Trusted Research Environment: Architecture and Implementation, a case study

Informatics programme

Internal report OR/25/065

BRITISH GEOLOGICAL SURVEY

INFORMATICS PROGRAMME

INTERNAL REPORT OR/25/065

The National Grid and other Ordnance Survey data © Crown Copyright and database rights 2025. OS AC0000824781.

Keywords

TRE; AI; ML; Cloud computing

Bibliographical reference

Kingdon, A., Kadochnikova, A., Joseph, T. 2025. Trusted Research Environment: Architecture and Implementation, a case study. *British Geological Survey Internal Report*, OR/25/065. 26pp.

Copyright in materials derived from the British Geological Survey's work is owned by UK Research and Innovation (UKRI) and/or the authority that commissioned the work. You may not copy or adapt this publication without first obtaining permission. Contact the BGS Intellectual Property Rights Section, British Geological Survey, Keyworth, email ipr@bgs.ac.uk. You may quote extracts of a reasonable length without prior permission, provided a full acknowledgement is given of the source of the extract.

Maps and diagrams in this book use topography based on Ordnance Survey mapping.

Trusted Research Environment: Architecture and Implementation, a case study

A Kingdon, A Kadochnikova, T Joseph

© UKRI 2025. All rights reserved

Keyworth, Nottingham British Geological Survey 2025

BRITISH GEOLOGICAL SURVEY

The full range of our publications is available from the BGS shop at Nottingham and Cardiff (Welsh publications only). Shop online at https://shop.bgs.ac.uk/

The London Information Office also maintains a reference collection of BGS publications, including fossils, for consultation.

We publish an annual catalogue of our maps and other publications; this catalogue is available online or from the BGS shop.

The British Geological Survey carries out the geological survey of Great Britain and Northern Ireland (the latter as an agency service for the government of Northern Ireland), and of the surrounding continental shelf, as well as basic research projects. It also undertakes programmes of technical aid in geology in developing countries.

The British Geological Survey is a component body of UK Research and Innovation.

British Geological Survey offices

Nicker Hill, Keyworth,
Nottingham NG12 5GG
Tel 0115 936 3100

BGS Central Enquiries Desk
Tel 0115 936 3143
email enquiries@bgs.ac.uk

BGS Sales
Tel 0115 936 3241
email sales@bgs.ac.uk

The Lyell Centre, Research Avenue South,
Edinburgh EH14 4AP
Tel 0131 667 1000
email scotsales@bgs.ac.uk

Natural History Museum, Cromwell Road,
London SW7 5BD
Tel 020 7589 4090
Tel 020 7942 5344/45
email bgslondonstaff@bgs.ac.uk

Cardiff University, Main Building, Park Place,
Cardiff CF10 3AT
Tel 029 2167 4280

Maclean Building, Crowmarsh Gifford,
Wallingford OX10 8BB
Tel 01491 838800

Geological Survey of Northern Ireland, 7th Floor,
Adelaide House, 39-49 Adelaide Street, Belfast, BT2 8FD
Tel 0289 038 8462
www2.bgs.ac.uk/gsni/

Natural Environment Research Council, Polaris House,
North Star Avenue, Swindon SN2 1EU
Tel 01793 411500 Fax 01793 411501
www.nerc.ac.uk

UK Research and Innovation, Polaris House,
Swindon SN2 1FL
Tel 01793 444000
www.ukri.org

Website: https://www.bgs.ac.uk
Shop online: https://shop.bgs.ac.uk/

Acknowledgements

This work was funded by a combination of activities within 2 inter-related UKRI / NERC projects funded from the Digital Research Infrastructure initiative

- BOOST: Environmental Data Service, Research Data Cloud Pilot Initiative UKRI Digital Research Infrastructure Phase 1b (BGS project code: NEE9442N)
- AMPLIFY: Environmental Data Service, UKRI Digital Research Infrastructure Phase 2 EDS (BGS project code: NEE9478N)

Contains OS data © Crown copyright and database rights 2025. OS AC0000824781'.

Contents

FIGURES

Summary

This report summarises a theoretical examination of the use of a Trusted Research Environment (TRE) for integrating static and dynamic environmental data sets. This work was funded through the UKRI Digital Research Infrastructure BOOST and AMPLIFY projects.

The project proposal was predicated that an environmental TRE would be implemented by the University of Manchester hosted Digital Solution infrastructure. For technical reasons beyond the scope of this report, this has not been implemented within the timescale of the AMPLIFY project.

Therefore, this report summarises the conceptual thinking for how a TRE might be implemented for a use case of provision of confidential water resource data. This requirement crosscuts data supplied by BGS, from the National Geoscience Data Centre (NGDC) and integrates it in combination with data supplied from a hypothetical and confidential water industry data source, and also a conceptual data source provided by an environmental regulator (such as the Environmental Agency for England or Scottish Environmental Protection Agency).

The project objective is to provide an understanding of how a hypothetical model incorporating both geospatial and sensor driven modelling data sources could be implemented within a specimen TRE to demonstrate how this could be developed from both an architecture and modelling standpoint.

1 Introduction

1.1 MANAGING SENSITIVE DATA

The combined environmental data centres that collectively constitute the NERC Environmental Data Service (EDS) are responsible for provision of environmental data derived from UKRI-funded scientific research. Some of these data centres, particularly the British Geological Survey (BGS) hosted National Geoscience Data Centres (NGDC) have other wider responsibilities. NGDC is the UK National Archive registered "place of deposit" for all industrially-derived geoscience data, a responsibility dating back to the Geological Survey Act, 1845. Data may be supplied under a number of legislative requirements with different standards of confidentiality imposed. Some dataset can be shared with researchers but only once the source data has been anonymised to prevent the confidential data being directly identified from the outputs. This report is a conceptual study on the applicability of using a Trusted Research Environment (TRE) to share scientifically valuable information without compromising its confidentiality.

There have been legislative requirements for geoscience data for the UK landmass to be collected to support the work of British Geological Survey (BGS) and its antecedents dating back to the 1845 Geological Survey Act. As national place of deposit for geoscience information, the National Geoscience Data Centre at the British Geological Survey hosts significant volumes of scientific data that have been collected by organisations other than the BGS. Today, NGDC hosts significant volumes of information which is treated as confidential. These data can be valuable for national good science but only in anonymised form. Such data cannot be supplied to external partners. For some more sensitive data even its exact location cannot be acknowledged publicly. However, the underlaying information could be of significant value to environmental researchers as it contains embedded data that is crucial to environmental understanding. Therefore, environmental researchers require a way to derive reliable but anonymised information for incorporation of these inferences.

As an example, there are particular issues around supply of water industry data which are location protected. Water industry data is an important component in understanding groundwater levels, a critical interest to understanding groundwater flooding. Another important parameter is water quality information. Such data are relevant to practical issues such as water supply, and environmental modelling for ecosystem health. Such information cannot be shared in ways that allows individual source data to be identifiable. Such confidential data can include point data, polygon data coverage or time-series sensor data inputs from sensor networks or observatories.

Trusted Research Environments are a concept that has emerged from the health domain. These are highly secure computing environments that provide remote access to otherwise confidential health data for use in research by approved researchers. The security and confidentiality requirements for health data mean that all summarised health data must be anonymised, with all metadata that can be used to identify patients removed, and upscaled so that the locations of individual patient data cannot be recognised.

Whilst these specific issues are limited to health-related data associated with individuals some of these issues are applicable to other disciplines. Trusted Research allow data to be input securely from diverse sources into a protected environment. This can only be accessed by verified environmental researchers with appropriate access permissions to interact with these data. The data can then be integrated, anonymised and upscaled so that insights can be extracted but at a resolution that doesn't endanger the integrity of the source data.

1.2 ENVIRONMENTAL TRE USE CASE

The original intention of the BGS use case for the TRE for environmental data was to address issues created by confidential water supply boreholes. Both boreholes that are directly involved in public water supply, and the network of observation boreholes used for monitoring the state of groundwater levels, are treated as confidential for security of supply reasons. In addition, the water companies provide commercially sensitive data (eg time-series) to BGS and regulators like the EA. These data are confidential and must be kept secure, but once data is anonymised and it can used for modelling.

For this example use case, a hypothetical scientist needs to utilise sets of static data from point locations which will be integrated with dynamic sensor data inputs in order to create a modelling output. This can then upscaled be into time-series spatial data outputs delivered at a resolution sufficient upscaled that the initial input data cannot be defined. This process is undertaken within a simulated Python notebook to highlight a likely implantation of a model within a TRE.

The second stage of this process is the development of a synthesised architecture within which the TRE could be created. This utilises a kubernetes setup for the models and system or intended to be implemented using commercial cloud implementation.

1.2.1 Presupposed TRE implementation of BGS use case

The original intention of the BGS use case that it would be implemented within an existing TRE. The University of Manchester Digital Solutions (UoMDS) team have been funded to create a TRE for integrating environmental data with confidential medical information to look at health effects of environmental issues.

The UoMDS team have had a series of issues with identifying the appropriate architecture upon which to implement their TRE. As their use case involved dealing with highly confidential medical data, their intention is to create a TRE that achieves the full SATRE standard for implementation. These issues mean that UoMDS have suffered significant delays in developing their TRE platform. These issues have now apparently been resolved but the completed TRE has not been delivered in a manner which allows the BGS use case to be implemented by the time of writing of this report. Discussion of these issues will be reported by UoMDS and beyond the scope of this report.

Therefore, it was decided as an alternative that BGS's use case be tested in a hypothetical study. This report summarises that response.

This report is structured into two sections. The first is focussed upon the technical architecture for implementing a TRE upon a commercial cloud serve using Kubernetes. The second simulates the actual internal modelling structure for resolving the proposed TRE use case.

2 Trusted Research Environment (TRE) implementation in commercial cloud infrastructure

A TRE, in BGS terms, is a controlled, secure platform that scientists can use to implement and run large data models where this environment complies with privacy, security and ethical standards as set out by BGS and NERC. This report presents a synthesised architectural design of the proposed TRE utilising a commercial cloud environment.

Key attributes:

Successful delivery of a TRE means that a clear set of requirements to be fulfilled:

1. **Secure Data Access:** The platform within the environment is closed off to external users. Scientists can work on their data models with the assurance that their data is secured.
2. **Privacy Protection:** The environment is protected by safeguards to ensure sensitive information isn't leaked.
3. **Controlled Computing Environment:** The environment itself is a controlled computing space, which could be either on-premises or cloud-based.
4. **Collaboration Support:** TRE provides support for multiple scientists to work on the same data model or several others.
5. **Compliance and Governance:** TRE ensures adherence to regulations like GDPR and other national laws. Data placed within the TRE will need to be signed off by BGS and NERC.

2.1.1 Schematic of proposed Trusted Research Environment

Figure 1 gives a high-level representation of how the proposed BGS TRE use case architecture that would need to be delivered

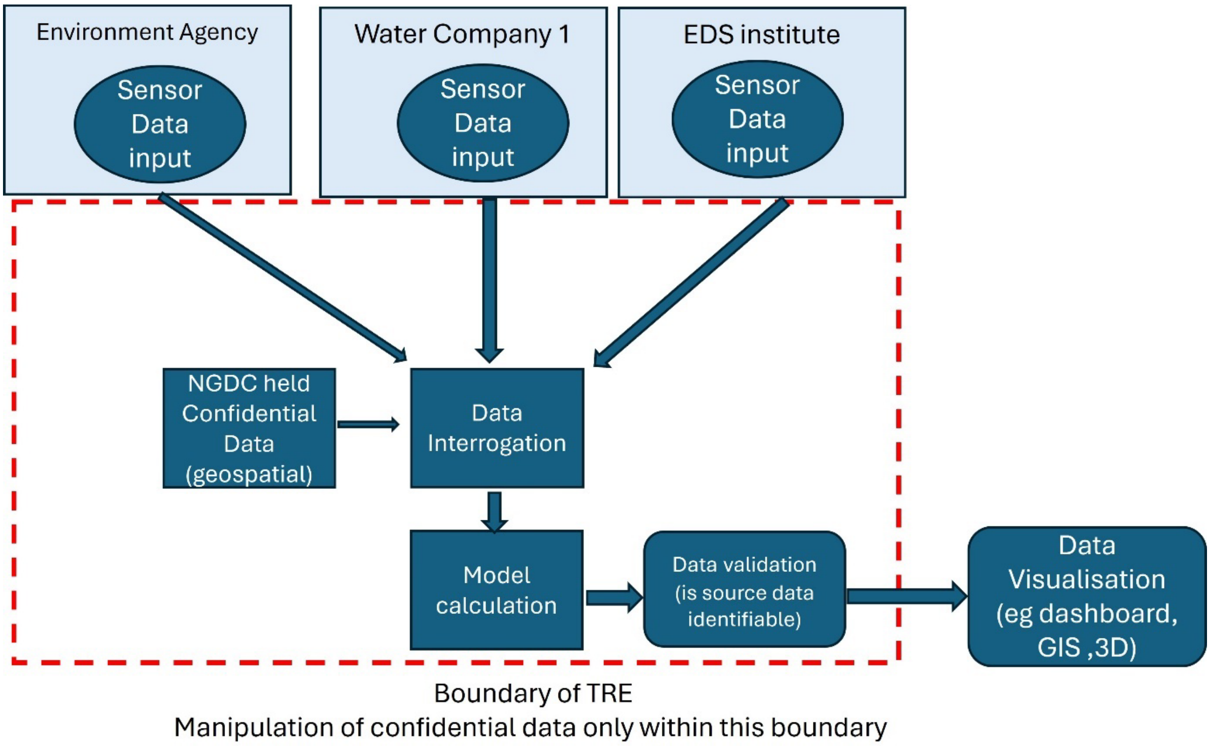
Environment Agency
Sensor Data input
Water Company 1
Sensor Data input
EDS institute
Sensor Data input
NGDC held Confidential Data (geospatial)
Data Interrogation
Model calculation
Data validation (is source data identifiable)
Data Visualisation (eg dashboard, GIS ,3D)
Boundary of TRE
Manipulation of confidential data only within this boundary

Figure 1: Visualisation of TRE workflow in BGS.

2.1.2 Explanation of requirements for BGS Trusted Research Environment use case

The BGS used case for a TRE for water industry data incorporates multiple confidential data inputs, including spatial and temporal data and modelling outputs

1. The diagram above shows three potentials points of origin for sensor data (each with different confidentiality requirements) to be input into the TRE. Such data will need to be ingested from external sources, using secure API feeds via API calls.
2. Sensor data (with different confidentiality states) will be ingested into the TRE so that they can be integrated into a process model. Environmental parameters are then calculated which synthesise parts of a geological system. Users have control over model behaviour, whose system specifications can be controlled by the scientists undertaking the modelling process.
3. Once model outputs are calculated, these data are then validated. These output data must be de-identified before they are exported outside the TRE and ingested to a data visualisation software such as GIS dashboard such that external users viewing the data are unaware of where the data originated from.

2.2 DETAILED REQUIREMENTS FOR TRE

The delays experienced by the UoMDS team (see above) meant there was no realistic prospect of delivering the functionality necessary for the BGS TRE use case as a fully operating pilot within the project lifespan. These delays made practically impossible to deliver the BGS TRE use case in any detail.

Therefore, there was no useful purpose served in undertaking a detailed requirements setting process for which there was no realistic prospect of delivery. Therefore, this document deals with a high-level conceptual implementation without getting distracted by the details of requirements for purely virtual capability.

3 Analysis of hosting options for BGS TRE use case

There are multiple computational platforms in which the TRE could be hypothetically established. Different methodologies for enabling this in a realistic architecture were considered, but the most flexible opportunity identified is a Kubernetes implementation which can be rapidly established on a commercial cloud platform. The specimen cloud provider was intended to be AWS, based on other work BGS was already undertaking. The familiarity with this platform led to the TRE use case being conceptually implemented within this environment.

Kubernetes is an open-source container orchestration system for automating software deployment, scaling, and management. It is platform independent, enabling systems to be created and trialled in one architecture which are then deployed on another architecture and installed within another environment, for example commercial cloud. Kubeflow is a family of utilities for setting up AI related codes for implementation through a Kubernetes installation.

3.1 SETTING UP AWS AS ZONE FOR TRE

3.1.1 Overview of Utilities necessary to deliver TRE

The following utilities are necessary for enabling implementation of a functional TRE in AWS using Kubernetes

- **BGS Virtual Private Network (VPN):** For controlled access to AWS
- **AWS Transit Gateway:** For secure networking within AWS
- **EKS with Kubeflow:** For machine learning workflows and data models
- **AWS IAM, VPC and Storage Services:** For security and data management

3.2 ARCHITECTURE OVERVIEW

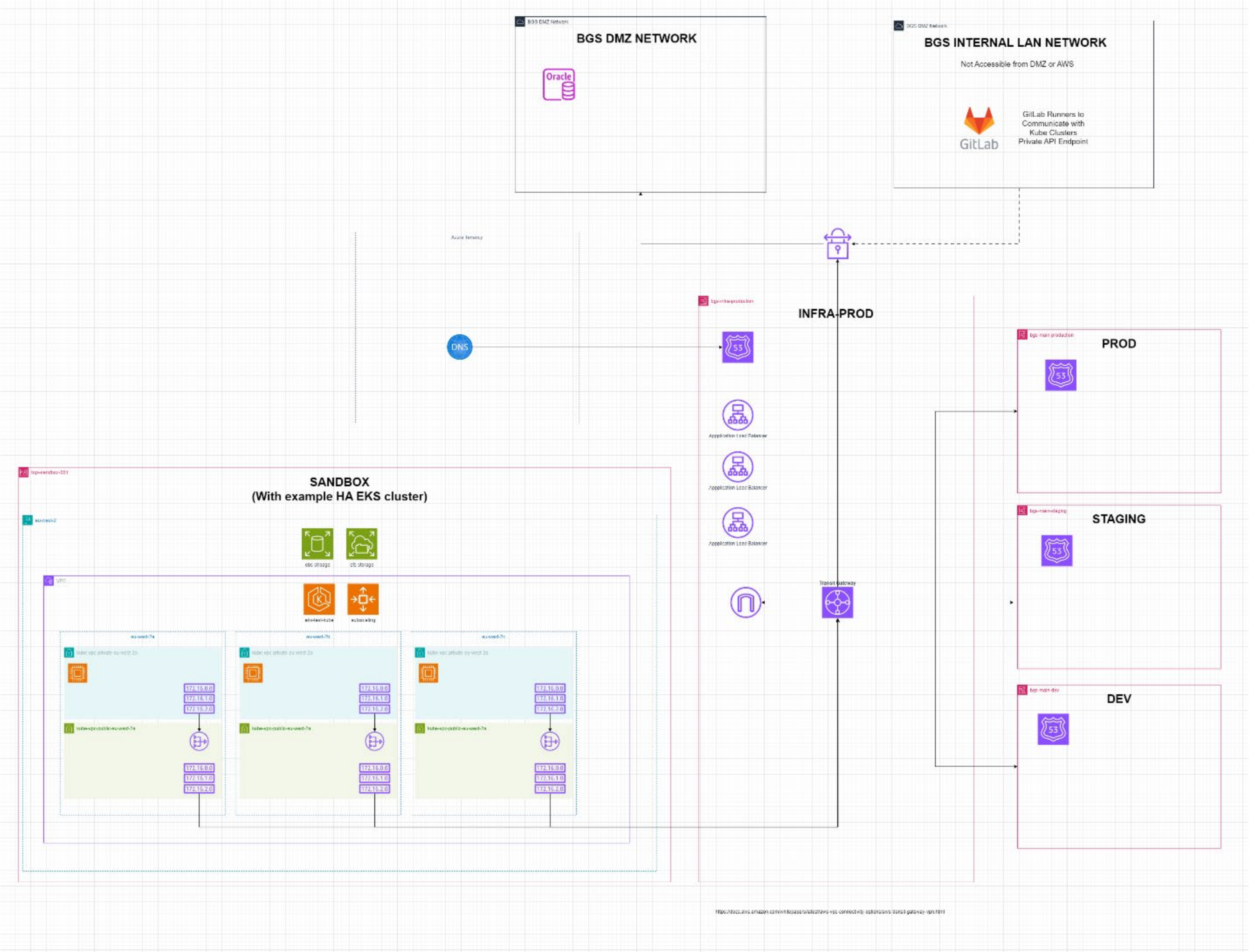
BGS DMZ Network
BGS DMZ NETWORK
Oracle
BGS INTERNAL LAN NETWORK
Not Accessible from DMZ or AWS
GitLab
GitLab Runners to Communicate with Kube Clusters Private API Endpoint
INFRA-PROD
DNS
Application Load Balancer
Application Load Balancer
Application Load Balancer
Transit gateway
PROD
bgs-main-production
STAGING
DEV
SANDBOX
(With example HA EKS cluster)
VPC
ebs storage
efs storage
autoscaling
172.16.0.0
172.16.1.0
172.16.2.0

Figure 2: AWS architecture diagram.

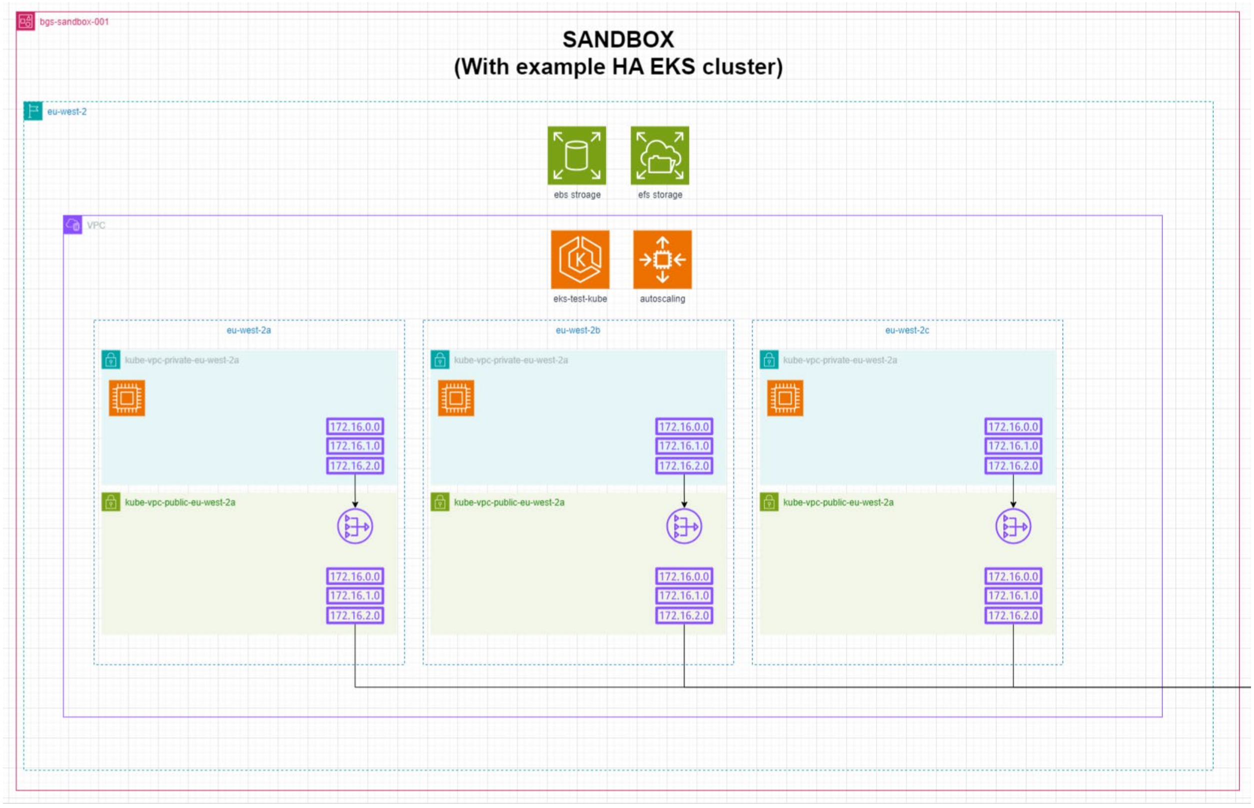
bgs-sandbox-001
SANDBOX
(With example HA EKS cluster)
eu-west-2
ebs stroage
efs storage
VPC
eks-test-kube
autoscaling
eu-west-2a
eu-west-2b
eu-west-2c
kube-vpc-private-eu-west-2a
kube-vpc-private-eu-west-2a
kube-vpc-private-eu-west-2a
172.16.0.0
172.16.1.0
172.16.2.0
kube-vpc-public-eu-west-2a
kube-vpc-public-eu-west-2a
kube-vpc-public-eu-west-2a
172.16.0.0
172.16.1.0
172.16.2.0

Figure 3: Detailed view of sandbox.

Figure 2: AWS architecture diagram shows the following capabilities that are inherent to solving these problems:

- **Infra-Prod:** This is the AWS environment that contains the AWS Transit Gateway and Internet Gateway, as well as Load Balancers and the DNS service for AWS, Route 53.
- **Lock (VPN):** The lock above Infra-Prod symbolises the BGS VPN, used to connect to the AWS environment as well as access to BGS LAN and DMZ networks.
- **Dev, Staging and Prod:** These are the individual environments accessible by scientists to run their notebooks and data models.
- **Sandbox:** This is the example of an environment in AWS, accessible via the Transit Gateway (see Figure 3). It contains the following:
    - ***Storage solutions:*** To provide High Availability data storage in AWS, there are two solutions that can be used: Elastic Block Storage (EBS) and Elastic File System (EFS).
    - **Elastic Kubernetes Service (*EKS):*** AWS EKS is used for machine learning workloads from large data models.
    - ***Auto-scaling:*** AWS dynamically adjusts resources based on demand, cost efficiency and performance.
    - ***Private and public VPCs:*** Each region selected in AWS has three availability zones within it (eu-west-2a/b/c). Each zone is separated into private and public VPC (Virtual Private Cloud) networks. These VPCs further isolates and secures data inside AWS. Access to the private VPC is only through the Public VPC, which is accessed via the Transit Gateway.

- **BGS LAN and DMZ Network:** These are the two BGS network zones that are accessible behind the BGS VPN. The LAN network contains GitLab and two Kubernetes environments: internal-dev and internal-prod. The DMZ network contains the other two Kubernetes environments: dmz-staging and dmz-prod. These two network zones are accessible by AWS only if the BGS VPN is used.

3.3 DETAILED ARCHITECTURE OVERVIEW

3.3.1 BGS VPN

The VPN used by BGS is accessed via Big-IP Edge Client. This VPN provides scientists secure access to the BGS LAN, DMZ and AWS environments. Big-IP Edge Client is a secure remote access solution provided by F5 Networks. It enables users to connect securely to enterprise networks from remote locations, ensuring encrypted communication and policy enforcement.

Working arrangements:

1. Scientists authenticate using their BGS credentials through the Big-IP Edge Client.
2. The VPN establishes an encrypted tunnel between the scientist’s device and AWS, LAN and DMZ networks.
3. Once connected, traffic is securely routed through AWS Transit Gateway to grant access to TRE services.
4. Security measures include multi-factor authentication (MFA) and network access control policies to prevent unauthorized access.

3.3.2 AWS Transit Gateway

AWS Transit Gateway is a centralised networking service that securely connects multiple VPCs, on-premises networks, and AWS services. Within the TRE, the Transit Gateway plays a crucial role by:

1. Connecting research environments (VPCs) securely and efficiently.
2. Routing VPN traffic to Amazon EKS (Kubeflow), data storage (S3, EBS, EFS), and security monitoring tools.
3. Reducing complexity by consolidating multiple VPC peering connections.
4. Enhancing security by enforcing routing policies and network segmentation.
5. Gateway allows access to different environments within AWS: Dev, Staging and Prod.
6. It also provides access to the Internet Gateway, providing communication to the internet for all AWS resources.

3.3.3 Kubeflow on Amazon EKS

Kubeflow is an open-source machine learning (ML) toolkit designed for Kubernetes. It provides a comprehensive framework for deploying, managing and scaling machine learning workflows in a cloud-native environment. Amazon EKS (Elastic Kubernetes Service) is used to host Kubeflow, ensuring scalability, security, and ease of deployment.

Kubeflow is structured into multiple interconnected components that help streamline the entire machine learning lifecycle:

- **Kubeflow Pipelines**: A robust workflow orchestration tool that enables users to build, deploy, and manage end-to-end ML workflows.
- **Jupyter Notebooks**: Provides an interactive development environment where researchers can write, test, and refine ML models.
- **TensorFlow, PyTorch, and Other Frameworks**: Kubeflow supports popular ML frameworks, making it adaptable for various research needs.

3.3.4 Role of Kubeflow within TRE

- **Security & Access Control**: Kubeflow integrates with AWS IAM to enforce role-based access control, ensuring that only authorized researchers can run ML workloads.
- **Scalability**: With EKS, researchers can dynamically scale computational resources based on workload demand.
- **Automation**: ML pipelines automate the training, tuning, and deployment processes, reducing manual intervention.
- **Monitoring & Logging**: Kubeflow provides real-time insights into ML job performance, leveraging AWS CloudWatch and Prometheus.

3.3.5 Amazon EKS (Elastic Kubernetes Service)

Amazon Elastic Kubernetes Service (EKS) is a managed Kubernetes service that allows scientists to run containerised machine learning workloads securely. Within the TRE, EKS provides:

1. ***Scalability:*** Scientists can dynamically deploy workloads without worrying about infrastructure management.
2. ***Security:*** Integrates with AWS IAM, ensuring granular access control.
3. ***High Availability:*** Distributes workloads across multiple availability zones.
4. ***Load Balancing with AWS ALB/NLB:*** Manages incoming traffic distribution to ensure optimal resource utilisation.
5. ***Integration with AWS Route 53:*** Ensures domain name resolution for applications hosted within the research environment.

Auto Scaling in EKS automatically adjusts the number of worker nodes based on computational demand. The key components of auto scaling include:

1. ***Cluster Autoscaler:*** Dynamically scales worker nodes in response to Kubernetes workload changes.
2. ***Horizontal Pod Autoscaler (HPA):*** Adjusts the number of running pods based on CPU/memory utilisation.
3. ***AWS Auto Scaling Groups:*** Manages EC2 instances to ensure cost efficiency.

3.3.6 Storage Solutions (EBS and EFS)

AWS provides multiple storage solutions within the TRE:

3.3.6.1 Amazon EBS (Elastic Block Store)

- High-performance block storage service for EC2 instances. In terms of performance, this storage solution supports provisioned IOPS (input/output operations per second) for low-latency applications. The limitations are that it can only be attached to one EC2 instance at a time. Ideal for databases, transactional workloads, and machine learning models requiring persistent, high-speed storage.

3.3.6.2 Amazon EFS (Elastic File System)

- These are shared file storage accessible by multiple EC2 instances. In terms of performance, it is scalable and optimised for workloads requiring concurrent access. Its limitations include higher latency compared to EBS, but ideal for shared research datasets. Ideal for machine learning pipelines, collaborative research data storage and shared scripts across research teams.

3.3.6.3 Choosing between EBS and EFS

- EBS should be used when scientists need high-speed, low-latency storage attached to a single EC2 instance, such as when running databases or individual ML model training.

- EFS should be used for shared storage, enabling multiple researchers to access the same datasets concurrently without needing multiple storage volumes.

3.4 CONCLUSIONS ABOUT USING AWS STRUCTURES FOR DEVELOPING A TRE

This project showed that creation of a TRE within AWS using Kubernetes to build the environment is a tractable objective. It provides a number of the key safeguards necessary to ensure that commercially sensitive data could be us to provide upscaled environmental data to researchers in a manner where the source data is provided in sufficiently abstracted form to protect direct loss of sensitive information.

The practicalities of setup and management of a TRE within AWS have been demonstrated. The cost and sustainability of this approach is not addressed in this conceptual pilot and would need to be assessed before implementation.

4 Sensor Data De-Identification within Trusted Research Environments

4.1 INTRODUCTION

The UKRI / NERC Digital Research Infrastructure funded AMPLIFY project funded by NERC aims to establish the mechanism to “deliver EDS services through consolidated data governance, processes, standards, policies and innovative new technologies”. The data commons are an important target for these integration efforts enabling data to be shared in common formats using common utilities across disciplines. AMPLIFY activities are intended to support the conceptualisation of end-to-end data sharing and processing pipelines, providing artificial intelligence (AI-based) Quality Assurance/Quality Control procedures on streaming data, and developing a portal offering access to sensor data and associated services.

The specific objectives on the AMPLIFY project that we seeking to address in this report are “providing support for computations that involve sensitive sensor data through the use of Trusted Research Environments (TREs).”

4.1.1 Motivation

There is a need to provide a unified TRE for researchers who wish to incorporate confidential data contained within Environmental Data Service (EDS) into their analysis. For example, some data supplied by the Environment Agency (EA) and commercial companies to the British Geological Survey (BGS) and the provide a more comprehensive insight into the environmental conditions, thus it would be beneficial to provide access to these third-party datasets. However, these datasets may contain information that can lead to unique identification of specific objects that have either commercial significance or pose a security risk for the data owners or the public. This poses the problem of de-identification of the third-party data by users of the EDS.

Two projects concerning third-party sensor data are relevant to the development the TRE:

- **BOOST**: amongst project aims, it sought to assign persistent identifiers (PIDs) to all sensors within NERC ecosystem, permitting users to check maintenance status and history of measurements associated with any given sensor. Individual sensors will be assigned a PID. One of the anticipated outcomes of this project is the NERC-wide adoption of a single PID assignment system for all the sensors within its capacity.
- **AMPLIFY**: an initiative to create a prototype of a unifying data-delivery environment that will interrogate all NERC sensor data and combine it with the third-party data. This initiative requires the adoption of PIDs across NERC and third-party sensors and instruments. Specific objectives relevant to this case study are:
  - Delivering a TRE for integration of EDS-held environmental data and measurements from other sources (e.g. companies and regulators) that allows combined outputs without risking exposure of confidential data sources.
  - Incorporation of sensor data which feeds directly into TRE to create dynamic data feeds to support dynamic process modelling, including sensitive commercial datasets.

The objective of the BOOST project indicates that it is essential to provide TRE users with output readings corresponding to individual sensors, so that sensor calibration metadata can be taken into account in the analysis of the data (Stocker et al., 2020). However, direct incorporation of third-party sensor feeds into the TRE may impose additional restrictions on the data sharing.

The expectation at project inception had been that the BGS user case study would be able to be implemented within the impending UoMDS environmental data hub, which is intended to incorporate a TRE. The presumed use case for the UoMDS TRE involved integration of environmental and medical data for studies of the environmental impacts upon health conditions e.g. relationships between disease prevalence and air pollution. Unfortunately, delivery of the Digital Solutions TRE has been delayed. Therefore, this document describes the envisaged BGS use case for implementation of a TRE, the intended implementation for this and a hypothesised future architecture.

4.1.2 User story

The development of this conceptual TRE was farmed by considering the following user example: an environmental scientist wishes to test a hypothesis about the relationship between a particular groundwater pollutant and environmental conditions within a region of interest (ROI). They may access the TRE to request water quality indicators collected from boreholes and pumping stations across ROI. They will want to extract the data in one of the commonly used geospatial formats (such as GeoJSON) to be able to relate it with spatial information from other data. Some of these indicators are collected from sites owned by commercial companies that wish to restrict public access to the exact locations of their boreholes. Much critical UK environmental scientific data is sourced from industry (e.g. water supply companies) or regulators (e.g. the Environment Agency in England or Scottish Environmental Protection Agency) and not from dedicated scientific endeavours. This poses the problem of i) providing accurate high-resolution sensor data to researchers while ii) satisfying restrictions of the third-party data owners.

This study identifies key problems and defines a set of technical requirements for the delivery of a Trusted Research Environment to delivery these capabilities.

4.1.3 Study objectives

The aim of this case study is to provide a conceptual model of de-identification process for the geospatial outputs within the TRE. Specifically, the following tasks will be addressed within this report:

- An overview of methods for obfuscating geospatial data from governmental and public datasets.
- A case study of a scenario where de-identification is performed after multiple datasets have been queries by their respective APIs. This is the case where the API to third party datasets is provided by data owners.
- Explore a scenario where de-identification is performed "behind the API". This is in case we have a unified API that relates all the datasets has been implemented within the TRE.
- A demonstration of geomasking borehole locations for an example ROI.

4.2 BACKGROUND ON GEOSPATIAL DATA DE-IDENTIFICATION

4.2.1 Statistical disclosure control

The problem of the de-identification of sensitive data has been encountered across multiple public institutions and national governments (Garfinkel et al. 2023). The continuous development of methods to de-identify various datasets has formed the basis of Statistical Disclosure Control (SDC) field (Abowd and Schmutte 2016; Domingo-Ferrer and Domingo-Ferrer, 2018; Griffiths et al. 2024). This discipline largely focuses on de-identifying microdata that is collected from individuals to preserve the privacy (Zandbergen 2014; Tiwari et al. 2023) or anonymity (Kalnis et al. 2007). The problem addressed in this report is simpler than the geoprivacy problem in the location aspect since all objects that must be de-identified are stationary. However, the challenge arises from the requirement that the readings from individual sensors are accurately reported on the sensor level (equivalent to microdata in survey and census data). This means data aggregation methods commonly used for SDC are not appropriate. Here we overview some approaches to geomasking that preserve the micro-level data reporting.

4.2.2 Geomasking on microdata level

The collection of SDC approaches to de-identify the geospatial data are referred to as geomasking. Common methods of geomasking include:

- **Affine transformation.** In this method, the geographic coordinates of spatial data are changed according to predefined transforms of the following types: translation, rotation, and scaling (Wang 2024). The disadvantage of applying this type of masking to the data is that the new coordinate may not have the same real-world context, and the data will not be easy to relate to other spatial datasets. This makes the affine approach to geomasking incompatible with the very purpose of the TRE which is meant to facilitate combination of cross-domain data.
- **Random perturbation.** In this method, geomasking is achieve by adding random noise to the coordinates of the spatial dataset. Widely used types of random perturbations are reviewed in (Zandbergen 2014): the coordinates may be displaced by a fixed distance in random direction, by a random distance in a random direction, or randomly moved to a superscribed donut shape around the true location. The distance and orientation can be sampled from various distributions. While these methods are reasonably effective for disclosing data that has a single instance in time (such as in a census), they incur and increased risk of disclosure in situations where a time series is reported from a single location. We provide and illustrative example of possible disclosure for several random perturbation methods in Figure 4. If the extracted dataset includes 100 readings from the same sensor and its location is obfuscated via random perturbation methods, it is not easily discovered. If a user requests reading for a longer period, they will have access to a large set of noisy coordinates for a single sensor. Then, a simple manipulation of these coordinates may lead to the high-precision estimation of the original coordinates. We thus conclude that this approach is inappropriate for geomasking live feed sensor data.
- **Gridding.** This method falls within a broader category of aggregation methods that usually provide summary statistics of microdata for a coarse grid of polygons that may correspond to administrative or geographic units, depending on the context (Armstrong et al, 1999; Skøien et al. 2024). To satisfy the requirement to provide individual sensor data, we propose adopting a simplified version of data gridding. For third-party sensor locations, coordinates of a sensor are replaced by coordinates of the geographic unit it falls within. The benefit of using this method is that the risk of disclosure does not increase with the growing number of readings from a single sensor.

4.2.3 Practical considerations for grid geomasking

While gridding sensor data satisfies requirements i) and ii) set in the user story, this approach implies existence of a predetermined grid covering UK territories that can be accessed at the stage of data interrogation. Several aspects of the grid must be defined and added to the metadata of all output files, such as the coordinate system within which the grid is created, the definition of the smallest regional unit, or the resolution in case the grid is regular.

An example of such grid can be found in the Open Geography Portal of the Office for National Statistics ('GEOSTAT (December 2011) Boundaries UK BGE' 2022). This grid is formed of identical

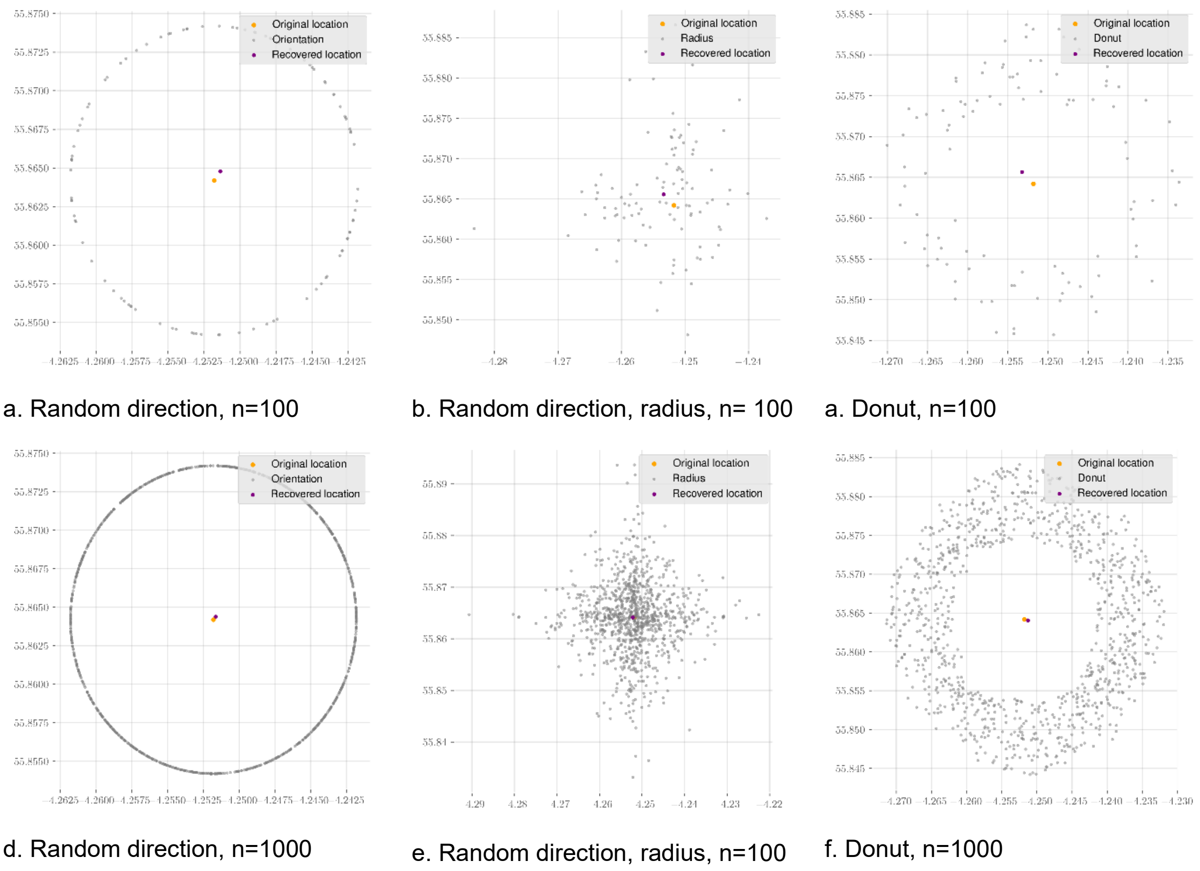

a. Random direction, n=100

b. Random direction, radius, n= 100

a. Donut, n=100

d. Random direction, n=1000

e. Random direction, radius, n=100

f. Donut, n=1000

Figure 4. Recovery of original stationary coordinates from the collection of perturbed values associated with long timeseries.

1km^2 areas which is the smallest resolution of a grid available. It remains to be determined whether this resolution is sufficiently large for the purpose of obfuscating the third-party sensor data. A regular grid such as this one may be beneficial for the purpose of relating cross-domain datasets. The grid information must be included in the output dataset data to ensure interpretability by the users, for example using the ArcGIS convention ('GEOSTAT Grid Layer', n.d.).

4.3 PROPOSED IMPLEMENTATION WITHIN TRE

Currently within the BGS sensor database there exists a dictionary related to sensor confidentiality. It has three definitions:

- Accessible - data from sensors with this definition is published in full.
- Confidential - data from sensors with this definition is completely removed from public access.
- Restricted - an intermediate definition that currently has no logic attached to it.

We proposed utilising these definitions for the sensors whose locations need to be de-identified, then the data can be made publicly available with geospatial information obfuscated from the TRE user. Note that the restricted flag may have to be made visible to the user to ensure interpretability of the output, especially if only some entries in the final dataset are geomasked while others are provided with their exact locations.

4.3.1 An example dataset

We considered a small collection of borehole locations available for an area within an arbitrary ROI. These locations are selected for a small rectangular area via the BGS Onshore GeoIndex tool and extracted to a shape file. An additional field titled “Restricted” is added to the dataset, with a Boolean data type indicating whether an individual borehole is marked as restricted. 1% of all boreholes within the selected area are randomly marked as restricted and the Boolean flag is toggled to True for those locations. In addition, a field titled “Pseudo PID” is added to the dataset to emulate the PID assigning procedure to be implemented within the BOOST project. A unique integer number is assigned to that field. The resulting example dataset is visualised in Figure 5a.

To demonstrate the principle of geomasking by gridding, we use the grid publicly available on the OSN geographic data portal. It utilises a 1km^2 base grid from the Eurostat set, which corresponds to the finest grid resolution permitted for census reporting in the European Commission. A subset of the grid for the considered ROI is demonstrated in Figure 5b.

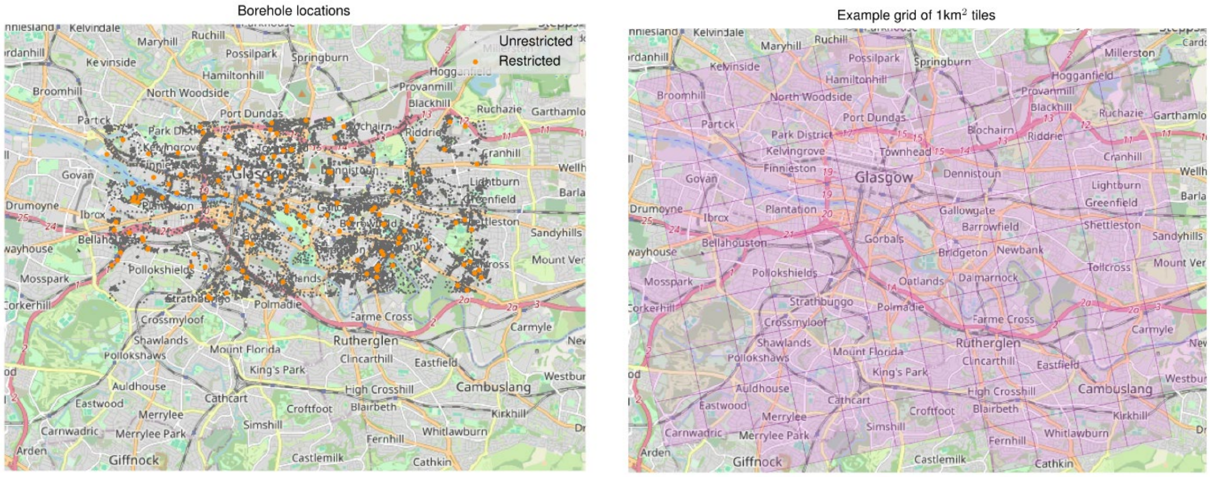

(a) An example dataset of boreholes. Grey dots marked unrestricted boreholes for which the locations may be revealed. Orange dots represent a hypothetical set of restricted boreholes for which geospatial data must be de-identified.

(b) An example of Eurostat base grid overlayed over the central area of Glasgow.

Figure 5 An example set of geolocations corresponding to borehole sensor collections and the base grid of central Glasgow region.

Although there are efforts being made to implement a unifying API for interrogation of all sensor data across different providers, currently each dataset is maintained by its respective owner and thus comes with its own API. Below we discuss how the geomasking can be perfoemed in the current setting (Scenario 1) and in the future setting whe a common API is available (Scenario 2). For both scenarios we consider the configuration of the TRE where the data is stored on a secure server and the users are accessing it via a virtual machine, similartly to the Geospatial Virtual Data Enclave framework suggested in (Richardson et al. 2015). However, we consider a less strict setting than originally proposed for ensuring individual data privacy where a TRE user may be permitted to extract the data after it has been de-identifyied and passed the SDC checks. This is will offer flexibility in situations where users wish to run computationally expensive analyses elsewhere or combine sensor data with sensitive data in other TREs (for example public healthrecords).

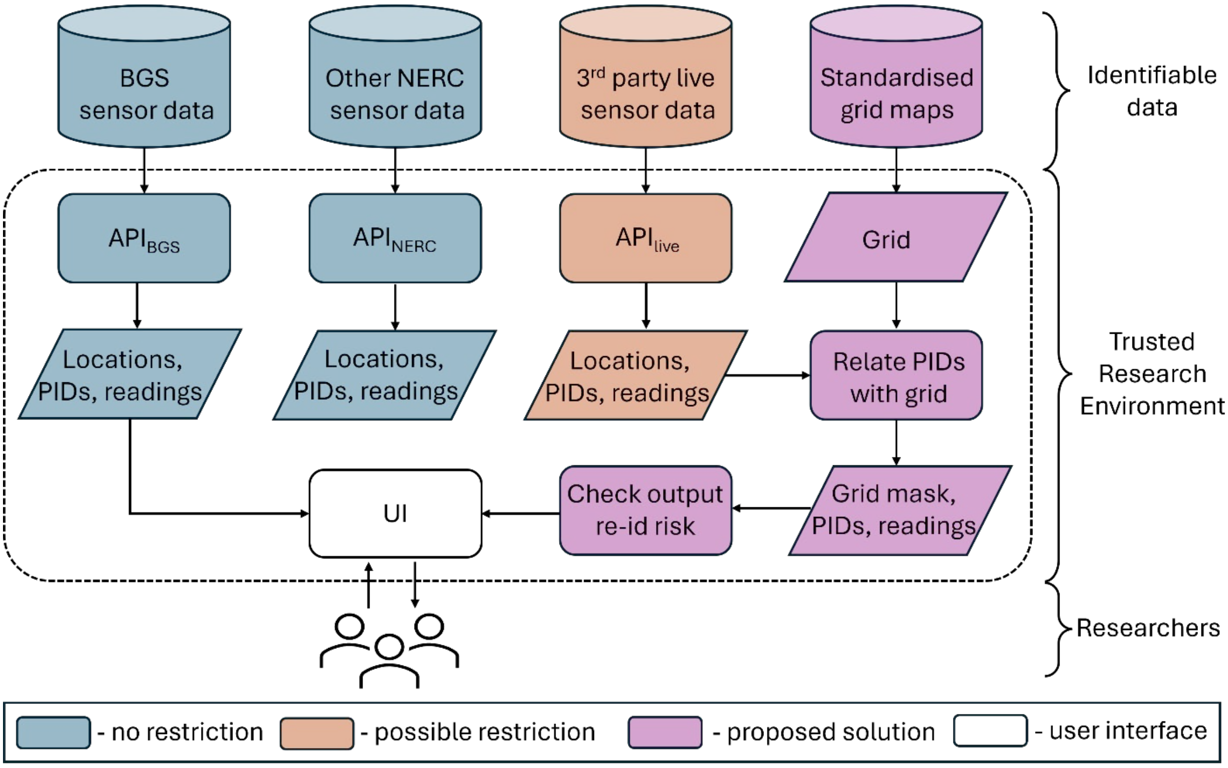

Figure 6. Proposed workflow of sensor data post-processing within the TRE in a scenario where all datasets relevant for a user are interrogated by their respective APIs independently.

4.3.2 Scenario 1: multiple APIs

In the scenario where all datasets will be interrogated by their respective APIs within the TRE, query results must be post-processed before being made available to users. An overview of this procedure can be found in Figure 6. Within the TRE, a trusted researcher accesses user interface (UI) to query multiple datasets available within the EDS, some of which may have a third-party restriction. The user-specified query is transformed to multiple queries that are passed to relevant APIs. Raw data is transferred to the TRE in response to the queries. If the restricted data has been accessed, it is related to the prespecified grid, and the locations of the restricted sensors are geomasked. The output dataset is constructed by combining the accessible and geomasked data. A definition is added to the output dataset to indicate whether the location has been geomasked. The metadata of the grid is added to the metadata of the output dataset. The output dataset is checked for the disclosure risk before being released to the user.

User-accessible output for the example dataset is presented in Figure 7. It contains exact locations for sensors marked as accessible, and corresponding tile locations if sensors marked as restricted. Individual measurements are provided without any manipulation and can be traced to the sensor that produced them via PIDs.

In the interest of lowering the participation barriers for various data providers, this configuration must be accompanied by a set of recommendations for the owners of the third-party API, including the specifications for creating the PIDs for each sensor and specification of database dictionaries that they should implement to mark their data as restricted. An agreement on types and definitions of various restriction settings must be reached with all participating data providers. Note that this may lead to the expansion of the original confidentiality dictionaries.

4.3.3 Scenario 2: a single API

This scenario is considered under the assumption that a single common API is developed for all the datasets that will be accessible from the TRE. This implies that all data can be extracted simultaneously via a single query and thus will include a mixture of restricted and unrestricted objects. A proposed workflow is shown in Figure 8. Within the TRE, a trusted researcher accesses

UI to query multiple datasets available within the EDS, some of which may have a third-party restriction. The query is passed through the common API and all raw data from the relevant datasets is passed to the TRE. If the restricted flag is implemented for all datasets, then the subset for which flag is toggled to true can be geomasked as was done in Scenario 1. Otherwise, all sensor locations are related to the pre-specified grid and the exact coordinates are replaced by the grid tiles. Grid metadata is added to the output metadata. The output dataset is checked for the disclosure risk before being released to the user.

User-accessible output for the example dataset is presented in Figure 9. It contains geomasked locations for all sensors. Individual measurements are provided without any manipulation and can be traced to the sensor that produced them via PIDs, but the locations of sensors in this case will be provided as a tile-shaped polygons.

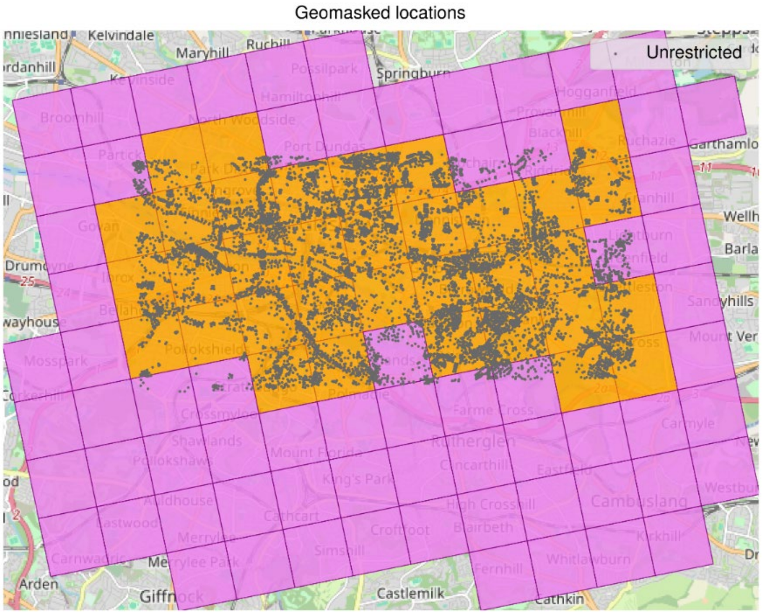

Figure 7. The example presentation of the packaged output that contains unrestricted borehole locations as exact points, and gasmasks restricted borehole locations to a pre-specified grid.

If a single API is implemented, it implies that some common dictionaries have been introduced for data across all providers. In this case, the third-party data can be incorporated in plug-and-play manner as the logic for each dataset will have been implemented within the API.

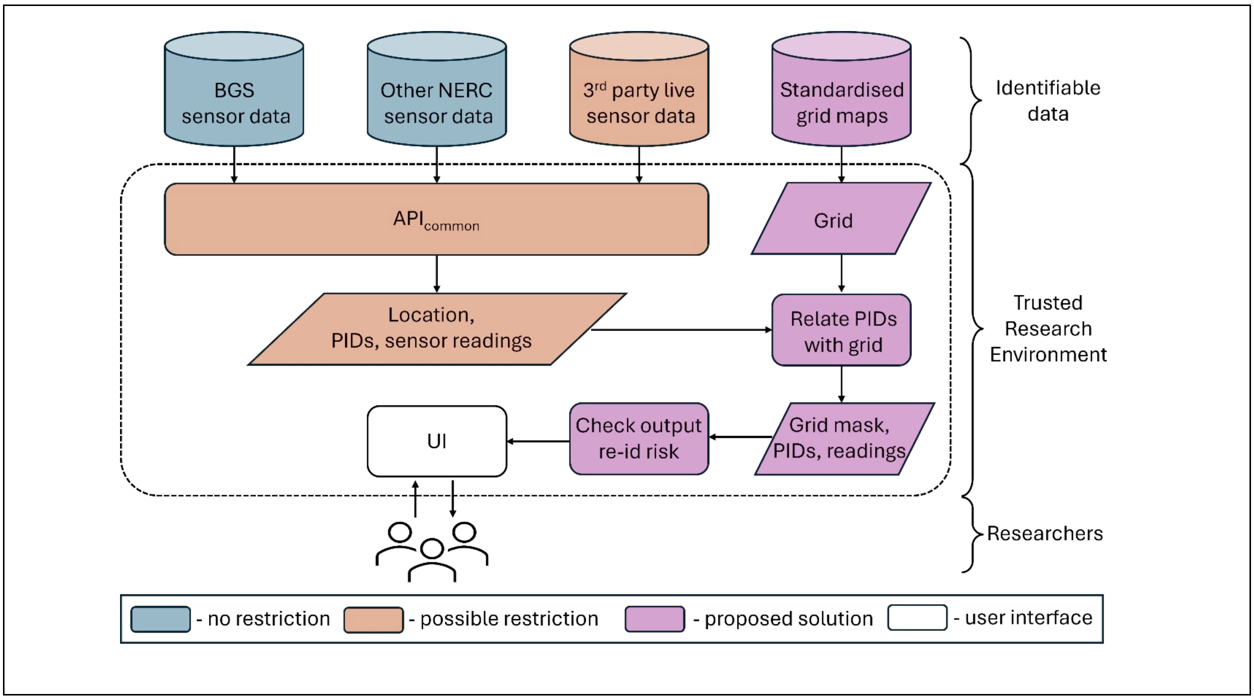

Figure 8. A proposed workflow for post-processing data acquired from the mixed datasets by a single API.

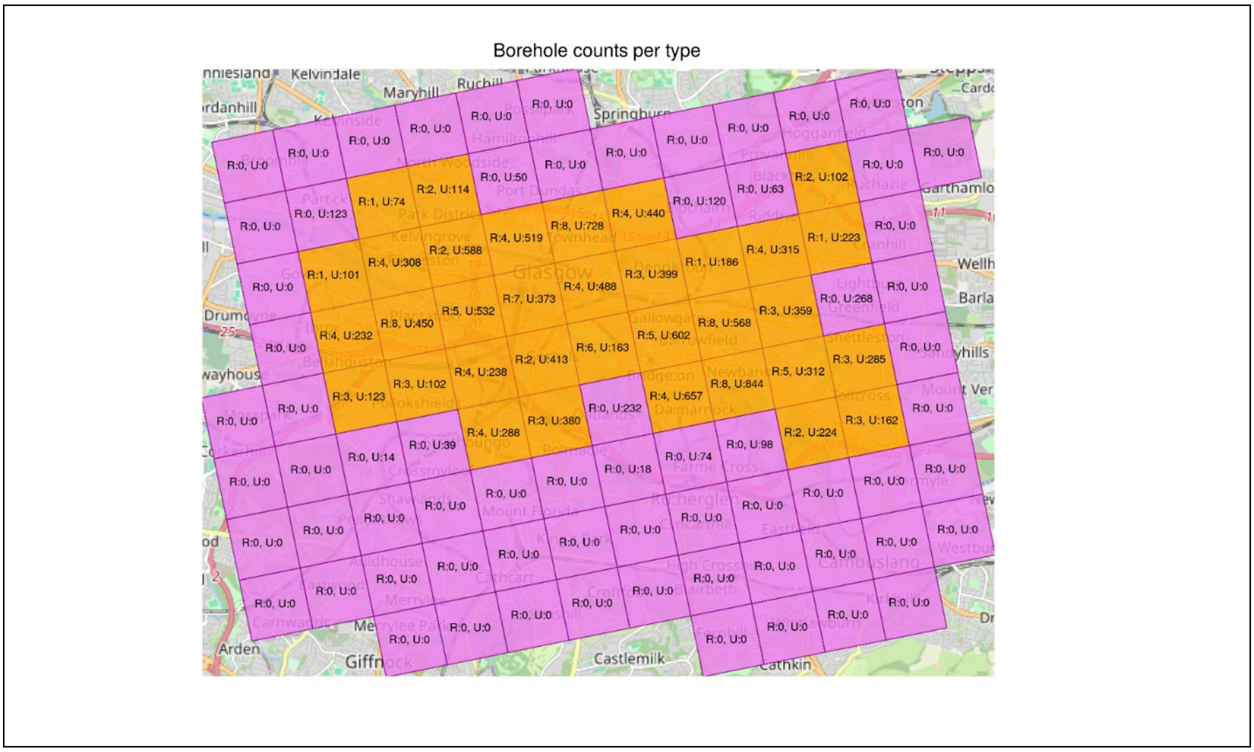

Figure 9. Masked output for the example dataset in case of a single API. For each tile, R restricted borehole locations and U of unrestricted boreholes are reported. The output dataset contains tile locations for all sensors, regardless of their confidentiality level.

4.4 CONSIDERATIONS FOR OUTPUT CHECKING BEFORE RELEASE TO USER

Geomasked data may carry a risk of disclosure if the masking method is not appropriate for the type of data, or if it is combines with other datasets carrying additional information. To ensure the safety of outputs, assessment of location disclosure risk from the output datasets must be performed. Several stages of checks may be appropriate once all types of restricted data are identified:

- Checking of the geomasking principle prior to implementation. The example here provides a conceptual model but does not include a model of potential “geohacking” attacks on the output data.
- Automated or semi-automated checking of the geomasked datasets before they are release to users.
- Manual approval by TRE moderators to approve release of geomasked datasets if cross-domain data is requested.

5 Conclusions

This report establishes a concept and defines the practicalities for creating a Trusted Research Environment that enables the sharing of modelling outcomes informed by hypothetical confidential water industry data to a range of stakeholders. It is predicated in meeting the fundamental requirements that the end-users cannot identify the individual input data points and streams. On the other hand, the modelling outcomes must be available and shareable with the source data embedded but hidden within the models, thus meeting data confidentiality requirements.

The first part of the report describes the architecture of a commercial cloud Kubernetes implementation using Amazon Web Services as a proposed host for such services. The choice of AWS is pragmatic based on prior BGS experience for such project rather than an advertisement. However, commercial cloud vendors do have the benefit of an easily configurable environment to allow such projects to be implemented quickly and easily, especially using the Kubernetes structures they support.

The second part of the report provides an overview of common approaches to de-identification of geospatial data, also called geomasking. The shortcomings of random perturbation approach for the long-running time series data identified in a simulation example led to the selection of gridding as the most appropriate geomasking technique. Two possible scenarios of de-identification are provided, depending on the type of API used to access varying data sources within the TRE.

Because of issues elsewhere in the wider project, it has not been possible to implement this TRE as was intended from the original proposal. Consequently, this report is the tangible output from this project. This establishes that the establishment of a spatial TRE that incorporates both spatial and time variant data, with the data embedded into a process model, then outputs delivered in an upscaled from to obscure the source data are a realistic deliverable from such a project.

The commercial cloud architecture into which this hypothetical TRE will be implemented is shown to be a readily configurable model for the task at hand. It contains the necessary safeguards and flexibility for implementation to be a realistic prospect. Due to issues with the UoMDS team's efforts to create a more fully configured TRE, the exact architecture within which their TRE might be developed is unknown. The BGS use case highlights how an AWS implementation provides an acceptable solution that might deliver the necessary functionality in a light-touch way. This route presents further opportunities for delivery of an effective TRE.

6 References

Abowd, John M., and Ian M. Schmutte. 2016. ‘Economic Analysis and Statistical Disclosure Limitation’. *Brookings Papers on Economic Activity* 2015 (1): 221–93. https://doi.org/10.1353/eca.2016.0004.

Armstrong, Marc P., Gerard Rushton, and Dale L. Zimmerman. 1999. ‘Geographically Masking Health Data to Preserve Confidentiality’. *Statistics in Medicine* 18 (5): 497–525. https://doi.org/10.1002/(SICI)1097-0258(19990315)18:5<497::AID-SIM45>3.0.CO;2-#.

Elliot, Mark J. and Josep Domingo-Ferrer. 2018. ‘The Future of Statistical Disclosure Control’. The National Statistician’s Quality Review. https://analysisfunction.civilservice.gov.uk/policy-store/privacy-and-data-confidentiality-methods-a-national-statisticians-quality-review-nsqr/.

Emily Griffiths, Carlotta Greci, Yannis Kotrotsios, Simon Parker, James Scott, Richard Welpton, Arne Wolters, and Christine Woods. 2024. ‘Handbook on Statistical Disclosure Control for Outputs’. 2. Safe Data Access Professionals Working Group. https://securedatagroup.org/wp-content/uploads/2025/03/sdc-handbook-v2.0.pdf.

Garfinkel, Simson, Joseph Near, Aref Dajani, Phyllis Singer, and Barbara Guttman. 2023. ‘De-Identifying Government Datasets: Techniques and Governance’. NIST SP 800-188. Gaithersburg, MD: National Institute of Standards and Technology (U.S.). https://doi.org/10.6028/NIST.SP.800-188.

‘GEOSTAT (December 2011) Boundaries UK BGE’. 2022. https://geoportal.statistics.gov.uk/search?q=BDY_GEOSTAT%20DEC_2011.

‘GEOSTAT Grid Layer’. n.d. Accessed 18 March 2025. https://services1.arcgis.com/ESMARspQHYMw9BZ9/ArcGIS/rest/services/GEOSTAT_(Dec_2011)_GEC_in_the_United_Kingdom/MapServer/0.

Kalnis, Panos, Gabriel Ghinita, Kyriakos Mouratidis, and Dimitris Papadias. 2007. ‘Preventing Location-Based Identity Inference in Anonymous Spatial Queries’. *IEEE Transactions on Knowledge and Data Engineering* 19 (12): 1719–33. https://doi.org/10.1109/TKDE.2007.190662.

Richardson, Douglas B., Kwan ,Mei-Po, Alter ,George, and Jean E. and McKendry. 2015. ‘Replication of Scientific Research: Addressing Geoprivacy, Confidentiality, and Data Sharing Challenges in Geospatial Research’. *Annals of GIS* 21 (2): 101–10. https://doi.org/10.1080/19475683.2015.1027792.

Skøien, Jon Olav, Nicolas Lampach, Helena Ramos, Rudolf Seljak, Renate Koeble, Linda See, and Marijn van der Velde. 2024. ‘Flexible Approach for Statistical Disclosure Control in Geospatial Data’. arXiv. https://doi.org/10.48550/arXiv.2410.17601.

Stocker, Markus, Louise Darroch, Rolf Krahl, Ted Habermann, Anusuriya Devaraju, Ulrich Schwardmann, Claudio D’Onofrio, and Ingemar Häggström. 2020. ‘Persistent Identification of Instruments’. *Data Science Journal* 19 (1). https://doi.org/10.5334/dsj-2020-018.

Tiwari, Alok, Sohail Ahmad, Emad Qurunflah, Mansour Helmi, Ayad Almaimani, Alaa Alaidroos, and Majed Mustafa Hallawani. 2023. ‘Exploring Geomasking Methods for Geoprivacy: A Pilot Study in an Environment with Built Features’. *Geospatial Health* 18 (2). https://doi.org/10.4081/gh.2023.1205.

Wang, Jue. 2024. ‘Geomasking to Safeguard Geoprivacy in Geospatial Health Data’. *Encyclopedia* 4 (4): 1581–89. https://doi.org/10.3390/encyclopedia4040103.

Zandbergen, Paul A. 2014. ‘Ensuring Confidentiality of Geocoded Health Data: Assessing Geographic Masking Strategies for Individual-Level Data’. *Advances in Medicine* 2014:567049. https://doi.org/10.1155/2014/567049.