



**HAL**  
open science

## rgpdOS: GDPR Enforcement By The Operating System

Alain Tchana, Raphael Colin, Adrien Le Berre, Vincent Berger, Benoît Combemale, Natacha Crooks, Ludovic Pailler

### ► To cite this version:

Alain Tchana, Raphael Colin, Adrien Le Berre, Vincent Berger, Benoît Combemale, et al.. rgpdOS: GDPR Enforcement By The Operating System. DSN 2023 - 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks, Jun 2023, Porto, Portugal. pp.100-104, 10.1109/DSN-S58398.2023.00032 . hal-04425768

**HAL Id: hal-04425768**

**<https://inria.hal.science/hal-04425768>**

Submitted on 4 Apr 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# rgpdOS: GDPR Enforcement By The Operating System

Alain Tchana  
ENS Lyon  
Lyon, France  
alain.tchana@ens-lyon.fr

Raphael Colin  
ENS Lyon  
Lyon, France  
raphael.colin@ens-lyon.fr

Adrien Le Berre  
ENS Lyon  
Lyon, France  
adrien.leberre@ens-lyon.fr

Vincent Berger  
University of Lyon 3  
Lyon, France  
vincent.berger@univ-lyon3.fr

Benoit Combemale  
University of Rennes  
Rennes, France  
benoit.combemale@irisa.fr

Natacha Crooks  
Berkeley  
California, USA  
ncrooks@berkeley.edu

Ludovic Pailler  
University of Lyon 3  
Lyon, France  
ludovic.pailler@univ-lyon3.fr

## ABSTRACT

The General Data Protection Regulation (GDPR) forces IT companies to comply with a number of principles when dealing with European citizens' personal data. Non-compliant companies are exposed to penalties which may represent up to 4% of their turnover. Currently, it is very hard for companies driven by personal data to make their applications GDPR-compliant, especially if those applications were developed before the GDPR was established. We present rgpdOS, a GDPR-aware operating system that aims to bring GDPR-compliance to every application, while requiring minimal changes to application code.

## CCS CONCEPTS

- Software and its engineering → Process management; File systems management.

## KEYWORDS

GDPR, Privacy, File System, Data Base, Operating System

## 1 INTRODUCTION

The General Data Protection Regulation (GDPR), introduced in May 2018, forces data operators to comply with a number of principles made to enforce personal data (PD) protection for European citizens (hereafter the subjects). Data operators failing to comply with the regulation can be fined up to 4% of their global turnover. For example, in March of 2022, the Irish Data Protection Commission sentenced Meta to a fine of €17M [5].




Figure 1: Relying on [2]: (left) total amount of penalties; (right) top 5 most sanctioned business sectors.




Figure 2: GDPR at the DB engine level in userspace.

As summarized in Fig. 1 left, the amount of penalties imposed to companies increases every year, topping 1.2 billion euros in 2021. Also, as shown in Fig. 1 right, companies of all types are impacted. For instance, in 2020 the CNIL in France penalized two doctors (€9K) for hosting medical images on a server which was freely accessible on the Internet [1].

However, bringing computer science and law together is not an easy task [13, 14]. The law mainly targets the new opportunities and dangers that information technologies create. Therefore, it is necessary to rethink applications for GDPR compliance, while maintaining usability, performance and velocity. Very few studies




Figure 3: Our vision

investigated GDPR compliance at application design time. Among existing works, some [4, 6] suggested guidelines, mainly in terms of documents which provide best practices. However, these efforts are lacking a true operational form. To our knowledge, Shastri et al. [17], Schwarzkopf et al. [16] and Istvan et al. [11] are the only operational works. The two former take into account the GDPR within the application’s data base (DB) engine in userspace, thus relying on a general purpose OS, as shown in Fig. 2. This approach has two main limitations: first, it is application-specific; second, the OS’s filesystem upon which the DB engine is running can take actions that contradict the latter, thus potentially violating GDPR compliance. For illustration, the filesystem’s logging mechanism can compromise the GDPR’s *right to be forgotten* as data deleted by the DB engine can still be present in the filesystem’s logs. Istvan et al. [11] studied how existing storage hardware technologies can be leveraged by applications to ensure the RGPD. Here, the developer deals with the RGPD, which is far from her expertises.

In this paper introduces rgpdOS, a fundamentally new approach to make applications GDPR compliant. rgpdOS mainly proposes a *GDPR-aware operating system (OS)* for two main reasons. First, by working at the OS level, any application can benefit from GDPR compliance ensured by rgpdOS. Second, we want to build an end-to-end solution that enforces GDPR at any software level. This is contrary to existing solutions which need to trust the huge code base of existing OSes (e.g., Linux has about 26M LOC). As a consequence, one of our main goal is to keep rgpdOS as *lightweight* as possible.

rgpdOS acts as a framework which forces the data operator to respect a number of *technical* rules, which in turn allows the OS to ensure GDPR compliance. This way, the data operator does not need to deal with the articles of the Law, but only technical rules, which are easier to comply with according to their background and more straightforward. According to the European Union Agency for Network and Information Security (ENISA)’s

definition, rgpdOS can be classified as a Privacy Enhancing Technology (PET) [3]. Using rgpdOS a data operator is demonstrating a conscious effort towards GDPR compliance like imposed by its 25th article. In addition, this certainly helps to reassure subjects about the respect of their rights.

We build rgpdOS following three new principles which go against current OS architecture, summarized in Figure 3. (Idea 1) In rgpdOS, we introduce the abstraction of *active data*, unlike current OSes in which data is *passive*. Active data is able to enforce the respect of consents defined by its subject. This is the role of the black membrane depicted in Fig. 3 right. (Idea 2) rgpdOS is *data-centric*, unlike current OSes which are *process-centric*. In the latter, the process brings data to its domain (virtual address space), as illustrated in Fig. 2 (see *pd1* and *pd2*). A function which should not access some PD (e.g., because the PD’s subject did not approve this usage) could still gain access to them (e.g., accidentally due to a use-after-free [9] vulnerability). Fig. 2 illustrates such a situation where function *f2* accidentally accesses *pd2*. rgpdOS reverses this power balance and runs the function in the *PD’s domain* as shown in Fig. 3 right. (Idea 3) In rgpdOS, we advocate for a structured database-oriented filesystem, unlike current OSes such as Linux which are file-based. The latter is coarse grained and not aware of the notion of *data* which can have different types. The GDPR requires systems to work at the granularity of individual PD pieces while traditional file systems see files, which are just bytes without particular meaning.

In our vision, the same server should still be able to process PD and NPD sequentially or at the same time (e.g., a single application may need to combine these two data types). Each data type should be processed using a distinct OS as we want to keep rgpdOS as lightweight and simple as possible, dedicated to PD processing. To achieve this goal, we introduce *purpose kernel*, a new kernel model that we follow to design the entire machine kernel.

To summarize, from the Systems research perspective, this project will investigate three novel research concepts: data centric, active data, database-oriented file system, and purpose kernel model. In this paper, we are laying the groundwork around these concepts, which raise several challenges. We do not list all of the latter in this paper as we are still discovering them.

The rest of this paper is organized as follows: § 2 presents the design of rgpdOS and the main implementation choices are presented in § 3, § 4 shows how rgpdOS can effectively enforce some example GDPR rules, finally § 5 presents the conclusion.

## 2 DESIGN

The main guiding principle of rgpdOS is as follows: *every access to PD must be controlled by rgpdOS*. rgpdOS therefore intervenes during the entire life cycle of each piece of PD: from collection to processing and destruction. Fig. 4 left shows the architecture of rgpdOS, described below.

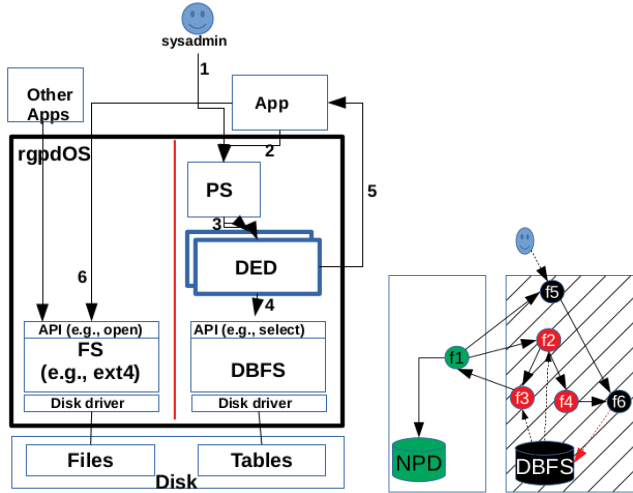


Figure 4: (left) Architecture of rgpdOS, followed by (right) its programming model and execution.

**Purpose kernel model.** In this model, the kernel is the aggregation of several sub-kernels where each sub-kernel achieves a specific purpose. The sub-kernels can be organized into three categories:

- IO driver kernels: every IO device is managed by a dedicated kernel which is mainly composed of the device driver. They are lightweight kernels.
- a general purpose kernel: it hosts and processes NPD. In addition, it does not include IO drivers.
- rgpdOS, a PD GDPR-aware kernel: it hosts and processes PD.

The different kernels cooperate to (dynamically) partition CPU and memory resources. We remove the IO devices under the management of the general purpose kernel because they are traversed by PD. Therefore, rgpdOS and IO driver kernels are the systems that we plan to prove in order to guarantee end-to-end GDPR compliance.

**File System.** rgpdOS contains at least two filesystems: the first one named DBFS is based on a structured DB and is used to store PD. Every PD has a precise *type* which corresponds to a table in the kernel’s DB. In rgpdOS, data types must be created in DBFS prior to use by a data processing application. PD structuration (typing)

is important for several reasons. First, it allows to apply the GDPR on different *views* of the same PD (see below). Second, it enables subjects to obtain their PD from the data operator in a structured and machine-readable format like the GDPR prescribes (see § 4). Using a DB as a filesystem is appropriate in this case, especially since DB engines have seen significant improvement over the last years[7].

The second filesystem in rgpdOS is used to store non-personal data. It can be implemented with a traditional filesystem (e.g., *ext4*) which works at the file granularity.

DBFS can only be accessed through the components of rgpdOS. In other words, DBFS is not visible from the outside and every direct access attempt from the outside is blocked by using a security mechanism (e.g., Linux Security Module). As for the second filesystem, it is accessible by any process running on the machine.

**PD membrane.** This is the first demonstration of the notion of *active data*. Each PD stored in DBFS includes a *membrane*. The configuration of the membrane and PD types is handled by the data operator (hereafter sysadmin). The membrane features different categories of metadata, among the most important ones are: the origin of the PD (it could be the subject, the sysadmin, or another data operator); consents relative to each data processing operation; time to live; level of sensibility; the interface to use for data collection (e.g., a web page to gather data directly from the subject) when the data is not present in DBFS; etc. Consent indicates the permissions of each operation on PD. A data type *view* corresponds to a specific representation or fragment of the data type. This notion enables our system to implement the GDPR’s *data minimization* principle. About the sensibility level, it is useful since the GDPR prescribes that sensitive data (e.g., a social security number) be stored separately from less sensitive data (e.g. a name). The time to live is directly requested by the GDPR and can be used to implement the *right to be forgotten*. Finally, the data collection interface will be used by rgpdOS to initialize DBFS with data collected from the appropriate data sources.

In order to illustrate how rgpdOS works, listing 1 shows a simplified example of a data type declaration (*user*), as well as its default membrane. This example also defines two views *v\_name* and *v\_anon*. The former makes only the name of a user visible while the latter view only allows access to the user’s age. The *consent* keyword indicates the default consent to apply when data of this type is created (collected). The GDPR allows PD processing when the latter relies on a legitimate basis, for example, the user’s consent.

The default consent allows the data operator to express operations that are backed by a legitimate basis, and thus do not need the specific subject's consent. In our example, operation *purpose1* can access the entirety of the *user* data type, *purpose3* can only manipulate the age, and *purpose2* is not allowed to process this data type at all.

**Listing 1: Example of a personal data definition type**

```

1 type user {
2   fields {
3     name: string,
4     pwd: string,
5     year_of_birthdate: int
6   };
7   view v_name {
8     name
9   };
10  view v_ano {
11    age
12  };
13  consent {
14    purpose1: all,
15    purpose2: none,
16    purpose3: ano
17  };
18  collection {
19    web_form: user_form.html,
20    third_party: fetch_data.py
21  };
22  origin: subject;
23  age: 1Y;
24  sensitivity: high;
25 }

```

**Programming model.** rgpdOS dictates that applications willing to process PD be programmed as a workflow of functions, as shown in Fig. 4 right. Functions can be organized into two categories: ( $F_{npd}$ ) functions that do not access PD (f1 in Fig. 4 right) and ( $F_{pd}$ ) functions that need to access PD (f2-f6 in Fig. 4 right).  $F_{npd}$  can be implemented with a single function, as shown in Fig. 4 right.  $F_{npd}$  is split into two subcategories: ( $F_{pd}^w$ ) functions that modify the state of DBFS (f5 and f6 in Fig. 4 right) and ( $F_{pd}^r$ ) functions that do not (f2-f4 in Fig. 4 right).  $F_{pd}^w$  functions are natively provided by rgpdOS (they are *built-in*) whereas  $F_{pd}^r$  functions are developed by the data operator. Built-in functions (described below) ensure that every PD is correctly wrapped, that is it always includes a membrane. Among built-in functions, we can list *update*, *delete*, *copy* and *acquisition* (for data collection).  $F_{pd}^r$  functions are forbidden to make syscalls that could leak PD (e.g., *write*).

Every  $F_{pd}$  function is the implementation of a unique data processing *purpose*. Functions can be implemented

in any programming language (e.g., C). As for the purpose declaration, we introduce a new very high level language as purposes should probably be written by the project manager while the actual implementation is written by developers. In rgpdOS's jargon, we call *data processing* a pair composed of one *purpose* and its *implementation*. Listing 2 shows an example where the implementation is done in C and the purpose is denoted by a C comment. In this example, this processing's goal is to compute the age of the input user. Since subjects can restrict each purpose to a limited set of PD's views, the implementation must be aware that some fields of the data type might not be available, as illustrated in line 4 of listing 2. In the latter, the computation is only possible for users for which the subject authorized access to the *age* field for *purpose3*.

**Listing 2: Example of data processing code**

```

1 #include "/etc/rgpdos/ps/types.h"
2 /* purpose3 */
3 struct age_pd compute_age(struct user_pd user) {
4   if (user.age) { //is age allowed to be seen?
5     return current_year() - user.year_of_birthdate;
6   }
7   else {
8     // error
9   }
10 }

```

**Listing 3: Example of main application**

```

1 #include "/etc/rgpdos/ps/ps.h"
2 int main() {
3   int age = ps_invoke(modpol, ref, "compute_age",
4     web_form, 0);

```

Any  $F_{pd}$  function takes as input the identifier of a PD or a PD type, which will be used latter by rgpdOS to automatically generate a request to extract data from DBFS. Finally, when a  $F_{pd}$  function wants to return some PD to the calling application, rgpdOS instead returns a reference or ID. Subsequently, the main application (app in Fig. 4 left, f1 in Fig. 4 right, or *main* in listing 3) never manipulates real PD within its address space.

**rgpdOS built-in functions.** In order to motivate rgpdOS's built-in functions, let us consider three examples:

- **copy:** rgpdOS must ensure membrane consistency across all copies of the same PD.
- **collection:** rgpdOS must ensure privacy and traceability (by keeping track of each PD's origin) from the moment PD enters the system. To this end, rgpdOS requests the needed metadata to fill the membrane with at data collection time. Thanks to this requirement, each entry in DBFS

is always correctly wrapped with its membrane. However, rgpdOS leaves the configuration of the collection interface (e.g., web form) to the data operator.

- *delete*: rgpdOS must ensure the GDPR’s *right to be forgotten*.

**Processing Store(PS)**. This is the second component of rgpdOS. It is the only rgpdOS entry point. Its public interface consists of two functions: `ps_register` and `ps_invoke`. Every  $F_{pd}$  function must be registered (`ps_register`) first in PS before they can be invoked. On call to `ps_register`, PS makes the following checks: if the function has no specified purpose, it is rejected; if the specified purpose does not “match” with the corresponding implementation, PS raises an alert that requires an explicit sysadmin approval<sup>1</sup>. PS is also the entry point for the invocation of  $F_{pd}$  functions, through the `ps_invoke` call, as shown in Fig. 4 left and illustrated in listing 3. `ps_invoke` takes as input the reference of a data processing operation, optionally a reference to PD, a data collection method and a boolean indicating whether or not the data collection function is to be called to initialize DBFS.

**Data Execution Domain (DED)**. This is the third and final component of rgpdOS. Any  $F_{pd}$  function is always executed as an instance of the DED, an environment that ensures GDPR compliance on manipulated PD. When PS receives a `ps_invoke` call, it instantiates a DED, which in turn follows a number of steps. (`ded_type2req`) the DED translates the processing’s input parameter type to requests at the destination of DBFS; (`ded_load_membrane`) the first request actually submitted to DBFS is fetching the membrane of the involved PD; (`ded_filter`) the DED filters the queried membranes in order to only keep references to PD which approve the current processing; (`ded_load_data`) the second request issued to DBFS fetches the actual data for PD which passed the previous filter; (`ded_execute`) the processing is executed on the fetched data; (`ded_build_membrane`) if any PD is generated by the processing, it is assigned a membrane; (`ded_store`) if any PD is generated by the processing, it is stored in DBFS with its newly computed membrane; (`ded_return`) the DED returns to the caller non PD as well as references to PD.

**Enforcement**. In order to guarantee any PD is only accessed under the supervision of rgpdOS, the latter enforces a number of restrictions that we summarize

<sup>1</sup>Let’s recall that the goal of rgpdOS is to guide the data operator in the path of GDPR compliance.

here: (1) PS is the only component that is able to access stored processings; (2) PS is the only entry point to invoke a processing; (3) every PD stored in DBFS must have a membrane attached to it; (4) DED is the only component that is able to access DBFS directly.

### 3 IMPLEMENTATION (IN PROGRESS)

This section presents the set of design choices that we make to implement rgpdOS.

(1) Besides the intrinsic challenges raised by the new paradigms that rgpdOS introduces, we impose on ourselves another challenge which is the implementation of rgpdOS by revisiting Linux. We do this to give rgpdOS a chance to be easily adopted. We are currently implementing rgpdOS following the semi-microkernel model [10, 12]. The latter runs some Linux kernel services as userspace processes for several reasons such as quick customization. In rgpdOS, DBFS, PS and DED run outside the Linux kernel while the second filesystem could be Linux’s native filesystem. To provide DBFS, we are rearchitecting uFS [12] in order to implement a database oriented filesystem. Notice that uFS is file based. The only part of uFS that we keep is the implementation of the inode concept. Starting from this point, PD in DBFS is represented by two major inode trees. The first tree gathers every PD from all subjects, with a separate set of inodes for each of them, grouping not only their personal data but also the membrane. Moreover, the second major tree provides the database structure, with a core inode or subset of inodes for each table describing the structure of the contained data, the different fields of the table, and a list of subject’s inodes, providing an easy link to quickly fetch the corresponding pieces of information. Finally, a dedicated set of inodes describes the general structure of the data encoded in the inode subtree of each subject: meant to be accessed only once by the filesystem during a given live session. This allows the system to properly format the pieces of data when returning them to DED.

(2) Concerning security enforcement, we rely on the Linux Security Module (LSM) framework. After the analysis of the major LSM implementations, we observed that SELinux and Smack can do the job. We leverage Linux Seccomp BPF to avoid functions which operate on PD to perform syscalls that can leak data.

(3) Different techniques can be used to ensure DED protection including TEEs like Intel SGX [8]. Furthermore, DED could be executed in multiple locations with the help of *Processing in Memory* [15] (e.g. UPMEM) and *Processing in Storage*.

(4) Finally, checking if a processing’s implementation matches its purpose is a challenging problem which is not yet addressed in rgpdOS. We plan to investigate approaches borrowed from several research domains such as Semantic and AI.

## 4 ILLUSTRATION

This section shows how rgpdOS can ensure compliance with two example GDPR rules: the *right of access* and the *right to be forgotten*.

**Right of access.** The way that data operators are asked by the GDPR to send PD and processings to the subject is unclear. As a matter of fact, the GDPR specifies that PD must be “structured and machine-readable”. Let’s say for example that a company stores the first and last names of data subjects with key value pairs (`first_name: "Chiraz", last_name: "Benamor"`). This pairing is structured, and is therefore machine readable, but what makes it exploitable is the fact that the keys make sense. According to the GDPR, nothing is preventing this company from sending this same data in a structured, but inadequate way (`Chiraz: "Benamor"`). Leveraging rgpdOS, which enforces PD structuration, an official authority could simply ask that data be sent to subjects as it is stored in DBFS. As for informing subjects about processings executed on their PD, this is easily obtained thanks to the DED, which logs every executed processing. This log is organized so that it can give information about executed processings for each piece of PD.

**Right to be forgotten.** This right is implemented as a built-in function in rgpdOS. When invoked by a subject, it doesn’t necessarily mean that her PD will be completely and definitively erased. In fact, it could be necessary to keep this data, for legal investigations for example. In order to implement this, rgpdOS assumes a model in which each data operator owns a public encryption key given to them by the authorities who keep the private key. When PD is to be deleted, rgpdOS will simply encrypt it using the public key, in this way, the data operator will not be able to access the data anymore, but the authorities will be able to decrypt it using their private key.

## 5 CONCLUSION

We presented rgpdOS, a GDPR-aware operating system, allowing every applications to be GDPR-compliant. To this end, rgpdOS introduces the concept of *active data*. Moreover, unlike current operating systems which are process-centric, rgpdOS focuses on data. We

are currently implementing rgpdOS following the semi-microkernel model, with Linux as a basis.

## ACKNOWLEDGMENTS

This project is funded by several french research programs: “Pack Ambition International” - Région Auvergne-Rhône-Alpes 2021, MSH Lyon St-Etienne “Projets thématique” 2021 and Amorçage Europe 2022.

## REFERENCES

- [1] [n.d.]. *CNIL*. <https://www.cnil.fr/en/node/120684>
- [2] [n.d.]. *Data Legal Drive*. <https://datalegaldrive.com/lentreprise/nos-outils/carte-sanctions-rgpd/>
- [3] [n.d.]. *Privacy Enhancing Technologies*. <https://www.enisa.europa.eu/topics/data-protection/privacy-enhancing-technologies>
- [4] 2019. *A Guide to Privacy by Design*. [https://www.aepd.es/sites/default/files/2019-12/guia-privacidad-desde-diseno\\_en.pdf](https://www.aepd.es/sites/default/files/2019-12/guia-privacidad-desde-diseno_en.pdf)
- [5] 2022. *Data Protection Commission announces decision in Meta (Facebook) inquiry*. <https://www.dataprotection.ie/en/news-media/press-releases/data-protection-commission-announces-decision-meta-facebook-inquiry>
- [6] 2022. *The Practical Guide to Privacy by Design Architecture*. <https://piiiano.com/blog/practical-guide-to-privacy-by-design-architecture/>
- [7] Michael Cafarella, David DeWitt, Vijay Gadepally, Jeremy Kepner, Christos Kozyrakis, Tim Kraska, Michael Stonebraker, and Matei Zaharia. 2020. DBOS: A Proposal for a Data-Centric Operating System. <https://doi.org/10.48550/ARXIV.2007.11112>
- [8] Victor Costan and Srinivas Devadas. 2016. Intel SGX Explained. Cryptology ePrint Archive, Report 2016/086. <https://ia.cr/2016/086>.
- [9] Márton Erdős, Sam Ainsworth, and Timothy M. Jones. 2022. MineSweeper: A “Clean Sweep” for Drop-in Use-after-Free Prevention. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (*ASPLOS 2022*). Association for Computing Machinery, New York, NY, USA, 212–225. <https://doi.org/10.1145/3503222.3507712>
- [10] Jack Tigar Humphries, Neel Natu, Ashwin Chaugule, Ofir Weisse, Barret Rhoden, Josh Don, Luigi Rizzo, Oleg Rombakh, Paul Turner, and Christos Kozyrakis. 2021. GhOST: Fast & Flexible User-Space Delegation of Linux Scheduling. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles* (Virtual Event, Germany) (*SOSP ’21*). Association for Computing Machinery, New York, NY, USA, 588–604. <https://doi.org/10.1145/3477132.3483542>
- [11] Zsolt István, Soujanya Ponnappalli, and Vijay Chidambaram. 2020. Towards Software-Defined Data Protection: GDPR Compliance at the Storage Layer is Within Reach. *CoRR* abs/2008.04936 (2020). arXiv:2008.04936 <https://arxiv.org/abs/2008.04936>

- [12] Jing Liu, Anthony Rebello, Yifan Dai, Chenhao Ye, Sudarsun Kannan, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. 2021. Scale and Performance in a Filesystem Semi-Microkernel. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles* (Virtual Event, Germany) (*SOSP '21*). Association for Computing Machinery, New York, NY, USA, 819–835. <https://doi.org/10.1145/3477132.3483581>
- [13] Denis Merigoux, Nicolas Chataing, and Jonathan Protzenko. 2021. Catala: A Programming Language for the Law. *Proceedings of the ACM on Programming Languages* 5, ICFP (Aug. 2021), 77:1–29. <https://doi.org/10.1145/3473582>
- [14] Denis Merigoux, Raphaël Monat, and Christophe Gaie. 2020. Étude formelle de l'implémentation du code des impôts. In *JFLA 2020 - 31ème Journées Francophones des Langages Applicatifs*. Gruissan, France. <https://hal.inria.fr/hal-02320347>
- [15] Joel Nider, Craig Mustard, Andrada Zoltan, John Ramsden, Larry Liu, Jacob Grossbard, Mohammad Dashti, Romaric Jodin, Alexandre Ghiti, Jordi Chauzi, and Alexandra Fedorova. 2021. A Case Study of Processing-in-Memory in off-the-Shelf Systems. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, 117–130. <https://www.usenix.org/conference/atc21/presentation/nider>
- [16] Malte Schwarzkopf, Eddie Kohler, M. Kaashoek, and Robert Morris. 2019. *Position: GDPR Compliance by Construction*. 39–53. [https://doi.org/10.1007/978-3-030-33752-0\\_3](https://doi.org/10.1007/978-3-030-33752-0_3)
- [17] Supreeth Shastri, Vinay Banakar, Melissa Wasserman, Arun Kumar, and Vijay Chidambaram. 2020. Understanding and benchmarking the impact of GDPR on database systems. *Proceedings of the VLDB Endowment* 13, 7 (mar 2020), 1064–1077. <https://doi.org/10.14778/3384345.3384354>