Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 2 January 2026

Y(J). Stein
Allot

1 July 2025

ECH Considered Harmful
draft-stein-tls-ech-considered-harmful-00

Abstract

Encrypted Client Hello is designed to enhance personal privacy, in particular obstructing the ability of communications service providers (but not Over The Top service providers) to map packet flows to applications. While mostly ineffective in attaining this goal, it does severely hamper network-based detection of malicious flows, thus exposing end-users to various security risks that were previously avoidable.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 2 January 2026.

Copyright Notice

Table of Contents

1. Introduction

Transport Level Security (TLS) 1.3 [RFC8446] secures communications between clients and servers, providing integrity, confidentiality and server authentication. However, certain cleartext extensions in the TLS Client Hello (CH) message actually facilitate application identification by third parties. This capability may be abused for pervasive monitoring [RFC7258].

Encrypted Client Hello (ECH) [I-D.ietf-tls-esni] has been proposed to remedy the perceived privacy infringement inherent in the TLS CH, in order to alleviate its pervasive monitoring aspects. However, RFC7258 recognizes that there are forms of monitoring that can be beneficial and not part of any attack, e.g., network management functions and malicious packet detection.

In this memo we explore precisely these two issues. While we support efforts to protect personal privacy and mitigate pervasive monitoring attacks, we will show that ECH is misguided in its rationale, largely ineffective in attaining its aims, and is harmful due to exposing clients to avoidable security risks.

2. ECH is misguided

The rationale behind ECH is the hiding of a client’s behavior, and in particular which servers a given client accesses, from the prying eyes of surveillance. However, the mechanism only impacts surveillance along the path between client and server; it obviously does not thwart information gathering by the Over The Top (OTT) entity operating the server itself.

OTT service providers routinely collect personal information of clients accessing their servers. In fact, their business model is based on knowing their clients better than these clients know themselves. Client personal information is packaged and monetized, and was indeed the main source of personal information exploited by governmental agencies. ECH plugs a pin-hole leak in an open-ended pipe.

On the other hand, Internet Service Providers (ISPs) need to identify the applications to which their client’s packet flows belong in order to economically guarantee a level of Quality of Experience (QoE). ISPs may be under contractual obligation to deliver some minimal QoE, and even if this level is not explicitly specified, underperformance will lead to unsatisfied clients and subsequent churn. Throwing additional bandwidth at this problem does not always ameliorate the problem, and even when it does that entails a price.

Many ISPs use smart traffic management based on application classification in order to optimally forward each packet - prioritizing packets requiring expedited delivery, delaying time insensitive packets, and under conditions of extreme congestion discarding packets belonging to rate-adaptive applications. Furthermore, ISPs may provide network-based security services that block malicious traffic. These use-cases require identifying the application but not the client, and thus do not involve personal information from a particular client.

Some use-cases do require identifying clients. Specific clients may have specific policies regarding applications or application classes; for example, first responders will generally be given priority over other clients. Another case is identifying and disconnecting clients belonging to botnets.

Unlike OTTs, even in these cases ISPs do not monetize client personal information, and thus have little incentive to collect and store flow-to-application mappings. Thus, ECH penalizes ISPs that need application classification in order to economically deliver the promised service level to their clients, and to protect these clients from infection with malware. At the same time OTTs continue to exploit much more intrusive collection of personal information.

It may be claimed that other actors (e.g., law enforcement, intelligence agencies, organized crime) may gain access to ISP client application information. That is true. But why coerce an ISP to hand over relatively lean and unorganized information when it is much easier and more rewarding to get this information from the OTTs or from client device(s)?

Mechanisms have been proposed to deliver application flow requirements to ISP network elements, such as explicit path signals [RFC8558]. On-path signaling of application requirements from client applications to ISP routers and middleboxes could obviate the need for application identification. However, no such signaling protocol exists, and were it to exist the economic motivation to implement is unclear (except for malicious actors), and use of such a protocol itself could equally be considered pervasive monitoring.

3. ECH is not effective

Deploying ECH is intended to deny ISPs the ability to map a packet flow to the application it serves. It fails in this in two distinct ways.

First, it is not difficult for an ISP to block ECH and force its client to send CH in the clear. Most clients are configured to employ DNS resolvers belonging to their ISP, thus assuring authoritative behavior for CDNs and enabling load balancing. In this case the ISP can simply decide not to distribute the ECH key in HTTPS RR responses, thus inhibiting the use of ECH. Even when the DNS resolver is not under direct control of the ISP, it is not difficult to excise the ECH SvcParam from unencrypted DNS, or even to completely discard HTTPS RR reponses, choking ECH without the client ever noticing.

Second, while Deep Packet Inspection (DPI) systems exploit fields in the CH when they are available, granular classification is certainly not impossible without them. By granular classification we mean not merely identifying the OTT (e.g., Google) but differentiating between the different services it provides (e.g., search, maps, docs, Youtube, etc.). In a recent experiment on a large mobile network, we trained a classifier on one week of flow metadata with no TLS CH

information, limiting ourselves to under one second of traffic. We then tested classification a week later, a month later, and several months later. We found that after one week we can granularly classify applications and web sites with over 90% F1 accuracy, approaching the accuracy of a DPI system that uses the TLS CH fields. Performance degraded by less than two percent per month, meaning that monthly retraining suffices to maintain high accuracy.

While accuracy was not severely impacted, the computational power, and hence the electric energy required for classification, markedly increased (and this without taking training computation into account). The increased computation influences the ISP’s CAPEX and OPEX, in addition to being environmentally detrimental.

4. ECH is harmful

While about 95 percent of packet flows were correctly identified, the long tail of extremely rare applications and web sites were classified as "other". These rare services either did not appear in the training data, or occurred insufficiently frequently to be reliably recognized. Unfortunately, it is precisely these unpopular sites that are most likely to be infected, either because their security is lax, or because they are specifically created to serve malware (frequently with domain names deceptively similar to those of popular services).

These small malicious sites are often cohosted on a single server IP address with hundreds of mostly benign sites. The inability to reliably distinguish between cohosted sites forces an ISP providing network-based security to either allow its clients to be attacked, or to block hundreds of completely legitimate sites.

One might ask whether network-based security is the correct approach, rather than leaving responsibility for security to the end-user. But why should the average Internet user be burdened with installation and constant upgrading of security solutions? People expect their water utility to supply safe water without having to install filters to remove dangerous substances! Similarly, they have the right to expect their ISP to supply them with safe Internet, without needing to install and maintain firewalls, antivirus software, and a plethora of other complex tools.

5. Summary

We have shown that ECH is misguided, mostly ineffective, and harmful. It heavy-handedly attempts to block application identification even when there is, in reality, no personal information involved, and despite other entities actively exploiting this personal information. It does not inhibit application identification, but merely renders it much more expensive and environmentally detrimental. In the security-vs.-privacy trade-off it comes down clumsily on the side of perceived privacy protection while trampling on security concerns.

While many (if not most) people are willing to relinquish control over some of their privacy rights, e.g., in order to receive free service, privacy remains tremendously important. The IETF must strive to develop protocols that are not needlessly privacy invasive. Security is another, no less important, goal. The IETF must ensure that its protocols do not needlessly endanger Internet users. Were you to ask a typical Internet user whether they want the behavioral anonymity supposedly provided by ECH, but at the cost of risking loss of all their digital assets to ransomware, what do you think their answer would be?

6. IANA

This memo includes no request to IANA.

7. Security

This document should not affect the security of the Internet.

8. References

8.1. Normative References

[I-D.ietf-tls-esni]
Rescorla, E., Oku, K., Sullivan, N., and C. A. Wood, "TLS Encrypted Client Hello", Work in Progress, Internet-Draft, draft-ietf-tls-esni-25, 14 June 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-tls-esni-25>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/info/rfc7258>.

[RFC8558] Hardie, T., Ed., "Transport Protocol Path Signals", RFC 8558, DOI 10.17487/RFC8558, April 2019, <https://www.rfc-editor.org/info/rfc8558>.

Author’s Address

Yaakov (J) Stein
Allot
22 Hanagar St.
Hod Hasharon 45240
Israel
Email: ystein@allot.com
URI: www.allot.com