UK Health Data Research Alliance

Trusted Research Environments (TRE)

A strategy to build public trust and meet changing health data science needs

UK Health Data Research Alliance

Draft Green Paper v1.0 dated 30 April 2020 – For consultation

Table of Contents

UK Health Data Research Alliance

Executive Summary

The UK Health Data Research Alliance is an independent alliance of data providers, custodians and curators dedicated to improving human health by maximising the potential of multiple forms of data at scale.

Our vision is that every health and care interaction and research endeavour will be enhanced by access to large scale data and advanced analytics.

1.Why does this matter?

Even before the COVID-19 pandemic, challenges to human health and health system sustainability have been increasing across the world. Whilst heart disease, stroke and cancer still account for nearly two thirds of all deaths globally, increasingly people are living with multiple diseases and long-term conditions. These affect us deeply – they change the lives of those with the diseases and those who care for them. By making health data available to researchers, we can develop a better understanding of these diseases and find ways to prevent, diagnose, treat and manage or cure them. It is also essential for tackling the direct and indirect impact of pandemics and other public health challenges.

2.Why is the UK the best place to do this?

The UK has some of the richest health data anywhere in the world. With the NHS it is feasible to collect longitudinal health data on a large and diverse population, and to make national-scale improvements to health and care. Combined with unique research expertise, outstanding talent in the NHS and universities, and vibrant life sciences and technology industries, the UK has an unprecedented opportunity to use data at scale to drive innovation, grow the UK industry base and improve the long-term health of the public. For example, the RECOVERY trial[1] is currently the world's largest trial of COVID-19 drugs[2].

3.**How do we ensure this happens in a safe way that retains and enhances public trust?**

Research conducted by Understanding Patient Data identified that people are generally comfortable with anonymised data from medical records being used for improving health, care and services, and for research, provided there is a public benefit. [3] The more informed people feel, the more they are likely to support these uses. However, people are more likely to be uncomfortable with the idea of commercial companies accessing their health data, and there are concerns about information being passed on for marketing or insurance purposes. There are recent examples of projects which have been widely reported in the media and may have increased public concern about health data use[4].

An unequivocal statement by the UK Health Data Research Alliance (the 'Alliance') committing to an approach to data access based around Trusted Research Environments, and with appropriate robust and independent TRE accreditation, monitoring and auditing, is a unique chance to address these public concerns and enhance public confidence in the use of health data for research in the UK.

---

[1] https://www.recoverytrial.net/
[2] https://www.bbc.co.uk/news/health-52478783
[3] https://understandingpatientdata.org.uk/how-do-people-feel-about-use-data
[4] For example: https://www.theguardian.com/commentisfree/2020/feb/16/our-personal-health-history-is-too-valuable-to-be-harvested-by-tech-giants, https://www.theguardian.com/technology/2020/feb/08/fears-over-sale-anonymous-nhs-patient-data

UK Health Data Research Alliance

This would mark a significant change for many researchers using health data. Rather than extracts of individual level data being distributed to researchers, TREs would provide access to a secure analytics environment (i.e. a safe setting) where researchers could bring analysis algorithms to the data. It is vital that users understand this approach has benefits to them as well and that they also develop trust in these environments. The benefits for researchers are discussed more fully in a blog by Rhoswyn Walker, Health Data Research UK's Chief Science Strategy Officer.[5]

There are now multiple examples of TREs operating successfully in this way, both for healthcare data and other potentially sensitive data. These include, but are not limited to:

- Scotland Data Safe Haven programme[6]
- UK Secure eResearch Platform in Wales[7]
- Genomics England Research Environment[8]
- UK Data Service Secure Lab[9]

This is also the approach being adopted in the development of a national health data research capability to support COVID-19 research questions[10].

There are a growing number of practical and cost saving benefits to this approach. It can maximise the utilisation of High-Performance Computing, whilst avoiding the costs of transferring and storing duplicates of increasingly large datasets, particularly imaging and genomic modalities. It also avoids the liabilities for users from having to ensure security of downloaded datasets.

However, questions remain and this document is intended to be the start of a dialogue with all stakeholders to develop an approach that meets the changing needs of health data science to deliver public benefit whilst protecting individual privacy.

Status of the document

The paper is currently a draft Green Paper. Comments are welcome as are contributions to its further development. It has been subject to initial review and feedback by representatives from the Trusted Research Environments Workstream of the UK Health Data Research Alliance and HDR UK's Public Advisory Board. It is now being circulated for wider consultation with patients, the public, researchers and innovators to determine best practice approaches for research on UK health data. The intention is to consolidate feedback into an updated paper and implementation plan for sharing at the HDR UK One Institute Event on 16th June.

[5] https://www.hdruk.ac.uk/news/why-take-the-risk/
[6] https://www.nhsresearchscotland.org.uk/research-in-scotland/data/safe-havens
[7] https://saildatabank.com/wp-content/uploads/UKSeRP_Brochure_v1.5.pdf
[8] https://www.genomicsengland.co.uk/about-genomics-england/research-environment/
[9] https://www.ukdataservice.ac.uk/use-data/secure-lab/how-it-works.aspx
[10] https://www.hdruk.ac.uk/wp-content/uploads/2020/04/200416-COVID19-Research-Data-Final.pdf

UK Health Data Research Alliance

Overview

Purpose

Health data used for research and innovation comes from a variety of sources, but most relates to peoples' interaction with the health and care system in some way – for example as an NHS patient, a participant in a clinical trial, being involved in a genomics initiative or as a blood donor. Therefore, achieving the **confidence and trust of patients and the public** in the use of this data is central to achieving our vision.

This paper defines criteria for **Trusted Research Environments** (TREs) which protect - by design - the privacy of individuals whose health data they hold, while facilitating large scale data analysis using High Performance Computing that increases understanding of disease and improvements in health and care. It also sets out processes for the transparent, independent accreditation of TREs that will make up the infrastructure of the UK Health Data Research Alliance.

Background

The UK Health Data Research Alliance is an alliance of leading health, care and research organisations united to establish best practice to enable the ethical **use of UK health data for research and innovation** at scale[11]. A central challenge in using health data is how to facilitate research while protecting privacy and so engendering public trust.

The Office of National Statistics (ONS) which facilitates research access to similarly sensitive administrative data described in 2017[12] its role as to *"find a way to maximise the use of the detailed data that ONS holds, while keeping them secure at all times; to let government, academics, businesses and others use these data, while being able to assure you [the public] that you will never be identified, your private details will never become public and that the information you have given us will only ever be used in ways that clearly serve the public good*". Their approach is summarised as "Five Safes": Safe people; Safe projects; Safe settings; Safe outputs; Safe data. These "Five Safes" should be considered as adjustable controls rather than binary settings. Risk is addressed by complementary adjustments on the implementation of each "Safe" to provide an appropriate context for research to occur which maintains an optimal balance between research benefit and overall risk management.

The term "Data Safe Havens" has been used to describe systems for providing researchers with access to data while managing risk of unauthorised re-identification of individuals from de-identified data, however

[11] https://ukhealthdata.org/

[12] https://blog.ons.gov.uk/2017/01/27/the-five-safes-data-privacy-at-ons/

implementations vary considerably[13]. When evaluated against the "Five Safes" defined by ONS, common features of Data Safe Havens are processes of evaluation of research proposals before granting access (Safe projects) and robust processes to de-identify and anonymise data being accessed to reduce the risk of re-identification (Safe data). However, in many cases the model of access is one of "Data Release", i.e. where processed datasets are distributed to researchers, rather than requiring them to carry out their analysis within a controlled environment operated by the Data Safe Haven (Safe setting). Once data is distributed to researchers, controls on who accesses the data (Safe people) and on what is publicly released as results (Safe outputs) and even exactly what research the data is used for (Safe projects) are out of direct control of the Data Custodians.

Historically organisations such as ONS that have operated safe settings have offered only a limited set of statistical analysis tools to researchers accessing data in their environment. These have been largely adequate given the structure and size of administrative and social science datasets being analysed. By contrast, health datasets, consisting of electronic health records, images, and omics data types, are typically richer and larger. Analysis is more likely to require complex custom analysis algorithms which correspondingly require more substantial compute resources.

Until recently it has been a challenge to provide a safe setting able to support analysis at such a scale and with such diverse tooling. There has also been a history, dating back to the open sharing of the first human genome, of data distribution for such biological datasets to researchers to allow them to carry out analysis on their own computer systems. As a result, with the rise of the generation of omics data from consented individuals, this data distribution model has become codified as the process for managed data sharing, exemplified by the archives database of Genotypes and Phenotypes (dbGaP) and European Genome-phenome Archive (EGA).

Critical to the success of the proposed TRE-based approach will be the user experience of the researcher. There will be a cultural reluctance to change; this will certainly impose additional controls on research activities. This has been the experience from some communities, that have moved towards this model.

It will therefore be essential for the UK Health Data Research Alliance TRE workstream to continue to engage the community to develop best practice for the full range of user requirements and experience. It will also be important to communicate the other benefits that will accompany a move away from the data release model such as improved data access request turnaround times and the potential for less stringent data minimisation requirements and therefore more potential for hypothesis-generating or agnostic analysis.

[13] Burton, P. R. *et al.* Data Safe Havens in health research and healthcare. *Bioinformatics* **31,** 3241–3248 (2015). https://europepmc.org/article/MED/26112289

UK Health Data Research Alliance

The case for TREs providing access to health data through safe settings

Recent events and developments have made the provision of safe settings for health data analysis both a desirable and technically practical alternative to data distribution.

Firstly, previous initiatives that have not effectively engaged or consulted on the rationale for data access have resulted in a lack of trust in the phrase 'data sharing'. It has become associated with the risk of jigsaw reidentification through the distribution of data to unknown third parties for unknowable types of analysis and potential unknown further distribution and linkage. All approaches to the de-identification of datasets that contain individual patient level data are limited and require controls enabled by TREs.

Secondly, the adoption of General Data Protection Regulation (GDPR) has made researchers and their organisations subject to serious financial consequences of failing to adequately protect personal health data distributed to them and hosted on computer systems they are responsible for. For many organisations, in particular universities and NHS trusts, it has becoming increasingly challenging to operate computing environments with the required level of security. This is particularly the case when large scale high-performance computing (HPC) is required to support the research community. Similarly, data custodians distributing data have become more risk averse because of potential shared responsibility with receiving organisations for any breach under GDPR.

Thirdly, for large datasets such as images and genomes, data distribution is both inefficient and costly. It results in funders directly or indirectly supporting the costs of storage of multiple copies of large datasets and the associated network costs for multiple large data transfers, when each copy may only be used for a limited period.

Fourthly, over the last 5-10 years the evolution of computer systems have made it practical for researchers to bring complex analysis pipelines to data held on centralised systems[14] and for these systems to be able to support both cost efficient and dynamic scalability of compute[15] for analysis and integral data security[16].

Finally, there are examples of operational systems that provide a safe setting as the only mode of access and that are being used at scale for these classes of health data. The SAIL DataBank has operated for 12 years with a 'no data leaves' or 'reading library' approach to data access. Using Swansea University's UK Secure e-Research Platform (UKSeRP) to deliver remote access to population-scaled linked data resources from over 400 partner organisations, UKSeRP's fully featured high powered analytical environment has supported hundreds of projects conducted by academics from across the UK over that time[17]. Similarly, Scotland has implemented a system of federated safe havens with secure analytics platforms, which are safe settings, as described in their Safe Havens Charter[18]. The more recent Genomics England Research Environment (GERE) is also a safe setting and, with >20Pb of genome data from the 100,000 genomes

[14] Through Virtual machines or lightweight containers such as Docker.

[15] On premise High Performance Compute (HPC), private or public cloud services

[16] Data encryption at rest under user control via Key Management Infrastructure.

[17] https://saildatabank.com/wp-content/uploads/SAIL_10_year_anniversary_brochure.pdf

[18] https://www.gov.scot/publications/charter-safe-havens-scotland-handling-unconsented-data-national-health-service-patient-records-support-research-statistics/pages/4/

project, operates at a much greater scale. It has >2,000 researchers onboarded to carry out analysis with a range of tools and a full HPC environment (see figure 1).

GERE has the advantages of being set up 1) with explicit research consent from each patient participant; 2) with substantial public and patient engagement and oversight and 3) with completely unequivocal published[19] and public statements that individual level data will not be distributed but will remain within the research environment. As such it has achieved a high level of trust despite dealing with individual genomes, a new and sensitive class of personal health data[20].

More recently, UKSeRP has been available as a private cloud offering to third parties wishing to take advantage of UKSeRP's "ready to go" platform to secure and provide access to their own data, under their own governance [21] [22]. As of the 1st January 2020 there were 25 UK-based organisations with UKSeRP tenancies in the UK, including Dementias Platform UK, Avon Longitudinal Study of Parents and Children (ALSPAC), and UK-CRIS (Clinical Record Interactive Search), with an increasing number of installations internationally.

It is apparent from general patient and public engagement work that there is much greater comfort with data being accessed through a safe setting than data distribution.[23] Where research access is via a safe setting with appropriate patient /public oversight of research activities an opt-out consent model may also be considered sufficient even where data is sensitive, such as the founder CRIS system providing a research environment with Natural Language Processing (NLP) tools on a de-identified copy of the EHR records of the South London and Maudsley NHS Foundation Trust (SLaM)[24]. Even during past media-fuelled public concern about data sharing, the SAIL Databank with its robust proven and robust approach to data curation and access remained uncriticised[25].

Requirements for a Trusted Research Environment

This paper proposes that the Alliance adopts a "safe setting" approach as part of the implementation of proportionate governance based on the "Five Safes".

Alliance members already implement a subset of these, however for the Alliance to maximise its potential, a common agreed specification will simplify processes for researchers lowering barriers to access to

[19] Turnbull, C. et al. The 100 000 Genomes Project: bringing whole genome sequencing to the NHS. BMJ 361, k1687 (2018); Genomics England Protocol https://www.genomicsengland.co.uk/library-and-resources/

[20] https://medconfidential.org/for-patients/loopholes/

[21] https://farrinstitute.org/wp-content/uploads/2018/03/UKSeRP_Case_Study.pdf

[22] https://saildatabank.com/wp-content/uploads/UKSeRP_Brochure_v1.5.pdf

[23] Great North Care Record 2018 Base: 824 North East representative

[24] e.g. CRIS system, allowing NLP over de-identified mental health records https://www.slam.nhs.uk/research/cris/

[25] Lyons, Ford and Jones, Care.data: why are Scotland and Wales doing it differently? https://www.bmj.com/content/348/bmj.g1702/rr/687637

multiple TREs. Over time, a common TRE specification combined with the adoption of common health data standards (see Appendix A) will facilitate federated analysis across multiple TREs.

Ensuring public trust is maintained across multiple TREs implemented in different ways and operated by different organisations will require both the adoption of a common TRE specification and independent accreditation and auditing (see accreditation below).

Most Alliance members manage health service data for which access through the safe setting model is most appropriate. A subset of members manage access to research cohorts where rules of data access may be significantly different. This is due to participants in research cohorts being volunteers who have consented to data access rules that were ethically agreed at the outset. For example, the consent of one of the largest UK research cohorts, UK Biobank allows data distribution to approved researchers ("safe people") who have an approved research plan ("safe projects"). For these members setting up a TRE instance based on the safe setting model need only be one method of data access. However, to maintain the integrity of the UK HDR Alliance TRE system, such a TRE instance would need to be completely isolated from systems that provide data distribution.

Safe people

Individuals allowed access to TREs should be researchers[26] able to demonstrate appropriate credentials. They are likely to be paid by research organisations that are prepared to take responsibility for their actions and vouch for each individual. However, the approach must not inadvertently constrain access for researchers from non-standard backgrounds. Researchers would be required to sign legally binding terms of use including:

- not trying to re-identify individuals[27]
- immediately reporting any security weakness found when using the system and not attempting to exploit it
- not sharing their login credentials with any other individual
- informing the TRE if they are changing institutions before they have done so

They would also be required to carry out information governance training and, potentially, training specific to the TRE and/or datasets.

Accredited TREs will require systems to track individuals and organisations, the status of their on-boarding progress. With a common definition for safe people across the Alliance, it would be possible to setup something like the "Approved Researcher Scheme" used by ONS, where researchers only have to be approved once to access multiple TREs. This is also the approach that is being developed through the Global Alliance for Genomics and Health (GA4GH) with their passport and visa standards[28].

[26] Researchers could be from academia, NHS and from industry. Any accreditation of research must cover all these communities.

[27] Data Protection Act 2018 Section 171 - Re-identification of de-identified personal data

[28] https://www.ga4gh.org/news/ga4gh-passports-and-the-authorization-and-authentication-infrastructure/

The Health Data Research Innovation Gateway (the 'Gateway') has the potential to provide the technical implementation to reduce burden on researchers, and their parent organisations, wanting to access multiple Alliance TREs. This would also simplify implementation for each TRE. This would be either through the Alliance managing identity of researchers itself or relying on third party national or international research directories that have been proposed should they implement appropriate review and management processes.

Safe projects

Despite the privacy protections offered by TREs, it remains essential to ensure that the use of data is appropriate and has the potential for public benefit. The TRE must also have the functionality to enable this use to be audited to ensure compliance. Alliance members already have systems in place to review proposed research projects, and will be guided by the streamlining work of the "Promoting participation and improving access" workstream which will involve representatives of patients and public, guided by the "Engaging and involving practitioners, patients and the public" workstream (see Appendix A).

An issue that patients, public and cohort research participants frequently raise in focus groups is the frequency of feedback or transparency around research that is being carried out on health data about them. Its therefore proposed that TREs require lay summaries to be provided as part of the project approval process with these being made public on approval. TREs should also implement systems to allow access by researchers of data held by the TREs to be linked back to projects so that research activity can be made transparent and reported back to participants if demanded.

Safe setting

As detailed above, there are multiple existing platforms providing research access to health data through the implementation of a safe setting. While one side of operating a safe setting is the need to ensure public trust through security and transparency, the other side is the need to ensure it is engineered to be as easy to use for research as possible.

At minimum a safe setting needs to implement:

- A system to hold data securely such that individual level data cannot be exported. For transparency the security design and implementations should be independently audited with reports reviewed by a patient/public oversight group and made public (see accreditation below).
- Systems to allow secure remote access by researchers to carry out analysis with the ability to keep track of researcher activity (to ensure compliance with "safe projects") and that ensures accounts cannot be shared (to ensure compliance with "safe people").
- A research environment containing a set of tools to allow data to be analysed.

Because the types of analysis that researchers wish to carry out will go beyond that which can be provided by standard statistical packages (e.g. SPSS), it must also be possible for researchers to bring their algorithms into the safe setting. This would require 'air lock' capability to ensure that imported tools are scanned to ensure that they will not compromise the security and integrity of the TRE and, in particular, do not facilitate the export of record level data.

Similarly, because researchers may wish to analyse data in the safe setting along with data held outside, the safe setting must provide mechanisms to support importing and linkage of data. As with tools, this will require 'air lock' capability to allow for the secure importing of user supplied data. This may require tools to ensure that this data does not compromise the environment nor enhance the risk of re-identification beyond that assessed at the time the researcher's request for access was approved.

TREs must implement a barrier between the safe setting environment and the outside world to control data and software import or export, referred to here as the 'air lock'. Processes and systems are required for export of summary data (Safe output below) and to support data or software import. In both cases it will be necessary to implement systems able to scan data files, such as for viruses hidden within software packages and for identifiable data that should not be imported into a TRE.

In practice, such barriers mean that researchers cannot operate inside the safe setting as they do on their own computer systems. While they will be able to access the safe setting remotely via a Virtual Desktop Interface to carry out their research, they will not be able to connect from the safe setting to the outside except via the 'air lock'. This means they will not be able to access external websites, unless these are whitelisted with appropriate security to prevent data export such as through proxies etc. It also means they cannot connect directly to software distributions such as GitHub.

To facilitate the development and configuration of software prior to import into the safe setting, it is likely to be beneficial for TREs to provide researchers with a separate test environment that behaves like the safe setting but is accessible from the internet. Such a test environment would need to have the data frameworks identical to those within the safe setting so processing of data formats and communication to APIs could be tested, but the data itself would need be synthetic[29]. Since software frequently assumes the existence of the internet and contains embedded external file requests etc. it is also sensible for this test environment to have a 'no-internet' mode so it can be checked that the software still runs in the absence of internet connectivity.

TRE safe settings will be multi-user environments, in most cases with the ability to run algorithms on high performance computing systems, either as on-premise native HPC, private cloud or public cloud (see safe computing below). In all cases systems to manage the competing demands of many researchers are essential, meaning imported software must be able to work with such systems. Similarly, regarding actual software import, software could be packaged in different ways including containers (such as Docker) to full virtual machines. Across the Alliance it may be appropriate for different TREs to consolidate on supporting one or a limited number of workflow management and software packaging solutions to reduce the complexity for researchers wanting to run their software in different TREs.

[29] Synthetic data is artificially generated to replicate the statistical components of real-world data but doesn't contain any identifiable information.

UK Health Data Research Alliance

Safe computing – an extension of Safe setting

Since the ONS definition was developed, a new issue has become important that is not explicitly covered by the "Five Safes" that needs to be addressed to build public trust. This is the outsourcing of provision of computing infrastructure for all or part of a safe setting to third parties through partnerships with commercial organisations or use of public cloud computing providers.

Previously, safe settings have been almost exclusively provisioned through "on-premise" computer hardware where physical security of equipment, network security, software maintenance etc. is the responsibility of the data custodian. Such systems can be configured as "private cloud" to support the use of software distributed as virtual machines and containers. However, use of third-party computing resources such as public cloud offers many potential advantages for TRE providers. This includes dynamic scalability of compute to enable short periods of intensive computation such as for AI training. Outsourcing layers of the hardware and software stack which have become commodities to cloud providers brings other potential benefits due to their greater capacity to engineer scalable platforms and implement robust security.

In order to build public trust, use of private sector computing infrastructure to provide a safe setting must be done in such a way that none of the hardware and software layers outsourced make it possible for the third-party provider to access any of the individual health data. This needs to be enabled through security design and engineering as well as contractual arrangements with the third-party provider to minimise the risk of a data security breach. It is accepted by some cloud providers that a security design that ensures they have no data access is a critical requirement to many organisations. Technical papers have been published about how to engineer this level of security where cloud provider administrators have no ability to access any customer data[30,31].

The security engineering and design required to make this possible is complex and involves encryption on rest of all health data and encryption key management infrastructure configuration such that only the data custodian controls the keys. It is proposed that TREs using public cloud should be engineered in this way and would be regarded as operating a "safe setting" that implements "safe computing".

Explaining this complex engineering and design in ways that data custodians, researchers and members of the public can understand and engenders trust represents a challenge that needs further consideration.

---

[30] AWS: https://d1.awsstatic.com/whitepapers/using_aws_context_nhs_cloud_security_guidance.pdf; https://d1.awsstatic.com/whitepapers/Security/AWS_Security_Best_Practices.pdf; https://d0.awsstatic.com/whitepapers/AWS_Securing_Data_at_Rest_with_Encryption.pdf

[31] Google: https://services.google.com/fh/files/misc/handling_healthcare_data_uk.pdf, https://cloud.google.com/solutions/setting-up-a-hipaa-aligned-project

UK Health Data Research Alliance

Safe data

As well as all the controls outlined above, TREs should also ensure that the data accessible to researchers within the safe setting is as non-disclosive as possible and in line with GDPR requirements. This means ensuring processes to import data into the safe setting carry out de-identification to ensure individuals cannot be directly identified and, where possible, anonymisation of fields where this will not impact research analysis but will reduce the risk of accidental re-identification of individuals. Requirements to implement this for TREs will also be guided by the "Data Standards and Quality" workstream.

There are significant limitations to the current approach to de-identification. These limitations are central to the need to move towards a TRE based model for future research and innovation. However, with the additional controls supported by the “Safes”, and using best-of-breed de-identification and encryption, it should be possible to explore whether the current conservative approaches to linkage and data minimisation can be adjusted therefore opening up new research opportunities and especially facilitating a more effective approach to hypothesis-generating research that has the potential for public benefit. Environments such as GERE which implements a safe setting with strong implementation of the other safes, allow researchers to analyse across the entire dataset, facilitating broad investigations of genome/phenome relationships, multimorbidity effects etc.

Safe outputs

As outlined in Safe Setting, TREs must implement a barrier (or “air lock”) between the safe setting environment and the outside world to prevent unauthorised data export. TREs must implement processes and systems to allow approved data to cross this barrier. Systems require functionality to track requests and decisions, supporting cycles of rejection and revision.

Current approaches to review requests to export summary data are based on manual review with typically final release being governed by an oversight committee. This is a potential bottleneck and could be one of the factors that will negatively impact the user experience of a researcher moving from a locally hosted analysis environment to a TRE. Work is therefore required to explore approaches to automation or partial automation where the risks of disclosure can be adequately controlled. This will also be beneficial for the data custodians to make management more scalable and sustainable. There are opportunities to establish a network of TRE airlock managers to share expertise and develop consistent approaches to definition of safe summary outputs.

UK Health Data Research Alliance

Safe return - Extending the Trusted Research Environment definition

While health data held within TREs is always de-identified to guard against accidental re-identification of individuals by researchers (safe data), there are differences as to whether it is technically possible or consented to send individual analysis results back to the clinical setting that originated the data and where identities are known. This will be for individual clinical care purposes and invitations to participate in trials and other research projects.

For example, in the Genomics England case (see figure 1), there is ethical approval and patient consent to pass analysis results for an individual generated in the research environment safe setting back to the clinical setting for re-identification, evaluation and return for clinical care. These 'outputs' supplement the results already generated by clinical analysis pipelines in the clinical setting. Given that the clinical analysis pipelines only produce diagnostic results in 20-25% of cases, there is considerable clinical value in additional individual diagnoses being proposed for undiagnosed patients from the research side. Making this possible requires completely robust and certified data paths for individuals to ensure that a result obtained in a research environment is always perfectly mapped back to that individual's clinical record.

On the other hand, in the case of research cohorts such as UK Biobank, which does contain clinical health data for each individual, there is no consent for return of results to individuals so a TRE based on UK Biobank data would not allow this.

For TREs where return of results is possible, there can be multiple benefits. It may be only a part of the research activity carried out within a TRE, but supporting this option has the potential of increasing the convergence of research and clinical care, bringing researchers and clinicians closer together. It may also provide an additional incentive to clinicians to ensure the clinical data they record is as complete as possible if research use could result in additional clinical feedback.

Implementing such a return path in a way that ensures no reports are returned to the wrong individual requires significant technical implementation. It is therefore proposed that such requirements are captured by an enhanced "safe output", namely "safe return". TREs that implement "safe return" would be regarded as "TRE+ environments."

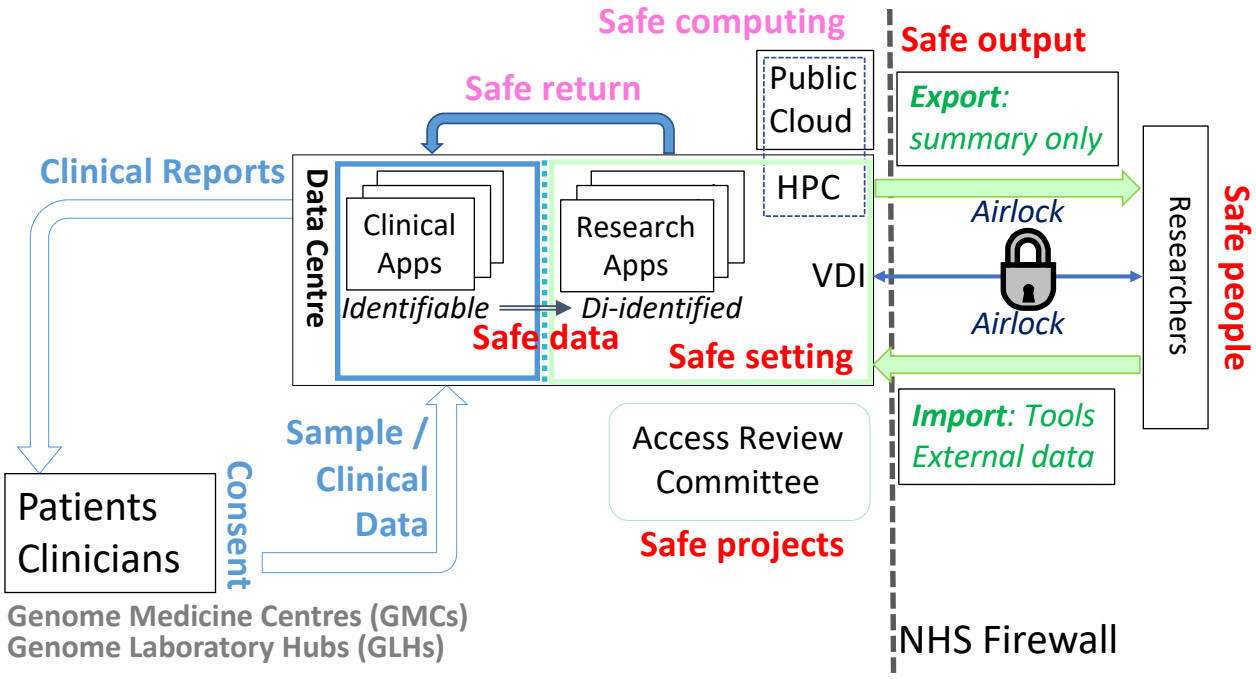

Figure 1: Schematic example of a TRE+ environment. The Genomics England TRE operates as a **Safe setting [green]** that approved researchers (**Safe people**) can only access via a virtual desktop interface (VDI). Only de-identified data is accessible by researchers (**Safe data**). Research is overseen by an Access Review Committee (**Safe projects**). Only summary data can be exported from the TRE through an Airlock and only after manual review (**Safe output**). Analysis of genome data requires High Performance Compute (HPC) resources, however scaling HPC to meet the needs of large numbers of researchers is challenging. One solution would be to also use public cloud resources, however this would need to meet the security requirements of **Safe computing**. In the Genomics England environment the data accessed by researchers is a de-identified version of real world data being analysed to produce reports for clinical care **[blue]**. This makes it possible to pass back research results that may be relevant for individual clinical care (**Safe return**).

Federating TREs and Data

It will be possible to undertake many projects within a single TRE or by running analysis separately on a number of different TREs and then combining the resultant summary outputs. Such an approach is already used where individual level data cannot be directly combined, such as meta-analysis of genome wide association studies. Similarly, there are approaches to allow TRE environments to provide programmatic interfaces that implement the summarisation rules of safe outputs, enabling researchers to carry out

limited types of federated queries remotely across multiple TREs. Implemented examples include the beacons network[32] developed by the Global Alliance for Genomics and Health[33].

However, there are classes of experiment that may benefit from a tighter form of data federation or the aggregation of datasets to allow linking within a single TRE or virtual TRE environment.

To achieve this, it will be necessary to build upon the accreditation of TRE to establish a common model of trust across a federation of TREs with common user identity. Not all TREs will implement the requirements in this paper to the same level and therefore a potential asymmetric hierarchical model of data flow will be required, ensuring that aggregated data is accessed in the safest environment across the federation. This could build on experience from the security and defence sector[34] and the implementation of the Digital Economy Act which categorises two levels of processing covering preparation (including linkage) and provision, and provision only.[35]

The move towards federated analytics and distributed Machine Learning will need to underpin the characteristics of such a federation of TREs. However, it is likely to be some time before this approach can be fully realised and in the interim the default approach will need to be through data federation utilising techniques such as from the HDR UK Sprint Project – "Graph-Based Data Federation for Healthcare Data Science".[36]

This would then support the controlled transient aggregation and linkage of data across TREs to support data federation. This will require auditability of data, secure transfer mechanism and controlled destruction of the transient data in line with the requirements of ISO 27001.

Accreditation of TREs

This paper has outlined a proposed approach for TREs. However, for such an approach to be implemented, it will need a widely accepted approach to accreditation that meets the requirements of data custodians, regulatory bodies and patients and public representatives. At the same time, this will need to be achieved without an excessive additional burden on the TRE who may already be undertaking multiple accreditation. There are particular challenges such as the highly technical oversight of security design and implementation required to support safe computing on public cloud.

Consensus is needed on who will provide accreditation, what is within scope (characteristics, processes, personnel, etc.) and how this will build on existing frameworks such as ISO 27001[37], the NHS Digital Data

[32] https://beacon-network.org/

[33] https://www.ga4gh.org/

[34] E.g. https://www.commoncriteriaportal.org/

[35] https://www.statisticsauthority.gov.uk/about-the-authority/better-useofdata-statistics-and-research/betterdataaccess-research/better-use-of-data/list-of-accredited-processors-under-the-research-strand-of-the-digital-economy-act/

[36] https://www.hdruk.ac.uk/projects/graph-based-data-federation-for-healthcare-data-science/

[37] https://www.iso.org/isoiec-27001-information-security.html

Security and Protection Toolkit (DSPT)[38], and the UK Statistics Authority accreditation process that supports implementation of the Digital Economy Act (DEA)[39].

Much work remains to be done in this area and will need focused discussion from the TRE workstream to bring recommendations forward.

Communications, Engagement and Involvement

Public Trust

Central to the adoption of the 'Five Safes' approach is the need to earn, build and sustain public trust. There is valid public concern over the control of data that is made available for research through data release and on the limits of de-identification. Communications, engagement and involvement with the public must be central to the Alliance's proposed approach to move towards data access via TREs, implement accreditation and support integration with the Health Data Research Innovation Gateway. This disruptive change must respond to key public questions and concerns such that potential benefits from health data research are achieved whilst protecting privacy.

Communications

The workstream will need to ensure that there are clear and tailored communications for all stakeholders. This will need to address the questions and interests that are specific to the audience and the benefits that will accrue from a robust TRE model for research on health data.

Public and Patients – Discussion needs to cover how the "Five Safes" augment de-identification, how this approach ensures that data remains under the control of the data custodians and not passed to private companies with the risk of use for unapproved purposes and that this will facilitate the UK as being the place for safe and secure health data research. Specifically, this will also need to cover the controls in place on data held in public cloud that ensures that the data cannot be accessed by the hosting organisation. Transparency regarding use and benefit of the data will remain paramount.

Data Custodians – It needs to be assured that this will provide security for the data they manage and that it remains within their controls and meets GDPR and Common Law Duty of Confidence requirements. They will need to be able to support the approach to accreditation with confidence that this builds upon the data management requirements of ISO 27001 and DSPT. This enhanced level of control, compared to the data

[38] https://www.dsptoolkit.nhs.uk/

[39] https://www.statisticsauthority.gov.uk/about-the-authority/better-useofdata-statistics-and-research/betterdataaccess-research/better-access-to-data-for-research-information-for-processors/

release model, should then encourage the data custodian to adjust their data access management processes to facilitate quicker access to richer data given the reduced risk of inappropriate disclosure or use.

Researchers and Innovators – It will need to be assured that their user experience has been considered in the development of the requirements and that there is focus on a first-class research and innovation experience. The workstream will need to address the concerns from this community that a TRE approach will impact research efficiency. The communications will need to highlight other longer-term benefits such as more rapid access to data and improved opportunities for linking data that has until now been restricted due to the data custodians risk positions. A TRE model should also provide a more cost effective approach to high scale compute and storage as environments move to a hybrid cloud model and the cloud providers' commercial models are refined to address some of the current issues for example around the costs of data egress.[40]

TRE Service Providers – The proposed approach offers service providers significant opportunities. But with these opportunities comes great responsibility to develop technical and governance systems that can protect privacy whilst providing a world class analytical experience.

Funders – TREs offer funders of research a range of potential benefits. These include more efficient research through improved utilisation of storage and compute resources; a proportionate approach to data access requests based on all the five safes; and audit trails on provenance of research outputs and data manipulation. This supports both transparency and replicability. These features may also benefit regulators.

[40] https://www.gartner.com/en/documents/3939969/the-art-of-taming-data-egress-charges-in-hybrid-and-publ

UK Health Data Research Alliance

Consultation Questions

This document outlines the position on the characteristics required for TREs to support safe and ethical research using health data assets from data custodians within the Alliance. We are seeking further input to guide the development of this approach, to understand the level of support and to identify particular concerns.

The key areas that have been identified as requiring further development through the work of the Alliance TRE workstream are:

- **Safe people:** How do we accredit researchers in a way that can support academia, NHS and industry (large and small, UK-based and beyond) to achieve the vision that every health and care interaction and research endeavour will be enhanced by access to large scale data and advanced analytics? What can we learn from the ONS Approved Researcher Scheme and other similar approaches?[41]
- **Safe setting:** How do we ensure that a move to including TREs based on hybrid/public cloud ensures that data remains safe and cannot be accessed by the hosting technology companies? How is this best explained to data custodians, researchers and the public and not just security experts?
- **Safe outputs:** How do we achieve a scalable and trustworthy approach to safe outputs? Should Safe Return be considered a separate 'safe' and represent a TRE + model?
- **User centred functionality:** What is the minimum set of analytical tools and functionality that a TRE must make available for researchers?
- **Accreditation:** What approaches should be considered to assess that TREs meet the characteristics required? Who should be the accrediting authority – and how should this be funded?
- **Public trust:** How can we engage patients and the public to demonstrate the benefits of health data research and build public trust around the use of trusted research environments for research and innovation at scale?'

Please provide your feedback and comments via the web-based form [https://ukhealthdata.org/trusted-research-environments-green-paper-consultation/].

[41] https://www.ons.gov.uk/aboutus/whatwedo/statistics/requestingstatistics/approvedresearcherscheme

UK Health Data Research Alliance

Appendix A: Dependencies with other Alliance workstreams

This workstream is one of five being led by the Alliance[42]. Defining various aspects and processes for providing TREs are therefore led by these other streams and are not discussed in detail here:

The processes for engaging with patients and public which feed into providing oversight of "safe projects" etc. is the responsibility of the "**Engaging and involving practitioners, patients and the public**" workstream.

Streamlining and standardising the process of approving projects "safe projects" is the responsibility of the "**Promoting participation and improving access**" workstream.

The data and quality standards adopted within TREs are the responsibility of the "**Data Standards and Quality**" workstream which should feed into operational definitions for "safe data".

The requirements for TREs to provide metadata descriptions to improve discoverability of data resources hosted by different TREs is the responsibility of the "**Supporting Health Data Research Innovation Gateway development and launch**" workstream.

In addition, there are Health Data Research Hub workstreams related to operating TREs, such as financing models for sustainability. The financing workstream will propose models for covering the costs of operating TRE environments for academic and commercial use, such as through a mixture of grants to provide baseline operation for academics and direct charges, such as for short term cloud compute.

[42] https://ukhealthdata.org/work/