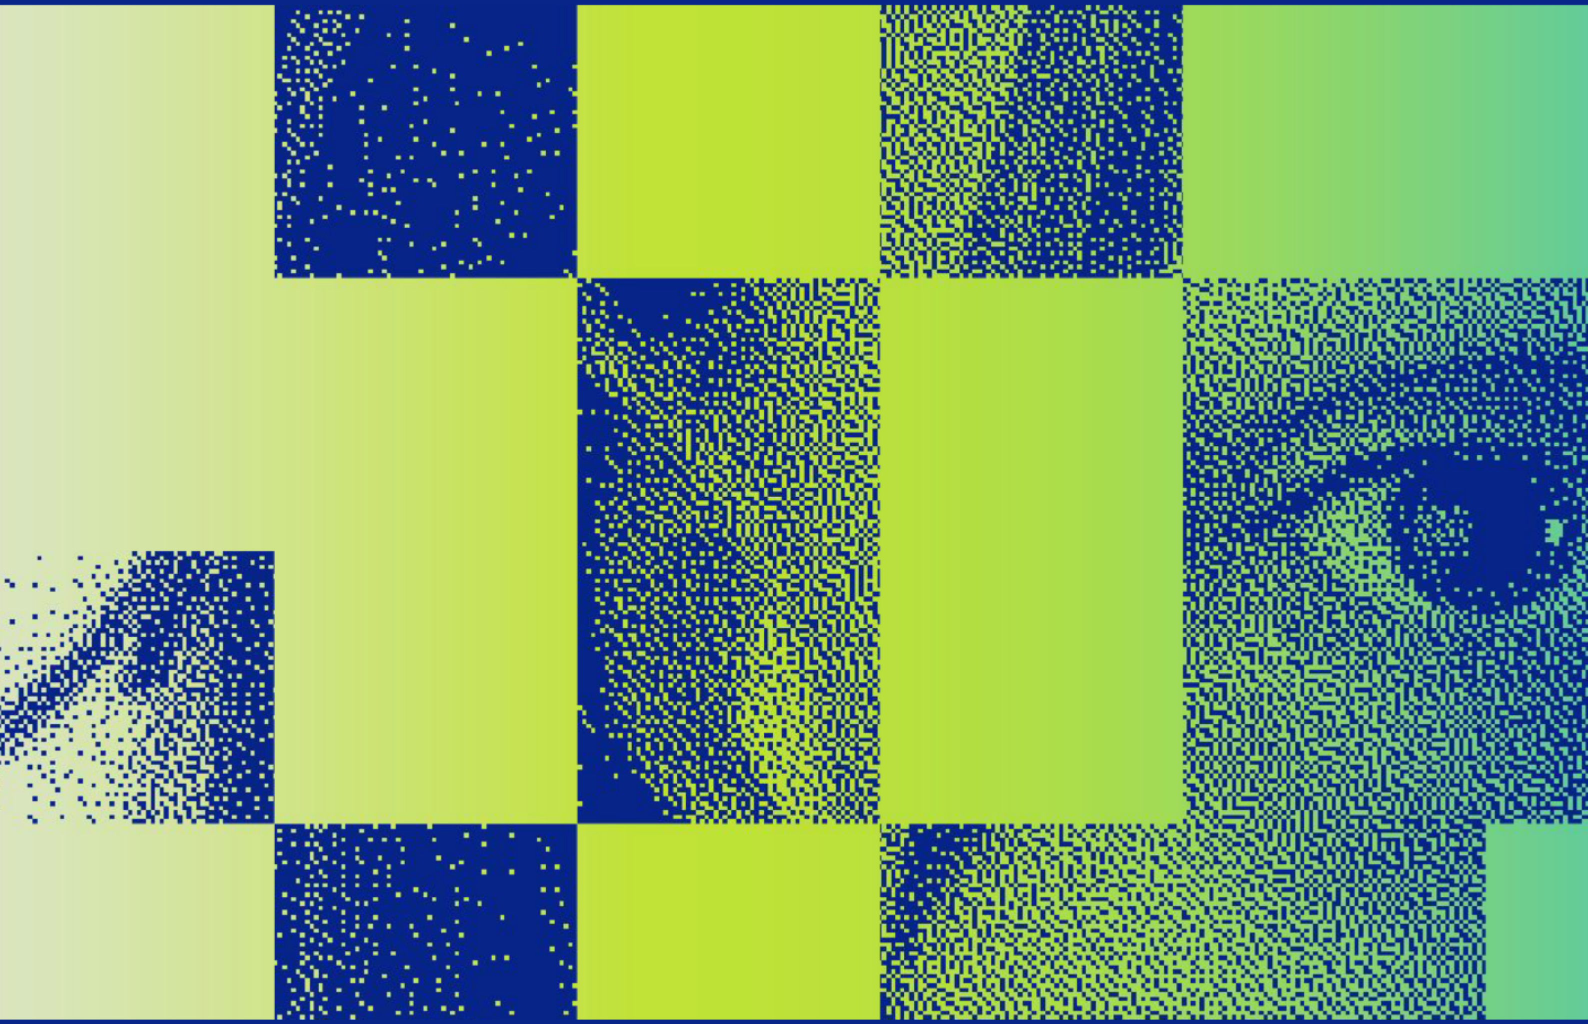

PETs in Practice: using trusted research environments to enable health data research

June 2024

Privacy Enhancing Technologies
ODI Research
ADVANCING TRUST IN DATA

Contents

About

This report has been researched and produced by the Open Data Institute, and published in June 2024. Its lead author was Neil Majithia, supported by Elena Simperl, Calum Inverarity, Jared Keller and Claudine Tinsman. If you want to share feedback by email or would like to get in touch, contact the PETs programme team at pets@theodi.org.

Suggested citation: Open Data Institute (ODI) (2024), ‘PETs in Practice: using trusted research environments to enable health data research’

To share feedback in the comments, highlight the relevant piece of text and click the ‘Add a comment’ icon on the right-hand side of the page.

PETs in Practice: using trusted research environments to enable health data research

Trusted research environments (TREs) are a type of privacy-enhancing technology (PET) that can enable analysis of sensitive data whilst preserving the privacy and security of the underlying data. TREs are flexible socio-technical architectures that can be used in a number of contexts. They are often deployed in healthcare settings to enable granular research that would otherwise be difficult, if not impossible.

At the Open Data Institute (ODI), we are interested in understanding the requirements that must be met for PETs to be adopted by organisations of all sizes in a wide array of use cases. This report outlines a case study documenting the development and application of SAIL Databank, a TRE which has been in operation since 2007 to enable researcher access to sensitive health and social data of individuals in the UK for specific research projects that are of benefit to the public.

There is value in increasing access to sensitive data for social benefit.[1] We hope that this case study will prove useful for potential adopters of TREs, policymakers exploring data access solutions, and technology providers developing and implementing PETs.

For this case study, we gained insights into the development and operation of SAIL Databank through publicly available materials, discussions with those at SAIL and responses to interview questions with those involved in the development and implementation of this TRE. This research was undertaken between November 2023 and March 2024.

[1] Understanding Patient Data (2018), '[Data for Public Benefit](#)'

Background

Impactful clinical research is ruled by an immutable law: in order to find the most useful insights for healthcare and medicine, clinical researchers require detailed clinical and demographic data.[2]

A study by Dadvand et al. (2014)[3] provides an example of the benefits that can be reaped through access to granular data from different sources. The researchers analysed sensitive data on newborn babies with low birth weight (LBW) and their mothers' residential locations and proximity to major roads to find a statistically significant result: for a pregnant person, living within 200 metres of a major road is associated with a 46% increased likelihood of LBW for their child, with a third of that relationship being the effects of heat and air pollution.

There are thousands of pieces of research[4] that follow this same methodology: analysing patients' sensitive clinical data against their identifiable demographic data to find valuable contributions to medical science. In the Dadvand et al. case, this led to the adaptation of advice for pregnant people to mitigate the risks of LBW, therefore benefiting newborns worldwide. Such impact highlights just how important access to granular clinical and demographic data is for modern day clinical research.

However, the more granular data is, the more sensitive it is; more granular datasets provide more detail about the individuals within them, for example their residential address or details of their family's health issues. If a data leak occurred, these details could be revealed to the world. Malicious actors, if they got hold of the data, could take the clinical information and location data and potentially use it to identify the individuals and families by linking them to records in other datasets obtained legally or otherwise. With this 're-identification', third parties could use the health data against the individuals for blackmail, predatory targeted advertisement or even raising insurance premiums.[5]

[2] Kornegay & Segal (2013) 'Developing a Protocol for Observational Comparative Effectiveness Research: A User's Guide; Chapter 8: Selection of Data Sources'

[3] Dadvand et al. (2014), 'Residential proximity to major roads and term low birth weight: the roles of air pollution, heat, noise, and road-adjacent trees'

[4] This is one of many cases collected in Residential proximity to environmental hazards and adverse health outcomes, in which all papers reviewed used residential data as a key demographic to understand health outcomes. Other papers use other demographics, including age and gender (Luo et al. 2016), smoking status (Chang et al. 2020), ancestry and genealogy (Stefansdottir et al. 2012), educational and legal history (Sharp et al. 2011), and others.

[5] Understanding Patient Data (2024), 'What are the risks around patient data?'

These risks provide the incentive for data stewards like the NHS to keep data as secure as possible, potentially limiting the extent to which it is shared with researchers in order to mitigate the risk of privacy loss.[6] These stewards must comply with data protection legislation, but risk aversion can have the consequence of slowing down and stifling research. Papers like that by Dadvand et al. wouldn't be possible without the granular data underpinning it, meaning medical care would suffer as a result.

Privacy-enhancing technologies (PETs) can tackle this challenge by facilitating safe access to sensitive data.[7]

A trusted research environment (TRE), also known as a data safe room or data haven,[8] is an example. A TRE is an infrastructure designed to provide researchers with access to sensitive data whilst simultaneously ensuring no privacy loss and no access for malicious or unintended actors. Built around the 'Five Safes' framework in the UK,[9] TRE host organisations modify data to certify that it is as anonymous and de-identified as possible, while concurrently subjecting researchers who want to use the data to rigorous authorisation protocols. Once authorised, these researchers can use only hardware administrated by the TRE hosts to access the anonymised data and analyse it, completing their research and exporting the results. This export is audited by TRE administrators to certify that no data leaves the TRE, ensuring impactful research can be conducted without privacy loss.

This report focuses on one specific TRE, SAIL Databank,[10] one of the first TREs introduced in the UK and an award-winning[11] architecture for the stewardship of sensitive clinical and demographic data. SAIL's thorough infrastructure and core focus on public benefit make it a prime example of a PET in practice with real impact.

[6] Davies & Collins (2006), Balancing potential risks and benefits of using confidential data

[7] ODI (2023), 'Privacy enhancing technologies (PETs)'

[8] As well as TRE, 'safe rooms' and 'data havens', this infrastructure can also go by a number of other names, including 'clean rooms', and 'secure', or sometimes 'virtual', research environments. While the degree to which they are similar varies, the commonality between these infrastructures is that they are designed to foster collaborations with sensitive data in a way that is secure.

[9] UK Data Service (2024), 'What is the Five Safes framework?'

[10] SAIL Databank (2021), 'Home - SAIL Databank'

[11] Health and Care Research Wales (2023), SAIL Databank recognised with prestigious Queen's Anniversary Prize

What is SAIL?

SAIL (Secure Anonymised Information Linkage) Databank is a rich source of information about individuals in Wales, containing over 10 billion[12] person-based records that enable clinical researchers to answer important questions for public benefit.[13]

SAIL Databank was developed and is now operated by the Health Information Research Unit (HIRU) housed in the Population Health Science department of Swansea University Medical School. SAIL was first released as a small-scale pilot in 2006 and has since facilitated over 700 published papers in medical journals[14] and supported research valued at over $50 million.[15] SAIL receives funding from Health and Care Research Wales and the UK Research and Innovation (UKRI) Economic and Social Research Council (ESRC).

What does a TRE involve?

A TRE is, by nature, closed, meaning the data it holds is only accessible to a select set of entities. SAIL Databank operates on a 'no data leaves' model[16] where data, once integrated into the Databank infrastructure, never exits the ecosystem. Researchers can read and analyse the data within the ecosystem, therefore ensuring a completely closed environment.

Such a closed environment means important insights in the medical field can be found. By guaranteeing data security, the most detailed clinical data can be used for research that requires it. This research is in fields such as infections, child health, chronic conditions, education and mental health, providing insights that directly contribute to healthcare planning and strategy, healthtech research and development, and medical diagnoses – all for public benefit.[17] Specific examples include Roberts et al. (2013),[18] which used SAIL Databank to identify the impact of demographic factors on the incidence of acute pancreatitis, and James et al. (2018),[19] which examined the effect of secondary school intervention on physical activity levels.

[12] SAIL Databank (2021), 'Data'

[13] SAIL Databank (2021), 'About us'

[14] SAIL Databank (2023), 'Publications'

[15] Swansea University Population Data Science (2023), 'Secure Anonymised Information Linkage (SAIL) Databank'

[16] UK Health Data Research Alliance (2021), 'Building Trusted Research Environments'

[17] SAIL Databank (2021), 'Data use register'

[18] Roberts et al. (2013), 'The incidence of acute pancreatitis: impact of social deprivation, alcohol consumption, seasonal and demographic factors'

[19] James et al. (2018), 'Active children through individual vouchers – evaluation (ACTIVE): protocol for a mixed method randomised control trial to increase physical activity levels in teenagers'

TREs sometimes employ trusted actors outside the host organisation. SAIL Databank is supported by Health Solutions Wales (HSW),[20] which acts as a trusted third party (TTP) for the anonymisation of the datasets as they enter the Databank. SAIL's team (the host organisation) is meanwhile concerned with stewardship of the datasets, including coordination of the authorisation protocol for prospective researchers and audits/supervision of user research.

Key parties external to SAIL's infrastructure are the data providing organisations (DPOs). DPOs are pillars of the SAIL Databank infrastructure given that SAIL does not collect any data itself, meaning all data stored within the system is from third parties. These third party DPOs are places where sensitive data is collected and stewarded, such as health services, social services, national registers and census organisations.

Initiating a partnership between a DPO and SAIL involves the DPO volunteering (or being asked to volunteer) to donate their data and telling SAIL how they would like it to be stewarded. There are three options for this. The most restrictive (**'project-specific'**) means the dataset isn't listed on the SAIL website and is only provided to research projects approved by both SAIL and the DPO. The least restrictive (**'core data'**) integrates the DPO's data into SAIL's ecosystem and allows it to be used by any SAIL-approved research project at the discretion of SAIL themselves. The middle option (**'core restricted'**) has the dataset listed on the SAIL website but still requires the DPO's consent for use on a project-by-project basis. The overall contract is determined by a data sharing agreement, completed by the DPO with their chosen option and other associated details such as refresh rates or metadata. These measures ensure that agency over the data remains firmly with the DPO, thereby providing assurances regarding the safety of their data as they enter a partnership with SAIL.

On the other side of the infrastructure lies the authorisation protocol for researchers wanting to use the data. A strict methodology is used by SAIL,[21] working with the potential researchers to understand their proposed projects and an independent review panel assessing them. CVs and proof of Safe Researcher Training[22] are required so that the panel can decide whether to provide access to the data. The process takes around 12 weeks, unless the researchers want to use non-core data, in which case DPO consent adds another layer to the process.

[20] Digital Health and Care Wales (n.d.), 'Home - Digital Health and Care Wales'

[21] For an in-depth walkthrough of this methodology, please see our corresponding Kumu presentation: ODI (2024), 'PETs in Practice: using trusted research environments to enable health data research'

[22] UK Data Service (2021), 'Safe Researcher Training'

When complete, researchers are provided access to the TRE via a USB dongle and password, which facilitates remote access to a SAIL Gateway environment that is custom made for their project, containing only the tools and datasets requested by the researchers.

This system maintains control of the data and trust in all parties that see it, facilitating safe access to the data while mitigating risks of data leakage as far as possible to a standard accredited by ISO 27001, the UK Statistics Authority and Cyber Essentials.

Incentives for actors in the SAIL Databank ecosystem

SAIL Databank has functioned without incident since 2007 because it meets the requirements of each of the parties involved in the ecosystem. Researchers get access to the sensitive data that is most useful for their research, with means to analyse it and support from SAIL to deliver insights that can shape the medical field. DPOs, who face the conflicting responsibilities of keeping data secure and facilitating research, reconcile this with the use of a TRE, which fulfils both obligations. In fact, DPOs are provided the ultimate control of their data even when it is held in SAIL: they can immediately take their data away whenever they like. Such control keeps stewardship in the hands of the DPOs and helps ensure that they retain sovereign control over the data that they have made available.

All parties that contribute towards the infrastructure have an ultimate responsibility towards the public, who SAIL keep front and centre of considerations at all times. The data being stewarded and shared is about the public, the medical research that comes out of SAIL affects the public, and if anything goes wrong, it is the public who would potentially suffer from re-identification attacks. SAIL represents the public as the major stakeholders in their authorisation process, involving everyday people in their review panels to ensure that every project for which access to the data is allowed is for the public benefit.

How SAIL works in practice

The figure below is a visual representation of one part of the SAIL Databank TRE. A fully annotated, interactive diagram, with a stage-by-stage narrative of the framework, is available on Kumu.

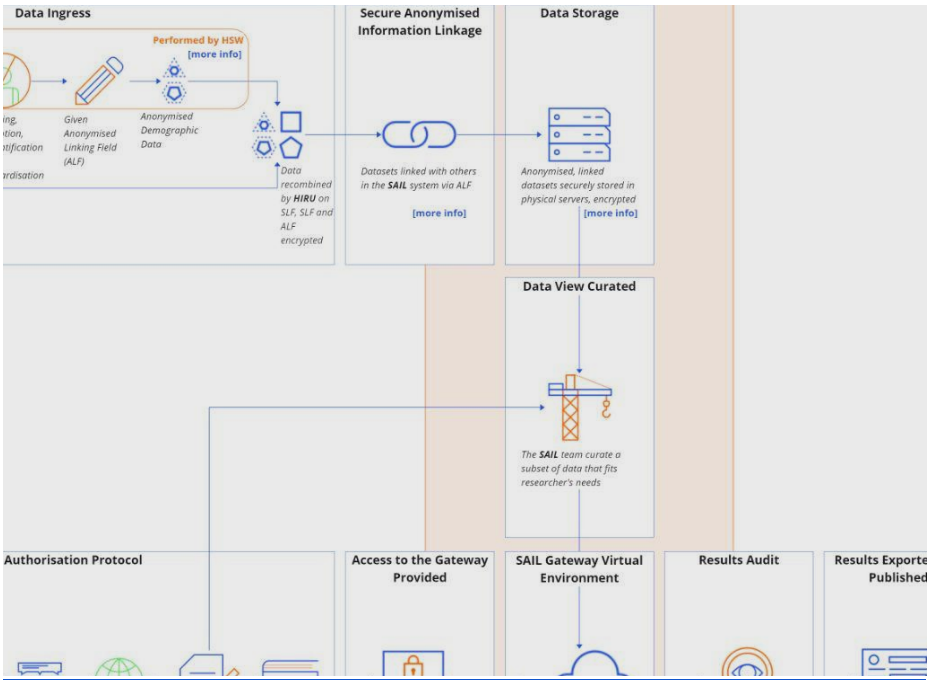
Data Ingress
Performed by HSW
[more info]
ing,
tion,
tification
rdisation
Given Anonymised Linking Field (ALF)
Anonymised Demographic Data
Data recombined by HIRU on SLF, SLF and ALF encrypted
Secure Anonymised Information Linkage
Datasets linked with others in the SAIL system via ALF
[more info]
Data Storage
Anonymised, linked datasets securely stored in physical servers, encrypted
[more info]
Data View Curated
The SAIL team curate a subset of data that fits researcher's needs
Authorisation Protocol
Access to the Gateway Provided
SAIL Gateway Virtual Environment
Results Audit
Results Exporte
Published

Figure 1: A snapshot of the deployment of SAIL Databank

Image source: Open Data Institute

SAIL as an early TRE adopter

The SAIL team at HIRU was established in 2006, with the first iteration of the Databank being a physical (on-site) TRE beginning in 2007 with a team of 10 people. Its entire infrastructure and design had been planned pre-implementation, with governance practices and use of TTPs built into SAIL's philosophy before any data was handled by the team. It expanded to being an online, remote-access TRE in 2011 due to user demand and has since facilitated thousands of research papers in the medical field. Remote access is facilitated by the SAIL Gateway, which was built and is now maintained by SeRP, a spin-off from HIRU.[23]

[23] SeRP (2020), 'Background and History'

SAIL's implementation of the Gateway in 2011 was only one year after the idea of remote-access TREs was first studied. In an initial paper on the subject, produced by funding agency JISC and published in 2010, candidate designs of remote-access TREs were discussed, with special attention to collecting all knowledge on the subject at the time in one place; a collection leveraged by SAIL's team to design their remote-access infrastructure.[24]

SAIL's original iteration as a physical TRE was arguably its safest iteration. With no remote access, researchers, once approved to access the data, had to physically travel to Swansea to carry out their research in a data clean room on SAIL-controlled computers. SAIL employees could watch what the researchers were doing and physically make sure that nothing malicious happened. Although remote access enables accessibility and convenience, the Databank loses a modicum of security: in Kavianpour et al. (2022),[25] TRE hosts acknowledge that remote-access frameworks mean they have no guarantee that the devices that approved researchers use to access their data are safe. A malicious actor could set up a screen capture or even steal the dongle/passwords in order to get hold of SAIL data. Ultimately, faith must lie in the authorisation protocol for this sort of situation not to occur.

Using TREs in other sectors

There is a general excitement in the health science community surrounding TREs. There are already so many in use for medical research that in 2023 DARE UK (Data Analytics Research Environments UK, set up by UKRI) awarded over £2 million to projects that could inform the standardisation of TRE processes.[26] Meanwhile, the Goldacre Review recommended that TREs be the new norm for research using electronic patient records (EPRs) in the UK, stating that "all analysis of NHS patient records should move to be done in a TRE".[27]

[24] Ohio College Library Center (OCLC) Research (2015), 'Virtual Research Environment (VRE) Study (OCLC/JISC)'

[25] Kavianpour et al. (2022), 'Next-Generation Capabilities in Trusted Research Environments: Interview Study'

[26] DARE UK (2023), 'Phase 1 Driver Projects'

[27] Goldacre, B & Morley, J. (2022), 'Better, Broader, Safer: Using health data for research and analysis. A review commissioned by the Secretary of State for Health and Social Care. Department of Health and Social Care.'

SAIL itself progressed from stewarding only clinical data to other social care data, including government and third sector datasets which have similar requirements to health data in terms of privacy and sensitivity.[28] Such requirements are present in data everywhere, to varying degrees; social media data, consumer data from e-commerce, and usage data from apps are all examples of where an individual's demographic data can be analysed against items such as their social media posts, online purchases or rideshare activities for either beneficial academic research[29] or malicious re-identification attacks.[30] Each of these types of data could benefit from the implementation of a TRE.

For social media data in particular, TREs are a way for online platforms to respond to Article 40 of the European Union's Digital Services Act (DSA), which as of 2024 introduces legislation to make researcher access to the online platforms' data mandatory.[31] These platforms now face a similar problem to that which health data stewards have: they must balance the security and safety of their users' data with transparency reporting requirements. Introducing a TRE or TRE-like structure would allow the platforms to maintain control of their users' data while conforming to the DSA's demands regarding access. Given the public's heightened awareness of how their social media data is used,[32] a transparent TRE architecture similar to that of SAIL, which prioritises public benefit, could go some way towards appeasing concerns of both users, researchers and law-makers.

Implementing a TRE for any type of data, whether it comes from big technology firms or small businesses, could be difficult for a potential host organisation. TREs have an associated set-up cost, where the infrastructure must be prepared before any data comes in, and the parties involved in the process must have capacity and be trained to support the workflow, making small-scale pilots resource-heavy and time-consuming. In the case of TREs, there are a number of pre-existing, successfully implemented examples that can provide inspiration or perhaps even a blueprint for tailored deployments. To make the development of a bespoke TRE easier, an organisation could follow examples of existing TREs such as SAIL and use pre-made parts of the process, for example the SeRP remote-access system[33] for SAIL Gateway, which has been used in other health sector TREs such as the Dementias Platform Australia.[34]

[28] For example, SAIL holds a dataset on the outcomes of Children and Family Courts proceedings (https://web.www.healthdatagateway.org/dataset/8ee61578-e298-423a-be22-cb0438023e5c). Likewise, they hold other information relevant to government or third sector work, like the dataset on 'Looked After Children Wales'.

[29] Abdul Ghani et al. (2019), 'Social media big data analytics: a survey'

[30] Wernke et al. (2014), 'A classification of location privacy attacks and approaches'

[31] EUR-Lex (2022), 'Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act)'

[32] Centre for Data Ethics and Innovation (2023), 'Public attitudes to data and AI: Tracker survey (Wave 3)'

[33] SAIL Databank (2021), 'Analytical Tools & Environment'

[34] SeRP (2020), 'Case Studies'

However, some components of the SAIL infrastructure might prove unnecessary for new adopters. SAIL operates as an intermediary and does not collect or own the data it stewards. Social media companies are the opposite and therefore would design their TREs with key differences to SAIL Databank, perhaps removing the long ingress protocol with DPOs and TTPs, and carrying out the matching, standardisation, encryption and linkage in-house.

Adapting SAIL's authorisation protocols is another way to reduce complexity. A TRE with less sensitive data may not require such a rigorous protocol, and although that protocol is there to ensure research projects are beneficial to the public, certain responsibilities can be redistributed. Credential systems like OpenID[35] that can be built on OAuth[36] and require third party certification are potential methods to speed up and/or centralise authorisation protocols, although theoretically loosening security depending on the implementation method. It is likely that further innovations will be found as research and literature on the subject grows.

The general TRE infrastructure is broadly flexible to the needs of the organisation it serves, but the idea behind it remains the same: keeping sensitive data safe while simultaneously facilitating access to it for researchers.

Other approaches to TREs

As already noted, SAIL's first incarnation as a physical TRE is arguably its most secure iteration. The risks that come from remote access are mitigated with this approach, and although making research less accessible, this implementation provides a potential approach for the most sensitive data out there.

TREs with different aims also require different infrastructures to SAIL. OpenSAFELY[37] was built as a way to deliver urgent results necessary for COVID-19 research with a core aim of preserving collaboration; economies of scale are vital for efficient global emergency response. To facilitate this, the environment in which researchers do their analysis is built on Git,[38] allowing all research and all analyses' code to be read, used and edited by others. Such a system increases efficiency but also comes with the side-effect of increasing transparency (as all analysis is open source) and accuracy (as all researchers can scrutinise each others' analysis methods).[39]

[35] OpenID (n.d.), 'What is OpenID Connect'

[36] OAuth (n.d.), 'OAuth 2.0'

[37] Bennett Institute for Applied Data Science (2024), 'About OpenSAFELY'

[38] OpenSAFELY (2024), 'OpenSAFELY'

[39] OpenSAFELY (2024), 'About'

To maintain security in the collaborative environment, OpenSAFELY also diverges from SAIL's architecture by providing researchers with synthetic data when they are building their analyses.[40] Researchers don't see the real EPRs when they are writing their code, they instead use synthetic data to build a robust methodology first. After passing an approval process, the researchers send their code to the secure environment where the EPRs are kept. The code is run there and the results are sent back to the researchers, who, again, never actually see the real data but nonetheless are provided with the outcomes of their analysis. This architecture therefore goes one step further than SAIL in terms of keeping data secure, albeit at the cost of a more complicated infrastructure.

Synthetic data is also used by the Office for National Statistics (ONS) Integrated Data Service (IDS) TRE, both as a means for researchers to become familiar with datasets pre-accreditation to access the IDS and for training researchers post-accreditation.[41] For the former, the process for obtaining the accreditation necessary to access certain ONS data can be lengthy and resource heavy for both applicants and the ONS. This can mean a significant and unnecessary cost for both parties in instances where it turns out that the requested dataset is not suitable for the analyses it was requested for. Therefore, pre-accreditation access to privacy-preserving synthetic data based on the requested data can provide prospective researchers with a sense as to whether the dataset meets their expectations and can inform the researcher in making the decision to proceed with the accreditation process required to access the original dataset, or to abort, thus saving all parties further time and expense.

For post-accreditation training, the IDS's user base contains governmental departments, local administrations and external researchers, each with varying levels of technical expertise and familiarity with the TRE. This is in contrast to SAIL and OpenSAFELY, whose users, as part of academic institutions or healthcare systems, tend to have an understanding of data query structures and TREs. Therefore, the IDS has the extra step of training researchers by providing synthetic data that resembles the real data within the TRE (in terms of structure) to build familiarity and make the infrastructure more accessible.[42]

Typical archival structures such as the Wellcome Collection[43] or the UK Medical Heritage Library[44] have similar stewardship roles to SAIL, although they are typically centred more on items that require physical storage such as

[40] Threaux, Olivier (2019), 'Anonymisation and synthetic data: towards trustworthy data'

[41] Office for National Statistics (2023), 'Enabling Data Access through Privacy Preserving Synthetic Data'

[42] Ibid.

[43] Wellcome Collection (n.d.), 'Archives at Wellcome Collection'

[44] UCL (n.d.), 'UK Medical Heritage Library'

rare books or genetic material. These sorts of items require surveying, organisation and preservation efforts, which are all analogous to similar roles within SAIL's ecosystem. After all, the obligations that archives face regarding security, preservation and facilitating researcher access are not any different from the obligations that stewards of health data face, and these are obligations that archival processes and TREs both solve.

Conclusion

This case study has explored how SAIL Databank is a valuable framework in the world of healthcare research which facilitates navigation of the privacy trade-off that comes with handling sensitive clinical data. The benefits of the TRE approach are outlined and alternatives to the relative stringency of SAIL's approach are discussed for contexts outside of healthcare, including social media research.

The design of the SAIL TRE framework is tailored to address SAIL's objectives as an organisation; requiring trusted third parties in data ingress and particularly strong authorisation protocols for researchers, but the underlying architecture is flexible and is a good option for organisations trying to facilitate safe research in the new era of increased data privacy awareness.

We believe that breaking down real-world implementations of PETs such as SAIL into their composite structures and components can bring increased awareness and understanding of novel PETs and how they work. It is hoped that explainers such as this and the corresponding annotated diagram can be useful in enabling their wider adoption and contribute towards the documentation of case studies.

Get in touch

Please **get in touch** if you want to share any feedback or need help sharing sensitive data. We welcome all forms of input.

Acknowledgements

The writing of this report would not have been possible without the help and insights provided by the team at SAIL Databank, including Jon Smart. We are incredibly grateful for the time they contributed to answering our questions on their experiences of developing and deploying SAIL.

We would also like to express our thanks to our colleagues at the ODI and Philpott Design who have contributed their time and assistance to the creation of this written report and the annotated diagram. This includes Matt Barber, Charlotte Mitchell, Emma Thwaites, David Dinnage, Lou Lynch, Zoe Philpott, Adrian Philpott and Ian Dutnall.