Your Censor is My Censor:
Weaponizing Censorship Infrastructure for Availability Attacks

Kevin Bock Pranav Bharadwaj Jasraj Singh Dave Levin

UNIVERSITY OF MARYLAND

Nation-state censorship

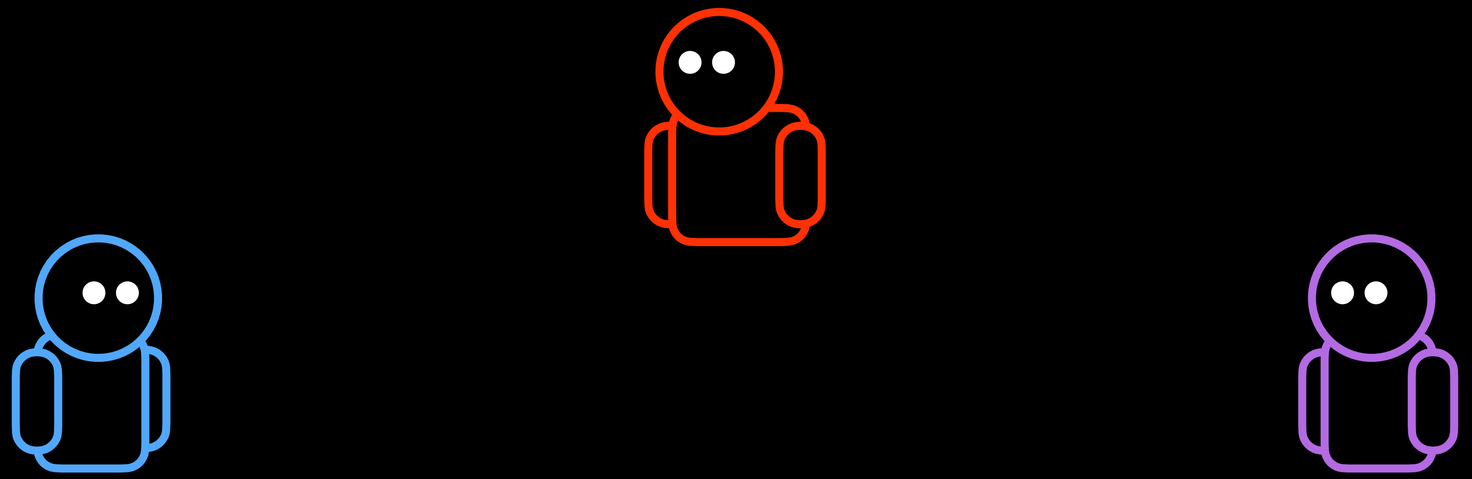

Nation-state censorship

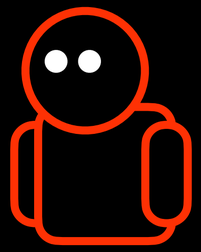

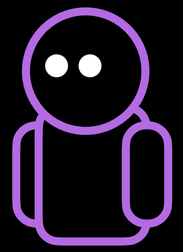

Web browser

Nation-state censorship

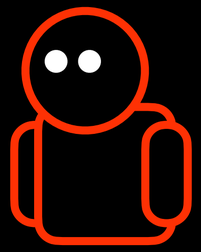

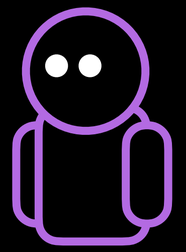

Web browser

Website

Nation-state censorship

Web browser

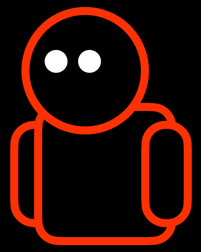

Censor

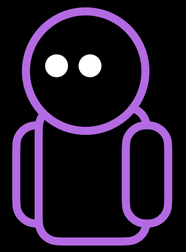

Website

Nation-state censorship

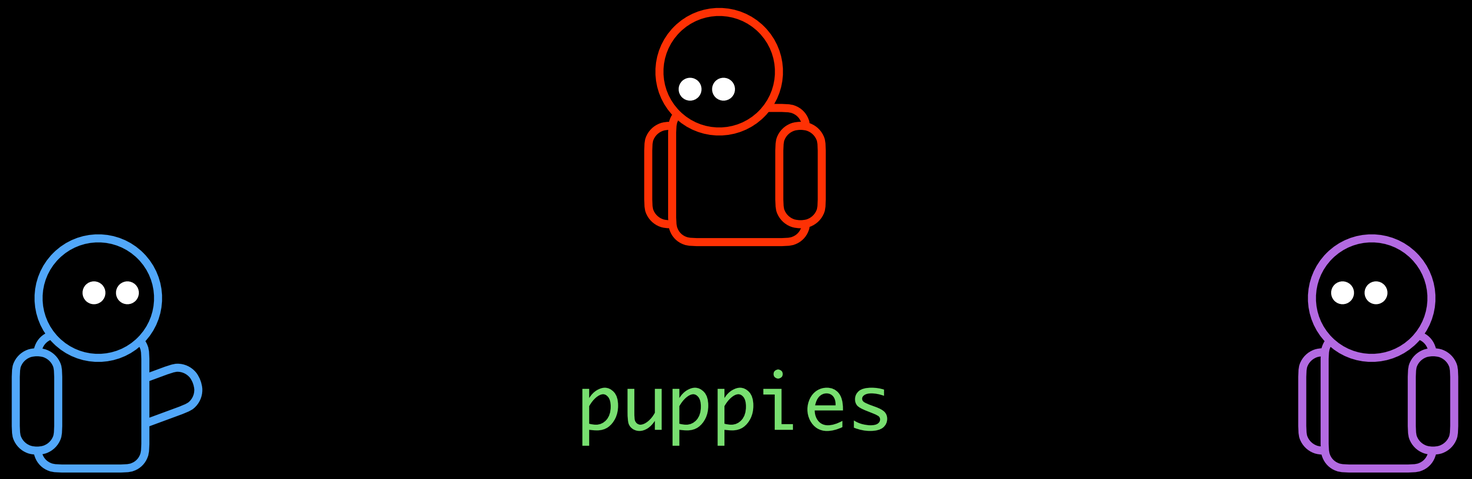

Nation-state censorship

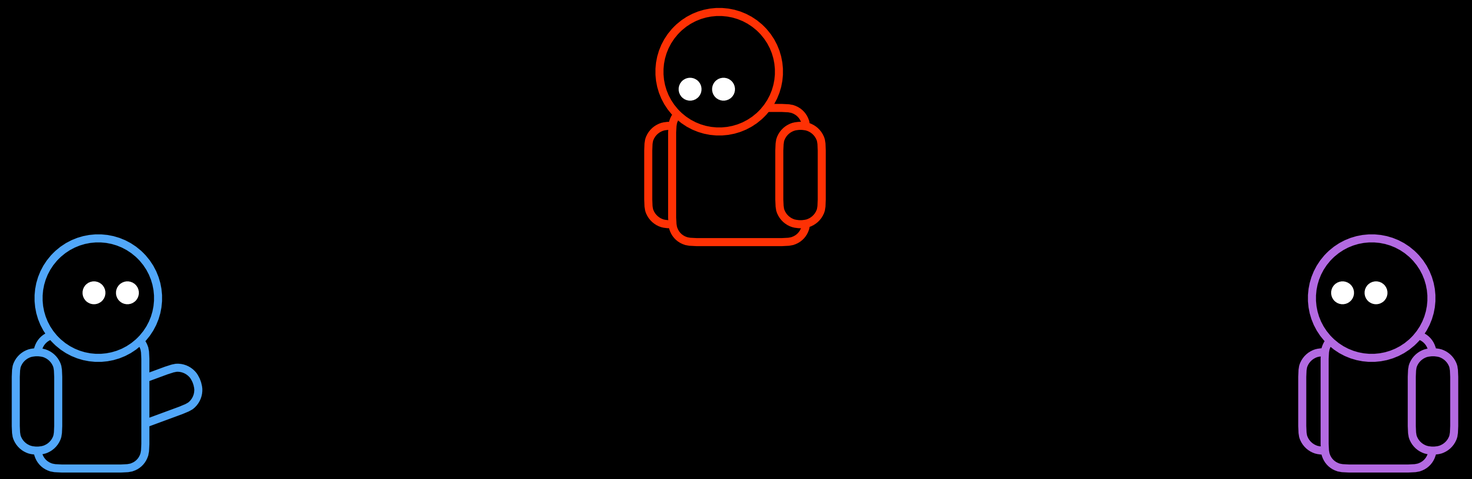

Nation-state censorship

kittens

Nation-state censorship

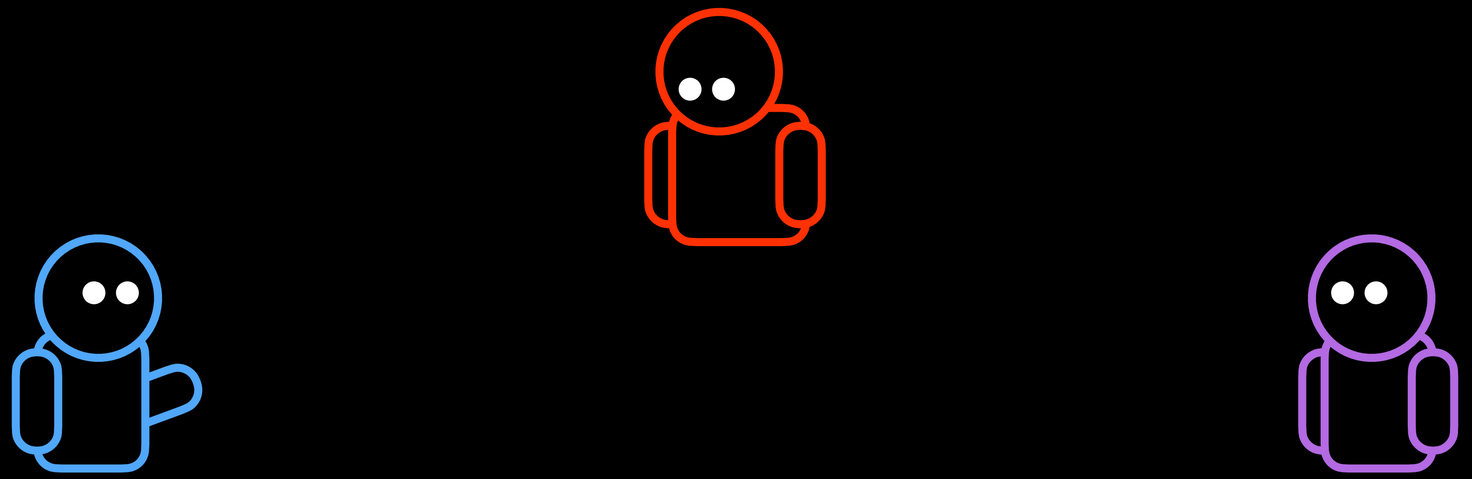

Nation-state censorship

Nation-state censorship

religion

Nation-state censorship

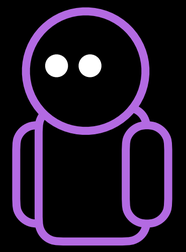

religion

Nation-state censorship

Nation-state censorship

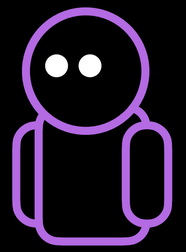

religion

Nation-state censorship

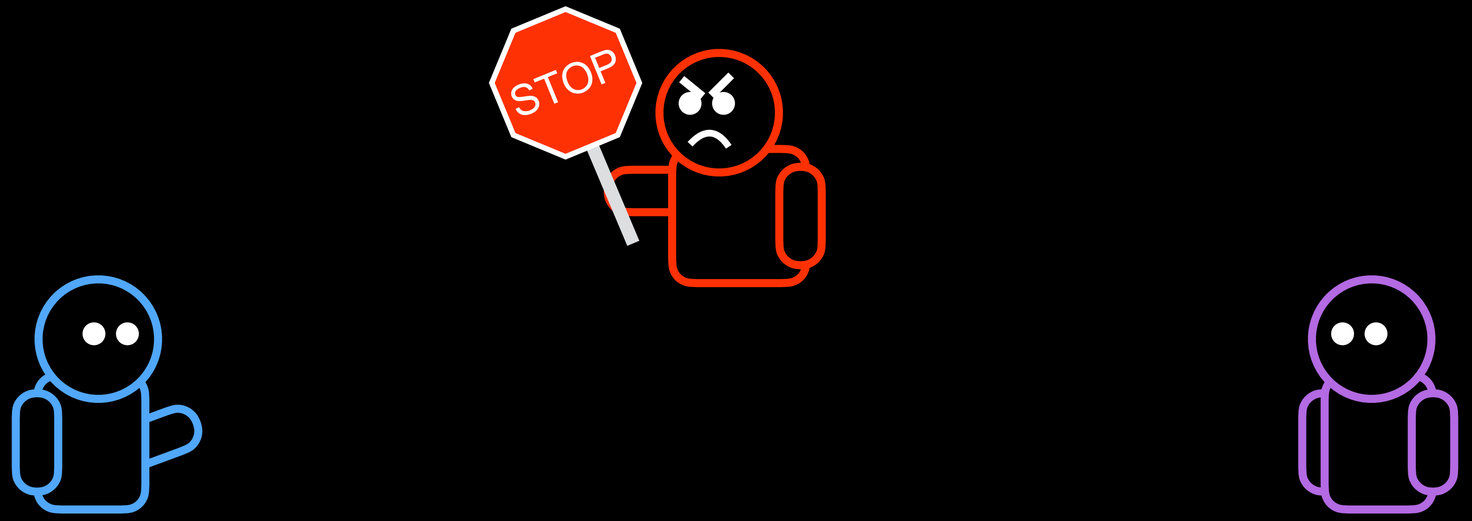

Residual censorship

Block *all* communication for some time

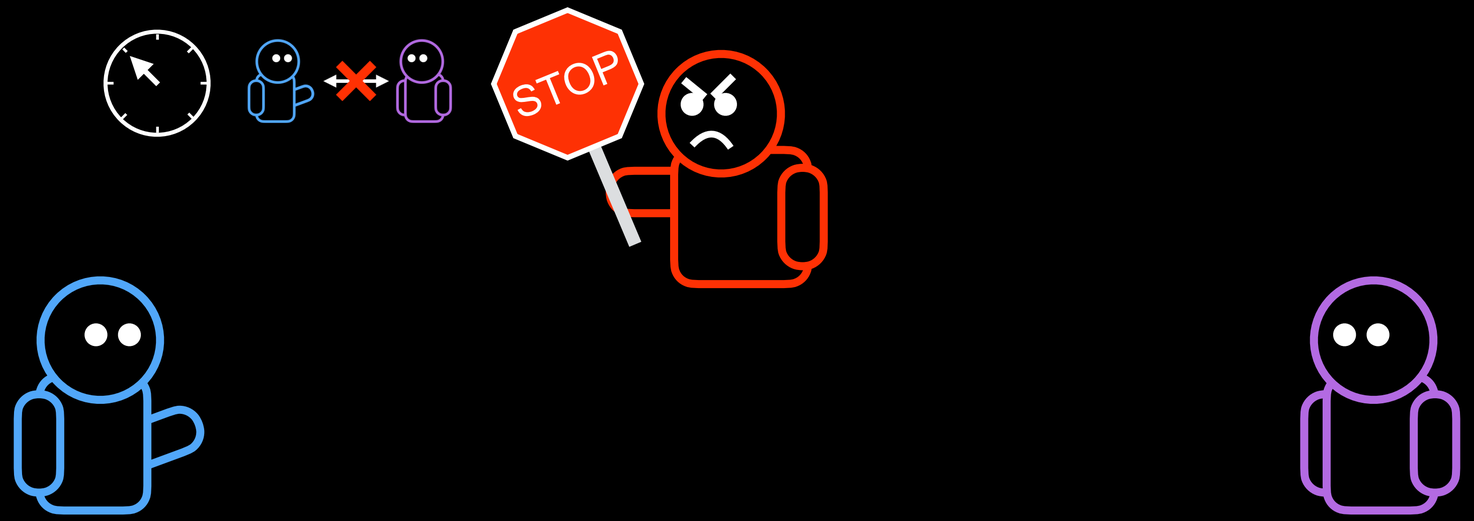

Residual censorship

Block *all* communication for some time

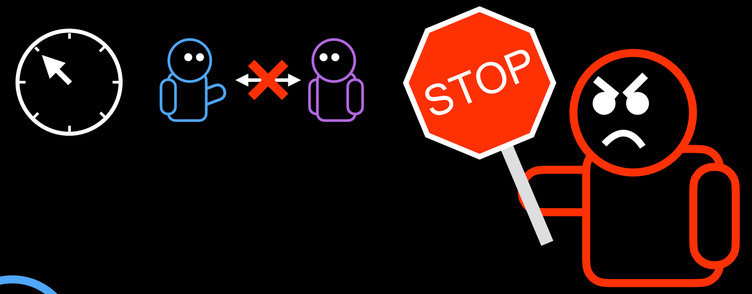

puppies

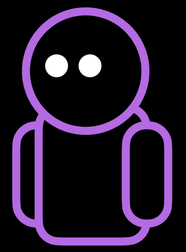

Residual censorship

Block *all* communication for some time

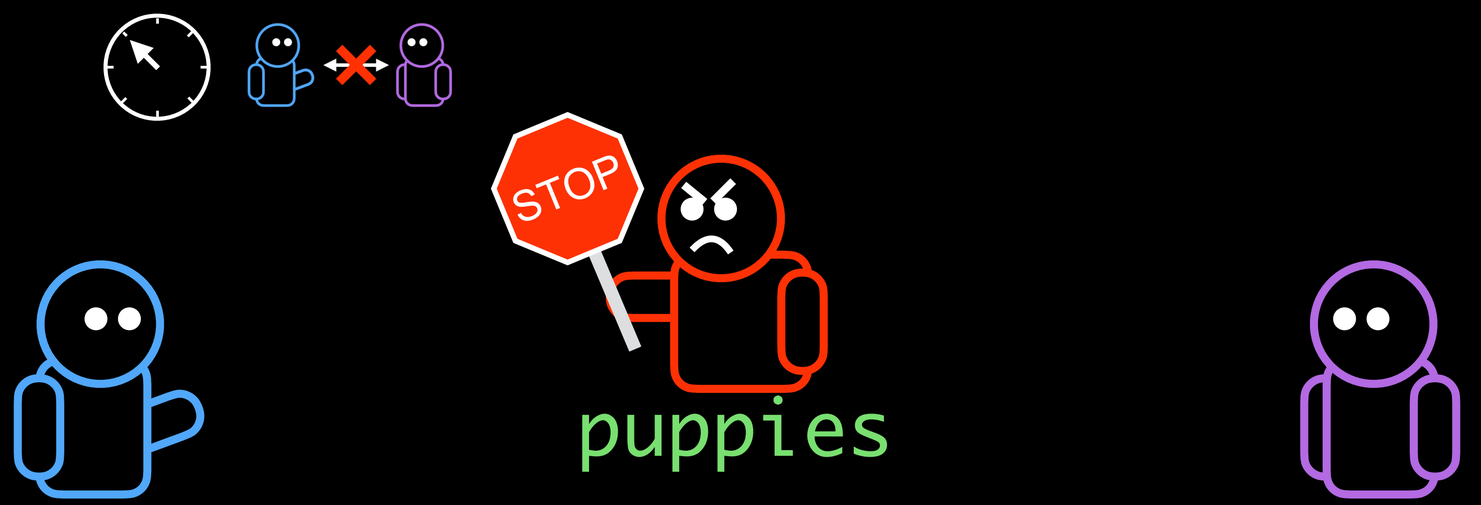

Residual censorship

Block *all* communication for some time

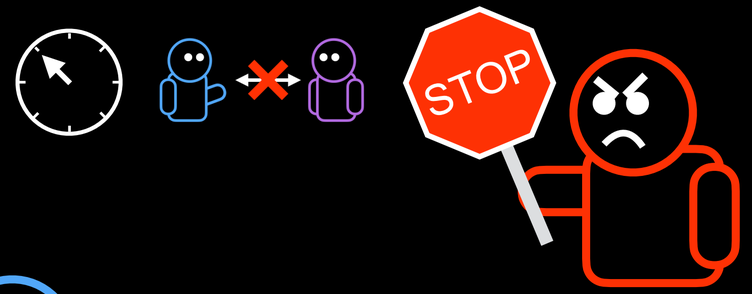

puppies

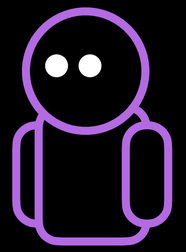

Residual censorship

Block *all* communication for some time

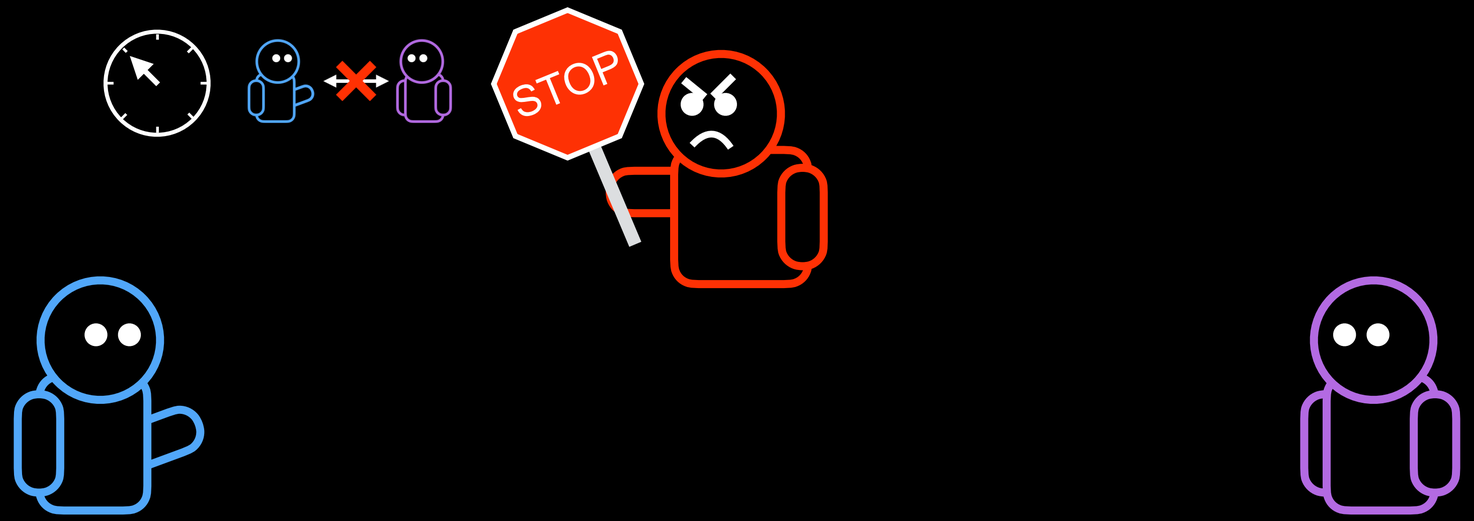

Residual censorship

Block *all* communication for some time

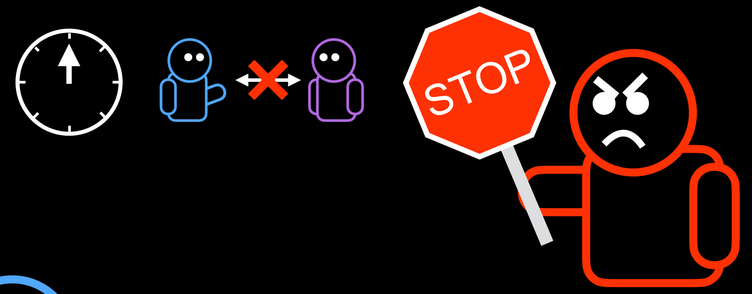

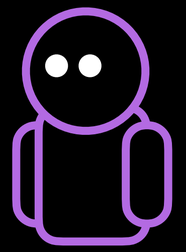

Types of residual censorship

Categorized by what information censor remembers

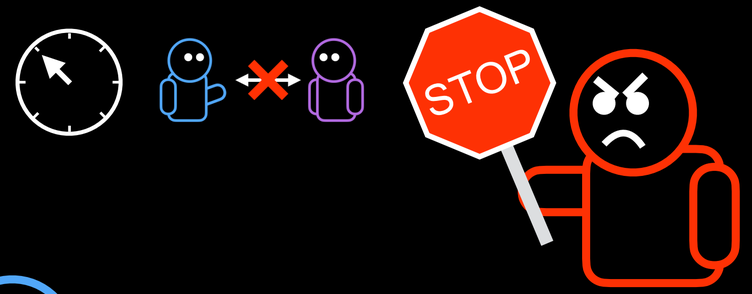

IP:port

IP:port

Types of residual censorship

Categorized by what information censor remembers

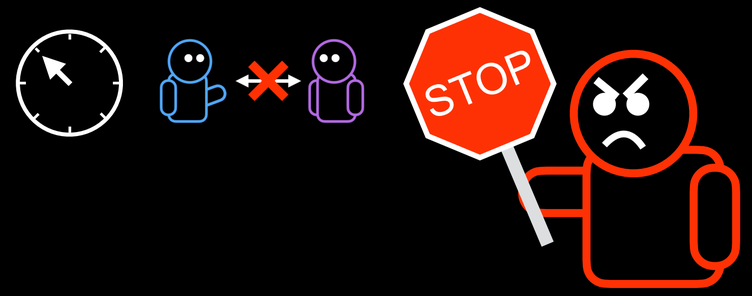

	Source	*Destination*
4-tuple	(IP, port,	IP, port)
3-tuple	(IP, IP,	port)
2-tuple	(IP,	IP)

More aggressive

Censorship infrastructure

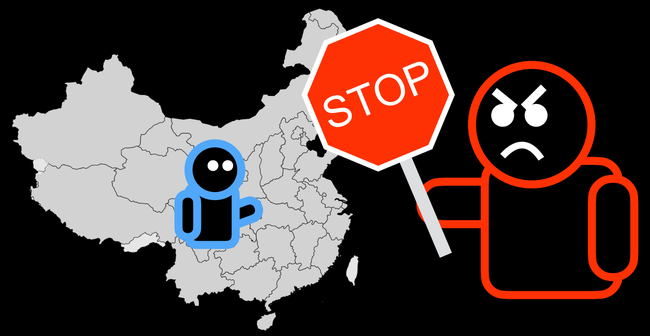

Censorship infrastructure

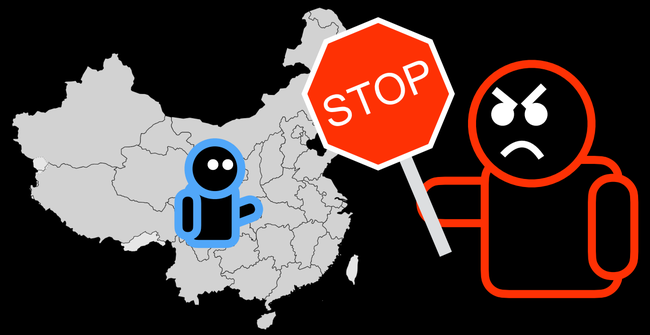

Censorship infrastructure can be weaponized

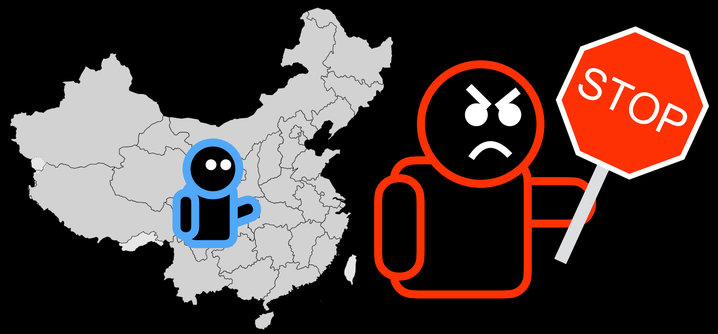

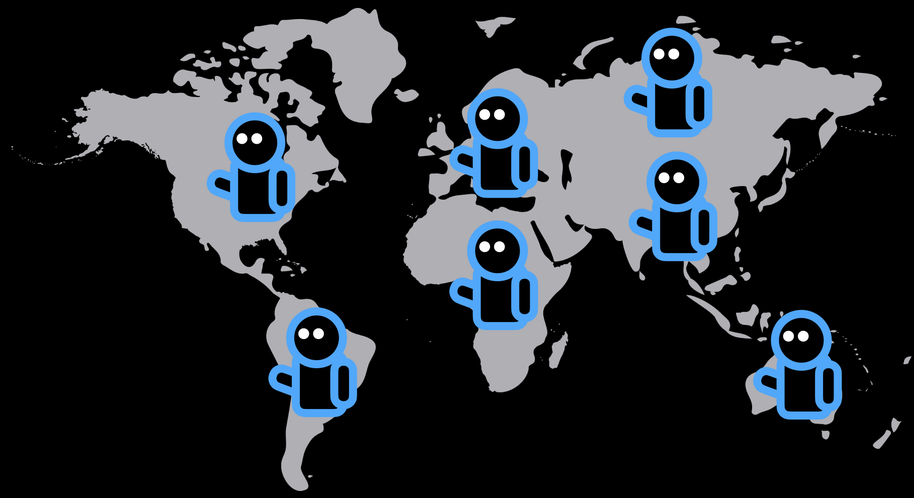

Weaponizing residual censorship

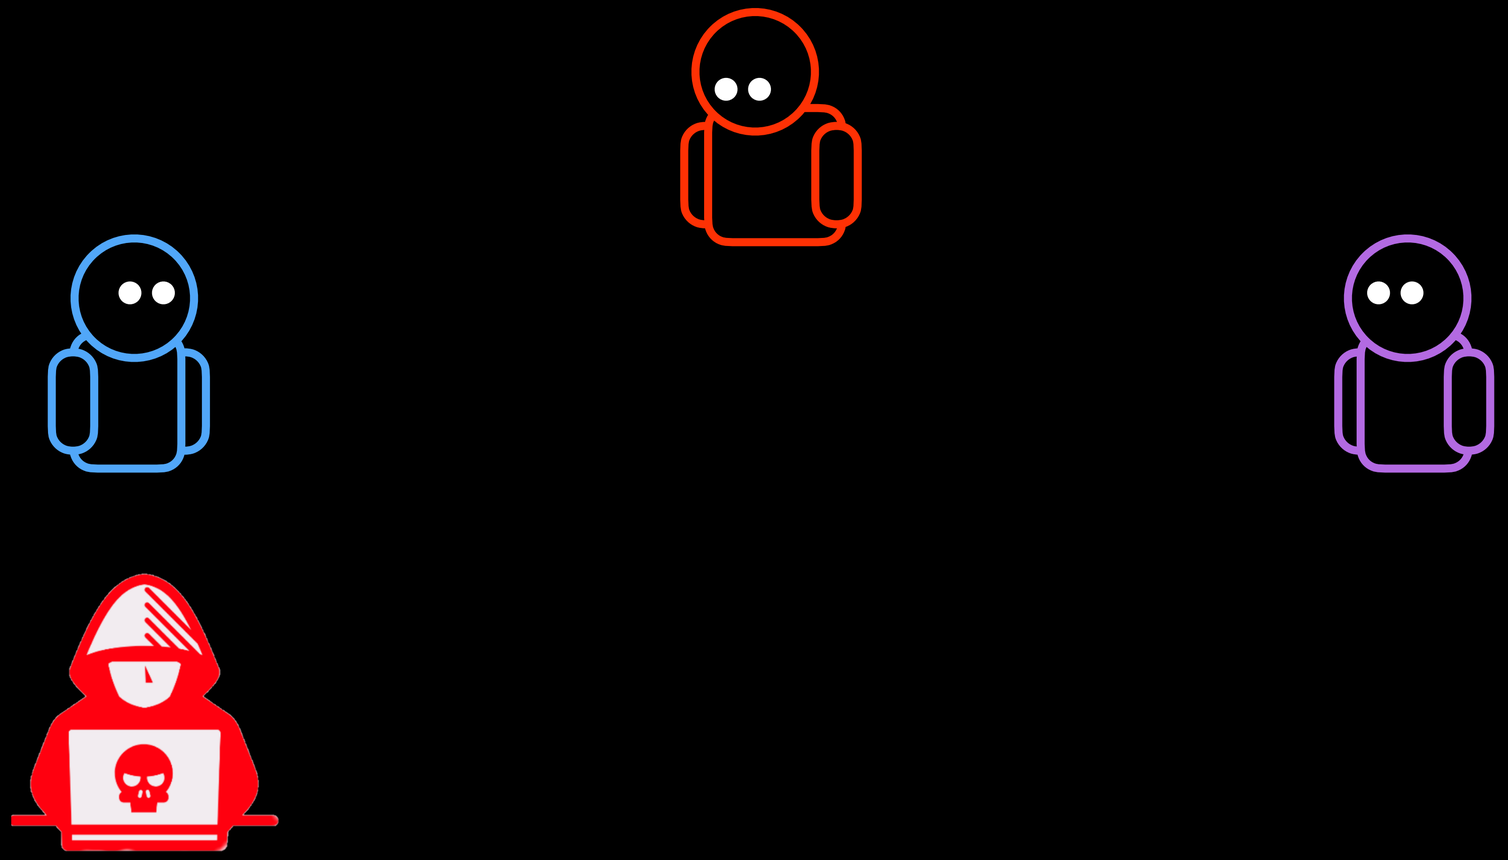

Weaponizing residual censorship

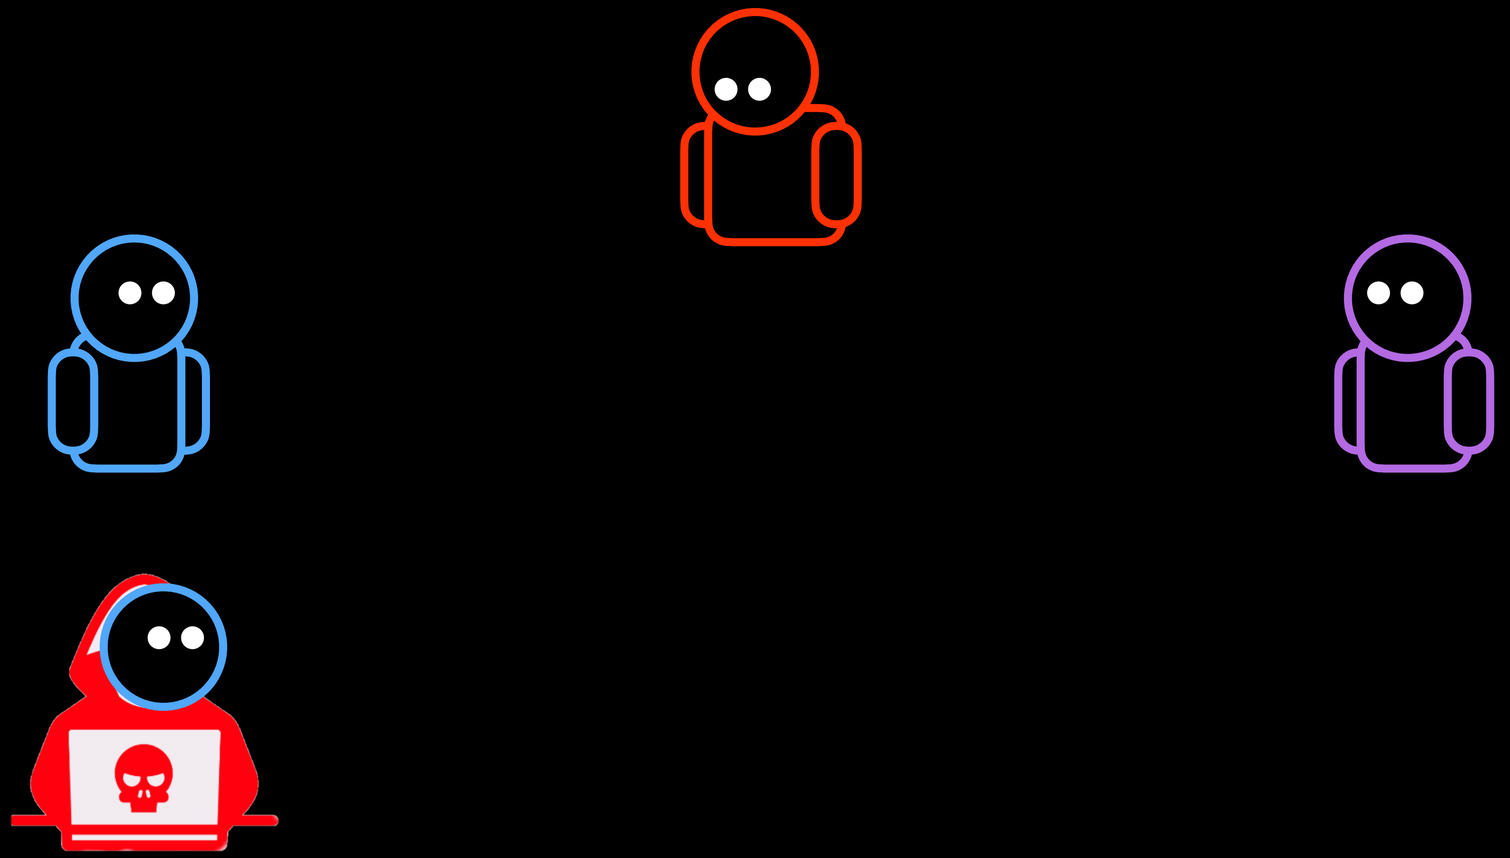

Weaponizing residual censorship

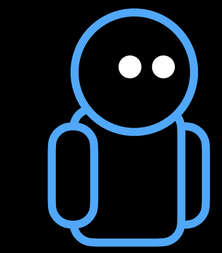

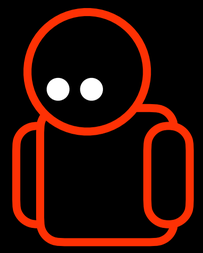

religion

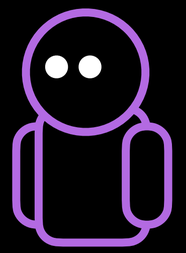

Weaponizing residual censorship

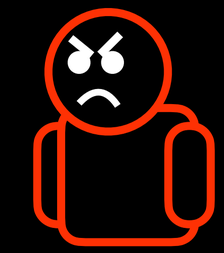

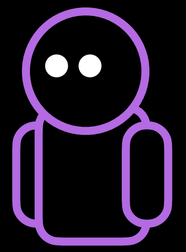

`religion`

Weaponizing residual censorship

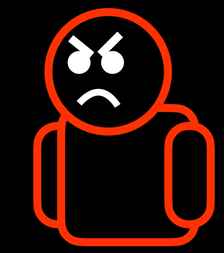

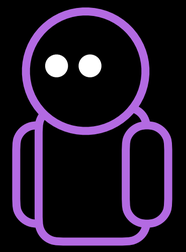

`religion`

Weaponizing residual censorship

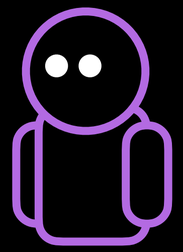

`religion`

Weaponizing residual censorship

Weaponizing residual censorship

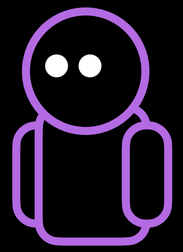

`religion`

Weaponizing residual censorship

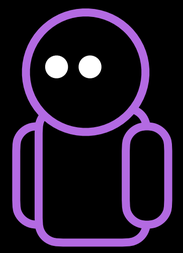

Weaponizing residual censorship

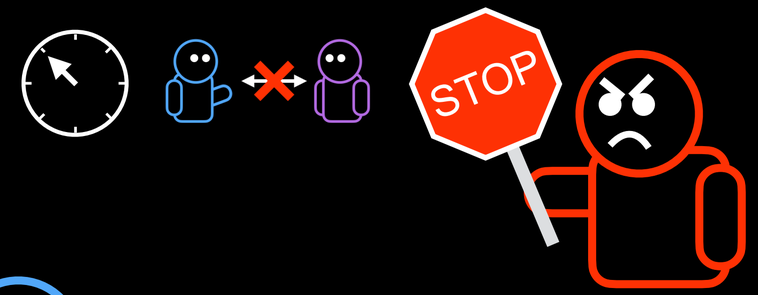

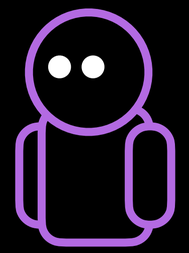

Weaponizing residual censorship

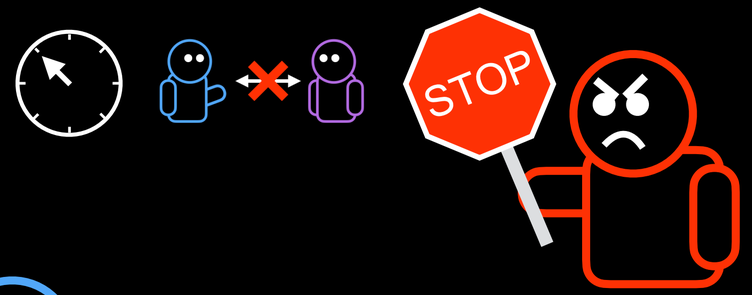

puppies

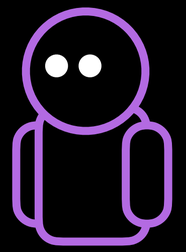

Weaponizing residual censorship

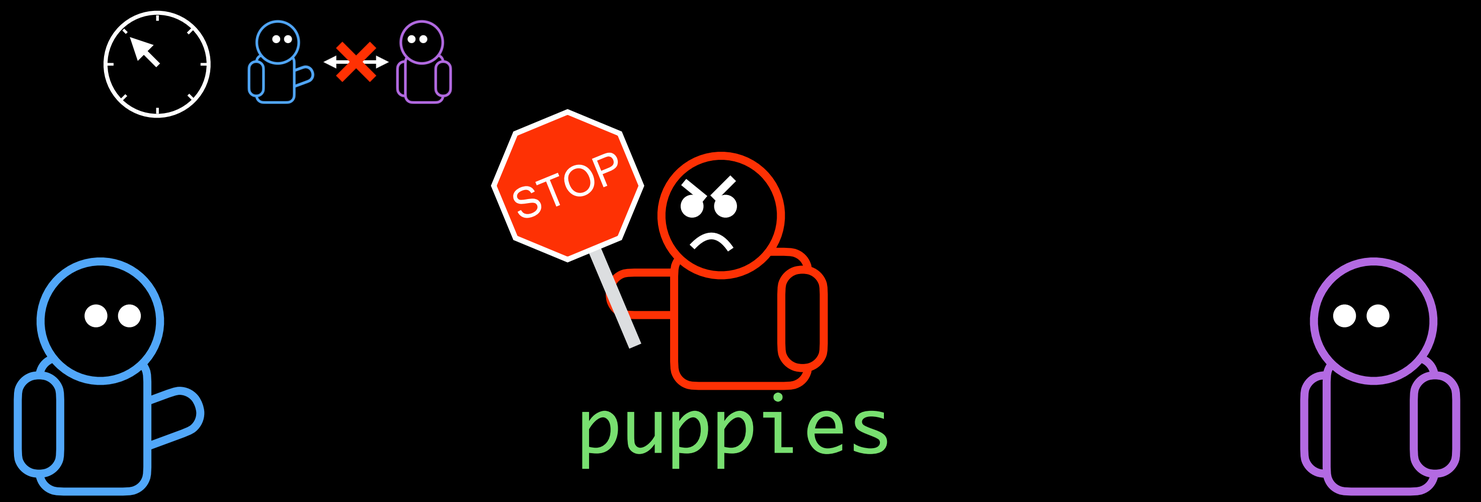

Weaponizing residual censorship

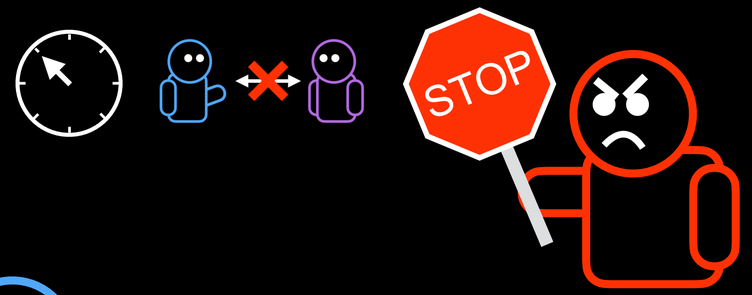

puppies

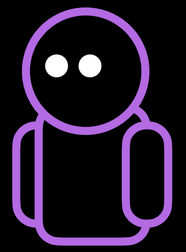

Weaponizing residual censorship

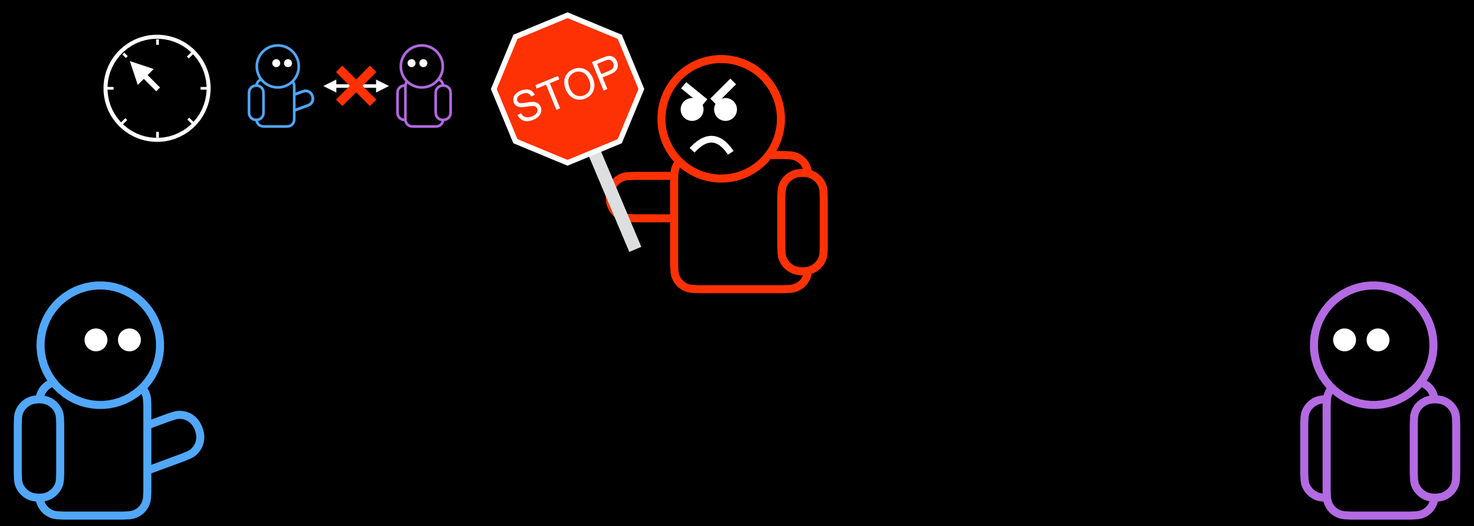

Weaponizing residual censorship

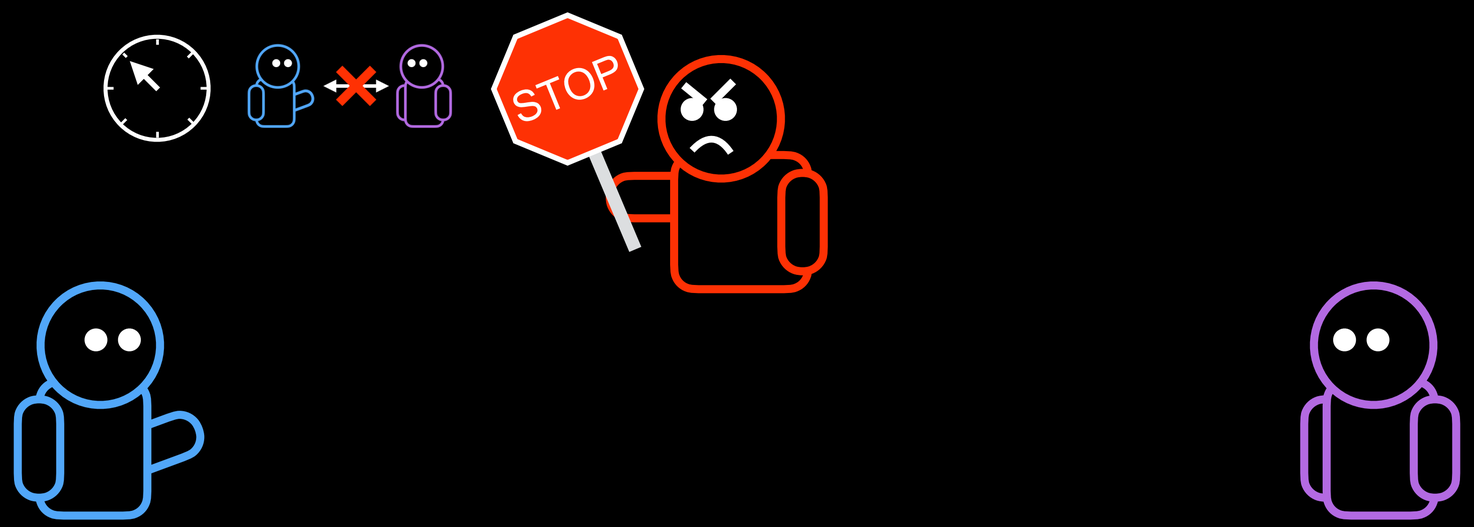

Attackers can restrict *benign* communication
from crossing the censors' borders

Weaponizing residual censorship

How long does it last?

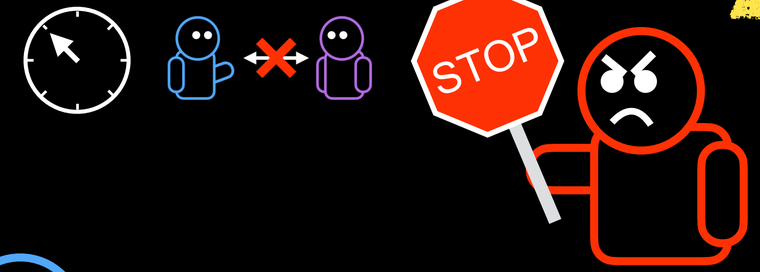

Limitations of residual censorship?

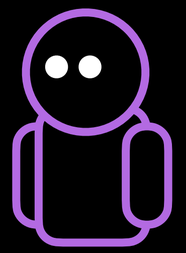

How can we trigger censors?

How can we evaluate ethically?

Weaponizing Residual Censorship

Current state of residual censorship?

- Experiments in Iran, China, and Kazahkstan
- Find differences in implementation and duration

How can it be weaponized?

- Ethical experiments with SP^3
- Attacked ourselves from a dozen vantage points

Weaponizing Residual Censorship

Current state of residual censorship?

- Experiments in Iran, China, and Kazahkstan
- Find differences in implementation and duration

How can it be weaponized?

- Ethical experiments with SP^3
- Attacked ourselves from a dozen vantage points

State of residual censorship

Diversity of censors

Diversity of protocols

		HTTP	SNI	ESNI	DNS	SMTP	Other
Injects TCP RSTs	 China	👮	👮*	👮			
Null routing	 Iran	👮	👮				👮
Null routing	 Kazakhstan	👮	👮				

Types of residual censorship

3-tuple or 4-tuple

	HTTP	SNI	ESNI	DNS	SMTP	Other
 China	3-tuple	3-tuple	both	✗	✗	✗
 Iran	4-tuple	4-tuple	✗	✗	✗	4-tuple
 Kazakhstan	4-tuple	4-tuple	✗	✗	✗	✗

Duration of residual censorship

How long does blocking last?

	HTTP	SNI	ESNI	DNS	SMTP	Other
 China	90s	60s	120s	✗	✗	✗
 Iran	180s	180s*	✗	✗	✗	60s
 Kazakhstan	120s	120s	✗	✗	✗	✗

Is residual censorship bidirectional?

Does it affect traffic entering the country?

	HTTP	SNI	ESNI	DNS	SMTP	Other
 China	✓	✓	✓*			
 Iran	✓	✓				✗
 Kazakhstan	✓	✓				

In all cases, censor tracks traffic direction

State of residual censorship

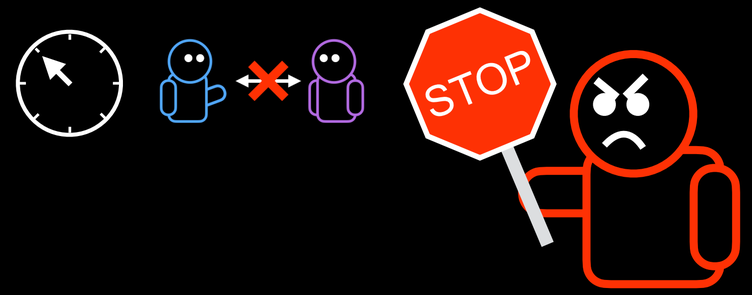

Residual censorship is implemented differently around the world

- Different censorship mechanisms (RSTs vs Null Routing)
- Different types of censorship, even within countries
- Bi-directional, but direction matters

Weaponizing Residual Censorship

Current state of residual censorship?

- Experiments in Iran, India, China, and Kazahkstan
- Find differences in implementation and duration

How can it be weaponized?

- Ethical experiments with SP^3
- Attacked ourselves from a dozen vantage points

Weaponizing Residual Censorship

- Current state of residual censorship?
  - Experiments in Iran, India, China, and Kazahkstan
  - Find differences in implementation and duration

- How can it be weaponized?
  - Ethical experiments with SP^3
  - Attacked ourselves from a dozen vantage points

How can we trigger censors?

Issue a request for forbidden content

How can we trigger censors?

Issue a request for forbidden content

Weaponizing Middleboxes for
TCP Reflected Amplification

to appear in USENIX Security
later this summer

How can we trigger censors?

Issue a request for forbidden content

Weaponizing Middleboxes for TCP Reflected Amplification

to appear in USENIX Security
later this summer

Packet sequences

- SYN *with Request*
- PSH
- PSH+ACK
- SYN ; PSH
- SYN ; PSH+ACK

How can we trigger censors?

Issue a request for forbidden content

Weaponizing Middleboxes for TCP Reflected Amplification

to appear in USENIX Security later this summer

Packet sequences

- SYN *with Request*
- PSH
- PSH+ACK
- SYN ; PSH
- SYN ; PSH+ACK

Censorship can be triggered without a proper 3-way handshake

How can we trigger censors?

Issue a request for forbidden content

Weaponizing Middleboxes for TCP Reflected Amplification

to appear in USENIX Security later this summer

Packet sequences

- SYN *with Request*
- PSH
- PSH+ACK
- SYN ; PSH
- SYN ; PSH+ACK

Censorship can be triggered without a proper 3-way handshake

Ethical evaluation

Attack ourselves *ethically* without affecting other hosts

- Need to spoof traffic to or from a censored regime
- Only between hosts we control
- Full control over packets we send

Solution: SP^3

Experiment Setup with SP^3

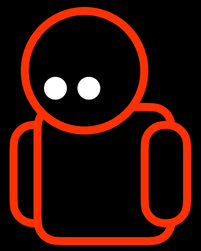

Censor

Client

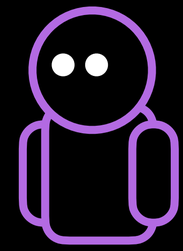

Server

SP^3

Experiment Setup with SP^3

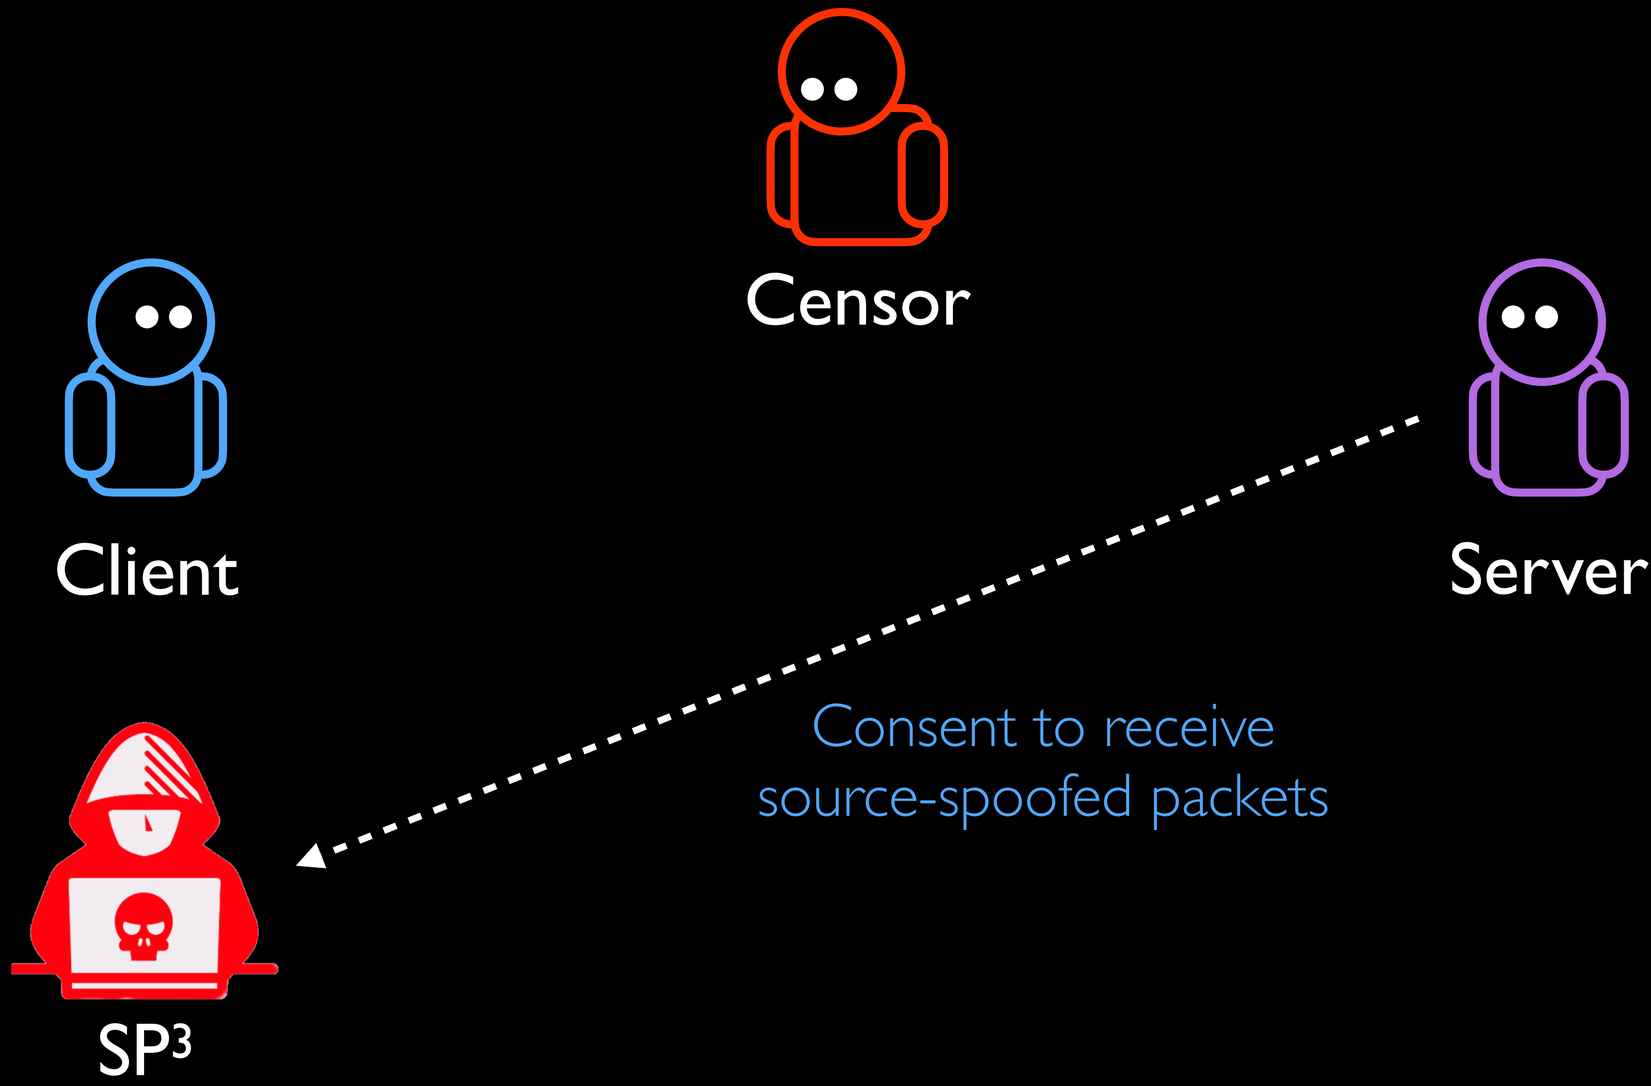

Experiment Setup with SP^3

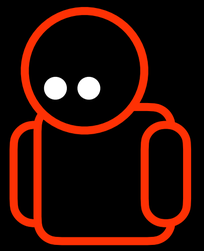

Censor

Client

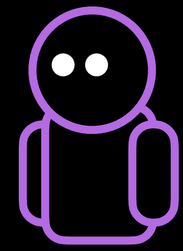

Server

SP^3

Experiment Setup with SP³

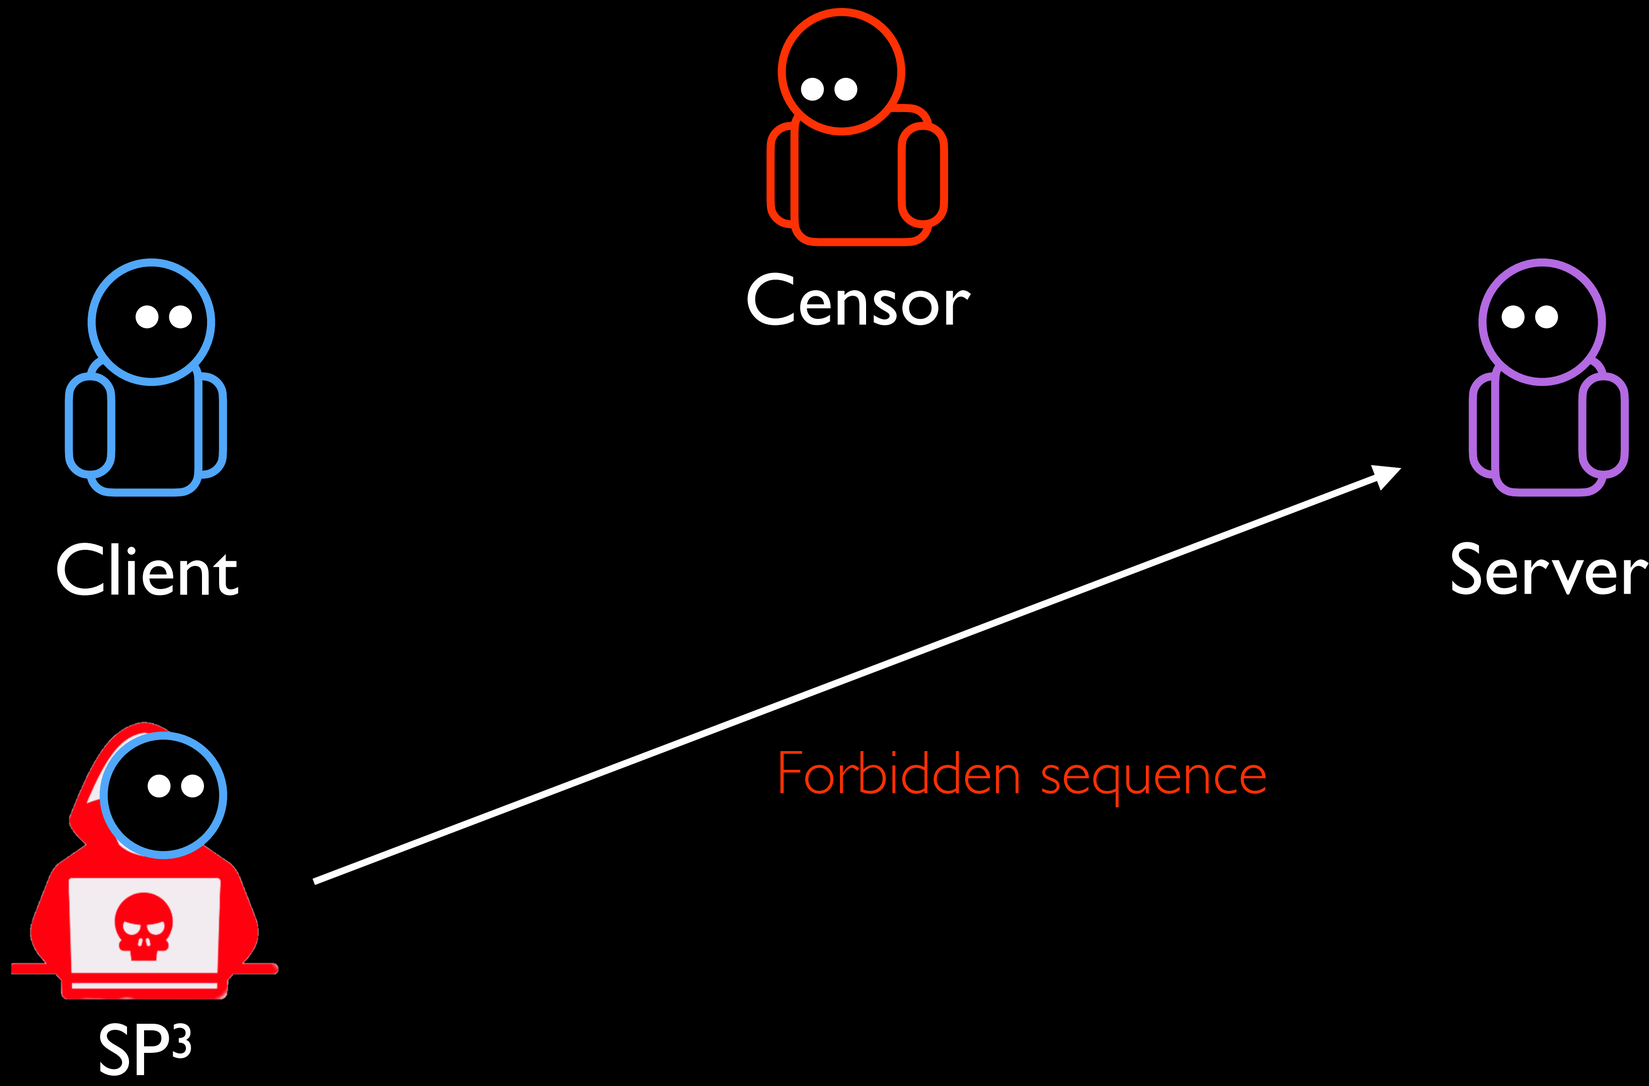

Experiment Setup with SP^3

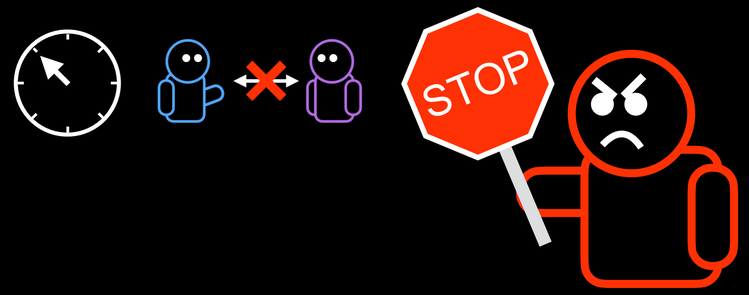

Censor

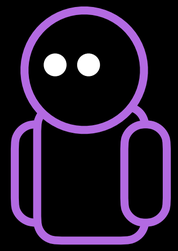

Experiment Setup with SP^3

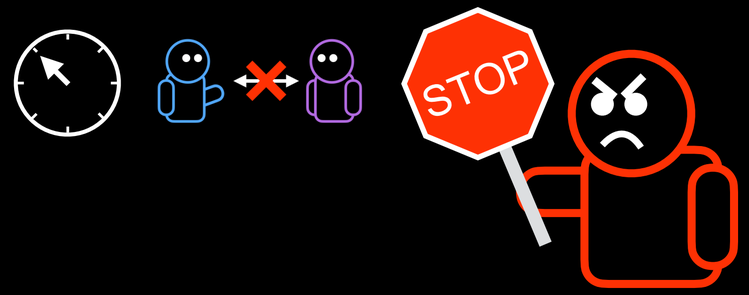

Censor

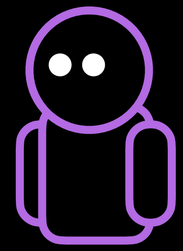

Innocuous request

Experiment Setup with SP^3

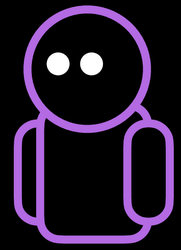

Innocuous request

Around the world

Censored Regimes

University of Washington

China

Kazakhstan

Iran

Experiment Setup with SP³

"Attacker"
(SP³)

Censoring
Nation-states

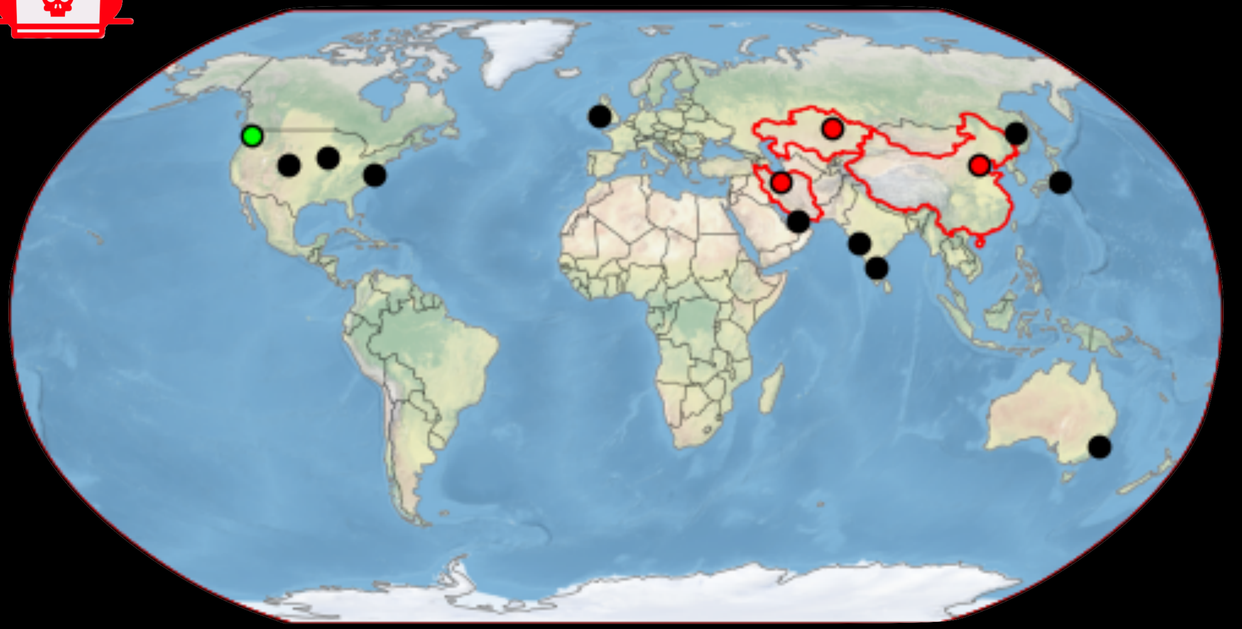

"Victims"
(Our clients)

"Victims"
(Our servers)

Tested from 16 external vantage points

Experiment Setup with SP³

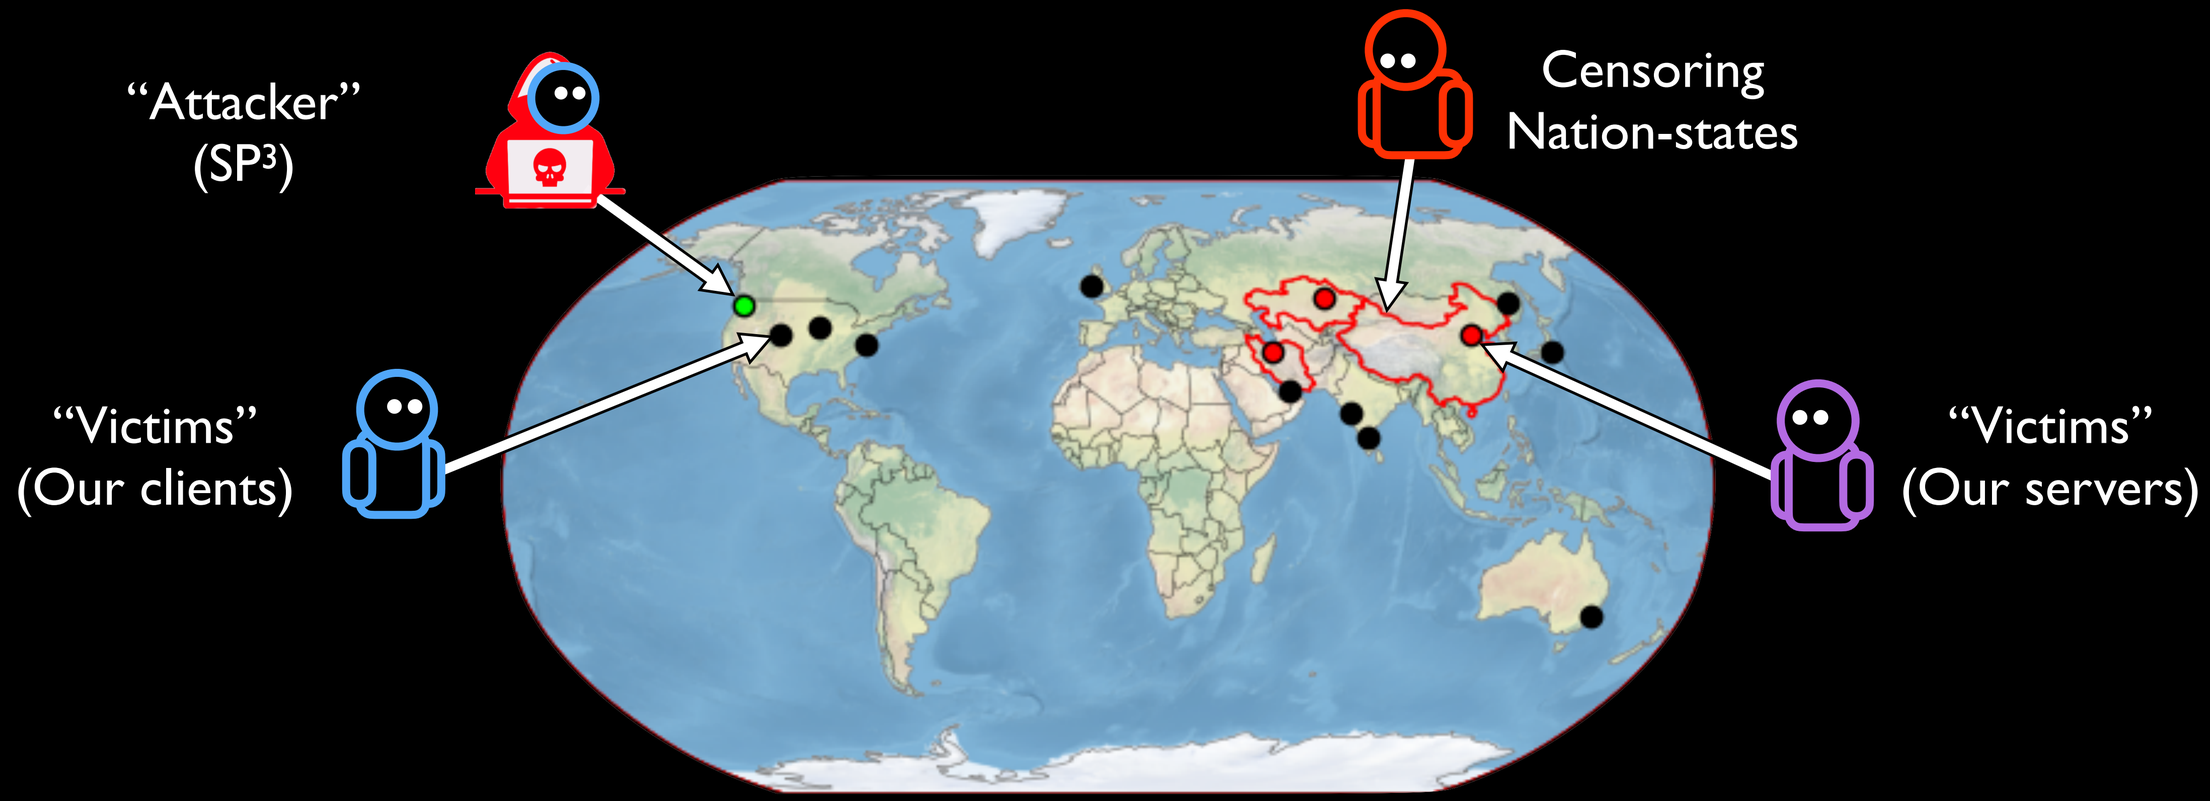

Tested from 16 external vantage points

Results

Victim Location		Destination Location							
		Kazakhstan		Iran		Beijing 1		Beijing 2	
		HTTP	HTTPS	HTTP	HTTPS	HTTP	ESNI	HTTP	ESNI
Australia	Sydney								
China	Beijing 1 Beijing 2								
India	Mumbai Bangalore 1 Bangalore 2								
Iran	Tehran								
Ireland	Dublin 1 Dublin 2								
Japan	Tokyo								
Kazakhstan	Qaraghandy								
Russia	Khabarovsk								
UAE	Dubai 1 Dubai 2								
USA	Colorado Iowa Virginia								

Results

Victim Location		Destination Location: Kazakhstan		Iran		Beijing 1		Beijing 2	
		HTTP	HTTPS	HTTP	HTTPS	HTTP	ESNI	HTTP	ESNI
Australia	Sydney	✓	✓	✓	✓	50%	10%	55%	✓
China	Beijing 1	✗	✓	✓	✓	N/A	N/A	N/A	N/A
	Beijing 2	✗	✓	✓	✓	N/A	N/A	N/A	N/A
India	Mumbai	✗	✓	✓	✓	✗	✗	✗	30%
	Bangalore 1	✓	✓	✓	✓	50%	10%	✓	✓
	Bangalore 2	✓	✓	✓	✓	25%	10%	✓	✓
Iran	Tehran	✓	✓	N/A	N/A	✗	50%	75%	✓
Ireland	Dublin 1	✗	✓	✓	✓	✗	✗	✗	5%
	Dublin 2	✗	✓	✓	✓	50%	✗	✗	✗
Japan	Tokyo	✓	✓	✓	✓	25%	✗	✗	✓
Kazakhstan	Qaraghandy	N/A	N/A	✓	✓	50%	✗	20%	✗
Russia	Khabarovsk	✓	✓	✓	✓	✓	✗	✓	✗
UAE	Dubai 1	✗	✓	✓	✓	85%	✗	95%	✗
	Dubai 2	✗	✓	✓	✓	✗	10%	✗	50%
USA	Colorado	✓	✓	✓	✓	✗	✗	✓	✗
	Iowa	✗	✓	✓	✓	✗	✗	✗	60%
	Virginia	✓	✓	✓	✓	85%	✓	55%	✗

Why does it fail?

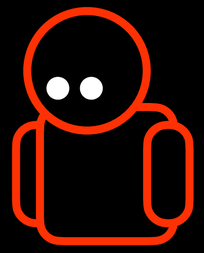

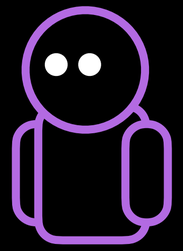

Source-spoofed traceroute from both to compare network path

Why does it fail?

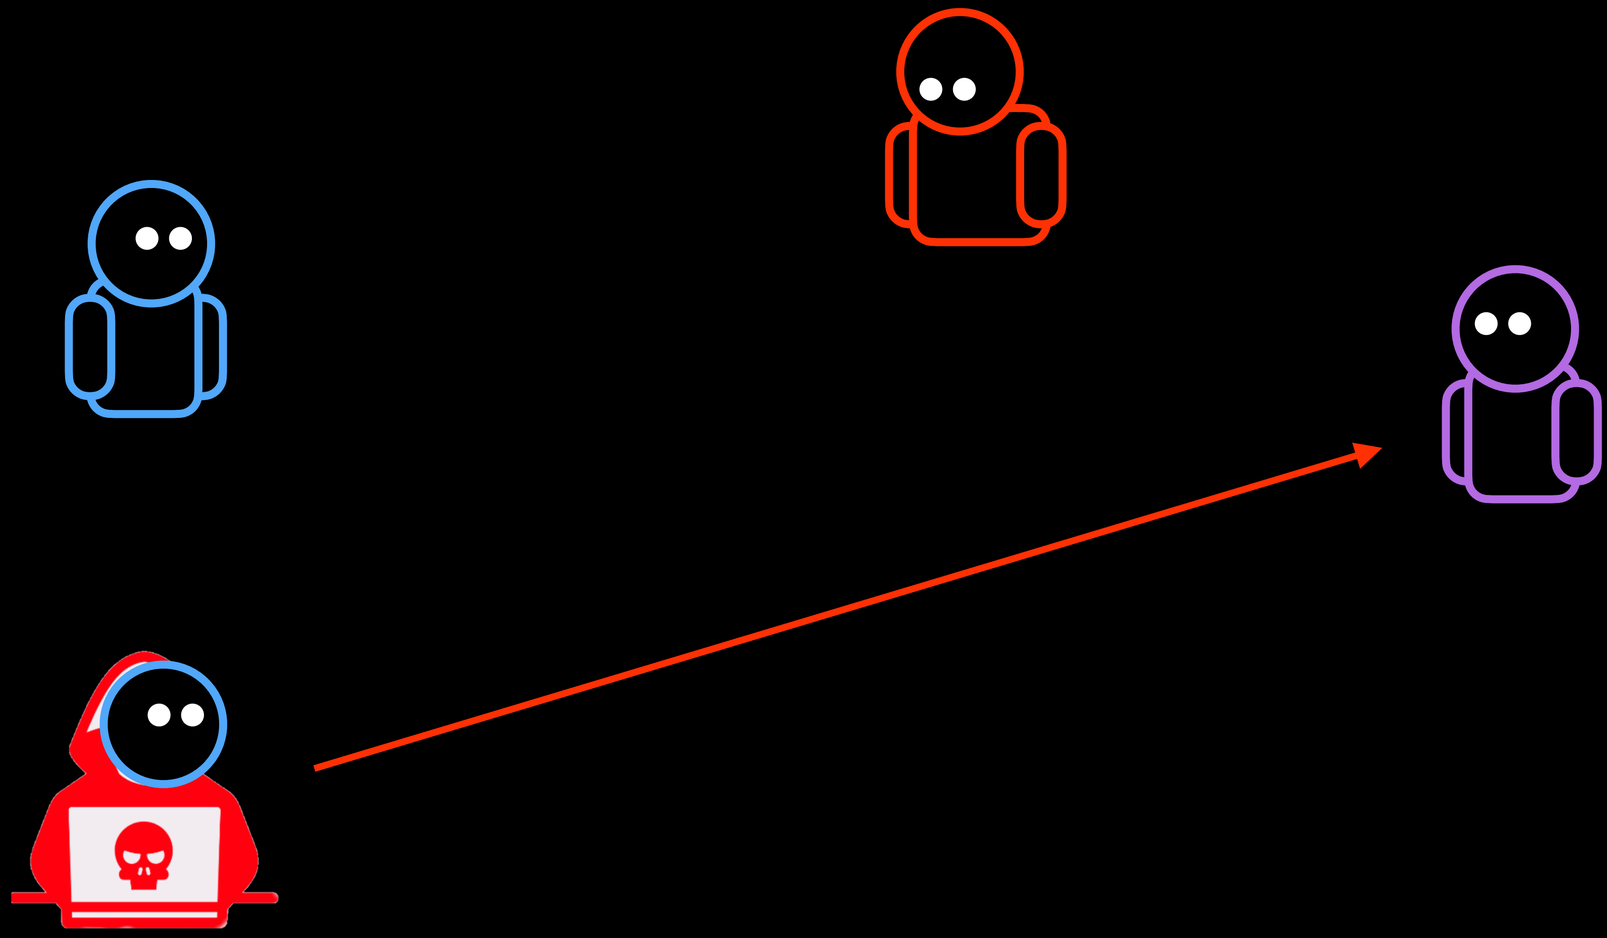

Source-spoofed traceroute from both to compare network path

Why does it fail?

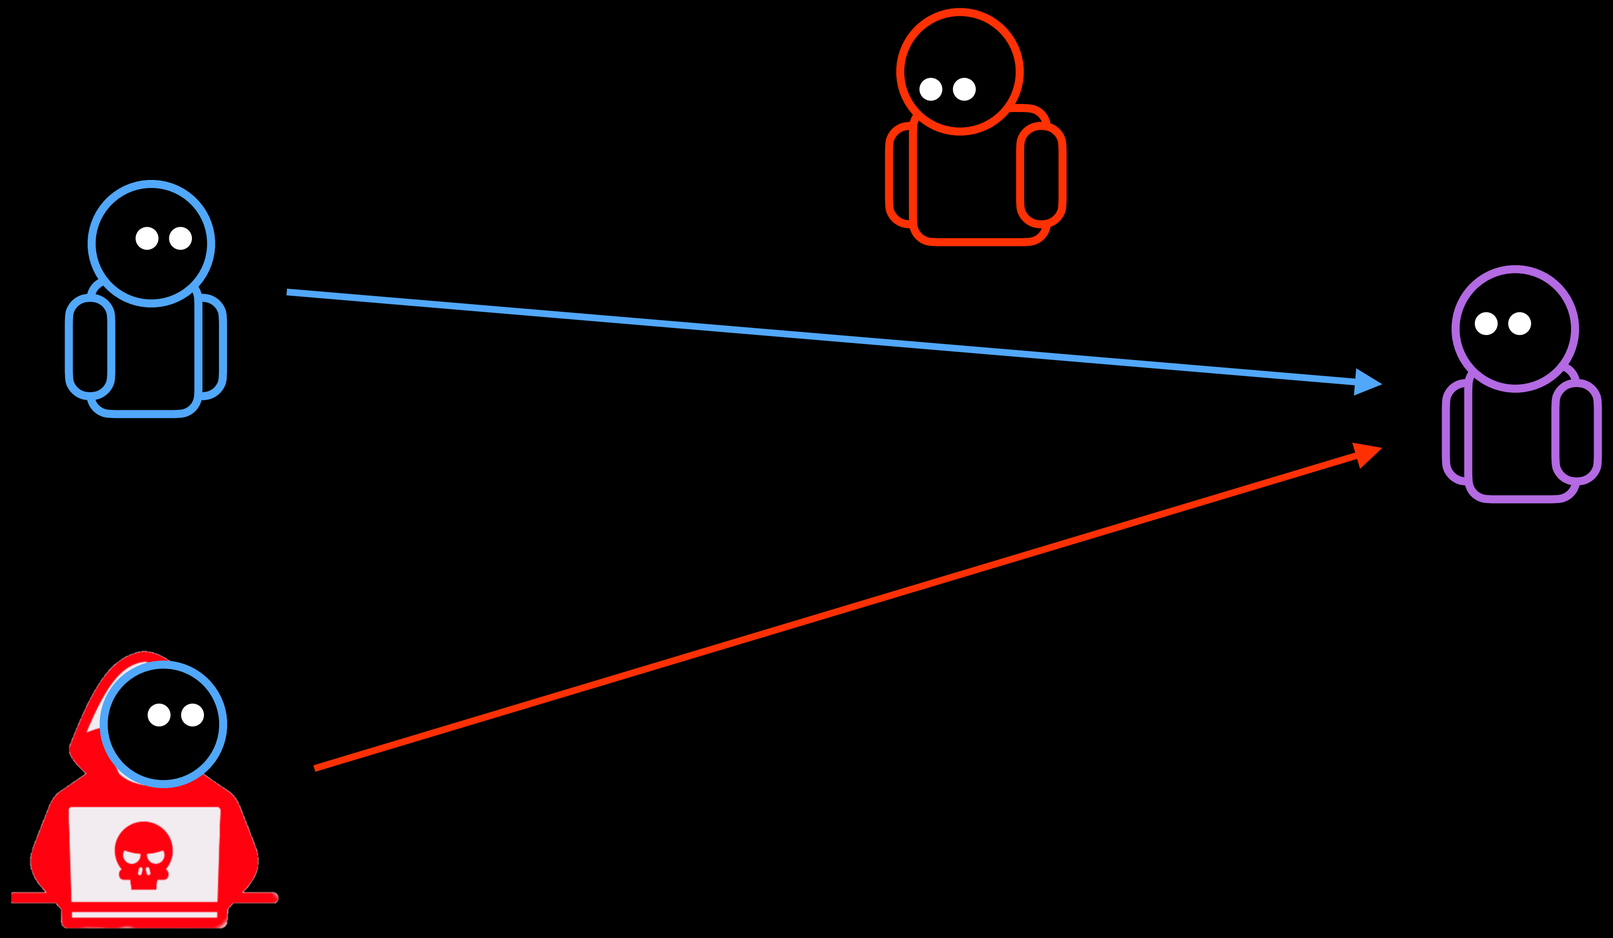

Source-spoofed traceroute from both to compare network path

Why does it fail?

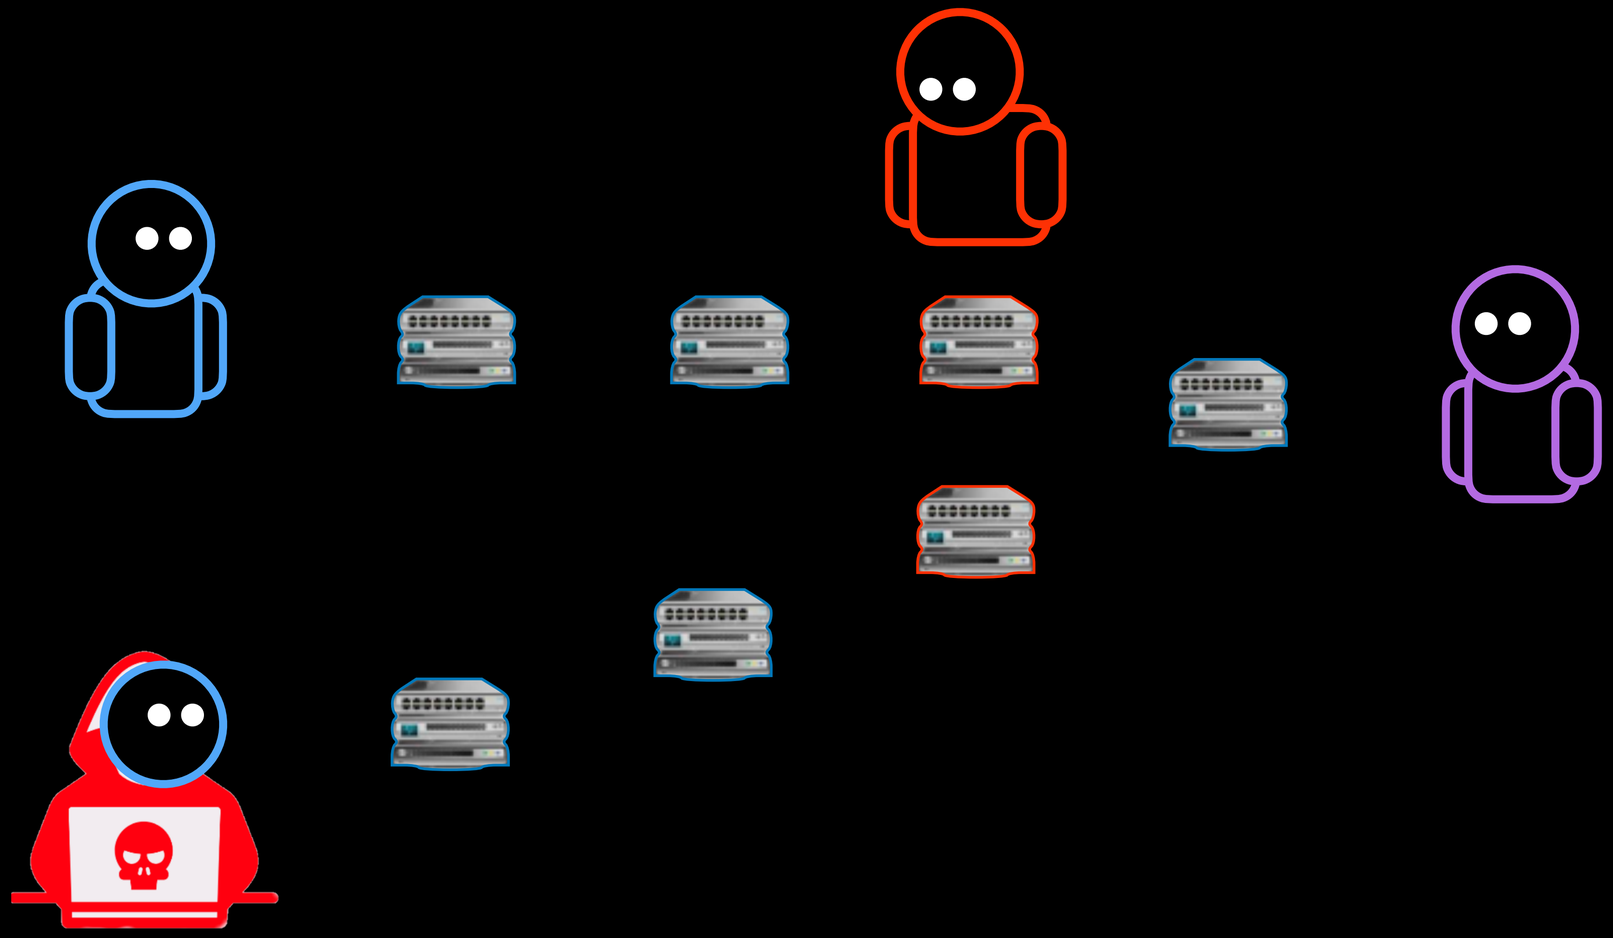

Why does it fail?

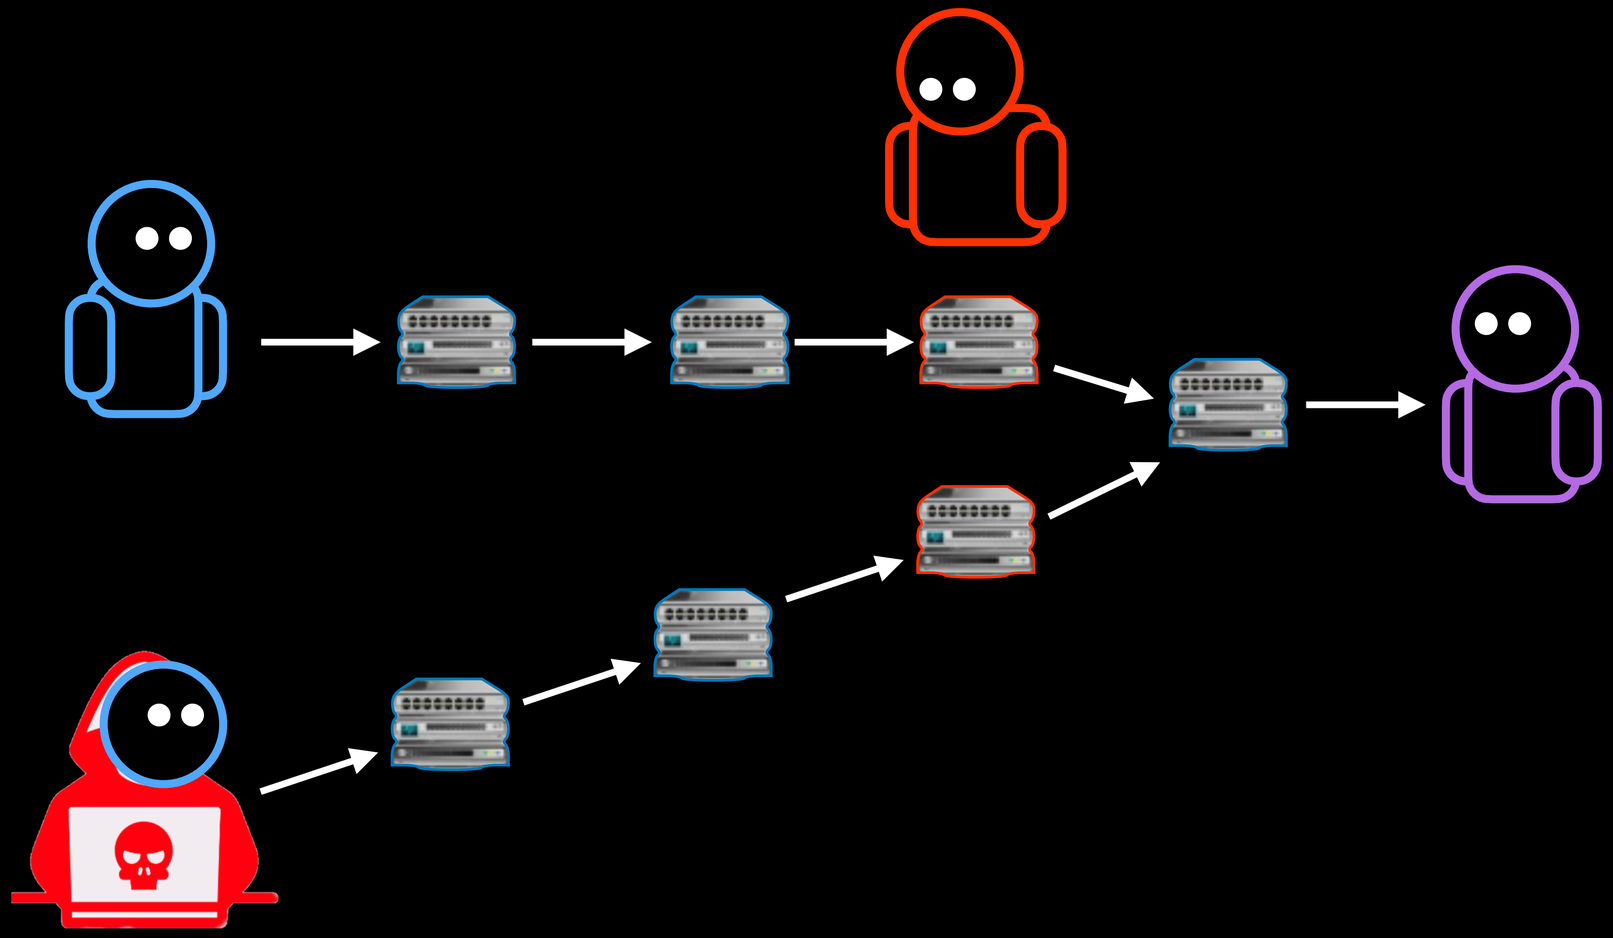

Why does it fail?

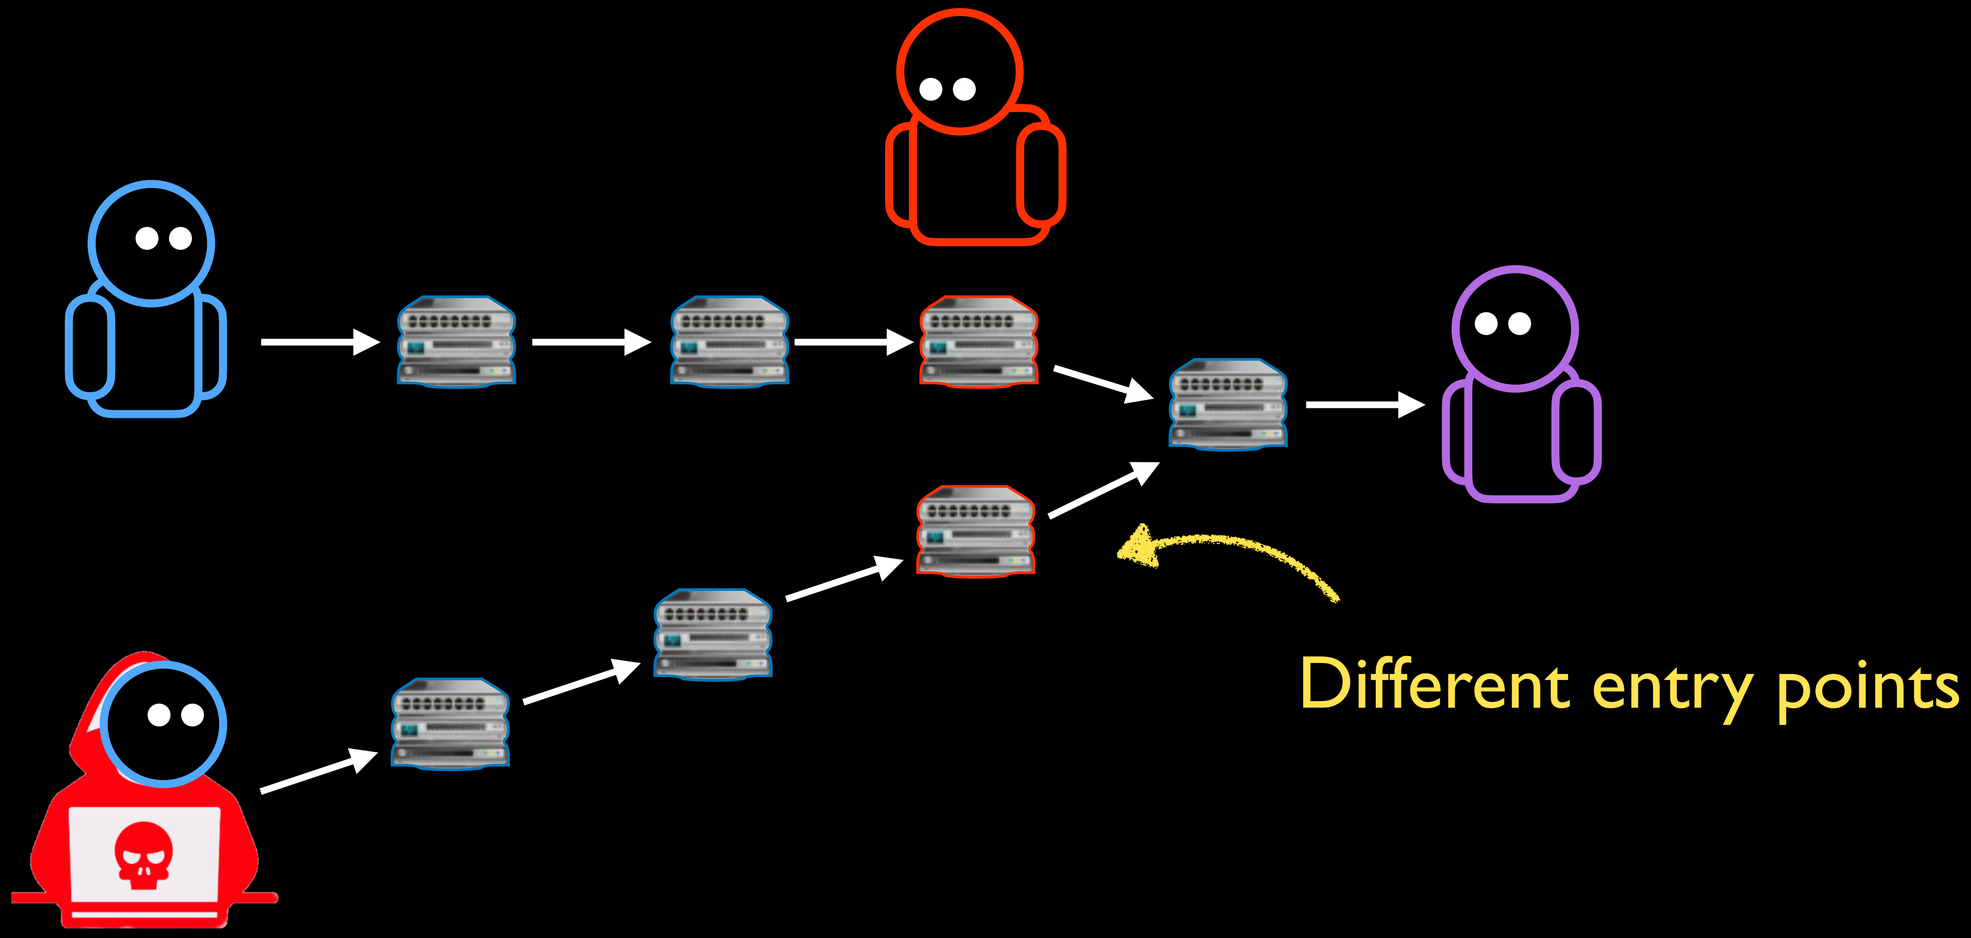

Why does it fail?

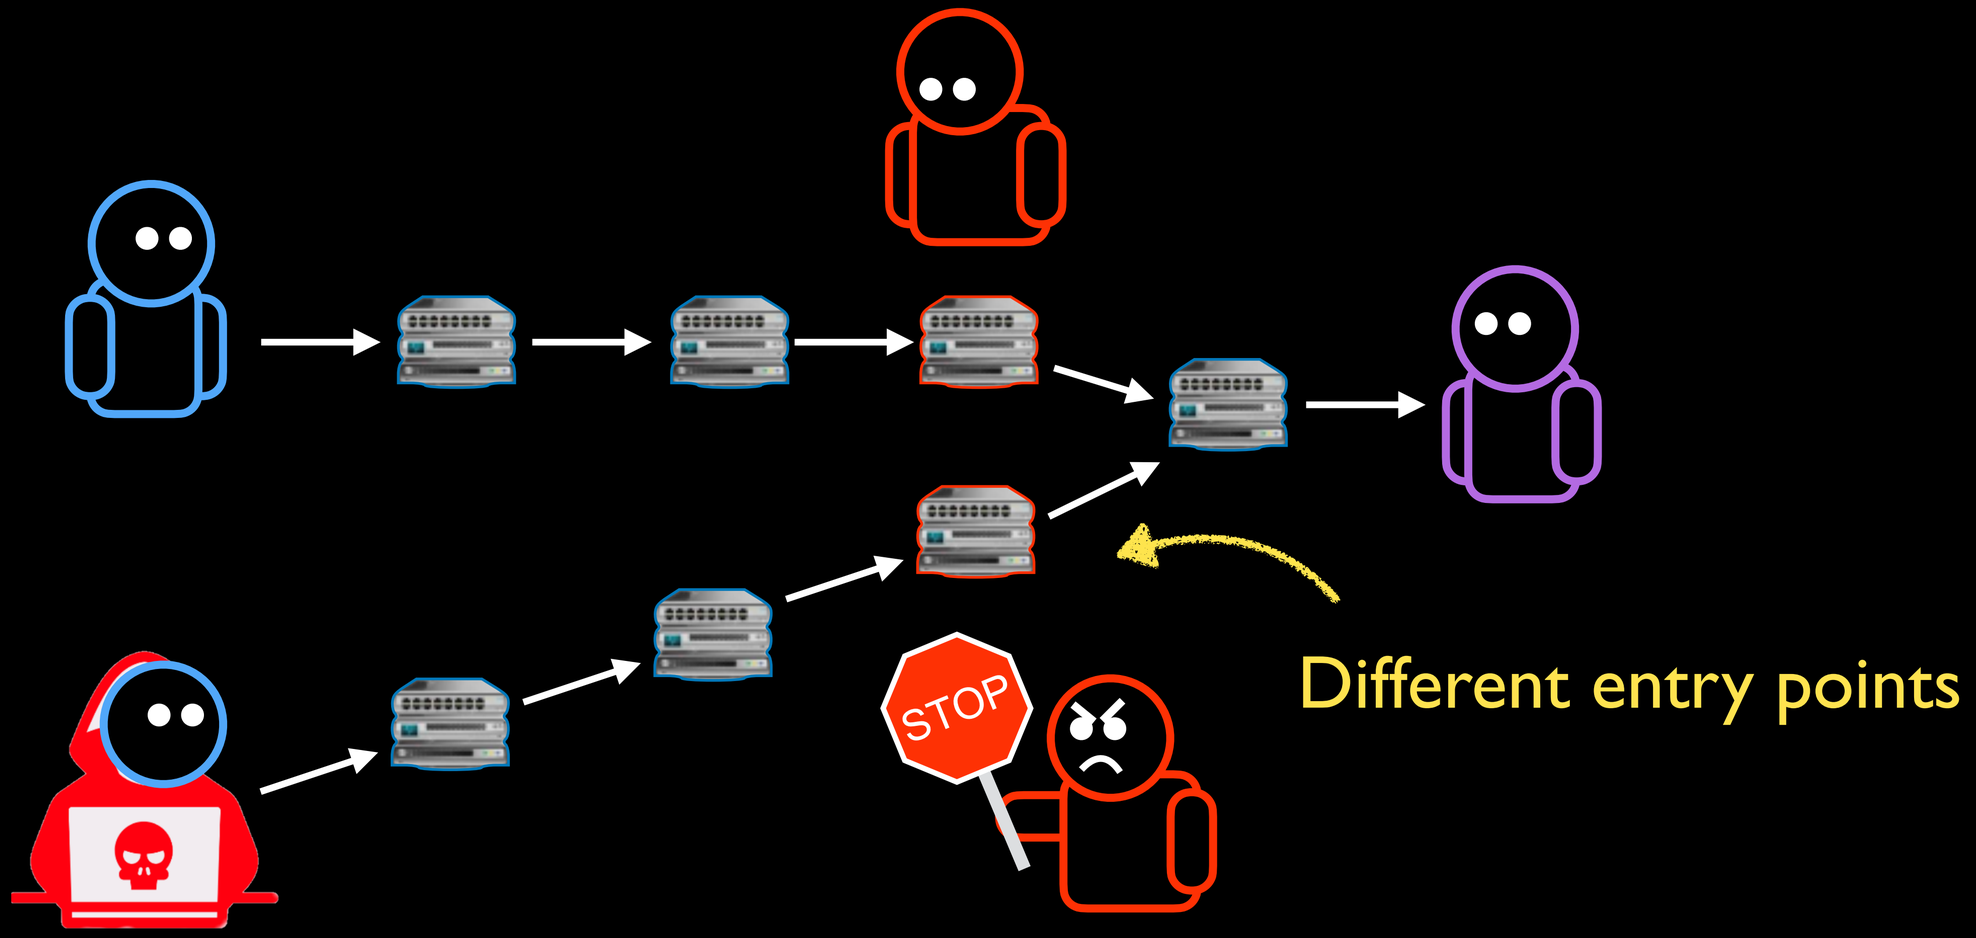

Sustaining the attack

Goal: block client IP to server IP:port

Sustaining the attack

Goal: block client IP to server IP:port

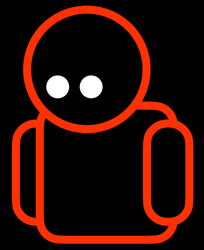

Censor

Client

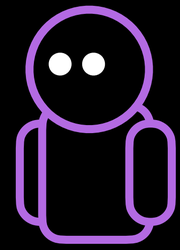

Server

Depends on type of residual censorship

Sustaining the attack

Goal: block client IP to server IP:port

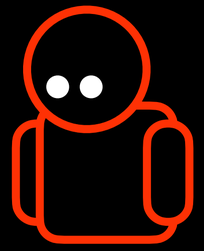

Censor

Client

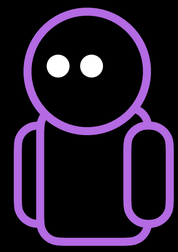

Server

Attacker can't guess source port

4-tuple (IP, port, IP, port)

3-tuple (IP, IP, port)

Sustaining the attack: 4-tuple

Goal: block client IP to server IP:port

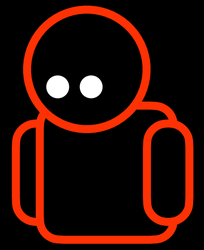

Censor

Client

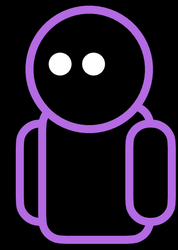

Server

Attacker can re-trigger from all 65,535 src ports

Sustaining the attack: 4-tuple

Goal: block client IP to server IP:port

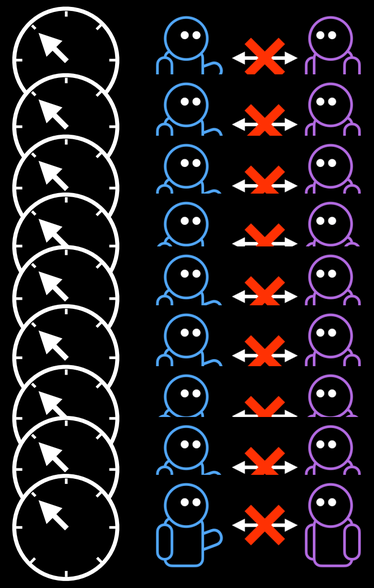

Censor

Client

Server

Attacker can re-trigger from all 65,535 src ports

Sustaining the attack: 4-tuple

Goal: block client IP to server IP:port

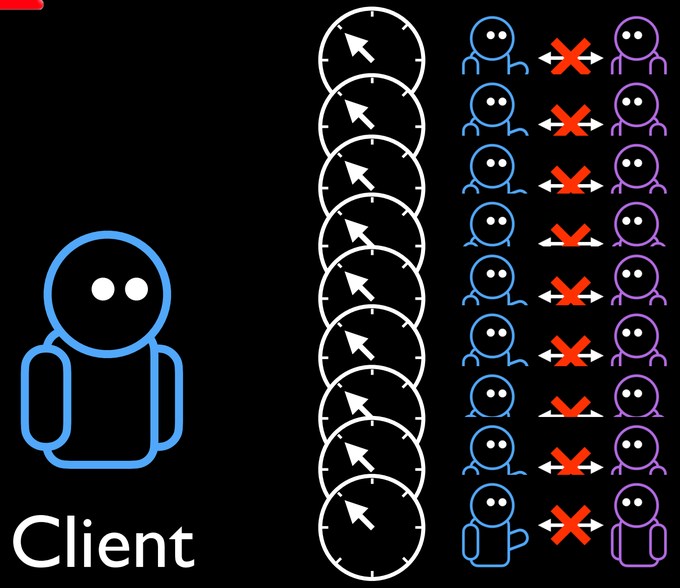

Client

Censor

Server

$$\text{Speed required} = \frac{|\text{Trigger packets}| \times 65{,}535}{\text{Duration}}$$

Sustaining the attack: 4-tuple

Goal: block client IP to server IP:port

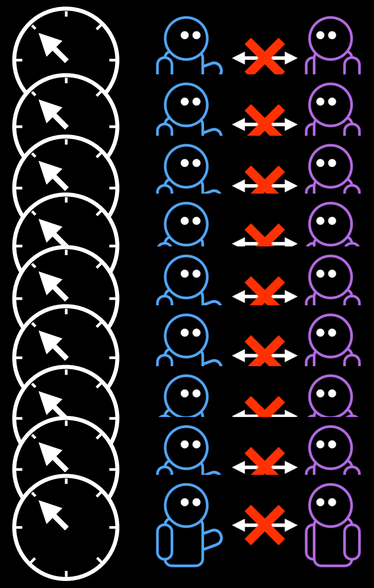

Censor

Client

Server

$$\text{Speed required} = \frac{|\text{Trigger packets}| \times 65{,}535}{\text{Duration}} = \frac{145 \text{ bytes} \times 65{,}535}{120 \text{ seconds}} = 634 \text{ kbps}$$

Kazakhstan

Sustaining the attack: 3-tuple

Goal: block client IP to server IP:port

$$\text{Speed required} = \frac{|\text{Trigger packets}|}{\text{Duration}} = \frac{145 \text{ bytes}}{90 \text{ seconds}} = 13 \text{ bps}$$

China

Weak attacker can launch this attack effectively

Sustaining the Attack

Victim helps sustain the attack

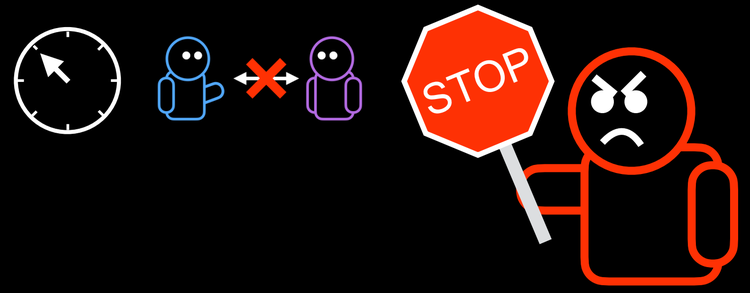

Censor

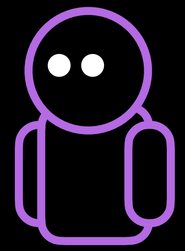

Sustaining the Attack

Victim helps sustain the attack

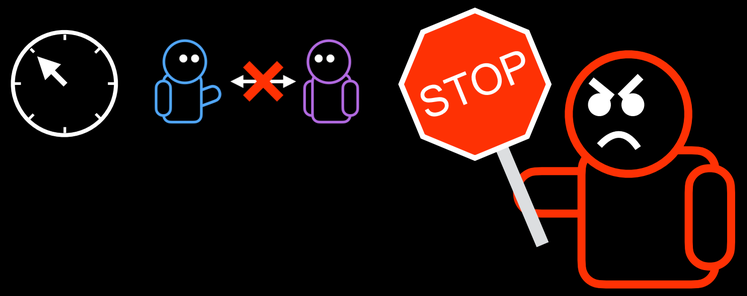

Censor

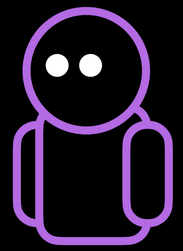

Innocuous request

Sustaining the Attack

Victim helps sustain the attack

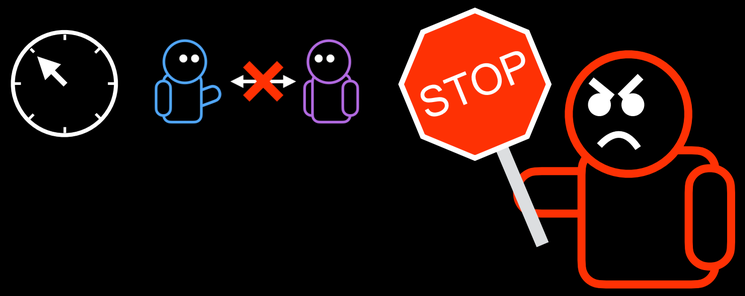

Censor

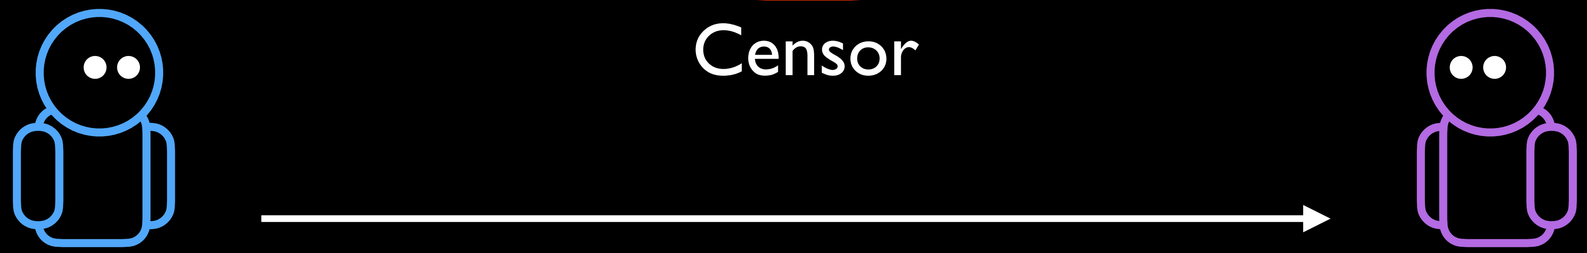

Innocuous request

Residual timer resets if the victim sends data

Victim retransmissions unknowingly sustain the attack on themselves

Can the server detect this?

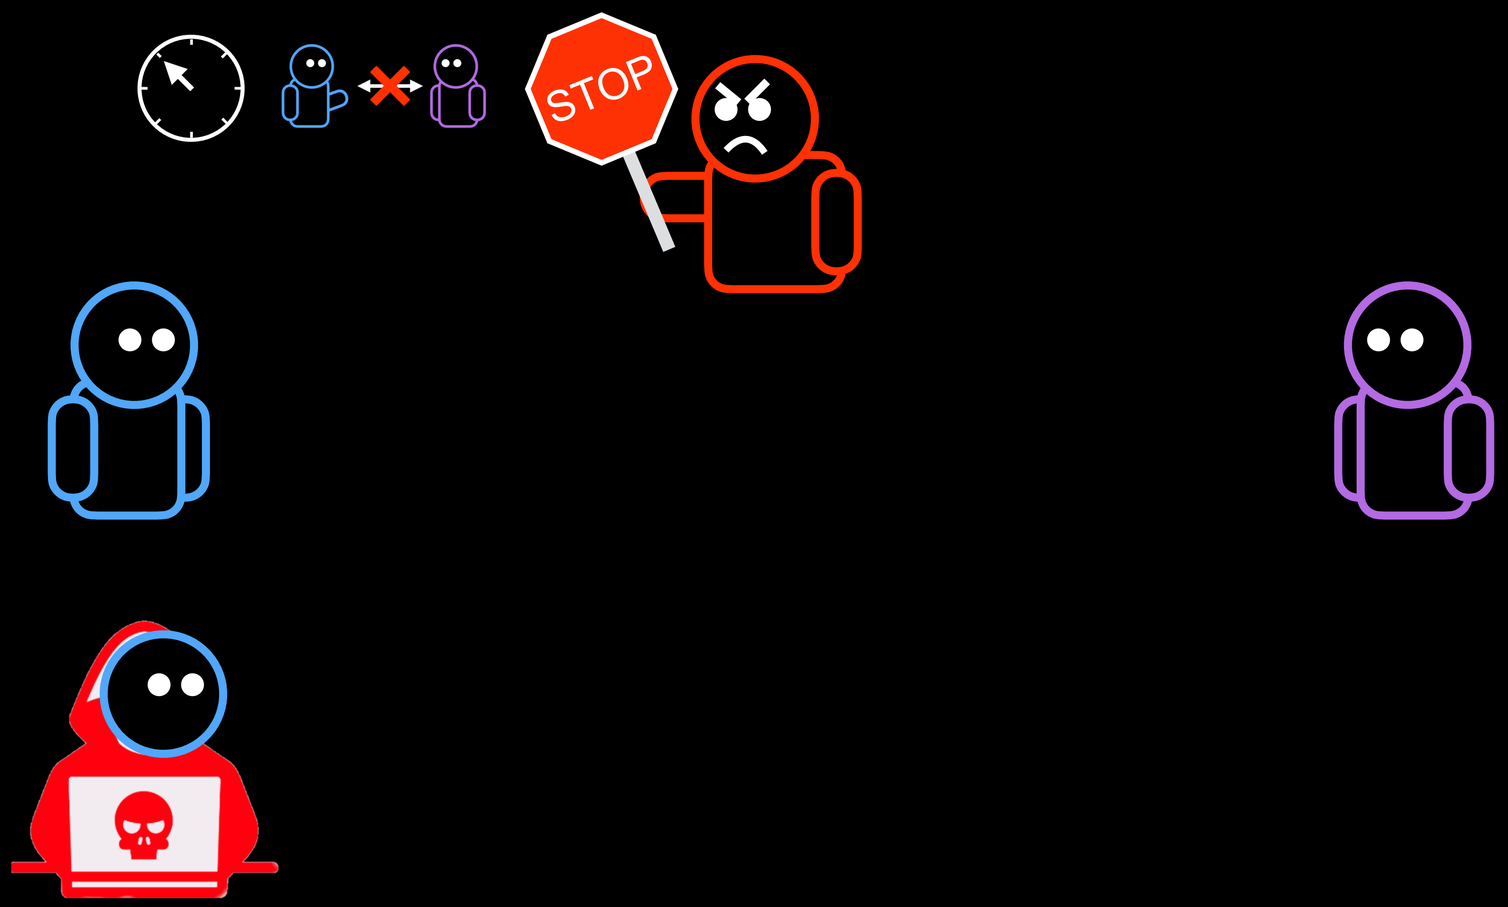

Attacker can limit TTL of packets to reach censor, but not server

Can the server detect this?

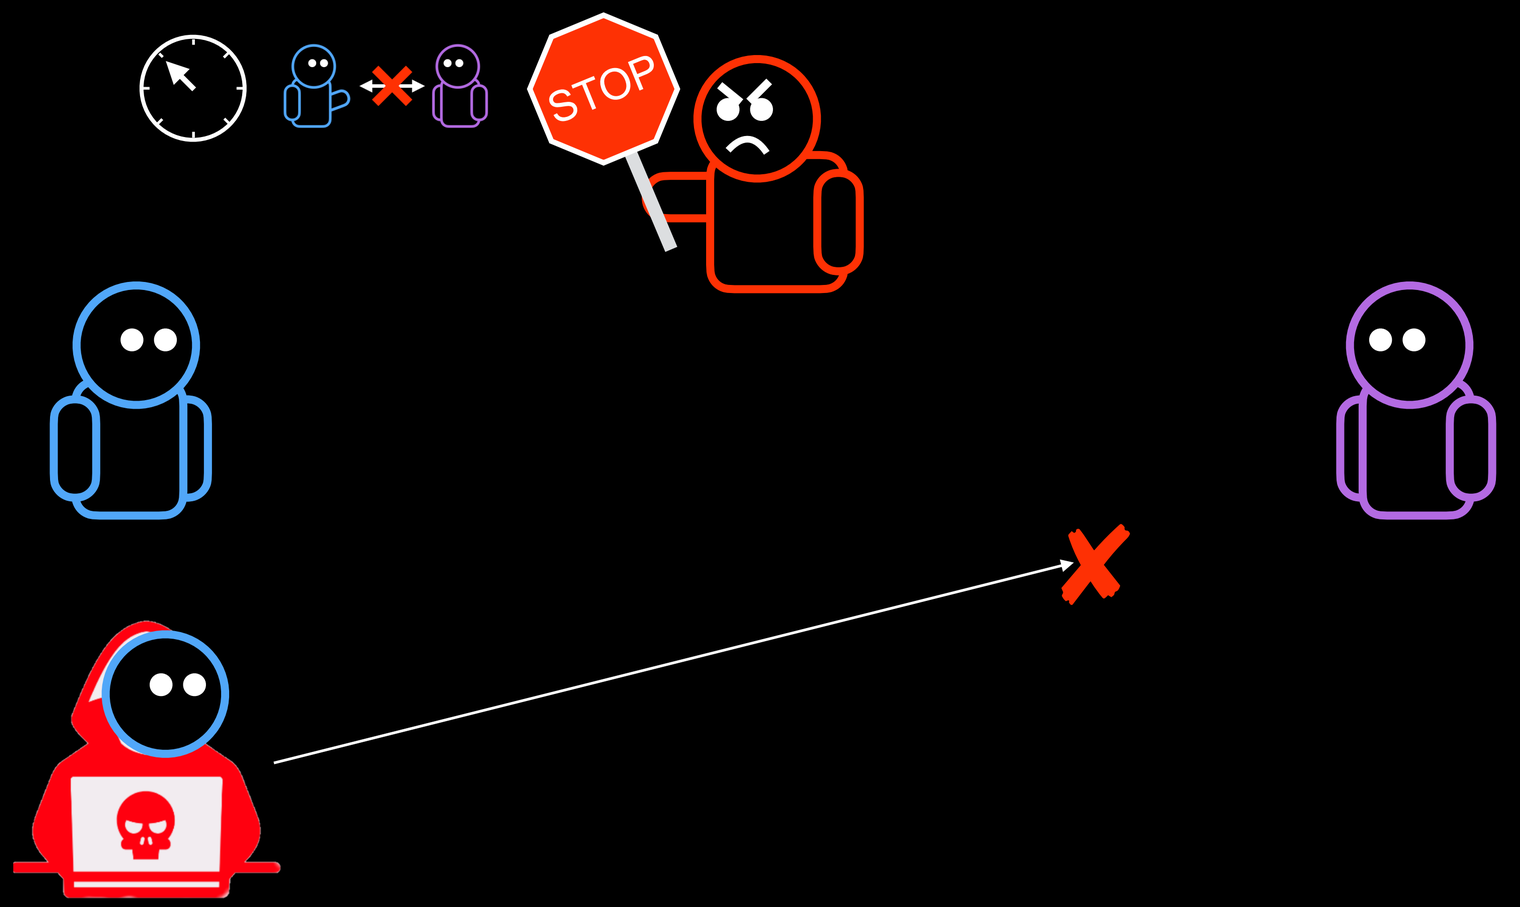

Attacker can limit TTL of packets to reach censor, but not server

Attack Limitations

Attacker must have a vantage point:

1. Without egress filtering
2. Shares a similar enough path with their victim
3. Traffic crosses a censor (with residual censorship)
4. Censor can be triggered statelessly

Surprisingly high number of shared network paths

What can be done?

Some mitigations available to censorship infrastructure:

- Abolish 3-tuple residual censorship
- Null routing should track sequence numbers
- Properly track 3-way handshake

Unfortunately, no good countermeasures available to victims

Other details in the paper

Port Experiments — Examine which ports are affected

Reliability Experiments — Studied the reliability of residual censorship

ESNI Weaponization — Details on how to weaponize ESNI in China

Attack Breadth — Analysis of other countries that might be affected

Weaponizing Middleboxes

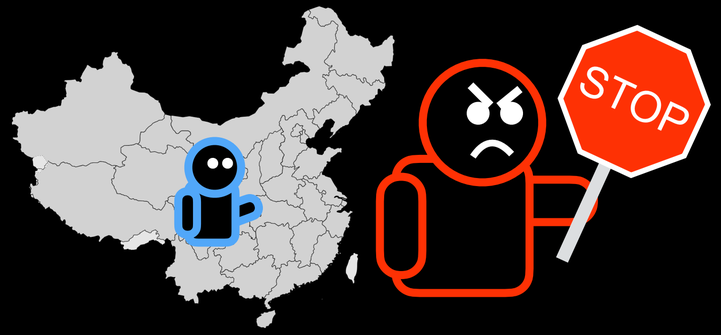

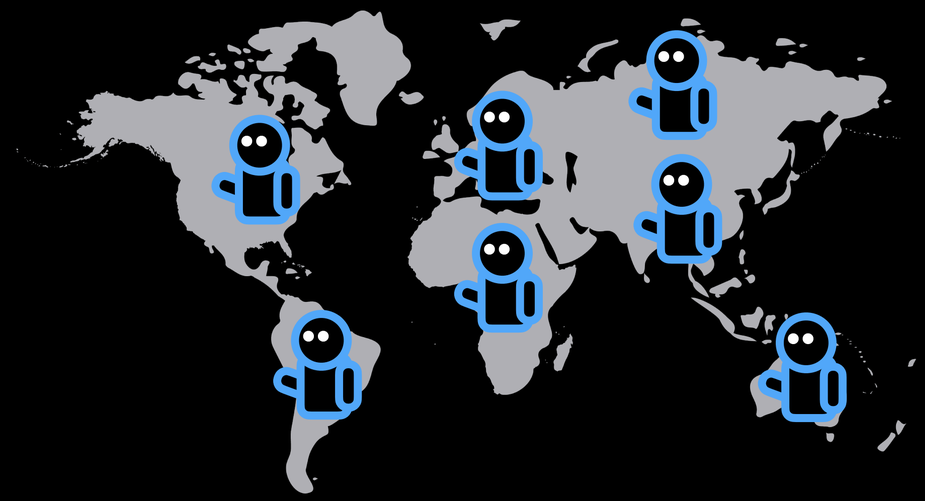

Censors can be weaponized to launch availability attacks

Can be done from a weak attacker

Censors pose a threat to the entire Internet

Code and website censorship.ai