




# White Paper

## **Balancing Privacy: Leveraging Privacy Budgets for Privacy-Enhancing Technologies**

By Mariya Georgieva Belorgey, PhD, VP of Cryptography, Inpher  
and Kevin Deforth, Senior Software Engineer, Inpher



# Introduction

---

In the past few years, the support for emerging privacy-enhancing technologies (PETs) that allow for data sharing and analytics while preserving privacy, has grown considerably.

A critical (and sometimes overlooked) aspect in ensuring robust privacy preservation is to account for the privacy budget allocated in any given PETs project. The privacy budget refers to the finite amount of privacy protection that can be allocated when performing various computations or data analysis using PETs.

In this work we give a comprehensive overview to privacy budget allocation for PETs by reviewing different privacy metrics and a number of relevant policy and governance guidelines that refer to best practices related to the privacy budget and show how the concept of a privacy budget can help to demonstrate compliance with the requirements of privacy by design.

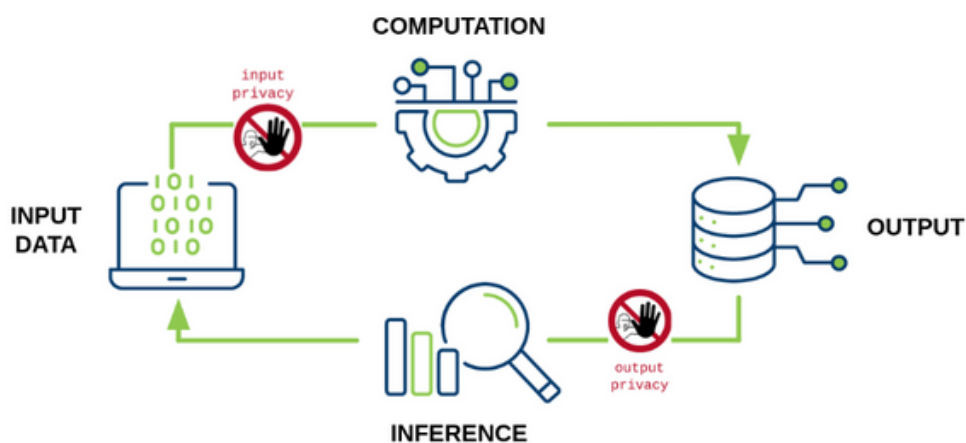
Secure multiparty computation (MPC), fully homomorphic encryption (FHE), federated learning (FL), trusted execution environments (TEEs) and differential privacy (DP) are prominent examples of emerging PETs. They enable the computation of a function without revealing the input data. By incorporating these techniques, organizations strike a balance between preserving the privacy of sensitive input data and deriving valuable insights from data analysis, optimizing the privacy-utility tradeoff. To achieve this, the protocols employ cryptographic techniques and algorithms.

Two important privacy goals for PETs are input and output privacy. On the one hand, input privacy allows forward computation from input data without disclosing it, while on the other hand, output privacy prevents backward inference from disclosed output results as shown in the Figure below.

PETs need to be applied in conjunction with Privacy Budget otherwise the output privacy can be compromised, even when the input data is protected with strong cryptographic guarantees.

In this white paper, our main focus is on output privacy problems, delving into strategies and techniques to enhance the protection of sensitive information in the results of data analysis. In this manner, the white paper aims to contribute to the effective communication and delivery of the privacy guarantees provided by PETs in light of the Fundamental Law of Information Recovery (see below) and increasing applications of PETs.

"Input privacy means that the Computing Party cannot access or derive any input value provided by Input Parties, nor access intermediate values or statistical results during processing of the data (unless the value has been specifically selected for disclosure)." - p.15 [UN Handbook for Privacy-Preserving Techniques](#)



"A privacy-preserving statistical analysis system implements output privacy to the extent it can guarantee that the published results do not contain identifiable input data beyond what is allowable by Input Parties." - p.15 [UN Handbook for Privacy-Preserving Techniques](#)

# The Concept of Output Privacy and Privacy Budget

---

The challenge of safeguarding sensitive information in a database while permitting statistical queries has been extensively [researched](#) for a long time. In the [Database Reconstruction Theorem](#) it was shown that too many statistics published too accurately from a confidential database exposes the entire database with near certainty.

Today, this phenomenon is coined as the [Fundamental Law of Information Recovery](#), saying that “overly accurate answers to too many questions will destroy privacy in a spectacular way”. The Fundamental Law of Information Recovery is a general rule which holds true for any technique employed to limit data disclosure.

This means that when using PET techniques, privacy and confidentiality of the input data are not automatically maintained in the output. Instead, with the repeated queries on the same data, [privacy risk](#) in the form of the [possibility to recover the information](#) used for the query is growing.

As an answer to this, the concept of a Privacy Budget was developed. The purpose of a Privacy Budget is to define a maximum tolerance for revealing information about each user and keep the total amount of revealed information within acceptable bounds (the “budget”).

Originally introduced in the [context of DP](#), this article expands the notion of a Privacy Budget to encompass various computational techniques. For example, PET protocols may involve the exchange of information in the course of computations that reveal partial data. When combined with the final result, this would further increase the privacy loss of the input data.

To communicate the privacy loss in all of these contexts, we suggest defining and using the notion Privacy Budget as a vehicle of tracking the amount of information that is revealed across different computations.

The concept of a Privacy Budget offers valuable benefits for organizations and stakeholders beyond differential privacy. Depending on the technology used, it can serve to quantify privacy loss in various contexts, promoting transparency and supporting reliable long-term privacy preservation. Broadening the concept of the Privacy Budget beyond DP also aligns with the more inclusive use of the term in practice. Instead of Privacy Budget, the term Privacy Metric is also common.



# The Origins of Privacy Budget in the Context of Differential Privacy

---

Differential privacy strives to balance privacy protection and meaningful data analysis by introducing small amounts of controlled randomness to a dataset, a model, or an output. This approach guarantees that models or outputs of data analytics remain indistinguishable, regardless of the inclusion or exclusion of any single data-point. However, every query on the underlying private data results in some amount of information being revealed and an increasing privacy loss. Given enough computations or queries on the same data, an attacker might be able to [learn about the input data over time](#).

To calculate and limit this privacy loss that accumulates over multiple computations, the concept of a privacy budget was established. The privacy budget defines a limit on the amount of information that may be revealed to a specific algorithm or analysis. The amount of revealed information can be estimated by measuring the so-called [sensitivity of a function](#). The sensitivity of a query is used to calculate the remaining privacy budget after execution of the query.

*Sensitivity quantifies the information a single query / function reveals about the underlying data. The **privacy budget** indicates how much information is allowed to be revealed cumulatively across all queries.*

For a detailed example of how to use **sensitivity** see Section **Examples of Privacy Metrics for Privacy Budget** on page 10 of this document. In operations or analysis involving sensitive data, such as counting, averaging, or aggregating information, each query reduces the allocated privacy budget by a certain amount. The privacy budget in differential privacy is typically represented as a parameter denoted as  $\epsilon$  (epsilon).  $\epsilon$  bounds the effect of an individual data point on the output of an analysis. A smaller value of  $\epsilon$  indicates a stricter privacy guarantee.

# Privacy Budget in the Context of PETs and MPC

---

The privacy budget provides significant advantages for organizations in managing privacy loss within acceptable limits. In this understanding, the concept of a privacy budget extends beyond differential privacy, and can be meaningfully applied to the broader range of PETs.

In the context of MPC, where multiple parties jointly compute a result without revealing their individual data to each other, it can be challenging to estimate the amount of revealed information by party and datasource. This is particularly the case when the data sources or the requests may be correlated and when the parties can collude. All the more important and useful is applying the concept of privacy loss and privacy budget to this complex setting to make sure that a threshold on the amount of information that is explicitly revealed during the joint computation is set.

This can be achieved by the following steps.

Initially, a privacy budget is defined per datasource and per party. Before running a computation, the privacy budget of each input datasource and party is checked in advance. After a computation is executed, the consumed privacy budgets are subtracted from the respective privacy allowances with permanent effect on the privacy budget.

Depending on the specific context of a joint computation, and taking into account intermediate results, meta-data and the final output, the defined privacy budget per datasource and per party may be partially or fully consumed. If the remaining budgets are sufficient for a specific operation, the operation is permitted; otherwise it is blocked.

In this scenario, the privacy budget acts as a form of access control, allowing data owners to limit the amount of revealed information on their datasources.

One possible approach to achieve this is to use a Privacy Budget Table (Table 1). In the beginning the table is initialized with the maximal privacy budget for all participants. The rows contain the Privacy Budget of each data source for each participant. In our example those are the Input parties (IP1, IP2), the Compute parties (CP1, CP2, CP3) and the Result parties (RP1). The collusion model (*defines a group of parties working together to achieve a common goal*) is defined and the privacy loss for the collusion group is aggregated. For example, if the compute party (CP3) colludes (for example by sharing data or outputs) with the result party (RP1), the budget will be decreased for both of them. The table is updated after each computation.

| Input | Op ID | IP1  | IP2  | CP1  | CP2  | CP3  | RP1  | Executed  |
|-------|-------|------|------|------|------|------|------|---|
| D1    | init  | 100% | 100% | 100% | 100% | 100% | 100% |   |
| D2    | init  | 100% | 100% | 100% | 100% | 100% | 100% |   |
| D1    | comp1 |      |      |      |      | -20% | -20% |  |
| D2    | comp1 |      |      |      |      | -30% | -30% |  |

*Table 1: Privacy Budget Table, the first two rows contain the initial privacy budget set to 100% per data sources, the third and fourth rows show how the privacy budget for inputs D1 and D2 decreases by 20 and 30 percent respectively, after execution of operation “comp1”. The last column shows if the operation can be executed, which is the case if the remaining budget is positive. We first estimate the remaining budget and after that we execute it if the budget is still positive.*



Two important fields of active research are privacy budget composition and privacy budget in multiparty settings.

- Privacy budget composition refers to how privacy budgets interact when multiple privacy-preserving operations are performed consecutively. Understanding how privacy budgets combine or accumulate during a sequence of operations (that can be correlated or not) is an open problem.
- The concept of multiparty explores how privacy budgets interact when different parties collaborate and have different access to the revealed data. This deeply depends on the collusion model between the different parties.

For example, when using differential privacy in a multiparty setting, methods for an optimal distribution of a given privacy budget has been suggested in [Optimal Distribution of Privacy Budget in Differential Privacy](#).

# Examples of Privacy Metrics for Privacy Budget

---

Next to privacy budgets for specific PETs as shown above, there is a variety of general privacy metrics that can be used in different scenarios to estimate the amount of revealed information. In the following section, we list examples of approaches that can be applied for MPC, FHE, DP, or TEEs, depending on the concrete use case and context.


## **Number of queries/data points**

In the query restriction approach, special rules are imposed on queries to prevent information leakage. This approach is usually combined with query auditing, where a log of queries is kept, and new queries are checked for potential compromise. What follows is a very simplistic example of query restriction.

Suppose an online survey platform is used to collect personal information, such as age, gender, nationality, job and education level, from respondents. The platform wants to ensure the privacy of its users while still providing valuable insights to survey creators. To manage privacy, the platform sets a privacy budget that limits the amount of information that can be disclosed. In this hypothetical setting, let's assume that a privacy budget is set to three data points/queries per respondent (e.g age, gender and education grade or age, nationality and job level). By using a privacy budget, the online survey platform can strike a balance between providing valuable insights to survey creators and limiting the exposed information to make it insufficient to identify a user.

An example is given in Table 2 below.

| Input | Query               | Budget in Number Query (RP1) | Remaining Budget | Executed |
|-------|---------------------|------------------------------|------------------|----------|
| User1 | init                | 3                            | 3                |          |
| User1 | Age                 | -1                           | 2                | ✓        |
| User1 | job level           | -1                           | 1                | ✓        |
| User1 | Gender, nationality | -2                           | 1                | ✗        |



time




Table 2: Privacy Budget Table for Survey


A limitation of this approach is that the number and type of authorized queries must be tightly related to the size and the distribution of the input data. For example, querying the gender of an individual in a gender-balanced population might give less information to identify the individual than querying for age. But if the population is gender-imbalanced and age-balanced, the inverse would be the case. This was observed in numerous papers and ultimately led to a precursor of differential privacy c.f. [stanford paper](#).

## Number of bits

Another simple example of an approach that can be taken into account for a privacy budget is related to the number of bits that are revealed and the size of a datasource. In the beginning, the privacy budget for each fresh source (S) is initialized at e.g., 2% of the number of bits of the source:  $Privacy\ Budget(S) := 2\% \text{ Number of bits of } S$ . That is equivalent to say, if during the computation more than two percent of the number of bits of S are revealed, the computation is not allowed. As in the case of DP, the effects on the privacy budget must be tracked across different computations.

For example, if we have a file of 1k rows and 20 columns populated with coefficients of 8 bits, this may represent 20kB, and we may initialize the privacy budget to 2%, that means 3200 bits. Revealing the average of a single row or column may expose 8 bits of information. Revealing the *colsum* (sum of each column) may expose  $20(8 + \log_2(1000) = 360)$  bits, whereas for the *rowsum* this may expose around 12322 bits.

| Input | Op ID   | Budget in bits (RP1) | Remaining Budget | Executed  |
|-------|---------|----------------------|------------------|---|
| File1 | init    | 3200                 | 3200             |   |
| File1 | rowsum  | -12322               | 3200             |  |
| File1 | average | -8                   | 3192             |  |
| File1 | colsum  | -360                 | 2832             |  |



time

Table 3: Privacy Budget Table based on number of bits

This measure is easy to compute, but overly pessimistic, because it does not take into account the entropy of the output and the dependance on the input (which is related to the sensitivity of the function).

# Entropy

Entropy of an output also can be used as a metric of a privacy budget.

Imagine that we compute the function  $f(x) = 00000 \text{ lsb}(x)$ , whose output has 6 bits with 5 leading zeros and a least significant bit equal to the least significant bit of the input. Even if the size of the output is 6, the output contains only one bit of information (entropy).

Entropy quantifies the amount of uncertainty or randomness in a dataset and can be leveraged to estimate the privacy loss, but unfortunately is complicated to quantify.

## Sensitivity and Epsilon

The sensitivity of a function / query quantifies the information revealed about the underlying input data. Mathematically speaking, the sensitivity of a function indicates the maximal change of the output value caused by [small variations of the input values](#).

For example, asking for the average age of a population is more sensitive than asking if the average is less than 40 years.

Consider a database (Table 4) containing personal information of two individuals, each at least one, at most eighty years of age. If I change the age of one individual by  $x$  years, it will affect the output of my function by  $x/2$  years, or at most  $\log_2(80/4) = 5.32$  bits. But if I ask if the average age of a population is greater or equal to 40 years of age, it will affect the output of my function by at most 1 bit.

| Individual | Age in years | Income in USD | Nationality |
|------------|--------------|---------------|-------------|
| id_1       | 56           | 150K          | US          |
| id_2       | 44           | 60K           | DE          |

*Table 4: Database with personal information of two individuals; age in years (at least 1, at most 80 years), Income in USD (between 1k and 240k, rounded to the nearest 1k), Nationality (197 possibilities)*

In Differential Privacy, we add noise to the output to counteract the sensitivity of a function and we use the privacy metric epsilon to measure the privacy loss. It bounds the probability that a particular output can be preserved by adding or removing an input value. The [NIST blogpost](#) discusses suggestions for values of epsilon based on current experience, while cautioning that a deeper understanding for the impact of the value of epsilon on privacy is required.

## Size of preimage

The size of the preimage of the computed functions for a given output measures how many possible inputs we can obtain from the given output. This measure is precise, but not always easy to compute for arbitrary functions. It is proportional to the number of bits not exposed by revealing this output.

Picking up the example from the previous paragraph and Table 4.

The entire database offers  $80^2 * 240^2 * 197^2 \cong 14.3$  trillion different options (age is between 1 and 80 years, Income is between 1k and 240k USD, rounded to the nearest 1k and there are 197 possibilities for Nationality), or about 43.7 bits. The information “*the average age is greater or equal to 40.0 years*” corresponds to a pre-image of size  $3239 * 240^2 * 197^2 \cong 7.2$  trillion different options. This equates to around 42.7 bits, meaning this query exposes **1 bit** of information.

The information “*the average age is exactly 40.0 years*” corresponds to a pre-image of size  $79 * 240^2 * 197^2 \cong 177$  billion different options, which equates to around 37.36 bits, meaning this query exposes around **6.34** bits of information.

*Explanation:*

*79 possibilities lead to an exact average age of 40.0: (1, 79), (2, 78), ..., (79, 1)  
3239 possibilities lead to average age greater or equal to 40.0: ([1-80], 80), ([2-80], 79) + ([3- 80], 78), ..., ([79-80], 1) → 80 + 79 + 78 + .. + 2 (arithmetic series, leading to 3239 possibilities).*

The size of the preimage is also related to the parameter k in the context of k-anonymity, used to quantify if a user is partially identifiable. For example k-anonymity requires that each individual's information must be indistinguishable from at least k-1 other individuals in the dataset.

## Challenges and active research

Finding meaningful (and precise) privacy metrics to estimate the amount of revealed information which is easy to compute for arbitrary algorithms is an ongoing and exciting research area. A potentially meaningful contribution to the space would be to explain the relationships between the different metrics of privacy (sensitivity, entropy, pre-image size, output size and k-anonymity).

What follows is a non-exhaustive collection of references to active and existing efforts for the research of privacy metrics and creation of tools for handling privacy budgets.

Recently MIT researchers developed a new data privacy metric, [Probably Approximately Correct \(PAC\) Privacy](#) that exploits the uncertainty or entropy of the sensitive data in a meaningful way and builds a tool that automatically determines the minimal amount of noise to be added to the output to protect sensitive data. PAC Privacy allows a user to specify the desired level of confidence. The method is described in the [research paper](#) presented at Crypto'23.

The [Sdcmciro K-anonymity Library](#) includes anonymization methods to achieve k-anonymity as well as estimation of various risks.

In [Privacy Panel: Usable and Quantifiable Mobile Privacy](#) the authors propose privacy metrics to quantify the privacy impact of an app accessing user data. They focus on three user data categories and define privacy metrics for Location, Contacts and Content.

The NIST project [Collaborative Research Cycle](#) (CRC) has as a goal to accelerate research, innovation, and understanding of data de-identification techniques. It classifies and compares some selected de-identification algorithms. For the purposes of classification, two metrics are used: one to measure privacy leakage (UEM: Unique Exact Match), and one to measure the utility (SsE: Subsample Equivalent). UEM is a simple privacy metric that counts the percentage of uniquely identifiable records in the deidentified data. SeS is a utility metric that uses an analogy between deidentification error and sampling error to communicate utility.

Some useful tools are already proposed and some of them are open-source available. For example:

[Privacy Meter Tool](#) is an open-source Python library for auditing and quantifying the privacy risks of statistics and machine learning models. The tool provides privacy risk scores and identifies records with high risk of being leaked. It applies state-of-the-art inference attacks through systematic method and aids to audit a wide range of machine learning algorithms.

[Galois on Measuring Privacy of Computations with SCALE-MAMBA](#) proposes a tool that computes privacy leakage that combines static and dynamic leakage based on Quantitative Information Flow.

[Meta's IPA End to End Protocol](#) introduces the privacy budget in the context of web platforms for advertising attribution. DP is used to ensure that for a specific period of time (epoch) the amount of information revealed about an individual person is bounded, providing each website with a privacy budget.

In [Combining Fingerprinting with Privacy Budget](#), Chrome has proposed a privacy budget as a mitigation of fingerprinting (stable information about a given user's browser).

[Privacy Budget Scheduling paper from Columbia University and Microsoft Research \(2021\)](#) A scheduler for queries on private data, that aims to incentivize actors to collaborate and overall reduce the collectively used privacy budget.



# Privacy Budget as Best Practice in Privacy Risk Management

---

Privacy risk mitigation is a collaborative effort, involving multiple stakeholders and necessitates effective communication and shared understanding. The concept of a privacy budget is particularly valuable in this context as it emphasizes the need to control amounts of information disclosed across various computations and across teams to ensure privacy risks remain within acceptable limits.

A number of relevant policy and governance guidelines refer to best practices that correlate to the extended notion of the privacy budget as understood in this article.

The [UN Handbook on Privacy-Preserving Computation Techniques](#) highlights output privacy and considers it a privacy goal in privacy-preserving computation, in addition to input privacy and policy enforcement. According to the handbook, “output privacy addresses the problem of measuring and controlling the amount of leakage present in the result of a computation.” Output privacy focuses on measuring and controlling the amount of information leakage present in the outcome of a computation, e.g., the extent to which the original data can be inferred from a published model. Information leaked during the computation process itself is usually not taken into account by the output privacy, but is included in the input privacy. In contrast, the privacy budget is more generic because it also includes all intermediate revealed information, even if they are not part of the output.

The [IAB Data Clean Rooms Guidance and Recommended Practices](#) for AdTech lays out the requirements of data clean rooms as secure collaboration environments. The IAB guidance includes the principle of least privilege to maintain privacy by minimizing unnecessary data exposure and ensuring that only essential information is accessed and disclosed. The required mechanisms to enable this include limiting the number of queries allowed, restricting the types or complexity of queries, preventing reuse of data sets with other participants, and limiting outputs to only the necessary insights required for a task at hand. Here as well, the notion of a privacy budget can be a helpful tool to engage in compliance efforts that satisfies these requirements while allowing for more flexibility regarding the re-use of data-sets with other participants.

In another prominent example - NIST and the U.S. National Science and Technology Council published the [US National Strategy to Advance Privacy- Preserving Data Sharing and Analytics](#) -, which proposes to accelerate research to develop metrics and effective measurement techniques - quantitative or qualitative - for privacy risks, accuracy, and associated unintended consequences or harms. The Strategy emphasizes that some techniques such as differential privacy already use a privacy-loss parameter  $\epsilon$  (epsilon) to capture the privacy disclosure that is acceptable. It highlights that there remains an inadequate level of understanding of how privacy parameter values should be set for different applications. Configuring privacy-related parameters becomes especially challenging when different types of techniques are combined in a specific application, underscoring the need for better cross-collaboration. It is exactly this gap in practice that the expanded notion of privacy budget - taking into account all potential and factual privacy loss parameters to quantify privacy risks - can help to solve.

The guidelines mentioned above demonstrate the rising awareness of the need to quantify privacy risk and take into account accumulating data loss during computations to meet the [legal privacy obligations](#) around anonymization and de-identification of personal data and to prevent re-identification of individuals.

# Privacy Budgets in the Context of Privacy Law

---

In the context of the General Data Protection Regulation (GDPR), the risk of re-identification of personal data that has been anonymized is based on determining if a specific data processing workflow falls within the scope of the GDPR or not. The GDPR does not apply to the processing of data in cases where the data subject is not or no longer identifiable. (see [Recital 26 GDPR](#)). Here, the privacy budget can be used as a powerful tool to track the amount of information that gets revealed during a computation or is included in the output.

When the GDPR is applicable, the concept of a privacy budget can help to demonstrate compliance with the requirements of [privacy by design](#). As a concept that was developed in the 1990, privacy by design has been formalized in data protection laws globally, and requires organizations to put in place security safeguards and engineer privacy in all products, services and their IT infrastructure. In this context, the privacy budget can be used to better communicate the technical and organizational measures that have been taken to protect privacy under [Article 25 GDPR](#).

The same is true in the US. The Federal Trade Commission has recently warned in a [blog post](#) that firms making inaccurate claims about anonymization or secure data aggregation should be on guard that this can be a deceptive trade practice and violate the FTC Act. The FTC points out that significant research has shown that “anonymized” data can often be re-identified. Defining and monitoring an overall privacy budget can substantially help to manage the risk of insufficient anonymization of personal data.

*The FTC points out that significant research has shown that “anonymized” data can often be re-identified.*

At the same time, in the US as well, privacy by design is required by the FTC. In its 2012 [Protecting Consumer Privacy in an Era of Rapid Change report](#), it stated as a clear baseline principle that companies should promote consumer privacy throughout their organizations and at every stage of the development of their products and services.

Beyond personal data protection, allocating a privacy budget for computations and ML projects is a crucial strategy in preventing data loss in general. Working with sensitive intellectual property data calls for enhanced data loss prevention (DLP) practices. By integrating the concept of a privacy budget into the DLP ecosystem, organizations can strengthen asset management, monitoring, and access controls, leading to improved data breach prevention and risk mitigation associated with data loss.



# Conclusion

---

While developing techniques and tools that allow for fine-grained control over the privacy budget is an active and challenging domain, the concept of a privacy budget has many advantages.

A privacy budget is a unique concept that allows stakeholders to make informed decisions and to demonstrate a commitment to privacy protection and transparency. In helping to optimize the trade-off between privacy and utility, it enables flexible data analysis while supporting long-term information preservation. When working with personal data, the concept can be used as a measure of the additional privacy risk that an individual might face.

Understanding and effectively communicating the implications of the privacy loss, measured with the privacy budget, helps stakeholders correctly manage the privacy of PETs applications, gauge the level of information loss in the data analytics process and demonstrate compliance with requirements with privacy guarantees and requirements from different domains such as Adtech, Healthcare, Government, and Finance.