PRIVITAR

The Right to be Forgotten

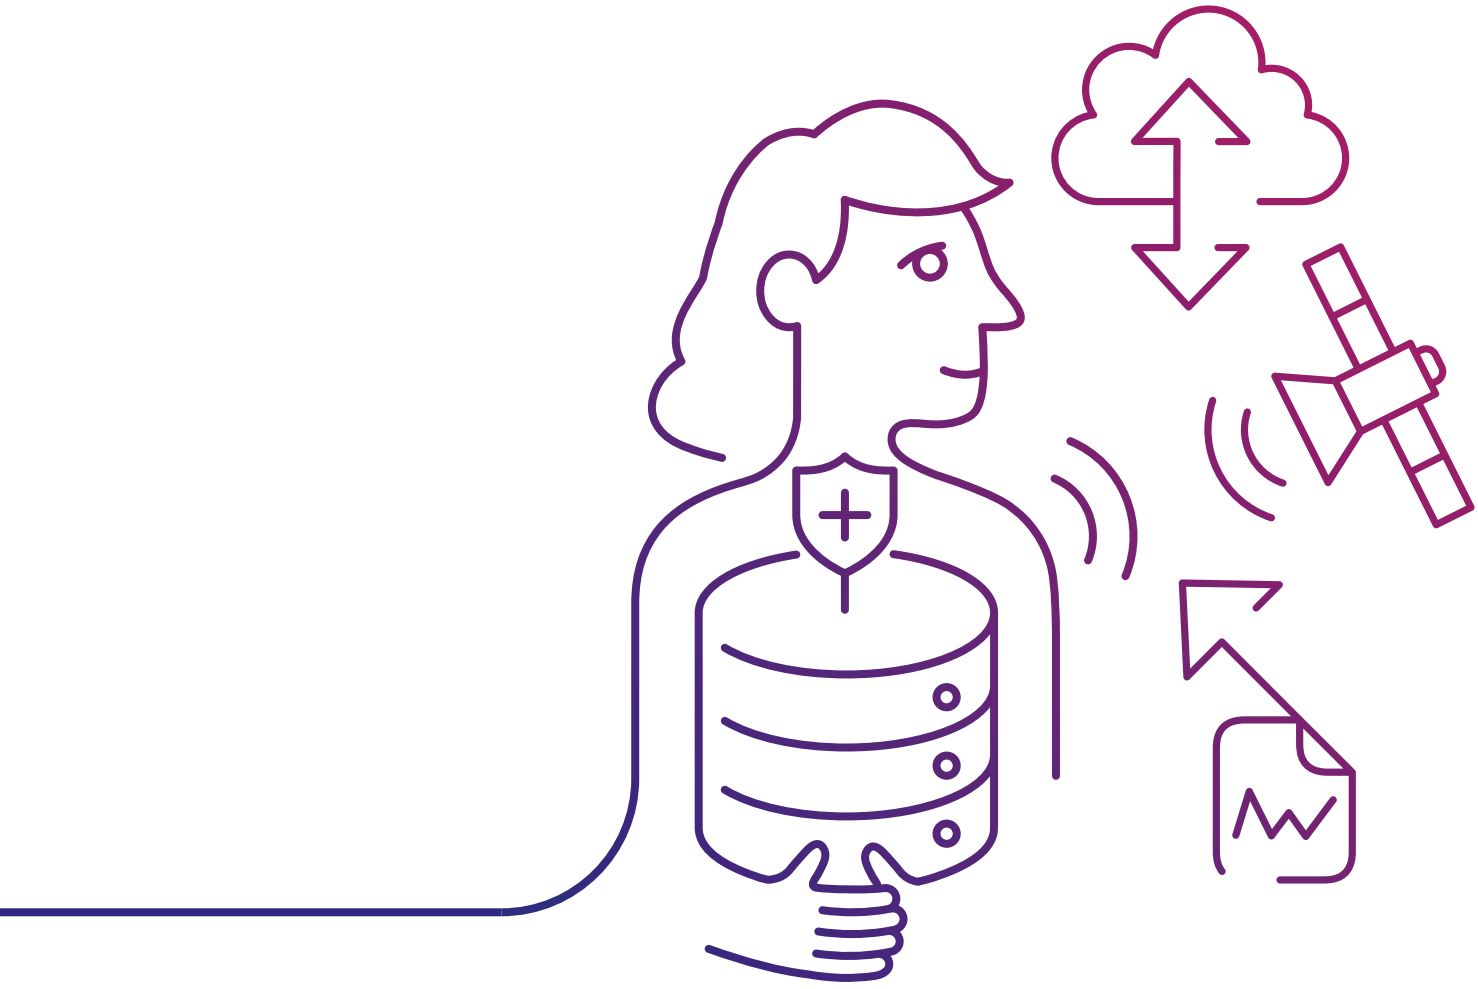

www.privitar.com

Abstract

This paper is intended to support those preparing their organisation for the General Data Protection Regulation (GDPR) and considering what the Right to be Forgotten (RtbF) means for their organisation.

It covers:

- What the RtbF is
- How it has been changed by the GDPR
- What guidance is available on the RtbF, and what that guidance says
- How organisations can prepare for, and manage, RtbF requests.

About the author:

Guy Cohen is Policy and Strategy lead at Privitar. In addition to his role at Privitar, Guy is a fellow at the University of Cambridge Centre for Science and Policy. Before joining Privitar, Guy worked in various roles in the Civil Service, Cabinet Office, the Department of Health and HMRC.

Introduction

The Right to be Forgotten (RtbF) is not new. Although the 2016 General Data Protection Regulation (GDPR) will be the first law to refer to it explicitly, it remains largely unchanged from the 1995 Data Protection Directive (95/46/EC – The Directive). Under the Directive, individuals have the right to object to processing of their personal data, and have that data deleted if their objection is valid.

What has changed in the new GDPR is the onus for justification. Under the Directive it is for the individual to demonstrate why their data shouldn't be processed. Under the GDPR this will be reversed, so controllers will have to demonstrate why the data should be processed.

This paper will explore what the RtbF is, highlight some key points for consideration, and look at how privacy engineering techniques can be used to respect individual privacy, whilst minimising disruption and cost.

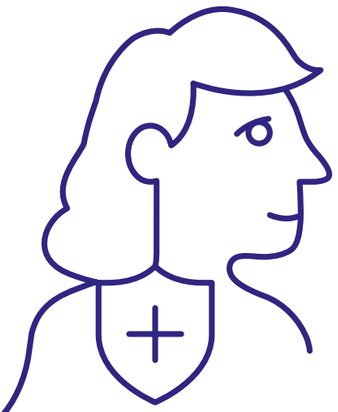

Background

The RtbF dates back to at least 1978, when the Organisation for Economic Co-operation and Development established seven principles of data protection.

The last of these was that data subjects should have the right to gain access to their data, to challenge such data, to request erasure, and to have the right to challenge any denial of these rights.

The 1995 Directive does not explicitly provide a RtbF, and so it was not until the 13th May 2014 that the Court of Justice of the EU ruled that the RtbF was a necessary result of Articles 6 and 14 of the 1995 Directive (see below).

It is worth noting that the right is not absolute, and will continue not to be; to date Google accepts roughly 42% of RtbF requests. Of those in the UK subsequently appealed, which is a relatively small proportion, roughly 25% are overturned by the Information Commissioner's Officer (ICO), implying that the ICO mostly agrees with Google's decisions.

Since May 2014 Google alone has received hundreds of thousands of RtbF requests. But whilst RtbF has been a right for several years now, some details are still in debate. For instance, the Commission Nationale de l'Informatique et des Libertés (CNIL), the French equivalent of the UK's ICO, is currently in a legal dispute with Google regarding the reach of the RtbF.

Google argues the right should be upheld by removing the listings from all searches from EU IP addresses.

CNIL argues that it should be from all searches worldwide, irrespective of where the search is taking place.

To date Google accepts roughly 42% of RtbF requests.

So what does the GDPR change?

Under both the 1995 Directive and the new GDPR, data subjects have the right to object to processing of their data where the controller's basis for that processing was under a public or legitimate interest, and to request the erasure of that data.

The main legal difference is that the onus of justification for processing has shifted from the data subject to the data controller.

As an example, imagine at the moment an individual wants to have a search result revealing information about them deleted. Currently they would need to write to the search engine explaining how their privacy was put at risk and the search engine would then weigh their request against the public interest.

If there continued to be a disagreement, then the ICO would examine the individual's case for why their privacy had been violated. Under the GDPR this will change so the search engine will have to explain their reasoning for why they thought the processing was acceptable and an individual will be able to raise an objection. If there is a disagreement the ICO will then weigh the search engine's case for processing.

What does the law actually say?

Under the Directive the data subject had the right:

"...to object at any time on compelling legitimate grounds relating to his particular situation..."
– Article 14, 95/46/EC

Which meant that if they could show that the data held was not:

"...accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that data which are inaccurate or incomplete, having regard to the purposes for which they were collected or for which they are further processed, are erased or rectified."
– Article 6, 95/46/EC

Then, the data would need to be erased. The new GDPR gives the same right, but the right to object is changed, so that now:

"The controller shall no longer process the personal data unless the controller demonstrates compelling legitimate grounds for the processing."
– Article 21, GDPR

Under the 1995 Directive, the 'compelling legitimate grounds' is implied to come from the data subject, but under the GDPR, it is for the controller to demonstrate why their processing is legitimate.

The balancing between a data subject's RtbF and a controller's interests are explained further in recital 69:

"Where personal data might lawfully be processed...a data subject should, nevertheless, be entitled to object... It should be for the controller to demonstrate that its compelling legitimate interest overrides the interests...of the data subject."

Additional changes

Under the GDPR, erasure can also be requested if the data "have been unlawfully processed" – Article 17. This is broader than just the basis for processing. For instance, one example would be Article 30, which describes the meta-data which must be recorded for all processing (e.g. the purpose of processing, who'll see the data, etc.).

Not recording this data would be unlawful, and so were a processor to have failed to record this meta-data then the data subject could potentially request their data be erased, even if the grounds for processing are legitimate.

More serious examples would include processing which infringed on copyright or breached a duty of confidence.

Finally, the rules for consent under the GDPR are more onerous on the controller. This may lead controllers to change the legal basis for some of their processing from consent to legitimate interest. The increase in legitimate interest usage may affect the importance of the balancing test (see below) for some organisations, which will have implications for consideration of RtbF requests.

How does the legal basis for processing affect RtbF requests?

Under the GDPR there are 6 different grounds for the processing of personal data. How a company responds to an RtbF request will be dependent upon what their legal basis for processing is, as shown by the table below:

Basis for processing	Required action
Consent	Must erase data - Article 17 (1 (b))
Contract	Can reject request, if contract still in effect - Article 17 (1 (a))
Vital interest	Can reject request, if still relevant - Article 6 (1 (d))
Compliance with legal obligation (inc. legal claims)	Can reject request, if still relevant - Article 6 (1 (c)) (N.B. individual countries have some discretion to change how this works - Article 6 (2))
Public interest	Controller must weigh public interest case against individual's privacy rights - Article 6 (1 (2)) (N.B. individual countries have some discretion to change how this works - Article 6 (2))
Legitimate interest	Controller must weigh their legitimate interest against individual's privacy rights- Article 21 (1 (1))

What guidance do organisations have or trying to understand how to weigh individual rights against other interests?

Companies can reject RtbF requests when the public interestor their legitimate interest outweighs the individual's privacy.

The balancing test is the name for the process by which a controller weighs the interest (public or legitimate) in the processing against the data subject's right to privacy.

Following the Google-Spain case, the Article 29 Working Group (A29WG) published guidance on the RtbF ruling,which can be found here:

http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2014/wp225_en.pdf

As well as the A29WG, Google has also published its own report on how it intends to implement the ruling. The committee which they assembled for the report included leading figures in data protection ethics. The report can be found here:

https://drive.google.com/file/d/0B1UgZshetMd4cEI3SjlvV0hNbDA/view

Both the A29WG and Google's guidance are of only limited relevance as they are focused on the Google case specifically.

The Google case primarily weighs the individual's privacy against the public interest, rather than a legitimate interest, which is more likely to be the relevant test for most companies where the data they're processing is not going to be published.

Fortunately, the A29WG also issued guidance in 2014 on the balancing test with legitimate interests, which can be found here:

http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2014/wp217_en.pdf

The ICO has also issued a range relevant guidance, including on the basis for processing:

https://ico.org.uk/for-organisations/guide-to-dataprotection/conditions-for-processing

As well as on the RtbF:

https://ico.org.uk/for-organisations/data-protection-reform/overview-of-the-gdpr/individuals-rights/the-right-to-erasure/

And most recently in their guidance on data protection and big data (see page 41):

https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf

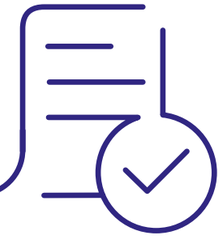

So what does the guidance say?

Both the A29WG and Google's RtbF guidance cannot be taken as direct guidance for those balancing a legitimate interest due to their focus on the public interest basis for processing.

The A29WG's legitimate interest guidance also requires a caveat, which is that the guidance is for the balancing test for general processing, not the balancing test for erasure, and so may differ. Whereas in the former there may be no consultation with an individual, in the latter, the data subject is actively objecting. This may mean that new information has been introduced which will affect the balancing, including the fact of the processing objection and erasure requests themselves.

The Google report makes some interesting points, such as highlighting the factors which should be considered in evaluating the impact on the data subject, and the information which the data subject should be asked to provide to enable the balancing test (although this may be something which changes under the GDPR due to the shift in the onus of justification).

The A29WG guidance on legitimate interests includes, amongst others, the following important points:

1. A legitimate interest must be clearly articulated, real and present. Speculative or vague interests are not valid.
2. Legitimate interests are potentially broad. Examples include conventional marketing, prevention of misuse of services and research (including for marketing purposes). Although the grounds may be broad, the processing allowed by any interest is not and must be strictly necessary for that interest.
3. If they appear equal, privacy rights generally trump legitimate interests. For legitimate interests to outweigh privacy rights, those privacy rights must clearly be more trivial. Evaluating this is the balancing test.
4. Factors to consider in the balancing test include:

> Whether the legitimate interest is also a public interest

> Quantity and level of invasiveness of data gathering

> Potential for adverse results on the data subject (e.g. damaging reputation, negotiating power or autonomy, exclusion, discrimination or defamation) and emotional impacts (e.g. irritation, fear and distress)

> Likelihood of risk materialising, and severity if it does

> The result of the balancing test can in some instances be changed by implementing appropriate mitigating safeguards, including pseudonymisation and other anonymisation techniques.

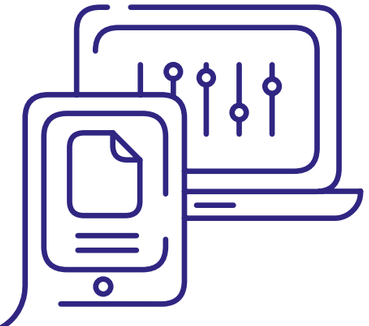

How should a CDO/CIO manage RtbF requests?

In modern data ecosystems identifying and deleting all the data relating to a specific individual might be difficult and have a high associated resource cost.

This is particularly trueof environments using software frameworks such as Hadoop,where the data about an individual may be widely dispersed.

As well as being challenging, deleting the data will reduce thevalue of the data set by removing information which may behelpful for analysis. Therefore, an organisation has an incentiveto preserve that data.

The RtbF stems from an individual's right to privacy. Addressing the privacy risk can be more beneficial than justdeleting data. The privacy risk posed by processing can bemitigated in a number of ways. Examples include: restricting access to the data, increasing data security through measures such as encryption at rest, and protecting privacy through pseudonymisation or anonymisation.

We believe that the best strategy is to take a privacy by design approach, so that all individuals have their privacy protected, and RtbF requests aren't unnecessarily arduous to comply with.

This process is reliant upon data being processed in a pseudonymised form, which is required where possible by the GDPR under Articles 25 and 32. Pseudonymisation ensures that individuals cannot be identified without additional information by replacing all direct identifiers (names, customer numbers, etc.) with pseudonyms. This substantially reduces the privacy risk for all members of the set. The relationship between the pseudonyms and the individual's identifying data is stored separately in what could be called a 'dictionary'.

If the dictionary is the only place where the pseudonymidentifier relationship is stored, then deleting the relationship entry for an individual will make it very difficult for the controller to be able to identify who the pseudonymised data relates to. Deleting the entry means the pseudonymised data can be left as it is, whilst the individual's privacy is respected, and their data is potentially anonymised, thereby taking it out of the scope of the GDPR and the RtbF.

We suggest considering the following process although would make it clear that this is not legal advice and should not be taken as such:

1. Establish the legal basis for processing and, if appropriate, carry out the balancing test, before you process the data to make sure your processing is legitimate. Pseudonymising may make processing, which would otherwise represent a risk to individual privacy, acceptable. This helps with the balancing test initially by reducing the risk posed to the individual, and aids compliance with Articles 25 and 32 of the GDPR.
2. Evaluate the RtbF request. Some RtbF requests may not be reasonable and will not need to be complied with, e.g. a customer requesting that their data be deleted, where the data is being processed in line with a contract which is still in effect.

3. Instead of deleting the pseudonymised data, delete the dictionary entry. Deleting the individual's entry in the file storing the relationship between the pseudonyms and the individuals makes it unlikely that they will be identifiable in the future by the controller, potentially meeting the RtbF requirement. It is important to note that whether or not this approach is appropriate will depend on factors such as the other variables in the data set, the strength of the legitimate interest and the potential privacy risk to the individual.

This approach reduces the potential burden of RtbF and may meet the requirement in one of two ways.

1. Anonymises the data. If there are no other quasiidentifiers in the data set which could likely or reasonably be used to re-identify the individual without the dictionary entry, then the data may be anonymous, and therefore out of the remit of the GDPR and RtbF requests.
2. 2. Sways the balancing test in favour of processing. If there still is a reasonable or likely way in which the individual could be re-identified, then the data will not be anonymous. However, when re-doing the balancing test, the fact that it is now much harder for the individual to be re-identified means that the data poses a significantly lower risk to the individual's privacy, meaning that the legitimate interest in processing may outweigh the reduced privacy risk, and therefore allow for the rejection of the request.

In many instances, this approach will not be sufficient to meet RtbF requirements. Whether this approach is sufficient or not, and why, will depend on the specifics of that data set and the processing, in particular what other variables are in the data set, what other relevant data exists, and who may have access to the data.

It is also worth noting that different supervisory authorities may have different views on what constitutes anonymised data and how interests and rights should be weighed in the balancing test. For international organisations it may be helpful to review their policy in light of the position of the local supervisory authority for each country they operate in.

RtbF Action plan – how to prepare

If you're trying to make a plan for how to deal with the RtbF, consider the following:

1. Map your data processing to understand what data you have and what you are using it for. Delete, and do not collect, any data you do not need.
2. Can your business accomplish its objectives using anonymised or pseudonymised data? If you can, do so, as this will both protect individual's privacy and give you more options for how you deal with RtbF requests.
3. What are your grounds for processing? Carry out the balancing test for any legitimate interests or public interests, mapping your position for possible scenarios.
4. Consider a range of potential objections and RtbF requests and carry out the balancing test to see if you think you would need to comply, before or after any further mitigations.
5. Use this thought exercise to draft a framework for how you might respond to different groups of requests, and then seek legal advice, to see if they agree with your framework. Google has a review board which looks at difficult cases, but for most requests they are able to match the request to a known type and therefore a known response, making it easier for them to manage the requests, and ensuring all applicants are treated the same way.

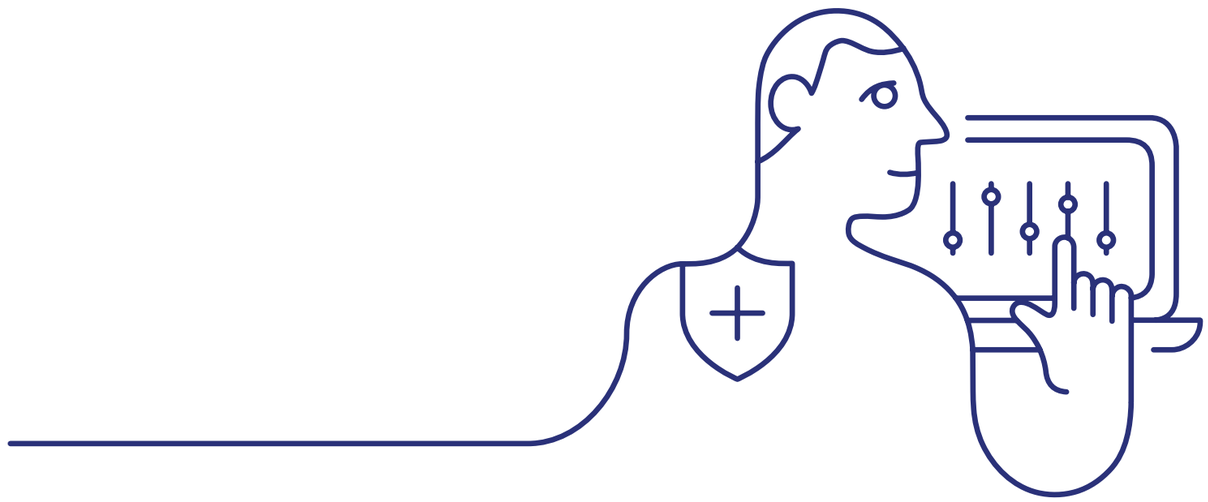

We're Privitar

We help organisations engineer privacy-preserving data operations, using advanced data privacy techniques that protect sensitive information while retaining data utility. Our software accelerates and automates privacy-safe data provisioning, helping our customers get more business value from their data, generate data-driven insights, and drive innovation.

Contact us:

e: info@privitar.com
t: +44 203 282 7136
w: www.privitar.com

PRIVITAR

@PrivitarGlobal

www.privitar.com