
PrivFramework: A System for Configurable and Automated Privacy Policy Compliance

Usmann Khan

Georgia Institute of Technology
ukhan35@gatech.edu

Lun Wang

University of California, Berkeley
lunwang@berkeley.edu

Jithendaraa Subramanian

National Institute of Technology, Tiruchirappalli
jithen.subra@gmail.com

Joseph P. Near

University of Vermont
jnear@uvm.edu

Dawn Song

University of California, Berkeley
dawnsong@cs.berkeley.edu

Abstract

Today’s massive scale of data collection coupled with recent surges of consumer data leaks has led to increased attention towards data privacy and related risks. Conventional data privacy protection systems focus on reducing custodial risk and lack features empowering data owners. As an end user there are limited options available to specify and enforce one’s own privacy preferences over their data. To address these concerns we present PRIVFRAMEWORK, a user-configurable framework for automated privacy policy compliance. PRIVFRAMEWORK allows data owners to write powerful privacy policies to protect their data and automatically enforces these policies against analysis programs written in Python. Using static analysis PRIVFRAMEWORK automatically checks authorized analysis programs for compliance to user-defined policies.

1 Introduction

Today our sensitive data is being collected and analyzed by our service providers at a scale that we are not, and cannot become, aware of. Though we often nominally own our data, granting others access to it is no simple task. Solutions like OAuth [9] exist to manage data sharing but limit the permission scopes that can be granted to those chosen and implemented by the service provider. Additionally these solutions do not give us control over what happens to our data after permission is granted. For example, Facebook users agreed to share data with app developers but were surprised when in 2018 it came to light that that data had been subsequently shared with political consulting firm Cambridge Analytica [11].

In an effort to combat this, governments have introduced regulations to limit the scope of unauthorized data sharing. These data privacy regulations include the European Union’s *General Data Protection Regulation (GDPR)* [6], California’s *California Consumer Privacy Act (CCPA)* [1] and the United States’ *Health Insurance Portability and Accountability Act (HIPAA)* [4]. Unfortunately, today’s data processing systems make compliance with these regulations very difficult. This is due in large part to the fact that many existing data processing systems were designed and deployed prior to the conception of these regulations.

These two problems: the difficulty involved with complying with data regulation and the inability to control our own sensitive data, are deeply interrelated. The root cause is that today’s approaches for collecting and processing sensitive personal data require data owners to *give up control* over their data. Once data has left the owner’s device, the data owner must *trust* that the organization collecting the data will act in their interest and the organization must go through extreme lengths to ensure their compliance. As can be seen by the numerous recent cases of data misuse [11][8][19], many organizations have demonstrated that they are not deserving of this trust.

If organizations could ensure verifiable and automatic compliance with privacy policies, it would be safe to trust them with our data. Additionally the ability to specify these policies themselves would extend data owners far greater control and ownership over their data.

PRIVFRAMEWORK: returning control to data owners. This paper presents PRIVFRAMEWORK, an end-to-end framework for building and deploying scalable systems that collect and process sensitive data while enforcing security and privacy policies specified by data owners. PRIVFRAMEWORK improves on previous approaches by allowing data owners to specify both *who* may process their data and *how* it may be used.

In PRIVFRAMEWORK, data owners submit encrypted data to data collectors, and the data can be decrypted and processed only within *Trusted Execution Environments (TEEs)* that keep the data confidential and ensure the integrity of the results. Data owners also submit policies governing their data, and PRIVFRAMEWORK ensures via static analysis that processing tasks performed by the data collector do not violate these policies. These two properties together ensure that data is used in ways that are consistent with its owners’ wishes, without requiring trust in the organizations which collect the data.

Related Work With the increasing visibility of data privacy issues, there has been increased attention paid to advancing technology in the data privacy space. Our work is most closely related to three directions of privacy research - (1) Data access control & de-identification, (2) - Pure differential privacy tools, (3) Policy enforcement systems. In Data access control systems there have been advances such as Google DLP [7] and Amazon Macie [3]. These systems exist to reduce custodial risk but do not allow end users to specify their privacy preferences. Pure differential privacy tools exist in both centralized and decentralized variants and function by adding noise to data. By reducing the signal-to-noise ratio in the data, these solutions often render the data too noisy to provide accurate results. Other privacy policy enforcement systems exist such as those presented by Chowdhury et. al [5] and Sen et. al [17] but these systems do both offer the flexibility of PRIVFRAMEWORK while also stopping violations *before they occur*.

2 The PRIVFRAMEWORK System

PRIVFRAMEWORK provides an end-to-end framework for building configurable privacy-sensitive systems. The framework is comprised of a *client-side API* for building applications that collect data and submit it for analysis, a *collection & analysis platform* for receiving submitted data, storing it, and executing analyses, and an *analysis API* for writing statistical analyses and machine learning pipelines against the collected data.

Consider the motivating example of a consumer financial application used to track personal budgeting. Currently this requires the use of an interface (e.g. Plaid [16] or Yodlee [20]) that requests end users’ banking username and passwords, two pieces of extremely sensitive information. Requesting these credentials just to obtain read-only access to transaction history is an overreach that creates a severe and unnecessary privacy hazard. In addition there is no way to limit the scopes of the services’ access to data; they have the ability to read and extract any and all information in the account in perpetuity. PRIVFRAMEWORK solves these problems by allowing individuals to control the amount of information shared with service providers while simultaneously enabling service providers to provide strong privacy guarantees to reduce custodial risk and build customer trust.

To accomplish this we implement the following workflow:

1. The Data Owner uses a client application to upload their sensitive data, along with a formal policy specifying how the data may be processed, to a Data Manager. We use the Oasis Parcel Platform [13] to store this data, encrypted with the Owner’s key.

2. The Data Owner uses the Oasis Steward application to authorize Analysts to submit analysis programs to the Data Manager that will be run over the data.
3. An authorized Analyst submits an analysis program to the Data Manager. This program is statically analyzed and constrained according to the specified policy by the Data Manager.
4. The Data Manager validates the analysis program and returns to the Analyst either the analysis results or a residual policy specifying unsatisfied constraints. At this point the Data Owner has not seen the details of the analysis program, and the analyst has not had direct access to the data.
5. If applicable, the Analyst composes a summary or result for the Data Owner.




Figure 1: Diagram of interactions in the PRIVFRAMEWORK system. The circled numbers refer to the above workflow. The Data Manager represents the TEE-enabled system where decryption, compliance checks and analysis occur.

Enabling a secure and private workflow like the one described is challenging in several domains. The tool must adhere to strict privacy guarantees, have ergonomic APIs that developers can leverage easily, and provide accessible means for end users to specify their privacy policies. To address these challenges we present PRIVFRAMEWORK. In the following sections we present our approach to specifying and enforcing policies, available APIs, the system architecture, and detail a demonstration of the previously described budgeting application designed using PRIVFRAMEWORK.

2.1 Specifying and Enforcing Privacy Policies

A major challenge in designing privacy sensitive applications is accurately foreseeing the system’s eventual use cases and breadth of end user preferences. To address this issue PRIVFRAMEWORK is designed to dynamically enforce user specified privacy policies. We have developed a static analysis tool and a Domain Specific Language (DSL) to allow users to configure their own privacy policies, avoiding the need to specify parameters and options in advance. The PRIVFRAMEWORK system relies on PRIVGUARD, our static analysis framework, and PRIVPOLICY, our DSL, to specify and enforce policies.

Data Capsules. Data submitted to PRIVFRAMEWORK is stored in *data capsules* [18], which pair submitted data with a formal policy governing its use. PRIVFRAMEWORK provides analysts with a standard Python environment for processing data capsules, and uses static analysis to ensure compliance with the relevant policies. Our policy analysis supports libraries including Pandas [14], NumPy [10], TensorFlow [2], and PyTorch [15] out of the box. Organizations can add extend the static analyzer by implementing tracing stubs for their library of choice.

The data capsule system is *compositional*: processing tasks can be split up into multiple steps that form a processing pipeline. Each program in the pipeline outputs a new data capsule with a new, residual, policy. As programs in the pipeline satisfy policy requirements, those changes are reflected in the new policy attached to the output data capsule. The contents of a data capsule cannot be viewed by the analyst until all of the requirements of its policy have been satisfied.

Policies. When submitting data to PRIVFRAMEWORK, data owners attach policies that govern how the data may be processed. Policies in PRIVFRAMEWORK are encoded using the PRIVPOLICY language [18], which allows policies to encode both security and access control requirements (e.g. “only my doctor may view my medical records”) and privacy requirements (e.g. “my data must be

aggregated with data from 100 others before it is processed” or “differential privacy must be used when processing my data”).

Static Analysis. We use PRIVGUARD to statically enforce the privacy policies over the Data Owner’s data. Only after verifying that the analysis program will not violate any policies is the analysis actually run with access to the data. In the case of noncompliance we also generate a residual policy. This residual policy is packaged with an intermediate analysis state into a new data capsule. If the analysis is compliant then the program’s outputs are returned. Otherwise the analyst is provided a reference to the new intermediate capsule. This analysis is done in a TEE to prevent the Data Manager from accessing the data at any time.

2.2 Developing Applications with PRIVFRAMEWORK

As a tool meant for developers, all of our components are designed for scalability and interoperability with existing systems. The client-side API is easily integrated with Android applications, and can leverage data stored on the device or drawn from 3rd party sources. The analysis API is implemented as a library for Python programs, so that analysts can re-use existing code and do not need to learn a new programming language. The collection & analysis platform is built on the Oasis Parcel platform and runs within the AMD SEV [12] Trusted Execution Environment. The collection & analysis platform is responsible for collecting and processing data, while enforcing data owners’ policies.

Client-side Application Development. PRIVFRAMEWORK’s client-side API provides application developers with tools for the rapid development of PRIVFRAMEWORK applications. Existing applications can be modified to submit data to PRIVFRAMEWORK in just a few lines of code. The API also provides tools to help data owners select a policy for submitted data. Applications can leverage the API to allow data owners to select from a list of possible policies, specify values for parameterized policies, or even design their own policies from scratch.

3 Demonstration

We have implemented a prototype of PRIVFRAMEWORK that is capable of demonstrating all of its main features. In this section we will describe the budgeting example discussed in Section 2 one of several such applications we have implemented.

We present an application where a Data Owner downloads their own financial data using the APIs that are generally leveraged by service providers. This data is then packaged with a specified policy, uploaded to a Data Manager running a PRIVFRAMEWORK agent, and encrypted at rest with their own key. Data Owners then authorize analysts to run analyses over their data. Analysts run these analyses and use the results to generate personal dashboards for the users.


In this process we display the key features of PRIVFRAMEWORK:

1. The Data Owner privately keeps possession of all secret tokens, passwords, and unencrypted data.
2. Data Owners may easily specify policies and authorize analysts.
3. Analysts submit existing python programs with minimal adjustments to run analyses.
4. Data Owners do not have access to analysis program details nor do Analysts have direct access to the target data.

4 Conclusion & Future Work

In this paper we have described PRIVFRAMEWORK, a new system for user-configurable and automated privacy regulation compliance. PRIVFRAMEWORK has the potential to give data owners vastly

Figure 2: Choosing a policy in the client application is simple for the user and can be customized by the developer.



increased control over their data while also easing organizations' compliance with local data privacy regulations. We implement PRIVFRAMEWORK as well as a number of proof of concept applications on top of it, including a personal budgeting system. In the future we hope to work on further use cases leveraging PRIVFRAMEWORK and to extend the list of supported back-end analysis libraries. Future work may also introduce additional advanced privacy features to the policy language such as specialized differential privacy implementations.

References

- [1] California consumer privacy act of 2018. Online at https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB1121, 2018.
- [2] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016. URL <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>.
- [3] Amazon Web Services. Online at <https://aws.amazon.com/maciek/> Accessed: 2020-10-05.
- [4] Centers for Medicare & Medicaid Services. The Health Insurance Portability and Accountability Act of 1996 (HIPAA). Online at <http://www.cms.hhs.gov/hipaa/>, 1996.
- [5] O. Chowdhury, A. Gampe, J. Niu, J. von Ronne, J. Bennatt, A. Datta, L. Jia, and W. H. Winsborough. Privacy promises that can be kept: a policy analysis method with application to the hipaa privacy rule. In *Proceedings of the 18th ACM symposium on Access control models and technologies*, pages 3–14, 2013.
- [6] European Commission. 2018 reform of eu data protection rules. Online at <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN>, 2016.
- [7] Google Cloud Platform. Online at <https://cloud.google.com/dlp> Accessed: 2020-10-05.
- [8] S. Gressin. The equifax data breach: What to do. *Federal Trade Commission*, 8, 2017.
- [9] D. Hardt. The OAuth 2.0 Authorization Framework. RFC 6749, Oct. 2012. URL <https://rfc-editor.org/rfc/rfc6749.txt>.
- [10] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. Fernández del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. *Nature*, 585:357–362, 2020. doi: 10.1038/s41586-020-2649-2.
- [11] J. Isaak and M. J. Hanna. User data privacy: Facebook, cambridge analytica, and privacy protection. *Computer*, 51(8):56–59, Aug. 2018. ISSN 0018-9162. doi: 10.1109/MC.2018.3191268. URL <https://doi.org/10.1109/MC.2018.3191268>
- [12] D. Kaplan, J. Powell, and T. Woller. Amd memory encryption. URL https://developer.amd.com/wordpress/media/2013/12/AMD_Memory_Encryption_Whitepaper_v7-Public.pdf.
- [13] Oasis Labs. URL <https://www.oasislabs.com/parcelsdk>
- [14] T. pandas development team. pandas-dev/pandas: Pandas, Feb. 2020. URL <https://doi.org/10.5281/zenodo.3509134>
- [15] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Álché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [16] Plaid. URL <https://plaid.com/>

- [17] S. Sen, S. Guha, A. Datta, S. K. Rajamani, J. Tsai, and J. M. Wing. Bootstrapping privacy compliance in big data systems. In *2014 IEEE Symposium on Security and Privacy*, pages 327–342. IEEE, 2014.
- [18] L. Wang, J. P. Near, N. Somani, P. Gao, A. Low, D. Dao, and D. Song. Data capsule: A new paradigm for automatic compliance with data privacy regulations. In V. Gadepally, T. Mattson, M. Stonebraker, F. Wang, G. Luo, Y. Laing, and A. Dubovitskaya, editors, *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*, pages 3–23, Cham, 2019. Springer International Publishing. ISBN 978-3-030-33752-0.
- [19] Z. Whittaker. URL www.zdnet.com/article/adultfriendfinder-network-hack-exposes-secrets-of-412-million-users/
- [20] Yodlee. URL <https://yodlee.com/>.