U.S. NAVAL RESEARCH LABORATORY

Toward Safer Tor Research

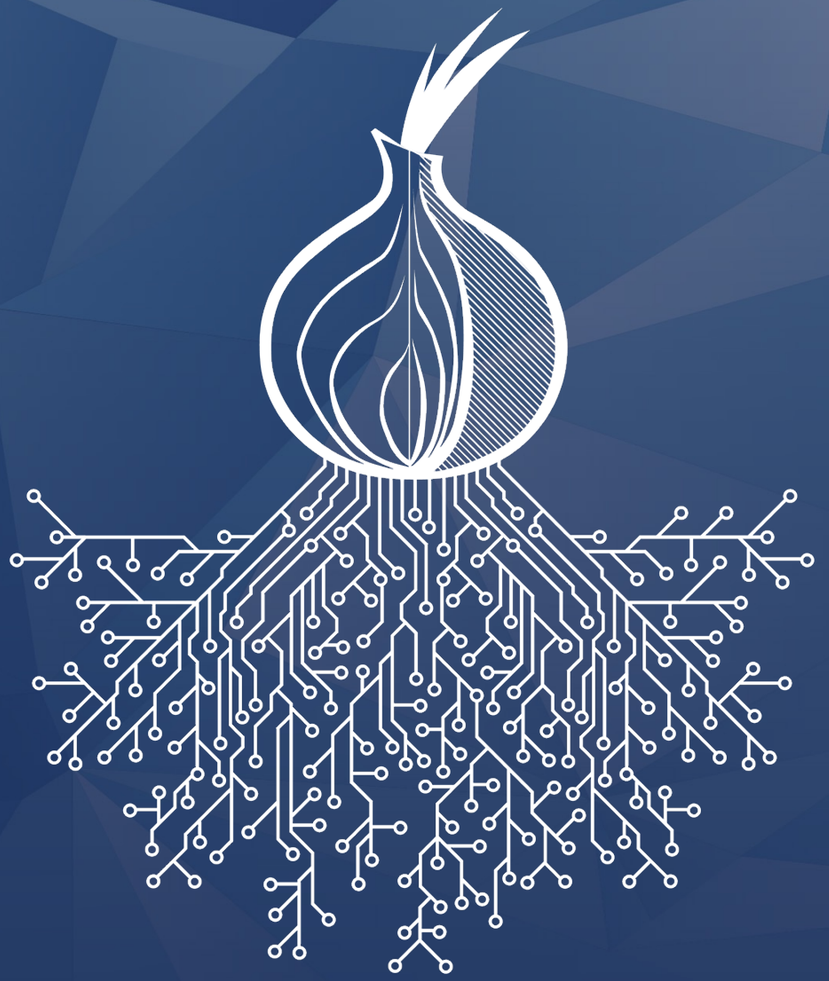

Rob Jansen, U.S. Naval Research Laboratory

Rob Jansen, Ph.D.
Computer Security Research Scientist
Center for High Assurance Computer Systems
U.S. Naval Research Laboratory

Workshop on Ethics in Computer Security
European Symposium on Security and Privacy
Genoa, Italy
June 10th, 2022

U.S. NAVAL RESEARCH LABORATORY

What is Tor?

Anonymous Communication

- Separates identification from routing
- Provides unlinkable communication
- Protects user privacy and safety online

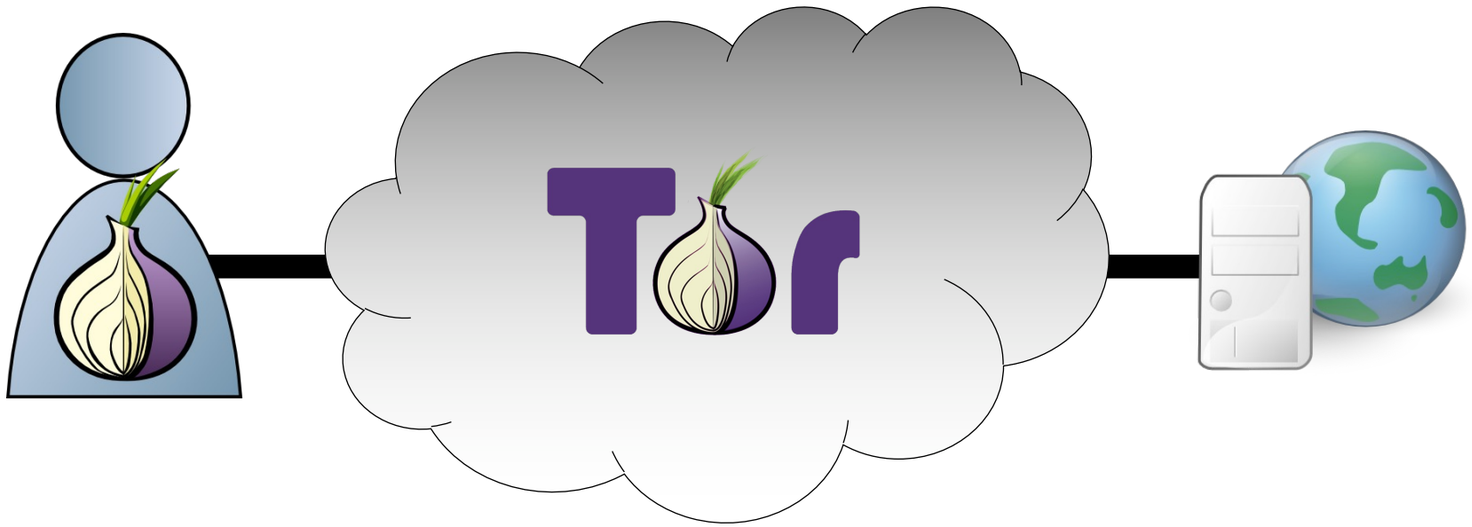

Tor is Popular

- ~2-8 million daily active users
- ~7,000 volunteer relays
- Transferring ~300 Gbit/s

U.S. NAVAL RESEARCH LABORATORY

Why is Tor Used?

Block Trackers

- isolate each website you visit so third-party trackers and ads can't follow you

Defend against surveillance

- prevent someone watching your connection from knowing what websites you visit

Resist Fingerprinting

- all users look the same, making it difficult to be fingerprinted based on browser/device

Browse Freely

- free to access sites that your local network may have blocked

https://www.torproject.org

Who Uses Tor?

Normal People	Journalists	Law Enforcement
Activists	Business Executives	Bloggers
Militaries	IT Professionals	Whistleblowers

https://2019.www.torproject.org/about/torusers.html.en

U.S. NAVAL RESEARCH LABORATORY

Who Uses Tor?

Normal People		ement
Activists	B	rs
Militaries	IT Professionals	Whistleblowers

- Protect privacy from marketers and identity thieves
- Protect comms from irresponsible corporations
- Protect their children online
- Research sensitive topics
- Skirt surveillance
- Circumvent censorship

https://2019.www.torproject.org/about/torusers.html.en

U.S. NAVAL RESEARCH LABORATORY

Who Uses Tor?

Normal People	Journalists	Law Enforcement
Activists	Business Exe	
Militaries	IT Professionals	Whistleblowers

- Online surveillance
- Sting operations
- Truly anonymous tip lines

https://2019.www.torproject.org/about/torusers.html.en

U.S. NAVAL RESEARCH LABORATORY

Who Uses Tor?

Normal People	Journalists	Law Enforcement
Activists	Business Executives	Bloggers
Militaries		

- Security breach information clearinghouses
- Seeing your competition as your market does
- Keeping strategies confidential
- Accountability

U.S. NAVAL RESEARCH LABORATORY

Who Uses Tor?

Normal People		forcement
Activists		ggers
Militaries	IT Professionals	Whistleblowers

- To verify IP based firewall rules
- To bypass their own security systems
- To access internet resources
- To work around ISP network outages
- To connect back to deployed services

https://2019.www.torproject.org/about/torusers.html.en

User Safety in Tor

Tor is a Privacy-Preserving Anonymous Communication Network in which User Safety is of Highest Priority

Outline

- **Tor Safety Board overview**
- **Research areas relevant to safety**
  - Simulating Tor with Shadow
  - Measuring Tor with PrivCount
  - Safe research applications
- **Ongoing struggles**

Outline

- **Tor Safety Board overview**
- **Research areas relevant to safety**
  - Simulating Tor with Shadow
  - Measuring Tor with PrivCount
  - Safe research applications
- **Ongoing struggles**

What is the Tor Research Safety Board?

A group of Tor researchers who want to **minimize privacy risks while fostering a better understanding of the Tor network and its users**

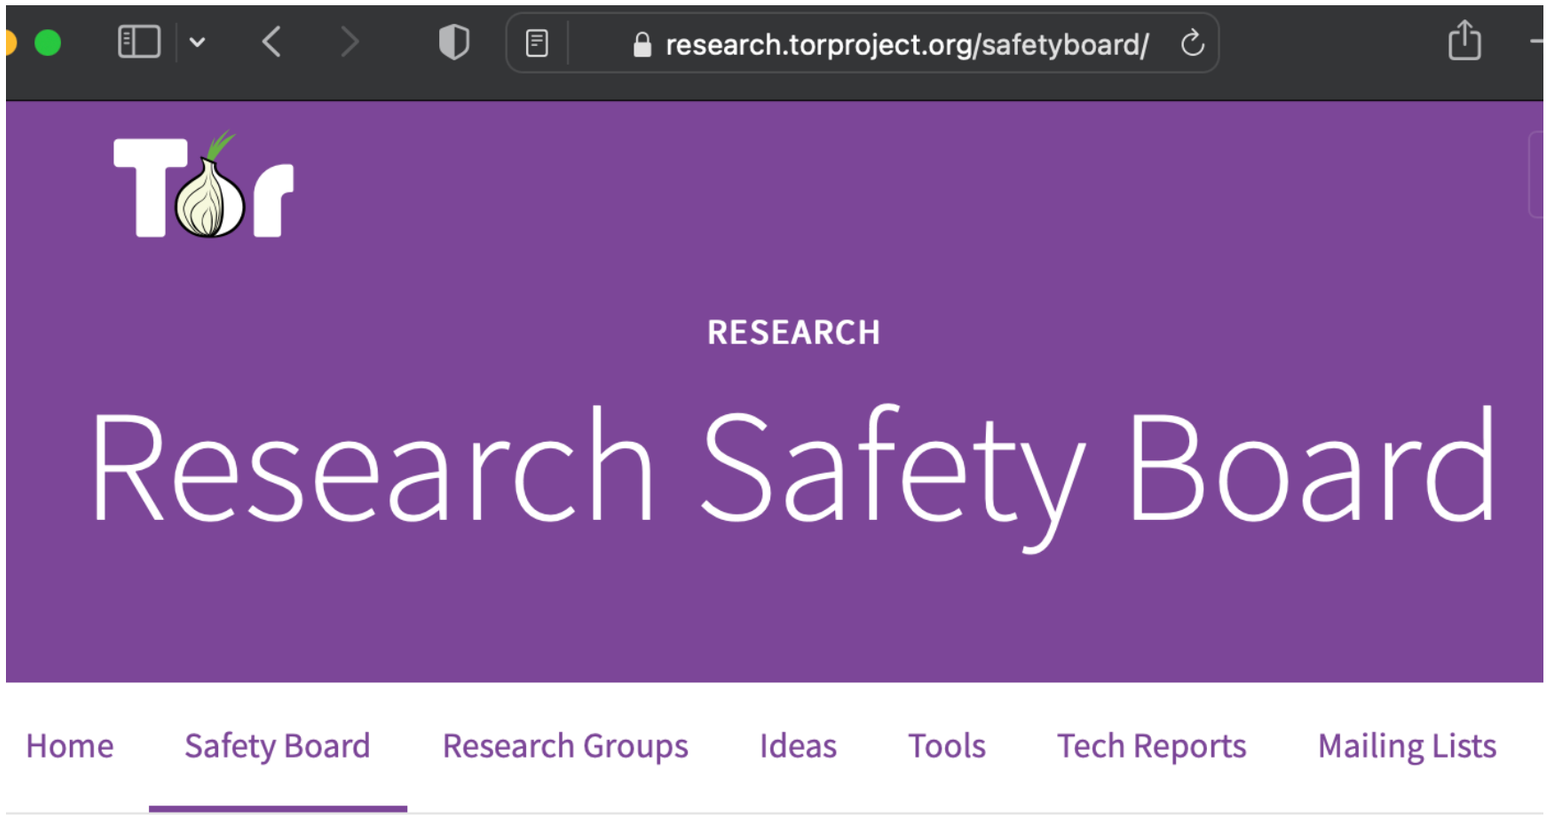

U.S. NAVAL RESEARCH LABORATORY

What is the Tor Research Safety Board?

A group of Tor researchers who want to **minimize privacy risks while fostering a better understanding of the Tor network and its users**

Main activities of the board:

1. develop and maintain a set of guidelines that researchers can use to assess the safety of their Tor research
2. give feedback to researchers who use our guidelines to assess the safety of their planned research
3. teach program committees about our guidelines, and encourage reviewers to consider research safety when reviewing Tor papers

U.S. NAVAL RESEARCH LABORATORY

What the Tor Research Safety Board is NOT

- The Safety Board is NOT a replacement for your IRB!
- You must still follow your organization's requirements for IRB
- We are not an ethics board
  - We help you think about how to make your work safer
  - We do not approve/deny research
  - We hope your safety analysis will be added to your paper

U.S. NAVAL RESEARCH LABORATORY

Tor Research Safety Board Guidelines

1. Use a test Tor network whenever possible.
2. It’s safest to only attack yourself / your own traffic.

U.S. NAVAL RESEARCH LABORATORY

Tor Research Safety Board Guidelines

1. Use a test Tor network whenever possible.
2. It’s safest to only attack yourself / your own traffic.
3. Only collect data that is safe to make public.
4. Don’t collect data you don’t need (minimization).
5. Limit the granularity of data (e.g. use bins or add noise).

U.S. NAVAL RESEARCH LABORATORY

Tor Research Safety Board Guidelines

1. Use a test Tor network whenever possible.
2. It's safest to only attack yourself / your own traffic.
3. Only collect data that is safe to make public.
4. Don't collect data you don't need (minimization).
5. Limit the granularity of data (e.g. use bins or add noise).
6. Take reasonable security precautions, e.g. about who has access to your data sets or experimental systems.

U.S. NAVAL RESEARCH LABORATORY

Tor Research Safety Board Guidelines

1. Use a test Tor network whenever possible.
2. It's safest to only attack yourself / your own traffic.
3. Only collect data that is safe to make public.
4. Don't collect data you don't need (minimization).
5. Limit the granularity of data (e.g. use bins or add noise).
6. Take reasonable security precautions, e.g. about who has access to your data sets or experimental systems.
7. The benefits should outweigh the risks.
8. Consider auxiliary data (e.g. third-party data sets) when assessing the risks.
9. Consider whether the user meant for that data to be private.

U.S. NAVAL RESEARCH LABORATORY

Questions to Guide Feedback Request

1. What are you trying to learn, and why is that useful for the world? That is, what are the hoped-for benefits of your experiment?

U.S. NAVAL RESEARCH LABORATORY

Questions to Guide Feedback Request

1. What are you trying to learn, and why is that useful for the world? That is, what are the hoped-for benefits of your experiment?

2. What exactly is your plan? That is, what are the steps of your experiment, what will you collect, how will you keep it safe, and so on.

U.S. NAVAL RESEARCH LABORATORY

Questions to Guide Feedback Request

1. What are you trying to learn, and why is that useful for the world? That is, what are the hoped-for benefits of your experiment?

2. What exactly is your plan? That is, what are the steps of your experiment, what will you collect, how will you keep it safe, and so on.

3. What attacks or risks might be introduced or assisted because of your actions or your data sets, and how well do you resolve each of them? Use the “safety guidelines” above to help in the brainstorming and analysis.

U.S. NAVAL RESEARCH LABORATORY

Questions to Guide Feedback Request

1. What are you trying to learn, and why is that useful for the world? That is, what are the hoped-for benefits of your experiment?

2. What exactly is your plan? That is, what are the steps of your experiment, what will you collect, how will you keep it safe, and so on.

3. What attacks or risks might be introduced or assisted because of your actions or your data sets, and how well do you resolve each of them? Use the “safety guidelines” above to help in the brainstorming and analysis.

4. Walk us through why the benefits from item 1 outweigh the remaining risks from item 3: why is this plan worthwhile despite the remaining risks?

Submitting your Feedback Request

- HotCRP instance: https://safetyboard.torproject.net/submit
- Submit, then assigned board reviewers who provide written feedback
- ~35 submissions in last ~5 years
- Examples: https://research.torproject.org/safetyboard/

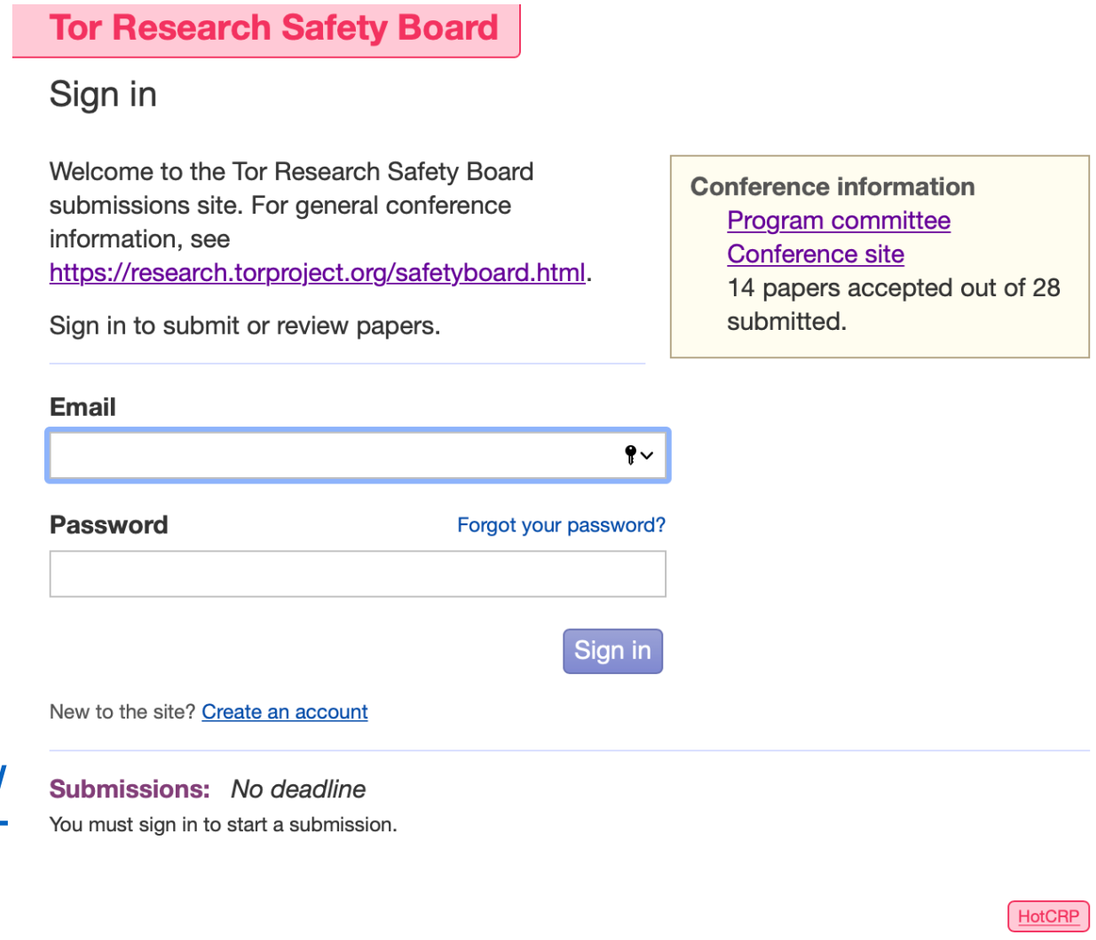

Outline

- **Tor Safety Board overview**
- **Research areas relevant to safety**
  - Simulating Tor with Shadow
  - Measuring Tor with PrivCount
  - Safe research applications
- **Ongoing struggles**

U.S. NAVAL RESEARCH LABORATORY

Test Tor Networks

- Completely private test network
- 100% safe: absolutely no safety or privacy risks
- Enables evaluation of attacks that potentially harm Tor or its users:
  - Website fingerprinting
  - End-to-end correlation
  - Other traffic analysis
  - Denial of service
  - Protocol attacks

Safety Guideline 1:

Use a test Tor network whenever possible

U.S. NAVAL RESEARCH LABORATORY

Contributing Research

- **Co-opting Linux Processes for High-Performance Network Simulation**. Rob Jansen, Jim Newsome, and Ryan Wails. USENIX Annual Technical Conference, 2022.
- **Once is Never Enough: Foundations for Sound Statistical Inference in Tor Network Experimentation**. Rob Jansen, Justin Tracey, and Ian Goldberg. USENIX Security, 2021.
- **High Performance Tor Experimentation from the Magic of Dynamic ELFs**. Justin Tracey, Rob Jansen, and Ian Goldberg. CSET, 2018.
- **Shadow-Bitcoin: Scalable Simulation via Direct Execution of Multi-threaded Applications**. Andrew Miller and Rob Jansen. CSET, 2015.
- **Methodically Modeling the Tor Network.** Rob Jansen, Kevin Bauer, Nick Hopper, and Roger Dingledine. CSET 2012.
- **Shadow: Running Tor in a Box for Accurate and Efficient Experimentation**. Rob Jansen and Nicholas Hopper. NDSS, 2012.

U.S. NAVAL RESEARCH LABORATORY

Network Experimentation Requirements

- The most important property of test networks:
  - Controllable: isolate important factors
  - Replicable: achievable with determinism
- Requirements for large distributed systems (e.g., Tor)
  - Accurate: directly execute system software (not an abstraction)
  - Scalable: can run studied system at scale
- Shadow
  - Network simulator with above design goals
  - Open source: https://shadow.github.io

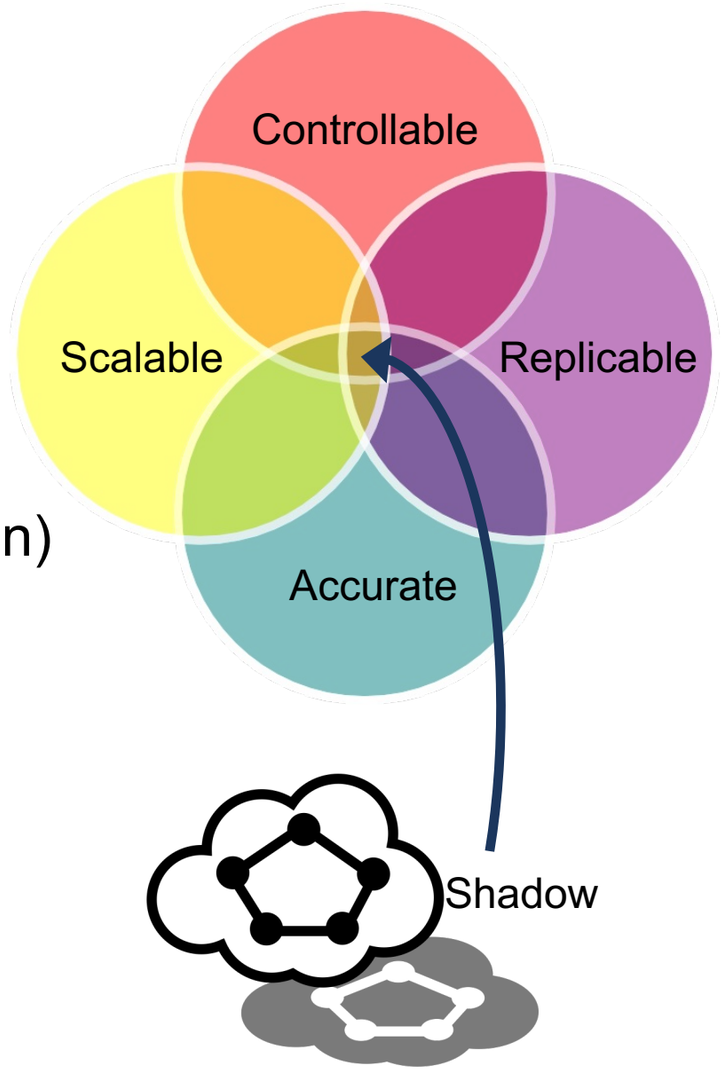

U.S. NAVAL RESEARCH LABORATORY

What is Shadow?

- A parallel, discrete-event, packet-level, hybrid network simulator/emulator

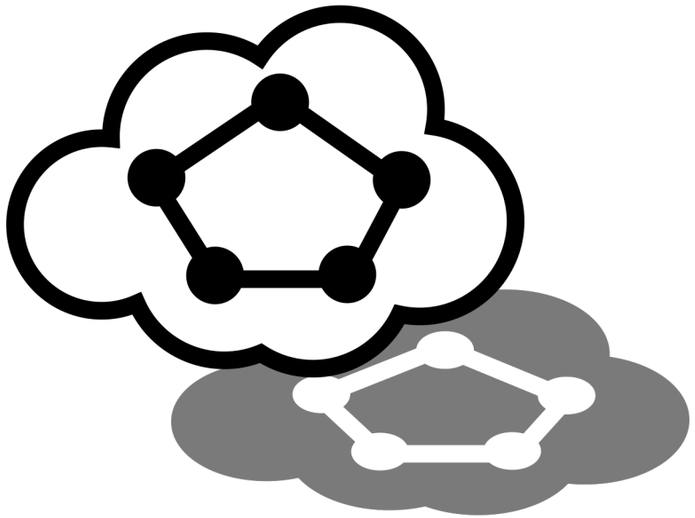

Direct Execution

- Directly executes applications as standard Linux processes

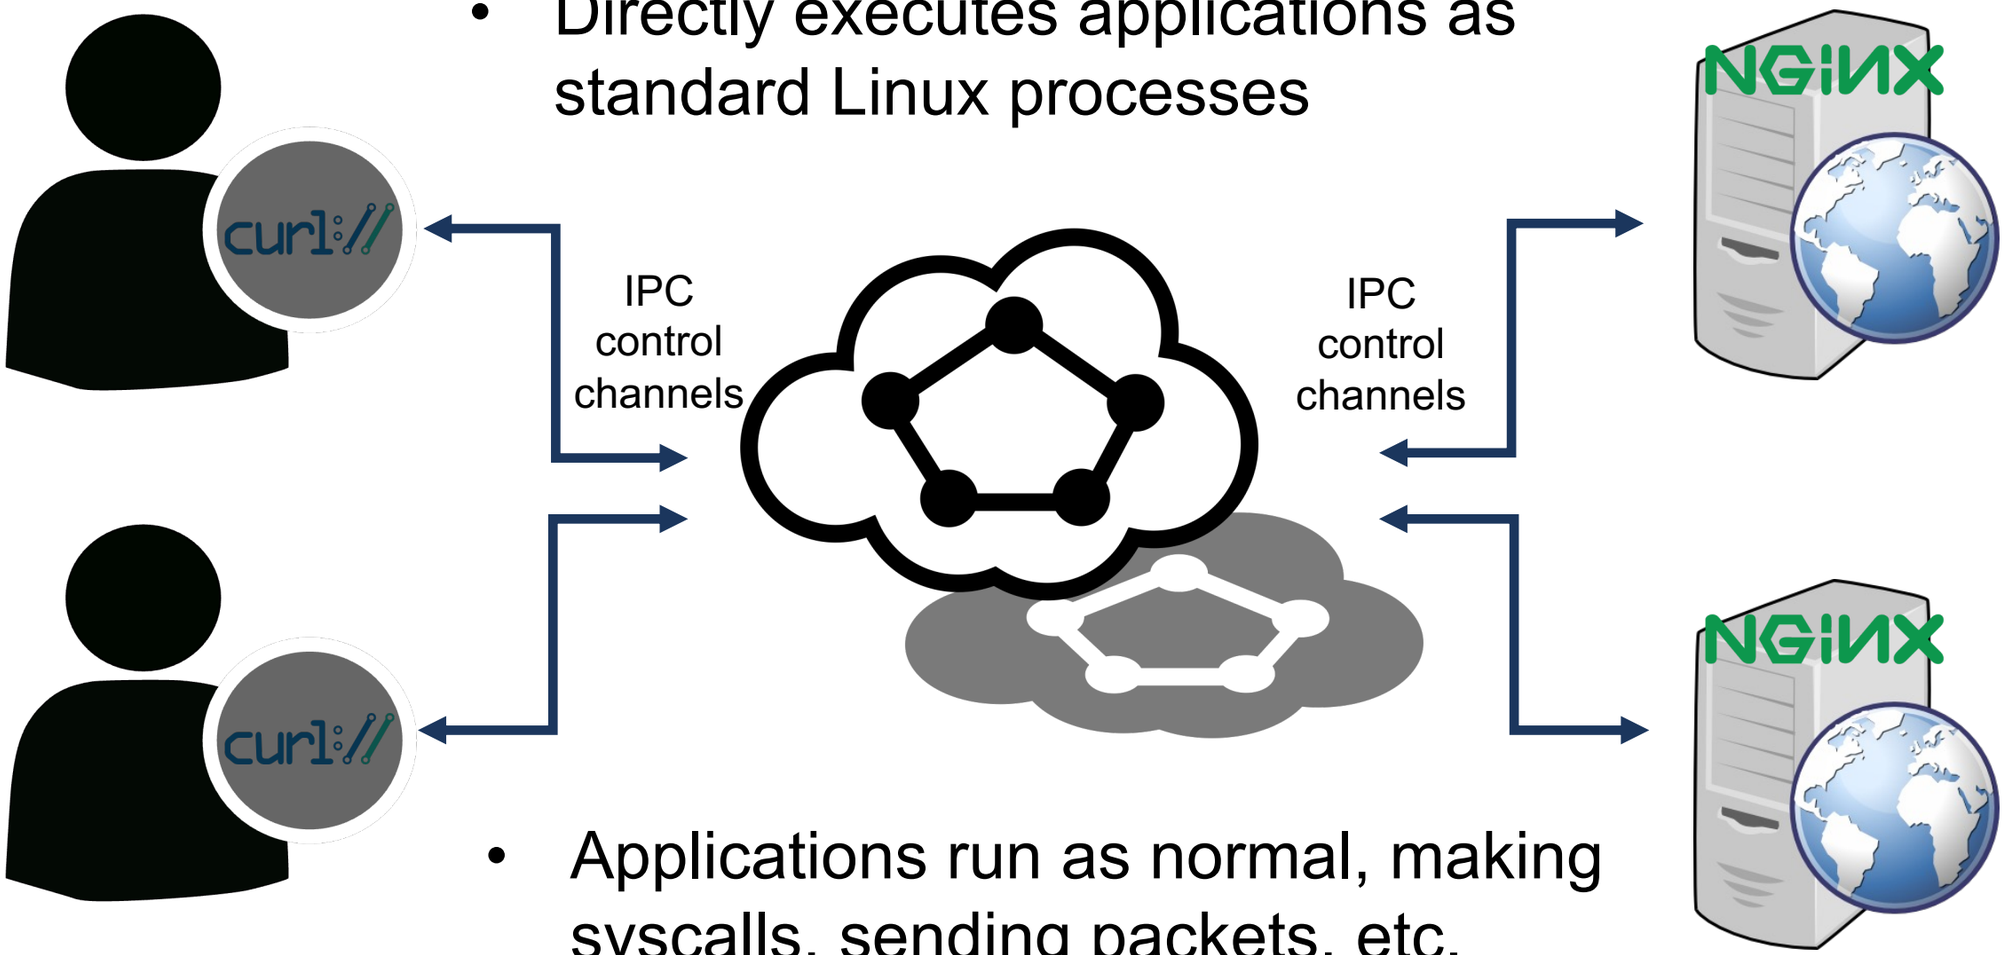

- Applications run as normal, making syscalls, sending packets, etc.

U.S. NAVAL RESEARCH LABORATORY

Processes in Linux

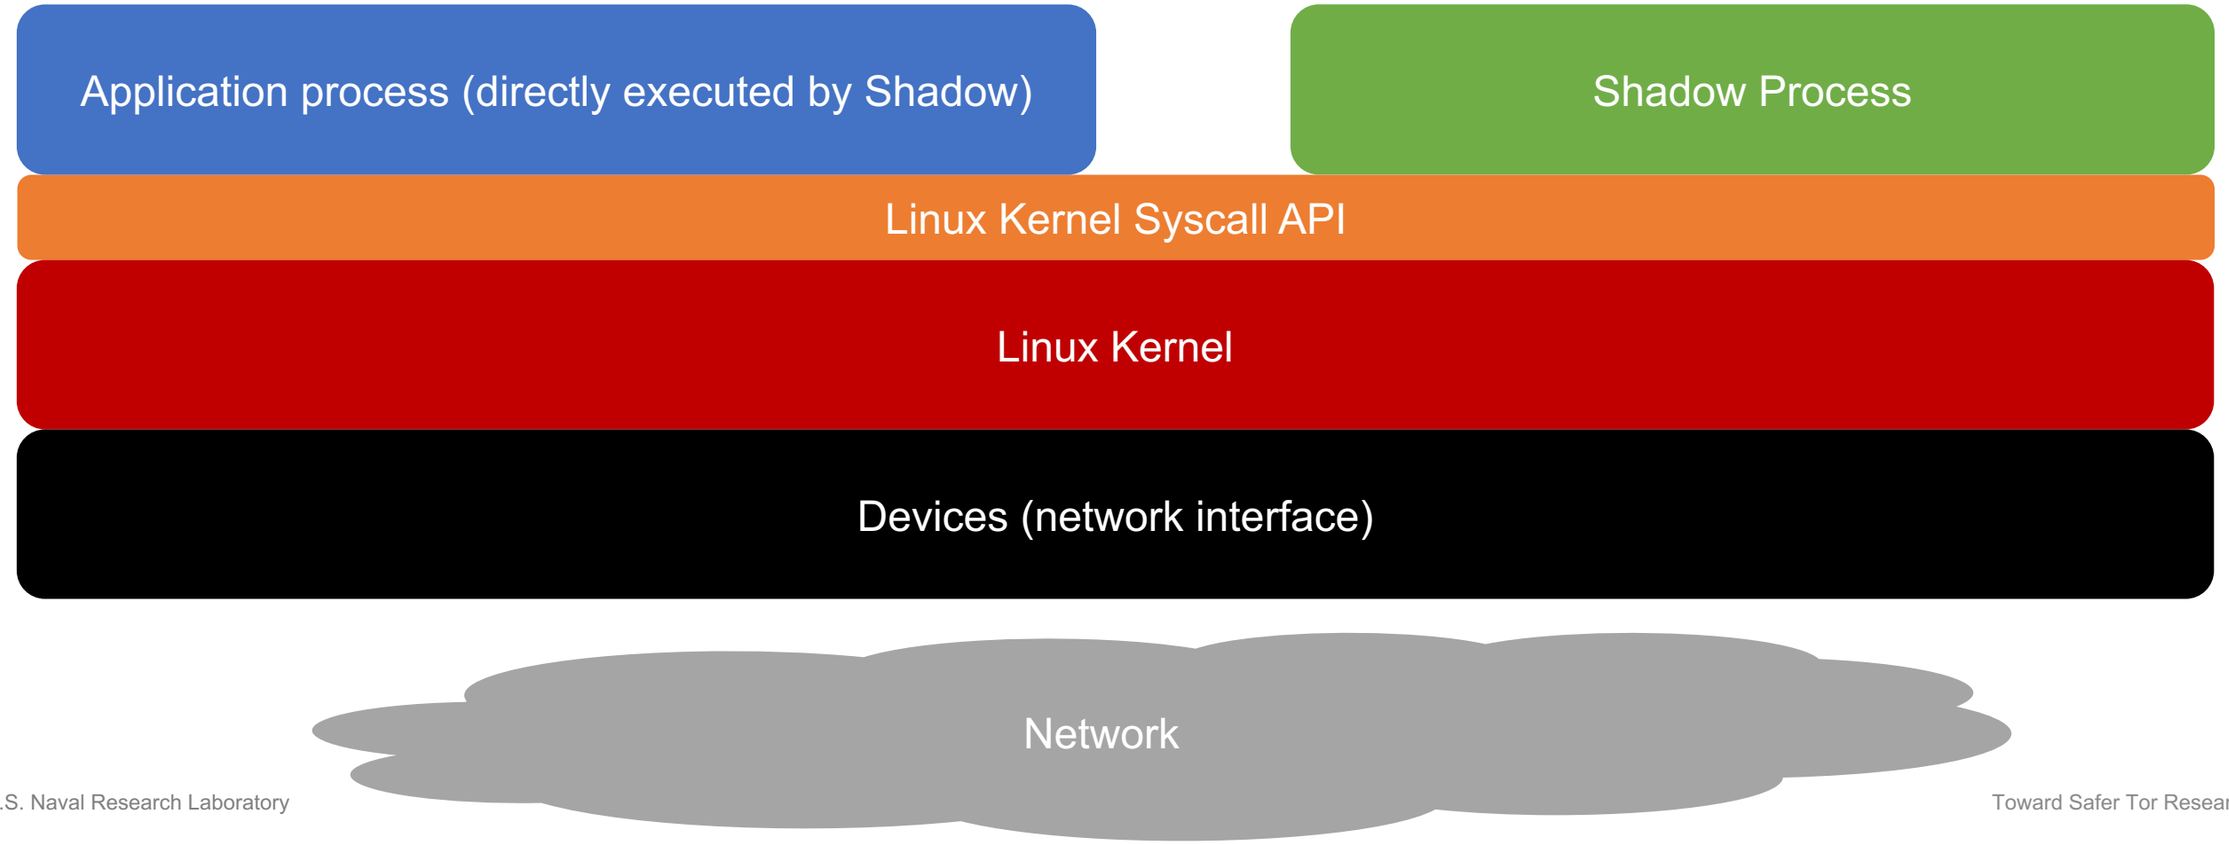

U.S. NAVAL RESEARCH LABORATORY

Resolving Library Symbols

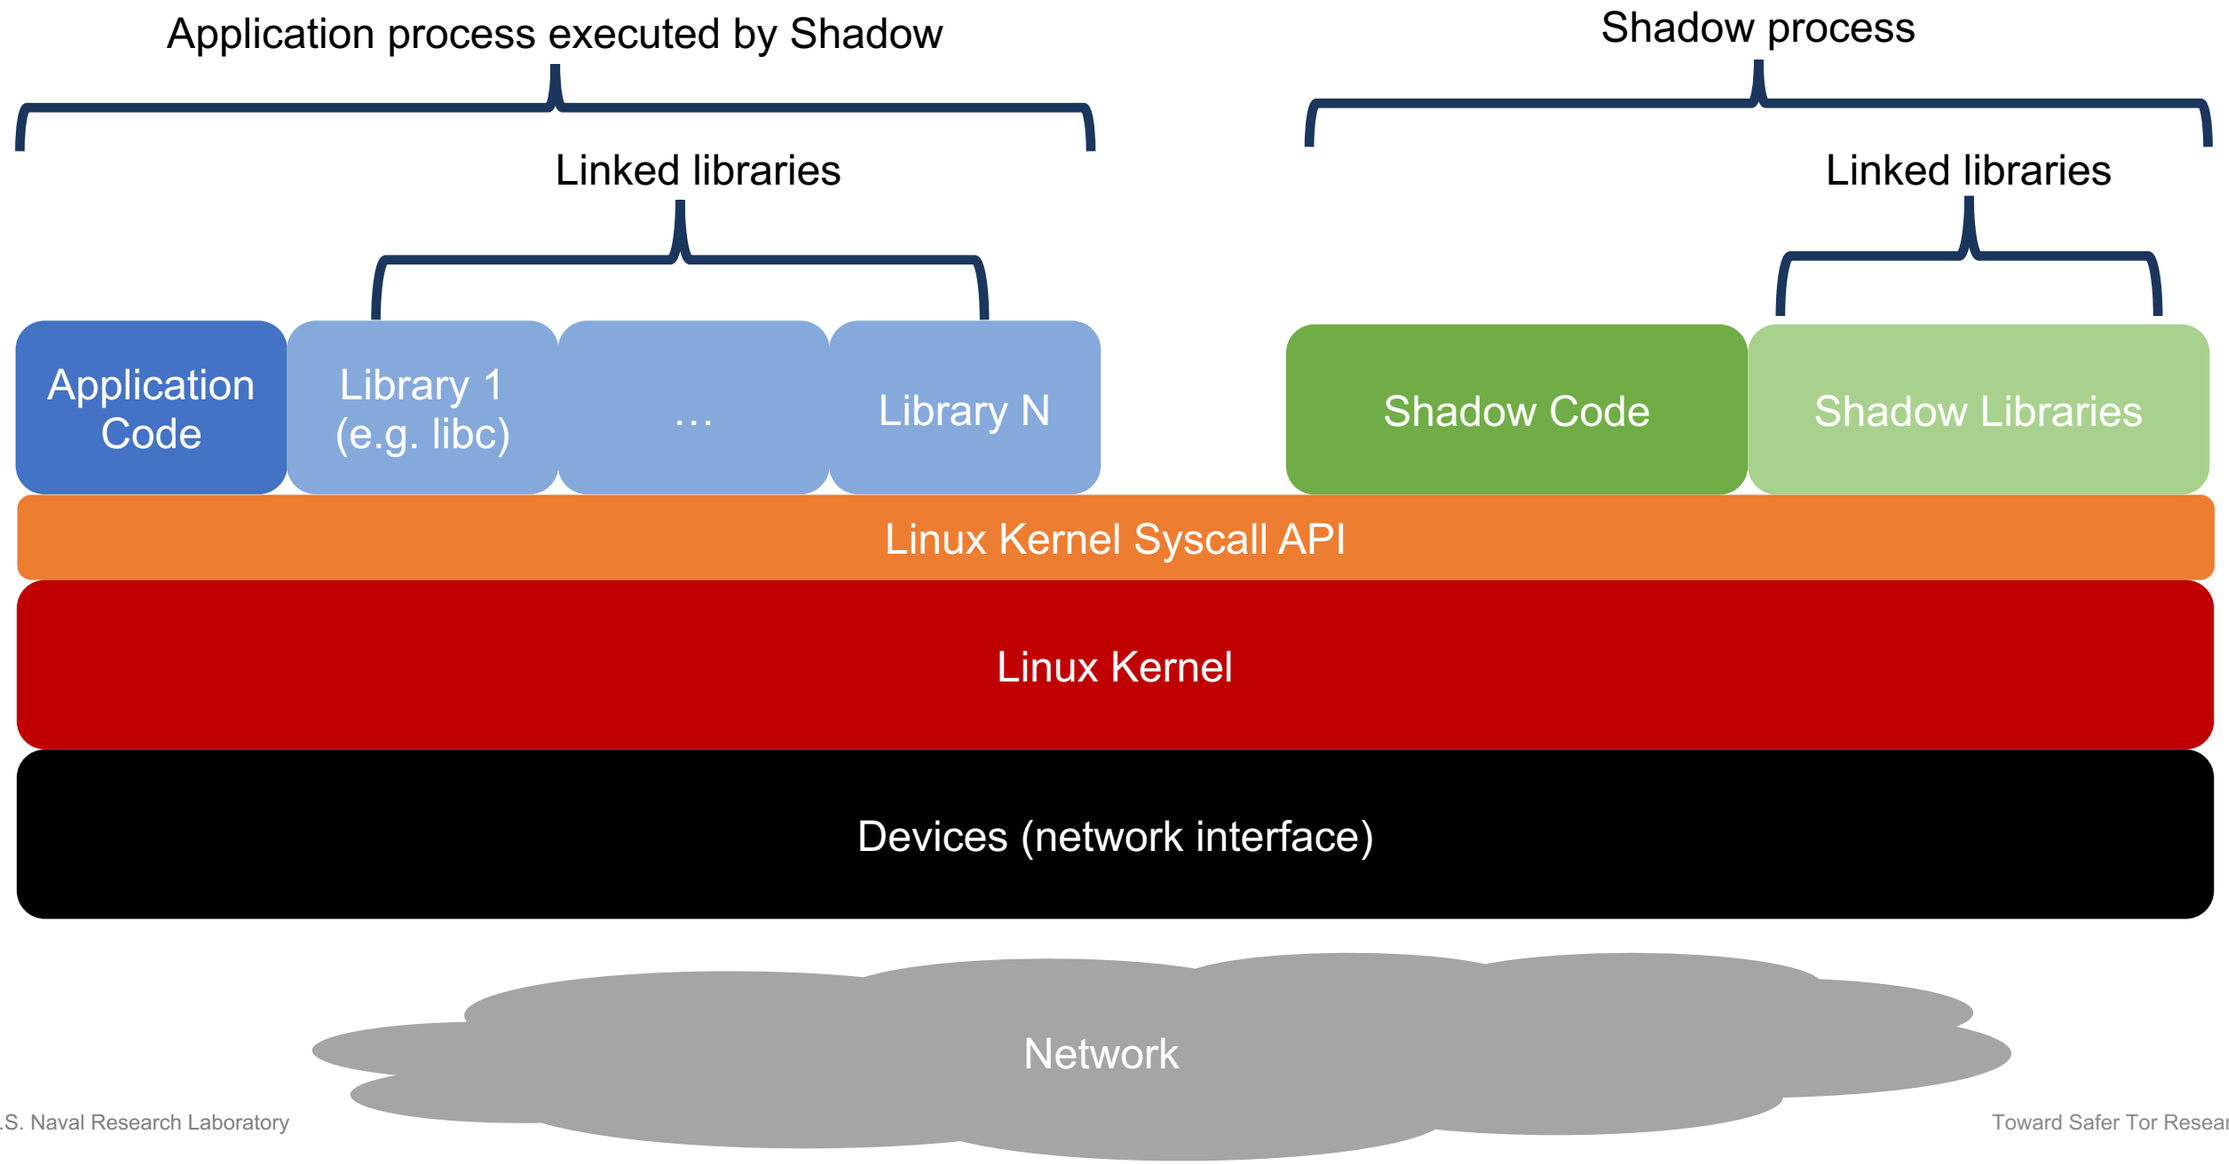

U.S. NAVAL RESEARCH LABORATORY

Use Shim to Control Process

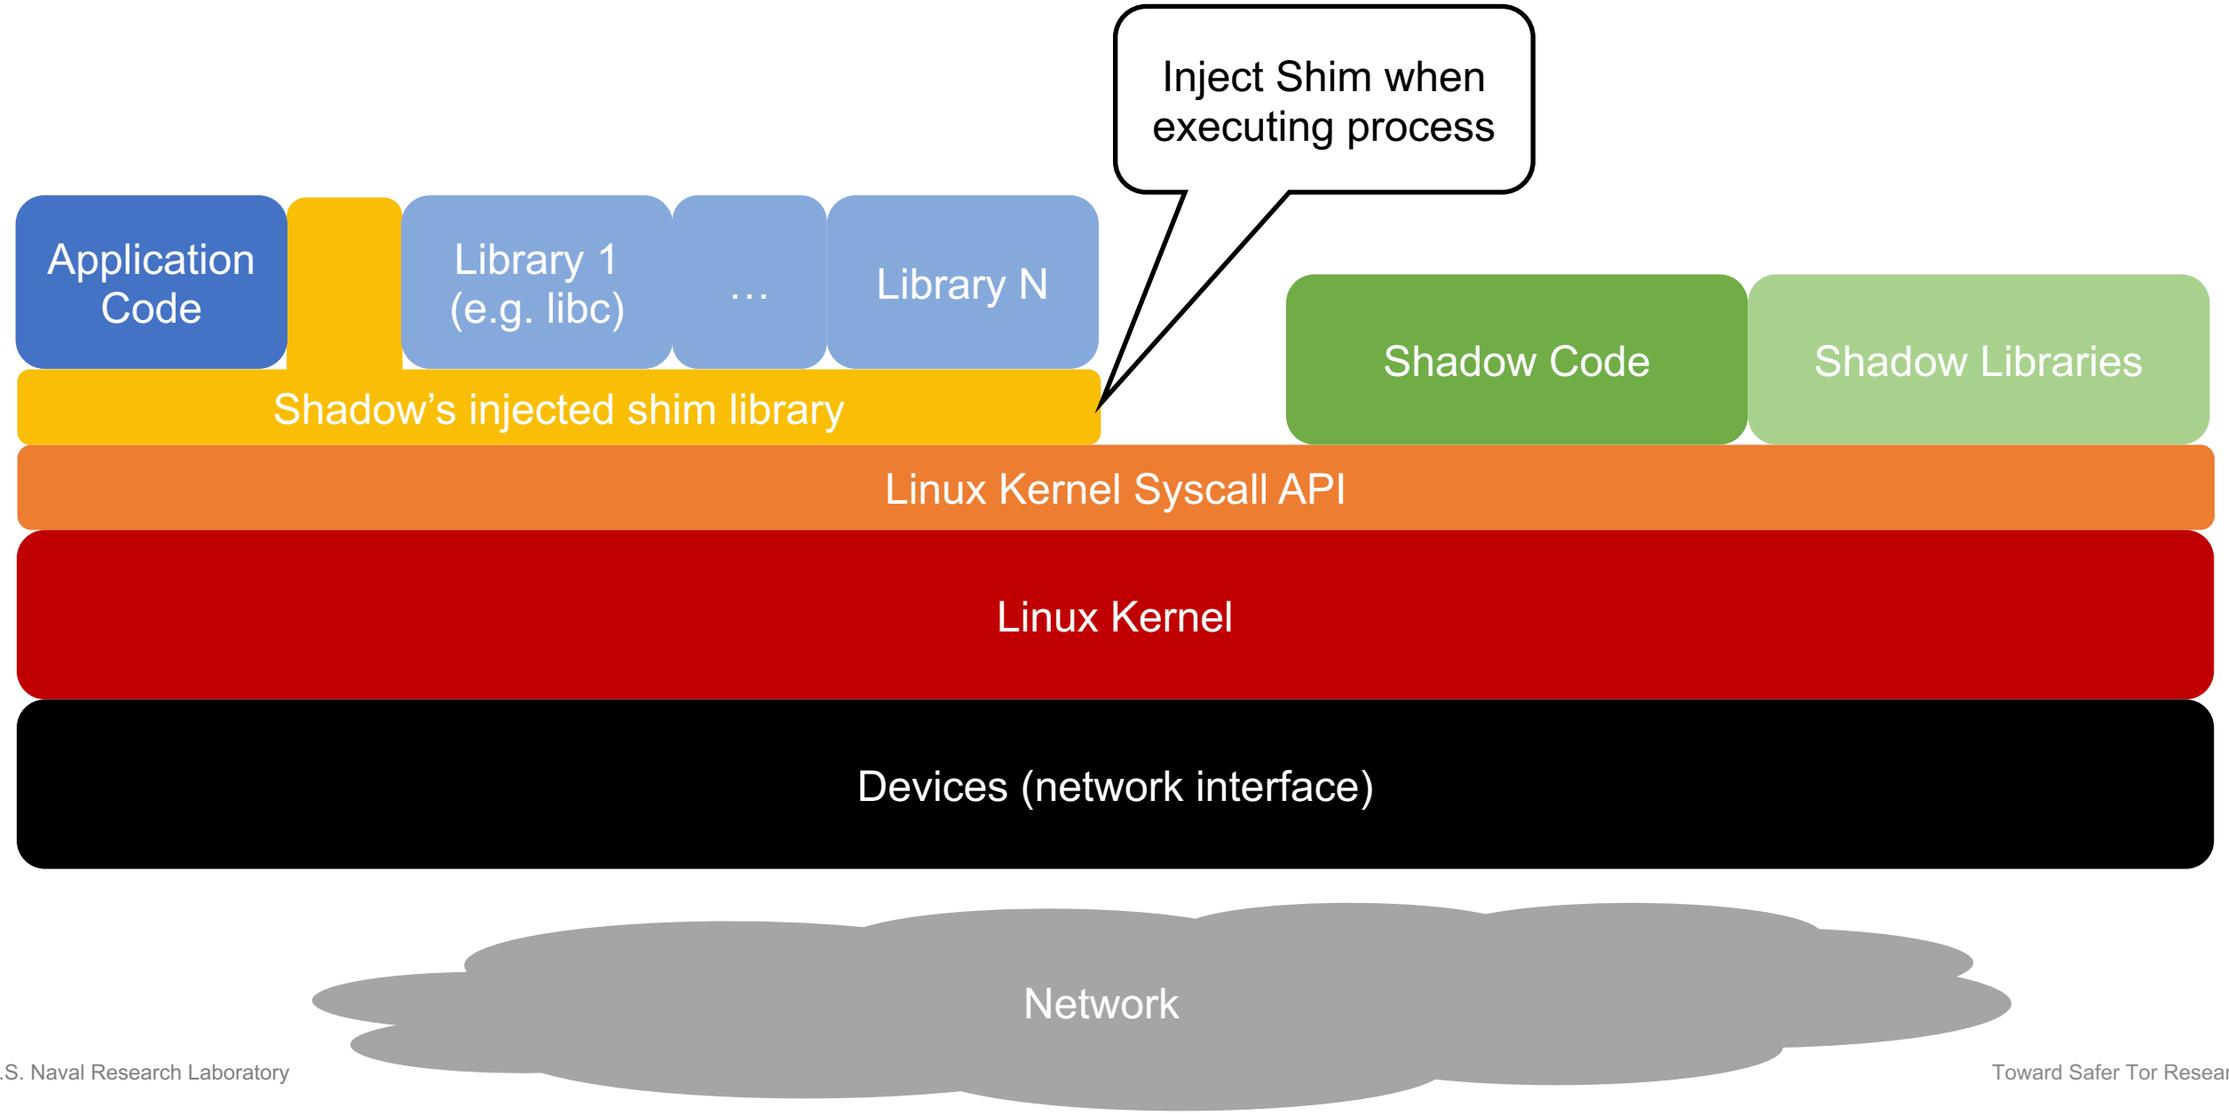

U.S. NAVAL RESEARCH LABORATORY

Intercepting System Calls

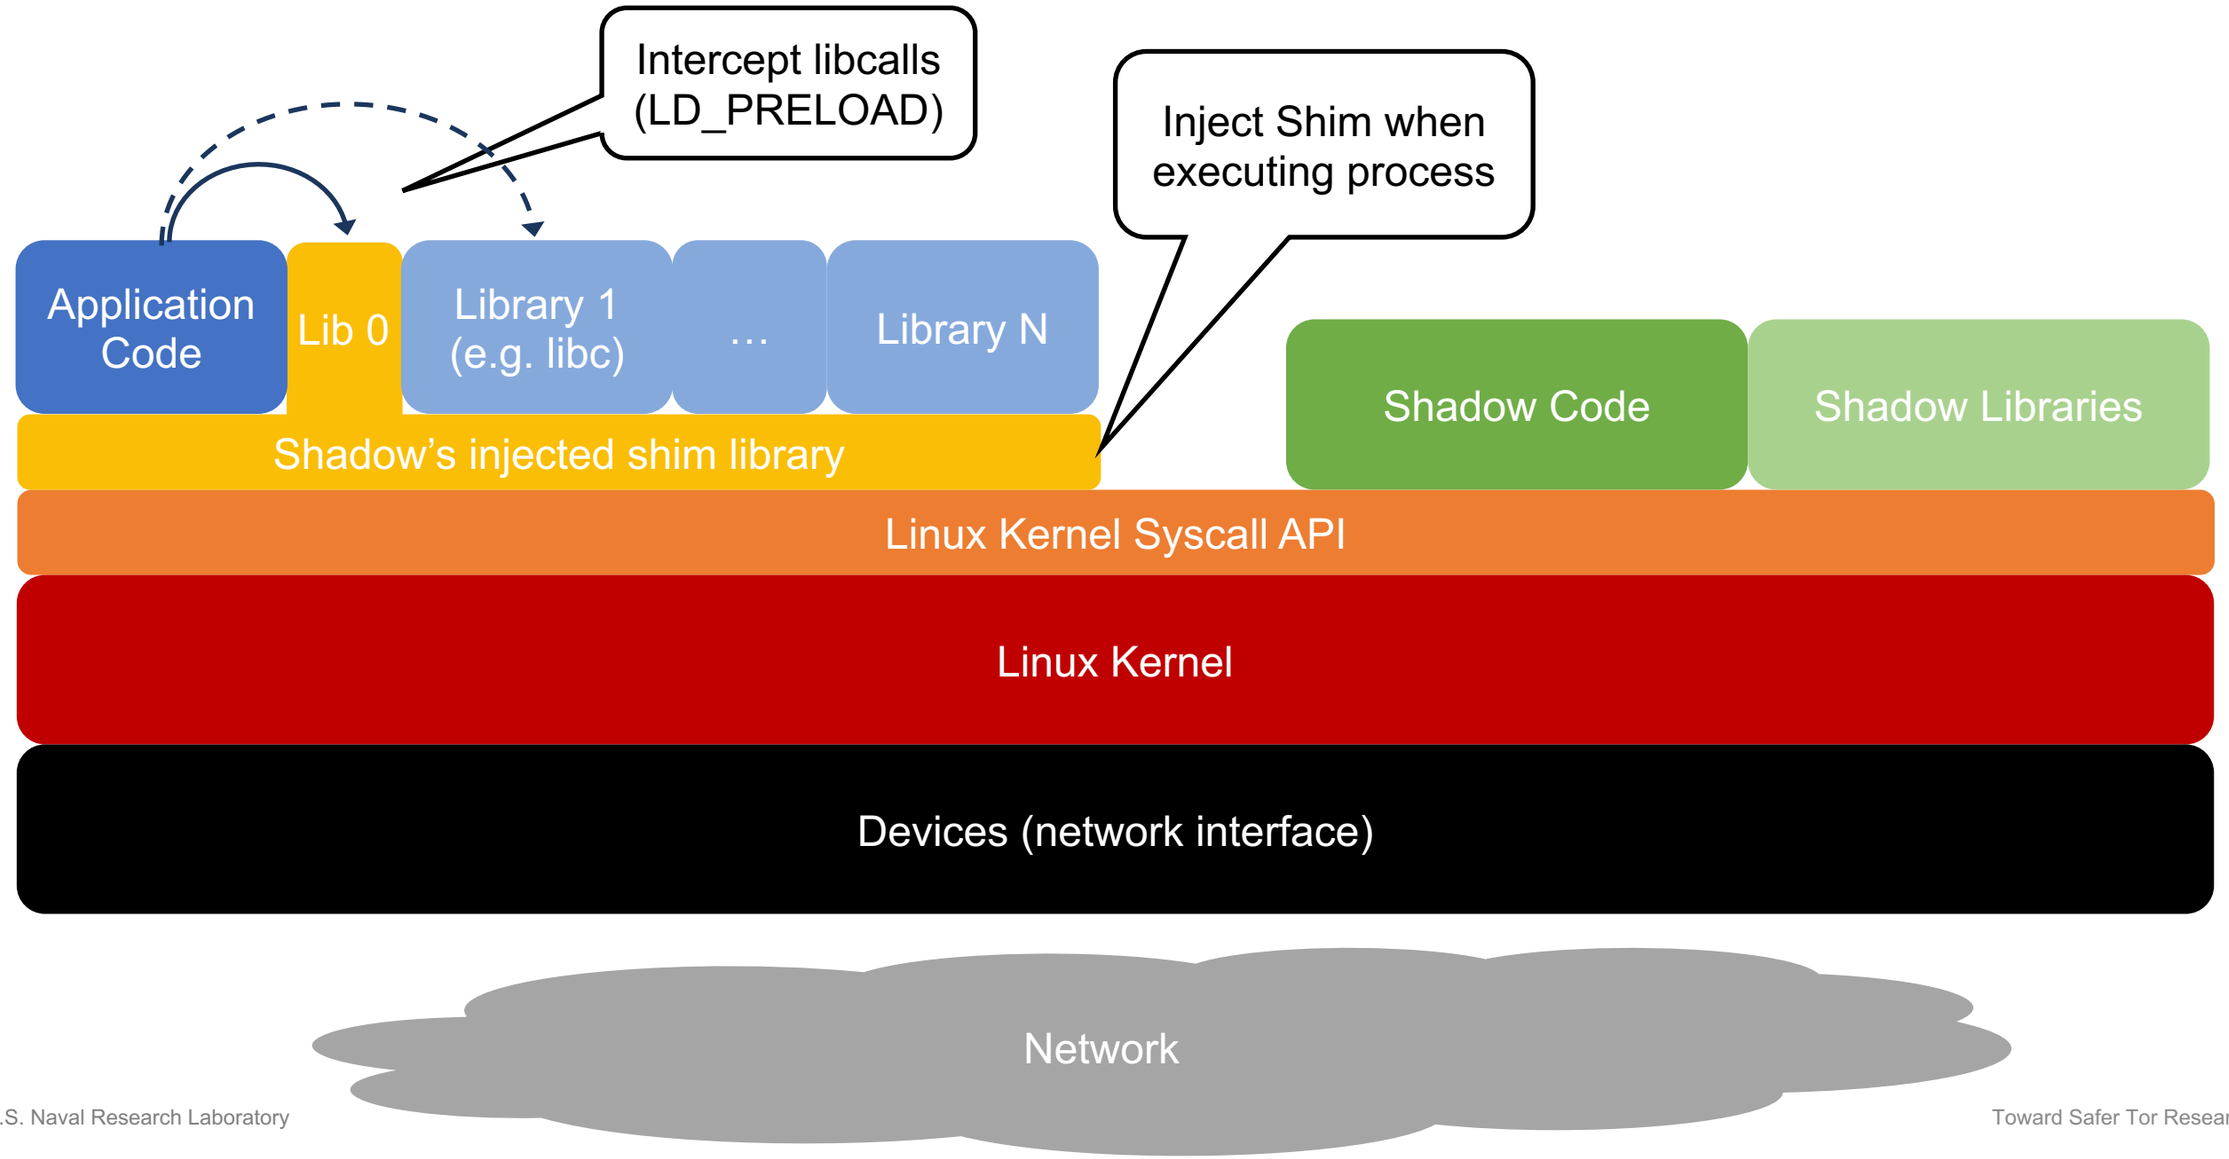

U.S. NAVAL RESEARCH LABORATORY

Intercepting System Calls

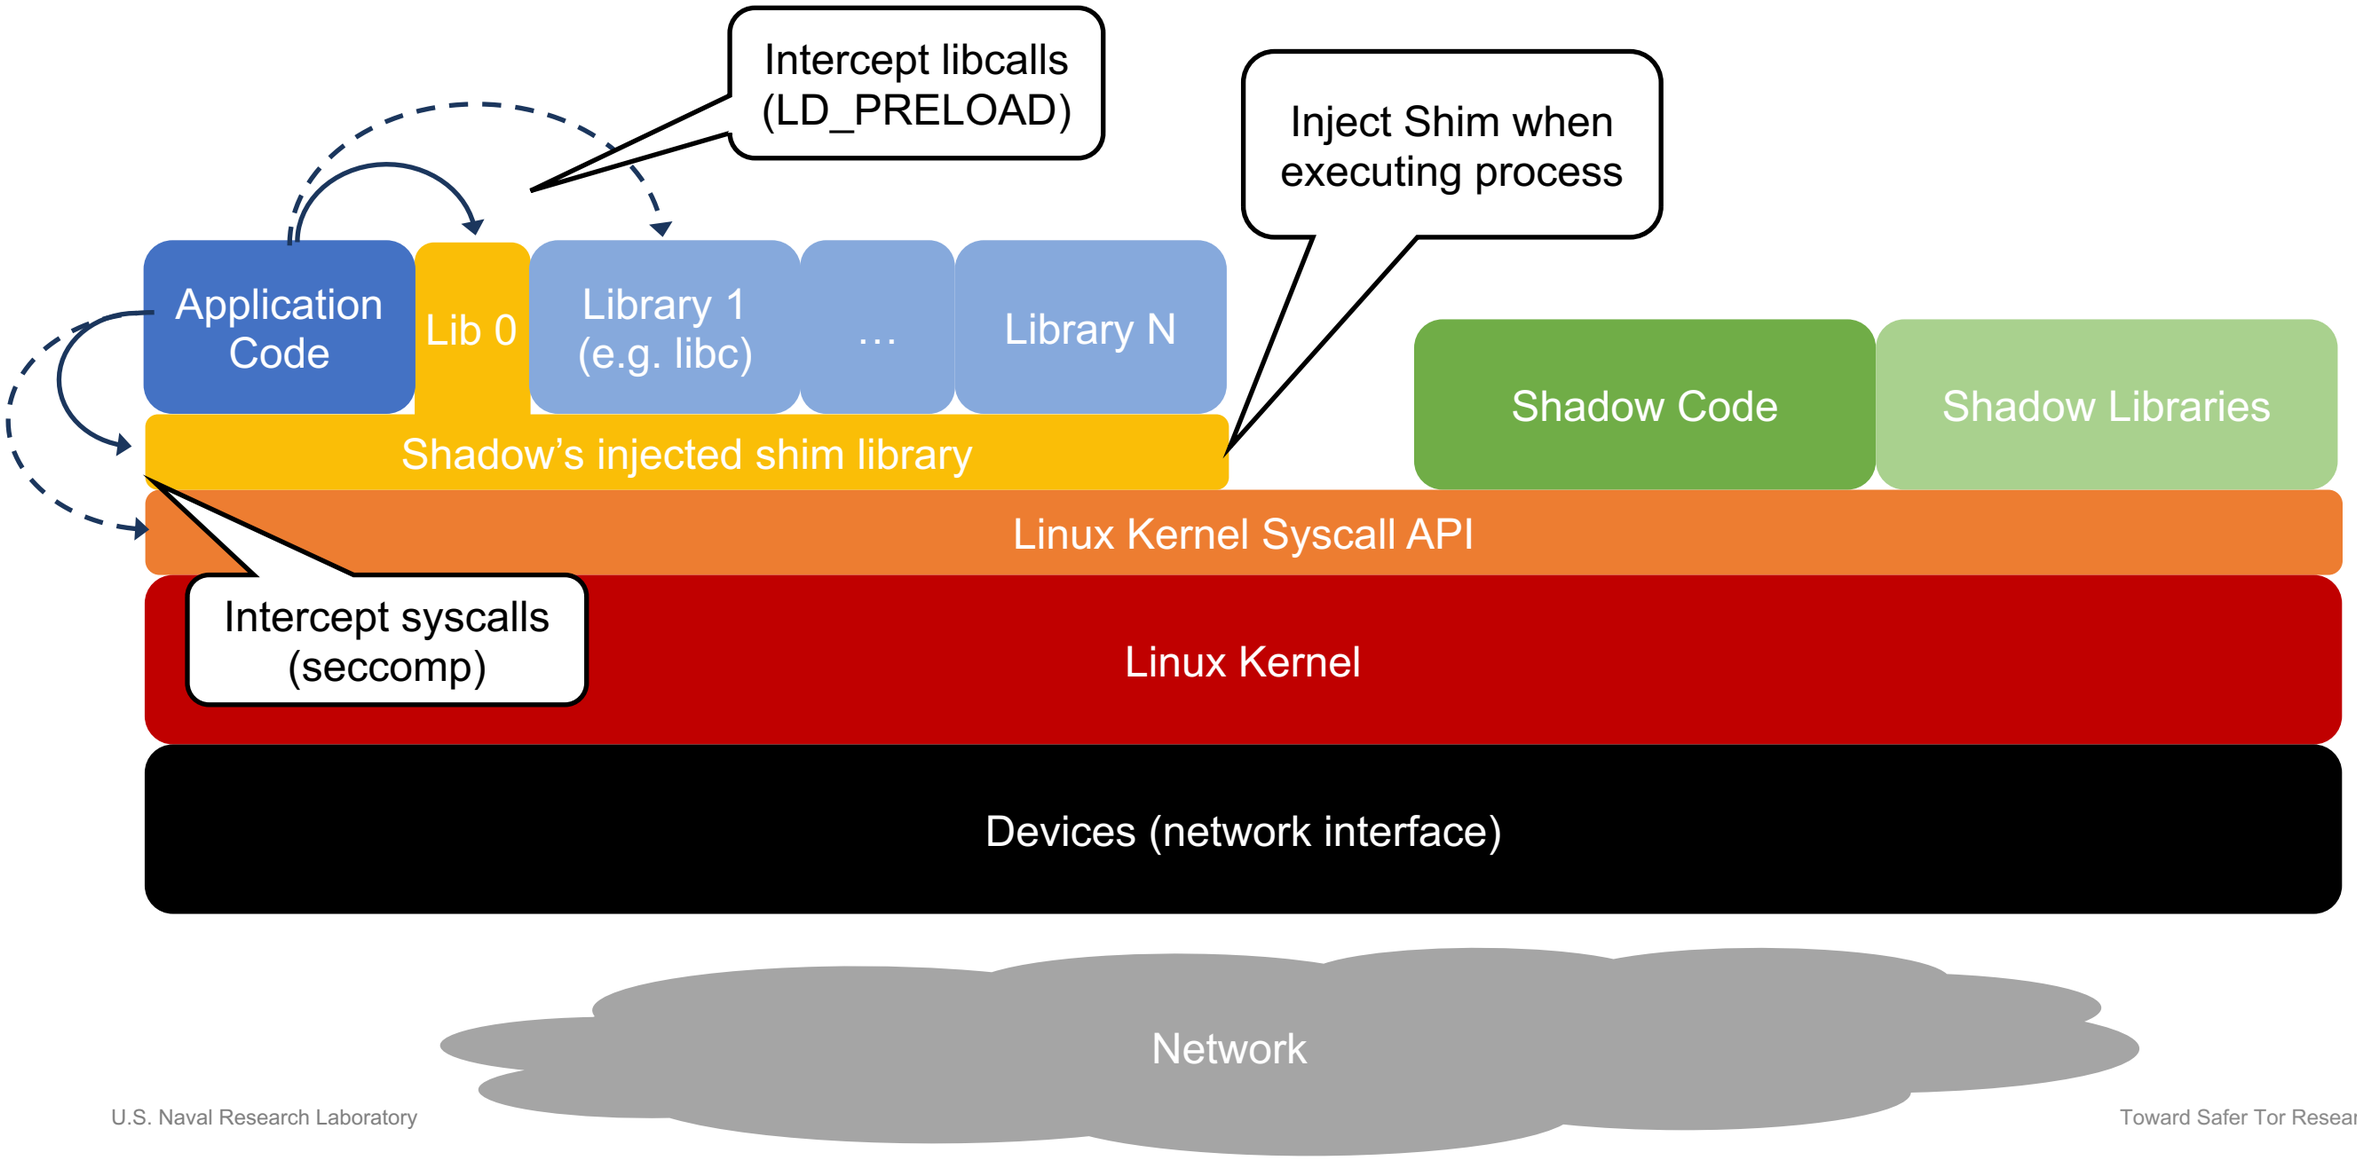

U.S. NAVAL RESEARCH LABORATORY

App-Process to Shadow-Process Communication

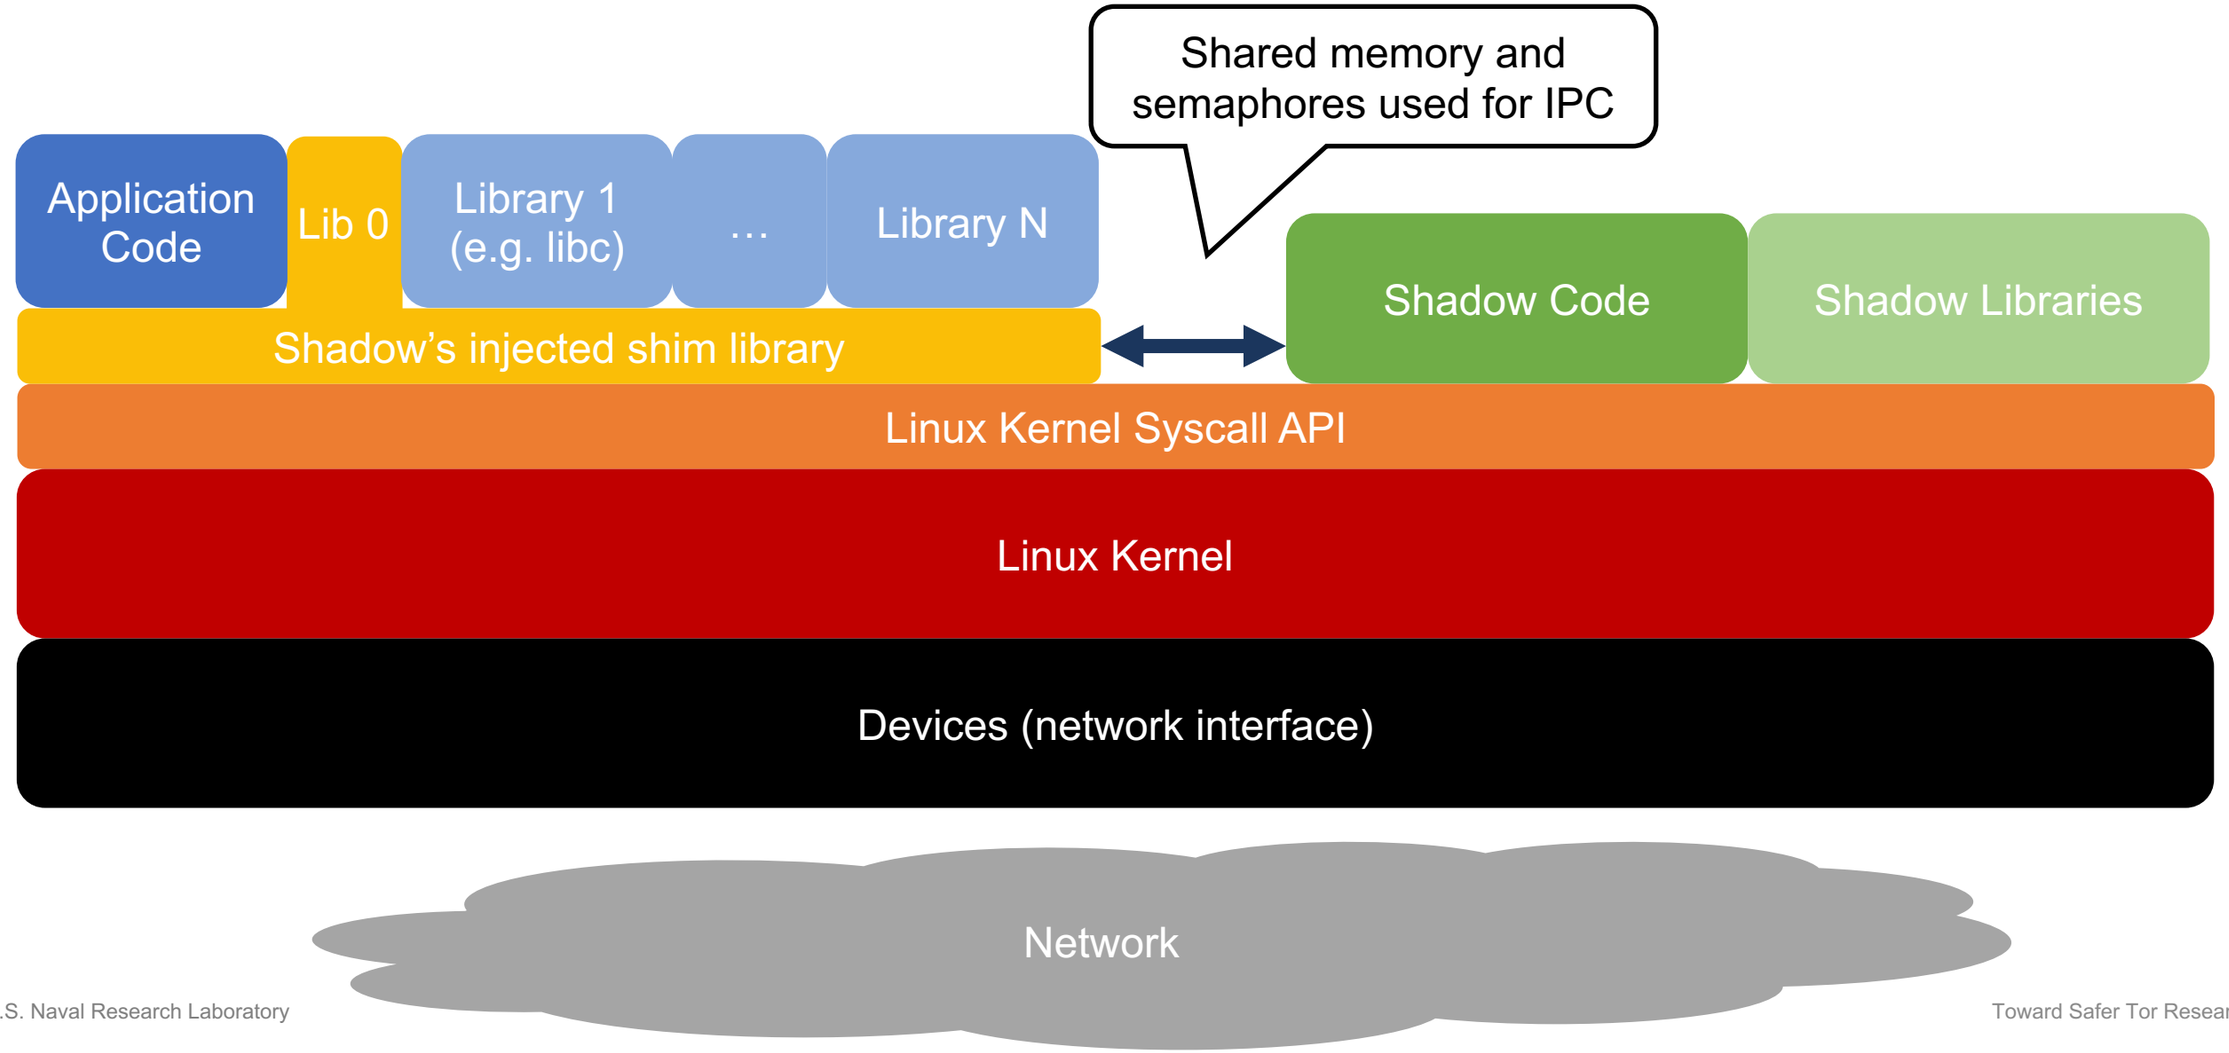

U.S. NAVAL RESEARCH LABORATORY

Syscall Emulation and Network Simulation

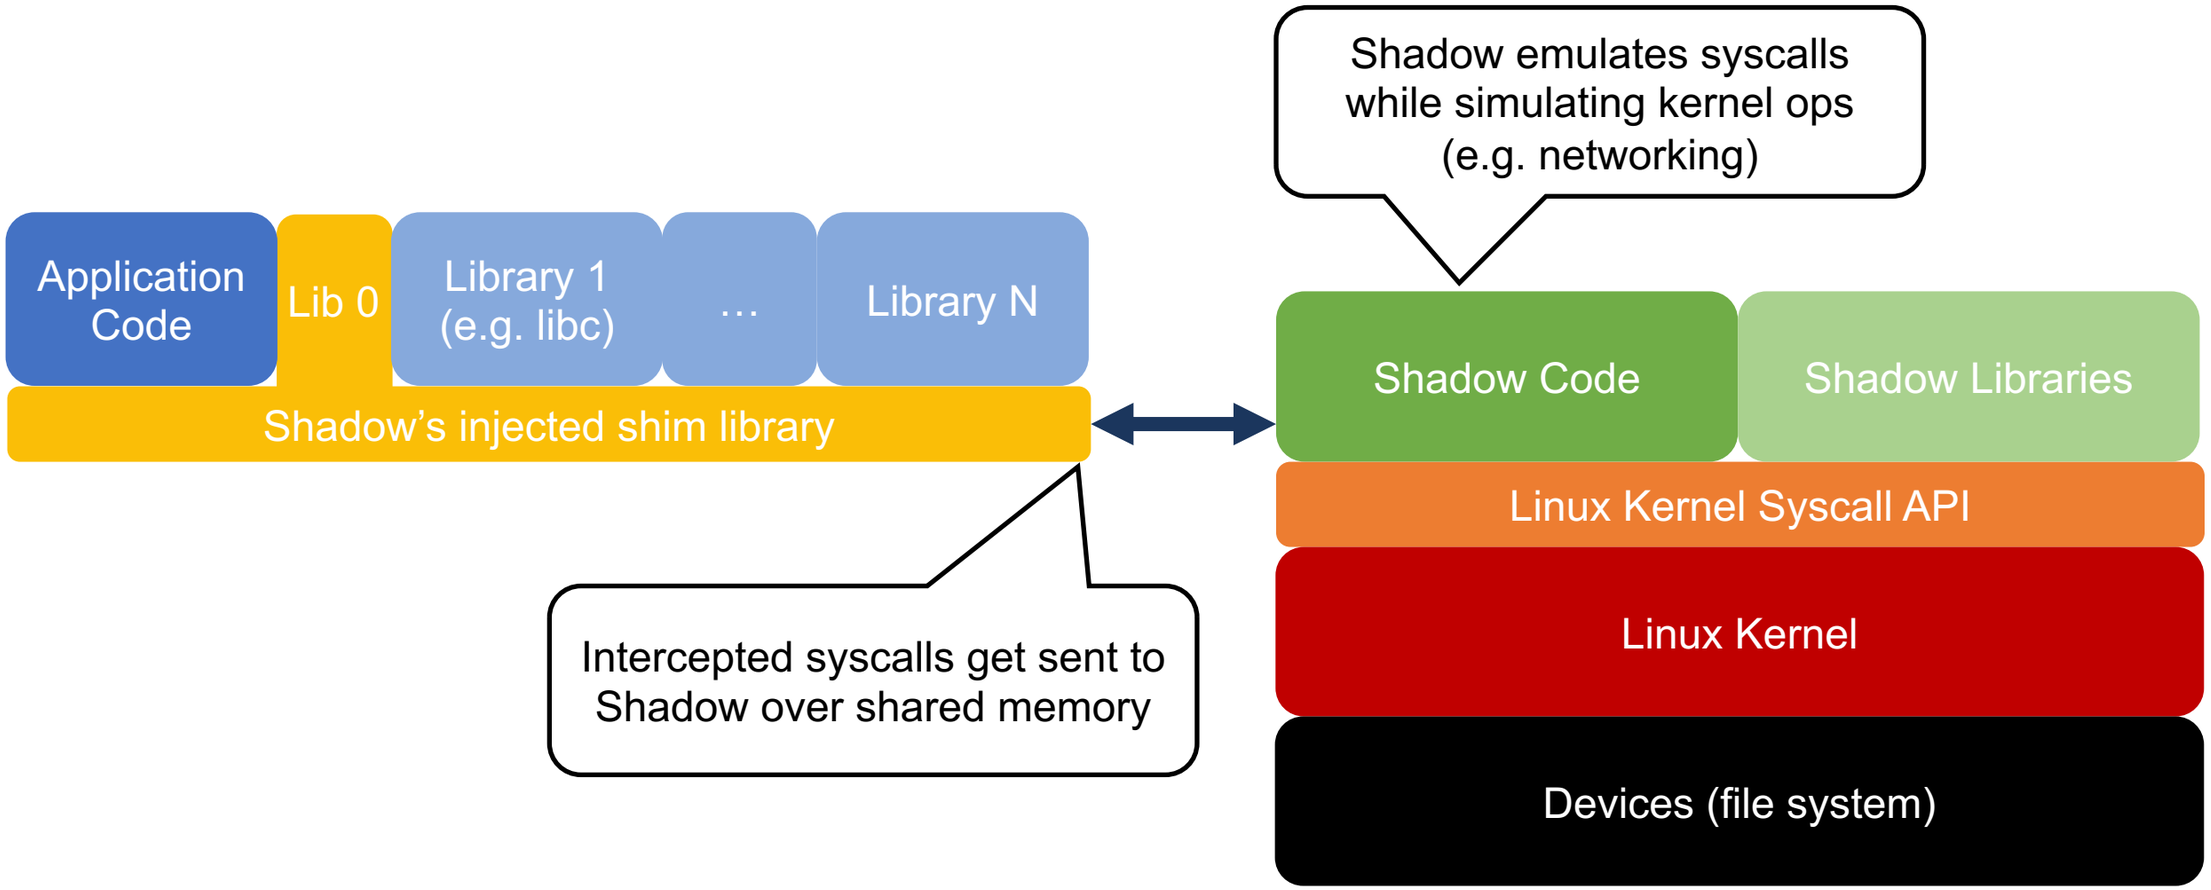

Outline

- **Tor Safety Board overview**
- **Research areas relevant to safety**
  - Simulating Tor with Shadow
  - Measuring Tor with PrivCount
  - Safe research applications
- **Ongoing struggles**

U.S. NAVAL RESEARCH LABORATORY

Modeling Tor Networks

Tor research depends on network experimentation tools to:

- Evaluate research design changes and trade-offs
- Test effects across a range of deployment scenarios and network conditions
- Reproduce research results

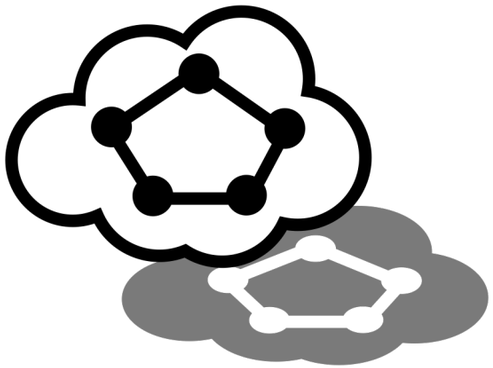

Shadow:
Network Simulation

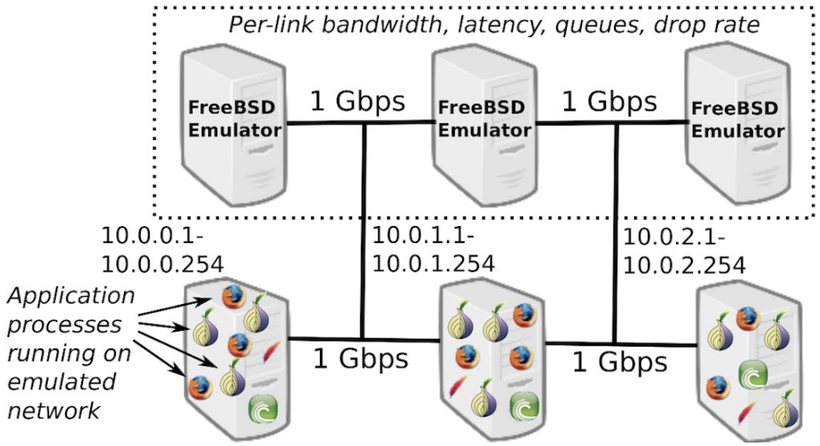

Chutney/NetMirage/ExpTor:
Network Emulation

Modeling Tor Networks

How should we construct a private Tor network from scratch?

- How many clients? How many servers?
  - What is the behavior? Download a file?
- How many relays?
  - Entries, middles, exit positions
  - How to sample relays?
- What are the node characteristics?
  - location, bandwidth, rate limits

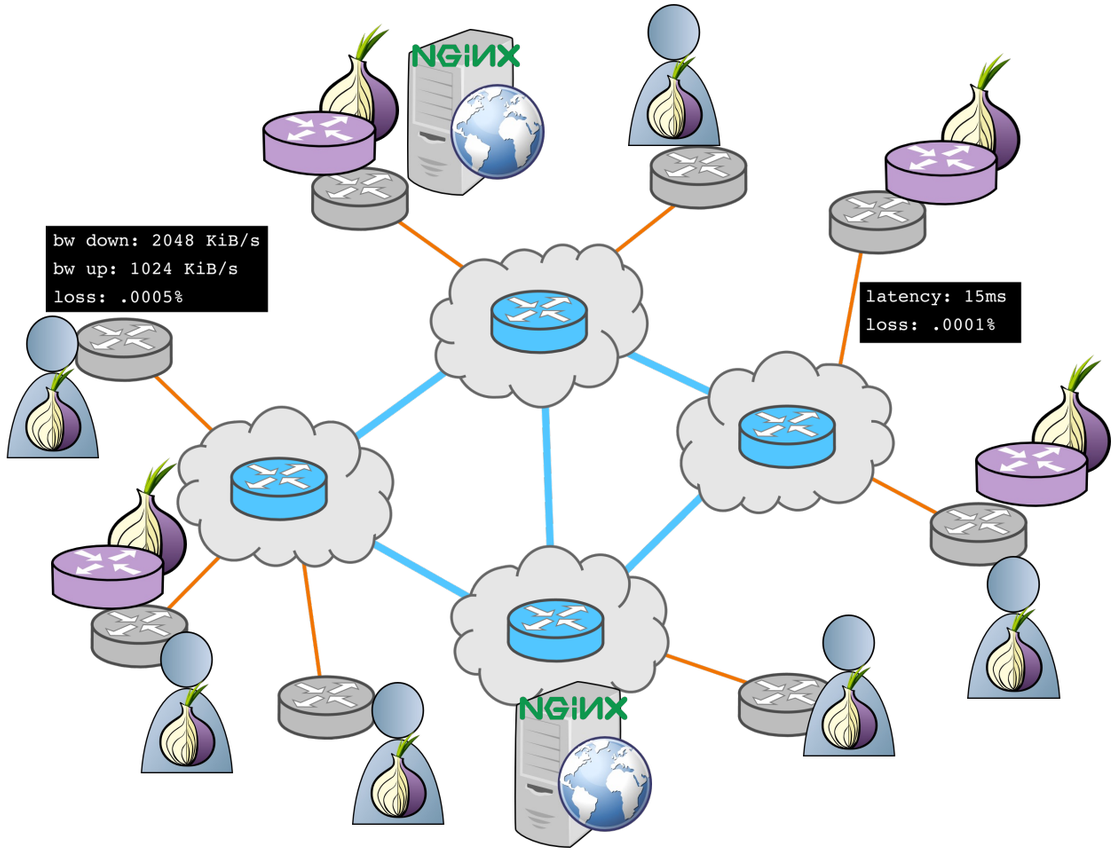

We develop methods and tools:

- To safely measure Tor
- To generate realistic traffic in test networks
- To construct realistic private Tor nets

U.S. NAVAL RESEARCH LABORATORY

Contributing Research

- **On the Accuracy of Tor Bandwidth Estimation**. Rob Jansen and Aaron Johnson. PAM, 2021.
- **Understanding Tor Usage with Privacy-Preserving Measurement.** Akshaya Mani, T Wilson-Brown, Rob Jansen, Aaron Johnson, and Micah Sherr. IMC, 2018.
- **Privacy-Preserving Dynamic Learning of Tor Network Traffic**. Rob Jansen, Mathew Traudt, and Nicholas Hopper. CCS 2018.
- **Safely Measuring Tor**. Rob Jansen and Aaron Johnson. CCS 2016.

U.S. NAVAL RESEARCH LABORATORY

PrivCount Measurement System

PrivCount: a privacy-preserving counting system

- Designed to safely collect useful Tor statistics
- Based on the PrivEx secret sharing protocol [CCS'14]

PrivCount security goals:

- Forward privacy: adversary cannot learn state of measurement before time of compromise
- Secure aggregation across all measurement nodes
- Measurement results are differentially private to protect user actions

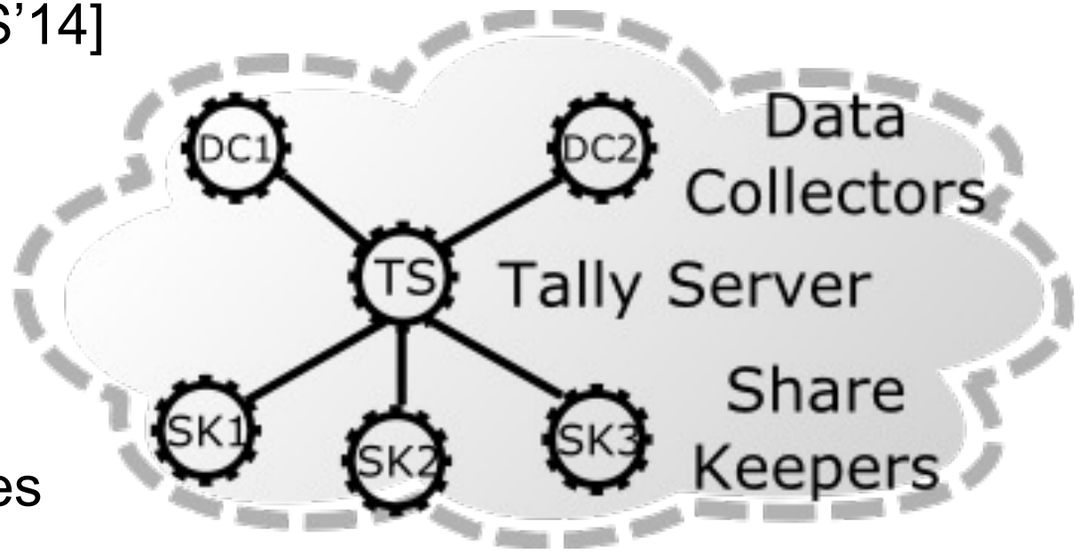

U.S. NAVAL RESEARCH LABORATORY

Deploying PrivCount on the Public Tor Network

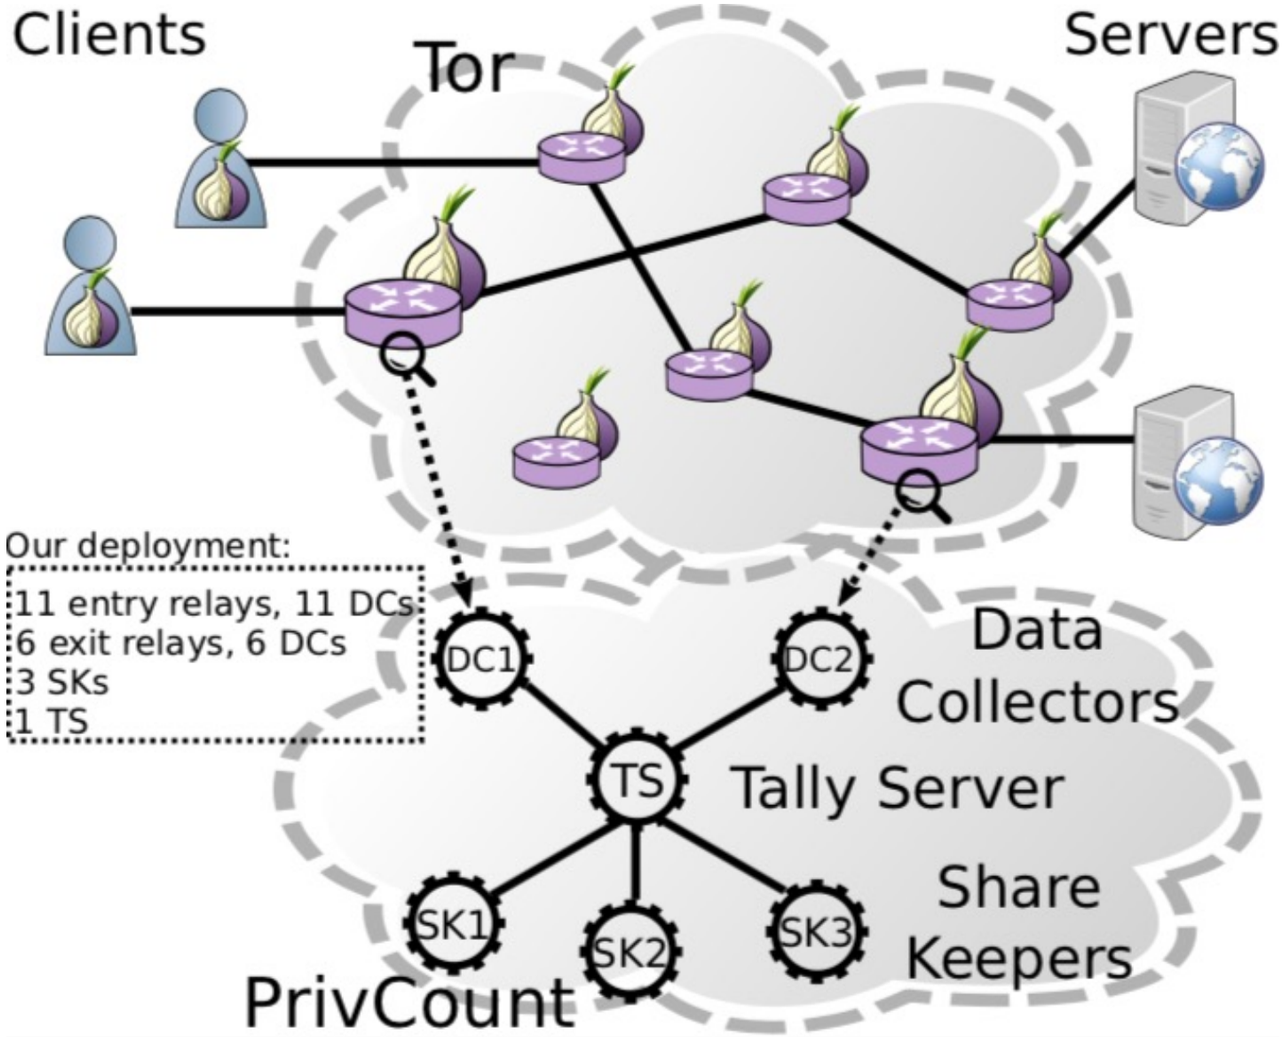

U.S. NAVAL RESEARCH LABORATORY

Deploying PrivCount on the Public Tor Network

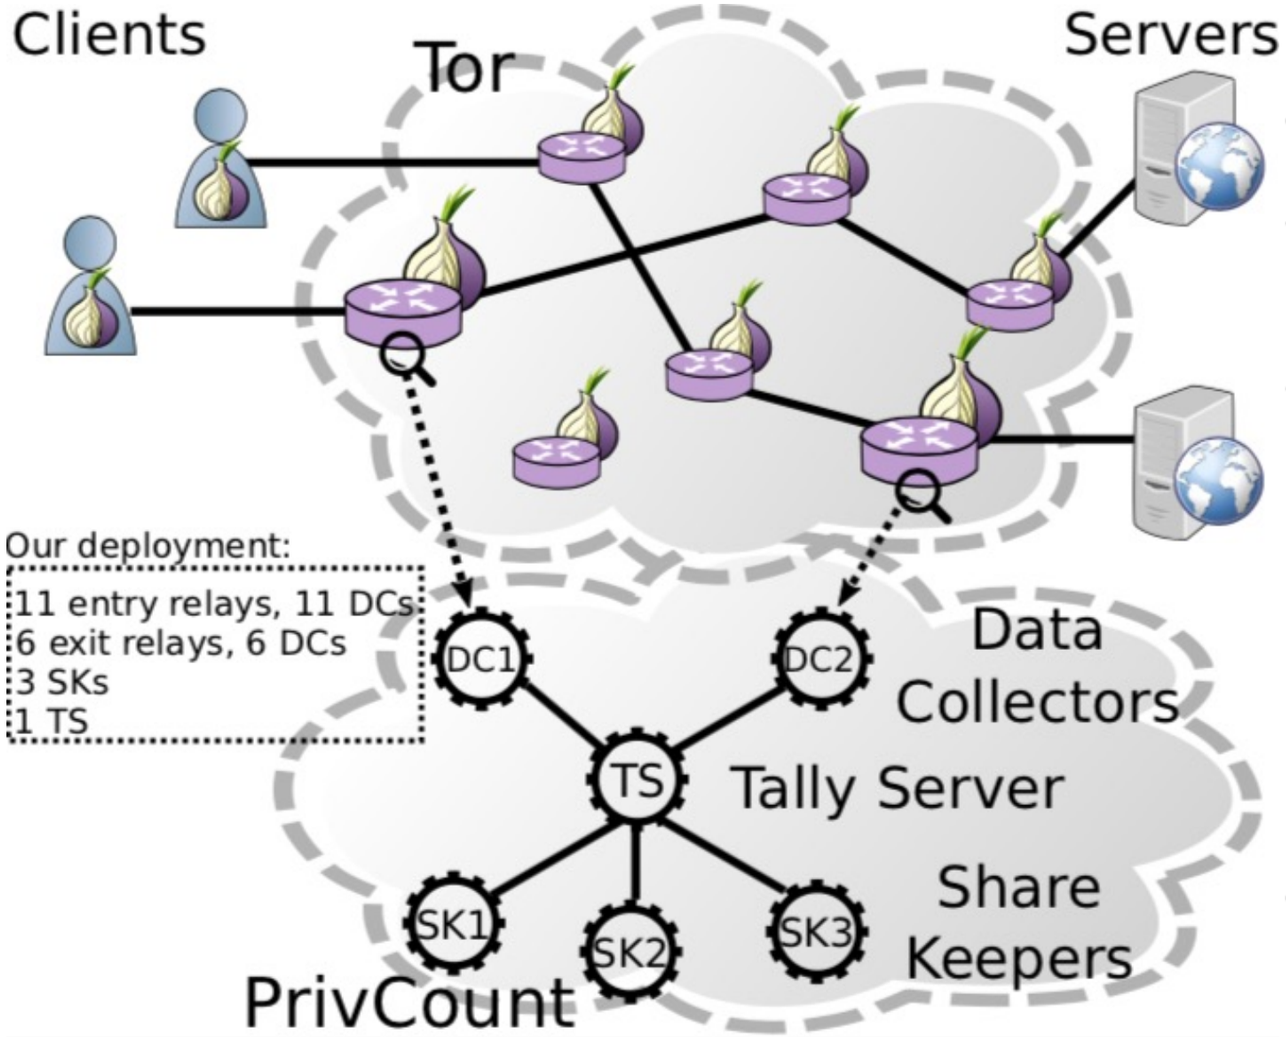

	#	Purpose of Measurement
Entry	1	Total clients and circuits
	2	Circuits per client
Exit	3	Total circuits and streams
	4	Total bytes on streams
	5	Streams per circuit, bytes per stream (All)
	6	Streams per circuit, bytes per stream (Web)
	7	Streams per circuit, bytes per stream (Other)
	8	Hidden Markov packet model
	9	Hidden Markov stream model

* Weights correspond to the relay measurement position.

U.S. NAVAL RESEARCH LABORATORY

Hidden Markov Modeling of Traffic

Use exit relay observations and PrivCount to safely learn HMM stream and packet models of live Tor traffic

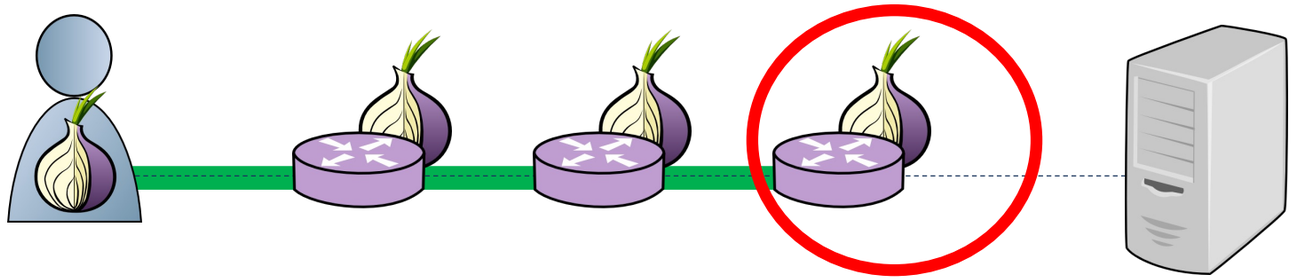

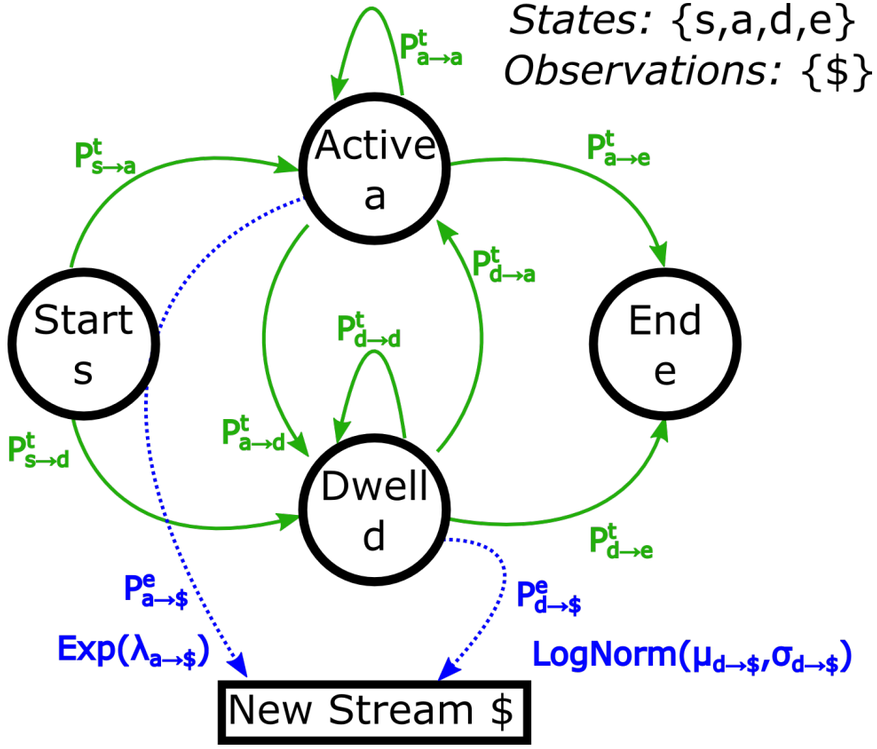

Exits can observe:

- Stream model events
  - Circuit opened, stream created, circuit closed
- Packet model events
  - Stream opened, packet transferred (directional), stream closed
- Both models
  - Inter-event timing (relative time since previous observed event)

Treat Markov model transition probabilities as PrivCount counters

U.S. NAVAL RESEARCH LABORATORY

Realistic Private Tor Networks

- Shadow:
  - Directly executes Tor and other apps
  - Use to run private Tor networks that are 100% safe
  - https://shadow.github.io
- PrivCount:
  - Measurement system for safely measuring the public Tor network
  - Use to measure Tor and train hidden Markov traffic models
  - https://github.com/privcount/privcount
- Tgen:
  - Produces network traffic flows from the Markov models
  - https://github.com/shadow/tgen
- TorNetTools:
  - Construct private Tor networks using PrivCount measurements and HMMs
  - Run sims using Shadow, Tgen, Tor
  - Process and plot sim results
  - https://github.com/shadow/tornettools

Outline

- **Tor Safety Board overview**
- **Research areas relevant to safety**
  - Simulating Tor with Shadow
  - Measuring Tor with PrivCount
  - Safe research applications
- **Ongoing struggles**

Exploring Tor Performance Attacks

Explore the costs and effects of bandwidth denial-of-service attacks on Tor

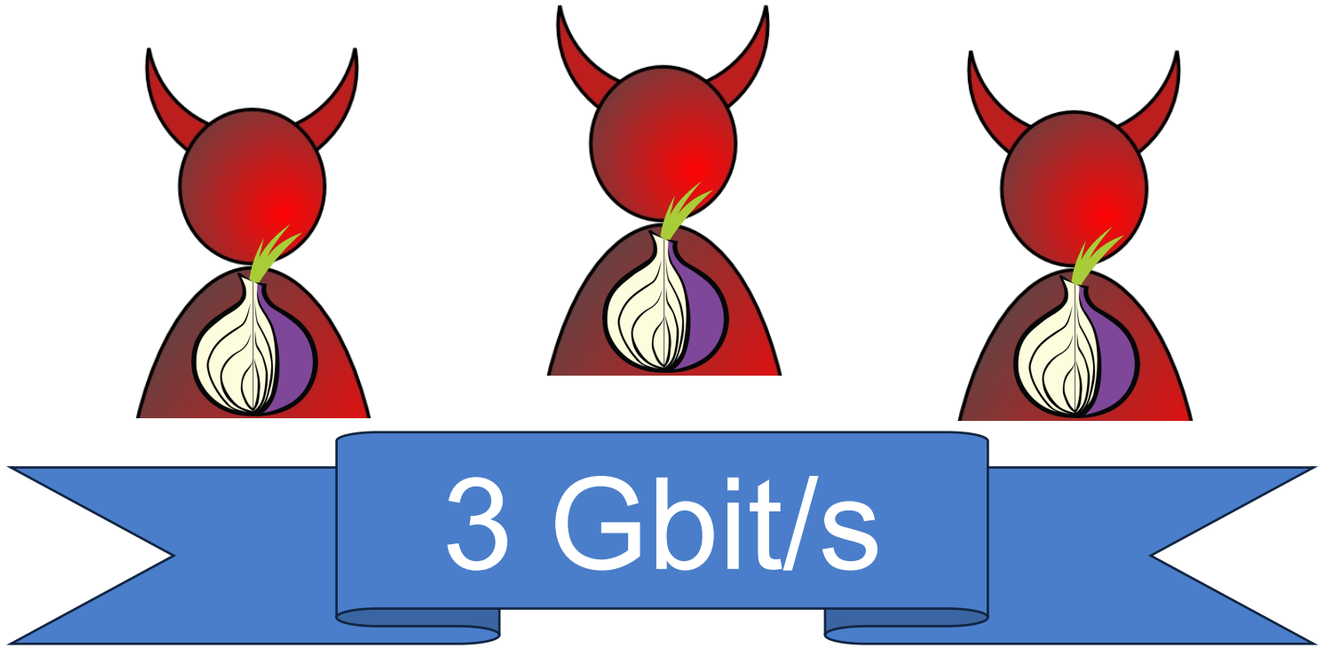

$140 - $1.6K / mo.

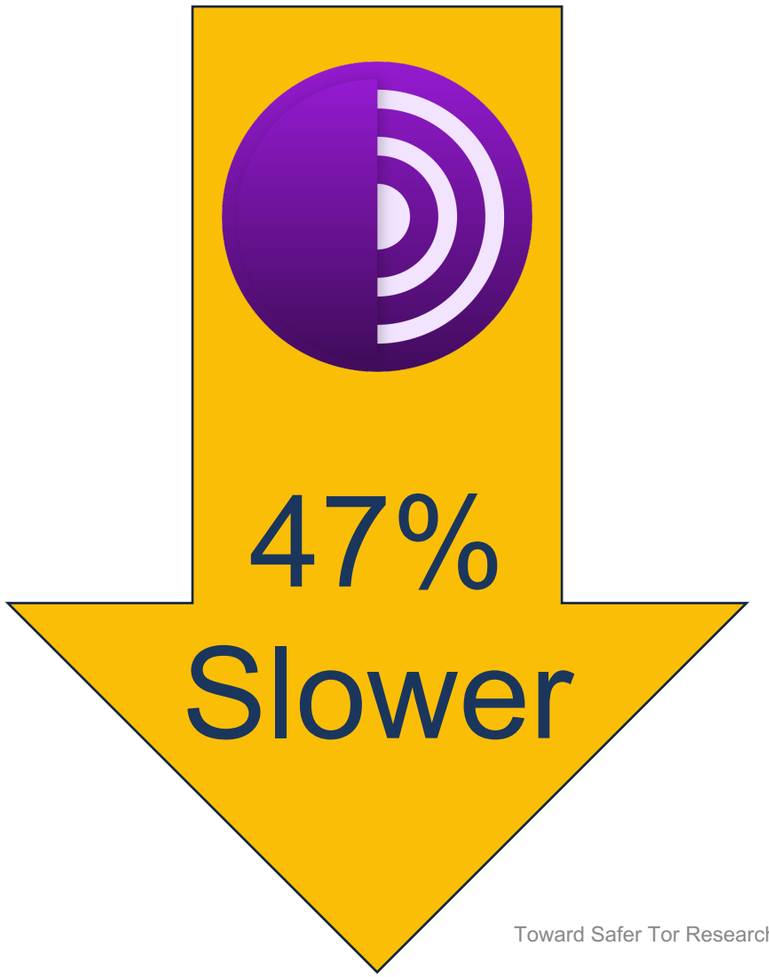

U.S. NAVAL RESEARCH LABORATORY

Contributing Research

- **Point Break: A Study of Bandwidth Denial-of-Service Attacks against Tor**. Rob Jansen, Tavish Vaidya, and Micah Sherr. USENIX Security, 2019.

- **The Sniper Attack: Anonymously Deanonymizing and Disabling the Tor Network**. Rob Jansen, Florian Tschorsch, Aaron Johnson, and Björn Scheuermann. NDSS, 2014.

The Relay Congestion Attack

Step 1:
Build 8-hop circuit

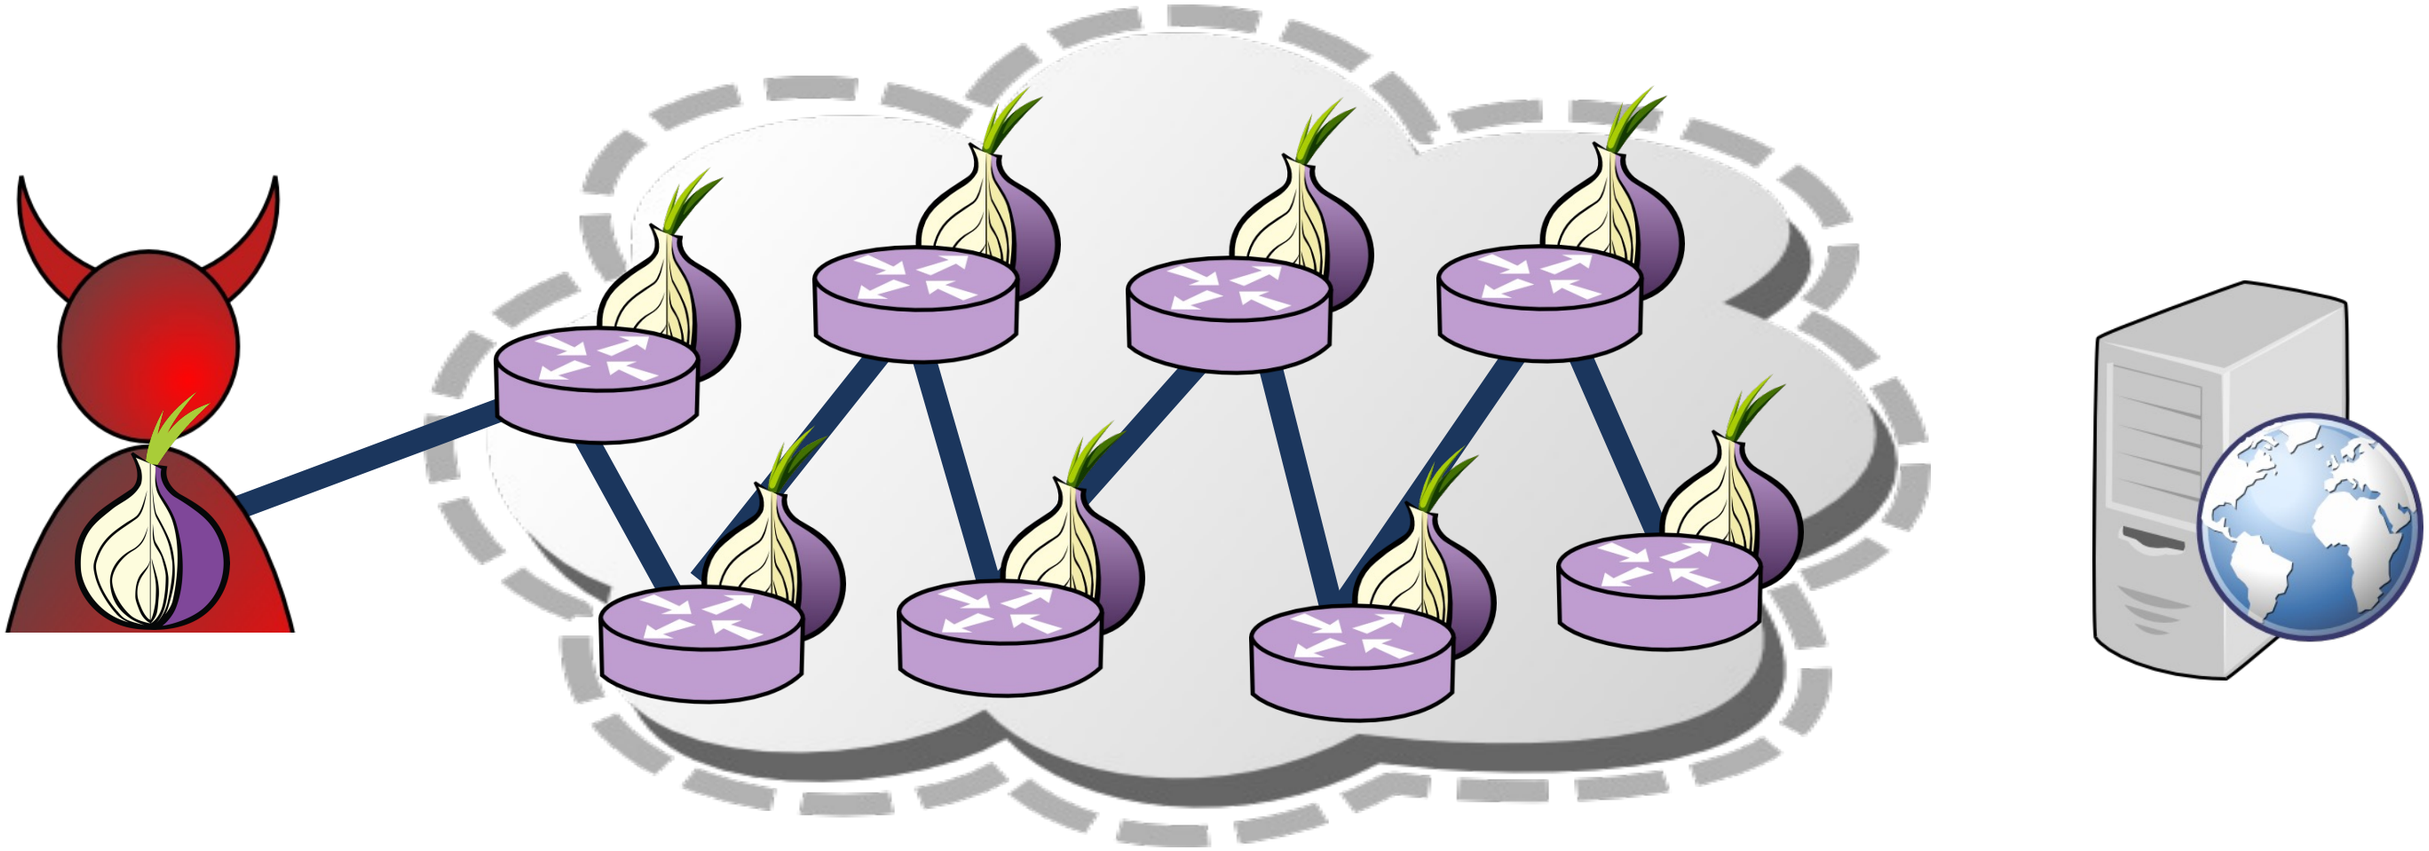

The Relay Congestion Attack

Step 1:
Build 8-hop circuit

Can be targeted or indiscriminate

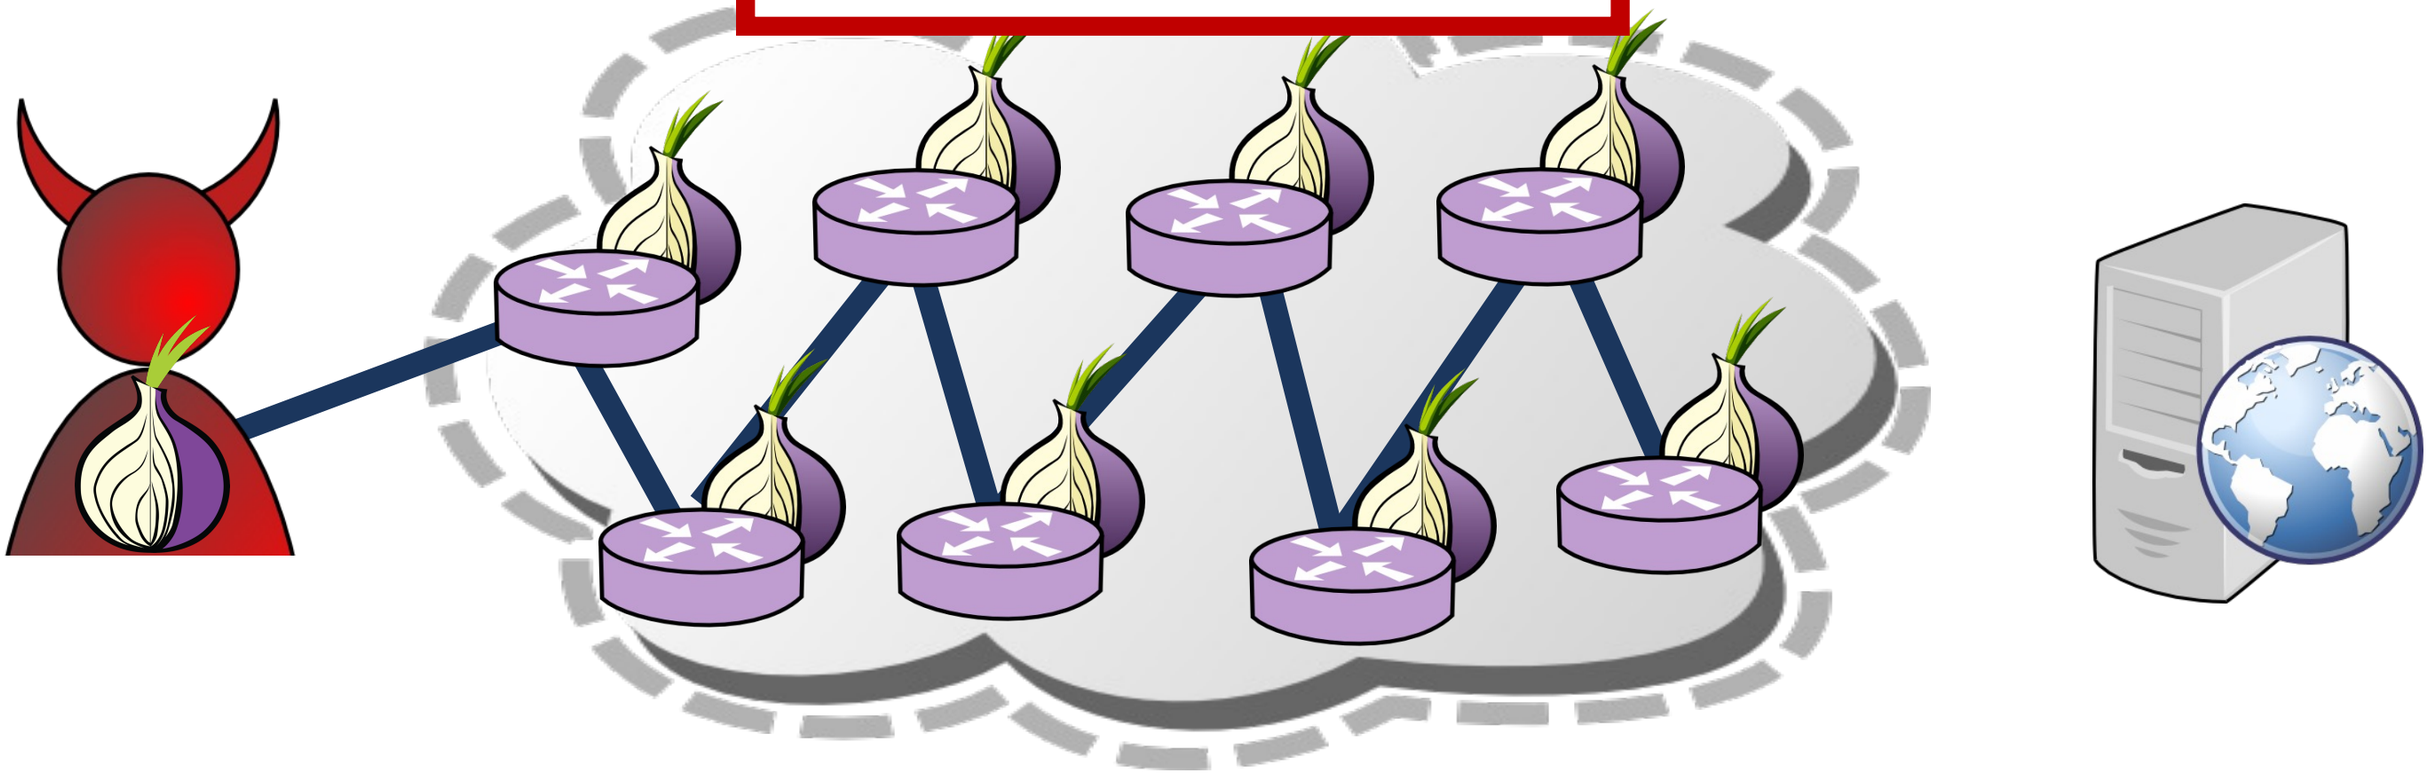

U.S. NAVAL RESEARCH LABORATORY

The Relay Congestion Attack

Step 1:
Build 8-hop circuit

Step 2:
GET large files

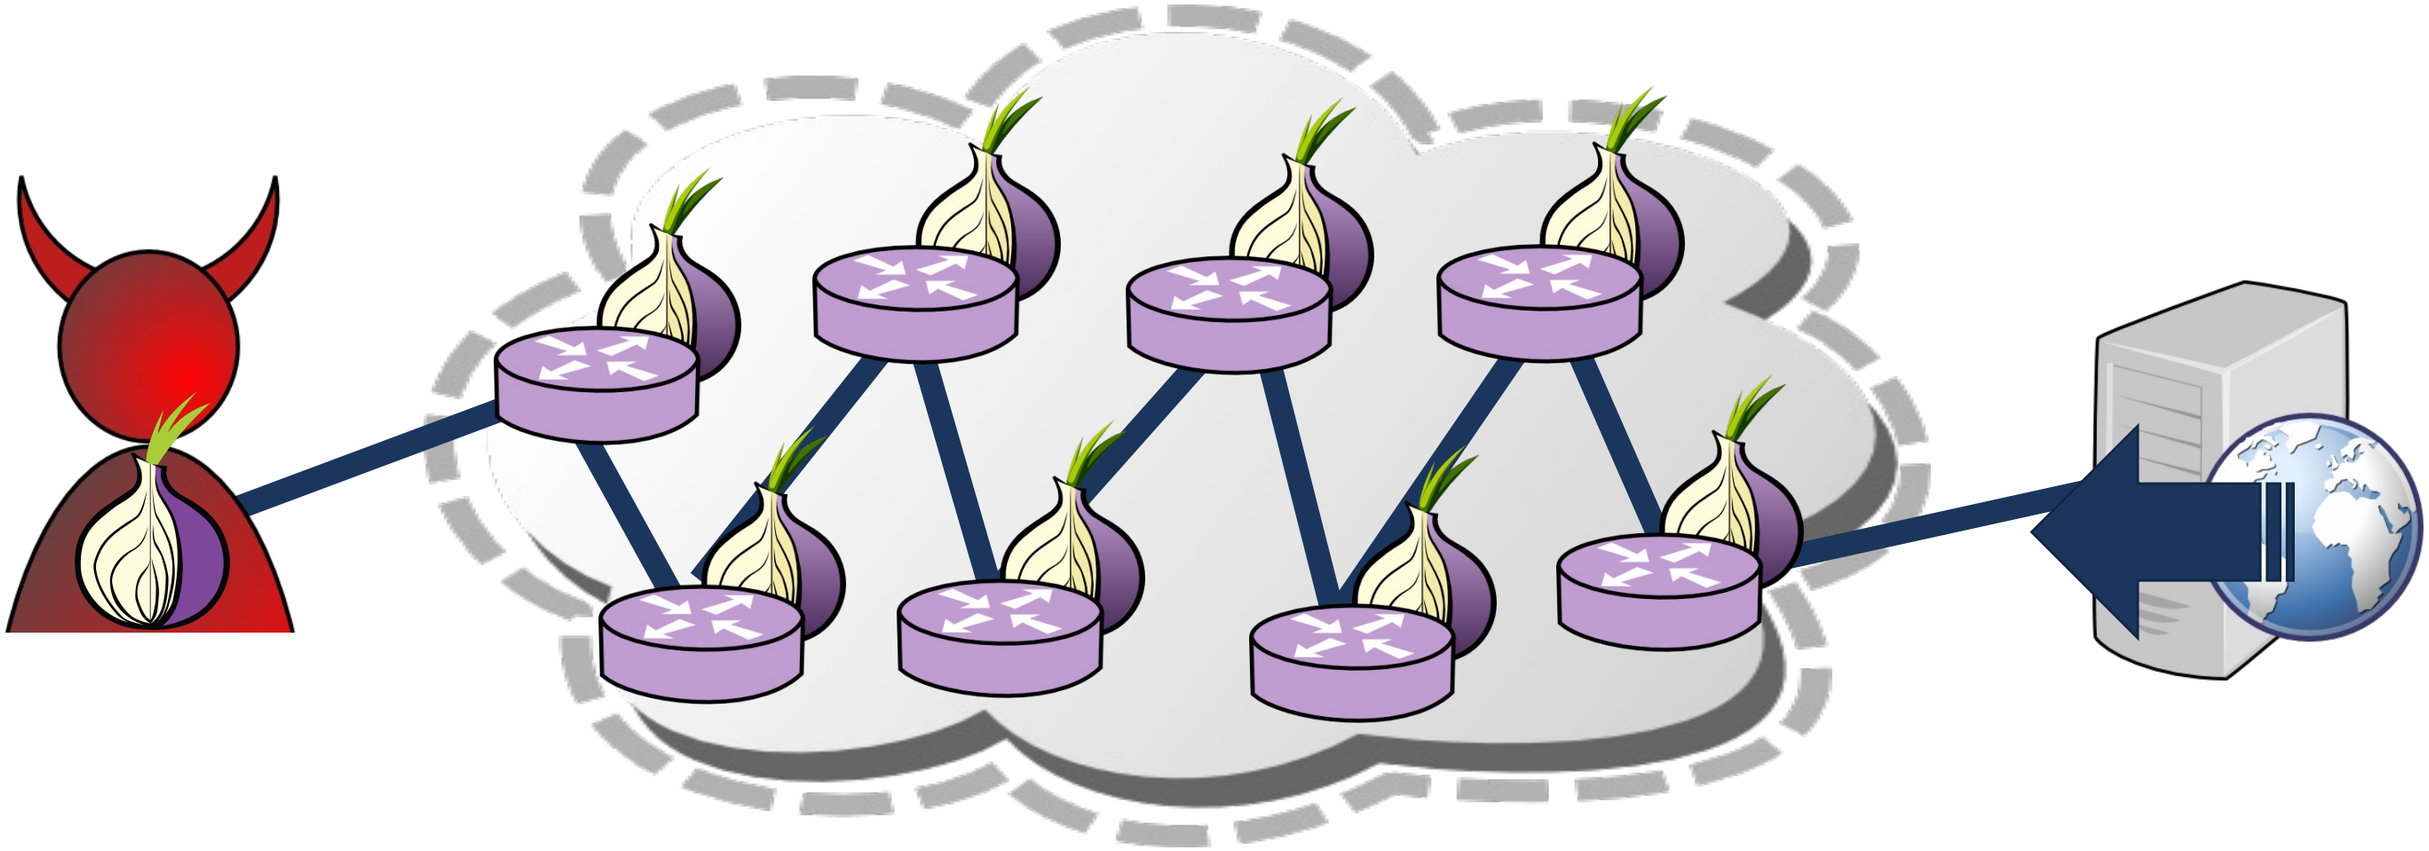

The Relay Congestion Attack

Step 1:
Build 8-hop circuit

Step 2:
GET large files

Step 3:
Stop reading

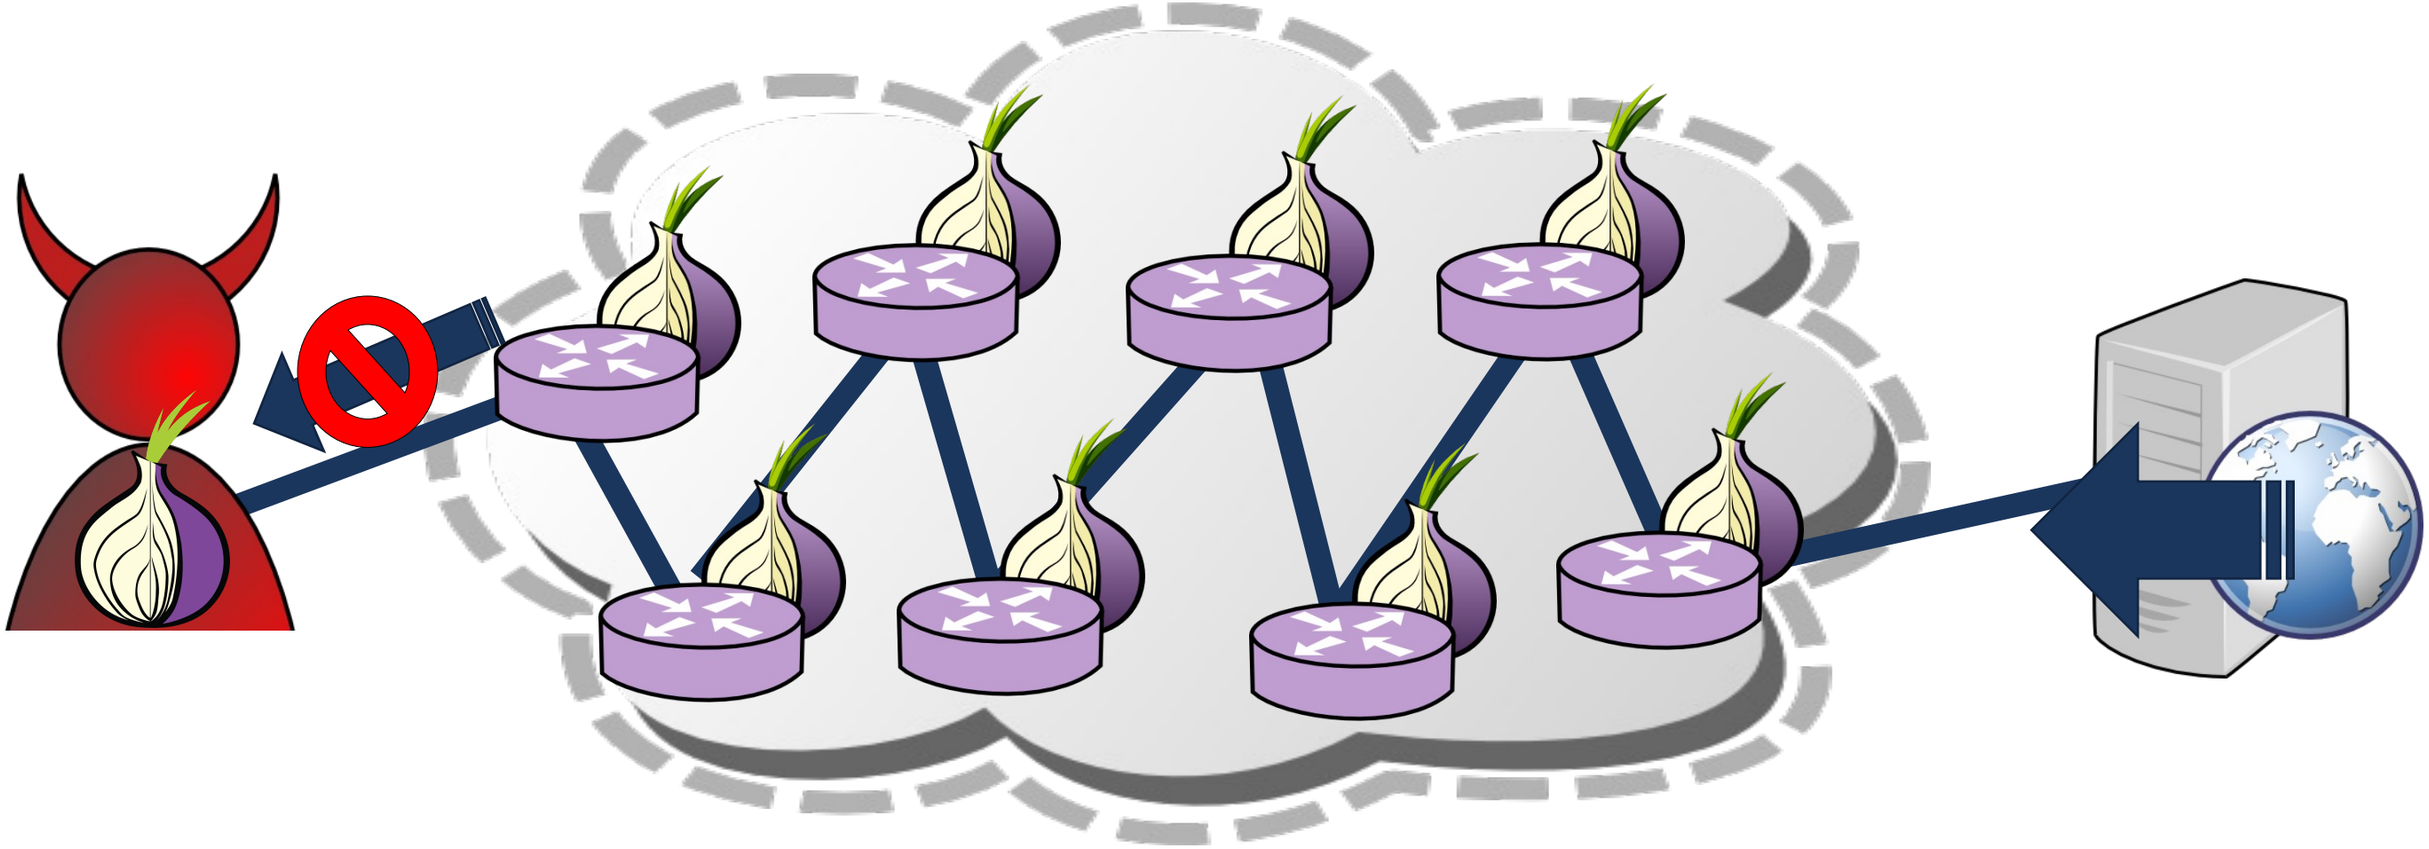

U.S. NAVAL RESEARCH LABORATORY

The Relay Congestion Attack

Step 1:
Build 8-hop circuit

Step 2:
GET large files

Step 3:
Stop reading

Step 4:
Send flow control cells

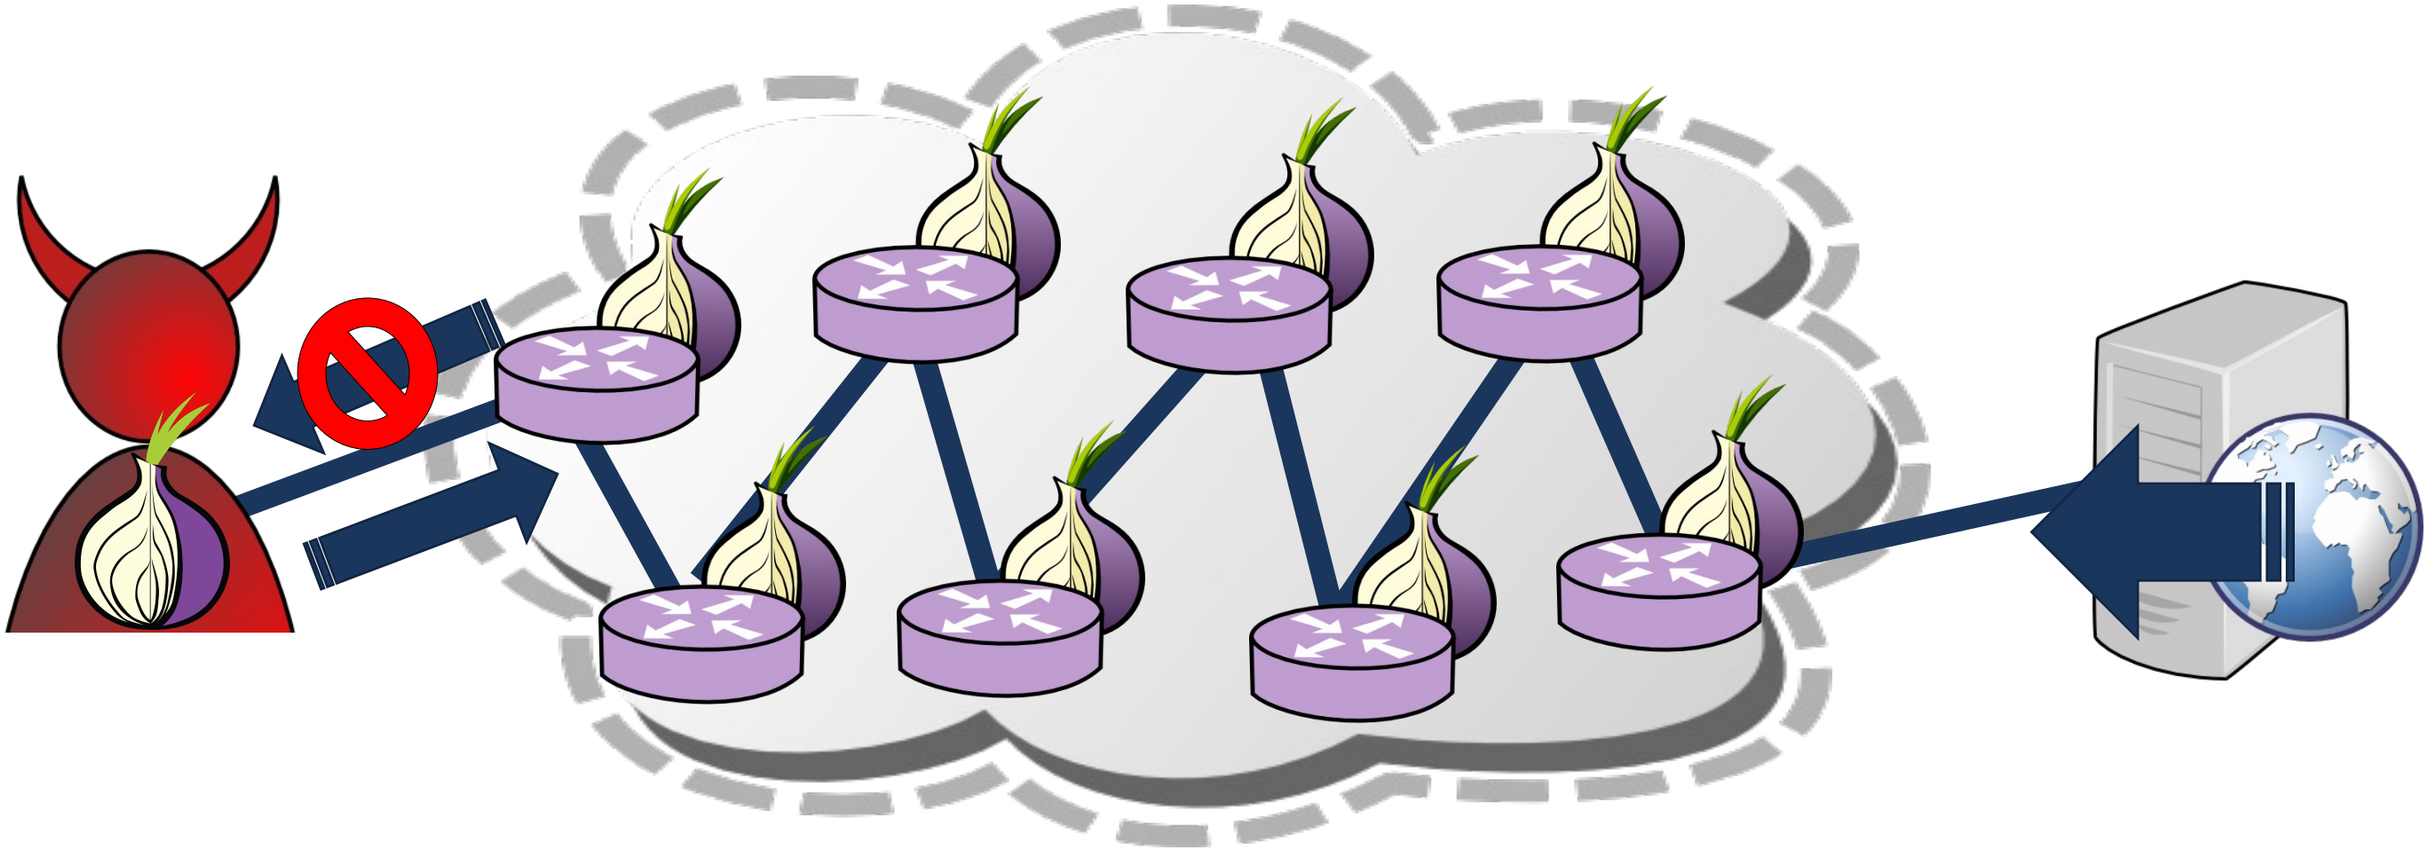

The Relay Congestion Attack

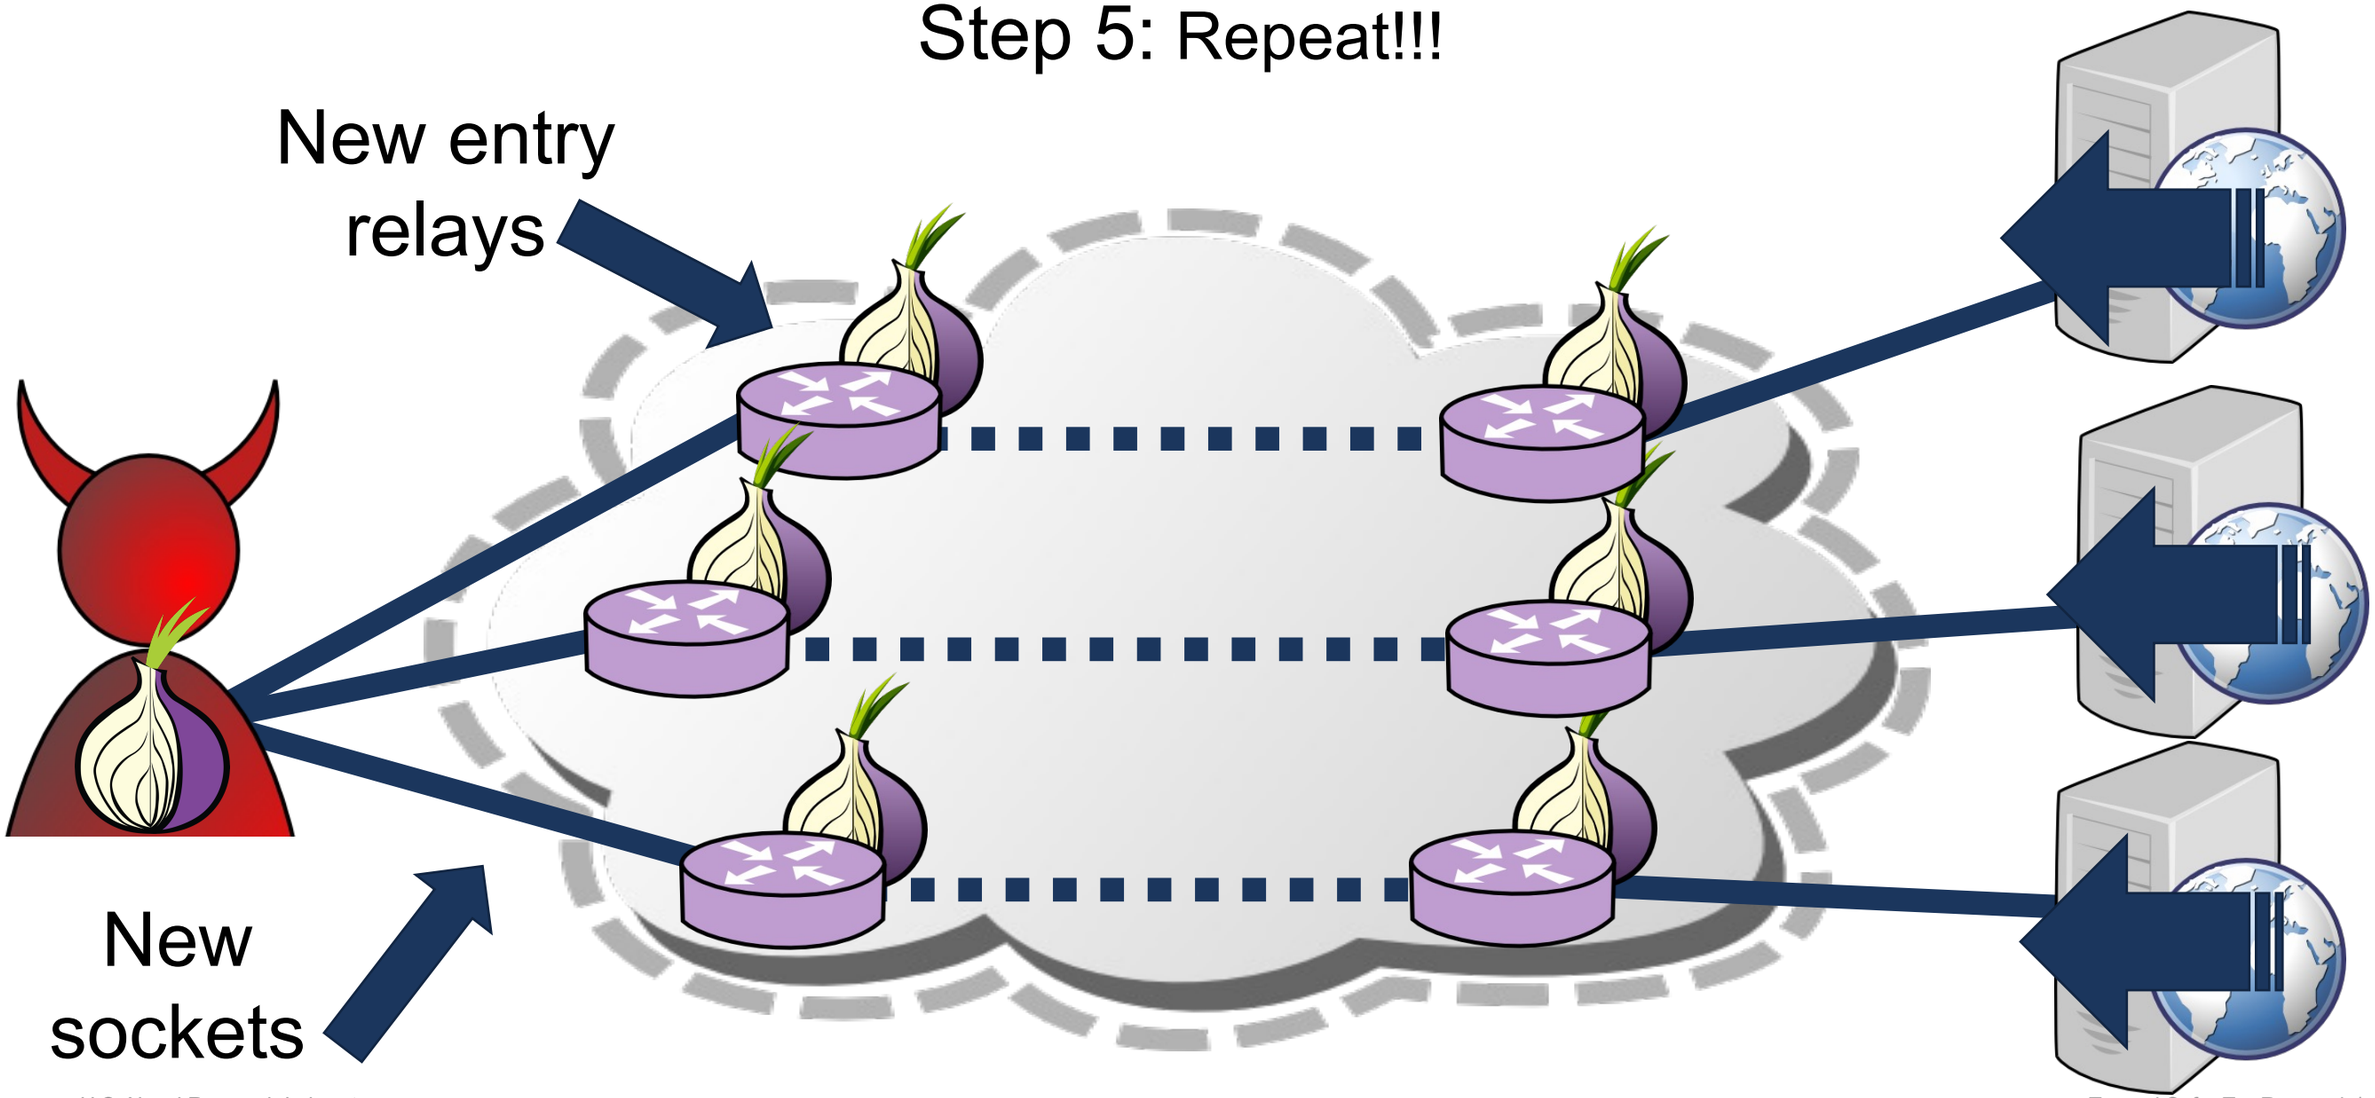

Evaluation Setup

Use Shadow for evaluation

- Private Tor network for safety
- 634 relays (10% size, capacity of Tor)
- 15,000 clients and 2,000 servers generating traffic through Tor

Explore network effects

- Attack strength (num. attack circuits)
- Network load, attacker resource usage, client performance

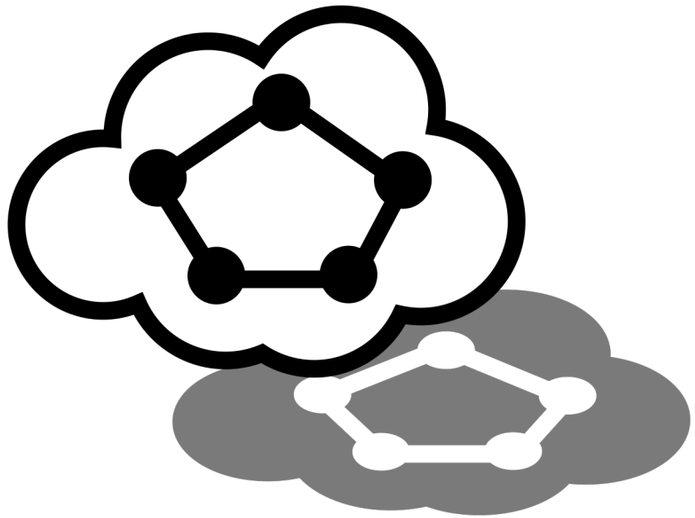

https://shadow.github.io

U.S. NAVAL RESEARCH LABORATORY

Bandwidth Used by Attacker and Tor Network

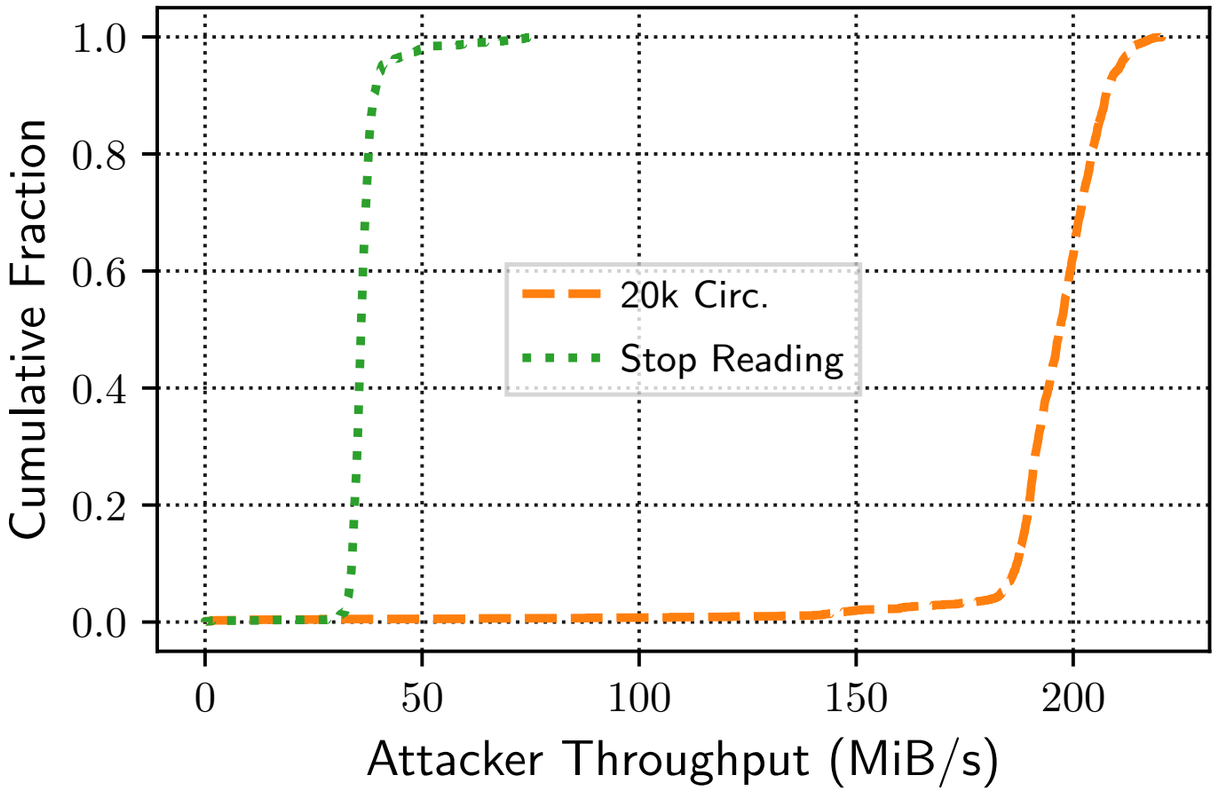

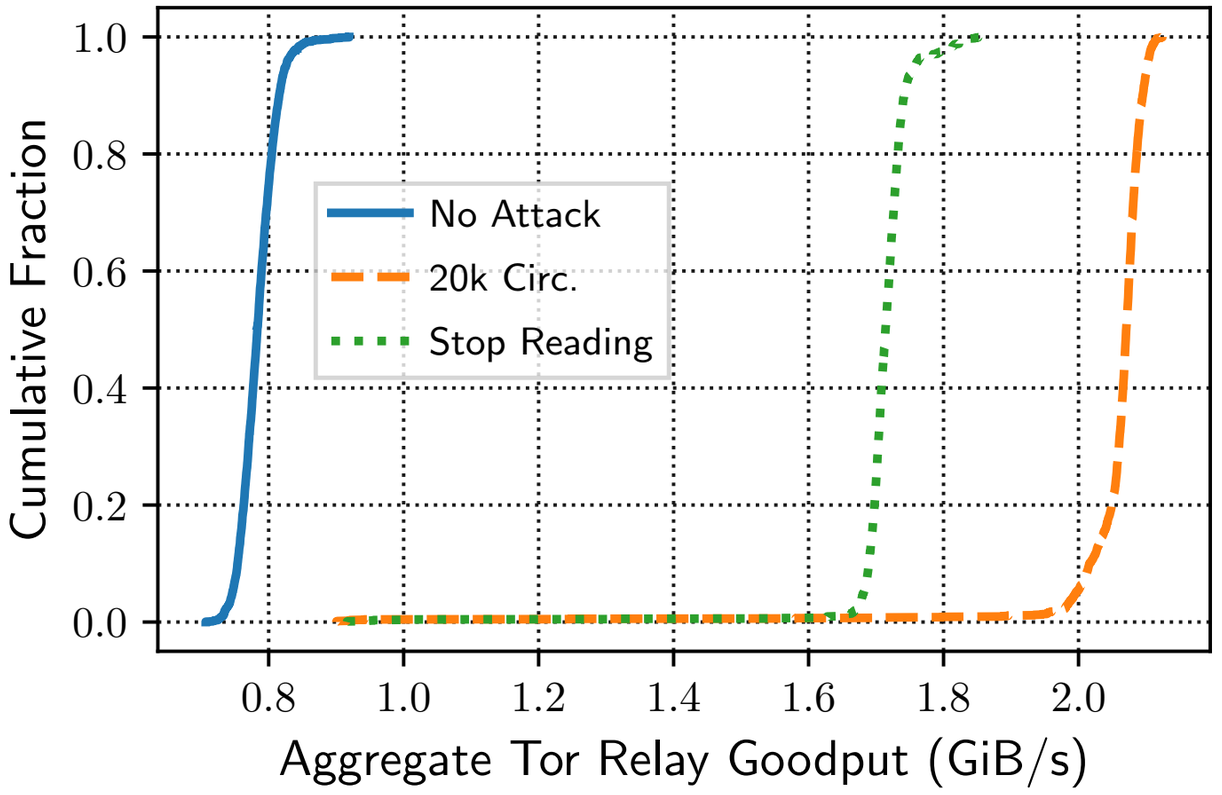

Bandwidth Used by Attacker and Tor Network

Bandwidth Amplification Factor:

Stop Reading

26

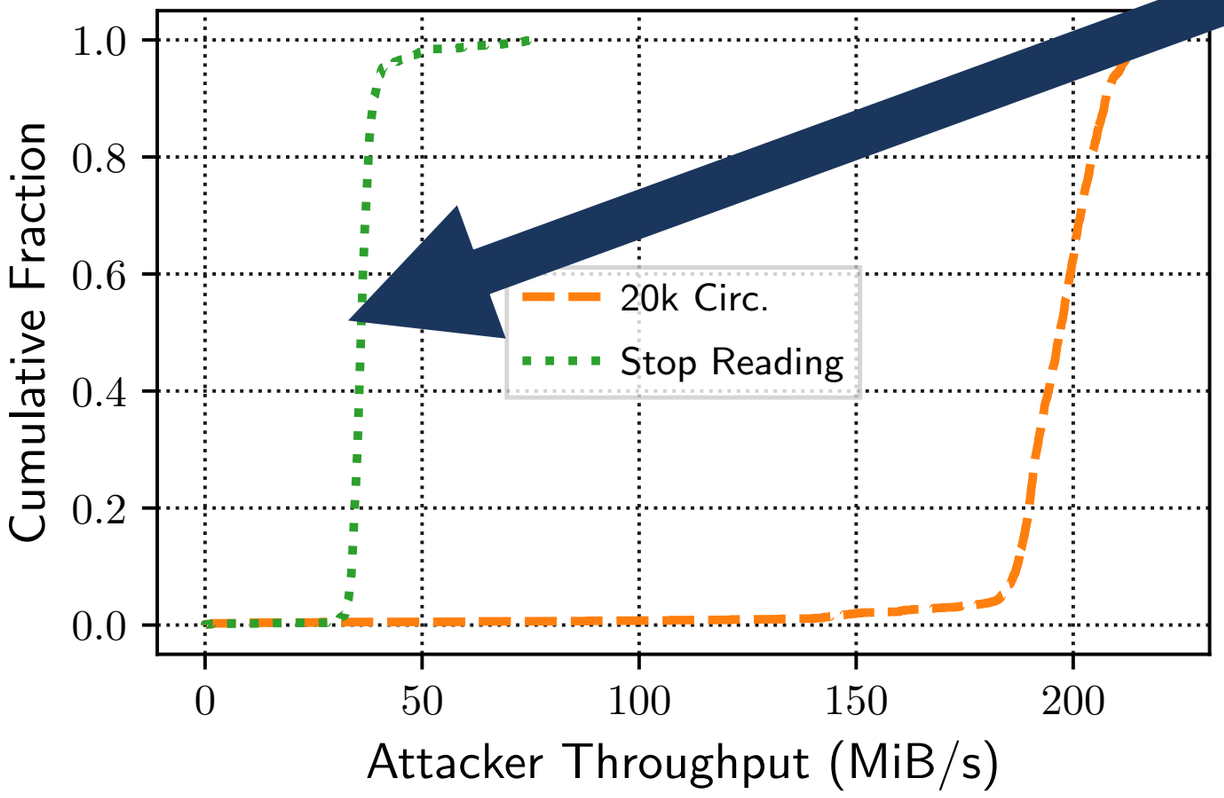

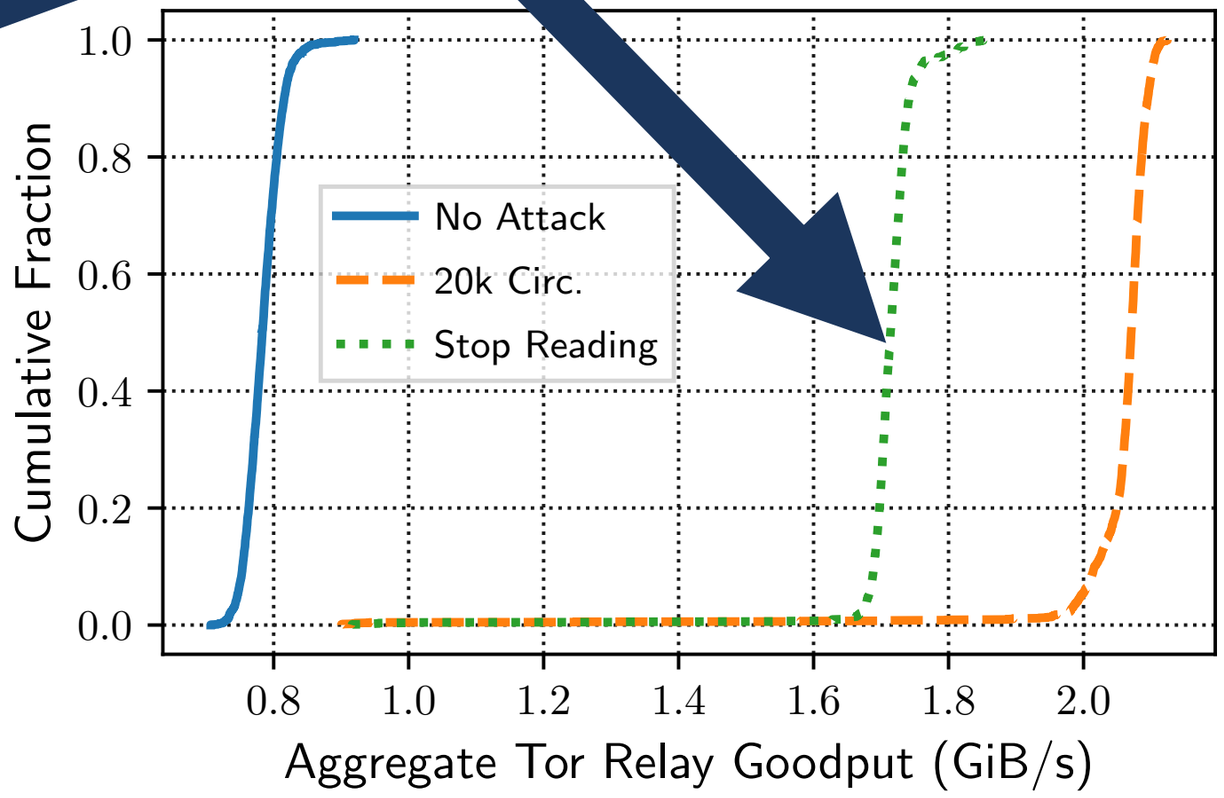

U.S. NAVAL RESEARCH LABORATORY

Effect on Client Performance

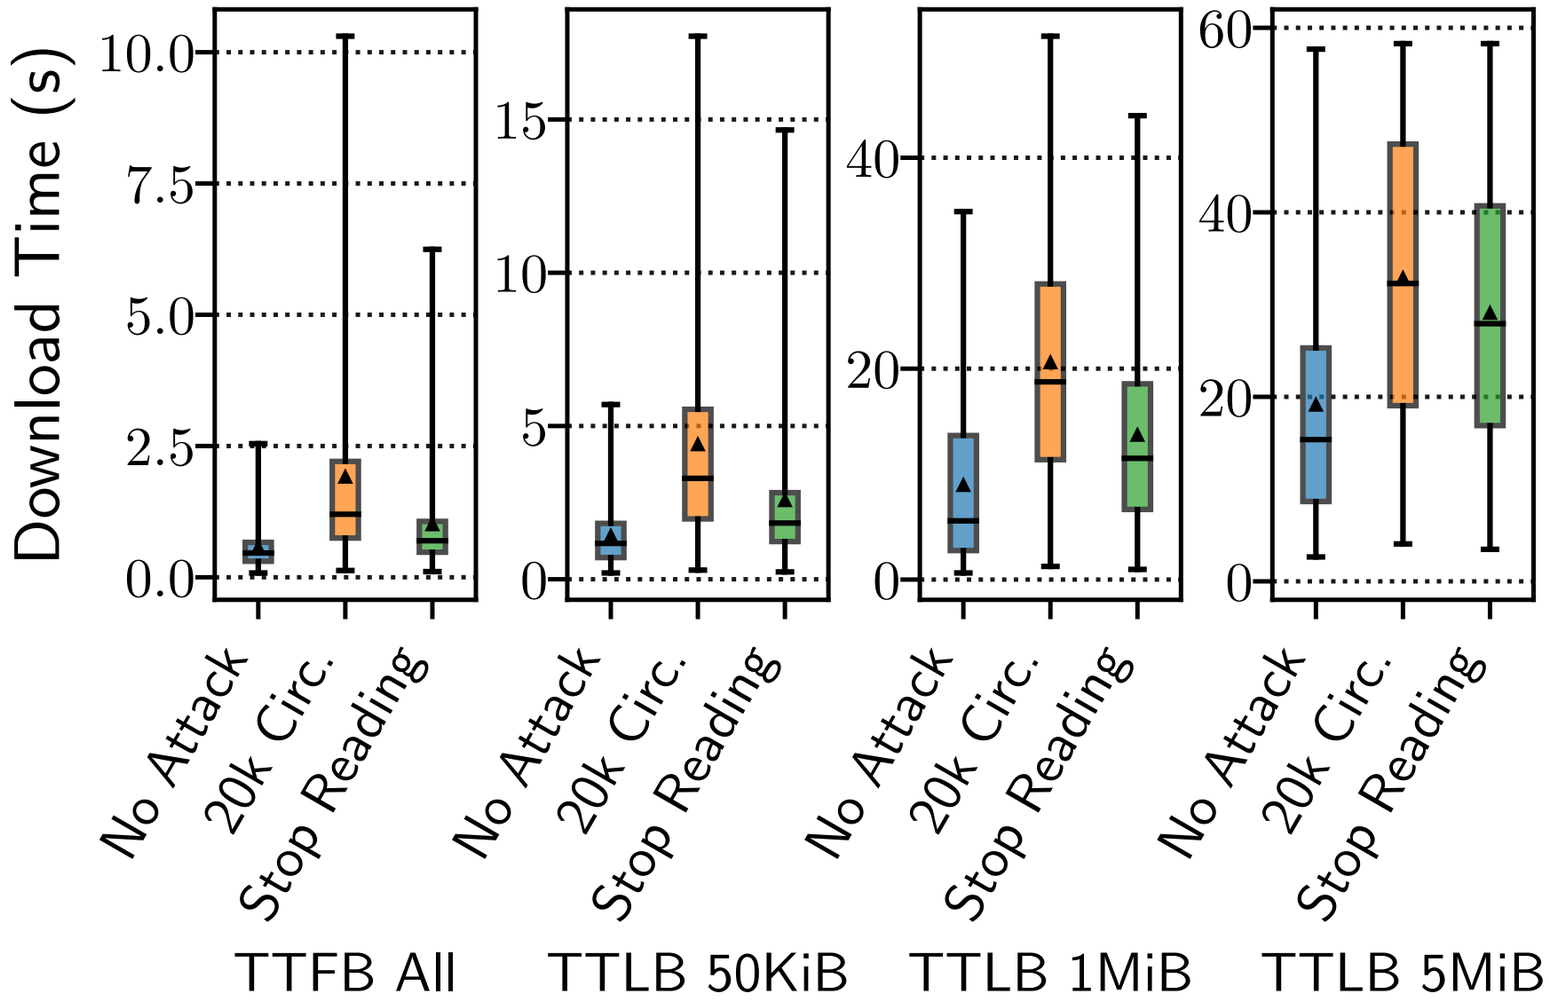

U.S. NAVAL RESEARCH LABORATORY

Effect on Client Performance

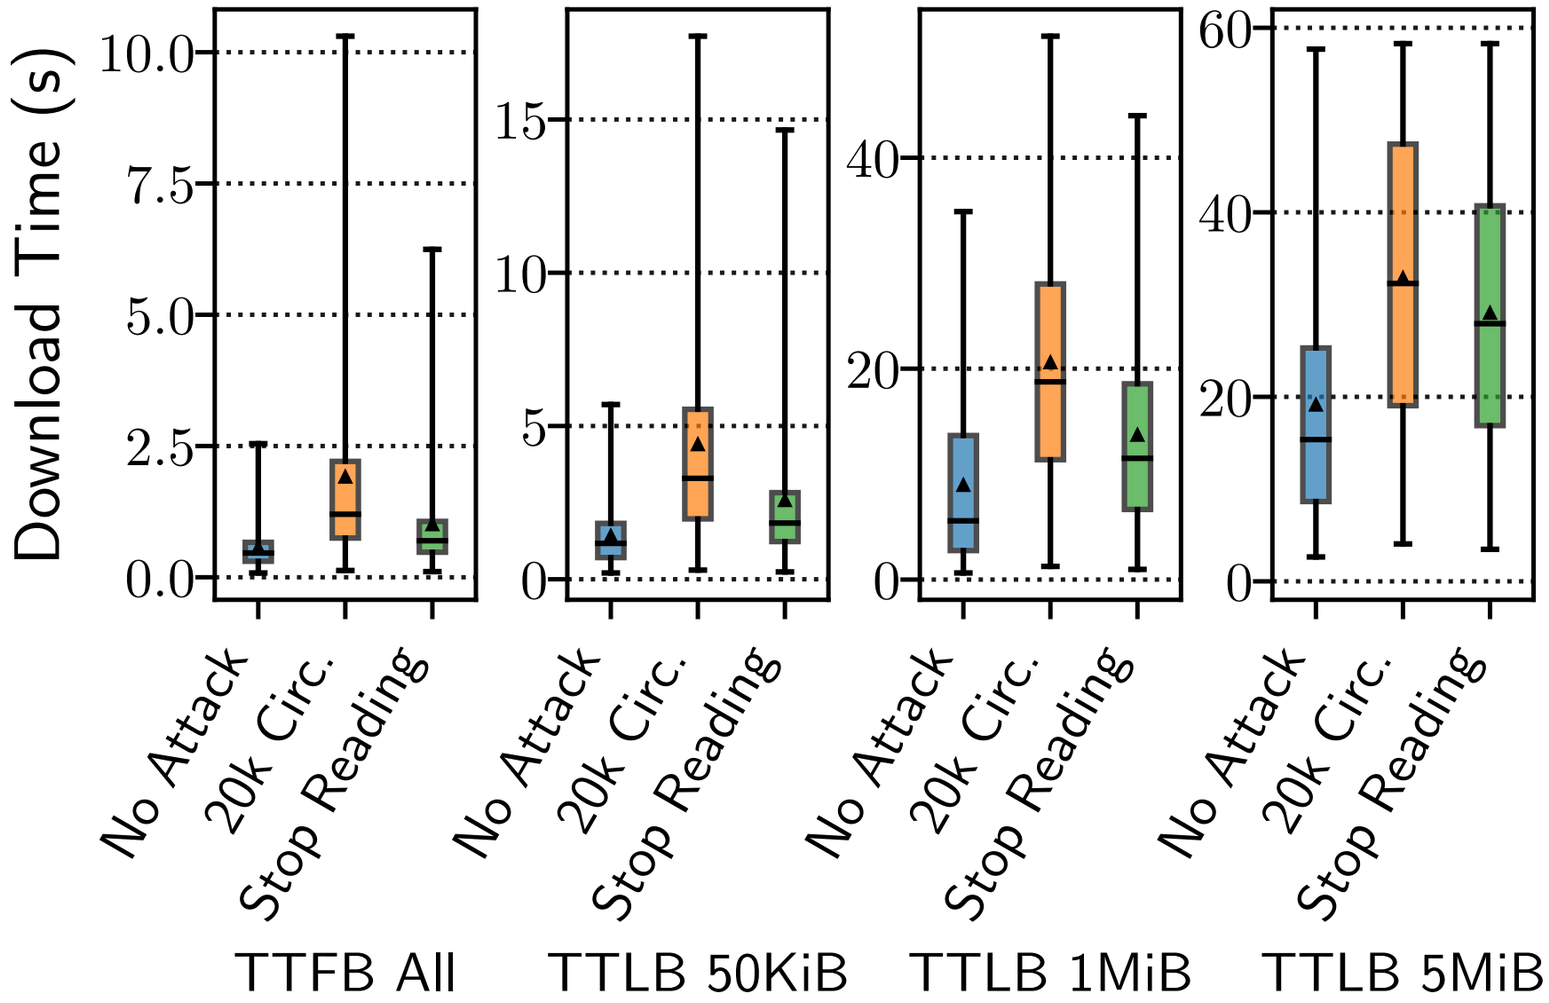

20k Circuits

TTLB: +120%

Stop Reading

TTLB: +47%

Mitigations to Relay Congestion Attack

Ability to stop reading from circuits

- Authenticated SENDMEs, Tor Proposal 289, implemented in 0.4.1.1-alpha

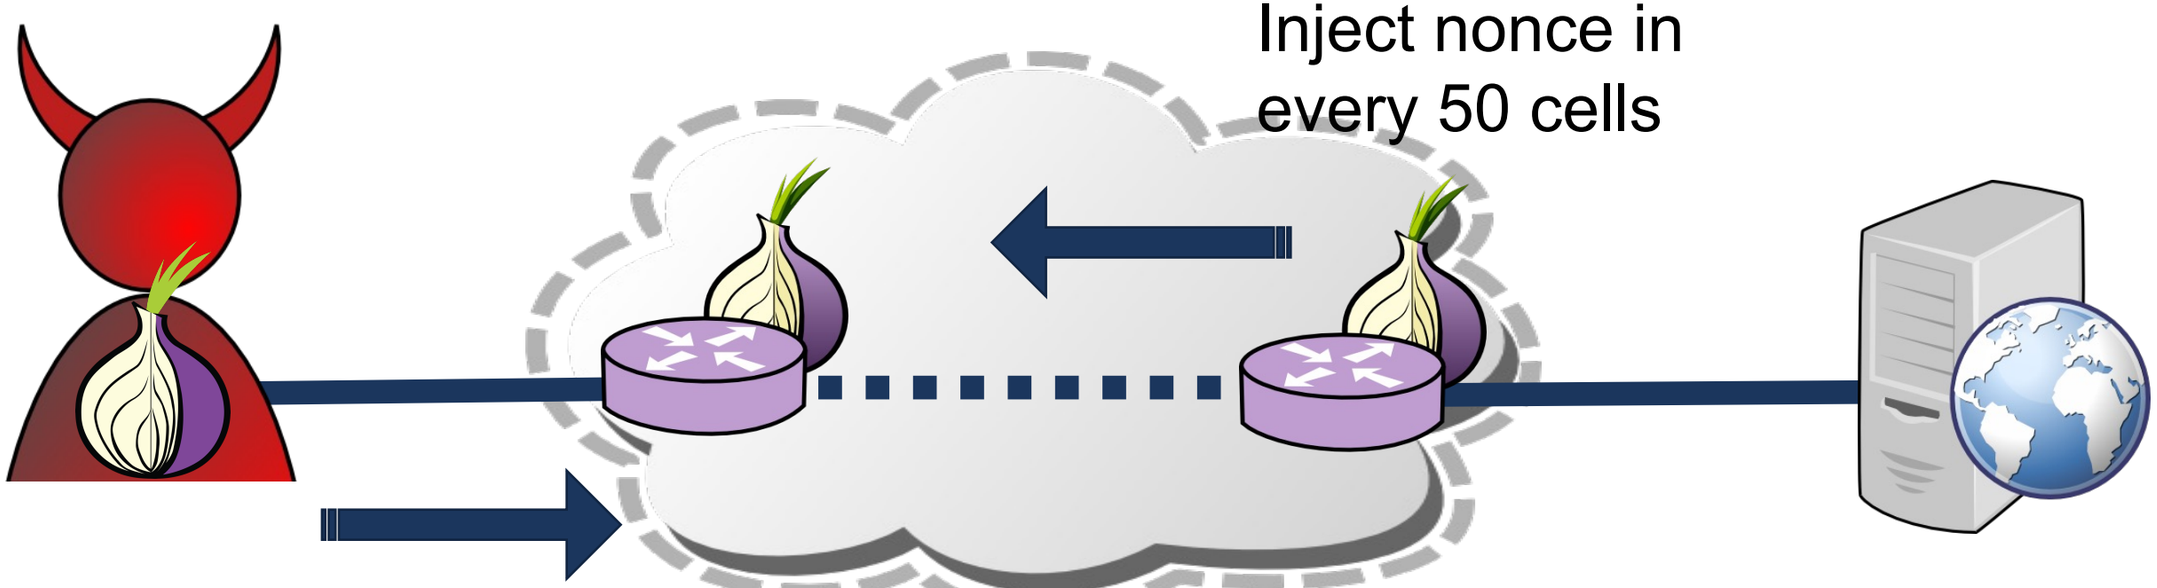

Outline

- **Tor Safety Board overview**
- **Research areas relevant to safety**
  - Simulating Tor with Shadow
  - Measuring Tor with PrivCount
  - Safe research applications
- **Ongoing struggles**

Paper from SigComm'15

Statement from the SIGCOMM 2015 Program Committee: The SIGCOMM 2015 PC appreciated the technical contributions made in this paper, but found the paper controversial because some of the experiments the authors conducted raise ethical concerns. The controversy arose in large part because the networking research community does not yet have widely accepted guidelines or rules for the ethics of experiments that measure online censorship. In accordance with the published submission guidelines for SIGCOMM 2015, had the authors not engaged with their Institutional Review Boards (IRBs) or had their IRBs determined that their research was unethical, the PC would have rejected the paper without review. But the authors did engage with their IRBs, which did not flag the research as unethical. The PC hopes that discussion of the ethical concerns these experiments raise will advance the development of ethical guidelines in this area. It is the PC's view that future guidelines should include as a core principle that researchers should not engage in experiments that subject users to an appreciable risk of substantial harm absent informed consent. The PC endorses neither the use of the experimental techniques this paper describes nor the experiments the authors conducted.

Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests

Sam Burnett
School of Computer Science, Georgia Tech
sam.burnett@gatech.edu

Nick Feamster
Department of Computer Science, Princeton
feamster@cs.princeton.edu

U.S. NAVAL RESEARCH LABORATORY

Quote from the “Encore” paper

Balancing the benefit and risk of measuring filtering with Encore is difficult.

- *This paper has made the benefit clear: Encore enables researchers to collect new data about filtering from a diversity of vantage points that was previously prohibitively expensive to obtain and coordinate. Ongoing efforts to measure Web filtering would benefit from Encore’s diversity and systematic rigor [8, 15, 35].*
- *The risk that Encore poses are far more nebulous: laws against accessing filtered content vary from country to country, and may be effectively unenforceable given the ease with which sites (like Encore) can request cross-origin resources without consent; there is no ground truth about the legal and safety risks posed by collecting network measurements.*

U.S. NAVAL RESEARCH LABORATORY

Ongoing Struggles

- How to balance benefits vs. risks??
- Technologists are good at the benefits
  - They are experts on their technology
  - Taught to market their work, to show it in the best light
  - Usually optimistic, best-case analysis
- Technologists are bad at risks
  - Should consider the worst-case
  - Even one person put in jail because of a search is already too much harm
- We should think about risks with our security hat on

U.S. NAVAL RESEARCH LABORATORY

Ongoing Struggles

- How do we incentivize lower utility?
  - There is ~always a way to do things safely (e.g., differential privacy)
  - But safety reduces utility
  - If I only had _____ then I could solve _____ !
- We need to break the desire for splashy headlines
- The adversary knows no ethical bounds
  - Will they always have the upper hand?

U.S. NAVAL RESEARCH LABORATORY

Departing Advice

- Think carefully about safety during research design
  - There will be clear choices, and not so clear choices
- Seek advice from colleagues
- Seek advice from experts when possible
- If still unsure, be conservative or choose a different problem