DOI: 10.1145/3657629

BY GENE TSUDIK

Tracing the Network Traffic Fingerprinting Techniques of OpenVPN

It is well known that a fundamental conflict exists between various entities (for example, ISPs, corporations, and governments) that want to control and/or monitor Internet activity and individual users (or groups thereof) who desire privacy and/or censorship circumvention. This conflict manifests itself as a global and seemingly never-ending arms race between privacy technologies (for example, Tor and VPNs) and traffic analysis methods that aim to identify (fingerprint) and block communication conducted using the former.

This is not a universal battle between good and evil. There are numerous examples where upstanding users (for example, citizen journalists, dissidents, whistle-blowers) fight the proverbial "good fight" to circumvent oppression. However, it is just as simple to turn the table and consider settings where malicious actors (for example, trolls, spies, terrorists, criminals) are pitted against "righteous" entities that attempt to prevent their activities. Also, sometimes the motivation for blocking cloaked traffic is not nefarious but is simply driven by the need to maintain acceptable QoS for other traffic.

Security researchers are not supposed to take sides in this battle. It is equally legitimate to explore privacy and anti-privacy techniques as well as to attack either. Both sides in the conflict must be aware of flaws and strengths of the tools they use. To this end, the accompanying paper investigates network traffic fingerprinting (simply *fingerprinting* hereafter) of a very popular privacy technique—OpenVPN.

Fingerprinting is the art of probabilistically identifying—ideally, with low error rates—traffic patterns that correspond to a particular targeted activity, VPN use in this case. As the paper illustrates, OpenVPN is susceptible to quite accurate fingerprinting via a two-stage process: passive traffic analysis (Filter), followed by active probing (Prober). It reports >85% success rate in identifying OpenVPN connections. Furthermore, even when coupled with an optional obfuscation layer (for example, Chameleon or Stealth), OpenVPN traffic remains detectable. These results were obtained in partnership with Merit, a Michigan-based ISP that serves about one million users. Such a partnership is necessary for any credible study/experiment of real-world traffic analysis since the experimenters must play the combined role of the ISP and the censor.

Interestingly, the Deep Packet Inspection (DPI) approach taken in this work is inspired by the infamous Great Firewall of China. Although DPI is well-known and widely used, its features and accuracy grew and improved over the years. The main issue is whether it can be used in real time and at scale. The authors show it can. Moreover, the proposed method outperforms prior ML-based techniques in terms of FPR.

One thorny issue in conducting this type of a study is ethics. Genuine VPN (especially obfuscated) traffic is, by its very nature, sensitive and any information collected about it must be treated with utmost care. The authors tried their best to make sure that no information (beyond server IP addresses and `port #-s`) is kept. Furthermore, the study must minimize influencing ISP traffic, which is achieved via passive logging. The only potential interference is the subsequent probing of suspected VPN servers. However, the authors claim: "Each server receives only 2–10 innocuous connection attempts, similar to those commonly used in Internet measurement tools like Nmap."

While it is debatable whether this is truly negligible or innocuous, any study with an active component must involve some probing.

The probing process first attempts "base probes" that identify possible OpenVPN traffic and follows up with additional probes that rule out certain other protocols that behave similarly to OpenVPN. The latter cleverly uses long(er) RST packets to test for RST thresholds unique to OpenVPN servers.

To conclude, designing VPN software is not rocket science. It requires carefully following accepted key management and secure communication techniques. In contrast, designing an effective VPN obfuscation layer is much more difficult, as the paper illustrates. The arms race between obfuscation and counter-obfuscation techniques is on-going and is unlikely to end soon. This is the bad news. However, the authors point out: "...We evaluated the practicality of our approach in partnership with a mid-size ISP, and we were able to identify *the majority* of vanilla and obfuscated OpenVPN flows with only negligible false positives..."

The majority does not mean all: for five OpenVPN configurations (Table 5 in the original paper), a detection rate of 0 (zero) was reported, which gives some hope. Interestingly, three of those use obfs4, which employs stronger padding (by randomizing packet sizes) than that its predecessors obfs2/3 used by other providers.[a]

This paper amply demonstrates formidable challenges facing successful obfuscation of VPN traffic and outlines several directions for both near-term and long-term mitigation measures. However, it is highly unlikely that any panacea will materialize in the near term.

Gene Tsudik is a Distinguished Professor of Computer Science at the University of California, Irvine, CA, USA.

---

[a] Diwen Xue, private communication.