

# Keeping the Watchers at Bay

Mark Nottingham @mnot linux.conf.au 2020

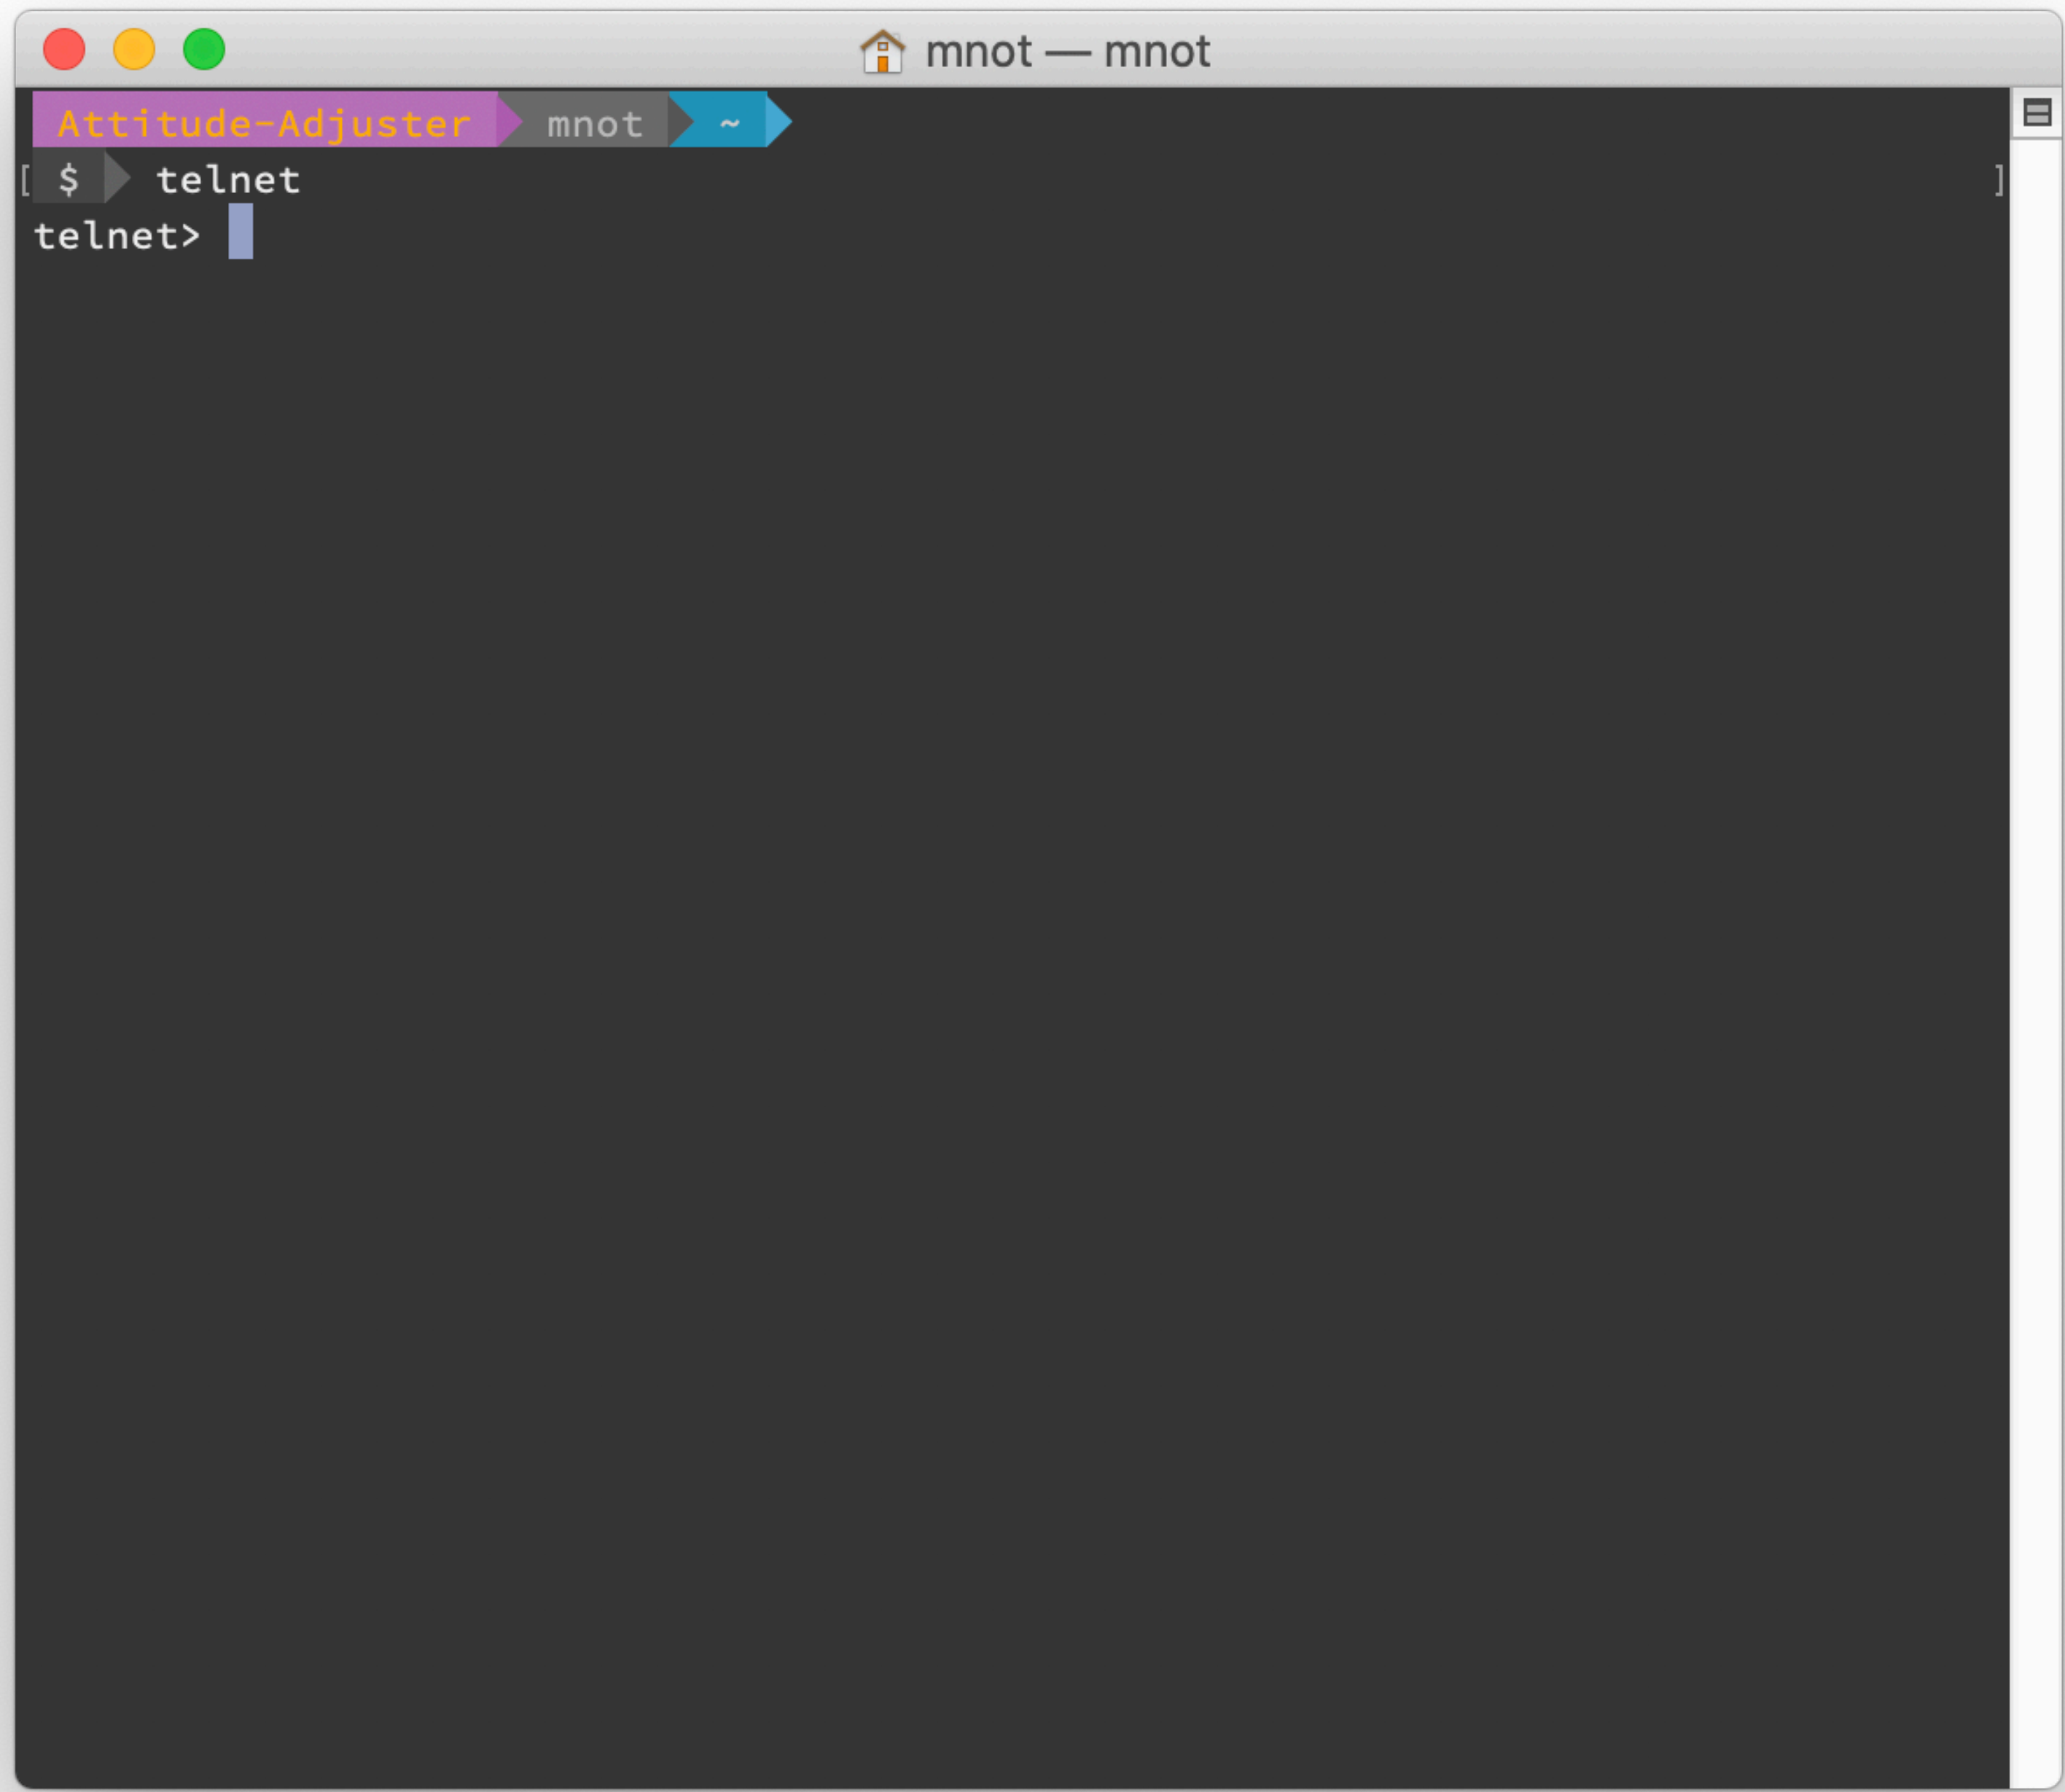
I am not a security person.



We are all security people.

1. Why secure the Internet?
2. What's happened so far?
3. What's left?
4. Some observations.
5. What can you do?

# 1. Why secure the Internet?



A terminal window titled "mnot — mnot" with a home icon. The window has three colored window control buttons (red, yellow, green) in the top-left corner. The terminal content shows a shell prompt "\$" followed by the command "telnet" being entered. The prompt changes to "telnet>" and a blue cursor is visible after the command. The terminal background is dark gray.

```
Attitude-Adjuster mnot ~  
[ $ telnet ]  
telnet> |
```



CC BY 3.0 Laura Poitras / Praxis Films

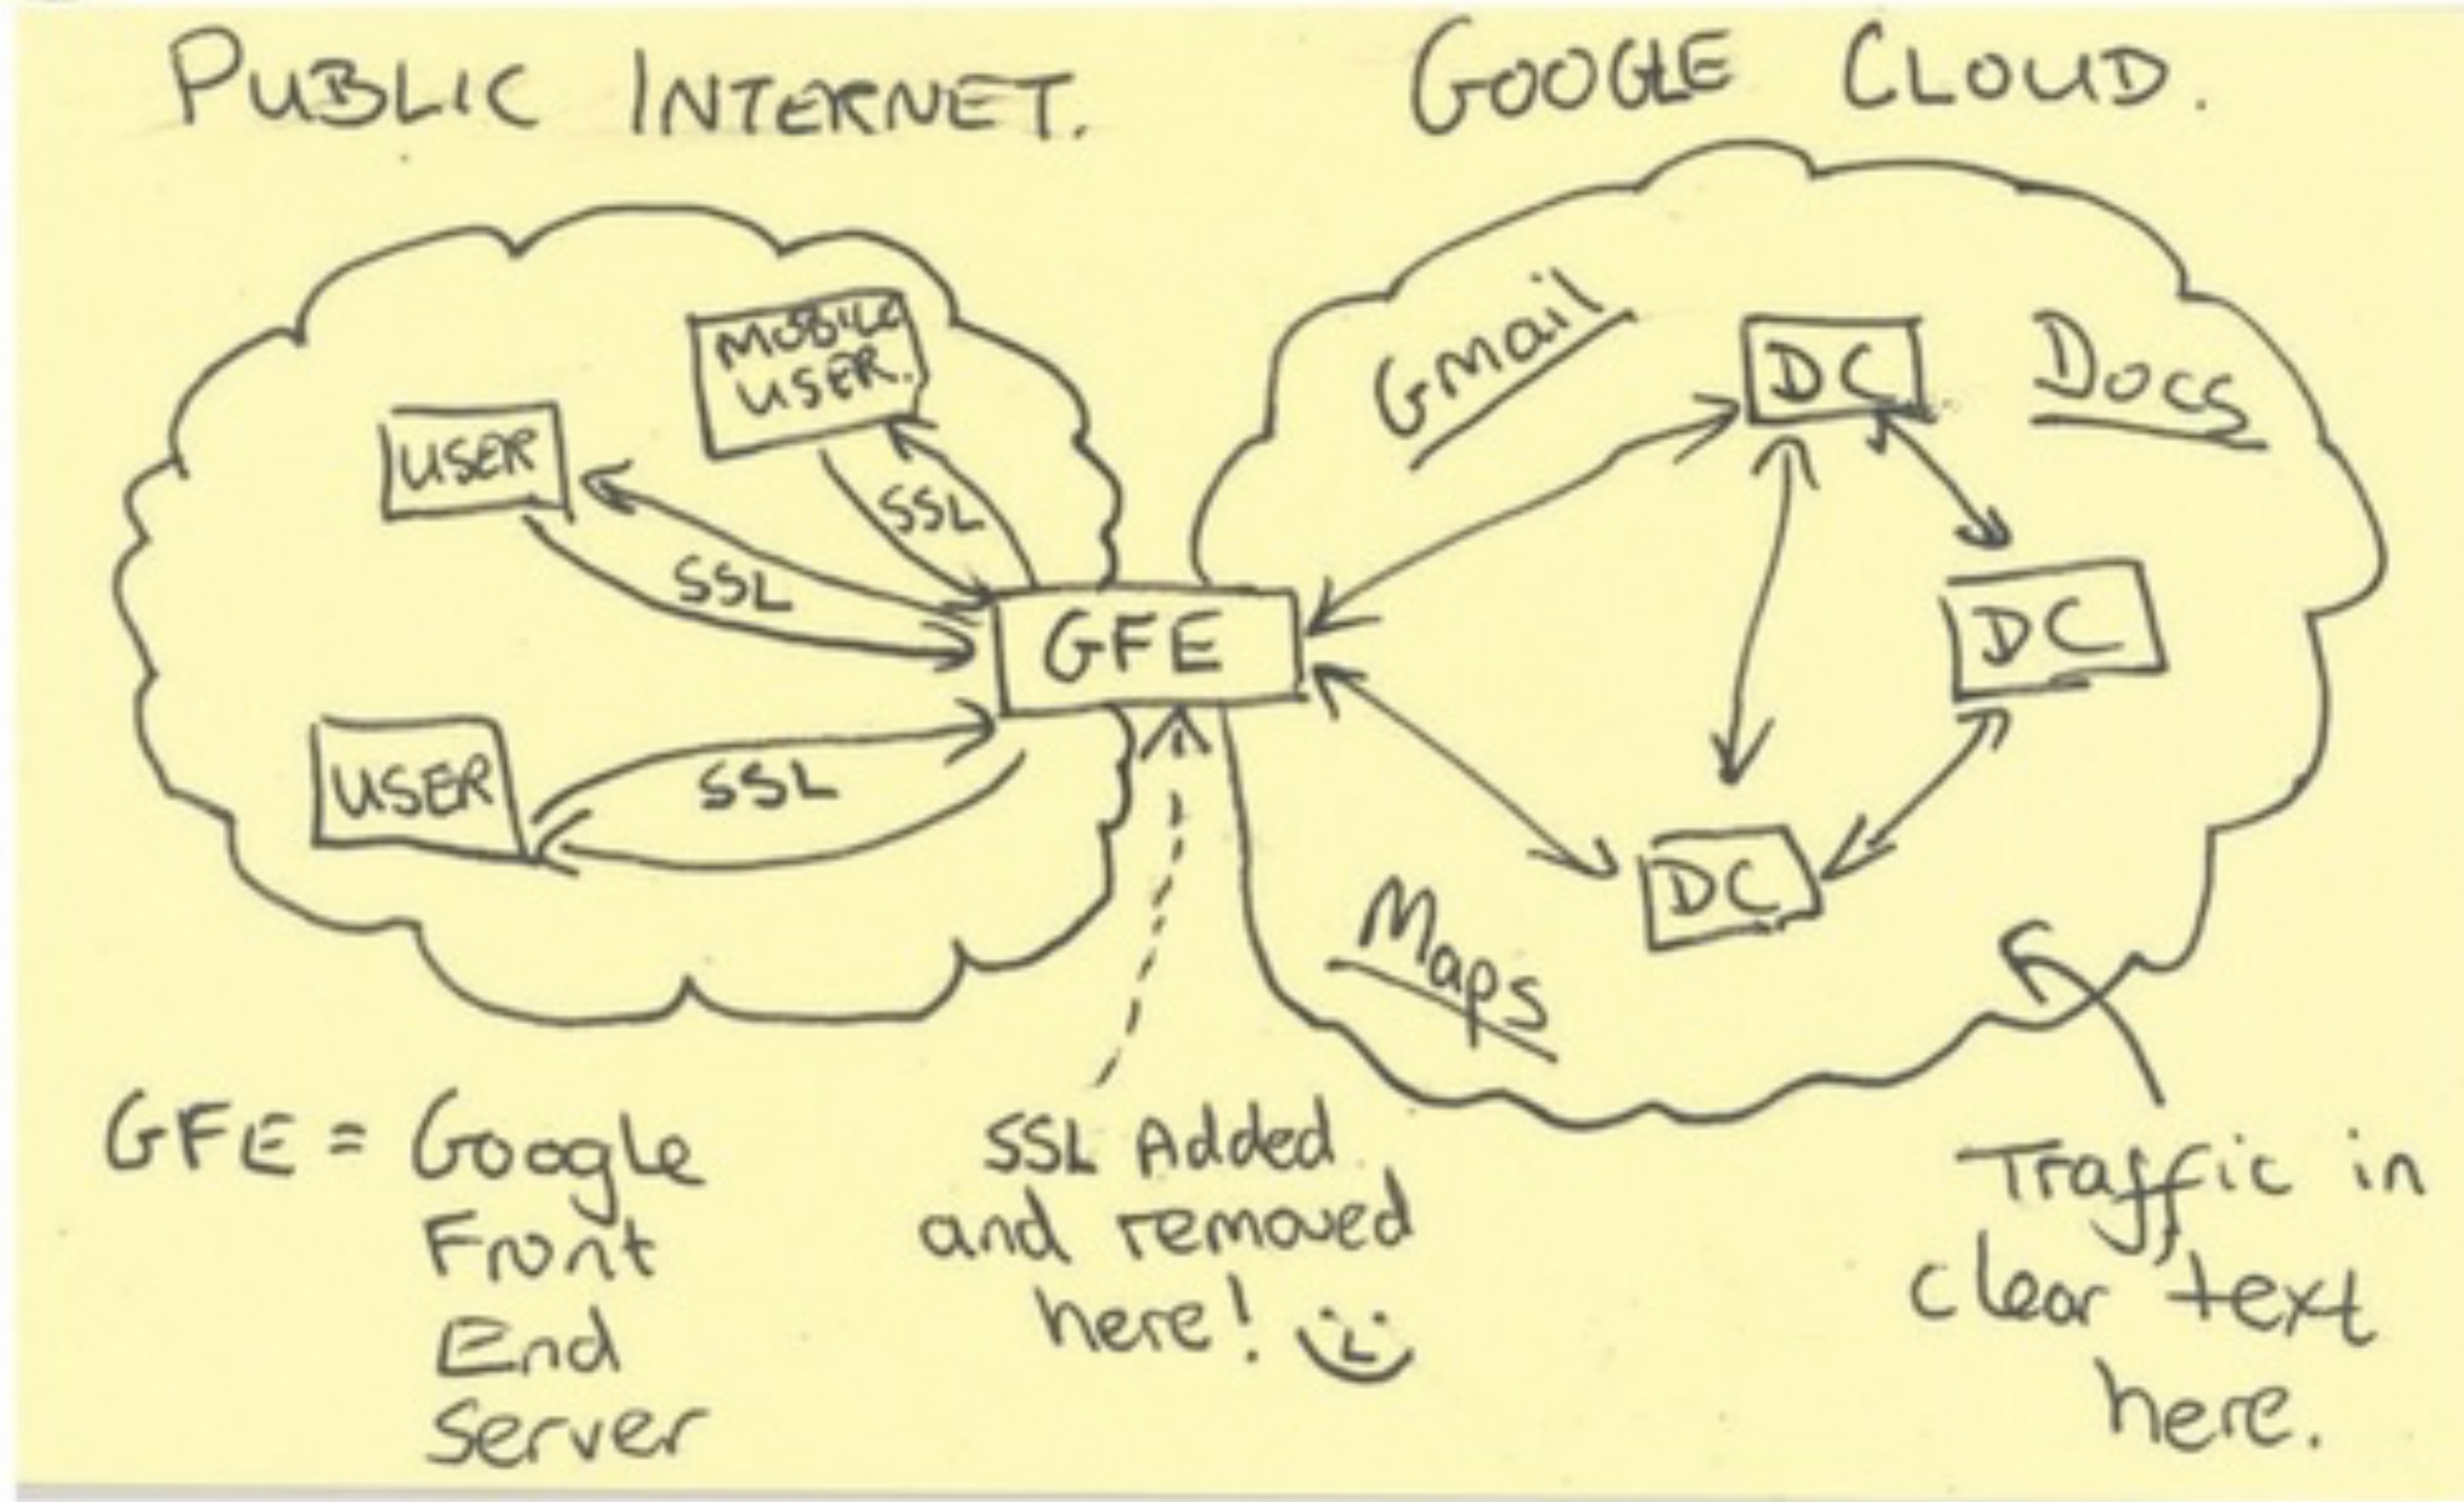
Keeping the Watchers at Bay

@mnot





# Current Efforts - Google



TOP SECRET//SI//NOFORN



# STRINT Workshop

## A W3C/IAB workshop on Strengthening the Internet Against Pervasive Monitoring (STRINT)

28 February – 1 March 2014, London



<a href="#">Home</a>	<a href="#">Agenda</a>	<a href="#">Report &amp; Papers</a>	<a href="#">How to Participate</a>	<a href="#">Logistics</a>	<a href="#">Program Committee</a>	<a href="#">... all workshops</a>
----------------------	------------------------	-------------------------------------	------------------------------------	---------------------------	-----------------------------------	-----------------------------------

The Vancouver IETF plenary concluded that pervasive monitoring represents an attack on the Internet, and the IETF has begun to carry out various of the [more obvious actions](#) required to try to handle this attack. However, there are additional much more complex questions arising that need further consideration before any additional concrete plans can be made.

The [W3C](#) and [IAB](#) will therefore host a one-day workshop on the topic of “Strengthening the Internet Against Pervasive Monitoring” before [IETF 89](#) in London in March 2014, with support from the EU

### Important Dates

**20 January 2014:**  
Deadline for Position Papers

**31 January 2014:**  
Acceptance notification and registration



[\[Docs\]](#) [\[txt|pdf\]](#) [\[draft-farrell-p...\]](#) [\[Tracker\]](#) [\[Diff1\]](#) [\[Diff2\]](#)

BEST CURRENT PRACTICE

Internet Engineering Task Force (IETF)

Request for Comments: 7258

BCP: 188

Category: Best Current Practice

ISSN: 2070-1721

S. Farrell

Trinity College Dublin

H. Tschofenig

ARM Ltd.

May 2014

## **Pervasive Monitoring Is an Attack**

### **Abstract**

Pervasive monitoring is a technical attack that should be mitigated in the design of IETF protocols, where possible.





SALE

# WIFI PINEAPPLE

\$99.99

The leading rogue access point and WiFi pentest toolkit for close access operations. Passive and active attacks analyze vulnerable and misconfigured devices.

The WiFi Pineapple® NANO and TETRA are the 6th generation pentest platforms from Hak5. Thoughtfully developed for mobile and persistent deployments, they build on over 10 years of WiFi attack expertise.

### WIFI PINEAPPLE

- ~~TETRA BASIC~~
- NANO BASIC**
- ~~TETRA TACTICAL~~
- ~~NANO TACTICAL~~

QTY

— 1 +

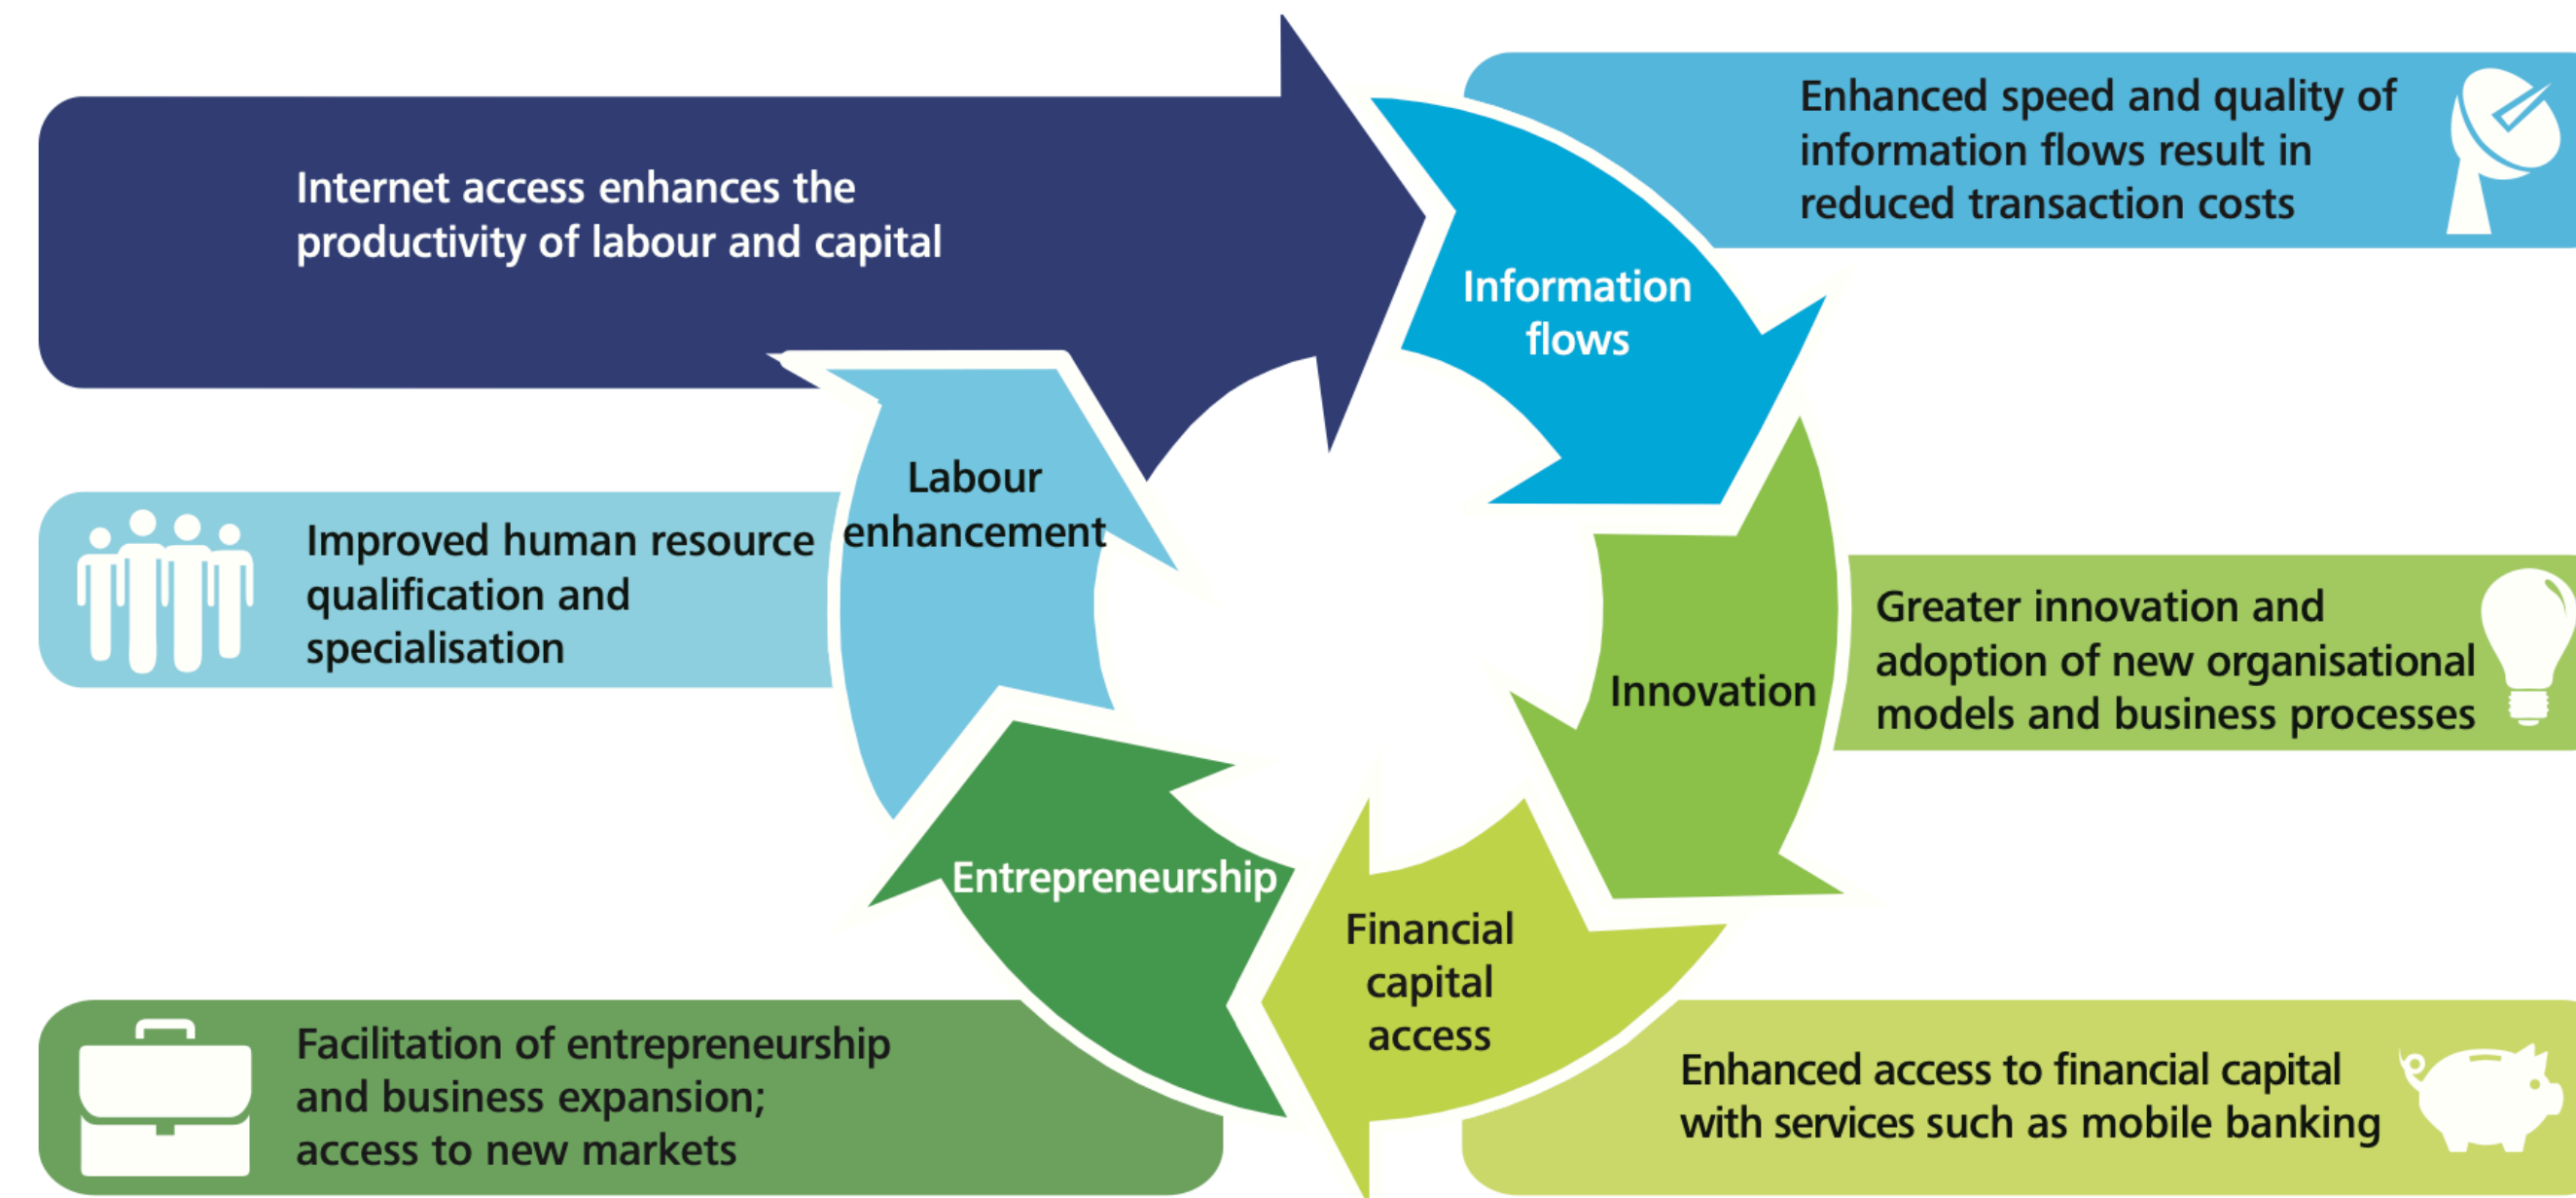
ADD TO CART



<https://shop.hak5.org/products/wifi-pineapple>

## How the internet enables economic growth

The internet offers unprecedented opportunities for economic growth in developing countries. By providing access to information, connecting people to businesses everywhere, and opening up new markets, the internet can transform the very nature of an economy and support economic development.



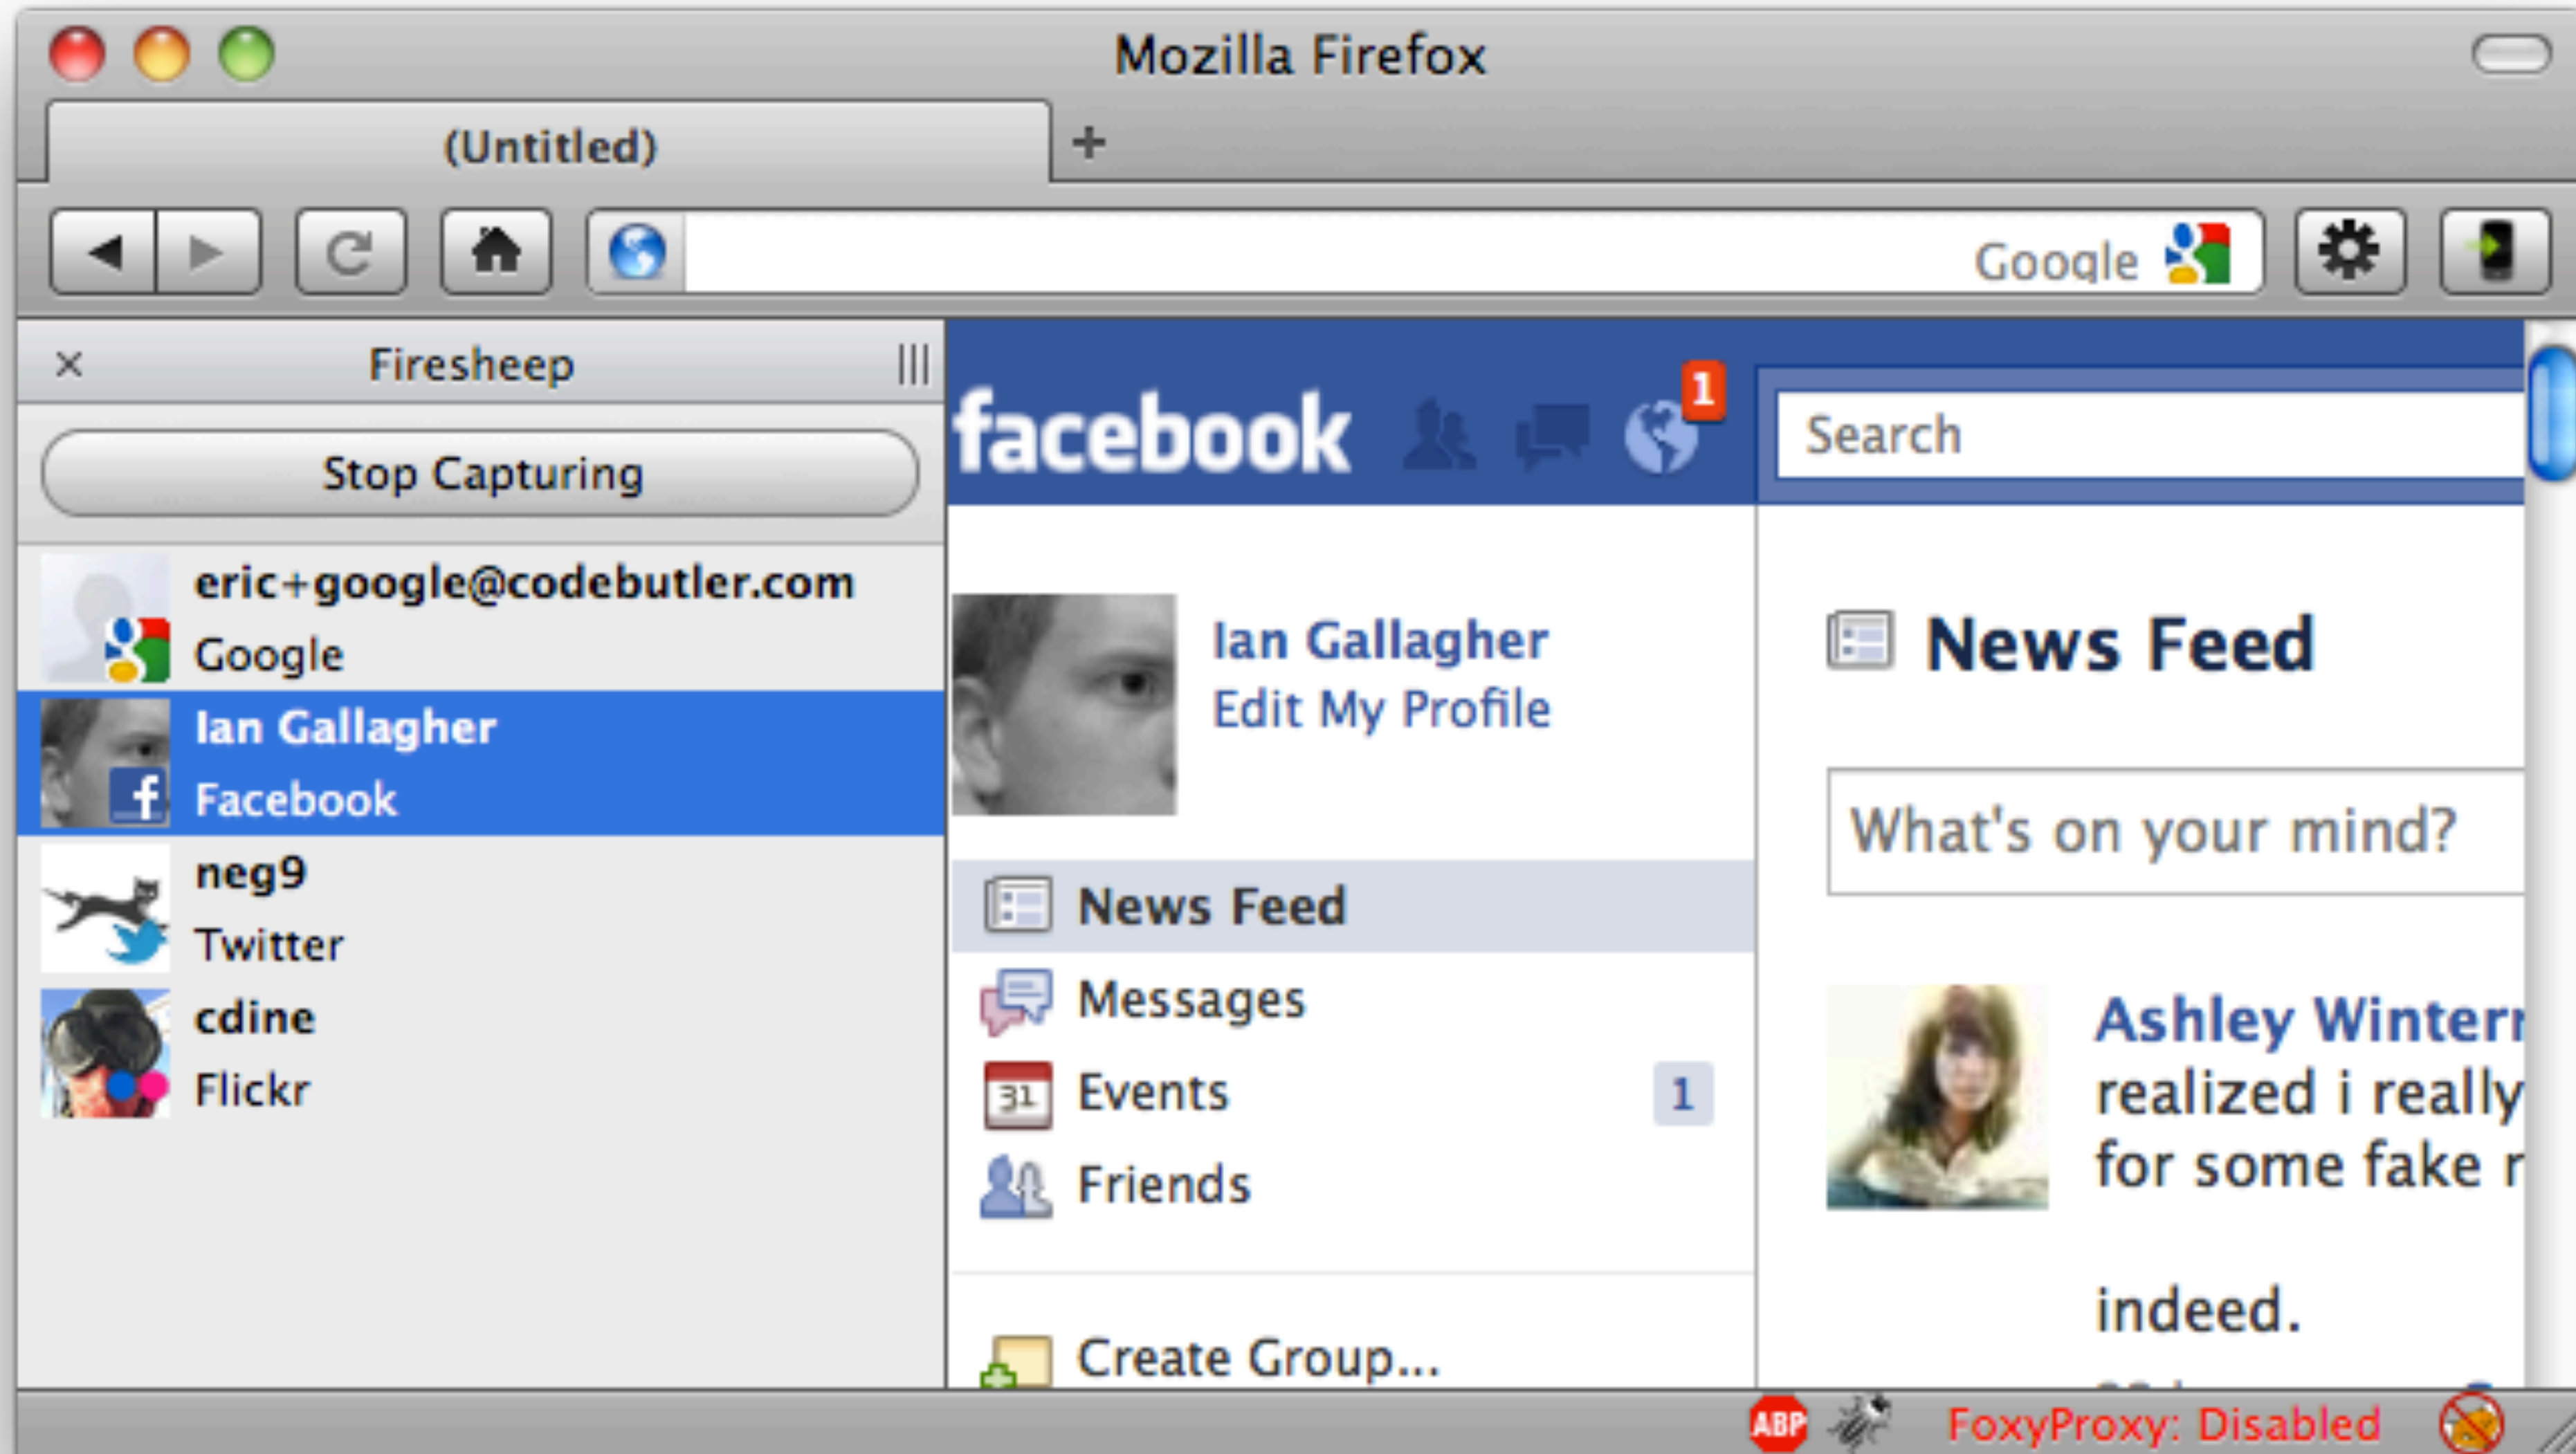
Deloitte estimates that the resulting economic activity could generate \$2.2 trillion in additional GDP, a 72% increase in the GDP growth rate, and more than 140 million new jobs.



## 2. What's happened so far

HTTP → HTTPS





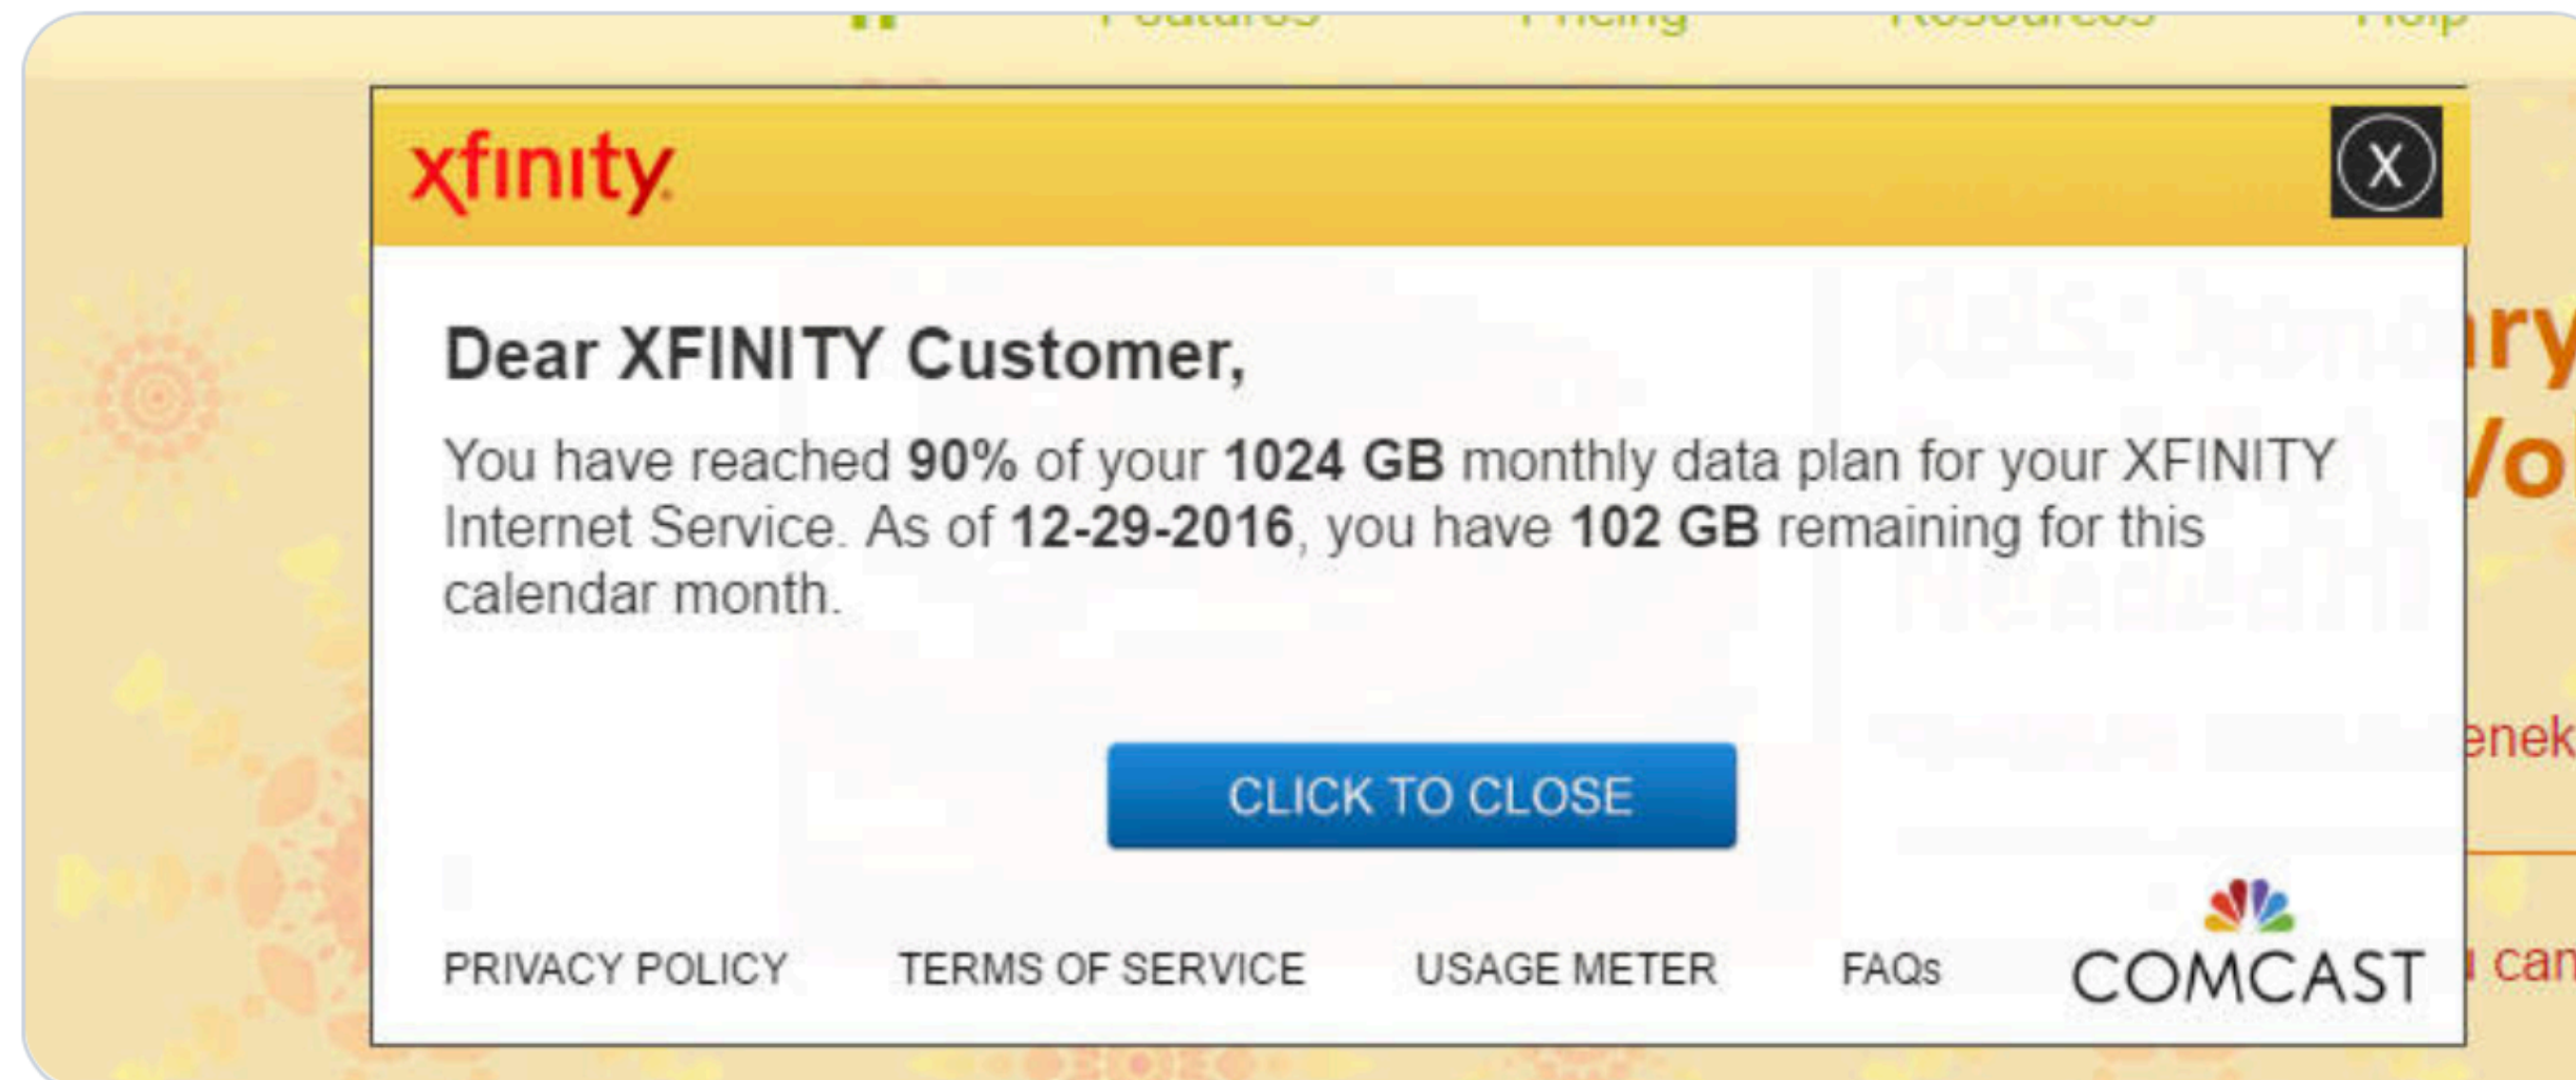
<https://codebutler.com/2010/10/24/firesheep/>



**Scott Manley**  
@DJSnM



Comcast is injecting Bandwidth cap warnings into websites. Remember, when I signed up for this I asked if there was a cap and they said no.



2:07 PM · Dec 29, 2016 · [Twitter Web Client](#)

# Verizon Injecting Perma-Cookies to Track Mobile Customers, Bypassing Privacy Controls

TECHNICAL ANALYSIS BY [JACOB HOFFMAN-ANDREWS](#) | NOVEMBER 3, 2014



Verizon users might want to start looking for another provider. In an effort to [better serve advertisers](#), Verizon Wireless has been silently modifying its users' web traffic on its network to inject a cookie-like tracker. This tracker, included in an HTTP header called X-UIDH, is sent to every unencrypted website a Verizon customer visits from a mobile device. It allows third-party advertisers and websites to assemble a deep, permanent profile of visitors' web browsing habits without their consent.



# China Uses Unencrypted Websites to Hijack Browsers in GitHub Attack

BY BILL BUDINGTON | APRIL 1, 2015



Over the past few weeks, China has been using its country's Internet infrastructure to attack political opponents by turning normal users' web browsers into Denial of Service tools. These attacks were a deep violation of the basic trust that allows the Internet to function smoothly, and an disquieting and unprecedented development in the history of state-orchestrated denial-of-service attacks. They exploited the fact that many enormous sites still use insecure HTTP rather than HTTPS, allowing the Great Firewall to modify those sites, and the fact that our web browsers are willing to run JavaScript code on an extremely liberal basis. These facts allowed China to marshal an incredible number of "zombie" systems both inside and outside of China, making billions of requests in an attempt to overwhelm the targets' servers.

<https://www.eff.org/deeplinks/2015/04/china-uses-unencrypted-websites-to-hijack-browsers-in-github-attack>

Internet Engineering Task Force (IETF)  
Request for Comments: 7540  
Category: Standards Track  
ISSN: 2070-1721

M. Belshe  
BitGo  
R. Peon  
Google, Inc  
M. Thomson, Editor  
Mozilla  
May 2015

## RFC 7540

1. Introduction
2. HTTP/2 Protocol Overview
  - 2.1. Document Organization
  - 2.2. Conventions and Terminology
3. Starting HTTP/2
  - 3.1. HTTP/2 Version Identification
  - 3.2. Starting HTTP/2 for "http" URIs
    - 3.2.1. HTTP2-Settings Header Field
  - 3.3. Starting HTTP/2 for "https" URIs
  - 3.4. Starting HTTP/2 with Prior Knowledge
  - 3.5. HTTP/2 Connection Preface
4. HTTP Frames
  - 4.1. Frame Format
  - 4.2. Frame Size
  - 4.3. Header Compression and Decompression
5. Streams and Multiplexing
  - 5.1. Stream States
    - 5.1.1. Stream Identifiers
    - 5.1.2. Stream Concurrency
  - 5.2. Flow Control
    - 5.2.1. Flow-Control Principles
    - 5.2.2. Appropriate Use of

# Hypertext Transfer Protocol Version 2 (HTTP/2)

---

## Abstract

This specification describes an optimized expression of the semantics of the Hypertext Transfer Protocol (HTTP), referred to as HTTP version 2 (HTTP/2). HTTP/2 enables a more efficient use of network resources and a reduced perception of latency by introducing header field compression and allowing multiple concurrent exchanges on the same connection. It also introduces unsolicited push of representations from servers to clients.

This specification is an alternative to, but does not obsolete, the HTTP/1.1 message syntax. HTTP's existing semantics remain unchanged.



<b>1</b>	<b>Introduction</b>
1.1	Top-level Documents
1.2	<u>Framed Documents</u>
1.3	Web Workers
1.4	Shared Workers
1.5	Service Workers
<b>2</b>	<b>Framework</b>
2.1	Intergration with WebIDL
2.2	Modifications to HTML
2.2.1	Sandboxing
2.2.2	Shared Workers
2.2.3	Feature Detection
<b>3</b>	<b>Algorithms</b>
3.1	Is the environment settings object <i>settings</i> a secure context?
3.2	Is <i>origin</i> potentially trustworthy?
3.3	Is <i>url</i> potentially trustworthy?
<b>4</b>	<b>Threat models and risks</b>
4.1	Threat Models
4.1.1	Passive Network Attacker
4.1.2	Active Network Attacker
4.2	Ancestral Risk
4.3	Risks associated with non-secure contexts
<b>5</b>	<b>Security Considerations</b>
5.1	Incomplete Isolation
5.2	localhost

# Secure Contexts

W3C Candidate Recommendation, 15 September 2016



## This version:

<https://www.w3.org/TR/2016/CR-secure-contexts-20160915/>

## Latest published version:

<https://www.w3.org/TR/secure-contexts/>

## Editor's Draft:

<https://w3c.github.io/webappsec-secure-contexts/>

## Previous Versions:

<https://www.w3.org/TR/2016/WD-secure-contexts-20160719/>

## Version History:

<https://github.com/w3c/webappsec-secure-contexts/commits/master/index.src.html>

## Feedback:

[public-webappsec@w3.org](mailto:public-webappsec@w3.org) with subject line “[secure-contexts] ... *message topic* ...” ([archives](#))

## Editor:

[Mike West](#) (Google Inc.)

## Former Editor:

Yan Zhu (Brave)

## Participate:

[File an issue](#) ([open issues](#))

Copyright © 2016 W3C® (MIT, ERCIM, Keio, Beihang). W3C [liability](#), [trademark](#) and [document use](#) rules apply.

## Abstract

This specification defines "secure contexts", thereby allowing user agent implementers and specification authors to enable certain features only when certain minimum standards of authentication and confidentiality are met.

# Problem: *Mixed Content*



# Upgrade Insecure Requests

W3C Candidate Recommendation, 8 October 2015

**This version:**

<http://www.w3.org/TR/2015/CR-upgrade-insecure-requests-20151008/>

**Latest version:**

<http://www.w3.org/TR/upgrade-insecure-requests/>

**Editor's Draft:**

<https://w3c.github.io/webappsec-upgrade-insecure-requests/>

**Previous Versions:**

<http://www.w3.org/TR/2015/WD-upgrade-insecure-requests-20150424/>

**Version History:**

<https://github.com/w3c/webappsec-upgrade-insecure-requests/commits/master/index.src.html>

**Feedback:**

[public-webappsec@w3.org](mailto:public-webappsec@w3.org) with subject line “[upgrade-insecure-requests] ... message topic ...” ([archives](#))

**Editor:**

[Mike West](#) (Google Inc.)

**Participate:**

[File an issue](#) ([open issues](#))

Copyright © 2015 W3C® ([MIT](#), [ERCIM](#), [Keio](#), [Beihang](#)). W3C [liability](#), [trademark](#) and [document use](#) rules apply.

---

## Abstract

This document defines a mechanism which allows authors to instruct a user agent to upgrade *a priori* insecure resource requests to secure transport before fetching them.

“HTTPS is Slow”

<https://istlsfastyet.com>



BIZ &amp; IT —

# It wasn't easy, but Netflix will soon use HTTPS to secure video streams

Netflix move leaves Amazon as the most visible no-show to the Web crypto party.

DAN GOODIN - 4/17/2015, 1:47 AM



Netflix will soon use the HTTPS protocol to authenticate and encrypt customer streams, a move that helps ensure what users watch stays secret. The move now leaves Amazon as one of the most noticeable no-shows to the Web encryption party.



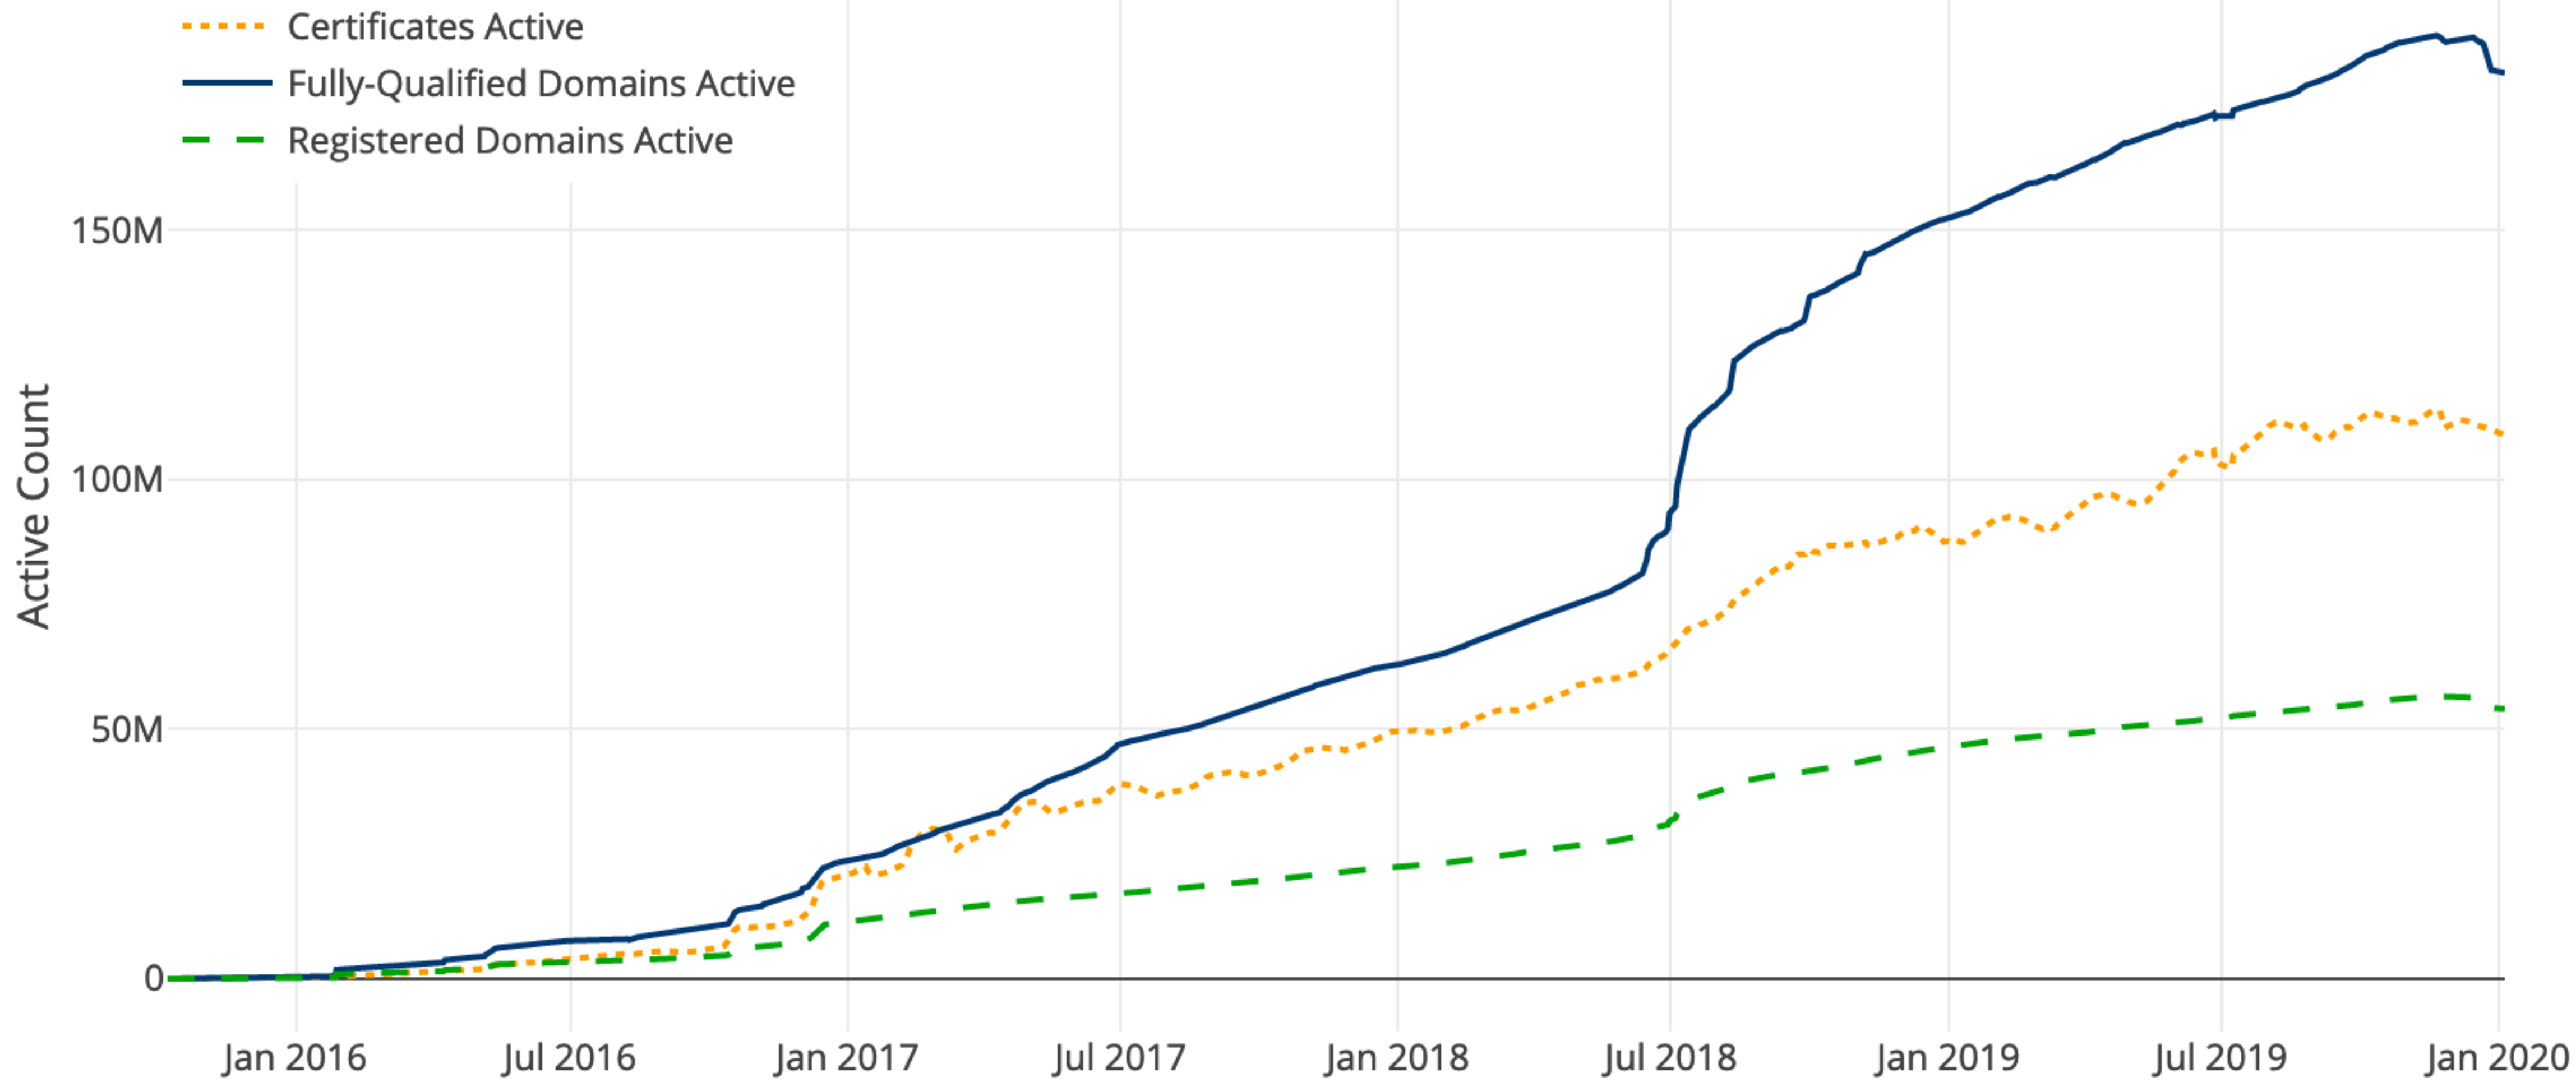
Flipping on the HTTPS switch on Netflix's vast network of OpenConnect Appliances (OCAs) has been anything but effortless. That's because the demands of mass movie streaming can impose severe penalties when [transport layer security \(TLS\)](#) is enabled. Each Netflix OCA is a server-class computer with a 64-bit Xeon CPU running the FreeBSD operating system. Each box stores up to 120 terabytes of data and serves up to 40,000 simultaneous, long-lived connections, a load that requires as much as 40 gigabits per second of continuous bandwidth. Like Amazon, Netflix has long encrypted log-in pages and other sensitive parts of its website but has served movie streams over unsecured HTTP connections. Netflix took the unusual step of announcing the switch in a [quarterly earnings letter](#) that company officials sent shareholders Tuesday.



**“Certificates are Expensive”**




# Let's Encrypt



<https://letsencrypt.org/stats/#growth>



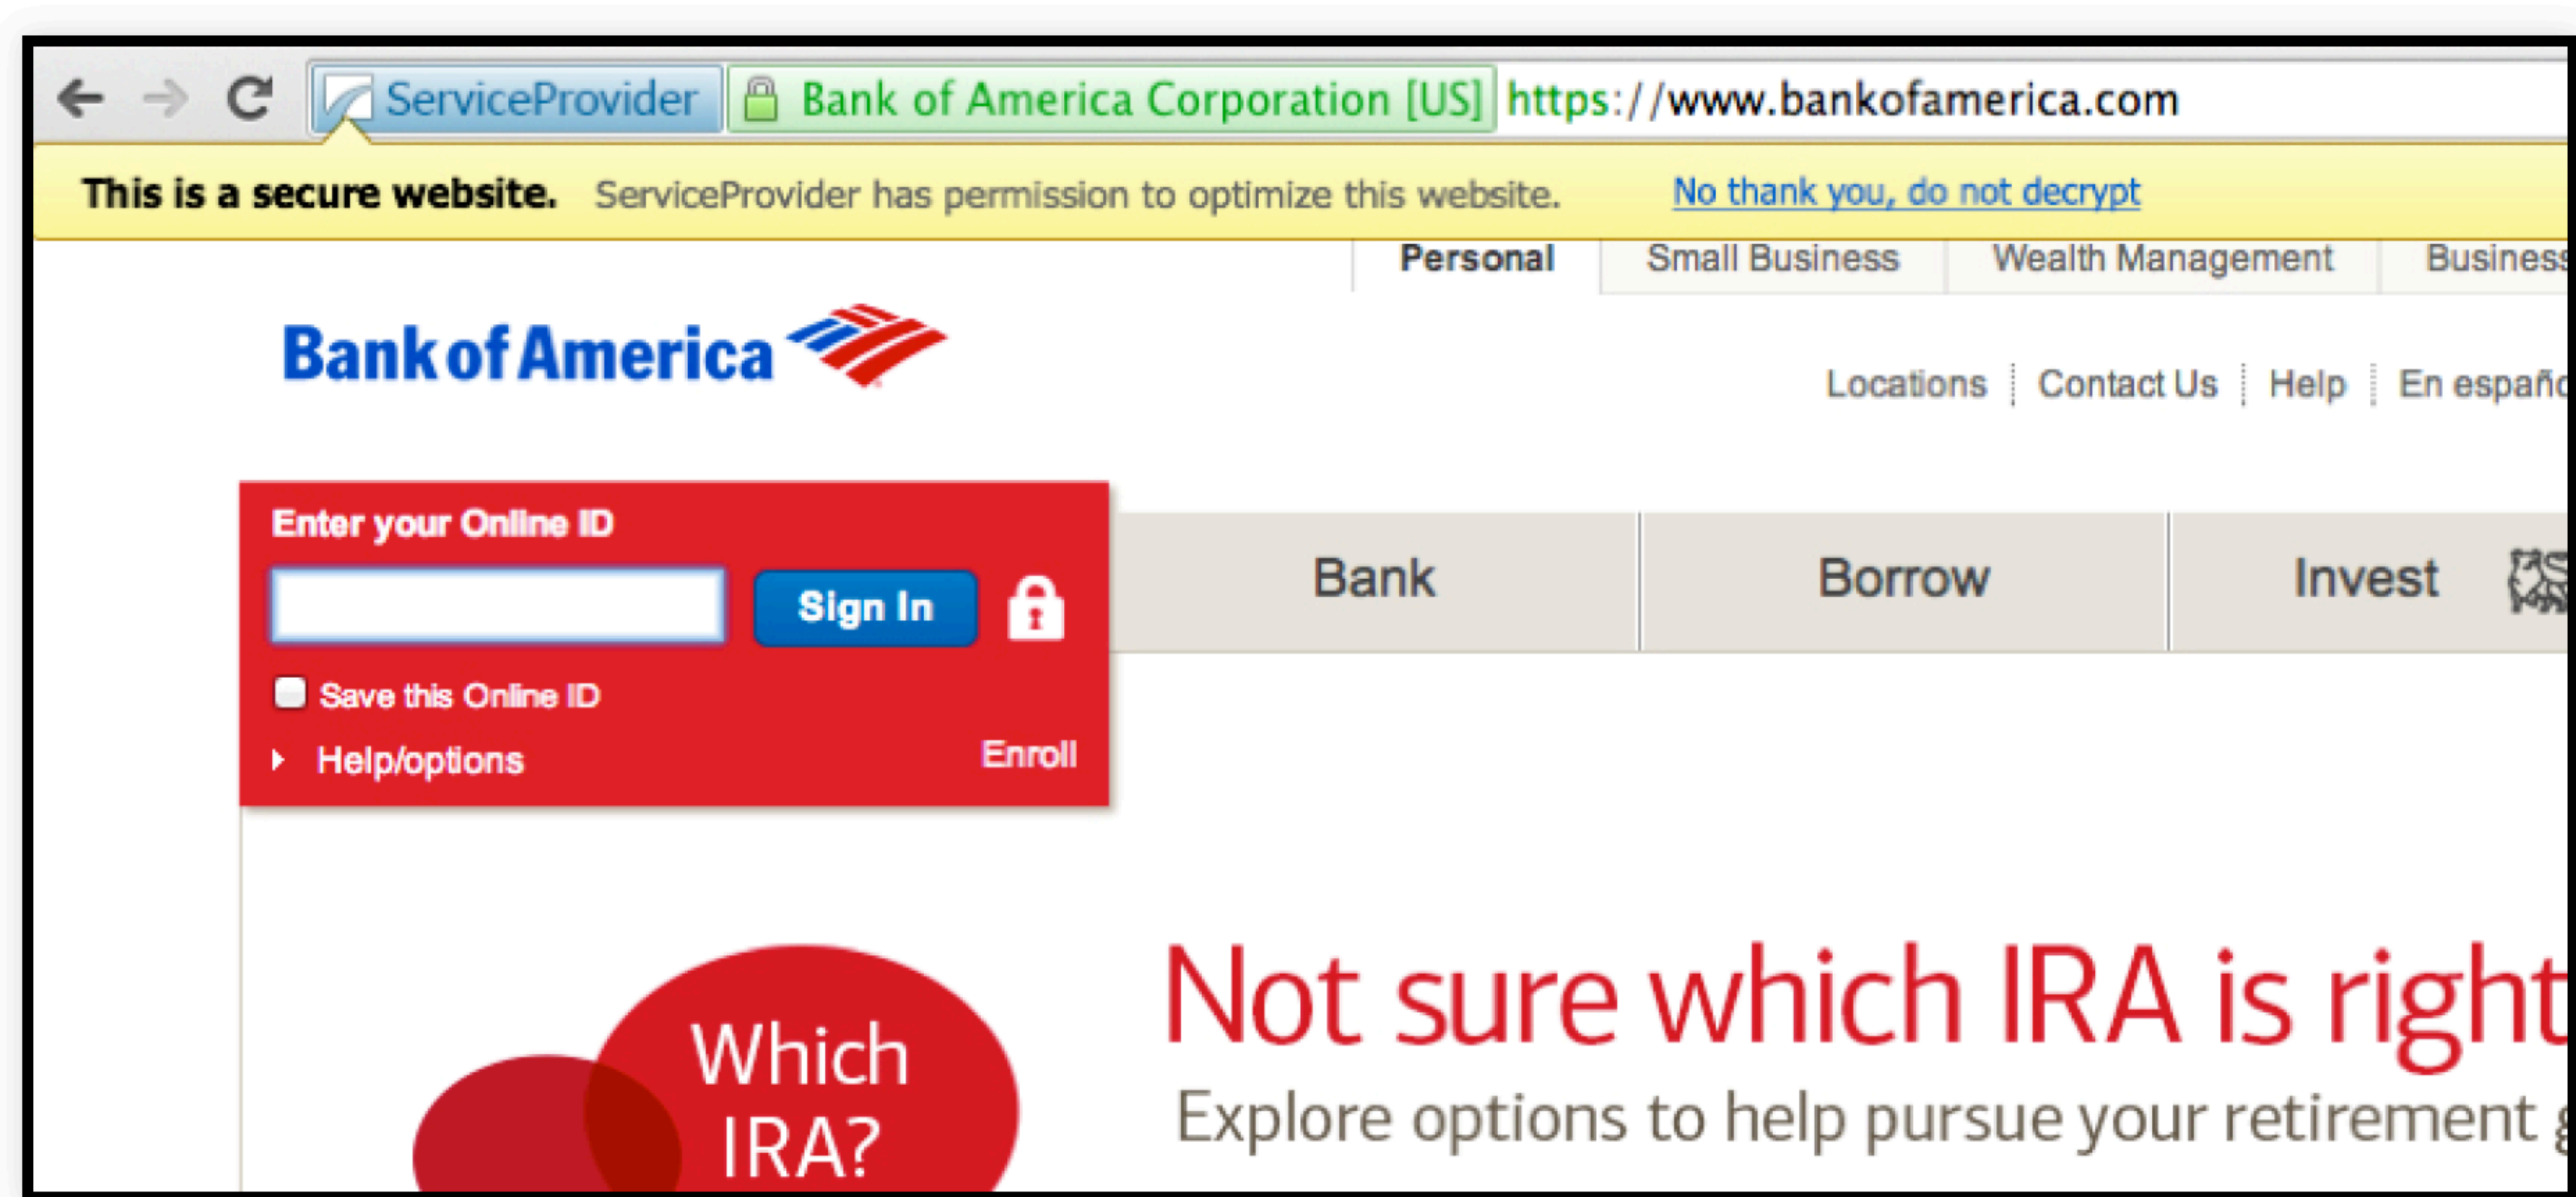
# Percentage of sites redirecting to HTTPS



Some were unhappy.



# User notification and opt-out



Address bar shows logo of trusted proxy entity to the left of the lock.





**TLS 1.2 → TLS 1.3**



<http://heartbleed.com/>

## FREAK: Factoring RSA Export Keys



Fig. 2: FREAK exploit on Safari

export key. By design, export RSA moduli must be less than 512 bits long; hence, they can be factored in less than 12 hours for \$100 on Amazon EC2.

Ironically, many US government agencies (including the NSA and FBI), as well as a number of popular websites (IBM, or Symantec) enable export ciphersuites on their server - by factoring their 512-bit RSA modulus, an attacker can impersonate them to vulnerable clients.

Among the various state machine problems we found, one is particularly interesting because it leads to a server impersonation exploits against several mainstream browsers (including Safari and OpenSSL-based browsers on Android).

This attack targets a class of deliberately weak **export cipher suites**. As the name implies, this class of algorithms were **introduced under the pressure of US governments agencies to ensure that they would be able to decrypt all foreign encrypted communication**, while stronger algorithms were banned from export (as they were classified as weapons of war).

Support for these weak algorithms has remained in many implementations such as OpenSSL, even though they are typically disabled by default; however, we discovered that several implementations incorrectly allow the message sequence of export ciphersuites to be used even if a non-export ciphersuite was negotiated.

Thus, if a server is willing to negotiate an export ciphersuite, a man-in-the-middle may trick a browser (which normally doesn't allow it) to use a weak



# Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice

David Adrian<sup>¶</sup> Karthikeyan Bhargavan<sup>\*</sup> Zakir Durumeric<sup>¶</sup> Pierrick Gaudry<sup>†</sup> Matthew Green<sup>§</sup>  
J. Alex Halderman<sup>¶</sup> Nadia Heninger<sup>‡</sup> Drew Springall<sup>¶</sup> Emmanuel Thomé<sup>†</sup> Luke Valenta<sup>‡</sup>  
Benjamin VanderSloot<sup>¶</sup> Eric Wustrow<sup>¶</sup> Santiago Zanella-Béguelin<sup>||</sup> Paul Zimmermann<sup>†</sup>

<sup>\*</sup> INRIA Paris-Rocquencourt    <sup>†</sup> INRIA Nancy-Grand Est, CNRS, and Université de Lorraine  
<sup>||</sup> Microsoft Research    <sup>‡</sup> University of Pennsylvania    <sup>§</sup> Johns Hopkins    <sup>¶</sup> University of Michigan

For additional materials and contact information, visit [WeakDH.org](http://WeakDH.org).

## ABSTRACT

We investigate the security of Diffie-Hellman key exchange as used in popular Internet protocols and find it to be less secure than widely believed. First, we present Logjam, a novel flaw in TLS that lets a man-in-the-middle downgrade connections to “export-grade” Diffie-Hellman. To carry out this attack, we implement the number field sieve discrete log algorithm. After a week-long precomputation for a specified 512-bit group, we can compute arbitrary discrete logs in that group in about a minute. We find that 82% of vulnerable servers use a single 512-bit group, allowing us to compromise connections to 7% of Alexa Top Million HTTPS sites. In response, major browsers are being changed to reject short groups.

We go on to consider Diffie-Hellman with 768- and 1024-bit groups. We estimate that even in the 1024-bit case, the computations are plausible given nation-state resources. A small number of fixed or standardized groups are used by millions of servers; performing precomputation for a single 1024-bit group would allow passive eavesdropping on 18% of popular HTTPS sites, and a second group would allow decryption

coded, or widely shared Diffie-Hellman parameters has the effect of dramatically reducing the cost of large-scale attacks, bringing some within range of feasibility today.

The current best technique for attacking Diffie-Hellman relies on compromising one of the private exponents ( $a$ ,  $b$ ) by computing the discrete log of the corresponding public value ( $g^a \bmod p$ ,  $g^b \bmod p$ ). With state-of-the-art number field sieve algorithms, computing a single discrete log is more difficult than factoring an RSA modulus of the same size. However, an adversary who performs a large precomputation for a prime  $p$  can then quickly calculate arbitrary discrete logs in that group, amortizing the cost over all targets that share this parameter. Although this fact is well known among mathematical cryptographers, it seems to have been lost among practitioners deploying cryptosystems. We exploit it to obtain the following results:

*Active attacks on export ciphers in TLS.* We introduce Logjam, a new attack on TLS by which a man-in-the-middle attacker can downgrade a connection to export-grade cryptography. This attack is reminiscent of the FREAK attack [7]

# SSL/TLS Blues

- Complex old protocol
- Monoculture
- Outdated crypto
- Unsure about how much worse it would get





# TLS 1.3

- Simplify where possible
- Encrypt most of the handshake
- Remove outdated cryptography
- Wide review (including formal verification)

# TRON Workshop Call For Papers

## TLSv1.3 – Ready or Not? (TRON)

### Background

The Transport Layer Security (TLS) protocol (RFC5246) formerly known as the Secure Sockets Layer (SSL) has evolved to become one of, or perhaps the, most important security protocol used on the Internet, providing the security layer that underpins the web but also many other Internet protocols. Over the two decades since its inception, many TLS implementation vulnerabilities and some protocol design flaws have been discovered, sometimes requiring vary large scale and urgent remediation which is costly, damages confidence and exposes hosts on the Internet to sometimes significant risk. At the same time, there are ongoing trends towards much greater use of encryption, and in particular for use of TLS, as the importance of security and privacy for Internet users becomes more apparent and as the Internet is more and more attacked from perse sources. As a result, ensuring that new versions of TLS are thoroughly analysed is more important than used be the case.

TLS is developed by the Internet Engineering Task Force (<https://www.ietf.org>) TLS working group (<https://www.ietf.org/wg/tls/>) who are currently developing a major revision of TLS, TLSv1.3, which aims to improve both the security properties of TLS and the efficiency with which the TLS protocol can be used in important use-cases. The expected timeframe for this workanticipates that the TLSv1.3 protocol should be finalised around November 2015, ideally with only editorial changes being needed subsequently. The latest version of the TLSv1.3 draft can be accessed at: <https://tools.ietf.org/html/draft-ietf-tls-tls13>

Normally this timeline would result in the final specifcaton (a new RFC) being issued very early in 2016. Given the importance of TLS, the likelihood that TLSv1.3 will be widely deployed very soon after it is complete, and the history of issues with earlier TLS versions and implementations, it is important that the security research communityhas the opportunity to analyse the new protocol before final

<https://bit.ly/tls-tron>



# Automated Analysis and Verification of TLS 1.3: 0-RTT, Resumption and Delayed Authentication

(Preprint, February 10, 2016)

Cas Cremers, Marko Horvat  
*Department of Computer Science*  
*University of Oxford, UK*

Sam Scott, Thyla van der Merwe  
*Information Security Group*  
*Royal Holloway, University of London, UK*

**Abstract**—After a development process of many months, the TLS 1.3 specification is nearly complete. To prevent past mistakes, this crucial security protocol must be thoroughly scrutinised prior to deployment.

In this work we model and analyse the latest draft of the TLS 1.3 specification, namely revision 10, using the Tamarin prover, a tool for the automated analysis of security protocols. We specify and analyse the interaction of various handshake modes for an unbounded number of concurrent TLS connections. **We show that revision 10 meets the goals of authenticated key exchange in both the unilateral and mutual authentication cases.**

[30], [33], [34], [36], [39], many representing great advances on both the manual and automated fronts and resulting in the discovery of many weaknesses.

The various flaws identified in TLS 1.2 [17] and below, be they implementation- or specification-based, have prompted the TLS Working Group to adopt an ‘analysis-before-deployment’ design paradigm in drafting the next version of the protocol, TLS 1.3 [47]. Most notably, the cryptographic core of the new TLS handshake protocol is largely influenced by the OPTLS protocol of Krawczyk and Wee [34], a protocol that has been expressly designed to offer zero Round-Trip Time (0-RTT) exchanges and ensure perfect

<b>Browser</b>	<b>Percentage of TLS 1.3</b>
Chrome	30%
Firefox	27%
Safari	27%

Table 1: Percentage of TLS 1.3 connections amongst web browsers.

# TLS 1.3 Implementations

name	language	role(s)	version	features/limitations
<a href="#">fizz</a>	C++	C/S	RFC 8446	Based on libsodium, includes secure design abstractions. Zero-copy for advanced performance.
<a href="#">NSS</a>	C	C/S	RFC 8446	Almost everything, except some crypto primitives
<a href="#">Mint</a>	Go	C/S	-18	PSK resumption, 0-RTT, HRR
<a href="#">nqsb</a>	OCaml	C/S	-11	PSK/DHE-PSK, no EC*, no client auth, no 0RTT -- live server at <a href="https://tls13test.nqsb.io">tls13test.nqsb.io</a> port 4433, records traces, ping <a href="#">@hannesm</a> , contains a static PSK/DHE_PSK token: id: 0x0000 ▶ secret:
<a href="#">ProtoTLS</a>	JavaScript	C/S	-13	EC/DHE/PSK, no HelloRetryRequest
<a href="#">miTLS</a>	F*	C/S	RFC 8446	EC/DHE/PSK/0-RTT, no RSA-PSS, no post-HS-auth, no ESNI
<a href="#">Tris</a>	Go	C/S	RFC 8446	ECDHE/PSK/0-RTT, no HelloRetryRequest
<a href="#">BoringSSL</a>	C	C/S	-23, -28, RFC 8446	P-256, X25519, HelloRetryRequest, resumption, 0-RTT, KeyUpdate
<a href="#">Wireshark</a>	C	other	-18 to -28, RFC 8446	Full decryption and dissection support for drafts 19-21 since 2.4.0 ( <a href="#">keylog format</a> ). Supports 18-21 since 2.4.2, -22 since 2.4.3, -23 since 2.4.5, -24 to -28 (+0RTT trial decryption) since 2.6.0. <a href="#">Tracking bug</a> .
<a href="#">picotls</a>	C	C/S	-18,-21,-23,-26	P-256, X25519, HelloRetryRequest, resumption, 0-RTT
<a href="#">rustls</a>	Rust	C/S	-28 (final on branch)	P-256/P-384/curve25519, HRR, resumption, 0-RTT client
<a href="#">Haskell tls</a>	Haskell	C/S	-28	ECDHE w/ P* and X*, full, HRR, PSK, 0RTT
<a href="#">Leto</a>	C#	S	-18	DHE, X25519, AES, no PSK no 0RTT. Tested against NSS
<a href="#">OpenSSL</a>	C	C/S	RFC 8446	P-256, P-384, P-521, FFDHE, X25519, X448, Ed25519, Ed448, HelloRetryRequest, resumption, PSK, 0-RTT, CCS, cookies, stateless server, Post-



Some were unhappy.

## [TLS] Industry Concerns about TLS 1.3

BITS Security <BITSSecurity@fsroundtable.org> | Thu, 22 September 2016 17:24 UTC | [Show header](#)

To: IETF TLS 1.3 Working Group Members

My name is Andrew Kennedy and I work at BITS, the technology policy division of the Financial Services Roundtable (<http://www.fsroundtable.org/bits>). My organization represents approximately 100 of the top 150 US-based financial services companies including banks, insurance, consumer finance, and asset management firms.

I manage the Technology Cybersecurity Program, a CISO-driven forum to investigate emerging technologies; integrate capabilities into member operations; and advocate member, sector, cross-sector, and private-public collaboration.

While I am aware and on the whole supportive of the significant contributions to internet security this important working group has made in the last few years I recently learned of a proposed change that would affect many of my organization's member institutions: the deprecation of RSA key exchange.

Deprecation of the RSA key exchange in TLS 1.3 will cause significant problems for financial institutions, almost all of whom are running TLS internally and have significant, security-critical investments in out-of-band TLS decryption.

Like many enterprises, financial institutions depend upon the ability to decrypt TLS traffic to implement data loss protection, intrusion detection and prevention, malware detection, packet capture and analysis, and DDoS mitigation. Unlike some other businesses, financial institutions also rely upon TLS traffic decryption to implement fraud monitoring and surveillance of supervised employees. The products which support these capabilities will need to be replaced or substantially redesigned at significant cost and loss of scalability to continue to support the functionality financial institutions and their regulators require.

<https://mailarchive.ietf.org/arch/msg/tls/KQlyNhPk8K6jOoe2ScdPZ8E08RE>

# ETSI TS 103 523-3 V1.1.1 (2018-10)



**CYBER;  
Middlebox Security Protocol;  
Part 3: Profile for enterprise network and  
data centre access control**

[https://www.etsi.org/deliver/etsi\\_ts/103500\\_103599/10352303/01.01.01\\_60/ts\\_10352303v010101p.pdf#page=15](https://www.etsi.org/deliver/etsi_ts/103500_103599/10352303/01.01.01_60/ts_10352303v010101p.pdf#page=15)

# 3. What's left

TCP → QUIC



### TCP Header

Offsets	Octet	0								1								2								3							
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Source port																Destination port															
4	32	Sequence number																															
8	64	Acknowledgment number (if ACK set)																															
12	96	Data offset				Reserved 0 0 0			N S	C W R	E C R	U R G	A C K	P S H	R S T	S S Y	F I N N	Window Size															
16	128	Checksum																Urgent pointer (if URG set)															
20	160	Options (if <i>data offset</i> > 5. Padded at the end with "0" bytes if necessary.)																															
...	...	...																															

# Identifying HTTPS-Protected Netflix Videos in Real-Time

Andrew Reed, Michael Kranch

Dept. of Electrical Engineering and Computer Science

United States Military Academy at West Point

West Point, New York, USA

{andrew.reed, michael.kranch}@usma.edu

## ABSTRACT

After more than a year of research and development, Netflix recently upgraded their infrastructure to provide HTTPS encryption of video streams in order to protect the privacy of their viewers. Despite this upgrade, we demonstrate that it is possible to accurately identify Netflix videos from passive traffic capture in real-time with very limited hardware requirements. Specifically, **we developed a system that can report the Netflix video being delivered by a TCP connection using only the information provided by TCP/IP headers.**

To support our analysis, we created a fingerprint database comprised of 42,027 Netflix videos. Given this collection of fingerprints, we show that our system can differentiate between videos with greater than 99.99% accuracy. Moreover, when tested against 200 random 20-minute video streams, our system identified 99.5% of the videos with the majority of the identifications occurring less than two and a half minutes into the video stream.

protected Netflix videos. We then improve upon the previous work by fully automating the fingerprint creation process, thereby enabling us to create an extensive collection of Netflix fingerprints which we then use to conduct a robust assessment of the attack. Finally, we developed a network appliance that can, in real-time, identify HTTPS-protected Netflix videos using IP and TCP headers obtained from passive capture of network traffic.

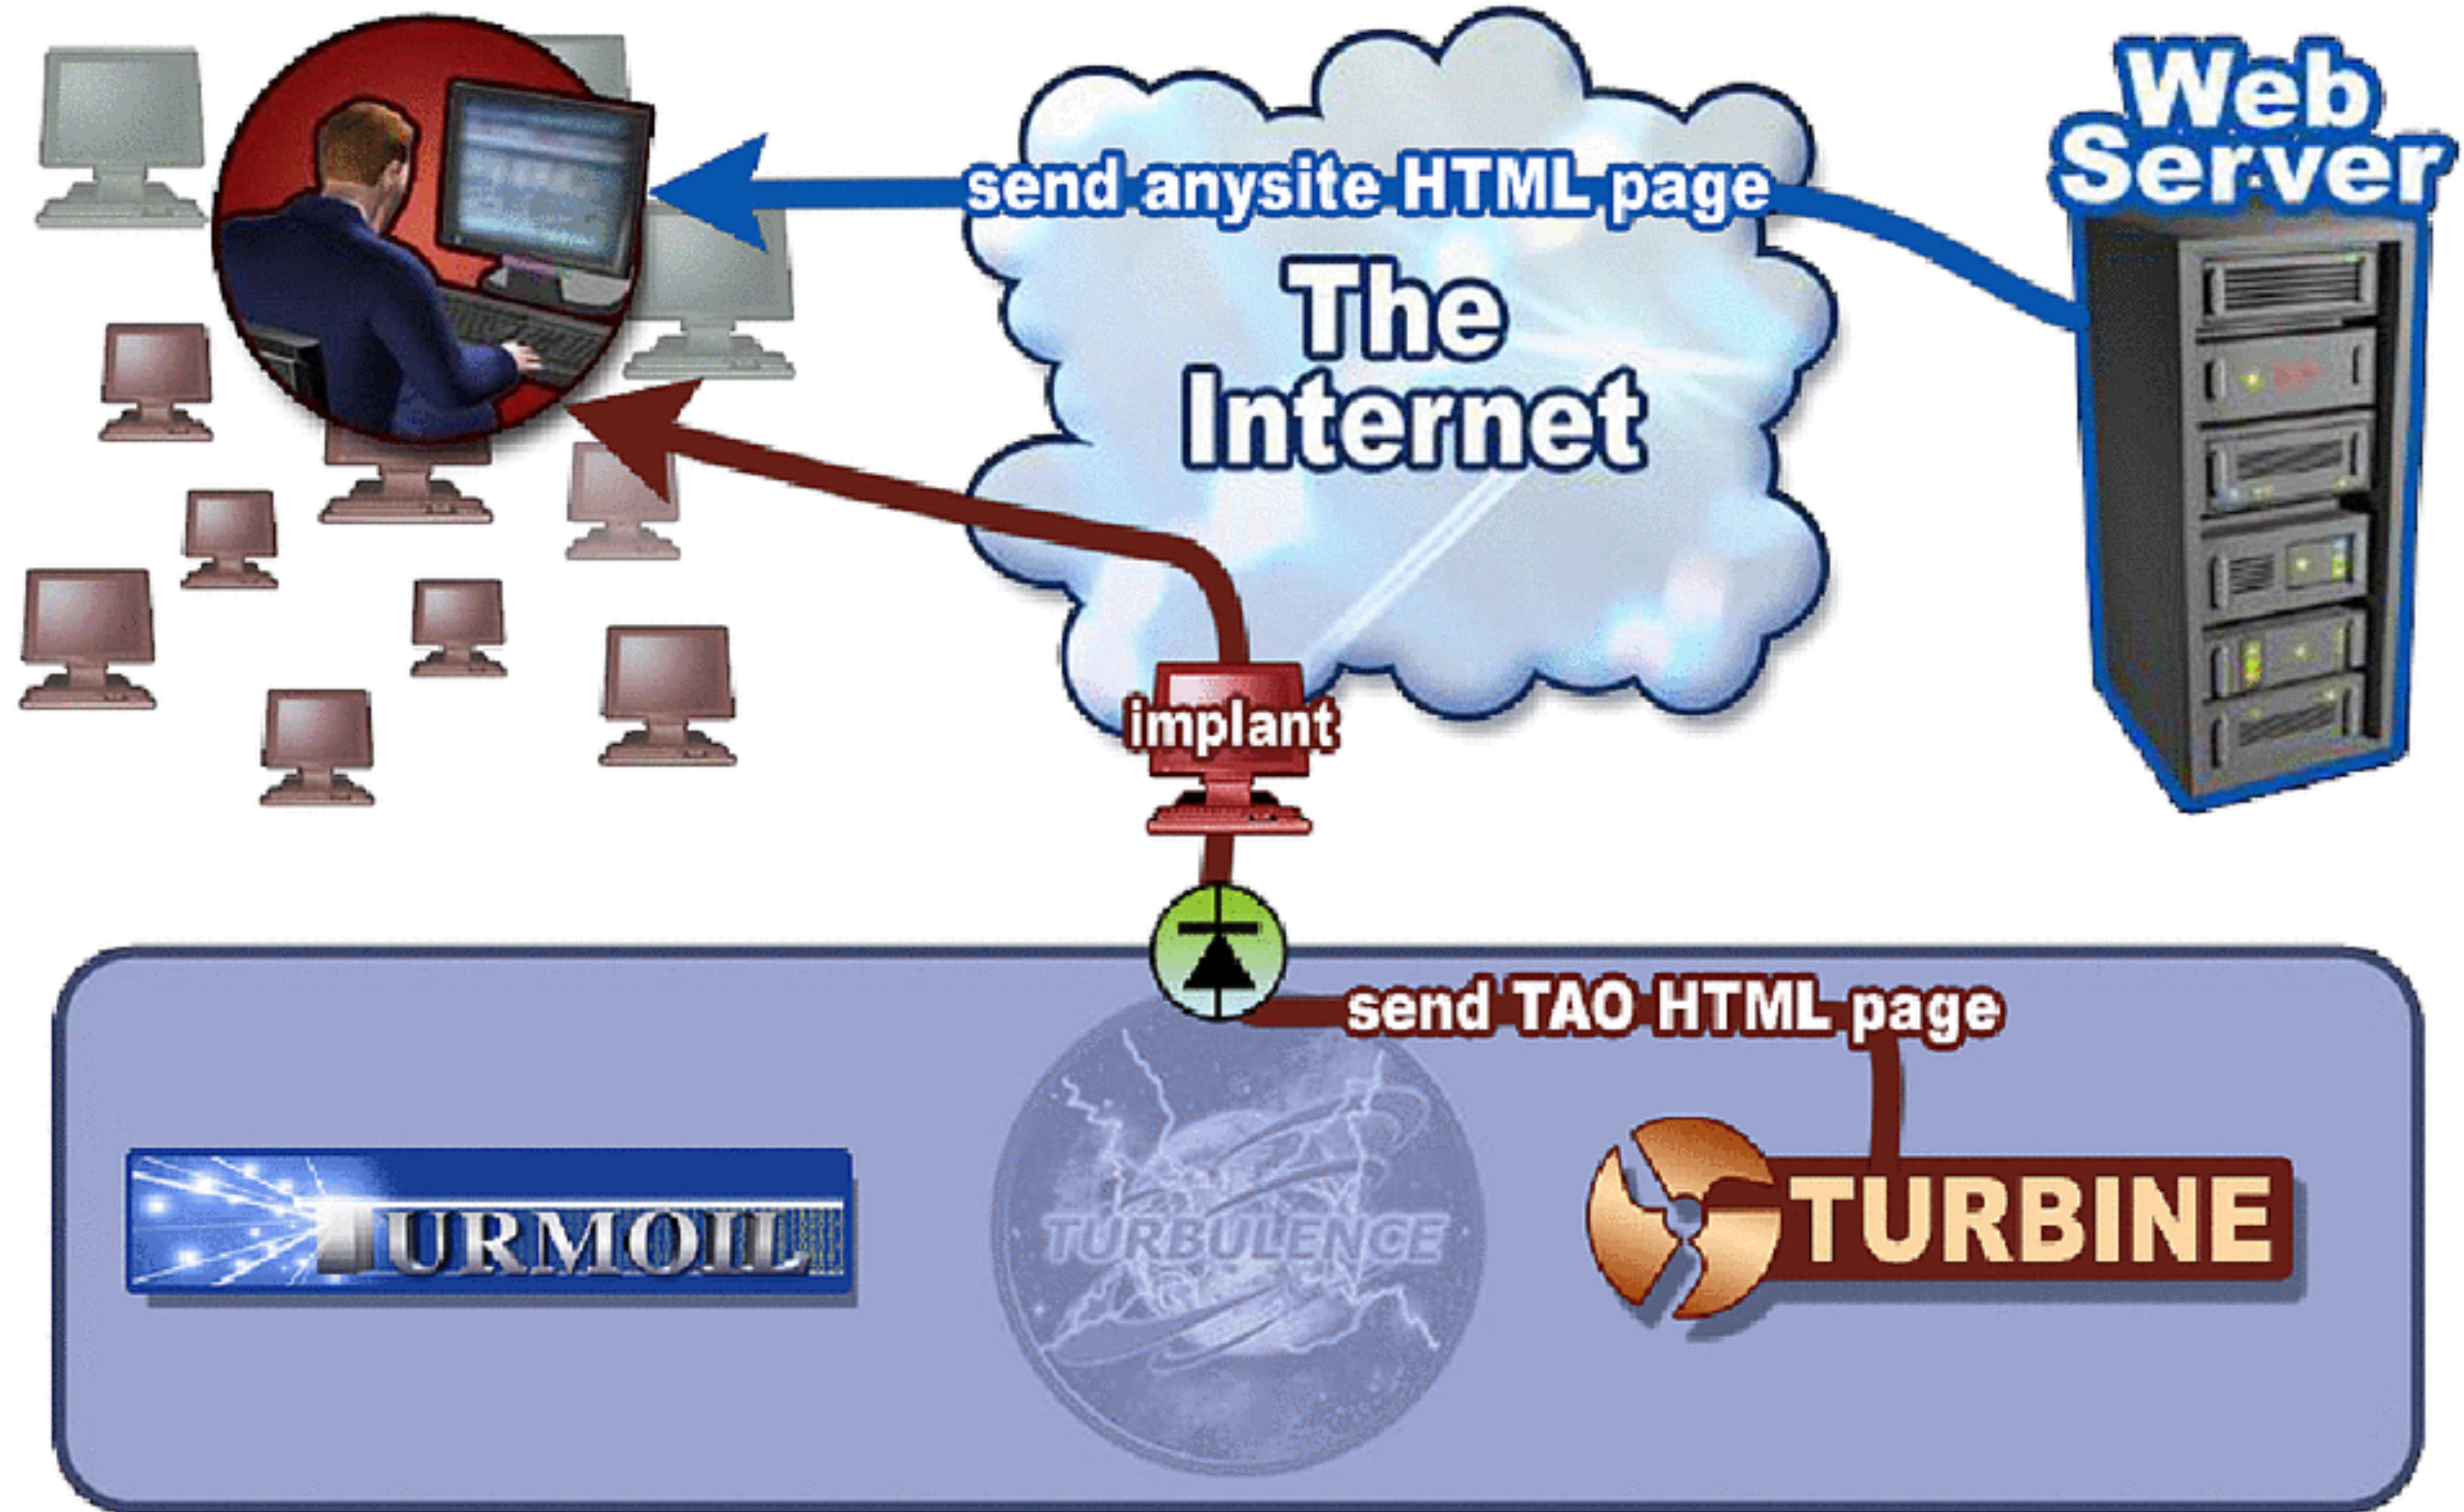
Our primary contributions are:

- A dataset that contains the fingerprints for 42,027 Netflix videos.
- An automated crawler that creates Netflix video fingerprints.
- A method to identify Netflix videos in real-time that does not rely on application-layer information.

We have made our code available at [4]. The rest of our paper is organized as follows. In Section 2, we describe the previous work that we leverage in our paper. In Section 3, we detail our method



# QUANTUMINSERT









Some were unhappy.



# Just one QUIC bit

By [Geoff Huston](#) on 28 Mar 2018

Category: [Tech matters](#)

Tags: [IETF](#), [NATs](#), [QUIC](#), [TCP](#)

[5 Comments](#)

[Like 77](#) [Share](#)

[Tweet](#) [Pocket](#)

[← Blog home](#)



[Original image credit: Eric Schmuttenmaer](#)

I'm never surprised by the ability of an IETF Working Group to obsess over what to any outside observer would appear to be a completely trivial matter. Even so, I was impressed to see a large-scale discussion emerge over a single bit in a transport protocol being standardized by the IETF. Is this an example of a severe overload of obsessive compulsive behaviour? Or does this single bit represent a major point of design principle, and was the extended discussion about that design principle rather than the use of the bit itself?

The transport protocol under consideration here is QUIC. QUIC was originally developed by Google, and is in use by their Chrome browser and by various Google servers. Given the extensive use of Chrome in the Internet, and the extensive use of Google services by Internet users, the obvious corollary is that QUIC is used extensively in the Internet.

<https://blog.apnic.net/2018/03/28/just-one-quic-bit/>



# How Google's QUIC Protocol Impacts Network Security and Reporting

QUIC is a new protocol designed by Google to make the web faster and more efficient. It's on by default in Google Chrome and used by a growing list of websites. Unfortunately, most, if not all, firewalls do not currently recognize QUIC traffic as 'web' traffic, therefore it is not inspected, logged or reported on, leaving a gaping hole in your network's security.

This article describes how QUIC works, its current consequences on network security and reporting, and how you can resolve the issues associated with QUIC.

## About QUIC

Google has always been obsessed with speed and over the years they have made numerous efforts to make the web more efficient and more performant. The new kid on the block for performance improvement is a protocol named **QUIC**. Where **SPDY** and **HTTP/2** were iterative improvements on HTTP over TCP, QUIC is a different approach using *UDP* as the transport protocol.

QUIC is essentially HTTP/2 over UDP which is a new layer4 protocol.

At the time of writing this article, QUIC is still 'experimental', but is enabled by default in Google Chrome, and can be enabled in Opera 16. Other browsers will surely follow once the protocol is finalized. It is implemented on all Google web properties such as Google Search, YouTube, Gmail, Drive etc, and is being adopted by a **growing list of other websites**.

## How QUIC Impacts Network Security

The issue is not with the protocol or the technology itself. The supposed upside of QUIC is that it makes web communications more efficient

DNS → DOH

# Redirecting DNS for Ads and Profit

Nicholas Weaver

*ICSI*

nweaver@icir.org

Christian Kreibich

*ICSI*

christian@icir.org

Vern Paxson

*ICSI & UC Berkeley*

vern@cs.berkeley.edu

## Abstract

Internet Service Providers (ISPs) increasingly try to grow their profit margins by employing “error traffic monetization,” the practice of redirecting customers whose DNS lookups fail to advertisement-oriented Web servers. A small industry of companies provides the associated machinery for ISPs to engage in this monetization, with the companies often participating in operating the service as well. We conduct a technical analysis of DNS error traffic monetization evident in 66,000 *Netalyzr* sessions, including fingerprinting derived from patterns seen in the resulting ad landing pages. We identify major players in this industry, their ISP affiliations over time, and available user opt-out mechanisms. **One monetization vendor, Paxfire, transgresses the error-based model and also reroutes all user search queries to Bing, Yahoo, and (sometimes) Google via proxy servers controlled or provided by Paxfire.**

In the *ICSI Netalyzr* [8], our widely used network debugging and diagnostic tool,<sup>2</sup> we have employed tests for various forms of NXDOMAIN wildcarding since we started offering the service in mid-2009. In this paper we illuminate the DNS error monetization market by combining *Netalyzr*’s measurements with an analysis of the redirection pages collected between January 2010 and May 2011, the location and content of the ad servers, and the marketing material provided by the companies involved. We identify ISPs employing DNS error monetization, their choice of monetization provider (including shifts of provider and apparent in-house realization), potential redirection policy customizations, as well as availability of opt-out mechanisms.

We also observe a more aggressive form of DNS-driven traffic manipulation, *search-engine proxying*.



# Compete CEO: ISPs Sell Clickstreams For \$5 A Month

Mar. 13, 2007 9:44 AM ET | [3 comments](#)



Henry Blodget 

Follow

(55 followers)

At the Open Data 2007 conference in New York today, David Cancel, the CEO of Compete Inc. revealed that ISPs happily sell clickstream data -- and that it's a big business. They don't sell your name -- just your clicks -- but the clicks are tied to you as a specific user (User 1, User 2, etc.).


How much are your clicks worth? About 40 cents a month per user (per customer)... and the Compete CEO estimates that there are 10-12 big buyers of this data. In other words, your ISP is probably making about \$5 a month (\$60 a year) off your clickstreams.

<https://seekingalpha.com/article/29449-competite-ceo-isps-sell-clickstreams-for-5-a-month>



# THE \$24 BILLION DATA BUSINESS THAT TELCOS DON'T WANT TO TALK ABOUT

Mobile Carriers Are Working With Partners to Manage, Package and Sell Data



By Kate Kaye. Published on October 26, 2015.

U.K. grocer Morrisons, ad-buying behemoth GroupM and other marketers and agencies are testing never-before-available data from cellphone carriers that connects device location and other information with telcos' real-world files on subscribers. Some services offer real-time heat maps showing the neighborhoods where store visitors go home at night, lists the sites they visited on mobile browsers recently and more.

Under the radar, Verizon, Sprint, Telefonica and other carriers have partnered with firms including SAP, IBM, HP and AirSage to manage, package and sell various levels of data to marketers and other clients. It's all part of a push by the world's largest phone operators to counteract diminishing subscriber growth through new business ventures that tap into the data that showers from consumers' mobile web surfing, text messaging and phone calls.




<https://adage.com/article/datadriven-marketing/24-billion-data-business-telcos-discuss/301058>

# Senate Republicans Vote to Allow ISPs to Sell Your Private Data

Privacy watchdogs responded to the Senate vote with outrage.

---

By [Sam Gustin](#)

Mar 24 2017, 5:50am  Share  Tweet  Snap

---

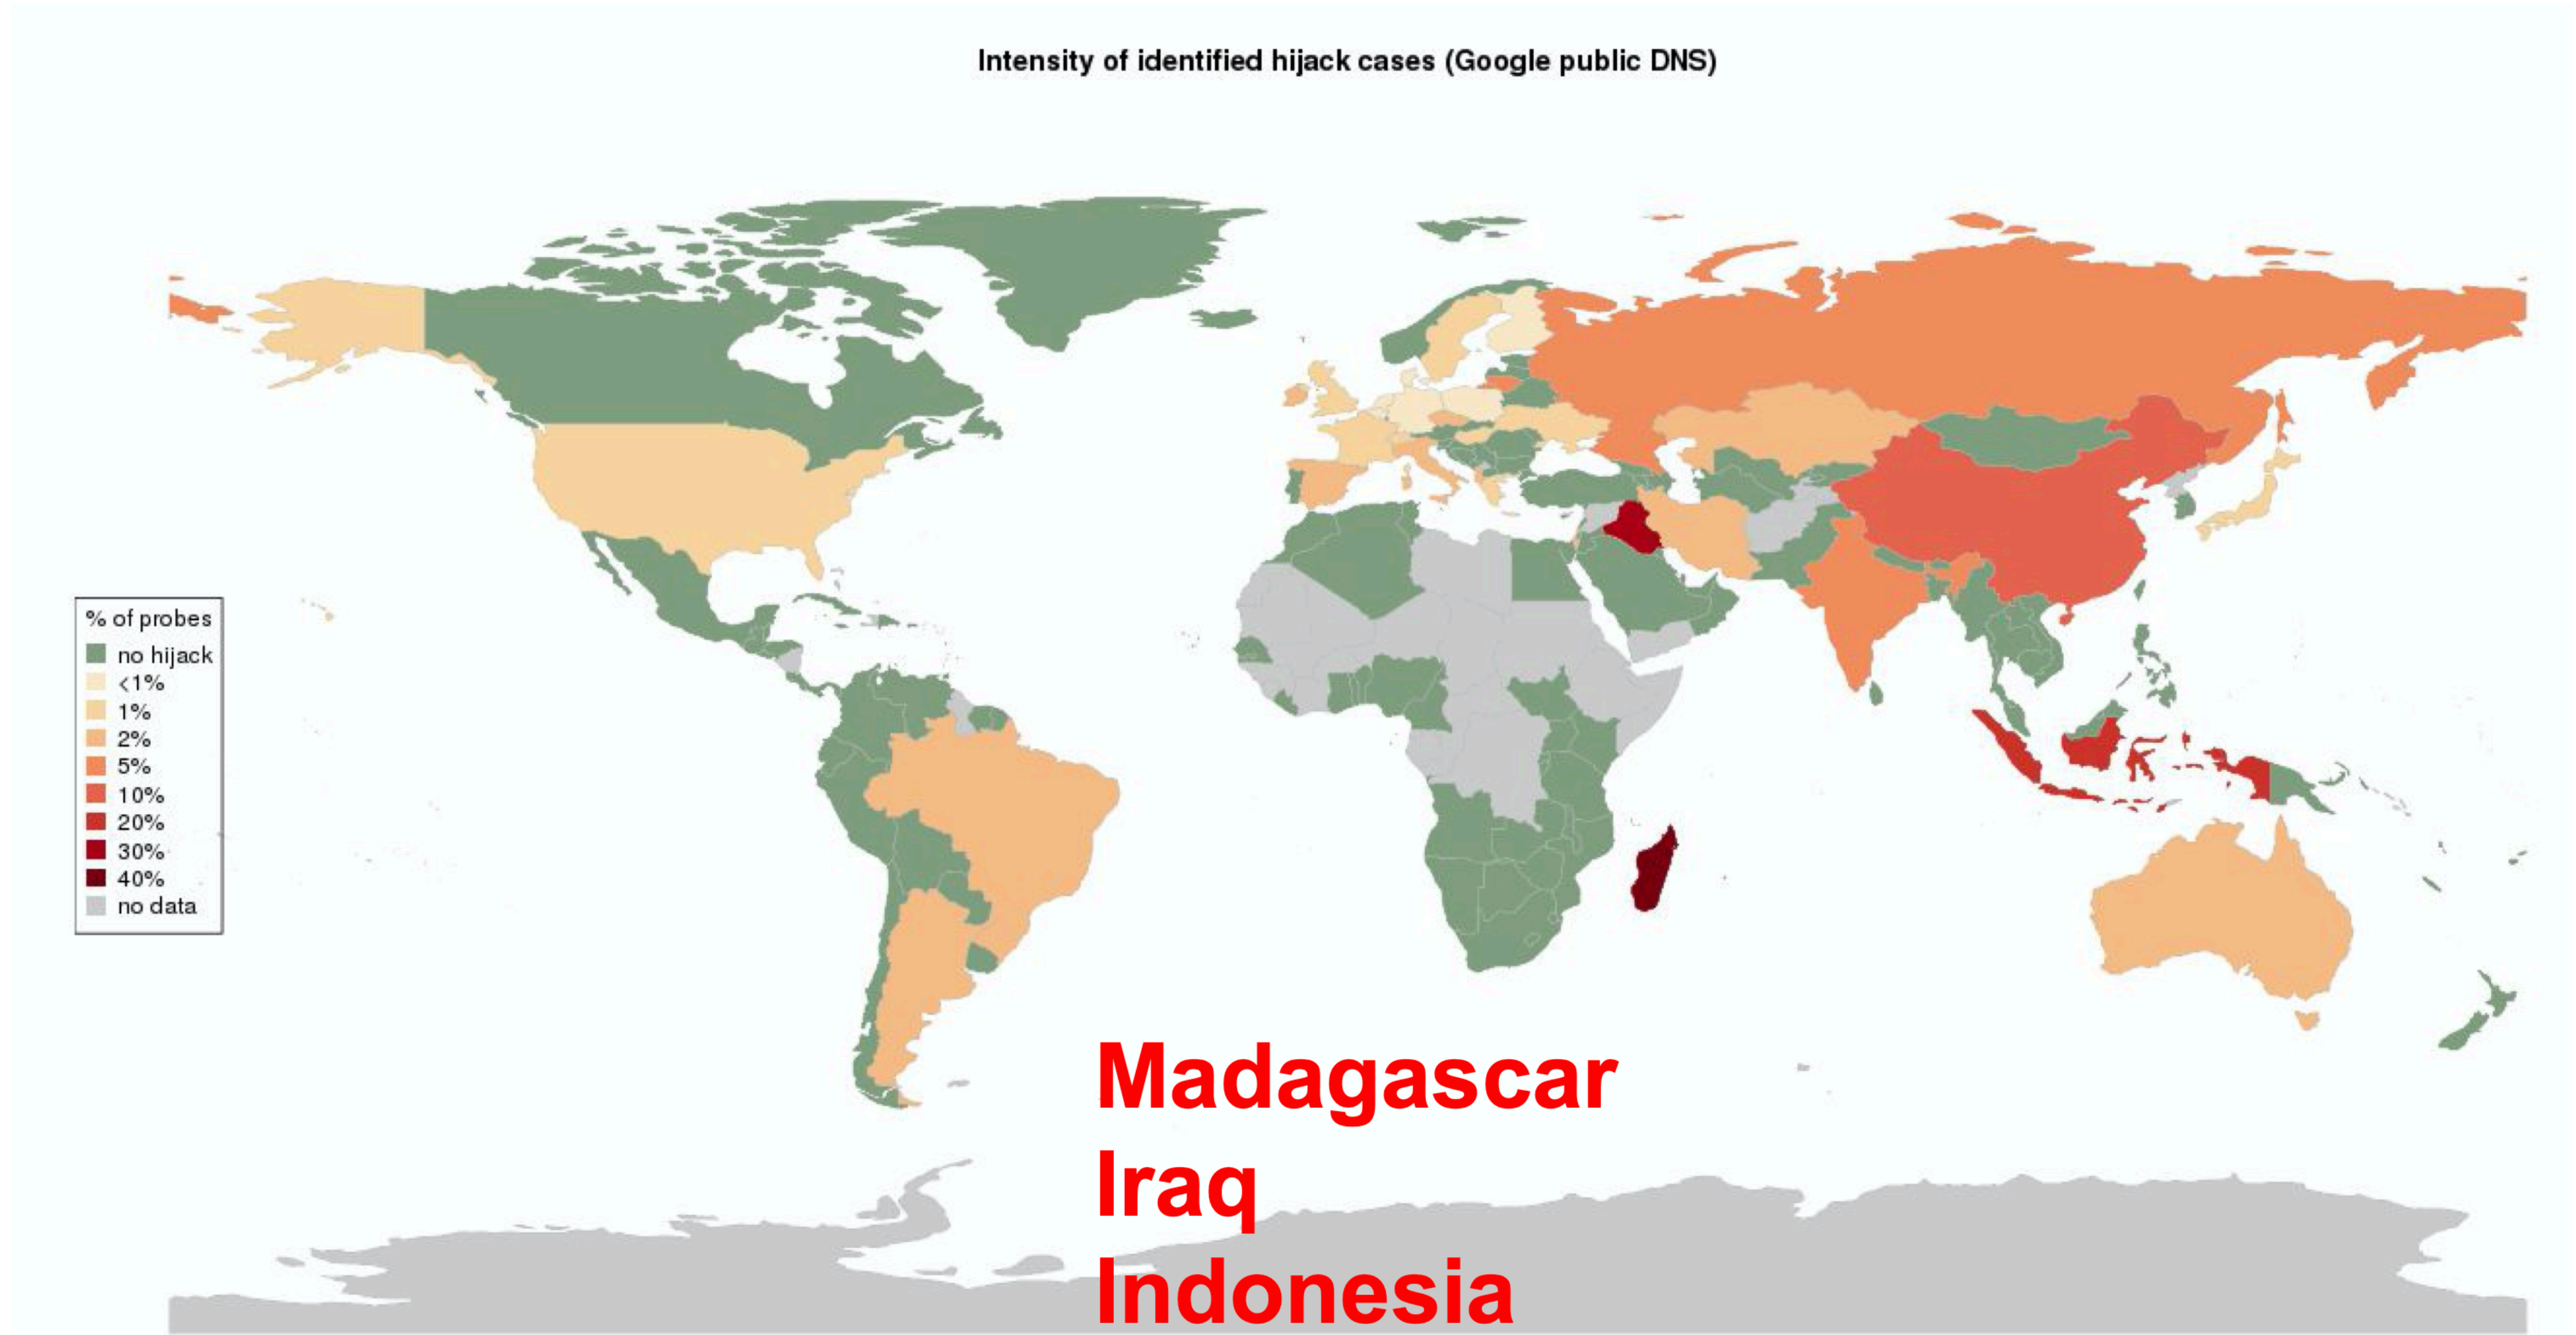
Republican lawmakers in the US Senate approved a measure on Thursday designed to kill **federal broadband privacy protections** and allow internet service providers like AT&T and Verizon to sell your sensitive private information to the highest bidder.

The vote represents the culmination of a **year-long campaign** by the nation's largest internet service providers (ISPs) and their GOP allies to torpedo Federal Communications Commission rules that require broadband providers to obtain "opt-in" consent before using, sharing, or selling private consumer data.

[https://www.vice.com/en\\_us/article/bmbkym/senate-republicans-vote-to-allow-isps-to-sell-your-private-data](https://www.vice.com/en_us/article/bmbkym/senate-republicans-vote-to-allow-isps-to-sell-your-private-data)



# Results: Google DNS hijacks (%)



**Madagascar**  
**Iraq**  
**Indonesia**  
**China**

<https://bit.ly/irtf-dns-hijacking>



[\[Docs\]](#) [\[txt|pdf\]](#) [\[draft-ietf-doh-...\]](#) [\[Tracker\]](#) [\[Diff1\]](#) [\[Diff2\]](#) [\[Errata\]](#)

PROPOSED STANDARD

**Errata Exist**

Internet Engineering Task Force (IETF)

P. Hoffman

Request for Comments: 8484

ICANN

Category: Standards Track

P. McManus

ISSN: 2070-1721

Mozilla

October 2018

## DNS Queries over HTTPS (DoH)

### Abstract

This document defines a protocol for sending DNS queries and getting DNS responses over HTTPS. Each DNS query-response pair is mapped into an HTTP exchange.

Some were unhappy.

# Warning over Google Chrome browser's new threat to children

Nicholas Hellen and Richard Kerbaj

April 21 2019, 12:01 am,  
The Sunday Times

Politics

Google

Television

Technology



The new version of Google Chrome will bypass most parental control systems  
ELVA ETIENNE/GETTY IMAGES

Internet safety watchdogs and intelligence agencies are holding crisis talks about a new version of Britain's most popular web browser, which they fear will endanger children.

They say Google's plans to encrypt Chrome will make it harder to block harmful material, including child-abuse images and terrorist propaganda. The new version will bypass most parental control systems and undermine the government's attempts to stop under-18s viewing pornography.

<https://bit.ly/times-doh-children>



## ISPA announces finalists for 2019 Internet Heroes and Villains: Trump and Mozilla lead the way as Villain nominees

Posted on 2nd July 2019

The Internet Services Providers' Association is pleased to announce the finalists for the 2019 Internet Hero and Villain.

At a time where technology and the Internet has become fully mainstream and a driver of innovation and growth, the policy challenges presented by this disruption are now some of the biggest issues facing policymakers around the world.

The Internet Hero nominations this year include those campaigning to improve trust and confidence online; mapping out the UK's evolving broadband landscape; and working on global internet governance issues. While, the Villain nominees take in the impact of new technical standards on existing online protections, the balance between freedom of expression and copyright online and the global telecoms supply chain.

Following weeks of consultation and a large range of nominations received via email and Twitter from the public, this year's nominations for the 2019 Internet Heroes and Villains in full are:

### ISPA Internet Hero

**Sir Tim Berners-Lee** – for spearheading the 'Contract for the Web' campaign to rebuild trust and protect the open and free nature of the Internet in the 30<sup>th</sup> anniversary of the World Wide Web

**Andrew Ferguson OBE**, Editor, Thinkbroadband - for providing independent analysis and valuable data on the UK broadband market since the year 2000

**Oscar Tapp-Scotting & Paul Blaker**, Global Internet Governance Team, DCMS – for leading the UK Government's efforts to ensure a balanced and proportionate agenda at the International Telecommunications Union Conference

### ISPA Internet Villain

**Mozilla** – for their proposed approach to introduce DNS-over-HTTPS in such a way as to bypass UK filtering obligations and parental controls, undermining internet safety standards in the UK

**Article 13 Copyright Directive** – for threatening freedom of expression online by requiring 'content recognition technologies' and

**President Donald Trump** – for causing a huge amount of uncertainty across the complex, global telecommunications supply chain and for failing to protect national security

It looks like your cookies are switched off. To ensure the best experience whilst visiting our website please consider allowing cookies.



## Summary

---

- However, Google has announced **unilateral plans** (along with Mozilla, which derives over 90% of its revenue from Google) to activate DoH in its Chrome browser as soon as October. Google also appears poised to activate DoT for devices using its Android mobile operating system in the near future.
- If activated, this feature would by default route all DNS traffic from Chrome and Android users to Google Public DNS, thus **centralizing a majority of worldwide DNS data with Google**.
- This change would mark a **fundamental shift** in the decentralized nature of the Internet's architecture and give one provider control of Internet traffic routing and vast amounts of new data about consumers and competitors.
- The unilateral centralization of DNS raises serious policy issues relating to **cybersecurity**, privacy, antitrust, **national security and law enforcement**, **network performance and service quality (including 5G)**, and other areas.





**Paul Vixie**  
@paulvixie



Replying to [@grittygrease](#)

Rfc 8484 is a cluster duck for internet security. Sorry to rain on your parade. The inmates have taken over the asylum.

8:49 AM · Oct 21, 2018 · [Twitter Web App](#)

---

**24** Retweets   **72** Likes





## ^ Mozilla Policy Requirements for DNS over HTTPs Partners

This document describes the minimum set of policy requirements that a party must satisfy to be considered as a potential partner for Mozilla's Trusted Recursive Resolver (TRR) program. It specifically describes data collection and retention, transparency, and blocking policies and is in addition to any contractual, technical or operational requirements necessary to operate the resolver service.

### Privacy Requirements

Mozilla's TRR is intended to provide better, minimum privacy guarantees to Firefox users than current, ad hoc provisioning of DNS services. As such, resolvers must strictly limit data collection and sharing from the resolver. More specifically:

1. The resolver may retain user data (including identifiable data, data associated with user IP addresses, and any non-aggregate anonymized data) but should do so only for the purpose of operating the service and must not retain that data for longer than 24 hours.
  - Only aggregate data that does not identify individual users or requests may be retained beyond 24 hours.
2. The resolver must not retain, sell, or transfer to any third party (except as may be required by law) any personal information, IP addresses or other user identifiers, or user query patterns from the DNS queries sent from the Firefox browser.
3. The resolver must not combine the data that it collects from queries with any other data in any way that can be used to identify individual end users.
4. The resolver must not sell, license, sublicense, or grant any rights to user data to any other person or entity.
5. The resolver must support DNS Query Name Minimisation as defined in [RFC 7816](#).
6. The resolver must not propagate unnecessary information about queries to authoritative name servers. In particular, the client subnet DNS extension in [RFC 7871](#) must not be sent to servers unless the connection to the authoritative server is encrypted and only to authoritative name servers operated by the domain owner directly or by a DNS provider pursuant to its contract with the domain owner.

### Transparency Requirements

[\[Docs\]](#) [\[txt|pdf\]](#) [\[Tracker\]](#) [\[Email\]](#) [\[Diff1\]](#) [\[Diff2\]](#) [\[Nits\]](#) [\[IPR\]](#)

Versions: [00](#) [01](#)

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: May 4, 2020

E. Kinnear  
T. Pauly  
C. Wood  
Apple Inc.  
P. McManus  
Fastly

November 01, 2019

**Adaptive DNS: Improving Privacy of Name Resolution**  
**draft-pauly-dprive-adaptive-dns-privacy-01**

Abstract

This document defines an architecture that allows clients to dynamically discover designated resolvers that offer encrypted DNS services, and use them in an adaptive way that improves privacy while co-existing with locally provisioned resolvers. These resolvers can be used directly when looking up names for which they are designated. These resolvers also provide the ability to proxy encrypted queries, thus hiding the identity of the client requesting resolution.

**SNI → Encrypted SNI**



### 3. Server Name Indication

TLS does not provide a mechanism for a client to tell a server the name of the server it is contacting. It may be desirable for clients to provide this information to facilitate secure connections to servers that host multiple 'virtual' servers at a single underlying network address.

In order to provide any of the server names, clients MAY include an extension of type "server\_name" in the (extended) client hello. The "extension\_data" field of this extension SHALL contain "ServerNameList" where:

```
struct {
    NameType name_type;
    select (name_type) {
        case host_name: HostName;
    } name;
} ServerName;

enum {
    host_name(0), (255)
} NameType;

opaque HostName<1..2^16-1>;

struct {
    ServerName server_name_list<1..2^16-1>
} ServerNameList;
```

[\[Docs\]](#) [\[txt|pdf\]](#) [\[Tracker\]](#) [\[WG\]](#) [\[Email\]](#) [\[Diff1\]](#) [\[Diff2\]](#) [\[Nits\]](#)

Versions: ([draft-rescorla-tls-esni](#)) [00](#) [01](#) [02](#)  
[03](#) [04](#) [05](#)

tls  
Internet-Draft  
Intended status: Experimental  
Expires: May 7, 2020

E. Rescorla  
RTFM, Inc.  
K. Oku  
Fastly  
N. Sullivan  
Cloudflare  
C. Wood  
Apple, Inc.  
November 04, 2019

## **Encrypted Server Name Indication for TLS 1.3** **draft-ietf-tls-esni-05**

### Abstract

This document defines a simple mechanism for encrypting the Server Name Indication for TLS 1.3.

Some were unhappy.



## South Korea is Censoring the Internet by Snooping on SNI Traffic

By [Sergiu Gatlan](#)

February 13, 2019 06:19 PM 1

South Korea has been blocking HTTP websites that are on their censor list for a while now and they have recently started using SNI filtering to block their counterparts served over HTTPS.

A warning page bearing the seals of the Korea Communications Standards Commission (KCSC) and the Korean National Police Agency is displayed for blocked HTTP websites, while TLS sites blocked using Server Name Indication (SNI) filtering will only throw a "This site can't be reached" error.

An [OpenNet Initiative report](#) from 2012 is still valid although quite dated given that the country has not updated its Internet surveillance approach since 2008, and it paints an accurate picture of the current state of Internet censorship in South Korea:

Despite the fact that South Korea has one of the most advanced information communication technology sectors in the world, online expression remains under the strict legal and technological control of the central government. The country is the global leader in Internet connectivity and speed, but its restrictions on what Internet users can access are substantial.

Also, [Reporters Without Borders included](#) South Korea on its list of countries "Under Surveillance" during 2011, and it also compared the level of Internet censorship to those experienced by citizens of Russia and Egypt in its "Enemies of the Internet" report, as described by [The New York Times in 2012](#).

### SNI filtering used to block websites

SNI is a TLS extension which allows browsers to inform a web server of the hostname they want to connect to at the beginning of the handshaking process, as detailed in [IETF's RFC3546](#).


As [reported](#) by JoongAng Media Network Group's Lee Min Jung and [other sources](#), South Korea has begun filtering the country's internet traffic to block TLS websites blacklisted by the KCSC.

# Traffic Analysis

- Packet length
- Packet timing
- Host(s) communicating



# Encrypted Traffic Analytics (ETA)



## Higher precision, faster investigation

Leverage the latest Cisco networking capabilities to avoid, stop, or mitigate threats faster than ever before. Cisco Digital Network Architecture (Cisco DNA) is the industry's first network with the ability to find threats in encrypted traffic.

[Watch video \(1:59\)](#)

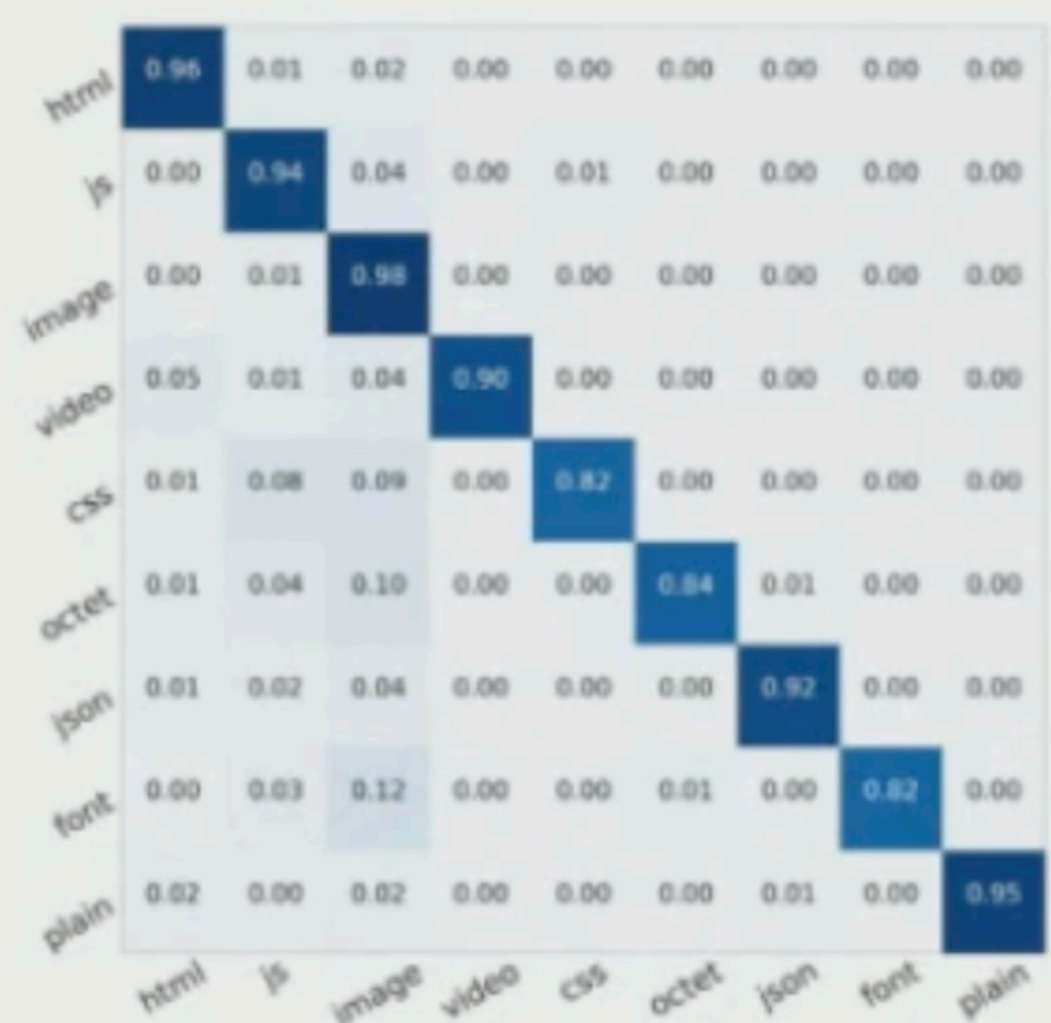


# IETF105

## ANRW: Network Functions and Middleboxes




### Results - Content-Type



(a) chrome



(b) malware



(c) tor



# Conclusions

- ▶ Detailed inferences about the encrypted HTTP protocol are possible with careful dataset construction and feature selection
- ▶ Multiplexing and fixed-length records provide a valuable defense against these techniques
- ▶ Results are client dependent; TLS fingerprinting can provide guidance



Some will be unhappy.

# The Other End (and Yours)








### 3. The Internet Threat Model

A THREAT MODEL describes the capabilities that an attacker is assumed to be able to deploy against a resource. It should contain such information as the resources available to an attacker in terms of information, computing capability, and control of the system. The purpose of a threat model is twofold. First, we wish to identify the threats we are concerned with. Second, we wish to rule some threats explicitly out of scope. Nearly every security system is vulnerable to a sufficiently dedicated and resourceful attacker.

The Internet environment has a fairly well understood threat model. In general, we assume that the end-systems engaging in a protocol exchange have not themselves been compromised. Protecting against an attack when one of the end-systems has been compromised is

extraordinarily difficult. It is, however, possible to design protocols which minimize the extent of the damage done under these circumstances.



Your browser fingerprint **appears to be unique** among the 209,562 tested in the past 45 days.

Currently, we estimate that your browser has a fingerprint that conveys **at least 17.68 bits of identifying information**.

The measurements we used to obtain this result are listed below. You can **read more about our methodology, statistical results, and some defenses against fingerprinting here**.

Browser Characteristic	bits of identifying information	one in <i>x</i> browsers have this value	value
User Agent	7.37	165.01	Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_2) AppleWebKit/605.1.15 (KHTML, like Gecko) Version/13.0.4 Safari/605.1.15
HTTP_ACCEPT Headers	9.14	564.86	text/html, */*; q=0.01 gzip, deflate, br en-au
Browser Plugin Details	6.44	86.56	Plugin 0: WebKit built-in PDF; ; ; (Portable Document Format; application/pdf; pdf) (Portable Document Format; text/pdf; pdf) (PostScript; application/postscript; ps).
Time Zone	6.58	95.47	-660
Screen Size and Color Depth	5.17	35.98	2560x1440x24
System Fonts	3.98	15.78	Andale Mono, Arial, Arial Black, Arial Hebrew, Arial Narrow, Arial Rounded MT Bold, Arial Unicode MS, Comic Sans MS, Courier, Courier New, Geneva, Georgia, Helvetica, Helvetica Neue, Impact, LUCIDA GRANDE, Microsoft Sans Serif, Monaco, Palatino, Tahoma, Times, Times New Roman, Trebuchet MS, Verdana, Wingdings, Wingdings 2, Wingdings 3 (via javascript)
Are Cookies Enabled?	0.22	1.16	Yes
Limited supercookie test	0.34	1.26	DOM localStorage: Yes, DOM sessionStorage: Yes, IE userData: No
Hash of canvas fingerprint	6.96	124.37	8053b1c9d7560455edd6254c3582727e
Hash of WebGL fingerprint	2.98	7.87	00000000000000000000000000000000
DNT Header Enabled?	1.11	2.15	False
Language	7.38	166.06	en-AU
Platform	3.23	9.36	MacIntel
Touch Support	0.75	1.68	Max touchpoints: 0; TouchEvent supported: false; onTouchStart supported: false



# Telecommunications and surveillance annual reports

The Minister for Home Affairs issues annual reports on the use of telecommunications interception and surveillance devices by Australian agencies under the [Telecommunications \(Interception and Access\) Act 1979](#) <sup>↗</sup> and [Surveillance Devices Act 2004](#) <sup>↗</sup>.

These reports include information and statistics on:

- the relevant legislation
- agencies that have intercepted or accessed telecommunications, or used surveillance devices
- the type of warrants applied for, or the type of surveillance devices used legal developments since the last reporting period
- the number of prosecutions and convictions resulting from the use of intercepted or accessed telecommunications information, or from the use of surveillance devices

## Telecommunications (Interception and Access) Act

- [Telecommunications \(Interception and Access\) Act 1979 – Annual Report for the year ending 30 June 2018 \(2MB PDF\)](#)
- [Telecommunications \(Interception and Access\) Act 1979 – Annual Report for the year ending 30 June 2017 \(3MB PDF\)](#)
- [Telecommunications \(Interception and Access\) Act 1979 – Annual Report for the year ending 30 June 2016 \(1076KB PDF\)](#)
- [Telecommunications \(Interception and Access\) Act 1979 – Annual Report for the year ending 30 June 2015 \(2MB PDF\)](#)
- [Telecommunications \(Interception and Access\) Act 1979 – Annual Report for the year ending 30 June 2014 \(2MB PDF\)](#)



↻ Internet of Shit Retweeted



**Rory Sutherland** ✓ @rorysutherland · Dec 1, 2019



I promise I am not making this up. I have just been asked to upgrade the firmware for my new toilet.



💬 161

↻ 341

♥ 1.7K






# How Your Phone Betrays Democracy

By Charlie Warzel and Stuart A. Thompson

DEC. 21, 2019

“By tracking specific devices, we followed demonstrators from the 2017 Women’s March back to their homes. We were able to identify individuals at the 2017 Inauguration Day Black Bloc protests. It was easy to follow them to their workplaces. In some instances — for example, a February clash between antifascists and far-right supporters of Milo Yiannopolous in Berkeley, Calif. — it took little effort to identify the homes of protesters and then their family members.”



<https://www.nytimes.com/interactive/2019/12/21/opinion/location-data-democracy-protests.html>



## *Barr Asks Apple to Unlock Pensacola Killer's Phones, Setting Up Clash*

The request set up a collision between law enforcement and big technology firms in the latest battle over privacy and security.



Attorney General William P. Barr is pushing Apple to unlock the phones of the gunman behind a deadly shooting at a naval air station in Pensacola, Fla.  
Calla Kessler/The New York Times



By **Katie Benner**

Jan. 13, 2020 Updated 5:42 p.m. ET



WASHINGTON — Attorney General William P. Barr declared on Monday that a deadly shooting last month at a naval air station in Pensacola, Fla., was an act of terrorism, and he asked Apple in an unusually high-profile request to provide access to two phones used by the gunman.

<https://www.nytimes.com/2020/01/13/us/politics/pensacola-shooting-iphones.html>

# Intelligent Tracking Prevention 2.3

Sep 23, 2019

by John Wilander

[@johnwilander](#)

Note: Read about past updates to this technology in [other blog posts about Intelligent Tracking Prevention, the Storage Access API, and ITP Debug Mode.](#)

Intelligent Tracking Prevention (ITP) version 2.3 is included in Safari on iOS 13, the iPadOS beta, and Safari 13 on macOS for Catalina, Mojave, and High Sierra.

## Enhanced Prevention of Tracking Via Link Decoration

Our previous release, [ITP 2.2](#), focused specifically on the abuse of so-called link decoration for the purposes of cross-site tracking. With ITP 2.2, when a webpage is navigated to from a domain classified by ITP and the landing URL has a query string or fragment, the expiry of persistent client-side cookies created on that page is 24 hours.

Unfortunately, we see continued abuse of link decoration, so ITP 2.3 takes two new steps to combat this.

## Capped Lifetime For All Script-Writeable Website Data

Since ITP 2.2, several trackers have announced their move from first-party cookies to alternate first-party storage such as LocalStorage. ITP 2.3 counteracts this in the following way:



Some will be unhappy.

APPLE

# Apple's Ad-Targeting Crackdown Shakes Up Ad Market

By [Tom Dotan](#) Dec. 9, 2019 7:01 AM PST • Comments by [Sutha Kamal](#) and [Richard Reisman](#)

[Subscribe now](#)

Two years ago, Apple launched an aggressive battle against ads that track users across the web. Today executives in the online publishing and advertising industries say that effort has been stunningly effective—posing a problem for advertisers looking to reach affluent consumers.

Since Apple introduced what it calls its Intelligent Tracking Prevention feature in September 2017, and with subsequent updates last year, advertisers have largely lost the ability to target people on Safari based on their browsing habits with cookies, the most commonly used technology for tracking. One result: The cost of reaching Safari users has fallen over 60% in the past two years, according to data from ad tech firm Rubicon Project. Meanwhile ad prices on Google's Chrome browser have risen slightly.



Apple CEO Tim Cook. Photo by Bloomberg

# 4. Some Observations



# Cost and Control

“The network is a dumb pipe” vs.  
“Those are the rules on *my* network.”

*We have* to design the Internet  
for the pessimal case.




Assuming latent availability of application data on-path is no longer viable.

What does  
“trust my [employer, ISP, government]  
to do X” look like?

Keychains

Category

- All Items
- Passwords
- Secure Notes
- My Certificates
- Keys
- Certificates



**AAA Certificate Services**  
 Root certificate authority  
 Expires: Monday, 1 January 2029 at 9:59:59 am Australian Eastern Standard Time  
 This certificate is valid

Name	Kind	Date Modified	Expires	Keychain
AAA Certificate Services	certificate	--	1 Jan 2029 at 9:59:59 am	System Roots
AC RAIZ FNMT-RCM	certificate	--	1 Jan 2030 at 10:00:00 am	System Roots
Actalis Authentication Root CA	certificate	--	22 Sep 2030 at 9:22:02 p...	System Roots
AddTrust Class 1 CA Root	certificate	--	30 May 2020 at 8:38:31...	System Roots
AddTrust External CA Root	certificate	--	30 May 2020 at 8:48:38...	System Roots
Admin-Root-CA	certificate	--	10 Nov 2021 at 5:51:07 pm	System Roots
AffirmTrust Commercial	certificate	--	1 Jan 2031 at 12:06:06 am	System Roots
AffirmTrust Networking	certificate	--	1 Jan 2031 at 12:08:24 am	System Roots
AffirmTrust Premium	certificate	--	1 Jan 2041 at 12:10:36 am	System Roots
AffirmTrust Premium ECC	certificate	--	1 Jan 2041 at 12:20:24 am	System Roots
Amazon Root CA 1	certificate	--	17 Jan 2038 at 10:00:00...	System Roots
Amazon Root CA 2	certificate	--	26 May 2040 at 10:00:00...	System Roots
Amazon Root CA 3	certificate	--	26 May 2040 at 10:00:00...	System Roots
Amazon Root CA 4	certificate	--	26 May 2040 at 10:00:00...	System Roots
ANF Global Root CA	certificate	--	6 Jun 2033 at 3:45:38 am	System Roots
Apple Root CA	certificate	--	10 Feb 2035 at 7:40:36 am	System Roots
Apple Root CA - G2	certificate	--	1 May 2039 at 4:10:09 am	System Roots
Apple Root CA - G3	certificate	--	1 May 2039 at 4:19:06 am	System Roots
Apple Root Certificate Authority	certificate	--	10 Feb 2025 at 10:18:14 a...	System Roots
ApplicationCA2 Root	certificate	--	13 Mar 2033 at 1:00:00 am	System Roots
Atos TrustedRoot 2011	certificate	--	1 Jan 2031 at 9:59:59 am	System Roots
Autoridad de Certificacion Firmaprofesional CIF A62634068	certificate	--	31 Dec 2030 at 6:38:15 pm	System Roots
Autoridad de Certificacion Raiz del Estado Venezolano	certificate	--	18 Dec 2030 at 9:59:59 a...	System Roots
Baltimore CyberTrust Root	certificate	--	13 May 2025 at 9:59:00 a...	System Roots
Belgium Root CA2	certificate	--	15 Dec 2021 at 6:00:00 pm	System Roots
Buypass Class 2 Root CA	certificate	--	26 Oct 2040 at 6:38:03 p...	System Roots
Buypass Class 3 Root CA	certificate	--	26 Oct 2040 at 6:28:58 p...	System Roots
CA Disig Root R1	certificate	--	19 Jul 2042 at 7:06:56 pm	System Roots
CA Disig Root R2	certificate	--	19 Jul 2042 at 7:15:30 pm	System Roots
Certigna	certificate	--	30 Jun 2027 at 1:13:05 am	System Roots
Certinomis - Autorité Racine	certificate	--	17 Sep 2028 at 6:28:59 pm	System Roots
Certinomis - Root CA	certificate	--	21 Oct 2033 at 7:17:18 pm	System Roots
Certplus Root CA G1	certificate	--	15 Jan 2038 at 10:00:00...	System Roots
Certplus Root CA G2	certificate	--	15 Jan 2038 at 10:00:00...	System Roots
certSIGN ROOT CA	certificate	--	5 Jul 2031 at 3:20:04 am	System Roots



# Well-defined interfaces and counterbalanced roles

Technology and policy need to work together – *and* keep each other in check.

**Making some people unhappy means  
you need some guiding principles.**



Internet Architecture Board (IAB)

Internet-Draft

Intended status: Informational

Expires: May 10, 2020

M. Nottingham

November 7, 2019

## Table of Contents

1. Introduction
2. What Are “End Users”?
3. Why The IETF Should Prioritise End Users
4. How The IETF Can Prioritise End Users
  - 4.1. Engaging the Internet Community
  - 4.2. Creating User-Focused Systems
  - 4.3. Designing for Positive User Outcomes
  - 4.4. Identifying Negative End User Impact
  - 4.5. Handling Conflicting End User Needs
  - 4.6. Deprioritising Internal Needs
5. IANA Considerations
6. Security Considerations
7. Informative References
- Appendix A. Acknowledgements
- Author's Address

# The Internet is for End Users

draft-iab-for-the-users-latest

## Abstract

This document explains why the IAB believes the IETF should consider end users as its highest priority concern, and how that can be done.

## Note to Readers

The issues list for this draft can be found at <https://github.com/intarchboard/for-the-users/issues>.

The most recent (often, unpublished) draft is at <https://intarchboard.github.io/for-the-users/>.

# 5. What you can do

- **Secure your protocols** even if you don't think you need to (for herd immunity / defining “normal” on the Internet)
- **Think carefully** about the software, hardware and services you use, and who they are working for
- **Raise your voice** to vendors, governments, ISPs
- **Stay informed** and counter misinformation where you see it
- **Get involved** in Internet and Web standards



# Keeping the Watchers at Bay

Mark Nottingham @mnot linux.conf.au 2020