



Data Anonymisation

What it is and Why it Matters

Executive Summary

We live in the golden age of data. Companies are seeking to exploit the data they gather on their customers in order to improve their services and develop new products. With that comes a vastly increased risk of data loss and accidental data breaches. The new GDPR regulations are designed to protect against that, but they threaten to reduce the ability to analyse these datasets. Aircloak have developed Insights, the first GDPR-approved product to dynamically anonymise data as it is queried.

A Data Gold Rush

Today we find ourselves in the middle of a data revolution. Exciting new technologies like artificial intelligence, the internet of things, and self-driving cars are all made possible by the new science of data.

A good analogy is that data is to this century as oil was to the last. Today, data drives growth and change, opening up new avenues of risk and innovation. Now everything, from trains to toasters, has the potential to be a source of valuable, analysable data.

Understandably, this trend shows no signs of slowing down. In fact, it's expected to experience explosive growth in the coming years, with the big data industry reaching \$66.8 billion by 2021.ⁱ

The Rise of Data Protection

This new data gold rush has brought with it increased concern around data security and privacy. High profile data breaches by large Corporations have helped make customers aware of what is happening to their personal data. And after Edward Snowden's revelations

about the post 9/11 erosions of citizens' rights, the average citizen is no longer willing to stand idly by as Corporations profit from vast oceans of personal data while putting it at risk.

The General Data Protection Regulation

The General Data Protection Regulation (or GDPR) is a new piece of legislation passed by the European Union to protect the personal data rights of EU citizens. As of May 25th, 2018, many new data protections will come into force. For example, individuals will be given the right to access any of their own personal data being held by companies. They will also gain the so-called "right to be forgotten," which is the right to have their personal data destroyed upon request.

Under GDPR, companies will need to implement "data protection by design and by default."ⁱⁱ They will also be required to promptly notify the Data Protection Authority and any affected individuals in the event of any data breach. Furthermore, consent to use personal data must be freely and specifically

given rather than buried in a sea of complex legal jargon.

One important thing to be aware of is the breadth of the GDPR's jurisdiction. The new regulation applies to all companies collecting, storing, or using the personal data of EU citizens, regardless of where they currently reside.

What happens if you fail to comply? The fines are no laughing matter: 20 million Euros or 4% of annual global turnover, whichever is higher, and the pain doesn't end there. The EU will also have the power to ban non-compliant organisations from trading with any nation that has adopted the GDPR into national law.ⁱⁱⁱ Not a bad motivation to update your organisation's data security policies.

Is the GDPR Bad for Business?

It doesn't have to be. For one thing, it makes EU data privacy legislation clearer and more consistent. In addition, with all the public focus on personal data security, there is a real opportunity for companies to differentiate themselves through superior data privacy practices. Smart companies will use the GDPR as the perfect time to revisit and improve their data privacy programs.

"The GDPR comes as a welcome piece of legislation, and in many ways the reform creates a strong and comprehensive set of rules that need to be applied in order to sufficiently protect data," says Phil Bindley, CTO of the service provider The Bunker. "In doing so, this creates more trust in the digital environment and means that a privacy-focused approach can work in conjunction with the use and analysis of data."^{iv}

Loss of Reputation from Data Breaches

If your organisation suffers a data breach, you may have more to worry about than massive regulatory fines. In a recent survey conducted by the Ponemon Institute, 75% of companies

said their reputation was a key asset but fewer than half thought that their reputation could withstand a negative event such as a data breach.^v

Additional research bears this out. According to data from Semafone, the overwhelming majority of people would stop doing business with a company that had suffered a breach. As Tim Critchley of Semafone put it, these numbers prove "what we should already know—that the reputational damage suffered by companies who fail to protect personal data can translate directly into loss of business."^{vi}

Anonymisation vs Pseudonymisation (and Why it Matters)

With the increased scrutiny around data privacy, tomorrow's winners will be those who figure out how to leverage data insights without compromising privacy.

Data protection laws exist to protect the personal identity of people whom the data describes. If the data subject is not identifiable in any way, data protection law does not apply. That is, when it becomes impossible to connect individuals to the data, people controlling and processing the sensitive data are not restricted in data use or sharing. On the other hand, an identifiable data subject can result in legal consequences including damage claims, loss of reputation, and fines or penalties. Consequently, data anonymisation becomes vitally important because it allows you to use your data for marketing insights without jeopardising individual data privacy.

The purpose of data anonymisation is privacy protection. It involves the modification of data sets so that no personally identifiable information remains. As a result, data can be used and transferred without individuals' identities being disclosed unintentionally.

Pseudonymisation differs in that personally identifying fields in the data are simply replaced by pseudonyms such as random numbers. Simply replacing these fields, however, does not eliminate the possibility of

identifying individuals in the data. In fact, a study of human mobility data revealed that four spatiotemporal points are enough to uniquely identify 95% of the individuals in a dataset, even if the granularity is low and you have no additional information about a user. ^{vii}

For a mobile call records dataset, pseudonymisation might involve replacing people's names and phone numbers with random (although not necessarily arbitrary) numbers. However, this is no protection against what's known as an inference attack. For example, if the analyst knows details of 2 or 3 calls that a friend made at certain times from certain places, they could (re)identify their friend's records from the pseudonymised data.

True data anonymisation is difficult to achieve, and many data controllers fail to do so properly and completely. Proper anonymisation often destroys the value of data to the point where the whole exercise becomes moot. This happens because when you anonymise a dataset before analysing it, you have to anonymise it across all dimensions that are identifying (or potentially identifying) in combination with other attributes or in combination with external knowledge. Understanding what could be identifying (beyond the obvious things like a social security number or phone number) is difficult and usually underestimated.

In fact, most companies use weaker pseudonymisation techniques to protect personal data, which also means many companies will be constrained by data privacy laws and subject to penalties when GDPR comes into effect.

The EU has developed a list of criteria for anonymisation published in WP29's Opinion on Anonymisation Techniques ^{viii}. First of all, WP29 makes it clear that pseudonymisation is not strong enough to protect personal data. Pseudonymisation permits data controllers to handle their data more liberally, but it does not abolish all risks due to the possibility of re-identification. Crucially, this means pseudonymous data is still subject to privacy regulations under GDPR.

Overall, WP29 is intent on educating organisations about proper anonymisation techniques and they are calling for certification so that companies follow their requirements. Data protection authorities (DPAs) can carry out this certification. Yet no clear guidance currently exists on how data anonymisation techniques should be certified.

Because there are no specific guidelines on how to ensure anonymisation, DPAs are certifying data anonymisation practices on an ad hoc basis. Instead of following specific guides, organisations are making certifications based on their own interpretations of the GDPR's loose criteria. There is no general anonymisation certification program that companies can follow before being certified. It is, for instance, not clear what happens legally when a certified anonymisation process is later discovered to be weak or was performed by someone who is unqualified. It is difficult to tell who is legally liable in such situations, and non-compliance can result in steep penalties after GDPR comes into effect.

There are a number of traditional data anonymisation techniques, like K-anonymity and differential privacy, that are simple and provide strong anonymity but unfortunately destroy data utility.

To retain data utility, data controllers typically choose from a complex variety of mechanisms that provide some anonymity but may not protect against re-identification in many cases. These include rounding, cell swapping, outlier removal, aggregation, sampling, and others. Getting this right requires substantial expertise both on the part of the data controller and the DPA. And again, if done right, this process typically destroys the utility of the data

Diffix: A new approach to data anonymisation

Aircloak, in research partnership with the Max Planck Institute for Software Systems, has spent many years researching Privacy-Enhancing Technologies (PETs) and have introduced a new approach to anonymising data sets called Diffix.

Diffix is the base of Aircloak's flagship product, Insights™, a solution for these privacy regulation problems. Aircloak Insights provides strong anonymity and good utility for a wide range of use cases and requires no special expertise to setup and configure. Aircloak Insights maintains the quality of the data set for analysis but also achieves a strong level of anonymity.

Because Aircloak Insights works "out-of-the-box" for a wide range of use cases, once a Data Protection Authority has certified data anonymised by Insights, any new use case requires little to no additional evaluation. CNIL, the French national data protection authority, has already evaluated Diffix against the GDPR anonymity criteria, and have stated that Aircloak delivers GDPR-level anonymity.

As a result, data that is anonymised by Aircloak's solution will not be constrained by privacy laws or be the subject of penalties when the GDPR comes into effect.

Case Study: A large European banking corporation

One of Aircloak's clients had traditionally been a provider of consumer credit products. They wanted to find ways to acquire high-value customers, better understand their existing customer base, and offer new products. This was the inspiration behind their creation of a mobile application that allowed end-users to aggregate their bank account information and financial transactions. In return, the application provided valuable financial insights and recommendations.

The client relies on continuous detailed customer analytics to better understand their customers' financial situations. This knowledge makes it possible to deliver the best recommendations and advice to customers and improve the usability and functionality of their app without creating unnecessary features.

While enhancing user experience is a business priority for the client, so is guaranteeing the

privacy of their customer data. It is necessary that the company ensures that nobody - not even their own analysts - can access individual user data. Therefore, all aggregate analysis of customer data must be carried out with the highest safety and anonymity possible.

Prior to using Aircloak, they achieved safe aggregate data analysis using a classical approach.

Whenever a new type of exploratory analysis was considered, an approval process had to be followed involving various internal stakeholders and privacy professionals. This process was complex and could take weeks to complete.

Following the approval process, the client would provide a snapshot of their data set to a third-party expert organisation who would anonymise it using traditional data-masking techniques. While safe, this process resulted in additional delays and a severe reduction in data quality. At times, the resulting data quality would be so low that a new analysis approach would need to be devised, beginning the process all over again. Additionally, whenever new data became available, the anonymisation procedure had to be repeated.

Pain Point: A Slow and Complex Anonymisation Process

A key problem for the client was the fact that the quality of free-form text information suffered greatly from anonymisation. This is because only white-listed terms were allowed through while everything else was immediately removed or made illegible. Changes to the white list required new approvals which cause additional delays.

The client's approach, although compliant, was far too slow and complex. By working with third-party data processors, they introduced risks to the system and were left with inaccurate data that impaired analysis accuracy.

The Client needed a new solution to analyse customer financial transaction data that:

- Satisfied the highest level of privacy protection available and was fully

compliant with current and future privacy regulatory requirements (GDPR).

- Was simpler and faster to execute.
- Did not introduce dependencies on external data processors.
- Did not destroy valuable information in their data.

Aircloak's Fast and Simple Solution

Aircloak Insights takes a fundamentally different approach to securing the privacy of databases. Architecturally, Aircloak Insights acts as a proxy between a database and an analyst. Data can be queried through the product's query editor, by using Aircloak Insights' API, or by using Tableau Desktop – software designed to help people visualise and understand their data.

Instead of anonymising or pseudonymising the source data to the point of potentially becoming unusable, analytics are done on the unprocessed data. The analytics output is then instantly and automatically anonymised as queries are run, thus guaranteeing by design the strongest level protection of a client's private customer data. The key to this process is the fact that the analytics and anonymisation steps are one and the same. If anonymisation is done before analytics, the analytics suffers, if it's done after then anonymisation cannot be ensured. Because Aircloak has the technology to do it all at the same time, it ensures both strong anonymity and good data quality.

Because Aircloak Insights automatically ensures that no personal information ever leaks, our client has been able to greatly optimise their existing process of analysing private customer data in the following ways:

- Any approved internal stakeholder can now run arbitrary queries over the raw data.
- There is zero dependence on a third-party processor, which eliminates a source of risk, cost and complexity.
- Delays for the entire process are minimised.

- The client can be sure that they are compliant with all current and future data protection regulations.

Client Benefits of Aircloak Insights

Aircloak Insights provided the client with a solution that ensures compliance with privacy and data protection regulations. In particular, a European data protection authority has verified that Aircloak satisfies the anonymisation requirements of the GDPR. So our banking client can now better understand and segment its customer because they are able to use the full transaction details for analysis in a secure and compliant way. This allows them to make better recommendations to the customer regarding their financial situation and potential relevant products.

Instant Privacy Compliance: Share Insights Not Data

Aircloak's first-in-class real-time database anonymisation solution provides instant privacy compliance and enables high-quality analytics for any data set and any use case. Aircloak has been evaluated and approved for GDPR-level anonymity for all use cases by CNIL, the French data protection authority.

Aircloak unlocks sensitive data for analysis, sharing, and monetisation. Aircloak's strong anonymisation satisfies the world's strictest anonymity laws, slashing compliance costs and all but eliminating data sharing risks.

Aircloak's patented breakthrough technology is the first anonymisation solution to provide high-quality analytics, unprecedented ease-of-configuration, and strong anonymity. Prior approaches that use pseudonymisation provide good analytics but weak anonymisation, falling well short of European standards. Prior forms of strong anonymisation, like K-anonymity, tend to destroy the data in all but the simplest use cases.

Aircloak's ease-of-configuration far exceeds all previous solutions. There is no need to flag or

remove Personally Identifiable Information (PII), pseudo-identifiers, or sensitive data. The existing primary-use database is not modified in any way, and Aircloak handles all data types including unstructured text. Aircloak is installed in front of the database, and relays queries and answers between the analyst and the database.

Aircloak offers analysts a rich explorative SQL database interface. Analysts submit SQL queries and interact with the existing database to extract the requested data. Both queries and answers are dynamically modified by Aircloak to ensure anonymity while still providing high accuracy.

To the analyst this has the look and feel of querying the database directly. The analyst can specify tables and columns, join tables, set ranges, run standard or custom SQL aggregate, maths, and string functions, and much more. The analyst receives answers at the speed of

SQL, thus allowing the same interactive and explorative analytics the analyst would expect from native access to the database.

Aircloak enables new use cases; for example, improving internal business intelligence, external data monetisation, and reducing the cost of mandatory reporting. Internally, divisions within large corporations can use Aircloak to share customer data with other divisions or out-sourced analysts. Organisations with valuable data, such as health, location, or financial data, can instantly and safely monetise the data regardless of its format. Aircloak gives organisations full visibility into how their data is being used, leading to better billing models.

Contact us today and let us show you how we can help you make wider use of your data by allowing you to share insights and not the data itself.

Aircloak GmbH

Gormannstrasse 14

10119 Berlin

Germany

solutions@aircloak.com

ⁱ <http://www.information-age.com/big-data-vs-privacy-big-balancing-act-123461795/>

ⁱⁱ <https://gdpr-info.eu/art-25-gdpr/>

ⁱⁱⁱ <https://www.red-gate.com/simple-talk/opinion/opinion-pieces/questions-gdpr-shy-ask/>

^{iv} <http://www.predictiveanalyticsworld.com/patimes/big-data-vs-privacy-balancing-act/7975/>

^v <https://www.experian.com/assets/data-breach/white-papers/reputation-study.pdf>

^{vi} <https://www.csoonline.com/article/3019283/data-breach/does-a-data-breach-really-affect-your-firm-s-reputation.html>

^{vii} <https://www.nature.com/articles/srep01376>

^{viii} http://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf