THE ROYAL SOCIETY

Centre for Data Ethics and Innovation

Privacy Enhancing Technologies: Market Readiness, Enabling and Limiting Factors in the UK public sector

Executive summary

This report summarises market research undertaken by the Open Data Institute[1] and London Economics[2] on behalf of the Royal Society in partnership with the Alan Turing Institute, and with the support of the Centre for Data Ethics and Innovation (CDEI). The objective of this research was to determine the state of PETs adoption in the UK public sector in early 2022, inhibiting factors and potential evidence-based opportunities to facilitate increased consideration of their adoption.

In conducting this research a number of key observations have been gathered around the level of adoption of PETs in the UK public sector. First and foremost is the observation that the majority of those spoken to from within the UK public sector for this market research are in the early stages of raising their personal and institutional awareness of these technologies. There are, however, some institutions within the sector that are forging ahead in taking PETs from ideation through to the development of proofs of concept and implementation, which holds promise for increasing the awareness and potential adoption of these technologies elsewhere within the UK public sector. This includes the use of synthetic data and proofs of concept of federated learning and federated analytics by the Office for National Statistics, as well as research into similar technologies at The National Archives.

The current, relatively low uptake of what may be considered "emerging PETs"[3] in the UK public sector can be partially attributed to the existing confidence and familiarity with existing privacy preserving techniques such as anonymisation, pseudonymisation, encryption and data minimisation, which respondents noted that there was general satisfaction with for the purposes of the work they are presently undertaking. Consequently, awareness and consideration of emerging PETs, such as federated learning, federated analytics and secure multi-party computation, has not extended far beyond familiarity with the terms and the underlying theory at present.

Furthermore, existing data governance concerns such as data quality within and across institutions and the cultural dimension of managing data remain fundamental concerns within the UK public sector, which are being heavily prioritised. This observation suggests that, institutionally, the exploration and usage of novel technologies such as PETs are not presently viewed as high priorities within the public sector.

[1] See https://theodi.org/about-the-odi/ (accessed 09 May 2022)
[2] See https://londoneconomics.co.uk/about-us/ (accessed 09 May 2022)
[3] "Emerging PETs" for the purpose of this study is used to differentiate between technologies such as federated learning, federated analytics and secure multi-party computation, which are at an earlier stage of adoption, as opposed to "traditional PETs" such as anonymisation, pseudonymisation and minimisation that have had greater adoption.

High barriers to deployment of PETs are also contributing to this current lack of prioritisation. These include factors such as the limited pool of technical expertise available to develop and communicate the value of PETs to senior decision makers within the public sector, uncertainty around the impacts the adoption of PETs might have on existing data governance practices and the absence of an overwhelming case for a return on investment to justify investment of limited financial resources.

Interviewees spoken to from within the UK public sector for this research did, however, express personal interest to explore the utilisation of newer privacy enhancing technologies and techniques, although at present there is a need for further evidence and examples of their applications in practice. This is necessary to develop the body of evidence that senior decision makers who have authority to commission applied PETs require in order to justify the exploration and potential adoption of what remain emerging technologies. This is based on concerns that remain amongst those interviewed within the sector on the underlying data infrastructure which must be addressed in order to foster an enabling environment for the adoption of PETs.

Contents

Project background

PETs are a growing area of interest for governments, companies, academia and other organisations such as data institutions[4] due to the potential they hold for unlocking greater value from sensitive data that both organisations and individuals may otherwise have been reluctant to share. This has accelerated in recent years, and PETs were identified within the UK's 2021 National Data Strategy[5] as an area of particular interest due to their abilities to facilitate the building of greater trust through providing additional layers of privacy.

This piece of research comes three years after the publication of the Royal Society's report: Protecting privacy in practice[6], which provided an overview of the uses and limitations of PETs at the time, as well as their potential for future applications. The Royal Society is now in the process of undertaking a refresh of this report and thus commissioned this market research to help inform the Royal Society's understanding of the current state of PETs usage and identify key user needs with regard to PETs in the UK public sector.

The research was initially intended to focus on the market readiness of a subset of PETs - specifically secure multi-party computation, federated learning and federated analytics - however this was broadened out to include other types of PETs, based on varying stages of consideration around those that are considered as more traditional and established, and those that are considered as emerging. Attention throughout the course of this research was placed not solely on the privacy enhancing benefits of PETs, but also upon the opportunities that these may be able to provide to organisations in extrapolating greater innovation and collaboration through greater data sharing.

Secure multi-party computation, federated learning and federated analytics

Since Yao's seminal work in 1982[7], 30 years of research on secure multi-party computation (SMC) has been conducted, proceeding from purely theoretical research into real-world applications. The first large-scale and practical application of SMC was an electronic double auction in 2008 designed to establish the market-clearing price for sugar beets in Denmark[8]. Playing the role of auctioneer in this scenario was an SMC scheme involving representatives of Denmark's only sugar beets processor (Danisco), the sugar beet growers' association (DKS), and a third group responsible for devising and implementing the system. Everyone wanted to know the outcome — the price where total supply equalled total demand — but the farmers wanted their bids to remain confidential.

---

[4] Hardinges, J. and Keller, J. 2021 What are data institutions and why are they important? (see https://theodi.org/article/what-are-data-institutions-and-why-are-they-important/, accessed 23 May 2022)
[5] Department for Digital, Culture, Media and Sport 2021 National Data Strategy (see https://www.gov.uk/government/publications/uk-national-data-strategy/national-data-strategy, accessed 09 May 2022)
[6] The Royal Society 2019 Protecting privacy in practice: The current use, development and limits of Privacy Enhancing Technologies in data analysis (see https://royalsociety.org/-/media/policy/projects/privacy-enhancing-technologies/privacy-enhancing-technologies-report.pdf?la=en-GB&hash=862C5DE7C8421CD36C105CAE8F812BD0, accessed 09 May 2022)
[7] Yao, A.C. 1982 Protocols for Secure Computations (see https://research.cs.wisc.edu/areas/sec/yao1982-ocr.pdf, accessed 23 May 2022)
[8] Gomi, K. 2021 Multi-Party Computation: Private Inputs, Public Outputs (see https://www.forbes.com/sites/forbestechcouncil/2021/10/26/multi-party-computation-private-inputs-public-outputs/?sh=1926e311bb0c, accessed 23 May 2022)

Practical applications of SMC include privacy-preserving machine learning, private set intersection, and secure genomic sequence comparison[9]. A group of Estonian researchers has developed a VAT fraud detection system prototype for the Estonian Tax and Customs Board that uses SMC to remove the companies' concerns over confidentiality with the declaration of purchase and sales invoices for automated risk analysis and fraud detection[10].

Some of the limitations of SMC highlighted by the literature include the lack of an easy-to-use secure sorting method to implement, of a privacy-preserving method to link records when the inputs are completely different datasets with different structures, and of best practices in delivering and administration[11]. Talviste[12] also poses some interesting legal questions for public organisations: when sensitive data that is secret gets shared into random pieces and distributed among many parties, is it still considered sensitive data from the legal point of view? Can secret sharing be considered a form of encryption and if so then where is the key? It is important that public entities answer these questions before deploying secure multi-party computation applications.

The general description of Federated Learning (FL) has been given by McMahan & Ramage[13], and its theory has been explored in Konečný et al.[14], McMahan et al.[15] and McMahan et al.[16]. Training statistical and machine learning models in heterogeneous and potentially massive networks introduces novel challenges that require a fundamental departure from standard approaches for large-scale machine learning, distributed optimization, and privacy-preserving data analysis. As summarised by Li et al.[17], some of the key challenges related to solving the distributed optimisation problem posed by FL are (1) overcoming expensive communication (and privacy concerns over sending raw data) developing communication-efficient methods that iteratively send small messages or model updates as part of the training process; (2) heterogeneous hardware and systems across the federated network; (3) the statistical heterogeneity of the data, as devices frequently generate and collect data in a non-identically distributed manner across the network; (4) addressing privacy concern related to transferring information from devices, without losing too much model performance or system efficiency, e.g. through SMC or differential privacy.

In May 2020 the Google AI team published an article on leveraging computing mechanisms of the distributed learning model training infrastructure to facilitate data analytics, namely Federated Analytics (FA)[18]. FA allows scientists to derive analytical insights of distributed

[9] Zhao, C. et al. 2019 Secure multi-party computation: Theory, practice and applications (see https://www.sciencedirect.com/science/article/abs/pii/S0020025518308338, accessed 23 May 2022)
[10] Bogdanov, D. et al. 2015 How the Estonian Tax and Customs Board Evaluated a Tax Fraud Detection System Based on Secure Multi-party Computation (see https://link.springer.com/chapter/10.1007/978-3-662-47854-7_14, accessed 23 May 2022)
[11] Talviste, R. 2016 Applying Secure Multi-party Computation in Practice (see https://core.ac.uk/download/pdf/144708931.pdf, accessed 23 May 2022)
[12] Ibid.
[13] McMahan, H.B. and Ramage, D. 2017 Federated learning: Collaborative machine learning without centralized training data (see https://ai.googleblog.com/2017/04/federated-learning-collaborative.html, accessed 23 May 2022)
[14] Konečný, J. et al. 2016 Federated Learning: Strategies for Improving Communication Efficiency (see https://arxiv.org/abs/1610.05492, accessed 23 May 2022)
[15] McMahan, H.B. et al. 2017 Communication-Efficient Learning of Deep Networks from Decentralized Data (see http://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf, accessed 23 May 2022)
[16] McMahan, H.B. et al. 2018 A General Approach to Adding Differential Privacy to Iterative Training Procedures (see https://arxiv.org/abs/1812.06210, accessed 23 May 2022)
[17] Li, T. et al. 2019 Federated Learning: Challenges, Methods, and Future Directions (see https://arxiv.org/pdf/1908.07873.pdf, accessed 23 May 2022)
[18] Ramage, D. and Mazzocchi, S. 2020 Federated analytics: Collaborative data science without data

datasets without the need of moving data to a central computing entity. FA allows to not only consider the distributed model training processes for improving model accuracy but also to exploit such collaboration architecture to evaluate the quality of the trained model at the user-level perspectives, i.e., the model performance at the user's end. Hence, without the learning part, it is possible to reuse the computing scheme of the learning architecture to perform statistical analysis on local data that may lead to building better products[19].

FA differs from the recent federated learning paradigm in the sense that federated learning emphasises collaborative model training, whereas federated analytics emphasises drawing conclusions from data. FA is susceptible to the same challenges highlighted above for FL, including architecture, heterogeneous edge devices, privacy and security management including peer management, raw data, and intermediate data protection, resource management (including computing power, communication, energy, and monetary cost), and analytics design and optimisation[20].

Approach

This short piece of market research took place in three phases: project inception, in which desk research was conducted to inform the development of semi structured interviews; primary research, in which interviews were conducted with senior officials from within the UK public sector and a backup survey was developed as a potential redundancy measure; and finally, the synthesis and reporting phase, in which the research findings were written up. A fuller account of the methodological approach to the research can be found in Appendix 2.

Together with the Royal Society, respondents from four types of UK institutions were identified from which to gather data to inform this report. These four types of institutions included: those that use data to generate insights; those who act as data gatekeepers; those that are developing infrastructure for data sharing and those publishing open data. Efforts were also made to canvas a cross-section of respondents from local and central government, those from collaborations between government and academia, as well as respondents from regulatory bodies.

The series of interviews explored three key areas as the roots from which the conversations developed and were tailored, depending upon the answers to these questions and - most significantly - the degree to which the interviewees were familiar with various PETs and the stage at which they and their institutions were at in considering their potential deployment. The three key areas were as follows:

- Their experiences of the UK public sector with PETs
- The concerns and hindering factors surrounding their adoption
- The underlying data access, infrastructure and governance issues in the UK public sector

Collection (see https://ai.googleblog.com/2020/05/federated-analytics-collaborative-data.html, accessed 23 May 2022)

[19] Pandey, S.R. et al. 2021 Edge-assisted Democratized Learning Towards Federated Analytics (see https://arxiv.org/pdf/2012.00425.pdf, accessed 23 May 2022)

[20] Wang. D. et al. 2021 Federated Analytics: Opportunities and Challenges (see https://www4.comp.polyu.edu.hk/~csdwang/Publication/FA-vision21.pdf, accessed 23 May 2022)

Interviewed institutions

The table below includes a list of the institutions from which interviewees were spoken to for this research. This includes a brief description of the institution and a summary of their data interest, contextualising the relevancy of PETs to their work.

Interviewed institution	**Brief description and data interest**
Competition and Markets Authority (CMA)	The CMA is the competition regulator in the United Kingdom. The CMA has responsibility for strengthening business competition and preventing and reducing anti-competitive activities. The Digital Markets Unit (DMU) was established within the CMA in 2021 to operationalise the future pro-competition regulatory regime for digital markets.
DataLoch	DataLoch is a collaboration between the University of Edinburgh and NHS Midlothian. This service provides access to de-identified data for research, service management purposes and innovation in a trusted research environment.
Department for Transport (DfT)	The DfT supports the UKs transport network through working with agencies and partners. Part of this involves creating an enabling environment to facilitate app developers access open data about services, including timetables and location data for the purpose of reducing complexity around planning journeys on public transport.
Government Digital Service (GDS)	GDS facilitates the development and delivery of products that enable personal data processing, including the Government Data Exchange, which will look to create a pan-government digital identity solution.
Greater London Authority (GLA)	The GLA is a strategic regional authority with powers over transport, policing, economic development and fire and emergency planning. The GLA is also home to the London Datastore, which is a data orchestration platform of over 700 datasets that helps to identify who holds what data where for the purpose of aiding understanding and the development of solutions to London's problems.
The National Archives	The National Archives is the official archive and publisher for the UK government, and for England and Wales. This includes responsibility for the management of the digital archive, which involves the long-term preservation of records and providing access to the public.

Office for National Statistics (ONS)	The ONS is the UK's largest independent producer of official statistics and the recognised national statistical institute of the UK. The ONS collects and publishes statistics at the national and local level on the economy, population and society. The ONS is also responsible for conducting the census in England and Wales.

Summary findings

Through interviews with senior officials from within the UK public sector, the ODI sought to explore the stage at which various institutions were at in the consideration or implementation of various PETs. This section of the report contains a synthesis of the findings that emerged from the interviews carried out, which explored the three key areas mentioned within the preceding section, covering the research approach. This section begins with an overview of the present experiences of respondents with PETs, including some potential applications, inhibiting factors and the existing practices and technologies used. This discussion then expands out to some of the more contextual considerations that feed into the awareness and adoption of PETs in the UK public sector.

The experiences of the UK public sector with PETs

One of the primary objectives of this research was to gauge the extent to which emerging PETs such as federated learning, federated analytics and secure multi-party computation were currently being either considered or deployed in the public sector at present. This section will provide an overview based on the interviews that were carried out, in which it was apparent that much of the thinking around these emerging PETs remains at an early stage.

An important caveat that must be noted at this stage is that interviewees were uncertain as to the utility of emerging PETs for the work that they undertake. This is to acknowledge that certain individuals who were interviewed, such as those from the CMA, emphasised that while PETs have relevance within their work, this would not necessarily require them to adopt these technologies themselves. Specifically, in conversations with interviewees from the CMA it was noted that their interest lies primarily in the impact that these PETs might have on competition and consumer protection issues[21]. It is therefore important to appreciate the indirect implications that PETs are having on the ways in which those working within the public sector are viewing data governance.

Interviewees from DataLoch and the GLA had awareness of the range of PETs that exist, however they noted that they are primarily at the beginnings of the journey in terms of considering how they might be deployed. In addition, the interviewee from the DfT noted that, in their specific area of work within the department, very little sensitive data is handled and that which is, is open. As a result, for this area of work there has not been significant consideration of PETs, whereas it is possible that there has been elsewhere in the DfT.

[21] See https://www.theguardian.com/commentisfree/2020/jul/01/apple-google-contact-tracing-app-tech-giant-digital-rights (accessed 09 May 2022)

Several of the interviewees did, however, discuss specific examples in which they are either considering or developing emerging PETs. These include differential privacy, the use of synthetic data and federated analysis and learning. It is worth noting that secure multi-party computation was not presently being considered by any of the interviewees' institutions, as far as the interviewees were aware.

Within GDS, discussions are presently ongoing around differential privacy and role-based access[22]. When asked as to whether other types of PETs are being considered, the interviewee noted that federated learning is now coming up in data exchange discussions. While still at an early stage, there is an acknowledgement that the creation of a single central data repository for governments is infeasible due to security[23] and data management concerns and that there is currently ongoing thinking around how it might be possible to create a federated model that would allow access and analysis without actually sharing the underlying data. As part of these considerations, the respondent cited that there was an underlying need to understand what data exists where, what quality this data is and if it might be possible for some of these datasets to talk to each other.

The Synthetic Data and Privacy Enhancing Technologies team within the ONS is relatively advanced in the development and trialling of several PETs. This included examples of utilising synthetic data, as well as the development of proof of concepts for both federated learning and federated analytics.

Gaining access to ONS data assets through the Secure Research Service (SRS) - to be replaced by the Integrated Data Service (IDS) - is a time consuming and resource intensive process for both the applicant and the ONS[24]. This process can be lengthy, requiring applicants to secure several levels of accreditation, including as an individual, for their project and for their institute. It is for this reason that the ONS has been exploring the use of synthetic data, to develop indicative data assets which individuals can access while applying for the accreditations required to access the original data asset. The intention behind this is that accessing these synthetic data assets can allow the individual to get a sense of the data asset and whether it has the expected utility that they had hoped for, without providing access to the original data asset. Doing so can allow the individual to make this judgement and determine whether it is worth continuing the accreditation process to then use the original data asset.

Furthermore, the team at the ONS has been involved in research on PETs with other national statistics offices as part of the UN PETs Lab[25], which has included partaking in the development of a proof of concept using trade data. At present, the ONS is at the stage of trying to build knowledge amongst colleagues internally of what might be possible and not yet at the stage of applying federated learning or analytics. As part of the shift to the IDS - which will be the way in

[22] Role-based access control (RBAC) is a means by which to assign permissions access within an organisation by role, rather than on an individual basis. This can be used as a means by which to limit employees' access to data and information that is of relevance to their job.
[23] See https://arxiv.org/pdf/1902.01046.pdf (accessed 09 May 2022) and https://www.sciencedirect.com/science/article/abs/pii/S0167739X20329848 (accessed 09 May 2022)
[24] For more on the Secure Research Services and the process of obtaining approval to access data, see: https://www.ons.gov.uk/aboutus/whatwedo/statistics/requestingstatistics/secureresearchservice (accessed 23 May 2022)
[25] Information Age 2022 UN launches privacy lab pilot to unlock cross-border data sharing benefits (see https://www.information-age.com/un-launches-privacy-lab-pilot-to-unlock-cross-border-data-sharing-benefits-123498572/, accessed 09 May 2022)

future that individuals external to the ONS will access ONS data - the ONS is onboarding other departments' data assets and other, third-party, data assets, which is where they are presently looking at whether federated learning methods can perhaps be used instead of onboarding datasets as a way of producing statistics.

When in conversation with the interviewee from The National Archives, discussion was had on research that they are presently undertaking on 'non-consumptive access'[26]. This arose in conversation around take-down rules[27] that The National Archives implements when they find that information cannot continue to be published due to privacy and sensitivity concerns. It was noted that researchers will often be interested in conducting analysis on records within The National Archives' web archive, but this would run the risk of undoing these rules. Elaborating, the interviewee raised concerns around this occurring "particularly [through] computational analysis over collections of records that we have, where we are concerned about the consequences of another party wandering off with the data and being free to process it without constraints [...] And we've got a whole bunch of use cases where - not just in relation to privacy - [...] where we may need to provide forms of non-consumptive access to enable researchers to be able to do things with our records, without giving them a copy of the records".

At present, the solutions at hand are either providing some way of querying this information without undoing the takedown rules, or processing the whole web archive collection through the take-down rules to produce a version that has these rules applied. The second of these two solutions is practically infeasible, due to time and resource considerations[28]. As such, the interviewee noted that thinking has shifted within The National Archives towards the idea of alternatives that fall within the realm of 'non-consumptive access' in order to facilitate greater access to National Archive data without compromising data privacy.

Potential applications

Following discussion of their institutions' current stances towards PETs, respondents were asked to consider the ways in which PETs could facilitate their work, or play a role in increasing access to sensitive data.

From the responses of interviewees, it appears that it could be an interesting exercise to bring together those with practical experience of taking emerging PETs from ideation through to deployment with individuals from throughout the public sector who hold an interest in PETs. While there is growing awareness and understanding of how PETs function and their potential benefits, translating this to context-specific use cases for deployment remains a challenge.

Two applications that were put forward by respondents were the use of PETs to facilitate evidence gathering and the use of PETs for the purpose of auditing algorithms. It was

[26] The respondent used the term 'non-consumptive access' when speaking about concerns held by The National Archives of other parties processing National Archives data without constraints. As noted in the following quote, this also extends to the aim of enabling researchers to carry out analysis on National Archives records without providing them a copy of the original records. On this basis, 'non-consumptive access' was used as a term that encompasses a variety of privacy preserving methods and technologies, including PETs such as federated learning and analytics, as well as trusted research environments.

[27] Takedown rules are used to suppress material that has been gathered and held by The National Archives, but cannot be made publicly available. For more, see: https://www.nationalarchives.gov.uk/legal/takedown-and-reclosure-policy/#:~:text=As%20a%20general%20rule%2C%20information,discretion%20of%20The%20National%20Archives. (accessed 23 May 2022)

[28] See https://www.adalovelaceinstitute.org/blog/privacy-enhancing-technologies-not-always-our-friends/ (accessed 09 May 2022)

suggested that the use of emerging PETs, such as federated analytics might facilitate the monitoring of large tech platforms, as this process could potentially involve the collection of personal data.

Inhibiting factors to the adoption and consideration of emerging PETs

Numerous reasons were suggested as to why consideration of emerging PETs remains at a relatively early stage within much of the UK's public sector. This can largely be attributed to overarching obstacles such as a lack of technical expertise and experience of transferring the conceptual benefits of many PETs into practical applications. This echoes previous research[29], which suggested that the barriers to adoption of emerging PETs are often economic, as well as due to a lack of skills.

Need for further examples

Related to the obstacle of translating emerging PETs as concepts into practical applications, most interviewees raised the need for further case studies as central to greater consideration of PETs in future. Benefits of further examples include greater evidence to support the development of proposals for consideration, as well as providing assurance to those responsible for making the decision to explore the deployment of a particular PET that such applications are feasible, safe, secure and facilitate greater access to data. The interviewee from DataLoch raised this point specifically, noting that the data controllers who they work with - particularly given the intended expansion of the service - would likely need to see examples of PETs working in their case ahead of seriously considering their deployment. Further materials and resources, similar to the guidance produced by the Goldacre Review for using health data for research and analysis, could be one means by which this could be addressed.

There was strong appetite amongst interviewees for evidence of the utilisation of emerging PETs such as federated learning and federated analytics within their specific context, that could serve as a proof of concept. While perhaps not presently practically feasible, it is certainly worth exploring the possibility as to whether the proofs of concept could have utility on smaller scales in which their functionality could be demonstrated using open data - as occurred in the development of the ONS proof of concept.

A related obstacle that emerged was the need for those with decision-making responsibilities within organisations being sufficiently informed about both the risks and benefits associated with emerging PETs in order for them to have the confidence to make the case for their consideration. As one interviewee succinctly stated, it is a case of "how are new PETs better than the old ones?", which in the case of PETs can involve a time-consuming process of developing a highly technical cost-benefit analysis. One suggestion that arose through the course of the interviews was the external pressure that can be exerted by researchers who are keen to access public sector data. If this is concerted, it was suggested that this might be sufficient leverage to encourage more serious consideration of the value of adopting PETs.

Additional guidance and harmonisation across regulators

As an extension of this line of discussion, interviewees were asked what resources on emerging PETs they were familiar with and how useful these are. Most respondents were aware of the

[29] Acquisti, A. *et al* 2020 Secrets and Likes: The Drive for Privacy and the Difficulty of Achieving It in the Digital Age (see https://www.cmu.edu/dietrich/sds/docs/loewenstein/secretslikes.pdf, accessed 09 May 2022)

PETs adoption guidance from CDEI[30] and appreciated the efforts at collating the repository of use cases. Interviewees did note, however, that the utility of the repository did not extend as far as to help them in generating ideas for potential applications of PETs within their respective areas of focus. Interviewees found these examples too-far-removed from their own work to be of greater practical use and again stressed the need for more examples that could help demonstrate the value that the use of PETs would have over existing practices and technologies.

On the subject of guidance, several interviewees also cited a lack of clarity on PETs guidance from regulators, such as the Information Commissioner's Office (ICO), as a source of hesitancy and uncertainty. This is an area in which activity is currently taking place[31]. Interviewees from the CMA noted that there is ongoing work between the regulators, however this is taking place at various speeds given that these are rapidly emerging technologies that have broad potential implications. Furthermore, the technical complexity of PETs was cited as an additional complication that is currently proving to be an impediment to greater cooperation between regulators. As such, there are ongoing attempts being made to translate the highly technical language surrounding PETs into something that is meaningful across the various regulators.

Technical barriers to adoption

Interviewees were almost unanimous in raising the technical complexity of PETs as the most significant inhibiting factor to the broader consideration and adoption of emerging PETs. This was noted as a particular disincentive that was stymying serious consideration by data owners from investing their time in developing their understanding of these technologies. During conversation with the interviewee from the ONS, it was noted that communication of the risk that data owners are taking in order for them to feel more comfortable with exploring the possibility of using PETs is central to allaying this reservation.

This was echoed in conversations with other interviewees, in which concerns around “what is good enough?”, “what is de-identified enough?” and “what is safe enough?” were repeatedly raised when contemplating the use of PETs. These are some of the fundamental questions that data controllers are looking for greater clarity on before they seriously consider some of the further steps towards using these emerging technologies. While specific answers to these questions were not proposed, one respondent did note that there is the added difficulty when working with data that these answers may be circumstantial, depending on the project, institution or dataset. There is an interconnection between levels of trust in the technologies themselves and the corresponding trust in organisations’ data governance and whether the current infrastructure and practices are adequate to facilitate their adoption. Further, more specific examples could be useful in illustrating how to address these questions.

PET-specific data governance concerns

Some explicit concerns were raised by interviewees regarding the impact that the use of PETs by others may have on their own data practices. This highlights the need for greater understanding and awareness of some of the secondary effects of the usage of PETs.

[30] Centre for Data Ethics and Innovation 2021 Privacy Enhancing Technologies Adoption Guide (see https://cdeiuk.github.io/pets-adoption-guide/, accessed 09 May 2022)
[31] OpenMined 2022 Classifying the Challenges of Privacy Enhancing Technologies (PETs) in IOT Data Markets (see https://blog.openmined.org/classifying-the-challenges-of-privacy-enhancing-technologies-pets-in-iot-data-markets/, accessed 09 May 2022)

The respondent from The National Archives singled out the complications that may arise out of the use of synthetic data. Specifically, an example was considered, in which a synthetic data asset, or one which contains synthetic data, is selected for permanent preservation within the archives. Currently there is no archival practice to address this issue, however this has encouraged those at The National Archives to revisit practices and procedures around how they describe records, what information they record about records and how they retain intellectual control. It has been necessary to do this as, through encounters with other parts of the public sector in the UK, they have found instances in which synthetic data has been used in records that they have received. This illustrates that additional guidance may be required around the practice of labelling and describing data assets, in light of the increased use of PETs.

Discussion with the interviewee from the ONS provided a different perspective on the current concerns that surround the adoption of PETs within the public sector. As a potential solution provider, following the development of the federated proofs of concept, the respondent noted that if they were keen to collaborate on similar efforts with other teams or departments within the public sector, they would likely be faced with a substantial degree of scrutiny as to whether the PET has been implemented correctly, so that the safety guarantees that have been promised are definitely going to hold. Therefore, those within the Synthetic Data and Privacy Enhancing Technologies team at the ONS are now in the process of beginning to consider the types of assurance that they might be able to provide to potential adopters, now that they have proofs of concept that can be used as examples of using synthetic data and federated models in practice.

Interactions with and expectations of technology partners and providers

Amongst interviewees, there was a limited record of interactions with tech companies who might facilitate the deployment of PETs. The majority of interviewees noted that they had not had any interactions with tech providers on the matter thus far, therefore it is difficult from the interviews carried out to gauge the preparedness of industry to address the concerns of clients within the public sector who may be looking to adopt PETs.

With this caveat in mind, several respondents were able to provide some insights on the matter. This ranged from conversation with the interviewee from the ONS, who mentioned that all consideration and development of PETs applications had been handled in-house. This includes the examples of utilising synthetic data, as well as the development of proof of concepts for both federated learning and federated analytics. The interviewee also noted that there was appetite in exploring the utilisation of these PETs not only within their institution, but also throughout HMG, however there are a number of obstacles that would likely need to be addressed ahead of this being practical. These include:

- User-friendly software that is sufficiently easy to use, which could lower the technical expertise threshold
- Ability to assure adopters that the PET has been implemented correctly so that the purported safety guarantees hold

The interviewee from The National Archives expressed scepticism about the present capabilities of tech providers to provide tailored solutions that took unique requirements into consideration. Of these interactions with providers, the respondent noted that, as there is not a one-size-fits all solution, they have found that there has been a tendency by those offering the solutions to oversimplify matters. It may be, therefore, that additional work is necessary to improve the communication around what is practically feasible and encouraging open discussion about the limitations and uncertainties that remain around PETs.

This experience does, however, contrast with that of the interviewee from GDS, who suggested that there are a lot of existing off-the-shelf offerings, albeit for the purpose of facilitating simpler privacy preserving methods such as data minimisation, that are easy to adopt and unproblematic. However, when asked about the availability of tailored emerging PETs, they commented - with the caveat that they are not closely involved in these discussions - that, "my view is that that customizable solutions are not there yet, or I don't know whether they are there, but the channel of conversation has not started. [...] And one issue is that some of those suppliers with those customizable options tend to be quite small and sometimes they find it difficult to get into places like the NHS or central government. So it may be that communication has not started yet."

When considered collectively, these perspectives and the relative lack of interaction between those working within the public sector and those who may provide the technological solutions indicate that there is room for greater dialogue between these parties. As noted within the previous section, one of the obstacles to adoption is the challenge of identifying where emerging PETs might be used creatively to unlock greater value from existing data. This could be approached through convening both groups with familiarity of implementing PETs and those who have yet to consider their practical application within their organisation.

Given the limited interactions that have taken place between interviewees and external providers who might supply the technical solutions, it was not possible to determine the types of assurance that is currently being provided to public sector institutions by commercial PETs suppliers. This should therefore be explored further through engagement directly with suppliers themselves, which was outside of the scope of this research. It would be worthwhile asking suppliers what types of requests for assurance they have received surrounding the adoption of PETs and the questions that they commonly receive from interested parties about their products.

Use of traditional privacy-preserving methods and technologies within the public sector

While discussing the obstacles that interviewees encounter with regards to sharing and extracting value and insights from data, there was some discussion of the current efforts and technologies that they use in order to maintain privacy and minimise some of the risks associated with sharing data. This provided some interesting context to the methods that those interviewed are currently comfortable with, as well as providing insight as to how well they feel these presently address the challenges that they have.

While speaking with the respondent from GDS, the centrality of data minimisation within their practices was stressed. The interviewee went on to provide an example from during the coronavirus pandemic, where GDS had a lot of programs that supported the public with efforts such as shielding. Following this anecdote, the interviewee then reflected on the centrality of data minimisation as the basis on which further practices should then be considered in a layered manner: "I think when people talk about privacy enhancement and the technologies that come with it, there's always that fear of how much it's going to cost and how complicated the technology is going to be. But the primary solutions, actually, data minimisation doesn't need that. I mean, it's just an early thinking exercise and can be done with little costs and little resources". This observation was shared amongst many of the other interviewees who expressed similar sentiments around the present focus within their institutions on embedding good data practices, where they believe headway is being made. As a result, there appears to be less focus on emerging practices and technologies in the public sector as opposed to building confidence around traditional privacy preserving methods, such as minimisation, anonymisation, pseudonymisation and the use of encryption.

In conversation with the interviewee from the ONS, in which there has been greater consideration of emerging technologies for privacy purposes, it was mentioned that scrutiny around the value of pursuing these emerging technologies remains and that it is important for organisations to understand whether the investment in time, and understanding the breadth and extent of the risks is worth the investment.

While perhaps not considered as a typical means of privacy preservation, or enhancement, several of the respondents cited the use of systems and technologies in order to facilitate the handling of contracts and agreements when dealing with sensitive data. This included reference by the interviewee from the GLA to the use of tools such as Information Sharing Gateway[32], noting that "I think quite a lot of the time, people see as privacy enhancing technologies, some data trust or some kind of data intermediary, but actually at this stage the one that's kind of really being used and really being developed by the market is stuff to ease the signing off of legal agreements to ensure that data is shared safely, ethically and securely". It therefore appears that interpretations of what privacy enhancing technologies and processes are remains quite broad at present.

Current data access, infrastructure, and governance concerns

Each of the interviews conducted throughout this research began with a discussion of the ways in which the interviewees' institutions use data and the related challenges that they face in carrying out their work. Through these conversations, a fuller understanding was gathered of the broader data governance concerns that are held within parts of the UK's public sector at present. These extend beyond solely privacy-specific considerations. However, they provide important insights that are of relevance to the adoption of PETs and also provide additional utility in situating privacy concerns amongst others.

Access to data

Barriers to accessing data featured to varying degrees as a concern amongst the interviewees, which can perhaps be attributed to the range of officials interviewed and the types of

[32] See: https://www.informationsharinggateway.org.uk/ (accessed 23 May 2022)

institutions they work within, the type of data that they hold, or try to access, and the functions that they perform.

That said, the respondent from the ONS singled out access to data as the most pressing concern within their work, given the sensitivity of the data that they are dealing with and the fact that this can often be at the level of named individuals. As a result, the respondent detailed the various accreditations that are necessary to access data assets held by the ONS. This includes the aforementioned process of receiving accreditation in order to access ONS data assets through the SRS. This process would historically take several months, after which, access to the data asset would be limited as the individual would be required to access this on a secure machine. These time and resource intensive constraints on accessing data appear to have partially influenced the decision by the ONS to explore alternative solutions, as covered earlier within the report.

The interviewee from The National Archives similarly noted that privacy was a central concern within their work, given that they facilitate access to sensitive data, which can often be very rich with information about people and events. An added dimension to this challenge is that they are publishing both historic and contemporary data, which requires separate considerations. As a result, the interviewee spoke of operating within an "interesting grey area where we're having to make decisions about what we publish, what other people might digitise and publish on our [National Archives] behalf, and how we comply with the obligations that we have, particularly under data protection law, UK GDPR and the data minimisation principle[33]". They went on to note that ensuring adherence to these obligations has implications on what providing reasonable facilities means for The National Archives in terms of public access to records particularly when it is unclear whether certain privacy exemptions apply.

Specifically when dealing with born-digital records, the interviewee noted that The National Archives has yet to build the level of sophistication that they would like around access, which has resulted in material being either closed or published. Ideally, there would be the possibility of "gradated access", that would allow some modality of access, such as a shorter publication. It was noted that they are currently in the phase of conducting discovery work around this possibility, exploring the ways in which they could build a "gradated access" approach that allows - in a risk based way - to make decisions against imperfect information about the level of access being provided.

While also commonly cited as a significant obstacle to data sharing in other contexts, it is worth noting that the culture of risk aversion based on fear of doing something wrong was cited by the interviewee from GDS as the most significant barrier to greater data sharing within the public sector. Specifically, the interviewee cited confusion around legislation and a lack of clarity on the purposes that would be served by sharing data as contributing concerns. It was noted that this hesitancy is particularly marked amongst senior leaders in the public sector and that this is the primary obstacle that must be overcome.

[33] The "data minimisation principle" refers to article 5(1)(c) of UK GDPR, which stipulates 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')", see (https://www.legislation.gov.uk/eur/2016/679/article/5, accessed 23 May 2022)

Quality of data

Interviewees also regularly cited quality of data as a considerable concern that had impacts for the ability to share data more readily. This is another concern that frequently features across data dependent practices, however interviewees did question whether this could have implications for the use of privacy enhancing technologies such as federated learning and federated analytics. In addition to the desire for the use of common data models and data standards, an interviewee from GDS stressed that "Often we assume that obviously large departments hold a lot of data - which is correct - but the cleanliness of that data, the quality of the data, is sometimes not as good as we assume. And for that reason, sometimes departments take time to actually agree to share that data". This concern has implications for the potential utilisation of federated learning or analytic models, as some stewards may not have the time or skills to standardise data, which could in turn perpetuate a reluctance to make data assets available for training models or conducting analysis[34]. While there are potential benefits to the use of PETs such as federated learning, like with existing machine learning technologies, they are largely reliant on the supply of good quality data[35].

Expansion of services

A couple of interviewees remarked that the expansion of their service is presenting them with some new challenges that are impacting upon their data governance practices and procedures. One interviewee, who works with DataLoch, spoke of the expansion of their organisation's work, both in terms of the regions that it covers and the types of data that it holds. This has brought challenges in that they have been receiving non-standardised data from more diffuse sources and are, at the same time, dealing with requests for more complex research projects. Specifically, the interviewee noted that they are now receiving a greater number of requests from researchers for the results of analysed data, rather than access to data, as the service was originally set up for. This shift from providing access to a trusted research environment towards now conducting what the interviewee described as 'commissioned analysis' has led the organisation to have to reassess the balance between public benefit and privacy.

An additional challenge related to the extension of DataLoch's services concerns the aspiration to extend their service to encourage "non-typical researchers", such as charitable organisations, to make more use of data. In conversation, the interviewee commented that the governance around this is very difficult - particularly when working with the commercial sector who have developed tools to facilitate these types of efforts. Specific challenges that were mentioned include securing people's trust, securing different access privileges and ensuring that outputs are not identifiable on the way. This speaks more broadly to concerns that were raised by several interviewees, who noted that they currently feel comfortable with the systems that they are using and are confident in how they work. Conversely, the risks of adopting new technologies such as those included within the suite of emerging PETs are still not sufficiently apparent, which continues to act as a disincentive.

Another of the interviewees, from the GLA, spoke similarly of the challenges they are now encountering as the role of their core offering - the London Datastore - as a data publishing

[34] See https://medium.com/codex/ai-privacy-and-why-you-should-care-1ef503a789b6 (accessed 09 May 2022)
[35] Rieke, N. *et al* 2020 The future of digital health with federated learning (see https://www.nature.com/articles/s41746-020-00323-1, accessed 09 May 2022)

platform to one that is more proactively linking different datasets to provide insights. This development has raised questions internally on how to apply consistent methodologies to data projects across the board so that the administrative, legal and technical friction points can be navigated while concurrently maintaining privacy and creating a good product regardless of who the user is.

Conclusions, recommendations and opportunities

This research into the market readiness for PETs within the public sector in the UK has attempted to provide insights into the current level of uptake of these technologies and the impacts that this has had upon the data governance practices of institutions within the public sector. As outlined at the beginning of this report, it is apparent that there is a much greater degree of confidence amongst the majority of those interviewed with what are considered to be 'traditional' privacy enhancing technologies, such as anonymisation, pseudonymisation and data minimisation.

That is not to say that discussion and development of 'emerging' PETs, such as synthetic data and the use of federated models is entirely absent, as has been evidenced by the responses of the interviews that were conducted for this research. It is, however, worth acknowledging that some of the barriers that are currently impeding the uptake or exploration of emerging PETs are difficult barriers to move. This is reflected in the concerns raised by respondents throughout discussions of the challenges that they are primarily dealing with when attempting to share and derive greater value from existing data assets within their respective institutions. In essence, at the heart of many of the fundamental concerns held and barriers to greater adoption of PETs within the private sector is high continued high costs - both human and financial resources - to entry and data infrastructure that is not quite at the stage of actively facilitating their adoption.

There are, however, opportunities and suggestions that have arisen through the course of the interviews.

Develop further guidance

A common suggestion from respondents was the development of additional guidance around the adoption of emerging PETs. Such guidance should look to complement that which is already available, such as the CDEI's PETs adoption guide[36]. If possible, efforts should be made to elaborate on the existing repository of cases as a matter of priority, as this evidence of practical applications is particularly desirable for those in the public sector who are eager to equip themselves with the evidence by which to make the case for their consideration within their institutions. As noted previously in this report, there is now a degree of familiarity and comfort with traditional PETs and an understanding of how they fit within many institutions' existing data governance practices. This appears to be the case particularly amongst those who are responsible for authorising the exploration or development of prototype uses of these PETs. Additional evidence of both benefit and value are therefore required in order to disrupt this status quo.

[36] Centre for Data Ethics and Innovation 2021 Privacy Enhancing Technologies Adoption Guide (see https://cdeiuk.github.io/pets-adoption-guide/, accessed 09 May 2022)

Convene experience sharing or co-creation workshops

Lack of familiarity with the technologies, uncertainty around the associated data governance risks and competing demands for resources are but a few of the blockers that interviewees cited as obstacles to their institutions considering the adoption of emerging PETs more seriously.

Building on recent efforts by actors such as the ICO to co-create use cases for the deployment of PETs based on health data[37], targeted workshops should be considered in which those from within the UK public sector who have experience of developing PETs from an idea through to a proof of concept can be brought into contact with those who hold a curiosity, but lack the means and experience of doing so themselves. While an interesting thought exercise, interviewees often found it difficult to imagine potential applications of the various emerging PETs within their work. Doing so is not straightforward and it is evident that there are a multitude of considerations that need to be borne in mind, which is why learning from those who have this experience within the public sector could prove particularly valuable.

Explore partnerships between the public sector and smaller tech providers

One interviewee suggested that there could be benefits to encouraging greater partnerships between smaller tech companies and the public sector in the UK. The rationale provided for this suggestion was that larger tech companies would be less willing or likely to tailor their solutions to cater to the specific needs of actors within the UK's public sector. This also chimes with existing concerns around the cornering of the PETs market at an early stage in the increased adoption of these technologies, which might run the risk of further concentrating power and market position for the large tech providers[38]. Similar sentiment was noted earlier within the report, however further research would be required in order to explore the feasibility behind this proposal, given the

[37] Information Commissioner's Office 2022 ICO consults health organisations to shape thinking on privacy-enhancing technologies (see https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2022/02/ico-consults-health-organisations-to-shape-thinking-on-privacy-enhancing-technologies/, accessed 09 May 2022)

[38] Renieris, E. 2021 Why PETs (privacy-enhancing technologies) may not always be our friends (see https://www.adalovelaceinstitute.org/blog/privacy-enhancing-technologies-not-always-our-friends/, accessed 09 May 2022)

Appendix 1: Acknowledgement of interviewed institutions

The writing of this report would not have been possible without the generous support of experts from the institutions listed below, who contributed their valuable time and shared their experiences with the research team:

- Competition and Markets Authority (CMA)
- DataLoch
- Department for Transport (DfT)
- Government Digital Service (GDS)
- Greater London Authority (GLA)
- National Archives
- Office for National Statistics (ONS)

Appendix 2: Research methodology brief

The market research took place in three phases: project inception; primary research and synthesis and reporting.

Phase 1 - project inception:

Preliminary desk research was undertaken for the purpose of informing the semi-structured interviews that were undertaken in the second phase of the research. The literature was also revisited, on occasion, during the second phase of the research as the team revised and adapted the structure of interviews to accommodate a greater variety of potential responses. This was necessary as it became apparent that familiarity with PETs varied somewhat between the engaged stakeholders.

The project inception phase also included the fine tuning of the project methodology in consultation with the Royal Society - including agreement on interview questions prior to data collation through engagement with identified stakeholders - and the completion of the preliminary review of relevant literature.

Phase 2 - primary research:

The ODI conducted semi-structured interviews with relevant stakeholders during this phase of the research as the primary means of collecting data. Details surrounding the selection of stakeholders are included in the following section of the methodology brief, which outlines the engagement rationale.

In anticipation of instances where it may have proven unviable for interviews to be conducted, the research team developed the basis of an online survey modelled on the semi-structured interview questions included in the appendix. The survey aimed to gather primarily comparable qualitative data, however this also involved efforts to convert elements of the original interview structure into questions that would allow for respondents to provide pre-formed answers, to increase the likelihood of completion by respondents.

As the interviews with stakeholders got underway, it became apparent that it would be unlikely that the use of the survey would be necessary. Practically, the likelihood that the survey would

have been able to capture a similar degree of nuance to the interviews was unrealistic, given the range of responses received from stakeholders throughout the course of the interviews, given the variety of uses of data by the interviews stakeholders and the varying levels of awareness of PETs. As a result, each interview therefore required a greater degree of tailoring, to account for the role of the institution. A positive of this need for greater tailoring of the interviews was that a diverse range of responses were received from interviewees.

Sampling

The Royal Society provided a preliminary list of stakeholders to be interviewed. This list provided a good starting point and was iterated upon by the ODI, which drew upon its network of contacts to identify stakeholders. Responses were sought primarily from four types of UK institutions: those that use data to generate insights; who act as data gatekeepers; those that are developing infrastructure for data sharing and those publishing open data. While this project aimed to generate qualitative insights, thus reducing the necessity for an entirely representative sample, efforts were made to interview a variety of stakeholders from across the four types of institution identified.

From the preliminary list provided by the Royal Society, the ODI aimed to make contact with 20-30 individuals with the aim of securing between 10-15 interviews. In practice, the ODI contacted 36 individuals and managed to secure 7 interviews. Efforts towards securing interviews were stymied by a number of factors, including a lack of familiarity with the subject matter amongst contacted stakeholders and low response rates from those who were approached. That being the case, fruitful conversations were had with eventual interviewees, who spanned the types of institutions that we had hoped to reach. This included a mix of respondents from local and central government, respondents from collaborations between government and academia and respondents from regulatory bodies.

Phase 3 - synthesis and reporting:

Upon completion of the interviews, the ODI synthesised the findings into this final report. This synthesis brings together the interview responses to provide insight into the role that PETs are currently or could potentially play in organisational use and governance of data in the public sector.