

Identifying Organizations Receiving Personal Data in Android Apps

David Rodriguez^a, Miguel Cozar^b and Jose M. Del Alamo^{*c}
ETSI Telecomunicación, Universidad Politécnica de Madrid, Madrid, Spain

Keywords: Privacy, Data Protection, Personal Data, Data Controller, First-Party, Corporation, Android, Apps.

Abstract: Many studies have demonstrated that mobile applications are common means to collect massive amounts of personal data. This goes unnoticed by most users, who are also unaware that many different organizations are receiving this data, even from multiple apps in parallel. This paper assesses different techniques to identify the organizations that are receiving personal data flows in the Android ecosystem, namely the WHOIS service, SSL certificates inspection, and privacy policy textual analysis. Based on our findings, we propose a fully automated method that combines the most successful techniques, achieving a 94.73% precision score in identifying the recipient organization. We further demonstrate our method by evaluating 1,000 Android apps and exposing the corporations that collect the users' personal data.

1 INTRODUCTION

Extensive research (Wongwiwatchai, 2020) (Gamba, 2020) has shown that mobile applications are eager collectors and leakers of their users' personal data. This goes unnoticed by most users (Balebako, 2013) (Kim, 2015) and even the app developers (Balebako, 2014), who are also unaware that many different organizations are receiving this data (Razaghpanah, 2018), even from multiple apps in parallel.

Properly identifying the organizations that receive personal data is becoming increasingly important for different stakeholders. For example, supervisory authorities must carry out investigations on the relationship between the source and destination of some personal data flows to understand a system's compliance with e.g., legal requirements for international transfers of personal data. Similarly, developers may want to check to which organizations they are sending their users' personal data.

Also, privacy researchers can leverage this information to audit what corporations are avidly collecting massive amounts of personal data.


However, the apps' privacy policies often fail to include the third parties with which personal data are shared. Although a dynamic analysis of the app can reveal the personal data flows and the destination


domains, identifying the organizations receiving the data may become challenging due to e.g., WHOIS accuracy and reliability issues (Ziv, 2021), or the lack of details in SSL certificates ("SSL survey", 2022).


In this context, we aim to advance the fundamental understanding of the accuracy of different methods to identify the organizations receiving personal data. To this end, we have first gathered a ground truth of domain holders, and then we have assessed three different methods against it. To improve the overall performance, we have integrated them into a new method, showing a high precision level (94.73%). Finally, to demonstrate its applicability at scale we have applied the new method to understand the organizations and corporations that are receiving personal data on a random sample of 1,000 Android apps.

2 RELATED WORK

WHOIS ("Technical Overview | ICANN WHOIS", 2022) is the standard protocol for retrieving information about registered domains and their registrants, including the domain holder's identity, contact details, domain expiration date, etc. However, several issues have been reported ("Current Issues |

^a  <https://orcid.org/0000-0002-0911-4608>

^b  <https://orcid.org/0000-0002-8958-6217>

^c  <https://orcid.org/0000-0002-6513-0303>

* Corresponding author

ICANN WHOIS”, 2022) that threatens WHOIS reliability and accuracy.

An SSL certificate is another source of information about the authority holding a domain as they digitally bind a cryptographic key, a domain, and, sometimes, the domain holder details. SSL certificates are usually issued by a Certification Authority (CA), who checks the right of the applicant organization to use a specific domain name and may check some other details depending on the certificate type issued. However, recent studies report that as little as 30% of the certificates available contain details about the applicant (“SSL survey”, 2022).

Due to the limitations of the WHOIS service and SSL certificates to provide information on a domain holder some researchers have looked at privacy policies as an alternative means of identifying an authority controlling a domain. In many jurisdictions, the privacy policy must include the identity and the contact details of the personal information handler, or first party/data controller in data protection parlance. For example, the European General Data Protection Regulation (GDPR) (“Regulation 2016/679 of the European Parliament and of the Council”, 2016) sets the requirements (Article 13) for the information to be provided to data subjects when their personal data are collected, including “*the identity and the contact details of the controller*”. It is reasonable to assume that the data controller for a given subdomain is also the authority holding that domain, and vice versa.

Previous work on privacy policy analysis has focused on determining the presence/absence of specific information in the text (Torre, 2020) (Harkous, 2018), mainly to assess compliance with legal requirements for transparency. We advance these works further extracting the controller’s identity by applying Named-Entity Recognition (NER) techniques. Hosseini et al. (Hosseini, 2020) also used NER techniques to identify third party entities on privacy policies, but their goal was to recognize all entities of a class (i.e. organization) in a policy, while we aim to get only one output (the data controller) from all possible organizations (i.e. first party, third parties) recognized in the policy text.

Similar to our work, WebXray (Libert, 2021) also provides information about the holder of a given domain. However, identification is achieved through information extracted from WHOIS and other methods (e.g. web search).

The works mentioned above focus on analyzing privacy policies to identify and extract specific information. Our work aims to further apply these

techniques to discover third-party organizations receiving personal data from Android apps. To achieve this, it is necessary to intercept and analyze the outgoing personal data flows. This has been the subject of previous works too, mainly through static, dynamic or hybrid analysis techniques (Guaman, 2020). Similar to these works, we have leveraged our previous research on a platform for the dynamic analysis of Android apps that allows carrying out assessments of data protection features in Android apps at scale (Guaman, 2021).

3 METHOD

WHOIS service consultation, SSL certificates inspection, and privacy policies analysis are three different methods to obtain information on an organization receiving personal data. In all cases, we depart from the information obtained when intercepting a personal data flow i.e. a Fully Qualified Domain Name (FQDN) receiving the data.

We leverage WHOIS records to learn the organization holding a given domain. Our queries ask for FQDNs as well as the corresponding Second Level Domains (SLDs), obtaining better results when asking for the latter. We tried to retrieve the information using the *python-whois*¹ (over eight million downloads, 180,000 in the last month) observing incomplete or missing fields that were not correctly parsed, particularly those related to the Registrant Organization identity. This is probably due to the absence of a consistent schema followed during the domain registration process, as noted by previous research. Thus, we developed our own code to query the WHOIS service using the command line tool and parsing the Registrant Organization details. We discarded values hidden for privacy reasons by detecting keywords such as “redacted” or “privacy”.

To leverage the SSL certificates, we set up HTTPS connections to the destination FQDN, intercepted the certificate sent by the server, and analysed it for the holder details.

Finally, we look for the privacy policy governing the FQDN and analyse it to extract the data controller identity. First, we get the SLD from the FQDN and compose an HTTP URL, to which we send an HTTP request aiming to be redirected to a valid home page. Otherwise, we leverage the Google search engine to find the domain’s home page URL. Once the home page is found we search for the privacy policy. We have followed two different approaches to get it: 1)

¹ <https://pypi.org/project/python-whois/>

Scraping the home page with Selenium and, in case a valid policy is not found, 2) searching the policy on the Google search engine. When a potential privacy policy is found its text is downloaded and kept for further analysis. We relied on Selenium to deal with dynamic JavaScript code displaying the privacy policy. In our experimental tests, these techniques correctly found 65% of the privacy policies governing the target domain.

Once a potential privacy policy is collected, its language is checked with the *langdetect*² python package, and non-English texts are discarded. Afterwards, a supervised Machine Learning model based on Support Vector Machines (SVM) checks whether the text is indeed a privacy policy. We trained the model with 195 manually annotated texts, achieving 98.76% precision, 97.56% recall and 98.15% F1 score when evaluated against 100 unseen English texts.

To identify the data controller in the privacy policy, we first select the paragraphs of the text where it is likely to appear. This selection is based on a bag of words that seeks keywords empirically demonstrated to be closer to the data controller mention (e.g. keywords such as “we”, “us” found in many privacy policies). Then, NER techniques are applied to the selected paragraphs to identify the data controller name. We have used *SpaCy*³ for this, which provides two different trained NER models, one prioritizing efficiency and another favouring accuracy. The efficiency model showed poorer results and therefore the accuracy-based model was implemented. We validated the combination of the bag of words and the NER performance with 140 privacy policies achieving 92.14% accuracy, 95.41% precision, 94.54% recall and 94.97% F1 score.

We used a ground-truth dataset to evaluate the performance of the three methods in identifying the controller behind a given domain. The dataset includes 100 unique domains manually annotated with the organization holding the domain. The domains were randomly selected from those obtained with the experimental setup described in section 4.

The evaluation result for each method is either 1) a given value for the domain holder, which can be right (i.e. true positive - TP) or wrong (i.e. false positive - FP), or 2) no value (i.e. Null), in case the method cannot determine a specific organization. For this reason, none of the results can be considered as a False Negative or True Negative. Thus, no accuracy, recall or F1-score can be measured and only precision will be considered to assess each method performance.

The SSL certificate inspection method retrieved 99 certificates out of the 100 domains fed, with 30 of them containing the organization name. The missing certificate could not be obtained because this domain uses HTTP protocol. After a manual check, only 20 of those organization names were correct and 10 were wrong. These results translate into a 66.67% precision score and only 20% identified organizations.

The WHOIS consultation method failed to retrieve information on 10 out of the 100 SLDs. From the remaining 90 registries, 37 were correctly obtained, 2 were wrong, 24 did not contain the Registrant Organization field, 2 had an empty value on this field, and 35 included hidden values due to privacy reasons. This results into a 94.87% precision and 37% identified organizations.

The privacy policy analysis method obtained the highest rate of TP (Table 1). The evaluation is applied to the whole pipeline including the extraction of the privacy policy associated with the FQDN and the extraction of the data controller name from the policy text.

Table 1: Evaluation results.

| | TP | FP | N | Precision |
|------------------------|----|----|----|-----------|
| SSL certificate | 20 | 10 | 69 | 66.67% |
| WHOIS | 37 | 2 | 61 | 94.87% |
| Privacy Policy | 56 | 4 | 40 | 93.34% |
| Privacy Policy + WHOIS | 72 | 4 | 24 | 94.73% |

Interestingly, the combination of the privacy policy analysis method as first choice and the WHOIS consultation as second choice outputs the best results, showing almost the same precision score (94.73% against 94.87%) while reducing considerably the number of null results. We have applied this combination to the evaluation of 1,000 Android apps.

4 ANDROID APPS EVALUATION

We developed a controlled experiment leveraging our previous work on personal data flow interception and analysis in Android apps (Guaman, 2021).

Basically, this is a pipelined microservices-based platform able to automatically 1) search, download, install, run, and interact with Android apps, and 2) intercept and analyse outgoing network connections.

² <https://pypi.org/project/langdetect/>

³ <https://spacy.io/>

The platform intercepts HTTP/HTTPS connections established by the app under assessment through a Man in the Middle (MITM) proxy, logging information on the destination FQDNs and the payload of the messages. The platform bypasses most certificate pinning protections using the Frida tool. The payload analysis accounts for different obfuscation, encoding, and hashing techniques (e.g. Base64, MD5, SHA1, SHA256, etc.) that might be used by the app developer to evade the detection of personal data. Finally, the IP address of the personal data recipient is geolocated (Cozar, 2022), providing more details for their subsequent analysis.

This platform was fed with a list of 1,000 random Google Play Store apps, from which 943 managed to execute. The apps were downloaded and tested between 23 February and 2 March 2022. They were installed and executed on a mobile device Xiaomi Redmi 8 running Android 9 (API 28). The Android Monkey (“UI/Application Exerciser Monkey”, 2022) was used to automatically stimulate the app.

4.1 Results

Our platform identified 99,300 personal data flows from 767 apps to 1,004 unique domains during the experiment. A huge portion (96.46%) of these data flows correspond to HTTPS connections. A subset of 1,849 HTTPS data flows could not be analysed due to their further security protection we could not break. Interestingly, we found 3,515 (3.54%) HTTP connections containing personal data, which is an insecure practice. Also, personal data flows were not identified in 176 apps.

Fig. 1 shows the number of connections sending personal data to the top-20 destination FQDNs. As expected, most of these domains serve analytics, marketing or monitoring purposes.

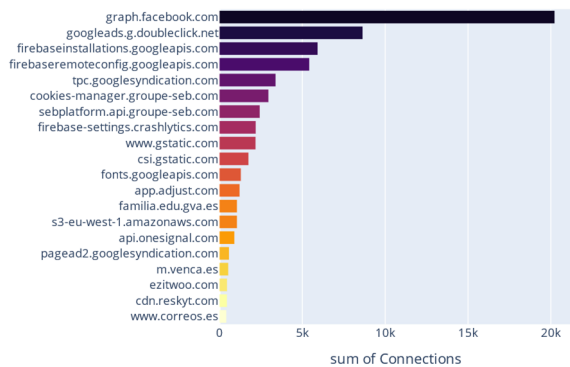


Figure 1: Personal data flows destinations.

⁴ <https://www.crunchbase.com/>

We further applied our method to identify the companies receiving personal data (Fig. 2). Overall, we were able to find them in 77.42% (76,878) of the personal data flows, representing 68.92% (692) of the unique destination domains. The top-6 companies to which most apps send personal data provide analytics and marketing services. The list is led by Google, receiving data from 646 apps.




Figure 2: Companies receiving the personal data.

We have leveraged the Crunchbase⁴ database to further understand which corporations are beneath these companies, showing the parent company and all the subsidiaries that collect personal data. Fig. 3 shows how some of them collect data from different subsidiaries, being the aggregated data higher than expected as for Fig 2. The example of AppLovin is quite representative, as it receives personal data through AppLovin (monetization tools), but also Adjust (developers’ support) and MoPub (advertisement). The result is a whole ecosystem of companies collecting data that situate the corporation on our top-3. Meta is another example of a company aggregating subsidiaries and collecting data from different sources, including Instagram and Branch.




Figure 3: Corporations receiving personal data.

5 CONCLUSIONS

This paper has described a new method that combines information from different sources to identify organizations receiving personal data. The method achieves a 94.73% precision and has been applied to identify the corporations receiving personal data from 1,000 Android apps. We are working on applying these results at scale to have a clearer picture of the personal data collectors in the mobile ecosystem.

ACKNOWLEDGEMENTS

This work was partially supported by the Comunidad de Madrid and Universidad Politécnica de Madrid through the V-PRICIT Research Programme Apoyo a la realización de Proyectos de I+D para jóvenes investigadores UPM-CAM, under Grant APOYO-JOVENES-QINIM8-72-PKGQ0J. The identification of the relationships between companies was possible thanks to Crunchbase, who kindly allowed us free access to its API for this research.

REFERENCES

- Cozar, M., Rodriguez, D., Del Alamo, J., Guaman, D. (2022). Reliability of IP Geolocation Services for Assessing the Compliance of International Data Transfers. In *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*.
- SSL Survey | Netcraft. Retrieved April 5, 2022, from <https://www.netcraft.com/internet-data-mining/ssl-survey/>.
- “Technical Overview | ICANN WHOIS.” *ICANN LOOKUP*, Internet Corporation for Assigned Names and Numbers, [whois.icann.org/en/technical-overview](https://www.icann.org/en/technical-overview). Accessed 27 May 2022.
- “Current Issues | ICANN WHOIS.” *ICANN | LOOKUP*, Internet Corporation for Assigned Names and Numbers, [whois.icann.org/en/current-issues](https://www.icann.org/en/current-issues). Accessed 27 May 2022.
- “UI/Application Exerciser Monkey.” *Android Developers*. Retrieved May 27, 2022, from <https://developer.android.com/studio/test/other-testing-tools/monkey>.
- Guaman, D. S., Del Alamo, J. M., & Caiza, J. C. (2021). GDPR Compliance Assessment for Cross-Border Personal Data Transfers in Android Apps. *IEEE Access*, 9, 15961-15982.
- Libert, T., Desai, A., & Patel, D. (2021). Preserving Needles in the Haystack: A search engine and multi-jurisdictional forensic documentation system for privacy violations on the web.
- Ziv, M., Izhikevich, L., Ruth, K., Izhikevich, K., & Durumeric, Z. (2021). ASdb: a system for classifying owners of autonomous systems. In *Proceedings of the 21st ACM Internet Measurement Conference (pp. 703-719)*.
- Guaman, D. S., Del Alamo, J. M., & Caiza, J. C. (2020). A systematic mapping study on software quality control techniques for assessing privacy in information systems. *IEEE access*, 8, 74808-74833.
- Torre, D., Abualhaija, S., Sabetzadeh, M., Briand, L., Baetens, K., Goes, P., & Forastier, S. (2020). An ai-assisted approach for checking the completeness of privacy policies against gdpr. In *2020 IEEE 28th International Requirements Engineering Conference (RE) (pp. 136-146)*.
- N. Wongwiwatchai, P. Pongkham, and K. Sripanidkulchai, (2020). “Detecting personally identifiable information transmission in android applications using light-weight static analysis,” *Comput. Secur.*, vol. 99, 2020, doi: 10.1016/j.cose.2020.102011.
- Gamba, J., Rashed, M., Razaghpanah, A., Tapiador, J., & Vallina-Rodriguez, N. (2020, May). An analysis of pre-installed android software. In *2020 IEEE Symposium on Security and Privacy (SP)* (pp. 1039-1055).
- Hosseini, M. B., Pragyani, K. C., Reyes, I., & Egelman, S. (2020, November). Identifying and classifying third-party entities in natural language privacy policies. In *Proceedings of the Second Workshop on Privacy in NLP (pp. 18-27)*.
- Harkous, H., Fawaz, K., Lebret, R., Schaub, F., Shin, K. G., & Aberer, K. (2018). Polisis: Automated analysis and presentation of privacy policies using deep learning. In *27th USENIX Security Symposium (USENIX Security 18) (pp. 531-548)*.
- Razaghpanah, A., Nithyanand, R., Vallina-Rodriguez, N., Sundaresan, S., Allman, M., Kreibich, C., & Gill, P. (2018, February). Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. In *The 25th Annual Network and Distributed System Security Symposium (NDSS 2018)*.
- “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation).” *EUR-Lex*, Publications Office, 27 Apr. 2016, eur-lex.europa.eu/eli/reg/2016/679/oj.
- Balebako, R., Marsh, A., Lin, J., Hong, J. I., & Cranor, L. F. (2014). The privacy and security behaviors of smartphone app developers. In *The 21st Annual Network and Distributed System Security Symposium (NDSS 2014)*.
- Balebako, R., Jung, J., Lu, W., Cranor, L. F., & Nguyen, C. (2013, July). “Little brothers watching you” raising awareness of data leaks on smartphones. In *Proceedings of the Ninth Symposium on Usable Privacy and Security (pp. 1-11)*.