Basic Course on Onion Routing

Paul Syverson

U.S. Naval Research Laboratory

paul.syverson@nrl.navy.mil

SAC Summer School

Mount Allison University Aug 10, 2015

Course Outline

- Lecture 1: Basics and Formalization
  - Usage examples, basic notions of traffic-secure communications, mixes and onion routers
  - Onion routing design basics: circuit construction protocols, network discovery
  - Formalization and analysis, possibilistic and probabilistic definitions of anonymity
- Lecture 2: Security for the real world
  - Simple demo of obtaining/using Tor
  - Security of obtain/using Tor
  - Adding network link awareness
  - Importance of modeling users
  - Importance of realistic and practical
    - Adversary models
    - Security definitions

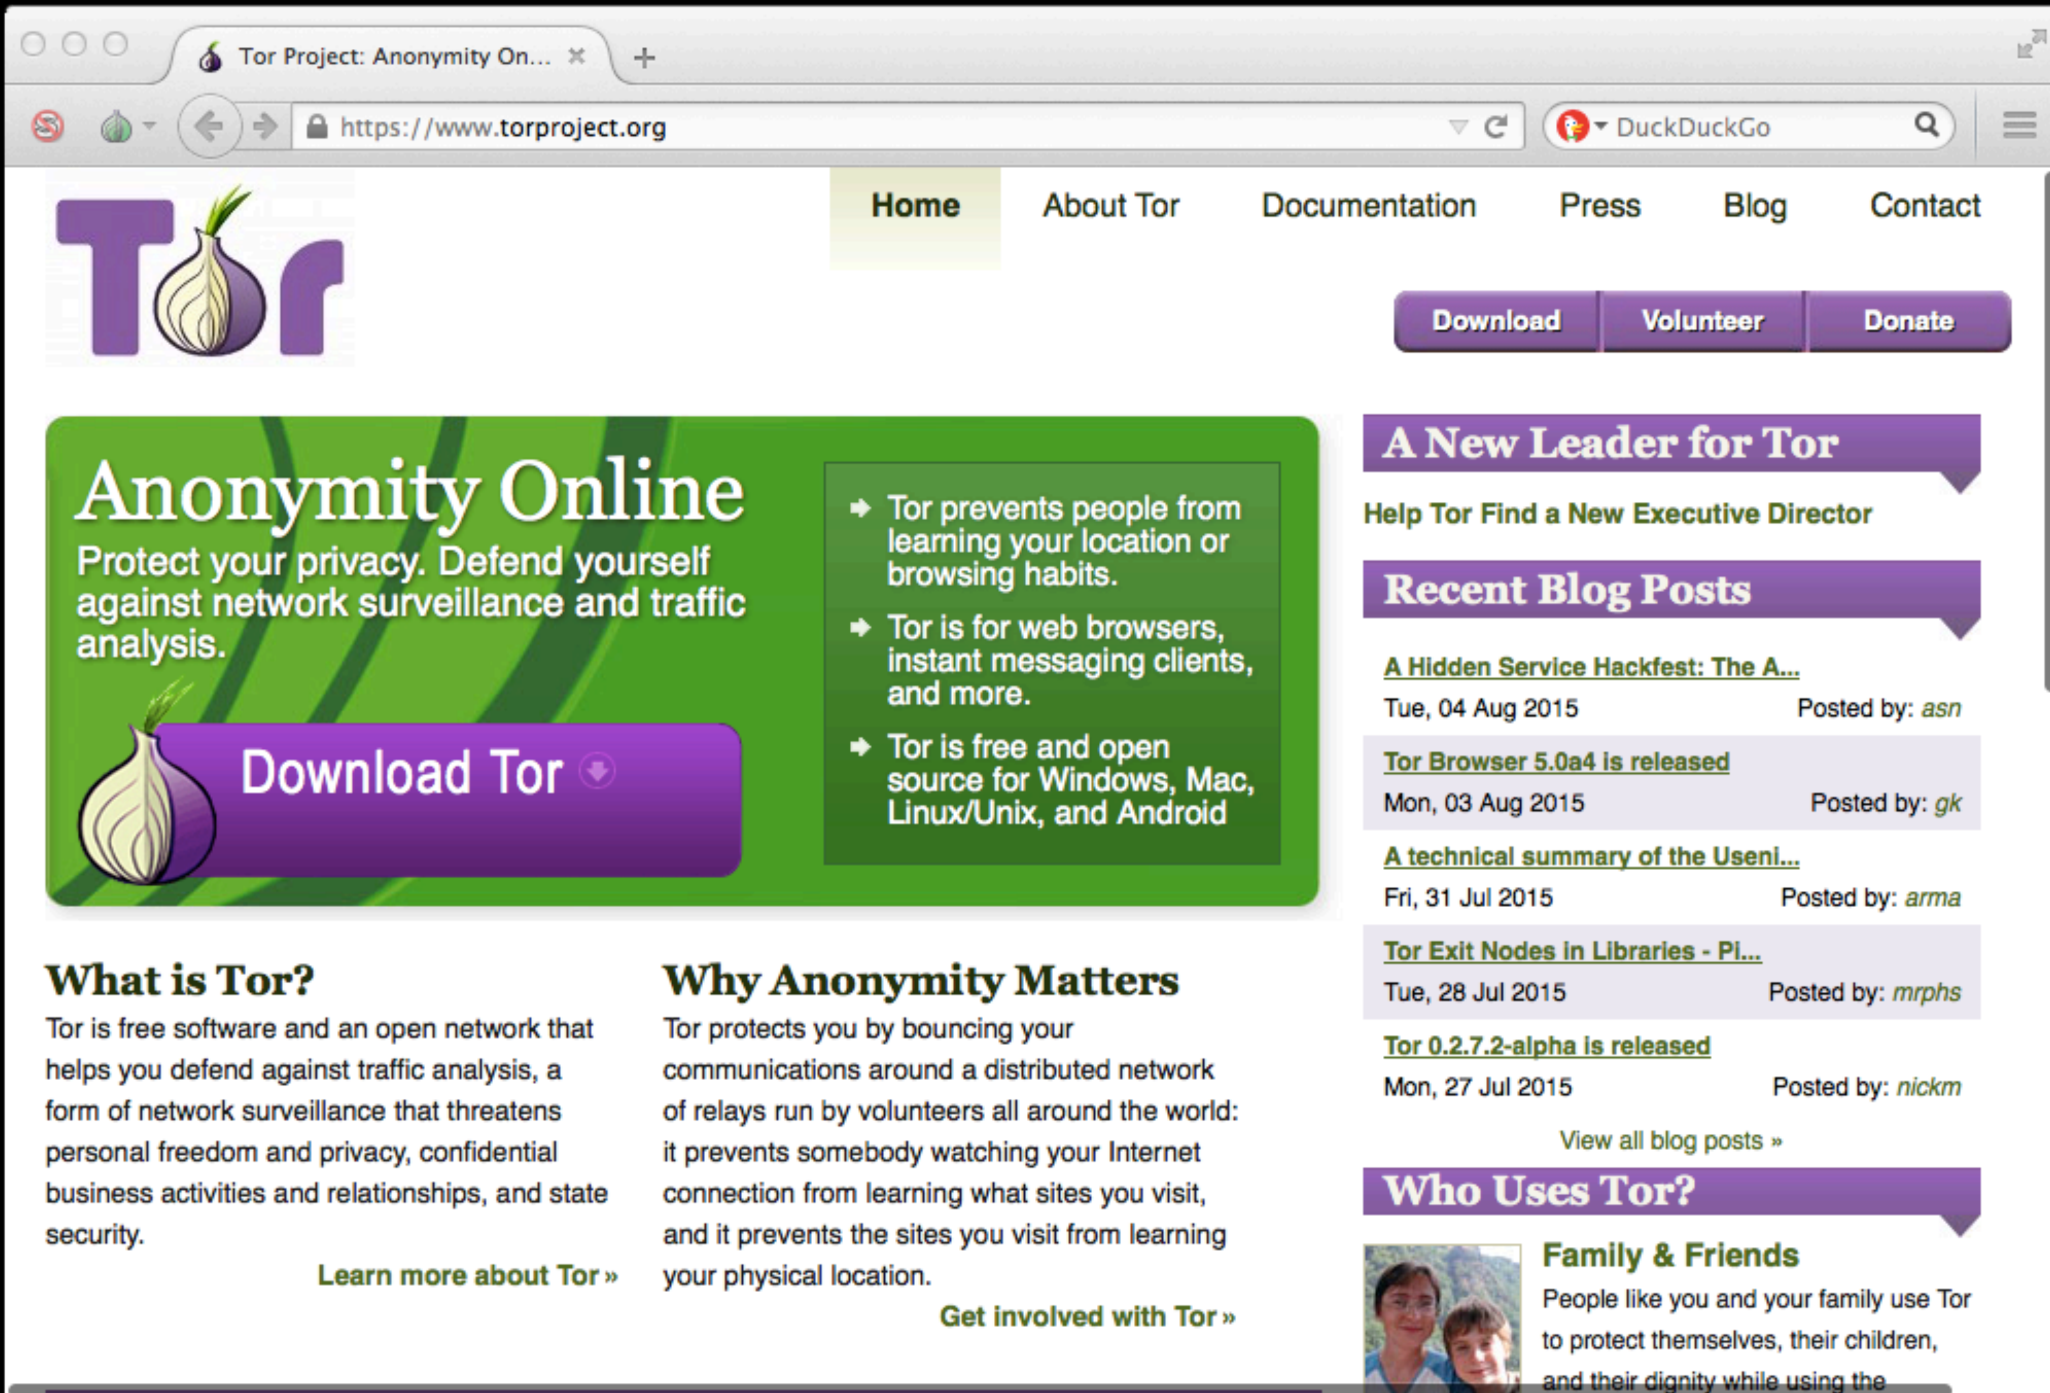
Tor Project: Anonymity On...
https://www.torproject.org
DuckDuckGo
Home
About Tor
Documentation
Press
Blog
Contact
Download
Volunteer
Donate
Anonymity Online
Protect your privacy. Defend yourself against network surveillance and traffic analysis.
Download Tor
Tor prevents people from learning your location or browsing habits.
Tor is for web browsers, instant messaging clients, and more.
Tor is free and open source for Windows, Mac, Linux/Unix, and Android
A New Leader for Tor
Help Tor Find a New Executive Director
Recent Blog Posts
A Hidden Service Hackfest: The A...
Tue, 04 Aug 2015 Posted by: asn
Tor Browser 5.0a4 is released
Mon, 03 Aug 2015 Posted by: gk
A technical summary of the Useni...
Fri, 31 Jul 2015 Posted by: arma
Tor Exit Nodes in Libraries - Pi...
Tue, 28 Jul 2015 Posted by: mrphs
Tor 0.2.7.2-alpha is released
Mon, 27 Jul 2015 Posted by: nickm
View all blog posts »
Who Uses Tor?
Family & Friends
People like you and your family use Tor to protect themselves, their children, and their dignity while using the
What is Tor?
Tor is free software and an open network that helps you defend against traffic analysis, a form of network surveillance that threatens personal freedom and privacy, confidential business activities and relationships, and state security.
Learn more about Tor »
Why Anonymity Matters
Tor protects you by bouncing your communications around a distributed network of relays run by volunteers all around the world: it prevents somebody watching your Internet connection from learning what sites you visit, and it prevents the sites you visit from learning your physical location.
Get involved with Tor »

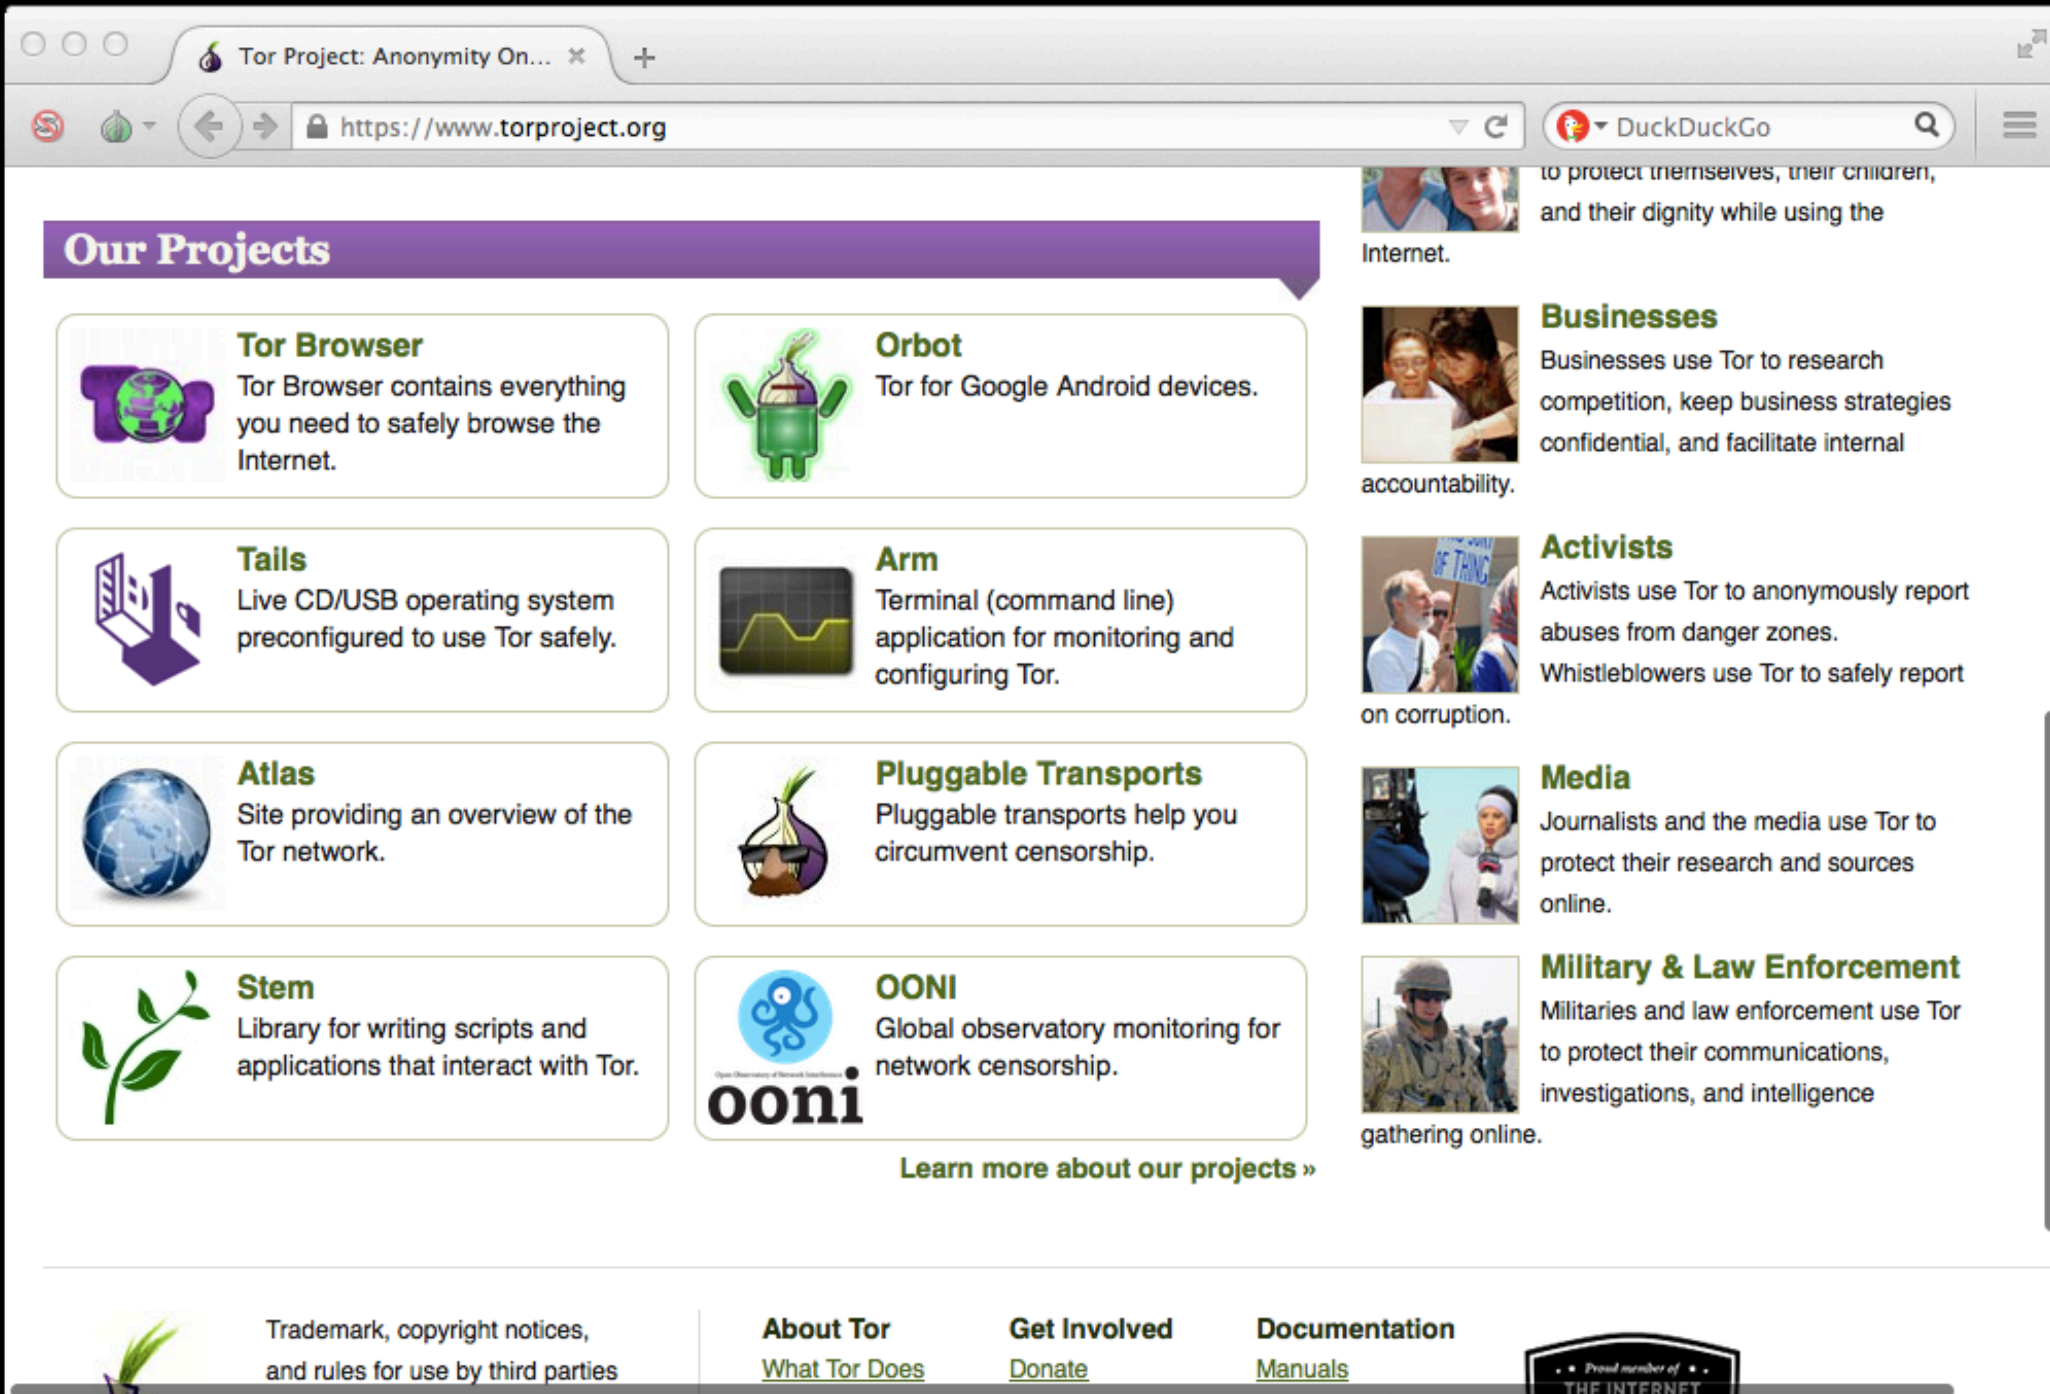
Tor Project: Anonymity On...
https://www.torproject.org
DuckDuckGo
Our Projects
Tor Browser
Tor Browser contains everything you need to safely browse the Internet.
Orbot
Tor for Google Android devices.
Tails
Live CD/USB operating system preconfigured to use Tor safely.
Arm
Terminal (command line) application for monitoring and configuring Tor.
Atlas
Site providing an overview of the Tor network.
Pluggable Transports
Pluggable transports help you circumvent censorship.
Stem
Library for writing scripts and applications that interact with Tor.
OONI
Global observatory monitoring for network censorship.
ooni
Learn more about our projects »
to protect themselves, their children, and their dignity while using the Internet.
Businesses
Businesses use Tor to research competition, keep business strategies confidential, and facilitate internal accountability.
Activists
Activists use Tor to anonymously report abuses from danger zones. Whistleblowers use Tor to safely report on corruption.
Media
Journalists and the media use Tor to protect their research and sources online.
Military & Law Enforcement
Militaries and law enforcement use Tor to protect their communications, investigations, and intelligence gathering online.
Trademark, copyright notices, and rules for use by third parties
About Tor
What Tor Does
Get Involved
Donate
Documentation
Manuals
THE INTERNET

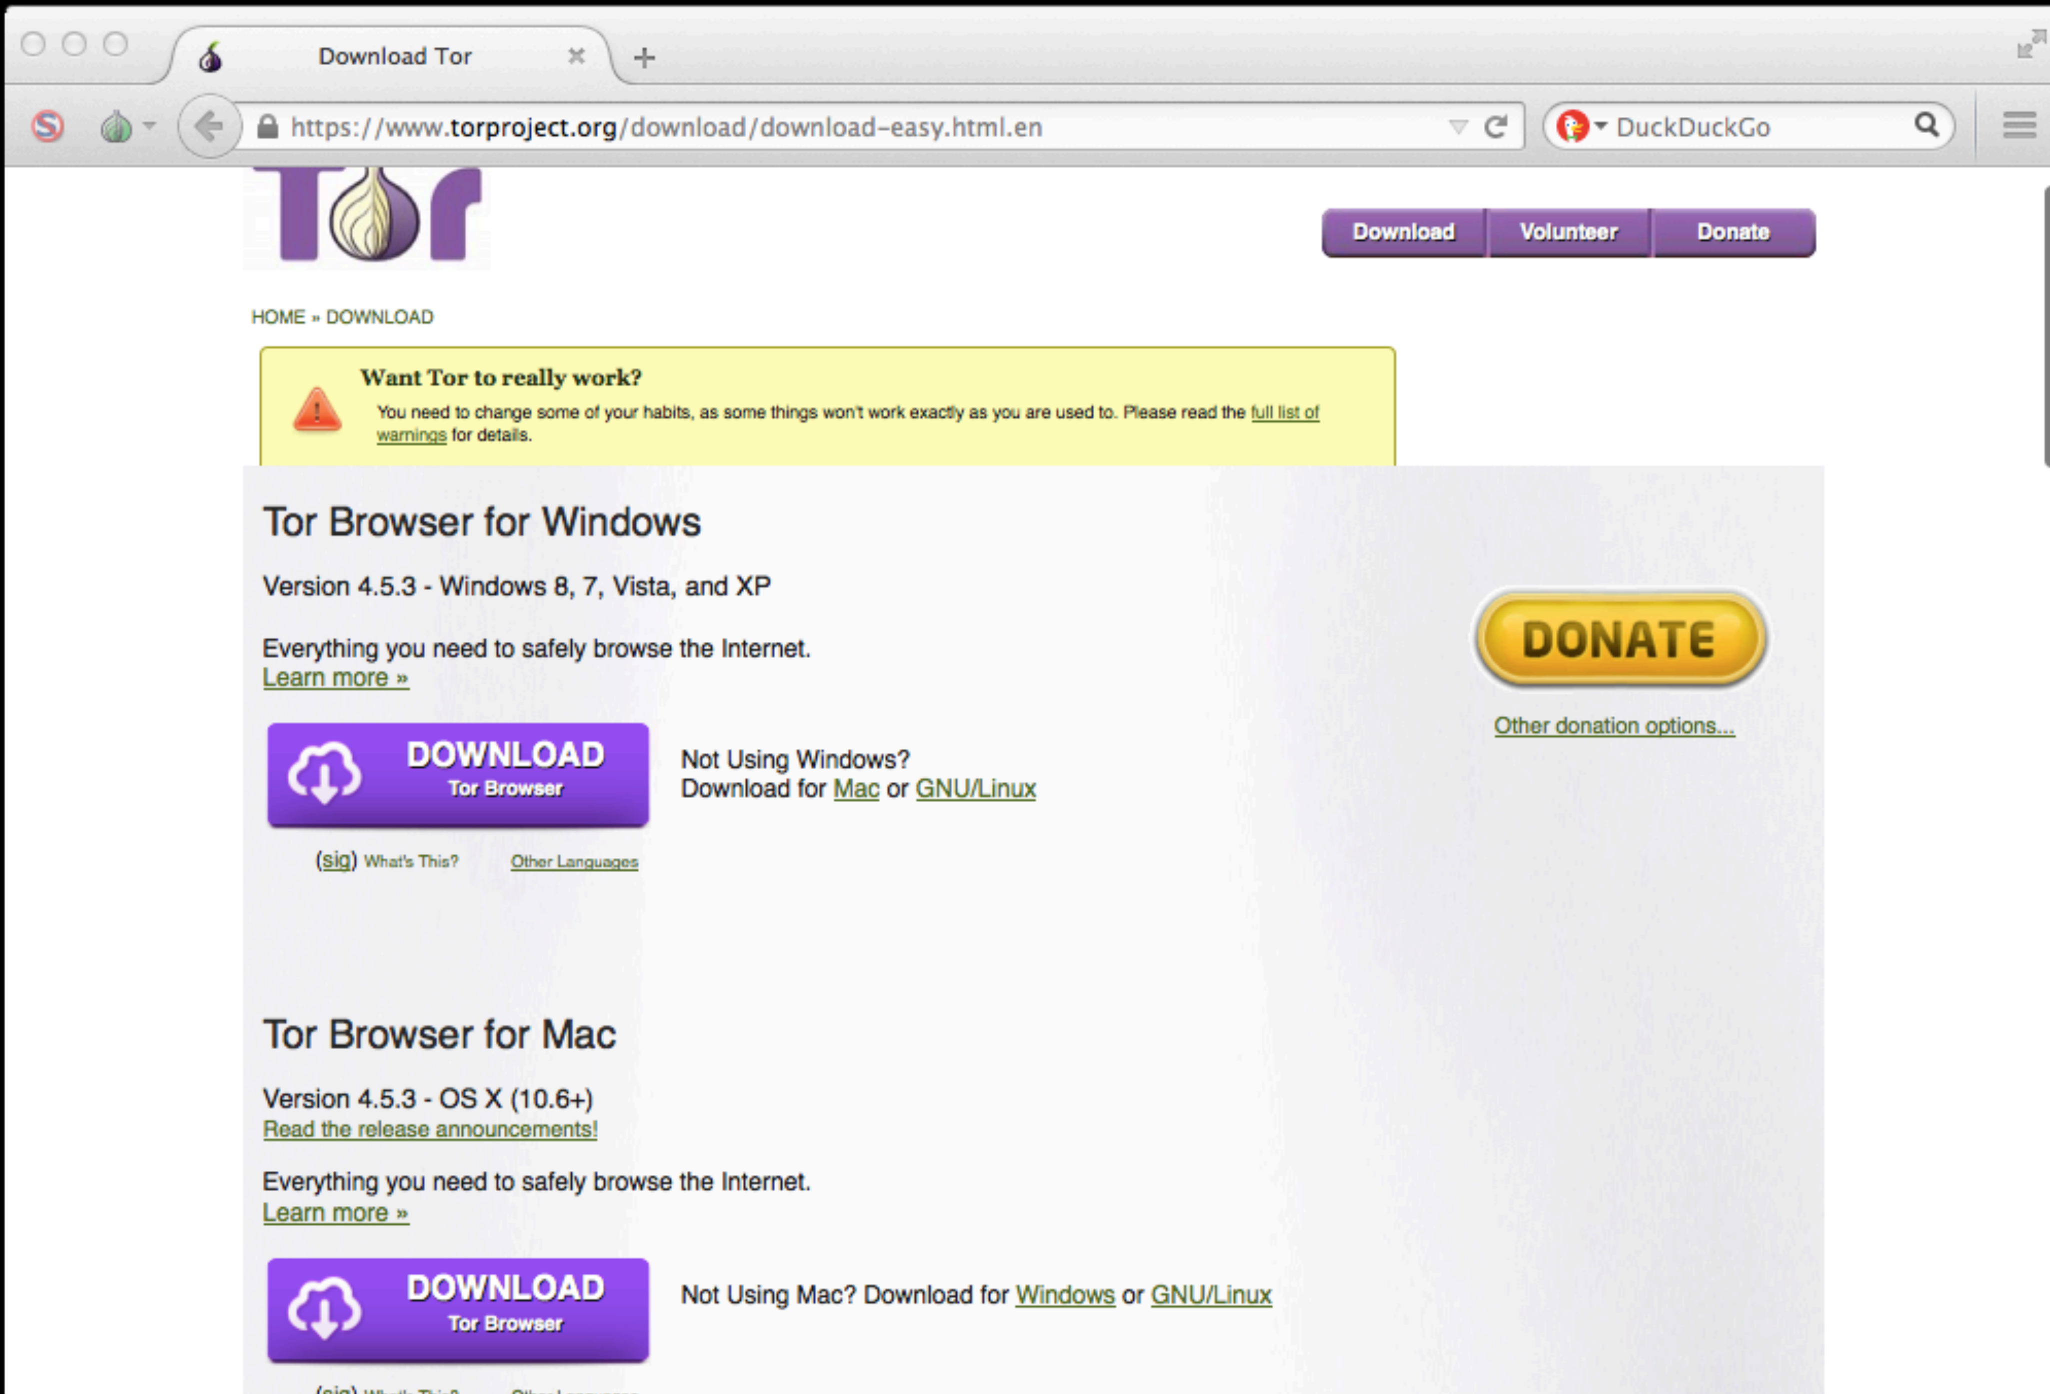
Download Tor
https://www.torproject.org/download/download-easy.html.en
DuckDuckGo
Download
Volunteer
Donate
HOME » DOWNLOAD
Want Tor to really work?
You need to change some of your habits, as some things won't work exactly as you are used to. Please read the full list of warnings for details.
Tor Browser for Windows
Version 4.5.3 - Windows 8, 7, Vista, and XP
Everything you need to safely browse the Internet.
Learn more »
DOWNLOAD
Tor Browser
Not Using Windows?
Download for Mac or GNU/Linux
(sig) What's This?
Other Languages
DONATE
Other donation options...
Tor Browser for Mac
Version 4.5.3 - OS X (10.6+)
Read the release announcements!
Everything you need to safely browse the Internet.
Learn more »
DOWNLOAD
Tor Browser
Not Using Mac? Download for Windows or GNU/Linux

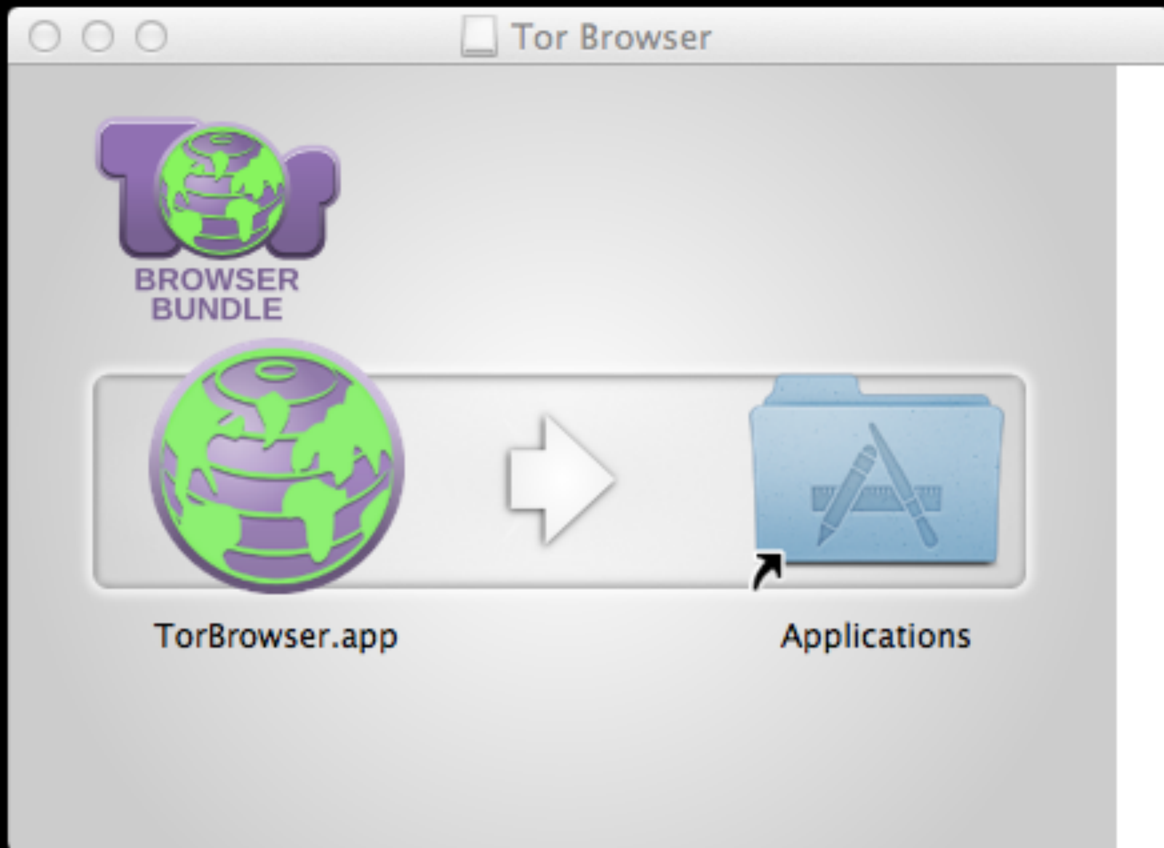
Tor Browser
BROWSER
BUNDLE
TorBrowser.app
Applications


STOP

How do we know I'm installing Tor?

- Known cases of bogus sites offering not-Tor
- HTTPS isn't enough here
  - Browser-recognized authorities have issued bogus certificates for torproject.org

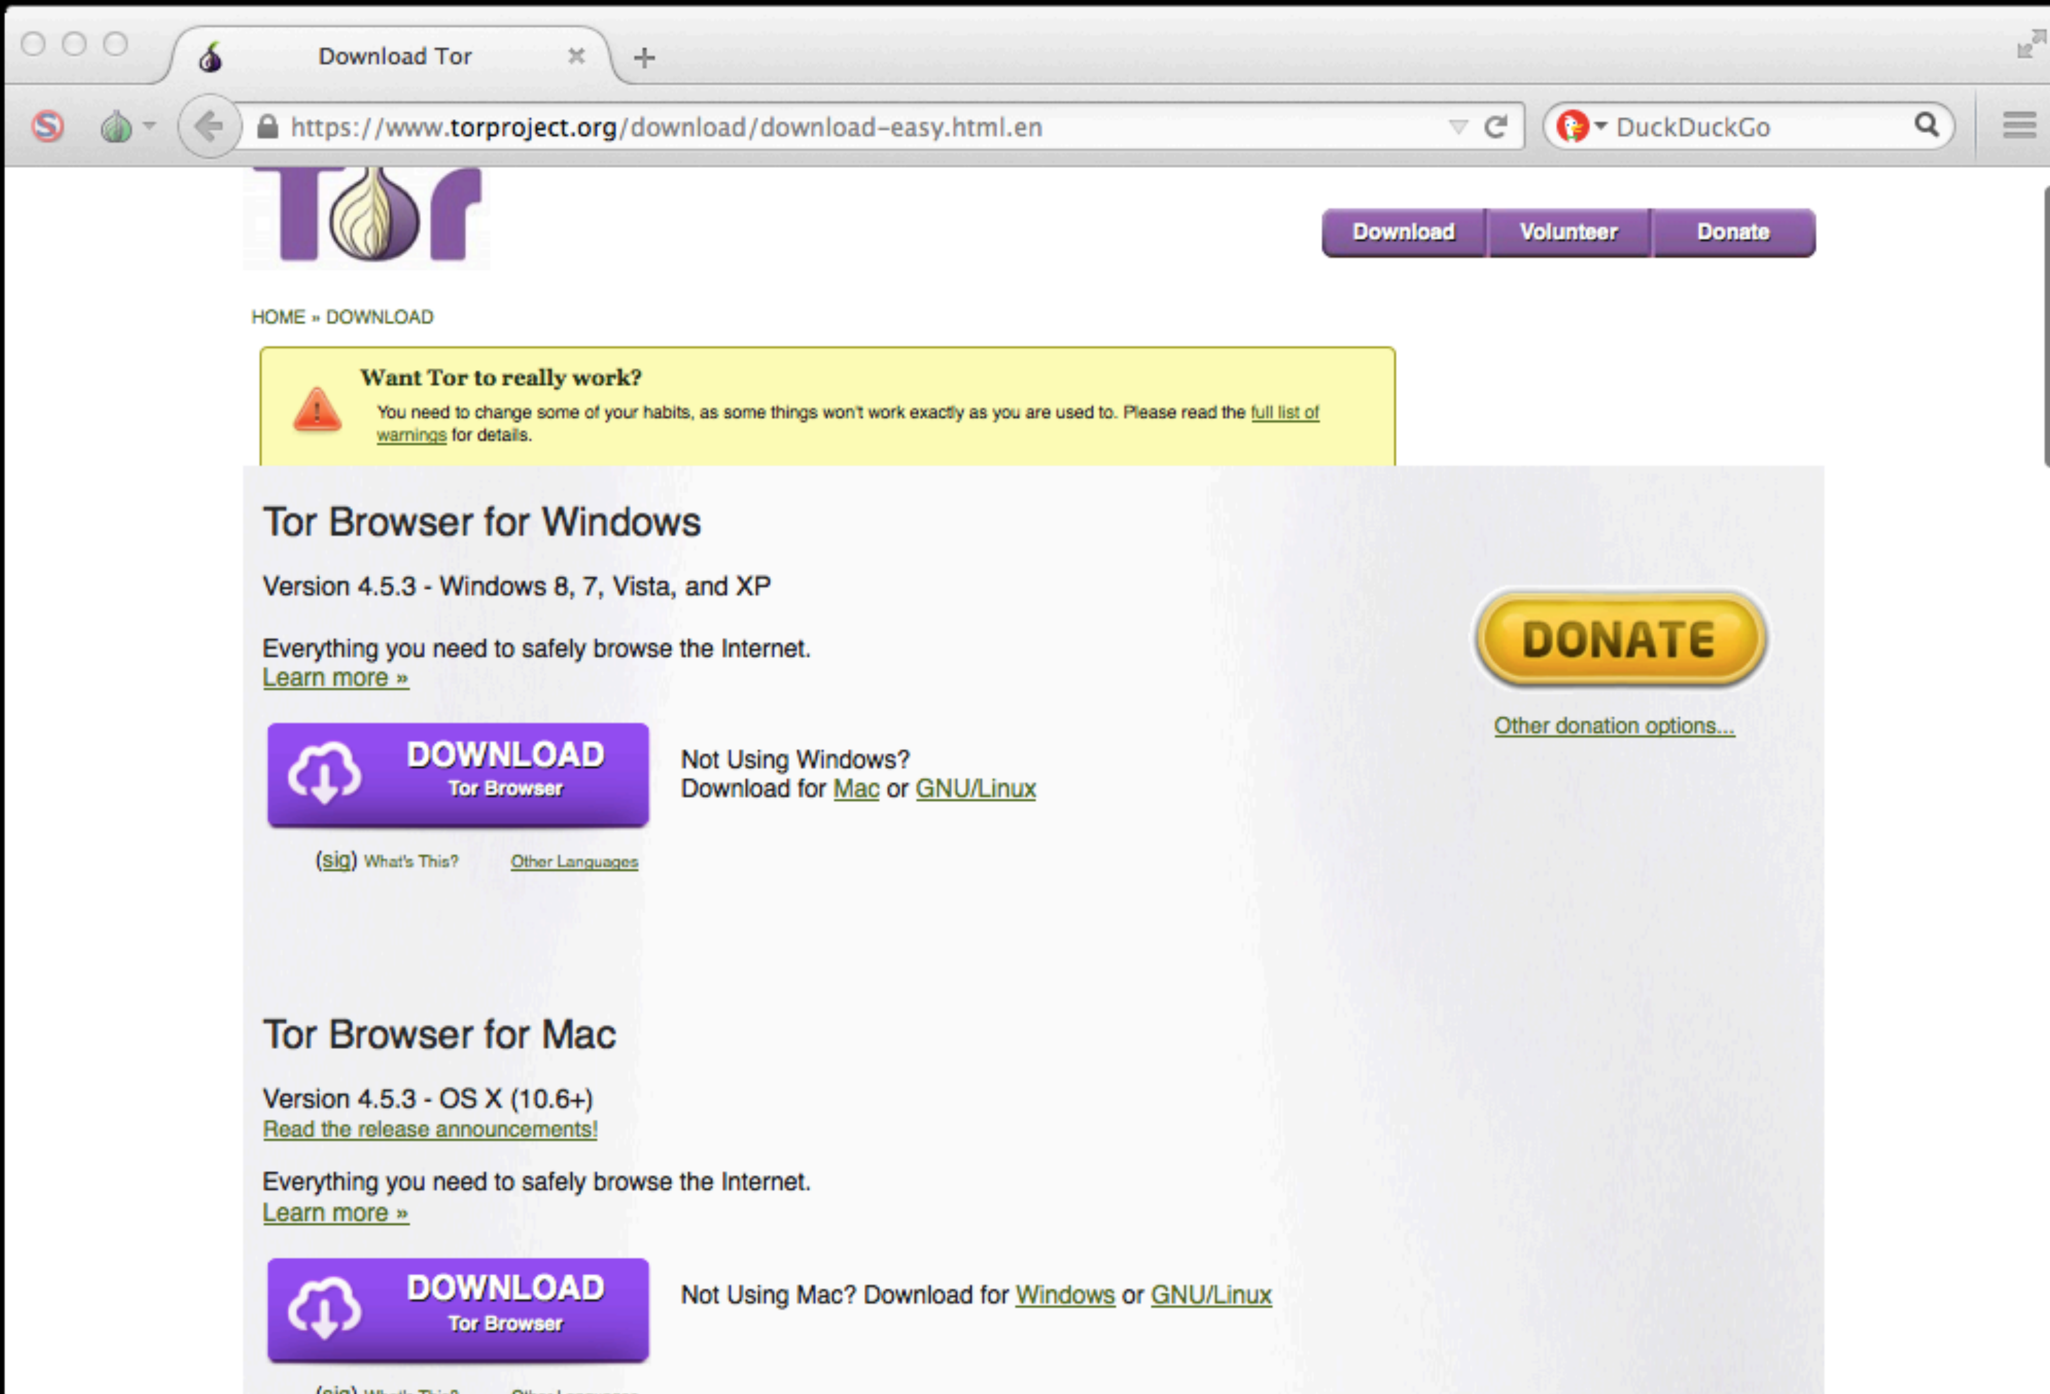
Download Tor
https://www.torproject.org/download/download-easy.html.en
DuckDuckGo
Download
Volunteer
Donate
HOME » DOWNLOAD
Want Tor to really work?
You need to change some of your habits, as some things won't work exactly as you are used to. Please read the full list of warnings for details.
Tor Browser for Windows
Version 4.5.3 - Windows 8, 7, Vista, and XP
Everything you need to safely browse the Internet.
Learn more »
DOWNLOAD
Tor Browser
Not Using Windows?
Download for Mac or GNU/Linux
(sig) What's This?
Other Languages
DONATE
Other donation options...
Tor Browser for Mac
Version 4.5.3 - OS X (10.6+)
Read the release announcements!
Everything you need to safely browse the Internet.
Learn more »
DOWNLOAD
Tor Browser
Not Using Mac? Download for Windows or GNU/Linux

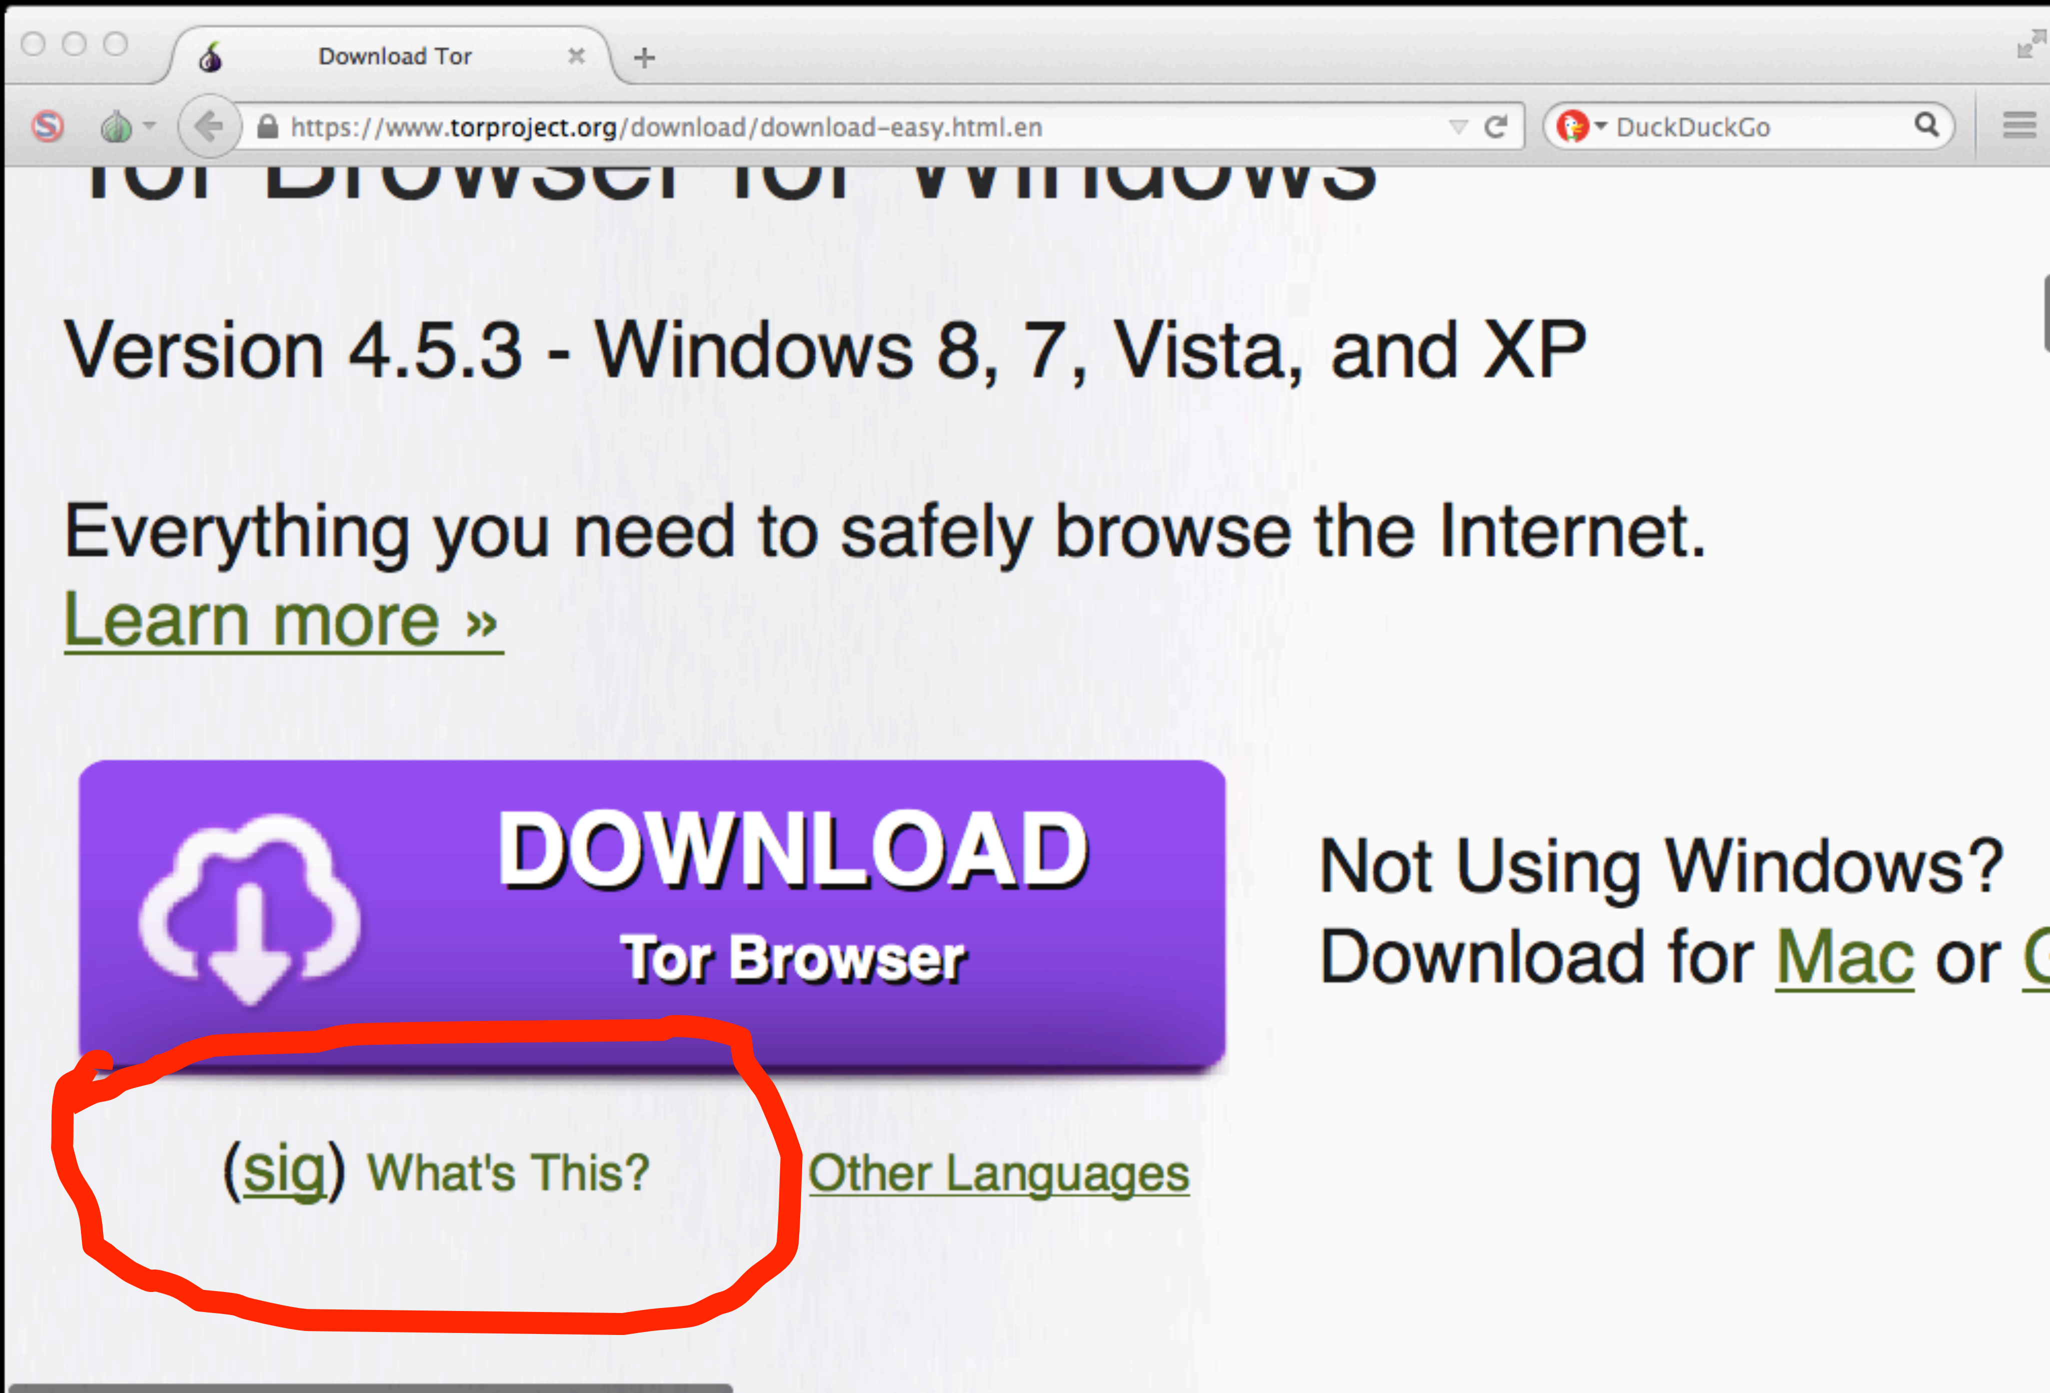
Download Tor
https://www.torproject.org/download/download-easy.html.en
DuckDuckGo
Version 4.5.3 - Windows 8, 7, Vista, and XP
Everything you need to safely browse the Internet.
Learn more »
DOWNLOAD
Tor Browser
Not Using Windows?
Download for Mac or G
(sig) What's This?
Other Languages

Tor Project: Verifying Sign...

https://www.torproject.org/docs/verifying-signatures.html.en

Download | Volunteer | Donate

HOME » VERIFYING SIGNATURES

- Documentation Overview
- ▼ Installation Guides
  - Installing on Windows
  - Installing on Linux/BSD/Unix
  - Installing Tor on Debian/Ubuntu
  - Installing Tor on Fedora/CentOS
  - Installing Tor on Mac OS X
  - Installing Tor on Android
  - Installing Tor on Maemo/N900
  - **Verify our GPG signatures**
- ▶ Manuals
- Tor Wiki
- General FAQ

How to verify signatures for packages

What is a signature and why should I check it?

How do you know that the Tor program you have is really the one we made? Many Tor users have very real adversaries who might try to give them a fake version of Tor — and it doesn't matter how secure and anonymous Tor is if you're not running the real Tor.

An attacker could try a variety of attacks to get you to download a fake Tor. For example, he could trick you into thinking some other website is a great place to download Tor. That's why you should always download Tor from **https**://www.torproject.org/. The https part means there's encryption and authentication between your browser and the website, making it much harder for the attacker to modify your download. But it's not perfect. Some places in the world block the Tor website, making users try somewhere else. Large companies sometimes force employees to use a modified browser, so the company can listen in on all their browsing. We've even seen attackers who have the ability to trick your browser into thinking you're talking to the Tor website with https when you're not.

Some software sites list sha1 hashes alongside the software on their website, so users can verify that they downloaded the file without any errors. These "checksums" help you answer the question "Did I download this file correctly from whoever sent it to me?" They do a good job at making sure you didn't have any random errors in your download, but they don't help you figure out whether you were downloading it from the attacker. The better question to answer is: "Is this file that I just downloaded the file that Tor intended me to get?"

Where do I get the signatures and the keys that made them?

Each file on our download page is accompanied by a file with the same name as the package and the extension ".asc".

How do we know I'm installing Tor?

- Known cases of bogus sites offering not-Tor
- HTTPS isn't enough here
  - Browser-recognized authorities have issued bogus certificates for torproject.org

How do we know the Tor from the Tor Project Inc. is OK?

- It's baked into our approach to offering onion routing to the public

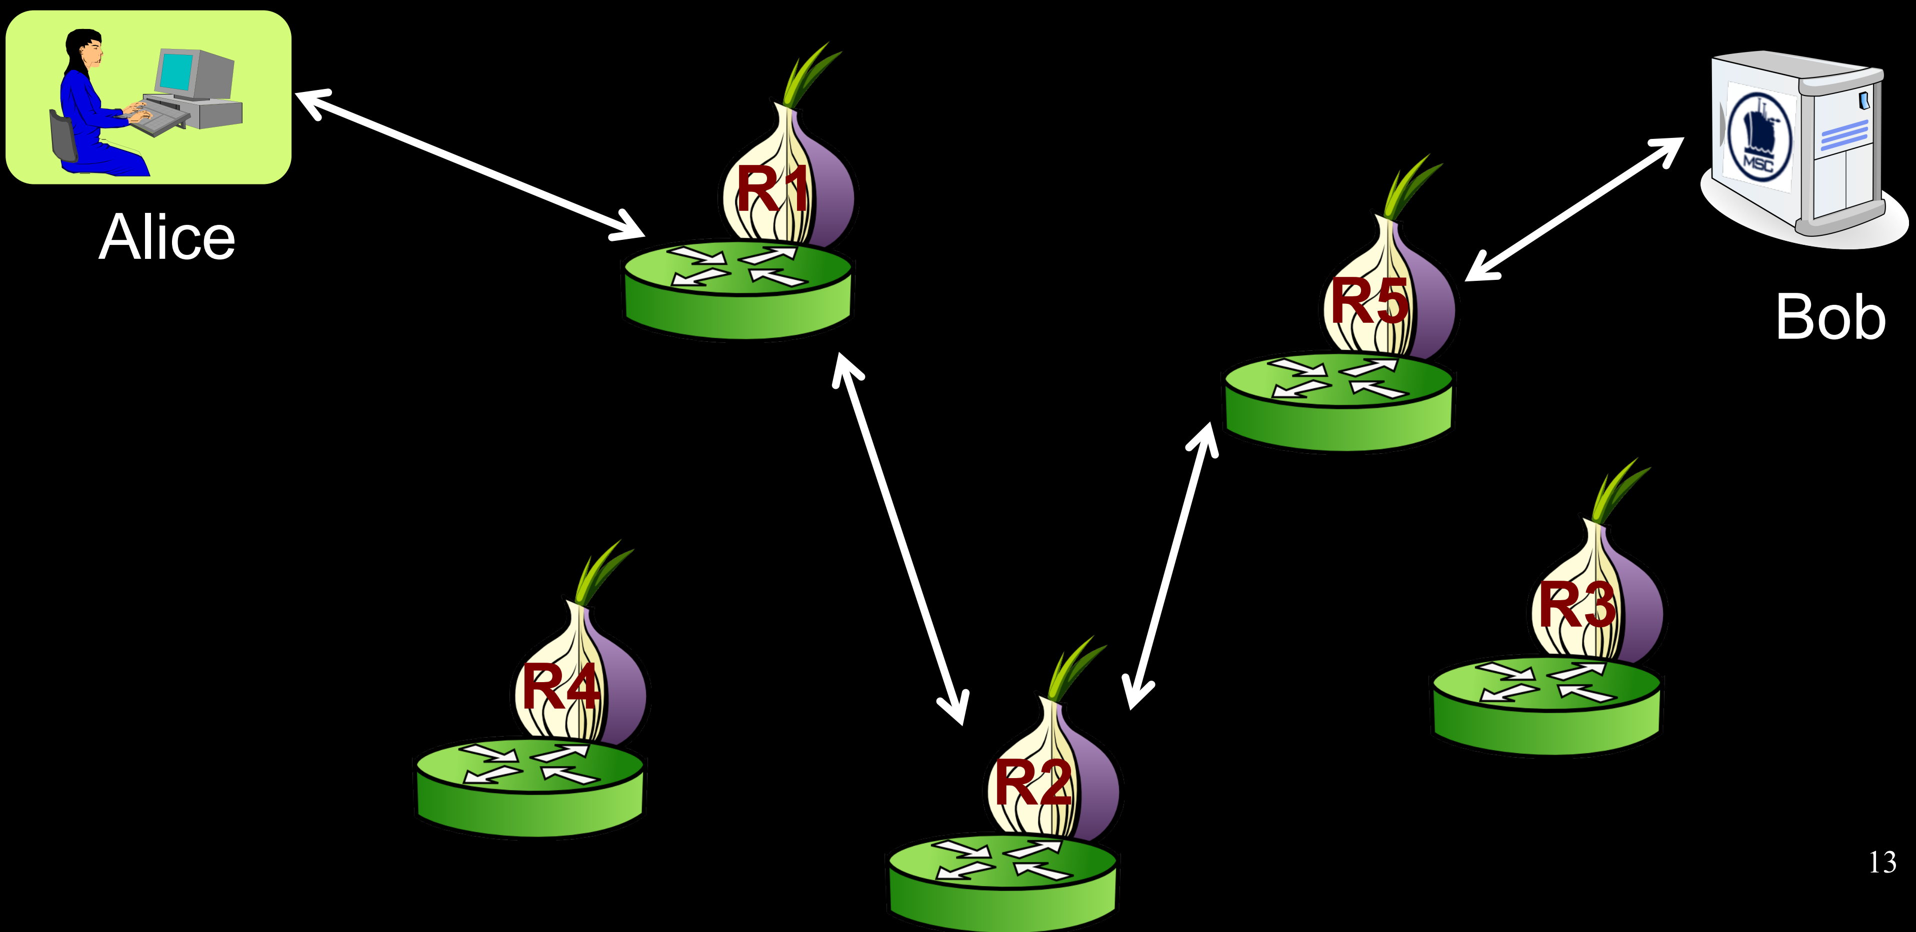

How do we know the Tor from the Tor Project Inc. is OK?

- It's baked into our approach to offering onion routing to the public

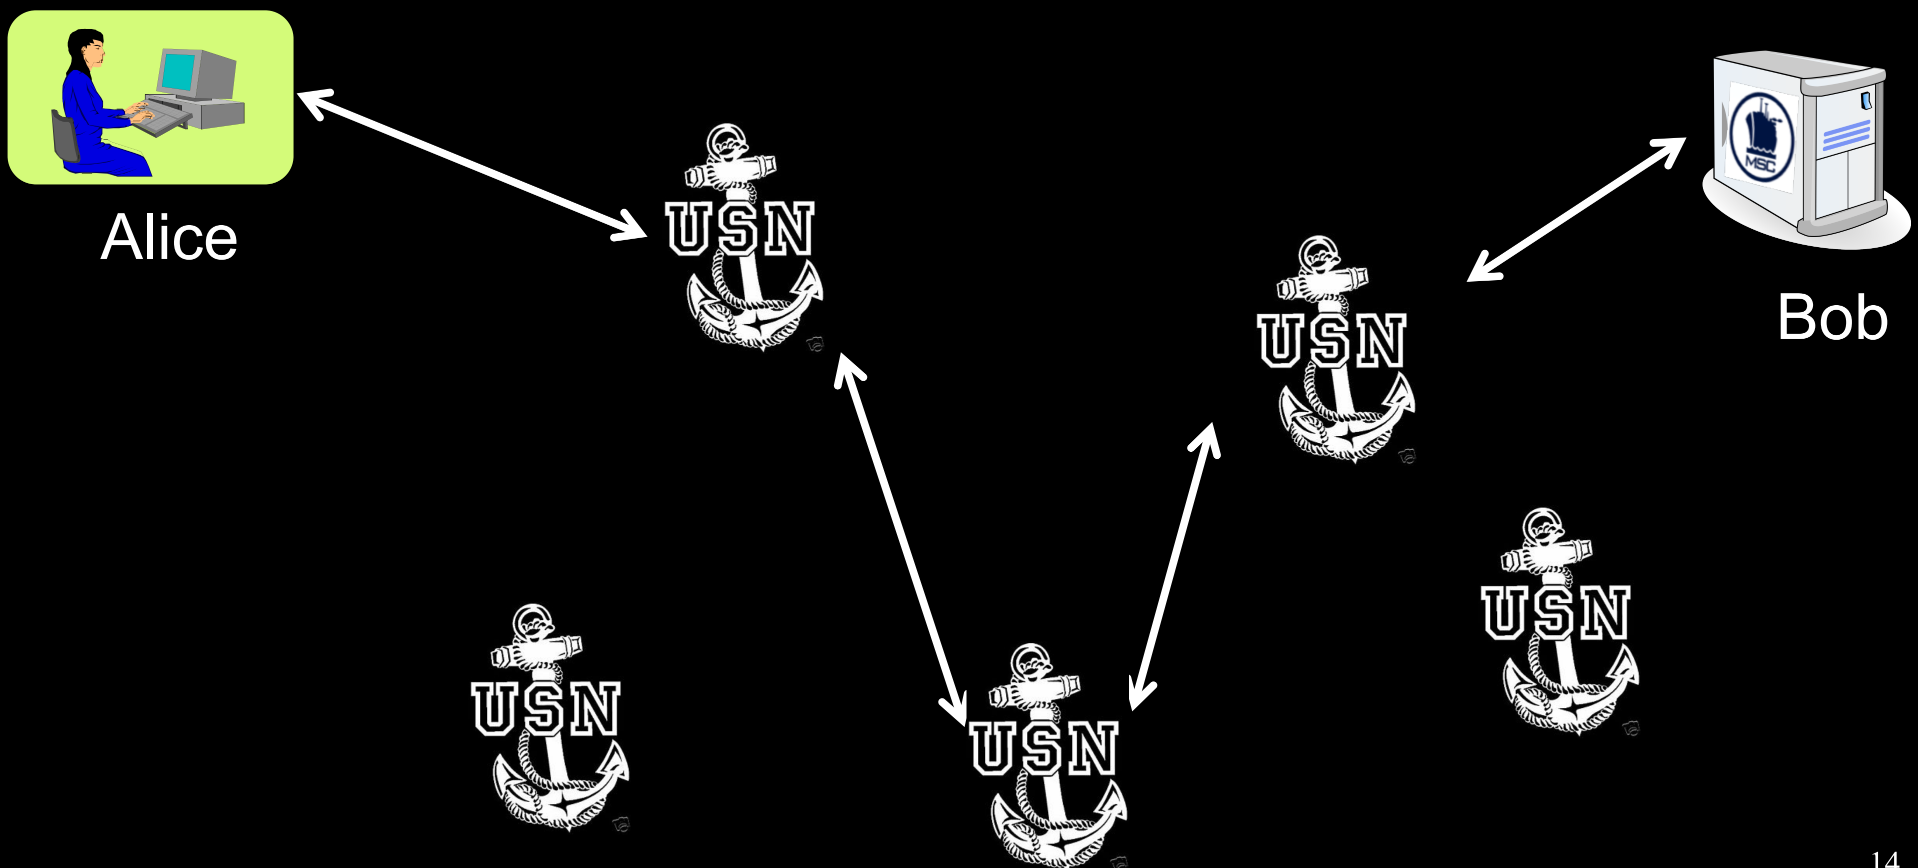

How do we know the Tor from the Tor Project Inc. is OK?

- It's baked into our approach to offering onion routing to the public
- Carry traffic for a diverse user population
  - not just Navy or U.S. govt.
  - cannot have single point of failure/trust for any type of user
    - Diversely managed infrastructure
      - Open source

How do we know the Tor from the Tor Project Inc. is OK?

- Design, Specifications, and Software amongst most scrutinized on the planet
- Can download (signed) source code and build the binaries yourself.

How do we know the Tor from the Tor Project Inc. is OK?

- Design, Specifications, and Software amongst most scrutinized on the planet
- Can download (signed) source code and build the binaries yourself.

But we’re not done yet!

How do we know the Tor from the Tor Project Inc. is OK?

- “Reflections on Trusting Trust” Thompson, Turing award lecture, 1984
- Problem: Creation of signed Tor binaries might have been attacked: compiler, libraries, etc.

How do we know the Tor from the Tor Project Inc. is OK?

- “Reflections on Trusting Trust” Thompson, Turing award lecture, 1984
- Problem: Creation of signed Tor binaries might have been attacked: compiler, libraries, etc.

- Why should you care if you verify source and compile yourself on a trusted system?

How do we know the Tor from the Tor Project Inc. is OK?

- Most users will be running TPI compiled binaries
- It would be good to protect them.
- What if you're purely self-interested?

How do we know the Tor from the Tor Project Inc. is OK?

- Most users will be running TPI compiled binaries
- It would be good to protect them
- What if you're purely self-interested?
- Relatively small handful of users with self-compiled Tor will stick out significantly
- Many relay operators may also run TPI compiled code

How do we know the Tor from the Tor Project Inc. is OK?

- “Reflections on Trusting Trust” Thompson, Turing award lecture, 1984
- Problem: Creation of signed Tor binaries might have been attacked: compiler, libraries, etc.
- Solution: Deterministic builds
  - Packages identical across software, hardware platforms
  - Distributes the trust in Tor binaries
  - Available for Tor Browser Bundle since 2013

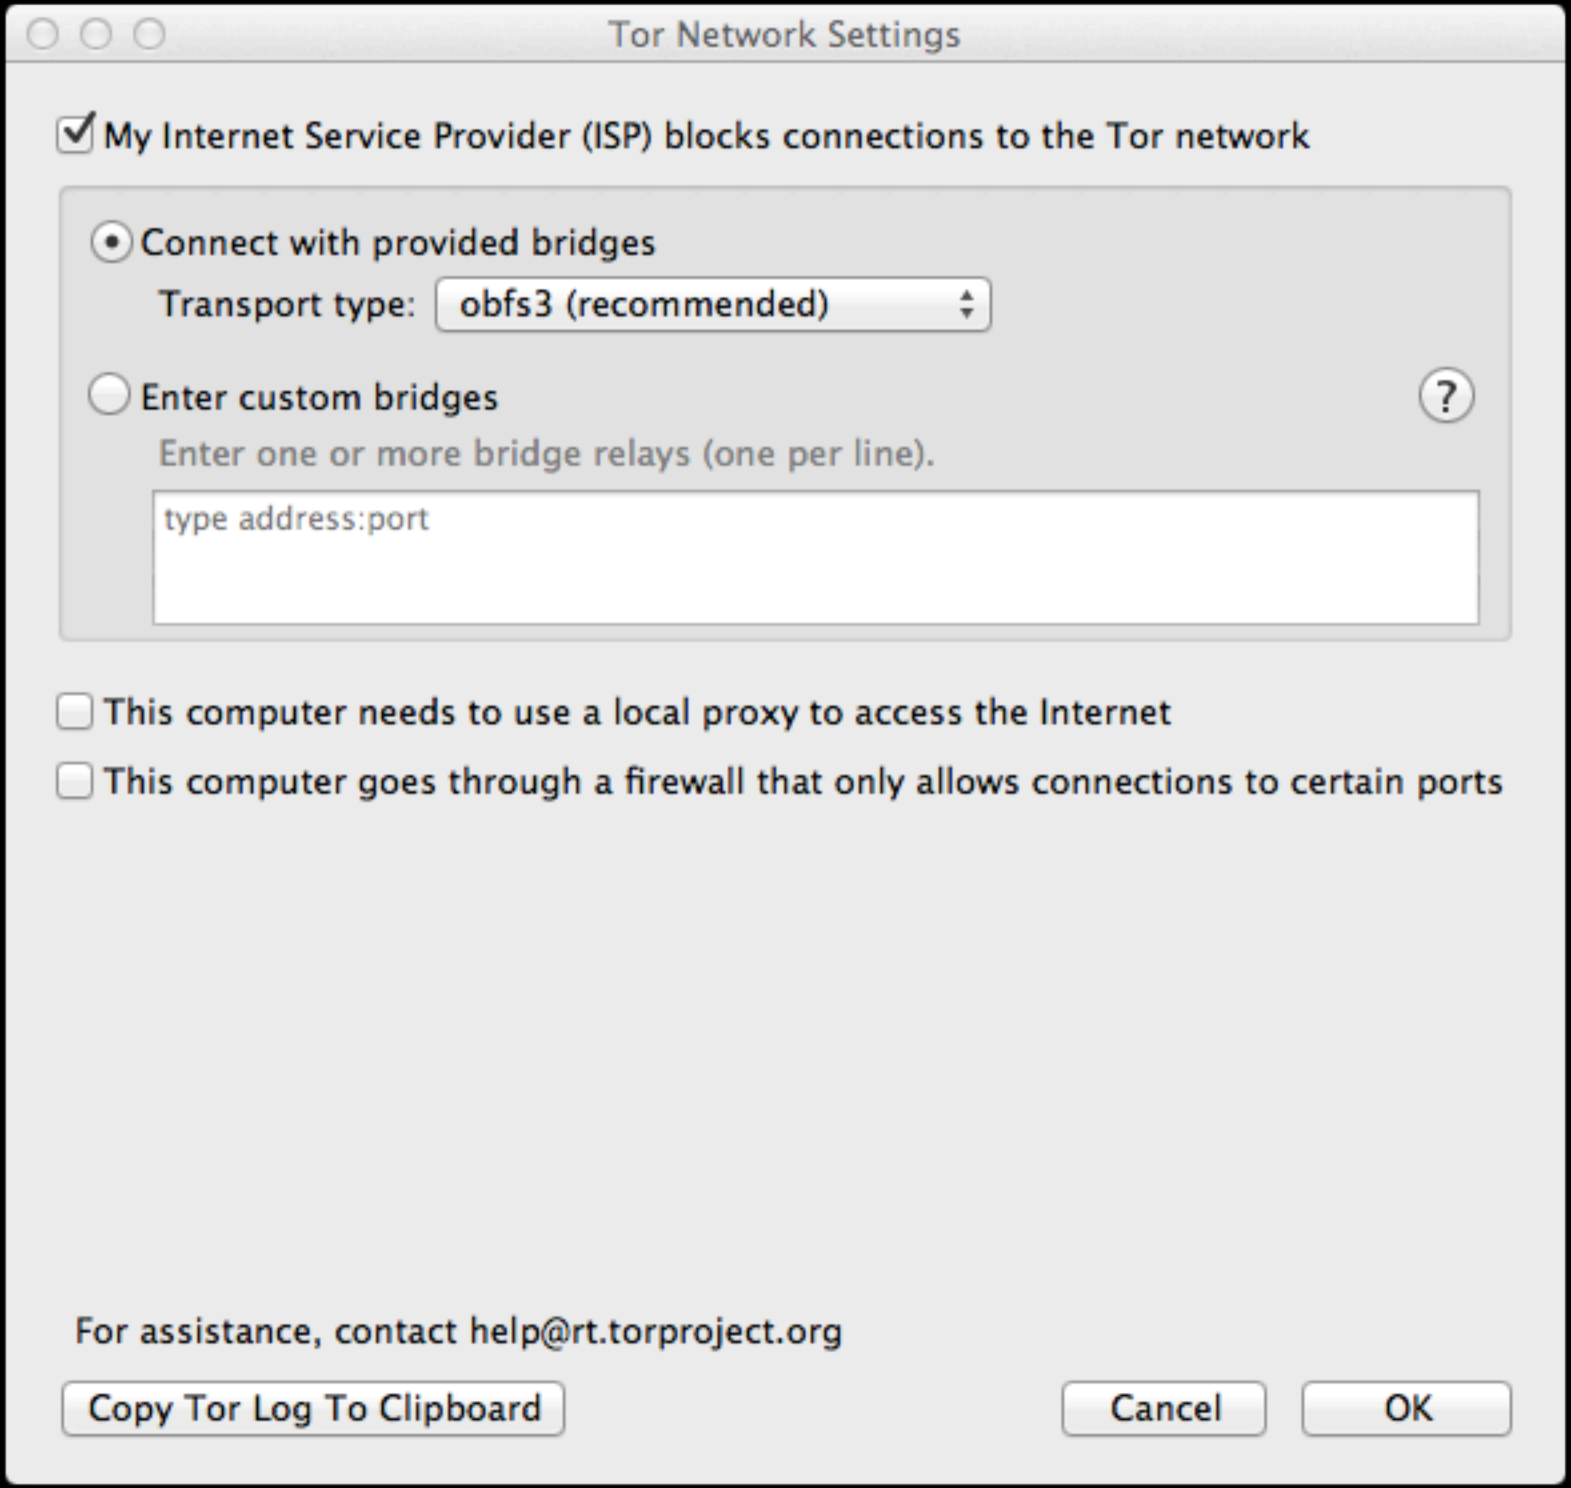
Tor Network Settings
My Internet Service Provider (ISP) blocks connections to the Tor network
Connect with provided bridges
Transport type:
obfs3 (recommended)
Enter custom bridges
?
Enter one or more bridge relays (one per line).
type address:port
This computer needs to use a local proxy to access the Internet
This computer goes through a firewall that only allows connections to certain ports
For assistance, contact help@rt.torproject.org
Copy Tor Log To Clipboard
Cancel
OK

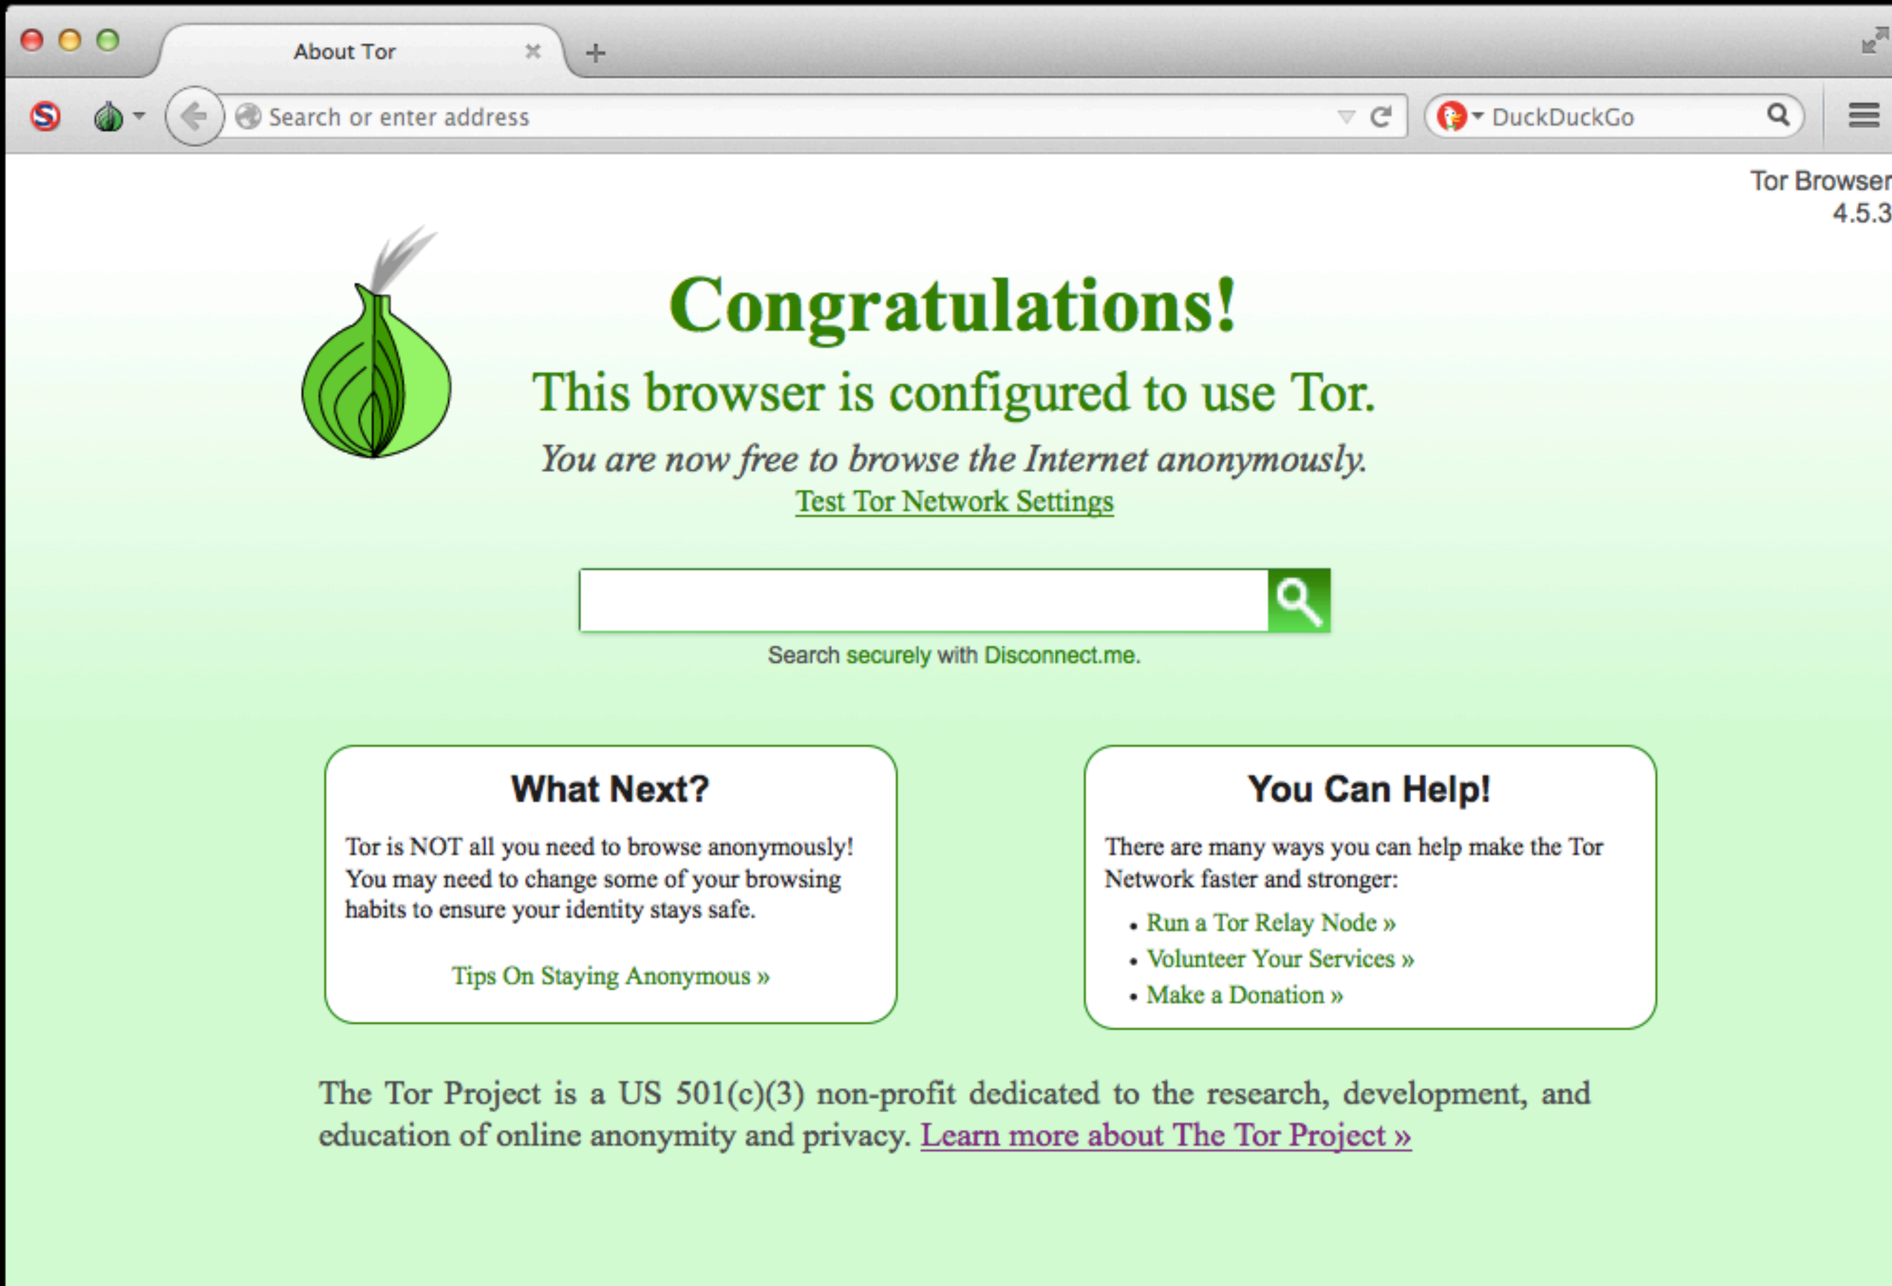
About Tor
Search or enter address
DuckDuckGo
Tor Browser
4.5.3
Congratulations!
This browser is configured to use Tor.
You are now free to browse the Internet anonymously.
Test Tor Network Settings
Search securely with Disconnect.me.
What Next?
Tor is NOT all you need to browse anonymously! You may need to change some of your browsing habits to ensure your identity stays safe.
Tips On Staying Anonymous »
You Can Help!
There are many ways you can help make the Tor Network faster and stronger:
• Run a Tor Relay Node »
• Volunteer Your Services »
• Make a Donation »
The Tor Project is a US 501(c)(3) non-profit dedicated to the research, development, and education of online anonymity and privacy. Learn more about The Tor Project »

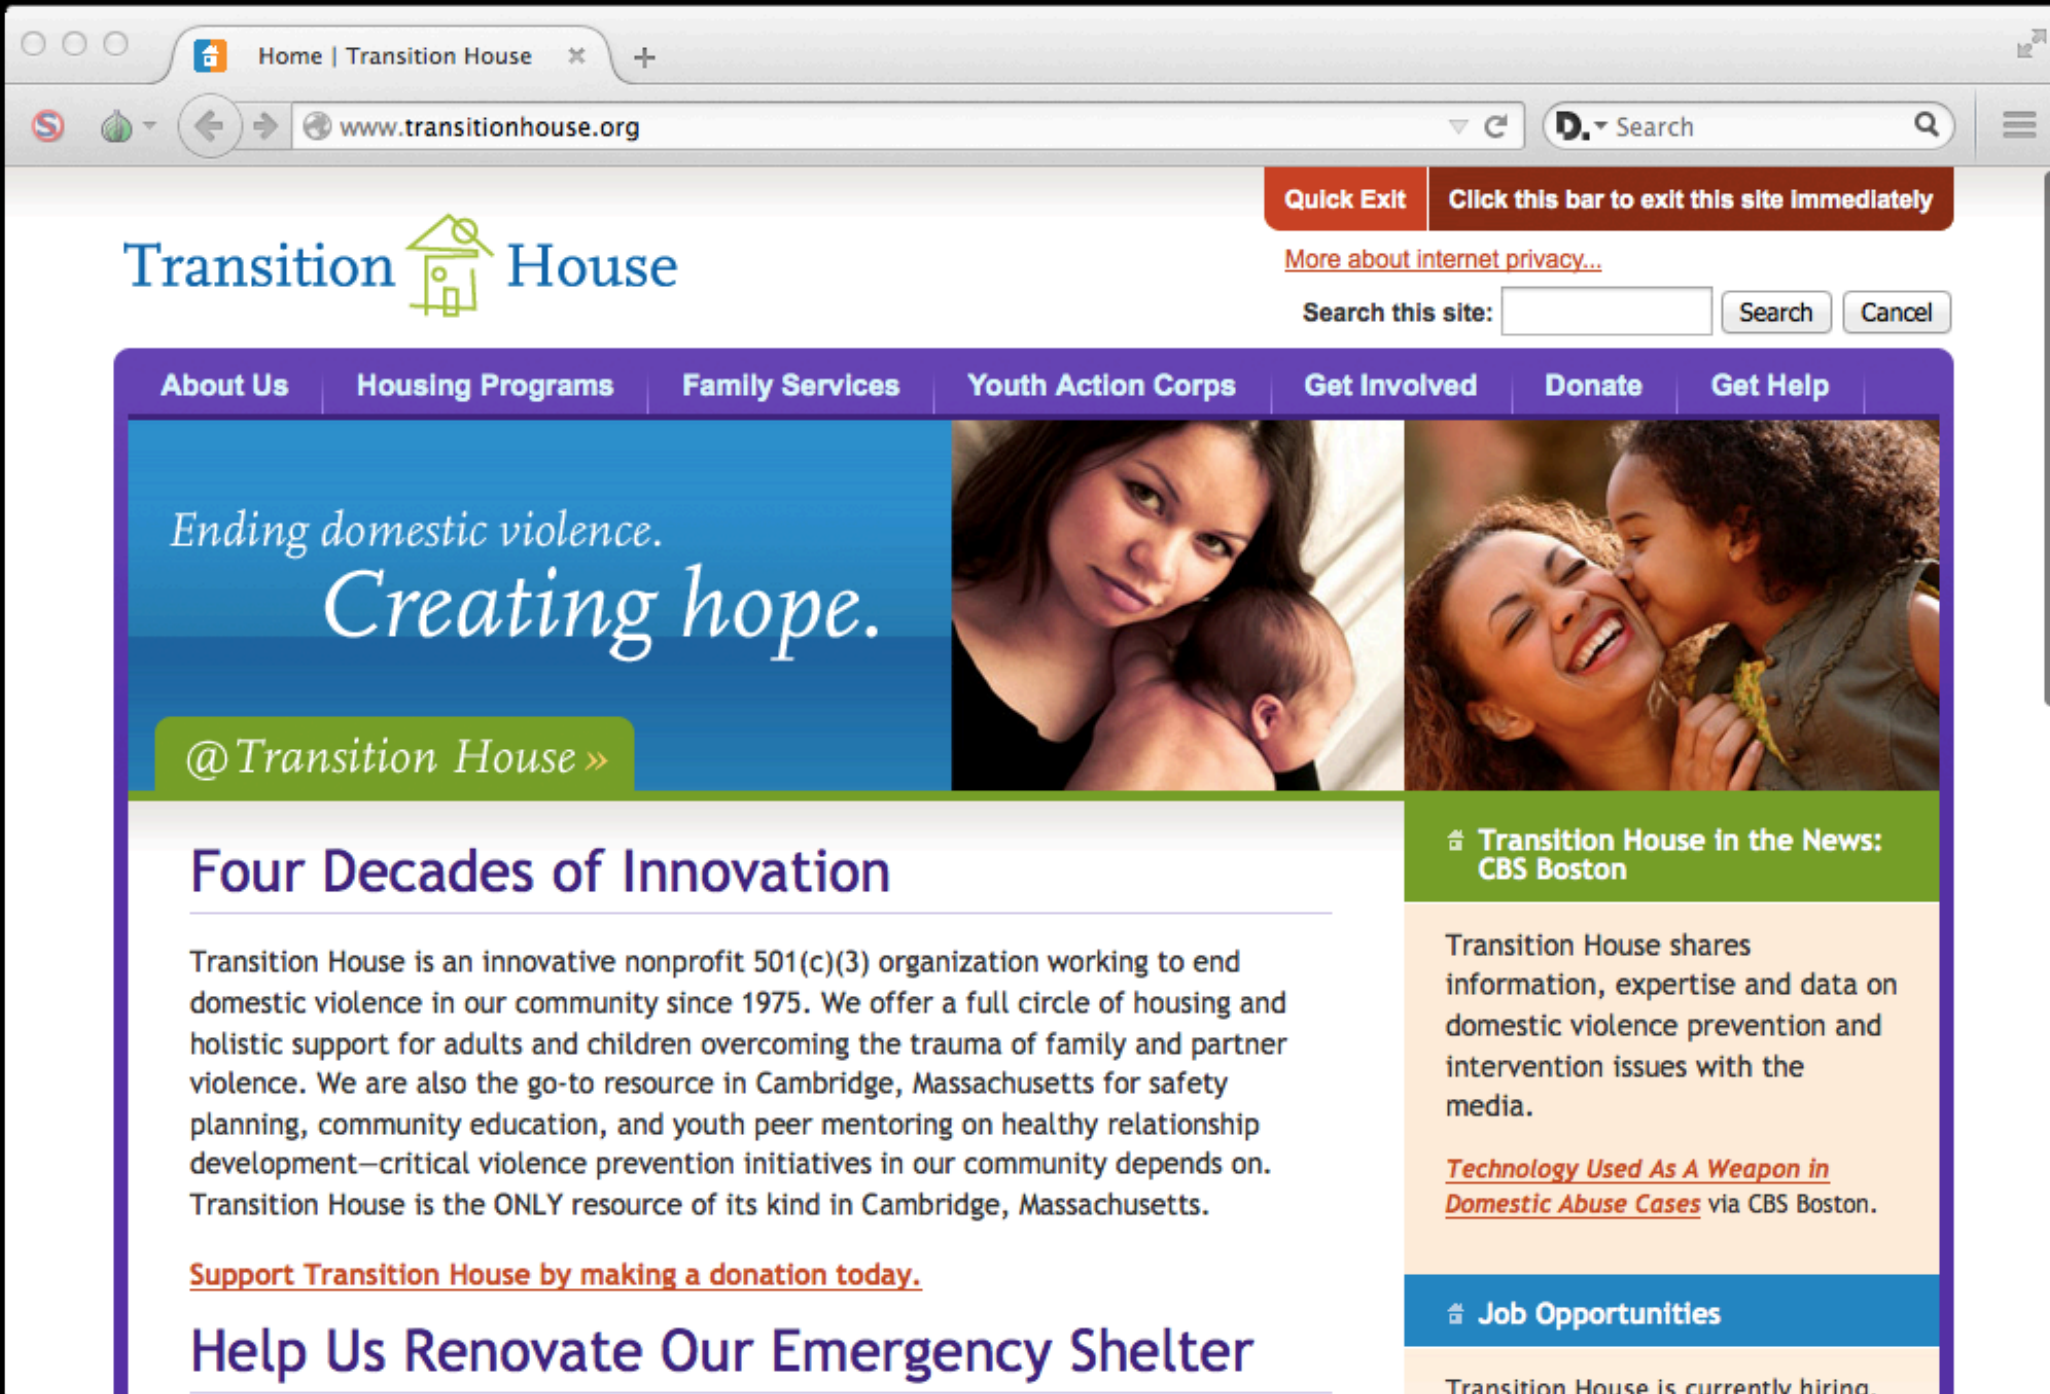
Home | Transition House
www.transitionhouse.org
Search
Quick Exit
Click this bar to exit this site immediately
More about internet privacy...
Search this site:
Search
Cancel
Transition House
About Us
Housing Programs
Family Services
Youth Action Corps
Get Involved
Donate
Get Help
Ending domestic violence.
Creating hope.
@Transition House »
Four Decades of Innovation
Transition House is an innovative nonprofit 501(c)(3) organization working to end domestic violence in our community since 1975. We offer a full circle of housing and holistic support for adults and children overcoming the trauma of family and partner violence. We are also the go-to resource in Cambridge, Massachusetts for safety planning, community education, and youth peer mentoring on healthy relationship development—critical violence prevention initiatives in our community depends on. Transition House is the ONLY resource of its kind in Cambridge, Massachusetts.
Support Transition House by making a donation today.
Help Us Renovate Our Emergency Shelter
Transition House in the News: CBS Boston
Transition House shares information, expertise and data on domestic violence prevention and intervention issues with the media.
Technology Used As A Weapon in Domestic Abuse Cases via CBS Boston.
Job Opportunities
Transition House is currently hiring.

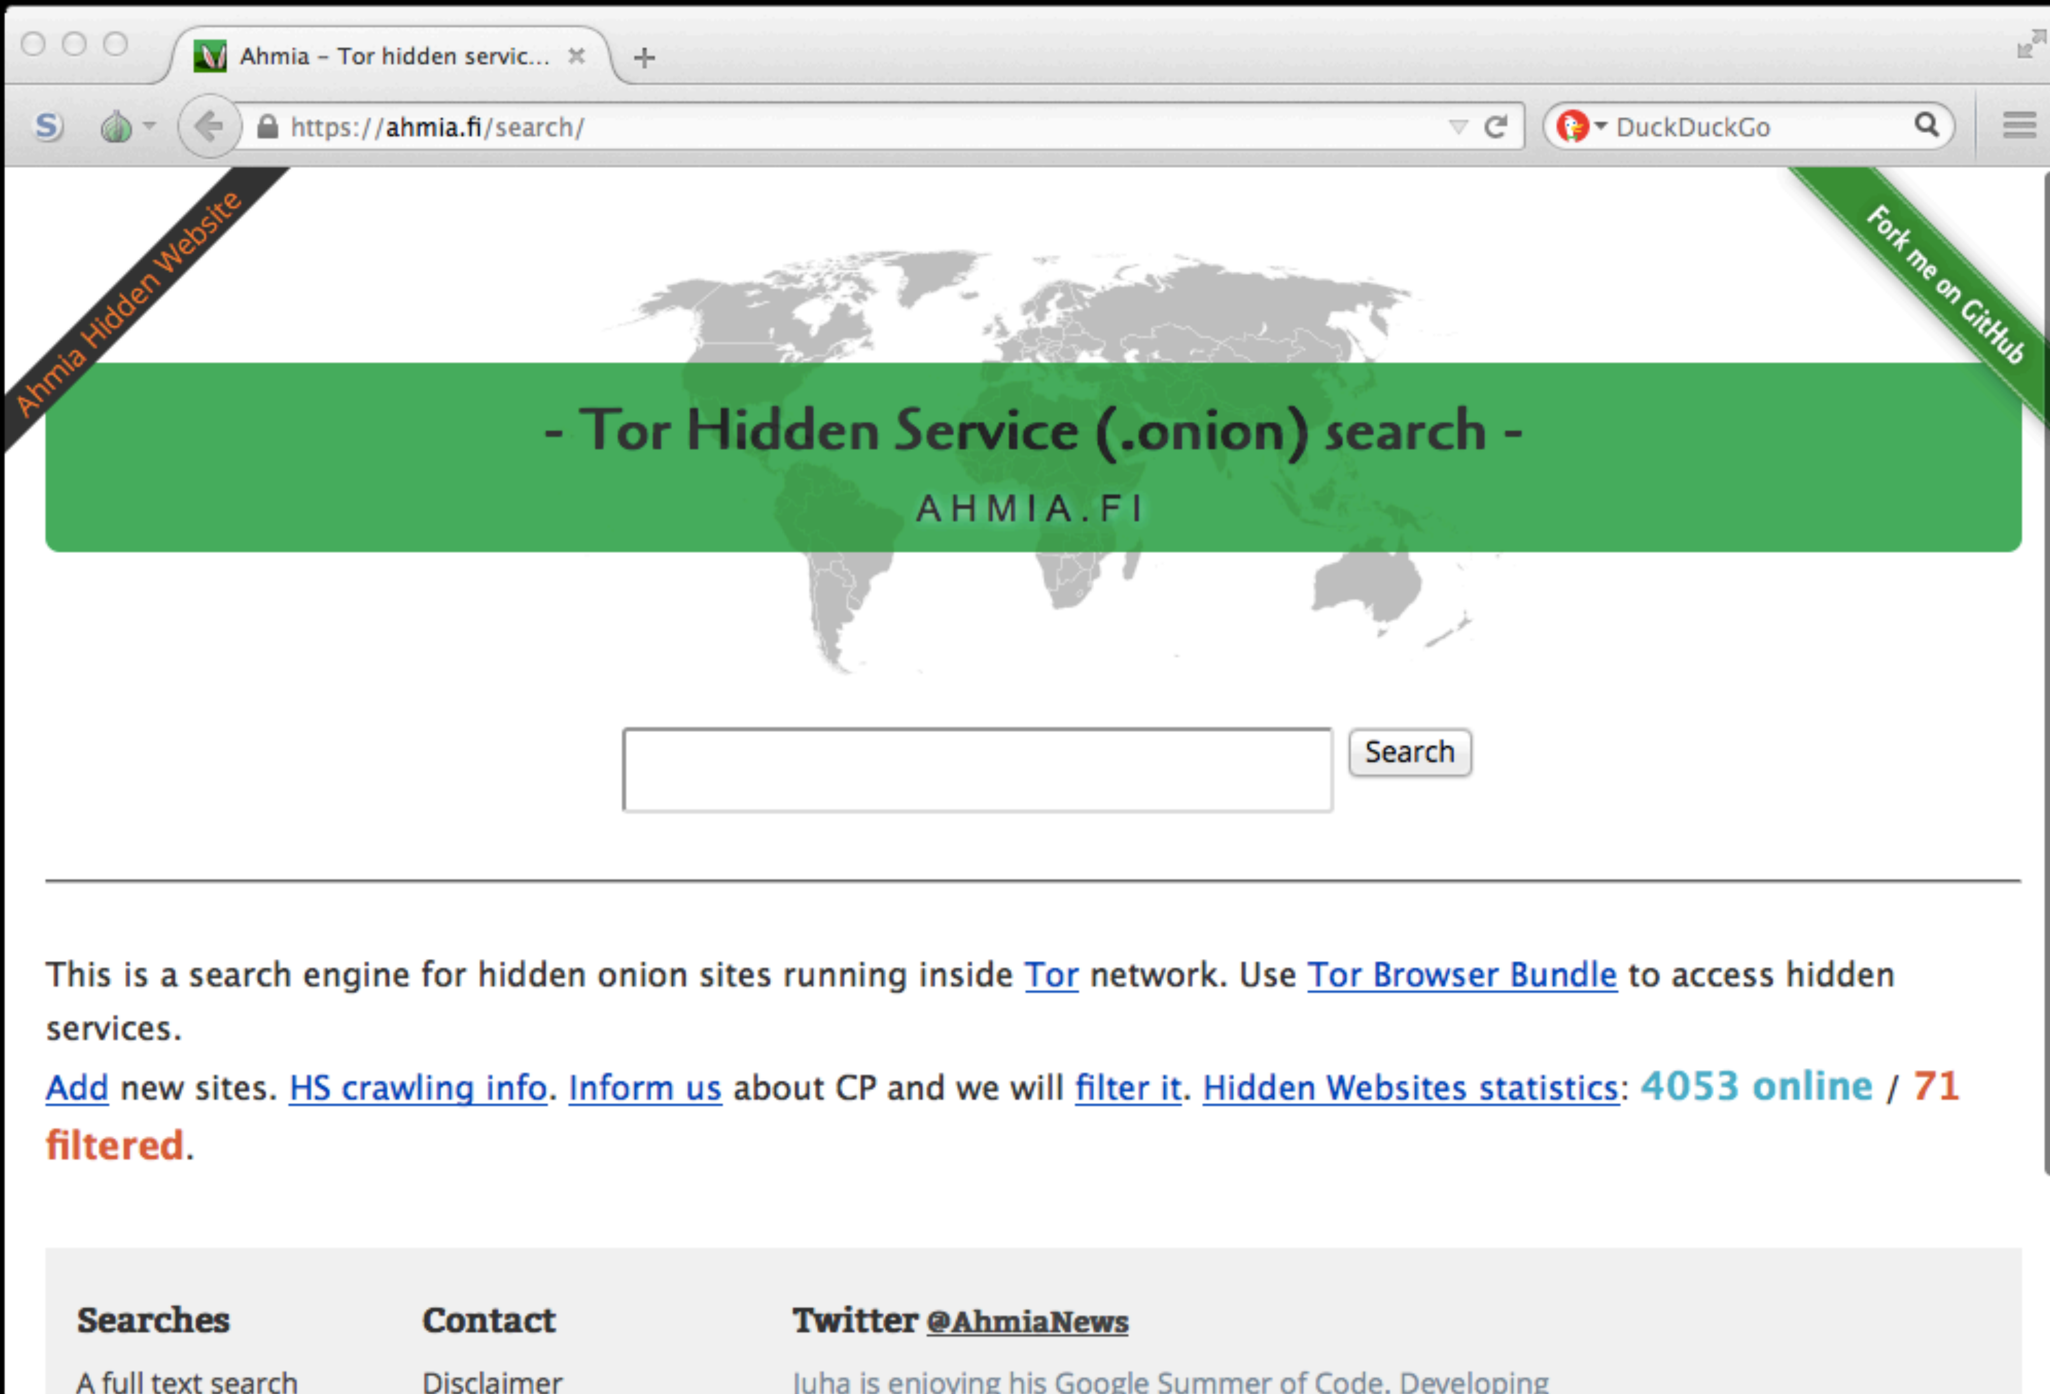
Ahmia - Tor hidden servic...
https://ahmia.fi/search/
DuckDuckGo
Ahmia Hidden Website
Fork me on GitHub
- Tor Hidden Service (.onion) search -
AHMIA.FI
Search
This is a search engine for hidden onion sites running inside Tor network. Use Tor Browser Bundle to access hidden services.
Add new sites. HS crawling info. Inform us about CP and we will filter it. Hidden Websites statistics: 4053 online / 71 filtered.
Searches
Contact
Twitter @AhmiaNews
A full text search
Disclaimer
Juha is enjoying his Google Summer of Code. Developing

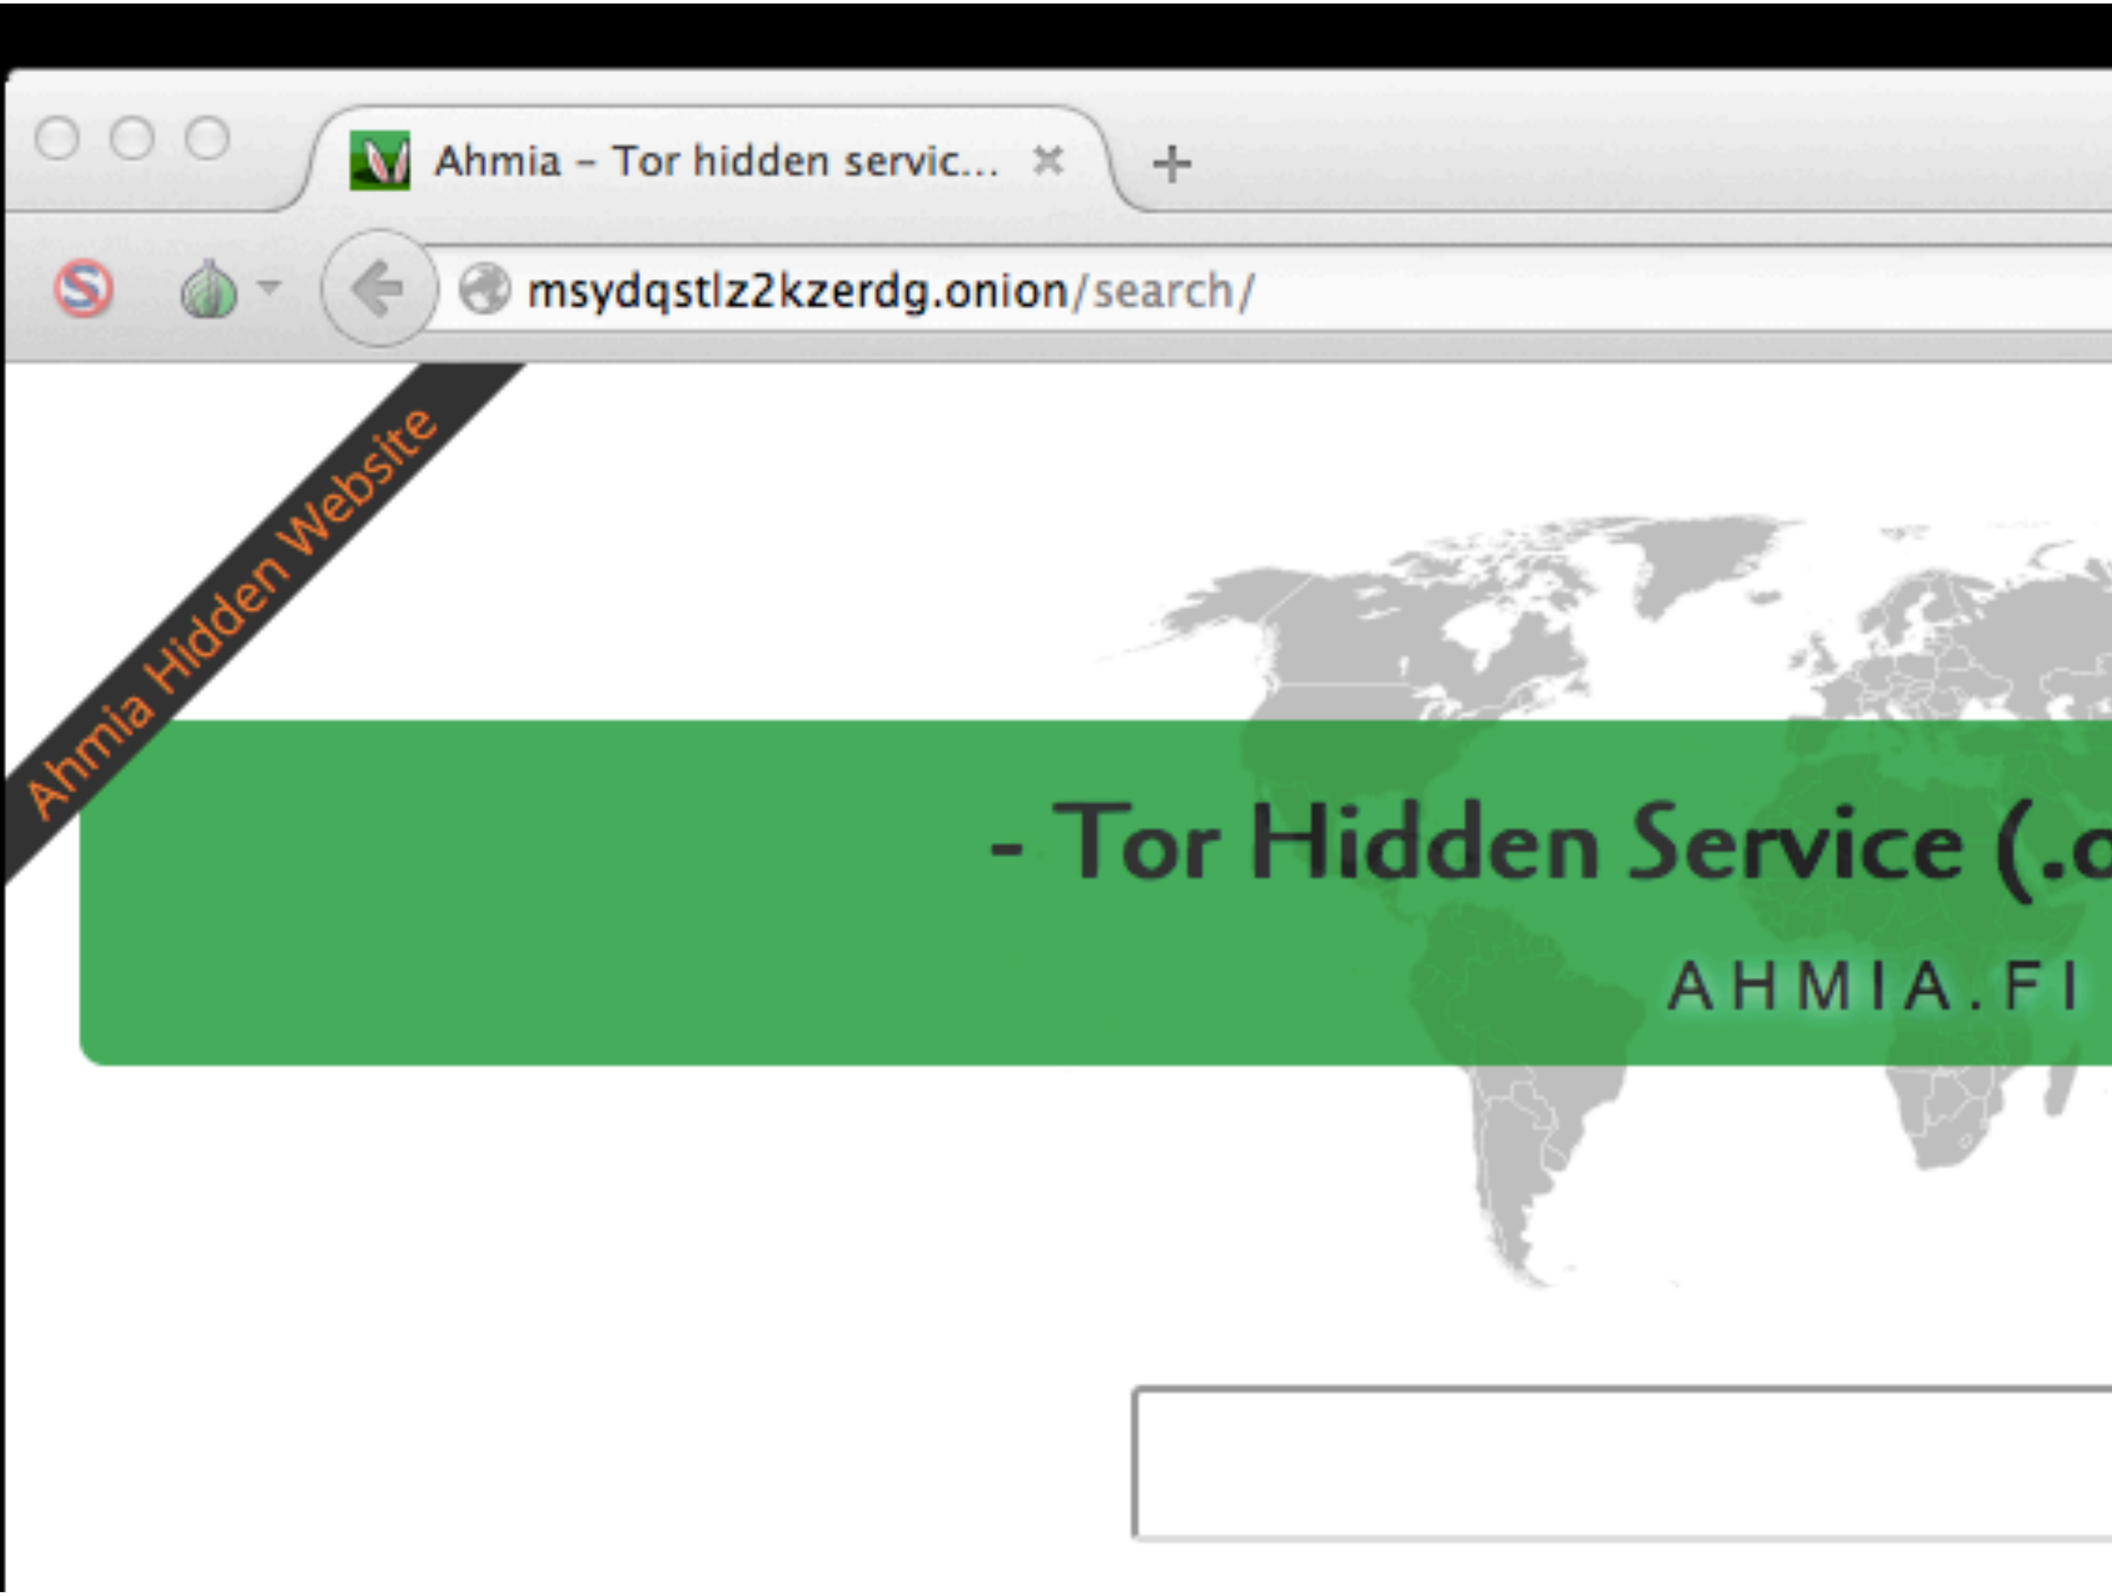
Ahmia - Tor hidden servic...
msydqstlz2kzerdg.onion/search/
Ahmia Hidden Website
- Tor Hidden Service (.o
AHMIA.FI

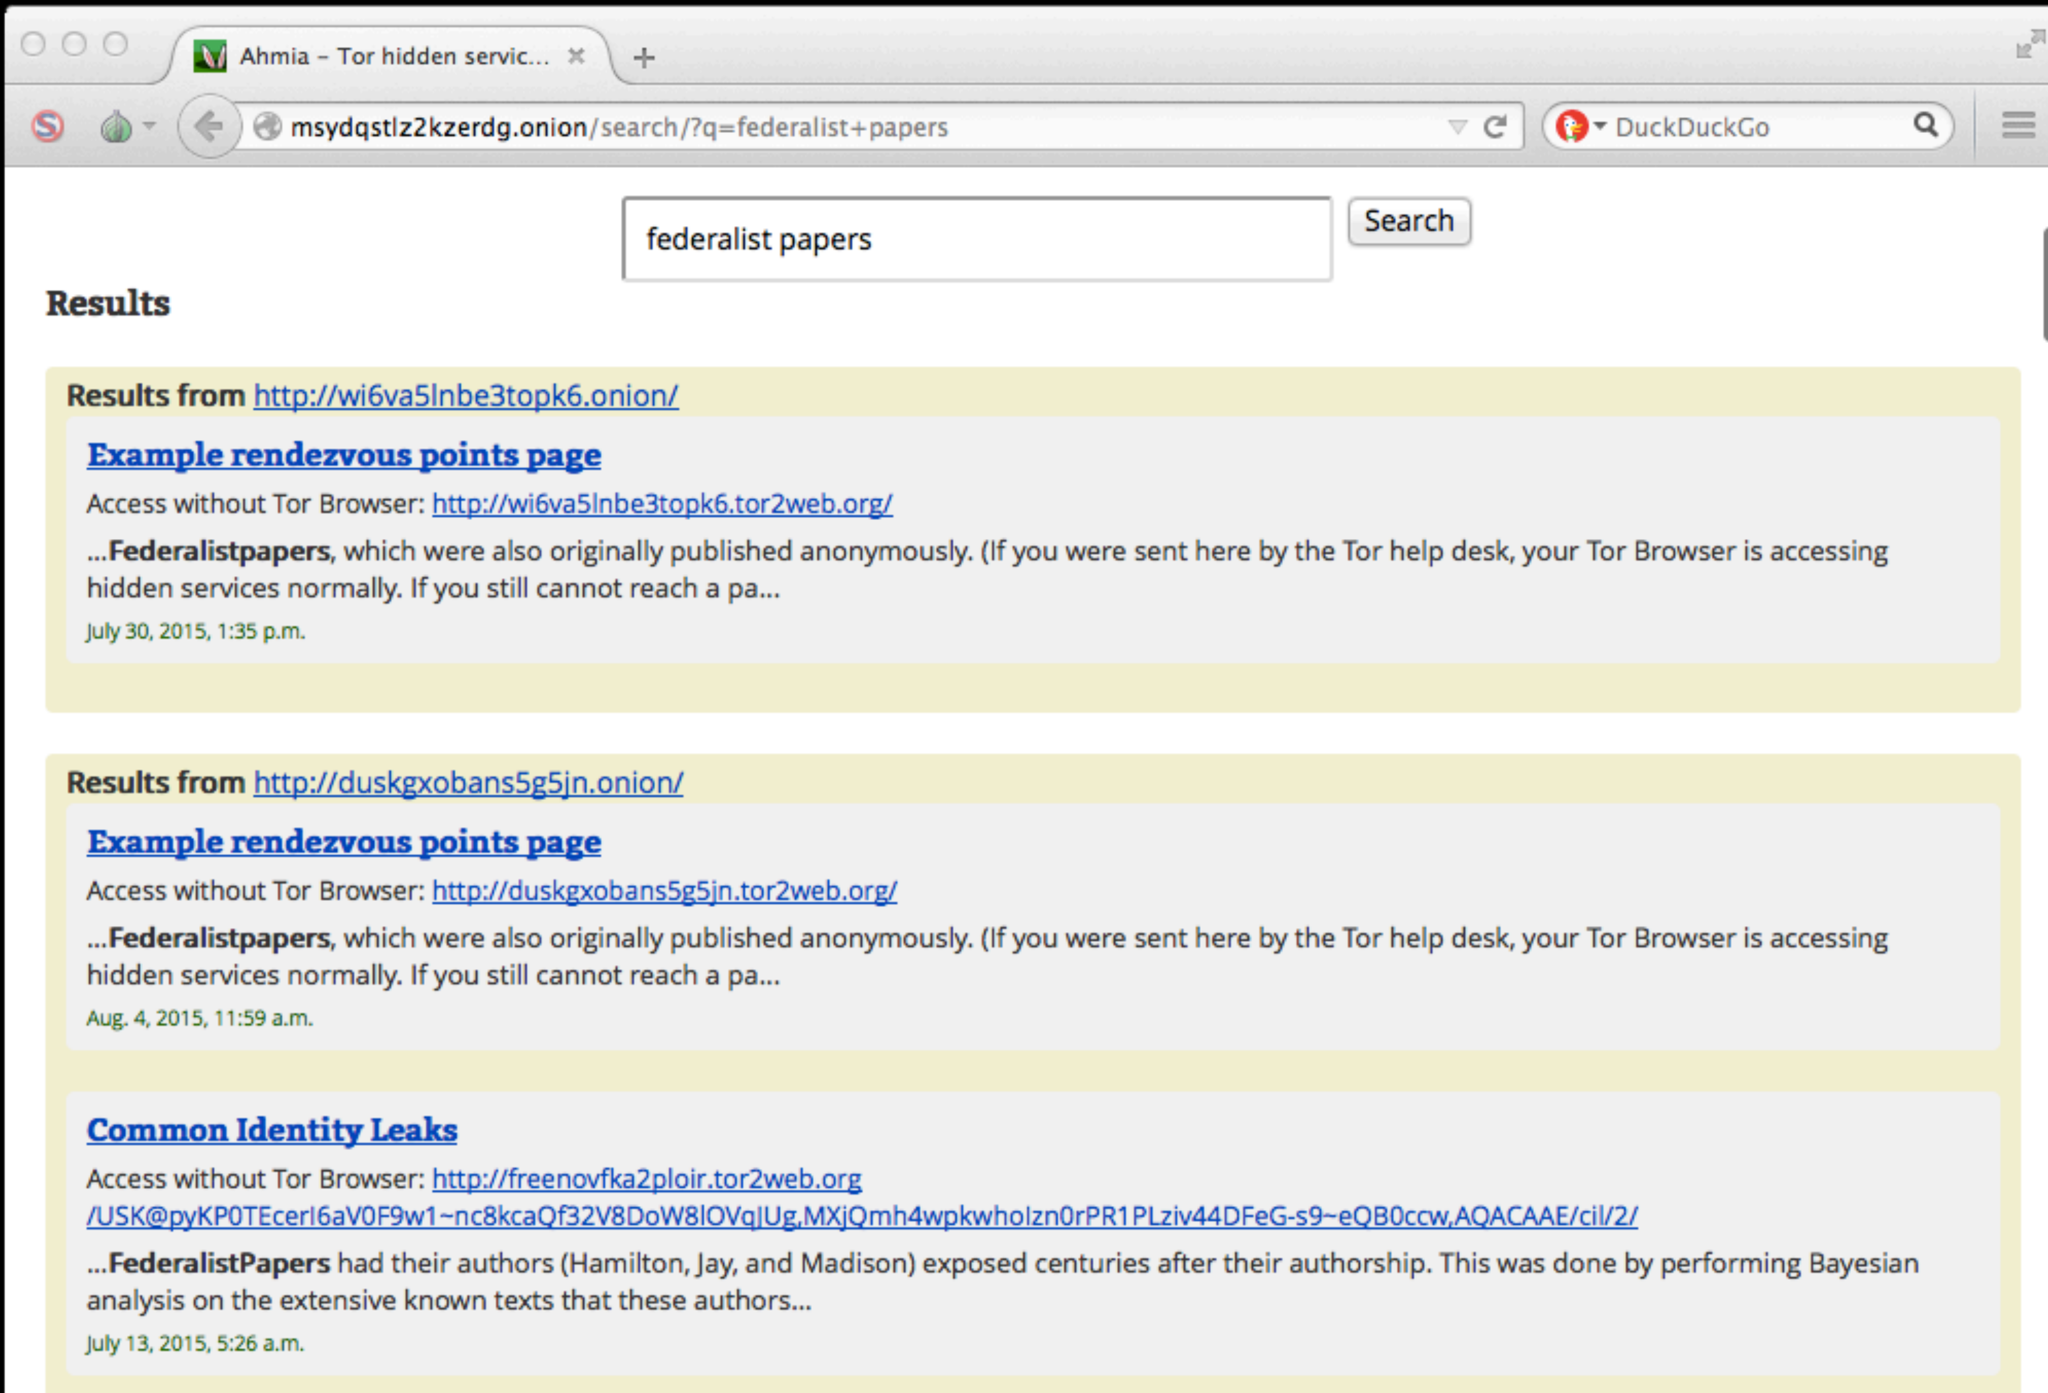

Ahmia - Tor hidden servic...
msydqstlz2kzerdg.onion/search/?q=federalist+papers
DuckDuckGo
federalist papers
Search
Results
Results from http://wi6va5lnbe3topk6.onion/
Example rendezvous points page
Access without Tor Browser: http://wi6va5lnbe3topk6.tor2web.org/
...Federalistpapers, which were also originally published anonymously. (If you were sent here by the Tor help desk, your Tor Browser is accessing hidden services normally. If you still cannot reach a pa...
July 30, 2015, 1:35 p.m.
Results from http://duskgxobans5g5jn.onion/
Example rendezvous points page
Access without Tor Browser: http://duskgxobans5g5jn.tor2web.org/
...Federalistpapers, which were also originally published anonymously. (If you were sent here by the Tor help desk, your Tor Browser is accessing hidden services normally. If you still cannot reach a pa...
Aug. 4, 2015, 11:59 a.m.
Common Identity Leaks
Access without Tor Browser: http://freenovfka2ploir.tor2web.org/USK@pyKP0TEcerl6aV0F9w1~nc8kcaQf32V8DoW8lOVqJUg,MXjQmh4wpkwholzn0rPR1PLziv44DFeG-s9~eQB0ccw,AQACAAE/cil/2/
...FederalistPapers had their authors (Hamilton, Jay, and Madison) exposed centuries after their authorship. This was done by performing Bayesian analysis on the extensive known texts that these authors...
July 13, 2015, 5:26 a.m.

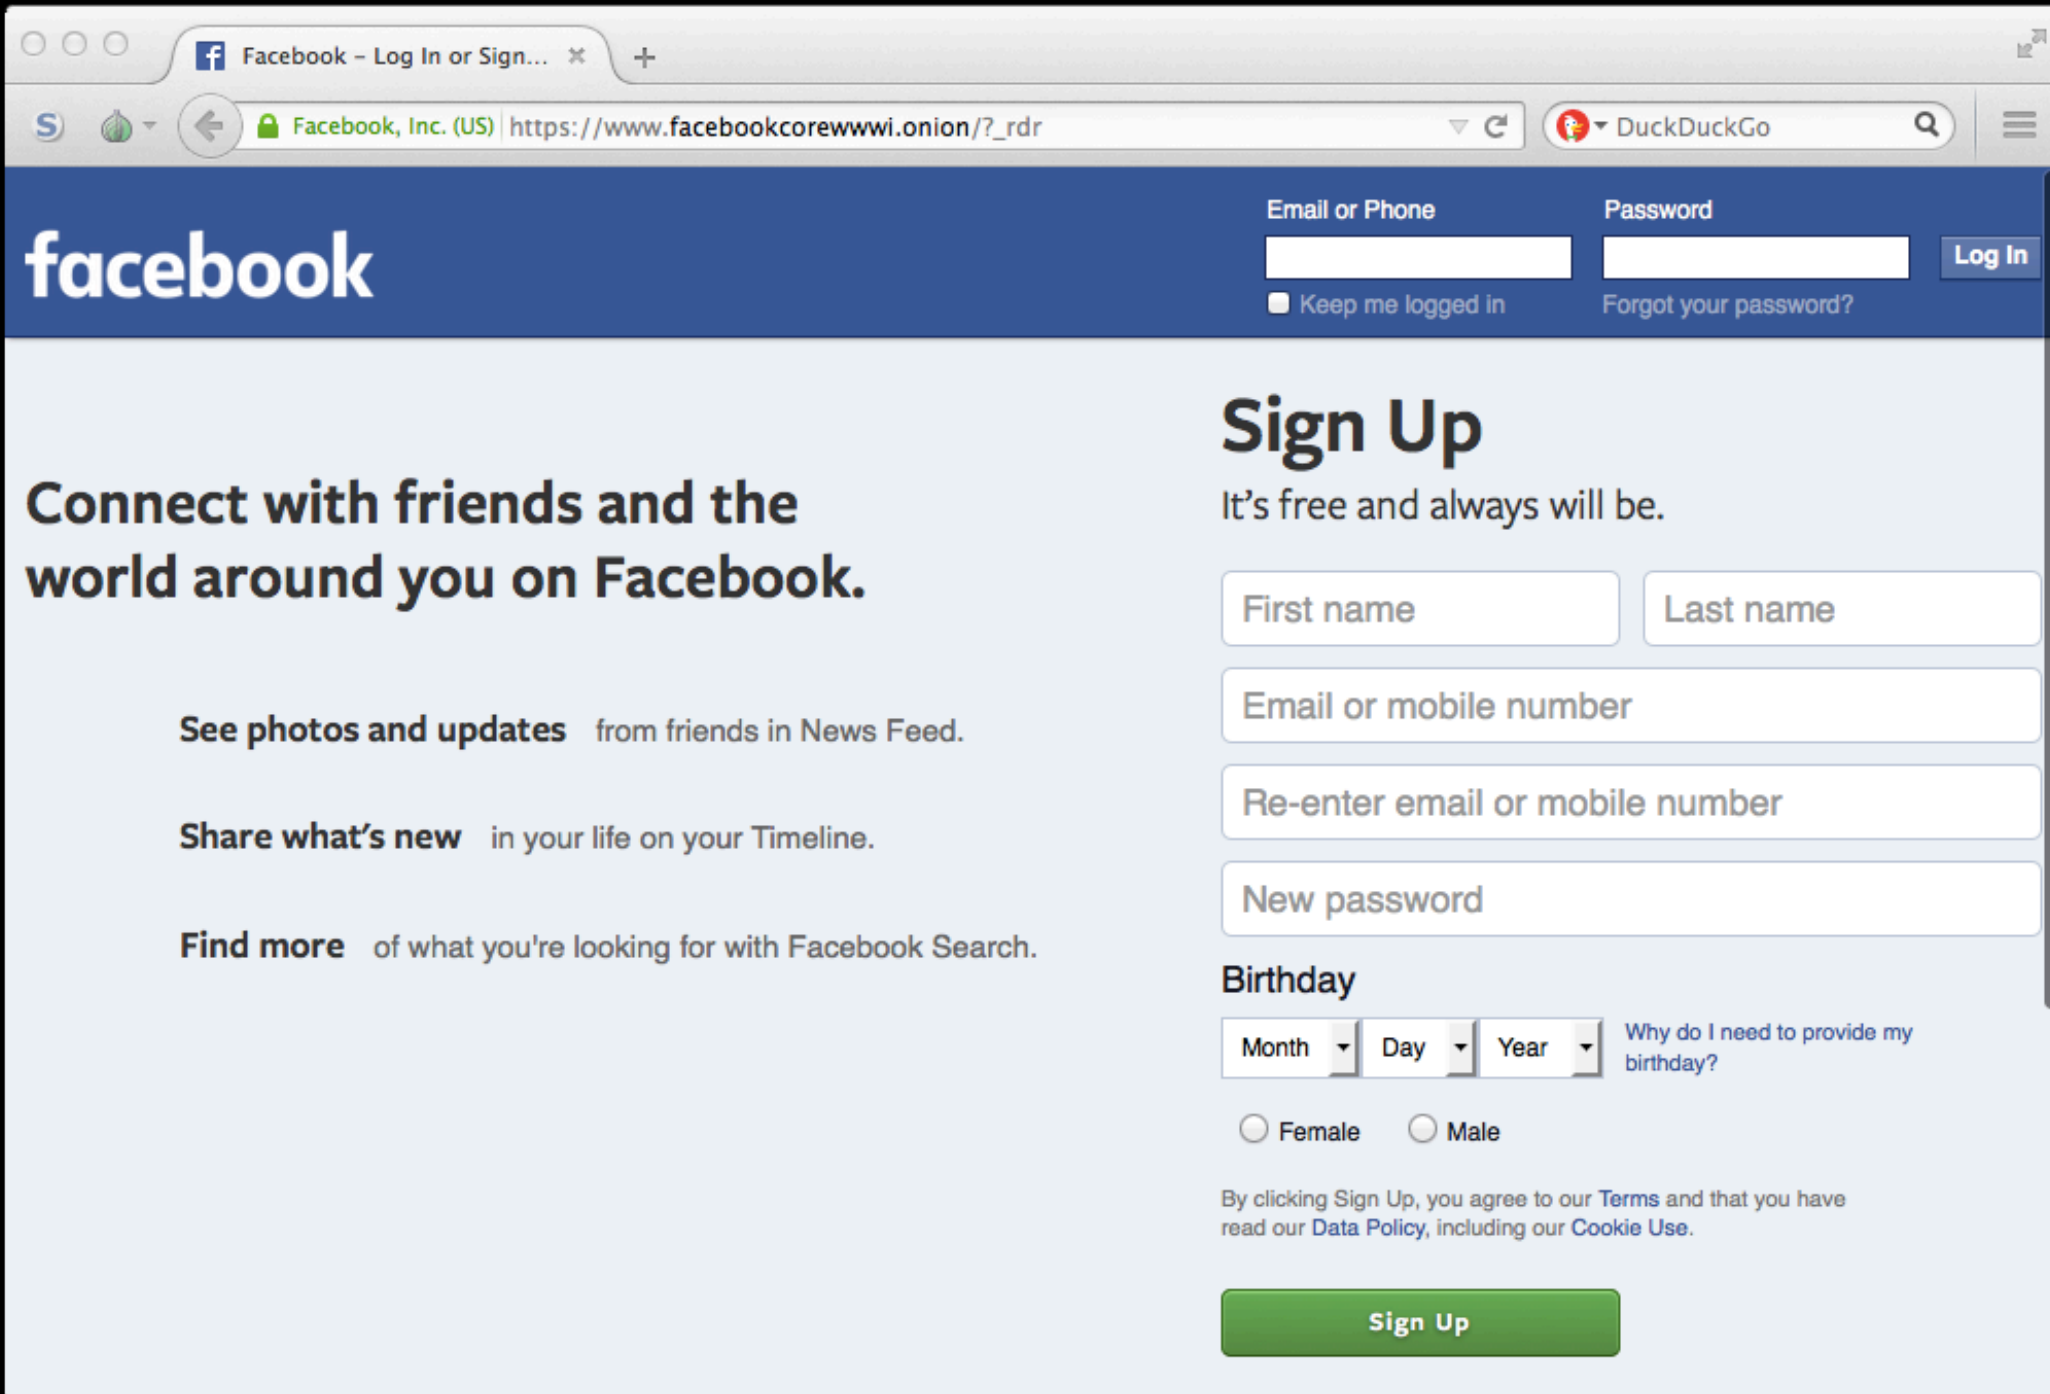
Facebook – Log In or Sign...
Facebook, Inc. (US) https://www.facebookcorewwwi.onion/?_rdr
DuckDuckGo
facebook
Email or Phone
Password
Log In
Keep me logged in
Forgot your password?
Connect with friends and the world around you on Facebook.
See photos and updates from friends in News Feed.
Share what's new in your life on your Timeline.
Find more of what you're looking for with Facebook Search.
Sign Up
It's free and always will be.
First name
Last name
Email or mobile number
Re-enter email or mobile number
New password
Birthday
Month
Day
Year
Why do I need to provide my birthday?
Female
Male
By clicking Sign Up, you agree to our Terms and that you have read our Data Policy, including our Cookie Use.
Sign Up

End: obtaining/using Tor

- Lecture 2: Security for the real world
  - Simple demo of obtaining/using Tor
  - Security of obtain/using Tor
  - Adding network link awareness
  - Importance of modeling users
  - Importance of realistic and practical
    - Adversary models
    - Security definitions

Adversary observing all traffic entering and leaving network breaks onion routing

- “Towards an Analysis of Onion Routing Security” Syverson et al. PETS 2000
- Presented and analyzed adversary model assumed in prior onion routing work
  - Network of n onion routers, c compromised onion routers
  - Security approx. c^2 / n^2

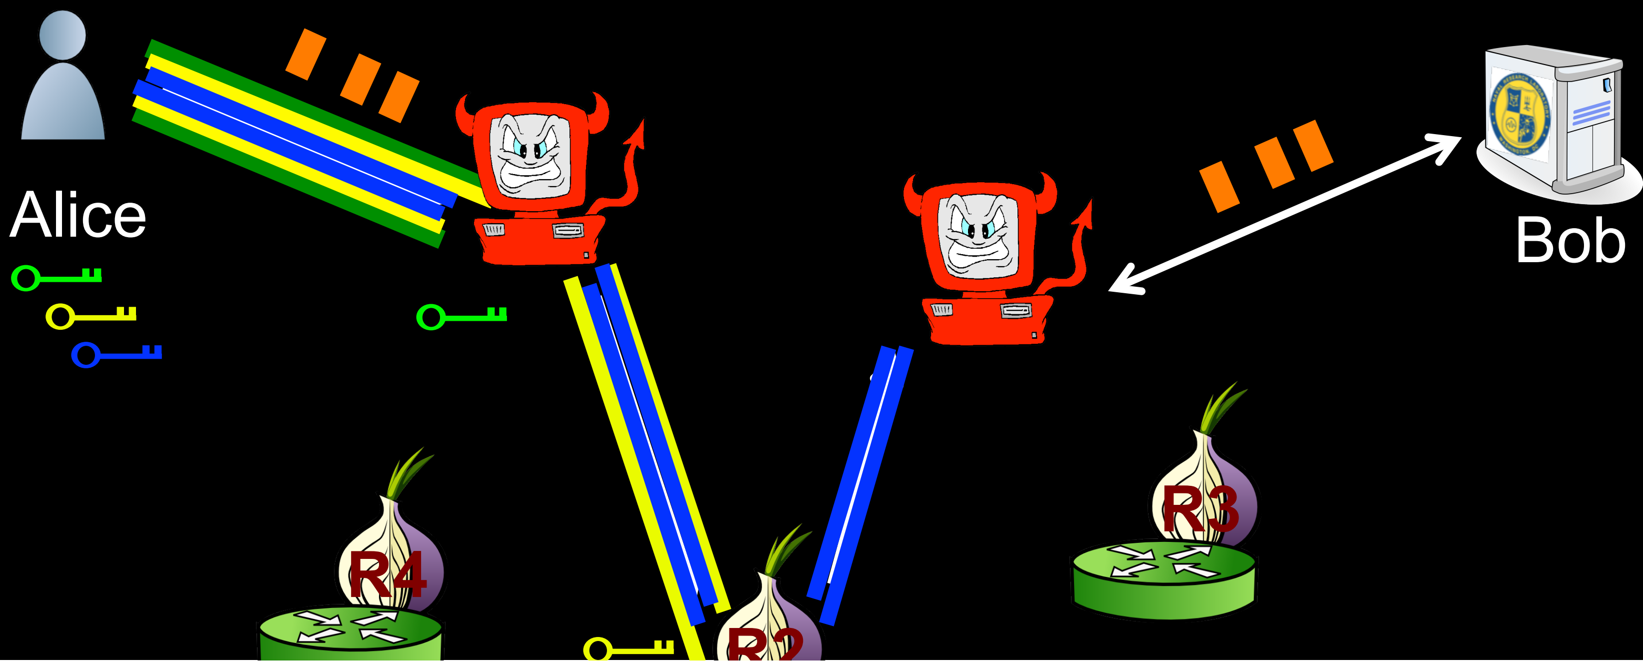

Adversary observing all traffic entering and leaving network?

- “Location diversity in anonymity networks” Feamster-Dingledine. WPES 2004
- Adversaries live on network links as well as onion routers

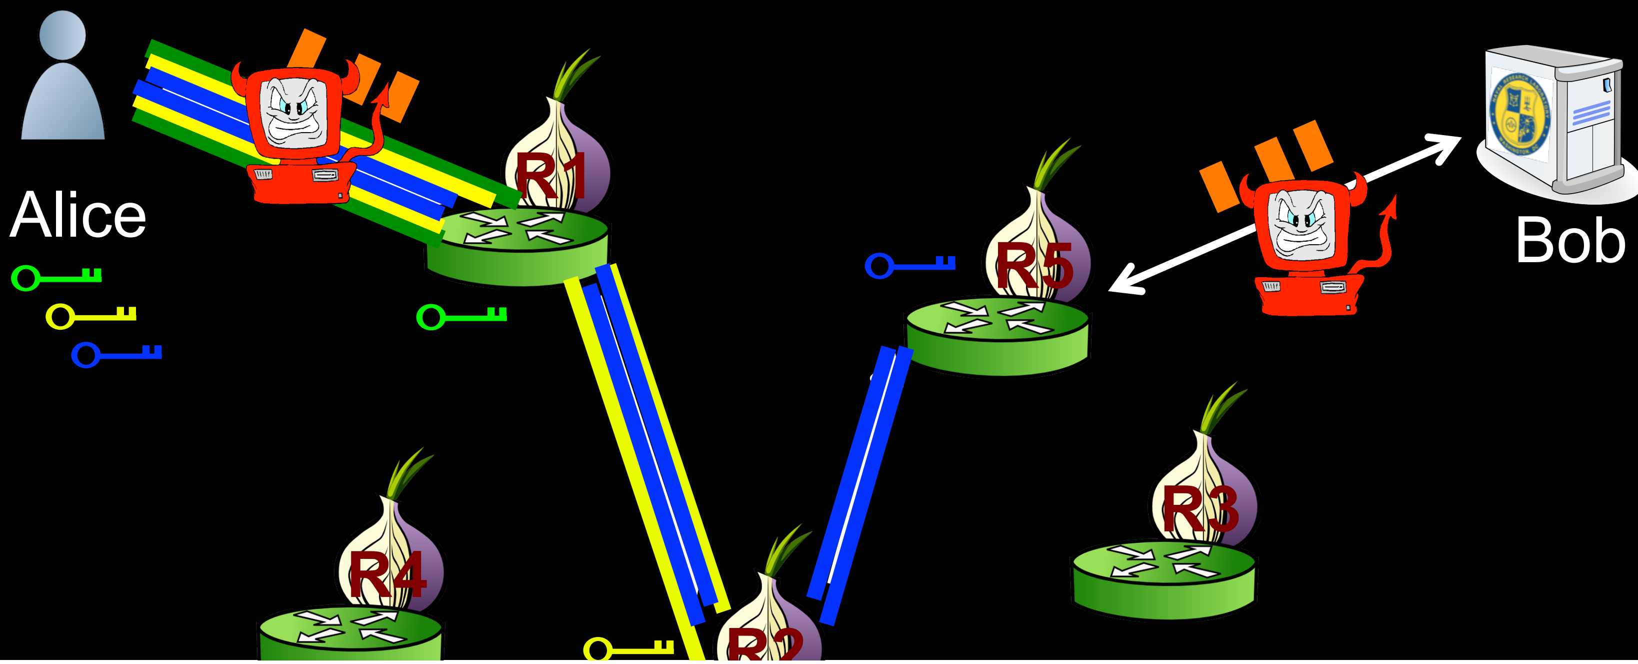

Onion Routers (Tor Relays) overlay underlying Internet

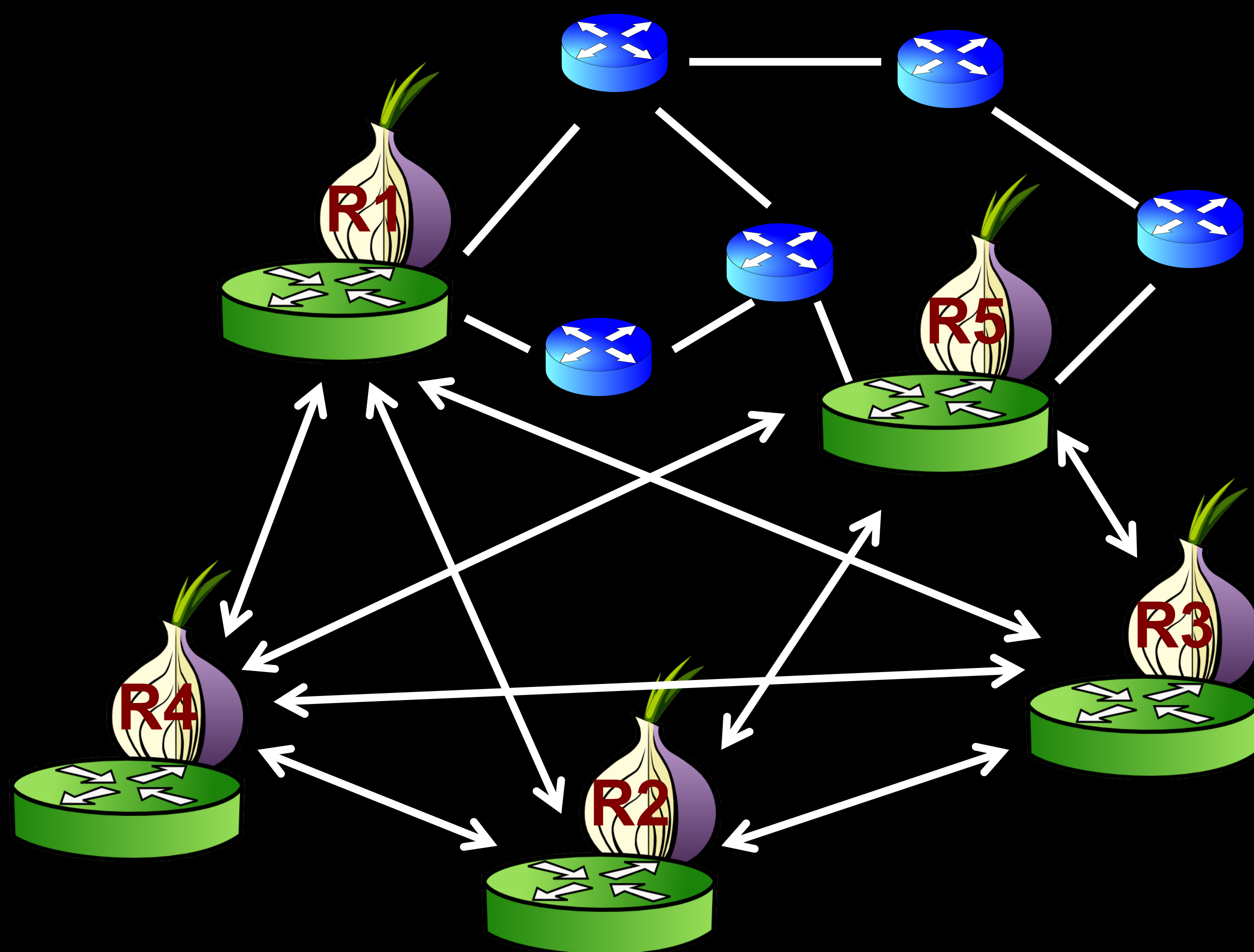

Adversaries can live on network links to/from onion routers too

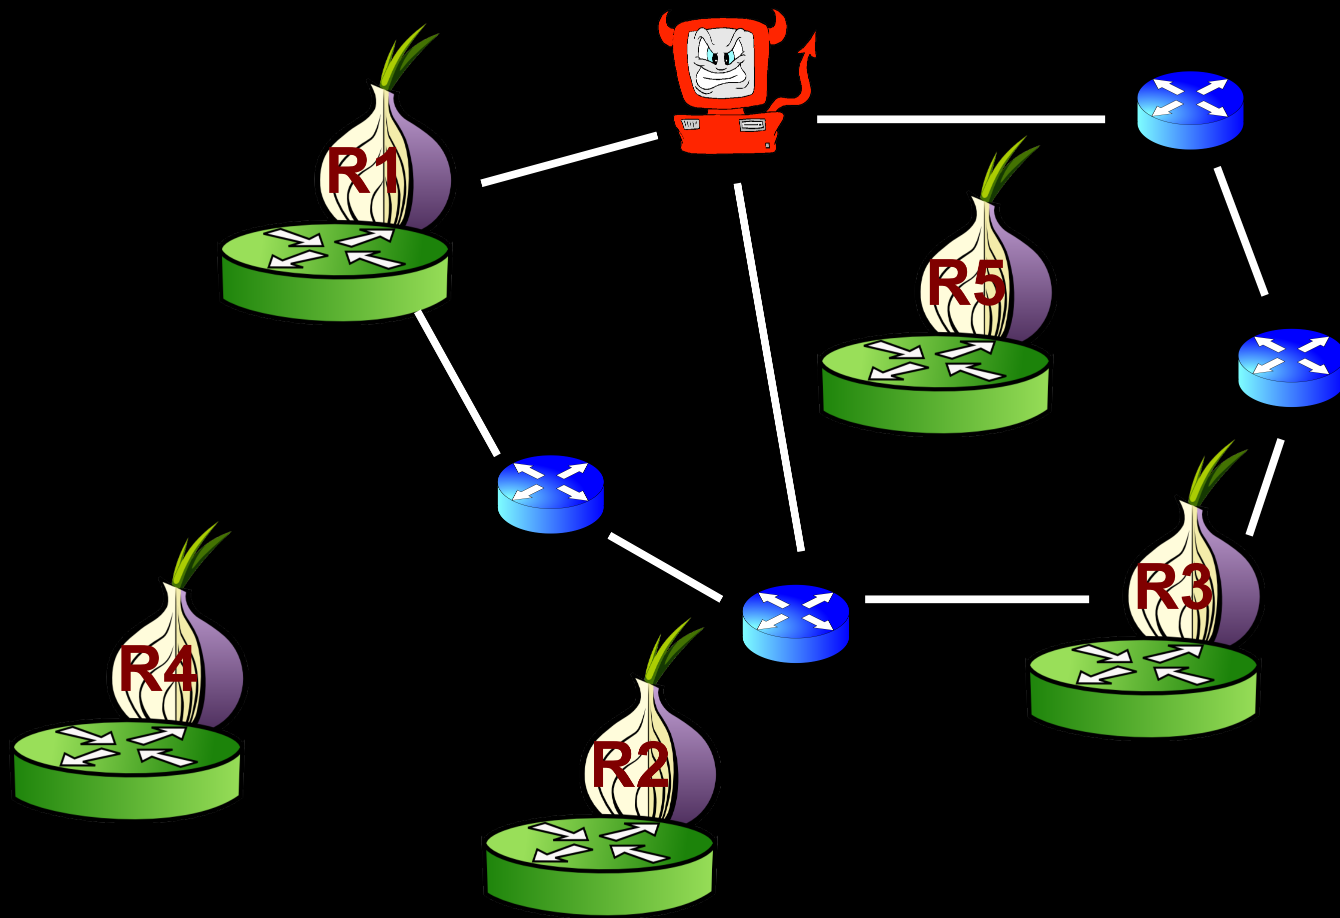

Adversaries can live on network links to/from onion routers too

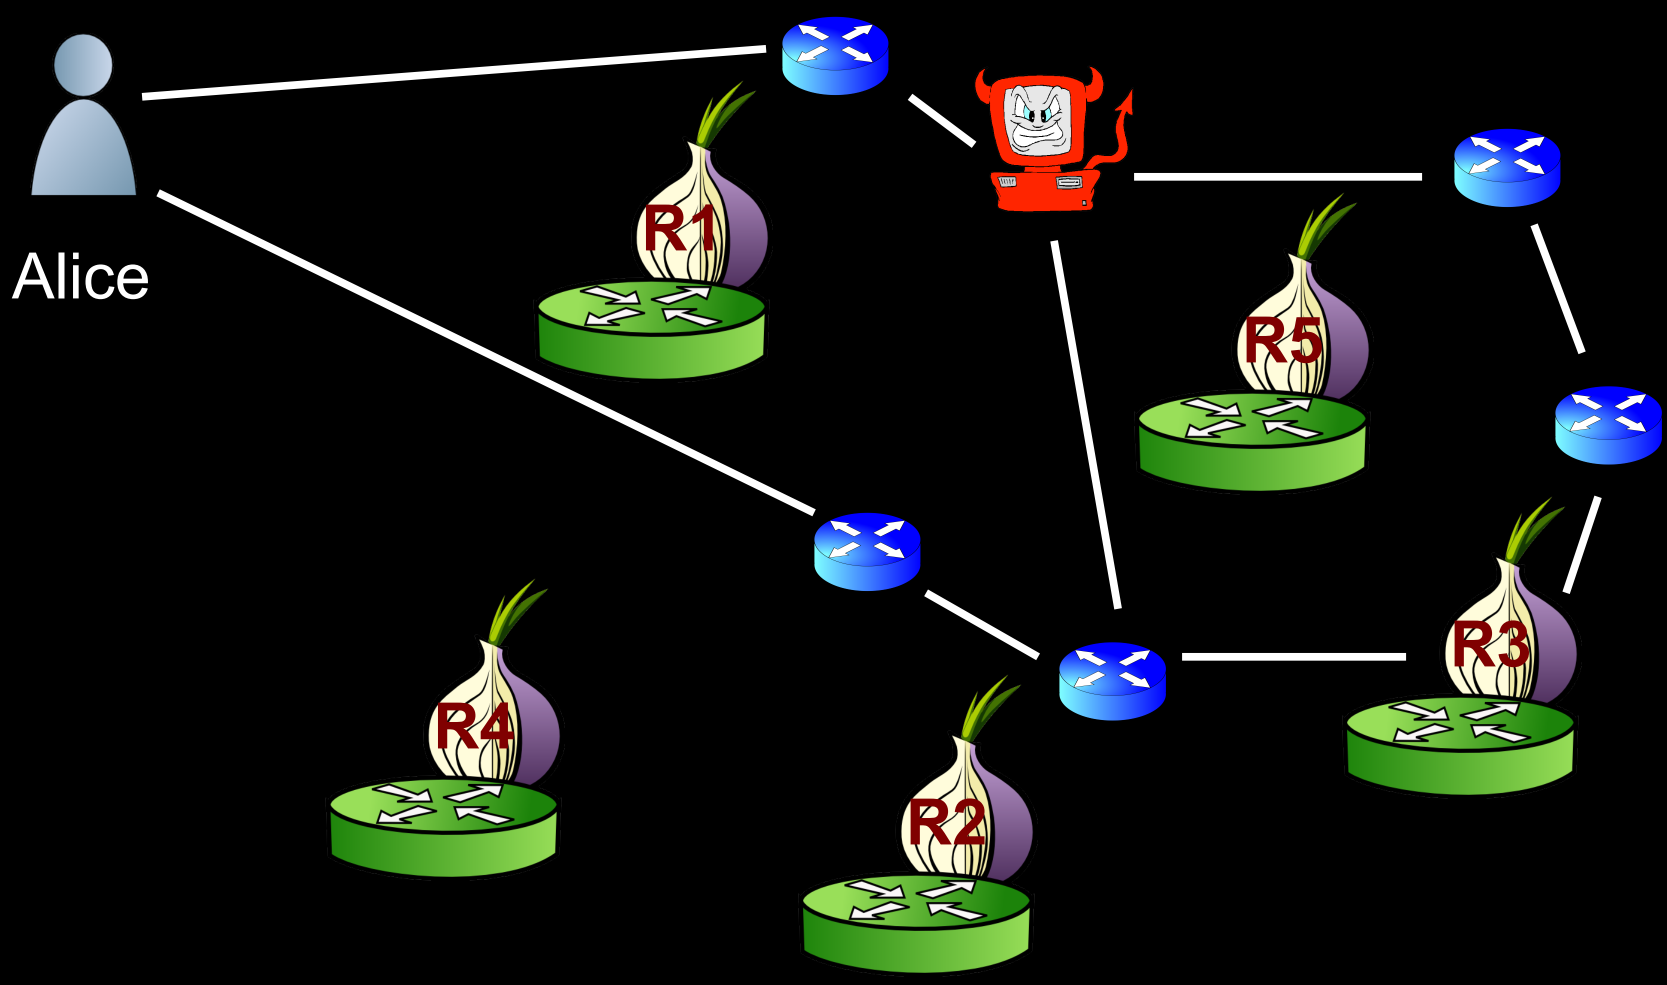

Link Adversary

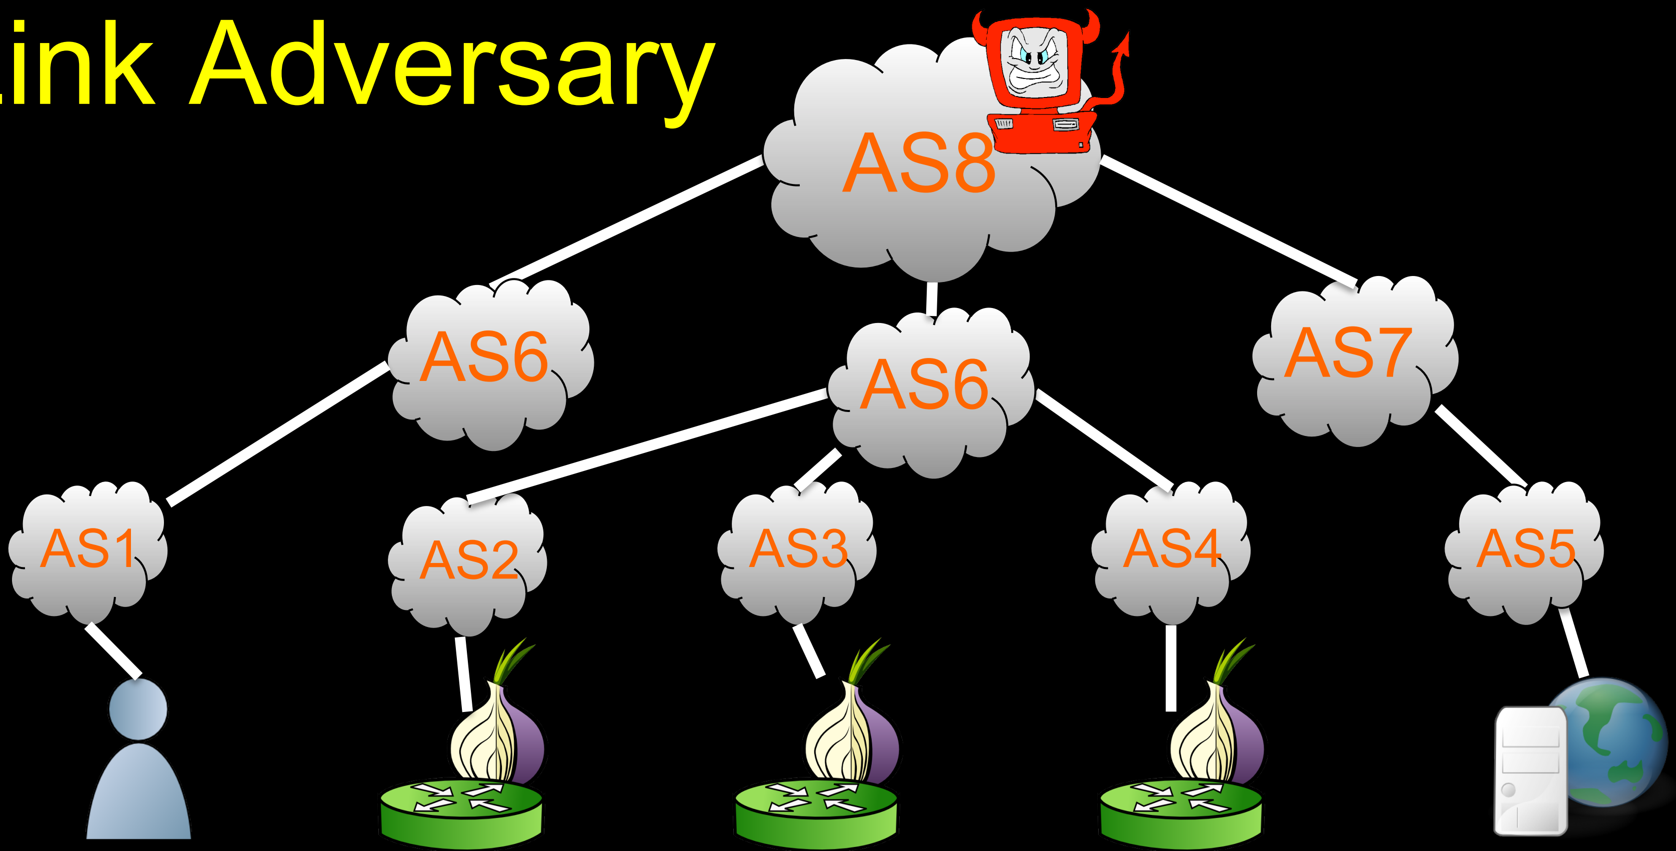

1. Autonomous Systems (ASes)

Adversary observing all traffic entering and leaving network breaks onion routing

- Recall Onion Routing security approach: Large, Diverse Network so adversary has to expend much resources in many places

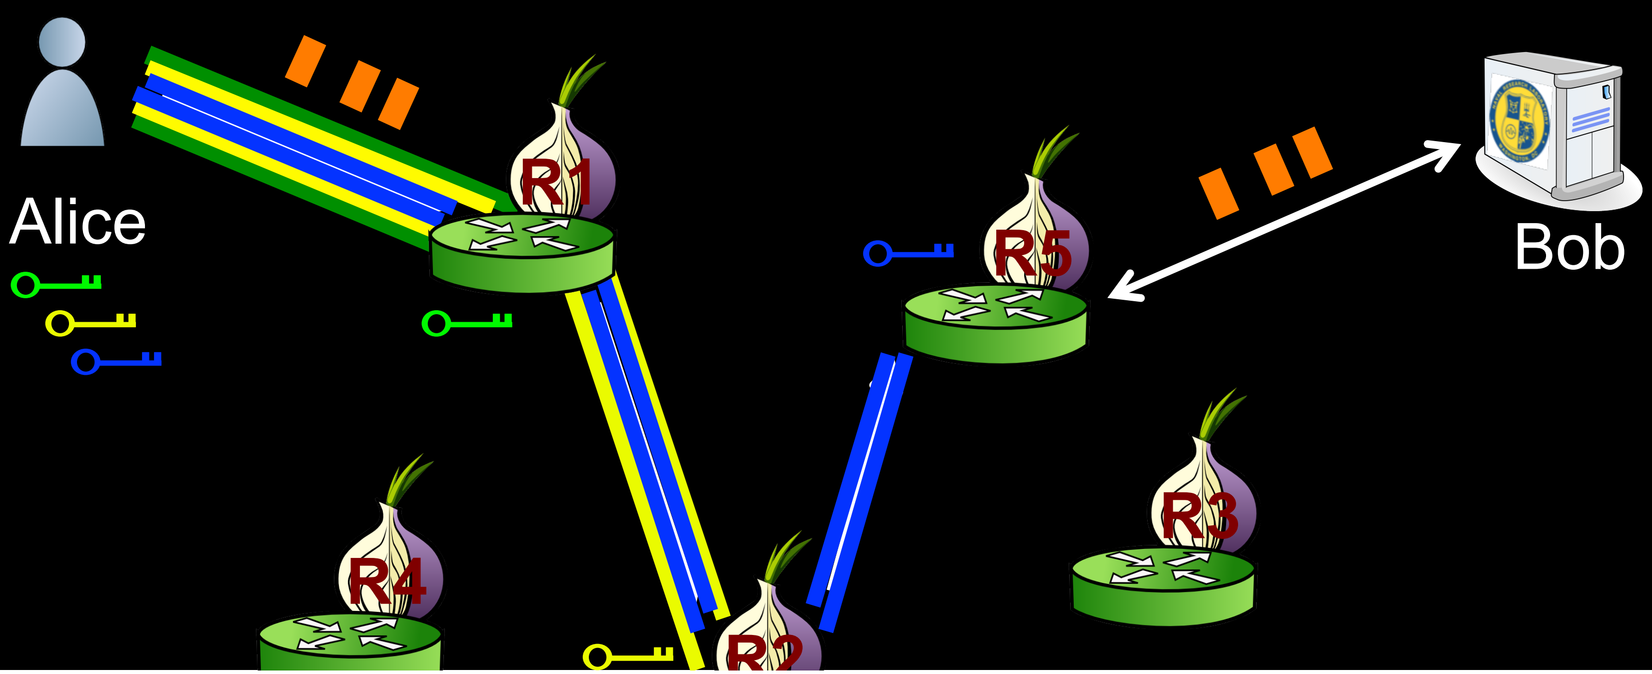

Adversaries can live on network links to/ from onion routers too

- "Location diversity in anonymity networks" Feamster-Dingledine. WPES 2004
- Model adversaries at Autonomous Systems (ASes)
  - Path Independence: No AS on both client and destination end of circuit

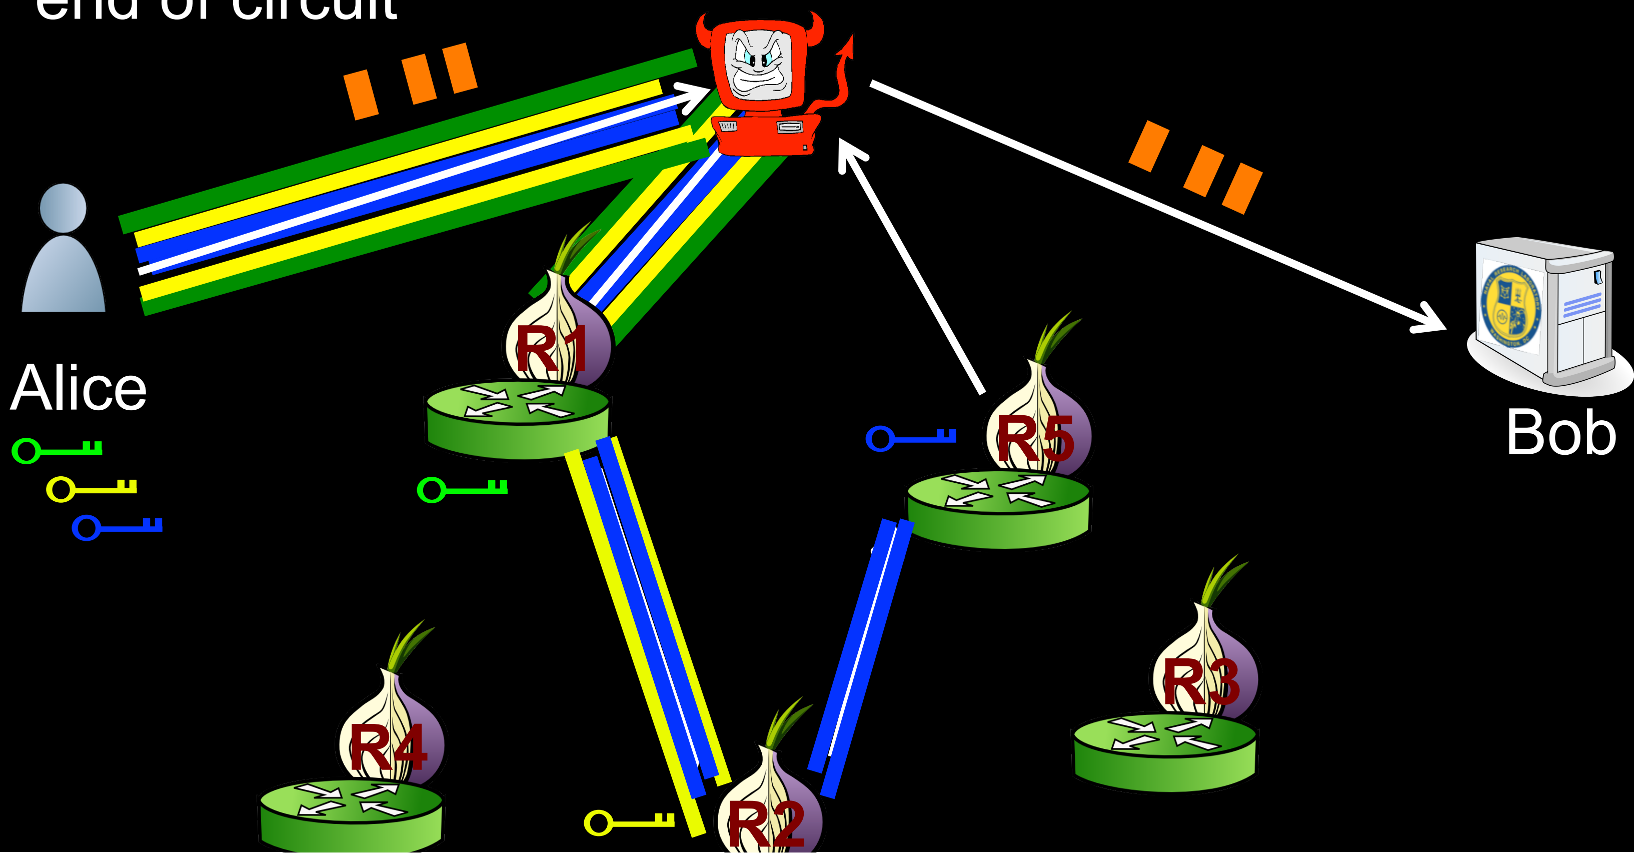

Adversaries can live on network links to/ from onion routers too

- "Location diversity in anonymity networks" Feamster-Dingledine. WPES 2004
- Model adversaries at Autonomous Systems (ASes)
  - Path Independence: No AS is on both client and destination end of circuit

- How bad is it?
- What can we do?

First pass look at link attacks

- “AS-awareness in Tor Path Selection” Edman-Syverson. CCS 2009
- Background
- **AS Path Inference**
- Analysis of Tor network growth
- Tor AS statistics
- Proposed path selection heuristics

AS Path Inference

- Tries to predict route packets will take on the Internet
- We do not have access to routing tables for the entire Internet
- We cannot traceroute from arbitrary hosts
- AS relationships are not often publicized for contractual reasons

AS Path Inference

Deriving AS Paths from Known Paths (Qiu & Gao 2006)

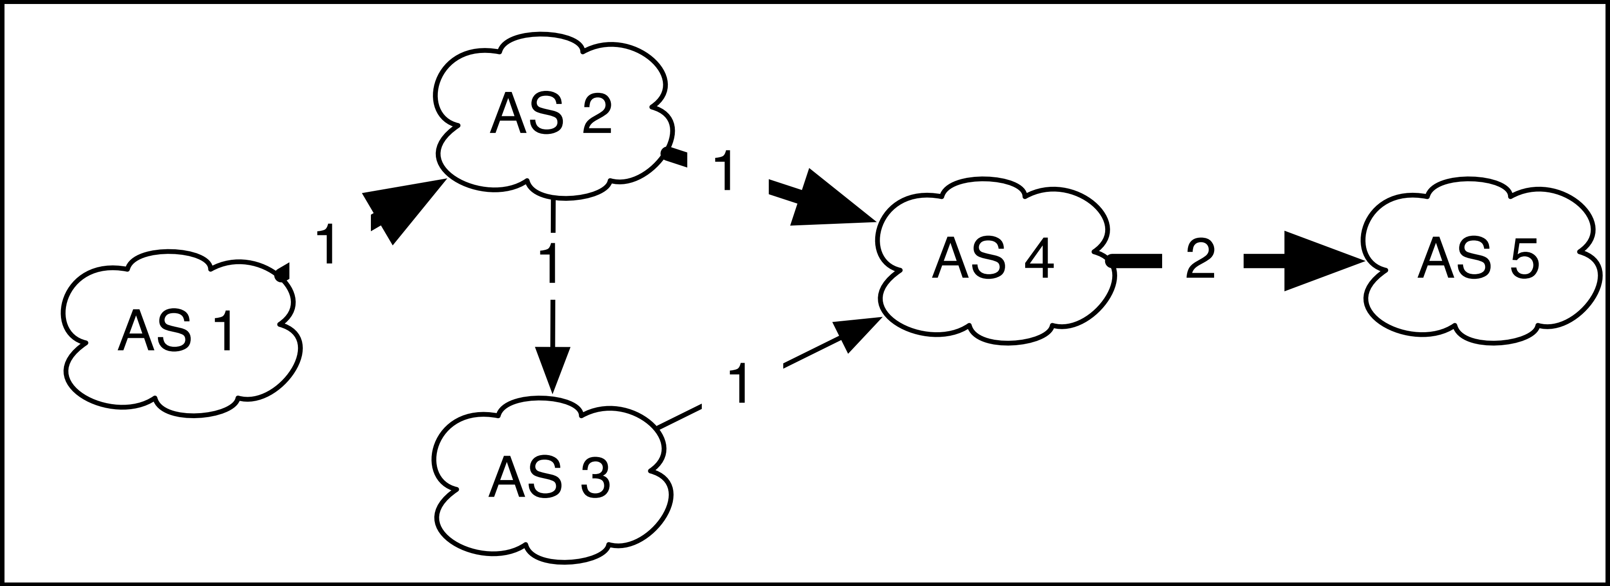

{1,2,3}, {2,4,5} and {3,4,5} are *known* paths

{1,2,4,5} is a *derived* path (must satisfy *valley-free* property)

AS Path Inference

- Used input routing tables from multiple Internet vantage points
  - OIX, Equinix, PAIX, KIXP, LINX, DIXIE
  - 1.47 GB, 15.7 million paths, 29,000 ASes, 132,000 edges
- Implementation
  - Implemented in C
  - Used Gao's (2000) algorithm for relationship inference
  - Modified slightly for better parallelization
  - All experiments done on a commodity Dell workstation

First pass look at link attacks

- Background
- AS Path Inference
- **Analysis of Tor network growth**
- Tor AS statistics
- Proposed path selection heuristics
- Conclusions & future work

Tor Grows Up

	June 2004 (33 relays)						**September 2008** (1239–1303 relays)					
Sender	2914	11643	12182	15130	15169	26101	2914	11643	12182	15130	15169	26101
209	0.49	0.45	0.40	0.39	0.19	0.30	0.17	0.26	0.19	0.51	0.23	0.25
1668	0.39	0.24	0.30	0.30	0.19	0.32	0.18	0.23	0.20	0.25	0.13	0.16
4355	0.38	0.27	0.28	0.27	0.43	0.51	0.13	0.29	0.12	0.20	0.19	0.14
6079	0.62	0.45	0.48	0.24	0.43	0.71	0.12	0.30	0.15	0.22	0.20	0.17
18566	0.39	0.42	0.41	0.32	0.56	0.73	0.18	0.36	0.20	0.31	0.20	0.16
22773	0.56	0.35	0.37	0.21	0.34	0.54	0.21	0.14	0.20	0.20	0.17	0.19
22909	0.21	0.24	0.26	0.22	0.22	0.37	0.19	0.30	0.24	0.25	0.21	0.19
23504	0.39	0.29	0.37	0.33	0.42	0.54	0.49	0.22	0.23	0.19	0.16	0.12

- Used 3 separate Tor consensus snapshots from September 2008
- Mean overall probability of an AS-level observer decreased from 37.74% to 21.86%
-

Tor Grows Up

	June 2004 (33 relays)						September 2008 (1239–1303 relays)					
Sender	2914	11643	12182	15130	15169	26101	2914	11643	12182	15130	15169	26101
209	0.49	0.45	0.40	0.39	0.19	0.30	0.17	0.26	0.19	0.51	0.23	0.25
1668	0.39	0.24	0.30	0.30	0.19	0.32	0.18	0.23	0.20	0.25	0.13	0.16
4355	0.38	0.27	0.28	0.27	0.43	0.51	0.13	0.29	0.12	0.20	0.19	0.14
6079	0.62	0.45	0.48	0.24	0.43	0.71	0.12	0.30	0.15	0.22	0.20	0.17
18566	0.39	0.42	0.41	0.32	0.56	0.73	0.18	0.36	0.20	0.31	0.20	0.16
22773	0.56	0.35	0.37	0.21	0.34	0.54	0.21	0.14	0.20	0.20	0.17	0.19
22909	0.21	0.24	0.26	0.22	0.22	0.37	0.19	0.30	0.24	0.25	0.21	0.19
23504	0.39	0.29	0.37	0.33	0.42	0.54	0.49	0.22	0.23	0.19	0.16	0.12

- Used 3 separate Tor consensus snapshots from September 2008
- Mean overall probability of an AS-level observer decreased from 37.74% to 21.86%
- ≈12.5% AS pairs were worse off than before

First pass look at link attacks

- Background
- AS Path Inference
- Analysis of Tor network growth
- **Tor AS statistics**
- Proposed path selection heuristics

Tor AS Distribution Model

- Data Collection
- Ran two relays for 7 days in early September 2008
- Mapped client and destination IP addresses to AS numbers
- Kept only aggregated statistics at AS level
  - *Never wrote IP addresses, timestamps or other metadata to disk*

Tor AS Distribution Model

- Results
- 20638 client connections
  - 2251 distinct ASes
  - 85% produced fewer than 10 connections
  - >50% produced only a single connection
- 116781 destination connections
  - 4203 distinct ASes
  - 72% produced fewer than 10 connections
  - 34% had only a single connection

Aside Moral: Privacy preserving statistics gathering is hard

Tor Client AS Distribution

Rank	#	CC	Description
1	2238	DE	Deutsche Telekom AG
2	701	CN	ChinaNet
3	672	EU	Arcor
4	576	IT	Telecom Italia
5	566	DE	HanseNet Telekommunikation
6	429	DE	Telefonia Deutschland
7	280	FR	Proxad
8	279	US	AT&T Internet Services
9	276	CN	CNC Group Backbone
10	272	TR	TTNet

Tor Destination AS Distribution

Rank	#	CC	Description
1	5203	CN	ChinaNet
2	**4960**	**US**	**Google Inc.**
3	3527	NL	NForce Entertainment
4	2824	TW	HiNet
5	**2085**	**US**	**AOL**
6	2029	US	ThePlanet.com
7	1530	CN	CNC Group Backbone
8	1104	CN	CNC Group Beijing Province
9	1083	US	Level3 Communications
10	1011	NL	LeaseWeb

First pass look at link attacks

- Background
- AS Path Inference
- Analysis of Tor network growth
- Tor AS statistics
- **AS-aware path selection algorithms**

Tor Path Selection Changes over time

- Weighted node selection
  - Relay bandwidth
  - Uptime
- Entry guards (motivation in c. 10 more slides)
- Distinct /16 subnets

Tor Path Selection Changes

- Effectiveness of Distinct /16 Subnets
  - Using mid-September 2008 Tor consensus
    - 876/1238 (≈70%) relays in same AS as at least one other relay, but in distinct /16 subnets
    - 850/1238 (≈68.7%) in same AS but distinct /8 subnet
  - Generated 15,000 paths using Tor's algorithm
    - 1 out of every 133 paths contained entry and exit node in same AS but distinct /16 subnet
    - All but four also in distinct /8 subnets

Proposed Path Selection Algorithms

- Unique Relay Countries (Unique-CC)
  - Do not permit multiple relays from the same country in a single circuit
  - Easy to implement with current Tor software
  - Has been informally suggested or requested on Tor mailing list

Proposed Path Selection Algorithms

- Unique Relay ASes (Unique-AS)
  - Do not permit multiple relays from the same AS in a single circuit
  - Requires clients or directory authorities to map a relay to an origin AS
  - Tor Proposal #144
    - Tor Proposals are the Tor equivalent of IETF RFCs (requests for comments)
    - Has not been revised since introduced 2008
      - (awaiting clearer research direction)

Proposed Path Selection Algorithms

- Approximate AS Paths
  - Directory authorities generate and distribute AS graph snapshot and prefix table files

- Prior to building a circuit, clients can
  1. Map self, entry node, exit node, destination to ASes in the topology
  2. Compute shortest length *valley-free* paths from
     - Client to entry node (and reverse)
     - Exit node to destination (and reverse)
  3. Sort in descending order by frequency value
  4. Compare the top *n* paths for intersections

Testing AS-aware routing
Results Summary

- Used same 3 consensus snapshots from Sept. 2008
- Generated 5,000 Tor circuits per snapshot per algorithm

	Forward	Reverse	**Total**
Uniform	12.79%	13.23%	20.49%
Weighted (Tor)	10.92%	11.14%	17.81%
Unique-CC	10.41%	11.24%	17.61%
Unique-AS	10.07%	10.14%	16.73%
Approx. AS Path $(n = 1)$	*6.29%*	*6.01%*	*11.09%*
Approx. AS Path $(n = 3)$	*3.17%*	*3.34%*	*6.23%*

Adversaries can live on network links to/from onion routers too

- “Location diversity in anonymity networks” Feamster-Dingledine. WPES 2004
- Model adversaries at Autonomous Systems (ASes)
  - Path Independence: No AS is on both client and destination end of circuit
- How bad is it?   What can we do?
- “AS-awareness in Tor Path Selection” Edman-Syverson. CCS 2009
- It’s fairly bad (for Path Indep.)
- Can design AS-aware routing algorithms

- Is that it? Any other link-level problems?

Link Adversary

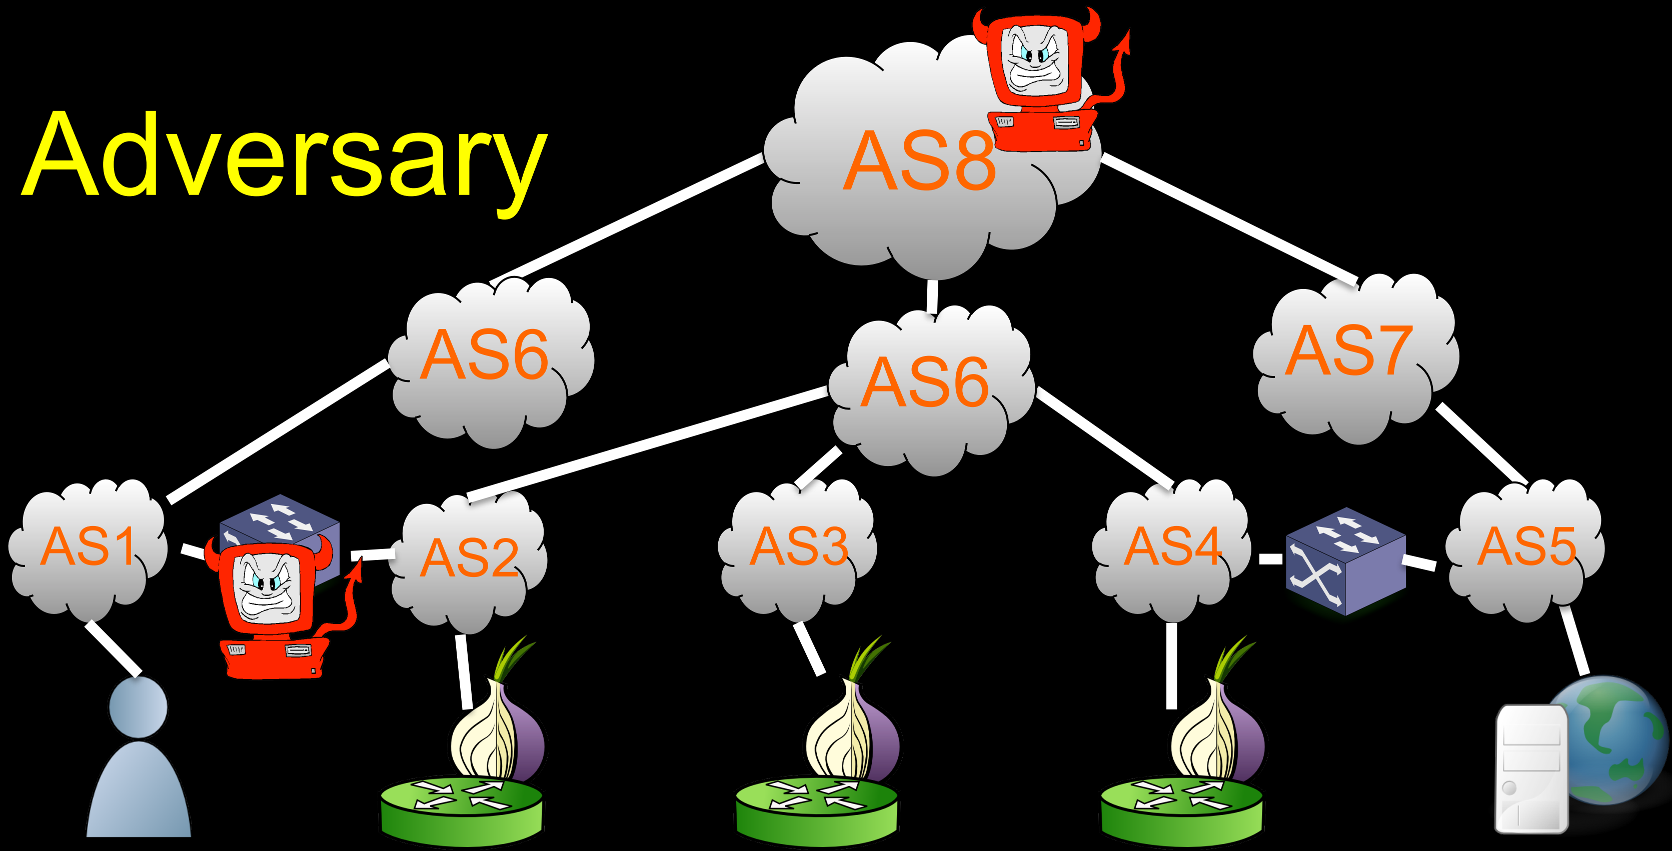

- For performance and cost, many ASes peer directly at Internet Exchange Points (IXPs)
  - Invisible to BGP and route inference
  - Can be found by traceroute

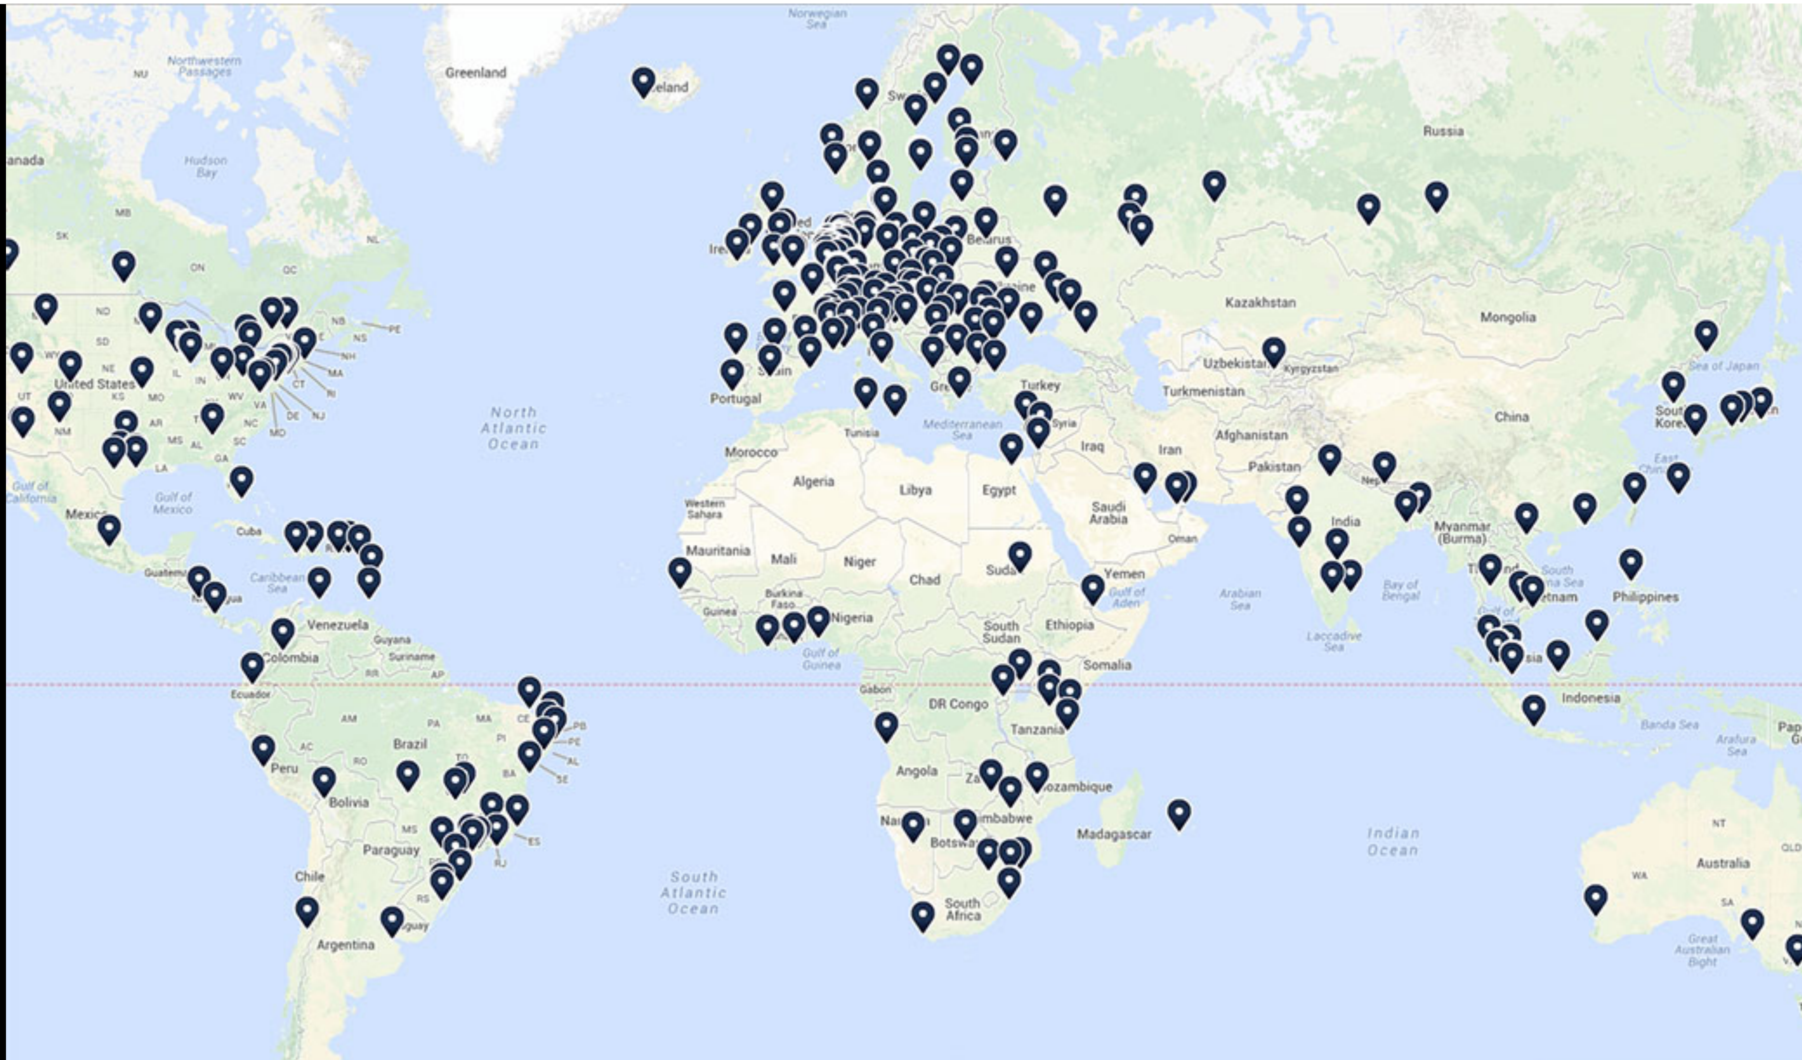

- Thousands of IXPs around the world
- Example: One company Equinix operates
  - 100+ IXPs, in 33 metro areas, in 15 countries, on 5 continents
  - Estimates itself to be on 90% Internet routes

London IX (LINX) main bldg

Inside Midwest IX, Indianapolis

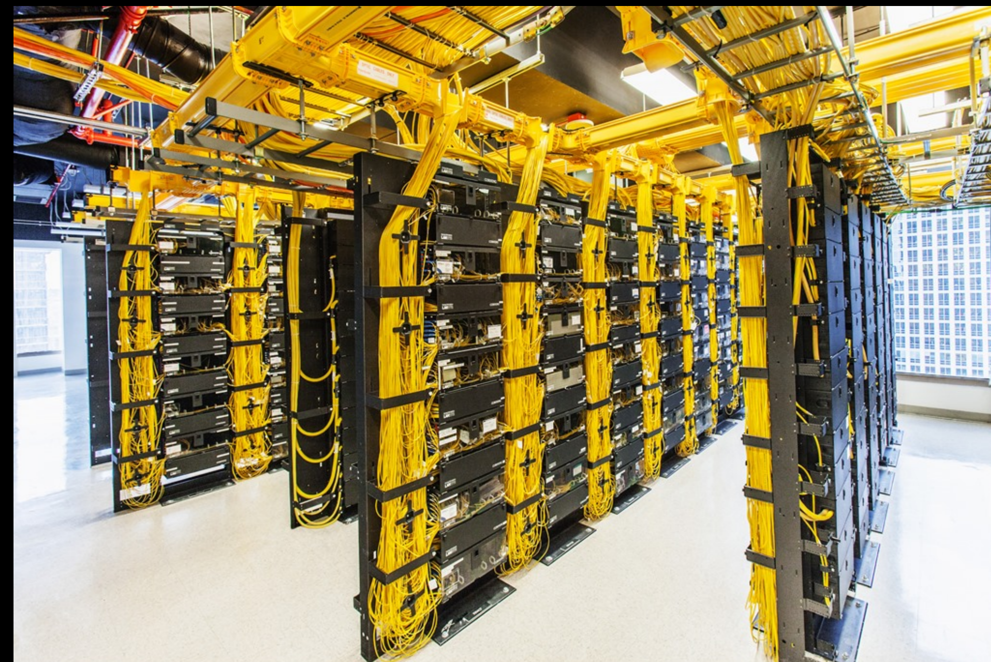

- Murdoch and Zielinski (PETS 2007) showed 27% of routes to UK Tor nodes passed through LINX

- Also showed can recognize Tor circuits at low sampling rates (c. 1/2000 packets) needed to cope with high volumes of IXPs

Adversaries can live on network links to/from onion routers too

- “Location diversity in anonymity networks” Feamster-Dingledine. WPES 2004
- Model adversaries at Autonomous Systems (ASes)
  - Path Indepence: No AS is both
    - between Alice and Tor Entry Guard
    - Between Bob and Tor Exit
- “AS-awareness in Tor Path Selection” Edman-Syverson. CCS 2009
- Empirical analysis of Path Independence on live Tor network
- First AS-aware path selection algorithm
- “Sampled Traffic Analysis by Internet-Exchange-Level Adversaries” Murdoch-Zielinski. PETS 2007
- Can correlate traffic at low sampling rate (1/2000) necessary for high volume locations IXPs (Internet Exchange Points)

How Bad Is It Really?
Putting it all together for correlating adversaries

"Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries" Johnson et al. CCS 2013

- Empirical analysis of security against adversaries controlling moderate fraction of resources on Tor network
  - Tor relays
  - Autonomous Systems
  - Internet Exchanges and families of Internet Exchanges

How Bad Is *WHAT* Really?

"Locating Hidden Servers" Øverlier and Syverson. IEEE S&P (Oakland) 2006

- Single hostile relay and client could find an onion service in a few seconds or minutes
- Note to anonymity geeks: First known intersection attack against production network

Attacking Hidden Servers
(Not Simulations)

Attacking Hidden Servers
(Actual Attacks on Servers in the Wild)

How Bad Is *WHAT* Really?

"Locating Hidden Servers" Øverlier and Syverson. IEEE S&P (Oakland) 2006

- Single hostile relay and client could find an onion service in a few seconds or minutes
  - Analysis of attack on onion service over live Tor network
  - Basis of introduction of guards
  - Onion services can be caused to create many circuits back to client
- Moral: Must consider client behavior when modeling adversary capabilities

User Models from “Users Get Routed”

Gmail/GChat

Gcal/GDocs

Facebook

Web search

IRC

BitTorrent

20-minute traces

User Models from "Users Get Routed"

- Gmail/GChat
- Gcal/GDocs
- Facebook
- Web search

Typical

- IRC
- BitTorrent

20-minute traces

User Models

User model	Session schedule
Gmail/GChat, Gcal/GDocs, Facebook, Web search } Typical	Sessions at 9:00, 12:00, 15:00, and 18:00 Su-Sa
IRC	Repeated sessions 8:00-17:00, M-F
BitTorrent	Repeated sessions 0:00-6:00, Sa-Su

20-minute traces

How Bad Is It Really? Putting it all together

"Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries" Johnson et al. CCS 2013

- Empirical analysis of security against adversaries controlling moderate fraction of resources on Tor network
  - Tor relays
  - Autonomous Systems
  - Internet Exchanges and families of Internet Exchanges

How Bad Is It Really? Putting it all together

"Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries" Johnson et al. CCS 2013

- Empirical analysis of security against adversaries controlling moderate fraction of resources on Tor network wrt various usage models
  - Tor relays
  - Autonomous Systems
  - Internet Exchanges and families of Internet Exchanges

How Bad Is *WHAT* Really? (Part 2)

What is the adversary trying to accomplish

- Prior metrics ask things like
  - How differentiated is the set of all users of this system by the adversary?
  - What fraction of circuits through the network are compromised at a given time?
- Users want to know how secure *they* are against a realistic adversary
  - If I use the network the way I use it, how long till I get a compromised connection by?
  - What fraction of my traffic will get compromised if I use the system the way I use it for T hours/months/etc. ?

Adversary Framework

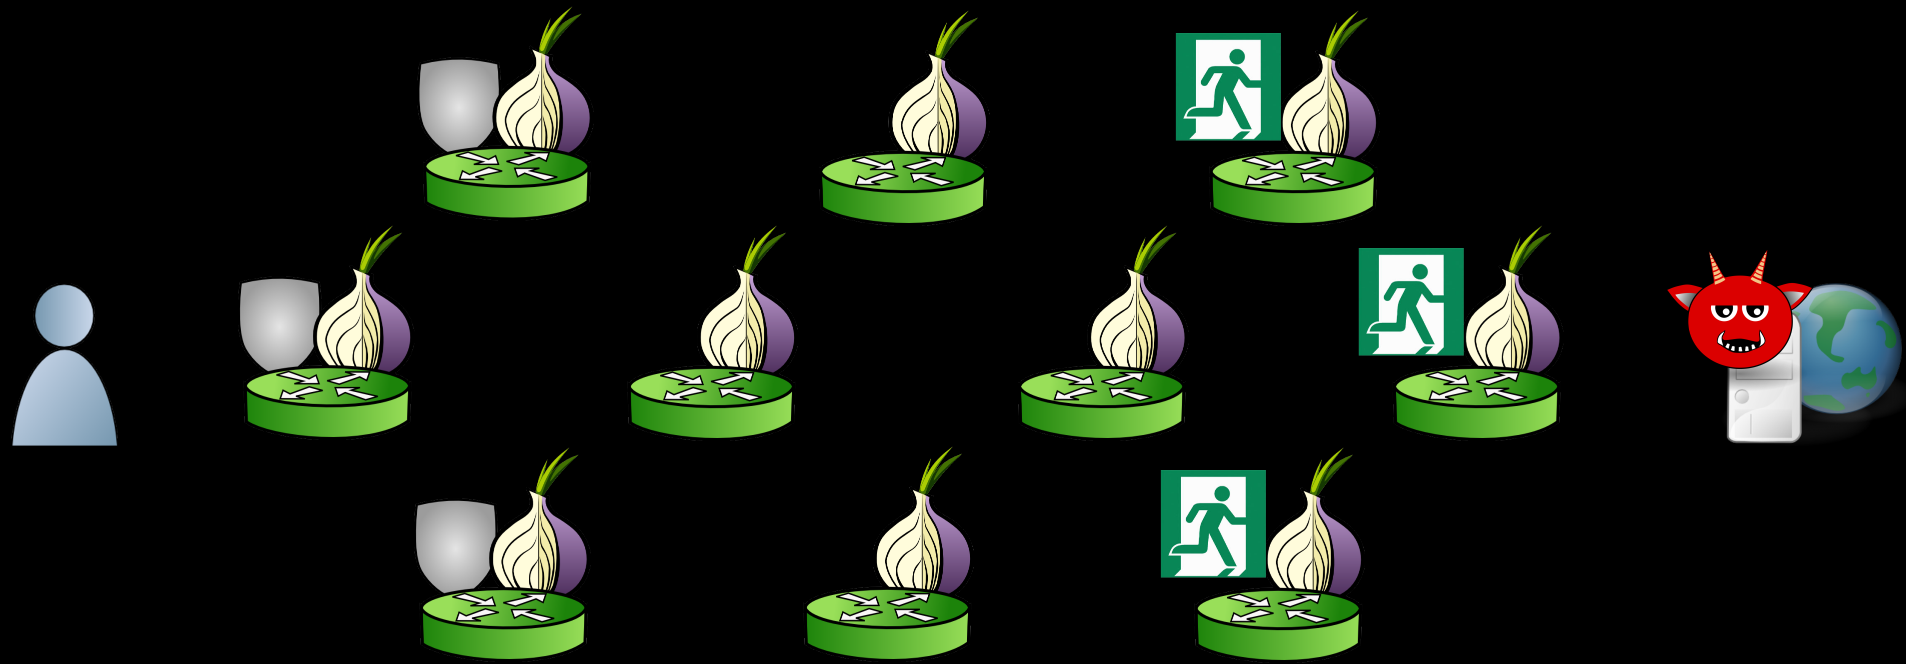

Adversary Framework

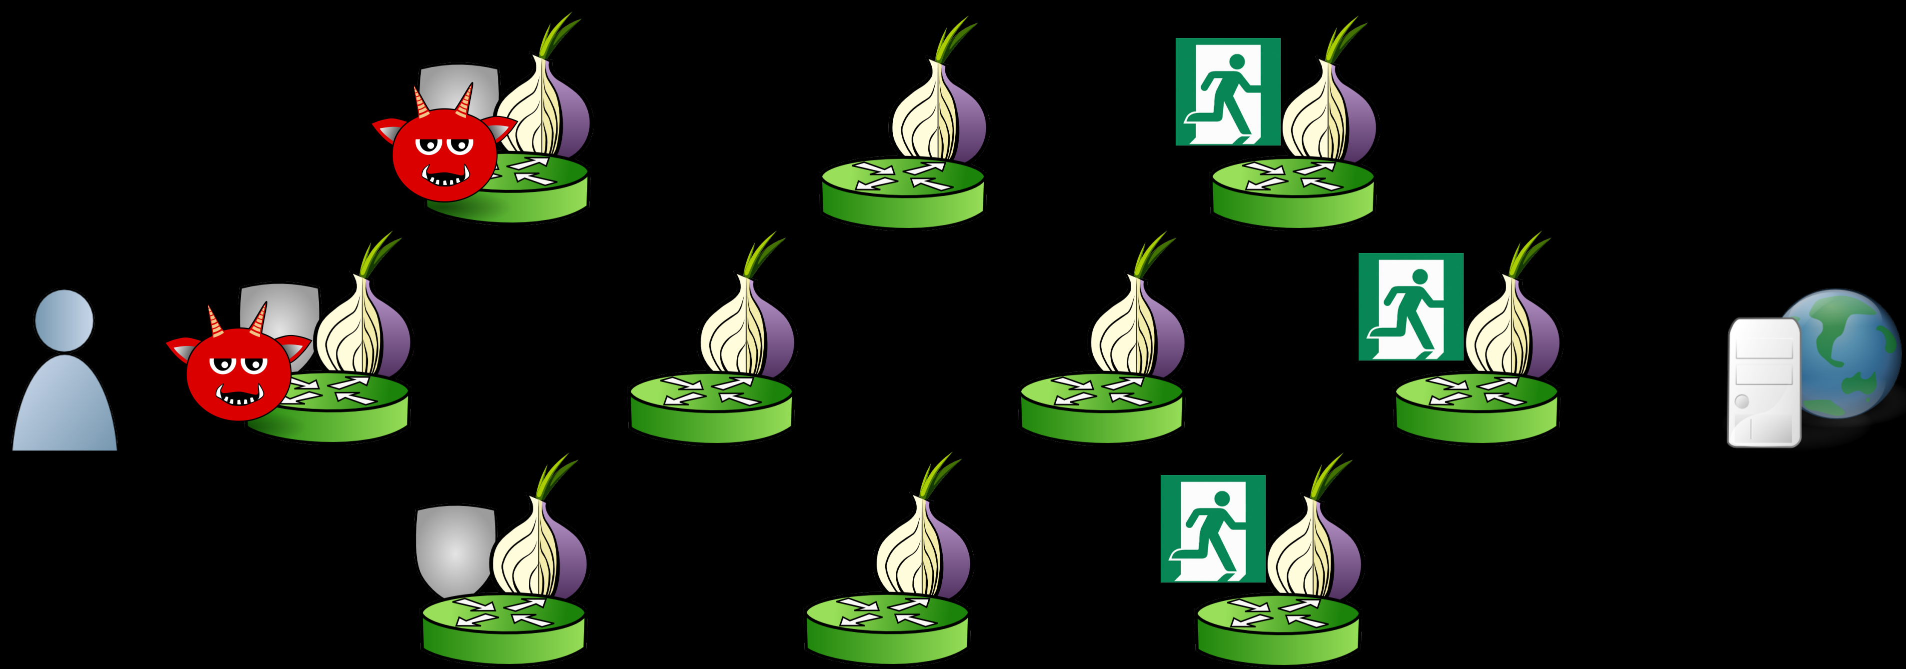

Adversary Framework

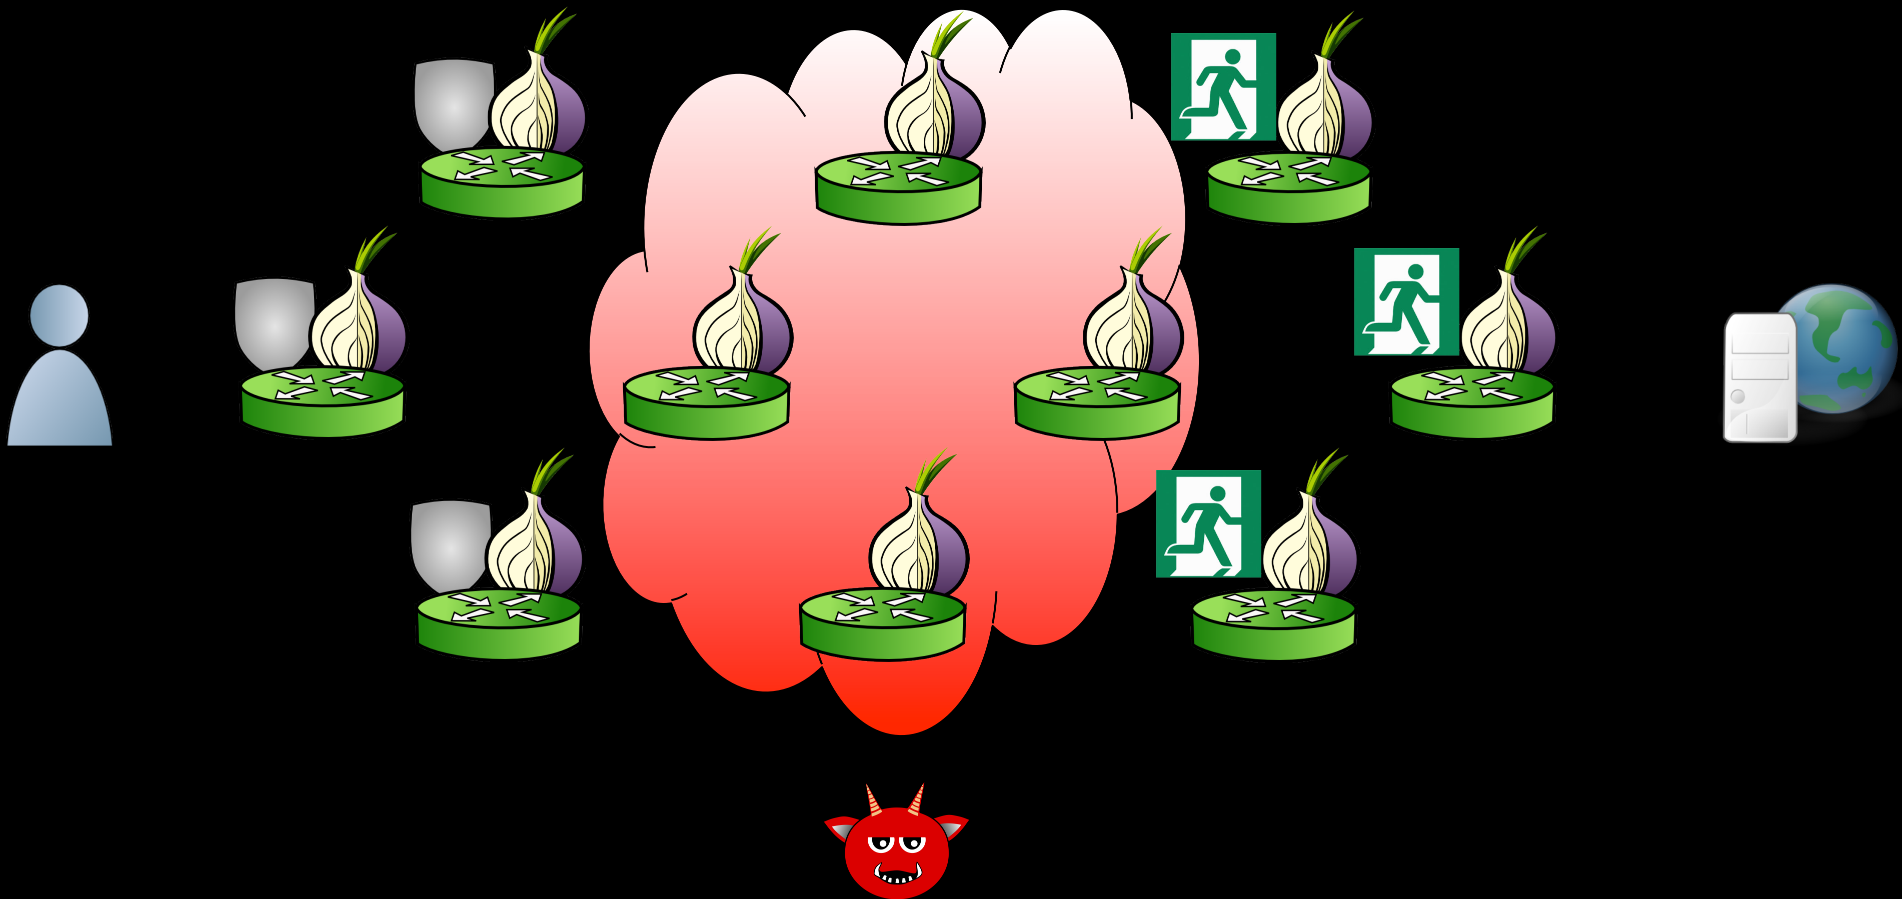

Adversary Framework

Resource Types

- Relays
- Bandwidth
- Autonomous Systems (ASes)
- Internet Exchange Points (IXPs)
- Money

Adversary Framework

Resource Types

- Relays
- Bandwidth
- Autonomous Systems (ASes)
- Internet Exchange Points (IXPs)
- Money

Resource Endowment

- Destination host
- 5% Tor bandwidth
- Source AS
- Equinix IXPs

Adversary Framework

Resource Types

- Relays
- Bandwidth
- Autonomous Systems (ASes)
- Internet Exchange Points (IXPs)
- Money

Resource Endowment

- Destination host
- 5% Tor bandwidth
- Source AS
- Equinix IXPs

Goal

- Target a given user's communication
- Compromise as much traffic as possible
- Learn who uses Tor
- Learn what Tor is used for

“Users Get Routed” Outline Summary

- Tor Security Analysis
  - Adversary Framework
  - Security Metrics
  - Evaluation Methodology
  - Node Adversary Analysis
  - Link Adversary Analysis

Security Metrics

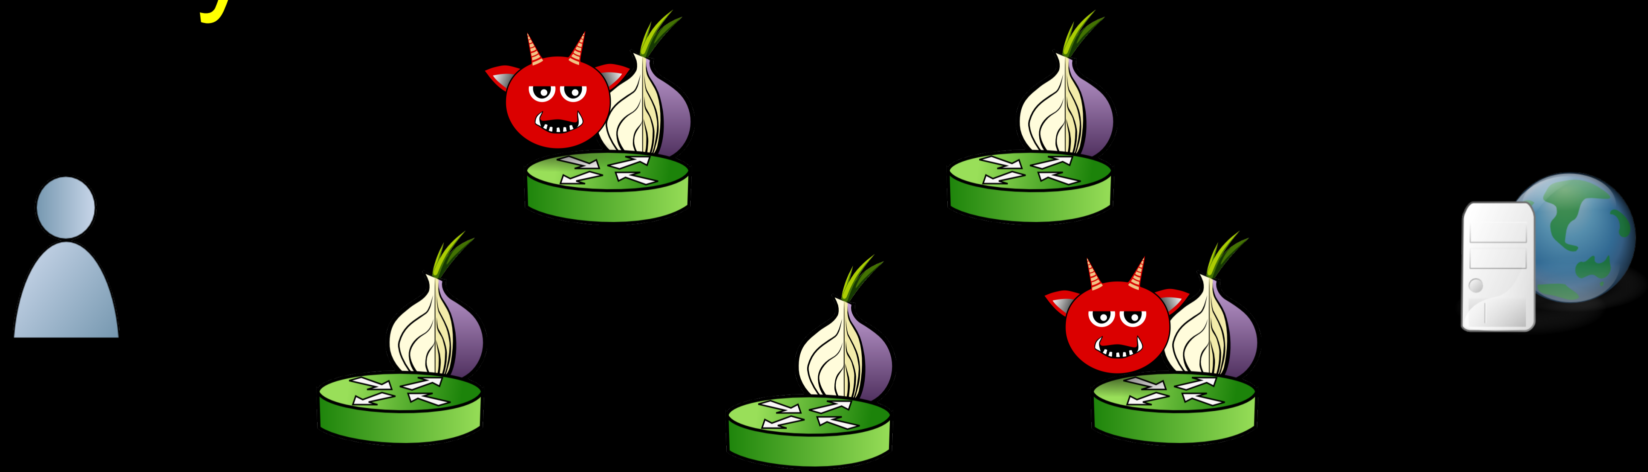

Prior metrics

Security Metrics

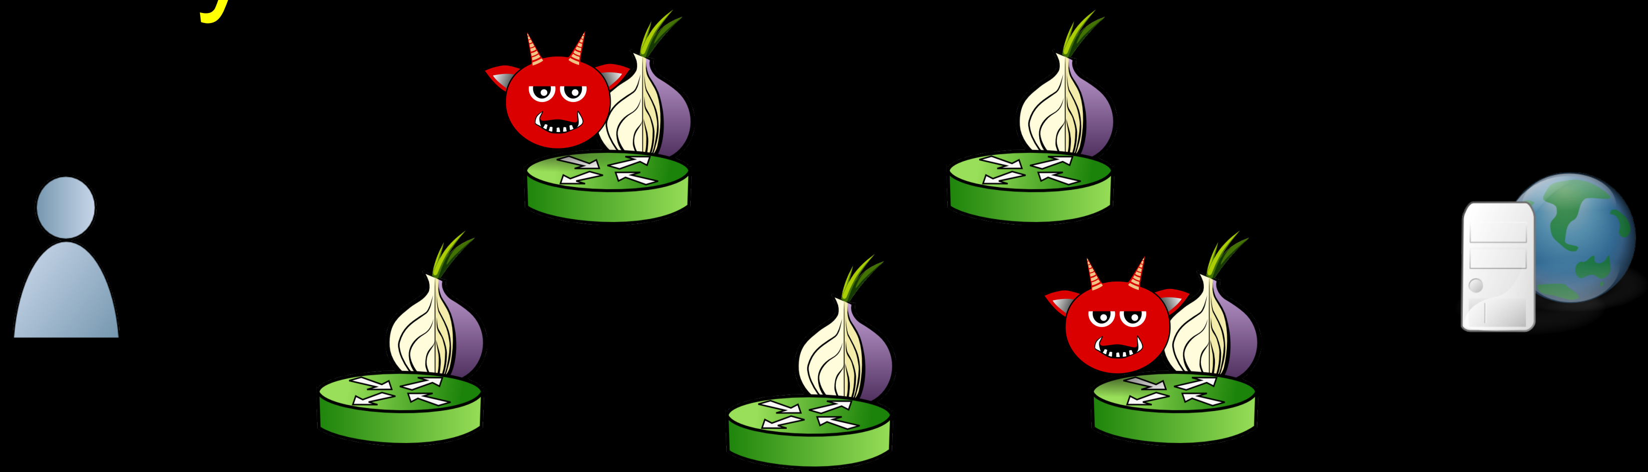

Prior metrics

1. Probability of choosing bad guard and exit
   a. c^2 / n^2 : Adversary controls c of n relays
   b. ge : g guard and e exit BW fractions are bad

Security Metrics

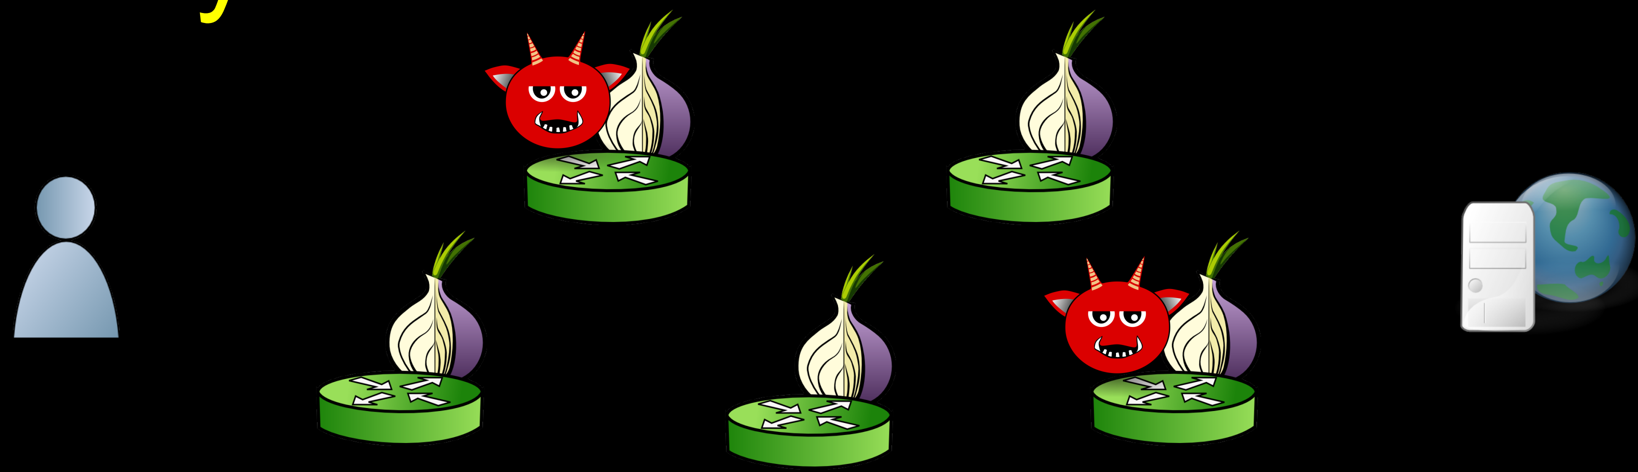

Prior metrics

1. Probability of choosing bad guard and exit
   a. c^2 / n^2 : Adversary controls c of n relays
   b. ge : g guard and e exit BW fractions are bad
2. Probability *some* AS/IXP exists on both entry and exit paths (i.e. *path independence*)

Security Metrics

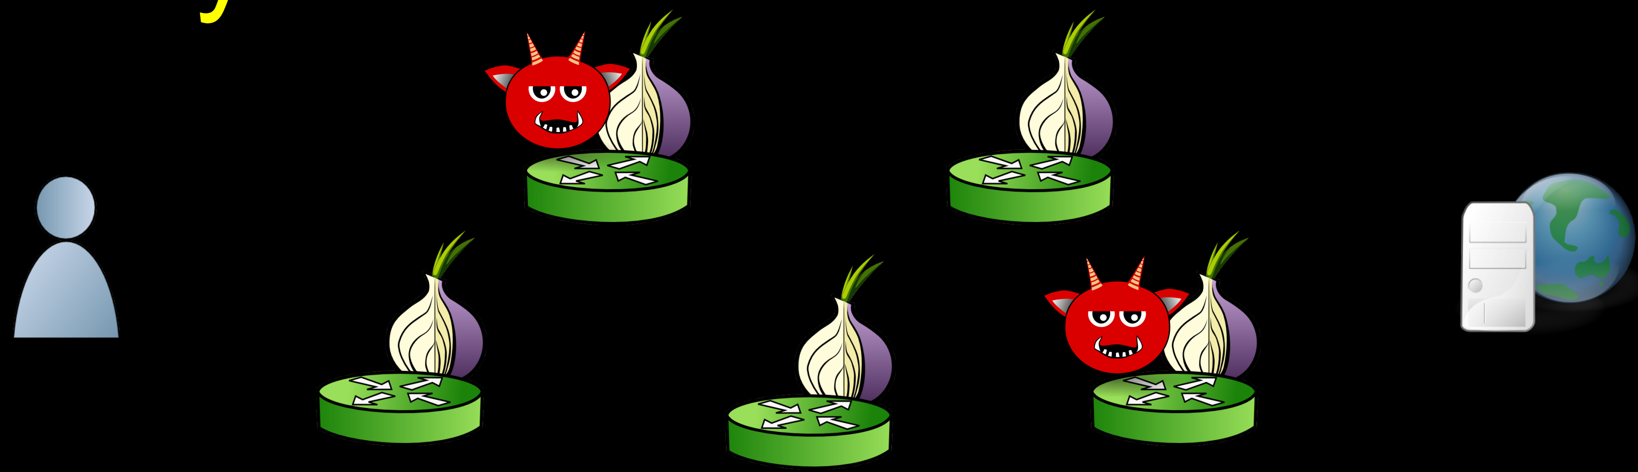

Prior metrics

1. Probability of choosing bad guard and exit
   a. c^2 / n^2 : Adversary controls *c* of *n* relays
   b. *ge* : *g* guard and *e* exit BW fractions are bad
2. Probability *some* AS/IXP exists on both entry and exit paths (i.e. *path independence*)

3. g_t : Probability of choosing malicious guard within time *t*

Security Metrics

Principles

1. Probability distribution
2. Measure on human timescales
3. Based on adversaries

Security Metrics

Principles

1. Probability distribution
2. Measure on human timescales
3. Based on adversaries

Metrics

1. Probability distribution of time until first path compromise
2. Probability distribution of number of path compromises for a given user over given time period

“Users Get Routed” Outline Summary

- Tor Security Analysis
  - Adversary Framework
  - Security Metrics
  - Evaluation Methodology
  - Node Adversary Analysis
  - Link Adversary Analysis

TorPS: The Tor Path Simulator

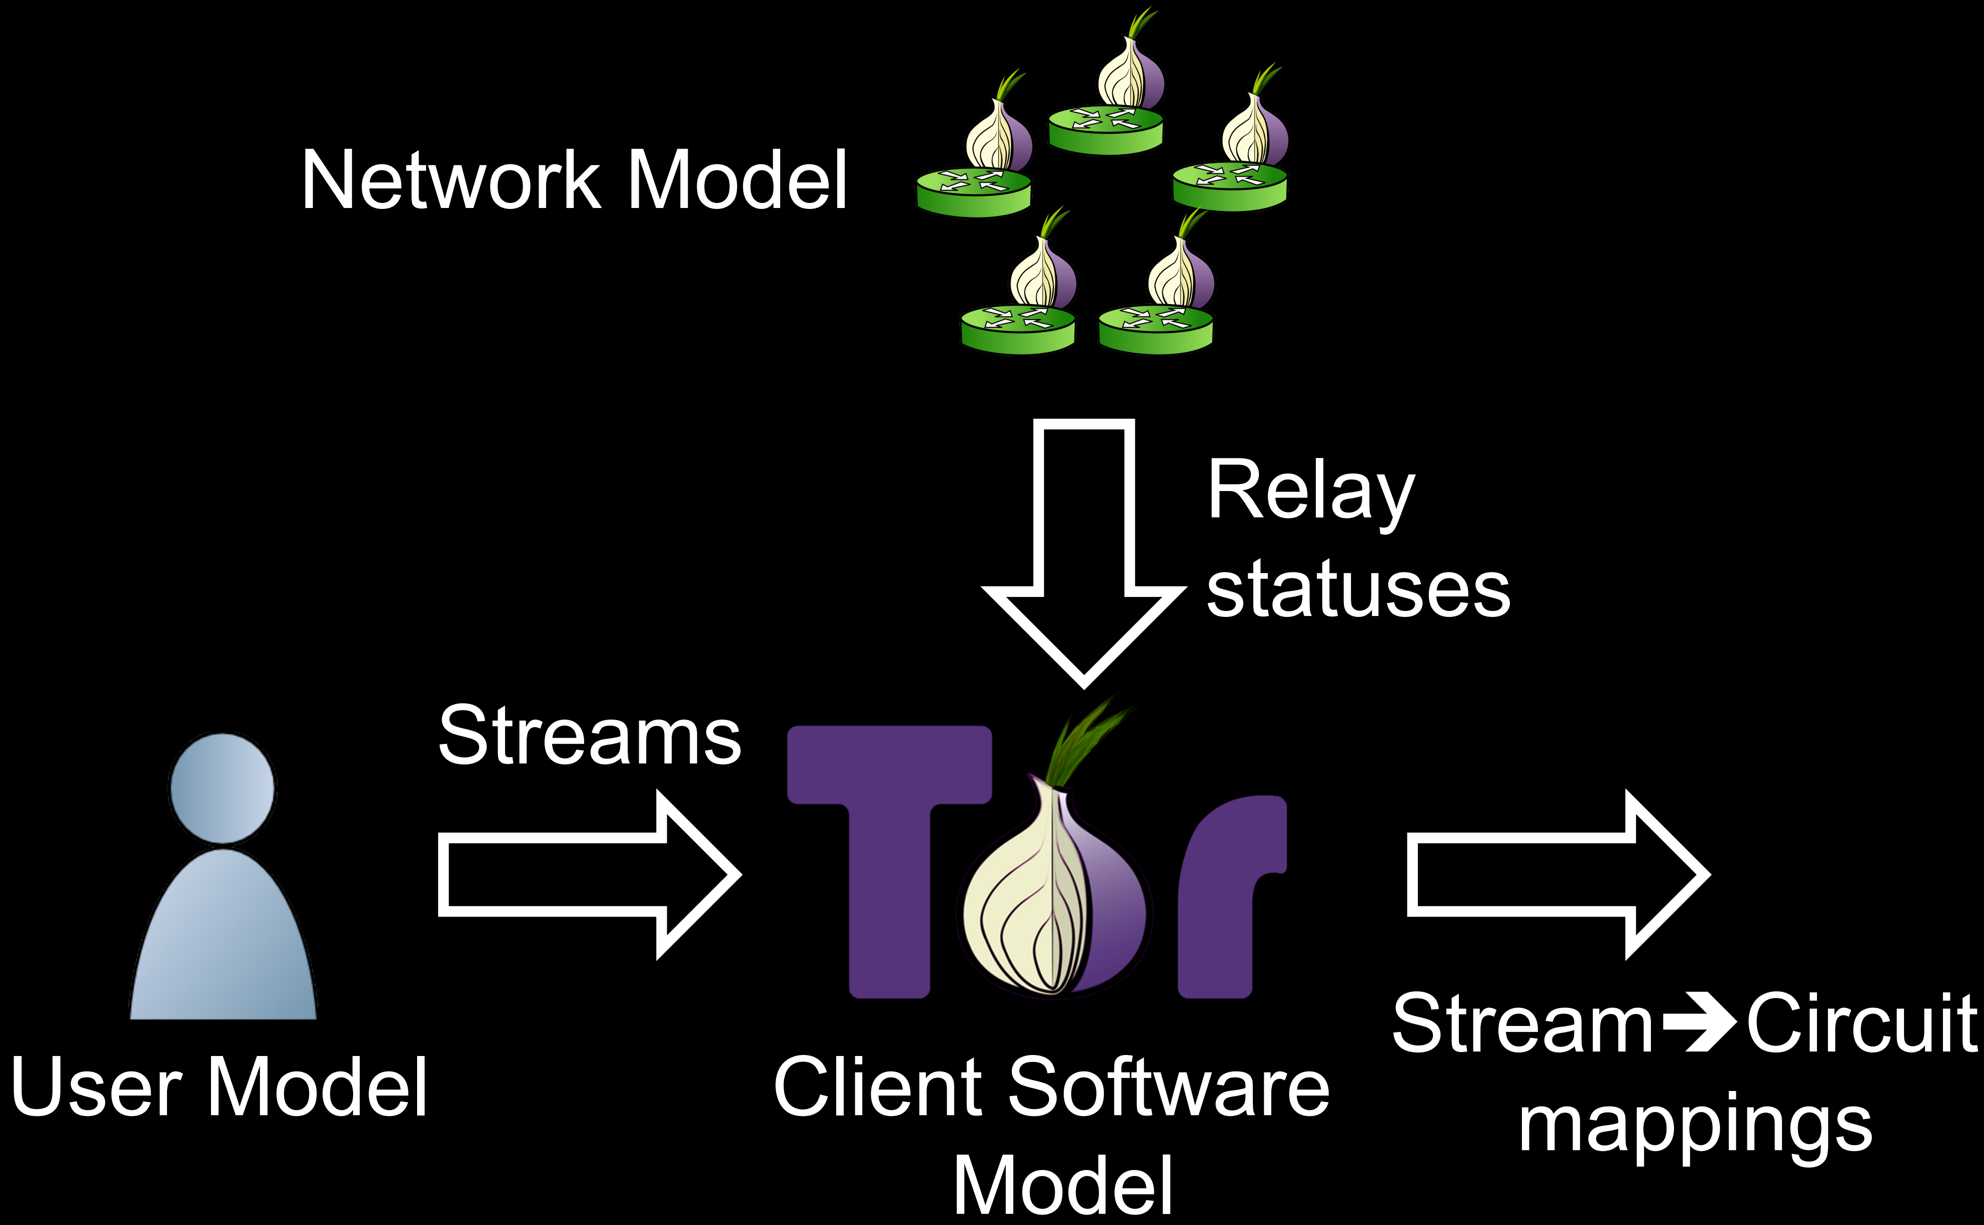

TorPS: The Tor Path Simulator

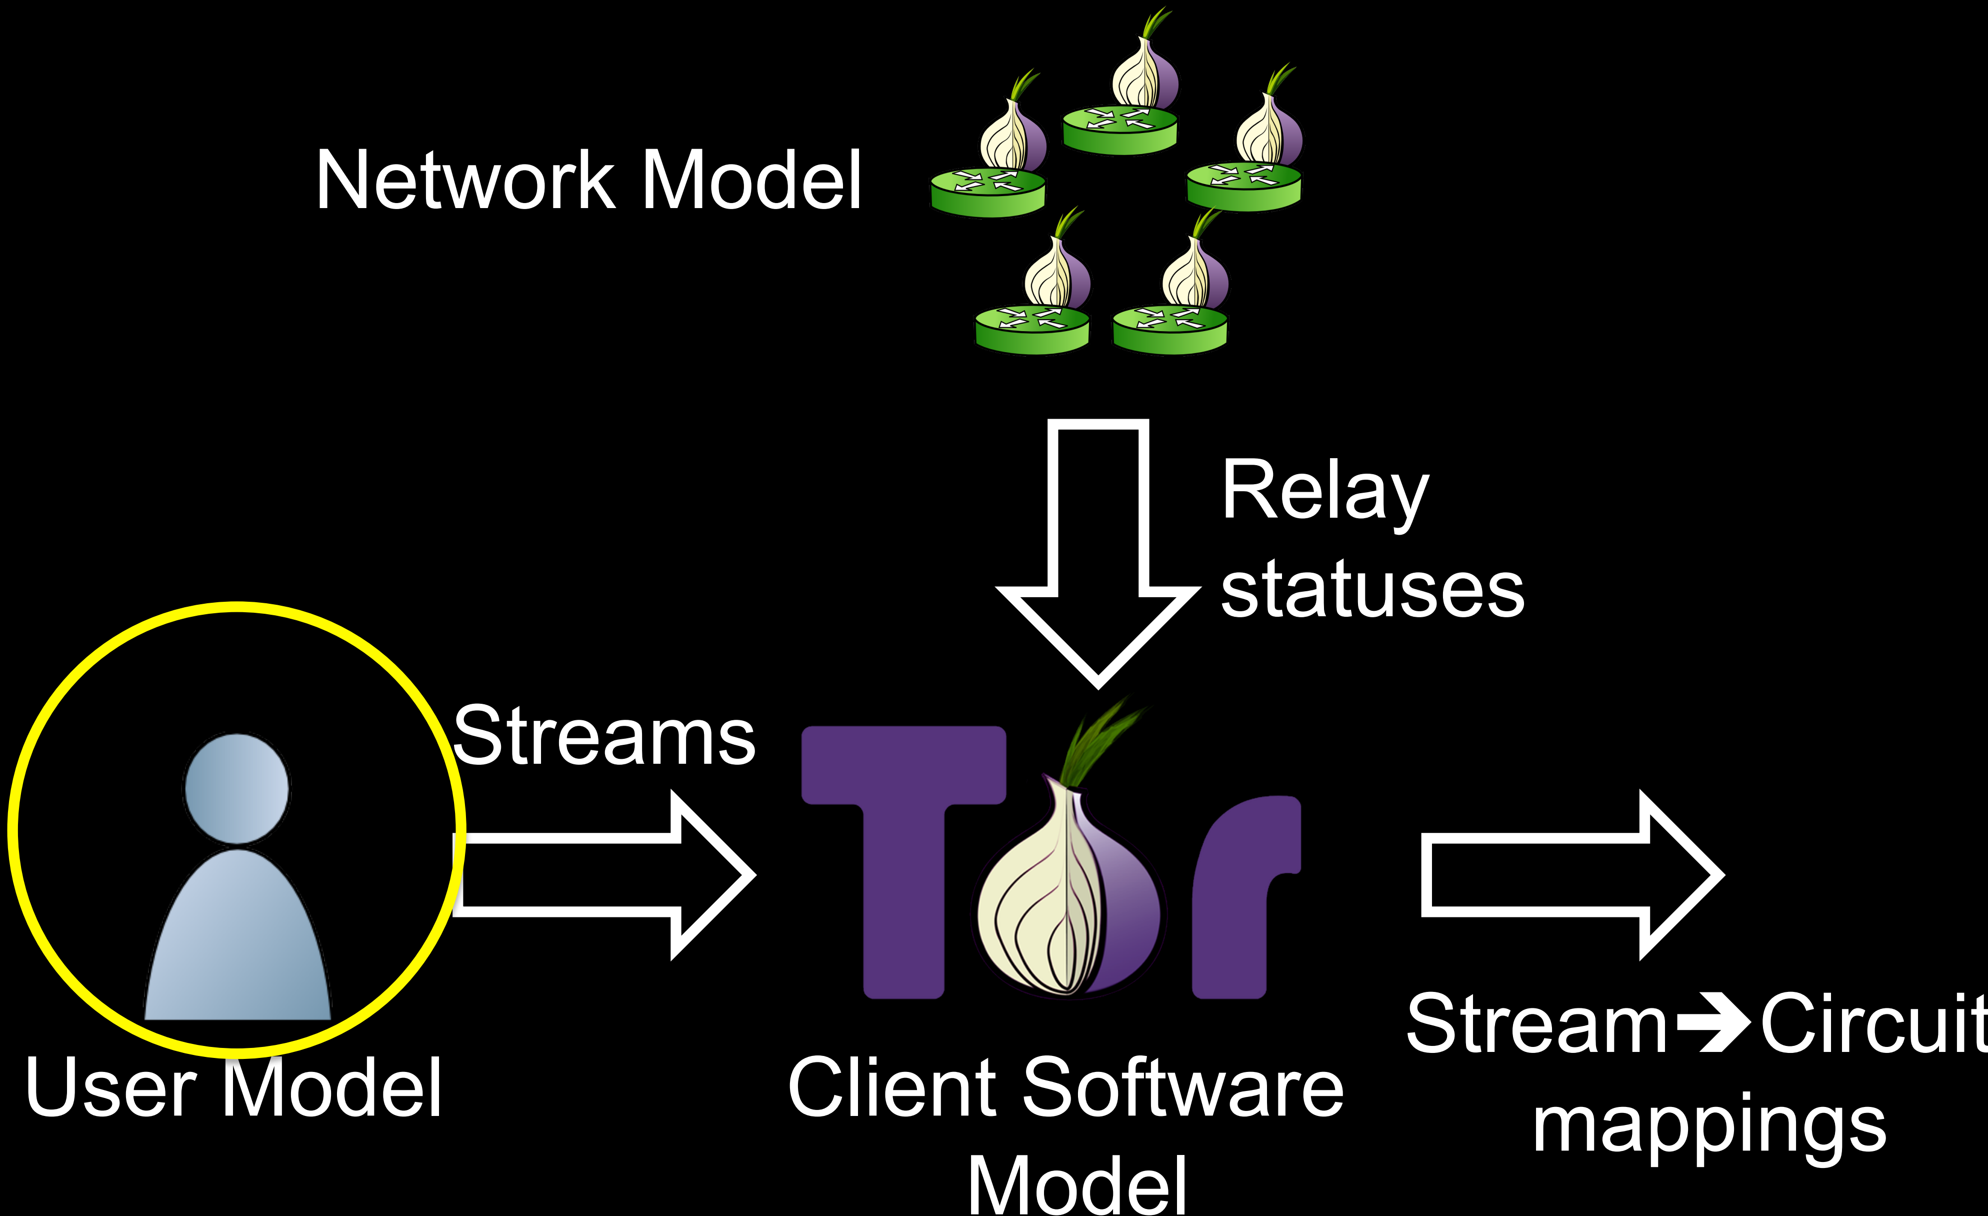

TorPS: User Model

- Gmail/GChat
- Gcal/GDocs
- Facebook
- Web search

} Typical

Worst Port (6523)

Best Port (443)

- IRC
- BitTorrent

Session schedule
Sessions at 9:00, 12:00, 15:00, and 18:00 Su-Sa
Repeated sessions 8:00-17:00, M-F
Repeated sessions 0:00-6:00, Sa-Su

20-minute traces

TorPS: User Model

Rank	Port #	Exit BW %	Long-Lived	Application
1	8300	19.8	Yes	iTunes?
2	6523	20.1	Yes	Gobby
3	26	25.3	No	(SMTP+1)
65312	993	89.8	No	IMAP SSL
65313	80	90.1	No	HTTP
65314	443	93.0	No	HTTPS

Default-accept ports by exit capacity.

TorPS: User Model

Model	Streams/ week	IPs	Ports (#s)
Typical	2632	205	2 (80, 443)
IRC	135	1	1 (6697)
BitTorrent	6768	171	118
WorstPort	2632	205	1 (6523)
BestPorst	2632	205	1 (443)

User model stream activity

TorPS: The Tor Path Simulator

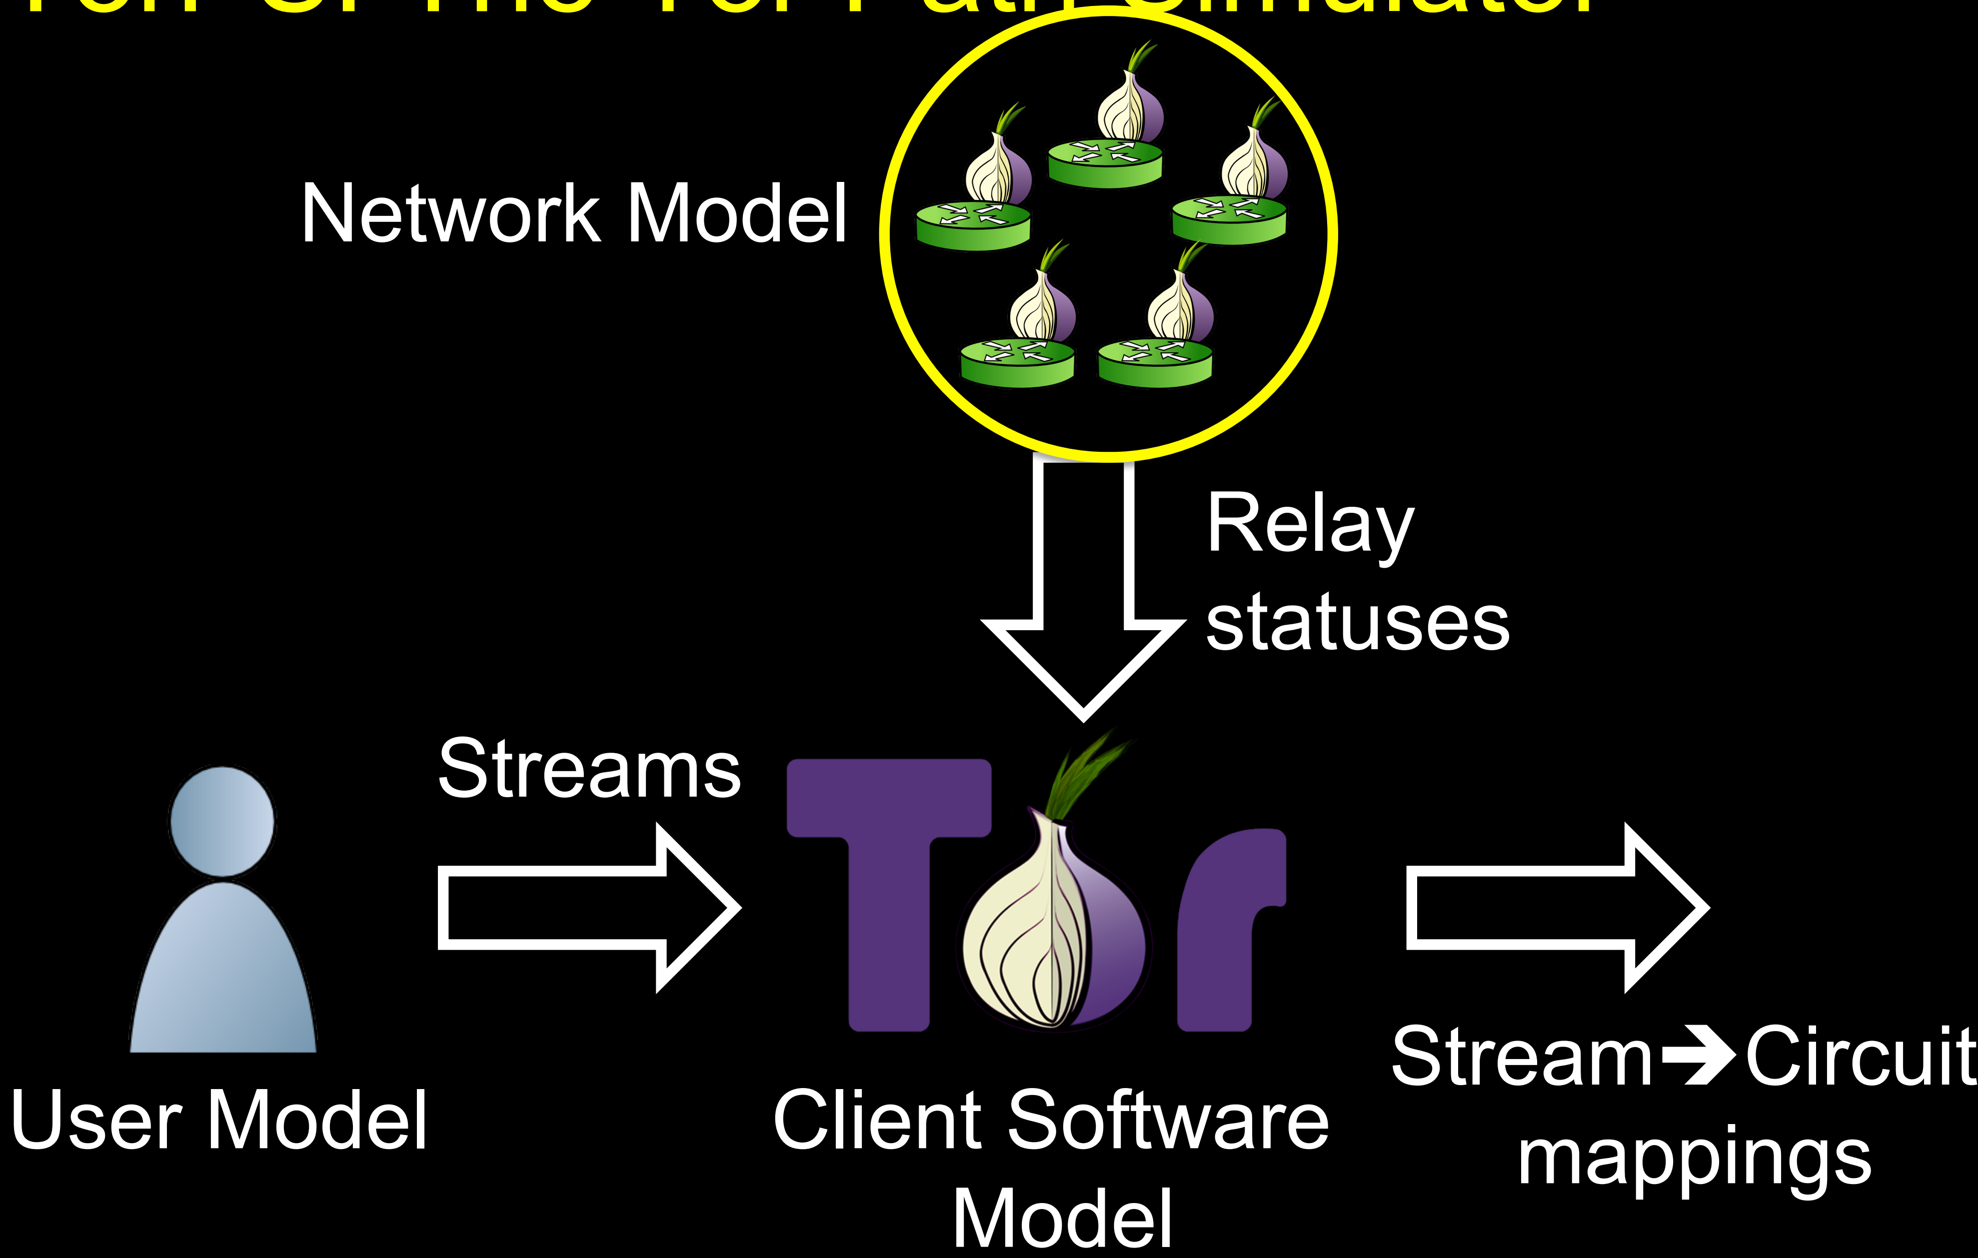

TorPS: The Tor Path Simulator

Network Model

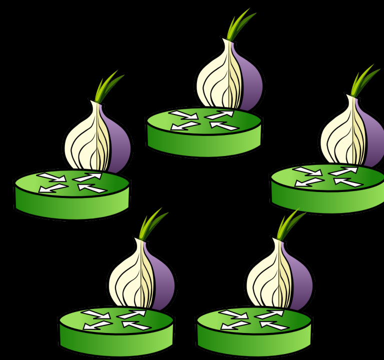

metrics.torproject.org

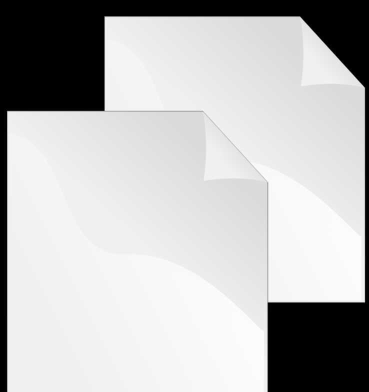

Hourly consensuses

Monthly server descriptors archive

TorPS: The Tor Path Simulator

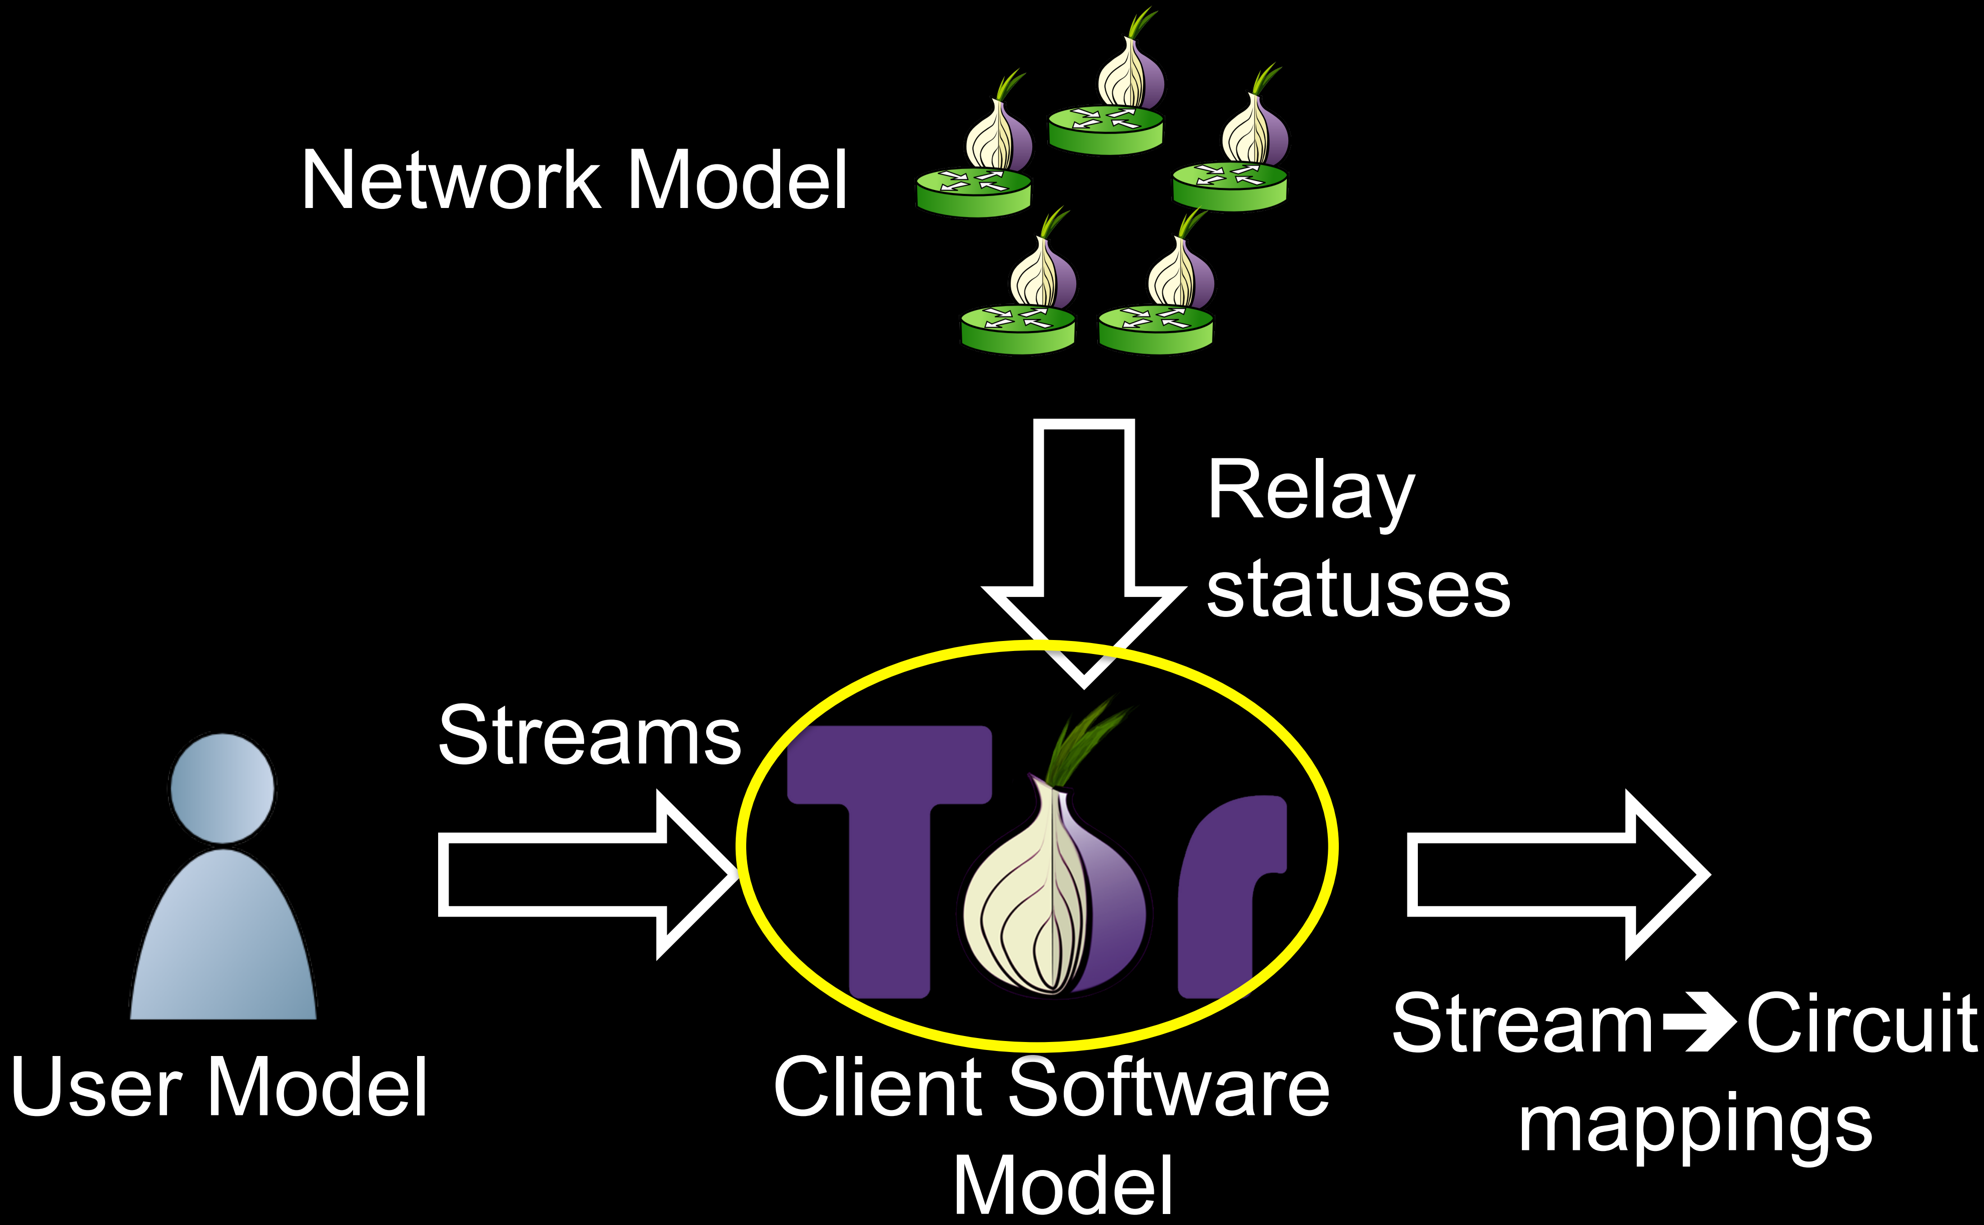

TorPS: The Tor Path Simulator

Client Software Model

- Reimplemented path selection in Python
- Based on current Tor stable version (0.2.3.25)
- Major path selection features include
  - Bandwidth weighting
  - Exit policies
  - Guards and guard rotation
  - Hibernation
  - /16 and family conflicts
- Omits effects of network performance

“Users Get Routed” Outline Summary

- Tor Security Analysis
  - Adversary Framework
  - Security Metrics
  - Evaluation Methodology
  - Node Adversary Analysis
  - Link Adversary Analysis

Node Adversary

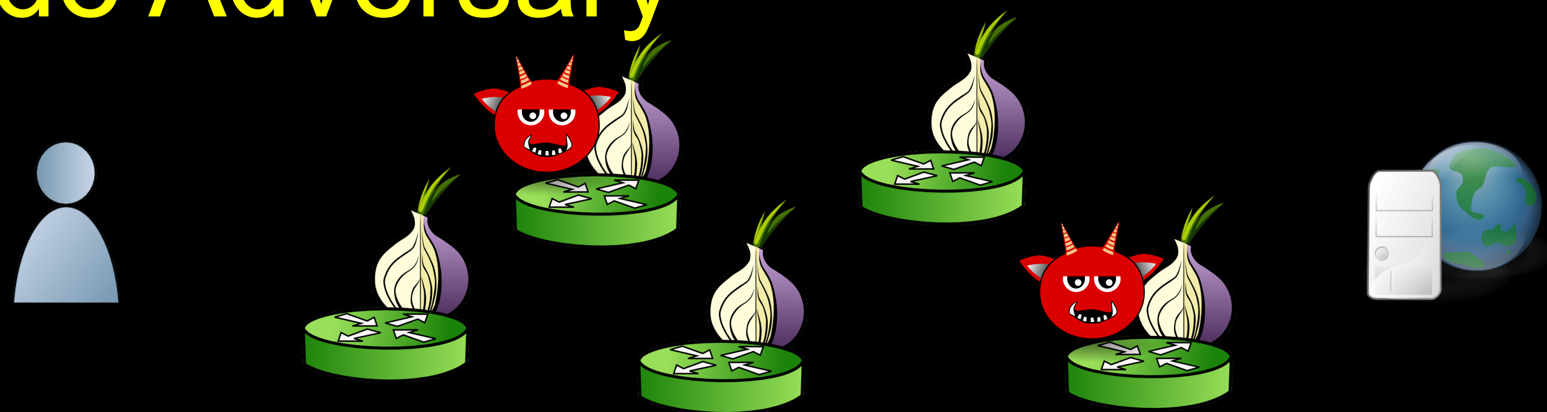

100 MiB/s total bandwidth

Relay Type	Number	Bandwidth (GiB/s)
Any	2646	3.10
Guard only	670	1.25
Exit only	403	0.30
Guard & Exit	272	0.98

Tor relay capacity, 3/31/13

Rank	Bandwidth (MiB/s)	Family
1	260.5	torservers.net
2	115.7	Chaos Computer Club
3	107.8	DFRI
4	95.3	Team Cymru
5	80.5	Paint

Top Tor families, 3/31/13

Node Adversary

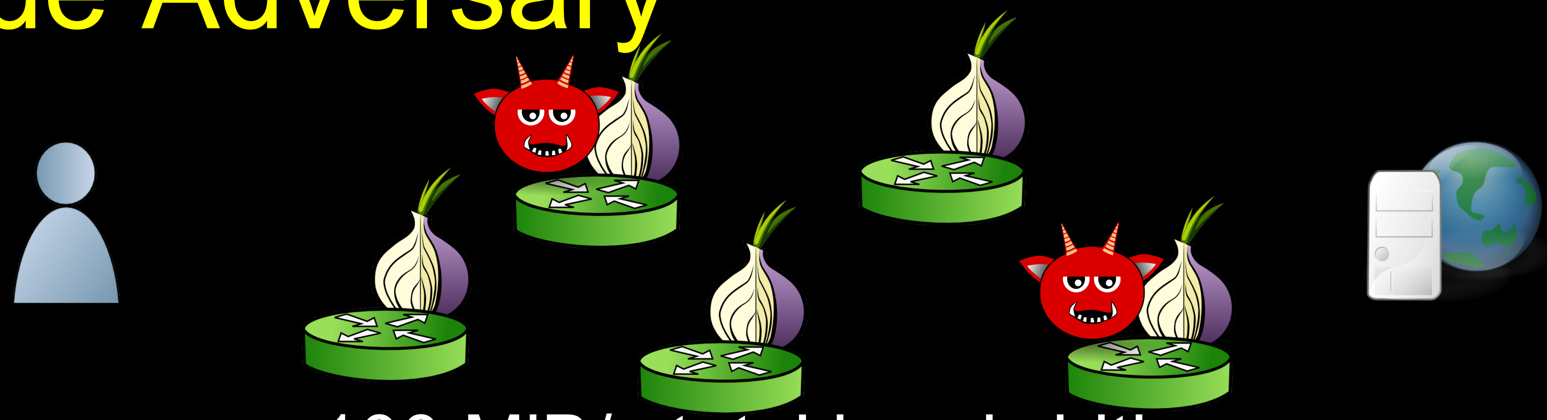

100 MiB/s total bandwidth

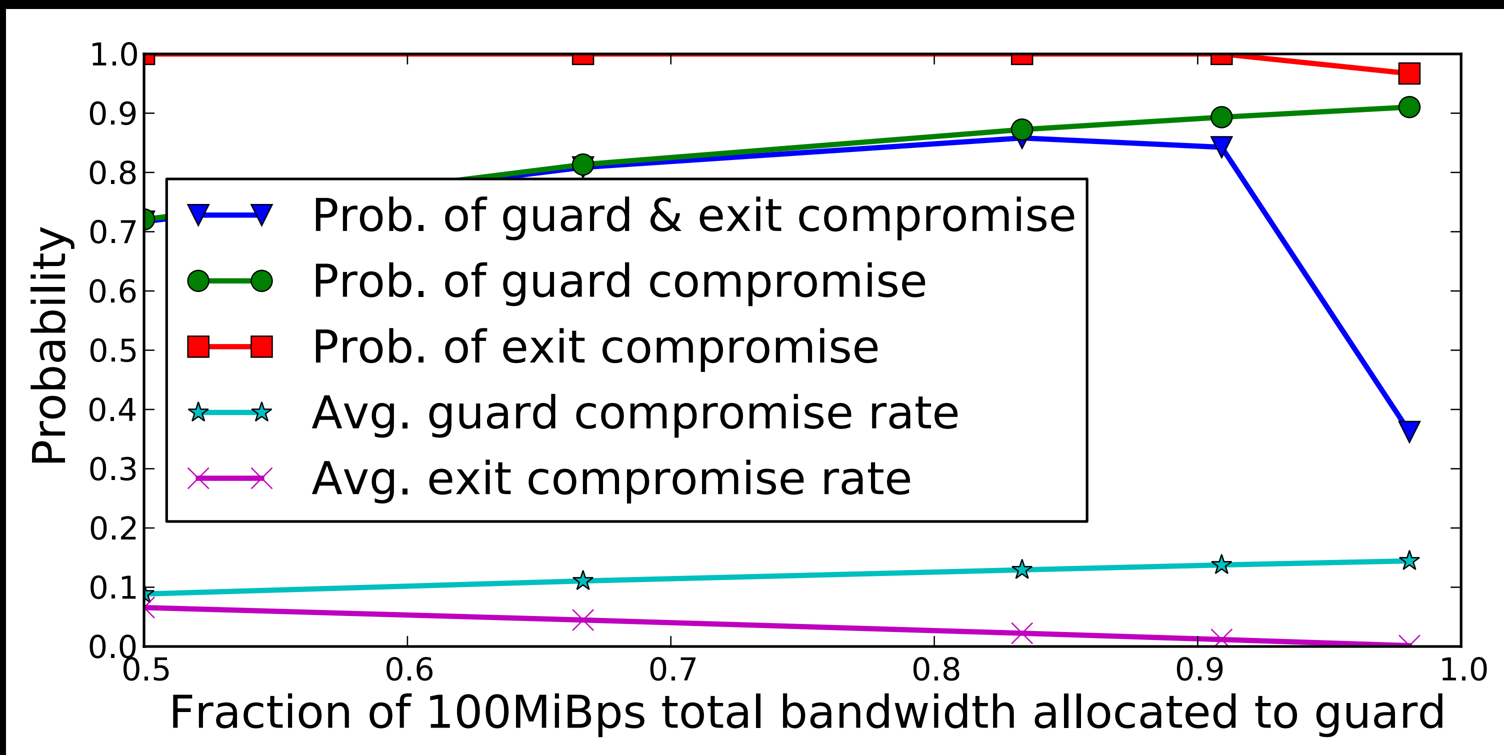

Probability to compromise at least one stream and rate of compromise, 10/12 – 3/13.

Node Adversary

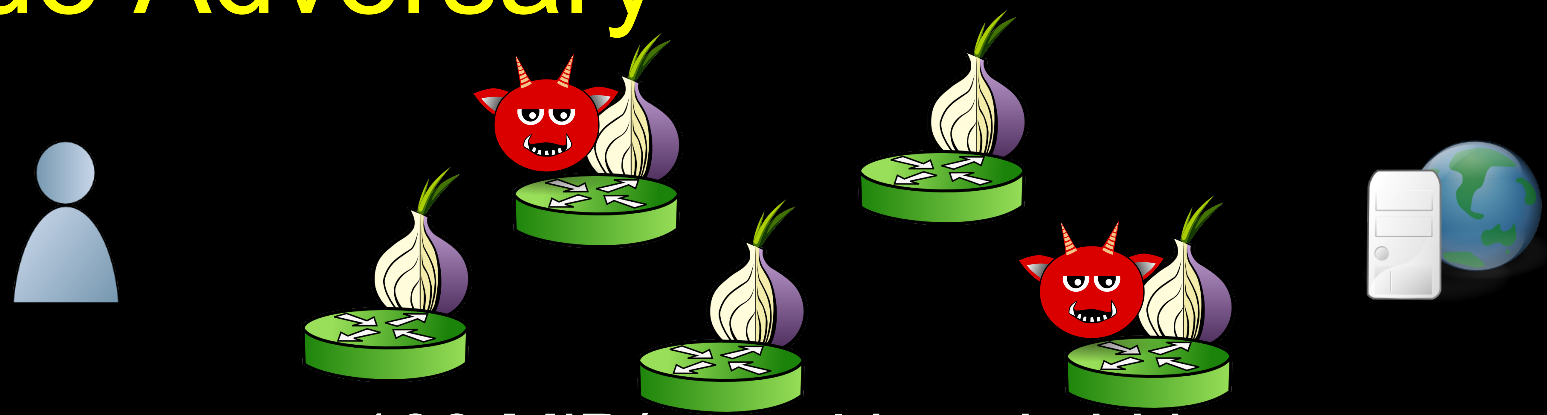

100 MiB/s total bandwidth

83.3 MiB/s guard,16.7 MiB/s exit

Node Adversary Results

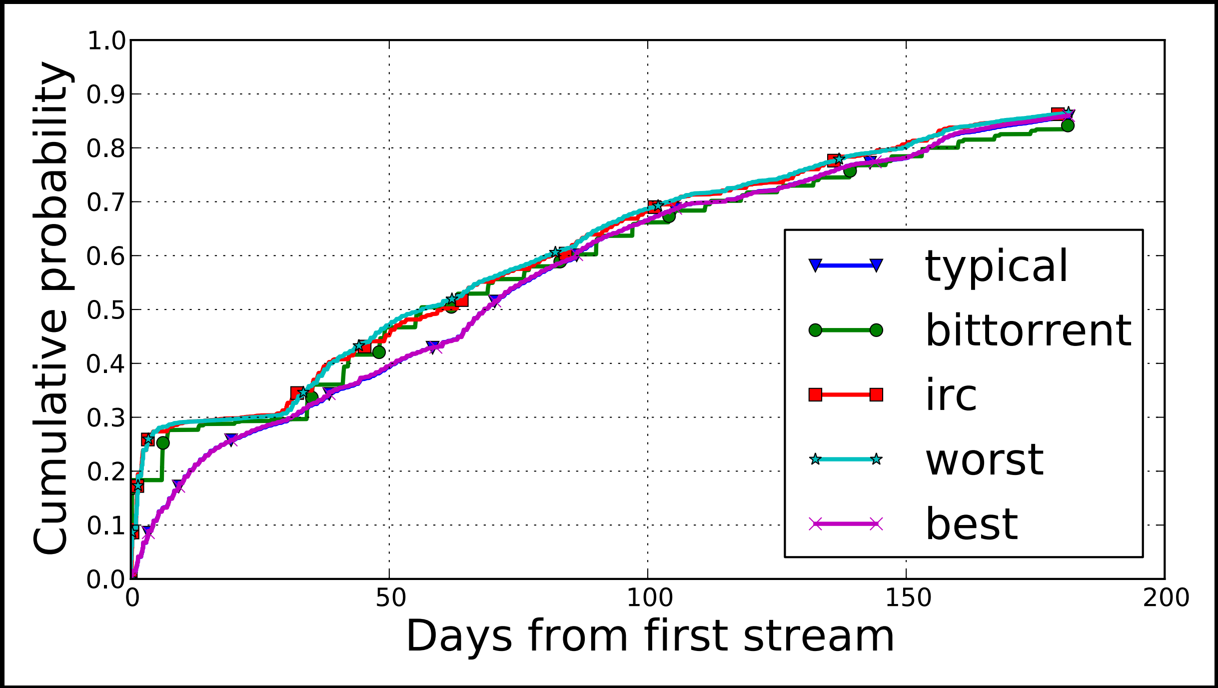

Time to first compromised stream, 10/12 – 3/13

Fraction compromised streams, 10/12 – 3/13

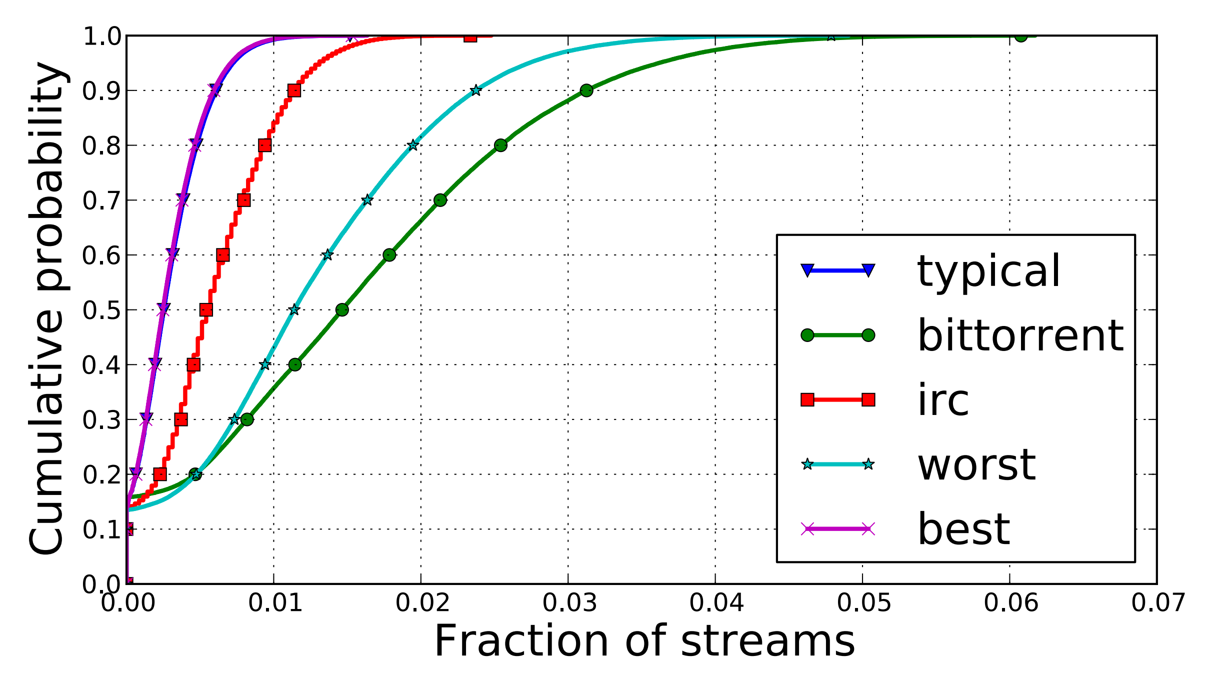

Node Adversary Results

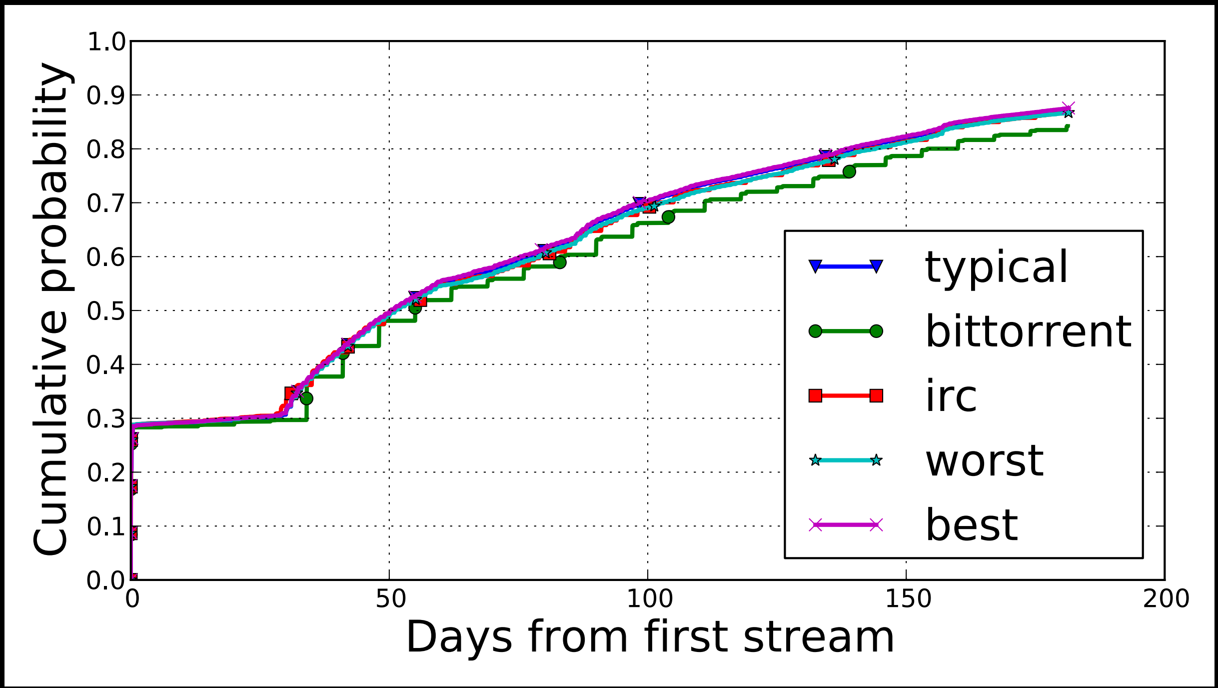

Time to first compromised guard, 10/12 – 3/13

Fraction streams with compromised guard, 10/12 – 3/13

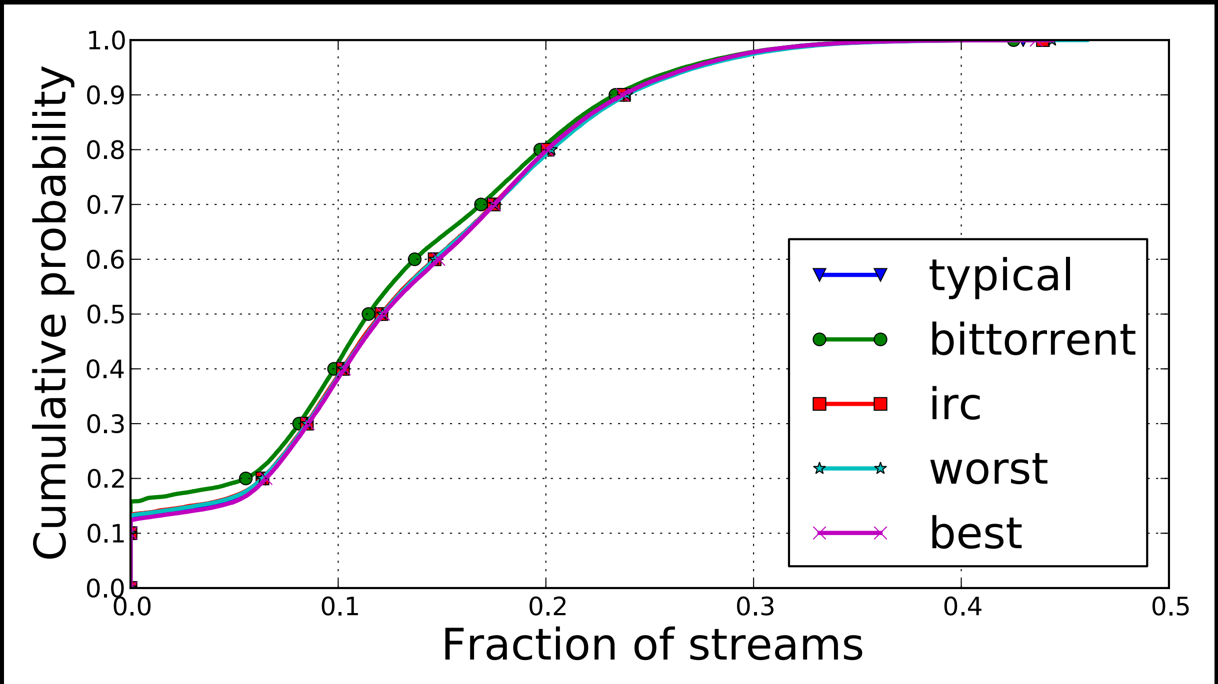

Node Adversary Results

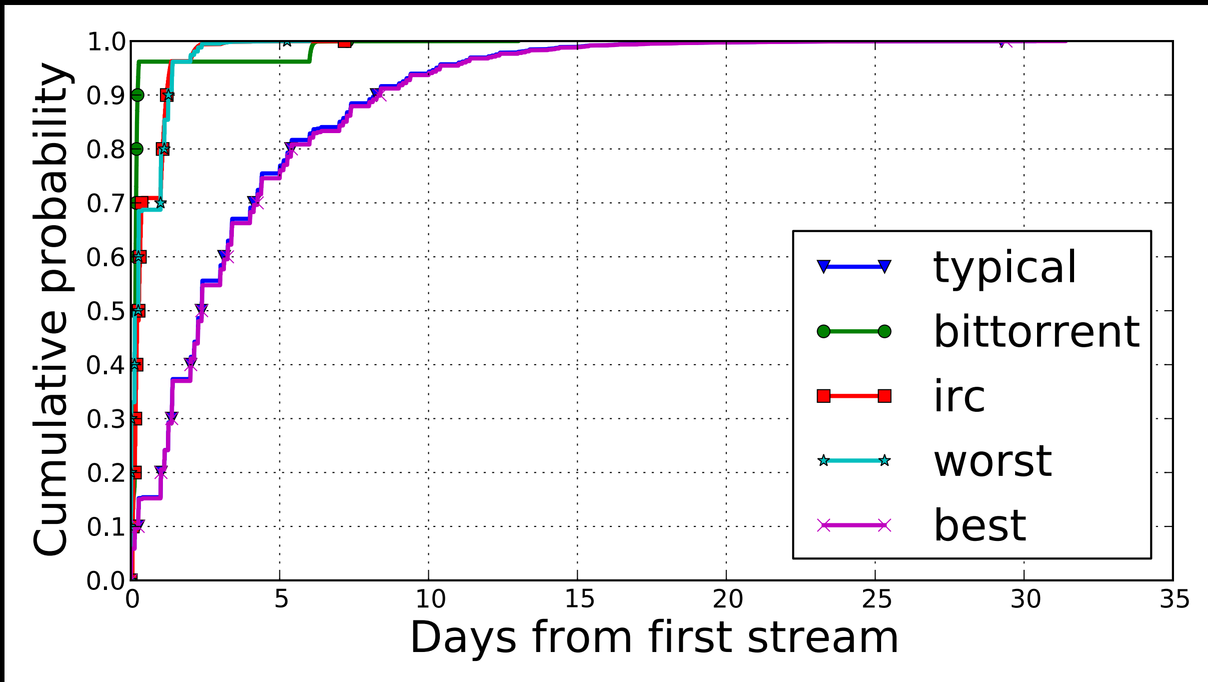

Time to first compromised exit, 10/12 – 3/13

Fraction compromised exits, 10/12 – 3/13

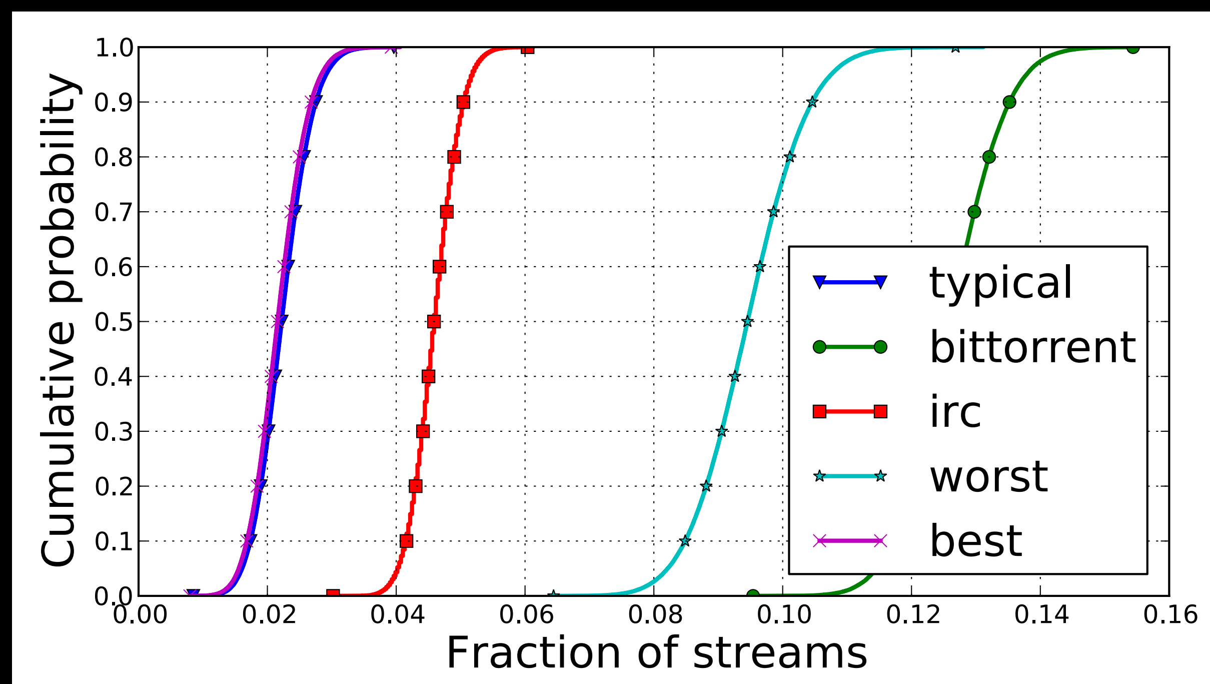

Node Adversary Results

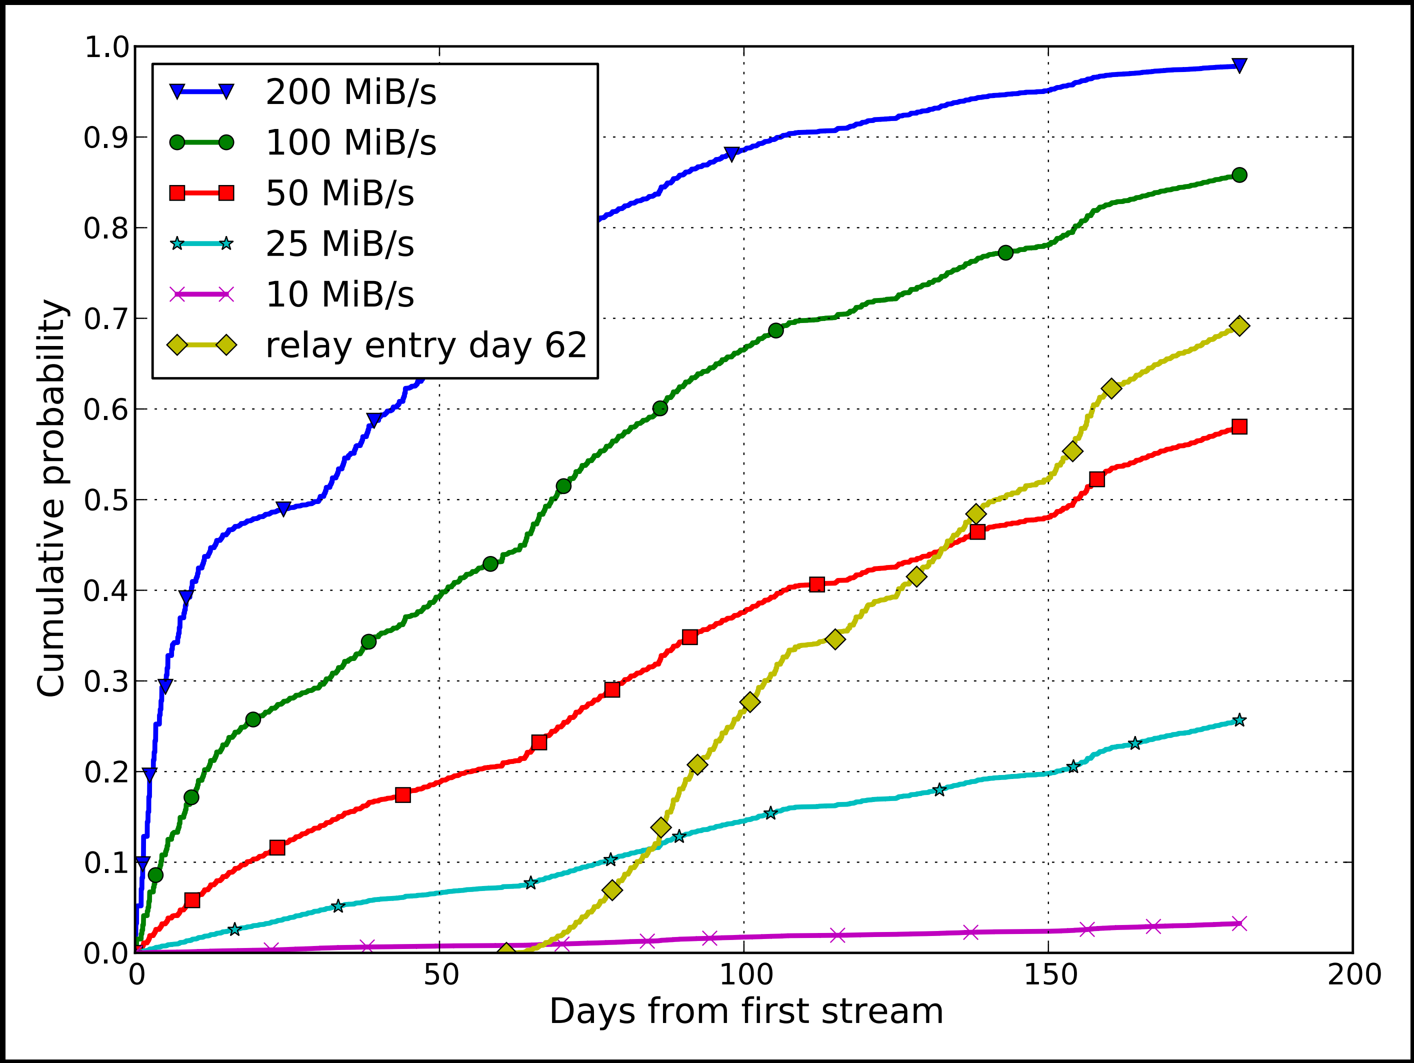

Time to first compromised circuit, 10/12-3/13

“Users Get Routed” Outline Summary

- Tor Security Analysis
  - Adversary Framework
  - Security Metrics
  - Evaluation Methodology
  - Node Adversary Analysis
  - Link Adversary Analysis

Link Adversary

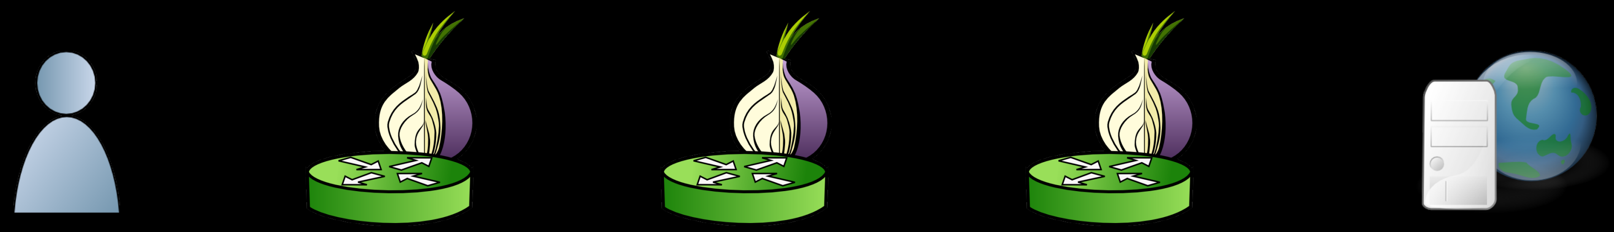

Link Adversary

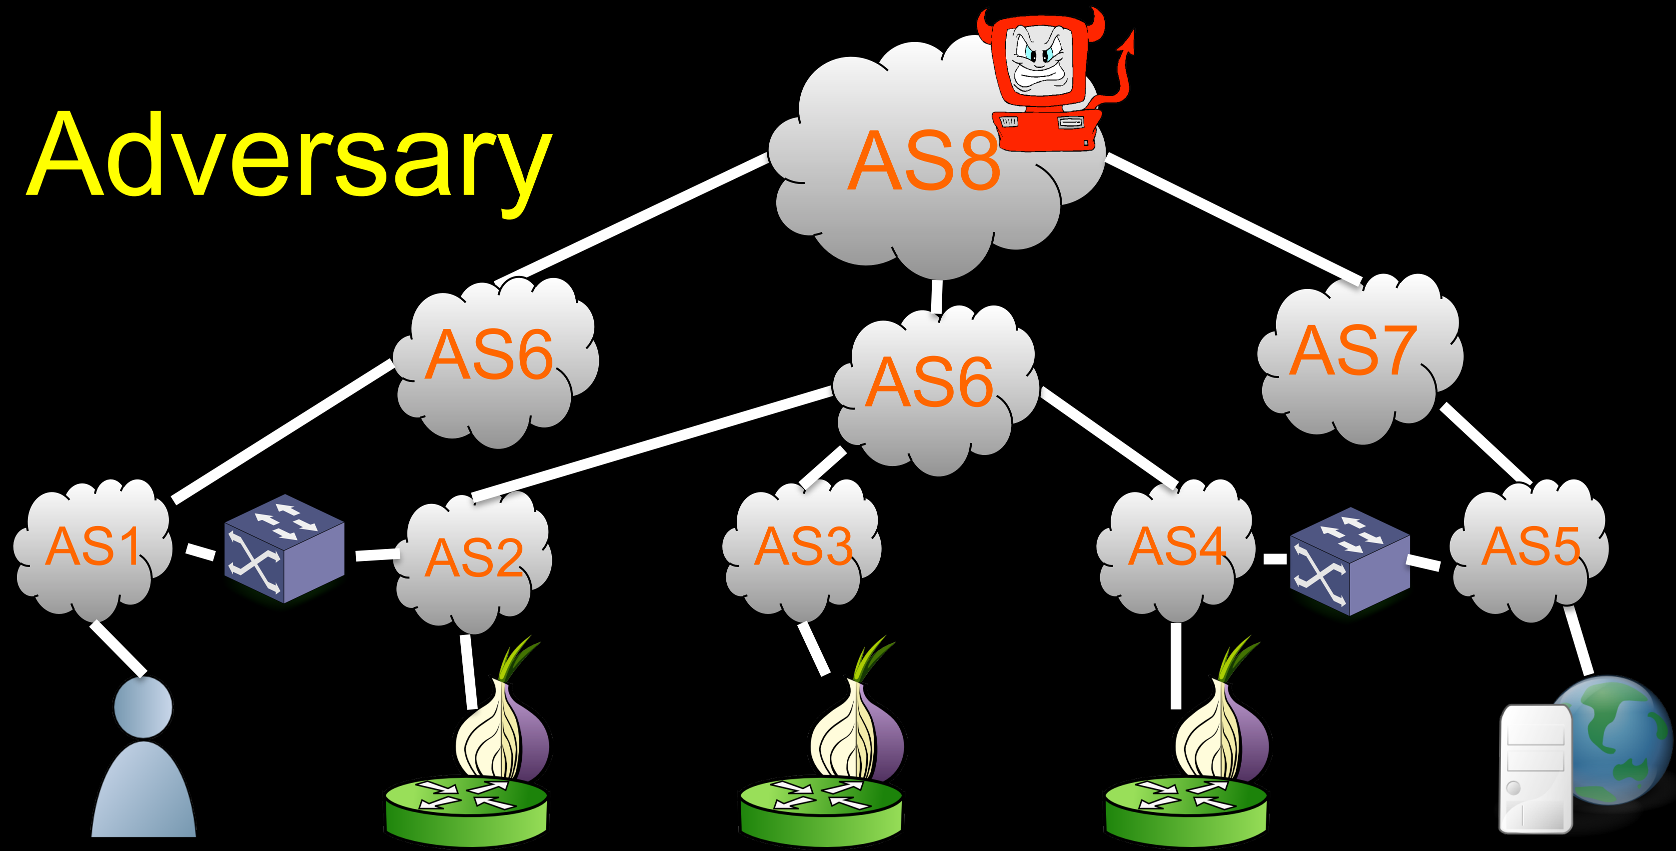

1. Autonomous Systems (Ases)
2. Internet Exchange Points (IXPs)
3. Adversary has fixed location (unlike Path Independence)
4. Adversary may control multiple entities
   – "Top" ASes
   – IXP organizations

Link Adversary

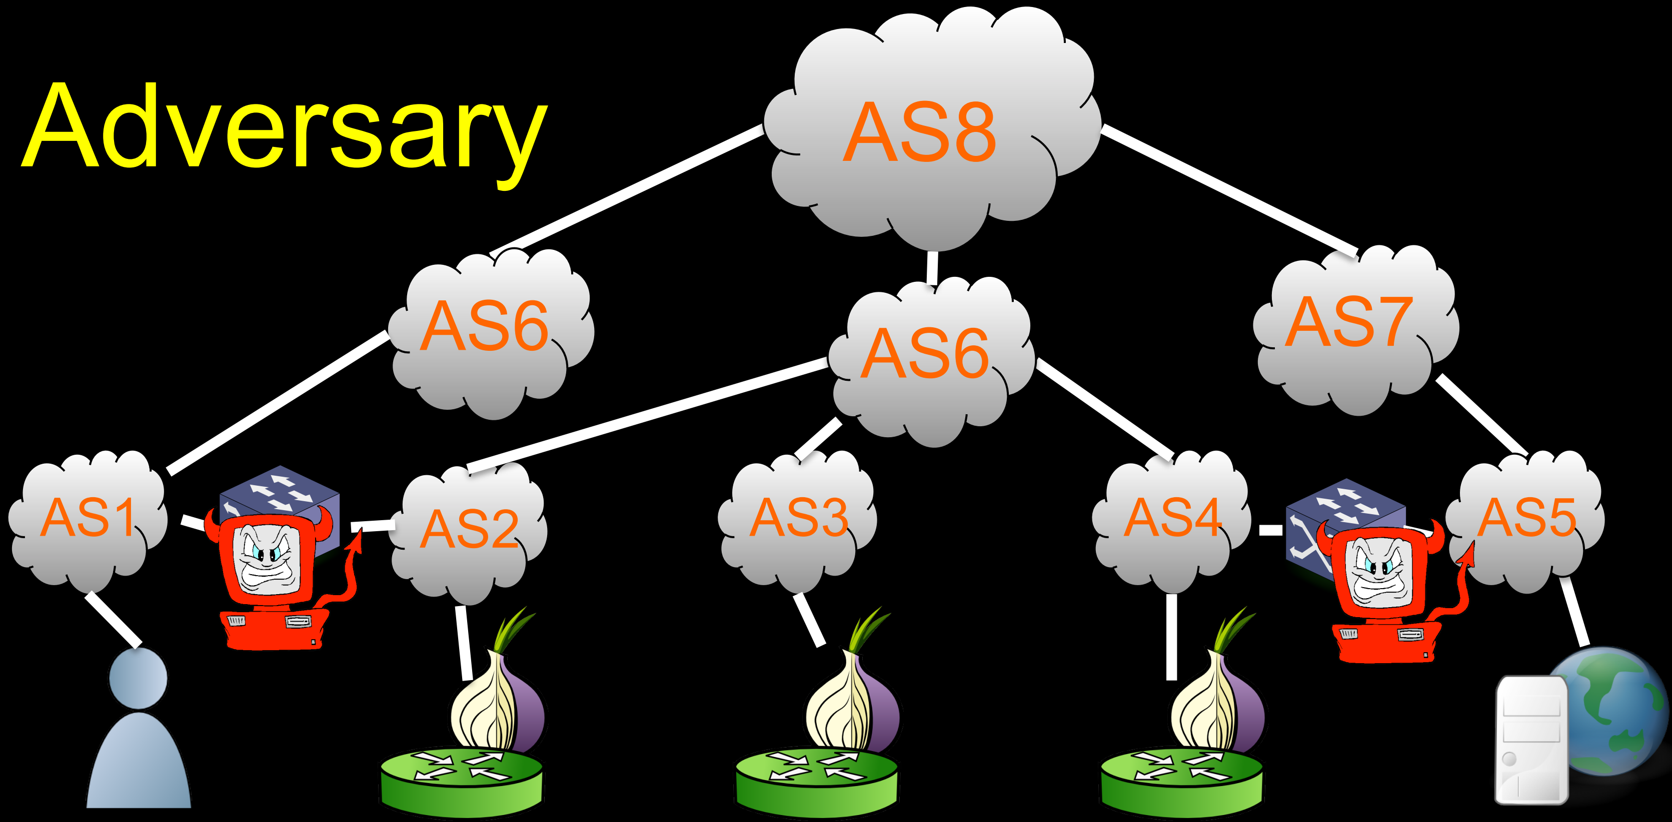

1. Autonomous Systems (Ases)
2. Internet Exchange Points (IXPs)
3. Adversary has fixed location (unlike Path Independence)
4. Adversary may control multiple entities
   – "Top" ASes
   – IXP organizations

Link Adversary

Client locations

- Top 5 non-Chinese source ASes in Tor (Edman&Syverson 09)

AS#	Description	Country
3320	Deutsche Telekom AG	Germany
3209	Arcor	Germany
3269	Telecom Italia	Italy
13184	HanseNet Telekommunikation	Germany
6805	Telefonica Deutschland	Germany

AS/IXP Locations

- Ranked for client location by frequency on entry or exit paths
- Exclude src/dst ASes
- Top k ASes /top IXP organization

Type	ID	Description
AS	3356	Level 3 Communications
AS	1299	TeliaNet Global
AS	6939	Hurricane Electric
IXP	286	DE-CIX Frankfurt
IXP Org.	DE-CIX	DE-CIX

Example: Adversary locations for BitTorrent client in AS 3320

Link Adversary

IXP organizations obtained by manual clustering based on PeerDB and PCH.

#	IXP Organization	Size	Country
1	Equinix	26	global
2	PTTMetro	8	Brazil
3	PIPE	6	Australia
4	NIXI	6	India
5	XChangePoint	5	global
6	MAE/VERIZON	5	global
7	Netnod	5	Sweden
8	Any2	4	US
9	PIX	4	Canada
10	JPNAP	3	Japan
11	DE-CIX	2	Germany
12	AEPROVI	2	Equador
13	Vietnam	2	Vietnam
14	NorthWestIX	2	Montana, US
15	Terremark	2	global
16	Telx	2	US
17	NorrNod	2	Sweden
18	ECIX	2	Germany
19	JPIX	2	Japan

IXP organizations ranked by size

Link Adversary

Adversary controls one AS,
Time to first compromised stream,
1/13 – 3/13
"Best": most secure client AS
"Worst": least secure client AS

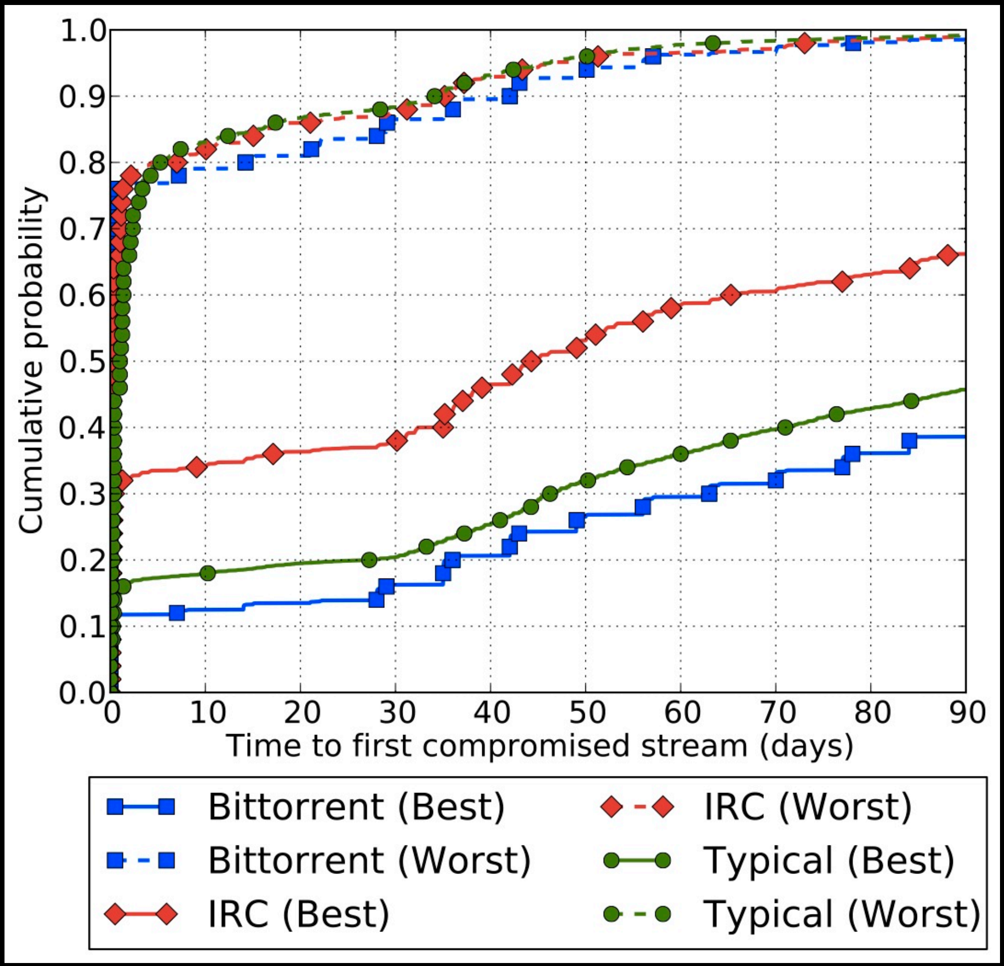

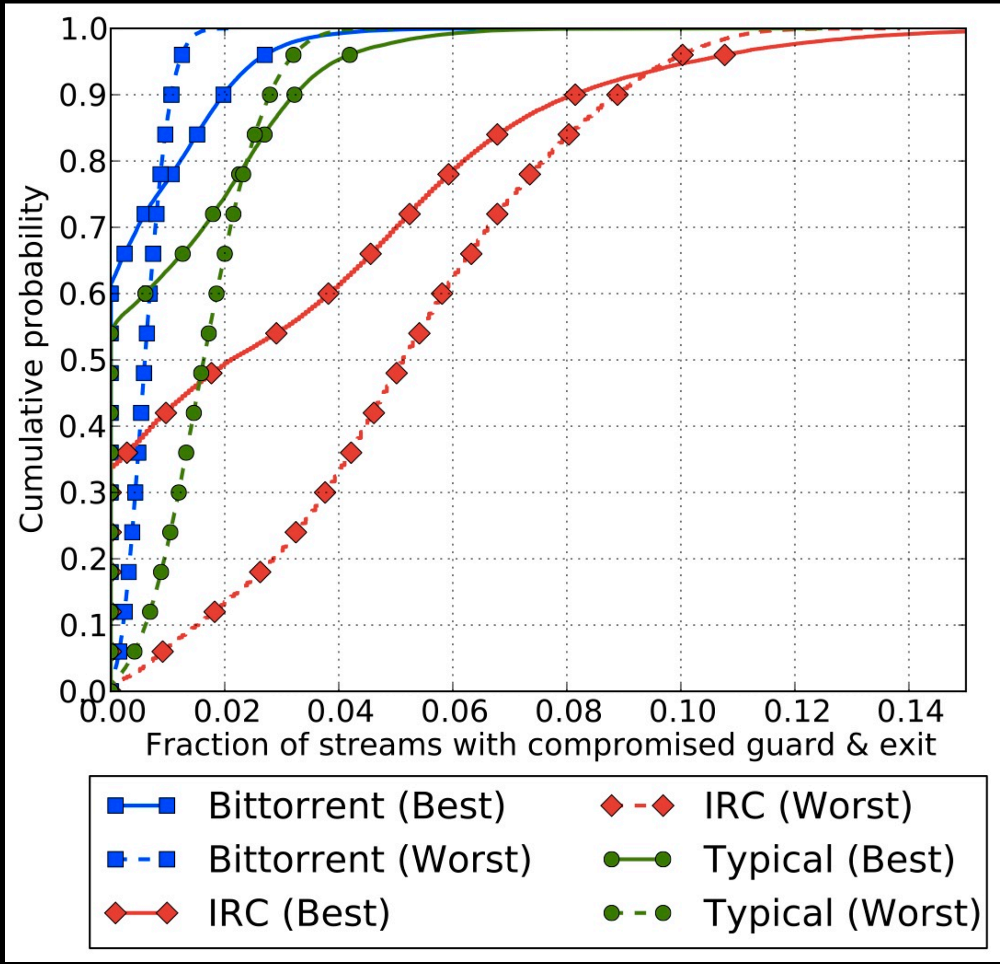

Adversary controls one AS,
Fraction comp. streams, 1/13 – 3/13
"Best": most secure client AS
"Worst": least secure client AS

Link Adversary

Adversary controls IXP organization,
Time to first compromised stream,
1/13 – 3/13,
“Best”: most secure client AS
“Worst”: least secure client AS

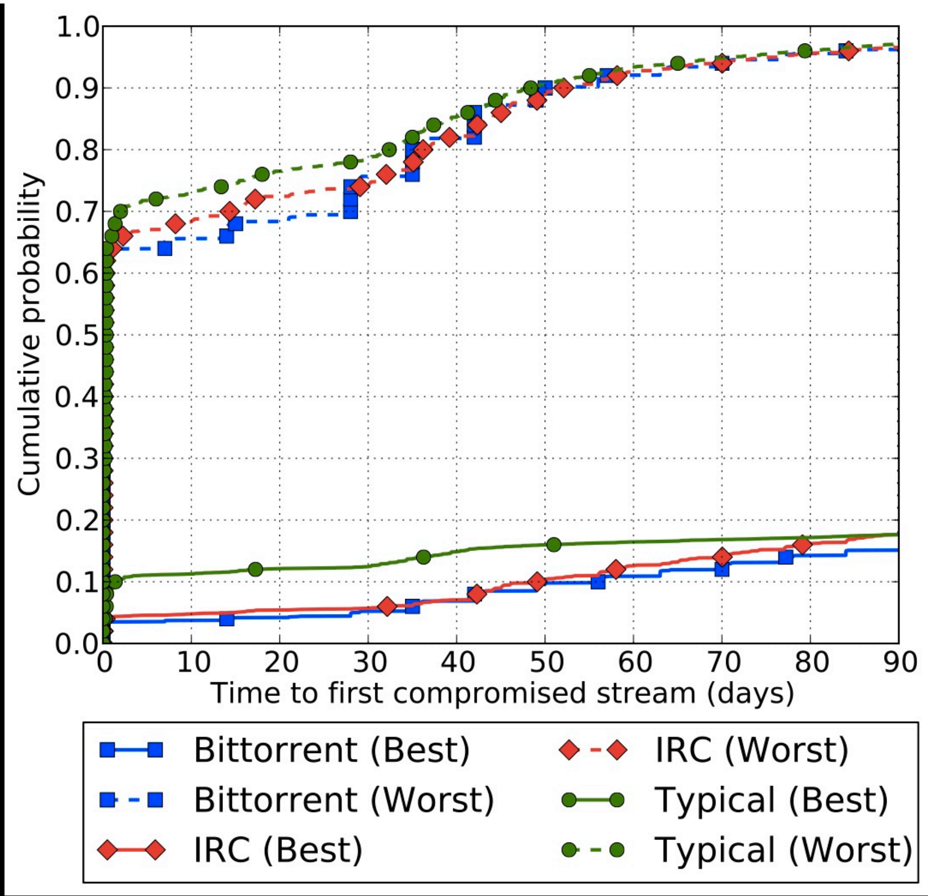

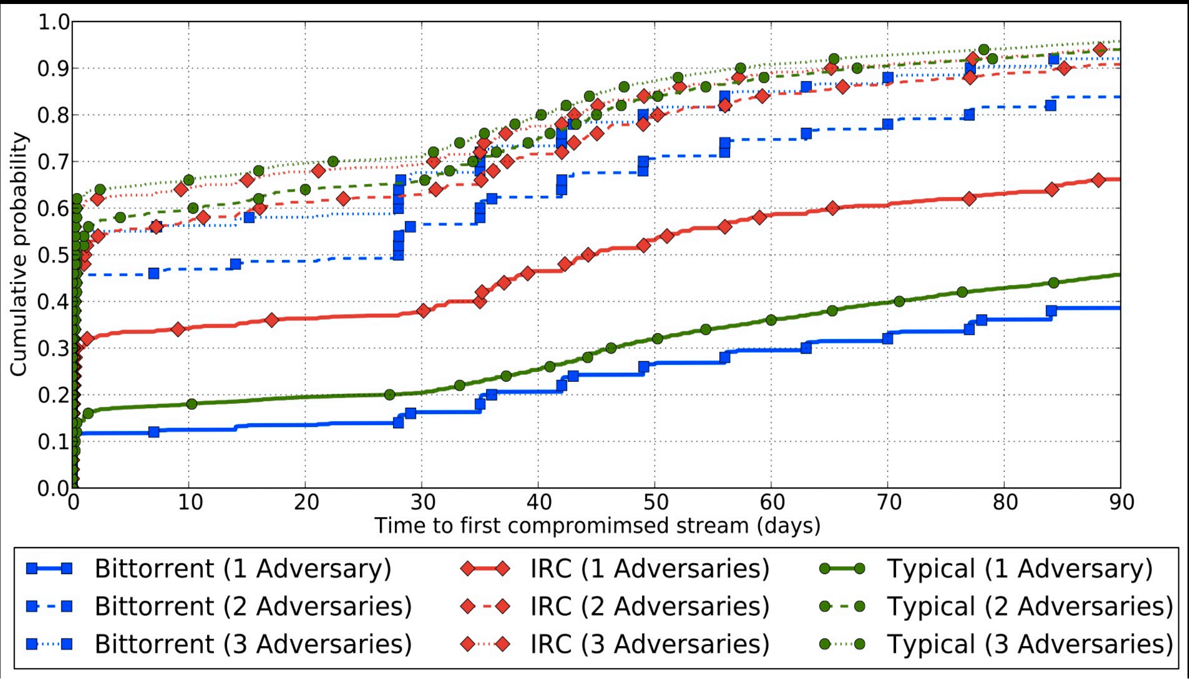

Adversary controls top ASes,
Time to first compromised stream, 1/13 – 3/13,
Only “best” client AS

“Users get routed”: bad news summary

- 80% of all types of users may be deanonymized by moderate Tor-relay adversary within 6 months
- Bittorrent user by far worst off for fraction of connections compromised by Tor-relay adversary
- Against a single-AS adversary roughly 100% of users in some common locations are deanonymized within three months
- (or 95% in 3 months for a single IXP)
- 2-AS adversary reduces median time to the first client deanonymization by an order of magnitude:
  - from over 3 months to only 1 day for typical web user
  - from over 3 months to c. 1 month for a BitTorrent user

Course Outline

- Lecture 1: Basics and Formalization
  - Usage examples, basic notions of traffic-secure communications, mixes and onion routers
  - Onion routing design basics: circuit construction protocols, network discovery
  - Formalization and analysis, possibilistic and probabilistic definitions of anonymity
- Lecture 2: Security for the real world
  - Simple demo of obtaining/using Tor
  - Security of obtain/using Tor
  - Adding network link awareness
  - Importance of modeling users
  - Importance of realistic and practical
    - Adversary models
    - Security definitions

Where to turn for further information

- Anonymity bibliography: http://freehaven.net/anonbib/
  - Best general source for papers on anonymous communication.
  - Strangely, many original onion routing developments not there so…
- Early onion routing publications list: http://www.onion-router.net/Publications.html
  - See also http://www.onion-router.net/History.html
  - And since it is not on any of the above lists, history https://www.acsac.org/2011/program/keynotes/syverson.pdf
- Privacy Enhancing Technologies Symposium: https://petsymposium.org/
  - Primary venue for research publications on anonymous communication and primary annual confluence of anonymous comms researchers
- My personal homepage: http://www.syverson.org/
  - Updated too infrequently, but has some useful relevant links

What to do if adversary can observe much of the network?

- Nation-state network observer
- Botnet or nation-state running many relays

What to do if adversary can observe much of the network?

- Nation-state network observer
- Botnet or nation-state running many relays

Come to the Stafford Tavares Lecture on Thursday for some analysis and possible answers.