

Watching the Watchers: Privacy Analysis of U.S. Smart TV Automatic Content Recognition

Ilgar Mammadov
Northeastern University
Boston, Massachusetts, USA
mammadov.i@northeastern.edu

Daniel J. Dubois
Northeastern University
Boston, Massachusetts, USA
d.dubois@northeastern.edu

Het Rutul Joshi
Northeastern University
Boston, Massachusetts, USA
joshi.hetr@northeastern.edu

David Choffnes
Northeastern University
Boston, Massachusetts, USA
choffnes@ccs.neu.edu

Abstract

As of 2025, 82% of U.S. households contain smart TVs that support automatic content recognition (ACR), a technology that captures audio or video snapshots of viewed content. ACR presents a particular privacy risk because it enables providers and their partners to learn otherwise private and sensitive content-viewing habits, as well as to monetize that data through targeted advertising and selling/sharing that information with other parties. Despite its importance, little is publicly known about the privacy implications of deployed ACR implementations—in part due to the opaque nature of their deployments and challenges with automated testing.

In this work, we conduct the first highly automated, at-scale, systematic measurement of ACR across the eight dominant smart-TV operating systems that account for approximately 95% of the U.S. market, on twelve commercial devices and ten viewing scenarios per device, under both opted-in and opted-out states. Using 5,943 experiments, we identify and validate ACR-attributable network endpoints through controlled content exposure, evaluate which viewing behaviors exhibit ACR-attributable traffic, determine the extent to which ACR opt-out decisions suppress such traffic, and identify exposure of ACR-based viewing data to other parties. We find that the advertised ACR opt-out fully suppresses ACR-attributable traffic on only four operating systems; on Samsung Tizen traffic consistent with ACR activity persists in more than 60% of evaluated scenarios, and on Vizio SmartCast a content-correlated endpoint hosted on Amazon Kinesis continues to receive traffic in 100% of opted-out scenarios. These behaviors indicate potential privacy violations; however, we cannot know for certain without access to payload content. Further, we reveal that viewing data is exposed over unencrypted channels for two providers; worse yet, we discover an unauthenticated MQTT broker that, in a three-hour ethically-bounded sample, exposed 57 million ACR detection results from more than one million unique TVs (we reported the vulnerability and as a result it was remediated).

Keywords

Smart TV, Automatic Content Recognition, IoT measurement, second-party tracking, network privacy, advertising

1 Introduction

As of 2025, 82% of U.S. TV households contain a smart TV [20], 90% of U.S. households use at least one smart TV device monthly [27], and streaming accounts for 43.8% of total TV viewing time, surpassing cable for the first time [30]. The widespread adoption of smart TVs has exposed consumers to potential privacy harms in the form of viewing information gathered via *automatic content recognition* (ACR). ACR monitors viewed content by capturing frequent audio and/or video snapshots, transferring them to remote servers, and identifying the content through comparison with a media library [8, 39]. However, despite increasing regulatory attention [13, 31, 47] and prior work on a small number of smart TV platforms [8], the community still has limited insight into how deployed ACR systems operate, how they differ across smart-TV OSes, how they relate to disclosures of ACR activity, and whether they expose sensitive viewing information to other parties. The goal of this work is to fill this gap.

Two key challenges limit progress on understanding ACR deployments. First, there is a large and diverse smart TV market [21], so characterizing ACR deployments requires strategies that scale to large numbers of devices and experiments. Second, smart TVs are locked down [8] (making techniques like rooting and TLS interception generally unavailable), so there is a need instead for techniques that reliably infer ACR activity. To address these challenges, our study focuses on bringing greater transparency to ACR behavior across the dominant smart-TV OSes in the U.S. market. Specifically, we address the following research questions:

- **RQ1:** Under what TV-viewing conditions is ACR-attributable traffic observed?
- **RQ2:** To what extent are ACR opt-outs offered and effective?
- **RQ3:** What TV-viewing information exposure is associated with ACR?

To answer these questions, we conduct a comprehensive study of smart TVs on the U.S. market, covering the eight dominant smart-TV OSes, representing approximately 95% of the U.S. market, and TVs from brands used by more than 80% of U.S. households [16, 33]. We identify the ACR implementation partner associated with each

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 YYYY(X), 1–18

© YYYY Copyright held by the owner/author(s).

<https://doi.org/XXXXXXX.XXXXXXX>



OS, present a systematic methodology for identifying potential ACR endpoints and characterize ACR behavior through network measurements, conduct large-scale automated experiments to support our findings with statistical confidence, and provide insight into specific ACR implementation details by analyzing a provider that uses unencrypted ACR traffic for its services.

Our results reveal several concerning findings. First, we observe persistent ACR-attributable traffic after opt-out on multiple smart-TV OSes, which suggests potential privacy violations (though confirmation of ACR transmissions via payload-level analysis is infeasible due to encryption). Further, we find that different OSes conduct ACR on different inputs with different regularity, meaning consumers have no easy way to tell when their content-viewing is being monitored. Finally, we discover multiple scenarios in which viewing information is exposed in plaintext and one vulnerability that exposed 57 million ACR classifications from more than one million unique TVs during a three-hour observation window. The vulnerability affects the ACR ecosystem associated with VIDAA, the second-largest smart-TV OS globally in 2024 [10]. The ACR classifications include potentially sensitive information, such as inferences regarding religion, demographics, and “health and medicine.”

Our contributions are as follows:

- We present a systematic methodology to identify and validate ACR-attributable network endpoints. Using this approach, we conduct the first large-scale study of ACR behavior across eight dominant smart-TV OSes, covering 95% of the U.S. market and devices from brands used by over 80% of U.S. households.
- We show that different smart-TV OSes apply ACR to different content sources and viewing scenarios. We identify and characterize which content sources and viewing scenarios trigger ACR-attributable traffic across different smart-TV OSes, offering consumers visibility into when ACR is active.
- We evaluate ACR opt-out effectiveness across all tested smart-TV OSes and find persistent ACR-attributable traffic after opt-out on two OSes representing approximately 52% of the U.S. smart-TV OS market [33]. While persistent traffic to an ACR endpoint is treated as an opt-out mechanism failure, we cannot determine whether the received data is processed server-side.
- To the best of our knowledge, we provide the first technical analysis of the specific data shared by an ACR provider by examining an implementation that transmits ACR-related information over unencrypted channels.
- We discover multiple vulnerabilities affecting ACR infrastructure associated with more than one million smart TVs globally. We disclosed these vulnerabilities to the affected manufacturer and they have implemented partial mitigations.¹

To support reproducibility and future research, we release our datasets, endpoint classifications, analysis scripts, and the first publicly accessible blocklist dedicated to ACR endpoints [23].

¹We evaluated mitigation in September 2026. We found that the manufacturer fixed two vulnerabilities by requiring authentication to connect to the ACR data server and preventing packet publishing to unauthorized endpoints. However, ACR traffic remained unencrypted.

2 Background and Related Work

This section reviews prior smart-TV privacy measurement research, summarizes how ACR systems operate in practice, describes how ACR deployments relate to smart-TV OSes and vendors, and introduces network measurement concepts used in our methodology.

2.1 Prior Work on Smart-TV Tracking

Prior smart-TV privacy measurement research initially focused on tracking performed by applications and external services. Mohajeri Moghaddam et al. [28] proxied Over-The-Top streaming devices and documented identifiers exposed to ad networks. Varmarken et al. [50] extended this analysis to Roku and Fire OS app ecosystems. Tileria and Blasco [48] audited Android TV apps, and Tagliaro et al. [46] analyzed the HbbTV broadcast-internet hybrid stack. These works share a common methodology (TLS interception or static APK analysis) and analyze tracking by individual applications.

Anselmi et al. [8] analyze *second-party* ACR, establishing that it is active during HDMI and streaming scenarios and that the platform-level opt-out suppresses it on Samsung Tizen and LG webOS. While this work provides the first systematic investigation of second-party ACR, its two-TV scope leaves open questions about generality across the broader ecosystem.

Hyman [21] provides complementary qualitative context on the role of ACR within the TV ad ecosystem. The work argues that ACR enables TV manufacturers to act as data intermediaries in digital advertising markets and helps explain the commercial incentives driving ACR deployment. However, it does not provide technical analysis of ACR implementations on contemporary smart-TV OSes.

To understand the awareness and concern of smart-TV users regarding tracking practices and privacy vulnerabilities, Wang et al. [56] interviewed 22 smart-TV users. Their findings show that while users are generally concerned about data collection, they commonly attribute tracking to third-party services collecting watch history and account details, reflecting a limited understanding of OS-level tracking mechanisms such as ACR. Users also expressed concern about privacy options being buried in deep settings menus, highlighting the importance of understanding the effectiveness and accessibility of ACR opt-out mechanisms.

Inspired by the second-party ACR perspective introduced by Anselmi et al. [8], we conduct a broader study of ACR deployments across eight smart-TV OSes. Beyond expanding OS coverage, we characterize previously unexplored properties of ACR implementations, evaluate ACR behavior across new viewing scenarios, and identify new privacy and security issues. By providing technical visibility into ACR behavior, disclosures, and opt-out mechanisms across the vast majority of the U.S. smart-TV market, our work also helps address the transparency and user-awareness challenges identified by Wang et al. [56]. A detailed claim-by-claim comparison with the most closely related work, Anselmi et al. [8], is reported in Table 1.

2.2 How ACR Works

ACR systems sample the TV content, compute a fingerprint, and transmit it to a remote server that matches the fingerprint against a content library (Figure 1). *Audio-based ACR* samples the audio stream and computes an audio fingerprint [25], following the family

Table 1: Claim-by-claim comparison against the only prior ACR-specific audit, Anselmi et al. [8]. *Extend*: we replicate and add new evidence; *Refine*: prior result is qualitatively correct but with exceptions; *New*: no analogue in prior work.

Analysis	Anselmi 2024	Our work	Status
OSes audited	Samsung Tizen, LG webOS (2 platforms).	Tizen, webOS, Roku, Fire OS, Google TV, SmartCast, Xumo, VIDAA (8 platforms).	Extend
ACR endpoints	acr-us-prd.samsungcloud.tv(Tizen); tkacrX.alphonso.tv(webOS).	Replicated on Tizen and webOS. Additionally identified endpoints for Roku, Sony Google TV, SmartCast, Fire OS, and VIDAA.	Extend
Opt-out completeness	Fully suppresses ACR on both platforms.	Fully effective on webOS, Sony Google TV, Fire OS, and VIDAA. Suspected ACR activity during opt out in more than 60% of measurements on Tizen, 100% on Vizio, and 100% Roku first-party measurements.	Refine
ACR during first-party streaming	Observed on Samsung TV Plus and LG Channels.	Observed on all platforms except Sony Google TV and Fire OS; persists after opt-out on Tizen, Roku, and SmartCast.	Refine
ACR during HDMI	Observed on Tizen and webOS.	Replicated; also observed on Sony Google TV, Fire OS, and VIDAA.	Extend
Sample modality	Audio vs. video contribution not isolated.	Antenna experiments identify audio fingerprinting on webOS & VIDAA, video fingerprinting on Tizen & Sony.	New
ACR during third-party streaming	Not observed during Netflix and YouTube on either platform.	Replicated; additionally observe ACR during Tubi on Tizen, Sony Google TV, and VIDAA, indicating per-app ACR policy differences.	New
ACR during screen mirroring	Not observed on either platform.	Replicated on Tizen and webOS; observed on VIDAA.	Extend
Vulnerability disclosure	None reported.	Authentication weakness in VIDAA’s MQTT stack permits passive reconstruction of all connected devices’ viewing-event streams; disclosed under 90-day policy.	New
Cleartext viewing-history exposure	Not analysed.	Program titles emitted in plaintext DNS queries and URL paths to Amagi (*.amg*.amagi.tv), leaking viewing history on-path independent of ACR opt-out.	New


of techniques established by Wang’s Shazam algorithm [54, 55] and Haitsma-Kalker fingerprints [18]; *Video-based ACR* samples the frame buffer; Samsung documents fingerprinting at two samples per second on Tizen [39], and Chen et al. [9] surveys the underlying techniques.

A matched viewing event is the unit of commercial value extracted by ACR. Vendors and their ACR partners monetize these events through advertising attribution and audience targeting, as acknowledged in their privacy policies and Securities and Exchange Commission (SEC) filings [38, 52].

Unlike traditional web and mobile tracking, which observe user activity within a particular application or service, ACR operates at the television level and may observe content delivered through multiple input paths, including streaming applications, HDMI-connected devices, antenna broadcasts, and screen mirroring. Consequently, ACR can monitor viewing activity independently of the application that delivers the content.

2.3 Smart-TV OSes and ACR Deployments

ACR is implemented by the smart-TV OS vendor or its contracted ACR provider. As a result, ACR behavior is determined by the OS and its associated ACR implementation rather than by individual apps running on the TV. Thus, the same Netflix app may be subject to different ACR behavior on different OSes. In lieu of app-specific controls, TV OSes typically offer only a system-level opt-out, rather than app- or modality-specific settings. For most smart-TV OSes,

**Figure 1: ACR system architecture: fingerprint capture, transmission, and content detection.**

the OS vendor (or its contracted provider) uses an identical ACR pipeline across devices running that OS. This means that if we analyze one TV with a given OS, we expect opt-out controls, privacy disclosures, security posture, and partner relationships to be the same as all devices running that OS.

However, Google TV is the exception. Google licenses its OS to multiple original equipment manufacturers (OEMs) without an ACR pipeline, and OEMs may add ACR as an optional customization.

In this paper, we treat the OS as the primary unit of analysis, except for Google TV where the OS/OEM pair is the relevant unit.

2.4 Network Measurements

Network measurement is a well-established methodology for studying privacy-related system behavior in locked-down environments where rooting or jailbreaking is infeasible [12, 24, 35]. Smart TVs transmit captured audio and visual fingerprints to remote ACR servers over the network. These ACR server destinations can be identified by systematically analyzing network endpoints contacted by the TVs, using destination Internet Protocol (IP) addresses, Domain Name System (DNS) requests, and Server Name Indication (SNI) fields from captured network packets. However, applications typically contact numerous fully qualified domain names (FQDNs) during operation, resulting in a large number of domains contacted by the TVs. For example, Netflix traffic includes both `api-global.netflix.com` and `eu-west-1.alb.pulse.perfmetrics.netflix.com`. To manage this complexity, domains can be grouped by second-level domain (SLD). For instance, Netflix FQDNs can be aggregated under `netflix.com`. Conversely, applications frequently use non-descriptive domain names, necessitating network traffic analysis based on data transfer volume, frequency, and other properties to identify network endpoints belonging to each application.

3 Research Questions

Smart TVs are effectively black-box platforms for users and researchers. While rooting, jailbreaking, and reverse engineering can provide valuable insight into specific devices, these techniques typically depend on device-specific vulnerabilities, firmware versions, and implementation details that may vary across products and change over time [8]. In this work, we instead adopt a network-based measurement approach that can be applied consistently across smart-TV OSes, manufacturers, and device models without requiring privileged access. Within this context, the goal of this paper is to answer the following research questions.

RQ1: Under what TV-viewing conditions is ACR-attributable traffic observed? ACR disclosures by TV vendors often provide limited information about when content monitoring occurs. Understanding which viewing activities are subject to ACR is necessary to determine the actual scope of data collection and its privacy impact. Our scenarios include first-party streaming apps, third-party streaming apps across three different categories, antenna, HDMI content, and screen mirroring. Using network measurements, we classify which scenarios exhibit ACR-attributable traffic and compare our results against each vendor’s privacy policy disclosures to assess whether the disclosed information accurately reflects the observed ACR behavior.

RQ2: To what extent are ACR opt-outs offered and effective? Although ACR-implementing OSes provide opt-out controls for at least some viewing scenarios, the scope and effectiveness of these controls are not well understood. Furthermore, some OSes do not allow users to opt out of ACR for first-party streaming. This research question investigates available opt-out controls and whether those controls suppress ACR-attributable traffic.

RQ3: What TV-viewing information exposure is associated with ACR? ACR is designed to identify viewed content, but vendors provide limited visibility into what TV-viewing information

Table 2: Smart TV OSes and their market share in the U.S. [33]. There is no reliable source providing the individual U.S. market share of VIDAA and Xumo OSes, and [33] provides combined market share information.

Operating System	U.S. Market Share	TV Manufacturer(s)
Tizen OS	34%	Samsung
Roku OS	18%	Roku, TCL
WebOS	15%	LG
SmartCast OS	11%	Vizio
Google TV OS	8%	2 x Sony, Hisense, TCL
Fire TV OS	5%	Amazon Fire
Others (including VIDAA, Xumo)	4%	Supersonic, Westinghouse

is exposed, which entities can observe it, and the privacy implications of that exposure. This research question investigates the nature of TV-viewing information exposure associated with ACR and the entities that can observe it. We also examine whether other smart-TV infrastructure exposes information that enables external entities to infer viewed content, even when that information is not transmitted through ACR itself.

4 Methodology

Our methodology consists of two steps: data collection and ACR endpoint identification. In this section, we describe our testbed and data collection infrastructure and explain our approach to answer the research questions in §3. In §5, we use empirical observations to develop strategies for identifying ACR endpoints, then we evaluate those strategies on the collected data.

4.1 Testbed and Data Collection Infrastructure

We built a testbed of 12 smart TVs spanning the eight dominant smart-TV OSes in the U.S. market (Table 2), covering approximately 95% of the U.S. smart-TV OS market [33], and more than 80% of TV brands in U.S. households with a market share above 3% [16]. Since the OS is the primary unit of analysis (§2.3), our selection prioritized OS coverage. Multiple TVs were included for selected OSes when that OS is deployed by multiple vendors. We captured all network traffic to and from each TV using Mon(IoT)r software [29] and SPHERE Research Infrastructure [45] (Figure 2). Collectively, we conducted 5,943 controlled experiments on TVs, demonstrating that our approach addresses the coverage limitation for prior work in this space.

4.2 Viewing-Scenario Experiments (RQ1)

To answer RQ1 (*Under what TV-viewing conditions is ACR-attributable traffic observed?*), we designed experiments covering common smart-TV viewing scenarios [40] (Table 3). We designed controlled sub-scenarios using TVs’ antenna input. In the “Channel” sub-scenario, we use an RF modulator to deliver custom broadcast TV content with both dynamic video and audio. In “Black screen + audio,” we deliver simulated content with dynamic audio but a consistent black screen. In “Dynamic content + no audio,” we deliver simulated content with dynamic video but no audio. By varying audio and visual content across these scenarios, we can differentiate ACR implementations that rely on audio fingerprints [25] from those using video fingerprints [39], and provide supporting evidence for

ACR endpoint validation (§5). We use the antenna input because broadcast television content delivered through the antenna input cannot be observed through an individual application’s tracking mechanisms, making it a natural ACR target [21]. We use HDMI-to-RF digital modulators to deliver custom content over the antenna input, and IP-to-HDMI converters receiving RTSP streams from the control machine to provide HDMI input to all TVs (Figure 2). An *Idle* scenario serves as a control condition to capture any baseline traffic differences between active and inactive TV usage.

We developed control scripts using OS-specific Python libraries for six OSes (provided in Appendix A, Table 6). Xumo OS and VIDAA OS, which lacked efficient remote control libraries, were controlled via their Android remote controller apps using ADB. The key challenge in automated control was that TV user interfaces are dynamic, requiring OS-specific command sequences that reliably launch the intended experiment. Once these sequences were established for each TV, automated tests ran successfully with few failures. We recorded experiments using IP cameras attached to each TV, extracted the corresponding frames, and reviewed them to confirm experiment success. Failed experiments due to Internet outages or unexpected TV behavior were repeated.

Since each OS supports a specific screen mirroring protocol and some protocols such as AirPlay have limited automation support due to iOS restrictions, we used Google Pixel phones and a Windows laptop with Pigeoncast [34] and Miracast [57], which collectively support the native screen mirroring protocol of each TV. Since VIDAA OS uses unencrypted ACR traffic, we can directly inspect the transmitted payload and confidently answer RQ1 without relying on statistical analysis over a large number of measurements. Therefore, three manual measurements were sufficient to determine whether screen mirroring triggered ACR on VIDAA OS.

For each viewing scenario, we use the ACR endpoint identification approach described in §5 to identify the endpoints associated with ACR. We then compare traffic observed at those endpoints across scenarios to determine which TV-viewing conditions exhibit ACR-attributable traffic.

4.3 Opt-Out Experiments (RQ2)

To answer RQ2 (*To what extent are ACR opt-outs offered and effective?*), we repeat the viewing scenario experiments described in §4.2 in the two ACR states: opted-in (ACR enabled) and opted-out (ACR disabled through the TV settings). To reduce the impact of transient failures and increase confidence in our results, we run each (scenario, ACR-state) experiment 30 times, following the methodology of a prior large-scale IoT measurement study [26]. To prevent ordering effects, we randomly shuffle both the sequence of scenarios and the ACR states, which reduces the risk of artifacts that a fixed order might introduce (e.g., delayed ACR reporting from one scenario always being flagged as ACR for a subsequent scenario).

For each scenario and ACR state, we identify the ACR endpoint using the methodology in §5. We then compare traffic observed at those endpoints between the opted-in and opted-out states to determine if opt-out controls suppress ACR-attributable traffic.

Table 3: Experimental scenarios and their descriptions.

Scenario	Sub-scenario	Description
Idle	–	Home screen, no interaction.
Antenna	Channel	Custom channel with dynamic visual and audio content.
	Black screen + audio	Constant black frame with valid audio over RF.
	Dynamic + no audio	Dynamic video content with no audio over RF.
First-party streaming	–	OS-provided free streaming service.
Third-party streaming	Netflix	Subscription-only service.
	Tubi TV	Free ad-supported service.
	YouTube	Hybrid service.
HDMI	–	Dynamic video with audio via HDMI from an external source.
Screen mirroring	–	External device screen mirrored via native protocol playing dynamic video with valid audio.

4.4 Viewing-Information Exposure (RQ3)

To answer RQ3 (*What TV-viewing information exposure is associated with ACR?*), we analyze the network traces collected during the viewing-scenario experiments described in §4.2.

We first examine traffic associated with the ACR endpoints identified in §5. Where traffic is unencrypted, we directly inspect the payload to characterize the transmitted viewing-related information and metadata. Where traffic is encrypted, we rely on observable network signals such as traffic patterns, timing characteristics, destination domains, and protocol metadata to infer the nature of the exchanged information. We additionally investigate whether smart-TV infrastructure associated with content delivery, such as the content-delivery services used by first-party streaming platforms, exposes viewing information through channels beyond the identified ACR endpoints. Using these observations, we characterize the TV-viewing information exposed through ACR and related infrastructure, the parties that receive or can observe that information, and the resulting privacy implications.

5 ACR Endpoint Identification

In this section, we present our approach for identifying and validating ACR endpoints and present the endpoints identified across all smart-TV OSes in our testbed (Figure 3).

5.1 ACR Endpoint Selection Steps

VIDAA OS (Supersonic TV) exposes unencrypted ACR traffic, while Tizen OS (Samsung TV) and webOS (LG TV) expose self-disclosing ACR domains containing keywords such as “acr” and ACR provider names (e.g., “alphonso,” “unruly”) [8]. However, the remaining OSes do not follow this convention, so we need a way to systematically identify ACR endpoints in these cases. We use the self-disclosing and unencrypted ACR domains as our reference cases to derive hypotheses about ACR behavior, which guide the filtering steps described below. We apply these steps to the network traffic captured from each TV to identify candidate ACR endpoints. Once potential ACR endpoints are identified, we validate them using




Figure 2: Data collection infrastructure. The orchestrator dispatches per-OS driver scripts over Wi-Fi while Mon(Io)Tr and SPHERE Research Infrastructure software captures all TV-originated traffic at a gateway. An RTSP-to-HDMI bridge supplies controlled HDMI content; an HDMI-to-RF modulator supplies controlled antenna stimuli; per-TV IP cameras record on-screen state for per-trial validation.

two traffic-based hypotheses and additional evidence derived from vendor documentation, patents, privacy policies, and evidence from the reference cases (§5.2).

IP-to-domain name mapping. We start by identifying all IP addresses that receive traffic from each TV and map them to domain names using DNS and SNI fields. IP addresses that do not map to a domain name are grouped by IP prefix and added to the candidate list. We additionally annotate IP addresses with their prefix owner information, which serves as a substitute for domain names for hardcoded IP addresses during subsequent filtering and validation steps.

Filter certain well-known service domains. We remove domains associated with well-known streaming services, advertising platforms, and content delivery infrastructure using a set of keywords provided in Appendix B (Table 7), such as “netflix” and “googlevideo.”

Group by SLD and filter by opt-in threshold. From our reference cases (self-disclosing and unencrypted ACR endpoints for Tizen OS, webOS, and VIDAA OS), we observed that known ACR services may use multiple FQDNs, while the corresponding second-level domain (SLD) remains stable across measurements. We therefore group the remaining domains by SLD. For example, webOS has 41 ACR-related FQDNs for Alphonso (such as tkacr719.alphonso.tv and tkacr725.alphonso.tv), all grouped under alphonso.tv. We evaluated the presence of these reference ACR endpoints in viewing scenarios subject to tracking and found that they receive traffic in at least 26 out of 30 measurements. Based on this observation, we retain only SLDs appearing in at least 80% of opted-in measurements (24 out of 30).

Filter cross-OS endpoints. From our reference cases, we observed that ACR services are OS-specific. We therefore remove domains that appear across multiple OSes, as they are more likely to correspond to shared services than OS-specific ACR functionality. For example, adrise.tv is an ad tech platform appearing in network traces from all smart-TV OSes, because it manages advertising for Tubi, a third-party streaming service in our test scenarios. We retain shared cloud infrastructure domains such as “amazonaws.com” and “amazon.com,” which may host OS-specific services. This heuristic

may not hold when the same ACR provider is deployed across multiple OSes; however, none of the OSes in our dataset share an ACR provider.

Filter by traffic direction and ratio. Since ACR frequently transmits audio/video fingerprints from the smart TV to remote servers, we hypothesize that ACR endpoints should receive significantly more traffic from the TV than they transmit back. This hypothesis is supported by our reference cases, where the identified ACR endpoints on Tizen OS, webOS, and VIDAA OS receive approximately 2x, 7x, and 7x more traffic from the TV than they transmit back, respectively (Figure 4). We therefore convert the remaining second-level domains (SLDs) back into fully qualified domain names to prevent artifacts caused by grouping, and remove endpoints with higher download than upload traffic from the candidate list.

Filter by traffic periodicity. Since ACR data is used for real-time targeted advertising [39], we expect frequent, regular transmissions to the endpoint to make targeted advertising actionable before the viewer switches to different content. Using a Discrete Fourier Transform (DFT)- and autocorrelation-based periodicity analysis [19], we confirmed that ACR endpoints in our reference cases (VIDAA OS, Tizen OS, and webOS) receive bursty periodic traffic every 7 seconds, 1 minute, and 14 seconds, respectively (Figure 5), consistent with prior work [8]. While these observations show ACR traffic is generally quite frequent, we choose a conservative (i.e., inclusive) threshold of 5 minutes that might include false positives filtered out by other steps, but is less likely to have false negatives if other ACR providers use periods greater than a minute.

5.2 ACR Endpoint Validation

We use two traffic-based indicators as supporting evidence for each candidate ACR endpoint. When neither indicator is observed, we supplement our analysis with manufacturer documentation, patents, privacy policies, and other public sources, which may provide additional support for endpoint validation.

Validate Content-Based Traffic Shape Since some OSes use audio-based ACR and others use video-based ACR [25, 39], we expect traffic to vary across antenna sub-scenarios. Specifically, scenarios with both audio and video should produce the highest traffic; scenarios with dynamic video but no audio should show a decrease for audio-based ACR implementations; and scenarios with audio but a black screen should show a decrease for video-based ACR implementations.

We therefore examine whether traffic volume responds to the available audio and video content across these sub-scenarios. Endpoints whose traffic varies with content provide supporting evidence that the endpoint is associated with ACR and may enable inference of the ACR modality. Endpoints whose traffic remains invariant to content provide weaker evidence, as some ACR implementations may transmit fixed-size records irrespective of the underlying content.

Opt-Out Traffic Analysis When ACR is opted out, an ACR endpoint should stop receiving viewing-related data from the smart TV. We therefore use the disappearance of traffic to a candidate ACR endpoint in the opted-out state as supporting evidence for ACR endpoint validation. However, two caveats apply. First, an endpoint may also carry non-viewing traffic, such as configuration




Figure 3: ACR endpoint identification and validation. The identification stage filters each TV’s destination set to candidate endpoints through well-known service domain removal, opt-in-frequency threshold, cross-OS filtering, and traffic-direction and periodicity tests. The validation stage corroborates each candidate against content-modality response, opt-out behavior, and OS-specific evidence.




Figure 4: CDF of upload-to-download traffic ratios across domains contacted by the reference smart-TV OSes (Tizen, webOS, and VIDAA OS) after cross-OS filtering step (§5.1), with the identified ACR endpoint annotated for each OS. The x-axis shows the ratio of total bytes uploaded from the smart TV to total bytes downloaded from each endpoint; the dashed vertical line marks a ratio of 1.0 (balanced traffic). The y-axis shows the cumulative fraction of domains with a ratio $\leq x$. ACR endpoints consistently fall to the right of most other domains, confirming that they receive significantly more traffic than they transmit back, consistent with our traffic direction hypothesis.

or telemetry data, even when ACR is opted out, in which case some residual traffic may remain after ACR-related traffic disappears. Second, traffic may not change after opt-out if the implementation does not fully suppress ACR-attributable traffic. We account for both cases in our analysis.

OS-Specific Evidence Analysis We examine OS-specific documentation, including SEC filings, ACR-related patents, privacy policies, and advertising documentation, as additional documentary evidence. Where neither traffic-based indicator is observed, this evidence becomes the primary basis for validation.

5.3 Identified ACR Endpoints

We present the ACR endpoints identified for each OS and the evidence supporting each identification.

Tizen OS (Samsung TV). After applying our methodology, one domain remained: `acr-us-prd.samsungcloud.tv` (Table 4), consistent with prior work [8]. Although the domain name suggests ACR usage (it contains the keyword “acr”), we use the content-dependent traffic shape as supporting evidence. Samsung’s documentation states that Tizen OS fingerprints visual content twice per second [39]. Figure 7(a) shows the fraction of opted-in measurements per TV and scenario in which traffic to the identified ACR endpoint was observed, and Figure 7(b) shows the same for the opted-out state. Consistent with Samsung’s video-based ACR modality, there is no traffic to this endpoint in all 30 “Antenna Black + Audio” scenario measurements, in both ACR states.

webOS (LG TV). Our methodology identifies `tkacrXXX.alphonso.tv` (where XXX is a number, such as 324) as the single ACR endpoint for webOS (Table 4), consistent with prior work [8]. The domain name provides additional support, as Alphonso is the ACR-provider company acquired by LG [1]. Alphonso uses audio-based ACR [25]. Consistent with this, “Antenna Dynamic No Audio” shows a mean traffic reduction of 97.3% ($\sigma=11.1\%$), while “Antenna Black + Audio” shows no decrease (Figure 6 (a)). There is no traffic to this endpoint in all opted-out measurements (Figure 7(b)).

Roku OS (Roku TV and TCL TV). Our methodology identifies `scribe.logs.roku.com` as the ACR endpoint for both Roku OS TVs (Table 4), a domain previously suspected by users reporting high traffic to it [11]. We further support this identification using Roku’s documentation. The TCL TV opt-out settings explicitly state: “This choice won’t affect info collected via Roku’s streaming service.” Consistent with this, Roku’s privacy policy [36] identifies only connected devices and antenna inputs as subject to ACR opt-out, with no mention of streaming services. Traffic to this endpoint persists for first-party streaming after opt-out on both TVs, while there is no traffic for antenna sub-scenarios in all opted-out measurements on Roku TV and 90% on TCL TV (Figure 7).

Sony Google TV (Samba TV ACR). Our methodology identifies one hardcoded IP address (216.183.117.80) as a candidate ACR endpoint (Table 4). WHOIS registration data confirms that the corresponding prefix belongs to Samba TV. We further validate this as a potential ACR endpoint using both traffic-based indicators (§5.2). “Antenna Dynamic No Audio” produces high periodic traffic, while “Antenna Black + Audio” results in a mean decrease of 52.9% ($\sigma=7.9\%$) (Figure 6 (b)) and there is no traffic in all opted-out measurements (Figure 7(b)).




Figure 5: Autocorrelation of upload traffic for identified ACR endpoints. The x-axis shows the lag in seconds and the y-axis shows the normalized autocorrelation coefficient; a strong peak at lag T indicates that the traffic repeats with period T . Dashed vertical lines mark the top-3 dominant periods detected per OS. All endpoints exhibit clear periodic behavior, with dominant periods ranging from 7 seconds (VIDAA OS) to approximately 135 seconds (Sony Google TV). This periodicity supports our hypothesis that ACR implementations transmit fingerprints at regular intervals for real-time content identification (§5.1).

Table 4: Number of candidate endpoints remaining after each analysis step per smart TV OS. After all filtering steps, one candidate ACR endpoint remains for Tizen, WebOS, Google TV (Sony Bravia), SmartCast, VIDAA, Roku, and Fire OS. Xumo OS, Sony W830K, and Hisense Google TV have no remaining candidate endpoints. TCL Google TV has one remaining endpoint, which is classified as a non-ACR endpoint as neither validation hypothesis holds (§5.2).

Step	Tizen	webOS	Roku (Roku)	Roku (TCL)	Google (Sony)	Fire	SmartCast	VIDAA	Xumo	Google (Sony W830K)	Google (Hisense)	Google (TCL)
Initial set	579	753	846	777	578	623	739	541	429	566	495	453
Filter services	307	390	551	465	368	389	434	378	220	309	281	217
Opt-in filter	26	23	12	3	14	4	24	13	22	17	11	15
Cross-OS filter	12	12	3	3	5	4	11	5	10	1	1	5
Traffic ratio	2	3	3	3	1	1	3	1	2	0	0	2
Periodicity	1	1	1	1	1	1	1	1	0	0	0	1

Fire OS (Fire TV). Our methodology identifies gateway-ink.amazon.com as the ACR endpoint (Table 4). We validate this identification using the opt-out traffic analysis (§5.2): there is no traffic to this endpoint for all scenarios in the opted-out state (Figure 7 (b)).

SmartCast OS (Vizio TV). Our methodology identifies kinesis.us-west-2.amazonaws.com as the endpoint most consistent with ACR behavior for SmartCast OS (Table 4). Although we did not observe either traffic-based indicator described in §5.2, several sources of documentary evidence support this identification. Vizio’s SEC filings state that certain third-party content providers, including Netflix and Disney+, restrict Vizio from collecting ACR data from their apps [52], and that screen mirroring is not subject to ACR. Consistent with this, traffic to the Kinesis domain disappears during third-party app and screen mirroring experiments. Furthermore, the endpoint corresponds to Amazon Kinesis Data Streams, a service used to collect and process large streams of data records in real time [5]. In an interview with an AWS representative, a Vizio

representative states that Amazon Kinesis Data Streams is used to collect logs, performance data, and user data [4], which could include ACR data.

VIDAA OS (Supersonic TV). VIDAA OS partners with Nexxen via its sub-service Unruly for ACR [49]. Our methodology identifies xx.f.p.us.acr.unruly.co (where xx is a number, e.g., 16) as the ACR endpoint (Table 4). Uniquely among all tested OSes, this endpoint transmits ACR data in plaintext, enabling direct payload inspection. The payload confirms ACR data, including audio fingerprints, client identifiers, app names, and input sources, representing to the best of our knowledge the first payload-level analysis of ACR data. Both traffic-based validation indicators (§5.2) hold for this endpoint.

Xumo OS (Westinghouse TV). Xumo OS has no ACR settings and no public ACR disclosures. Applying our identification methodology yielded no candidate endpoints (Table 4). We manually reviewed network traffic to all endpoints for a subset of measurements and




Figure 6: CDF of Antenna sub-scenarios on webOS (LG TV) and Sony Google TV. The x-axis shows the upload traffic rate to the ACR endpoint; the y-axis shows cumulative fraction of seconds with traffic rate $\leq x$. webOS employs audio-based ACR: “Antenna Dynamic No Audio” shows a 97.3% mean traffic reduction, while “Antenna Black + Audio” shows no decrease. Samba TV ACR on Sony Google TV employs video-based ACR: “Antenna Dynamic No Audio” produces high periodic traffic, whereas “Antenna Black + Audio” results in a 52.9% mean reduction.

found no evidence of ACR-related traffic. The manufacturer further confirmed that ACR is not implemented.

Sony, Hisense, TCL Google TV (No ACR). Sony provides a list of ACR-enabled models [43]; however, all listed models are outdated. We obtained Sony model W830K, which had no ACR-related settings. Our methodology found no candidate ACR endpoint for this model (Table 4). Sony further confirmed that ACR is implemented only on Bravia N models (e.g., Bravia 2), not on other models.

Similarly, the Hisense Google TV had no ACR settings and no candidate ACR endpoints were identified (Table 4).

The TCL Google TV also had no ACR settings. Although our methodology produced one candidate endpoint (firehose.us-east-1.amazonaws.com), neither validation indicator (§5.2) was observed, and the endpoint appeared only during first-party streaming. We therefore consider this domain unlikely to be an ACR endpoint.




Figure 7: Share of measurements (out of 30) in which the identified ACR endpoint received content-correlated traffic, per TV and scenario, in the opted-in (top) and opted-out (bottom) states. The opted-in panel shows the bimodality our classification thresholds rely on: scenarios cluster near 100% or near 0% with no ambiguous cases. The opted-out panel shows opt-out effectiveness: cells that remain dark indicate persistent traffic after opt-out. Plaintext ACR data is accessible on VIDAA OS and a single-measurement analysis is sufficient to characterize ACR-subject scenarios.

6 Results

We present the results of our evaluation across the smart-TV OSeS for which we identified and validated an ACR endpoint. We identify the viewing scenarios in which ACR-attributable traffic is observed, evaluate the effectiveness of ACR opt-out mechanisms, and characterize the viewing information exposed through ACR and related infrastructure. Unless otherwise stated, in this section we use “ACR tracking” as shorthand for scenarios classified as exhibiting ACR-attributable traffic by our methodology and for opted-out measurements classified as exhibiting traffic consistent with ACR activity.

VIDAA OS is the exception, where plaintext payloads allow direct inspection of ACR data.

6.1 Observed ACR-Attributable Traffic (RQ1)

We classify a scenario as exhibiting ACR-attributable traffic if consistent periodic outgoing traffic to the ACR endpoint is observed in more than 80% of opted-in measurements, and as not exhibiting ACR-attributable traffic if observed in fewer than 30% of measurements. These thresholds are derived from our observations on the reference ACR endpoints for Tizen OS, webOS, and VIDAA OS (§5.1), consistent with prior work [8]. Any scenario falling between these thresholds would be flagged as inconclusive; however, no such cases were observed. In practice, our results show a clear bimodal distribution (Figure 7): scenarios exhibiting ACR-attributable traffic cluster near 100% of measurements, while scenarios not exhibiting ACR-attributable traffic cluster near 0%.

Figure 7 (a) presents a binary classification of scenarios with and without ACR activity across smart-TV OSes. Figure 7 (b) shows where ACR opt-out is effective or ineffective. In this section, we highlight notable findings for each smart-TV OS.

Tizen OS (Samsung TV). We confirm the absence of ACR tracking for Netflix and YouTube reported by prior work. In addition, we find that Tubi, a free ad-supported streaming service, not previously studied, is subject to ACR tracking in 29 out of 30 measurements (Figure 7).

webOS (LG TV). Unlike Tizen OS, the ACR endpoint continues to receive traffic in scenarios classified as not subject to ACR, showing a mean traffic reduction of 95% relative to tracked scenarios. This pattern aligns with the no-audio antenna sub-scenario. This residual traffic suggests that the endpoint may also support functions beyond ACR, although we cannot determine their purpose from encrypted traffic alone.

Roku OS (Roku TV and TCL TV). Unlike the other evaluated OSes, Roku OS is the only one that does not exhibit ACR-attributable traffic for HDMI. Similar to webOS, however, the identified ACR endpoint continues to receive traffic in non-tracked scenarios. Relative to the tracked scenarios, traffic decreased by 90% for most non-tracked scenarios and by 69% for Netflix. We cannot determine the purpose of this residual traffic from encrypted traffic alone.

SmartCast OS (Vizio TV). Unlike the other evaluated OSes, the identified ACR endpoint also exhibits consistent ACR activity in the Idle scenario (100% of measurements), although its purpose remains unclear.

VIDAA OS (Supersonic TV). Because the identified ACR endpoint is contacted in all scenarios and we have access to plaintext ACR data rather than endpoint activity alone, we directly examine the transmitted data to determine whether ACR tracking is active. Antenna sub-scenarios with audio, the first-party streaming service (VIDAA Channels), HDMI, Tubi, and screen mirroring are subject to ACR tracking, while the other scenarios are not. Although the ACR endpoint is contacted in all scenarios, ACR data is transferred only in the scenarios that we classify as subject to ACR tracking (Table 5). In the remaining scenarios, we observe only MQTT broker subscription traffic, with no viewing-related metadata being

published. This produces a substantial traffic-volume reduction relative to tracked scenarios, with residual traffic accounting for less than 5% of the tracked-scenario traffic volume.

Takeaways All analyzed smart TV OSes exhibit ACR-attributable traffic during antenna scenarios. Although most OSes also exhibit ACR-attributable traffic during first-party streaming, Fire OS and Sony Google TV do not. Consistent with reported agreements [52], no analyzed smart TV OS exhibits ACR-attributable traffic for Netflix. Although no public source discloses whether similar agreements exist for other third-party streaming apps, no analyzed smart TV OS exhibits ACR-attributable traffic for YouTube. Tubi exhibits ACR-attributable traffic on Tizen OS, Sony Google TV, and VIDAA OS. While ACR-attributable traffic for HDMI varies across OSes, screen mirroring is subject to ACR tracking only on VIDAA OS.

6.2 ACR Opt-Out Effectiveness (RQ2)

We evaluate the effectiveness of ACR opt-out settings using the 30 opted-out measurements collected for each viewing scenario. Because identified ACR endpoints may continue to receive traffic after opt-out that is unrelated to ACR, we first distinguish between residual traffic and traffic consistent with ACR tracking. Empirical observations on VIDAA OS support this distinction: plaintext visibility reveals that the TV periodically subscribes to the ACR broker without publishing any ACR data, which amounts to less than 5% of the ACR data traffic. Based on this VIDAA OS observation, we classify residual traffic as traffic whose volume is less than 5% of that observed in the corresponding opted-in measurement. This threshold is intended to distinguish low-volume background traffic from persistent endpoint traffic consistent with ACR activity. We then classify an opted-out measurement as exhibiting suspected ACR activity when traffic consistent with ACR tracking is present, and as not exhibiting suspected ACR activity otherwise.

Fully Effective ACR Opt-Out Our results in Figure 7 show that opt-out settings on webOS (LG TV), Sony Google TV, VIDAA OS (Supersonic TV), and Fire OS (Fire TV) are fully effective in our measurements: no traffic consistent with ACR activity is present in opted-out measurements across all applicable scenarios.

Suspected ACR After Opt-Out The remaining analyzed smart TV OSes exhibit traffic patterns that we suspect may contain ACR data after opt-out occurs (Figure 7).

Tizen OS (Samsung TV). Traffic consistent with ACR tracking remains present in all 30 first-party streaming measurements, 41 out of 60 antenna measurements, and 18 out of 30 HDMI measurements after opt-out. Traffic consistent with ACR tracking is observed in 6 out of 30 opted-out Tubi measurements. Additionally, traffic consistent with ACR tracking is observed in 9 Idle, 2 Netflix, and 3 YouTube opted-out measurements despite the absence of ACR-attributable traffic in the corresponding opted-in scenarios. We interpret the continued presence of traffic consistent with ACR tracking after opt-out as suspicious, though we cannot confirm the contents of the network traffic (due to encryption), nor whether data is processed or retained on the server side.

Roku OS (Roku TV and TCL TV). ACR opt-out does not apply to the first-party streaming service. TCL TV explicitly discloses this limitation in its settings menu, while Roku TV does not. For antenna

Table 5: ACR activity and disclosure status per scenario and smart TV OS. Each scenario reports two values: Obs (ACR observed: ✓ ACR active, × not active) and Cls (classification against the vendor’s current online privacy policy and ACR-specific notice: D disclosed – policy explicitly enumerates this scenario; U undisclosed – policy silent on this input class; C contradicted – policy excludes this scenario or claims opt-out stops it; em-dash – ACR not observed, no claim to evaluate).

TV	Antenna Channel		Antenna Black+Audio		Antenna Dynamic No Audio		First-party Streaming		HDMI		Netflix		YouTube		Tubi		Screen Mirroring		Idle	
	Obs	Cls	Obs	Cls	Obs	Cls	Obs	Cls	Obs	Cls	Obs	Cls	Obs	Cls	Obs	Cls	Obs	Cls	Obs	Cls
Tizen OS	✓	D	×	—	✓	U	✓	D	✓	U	×	—	×	—	✓	U	×	—	×	—
webOS	✓	D	✓	C*	×	—	✓	D	✓	D	×	—	×	—	×	—	×	—	×	—
Roku OS (Roku TV)	✓	D	✓	D	×	D	✓	C	×	×	×	—	×	—	×	—	×	—	×	—
Roku OS (TCL TV)	✓	D	✓	D	✓	D	✓	C	×	×	×	—	×	—	×	—	×	—	×	—
Sony Google TV	✓	D	✓	D	✓	D	×	—	✓	D	×	—	×	—	✓	C	—	—	×	—
Fire OS	✓	D	✓	U	✓	U	×	—	✓	D	×	—	×	—	×	—	×	—	×	—
SmartCast OS	✓	D	✓	D	✓	D	✓	D	×	—	×	—	×	—	×	—	×	—	✓	D
VIDAA OS	✓	D	✓	D	×	—	✓	D	✓	U	×	—	×	—	✓	D	✓	U	×	—

*Alphonso’s policy states it does not recognize “ambient sounds and other non-commercial audio,” which is in tension with observed fingerprint exfiltration on a black-screen + audio feed. Whether this is contradiction or merely a distinction between *capture* and *match* depends on the level at which we measure ACR activity; we discuss this in §7.

scenarios, ACR opt-out suppresses ACR-attributable traffic on Roku TV, with only 1.7% residual traffic remaining after opt-out. On TCL TV, ACR opt-out suppresses ACR-attributable traffic in 81 out of 90 antenna measurements. However, traffic consistent with ACR tracking remains present in the remaining 9 measurements (Figure 7).

SmartCast OS (Vizio TV). Traffic to the identified Amazon Kinesis endpoint does not disappear or decrease in any opted-out measurement (Figure 7), indicating that the opt-out setting has no observable effect on the traffic sent to this endpoint.

Takeaways We observe no traffic consistent with ACR activity after opt-out on webOS, Sony Google TV, Fire OS, and VIDAA OS. Roku OS provides effective opt-out for most scenarios with ACR-attributable traffic, with the notable exception of the first-party streaming service; this limitation is explicitly disclosed on TCL TV but not on Roku TV. In contrast, Tizen OS exhibits persistent traffic consistent with ACR activity across most opted-out measurements, while all traffic to the identified SmartCast OS ACR endpoint is unaffected by opt-out, indicating that the opt-out mechanism may not be fully effective. For encrypted ACR implementations, these conclusions are based on observed traffic behavior rather than direct inspection of ACR payloads.

6.3 Viewing Information Exposure (RQ3)

We characterize the TV-viewing information exposed through the ACR implementations identified in §5 and evaluate which entities can observe that information. We further investigate whether smart-TV infrastructure outside the identified ACR endpoints exposes viewing information to third parties.

Encrypted Viewing Information

Tizen OS (Samsung TV). Traffic to the ACR endpoint disappears when a black screen with valid audio is displayed, while dynamic content with no audio does not affect the traffic shape. This is consistent with Samsung’s use of video-based ACR, as stated in its advertising documentation [39].

webOS (LG TV). Traffic to the ACR endpoint decreases by 97.3% ($\sigma=11.1\%$) when dynamic content is displayed without audio, compared with scenarios with audio. This is consistent with Alphonso’s audio-based ACR used on LG TVs [25].

Sony Google TV (Samba TV ACR). Traffic to the ACR endpoint decreases by 52.9% ($\sigma=7.9\%$) when a black screen with audio is displayed, consistent with Samba TV’s use of video-based ACR [44]. *Fire OS (Fire TV)*, *Roku OS (Roku TV, TCL TV)*, and *SmartCast OS (Vizio TV)*. The traffic shape does not vary across antenna sub-scenarios, preventing us from inferring whether these ACR implementations rely primarily on audio-based or video-based fingerprinting from network traffic alone.

Unencrypted Viewing Information The VIDAA OS (Supersonic TV) ACR implementation exposes privacy-sensitive viewing data through unencrypted communication and multiple security vulnerabilities, which could allow an adversary to access the viewing data of millions of users globally. The OS uses an unencrypted MQTT protocol to transfer ACR data to Unruly ACR servers. Through this vulnerability, any on-path observer between the smart TV and ACR servers can intercept sensitive personal data. In addition to ACR-related metadata, including timezone, app name, input source (e.g., antenna, HDMI), tuner channel for antenna input, and the audio fingerprint, the Client ID (a non-resettable device-unique identifier) is transmitted in plaintext. Such device identifiers are regularly considered personally identifiable information (PII).

Testing multiple applications revealed several concerning behaviors arising from the unencrypted traffic. First, the first-party app “HOME OS” transmits show or film titles rather than app names in the “App Name” field. Second, the “App Name” field itself can reveal sensitive information. For example, apps such as “VIDAA Kids” and “Christianity and Money” are subject to ACR tracking in plaintext. In the former case, ACR is performed on well-known child-directed content, raising potential COPPA [15] concerns in the US. In the latter case, content on sensitive topics such as religion is explicitly categorized by the ACR servers. More generally, we observed categories including “Religious,” “Health and Medicine,”

and “Political.” Combined with the lack of encryption, this allows an on-path observer to infer potentially sensitive religious, political, and demographic information about smart-TV users.

VIDAA OS claims to use ACR data to enhance audio and video output, and the corresponding setting is named “Enhanced Viewing Options.” Once content is detected by the ACR servers, the TV receives two levels of genre classification (e.g., Genre 1: “Movie,” Genre 2: “Romance”) in plaintext JSON format. We tested multiple sensitive and protected content types over HDMI, including religious (“The Story of Moses”) and children’s (“Johnny Johnny Yes Papa”) content, and confirmed that the ACR servers correctly identified them as “Religious” and “Entertainment,” respectively.

Beyond the unencrypted communication described above, we identified two additional security vulnerabilities that could expose viewing information from other VIDAA OS users or enable unauthorized manipulation of ACR-related communications:

Unauthenticated access to ACR broker. When an MQTT broker lacks authentication, any client can subscribe to a topic and receive messages published by other subscribers. The Unruly MQTT broker lacked such checks, allowing us to subscribe from an author’s laptop and receive ACR detection results from other VIDAA TV users. To assess the severity of this vulnerability, we subscribed for approximately three hours and collected data from more than 1 million unique TVs, resulting in more than 57 million non-empty ACR detection results across 82 categories. Several categories are sensitive, including “Politics,” “Health and Medicine,” and “Religious.” Each record included a unique client identifier transmitted in plaintext. After we disclosed the vulnerability, VIDAA confirmed it, and indicated they would address it. We find that the MQTT broker now requires authentication, but all transmissions to the broker remain in plaintext.

Packet injection vulnerability. We successfully republished a previously observed ACR message from an author’s laptop to the MQTT broker, which was received by our subscribed TV. This demonstrates that an adversary could publish misleading or harmful content to subscribed devices, though for ethical reasons we did not attempt to do so.

Viewing-History Exposure Outside ACR We document a separate, previously unreported channel that exposes viewing information independently of the ACR opt-out state. In the WatchFree+ first-party streaming service on SmartCast OS, the program or channel title is directly embedded into unencrypted DNS queries and SNI fields of the form {title-slug}-{platform-token}.amagi.tv or amgNNNNN-{title-slug}-{platform-token}.amagi.tv, where amgNNNNN is a stable channel identifier assigned by Amagi [3] and {platform-token} identifies the requesting platform (e.g., soul-5mincrafteng-vizio.amagi.tv). Program titles appear in URL-safe-encoded form (spaces hyphenated, punctuation stripped) but are otherwise directly human-readable. Any on-path observer between the television and its recursive resolver can reconstruct an ordered viewing history without access to any ACR payload, fingerprint, or vendor backend.² We note that no OS we audited mentions Amagi [3] in its privacy policy, settings, or opt-out

²We also observed a second URL format in which the title appears in the HTTPS path (e.g., cdn-{region}-prod.tsv2.amagi.tv/linear/{title-slug}/playlist.m3u8). This form is protected by TLS and is therefore not visible to passive on-path observers, although TLS-terminating middleboxes would see it.

controls. Furthermore, users are provided with no UI mechanism to disable or opt out of this behavior.

Takeaways Although encrypted traffic prevents direct inspection of ACR payloads for most OSes, our methodology provides insight into the viewing information exposed by ACR. Tizen OS (Samsung TV) and Sony Google TV (Samba TV ACR) are consistent with video-based fingerprinting, while webOS (LG TV) and VIDAA OS (Supersonic TV) are consistent with audio-based fingerprinting. Additionally, the VIDAA OS ACR implementation exposes important privacy and security vulnerabilities, including unencrypted transmission of privacy-sensitive metadata and unauthenticated access to the ACR broker, exposing the viewing data of its users to unauthorized access. Finally, the plaintext DNS leakage through Amagi, though independent of ACR, shows that even users who fully opt out of ACR remain exposed through the content-delivery path (we discuss mitigations and regulatory implications in §7).

7 Discussion

We now summarize findings with regulatory impact, discuss the associated implications, describe suggested mitigations, and identify limitations and future work.

Disclosure-Behavior Gaps Table 5 shows four recurring gaps.

Undisclosed input paths. Tizen OS, Fire OS, and VIDAA OS policies do not address HDMI input, yet we observe traffic consistent with ACR activity under HDMI on five of eight OSes.

Contradicted opt-out scope. On Tizen OS and SmartCast OS, the opt-out toggle does not fully suppress ACR-attributable traffic. We interpret persistent opted-out traffic as a potentially ineffective opt-out, while acknowledging no visibility into server-side processing.

Ambiguous modality claims. webOS’s policy states that the matcher discards non-commercial audio; our black-screen-with-audio condition produces no traffic decrease. From outside the matcher, we cannot tell whether the policy is consistent with capture-then-discard or contradicted by capture-and-retain behavior.

Outdated model lists. Sony’s published ACR-enabled model list is from 2017 and omits our 2025 N-series test unit, which Sony confirmed by direct contact does implement ACR.

Opt-Out Friction and Dark Patterns Five of eight OSes label the opt-out toggle with vendor-specific phrasing that does not contain “Automatic Content Recognition” (*Viewing Information Services* on Tizen OS, *Live Plus* on webOS, *Samba Interactive TV* on Sony Google TV, *Viewing Data* on SmartCast OS, *Use info from TV inputs* on Roku OS), consistent with the *Interface Interference* dark-pattern category [17].

Regulatory Implications *FTC consent decree.* In 2017, VIZIO paid \$2.2 million to settle FTC charges of undisclosed ACR collection [14], and agreed to ensure a working opt-out for ACR going forward. Our results identify kinesius.us-west-2.amazonaws.com as the SmartCast OS endpoint most consistent with ACR activity. The opt-out setting has no observable effect on traffic sent to this endpoint; however, we cannot inspect the payload due to encryption so we cannot confirm whether this violates the FTC settlement. *Texas Deceptive Trade Practices Act.* The 2025 Texas Attorney General complaint against Samsung, Sony, LG, TCL, and Hisense [31]

alleges ACR-based profiling without adequate consent; our measurements characterize the technical surface against which it was filed. The 2026 Samsung settlement [32] requires explicit opt-in consent, and our methodology can help assess compliance by measuring whether traffic consistent with ACR activity is present.

COPPA. In 2025, Michigan sued Roku over COPPA [15] violations related to ACR [47]. For a related concern (but different TV provider), we found that the VIDAA OS classifies child-directed programming and transmits its details in plaintext, alongside TV-unique identifiers. A practicing lawyer experienced with COPPA litigation confirmed that this activity certainly implicates COPPA (though only a court can decide if there is a violation).

Infrastructure outside the consent surface. Amagi content-delivery domains (§6.3) leak first-party streaming titles in plaintext regardless of ACR opt-out state. Thus, even a perfectly enforced ACR opt-out does not prevent exposure of viewing history when it happens via non-ACR channels. To our knowledge, this is the first documented case of content-delivery infrastructure acting as a viewing-history channel on smart TVs.

Mitigations We identify the following opportunities to mitigate the privacy harms from our findings.

For users. To avoid ACR, disconnect the TV from the network and use an external streaming device that does not perform ACR, or use the ACR opt-out on OSes where it was effective in our measurements (webOS, Fire OS, Sony Google TV, VIDAA OS); use the blocklist we release [22] on remaining OSes. We implemented a prototype blocklist in our lab, and verified that the blocked endpoints do not interfere with smart-TV functionality, nor do they affect quality of experience for the scenarios we tested. However, neither approach addresses the Amagi plaintext-DNS channel (§6.3).

For OS vendors. Enumerate input paths subject to ACR in privacy policies; authenticate and encrypt ACR traffic; and close the Amagi channel by adopting encrypted DNS together with delivery architectures that avoid embedding viewing information in hostnames.

For regulators. Auditing ACR disclosures and opt-out effectiveness may help regulators assess whether deployed behavior aligns with user-facing disclosures and consent mechanisms. Our methodology enables automated collection of compliance evidence to conduct such audits. In addition, regulators may also have the authority to subpoena additional information not accessible to researchers (e.g., data collected and stored at ACR servers).

Limitations and Future Work Except for VIDAA OS, our conclusions about ACR activity are inferred from externally observable traffic behavior rather than direct inspection of ACR payloads.

No payload ground truth on six platforms. We identify traffic attributable to ACR by the convergence of four indicators (traffic ratio, periodicity, content-modality response, opt-out behavior) rather than direct payload inspection. The VIDAA case, where payload visibility is available, provides supporting evidence for this four-indicator approach. GDPR Article 15 or CCPA data-access requests may provide better ground-truth information about data collected from smart-TV viewers; however, such requests are most actionable in jurisdictions with enforceable access rights, such as California. We ran our experiments in Massachusetts, which has no data subject access rights at the time of writing, and our data requests to Tizen and webOS went unanswered.

Cloud endpoint drift. As ACR operators evolve, our specific detection approach and thresholds may lose accuracy, particularly if ACR operators use an adversarial approach. However, evolution may be slow: we note that ACR domains found two years ago remained unchanged in this study.

Partial dark-pattern coverage. A systematic UI-flow audit against a consent-friction taxonomy is follow-up work.

U.S.-only scope. Prior work [8] comparing Tizen and webOS in the U.K. and U.S. found an additional ACR-related domain in the U.K. (unused for data collection), and that first-party streaming was not an ACR-subject scenario there. This suggests that regulatory differences may determine ACR-subject scenarios and ACR behavior under EU privacy regulation (or other jurisdictions) may differ.

VIDAA OS MQTT failures. VIDAA OS's unencrypted plaintext ACR communication allowed us to observe its ACR behavior. However, we found that the OS fails to publish ACR data to the MQTT broker in approximately 50% of measurements after successfully subscribing. While we hypothesize this is a software bug, we lack visibility into the implementation details to confirm.

8 Conclusion

We developed a scalable, automated, and rigorous method for identifying ACR-attributable traffic across a wide range of TV-viewing scenarios. We applied this to 12 TVs to characterize ACR deployments across the eight smart-TV operating systems that account for approximately 95% of the U.S. market. Of the seven OSes for which we identified an ACR endpoint, four showed no evidence of ACR-attributable traffic after opt-out, while three exhibited behavior consistent with ACR activity after opt-out. In addition, we identified an ACR provider that not only exposed ACR traffic (including fingerprints and classifications) in plaintext on one of our TVs, but also ran an unauthenticated MQTT server that exposed classifications for more than a million TVs in a short three-hour evaluation. These findings have been disclosed and acknowledged. Along with other findings (e.g., classification of child-directed content and sensitive topics, exposure of TV-viewing information by non-ACR third parties), our results suggest that many smart TVs present significant privacy concerns. We recommend that regulators, TV OS providers, and ACR providers adopt better standards and practices around this topic.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback. We are also grateful to the SPHERE Research Infrastructure [45] team, Kevin Aldana, and Sagar Mohan for their contributions. This research was partially funded by the NSF under the ProperData Frontier award (SaTC-1955227) and the Ludlow Institute. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsors.

Ethical Considerations

Our experiments on our own devices did not entail human subjects or ethical concerns.

As part of our analysis, we identified an MQTT broker that allowed us to subscribe without authentication. The unauthenticated

broker exposed live viewing-event records to any client; no credential, exploit, or circumvention of an access control was required to receive them. To establish the severity of the exposure for our disclosure, we subscribed for a single bounded window of approximately three hours and recorded only the fields necessary to characterize the leak. We made no attempt to deanonymize any household, to link the per-record client identifiers to real-world identities, or to correlate records across time beyond counting unique identifiers. The collected data were stored encrypted on access-controlled infrastructure, used solely to quantify the exposure and to support disclosure, and will not be redistributed; we report only aggregate statistics. We did not run the collection longer than required to establish the proof of concept, as required by the VIDAA vulnerability disclosure policy.

We privately disclosed the MQTT vulnerabilities that we observed over a month prior to the submission of this manuscript and the report has been acknowledged. We observed that the unauthorized subscription to MQTT broker has been removed, but the ACR data transfer is not yet end-to-end encrypted. We were informed that the mitigation efforts are in the final stages.

Open Science

We are committed to reproducibility and release the following research artifacts via [23]:

- the controlled-testbed network captures (PCAPs) for all TVs (except Supersonic TV) across all scenarios and opt states;
- the smart-TV control scripts;
- the data processing and analysis codes implementing the endpoint identification and four-indicator validation stages;
- the per-TV record of candidate endpoints remaining after each filtering step (the CSVs underlying Table 4), so that each classification is independently auditable;
- the keyword lists used for certain well-known service filtering (Appendix B, Table 7);
- the ACR endpoint blacklist [22];

We deliberately do not release the data collected from the exposed MQTT broker. Those records concern other households, carry per-record identifiers that constitute PII, and were collected only to quantify an exposure under coordinated disclosure; releasing them would re-create the very harm we reported. We instead release the aggregate statistics and category counts reported in §6.3.

AI Use

Generative AI-based tools were used during manuscript preparation for editorial assistance, including language refinement, clarity improvements, consistency checks, editing suggestions, and assistance with manuscript organization and presentation. AI-based tools were also used to assist with software development tasks supporting the research workflow, including implementation of codebase components, data-processing scripts, and generation of tables. We used Claude (Sonnet 5) via its web interface and did not use AI agents. The codebase, including the data-processing pipeline and TV-control algorithms, was developed manually, with AI providing code suggestions we always checked before incorporating, and third-party control libraries were similarly incorporated.

AI-based tools were not used to collect data, determine experimental design, perform scientific analyses, derive conclusions, interpret results, or make research decisions.

All scientific content, methodology, analyses, results, figures, tables, and conclusions were produced and verified by the authors. All AI-generated suggestions were reviewed by the authors before incorporation into the manuscript or codebase. The authors have manually verified and take responsibility for the accuracy, originality, and integrity of all content included in the paper.

References

- [1] Alphonso Inc. 2026. Alphonso: TV Data Company. <https://alphonso.tv/>. Accessed: September 2026.
- [2] Alphonso Inc. 2026. Smart TV Privacy Policy (US). <https://alphonso.tv/privacy/smart-tvs/smart-tv-privacy/>. Governs LG webOS ACR. Last updated May 2026. Accessed: September 2026.
- [3] Amagi. 2026. Amagi: Cloud-Native Broadcast and Streaming TV Technology. <https://www.amagi.com/>. Accessed: September 2026.
- [4] Amazon Web Services. 2020. Vizio and AWS: Building a Data-Driven TV Platform. https://www.youtube.com/watch?v=BLCXEmp_1qY. Accessed: September 2026.
- [5] Amazon Web Services. 2026. Amazon Kinesis Data Streams: Developer Guide. <https://docs.aws.amazon.com/streams/latest/dev/introduction.html>. Accessed: September 2026.
- [6] Amazon.com, Inc. 2026. Automatic Content Recognition (Fire TV). US Customer Service. <https://www.amazon.com/gp/help/customer/display.html?nodeId=GPFFD7F3EVCLUD3C>. Accessed: September 2026.
- [7] Amazon.com, Inc. 2026. Privacy Settings FAQs for Fire TV, Fire tablets and Kindle e-readers. US Customer Service. <https://www.amazon.com/gp/help/customer/display.html?nodeId=GQFYXZHB2H629WN>. Accessed: September 2026.
- [8] Gianluca Anselmi, Yash Vekaria, Alexander D'Souza, Patricia Callejo, Anna Maria Mandalari, and Zubair Shafiq. 2024. Watching TV with the Second-Party: A First Look at Automatic Content Recognition Tracking in Smart TVs. In *Proceedings of the 2024 ACM Internet Measurement Conference (IMC '24)*. ACM, Madrid, Spain. doi:10.1145/3646547.3689013
- [9] Wendi Chen, Wensheng Gan, and Philip S. Yu. 2024. Digital Fingerprinting on Multimedia: A Survey. arXiv preprint arXiv:2408.14155. arXiv:2408.14155 [cs.MM] <https://arxiv.org/abs/2408.14155>. Accessed: September 2026.
- [10] CTVMAO. 2024. CTV Home Screen Open Data Partnerships. <https://ctvmao.org.blogspot.com/2024/06/ctv-home-screen-open-data-partnerships.html>. Accessed: September 2026.
- [11] William Curran. 2024. Why I'm Looking at Roku's Data Practices. <https://www.linkedin.com/pulse/why-im-looking-rokus-data-practices-william-curran-ww81c>. Accessed: September 2026.
- [12] Daniel J. Dubois, Roman Kolcun, Anna Maria Mandalari, Muhammad Talha Paracha, David Choffnes, and Hamed Haddadi. 2020. When Speakers Are All Ears: Characterizing Misactivations of IoT Smart Speakers. In *Proc. of the Privacy Enhancing Technologies Symposium (PETS)*.
- [13] Federal Trade Commission. 2017. FTC, New Jersey Attorney General Settle with TV Maker VIZIO for Tracking Consumers' Viewing Habits Without Consent. <https://www.ftc.gov/news-events/press-releases/2017/02/vizio-pay-22-million-ftc-state-new-jersey-settle-charges-it>
- [14] Federal Trade Commission. 2017. VIZIO to Pay \$2.2 Million to FTC, State of New Jersey to Settle Charges It Collected Viewing Histories on 11 Million Smart Televisions Without Users' Consent.
- [15] Federal Trade Commission. 2026. Children's Online Privacy Protection Rule (COPPA). <https://www.ftc.gov/legal-library/browse/rules/childrens-online-privacy-protection-rule-coppa>. Accessed: September 2026.
- [16] Grand View Research. 2026. U.S. Television Market Brand Share Landscape. <https://www.grandviewresearch.com/market-trends/us-television-market-brand-share-landscape>. Accessed: September 2026.
- [17] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. 2018. The Dark (Patterns) Side of UX Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, Montreal QC, Canada, Article 534, 14 pages. doi:10.1145/3173574.3174108
- [18] Jaap Haitisma and Ton Kalker. 2002. A Highly Robust Audio Fingerprinting System. In *Proceedings of the 3rd International Conference on Music Information Retrieval (ISMIR)*. Paris, France, 107–115.
- [19] Tianrui Hu, Daniel J. Dubois, and David Choffnes. 2023. BehaviorIoT: Measuring Smart Home IoT Behavior Using Network-Inferred Behavior Models. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC '23)*. ACM, Montreal, QC, Canada. doi:10.1145/3618257.3624829

- [20] Hub Entertainment Research. 2025. Connected Home 2025: Smart TV Penetration Reaches 82% of U.S. TV Households. Reported via Marketing Charts. <https://www.marketingcharts.com/industries/media-and-entertainment-235609>
- [21] Owen Hyman. 2025. Automated Content Recognition (ACR), Smart TVs, and Ad-Tech Infrastructure, 2012–2022. *Media, Culture & Society* (2025). doi:10.1177/13548565251327885
- [22] Ilgar Mammadov, Het Rutul Joshi, Daniel J. Dubois, David Choffnes. 2026. ACR Endpoint Blocklist. <https://github.com/IlgarNEU/Watching-the-Watchers/blob/main/Blocklist/Blocklist.txt>. Released ACR endpoint blocklist. Accessed September 2026.
- [23] Ilgar Mammadov, Het Rutul Joshi, Daniel J. Dubois, David Choffnes. 2026. Watching the Watchers: ACR Endpoint Blocklist and Artifacts. <https://github.com/IlgarNEU/Watching-the-Watchers.git>. Artifact repository; includes the ACR endpoint blocklist (Blocklist.txt), per-OS candidate endpoints after each filtering step, TV-control algorithms, the analysis pipeline, and filter keyword lists. Accessed: September 2026.
- [24] Umar Iqbal, Pounch Nikkha Bahrami, Rahmadi Trimnanda, Hao Cui, Alexander Gamero-Garrido, Daniel J. Dubois, David Choffnes, Athina Markopoulou, Franziska Roesner, and Zubair Shafiq. 2023. Tracking, Profiling, and Ad Targeting in the Alexa Echo Smart Speaker Ecosystem. In *ACM Internet Measurement Conference (IMC)*.
- [25] LG Ads. 2026. Site Privacy Policy. <https://lgads.tv/site-privacy-policy/>. Accessed: September 2026.
- [26] Anna Maria Mandalari, Daniel J. Dubois, Roman Kolcun, Muhammad Talha Paracha, Hamed Haddadi, and David Choffnes. 2021. Blocking Without Breaking: Identification and Mitigation of Non-Essential IoT Traffic. *Proceedings on Privacy Enhancing Technologies* 2021, 3 (2021). doi:10.2478/popets-2021-0075
- [27] MNTN. 2025. 2025 Connected TV Statistics: 90% of U.S. Households Use CTV.
- [28] Hooman Mohajeri Moghaddam, Gunes Acar, Ben Burgess, Arunesh Mathur, Danny Yuxing Huang, Nick Feamster, Edward W. Felten, Prateek Mittal, and Arvind Narayanan. 2019. Watching You Watch: The Tracking Ecosystem of Over-the-Top TV Streaming Devices. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*. ACM.
- [29] Mon(IoT)r Lab. 2024. Research artifacts. <https://moniotr-lab.khoury.north-eastern.edu/publications/iot-ipv6/>
- [30] Nielsen. 2025. The Gauge: Streaming Reaches 43.8% of Total TV Viewing.
- [31] Office of the Texas Attorney General. 2025. Attorney General Paxton Sues Five Major TV Companies, Including Some with Ties to the CCP, for Spying on Texans. Filed December 15, 2025.
- [32] Office of the Texas Attorney General. 2026. Attorney General Paxton Secures Major Agreement with Samsung to Ensure that Texans are Protected from Smart TVs Collecting Their Data Without Their Knowledge.
- [33] Parks Associates. 2026. Tech Ecosystem Consumer Insights Dashboard. <https://www.prnewswire.com/news-releases/sixty-one-percent-of-us-in-ternet-households-use-their-smart-tv-as-the-primary-streaming-device-according-to-parks-associates-302653704.html>. Released at CES 2026. Market share figures are derived from the pie chart provided.
- [34] Pigeoncast. 2026. Pigeoncast: Screen Mirroring Software. <https://pigeoncast.com/>. Accessed: September 2026.
- [35] Jingjing Ren, Daniel J. Dubois, David Choffnes, Anna Maria Mandalari, Roman Kolcun, and Hamed Haddadi. 2019. Information Exposure for Consumer IoT Devices: A Multidimensional, Network-Informed Measurement Approach. In *Proc. of the Internet Measurement Conference (IMC)*.
- [36] Roku, Inc. 2026. Roku ACR Service Policy. <https://docs.roku.com/published/acrservicepolicy/en/CA>. Accessed: September 2026.
- [37] Roku, Inc. 2026. Roku ACR Service Policy. <https://docs.roku.com/published/acrservicepolicy/en/CA>. Accessed: September 2026.
- [38] Samsung. 2026. Samsung Privacy Policy. <https://policy.samsung.com/>. Accessed: September 2026.
- [39] Samsung Ads. 2022. Samsung Ads ACR Guide. <https://iabcanada.com/wp-content/uploads/2022/09/Samsung-Ads-ACR-Guide-CA.pdf>. Accessed: September 2026.
- [40] Samsung Ads. 2025. State of CTV Report. https://image-us.samsung.com/SamsungUS/samsungbusiness/samsung-ads/insights/11142025/State_of_CTV_Report_Comms.pdf. Accessed: September 2026.
- [41] Samsung Electronics America. 2026. Samsung Privacy Policy (US). <https://www.samsung.com/us/account/privacy-policy/>. Binding US privacy policy; does not mention ACR. Accessed: September 2026.
- [42] Samsung Electronics America. 2026. Samsung Smart TV Automatic Content Recognition (ACR) Feature. Support FAQ (US). <https://www.samsung.com/us/support/answer/ANS10010616/>. Accessed: September 2026.
- [43] Sony Electronics. 2026. Information About ACR (Automatic Content Recognition) Technology on Sony TVs. <https://www.sony.com/electronics/support/articles/00182856>. Accessed: September 2026.
- [44] Sony Electronics Inc. 2026. Information about Samba TV / Samba Interactive TV. US support article 00182856. <https://www.sony.com/electronics/support/articles/00182856>. Accessed: September 2026.
- [45] SPHERE Project. 2024. SPHERE Research Infrastructure. <https://sphere-project.net/>
- [46] Carlotta Tagliaro et al. 2023. Large-Scale Security and Privacy Analysis of HbbTV. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*.
- [47] The People of the State of Michigan v. Roku, Inc. 2025. The People of the State of Michigan v. Roku, Inc. <http://michigan.gov/ag/-/media/Project/Web/sites/AG/releases/2025/April/Roku-Complaint.pdf>
- [48] Marcos Tileria and Jorge Blasco. 2022. Watch Over Your TV: A Security and Privacy Analysis of the Android TV Ecosystem. *Proceedings on Privacy Enhancing Technologies* (2022).
- [49] Unruly. 2026. Unruly and VIDAA Partnership. <https://unruly.co/vidaa-partnership/>. Accessed: September 2026.
- [50] Janus Varmarken, Hieu Le, Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. 2020. The TV is Smart and Full of Trackers: Measuring Smart TV Advertising and Tracking. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (2020).
- [51] VIDAA USA Inc. 2023. VIDAA OS End User License Agreement (US), Section 10. <https://www.vidaa.com/os-end-user-license-agreement/>. Accessed: September 2026.
- [52] Vizio Holding Corp. 2021. Form S-1 Registration Statement. <https://www.sec.gov/Archives/edgar/data/1835591/000119312521062802/d87723ds1.htm>. Accessed: September 2026.
- [53] VIZIO, Inc. 2026. Viewing Data Supplement to the VIZIO Privacy Policy (US). <https://www.vizio.com/en/terms/privacy-policy/viewing-data-privacy-policy>. Accessed: September 2026.
- [54] Avery Wang. 2006. The Shazam Music Recognition Service. *Commun. ACM* 49, 8 (Aug. 2006), 44–48. doi:10.1145/1145287.1145312
- [55] Avery Li-Chun Wang. 2003. An Industrial-Strength Audio Search Algorithm. In *Proceedings of the 4th International Conference on Music Information Retrieval (ISMIR)*. ISMIR, Baltimore, MD, USA, 7–13.
- [56] Synthia Qia Wang, Alan Gao, Marshini Chetty, and Nick Feamster. 2025. Understanding User Privacy Concerns of Shared Smart TVs. *Proceedings of the ACM on Human-Computer Interaction* (2025). doi:10.1145/3757696
- [57] Wi-Fi Alliance. 2026. Wi-Fi CERTIFIED Miracast Specification. <https://www.wi-fi.org/file/miracast-specification>. Accessed: September 2026.
- [58] Xumo (Comcast Cable Communications, LLC). 2026. Xumo TV privacy features (US support). <https://support.xumo.com/hc/en-us/articles/29032870982548-Xumo-TV-privacy-features>. Updated December 2025. Accessed: September 2026.

Table 6: Control software and protocols used for automated experiment execution per smart TV OS.

Operating System	TV Platform	Control Software	Protocol
Tizen OS	Samsung	samsungtvws	WebSocket
WebOS	LG	aiowebostv	WebSocket
Roku OS	Roku / TCL	aiohhttp	HTTP (ECP)
Google TV	Sony Bravia	Custom implementation	TCP (SSIP)
Fire OS	Amazon Fire TV	adb CLI	ADB over TCP
SmartCast OS	Vizio	Custom implementation	HTTPS (SmartCast API)
Google TV	TCL / Hisense	adb-shell	ADB over TCP

A Data Collection Infrastructure

Table 6 presents the software tools and protocols used to communicate with each smart TV OS for automated experiment control. Samsung TV and LG TV were controlled using OS-specific third-party Python libraries (samsungtvws via WebSocket and aiowebostv via WebSocket, respectively). TVs running Roku OS (Roku TV and TCL TV) were controlled using Roku’s External Control Protocol (ECP) via aiohttp. Vizio TV was controlled using a custom implementation of the SmartCast API over HTTPS. Although Sony and TCL/Hisense TVs both run Google TV, they required different control approaches: Sony required a custom implementation using the SSIP protocol over TCP, while TCL and Hisense were controlled via ADB over TCP using adb-shell. Fire TV was also controlled via ADB over TCP using the adb CLI.

B Filter Keywords and ACR Endpoint Blocklist

We present the keywords used to filter well-known service domains during ACR endpoint identification (§5.1) and a blocklist of identified ACR endpoints that users can apply to limit ACR tracking.

B.1 Keywords Used to Filter Well-Known Service Domains

To filter the initial list of all endpoints contacted by smart TVs, we remove domains belonging to well-known streaming services (e.g., Netflix, YouTube), content delivery infrastructure (e.g., Google static content, Google user content delivery), and advertising platforms (e.g., Facebook advertising). The full list of keywords used to identify these domains is provided in Table 7.

Table 7: Keywords used to filter well-known service domains.

Keyword	Purpose
netflix, nflx	Netflix streaming service
youtube, ytim, yt3	YouTube streaming service
googlevideo	Google video delivery
tubi	Tubi streaming service
pluto	Pluto TV streaming service
haystack	Haystack News streaming service
amagi	Amagi content delivery
doubleclick	Google advertising platform
facebook	Facebook advertising platform
apple	Apple services
gstatic	Google static content delivery
googleusercontent	Google user content delivery
arpa	DNS infrastructure

B.2 ACR Endpoint Blocklist

We provide a blocklist (Table 8) that smart TV users can use to limit ACR tracking. For each OS, we list the ACR partner and the domains identified as ACR viewing data collection endpoints under “Validated endpoints.” We also provide additional recommended blocklist rules for domains associated with ACR operators but not identified as viewing data collection endpoints through our methodology. Finally, we note that although Amagi domains are listed due to the plaintext channel name leakage they enable, we do not recommend blocking them as they are required for first-party streaming content delivery.

C Comparison with Policy Disclosures

Table 9 and Table 10 present the classification of privacy policy disclosures for each smart TV OS and the corresponding verbatim disclosure language from each source document. We categorize disclosures into three tiers: T1 (main U.S. privacy policy), T2 (binding supplement or EULA section), and T3 (support or FAQ page outside the privacy policy), and classify each disclosure dimension as disclosed (D), quasi-disclosed (Q), undisclosed (U), or contradicted (C). Our analysis reveals that several OSes contain undisclosed or contradicted details across multiple dimensions. Notably, ACR-related disclosures are frequently found only in T2 or T3 sources rather than in the main privacy policy, meaning many users are unlikely to encounter them during routine privacy policy review.

Table 8: ACR endpoint blacklist released with this work. For each operating system we list the ACR operator, the validated viewing-data endpoint(s), the durable second-level-domain (SLD) or wildcard rule we recommend for blocklists (ACR operators rotate sub-hosts), and the classification status. *Validated* endpoints are confirmed by our four-indicator methodology (§5.2); *Suspected* endpoints have strong supporting evidence but no payload-level confirmation. We additionally release the per-OS exact-host lists and the Amagi plaintext-DNS leak channel, which is excluded from the recommended blacklist because its hostnames also deliver First-Party (FAST) streaming content and blocking them breaks playback. The complete machine-readable blacklist is released in our artifact repository [22].

Operating System	ACR Operator	Validated Endpoint(s)	Recommended Blocklist Rule
Tizen (Samsung)	Samsung	acr-us-prd.samsungcloud.tv	*.samsungcloud.tv, *.samsungacr.com
webOS (LG)	Alphonso	tkacrNNN.alphonso.tv	*.alphonso.tv
Roku OS (Roku, TCL)	Roku	scribe.logs.roku.com	scribe.logs.roku.com
Google TV (Sony)	Samba TV	*.cid.samba.tv; IP 216.183.117.80	*.samba.tv (+ block IP at L3)
Fire OS (Amazon)	Amazon	gateway-ink.amazon.com	gateway-ink.amazon.com
SmartCast (Vizio)	Inscape	Kinesis [†]	*.tvinteractive.tv
VIDAA OS (Supersonic)	Nexxen/Unruly	NN.fp.us.acr.unruly.co	*.acr.unruly.co
Xumo OS (Westinghouse)	—	none observed	—
Parallel plaintext channel (independent of ACR; <i>not</i> in recommended blacklist)			
Amagi first-party streaming path	Amagi	*.amagi.tv	*.amagi.tv (breaks FAST streaming)

[†] Vizio’s content-correlated traffic terminates at `kinesis.us-west-2.amazonaws.com`, shared Amazon Kinesis infrastructure. Because the host is shared cloud infrastructure, we do not include the bare Kinesis hostname in the recommended blacklist; the Inscape control domains provide a safer block target.

Table 9: Privacy-policy disclosure of ACR across the eight smart-TV OSes, decomposed into eight disclosure dimensions. *Tier* indicates where ACR is disclosed: T1 main US privacy policy, T2 a binding supplement/EULA section, T3 a support/FAQ page outside the policy. Per-dimension scale: **D disclosed (explicit), **Q** quasi-disclosed (vague, euphemistic, or only in a T3 page), **U** undisclosed (policy silent), **C** contradicted (policy asserts the opposite of our observed behavior). The on-device opt-out label is each vendor’s consumer-facing term; only Vizio and VIDAA pair it with the literal term “ACR” in the same document. Xumo is the no-ACR control.**

OS (operator)	Opt-out label	Tier	ACR named	Capture mech.	Input paths	Data types	Partner named	Opt-out scope	Sens. infer.	Source
Tizen (Samsung)	Viewing Information Services	T3	Q	D	U	D	U	C	U	[41, 42]
webOS (Alphonso)	Live Plus	T2	D	D	Q	D	D	D	U [‡]	[2]
Roku OS (Roku)	Use info from TV inputs	T2	D	Q	D	D	Q	D	U	[37]
Google TV (Samba/Sony)	Samba Interactive TV	T3	Q	U	U	U	D	D	U	[44]
Fire OS (Amazon)	Automatic Content Recognition	T3	D	Q	Q	Q	U	D	U	[6, 7]
SmartCast (Inscape/Vizio)	Viewing Data	T2	D	D	D	D	Q	C	Q	[53]
VIDAA OS (Nexxen)	Enhanced Viewing Service	T2	D	D	Q	D	D	D	U	[51]
Xumo OS (none)	—	—	—	—	—	—	—	—	—	[58]

[‡] Alphonso’s policy categorically promises it does not collect religious, health, or sexual-orientation data. webOS ACR traffic is encrypted, so we have no payload-level ground truth and do not score this as *contradicted*; we record the promise. Separately, Alphonso’s capture-side carve-out (“does not recognize ... ambient sounds, and other non-commercial audio”) is in tension with the absence of any traffic decrease we observe on a black-screen-with-valid-audio feed (§5.3).

Table 10: Verbatim disclosure language underlying Table 9, with the operative quotation from each vendor’s US-facing source. Sources are the main US privacy policy (T1), a binding supplement/EULA section (T2), or a support/FAQ page outside the policy (T3). Quotations are reproduced exactly as published at the cited URLs (accessed May 2026).

OS (operator)	Tier	Operative disclosure language (verbatim)	Src
Tizen (Samsung)	T3	“Your Samsung Smart TV may use Automatic Content Recognition (“ACR”) technology to enhance your viewing experience.” The term appears only on a support FAQ; the binding US privacy policy does not mention ACR. The FAQ also states: “Please note that Samsung does not record or watch any of the content displayed on your Smart TV.”	[41, 42]
webOS (Alphonso)	T2	“the Service may collect video and/or audio information from what you are watching through your device, regardless of the source of the information ... in order to determine the programs you have watched and the advertisements you have seen.” Capture carve-out: “Alphonso’s Video AI does not recognize human conversation, ambient sounds, and other non-commercial audio or video content.” Categorical promise: “Under no circumstances does Alphonso collect any sensitive personal information, including ... data revealing racial or ethnic origin, religious beliefs, mental or physical health diagnosis, sexual orientation ...”	[2]
Roku OS (Roku)	T2	“Roku will, through ACR technology, collect information about what you watch and when you watch (e.g., the programs, commercials and channels you viewed, and the date, time and duration of the viewing) via your Roku TV’s antenna, and devices connected to your Roku TV.” Scope limit: “Your selection will not affect data collected when you access or stream content over the Internet.”	[37]
Google (Samba/Sony)	TV T3	“You may enable Samba Interactive TV only if you accept Samba TV’s terms of service and acknowledge notice of its privacy policy.” Sony’s page uses the brand name “Samba Interactive TV” and does not use the term “ACR”; content sensing is deferred to the third-party Samba TV policy.	[44]
Fire OS (Amazon)	T3	“When ACR is enabled, Amazon collects viewing data about content you view on your Fire TV smart TV through connected devices.” Separately: “you may opt out of providing information to Amazon on your use of over-the-air TV content received through an antenna connected directly to your TV.”	[6, 7]
SmartCast (In-scape/Vizio)	(In- T2	“VIZIO OS products ... have built-in technology, such as Automated Content Recognition (“ACR”), that can understand the VIZIO OS product viewing behavior and usage, including audio and video programming, ads, gaming content, devices connected to the VIZIO OS product (e.g., a streaming stick plugged into an HDMI port), and third-party apps, in real-time.”	[53]
VIDAA (Nexxen)	OS T2	“The EV Service uses Automated Content Recognition (“ACR”) technology, which works by analyzing and matching onscreen audio sounds that come from your Smart TV ... against an extensive database of programming content.” Partner: “Viewing Data will be used to help our advertising partner Nexxen and its partners select, deliver and measure personalized ads.”	[51]
Xumo OS (none)	—	No ACR, fingerprint, or content-matching language appears. The privacy page lists “Recent history,” “Personalized recommendations,” and app-data sharing; governing policy is Comcast/Xfinity.	[58]