DAC BEACHCROFT

Perspectives on the Data, Privacy and Cyber Landscape

January
2026

Welcome to our Perspectives on the **Data, Privacy and Cyber Landscape 2026**

Change has arrived. The question now is; how to respond? In the year ahead, the data, privacy, and cybersecurity landscape will increasingly centre on the role of regulation in fostering innovation, while simultaneously addressing the need for data protection and growing cyber threats.

In the past 12 months, governments and regulators have been evaluating the benefits of economic growth and competitiveness against the traditional regulatory questions of privacy, safety, and accountability. The broader struggle between innovation and regulation has come into sharp focus, with geopolitical considerations illustrating the challenges for existing global frameworks such as the GDPR and the EU AI Act. Our keynote piece considers the longstanding dilemma faced by policymakers and regulators: the delicate balance to be achieved between innovation and regulation when presented with a transformative technology such as AI.

We also consider the relationship between technology and data, the 'power couple' of AI-related research, and whether existing regulatory frameworks are capable of accommodating this relationship. Against this backdrop, the piece also addresses the contractual tensions that can arise in AI-related research.

Unlocking the value of data is a crucial issue for organisations grappling with the intense pressure to utilise and monetise data. Effective data utilisation can drive competition, operational efficiencies, and the wider need to support AI adoption and training. The principles of anonymisation and pseudonymisation can support these ambitions and our team have considered these concepts.

As the UK data protection regulator, the Information Commissioner's Office is expected to give organisations clear guidance on best practices and when it can be expected to intervene and enforce in the face of non-compliance. Our team have considered recent criticisms levied at the ICO and whether upcoming structural changes to the regulator itself will change its approach to public and private sector enforcement.

Continuing the theme of regulation, the past year has seen significant legal and regulatory changes related to Data Subject Access Requests and data protection complaints. We review these developments and discuss the key takeaways for controllers looking ahead to 2026. We also examine the assessment of the risk of harm in a post-breach environment, reflecting on the legislative background and regulatory guidance, and their impact on the role of discretion in this assessment.

We invite you to explore these articles and help shape the conversation on data, privacy, and cybersecurity in the year ahead.

Jade Kowalski
Partner
Co-lead of DACB's Data, Privacy and Cyber Practice

To prioritise innovation or regulation?: Global tensions and pressures on regulatory models

01

Jade Kowalski
Partner
jkowalski@dacbeachcroft.com

Louis-Axel Batiste
Associate
labatiste@dacbeachcroft.com

María José Sánchez
Partner
msanchez@dacbeachcroft.com

Joshua Chan
Senior Associate
jchan@dacbeachcroft.com

Policymakers and regulators in many jurisdictions face a recurring dilemma: how to encourage innovation while safeguarding against emerging risks. For many years, regulation appeared to be the prevailing priority, with certain jurisdictions (notably, the EU) racing to set global standards for legal frameworks for governance of innovation. But the dial is shifting. As technological developments and geopolitics take centre stage, governments are reassessing their approach. Few issues reveal the strain between innovation and regulation as starkly as the ongoing debate over artificial intelligence.

Differing approaches to the regulation of AI amplify longstanding challenges and illustrate how difficult it has become for global governance frameworks to keep pace with transformative technologies. As governments weigh the benefits of economic growth, scientific progress, and global competitiveness against concerns about privacy, safety, and accountability, the broader struggle between innovation and regulation is coming into sharp focus, and the innovation versus regulation pendulum is swinging.

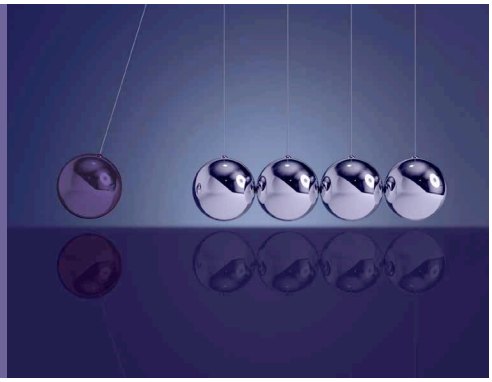

The development of contemporary regulatory data protection frameworks

Data protection in the UK and EU

The EU and UK have been at the forefront of the growth of data related legislation and regulation, with each stage of development reflecting advances in technology, from early computing to email and the internet, through to big data and artificial intelligence. Early laws in the 1980s and 1990s established key concepts such as fair and lawful processing, purpose limitation and data minimisation. However, it was the drafting of the EU General Data Protection Regulation (**GDPR**) which supercharged the data protection debate.

The final text was agreed in 2016. At that time, the recitals highlighted that *"Rapid technological developments and globalisation have brought new challenges for the protection of personal data"* and the GDPR was heralded as *"a strong and more coherent data protection framework… backed by strong enforcement"*. Taking effect from May 2018, the GDPR remains the foundation of data protection regimes in both the EU and the UK.

From global benchmark to over-regulated framework?

For many years, the GDPR was accepted as the global standard for data protection law. It is frequently cited as an example of the 'Brussels effect', being an EU legal instrument which serves as a standard for other jurisdictions to follow or is voluntarily adopted by companies in countries without similar regulations. For several years, a company could assert with confidence that being GDPR-compliant meant generally meeting regulatory standards on a global basis.

Many jurisdictions, particularly in Latin America, followed the EU's lead. A number of countries implemented their own data protection regimes, borrowing principles from the GDPR including Brazil in 2018, Mexico in 2021 and Chile in 2024, largely aligning their local regimes with that of the EU. Additional efforts are also underway in Argentina to update the existing data protection law to bring closer alignment to GDPR provisions, despite already being in receipt of an adequacy decision from the European Commission.

The GDPR also serves as a tool for trade, motivating other countries to raise the bar of local laws in order to ensure the free flow of personal data from the EU. Adequacy decisions recognising those jurisdictions offering an 'essentially equivalent' level of data protection to that provided by GDPR have been adopted in respect of geographically varied locations such as Japan, South Korea, Argentina and Uruguay. The European Commission also recently published a draft adequacy decision for Brazil.

However, there are key outliers on all sides; those with less developed data protection laws (for example, the United States) and those which have more stringent requirements (for example, data localisation requirements in China).

At opposite ends of the regulatory spectrum, both approaches can be understood as serving protectionist economic objectives, whether through a light touch regulatory approach or state control. In the short term, jurisdictions adopting a minimalistic approach to regulation may confer cost-based competitive advantages on domestic firms. However, we have seen historically that compliance with globally recognised frameworks, such as GDPR can, over time, confer significant advantages in terms of trust, market access and cross border operability. Whether these longer-term advantages continue to hold in an era of accelerating technological innovation remains an open and increasingly important question.

A new path for the UK

Against a backdrop of growing concern about overregulation hampering digital transformation activity in organisations, the UK has spent several years considering its place in the global ecosystem. Following Brexit, the UK has reassessed its economic position and its degree of divergence from the EU rulebook. In 2025, the government published its Regulation Action Plan focussing on (i) tackling the complexity and burden of regulation; (ii) reducing uncertainty across the regulatory system; and (iii) challenging and shifting excessive risk aversion in the system – a clear indication of the UK's positioning in the ongoing balance between innovation and regulation.

Throughout the DUAA legislative process, a central consideration was the desire for the UK to retain its adequacy decision issued by the European Commission. Earlier, subsequently abandoned, attempts in the form of the Data Protection and Digital Information Bills were criticised for including proposals which would potentially jeopardise the UK's adequacy status.

*In respect of data protection regulation, consideration of how and where to break away from certain GDPR requirements spanned multiple governments and draft bills and ultimately resulted in the Data (Use and Access) Act 2025 (**DUAA**).*

The DUAA introduces a number of amendments to the UK GDPR, the Data Protection Act 2018 and the Privacy and Electronic Communications Regulations 2003. In the words of the UK data protection regulator, the Information Commissioner's Office (**ICO**), the DUAA will *"promote innovation and economic growth and make things easier for organisations, whilst it still protects people and their rights".* In reality though, the reforms were relatively modest, and may represent something of a missed opportunity, particularly in light of recent developments in the EU.

Arguably, the UK has fallen victim to unfortunate timing and external circumstances. In November 2025, the European Commission introduced a significant package of measures known as the "Digital Omnibus" to amend its own data and privacy laws including the GDPR. Had this occurred 6-12 months prior, the UK Government might have considered more ambitious reforms (and of course may still consider doing so in the future).

Time for the EU pendulum to swing?

The EU is clearly now considering its position in the global economy and the impact of balancing regulation with innovation.

In 2024, the former European Central Bank president and former Prime Minister of Italy Mario Draghi issued a report on European competitiveness (known as the Draghi report). The report raised concerns that capacity to innovate within the EU is being *"hindered at every stage by inconsistent and restrictive regulations"* and urged reflection on the regulatory burden placed on companies across the entirety of the EU rulebook. By way of one example, it highlighted that *"there is no EU company with a market capitalisation over EUR 100 billion that has been set up from scratch in the last fifty years, while all six US companies with a valuation above EUR 1 trillion have been created in this period."* As a result, *"Member States are already acting individually and [protectionist] industrial policies are on the rise"* in the EU itself.

In response to these and related concerns, in November 2025 the European Commission published the 'Digital Omnibus' package of measures proposed to amend EU data and privacy laws with objectives including 'simplification', 'streamlining rules' and 'making it easier to do business'.

The proposed changes include:

- An updated definition of personal data, narrowing its scope.
- An expansion of 'legitimate interest' to allow it to be used as a legal basis for AI training and use.
- Amendments to the regime governing tracking technologies such as cookies.
- Allowing data controllers to refuse subject access requests where the request is considered to be an abuse of the rights conferred by GDPR.

Although these proposals are far from finalised, many consider that the essence of change is likely to be retained in any final provisions agreed by the European Parliament and Council.

AI as a microcosm for the regulation vs innovation divide

The tensions between innovation and regulation are particularly pronounced when it comes to AI.

European Union – amendments to requirements before they come into effect?

Unsurprisingly, given the intent to set a global standard once again, the EU won the race to establish the first AI specific legislative framework in the form of the EU AI Act. The Act touches almost everyone who handles or uses AI, dividing them into 'providers'; 'deployers'; 'importers'; and 'distributors' and has broad extra-territorial effect.

Utilising a risk-based approach, the relevant requirements depend on the risk of an AI system or use case. Whilst certain 'high-risk AI systems' can be used with appropriate safeguards, prohibited AI practices are banned due to posing an 'unacceptable risk'. Prohibitions on AI practices with unacceptable risks and the obligations for general-purpose AI models are already applicable, but wide-ranging provisions relating to high-risk AI systems are yet to apply. Could it be that they won't? Or at least not in their current form/for some time yet?

As part of the Digital Omnibus proposals, the Commission has put forward a number of simplification measures including the delay of the introduction of rules relating to high-risk AI systems until December 2027. Again, throughout the proposals there are numerous references to innovation and the need for implementation to be innovation-friendly. The proposals clearly demonstrate the EU's current perspective on the innovation versus regulation debate; efforts are underway to simplify measures before they are even implemented.

> *It can be argued that the motive for these measures is not limited to efforts to promote innovation. Geopolitical pressures have played a role.*

Prior to the publication of these proposals, the EU (and the UK) faced challenges from across the pond, both from US companies and government officials, over digital policy.

In particular, the EU AI Act has faced specific criticism due to the extraterritorial reach it wields. In late 2024, Senator Ted Cruz wrote to the then-Attorney General warning of *"heavy-handed regulation of U.S-developed internet technologies,"* and warning against *"artificial roadblocks"* which would prevent competition with China. President Trump also threatened trade action against the EU following the issue of a $2.95bn fine to Google by the European Commission over abusive practices in online advertising.

United Kingdom – the AI bill that never was

Although the UK's data protection framework is based on the same foundation as the EU, it has (thus far) opted for an independent approach to the regulation of AI. Currently, the UK is operating an expressly pro-innovation, principles-based approach. Adopted under the previous government, this approach is underpinned by five principles to inform the responsible development and use of AI in all sectors of the economy:

- **Safety, security and robustness**
- **Appropriate transparency and explainability**
- **Fairness**
- **Accountability and governance**
- **Contestability and redress**

These principles are supplemented by sector-specific guidelines from regulators such as the ICO, Financial Conduct Authority and the Competition and Markets Authority. Although the first King's Speech of the Labour Government in 2024 stated that the government would *"seek to establish the appropriate legislation to place requirements on those working to develop the most powerful artificial intelligence models",* no AI-specific legislation has been advanced by government to date. Reporting from mid-2025 indicates that any legislation will not be advanced until mid-2026 at the earliest, if at all.

In the meantime, the UK Government continues to place AI at the heart of its growth strategy, believing that the *"current pro-innovation approach to regulation is a source of strength relative to other more regulated jurisdictions".* The DUAA further underpinned the UK's existing approach by providing a clearer legal foundation for AI-driven data processing and model training. It is clear that, for the moment, innovation is taking precedence over regulation in the UK approach.

United States – a microcosm within a microcosm?

The current relationship between US state and federal bodies on the regulation of AI is itself a microcosm of the overall tension between innovation and regulation. An Executive Order was issued by President Trump shortly after his inauguration which ordered the removal of barriers to 'American Leadership in Artificial Intelligence'. The order explicitly revoked *"certain existing AI policies and directives that act as barriers to American AI innovation."*

Because of distinct state and federal regulations in the US, several states created their own AI laws. States such as Texas, California and Utah introduced laws intended to be effective in the first months of 2026, differing in their priorities, but sharing a common focus on transparency, accountability, and risk mitigation.

In response, in December 2025, President Trump issued a further Executive Order mandating a *"minimally burdensome national policy framework for AI"* and seeking to challenge inconsistent state laws, effectively blocking states from enforcing their own AI regulations. The issue of the Executive Order was not a surprise following the rejection of proposals for a Trump-supported 10 year federal moratorium on AI state regulation in the Senate earlier in 2025. At the federal level in the United States, the focus is clearly on the development of and innovation in artificial intelligence.

Further afield

Looking to other countries, to date, Singapore has adopted a similar approach to the regulation of AI to that of the UK. Although the country is committed to wider discussions about the development of AI, it has elected not to adopt a single AI law as of yet. This approach relies upon a series of frameworks and guidelines, tailored to specific sectors.

Whilst formal, binding legislation remains rare, several jurisdictions are clearly modelling binding frameworks on the EU AI Act, in a manner comparable to their adoption of domestic data protection frameworks using the GDPR as a guide. It is not yet clear whether those countries who have used the GDPR and AI Act as models will consider 'simplifications' of their own in the future. What is clear is that there is currently no uniform approach across Latin America to the regulation of AI.

Across Latin America, the regulatory approach to AI varies significantly from country to country, and cross-sectoral and foundational legislation akin to the EU AI Act remains limited. Peru recently approved the creation of a general framework for the use of AI in the country, one of the first of its kind in the region. Similar to the EU AI Act, the framework is structured into three tiers, with risks classified as 'misuse' being prohibited (e.g. manipulation and biometric surveillance), 'high-risk' (e.g. credit scoring and critical infrastructure) and 'acceptable risk'.

In Brazil, a wide-ranging bill to regulate AI is progressing, although with no clear timescales for implementation in sight yet. If passed, the Bill will adopt a similar tier system to that established by the EU AI Act, with 'excessive' systems being prohibited and 'high risk' systems being subject to heavy regulation. Chile's AI Bill, currently under legislative review, will create a similar classification model with four tier classifications. A draft Federal Law Regulating AI has also been introduced in Mexico, again progressing through the legislative process.

Other countries, including Argentina and Colombia, are taking a softer approach, being guided by a variety of measures such as national AI strategies and sector-led guidelines. Similar to the UK, these are underpinned by principles such as transparency.

The innovation and regulation cycle?

At the heart of the differences in approach lies the pivotal question, is the correct course fostering innovation, or establishing regulation?

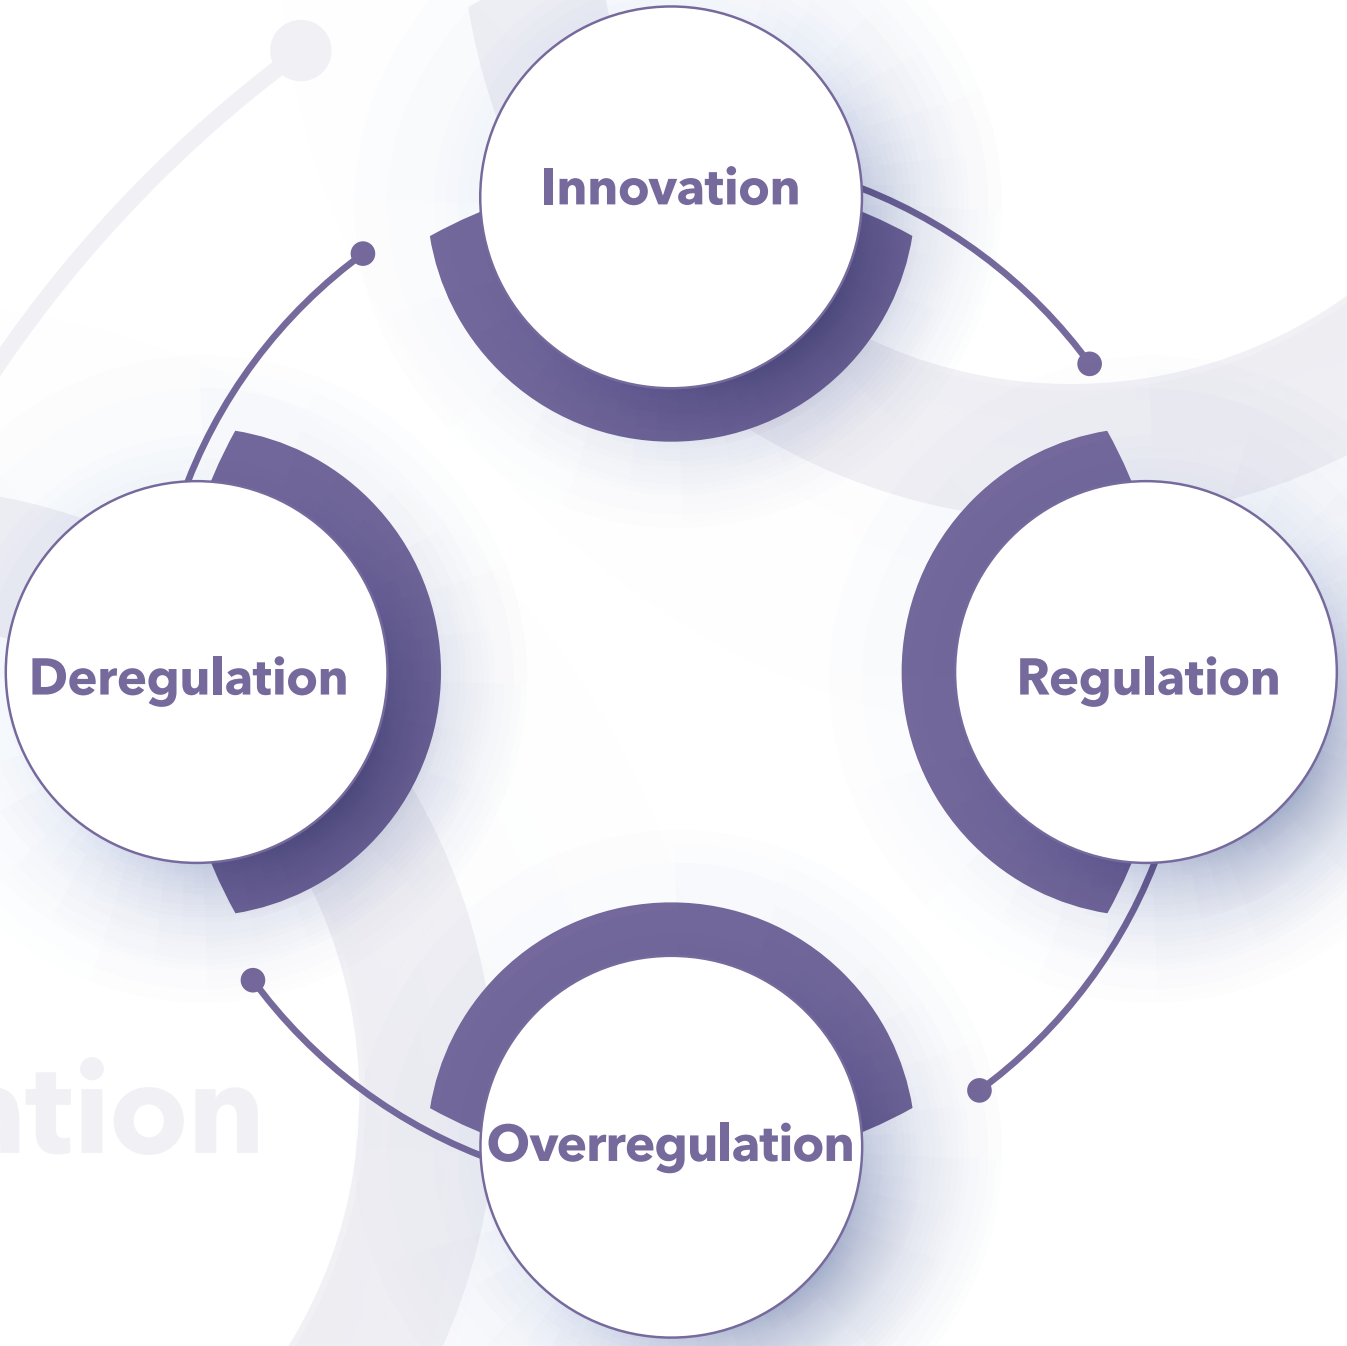

No single approach to the question of innovation and regulation will satisfy all stakeholders or protect against future developments. This is especially true in respect of technological advances.

As noted above, the GDPR was (and in some locations still is) considered to be the global standard for data protection frameworks. Recent years have seen a backlash to that perspective. The UK, and to a lesser degree the EU, have worked to develop more adaptable data protection rules in response to concerns regarding innovation and economic opportunity.

However, the primary focus of the tension between innovation and regulation crystalises when considering AI and is most starkly observed when comparing the US and the EU, with the US position clearly aimed at promoting innovation. Conversely, the EU, having opted for a risk-based legislative framework is already taking steps to 'simplify' measures not yet in force in order to foster innovation and economic growth. These two positions illustrate the regulation and innovation paradox; the pro-innovation jurisdiction remains in a state of flux due to a lack of regulation, and vice versa.

It is of course an oversimplification to frame the decisions faced by policymakers as choosing to head down one of two paths; one marked 'innovation' and another marked 'regulation'. There is a delicate balance to be achieved, with a number of multi-faceted considerations to consider.

So, what does this mean for organisations trying to advance AI and data-driven initiatives amid such uncertainty? In our view, a principles-based approach offers the most effective way to navigate the space between regulation and innovation. Core principles such as accountability, transparency, and fairness provide stable guidance even as more prescriptive aspects of data protection law continue to shift.

When it comes to AI specifically, this focus on principles over rigid rules creates the agility and adaptability organisations need as AI technologies and regulatory expectations continue to evolve.

In the end, a principles-based approach offers the flexibility to adapt, the clarity to guide, and the resilience to endure as technologies evolve. By focusing on outcomes rather than rigid requirements, businesses are empowered to build and use AI responsibly. The challenge now is not to predict every future development (or legislative solution), but to commit to the values that will shape trustworthy AI for decades to come.

Regulation

Unlocking the value of data:
Navigating anonymisation, pseudonymisation & PETs

02

Charlotte Halford
Partner
chalford@dacbeachcroft.com

Pete Given
Partner
pgiven@dacbeachcroft.com

In 2006, British mathematician Clive Humby stated that ***"data is the new oil"***. Twenty years later and this statement holds truer than ever. Organisations sit on a goldmine of personal data; from customer interactions and HR records to usage data from Internet of Things devices and digital platforms, personal data is at the core of how an organisation develops, markets and sells its products and services, and operates its business.

"Personal data powers our economy, from retail to hospitality to healthcare. Unlocking the potential of this data is key to encouraging economic growth and investment – as long as the public can trust it will be appropriately protected."

UK Information Commissioner, March 2025

There is an intense pressure to utilise and monetise data, driven largely by a desire to:

- **Have a competitive advantage:** data-driven insights and predictive analysis fuel smarter and faster decisions.
- **Implement operational efficiencies:** data can be used to streamline and automate processes.
- **Support AI adoption and training:** large, diverse datasets are essential for developing and refining AI models, enabling organisations to innovate responsibly and competitively.
- **Develop new products and services:** analysing data can lead to more desirable and personalised offerings.
- **Comply with regulatory incentives and mandates:** data-sharing mandates are increasingly being legislated for in sectors such as health, energy and finance.
- **Conduct evidence-based decision-making:** insights from analytical data drive stronger risk management, benchmarking and reporting.

However, the benefits of data utility come with an inherent challenge: data privacy compliance.

Anonymisation: unlocking value safely

Anonymisation is a method by which an organisation can harness the potential of data, but in a privacy-friendly way.

What is anonymisation?

Anonymisation is the process of turning personal data into anonymous information so that the individual is no longer identifiable. This limits the risk to individuals and can allow organisations to share the information more freely with other organisations or the public.

Truly anonymised data falls outside the scope of data protection laws. Consequently, it is important to ensure that where an organisation intends to rely upon anonymisation, it is correctly anonymising the data.

Organisations should consider the nature of the data; the purpose(s) it is used for; and the surrounding context. Key indicators of whether information is personal data include 'singling out' and 'linkability', as further examined in this piece.

What does the ICO say about anonymisation?

According to new anonymisation guidance issued by the Information Commissioner's Office (**ICO**), data is only considered anonymised if no individual can be identified by any party using means that are reasonably likely to be employed.

This test is contextual and risk-based, focusing on:

- **Means of reidentification:** The anonymisation must account for *"all the means reasonably likely to be used"* by any party, including the controller.
- **Nature of controls:** Technical safeguards such as access controls (e.g., firewalls, user permissions) may mitigate risk but do not constitute anonymisation if the controller retains the ability to reidentify individuals.
- **Context and role of the data recipient:** The likelihood of reidentification is assessed in light of the recipient's role, knowledge, and access to other data or tools.

The ICO's guidance sets out the two primary anonymisation techniques:

- **Generalisation:** which reduces the specificity of the data. This changes information that may identify someone so that it relates to multiple people. This means members of that group can't be identified or are no longer identifiable (for example, listing age ranges rather specific ages).
- **Randomisation:** which can be used to reduce the certainty that a record relates to a particular person. This changes information that may identify someone so that it cannot be definitively attributed to one person (for example, noise injection).

In addition, the ICO references the following techniques as supporting the above anonymisation techniques, but which will not themselves constitute anonymisation:

- Masking can reduce identifiability by deleting or supressing certain values or data records. While masking can be effective when used alongside generalisation and suppression, it is not considered an anonymisation technique on its own.
- Suppression is a disclosure control process where parts of the data are made unavailable to the user. The term is usually used to describe approaches like cell suppression, the removal of outliers and local suppression of particular values within microdata records.

The ICO emphasises that identifiability is not binary but exists on a spectrum.

> *At one end, data is fully identifiable; at the other, it is effectively anonymised. Between these extremes lies a grey area where re-identification risk depends on context, available resources, and technical measures.*

As such, when determining whether data is anonymous, organisations must assess the risk of:

- **Singling Out:** Can an individual be isolated within a dataset?
- **Linkability:** Could separate datasets be combined to reveal identity?
- **Inference Risk:** Could additional attributes allow identification indirectly?

This spectrum approach means anonymisation is rarely absolute and the risk of identification must be evaluated on a case by case basis, considering who holds the data and what auxiliary information they might access.

To operationalise this assessment, the ICO proposes the 'Motivated Intruder Test'.

> *ICO's "Motivated Intruder" test: Could a reasonably competent person, who has access to publicly available resources, use investigative practices to re-identify individuals?*

If the answer to the above question is yes, the data is not truly anonymised. This test reflects real-world conditions, where attackers may exploit open-source data, social media, or public records. It's a practical benchmark for organisations to gauge whether anonymisation techniques withstand plausible attempts at re-identification.

What does the case law say?

Two landmark cases illustrate the evolving interpretation of 'personal data' and identifiability under EU law:

- **SRB v EDPS (CJEU, 2025):** This case considered whether pseudonymised comments shared by the Single Resolution Board with Deloitte were 'personal data.' The Court clarified that pseudonymised data is not automatically personal data for every recipient; identifiability must be assessed contextually and from the recipient's perspective, considering whether the recipient has *"reasonable means"* to re-identify individuals. This is a pragmatic approach from the CJEU which requires assessment on a case-by-case basis. If organisations can evidence genuinely minimal re-identification risks, a contextual approach to determining whether pseudonymised data is personal data could enhance innovation.
- **Breyer v Germany (CJEU, 2016):** This case considered whether dynamic IP addresses could constitute personal data for a website operator if the operator has legal means to obtain additional information from an internet service provider to identify the user. The Court considered the *"means reasonably likely to be used"* test and concluded that they could. The test applied is the same concept that underpins the ICO's anonymisation guidance. It highlights that identifiability is not absolute but depends on practical and legal access to linking information.

Navigating the legal and contractual hurdles for anonymisation

Notwithstanding the seeming pragmatism of regulators and the courts, in the increasingly digital and data-driven world in which we live, effective anonymisation remains difficult to achieve in practice.

Where data that appears effectively 'anonymised' can be linked back to the underlying data subjects by combining it with other information which could enable reidentification, it will not meet the high standard for anonymisation required to fall outside data protection law. This might be the case where, for example, the organisation retains supplementary information, such as the original dataset or auxiliary data, or has access to additional data, such as a publicly available dataset. This principle was reinforced in *Breyer v Germany,* where dynamic IP addresses were deemed personal data because the website operator could legally access additional information to identify users. In practice, if an organisation retains linkage keys or supplementary datasets, the data remains pseudonymised and therefore subject to data protection obligations.

As highlighted in *SRB v EDPS*, identifiability is contextual: data may be personal in one organisation's hands but effectively anonymous in another's. This opens strategic opportunities. For example, creating a separate, dedicated entity for analytics, operating on segregated systems and dedicated employees without access to original identifiers, can transform pseudonymised data into anonymous data in that entity's hands, potentially even within the same group. This approach mitigates regulatory risk and unlocks data utility.

However, even when organisations have the technical capability to anonymise or pseudonymise data, they often face legal and contractual constraints. It is worth noting that the act of anonymisation itself, even if it produces effectively anonymised data, is an act of processing, which an entity may not be entitled to do, whether under contract or from a regulatory perspective.

A common scenario arises where an entity gains access to a dataset under a client contract that strictly limits its use to delivering services. These agreements may prohibit creating derivative datasets for internal purposes. Similarly, if the entity acts as a processor, it cannot lawfully repurpose the data without meeting controller obligations, such as establishing a lawful basis or complying with transparency requirements. This dual challenge means that, in practice, organisations may be unable to leverage data for innovation despite having the tools to anonymise it.

One solution is to address these hurdles at the contracting stage. Organisations can negotiate provisions that allow the creation of a dataset which, once the original data is no longer accessible (for example, at the end of the engagement), would be considered anonymous and therefore fall outside data protection law. This dataset could then be used for the organisation's own purposes, such as analytics, service optimisation, or product development. While residual issues remain, such as ensuring the anonymisation is robust and overcoming reluctance from data providers, this approach offers a pragmatic pathway to unlock data utility without breaching regulatory or contractual obligations.

Beyond anonymisation: strategic pathways to unlock data value

Even where full anonymisation proves elusive, organisations are not without options; there are still ways to extract value from the data they hold.

Pseudonymisation: a bridge between privacy and utility

When full anonymisation is not achievable, pseudonymisation offers a pragmatic alternative. It involves replacing or transforming identifiers so that data cannot be attributed to a specific individual without additional information, which is typically held separately.

Common techniques include hashing, encryption, and tokenisation. Unlike anonymisation, pseudonymisation preserves much of the dataset's richness, enabling advanced analytics, AI training, and cross-functional insights while reducing direct identifiability.

However, pseudonymisation is not a silver bullet. Data protection laws still apply because the possibility of re-identification remains if linkage keys or auxiliary datasets exist. Organisations must assess the robustness of their approach against realistic attack scenarios, such as exhaustive searches or dictionary attacks. The ICO's guidance on pseudonymisation stresses that pseudonymisation should not be treated as anonymisation; it is a risk-reduction measure, not an exemption from compliance.

Despite these constraints, pseudonymisation can unlock significant value. It enables organisations to share data internally or with partners under controlled conditions, supporting innovation without exposing raw identifiers. For example, pseudonymised datasets can power predictive models, benchmarking, and service optimisation while maintaining a privacy-conscious posture. Combined with strong governance, such as segregating linkage keys, implementing contractual safeguards, and conducting Data Protection Impact Assessments (**DPIAs**), pseudonymisation becomes a strategic enabler.

In a regulatory environment increasingly focused on 'privacy by design,' organisations that master pseudonymisation can position themselves ahead of competitors, balancing compliance with data-driven growth.

Legislative reform: a new hope for data utility?

If anonymisation and pseudonymisation alone are not the panacea, could legislative reform offer a lifeline for organisations seeking to unlock data value? Recent policy developments suggest that lawmakers are beginning to recognise the tension between privacy protection and innovation and are taking steps to recalibrate the balance.

The UK's Data (Use and Access) Act 2025 (**DUAA**) is a prime example. By introducing broad consent provisions for scientific and statistical research, it reduces friction for organisations that want to share and analyse data for socially beneficial purposes. The DUAA also streamlines compliance by relaxing certain privacy notice obligations, signalling a shift toward enabling responsible data use rather than simply restricting it. For organisations operating in research-heavy sectors - healthcare, energy and financial services - this could be transformative, provided they embed robust governance to maintain trust.

Across the Channel, the EU's Digital Omnibus Package, published in November 2025, proposes amendments to the EU GDPR that could redefine the status of pseudonymised data. Under the draft, pseudonymised datasets may, in certain contexts, fall outside the definition of personal data, potentially unlocking new opportunities for cross-border collaboration and AI development. While implementation remains uncertain, the direction of travel is clear: regulators are exploring pragmatic solutions to reconcile privacy with innovation.

> *"Pseudonymisation is a way of reducing risk and improving security. It is not a way of transforming personal data to the extent the law no longer applies."*
>
> ICO Guidance on Pseudonymisation

Privacy-enhancing technologies: unlocking value without compromise

When anonymisation and pseudonymisation reach their limits, organisations are increasingly turning to Privacy-Enhancing Technologies (**PETs**) as a way to reconcile two competing priorities: extracting insight from data and safeguarding individual privacy. PETs represent a shift from traditional compliance tools to advanced technical solutions that enable collaboration, analytics, and AI development without exposing raw personal data. These technologies are not just about compliance, they are about embedding trust into data strategies and enabling responsible growth in a data-driven economy.

PETs encompass a range of techniques designed to minimise personal data use, maximise information security to preserve privacy, and empower individuals. According to the ICO's guidance, PETs can:

- **Reduce the identifiability of individuals:** through synthetic data generation and differential privacy.
- **Hide or shield information:** using zero-knowledge proofs, homomorphic encryption and trusted execution environments.
- **Split datasets for secure computation:** through secure multi-party computation and federated learning.

The benefits are significant. PETs help organisations demonstrate data protection by design and default; comply with data minimisation principles and maintain robust security whilst enabling analytics and AI training. They can also support anonymisation or pseudonymisation strategies, control access to sensitive datasets; and reduce breach risks. For AI-driven initiatives, PETs are particularly valuable, allowing models to learn from distributed or sensitive data without compromising privacy.

However, PETs are not without limitations. Most PETs still involve processing personal data, meaning organisations must ensure lawful, fair, and transparent handling.

In addition, they face practical challenges such as scalability, lack of mature standards, vulnerability to sophisticated attacks, and implementation errors. Expertise gaps and insufficient organisational measures can undermine effectiveness.

To mitigate these risks, organisations should conduct DPIAs and integrate PETs into broader governance frameworks. The strategic question is not whether PETs will become mainstream, but how quickly organisations can adopt them to turn privacy into a competitive advantage.

The ICO's guidance is clear: *"the purpose of many PETs is to enhance privacy and protect the personal data you process, rather than to anonymise that data. This means that:*

- many PET use-cases still involve personal data; and
- when you deploy such techniques, you still need to meet your data protection obligations".

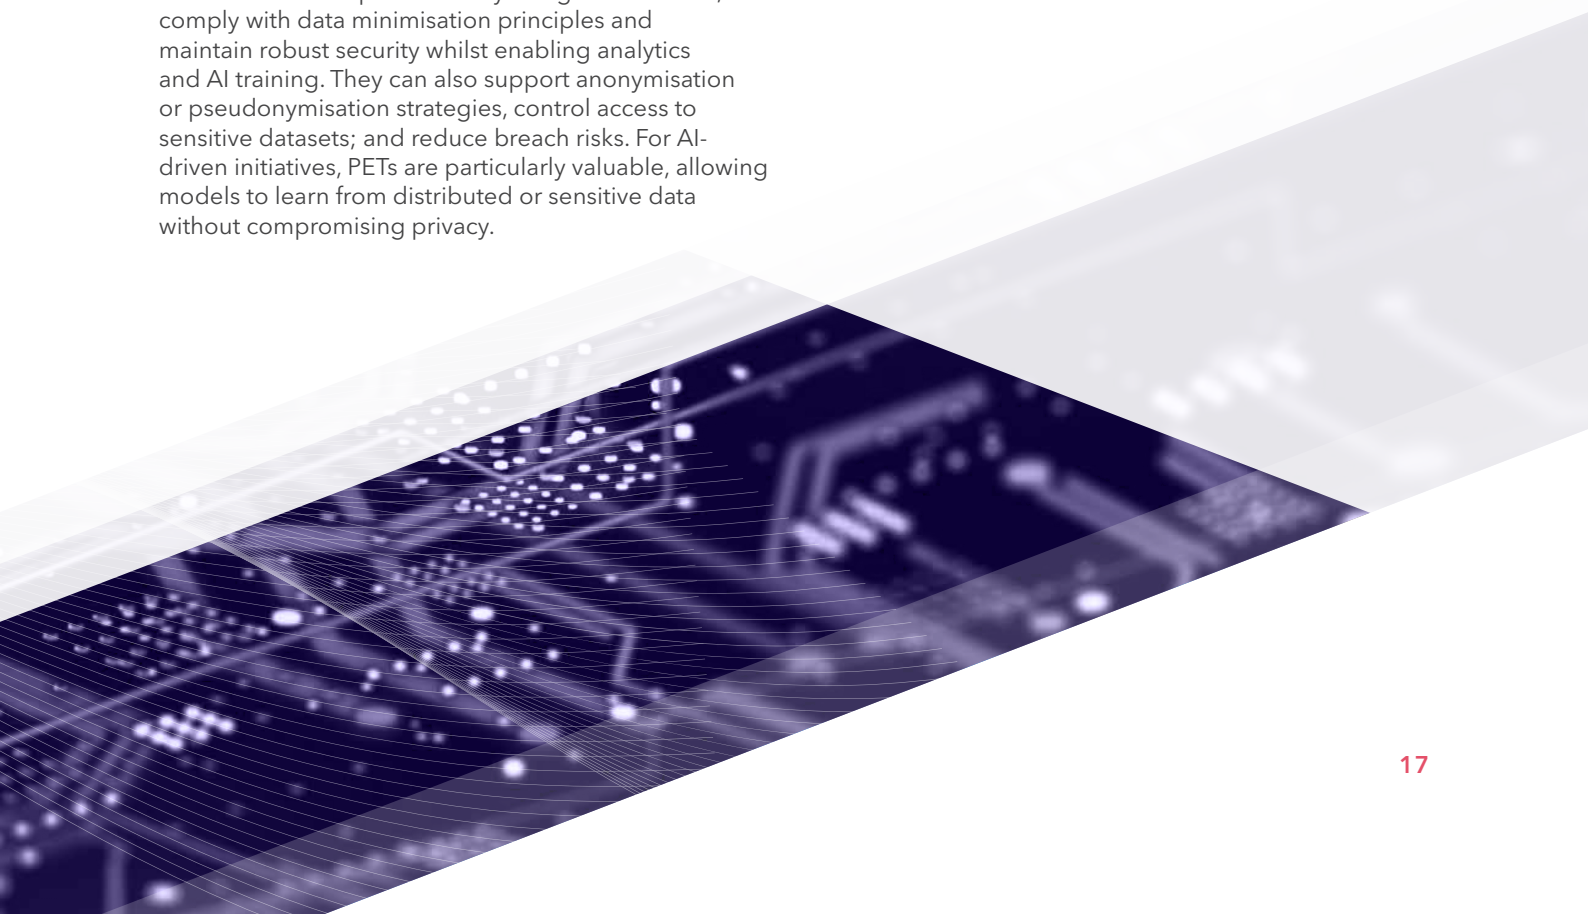

Synthetic data

Synthetic data is a particular PET that we envisage will see increased uptake with the rise in creation and deployment of AI tools. Synthetic data is 'artificial' data generated by data synthesis algorithms. There are two types of synthetic data:

- **'Partially' synthesised data:** where only some variables of the original data are synthesised (for example, just synthesising location in an A&E admission dataset).
- **'Fully' synthesised data:** where all variables of the original dataset are synthesised (for example, synthesising name, location, admission time and reason for admission in an A&E admission dataset).

Synthetic data is a useful tool for training artificial intelligence models in environments where real data is scarce or sensitive. For example, Mastercard found that using synthetic data gave them a competitive advantage when performing analysis in emerging markets. Synthetic data benefits from preserving statistical integrity, eliminating directly identifiable data, and being cost-efficient. The ICO advises that organisations should consider using synthetic data as a tool for generating non-personal data in situations where it does not need to, or cannot, share personal data.

The challenge for organisations is whether the synthetic data used is an accurate substitute for the original data. Quality controls should be put in place to ensure that poorly generated synthetic data does not skew models and datasets (for example, through biases).

Practical measures for organisations to unlock data value

Turning data into a strategic asset requires more than technical capability, it demands foresight, governance, and contractual clarity. To unlock value effectively and compliantly, organisations should:

- **Define a data vision:** Go beyond short-term needs. Map current and future data requirements, identify sources, and assess whether internal or external datasets can deliver competitive advantage. This is the foundation for any data-driven strategy.
- **Embed governance by design:** Implement robust frameworks that integrate privacy, security, and ethical considerations from the outset. Use DPIAs not as a tick-box exercise, but as a strategic tool to evaluate risk and opportunity.
- **Master advanced techniques:** Understand the spectrum of options, anonymisation, pseudonymisation, and PETs, and when each is appropriate. Keep pace with evolving case law and legislative reforms that may redefine what counts as personal data.
- **Engineer contractual flexibility:** Contracts often dictate what you can and cannot do with data. Negotiate provisions that allow for anonymisation or pseudonymisation for analytics and innovation and consider future-proofing agreements to accommodate emerging technologies and regulatory changes.
- **Invest in capability and culture:** Technology alone won't deliver value. Build internal expertise in data ethics, privacy engineering, and AI governance. Foster a culture where compliance is seen not as a constraint but as a catalyst for trust and innovation.

Concluding thoughts

Unlocking the value of data requires careful consideration of the legal, regulatory and contractual controls that may attach to that data. Anonymisation, pseudonymisation and PETs can all be used to unlock the value of data in a manner that complies with those controls, but there is no one size fits all approach and each has its advantages and disadvantages.

Organisations that adopt robust governance and contractual safeguards will be best positioned to harness data responsibly and competitively in an evolving regulatory landscape.

Data Utility v Privacy Risk

Technology and Data:
Analysing the relationship between the power couple of AI-related research

03

Darryn Hale
Partner
dahale@dacbeachcroft.com

Christopher Air
Partner
cair@dacbeachcroft.com

Modern society loves power couples – from Posh and Becks through to Will and Kate, we are obsessed with glamorous celebrity couples who are collectively more successful than the sum of their parts and who are indispensable to each other's success, fame and happiness. Technology and data are the power couple of AI-related research – never out of the spotlight, the subject of endless discussion and often the source of controversy (but no paparazzi thankfully).

From training machine learning models to refining algorithms for natural language processing and computer vision, data plays a hugely influential role in determining the performance and reliability of AI applications. In research contexts, data is not merely a passive input; it is actively curated, pre-processed, and analysed to uncover patterns that drive AI innovation. However, the relationship between the commercial drivers for undertaking AI related research, and the legal limitations around using data for such purposes, have something of an uncomfortable history under UK law (not quite Meghan and Harry but you get the idea!). Given the clear desire for the nation to be at the forefront of AI discovery, the question arises as to whether our current legal framework will be made more flexible and accommodating for conducting AI research or whether counter pressures such as needing to maintain adequacy in the eyes of the European Commission for unhindered data flows with the EU result in us leaving UK data protection law largely unchanged in this respect? We therefore reflect on recent changes to UK data protection law made by the Data (Use and Access) Act 2025 (**DUAA**), as those changes specifically relate to research.

We have considered the legal environment in which our power couple operate, providing an overview of some of the legal and ethical challenges relating to use of data, particularly personal data, for AI related research, and explain the current legal position under UK data protection law in this context, which includes a look at the accountability framework of the Information Commissioner's Office (**ICO**).

We have also reflected on sources of flare ups and spats which our couple face daily; tensions that can arise between contracting parties as to their roles as either joint controllers, controllers and/or processors and wider concerns regarding supply chain risks. We then conclude by setting out some recommendations for managing such contractual tensions in practice – a guidebook for our power couple to live happily ever after!

AI-related research – the Legal and Ethical Challenges

Let's start by looking briefly at a few examples of some of the legal and ethical challenges posed by use of data for AI research.

- **Privacy:** AI systems often rely on large datasets containing personal data and even special category data. The party designated as the data controller will need to ensure that such processing fully complies with the UK General Data Protection Regulation (**UK GDPR**), Data Protection Act 2018 (**DPA 2018**) and relevant ICO guidance. Of course, fully anonymised data falls outside the remit of such legislation and we explore the concept of anonymisation in this collection.
- **Bias and Fairness:** If datasets are non-representative or skewed, AI models can perpetuate or amplify social biases, leading to discriminatory outcomes. This is particularly problematic in areas like hiring, lending, or law enforcement, where biased algorithms can harm marginalised groups.
- **Transparency and Accountability:** Many AI systems operate as 'black boxes,' making it difficult to explain decisions. Lack of transparency undermines trust and accountability, especially when AI influences critical decisions in healthcare, finance, or criminal justice for instance.
- **Data Ownership and Commodification:** Increasingly, personal data is treated as a commodity, raising ethical questions about ownership and exploitation. Individuals often have little control over how their data is monetised or shared.

A UK data protection law perspective

The UK GDPR and DPA 2018 set out specific provisions relating to the use of personal data for research purposes, breaking research down into the following types: (i) archiving purposes in the public interest; (ii) scientific or historical research purposes; and (iii) statistical purposes. Helpfully, ICO guidance explicitly lists the development of AI as falling within the scope of what constitutes 'research'. Furthermore, the DUAA has also introduced a new statutory definition of "scientific research" that covers any research reasonably described as scientific, whether publicly or privately funded, and whether commercial or non-commercial. The research provisions within the UK GDPR and DPA 2018 are scattered across several far flung parts of the legislation and require a reader to cross refer to, and understand the interplay between, data protection principals in Article 5 UK GDPR, data subject rights under Chapter 3 UK GDPR and safeguards under Article 89 UK GDPR, as well as exemptions under Schedule 2, DPA 2018. Understandably therefore, the research provisions are notoriously difficult to understand and apply in practice (at least with a sufficient degree of certainty).

Below we summarise the key provisions as they currently stand and look forward to how these are set to change by virtue of the DUAA.

- **Lawful Basis for Processing:** Research involving personal data must have a lawful basis under Article 6 UK GDPR, the most common being "public task", "legitimate interests" or "consent". Moreover, as data used for research is often special category data (e.g. information relating to health) a condition for processing is also required under Article 9 UK GDPR or Schedule 1, DPA 2018. One such condition is research related purposes under Article 9 UK GDPR, which is expanded upon under Schedule 1, DPA 2018. It restricts use of the data for measures or decisions about particular people, except for approved medical research; and requires that the processing be in the public interest. Furthermore, the DUAA introduced amendments allowing "broad consent" for processing personal data for scientific research purposes. This means consent remains valid even if specific research purposes cannot be fully identified at the time of collection, provided ethical standards are followed and participants can consent to parts of the research where possible.
- **Purpose Limitation & Compatibility:** Usually, personal data must only be processed for the original purpose for which it was collected. However, the UK GDPR allows further processing for scientific or historical research or statistical purposes if appropriate safeguards are in place in accordance with Article 89. These safeguards focus in particular on measures to ensure respect for the principle of data minimisation, including where possible, anonymising or pseudonymising data. The DUAA introduced some welcome flexibility around the purpose limitation principle, enabling reuse of personal data for research without requiring a new lawful basis, as long as safeguards are maintained. This is particularly helpful for secondary research using existing datasets.
- **Exemptions:** Schedules 2–4, DPA 2018 set out exemptions which apply to the use of personal data for research purposes. Such research related uses are exempt from certain data subject rights (e.g. access, erasure, objection) if applying those rights would seriously impair or prevent the research, and the required Article 89 UK GDPR safeguards are in place.
- **Storage Limitation:** Normally, personal data should not be kept longer than necessary. However, for research, archiving, or statistical purposes, data can be retained indefinitely if the Article 89 UK GDPR safeguards are in place.
- **Transparency:** Despite the challenges around explainability of AI systems and their outputs, transparency remains an important principle under Article 5 UK GDPR, so organisations will need to provide information to data subjects, informing them that their data will be used in AI related research and the resulting decisions. Indeed, the ICO has published specific guidance on this topic. However, the DUAA disapplies certain transparency requirements. Ordinarily, if using data for a purpose other than the original purpose, the controller needs to inform affected data subjects. However, under the DUAA this does not apply to scientific research if doing so is impossible or would involve disproportionate effort.

The DUAA therefore goes some way towards relaxing the rules around research. In terms of how an organisation complies with, and demonstrates its compliance with, these legal requirements, specifically in the context of AI research, it is useful to turn to the ICO's accountability framework, which we explore next.

A regulatory perspective

The ICO has produced an extensive accountability framework which is supplemented by specific guidance on the accountability and governance implications of AI. In essence, accountability is a matter of actively demonstrating and evidencing how the above legal compliance requirements have been met. Neither the framework nor the guidance necessarily focuses in on the implications of AI-related research, but nonetheless the principles set out are of sufficiently broad scope and effect to be relevant. The documents are extensive; headline points include:

1. A Data Protection Impact Assessment DPIA is required for any processing of personal data which involves the use of new technologies including AI, machine learning and deep learning.
2. There is a tension between data minimisation and statistical accuracy which, on the face of it, can be difficult to reconcile. In particular, the fairness principle under the UK GDPR requires any AI to be statistically accurate but at the same time the AI needs lots of data to properly train which potentially conflicts with data minimisation requirements. Ultimately this comes down to technical analysis of the minimum dataset required to achieve sufficient accuracy. This may need to be kept under review, particularly once AI moves beyond research into live deployment (i.e. checking whether the outputs remain within expected accuracy levels).
3. Careful consideration needs to be given to the nature of any training data used to research or develop AI, both to ensure that it is not imbalanced but also to correct for any past discrimination inherent in the dataset. This is to ensure compliance both with the fairness requirement under the UK GDPR but also broader anti-discrimination laws and may include steps such as adding or removing data relating to under/over-represented groups from the dataset.

In addition, the ICO has also recently focussed on AI supply chain contract management, including that related to generative AI technology (which is often more nuanced). It recently conducted a five-part consultation to clarify how UK data protection law applies to generative AI. The series addressed five key areas, namely: (i) lawful basis for web scraping; (ii) purpose limitation; (iii) accuracy of training data and outputs; (iv) engineering individual rights into AI models; and (v) allocating controllership across the AI supply chain. One point of particular interest is that while the consultation itself does not indicate or settle a formal legal position, it does challenge the traditional position of AI developers only ever being processors. This is a point we consider further below in the context of the inevitable constraints which arise from a controller/processor relationship being adopted.

Contractual tensions

Research is one of the areas in which we see most variation in approaches to contracting and in particular the assessment and attribution of roles under the UK GDPR i.e. whether a party is acting as a controller, joint controller or processor. This, to some extent, reflects the different ways in which AI-related research may be conducted: for instance, it could be a standalone research project in which one party sponsors another to gather data on their behalf in relation to a specific hypothesis or issue. It could also be a collaborative arrangement in which several parties agree to pool their data in order to enrich the datasets they individually hold and they conduct research on that pooled data. Finally, research could also be an adjunct to the delivery of services – for instance, a tech supplier delivering a particular product or service to a customer and wanting to use the data they collect when delivering those services in order to research potential new products or services.

As most will know, the question of designation under the UK GDPR is a factual one i.e. looking at who does what with the data and then assessing whether that is indicative of a controller, joint controller or processor role. In our experience, however, that is often not interrogated as robustly as it ought to be and instead habit creeps in – in the health sector, for example, a supplier of a tech-based solution or platform is usually deemed a processor by default because this is felt to be a way of safeguarding highly sensitive patient data. This then, inevitably, leads to a very tightly restricted set of instructions to the supplier in the contract which, realistically, stifle any possibility for them to use the data for other innovative purposes (notably research).

It is not only controller-processor contacts which are prone to contractual constraints, whether express or implied, as controller-controller and joint controller arrangements are often so sparse with detail that it is completely unclear as to what is deemed legitimate use of data under the contract.

We consider some of the common dynamics and tensions commonly encountered in the context of the above scenarios below.

1. Controller to processor

This scenario most commonly arises where one party is providing services to another, which may themselves be based on AI. Equally they may not but the data processed in order to deliver the services is sufficiently interesting to mean that it could be used to research and develop AI in the future. In turn, one of two things will usually happen: first, the processor is expressly prohibited from using personal data processed on the controller's behalf for any purposes other than delivery or second, the possibility of research is not contemplated when the contract is agreed and so is not addressed one way or another.

Either scenario would, effectively, constrain the possibility of using the relevant data to research and develop AI. It can also be further compounded, in our experience, by tech suppliers contracting with customers on the terms they supply and ending up with hugely variable approaches to data protection compliance (and which, self-evidently, is very difficult to unpick).

2. Separate controller to controller

In a separate controller to controller relationship, the picture is arguably somewhat more straightforward than with controller to processor, but not always. Essentially, each party bears its own responsibility for UK GDPR compliance, and either agrees to share data for different/shared purposes or simply does their own thing. The degree of co-operation or assistance between the parties to aid the other's compliance is variable. Either way, responsibility for use of data for aspects such as AI research is usually borne by each party.

In principle, therefore, this is a more favourable scenario for the purpose of enabling AI-related research. However, the common mistake is often that the data sharing provisions between the controllers are so vague or, at the other end of the spectrum, so tight that it is unclear whether use of shared personal data for research into AI was contemplated.

3. Joint controller scenario

This applies where there is collaboration between two or more parties, often with pooling of data. Article 26 UK GDPR sets out high level requirements for this scenario, but these aren't prescriptive. In practice, this means agreeing which party is responsible for providing the fair processing notice and handling data subject rights requests, and which acts as point of contact for requests from data subjects.

Concluding thoughts

Considering the above, our take from a contractual perspective is that in order to enable AI-related research (i.e. our recipe for a happy, harmonious and lasting power couple relationship) it is imperative to:

(i) consider whether the purpose can be achieved through the use of anonymised data;

(ii) properly analyse roles under data protection law and in particular whether a tech supplier is genuinely always a data processor;

(iii) develop consistent contractual template terms which either reflect a controller designation or, where a processor, contain authorisation from the controller to enable research processing; and

(iv) have honest up front discussions before agreeing contracts about the potential AI-related research use cases.

From DSARs to data protection complaints:

Implementing lessons from 2025

04

Rebecca Morgan
Legal Director
rebeccamorgan@dacbeachcroft.com

Amanda MacKenzie
Associate
ammackenzie@dacbeachcroft.com

Heasha Wijesuriya
Solicitor (Australian Qualified)
hwijesuriya@dacbeachcroft.com

Kate Galloway
Partner
kgalloway@dacbeachcroft.com

Data Subject Access Requests (**DSARs**) remained in the spotlight in 2025. Indeed, developments picked up pace; we saw a number of notable legal and regulatory interventions and trends in technology and public sentiment impact upon the volume and complexity of such requests. What are the lessons we can take into 2026? And how will those lessons serve to support controllers manage data protection complaints in a year when the new statutory complaints handling requirements will come into effect? We examine these issues.

The evolving DSAR legal and regulatory landscape

2025 started with the High Court ruling in *Ashley v HMRC*, notable because recent DSAR related case law is sparse.

This case arose from His Majesty's Revenue and Customs (**HMRC**) challenge to Mr Ashley's property valuations following a 2012 sale, which led to a £13.6m tax bill. Although the tax bill was later withdrawn, Mr Ashley exercised his right under Article 15 UK General Data Protection Regulation (**UK GDPR**) to access all personal data held by HMRC relating to its tax enquiry. HMRC adopted a narrow approach to the DSAR, limiting searches to one department and withholding most data under the "tax" and "legal privilege" exemptions set out in Schedule 2, Data Protection Act 2018 (**DPA 2018**), initially providing very limited information. After prolonged correspondence, Mr Ashley issued court proceedings. In January 2025, the Court ruled largely in Mr Ashley's favour, finding HMRC had failed to meet its UK GDPR obligations.

This case provides valuable guidance on how to manage DSARs. The key takeaways can be summarised as follows:

- **Adequately define the scope of the search:** Controllers should apply a holistic, organisation-wide approach to searches which should not be limited by internal policies or departmental boundaries.
- **Carry out reasonable and proportionate searches:** The Court found that HMRC had incorrectly asserted that the searches were disproportionate based on hours taken (150 hours). Time alone does not make a search disproportionate; consider size, resources and nature of the request.
- **Carefully consider the meaning of personal data:** The Court applied a broad interpretation of the definition of personal data, finding that the valuations of Mr Ashley's properties were to be regarded as his personal data because they were, by reason of *"their content, purpose or effect",* linked to Mr Ashley. Controllers should assess each piece of information individually to determine if it relates to the data subject, maintain consistency across departments and keep an audit trail of decisions.
- **Use of exemptions:** The "tax" exemption allows a controller to withhold personal data if disclosure is likely to prejudice the assessment or collection of tax. It was held that "likely" means a *"very significant and weighty chance of prejudice",* supported by evidence. The application of exemptions requires strong evidence and must be applied granularly, not on a blanket basis.
- **Provision of data:** Responses must be concise, transparent, and intelligible. Controllers should avoid excessive redaction that renders data meaningless; provide context where necessary to ensure intelligibility.

From DSARs to data protection complaints:
Implementing lessons from 2025

In June 2025, the Data (Use and Access) Act 2025 (**DUAA**)received royal assent, making minor changes to the existing DSAR regime. Specifically, it inserted a new Article 15(1A) UK GDPR which provides that a data subject is only entitled to personal data based on a *"reasonable and proportionate"* search. Unusually this provision has retrospective effect and is treated as having come into force on 1 January 2024. The DUAA also inserted a new Article 12A UK GDPR which sets out the meaning of *"applicable time period"*, bringing the "stop the clock" mechanism formally into the wording of the UK GDPR. Whilst these changes merely codify existing guidance from the Information Commissioner's Office (**ICO**), putting them on a statutory footing may encourage controllers to be more robust in their approach to unwieldly DSARs.

The ICO also kept DSARs firmly on its radar with some notable enforcement action in 2025, including those against Bristol City Council (**BCC**) and South West Police (**SWP**). The facts of these cases were similar; both public sector bodies failed to respond to hundreds of DSARs within the statutory time frames over several years. This resulted in numerous complaints to the ICO and some data subjects expressing distress and detriment as a result of the delays. For both BCC and SWP, the DSARs often involved large volumes of data including sensitive information and, in the case of BCC, children's data.

Both controllers cited lack of resources and experienced staff to deal with backlogs. Whilst the ICO acknowledged these difficulties, stating *"Although the Commissioner understands the difficult impact under-resourcing has on organisations, this is not an issue for the Commissioner to rectify. The Commissioner is of the view that it is for organisations to demonstrate compliance with their data protection obligations and to judge what resources are appropriate for their organisation as they alone will best understand the pressures they face and the nature of their business",* it found that infringements had occurred and issued formal enforcement notices requiring urgent remedial action within specified timescales. This aligns with the ICO's public sector approach to enforcement; had the controllers been private sector companies there may well have been an outcome which included a monetary penalty notice.

These actions illustrate the importance the ICO places on the enforcement of this fundamental right and serve as a broader warning to all organisations that failure to prioritise DSAR compliance can lead to reputational damage, regulatory scrutiny, and significant operational disruption.

An unreported ICO prosecution in September 2025 signalled a possible shift in its approach to DSAR enforcement with what we understand to be its first criminal prosecution under section 173 Data Protection Act 2018 (DPA 2018). While DSAR non-compliance is usually handled as a civil matter, with the ICO able to take enforcement action including the issuing of a monetary penalty notice, section 173(3) DPA 2018 provides an alternative route of action. Under this section, it is a criminal offence to intentionally alter, erase, block, or conceal information to prevent disclosure in response to a DSAR, designed to deter deliberate obstruction of the right of access.

In this case a director of a care home was convicted for deliberately obstructing a DSAR submitted on behalf of a resident. Instead of responding to the request, the director concealed and erased records, prompting an ICO investigation and resulting in a criminal prosecution leading to a conviction and a fine of £1,100 (plus costs of £5,440). This case illustrated that deliberate obstruction can lead to criminal liability of the controller, as well as its employees, officers or individuals. This offence is punishable not only by a fine but, in serious cases, imprisonment.

The implications of the ICO's step change in its decision to bring a criminal prosecution is potentially significant, particularly in light of the new enhanced investigatory and enforcement powers under the DUAA. At the time of writing, these powers were expected to come into force in early 2026 or swiftly thereafter, and enable the ICO to require the production of documents under existing powers to issue "information notices"; require individuals to attend interviews and answer questions via new provisions relating to "interview notices"; and require controllers or processors to commission and pay for a report by an "approved person" as part of an investigation into a specific matter. These broad powers will enable the ICO to conduct more thorough and effective investigations.

Trends in technology and public sentiment

2025 seems to have been the year that DSARs entered the mainstream. Public awareness has steadily increased, such that any contentious employment matter, data breach, or other issue causing dissatisfaction to a data subject will inevitably see the submission of a DSAR. This rise in public consciousness, coupled with the accessibility of technology, has led to a perfect storm for controllers on the receiving end of requests.

The rise of AI-supported DSARs

With the use of AI technology such as ChatGPT now commonplace in our daily lives, it is clear to see its growing utilisation in data protection issues, and especially in relation to DSARs. The ease of use and free access to these tools has brought about a commensurate rise in the number of DSARs organisations are receiving. From a data subject perspective, any technology that makes exercising individual rights easier is to be welcomed. But what of the impact this is having on controllers?

We have observed AI-supported DSARs falling squarely into two camps. On the one hand, some data subjects are clearly using AI to generate what will often be very long and confused correspondence, frequently containing irrelevant or out of context UK GDPR references, without applying much in the way of critical thinking to the request. This makes the task of deciphering the request itself time-consuming for the controller, before they have even begun to respond.

On the other hand, and particularly in the employment field, we are seeing a large increase in DSARs that are becoming increasingly sophisticated. There is clear evidence of the considered use of AI technology to carefully scope DSARs in ways which can make the response more complex for the employer with clear links being made to the issues in prospective or ongoing employment disputes.

Further to the initial DSAR, AI-generated complaints are now commonplace, frequently being sent to controllers within minutes of the parties' prior communications. As well as the obvious impact on resources involved in responding, the effect on the data protection team's morale should also be acknowledged. Having to dedicate additional time to respond to largely unfounded or fallacious queries which have been AI-generated, as opposed to genuinely-held concerns of the data subject, can be frustrating as well as time-consuming, especially where it is obvious that the data subject has not taken the time to properly consider the response they have received.

These issues highlight that data subjects themselves are often misinformed and may not even fully understand their own requests or correspondence. For instance, the accuracy of the AI output will be guided by the quality of the prompts made by data subjects. For controllers, we can expect to see a rise in the submission of complex DSARs; either being complex in the sense that they are difficult to decipher, or because the DSARs are very focused on the key issues in hand, which are in themselves complex in nature.

The rise in use of AI technology is clearly benefiting data subjects by making it easier to exercise their right of access, with controllers across the board feeling the additional burden of both the uptick in volumes of requests, and the increasing sophistication and complexity of the DSARs being submitted.

There is certainly a role for controllers to make use of AI tools themselves to support DSAR responses, but we are some way from AI being a panacea for the headache that is a complex or bulk DSAR.

Mass DSARs

Throughout 2025, the trend for DSARs stemming from circumstances of a cyber-attack has continued. Given the constant threat of cyber incidents, as evidenced by several high-profile attacks in the UK this year, we do not expect this trend to slow in the coming year.

AI-generated DSARs can only serve to fuel this trend, often being utilised as a vehicle to overwhelm controllers where a large group of data subjects have been impacted by the same data breach. For instance, we have observed a vocal community of data subjects sharing AI-generated DSARs via social media (e.g. Reddit) for onward transmission to the relevant controller following a personal data breach involving over 4,000 data subjects.

The combined use of AI and social media by data subjects exposes controllers to the risk of mass DSARs which have the potential to overwhelm organisations, taking significant resources to respond.

Whilst by no means a new development, we are continuing to observe the submission of mass DSARs via Claims Management Companies (**CMCs**) as a pre-cursor to bringing a regulated complaint or litigation. This frequently raises questions about the authenticity of the DSARs; whilst clients of CMCs may have signed letters of authority, they are not always fully aware that this right will be exercised on their behalf. Indeed, we are aware of cases where CMCs have insisted that DSARs are pursued, failing which the data subject is told they will be required to pay significant costs incurred by the CMC on their behalf. Although such pressure tactics displayed by CMCs may not be considered 'enforced' DSARs within the meaning of the DPA 2018, it is clear that in some cases it is not the data subject that genuinely wants to exercise their rights. Whilst the ICO guidance is silent on this point, controllers might therefore look to other avenues, such as refusing a request on the basis it is manifestly unfounded.

In any event, controllers must generally respond to DSARs despite the use of AI. While there are existing protections against DSARs considered manifestly unfounded or excessive, guidance from the Information Commissioner's Office explicitly refers to an expectation that DSARs are reviewed on a case-by-case basis. Whilst at one point, there were proposals to lower the threshold for refusing to deal with such requests to those which were 'vexatious', this proposal to lower the bar did not find its way into the final text of the DUAA.

Even if the burden on controllers in handling individual AI-generated DSARs does not seem onerous, one can appreciate the issues caused when DSARs are brought to controllers en mass. Obvious difficulties emerge in relation to the capacity of controllers to review mass DSAR requests which are increasingly being weaponised by data subjects given ease of access via AI.

Strategies for dealing with DSARs

Having painted a pretty bleak picture, you may be wondering what practical steps you can take to get a handle on DSARs. A controller's guiding principles for approaching DSARs might look something like this:

- **If in doubt, ask the data subject:** this applies to seeking clarification if a request is unclear (especially if it is difficult to decipher what the data subject is genuinely seeking when they have submitted an AI-generated DSAR), requesting proof of ID if there is any uncertainly, or requiring a letter of authority where a DSAR is submitted on behalf of (rather than directly by) the data subject. Where clarification or ID is required, the controller may adjust the deadline in accordance with the relevant provisions.
- **Create an appropriate search strategy:** this should be considered early on in the process, and documented, particularly as searches are refined to create a set of responsive documents that are relevant to the data subject's request as well as being reasonable and proportionate. The controller should think about the background that has given rise to the DSAR and use that context to identify key data custodians across the organisation and to generate search terms, along with any key phrases that the data subject may have used in the request.
- **Extend the deadline where a request is complex:** it is perfectly legitimate to extend the deadline by up to a total of three months if there are complexities in handling the DSAR. This will perhaps be even more likely with the submission of AI-supported DSARs. Use the ICO's guidance to help identify all relevant factors and document the decision. The data subject should be made aware of the application of the extension, along with the accompanying reasons. In any event, the controller should respond "without undue delay" which may be earlier than within the full three months where appropriate.
- **Refuse to handle manifestly unfounded requests:** where the DSAR appears not to be a true exercise of the data subject's right, consider whether the request can be refused on the basis that it is manifestly unfounded. This should be considered on a case-by-case basis, but if a data subject appears to be under pressure from a CMC to make or continue with a DSAR, then this may be one option to refuse the request.
- **Use AI and other tech to your advantage:** careful use of AI tools can save time and reduce manual involvement. Current use cases predominantly focus on activities such as de-duplication, batching and transcribing. We expect to see further uses of AI as technologies develop and become more sophisticated. Might we see an AI tool that could assess the tone of a data subject's correspondence to predict when a DSAR might be made, to enable a controller to be on the front-foot in terms of recognising DSARs promptly and predicting peaks in resourcing requirements?
- **Use AI with care:** don't assume that references to legislation and caselaw are correct and train your staff to be healthy sceptics regarding any AI-generated content.
- **Accountability:** always keep clear records of your DSAR handling especially in relation to search parameters and the application of any exemptions. You should ensure that this records the use of exemptions on a case-by-case basis, and not as a blanket approach. Records will prove useful if you receive a data subject complaint (on which, see below) or ICO enquiries.
- **Peer review:** for particularly contentious DSARs, consider whether the DSAR response should be peer reviewed by someone within the business (with skills beyond data protection), who is aware of the broader context in which the DSAR has been submitted. In an employee context, carefully consider who will review the response bearing in mind the potential sensitivity of HR related issues and the need to have knowledge of any wider anticipated or ongoing employment dispute.
- **Respond but don't pander:** Try to avoid being drawn into protected and lengthy correspondence after responding to the DSAR. Refer back to the original comprehensive response and remind the data subject of their right to submit a data protection complaint, following which they may complain to the ICO. Ensure you keep a record of all your correspondence in case it is required by the ICO.

DUAA – Data Subject Complaints

The DUAA introduces a new section 164A DPA 2018 which requires controllers to establish processes to handle data protection complaints. The requirement is expected to come into force in or around June 2026. A data protection complaint may arrive in a number of guises: it may relate to a stand-alone data protection issue, be a follow-up to a DSAR where the data subject is dissatisfied with the response or could even be submitted alongside the DSAR itself.

Controllers will be required to facilitate the making of such complaints, for example by making complaint forms available online or by hard copy, noting that a form is not mandatory but is suggested as one option. An acknowledgement must be issued within 30 days of receipt, and the controller must take appropriate steps to respond to the complaint without undue delay and inform the data subject of the outcome. The data subject should be kept informed if the final response is going to take some time.

Aside from the mandatory 30 day complaint acknowledgement and the obligation to investigate and respond without undue delay, the DUAA is not unduly prescriptive, and it therefore leaves controllers with a healthy degree of flexibility to tailor their complaints process to their own business.

It is also worth being aware of related powers issued to the Secretary of State to issue secondary legislation establishing mandatory complaint volume reporting to the ICO. When in planning phase, controllers would be wise to think about what management information they should collect, both for internal reporting purposes but also with one eye to the future if the mandatory reporting requirements do come to fruition.

ICO guidance

Noting that this is a new requirement, albeit many controllers will already have some form of complaints process in place, the ICO has already produced and consulted on draft guidance. We are expecting the final guidance to be published this year, to give controllers time to take into account the ICO's views when establishing or finessing a data protection complaints process. As well as addressing specific legislative requirements, the draft guidance also offers examples of practical steps.

Creating a data protection complaints process

Key to ensuring that your organisation is ready for this new requirement is to create a complaints procedure, and make your staff aware of the right to complain and who will be responsible for handling any complaints. Crucially, you should consider any overlap with other complaints obligations you may have, especially if your organisation is in a regulated sector, since data protection issues are most often raised as a precursor to or as a result of a bigger issue. Do you want your complaints processes to be conjoined, or separate but aligned? Do other complaints regimes have stricter timescales attached to them?

The key steps in preparing to deal with data protection complaints include:

- Drafting or updating a formal complaints procedure, which data subjects can easily access.
- Assign roles and responsibilities – will your process facilitate an independent review by a colleague who was not involved in any previous correspondence with the data subject?
- Consider how data subject complaints will sit with any other complaints obligations you are subject to.
- Ensure that you keep adequate records to be able to respond to complaints.
- Inform and train staff – this will involve specific training for data protection staff, as well as general awareness-raising across your organisation so that complaints are recognised and dealt with promptly.
- Establish a process to identify any trends and to reflect so that any lessons are learned.

Whilst a new data protection complaints regime may feel like an additional burden, it might also help to draw a line under protracted correspondence with dissatisfied data subjects. Once you have issued a final complaint response, an ICO referral is the next logical step. If you are confident that you have addressed the data subject's concerns, and have documented any decisions taken, then you should take a robust approach to using the complaints process as the end point and let the complaint run its course via the ICO. When crafting your complaints processes and any associated management information you will collect, think carefully about how this data can be used to your advantage; allow the information to give you meaningful insights into your organisation's compliance so you can get ahead of any more systemic issues; if you have good oversight of your complaints you will be able to identify recurring themes, and can intervene before they become big-ticket problems.

It's clear that 2026 will be a busy year for data protection practitioners, not least because we are unlikely to see any decline in the number of data subjects exercising their rights and thereafter submitting complaints about the handling of those requests, or for broader data protection concerns. Whilst there is no effective short-cut to compliance, controllers can begin to make smart use of technology to assist where possible, and should shore-up internal processes, training, and strategies to ease the burden of handling DSARs and data protection complaints.

PR and penalties: Behind the ICO regulatory strategy

05

Hans Allnutt
Partner
hallnutt@dacbeachcroft.com

Lara Maslowska
Associate
lmaslowska@dacbeachcroft.com

Published in November 2022, the ICO25 plan stated that the ICO aims to provide *"certainty on what the law requires, what represents good practice and* [its] *approach as the* [data protection] *regulator."* As a regulator, it aims to *"intervene proportionately, clearly and only where needed."* Recent criticism suggests that there is a great need for the ICO to intervene with enforcement but is failing to do so.

An open letter issued by civil society organisations, academics and data protection experts called out a 'collapse' in ICO enforcement activity. The letter argues that current ICO enforcement strategies in respect of both public and private sector, deprioritising enforcement, are resulting in more data breaches than the regulations are supposed to prevent. The open letter claims a *"strong correlation between the ICO lack of formal regulatory action and a surge in, sometimes egregious, data breaches in the UK."*

The open letter points to the ICO's interpretation of its duty to promote growth and its softer public sector approach as reasons why the ICO is not enforcing. The open letter recognises the recently concluded ICO call for views on how it investigates infringements, but notes that this stops short of consulting on the ICO's overarching approach to enforcement, claiming that alleged deeper structural failures should be the subject of a public inquiry.

The open letter is scathing in its criticism but is by no means an exceptional critique of the ICO's approach to enforcement. As in the open letter, the public sector approach is often cited as an example of the ICO being a toothless regulator. For the private sector, of those precious few fines that are levied, a disproportionate number are reduced through appeal.

We do see, anecdotally, an effect on an organisations' assessment of risk. In the immediate aftermath of a breach, clients will often ask whether they should expect to receive a fine from the ICO. Irrespective of the circumstances of the breach, the reality is that a fine is statistically unlikely, if not very unlikely. Over 4,500 ransomware incidents have been reported to the ICO since 2019 with a small handful resulting in a fine. It is no surprise that critics, such as the signatories to the open letter, say this has led to organisations underfunding and under-resourcing data protection with a resulting rise in breaches.

If the probability of a fine being issued weighs in favour of a breached organisation, what then determines if the ICO will issue a fine? The ICO has published a range of factors in its Regulatory Action Policy (**RAP**) for all to see. However, this has arguably highlighted a significant level of inconsistency between comparable breaches, with some facing enforcement whilst others do not (ransomware, being one such example). Could the answer lie in one factor, which is notable by its absence in the RAP despite appearing in other policy documents that drive the ICO's objectives; that of the ICO's reputation. Before exploring this further, we delve into some of the examples cited as ICO enforcement inactivity.

The ICO public sector approach

In November 2025, the ICO published an update to the much-discussed public sector approach, clarifying those organisations in scope, and crucially, the circumstances in which a fine may be issued. Introduced in 2022, the public sector approach saw the increased use of the Commissioner's discretion to reduce the impact of fines on public bodies, with the additional aim of improving data protection standards.

Following a two-year trial, and subsequent consultation on its extension, the ICO confirmed the approach would continue, publishing clearer guidance on the circumstances in which a fine may be considered. Only in the 'most egregious' cases will a fine be issued to a public authority, following an assessment of a number of factors including:

- Actual or potential harm to people: physical or bodily harm, psychological harm, economic or financial harm, discrimination, reputational harm or loss of human dignity.
- Intentional or negligent character of the infringement.
- Relevant previous infringements, or recent infringements, by the controller or processor.

The value of any fine will be determined by the application of the five-step approach set in the ICO Data Protection Fining Guidance.

Reflecting on those fines issued since the public sector approach was introduced in June 2022, the £750,000 fine issued to the Police Service of Northern Ireland (**PSNI**) remains the high water mark. The PSNI was fined for exposing the personal information of its entire workforce, with individuals affected leaving the police, taking steps to remove their online presence and expressing concern for not only their safety but that of their families.

The ICO applied its public sector approach in issuing the fine of £750,000 as opposed to a fine of £5.6 million had the usual fining guidance applied. The Information Commissioner commented that PSNI's financial pressures had to be balanced against need to protect people's rights. Clearly to the casual observer, this incident fell within the definition of 'most egregious'.

However, other actions indicate that that further fines for public authority data breaches remain unlikely, even after clarification on what is considered a 'most egregious' case.

The accidental disclosure of personal details of thousands of Afghans who had worked with British forces by the Ministry of Defence (**MoD**) was described as *"unacceptable and should never happen again – the stakes are simply too high."* This sounds not too dissimilar to the ICO's criticism following the PSNI breach, yet the ICO elected to take no regulatory action in this instance. Unsurprisingly, accusations of inconsistency in the ICO's approach were raised.

The Information Commissioner published his comments on that decision, emphasising the wide range of issues in play. It was made clear that the impact of the data breach brought it within the range of a 'most egregious' incident justifying a fine. However, other circumstances justified no further investigation such as the considerable expense to the public purse due to the MoD's prompt response.

The ICO record of the data breach concluded that *"while no formal investigation had been commenced, it was considered that, in line with the ICO public sector approach, further engagement with the MoD would be more effective than seeking to impose additional cost to the tax payer at this time."*

Interestingly, this reason is not included in the factors listed in the RAP which the ICO has stated will determine whether it selects a breach for enforcement. The list is, admittedly, not exhaustive but this response highlights that the ICO's reasons for not acting are potentially unlimited and unknown.

The ICO's messaging appears to indicate that organisations are fortunate to attract leniency and that the policy serves the public good.

However, critics have suggested that the policy is perhaps more self-serving. It has not been lost on a number of organisations facing reprimands that it enables the ICO to avoid the scrutiny, cost and uncertainty of the appeals process that could be followed if a financial sanction were issued. It is not open to an organisation to appeal a reprimand to the First-tier Tribunal, only the unconventional and higher threshold of judicial review. Unsurprisingly, no public organisation has sought to subject a reprimand to judicial review.

The public sector approach therefore arguably allows the ICO to act with a wider licence in actions where an appeal is possible. Indeed, one of the responses to the ICO's consultation on the continued use of the public sector approach suggested "formalising public reprimands as formal enforcement tools in legislation and putting in place an accompanying appeals process." Concern is rising that an ICO reprimand alone does not drive the systemic change the ICO seeks in the public sector.

In support of the public sector approach, the ICO has maintained that reprimands are an effective deterrent, mainly due to the reputational damage and public trust relating to the reprimanded organisation. Reputational consequences, both for the investigated organisation and the ICO, is a lever considered in greater detail below. However, the public sector approach is increasingly being seen as a watering down of enforcement and it is now questionable whether the reputational and public trust damage is that of the ICO's.

This watering down can be seen in some of the more recent reprimand language. Whereas previously the ICO has listed the fine that would have been applied, but for the reprimand, the ICO has taken to using "up to" language. This reflects the fact that if it were to issue a fine, it would have to offer the discounts that it has afforded to other organisations for mitigating circumstances, cooperation, early payment, and general negotiation (see Capita below). One might question the effectiveness of a deterrent if the ICO is having to implement caveats on its own predicted enforcement.

Finally, in another notable reprimand against a public sector organisation, that organisation simply responded with a statement to the effect that it disagreed with the ICO's finding but that it would not use its limited resources to challenge it. Again, an effective appeals process may have enabled a judicial finding of a winner in this particular example but unless and until the ICO adopts this recommendation, the public sector approach will continue have these outcomes.

All fine but no fines?

When it comes to the fines themselves, a review of the ICO's Annual Reports indicates a clear direction of travel for the ICO's enforcement strategy against the private sector.

Between 2021–22 and the present, the ICO's enforcement trends have shifted from a model of occasional but significant financial penalties toward a more collaborative and proportionate approach. While early reports emphasized large monetary penalties for serious breaches, subsequent years saw a marked increase in the use of reprimands and negotiated settlements (as noted for public sector bodies). The ICO has arguably moved towards balancing deterrence with practical compliance support, aligning enforcement with broader objectives of trust, innovation, and proportionality.

The difficulty with this approach in reality is that few breaches impacting large numbers of individuals and/or affecting individuals in significant ways, will be subject to enforcement. Those few matters that do progress to enforcement face potential leniency which undermines their effect.

A recent regulatory fine issued highlights the signals the ICO is sending, and perceptions of leniency even where fines are levied. The outsourcing organisation Capita received a fine of £14 million in October 2023 for failing to ensure the security of personal data during a 2023 data breach. The fine covered Capita's failings as both data controller and processor.

Interestingly, Capita argued it should be treated like its public-sector clients, who would typically expect to receive reprimands instead of fines per the public sector approach. This could reflect a view that the public sector receives enforcement leniency, encouraging Capita to make such submissions as a result.

Nonetheless, Capita successfully achieved a significant reduction on the fine proposed. The proposed £45 million fine was reduced to £14 million after Capita admitted liability, agreed not to appeal, and demonstrated remediation efforts, victim support, and cooperation with regulators.

The significant reductions given to Capita, received in response to the above factors, totalled almost 70%. Notwithstanding the reasons given in the Monetary Penalty Notice, the Commissioner noted the wider consideration of the *"impact a penalty of this nature will have on the growth of the UK economy, and the desirability to promote economic growth, innovation, and competition."* Furthermore, Capita's voluntary admission and commitment to refrain from appealing contributed to the overall context.

Of course, a reduction of 70% is still arguably better than the overturning of a fine altogether. The £7.5 million fine imposed on Clearview AI by the ICO in 2022, one of the first multi-million pound penalties, is now the subject of a further appeal to the Court of Appeal. The most recent ruling from the Upper Tribunal has supported the ICO's initial stance, but the lengthy appeals process reveals vulnerabilities in the ICO's enforcement system, especially when dealing with matters involving extraterritorial jurisdiction.

Aspect	Public Sector	Private Sector
Policy Basis	Public Sector Approach (permanent since Dec 2024)	Standard enforcement under UK GDPR and PECR
Primary Enforcement Tools	Reprimands and Enforcement Notices Corrective action plans	- Monetary Penalties (fines) - Enforcement Notices - Reprimands
Use of Fines	Rare – reserved for 'egregious' cases. Hypothetical fine amount now published for transparency (Post Office)	Common for serious breaches Fines calculated based on severity and turnover
Transparency	ICO may publish what amount the fine would have been "up to", had a fine been levied	Actual fine amounts published, including details of reductions
Impact on Budgets	Designed to avoid diverting taxpayer funds from essential services	No special consideration – fines intended as a deterrent
Scope	Government departments, local authorities, NHS bodies	All private organisations, including charities and social enterprises
Criticism	Concerns about reduced deterrent effect and accountability	Criticism usually relates to proportionality of fines

Reputation, reputation, reputation

One would be forgiven for thinking that the ICO is more frequently explaining why it has not enforced a fine, or reduced it, than pursuing enforcement. Further, one might ask what drives the ICO to decide take enforcement action in one breach over another.

As noted above, the RAP includes a range of factors that the ICO will consider when deciding to take enforcement action. This is a non-inclusive list and, also noted above, the ICO has not shied away from citing previously unpublished factors in its enforcement decision making.

One further factor, however, that isn't immediately obvious is the ICO's own reputation and public relations image.

In the ICO's former regulatory action policy of 2013, the ICO was more candid in including the impact on its own reputation within the criteria for deciding whether to take enforcement. The criteria included:

> ***"What is the risk to the credibility of the law or to our reputation and influence of taking or not taking action as a reputational factor when considering enforcement."***

More recently, examples are a little harder to find although the ICO's own disclosures of internal ICO messaging between ICO senior management following the MoD Afghan incident, are instructive:

> ***"There is a reputational risk we face if we can't sufficiently explain why we took the course of action we took (or in this case - didn't take)."***

In some ways, the PR influence on enforcement is not surprising. Following a well publicised breach in the mainstream media, public attention is directed at the ICO for comment. Public expectations of the ICO will be higher in a publicised breach compared to a breach that has not garnered wider publicity.

That being said, it may not be fully appreciated by organisations that a factor in whether it will face enforcement could be the level of publicity and the ICO's own reputation as opposed to the pure technicalities of the breach itself.

The 'new' Information Commission, will it prompt a change in approach?

The Data (Use and Access) Act (**DUAA**) introduced a number of significant changes to the structure and operation of the ICO. The ICO as it is currently structured is to be abolished and replaced with the Information Commission. At the time of writing, the changes (identified as Stage 4 of the commencement provisions) are expected to take effect from early 2026.

The change moves the regulator to a body corporate with a statutory board with a chair and chief executive. The Information Commissioner responded positively to these changes, stating the *"refreshed governance arrangements will maintain our independence and enhance our accountability."*

The DUAA also provides the Information Commission with additional powers in respect of a breach of the Privacy and Electronic Communications Regulations 2003 (**PECR**) as provided under the UK GDPR or Data Protection Act 2018. However, the question remains whether this organisational change will prompt any change in the ICO's approach to either public or private sector enforcement.

This is where politics and performance may play their role. The DUAA introduces a new principal objective for the Information Commissioner as follows:

> ***(a) to secure an appropriate level of protection for personal data, having regard to the interests of data subjects, controllers and others and matters of general public interest, and***
>
> ***(b) to promote public trust and confidence in the processing of personal data.***

As well as the principal objective, the Information Commissioner must also have regard to 'other matters' such as promoting innovation, competition, safeguarding public and national security, children's interests, and the importance of dealing with criminal offences.

Against this backdrop, sections 95 and 102 of the DUAA introduces reporting requirements on the Commissioner:

1. **To prepare and publish an annual analysis of the Commissioner's performance using key performance indicators.**
2. **To prepare and publish an annual report on UK GDPR investigations including details on the number of investigations undertaken during the reporting period, the acts or omissions compelling the investigations, any enforcement powers exercised, the duration of investigations and details of the types of outcomes.**

Of course, the ICO already issues Annual Reports as discussed above, but the new reporting requirements are a change from purely operational reporting to strategic accountability. It should be noted that the DUAA removed provisions previously contained in the Data Protection and Digital Information Bill. This would have had the Information Commissioner consider a statement issued by the relevant Secretary of State, which would have outlined the government's strategic priorities with respect to data protection.

That proposal raised concerns about the independence of the ICO both in the UK and EU; unsurprisingly the proposal was withdrawn. Nonetheless, the requirements on the Commissioner to prepare performance and investigation reports may impact the enforcement approach going forward.

More detailed public reporting may make the ICO's regulatory priorities and sector-specific concerns more transparent for organisations. The requirements on the Information Commissioner to take account of 'other matters' may also mean that enforcement focus may shift towards areas such as children's privacy and industries with a high level of innovation such as AI.

At present, fines for breaches of the UK GDPR have, arguably, lost much of their impact due to how infrequently they are imposed, the amounts involved, and challenges with enforcement. The ICO still appears hesitant to fully exercise its authority to issue financial penalties against both public sector bodies and private companies.

It is unclear if the DUAA reporting requirements will affect the ICO's approach. Overall, the remit of the Information Commissioner to consider economic factors such as innovation and competition, alongside its traditional focus on data protection with an undercurrent of its own public reputation, suggests there will continue to be a nuanced regulatory approach from the ICO.

Post-breach: The discretion in assessing the risk of harm

06

Justin Tivey
Partner
jtivey@dacbeachcroft.com

Becky Lea
Legal Director
blea@dacbeachcroft.com

Protection from the risk of harm is central to most data protection regimes. The clue is very much in the name after all: the General Data Protection Regulation, the Data Protection Act 2018 and so on.

Personal data breaches occur when there is a security breach leading to accidental or unlawful loss of confidentiality, integrity or availability of personal data being processed. A breach does not have to result from third party activity; accidental disclosures are as reportable as hacking incidents.

Preparatory tools such as Data Protection Impact Assessments (**DPIAs**) help identify, analyse and mitigate the risks to data subjects likely to arise from processing activities relating to their personal data.

However, if the worst happens, and a data breach occurs, then the assessment of risk arises again. Whether a controller must notify the regulator or affected individuals is determined by assessing the potential risk of harm to affected data subjects. This assessment is necessary to ensure compliance with regulatory requirements and to avoid the risk of possible enforcement activity by the regulator. Failures to notify can attract regulatory investigations, reprimands and potentially fines of up to £17.5 million or 4% of total annual worldwide turnover, whichever is higher.

The legislative background and guidance

The two key statutory provisions within the UK GDPR for post-breach considerations are as follows:

Article 33(1):

"In the case of a personal data breach, the controller shall without undue delay and, where feasible, not later than 72 hours after having become aware of it, notify the personal data breach to* [the regulator], *unless the personal data breach is unlikely to result in a risk to the rights and freedoms of natural persons."

Article 34(1):

"When the personal data breach is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall communicate the personal data breach to the data subject without undue delay."

How does a controller decide if a breach is unlikely to result in a risk and thus is not notifiable to the regulator or, in the alternative, that the breach is likely to result in a high risk and should be communicated to data subjects themselves?

Both the UK data protection regulator, the Information Commissioner's Office (**ICO**) and the European Data Protection Board (**EDPB**) have issued guidance about how to make this assessment, even including a mathematical formula! Controllers need to consider how straightforward this guidance is to use and whether risk and harm assessments allow for subjectivity or are strictly formulaic.

The GDPR Recitals indicate the types of harm that might result from a data breach. Recital 85 highlights *'material'* and *'non-material'* damage and lists *'loss of control'* over personal data, *'limitation of rights'* and more specifically *'discrimination, identity theft or fraud, financial loss, reversal of pseudonymisation, reputational damage, loss of confidentiality of data protected by professional secrecy or any other significant economic or social disadvantage'*.

Recital 88 states that consideration should be given to the circumstances of the breach, including whether technical measures protecting data limits the likelihood of identity fraud or other misuse of the data.

These factors are good starting points in indicating the types of harm that controllers need to think about when assessing the risk to data subjects. However, there is a lot of room for interpretation in some of these concepts; limitation of rights, reputational damage and economic and social disadvantage in particular are capable of being interpreted widely or narrowly. One controller or data subject may see information as reputation-related, while another may not. How does the controller determine whether the risk of harm is unlikely, high, or somewhere in between?

Recital 75 also lists types of data which might give rise to risk including special category data, data evaluating data subjects, data relating to vulnerable people and large datasets. Recital 76 refers to the likelihood and severity of the risk to data subjects but only to state that the risk should be determined *'by reference to the nature, scope, context and purposes of the processing'* and that risk *'should be evaluated on the basis of an objective assessment, by which it is established whether data processing operations involve a risk or a high risk'*.

Recital 77 goes on to suggest that guidance as regards the identification of the risk related to processing and assessment of the likelihood and severity could be provided by means of guidelines provided by the EDPB. The EDPB may also issue guidelines on processing operations unlikely to pose high risks to individuals' rights and freedoms. To date, the EDPB has issued guidance on these types of operations via binding opinions such as Opinion 6/2024 on processing operations exempt from the data protection impact assessment requirements.

The regulatory guidance

The ICO advises controllers to implement strong breach detection and reporting procedures to help determine if they must notify authorities or affected individuals about a data breach. The guidance starts by referencing the Recitals of the UK GDPR mentioned above. The ICO recommends that the focus should be on any potential negative consequences for individuals, particularly if the breach is not addressed appropriately. On becoming aware of a breach, the controller should contain it and assess the potential adverse consequences for individuals, based on how serious or substantial these are, and how likely they are to happen.

The ICO also recommends consulting risk section IV of the Article 29 Working Party Guidelines (now the EDPB) on personal data breach notification, now identified as Guidelines 9/2022. Drafted in anticipation of the implementation of the GDPR and adopted on 3 October 2017 and updated in 2018. 2022 and 2023, these are a useful reference point.

The Guidelines repeat the Articles and Recitals referred to above but add assumptions that where special category data is involved, then damage to data subjects should be considered "likely to occur". The need to evaluate likelihood and severity of risk, and to make an objective assessment, is also emphasised. The difference between a DPIA and breach assessment is also highlighted, noting that the DPIA considers hypothetical situations, whereas a breach scenario is in the context of an actual event.

The EDPB Guidelines also highlight assessment criteria to help the thought process for the controller. The key criteria being:

- **The type of breach:** For example, a breach involving the disclosure of data to a third party may create more risk than one where data is simply lost.
- **The nature, sensitivity and volume of data:** The relevance of these criteria is self explanatory, but the guidance highlights greater risk if different types of data can be combined to create an identity theft risk. Similarly, a large variety of data about one person or a large volume affecting many people will also increase risk.
- **The ease of identification of individuals:** Can a specific individual be identified from the data or with the data and other available information? If not, the risk is lowered. Similarly, encryption and pseudonymisation of data will lower the risk.
- **The severity of consequences to the data subjects:** Again, this speaks for itself, but the guidance also flags that if data is in the hands of a malicious party the risk will likely increase. Accidental disclosure to a third party who agrees to delete the data being at one end of the scale and unknown hackers at the other end. The duration of any adverse consequences is also a factor.
- **The characteristics of data subjects and controller:** These considerations are contextual. Children and vulnerable adults are obvious examples of data subjects at inherently greater risk. Controllers such as medical establishments are likely to have data which could cause harm if mishandled.
- **The number of affected individuals:** Although one individual can be seriously impacted, the guidance is that generally the higher the number of affected data subjects, the greater the impact of the breach.

Overall, the guidance confirms that, if in doubt, a controller should err on the side of caution and notify. Additional examples of scenarios and suggested appropriate notification outcomes are provided in the guidance, but it is far from a clear decision tree or formula.

The formulaic approach

Mathematics geeks may want to refer to the European Union Agency for Network and Information Security (**ENISA**) for guidance. ENISA has published a recommended methodology for the assessment of severity of personal data breaches. The methodology was developed in 2013 based on the legislation at that time, but with an eye to the then-developing GDPR.

The aim was to develop a quantitative tool to assess the severity of data breaches to assist in notification decisions. The tool identifies factors, scores them, then uses a formula to produce a 'severity score'. The factors fed in are the "Data Processing Context" (DPC); "Ease of Identification" (EI); and "Circumstances of the Breach" (CB). Each factor has a score. For example, a DPC score of 1 is given for simple biographical data but it could be increased to 2, 3 or 4 if that biographical data revealed more sensitive information such as behavioural data, financial data or special category data. By way of example, a database of professional CVs might be given a score of 1, but that database is owned by an organisation (i) helping the unemployed, then the score might increase to 2; or (ii) supporting recovering drug addicts find work, then the score might be increased to 4.

The Ease of Identification (EI) scores are 0.25, 0.5, 0.75 or 1, with the scores rising as identification of the data subject is deemed easier.

Finally, the Circumstances of the Breach (CB) scores are 0, 0.25 or 0.5. These are applied depending on whether the breach is a confidentiality breach, an integrity breach or an availability breach. A lost file might score 0, an email sent to a few known recipients in error might score 0.25 and data published to an unknown number of recipients, say on a webpage, might score 0.5.

The formula is: DPC x EI + CB = Severity score

Severity score of less than 2 are suggested to be low risk breaches, score of up to 3 are suggested to be medium risk breaches, scores of up to 4 are suggested to be high risk breaches and scores of 4 and above are suggested to be very high risk.

However, even such a formulaic approach should be subject to review to consider the specific circumstances in question. For example, other features of the breach such as the number of data subjects affected (larger numbers, in the methodology put at over 100, being seen as inherently greater risk) or unintelligibility of data (for example by encryption) before reaching a final assessment.

The role of discretion

Regulators warn against of over-notification, arguing it causes notification fatigue and unnecessary alarm. Unnecessary notifications can invite claims for compensation which would otherwise be avoided. Of course, seeking to avoid claims is in no way a legitimate basis for withholding notification when the established threshold has been reached.

In reality, there is a wide role for discretion to play in post-breach risk assessment. How risk is evaluated is often still a subjective exercise straining to be objectively reasonable.

This can be illustrated by the frequently encountered problem of exfiltration of the contents of HR files. These will comprise a range of data varying wildly from data subject to data subject. Although some studies have examined the market value of personal data on the dark web, the findings generally align with expectations; for example, scans of identity documents hold significant value, whereas information such as a national insurance number is typically less valuable. Similarly, your bank account number has little value but your bank account log on credentials does.

The problem for controllers is that there is no regulatory advice on what combinations of data will likely to give rise to identity theft. It is left to organisations to determine that question for themselves.

If a controller's data is encrypted, how strong does the encryption need to be for the data to be considered secure? Is a complex password necessary, or is higher-level security required? What is the level of effort hackers would invest in accessing numerous documents with uninformative labels if the password is just "1234"? Controllers and regulators often lack the expertise or motivation typical of cyber criminals.

It is also the case that there is often no proof of a correlation between say increased phishing emails being received by a data subject and a breach that has affected their data.

This means there is still a large element of interpretation of what constitutes a 'high risk' and what is 'likely.' The decision about whether to notify or not, involving reaching a rational decision about risk, should not prioritise the controller's interests over the data subject's, and must be capable of being justified if the decision is later called into question.

Alternative approaches

In some jurisdictions, for example some US states, certain items of data or combinations are designated as requiring notification. If a breach involves that data, then those data subjects must be notified. Discretion is all but eliminated.

Evolving standards & flexibility v certainty

The future of risk assessment is uncertain. We are not aware of any radical shifts in the pipeline, certainly in the UK. The current regime leads towards more detailed data and risk analysis, whereas a broad brush or formulaic process may miss the mark for data subjects or controllers. Currently, the main emphasis is on achieving the appropriate outcome for each data subject, which aligns with the core principles underpinning the data protection framework in the UK.

Meet Our Team - United Kingdom

United Kingdom

Jade Kowalski
Partner
jkowalski@dacbeachcroft.com

Rebecca Morgan
Legal Director
rebeccamorgan@dacbeachcroft.com

Hans Allnutt
Partner
hallnutt@dacbeachcroft.com

Tessa Davies
Legal Director
tedavies@dacbeachcroft.com

Pete Given
Partner
pgiven@dacbeachcroft.com

Becky Lea
Legal Director
blea@dacbeachcroft.com

Charlotte Halford
Partner
chalford@dacbeachcroft.com

Alistair Cooper
Legal Director
alcooper@dacbeachcroft.com

Patrick Hill
Partner
phill@dacbeachcroft.com

Fran Treemer
Senior Associate
ftremeer@dacbeachcroft.com

Christopher Air
Partner
cair@dacbeachcroft.com

Yassar Lodhi
Senior Associate
ylodhi@dacbeachcroft.com

Justin Tivey
Partner
jtivey@dacbeachcroft.com

Ben Gwyther
Senior Associate
bgwyther@dacbeachcroft.com

Darryn Hale
Partner
dahale@dacbeachcroft.com

Owen Newcombe
Chartered Legal Executive
onewcombe@dacbeachcroft.com

Kate Galloway
Partner
kgalloway@dacbeachcroft.com

Abigail Gray
Associate
abgray@dacbeachcroft.com

Amanda MacKenzie
Associate
ammackenzie@dacbeachcroft.com

Astrid Hardy
Associate
ahardy@dacbeachcroft.com

Calum Doherty
Associate
caldoherty@dacbeachcroft.com

Eboni Beckford-Chambers
Associate
ebeckfordchambers@dacbeachcroft.com

Jonathan Hopkins
Associate
jonhopkins@dacbeachcroft.com

Lara Maslowska
Associate
lmaslowska@dacbeachcroft.com

Calum Glover
Associate
cglover@dacbeachcroft.com

Claudia George
Associate (Australian Qualified)
cgeorge@dacbeachcroft.com

Ethan Telford-Cooke
Associate
etelfordcooke@dacbeachcroft.com

Georgina Jones
Associate
gejones@dacbeachcroft.com

Ornela Markaj
Associate
omarkaj@dacbeachcroft.com

Hannah Clements
Associate
hclements@dacbeachcroft.com

Issac Jong
Solicitor
ijong@dacbeachcroft.com

Millie Elliot
Solicitor
celliot@dacbeachcroft.com

Ellen McWhirter
Solicitor
emcwhirter@dacbeachcroft.com

Heasha Wijesuriya
Solicitor (Australian Qualified)
hwijesuriya@dacbeachcroft.com

Isabel Becker
Solicitor
ibecker@dacbeachcroft.com

Maia Crockford
Solicitor
mcrockford@dacbeachcroft.com

Meet Our Team - Rest of the World

Argentina

Maria Jose Sanchez
Partner
msanchez@dacbeachcroft.com

Oliver Peebles
Associate
opeebles@dacbeachcroft.com

Chille

Andrés Amunátegui
Partner
aamunategui@dacbeachcroft.com

Nicolas Le Blanc
Partner
nleblanc@dacbeachcroft.com

France

Christophe Wucher-North
Partner
cwuchernorth@dacbeachcroft.com

Louis-Axel Batiste
Associate
abatiste@dacbeachcroft.com

Germany

Dr. Alexander Beyer
Partner
alexander.beyer@bld.de

Dr. Franz König
Partner
franz.koenig@bld.de

Hong Kong

David Kwok
Legal Director
davidkwok@ckleelaw.com

Italy

Guido Foglia
Partner
gfoglia@dacbeachcroft.com

Monica Spiga
Associate
mspiga@dacbeachcroft.com

Giulio Di Fabio
Associate
gdifabio@dacbeachcroft.com

Ireland

Rowena McCormack
Partner
rmccormack@dacbeachcroft.com

Aidan Healy
Legal Director
ahealy@dacbeachcroft.com

Charlotte Burke
Senior Associate
cburke@dacbeachcroft.com

Netherlands

Astrid van Noort
BLD Ekelmans
astrid.vannoort@bld-ekelmans.com

Singapore

Summer Montague
Partner
smontague@dacbeachcroft.com

Andrew Robinson
Partner
arobinson@dacbeachcroft.com

Joshua Chan
Senior Associate
jchan@dacbeachcroft.com

Spain

Pilar Rodríguez
Partner
prodriguez@dacbeachcroft.com

Igor Pinedo Garcia
Associate
ipinedo@dacbeachcroft.com

USA

Greg Lahr
Partner
glahr@dacbeachcroft.com

Aaron Mandel
Partner
amandel@dacbeachcroft.com

insurance.dacbeachcroft.com

dacbeachcroft.com

Connect with us: DAC Beachcroft LLP

Follow us: DACBeachcroft

Data

Priva

Cybe