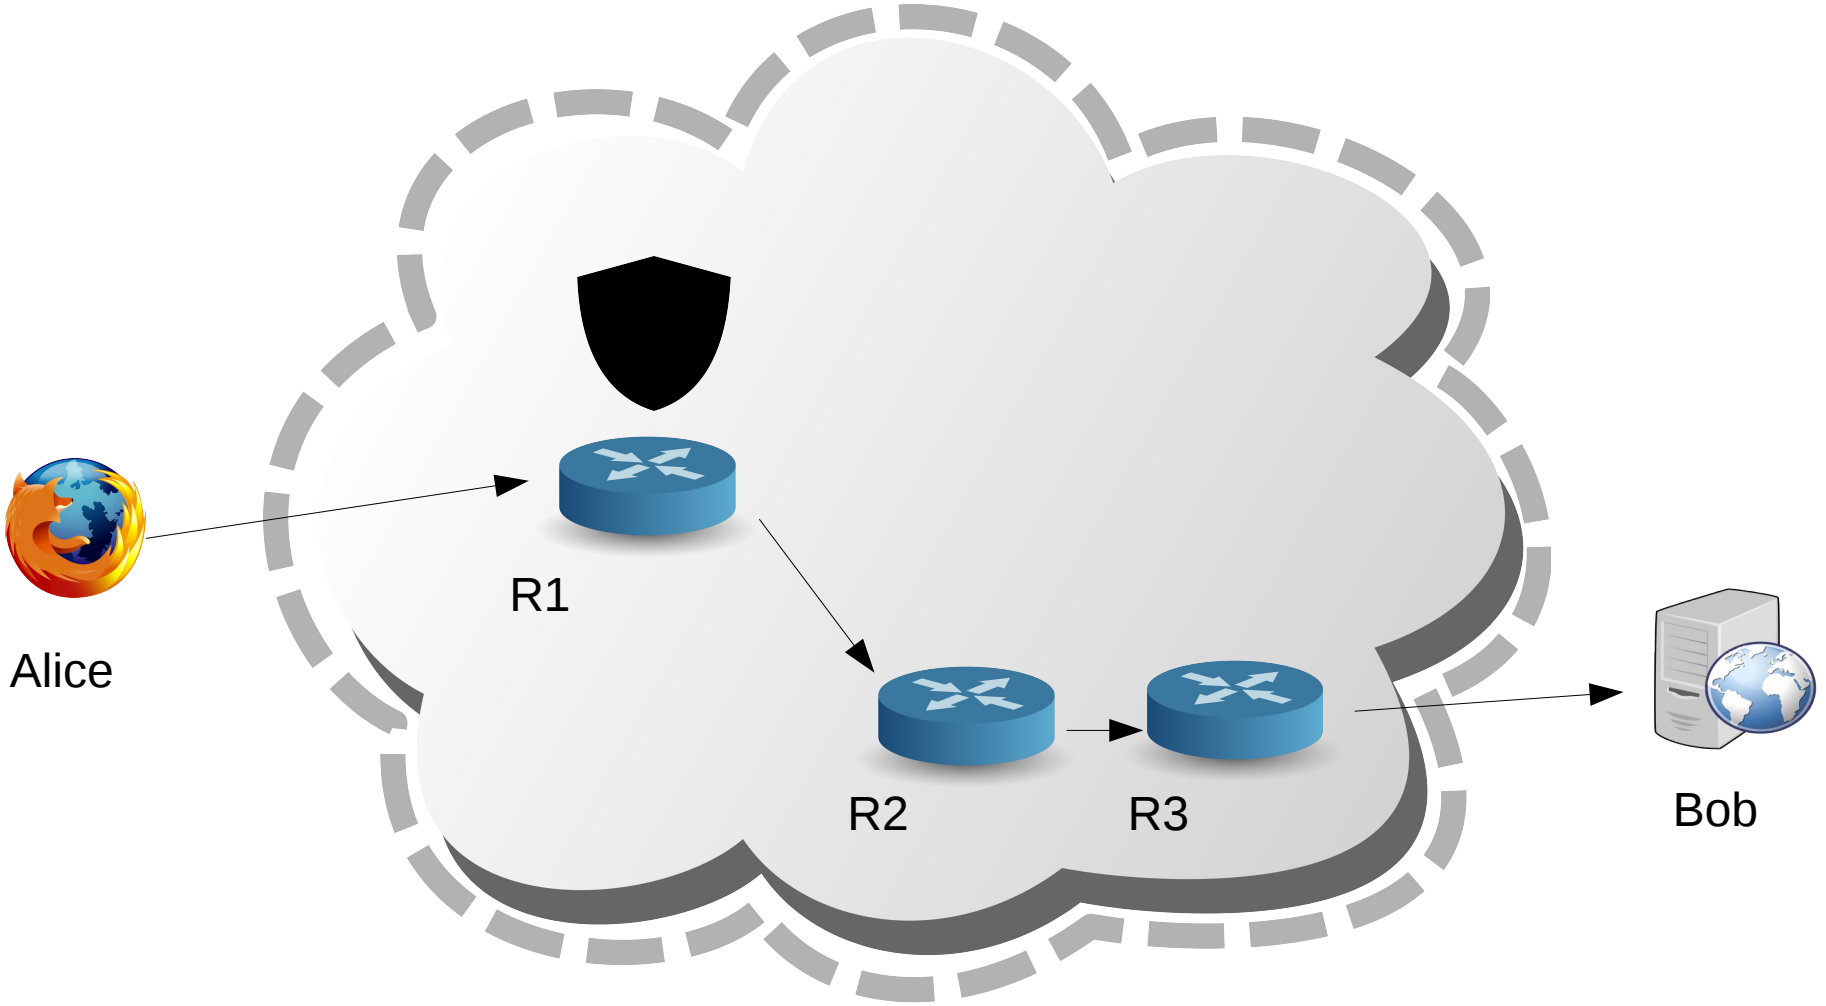

Guard = high uptime + high bandwidth

Every client has 3 Guard nodes

Guards

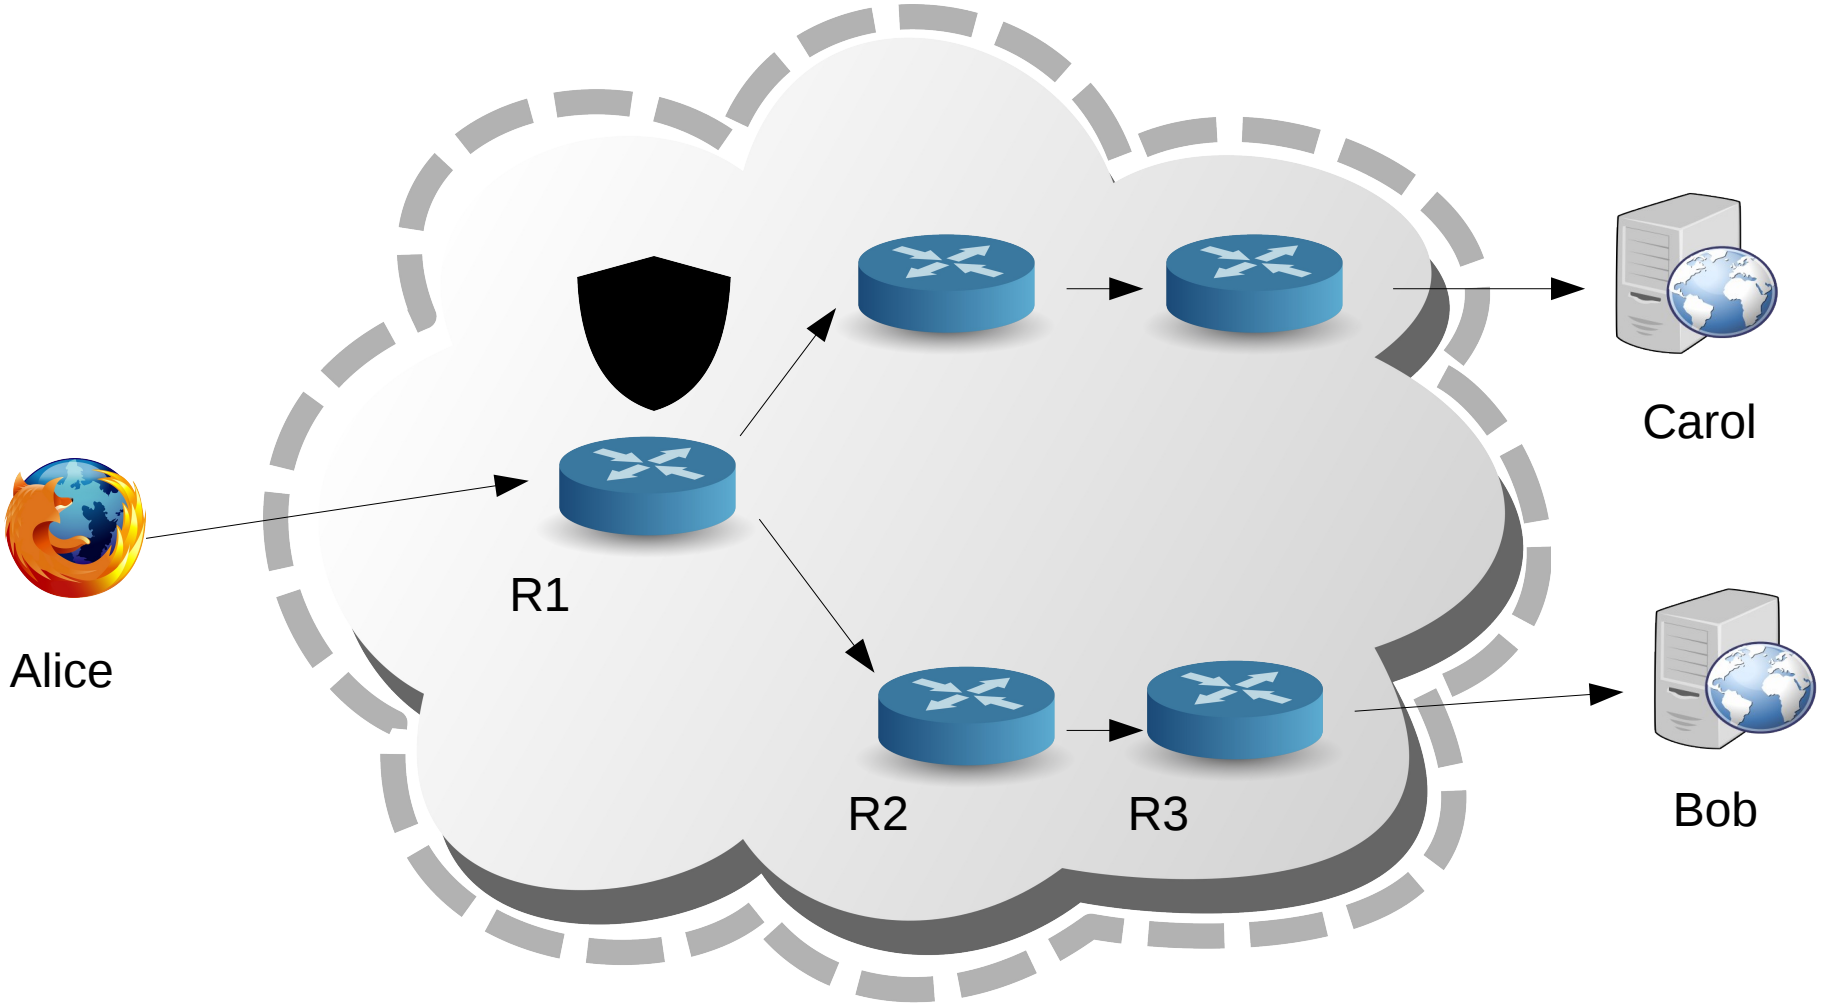

Guard = high uptime + high bandwidth

Every client has 3 Guard nodes

Examples of Tor HS

Public Library of US Diplomacy: Kissinger Cables

2013-04-08

The Kissinger Cables are part of today's launch of the WikiLeaks Public Library of US Diplomacy (PlusD), which holds the world's largest searchable collection of United States confidential, or formerly confidential, diplomatic communications. As of its launch on April 8, 2013 it holds 2 million records comprising approximately 1 billion words.

Detainee Policies

2012-10-24

WikiLeaks has begun releasing the 'Detainee Policies': more than 100 classified or otherwise restricted files from the United States Department of Defense covering the rules and procedures for detainees in U.S. military custody. Over the next month, WikiLeaks will release in chronological

In Wikile

U.K. (2008)
contacts li

WikiLeaks
individuals
whose men
number of
details of al
individuals
fascists" wh

China (200

Examples of Tor HS

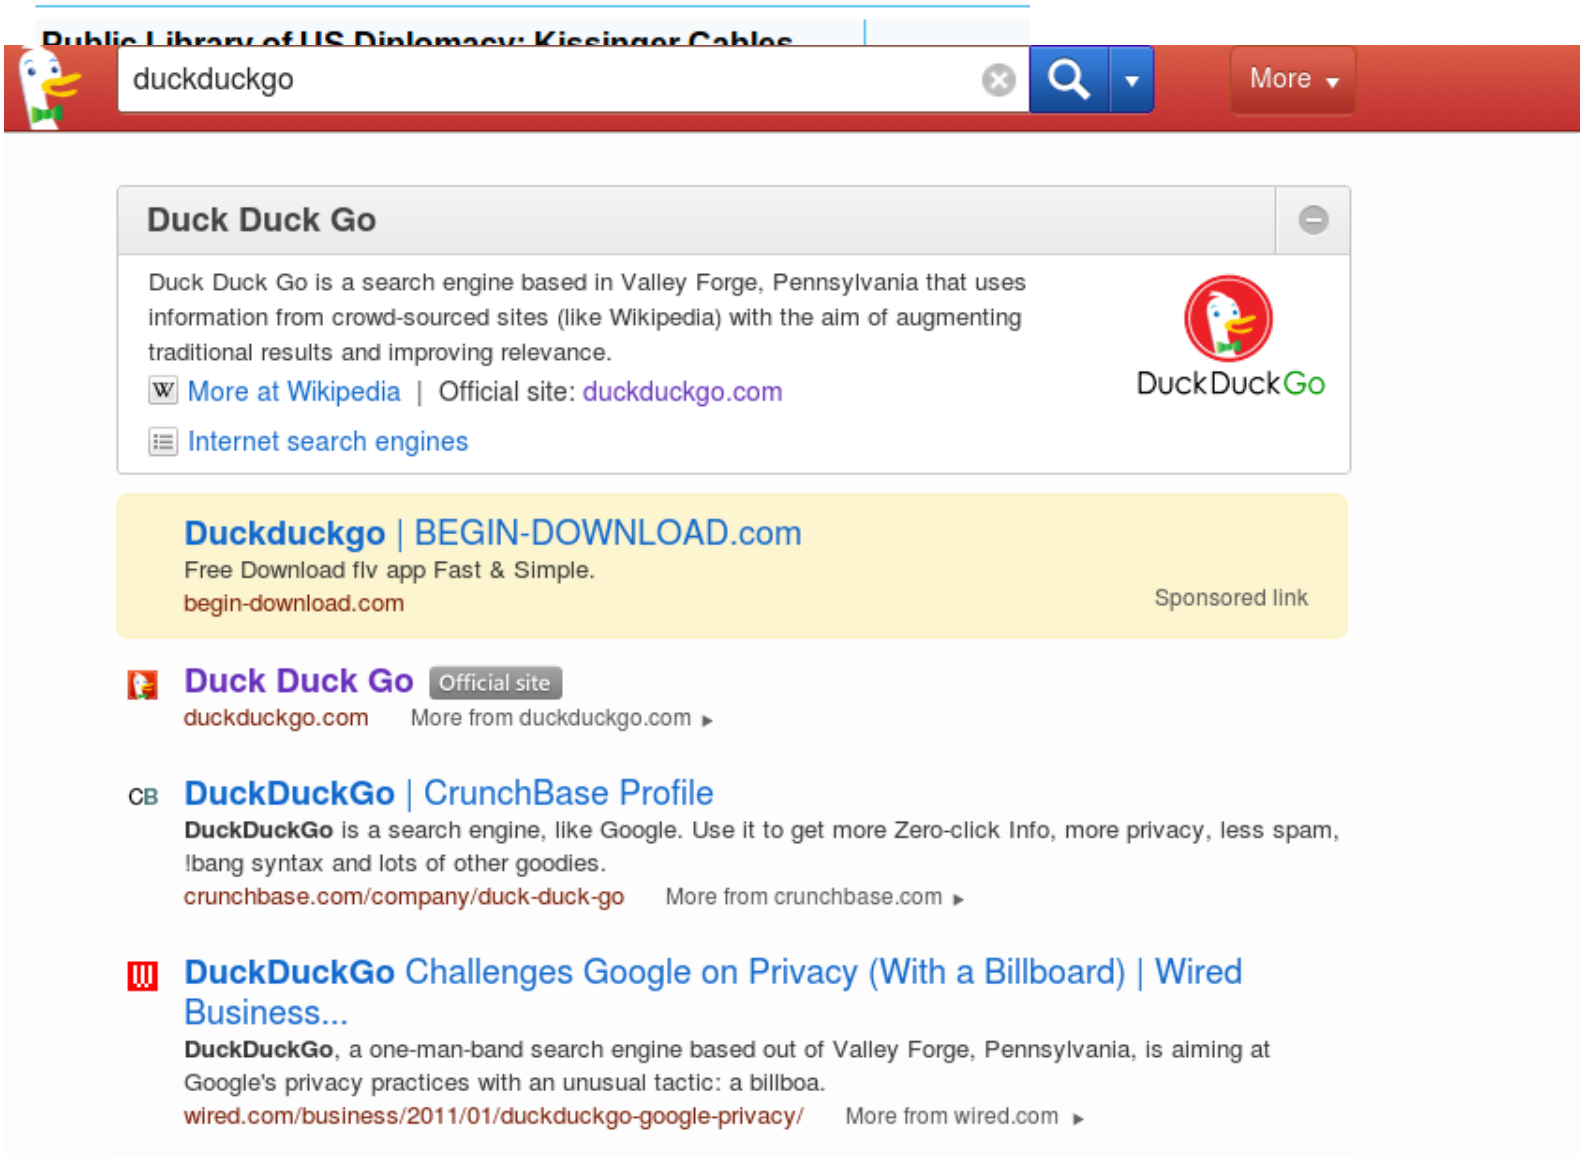

Examples of Tor HS

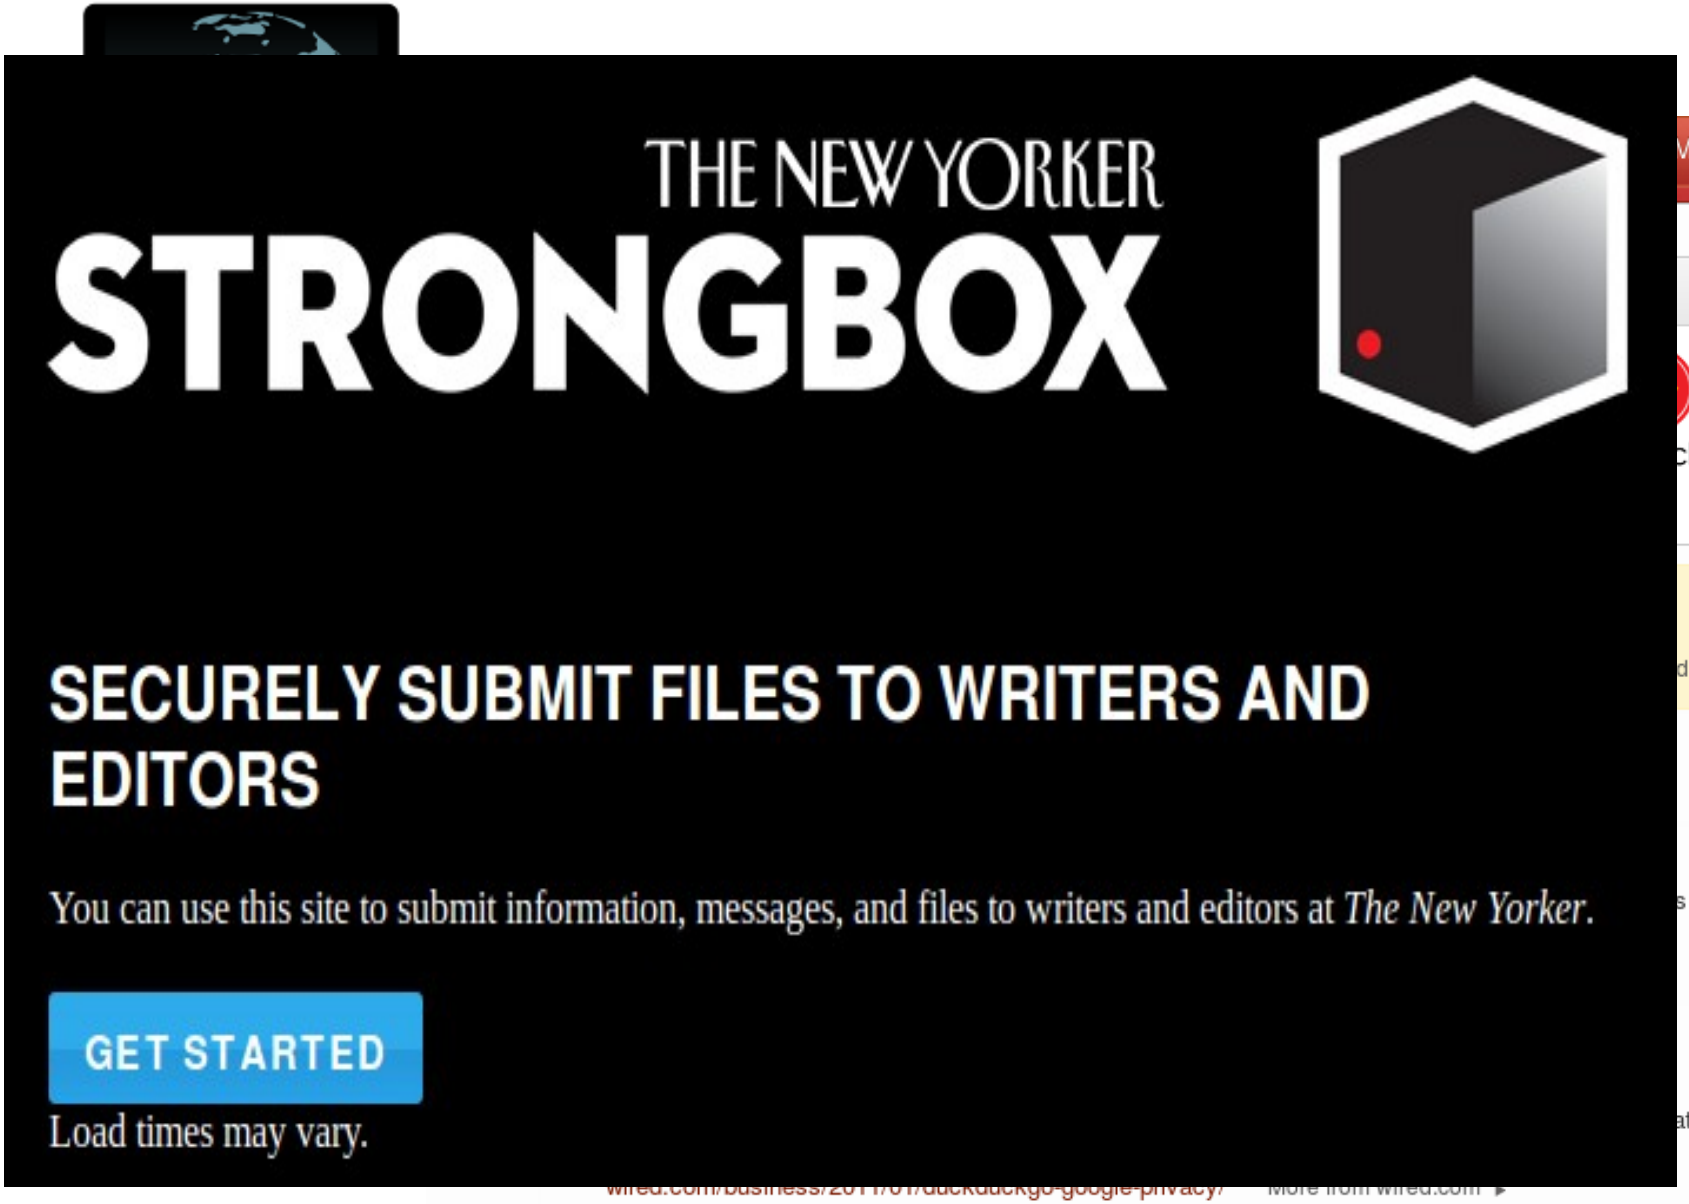

Examples of Tor HS

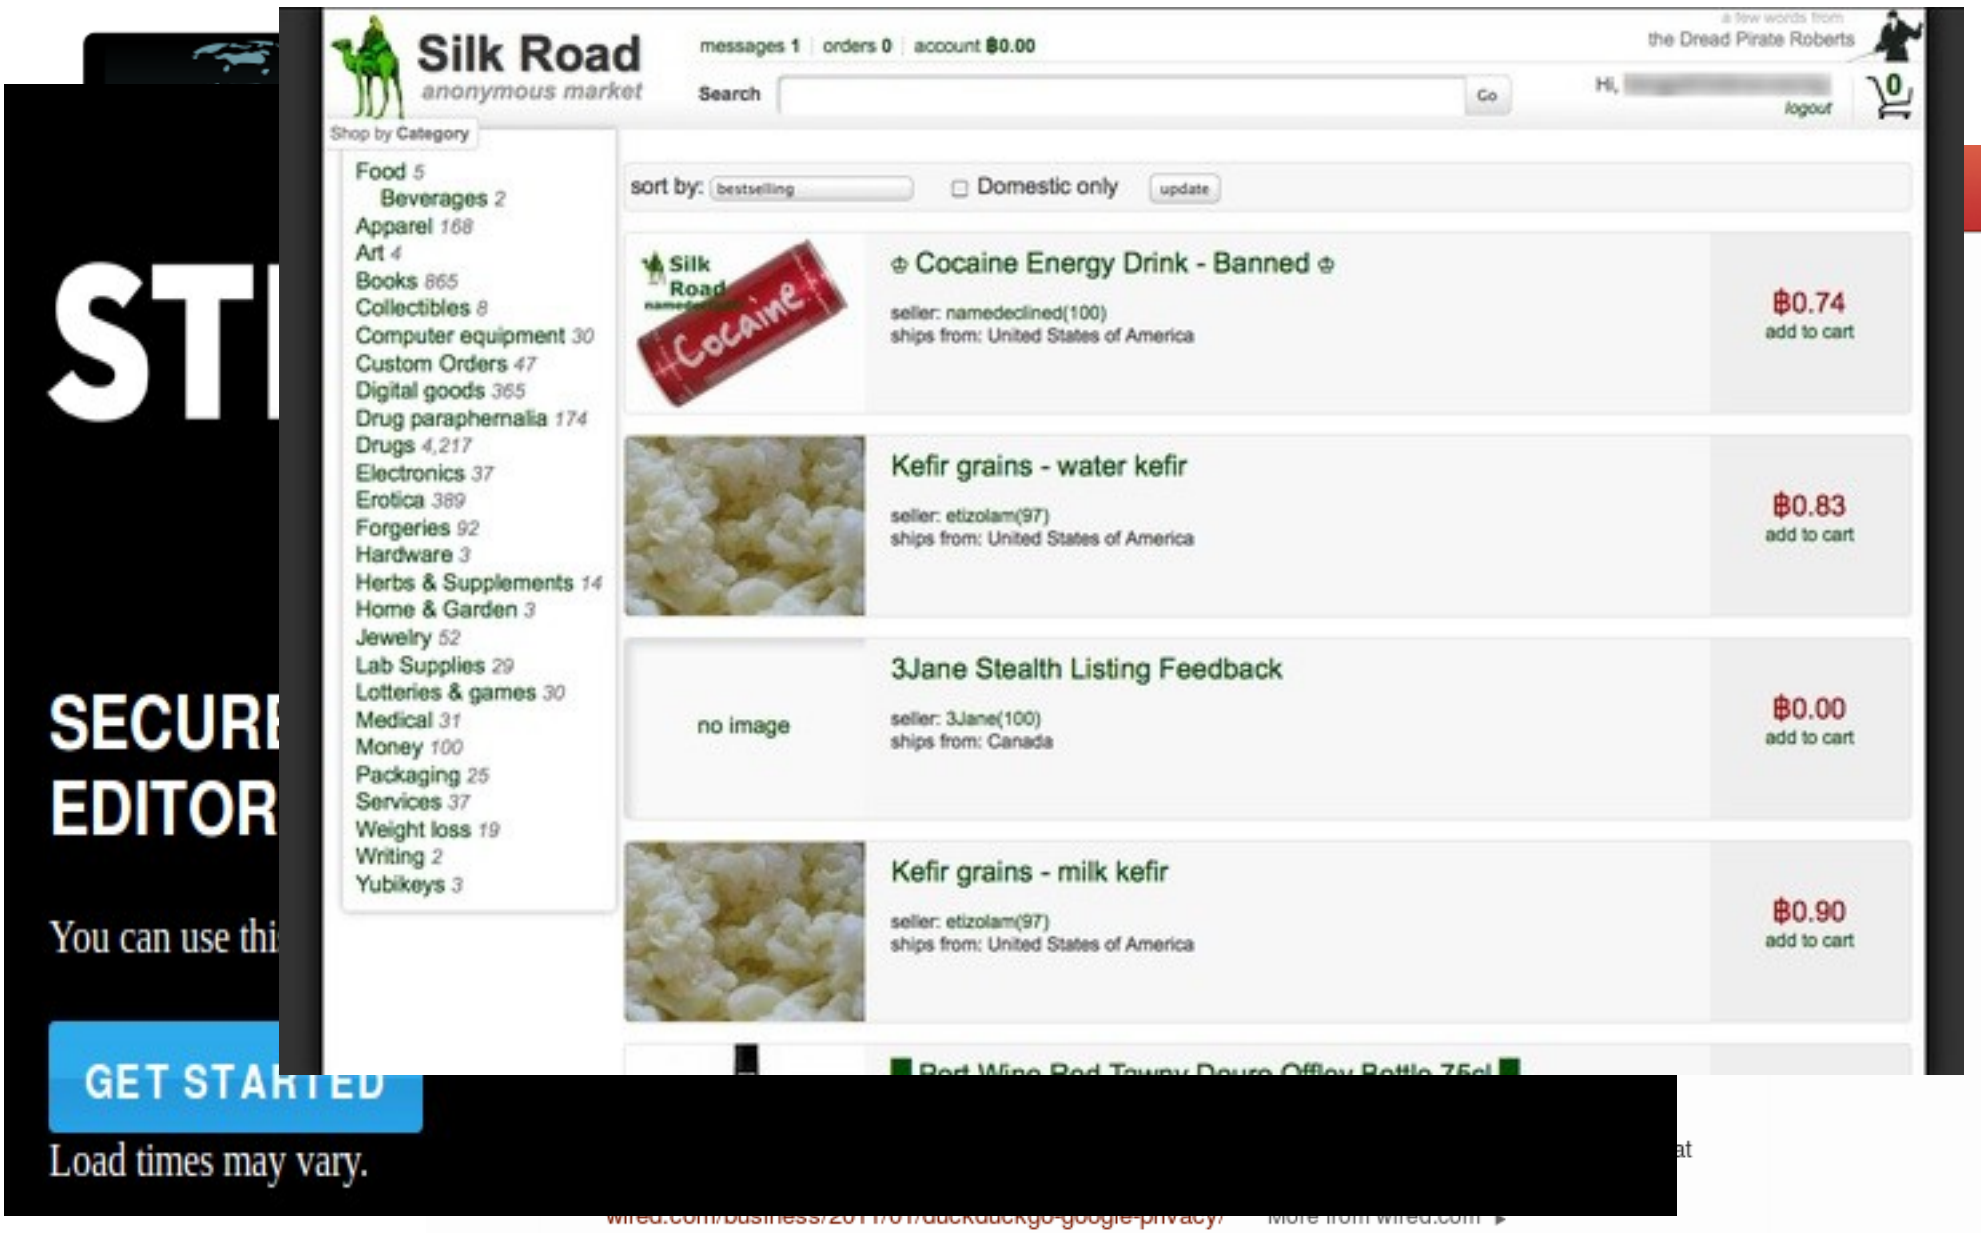

Examples of Tor HS

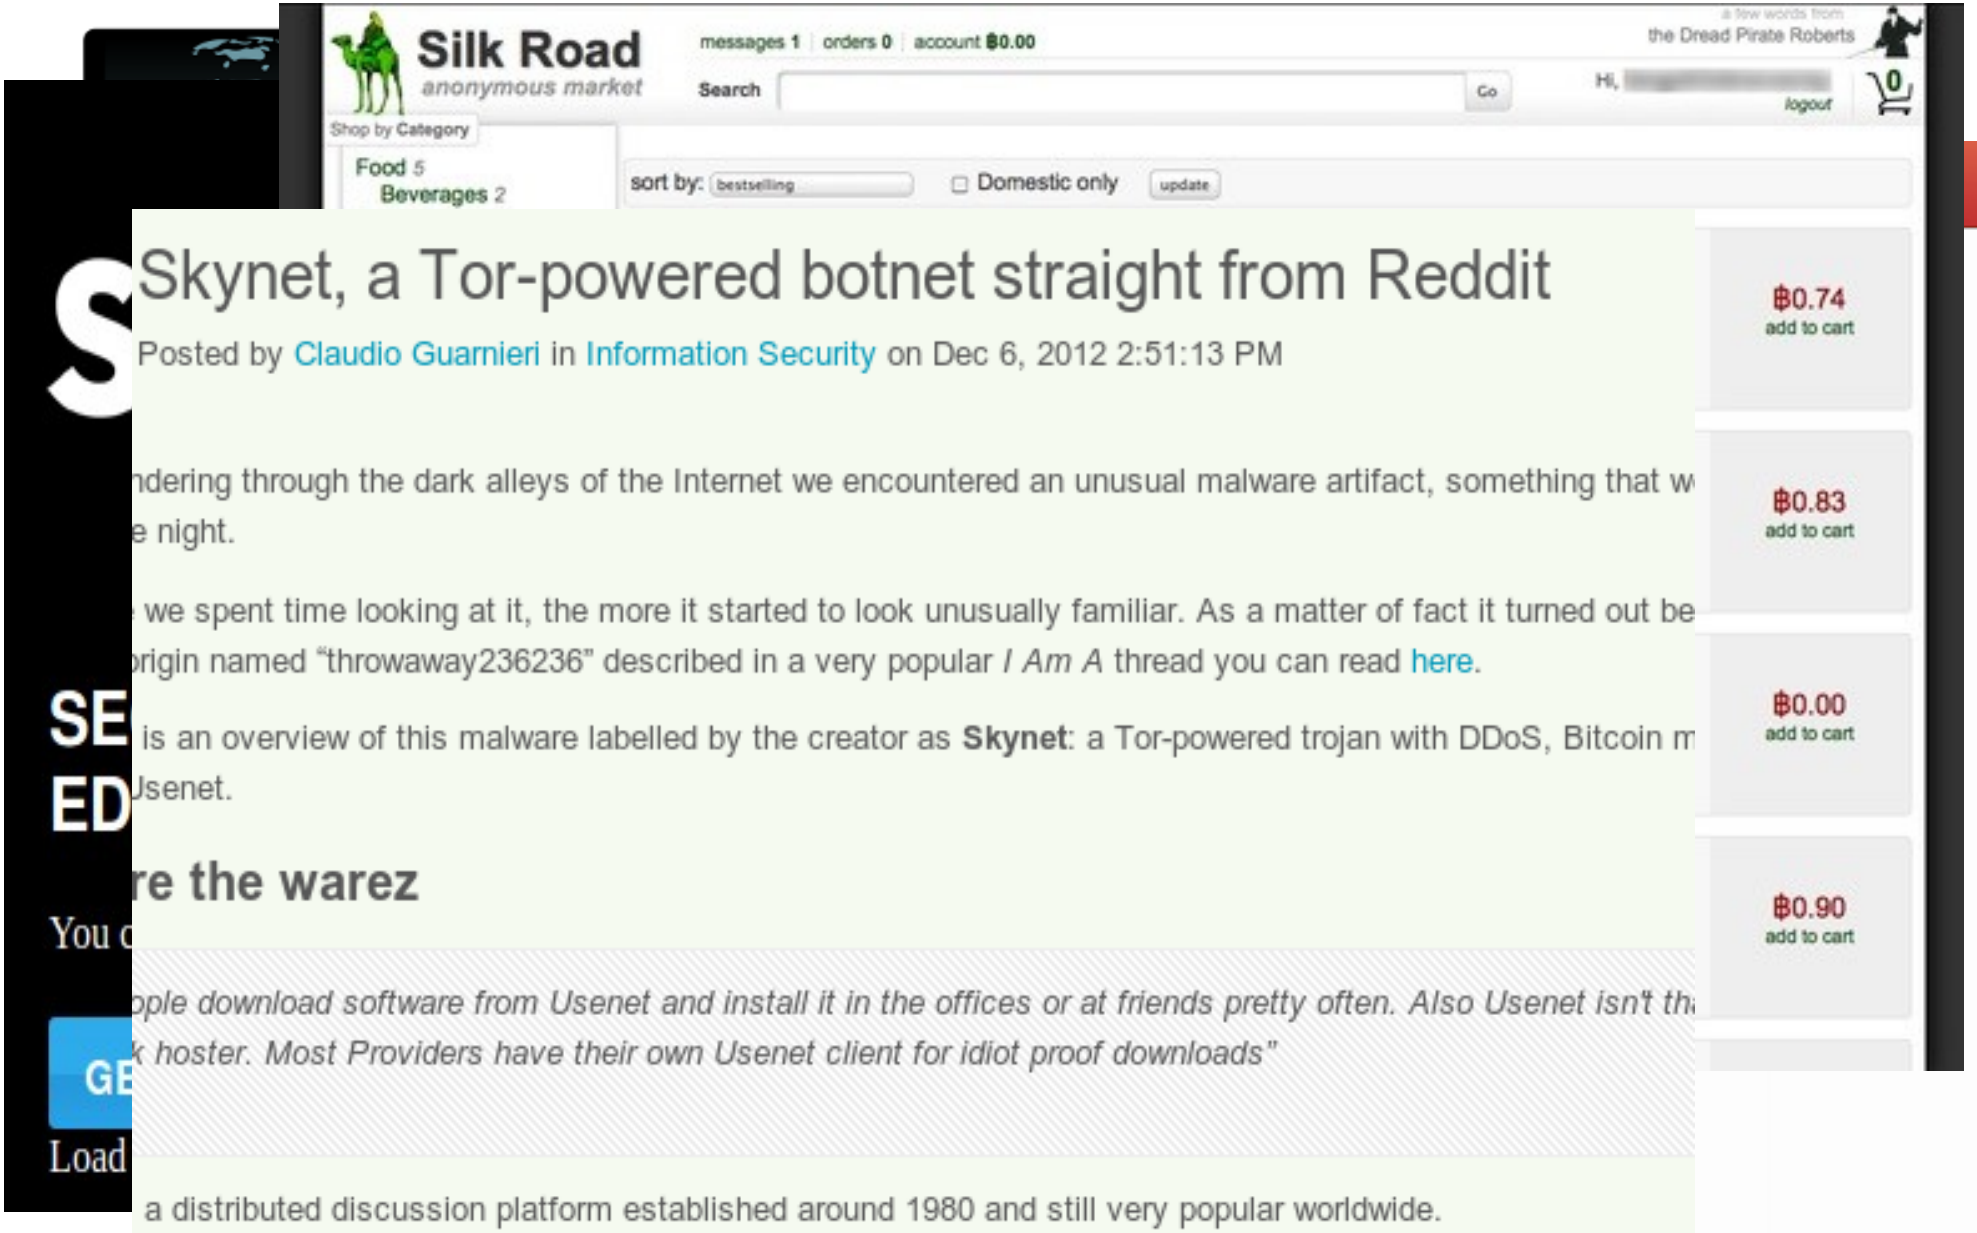

Tor rendezvous protocol

Step1: Bob picks some introduction points and builds circuits to them.

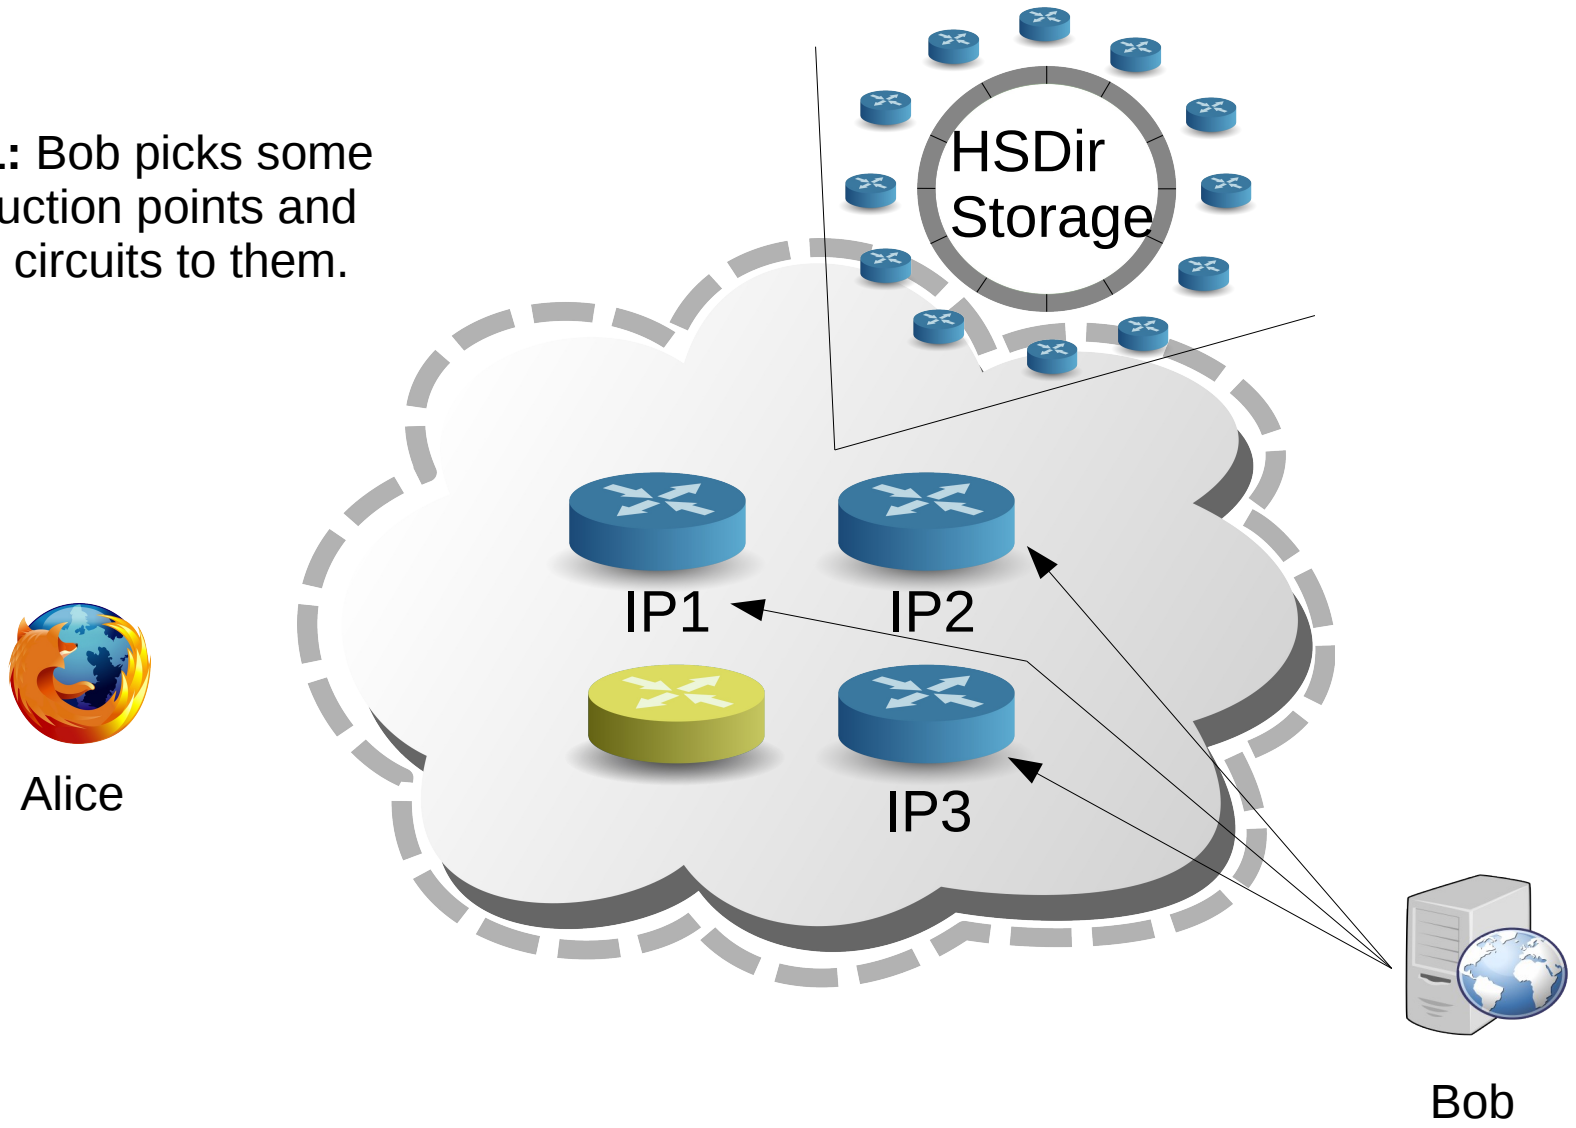

Tor rendezvous protocol

Step2: Bob advertises his hidden service – <z>.onion – at the database.

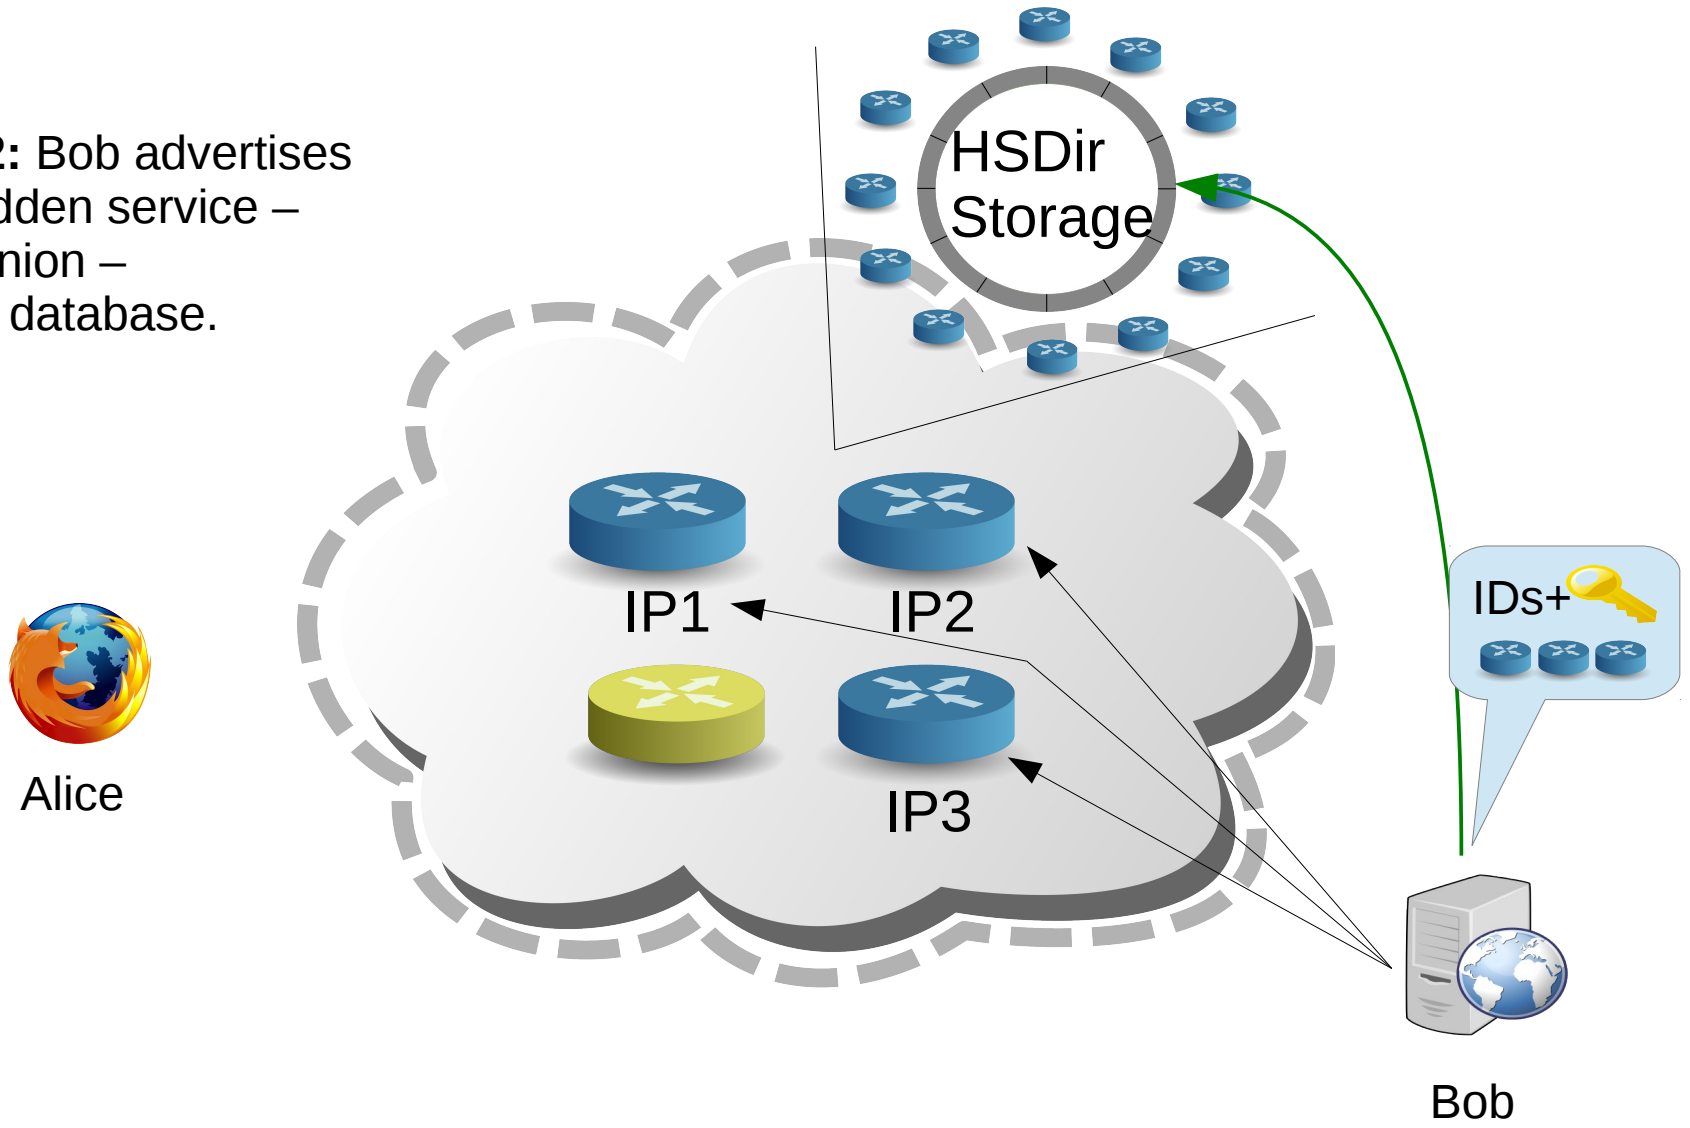

Tor rendezvous protocol

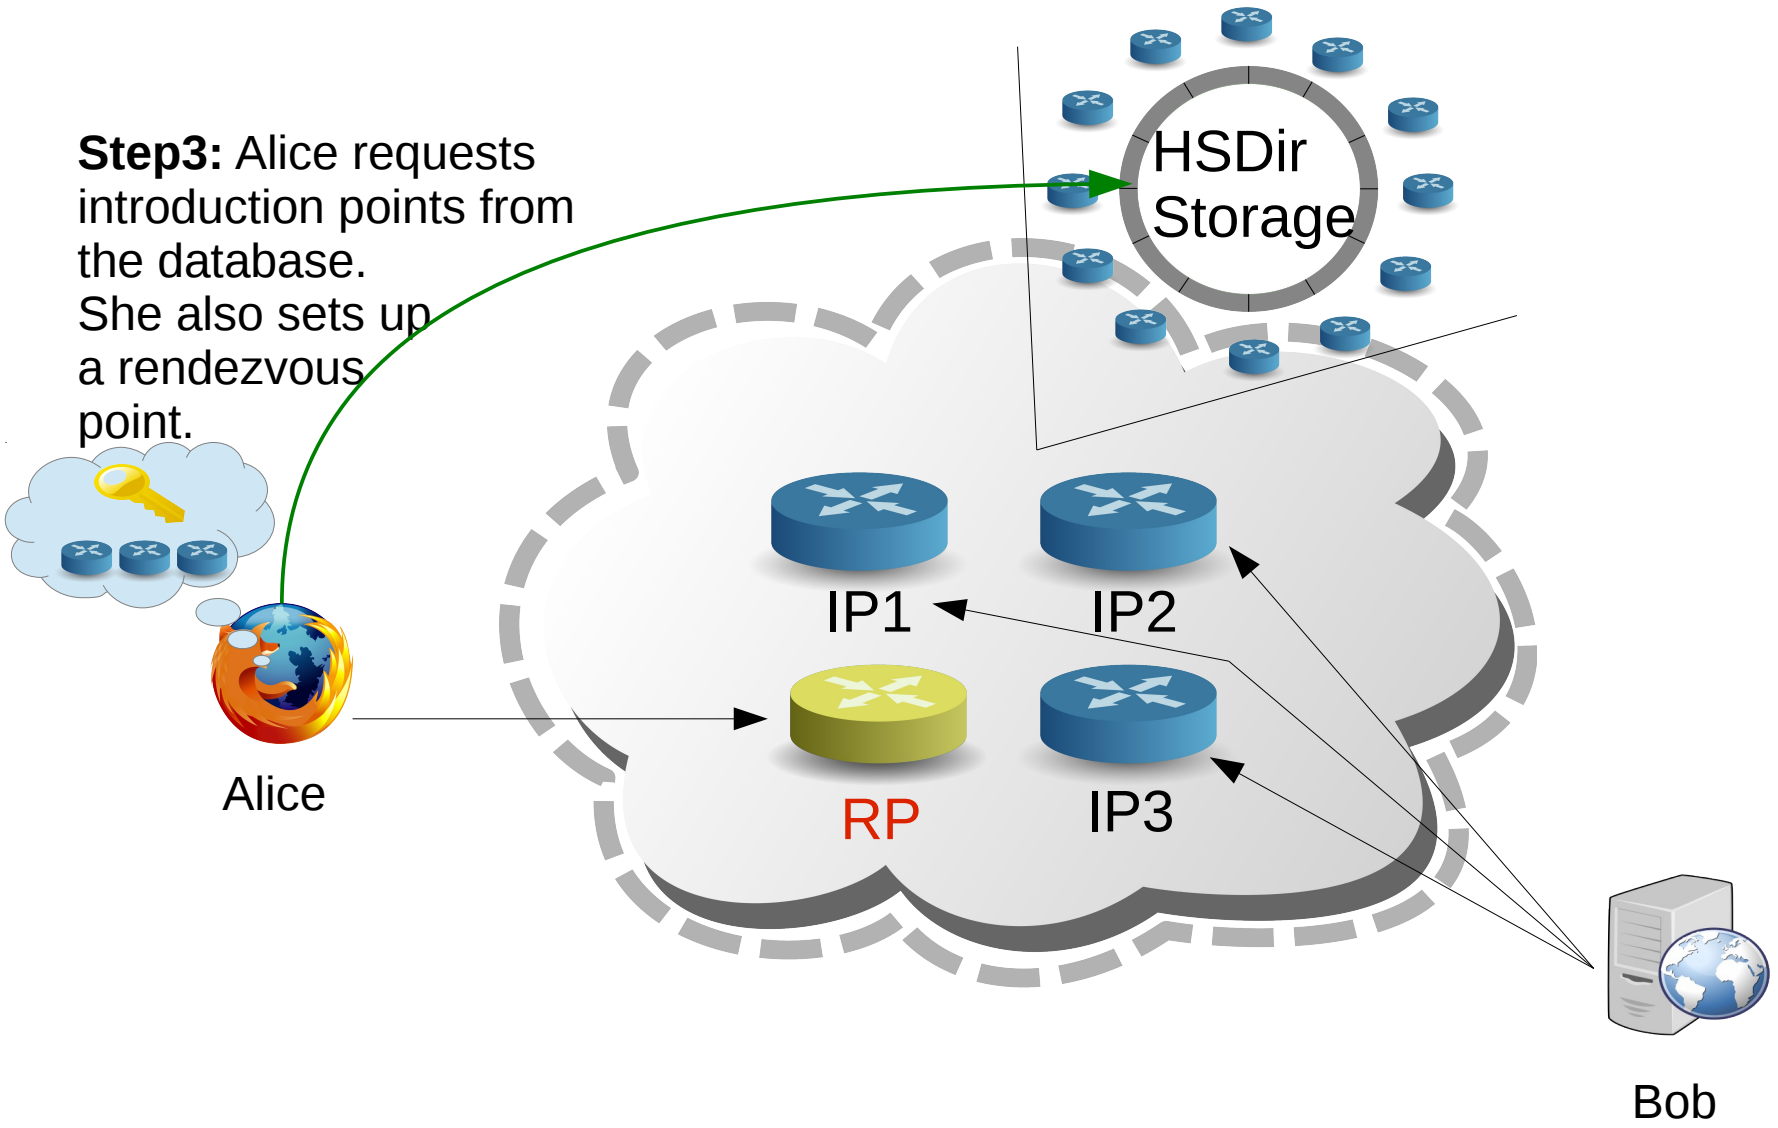

Tor rendezvous protocol

Step4: Alice sends a message to Bob listing the rendezvous point and asks the introduction points from to deliver it.

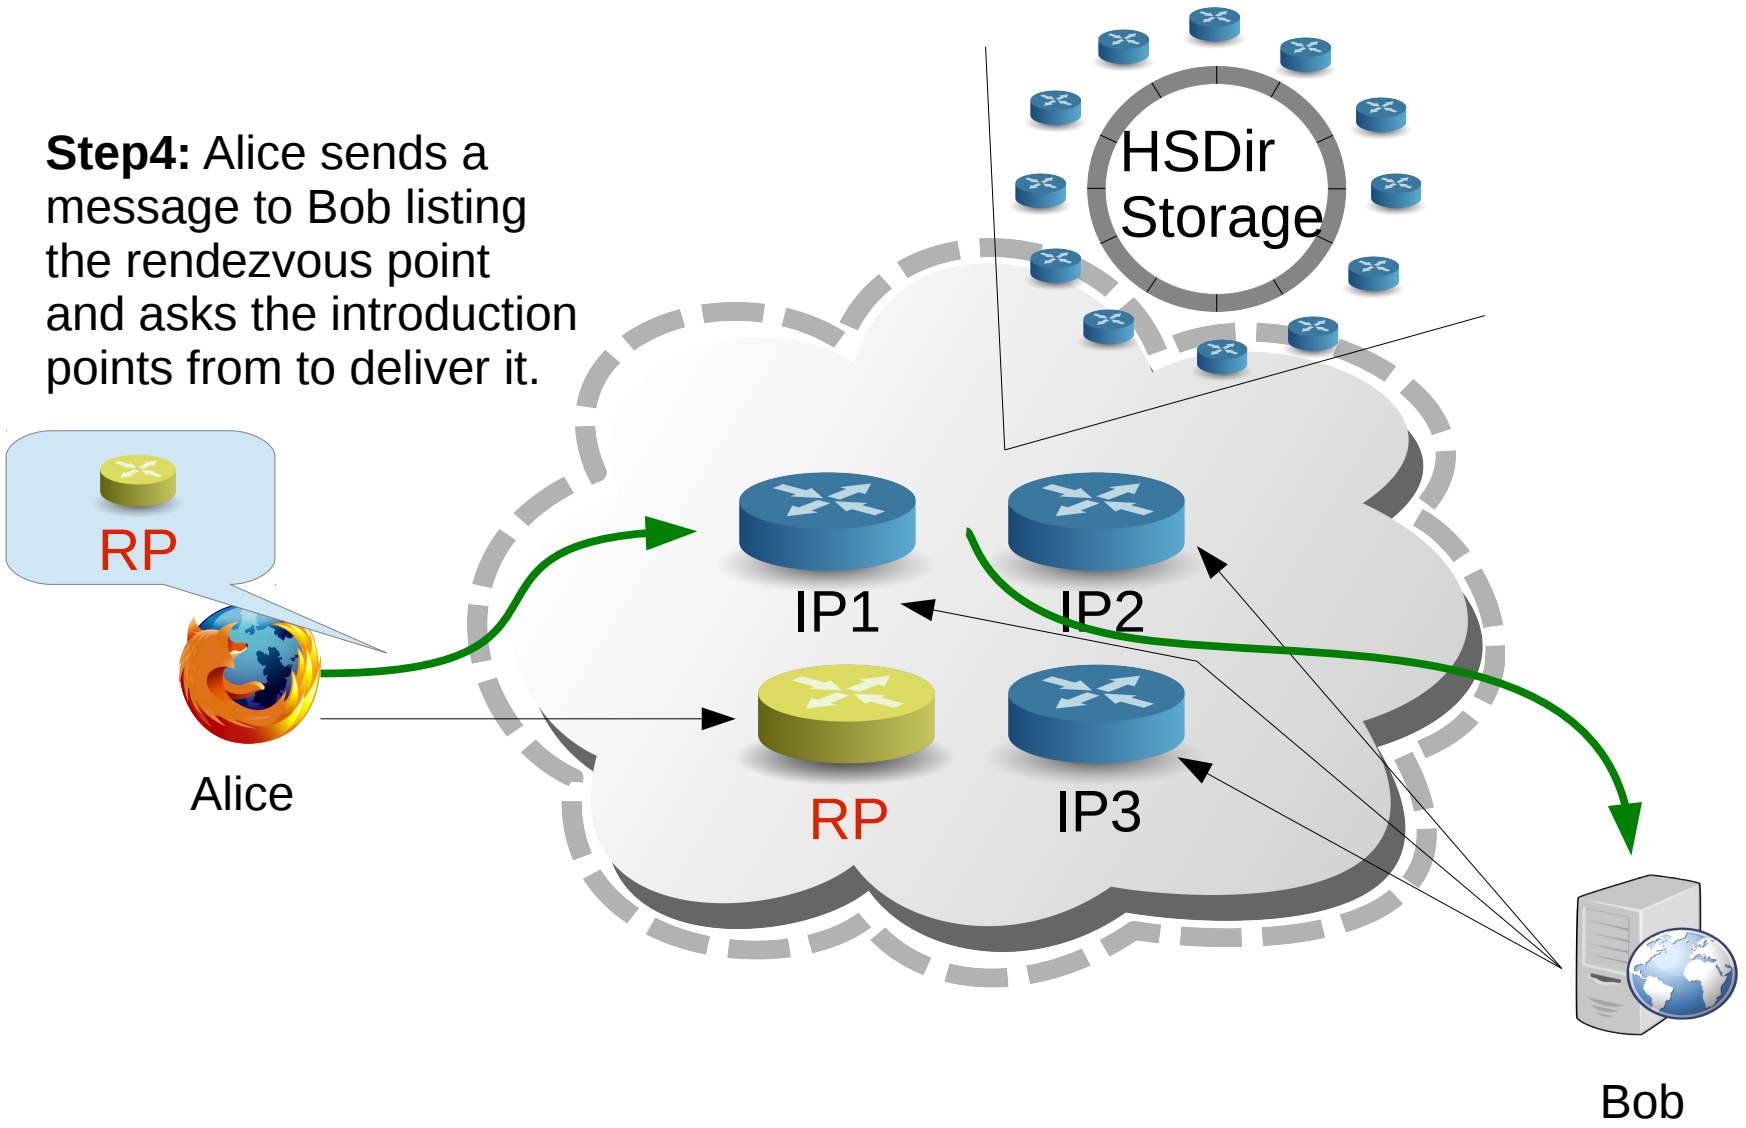

Tor rendezvous protocol

Step5: Alice and Bob Connect at the Rendezvous point

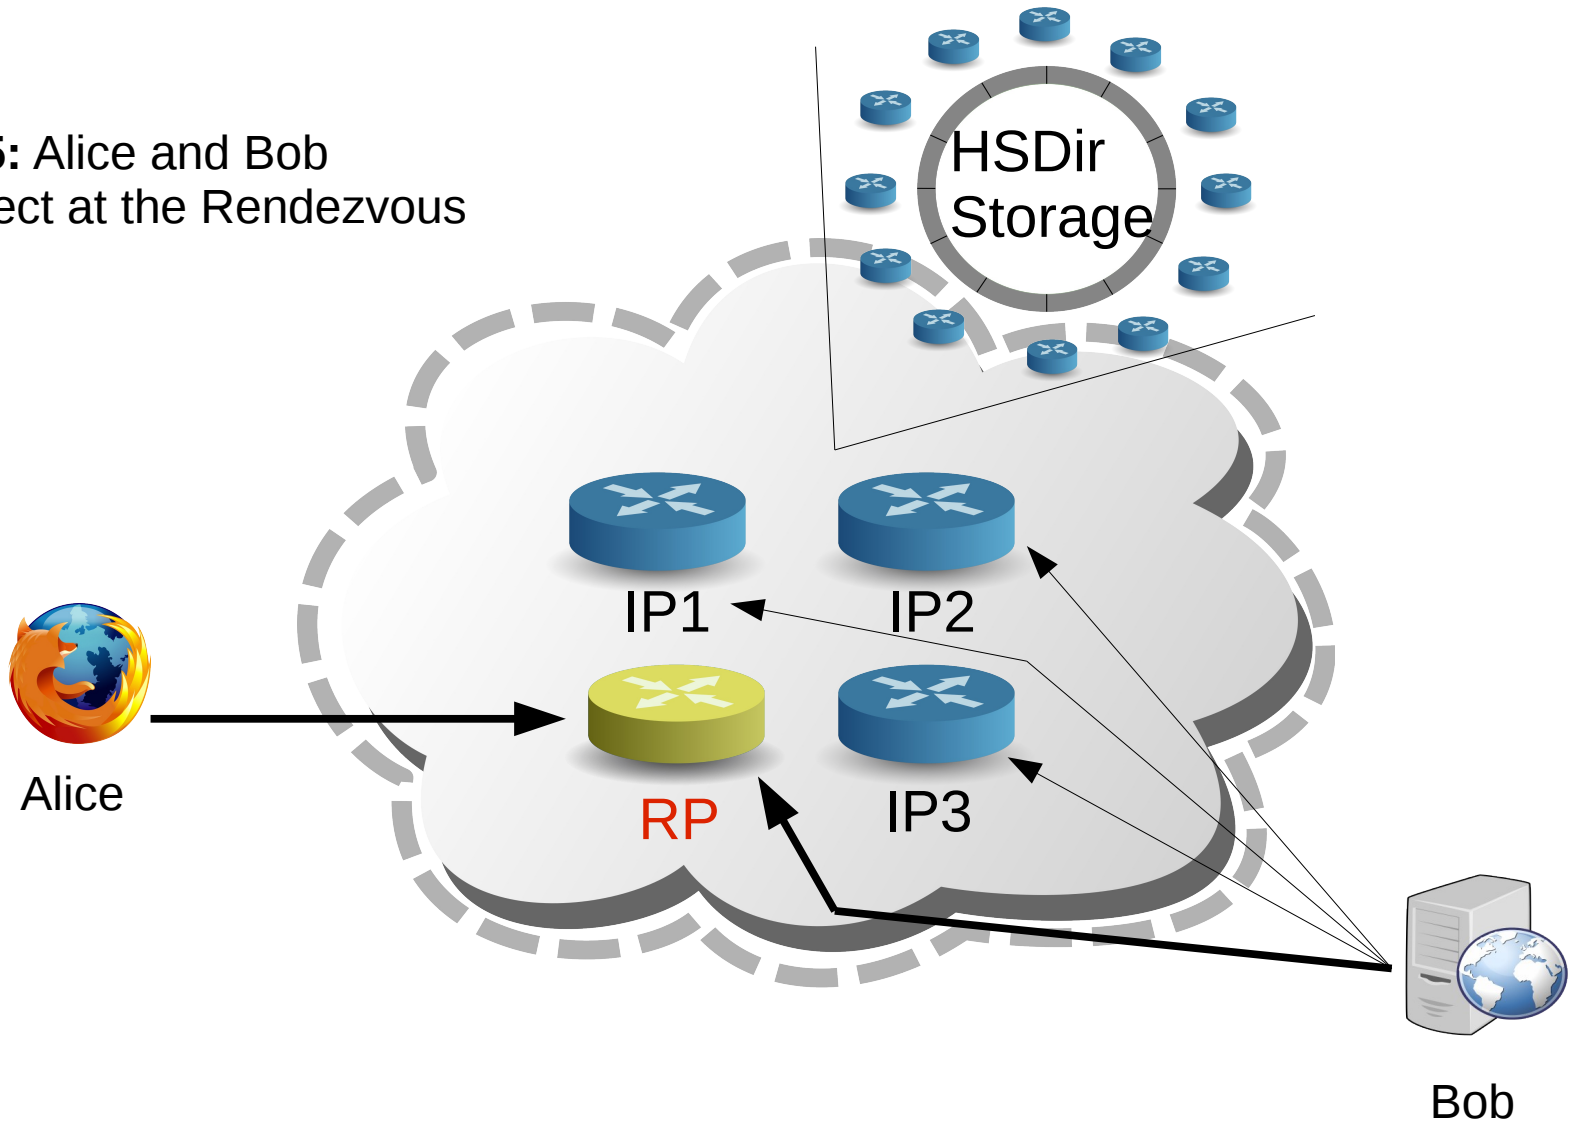

Tor rendezvous protocol

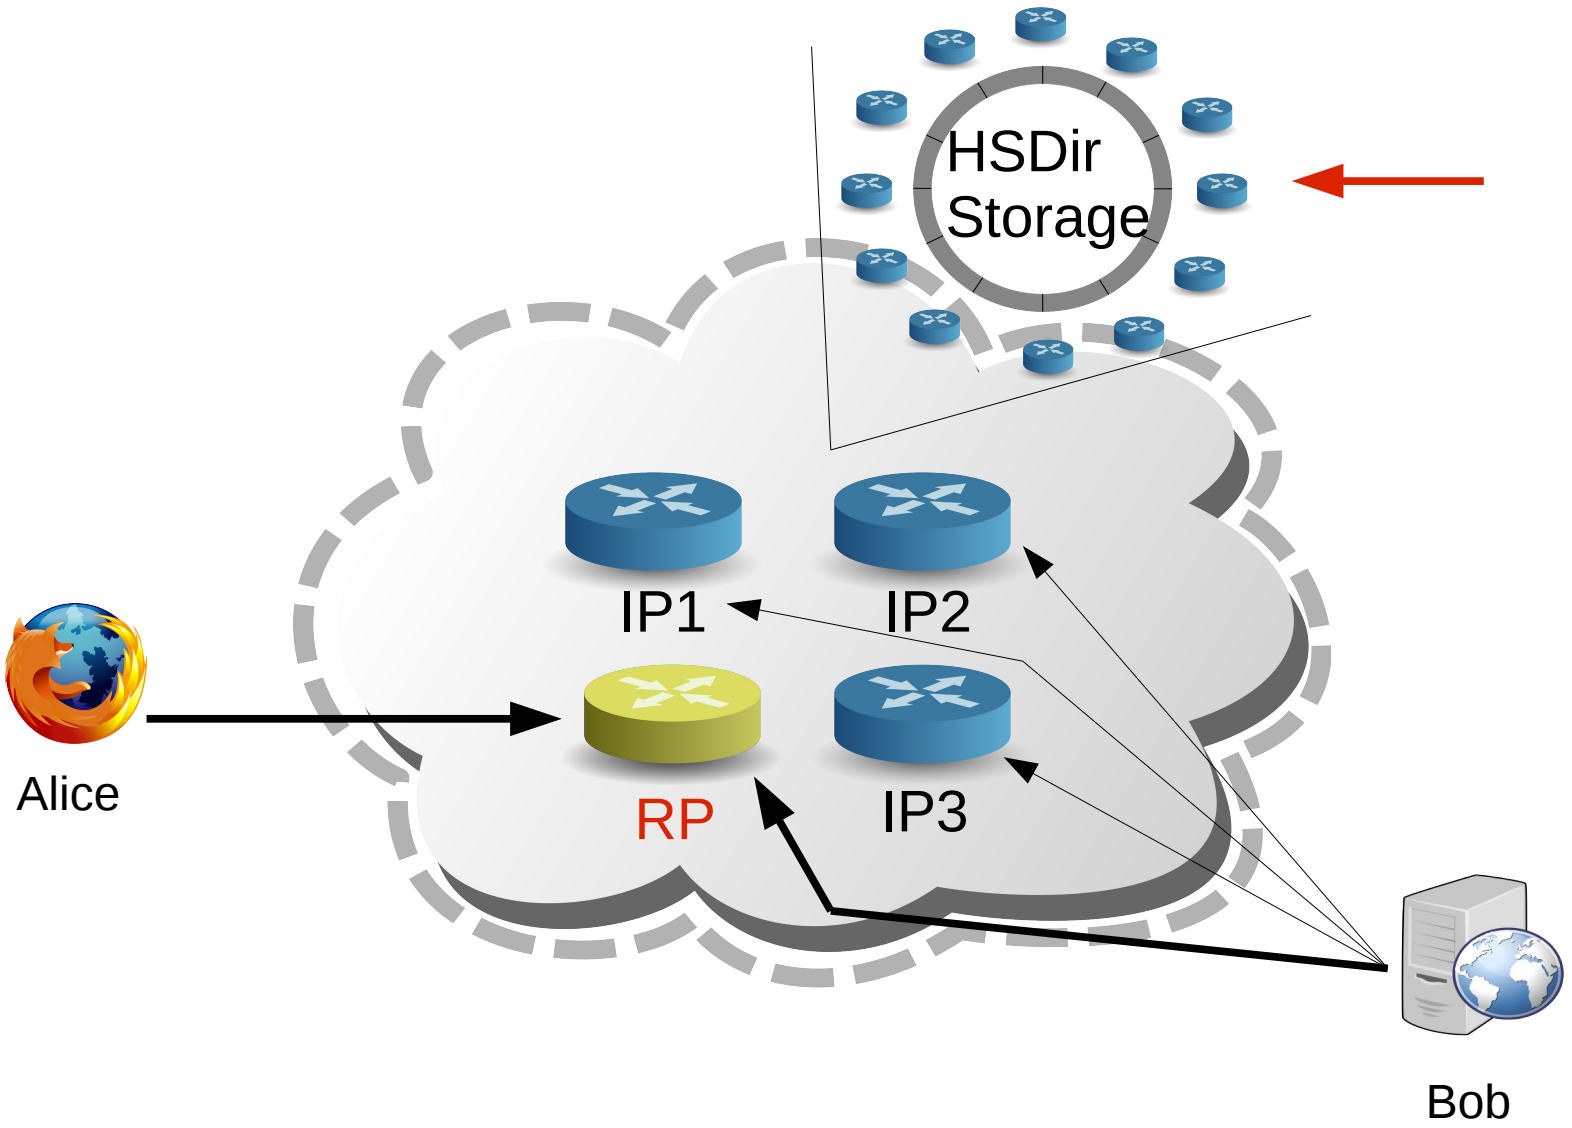

Responsible hidden service directories

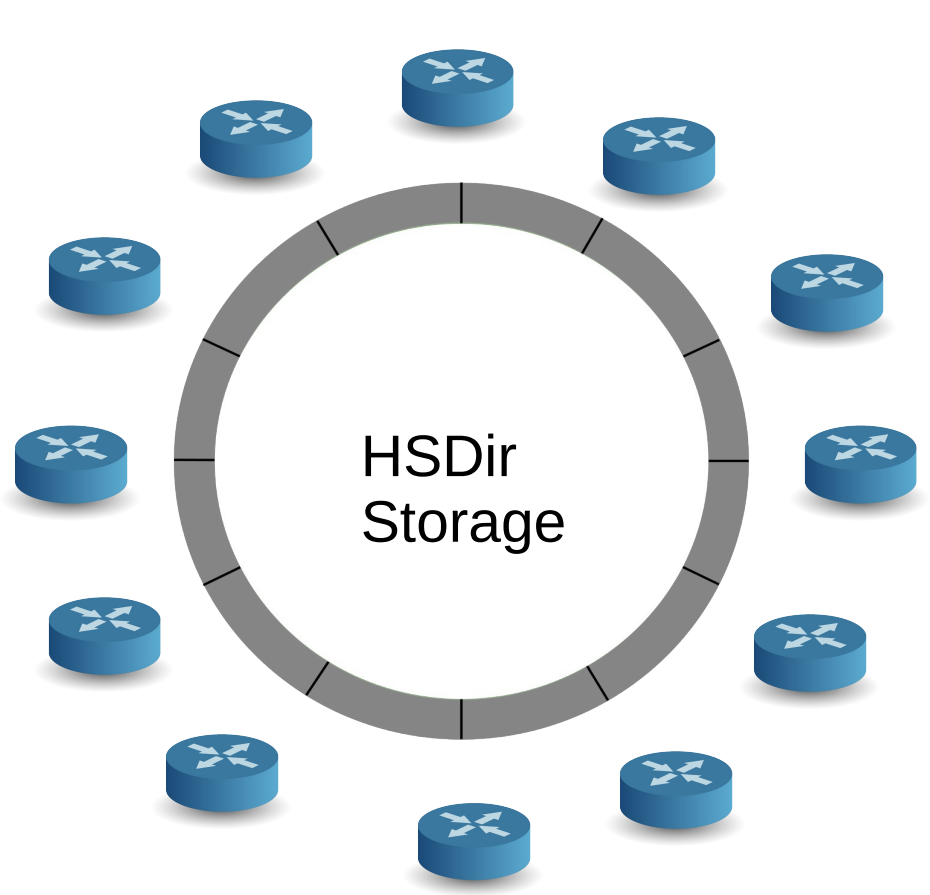

= HSDir = 25 hours of uptime

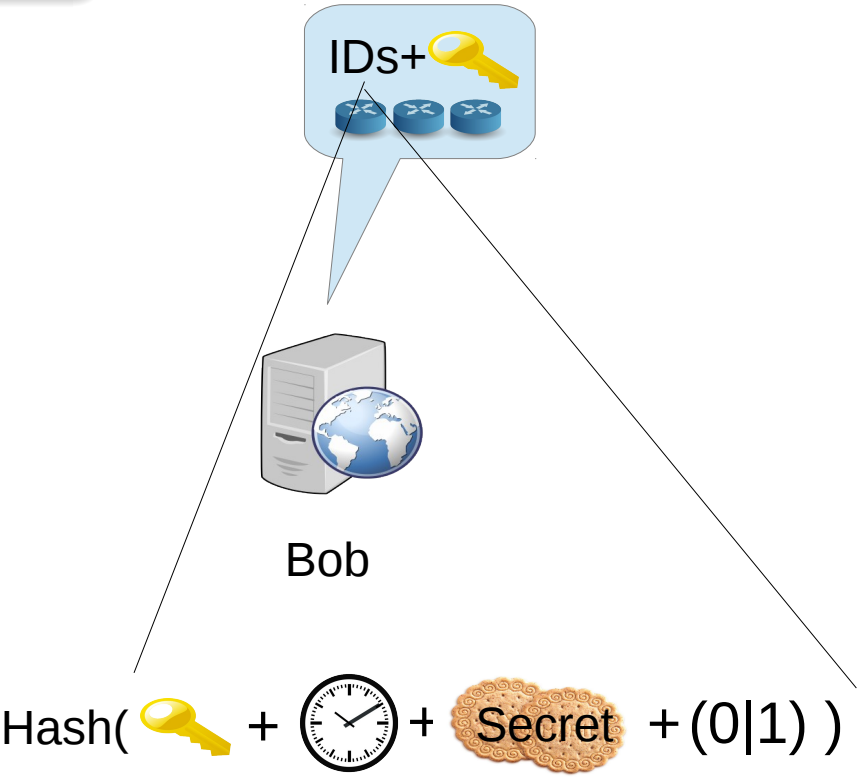

Responsible hidden service directories

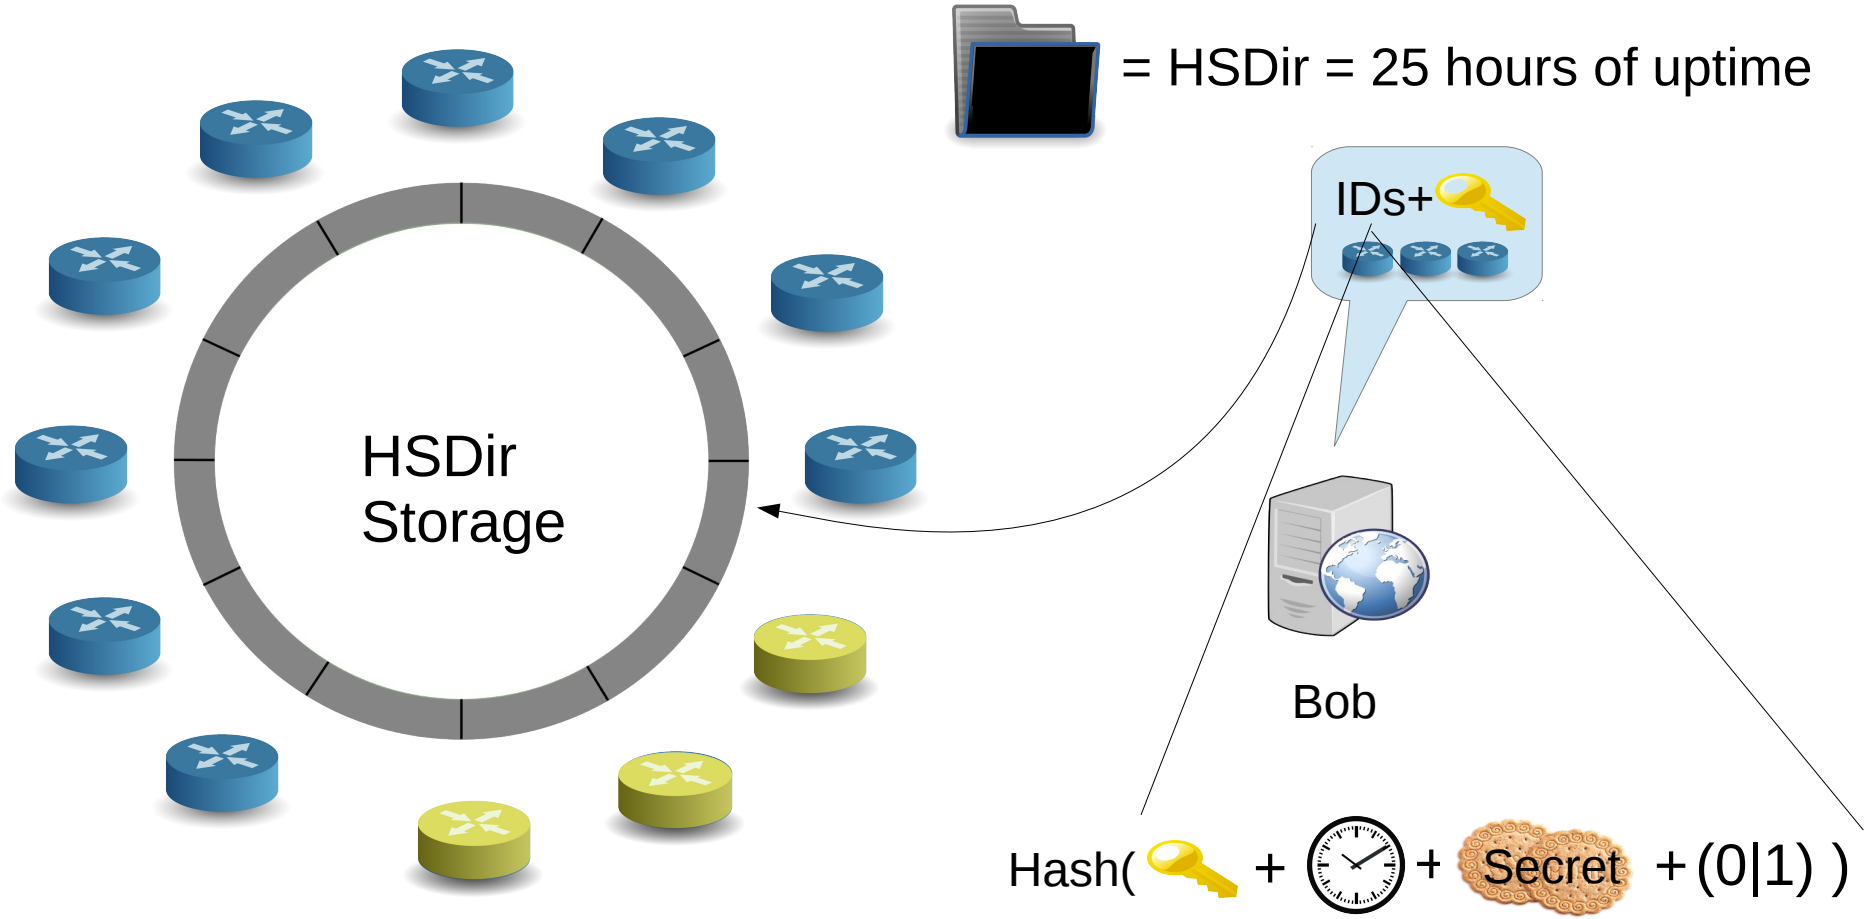

Responsible hidden service directories

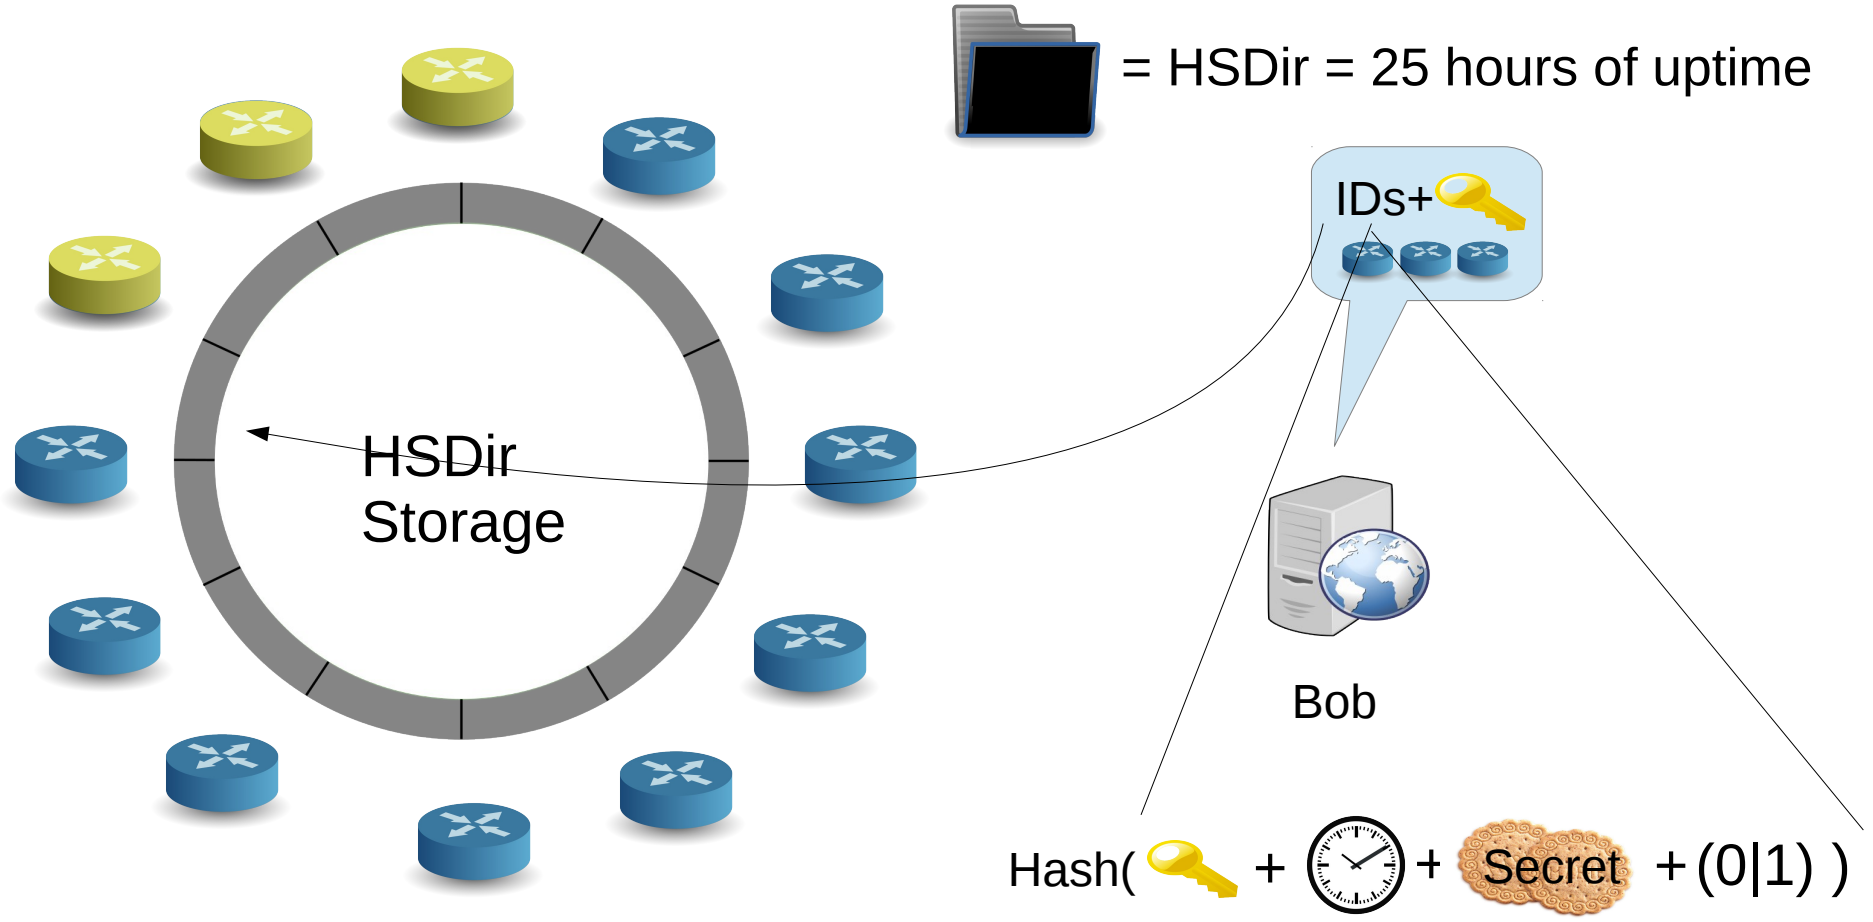

- IDs change every 24 hours at some time during the day
- Re-upload every hour

Overview

- Background
- **Measuring the popularity of hidden services**
- **DoSing hidden services.**
- Harvesting onion addresses.
- Revealing the guards.
- Opportunistic deanonymisation.

Impersonating Hidden service directory

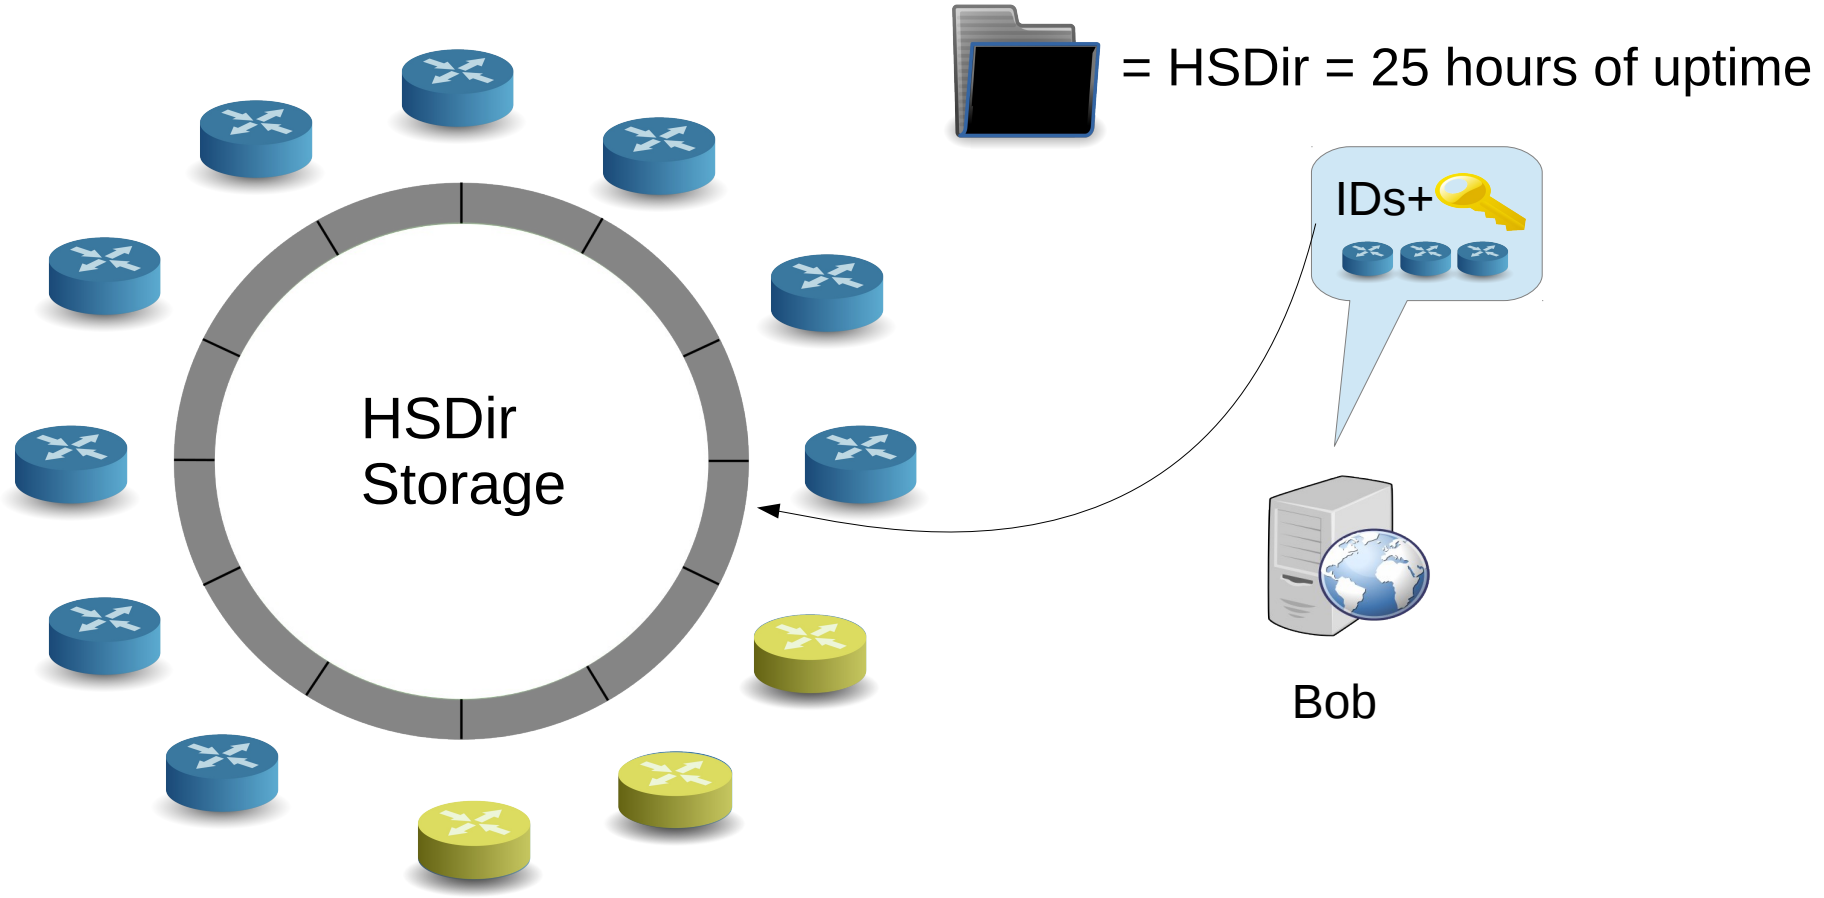

Impersonating Hidden service directory

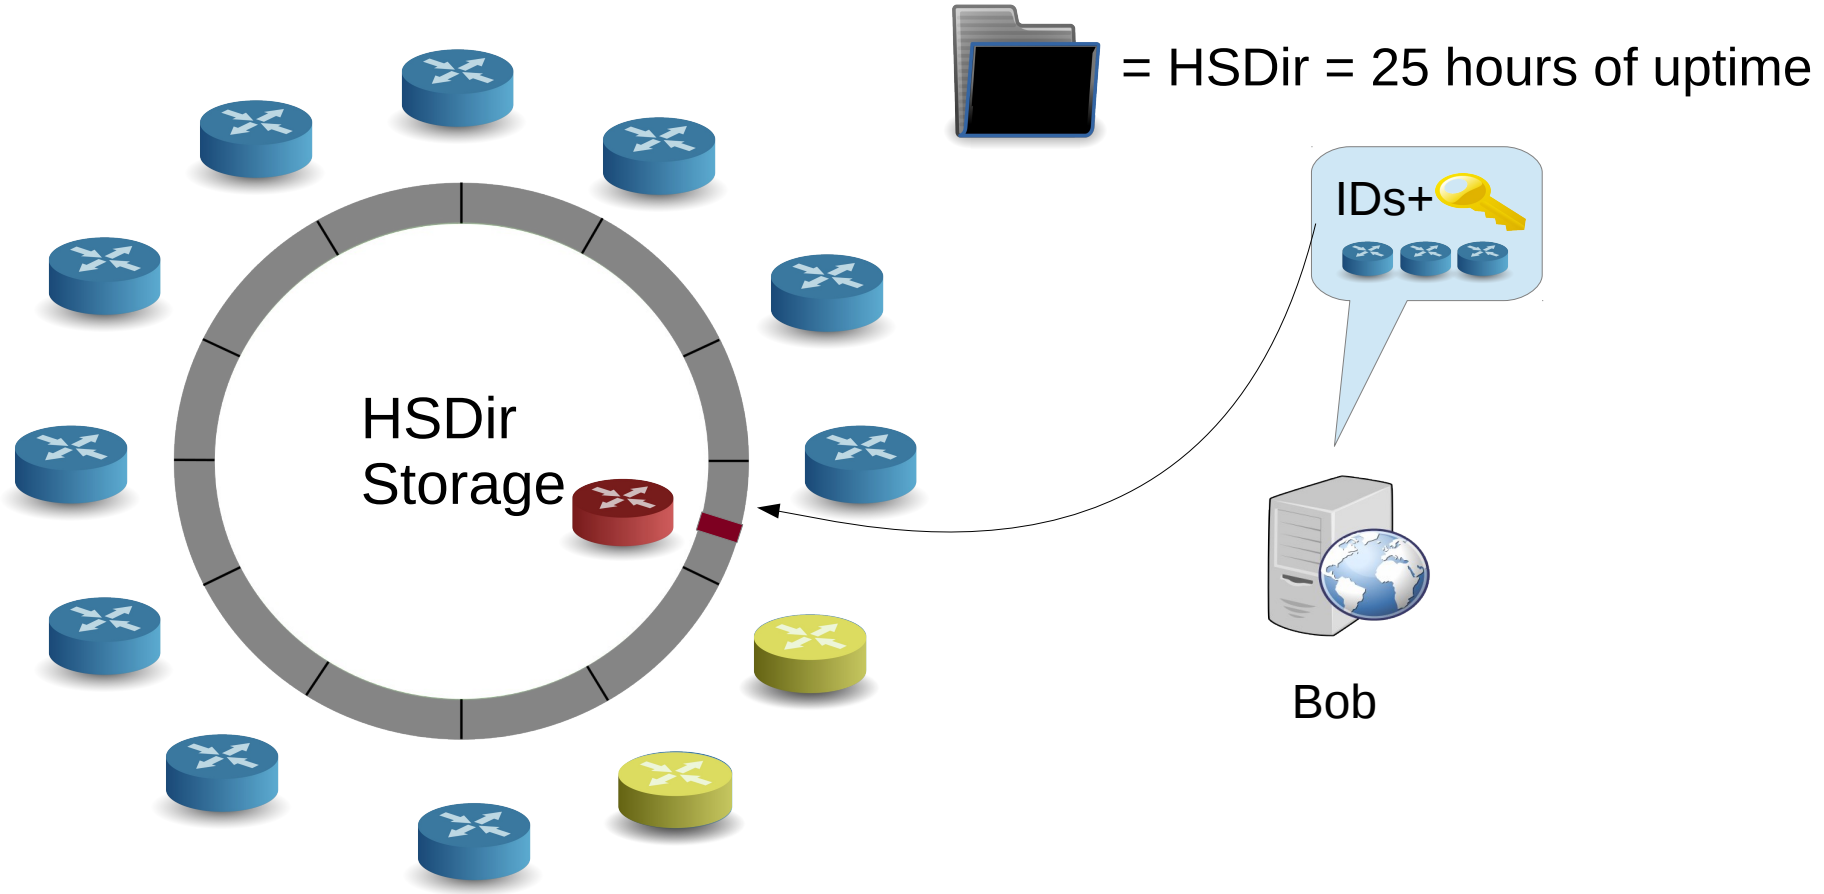

Impersonating Hidden service directory

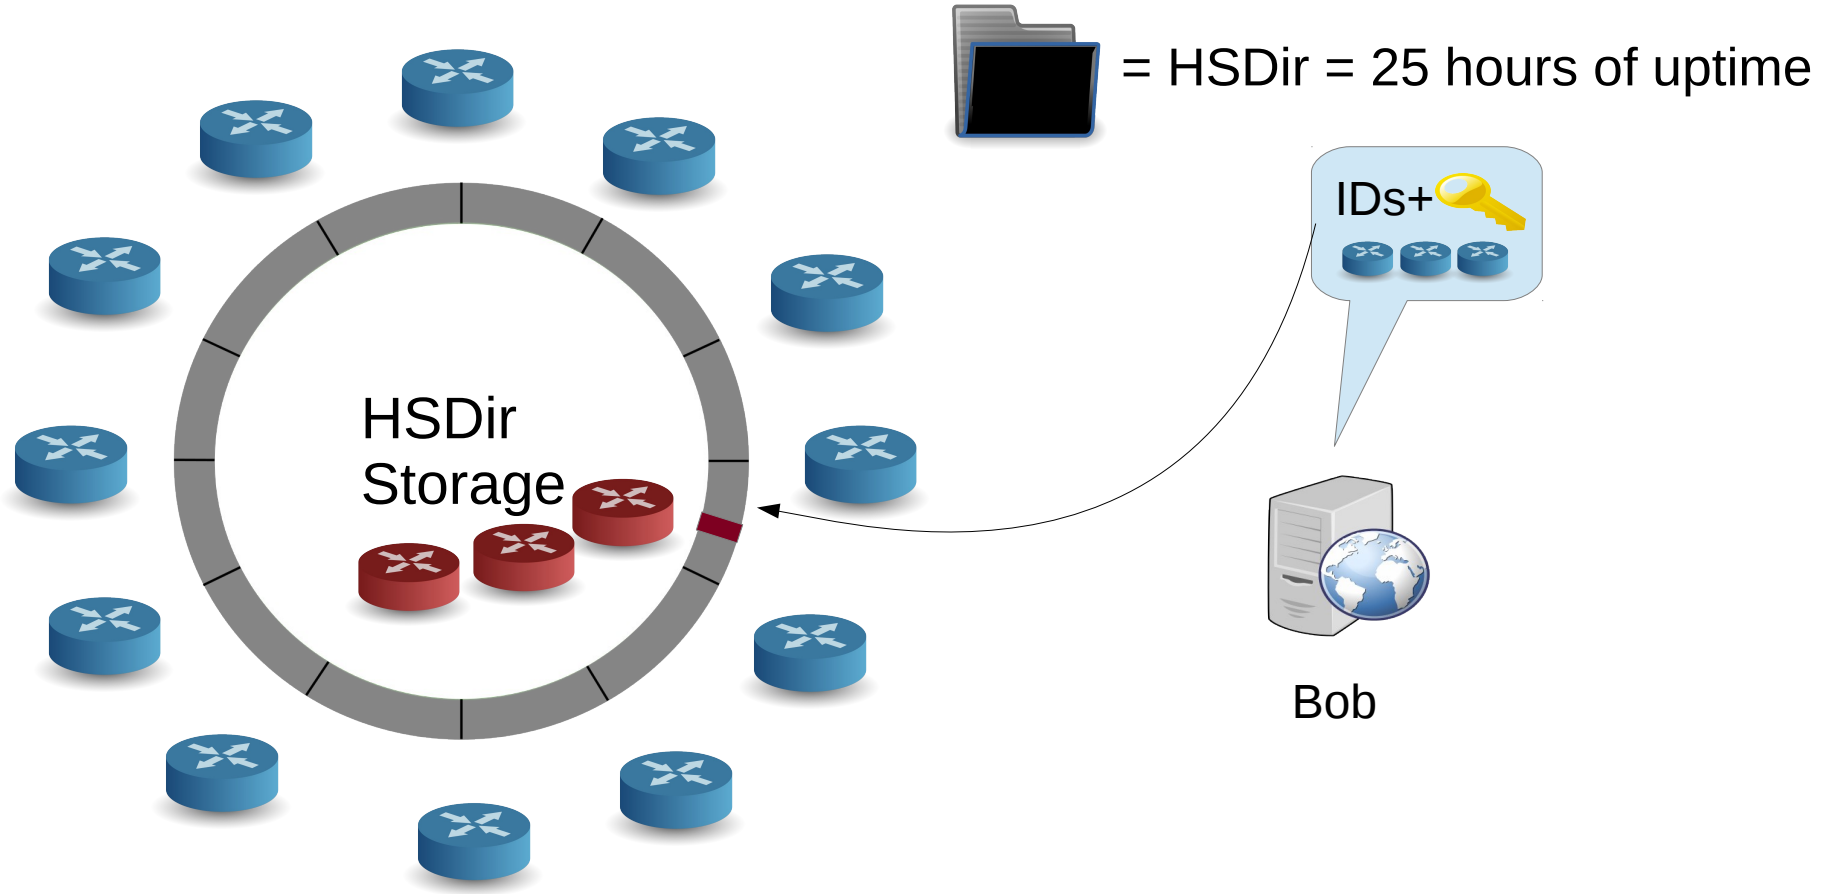

Impersonating Hidden service directory

- By impersonating 1 directory, we can track the popularity
- By impersonating all 6 directories, we can DoS.

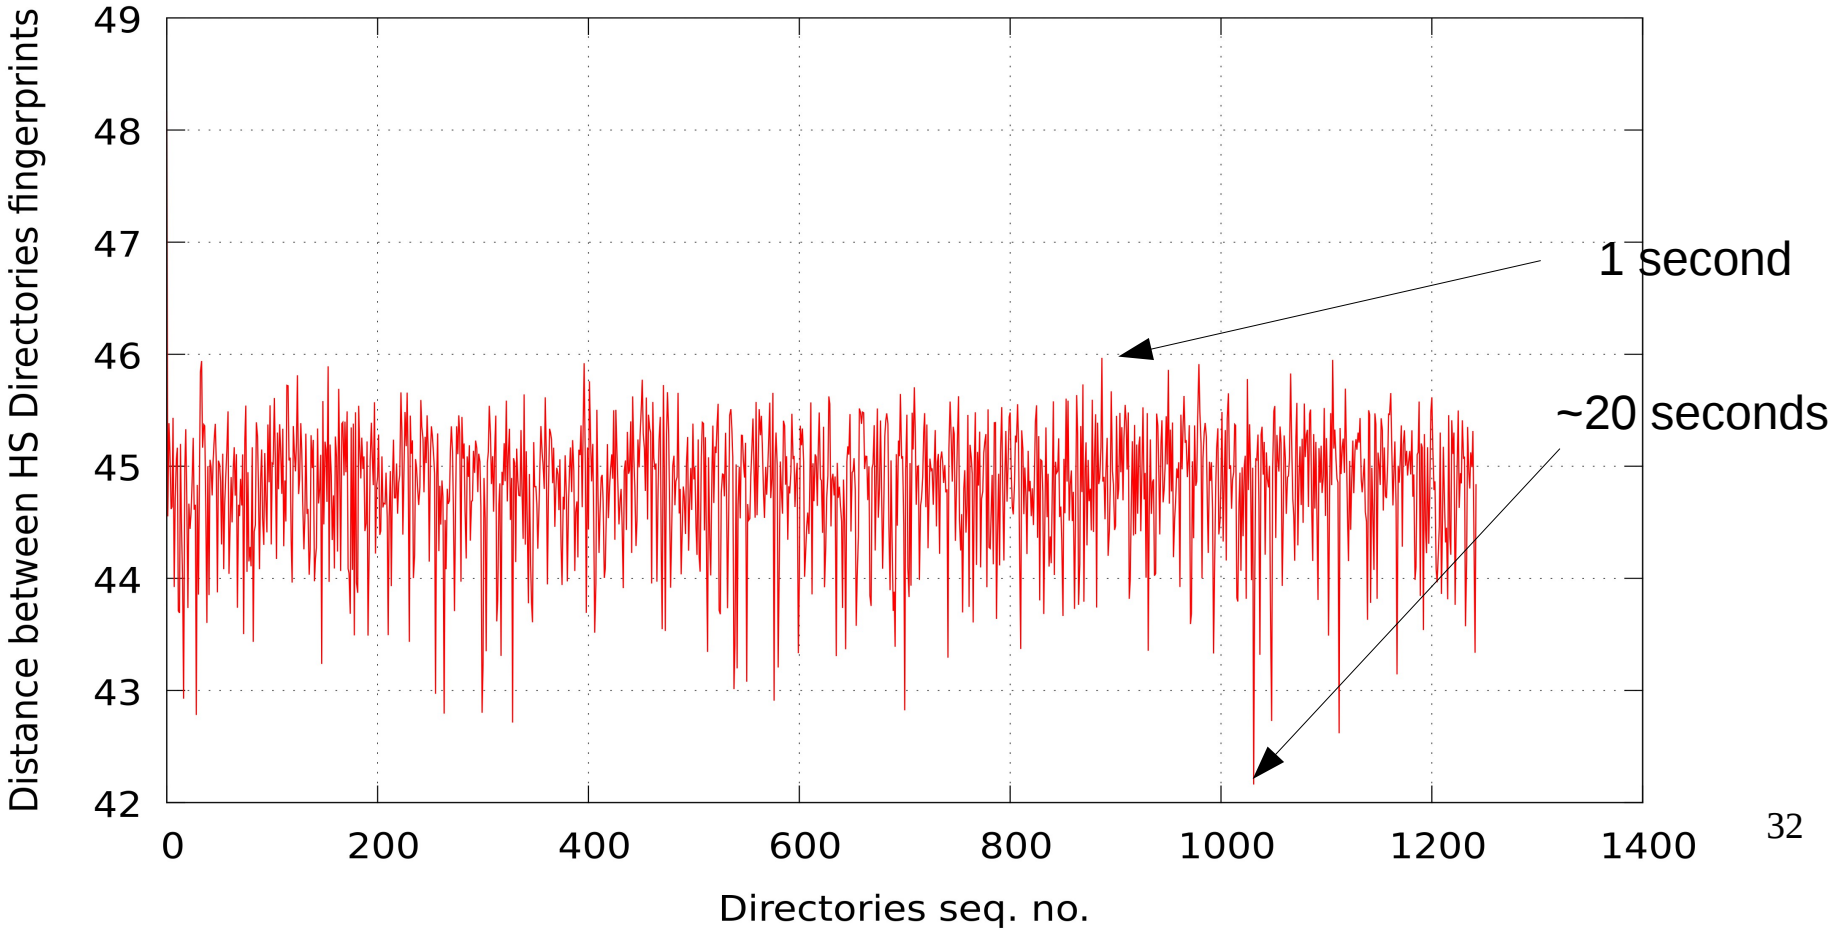

Tracking popularity

- We tracked popularity of Skynet C&C, Silkroad, and DuckDuckGo

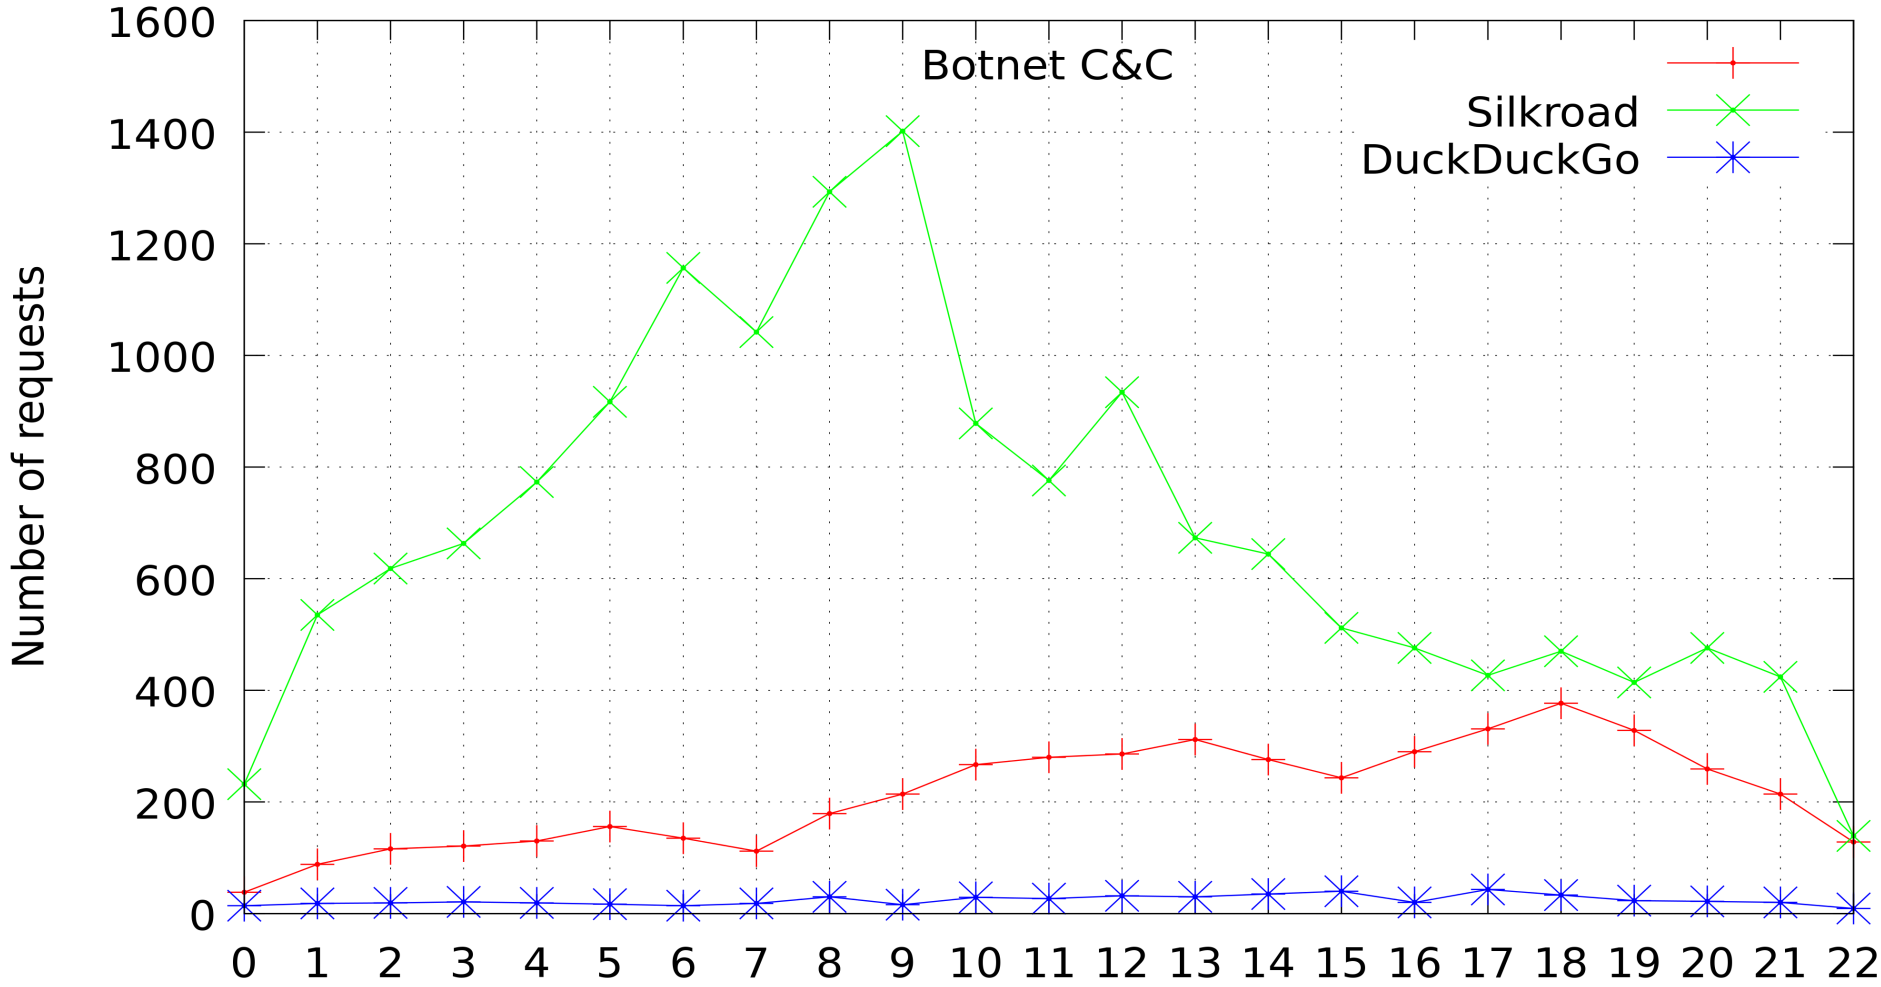

Overview

- Background
- Measuring the popularity of hidden services
- DoSing hidden services.
- **Harvesting onion addresses.**
- Revealing the guards.
- Opportunistic deanonymisation.

Shadowing

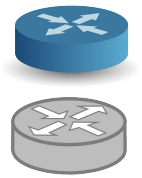

- Active

- Shadow

Consensus

menTor	1737	67 d	55863896.cust.multi.fi [85.134.56.150]
microshaft	2820	66 d	tor-exit.microshaft.org [208.201.249.3]
minisausage	3348	35 d	50.7.184.58 [50.7.184.58]
morphium	298	51 d	this.is.a.Tor.server.please.see.tor.morphium.info [91.143.90.25]
NetromAc	2115	47 d	1385160742.business.dbnet.dk [82.143.224.38]
Nitr0x	175	78 d	50.97.1.36-static.reverse.softlayer.com [50.97.1.36]
ObCanada	419	51 d	van1.zworg.com [209.17.191.117]
onconnex80	392	213 d	tor01.onconnex.com [184.105.231.11]
PasToutAFaitNet1	261	176 d	91.229.20.159 [91.229.20.159]
PasToutAFaitNet2	763	196 d	tor2.pastoutafait.net [95.130.11.247]
plebia	3599	79 d	tor-exit.plebia.org [37.59.162.218]
pps	9	59 d	184-22-164-107.static.hostnoc.net [184.22.164.107]
PrivaTOReu	4229	100 d	torexit.privator.eu [88.208.90.1]
programmercpp	149	36 d	proxy [213.171.220.40]
PsyNetNP	155	52 d	broadband-95-84-148-164.nationalcablenetworks.ru [95.84.148.164]
Qwerty	91	157 d	93.167.245.178 [93.167.245.178]

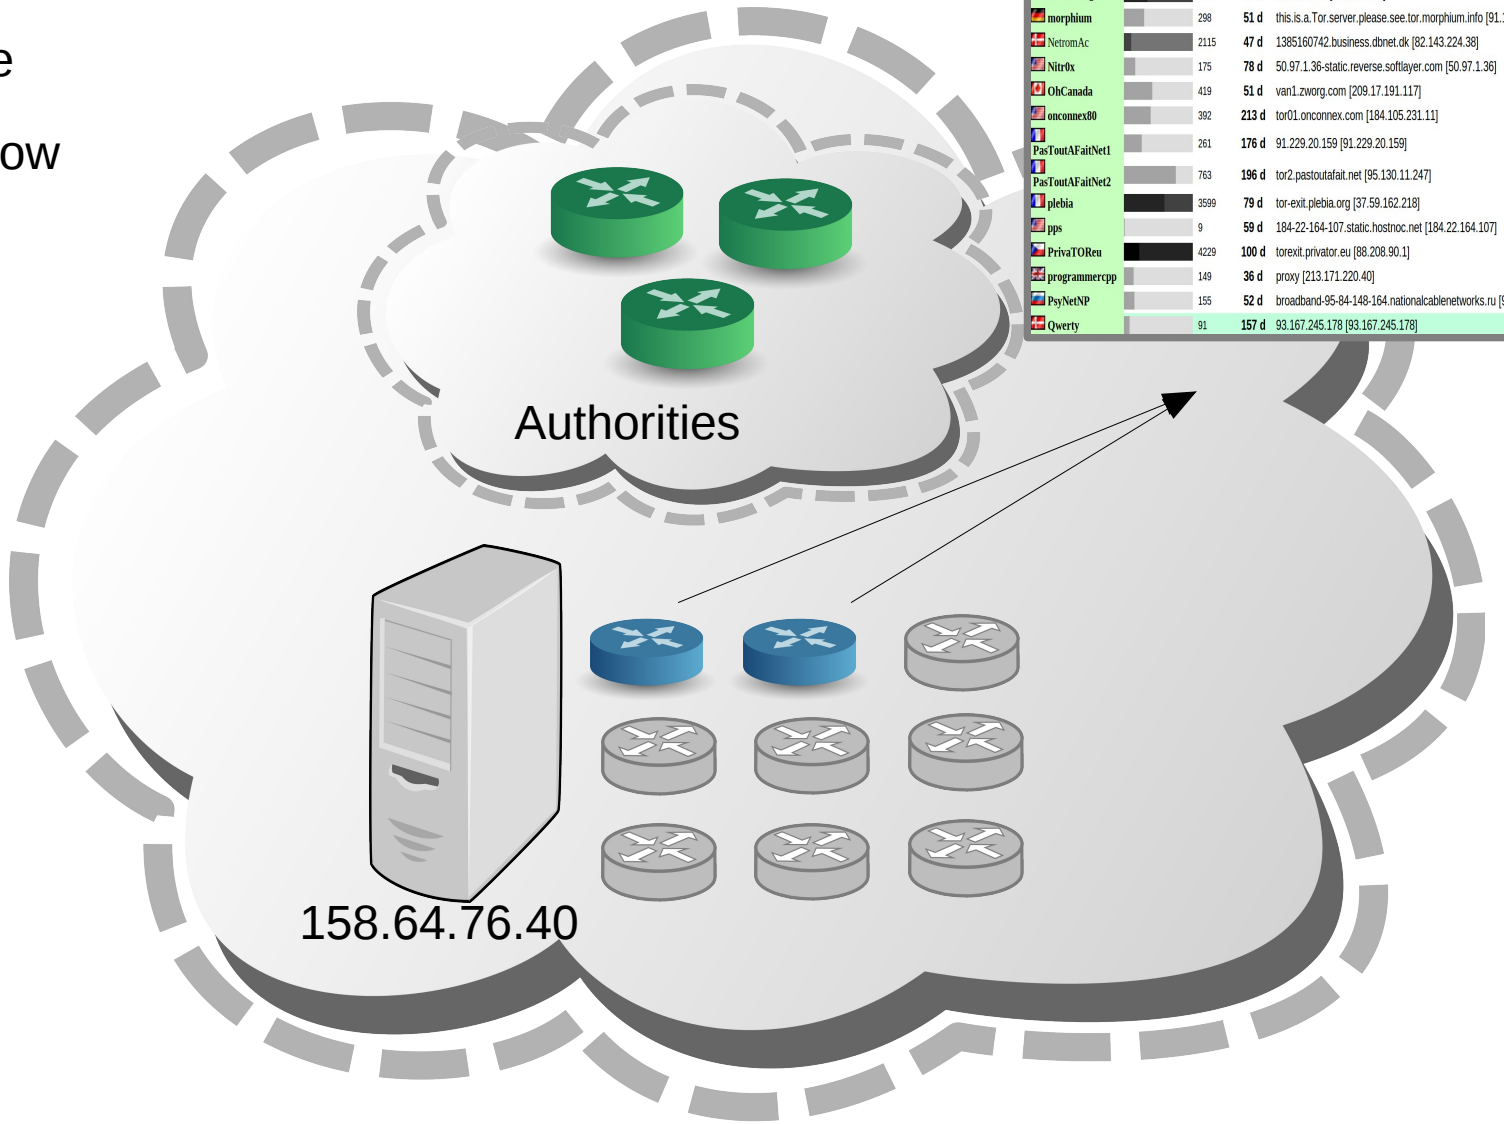

Shadowing

- Active
- Shadow

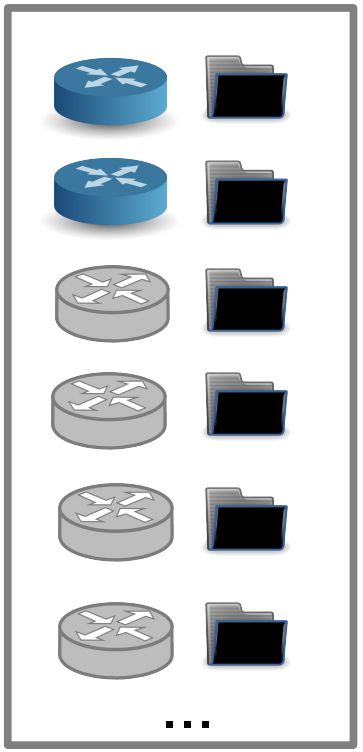

Authorities
Internal database

Consensus

menTor	1737	67 d	55863896.cust.multi.fi [85.134.56.150]	
microshaft	2820	66 d	tor-exit.microshaft.org [208.201.249.3]	
minisausage	3348	35 d	50.7.184.58 [50.7.184.58]	
morphium	298	51 d	this.is.a.Tor.server.please.see.tor.morphium.info [91.143.90.25]	
NetromAc	2115	47 d	1385160742.business.dbnet.dk [82.143.224.38]	
Nitr0x	175	78 d	50.97.1.36-static.reverse.softlayer.com [50.97.1.36]	
ObCanada	419	51 d	van1.zworg.com [209.17.191.117]	
onconnex80	392	213 d	tor01.onconnex.com [184.105.231.11]	
PasToutAFaitNet1	261	176 d	91.229.20.159 [91.229.20.159]	
PasToutAFaitNet2	763	196 d	tor2.pastoutafait.net [95.130.11.247]	
plebia	3599	79 d	tor-exit.plebia.org [37.59.162.218]	
pps	9	59 d	184-22-164-107.static.hostnoc.net [184.22.164.107]	
PrivaTOReu	4229	100 d	torexit.privator.eu [88.208.90.1]	
programmercpp	149	36 d	proxy [213.171.220.40]	
PsyNetNP	155	52 d	broadband-95-84-148-164.nationalcablenetworks.ru [95.84.148.164]	
Qwerty	91	157 d	93.167.245.178 [93.167.245.178]	

Authorities

158.64.76.40

Shadowing

- Active

- Shadow

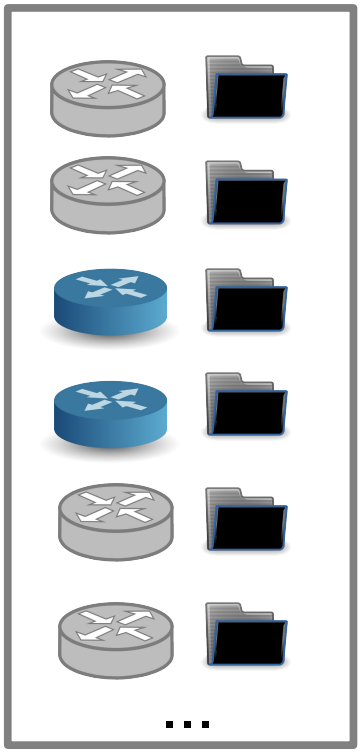

Authorities
Internal database

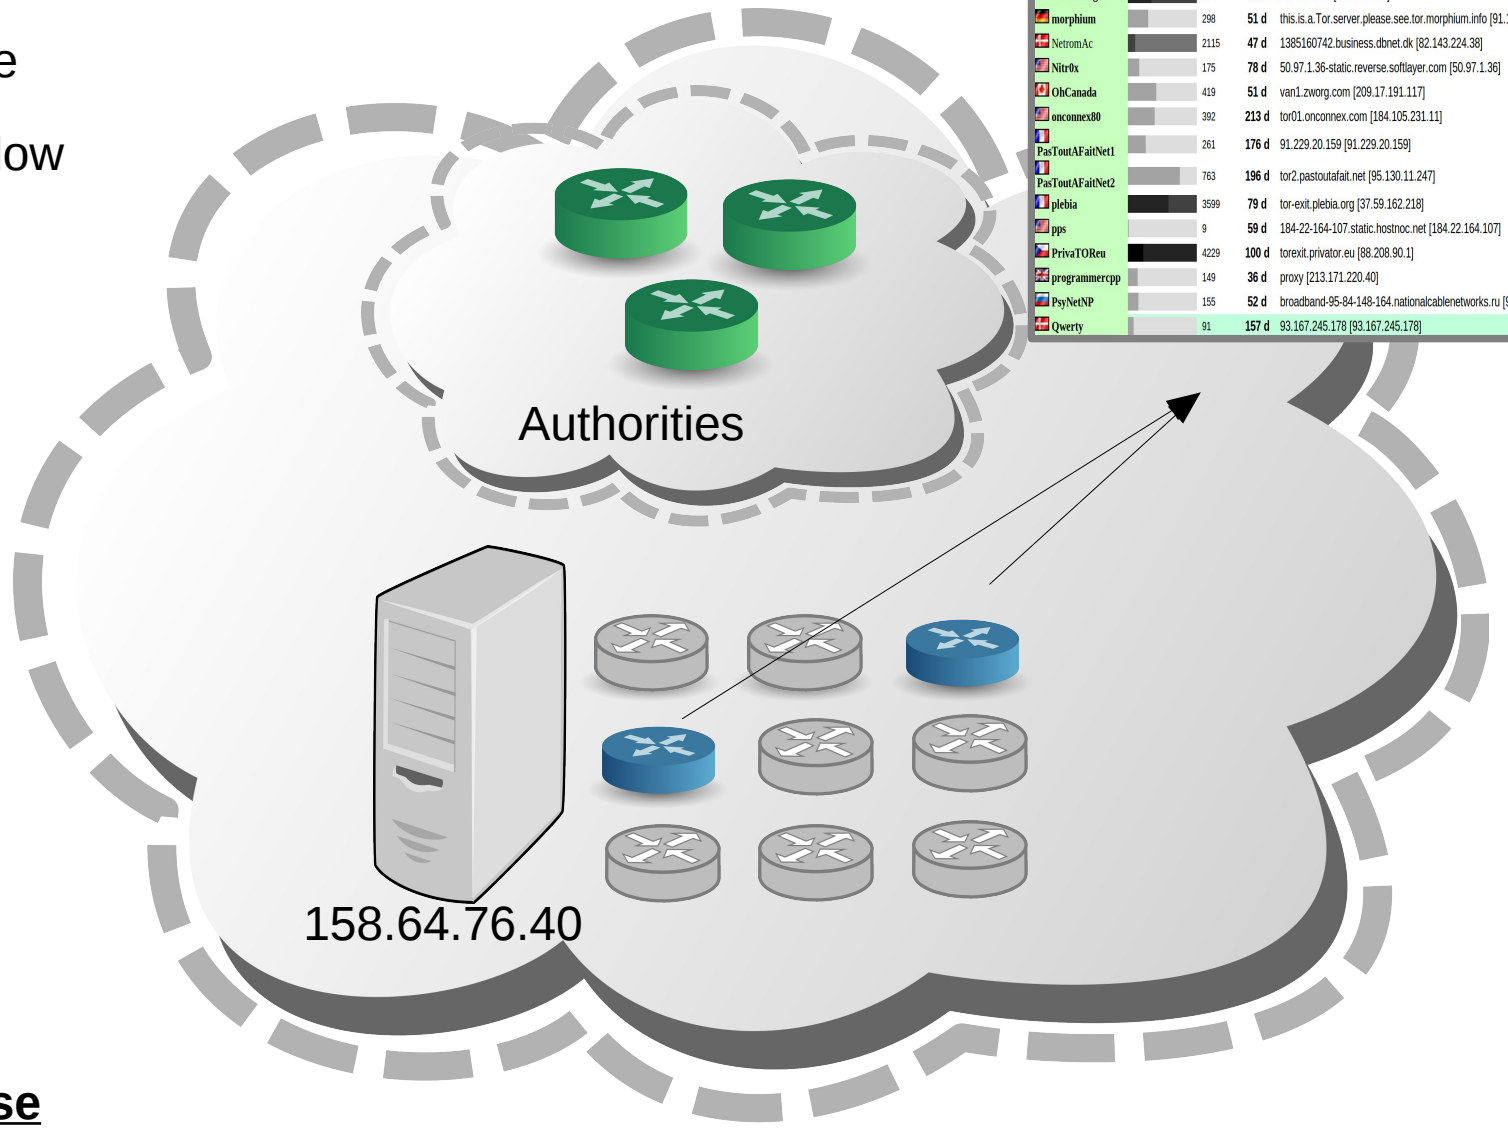

Consensus

Shadowing

- Active

- Shadow

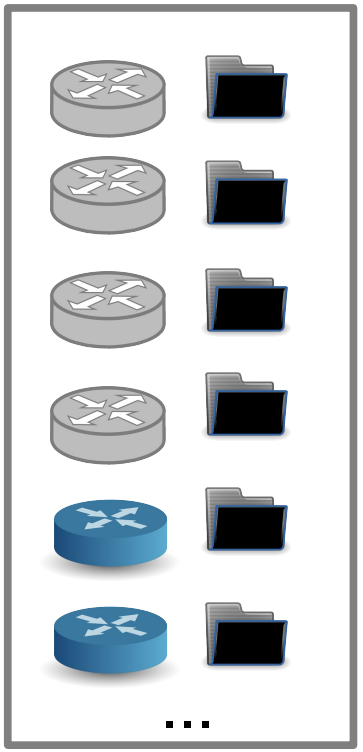

Authorities
Internal database

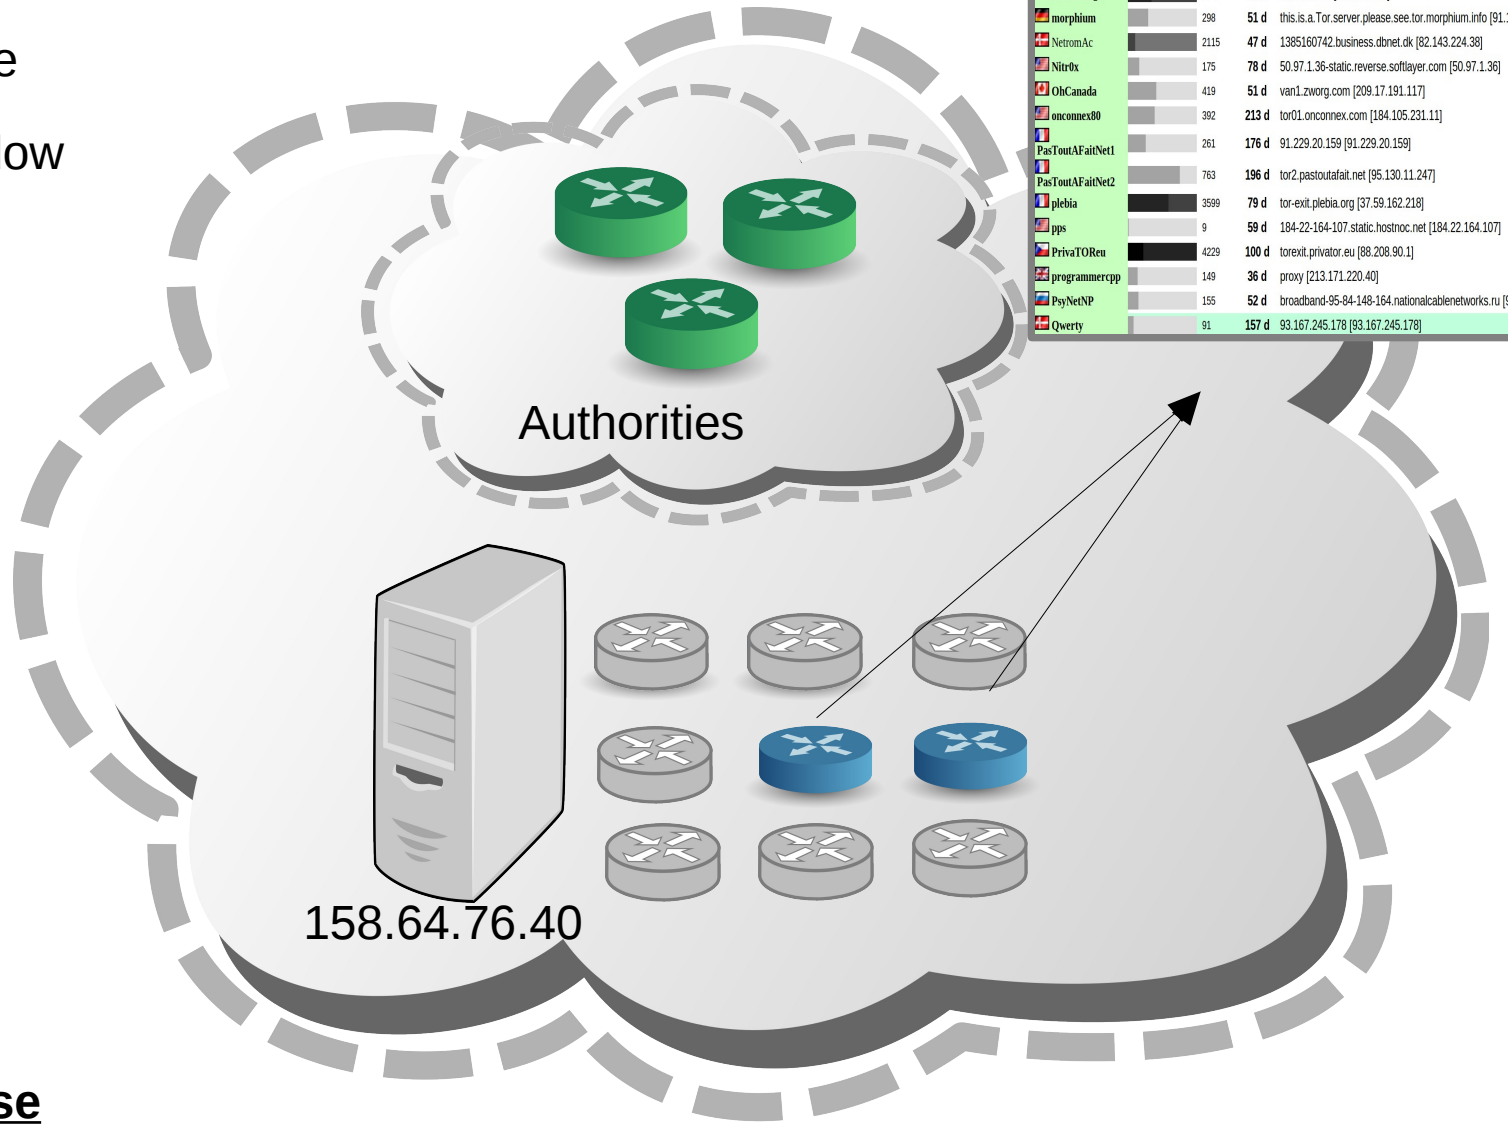

Consensus

Collecting onion addresses

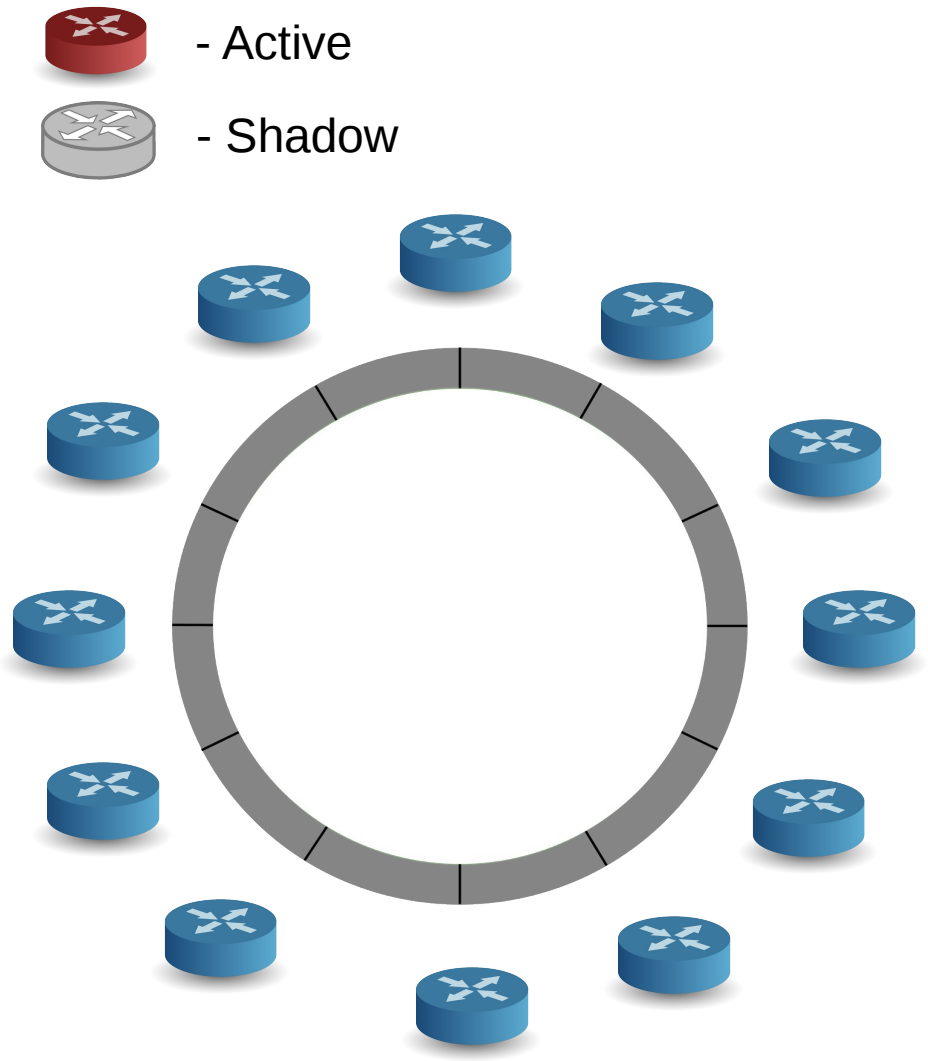

Collecting onion addresses

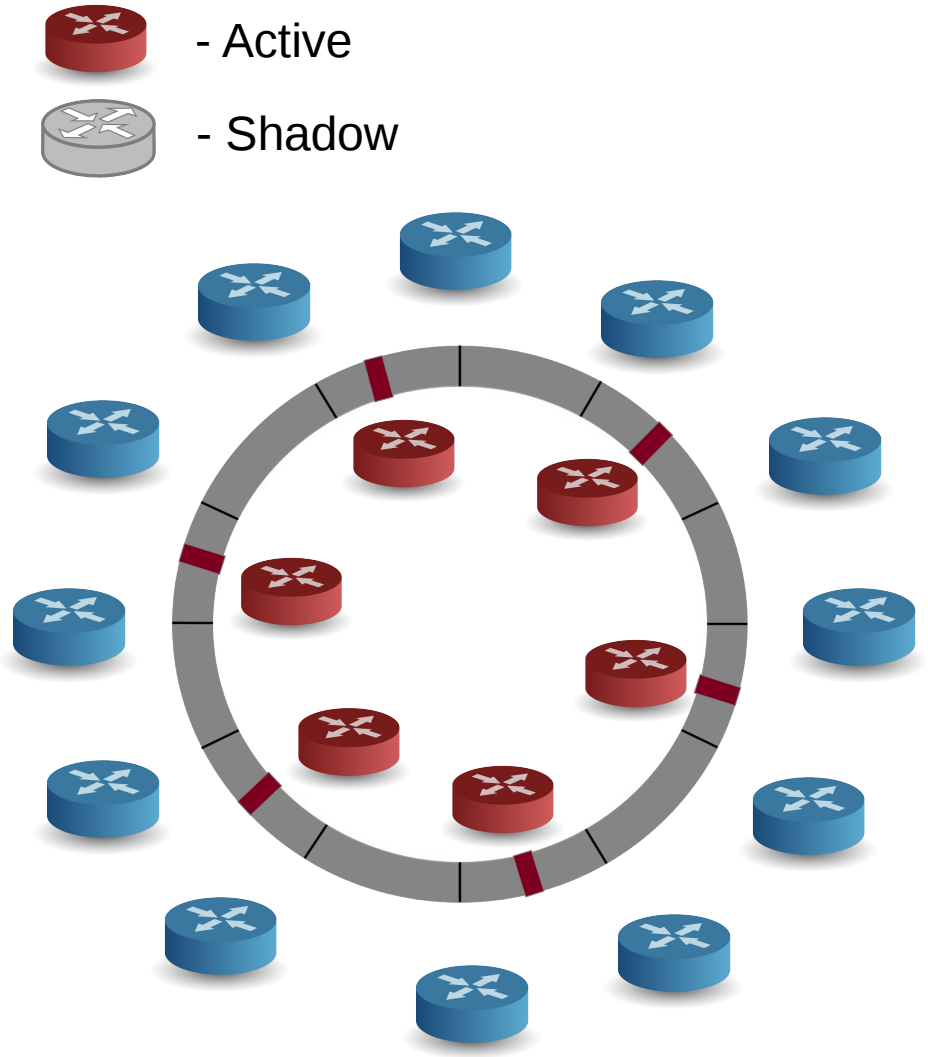

- Naive approach will require ~350 IP addresses.

Collecting onion addresses

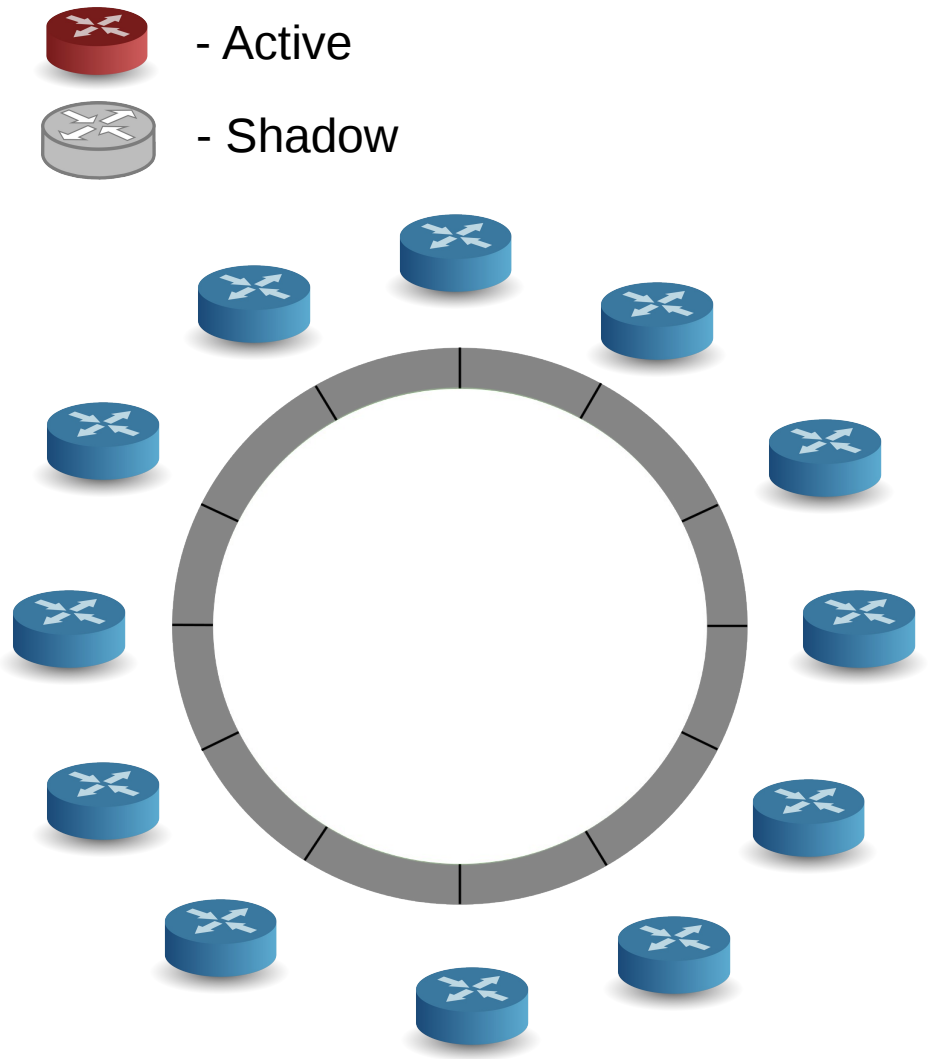

- Naive approach will require ~350 IP addresses.
- Descriptors don't relocate within 24 hours.
- Prepare shadow HSDir relays and gradually pull to consensus.

Collecting onion addresses

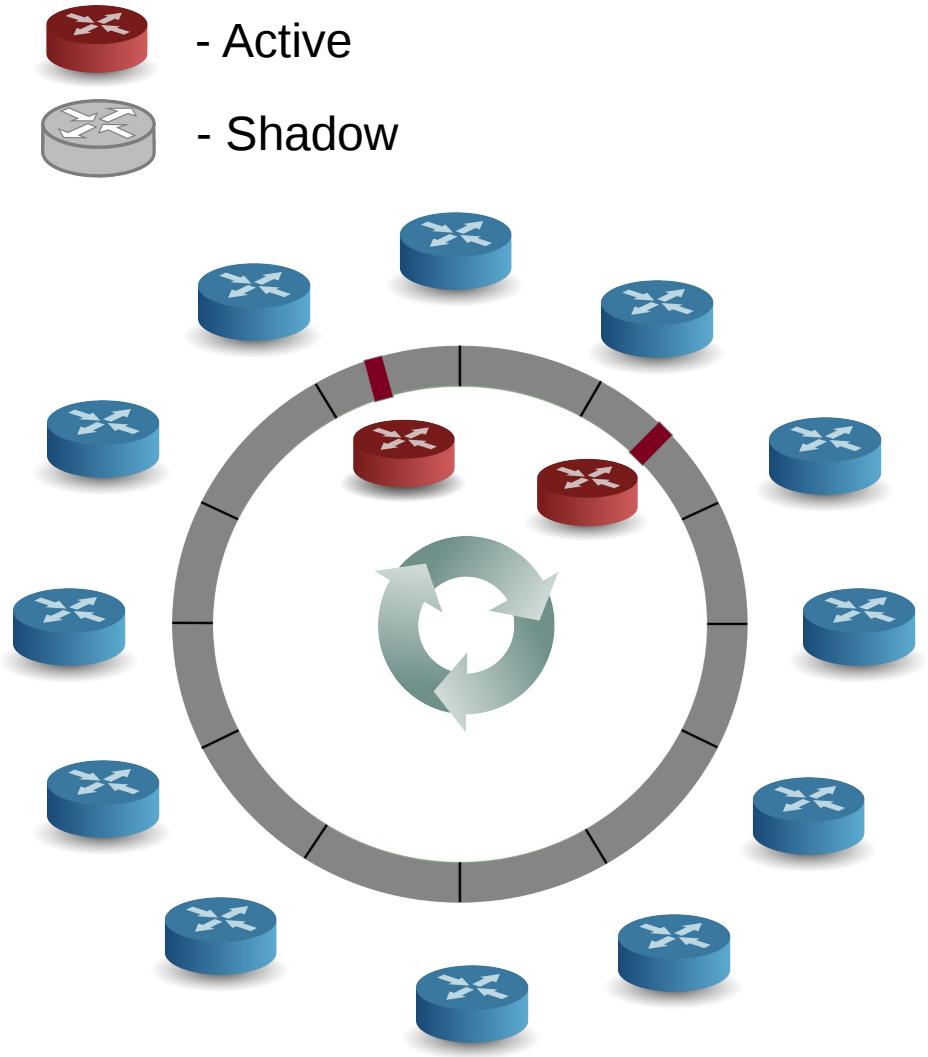

- Naive approach will require ~350 IP addresses.
- Descriptors don't relocate within 24 hours.
- Prepare shadow HSDir relays and gradually pull to consensus.

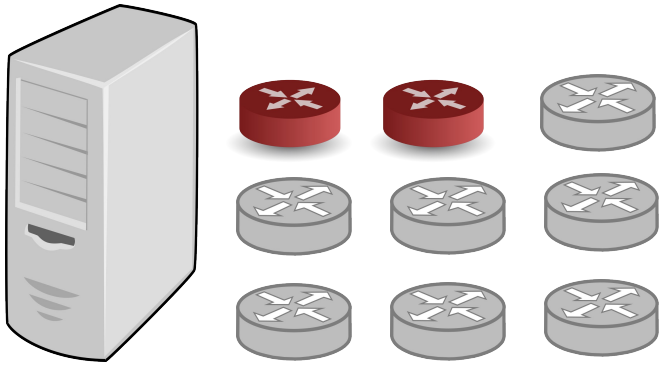

158.64.76.40

Collecting onion addresses

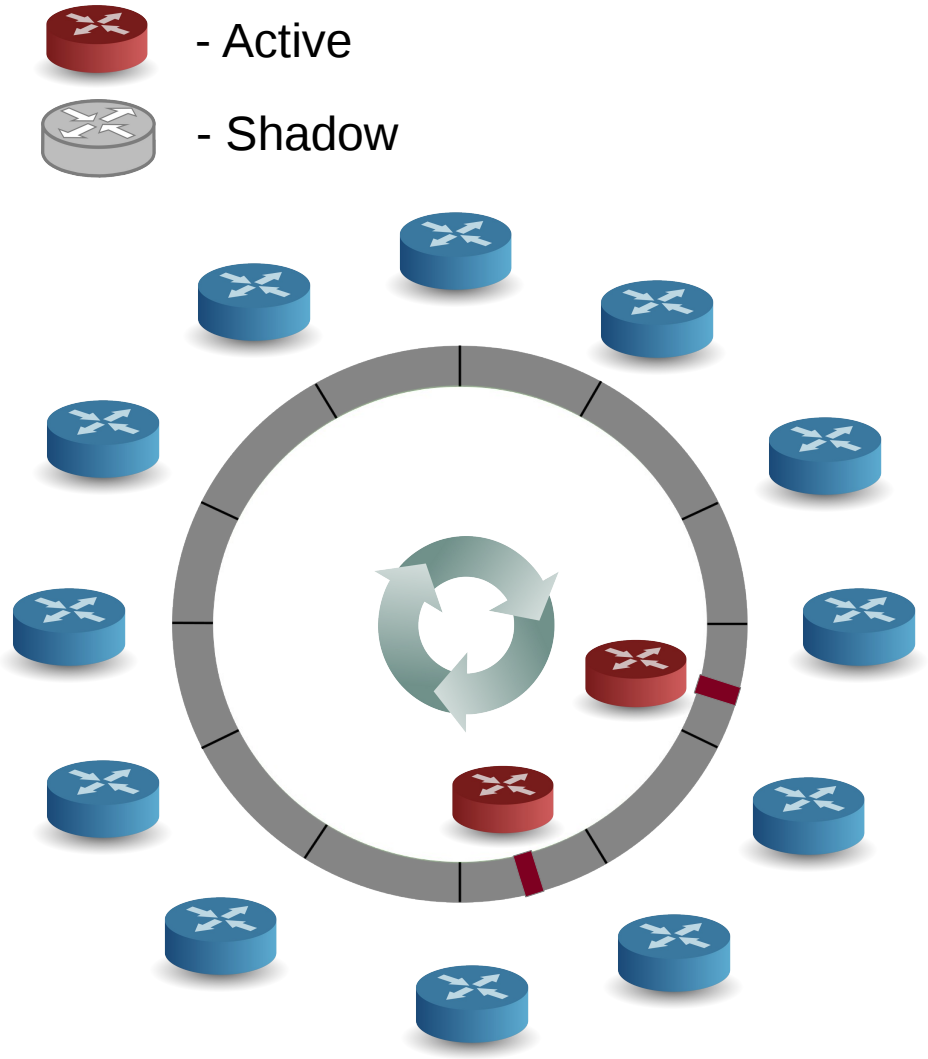

- Naive approach will require ~350 IP addresses.
- Descriptors don't relocate within 24 hours.
- Prepare shadow HSDir relays and gradually pull to consensus.

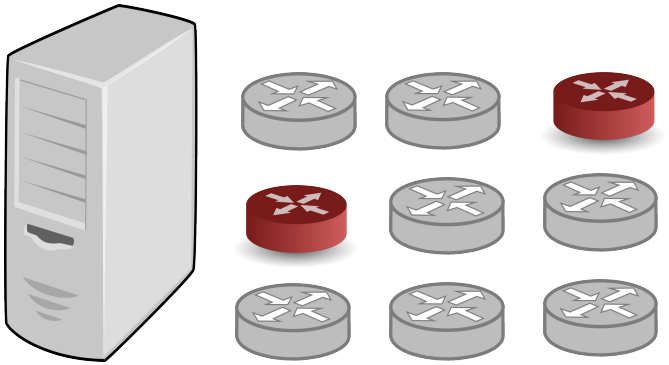

158.64.76.40

Collecting onion addresses

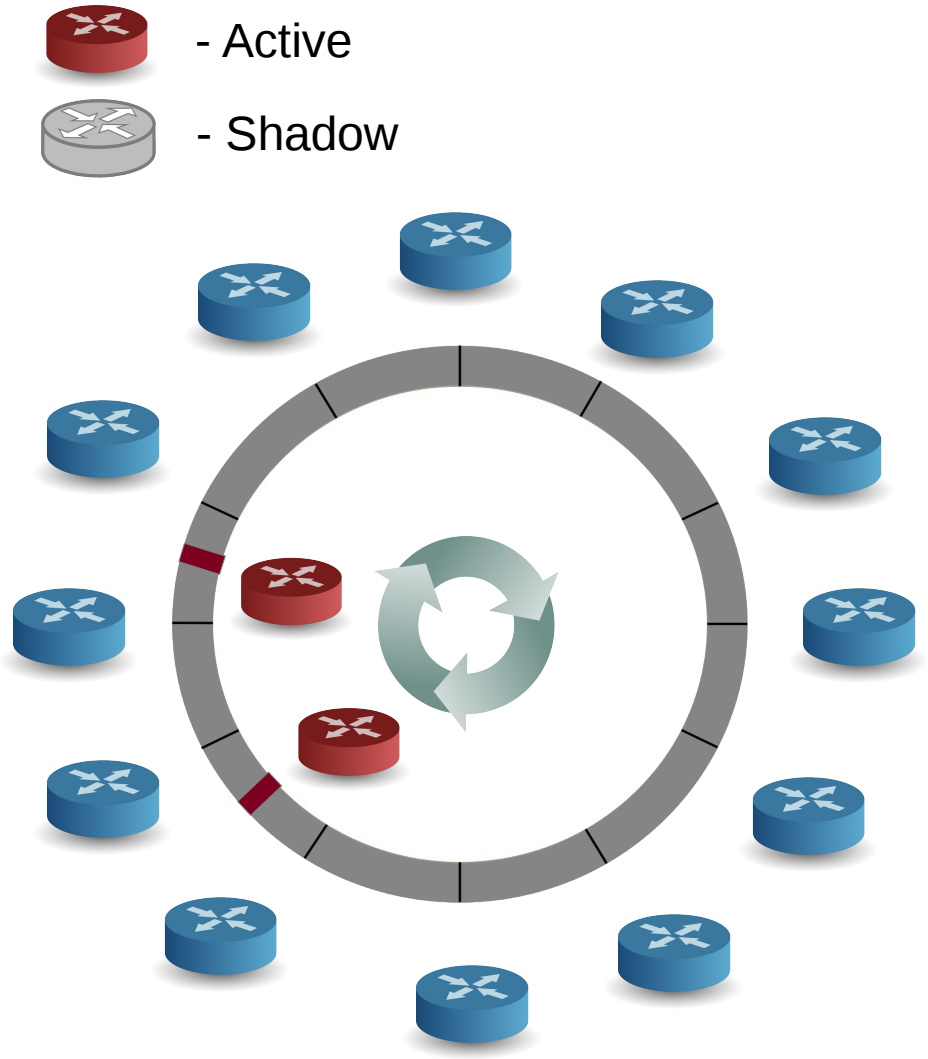

- Naive approach will require ~350 IP addresses.
- Descriptors don't relocate within 24 hours.
- Prepare shadow HSDir relays and gradually pull to consensus.

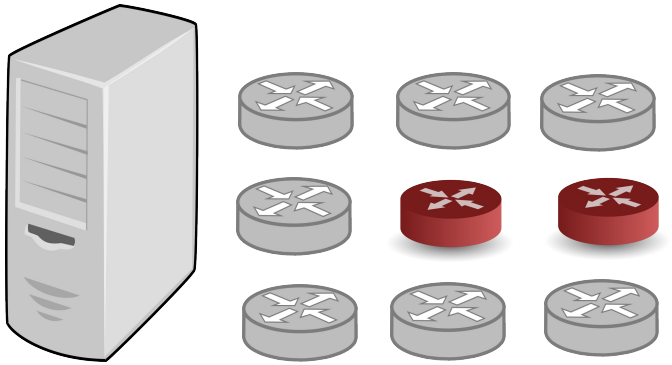

158.64.76.40

Harvest results

- We used 58 IP addresses from Amazon EC2 and spent 57 USD
- We collected 39824 unique onion addresses in 49 hours (on hidden wikis one can find ~2500 addresses only)
- Some interesting note: 12 onion addresses in the form silkroad*****.onion.

Overview

- Background
- Measuring the popularity of hidden services
- DoSing hidden services.
- Harvesting onion addresses.
- **Revealing the guards.**
- **Opportunistic deanonymisation.**

Revealing Guard Nodes

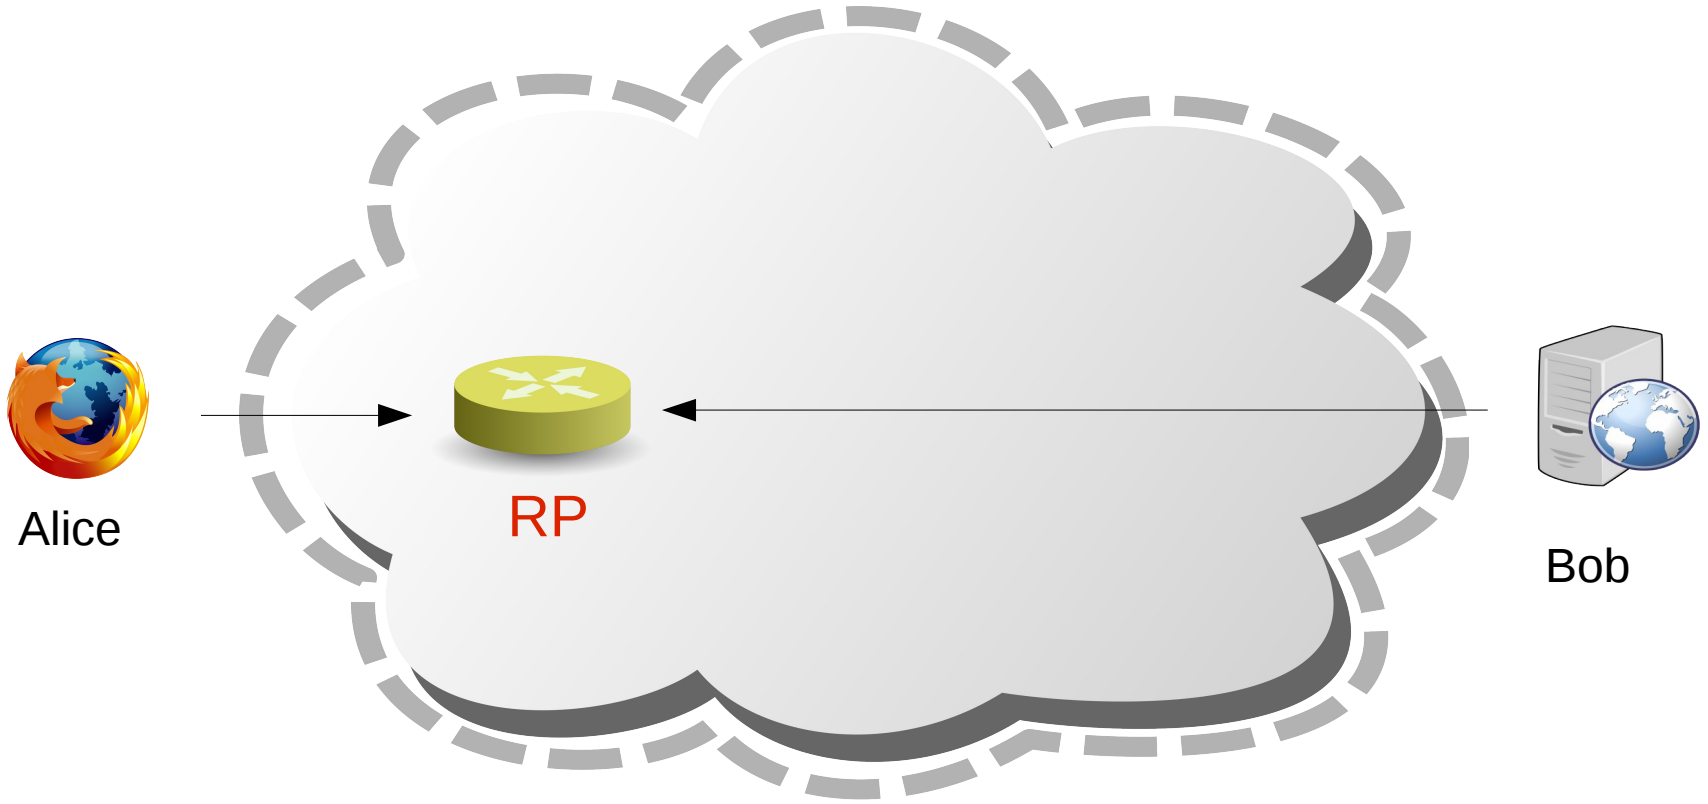

Revealing Guard Nodes

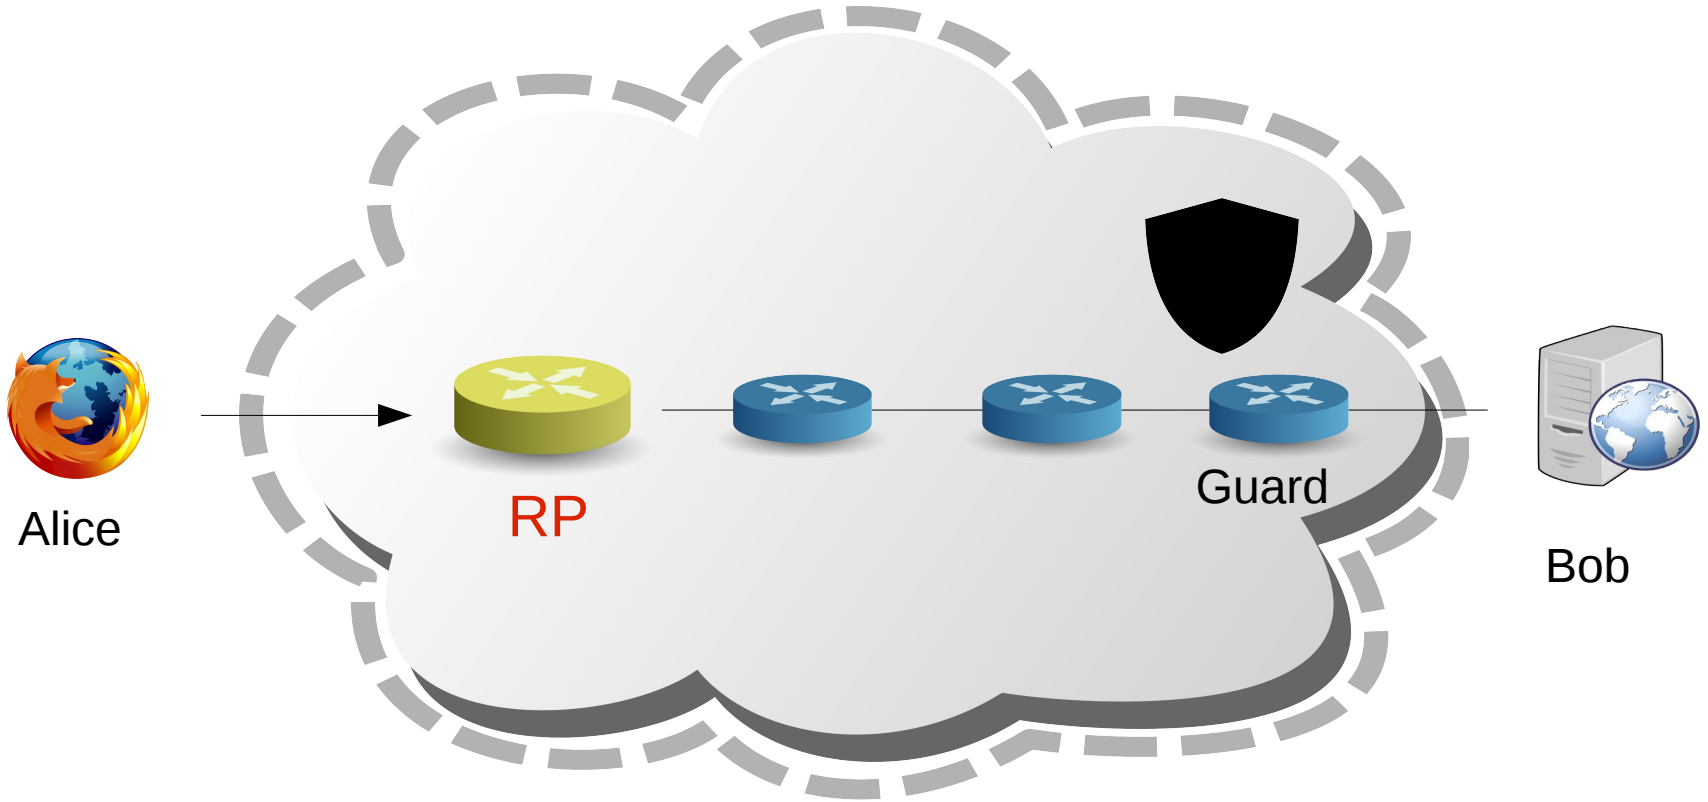

Revealing Guard Nodes

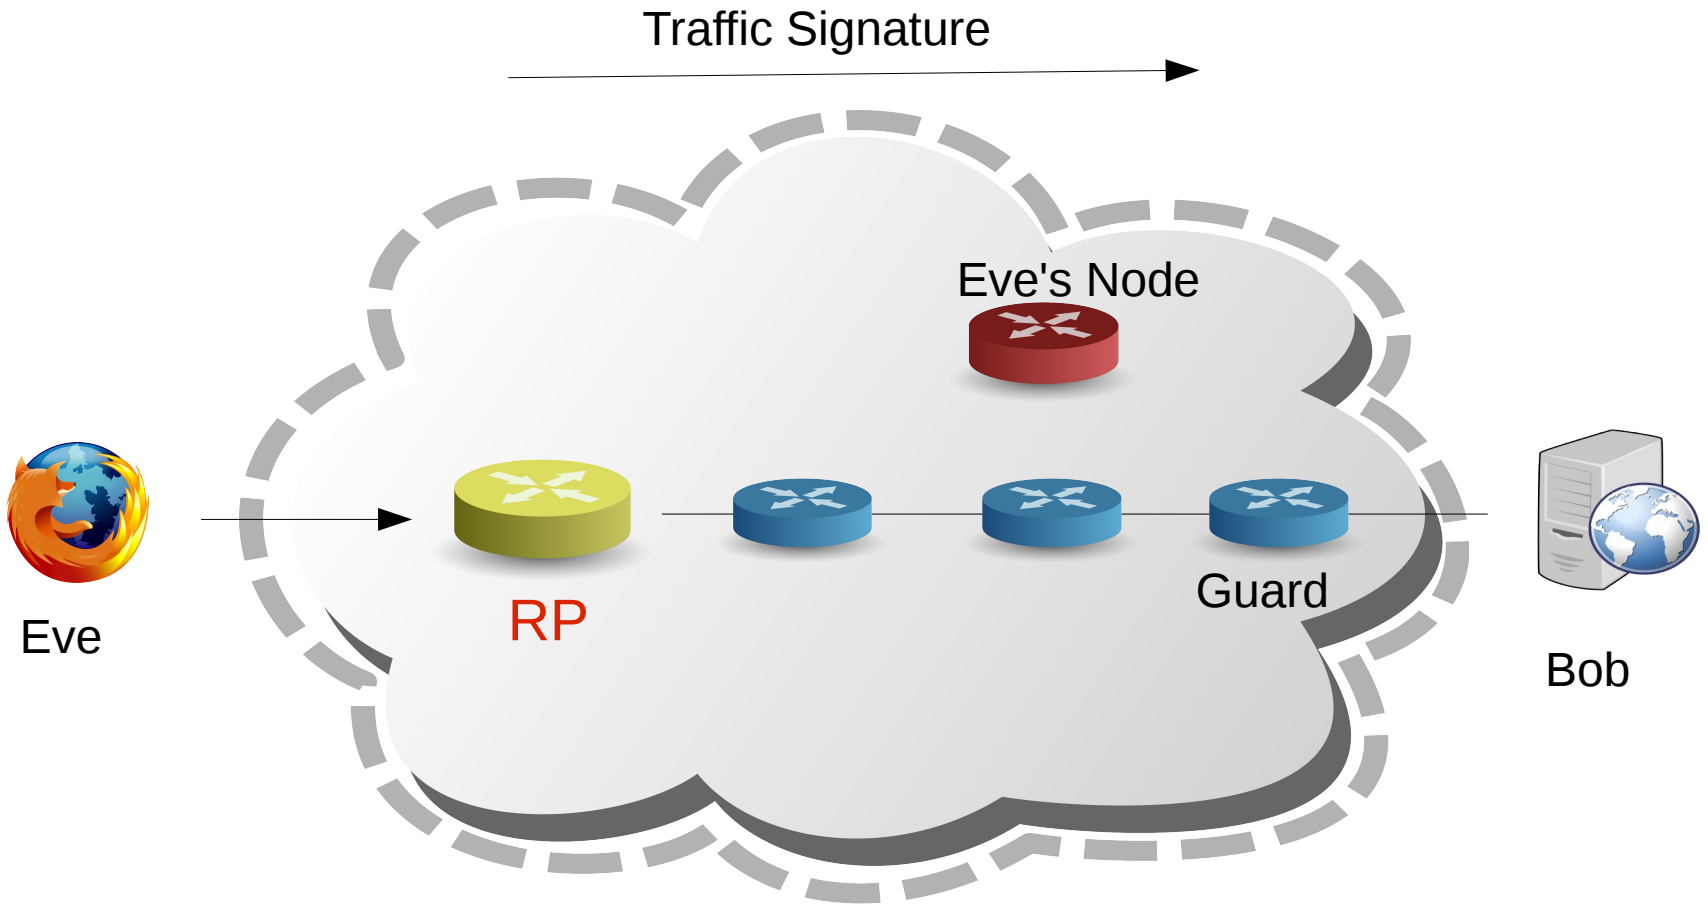

Revealing Guard Nodes

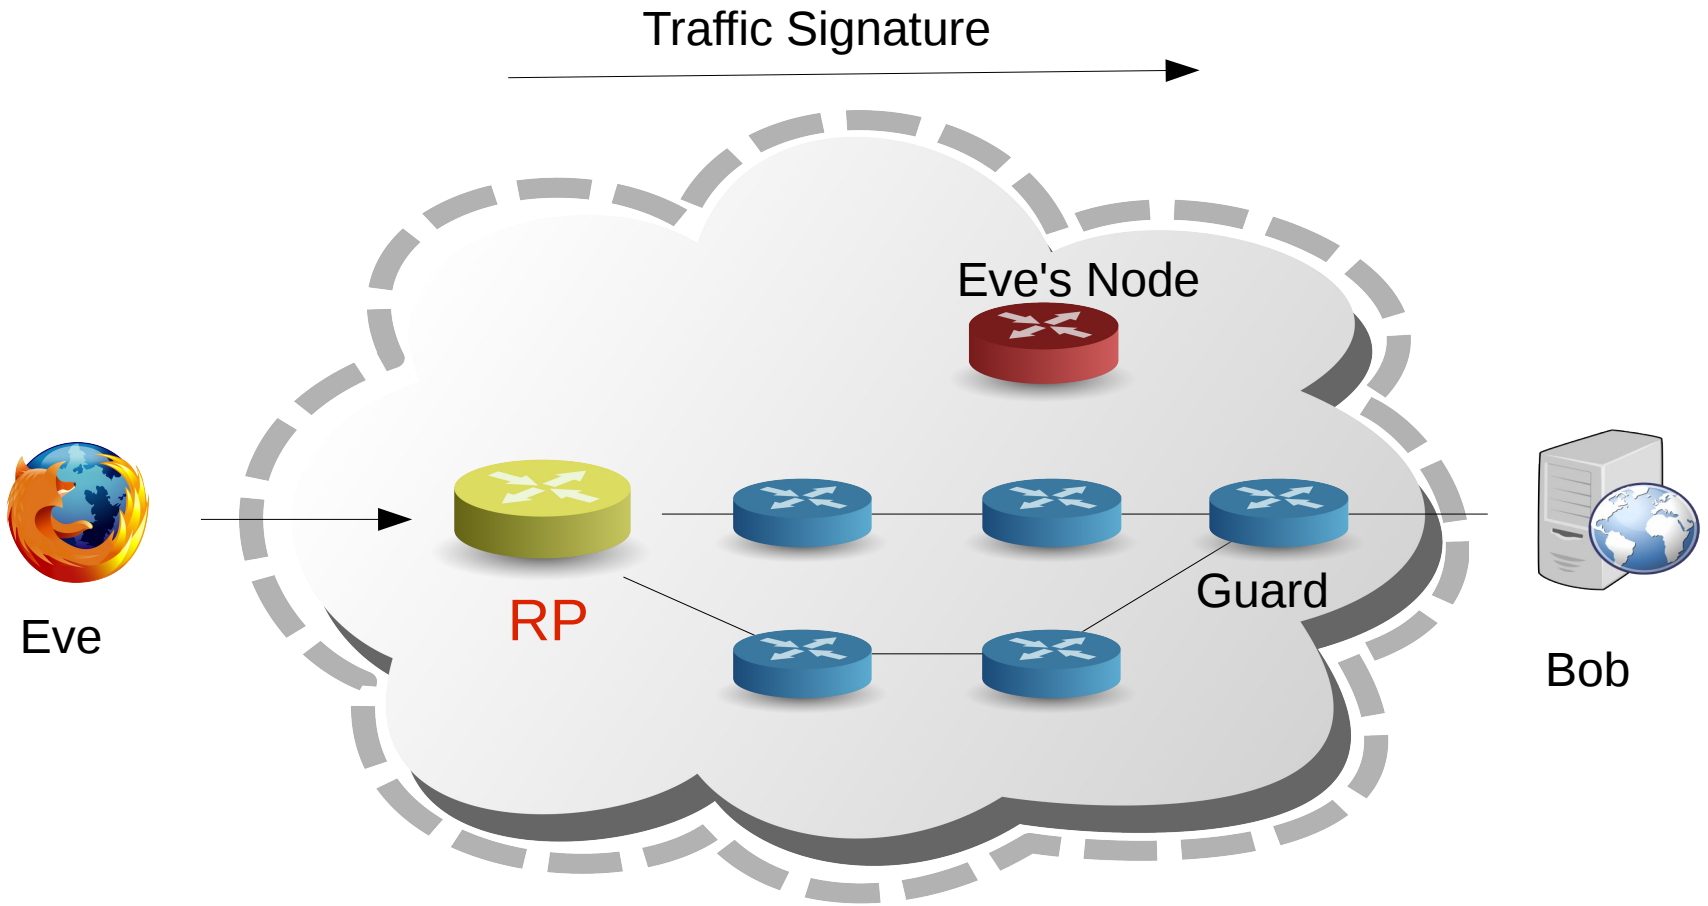

Revealing Guard Nodes

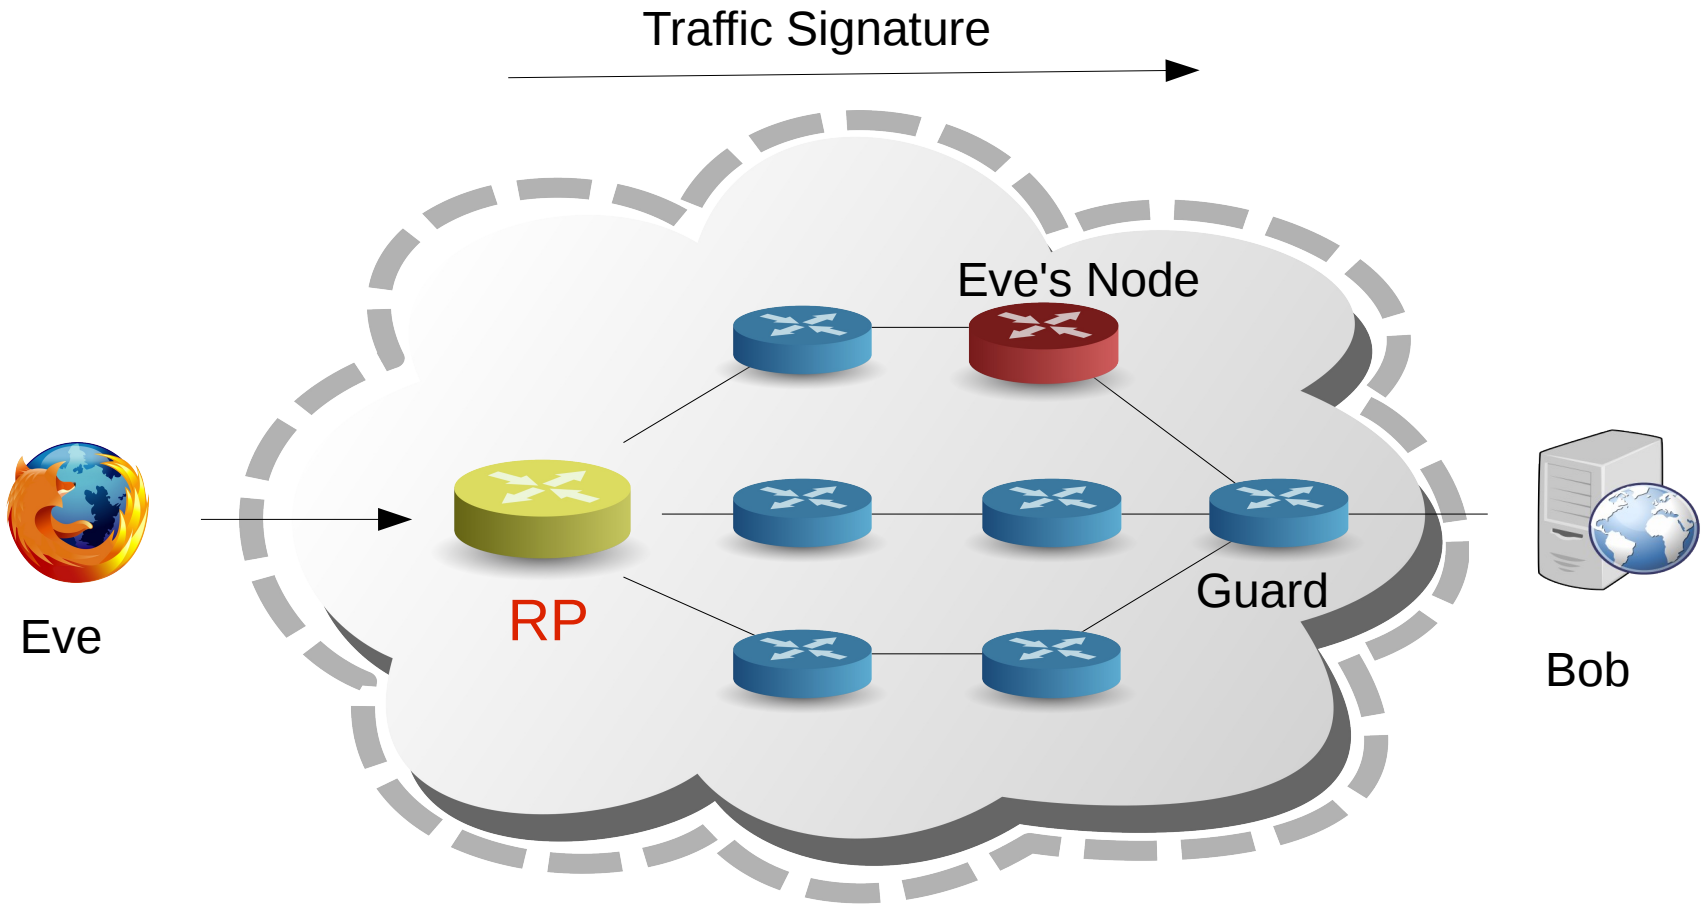

~40 minutes to reveal the guard nodes for a 5Mb/s node

Opportunistic deanonymisation

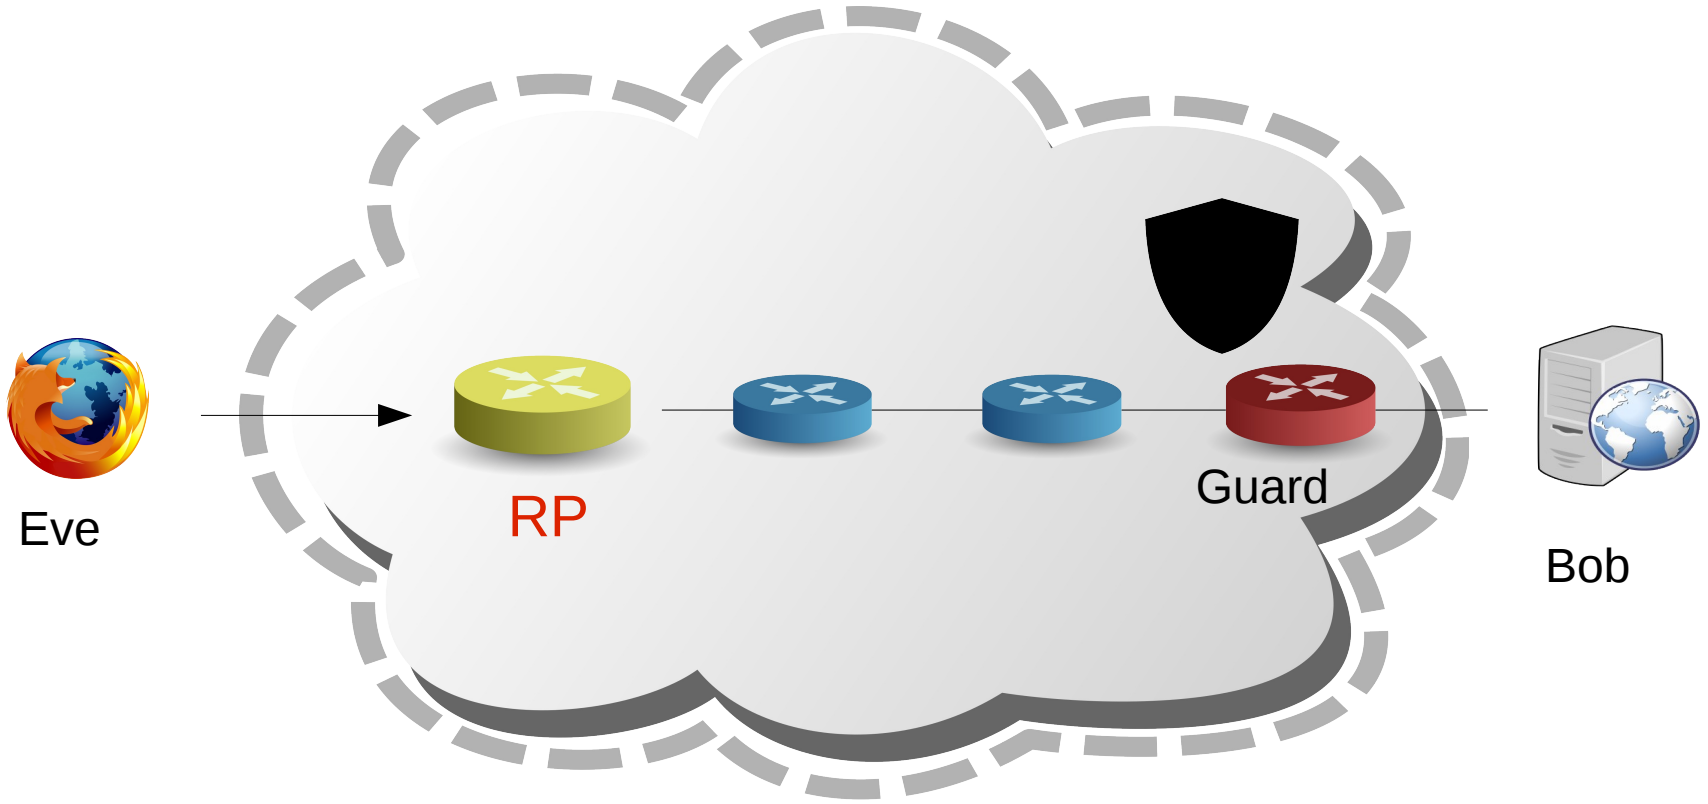

Opportunistic deanonymisation

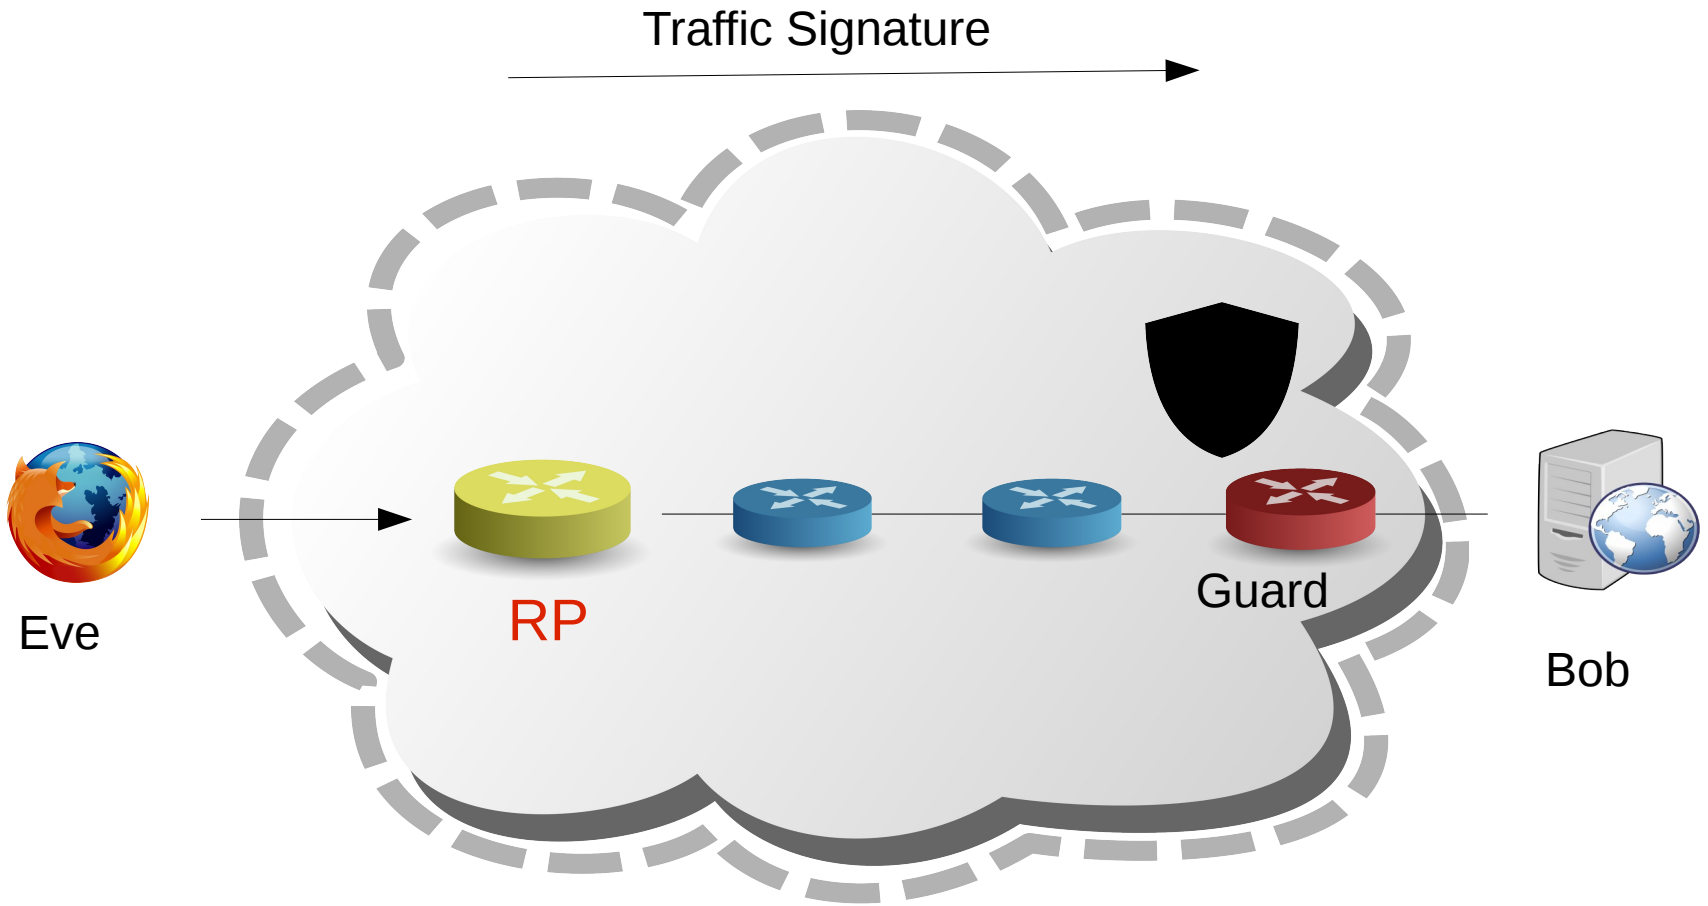

How long does it take to become a Guard of a hidden service?

Opportunistic deanonymisation

- Rent a server for 60 USD per month => 0.6% probability to be chosen as a Guard.
- Deanonymisation **~150 hidden services per month** (for **60 USD** per month)
- By running 23 such servers, the probability to deanonymize **any** long-running **hidden service** within **8 months is 99%**. (~11 000 USD total).

Side effect (flag assignment)

- Large number of shadow relays with bw <= 1 accelerated flag assignment.

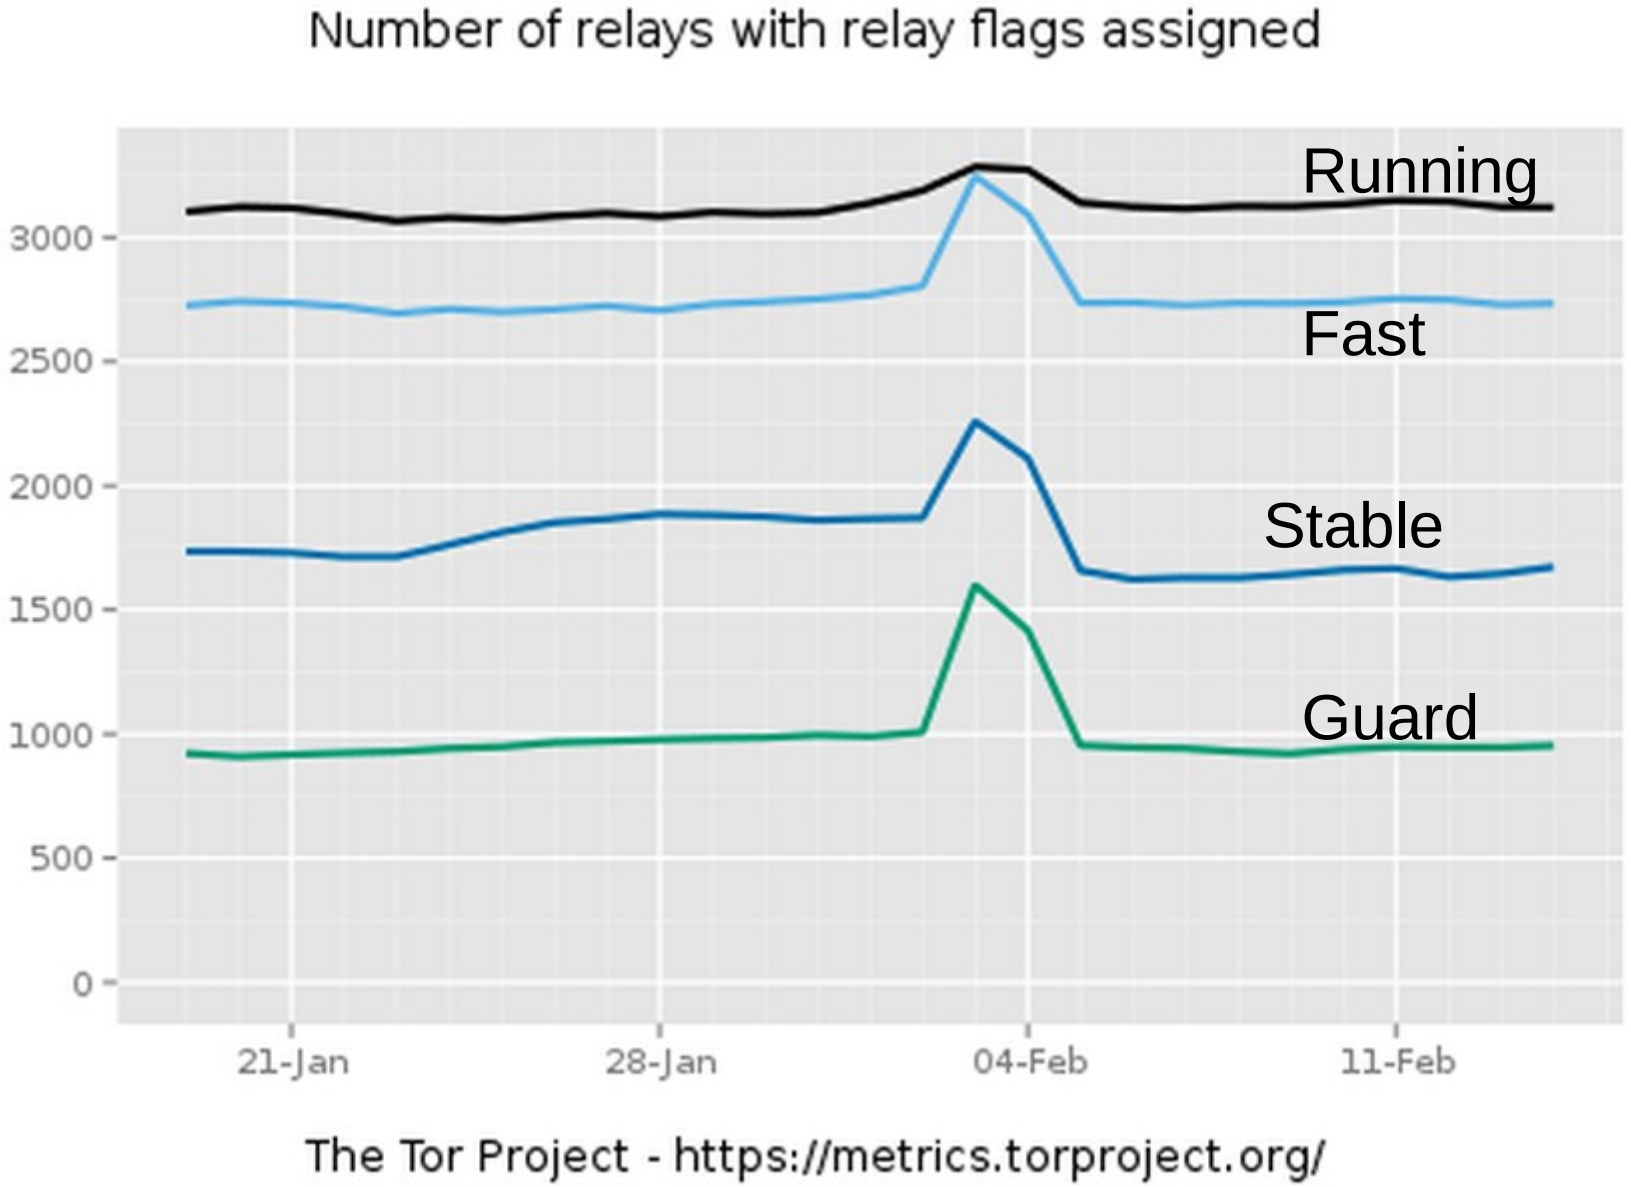

Conclusions

Tracking	
Denial of Service	
Collecting onion addresses	
Revealing Guard Nodes	
Deanonymisation	• 150 addresses per month (60 USD) •Any HS (8 months+11000 USD) 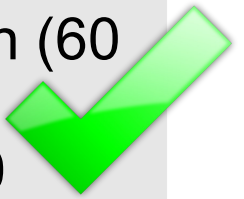

Support slide 1

- Triggered
  - #8243: Getting the HSDir flag should require more effort
  - #8243: Getting the HSDir flag should require more effort
- Related
  - Changing of the Guards: A Framework for Understanding and Improving Entry Guard Selection in Tor", WPES 2012
  - #8240: Raise our guard rotation period

    (patch to raise it to 9.5 month still pending)

Support slide 2

- Not included into the presentation
  - Finding guard nodes using topological properties
  - Bandwidth inflation