

Differential Privacy

Guidance for Digital Advertising

Please email support@iabtechlab.com for questions, public comments and feedback. This document is available online at <https://iabtechlab.com/diffprivacy>

Public comments close on December 14, 2023.

About this document

Many advertising technology workflows and necessary analysis for e.g. measurement and attribution require linking and sharing of multiple data sets both within and outside the organization that collected the data. Linking identifiable attributes can easily reveal the identity of individuals, for e.g. gender, age and zip code are uniquely sufficient to identify the vast majority of individuals in the USA. Different techniques and heuristics have been applied to protect identity of individuals and prevent identification of individuals by linking multiple data sets, for e.g. anonymization of identifiable attributes. But none of these provide a robust or rigorous guarantee of privacy.

Differential Privacy, a rigorous mathematical definition of privacy has emerged as a leading technique to analyze and draw inferences from data sets in a way where one cannot determine if a particular individual was present in the data or not. This document explores the application of Differential Privacy for ad tech use cases and provides guidance on:

- What is Differential Privacy
- How is it applied to ad tech use cases with deeper dive into attribution
- Privacy vs. utility considerations
- How is it different from other anonymization techniques
- What are the considerations when using Differential Privacy in advertising

This document is intended to be an informational guide for decision makers, analysts and product developers working with advertisers, publishers and ad tech providers to demystify the technology, scope and application of Differential Privacy.

This document is developed by the IAB Tech Lab [Privacy Enhancing Technologies \(PETs\) Working Group](#).

Note: The use of words or phrases “Privacy”, “Private”, “Security”, “Control”, “Processing”, “Personal Data”, “PII” in this document is generic and does not refer to definitions in any specific regulation e.g. GDPR or CCPA.

Throughout the document the word or phrases “ID”, “user ID”, “Consumer ID”, are used interchangeably referring to a unique identifier associated with a user of a service.

License

Differential Privacy Guidance document is licensed under a [Creative Commons Attribution 3.0 License](#). To view a copy of this license, visit creativecommons.org/licenses/by/3.0/ or write to Creative Commons, 171 Second Street, Suite 300, San Francisco, CA 94105, USA.

Significant Contributors

Andrei Lapets, *Magnite*; Andrew Knox, *Decentriq*; Brian May, *Dstillery*; Chris Watts, *NumberEight*; Jordan Cauley, *Mediavine*; Joshua Attardo, *Pea Pod Digital Labs*

IAB Tech Lab Lead

Shailley Singh, EVP Product & COO, IAB Tech Lab

Miguel Morales, Director Addressability & Privacy Enhancing Technologies (PETs)

About IAB Tech Lab

The IAB Technology Laboratory is a nonprofit research and development consortium charged with producing and helping companies implement global industry technical standards and solutions. The goal of the Tech Lab is to reduce friction associated with the digital advertising and marketing supply chain while contributing to the safe growth of an industry.

The IAB Tech Lab spearheads the development of technical standards, creates and maintains a code library to assist in rapid, cost-effective implementation of IAB standards, and establishes a test platform for companies to evaluate the compatibility of their technology solutions with IAB standards, which for 18 years have been the foundation for interoperability and profitable growth in the digital advertising supply chain. Further details about the IAB Technology Lab can be found at <https://iabtechlab.com>.

Disclaimer

THE STANDARDS, THE SPECIFICATIONS, THE MEASUREMENT GUIDELINES, AND ANY OTHER MATERIALS OR SERVICES PROVIDED TO OR USED BY YOU HEREUNDER (THE “PRODUCTS AND SERVICES”) ARE PROVIDED “AS IS” AND “AS AVAILABLE,” AND IAB TECHNOLOGY LABORATORY, INC. (“TECH LAB”) MAKES NO WARRANTY WITH RESPECT TO THE SAME AND HEREBY DISCLAIMS ANY AND ALL EXPRESS, IMPLIED, OR STATUTORY WARRANTIES, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AVAILABILITY, ERROR-FREE OR UNINTERRUPTED OPERATION, AND ANY WARRANTIES ARISING FROM A COURSE OF DEALING, COURSE OF PERFORMANCE, OR USAGE OF TRADE. TO THE EXTENT THAT TECH LAB MAY NOT AS A MATTER OF APPLICABLE LAW DISCLAIM ANY IMPLIED WARRANTY, THE SCOPE AND DURATION OF SUCH WARRANTY WILL BE THE MINIMUM PERMITTED UNDER SUCH LAW. THE PRODUCTS AND SERVICES DO NOT CONSTITUTE BUSINESS OR LEGAL ADVICE. TECH LAB DOES NOT WARRANT THAT THE PRODUCTS AND SERVICES PROVIDED TO OR USED BY YOU HEREUNDER SHALL CAUSE YOU AND/OR YOUR PRODUCTS OR SERVICES TO BE IN COMPLIANCE WITH ANY APPLICABLE LAWS, REGULATIONS, OR SELF-REGULATORY FRAMEWORKS, AND YOU ARE SOLELY RESPONSIBLE FOR COMPLIANCE WITH THE SAME, INCLUDING, BUT NOT LIMITED TO, DATA PROTECTION LAWS, SUCH AS THE PERSONAL INFORMATION PROTECTION AND ELECTRONIC DOCUMENTS ACT (CANADA), THE DATA PROTECTION DIRECTIVE (EU), THE E-PRIVACY DIRECTIVE (EU), THE GENERAL DATA PROTECTION REGULATION (EU), AND THE E-PRIVACY REGULATION (EU) AS AND WHEN THEY BECOME EFFECTIVE.

Glossary

Addressability Ability or extent of capability to uniquely identify an individual or a device between data sets of two or more parties in a given context e.g. targeting individuals with advertisements

Attribution Attribution in advertising is the way advertisers determine how advertising and subsequent customer interaction contributed to sales, conversions, or other goals. These metrics are used to identify the websites, apps and channels and messages that resulted in buyers taking action the advertiser wanted

Audience Group of people with a common set of characteristics whom an advertiser wants to show an ad. More specifically this is a list or group of customers or individuals that is most likely to purchase a given product or service from an advertiser

Bid Request When a publisher wants to show an ad to a user that visits their website or app, bid request is the message in programmatic advertising sent on behalf of the publisher that contains all the information required by advertiser to bid on for buying that advertising opportunity

CCPA The California Consumer Privacy Act (CCPA) is a state-wide data privacy law that regulates how organizations handle the personal information (PI) of California consumers

Click When a visitor to a website or an app takes action an ad shown to them by pressing or tapping a button inside the ad that results in them being directed to an advertiser's website or app, its called a click

<i>Conversion Event</i>	Any user action as a result of advertising that is valuable for a business for e.g. making a purchase or subscribing to newsletter or service
<i>GDPR</i>	The European Union general data protection regulation (GDPR) governs how the personal data of individuals in the EU may be processed and transferred
<i>Homomorphic Encryption</i>	Homomorphic encryption is the conversion of data into ciphertext that can be analyzed and worked with as if it were still in its original form. Homomorphic encryption enables complex mathematical operations to be performed on encrypted data without compromising the encryption.
<i>Impression</i>	An impression is when a user sees an advertisement. In practice, an impression occurs any time a user opens an app or website and an advertisement is visible.
<i>Machine Learning</i>	A mechanism and technology by which a computer can be trained to use existing data and learn how to perform a specific task
<i>PETs</i>	Privacy enhancing technologies (PETs) are technology solutions that use one or more of the privacy technologies (differential privacy, secure multi party compute and on device learning) to accomplish complex data processing functions in digital advertising without revealing the individual, household or device level personal information to parties that do not already have them
<i>Privacy Leakage</i>	It is disclosure of information that exposes sensitive details which define your identity. This can include information such as your date of birth, your SSN, your emails, usernames and passwords, home

address, phone number, and medical history

Reach

It is the number of people (or households) exposed to a given medium at a given point in time.

Re-Identification

It is the practice of matching anonymous data (also known as de-identified data) with publicly available information, or other data sets, in order to discover the person the data belongs to

Reverse engineer

It is a process of extracting information from a data set specifically identity of individuals in the data and their personal information

Secure Multiparty Computation

It is a cryptographic technique that allows multiple parties to jointly compute a function by distributing the computation across multiple parties where no individual party can see the other parties' data. Also called MPC or SMPC

Source Events

These are events that happen on the source or where advertisements are shown to the users. Typically these are impressions or clicks or other engagement the user has with the advertisement

Trigger Events

These are similar to conversion events or events that trigger conversion for e.g. a purchase made on a website

Trusted Execution Environment

A Trusted Execution Environment is a secure environment where code is executed and data is processed in an isolated private server that is inaccessible to external parties. The technology protects data by ensuring no other application can access it, and both insider and outsider threats can't compromise it

Table of Contents

[About this document](#)

[Glossary](#)

[Table of Contents](#)

[Why Differential Privacy: Use Cases in Ad Tech](#)

[What is Differential Privacy?](#)

[Differential Privacy vs Other Anonymization Techniques](#)

[Key Differential Privacy Features](#)

[Risks & Benefits](#)

[Vulnerabilities](#)

[Deep Dive: Attribution](#)

[Scenario Description](#)

[Risks of Querying Data](#)

[Protecting Data in Query Results using DP](#)

[Considerations](#)

[Output Privacy vs. Input Security](#)

[Differential Privacy vs. Secure Computation PETs](#)

[Privacy vs. Utility](#)

[Desired Business Outcomes](#)

[Other Considerations](#)

Why Differential Privacy: Use Cases in Ad Tech

Many ad tech business and data workflows require that organizations disclose to one another, data that is directly sourced -- or indirectly derived -- from data containing identifiable attributes of individual users or customers. Such workflows can be split into two categories:

1. Those that can be realized even if the data at the event level tied to identifiable attributes is altered in some way (in terms of its granularity, accuracy, or precision)
2. Those that require raw event-level data tied to identifiable attributes in its unaltered form

For an example of a workflow from (1), consider that it may be acceptable for a site level impression and conversion report for a campaign to have an error of 1% in either direction. For an example of a workflow from (2), consider a retargeting activation workflow that displays an advertisement to a specific user that exposes to the publisher *which* advertisement was seen by that individual user. *If* a use case falls into (1), it is possible that Differential Privacy can be applied to limit the amount of information that is disclosed about individual users.

Examples of digital advertising workflows into which Differential Privacy has been or may be incorporated, include:

- **Measurement and attribution analytics** workflows can incorporate Differential Privacy to limit the amount of information that reports reveal about individual users. See the [Deep Dive: Attribution](#) section below for a simplified example of how Differential Privacy might be applied in an attribution workflow.
 - The **Interoperable Private Attribution proposal** stipulates the use of Differential Privacy in aggregate queries that report on the effectiveness of ads displayed to users.
 - **AppsFlyer offers** differentially private aggregate-level attribution and reporting
- **Audience discovery** workflows that rely on aggregated user activity data can return differentially private query results.
 - **Privacy Sandbox may be augmented** with Privacy Budget, allowing sites to receive differentially private information about visiting users.
 - **LinkedIn's Audience Engagements API** provides [differentially private responses](#) to queries that can be used to identify audiences engaging with specific content on the platform.

- **Models of users or their activities** can be built (either via basic data aggregation techniques or more sophisticated machine learning algorithms) in a differentially private way before they are used to enhance user experiences.
 - **Apple** [uses Differential Privacy when leveraging user data](#) for improving typing suggestions for users and for assessing the quality of the user experience in the Safari browser.
 - **Google** has been [working on employing Differential Privacy in machine learning](#) workflows to build features that enhance user experiences.

What is Differential Privacy?

Differential Privacy is a specific approach to protecting individual privacy while sharing information about a group of individuals. This is done by adding carefully calibrated noise to the data. The amount of added noise is large enough to “wash out” sensitive individual information, but small enough so that patterns within the data can be identified with statistical analysis.

Differential Privacy vs Other Anonymization Techniques

Differential privacy is unique among anonymization techniques, some of which are described below, because the carefully calibrated noise provides a strong mathematical guarantee for the worst-case chance of revealing personal data. Other techniques provide protection that can help protect individual privacy in some scenarios, but always have the potential for unexpected leaks of information.

- **Pseudonymization:** Pseudonymization is replacing identifiable data with pseudonyms, such as a randomly generated ID linked to the data. This offers some privacy benefits, especially against accidental disclosure by honest data analysts that are not trying to reverse engineer individual identity. However, it is usually still possible to re-identify individuals by combining with outside information, and there are many avenues for accidental disclosure of sensitive information.
- **Naive Anonymization:** Slightly stronger than pseudonymization, anonymization removes identifiable data entirely, such as entirely removing information like name, email address, and user id. Similar to pseudonymization, it can provide some protection against accidental disclosure when only honest people have access to the output data, but it is often poor protection against concerted efforts to reverse engineer the inputs, and does not provide comprehensive protection against accidentally revealing sensitive information.
- **Aggregation:** Aggregation is grouping data about multiple people into a summarized form, for example total site visitors in a day or total purchases for a week. Aggregation can provide useful privacy protection if the number of people aggregated is large (e.g. thousands or more), and the range of the data being aggregated is small (e.g. includes only a few attributes). However, it is usually

possible to identify individual contributions with rare combinations of attributes in aggregated data, and there are many ways that identifying information may be accidentally disclosed, such as large outliers still being easily identified in a sum.

- **Cohorts:** It is a method that groups or labels individuals together based on shared characteristics. This contains elements of both pseudonymization and aggregation, with multiple individuals sharing the same pseudonym based on a meaningful characteristic. For example, a cohort named “Sports Lovers” that contains people who watch sports content regularly. This generally provides protection similar to pseudonymization described above, though slightly stronger due to the aggregation protection in addition. Yet, there are many avenues for reverse-engineering or accidental disclosure of sensitive personal information due to small cohort size or very unique and easily identifiable shared characteristics.
- **k-anonymity:** In k-anonymity, attribute groups (such as age range and zip code) are carefully constructed such that every possible combination of attributes identifies a group of no less than 'k' individuals, where 'k' determines how difficult it is to re-identify an individual in a data set. This is similar to regular aggregation, but more restrictive. Because k-anonymity is difficult to guarantee in an automated system, it is almost never implemented optimally in practice. Many systems claiming to provide k-anonymity actually use simple aggregation with a minimum threshold. While stronger than most among techniques that do not employ noise, there are still numerous avenues for accidental identification and reverse engineering when applied non-optimally.
- **Noise:** Adding noise is making random, usually small, changes to the data, generally after an aggregated calculation, such as adding or subtracting a bit after counting how many people saw an ad. This often provides very broad protection against a wide range of accidental leaks and intentional reverse engineering. The exact protection depends on how much noise is added and how it is randomly selected. While anonymization techniques may lead to a loss of analytical accuracy, adding noise always does, and may result in greater loss of accuracy than other methods.

- **Differential Privacy:** Differential privacy is a specific approach to adding noise where the type and amount of noise is carefully calibrated to provide mathematical proof about how unlikely it is to accidentally release identifying data and how difficult it would be to reverse engineer identifying data. The distinguishing characteristic of Differential Privacy is that it is the only approach that provides this type or worst-case guarantee.

Key Differential Privacy Features

- **Aggregation:** You need to aggregate a lot of individual data to get meaningful statistical results when using Differential Privacy because the noise washes out the contribution of any individual contribution. It is possible to apply Differential Privacy on individual level information before aggregation, but it is generally more accurate to apply the noise after aggregation. This is because the noise is proportional to an individual's data – if you add it before aggregation you have to add noise many times which will tend to cancel out but have a broader distribution than adding the noise once after aggregation. Differential privacy and aggregation are always used together in some respect.
- **Injecting Noise:** The key feature of Differential Privacy is adding a controlled amount of noise to the data. A deterministic system can not be differentially private. The noise is carefully calibrated to prevent significant changes in the aggregated results while maintaining a certain level of privacy. More noise has to be added to reach the same level of privacy if there are many variables with large ranges of possible inputs. There is a fundamental tradeoff between privacy and accuracy of the results. One of the main challenges of designing differentially private systems is finding an appropriate balance between privacy and accuracy for a given use case. An important question you can ask about a differentially private system is “Who set the level of privacy, and how did they arrive at that decision?”
- **Epsilon:** You may have heard the term “epsilon” (ϵ) in relation to Differential Privacy. Epsilon is a critical parameter in the mathematical definition of Differential Privacy that quantifies the level of privacy protection. A smaller value of epsilon indicates more noise and stronger privacy guarantees, but can also mean lower accuracy of the statistical results. A zero value of epsilon means complete privacy and no data or information is available. As the value of epsilon

increases, accuracy increases and privacy reduces. Some versions of Differential Privacy have other critical mathematical terms like “sensitivity”, or “delta” (δ) i.e. the probability of leakage and is related to size of the data– these indicate other specific properties of the mathematical guarantee being made.

- **Limited Number of Queries Against the Same Data:** Because more potentially identifying information might be leaked if more queries are made against the same data, differentially private systems often limit repeat queries against the same underlying data. For this reason, Differential Privacy can be challenging to work with in an interactive environment where you may wish to run followup queries, and you don’t know all the queries you want to run ahead of time. Differentially private systems are easier to design if you know all the queries in advance.
- **Sharing w/ Other Parties:** Applying Differential Privacy to data before releasing it or sharing it can greatly reduce the chance that any identifiable data can be gleaned by people reading the results of a calculation. This can make it easier to collaborate or publicly release information, because the participating organizations can be confident that they have not shared sensitive information. Who the data is shared to is one important consideration when choosing an appropriate level of privacy.
- **Types of Data:** Differential privacy can be applied to most types of data, including integers, continuous numerical, and categorical data. However, the more possible input values that data can take, the more noise is needed to provide the same level of privacy. Therefore it is often helpful to “clamp”, a popular technique to restrict numerical values to a specified range, reduce the number of categories, or eliminate entire non-essential variables from a statistical output. This can provide greater accuracy for a given privacy budget, and designing this well is one of the key technical tasks for building a differentially private system. It is also very important to specify which facets of the data are being protected. For instance, in the case of categorical data, it can be more difficult to design a system if the number of categories or their exact names need to remain private.

Risks & Benefits

Differential privacy provides statistical protection of privacy: the noise makes sure that you can never be quite certain what the true information about any one person in a data set is. All privacy protections have a statistical element to them – it's always possible to make a guess (potentially very well educated) about secret information and be right.

For instance, you might guess that your young neighbor plays a specific video game, “Fortress Battleground Go” and you have a certain chance of being right even if they never told you.

If the taxi service they use released aggregated ride heatmaps and you see that it's common for pickups made near your house to terminate near a game store that hosts “Fortress Battleground Go” tournaments every week, you might become more certain. If “Fortress Battleground Go” released aggregated geographical playing statistics and that shows that there are a lot of hours played in your zip code, that could give you more certainty.

If you saw over their shoulder that they were reading a website about “Fortress Battleground Go” on their phone, you might arrive at near absolute certainty.

You can always guess, through dumb luck, that someone plays a specific video game and be right. Combining more and more data points increases the certainty that your guess is right. The purpose of privacy and anonymization techniques is to limit how much information a data release reveals about a specific person.

Vulnerabilities

There are three key types of reidentifying vulnerabilities, and while they overlap and are not exhaustive, it can be helpful to understand and review your deployment of Differential Privacy for these three categories:

- **Membership Vulnerability:** A membership vulnerability is when someone observing the released dataset can learn that a specific individual's data is present – revealing their membership in the group. For instance, if someone created a custom audience based on one real person's email address and the rest is of dummy or fake emails and then they try to buy ads on a specific website with that audience. If they get any bid requests or impressions at all, they know that that real person uses that website, since none of the other “audience members” could have visited since they don't exist.

- **Differencing Vulnerability:** A differencing vulnerability is when multiple different releases of the same data reveal private information. Systems that rely on aggregation alone can often be very vulnerable in this way. For instance, let's say someone can make attribution queries in a data clean room, but each query has to return at least 1000 users. If an analyst can make multiple queries and has no other limitations, they can set up one that has 1000 users in it, and then another query that adds just one person of interest. If anything about the query changes, they can learn the exact information about a specific person.
- **Linkage Vulnerability:** Linkage vulnerability is when released data can be correlated with external sources to identify individuals or reveal more information about them. Linkage vulnerability is one of the deepest sources of overall vulnerability, because it can turn a seemingly innocuous piece of information into a serious data breach. Consider a very basic 3rd party cookie example, where the only information being shared is that the same (anonymous) person visited two specific websites on the same day. That information doesn't seem that interesting, but by correlating timestamps it might be possible to correlate one visit to a specific credit card purchase with real personal information, and connect it to political comments on the other website. Linkage vulnerability is often particularly devastating to naive anonymization and pseudonymisation, but does threaten all anonymization techniques.

These vulnerabilities interact and overlap in uncountable ways in most real systems – most defenses make certain types of vulnerabilities less likely under common scenarios, but no system that releases information can prevent them. Differential privacy provides a robust defense against all of these types of attacks, by adding enough noise that someone trying to reverse engineer the data or deanonymize someone else can never be sure that they got it right. The noise is calibrated to provide a hard limit on that certainty, no matter what technique they use to reverse engineer. Injecting noise is the only method that can have the property of providing measurable protection against all output privacy vulnerabilities, and Differential Privacy is the specific approach to injecting noise that does it efficiently with no gaps.

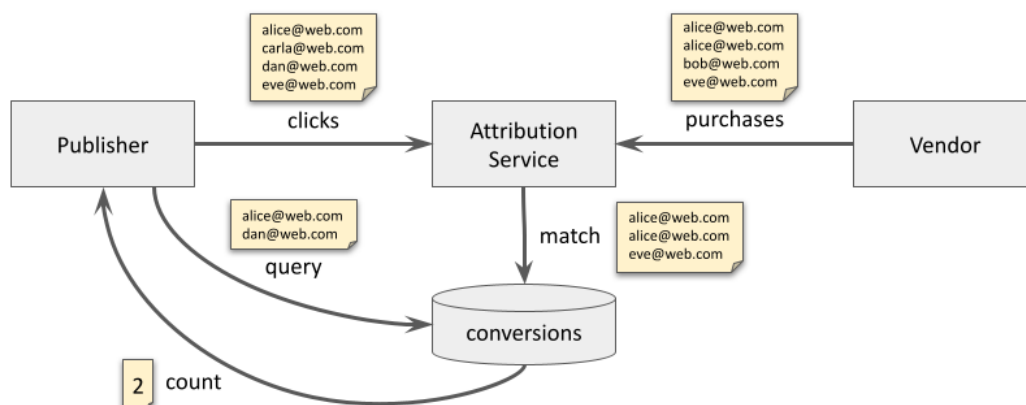
Deep Dive: Attribution

We consider an example use case (inspired by an [Ad Tech Explained article](#)) to illustrate and motivate the basic and essential concepts and features associated with Differential Privacy (DP).

Scenario Description

In this scenario, a **publisher** operates a website on which users view content after logging in using their email address. This website also displays advertisements from a **vendor**. The publisher records **source events** (impressions and clicks) associated with each email address.

When a user clicks an advertisement, they visit the vendor's website. Once there, the user has the option to log in with their email address and make a purchase. The vendor website records **trigger events**: purchases of products associated with a user's email address.



Source and trigger events are submitted to an attribution service that matches them up by email address. This can be either a trusted service such as a clean room or an ensemble of organizations running a secure multi-party computation protocol (as in the Interoperable Private Attribution protocol). The exact manner in which this is accomplished is not important for understanding the basics of DP (but any such organization or infrastructure can support DP queries against the matched data).

Risks of Querying Data

Suppose that a publisher wants to request the **count of conversion events** from the attribution service using a *query* consisting of a **list of recorded source event identifiers** from the publisher's own data. Given such a query, the attribution service sends back a *reply*: a **count of the total number of matching conversion events** (which may be approximate).

If the publisher is allowed to make any number of queries of any size to the attribution service, they can submit queries containing individual source event records to learn exactly which individual users' clicks led to conversions (and, potentially, those individual users' interests and purchasing histories) Even when limited to only a few queries, the publisher can submit two queries that differ by exactly one record (for example, "**alice@web.com**"). The difference between the two replies (*i.e.*, the two counts) can then reveal whether Alice's individual click led to a conversion.

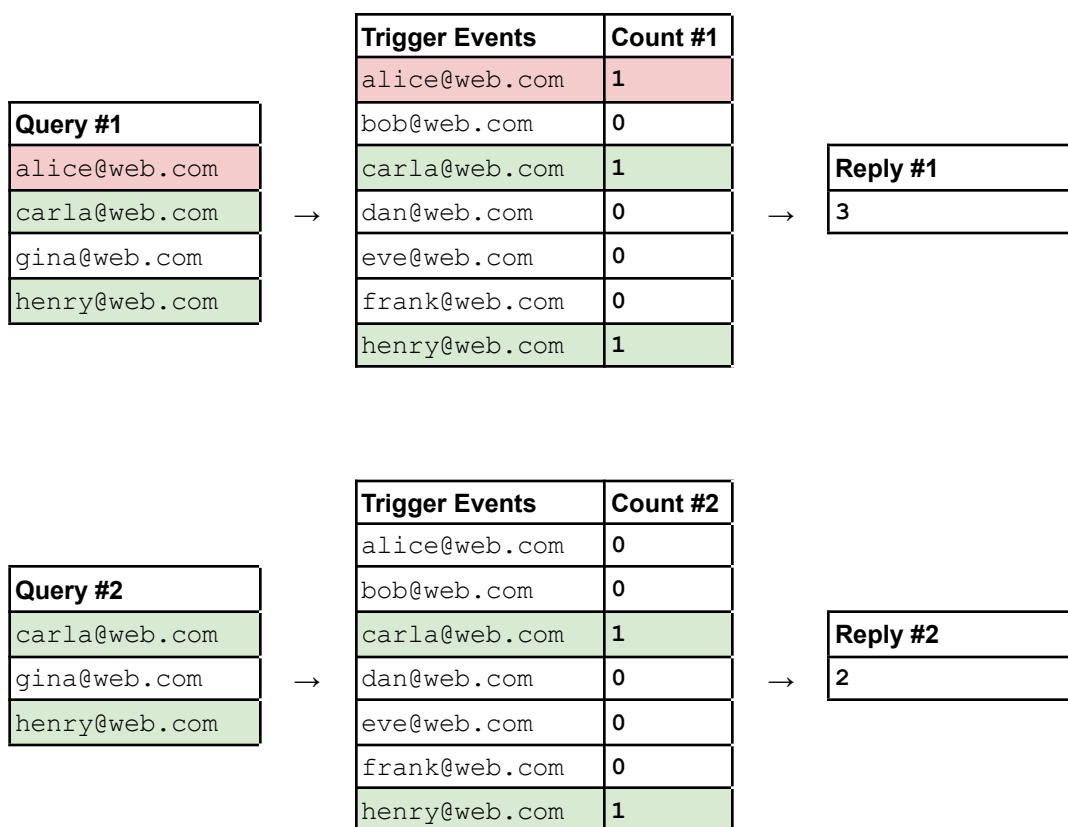


Figure 1. Diagram illustrating two attribution queries and their corresponding replies.

Is it possible to allow the publisher to submit queries that return an approximate count of matching conversion events, but also to protect individual users found in the data (such

as “alice@web.com”) from having their activity revealed in the manner above? Differential privacy techniques aim to allow just that.

Protecting Data in Query Results using DP

To mitigate the publisher's ability to learn about its individual users' conversion histories with just a few queries, the attribution service can instead return *differentially private* counts in its replies.

As an example, suppose that instead of using a fixed constant “1” for every source event that matches a conversion event, the attribution service returns “1” with a 60% probability and “0” with a 40% probability. Informally, this can be viewed as a "noisy" input to the summation operation that is used to calculate the count.

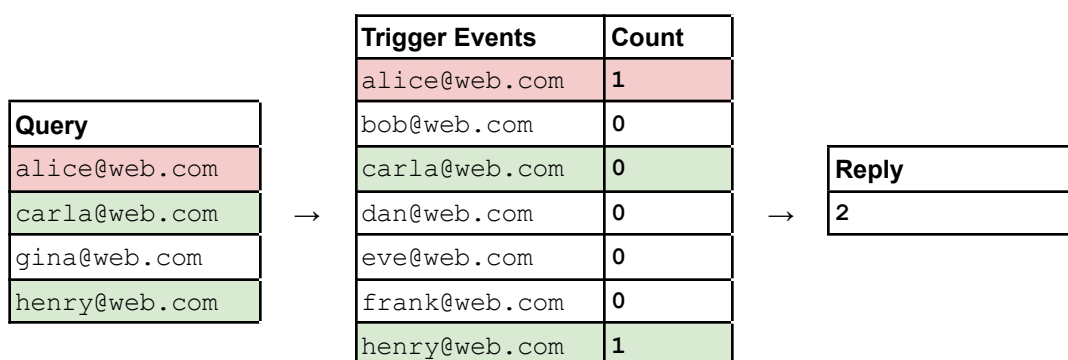


Figure 2. Diagram illustrating a query and reply when the count is a sum of “noisy” values.

When using this implementation of the workflow, the same query does not necessarily yield the same result each time that it is submitted. Figure 3 illustrates a distribution of replies for 10,000 queries submitted to the modified workflow that is illustrated in Figure 2.




Figure 3. Distribution of 10,000 replies from the modified workflow.

However, with enough pairs of queries, the difference becomes easier to discern. Figure 4 compares the distribution of replies over 5 queries to the distribution of replies over 10,000 queries. Notice that the latter provides more information. This difference is what motivates the concept of a *budget* in Differential Privacy: a limit on the number of queries that can be executed before the features of Differential Privacy no longer apply.




Figure 4. Distribution of reply differences (for Query #1 and #2) for the modified workflow over 5 trials (left) and over 10,000 trials (right).

In this section, we have seen that data belonging to individual users can be protected by adding noise in a specific way within the overall data workflow. We have also seen that even with the introduction of such noise, there is a budget of queries beyond which that noise may no longer be enough. In rigorous implementations of Differential Privacy, a precise mathematical relationship is defined between the amount of noise and the limit on the number of queries. More details about various aspects of Differential Privacy implementations are discussed in the sections below.

Considerations

There rarely is a one-size-fits-all implementation of Differential Privacy which applies to all use cases and policy interpretations across organizations. When considering whether Differential Privacy is an appropriate privacy-enhancing technique to apply to your data, it's important to understand your needs in terms of:

- Input versus output privacy protection guarantees
- Privacy versus utility tradeoff requirements
- Use case and desired business outcomes

Based on your needs on these three dimensions, you can interrogate whether a given vendor's implementation of Differential Privacy will meet your requirements.

Output Privacy vs. Input Security

Most implementations of anonymization techniques in industry focus on preserving *output privacy*, or preventing users from reverse engineering results of an analysis or process to glean information about its input data. This is most relevant in cases where a data custodian wishes to ensure collaborators or even internal employees are not able to tie results from a process back to individual data subjects with a certain degree of confidence. Differential privacy is a powerful tool for this if configured optimally because of its mathematical guarantees of privacy throughout the processing chain of events.

However, Differential Privacy is not generally applied to protect *input security*, or the guarantee that parties storing, transporting, or working with the data cannot glean any information about the inputs. An illustrative example of input security is what happens when you send a letter in the mail. By putting the letter in an envelope and sealing it, the contents of the letter are not readable by the postal staff while it is in transit. In technical terms, input security is typically ensured by encrypting data. This can be accomplished using standard, widely deployed encryption techniques (such as HTTPS when transmitting data or block ciphers when storing data and relying on standard access controls) or alternative privacy-enhancing technologies (like homomorphic encryption or secure multi-party computation to compute using that data).

Based on the sensitivity of your data and your organization's security or privacy policies, you will want to understand where it is most important to protect input security, output privacy, or both. If both are required, it is often recommended to use Differential Privacy in combination with other privacy-enhancing techniques to satisfy the requirements. Vendors offering Differential Privacy options often also incorporate mechanisms to

protect input security and should be able to answer any questions you have about requirements.

Differential Privacy vs. Secure Computation PETs

Differential privacy and other anonymization techniques aim to provide **output privacy** – ensuring that people that can see the output can't reverse engineer individual information from the input. This is distinct and complementary to input security. The strongest systems that operate on sensitive individual data provided by multiple parties will use a secure computation PET and Differential Privacy together.

There are several popular approaches to secure computation with multiple parties, including homomorphic encryption, secure multiparty computation, and trusted execution environments. These each have unique benefits and limitations, but in general they provide a secure way for multiple parties to combine data for computation. This provides **input security** – ensuring that no one can see the raw inputs to a calculation. Even when combined, these techniques still generally “only” provide input security, and are appropriate for systems where more than one party will contribute data.

Multiple Data Contributors?	Sensitive Data?	Appropriate Technology
No	No	Traditional techniques
No	Yes	Differential Privacy
Yes	No	Input Security PET
Yes	Yes	Input Security PET and Differential Privacy

Privacy vs. Utility

For some use cases, you may not want or need to protect output privacy. In these instances, Differential Privacy may not be a solution to your problem. As discussed above, use of Differential Privacy represents a tradeoff between achieving privacy i.e. prevention of identification of individual information and utility i.e. extracting useful and accurate information from the data. In cases where you have a high level of trust in those analyzing the data, low likelihood of adversarial attacks (people trying to

maliciously leak privacy from your data), or a need for close to perfect accuracy of results, applying Differential Privacy might be unnecessary or even counterproductive. In these cases, you might consider common industry practices around privacy and security measures sufficient to protect data while still providing the flexibility and accuracy necessary for your use cases.

However, in many cases, data contains sensitive and identifying information about individuals, trust between parties that are sharing or collaborating is low, or regulatory compliance requires that individual information is not disclosed. In such cases, the risk of privacy leakage becomes much more important than having particularly accurate results from analysis. In these cases, Differential Privacy is likely one potential solution which will meet your needs. Solutions exist in-market for both approximate application of Differential Privacy, without the need to provide access to raw data, and optimal privacy/accuracy tradeoff calculations, which does require access to underlying data or the introduction of additional PETs. Which type you choose will depend on your own organization's risk tolerance, privacy policies and legal requirements regarding data access and analysis.

It's important to understand from vendors whether you are able to fine-tune the privacy-utility tradeoff based on the sensitivity of the data involved, the use case needs, your requirements around data access, and the parties involved in use of the data.

Desired Business Outcomes

Finally, the use of Differential Privacy can make certain use cases or collaborations possible where they otherwise might have been considered too risky. However, it's important to understand whether a vendor's implementation of Differential Privacy still allows you to achieve your desired business outcomes. You can determine this by asking vendors where in the process Differential Privacy is applied, how it will typically affect expected outputs, and what downstream processes need to occur to achieve the desired outcome.

In some cases, Differential Privacy might be applied at a point in the data processing workflow which might not make sense for your goals. In these cases, you can often work with the vendor to understand how to configure your implementation such that you can benefit both from Differential Privacy guarantees and the ability to execute on data in the manner you need to achieve the desired outcomes.

Other Considerations

Solution providers with Differential Privacy application options can typically help you navigate the three dimensions outlined above for your datasets or use cases and understand whether their solutions are appropriate for your business.

Example questions you may wish to ask when evaluating solutions are:

- What is the end-to-end process flow for how Differential Privacy is applied?
- Explain the features included in your Differential Privacy solution and how the technique is applied? What epsilon values are used and how they are determined.
- To which use cases is your implementation of Differential Privacy most applicable?
- What are the tuning options for your application of Differential Privacy? Are there different options depending on different use cases or policy needs?
- How does your solution prevent privacy leakage? What attacks are contemplated in your approach, and how do you mitigate the risk of privacy leakage? E.g. limit on no. of queries permitted.
- What assumptions does your implementation of Differential Privacy make? What are its limitations?