

The Empire Strikes Back (at Your Privacy): An Archaeology of Tracking on Government Websites

Sachin Kumar Singh
University of Utah
Salt Lake City, Utah, USA

Robert Ricci
University of Utah
Salt Lake City, Utah, USA

Faisal Mahmud
New York University Abu Dhabi
Abu Dhabi, UAE

Sandra Siby
New York University Abu Dhabi
Abu Dhabi, UAE

Abstract

Citizens rely on government websites for a wide array of essential services. However, these websites may embed third-party trackers, raising questions on privacy, security, and data sovereignty. We conduct a large-scale longitudinal study of tracker adoption and evolution on government websites worldwide. Our study spans 61 countries and nearly three decades (1996-2025), using historical snapshots from the Internet Archive. We find that tracking has shifted from rare to routine: by 2025, third-party trackers appear on $\approx 50\%$ of studied government websites, with growth overwhelmingly driven by external (third-party) services. This tracking is dominated by a few large US-based organizations, with a long tail of lesser-known players. We also find great heterogeneity in tracker adoption and presence across regions, indicating that users around the world have differing exposure to privacy and security issues when accessing essential services. Our findings highlight the need for approaches to improve privacy and data sovereignty on public-service platforms.

Keywords

Government Websites, Web Tracking

1 Introduction

When citizens access government websites for various services—such as applying for a passport, filing taxes, or obtaining healthcare information, they generally expect privacy [74]. Unlike commercial platforms, government service portals are often the only available option to access essential public services, leaving individuals with no ability to opt out. While it may seem reasonable to assume these portals would prioritize privacy, in practice, many government websites incorporate tracking technologies [34, 68, 78].

Trackers gather and analyze user data, including browsing behavior [28, 66], unique identifiers [5, 73], and digital fingerprints [10, 27, 57], to build detailed user profiles for commercial and analytical purposes, raising significant privacy concerns when the data is sensitive [44, 48] or subject to regulations [15, 29, 30]. Trackers are heavily used across the modern web as part of a largely invisible tracking infrastructure, often without users' explicit awareness,

limiting meaningful control and consent [51]—especially on government sites where users have no alternative and where foreign third-party can introduce cross-border tracker-data flows [41, 71] that raise questions about sovereignty and surveillance.

Despite the sensitivity of government websites and the rapid expansion of national governments' digital presence [77]—bringing more citizens into routine online interactions with public institutions—few studies have examined the use of tracking technologies on these sites or how these practices evolve over time. Most existing research on web tracking focuses on commercial websites [28, 66] or on smaller, targeted sets of government websites [34, 71, 78]. Further, prior work is often not longitudinal [34, 68, 71, 83] and frequently treats the web as geographically uniform rather than tracing country-specific trajectories over time [28, 46].

Understanding the use of tracking technologies on government websites at the country level provides a valuable view for linking deployment patterns to national policies, governance models, and technological ecosystems. Examining these practices over time also helps characterize regional differences in tracking exposure, since government websites primarily serve domestic audiences. In this work, we present the first large-scale historical study of tracking on government websites. We use historical snapshots from the Internet Archive [40] to identify and track the presence of trackers on government websites, allowing us to reconstruct a timeline of tracking growth and distribution across regions. With our study, we answer the following research questions:

- **RQ1 — Longitudinal Trajectory.** How has tracking on government websites evolved over time?
- **RQ2 — Geographic Heterogeneity.** How do countries/regions differ in adoption levels, timing (surges, plateaus, reversals)?
- **RQ3 — Tracking Actors.** Which organizations (parent companies) dominate government website tracking, and how concentrated is ownership?
- **RQ4 — Inferring Adoption From Internet Archives.** Can we turn irregular archives of government websites into credible tracker adoption events on these websites?

Our study provides the first global, longitudinal view of trackers on government websites across 61 countries (1996–2025). **(RQ1)** We find that tracking has shifted from rare to routine: by 2025, 50% of sites embed trackers, driven overwhelmingly by third parties. **(RQ2)** Adoption is highly uneven, with very high 2025 levels in Singapore/Australia/Bangladesh ($\approx 80\text{--}95\%$) and persistently low levels in China/Germany/France ($\approx 6\text{--}12\%$). **(RQ3)** Ownership is

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2026(2), 1–19
© 2026 Copyright held by the owner/author(s).
<https://doi.org/XXXXXXX.XXXXXXX>

concentrated and dominated by US-based corporations (Google and Meta). **(RQ4)** Finally, we show that irregular web archives can credibly yield adoption events using a *First-Observed Uptake* (FOU) measure.

We first review prior work on government tracking and archive-based web studies (Section 2), then describe our methods (Section 3). We next present results on tracker prevalence and dynamics over time (Section 4), tracker ownership and concentration (Section 5), and a First-Observed Uptake (FOU) approach for inferring adoption events from irregular archives (Section 6).

2 Background and Related Work

Trackers and Government Services: Foundational work on digital government highlighted enduring tensions between efficient service delivery and the protection of citizens’ privacy, and showed that trust in government and its surveillance intentions strongly shapes citizens’ privacy concerns about online tracking [20, 43]. Subsequent empirical studies reveal that tracking on government digital services is widespread [34, 68, 78, 83]. An analysis of Dutch government websites in 2012 found third-party trackers on roughly 60% of sites and documented centralized deployment patterns, such as shared analytics across hundreds of embassy and consulate domains [78]. A large-scale study of 150,244 government websites in 206 countries and 1,166 Android apps from 71 countries reported that 17% of websites and 37% of apps embedded Google trackers, and 27% of apps leaked sensitive information to third parties, including in jurisdictions with stringent privacy laws such as the EU and California [68]. Another analysis of more than 5,500 government, international-organization, and COVID-19 information websites in G20 countries found extensive use of third-party tracking cookies even in GDPR jurisdictions; for example, approximately 55%, 62%, and 90% of government sites in France, China, and Russia contained persistent third-party trackers, and official sites in 77–100% of G20 countries set cookies without user consent [34]. Work focusing on health services in the United States found that government health websites still include advertising and tracking technologies, though at lower levels than commercial health sites (2.11 vs. 15.84 ad/trackers per site on average) [83]. In this ecosystem, tag management systems play a central role by enabling site operators to embed and control large numbers of third-party scripts through a single interface. For example, Google Tag Manager appears on about 52% of the top million websites and can collect personal data via “cookieless pings” and trigger additional, undisclosed third-party requests even when users appear to opt out [52].

Consent, Choice, and Dark Patterns: Although consent banners and privacy controls aim to give users control over tracking, meaningful choice is often elusive. Users rarely have the time or expertise to read policies or assess how settings map onto concrete risks [6, 18, 50]. Studies of cookie banners show that pre-selected options, ambiguous wording, and higher effort to refuse than to accept systematically slide users toward tracking [19, 59], with interface asymmetries such as extra clicks to refuse, reducing opt-out rates from 17% to 4% [12].

Re-identifiability risks: Tracking data also raises substantial re-identifiability risks. Browser fingerprints and combinations of device and network attributes are often highly distinctive, enabling

persistent recognition even without cookies [10, 27, 57]. Large-scale third-party trackers can aggregate these identifiers with long-term browsing histories across sites [28, 66]. Even pseudonymized datasets can frequently be linked back to individuals or small groups and used to infer sensitive traits when combined with auxiliary information [44, 48]. On government websites, where page views often correspond to specific services, such re-identification and inference directly implicate sensitive categories of personal data.

The Wayback Machine/Internet Archive (IA): The Internet Archive’s (IA) Wayback Machine [40] is one of the largest publicly available web archives, providing historical snapshots of websites dating back to 1996 (<https://archive.org>). It continuously crawls and preserves web content, creating snapshots that include not only the HTML pages but also embedded resources such as scripts, images, and style files, enabling users to reconstruct the historical state of websites. While there are other archival data resources, e.g., commoncrawl.org, httparchive.org, and Memento Time Travel, previous research has shown that the Internet Archive is the largest such dataset and using other archive data with it provides only a marginal advantage [10, 37]; thus we use only IA. While coverage is not uniform across countries or domains and the frequency of captures varies, it remains the most comprehensive longitudinal dataset of the public web [37].

Studies based on Internet Archive data: The use of Internet Archive data to study web tracking was pioneered by work that analyzed the 500 most popular websites from 1996–2016 using Wayback Machine snapshots, leveraging preserved JavaScript to retrospectively measure third-party requests, and documenting a sharp rise in such requests over time, with Google Analytics appearing on nearly one-third of sites by 2016 [47]. This established both the growth of the surveillance ecosystem and a methodological basis for archive-based tracking studies. Subsequent research applied Internet Archive data to government domains, showing, for example, substantial content changes on U.S. federal agency sites between 2016 and 2020, including a 38% reduction in the term “climate change” and loss of access to about 20% of EPA webpages [58]. Other work used Wayback snapshots to detect browser fingerprinting, revealing API abuse years before public disclosure and demonstrating that archives capture both static, filter-list-detectable trackers and stateless tracking techniques, thereby complementing conventional filter-list-based measurement [10].

Our Study: Unlike studies that focus on commercial websites [47] or a single country, our work uniquely centers on the government website across multiple countries, where issues of transparency, sovereignty, and citizen privacy are especially critical. Our study provides the first large-scale, longitudinal analysis of trackers on the government websites across 61 countries. By using the Internet Archive’s historical data, we are able to examine the global presence of trackers and the regional differences in adoption over nearly three decades (1996–2025). We also want to note that our observations are scoped strictly to government websites in 61 countries and we do not claim that trackers or their growth that we document are unique to government websites, as tracking is widely deployed across the broader web [28, 32, 47, 71].




Figure 1: World map showing countries studied. Colored countries indicate those included in the dataset.

3 Method

We build a system to conduct a large-scale longitudinal study of web trackers on government websites. Our system downloads and structures Internet Archive snapshots from 1996-2025 (the full available timeframe) for government domains across 61 countries. It extracts all the URLs embedded in a webpage—including web links, API endpoints, and resource links such as JavaScript, CSS and image files. We then use these URLs to identify tracking-related instances and build a regional, longitudinal picture of tracking. While our focus is on government websites, our system can be extended to other contexts, and it is available at [75] to support future research. We discuss the details of our methodology in this section.

3.1 Data Collection

Government Website List: We start with a list of government websites. Prior studies [34, 42] have relied on top-level domains (TLDs) to identify government websites. For a more extensive and up-to-date list, we use Kumar et al.’s IMC 2024 study [45], which constructed a dataset of government websites by using multiple approaches, including the use of TLDs and deeper scraping of government websites. Since our dataset is drawn from this work, we also adopt their definition of what constitutes a government website. Particularly, they focus on “*federal-level (or equivalent) resources, including various segments of the federal administration (e.g., the presidency, ministries, and secretaries), federal agencies, often referred to as decentralized agencies (e.g., the US National Science Foundation and the US Internal Revenue Service) and state-owned enterprises. To consider State-Owned Enterprises (SOEs), we follow the International Monetary Fund (IMF) guidelines and only include companies where the federal government holds more than 50% of the shares*” [45].

Their dataset covers 61 countries distributed across all continents: Africa (5), Asia (18), Europe (26), North America (4), Oceania (2), and South America (6). Figure 1 shows these countries on the world map, with the studied 61 countries highlighted in red. In total, the dataset comprises approximately 18.9K government websites, with an average of 310 websites per country. The distribution is skewed, with a median of 174, a standard deviation of 447.3, a 25th percentile of 104, and a 75th percentile of 303. In the dataset only two countries have fewer than 50 websites (Kazakhstan with 17 and Uruguay with 35). Twelve countries fall between 50 and 100 websites, while the majority (37 countries) lie between 100 and 500. Seven countries have between 500 and 1000 websites, and only

three countries exceed 1000 — South Korea (1184), the Netherlands (1413), and the United States (3021).

Sampling Bias: Our method also inherits limitations from Kumar et al. [45]: it aggregates *self-reported* government sites from heterogeneous registries, focuses on public-facing pages (excluding authenticated portals and many state-owned enterprises), and prioritizes federal-level domains. As a result, subnational and operational services may be omitted, and decentralized countries with many regional/municipal sites are undersampled—so our findings should be interpreted as conservative lower bounds.

Data from the Internet Archive: After obtaining the list of government websites from Kumar et al.’s work [45], we download the snapshot records for all websites to determine how many had at least one archived snapshot in the Internet Archive between 1996 (earliest possible) and 2025 [June 2025]. In total, this amounts to 16.5K archived websites (out of 18.9K). On average, countries had 270 archived websites, with a median of 151. The distribution shows the expected variation, with a standard deviation of 385, a 25th percentile of 90, and a 75th percentile of 261. At the extremes, Kazakhstan (KZ) has the fewest archived websites (16), while the United States (US) has the most (2597). Appendix 9.2 shows a bar plot of the number of government websites archived on the Internet Archive. In terms of percentage coverage *i.e.*, the proportion of websites from the dataset that have at least one archived snapshot, the average coverage across countries is 88%. The lowest coverage is observed in Russia at 75%, and Uruguay achieves full coverage of 100%. These results indicate that the Internet Archive has strong coverage of government websites overall, providing a reliable source of data for analyses of web tracking on government domains.

We download website homepage snapshots of all the identified websites from the Internet Archive. For each website, we focus on the HTML of the homepage, as this serves as the primary entry point for visitors and is the most consistent component across time. We acknowledge that this scope can miss trackers that appear only on deeper subpages. This trade-off is common in archive-based longitudinal work [10, 47]. For each website, we select two snapshots per year: the first available snapshot of the year (to capture the earliest state of the website) and one additional snapshot chosen at random from the same year (to account for potential variations or updates during the year). We analyze the collected data at yearly granularity. For example, a website is marked *no tracker* in 2018 only if none of its 2018 snapshots contain a tracker; if any snapshot does, we mark that it has *tracker* for 2018.

We then process all the website snapshots and extract embedded URLs from the HTML content. Since the Internet Archive rewrites URLs to point to cached versions (e.g., <https://web.archive.org/web/20220101000000/https://a.com/script.js>), we clean these URLs to recover their original form (e.g., <https://a.com/script.js>). This step was essential for accurately identifying original resources and detecting embedded advertising/tracking domains. In total, we extracted 7.3 million URLs, of which 5.6 million were unique, across the 16.5K government websites from 61 countries.

Best Practices: Given the large number of websites and snapshots to be downloaded, we ensured that our data collection process was respectful to the Internet Archive’s infrastructure. Specifically,

we avoided making multiple simultaneous requests and introduced a delay of 3 to 5 seconds between consecutive download attempts.

3.2 Tracker Identification

To determine whether the embedded URLs are trackers, we use the EasyList community maintained filters lists [23–26][June 2025] i.e. EasyList Standard (including both EasyList [23] and EasyPrivacy [25]) and the regional filters (Affiliated Filter Lists [26]), which are subsets focused on specific contexts, e.g., ABPindo (Indonesia), Liste FR (France), and ROList (Romania). These lists have been used in previous studies [4, 7, 14, 16, 32, 72, 81, 82].

EasyList (the standard list) is designed for ad blockers and targets global tracker sites [24], while regional/affiliated lists focus on specific languages and regions. Additionally, to evaluate the impact of regional or affiliated filters on URLs from the Wayback Machine, we first check that these lists are used as add-ons to EasyList Standard rather than standalone filters. When we compared all the rules from 21 regional/affiliated filter lists to the EasyList standard list and we found very small overlap. Out of 113K rules across the regional lists, only 712 were present in the EasyList standard list which is just 0.006%. This indicates that regional lists can block trackers that the standard list misses and because they are tailored to specific contexts, they should be used as add-ons to the standard list, not as standalone lists. Additionally, prior research has shown that regional filter lists underperform in their intended contexts and combined lists can improve blocking effectiveness [14]. Thus, we use all available filter lists to identify tracking-related URLs. If any list flags a URL, we mark it as a tracker.

In our dataset, the majority of URLs were flagged by EasyPrivacy (19K), followed by EasyList (2K). Within the regional lists, only Liste FR (French, 1.9K) and the Nordic list (3.4K) identified more than 1K URLs as trackers, while 13 of the regional lists flagged fewer than 100 URLs each. This suggests the EasyList standard list captures most tracking activity, while regional lists add only marginal coverage on government websites, likely because regional trackers are less common or regional lists have narrower coverage.

3.2.1 Validation of EasyList Snapshot Choice. A natural concern is whether we should use historical EasyList versions rather than a single recent snapshot. Prior work finds that the most important rules are stable over many years, and that most list growth comes from cosmetic or rarely used entries [72]. We therefore label trackers using the June 2025 EasyList and EasyPrivacy snapshot for all years, keeping the labeling rule fixed and avoiding confounding from blocklist changes. To empirically validate this design choice, we retrieved historical snapshots of EasyList and EasyPrivacy from the GitHub repository [63] and the Internet Archive for three anchor years (2006—the oldest archived version we could find, 2010, and 2020). For each anchor year y , we drew a simple random sample of 10,000 embedded URLs from our dataset for that year and applied both (i) the EasyList and EasyPrivacy versions from year y (EasyPrivacy is available via the archive only for 2020) and (ii) the June 2025 snapshot of EasyList and EasyPrivacy.

For URLs detected as trackers in our anchor-year validation, we manually inspect them to assess whether they are true trackers and to compare the anchor-year lists with the June 2025 list. URLs that clearly serve only content (e.g., background images or paths

like /advertisement-for-vacancies/) but are flagged as trackers are treated as false positives. Because resource purpose is often ambiguous, we apply a conservative rule: unless we find concrete evidence that a detected URL is *not* used for tracking, we keep its tracker label. This is especially important for JavaScript files and API endpoints, where usage (e.g., analytics or tracking) is hard to infer from the path alone. Where useful, we also consult the Ghostery WhoTracksMe portal [33] to identify known tracker domains.

For data from 2006, 2010 and 2020, after labeling every URL that any of the lists identified as a tracker as either tracker or non-tracker. On a random sample from 2006, the 2006 EasyList flagged 67 URLs and the June 2025 snapshot flagged 50, with 6 in common. We manually inspected all 111 distinct URLs: the 2006 list correctly identified 13 trackers but produced 54 false positives, whereas the 2025 list correctly identified 49 trackers with a single false positive, corresponding to false-positive rates of 0.54% and 0.01% on the 10,000-URL sample and about 3.8× more true trackers for the 2025 list. For 2010 data, the 2010 list detected 18 trackers (4 false positives) and the 2025 list detected 26 (1 false positive), with 6 common detections. For 2020 data, the 2020 list flagged 54 URLs (15 false positives) and the 2025 list flagged 41, all genuine trackers, with 25 URLs in common.

Taken together, these results justify our choice to apply a single June 2025 snapshot to all years in our dataset. The empirical comparison shows that the modern list has higher precision and at least comparable ability to detect trackers on old data, using it uniformly across time is conservative: it is more likely to undercount trackers than to overcount them. Additionally, our EasyList/EasyPrivacy validation aligns with and extends prior list evaluations. Englehardt & Narayanan [28] emphasize that blocklists are incomplete—particularly due to false negatives and limited long-tail coverage—so EasyList-based estimates should be interpreted as conservative lower bounds of tracking prevalence. Consequently, our longitudinal curves should be interpreted as a lower bound on the historical prevalence of tracking.

3.2.2 First-party and Third-party Trackers: To distinguish between first-party and third-party trackers, we classify each identified tracker based on its effective top-level domain plus one label (eTLD+1). Specifically, if the domain of the tracker matched the eTLD+1 of the government website it is embedded in, it is labeled as first-party tracker, if not, then it is labeled as a third-party tracker. The classification is domain-based and does not infer ownership, so a URL embedded from another government agency with a different eTLD+1 is counted as third-party.

3.3 Organization Mapping

Mapping identified trackers to their parent organizations is an important part of our analysis. Multiple trackers can be owned and managed by large corporations or advertising networks that control extensive tracking infrastructures across the web. For example, Google operates a wide range of tracker domains, including *google-analytics.com*, *doubleclick.net*, *googletagmanager.com*, *googlesyndication.com*, and others. Together, these services allow Google to collect user data across websites, showing how a single organization can dominate tracking by using multiple domains. By linking tracker domains to their parent organizations, we move from a




Figure 2: Presence of trackers on government websites across all studied countries over time.

view of individual domains to a broader view of organizational influence. This allows us to assess which organizations dominate the tracking space on government websites and how their influence has changed over time, how concentrated tracker ownership is, and whether certain organizations disproportionately control user data across different countries. To perform this mapping, we rely on the DuckDuckGo Tracker Radar dataset [22], which maintains a repository of tracker domains and their associated parent organizations. We match each identified tracker domain (eTLD+1) from our dataset to the Tracker Radar data. When a domain appears in the Tracker Radar dataset, we assign it to the corresponding parent organization. This ensures that multiple subdomains or specific URLs (e.g., *doubleclick.net* or *googleanalytics.com*) are consistently grouped under their parent company (e.g., Google). We identified 1098 tracker domains (eTLD+1). Using Tracker Radar, we map 365 of these tracker domains to 265 organizations ($\approx 33\%$ coverage).

4 Evolution of Trackers Over Time

Our first set of research questions interrogates how the presence of trackers on government websites has evolved over the thirty-year period of our study. We quantify the presence of trackers by calculating, for each country-year (1830), the percentage of archived websites containing trackers hosted by the site itself (first-party), by other sites (third-party), and *any* tracker (first- or third-party). (In Section 5 we take a deeper look at the nature of first- and third-party trackers and their operators.) We exclude country-year pairs when no data is available in the archive (39).

How has the presence of trackers on government websites evolved globally over the past three decades? We measure tracker presence via the percentage of websites containing at least one tracker: these percentages are weighted by the number of websites archived in that country-year; this enables the identification of broader temporal patterns and shifts in tracker presence across the observed period. Figure 2 shows the global presence of trackers on government websites from 1996 to 2025. For most of the observed period, first-party tracker use remains consistently low, starting near zero in 1996 and only gradually increasing to around 4.4% by 2025. In contrast, third-party tracker and overall tracker presence follow a similar and much steeper trajectory. Both rise modestly through the late 1990s and early 2000s, then experience a sharp increase between 2005 and 2008, peaking at around 15% before declining slightly during the early 2010s. From 2016 onward, we observe a sustained and dramatic rise in both third-party and overall tracker presence, with values almost tripling in less than a decade.

Between 2023 and 2025, growth levels off and the curve flattens. By 2025, third-party trackers are present on nearly half (48%) of all observed government websites, and tracker presence slightly exceeds 50%. Previous studies [68, 78] observed known trackers on 29.9% of government websites’ landing pages and 49% of EU government websites. The consistently wide gap between the first-party and third-party in Figure 2 indicates that the global growth in tracking is overwhelmingly driven by external, third-party services rather than trackers operated by the government websites themselves.

RQ1: Tracking on government websites has shifted from rare to routine over the past three decades. First-party use remains marginal (4.4% in 2025), while third-party services drive nearly all growth, pushing overall presence slightly above 50% by 2025, following a steep post-2016 surge.

Do individual regions and countries show patterns distinct from global trends? To answer this question, we dig deeper into country-level data for 61 countries. Based on 1830 country-year pairs, an average of 2.7% (standard deviation = 14.5) of websites contained first-party trackers, 24.8% (SD = 23.5) contained third-party trackers, and 26.4% (SD = 23.6) contained at least one tracker. In 2025, the countries with the highest presence of any trackers were Singapore (SG, 95.5%), Australia (AU, 89.0%), and Bangladesh (BD, 82.5%). The lowest presence was observed in China (CN, 6.1%), Germany (DE, 9.2%), and France (FR, 11.6%). This shows that while some countries consistently use trackers, others show minimal use.

To understand this variability, we look more closely at individual countries over time. Figure 3 plots the tracker presence for all countries, grouped by continent, showing individual country-level and broader trends. Across all continents, third-party trackers are more prevalent than first-party trackers, and their presence has increased more sharply over time as compared to first-party.

In Africa, we observe two distinct groups in the presence of third-party trackers. From the early 2000s until around 2018, all five countries follow a broadly similar trajectory. After this point, the pattern diverges: Nigeria (NG) and Algeria (DZ) show similar behavior, while Egypt (EG), Morocco (MA), and South Africa (ZA) form a separate cluster with higher presence. A similar trend is visible in South America. Initially, third-party tracker presence rises across countries at roughly the same pace, but more recently, two clusters emerge. In Asia and Europe, the picture is far more heterogeneous, with highly divergent patterns of tracker presence. For example, in the context of third-party trackers in Asia, we observe high presence early in Georgia (GE) in 2010 and Kazakhstan (KZ) in 2010, while more recent years show aggressive growth in Bangladesh (BD) and Singapore (SG), both of which surge sharply around 2020. In the case of Bangladesh we observed all websites were using the same Google Tag Manager script URL with the same Measurement ID. This pattern suggests that the rollout was planned and centrally coordinated, rather than ad hoc. One plausible interpretation is that a single GA4 property is being used to aggregate data from multiple government domains, effectively enabling cross-site tracking of users as they move between different services. At the same time, the reuse of a single Measurement ID across multiple, distinct websites could also reflect a configuration error or reliance on a shared template. In that scenario, implementers may simply




Figure 3: Tracker prevalence across 61 countries on six continents, 1995–2025. The top subplot shows the percentage of websites containing first-party trackers, and the bottom shows third-party trackers.

have copied the same Google Tag Manager snippet without fully understanding its implications, inadvertently merging analytics for unrelated services into one property.

At present, the percentage of websites with third-party trackers in Asia spans a wide range—from very low levels in China (CN) (6%) to very high levels in Singapore (SG) (95%). We also observe rise-and-fall, for example, Sweden (SE) peaks in 2020 at 65% of government sites with trackers, then declines despite stable Internet Archive coverage (291 sites in 2020 vs. 314 in 2025). Similar trajectories appear in France (FR), Estonia (EE), and Italy (IT). In Oceania, Australia (AU) and New Zealand (NZ) exhibit mostly similar trajectories for both first- and third-party trackers. However,

after 2021, first-party tracker presence in Australia increases and then stabilizes at a higher level, whereas in New Zealand it remains relatively low and stable, continuing the same pattern as earlier years. All countries in North America follow similar trends.

RQ2: Taken together, regional and country trajectories diverge significantly. Third-party trackers dominate everywhere, but adoption clusters by geography vary (e.g., SG/BD surging, DE/FR remaining low). Thus, while most countries trend upward, the timing and levels vary, underscoring strong geographic heterogeneity in government website tracking.

At the present time, is tracker prevalence increasing? Given the overall trend towards increasing prevalence of trackers, a relevant question is whether usage is still increasing. We answer this question in two ways: (i) by determining *when* tracker use reached its peak for each country—specifically, whether maxima occur in 2025 or earlier—and (ii) by characterizing the *shape* of change over time (increase or decline).

We find that peaks in tracker presence (visible in Figure 3) tend to occur in the most recent years of the study period, especially for third-party trackers. At the country-level, 2025 registers 26 countries having their highest observed percentage of tracker-enabled websites, followed by 2024 (14 countries), 2023 and 2021 (6 countries each) and in 2020 (3 countries). A few countries buck this trend, however, reaching their peaks before 2020. For example, Kazakhstan peaked in 1997, France in 2015, China in 2018, and Germany in 2019. These countries may have reached a plateau in tracker adoption, or may in fact be in a period of decreasing usage. At the continental level, mixed patterns exist. In Asia, 10 of 18 countries peaked in 2025, as did 5 of 6 countries in South America. Europe shows a more mixed pattern, with 6 countries peaking in 2025, another 6 in 2024, and 14 peaking earlier. In North America, two countries peaked in 2024 and two in 2025. Africa shows a similar split, with two peaks in 2023 and two in 2025. Oceania also divides evenly, with one peaking in 2024 and one in 2025. This suggests that certain regions may tend to have different attitudes or policies around trackers on government websites.

While global and regional summaries reveal long-run growth in tracker presence, they can mask shorter periods of acceleration or decline. To assess whether the direction of change within shorter intervals is statistically monotonic (consistently increasing or decreasing), we apply the non-parametric *Mann–Kendall* (MK) test on disjoint 5-year buckets (i.e., 1996–2000, 2001–2005, ...). We use the MK test because it does not assume normality, tolerates outliers and missing years, and remains valid in the presence of ties (identical values at different time points) – features common in archival web measurements where many early years have zeros or plateaus. We use five-year buckets to reduce year-to-year noise due to sampling variation, missing data, and one-off events. For each country and five-year bucket, we run the Mann–Kendall test on the yearly percentage of sites with trackers in that span. We label a bucket *increasing* or *decreasing* only when the two-sided *p*-value is below 0.05, otherwise we report *no-trend*. For websites with third-party trackers, 56 of 61 countries show at least one five-year bucket with a statistically significant trend (*increasing* or *decreasing*). The remaining five have only no-trend or insufficient-data buckets. For first-party trackers, 26 countries show at least one significant trend.

Building on Mann–Kendall, we look more closely at statistically significant *decreases* because they are more anomalous: they do not follow the long-term trend showing an increase in tracker presence. Focusing on decreases helps isolate meaningful local pullbacks rather than restating the overall rise in websites with tracker presence. Among third-party trackers, 15 countries have at least one decreasing bucket, Africa (1: EG), Asia (3: CN, JP, SG), Europe (9: BE, BG, CH, CZ, FR, IT, LV, MD, SE), North America (1: US), and South America (1: AR). As illustrated in Figure 3, these declines are temporary and occur in separate periods. For example, China (CN) shows decreases in 2001–2005 and 2021–2025, the United States (US) in

2011–2015, Sweden (SE) in 2021–2025, France (FR) in 2016–2020 and 2021–2025, and Czechia (CZ) across 2011–2015. For first-party trackers, 16 countries have at least one decreasing bucket: Asia (7: BD, CN, GE, IN, MY, TW, VN), Europe (6: CH, DK, GB, IT, MD, SE), North America (1: CR), and Africa (2: DZ, NG). Fig. 3 shows examples, including Vietnam (VN) in 2001–2005, Bangladesh (BD) in 2001–2005 and 2011–2015, Malaysia (MY) in 2011–2015, Algeria (DZ) in 2011–2015, and Sweden (SE) in 2011–2015.

RQ1 and RQ2: Overall, tracker prevalence continues to rise in most countries, with peaks concentrated in the most recent years; however, we also detect statistically significant, localized pullbacks indicating that plateaus or declines can occur even amid a global upward trend. In short, the global direction remains upward, but few reversals occur and are geographically uneven.

Do trackers ever disappear entirely from government websites? Having seen that tracker use sometimes shows a decrease in certain countries, we next ask the question of whether it ever disappears entirely. Indeed, in several countries, we observe tracker usage disappearing entirely for a time period. We look for *tracker gaps*: sequences of one or more consecutive years in which the share of sampled government websites using trackers drops to zero, *sandwiched* between years with non-zero values. This sandwich definition is appropriate because, in recent years, no country’s share drops to zero. Additionally, years with no Internet Archive data (common in the early period) are treated as missing and excluded from gap detection. We apply this test to each country over 1996–2025 to capture disappear-and-return patterns in tracker use. We find 26 countries that have at least one *tracker gap*.

These tracker gaps are important because they may indicate meaningful behavioral changes. A temporary disappearance of trackers could suggest that governments deliberately opted out of tracking for a time, before later resuming. Alternatively, they could reflect limitations in archival coverage. A gap is recorded only when every sampled site in a given country-year is tracker-free. Because we have two snapshots per country-year, the inferred timing of gaps is an *upper bound*: if a gap appears between yearly snapshots, trackers may have disappeared and reappeared at any point in between. Our within-year design, which includes one randomly chosen snapshot per site, makes it unlikely that routine site-level maintenance would produce a synchronized, country-wide absence. We therefore interpret tracker gaps as country-level pauses or reversals in deployment. Nonetheless, the presence of such patterns demonstrates that these instances do exist.

In total, we identify 39 tracker gaps across the dataset. To ensure these gaps are not artifacts of missing data (i.e., a site observed with a tracker in year y_n but absent from the archive in y_{n+1}), we perform an additional check. For each gap spanning years i through j ($i \leq j$), every website observed with a tracker in the year immediately before the gap ($i - 1$) must have an archived snapshot in *every* year $t \in [i, j]$. This makes sure that a gap is not detected merely because previously websites with trackers are missing from the archive during the gap interval. In 37 cases, the same government websites that previously used trackers were still present in the archives during the gap years, but without trackers,

indicating decision to suspend their use. Only in two cases was the drop likely due to missing archival data. We note that page dynamism may also be a factor contributing to tracker gaps, i.e., the tracker might have been present during a particular year but was just not present in the archived snapshot we collected. While we collect two snapshots per year to alleviate this issue to some extent, we cannot fully confirm that tracker gaps are caused by intentional decisions to suspend trackers without having a large number of snapshots per year. However, the presence of tracker gaps spanning multiple consecutive years (and snapshots), followed by consecutive years of tracker presence and growth in countries such as the United Arab Emirates (AE), Nigeria (NG), Singapore (SG), hints at the possibility of deliberate temporary suspension. Figure 8 in Appendix 9.6 shows the percentage of government websites with trackers over the years and also shows their *tracker gaps* for countries with tracker gaps and their timing. Figure 6 in Appendix 9.4 shows the percentage of government websites with trackers as a heatmap for all 61 countries.

RQ1 and RQ4: We observe *tracker gaps*, periods where trackers disappear entirely and later return—across many countries. Most gaps occur while the same sites remain archived, suggesting the possibility of intentional, temporary suspensions.

5 Organizations Owning the Trackers

Having looked at global trends, we next investigate which organizations own or operate the observed trackers. This has important privacy implications, because first-party trackers (those operated by the government websites itself) have significantly different privacy risks than those operated by third parties. The identities of those third parties are also important, as their locations may affect the privacy laws with which they need—or need not—comply.

Do government websites tend to deploy first- and third-party trackers together, or are these largely independent? We classify trackers as first-party or third-party based on eTLD+1. If a tracker’s eTLD+1 matches the eTLD+1 of the government site it appears on, we label it first-party, otherwise, we label it third-party. For each country, we include all government websites with at least one archived snapshot in our dataset (1996–2025). A site counts toward a country’s first-party (or third-party) percentage if it ever exhibits at least one first-party (or third-party) tracker at any point in the study timeline, otherwise, it does not.

Across all countries and years, we observe 16.5K websites in total. Of these, only 1990 ever use a first-party tracker and 10.7K use a third-party tracker. By country, the percentage of websites with any first-party tracker averages 13.96% (SD 7.2; 25th percentile 8.33%; 75th percentile 18.38%) and third-party averages 66.88% (SD 17.8; 25th percentile 58.15%; 75th percentile 78.70%). Uruguay (UY) has 0% of sites with first-party trackers (35 sites total) and 62% with third-party trackers, followed by Singapore (SG) with 1% first-party, 90% third-party, Bangladesh (BD, 2.2%, 88%), Russia (RU, 2.7%, 80%) and Switzerland (CH, 3%, 24%). The highest first-party shares appear in Australia (AU) (33%), Malaysia (MY) (30%), and Latvia (LV) (26%). For third-party presence, the lowest countries are Germany (DE) (22%), Switzerland (CH) (24%), and South Korea (KR) (32%) and the

highest are Kazakhstan (KZ) (100%), Australia (AU) (93%), Moldova (MD) (90%), New Zealand (NZ) (89%), and Singapore (SG) (90%).

To further investigate whether the same government websites simultaneously adopt first-party and third-party trackers over the study timeline, we compute the Jaccard Similarity (JS) for each country-year pair. For a given year and country, we define two sets: the set of websites with first-party trackers (f_{set}) and the set with third-party (t_{set}). The Jaccard Similarity is then calculated as: $JS = |f_{set} \cap t_{set}| / |f_{set} \cup t_{set}|$.

This measure captures the degree of overlap, with zero indicating no overlap (disjoint sets), while a value of 1 means perfect overlap (the same websites use both first- and third-party trackers). We exclude cases where either f_{set} or t_{set} is null (721 pairs) and also remove 39 instances with missing data, leaving 1070 valid country-year pairs. Across all country-year pair, we observed very low JS values: the mean JS is 0.04 (SD = 0.08), with the 25th percentile at 0.00 and the 75th percentile at 0.06. In 457 of the 1070 country-year pairs the JS equals 0, meaning there was no overlap at all. This indicates that, in most cases, websites using first-party trackers are not the same as those using third-party trackers.

We observe only 9 instances where $JS \geq 0.5$, typically when the number of tracked websites is very small. For example, in Bangladesh in 2002 and 2003, the JS is 1.0, but in both years only a single government website used any tracker. Another example is Nigeria in 2014, where the JS is 0.5: among 114 archived websites, four contained trackers, and two of them used both first- and third-party trackers. Overall, these results show that overlap between first- and third-party trackers on government websites is limited to countries with a small number of tracking websites. Most of the global growth in tracking is therefore not driven by websites layering first- and third-party trackers together, but rather by the widespread adoption of third-party trackers alone.

RQ2 and RQ1: Across countries, first- and third-party trackers are typically deployed by different sets of government websites, with meaningful overlap appearing only in edge cases with very few tracked sites.

How many unique third-party organizations does each website use? Our next question looks at whether each government website sends all tracking information to the *same* third-party tracking organization, or whether there are cases in which one website uses multiple trackers. As described in Section 3, third-party domains were mapped to organizations using DuckDuckGo’s Tracker Radar. Since its coverage is not exhaustive, domains that could not be mapped were labeled as *Unknown*. In our analysis, all *Unknown* entries were treated as a single organization, which makes the reported counts a conservative lower bound. We analyze this data on a country-by-country basis, looking at the average number of third-party tracking organizations over all of the country’s websites. Across 1,490 country-year pairs with non-zero third-party tracker data, the mean number of unique organizations per website per pair is 1.11, with a median of 1.02. At the country level, averaged across years, the highest values are observed for Russia (RU, 1.44) and Ukraine (UA, 1.34) followed by Singapore (SG, 1.28), Canada (CA, 1.23), Vietnam (VN, 1.22), Israel (IL, 1.21), and Norway (NO, 1.21). Around 25 countries have an average of ≈ 1.10 , and

most of the remainder are close to ≈ 1.00 , with the lowest averages found in Bulgaria (BG, 1.00) and Albania (AL, 1.02). These numbers show that most of the time, each website uses a single tracking organization—though, notably, these numbers do *not* mean that all government websites within a country use the *same* organization's trackers.

There are some outliers: the largest maximum within a year occurs in New Zealand (NZ) in 2013, where at least one website sent data to 10 distinct organizations (year mean 1.37) followed by 8 in New Zealand (NZ) 2012 (mean 1.27, median 1.00) and Turkey (TR) 2020 (mean 1.21, median 1.00). Seven country-year pairs reached a maximum of 6 organizations, 3 of these in New Zealand (NZ), 2 in France (FR), and 1 each in Australia (AU) and Malaysia (MY).

RQ3 and RQ2: Per website, third-party exposure is typically narrow, most government websites send data to a single third-party organization, yet a small tail of outlier sites sends data to many third-party organizations.

To what extent is tracker centralized across organizations, and how has this centralization changed over time? We quantify centralization using the Herfindahl–Hirschman Index (HHI) computed separately at the tracker-domain level (HHI_{tracker}), using eTLD+1, and at the parent-organization level (HHI_{org}). HHI ranges from 0 (many providers with equal presence) to 1 (all incidences belong to a single provider). HHI is a standard concentration measure in industrial organization and is well suited here to summarize, per country-year, whether tracker presence is dominated by a few actors or spread across many. We compute HHI_{tracker} over eTLD+1 trackers and HHI_{org} after aggregating those trackers by their owning organizations. For each country-year, we tabulate the number of sampled government sites that embed each eTLD+1 (and its mapped owner), convert these counts to shares, and apply the HHI. To ensure stability, we exclude country-years with fewer than 10 websites that have any tracker.

Across all observations, HHI_{tracker} has mean 0.374, standard deviation 0.180, median 0.335, interquartile range 0.254–0.456, with minimum 0.055 and maximum 1.000. By contrast, HHI_{org} has mean 0.601, standard deviation 0.203, median 0.607, interquartile range 0.466–0.748, with minimum 0.084 and maximum 1.000. The systematically higher HHI_{org} indicates that multiple tracker domains often roll up to the same owner, yielding greater concentration once ownership is taken into account. Looking at country-level means, eight countries exceed 0.5 on HHI_{tracker} (AL 0.669, TW 0.580, MX 0.564, BG 0.554, BR 0.543, FR 0.541, MD 0.524, GR 0.520), while the lowest values occur in VN 0.237 and SE 0.233. In the cross-country distribution of mean HHI_{tracker} , 14 countries lie below 0.3, 23 between [0.3, 0.4), and 16 between [0.4, 0.5).

We observe country-years in which both HHI_{tracker} and HHI_{org} equal 1, meaning all observed tracker incidences belong to a single eTLD+1 owned by a single organization (BG 2009; EE 2008; GR 2017; MX 2007, 2014). More commonly, we also find 31 country-year pairs where $HHI_{\text{tracker}} < 1$ but $HHI_{\text{org}} = 1$: multiple distinct tracker domains appear, yet ownership is fully consolidated under one parent (e.g., a provider operating several eTLD+1s). Over time, concentrations are non-monotonic rather than steadily rising or falling, for example, the United Arab Emirates fluctuates from

$HHI_{\text{tracker}} = 0.33$ (2017) to 0.23 (2018), then oscillates upward with later increases (2022–2025).

RQ3 and RQ1: Government websites tend to rely on a relatively small set of third-party tracking providers, and this concentration becomes even more pronounced once tracker domains are grouped by their owning organizations.

Within each country, do the counts of unique organizations that own the trackers on government websites differ significantly from one year to the next? Within each country, we want to analyze whether the distribution of website-level counts of unique third-party organizations differs from one year to the next. We treat each calendar year as an independent sample (all websites measured that year) and compare the current year to the immediately preceding year using the Mann–Whitney U test (Wilcoxon rank-sum), two-sided. This nonparametric test evaluates whether two samples come from the same distribution without assuming normality or equal variances—properties that suit our discrete, skew-prone counts and the presence of many ties at small integers (often 1 or 2). We require at least three observations in both years to run a comparison and used $p < 0.05$ as the per-comparison threshold.

Across all countries, most adjacent year pairs are not statistically distinguishable, indicating broadly stable year-to-year distributions. Among the countries with at least one significant shift, the United States (US) shows 5 significant year-to-previous differences out of 29 comparisons (17%), France (FR) 4/25 (16%), Estonia (EE) 3/20 (15%), and Singapore (SG) 2/24 (8%). Several others exhibit a single significant change over their time series, Sweden (SE) 1/28 (4%), Bangladesh (BD) 1/20 (5%), Malaysia (MY) 1/27 (4%), Israel (IL) 1/21 (5%), the Netherlands (NL) 1/24 (4%), and Russia (RU) 1/27 (4%). In particular, Sweden (SE) shows change between 2011 and 2012, in 2011, its government websites embedded trackers from just five organizations (WebTrends, Verizon Media, Google, Magnetic Media Online, and CacheNetworks), but by 2012 this expanded to ten, with the addition of Akamai, Oracle, Chartbeat, Adform, and ClickTale. Similarly, Estonia (EE) exhibits a sharp jump between 2015 and 2016, moving from a single tracker organization (Yoast) to four (Google, Booking.com, Yoast, and Facebook).

RQ3 and RQ1: Most countries are stable, adjacent years rarely differ statistically. But a minority of instances in countries exhibit abrupt, statistically detectable jumps, years in which government sites begin sending data to a wider set of third-party organizations.

Which third-party domains and organizations dominate tracker usage? Across our study period, we observed that the distribution of third-party tracker domains is highly skewed. Two domains, *google-analytics.com* and *googletagmanager.com* were observed in all 61 countries in our dataset, showing the global reach of Google's tracking infrastructure, which has also been observed before [32, 71]. Following these, *facebook.com* appeared in 58 countries and *facebook.net* in 57. Beyond these, other notable tracker domains include *wp.com* (45 countries), *cloudflareinsights.com* (44), *polyfill.io* (41), *google.com* (40), *googlesyndication.com* (39), *doubleclick.net* (36),

go-mpulse.net (34), and *clarity.ms* (33). After this top tier, the long tail becomes extremely thin, 901 organizations appear in only a single country, 67 appear in two, and just 37 are present in more than ten. This shows a highly concentrated ecosystem dominated by a handful of tracker domains and two organizations: Alphabet (Google) and Meta (Facebook).

For every site in all 61 countries, we count each time a website embeds a tracker domain (for example, *google-analytics.com*). We then added those counts across all sites and countries. Finally, we grouped related domains under the same company (e.g., Google, Facebook). Google’s trackers were embedded the most 68K times—far more than anyone else followed by Facebook/Meta (6,635), AT Internet (6227), Automattic (1901), Microsoft (1717), Rijksoverheid (1513), Akamai (1433), Yandex (1232), and EPiServer (1,080). All other organizations had fewer than 1,000 loads. In short, Google dominates both how widely and how often trackers appear, with a very long tail after the top few companies. Surprisingly, Rijksoverheid appears only on the Netherlands (NL) government websites yet still ranks among the top organizations.

We also examine how organizations’ footprints changed across countries over time and find notable differences. For example, Facebook’s trackers peaked at 22% of websites in Israel (IL) in 2021, while in Sweden (SE) they peaked in 2019 and declined by 2021. Automattic (owner of WordPress.com) shows similarly varied patterns, some countries peak at different times and at different levels, Paraguay (PY) reached 8% in 2025 whereas Ukraine (UA) fell from 2% to 1% over the same period. Some countries have used Automattic since 2010 (e.g., the United States, US, and the Netherlands, NL), whereas the United Arab Emirates (AE) only begins in 2024. Figure 5 in Appendix 9.3 shows the percentage of websites embedding Automattic owned trackers across countries over the study timeline.

To analyze change over time, we examined the temporal persistence of the most common trackers within countries. For each year, we identified the top ten tracker domains (measured as the percentage of government websites embedding them) and then computed the Jaccard similarity between consecutive years using a sliding window. This provides a measure of stability, a higher score indicates that the dominant trackers remain largely the same across adjacent years. Overall, we find an average Jaccard similarity of 0.50 (SD = 0.08, 25th percentile = 0.52 and 75th = 0.63). Some countries exhibit high stability, Greece (0.75), Albania (0.74), and Uruguay (0.74) are among those with the most consistent top-tracker sets, closely followed by Ukraine, Russia, and Romania (0.73). By contrast, few countries show low stability. The lowest stability is found in the United Arab Emirates (0.46), South Korea (0.47), and Costa Rica (0.45).

We also observe a strong temporal trend toward increasing stability at the global level. In the late 1990s, average Jaccard scores were very low, 0.15 between 1997–1998. By the mid-2000s this had risen to around 0.50 (2007–2008), reached 0.72 in 2020–2021, and further increased to 0.73 by 2024–2025.

Even though we observe a high overall number of distinct tracker domains across countries, the top ten organizations remain remarkably stable, and their stability has increased over time. Our Jaccard analysis shows that, despite the dynamic nature of web technologies and the rise of real-time bidding (RTB) ad ecosystems, the majority of the most prevalent trackers persist from year to year.

In other words, while the long tail of trackers reflects expansion in tracker domains, the core set of dominant tracker domains remains largely unchanged. The core set are dominated by tracking organizations headquartered in the US, which a presence on the government websites of all countries studied.

RQ3, RQ1 and RQ2: Overall, the government website tracking is highly concentrated, a small set of largely U.S.-based firms—led by Alphabet and Meta—dominates global coverage while the long tail is thin. Over time, the leaders’ footprints have grown increasingly stable.

6 Tracker First-Observed Uptake (FOU)

Our final set of research questions continues our tracker-centric direction, this time asking, for each tracker: “*when did a website/-country start using it?*” Presence alone can be misleading, as a website’s first archived snapshot may appear only after trackers were already added. Tracker *First-Observed Uptake* (FOU), by contrast, helps pinpoint the first time a tracker was used on a site (uptake).

Specifically, we only consider FOU when we observe a government website without a given tracker at year y_n , and then observe the same website with that tracker at a later year y_{n+i} (with $i \geq 1$); y_{n+i} may be consecutive ($i = 1$) or occur further in the future. We count a FOU event only when we observe both states—*absent* at y_n and *present* at y_{n+i} —so the timing is supported by before/after evidence. This avoids misdating cases where a site’s first-ever archived snapshot already contains a tracker. This distinction matters for interpreting the historical data from the Internet Archive. Prior work shows tracker presence as early as 1996 [47], which is accurate for visibility in the archive, but does not by itself identify *when* specific sites first took up a tracker. For example, if a website’s first archived snapshot is in 2015 and already includes Google Analytics, we cannot conclude uptake happened in 2015; it may have occurred years earlier. Our FOU rule requires a documented transition from y_n to y_{n+i} , reducing such left-censoring errors. FOU dates are therefore *first-observed* lower bounds given incomplete archival coverage. Based on this evidence, we compute country-level statistics derived from these websites to understand tracker uptake in each country. Additionally, we focus on eTLD+1 for identifying trackers, as comparing exact URLs may lead to inconsistencies due to changes in URLs over time or across versions (e.g., Google Analytics). This approach allows us to analyze uptake and provide reliable evidence that a particular tracker has been actively used in a country that had not previously employed it.

When did each country first show First-Observed Uptake?

The timeline of tracker FOU across countries shows how web tracking spread over time. Using our $y_n \rightarrow y_{n+i}$ criterion, we identify the first year in which at least one government website in a country shows a documented transition from no tracker to a tracker. The earliest FOU we observe is 1997 (Argentina, AR). By 1998, ten countries—Belgium (BE), Estonia (EE), Germany (DE), Israel (IL), South Korea (KR), Russia (RU), Turkey (TR), Australia (AU), United States (US), and Taiwan (TW) show First-Observed Uptake. In 1999, seven more countries show their first FOU events, followed by ten in 2000, six in 2001, and eight in 2002. After that, fewer than five

countries per year show first FOU between 2003 and 2011, with the last FOU instances in our data appearing in 2010 (Nigeria, NG) and 2011 (Serbia, RS). Over the full timeline, without applying the FOU criterion, we identify 2457 non-unique eTLD+1 tracker occurrences across 61 countries (mean 40 per country, SD 31, 25th percentile 20 and 75th percentile 46). After applying FOU (i.e., requiring a before/after transition), we observe 2248 non-unique eTLD+1 with a mean of 36.85 eTLD+1 occurrences per country (SD 28.89, 25th percentile 18, 75th percentile 44). This corresponds to a median coverage of 91.21% (25th percentile 87.10%; 75th percentile 96.30%) of the trackers we see in presence data. In practical terms, this means that in most countries, we can verify a first uptake for the vast majority of trackers. The remaining 9% are cases where we see a tracker but cannot observe the transition (e.g., the first archived snapshot already includes it). For example, Uruguay (UY) has the lowest coverage, with FOU detected for 6 of 8 eTLD+1 trackers. At the other end, six countries—Israel (IL), Romania (RO), Czechia (CZ), Switzerland (CH), Georgia (GE), and Taiwan (TW) show FOU for 100% of the eTLD+1 trackers observed in IA data.

FOU dates are *first observed* lower bounds: if a site already had a tracker in its first archived snapshot, the true uptake likely occurred earlier. The early concentration of first uptakes (1997–2002) is consistent with an initial adoption of trackers. The small number of new FOU events after 2003 reflects that many countries had already adopted trackers by then or that later adoptions happened before our first clean “absent-then-present” evidence. Compared with presence-only timelines (which can begin in 1996 [47]), FOU provides a simpler, more interpretable signal of *when* uptake occurred, rather than merely *whether* a tracker is visible in the archive.

RQ4, RQ1 and RQ2: First-Observed Uptake (FOU) yields credible, lower-bound adoption dates that mitigate the “first-snapshot” bias, when a site’s earliest archived page already contains a tracker. Applied across countries, FOU shows first uptakes clustering in the late 1990s–early 2000s, with later years characterized more by expansion of existing tracking than by new adoptions.

After first uptake, do websites continue using the same trackers? To explore the stability of a tracker, we measure the *persistence* of each tracker (eTLD+1) after it first appears on a country’s government websites e.g., the share of trackers (eTLD+1) that *continue to appear in every subsequent year* after their first observation (FOU). Across countries, on average, only 30% of trackers persist year-over-year (SD = 9.61, 25th percentile = 25.0, 75th percentile = 36.11). In simple terms, most trackers either disappear, are replaced, or show up inconsistently after their first uptake. The highest persistence is observed in Bulgaria (BG) and Greece (GR), where 56% of adopted trackers continue to appear every year after FOU. Other countries with high persistence include Vietnam (VN) at 47%, Singapore (SG) at 46%, and Estonia (EE) at 44%. At the other end, China (CN) shows only 8% persistence, followed by Thailand (TH) at 10.8%, Uruguay (UY) at 11.1%, and the Netherlands (NL) at 15.5%.

Overall, about 86% of trackers (eTLD+1) are not used consistently after their first appearance; they do not show up every year and stability varies widely across countries. This churn has several causes. One is tracker (eTLD+1) changes and rebranding (domains retire,

migrate, or merge). Many tags are also conditional, triggered only through tag managers under certain triggers (consent choices, page types, user segments) or rotated through ad auctions. In contrast, the core set of dominant trackers (top 10 by prevalence) has grown more stable over time (Section 5), reaching a Jaccard similarity of 0.73 in 2025. This suggests that most inconsistency lies in the long tail rather than among the leaders.

RQ1 and RQ2: After first uptake, most trackers do not persist year-over-year, indicating substantial churn and persistence that varies widely across countries. The instability is concentrated in the long tail, whereas the most prevalent trackers are comparatively more stable over time.

How many new tracker uptakes occur globally each year?

Using FOU at the eTLD+1 level, we count the number of new trackers (eTLD+1) appearing for the first time each year worldwide (1996–2025). Across the study period, the average is 77 new uptakes per year (SD 46.7, 25th percentile 41.5, 75th percentile 118.8). The series starts at one in 1997, rises through the mid-2000s (36 in 2003), and surges in the mid-2010s (77 in 2013), peaking in 2017 (161) and staying elevated through 2018–2024 (100 to 160 per year). The lower value in 2025 (73) likely reflects partial archival coverage. At the country level (per-year averages over the study period), the United States averages 6.28 new uptakes per year (182 total), the United Kingdom 3.55 (103), and Australia 3.31 (96). The lowest averages are Japan 0.38 (11), Albania 0.34 (10), and Uruguay 0.24 (7). Overall, countries average 1.31 new uptakes per year (SD 1.00, 25th percentile 0.66, 75th percentile 1.55). The first-observed uptake detection also depends on the trackers list used (e.g., Easylist) and is often stronger for English-language vendors, which helps explain why English-speaking countries dominate absolute counts.

The tracker ecosystem is high-churn, with new tracking domains appearing due to page dynamism and real-time bidding (RTB) rotations that swap them over time. As a result, “new” trackers (eTLD+1) will continue to show up even if overall practices are steady. Importantly, each additional eTLD+1 typically corresponds to a distinct organization, so more first-observed uptakes generally mean users’ data is flowing to a larger set of organizations (though some domains may represent rebrands or are owned by one organization).

RQ1 and RQ4: Annual FOU accelerate through the 2000s, surge in the mid-to-late 2010s, and remain elevated into the 2020s, indicating a steady inflow of new tracking domains even when overall practices appear stable. Because FOU relies on blocklists and vendor coverage, counts likely tilt toward English language countries/trackers, but in general, more uptakes mean users’ data flow to a broader set of organizations.

7 Discussion

Triangulation with browsing- and flow-based work: Bird et al. [13] analyze anonymous Firefox browsing histories from mid-2019 over top-list domains and find third-party tracking to be highly concentrated, led by Alphabet and Meta. More recently, Singh et al. [71] directly measure tracker data flows from in-country user vantage points across 23 countries and find that foreign (non-local)

trackers occur in 91% of examined countries, with U.S. organizations dominating the set of foreign trackers and much hosting routed through European infrastructure. Although both works study the general web rather than only government sites specifically, their observations support our interpretation that the rise we observe on government portals is coupled with a highly concentrated tracking ecosystem led by large U.S.-based organizations.

Real-world events: Our timelines reveal statistically significant local pullbacks within the broader upward trend, including temporary *tracker gaps* where use falls to zero before returning. Linking such shifts directly to real-world events is inherently challenging. Nevertheless, we find instances where declines align with documented external incidents.

- **The polyfill.io compromise:** In June 2024, the domain and GitHub account for *polyfill.io* were acquired by a Chinese company, Funnnull. Shortly thereafter, the *polyfill.js* library was modified to inject malicious code into any website embedding scripts from *cdn.polyfill.io*. This turned a widely used compatibility library into a vector for attack. Documented outcomes included redirecting users to scam sites, stealing sensitive data, and potentially enabling remote code execution [11, 69]. Unsurprisingly, we observe a sharp decline in *polyfill.io* usage across our country sample during 2024–2025. In many cases (Egypt (EG), Israel (IL), Singapore (SG), Denmark (DK), Norway (NO), Estonia (EE), and Canada (CA)) government sites stop embedding *polyfill.io* scripts after 2024. Norway, in particular, shows a clear decline. Egypt is a notable exception, with usage unexpectedly increasing post-2024. Even more surprising, Kazakhstan begins embedding *polyfill.io* for the first time in 2024, despite widespread publicity around the compromise. Figure 7 in Appendix 9.5 shows the trajectories of *polyfill.io* usage across selected countries.
- **Russia after February 2022:** In early March 2022, after Russia’s invasion of Ukraine, Russian authorities blocked major Western platforms (e.g., Facebook and Instagram), while Google suspended advertising in Russia and tightened monetization policies. These incidents plausibly reduced or removed Western-tracker embeds and ad tags on .ru government sites [17, 53]. These moves occurred amid Russia’s broader, ongoing internet isolation [38]. Empirically, we observe a clear effect; the share of RU sites with third-party trackers falls from 62.9% to 44% in 2023 and remains stable at 46% in 2024 (Figure 3). Internet Archive coverage remains comparable (105 sites in 2022 vs. 91 in 2024), suggesting the pattern is not driven by missing data. Tracker-domain patterns also shifted, we see no instances of *google-analytics.com* on RU sites after 2022, and no use of *facebook.net* after 2024.
- **Regulatory Signals and the Rise of Privacy-Centric Analytics:** Our data show that privacy-centric analytics services (e.g., Matomo Cloud, Plausible, Fathom) appear on government websites only from 2021 onwards, and mainly in jurisdictions where TCF-based adtech and Google Analytics have faced the strongest legal challenges. We first see *matomo.cloud* on French government sites in 2021, then in Belgium, Italy, Norway, and Australia, coinciding with CNIL guidance that Matomo, when configured in specific ways, can be exempt from consent as an audience-measurement tool [49, 56]. Plausible (*plausible.io*) is almost entirely confined to European countries (BE, DE, DK, EE,

GB, LV, NO, SE, plus DZ and MY), with adoption starting in 2021 (first in EE) and growing from about 1% to 4% of Norwegian government sites between 2022 and 2025. Similarly, *usefathom.com* first appears in Belgium in 2021 and then in GB, NL, NO, and NZ; its documentation foregrounds GDPR compliance and a Schrems II–oriented “EU isolation” mode [8, 9]. These patterns align with mounting regulatory pressure on dominant adtech: the ICO’s 2019 finding that real-time bidding is systemically non-compliant [39, 79], the Belgian DPA’s 2022 decision that the Transparency and Consent Framework itself generates personal data, the *Schrems II* ruling on EU–US data transfers [62], and, crucially for analytics, post-2022 decisions by several DPAs (first in Austria, then in France and Denmark) that standard Google Analytics deployments and associated transfers to US servers are unlawful under the GDPR [21, 31, 64, 70]. Overall, our observations suggest that such privacy-centric analytics tools were available long before they appear in our dataset (Matomo dates back to 2007 as Piwik), and that their deployment on government sites follows prominent GDPR- and Schrems II–related guidance with a substantial lag rather than through abrupt technical shifts.

- **Gradual Adoption of Cookie-Consent Infrastructures:** Our data show that cookie-consent tools (e.g., *cookielaw.org*, *cookiepro.com*, *onetrust.com*) diffuse gradually across countries rather than appearing everywhere at once. *cookielaw.org* appears on government sites in 11 EU countries, first in Spain in 2020, then in Brazil from 2021 and in the United States by 2024, indicating that a solution originally linked to European “cookie law” has been exported to non-EU jurisdictions. *onetrust.com* exhibits the clearest growth curve: in Japan it is present on about 1% of government sites in 2021 and around 4% by 2025, in Thailand it rises from roughly 1% to 3% between 2022 and 2025, and it also appears in Great Britain (2022), Portugal (2023), and Australia (2024). Overall, these patterns point to gradual, uneven adoption of consent-management platforms, with different tools gaining traction at different times and in different regions.

Exposure Equity: We observe substantial cross-country heterogeneity in third-party tracker prevalence, even among neighbors on the same continent and among countries with similar digital capacity. For example, in 2015, France shows 68% of government sites embedding third-party trackers, while the Netherlands shows just 5%. In 2024, Denmark, Estonia, and Singapore all have high and similar *EGDI* scores (the United Nations *E-Government Development Index*, a 0–1 composite of online service provision, telecommunications infrastructure, and human capital) [77], at roughly 0.98–0.96, yet their tracker prevalence differ sharply: 32%, 72%, and 88% of government sites with third-party trackers, respectively. This suggests that governance choices, rather than digital capacity alone, drive adoption. Because these websites are used by local residents for essential services, these gaps translate into unequal exposure—different populations encounter substantially different *shares of government websites* embedding third-party trackers, exposing citizens to different trackers and creating distinct privacy and security risks (e.g., cross-border data flows, personal data exposure).

Study on Government Websites vs. Top Websites: We compare our longitudinal results to a previous study that analyzed tracker presence on the top 500 websites using Internet Archive

data (1996–2016) [47]. We can fairly compare only selected tracker domains (eTLD+1) based on the graphs they provide. For example, they show *gstatic.com* rising after 2011 until 2016. In our data, *gstatic.com* is never recorded in three continents (Asia, Oceania, Africa), appears in only six countries overall, and is present on fewer than 1% of government websites in any year. Another example is *scorecardresearch.com*. We observe it in just 10 countries and not for most years. It is most visible in MX from 2015 to 2025; in the USA from 2010 to 2017; in GB from 2016 to 2017. In all years it remains very low (under 2% on average), except in MX where it peaks at 14% in 2020. In their results it ranks as a top tracker domain. They also report a pattern for *google-analytics.com*, a major presence by 2011 that declines through 2015. In our government-site data (across all countries), we see a small peak around 2009, a period of normalization (decrease), and then another peak around 2019–2020. The takeaway is that results derived from global or top websites are not reliable proxies for the government websites and vice versa. Future work should avoid generalizing from top site or regional top sites samples to government websites.

Policy and Regulatory Impact: Several jurisdictions have introduced privacy laws—EU *GDPR* (Regulation (EU) 2016/679; applicable 25 May 2018) [30], Egypt’s Personal Data Protection Law No. 151 (2020; in force 17 Oct 2020) [67], India’s Digital Personal Data Protection Act (No. 22 of 2023; enacted 11 Aug 2023; commencement by notification) [3], Russia’s Personal Data Law No. 152-FZ (2006) and subsequent data-localization amendments via 242-FZ (effective 1 Sep 2015) [1, 2] and related amendments. We note that these laws do not directly regulate whether a site includes tracking scripts. They mostly set rules for how data is handled (e.g., consent, transparency, and data-localization), which we do not analyze in our study. In our data, we do not see a consistent change in the percentage of government sites using third-party trackers around the dates these laws took effect.

Web Templates and Shared Codebase: Government web templates and shared codebases can make third-party analytics effectively default-on once embedded in a site’s standard *skeleton*. For example, the U.S. Department of Agriculture’s *Analytics-Plays* treat analytics tags as required infrastructure, directing site owners to include both the government-wide Digital Analytics Program (DAP) code and a USDA Google Analytics ID, supported by a government-wide JavaScript snippet distributed for reuse [60, 65]. Similar official templates and frameworks exist in California [61], Canada [76, 80], Australia [35, 55], the United Kingdom [36], and India [54]. Once embedded in these defaults, analytics snippets are replicated with little scrutiny, causing widespread and durable adoption across agencies and jurisdictions.

Recommendations: We recommend that public authorities adopt centrally maintained web templates with third-party trackers disabled by default, requiring explicit, documented opt-in for any additions, and embed safeguards in procurement and contracts that mandate full visibility into third-party components, documented data flows and roles, and a standing right to remove non-essential trackers. Authorities should prioritize privacy-preserving analytics, such as self-hosted or log-based solutions with strict retention limits, and run periodic automated checks of official domains for third-party requests, reporting results internally and publishing

high-level summaries alongside a simple tracker notice on official sites. At the same time, users can reduce unnecessary exposure by enabling privacy-protective modes (blocking third-party cookies, using built-in tracking protection or a reputable content blocker), separating contexts for government tasks (e.g., using a private window or dedicated browser/profile and avoiding concurrent logins to large platforms), and limiting ambient permissions and extensions by setting sensitive permissions to “ask every time” and disabling non-essential extensions when accessing government services.

Limitations: While our study provides valuable insights, it also has several limitations. The main limitations arise from the use of archival data for analysis. Prior work has demonstrated that Internet Archive data is reliable enough to be used in similar studies on security and privacy [37, 47]. Our study suffers from similar issues to prior work. Archiving is not uniform across countries or websites, leading to potential sampling bias where some sites are overrepresented while others are sparsely covered. Also, snapshots are often irregular in frequency, which means that important changes (such as the introduction or removal of trackers) may not be captured. Archived pages may occasionally be incomplete or missing resources (e.g., scripts or embedded content) due to crawler restrictions, robots.txt exclusions, or technical errors during capture. Page dynamism, influenced by various factors including run-time ad-tech protocols, may lead to variations in tracker behavior captured by snapshots. We note that tracker lists such as EasyList and DuckDuckGo Tracker Radar provide extensive coverage of commonly-observed advertising and tracking domains, but they may not capture the full universe of trackers. In addition, we apply a single snapshot of lists (June 2025) to all historical snapshots. While our validation suggests that this snapshot has higher precision and at least comparable ability to detect trackers on old traffic, it also introduces a temporal bias: our labels reflect what is known to be tracking in 2025, and may fail to recover some trackers that were active only in earlier periods or never encoded in the current lists. As a result, our longitudinal measurements should be interpreted as a conservative lower bound on historical tracking prevalence. Despite these limitations, we think that our study manages to capture the main trends in tracker adoption and evolution, and contributes to understanding the privacy of government sites.

8 Conclusion

Our study shows that government websites, across the world, have been increasingly adopting third-party trackers, posing privacy and security risks for users who rely on these sites for essential services. Third-party tracking is dominated by a few large US-based organizations, but there is a long tail of smaller, lesser-known tracking domains present on these sites. We observe large heterogeneity in tracker adoption, presence, and evolution across countries, resulting in deeply unequal tracking exposure for citizens.

Acknowledgments

The authors used generative AI-based tools to revise the text, improve flow and correct any typos, grammatical errors, and awkward phrasing. ChatGPT was also used to assist background research on notable incidents involving trackers and broader tracking-related trends. Specifically, we used ChatGPT-5.1 in November 2025 to

map tracker domains to publicly reported events that could plausibly relate to adoption or trend changes. Every AI-suggested lead was independently verified through additional web research, and only information confirmed by reliable sources was included in the paper.

This work was partly funded by the National Science Foundation (NSF) under Award No. CNS 2027208.

References

- [1] 2006. Federal Law No. 152-FZ of 27 July 2006 On Personal Data. ICRC, National Implementation Database. <https://ihl-databases.icrc.org/en/national-practice/federal-law-no-152-fz-personal-data-2006>
- [2] 2014. Data Localization Amendments to Russia's Personal Data Law (Federal Law No. 242-FZ, effective 1 Sep 2015). Morgan Lewis, *Data Localization Laws: Russian Federation*. <https://www.morganlewis.com/-/media/files/publication/outside-publication/article/2021/data-localization-laws-russian-federation.pdf> Amends Law No. 152-FZ; effective date 1 Sep 2015.
- [3] 2023. The Digital Personal Data Protection Act, 2023 (No. 22 of 2023). Ministry of Electronics and Information Technology (MeitY). <https://www.meity.gov.in/static/uploads/2024/06/2bf1f0e9f04e6fb4f8fe35e82c42aa5.pdf> Enacted 11 Aug 2023; commencement by notification.
- [4] Zainul Abi Din, Panagiotis Tigas, Samuel T King, and Benjamin Livshits. 2020. {PERCIVAL}: making {In-browser} perceptual Ad Blocking practical with deep learning. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*. 387–400.
- [5] Gunes Acar, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. 2014. The web never forgets: Persistent tracking mechanisms in the wild. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. 674–689.
- [6] Alessandro Acquisti and Jens Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE security & privacy* 3, 1 (2005), 26–33.
- [7] Mshabab Alrizah, Sencun Zhu, Xinyu Xing, and Gang Wang. 2019. Errors, misunderstandings, and attacks: Analyzing the crowdsourcing process of ad-blocking systems. In *Proceedings of the Internet Measurement Conference*. 230–244.
- [8] Fathom Analytics. 2021. GDPR compliant website analytics. <https://usefathom.com/legal/compliance/gdpr-compliant-website-analytics> Accessed: 2025-11-16.
- [9] Fathom Analytics. 2021. Schrems II compliant website analytics. <https://usefathom.com/legal/compliance/schrems-ii-compliant-website-analytics> Accessed: 2025-11-16.
- [10] Pounh Nikkhhah Bahrami, Umar Iqbal, and Zubair Shafiq. 2022. FP-radar: Longitudinal measurement and early detection of browser fingerprinting. *Proceedings on Privacy Enhancing Technologies* 2022 (2):557–577 (2022).
- [11] Zbigniew Banach. 2024. Polyfill Supply Chain Attack: What to Do When Your CDN Goes Evil. <https://www.invisi.com/blog/web-security/polyfill-supply-chain-attack-when-your-cdn-goes-evil/>. Invisi Web Security Blog. Accessed: 2025-09-01.
- [12] Nataliia Bielova, Laura Litvine, Anyisia Nguyen, Mariam Chammam, Vincent Toubiana, and Estelle Hary. 2024. The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 2813–2830. <https://www.usenix.org/conference/usenixsecurity24/presentation/bielova>
- [13] Sarah Bird, Ilana Segall, and Martin Lopatka. 2020. Replication: Why we still can't browse in peace: On the uniqueness and reidentifiability of web browsing histories. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. 489–503.
- [14] Christian Böttger, Nurullah Demir, Jan Hörnemann, Bhupendra Acharya, Norbert Pohlmann, Thorsten Holz, Matteo Große-Kampmann, and Tobias Urban. 2025. Understanding Regional Filter Lists: Efficacy and Impact. (2025).
- [15] California Attorney General. 2024. California Consumer Privacy Act (CCPA) – Overview and Regulations. <https://oag.ca.gov/privacy/ccpa>
- [16] Quan Chen, Peter Snyder, Ben Livshits, and Alexandros Kapravelos. 2021. Detecting filter list evasion with event-loop-turn granularity javascript signatures. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1715–1729.
- [17] Elizabeth Culliford. 2022. Russia Blocks Facebook, Accusing It of Restricting Access to Russian Media. <https://www.reuters.com/business/media-telecom/russia-blocks-facebook-accusing-it-restricting-access-russian-media-2022-03-04/>. Reuters. Published: March 5, 2022. Accessed: 2025-09-01.
- [18] Sourya Joyee De and Daniel Le Métayer. 2018. Privacy Risk Analysis to Enable Informed Privacy Settings. In *Proceedings of the 4th IEEE International Workshop on Privacy Engineering (IWPE 2018) (Proceedings of the 4th IEEE International Workshop on Privacy Engineering (IWPE 2018))*. London, United Kingdom, 1–8. <https://hal.science/hal-01939845>
- [19] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2019. We value your privacy... now take some cookies: Measuring the GDPR's impact on web privacy. *Network and Distributed System Security Symposium (NDSS)* (2019).
- [20] Tamara Dinev, Paul Hart, and Michael R Mullen. 2008. Internet privacy concerns and beliefs about government surveillance—An empirical investigation. *The Journal of Strategic Information Systems* 17, 3 (2008), 214–233.
- [21] DLA Piper. 2022. France: The CNIL provides further insights following its formal notices against the use of Google Analytics. <https://privacymatters.dlapiper.com/2022/06/france-the-cnil-provides-further-insights-following-its-formal-notice-against-the-use-of-google-analytics/> Accessed: 2025-11-16.
- [22] DuckDuckGo. 2025. Tracker Radar. <https://github.com/duckduckgo/tracker-radar>. GitHub repository. Accessed: November 20, 2025.
- [23] EasyList. 2025. EasyList Filter List. <https://easylist.to/easylist/easylist.txt> Accessed: 2025-06-25.
- [24] EasyList. 2025. EasyList: Main Page. <https://easylist.to/index.html> Accessed: 2025-06-25.
- [25] EasyList. 2025. EasyPrivacy Filter List. <https://easylist.to/easylist/easyprivacy.txt> Accessed: 2025-06-25.
- [26] EasyList. 2025. Supplementary Filter Lists and EasyList Variants. <https://easylist.to/pages/other-supplementary-filter-lists-and-easylist-variants.html> Accessed: 2025-06-25.
- [27] Peter Eckersley. 2010. How unique is your web browser?. In *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 1–18.
- [28] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 1388–1401.
- [29] European Union. 2009. Directive 2009/136/EC amending Directive 2002/58/EC (ePrivacy; cookie consent). <https://eur-lex.europa.eu/eli/dir/2009/136/oj/eng>
- [30] European Union. 2016. General Data Protection Regulation, Article 9: Processing of special categories of personal data. <https://gdpr-info.eu/art-9-gdpr/>
- [31] European Center for Digital Rights. 2022. Austrian DSB: EU-US data transfers to Google Analytics illegal. <https://noyb.eu/en/austrian-dsb-eu-us-data-transfers-google-analytics-illegal> Accessed: 2025-11-16.
- [32] Alexander Gamero-Garrido, Kicho Yu, Sumukh Vasisht Shankar, Sachin Kumar Singh, Sindhya Balasubramanian, Alexander Wilcox, and David Choffnes. 2025. Empirically Measuring Data Localization in the EU. *Proceedings on Privacy Enhancing Technologies* 2025 (2025).
- [33] Ghostery. 2025. WhoTracks.Me Search Tool. <https://www.ghostery.com/whotracksme/search> Accessed: 2025-11-12.
- [34] Matthias Gotze, Srđjan Matic, Costas Jordanou, Georgios Smaragdakis, and Nikolaos Laoutaris. 2022. Measuring web cookies in governmental websites. In *Proceedings of the 14th ACM Web Science Conference 2022*. 44–54.
- [35] GovCMS. 2025. Get Started: Onboarding. <https://www.govcms.gov.au/get-started/onboarding>. Accessed: 2025-11-16.
- [36] Government Digital Service (GDS). 2025. GA4 Resources and Guidance. <https://docs.data-community.publishing.service.gov.uk/processes/ga4-resources/>. Accessed: 2025-11-16.
- [37] Florian Hantke, Stefano Calzavara, Moritz Wilhelm, Alvis Rabitti, and Ben Stock. 2023. You call this archaeology? evaluating web archives for reproducible web security measurements. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 3168–3182.
- [38] Human Rights Watch. 2025. Russia: Internet Blocking, Disruptions and Increasing Isolation. <https://www.hrw.org/news/2025/07/30/russia-internet-blocking-disruptions-and-increasing-isolation>. News release. Accessed: 2025-09-01.
- [39] Information Commissioner's Office. 2019. Update report into adtech and real time bidding. (2019). <https://ico.org.uk/media2/migrated/2615156/adtech-real-time-bidding-report-201906-dl191220.pdf> Accessed: 2025-11-16.
- [40] Internet Archive. 2025. The Internet Archive's Wayback Machine. <https://archive.org/web/>. Accessed: 2025-12-12.
- [41] Costas Jordanou, Georgios Smaragdakis, Ingmar Poese, and Nikolaos Laoutaris. 2018. Tracing cross border web tracking. In *Proceedings of the internet measurement conference 2018*. 329–342.
- [42] Bernardus Jansen, Natalia Kadenko, Dennis Broeders, Michel van Eeten, Kevin Borgolte, and Tobias Fiebig. 2023. Pushing boundaries: An empirical view on the digital sovereignty of six governments in the midst of geopolitical tensions. *Government Information Quarterly* 40, 4 (2023), 101862.
- [43] James BD Joshi, Arif Ghafour, Walid G Aref, and Eugene H Spafford. 2002. Security and privacy challenges of a digital government. In *Advances in digital government: Technology, human factors, and policy*. Springer, 121–136.
- [44] Michal Kosinski, David Stillwell, and Thore Graepel. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the national academy of sciences* 110, 15 (2013), 5802–5805.
- [45] Rashna Kumar, Esteban Carisimo, Lukas De Angelis Riva, Mauricio Buzzzone, Fabián E Bustamante, Ihsan Ayyub Qazi, and Mariano G Beiró. 2024. Of Choices and Control-A Comparative Analysis of Government Hosting. In *Proceedings of the 2024 ACM on Internet Measurement Conference*. 462–479.
- [46] Victor Le Pochat, Tom Van Goethem, and Wouter Joosen. 2019. Evaluating the long-term effects of parameters on the characteristics of the franco top sites ranking. In *12th USENIX Workshop on Cyber Security Experimentation and Test*

- (CSET 19).
- [47] Ada Lerner, Anna Kornfeld Simpson, Tadayoshi Kohno, and Franziska Roesner. 2016. Internet jones and the raiders of the lost trackers: An archaeological study of web tracking from 1996 to 2016. In *25th USENIX Security Symposium (USENIX Security 16)*.
 - [48] Timothy Libert, David Grande, and David A Asch. 2015. What web browsing reveals about your health. *Bmj* 351 (2015).
 - [49] Matomo. 2021. *ePrivacy Directive, National Implementations and Website Analytics*. <https://matomo.org/faq/general/eprivacy-directive-national-implementations-and-website-analytics/> Accessed: 2025-11-16.
 - [50] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543.
 - [51] William Melicher, Mahmood Sharif, Joshua Tan, Lujo Bauer, Mihai Christodorescu, and Pedro Giovanni Leon. 2016. (Do Not) track me sometimes: Users' contextual preferences for Web tracking. *Proceedings on Privacy Enhancing Technologies* (2016).
 - [52] Gilles Mertens, Nataliia Bielova, Vincent Roca, and Cristiana Santos. 2025. You Can't Trust Your Tag Neither: Privacy Leaks and Potential Legal Violations within the Google Tag Manager. In *EuroS&P 2025 - 10th IEEE European Symposium on Security and Privacy*. Venice (Ca' Foscari University), Italy, 1–20. <https://hal.science/hal-05032798>
 - [53] Dan Milmo. 2022. Russia Blocks Access to Facebook and Twitter. <https://www.theguardian.com/world/2022/mar/04/russia-completely-blocks-access-to-facebook-and-twitter>. The Guardian. Accessed: 2025-09-01.
 - [54] Ministry of Electronics and Information Technology (MeitY), Government of India. 2025. Centralized Monitoring Framework (CMF). <https://cmf.gov.in/>. Accessed: 2025-11-16.
 - [55] Morpht. 2017. GovCMS and digital transformation. Blog post, <https://www.morph.com/blog/govcms-and-digital-transformation>. Accessed: 2025-11-16.
 - [56] Commission nationale de l'informatique et des libertés (CNIL). 2020. Use analytics on your websites and applications. <https://www.cnil.fr/en/sheet-ndeg16-use-analytics-your-websites-and-applications> Accessed: 2025-11-16.
 - [57] Nick Nikiforakis, Alexandros Kapravelos, Wouter Joosen, Christopher Kruegel, Frank Piessens, and Giovanni Vigna. 2013. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *2013 IEEE Symposium on Security and Privacy*. IEEE, 541–555.
 - [58] Eric Nost, Gretchen Gehrke, Grace Poudrier, Aaron Lemelin, Marcy Beck, Sara Wylie, and Environmental Data & Governance Initiative. 2021. Visualizing changes to US federal environmental agency websites, 2016–2020. *Plos one* 16, 2 (2021), e0246450.
 - [59] Midas Nouwens, Ilaria Liccadi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
 - [60] U.S. Department of Agriculture (USDA). 2019. Analytics Plays | USDA. <https://www.usda.gov/about-usda/policies-and-links/digital/digital-strategy/analytics/analytics-plays>. Accessed: 2025-11-16.
 - [61] California Office of Digital Services. 2025. California-State-Web-Template-Website: The California State Web Template. <https://github.com/Office-of-Digital-Services/California-State-Web-Template-Website>. Accessed: 2025-11-16.
 - [62] Court of Justice of the European Union. 2020. The CJEU judgment in the Schrems II case. [https://www.europarl.europa.eu/RegData/etudes/ATAG/2020/652073/EPRS_ATA\(2020\)652073_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/ATAG/2020/652073/EPRS_ATA(2020)652073_EN.pdf). Accessed: 2025-11-16.
 - [63] Easylist Organization. 2025. EasyList. GitHub repository. <https://github.com/easylist/easylist> Accessed: 2025-11-16.
 - [64] Piwik PRO. 2022. *The Danish Data Protection Agency: Google Analytics is illegal under GDPR*. <https://piwik.pro/blog/google-analytics-illegal-in-denmark/> Accessed: 2025-11-16.
 - [65] Digital Analytics Program. 2025. Digital Analytics Program (DAP) Code for US Federal Websites. <https://github.com/digital-analytics-program/gov-wide-code>. Accessed: 2025-11-16.
 - [66] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. 2012. Detecting and defending against {Third-Party} tracking on the web. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. 155–168.
 - [67] George Sadek. 2020. Egypt: Country's First Law on Protection of Personal Data Enters into Force. Law Library of Congress, Global Legal Monitor. <https://www.loc.gov/item/global-legal-monitor/2020-10-29/egypt-countrys-first-law-on-protection-of-personal-data-enters-into-force/>
 - [68] Nayanamana Samarasinghe, Aashish Adhikari, Mohammad Mannan, and Amr Youssef. 2022. Et tu, brute? Privacy analysis of government websites and mobile apps. In *Proceedings of the ACM web conference 2022*. 564–575.
 - [69] Sheela Sarva. 2025. Understanding the Polyfill.io Supply Chain Attack and Its Impact. <https://blog.qualys.com/vulnerabilities-threat-research/2024/06/28/polyfill-io-supply-chain-attack>. Qualys Blog. Accessed: 2025-09-01.
 - [70] Simmons & Simmons. 2022. *Google Analytics and data transfers: the French position*. <https://www.simmons-simmons.com/en/publications/cl4clgd2m17ge0a447gm3m9xn/google-analytics-and-data-transfers-the-french-position> Accessed: 2025-11-16.
 - [71] Sachin Kumar Singh, Robert Ricci, and Alexander Gamero-Garrido. 2025. Where in the World Are My Trackers? Mapping Web Tracking Flow Across Diverse Geographic Regions. In *Proceedings of the ACM Internet Measurement Conference (IMC '25)*. ACM, USA.
 - [72] Peter Snyder, Antoine Vastel, and Ben Livshits. 2020. Who filters the filters: Understanding the growth, usefulness and efficiency of crowdsourced ad blocking. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 2 (2020), 1–24.
 - [73] Ashkan Soltani, Shannon Canty, Quentin Mayo, Lauren Thomas, and Chris Jay Hoofnagle. 2010. Flash Cookies and Privacy.. In *AAAI spring symposium: intelligent information privacy management*, Vol. 2010. 158–163.
 - [74] Salvatore Stolfo, Eric J Johnson, Tomislav Pavlicic, and Stephen Jan. 2003. Citizen's Attitudes about Privacy While Accessing Government Websites: Results of an Online Study. (2003).
 - [75] such-in. 2025. archive-based-tracking. <https://github.com/such-in/archive-based-tracking>. GitHub repository. Accessed: 2025-12-12.
 - [76] Treasury Board Secretariat of Canada. 2025. GCWeb theme: Implementation of the Canada.ca design system. <https://wet-boew.github.io/GCWeb/index-en.html>. Accessed: 2025-11-16.
 - [77] United Nations Department of Economic and Social Affairs. 2024. United Nations E-Government Survey 2024. 13th edition. Available at <https://publicadministration.un.org/publications/un-e-government-survey-2024-0>.
 - [78] Lonneke Van der Velden. 2014. The Third Party Diary—Tracking the trackers on Dutch governmental websites. *NECSUS. European Journal of Media Studies* 3, 1 (2014), 195–217.
 - [79] Michael Veale and Frederik Zuiderveen Borgesius. 2022. Adtech and Real-Time Bidding under European Data Protection Law. *Computer Law & Security Review* 43 (2022), 105632. https://discovery.ucl.ac.uk/id/eprint/10127120/7/Veale_adtech-and-real-time-bidding-under-european-data-protection-law.pdf
 - [80] Web Experience Toolkit (WET) Team. 2025. gcweb-jekyll: Jekyll theme for the Government of Canada Web presence. <https://github.com/wet-boew/gcweb-jekyll>. Accessed: 2025-11-16.
 - [81] Ahsan Zafar and Anupam Das. 2023. Comparative privacy analysis of mobile browsers. In *Proceedings of the thirteenth ACM conference on data and application security and privacy*. 3–14.
 - [82] Ahsan Zafar and Anupam Das. 2025. Assessing Compliance in Digital Advertising: A Deep Dive into Acceptable Ads Standards. In *Proceedings of the ACM on Web Conference 2025*. 1536–1547.
 - [83] Alexander R Zheutlin, Joshua D Niforatos, and Jeremy B Sussman. 2022. Data-tracking on government, non-profit, and commercial health-related websites. *Journal of general internal medicine* 37, 5 (2022), 1315–1317.

9 Appendix

9.1 Countries in the Study Sample

The table 1 lists all countries included in our dataset, grouped by their respective continents. Each entry shows the ISO-2 country code alongside the official country name.

Continent	Code	Country	Continent	Code	Country
Africa	DZ	Algeria	Europe	AL	Albania
	EG	Egypt		BE	Belgium
	MA	Morocco		BA	Bosnia and Herzegovina
	NG	Nigeria		BG	Bulgaria
	ZA	South Africa		CZ	Czechia
Asia				DK	Denmark
	BD	Bangladesh		EE	Estonia
	CN	China		FR	France
	GE	Georgia		DE	Germany
	HK	Hong Kong		GR	Greece
	IN	India		HU	Hungary
	ID	Indonesia		IT	Italy
	IL	Israel		LV	Latvia
	JP	Japan		MD	Moldova, Republic of
	KZ	Kazakhstan		NL	Netherlands
	KR	Korea, Republic of		NO	Norway
	MY	Malaysia		PL	Poland
	PK	Pakistan		PT	Portugal
	SG	Singapore		RO	Romania
	TW	Taiwan, Province of China		RU	Russian Federation
	TH	Thailand		RS	Serbia
	TR	Turkey		ES	Spain
	AE	United Arab Emirates		SE	Sweden
	VN	Viet Nam		CH	Switzerland
North America				UA	Ukraine
	CA	Canada		GB	United Kingdom
	CR	Costa Rica	South America		
	MX	Mexico		AR	Argentina
Oceania	US	United States		BO	Bolivia, Plurinational State of
				BR	Brazil
	AU	Australia		CL	Chile
	NZ	New Zealand		PY	Paraguay
				UY	Uruguay

Table 1: Study sample: countries grouped by continent

9.2 Archived Government Websites




Figure 4: Count of government websites archived across 61 countries

9.3 Percentage of websites embedding Automatic owned trackers across countries

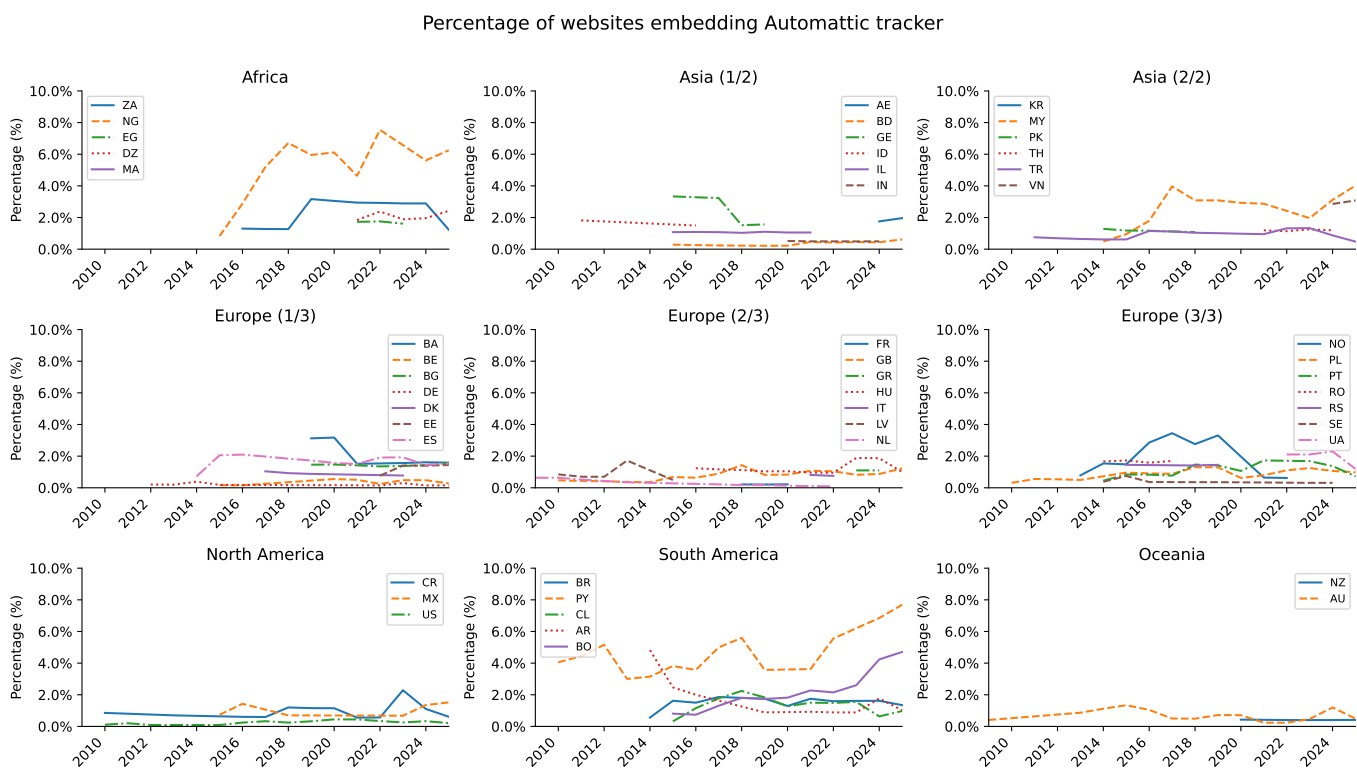


Figure 5: Percentage of websites embedding Automatic owned trackers across countries over the study timeline.

9.4 Government Websites With Trackers




Figure 6: Heatmap showing the fraction of websites with trackers.

9.5 Polyfill.io usage across selected government sites (EG, IL, SG, DK, NO, EE, CA, KZ)




Figure 7: Polyfill.io usage across selected government sites (EG, IL, SG, DK, NO, EE, CA, KZ).

9.6 Heatmap showing the fraction of websites with trackers for countries with tracker gaps




Figure 8: Heatmap showing the fraction of websites with trackers for countries with tracker gaps.