Data, Consent and Control – UK GDPR After the Data (Use and Access) Act 2025

Procurement and Supply Chain Management Made Simple.

Section Index

Foreword

Personal data has become one of the defining resources of the digital economy. It underpins financial services, healthcare, employment, retail, public administration, artificial intelligence and countless everyday transactions. Yet the value created from data depends on information belonging to real people whose privacy, choices and opportunities can be affected by its use. Modern data protection law must reconcile two legitimate aims: enabling productive use of information while ensuring individuals are not reduced to exploitable datasets.

The General Data Protection Regulation placed individual rights, transparency and accountability at the centre of that balance. Its introduction in 2018 raised expectations around consent, lawful processing, minimisation, security and access to personal information. Following Brexit, those principles continued under the UK GDPR and the Data Protection Act 2018. The underlying philosophy stayed consistent: organisations could derive value from information, but only within boundaries designed to preserve individual protection and meaningful control.

Technology has since transformed the environment in which those principles operate. Artificial intelligence can infer characteristics never deliberately disclosed, automated systems can shape employment or lending decisions, and digital services can analyse behaviour across enormous populations in real time. Government and industry increasingly treat data as infrastructure capable of improving productivity, research, competition and public services. Regulation consequently faces pressure to protect individuals without needlessly blocking beneficial uses that earlier legislation did not anticipate.

The Data (Use and Access) Act 2025 is the United Kingdom's answer to that pressure. It does not discard UK GDPR, but adjusts parts of the framework to make certain uses of information easier, clearer or more proportionate. Recognised legitimate interests, wider automated decision-making, research provisions, Smart Data and targeted cookie exceptions reflect a deliberate policy choice: privacy protection should coexist with innovation and economic growth rather than automatically obstruct them.

The resulting question is no longer simply whether an individual consented to processing. Meaningful control now depends on whether organisations act transparently, minimise what they collect, justify their purposes, secure information and remain accountable for consequences. The post-DUAA framework is therefore a genuine test of modern data governance: whether greater freedom to use personal information can generate economic value without eroding the rights and confidence that responsible data use ultimately depends on.

Introduction – Data, Consent and Control

The distinction between consent and control matters more than it first appears. Consent is only one lawful basis under data protection law, and many legitimate activities occur entirely without it. Control is broader: it rests on transparency, lawful purpose, proportionality, security and enforceable rights. An individual may retain meaningful control even where consent is unnecessary. In contrast, a superficially obtained consent can offer little real control if the choices were manipulative or the later uses unexpected.

Tesco Clubcard illustrates the practical exchange at the heart of modern data use. The ICO records that the scheme has around 20 million users, who receive discounts and offers in return for purchasing behaviour that reveals where, how often, and what they spend on. Such information sharpens forecasting and customer understanding, but it deepens organisational responsibility. A loyalty relationship of that scale stays sustainable only where customers understand the bargain and their information is handled securely.

The Data (Use and Access) Act 2025 shifts parts of that balance towards easier, more productive data use while keeping the core UK GDPR framework intact. It clarifies existing rules, opens selected processing routes, and supports a wider data-sharing infrastructure, but it does not grant unrestricted access to personal information. The practical test is whether greater organisational flexibility can coexist with transparency, safeguards and credible regulatory enforcement for the people whose data creates that value.

From EU GDPR to UK GDPR

Modern UK data protection originated within a European framework. Directive 95/46/EC, adopted in 1995, required Member States to implement national rules, but differing national approaches proved difficult to reconcile with cross-border digital services. The EU responded with Regulation 2016/679, the General Data Protection Regulation, adopted on 27 April 2016. GDPR became directly applicable on 25 May 2018, replacing the Directive with harmonised rules intended to strengthen rights across the European Economic Area.

In the United Kingdom, GDPR arrived alongside the Data Protection Act 2018. The two instruments operated together: GDPR supplied the central principles, rights, and controller obligations, while the 2018 Act supplemented areas requiring domestic legislation, including exemptions, law-enforcement processing, and regulatory powers. Together, they marked a major strengthening of accountability, backed by substantially higher potential penalties and clearer rights regarding access, portability, erasure, and automated decision-making.

The Cambridge Analytica affair showed why stricter rules had become politically unavoidable. The ICO found serious failures involving Facebook users' information and imposed the maximum £500,000 penalty available at the time under the Data Protection Act 1998. The investigation exposed how ordinary social media activity could feed political profiling far beyond what users expected. GDPR was not created because of that case, but the scandal showed the scale of behavioural data use that older legislation could not govern.

Brexit changed the legal source of these rules rather than their substance overnight. The transition period ended on 31 December 2020, after which EU GDPR was retained and adapted into domestic law as the UK GDPR. From 1 January 2021, organisations operated under the UK GDPR alongside the 2018 Act, while the EU GDPR could still apply separately where activities fell within its territorial scope. Continuity mattered more than divergence at first.

That continuity became the platform for independent UK reform. Parliament can now amend domestic data rules without EU legislation dictating the outcome, although international consequences remain real. The European Commission renewed the United Kingdom's adequacy status on 19 December 2025, having assessed the DUAA amendments and concluded UK protection remained essentially equivalent. The relationship combines regulatory independence with practical interdependence: UK rules evolve on their own terms,

while cross-border flows still rest on comparable safeguards.

Why Data Protection Reform Was Considered Necessary

Pressure for reform grew because certainty of compliance and technological change did not always advance together. Organisations faced detailed documentation, complex interpretations of legitimate interests, uncertainty over secondary data use and caution around automated decision-making. Smaller organisations often felt these burdens most acutely, having fewer specialist legal and governance resources to draw on. Government policy sought to reduce friction where risks were genuinely manageable, while preserving stronger controls for sensitive processing and real harm.

Innovation created a second, related pressure. Artificial intelligence, digital identity, fraud detection, personalised services and large-scale research all depend on data being reusable at speed, sometimes for purposes not fully foreseeable when the information was first collected. A regime interpreted too defensively can discourage beneficial processing; one interpreted too loosely can normalise surveillance or discriminatory profiling. Reform therefore aimed to clarify reuse and create defined legitimate-interest routes without abandoning lawful-basis requirements or meaningful safeguards.

The Government attached substantial economic ambitions to reform. Its October 2024 announcement projected the legislation would add approximately £10 billion to the UK economy over ten years, freeing up around 1.5 million hours of police time and 140,000 hours of NHS staff time annually. Reduced police administration alone was estimated to save roughly £42.8 million each year, showing that data governance had come to be treated as economic infrastructure rather than a narrow compliance exercise.

The Data (Use and Access) Act 2025

The Data (Use and Access) Act 2025 received Royal Assent on 19 June 2025. Its significance reaches well beyond amendments to privacy law. The legislation creates a broad framework for improving how personal and non-personal data can be accessed, verified, shared and reused across commercial activity and public services. Data protection is therefore only one component of a wider programme to modernise digital infrastructure, reduce administrative friction and improve trusted information exchange across the economy.

Smart Data is a major part of that architecture. The Act gives the government powers to establish schemes that require specified data holders to securely share customer or business information with authorised recipients under defined conditions. Open Banking provides the established model, but the powers can extend similar arrangements into other sectors. Well-designed schemes can increase switching and competition because customers can direct regulated services to use information previously locked inside closed provider systems.

Digital verification services form another strand of the Act. It supports a statutory trust framework, certification arrangements, and a register of providers that meet prescribed standards. Registered providers may use a government-backed trust mark, while information gateways let public authorities disclose relevant information for identity or eligibility checks. The aim is to make proving identity, renting a home, or starting a job more efficient, without every organisation having to reproduce the same documentary checks.

The National Underground Asset Register shows how the Act also addresses infrastructure data. The United Kingdom holds around four million kilometres of buried pipes and cables, with excavation taking place roughly every seven seconds nationwide. Approximately one in 65 excavations causes an accidental asset strike, equating to some 60,000 incidents annually and an estimated cost of £2.4 billion. More than 600 asset owners previously held this information in inconsistent formats, creating avoidable complexity and safety risk.

Other provisions modernise public administration and information exchange more broadly. The Act supports the electronic registration of births and deaths in England and Wales and introduces measures that facilitate the use of data for law enforcement and national security

purposes. It also underpins common standards for health and care information systems. These reforms confirm the legislation is not a privacy statute alone: it is an enabling framework for digital infrastructure, public administration and trusted verification.

Within that wider framework sit targeted amendments to the UK GDPR, the Data Protection Act 2018 and the rules governing electronic communications. Changes cover recognised legitimate interests, scientific research, automated decision-making, subject access, complaints handling, international transfers and cookies. Commencement was phased rather than immediate: the ICO confirmed on 19 June 2026 that all DUAA provisions affecting data protection law were finally in force, completing the transition from legislative reform to an operating regulatory framework.

What the DUAA Changes – and What It Does Not

The most important boundary is what the DUAA does not do. It does not repeal the UK GDPR, replace the Data Protection Act 2018 or displace the Privacy and Electronic Communications Regulations 2003. Organisations remain subject to the familiar structure of principles, lawful bases, individual rights, security requirements and accountability obligations. Existing compliance programmes have not become obsolete; specific processes and legal tests require adjustment where Parliament has clarified, widened or modified established rules.

Some changes are nevertheless substantive. A new recognised-legitimate-interests basis permits specified processing without the balancing test used for ordinary legitimate interests, although necessity remains required throughout. The Act broadens the scope for significant automated decisions, clarifies reasonable and proportionate searches for subject access, adjusts aspects of purpose compatibility, introduces complaint-handling requirements, and relaxes consent for certain storage technologies. These reforms operate inside a system that continues to regulate processing rather than presume unrestricted discretion.

Individuals retain the central architecture of protection. Personal data must still be processed lawfully, fairly and transparently; rights of access, rectification, erasure, restriction and objection continue where legal conditions are met; and stronger controls remain for special-category information. Controllers must still secure information and demonstrate compliance. The reforms may alter when an organisation can act without consent, but absence of consent does not remove duties concerning necessity, transparency, rights handling or protection against unlawful access.

The Continuing Data Protection Principles

The foundational principles remain the organising logic of UK data protection. Lawfulness requires processing to rest on a valid legal basis and, where relevant, satisfy additional conditions for special-category or criminal-offence data. Fairness asks whether use would be unjustifiably detrimental, misleading or unexpected. Transparency requires organisations to explain processing intelligibly. Together, these principles prevent legal technicalities substituting for treating individuals predictably and openly whenever personal information is collected, analysed, shared or reused.

Purpose limitation requires data to be collected for specified, explicit and legitimate purposes, subject to rules governing compatible further use. Data minimisation requires that information be adequate, relevant, and limited to what is necessary. Accuracy requires steps to correct or erase inaccurate data, while storage limitations prevent identifiable information from being retained longer than necessary. Together these constrain organisational appetite: cheap storage or speculative future value does not justify collecting or retaining every data point indefinitely.

Accountability converts the principles from aspirations into demonstrable governance. Controllers must show why processing is lawful, how risks are controlled and how responsibilities are allocated through policies, records, contracts, assessments and oversight. The stakes remain substantial: infringements of the basic principles can attract the highest UK

GDPR tier, up to £17.5 million or 4% of total worldwide annual turnover, whichever is higher. Greater DUAA flexibility strengthens, rather than removes, the need for defensible decisions.

Lawful Bases for Processing Personal Data

Every use of personal data requires a lawful basis. The traditional UK GDPR framework offered six bases:

- *consent,*
- *contract,*
- *legal obligation,*
- *vital interests,*
- *public task, and*
- *legitimate interests.*

The DUAA has added recognised legitimate interest as a seventh basis, applying only to specified public-interest purposes. Selecting a basis is not an administrative formality: it determines why processing is lawful, which rights apply, and what justification an organisation must be able to demonstrate if challenged by a regulator, a court or the individual concerned.

Consent applies where an individual genuinely agrees to processing for a specified purpose. In contrast, contract applies where processing is objectively necessary to perform a contract with that person or take requested pre-contractual steps. An online retailer, for example, may use a customer's delivery address to fulfil an order under contract rather than consent. Asking permission would add little control, because refusing the necessary processing would make delivery impossible.

A legal obligation applies where processing is necessary to comply with statute or common law, rather than to a contractual requirement. Employers, for example, process salary and tax information because legislation requires them to report to HM Revenue & Customs. Vital interests are narrower and typically concern the protection of life: an unconscious casualty arriving at accident and emergency may have information shared urgently between clinicians because consent cannot realistically be obtained.

Public task supports processing necessary for functions carried out in the public interest or under official authority, provided those functions have a legal basis. Local authorities, government departments and NHS bodies rely on it for substantial areas of administration. Legitimate interests serve a different function for many private and third-sector organisations, permitting necessary processing in the organisation's legitimate interests unless the individual's interests, rights or freedoms override the organisation's justification.

The appropriate basis must be identified before processing begins and should reflect the real purpose rather than whichever option seems most convenient. Special-category data, such as health, biometric, religious or ethnicity information, requires an Article 6 lawful basis plus an additional Article 9 condition, and criminal-offence data has further restrictions. A lawful basis opens the legal gateway to processing; it does not, by itself, make excessive or inadequately explained processing compliant.

Consent – What Does Valid Consent Mean?

Consent remains the lawful basis most closely associated with personal choice, but UK GDPR sets a demanding standard. It must be freely given, specific, informed and unambiguous, resulting from a statement or a clear affirmative act. Silence, inactivity and pre-ticked boxes are insufficient. A person must understand who is seeking agreement, the purpose for which information will be used, and the processing activities involved, before any apparent permission can constitute valid consent.

Freely given consent requires genuine choice. An organisation should not make access to a

service conditional on agreeing to unrelated processing where that information is unnecessary for the service itself. Power imbalances matter too: employers and public authorities may struggle to rely on consent where refusal could appear disadvantageous. Consent should therefore be separated from general terms and, where distinct purposes exist, made granular enough for individuals to choose rather than accept one bundled proposition.

Valid consent must also be provable and reversible. Controllers should retain records showing who consented, when, how and what information they were given. Individuals must be told that consent can be withdrawn and should be able to withdraw it as easily as they gave it. Withdrawal does not retrospectively make earlier lawful processing unlawful, but processing dependent solely on that consent must normally stop unless another lawful ground genuinely applies.

ZMLUK Limited shows the consequences of treating nominal permission as meaningful consent. The ICO fined the company £105,000 in January 2026 after it sent more than 67 million marketing emails using third-party data. People registering on the source website saw 361 partner companies, with no mechanism to select which ones could contact them. The ICO concluded this could not produce informed, specific consent, showing that volume of disclosure is not the same as genuine choice.

Is Consent Becoming Less Important?

Consent is highly visible because websites, applications and marketing services repeatedly ask users to click acceptance buttons. That visibility has encouraged the misconception that data protection law requires permission before most personal information can be used. It does not. Consent is one of seven lawful bases under the amended UK GDPR, and an organisation should not select it merely because it appears reassuring; the correct basis depends on purpose, necessity and context.

Routine organisational functions often work better under another basis. An employer does not need an employee's consent to report taxable pay to HMRC where a legal obligation applies. A bank does not need consent to process information required to operate an account under contractual necessity. An NHS body may rely on public task for functions grounded in law. Treating consent as universal can obscure the real justification and imply a choice that does not exist.

Consent can also be operationally fragile because it may be withdrawn at any time. Where processing is genuinely necessary regardless of preference, building the activity around consent creates contradictions: an organisation cannot credibly tell an individual that processing is optional and then continue the same activity after permission is refused. The ICO therefore advises organisations to choose the lawful basis that genuinely applies, rather than defaulting to consent because it feels safer.

A better description is that consent is becoming more specialised rather than less important. It remains valuable for individuals deciding whether to opt in, and explicit consent serves additional functions for certain sensitive activities. Elsewhere, contract, legal obligation, public task, legitimate interests or recognised legitimate interest may more accurately describe why information is needed. Policy now favours honest allocation of lawful bases over presenting unavoidable processing as though individuals could freely refuse it.

Legitimate Interests

Legitimate interests are the most flexible general lawful basis because they are not confined to a predetermined activity. Commercial interests, fraud prevention, network security, administrative efficiency and wider societal benefits may all qualify. Flexibility does not create a presumption of lawfulness: Article 6 requires that processing be necessary for interests pursued by the controller or a third party, unless the individual's rights or freedoms override them, particularly where the individual is a child.

The ICO expresses this assessment through a three-part test. First, the purpose test identifies

a specific legitimate interest. Second, the necessity test asks whether the processing is targeted and proportionate, and whether a less intrusive method could achieve the objective. Third, the balancing test weighs reasonable expectations, information sensitivity, likely effects and safeguards. Organisations should document this reasoning through a legitimate interests assessment before processing begins.

Credit reference activity demonstrates why the basis can serve both commercial and societal purposes. Lenders share repayment information with credit reference agencies so future lenders can assess whether applicants are likely to repay. The lender has an interest in reducing bad debt, the applicant has an interest in responsible lending, and the wider market benefits from reliable risk assessment. Those benefits may support legitimate interests, but necessity and transparency still require separate consideration.

Experian illustrates the limits of treating legitimate interests as a convenient commercial gateway. Litigation arising from ICO enforcement concerned large-scale use of personal information for direct marketing. The Upper Tribunal dismissed the ICO's appeal in 2024, but the proceedings reinforced a critical distinction: legitimate interests may support commercial processing where the required assessment succeeds, yet transparency and fairness remain independent obligations rather than optional consequences of finding a profitable purpose.

Recognised Legitimate Interests – A New Lawful Basis

The DUAA introduced recognised legitimate interest as a distinct seventh Article 6 lawful basis. It is narrower than ordinary legitimate interests because Parliament has predetermined that certain purposes carry sufficient public value to justify processing without the conventional balancing test. The basis is not a licence for businesses to declare their own interests recognised; processing must fall within a statutory condition in the new Annex 1 to the UK GDPR and remain necessary.

Five categories are covered. They concern disclosure requested for another controller's legally grounded public task or official function; national security, public security and defence; responding to emergencies; preventing, detecting or investigating crime, including apprehending or prosecuting offenders; and safeguarding vulnerable individuals. The change removes hesitation where socially valuable processing previously required a careful balance between organisational interests and individual rights, potentially delaying urgent or protective decisions.

The crime condition can support the necessary use of information where fraud, theft, or other offending is suspected, while safeguarding can assist organisations in protecting children or adults who meet the statutory vulnerability criteria. Emergency processing may apply where circumstances threaten serious harm to people, property or the environment. These routes recognise that requiring a conventional legitimate-interests balance during urgent events could add delay without materially improving protection, provided processing remains genuinely necessary.

The public-task disclosure condition is carefully framed. A controller may disclose information when another controller requests it, provided the recipient needs the data to perform a legally grounded public task or to exercise official authority. The requesting public authority cannot simply invoke a recognised legitimate interest for its own official processing; its basis will ordinarily be public task. The new route instead helps the information holder justify a necessary disclosure in response.

The removal of the balancing test does not remove the UK GDPR. Purpose limitation, minimisation, transparency, accuracy, security, accountability and applicable individual rights all continue. The right to object also applies to processing under new Article 6(1)(ea). Government guidance explains that the reform gives non-public bodies greater confidence in socially valuable processing, not unrestricted access: recognised legitimate interest removes one layer of assessment while leaving necessity and safeguards intact.

Purpose Limitation – Why Data Was Collected Matters

Purpose limitation requires organisations to decide why personal information is being collected before processing begins, and to keep later use within those stated purposes. Article 5(1)(b) requires purposes to be specified, explicit and legitimate, preventing information gathered for one reason becoming a general asset available for any later opportunity. The principle supports predictability: individuals should not discover that data supplied for an ordinary service has been used for unrelated profiling or marketing.

The restriction is not absolute, because organisations often need information to serve connected purposes over time. A retailer may use transaction records for fulfilment, accounting, fraud prevention, and customer service, provided those activities are properly justified and communicated. Difficulty arises when a proposed use materially changes the relationship. Information gathered to provide healthcare, employment, education or financial services can reveal far more than the immediate transaction, making unexpected secondary exploitation particularly intrusive.

The Royal Free London NHS Foundation Trust illustrates the danger. In 2017, the ICO concluded that approximately 1.6 million partial patient records had been processed by DeepMind during clinical safety testing of the Streams application without patients being adequately informed. The regulator identified failures under several then-applicable data protection principles. The case showed that an apparently beneficial healthcare objective does not remove the need to define, communicate and justify how existing records are repurposed.

Purpose limitation therefore works alongside transparency rather than merely restricting databases. Privacy information should explain purposes with enough precision for people to understand the intended use, while internal records provide corresponding governance. Broad statements, such as improving services or supporting business purposes, may be insufficient if they conceal materially different processing. The principle encourages organisations to design data flows around defined needs rather than collecting information first and then searching for convenient uses.

Reusing Personal Data for New Purposes

When an organisation wants to use personal information for a new purpose, it must first determine whether the proposed processing is compatible with the purpose for which the information was collected. Compatibility is a legal assessment rather than a judgement that the new activity appears useful. If the new purpose falls within one of the statutory compatible categories, specific treatment may apply; otherwise, the controller must assess the relationship between the original and proposed uses.

The compatibility assessment considers several factors: the links between the original and new purposes, the circumstances in which the information was collected, the relationship with the individual, reasonable expectations, the nature of the data, possible consequences, and safeguards such as encryption or pseudonymisation. A closely related, low-impact reuse is easier to justify than an unexpected activity involving sensitive information, behavioural profiling or disclosure to an organisation with which the individual has no existing relationship.

A compatible new purpose still requires a lawful basis. The original basis may continue to support the new activity, but organisations must confirm that it remains appropriate. Where compatibility cannot be demonstrated, further processing may require fresh consent or another lawful route, where permitted by law. The DUAA makes this structure clearer by expressly addressing further processing, but compatibility and lawfulness remain separate questions rather than alternative ways of legitimising the same activity.

The DUAA and Assumed Compatibility

The DUAA creates greater certainty by identifying reuses treated as compatible with the original purpose. Schedule 5 inserted Annex 2 into the UK GDPR, covering defined public-interest circumstances. Where the statutory conditions are met, an organisation does not

undertake the ordinary compatibility assessment. The change addresses situations where socially necessary reuse previously required uncertain case-by-case analysis, while preserving separate requirements for a lawful basis, necessity and compliance with wider principles.

The list includes necessary reuse for public security, emergencies, crime prevention, investigation and prosecution, protection of vital interests, safeguarding vulnerable people, taxation and compliance with legal obligations. It also addresses qualifying disclosures requested for public tasks or official authority and certain disclosures for archiving in the public interest. These categories are specific: a controller cannot label an ordinary commercial initiative as a public-interest initiative merely because the resulting service produces wider social benefit.

Consent receives additional protection where it supported the original collection. If information was originally obtained through consent, an organisation relying on an Annex 2 compatibility condition must consider whether it would be reasonable to obtain consent for the new use instead. This prevents statutory compatibility routinely overriding an earlier relationship based on individual choice, recognising that changing purposes after consent-based collection can undermine the expectations that made the initial agreement meaningful.

The crime provisions show the practical effect. Information collected during ordinary commercial activity may later reveal suspected fraud or other offending, and necessary reuse for detecting, investigating or preventing crime can be treated as compatible. The same concept applies to fraud, money laundering and terrorist financing. This reduces uncertainty where delay could frustrate legitimate investigation, but it does not permit disproportionate trawling through customer information merely because criminal behaviour is theoretically possible.

Data Sharing and the Public Interest

Data sharing can serve substantial public interests when organisations hold information that public authorities need to perform legally grounded functions. The DUAA facilitates voluntary disclosure where a requesting body needs personal information for a public task or official authority. If no statutory power compels disclosure, the organisation receiving the request may rely on recognised legitimate interest. The requesting authority itself will ordinarily rely on public task for its subsequent processing.

The request should identify the relevant task or official function and state that the information is needed for that purpose. Necessity remains important: public status does not create an unrestricted entitlement to private-sector databases. Where an authority already possesses statutory information-gathering powers, those powers should generally continue to be used, with the disclosing organisation likely relying on legal obligation instead. The DUAA primarily helps where legitimate public functions depend on voluntary rather than compulsory sharing.

The National Fraud Initiative demonstrates the value of lawful data matching at scale. More than 1,100 public- and private-sector organisations supply datasets for comparison to identify anomalies indicating fraud or error. Between April 2022 and March 2024, the exercise prevented, detected or recovered £510.1 million across the United Kingdom, its best result on record, including £477.5 million in England alone. The Public Sector Fraud Authority puts the initiative's cumulative outcomes since 1996 at £2.9 billion.

Those benefits do not mean every data match is justified. Responsible sharing requires a defined objective, appropriate lawful bases, minimised datasets, secure transmission, retention controls and clarity about responsibilities. Data-sharing agreements can record purpose, roles, information standards, security and review arrangements even where legislation does not mandate a document. Strong governance matters because matching can produce false positives, and an apparent anomaly should trigger proportionate investigation rather than an automatic adverse conclusion.

Scientific Research and Secondary Data Use

The DUAA clarifies that scientific research can include commercial research, technological development or demonstration, fundamental or applied research, and qualifying public health research. The activities must still reasonably be described as scientific, preventing the label becoming a device for ordinary product analytics or marketing. The clarification matters because modern research frequently crosses universities, charities, technology companies and pharmaceutical businesses, making distinctions based solely on profit motive increasingly artificial.

Research can also involve broad consent, where the exact future study cannot be identified when information is collected. Under the amended UK GDPR, a person may consent to an area of scientific research if defining the precise purpose is initially impossible, the approach is consistent with recognised ethical standards, and participants can consent only to parts of the research where appropriate. Broad consent permits realistic scientific uncertainty without becoming unrestricted permission for any future use.

UK Biobank illustrates the scale of secondary research use. It recruited 500,000 participants between 2006 and 2010 and combines genetic, lifestyle and health information for approved researchers from academia and industry. By 2026, more than 22,000 researchers were using the resource worldwide. Earlier genome sequencing work attracted £200 million of funding, including £50 million each from UK Research and Innovation and Wellcome and £100 million from four pharmaceutical and healthcare companies.

Research flexibility is conditional on safeguards. Processing must remain necessary for scientific, historical, statistical or public-interest archiving purposes and comply with Article 84B. Safeguards include measures designed to protect individuals, with pseudonymisation relevant wherever those purposes can be achieved that way. Research exemptions and compatibility rules therefore recognise long-term societal value while preventing research status from becoming a blanket justification for retaining identifiable information indefinitely.

Special Category and Sensitive Personal Data

Some personal information receives enhanced protection because misuse can expose individuals to discrimination, stigma, exclusion or deeply intrusive inference. Article 9 special-category data covers racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data used for unique identification, health information, sex life and sexual orientation. The category concerns both explicit facts and, in some circumstances, inferences deliberately drawn from apparently ordinary behavioural information.

An Article 6 lawful basis is not enough on its own. A controller processing special-category data must also identify an Article 9 condition, such as explicit consent, employment and social-protection law, vital interests, healthcare, substantial public interest or qualifying research. Some conditions also depend on Schedule 1 to the Data Protection Act 2018 and may require an appropriate policy document. The additional gateway reflects the greater consequences that disclosure or misuse can create.

Biometric information illustrates the distinction. A photograph is not automatically special-category data merely because a face is visible. If technology processes facial measurements to uniquely identify a person, the resulting biometric processing can fall within Article 9. Similar care applies to inferred information: where profiling deliberately infers ethnicity, health status, political opinion, religion or sexual orientation and uses that inference, the organisation may be processing special-category information even where predictions are uncertain.

The 23andMe cyber-attack demonstrates the stakes. In June 2025, the ICO fined the genetic testing company £2.31 million after finding inadequate security measures affecting 155,592 UK users, following a joint investigation with Canada's privacy regulator. Information that was potentially exposed included names, birth years, locations, profile images, ethnicity, family

trees, and health reports. Genetic and health information can remain sensitive throughout a person's life, making compromise materially different from losing replaceable account credentials.

The Right to Be Informed

The right to be informed gives practical effect to the UK GDPR's transparency principle. Organisations must clearly explain when they collect or use personal data, why it is needed, the lawful basis relied upon, how long it will be retained, who may receive it, and what rights individuals can exercise. Transparency is not satisfied by publishing dense legal wording somewhere on a website; information must be concise, intelligible, accessible and written in clear language.

Where information is collected directly from an individual, privacy information should normally be supplied when the data is obtained. Where it comes from another source, the organisation generally has up to one month to provide the required information, subject to limited exceptions. Privacy notices should also identify overseas transfers, relevant safeguards, the data source where applicable, complaint rights and the existence of qualifying automated decision-making or profiling.

Transparency must continue after the initial collection. If an organisation develops a materially new use for existing information, it should bring that use to the individual's attention before processing begins. Layered notices, dashboards and just-in-time explanations can be more effective than a single lengthy document, particularly where digital services collect information incrementally. Good privacy design therefore concerns when and how explanations are delivered, not merely whether statutory wording appears somewhere.

TikTok demonstrates the consequences of weak transparency involving vulnerable users. In April 2023, the ICO fined TikTok £12.7 million for breaches of UK data protection law involving children's information, estimating that around 1.4 million UK under-13s had used the platform without parental consent. The penalty found infringements of Articles 8, 12, 13 and 5(1)(a), linking age-related lawfulness with transparency obligations. Online services cannot treat privacy information as a peripheral notice where children's data is processed.

Subject Access Requests

A subject access request, commonly called an SAR, allows an individual to ask whether an organisation is processing their personal data and to obtain a copy of it. The right also includes supplementary information about purposes, categories, recipients, retention, rights and certain automated decision-making. A request does not require legal terminology, a prescribed form or professional representation; it may be made verbally, in writing or through channels such as social media.

Organisations must usually respond without undue delay and within one month. They may extend the period by up to two further months if the request is complex or numerous, provided the individual is appropriately informed. In most cases no fee can be charged, although a reasonable administrative fee may apply to manifestly unfounded or excessive requests or additional copies. Identity checks are permitted where necessary but should remain proportionate to the risk of wrongful disclosure.

The operational challenge is often retrieval rather than legal interpretation. Relevant information may be stored across email accounts, customer systems, personnel files, archived databases, messaging platforms, and records held by processors. Organisations need systems capable of locating information efficiently, separating the requester's data from third-party information and applying exemptions correctly. Poor information architecture can turn a routine legal right into a resource-intensive exercise spanning legal, IT and operational teams.

South Wales Police illustrates the consequences of persistent delay. The ICO found that between April 2023 and March 2024 only 29% of its SARs were answered within the statutory timeframe. By August 2025, 352 requests were overdue, one approaching two years old. An

enforcement notice required the force to clear the backlog by June 2026, showing that staffing or workload pressures do not remove the obligation to make access rights effective in practice.

Reasonable and Proportionate Searches

The DUAA expressly confirms that organisations responding to subject access requests need only undertake searches that are reasonable and proportionate. This wording codifies an approach already established through case law and regulatory practice. It does not allow controllers to ignore inconvenient systems or conduct token searches; rather, it recognises that the right of access requires serious efforts to locate relevant information, without demanding searches whose burden is clearly disproportionate.

Reasonableness depends on context. A focused request concerning a defined transaction may justify searches of specific systems, mailboxes or date ranges, while a broad request covering many years may require a wider exercise. Relevant factors can include the quantity of information held, how records are organised, the age of the material, the resources required to retrieve it and the likelihood that further searches would locate additional personal data within scope.

The reform matters particularly for organisations with fragmented or historic records. A controller may hold millions of emails, legacy databases and backups that cannot all sensibly be searched manually for every request. Poor record management should not, however, automatically become a justification for doing less. Organisations should be able to explain their search methodology, the systems examined, the terms used, and why further steps would have been disproportionate if later challenged.

Reasonable and proportionate does not mean individuals can be forced to narrow requests. ICO guidance states that an organisation may request clarification where reasonably required, but the requester may decline to narrow the scope. If the request remains broad, the controller must still carry out a reasonable search. The statutory change therefore moderates the burden of retrieval while preserving genuine entitlement to access rather than converting clarification into a precondition.

The New 'Stop the Clock' Rule

The DUAA introduces a formal ability to pause, or "stop the clock" on, the response period for certain information-rights requests where clarification or additional information is reasonably required. For subject access, the one-month period pauses on the day the organisation asks for clarification and resumes the day after the individual provides it. The change prevents controllers losing response time while genuinely necessary questions remain unanswered.

The power is not a general mechanism for extending difficult requests. Clarification must be reasonably required to identify the information or processing activities being requested, and organisations should ask promptly rather than waiting until the deadline approaches. They also cannot force a person to narrow a broad but intelligible SAR. If the individual confirms the original request unchanged, the organisation must resume the timetable and undertake a reasonable, proportionate search.

Identity verification operates differently but has a similar practical effect. Where an organisation reasonably needs information to confirm the requester's identity, the response period does not begin until that information is received. The requirement should be proportionate: requesting passports or driving licences unnecessarily can itself create privacy risk. Where an employer already recognises a worker through a corporate email account, demanding extensive additional identification may be hard to justify.

For governance teams, the reform makes accurate chronology essential. Case records should show when the request arrived, when clarification was sought and why, when the requester replied, and when the clock resumed. Automated workflow systems can help calculate deadlines, but staff still need legal judgement about whether pausing is justified. Misusing stop-the-clock provisions as routine delay tactics could undermine access rights and expose

the organisation to complaints.

Rectification, Erasure and Restriction

The right to rectification allows individuals to require that inaccurate personal data be corrected and that incomplete information be completed where appropriate. Accuracy is contextual: an organisation may need to distinguish an incorrect fact from a recorded opinion. Once challenged, the controller should take reasonable steps to examine evidence and determine whether amendment is required. Rectification reinforces the accuracy principle by providing a mechanism for individuals to challenge information affecting them.

Erasure, often called the right to be forgotten, applies only in specified circumstances, not as a universal right to deletion. It may arise where information is no longer necessary, consent is withdrawn, and no other lawful basis applies, processing is unlawful, or an objection succeeds. Exceptions remain for freedom of expression, legal obligations, public-interest tasks, certain health purposes, qualifying research or archiving, and the establishment, exercise or defence of legal claims.

Restriction provides an intermediate remedy where deletion would be premature or inappropriate. An individual may require that processing be limited, for example while accuracy is checked, where processing is unlawful but deletion is opposed, or where an objection is being assessed. During restriction, the organisation may generally retain the information but should not actively use it except in defined circumstances, so systems must support practical suppression rather than policy statements alone.

The ICO reprimanded the Labour Party in August 2024 for failures involving both subject access and erasure requests. The action showed that individual rights cannot be treated as isolated administrative correspondence. Organisations need coordinated processes capable of identifying which right has been exercised, locating associated records, applying exemptions consistently and implementing outcomes across operational systems, since a deletion that leaves duplicate records active elsewhere fails the right in practice.

The Right to Object

The right to object gives individuals a mechanism to challenge processing carried out for public tasks, official authority, legitimate interests and recognised legitimate interests, as well as certain research or statistical purposes. In most of these situations the right is not absolute. An organisation may continue where it demonstrates compelling legitimate grounds that override the individual's interests, rights and freedoms, or where processing is required to establish, exercise or defend legal claims.

Direct marketing is different. Individuals have an absolute right to object to the use of their personal data for direct marketing, including related profiling, and once an objection is received the organisation cannot continue that processing. Deleting every record is usually unnecessary and can be counterproductive; retaining minimal information on a suppression list is often appropriate because it prevents the person being accidentally re-added through later imports or refreshed databases.

The right applies across marketing channels and must be brought clearly to people's attention. Electronic marketing can also trigger PECR requirements governing consent, soft opt-ins and communications methods. In January 2026, the ICO fined Allay Claims Limited £120,000 after more than 4 million unlawful marketing texts, and ZMLUK Limited £105,000 after more than 67 million unlawful emails, illustrating the financial consequences of disregarding people's communication choices at scale.

O'Carroll v Meta has also highlighted the relevance of Article 21 to personalised advertising. In March 2025, the ICO stated that targeted online advertising should be treated as direct marketing and that people must have a clear way to object to their information being used for it. The dispute shows how an apparently simple right becomes strategically important once business models depend on behavioural profiling and advertising revenue rather than

occasional promotional messages.

Data Portability and Personal Control

Data portability allows individuals to obtain personal data they have provided to an organisation in a structured, commonly used and machine-readable format and, where feasible, have it transmitted directly to another controller. The right is narrower than the subject access right. It generally applies to automated processing based on consent or contract, and to information provided by the individual, including observed data from their use of a service, rather than to every inference the organisation creates.

Formats such as CSV, XML and JSON can satisfy portability requirements because software can process their structure; the objective is practical reuse, not an inaccessible document dump. Controllers should not create legal, technical or financial barriers to transfer, although UK GDPR does not require organisations to redesign systems for complete interoperability. Secure transmission remains essential, because portability increases the movement of information and can have serious consequences if data reaches the wrong recipient.

Portability gives control greater economic significance because information can help individuals switch providers or use services that analyse their behaviour. Transaction histories, consumption records and account information may support price comparisons, budgeting tools or alternative service offers. The right, therefore, differs from privacy protections that merely limit organisational activity: it creates positive agency, allowing the individual to deploy their data elsewhere rather than leaving it locked within the original provider's systems.

Open Banking shows what systematic portability can achieve beyond individual requests. Open Banking Limited reported that by June 2026 the ecosystem had recorded 18.81 million active user connections and 2.81 billion API calls that month alone, with the sector crossing 100 billion cumulative API calls and one billion cumulative payments in July 2026. Standardised interfaces enable authorised services to access data securely with customer permission, turning portability from an occasional compliance exercise into a genuine competitive infrastructure.

The DUAA provides powers for further Smart Data schemes, potentially extending similar models beyond banking into energy and retail. This represents a broader concept of control: individuals and businesses may be empowered not merely to obtain a copy of information but to direct trusted, standardised flows between providers. Success depends on authentication, security and clear permissions; poorly designed schemes could increase exposure, while good interoperability can reduce switching costs and weaken proprietary data silos.

Complaining About the Use of Personal Data

The DUAA creates an explicit statutory duty for controllers to maintain a process through which people can make data protection complaints directly. Organisations must accept complaints through appropriate channels, acknowledge receipt within 30 days and, without undue delay, take suitable steps to investigate, keep complainants informed and communicate the outcome. The duty became operational in June 2026 and applies across all organisations that handle personal data, not only to sectors with existing formal complaints schemes.

An organisation may provide an electronic form, a dedicated email address, a telephone route, a portal, live chat or another suitable mechanism, and may integrate data protection complaints into an existing complaints system. Individuals are not required to use the preferred route, however, so frontline employees need to recognise complaints received elsewhere. Published procedures should explain how concerns can be raised, what supporting information may help, and how the organisation will investigate and report back.

The scale of complaints demonstrates why internal resolution matters. In February 2026, the Information Commissioner reported that the ICO had already received around 66,000 data protection complaints during 2025/26, compared with more than 40,000 during 2024/25, and that the annual total could reach 75,000. Rising volumes strengthen the case for organisations

to resolve legitimate concerns themselves rather than treat the regulator as the first practical route to redress.

The new duty encourages problems to be resolved closer to their source. ICO guidance states that in most cases, individuals who complain to the regulator will first be asked to raise the matter with the organisation. Effective internal investigation can identify missing disclosures, incorrect data, poor explanations, or process failures before escalation, so complaints should be treated as governance intelligence capable of revealing recurring weaknesses, rather than correspondence to be closed defensively within a target period.

Automated Decision-Making Before the DUAA

Before the DUAA, Article 22 of the UK GDPR imposed a general restriction on decisions based solely on automated processing, including profiling, that produce legal or similarly significant effects. Examples included an automatic refusal of an online credit application or an entirely automated recruitment decision. The rule targeted consequential decisions made without meaningful human involvement, rather than every use of algorithms, analytics or decision-support tools within an organisation.

Such decisions were permitted only in limited circumstances: where necessary for entering into or performing a contract, where authorised by UK law with appropriate safeguards, or where based on the individual's explicit consent. Additional restrictions applied to special-category information. The architecture therefore began from prohibition and then created exceptions, requiring controllers to determine both whether a process was solely automated and whether its effect crossed the statutory significance threshold.

The 2021 Amsterdam proceedings involving British and other Uber drivers illustrated the factual difficulty. Drivers argued that fraud-related account deactivations were automated decisions. Uber described an investigation by its EMEA Operational Risk team, and the court found the deactivations were not based solely on automated processing, although temporary app blocking after a fraud signal did occur automatically. The dispute showed Article 22 turned on genuine human involvement rather than mere algorithmic use.

The pre-DUAA structure was criticised for its uncertainty, because organisations struggled to determine whether innovative automated services fell within the prohibition. That caution could discourage useful automation or encourage artificial human checkpoints designed mainly to avoid Article 22. Government reform therefore targeted the legal architecture rather than abandoning safeguards. Automation was never unlawful in general; only significant decisions made solely by automation were restricted unless a specified exception applied.

The New Automated Decision-Making Framework

The DUAA replaces the former Article 22 structure with new Articles 22A to 22D, creating a more permissive framework for significant decisions based solely on automated processing. For ordinary personal data, organisations can now make such decisions in wider circumstances where they have an appropriate lawful basis, including legitimate interests, provided statutory safeguards apply. The reform moves away from general prohibition with narrow exceptions towards regulated permission accompanied by enforceable protections.

A decision is solely automated where there is no meaningful human involvement, and significance remains central because routine automation without legal or similarly significant effects does not trigger the special framework. Relevant examples may include automated credit decisions, recruitment screening or account actions that materially affect employment or access to services. Controllers must still satisfy ordinary UK GDPR requirements, including lawfulness, fairness, transparency, minimisation, and, where high risks arise, an appropriate impact assessment.

Special-category data remains more constrained. The broader permission for significant solely automated decisions does not simply authorise organisations to use health, ethnicity, religion, biometric or sexual orientation information under ordinary legitimate interests. Stronger

conditions continue to apply, reflecting the heightened discrimination and privacy risks associated with sensitive characteristics. The distinction prevents the new framework becoming a blanket legal basis for high-impact AI decisions built on particularly intrusive or protected information.

What Counts as Meaningful Human Involvement?

Meaningful human involvement exists where a person genuinely participates in taking a decision rather than simply appearing somewhere in the process. Under the DUAA, a decision is deemed solely automated when there is no meaningful human involvement, and organisations must consider the extent to which profiling contributed to the outcome. The legal focus, therefore, rests on substance: whether a human can understand, question, alter, or reject the recommendation before the significant decision is finalised.

A reviewer who routinely accepts an algorithmic score without examining its reasoning is unlikely to provide meaningful involvement. Genuine intervention requires authority, competence and sufficient information to reach a different conclusion. The reviewer should weigh the recommendation against other relevant evidence, understand the factors influencing it, and remain free from targets or workflow pressures that effectively require approval. Human involvement should change the decision-making process, not merely add a signature or final click.

Recruitment provides a practical example. The ICO's 2026 Recruitment Rewired work found that many employers using automated recruitment were likely making decisions solely through automation, with meaningful human involvement absent. Where employers introduced human review, the ICO stressed it should be applied consistently to candidates at the same hiring stage, since apparently protective intervention can itself create unfairness by giving some applicants human reconsideration. In contrast, others remain governed entirely by scoring.

Meaningful involvement also requires appropriate timing. A person who reviews an outcome only after rejection has already taken effect may be providing a safeguard or appeal rather than participating in the original decision. Organisations should map the point at which significant effects occur. If an automated system rejects a loan, job application or benefit claim before any genuine human assessment, the original decision remains solely automated, even if a later appeal permits reconsideration.

Governance should make human responsibility observable. Procedures can identify who reviews outputs, what training they receive, what evidence they consider, when they may depart from recommendations and how disagreements are recorded. Monitoring should test whether reviewers actually overturn automated outputs where justified. A process in which human decisions never differ from algorithmic recommendations may indicate exceptional model accuracy, or, more concerningly, automation bias that has converted nominal oversight into routine endorsement.

The Right to Challenge Automated Decisions

The DUAA requires safeguards whenever a significant decision is based solely on automated processing. Controllers must provide information about the decision, enable the individual to make representations, allow that person to obtain human intervention and provide a means of contesting the outcome. These safeguards recognise that automation can make decisions quickly and consistently, but speed cannot eliminate opportunities to correct inaccurate data, misunderstood circumstances, model errors, or unanticipated consequences.

Information should be useful enough to support a challenge, rather than merely stating that automation occurred. The individual should understand that a significant automated decision has been taken, the context in which the system operated, and how to exercise available rights. ICO guidance on explaining AI distinguishes rationale, responsibility, data, fairness, safety and impact explanations. Not every case requires every category, but explanations should let the affected person identify why reconsideration may be justified.

Human intervention should involve a fresh, competent assessment rather than a repetition of the automated result. The reviewer needs authority to change the outcome, access to relevant information and an understanding of the system's limitations. Where a job applicant challenges an automated rejection, for example, the reviewer should examine the applicant's evidence and the model's reasoning, rather than confirm that the score fell below a predetermined threshold.

The safeguards do not guarantee that every challenge will succeed; they ensure instead that consequential automated decisions remain contestable by the person affected. Organisations should establish accessible routes, record representations, set review standards and communicate outcomes clearly. Where automation affects employment, credit, insurance, or housing, a weak appeals mechanism can turn technical efficiency into procedural unfairness, so effective challenge rights are part of the governance needed to justify broader automation.

Special Category Data and Automated Decisions

Special-category information remains subject to stronger restrictions because automated use can amplify harms associated with health, ethnicity, religion, political opinions, trade union membership, genetic information, biometric identification and sexual orientation. The DUAA liberalised significant automated decision-making for ordinary personal data, but deliberately retained tighter conditions where special-category information contributes to the decision, distinguishing between expanding automation generally and allowing algorithms to determine outcomes based on particularly sensitive characteristics.

A significant automated decision based entirely or partly on special-category data is permitted only through limited routes. One is explicit consent. The alternative requires that the decision be necessary for entering into or performing a contract, or be required or authorised by law, together with processing necessary for reasons of substantial public interest under Article 9(2)(g). The recognised legitimate interests basis cannot be used for such decisions, thereby preserving an additional statutory boundary.

The distinction matters because machine-learning systems can infer sensitive characteristics even where organisations never directly request them. The ICO's audits of AI recruitment providers found examples of systems inferring gender and ethnicity from candidates' names, prompting almost 300 recommendations across the audited providers on fairness, minimisation, transparency, and bias, all of which were accepted or partially accepted. Such inferences can create special-category issues when deliberately used to evaluate identifiable individuals.

Organisations should identify sensitive inputs and inferred attributes before deployment, not after an adverse decision is challenged. Data protection impact assessments should examine whether special-category information is necessary, whether proxies reproduce the same effects, and whether less intrusive alternatives exist. Technical teams also need controls to prevent models from quietly incorporating sensitive variables during retraining, since broader automation is only defensible where sensitive-data safeguards genuinely hold in practice.

Artificial Intelligence, Profiling and Algorithmic Control

Artificial intelligence and profiling now influence decisions across employment, lending, insurance, fraud prevention, marketing, housing and public administration. Recruitment systems can rank applicants, lenders can estimate default risk, insurers can price policies, retailers can segment customers and landlords can prioritise arrears or fraud cases. These applications can improve speed and consistency, yet they also concentrate power in models whose outputs may reflect incomplete data, historical patterns or assumptions individuals cannot easily inspect.

Public-sector adoption demonstrates the scale of the shift. By September 2026, the Government's Algorithmic Transparency Recording Standard repository contained 143 published records describing tools used by public organisations. The Department for Work

and Pensions' Universal Credit Advances model, for example, risk-assesses advance requests for fraud; in 2024/25, 1.4 million advances worth £0.8 billion were issued, while the model identified fraud risk around three times more effectively than a random control sample.

Commercial uses raise different incentives. Credit scoring may expand lending decisions, insurance models can refine risk selection, fraud systems can block suspicious transactions, and customer profiling can improve targeting. Housing providers may use analytics to prioritise interventions for arrears, repairs, tenancy fraud or vulnerability. The same efficiency becomes problematic where a risk score silently determines access to essential opportunities, so accuracy, explainability, bias testing and meaningful review matter wherever automated classifications carry material consequences.

Algorithmic control should therefore be assessed as a governance issue rather than simply a technology procurement. Controllers need to understand training data, model objectives, error rates, protected-characteristic impacts, human roles and routes for challenge. Central government requires qualifying algorithmic tools with public impact to be recorded under the transparency standard. Comparable discipline is valuable in private organisations, because outsourcing an AI system does not outsource responsibility for decisions made with personal data.

Does AI Change the Meaning of Consent?

Artificial intelligence complicates consent because future uses of personal information may be hard to describe at the time data is collected. Conventional consent assumes an individual can understand a specific purpose and make an informed choice. Machine-learning systems may later discover correlations or infer characteristics that neither the organisation nor the individual anticipated. A broad statement agreeing to "AI use" therefore risks becoming too vague to satisfy the requirements of specific, informed consent.

Complexity does not lower the legal standard. Where an organisation relies on consent, it must explain the relevant processing in accessible terms, separate optional uses where appropriate and allow withdrawal. ICO guidance on AI transparency emphasises explaining the purposes, retention, sharing, and use of automated systems. An organisation need not reveal source code, but should give enough information for individuals to understand what data is used and why AI is involved.

Inference creates a deeper control problem because AI can generate information the individual never consciously supplied. Browsing, purchasing, location or interaction data may be used to predict interests, reliability, health or other characteristics. Consent to the original collection does not automatically amount to informed agreement to every inference later drawn from it. Purpose limitation, fairness, lawful basis and special-category rules therefore remain important even where an earlier consent interaction can be pointed to.

Consent can also become illusory where refusing AI processing means losing access to an important service or opportunity. Employment candidates, tenants, borrowers or insurance customers may have little negotiating power if automated assessment is embedded within the standard process. In those circumstances, another lawful basis may be more appropriate than pretending that participation is voluntary, forcing the organisation to justify processing openly under the basis that actually supports the activity.

AI therefore changes the practical context of consent more than its legal definition. Meaningful agreement remains possible for optional, well-defined processing, but it becomes less convincing as systems grow more complex, adaptive and consequential. Individual control must consequently rely on a broader package of safeguards: transparent purposes, minimised data, explainable outcomes, objections where applicable, human intervention and effective complaints. Consent remains valuable, but it cannot carry the entire burden of governing data-intensive automated systems.

Children's Data and Online Services

Children receive particular protection because they may be less able to understand how data collection, profiling and design choices affect them. The DUAA strengthens this approach by requiring providers of online services likely to be used by children to consider how technical and organisational measures can best protect and support them. Providers must recognise that children merit specific protection and that their needs differ according to age and developmental stage when designing compliant services.

The requirement builds on the ICO's Children's Code, which sets out 15 standards for online services likely to be accessed by under-18s. These include treating the child's best interests as a primary consideration, high-privacy defaults, data minimisation, restrictions on unnecessary sharing, geolocation controls and limits on nudge techniques that encourage children to weaken privacy settings. A service cannot avoid these expectations merely by stating in its terms that it was designed primarily for adults.

Reddit demonstrates the enforcement consequences. In February 2026, the ICO fined Reddit £14.47 million after finding it had failed to apply robust age assurance and had no lawful basis for processing information relating to under-13s. A smaller platform, Imgur's owner MediaLab, was fined £247,590 for near-identical failings weeks earlier. Reddit appealed its penalty in April 2026, but the actions confirm that platforms cannot rely on stated age restrictions without effective measures to back them up.

Age assurance must remain proportionate, as verifying age may require additional personal information. Providers may apply Children's Code protections to all users or use risk-based age assurance to tailor safeguards, depending on the service and potential harm. The objective is not maximum identification of every child, but a reliable enough age assessment to deliver suitable protection without creating an unnecessarily intrusive identity database that introduces new security and privacy risks.

Cookies and Tracking Technologies

Cookies are small files stored on a user's device, but the regulatory concept is wider than cookies alone. PECR governs technologies that store information on, or access information from, terminal equipment, including tracking pixels, local storage, device fingerprinting and some software development kits. The UK GDPR applies alongside PECR when the resulting information constitutes personal data, so organisations need to consider both access to the device and subsequent processing of identifiable information.

Historically, Regulation 6 of PECR prohibited storing or accessing information on a user's device unless the user received clear information and gave consent, subject to limited exceptions. Consent had to meet the UK GDPR standard: freely given, specific, informed and an unambiguous affirmative act. Strictly necessary technologies did not require consent, allowing functions such as maintaining shopping baskets, authenticating logged-in users or providing security features essential to a requested service.

The distinction became commercially important because advertising and behavioural tracking were generally not strictly necessary to provide a website. Technologies capable of tracking users across pages or services could build detailed profiles for targeted advertising, attribution or audience measurement, often involving multiple organisations and opaque data flows. Consent banners therefore became the visible mechanism by which websites attempted to comply with PECR before placing non-essential advertising and analytics technologies on devices.

The DUAA changes this framework without removing PECR's basic prohibition. Storage or access remains unlawful unless consent or another statutory exception applies; what has changed is the range of circumstances in which consent is unnecessary. Organisations must therefore classify technologies by purpose rather than treating every cookie identically, since advertising trackers, essential authentication tools, statistical analytics and appearance preferences can fall under different rules despite relying on the same device storage.

New Cookie Consent Exceptions

The DUAA expands exceptions to PECR's consent requirement through a new Schedule A1. Existing exceptions for transmitting communications and providing services strictly necessary at the user's request remain, while additional categories cover certain statistical purposes, appearance or functionality preferences and emergency assistance. The policy objective is to remove low-value consent prompts where privacy intrusion is limited, while retaining consent for uses, such as behavioural advertising, that pose more significant tracking and profiling risks.

The statistical exception permits an online service provider to collect information about how its service is used solely for statistical purposes with a view to improvement. Users must receive clear information and a simple, free means of objecting, and individual-level personal data must not be retained longer than needed for aggregation. The exception does not apply where analytics are used to infer characteristics, take decisions about people or pursue unrelated advertising or profiling.

The appearance exception can cover technologies used to adapt how a service appears or functions to a user's preferences, again subject to a simple, free objection mechanism. The emergency-assistance exception permits storage or access used solely to determine a device's geographic location to provide emergency assistance. These targeted categories recognise that repeatedly requesting consent for low-risk functionality creates friction without a proportionate improvement in privacy or genuine understanding.

The boundaries remain important when third-party technology is involved. A service may use an analytics provider under the statistical exception only where information serves the permitted improvement purpose. If the provider also combines data across customers for advertising, profiling or independent commercial purposes, the exception may fail. Organisations therefore need contractual and technical assurance about downstream use, since labelling a tool "analytics" does not determine its legal status; actual purpose and handling do.

Analytics Without Consent – Convenience or Reduced Control?

Fewer cookie banners can improve usability because constant requests may encourage automatic acceptance rather than considered choice. Visitors encounter repetitive prompts before reading an article, checking a timetable or completing a transaction. Removing consent requirements for genuinely low-risk statistical or functional technologies can reduce friction while still allowing individuals to object. Regulatory simplification is defensible where the technology is narrowly limited and the privacy impact is materially lower than behavioural tracking or personalised advertising.

The efficiency argument should not be confused with permission for invisible surveillance. Statistical analytics may operate without consent only within defined limits, including service improvement, information provision, objection mechanisms and aggregation requirements. If an organisation begins constructing user-level profiles, linking behaviour across services or making inferences about individuals, the lower-risk justification weakens. Convenience stays legitimate only while the organisation preserves the boundary between aggregate understanding and persistent observation of identifiable users.

ICO enforcement provides useful context. In December 2025, 979 of the UK's top 1,000 websites met the regulator's cookie compliance checks, protecting an estimated 40 million users from non-compliant personalised-advertising tracking; of those, 564 improved only after direct ICO engagement, with 17 preliminary enforcement notices issued along the way. By April 2026, that figure had risen to 99% of the top 1,000 websites, demonstrating substantial behavioural change across major online services.

Those figures show that consent design can materially affect control. The ICO assessed whether advertising cookies were placed before choice, whether rejection was as easy as

acceptance, and whether cookies continued after refusal. A framework that reduces unnecessary prompts could make remaining choices more meaningful if websites reserve banners for genuinely intrusive processing. Expanding exceptions too broadly, conversely, could normalise unseen data collection and leave individuals less aware of how their digital behaviour is observed.

Direct Marketing and Electronic Communications

Direct marketing sits at the intersection of UK GDPR and PECR. Where personal information is used to select or contact people, the organisation needs an appropriate UK GDPR lawful basis, commonly consent or legitimate interests. PECR imposes additional channel-specific rules for electronic communications such as email, text and social media direct messages. A legitimate interest under UK GDPR does not override a PECR consent requirement, so both regimes need considering before any campaign begins.

For unsolicited electronic mail to individual subscribers, organisations generally need specific consent unless a statutory soft opt-in applies. The traditional customer soft opt-in can cover similar products or services for which contact details were obtained during a sale or negotiation, provided an opt-out was offered when the details were collected and in every message. Corporate subscribers are treated differently, although named business contacts still involve personal data and remain protected by UK GDPR obligations.

The DUAA also strengthened regulatory consequences by aligning most PECR enforcement with UK GDPR levels, allowing maximum fines of £17.5 million or 4% of annual worldwide turnover, whichever is higher, and introduced a charitable soft opt-in for qualifying electronic marketing. Organisations should therefore separate three questions: whether marketing uses personal data lawfully, whether the chosen channel complies with PECR, and whether objections are operationally suppressed across every relevant marketing system.

The New Charitable Purposes Soft Opt-In

The DUAA creates a charitable purposes soft opt-in that allows qualifying charities to send electronic mail marketing without obtaining prior consent in every case. The rule covers communications intended solely to further charitable purposes, including fundraising or campaigning. It applies where the charity obtained the person's contact details when that individual expressed interest in, offered support to, or provided support for those purposes, subject to continuing safeguards against unwanted communications.

The exception is narrower than a general permission to market supporters. The charity must have collected the contact details itself, must offer a clear opportunity to refuse marketing when those details are obtained, and must provide a simple opt-out in every subsequent message. The sole purpose of the communication must be to further the charity's charitable purposes; purchased lists or unrelated commercial marketing cannot simply be brought within the new exception's scope.

The reform does not retroactively convert historical supporter databases into soft-opt-in lists. ICO guidance makes clear that charities need to have offered the required opt-out when the contact details were originally collected. A charity holding thousands of older email addresses cannot therefore assume that past donations or enquiries automatically permit electronic marketing under the new rule; existing contacts must be assessed against the permissions that applied when their details were collected.

The change places charities closer to commercial organisations that have long used the customer soft opt-in, but the triggers differ. A supporter need not purchase anything: expressing interest in, or offering or providing support for, a charitable purpose may be sufficient. A person who registers interest in a fundraising event could therefore receive related communications if the collection process included the required initial opt-out and each later message preserved that choice.

The practical advantage is reduced friction in supporter engagement, particularly where a

separate marketing consent would interrupt a genuine charitable interaction. The privacy safeguard is continuing control rather than advance permission alone. Charities should retain evidence of how contact details were obtained, the notice shown at collection and subsequent opt-outs, because a soft opt-in is not an exemption from UK GDPR: lawful basis, transparency, fairness and the absolute marketing objection right still apply.

International Transfers of Personal Data

Personal data transferred outside the United Kingdom can become harder for individuals and regulators to protect, because the recipient may operate under different laws, surveillance powers, remedies and enforcement systems. UK GDPR therefore treats certain overseas disclosures as restricted transfers. The objective is not to prevent international commerce, cloud computing or multinational operations, but to ensure protection does not disappear merely because information crosses a border or is accessed from another jurisdiction.

A restricted transfer can proceed where the destination benefits from UK adequacy regulations, where appropriate safeguards are used, or where a specific Article 49 exception applies. Safeguards can include the UK International Data Transfer Agreement, the UK Addendum to EU standard contractual clauses, binding corporate rules and certain approved codes or certifications. Where safeguards are relied upon, the sender must also consider whether the transferred information will receive sufficiently protective treatment in practice.

The 2026 UK Business Data Survey demonstrates how embedded international transfers have become. Among businesses sending personal data outside the UK and trading beyond the EU and EEA, 34% reported using EU standard contractual clauses, 20% binding corporate rules, 14% International Data Transfer Agreements and 22% adequacy. The figures also reveal ongoing complexity, since some businesses reported contractual safeguards even where an adequacy route was available for the destination in question.

The Schrems II judgment provides the wider European context for these controls. In 2020, the Court of Justice of the European Union invalidated the EU-US Privacy Shield and required closer examination of the protections surrounding transfers made under contractual clauses. Although the United Kingdom now operates its own regime, the underlying problem remains relevant: contractual wording alone cannot protect information if destination-country law or practice materially undermines the safeguards promised to individuals.

The DUAA's New International Transfer Test

The DUAA replaces older terminology surrounding international transfer assessments with a statutory data protection test. For transfers relying on appropriate safeguards, the sender must decide, acting reasonably and proportionately, that protection after transfer will not be materially lower than the protection provided in the United Kingdom. The test emphasises the practical standard observed by individuals rather than requiring overseas legal systems to reproduce UK GDPR provisions word for word.

The same "not materially lower" concept now governs the Secretary of State's assessment when making adequacy regulations for countries or international organisations. Protection is considered holistically, taking into account the destination's legal framework, respect for rights, access by public authorities, effective remedies and independent oversight. A different constitutional or regulatory model may therefore qualify if, taken as a whole, it provides a sufficiently comparable level of protection for people whose personal information is transferred.

For organisations using contractual safeguards, the reform does not eliminate the need for transfer risk assessment. The ICO still uses the term transfer risk assessment, although legislation now frames the requirement as the data protection test. Exporters must examine whether the chosen safeguard operates effectively in the destination and whether supplementary measures are needed; if protection remains materially lower and additional controls cannot correct the problem, the organisation should not rely on that safeguard.

The revised wording is intended to make international assessment more proportionate. A point-by-point comparison can produce unnecessary complexity where overseas law achieves comparable outcomes through different institutions or legal traditions. The new test instead asks whether any difference is material to protection. That flexibility supports international data flows, but it also places greater judgement on exporters, requiring decisions that can explain why a different regime remains sufficiently protective in practice.

UK–EU Data Adequacy and Regulatory Divergence

EU adequacy is economically important because it allows personal data to flow from the European Economic Area to the United Kingdom without organisations establishing additional transfer safeguards for each relationship. The European Commission renewed the UK's GDPR and Law Enforcement Directive adequacy decisions on 19 December 2025, following an assessment of the framework as amended by the DUAA, and concluded UK safeguards remained essentially equivalent to those required under EU law until December 2031.

The scale of the trading relationship explains why that decision matters. House of Commons Library analysis puts UK exports of goods and services to the EU at £384 billion in 2025, some 41% of all UK exports, with imports from the EU at £472 billion, or 50% of the UK total, giving combined trade of roughly £856 billion. Banking, insurance, technology and logistics routinely depend on data moving alongside goods and payments.

Adequacy does not require the United Kingdom to copy every future EU rule. The Commission assesses whether protection is essentially equivalent overall rather than identical in wording, giving Parliament room to develop distinct approaches to automated decision-making, legitimate interests, cookies and international transfers. Divergence nevertheless has limits: if reforms materially weaken safeguards, the Commission can monitor, amend, suspend or withdraw adequacy, creating potential compliance costs for organisations receiving EEA personal data.

The post-DUAA position therefore reflects managed regulatory independence rather than complete separation. The United Kingdom can pursue growth-oriented reforms and international data partnerships while preserving protection acceptable to its largest neighbouring market. Businesses consequently have an interest in both flexibility and continuity: excessive alignment could restrict domestic policy choices, while excessive divergence could force thousands of cross-border relationships to adopt contractual mechanisms that adequacy currently avoids.

The Information Commissioner's Changing Role

The DUAA changes not only the rules organisations follow but also the governance of the regulator enforcing them. It establishes a new body corporate, the Information Commission, to replace the existing corporation-sole model, in which statutory powers are vested in a single Information Commissioner. The future Commission will operate through a chair, chief executive, executive members and non-executive members, distributing governance more broadly while retaining the regulator's substantive statutory functions.

The institutional transition was still underway in September 2026. Seven non-executive members were appointed to the new Information Commission Board in July, while recruitment for the chair had also begun. The ICO stated that the Information Commission would take over its functions later in 2026. The change is therefore structural rather than the creation of a second regulator: responsibilities, property and ongoing work transfer to the new body once the replacement arrangements become operative.

The DUAA also introduces a strategic framework for regulatory decision-making. The regulator's principal objective continues to centre on securing appropriate protection for personal data while promoting public trust and confidence. When carrying out relevant functions, it must also have regard to matters including innovation, competition and the particular protection merited by children, embedding wider economic considerations without

converting the regulator into an agency primarily responsible for promoting business growth.

Enforcement Powers and Financial Penalties

The regulator retains a graduated enforcement toolkit rather than relying solely on fines. It can investigate complaints and suspected infringements, issue information or assessment notices, deliver warnings and reprimands, require corrective action through enforcement notices, and impose monetary penalties where statutory conditions are met. This range matters because an effective regulator must stop ongoing unlawful processing and improve future compliance, rather than treating every infringement as an offence resolved purely by payment.

The DUAA strengthens investigations by allowing the regulator to compel relevant current or former workers to attend interviews and answer questions where statutory conditions are satisfied. It can also require an organisation to commission a report from an approved person addressing specified technical matters, and provide reasonable assistance to that reviewer. These powers respond to increasingly complex systems, where understanding an alleged breach may require specialised evidence about cyber-security, algorithms or organisational controls.

Financial exposure under UK GDPR remains substantial. The standard maximum is £8.7 million or 2% of an undertaking's worldwide annual turnover, whichever is higher, while the higher maximum is £17.5 million or 4%, whichever is higher. Higher-tier penalties may apply to infringements of basic principles, individual rights and international transfer requirements. The statutory ceiling is not an automatic tariff: fines must be effective, proportionate and dissuasive, and are determined by the circumstances of each case.

Advanced Computer Software Group provides a practical example. In March 2025, the ICO imposed a final penalty of £3,076,320, reduced from a provisional £6.09 million, after a ransomware incident put information relating to 79,404 people at risk. Hackers accessed systems via an account without multi-factor authentication, and stolen data included details of how to enter the homes of 890 people receiving care, linking a basic technical failure to serious risk for vulnerable individuals.

PECR Enforcement After the DUAA

PECR historically carried much lower financial penalties than UK GDPR, creating an enforcement gap even though unlawful electronic marketing and tracking could affect millions of people. The DUAA largely aligns PECR enforcement mechanisms and monetary penalties with the UK GDPR regime. Serious infringements can now attract fines up to £17.5 million or 4% of an undertaking's annual worldwide turnover, whichever is higher, making electronic communications compliance a materially greater board-level financial risk.

The alignment matters particularly for large-scale marketing. Before the change, a business generating substantial revenue from aggressive campaigns could weigh relatively modest PECR penalties against commercial returns. The higher ceiling changes that calculation, because percentage-based penalties can scale with global turnover. It also brings the regulator's sanctions closer to the seriousness of modern tracking ecosystems, where a single unlawful practice can place cookies, send messages or profile users across millions of devices.

The DUAA also broadens the regulator's PECR enforcement framework beyond fines. The updated regime brings enforcement mechanisms into line with data protection law in most cases and permits stronger investigatory approaches. It also introduces provisions supporting approved PECR codes of conduct, which organisations may use as evidence of compliance, encouraging sector-specific standards while preserving meaningful sanctions for those who ignore consent, opt-outs or security requirements.

Recent enforcement already shows the scale of conduct involved. In May 2026, KRA Consultancy Limited was fined £300,000 after sending more than 5.5 million unsolicited marketing and fake bailiff text messages, which generated more than 60,000 complaints to

the 7726 spam-reporting service. Under the strengthened regime, comparable future misconduct by larger undertakings carries potentially much greater exposure, particularly where deliberate practices continue after warning signs or objections are ignored.

Accountability After the DUAA

Accountability requires organisations to take responsibility for compliance and demonstrate how they meet data protection requirements. The DUAA does not remove that principle; indeed, greater flexibility around recognised legitimate interests, secondary use, automated decisions and cookies increases the importance of documented judgement. Where the law gives a controller more discretion, the organisation should be able to show why a route was applied, what risks were considered and which safeguards prevented that flexibility becoming excessive processing.

The ICO's Accountability Framework offers a practical structure for that evidence, covering leadership and oversight, policies, training, transparency, records management, risk assessment, breach response and contracts with processors. Organisations can use it for self-assessment rather than treating accountability as a single annual audit. Mapping actual practice against the framework's categories tends to surface gaps between what a policy document claims and what operational systems actually do day to day.

The DUAA's policy objective of reducing unnecessary bureaucracy should not be interpreted as encouragement to document less. The ICO describes most reforms as opportunities to work differently while continuing to protect people. Mature organisations can simplify low-value processes while strengthening records around risky decisions, so accountability should become more proportionate rather than thinner: evidence should concentrate on consequential processing, vulnerable people, large datasets, novel technology and activities capable of significant individual effects.

Privacy Notices and Transparency

Privacy notices should be reviewed where DUAA changes alter what an organisation actually does with personal information. The continuing UK GDPR requirements include explaining purposes, lawful bases, recipients, retention, international transfers, relevant rights and significant automated decision-making. Updating a notice does not mean inserting generic references to the DUAA; the document should change only where processing, legal justification, rights or disclosures have genuinely changed, keeping information concise enough that individuals understand the practical effect.

Recognised legitimate interests may require organisations to revise their descriptions of the lawful basis when qualifying processing moves from ordinary legitimate interests to the new statutory route. The explanation should identify the real purpose, such as crime prevention or safeguarding, rather than merely naming an Article number. Where ordinary legitimate interests remain in use, the underlying interests should still be explained in language linking the legal basis to what the organisation is trying to achieve.

Automated decision-making deserves particular attention because the DUAA permits significant solely automated decisions on a wider range of lawful bases. Privacy information should say when such decision-making occurs and provide meaningful information about the logic, significance and envisaged consequences. An organisation deploying automated credit assessment, recruitment screening or fraud controls should not rely on a vague statement that artificial intelligence may be used; individuals need enough context to understand when automation can materially affect them.

Research and cookie reforms create different transparency issues. Qualifying scientific research may, in limited circumstances, proceed without directly supplying a privacy notice where doing so would involve disproportionate effort, provided other safeguards apply, and information is made publicly available. Some lower-risk cookies can now be used without prior consent, but statutory information and objection requirements may remain, so removing a

consent banner does not necessarily remove the obligation to explain what the technology does.

Data Protection Impact Assessments

A data protection impact assessment is required before processing likely to result in a high risk to individuals' rights and freedoms. Its purpose is preventive: the organisation describes the proposed processing, assesses necessity and proportionality, identifies potential harms and determines measures to reduce them before deployment. The DUAA has not removed this obligation; greater freedom to use data or automation does not mean high-risk projects can proceed without structured consideration of their consequences.

UK GDPR expressly requires DPIAs for systematic and extensive profiling that produces legal or similarly significant effects, large-scale processing of special-category or criminal-offence data, and large-scale systematic monitoring of publicly accessible areas. ICO guidance identifies further risk indicators, including innovative technology, large-scale profiling, data matching, biometric or genetic processing and decisions affecting access to services. Artificial intelligence can therefore trigger DPIA requirements through both its technology and the consequences of its use.

A meaningful DPIA should influence design decisions rather than merely justify an already approved project. Mitigations might include reducing the number of data fields, shortening retention periods, removing sensitive variables, introducing human review, improving explanations, strengthening access controls, or changing the supplier. Where residual high risk remains after reasonable mitigation, the controller may need to consult the ICO before processing begins, and the assessment should be revisited when systems, datasets or risks materially change.

Biometric workplace monitoring illustrates why early assessment matters. The ICO states that facial recognition used to identify workers constitutes special-category biometric data and requires a DPIA because of its high risk. An employer considering facial recognition for attendance should examine whether less intrusive methods can achieve the same objective, while weighing accuracy, bias, security and power imbalances, testing necessity before workers become subject to a system whose biometric identifiers cannot be replaced after compromise.

Data Protection by Design and Default

Data protection by design and by default requires that privacy be built into systems, services, products and processes from the planning stage and throughout their lifecycle. Article 25 requires appropriate technical and organisational measures that implement the data protection principles and protect individual rights. The objective is preventative rather than corrective: organisations should identify the necessary data, risks, access arrangements, and safeguards before a database, application, AI model, or digital service becomes operational.

Data protection by default narrows processing to what is necessary for each stated purpose, applying to the amount of information collected, the extent of processing, storage periods and accessibility. Practical controls may include field minimisation, role-based access, pseudonymisation, encryption and automated deletion. Default settings should not expose information more widely than necessary or require users to discover and disable intrusive options after a service has already begun collecting unnecessary data.

Artificial intelligence makes early design particularly important because training choices can become difficult and expensive to reverse once models are deployed. Development teams should determine which datasets are necessary, whether sensitive attributes or proxies are present, how outputs will be tested, what explanations can be provided and where human review is required. Privacy-enhancing techniques and staged deployment can reduce risk before an automated system begins influencing consequential decisions about identifiable people.

The ICO's 2024 audits of AI recruitment providers illustrate the value of intervening before harmful practices embed, producing almost 300 recommendations on fairness, minimisation and transparency, all accepted or partly accepted. Recruiters cannot assume that a developer has resolved every compliance issue and should complete impact assessments before processing candidate information. Procurement specifications should incorporate privacy-by-design requirements from the outset, since retrofitting them later is far more costly.

Data Security and Personal Data Breaches

Data security remains a core UK GDPR obligation. Controllers and processors must implement technical and organisational measures appropriate to the risk, taking into account technological capability, implementation cost, processing context and potential harm. Effective security combines cyber controls with organisational and physical measures, including access management, patching, encryption, resilient backups, staff training, vulnerability management, and regular testing, thereby protecting confidentiality, integrity, and availability rather than focusing solely on external hackers.

A personal data breach includes more than theft or disclosure. It can involve accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal information. Where a breach is likely to pose a risk to individuals' rights and freedoms, the controller must normally notify the ICO without undue delay and, where feasible, within 72 hours of becoming aware of it. High-risk breaches can also require direct communication to affected individuals without undue delay.

Effective breach management therefore requires preparation before an incident occurs. Organisations should establish internal escalation routes, incident-response responsibilities, forensic preservation procedures, containment procedures and criteria for regulatory notification. Processors must inform controllers without undue delay after becoming aware of a breach, allowing the controller to assess its own reporting duties. Every personal data breach should be documented, including its facts, effects and remedial action, even where ICO notification is ultimately judged unnecessary.

South Staffordshire demonstrates how weak monitoring can magnify cyber harm. In May 2026, the ICO fined South Staffordshire Plc and South Staffordshire Water Plc £963,900 after information relating to 633,887 customers and employees was compromised and later published on the dark web. Initial attacker access went undetected for nearly two years, while only 5% of the IT environment was monitored, exposing fundamental weaknesses in detection, access control and security governance.

Data Processors, Suppliers and Contracts

The UK GDPR distinguishes between controllers and processors because responsibility follows decision-making power. A controller determines why and how personal data is processed, while a processor acts on the controller's behalf. Outsourcing payroll, cloud hosting, customer contact, analytics or AI does not outsource the controller's accountability. Controllers must select processors that provide sufficient guarantees of compliance and continue to monitor them after appointment, rather than treating supplier certification as permanent proof that the risk is controlled.

Article 28 requires a written contract or other binding legal act whenever a controller appoints a processor. Mandatory terms cover processing subject matter, duration, nature and purpose, data types and affected individuals. Processors must follow documented instructions, preserve confidentiality, implement appropriate security measures, assist with individual rights and breaches, support audits, and delete or return data when services end. These requirements create an operational control framework rather than adding standard boilerplate.

Sub-processing requires particular attention in cloud and software supply chains. A processor cannot appoint another processor without the controller's prior specific or general written authorisation, and equivalent Article 28 protections must flow down contractually. Controllers

should understand where information is hosted, which third parties can access it, and whether overseas processing results in restricted transfers, since a primary supplier's contract is only one layer of governance across multiple infrastructure, analytics or AI providers.

Artificial intelligence suppliers can blur conventional roles, because a provider may process customer data on instructions while separately using prompts, outputs or interaction records for its own product development. Where a supplier determines the purposes independently, it may become a controller for that processing rather than remain solely a processor. Procurement teams should establish ownership of purposes, model-training restrictions, retention, confidentiality, audit rights and deletion before personal information enters an AI service.

Contractual protection must be supported by due diligence and ongoing assurance. ICO guidance requires controllers to assess processors' expertise, resources and reliability, and monitor compliance based on processing risk. Evidence can include security policies, independent assurance, penetration testing, certifications and audit findings. Where a supplier cannot explain how it protects, locates or deletes personal information, contractual promises alone should not overcome the governance weakness exposed during procurement.

Data Retention and Deletion

Storage limitation requires personal data to be kept no longer than necessary for the purposes for which it is processed. The UK GDPR does not prescribe universal retention periods because the appropriate duration depends on context, legal requirements, and operational needs. Organisations must therefore justify their own periods rather than retain information indefinitely "just in case"; cheap storage is not a lawful reason to preserve identifiable records once the purpose supporting retention has ended.

Retention schedules should translate legal and business requirements into practical deletion rules. Different records may legitimately require different periods because employment, taxation, contracts, safeguarding, complaints and litigation create distinct needs. At the end of a scheduled period, information should be reviewed and erased or anonymised unless a continuing justification exists. Automated deletion can improve consistency, but organisations need controls to prevent legal holds or active disputes being destroyed prematurely.

Unnecessarily retaining information creates measurable operational risk even where storage costs are negligible. Old records can become inaccurate, appear in subject access searches, increase breach exposure and complicate migrations or system replacement. A cyber-attack affecting ten years of unnecessary customer data can expose far more people than one involving a disciplined current dataset, so retention links directly to minimisation, accuracy, security and the practical cost of responding to access and erasure requests.

Longer retention is permitted for public-interest archiving, scientific or historical research, and statistical purposes, provided statutory conditions and safeguards are satisfied. That exception should not be confused with an ordinary organisational preference to preserve data for possible future analysis. Good governance distinguishes records with continuing legal, evidential, or research value from information that has never been deleted; disposal should be secure, documented, and capable of operating across live systems, backups, and outsourced environments.

Records, Policies and Governance

Accountability depends on records showing what an organisation actually does with personal data and why. Article 30 requires most controllers and processors to maintain records of processing activities covering purposes, categories of data and individuals, recipients, transfers, retention and security measures. These records should remain current: a data inventory written during an implementation project and left unchanged for years may demonstrate documentation existed, but not that present-day processing is understood or governed.

The record of processing activities should connect with supporting decisions rather than operate in isolation. Organisations may need documented lawful bases, Article 9 conditions, legitimate interests assessments, transfer assessments, impact assessments and processor arrangements. Consent-based activities require evidence showing who consented, when, how and what they were told. When the legal justification changes, records should be updated so staff do not continue to rely on obsolete permissions or assessments that no longer reflect reality.

Operational rights procedures form another part of the governance record. Subject access processes should define intake channels, identity checks, search responsibilities, exemptions, redaction, deadlines and escalation. Complaint procedures must now reflect the DUAA requirement to acknowledge data protection complaints within 30 days and respond without undue delay. Rectification, erasure, restriction and objection also require workflows capable of updating information across multiple systems, rather than relying on individual staff memory.

Breach registers provide a further accountability trail. The UK GDPR requires controllers to document personal data breaches, including facts, consequences and remedial action, regardless of whether each incident is reportable to the ICO. Records allow organisations to identify recurring weaknesses such as misdirected emails, excessive permissions, or supplier failures, thereby converting isolated incidents into preventive controls, training, or system changes rather than evidence retained purely for regulators.

Policies should ultimately assign ownership and decision rights. Senior management, information security, procurement, HR, legal teams, operational managers and data protection officers may each control different parts of the information lifecycle. Governance becomes ineffective where responsibility is distributed, but accountability is unclear. Regular audits, management reporting, staff training, and documented exceptions can test whether policies operate in practice, producing a traceable system of decisions that can withstand changes in personnel and technology.

Who Really Controls Personal Data?

Legal rights do not automatically create practical control. An individual may possess rights to information, access, objection, correction, and erasure, yet still struggle to understand how hundreds of organisations collect, combine, and infer information across websites, applications, payment systems, devices, and data brokers. Modern data ecosystems are technically interconnected and commercially layered, while privacy choices are often fragmented across lengthy notices, dashboards and settings that require time, knowledge and persistence to navigate.

Control is particularly difficult where data moves beyond the organisation with which the individual originally interacted. Advertising networks, cloud providers, analytics companies, fraud platforms and artificial intelligence suppliers can each process information within a single service chain. Contracts and transparency duties govern those relationships, but an individual may never recognise every participant, so the legal framework places substantial responsibility on controllers to constrain processing on behalf of people who cannot realistically supervise every downstream use themselves.

Artificial intelligence widens that gap because systems can create predictions from information that appears innocuous when collected. The ICO reports that 64% of adults believe employers may rely too heavily on AI in recruitment, while 59% are concerned about AI determining welfare eligibility. These figures reflect more than fear of technology; they point to genuine uncertainty about whether individuals can understand, challenge or meaningfully influence systems that increasingly classify them before significant decisions are made.

The imbalance does not mean rights are ineffective. Subject access, objections, complaint procedures, automated-decision safeguards and regulatory enforcement can expose or correct harmful processing, but their effectiveness depends on awareness and accessibility. A person who does not know that profiling occurs cannot object intelligently. At the same time,

someone unable to understand an explanation cannot effectively challenge a decision, so practical control depends as much on organisational transparency as on the mere existence of statutory rights.

Innovation Versus Privacy

Regulatory flexibility can support innovation where uncertainty previously discouraged legitimate data use. The Government estimates that the DUAA could contribute approximately £10 billion to the UK economy over ten years through measures that support data sharing, digital verification, Smart Data, and more efficient public services. Clearer rules on research, automated decisions, and recognised legitimate interests can reduce unnecessary legal hesitation, allowing organisations to develop services without presuming that every unfamiliar use of data is prohibited.

Smart Data illustrates the potential economic upside more concretely. Government-commissioned analysis published in March 2026 estimated that four future Smart Data schemes could generate £71.2 billion in social net present value between 2028 and 2043, contributing £9.6 billion annually to GDP by 2043. Such benefits depend on secure interoperability and trusted access, showing that privacy protection and innovation are not necessarily opposed once governance enables controlled, rather than unrestricted, data sharing.

The risk is that flexibility becomes cumulative rather than targeted. Broader automated decision-making, additional cookie exceptions, secondary-use rules and alternatives to consent can each appear modest, yet together permit more processing without active individual agreement. If organisations treat every new permission as an invitation to expand collection, profiling and retention, regulatory simplification could normalise data intensity rather than genuine innovation, since valuable innovation should solve problems rather than merely create opportunities to exploit information.

Responsible regulation should therefore reduce uncertainty without lowering the cost of harmful behaviour. The ICO's regulatory sandbox already supports organisations testing data-intensive services before full deployment, while its 2026 AI programme focuses on foundation models, automated decision-making and biometrics. Innovation is more sustainable when legal expectations are understood early: a system that launches quickly but later requires redesign or enforcement may prove less economically valuable than one built with privacy constraints from inception.

The Cost of Compliance Versus the Cost of Failure

Compliance creates visible expenditure: organisations fund data protection staff, cyber-security controls, legal advice, audits, training, impact assessments, supplier assurance and technology capable of supporting individual rights. Those costs can appear burdensome because successful prevention produces no dramatic event to measure against them. Yet the relevant comparison is not between compliance spending and zero cost; it is between proportionate governance and the expected financial, operational and reputational consequences of uncontrolled data risk over time.

The 2025/26 Cyber Security Breaches Survey found that 43% of UK businesses reported a cyber breach or attack in the previous 12 months, equivalent to approximately 612,000 businesses. Among medium and large businesses, the figures reached 65% and 69% respectively. Most incidents incurred modest direct costs, but the distribution was highly uneven, showing why averages can understate exposure: serious cyber events are uncommon relative to phishing, yet potentially far more damaging when they occur.

Among businesses reporting a material financial cost, the survey placed the median cost of the most disruptive breach at £200 overall and £300 for medium and large organisations. For the costliest 5% of medium and large cases, the figure reached £28,200. These survey estimates still exclude much of the exceptional tail risk associated with major ransomware

attacks, prolonged outages, regulatory investigations and large-scale remediation, which can shift losses from thousands to millions of pounds.

Capita provides that contrast in concrete terms. Following its March 2023 cyber incident, the company recognised £21.8 million of exceptional costs in the first half of 2023 alone, against total incident-related costs later put at over £25 million. In October 2025, the ICO separately imposed combined penalties of £14 million, reduced from a provisional £45 million, after information relating to 6.6 million people was stolen, extending the financial consequences well beyond the initial technical response.

Failure also creates costs that do not appear neatly in regulatory penalties. Litigation, customer remediation, forensic work, higher insurance premiums, diverted management time and reputational damage can continue long after systems are restored. The 2025/26 government survey found that 5% of businesses hit by cyber incidents reported lost revenue or share value, up from 2% previously, while 3% reported reputational damage, up from 1% - preventive governance is best assessed as a genuine risk investment.

Building Trust Through Responsible Data Use

Legal compliance establishes the minimum conditions under which personal data may be used; it does not guarantee individuals will regard the use as acceptable. Trust depends on whether an organisation behaves consistently with reasonable expectations, explains its purposes, limits data collection, and responds constructively when problems arise. A technically lawful activity can still damage confidence if people believe information has been obtained through obscurity, retained unnecessarily or used disproportionately to the relationship.

Public attitudes demonstrate the consequences of poor stewardship. ICO reporting in 2025 stated that 55% of people surveyed had experienced a data breach and 69% of those affected reported a negative impact, including loss of trust, emotional distress or financial loss. Separate 2025 research found psychological harm was the most commonly reported consequence among affected respondents who described an impact. Hence, data failures shape behaviour and confidence as well as creating measurable financial liabilities.

The NHS App offers the opposite case: evidence that people will share sensitive information at scale where the exchange feels secure and genuinely useful. NHS England recorded more than 41.8 million all-time app registrations by June 2026, with over 77 million login sessions that month alone, driven by prescription ordering and access to GP health records. That scale suggests people will trust an organisation with health data once authentication and everyday usefulness are visibly aligned.

Responsible data use therefore requires organisations to treat trust as an operating asset. Transparency should explain genuine practices; minimisation should restrain unnecessary ambition; security should reflect potential harm; and complaints should inform improvement. Where a proposed use would be difficult to defend openly to the people affected, legal permission alone may be an insufficient standard. Sustainable data strategies align commercial or public value with a continuing social licence to collect, analyse and reuse personal information.

Best Practice for Organisations After the DUAA

Best practice after the DUAA begins with accurate data mapping. Organisations should know what personal information they hold, where it comes from, why it is processed, which lawful basis applies, who receives it, how long it is retained and whether it leaves the United Kingdom. New recognised legitimate interests and compatibility routes should be documented precisely, rather than used as generic alternatives whenever consent or ordinary legitimate interests appear administratively inconvenient.

Transparency should then match the actual processing environment. Privacy notices need to explain material purposes, lawful bases, sharing, retention, overseas transfers and significant

automated decisions in accessible language. Consent should be reserved for situations involving genuine choice and recorded so that withdrawal can be honoured. Data minimisation should operate at collection and throughout the lifecycle, with unnecessary fields, duplicate repositories and speculative retention removed before they become permanent organisational practice.

Artificial intelligence requires additional governance. Organisations should identify training and input data, sensitive attributes and proxies, model objectives, error risks, human decision points and routes for challenge before deployment. High-risk uses should undergo a DPIA, while automated significant decisions need appropriate safeguards. Procurement teams should establish whether an AI provider acts as a processor or an independent controller, whether customer information is used to train models, and how deletion and security are contractually controlled.

Individual-rights processes should be operational rather than theoretical. SAR teams need reliable search methods, deadline controls and escalation routes; rectification, erasure, restriction and objections must propagate across relevant systems. Complaint arrangements should acknowledge data protection complaints within 30 days and investigate without undue delay. Records should capture recurring themes, so that complaints and rights requests reveal weaknesses in notices, retention or system design rather than being treated as isolated administrative cases.

Security, retention and supplier oversight complete the framework. Access controls, multi-factor authentication, patching, monitoring, tested backups and incident response should reflect the sensitivity and scale of data held. Retention schedules should trigger review, deletion or anonymisation, not indefinite storage. Suppliers should be assessed before appointment and monitored afterwards for security, sub-processing, transfer and deletion obligations, with governance tested through audits, management reporting and periodic reviews as technology changes.

The Future of UK Data Protection

UK data protection will increasingly be shaped by technologies that blur the boundary between information supplied by an individual and information inferred about them. Foundation models, agentic AI, facial recognition, neurotechnology and biometric identification can generate or act upon revealing data at scale. The ICO's 2026 programme includes automated decision-making, foundation models, agentic AI and neurotechnology, signalling that future regulation will increasingly focus on inference, autonomy and accountability rather than on databases alone.

Digital identity will create another major data layer. The DUAA establishes statutory foundations for digital verification services, while the Government's framework sets certification, trust and governance requirements for providers. DSIT has stated that measures in the Act could contribute a further £4.3 billion to the economy over the next decade, on top of a sector that already generates more than £2 billion in annual revenue and supports over 10,000 jobs across certified providers.

Smart Data is likely to extend controlled portability beyond banking into sectors such as energy, retail, property and trade. The 2026 Smart Data Strategy envisages interoperable schemes connected with the wider data economy and artificial intelligence. This could strengthen individual agency by enabling authorised services to act on customer-directed data. Interoperability also expands the number of relationships that require secure authentication, transparent permissions, effective revocation, and protection against misuse by authorised recipients.

Internationally, the United Kingdom must balance regulatory independence with interoperability. The European Commission renewed the UK's adequacy status from December 2025 to December 2031, allowing EEA personal data to continue to flow without additional transfer safeguards. Future divergence remains possible, particularly around AI,

automated decisions, and international transfers, but material weakening could threaten that status, so UK policy will continue to operate between distinctive pro-innovation rules and sufficient comparability with international privacy regimes.

Does the Individual Still Have Control?

The individual still has legal control, but it is no longer sensible to equate control primarily with consent. Post-DUAA rights to information, access, correction, erasure, restriction, objection, portability, automated-decision challenge and regulatory complaint remain substantial. Those rights can constrain organisations after data has been collected and, in some situations, stop processing altogether. The framework therefore preserves agency but distributes it across continuing rights and organisational duties, rather than relying on initial permission alone.

Practical agency is weaker where individuals cannot see or understand the processing taking place. Data can move through suppliers, be reused under statutory compatibility rules, generate inferences through AI or support decisions under lawful bases that do not require consent. Control then depends on organisations explaining their activities and providing realistic routes for challenge. A right that exists only after someone discovers hidden processing offers less protection than transparent design that prevents surprise altogether.

The DUAA also gives organisations greater judgement. Recognised legitimate interests, broader automated decision-making and selected consent exceptions reduce procedural barriers in deliberately defined areas. That can improve services and remove artificial consent exercises, but it increases the importance of accountability: when fewer decisions require active permission, more responsibility shifts to controllers to correctly determine necessity, proportionality and safeguards, making control partly dependent on institutional competence rather than personal choice alone.

Technology further complicates the balance because personal information increasingly includes predictions rather than facts the individual volunteers. A person can refuse to disclose a characteristic yet still have an algorithm infer it from behaviour, location, purchases or associations. Traditional privacy controls were designed around collection and disclosure; modern systems can create new information internally, so meaningful agency now requires governance of inference and automated consequences alongside conventional rights over the source data itself.

Personal data is consequently controlled through a shared but unequal structure. Individuals retain legal agency, regulators can intervene, and organisations remain accountable, yet controllers usually possess superior information, technology and resources. The DUAA does not remove that imbalance and, in some areas, increases organisational discretion. Its legitimacy therefore depends on whether controllers use that discretion narrowly and responsibly, because meaningful control cannot rest on expecting every individual to monitor complex data flows personally.

Summary – Data Use and Responsibility

The Data (Use and Access) Act 2025 represents an evolution of UK data protection rather than a rejection of GDPR principles. It creates recognised legitimate interests, clarifies secondary use and research, broadens significant automated decision-making, simplifies selected cookie rules and strengthens Smart Data and digital verification infrastructure. Organisations gain flexibility in defined circumstances, but the UK GDPR, Data Protection Act 2018 and PECR continue to regulate how personal information is collected, secured, shared and challenged.

The central change, therefore, is where decision-making responsibility sits. Consent remains important but is less central in some forms of processing, while organisational judgement, statutory safeguards and accountability carry greater weight. Individuals still possess rights to transparency, access, correction, erasure, restriction, objection, portability and complaint. Easier data use has not eliminated control, but it has made organisational governance more

important wherever processing can proceed without an individual first giving explicit permission.

The lasting principle is that permission to use data is not permission to use it carelessly. Organisations that understand purposes, minimise collection, design privacy into technology, govern AI, manage suppliers, delete unnecessary records, secure systems and respect individual rights can use the DUAA's flexibility constructively. Those that treat reform as deregulation increase legal, financial and reputational risk. Data can now be used more freely in selected circumstances, but responsibility remains the price of that freedom.

Additional articles can be found at Procurement and Supply Chain Management Made Simple. This site examines procurement and supply chain management issues to assist organisations and individuals in improving the quality, efficiency, and effectiveness of their product and service supply, to the customers' delight. © Procurement and Supply Chain Management Made Simple. All rights reserved.

Sources and Further Reading

Legislation and official guidance

- *Data (Use and Access) Act 2025, c. 18, and accompanying Explanatory Notes (legislation.gov.uk)*
- *UK General Data Protection Regulation (Retained Regulation (EU) 2016/679) and Data Protection Act 2018, as amended*
- *Privacy and Electronic Communications (EC Directive) Regulations 2003 (PECR), as amended*
- *Information Commissioner's Office, "Data (Use and Access) Act 2025" guidance hub, ico.org.uk*
- *Information Commissioner's Office, Accountability Framework, Children's Code, Recruitment Rewired (2026) and "Explaining decisions made with AI" guidance*
- *Department for Science, Innovation and Technology, press notices on the Data (Use and Access) Bill and Act, including "New data laws unveiled to improve public services and boost UK economy by £10 billion" (October 2024) and "New data laws will make life easier for British people" (June 2025), gov.uk*
- *Department for Business and Trade / Public Service Consultants, "Understanding the costs and benefits of Smart Data use cases" (March 2026), gov.uk*
- *European Commission, Implementing Decisions renewing UK adequacy under the GDPR and Law Enforcement Directive (19 December 2025)*

Regulatory enforcement notices and reports

- *ICO monetary penalty notices and enforcement updates: TikTok (2023), Advanced Computer Software Group (2025), Capita plc and Capita Pension Solutions (2025), South Staffordshire Plc (2026), Reddit Inc and MediaLab.AI Inc (2026), 23andMe (2025), Allay Claims Limited, ZMLUK Limited and KRA Consultancy Limited (2026), all published at ico.org.uk/action-weve-taken/enforcement*
- *Public Sector Fraud Authority, National Fraud Initiative Report 2022-2024 and National Fraud Initiative Strategy 2024-2028*
- *Department for Science, Innovation and Technology, Cyber Security Breaches Survey 2025/26*
- *Open Banking Limited, monthly ecosystem performance updates and "UK Open Banking ecosystem surpasses one billion payments and 100 billion API calls" (July 2026), openbanking.org.uk*

- *NHS England Digital, NHS App Management Information statistical releases*

Case law and academic commentary

- *Information Commissioner v Experian Limited, Upper Tribunal (2024)*
- *Uber B.V. and Ors v Aslam and Ors related automated-decision proceedings, Amsterdam District Court (2021)*
- *Data Protection Commission (Ireland) and CJEU, Data Protection Commissioner v Facebook Ireland Ltd and Maximillian Schrems (Schrems II), Case C-311/18 (2020)*
- *House of Commons Library, research briefings on the Data (Use and Access) Bill/Act and on UK-EU trade*

This article draws on official legislative text, regulatory guidance and enforcement decisions current as at September 2026. Because implementation of the DUAA continued in stages through to June 2026, and because ICO guidance is updated periodically, readers relying on this material for live compliance decisions should verify current positions against ico.org.uk and legislation.gov.uk.