A Network-Level Comparison of Privacy Risks for Children vs. Adults in Virtual Reality

Eleanor Brunskog
Linköping University
Linköping, Sweden

Sofia Knutas
Linköping University
Linköping, Sweden

Sheyda Mirzakhani
Linköping University
Linköping, Sweden

Niklas Carlsson
Linköping University
Linköping, Sweden

Abstract

As virtual reality (VR) platforms grow in popularity—especially among children—they raise pressing concerns about the nature and extent of user data exposure. While providers such as Meta claim to implement age-based privacy protections, it remains unclear whether these measures result in observable differences in network-level behavior. In this paper, we present the first systematic network-level analysis comparing traffic patterns associated with child and adult accounts in VR. Using a controlled and replicable measurement framework, we analyze encrypted traffic from 96 matched sessions across six popular Meta Quest applications. Although encryption prevents us from observing payload contents, our analysis of metadata—including traffic volume, contacted domains, and geographic routing—reveals that child accounts often generate more outbound traffic, connect to a broader range of third-party and platform-party domains, and in several cases, contact advertising and tracking services not reached by adult profiles. We further map outbound connections to organizational ownership and geographic location, uncovering frequent international transfers, sometimes to jurisdictions with limited privacy safeguards. These findings raise critical questions about whether existing age-based privacy controls meaningfully uphold legal obligations, such as those under GDPR and CRC, and point to a need for greater transparency and stronger enforcement in immersive ecosystems.

CCS Concepts

• **Security and privacy → Web application security**; **Network security**; **Human and societal aspects of security and privacy**; • **Information systems → Multimedia information systems**.

Keywords

Virtual Reality (VR), Network Traffic Analysis, Children's Privacy, GDPR, Age-Based Privacy Controls, Cross-Border Data Flows

ACM Reference Format:
Eleanor Brunskog, Sofia Knutas, Sheyda Mirzakhani, and Niklas Carlsson. 2025. A Network-Level Comparison of Privacy Risks for Children vs. Adults in Virtual Reality. In *Proceedings of the 2025 Workshop on Privacy in the Electronic Society (WPES '25), October 13–17, 2025, Taipei, Taiwan.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3733802.3764049

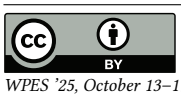

WPES '25, October 13–17, 2025, Taipei, Taiwan

ACM ISBN 979-8-4007-1898-4/2025/10
https://doi.org/10.1145/3733802.3764049

1 Introduction

Virtual reality (VR) platforms are rapidly reshaping how users engage with digital content, offering immersive experiences in gaming, education, and social interaction. However, these systems inherently collect rich telemetry, including motion patterns, biometric signals, and behavioral traces; data that often exceeds what is gathered by conventional web or mobile applications.

As adoption accelerates, particularly among youth, VR systems are becoming deeply integrated into children's digital lives. This raises urgent questions about the extent to which these environments expose children to privacy risks, especially given their legal entitlement to heightened protections under laws such as the General Data Protection Regulation (GDPR) [22] and the Convention on the Rights of the Child (CRC) [19].

While platform providers like Meta claim to enforce age-based privacy controls, it is unclear whether these measures manifest in measurable differences in how applications behave at the network level. Regulatory frameworks (e.g., GDPR and CRC) mandate limits on tracking and profiling of minors; yet, no prior empirical work has evaluated whether those limits translate into observable reductions in third-party exposure or data transmission patterns during real-world VR usage.

In this paper, we present the first systematic, network-level measurement study comparing traffic patterns between child and adult accounts in VR. Using a controlled and replicable experimental framework, we analyzed 96 matched sessions across six popular applications on the Meta Quest platform. Each session follows an identical usage script to isolate differences attributable to account type, rather than user behavior.

Although most of the observed traffic is encrypted, this lens is necessary in VR where platform services are largely closed-source and resist static inspection; by analyzing metadata—including packet volume, contacted domains, organizational ownership, and geographic routing—we can empirically assess whether age-based protections manifest in practice. Our results show that child accounts often generate more outbound traffic, reach a wider set of platform and third-party domains, and in some cases contact advertising and tracking services not seen with adult accounts, with frequent international transfers, including to jurisdictions with limited oversight. These findings raise questions about the effectiveness of self-declared privacy protections for children and highlight persistent gaps between regulatory expectations and what is observed at the network level in practice.

At a high level, this work makes the following contributions:

- We present the first systematic network-level study comparing privacy exposure between child and adult accounts in virtual reality, using matched usage patterns across six popular Meta Quest applications.

- We develop a controlled and replicable measurement framework that isolates traffic differences attributable solely to user age, capturing packet-level data under consistent and realistic usage conditions.
- We perform a differential traffic analysis showing that child accounts often generate more outbound traffic and connect to a broader set of destinations—including ATS-flagged and third-party domains—than adult accounts.
- We map outbound connections to organizational ownership, geographic jurisdiction, and domain type, revealing that child accounts routinely interact with international and potentially non-compliant entities.
- We interpret these findings through the lens of GDPR and CRC, demonstrating that purported age-based privacy protections are not meaningfully enforced in practice, and highlighting gaps in transparency and regulatory compliance.

Our results suggest that current VR platforms may not meaningfully differentiate the data collection behavior of child and adult accounts, potentially exposing children to equal or even greater privacy risks despite legal obligations for heightened safeguards. This raises concerns about whether existing self-regulatory mechanisms are sufficient to enforce age-appropriate data practices in immersive environments. By systematically measuring real-world traffic, our study offers concrete evidence of these gaps and provides actionable insights for developers, platform providers, and policymakers seeking to strengthen privacy protections for younger users.

Outline: Section 2 reviews related work. Section 3 details our measurement methodology. Sections 4 and 5 present our traffic and destination-level analyses. Section 6 examines a case study of a privacy-oriented app update. Finally, we discuss broader implications in Section 7 and conclude in Section 8.

2 Related Work

We first review prior research on privacy risks in VR environments and child data protection across digital platforms. While existing work highlights significant transparency gaps, over-collection of data, and regulatory non-compliance, no previous study has systematically compared child and adult privacy exposure in VR through network-level traffic analysis; an important gap addressed here.

Privacy and Data Collection in VR: Virtual reality platforms inherently collect rich telemetry, including behavioral, biometric, and environmental data. Trimananda et al. [25] introduced the OVRseen framework to examine data practices in 140 VR applications, revealing that 70% of collected data types were not disclosed in app privacy policies. Even accounting for third-party policies, only 74% of flows were documented. Zhan et al. [27] found that nearly half (48.1%) of 1,726 apps collected excessive sensitive data, often violating data minimization principles.

Other studies have highlighted systemic issues in VR such as weak privacy controls, missing incognito modes, and opaque platform behavior [8, 16, 20]. Behavioral and biometric signals—such as motion, eye gaze, and voice—have been shown can uniquely identify users [21, 23], while some applications have embedded covert data collection mechanisms [17]. These findings collectively raise concerns about how VR platforms handle user data, especially in the absence of effective oversight or technical safeguards [10, 14, 18, 26].

Children's Privacy on Digital Platforms: Prior work on child privacy has mostly focused on mobile and web ecosystems. Reyes et al. [24] found that 19% of child-directed apps collected personally identifiable information (PII) in violation of Children's Online Privacy Protection Act (COPPA) and failed to disable tracking. Figueira et al. [7] reported widespread data collection before age or consent checks. Others cite weak disclosures, poor defaults, and lax enforcement as compounding risks [3, 11, 15]. Carlsson et al. [4] showed that 87% of apps shared data with third parties—often without parental consent—breaching GDPR. Lyu et al. [12] call for immersive-specific protections for children. Meanwhile, Fiani et al. [6] found that nearly half of surveyed parents reported underage children actively using social VR apps. These patterns underline the urgency of evaluating whether immersive platforms offer children the elevated privacy protections they are legally entitled to.

Our Contribution in Context: While prior work has examined VR privacy [16, 20, 25, 27] and children's data practices on mobile and web platforms [3, 7, 11, 12, 15, 24], no study has compared child and adult profiles at the network level in VR. We present the first controlled analysis using matched Meta Quest accounts, capturing encrypted traffic and classifying domains by ownership, type, and tracking affiliation. Our results reveal systematic differences in outbound communication and third-party exposure, raising compliance concerns under GDPR and the CRC [19, 22] and highlighting the need for stronger age-aware design and transparency.

3 Methodology

Our study is designed to answer one core question: do VR applications exhibit measurable differences in network-level behavior between child and adult accounts? To address this, we developed a four-step methodology designed for consistency and control: (1) selecting representative applications from the Meta VR ecosystem, (2) collecting network traffic data using a controlled setup, (3) processing raw captures to extract relevant metadata, and (4) analyzing communication patterns to identify potential differences in third-party exposure and data handling practices. This design isolates backend treatment attributable to account age while holding user actions constant. We next outline each step in more detail.

3.1 Selection of VR Applications (Step 1)

The selection of VR applications was guided by four key criteria: popularity, availability to both children and adults, the applications' reliance on internet connectivity, and cost. By prioritizing widely used and free applications that support child users, the study aimed to ensure real-world relevance while maintaining accessibility.

Given Meta's dominance in the VR market [2], the Meta Quest ecosystem was selected as the platform for evaluation. Six free applications were chosen from the Meta Horizon Store (Table 1). All were required to allow usage by individuals aged 13 and above, in line with Meta's age policy [13]. Although applications do not prompt for age during login, Meta's backend uses the Get Age Category API to distinguish between users based on their account information [9]. This ensures that differential treatment–if any–is applied at the system level, rather than through user prompts.

Pre-Usage Analysis (and Further Motivation): Before gameplay, each application was examined to assess whether any privacy

Table 1: Selected VR applications used in the study. None of the applications required an account.

Application	Developer	Category
Gorilla Tag	Another Axiom	Action
Gym Class	IRL Studios	Sports/Simulation
Penguin Paradise	Sava	Action
Scary Baboon	Aura Vision LLC	Action/Social/Strategy
VRFS - Football Simulator	Immersification	Sports/Simulation
Yeeps: Hide and Seek	Trass Games	Action/Sandbox

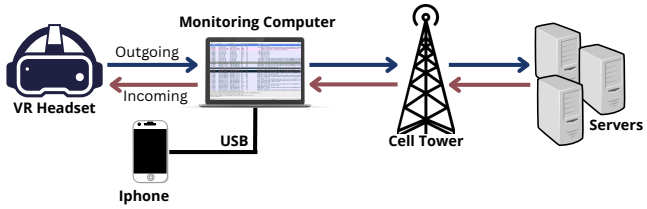

Figure 1: Network packet flow during data collection.

policies or consent mechanisms were presented differently to children and adults. Interestingly, no differences were observed: all applications presented the same interface and information to both profiles. This absence of visible differentiation suggests that any age-based privacy controls operate opaquely, if at all, highlighting the need to empirically examine how these applications behave at the network level for child versus adult users.

3.2 Data Collection (Step 2)

We designed a controlled setup that captured all inbound and outbound traffic from the VR headset, isolating session-specific flows and allowing direct comparison between matched child and adult Meta accounts.

Experimental Setup: Two Meta accounts were created: one for an adult, and one for a 13-year-old child (DOB: Jan. 1, 2012). To capture traffic in a controlled environment, a dedicated monitoring computer was configured to act as both a Wi-Fi access point and a data capture device. The setup is illustrated in Figure 1 and included:

- **Wi-Fi Hotspot:** We created a private Wi-Fi network using the monitoring computer, which shared its internet connection via USB tethering to an iPhone's 4G data link. Only the Meta Quest 3 headset connected to this network, minimizing background noise. Average download and upload speeds were 117.5 Mbit/s and 33.3 Mbit/s, respectively.
- **Monitoring Computer:** A MacBook Pro functioned as both network gateway and packet capture device. It ran Wireshark to log all headset traffic via the virtual interface *bridge100*.
- **VR Headset:** A Meta Quest 3 was used for all testing. As a standalone device widely used in consumer VR, it ensured compatibility with current Meta applications.
- **Network Analysis Tool:** Wireshark captured raw packet data. Traffic was filtered by IP addresses using expressions like `ip.src == XXX.XXX.X.X` and `ip.dst == XXX.XXX.X.X` to isolate traffic generated and received by the headset.

Experimental Design: To ensure a fair and repeatable comparison of traffic patterns, we implemented controlled experiments with standardized usage across account types. Each application was tested eight times per profile, yielding 16 sessions per application and 96 sessions in total. To reduce temporal bias, sessions were conducted in alternating pairs (e.g., child → adult, then adult → child) and we tried to make each pair of sessions as similar as possible. Each session lasted five minutes and involved standardized interactions such as menu navigation, environment exploration, and basic gameplay. We performed similar actions across account types within each application to balance realism and repeatability, while mimicking typical short-duration use. No UI or functional differences were observed between child and adult accounts, ruling out front-end factors as explanations for traffic variation.

Structure and Content of Captured Network Data: All sessions were recorded as `.pcapng` files, each containing a timestamped sequence of packets with metadata including size, IP addresses, protocol types (e.g., TCP, UDP, TLS), and direction. The dataset comprises 96 total recordings (6 applications × 8 sessions × 2 profiles), totaling 5.3 GB in size. File sizes varied from 10 MB to 190 MB depending on application behavior.

3.3 Post-Processing (Step 3)

Following data collection, packet-level data from the network captures was organized and enriched to support deeper analysis. First, Wireshark's built-in tools were used to generate session-level statistics such as IP conversations, packet counts, and data volumes. Second, a custom Python script was developed to: (1) extract unique IP addresses per session, (2) query external services to retrieve Autonomous System Number (ASN), organization name, and server location for each IP, (3) associate each IP with fully qualified domain names (FQDNs) using TLS Server Name Indication (SNI) values (where available), and (4) derive effective second-level domains (eSLDs) for domain classification. Finally, domains were subsequently (manually) checked against known advertising and tracking service (ATS) blocklists to flag potential privacy concerns.

3.4 Data Analysis (Step 4)

To identify and evaluate potential differences in network-level communication patterns between child and adult profiles, we first analyzed a set of key traffic-level metrics, including:

- Total number of unique IP addresses contacted
- Contacted organizations and number of third-party domains
- Volume of transmitted and received data
- Geographic distribution of contacted servers

When applicable, we applied statistical tests to evaluate whether the metrics differed significantly by account type. Pairwise t-tests were used for key metrics, with 95% confidence intervals reported to assess significance, and we applied binomial hypothesis testing to assess statistical significance in sets of pairwise winners.

Second, we performed a domain-level analysis of the SNI values from TLS handshakes, mapping FQDNs to organizations and comparing profiles to assess third-party exposure. Specifically, we:

- Classified domains into first-, platform-, or third-party
- Identified ATS-flagged domains and their traffic volumes

Further methodological details, including classification rules for domain ownership, are presented in Section 5.

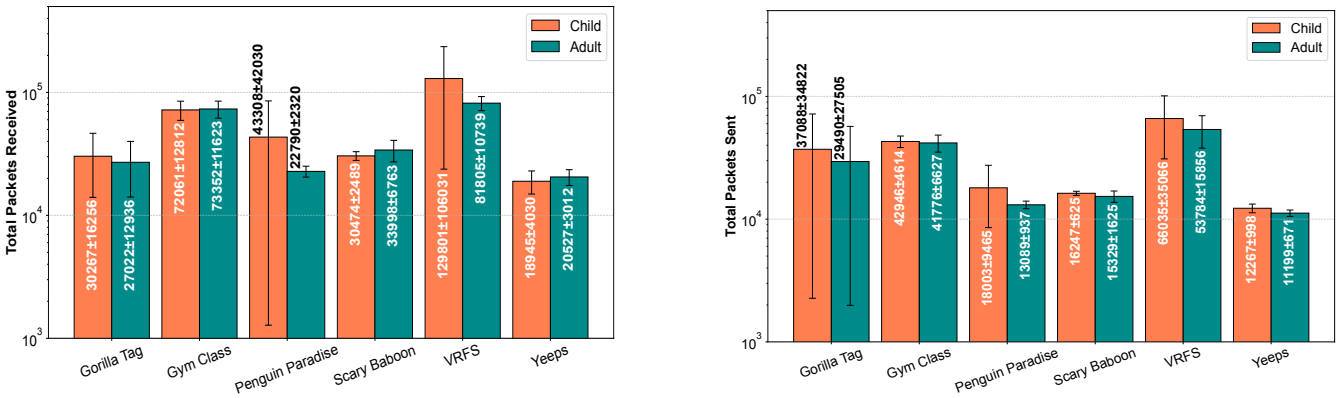

(a) Incoming traffic. **(b) Outgoing traffic.**

Figure 2: Total number of packets sent and received per application.

3.5 Methodological Limitations

While the methodology was robust and tailored to the study's goals, several limitations should be acknowledged:

- **Short Session Duration:** Each session lasted five minutes, which allowed for consistent comparisons but may not reflect longer-term data sharing behaviors such as those triggered by matchmaking, updates, or gameplay progression.
- **Application Selection:** By selecting six free, popular consumer VR apps on Meta Quest, we capture realistic usage in today's most accessible VR ecosystem. Although this scope does not cover all contexts (e.g., enterprise or educational VR), it establishes a foundation that future studies can extend to assess generality across platforms.
- **Application Stability:** Some of the tested VR applications were in early access or exhibited instability. Glitches or inconsistent behavior may have affected the type or volume of data transmitted.
- **Encrypted Traffic:** The majority of network traffic was encrypted (via TLS), limiting visibility into the full extent of data exchanges. While some FQDNs could be extracted via the SNI field, many were unavailable—either due to encryption or omission—reducing coverage. Finally, we note that our IP-based analysis also provides some insights into the diversity of services contacted. While complementary insights may be obtained using alternative techniques (e.g., static or dynamic code analysis), many VR apps rely on opaque platform-side services and dynamically integrated third parties that are only visible through network-level monitoring such as the one used in this study.
- **Lack of Action-Level Correlation:** The study did not map specific in-app actions to traffic events. Doing so could improve attribution of data flows to particular features or user behaviors. To provide fair comparison, we instead ensured that we performed the same actions for the two user groups.

Despite these constraints, the methodology enabled consistent data capture and reliable identification of key communication endpoints. By focusing on differences in observable metadata and domain contact behavior, the analysis offers meaningful insights into how VR applications interact with child versus adult accounts. Even without access to payload contents (due to encryption), our methodology offers a practical, privacy-respecting approach to auditing immersive technologies, and our findings provide important signals of age-based variation in backend behavior.

4 High Level Comparison

Given the sensitivity of behavioral and sensory data in VR, child profiles—subject to heightened legal protections—might be expected to communicate less and with fewer external entities than adult profiles. Although encryption limits visibility into payload contents, metadata such as packet volume, size, and destination diversity serve as a useful proxy for backend behavior. In this section, we begin with a high-level comparison of incoming and outgoing traffic across three key metrics: total packet count, average packet size, and the number of unique destination IPs contacted.

For this analysis, we present aggregated results per application and user category, showing means with confidence intervals (Figures 2–4) and binomial pairwise tests (Table 2) to assess consistency. From a privacy-by-design perspective, we expect child accounts to generate less traffic and fewer contacts. In the following three subsections, we examine each metric one-by-one.

4.1 Packets Sent and Received

Figure 2 shows the total number of packets sent and received by child and adult profiles across six applications, plotted on a log scale. From these results, we note that (contrary to the above hypothesis) the child profiles consistently generated more outgoing traffic across all six applications. For incoming traffic, the pattern was less uniform, with child accounts receiving more packets than adults in half of the applications (Gorilla Tag, Penguin Paradise, VRFS). This suggests that while outbound communication differs systematically by account type, inbound traffic varies more by application.

The trend that there is a higher degree of outgoing traffic among children is further supported by the pairwise analysis results summarized in Table 2. Across the eight trials per application, child profiles generated more outgoing traffic in 5 out of 6 applications, winning 6, 7, 6, 5, and 6 out of 8 trials, respectively, while one (Yeeps) was evenly split (4 out of 8). For incoming traffic, child dominance was more evenly distributed and less consistent.

Table 2: Combined child-win counts and p-values (based on one-sided binomial tests, with null-hypothesis $p = 0.5$) across six traffic characteristics, color-coded by: red (significant child dominance), black (similar values), blue (significant adult dominance). In all cases we performed eight pairwise trials.

Application	Packet Transmissions				Packet Sizes				Unique Destinations			
	Incoming		Outgoing		Incoming		Outgoing		Incoming		Outgoing	
	Child >	p-value	Child >	p-value	Child >	p-value	Child >	p-value	Child >	p-value	Child >	p-value
Gorilla Tag	5	0.3633	6	0.1445	1	0.9961	7	0.0352	5 (+1 tie)	0.3633	(+2 ties) 5	0.3633
Gym Class	2	0.9648	4	0.6367	3	0.8555	6	0.1445	5	0.3633	5	0.3633
Penguin Paradise	5	0.3633	7	0.0352	3	0.8555	3	0.8555	2 (+1 tie)	0.9648	2 (+1 tie)	0.9648
Scary Baboon	2	0.9648	6	0.1445	2	0.9648	7	0.0352	4 (+1 tie)	0.6367	4 (+1 tie)	0.6367
VRFS	5	0.3633	5	0.3633	3	0.8555	3	0.8555	5	0.3633	5	0.3633
Yeeps	3	0.8555	6	0.1445	3	0.8555	3	0.8555	2 (+3 ties)	0.9648	3 (+3 ties)	0.8555

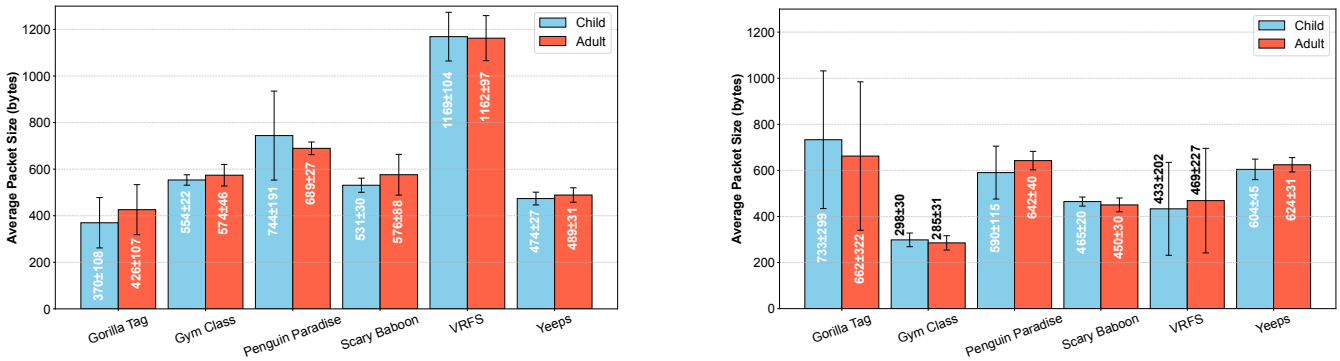

(a) Incoming traffic. **(b) Outgoing traffic.**

Figure 3: Average packet size (in bytes) for each application.

While overall packet transmission rates were relatively balanced for most applications, VRFS stood out with particularly heavy traffic and a pronounced difference between child and adult profiles. On average, the child profile sent 66,035 packets, compared to 53,784 by the adult. This difference was accompanied by much higher variability in the child profile (CI ±35,066 vs. ±15,856 for adults), resulting in overlapping confidence intervals, though the adult CI was entirely contained within the child's.

The child profile also observed greater variability for both outgoing and incoming traffic with Gorilla Tag, Penguin Paradise, and Yeeps, though in most cases with overlapping intervals, making statistical separation difficult at a 95% confidence level.

Despite these overlapping confidence intervals, there was two clear statistical differences associated with the outgoing traffic: (1) with Yeeps, the child profile sent significantly more outgoing packets than the adult profile, as confirmed by a paired t-test ($t(7) = 2.68$, $p = 0.031$) and (2) with Penguin Paradise, the child profile sent more outgoing packets than the adult profile in 7 out of 8 cases ($p = 0.0352$ when applying one-sided binomial testing).

Overall, these results suggest that while there are consistent trends of higher outgoing traffic in child accounts—especially in applications like VRFS and Penguin Paradise—statistical significance is limited due to variability and sample size. Still, the directionality and consistency of packet transmission differences, especially in outgoing traffic, suggest systematic behavioral or design differences in how these applications interact with child versus adult profiles.

4.2 Packet Sizes

Figure 3 presents the average packet sizes (in bytes) for incoming and outgoing traffic across the six applications, comparing child and adult profiles. Overall, no consistent directional pattern emerges across applications or traffic types.

For incoming traffic, child profiles received larger packets in 2 out of 6 applications (Penguin Paradise and VRFS) and for outgoing traffic we observed an even split: the child profile sent larger packets in 3 out of 6 applications (Gorilla Tag, Gym Class, and Scary Baboon), while the adult profile sent larger packets in the remaining three (Penguin Paradise, VRFS, and Yeeps). These trends are also reflected in the pairwise comparisons shown in Table 2, where we only observe significant differences for two applications: (1) for Gorilla Tag, the adult profile see bigger incoming packets in 7 out of 8 pairwise trials ($p = 0.0352$) and the child profile see bigger outgoing packets in 7 out of 8 pairwise trials ($p = 0.0352$), and (2) for Scary Baboon, the child profile see bigger outgoing packets in 7 out of 8 pairwise trials ($p = 0.0352$).

While some directional differences emerge, confidence intervals for child and adult profiles largely overlap across applications and traffic directions, limiting statistical significance in most cases. Only Gorilla Tag's incoming traffic showed a significant difference, with adults receiving larger average packets ($t(7) = -2.94$, $p = 0.022$).

Beyond statistical tests, variability in packet sizes varied notably across applications. Gorilla Tag and VRFS exhibited the widest CIs, indicating substantial variation in traffic characteristics across trials. In contrast, Gym Class, Scary Baboon, and Yeeps had narrower CIs, reflecting more stable average packet sizes seen across sessions.

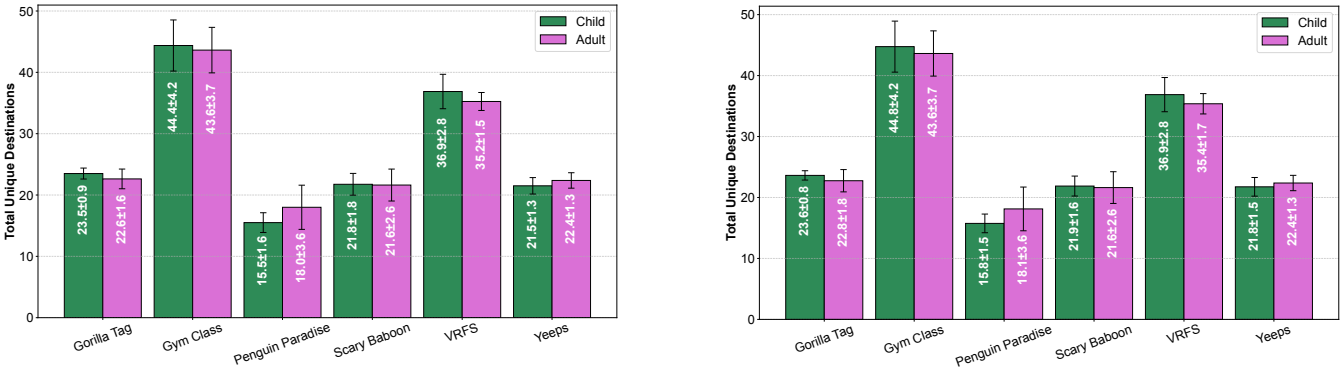

(a) Incoming traffic. **(b) Outgoing traffic.**

Figure 4: Total unique destinations per application.

Some applications also had significantly different average packet sizes. For outgoing traffic, Gorilla Tag had the largest packets (733 bytes for child profiles and 662 bytes for adult), while Gym Class had the smallest (298 bytes for child and 285 bytes for adult). For incoming traffic, VRFS recorded the largest packets for both groups (averaging 1,169 bytes for children and 1,162 bytes for adults), while Gorilla Tag again stood out on the opposite end, receiving the smallest packets (370 bytes for children vs. 426 bytes for adults).

In summary, while packet size differences exist across applications and profiles, the differences are generally small and highly variable, with little evidence of systematic variation based on child versus adult usage, aside from the isolated statistical significance in Gorilla Tag's incoming traffic (significant using both t-test and binomial test) and (to a smaller extent) the outgoing traffic of Gorilla Tag and Scary Baboon (only binomial tests significant). Although we cannot inspect encrypted payloads, variation in packet size may still reflect differences in the volume or structure of transmitted information. These findings reinforce the need for continued investigation into how age-based profiles may influence backend data handling, particularly in immersive environments where large volumes of behavioral telemetry are exchanged.

4.3 Number of Contacted IP Addresses

Figure 4 shows the number of unique destination IPs contacted per session for each application, for incoming and outgoing traffic. Overall, the results suggest moderate directional differences (indicating that the communication typically is bi-directional) but limited statistical significance between child and adult profiles.

In both traffic directions, the child profile contacted more unique destinations in 4 out of 6 applications; specifically in Gorilla Tag, Gym Class, Scary Baboon, and VRFS. The remaining two applications, Penguin Paradise and Yeeps, showed slightly higher or equal destination counts for the adult profile. These trends are mirrored in the pairwise comparisons shown in Table 2, where the "child wins" counts lean slightly in favor of the child profile, but no clear or consistent dominance emerges across the applications and none of the observed differences is statistically significant.

Importantly, most applications display only small differences in the number of unique destinations, often differing by just one or two IPs per session between the profiles. Penguin Paradise exhibits the largest directional difference, with the child profile averaging 15.8 destinations (in the outgoing traffic) compared to 18.1 for the adult, resulting in a modest absolute difference of 2.3 destinations.

In addition to small profile differences, all applications contact many servers. Gym Class consistently ranks highest (child profile averaging 44.8 outgoing destinations and the adult 43.6), followed by VRFS (36.9 for child profile and 35.4 for adult), with the other applications exhibit lower destination counts (typically 17-23).

From a statistical standpoint, however, these differences do not reach significance. The 95% confidence intervals overlap in all applications, for both incoming and outgoing traffic. This overlap suggests that variability across sessions is large enough to mask systematic differences between profiles. The accompanying paired t-tests confirm this, with the lowest observed p-values being 0.160 (outgoing) and 0.163 (incoming), well above the conventional threshold of 0.05. Similarly, we observe no significant differences when applying the binomial test on the pairwise winners.

In summary, although the child profile often contacts slightly more destinations, especially in high-traffic applications like Gym Class and VRFS, these differences are small in magnitude and statistically non-significant, reflecting a generally similar pattern of network diversity across both child and adult usage. This lack of differences is concerning, especially when combined with the high number of servers contacted per (short) five-minute session.

4.4 Cross-Metric Patterns and Discussion

Combined, the high-level metrics (packet count, packet size, and number of contacted destinations) reveal several patterns that may have important privacy implications, especially for child accounts.

First, child accounts consistently generate more outgoing traffic across all six applications, challenging the expectation that adult profiles would be more data-intensive (e.g., due to broader access to settings or features, or fewer usage restrictions). Instead, the elevated traffic from child profiles The elevated traffic in child profiles raises questions about whether it is strictly necessary for gameplay or reflects background telemetry, behavioral analytics, or third-party data flows; issues of particular concern under child privacy regulations. While higher traffic volumes in child accounts may partly reflect benign factors (e.g., session management, keep-alive messages), the consistent elevation suggests that children are not subject to stricter minimization in practice. We therefore

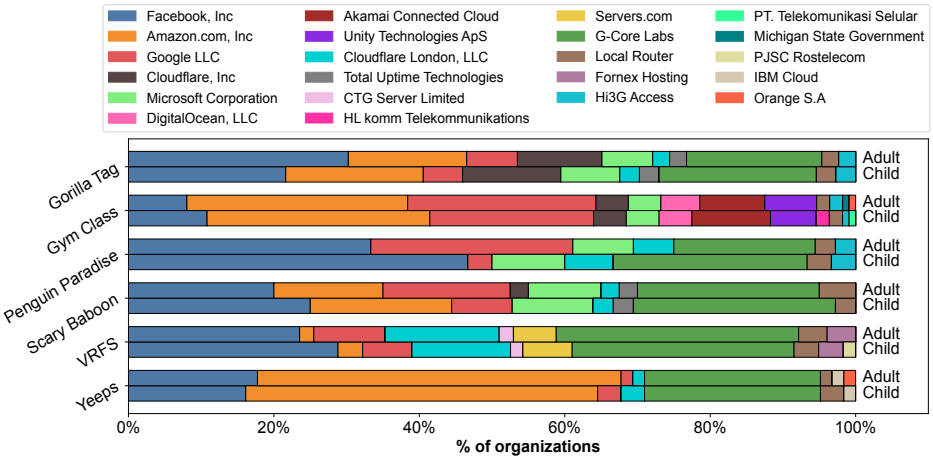

Figure 5: Distribution of contacted organizations per application.

interpret these results as risk signals, rather than definitive proof, of weaker privacy enforcement.

Second, although differences in average packet sizes are less consistent, some applications—most notably Gorilla Tag, Gym Class, and Scary Baboon—show both larger average packet sizes and higher packet counts for child profiles. This suggests not just a greater number of communications, but also increased total data payloads, potentially indicating more detailed or frequent data capture. Given that both profiles performed the same actions under controlled conditions, these differences may reflect divergent backend treatment of child versus adult users. Alternatively, they could stem from how features are architected or how child-specific content is implemented and delivered, with possible implications for profiling or tracking. Regardless of the underlying reason, these observations are concerning as they suggest that child account may be sending more information to external servers, or at least no less information, than the corresponding adult users.

Third, the number of unique destinations contacted appears less sensitive to profile type. The counts are generally similar across users, with only small absolute differences in most applications. This suggests that the breadth of external services or third-party domains contacted by the application may be primarily dictated by game design rather than user identity. However, games with higher traffic volumes—such as Gym Class, VRFS, and Gorilla Tag—also tend to contact more unique destinations, suggesting a potential link between data intensity and exposure to a wider set of servers. While lower-volume applications deviate from this trend, the privacy implications of contacting a larger number of distinct endpoints, especially for child accounts, remain worth scrutinizing.

Altogether, these findings highlight the importance of not only measuring data volume and patterns but also assess how data granularity and network reach differ between child and adult profiles. Even when overall differences are modest, the consistent directionality and elevated traffic from child accounts highlight a potential mismatch between user protections and data behavior. To investigate these concerns further, the next section presents a destination-based analysis, comparing the specific IP addresses and domains contacted by applications in child versus adult modes.

5 Destination-Based Analysis

We next examine the digital trail of outbound traffic, mapping destinations to organizations and countries. By distinguishing private, first-party, and third-party entities we uncover both who receives the data and how recipients differ between child and adult profiles. Throughout this section, we focus solely on outgoing traffic.

5.1 Contacted Organizations

Figure 5 shows the distribution of network connections by organization, while Table 3 details connection counts per application.

Across all applications and both profile types, connections frequently involve major infrastructure and platform providers, including Facebook, G-Core Labs, Amazon, Google, and Cloudflare. While we observe bigger differences between the applications than between the child and adult profiles, several privacy-relevant differences emerge between the two profiles, particularly in terms of organizational reach and connection asymmetries. We next discuss these on a per-application basis.

Gorilla Tag: Both profiles contacted the same set of organizations, with only minor differences in frequency: the adult connected more to Facebook (13 vs. 8), while the child made slightly more Google connections (3 vs. 2). This endpoint uniformity suggests similar backend behavior, but the increased Facebook traffic for adults may reflect greater telemetry, authentication, or social features.

Gym Class: This application shows one of the starkest contrasts. The child profile reached more Facebook and Akamai servers and uniquely contacted HL komm Telekommunikations and PT. Telekomunikasi Selular, entities not accessed by the adult profile. In contrast, the adult connected to Orange S.A. and the Michigan State Government, and had slightly higher counts to Google, Unity, and DigitalOcean. These divergences suggest that backend services are not only profile-dependent but may expose child users to a wider array of third-party infrastructure.

Penguin Paradise: While the adult profile connected far more often to Google (10 vs. 1), the child profile had higher Facebook (14 vs. 12) and G-Core Labs (8 vs. 7) interactions. The bigger difference in Google connections may reflect differing ad, analytics, or localization mechanisms depending on the user type.

Table 3: Top contacted organizations across applications (Org name and number of connections: # child / # adult). For ease of readability, we use red and blue to indicate cases where there are more or less connections when using the child account.

Rank	Gorilla Tag		Gym Class		Penguin Paradise		Scary Baboon		VRFS		Yeeps	
	Org.	#C/#A	Org.	#C/#A	Org.	#C/#A	Org.	#C/#A	Org.	#C/#A	Org.	#C/#A
1	Facebook	8/13	Amazon	34/34	Facebook	14/12	G-Core Labs	10/10	G-Core Labs	18/17	Amazon	30/31
2	G-Core Labs	8/8	Google	25/29	G-Core Labs	8/7	Facebook	9/8	Facebook	17/12	G-Core Labs	15/15
3	Amazon	7/7	Facebook	12/9	Microsoft	3/3	Amazon	7/6	Cloudflare London	8/8	Facebook	10/11
4	Cloudflare	5/5	Akamai	12/10	Cloudflare London	2/2	Microsoft	4/4	Servers.com	4/3	Local Router	2/1
5	Microsoft	3/3	Unity	7/8	Google	1/10	Google	3/7	Google	4/5	Cloudflare London	2/1
6	Google	3/2	Microsoft	5/5	Hi3G Access	1/1	Local Router	1/2	Fornex Hosting	2/2	Google	2/1
7	Hi3G Access	1/1	Cloudflare	5/5	Local Router	1/1	Hi3G Access	1/1	Amazon	2/1	Hi3G Access	1/1
8	Cloudflare London	1/1	DigitalOcean	5/6			Total Uptime Tech	1/1	Local Router	2/2	IBM Cloud	1/1
9	Local Router	1/1	Local Router	2/2			Cloudflare London	1/1	Hi3G Access	1/1	Orange S.A.	0/1
10	EastLink	1/1	HL komm	2/0			Cloudflare	0/1	CTG Server Ltd.	1/1		
11	Total Uptime Tech.	1/1	PT. Telkom Selular	1/0					Rostelecom	1/0		

Scary Baboon: The child profile again made more connections to Facebook (9 vs. 8) and Amazon (7 vs. 6), while the adult profile accessed Google more (7 vs. 3) and uniquely reached Cloudflare.

VRFS: This application shows consistent elevation for child profiles: more connections to Facebook (17 vs. 12), G-Core Labs (18 vs. 17), Servers.com (4 vs. 3), and Amazon (2 vs. 1). Uniquely, the child profile also contacted PJSC Rostelecom, while the adult reached one more Google server. The breadth of third-party contact for children is notable and raises concerns about data exposure.

Yeeps: Patterns here are more balanced. The adult profile had slightly more Amazon and Facebook connections and uniquely contacted Orange S.A. The child profile showed elevated counts to Cloudflare London, Local Router, and Google (each 2 vs. 1), though differences are minor overall.

Summary and Implications: Across applications, the child profile often shows equal or greater exposure to major third parties, and sometimes contact organizations not reached by the adult profile. This raises open questions about the backend logic driving these differences; e.g., whether it is tied to feature gating, telemetry, localization, or user analytics. Importantly, such asymmetries may have privacy implications; e.g., expanded third-party contact could increase tracking potential, especially if data collection practices differ by age group. Further analysis is needed to determine whether these patterns align with frameworks such as COPPA or GDPR-K.

5.2 Geolocation of Destination Servers

Understanding where user data travels is critical for assessing privacy and regulatory risks, especially for children. Cross-border transfers can expose data to inconsistent legal protections, with non-EU/EEA destinations posing heightened concerns under frameworks like GDPR and COPPA.

To gain some insights into the potential transfer of data, Figure 6 shows heatmaps of the number of connections to different countries, grouped by application and user profile.

With the exception of VRFS, the United States (US) consistently stands out as the dominant destination across nearly all applications and both profiles. Other countries that consistently appear high on the rankings are Sweden (SE), where the experiments were conducted, and the Netherlands (NL), a nearby country through which much cross-Atlantic traffic from Sweden is directed. There are, however, servers in many other countries contacted, including outside EU/EEA; e.g., United States (US), Canada (CA), Singapore

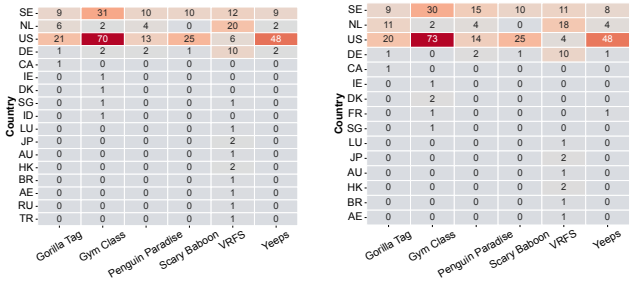

(a) Child profile. **(b) Adult profile.**

Figure 6: Geographic distribution of contacted servers, broken down per profile type. Our experiments were performed in Sweden (SE), which is part of the European Union (EU) and the European Economic Area (EEA).

(SG), Indonesia (ID), Japan (JP), Australia (AU), Hong Kong (HK), Brazil (BR), United Arab Emirates (UE), Russia (RU), and Turkey (TR). Yet, with exception for the United States, connections beyond the EU/EEA generally remain fewer in numbers and volume.

Several applications exhibit wider international footprints. Notably, *Gym Class*, *Penguin Paradise*, and *Yeeps* route a majority of traffic to the United States (US). Both *Gym Class* and *VRFS* show high geographic diversity, with the child profiles contacting more countries overall. For example, in *VRFS*, the child profile communicated with servers in 13 countries (vs. 9 for the adult), including unique access to Russia (RU), Turkey (TR), and Singapore (SG). Similarly, in *Gym Class*, the child profile uniquely contacted Indonesia and showed slightly higher activity in Germany, while the adult profile reached France (FR) and Denmark (DK).

In contrast, applications like *Gorilla Tag*, *Scary Baboon*, and *Penguin Paradise* exhibited highly localized behavior, with nearly identical country distributions between profiles and communications concentrated in three to five countries.

These observations highlight how data routing paths, and thereby potential legal exposure, can vary based on user profile and application. The fact that child accounts sometimes contact more diverse or less common jurisdictions may raise additional privacy concerns regarding oversight and data handling obligations across borders.

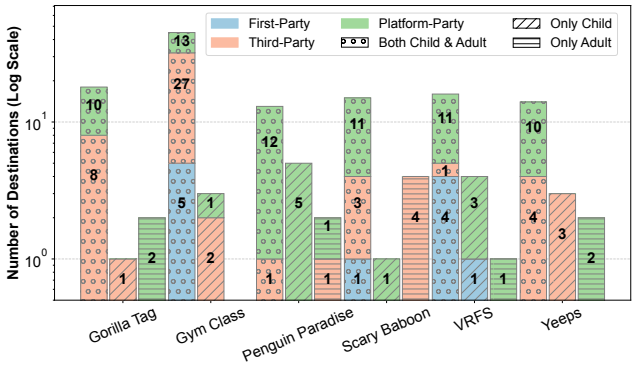

Figure 7: Distribution of destinations, categorized as First-, Third, or Platform-Party. (Note: y-axis on log scale.)

5.3 Domain Classification using TLS SNI + eSLD

To assess potential privacy risks such as user tracking and third-party exposure, we analyzed unencrypted TLS SNI fields to extract destination domain names. While most traffic was encrypted, the remaining plaintext SNI data revealed fully qualified domain names (FQDNs) contacted by each application and account type.

Method: We mapped FQDNs to organizations using effective second-level domains (eSLDs) and categorized them as first-, platform-, or third-party. A domain was labeled *first-party* if its eSLD matched the app or developer; platform-owned domains (e.g., containing *facebook* or *meta*) were marked as *platform-party*; all others were classified as *third-party*. For cloud-hosted services (e.g., *scarybaboon.azurewebsites.net*), we attributed the subdomain (*scarybaboon*) to the app developer.

To further refine classification accuracy, we cross-referenced all domains with Firebog's non-deprecated blocklists [1], part of the "Big Blocklist Collection" targeting advertising, telemetry, and malicious activity. This approach aligns with prior work [7, 25].

Analysis Based on Destination Type: Figure 7 presents the distribution of destination domains by account type and category, using a logarithmic y-axis. Since most domains were shared, they are grouped under *"Both Child & Adult"*, while profile-specific domains are shown in *"Only Child"* and *"Only Adult"*.

Across the six applications, the child profile contacted more first-party domains in 1 case, more platform-party domains in 4 cases, and more third-party domains in 3 cases.

First-party domains were observed only in Gym Class, Scary Baboon, and VRFS. Third-party domain activity was generally similar between profiles, with minor variations. Notably, platform-party domains dominated traffic across all applications except Gym Class, where third-party domains were most prominent.

These findings highlight the predominance of platform-party communications—often associated with major tracking ecosystems, and suggest that children may be exposed to even more of these domains than adults in several apps.

5.4 Analysis Based on ATS Blocklist Flags

To better understand the potential for tracking and behavioral profiling, we analyzed outbound connections against domains flagged by the Ad, Tracking, and Surveillance (ATS) blocklists.

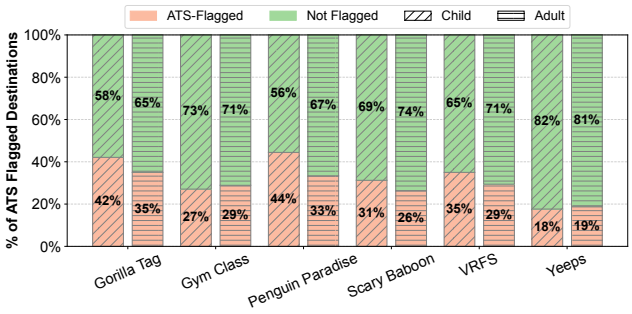

Figure 8: Share of ATS-flagged destinations per application.

Table 4: ATS-flagged and not flagged destinations by party type and user profile across all applications. Notation and color coding: Only child (red) / Both child and adult (black) / only adult (blue).

	First-Party		Third-Party		Platform-Party	
	ATS	Not	ATS	Not	ATS	Not
Gorilla Tag	–	–	1/3/-	-/5/-	-/4/-	-/6/2
Gym Class	–	-/5/-	-/9/-	2/19/-	-/4/-	1/8/-
Penguin Paradise	–	–	–	-/1/-	3/5/-	2/7/1
Scary Baboon	–	-/1/-	-/1/-	-/2/4	-/4/-	1/7/1
VRFS	–	1/4/-	-/1/-	–	2/4/-	1/7/1
Yeeps	–	–	-/1/-	3/3/-	-/2/2	-/8/-

Table 5: Packet count and average packet size per eSLD among ATS-flagged domains.

eSLD	Child Account		Adult Account	
	# Pkts	Size (bytes)	# Pkts	Size (bytes)
facebook.com	4,386	491.6	3,278	464.3
clarity.ms	69	654.8	74	655.2
fbcdn.net	63	676.3	59	702.7
bugsnag.com	38	897.2	38	893.5
gtag-cf.com	33	378.5	33	378.5
bing.com	16	628.6	16	609.8
gameanalytics.com	16	291.5	16	583.0
mxpnl.com	10	747.2	8	626.6
googletagmanager.com	8	627.3	8	626.6
google-analytics.com	8	381.0	8	256.0
mixpanel.com	2	594.0	2	594.0

Figure 8 summarizes the proportion of ATS-flagged domains contacted by each account type. Across the six applications, between 20–45% of contacted domains were ATS-flagged, with the child profile showing a higher share in 4 out of 6 cases.

Table 4 summarizes flagged domains by party classification. Most were platform-party domains contacted by both profiles. However, Gorilla Tag, Penguin Paradise, and VRFS included ATS-flagged domains contacted *only* by the child profile, while Yeeps had two unique to the adult. Gym Class and Scary Baboon showed no profile-exclusive flagged domains.

Table 5 lists the eSLDs for the flagged top-level domains together with the number of packets and average packet sizes. (For full list of all flagged FQDNs, including subdomains that here are combined, we refer to Table 7 in Appendix A.)

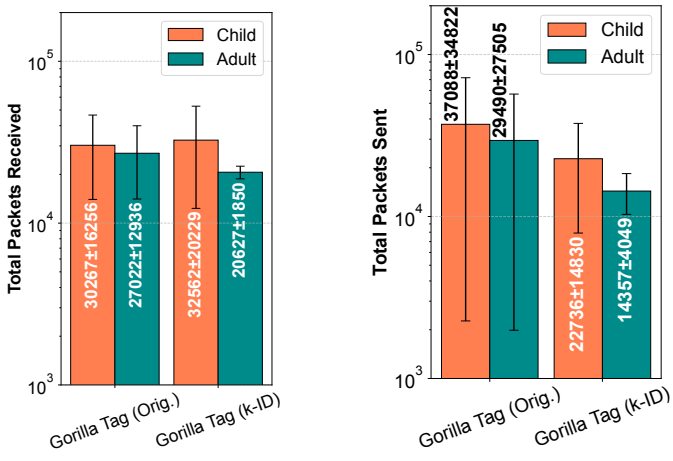

(a) Incoming traffic. **(b) Outgoing traffic.**

Figure 9: Total packets sent and received by Gorilla Tag before and after the k-ID integration update.

Two platform-party domains, *facebook.com* and *fbcdn.net*, accounted for significant traffic, with *facebook.com* alone receiving 1,108 more packets from the child profile than the adult profile. The remaining eSLDs were third-party domains. While traffic volume and average packet size were similar across profiles for most destinations, in two cases the child profile sent packets approximately 125 bytes larger on average.

These findings suggest that not only are ATS-associated domains prevalent, but in some applications children may be exposed to more of them, raising concerns about disproportionate tracking risks.

5.5 Discussion: Destination-Level Privacy Risks

Our destination-level analysis reveals some structural concerns about data flows in the VR ecosystem, particularly for child users. Across applications, both child and adult profiles routinely connected to major infrastructure and content delivery providers such as Facebook, Google, Amazon, and Cloudflare. While expected on a Meta-owned platform, the dominance of these platform-party connections complicated distinguishing necessary backend functionality from traffic potentially linked to profiling or advertising.

Facebook, in particular, was among the most frequently contacted domains, consistent with its role in service delivery. Yet, the high traffic volumes—especially from child profiles in several apps—raises questions about the extent and necessity of all such communication. While some variation may reflect differences in feature use or authentication mechanisms, the overall similarity in destination patterns suggests limited differentiation in backend handling between child and adult accounts, and we have not observed any data suggesting age-based data minimization enforcement efforts. Instead, the child profile often sends more data.

Our analysis using ATS blocklists shows that 20–45% of contacted domains are potentially associated with tracking. In four of six applications, child profiles contacted a greater proportion of ATS-flagged domains, and some were exclusive to child accounts. While blocklist inclusion does not confirm tracking, the presence of these domains, coupled with comparable or larger packet sizes for children, indicates a potential for meaningful data exchange and raises concerns about the practical enforcement of age-based privacy controls and lack of minimized data sharing for child users.

The presence of flagged domains like *graph.facebook.com* but not others like *graph.oculus.com* and *meta.graph.meta.com*, despite likely providing similar functions, highlights the difficulty of accurately capturing privacy risk in VR using conventional web-based blocklists. This further underscores the need for VR-specific tracking detection tools and greater transparency into traffic handling by platform-controlled endpoints.

Finally, the use of globally distributed infrastructure raises regulatory challenges. Cross-border traffic involving child profiles was common, sometimes exceeding that of adults, potentially exposing data to jurisdictions with weaker privacy protections or enforcement. This is particularly relevant in light of child-focused regulations such as GDPR and COPPA.

Combined, these findings suggest that current destination-level behaviors fall short of best practices in child privacy protection. The lack of clear separation in data flows, combined with platform opacity and potential tracking activity, points to a need for stronger safeguards, better transparency tools, and more effective implementation of child-specific data governance frameworks.

6 Case Study: Gorilla Tag Update

In May/June 2025, Gorilla Tag integrated k-ID[1], a company offering privacy solutions aimed at protecting children's digital identities. This update offered a unique opportunity to evaluate whether it altered the app's communication patterns; specifically in data volume, destination diversity, and third-party exposure. To evaluate the impact, we repeated our experiments on the updated version to determine whether the change reduced data exposure or shifted traffic behavior between child and adult profiles.

To ensure consistency and comparability with our previous findings, followed the same data collection steps and balanced the number of pairwise tests, resulting in 16 (2×8) additional user sessions. This allows us to analyze any differences in behavior or data flow based on user type, now also considering the possible effects of the new update.

6.1 Changes in High Level Statistics

Changes in Traffic Volume and Packet Size: Figure 9 shows that across both versions, the child profile consistently sends and receives more packets than the adult, indicating that k-ID integration did not reduce traffic volume for children.

As shown in Figure 10, the updated version introduces shifts in packet size: children receive larger incoming packets but send smaller ones compared to adults. Notably, the pre-update version had higher overall outgoing traffic, suggesting changes in transmission efficiency or payload structure post-update.

Changes in Communication Scope (Unique Destinations): Figure 11 shows an increase in unique contacted destinations for both profiles post-update, in both directions. This broader communication scope may reflect greater backend complexity or third-party involvement, raising questions about data sharing and exposure.

[1]https://www.k-id.com/

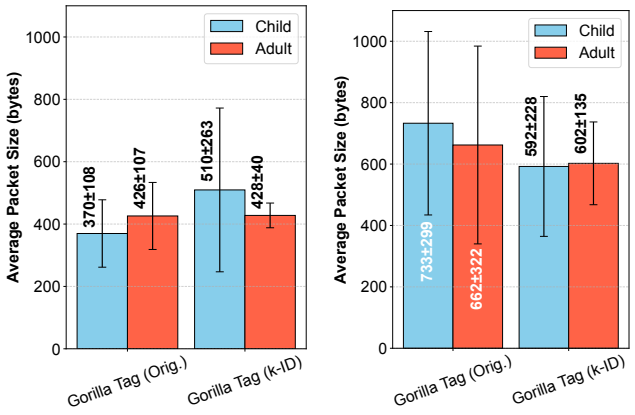

(a) Incoming traffic. (b) Outgoing traffic.

Figure 10: Average packet size (in bytes) in Gorilla Tag before and after the k-ID integration update.

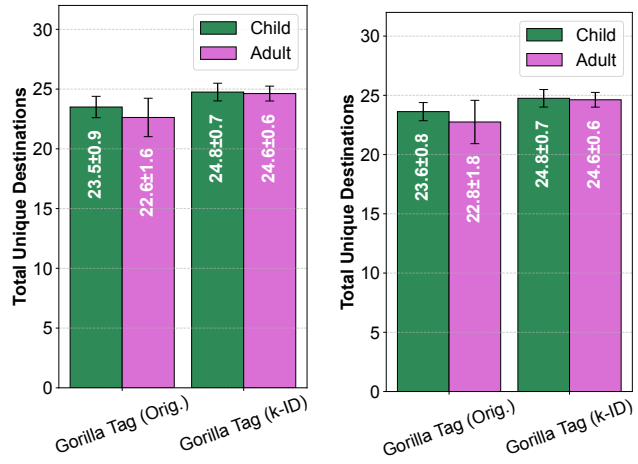

(a) Incoming traffic. (b) Outgoing traffic.

Figure 11: Total unique destinations in Gorilla Tag before and after the k-ID integration update.

Preliminary Privacy Implications: Despite the inclusion of k-ID, which emphasizes child privacy protection, the updated version of Gorilla Tag still results in: (1) higher traffic volume for child users compared to adults, (2) larger incoming packets to the child profile post-update, and (3) an increase in the number of unique destination servers for both profiles. These trends warrant further analysis into the nature of the contacted domains (e.g., third-party or ATS-flagged). We next look closer at these aspects.

6.2 Analysis of Traffic Destination Changes

To assess whether the Gorilla Tag update, introducing k-ID, has an impact on data exposure, we analyzed changes in the organizations contacted and the geographic locations of destination servers.

Contacted Organizations: Figure 12 compares the distribution of contacted organizations across app versions and user profiles. In both versions, traffic patterns were largely similar, though the k-ID version showed slightly greater divergence between profiles. In the original version, Facebook (8 vs. 13 connections), G-Core Labs (8 each), and Amazon (8 each) were the most contacted, and the main differences where that Facebook had more adult connections (8 vs. 13) and Google more child connections (3 vs. 2). The remaining organizations—including Microsoft (3), Cloudflare (5), and others (each with 1)—were contacted similarly by both profiles. In the updated version, the child profile contacted Facebook more frequently than the adult (10 vs. 9) but less with Google (2 vs. 4); otherwise, the set and number of connections per service remained the same for the two profiles. These very small shifts suggest that the integration of k-ID did not reduce the set of contacted organizations.

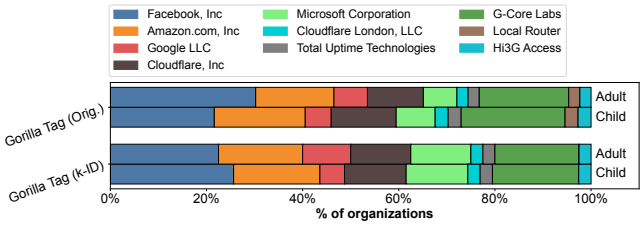

Figure 12: Organizations contacted with Gorilla Tag before and after the k-ID integration update.

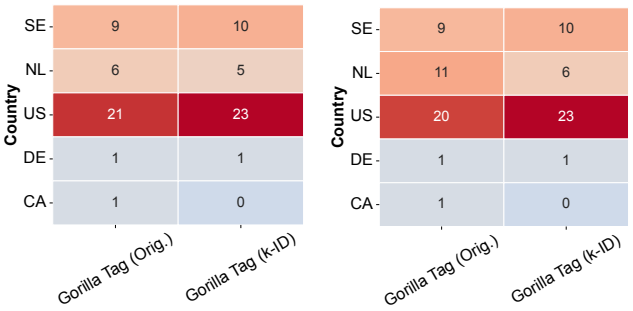

(a) Child profile. (b) Adult profile.

Figure 13: Total unique destinations by country for Gorilla Tag before and after the k-ID integration update.

Server Locations: Figure 13 shows the distribution of unique server locations by country. In both versions, the United States (20–23 servers), Sweden (9–10), and the Netherlands (5–11) hosted the most destinations. Notable changes include a slight increase in U.S. connections (21/20 to 23/23), a decrease to the Netherlands (6/11 to 5/6), and the disappearance of prior connections to Canada (1/1 to 0/0). The continued dominance of U.S.-hosted servers highlights persistent concerns about cross-border data flows, especially for child accounts.

Overall, the Gorilla Tag update does not appear to reduce data exposure or improve data localization from a privacy standpoint.

6.3 Domain and ATS Analysis

To assess whether the k-ID integration altered Gorilla Tag's communication patterns in ways that reduce third-party exposure or improve privacy, we compared domain classifications and ATS-flagged destinations before and after the update. These results are summarized in Table 6.

Table 6: Number of ATS-flagged domains in each domain category before and after Gorilla Tag's k-ID integration update. Here, we show: "flagged" / "total in category".

	First-Party		Third-Party		Platform-Party		Total	
	Child	Adult	Child	Adult	Child	Adult	Child	Adult
Before	–	–	4/9	3/8	3/10	3/12	8/19	7/20
After	–	–	0/11	0/12	5/13	3/8	5/24	3/20

Third- and Platform-Party Domains: Post-update, the number of third-party domains contacted increased slightly from 9 (8 seen by both + 1 seen by child only) to 12 (11 seen by both + 1 seen by adult only). Similarly, the number of platform-party domains contacted increased from 12 (10 both + 2 adult only) to 13 (8 both + 5 child only). Although not large, this growth suggests that backend complexity or third-party integration may have expanded post-update. While there is a shift in the set of domains seen only by the child account from third-party (+1 before) to platform-party (+5 after), also these differences are small, and we note that some platform domains provide similar services as third-party domains.

ATS-Flagged Domains: The share of ATS-flagged domains dropped notably after the update—from 7 out of 20 (35%) to 3 out of 20 (14%) for adult profile, and from 8 out of 19 (42%) to 5 out of 24 (21%) for child profile; indicating reduced exposure to known advertising, tracking, or surveillance endpoints. However, the child profile continued to contact more flagged domains than the adult.

Notably, all remaining flagged domains contacted by both profiles are Facebook operated (*www.facebook.com*, *graph.facebook.com*, *b-www.facebook.com*). While expected on a Meta platform, traffic to these domains was consistently higher, or equal, for the child profile, with substantial increases post-update. For example, packet counts for *www.facebook.com* rose from 240 (child) vs. 64 (adult) to 602 vs. 73; *graph.facebook.com* rose from 185 vs. 160 to 379 vs. 323; and *b-www.facebook.com* rose from 4 vs. 4 to 8 vs. 8.

Summary: Overall, the k-ID update corresponds with a modest reduction in ATS exposure and third-party contact, particularly for child profiles. However, the increased use of platform-party domains—some unique to children—raises new questions about where and how child data is being routed. These findings illustrate both progress and persistent opacity in post-update traffic patterns.

7 Discussion: Broad Legal Perspective

Our findings raise substantial concerns about whether child-specific privacy protections are meaningfully enforced in real-world VR applications. Despite claims of age-aware data handling, we observe that child accounts often generate greater volumes of network traffic, contact more unique destinations, and reach a comparable or higher number of ad- and tracking-related domains (as flagged by ATS blocklists). These patterns persist even when application behavior is matched under controlled conditions.

Such results challenge compliance with key provisions of the GDPR and the Convention on the Rights of the Child (CRC). Under GDPR, children are entitled to specific protections, including data minimization, purpose limitation, and transparency (Art. 5, 6, 8). Our analysis suggests that these principles are not consistently upheld. For instance, child accounts contacting more domains and generating higher traffic volumes — despite identical usage — which may indicate excessive or unnecessary data processing. This is especially problematic when traffic is directed to third-party or poorly identified services, undermining accountability and transparency.

Moreover, international data transfers compound these concerns. Many observed connections route traffic to countries outside the EU/EEA, including the United States, Russia, and Singapore. These cross-border flows are particularly significant given the *Schrems II* ruling [5], which emphasizes the risks associated with inadequate data protection in third countries. Since children's data is subject to heightened scrutiny, such transfers demand robust safeguards, none of which are visible at the network level.

Our domain classification and SNI analysis further reveal that platform- and third-party domains dominate outbound communication, often without obvious functional justification. While we cannot decrypt the contents of traffic, the prevalence of ATS-flagged domains — some exclusive to child profiles — indicates possible profiling, behavioral analytics, or telemetry beyond what is necessary for core functionality. The fact that such differences occur in the absence of user-visible consent mechanisms casts doubt on informed processing.

In summary, our results suggest a disconnect between age-based regulatory expectations and practical enforcement in VR ecosystems. The opaque backend behavior observed here reinforces the broader concern that self-regulation — without independent auditing or transparency requirements — may be insufficient to protect vulnerable users in immersive environments.

8 Conclusion

In this paper, we have presented the first network-level study to systematically compare child and adult accounts across six widely used VR applications. Our results show that child profiles often generate more traffic, contact a broader range of destinations, and reach a similar or greater share of tracking-related domains — despite using the applications in controlled, near identical ways.

These findings point to systemic shortcomings in how VR applications implement child privacy protection. Elevated traffic and minimal differentiation between account types suggest that child profiles may not receive the heightened safeguards required under frameworks like GDPR and the CRC. In some cases, children appear equally or more exposed to profiling and cross-border data flows, without meaningful notice, choice, or constraint.

From a regulatory perspective, our findings highlight a need for stronger auditing, transparency, and age-aware design in immersive platforms. Current practices often fall short of legal requirements around data minimization, consent, and child-specific protections. To safeguard children in VR ecosystems, regulators, platform providers, and developers must ensure that privacy claims are reflected not only in policies but in practice.

As VR adoption accelerates, embedding child privacy protections across protocols, applications, and governance frameworks is essential. Our findings show that current practices fall short but also reveal actionable opportunities. By strengthening platform-side data minimization, integrating parental control dashboards, and introducing VR-specific transparency tools, platforms could better align with GDPR-K and COPPA and move toward safer, more responsible immersive experiences for younger users.

Acknowledgments

This work was funded by the Swedish Research Council (VR) and the Graduate School in Computer Science (CUGS) at Linköping University. We acknowledge the use of ChatGPT-4o to assist with revising text and correct grammar, typos, and awkward phrasings.

References

[1] The Firebog. 2025. *The Big Blocklist Collection.* The Firebog. https://firebog.net
[2] Thomas Alsop. 2025. *Augmented reality (AR) and virtual reality (VR) headset companies shipment share worldwide from 2023 to 2024, by quarter.* Statista, Inc. https://www.statista.com/statistics/1407105/ar-vr-headset-companies-shipment-share/
[3] Duncan H Brown and Norma Pecora. 2014. Online data privacy as a children's media right: Toward global policy principles. *Journal of Children and Media* 8, 2 (2014), 201–207.
[4] Robin Carlsson, Sampsa Rauti, Samuli Laato, Timi Heino, and Ville Leppänen. 2023. Privacy in popular children's mobile applications: A network traffic analysis. In *2023 46th MIPRO ICT and Electronics Convention (MIPRO).* IEEE, 1213–1218.
[5] Court of Justice of the European Union. 2020. Data Protection Commissioner v Facebook Ireland Ltd and Maximillian Schrems. Case C-311/18, Judgment of 16 July 2020, ECLI:EU:C:2020:559.
[6] Cristina Fiani, Pejman Saeghe, Mark McGill, and Mohamed Khamis. 2024. Exploring the perspectives of social VR-aware non-parent adults and parents on children's use of social virtual reality. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–25.
[7] Olivia Figueira, Rahmadi Trimananda, Athina Markopoulou, and Scott Jordan. 2024. DiffAudit: Auditing Privacy Practices of Online Services for Children and Adolescents. In *Proceedings of the ACM on Internet Measurement Conference (IMC).* 488–504.
[8] Gonzalo Munilla Garrido, Vivek Nair, and Dawn Song. 2023. Sok: Data privacy in virtual reality. *arXiv preprint arXiv:2301.05940* (2023).
[9] Meta Horizon. 2024. *Get Age Category API.* Meta Platforms, Inc. https://developers.meta.com/horizon/documentation/native/ps-get-age-category-api/
[10] Ismat Jarin, Yu Duan, Rahmadi Trimananda, Hao Cui, Salma Elmalaki, and Athina Markopoulou. 2025. BehaVR: User Identification Based on VR Sensor Data. *Proceedings on Privacy Enhancing Technologies (PoPETs)* 2025, 1 (2025), 399–419.
[11] Sonia Livingstone, Mariya Stoilova, and Rishita Nandagiri. 2019. *Children's data and privacy online: growing up in a digital age: an evidence review.* London School of Economics and Political Science.
[12] Minzhao Lyu, Rahul Dev Tripathi, and Vijay Sivaraman. 2023. Metavradar: Measuring metaverse virtual reality network activity. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7, 3 (2023), 1–29.
[13] Meta. 2025. *Supplemental Meta Platforms Technologies Terms of Service.* Meta Platforms, Inc. https://www.meta.com/se/en/legal/supplemental-terms-of-service/
[14] Mark Roman Miller, Fernanda Herrera, Hanseul Jun, James A. Landay, and Jeremy N. Bailenson. 2020. Personal identifiability of user tracking data during observation of 360-degree VR video. *Scientific Reports* 10, 17404 (2020).
[15] Tehila Minkus, Kelvin Liu, and Keith W Ross. 2015. Children seen but not heard: When parents compromise children's online privacy. In *Proceedings of the international conference on World Wide Web.* 776–786.
[16] Gonzalo Munilla Garrido, Vivek Nair, and Dawn Song. 2023. SoK: Data Privacy in Virtual Reality. In *Proceedings on Privacy Enhancing Technologies (PoPETS)*, Vol. 2024. 21–40. Issue 1.
[17] Vivek Nair, Gonzalo Munilla Garrido, Dawn Song, and James F O'Brien. 2022. Exploring the privacy risks of adversarial VR game design. *arXiv preprint arXiv:2207.13176* (2022).
[18] Vivek Nair, Wenbo Guo, Justus Mattern, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2023. Unique Identification of 50,000+ Virtual Reality Users from Head & Hand Motion Data. In *Proceedings of the 32nd USENIX Security Symposium (USENIX Security).* 895–910.
[19] United Nations. 1989. Convention on the Rights of the Child. UN Treaty Series. United Nations General Assembly Resolution 44/25.
[20] Naheem Noah, Sommer Shearer, and Sanchari Das. 2022. Security and Privacy Evaluation of Popular Augmented and Virtual Reality Technologies. In *Proceedings of the IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence, and Neural Engineering (IEEE MetroXRAINE).*
[21] Ilesanmi Olade, Charles Fleming, and Hai-Ning Liang. 2020. Biomove: Biometric user identification from human kinesiological movements for virtual reality systems. *Sensors* 20, 10 (2020), 2944.
[22] European Parliament and Council. 2016. General Data Protection Regulation. Official Journal of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council.
[23] Ken Pfeuffer, Matthias J. Geiger, Sarah Prange, Lukas Mecke, Daniel Buschek, and Florian Alt. 2019. Behavioural Biometrics in VR: Identifying People from Body Motion and Relations in Virtual Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems.* 1–12.
[24] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, Serge Egelman, et al. 2018. "Won't somebody think of the children?" examining COPPA compliance at scale. In *Proceedings of the Privacy Enhancing Technologies Symposium (PETS).*
[25] Rahmadi Trimananda, Hieu Le, Hao Cui, Janice Tran Ho, Anastasia Shuba, and Athina Markopoulou. 2022. OVRseen: Auditing network traffic and privacy policies in oculus VR. In *Proceedings of the 31st USENIX security symposium.* 3789–3806.
[26] Yi Wu, Cong Shi, Tianfang Zhang, Payton Walker, Jian Liu, Nitesh Saxena, and Yingying Chen. 2023. Privacy leakage via unrestricted motion-position sensors in the age of virtual reality: A study of snooping typed input on virtual keyboards. In *2023 IEEE Symposium on Security and Privacy (SP).* IEEE, 3382–3398.
[27] Yuxia Zhan, Yan Meng, Lu Zhou, and Haojin Zhu. 2023. Vetting privacy policies in VR: a data minimization principle perspective. In *Proceedings of the IEEE Conference on Computer Communications INFOCOM Workshops.* 1–2.

Table 7: Packet count and average packet size per ATS-flagged domains that both child and adult account shared.

ATS flagged Domain	Child Account		Adult Account	
	# pkt	Avg. Size (bytes)	# pkt	Avg. Size (bytes)
www.facebook.com	1333	289.3	1230	286.9
graph.facebook.com	3007	583.0	1982	583.3
scontent-iad3-1.xx.fbcdn.net	63	676.3	59	702.7
b-www.facebook.com	33	336.6	30	321.0
k.clarity.ms	31	676.0	34	671.6
www.clarity.ms	30	643.6	32	646.9
sessions.bugsnag.com	30	924.9	30	922.0
c.bing.com	16	628.6	16	609.8
title-data.gtag-cf.com	16	379.0	16	379.0
cdn.mxpnl.com	10	747.2	10	755.2
temp-prod.gtag-cf.com	9	378.0	9	378.0
web.facebook.com	9	583.0	4	583.0
www.google-analytics.com	8	381.0	8	381.0
c.clarity.ms	8	614.3	8	618.3
notify.bugsnag.com	8	793.4	8	786.8
www.googletagmanager.com	8	627.3	8	626.6
auth-prod.gtag-cf.com	8	378.0	8	378.0
api.gameanalytics.com	8	583.0	16	583.0
edge-mqtt.facebook.com	4	267.0	3	267.0
api-js.mixpanel.com	2	594.0	2	594.0

A Additional Statistics for ATS-Flagged Domains

Table 7 shows the packet count and average packet size, with separate columns for the child and adult account, for each ATS-flagged domain that communicated with both account types.