WebSentinel

FORENSIC PRIVACY AUDIT - AGGREGATE REPORT

We take your privacy seriously.

A forensic report on website surveillance on the top 100 English-content websites in the world.

518,770
EVENTS THAT IMPACT PRIVACY

100
TOP 100 ENGLISH-CONTENT WEBSITES GLOBALLY

0
CONSENTS GIVEN

Captured from Malta, 1st to 3rd October 2026 - published 17th November 2026
That Privacy Guy! - websentinel.eu

Contents

1. Summary of findings
    Key points
2. Scope
3. Methodology
    3.1 Site selection
    3.2 The capture instrument
    3.3 Capture protocol
    3.4 Evidence integrity
    3.5 Limitations
4. Findings
    4.0 How to read the findings
    4.00 Who is responsible: joint controllership
    4.1 Device fingerprinting
    4.2 CNAME cloaking
    4.3 Script loads and sideloaded scripts
    4.4 Cookies
    4.5 Local storage, session storage and other browser storage
    4.6 Beacons, keep-alive requests, tracking pixels and link decoration
    4.7 Referrer leakage
    4.8 Cross-origin messaging, tracking frames and other browser mechanisms
    4.9 Event listeners, form reads and clipboard monitoring
    4.10 Client hints, ETags, HSTS, Alt-Svc and pre-emptive connections
    4.11 International transfers
    4.12 Websites that reveal special categories of data
    4A. Disclosures: privacy and cookie notices
    4B. Built-in AI interfaces
    4C. Dark patterns in the consent banners
5. Evidence
    Exhibits in this report
6. Scale: the impact on people in the EU
    6.1 Energy
7. Next steps
    7.1 Disclosure
    7.2 Recommendations

Appendix A. Event types recorded

Appendix B. Third-party cookie domains

Appendix C. Receiving networks

Appendix D. Sideloaded scripts by serving domain

Appendix E. Results by website

Appendix F. Notices not measured for length and readability

1. Summary of findings

This report sets out the findings of a forensic privacy audit of the 100 most visited English-language content websites - news, reference, academic publishing, social and video platforms and more - selected from the Tranco ranking for September 2026.

The audit was conducted on the home page, privacy notice, (and where present) the cookie notice - 265 pages in total; and the audit was conducted from a domestic ISP connection in Malta without any interaction from the user - the page was simply loaded and left for 5 minutes. Every event in this report therefore took place without any consent.

The report follows the same structure as the individual report each website receives, with the figures for all 100 websites combined.

518,770
events that impact privacy, without any consent

All 100
websites with tier 1 events

2,695
events on the median website

92.4%
of servers on networks registered outside the EEA

Measure	The 100 websites
Events that impact privacy	**518,770**
Tier 1 - interception, exfiltration, high-entropy fingerprinting	14,675
Tier 2 - storage on the device, device characteristics and script loads	284,463
Tier 3 - protocol mechanisms and passive reads used for tracking	219,632
Events rated high or critical	28,695
Websites with tier 1 events	100
Cookies set without any consent	1,492 on 96 websites, of which 372 on third-party domains
Servers contacted / of which third-party	1,400 / 1,034
Servers on networks registered outside the EEA	1,293

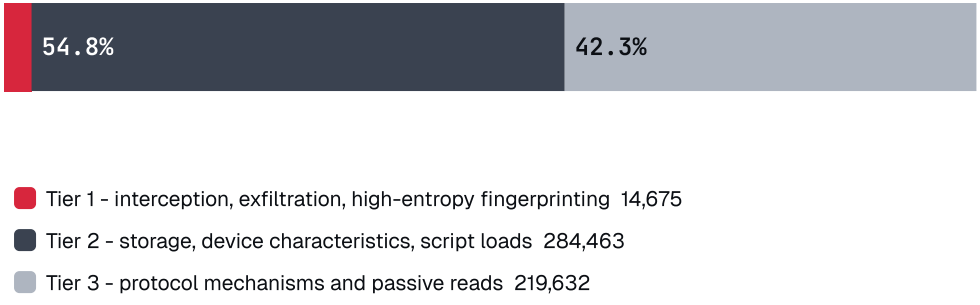


Figure 1. Events that impact privacy without any consent, by tier.

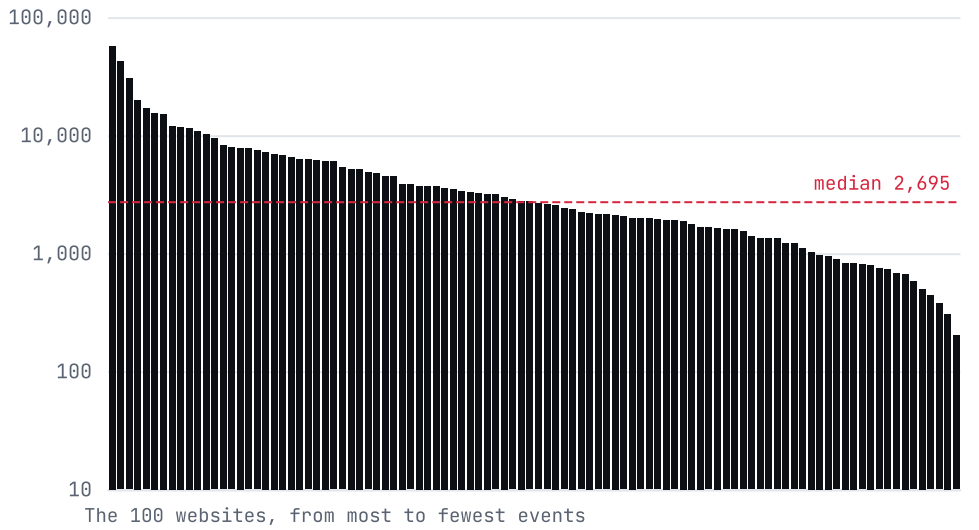


Figure 2. Events that impact privacy without any consent, per website, from most to fewest. The three highest are espn.com (57,194), go.com (42,253) and nytimes.com (30,158).

Key points

- All 100 websites stored information on, or read information from, the visitor's device without any consent - and every one of them produced tier 1 events, the most serious kind: interception, exfiltration or high-entropy fingerprinting.
- There were **518,770 events that impact privacy** without any consent across the 265 pages, each a separate access to a visitor's device and a separate breach - the median website produced 2,695, and 28,695 events were rated high or critical.
- Every website fingerprinted its visitors - **89,437 fingerprinting events** across all 100, with high-entropy techniques such as canvas, WebGL, WebGPU, WebRTC, audio, fonts, battery and sensors on 75 of them.
- The privacy and cookie notices, the very pages written to tell visitors how their data is used, produced **251,021 of the events** - 143,345 on privacy notices and 107,676 on cookie notices.
- **1,492 distinct cookies** were set without any consent on 96 websites, 372 of them on third-party domains across 60 websites, and the longest of them is set to expire in 2094.
- 98 websites loaded scripts which went on to load **7,456 further scripts** from other origins - code the websites themselves never named.
- 18 websites had subdomains of their own pointing, through DNS aliases, at **another company's servers**, so the browser treats that company's collection endpoint as part of the website.
- **92.4% of the 1,400 servers** contacted by the pages are on networks registered outside the EEA, almost all of them in the United States.

2. Scope

The 100 websites and the 265 pages audited on them are listed, with each website's results, in Appendix E. Each page was audited at the address reached after any redirects. Where a website has no separate cookie notice, or no notice could be located, fewer than three pages were captured.

Page	Pages audited	Events that impact privacy
Homepage	100	267,749
Privacy notice	97	143,345
Cookie notice	68	107,676

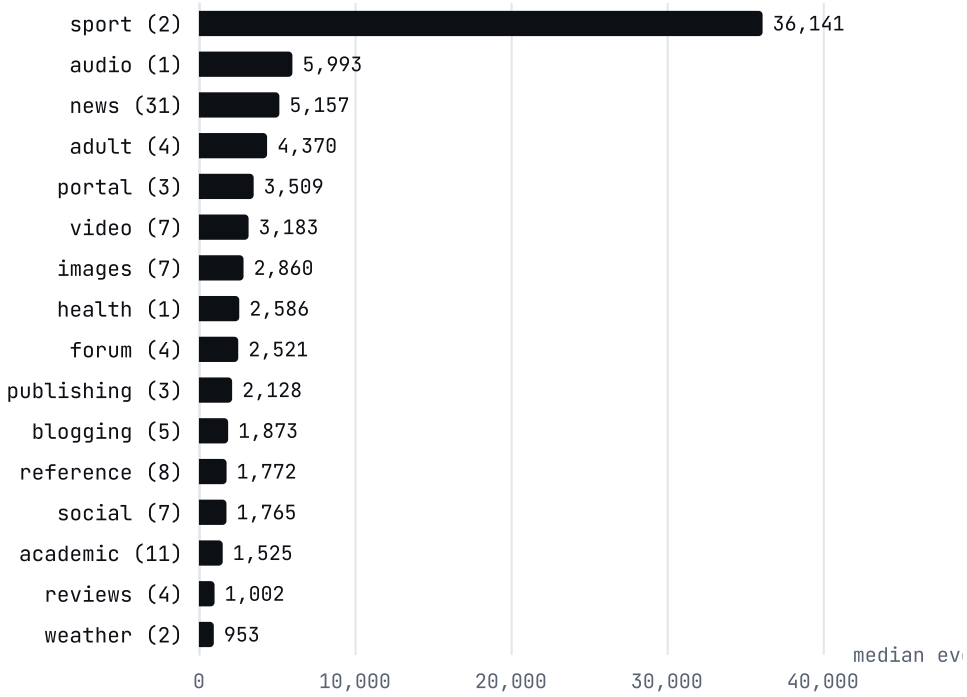


Figure 3. Median events that impact privacy per website, by type of website. The number of websites of each type is in brackets.

3. Methodology

3.1 Site selection

The websites were selected from the Tranco list, a research ranking of popular domains designed to be hard to manipulate and reproducible: list ID Q2K34, which combines the rankings published by the Chrome User Experience Report, Farsight, Majestic, Cloudflare Radar and Cisco Umbrella for 1st to 30th September 2026. A domain was included when English is its primary language, its purpose is content (news and media, reference, academic publishing, weather, sport, health information, reviews, blogging and publishing platforms, and user-generated content and social platforms), and its content can be read without logging in. Search engines, e-commerce, software services, infrastructure domains, subscription streaming, messaging, government, university and corporate websites were excluded. The first 100 domains meeting those criteria form the study.

On each website, three pages were audited where they exist: the homepage, the privacy notice and the cookie notice. Notices were located on 1st October 2026 by following the links a visitor would use - the homepage, and its footer in particular, then the privacy notice. Links inside consent banners were ignored.

3.2 The capture instrument

The capture was made with WebSentinel, a forensic website audit tool for data protection, privacy and security compliance work. It runs as an extension inside an ordinary, visible Google Chrome browser and records what a website does in that browser, at the point where the website's code runs. It only observes and records, and makes no change to what a website is able to do.

WebSentinel records evidence from several independent layers of the browser at once:

- **Instrumented browser interfaces.** Every call a script makes to the interfaces websites use to interact with the device - reading screen, time zone or graphics characteristics, storing data, listening to the mouse and keyboard, sending data to servers - is recorded as an event.
- **The network layer.** Every network request the browser makes, from the browser's own request log and from the Chrome DevTools Protocol, which adds the code that initiated each request, together with a full HTTP archive of requests and responses.
- **Storage state.** The contents of cookies, local storage and session storage as they change.
- **Scripts.** Every script loaded by the page, with its address, a hash of its content, its size and where it was inserted; scripts loaded by other scripts from a different origin are recorded with the script that loaded them.
- **Page state.** A copy of the rendered page, full-page screenshots, the text of the legal notices the page links to, and an inventory of software libraries.

Every fingerprinting, storage and tracking event carries the full call stack at the moment it happened - the script, the function, and the line and column in the code - so every behaviour can be traced to the code, and the company, that caused it, and section 5 reproduces a selection of these call stacks as exhibits.

Each event is also given a severity (info, low, medium, high or critical) by the engine's rules for the technique, which describes how intrusive a technique is and plays no part in deciding whether an event breaches the law (section 4.0).

3.3 Capture protocol

- **Location.** The capture was made from Malta, in the European Union, through a Maltese broadband connection (GO p.l.c.), whose exit address was checked before the capture. Each website saw an ordinary visitor from the EU.
- **Browser.** A visible Google Chrome browser (version 154; the browser's own user-agent string in every record), on a screen of 1512 by 982 pixels with a page viewport of 1290 by 857 pixels, with background Google services switched off - background networking, component updates, synchronisation and the download of on-device AI models - so that the only traffic captured was caused by the website.
- **No interaction, no consent.** The browser opened each page and did nothing else - nothing was clicked, nothing was typed and no banner was answered - so every event took place without any consent, which is exactly where every person stands when they first arrive on a website. The signed records contain one click across all 265 pages, and it was not the browser's: on slideshare.net's homepage, the website's own script clicked a toggle inside its Osano consent dialog, as the call stack recorded with the event shows.
- **A clean profile for every website.** Before each website was captured, the browser profile was reset to a fixed, verified baseline with no cookies, no website storage, no history and no cache, and the capture software confirmed that no cookies were present before the first page was opened. Within a website, storage was kept from page to page in the order a visitor would see them: homepage, privacy notice, cookie notice. A snapshot of the whole profile was committed to version control after every page.
- **One page at a time.** Only one page was open at any time, and the capture software confirmed that before each page.
- **Dwell time.** Each page was kept open for five minutes, capturing activity that starts after loading - timers, delayed scripts, periodic beacons and lazily loaded advertising.
- **Acceptance.** The capture of a website was accepted only when, after the browser had closed and all records had arrived, there was exactly one recorded session for each page, every session corresponded to one of the website's pages, and every session's record chain verified as complete and unbroken (section 3.4).

3.4 Evidence integrity

The evidence was captured so that anyone given it can verify that it is complete and has not been altered since capture.

1. **Hash-chained records.** Evidence is sent in a sequence of records while each page is open; each record carries a SHA-256 hash of its content and the hash of the record before it. Removing, reordering or altering any record breaks the chain detectably.
2. **Hardware-backed signatures.** Each record is signed on receipt with ECDSA on the P-256 curve with SHA-256, using a key held by a Secure Enclave. Each signature can be checked against the published public key.
3. **Trusted timestamps.** Each record's hash is timestamped by an independent RFC 3161 time-stamping authority, FreeTSA (freetsa.org).
4. **Bitcoin anchoring.** Each record's hash is submitted to two independent OpenTimestamps calendars (a.pool.opentimestamps.org and alice.btc.calendar.opentimestamps.org), which commit it to the Bitcoin blockchain.
5. **Chain verification.** Every session's chain was verified after capture; only complete, unbroken chains with no flags were accepted.

All 265 accepted sessions, holding 3,088 records, verified as complete and unbroken with no flags. Capture attempts that did not meet the acceptance conditions - on 13 websites - were discarded as a whole, never repaired or edited, and the website was captured again from the clean baseline. The discarded sessions are retained separately and are not part of this evidence.



Figure 4. How every record is sealed. Changing one byte of a record breaks its hash, every later link in the chain, its signature, its timestamp and its Bitcoin proof.

3.5 Limitations

- **One location.** The capture was made from Malta and shows what a visitor in the EU receives. Claims under the California Invasion of Privacy Act rest on that evidence: a website may be configured differently for Californian visitors, but such configuration normally reaches only the consent platform and the tag manager, and every audited website produced events rated high or critical without any consent.
- **One visit, no interaction.** Each page was opened once, by a first-time visitor who did nothing. Interaction, returning visits and consent decisions may produce different activity; the figures show what happens without any consent.
- **Desktop Chrome, five minutes per page.** Other browsers and mobile devices may receive different code; activity starting after five minutes is not captured.
- **Built-in AI interfaces.** The version of WebSentinel used did not record calls to the browser's built-in AI interfaces. The full source of every script delivered to the pages was searched for code using them instead (section 4B); that shows the code which was delivered to the device, and does not show whether it ran.
- **Recorded behaviour.** WebSentinel records what code did and makes no finding about why it did it. Where this report describes the purpose of a technique, it describes what the technique is capable of and how it is generally used.
- **Cookies sent to aliased subdomains.** The HTTP archive records request headers before the browser adds its cookies, so it does not show which cookies were sent with each request.
- **Notice text.** Where a notice page loads its text after the page or holds it in collapsed sections, the captured page may not contain the full text; this is stated where it affects a finding.

4. Findings

4.0 How to read the findings

Every event in this report happened **without any consent** - the browser opened each page and did nothing else (section 3.3).

This report is organised by the kind of thing the websites did to the visitor's device. Each section explains what the technique is and why it matters, sets out what we found, and sets out the provisions of EU and California law that apply to it.

HOW WE COUNT

Every event that impacts privacy is counted, because each time a script reads a characteristic of the device, stores an identifier on it or sends data off it, that is a separate access to the visitor's terminal equipment and a separate act of processing - and a separate breach of the law, so if a script reads the screen width fifty times, we count fifty events.

By page: 267,749 on homepages, 143,345 on privacy notices, 107,676 on cookie notices.

	Events
The 100 websites	**518,770**
Tier 1 - interception, exfiltration and high-entropy fingerprinting	14,675
Tier 2 - storage on the device, reading device characteristics and loading scripts	284,463
Tier 3 - protocol mechanisms and passive reads that can be used for tracking	219,632

Short form	Provision
ePD 5(1)	Directive 2002/58/EC, Article 5(1): confidentiality of communications - no listening, tapping, storage or other interception or surveillance without consent
ePD 5(3)	Directive 2002/58/EC, Article 5(3): storing information, or gaining access to information already stored, in the terminal equipment of a user, only with consent unless strictly necessary for a service the user explicitly requested
GDPR 5	Regulation (EU) 2016/679, Article 5(1): lawfulness, fairness and transparency; purpose limitation; data minimisation
GDPR 6	Article 6(1): lawful basis
GDPR 7	Article 7: conditions for consent - consent must be given before processing and be freely given, specific, informed and unambiguous
GDPR 9	Article 9: special categories of personal data
GDPR 12-14	Articles 12 to 14: transparency and information duties
GDPR 22	Article 22: decisions based solely on automated processing producing legal or similarly significant effects
GDPR 25	Article 25: data protection by design and by default
GDPR 26	Article 26: joint controllers
GDPR 32	Article 32: security of processing
GDPR Chapter V	Articles 44 to 50: transfers to third countries
Dir 2013/40	Directive 2013/40/EU on attacks against information systems: Article 3 illegal access, Article 4 illegal system interference, Article 5 illegal data interference, Article 6 illegal interception, as transposed in each Member State (for example Article 337C of the Maltese Criminal Code)
Budapest	Council of Europe Convention on Cybercrime: Article 2 illegal access, Article 3 illegal interception, Article 4 data interference, Article 5 system

TIERS

The tiers express how serious a technique is, and every tier is counted. Tier 3 covers mechanisms such as ETags and HSTS, which have legitimate functions and are also used to track people across visits; the reasons they are still counted are set out in section 4.10.

WHAT IS NOT COUNTED

Three kinds of event are excluded from every count:

- the detection of consent management platforms and banners, which we use only to show that a consent mechanism was present while everything else happened;
- two event types the capture engine derives by matching addresses against lists of known tracking domains, because every figure in this report is based on observed behaviour alone;
- WebSentinel's own instrumentation: everything the capture extension itself did or loaded (882 events, almost all the browser loading the extension's own files, are removed before any figure is calculated).

Every script loaded into a page counts, as tier 2: loading code into the visitor's browser stores it in the terminal equipment and runs it there (section 4.3).

WebSentinel's audit engine identifies Consent Management Platforms (CMPs) from a list of known vendors, recognised by their scripts.

Where a website runs its own in-house consent mechanism, or uses a platform its scripts do not identify, the engine identifies it from the screenshots it takes during the capture: 44 pages on 25 websites showed their own in-house consent banner, and 19 pages on 11 websites a platform identified this way (CookieChimp, OneTrust, Transcend and Usercentrics). A consent banner was in the page on 165 of the 265 pages, and none was ever answered.

THE LEGAL PROVISIONS

Article 5(3) of the ePrivacy Directive applies whether or not the information concerned is personal data (CJEU, Planet49, C-673/17, ECLI:EU:C:2019:801). Where it is, the GDPR applies as well.

HOW THE EPRIVACY DIRECTIVE AND THE GDPR FIT TOGETHER

The two apply side by side. Article 5(3) of the ePrivacy Directive is a special rule for the act of storing information on, or gaining access to information in, the visitor's device, and it takes precedence over the GDPR for that act; it contains no special rule for the processing that comes before or after it, which therefore needs a legal basis under Article 6 of the GDPR as well (EDPB Opinion 5/2019 on the interplay between the ePrivacy Directive and the GDPR, 12th March 2019, paragraphs 38, 40 and 75).

Consent under Article 5(3) has the same meaning as consent under the GDPR (Article 2(f) of the ePrivacy Directive; Planet49, paragraphs 62 and 63). It requires active behaviour: a pre-ticked box, silence, inactivity or simply carrying on using a service is not consent (Planet49, paragraphs 52 to 54;

EVERY WEBSITE, EVERY CATEGORY

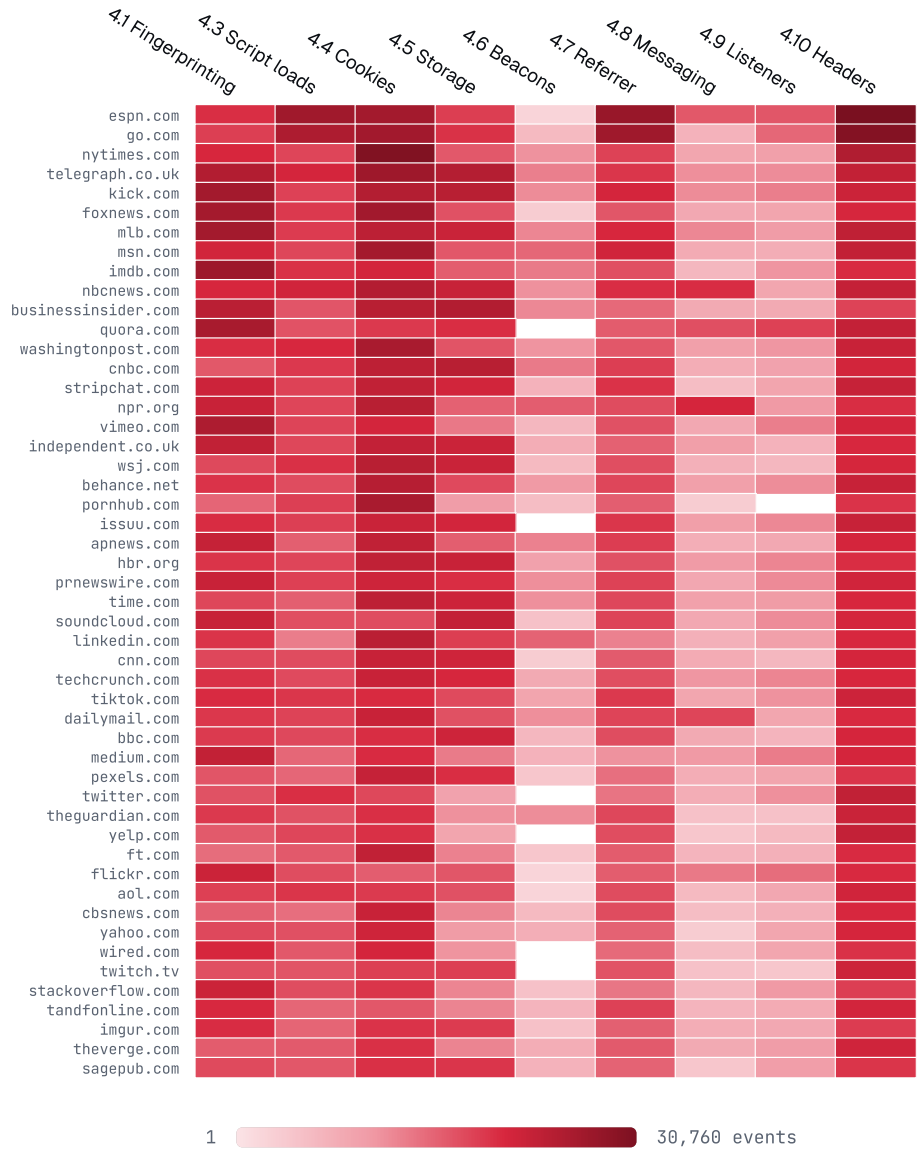


Figure 5. Events without any consent by website and category: websites 1 to 50 by number of events. White means none.

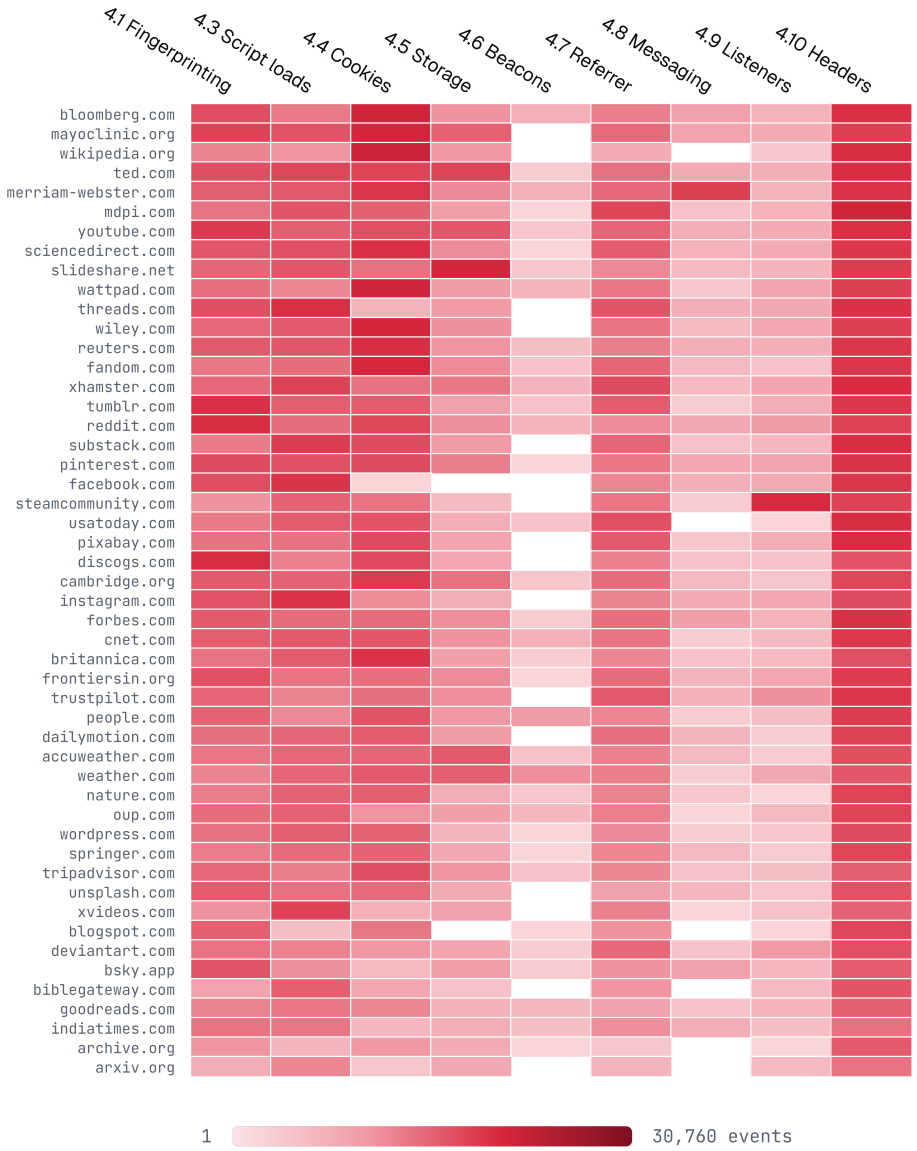


Figure 6. Events without any consent by website and category: websites 51 to 100 by number of events. White means none.

EDPB Guidelines 05/2020 on consent, version 1.1, 4th May 2020, paragraph 79), and scrolling or swiping through a page will not under any circumstances satisfy the requirement of a clear and affirmative action (paragraph 86). Consent must always be obtained before the processing it covers starts (paragraph 90).

No consent was given for any event in this report. GDPR Articles 5(1), 6(1), 7 and 25 are therefore engaged by every event counted: there was no lawful basis, no consent meeting the conditions of Article 7, and a system that does all of this without any consent is the opposite of data protection by design and by default - by default, only processing strictly necessary for a lawful purpose should take place (EDPB Guidelines 4/2019 on Article 25, version 2.0, 20th October 2020, paragraph 42). Articles 28, 32 and Chapter V depend on the facts of each event.

Where this report says a provision applies, it means the observed conduct falls within the provision's scope - whether any exemption could apply, such as the strict necessity exemption in Article 5(3), is addressed in each section, and no consent was given at any point during the capture.

The European Data Protection Board, and the Article 29 Working Party before it, have set out how Article 5(3) is to be read, and this report relies on that guidance throughout:

- Article 29 Working Party, Opinion 04/2012 on Cookie Consent Exemption (WP194, 7th June 2012) - the two exemptions from consent are narrow, the service has to have been explicitly requested by the user, and third-party advertising and analytics fall outside them;
- Article 29 Working Party, Working Document 02/2013 providing guidance on obtaining consent for cookies (WP208, 2nd October 2013) - consent has to be obtained before anything is stored on or read from the device;
- Article 29 Working Party, Opinion 9/2014 on the application of Directive 2002/58/EC to device fingerprinting (WP224, 25th November 2014) - fingerprinting is covered by Article 5(3) in the same way as cookies;

- European Data Protection Board, Guidelines 2/2023 on the Technical Scope of Article 5(3) of the ePrivacy Directive (version 2.0, 7th October 2024) - Article 5(3) covers every technique that stores information on, or gains access to information in, the device, including scripts instructing the browser to send information, tracking pixels, tracking links and processing carried out locally on the device;
- European Data Protection Board, Guidelines 05/2020 on consent under Regulation 2016/679 (version 1.1, 4th May 2020) - the conditions consent has to meet, which apply to consent under Article 5(3) as well.

4.00 Who is responsible: joint controllership

Much of the activity in this report was carried out by code the website operators did not write: advertising libraries, analytics tags, consent tools, video players, fonts, bot detection. A common response from website operators is that this is someone else's code and therefore someone else's responsibility - a position the Court of Justice of the European Union has rejected.

THE CASE LAW

In **Fashion ID** (C-40/17, judgment of 29th July 2019, ECLI:EU:C:2019:629), a retailer embedded Facebook's "Like" button on its website. The button caused visitors' browsers to send personal data to Facebook as soon as the page loaded, whether or not the visitor was a Facebook user and whether or not they clicked the button. The Court held that the website operator is a **joint controller** with the third party for the collection of the data and its transmission to the third party, because by embedding the code it determines, together with the third party, the purposes and means of that processing. It does not need to have access to the data to be a controller.

Fashion ID built on two earlier Grand Chamber judgments: **Wirtschaftsakademie** (C-210/16, 5th June 2018, ECLI:EU:C:2018:388), in which the administrator of a Facebook fan page was held to be a joint controller with Facebook for the processing of visitors' data, and **Jehovan todistajat** (C-25/17, 10th July 2018, ECLI:EU:C:2018:551), which confirmed that joint control does not require each party to have access to the data or an equal share of responsibility.

The EDPB applies the same reading: the choice to use for its own purposes a tool developed by another entity, which allows the processing of personal data, "will likely amount to a joint decision on the means of that processing", as in Fashion ID, where by embedding the Like button the website operator exerted a decisive influence over the collection and transmission of its visitors' data (Guidelines 07/2020 on the concepts of controller and processor in the GDPR, version 2.1, 7th July 2021, paragraph 67).

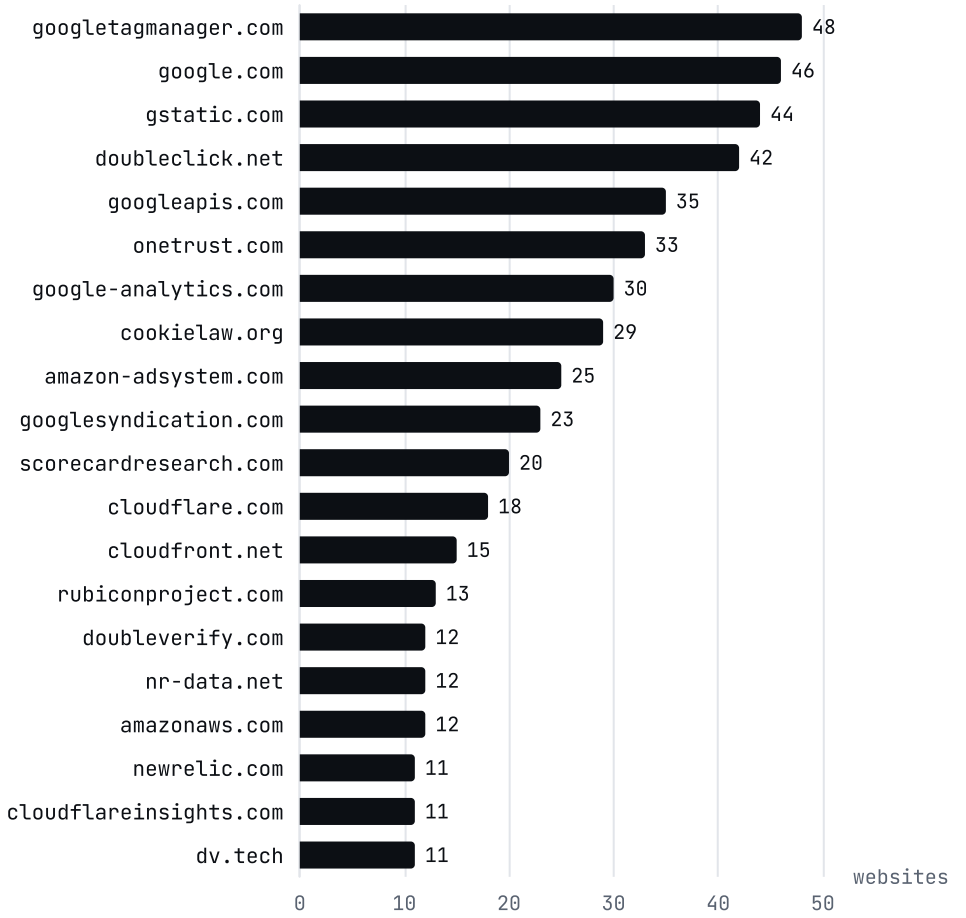


Figure 7. Third-party domains whose code or servers caused events without any consent, by the number of the 100 websites on which they did so.

WHAT IT MEANS FOR THE 100 WEBSITES

Every third-party script on these websites was there because the website operator chose to put it, or the code that loads it, on the page. For every collection of data by that script, and every transmission of data from the visitor's device to its operator, **the website operator is a joint controller** under Article 26 GDPR, alongside the company that supplied the script.

That has three consequences for the findings that follow:

1. **The website operator answers for every event on its pages**, including those carried out by third-party code, at least for the collection and transmission of the data.
2. **The obligations of Articles 5, 6, 7, 12 to 14, 25 and 32 GDPR rest on the website operator** for that processing - including the obligation to obtain valid consent before it takes place, and to tell visitors about it.
3. **Sideloading does not move responsibility elsewhere** - where a script the operator loaded goes on to load further scripts from other companies (section 4.3), those scripts are on the page because of the operator's choice.

Article 5(3) of the ePrivacy Directive points the same way: it prohibits storing information on, or gaining access to information in, the user's device without consent, whoever's code does it, and it is the website that delivers that code to the device.

THE THIRD PARTIES ON THE 100 WEBSITES

Code served from googletagmanager.com acted on 48 of the 100 websites without any consent, and under Fashion ID every one of those website operators is a joint controller with the company that supplied the code for the collection and transmission of the data it caused.

4.1 Device fingerprinting

WHAT IT IS

Fingerprinting identifies a device by reading its characteristics, without needing to store an identifier on it. Each characteristic - the screen size, the time zone, the number of processor cores, the list of graphics extensions, the way the device draws a hidden image - says little about the device by itself, but read together they are distinctive enough to recognise the same device again, across websites and across visits, without a cookie and without the person being able to see it, block it or delete it.

Measure	Figure
Events	**89,437**
Websites	100 of 100
By tier	tier 1: 5,606, tier 2: 43,290, tier 3: 40,541
Where	49,546 on homepages, 23,960 on privacy notices, 15,931 on cookie notices

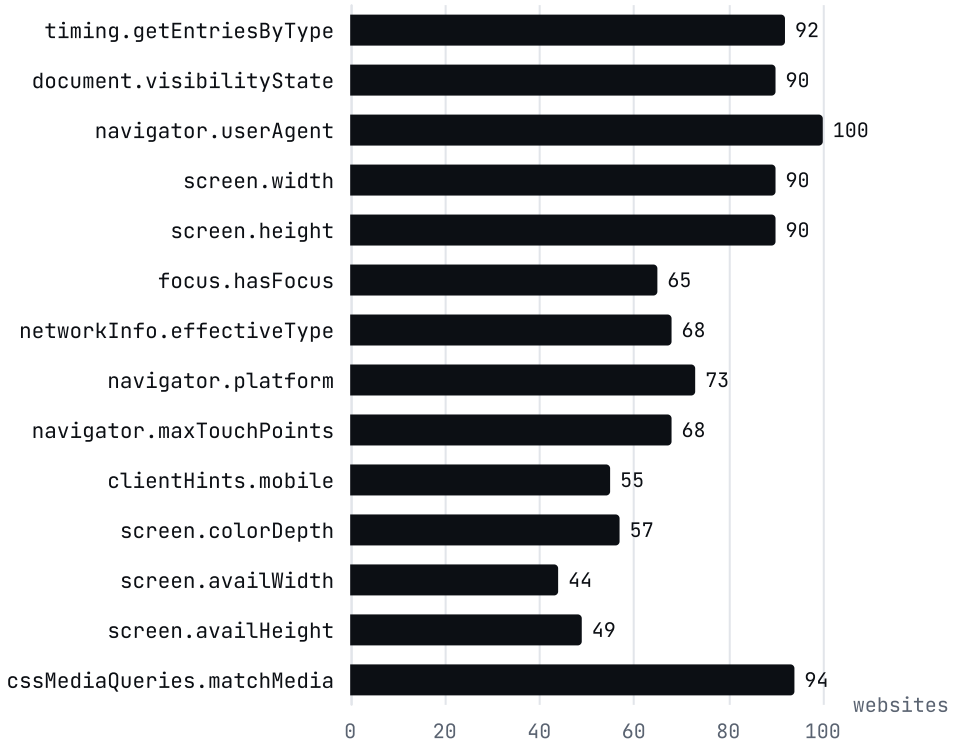


Figure 8. Fingerprinting: the techniques used, by the number of websites on which each was recorded.

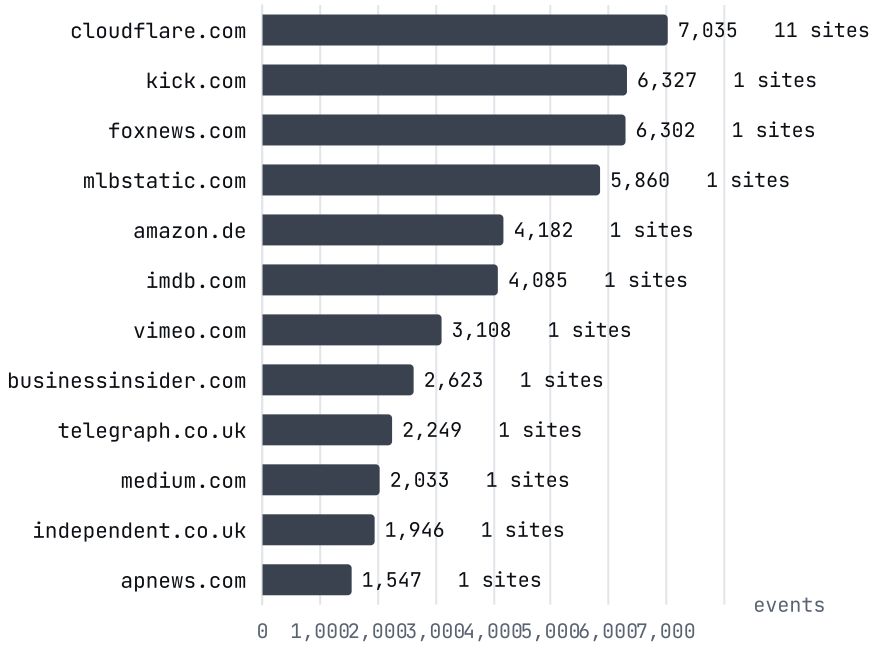


Figure 9. The scripts and servers responsible, by registrable domain: events, and the number of websites on which they caused them. "server" figures are events caused by a server response with no script involved.

Fingerprinting works by script: code on the page asks the browser for information about the device and sends the results to a server, and every one of those requests gives the script access to information stored on the visitor's device.

We group the techniques by how much they reveal:

- **High-entropy techniques (Tier 1)** read characteristics that vary widely between devices and are used specifically to tell devices apart: drawing to a hidden canvas and reading the pixels back, querying the graphics hardware through WebGL and WebGPU, generating audio signals, measuring installed fonts, setting up WebRTC connections that expose network addresses, reading the battery state, the keyboard layout, the installed speech voices, connected media devices and motion sensors, and probing for browser extensions.

- **Device characteristics (Tier 2)** are the standard components of a fingerprint: user agent, platform, language, screen dimensions and colour depth, time zone, processor cores, device memory, network connection type and the browser's client hints.
- **Environmental reads (Tier 3)** report the state of the page and the device - whether the page is visible or in focus, layout and timing measurements. They carry little identifying information by themselves, add to a fingerprint when combined with the techniques above, and are also used to measure attention.

WHAT WE FOUND ACROSS THE 100 WEBSITES

High-entropy fingerprinting (tier 1) was recorded on 75 of the 100 websites.

EXHIBIT E1 linkedin.com - homepage **CRITICAL**

`canvas.toDataURL` at 2026-10-01T13:22:39.807Z

https://www.linkedin.com/

api	HTMLCanvasElement.prototype.toDataURL
canvasSize	2000x200

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at t.getCanvasFp (https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:28929)
at Object.canvasKey [as getData] (https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:20758)
at https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:40398
at Array.forEach (<anonymous>)
at t.getFeatures (https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:40196)
```

Record `sess_a5b04afc-3360-46cd-9008-c695febcb3bd/rec_000001_764355cb57d7.json` - chain position 1

SHA-256 of the record file: `c3a86091a9c4d034c7faedf967cc685affa4872eac724dc57fc76c68d3a521e3`

EXHIBIT E2 linkedin.com - homepage **CRITICAL**

`webrtc.createOffer` at 2026-10-01T13:22:39.797Z

https://www.linkedin.com/

api	RTCPeerConnection.prototype.createOffer

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at t.getIPs (https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:14948)
at Object.webRtcKey [as getData] (https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:15540)
at https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:40398
at Array.forEach (<anonymous>)
at t.getFeatures (https://static.licdn.com/aero-v1/sc/h/azyj0u7xl6tpxe5fn3t3djk7p:2:40196)
```

Record `sess_a5b04afc-3360-46cd-9008-c695febcb3bd/rec_000001_764355cb57d7.json` - chain position 1
SHA-256 of the record file: `c3a86091a9c4d034c7faedf967cc685affa4872eac724dc57fc76c68d3a521e3`

Exhibits E1 and E2 show LinkedIn's homepage drawing a hidden 2000 by 200 pixel canvas and setting up a WebRTC connection inside functions its own code names `getCanvasFp` and `getIPs`; exhibit E3 shows Google's reCAPTCHA reading the battery state of a visitor to forbes.com.

THE LAW

Reading a characteristic of the visitor's device is gaining access to information stored in their terminal equipment, so Article 5(3) of the ePrivacy Directive applies to every fingerprinting event in this section - whether or not the information is personal data (Planet49). The Article 29 Working Party confirmed in Opinion 9/2014 that fingerprinting is covered by Article 5(3) in exactly the same way as cookies, and the EDPB's Guidelines 2/2023 confirm that a script instructing the browser to hand over information about the device is gaining access to it. Consent is required unless the access is strictly necessary to provide a service the visitor explicitly asked for, and nobody who opens a news article has asked for their graphics hardware, battery, installed voices or keyboard layout to be read.

The largest single source of fingerprinting was a bot challenge, so bot detection has to be dealt with directly. That bot challenge still needed consent - Opinion 04/2012 accepts security as exempt only where it is part of a service the user explicitly requested, such as logging in, and only to the extent that service needs it - and covertly collecting dozens of device characteristics from every visitor, on every page including the privacy notice, goes far beyond anything needed to show someone a page.

The fingerprints these websites collected single out a device, and through it a person, so they are personal data under the GDPR, and collecting them without any consent left that processing without a lawful basis under Article 6 and unable to meet the conditions of Article 7. It collected far more than any purpose needs (Article 5(1)(c)), it is incompatible with data protection by design and by default (Article 25), and because visitors could neither see it nor clear it, it breached the principle of fairness and transparency in Article 5(1)(a).

The scripts on these websites interrogated the hardware, software and configuration of visitors' devices without their knowledge or consent, accessing those systems without right, which is illegal access under Article 3 of Directive 2013/40/EU and Article 2 of the Budapest Convention, as transposed in every Member State - and in California, knowingly accessing a computer or its data without permission is an offence under Penal Code section 502.

Section 4B sets out how the availability of the browser's built-in AI interfaces was itself used as one input to a fingerprint.

EXHIBIT E3 forbes.com - homepage **HIGH**

`battery.getBattery` at 2026-10-01T17:41:35.320Z

https://www.forbes.com/

api	Navigator.prototype.getBattery

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at ud.<anonymous> (https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjBELZirRCZG7oY/recaptcha__en.js:468:446)
at ud.j.H.<computed>.P.zA (https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjBELZirRCZG7oY/recaptcha__en.js:1666:282)
at ud.j.H.<computed>.ud.y6 (https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjBELZirRCZG7oY/recaptcha__en.js:1663:121)
at https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjBELZirRCZG7oY/recaptcha__en.js:510:405
at MR (https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjBELZirRCZG7oY/recaptcha__en.js:1041:372)
```

Record `sess_b4e763c2-53ee-4c12-9fa8-c8036d75b449/rec_000002_be38c3d2cea3.json` - chain position 2

SHA-256 of the record file: `a165968a038f1d32e26ec3900dda47e2edc970f8aba31d7e8375ec0fb67db69c`

4.2 CNAME cloaking

WHAT IT IS

Browsers treat a website's own subdomains as part of the website. A request to `metrics.example.com` from a page on `example.com` is a first-party request: the browser sends it the cookies set for `example.com`, lets it set cookies of its own, and does not apply the protections it applies to third parties.

CNAME cloaking exploits this: the website points one of its own subdomains, through a DNS alias (a CNAME record), at a server run by another company - typically an analytics or advertising provider. To the browser, and to the visitor, the third party now looks like the website itself, and tracking that browsers and privacy tools would otherwise block or limit as third-party goes through under the website's own name.

HOW WE MEASURED IT

The capture engine has a CNAME detection of its own, which works by matching addresses against a list of known tracking domains, and as every figure in this report is based on observed behaviour we did not use it - instead we resolved the DNS records of every subdomain of each website that the audited pages contacted, and followed each chain of aliases to its end.



Figure 10. CNAME cloaking as recorded on nbcnews.com: the subdomain is a DNS alias of a third party's collection server, so the browser treats the request as the website's own and sends the website's cookies with it.

Subdomain	Alias resolves to	Requests without any consent
a1.api.bbc.com	a1-api-bbc-com-cddc.at-o.net	98
cmpv2.independent.co.uk	cdn-1746.privacy-mgmt.com	33
consent-manager.ft.com	cdn-1906.privacy-mgmt.com	28
tagdelivery.mayoclinic.org	tagdelivery.mayoclinic.org.pc-mc.greylabeldelivery.com	28
smetrics.telegraph.co.uk	telegraph.co.uk.ssl.sc.omtrdc.net	23
dwt.soundcloud.com	dwt.soundcloud.com.first-party-js.datadome.co	12
ans.mlb.com	mlb.com.ssl.sc.omtrdc.net	10
consent.dailymotion.com	cdn-1945.privacy-mgmt.com	10
dd.nytimes.com	dd.nytimes.com.first-party-js.datadome.co	9
smetrics.foxnews.com	foxnews.com.ssl.d1.sc.omtrdc.net	5
aamt.nbcnews.com	nbcnews.com.ssl.sc.omtrdc.net	3
dcf.espn.com	edge-geo.nc0.co	3
dd.reuters.com	dd.reuters.com.first-party-js.datadome.co	3
smetrics.hbr.org	hbr.org.ssl.sc.omtrdc.net	3
trk.prnewswire.com	p04f.t.eloqua.com	3
sw88.go.com	go.com.ssl.d1.sc.omtrdc.net	2
trk.cision.com	p04f.t.eloqua.com	2
auth.time.com	auth-time-com.tinypass.com	1
c2.time.com	c2-time-com.tinypass.com	1
verticalvideo.businessinsider.com	v.bytes.media	1
vx.time.com	vx-time-com.tinypass.com	1

The records were resolved on 7th October 2026, after the capture, and as DNS records can change, the results show how each subdomain was configured on that date.

WHAT WE FOUND ACROSS THE 100 WEBSITES

21 subdomains of 18 websites were DNS aliases of a service under another company's domain and were contacted without any consent:

THE LAW

The requests and cookies that went through these aliases were storage on, and access to, the visitor's device exactly as they would have been under the third party's own name, so Article 5(3) of the ePrivacy Directive applies to them in full and the alias has no effect on the need for consent - the EDPB's Guidelines 2/2023 make clear that Article 5(3) applies to any storage of, or access to, information on the device, whoever carries it out.

The 18 websites that dressed a third party's data collection up as their own breached the principle of fairness and transparency in Article 5(1)(a) of the GDPR, because a visitor, a browser or a privacy tool looking at the requests sees the website's name and has no way of knowing which company is receiving the data, and setting things up so that a third party's tracking gets past the protections the browser applies to third parties is incompatible with data protection by design under Article 25. Each of those website operators created the alias, so it is at least a joint controller for the collection and transmission under Article 26 (section 4.00), and under Articles 12 to 14 it has to name the company behind each aliased subdomain in its notice, as nothing in the request itself tells the visitor - section 4A sets out whether the notices do.

The cookies set and read through these aliases are subject to the same criminal law as any other cookie (section 4.4).

4.3 Script loads and sideloaded scripts

WHAT IT IS

Every script a page loads is code downloaded to the visitor's device and run there - the website's own code, written into the page or served from its own domains, and code supplied by third parties. It is stored in the browser, executed with access to the page and to the browser's interfaces, and can do everything described in the rest of this report.

A website chooses which scripts to put on its pages, and many of those scripts then load further scripts of their own, from other companies, which the website never named. A tag manager loads a dozen analytics and advertising tags; an advertising library loads bidders; a consent tool loads its own components. We call these **sideloaded scripts**: scripts loaded by another script from a different origin.

The visitor did not ask for any of these scripts, and in most cases neither did the website operator by name.

HOW WE MEASURED IT

The capture engine records every script loaded into a page, and separately records every script loaded by another script from a different origin, together with the script that loaded it. Every script load is counted, as tier 2, because loading code into the visitor's browser stores it in the terminal equipment and runs it there. The sideloading figures show how much of that code arrived without the website naming it. The engine's definition of sideloading is technical, so it includes some scripts a company loads from its own content domains - x.com loading code from twimg.com, for example.

WHAT WE FOUND ACROSS THE 100 WEBSITES

Measure	Figure
Script loads counted (tier 2)	**41,341**, on all 100 websites
Scripts sideloaded by other scripts from a different origin	**7,456**, on 98 websites

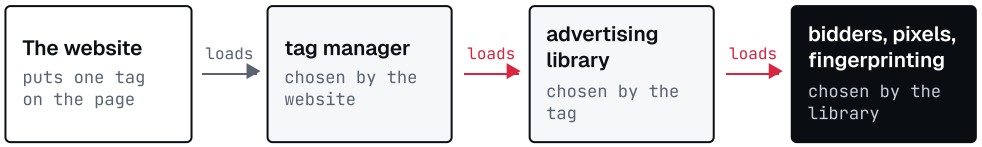


Figure 11. How sideloading works. Only the first script was placed on the page by the website; everything after it arrived because the script before it decided so.

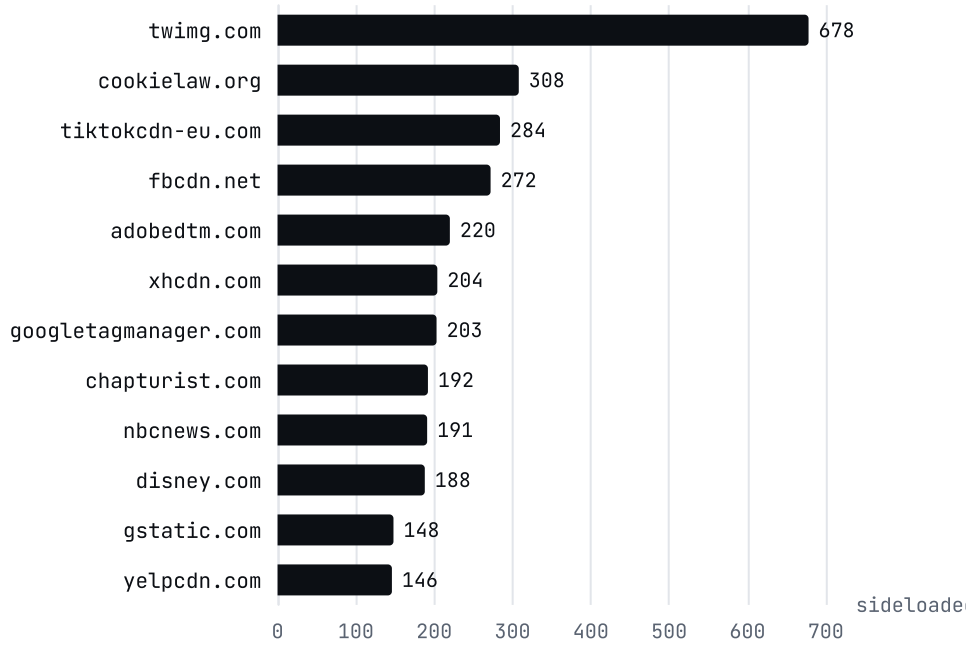


Figure 12. The domains serving the most sideloaded scripts across the 100 websites.

THE LAW

Every script these pages downloaded to a visitor's device and ran there is information stored in their terminal equipment, and everything it then read through the device's interfaces is access to information stored there - the EDPB's Guidelines 2/2023 treat code that instructs the browser to collect and send information as gaining access to the device. Article 5(3) of the ePrivacy Directive therefore applied as soon as third-party code was loaded onto a visitor's device without any consent, before that code had done anything else, and nobody needs an advertiser's bidding script on their device to read a news article.

On 98 of these websites, scripts went on to load further scripts the operator never chose, so those operators have no way of knowing what code runs on their visitors' devices, which processors and third parties receive their data, or whether that code is secure - which is incompatible with data protection by design (Article 25 GDPR), with the requirement to use processors only under contracts that bind them (Article 28) and with security of processing (Article 32). Under Fashion ID each of them is a joint controller for the collection and transmission this code carries out (section 4.00, Article 26), and none of them can tell visitors who receives their data, as Articles 12 to 14 require, without knowing which scripts their own pages will load.

Loading code onto visitors' devices and running it there without their knowledge or consent is access to an information system without right (Directive 2013/40/EU, Article 3; Budapest Convention, Article 2), and in California it is an offence under Penal Code section 502.

4.4 Cookies

WHAT IT IS

A cookie is a small piece of data a website stores in the visitor's browser and receives back with later requests.

Measure	Figure
Events	**154,952**
Websites	100 of 100
By tier	tier 1: 23, tier 2: 154,929
Where	97,906 on homepages, 33,677 on privacy notices, 23,369 on cookie notices

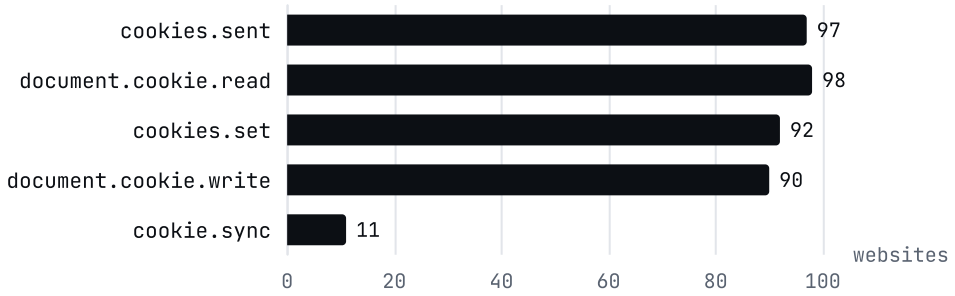


Figure 15. Cookies: the techniques used, by the number of websites on which each was recorded.

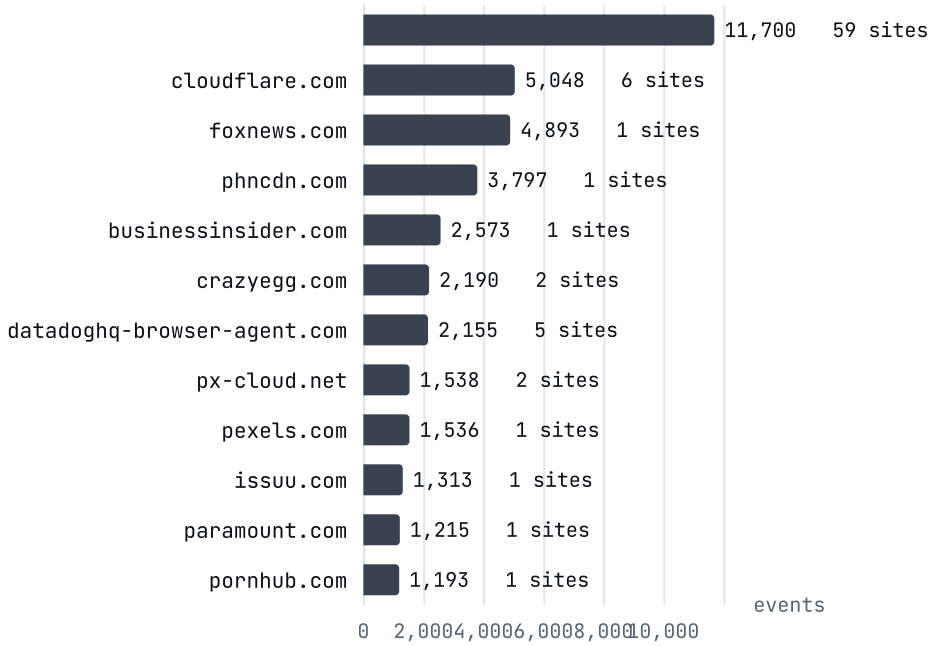


Figure 16. The scripts and servers responsible, by registrable domain: events, and the number of websites on which they caused them. "server" figures are events caused by a server response with no script involved.

Cookies can hold settings and session state; they can also hold an identifier that recognises the same browser across pages, visits and websites. Advertising and analytics cookies are surveillance technology: their purpose is to keep state that lets a company follow behaviour over time.

Cookies are set in two ways - by a server, in its response (a Set-Cookie header), and by a script on the page, through `document.cookie`.

Scripts can also read the cookies set for a site, and the browser sends a site's cookies back with every request to it.

WHAT WE FOUND ACROSS THE 100 WEBSITES

1,492 distinct cookies were set without any consent on 96 websites, 372 of them on third-party domains, on 60 websites.

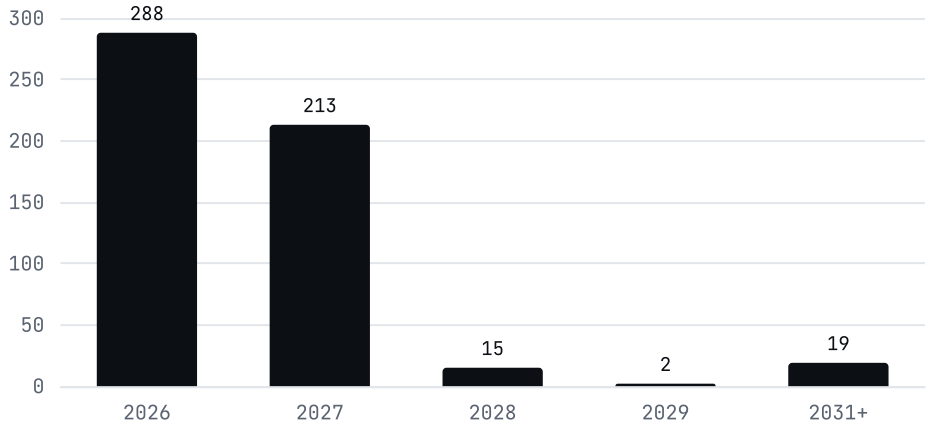


Figure 13. Cookies set without any consent, by the year in which they are set to expire.

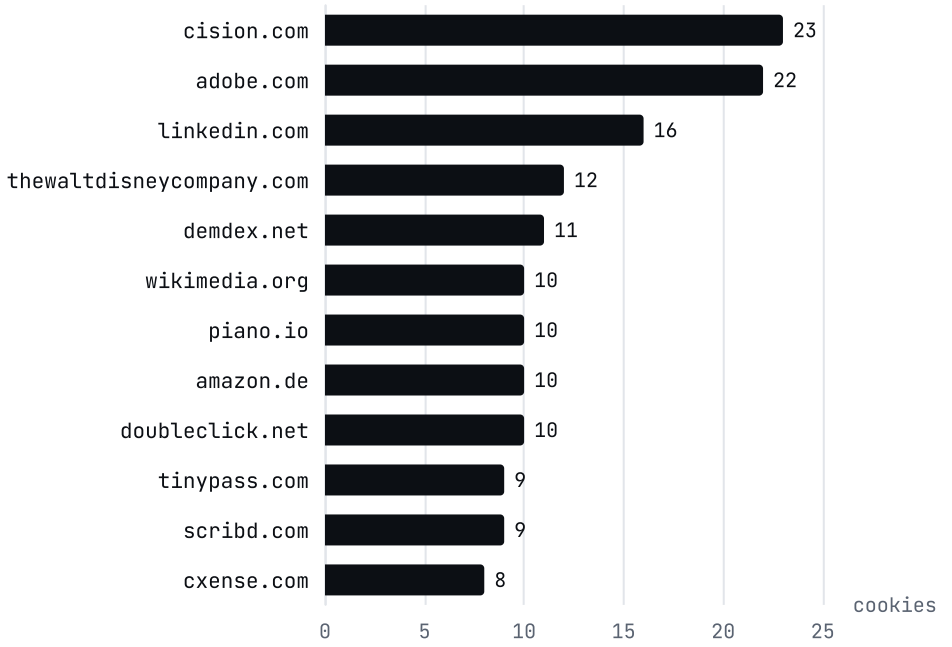


Figure 14. Third-party cookies set without any consent, by the domain they were set on.

THE LONGEST-LIVED COOKIES SET WITHOUT ANY CONSENT

Website	Cookie	Domain	Expires
pixabay.com	anonymous_user_id	pixabay.com	Wed, 20 Oct 2094 02:47:39 GMT
goodreads.com	ccsid	www.goodreads.com	Tue, 02 Oct 2046 00:02:43 -0000
slideshare.net	scribd_ubtc	scribd.com	Wed, 02 Oct 2041 02:21:04 GMT
slideshare.net	ssuuid	scribd.com	Wed, 02 Oct 2041 02:21:04 GMT
vimeo.com	vuid	vimeo.com	Sun, 28-Sep-2036 14:51:05 GMT
imdb.com	session-id	imdb.com	Tue, 01 Jan 2036 08:00:01 GMT
imdb.com	session-id-time	imdb.com	Tue, 01 Jan 2036 08:00:01 GMT
pixabay.com	lang	pixabay.com	Sun, 28 Sep 2036 23:39:18 GMT

The longest-lived was set to expire in 2094.

THE LAW

Setting a cookie is storing information in the visitor's terminal equipment, and reading it - by script or with a request - is gaining access to information already stored there, so both require consent under Article 5(3) of the ePrivacy Directive unless they are strictly necessary for a service the visitor explicitly asked for.

The Article 29 Working Party set out in Opinion 04/2012 that third-party advertising cookies and analytics cookies fall outside that exemption, and that a cookie which is exempt should have a lifespan in line with its purpose - a cookie holding an identifier for months, let alone until 2094, is nowhere near that - and in Working Document 02/2013 it made clear that cookies should only be set once the visitor has given consent. Every cookie in this section was set without any consent.

The cookies holding identifiers used to follow behaviour across visits and websites are instruments of surveillance, and Article 5(1) prohibits surveillance of a person's communications - here, the pages they request and read - without their consent, in addition to Article 5(3).

EXHIBIT E4 npr.org - privacy notice — MEDIUM

`cookies.set` at 2026-10-02T01:02:02.257Z

https://www.npr.org/about-npr/179878450/privacy-policy

api	CDP.Network.responseReceivedExtraInfo
name	cto_bundle
domain	criteo.com
path	/
sameSite	none
secure	True
httpOnly	False
expires	Wed, 27 Oct 2027 01:02:01 GMT

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
(no script on the stack: the event was caused by a server response - server:criteo.com)
```

Record `sess_314b66d1-e00d-4222-8abf-524173727fed/rec_000003_e640d0eabe09.json` - chain position 3

SHA-256 of the record file: `4bcc58780021883182f4c5df040d3ebb865e322cbb2af263370e0f6ed0fec171`

EXHIBIT E5 bbc.com - privacy notice MEDIUM

`document.cookie.write` at 2026-10-01T17:57:49.380Z

https://www.bbc.com/pages/privacy-policy

api	Document.prototype.cookie (setter)
name	https://www.bbc.com_oeu1790877147988r0.3170336110737656$$27302320011$$session_state
secure	False
httpOnly	False
sameSite	

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at https://a4621041136.cdn.optimizely.com/client_storage/a4621041136.html:7:1483
```

Record sess_2ced1792-7f88-419c-bc1f-688e491256ab/rec_000001_d94d98682f42.json - chain position 1
SHA-256 of the record file: 52cedc2c09e15771f44ff91303955943933ede1d2f9775b981ee15004c9cd675

The identifiers in those cookies are personal data, and processing them without any consent had no lawful basis under Article 6 of the GDPR and could not meet the conditions of Article 7.

Writing data to visitors' devices and reading it back without their knowledge or consent is access to an information system without right (Directive 2013/40/EU, Article 3; Budapest Convention, Article 2) and alteration of data on it (Directive 2013/40/EU, Article 5; Budapest Convention, Article 4). In California, knowingly accessing a computer or its data without permission is an offence under Penal Code section 502, and identifiers sent to third parties along with the address of the page being read are covered by section 631 (section 4.6).

4.5 Local storage, session storage and other browser storage

WHAT IT IS

Browsers give websites several places to store data on the visitor's device besides cookies: local storage, which persists until it is deleted; session storage, which lasts until the tab is closed;

WHAT WE FOUND ACROSS THE 100 WEBSITES

Measure	Figure
Events	**44,494**
Websites	98 of 100
By tier	tier 2: 44,494
Where	26,726 on homepages, 13,241 on privacy notices, 4,527 on cookie notices

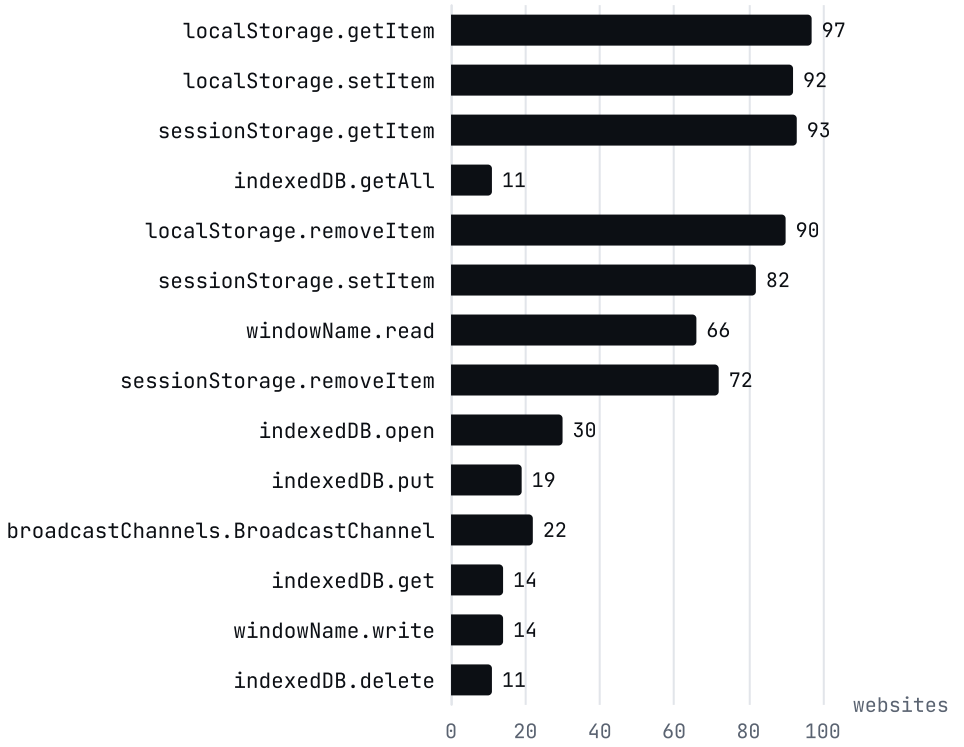


Figure 17. Storage: the techniques used, by the number of websites on which each was recorded.

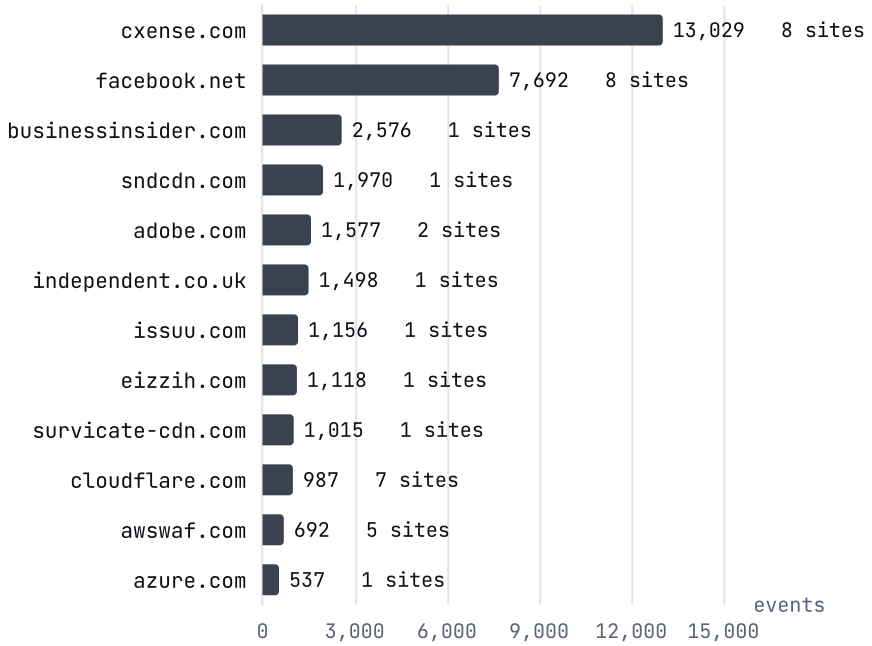


Figure 18. The scripts and servers responsible, by registrable domain: events, and the number of websites on which they caused them. "server" figures are events caused by a server response with no script involved.

IndexedDB, a database in the browser; the cache storage and service workers used by web applications; and `window.name`, a value that survives navigation from one site to another. Scripts can read these stores at any time and send what they contain to a server. They are widely used to keep identifiers, consent state and behavioural data, including by companies that want an identifier that survives the deletion of cookies.

THE LAW

Writing to any of these stores is storing information in the visitor's terminal equipment, and reading from them is gaining access to it - Article 5(3) of the ePrivacy Directive applies to every technology that does either, cookies included, as the EDPB's Guidelines 2/2023 confirm.

Where the stored data is an identifier or a behavioural record used to follow the visitor, it is a surveillance mechanism and Article 5(1) applies as well (section 4.4), and the identifiers and behavioural data held in these stores are personal data, processed here without any consent and without a lawful basis under Articles 5, 6 and 7 of the GDPR.

As with cookies, writing to and reading from the device's storage without right is illegal access and data interference (Directive 2013/40/EU, Articles 3 and 5; Budapest Convention, Articles 2 and 4) and an offence under California Penal Code section 502. Registering a service worker - code that keeps running in the browser after the page is closed - adds interference with the functioning of the system (Directive 2013/40/EU, Article 4; Budapest Convention, Article 5).

EXHIBIT E6 cnn.com - homepage **HIGH**

`localStorage.setItem` at 2026-10-01T17:09:10.108Z

https://edition.cnn.com/

api	Storage.prototype.setItem
key	_cX_check
valuePreview	_cX_check

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at Or (https://cdn.cxense.com/cx.js:6:43531)
at uc (https://cdn.cxense.com/cx.js:6:86770)
at Rc (https://cdn.cxense.com/cx.js:6:116816)
at https://cdn.cxense.com/cx.js:6:117692
at https://cdn.cxense.com/cx.js:6:117726
```

Record `sess_8afe9e2b-1a3c-4cde-a867-8981f77de8c1/rec_000001_21541b13260c.json` - chain position 1

SHA-256 of the record file: `0c7e082a71d5a84e7021b4b8732167dcc50ccb02c60331bd3171d14f2de7d7e7`

4.6 Beacons, keep-alive requests, tracking pixels and link decoration

WHAT IT IS

These are the techniques used to get data off the visitor's device.

- **Beacons** (`navigator.sendBeacon`) send data to a server in the background. They are designed to be delivered even as the page closes, so that nothing is lost when the visitor leaves.

WHAT WE FOUND ACROSS THE 100 WEBSITES

Measure	Figure
Events	**1,632**
Websites	78 of 100
By tier	tier 1: 1,632
Where	1,030 on homepages, 344 on privacy notices, 258 on cookie notices

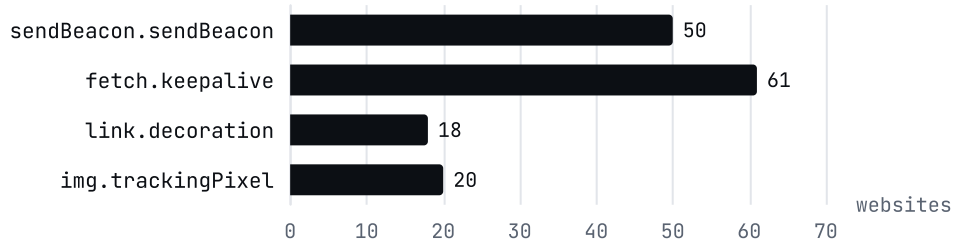


Figure 20. Beacons: the techniques used, by the number of websites on which each was recorded.

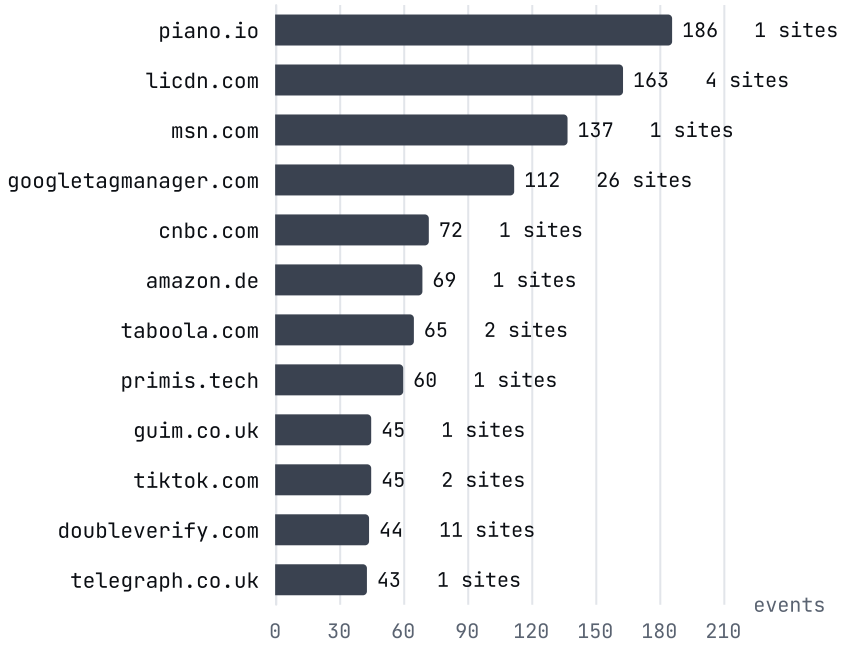


Figure 21. The scripts and servers responsible, by registrable domain: events, and the number of websites on which they caused them. "server" figures are events caused by a server response with no script involved.

- **Keep-alive requests** (`fetch` with `keepalive`) do the same with an ordinary request.
- **Tracking pixels** are invisible images whose address carries data to the server that serves them.
- **Link decoration** adds identifiers to the addresses of requests or links so that the receiving server can recognise the visitor.

THE LAW

Every beacon, pixel and keep-alive request in this section collected data from the visitor's device and sent it to a server, which is gaining access to information stored in the terminal equipment - the EDPB's Guidelines 2/2023 deal with tracking pixels and tracking links specifically, as techniques to which Article 5(3) of the ePrivacy Directive applies.

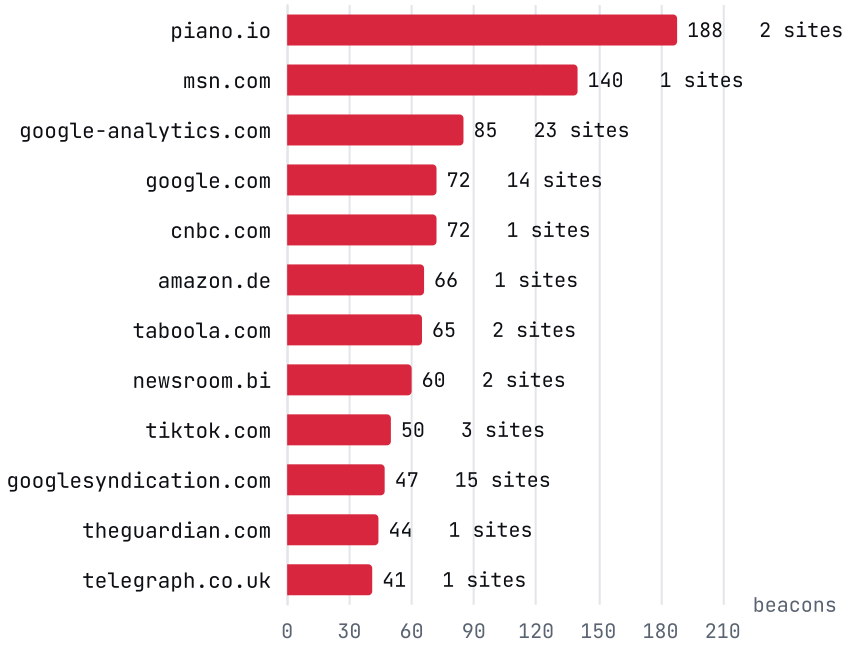


Figure 19. Where the beacons went: receiving domains, by beacons and the number of websites sending them.

EXHIBIT E7 npr.org - privacy notice HIGH

`sendBeacon.sendBeacon` at 2026-10-02T01:01:56.481Z

https://www.npr.org/about-npr/179878450/privacy-policy

api	Navigator.prototype.sendBeacon
urlHost	c2.piano.io
bodyKind	null
bodySize	0
bodyFormat	opaque
sensitiveKeysPresent	None

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at I (https://cdn.piano.io/api/tinypass.min.js:1:174948)
at Object.handleEvent (https://cdn.piano.io/api/tinypass.min.js:1:176305)
at Object.event (https://cdn.piano.io/api/tinypass.min.js:1:193298)
at Q (https://cdn.piano.io/api/tinypass.min.js:1:397362)
at https://cdn.piano.io/api/tinypass.min.js:1:408369
```

Record `sess_314b66d1-e00d-4222-8abf-524173727fed/rec_000002_c04792e0347d.json` - chain position 2
SHA-256 of the record file: `38c9d683a285ebac2560bc5dd4f3eed2a3f8e332912a3f42e45f910aa8787802`

That is the moment personal data leaves the visitor's control, and doing it without any consent had no lawful basis under Articles 5, 6 and 7 of the GDPR - section 4.11 sets out where the receiving servers are registered and located.

Taking that data from visitors' devices without their knowledge or consent is illegal access under Article 3 of Directive 2013/40/EU and Article 2 of the Budapest Convention, and an offence under California Penal Code section 502. Beacons and pixels that send a third party the full address of the page a visitor is reading - the article, the search terms, the product - together with an identifier for that visitor, let the third party learn the contents of the visitor's communication with the website while it is taking place, without the consent of all parties, and section 631 applies to that in the same way it applies to session replay.

4.7 Referrer leakage

WHAT IT IS

When a page loads a resource from another company's server - a script, a font, an image - the browser normally tells that server which page asked for it, in the Referer header. Depending on the website's settings, that is either just the website's address or the full address of the page, including the article or search being viewed.

WHAT WE FOUND ACROSS THE 100 WEBSITES

Measure	Figure
Events	**40,163**
Websites	100 of 100
By tier	tier 3: 40,163
Where	15,671 on homepages, 13,491 on cookie notices, 11,001 on privacy notices

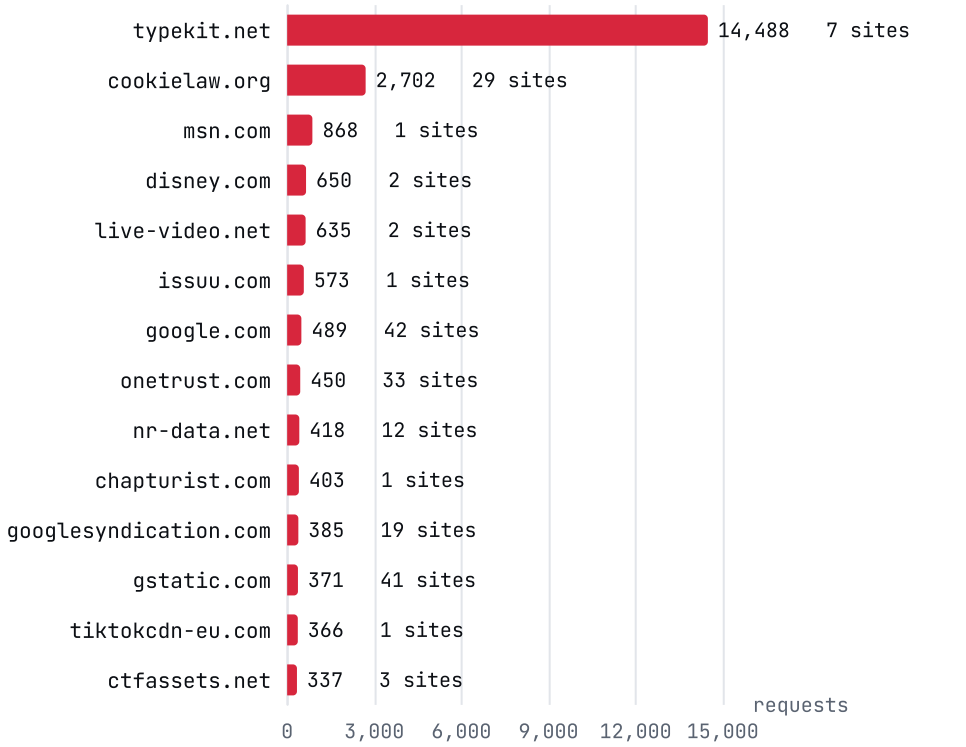


Figure 22. Domains told which page the visitor was on, by requests and the number of websites that told them. 621 domains in all.

EXHIBIT E8 npr.org - privacy notice HIGH

`referer.crossOrigin` at 2026-10-02T01:01:53.556Z

https://www.npr.org/about-npr/179878450/privacy-policy

api	CDP.Network.requestWillBeSent
requestHost	cdn.cookielaw.org
refererHost	www.npr.org
refererPathPresent	True

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
(no script on the stack: the event was caused by a server response - server:cdn.cookielaw.org)
```

Record `sess_314b66d1-e00d-4222-8abf-524173727fed/rec_000001_7ad8efbabdbb.json` - chain position 1
SHA-256 of the record file: `95f0eb94bac8f2a773007d8a7af2d4b95b6a64918e67131b816da411533e5032`

Each such request tells a third party that this browser is on this website at this moment, and on a health or adult website the fact of the visit alone is sensitive (section 4.12).

THE LAW

Every one of these requests told a third party that the visitor was on the website, along with their IP address, which is a disclosure of personal data to that third party, and doing it without a lawful basis, on pages configured to give third parties more than they need, breaches the principles of lawfulness and data minimisation and the duty of security under Articles 5, 6 and 32 of the GDPR - section 4.11 sets out where the receiving servers are registered and located.

Where the referrer carries the full address of the page the visitor is reading, including any search terms in it, it hands a third party the contents of the visitor's communication with the website at the moment it takes place, without the consent of all parties, and section 631 of the California Penal Code applies to it.

4.8 Cross-origin messaging, tracking frames and other browser mechanisms

WHAT IT IS

Browsers keep the code of different websites apart. Several mechanisms exist to let them talk to each other, and they are used to move data between the website and third parties inside the browser, out of sight of the network:

- **postMessage** sends data from one window or frame to another. With a target of `*` (a wildcard), the data can be received by any origin.

WHAT WE FOUND ACROSS THE 100 WEBSITES

Measure	Figure
Events	**4,344**
Websites	94 of 100
By tier	tier 1: 3,935, tier 2: 409
Where	2,522 on homepages, 1,214 on privacy notices, 608 on cookie notices

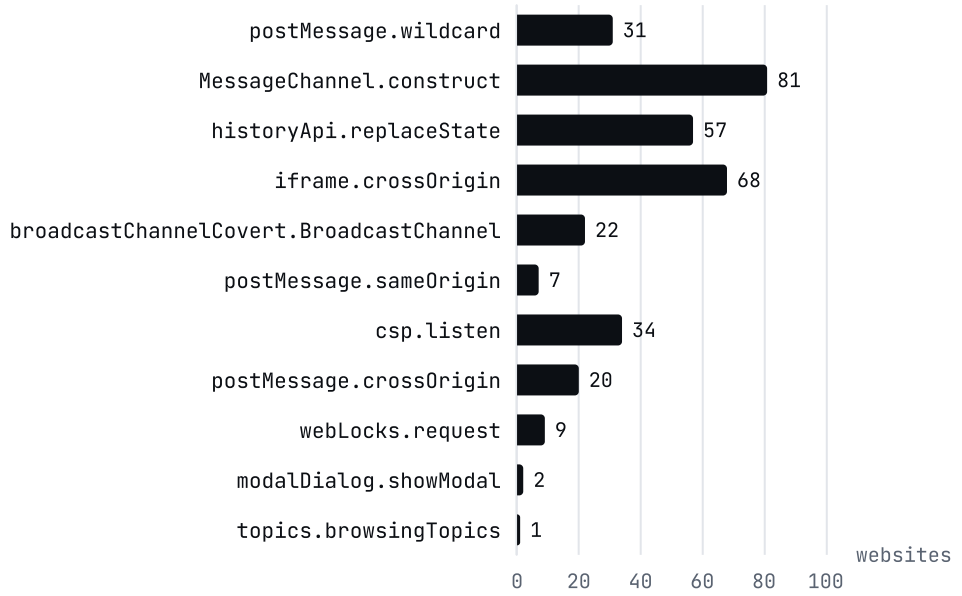


Figure 23. Messaging: the techniques used, by the number of websites on which each was recorded.

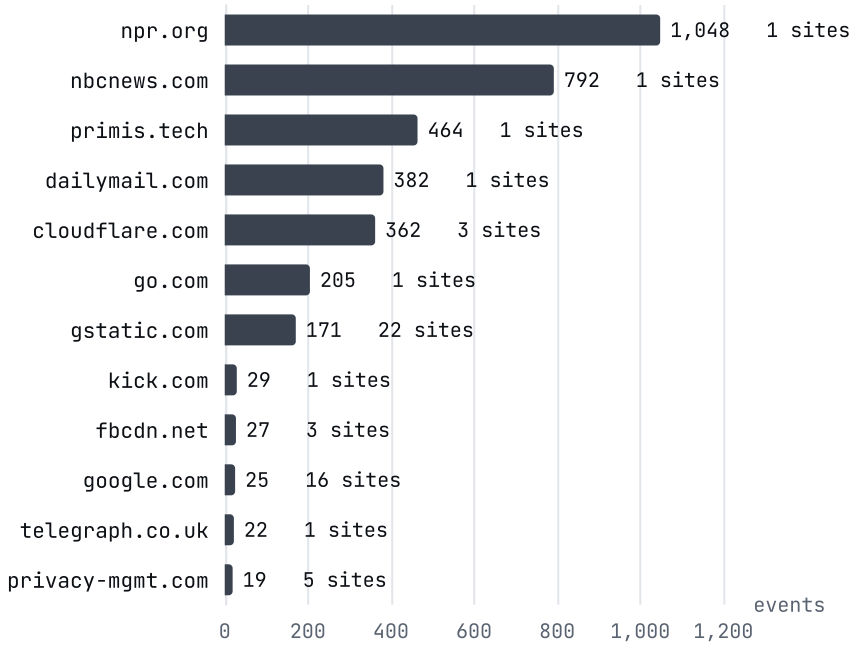


Figure 24. The scripts and servers responsible, by registrable domain: events, and the number of websites on which they caused them. "server" figures are events caused by a server response with no script involved.

- **MessageChannel** and **BroadcastChannel** open channels between frames and between tabs.
- **Hidden third-party frames** embed another company's page invisibly, giving it its own access to the browser.
- **History rewriting** (`history.replaceState`) changes the address shown in the browser without a new page load.
- **Content Security Policy violation listeners**, **Web Locks** and the **Topics** interface can be used to coordinate between scripts or, in the case of Topics, to request the browser's own profile of the visitor's interests.

THE LAW

Each of these mechanisms is used by scripts running on the visitor's device to move information obtained there, which is access to information in the terminal equipment under Article 5(3) of the ePrivacy Directive. Where these websites sent data with a wildcard postMessage, to any origin that happened to be listening, that is a security failure as well as a disclosure, and the hidden third-party frames they embedded gave third parties a presence of their own in the visitor's browser - neither is compatible with data protection by design (Article 25 GDPR) or with security of processing (Article 32), and requesting the browser's profile of the visitor's interests through the Topics interface is processing of personal data about those interests with no lawful basis under Articles 5, 6 and 7.

The scripts doing this are accessing the visitor's device without right, which is illegal access under Article 3 of Directive 2013/40/EU and Article 2 of the Budapest Convention, and an offence under California Penal Code section 502.

EXHIBIT E9 dailymail.com - homepage **HIGH**

`postMessage.wildcard` at 2026-10-02T07:32:19.252Z

https://www.dailymail.com/home/index.html

api	Window.prototype.postMessage
targetOrigin	*
identityPayloadHint	False

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at i (https://www.dailymail.com/i/html_modules/DataModule/data-module-monthly
-carousel/index.html:88:36)
at o (https://www.dailymail.com/i/html_modules/DataModule/data-module-monthly
-carousel/index.html:92:13)
at https://www.dailymail.com/i/html_modules/DataModule/data-module-monthly-ca
rousel/index.html:97:15
```

Record `sess_c683250b-4983-4a2f-adf4-66ddf66669d1/rec_000001_74017131d142.json` - chain position 1

SHA-256 of the record file: `b787c3259b2362b8e1e7b1f496133b70ed8443c5e8585bbbf90033e8b58abfbd`

4.9 Event listeners, form reads and clipboard monitoring

WHAT IT IS

Scripts can listen to what a visitor does - every movement of the mouse, every scroll, every touch, every key - and read what is typed into forms, including before it is submitted. Recorded together and sent to a server, this is how "session replay" tools reconstruct a visit, and how behavioural analysis and bot-detection products profile a visitor by the way they move.

Measure	Figure
Events	**3,479**
Websites	99 of 100
By tier	tier 1: 3,479
Where	1,990 on homepages, 850 on privacy notices, 639 on cookie notices

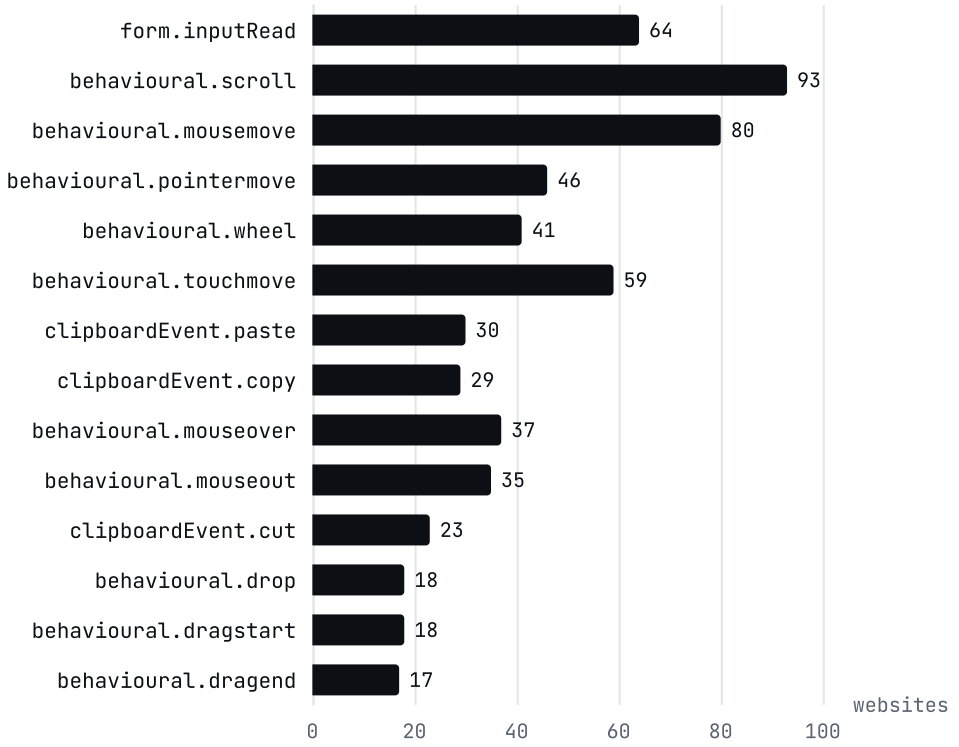


Figure 25. Listeners: the techniques used, by the number of websites on which each was recorded.

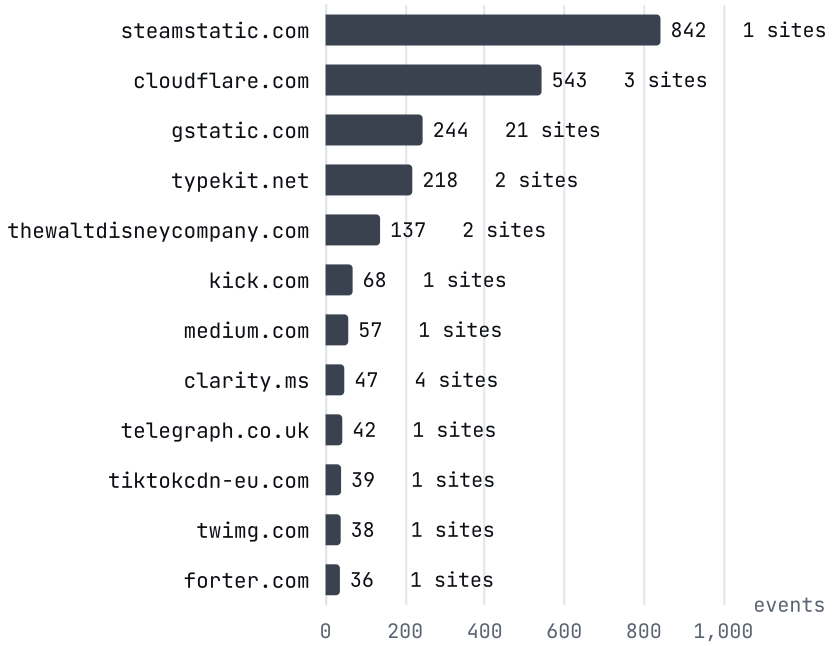


Figure 26. The scripts and servers responsible, by registrable domain: events, and the number of websites on which they caused them. "server" figures are events caused by a server response with no script involved.

Scripts can also watch the clipboard as content is copied, cut or pasted.

The visitor's interaction with a page is a communication between the visitor and the website, and recording it for a third party is interception.

WHAT WE FOUND ACROSS THE 100 WEBSITES

Scripts read the contents of form fields on 64 websites without any consent (1,469 reads). Exhibit E10 shows Google's reCAPTCHA reading the search box on reddit.com's privacy notice.

THE LAW

A visitor's interaction with a page is part of their communication with the website, and Article 5(1) of the ePrivacy Directive prohibits listening to, recording or otherwise intercepting a communication without consent - capturing that interaction for a third party is interception, and recording behaviour systematically is surveillance. Reading form values and registering listeners on the device is also access to information in the terminal equipment under Article 5(3), and as behavioural data and form contents are personal data that can include anything a visitor types, Articles 5, 6, 7, 25 and 32 of the GDPR apply to all of it.

Intercepting non-public transmissions of computer data is an offence under Article 6 of Directive 2013/40/EU and Article 3 of the Budapest Convention. In California, reading or attempting to learn the contents of a communication in transit without the consent of all parties is covered by Penal Code section 631, the provision on which session-replay claims in California are based, and accessing the device's data without permission is an offence under section 502.

EXHIBIT E10 reddit.com - privacy notice LOW

`form.inputRead` at 2026-10-01T15:38:17.314Z

https://www.reddit.com/policies/privacy-policy?solution=34113f85f10b7e8c34113f85f10b7e8c&js_challenge=1&jsc_token=2824be10929bdc604753c70a67a1c3310d19f61a7f57d76224aa86318424b00a&jsc_orig_r=

api	HTMLInputElement.prototype.value (getter)
fieldName	q
fieldType	text
formAction	https://www.reddit.com/search/?q=

CALL STACK CAPTURED AT THE MOMENT OF THE EVENT

```
at https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjgsGdhRslyEPwU/recaptcha__en.js:1266:62
at Array.<anonymous> (https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjgsGdhRslyEPwU/recaptcha__en.js:772:24)
at https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjgsGdhRslyEPwU/recaptcha__en.js:1265:495
at https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjgsGdhRslyEPwU/recaptcha__en.js:587:458
at Array.<anonymous> (https://www.gstatic.com/recaptcha/releases/kemdRjWFxNjgsGdhRslyEPwU/recaptcha__en.js:758:432)
```

Record `sess_f67430ad-399d-4145-b613-737647594ac0/rec_000001_6a46fa9e296e.json` - chain position 1

SHA-256 of the record file: `8b035f77e680ab5b01c9c28d28f2f9e56afe62938778db30e6f163865c74b490`

4.10 Client hints, ETags, HSTS, Alt-Svc and pre-emptive connections

WHAT IT IS

These mechanisms work at the level of the web's protocols, without any script being involved. Each has a legitimate function and each can also be used to recognise a visitor - some are documented ways of doing exactly that when cookies are blocked or deleted - which is why they are counted, in tier 3.

- **Client hints.** A server can instruct the browser (with an Accept-CH header) to send detailed information about the device - platform, model, architecture, version - with every subsequent request.

WHAT WE FOUND ACROSS THE 100 WEBSITES

Measure	Figure
Events	**138,928**
Websites	100 of 100
By tier	tier 3: 138,928
Where	55,822 on homepages, 45,143 on privacy notices, 37,963 on cookie notices

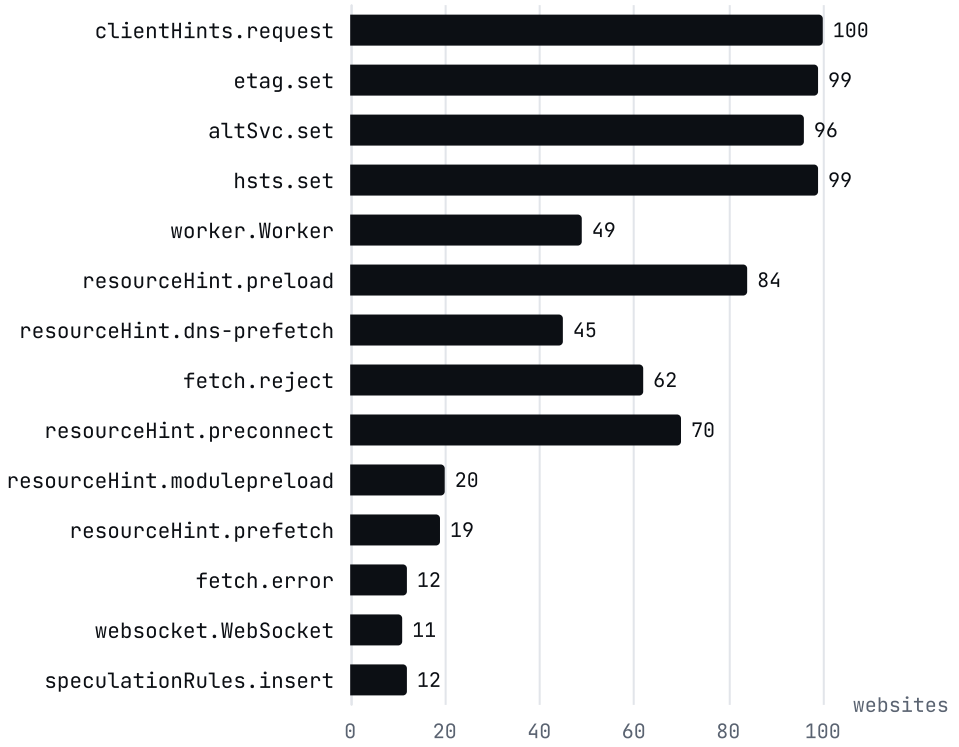


Figure 27. Headers: the techniques used, by the number of websites on which each was recorded.

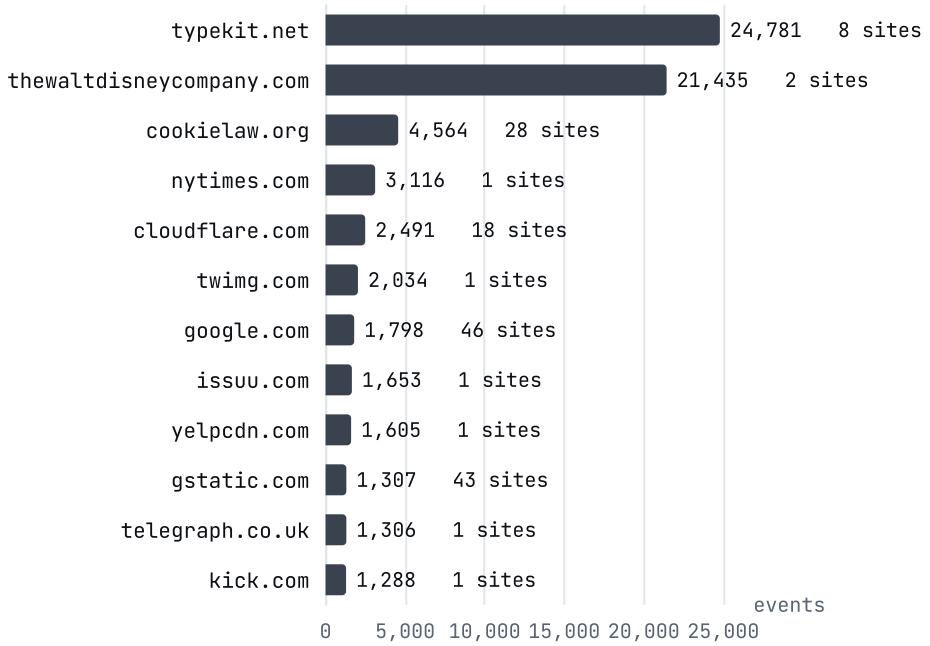


Figure 28. The scripts and servers responsible, by registrable domain: events, and the number of websites on which they caused them. "server" figures are events caused by a server response with no script involved.

- **ETags.** A server attaches a tag to a resource; the browser stores it in its cache and sends it back the next time it requests the resource. A unique ETag works as an identifier that survives the deletion of cookies.
- **HSTS.** A server tells the browser to use secure connections in future; the browser stores that instruction per host. The pattern of stored instructions across many hosts can be used as a "supercookie".
- **Alt-Svc.** A server tells the browser about alternative ways to connect to it, which the browser stores and can reuse to recognise the same browser.
- **Pre-emptive connections and background workers.** Pages instruct the browser to resolve, connect to or fetch from third-party servers before anything is needed, and start workers and WebSockets that keep connections open.

THE LAW

An ETag is stored in the browser's cache and sent back with the next request, HSTS and Alt-Svc state is stored by the browser, and a client hints instruction makes the browser hand over information about the device with every request - each of these is storage in, or access to information from, the terminal equipment under Article 5(3) of the ePrivacy Directive. Where these mechanisms are used to recognise a visitor, the data they carry is personal data, and tracking carried out by a server needs a lawful basis under Articles 5, 6 and 7 of the GDPR in exactly the same way as tracking carried out in the browser.

4.11 International transfers

WHAT IT IS

Chapter V of the GDPR (Articles 44 to 50) allows personal data to be transferred to a country outside the European Economic Area only on one of the bases it sets out - an adequacy decision, appropriate safeguards such as standard contractual clauses, or one of the narrow derogations. Every request a visitor's browser makes to a server carries at least their IP address, and usually much more, so every request to a server outside the EEA is a candidate transfer.

HOW WE MEASURED IT

The capture records the IP address of every server each page connected to. For each server address the website's pages contacted we first obtained the network it belongs to, the organisation that operates that network, and the country in which the address block is registered, from the Team Cymru IP-to-ASN mapping service, on 7th October 2026.

This is registry data, which identifies the jurisdiction of the organisation operating the receiving network without proving where a particular request was processed. Content delivery networks, in particular, answer from servers close to the visitor, so many requests to US-registered networks will have been served from inside the EEA.

Measure	Figure
Servers contacted	1,400
Of which third-party	1,034
On networks registered outside the EEA	1,293 (92.4%)
Located outside the EEA, of those whose location could be determined	380 of 954

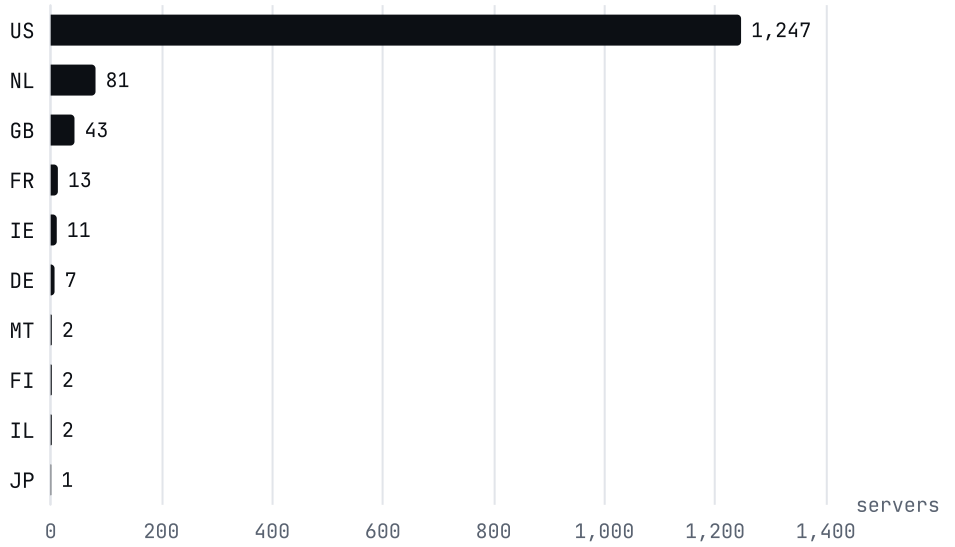


Figure 29. Servers contacted without any consent, by the country in which the receiving network is registered (Team Cymru IP-to-ASN).

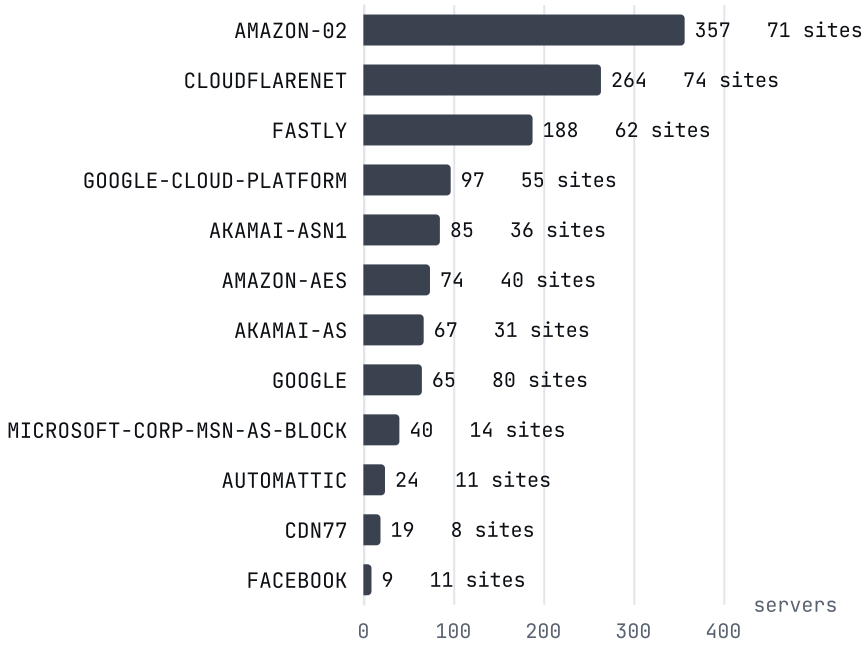


Figure 30. The organisations operating the receiving networks, by servers and the number of websites whose pages contacted them.

They remain within the scope of this section, because the operator receiving the data is subject to the law of the country where it is established - the concern the Court of Justice addressed in Schrems II.

WHAT WE FOUND ACROSS THE 100 WEBSITES

Geolocation: DB-IP IP to Country Lite, October 2026 (IP geolocation by DB-IP, licensed under CC BY 4.0). Cloudflare and Fastly answer from shared anycast addresses and are treated as indeterminate.

THE LAW

Every transfer counted here happened without any consent, and where consent was required for the processing in the first place (sections 4.1 to 4.10) there was no lawful basis under Article 6 of the GDPR - so the question of which Chapter V mechanism applies never arises, because no transfer mechanism can make lawful a processing operation that has no lawful basis.

Transfers to US organisations certified under the EU-US Data Privacy Framework rely on Commission Implementing Decision (EU) 2023/1795. The General Court dismissed a challenge to that decision in Latombe v Commission (T-553/23, judgment of 3rd September 2025, ECLI:EU:T:2025:831); an appeal to the Court of Justice was lodged on 31st October 2025 (C-703/25 P) and was pending when this report was written. Its predecessor, the Privacy Shield, was annulled by the Court of Justice in Schrems II (C-311/18, judgment of 16th July 2020, ECLI:EU:C:2020:559) because of the access US intelligence agencies have to data held by US companies. Whether each receiving organisation is certified under the Framework, and for which data, is something each website has to be able to show, and certification of a content delivery network does not extend to the advertising and analytics companies whose servers the data reaches.

Articles 13(1)(f) and 14(1)(f) also require a privacy notice to state the intention to transfer data and the mechanism relied on, and section 4A sets out how many notices do.

4.12 Websites that reveal special categories of data

WHY THESE WEBSITES ARE DIFFERENT

Article 9 of the GDPR prohibits the processing of personal data revealing, among other things, a person's health, religious beliefs, sex life or sexual orientation, unless one of a narrow set of conditions applies - in this context, explicit consent.

Six of the websites in this study deal with subjects of that kind: mayoclinic.org (health), biblegateway.com (religious belief), and four adult websites - pornhub.com, xvideos.com, xhamster.com and stripchat.com (sex life and sexual orientation).

A visit to one of these websites does not by itself prove anything about the visitor, and not every event on them involves special category data, but the Court of Justice has made clear that data can reveal a special category indirectly, through what can be inferred from it:

- In **OT** (C-184/20, 1st August 2022, ECLI:EU:C:2022:601), the Court held that personal data liable to disclose a person's sexual orientation indirectly, by an intellectual operation of comparison or deduction, falls within Article 9.
- In **Meta Platforms v Bundeskartellamt** (C-252/21, 4th July 2023, ECLI:EU:C:2023:537), the Court held that where a user visits websites or apps relating to one of the Article 9 categories, the collection of data about those visits can amount to processing of special category data.
- In **Lindenapotheke** (C-21/23, 4th October 2024, ECLI:EU:C:2024:846), the Court held that information entered when ordering pharmacy-only medicines online is health data even though it is not certain that the medicine is for the person ordering it.

The European Court of Human Rights has long treated a person's sex life as a most intimate part of private life protected by Article 8 of the Convention (Dudgeon v United Kingdom, 1981).

What a person reads on a health website, or watches on an adult website, is exactly the kind of information from which health or sexual orientation is inferred.

Website	Subject	Events	Tier 1	Fingerprinting	Domains told of the visit	Third-party servers	Of which outsid the EEA
stripchat.com	sex life and sexual orientation	7,867	43	1,455	7	8	8
pornhub.com	sex life and sexual orientation	6,805	13	150	7	12	7
mayoclinic.org	health	2,586	68	415	6	10	8
xhamster.com	sex life and sexual orientation	1,934	36	139	6	11	9
xvideos.com	sex life and sexual orientation	730	8	39	3	7	5
biblegateway.com	religious belief	498	6	19	2	2	2

A visitor reading about a lump or a mole on Mayo Clinic, or watching non-heterosexual pornography on xHamster, discloses something about themselves, and once that is collected and sent to third parties it is special category data in their hands.

WHAT WE FOUND

"Domains told of the visit" counts the other domains that received the visitor's page in a referrer (section 4.7). All of this happened without any consent - let alone the explicit consent Article 9(2)(a) requires.

THE LAW

Where the data collected reveals a visitor's health, religious belief, sex life or sexual orientation, directly or by inference, processing it needs one of the conditions in Article 9(2) of the GDPR, and none was present because no consent of any kind was given. These six websites deal with exactly those subjects, so they have to design their processing with that in mind (Articles 5 and 25), and high-risk processing of this kind needs a data protection impact assessment (Article 35) - sending visitors' presence and device fingerprints to third parties outside the EEA without any consent is incompatible with both.

Every provision that applies in sections 4.1 to 4.11, including the criminal law, applies to these websites as well, and the Article 9 analysis is in addition to them.

4A. Disclosures: privacy and cookie notices

Articles 12 to 14 of the GDPR require a controller to tell people, in a concise, transparent, intelligible and easily accessible form, using clear and plain language, who is processing their data, why, on what basis, who receives it, where it goes, how long it is kept and what their rights are. The ePrivacy Directive's consent requirement depends on that information too: consent is only valid if it is informed.

This section looks at whether the 100 websites provide notices at all, how readable they are, what they contain, and whether what they say matches what the websites did.

LENGTH AND READABILITY

Length and readability were measured on the full text of each website's privacy notice, retrieved with every collapsed or hidden section opened (13th November 2026), and on the text of each cookie notice as captured during the audit; cookie notices whose captured page did not hold the notice text were excluded.

COVERAGE

	Websites
Privacy notice found	97
No privacy notice found	3 - archive.org and discogs.com (no link we could locate), indiatimes.com (the privacy notice link in its footer returned "page not found")
Cookie information found	84 - a separate page on 68, a section of the privacy notice on 16
No cookie information found	16

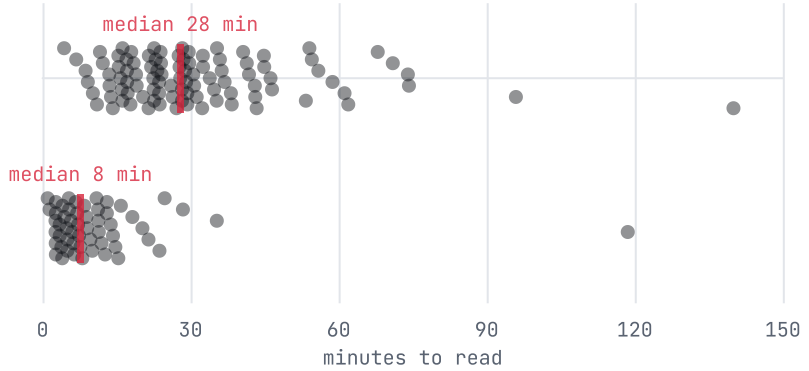


Figure 31. Time to read each privacy and cookie notice whose full text was captured, in minutes at 238 words a minute. The red line marks the median.

	Privacy notices (97)	Cookie notices (55)
Median length	6,641 words	1,824 words
Longest	33,320 words	28,222 words
Median time to read	about 28 minutes	about 8 minutes
Longest time to read	about 2 hours 20 minutes	about 1 hour 59 minutes
Median reading age	20.6 years	19.7 years
Median Flesch Reading Ease	31.7	36.3
Rated difficult or very difficult (below 50)	97 of 97	52 of 55
Reading age of 18 or over	79 of 97	41 of 55

The measurement covers 97 privacy notices and 55 cookie notices.

Readability was measured with the Flesch Reading Ease score (0 to 100, where below 50 is difficult and below 30 is very difficult) and the Flesch-Kincaid grade level, converted to a reading age by adding five years. Reading time assumes an average adult silent reading speed of 238 words per minute for non-fiction (Brysbaert, 2019).

Every privacy notice we could measure is written at a level classed as difficult or very difficult, and the median notice needs a reading age above university entry and half an hour to read.

CONTENT AGAINST ARTICLES 13 AND 14

We screened the text of 97 privacy notices for wording corresponding to each item of information Articles 13 and 14 require, using the full text of each notice and, for the one website where no full text was held, the text captured during the audit.

This screen is **indicative only**.

Finding the wording does not show that the information is accurate or complete, and not finding it does not prove it is absent - a notice can express a requirement in words the screen does not recognise. We therefore report the screen in aggregate and do not state that any named website's notice omits an item.

WHAT THE NOTICES SAY COMPARED WITH WHAT THE WEBSITES DID

For 99 websites we hold the text of the notices - each privacy notice in full and each cookie notice as captured. For those, we compared what the notices say with what the websites did without any consent. Every match described below was confirmed by reading the passage concerned.

Of the 99 websites, 59 used high-entropy fingerprinting techniques without any consent (section 4.1), and only two of them disclose device fingerprinting in their notices: dailymail.com states that it uses "digital fingerprint technology", also known as "Machine Identification", for the Ipsos Iris audience measurement, and issuu.com lists "fingerprinting technique" among the tracking technologies its privacy policy treats as cookies.

Information required	Notices where wording was found (of 97)
Identity and contact details of the controller	97
Purposes of processing	97
Recipients	96
Retention period	95
Right to complain to a supervisory authority	92
Legal basis	90
Legitimate interests	83
Right to erasure	83
Right to restriction	83
International transfers	72
Right to object	72
Right of access	70
Right to withdraw consent	70
Right to data portability	57
Data protection officer contact	55
Right to rectification	55
Automated decision-making and profiling	46
Source of data not obtained from the person	35

Website	What the notice says	What the website did without any consent
npr.org	"we only set these non-Essential cookies with your consent"	Set 28 cookies, including advertising cookies on criteo.com (uid, cto_bundle) and doubleclick.net (test_cookie); 49 storage writes; 1,699 fingerprinting events
theguardian.com	"With your consent we share and receive online data collected through cookies with our advertising partners"	Set 10 cookies; 11 storage writes; 560 fingerprinting events
imdb.com	Cookies remain "for 13 months after you provide us with your consent"	Set 17 cookies; 84 storage writes; 8,410 fingerprinting events
unsplash.com	Non-essential cookies "are usually placed on your device with your consent"	Set 3 cookies; 3 storage writes; 200 fingerprinting events
pornhub.com	"We will ask for your consent annually, or whenever we make significant changes to our cookies"	Set 8 cookies; 4 storage writes; 150 fingerprinting events
flickr.com	"Where required by law, we use advertising and analytics cookies only with your consent"	Set 19 cookies; 53 storage writes; 1,547 fingerprinting events
oup.com	We "may ask for your consent to send you marketing and to place cookies on your device"	Set 2 cookies; 6 storage writes; 119 fingerprinting events
independent.co.uk	Non-essential cookies	Set 12 cookies; 26 storage writes

Four other notices use the word "fingerprint", but only about biometric fingerprints - a sign-in method, fingerprints collected at offices or theme parks, or biometric data listed as a category of personal data - and not about the device fingerprinting their websites carried out.

The 99 websites set 1,332 distinct cookies without any consent, and only 83 of them (6.2%) are named anywhere in the website's own privacy or cookie notices - 68 of the 96 websites that set cookies name none of them.

Eight websites state in their notices that cookies or similar technologies are used with consent, and set them without any consent:

NOTICES THAT MATCH WHAT THE WEBSITES DID

Not one of the websites has a notice that fully describes what its pages did without any consent. Of the 95 websites whose notice text was captured, only two passed the automated screen - every cookie the website set is named in its notices, and fingerprinting is disclosed where it was used - and both fail when the notices are read: facebook.com mentions fingerprints only as a biometric sign-in method, and arxiv.org does not name Adobe (Typekit) or GitHub, whose servers its pages contacted.

THE TIME AND COST OF READING THE NOTICES

To find out what happens to their data on a website, people have to read the website's privacy notice and the privacy notice of every third party its pages send data to. To measure that, we replicated the method of McDonald and Cranor, "The Cost of Reading Privacy Policies" (I/S: A Journal of Law and Policy for the Information Society, volume 4, number 3, 2008). That research was done in 2008, and the reading speed it uses - 250 words a minute - is the one from its own model and data, kept here so that our results can be compared with it; it is not the 238 words a minute (Brysbaert, 2019) used elsewhere in this report. Following the same model, time spent reading at work is valued at twice the average hourly wage and time at home at a quarter of it, weighted by the paper's own estimate of 59 policies read at work and 112 at home each month. For the average hourly wage we used the mean gross hourly earnings in the EU in 2022, €17.69 (Eurostat, earn_ses22_15), which values each hour of reading at €15.10.

The third parties are the companies operating the servers each website's pages contacted, identified from the DuckDuckGo Tracker Radar entity map and checked by hand; companies belonging to the same group as the website are left out. Each company's privacy notice was retrieved on 12th November 2026 and its words counted, including text held in collapsed sections.

Website	Third parties	Words	Time to read	Value of that time
techcrunch.com	40	242,753	16 hours 11 minutes	€244.43
mlb.com	28	209,223	13 hours 57 minutes	€210.67
nbcnews.com	28	173,633	11 hours 35 minutes	€174.83
apnews.com	31	169,604	11 hours 19 minutes	€170.78
prnewswire.com	24	166,246	11 hours 5 minutes	€167.40
merriam-webster.com	24	164,142	10 hours 56 minutes	€165.28
hbr.org	17	154,921	10 hours 20 minutes	€155.99
stackoverflow.com	16	129,035	8 hours 36 minutes	€129.93
cnn.com	23	128,560	8 hours 34 minutes	€129.45
time.com	25	110,361	7 hours 22 minutes	€111.12

Website	EU monthly active recipients	Time per person	Value per person	Total time	Total value
pinterest.com	87,497,257	2 hours 2 minutes	€30.87	178.5 million	€2.70 billion
wikipedia.org	157,324,000	34 minutes	€8.42	88.1 million	€1.32 billion
linkedin.com	45,000,000	1 hour 35 minutes	€23.84	71.1 million	€1.07 billion
twitter.com	45,000,000	1 hour 21 minutes	€20.41	60.8 million	€0.92 billion
reddit.com	45,000,000	1 hour 14 minutes	€18.72	55.8 million	€0.84 billion
pornhub.com	45,000,000	1 hour 13 minutes	€18.42	54.9 million	€0.83 billion
tiktok.com	45,000,000	1 hour 11 minutes	€17.96	53.5 million	€0.81 billion
facebook.com	45,000,000	1 hour 11 minutes	€17.80	53.1 million	€0.80 billion
instagram.com	45,000,000	1 hour 11 minutes	€17.85	53.1 million	€0.80 billion
youtube.com	45,000,000	59 minutes	€14.76	44.1 million	€0.66 billion
xhamster.com	35,000,000	1 hour 7 minutes	€16.82	38.9 million	€0.59 billion
xvideos.com	45,000,000	28 minutes	€6.91	20.7 million	€0.31 billion
yelp.com	1,100,000	2 hours 7 minutes	€32.04	2.3 million	€0.04 billion
flickr.com	197,000	4 hours 37 minutes	€69.58	0.9 million	€0.01 billion

On the median website, reading the website's notice and the notices of the third parties on its pages takes 2 hours 39 minutes - 39,821 words - and that time is worth €40.10 for each visitor. Reading them for all 100 websites would take one person 353 hours.

Using the EU monthly active recipients the very large online platforms publish under the Digital Services Act (section 6), and counting one reading per person as the model does:

For people in the EU to read the notices of these 14 platforms and of the third parties on their homepages would take at least 776 million hours, worth €11.7 billion. These are minimums: only the third parties contacted from the pages we audited are counted, and the notices listed below could not be counted.

DATES OF LAST CHANGE

278 of the 320 notices retrieved state when they were last changed; 42 do not: Ad Lightning, Inc.; Akamai Technologies; Almondnet Group; Bending Spoons S.p.A.; Blockthrough Inc.; Brandmetrics; BrightLine Partners LLC; Collective Roll; Conversant LLC; Dianomi Ltd; Dotmetrics; Hexagon Data, S.A.P.I. de C.V.; ID5 Technology Ltd; Marfeel Solutions, S.L.; Outbrain; Perfect Sense Operations; Sharethrough, Inc.; Smartadserver S.A.S; The Ozone Project; Times Internet Ltd; Unity Software Inc.; Usercentrics GmbH;

Woorank sa; aditude.cloud; archive.org; cnn.com; consentmanager AB; datazoom.io; discogs.com; htlbid.com; kueezrtb.com; media-lab.ai; mediamelon.com; minutemedia-prebid.com; newsroom.bi; polymarket.com; programmaticx.ai; recombee.com; sagepub.com; sitefinity.com; trueanthem.com; yourbow.com.

THIRD PARTIES WITH NO PRIVACY NOTICE

For 47 third parties, receiving data from 30 of the websites, there is no privacy notice at all - the domain serves no website, or a website with no privacy notice, and where a company is known to operate it, that company publishes no notice we could find. Visitors have no way of finding out who these companies are, let alone what they do with their data, and the websites sending data to them cannot name them as recipients as Article 13(1)(e) of the GDPR requires, while the companies themselves give no information under Article 14.

A further 16 companies publish a notice we could not retrieve in full and are not counted in the figures above: Ad Lightning, Inc., Bending Spoons S.p.A., Collective Roll, Conversant LLC, Dianomi Ltd, Dotmetrics, Hotjar Ltd, ID5 Technology Ltd, Polar, Sharethrough, Inc., Smartadserver S.A.S, Woorank sa, aditude.cloud, htlbid.com, kueezrtb.com, minutemedia-prebid.com.

THE LAW

Article 12(1) of the GDPR requires the information to be concise, transparent, intelligible and easily accessible, in clear and plain language - the median notice here takes half an hour to read and needs a reading age above 20, which meets none of those requirements.

Articles 13 and 14 make every item in the table above mandatory, and if a notice does not name the cookies a website sets, the fingerprinting it carries out or the companies receiving the data, visitors have no way of making an informed decision about any of them. Consent has to be informed to be valid (Articles 4(11) and 7 of the GDPR, which also govern consent under Article 5(3) of the ePrivacy Directive), and here no consent was obtained at all.

Third party	Websites
easyAd Deutschland GmbH	accuweather.com, apnews.com, cnn.com, imgur.com, msn.com, theverge.com, time.com
dns-finder.com	accuweather.com, cnn.com, imgur.com, msn.com, theverge.com, time.com
CacheNetworks, LLC	imgur.com, independent.co.uk, washingtonpost.com
prodregistryv2.org	bloomberg.com, slideshare.net, soundcloud.com
wknd.ai	apnews.com, imgur.com, theverge.com
unpkg	techcrunch.com, washingtonpost.com
CDNWIDGET.COM	apnews.com
Clicksco FZ	businessinsider.com
Podsights	hbr.org
Snagfilms	cnbc.com
amspbs.com	stackoverflow.com
arttrk.com	mlb.com
assemblyexchange.com	imgur.com
atmtd.com	businessinsider.com
bkcdn.net	xvideos.com
bmtrcs.com	washingtonpost.com
brilliantchap.com	cnn.com
cdn-vitals.com	apnews.com
cdndex.io	kick.com
chapturist.com	stripchat.com
claydar.com	prnewswire.com
coalmerchants.com	aol.com
coconts.com	stripchat.com
drivetheweb.com	tripadvisor.com
dualping.com	apnews.com

The five websites whose notices say cookies are used with consent, and which set them without any consent, have misleading notices. That is how the FTC Act treats a statement about data practices that does not match what a company does (section 5, 15 U.S.C. 45), and how EU consumer protection law treats it, as a misleading commercial practice (Directive 2005/29/EC, Article 6).

4B. Built-in AI interfaces

WHAT THEY ARE

Google Chrome now ships interfaces that let websites use AI models built into the browser and running on the visitor's device: the Prompt API (`LanguageModel`), the Summarizer API (`Summarizer`) and related interfaces for writing, rewriting, translation, language detection and proofreading. A website can check whether the model is available on the device, create a session with it and give it text to process.

The availability of these interfaces depends on the device, the browser version and its settings, which makes it one more characteristic that can be read to tell devices apart.

HOW WE MEASURED IT

The version of WebSentinel used for this research did not record calls to these interfaces (section 3.5). Instead, we searched the full source of every script delivered to every audited page for code that uses them, and read the code around every match. In addition, the on-device model was switched off in the audit browser, so any attempt to use it could only have found it unavailable.

These findings therefore show **code delivered to visitors' devices** without any consent and written to use the built-in AI interfaces - the capture could not show whether that code ran.

WHAT WE FOUND

37 of the 100 websites delivered code that uses the browser's built-in AI interfaces.

Website	Script
cnn.com	https://ob.brilliantchap.com/i/674ef2f97b780fb53e1c2cddf38dbf22.js
prnewswire.com	https://euob.segreencolumn.com/sxp/i/a13c666c09e9671170b9d9245

ADVERTISING: 36 WEBSITES

Google Publisher Tag - Google's advertising library, served from securepubads.g.doubleclick.net and pagead2.googlesyndication.com - contains code that:

1. checks whether the `LanguageModel` and `Summarizer` interfaces exist and can be called;
2. checks whether the model is available;
3. creates a session with each - `LanguageModel.create({})` and `Summarizer.create({type: "key-points", format: "plain-text", length: "medium"})`;
4. measures how long creating the session takes, and records the outcome - unavailable, or success and the time taken.

This code was delivered on 36 websites: accuweather.com, apnews.com, bbc.com, businessinsider.com, cbsnews.com, cnbc.com, cnet.com, cnn.com, espn.com, fandom.com, flickr.com, ft.com, hbr.org, imgur.com, independent.co.uk, indiatimes.com, kick.com, merriam-webster.com, mlb.com (including on its privacy notice), nature.com, nbcnews.com, npr.org, nytimes.com, people.com, soundcloud.com, springer.com, stackoverflow.com, tandfonline.com, techcrunch.com, ted.com, telegraph.co.uk, theverge.com, time.com, washingtonpost.com, wired.com and wsj.com.

Whether a device has an on-device AI model, and how quickly a session with it can be created, varies between devices, so an advertising library measuring both on every page view is collecting device characteristics without any consent.

FINGERPRINTING: 2 WEBSITES

On two websites, scripts served from randomised-looking domains include the availability of the `LanguageModel` interface as one item in a list of feature tests - alongside the battery interface and others - and combine the results into a string of ones and zeros. That string is a fingerprint, and the presence of on-device AI is one bit of it.

THE LAW

Checking whether an AI interface or model is available on the visitor's device, creating a session with it and timing it are all access to information in the terminal equipment under Article 5(3) of the ePrivacy Directive, and the EDPB's Guidelines 2/2023 deal specifically with processing carried out locally on the device. Doing this for advertising measurement or for fingerprinting is not strictly necessary for any service the visitor asked for, and the results are characteristics of the visitor's device, collected and used without any lawful basis under Articles 5, 6, 7 and 25 of the GDPR.

Article 22 of the GDPR protects people against decisions based solely on automated processing, including profiling, which produce legal effects concerning them or similarly significantly affect them. Fingerprinting a person's device without consent infringes their fundamental rights to privacy and data protection under Articles 7 and 8 of the Charter, and a fully automated system that profiles a device, assigns it an identifier and so decides how that person is tracked and treated produces a legal effect on them, because it changes their legal position by subjecting them to processing that interferes with those rights. The Court of Justice has read the notion of a "decision" under Article 22 broadly, holding that an automated score on which others rely in deciding how to treat a person can itself fall within the provision (SCHUFA, C-634/21, 7th December 2023, ECLI:EU:C:2023:957).

The same interfaces being probed here let websites run AI processing on the visitor's own hardware, using the visitor's processor, memory and power, without any consent - and using a person's device to do the website's work is a further access to their terminal equipment that needs consent under Article 5(3).

It is a criminal matter as well, on the same basis as the techniques in the earlier sections. Probing and using the device's AI interfaces without the visitor's knowledge or consent is access to an information system without right (Directive 2013/40/EU, Article 3;

Budapest Convention, Article 2), and running processing on the visitor's hardware for the website's own purposes interferes with the functioning of that system (Directive 2013/40/EU, Article 4; Budapest Convention, Article 5). In California, knowingly accessing a computer without permission, and knowingly using its computer services without permission, are offences under Penal Code section 502(c)(7) and (c)(3).

4C. Dark patterns in the consent banners

WHAT THEY ARE

Dark patterns - the European Data Protection Board now calls them deceptive design patterns - are interfaces and user journeys that try to push people into decisions about their personal data that they did not intend and would not otherwise make, decisions that serve the company and work against the person (EDPB Guidelines 03/2022, version 2.0). The EDPB groups them in six categories: Overloading, Skipping, Stirring, Obstructing, Fickle and Left in the dark.

A consent banner is where they do the most damage. It is the one moment a website asks for permission before it tracks a visitor, and the way that question is designed decides the answer. The EDPB's Cookie Banner Taskforce set out the banner practices the supervisory authorities agree do not produce valid consent - no option to reject on the first layer, pre-ticked boxes, a reject option hidden in a link, an accept button made to stand out over the reject option, legitimate interest claimed for tracking (report of 17th January 2023).

Pattern	What it is	Source
P1 No option to reject on the first layer	The banner offers a way to accept but none to refuse on its first layer, so refusing takes more steps than accepting.	Taskforce report, Type A (paras 6-8); Guidelines 03/2022, Obstructing - Longer than necessary (4.4.2)
P2 The reject option is only a link	Accepting is a button; refusing is a link or plain text with far less visual weight.	Taskforce report, Type C (paras 11-14); Guidelines 03/2022, Stirring - Hidden in plain sight (4.3.2)
P3 The accept button is made to stand out	The accept button is given more visual weight than the reject option - a filled or brighter colour, stronger contrast, a larger size or a more prominent place.	Taskforce report, Types D and E (paras 15-18); Guidelines 03/2022, Stirring - Hidden in plain sight (4.3.2)
P4 Pre-ticked boxes	Purposes are already switched on in the banner before the visitor has done anything.	Taskforce report, Type B (paras 9-10); recital 32 GDPR; Planet49 (C-673/17); Guidelines 03/2022, Skipping - Deceptive snugness (4.2.1)
P5 Legitimate interest claimed	The banner tells the visitor that data will be processed on the basis of legitimate interest. Article 5(3) of the ePrivacy Directive allows information to be stored on or read from a device only with consent, or where it is strictly necessary for the service the user asked for or for transmitting a communication - there is no legitimate interest basis, so a banner claiming one for storing or reading information on the device is in breach. It also misleads the visitor into believing that refusing consent does not stop the processing, and that they must object separately.	Article 5(3) ePrivacy Directive; Taskforce report, Type H (paras 20-24), para 24: the legal basis for storing or reading information on a device cannot be legitimate interest
P6 Consent taken from inaction	The banner treats using the website, continuing or closing the banner as consent. None of these is a clear affirmative action.	Article 4(11) and recital 32 GDPR; Taskforce report, para 7
P7 Consent wall	The banner blocks the website and offers no way through other than accepting, or makes paying the price	Taskforce report, para 13

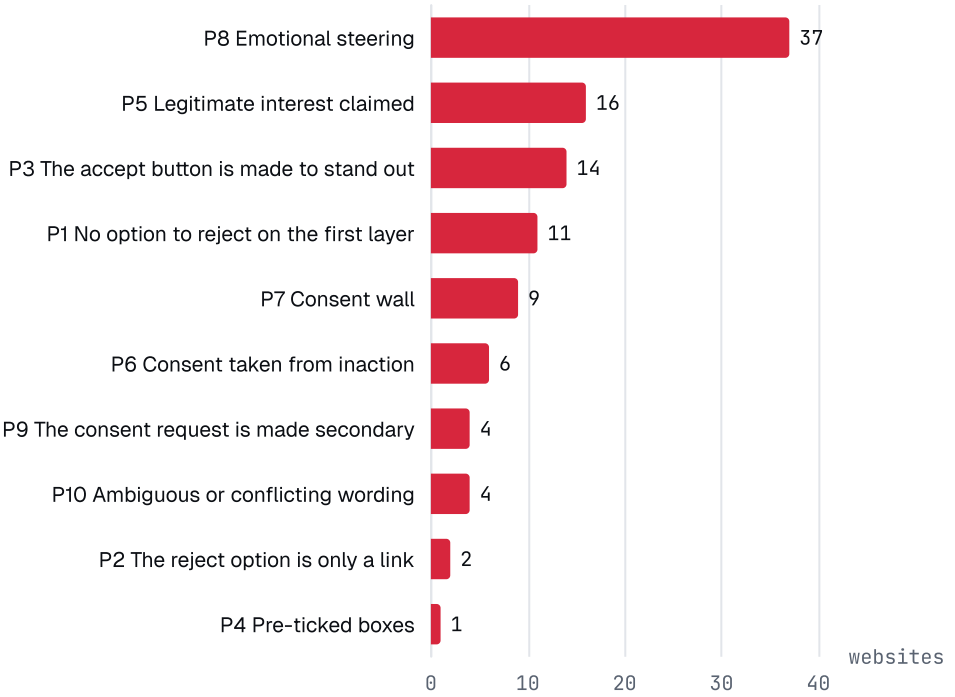


Figure 32. Websites whose consent banner used each dark pattern, of the 81 websites with a banner on at least one audited page.

HOW WE DETECTED THEM

Every consent banner on the audited pages was assessed from the capture engine's own screenshots, which are part of the signed record: the first screenshot of the screen taken after the banner appeared, and the engine's close-up of the banner where it took one. Nothing was clicked, so what was assessed is exactly what a visitor saw on arrival - the first layer of the banner. Each banner was checked against the ten patterns below. Two reviews were made independently and compared, and every difference between them was settled by examining the screenshot again.

Two things cannot be assessed without clicking into a banner - its further layers (including any pre-ticked boxes there) and the means of withdrawing consent once given - and they are not assessed here.

WHAT WE FOUND

A consent banner was in the page on 165 of the 265 audited pages, on 81 of the 100 websites. 115 of those 165 banners - 69.7% - used at least one dark pattern, on 59 of the 81 websites (72.8%); the worst used 5 at once.

The websites whose banners used the most patterns: stripchat.com, instagram.com, wired.com, cnbc.com, dailymotion.com, fandom.com, imgur.com, nbcnews.com, slideshare.net, tripadvisor.com. Each individual report sets out, banner by banner, what the screenshot shows, the banner's own wording and the screenshot itself (section 4E of each report).

THE LAW

Consent under Article 5(3) of the ePrivacy Directive has the meaning given in the GDPR: a freely given, specific, informed and unambiguous indication of the person's wishes, by a statement or by a clear affirmative action (Article 4(11) GDPR). Silence, pre-ticked boxes or inactivity do not constitute consent (recital 32 GDPR), and the Court of Justice confirmed in Planet49 (C-673/17) that a pre-ticked checkbox does not produce valid consent to cookies. The Cookie Banner Taskforce recalled that no cookie which requires consent can be set without it and that consent must be expressed by a positive action of the user (para 7), and that the legal basis for storing or reading information on a device can never be the legitimate interests of the controller (para 24).

The fairness principle in Article 5(1)(a) GDPR applies to every pattern: the EDPB treats it as the starting point for assessing deceptive design, and a deceptive design pattern fails it whatever else the controller has in place (Guidelines 03/2022, Annex I). A banner designed to steer the visitor towards accepting cannot collect consent that is freely given and unambiguous, and where no valid consent is obtained the processing that follows cannot comply with the GDPR (Taskforce report, para 24).

Pattern	Banners	Websites	Which websites
P1 No option to reject on the first layer	21	11	cnbc.com, cnn.com, fandom.com, imgur.com, instagram.com, mlb.com, nbcnews.com, prnewswire.com, stripchat.com, tripadvisor.com, wired.com
P2 The reject option is only a link	3	2	dailymotion.com, pexels.com
P3 The accept button is made to stand out	21	14	dailymotion.com, facebook.com, imdb.com, instagram.com, mdpi.com, pexels.com, reddit.com, slideshare.net, steamcommunity.com, threads.com, tumblr.com, twitter.com, unsplash.com, wordpress.com
P4 Pre-ticked boxes	1	1	slideshare.net
P5 Legitimate interest claimed	25	16	accuweather.com, biblegateway.com, businessinsider.com, cbsnews.com, dailymail.com, espn.com, hbr.org, imgur.com, people.com, slideshare.net, techcrunch.com, ted.com, tumblr.com, weather.com, wired.com, wsj.com
P6 Consent taken from inaction	12	6	cnbc.com, medium.com, mlb.com, nbcnews.com, stripchat.com, unsplash.com
P7 Consent wall	12	9	cnn.com, dailymail.com, fandom.com, imgur.com, instagram.com, stripchat.com, theguardian.com, tripadvisor.com, wired.com
P8 Emotional steering	78	37	accuweather.com, apnews.com, biblegateway.com, businessinsider.com, cnet.com, dailymotion.com, facebook.com, fandom.com, ft.com, independent.co.uk, instagram.com, issuu.com, kick.com, linkedin.com, mayoclinic.org, medium.com, npr.org, people.com, prnewswire.com, quora.com, reuters.com, soundcloud.com, stackoverflow.com, steamcommunity.com, theguardian.com, theverge.com, tripadvisor.com, tumblr.com, twitter.com, washingtonpost.com, wattpad.com, wired.com, wordpress.com, wsj.com, xhamster.com, yahoo.com, yelp.com
P9 The consent	5	4	deviantart.com, pornhub.com, stripchat.com, xhamster.com

5. Evidence

Each record is SHA-256 hash-chained, signed with ECDSA P-256 using a key held by a Secure Enclave, timestamped by FreeTSA under RFC 3161, and submitted to two OpenTimestamps calendars for anchoring in the Bitcoin blockchain (section 3.4).

Exhibits in this report

The exhibits in section 4 are a small selection of the evidence, chosen to show what the records contain. Each gives the page, the moment of the event to the millisecond, the browser interface or network mechanism involved, the call stack the engine captured as it happened, and the record it is held in with that record file's SHA-256 hash, so that it can be checked against the sealed evidence.

Measure	Figure
Sessions (one per page)	265
Records	3,088
Sessions whose chain verified complete and unbroken	265 of 265
Sessions with integrity flags	0

Exhibit	Website	Page	Technique	Section
E1	linkedin.com	homepage	`canvas.toDataURL`	fingerprinting
E2	linkedin.com	homepage	`webrtc.createOffer`	fingerprinting
E3	forbes.com	homepage	`battery.getBattery`	fingerprinting
E4	npr.org	privacy notice	`cookies.set`	cookies
E5	bbc.com	privacy notice	`document.cookie.write`	cookies
E6	cnn.com	homepage	`localStorage.setItem`	storage
E7	npr.org	privacy notice	`sendBeacon.sendBeacon`	beacons
E8	npr.org	privacy notice	`referer.crossOrigin`	referrer
E9	dailymail.com	homepage	`postMessage.wildcard`	messaging
E10	reddit.com	privacy notice	`form.inputRead`	listeners
E11	theguardian.com	privacy notice	`clientHints.request`	headers

6. Scale: the impact on people in the EU

The figures in section 4 describe what happened to one visitor, once, on the pages of each website. This section shows what that means at the scale at which these websites are used.

Under Article 24(2) of the Digital Services Act, online platforms must publish the average number of monthly active recipients of their service in the EU every six months, and the Commission designates any platform with at least 45 million as a very large online platform (Official Journal C/2026/5181, 6th October 2026).

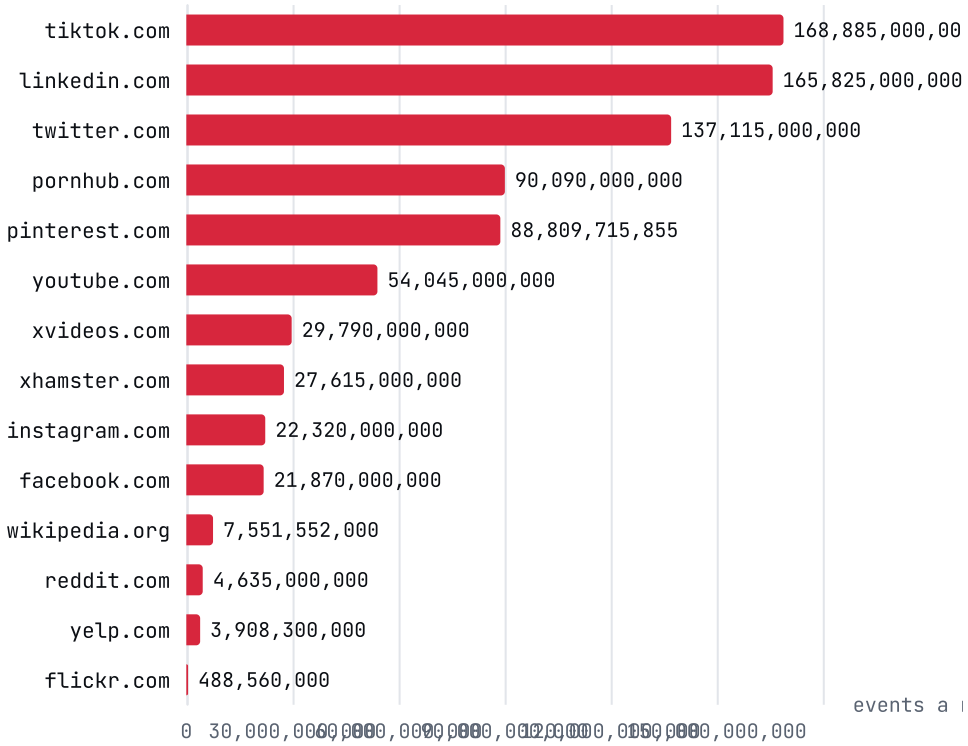


Figure 33. The minimum number of events that impact privacy each month on the devices of people in the EU, without any consent, assuming one homepage visit per monthly active recipient.

Website	EU monthly active recipients	Period	Events in one homepage visit	Minimum events a month
tiktok.com	at least 45,000,000	designated VLOP	3,753	168,885,000,000
linkedin.com	at least 45,000,000	designated VLOP	3,685	165,825,000,000
twitter.com	at least 45,000,000	designated VLOP	3,047	137,115,000,000
pornhub.com	at least 45,000,000	designated VLOP	2,002	90,090,000,000
pinterest.com	87,497,257	Jan-Jun 2026	1,015	88,809,715,855
youtube.com	at least 45,000,000	designated VLOP	1,201	54,045,000,000
xvideos.com	at least 45,000,000	designated VLOP	662	29,790,000,000
xhamster.com	35,000,000	6 months to 15 Aug 2026	789	27,615,000,000
instagram.com	at least 45,000,000	designated VLOP	496	22,320,000,000
facebook.com	at least 45,000,000	designated VLOP	486	21,870,000,000
wikipedia.org	157,324,000	Feb-Jul 2026	48	7,551,552,000
reddit.com	at least 45,000,000	designated VLOP	103	4,635,000,000
yelp.com	1,100,000	6 months to	3,553	3,908,300,000

Where a designated platform's own figure could not be retrieved, the legal threshold of 45 million is used as a minimum. A monthly active recipient has, by definition, used the service at least once in the month, so multiplying the recipients by **one homepage visit a month** gives the minimum number of events.

Figures given as "no more than" or "less than" are upper limits and are not used in any total; neither are figures from 2024 or figures published without a reporting period.

At least 822,948,127,855 events a month - 9,875,377,534,260 a year - take place on the devices of people in the EU on these platforms' homepages alone, without any consent. News, reference and academic publishers are generally not online platforms under the DSA and publish no comparable EU figure, so most of the 100 websites are not in this total.

6.1 Energy

The same arithmetic applies to energy. Every byte a page loads has to be carried across the network and processed on the device that receives it, so loading third-party content without any consent also uses energy on the visitor's device and across the network, for the benefit of the companies receiving the data. We took the bytes each homepage loaded from third parties during the visit and multiplied them by one homepage visit for each monthly active recipient, using the model in the WebSentinel environmental calculator: 0.06 kWh for every gigabyte transferred, the mid-point of the range in Parssinen and others, "Environmental impact assessment of online advertising" (2018), and 0.25 kg of CO2e for every kWh, the mid-point of the EU electricity grid average.

These are client-side figures - they cover the bytes delivered to the visitor's browser, carried across the network and processed on their device.

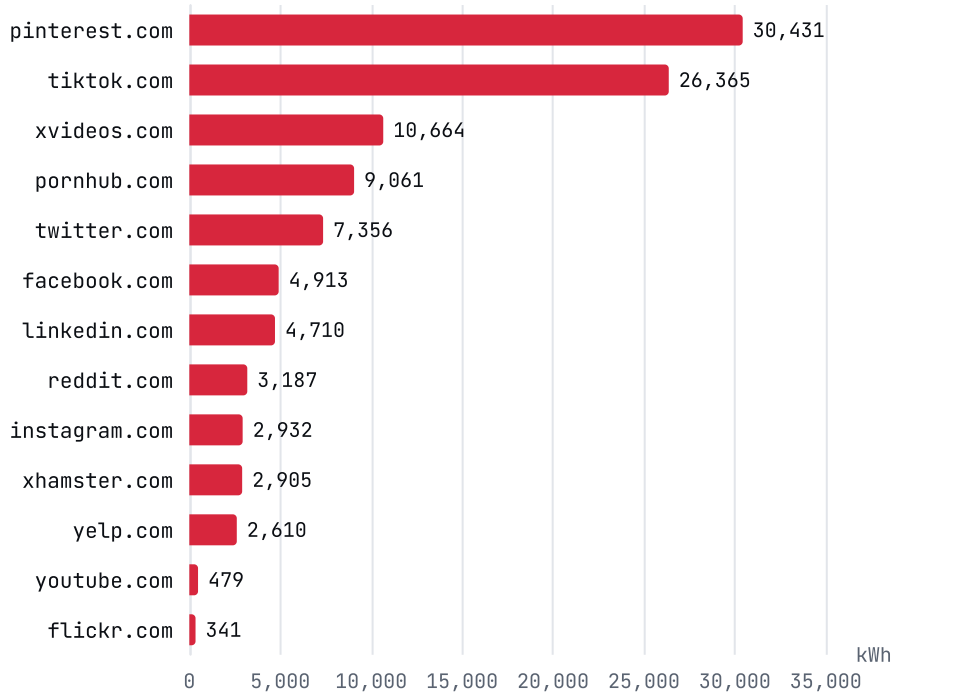


Figure 34. The minimum energy used each month to load third-party content without any consent on the homepages of these platforms for people in the EU, assuming one homepage visit per monthly active recipient.

Website	Third-party content in one homepage visit	kWh in 1 month	kWh in 1 year	kWh in 5 years	Tonnes of CO2e in 5 years
pinterest.com	5.80 MB	30,431	365,168	1,825,842	456
tiktok.com	9.76 MB	26,365	316,384	1,581,922	395
xvideos.com	3.95 MB	10,664	127,965	639,823	160
pornhub.com	3.36 MB	9,061	108,731	543,657	136
twitter.com	2.72 MB	7,356	88,269	441,347	110
facebook.com	1.82 MB	4,913	58,955	294,774	74
linkedin.com	1.74 MB	4,710	56,521	282,603	71
reddit.com	1.18 MB	3,187	38,246	191,230	48
instagram.com	1.09 MB	2,932	35,180	175,900	44
xhamster.com	1.38 MB	2,905	34,860	174,300	44
yelp.com	39.55 MB	2,610	31,325	156,625	39
youtube.com	0.18 MB	479	5,747	28,737	7
flickr.com	28.85 MB	341	4,092	20,460	5

They do not include the energy the data centres of these websites and of the third parties use to serve that content, or to receive, store and analyse the data collected, so the real figures are higher.

The Wikipedia homepage loaded nothing from third parties and is not in the chart. The same figures are left out as in the table of events above.

At least 105,954 kWh in 1 month, 1,271,444 kWh in 1 year and 6,357,219 kWh in 5 years - 26, 318 and 1,589 tonnes of CO2e - is used loading third-party content on these platforms' homepages alone, without any consent. That is one visit a month, to one page - every further page and every further visit adds to it.

To put that in context, the average home in the EU uses 3.6 MWh of electricity a year (Odyssee-Mure, electricity consumption per dwelling), so the energy used in a year is enough to power 353 homes for a whole year - and over five years, the electricity of 1,766 homes for a year. The greenhouse gas footprint of a person in the EU was 9.0 tonnes of CO2e in 2023 (Eurostat, 19th February 2026), so the emissions in one year match the entire annual footprint of 35 people, and over five years 177 - all of it for content that these websites loaded without any consent, on one page, for one visit a month.

7. Next steps

7.1 Disclosure

ADVANCE NOTICE

Before publishing, we gave every organisation whose website was audited the opportunity to discuss the findings for its own website - where the organisation had published an address through which it could be reached.

On 6th October 2026, a notice was sent by email to 102 privacy, data protection officer and EU or UK representative addresses, published by 75 of the 100 websites. Every address was taken from the website's own privacy or cookie notice, where the organisation itself names it as a contact for privacy and data protection matters. Addresses belonging to other organisations named in those notices - vendors' data protection officers, for example - were excluded.

The notice told each organisation that its website was one of the 100 audited, pointed to a public description of the audit, stated that the findings would be published in November and that legal complaints would be filed, and invited the organisation to discuss the findings for its website before then. No findings were included in the notice; they are provided only once a discussion begins.

ORGANISATIONS THAT COULD NOT BE NOTIFIED

25 websites publish no email address for privacy or data protection matters in their privacy or cookie notices, or publish no notice that we could find:

facebook.com, youtube.com, instagram.com, x.com, linkedin.com, tiktok.com, msn.com, wordpress.com, blogspot.com, archive.org, sciencedirect.com, weather.com, trustpilot.com, issuu.com, arxiv.org, goodreads.com, indiatimes.com, usatoday.com, tripadvisor.com, threads.com, xhamster.com, sagepub.com, apnews.com, discogs.com and hbr.org.

Several of them, including the largest platforms, direct privacy enquiries to web forms, and they were not notified. The GDPR allows a person to exercise their rights and to raise concerns with a controller by any reasonable means; a controller cannot require them to use a particular form, least of all one served on a website whose behaviour is the subject of the concern.

7.2 Recommendations

FOR SUPERVISORY AUTHORITIES

1. Count every event - each access to a person's device is a separate act, and enforcement that treats a website as one infringement, however many times and however many people it accesses, fails to reflect what the law prohibits or the scale at which it happens.
2. Enforce Article 5(3) beyond cookies - fingerprinting, browser storage, sideloaded scripts, beacons, ETags and client hints all store information on, or read information from, the user's device, as the EDPB set out in its Guidelines 2/2023, and the obligation to obtain consent applies to all of them.
3. Hold website operators to account for the code they load - under Fashion ID the operator is a joint controller for the collection and transmission carried out by third-party code on its pages, including code loaded by other code.
4. Examine CNAME aliasing, which runs third-party data collection under a website's own domain name, defeats the protections browsers apply to third parties and hides the recipient from the visitor.

5. Treat the notice pages as part of the processing - the pages that explain the processing carried out the same processing, including fingerprinting the people reading them.
6. Look at Chapter V in every case - almost every website sent requests to networks registered in the United States without any consent, and no transfer mechanism can make lawful a transfer of data that should never have been collected.
7. Refer conduct to the criminal authorities where it meets the definitions of illegal access, interception and data interference in national law transposing Directive 2013/40/EU.

FOR LEGISLATORS

1. Close the gap between the rules and their enforcement - Article 5(3) has required consent for every one of these techniques since it was amended in 2009, and the Article 29 Working Party and the EDPB have explained how it applies to them since 2012, so the rules are already in place and need to be enforced.
2. Make fingerprinting impossible to hide, by requiring browsers and platforms to expose, and allow users to block, the interfaces used for high-entropy fingerprinting.
3. Regulate access to on-device AI - websites can now probe and use AI models running on the visitor's own hardware, and that access should need consent in its own right, under the same rules as any other access to the device.
4. Fund regulators to audit at scale - every finding in this report was produced by automated, reproducible capture, and regulators can do the same across thousands of websites if they are equipped to.

FOR WEBSITE OPERATORS

1. Load nothing before consent that is not strictly necessary, including tag managers, advertising libraries, analytics, session recording, fingerprinting-based bot detection and fonts served by third parties that receive your visitors' page addresses.
2. Know what your pages load - keep an inventory of every script, including those loaded by other scripts, and remove what you cannot account for.
3. Remove CNAME aliases to third-party data collection.
4. Make your notices true - name the cookies you set, the techniques you use, including fingerprinting, and the companies that receive the data, and make sure the website does what the notice says.
5. Keep the notice pages clean, so that a privacy or cookie notice does not run the tracking it describes.

Appendix A. Event types recorded

The most frequent techniques in each category, with the number of websites on which each was recorded. The full list is in the data release.

Category	Technique	Events	Websites
beacons	`sendBeacon.sendBeacon`	683	50
beacons	`fetch.keepalive`	634	61
beacons	`link.decoration`	216	18
beacons	`img.trackingPixel`	99	20
cookies	`cookies.sent`	92,906	97
cookies	`document.cookie.read`	34,512	98
cookies	`cookies.set`	18,691	92
cookies	`document.cookie.write`	8,820	90
cookies	`cookie.sync`	23	11
fingerprinting	`timing.getEntriesByType`	13,641	92
fingerprinting	`document.visibilityState`	11,876	90
fingerprinting	`navigator.userAgent`	9,558	100
fingerprinting	`screen.width`	4,815	90
fingerprinting	`screen.height`	4,788	90
fingerprinting	`focus.hasFocus`	3,396	65
fingerprinting	`networkInfo.effectiveType`	3,302	68
fingerprinting	`navigator.platform`	1,888	73
fingerprinting	`navigator.maxTouchPoints`	1,787	68
fingerprinting	`clientHints.mobile`	1,692	55

Category	Technique	Events	Websites
fingerprinting	`screen.colorDepth`	1,524	57
fingerprinting	`screen.availWidth`	1,463	44
fingerprinting	`screen.availHeight`	1,456	49
fingerprinting	`cssMediaQueries.matchMedia`	1,449	94
fingerprinting	`screen.pixelDepth`	1,313	27
fingerprinting	`timezone.getTimezoneOffset`	1,265	77
fingerprinting	`intersectionObserver.IntersectionObserver`	1,155	89
fingerprinting	`visibility.listen.visibilitychange`	1,108	99
fingerprinting	`timezone.resolvedOptions`	1,045	52
fingerprinting	`navigator.language`	956	89
fingerprinting	`networkInfo.rtt`	929	54
fingerprinting	`mutationObserver.MutationObserver`	911	85
fingerprinting	`networkInfo.downlink`	903	52
fingerprinting	`visibility.listen.pagehide`	886	90
fingerprinting	`visibility.listen.pageshow`	610	90
headers	`clientHints.request`	52,124	100
headers	`etag.set`	45,955	99
headers	`altSvc.set`	20,103	96
headers	`hsts.set`	18,708	99

Category	Technique	Events	Websites
headers	worker.Worker	588	49
headers	resourceHint.preload	486	84
headers	resourceHint.dns-prefetch	308	45
headers	fetch.reject	298	62
headers	resourceHint.preconnect	153	70
headers	resourceHint.modulepreload	46	20
headers	resourceHint.prefetch	30	19
headers	fetch.error	30	12
headers	websocket.WebSocket	28	11
headers	speculationRules.insert	25	12
headers	xhr.reject	20	15
headers	sharedWorker.SharedWorker	17	4
headers	eventsource.EventSource	8	2
headers	etag.ifNoneMatch	1	1
listeners	form.inputRead	1,469	64
listeners	behavioural.scroll	466	93
listeners	behavioural.mousemove	394	80
listeners	behavioural.pointermove	274	46
listeners	behavioural.wheel	274	41
listeners	behavioural.touchmove	120	59
listeners	clipboardEvent.paste	81	30
listeners	clipboardEvent.copy	76	29
listeners	behavioural.mouseover	70	37

Category	Technique	Events	Websites
listeners	behavioural.mouseout	62	35
listeners	clipboardEvent.cut	62	23
listeners	behavioural.drop	32	18
listeners	behavioural.dragstart	30	18
listeners	behavioural.dragend	29	17
listeners	behavioural.drag	28	16
listeners	form.submit.programmatic	12	6
messaging	postMessage.wildcard	2,891	31
messaging	MessageChannel.construct	390	81
messaging	historyApi.replaceState	300	57
messaging	iframe.crossOrigin	236	68
messaging	broadcastChannelCovert.BroadcastChannel	231	22
messaging	postMessage.sameOrigin	126	7
messaging	csp.listen	81	34
messaging	postMessage.crossOrigin	60	20
messaging	webLocks.request	22	9
messaging	modalDialog.showModal	6	2
messaging	topics.browsingTopics	1	1
referrer	referer.crossOrigin	40,163	100
scripts	scripts.inserted	39,651	100
scripts	sideloaded (no insertion record)	1,690	0
storage	localStorage.getItem	16,700	97
storage	localStorage.setItem	11,982	92

Category	Technique	Events	Websites
storage	`sessionStorage.getItem`	9,207	93
storage	`indexedDB.getAll`	1,453	11
storage	`localStorage.removeItem`	1,387	90
storage	`sessionStorage.setItem`	947	82
storage	`windowName.read`	582	66
storage	`sessionStorage.removeItem`	522	72
storage	`indexedDB.open`	260	30
storage	`indexedDB.put`	231	19
storage	`broadcastChannels.BroadcastChannel`	231	22
storage	`indexedDB.get`	222	14
storage	`windowName.write`	203	14
storage	`indexedDB.delete`	148	11
storage	`indexedDB.count`	126	4
storage	`cacheStorage.match`	63	37
storage	`serviceWorkers.register`	50	23
storage	`storageAccess.hasStorageAccess`	35	21
storage	`storageAccess.requestStorageAccess`	25	17
storage	`sharedWorkers.SharedWorker`	17	4
storage	`serviceWorkers.getRegistration`	16	7
storage	`indexedDB.deleteDatabase`	15	11
storage	`serviceWorkers.getRegistrations`	15	9
storage	`indexedDB.openCursor`	14	6
storage	`cacheStorage.open`	10	4

Appendix B. Third-party cookie domains

Domain	Cookies set without any consent	Websites
cision.com	23	1
adobe.com	22	1
linkedin.com	16	4
thewaltdisneycompany.com	12	2
demdex.net	11	7
wikimedia.org	10	1
piano.io	10	8
amazon.de	10	1
doubleclick.net	10	6
tinypass.com	9	7
scribd.com	9	1
cxense.com	8	8
bing.com	8	2
youtube.com	7	1
elsevier.com	7	1
twitter.com	6	4
thomsonreuters.com	6	1
steampowered.com	6	1
everesttech.net	6	3
google.com	5	3

Appendix C. Receiving networks

Network operator	Servers	Websites
AMAZON-02	357	71
CLOUDFLARENET	264	74
FASTLY	188	62
GOOGLE-CLOUD-PLATFORM	97	55
AKAMAI-ASN1	85	36
AMAZON-AES	74	40
AKAMAI-AS	67	31
GOOGLE	65	80
MICROSOFT-CORP-MSN-AS-BLOCK	40	14
AUTOMATTIC	24	11
CDN77	19	8
FACEBOOK	9	11
OVH	9	9
REFLECTED	9	1
CLOUDFLARESPECTRUM	8	4
YAHOO-1	7	3
DIGITALOCEAN-ASN	7	3
HETZNER-AS	6	9
WIKIMEDIA	5	1
RUBICONPROJECT	5	9

Registry country	Servers
US	1247
NL	81
GB	43
FR	13
IE	11
DE	7
MT	2
FI	2
IL	2
JP	1
SG	1
LU	1
SE	1

Appendix D. Sideloaded scripts by serving domain

Domain	Sideloaded scripts
twimg.com	678
cookielaw.org	308
tiktokcdn-eu.com	284
fbcdn.net	272
adobedtm.com	220
xhcdn.com	204
googletagmanager.com	203
chapturist.com	192
nbcnews.com	191
disney.com	188
gstatic.com	148
yelpcdn.com	146
guim.co.uk	133
twitch.tv	130
staticflickr.com	128
bbci.co.uk	121
substackcdn.com	112
behance.net	106
foxnews.com	105
typekit.net	100

Appendix E. Results by website

#	Website	Type	Events	Tier 1	Tier 2	Tier 3	High or critical	Third-party servers
1	espn.com	sport	57,194	469	15,440	41,285	150	21
2	go.com	news	42,253	170	12,664	29,419	45	18
3	nytimes.com	news	30,158	87	24,972	5,099	177	27
4	telegraph.co.uk	news	19,770	250	13,722	5,798	2,036	44
5	kick.com	video	16,807	256	8,021	8,530	270	22
6	foxnews.com	news	15,475	38	14,148	1,289	395	20
7	mlb.com	sport	15,088	162	11,247	3,679	248	78
8	msn.com	portal	11,956	167	7,782	4,007	250	23
9	imdb.com	reference	11,711	139	9,639	1,933	240	15
10	nbcnews.com	news	11,481	901	7,400	3,180	1,760	82
11	businessinsider.com	news	10,820	76	9,919	825	122	34
12	quora.com	forum	10,281	4,049	3,311	2,921	3,426	8
13	washingtonpost.com	news	9,409	114	6,805	2,490	238	44
14	cnbc.com	news	8,193	112	6,309	1,772	1,162	36
15	stripchat.com	adult	7,867	43	5,095	2,729	1,188	8
16	npr.org	news	7,808	1,347	4,115	2,346	1,478	33
17	vimeo.com	video	7,691	91	2,958	4,642	98	17
18	independent.co.uk	news	7,432	39	4,186	3,207	185	26
19	wsj.com	news	7,200	25	5,629	1,546	868	33
20	behance.net	images	6,928	174	4,450	2,304	233	30
21	pornhub.com	adult	6,805	13	5,904	888	55	12
22	issuu.com	publishing	6,561	74	3,794	2,693	37	8

#	Website	Type	Events	Tier 1	Tier 2	Tier 3	High or critical	Third-party servers
23	apnews.com	news	6,281	100	2,971	3,210	141	52
24	hbr.org	news	6,236	112	4,725	1,399	944	56
25	prnewswire.com	news	6,171	137	3,875	2,159	562	66
26	time.com	news	6,061	96	4,373	1,592	799	55
27	soundcloud.com	audio	5,993	95	3,014	2,884	2,002	29
28	linkedin.com	social	5,325	233	3,941	1,151	237	14
29	cnn.com	news	5,157	39	3,676	1,442	814	65
30	techcrunch.com	news	5,146	177	3,357	1,612	304	78
31	tiktok.com	video	4,811	100	2,328	2,383	259	16
32	dailymail.com	news	4,766	516	2,685	1,565	640	26
33	bbc.com	news	4,527	28	2,703	1,796	768	28
34	medium.com	blogging	4,517	119	1,162	3,236	97	6
35	pexels.com	images	3,851	33	2,927	891	53	12
36	twitter.com	social	3,819	83	1,268	2,468	65	8
37	theguardian.com	news	3,700	55	1,083	2,562	413	11
38	yelp.com	reviews	3,697	28	1,215	2,454	113	7
39	ft.com	news	3,692	21	2,547	1,124	49	8
40	flickr.com	images	3,558	1,015	1,288	1,255	870	21
41	aol.com	portal	3,509	19	1,807	1,683	73	8
42	cbsnews.com	news	3,358	21	2,014	1,323	33	28
43	yahoo.com	portal	3,299	30	1,890	1,379	91	8
44	wired.com	news	3,236	21	1,541	1,674	40	11
45	twitch.tv	video	3,183	34	1,341	1,808	73	13
46	stackoverflow.com	forum	3,155	48	1,077	2,030	140	28
47	tandfonline.com	academic	2,983	28	538	2,417	81	21
48	imgur.com	images	2,860	29	1,425	1,406	81	36

#	Website	Type	Events	Tier 1	Tier 2	Tier 3	High or critical	Third-party servers
49	theverge.com	news	2,766	53	1,094	1,619	50	37
50	sagepub.com	academic	2,753	35	1,703	1,015	247	20
51	bloomberg.com	news	2,637	40	1,527	1,070	150	15
52	mayoclinic.org	health	2,586	68	1,726	792	148	10
53	wikipedia.org	reference	2,554	3	1,701	850	13	5
54	ted.com	video	2,401	27	1,328	1,046	67	13
55	merriam-webster.com	reference	2,378	488	983	907	545	32
56	mdpi.com	academic	2,206	20	485	1,701	30	11
57	youtube.com	video	2,203	44	1,028	1,131	290	14
58	sciencedirect.com	academic	2,134	22	1,194	918	37	14
59	slideshare.net	publishing	2,128	17	1,451	660	179	21
60	wattpad.com	publishing	2,096	43	1,455	598	47	14
61	threads.com	social	2,044	41	894	1,109	150	5
62	wiley.com	academic	1,991	25	1,314	652	20	6
63	reuters.com	news	1,967	65	1,155	747	141	18
64	fandom.com	reference	1,962	13	1,172	777	24	11
65	xhamster.com	adult	1,934	36	643	1,255	597	11
66	tumblr.com	blogging	1,908	17	423	1,468	21	3
67	reddit.com	forum	1,886	59	650	1,177	37	4
68	substack.com	blogging	1,873	9	885	979	14	8
69	pinterest.com	social	1,765	42	806	917	99	14
70	facebook.com	social	1,660	48	788	824	218	5
71	steamcommunity.com	forum	1,655	844	273	538	6	7
72	usatoday.com	news	1,641	5	512	1,124	10	11
73	pixabay.com	images	1,592	15	518	1,059	7	7
74	discogs.com	reference	1,582	10	442	1,130	14	7

#	Website	Type	Events	Tier 1	Tier 2	Tier 3	High or critical	Third-party servers
75	cambridge.org	academic	1,525	12	882	631	60	7
76	instagram.com	social	1,389	40	826	523	118	5
77	forbes.com	news	1,338	49	392	897	40	17
78	cnet.com	news	1,337	17	560	760	38	29
79	britannica.com	reference	1,328	17	926	385	22	22
80	frontiersin.org	academic	1,210	28	358	824	52	12
81	trustpilot.com	reviews	1,210	77	269	864	52	14
82	people.com	news	1,094	34	426	634	43	8
83	dailymotion.com	video	1,025	9	424	592	39	13
84	accuweather.com	weather	962	15	548	399	215	17
85	weather.com	weather	944	59	544	341	148	8
86	nature.com	academic	883	4	391	488	12	3
87	oup.com	academic	824	13	282	529	30	7
88	wordpress.com	blogging	822	8	388	426	12	6
89	springer.com	academic	802	6	345	451	12	6
90	tripadvisor.com	reviews	793	32	484	277	61	8
91	unsplash.com	images	740	17	358	365	31	0
92	xvideos.com	adult	730	8	482	240	23	7
93	blogspot.com	blogging	672	2	226	444	12	6
94	deviantart.com	images	670	38	151	481	20	5
95	bsky.app	social	576	28	96	452	30	3
96	biblegateway.com	reference	498	6	206	286	4	2
97	goodreads.com	reviews	437	23	174	240	25	7
98	indiatimes.com	news	375	23	149	203	23	21
99	archive.org	reference	305	2	65	238	4	0
100	arxiv.org	academic	200	6	80	114	5	3

Appendix F. Notices not measured for length and readability

The captured page did not hold the full text of these notices - for example because the text is loaded after the page or held in collapsed sections - so they are not part of the length and readability figures in section 4A. Fewer than 150 words of notice text were captured from each of them.

Website	Notice	Words captured
flickr.com	cookie notice	0
frontiersin.org	cookie notice	52
ft.com	cookie notice	0
hbr.org	cookie notice	7
issuu.com	cookie notice	0
kick.com	cookie notice	0
medium.com	cookie notice	7
quora.com	cookie notice	0
sagepub.com	cookie notice	0
time.com	cookie notice	0
trustpilot.com	cookie notice	0
washingtonpost.com	cookie notice	0
youtube.com	cookie notice	20