Emerging Tech

ashurst

OneTrust DataGuidance™
REGULATORY RESEARCH SOFTWARE

CONTRIBUTORS

Gita Shivarattan, Ashurst
Gita Shivarattan is a Counsel in the Digital Economy Transactions group at Ashurst LLP. Gita specialises in UK data protection law, and has extensive experience on advising on a range of technology, commercial and data protection law matters including IT outsourcing, business process outsourcing, development and licensing arrangements, and IT-related issues in mergers and acquisitions. Gita also provides practical guidance on the legal and regulatory aspects of digital transformations and implementing dynamic technologies such as cloud, SaaS and automation. Gita has a wide range of experience in advising clients in relation to data protection compliance and has recently supported a number of clients on GDPR compliance projects.
gita.shivarattan@ashurst.com

Tom Brookes, Ashurst
Tom Brookes is a Solicitor in Ashurst's Digital Economy Transactions group and is based in London. Tom has recently completed a six month secondment in the legal team of a global technology company, and has advised a number of multinational companies on data protection matters relating to data breach response, GDPR compliance projects and corporate acquisitions. He also has experience advising on the strategic issues relating to the adoption and use of emerging technologies, such as virtual shopping assistants. Prior to training as a solicitor, Tom was an Analyst at DataGuidance by OneTrust, where he was responsible for managing content for Africa, Middle East and Asia Pacific.
tom.brookes@ashurst.com

Tara Waters, Ashurst
Tara Waters is a partner in Ashurst's Corporate team and Co-CEO of Ashurst Digital Ventures. She advises on US and UK law for a wide range of corporate and financing transactions, with a particular focus on the technology sector. Tara leads Ashurst's high growth & VC team in London and is a key member of the firm 's fintech and distributed ledger technology & crypto asset practices.
In her role as Co-CEO of Ashurst Digital Ventures, Tara is responsible for the firm's in-house development and investment arm of Ashurst Advance, providing innovative technology-led solutions to clients.
tara.waters@ashurst.com

Image production credits
Just_Super / Signature collection / istockphoto.com

Published by OneTrust DataGuidance Limited, Dixon House, 1 Lloyd 's Avenue, London EC3N 3DS

Website **www.dataguidance.com**

 ISSN 2398-9955

Managing Editor **Alexis Kateifides**
akateifides@onetrust.com

Editorial Assistant **Victoria Ashcroft**
vashcroft@onetrust.com

OneTrust DataGuidance partnered with Ashurst LLP to present Emerging Tech, a four-part series of articles and videos on the data protection issues relating to novel forms of technology.

Alexis Kateifides was joined by Gita Shivarattan and Tom Brookes, of Ashurst LLP, for the first instalment of the series. Gita and Tom introduce some of the key issues clients should consider when incorporating new technology solutions, such as data subject rights, transparency, and discussions around innovation and regulation.

In the second installment of the series, Gita introduces some of the topics that clients should consider when incorporating artificial intelligence ('AI') technology solutions, such as lawful bases, data minimisation, data subject rights, and Privacy by Design, alongside discussing some uses and applications of AI technologies.

Tara Waters and Gita joined us for the third instalment of the series. They outline some of the main areas users should consider when working with blockchain, particularly its coexistence with data privacy law.

In the final part of the Emerging Tech series, Gita and Tom consider some of the concerns faced by the ad tech industry and discuss which regulations have had the most impact on digital advertising.

Emerging technologies

Gita Shivarattan Counsel
gita.shivarattan@ashurst.com

Tom Brookes Solicitor
tom.brookes@ashurst.com

Ashurst LLP, London

Emerging technologies, such as artificial intelligence ('AI'), blockchain, and adtech, provide limitless opportunities for businesses to develop innovative, tailored, and targeted offerings to consumers, and to engage a wider audience. Increasingly sophisticated data monitoring and analytics can provide businesses with a competitive advantage due to improved knowledge about service users, such as where they go, what they buy, and what their habits and preferences are.

Whilst these practices are not new, it is the emerging technologies that enable the scale of connectivity and the volume of data to be processed, which sets them apart from traditional technologies. These emerging technologies are able to harness vast volumes of data in real-time, and effortlessly transfer that data around a complex network of other connected systems and applications. The innovative ways in which data is used in these technologies, and the increasingly complex data flows, need to be considered in the context of the 'new world order' of the General Data Protection Regulation (Regulation (EU) 2016/679) ('GDPR').

Innovation v. regulation

In the last 24 months, we have seen a plethora of digital technologies that are already inspiring competition in consumer markets, for example; 'digital assistant' technologies (e.g. Apple's 'Siri,' Google's 'Home,' Amazon's 'Alexa,' or Microsoft's 'Cortana'), wearable computing, sensor technology, and drones. Many businesses, too, are taking advantage of the flexibility, learning capabilities, and connectivity of new technologies to improve and digitise their services to engage a wider audience.

This development of intelligent software and digital AI technologies, that are capable of simulating human reasoning and learning, is industry-agnostic. Consequently, the related privacy risks will depend on the type of data that is being processed and the purpose of the processing.

In parallel with these developments, the GDPR came into effect in May 2018 and introduced the highest standard for data protection laws globally, as well as new rights for individuals that are aimed at enabling greater control in respect of their data.

One of the main drivers behind the GDPR was to modernise European data protection laws, and to introduce a regime which was created with technology in mind. The GDPR replaced the Data Protection Directive (Directive 95/46/EC), which was over 20 years old (i.e. before social media networks became one of the main mediums for communication and some of the world's largest data hubs and exchanges). That said, given the pace of development and the innovation of emerging technologies, some aspects of the GDPR may already seem antiquated, and complying with its obligations can prove challenging. It follows as no surprise that data protection regulators acknowledge the tension between technology innovation and regulatory compliance, often emphasising that data protection should be seen as an opportunity, rather than a barrier to innovation. The UK 's Information Commissioner's Office ('ICO') has published its Technology Strategy 2018-2021[1], in which it states:

'The most significant data protection risks to individuals are now driven by the use of new technologies. The risks are broad - from cyber-attacks to the growth of [AI] and machine learning [...] [t]hese advances need not come at the expense of data protection and privacy rights - the ICO 's approach to technology will be underpinned by the concept that privacy and innovation are not mutually exclusive. When they both work together this creates true trust and data confidence. Technology is therefore viewed by the ICO as a risk and an opportunity.'

Data privacy - a right challenged by technology?

In Europe, privacy is seen as a fundamental human right, and it is from this premise that the GDPR was drafted. However, large data flows and connectivity of everyday devices raise difficult questions about how to best protect this right. Smart devices are able to record habits, lifestyles and health patterns, providing useful information to companies to improve and tailor products and services. But who else has access to this information? And what related assumptions and biases are being made?

Emerging, data-heavy technologies therefore carry a unique set of privacy challenges. Digital boundaries are regularly poorly defined and communication between smart devices can often be triggered

automatically. Couple this with the number of different stakeholders, each carrying out separate activities within the data processing lifecycle, from data-aggregators to device manufacturers and application developers, who may seek to repurpose that data for entirely different uses. Sophisticated data technologies can mean that what might have once been anonymised or insignificant user data, may now be used to make more consequential inferences.

Transparency and control

Therefore, the lack of transparency and user control is a huge issue, and is pervasive across all technological developments.

Businesses need to carefully assess the purpose for which they are using the data, and whether they need to obtain specific, informed consent from individuals for ancillary (and unnecessary) processing of their data. This is likely to present challenges where traditional consent mechanisms are not up to the task.

The solution, in many cases, is to embed privacy pop-ups or defaults at all stages within the design of devices and applications, and as close as possible to the point of data collection. Businesses will need to design and implement user-friendly privacy notices around the use of data, and seek to develop new methods of giving information to users to allow for greater transparency and individual control.

Controller or processor conundrum

Often in these complex relationships between stakeholders, it is not black or white as to which party is acting as a controller or a processor. Indeed, in some scenarios, a party may take multiple roles depending on the processing activity. Nevertheless, being clear about the roles of a party in respect of processing is key to assessing both risk and the applicable obligations under the GDPR. For example, controllers have to comply with the transparency principle, part of which requires that they provide individuals with a fair processing notice stating (amongst other things) the ways in which their data is used, who else may have access to it, how long it is kept for, and so on. However, it may be that a party that it quite far removed from the end-user is acting as a controller, and it would be impracticable for that party to provide the individual with a notice. The solution, more often than not, is to contractually delegate the responsibility to notify the end-user to the party with the direct relationship. We now see many businesses providing template wording for counter-parties to include in their privacy notices to address this issue.

Data subjects - rights and hype

The GDPR introduces a number of enhanced rights for individuals in respect of their personal data, such as the right to be forgotten, the right to data portability, the right to not be subject to automated decision-making, and the right to object. As mentioned above, the purpose of introducing these rights was to ensure that individuals have control over their data and the way in which it is used.

The hype around the introduction of the new data subject rights, as well as the abolishment of the previous subject access fee (although it was nominal) are likely to result in an increased number of such requests, which organisations will need to process within the one month period, unless there are extenuating circumstances.

Fines associated with infringements by an organisation of its obligations to comply with a data subject request to exercise his or her rights, may result in the higher tier of fine available under the GDPR of up to 4% of global turnover, or €20 million. Organisations should consider, at the outset, how a data subject request will be handled and implemented, as well as, how to ensure that the network of stakeholders, which may be in receipt of the data, are also notified, and assist with such requests.

Security - one size fits all?

Security and privacy principles go hand-in-hand in ensuring organisations stay on the right side of both regulators and consumers. Under the GDPR, both controllers and processors of data are required to implement appropriate technical and organisational measures to protect against unlawful or unauthorised processing of personal data. If organisations fall short, affected individuals have various rights to demand resolution and, in some situations, to receive compensation. In addition, regulators have the power to issue sanctions and fines of up to 2% of global annual turnover, or €10 million. When assessing what is an appropriate level of security, organisations need to give consideration to the nature of the personal data, the technology available in the market, the cost of implementation, the risk of processing, as well as the associated risk of varying likelihood, and the severity of the rights and freedoms of natural persons. It is easy to see from this list, that this is not a one size fits all determination and will require that the right internal stakeholders are involved in informing this decision.

In an era of connected technologies, the quality and safety of digital services is key. The security threat landscape is changing rapidly and the advent of smart connected technologies means any small, connected element might become a potential point of vulnerability to the whole system. A good illustration is the attack on domain name system provider, Dyn. The sustained cyber attack was launched through various ancillary (but connected) parts of the system, including CCTV cameras, printers, and even baby monitors.

To address the cybersecurity risk, organisations cannot assess risk in a silo of the individual pieces of technology or processing activities. A holistic review of the full system architecture, from workstations, communication links and storage infrastructure, to anything with the potential to connect, is key to ensuring that the risk has been appropriately quantified.

Whilst employing state of the art security measures may go some way to protecting data in the system, it is well known that most security breaches arise out of simple human error. Regular and effective staff training, supported by appropriate and detailed policies, can ensure cybersecurity becomes an embedded culture among employees. In addition, bringing in external advice where there are gaps in a team's expertise will help, and integrating cybersecurity into digital business models from the outset will increase trust, efficiencies and confidence across an organisation.

1. Available at: https://ico.org.uk/media/about-the-ico/documents/2258299/ico-technology-strategy-2018-2021.pdf

Artificial intelligence

Gita Shivarattan Counsel
gita.shivarattan@ashurst.com
Ashurst LLP, London

AI is emerging in its own right as a nascent industry with the potential to raise the productivity of a diverse range of sectors and create entirely new jobs. PWC estimates that 'AI could contribute up to $15.7 trillion to the global economy in 2030, more than the current output of China and India combined. Of this, $6.6 trillion is likely to come from increase productivity and $9.1 trillion is likely to come from consumption side effects.'[1]

The growth and adoption of this technology is inevitable. As a technology, albeit at an embryonic stage, AI has already proven that when strategically deployed it drives operational efficiencies and can lower costs. For example, chatbots, smart reply and predictive technologies being used as a first response customer service tool and warehouse and distribution analytics for network optimisation.

AI solutions are dependent on access to large and diverse datasets. These datasets are required to shape, train and direct AI towards the required outcomes. It follows that employing these technologies requires an understanding of data licensing, data sharing, and digital trust models, as well as inherent data protection challenges. In the second part of the Emerging Tech Series, we look at AI and related data protection considerations.

What is AI?

AI is a form of computing that allows machines to perform cognitive functions, such as reacting to input, in a similar way to humans. This is different to traditional computing functions, which also react to data, as in traditional computing all the responses are hand coded meaning that there is a finite set of responses, and unexpected inputs cannot be computed.

In comparison, current AI enabled technologies are able to modify the response based on an analysis and interpretation of data. This is known as 'machine learning,' the capacity for machines to learn and take independent decisions. Before we delve into the data protection considerations, for context we have set out some AI applications and examples of use cases in which are already present in our everyday lives (see **Figure 1**).

There are other legal challenges around AI which need to be grappled with, including:

- potential discrimination or bias;
- the impact on resourcing and labour markets; or
- antitrust issues.

Each of these important issues requires thoughtful consideration, but they are beyond the scope of this article which focuses exclusively on data protection legal issues and AI.

AI and data protection

Personal data

Data protection laws govern the use of personal data. The definition of personal data can vary by jurisdiction and by statute, therefore ascertaining whether personal data is involved is not a simple task. The line between what is 'personal' and what is not has been blurred by the correlations and inferences that can be made from aggregated data sets. Today, information that once seemed to be non-personal now has the potential to be personal data, particularly where distinct data elements are joined together. The General Data Protection Regulation (Regulation (EU) 2016/679) ('GDPR') defines personal data as:

'any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.'

The very nature of AI, including the variety of data sets on which it often depends to expand the capability for linking data or recognising patterns of data that may render non-personal data identifiable, seems to be in constant tension with the challenge of determining when data protection laws apply.

Businesses looking at investing in or developing AI solutions will need to ensure that they are able to clearly define the scope of any personal data that is collected by and processed through the proposed solution.

In addition, businesses should consider whether personal data is in fact necessary for the purpose of processing, or if anonymised data sets can achieve the same outcome.

Technology Process	Description	Use Case
Data mining	Discovering patterns or extrapolating trends from data	• Anomaly detection, identifying fraudulent entries or transactions; • association rules, detecting supermarket purchasing habits by looking at a shopper's typical shopping basket; and • predictions, predicting a variable from a set of other variables to extrapolate for example a credit score.
Image processing and tagging	Analyse images to get data or to perform transformations	• Identification/image tagging. Algorithms learn facial recognition in photos on social media to unlock smartphones and search images. This leads to the ability to ascertain other data from a visual scan, such as health of an individual or location recognition for geodata; and • optical character recognition. Algorithms learn to read handwritten text and convert documents into digital versions.
Text analysis	Extract information or apply a classification to items of text-based data	• Filtering, used in email exchanges to identify spam; • information extraction, for example to pull out particular pieces of data such as names and addresses; • sentiment analysis to identify the mood of the person writing, as Facebook has recently implemented in relation to postings which are potentially suicidal; and • chatbot technology allowing for interaction on line with customer service messaging services.
3D environment processing	An extension of the image processing and tagging skill where analysis is carried out on images presented in the real world to create spatial and relational images	• The learned skill required by an algorithm in 'Connected and Autonomous' vehicles to understand its location and driving environment; and • free roaming robot devices including pilotless drones.
Speech analysis	Equivalent skills to those used for textual documents and applies them to the spoken word	• Personal digital home assistants from the likes of Amazon's Echo device, Microsoft's Cortana, Google's Home device and Apple's Siri.

Figure 1: Key applications of AI technologies in everyday uses

Fairness and transparency

Under Article 5(1)(a) of the GDPR, the processing of personal data must be fair and lawful. Fairness is a fluid concept which is determined through several elements:

- transparency, i.e. that individuals are presented with information about the processing;
- the effect of the processing on individuals, i.e. does the processing determine an outcome; and
- what the expectations of the individuals are regarding how their data is used, i.e. is the intended processing in the 'reasonable expectations' of the individual.

The effect of processing the individual's data should be built into the design and implementation phases of AI solution deployment. At the design phase, datasets and any inferred data should be clearly defined and tested to ascertain any biases created by the algorithm.

Often, AI results in a form of profiling or automated decisioning which, subject to the circumstance, can have a more intrusive effect on individuals. An everyday example of this is credit scoring algorithms used to determine credit limits. Credit applicants are likely to be aware that an online application with 'instant credit decisions' will be processed through automated means, however, what they may not be aware of is that the decision to provide credit will be based on their assignment to a 'group' and the factors identified by the analytics, which are common to that group. It is this potential bias or discrimination, which is a result of inferred data, that individuals are often unaware of.

Under Article 22 of the GDPR, individuals have a right not to be subject to decisions carried out purely by automated profiling or decision making. If AI solutions are being relied on to solely determine a decision, that is without human interference, which will have an impact on an individual, under the GDPR the individual has a right to:

- object to the processing;
- request that a human carries out the analysis; and
- request further information about the way the automated decision was arrived at.

From a practical perspective, businesses should review the organisational measures implemented to ensure a request for information about the AI decision process, or a review of the AI decision, is adequately handled.

As AI solutions should be a more efficient and accurate method of arriving at a response, businesses should regularly spot check outputs against human decisioning on the same input in order to test whether the outcomes are aligned with a human review, and if any bias has been created through the learned data. This regular review should be built into the governance process for the technology.

The right to be informed and the principle of transparency require that businesses are able to clearly, and using plain language, describe:

- how the personal data is processed;
- any other sources of personal data used in the processing;
- the lawful condition of processing;
- retention periods; and
- the details of any automated decisioning, including a description of the algorithm etc.

In order to effectively discharge this obligation, businesses will need to have fully considered the privacy impact of employing the AI solution.

Lawful basis

In order for processing to be 'lawful' under the GDPR, it will need to satisfy one of the conditions for processing in Article 6 of the GDPR. The most likely conditions to be relevant to processing through AI are:

- consent;
- legitimate interest; or
- performance of a contract.

The standard of consent under the GDPR requires that it is 'freely given, specific, and informed.' In addition, consent needs to be able to be withdrawn. Businesses will need to assess whether this threshold can be met given the opaque nature of AI technologies. 'Just in time' notices and consents may be a practical way in which businesses can employ consent to ensure that consents are tailored to specific processing activities.

Legitimate interest is often lauded as the most flexible lawful basis, however the onus is on the businesses to consider any unwarranted impact on the rights and freedoms of individuals, and that the required safeguards and governance have been implemented to meet the GDPR obligations. Furthermore, in order to satisfy this condition, the processing is to be 'necessary' for the stated legitimate interest, meaning that if there is another way to process the personal data to meet the stated legitimate interest, which is less intrusive into people's privacy, the method of processing will not be considered necessary. A decision to rely on legitimate interest will need to be documented in a legitimate interest assessment.

Whilst legitimate interest is an alternative to seeking consent, individuals will have a qualified right to object to processing based on legitimate interest, and therefore businesses will be responsible for implementing appropriate processes and procedures to handle such requests.

The performance of a contract condition is of more limited applicability with regards to processing through AI solutions, and requires a case by case analysis. In general, the processing carried out by AI often goes beyond what is required in order to sell or deliver a product or service and therefore it may be difficult to evidence that the processing is strictly necessary in order to 'perform the contract.'

Purpose limitation

The purpose limitation, i.e. personal data that is only used for the purpose as notified to the individual when the data is initially collected, with certain exceptions, is a central tenant of the GDPR and transparency principle. However, the enhancement of AI solutions entail a material issue with reference to the purposes of data processing, such as:

- the ability of an AI solution to interact with the surrounding environment;
- to learn from the experience; and
- to address future behaviours based on such interactions and learnings.

Whilst the GDPR does not prohibit the use of personal data for an additional purpose, the secondary purpose must not be incompatible with the original purpose. This is a further assessment of fairness. Guidance from the Article 29 Working Party's opinion on purpose limitation states that where processing is carried out for a secondary purpose which involves making a decision which affects the individual, 'free, specific, informed and unambiguous consent would almost always be required, otherwise further use cannot be considered compatible[2].'

AI and machine learning features may cause the processing of personal data to be carried out in different ways and for different purposes than those for which it was originally set. Businesses will need to constantly review the potential outputs and use cases for derived and inferred data, ensuring that the way the data is used is consistent with the original purpose, or carrying out further analysis to determine related use cases, their related lawful basis and consider whether additional notifications, and related consents, are required.

Data minimisation

Data minimisation is the principle that '[p]ersonal data must be adequate, relevant and limited to what is necessary in relation to the purposes for which those data are processed' and 'personal data should not be kept for longer than is necessary for that purpose.' This stands in tension with developing AI technologies, as it is difficult to know in advance 'what is necessary' in a world of 'surprising correlations' and computer-generated discoveries. The challenges of defining a purpose for processing and only keeping data for that purpose are exacerbated because the nature of machine learning and AI means businesses are not able to predict what the algorithm will learn and secondly, the purpose may also be changed through advances in the algorithm.

AI also challenges retention limits because deleting or restricting the use of data after its original purpose has been fulfilled, or upon request by an individual, could strip businesses of the potential benefits of using that data for AI development, deployment, and oversight. Data is essential if these models are to perform optimally. Yet, keeping data for longer periods, or indefinitely, may fall foul of current data protection laws.

There is clearly no hard and fast rule as to how to apply the data minimisation principles to this evolving technology, but businesses will need implement appropriate controls to ensure that data used to teach these models are up-to-date and accurate, as well as reviewing processes for deletion, where technically possible, as appropriate.

Accuracy

Data quality plays a central role in the effectiveness, or ineffectiveness, of AI technologies. Data quality is relevant at all stages of the processing cycle, collection, analysis and application. Under the GDPR, businesses have an obligation to ensure that data is accurate and up-to date. Businesses employing AI will need to consider not only data accuracy on initial collection, but also the possibility that collected data may become out of date or be inaccurate. Predefining a process for how such inputs can be corrected, and considering the implication of correcting data on previous outcomes, will be crucial to ensuring that results are not tainted by inaccurate data.

Further considerations relating to the training datasets should also be carried out to ensure that the 'sample data' is representative of the population as a whole. This is particularly important where AI is used to perform a level of profiling or automated decisioning.

Finally, businesses should also assess and test any hidden biases contained in datasets which, subject to the purpose of the technology, may lead to inaccurate predictions based on inferred or derived data.

Data subjects' rights

The GDPR introduces a number of enhanced rights for individuals with regards to their personal data, such as the right of access to processed personal data, the right to be informed about the processing, the right to restrict the processing, the right to erase the personal data concerning the data subject, the right to object to the processing of personal data and the right to data portability. In practice, employing new AI technologies, are likely to require that processes relating to handling data subject requests will need to be reviewed and amended to take into account the new method of processing. It is critical that such processes are clearly defined and documented before new technologies are deployed to ensure that businesses do not fall foul of the mandated response periods as set out in the GDPR.

Privacy by Design and Privacy by Default

The GDPR codifies the concept of Privacy by Design and Privacy by Default (Article 25). Businesses looking to implement AI solutions should carefully consider how to adhere to these principles, which may not naturally fit with the nature of data processing by AI systems. The concept of Privacy by Design is important to ensure that data protection principles, such as minimisation, proportionality, etc., are considered at the design phase of AI solutions. It requires that businesses carefully consider and integrate available safeguards into the processing to ensure that the requirements of the GDPR are adhered to. It follows that under the Privacy by Default, businesses should set technical and organisational measures which typically only permit the processing of what is necessary for each specific purpose.

Data Protection Impact Assessment

Under Article 35(3) of GDPR, the Data Protection Impact Assessment ('DPIA') is required, in case of 'a systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person.'

Accordingly, most AI solutions would require a DPIA before carrying out any personal data processing. This will require a detailed assessment of AI solution from a data protection perspective, and an assessment of the relevant security measures which are applied. The Information Commissioner's Office Privacy Impact Assessment Code of Practice sets out the requirements and practical steps on how to complete a DPIA.

Prior consultation of the supervisory authority

In addition, 'where a data protection impact assessment indicates that the processing would result in a high risk in the absence of measures taken by the controller to mitigate the risk,' the controller shall consult the relevant data protection supervisory authority under Article 36 of the GDPR. Open engagement and fluid dialogue with data protection supervisory authorities is key to ensuring that businesses and regulators are informed about any challenges being faced when employing new technologies and determining how the legal obligations can be met.

AI and ethics

There is also a trend towards developing ethical frameworks, i.e. beyond legal compliance, to govern the use of data in AI and other data analytics technologies. This ethical approach addresses the concern that whilst use cases may be legal, is the use case 'responsible.' Currently there is no widely adopted ethics framework or harmonised set of principles for the ethical approach to AI and data, however, if businesses are looking to leverage emerging technologies which are more opaque regarding the way data is used and the outcome, they should be cognisant of the ethical considerations which are typically fairness and transparency.

Conclusion

AI technologies have great potential to offer insights to businesses and the public sector, however, these technologies need to be designed and tested to ensure that they adhere to the data protection legal framework. It is evident that the challenge will be applying the framework in the face of rapidly changing AI solutions and businesses should implement appropriate data protection governance frameworks to meet this challenge.

1. PWC: sizing the prize: what the real value of AI for your business and how can you capitalise? June 2017, available at: http://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf
2. Article 29 Data Protection Working Party. Opinion 03/2013 on purpose limitation, European Commission, 2 April 2013, available at: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2013/p03_en.pdf (pg. 46)

Blockchain

Tara Waters Partner
tara.waters@ashurst.com

Gita Shivarattan Counsel
gita.shivarattan@ashurst.com

Ashurst LLP, London

Heralded as the great disrupter across hyperbolic headlines over the recent years, some may say that blockchain, or more correctly, distributed ledger technology, has underwhelmingly failed to deliver on its promise. Opinions on blockchain are often emotionally charged and divisive. However, most would agree that its benefits remain elusive and enigmatic for the everyday person.

Those actively working with blockchain still hold out hope, although the practical challenges of mass adoption remain many. One of the key challenges that blockchain faces is the conundrum of how a technology which purports to store data permanently and immutably can exist in a world of increasing regulatory obligations relating to data, particularly those dictating its amendment, correction, and deletion.

Can these seemingly irreconcilable concepts be reconciled? In the third part of the Emerging Tech Series, we consider the question of the coexistence of blockchain and data privacy law.

What is blockchain?

Blockchain is a technology that enables the secure validation, recording, and sharing of data. The data is stored in a distributed database, meaning there is not one centralised database controlled by a single person, but rather multiple copies of the database which are continuously updated in real time across the network of participants. Not only does this eliminate one single point of failure risk, but it also makes tampering with the data a significantly more onerous task, as a person would need to tamper with all copies of the data near simultaneously.

The emergence of blockchain is considered notable, if not revolutionary, due to manner in which the technologies underpinning it, none of which are new, are used together to allow parties to transact directly with one another without the need for a trusted third-party intermediary, and for data relating to that transaction to be stored in an extremely secure manner.

Fundamental to this security are various cryptographic methods that:

- firstly, convert the data into a coded form, which is referred to as a hash, that bears no resemblance to the original data;
- secondly, are designed such that the hash cannot be reverse-engineered, meaning it can only be decoded by guessing the underlying original data; and
- thirdly, store hashes in a manner which enables a user to easily confirm whether any of the original underlying data or hashes have been tampered with. This enables a user to trust the integrity of the data once stored, without necessarily having to trust its counterparts.

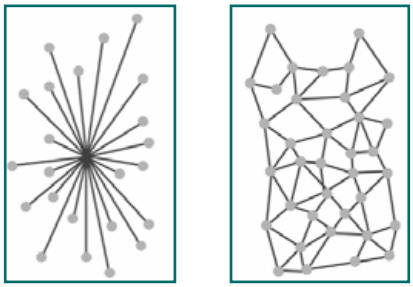

***Figure 1**: Centralised and distributed databases*

One of the key features of blockchain is its purported immutability, meaning that data stored in a blockchain-based database cannot be subsequently altered. Technically, this is not exactly correct, but because the hashes cannot be reverse-engineered that means it would generally require more computing power than is commercially available at present to not only guess, over and over, what the underlying data is, but to also then make the relevant changes across all copies of the database in the network as near to simultaneously as possible.

Moreover, blockchain networks implement specific rules, known as consensus protocols, most of which are designed to ensure that a single actor cannot unilaterally effect changes to the data. This is particularly true for widely distributed networks, such as the Bitcoin network and the Ethereum network, which were the first networks to be formed.

As participants in a blockchain network are working together to operate the network, there are limited incentives for those participants to join forces to enable tampering with the data. For all of the above reasons, blockchain-based databases are considered to be one of the most secure means of recording information.

Blockchain and data privacy

The fallacy of immutability

The immutability of blockchain is often held out as antithetical to data privacy laws, such as the General Data Protection Regulation (Regulation (EU) 2016/679) ('GDPR'), that empower data subjects to have control overw their personal data, including how it is collected and stored, and dictates that

persons collecting and storing such data must agree to hand over, correct, and delete that data on request.

data privacy compatible requires a case-by-case analysis.

that the relevant hashes are updated across the network. This would involve the participants in the private network

***Figure 2**: Blockchain*

In addition to the data subject rights, the GDPR also contains a principle on data minimisation, whereby organisations should only process personal data that is relevant and necessary for the defined purpose, and the principle of storage limitation, whereby organisations should only keep personal data for as long as necessary for the purposes for which it was collected. Both of these principles result in an obligation on the organisation that collected the data to either delete or anonymise any personal data once it is no longer necessary for the purpose.

However, focussing exclusively on the immutability of blockchain is an over-simplistic view. The truth is that there are many types of distributed networks that implement blockchain technology in a variety of ways.

Therefore, understanding whether a particular network is

And, it is possible that technological mechanisms can be built into the blockchain network and relevant consensus protocol to facilitate regulatory compliance.

Public vs. private blockchain

At the highest level, there are two type of blockchain networks: public and private. Public networks can be accessed by anyone and typically utilise consensus protocols that make modification of data near impossible. Private networks are limited to invited participants, and thus are more likely to apply consensus protocols that are more flexible in terms of enabling modifications.

In a private network context with a limited number of participants and other active users, it may be deemed appropriate, subject to the consensus protocol, to allow the computing power of the network to be used to enable the amendment of the database such

encoding specific rules for amending the database in certain agreed circumstances, which can be affected if the agreed consensus protocol threshold for those amendments is met.

Can hashed data be personal data?

Another threshold question when considering the compatibility of blockchain with data privacy is whether personal data stored in hash form would meet the definition of 'personal data' under applicable data privacy legislation.

In the EU, the answer to this question hinges on whether hashing the input data achieves anonymisation, which is not covered by data privacy laws, or only pseudonymisation, which is covered by data privacy laws, of that personal data. The Article 29 Working Party, in Opinion 05/2014 on Anonymisation Techniques, provided guidance that anonymisation, 'results from processing personal data in order to irreversibly prevent identification.'

In public/permisionless networks:	In private/permissioned networks:
• anyone can: – view the information on the ledger; – submit information to be recorded on the ledger; and – host the ledger; • participants are more likely to participate pseudonymously; • control is more likely to be fully decentralised; • there is an increased risk that a network participant may have malicious intent; and • there is a decreased risk that a network participant would have the computing power to carry out an effective attack.	• participants pre-selected or subjected to specified participation criteria or approval by an administrator (group); • control is likely to be more concentrated amongst certain participants or an administrator (group); • it is more likely to only be accessible by participants (but could be made available to the public or specific external entities); • it is expected to be more commonly adopted where recording sensitive/private information; • there is an increased risk if a participant has malicious intent because also more likely to have greater computing power vis-à-vis the network; and • there is a lower incentivisation to abuse any computing power.

***Figure 3:** Public v. Private networks*

The GDPR defines pseudonymisation as 'the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person.'

If a hash of personal data cannot be reverse engineered but only guessed by chance, does the hash function irreversibly prevent identifying the data subject by reference to that hash? The above guidance and definitions are not entirely helpful. Provided it remains possible to identify the underlying input data, then under the GDPR hashing results is pseudonymisation and not anonymisation. However, at the time of publication, this remains untested.

It is worth noting, however, that the GDPR clearly supports pseudonymisation as a security measure and risk mitigation technique and, therefore, there is a good argument that blockchain is a 'Privacy by Design and Privacy by Default' technology.

Options to avoid storing personal data on blockchain

It may be decided that any personal data is not stored in the blockchain network, but somewhere off-chain, to make regulatory compliance more straightforward and to avoid complicating the operation of the network. Utilising separate off-chain databases to store private information, not just personal data, is not uncommon. However, data privacy laws will apply to the collection and storage of that data, even if in an encrypted or pseudonymised form. It also follows that storing personal data in a single location may be more vulnerable to breach as well.

Controllers and processors

Under the GDPR, organisations that are processing personal data are categorised as either a controller, therefore they are deciding the means and purposes of the processing, or a processor, as they are only undertaking processing on behalf of a controller and under the controller's instructions. If the entry in a block contains personal data, participants may be acting as both a controller as they are writing on the chain, and miners acting as processors, in validating the entry. As a result, careful analysis is required to determine the roles of participants as controllers or processors on a case by case basis. The French data protection authority ('CNIL'), has released guidance in which it sets out an assessment of when blockchain participants are acting as data controllers and processors, which is a useful reference when assessing these roles.

The terms and conditions for participation in the blockchain will need to accurately identify the roles of the parties and allocate responsibility for issuing fair processing notices, responding to data subject requests, handling data breaches, Article 28 of the GDPR processor clauses, and liability.

Data Protection Impact Assessments

Prior to setting up a blockchain or entering into one, organisations should undertake a Data Protection Impact Assessment to help identify potential risks in respect of the technology and solution.

Conclusion

If blockchain and data privacy are not irreconcilable, can they coexist?

Hopefully, the above makes clear that there are several considerations when seeking to understand how a blockchain network must be set-up and operated in order to enable regulatory compliance. By design, certain types of networks, such as public networks, are less compatible with data privacy principles, but compliance is still possible.

However, the trade-off when implementing a consensus protocol that allows for the editing and deletion of data may be that the network is more readily interfered with by a single actor or group of actors. Ultimately, though, to achieve compliance enabling such measures will be necessary.

So, there may remain an uneasy coexistence until such time as regulators provide more clarity about how to implement blockchain in a fully compliant manner. Whether that clarity will be forthcoming remains to be seen.

1. Available at: https://www.cnil.fr/sites/default/files/atoms/files/blockchain.pdf

Global Regulatory Research Software

40 In-House Legal Researchers

500 Lawyers Across 300 Jurisdictions

With focused guidance around core topics, Comparison Charts, a daily customised news service and expert analysis, OneTrust DataGuidance provides a cost-effective and efficient solution to design and support your privacy program

- Legal Guidance & Opinion
- Law Comparison Tools
- Breach & Enforcement Tracker
- Ask-An-Analyst Service

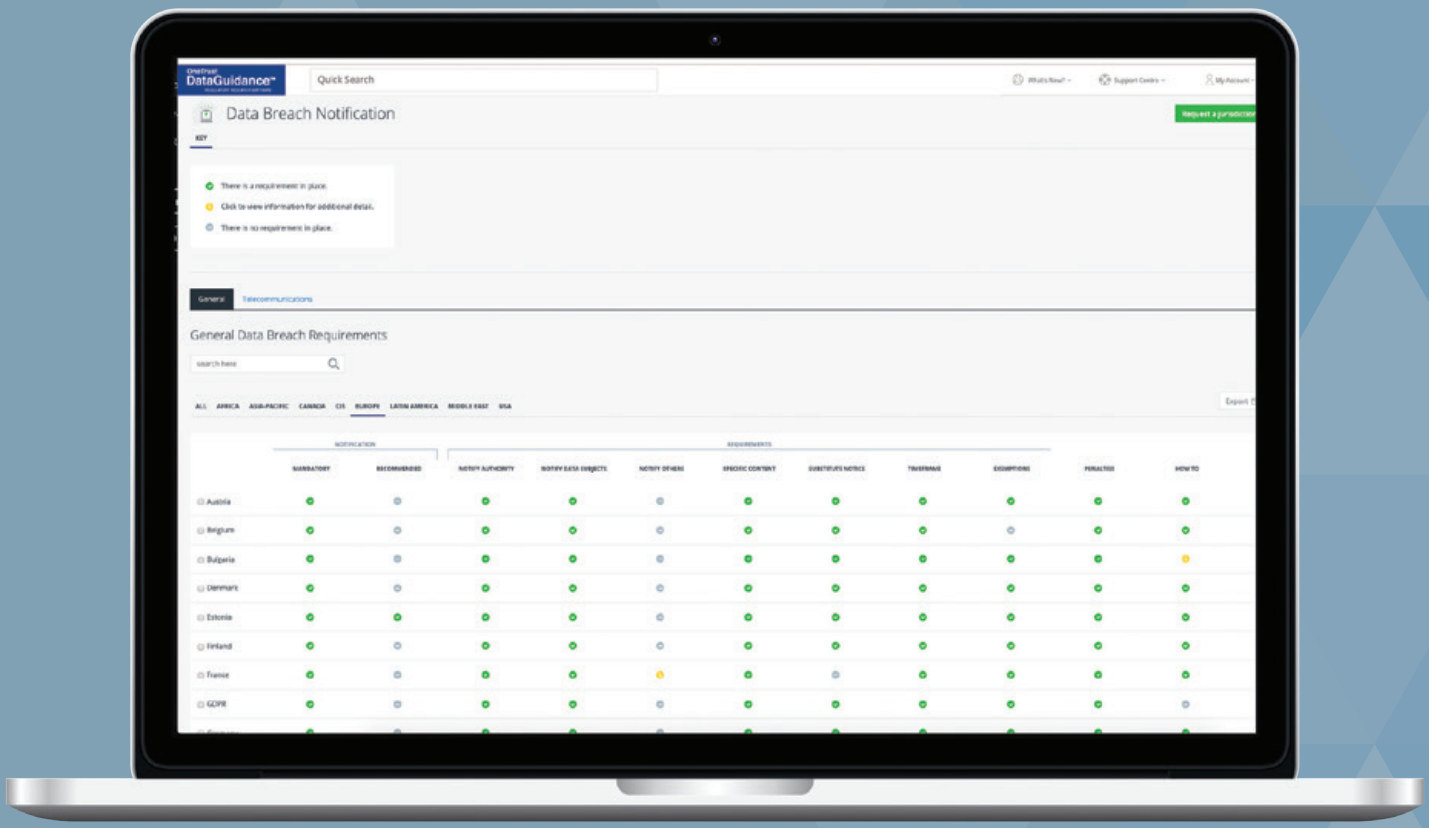

SCAN TO ACCESS
FREE TRIAL
Use your camera or a QR code reader

OneTrust DataGuidance™
REGULATORY RESEARCH SOFTWARE

Digital advertising

Gita Shivarattan Counsel
gita.shivarattan@ashurst.com

Tom Brookes Solicitor
tom.brookes@ashurst.com

Ashurst LLP, London

Fuelled by consumer expectations of receiving free services, information, and products, global digital advertising spend reportedly surpassed $100 billion dollars in 2018[1] with its revenue predominantly driven by 'ad tech.' At its core, ad tech refers to tools that analyse and manage information for online advertising companies and automated processing of advertising transactions.

Despite being a relatively mature industry, ad tech is one of the sectors to be hardest hit under the General Data Protection Regulation (Regulation (EU) 2016/679) ('GDPR') because it is an industry powered by cookies, the use of which often involves the processing of personal data. The current lack of clarity in relation to the interplay between the Privacy and Electronic Communications (EC Directive) Regulations 2003 ('PECR') and the GDPR has further had a disproportionately problematic impact on the industry.

Ad tech is in the regulatory spotlight as a result of several high-profile investigations into the technology, and the *Financial Times* reported that 12 data protection supervisory authorities are investigating complaints relating to ad tech. Earlier this year, the Information Commissioner's Office ('ICO') issued an updated report into ad tech and real-time bidding[2] ('RTB') ('the ICO Report') which has become a regulatory priority area for data protection, competition, and consumer protection.

In this fourth and final part of the Emerging Tech Series, we examine some of the main data protection challenges companies operating in the ad tech space are grappling with and look ahead to what the future may hold.

RTB – what is the big deal?

In its most basic form, advertising involves providing information about new products and services and digital advertising delivers this information through the medium of the internet. Whilst the first webpage banner advertisement was created back in 1994[3], the explosion of internet use due to increasingly ubiquitous connectivity has led to exponential growth in digital advertising.

RTB is an auction process that occurs in 'real-time' to sell visual advertising on websites and apps. Whilst there exist a variety of other forms of digital advertising and ad tech, this article focuses on programmatic advertising and RTB due to its widespread adoption[4] and the level of regulatory attention that RTB, in particular, has received to date.

In essence, RTB involves an operator of an online service (a 'Publisher') selling space on their website to be filled by the content of advertisers as a result of successful bids on a per end user basis.

The bid process relies on cookies or other technologies, such as pixels, that collect information from an end user's device when they visit a Publisher's website. This information is often enriched by a data management platform ('DMP'), which collates known information or infers

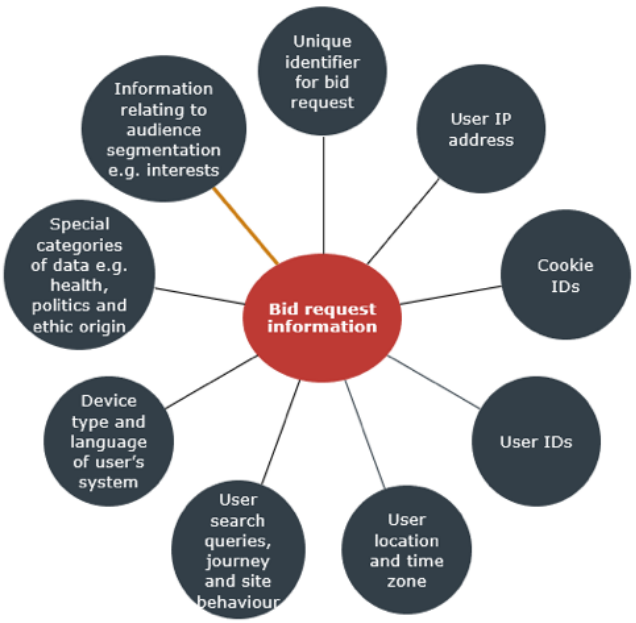

***Figure 1**: Information included in bid request*

<table>
<tr><th colspan="4">RTB Participants</th></tr>
<tr><td>Publisher

Organisations who have advertising space (inventory) on their websites and platforms and apps to sell to Advertisers</td><td colspan="2">Advertising Exchanges and Severs

The location where bidding takes place. Mediates between Publishers and Advertisers and operates on both buy and sell sides.</td><td>Advertiser

Organisations who want to broadcast information on their products and services to consumers on a Publisher's advertisement space</td></tr>
<tr><td></td><td>Sell-side platform ('SSP')

Platform to help Publishers manage and sell their inventory</td><td>DSP

Platform used by Advertisers to place bids for inventory space on websites and platforms of Publishers</td><td></td></tr>
<tr><td colspan="4">DMP

Platform which analyses and combines data including personal data from multiple sources to facilitate targeted advertising personalised to an individual consumer</td></tr>
<tr><td colspan="4">Consent management platform ('CMP')

A tool which manages consents e.g. of individuals using a Publisher's website or platform</td></tr>
</table>

***Figure 2**: RTB Participants*

additional information about the end user, thereby making the 'bid' more accurate, and therefore more valuable.

Multiple advertisers then place bids for the opportunity for their digital advertisement to be displayed to the end user in the advertisement inventory space on the Publisher's website. Advertisers are aiming for their adverts to be displayed to individuals who are most likely to purchase their products, which is reliant on the information collected about the individual in the RTB process.

The information contained in the bid request may contain varying amounts of known and inferred personal data relating to the end user. According to the ICO Report[5], the categories of information included in **Figure 1** (above) can potentially be included in a bid request and will constitute personal data as defined under Article 4(1) of the GDPR, where they enable an individual to be identified either directly or indirectly.

The ad tech ecosystem

Set out in **Figure 2** (above) is an overview of the range of parties who are commonly involved in the RTB process ('RTB Participants'). The ad tech operating model is complicated by the fact that one organisation could potentially wear 'a number of hats' in the ad tech process, for example, having a demand-side platform ('DSP'), an Advertising Exchange, and a DMP.

Core data protection issues highlighted by the ICO Report

In the UK, there are two main data protection legislations which RTB Participants are required to adhere to, PECR and GDPR. PECR governs the use of cookies and other technologies on an end user 's device, as this constitutes the use of an electronic communications network to store information. The GDPR governs personal data, and cookies will often involve the processing of personal data.

In July 2019, the ICO published Guidance on the Use of Cookies and Similar Technologies[6], which explains how PECR applies to the use of cookies and how PECR interacts with the GDPR. However, the EU data protection regulatory landscape is currently in a state of flux, with the underlying directive of PECR (Directive 2002/58/EC) set to be replaced by a new e-privacy regulation, once finalised by the EU institutions.

Lawful basis

Where personal data is processed, a lawful basis is required under the GDPR. In the context of RTB, the lawful bases traditionally relied on are either consent of the individual whose personal data is being processed, or legitimate interests of the RTB Participant. However, the ICO Report noted there was a lack of clarity over which lawful basis many RTB Participants were relying on.

For any processing of special categories of personal data, such as information about an individual's political opinions, religion, health

information, or ethnic group, the GDPR requires the explicit consent of the individual to be obtained, unless specific exemptions apply[7]. These exemptions are not applicable in the context of RTB. The ICO Report noted[8] that a proportion of RTB bid requests involve the processing of special categories of personal data, and found that consent requests which they had reviewed were not compliant with the GDPR, and needed to be modified to ensure explicit consent was collected prior to the processing of personal data.

In relation to non-special categories of personal data processed in the RTB process, the ICO stated that it believes the nature of the processing makes it impossible to meet the legitimate interests lawful basis requirements[9] for the main bid request processing[10].

This means that the lawful basis for processing involved in collecting personal data from an end user, and the onward transfer of that personal data in the bid request, is consent. In contrast with the GDPR, consent is required under PECR in order to drop cookies on the user's device, unless they are strictly necessary[11].

The consent requirements under both the GDPR and PECR are very specific and pose a number of challenges for RTB Participants. In particular, consent must be[12]:

- unambiguous, meaning pre-ticked boxes surfaced on a Publisher's website cannot be used;
- specific, meaning the consent must be obtained for each processing operation and be clearly distinguishable from other matters. Therefore one 'bundled' consent for each aspect of processing involved in the RTB process is not possible; and
- freely given, meaning that the individual must be given the choice of accessing a website or application without tracking cookies being dropped on their device.

It is common for CMPs to be used by RTB Participants, and in particular Publishers, to manage consents from users of their websites and applications. The Interactive Advertising Bureau's Transparency and Consent Framework[13], which is an industry initiative to assist RTB Participants in complying with the GDPR, relies on CMPs to obtain, store, and signal consents. However, RTB Participants need to carefully consider whether the consents obtained by CMPs meet the GDPR requirements discussed above if they are relying on these consents for their lawful basis to process personal data.

The ICO did not provide guidance regarding the lawful basis which could be relied on in relation to other aspects of processing which take place in the RTB process after a bid request is made, such as processing by advertisers, DSPs, and SSPs, who have no direct relationship with the end user whose personal data is being processed. Regardless of the lack of guidance, all RTB Participants need to undertake a careful assessment of the lawful basis they choose to rely on and ensure this is documented.

Transparency and fair processing notices

The GDPR requires transparency in relation to how personal data is processed and Articles 13 and 14 of the GDPR set out specific information which must be provided to individuals. PECR also requires clear and comprehensive information about the cookies and other technologies which are dropped on a device to be provided to the relevant individual[14].

In the context of RTB, the primary challenge faced by RTB Participants is being able to clearly describe to individuals, using clear and plain language, the complex processing operations and data flows which are taking place. These data flows often involve automated processing of large volumes of data for various purposes such as targeting, fraud prevention, analysis, and measurement, which requires both careful explanation and presentation in the privacy notice.

There is a tension between providing information which is either too granular or too high level, and in order to comply with these transparency requirements, it is essential that RTB Participants have clarity about:

- how their processing operations work, including what is the purpose for collecting each type of personal data;
- who they share any personal data with; and
- how they are enabling individuals to exercise their rights in relation to this processing.

Intrusive and unfair processing

The ICO Report also noted that during the RTB process, bid request information is often combined and enriched by creating a profile of an end-user using information gathered from other sources, such as DMPs.

The ICO's main concern with these activities is that this may constitute unfair and intrusive processing due to the quantity and nature of the personal data being processed, which appears to be disproportionate to the purpose of delivering targeted advertising[15]. Another key concern is the fact that individual users may not be aware that this combining and enrichment is taking place if fair processing notices do not clearly inform individuals what is happening.

Pursuant to Article 35(4) of the GDPR, the ICO has published a list of processing operations[16] likely to require a data protection impact assessment ('DPIA') to be undertaken, which includes data matching for the purposes of direct marketing. RTB Participants involved in information enrichment processes will need to conduct a DPIA in order to identify and minimise the data protection risks relating to this processing.

Other processing activities for which the ICO deems a DPIA mandatory include large scale profiling of individuals, tracking of an individual's location and behaviour, and invisible processing where personal data which is being processed was not obtained directly from the individual

and the organisation does not notify individuals of the processing due to a perceived disproportion effort[17].

Given the activities included on this mandatory list, all RTB Participants need to carefully consider whether they need to undertake DPIAs in relation to their processing of personal data.

Next steps

The ICO Report served notice on the ad tech industry that there are serious concerns about data protection compliance in relation to the RTB process, and highlights that the ICO wants the industry to take the initiative of reforming their activities. This raises difficult questions for each and every RTB Participant, such as how much personal data is actually necessary for the purposes in which they are using it, how can this personal data be collected and shared lawfully, and how can individuals be clearly informed about what is happening.

Conducting detailed DPIAs should form the starting point in assessing how to deal with these issues, however, time appears limited for the ad tech industry to come up with the answers. The ICO made it clear when the ICO Report was published in June 2019 that it would conduct a further industry review in six months' time. This could result in enforcement activity and potentially sanctions (including fines) given the nature of the non-compliance which was revealed in the ICO Report. Such a review could also focus on other aspects of non-compliance with the GDPR, which were referenced in the ICO Report but not considered in detail, such as data minimisation and data retention.

1. Available at: https://www.iab.com/wp-content/uploads/2019/05/Full-Year-2018-IAB-Internet-Advertising-Revenue-Report.pdf
2. Available at: https://ico.org.uk/media/about-the-ico/documents/2615156/adtech-real-time-bidding-report-201906.pdf
3. Available at: https://www.wired.com/2010/10/1027hotwired-banner-ads/
4. Almost 90% of digital display advertising in the UK is programmatic according to emarketer, available at: https://www.emarketer.com/content/programmatic-ad-spending-in-the-uk-2019
5. Available at: https://ico.org.uk/media/about-the-ico/documents/2615156/adtech-real-time-bidding-report-201906.pdf (section 2.6 pages 12-13)
6. Available at: https://ico.org.uk/for-organisations/guide-to-pecr/guidance-on-the-use-of-cookies-and-similar-technologies/
7. Article 9 of the GDPR.
8. ICO Report page 16 section 3.2.
9. To rely on legitimate interests, organisations need to identify a legitimate interest, show that the processing is necessary to achieve that interest, and balance that interest against the individual's interests, rights, and freedoms.
10. ICO Report page 17 section 3.3.
11. Regulation 6 of PECR.
12. Article 7 and Recital 32 of the GDPR.
13. Available at: https://iabeurope.eu/transparency-consent-framework/
14. Regulation 6(2) of PECR.
15. ICO Report page 20 section 3.4.
16 ICO Examples of processing 'likely to result in a high risk,' available at: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/data-protection-impact-assessments-dpias/examples-of-processing-likely-to-result-in-high-risk/
17. Ibid.

To watch all four episodes of the Emerging Tech series, in partnership with Ashurst, and other video content featuring regulators, privacy professionals, private practice lawyers and more visit the OneTrust DataGuidance video hub for free at:

platform.dataguidance.com/videohub

ABOUT ASHURST

A leading international firm

Ashurst is a leading international law firm with world class capability and a prestigious global client base.

Ashurst has **27 offices** in **16 countries** and offers the reach and insight of a global network, combined with the knowledge and understanding of local markets. With **417 partners** and a further **1300 lawyers** working across **10 time zones**, we are able to respond to our clients wherever and whenever required. As a global team, Ashurst has a reputation for successfully managing large and complex multi-jurisdictional transactions, disputes and projects while delivering outstanding outcomes for our clients.

Our vision is to be the **most progressive global law firm**. For us, 'progressive' is a mindset, an approach to how we do things. We instinctively take a fresh perspective on situations, exploring whether there are better ways of delivering practical, commercial solutions to the challenges our clients face in today's rapidly changing business landscape.

DIGITAL ECONOMY GROUP

Our Digital Economy Group works to cut through the complexities of digitalisation to provide practical, commercial legal solutions to deal with your business problems regardless of the uncharted nature of the legal issues they present.

We adopt a tailored approach to the digitalisation of particular industries with our core Digital Economy Group working hand in hand with our sector specialists. We have a key focus on fintech, infratech, proptech and TMT, but recognise that digitalisation is affecting all sectors of the economy from health to agriculture.

We provide legal solutions across all key disciplines including strategic commercial arrangements, M&A, IPOs, financing, joint ventures, outsourcing, disputes, exploitation of IP, data protection, tax and other key commercial activities.

We operate across the globe – so we can help you wherever the strategy applies or the problem arises.

Nick Elverston
Practice Group Head, Digital Economy Transactions

T +44 20 7859 3143
M +44 7823 340 890

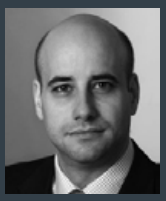

David Futter
Partner

T +44 20 7859 1594
M +44 7823 340 950
david.futter@ashurst.com

"Professional, approachable, collaborative, and they consistently exceed expectations."

CHAMBERS AND PARTNERS, 2020