OpenVPN is Open to VPN Fingerprinting

USENIX Security 2022

Diwen Xue, Reethika Ramesh, Arham Jain, Michalis Kallitsis
J. Alex Halderman, Jedidiah R. Crandall, Roya Ensafi

University of Michigan, Merit Network, Inc., Arizona State University

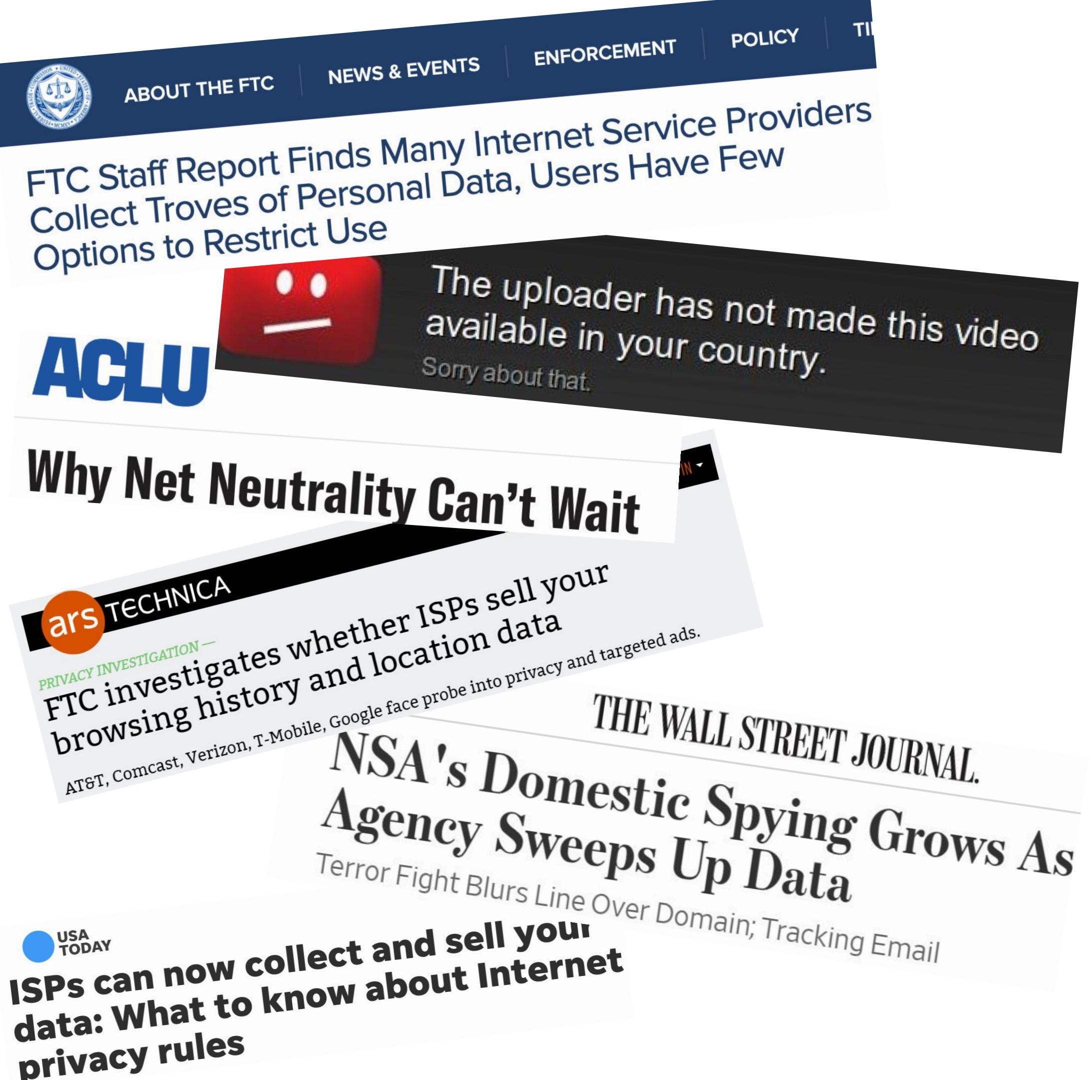

FTC Staff Report Finds Many Internet Service Providers Collect Troves of Personal Data, Users Have Few Options to Restrict Use

The uploader has not made this video available in your country.
Sorry about that.

ACLU

Why Net Neutrality Can't Wait

ars TECHNICA

PRIVACY INVESTIGATION —

FTC investigates whether ISPs sell your browsing history and location data

AT&T, Comcast, Verizon, T-Mobile, Google face probe into privacy and targeted ads.

THE WALL STREET JOURNAL.

NSA's Domestic Spying Grows As Agency Sweeps Up Data

Terror Fight Blurs Line Over Domain; Tracking Email

USA TODAY

ISPs can now collect and sell your data: What to know about Internet privacy rules

Internet traffic
is increasingly being **disrupted, tampered with, and monitored** by ISPs, advertisers, and other threat actors.

VPNs are on the Rise

"From 2010 to year-end 2019, the use of VPNs has **increased by approximately four times"**
Cybersecurity Company PC Matic, 2020

From Enterprise Security To Privacy and Censorship Circumvention

- Create private network across the public Internet through Encrypted Tunneling.
- Increasingly being used in non-enterprise setting.

OPENVPN

WIREGUARD

IKEv2 / IPSec

An Evolving Threat Model

- Most of past research focused on the **Integrity** and **Confidentiality** of the tunnel.
    - Tunnel Penetrating Attacks
    - Data Injection
    - Traffic Leaks
- Threat actors now attacking **Availability**.

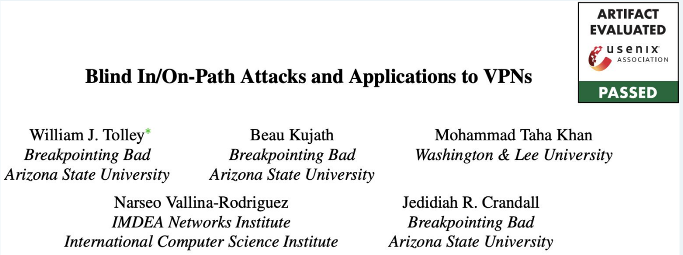

ARTIFACT EVALUATED usenix ASSOCIATION PASSED

Blind In/On-Path Attacks and Applications to VPNs

William J. Tolley*, *Breakpointing Bad*, *Arizona State University*
Beau Kujath, *Breakpointing Bad*, *Arizona State University*
Mohammad Taha Khan, *Washington & Lee University*
Narseo Vallina-Rodriguez, *IMDEA Networks Institute*, *International Computer Science Institute*
Jedidiah R. Crandall, *Breakpointing Bad*, *Arizona State University*

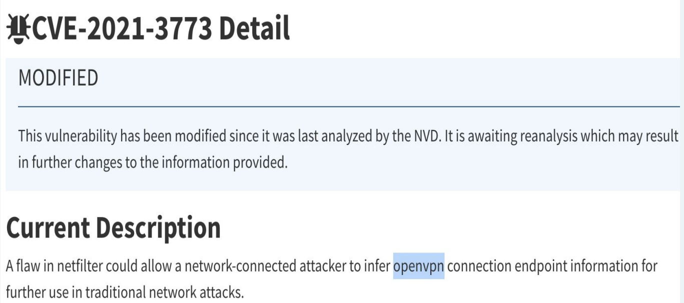

CVE-2021-3773 Detail

MODIFIED

This vulnerability has been modified since it was last analyzed by the NVD. It is awaiting reanalysis which may result in further changes to the information provided.

Current Description

A flaw in netfilter could allow a network-connected attacker to infer openvpn connection endpoint information for further use in traditional network attacks.

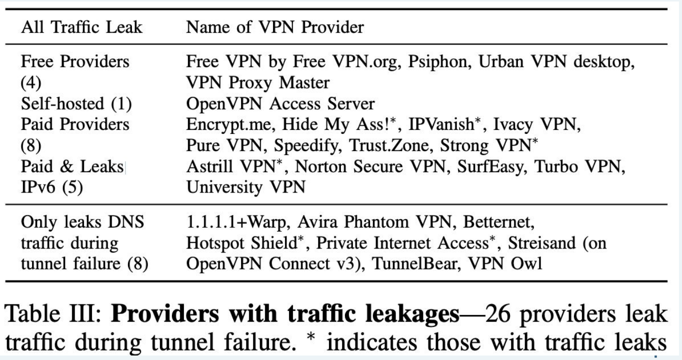

All Traffic Leak	Name of VPN Provider
Free Providers (4)	Free VPN by Free VPN.org, Psiphon, Urban VPN desktop, VPN Proxy Master
Self-hosted (1)	OpenVPN Access Server
Paid Providers (8)	Encrypt.me, Hide My Ass!*, IPVanish*, Ivacy VPN, Pure VPN, Speedify, Trust.Zone, Strong VPN*
Paid & Leaks IPv6 (5)	Astrill VPN*, Norton Secure VPN, SurfEasy, Turbo VPN, University VPN
Only leaks DNS traffic during tunnel failure (8)	1.1.1.1+Warp, Avira Phantom VPN, Betternet, Hotspot Shield*, Private Internet Access*, Streisand (on OpenVPN Connect v3), TunnelBear, VPN Owl

Table III: **Providers with traffic leakages**—26 providers leak traffic during tunnel failure. * indicates those with traffic leaks

Not a hypothetical threat...

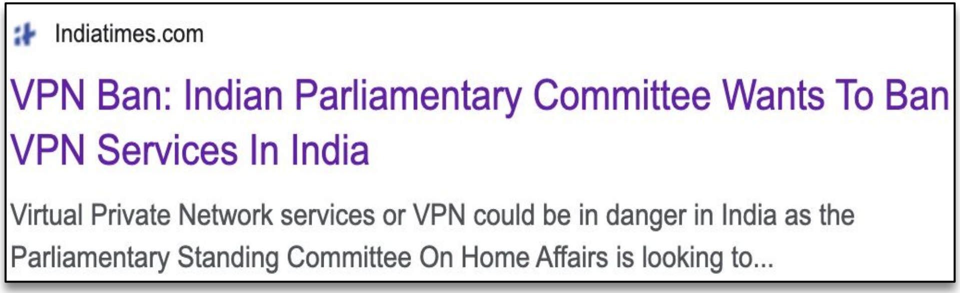

Indiatimes.com

VPN Ban: Indian Parliamentary Committee Wants To Ban VPN Services In India

Virtual Private Network services or VPN could be in danger in India as the Parliamentary Standing Committee On Home Affairs is looking to...

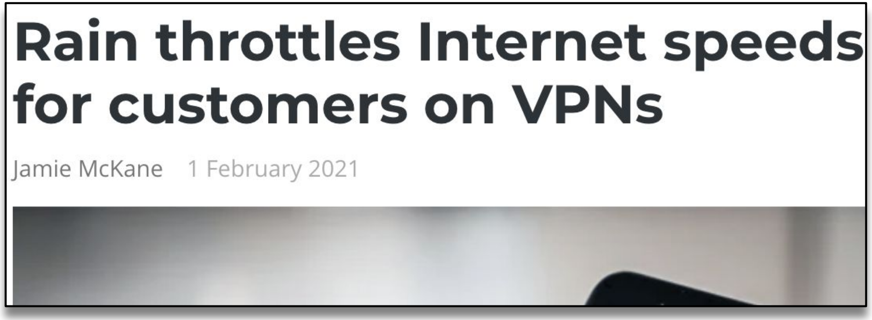

Rain throttles Internet speeds for customers on VPNs

Jamie McKane 1 February 2021

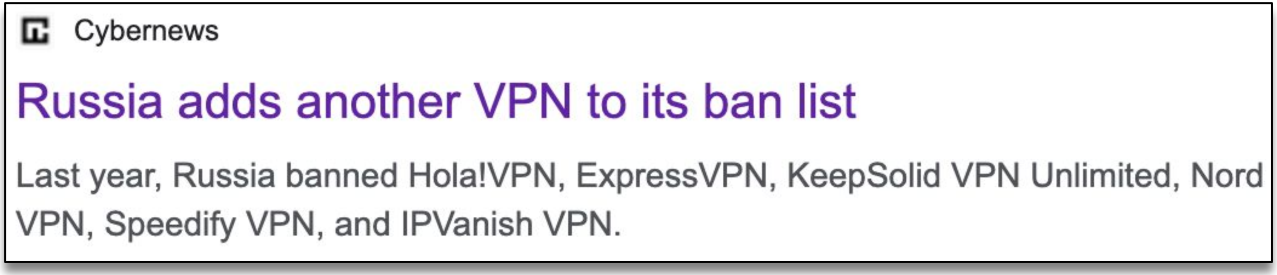

Cybernews

Russia adds another VPN to its ban list

Last year, Russia banned Hola!VPN, ExpressVPN, KeepSolid VPN Unlimited, Nord VPN, Speedify VPN, and IPVanish VPN.

“Obfuscated” VPN services

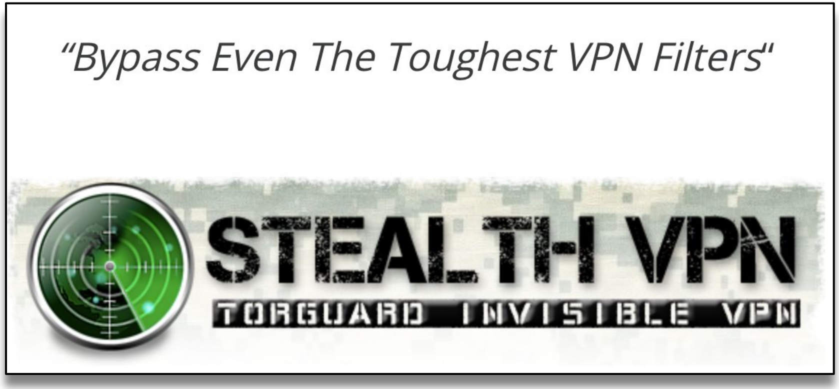

“Bypass Even The Toughest VPN Filters“

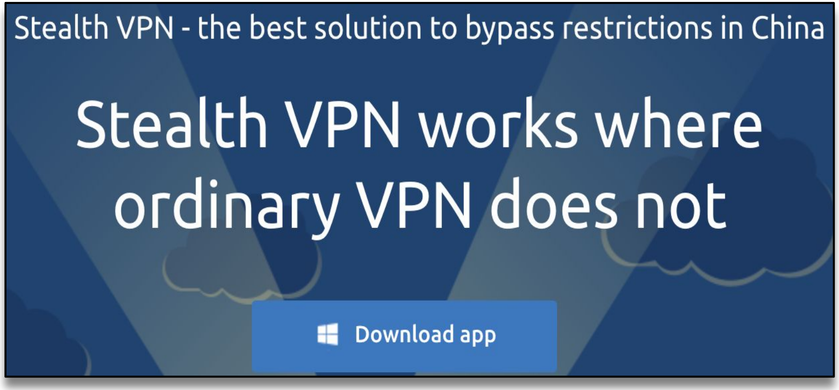

Stealth VPN - the best solution to bypass restrictions in China

Stealth VPN works where ordinary VPN does not

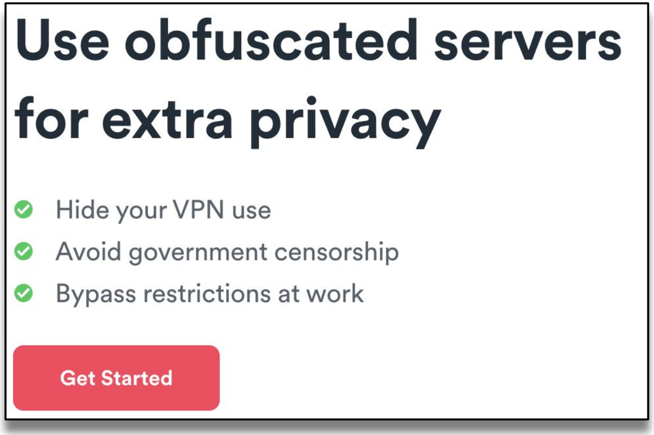

Use obfuscated servers for extra privacy

- ✓ Hide your VPN use
- ✓ Avoid government censorship
- ✓ Bypass restrictions at work

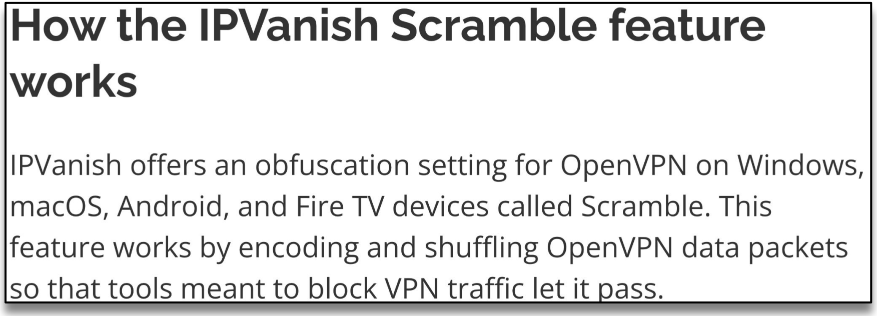

How the IPVanish Scramble feature works

IPVanish offers an obfuscation setting for OpenVPN on Windows, macOS, Android, and Fire TV devices called Scramble. This feature works by encoding and shuffling OpenVPN data packets so that tools meant to block VPN traffic let it pass.

Can ISPs and governments identify VPN traffic in near real time?

Can they do so at-scale, without incurring significant collateral damage from false positives?

We focus on OpenVPN and its variants!

The most popular protocol for commercial VPN services

“Obfuscated” VPN services built on top of OpenVPN

Mechanisms in place to impede fingerprinting attempts

Is OpenVPN Open to fingerprinting, *in practice*?

- Previous work used machine learning models on flow-level statistics
  - Connection duration
  - Inter-packet latency
  - Traffic symmetry

- Do these approaches work in practice?
  - Real-world ML-based censorship system not documented
  - Synthetic dataset, lab-based evaluations
  - Seemingly low false-positive can still be economically impractical

Effective investigation of Fingerprintability requires
not only to identify vulnerabilities,
but also to **demonstrate practical exploits**
under the constraints of
how ISPs and censors operate
in the real world.

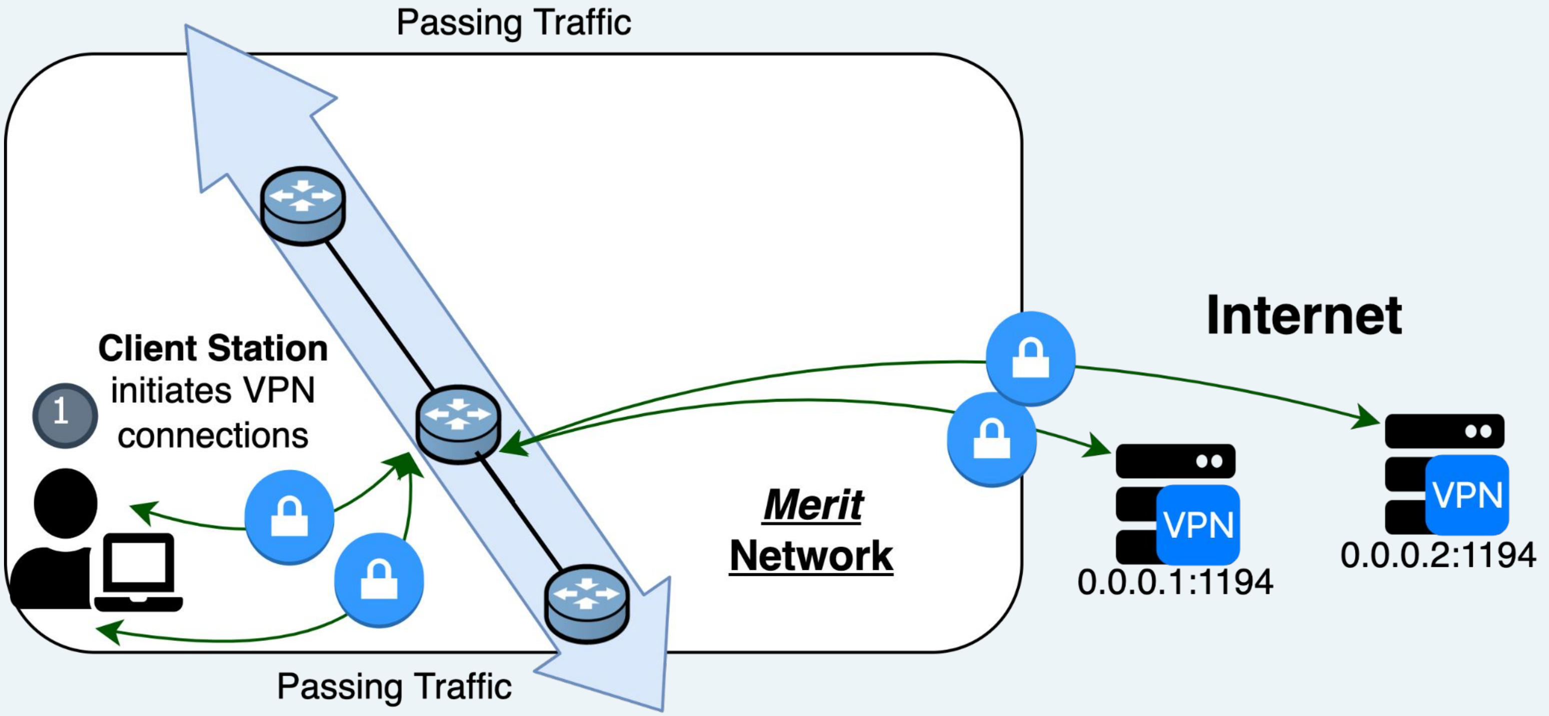

Deployment inside Merit Network

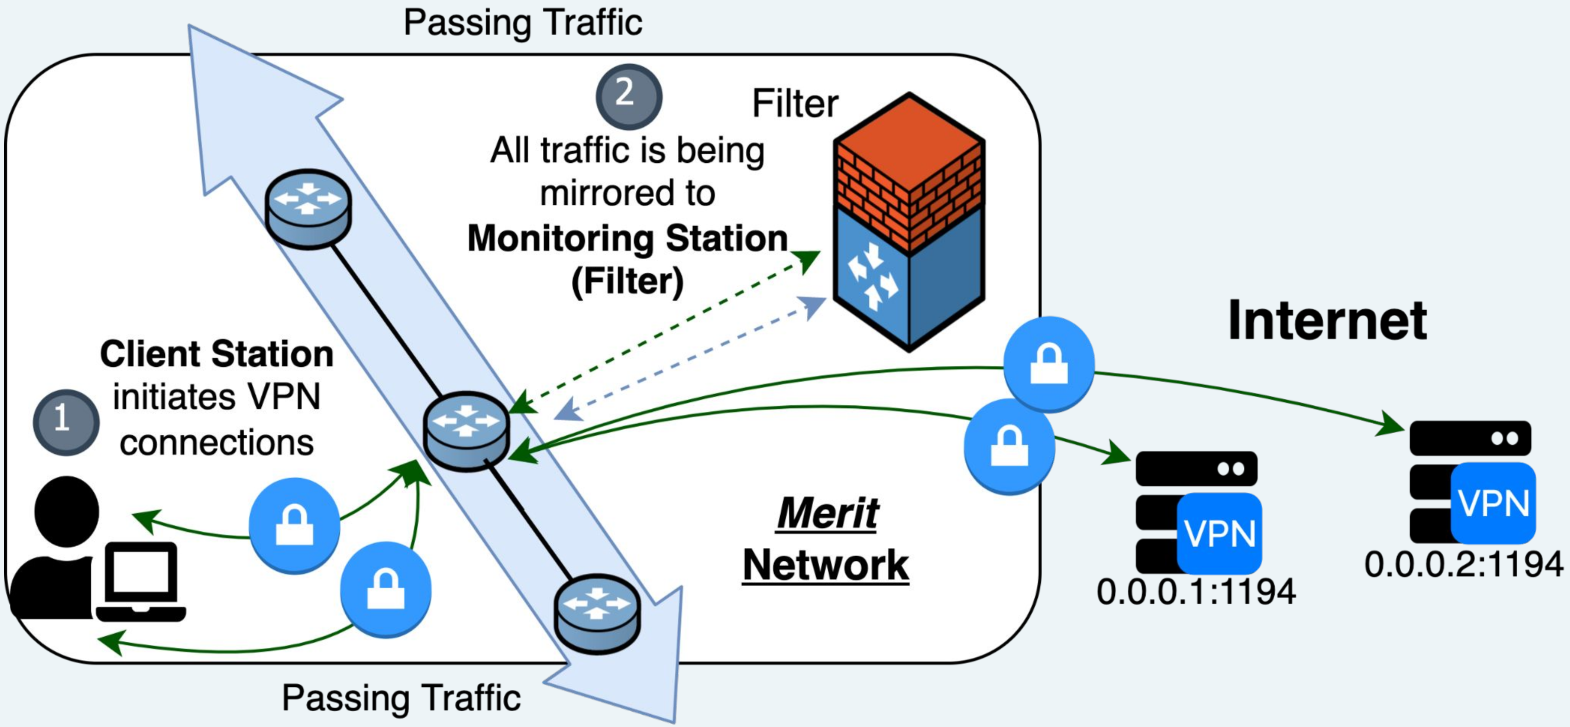
Passing Traffic
2
Filter
All traffic is being mirrored to
Monitoring Station (Filter)
Internet
Client Station initiates VPN connections
1
Merit
Network
VPN
VPN
0.0.0.1:1194
0.0.0.2:1194
Passing Traffic

Deployment inside Merit Network

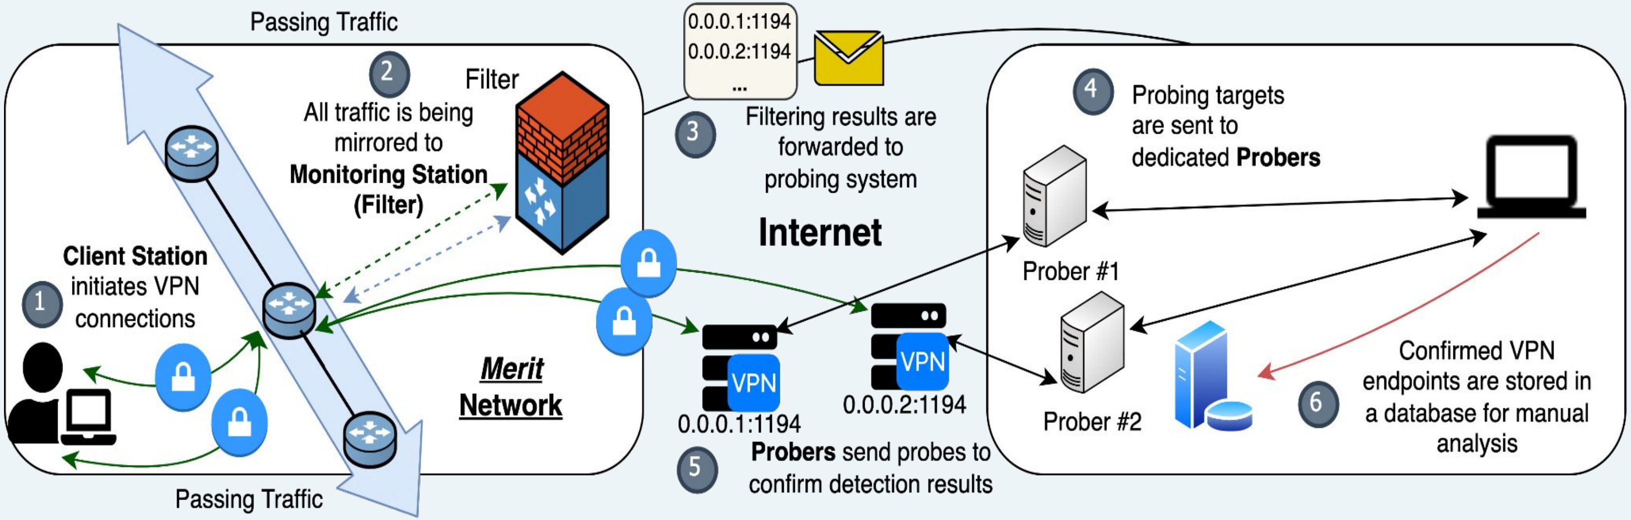
Passing Traffic
0.0.0.1:1194
0.0.0.2:1194
...
2
Filter
All traffic is being mirrored to **Monitoring Station (Filter)**
3
Filtering results are forwarded to probing system
4
Probing targets are sent to dedicated **Probers**
Internet
Client Station initiates VPN connections
1
Prober #1
Merit Network
VPN
VPN
0.0.0.1:1194
0.0.0.2:1194
Prober #2
Confirmed VPN endpoints are stored in a database for manual analysis
6
5
Probers send probes to confirm detection results
Passing Traffic

Deployment inside Merit Network

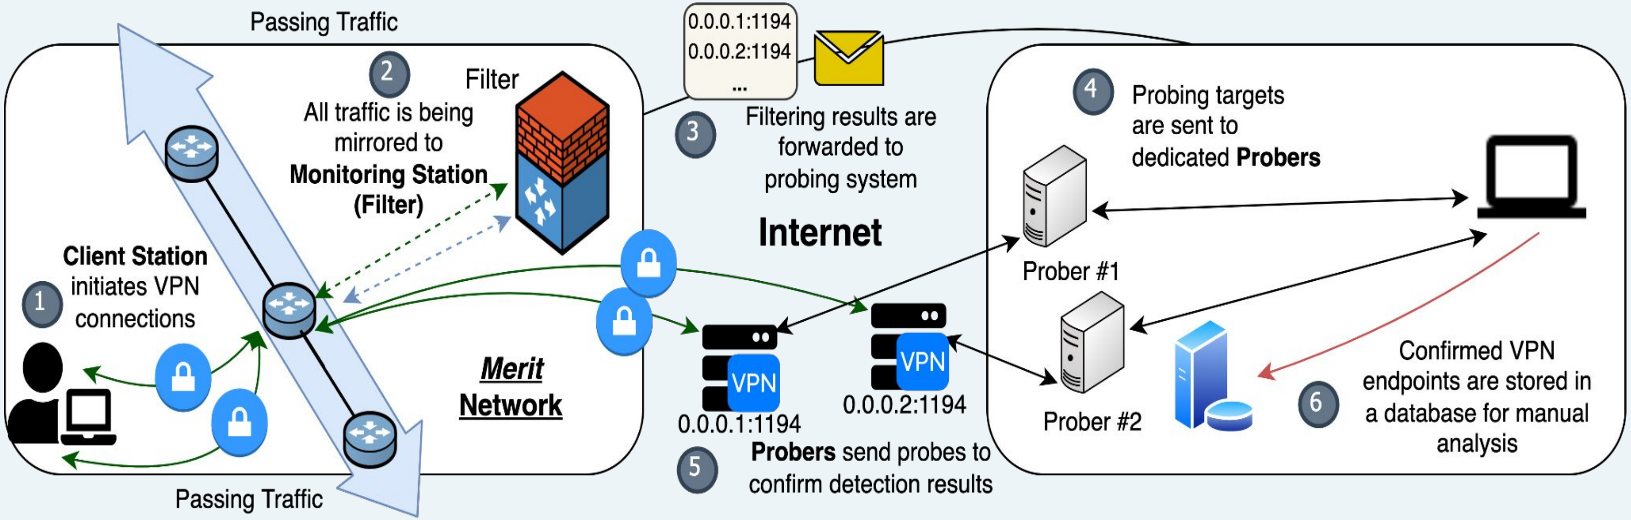
Passing Traffic
0.0.0.1:1194
0.0.0.2:1194
...
2
Filter
All traffic is being mirrored to **Monitoring Station (Filter)**
3
Filtering results are forwarded to probing system
4
Probing targets are sent to dedicated **Probers**
Internet
Prober #1
Client Station initiates VPN connections
1
***Merit* Network**
VPN
VPN
0.0.0.1:1194
0.0.0.2:1194
Prober #2
Confirmed VPN endpoints are stored in a database for manual analysis
6
5
Probers send probes to confirm detection results
Passing Traffic

Deployment inside Merit Network

Examining how the Great Firewall of China discovers hidden circumvention servers. IMC'15
Analyzing China's blocking of unpublished Tor bridges. FOCI'18
How China detects and blocks Shadowsocks. IMC'20

Fingerprinting OpenVPN

Filtering Phase:

- **Opcode Evolution** (Byte Pattern)
- **ACK Repetition** (Packet Size)

Probing Phase:

- **Customized Probes**

  (Server Behaviors)

Fingerprint 1: Opcode

- Opcode is a fixed value in the header which denotes each stage of the session
- Opcode evolution of a new OpenVPN session is unique and can be used to fingerprint OpenVPN.
- Flexible enough to catch certain “obfuscated” variants.

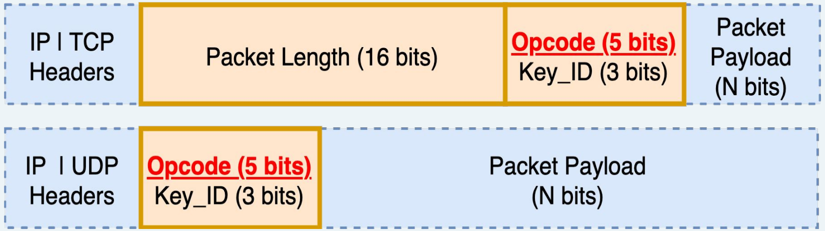

Opcode message types:

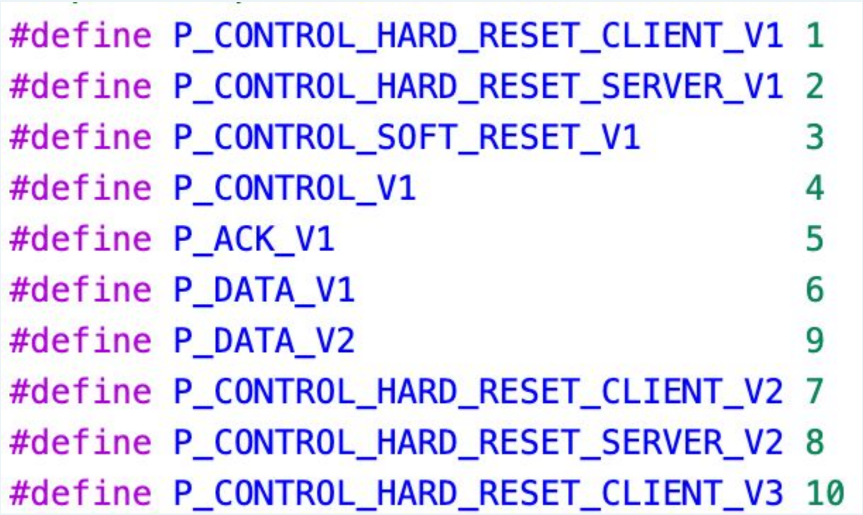

```
#define P_CONTROL_HARD_RESET_CLIENT_V1 1
#define P_CONTROL_HARD_RESET_SERVER_V1 2
#define P_CONTROL_SOFT_RESET_V1        3
#define P_CONTROL_V1                   4
#define P_ACK_V1                       5
#define P_DATA_V1                      6
#define P_DATA_V2                      9
#define P_CONTROL_HARD_RESET_CLIENT_V2 7
#define P_CONTROL_HARD_RESET_SERVER_V2 8
#define P_CONTROL_HARD_RESET_CLIENT_V3 10
```

Fingerprint 2: ACK Packets

- Explicit acknowledgement and retransmission model for "control" messages.
- Uniform in size for each session; not the same as TCP ACK flag;
- Quantify "ACK Fingerprint" as a set of threshold-based detection rules.

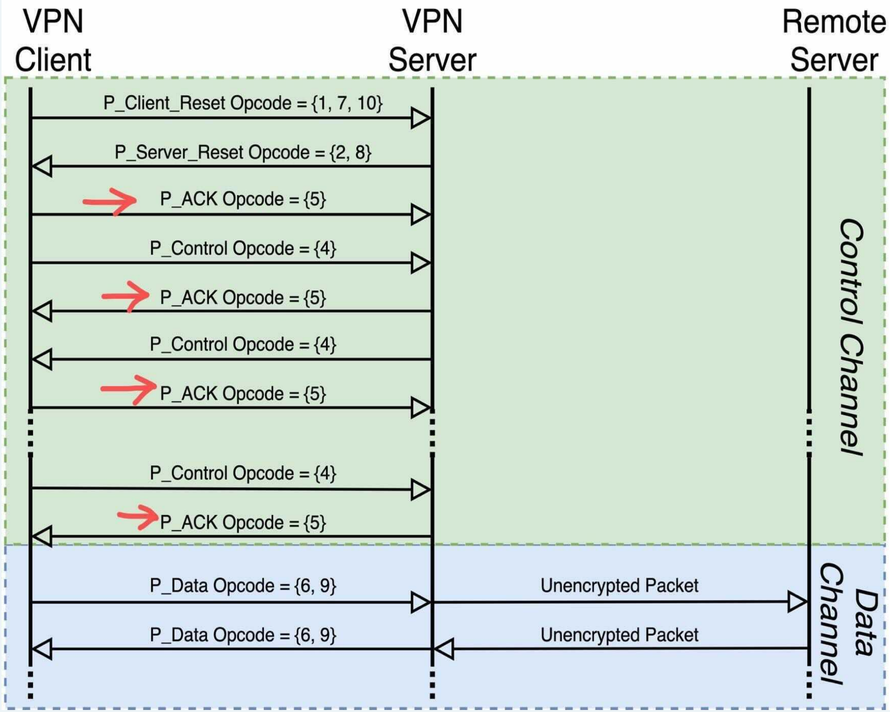

Detection accuracy of Filtering phase

Filtering Phase:

- **Opcode Evolution** (Byte Pattern)
- **ACK Repetition** (Packet Size)

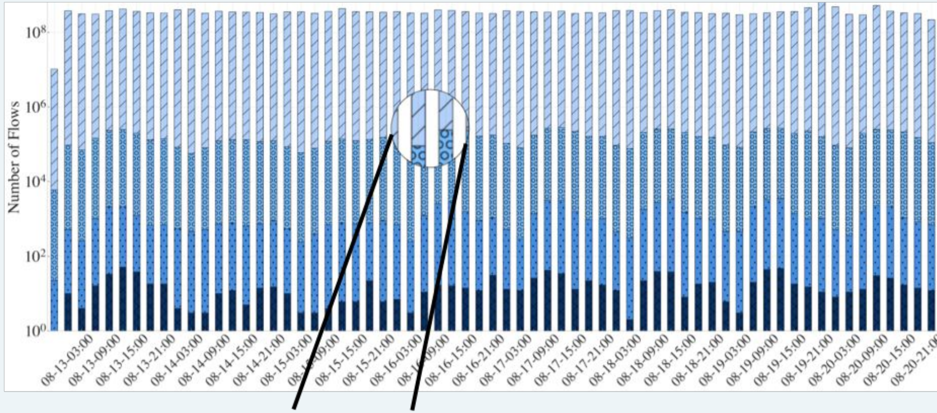

🟦	All Flows	3.3 Billion
🟦	Persistent Flows	1.9 Million
🟦	Filter Outputs	15835
🟦	Prober Outputs	519

Detection accuracy of Filtering phase

Filtering Phase:

- **Opcode Evolution** (Byte Pattern)
- **ACK Repetition** (Packet Size)

Increasing accuracy to prevent significant collateral damage requires **active probing**

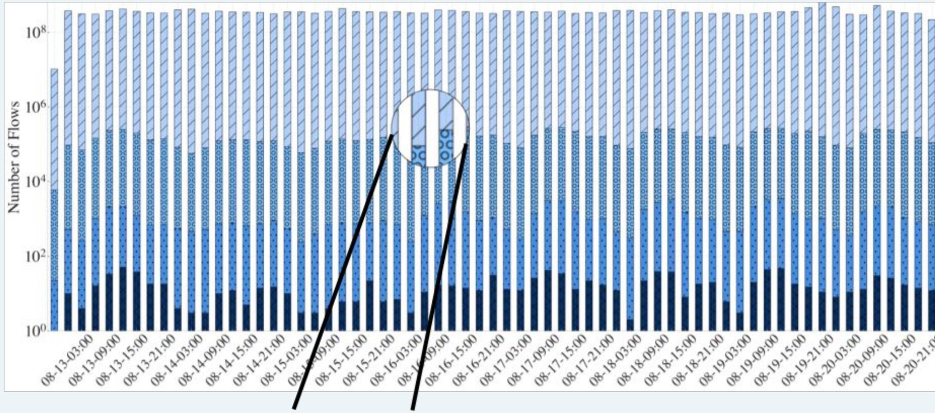

All Flows	3.3 Billion
Persistent Flows	1.9 Million
Filter Outputs	15835
Prober Outputs	519

Active Probing

- Defense mechanisms “*tls-auth*” and “*tls-crypt*” enable a firewall-like protection.
    - ○ OpenVPN remains silent until the client proves knowledge of a shared secret.
- Application may stay silent, but application-specific behaviors can still be observed at network level.

(related: *Detecting Probe-resistant Proxies* NDSS’20)

Our customized probes:

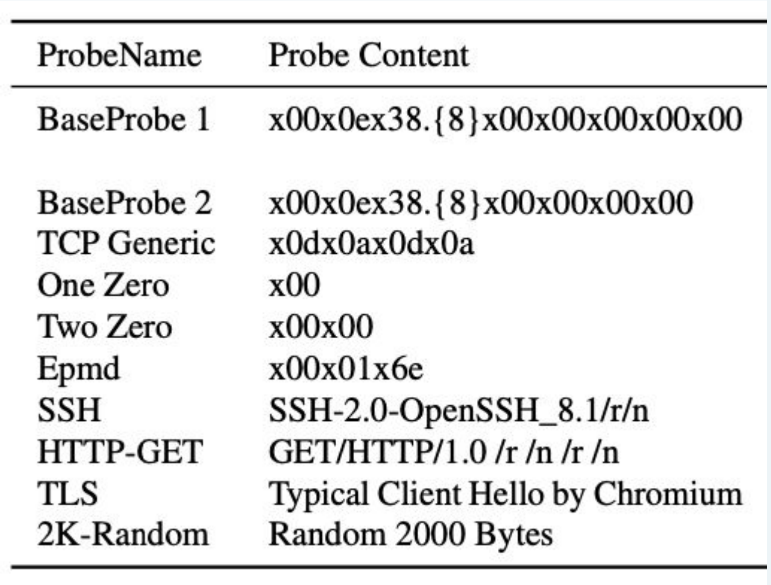

ProbeName	Probe Content
BaseProbe 1	x00x0ex38.{8}x00x00x00x00x00
BaseProbe 2	x00x0ex38.{8}x00x00x00x00
TCP Generic	x0dx0ax0dx0a
One Zero	x00
Two Zero	x00x00
Epmd	x00x01x6e
SSH	SSH-2.0-OpenSSH_8.1/r/n
HTTP-GET	GET/HTTP/1.0 /r /n /r /n
TLS	Typical Client Hello by Chromium
2K-Random	Random 2000 Bytes

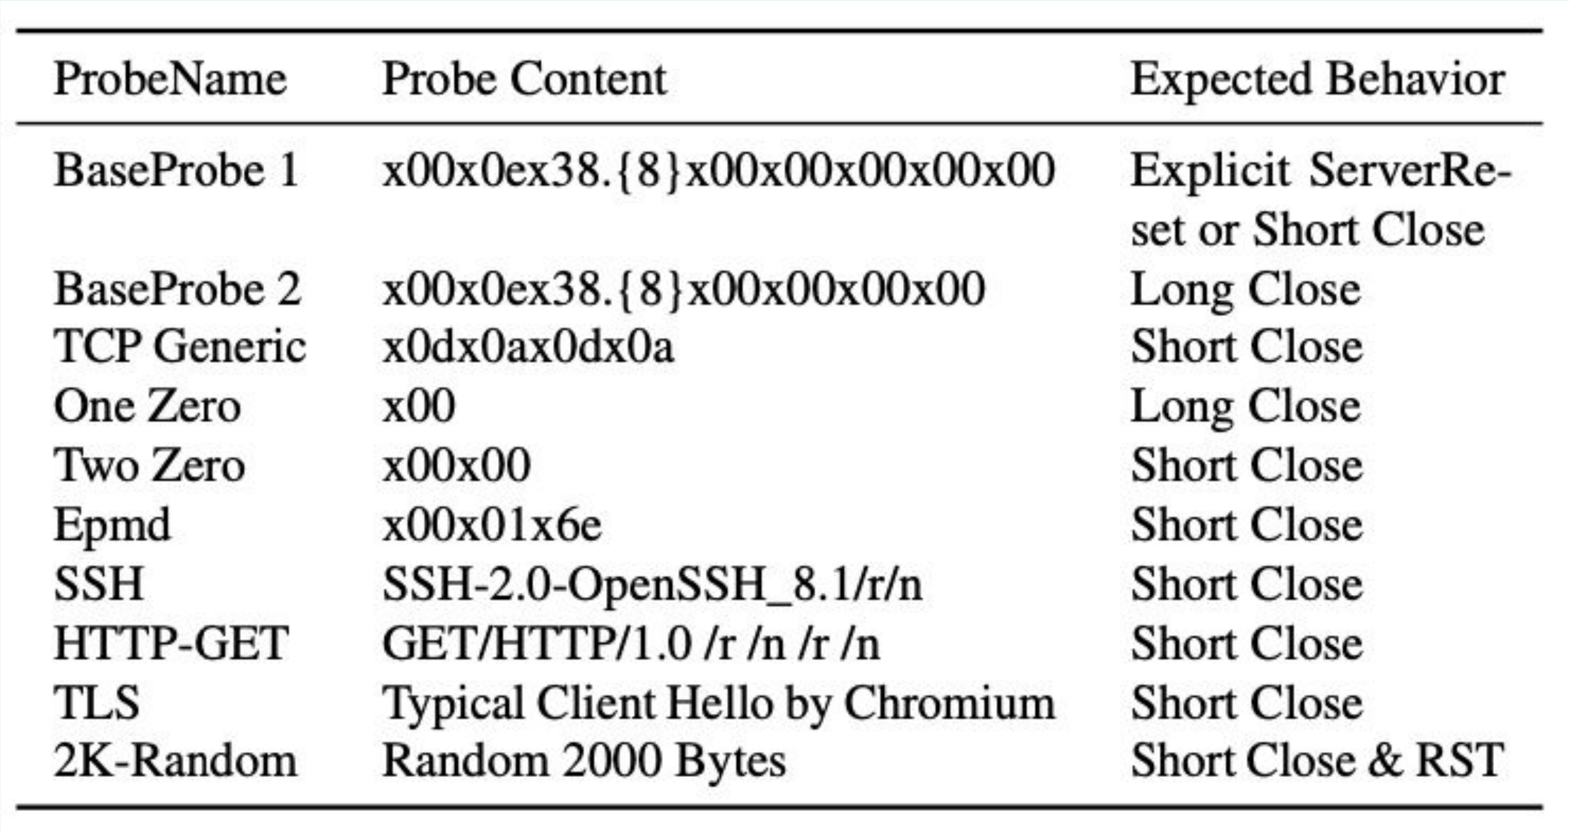

ProbeName	Probe Content	Expected Behavior
BaseProbe 1	x00x0ex38.{8}x00x00x00x00x00	Explicit ServerReset or Short Close
BaseProbe 2	x00x0ex38.{8}x00x00x00x00	Long Close
TCP Generic	x0dx0ax0dx0a	Short Close
One Zero	x00	Long Close
Two Zero	x00x00	Short Close
Epmd	x00x01x6e	Short Close
SSH	SSH-2.0-OpenSSH_8.1/r/n	Short Close
HTTP-GET	GET/HTTP/1.0 /r /n /r /n	Short Close
TLS	Typical Client Hello by Chromium	Short Close
2K-Random	Random 2000 Bytes	Short Close & RST

Testing on Commercial VPNs

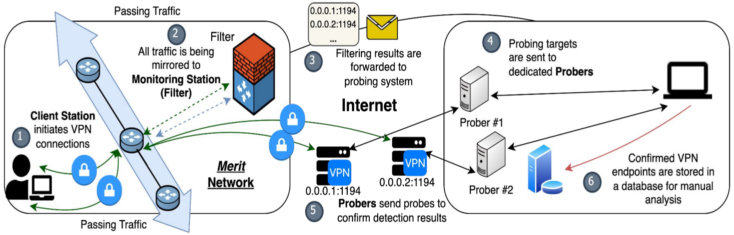
Passing Traffic
2
Filter
All traffic is being mirrored to
Monitoring Station (Filter)
Client Station initiates VPN connections
1
Merit Network
Passing Traffic
0.0.0.1:1194
0.0.0.2:1194
...
3
Filtering results are forwarded to probing system
Internet
4
Probing targets are sent to dedicated Probers
Prober #1
Prober #2
VPN
VPN
0.0.0.1:1194
0.0.0.2:1194
5
Probers send probes to confirm detection results
Confirmed VPN endpoints are stored in a database for manual analysis
6

- Effective in detecting vanilla OpenVPN flows. (39/40 vanilla configurations)

- **Surprisingly, 72.67% obfuscated flows also detected.** (34/41 obfuscated configurations).
  - “Obfuscated” VPN services use OpenVPN as backbone protocol
  - Insufficient obfuscation failing to mask fingerprints.

Fingerprinting “Obfuscated” VPNs

XOR Obfuscation

Additional Encrypted Tunneling

Obfuscated Servers

Fingerprinting “Obfuscated” VPNs

XOR Obfuscation

1:1 correspondence between opcodes and ciphertext

Additional Encrypted Tunneling

Lack of random padding

Obfuscated Servers

Co-location of Bridges and vanilla servers.

XOR Obfuscation

- Unofficial patch that scrambles payloads by a series of XOR operations.
- **Opcode excluded from reversal, therefore always mapped to the same ciphertext.** Behavior preserved in multiple implementations.

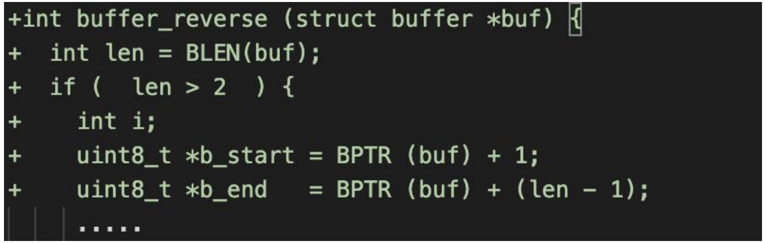

```
+int buffer_reverse (struct buffer *buf) {
+  int len = BLEN(buf);
+  if (  len > 2  ) {
+    int i;
+    uint8_t *b_start = BPTR (buf) + 1;
+    uint8_t *b_end   = BPTR (buf) + (len - 1);
     .....
```

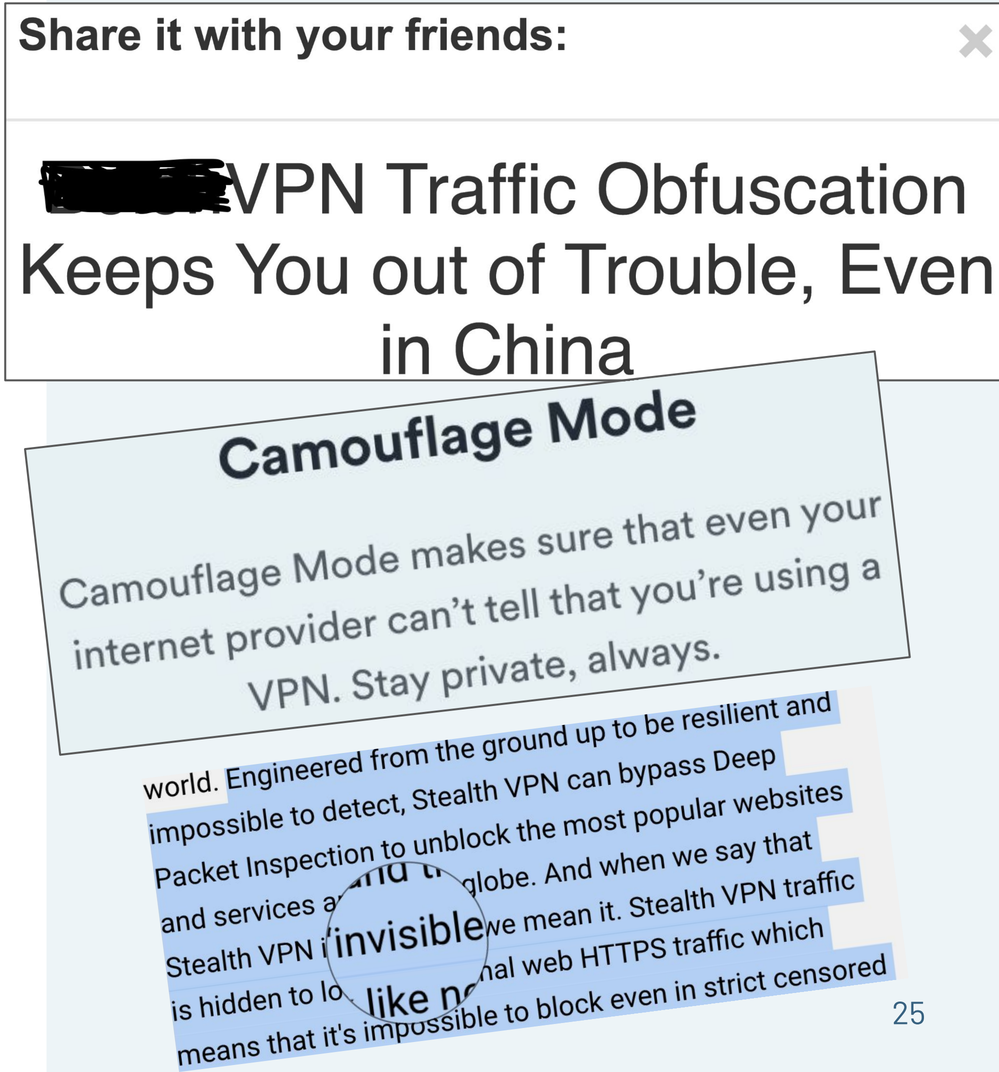

Share it with your friends:

[redacted] VPN Traffic Obfuscation Keeps You out of Trouble, Even in China

Camouflage Mode

Camouflage Mode makes sure that even your internet provider can't tell that you're using a VPN. Stay private, always.

world. Engineered from the ground up to be resilient and impossible to detect, Stealth VPN can bypass Deep Packet Inspection to unblock the most popular websites and services and t globe. And when we say that Stealth VPN i invisible we mean it. Stealth VPN traffic is hidden to lo like n nal web HTTPS traffic which means that it's impossible to block even in strict censored

Accuracy

- Collateral damage as the fundamental measure of practicality.
  - Week-long evaluation, aggregated 20 Gbps of mirrored traffic.
  - 3,638 flows flagged. (0.0039%)
  - Manual analysis found supporting evidence for 90% of flagged connections.

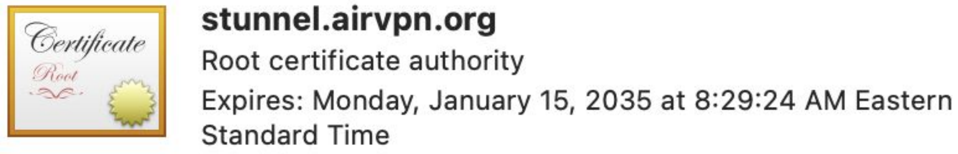

stunnel.airvpn.org

Root certificate authority

Expires: Monday, January 15, 2035 at 8:29:24 AM Eastern Standard Time

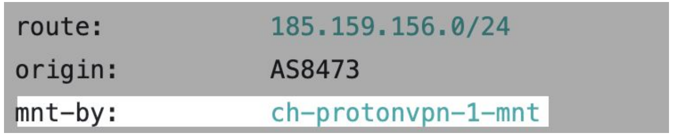

```
route:          185.159.156.0/24
origin:         AS8473
mnt-by:         ch-protonvpn-1-mnt
```

Conclusion

- Fingerprinting OpenVPN is within the reach of any network operator.
  - Even with obfuscation patches deployed in the wild.
  - Risk of throttling, blocking, and even follow-up attacks on VPN tunnel.
  - Users should *NOT* expect unobservability, even with “stealth” VPN.

Conclusion

- Fingerprinting OpenVPN is within the reach of any network operator.
    - Even with obfuscation patches deployed in the wild.
    - Risk of throttling, blocking, and even follow-up attacks on VPN tunnel.
    - Users should *NOT* expect unobservability, even with “stealth” VPN.
- Moving forward…
    - Short-term defense.
    - A gap between obfuscation research and implementation.

OpenVPN is Open to VPN Fingerprinting

USENIX Security 2022

VPNalyzer.org

Diwen Xue, Reethika Ramesh, Arham Jain, Michalis Kallitsis
J. Alex Halderman, Jedidiah R. Crandall, Roya Ensafi

University of Michigan, Merit Network, Inc., Arizona State University

Backup

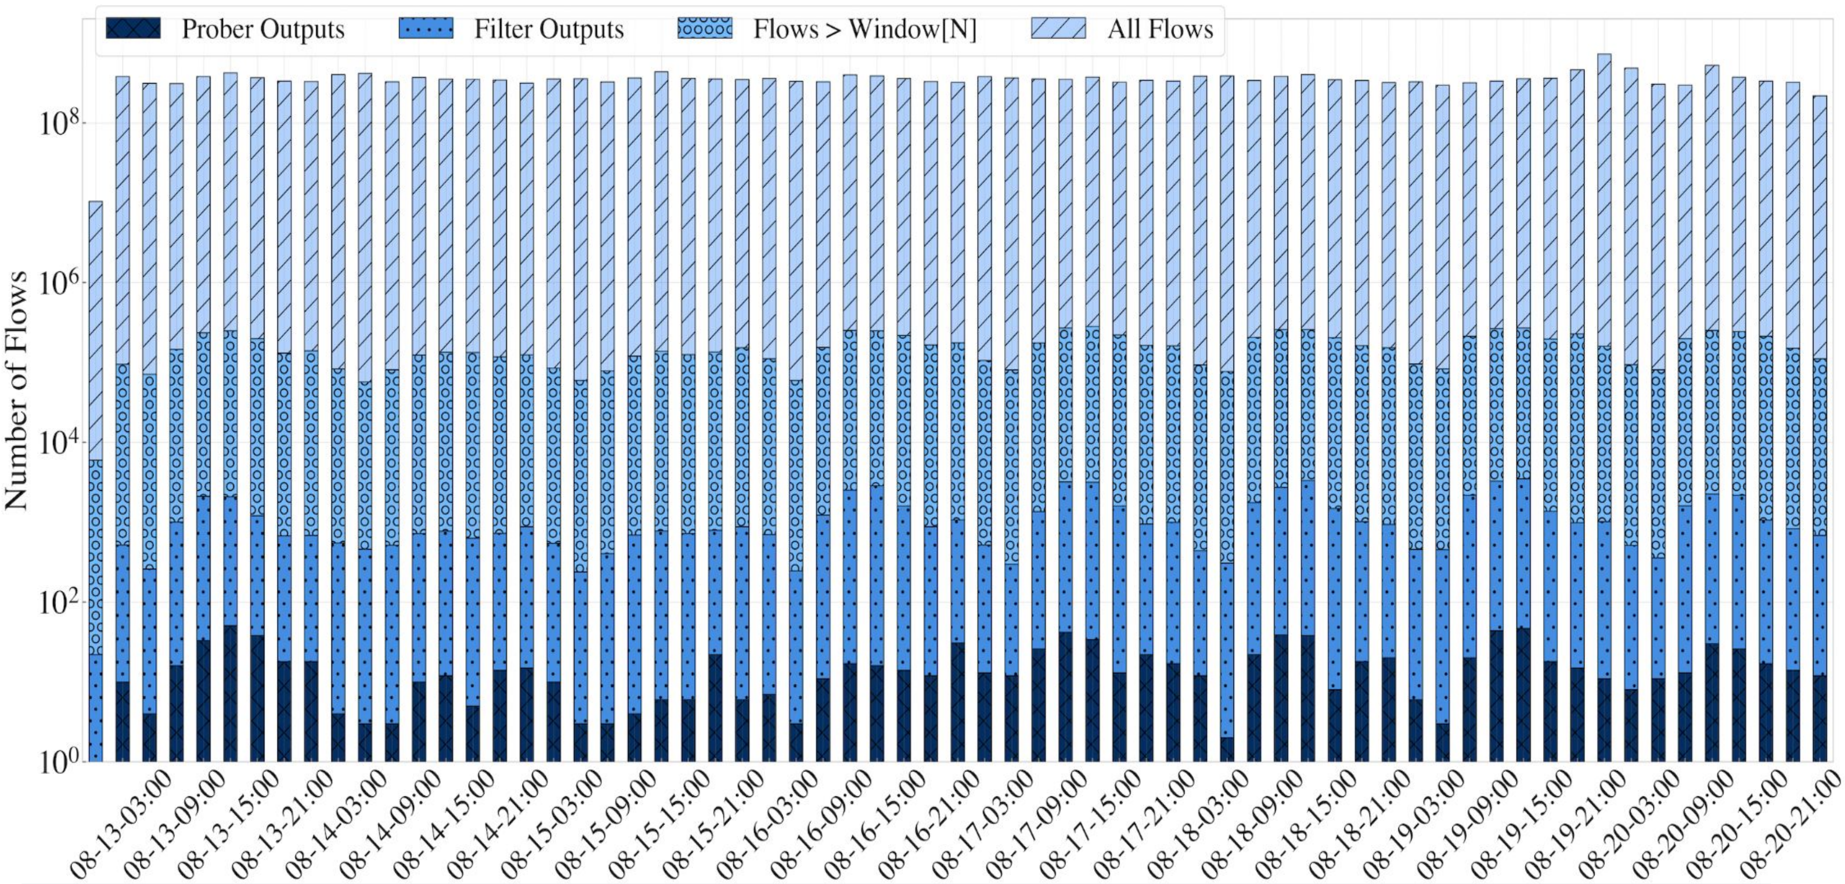
Prober Outputs
Filter Outputs
Flows > Window[N]
All Flows
Number of Flows
10^8
10^6
10^4
10^2
10^0
08-13-03:00
08-13-09:00
08-13-15:00
08-13-21:00
08-14-03:00
08-14-09:00
08-14-15:00
08-14-21:00
08-15-03:00
08-15-09:00
08-15-15:00
08-15-21:00
08-16-03:00
08-16-09:00
08-16-15:00
08-16-21:00
08-17-03:00
08-17-09:00
08-17-15:00
08-17-21:00
08-18-03:00
08-18-09:00
08-18-15:00
08-18-21:00
08-19-03:00
08-19-09:00
08-19-15:00
08-19-21:00
08-20-03:00
08-20-09:00
08-20-15:00
08-20-21:00

Encrypted Tunneling

- Tunnel-based obfuscation wraps OpenVPN traffic through encryption.
    - SSL/SSH Tunnel, obfs234 ...
- ACK fingerprints are still observable outside **tunnels that lack random padding**.

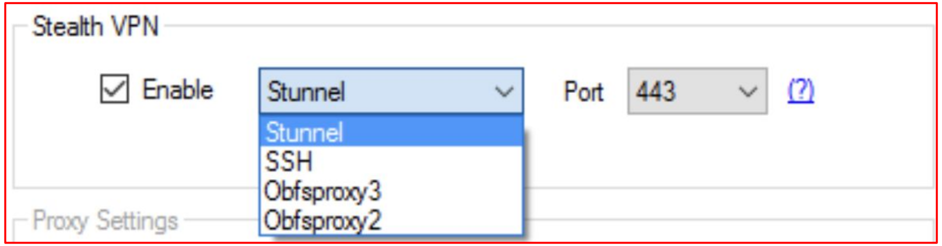
Stealth VPN

☑ Enable

Stunnel

Port 443

Obfuscation Servers

- In practice, most of obfuscation servers — “Bridges” — are co-located with vanilla TCP servers. (34/41 for /29 subnet)
- Infrastructures are shared between obfuscated and vanilla services from different providers.

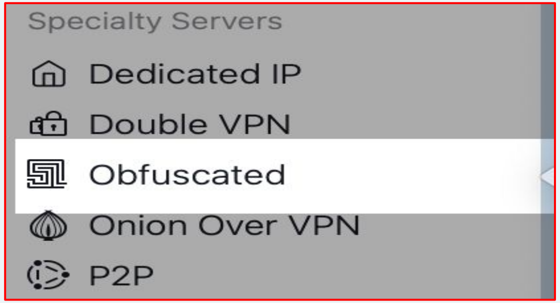

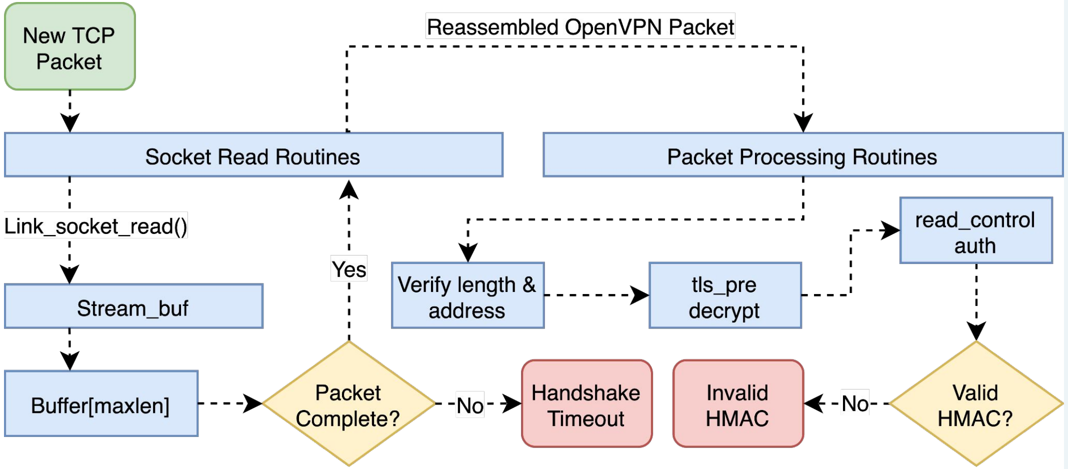

New TCP Packet

Socket Read Routines

Reassembled OpenVPN Packet

Packet Processing Routines

Link_socket_read()

Stream_buf

Buffer[maxlen]

Packet Complete?

Yes

No

Handshake Timeout

Verify length & address

tls_pre decrypt

read_control auth

Valid HMAC?

No

Invalid HMAC

Probe 1 & Probe 2

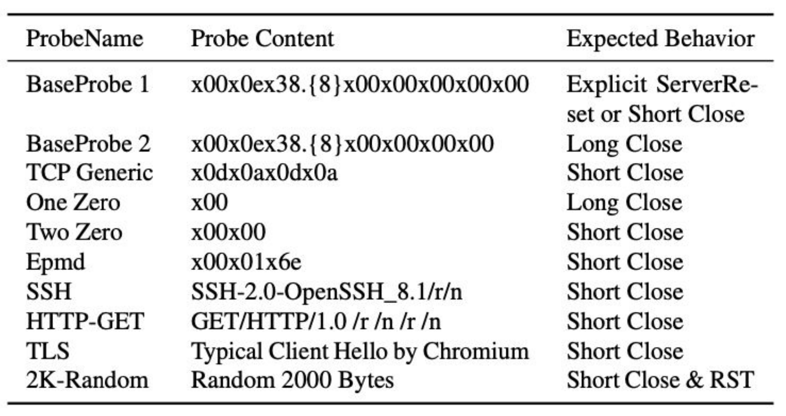

ProbeName	Probe Content	Expected Behavior
BaseProbe 1	x00x0ex38.{8}x00x00x00x00x00	Explicit ServerReset or Short Close
BaseProbe 2	x00x0ex38.{8}x00x00x00x00	Long Close
TCP Generic	x0dx0ax0dx0a	Short Close
One Zero	x00	Long Close
Two Zero	x00x00	Short Close
Epmd	x00x01x6e	Short Close
SSH	SSH-2.0-OpenSSH_8.1/r/n	Short Close
HTTP-GET	GET/HTTP/1.0 /r /n /r /n	Short Close
TLS	Typical Client Hello by Chromium	Short Close
2K-Random	Random 2000 Bytes	Short Close & RST

Is OpenVPN Open to fingerprinting, in practice?

Real-world ML-based censorship system not documented

Synthetic dataset, lab-based evaluation

Same dataset *ISCXVPN2016* [3,14,15,17,24,26,68]

Seemingly low false positive rate can be misleading.

(1% FPR, 0.01% Base Rate 1 in 100 blocked is actually VPN)