Regulatory Sandbox Insights Report 2024

Date: July 2024

Information Commissioner's Office

Contents

1. Introduction

What is the Regulatory Sandbox?

1.1 The Regulatory Sandbox ('the Sandbox') is a free service developed by the ICO to support organisations in the development of products and services that use personal data in innovative and safe ways, while ensuring compliance with the UK General Data Protection Regulation (UK GDPR) and the Data Protection Act 2018 (DPA 18).

1.2 Participants have the opportunity to engage with our dedicated Sandbox team, and to draw upon our wider ICO expertise and advice on embedding 'data protection by design' and mitigating data protection risks associated with their innovation.

1.3 The Sandbox is a free, professional service for organisations, of varying types and sizes, across a number of sectors. You can read more about how the Sandbox operates, its benefits, and how to apply in our guide to the Sandbox.

Where is the Sandbox now?

1.4 Since August 2021, the Sandbox selected the following areas of focus to explore the specific data protection challenges encountered by innovators:

- Innovations related to the Children's code.
- Innovations related to data sharing (particularly within the health, central government, finance, higher and further education, or law enforcement sectors).
- Products and services exploring the use of innovative technologies, such as privacy-enhancing technologies and distributed ledgers.

You can read more about the reasons why we selected the above areas of focus in the 'What's next?' section of our beta phase review. Following a call for expressions of interest in Autumn 2023, the Sandbox has successfully onboarded new projects under our updated areas of focus which include emerging technologies, biometrics and exceptional innovations. You can read more about our new areas of focus by visiting Our current areas of focus for the Regulatory Sandbox page. Please see our webpage on how to apply for the Sandbox, , if you are interested in participating.

1.5 This report details the work we have completed in relation to the above areas of focus. There were 14 participants in total from the private, public and third sectors.

2. Overview of Participants

2.1 The 14 participants selected to participate in the Sandbox were as follows:

- **Seers**

Seers entered the Sandbox with its privacy and consent management platform which intends to respect children's privacy. The Child Privacy Consent Management Platform (CPCMP) aimed to enable children, or their parents or guardians dependent on their age, to provide informed consent for the child's data to be processed by the cookies and scripts operating on the website they visited.

- **Gambling Commission**

The Gambling Commission entered the Sandbox to explore the concept of a Single Customer View (SCV). An SCV would allow data, which already exists about individual player behaviours, to be aggregated to drive better decision making, actions and evaluation with the ultimate aim of reducing gambling related harms.

- **Global Cyber Alliance**

The Global Cyber Alliance (GCA) entered the Sandbox to explore the next steps in developing their DomainTrust concept and the potential benefits of processing personal data in the platform to help fight against cybercrime.

- **Yoti**

Yoti entered the Sandbox to consider how they could extend the use of their age estimation technology to young people, aged between six and 12 years old. This would ensure that the providers of children's only services, such as gaming websites and forums, can create safe virtual environments and online spaces.

- **FlyingBinary**

FlyingBinary entered the Sandbox during the development of an online service which seeks to assist with the traditional mental healthcare of patients with ongoing pathologies such as eating disorders. 'lookafterme' is an online web intervention service, recommended for use by the patient's clinician, which aims to support the management of a patient's interaction with online content.

- **CDD Services**

CDD Services entered the Sandbox to examine SafeGuarden, a digital-ID centred online platform, which uses CDD Services' Spotlite Compliance Platform. The product intended to enable former military personnel to prove their identity and grant permission to share their personal data with Service Providers, who would then provide support to them and their families in relation to housing, employment and training.

- **Good With**

Good With entered the Sandbox whilst developing a Financial Virtual Assistant (FVA) app which intends to produce a 'financial readiness score' that helps to provide the FVA's users (18-24-year-olds) with fairer access to financial products and services. The readiness score will be informed by insights drawn, using Artificial Intelligence (AI), from the user's interactions with the FVA's chatbot, their open banking data and their progression through a bespoke educational pathway.

- **Betting and Gaming Council**

The Betting and Gaming Council (BGC) entered the Sandbox to explore the gambling industry's development and trial of an SCV solution. The trial built on the work previously undertaken by the Gambling Commission within the Sandbox. The SCV solution developed by the BGC and the participating operators (Entain, William Hill, 888, Gamesys, Bet365 and Flutter) aims to enable a more unified and proactive intervention by gambling operators to reduce incidents of gambling related harm.

- **Crisis**

Crisis entered the Sandbox as the lead organisation in a consortium with representatives from homelessness services in local authorities. This group has a broad objective of developing and delivering county-wide data sharing to support the prevention and relief of homelessness. Crisis believe that the creation and use of a By-Name-List (BNL) is one way of delivering this objective.

- **Smart Data Foundry**

Smart Data Foundry entered the Sandbox while developing a National Data Utility (NDU) where data will be maintained within a 'safe-haven' infrastructure. It aims to unlock the potential of financial information to enable research, innovation and skills development in financial services and FinTechs. Alongside the NDU, Smart Data Foundry aims to develop an innovation environment to support further research using synthetic data.

- **Our Future Health**

Our Future Health entered the Sandbox while aiming to develop the UK's largest ever health research programme. It intends to recruit a large cohort, who will be demographically representative of the wider UK population, to participate in the research programme by voluntarily providing their personal data. This includes a sample of their blood, which will be genetically analysed, alongside the answers to a detailed health and lifestyle questionnaire. Our Future Health will make that information available to researchers, via Trusted Research Environments, to improve the prevention, detection and treatment of common diseases such as cancer, stroke, dementia and diabetes.

- **Zamna**

Zamna entered the Sandbox to examine how it uses novel cryptography, decentralisation principles and blockchain technology to deliver an inclusive, privacy-preserving solution for pre-airport checks (eg identity and health status verification), thereby enabling passengers to travel seamlessly and securely.

- **Thames Valley Police**

Under the Serious Violence Duty (SVD), as found in the Police, Crime, Sentencing and Courts Act 2022 (PCSC 2022), police services, local authorities, fire and rescue authorities, specified criminal justice agencies and health authorities are required to work together to formulate an evidence-based analysis of the problems and causes of serious violence in a local area. Thames Valley Police (TVP), as a lead authority, entered the Sandbox with Thames Valley Together (TVT) which is an early example of a local partnership aimed at operationalising the SVD. Combining information shared by specified authorities across the Thames Valley region, TVT aims to identify the causes of serious violence and develop products designed to tackle them. TVT is likely to inform and influence similar violence reduction partnerships in England and Wales.

- **Financial institutions – facilitated by the Home Office and UK Finance**

A group of participating banks (financial institutions) entered the Sandbox whilst developing an 'information sharing pilot'. Facilitated by the UK Home Office and UK Finance, the participating banks intended to share certain personal data of customers who were deemed to pose a risk of financial crime to a central database in a peer-to-peer way. This aimed to help the group identify information relating to the potential financial crime risk presented by their customers, or potential customers, that they would have otherwise not known. The 'information sharing pilot' was intended to supplement existing business-as-usual processes to better manage financial crime risks.

2.2 The exit reports for these participants can be found on our website.

3. Key Data Protection Considerations

3.1 In this section of the report we discuss some of the key data protection considerations that were prominent among this cohort of Sandbox participants.

Joint Controllership

3.2 It is crucial, for any organisation seeking to comply with the UK GDPR, to appropriately establish its data protection roles and responsibilities for each processing activity it undertakes. The Sandbox team helped a number of its participants to assess whether they will be acting as controllers or processors for certain processing activities. In particular, various participants that entered the Sandbox under our data sharing area of focus intended to use intricate personal data ecosystems, involving a number of parties with differing degrees of influence over the determination of the purposes and means of processing activities. As a result, the Sandbox regularly determined that, in relation to certain processing activities, participants were likely acting as joint controllers with other organisations.

Why is it important for innovators to break the processing of personal information into granular activities?

3.3 A key insight gained from the Sandbox's work on joint controllership underscores the importance of breaking down the processing of personal data into specific stages for controller(s). Breaking down the activities can help controller(s) establish if, and where, they are independently or jointly determining the purposes and means of the processing of personal data. This is particularly important in scenarios where more than one organisation is involved in the activity. It can also help to ascertain where that assessment may change between activities.

3.4 The Sandbox's work with the financial institutions highlighted how assessments of controllership may change in relation to different processing activities where more than one controller is involved in the wider processing of personal data. First, the Sandbox assisted the participants to break down the processing activities in a granular way. We then helped the participants to identify which processing activities were likely to result in joint controllership relationships, where the participating banks

jointly determine the purposes and means of the processing. We also helped the participants identify which activities were more likely to be carried out by an independent controller. For example, a participating bank may use information from the central database separately, as part of its own existing due diligence processes, to help manage the risk of financial crime. This constitutes a specific and separate processing activity, distinct from the joint task of storing information on the central database to identify individuals alongside the other banks. For this activity, where that bank uses that information for its own purposes, such as to adhere to its own regulatory requirements, it is likely to be acting as an independent controller. This work was key to help emphasise how assessments of controllership can differ and transition between individual processing activities.

3.5 As we also found with our work with TVP as part of its TVT project, it is important for controllers involved in complex data sharing projects to clearly differentiate their processing activities and explain where exactly processing is subject to a joint controllership arrangement. This is particularly true in circumstances where one or more controllers within a project process personal data for law enforcement purposes, under Part 3 of the DPA 18. It is important to remember that joint controllership arrangements under the UK GDPR/Part 2 DPA 2018 do not apply to the processing of personal data by a competent authority for any of the law enforcement purposes. Similarly, Part 3 DPA 18 joint controllership arrangements apply only where "*two or more competent authorities jointly determine the purposes and means of processing personal data*". Organisations that process personal data under both UK GDPR/Part 2 DPA 18 and Part 3 DPA 18 should take extra care to separate their processing activities and ensure that their joint controllership arrangements apply to the appropriate data protection regime, particularly in data sharing projects which involve competent and non-competent authorities.

Do joint controllers need to share the exact same purpose for processing?

3.6 Our work with Zamna allowed us to explore the role of 'purpose' in joint controllership. For joint controllership to apply, organisations do not necessarily need to process personal data for the same purpose if the purposes of the joint controllers are complementary and inextricably linked. Controllers may have different motivations for participating in a joint processing activity that are complementary to each other although separate. For instance, Zamna's purpose for persisting cryptographic signals in their cloud was to enhance the commercial marketability of their product. Meanwhile, the airlines in Zamna's

Network were persisting signals in the cloud for the purpose of verifying the consistency of their passengers' travel documentation. Although these joint controllers have different purposes, the processing is also inextricably linked, meaning it cannot take place without the involvement of both parties. In this example, the airlines provided the identifier to retrieve the information from the signals, while Zamna persisted the signals and hosted them in its cloud. Zamna's signals would not be of value without the airlines' identifier and the airlines would not be able to verify subsequent travel unless the signals were persisted by Zamna.

3.7 We can distinguish the situation above from a scenario in which two organisations independently acquire access to the same personal information and utilise it for their own purposes. In this context, each organisation could carry out the processing without the need for involvement from the other, making them less likely to be considered joint controllers.

How important is the 'means' in determining joint controllership?

3.8 In the event that an organisation does not decide the purpose of the data processing activity, but has control over some of the essential means of processing, a joint controllership arrangement is likely required. We explored this in our work with Smart Data Foundry which looked at how joint controllership might arise in a research context. In particular, our engagement looked at the question, what was Smart Data Foundry's role when they have been commissioned to carry out research on behalf of an organisation that is already a controller for that data and has decided the purpose of the research? The key focus was understanding the extent of Smart Data Foundry's influence over some of the essential means of processing. Smart Data Foundry had a lot of control over the research design such as the categories of personal data that would be collected and how this information would be used to conduct the research. These are the types of decisions made by a controller. Meanwhile, the commissioning organisation also acts a controller because it retains control of the data and determines the purpose by commissioning the research. As both these organisations exercise influence over either the purpose or means of the research activity, they are acting as joint controllers.

Transparency

3.9 Article 5(1)(a) of the UK GDPR requires that personal data is processed in compliance with the lawfulness, fairness and transparency principle. Articles 13 and 14 establish an individual's right to be informed about how their personal data will be processed. These Articles relate to the UK GDPR's transparency requirements when personal data has been collected from an individual, and when it has not been obtained directly from an individual respectively. Transparency requirements have regularly formed part of the work the Sandbox has carried out with its participants.

How should innovators consider their consumers' expectations when implementing transparency?

3.10 Transparency requirements formed an important part of the Sandbox's work with Good With. Good With's specific use case makes use of novel indicators of 'financial readiness', such as the user's progression through a bespoke educational pathway, which people may not traditionally expect to be used to assess their suitability for access to financial products and services. As a result, our work reiterated the importance of providing effective transparency messaging because users who would not reasonably expect information to be used in a certain way may be prevented from making an informed decision about which elements of the service to engage with and, consequently, the amount of information they provide. Good With's model of offering users choice over engaging different elements of its FVA also highlighted the importance of providing a layered approach at key touch points in the user journey. This will help to supplement transparency messaging and provide users with awareness as to how their engagement with different elements may change the way their personal information is processed.

3.11 Transparency was also a key issue in the TVT project. Transparency is crucial in building public trust and helping the public to understand why a processing activity is taking place, particularly where there are multiple partners and the processing is new or complex. Organisations should work on the principle that the newer the processing, the less likely the public will expect its occurrence. While the TVT project was aligned conceptually with the new SVD under the PCSC 2022, the Sandbox highlighted the importance of clearly communicating to the public the underlying objectives of the processing. Consequently, we collaborated with TVP to devise communication strategies that aimed to effectively reach all people who might be impacted or included in the new processing activities of TVT. Given the nature and scale of the data sharing, we advised

against relying solely on a single privacy notice. Instead, we recommended adopting a layered approach to delivering privacy information, including making information available at key service touchpoints and holding public engagements.

How can innovators address transparency when unsolicited personal information is collected?

3.12 The Sandbox's work with Good With identified challenges in providing transparency messaging related to the processing of unsolicited personal data. Good With's FVA provides a free form chatbot, which means the user has discretion over entering unstructured information into it. As a result, Good With identified a risk that users may provide information to the chatbot which includes special category data; information that Good With neither required or wanted. To mitigate that risk Good With intends to use an automatic redaction procedure, which would likely involve the processing of personal data itself, to remove any information that is not required for the FVA's purpose. The Sandbox advised Good With to also ensure that its transparency messaging informs users what information they should provide to the chatbot and why, as well as detail on the types of information that are not required and will be redacted. These transparency messages could be supplemented by using techniques such as just-in-time notices. This use case highlighted the importance of organisations considering how they will comply with data protection requirements should they process personal data they do not necessarily expect to.

Lawful Basis and Special Category Data

3.13 To process personal data lawfully a controller must identify an appropriate lawful basis for processing, under Article 6 of the UK GDPR, for its various processing activities. Where that processing includes special category data an additional condition for processing, under Article 9, must also be appropriately identified.

When might consent not be an appropriate lawful basis?

3.14 It should be remembered that no single lawful basis is 'better' or more important than the others – which basis is most appropriate to use will depend on an organisation's processing purposes and its relationship with the individual. For example, the lawful basis of consent may not be appropriate in the context of children and safeguarding matters. The Sandbox's work with Crisis UK explored this matter and found that consent may not be the most appropriate lawful basis to

use for processing personal data in the context of a BNL. As the BNL is a list accounting for the journey of every person experiencing homelessness within a region, it relies on a number of parties contributing real-time data. As such, we identified a risk that obtaining genuine consent could be challenging for local authorities. Some of the challenges of consent include:

- There may be an imbalance of power between the organisations collecting personal data and providing a much-needed service on the one hand, and the people relying on these services to secure safe accommodation/housing on the other. Where people feel unable to consent without detriment (eg losing access to housing services), the consent is unlikely to be valid as it will not be considered freely given.
- While consent might appear to provide people with greater control, the organisations engaged in the BNL processing may prioritise maintaining their relationship and providing support, even if consent is declined. This can lead to situations where the individual's preference is overridden.
- BNL processing involves a lot of onward data sharing. In such a context, consent could become easily fragmented and challenging to manage, particularly if a person decides to withdraw consent.

3.15 The Sandbox's work with Our Future Health provided a use case to consider that consent may not always be appropriate in a research context. Our Future Health intended to rely on consent, and explicit consent, to collect and store personal data which would subsequently be used in health research. Referring to the ICO's research provisions guidance on consent, the Sandbox informed Our Future Health that requiring other consents (such as separate ethical or legal to participate in research) does not necessarily mean it should use consent to process the personal data of its research participants. In fact, the Sandbox helped identify some reasons within this specific use case why consent may not appropriate. They include, but are not limited to, that:

- There could be a risk that the various consents Our Future Health sought to obtain were 'bundled', depending on how they were collected. This could result in consent not being valid if its participants were confused about what they were consenting to; and

- Our Future Health advised the Sandbox that it would be unable to remove a research participant's pseudonymised personal data from research that was carried out before a person decided to withdraw from the research programme. This could also result in consent being invalid if people could not fully withdraw their consent.

These observations helped Our Future Health reassess its intention to use consent to collect and store its research participants' personal data. It now intends to use legitimate interests (Article 6(1)(f)) and scientific research purposes (Article 9(2)(j)) for this processing activity.

How can innovators consider the UK GDPR's requirement for processing to be generally lawful?

3.16 During the financial institutions' Sandbox project we stated that, for any processing to be lawful under the UK GDPR, personal data must not be used in a way that is unlawful in a general sense. For example, it must not be processed in a way that is in contravention of any other statute or common law obligations, whether criminal or civil. During their time in the Sandbox, the participants informed us that the participating banks owe a general 'duty of confidentiality' to their customers (subject to certain exemptions). However, the group informed us that, in this specific scenario, the duty of confidentiality is not applicable and they would not be in breach of it requirements. While the duty of confidentiality does not fall under the remit of the ICO, these discussions helped the participants consider the UK GDPR's requirement for the processing of personal data to be generally lawful.

Can special category data be inferred?

3.17 Our work has helped organisations overcome some of the challenges posed by processing special category data. The Sandbox and the BGC collaborated to determine that special category data was being processed in the gambling industry's development of an SCV. The SCV is a mechanism by which gambling operators can share personal data about customers considered to be experiencing gambling related harm. The success of the SCV relies on a limited amount of personal data being shared between gambling operators. Although the data items in isolation do not constitute special category data (eg name, date of birth, contact details etc), it is important to consider the potential inferences that can be drawn. Since problem gambling is recognised by medical literature as a mental health issue, coupled with the disclosure that gambling

operators were being informed about customers deemed unable to control their gambling to such an extent that a platform ban was necessary, we determined that health data was being shared. This meant that the data in the SCV would need to have an Article 9 basis, as well as an Article 6 basis.

What happens if innovators process special category data that has been unexpectedly provided?

3.18 An earlier section of this report (3.12) details the challenges Good With faced in relation to the processing of unsolicited personal data it neither required or wanted, provided by its FVA's users via the free form chatbot. The ICO helped Good With to understand that its automatic redaction procedure would still involve the processing of personal data as it would be collecting, and subsequently deleting it. Where this includes special category data an appropriate condition for processing under Article 9 of the UK GDPR must be identified. This work helped demonstrate that where it is reasonably foreseeable that people may provide special category data an organisation should think in advance about how to justify its collection of sensitive information, including identifying an appropriate condition for processing. It must also be transparent and provide people with appropriate information as mentioned earlier in this report.

Children's Code

3.19 The Children's code contains 15 standards of age appropriate design which seek to ensure and safeguard the privacy of children online. The code applies to "information society services likely to be accessed by children" in the UK. This includes, but is not limited to, apps, programs and many websites including search engines, social media platforms, online messaging or internet based voice telephone services etc. Electronic services for controlling connected toys and other connected devices are also information society services. During the period covered by this report the international discussion around children's online privacy was growing and data protection authorities in other jurisdictions considered, or implemented, similar frameworks of their own. The Sandbox has provided valuable, timely support to projects seeking to build these new regulatory requirements into product and service designs.

How can innovators address the 'age appropriate application' standard of the Children's code?

3.20 The Sandbox helped FlyingBinary consider the 'age appropriate application' standard of the code. FlyingBinary had identified the expected age range of lookafterme's users and designed different user journeys for each cohort. Instead of relying on the physical age of the user to place them into respective user journeys, FlyingBinary intended to use the assessed development age of the user, which was provided by their clinician. We considered that this would provide FlyingBinary with a good foundation to design lookafterme in a way that considered the different capacities, skills and behaviours of children at different stages of development. It also served as a good opportunity for FlyingBinary to pitch transparency messaging at an appropriate level for the user. Certain risks were highlighted by the Sandbox that FlyingBinary would need to mitigate, such as the possibility that the clinician may provide an inaccurate assessed development age.

3.21 Whilst not all organisations will have access to a clinician's assessed development age, it shows that organisations should take into account other factors where they are available as age ranges may not be a perfect guide to the interests, and evolving capacity of each child and their developmental needs. Opting to use the clinical developmental age which may more accurately reflect the needs of the child is a good example of implementing a risk based approach. It enabled FlyingBinary to choose the appropriate design and transparency information and reduce the likelihood that a child is presented with information that they do not understand. Any such approaches to processing must also be compliant with data minimisation requirements.

How can innovators address the 'transparency' standard of the Children's code?

3.22 The Sandbox worked with Seers to apply the transparency standard of the Children's code to its initial design for its CPCMP. Analysing Seers' proposed wording for the provision of relevant information to end users, before seeking their consent to process their personal data, revealed some good practice examples. For example, we noted that the use of images and layered banners, as part of the messaging within the CPCMP, aligned with the principles of prominence and accessibility outlined within the transparency standard. We also noted that Seers could improve their provision of the information by ensuring these methods feature as early as possible in user interactions as a safeguard to help ensure important information is not hidden behind subsequent interactions which the user may not read.

3.23 The ICO's guidance in relation to the transparency standard states that children should be prompted to talk to an adult before they activate any new use of their personal data ie changes to a privacy setting, and to not proceed if they are uncertain. The ICO suggested that Seers create more friction between the child using the platform and the parent or guardian granting consent. This would provide an opportunity for Seers to implement additional steps to help ensure that a child is not pretending to be their parent. For example, the platform could then request additional information (such as the parent's age and year of birth) to help corroborate the consent that is being provided.

How can innovators meet the data sharing standard of the Children's code?

3.24 Standard nine of the Children's code relates to data sharing. It states that organisations must not disclose children's data unless they can demonstrate a compelling reason to do so, taking into account the best interests of the child. In the early stages of its participation, FlyingBinary intended for all of lookafterme's data sharing functions to be 'off by default' with granular activation controls provided to the user. However, in 'acute' patient cases, data sharing with the user's clinician would be mandatory. Following constructive engagement with the Sandbox, FlyingBinary redesigned its approach so that all data sharing options would be 'off by default' as it intended to pursue a 'high privacy by default' approach. This demonstrates that the bar for 'a compelling reason' to disclose children's data is high. Compelling reasons include data sharing for safeguarding purposes, preventing child sexual exploitation and abuse online, or for the purposes of preventing or detecting crimes against children such as online grooming.

Anonymisation and Pseudonymisation

3.25 To coincide with the ICO's production of anonymisation guidance, the Sandbox looked at anonymisation and pseudonymisation as a key area of focus in 2022. Effective anonymisation can be useful for innovators as it enables them to use or share data in a non-identifiable way, safeguards individuals' privacy and is a practical example of 'data protection by design'. However, innovators need to be confident that the data is actually anonymous, otherwise it could lead to an inappropriate disclosure of personal data, eg through 're-identification'. A key challenge that some of our participants faced was determining the status of information that they processed.

How can innovators assess that they have effectively anonymised the information they process?

3.26 The Sandbox worked with Smart Data Foundry to consider the various tools available to them which could be used to assess and verify whether data had been anonymised within their data facility. Smart Data Foundry used identifiability assessments as a qualitative anonymisation tool to help identify the risk of re-identifiability for datasets held in their research data facility. The Sandbox advised Smart Data Foundry to develop their assessment of the three key indicators of identifiability from Chapter 2 of the ICO's draft anonymisation guidance: singling out, linkability, and inferences. In particular, we learned that innovators need to take a broader view of 'singling out'; an individual may still be singled out even if the data does not identify that individual by name. If a person's information can be individuated from others in a dataset, it is still singling out. Smart Data Foundry need to make sure that their assessments consider the richness of the data and how potentially identifying different categories are.

3.27 We also were able to apply an example of how to assess for inferences that can be drawn from the information held in Smart Data Foundry's research facility. As Smart Data Foundry was processing financial information, they were recommended to pay particular attention to how they deal with outliers. For example, high earners or individuals with large debts could be outliers in a dataset and there could be a risk that such an outlier could be connected to a high profile bankruptcy or wealth in the media. This shows that innovators will need to consider the industry context in which they operate when considering how inferences might be drawn about the data they are trying to anonymise.

3.28 Anonymisation can be a useful tool for sharing information while respecting individual privacy. However, information may be anonymous to the sharing organisation and not to the other organisation receiving the information. Innovators should be aware that both organisations will need to assess the risk of identifiability when processing information that has been anonymised. We advised Smart Data Foundry that they will need to consider in the conclusion of their identifiability assessments whether the data is identifiable both in the hands of Smart Data Foundry and also in the hands of the researcher seeking to use the outputs of the research database. It may be appropriate, and indeed more practical, for both organisations to undertake this assessment jointly.

How can innovators identify pseudonymised personal information?

3.29 Pseudonymisation is distinct from anonymisation; it is a technique that replaces or removes information in a data set that identifies an individual, so that the data is no longer attributable to a specific data subject without the use of additional information. Pseudonymised data remains personal data and is subject to UK data protection law.

3.30 Our work with Zamna highlighted the challenges that innovators may face when trying to identify whether data has been anonymised or pseudonymised. Zamna hosts cryptographic signals in its cloud and does not have access to the named passenger information that the signals represent. Our work through this issue showed that a key point of differentiation between anonymisation and pseudonymisation is the existence of additional information that can be used to identify an individual. Through the airlines in Zamna's network, Zamna can receive the additional information needed to retrieve the signals related to a passenger from their Zamna Cloud. The passenger's passport data acts as additional information, stored separately on the airlines' systems, which when processed by the Zamna Client App generates the identifier to locate signals for that identifier within the Zamna Cloud. Airlines are then able to reidentify the pseudonymised personal information to a named passenger on their own systems. This processing therefore qualifies as pseudonymisation under Article 4(5) of the UK GDPR.

3.31 Innovators should still consider these technical and organisational measures as privacy enhancing techniques that provide an enhanced level of confidentiality and security to personal information, also representing a good example of data protection by design and default. However, these measures described above do not make data anonymous because an individual can be identified with the use of additional information under pseudonymisation.

Data Minimisation

3.32 The Sandbox helped various participants consider how they will apply the data minimisation principle to their proposed processing of personal data. The UK GDPR's data minimisation requirements are set out in Article 5(1)(c). It states that "*Personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed* ('data minimisation')." The UK GDPR does not define the terms adequate, relevant and limited to what is necessary. Therefore, the controller will need to assess them on a case-by-case basis, considering the specified purpose for collecting and using the personal data. The assessment may also differ depending on how one individual's personal data is used compared to another.

How can data mapping help innovators to apply the data minimisation principle?

3.33 The Sandbox and TVP collaborated to explore how TVT could fulfil the objectives of the SVD, while ensuring data minimisation requirements were met. A crucial aspect of this engagement, therefore, involved assessing the necessary amount of personal data essential for achieving its purposes. In complex data sharing projects, it is common for controllers to initially consider processing all personal data before scaling down. However, it is important to remember that controllers must not collect personal data on the off chance that it might be useful in the future. As demonstrated in the TVP participation, a good starting place is for controllers to undertake a data mapping exercise which illustrates personal data flows, the personal data lifecycle and the type of personal data expected to be included. Controllers should also periodically review their data maps, as well as the personal data they process, to ensure their ongoing relevance with their processing purposes.

3.34 FlyingBinary's Sandbox participation further underlined the benefits of data mapping in order to comprehensively understand purposes for processing and determine the personal information that is necessary to achieve them. In the data mapping process, FlyingBinary itemised the individual items of personal data they would be processing, recorded the source of the data and mapped the items of personal data to its specified purposes. Data mapping also highlighted the importance of defining the purposes for processing in a more detailed and granular way so that personal data could be mapped against them. The map provided a solid foundation on which FlyingBinary could assess its data minimisation requirements. In

addition, the Sandbox helped FlyingBinary to use this approach to identify and consider additional information that it will process, such as the URL and web content that the lookafterme application will scan, which may constitute personal data.

What factors do innovators need to assess when applying the data minimisation principle?

3.35 The Sandbox's work with the financial institutions helped them to assess their data minimisation obligations against two different proposed datasets which would dictate what personal data the participating banks would process during the information sharing pilot. In particular, the work explored the cross-section between the 'adequate' and 'limited to what is necessary' elements of the principle. The Sandbox informed the participants that they must process the least amount of personal data possible in order achieve their stated purposes, yet it must also be sufficient to properly achieve those purposes. These are two of the data minimisation factors against which innovators should assess the personal data they intend to collect in order to apply the data minimisation principle. This involves finding an appropriate balance between these requirements, whilst also ensuring that personal data is relevant to the purpose for processing. This work highlighted the importance of having a strong grasp of the purposes for processing personal data from the outset and continuing to use them to assess the appropriate amount of personal data to process in line with the UK GDPR's data minimisation requirements.

4. Impact of the Sandbox

4.1 By engaging with the participants listed earlier in this report, the Sandbox has made a positive impact for a range of different stakeholders. This includes the wider public, the participants that enter the Sandbox and other external organisations, as well as the ICO itself. The section below details some of those positive impacts.

Impact of the Sandbox for organisations and the public

Delivering a public benefit, and safeguarding and empowering people

4.2 The Sandbox's work with various projects has helped to further demonstrate the ICO's strategic enduring objective to safeguard and empower people as part of its ICO25 strategic plan. A key entry requirement for potential Sandbox participants is that they can demonstrate a breadth or depth of public benefit within their proposals. As a result, Sandbox participants' processing of personal data will often result in benefits for different groups in society. For example:

- The financial institutions' (facilitated by the Home Office and UK Finance) work in the Sandbox seeks to reduce the impact of financial crime on the UK economy. It also, in turn, seeks to **reduce the degree of harm suffered by members of the public by improving the prevention and detection of financial crime.**
- Our work with the Gambling Commission and BGC helped gambling operators to **identify and support people experiencing harm from gambling**. Our work with both organisations stemmed from a recommendation from the House of Lords Gambling Select Committee in 2019, which recommended that the ICO collaborate with the Gambling Commission and BGC to remove perceived data protection barriers for gambling operators when sharing personal data for the purpose of protecting customers from gambling-related harms. Prior to the SCV, it was evident that the gambling industry's efforts to reduce the impact of harm did not go far enough. The Sandbox's work on supporting the development of the SCV demonstrates our commitment to helping people in the most need of support.

- Furthermore, our collaboration with TVP demonstrated the importance of identifying and mitigating data protection risks. The Sandbox identified additional risks around data retention and mis-matching of data, and assisted TVP in reducing the volume of personal data to be processed in its TVT project. In addition, the Sandbox provided operational guidance, namely the inclusion of a risk matrix within its Data Protection Impact Assessment (DPIA), to facilitate better risk identification and ensure consistency in risk assessment. As is the case with a number of engagements, the Sandbox often provides a review of a participant's DPIA, which ultimately serves to **safeguard the public from potential harm(s)**.

- The ICO's commitment to maximising the impact of the Sandbox is also demonstrated by our collaboration with the Home Office to update its passages on data protection within the Serious Violence Duty Guidance. This collaboration aimed to enhance the guidance provided to specified authorities that fall under the Duty, helping them to fulfil their obligations to the public effectively, and adhere to data protection law.

Empowering responsible innovation and sustainable economic growth

4.3 Another of the ICO's strategic enduring objectives within ICO25 is to empower responsible innovation and sustainable economic growth. The Sandbox is an important mechanism to help the ICO achieve this wider objective. Part of this objective is to provide regulatory certainty about what the law requires, reducing the cost of compliance and clarifying what the ICO will do if things go wrong. Each Sandbox participant has directly benefitted by receiving support from the ICO to implement 'data protection by design and default' into the development of their product or service. As the Sandbox is a free, professional service which welcomes participants from all sectors and sizes of organisation, this means that it contributes to various areas of the economy. It also helps individual companies and organisations grow and develop. Those benefits should be felt by different members of society as and when products or services are brought to market. Additionally, successful innovation drives economic growth.

4.4 Empowering responsible innovation and sustainable economic growth includes the ICO providing advice and guidance that can be relied upon to provide regulatory certainty for a wide range of external stakeholders. Sandbox projects operate at the

cutting edge of innovation and help the ICO to **understand where gaps in guidance exist, strength test existing guidance and contribute to new guidance**. For example:

- The Sandbox's work with Yoti helped the ICO reassess and refine its guidance on biometric data. Specifically, it helped the ICO to amend its definition of when biometric data constitutes special category data. This provides greater clarity on the ICO's position in our external facing guidance. In turn, it produces a greater degree of regulatory certainty for organisations processing, or intending to process, biometric data.

- FlyingBinary's Sandbox project helped the ICO further develop guidance aimed at keeping children safe online. FlyingBinary's input contributed to the ICO's development of the best interests of the child framework during its early design. That feedback has helped to influence the final product that is available on the ICO's external website. The framework helps organisations consider the best interests of children (established by the United Nations Convention on the Rights of the Child), alongside data protection requirements and the standards of the Children's code, when they design online services. It helps to provide regulatory certainty for organisations operating in this space, whilst aiming to help ensure the safety of children online.

- As detailed in section three of this report, the Sandbox team has assisted various participants in understanding their data protection roles and responsibilities. This particularly relates to contexts where joint controllership relationships might be established. The exit reports for these projects combine to provide a body of information that may be useful to organisations assessing data protection roles and responsibilities within complex data sharing ecosystems. As a result, organisations may gain a better understanding of when joint controllership relationships are established, and what they need to do when they are.

- Zamna and Smart Data Foundry's Sandbox projects helped to inform our consultation and draft guidance on anonymisation, pseudonymisation and privacy enhancing technologies. Working with these real use cases contributed to the ICO's deepening knowledge of this area during the consultation period. The outcomes, such as those in relation to whether information has been effectively anonymised and the difference between anonymisation and pseudonymisation, will likely be reflected in the ICO's finalised guidance. It may also include a case study that has been discussed with Our

Future Health following the completion of that project. That guidance is expected to be published, and available for organisations to use, later in 2024.

4.5 Data protection law is an enabler for fair and proportionate data sharing, rather than a blocker. Responsible data sharing plays a pivotal role in empowering innovation and fostering economic growth across various sectors. Whether that be Zamna's Verified Passport product which helps airlines to securely verify passenger information for travel, or an SCV which supports people experiencing gambling-related harms, data sharing enables organisations to unlock valuable insights, develop transformative solutions to real needs and reduce inefficiencies.

4.6 The Sandbox engages with projects that operate in new, emerging areas of data protection. As a result, it plays a key role in supporting organisations to innovate responsibly in evolving spaces. For example, the Sandbox helped Yoti and FlyingBinary to design services falling within the scope of the Children's code. The ICO provided both organisations with crucial advice on how to operationalise new requirements to protect the rights and privacy of children online.

Sharing our knowledge within the public sector at home and abroad

4.7 As the first dedicated GDPR-territory data protection Sandbox, the ICO's Sandbox is often seen as a leader in helping to promote responsible innovation using personal data. The Sandbox team has regularly accepted requests to meet with external organisations that are considering operationalising their own Sandboxes. Those organisations include other data protection authorities, UK government departments and public bodies in foreign countries. The proposed Sandboxes include both those with and without a data protection focus. These engagements have provided the Sandbox team with an opportunity to discuss our key learnings and experiences with organisations seeking to develop similar interventions from the ground up. Those learnings and experiences relate to both data protection themes, and the operational challenges and processes of running a Sandbox. The engagements have also established positive relationships, underlined the Sandbox team's collaborative values and have been well received externally.

Impact of the Sandbox for the ICO

Supporting ICO knowledge

4.8 The Sandbox continues to engage with projects at the cutting edge of data protection. This has helped to develop knowledge of new innovations and their fundamental data protection challenges. As a result, the Sandbox has often been used as a source of internal knowledge which can be leveraged by cross-office colleagues when they are working on projects such as the development of new guidance or opinions. That knowledge has also been used to communicate internal learnings which can be used by ICO staff when engaging with their respective stakeholders, contributing to a consistent approach. For example:

- The Sandbox's work with FlyingBinary helped the ICO to understand more about the challenges faced by industry where they are caught by the scope of the Children's code, which aims to protect the privacy of children online.
- The Sandbox's work with Zamna provided an opportunity to look at a practical example of complex cloud processing systems and examples of organisational considerations which provided feedback to the ICO's ongoing work on cloud computing. In addition, our work with Smart Data Foundry has provided a practical example of the use of synthetic data that has fed into the development of our guidance on privacy enhancing technologies after the publication of the draft.
- The Good With project helped the ICO develop knowledge in relation to the processing of unsolicited special category data.

4.9 Our work with Our Future Health, Zamna and the Financial Institutions has also allowed us to identify common themes in the challenges that innovators face in relation to identifying potential joint controllership relationships. We shared these themes with the cross office policy project that is currently reviewing the ICO policy on joint controllership in complex processing ecosystems.

Improving our internal processes

4.10 As the Sandbox continues to evolve, it regularly refines and amends its operational processes and procedures, and its governance structures. For example, the Sandbox team has recently revised its governance group and method of oversight. This helps to ensure the Sandbox operates effectively and efficiently to maintain its external impacts. It also produces positive impacts for the ICO. Engaging key decision makers and subject matter experts at the correct time helps to mitigate the risk of any high profile issues that may occur during Sandbox projects. Additionally, this acts as an effective forum for discussion that provides assurance for a consistency of approach across the Sandbox and the ICO.

Maintaining our reputation as a forward thinking regulator

4.11 The Sandbox contributes to the ICO's reputation as a trusted information rights and digital regulator. For example, the previous section details how the Sandbox contributes to some of the ICO's commitments to the public in its ICO25 strategic plan. Having implemented the first GDPR territory Sandbox, the ICO is seen as innovative and leading the way in its approach to upstream regulation. Its continued success underlines that the ICO provides interventions which are both useful to industry, prioritise information rights and pursue the long term outcome of empowering people through information.

Feedback from our participants

4.12 During the lifecycle of a Sandbox project, and following its conclusion, the Sandbox team seeks feedback from its participants. The purpose of this feedback is to help measure the impact of the Sandbox and to identify any areas for improvement. The results of the feedback from the 14 participants identified within this report have been overwhelmingly positive.

4.13 For example, during the period when the 14 participants were applying to enter the Sandbox their feedback scored the Sandbox team at an average of 9.1 out of 10 across the criteria that were measured. During the phase in which the ICO and the participants collaborated to plan the scope of the Sandbox projects feedback was received from 13 of the participants.

Within this phase, the average score given to the Sandbox team grew to 9.7 out of 10. This means that the Sandbox team consistently received very strong feedback in relation to important criteria including:

- responding to application queries quickly and efficiently;
- participants indicating that they would recommend the Sandbox to another organisation;
- Sandbox team members providing useful, clear and easy to understand input during the planning phase;
- Sandbox team members having a good understanding of the participant's product or service; and
- the participant receiving adequate support from the Sandbox team to effectively develop the participant's Sandbox plan.

4.14 Following the completion of their projects, the eight participants that completed the ICO's recently established innovation service survey provided important feedback. Each participant stated that it either agreed or strongly agreed that its organisation has a better understanding of what it needs to do for its innovation to be compliant with data protection law. They also either agreed or strongly agreed that they are more confident to develop innovations in a privacy compliant way. A key aim of the Sandbox is to help participants embed data protection by design and default and this feedback is a strong indicator that the Sandbox is having a positive impact on the data protection practices of its participants.

4.15 The same participants also provided further feedback on the innovation service survey which clearly demonstrates the positive impacts of engaging with the Sandbox. For example, nearly all of the participants indicated that engaging with the Sandbox, which is a free service, resulted in significant cost savings. Cost savings included overheads related to having to pay for legal expertise or consultancy services, realising operational efficiencies, time and resource savings, and facilitating the onboarding of partners. Two participants estimated these savings to be between £100,000 and £499,999. Another two estimated their savings to be between £20,000 and £49,999. This feedback demonstrates the tangible support the Sandbox provides to organisations seeking to bring products or services to market. It also further underlines that the Sandbox has had a significant impact in helping the ICO to empower responsible innovation and sustainable economic growth.

5. What's Next in the Sandbox?

5.1 In keeping with the Sandbox's mission to focus on the cutting edge of innovation, our areas of focus remain emerging technologies, exceptional innovations and biometrics. Moving forward, we will continue to align our areas of focus with the latest edition of our Tech Horizons report to support the wider ICO with practical use cases related to emerging technological trends and other exceptional innovations that we have identified as in need of data protection consideration. The Sandbox's current projects are detailed on our website.

5.2 If your organisation is developing an innovative product or service that involves the processing of personal data, we strongly encourage you to submit an expression of interest to us. Please see our webpage on how to apply for the Sandbox, and our current areas of focus, if you are interested in participating.

5.3 The Sandbox remains focussed on supporting participants who are navigating challenging data protection questions or themes. Our aim is to help innovators create products and services which utilise personal data in innovative and safe ways.