Geneva:
Evolving Censorship Evasion Strategies

Kevin Bock George Hughey

Louis-Henri Merino, Tania Arya, Daniel Liscinsky, Regina Pogosian
Xiao Qiang, Dave Levin


UNIVERSITY OF MARYLAND

Berkeley SCHOOL OF INFORMATION

Why study censorship evasion?

We have Tor!

We have VPNs!

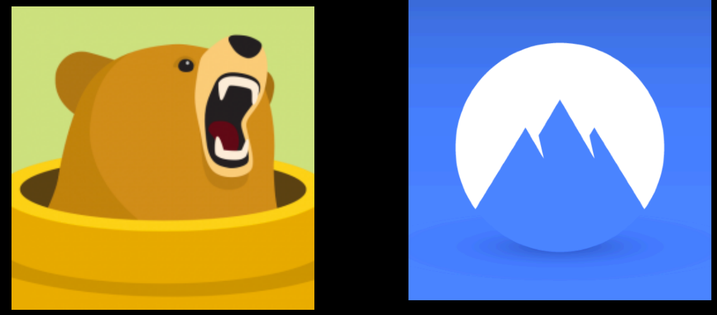

We have secure HTTPS!

Why study censorship evasion?

We have Tor

We have VPNs

We have secure HTTPS

Why study censorship evasion?

We have Tor

We have VPNs

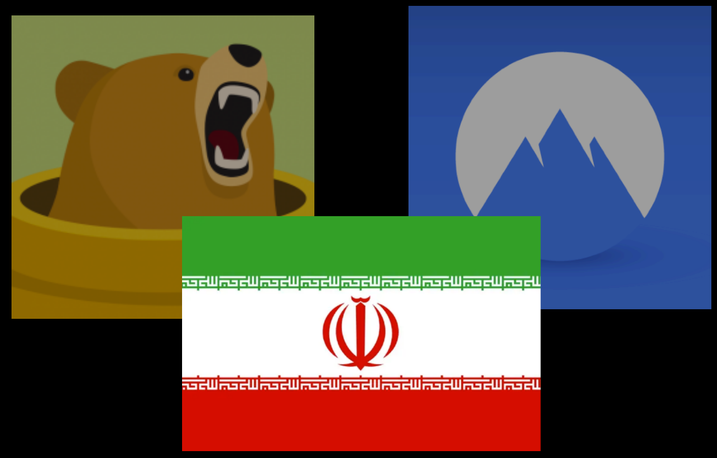

We have secure HTTPS

In-network censorship by nation-states

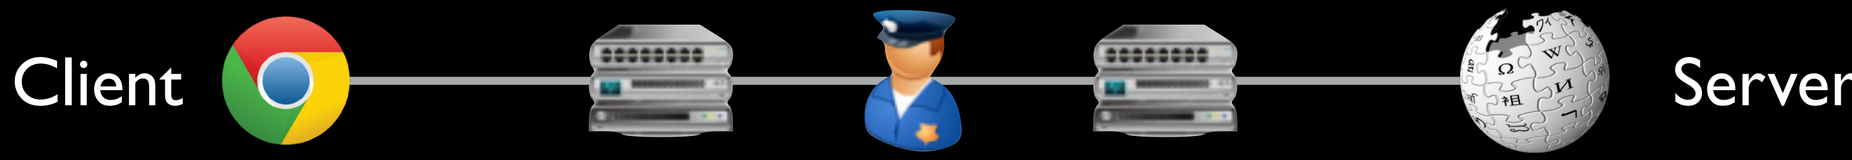

In-network censorship by nation-states

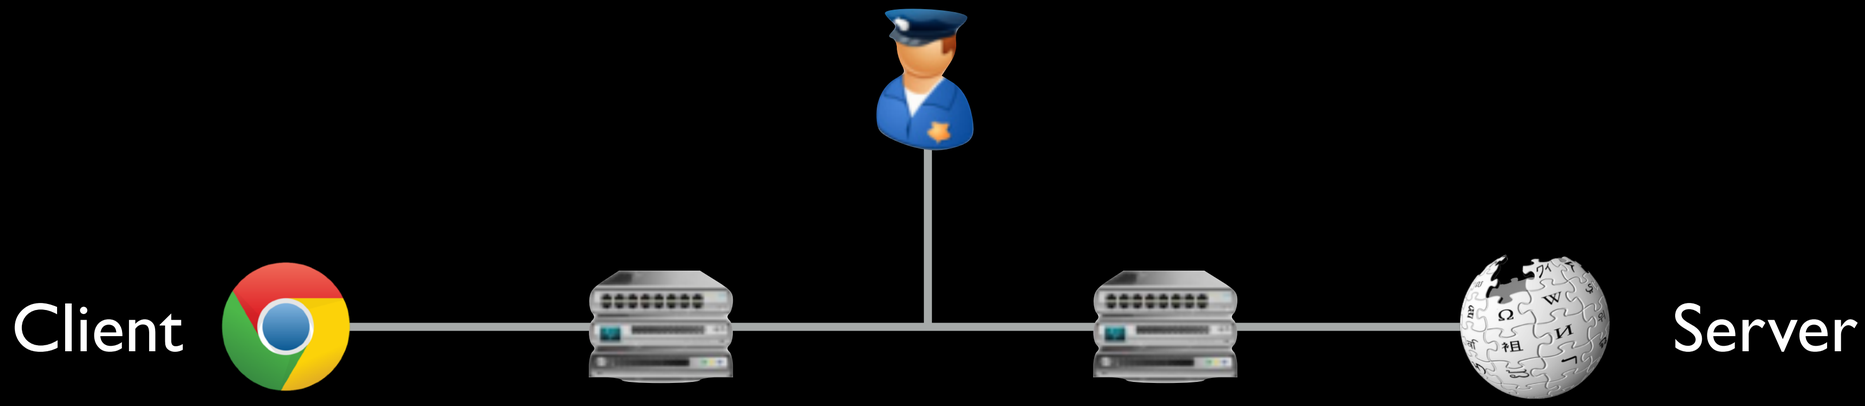

In-network censorship by nation-states

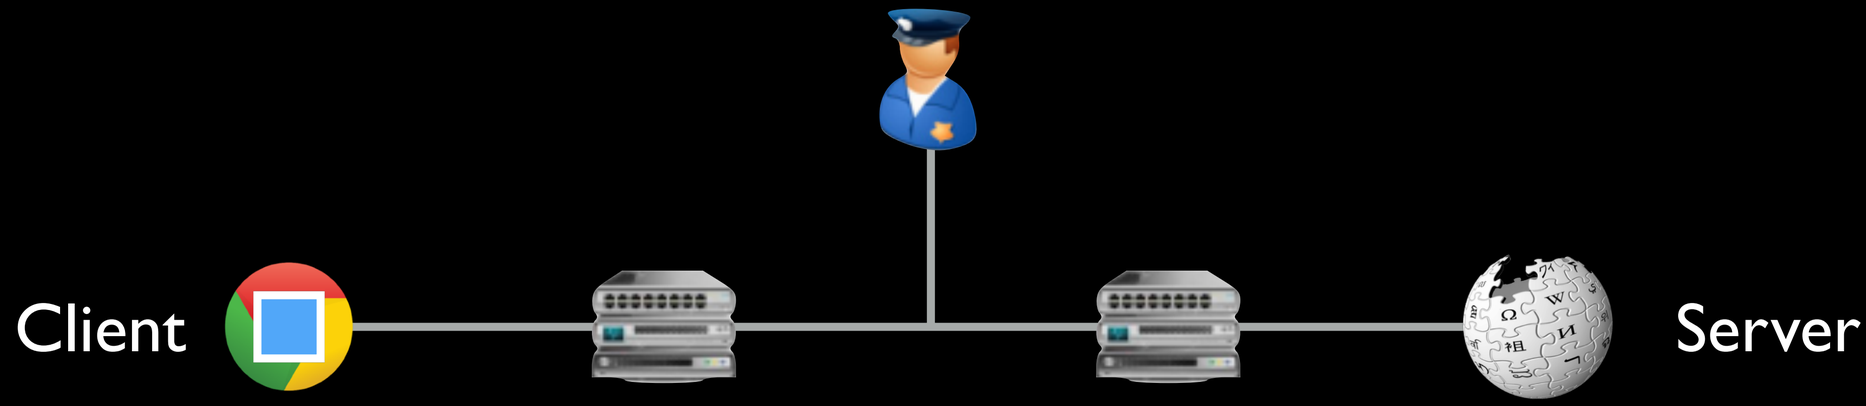

In-network censorship by nation-states

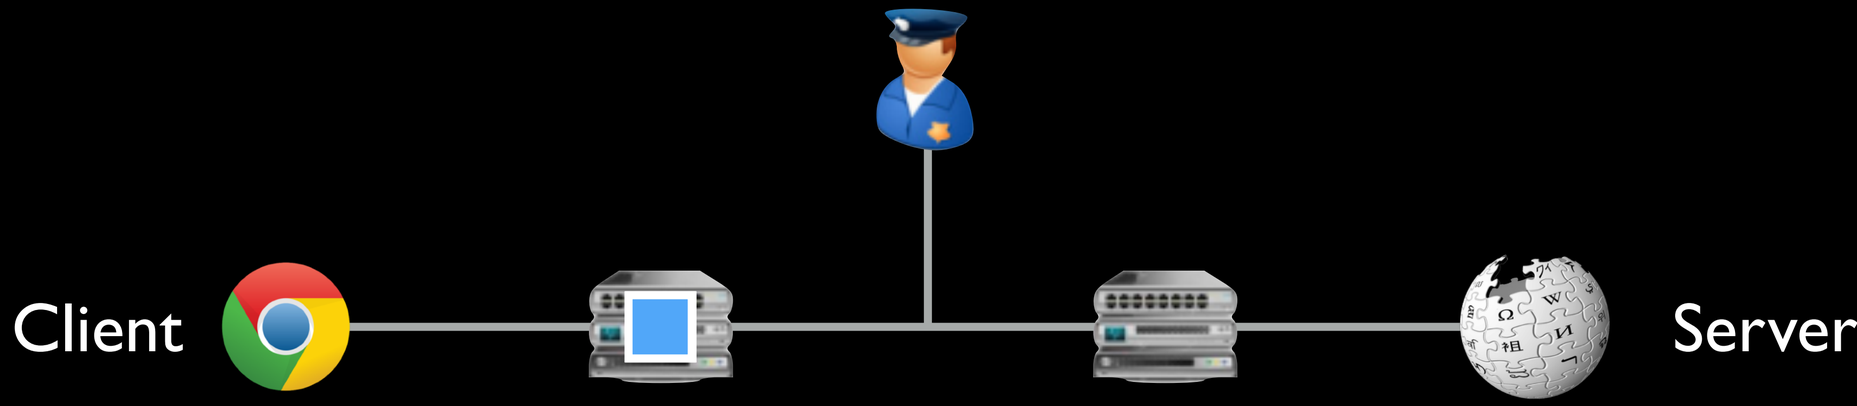

In-network censorship by nation-states

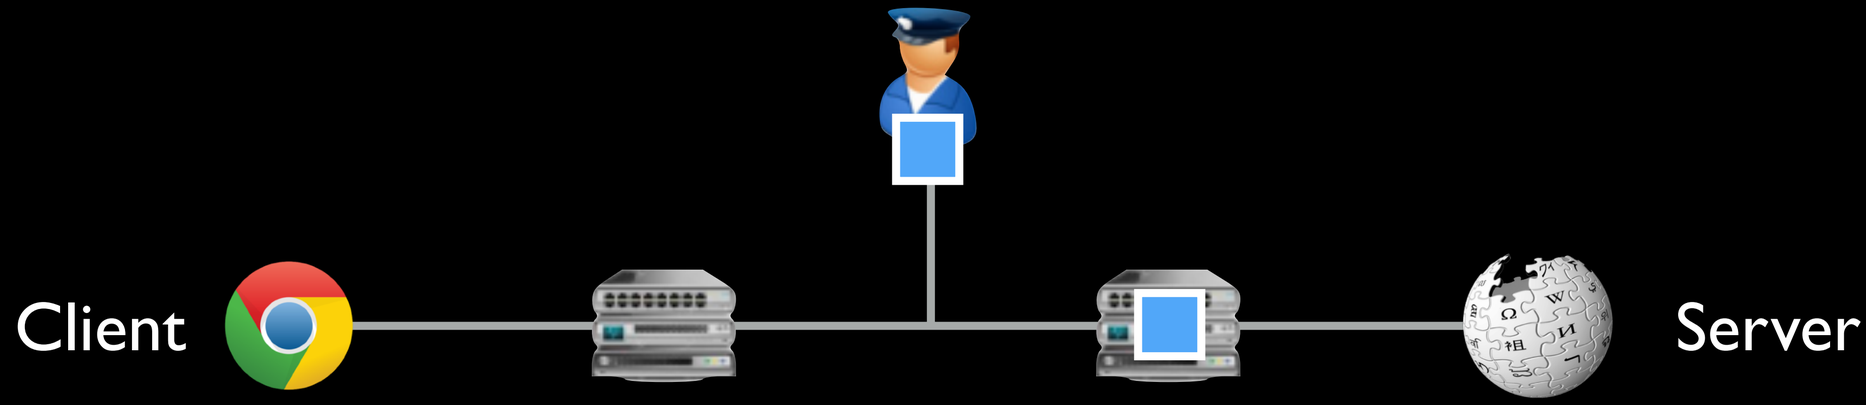

In-network censorship by nation-states

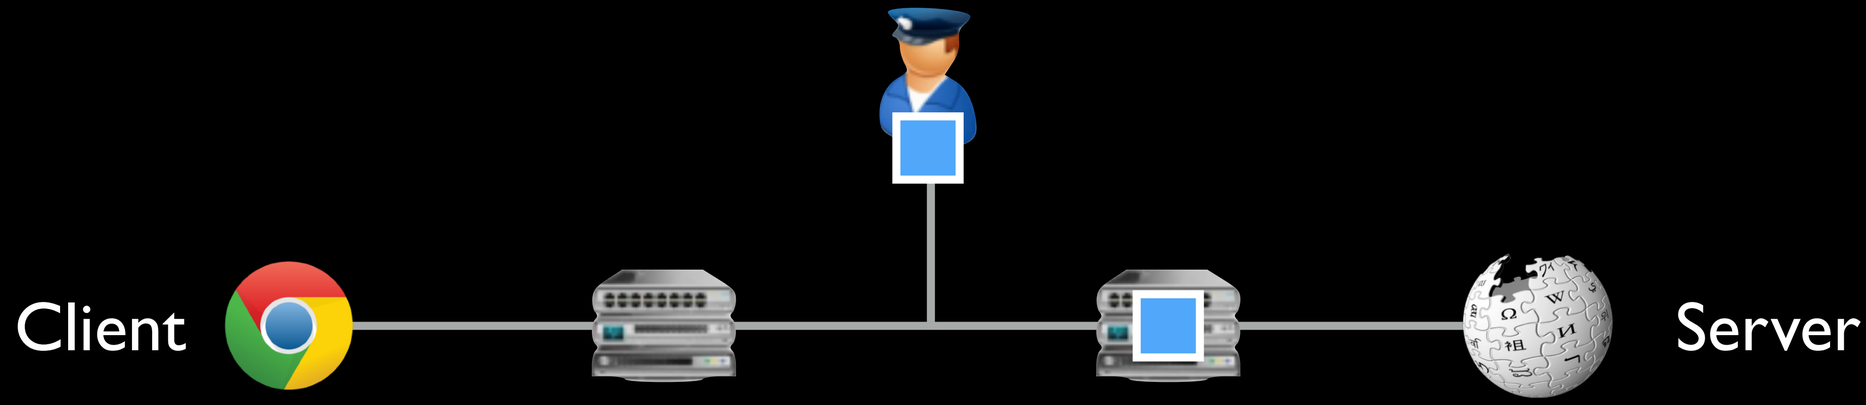

In-network censorship by nation-states

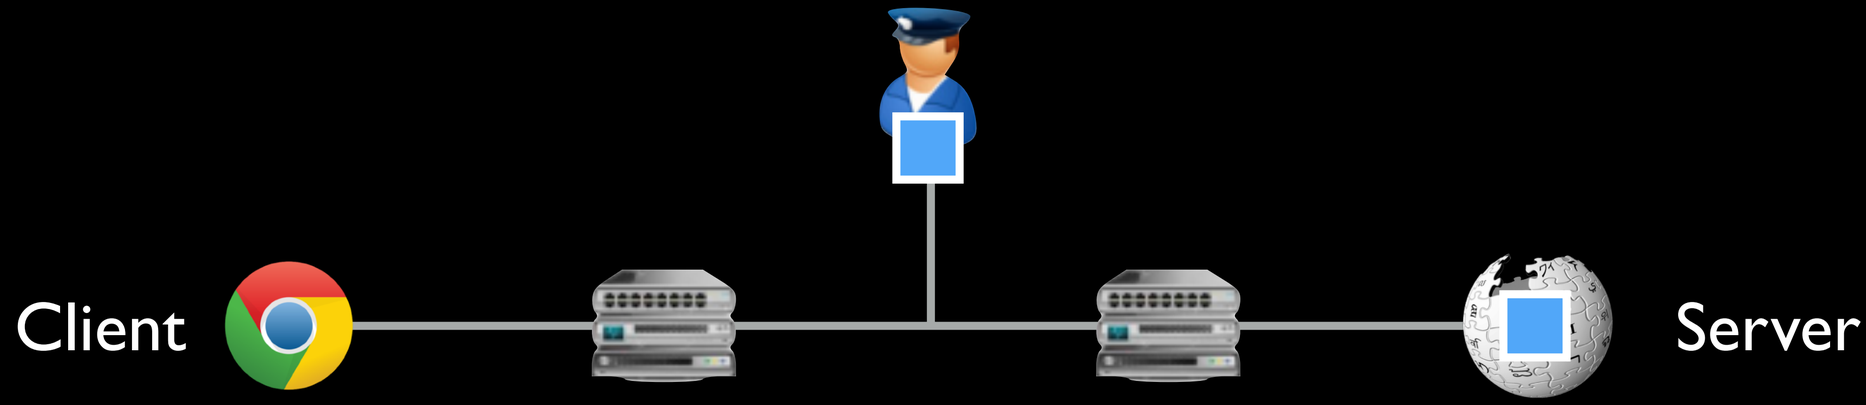

In-network censorship by nation-states

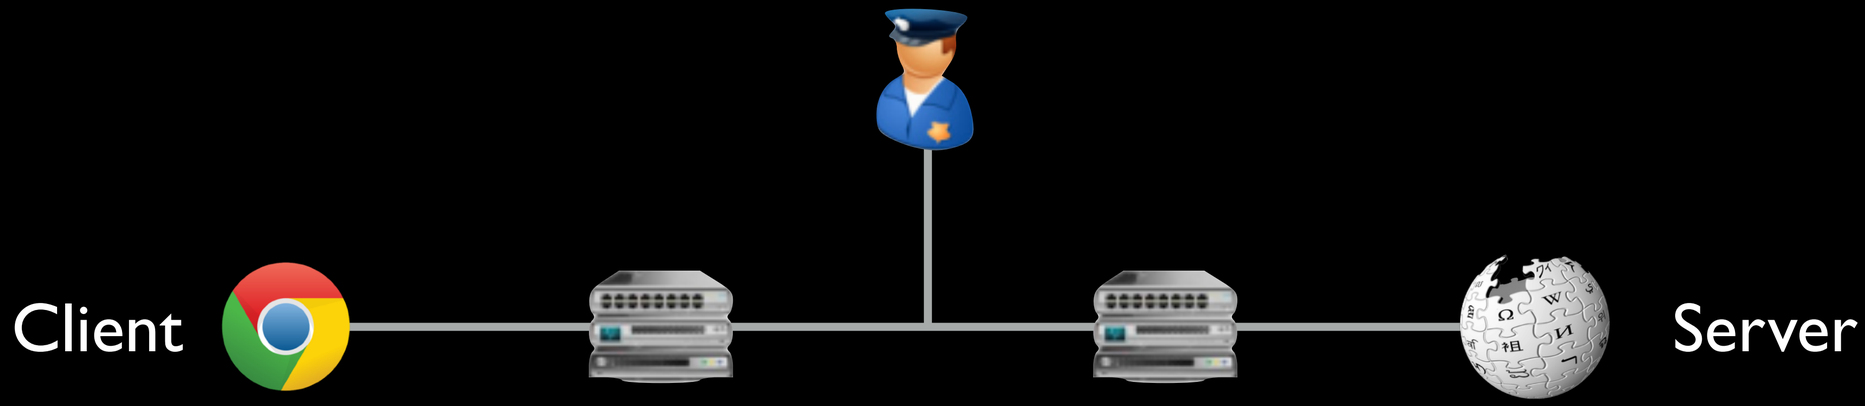

In-network censorship by nation-states

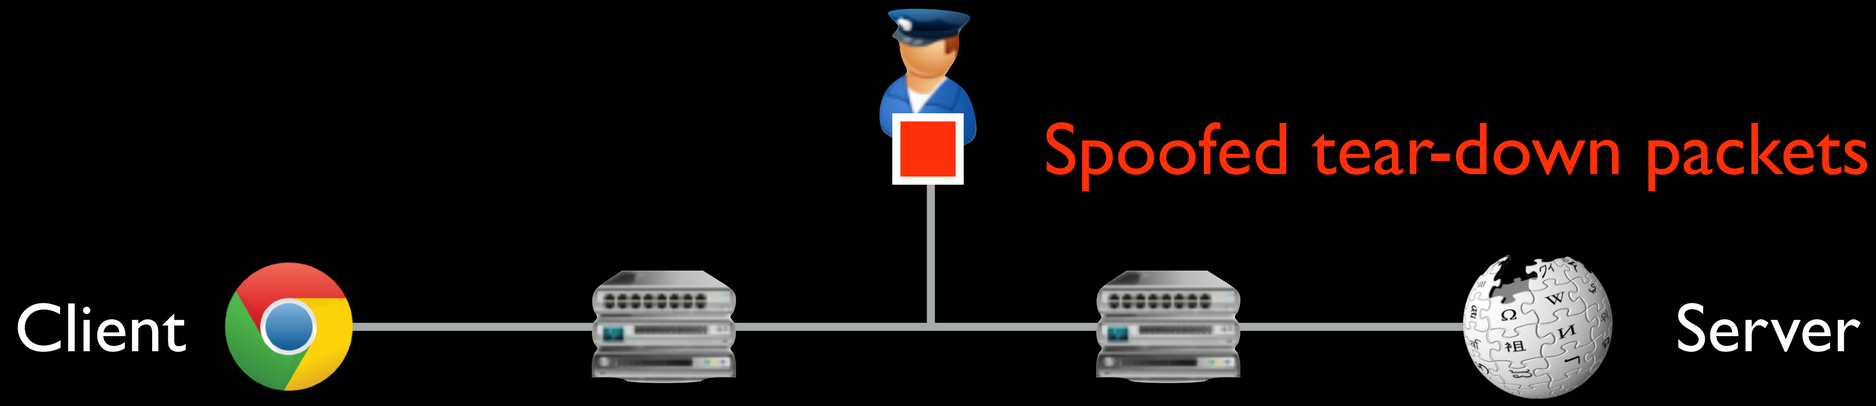

In-network censorship by nation-states

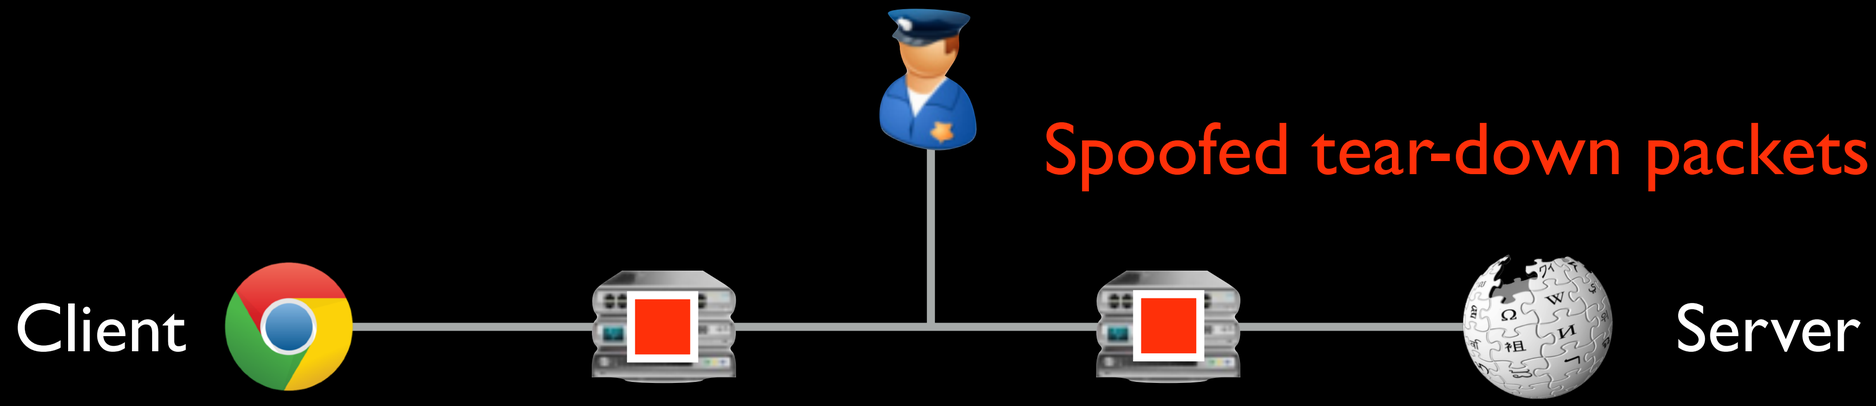

In-network censorship by nation-states

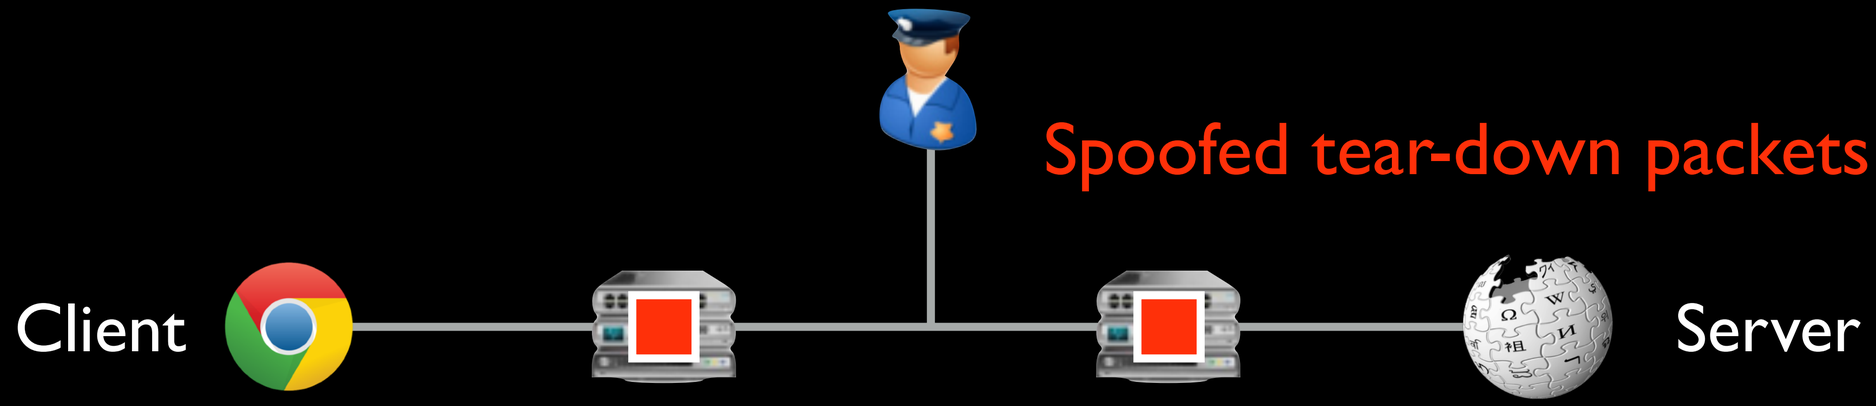

In-network censorship by nation-states

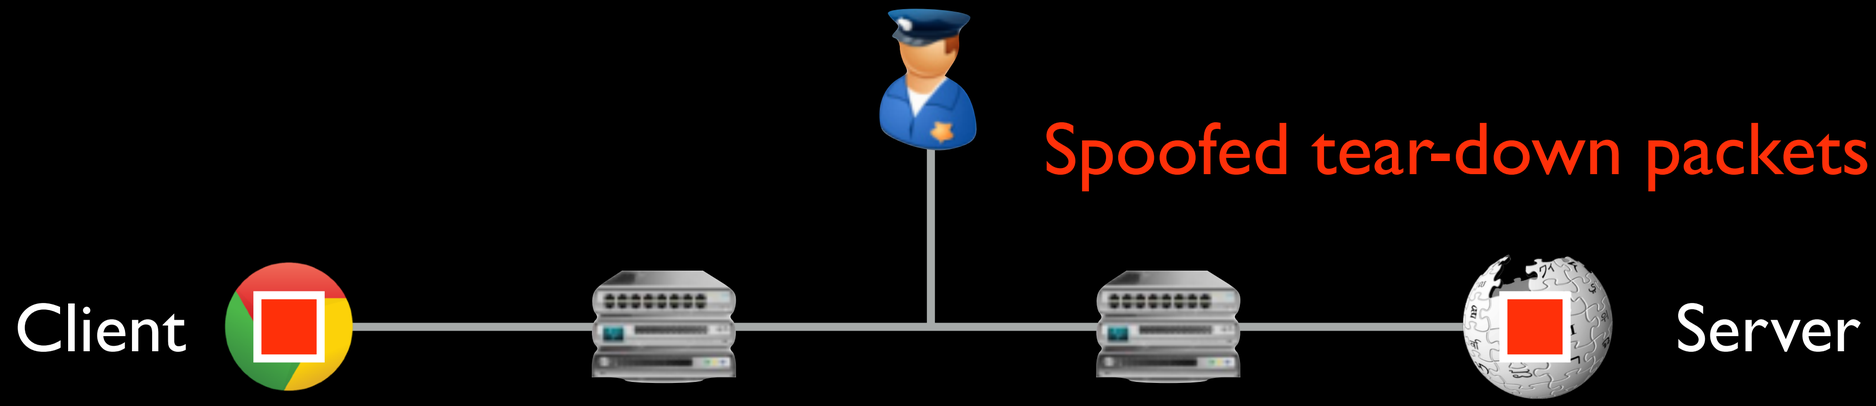

In-network censorship by nation-states

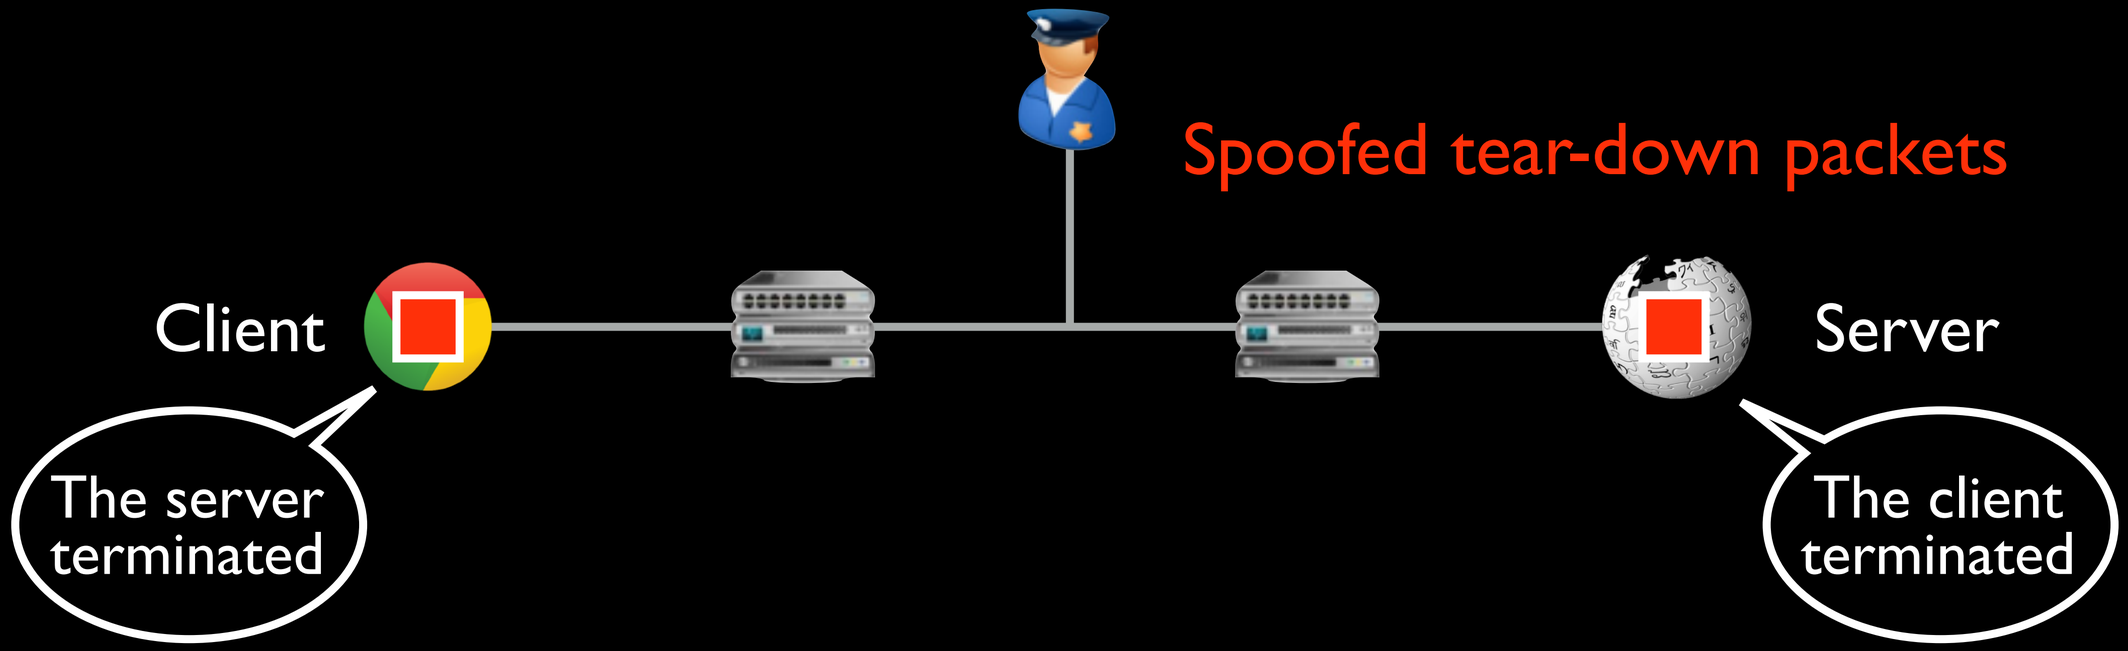

In-network censorship by nation-states

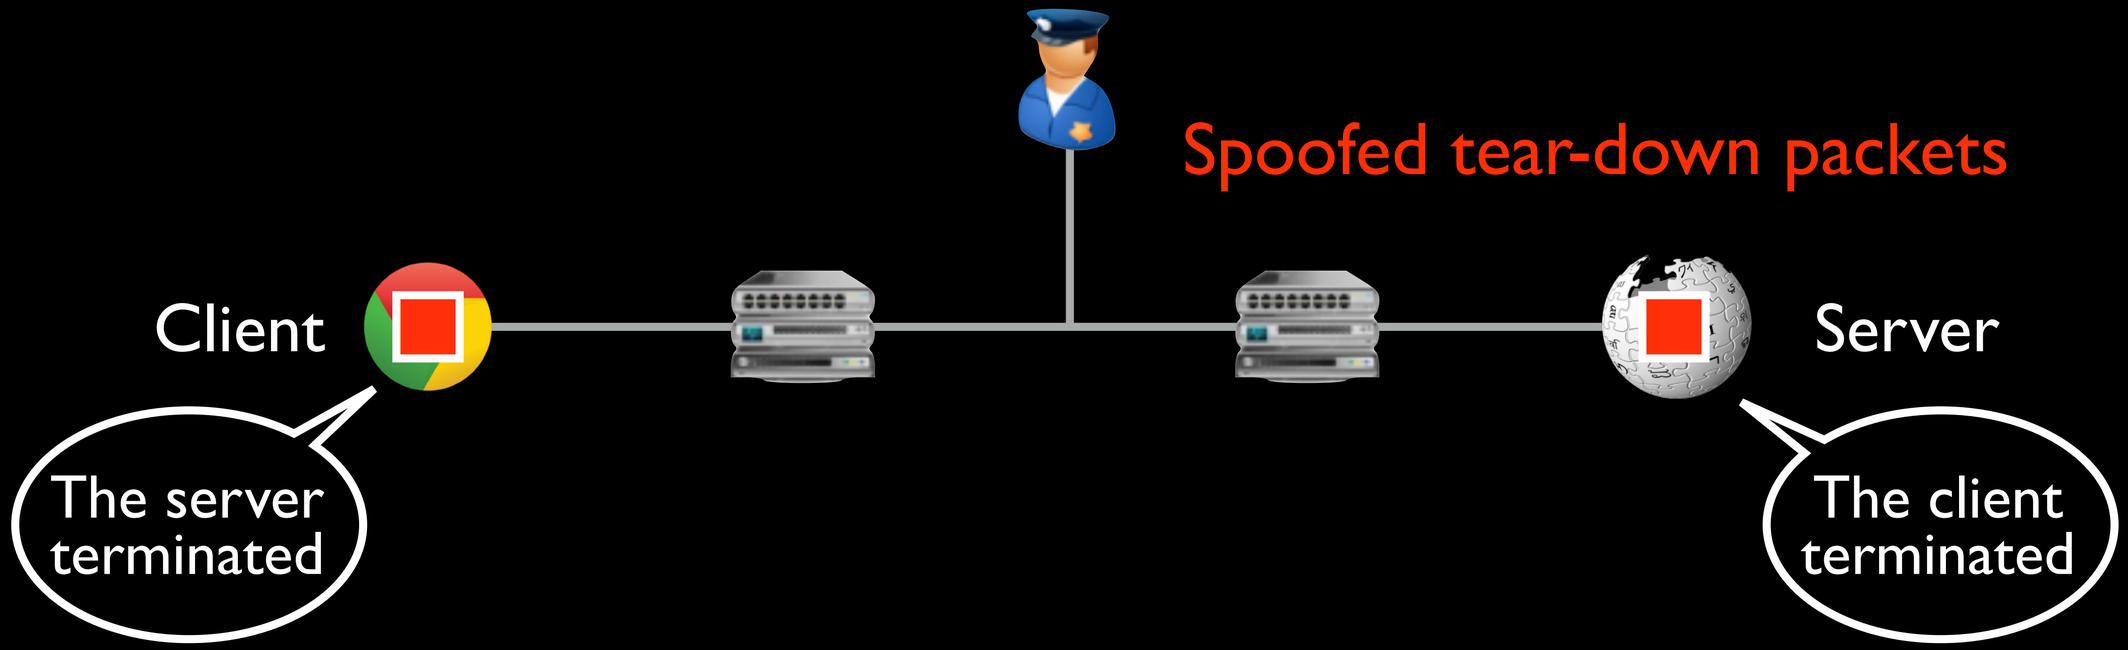

Requires *per-flow state*

In-network censorship by nation-states

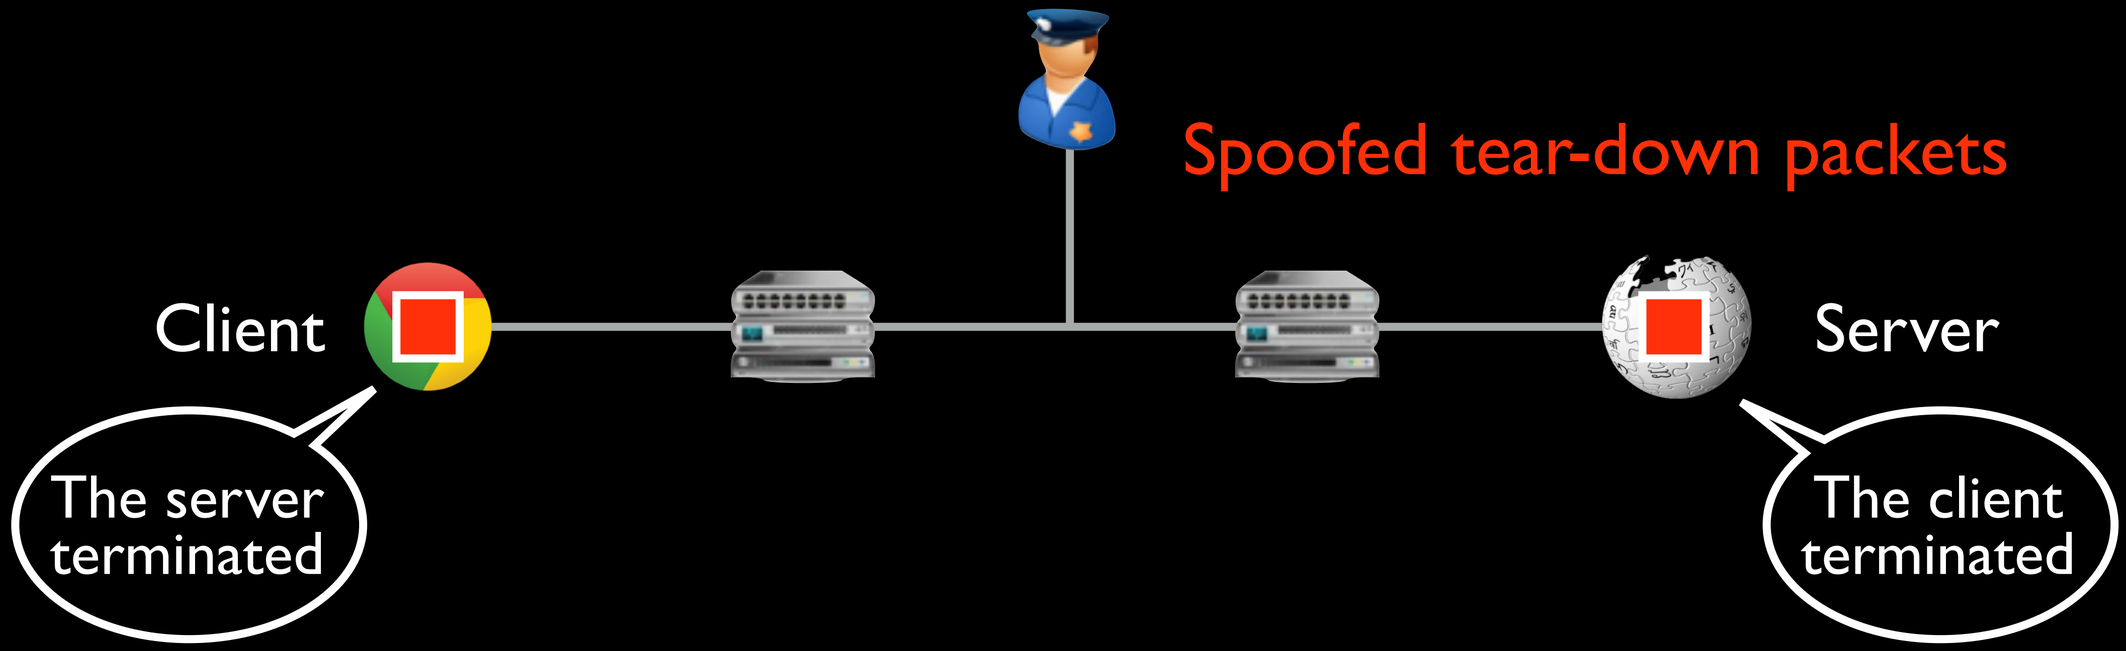

Requires *per-flow state*

Censors necessarily *take shortcuts*

In-network censorship by nation-states

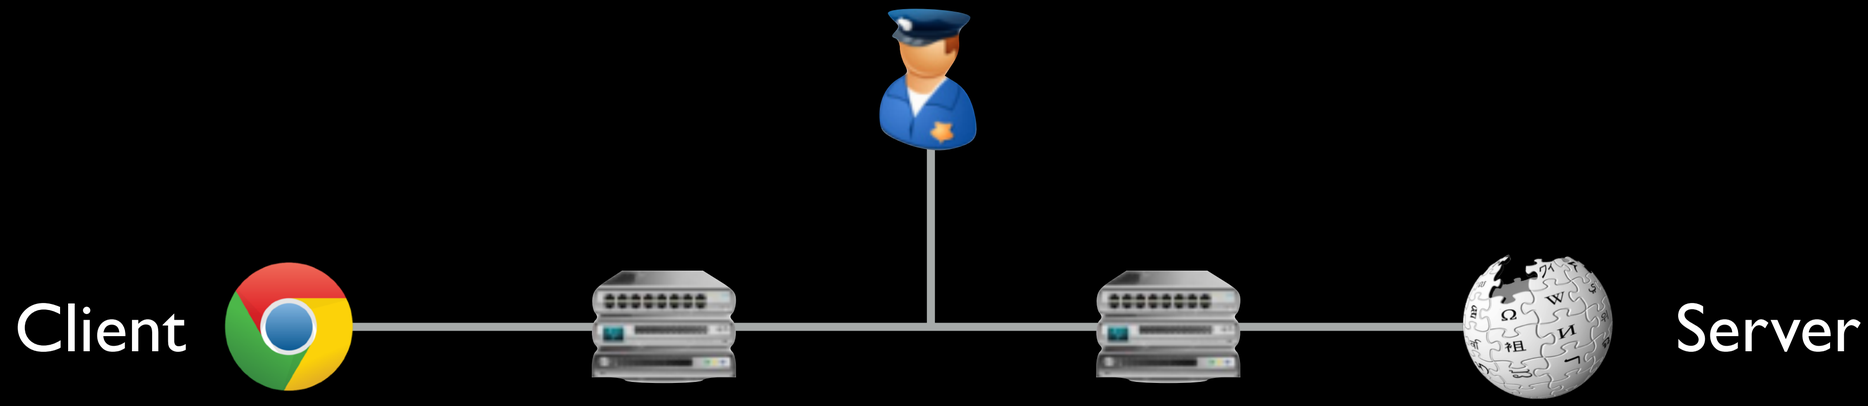

Requires *per-flow state*

Censors necessarily *take shortcuts*

Evasion can take advantage of these shortcuts

In-network censorship by nation-states

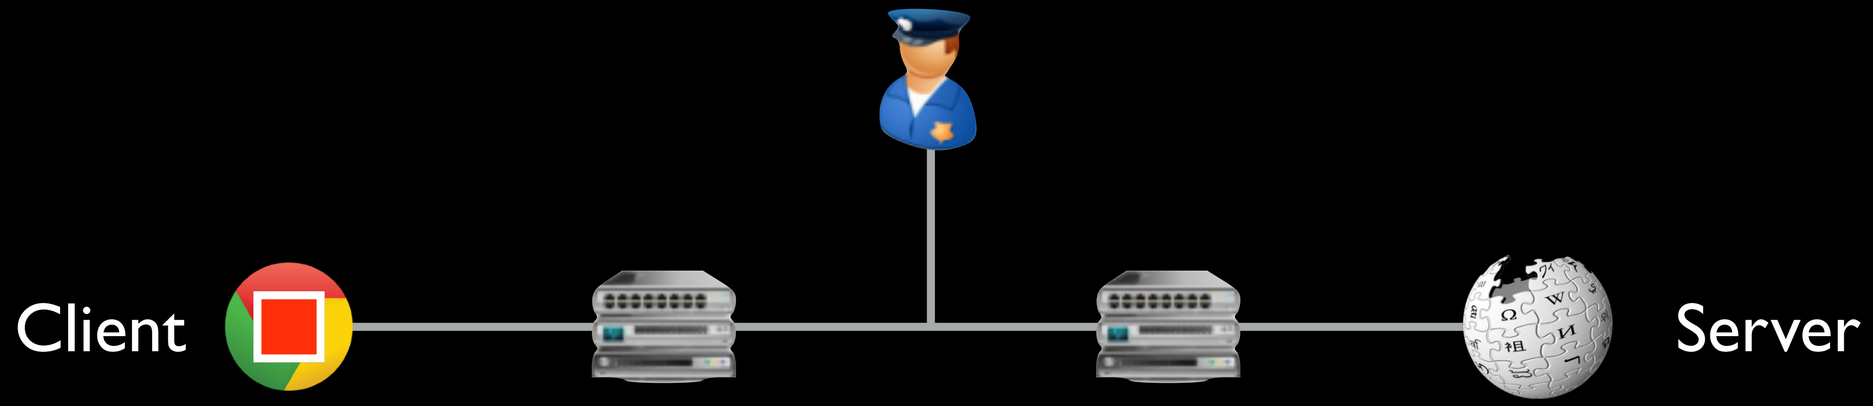

Requires *per-flow state*

Censors necessarily *take shortcuts*

Evasion can take advantage of these shortcuts

In-network censorship by nation-states

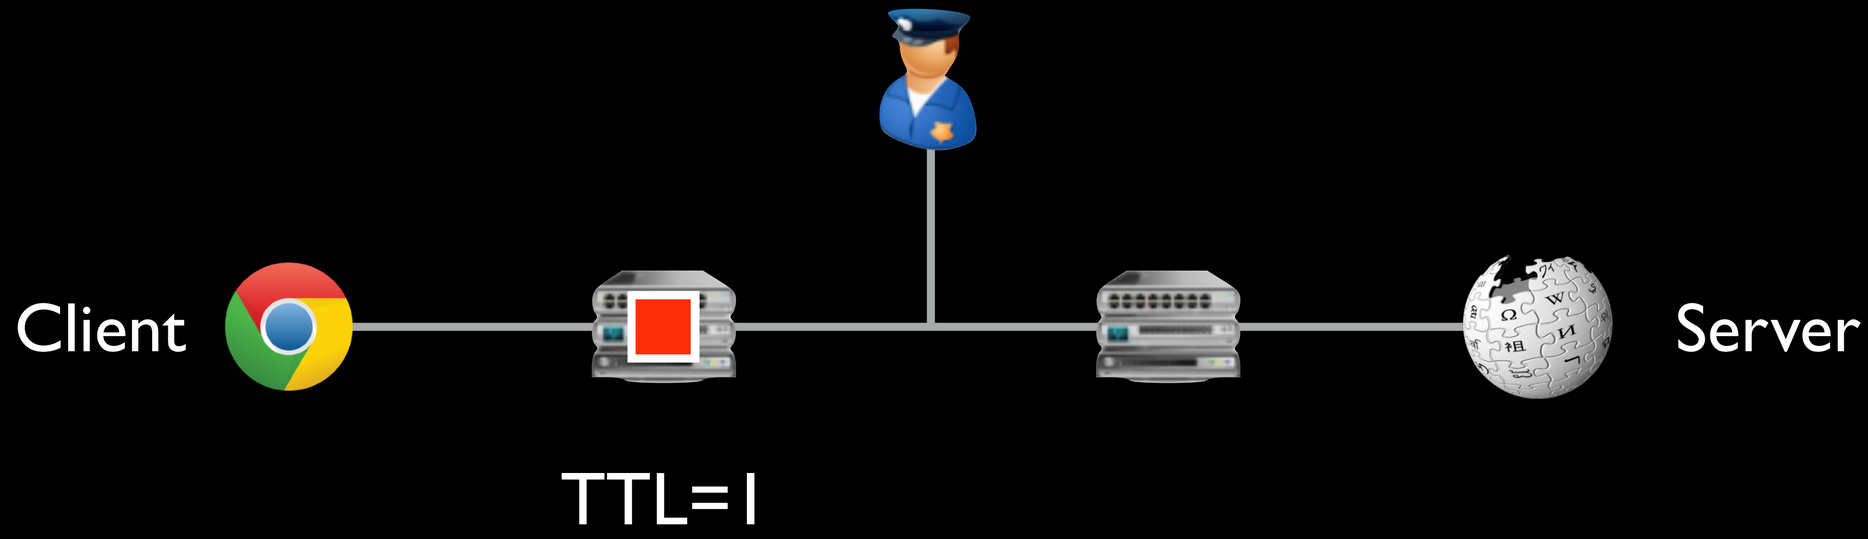

Requires *per-flow state*

Censors necessarily *take shortcuts*

Evasion can take advantage of these shortcuts

In-network censorship by nation-states

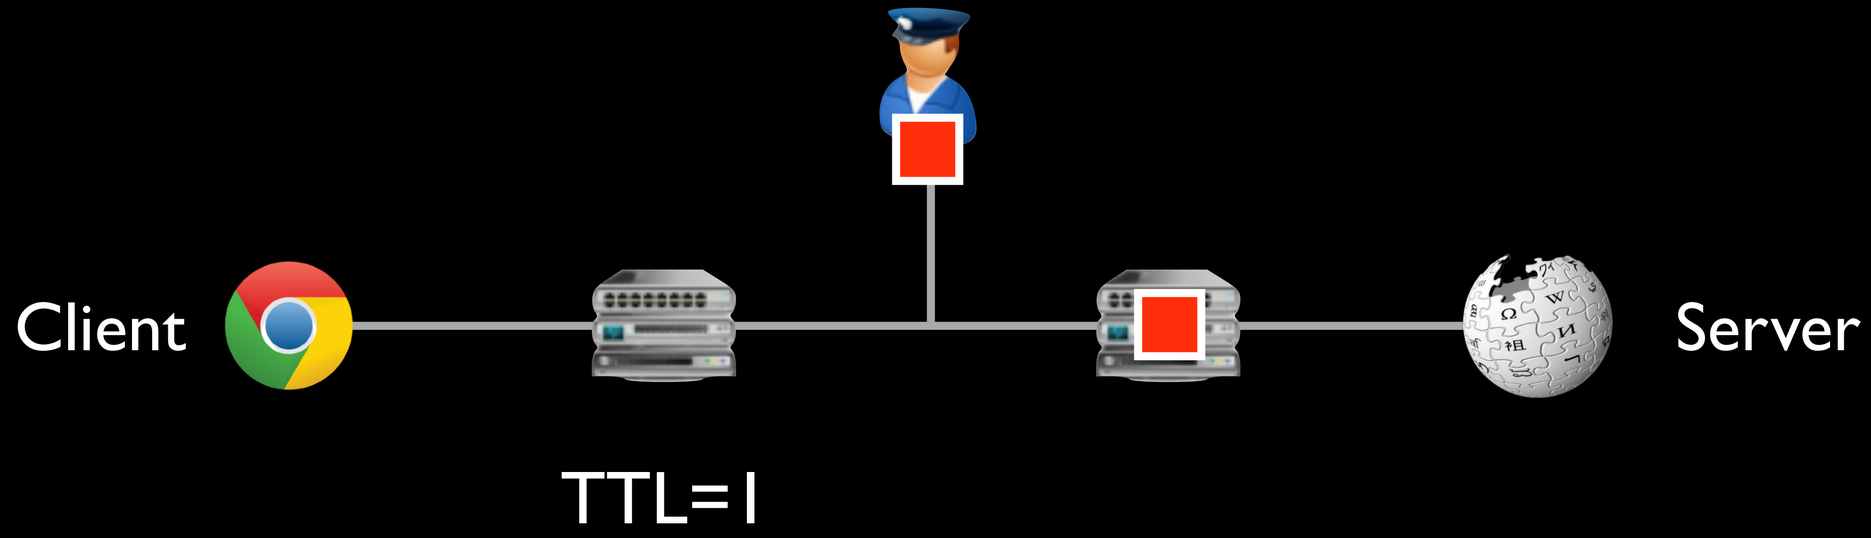

Requires *per-flow state*

Censors necessarily *take shortcuts*

Evasion can take advantage of these shortcuts

In-network censorship by nation-states

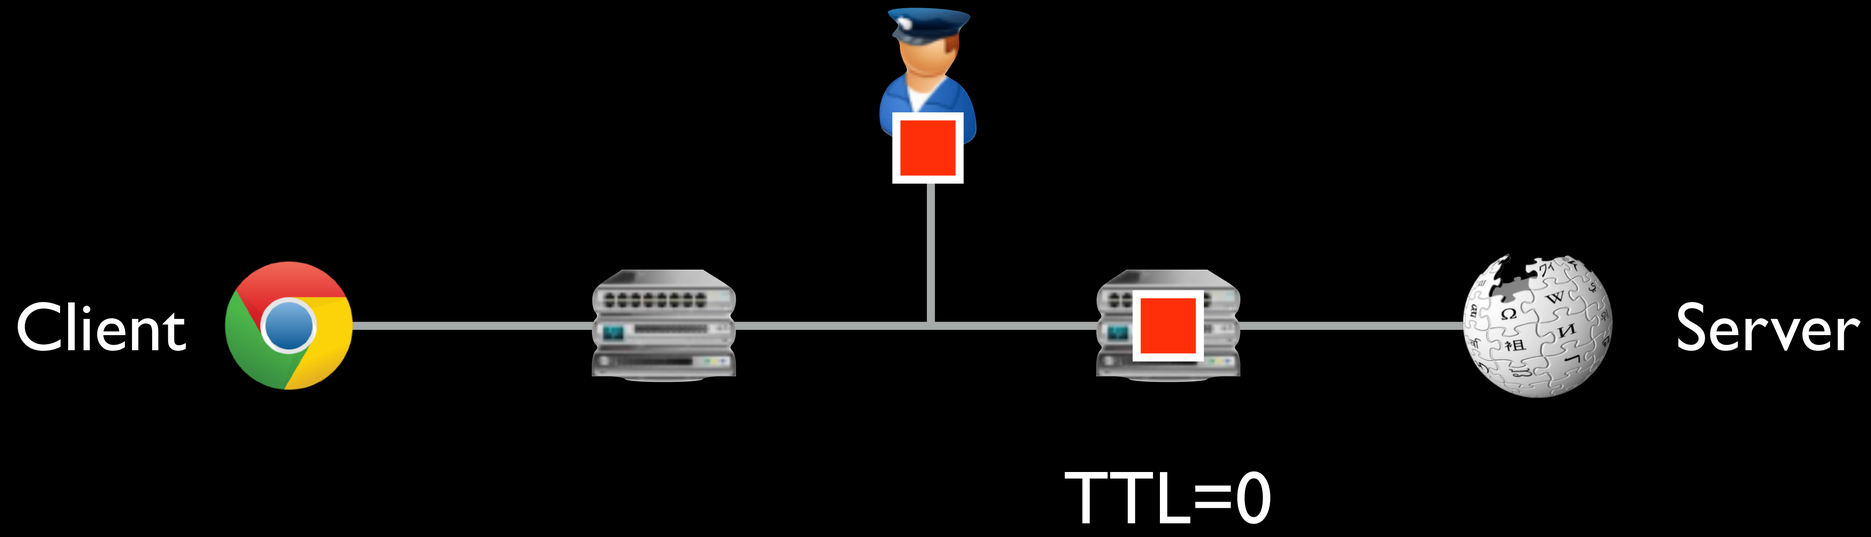

Requires *per-flow state*

Censors necessarily *take shortcuts*

Evasion can take advantage of these shortcuts

In-network censorship by nation-states

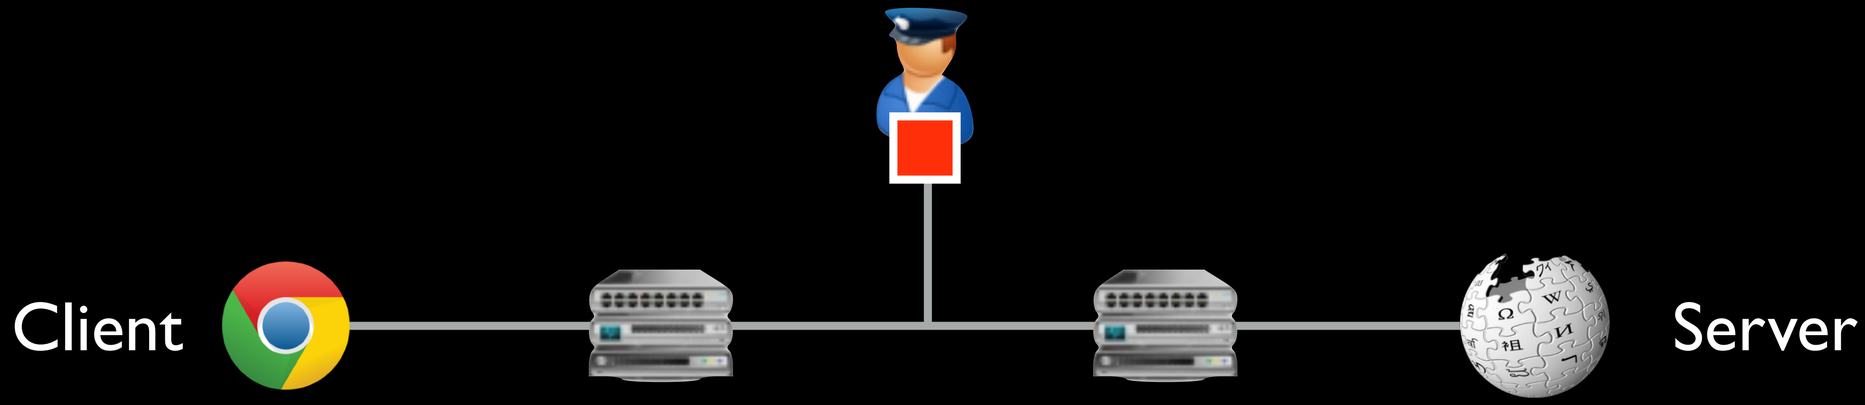

Requires *per-flow state*

Censors necessarily *take shortcuts*

Evasion can take advantage of these shortcuts

In-network censorship by nation-states

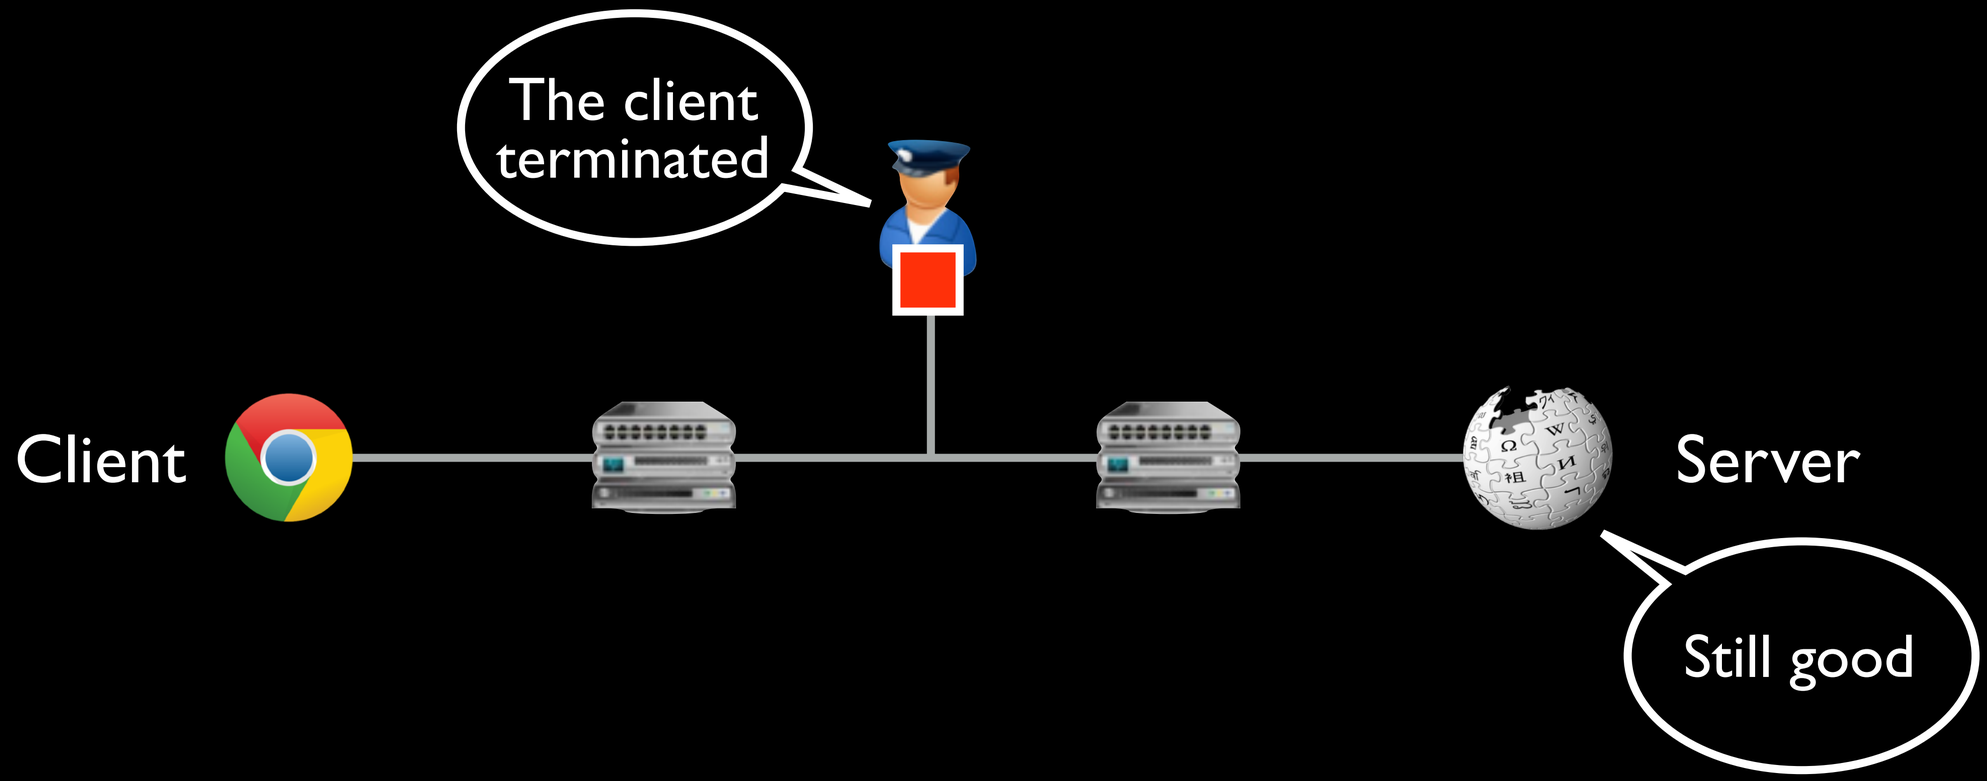

Requires *per-flow state*

Censors necessarily *take shortcuts*

Evasion can take advantage of these shortcuts

Censorship evasion research

Censorship evasion research

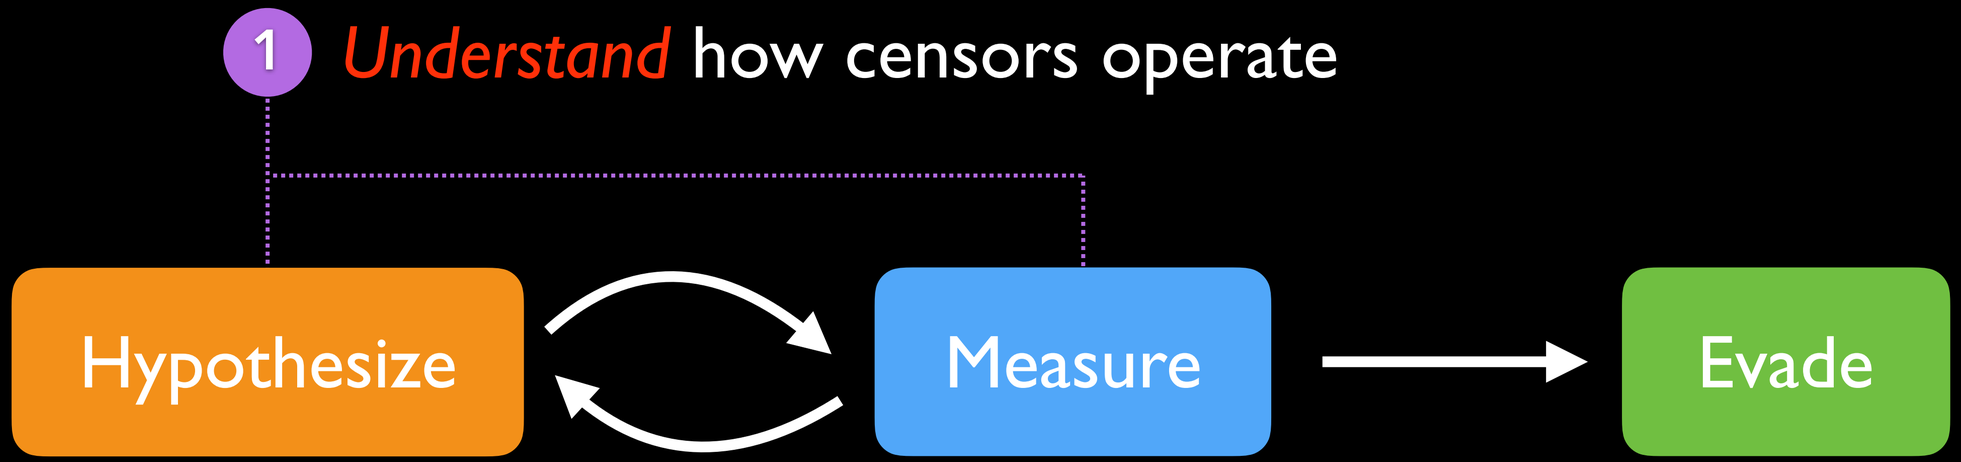

Censorship evasion research

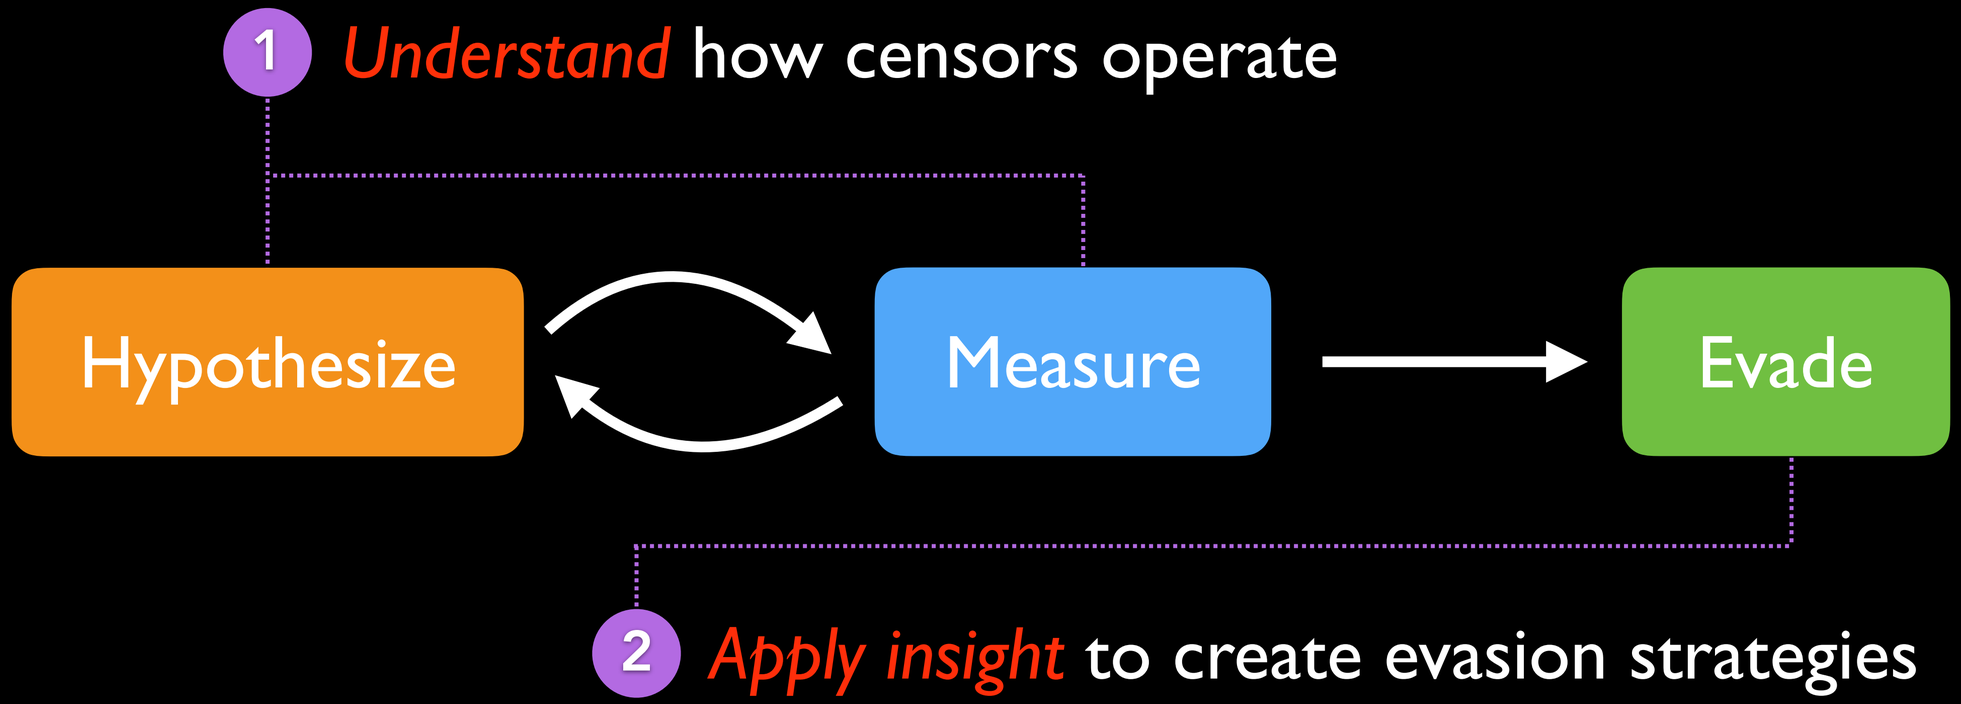

Censorship evasion research

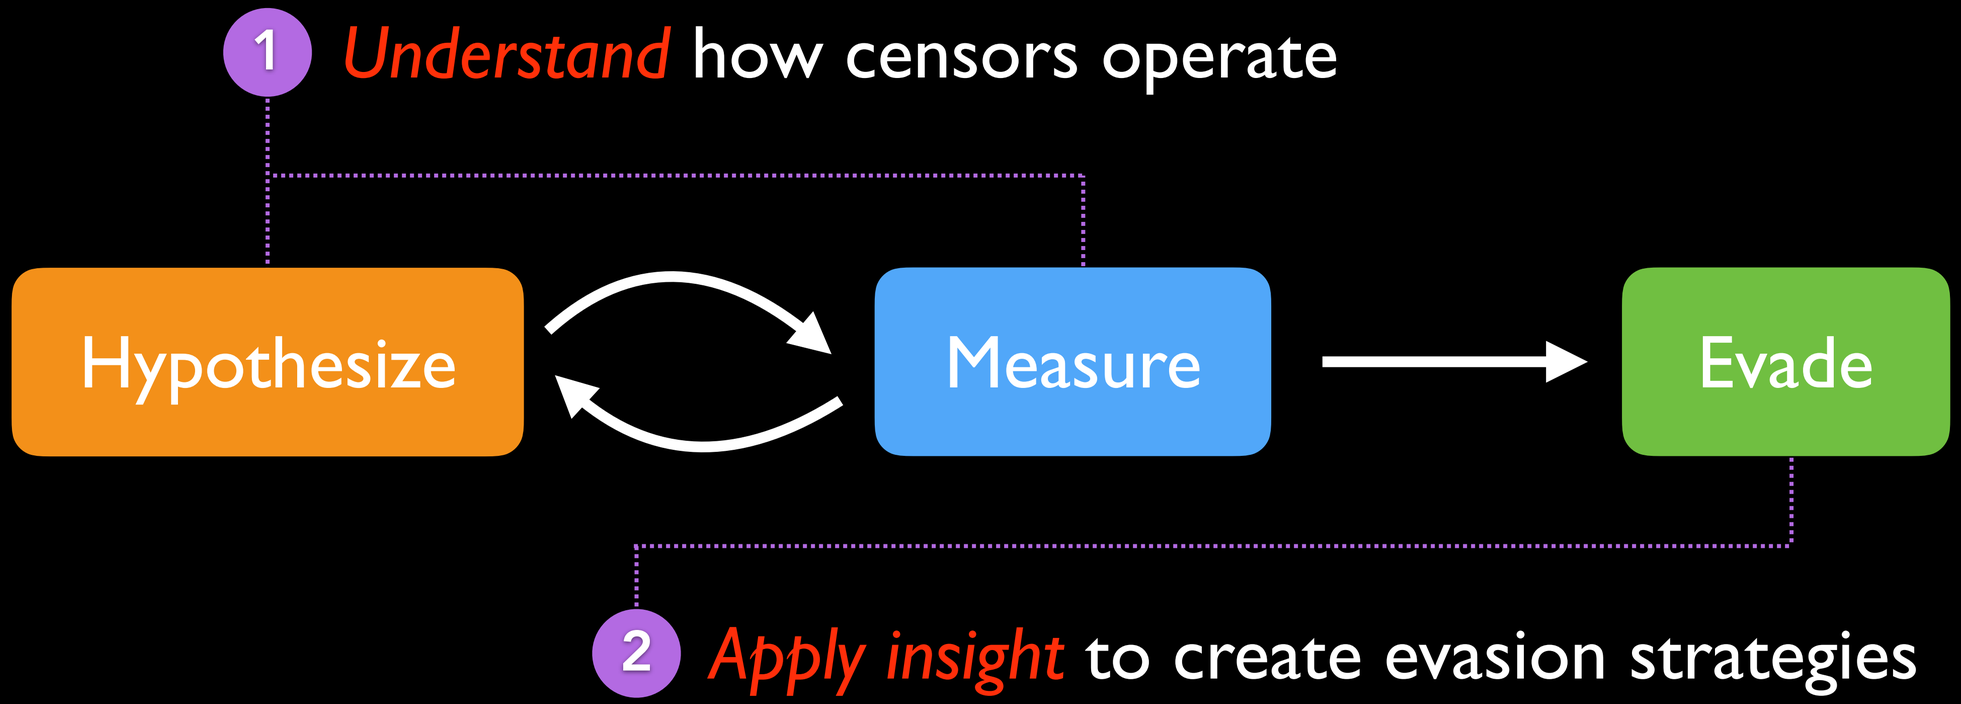

Largely manual efforts give censors the advantage

Censorship evasion research

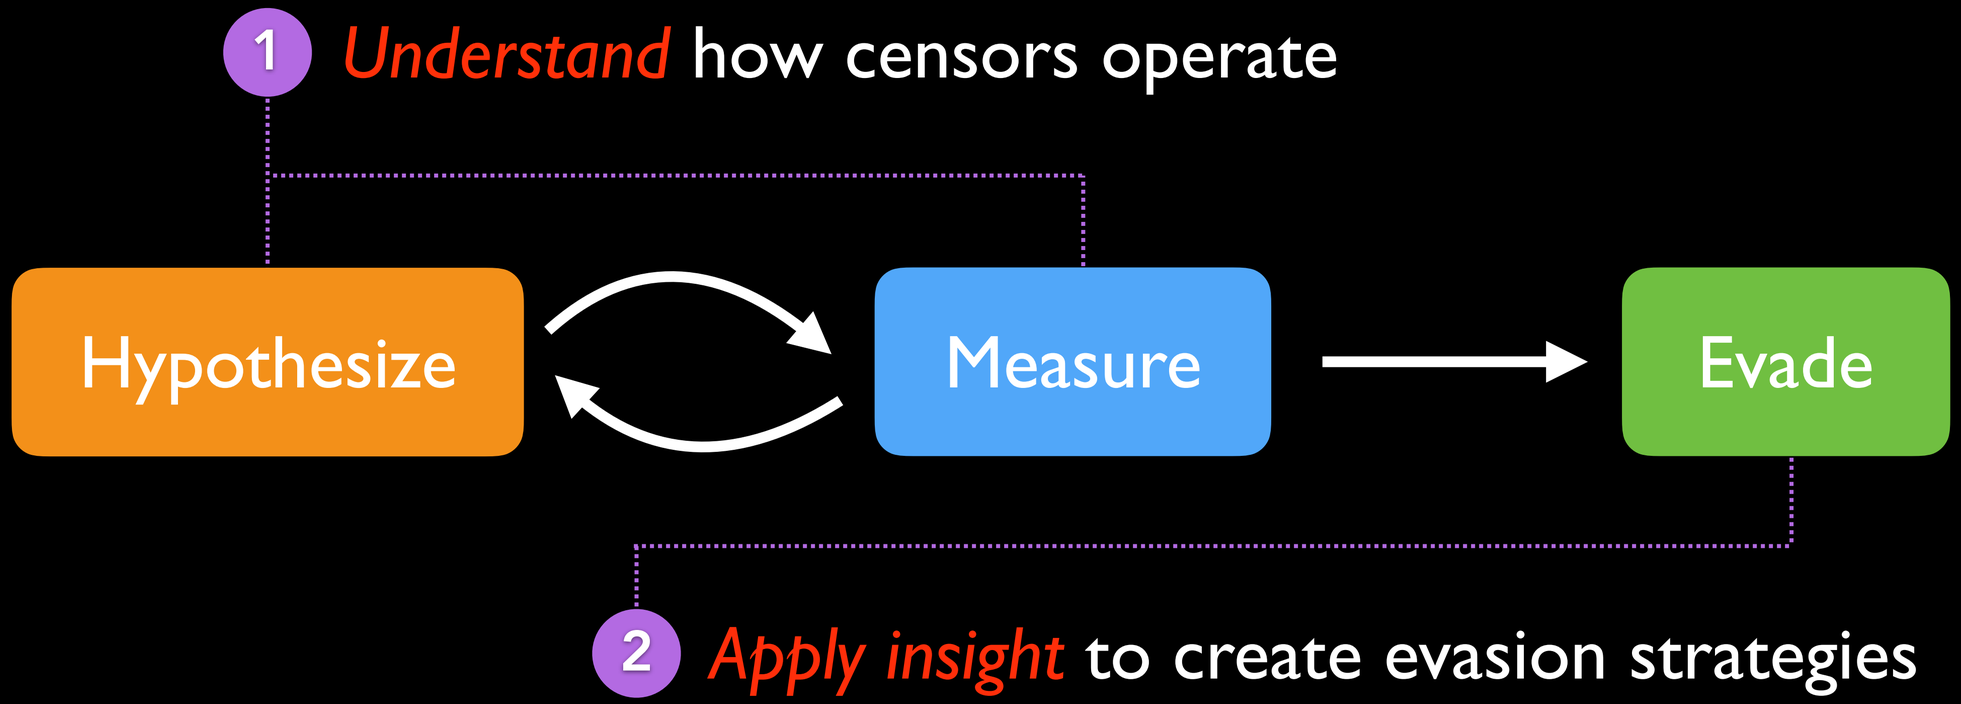

Largely manual efforts give censors the advantage

Our work gives evasion the advantage

AI-assisted censorship evasion research

AI-assisted censorship evasion research

1 Use AI to automatically learn new evasion strategies

AI-assisted censorship evasion research

1. Use AI to automatically learn new evasion strategies

2. Use the strategies the AI finds to understand how the censor works

Geneva

Genetic Evasion

1. Use AI to automatically learn new evasion strategies

2. Use the strategies the AI finds to understand how the censor works

Geneva

Genetic Evasion

Building Blocks

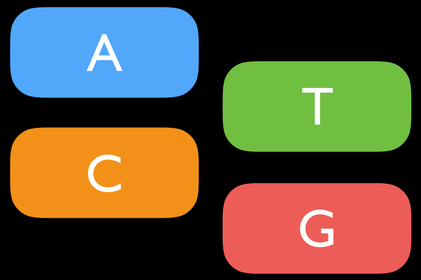

Composition

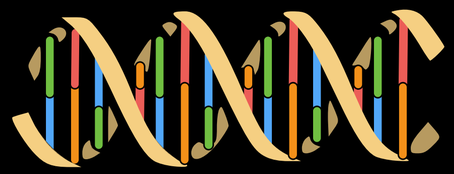

Mutation

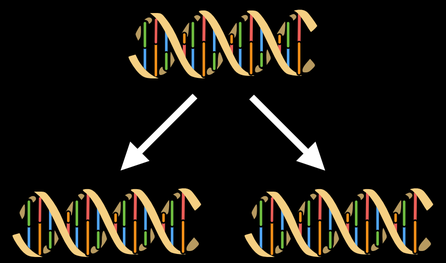

Fitness

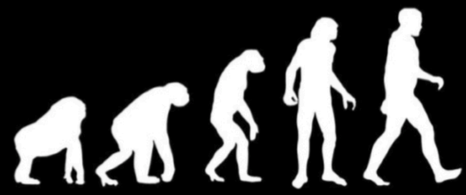

Geneva

Genetic Evasion

Building Blocks

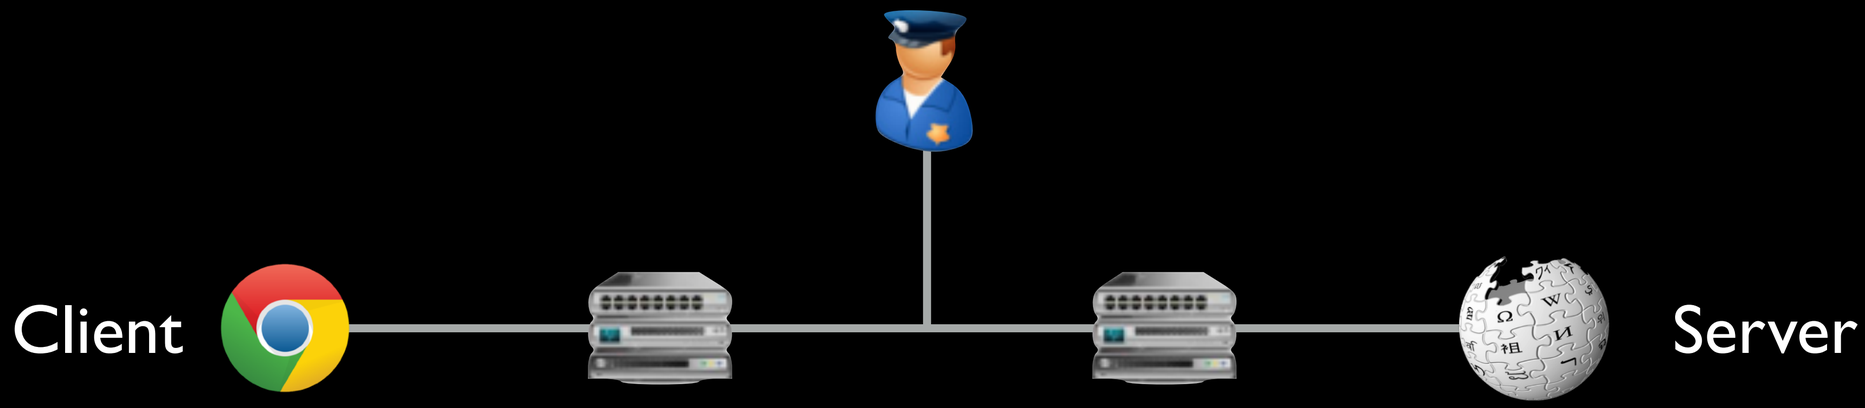

Geneva

Genetic Evasion

Building Blocks

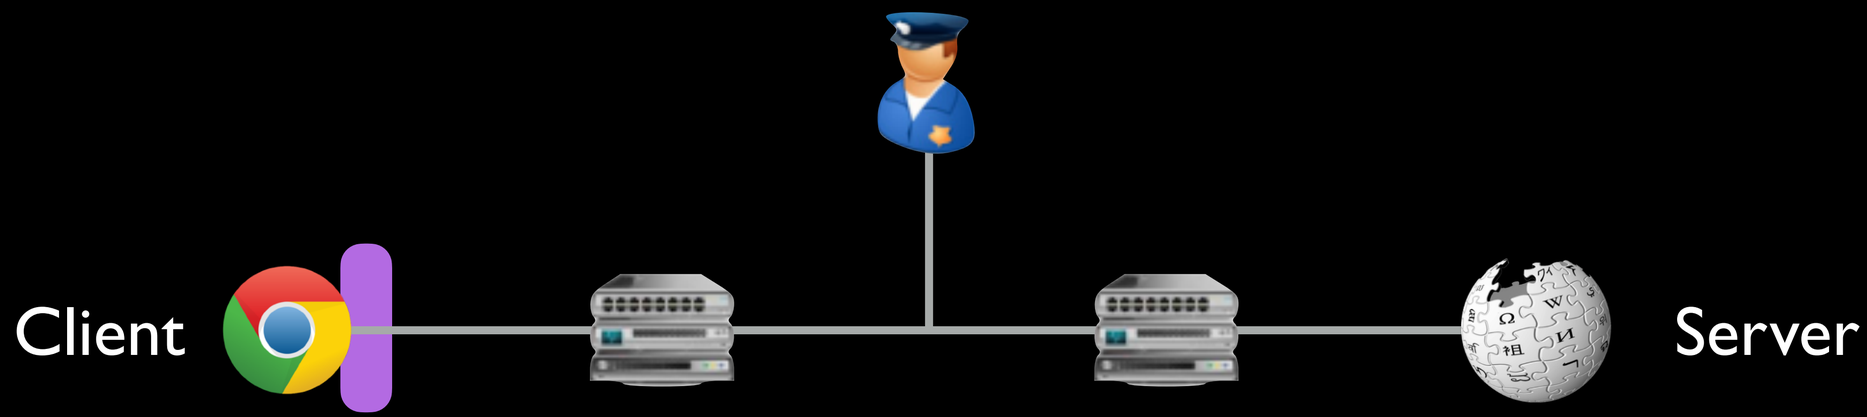

Geneva runs strictly at the client

Manipulates packets to and from the client

Geneva

Genetic Evasion

Building Blocks

Manipulates packets to and from the client

Geneva

Genetic Evasion

Building Blocks

Manipulates packets to and from the client

Bit manipulation

Versatile but inefficient

Geneva

Genetic Evasion

Building Blocks

Manipulates packets to and from the client

Bit manipulation

Versatile but inefficient

Known strategies

Efficient but limited

Geneva

Genetic Evasion

Building Blocks

Manipulates packets to and from the client

- Duplicate
- Tamper
- Fragment
- Drop

Geneva

Genetic Evasion

Building Blocks

Manipulates packets to and from the client

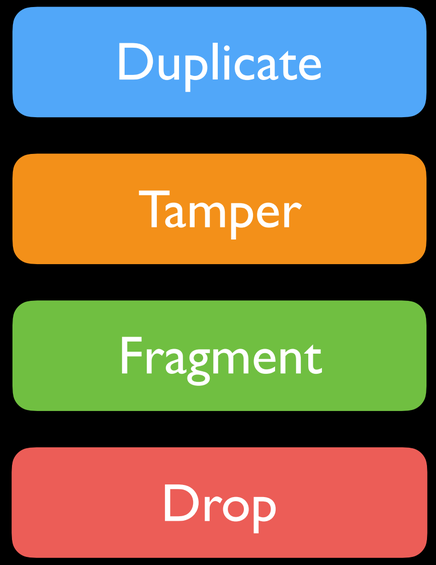

Alter or corrupt
any TCP/IP header field

No semantic understanding of what the fields mean

Geneva

Genetic Evasion

Building Blocks

Manipulates packets to and from the client

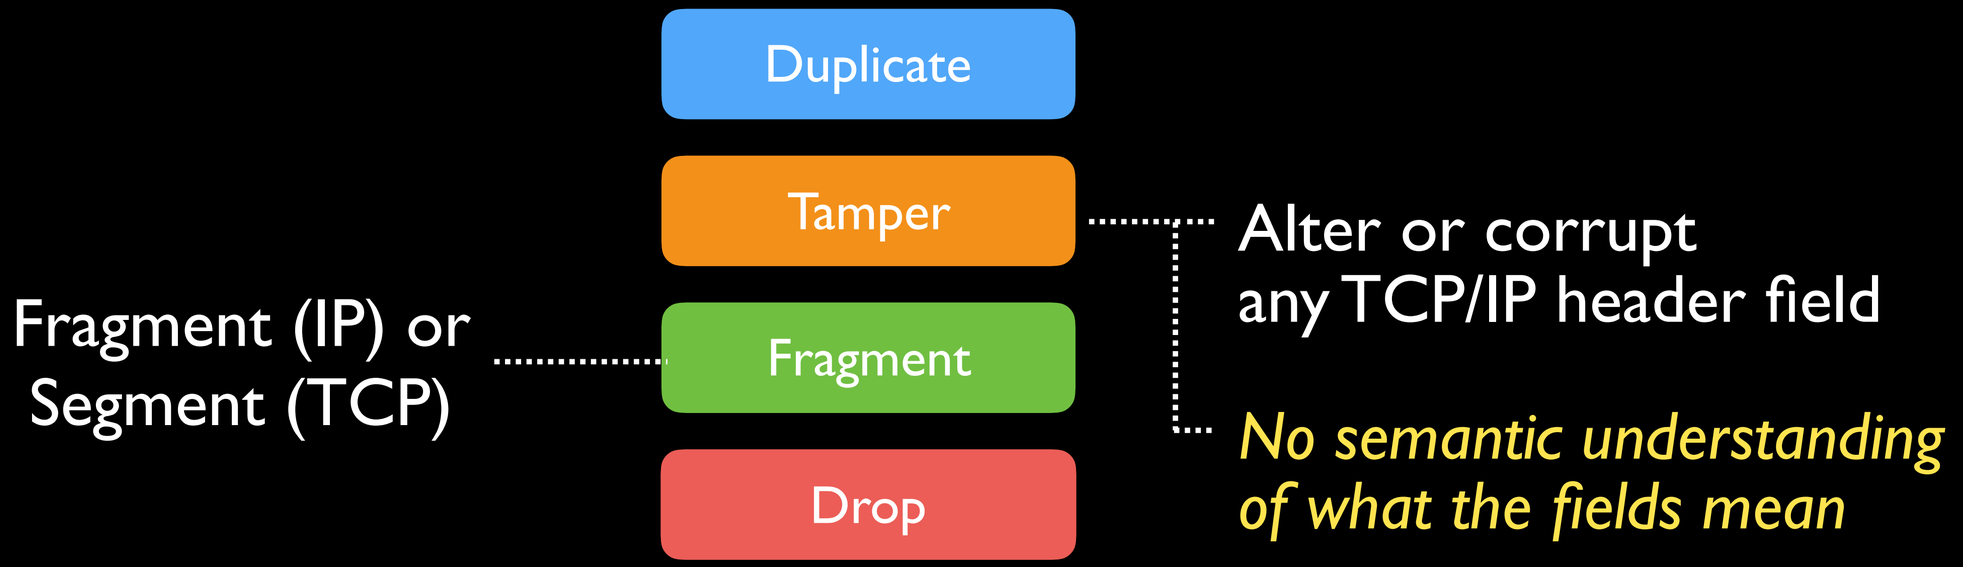

Geneva

Genetic Evasion

Building Blocks

Actions manipulate individual packets

- Duplicate
- Tamper
- Fragment
- Drop

Composition

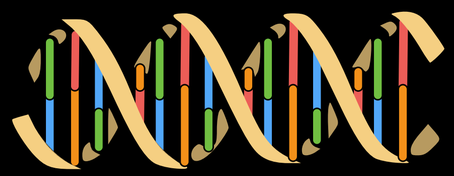

Mutation

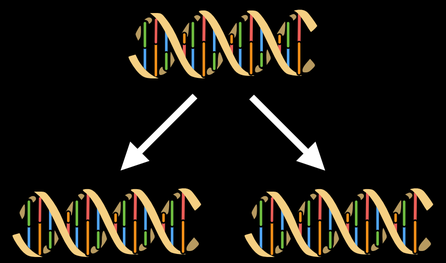

Fitness

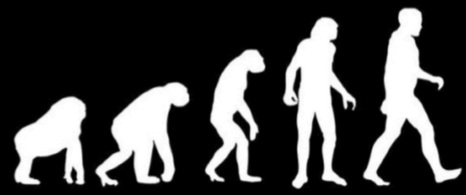

Geneva

Genetic Evasion

Building Blocks

Actions manipulate individual packets

- Duplicate
- Tamper
- Fragment
- Drop

Composition

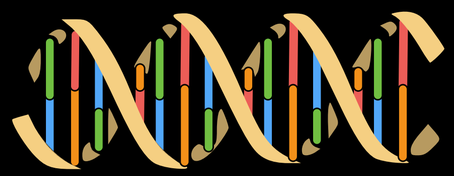

Mutation

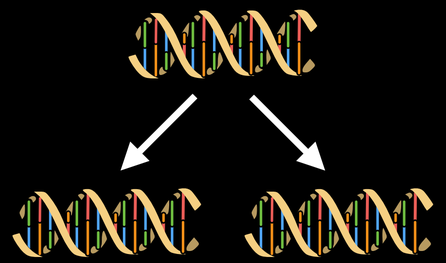

Fitness

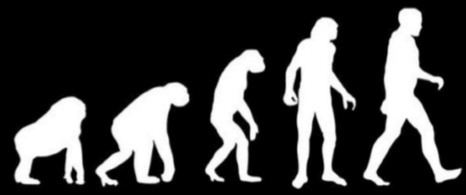

Geneva

Genetic Evasion

Composition

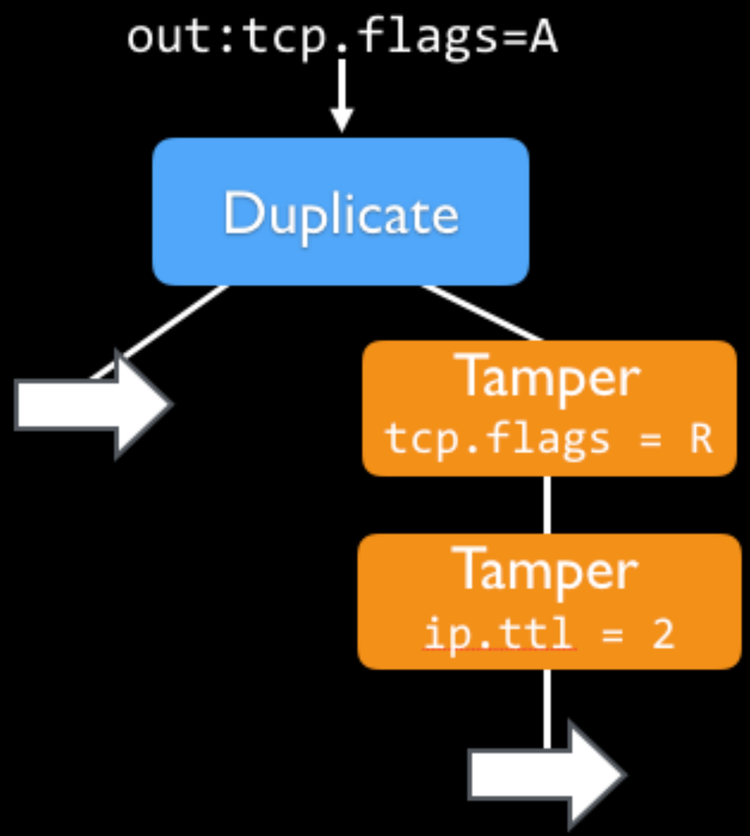

Running a Strategy

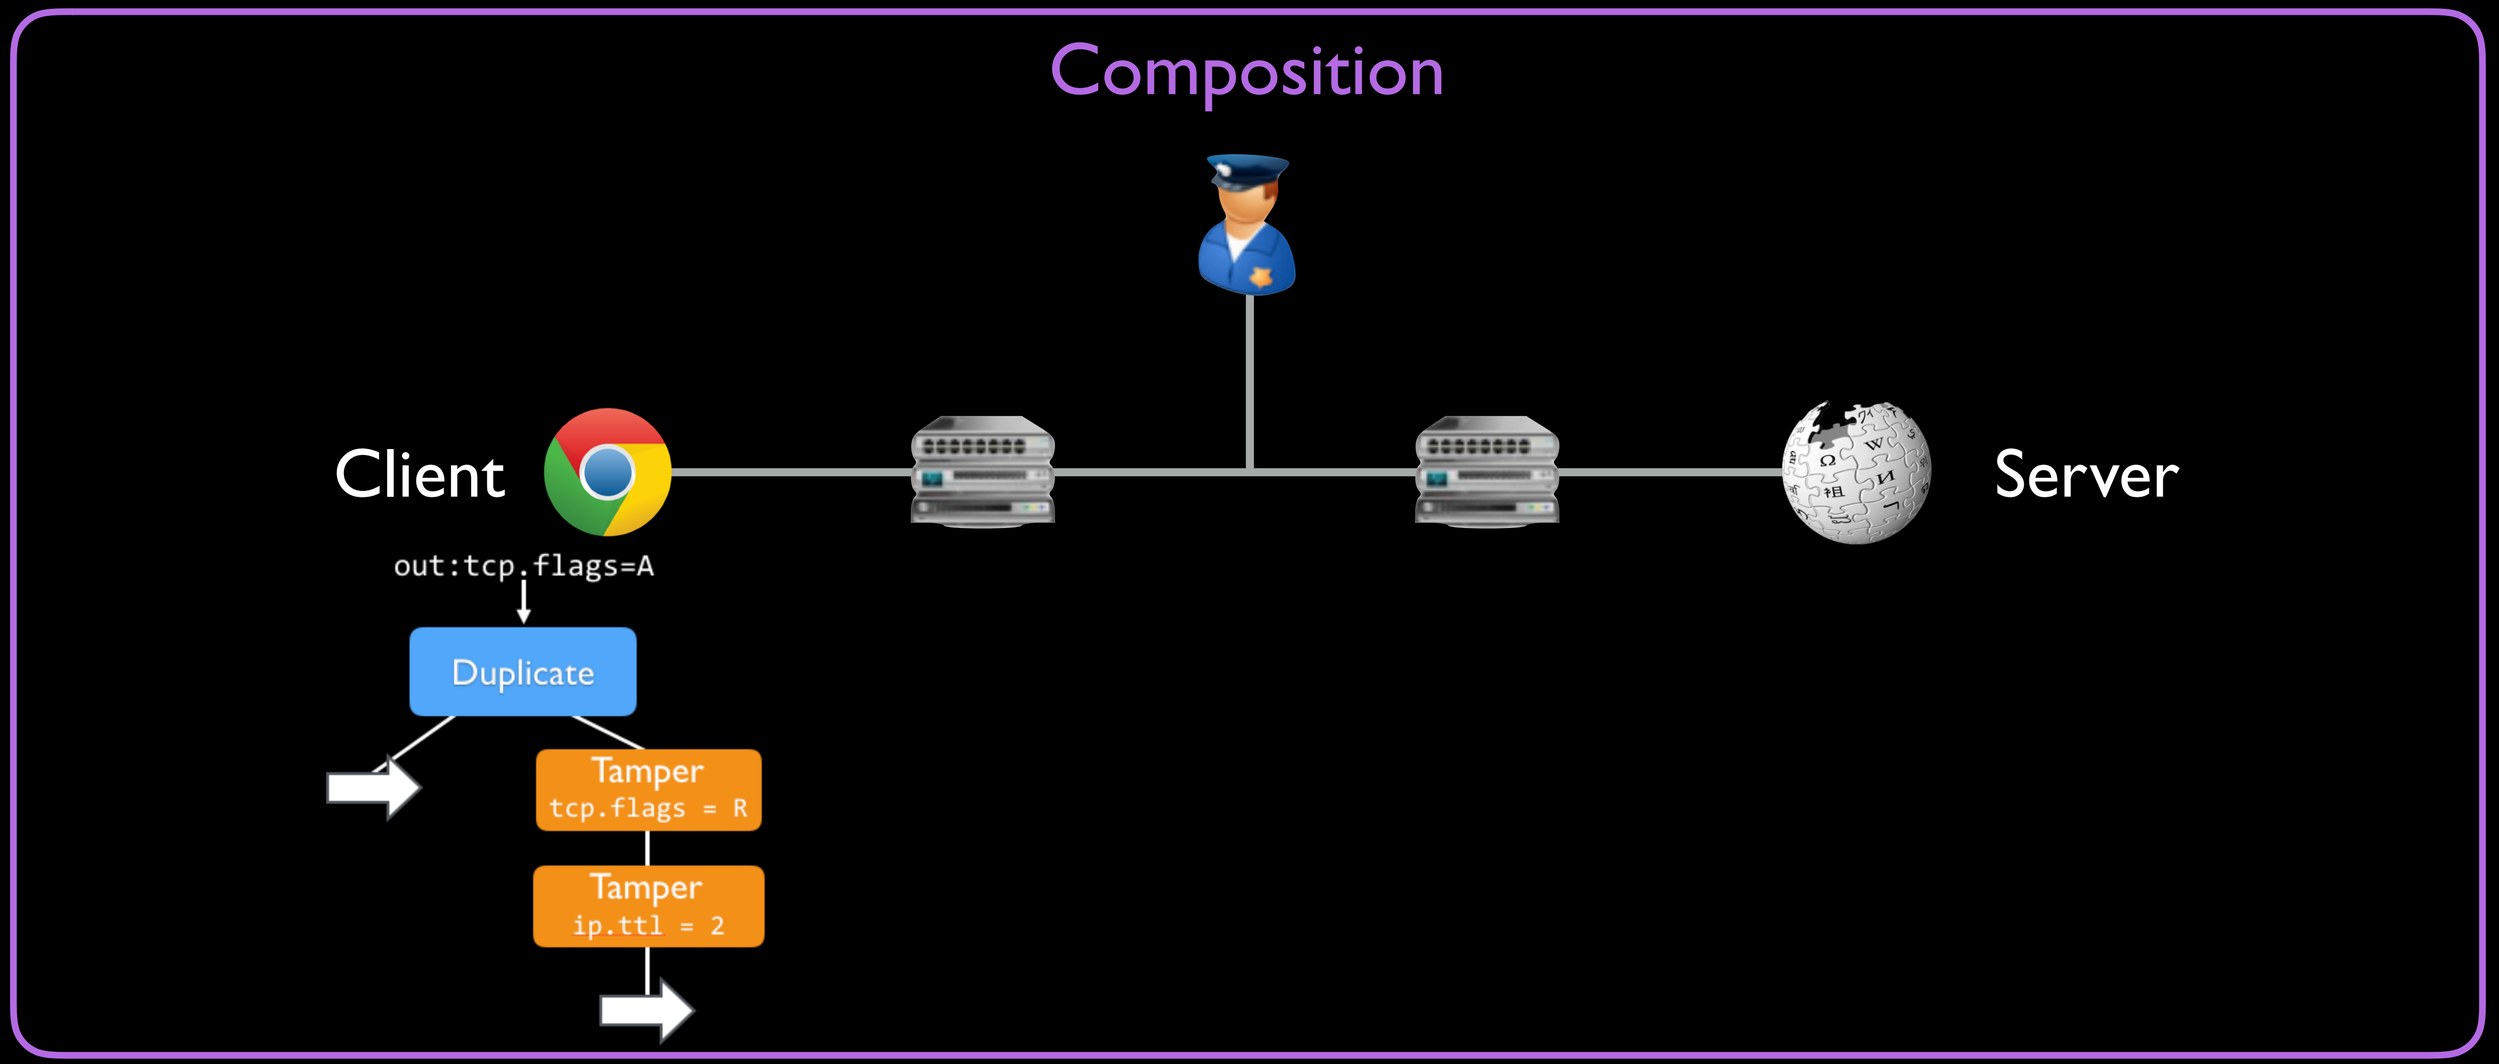

Running a Strategy

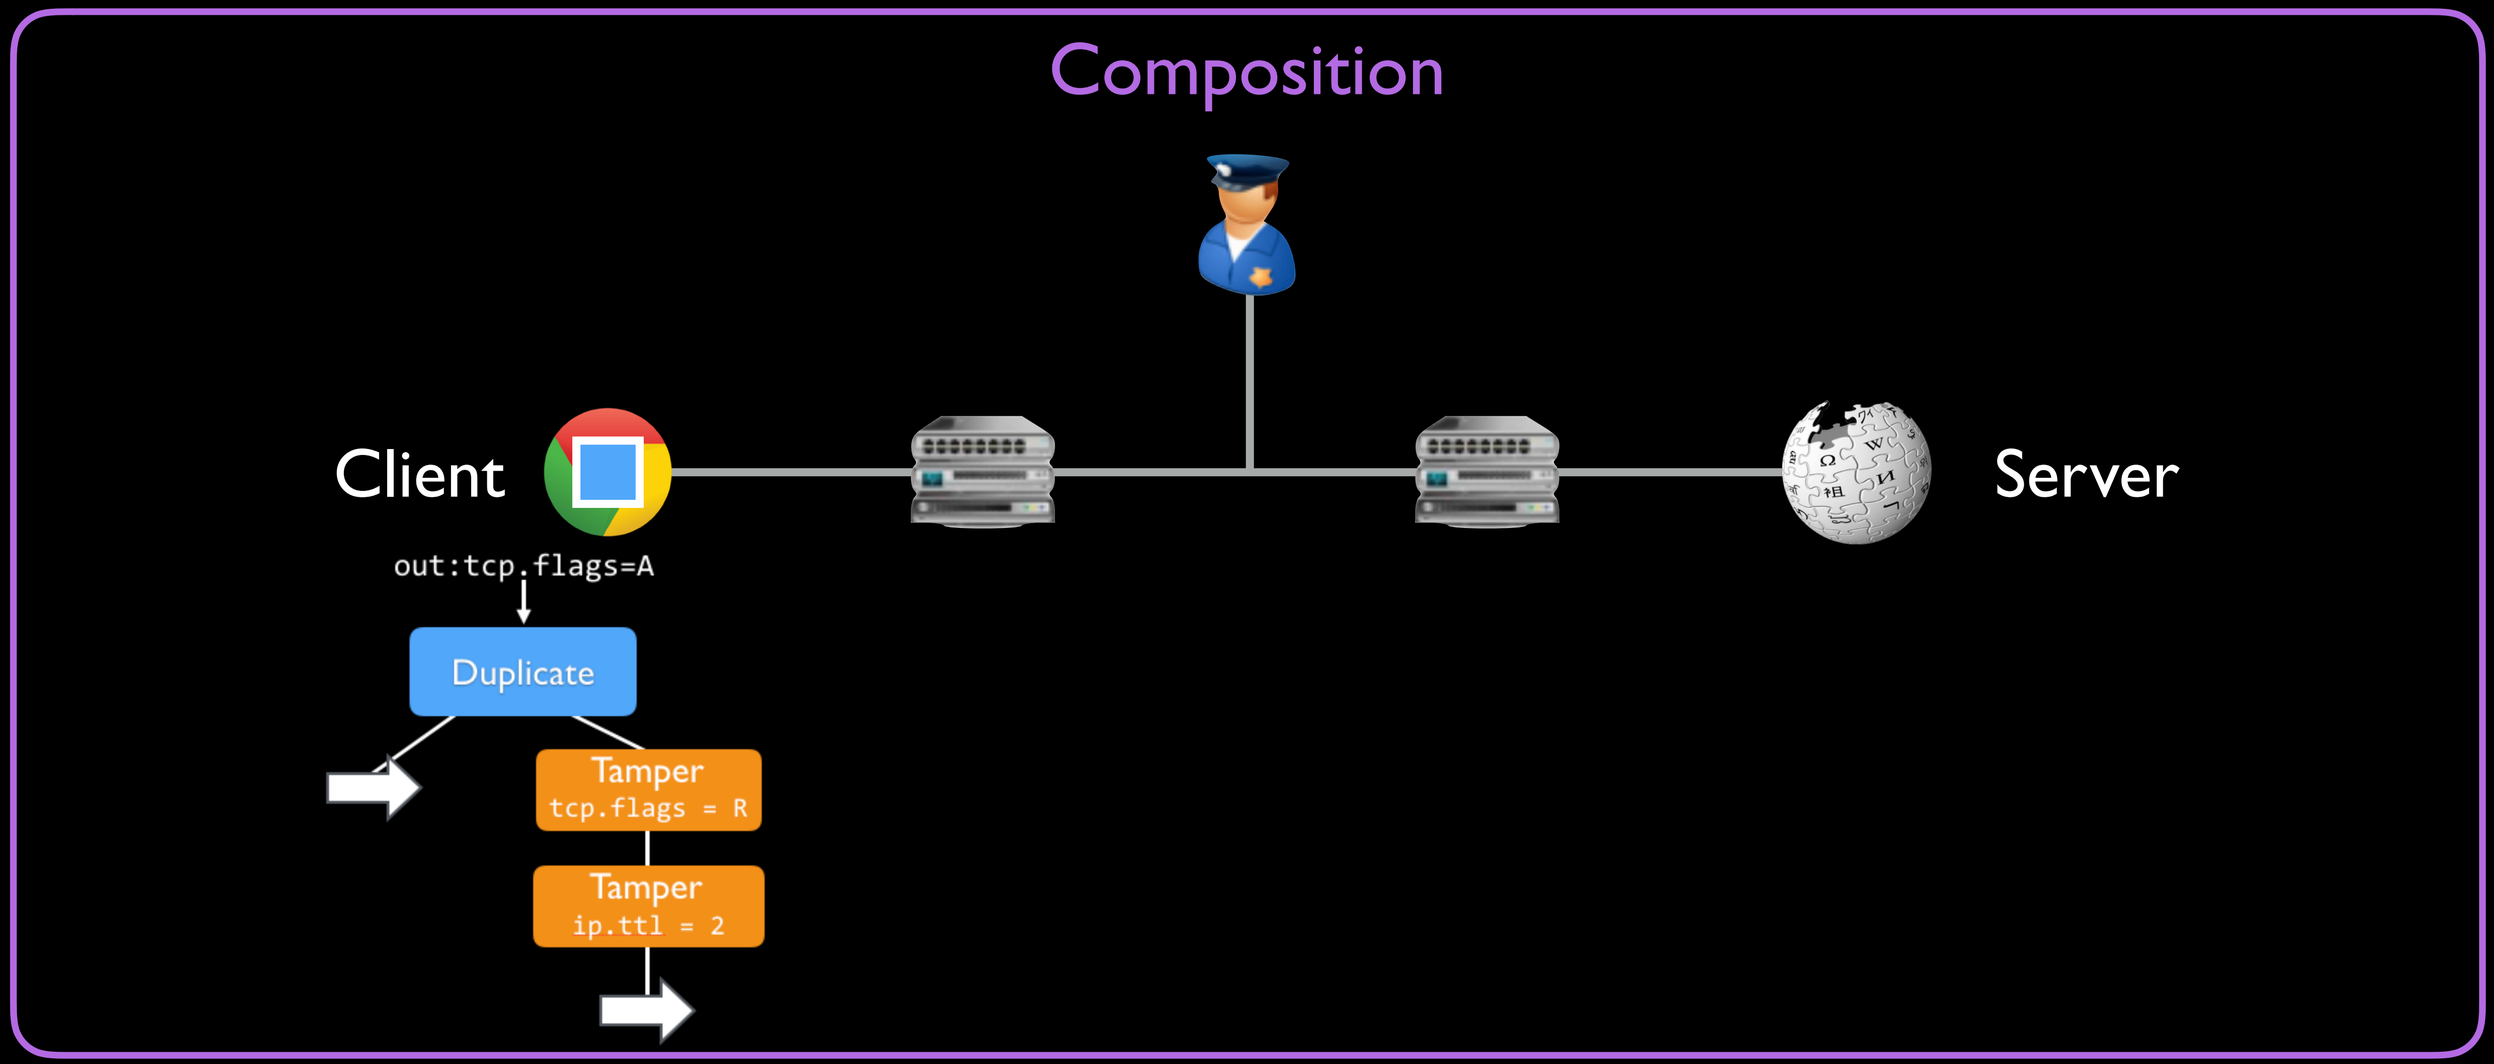

Running a Strategy

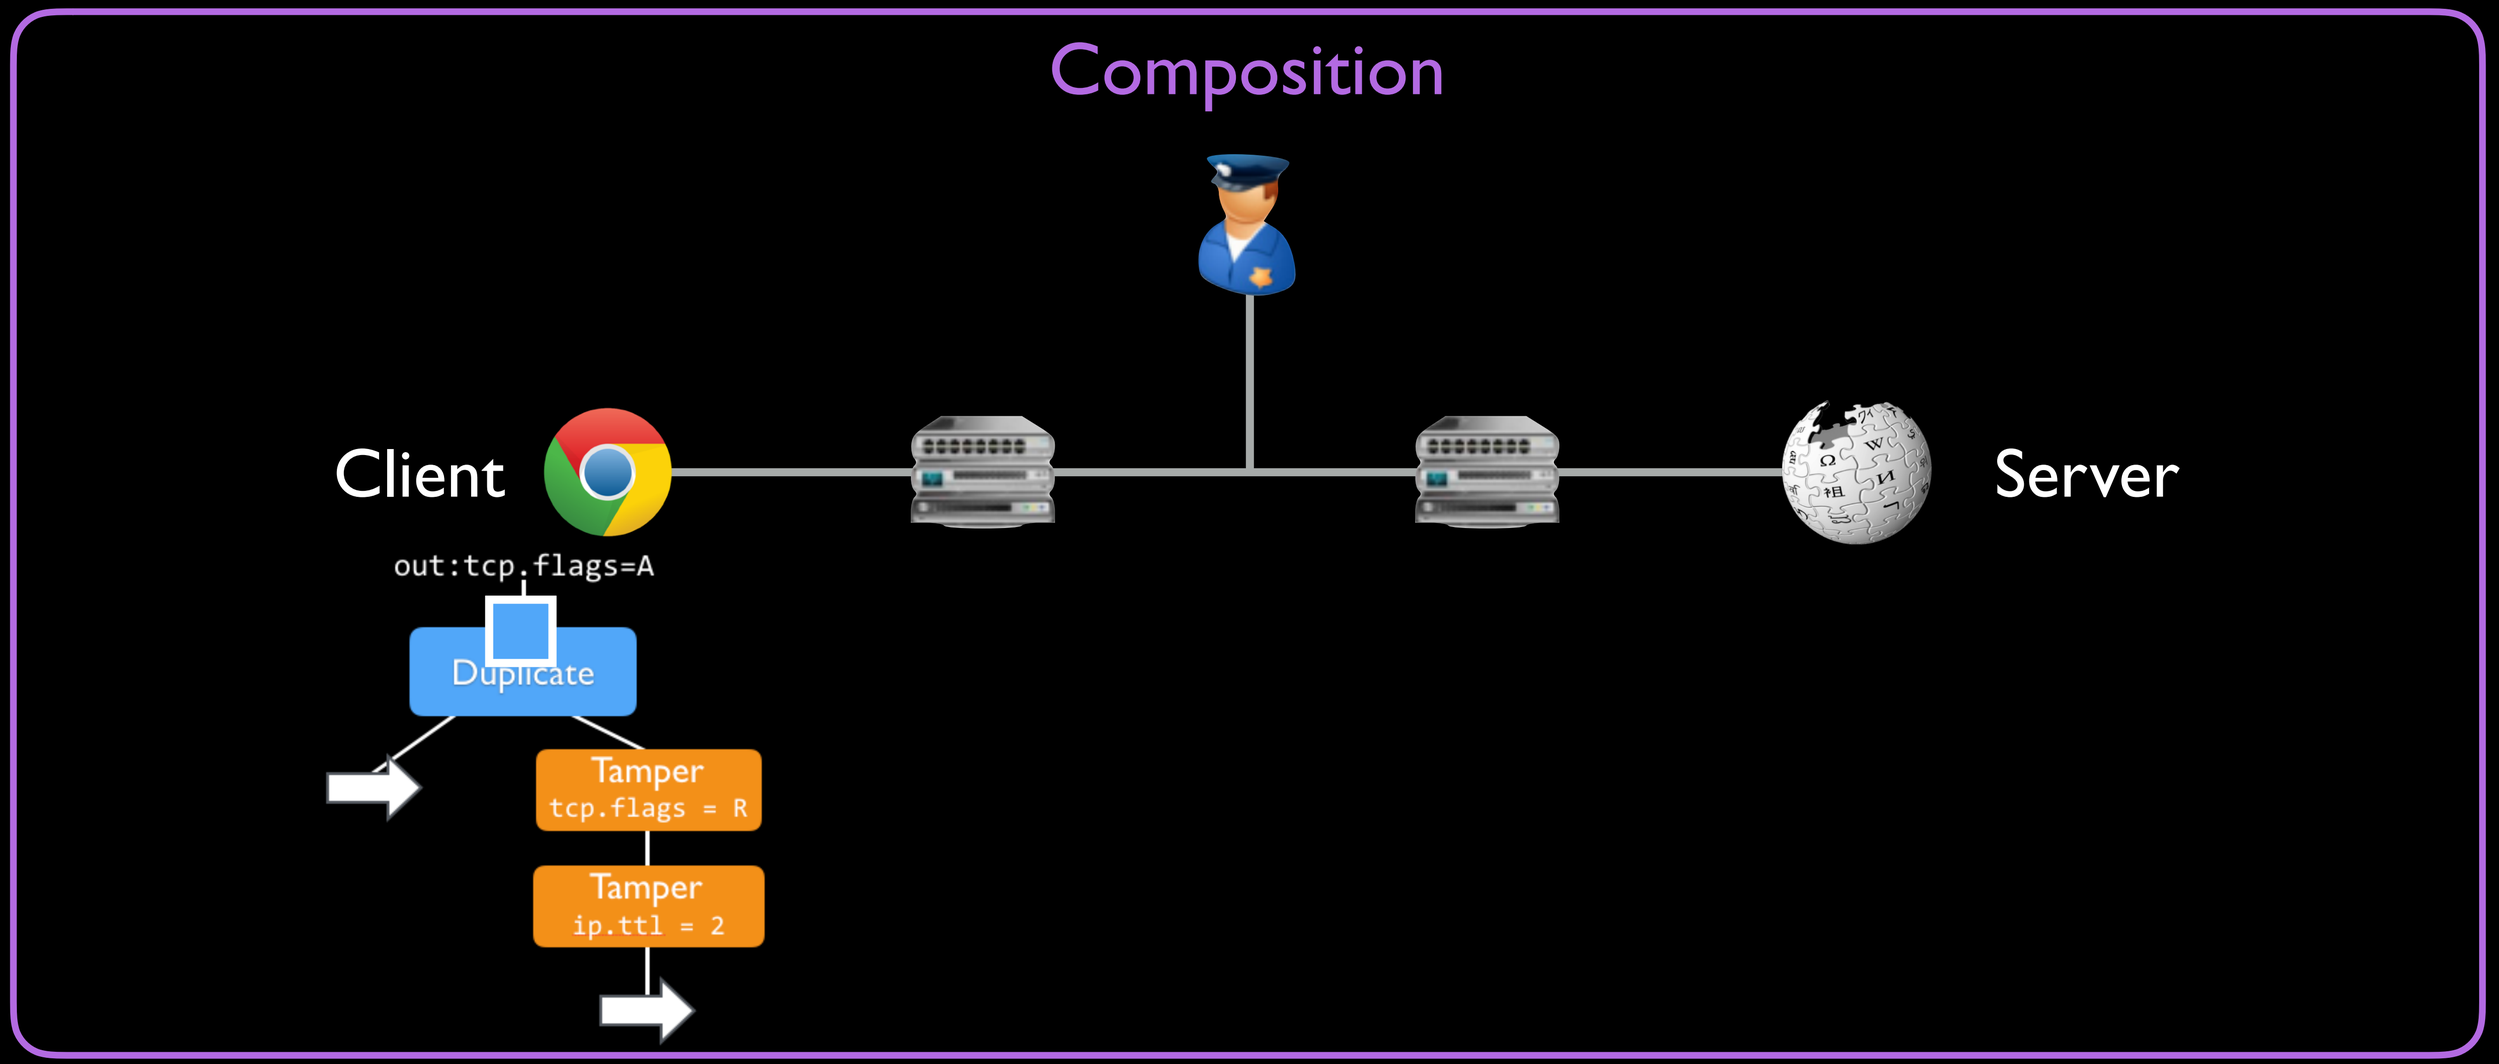

Running a Strategy

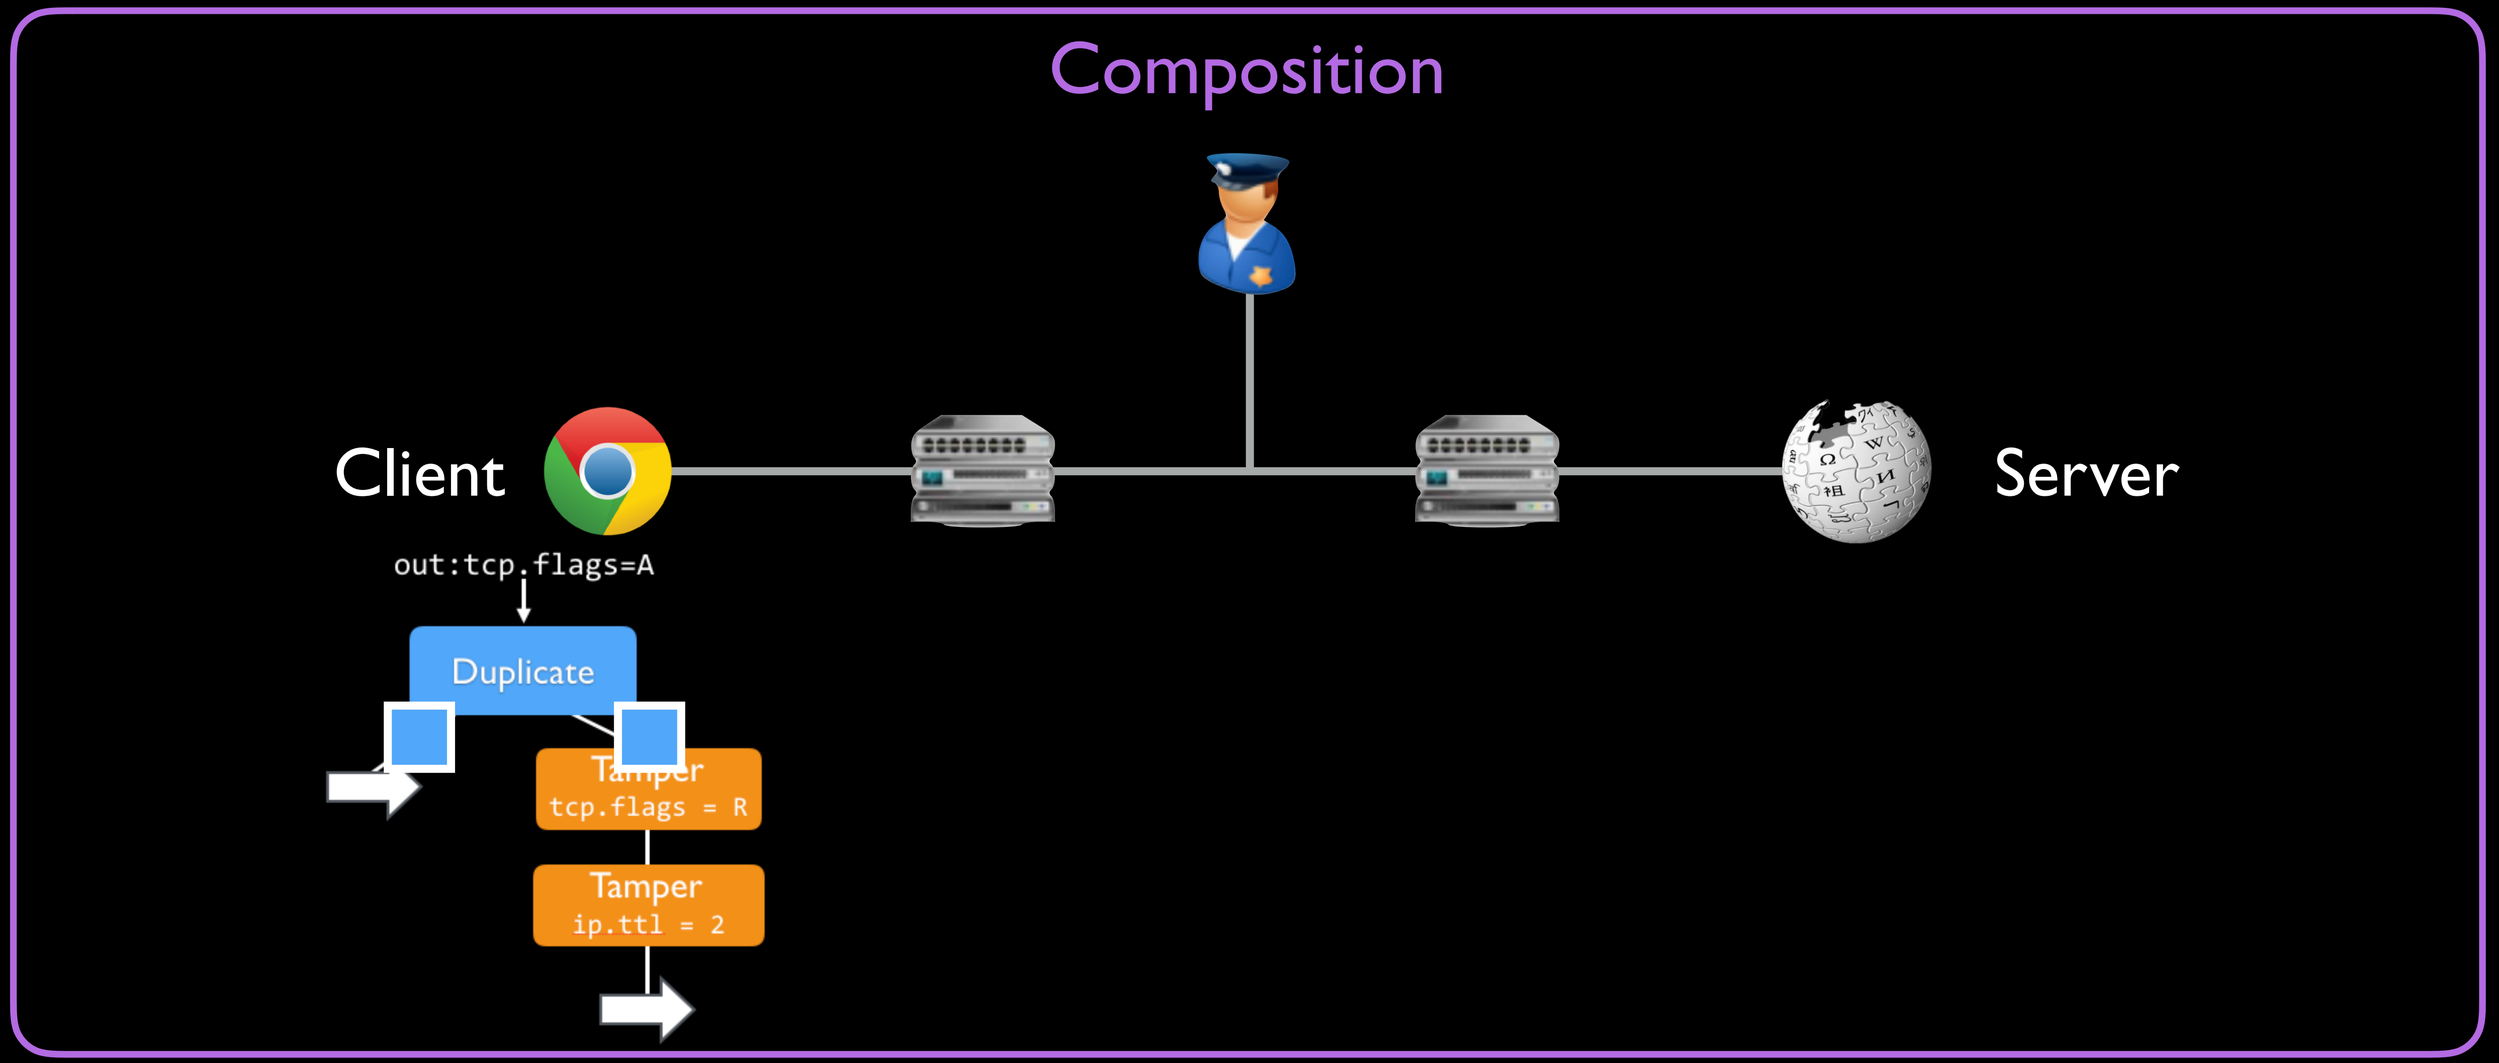

Running a Strategy

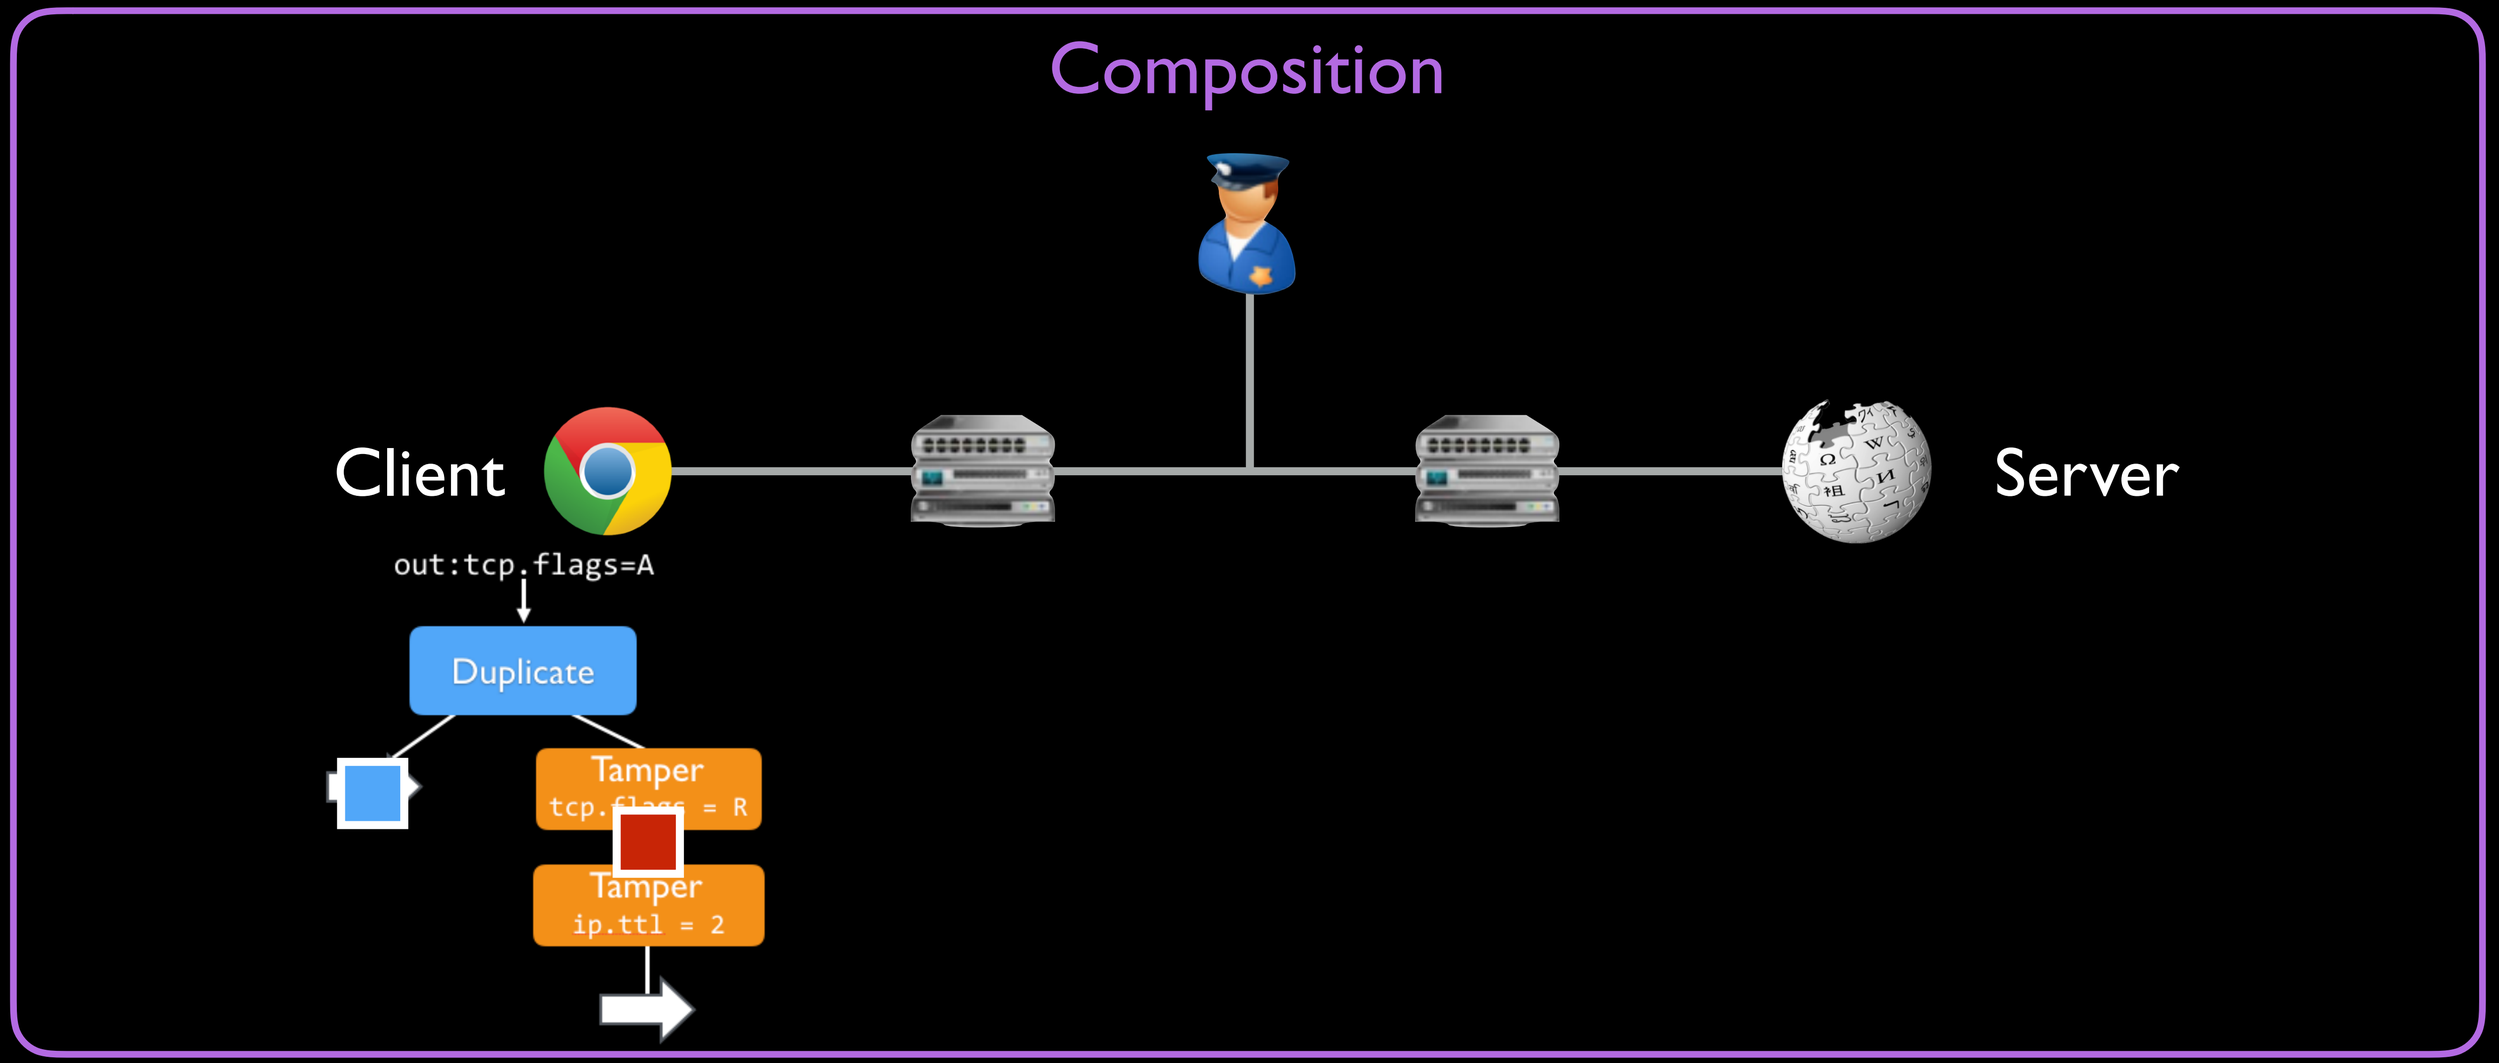

Running a Strategy

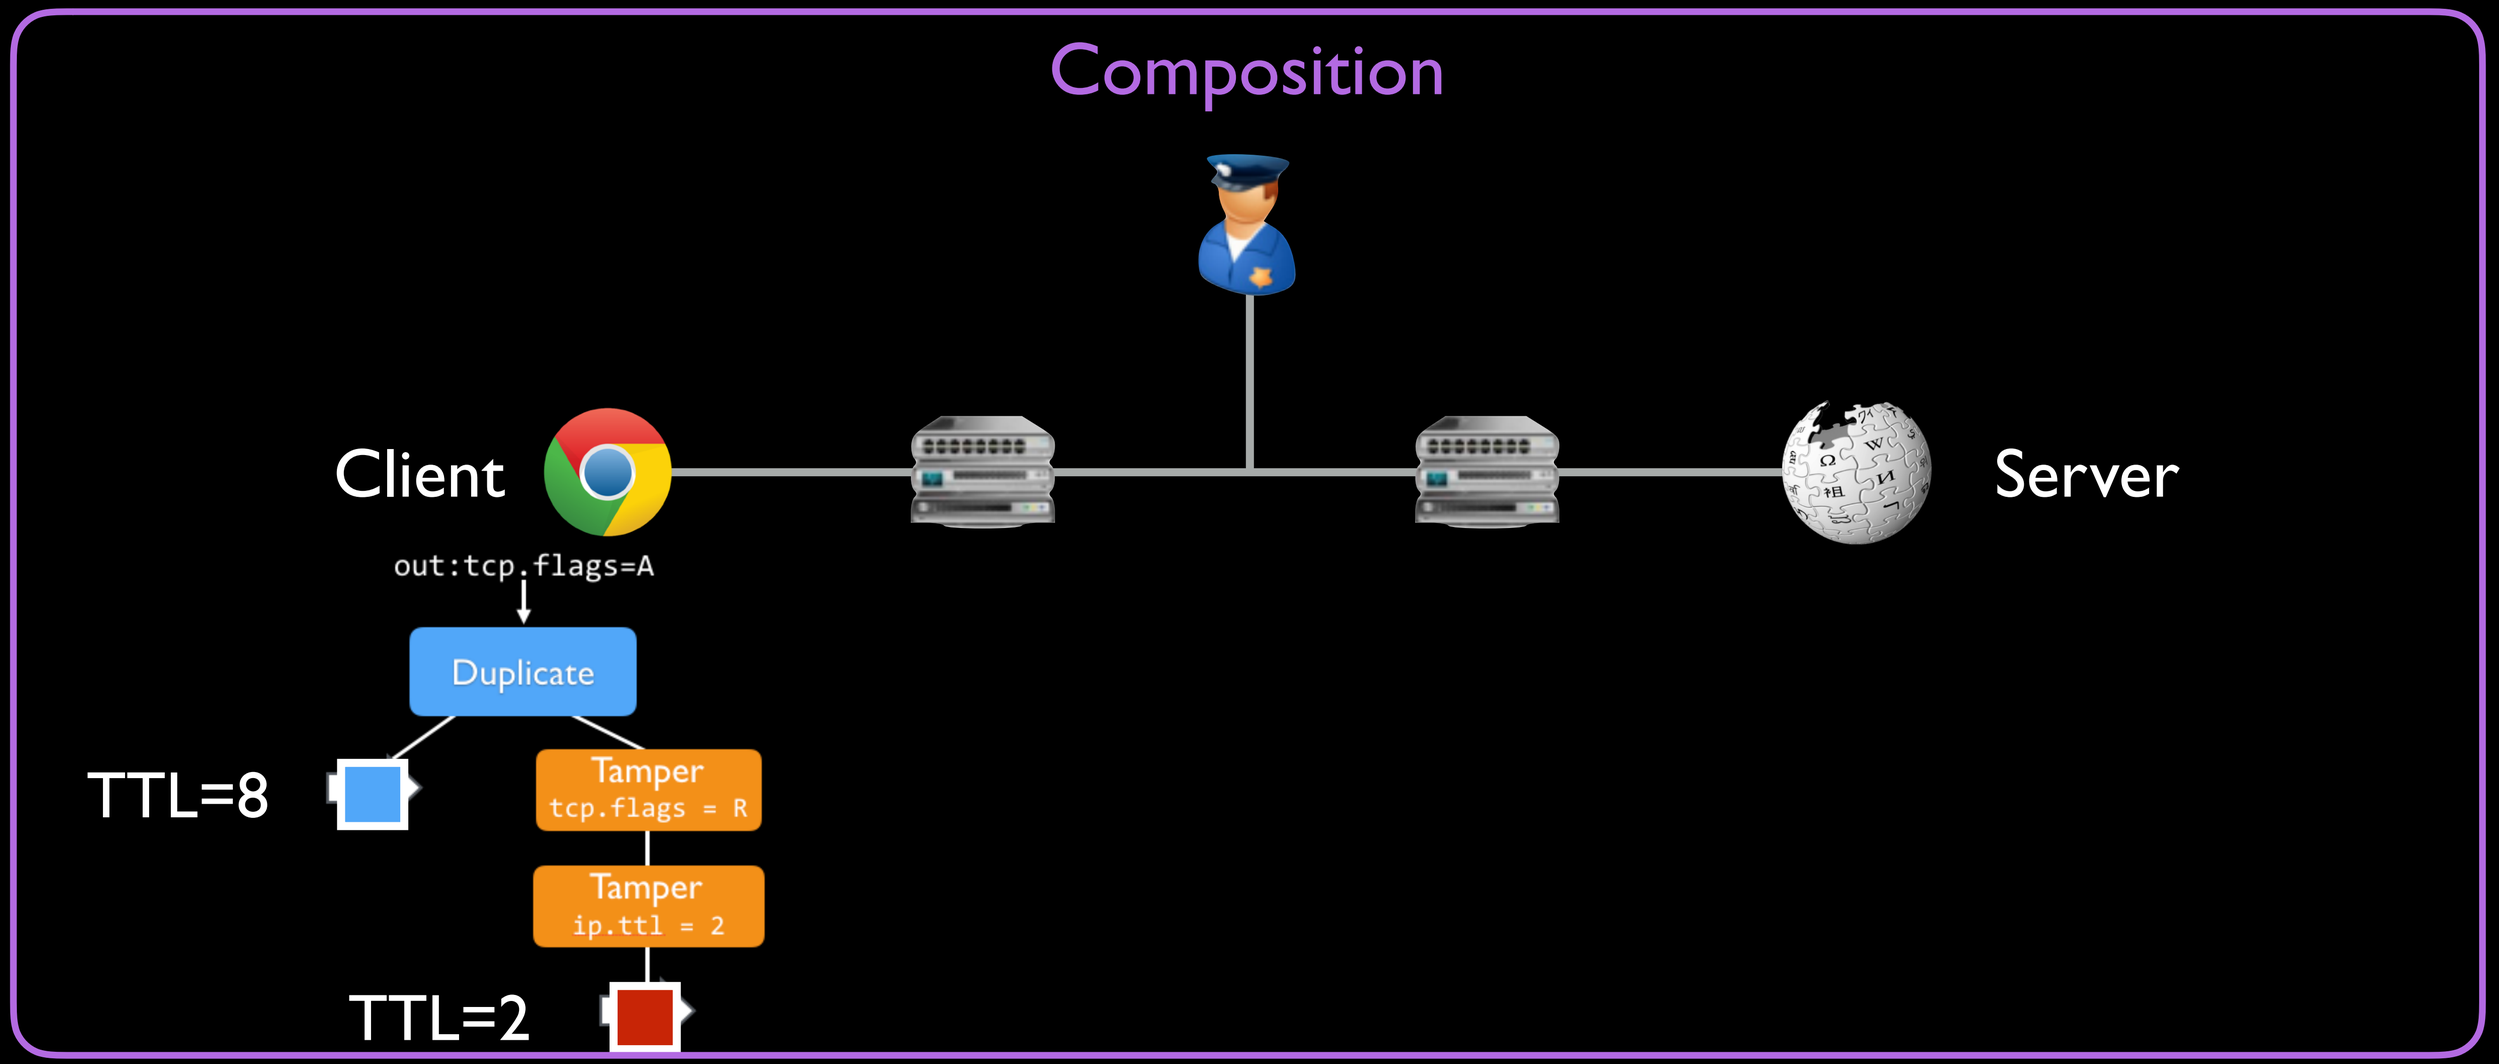

Running a Strategy

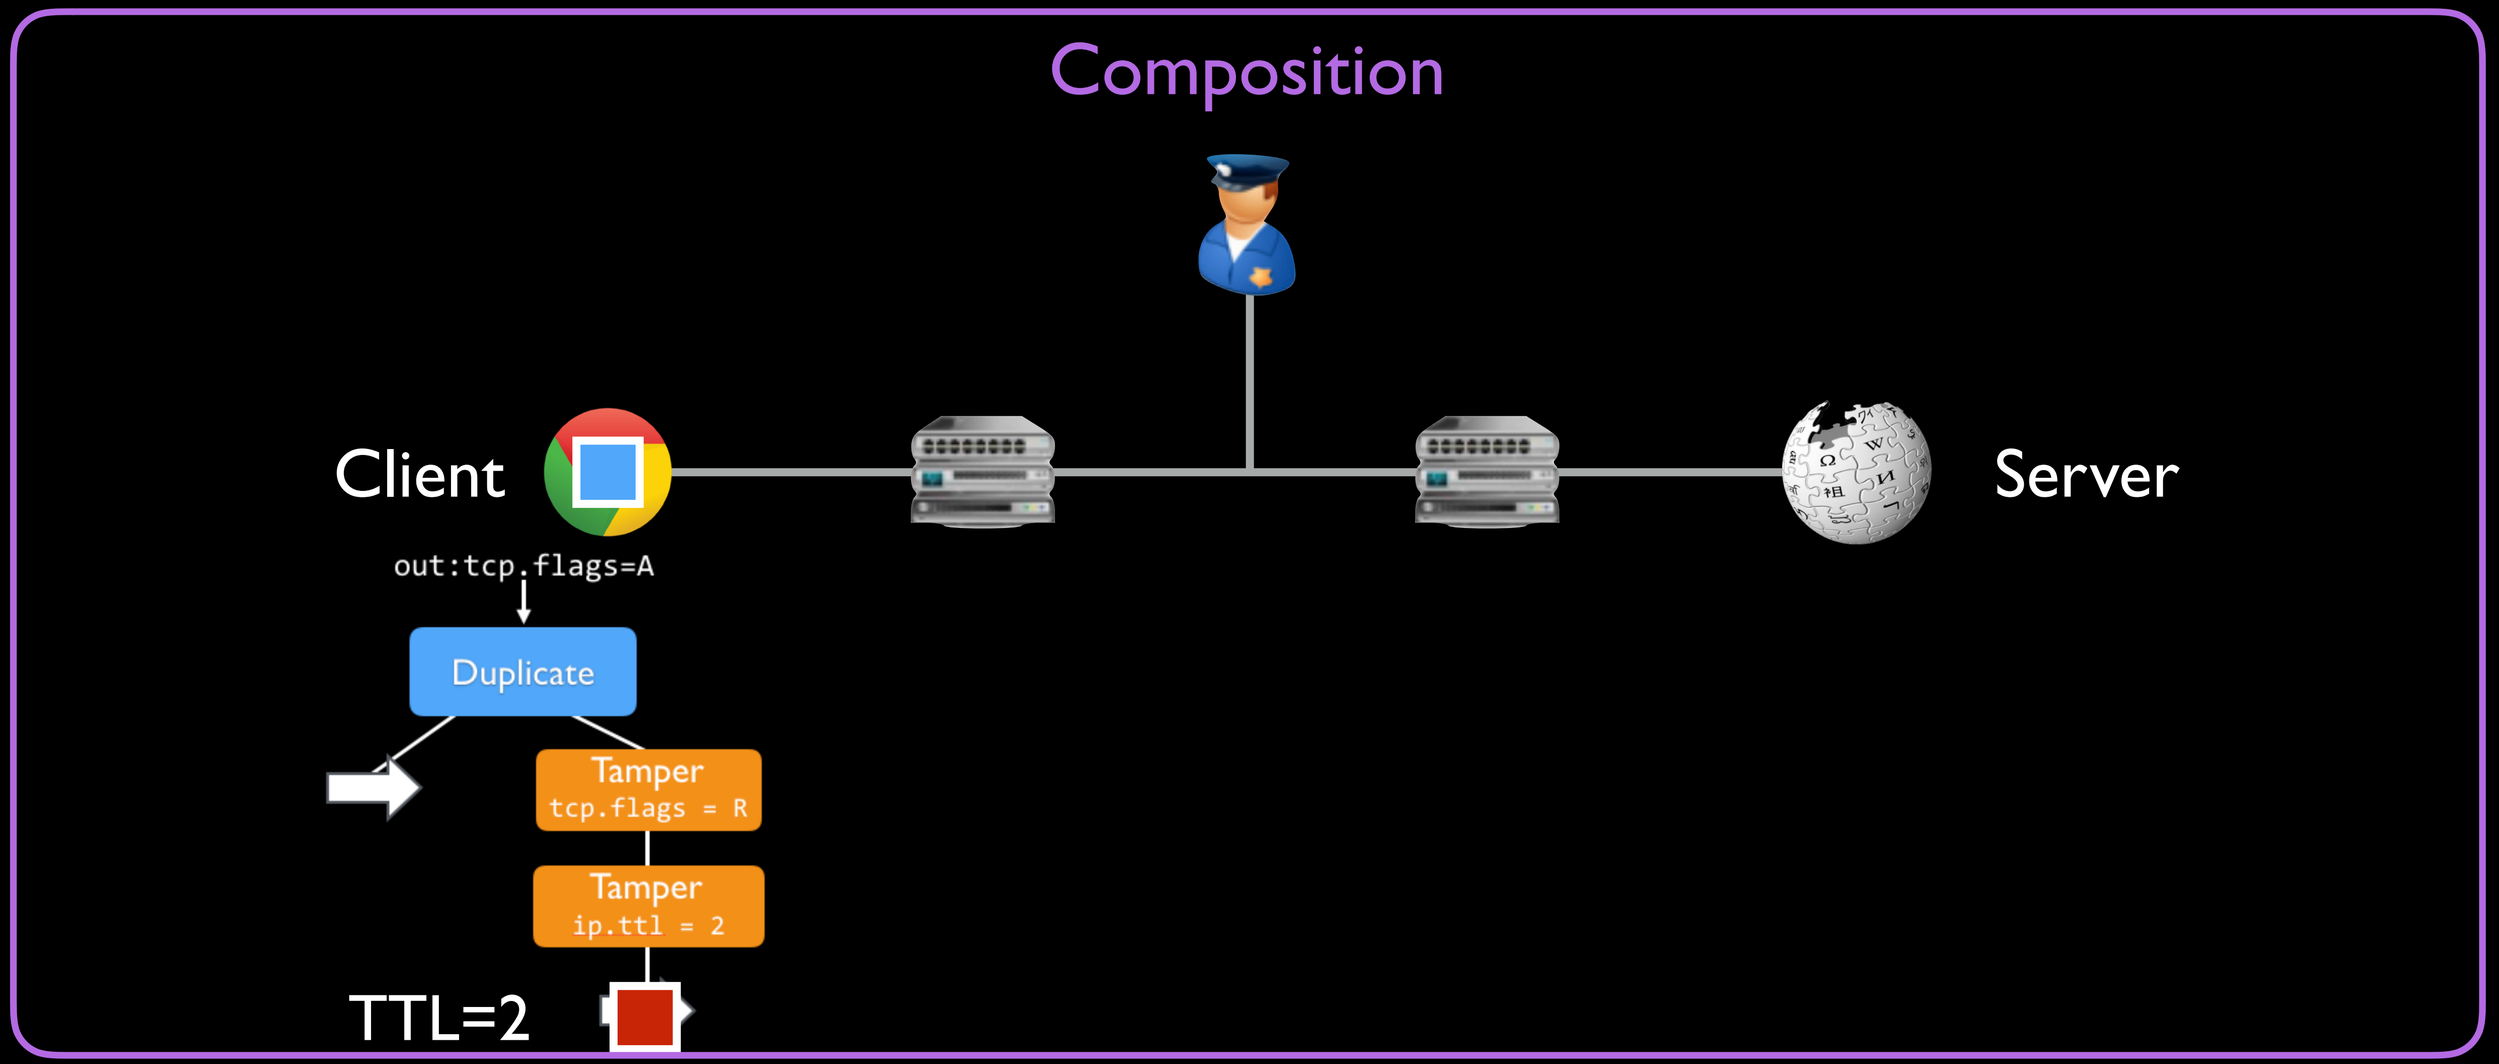

Running a Strategy

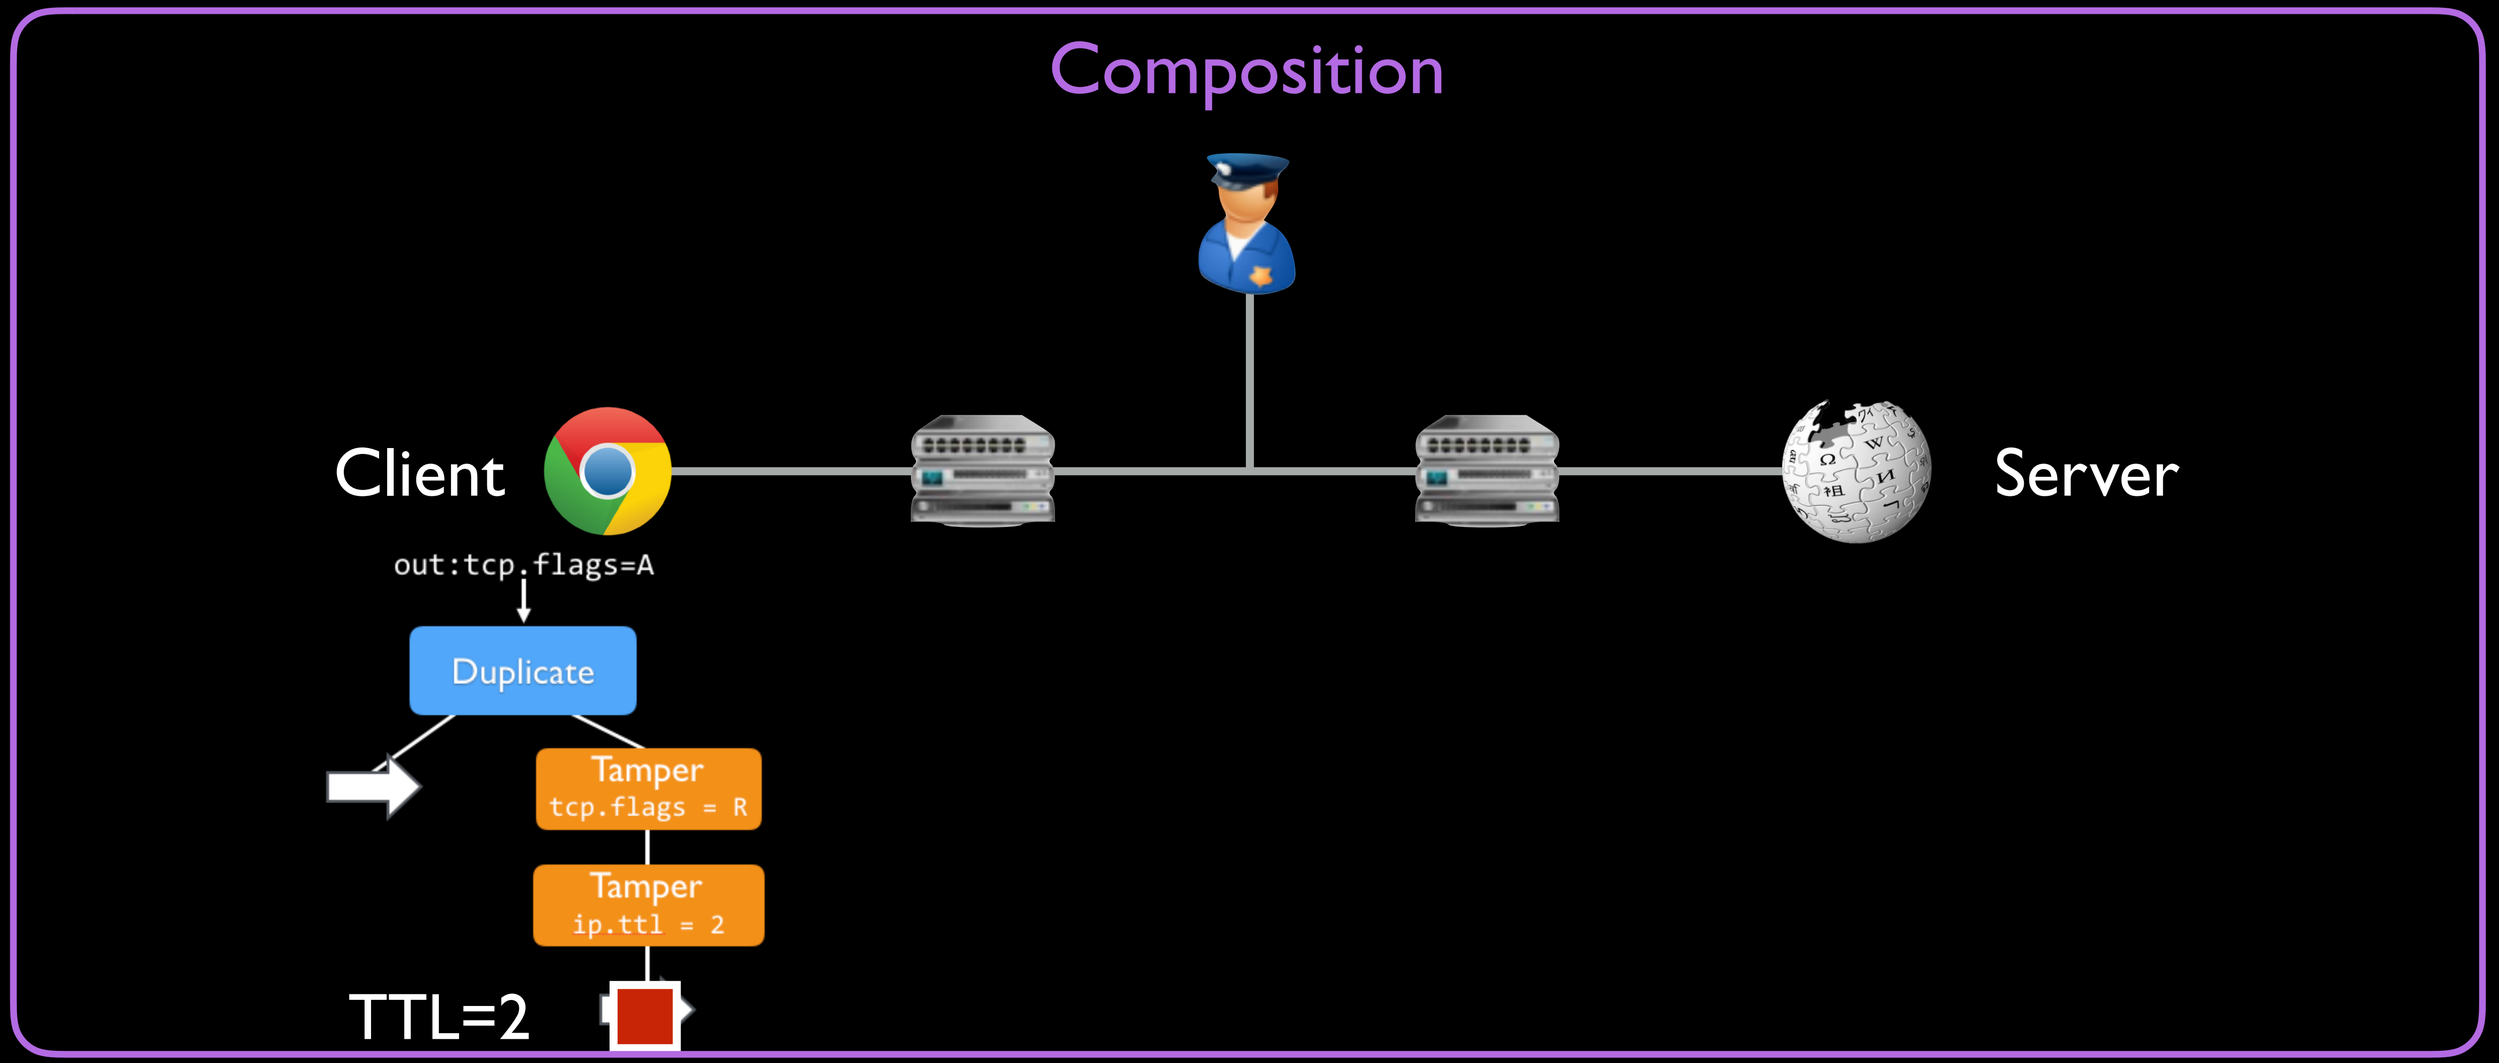

Running a Strategy

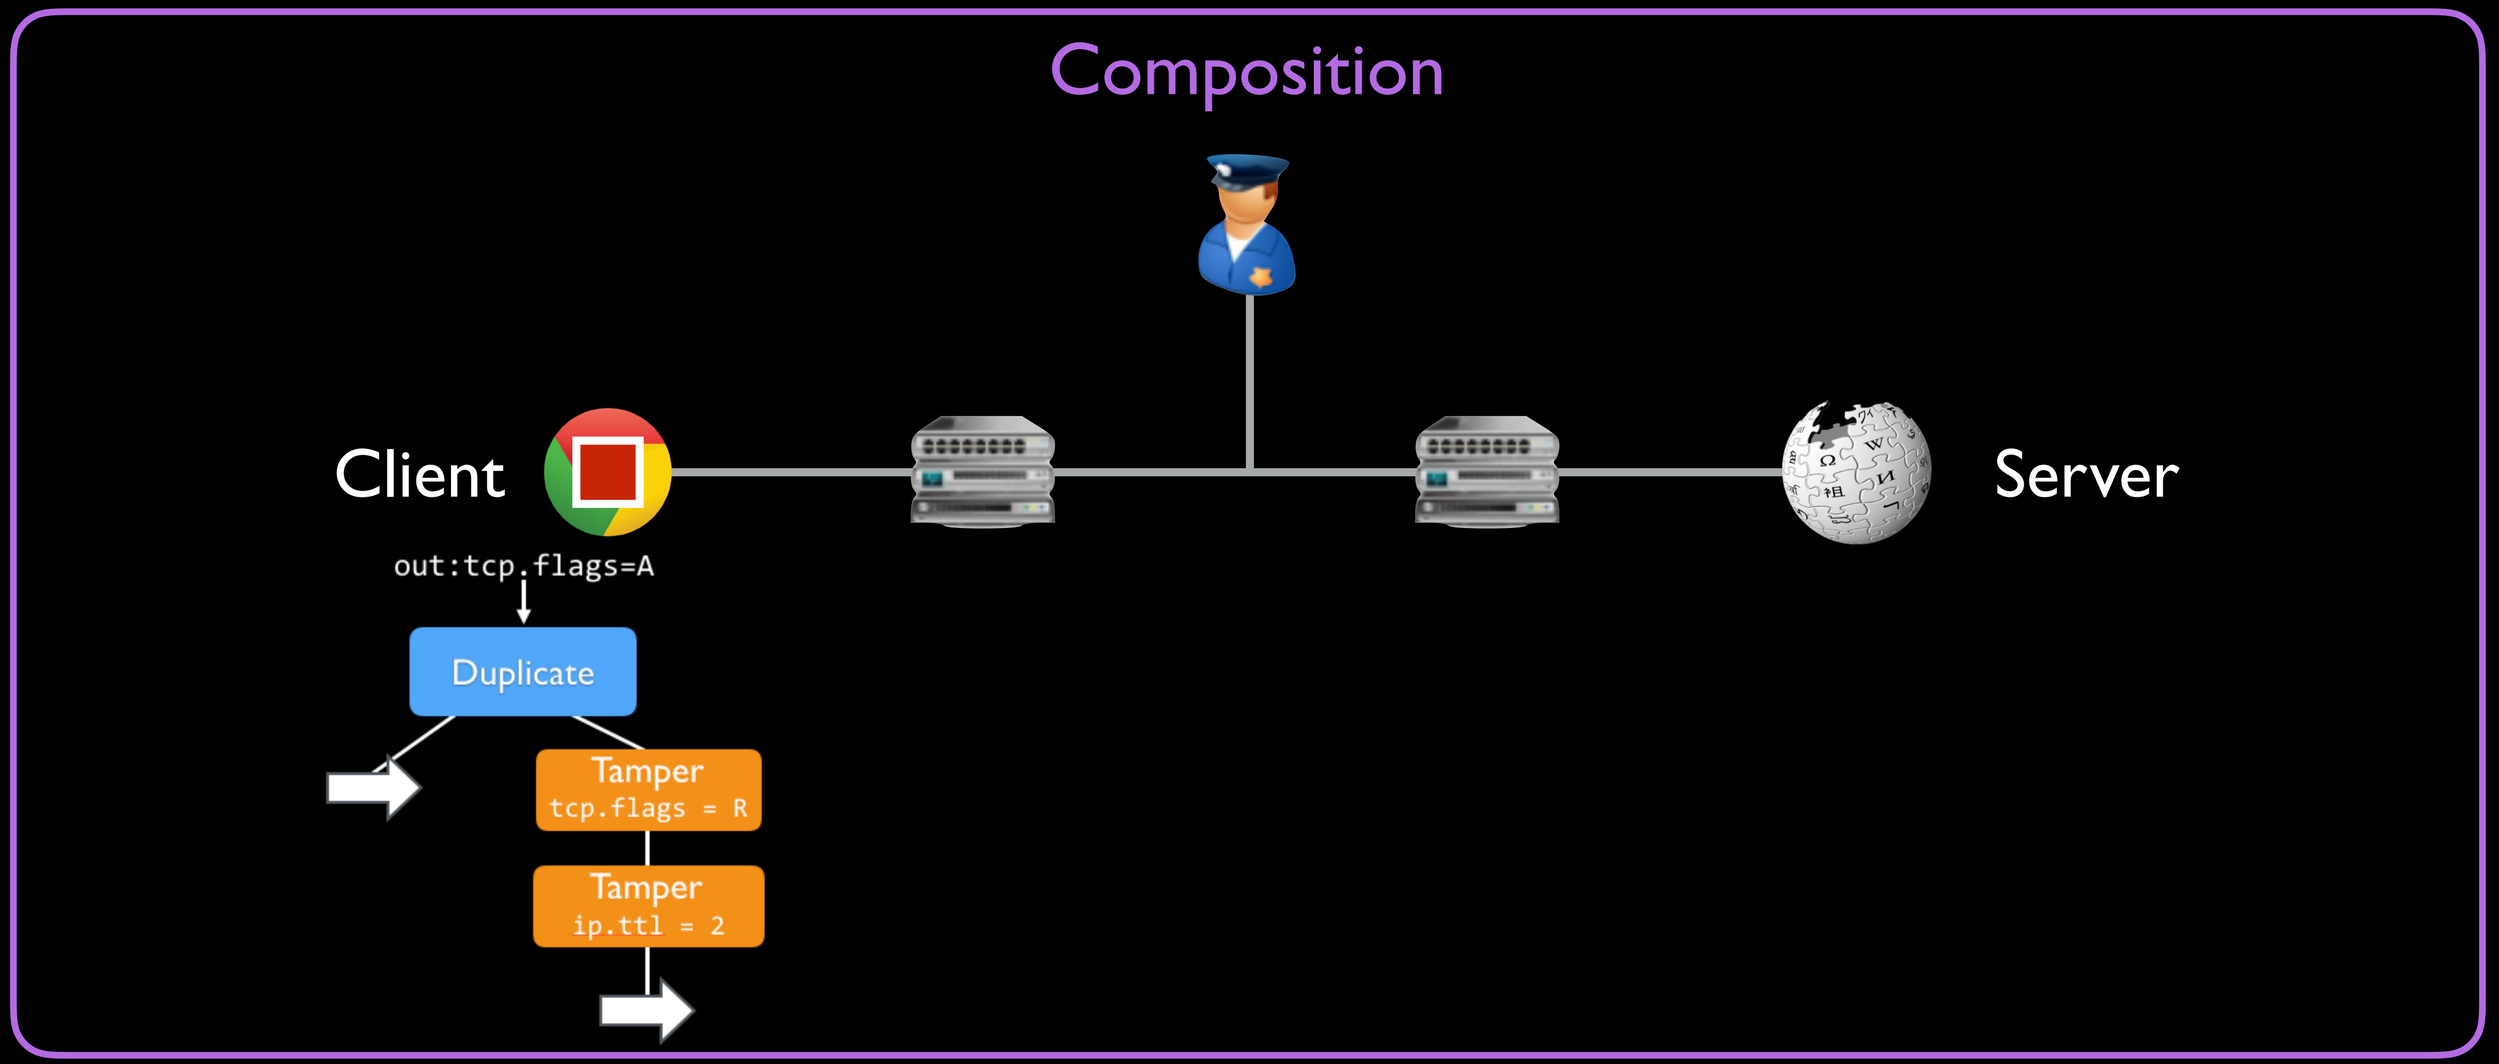

Running a Strategy

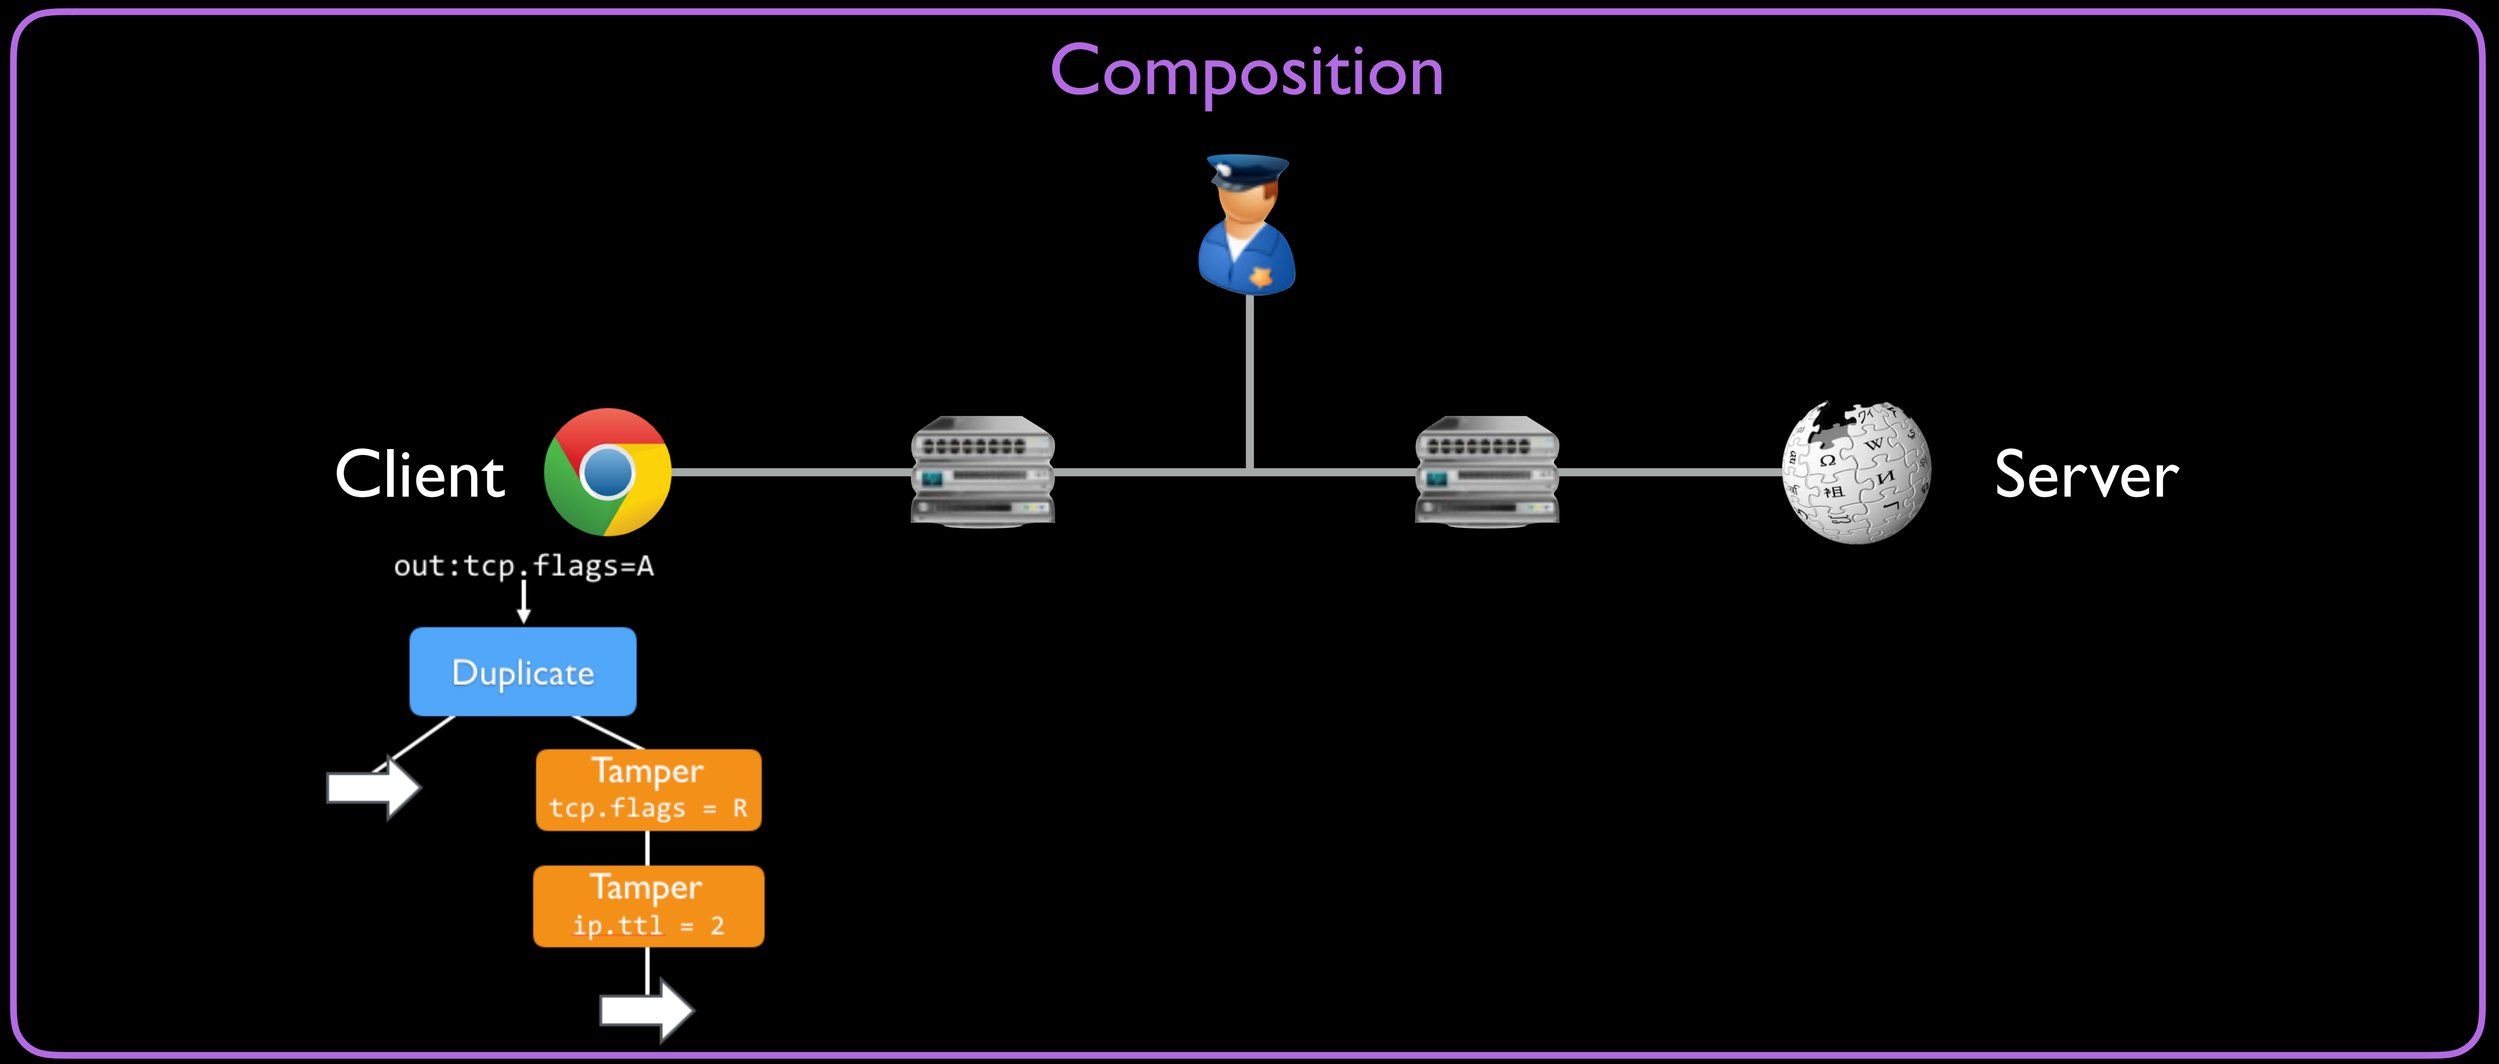

Geneva

Genetic Evasion

Building Blocks

Actions manipulate individual packets

- Duplicate
- Tamper
- Fragment
- Drop

Composition

Actions compose to form trees

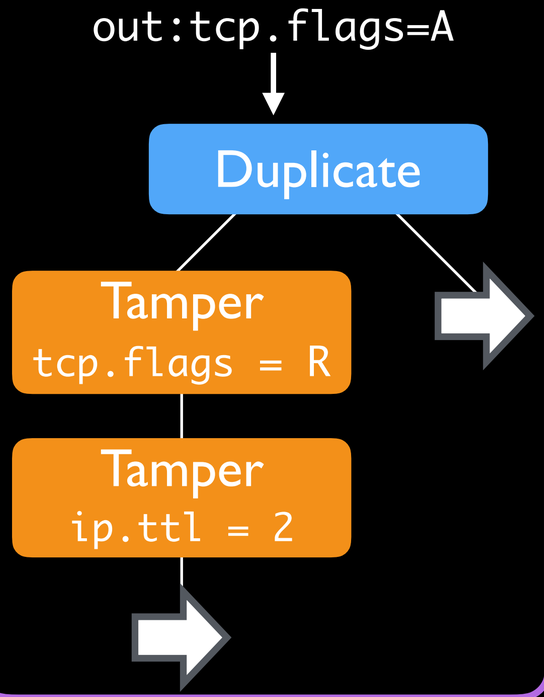

Mutation

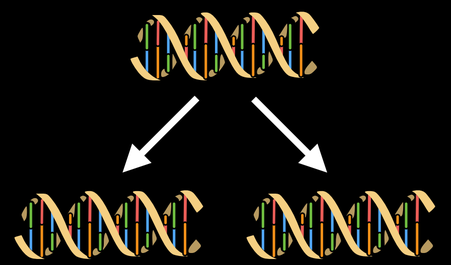

Fitness

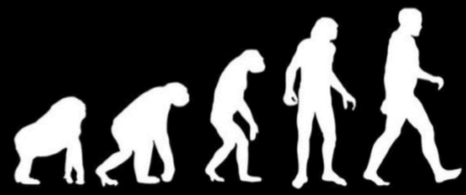

Geneva

Genetic Evasion

Building Blocks

Actions manipulate individual packets

- Duplicate
- Tamper
- Fragment
- Drop

Composition

Actions compose to form trees

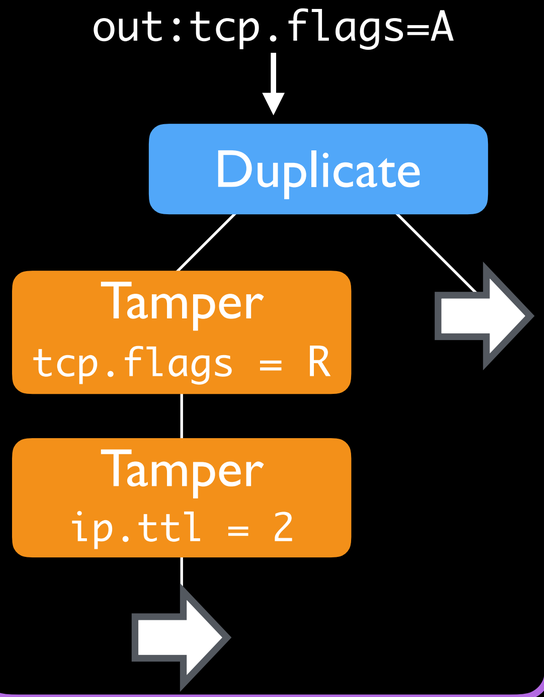

Mutation

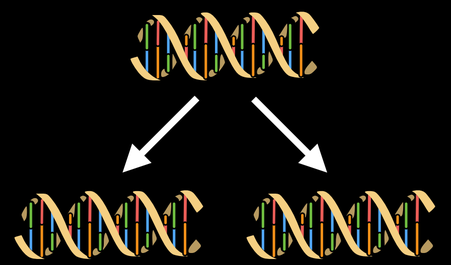

Fitness

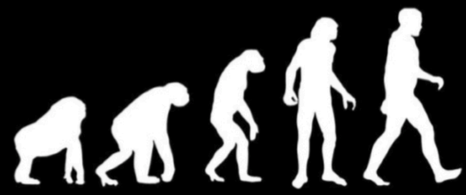

Geneva

Genetic Evasion

Building Blocks

Actions manipulate individual packets

- Duplicate
- Tamper
- Fragment
- Drop

Composition

Actions compose to form trees

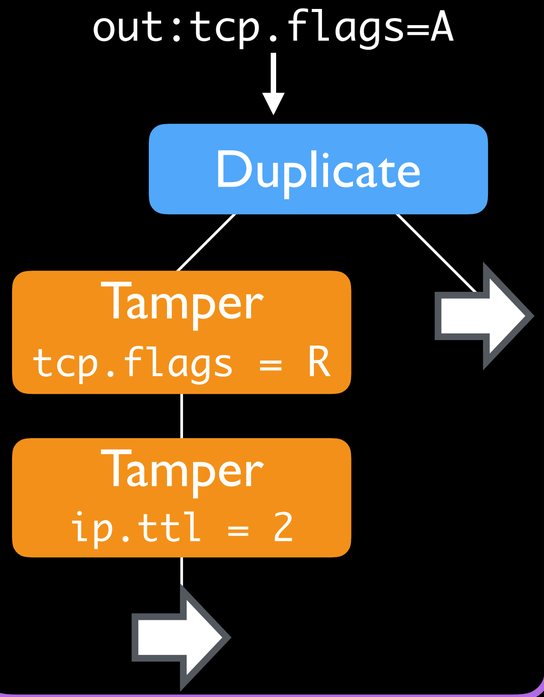

Mutation

Randomly alter types, values, and trees

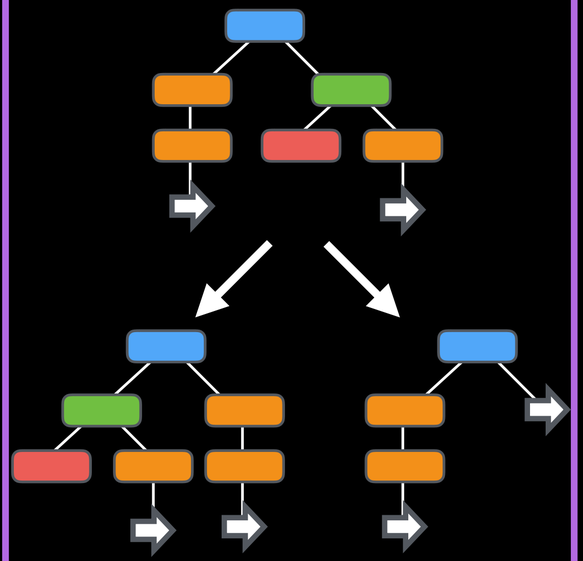

Fitness

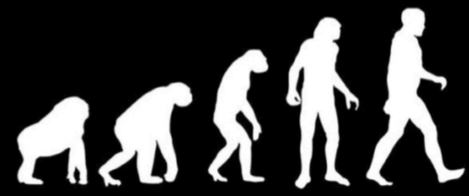

Geneva

Genetic Evasion

Building Blocks

Actions manipulate individual packets

- Duplicate
- Tamper
- Fragment
- Drop

Composition

Actions compose to form trees

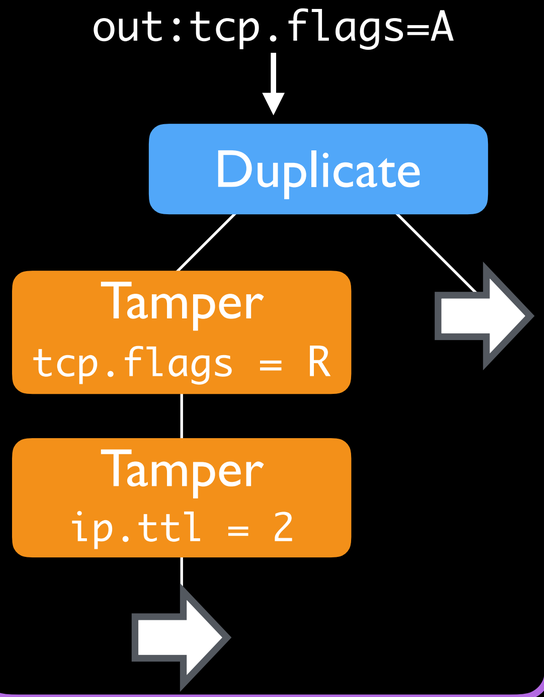

Mutation

Randomly alter types, values, and trees

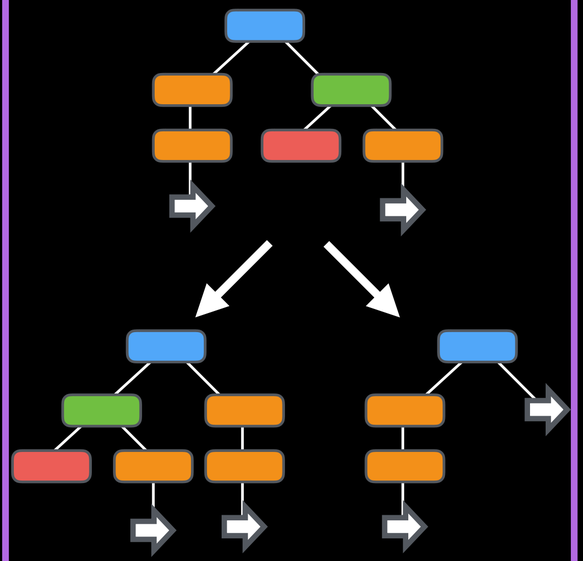

Fitness

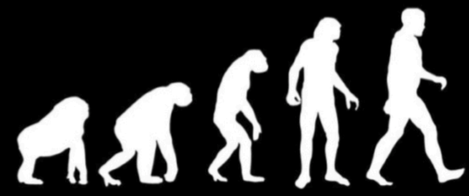

Geneva

Genetic Evasion

Fitness

Which individuals should survive to the next generation?

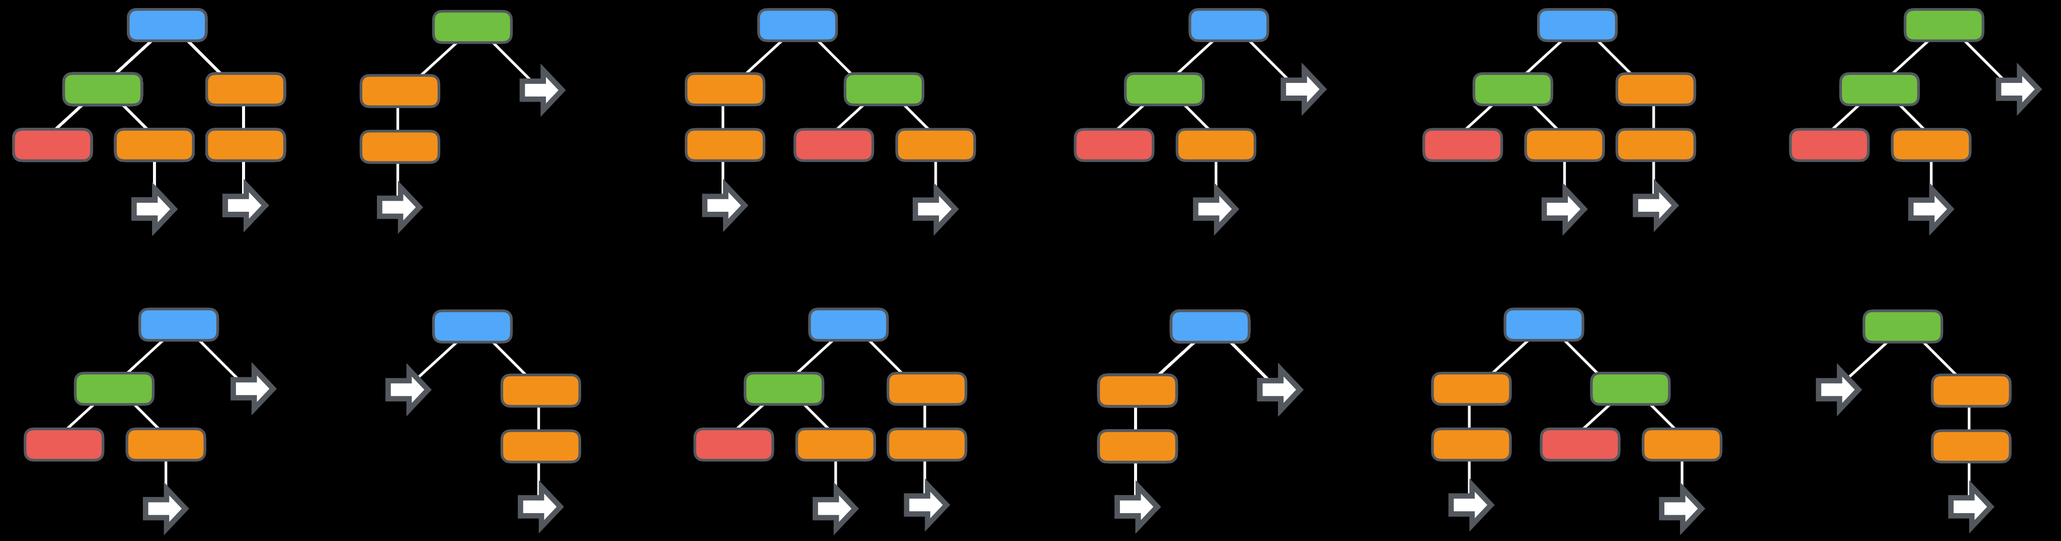

Geneva

Genetic Evasion

Fitness

Which individuals should survive to the next generation?

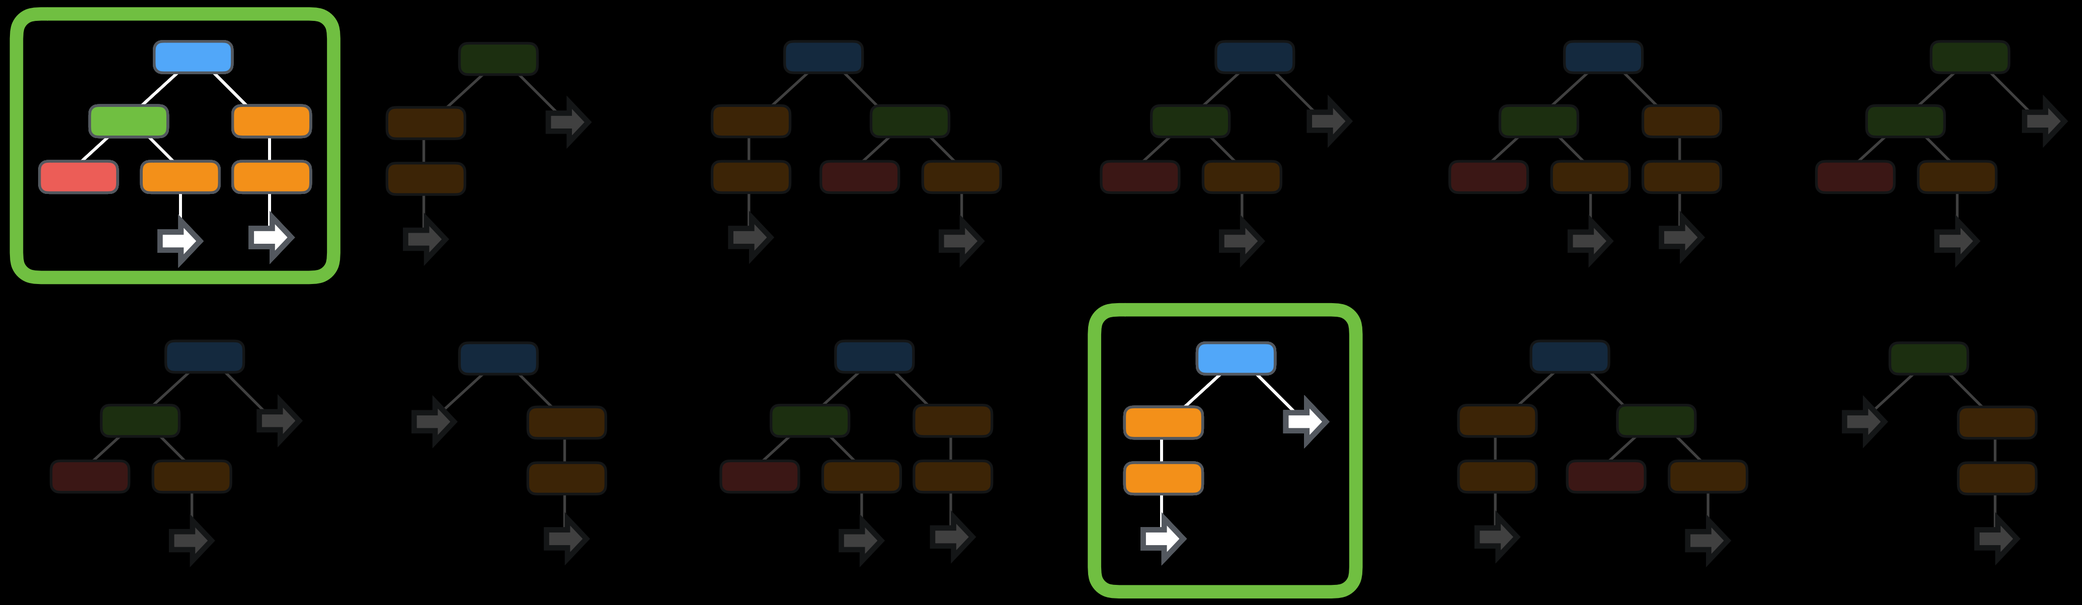

Geneva

Genetic Evasion

Fitness

Which individuals should survive to the next generation?

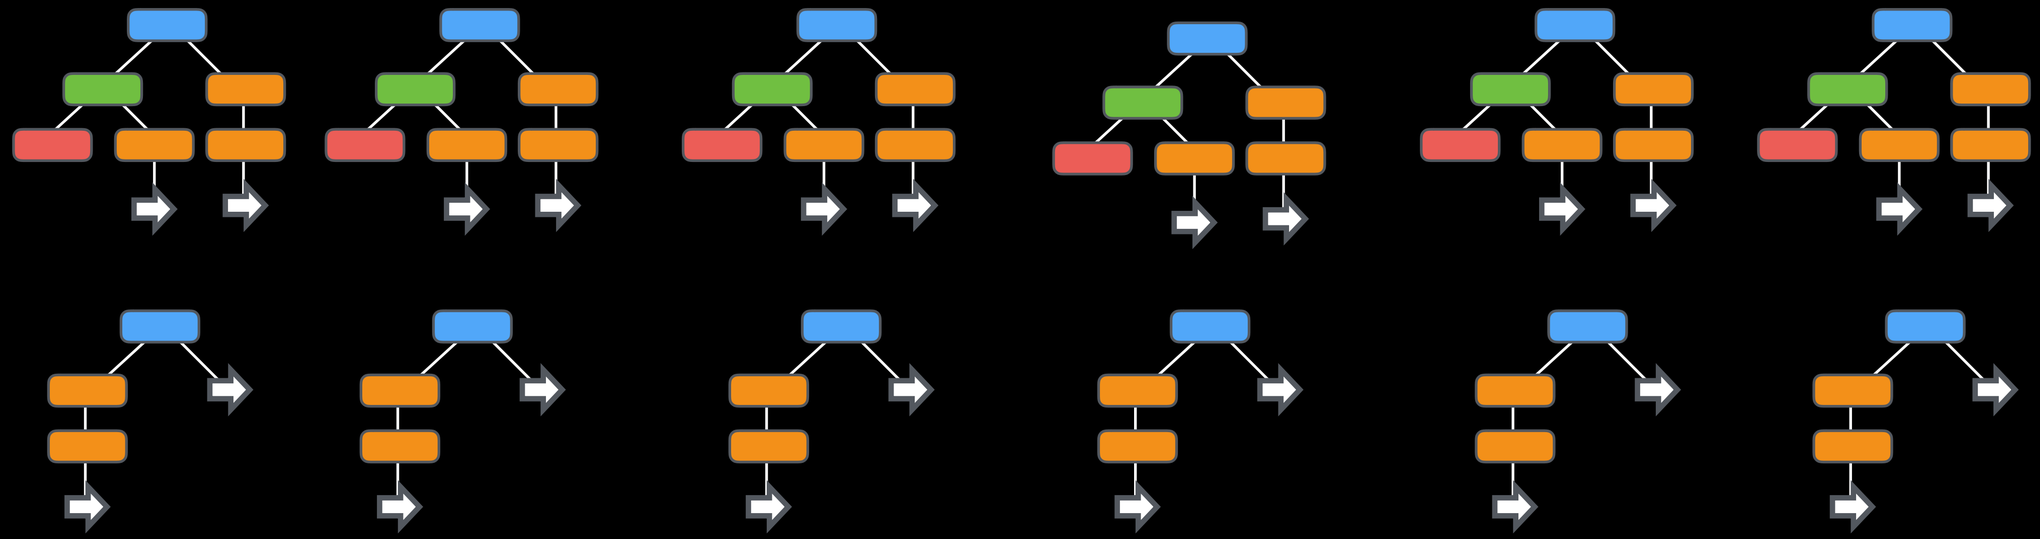

Geneva

Genetic Evasion

Fitness

Which individuals should survive to the next generation?

- Not triggering on any packets
- Breaking the TCP connection

+ Successfully obtaining forbidden content
+ Conciseness

Geneva

Genetic Evasion

Building Blocks

Actions manipulate individual packets

- Duplicate
- Tamper
- Fragment
- Drop

Composition

Actions compose to form trees

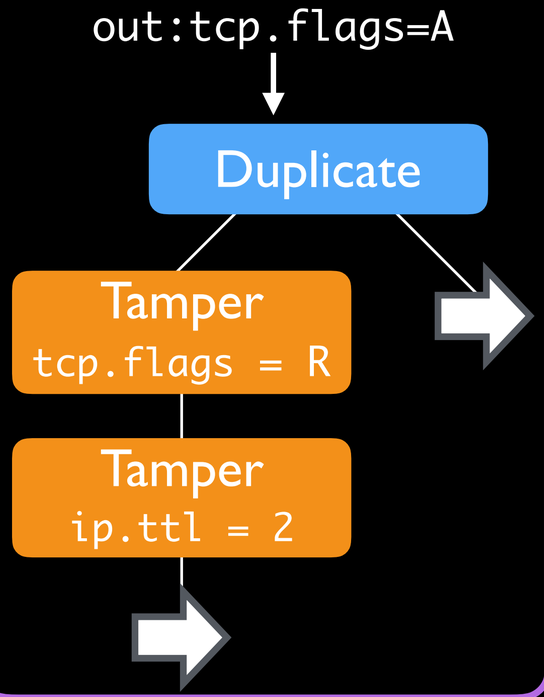

Mutation

Randomly alter types, values, and trees

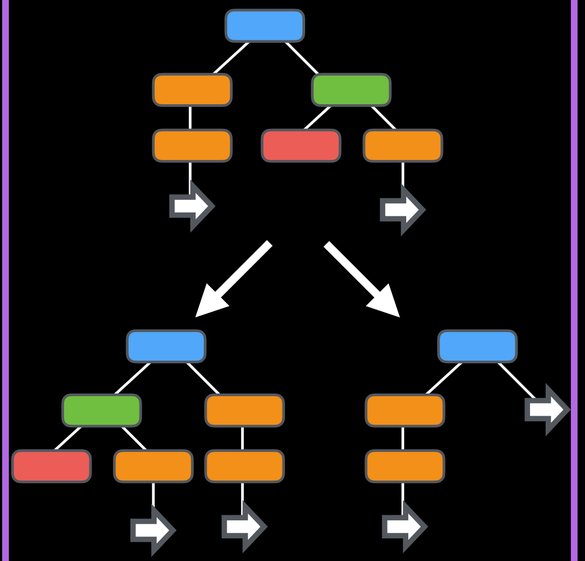

Fitness

Goal: Fewest actions needed to succeed

- No trigger
- Break TCP
- Successful
- Concise

Geneva's results

Real censor experiments

China India Kazakhstan

Geneva's results

Real censor experiments

China

India

Kazakhstan

Geneva's results

Real censor experiments

FTP
DNS
SMTP
HTTP
HTTPS

Injects TCP RSTs

China

India

Kazakhstan

Geneva's results

Real censor experiments

FTP
DNS
SMTP
HTTP
HTTPS

Injects TCP RSTs

China

HTTP

Injects a block page

India

Kazakhstan

Geneva's results

Real censor experiments

FTP
DNS
SMTP
HTTP
HTTPS

Injects TCP RSTs

China

HTTP

Injects a block page

India

HTTP
HTTPS

Injects & blackholes

Kazakhstan

Geneva's results

Real censor experiments

China

India

Kazakhstan

Geneva's results

Real censor experiments

6 Species

13 Sub-species

36 Variants

China

India

Kazakhstan

Geneva's results

Real censor experiments

6 Species ..... The underlying bug

13 Sub-species ..... How Geneva exploits it

36 Variants ..... Functionally distinct

China

India

Kazakhstan

Geneva's results

Real censor experiments

6 Species — The underlying bug

13 Sub-species — How Geneva exploits it

36 Variants — Functionally distinct

31 | 6 | 13

China

India

Kazakhstan

Turnaround species

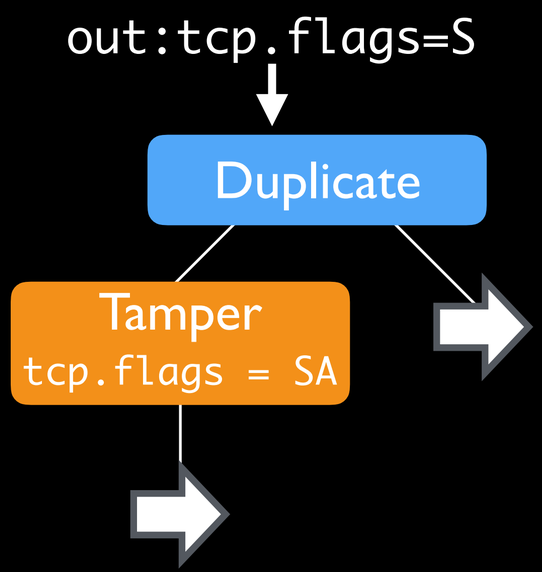

Trick the censor into thinking the client is the server

Turnaround species

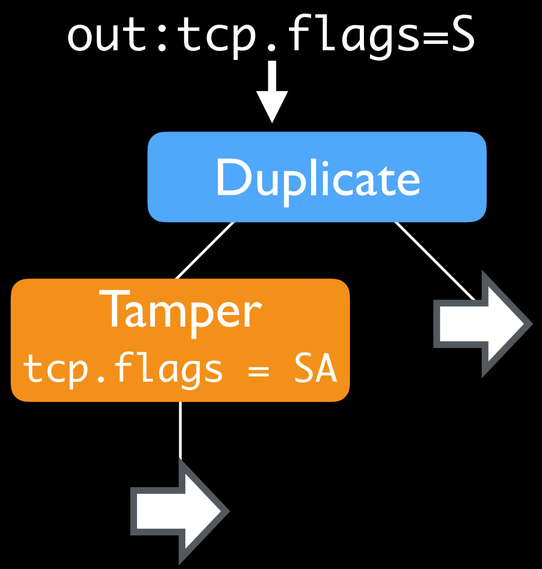

Trick the censor into thinking the client is the server

Segmentation species

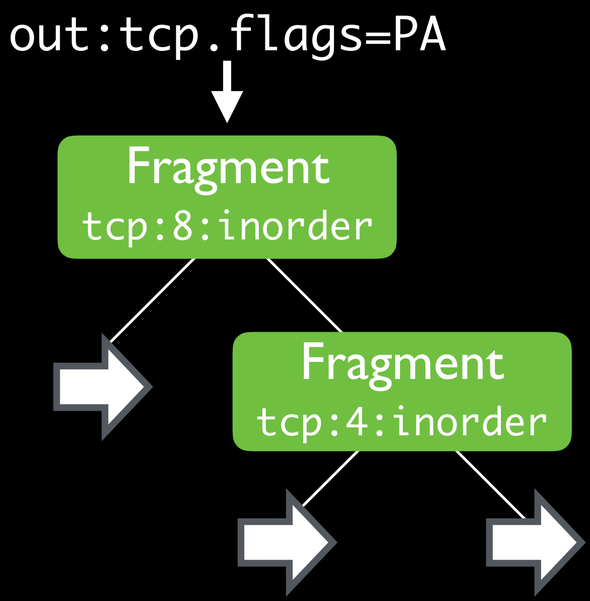

Segment the request

Turnaround species

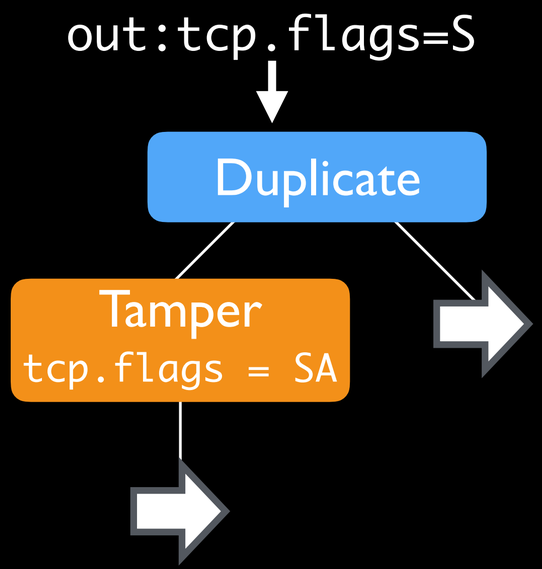

Trick the censor into thinking the client is the server

Segmentation species

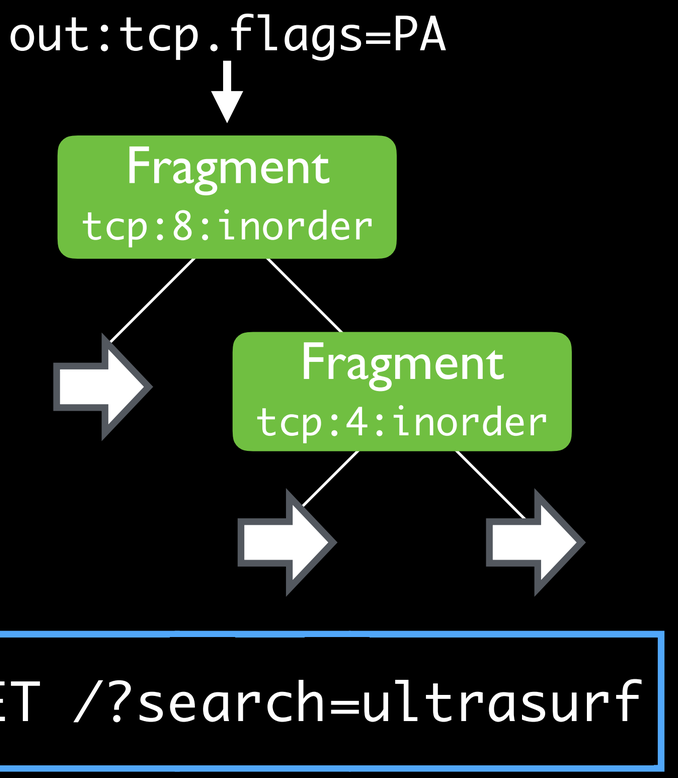

Segment the request

Turnaround species

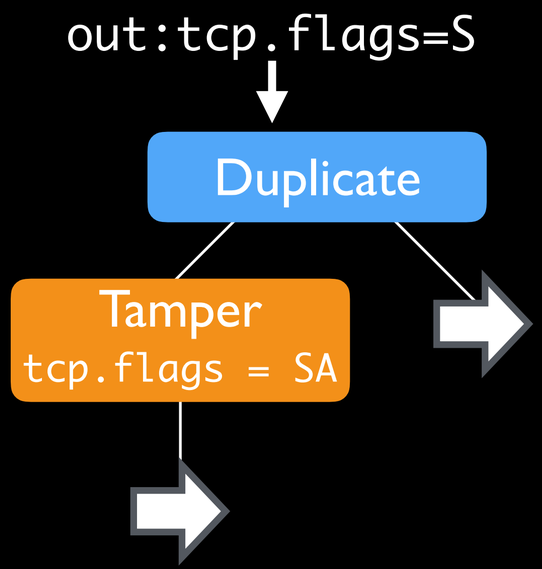

Trick the censor into thinking the client is the server

Segmentation species

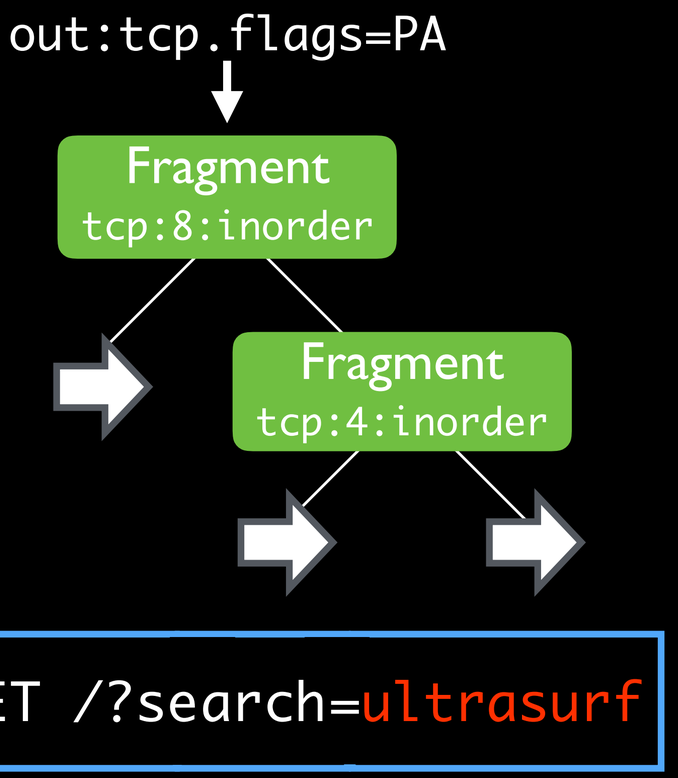

Segment the request

Turnaround species

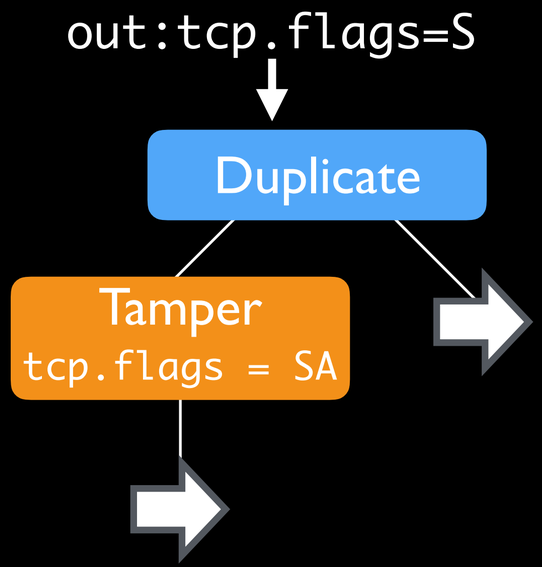

Trick the censor into thinking the client is the server

Segmentation species

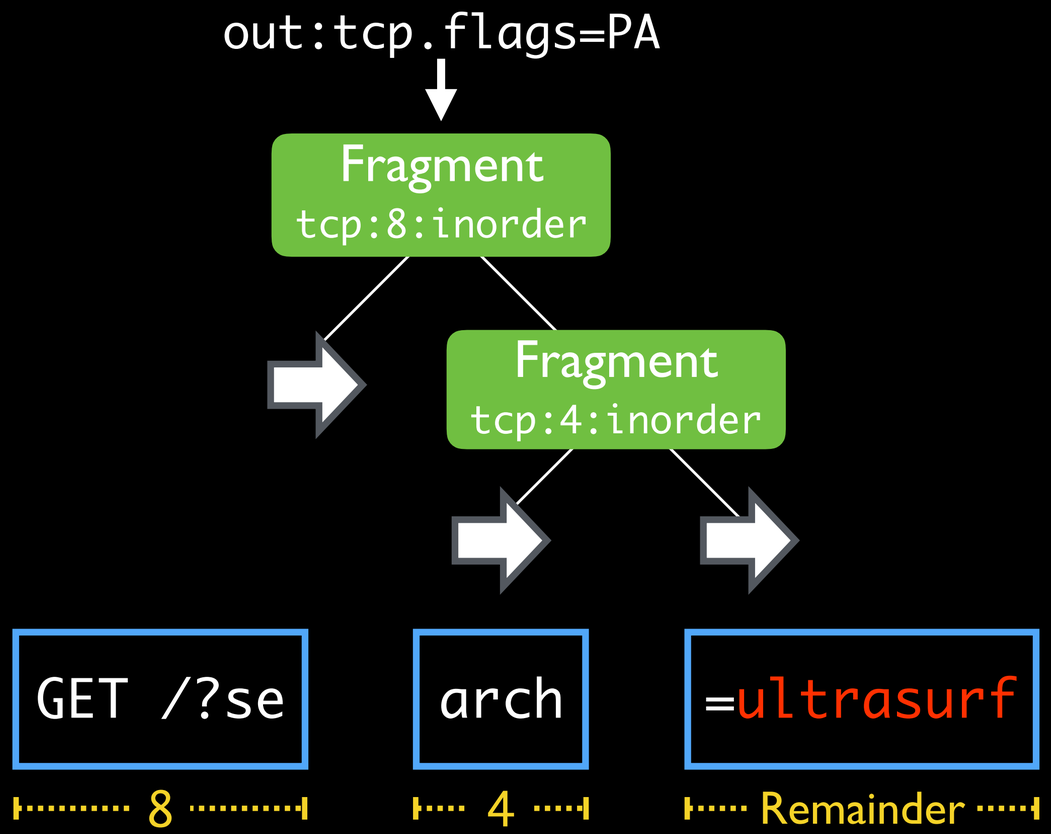

Segment the request, but *not the keyword*

Turnaround species

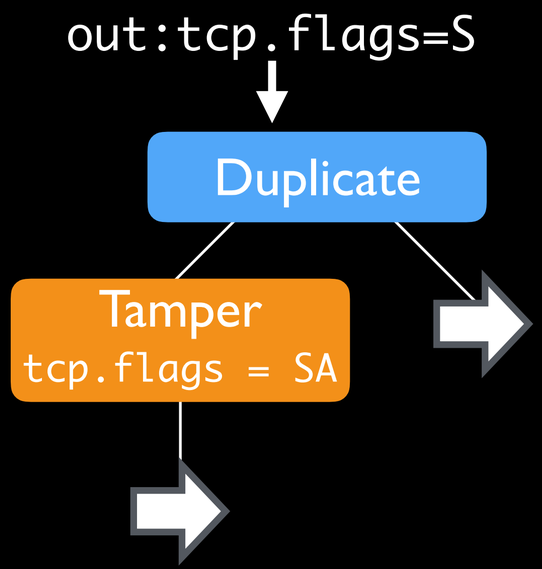

Trick the censor into thinking the client is the server

Segmentation species

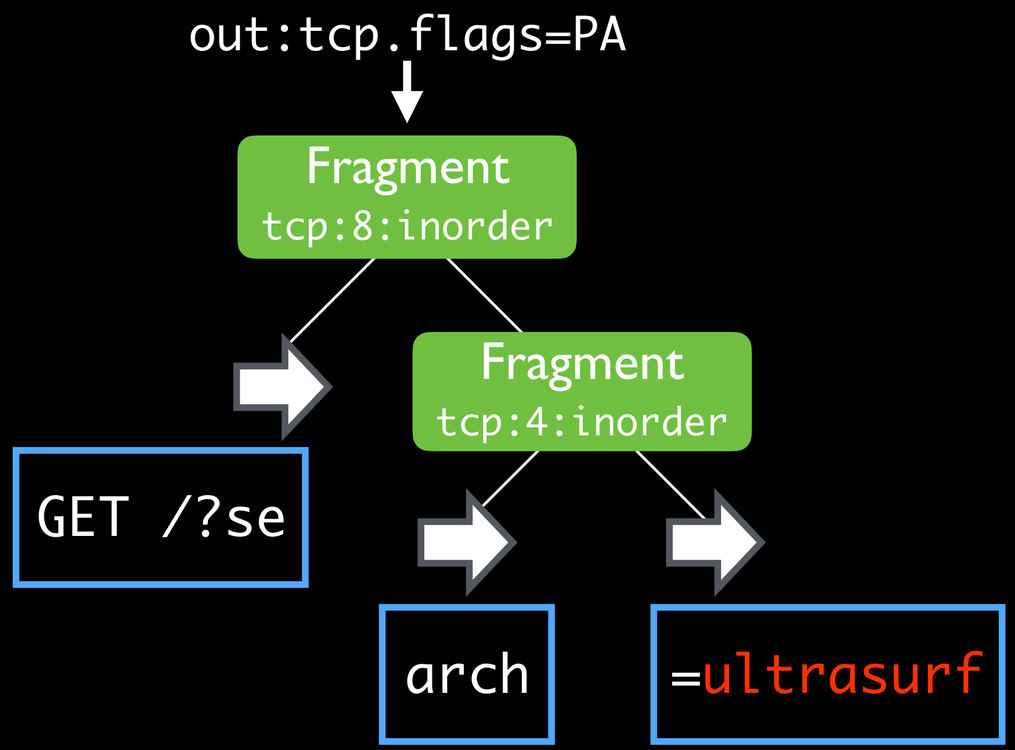

Segment the request, but *not the keyword*

Turnaround species

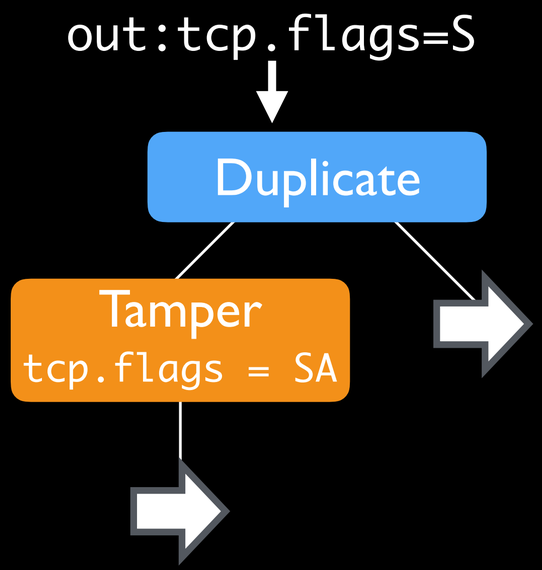

Trick the censor into thinking the client is the server

Segmentation species

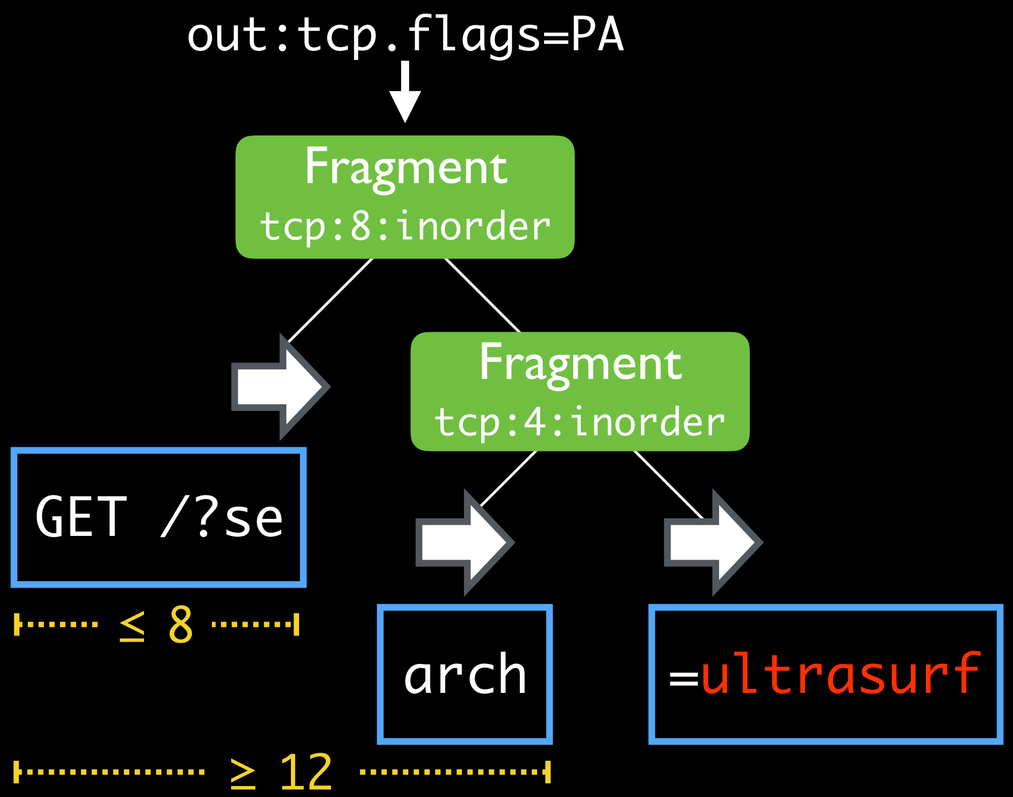

Segment the request, but *not the keyword*

Geneva is fast

July 2019

Kazakhstan launched an HTTPS
man-in-the-middle attack
that lasted several weeks

Geneva is fast

July 2019

Kazakhstan launched an HTTPS
man-in-the-middle attack
that lasted several weeks

Within 1 hour

Geneva found strategies to circumvent it

Censoring regime

Client

Geneva

Server

Server-side evasion

Censoring regime

Client

Server

Geneva

Server-side evasion

Censoring regime

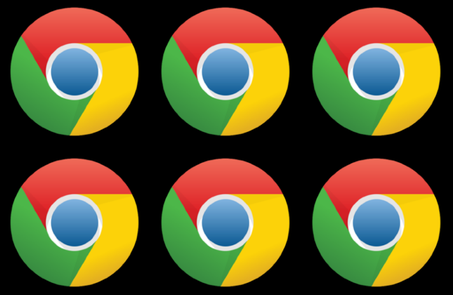

Clients

Server

Geneva

Potentially broadens reachability
without *any* client-side deployment

Server-side evasion “shouldn’t” work

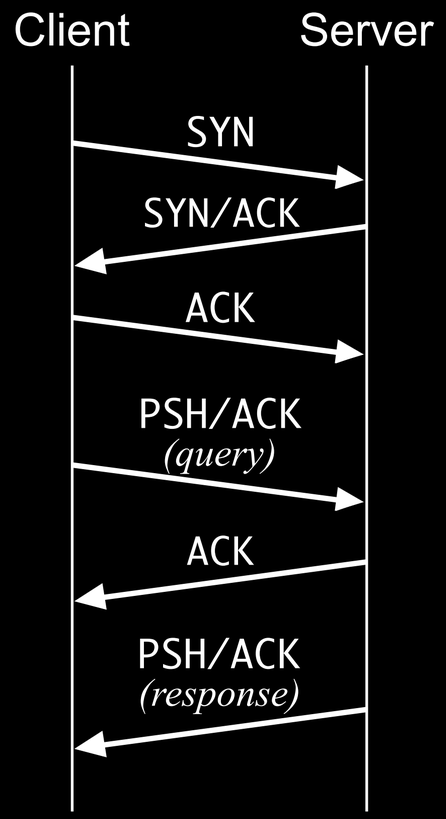

Server-side evasion “shouldn’t” work

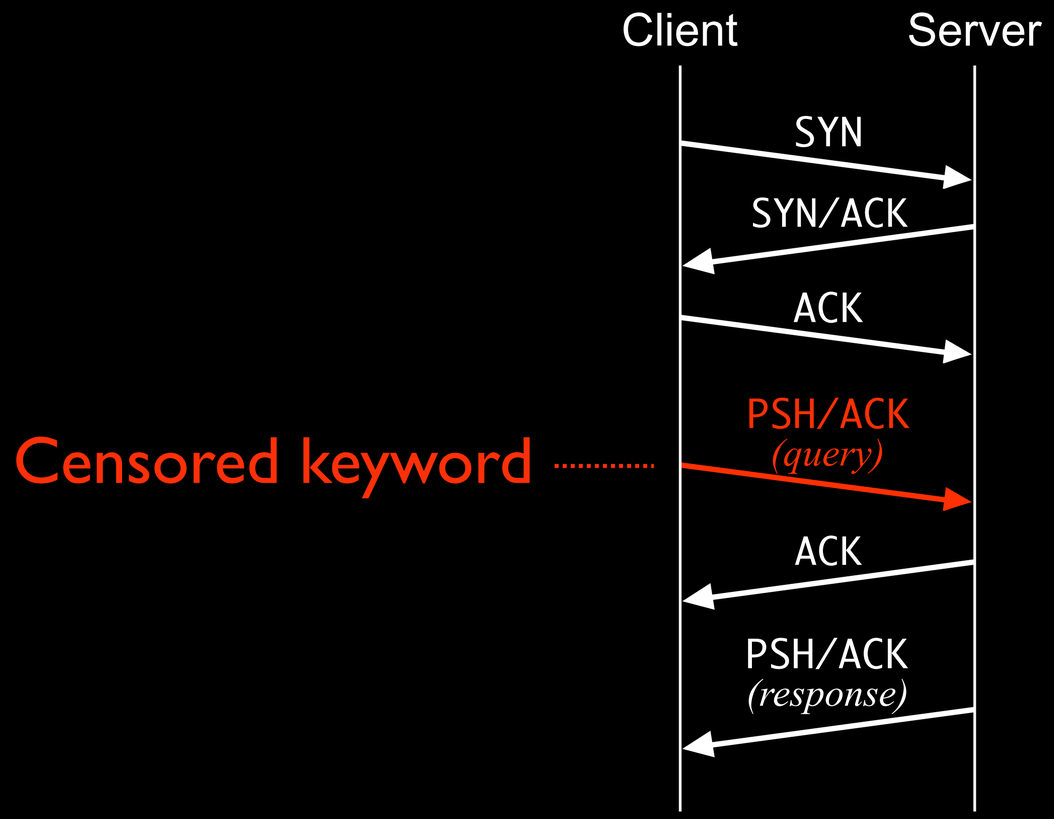

Server-side evasion “shouldn’t” work

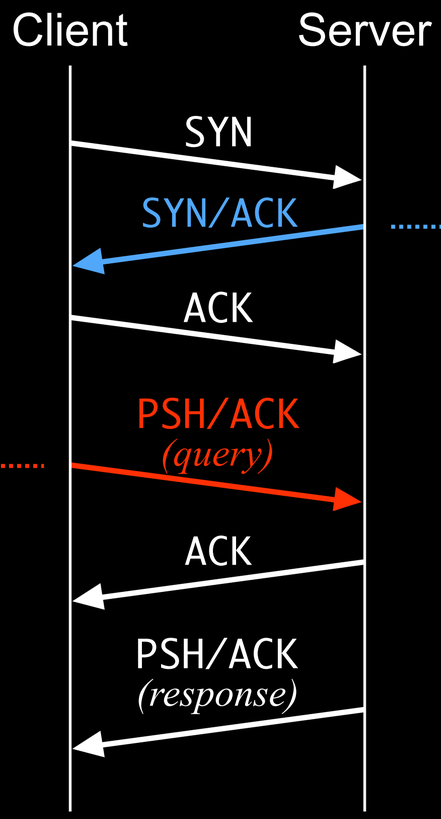

All a server does
before client is censored

Censored keyword

Server-side evasion "shouldn't" work

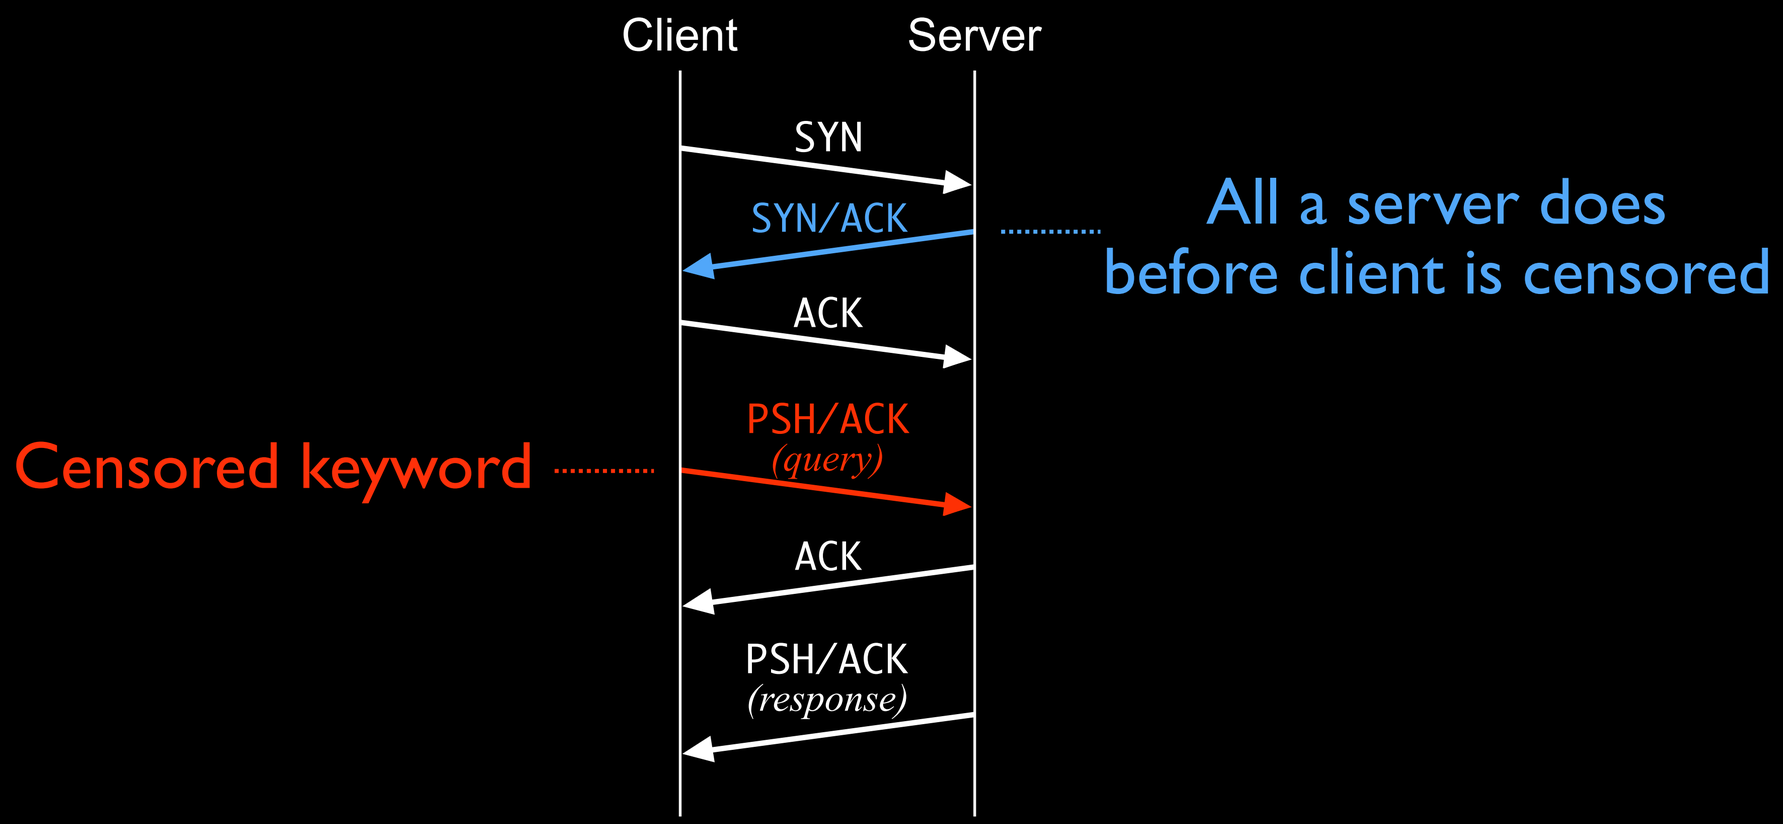

Fortunately, the AI doesn't know it "shouldn't" work

Server-side evasion “shouldn’t” work

Server-side results

Server-side evasion “shouldn’t” work

Server-side results

China

8 strategies

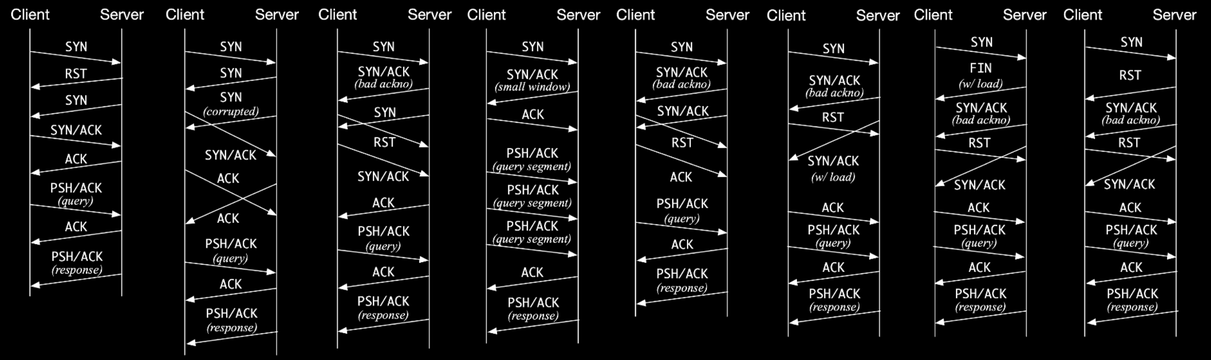

Server-side evasion “shouldn’t” work

Server-side results

China

8 strategies

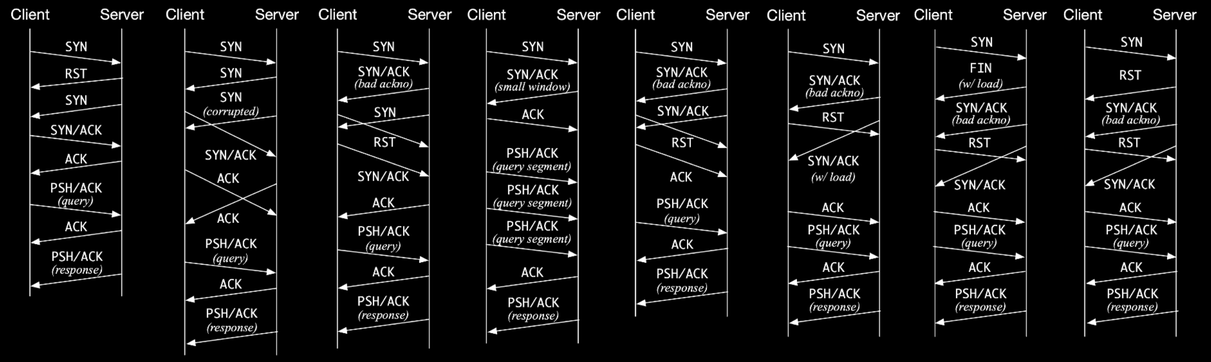

India

1 strategy

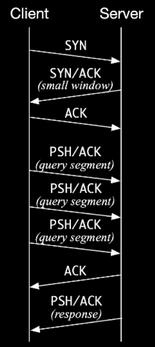

Server-side evasion “shouldn’t” work

Server-side results

China

8 strategies

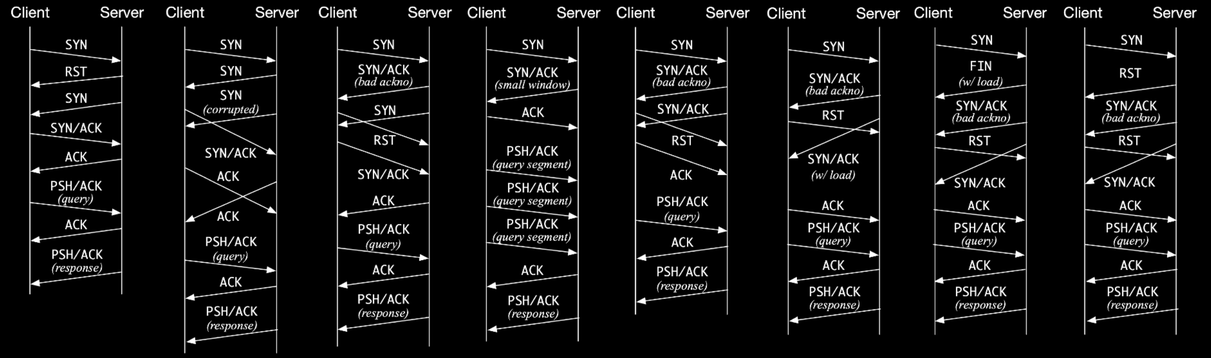

India

1 strategy

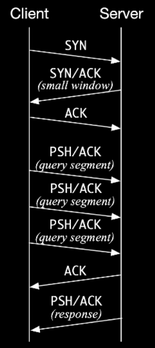

Kazakhstan

3 strategies

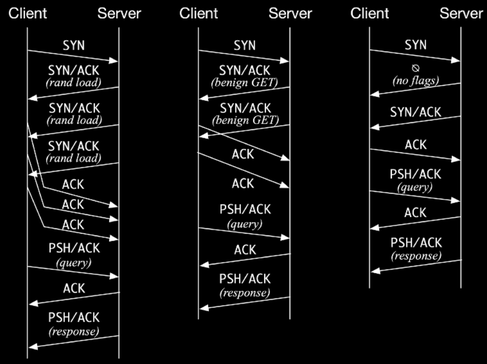

Server-side evasion “shouldn’t” work

Server-side results

China

8 strategies

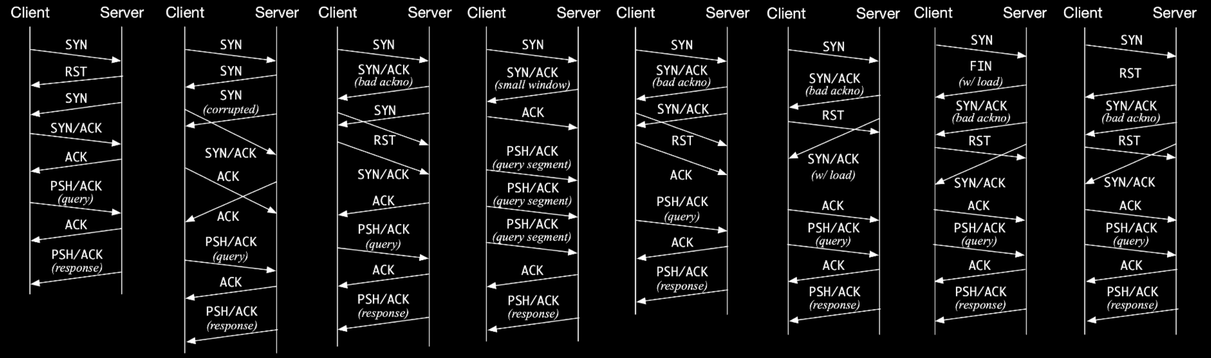

India

1 strategy

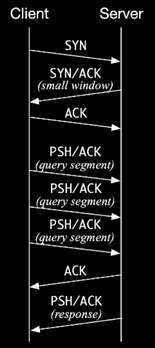

Kazakhstan

3 strategies

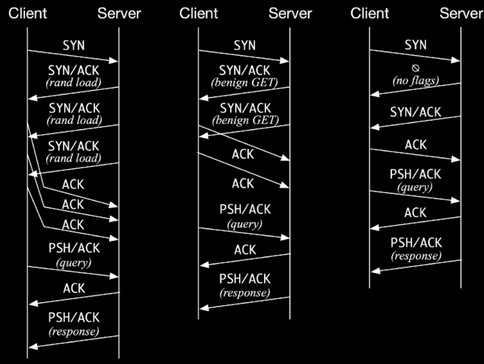

None of these require *any* client-side deployment

Server-side evasion “shouldn’t” work

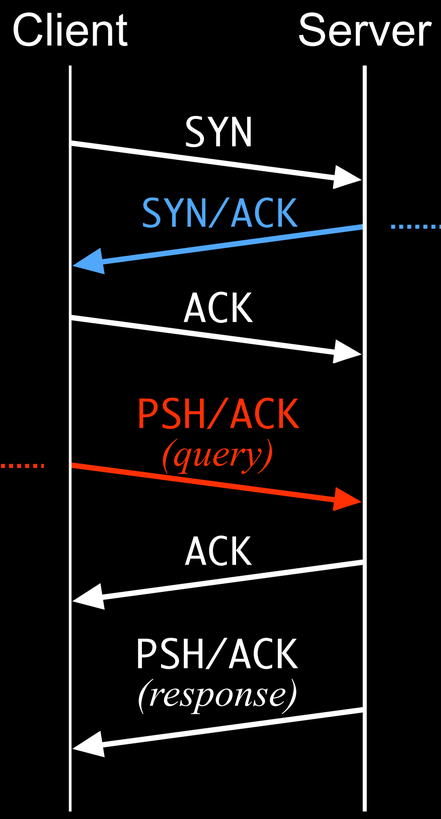

All a server does
before client is censored

Censored keyword

Simultaneous Open-based Desynchronization

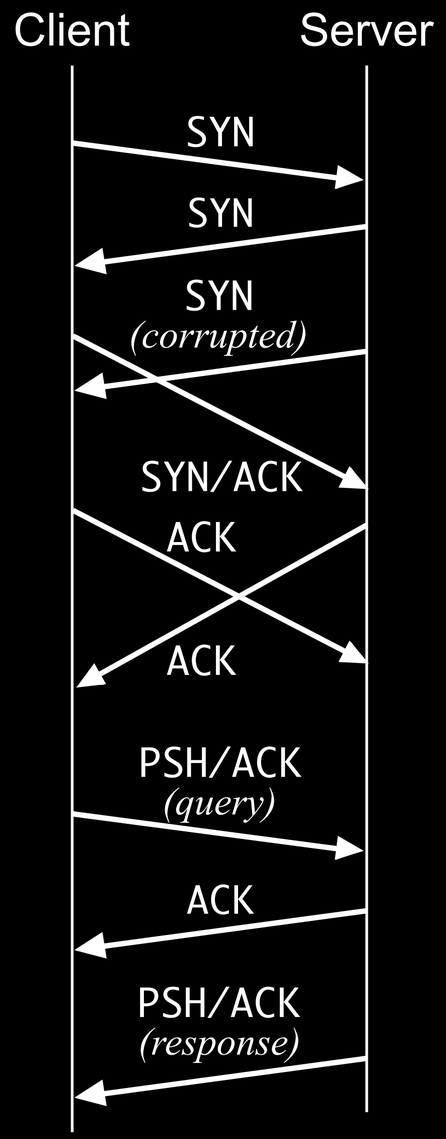

Simultaneous Open-based Desynchronization

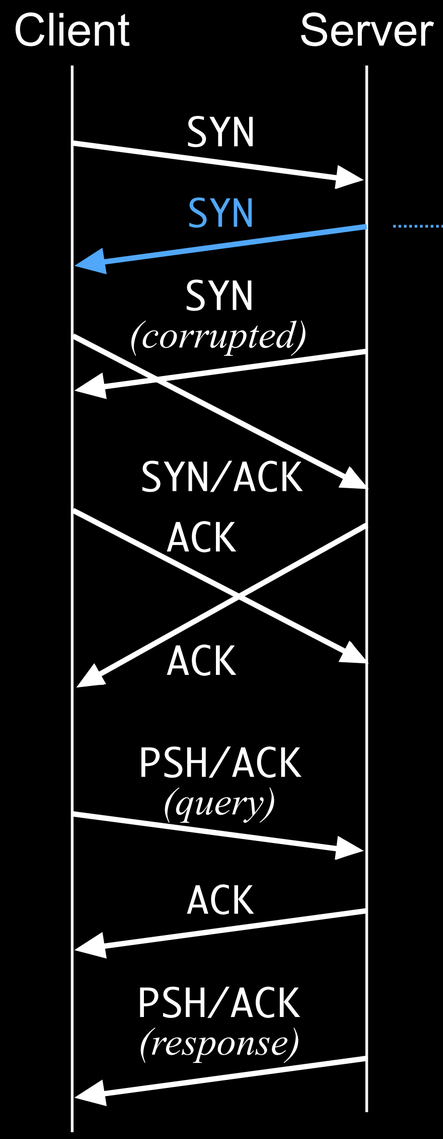

TCP simultaneous open

Simultaneous Open-based Desynchronization

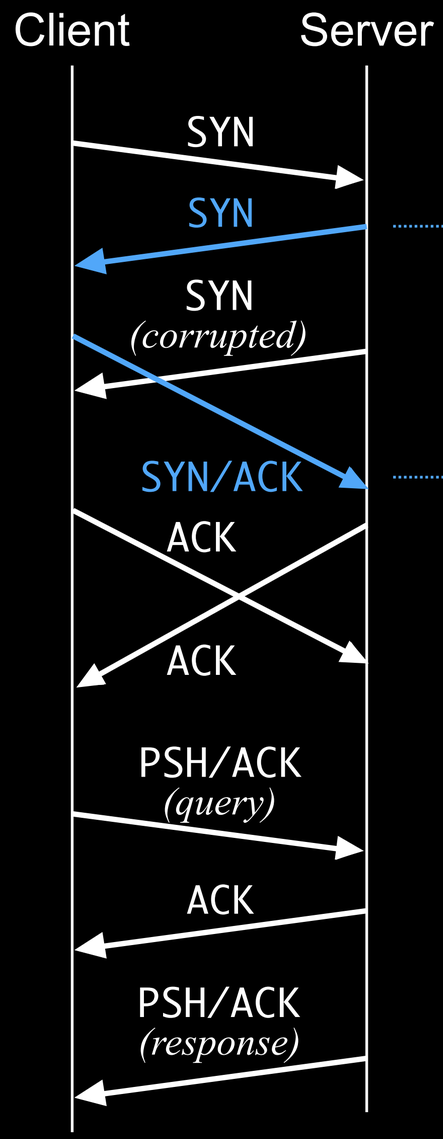

TCP simultaneous open

Client sends a SYN/ACK

Simultaneous Open-based Desynchronization

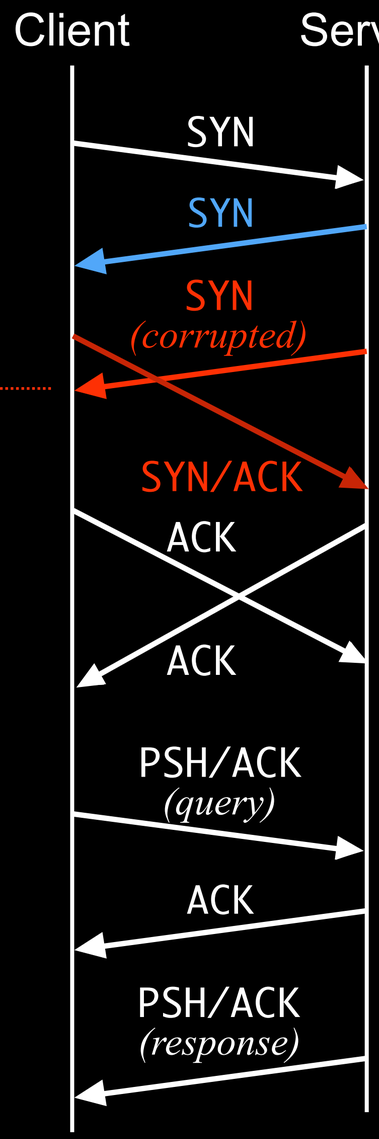

TCP simultaneous open

Censor de-synchronizes

Client sends a SYN/ACK

Double-benign GETs

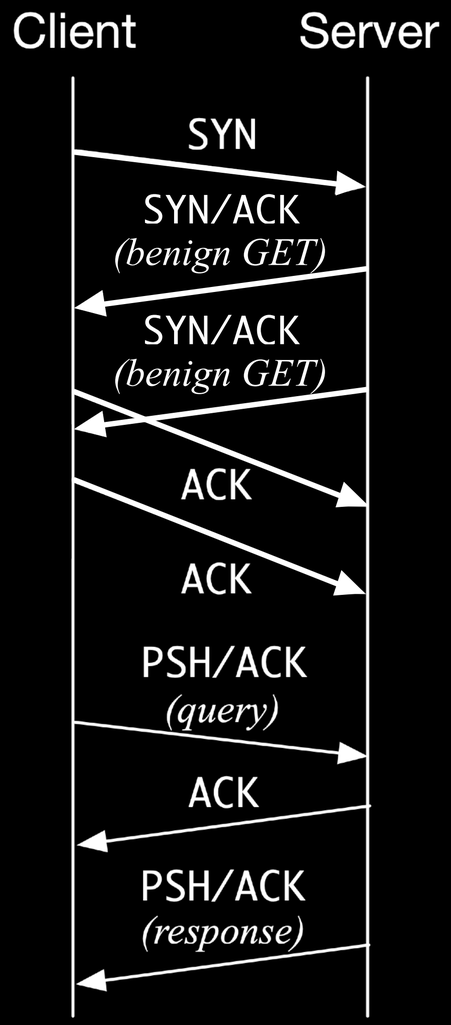

Double-benign GETs

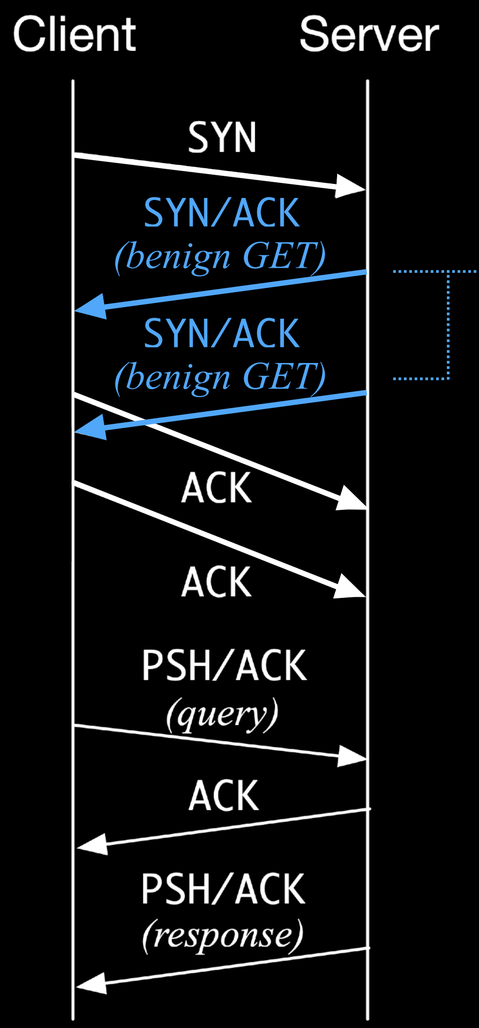

Server sends uncensored GETs inside *two* SYN/ACKs

Double-benign GETs

Censor confuses connection direction

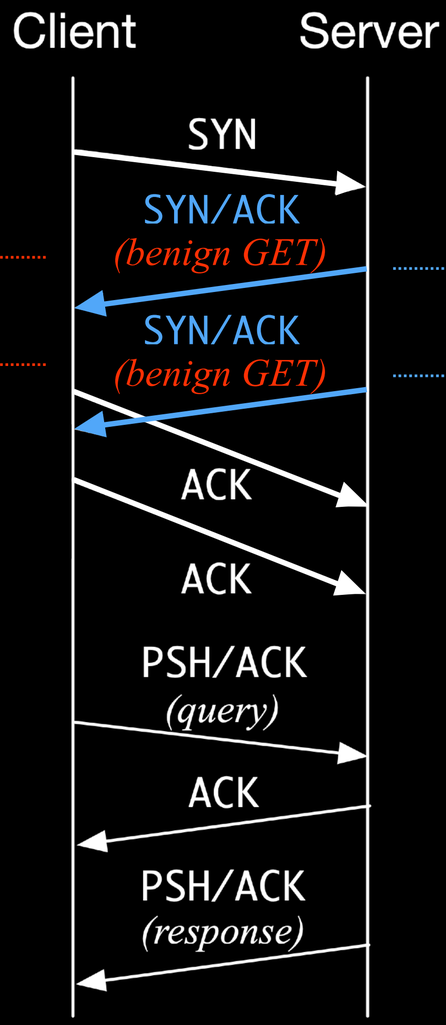

Server sends uncensored GETs inside *two* SYN/ACKs

Automating the arms race

AI has the potential to fast-forward the arms race *for both sides*

Automating the arms race

AI has the potential to fast-forward the arms race *for both sides*

Bugs in implementation	Easy for censors to fix the low-hanging fruit
Gaps in logic	Harder for censors to fix systemic issues

Automating the arms race

AI has the potential to fast-forward the arms race *for both sides*

Bugs in implementation	Easy for censors to fix the low-hanging fruit
Gaps in logic	Harder for censors to fix systemic issues

What is the *logical conclusion* of the arms race?

Automatically learning how to evade censorship

Geneva

Genetic Evasion

Server-side evasion

Finds strategies quickly

Dozens of strategies

Evasion advantage

Geneva code and website geneva.cs.umd.edu