How China Detects and Blocks Shadowsocks

Alice, Bob, Carol (GFW Report)
Jan Beznazwy
Amir Houmansadr (University of Massachusetts Amherst)

https://gfw.report/publications/imc20/en/

ACM Internet Measurement Conference 2020

Overview

The Great Firewall of China detects and blocks **Shadowsocks** using a combination of **passive traffic analysis** and **active probing**.

Shadowsocks

Shadowsocks is an encrypted proxy protocol, designed to be difficult to detect.

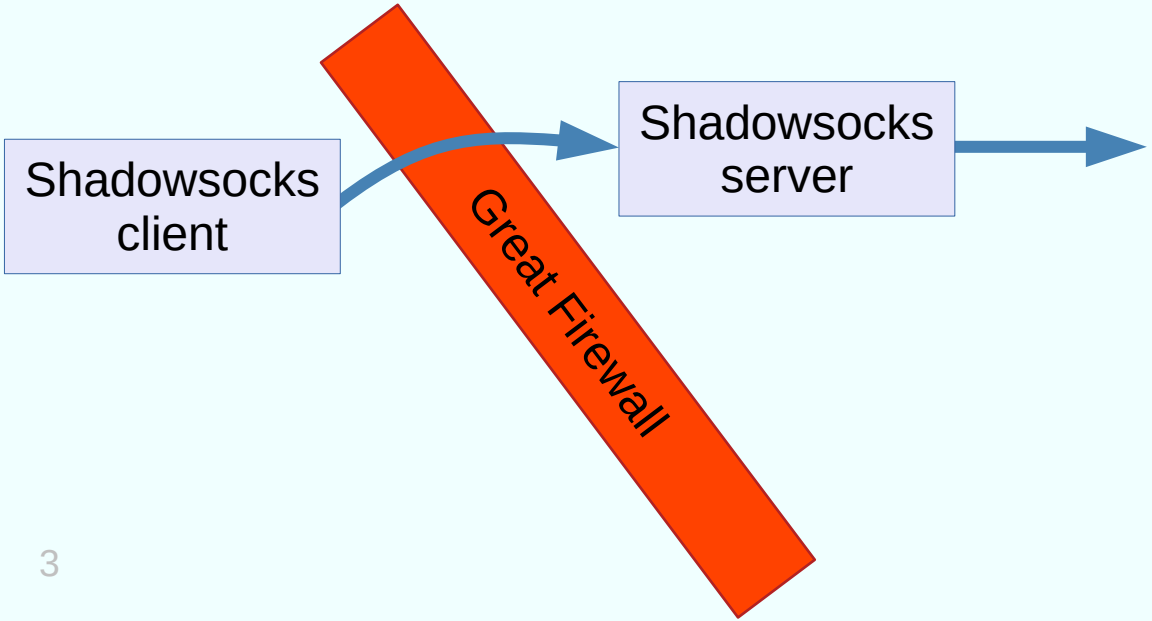

Active probing

1. Identify *possible* Shadowsocks connections.
2. Send probes to the server to confirm.

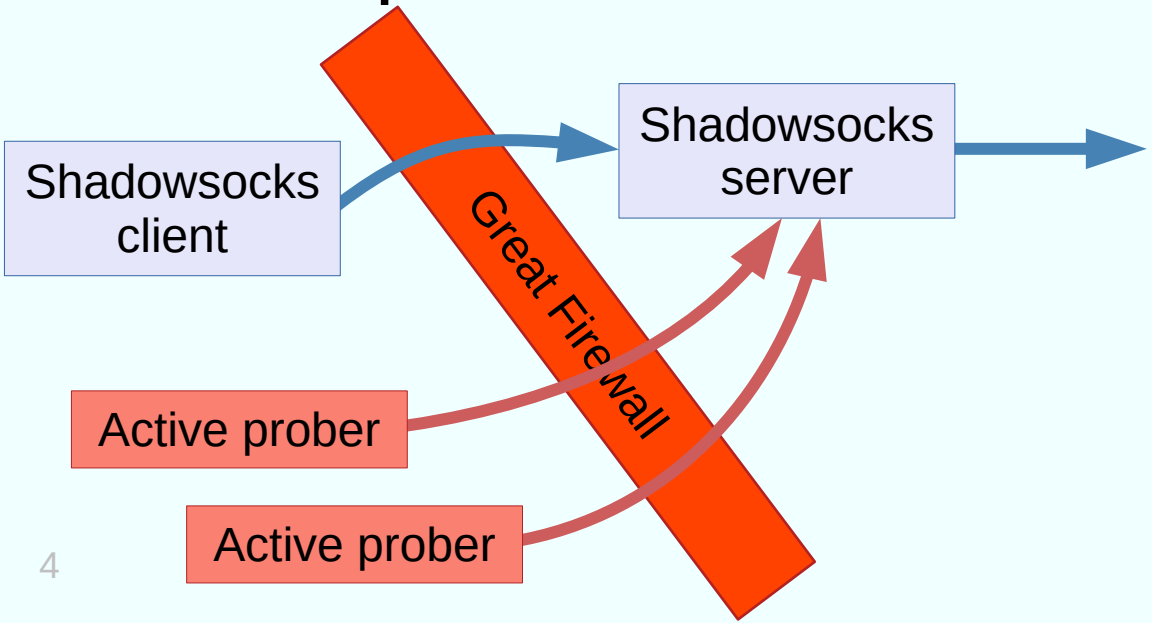

Live server experiment

- Run Shadowsocks servers outside China, connect to them from inside.
- Shadowsocks-libev and OutlineVPN.
- September 2019 to January 2020.

Server experiment: main observations

- Active probers send a variety of probe types, some using replay and some apparently random.
- Legitimate connections may be stored and replayed days later.
- Non-replay probes have a distinctive distribution of payload lengths.
- Active probes come from thousands of IP addresses.

Replay-based probes

- Derived from the first packet in a legitimate connection – perhaps with some bytes changed.

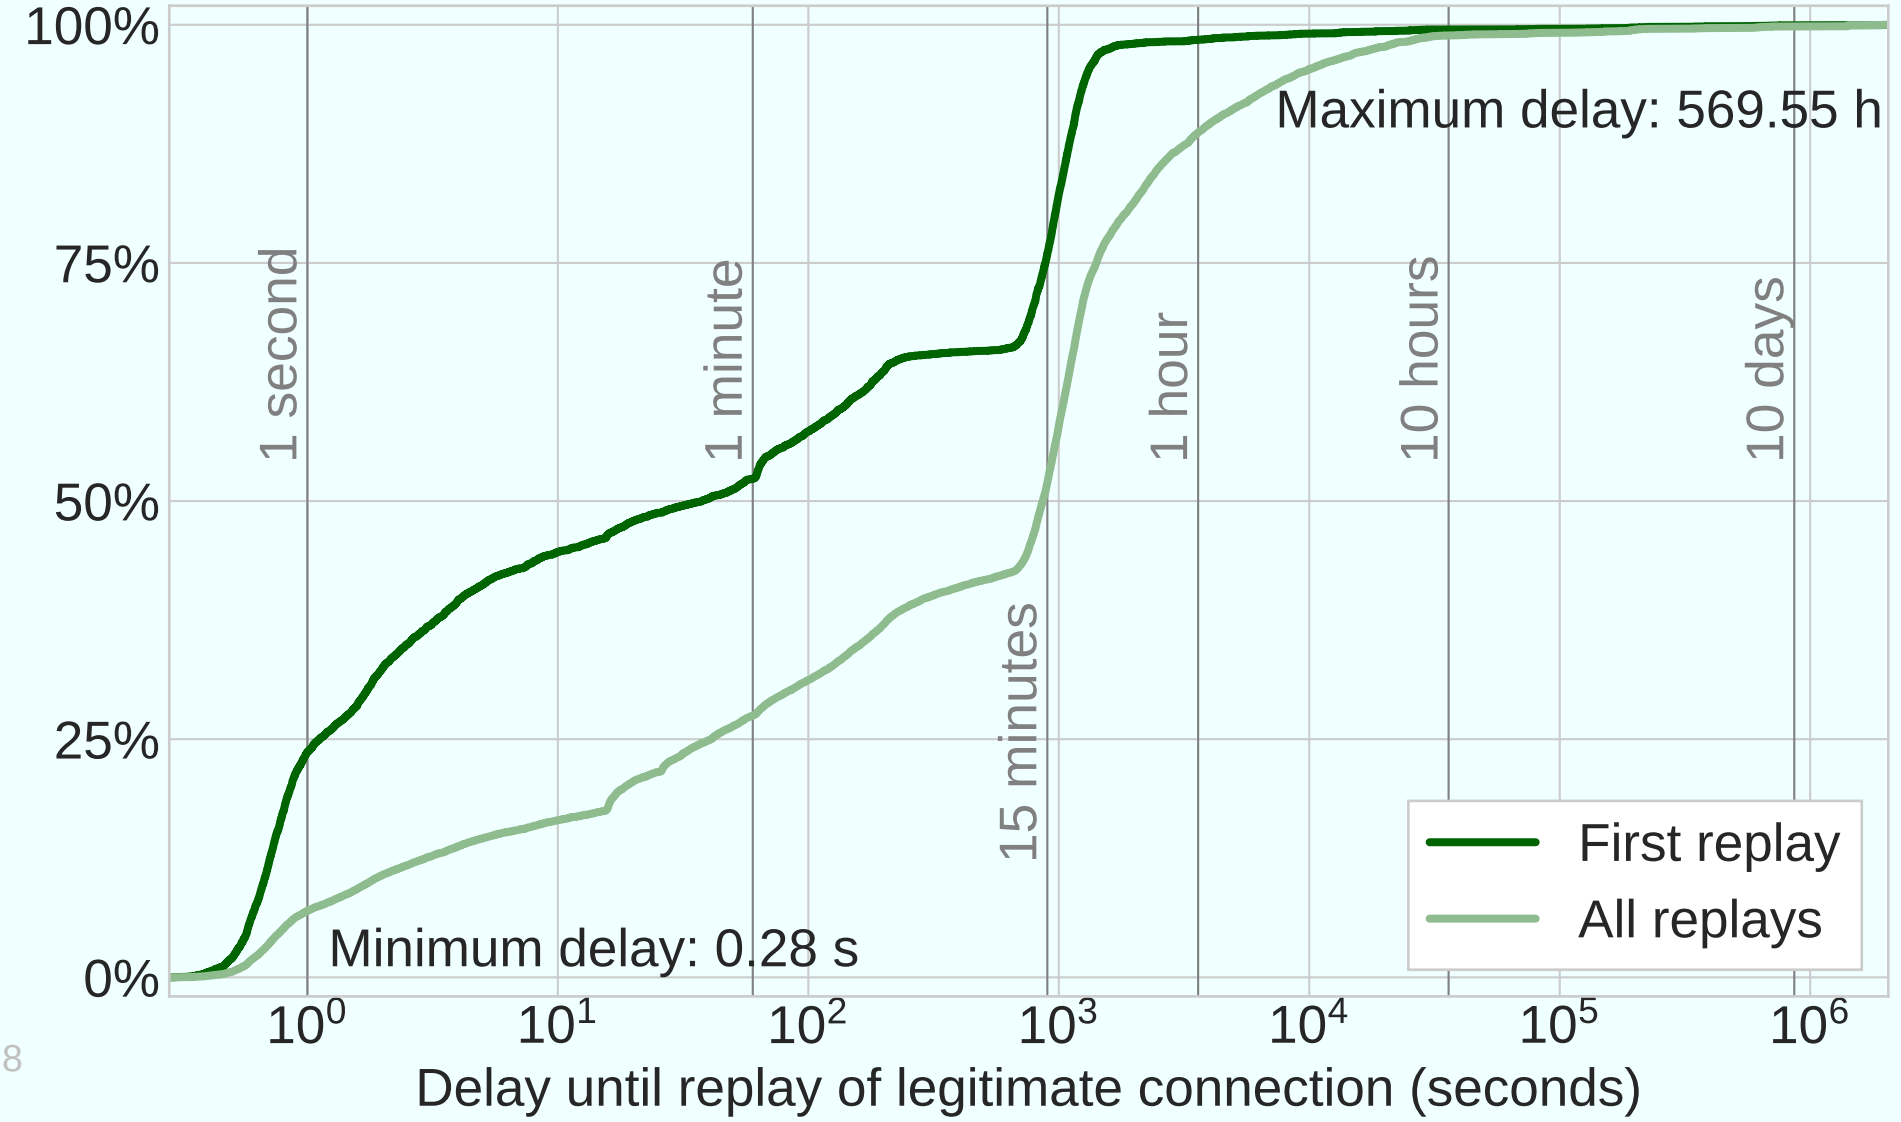
100%
75%
50%
25%
0%
1 second
1 minute
15 minutes
1 hour
10 hours
10 days
Maximum delay: 569.55 h
Minimum delay: 0.28 s
First replay
All replays
10^0
10^1
10^2
10^3
10^4
10^5
10^6
Delay until replay of legitimate connection (seconds)

Non-replay probes

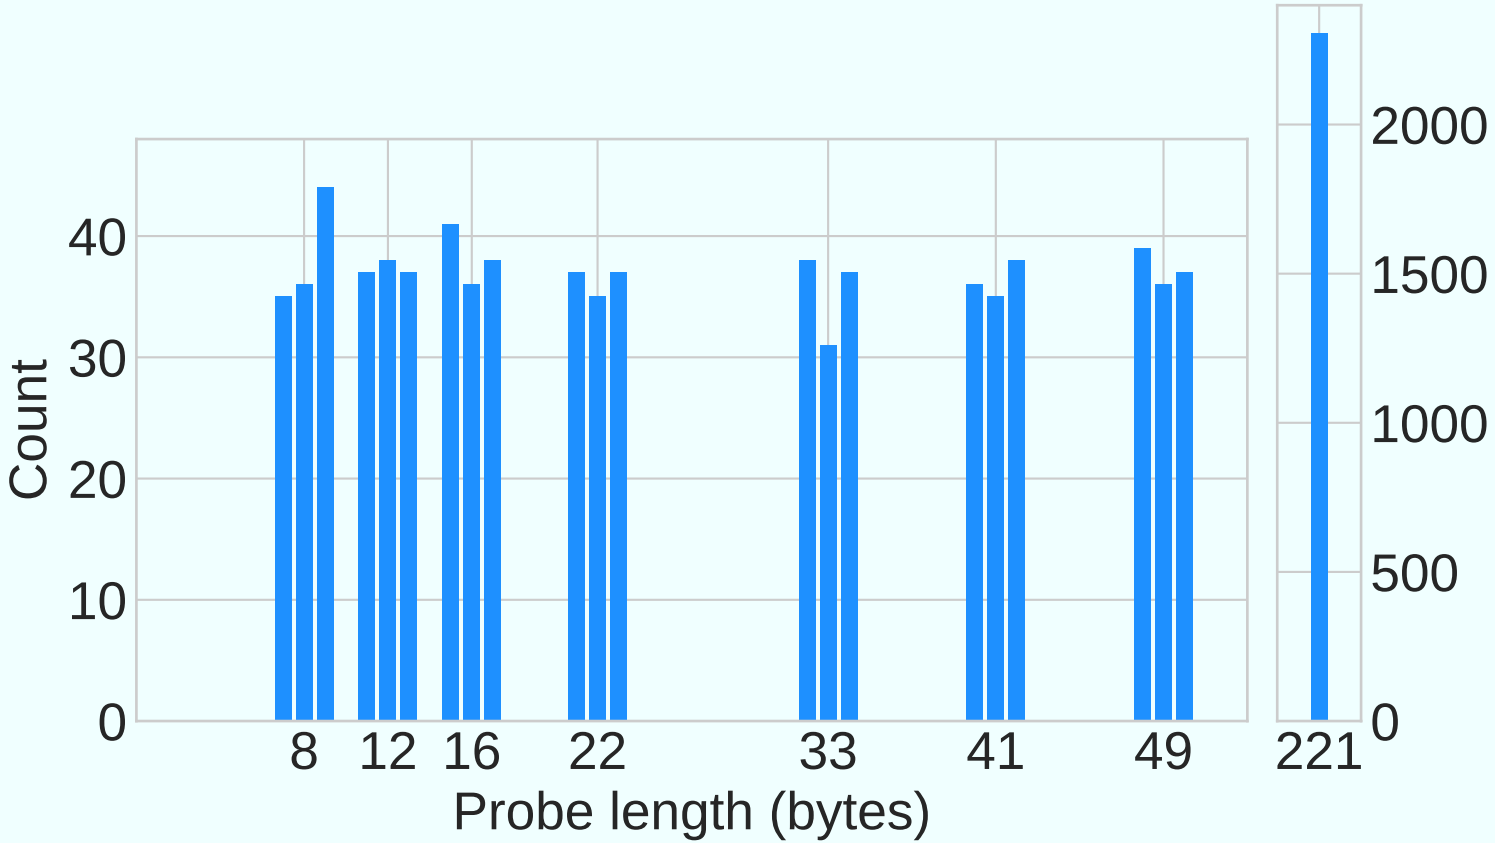

How Shadowsocks servers react to random probes

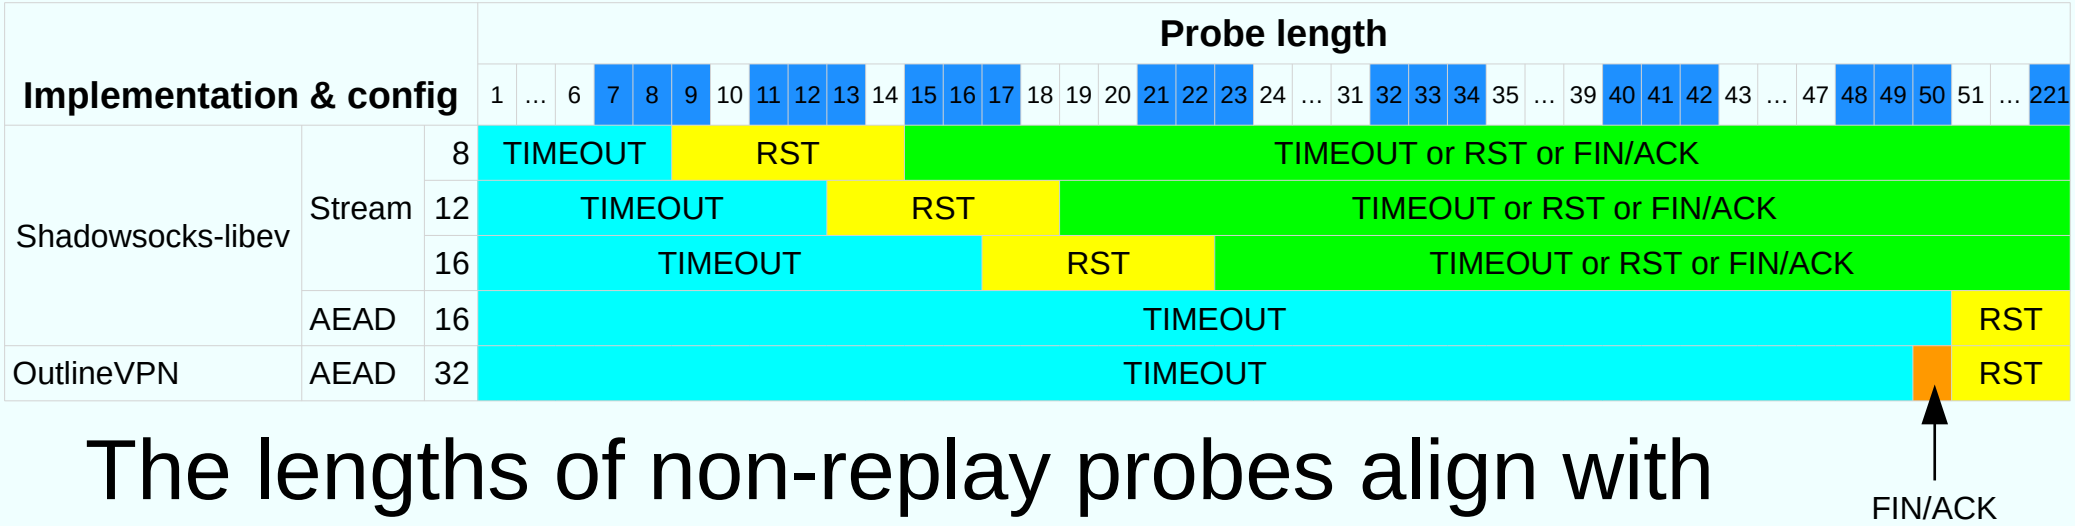

The lengths of non-replay probes align with thresholds at which servers switch from timing out to closing the connection.

Active prober source IP addresses

IP address	ASN	count
175.42.1.21	4837	44
223.166.74.207	17621	38
113.128.105.20	4134	36
124.235.138.113	4134	36
221.213.75.88	4837	33
112.80.138.231	4837	32
116.252.2.39	4134	32
124.235.138.231	4134	32
221.213.75.126	4837	32
223.166.74.110	17621	31
…12,288 additional rows…		
223.166.75.225	17621	1
223.166.75.226	17621	1

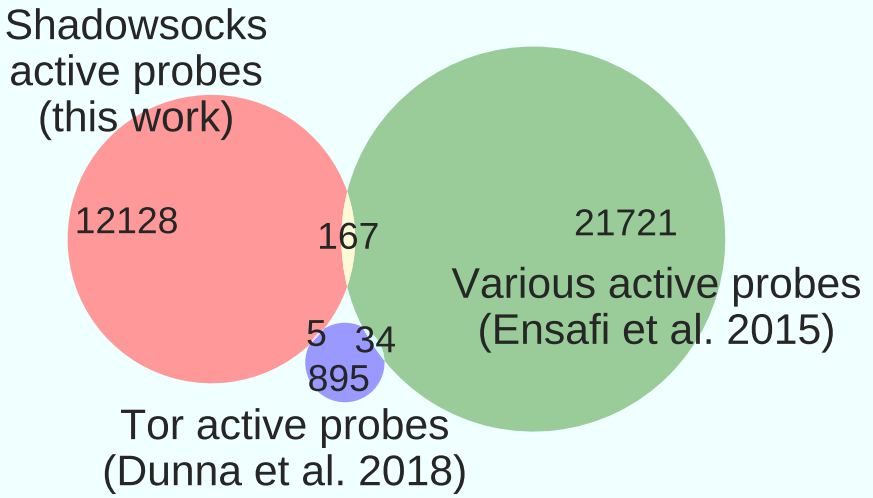

Shared TCP timestamp sequences

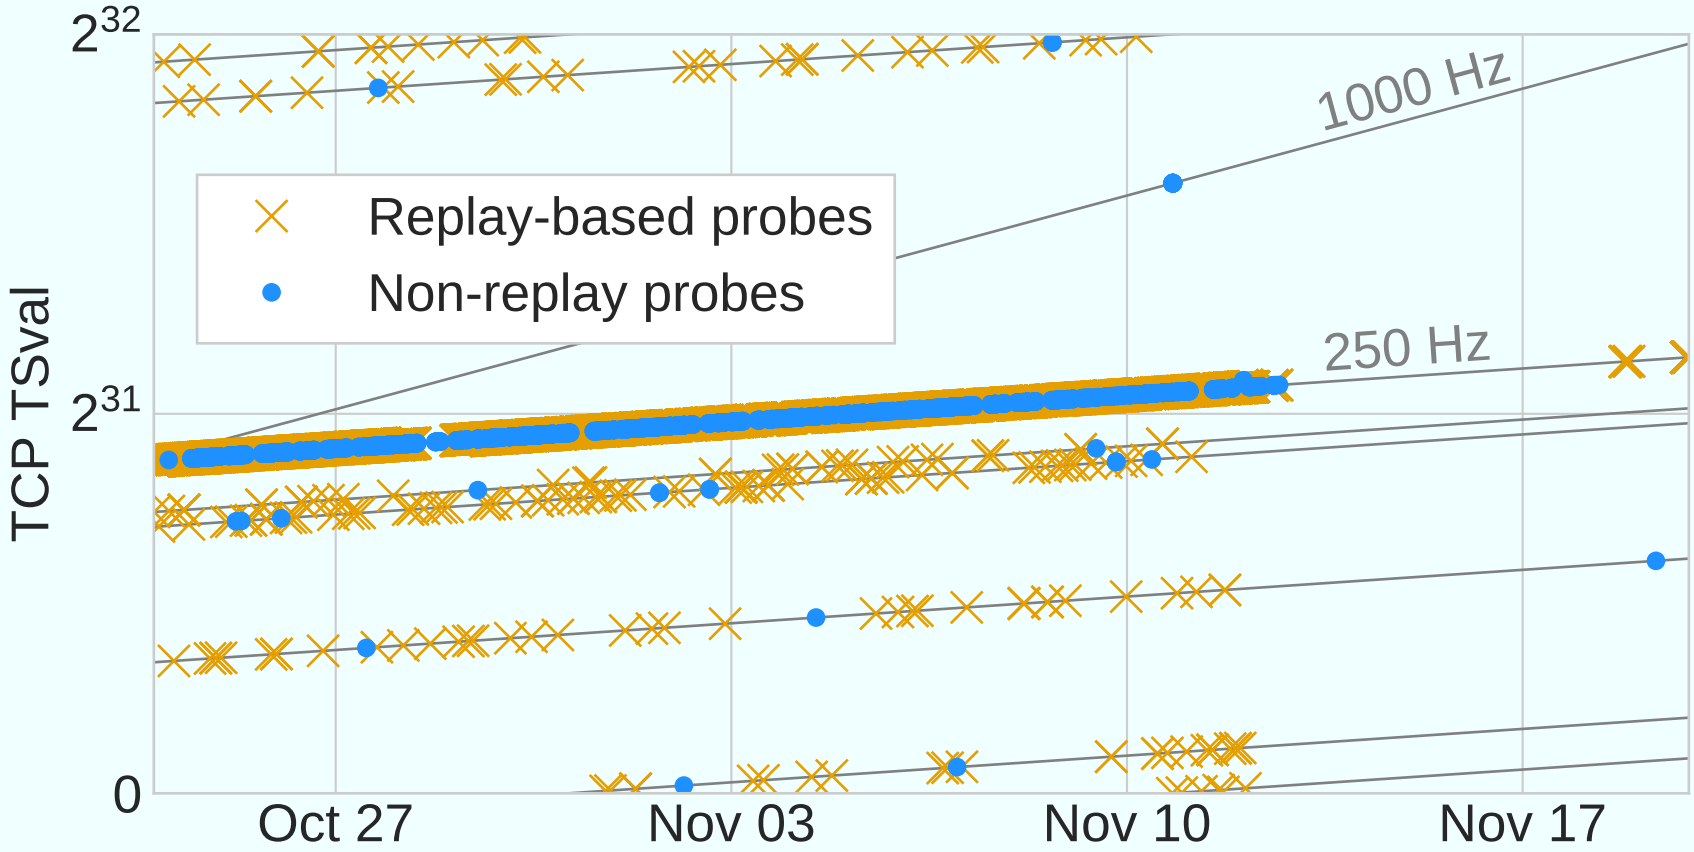

Likelihood of replay by entropy

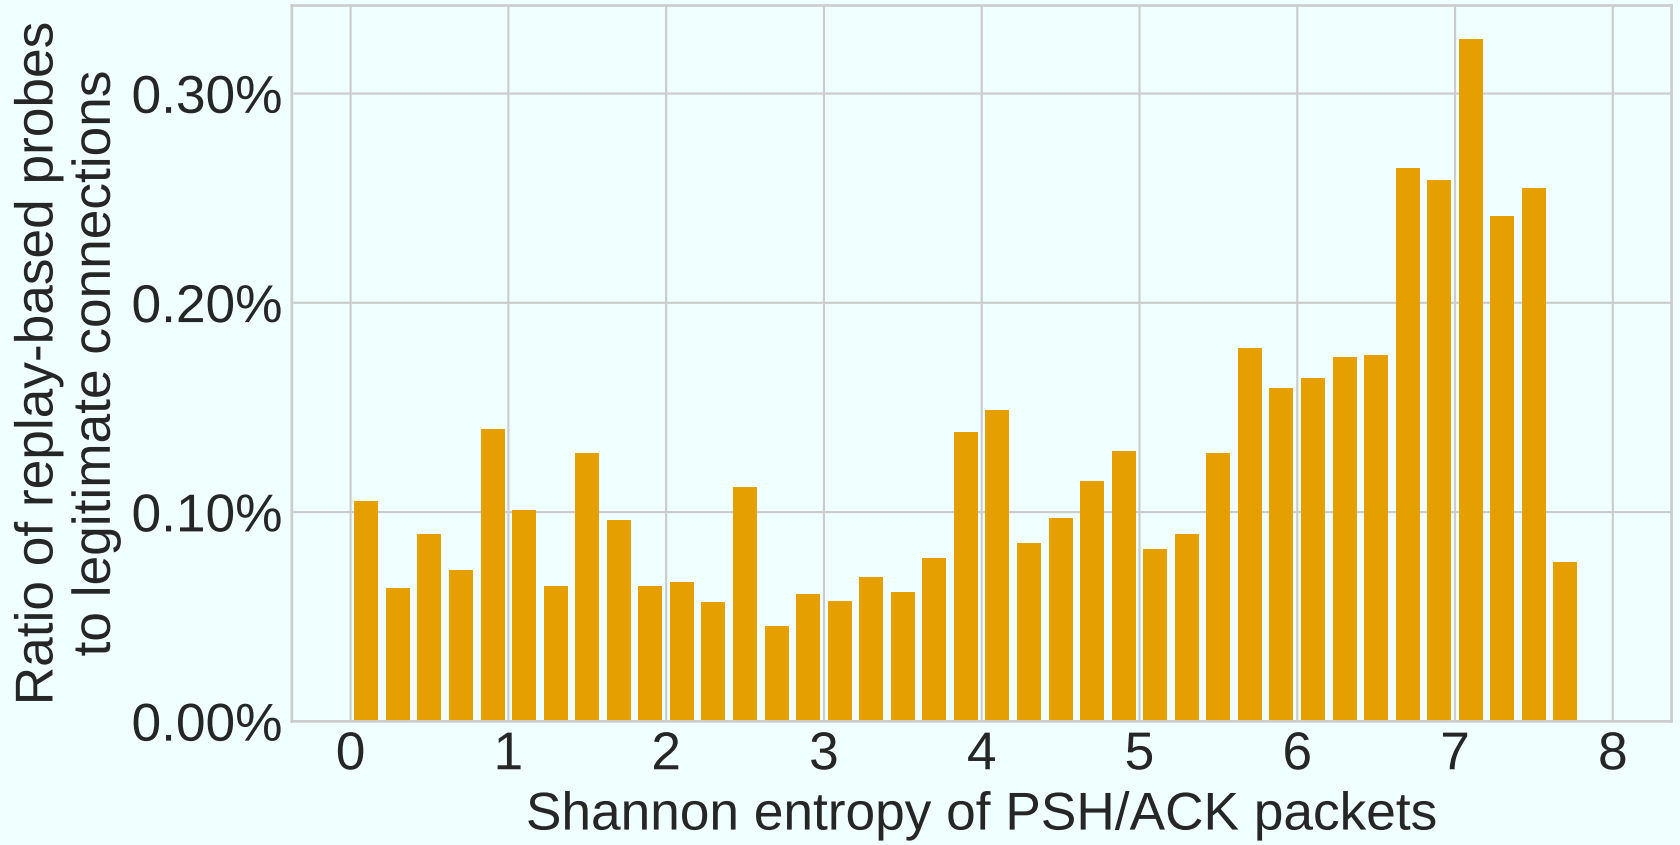

Active probe length distribution

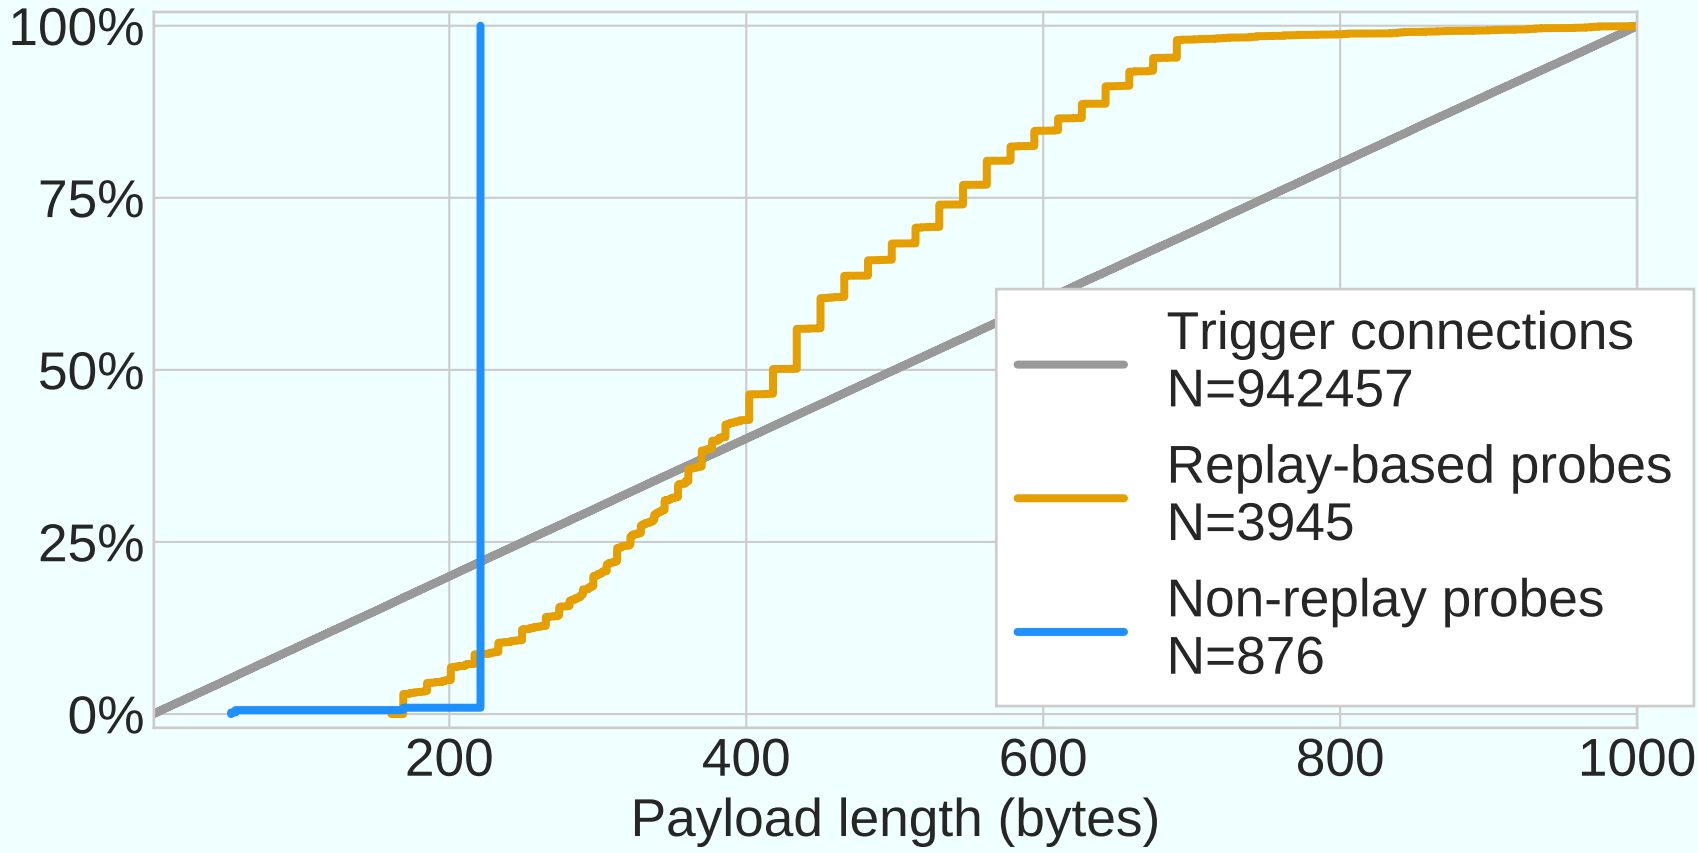

Active probe length distribution

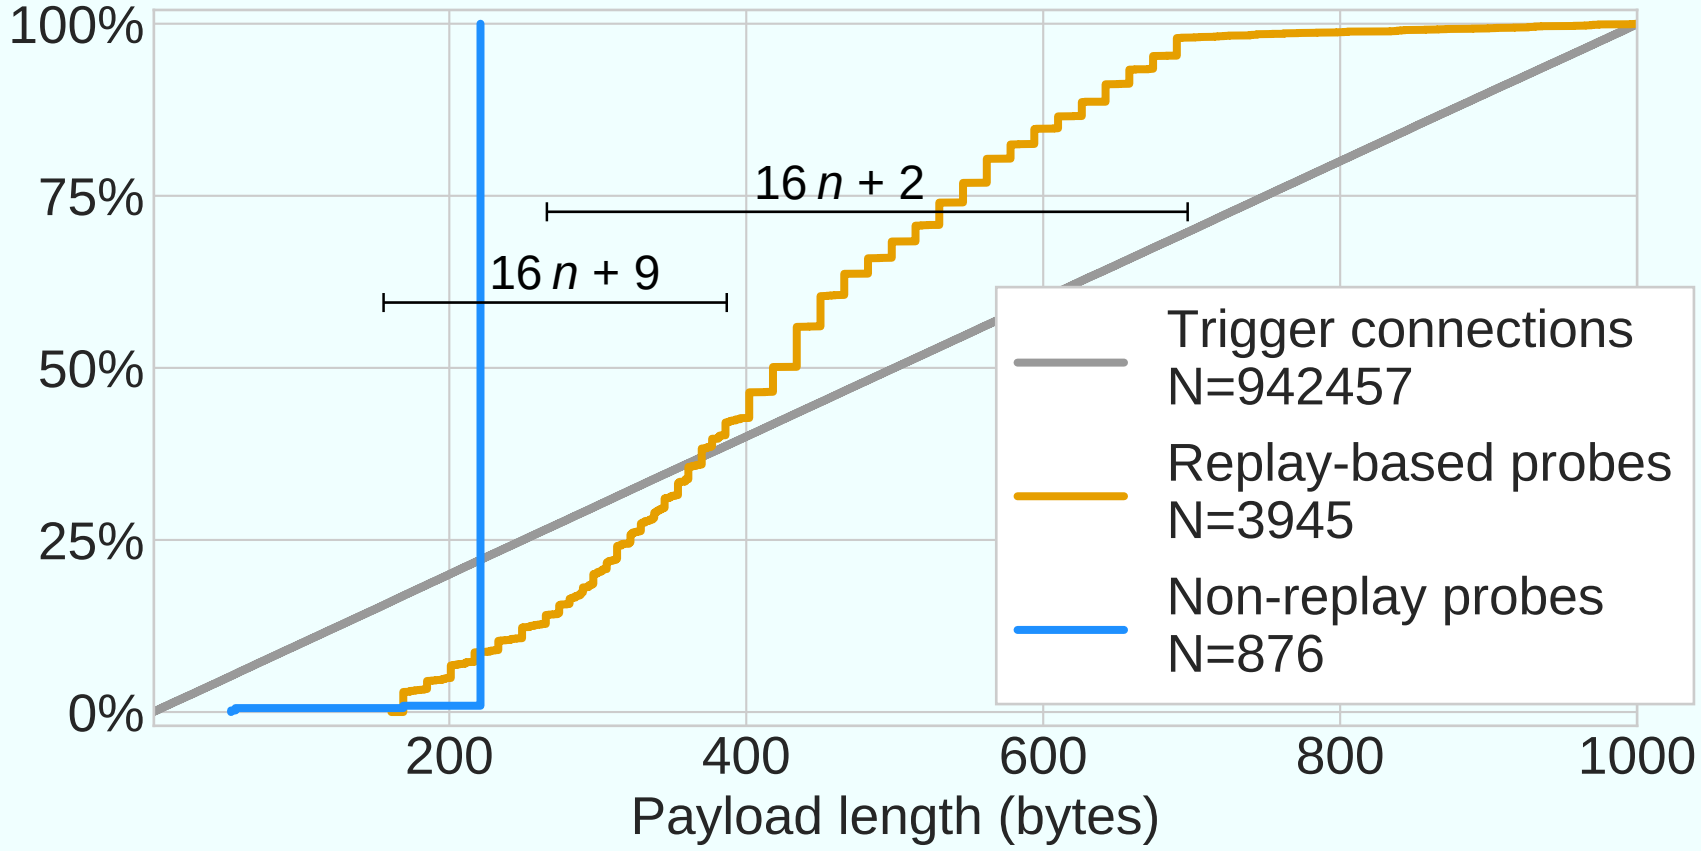

Mitigation and circumvention

- Evade passive traffic analysis
  (change entropy or packet lengths), or
- Change responses to unauthenticated probes.

Brdgrd

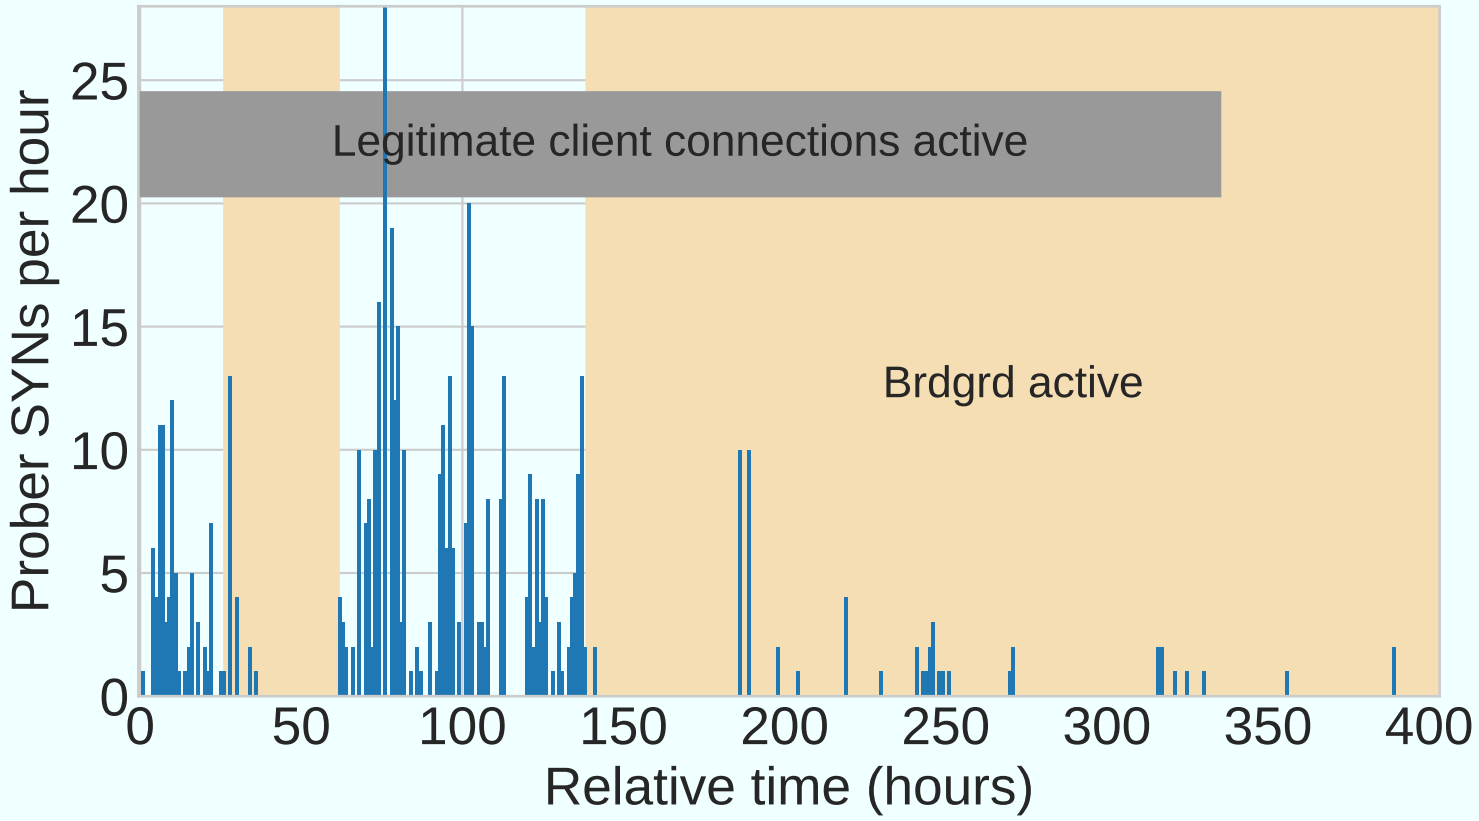
Legitimate client connections active
Brdgrd active
Prober SYNs per hour
0
5
10
15
20
25
0
50
100
150
200
250
300
350
400
Relative time (hours)

How (old) Shadowsocks servers react to random probes

Implementation & config			Probe length
			1 … 6 **7 8 9** 10 **11 12 13** 14 **15 16 17** 18 19 20 **21 22 23** 24 … 31 **32 33 34** 35 … 39 **40 41 42** 43 … 47 **48 49 50** 51 … **221**
Shadowsocks-libev	Stream	8	TIMEOUT (1 … 6 **7 8 9**) · RST (10 **11 12 13** 14 **15 16 17**) · TIMEOUT or RST or FIN/ACK (18 19 20 **21 22 23** 24 … 31 **32 33 34** 35 … 39 **40 41 42** 43 … 47 **48 49 50** 51 … **221**)
		12	TIMEOUT (1 … 6 **7 8 9** 10 **11 12 13** 14) · RST (**15 16 17** 18 19 20 **21 22 23**) · TIMEOUT or RST or FIN/ACK (24 … 31 **32 33 34** 35 … 39 **40 41 42** 43 … 47 **48 49 50** 51 … **221**)
		16	TIMEOUT (1 … 6 **7 8 9** 10 **11 12 13** 14 **15 16 17**) · RST (18 19 20 **21 22 23**) · TIMEOUT or RST or FIN/ACK (24 … 31 **32 33 34** 35 … 39 **40 41 42** 43 … 47 **48 49 50** 51 … **221**)
	AEAD	16	TIMEOUT (1 … 6 **7 8 9** 10 **11 12 13** 14 **15 16 17** 18 19 20 **21 22 23** 24 … 31 **32 33 34** 35 … 39 **40 41 42** 43 … 47 **48 49 50** 51 … **221**) · RST
OutlineVPN	AEAD	32	TIMEOUT (1 … 6 **7 8 9** 10 **11 12 13** 14 **15 16 17** 18 19 20 **21 22 23** 24 … 31 **32 33 34** 35 … 39 **40 41 42** 43 … 47 **48 49 50**) · FIN/ACK · RST (51 … **221**)

How (new) Shadowsocks servers react to random probes

Implementation & config			Probe length
			1 … 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 … 31 32 33 34 35 … 39 40 41 42 43 … 47 48 49 50 51 … 221
Shadowsocks-libev	Stream	8	TIMEOUT / TIMEOUT or RST or FIN/ACK
		12	TIMEOUT / TIMEOUT or RST or FIN/ACK
		16	TIMEOUT / TIMEOUT or RST or FIN/ACK
	AEAD	16	TIMEOUT
OutlineVPN	AEAD	32	TIMEOUT

Summary

- The Great Firewall of China detects Shadowsocks servers using a combination of passive traffic analysis and active probing.
- Probing is triggered by the first data packet in a TCP connection, and is more likely when the packet has high entropy and certain lengths.
- There are several probe types, some based on replay and some not.
- Probes come from many source IP addresses, but are evidently centrally managed.
- It is possible to mitigate the effects of active probing by altering packet lengths or changing how servers respond to unauthenticated probes.

gfw.report@protonmail.com

https://gfw.report/publications/imc20/en/