D7.2 Guideline for Health Data Access Bodies on data minimisation, pseudonymisation, anonymisation and synthetic data

TEHDAS2 – Second Joint Action Towards the European Health Data Space

24 March 2026

Co-funded by the European Union

0 Document info

Disclaimer
Views and opinions expressed in this deliverable represent those of the author(s) only and do not necessarily reflect those of the European Union or HaDEA. Neither the European Union nor the granting authority can be held responsible for them.

0.1 Authors

Author(s)	Organisation
Pia Brinkmann	BfArM, Germany
Luca Augello	ARIA/RL, Italy
Petr Holub	MU, Czech Republic
Jaakko Lähteenmäki	VTT, Finland
Eva Anjo	Serviços Partilhados do Ministério da Saúde (SPMS), Portugal
Justin Ansotte	Health Data Agency, Belgium
Jos Hendrikx	Nictiz, The Netherlands
Tamás Kovács	ESZFK, Hungary
Victor Leandre-Chevalier	Health Data Hub, France
Peter van Meerendonk	Nictiz, The Netherlands
Farzaneh Michaud	Health Data Hub, France
Juha Pajula	VTT, Finland
Lea Rizzuto	Health Data Hub, France
Lise Skovgaard Svingel	RM, Denmark
Katharina Schneider	BfArM, Germany
Hanna Tervonen	Findata, Finland
Steven Wolter	BfArM, Germany

0.2 Keywords

Keywords	TEHDAS2, Joint Action, Health Data, European Health Data Space

0.3 Document history

Date	Version	Editor	Change	Status
08/04/2025	0.1	Pia Brinkmann	Initial document creation	Draft
12/05/2025	0.2	All contributors	First draft	Draft
01/07/2025	0.3	Luca Augello, Pia Brinkmann, Jaakko Lähteenmäki, Petr Holub	Draft for consortium feedback	Draft
05/09/2025	1.0	Luca Augello, Pia Brinkmann, Jaakko Lähteenmäki, Petr Holub	Document to be submitted for public consultation	Draft
13/02/2026	2.0	All contributors	Deliverable draft after integration of public consultation feedback	Draft
10/03/2026	2.1	All contributors	Deliverable draft after integration of EC feedback, submitted for consortium	Draft
24/03/2026	2.2	All contributors	Deliverable draft, ready for acceptance	Final

Accepted in Project Steering Group on 24 March 2026.

Contents

1 Executive summary

This guideline focuses on processing electronic health data within the European health data space (EHDS), by detailing methods for **data minimisation**, **pseudonymisation**, **anonymisation**, and **synthetic data generation**. The goal is to create a secure, interoperable, and efficient health data ecosystem for secondary use in compliance with the EHDS and General Data Protection (GDPR) regulations, which means using health data beyond direct patient care.

One foundational principle for handling health data is **data minimisation**. This means that only the **minimum amount of personal health data** that is adequate, relevant, and limited to what is necessary for a specific purpose should be processed. This principle applies throughout the entire lifecycle of the data, from when it is first collected and prepared by the health data holder, to when it is assessed by a health data access body (HDAB), and finally, during its use and processing by the health data user. Data minimisation can involve reducing the volume of data, limiting its detail (granularity), making sensitive information less specific, or restricting geographical or temporal scopes. In addition, it can be applied in five dimensions (“Who”, “What”, “When”, “Where”, “How”). This helps to significantly reduce risks related to confidentiality, integrity, and availability of data.

Pseudonymised data remain personal data under Regulation (EU) 2016/679 (GDPR) and reduce the likelihood of direct identification, while preserving high data utility. Pseudonymised data is one data format HDABs may permit access to, if the re-identification risk is justified and appropriately mitigated (see Article 66(3), Regulation (EU) 2025/327 (EHDS)). The information needed to link these pseudonyms back to the original individuals is kept entirely separate and secure. Pseudonymisation is particularly valuable because it allows for the **linkage of different health datasets.** This is vital for comprehensive research. It also supports the **rights of data subjects**, such as the ability to opt-out of data use for future projects, or to be informed of significant findings related to their health data. The HDAB plays a key role in defining and overseeing the pseudonymisation process.

Finally, **anonymisation** and **synthetic data generation** offer strong privacy protection, often used when data or analysis results are intended to be exported or made publicly available. **Anonymisation** (Regulation (EU) 2016/679 Recital 26 (GDPR), Article 29 Data Protection Working Party, Opinion 05/2014 on Anonymisation Techniques, WP216, adopted on 10 April 2014) transforms original personal data so that it no longer relates to an identified or identifiable person, meaning that the individual cannot be re-identified by any reasonable means. **Synthetic data generation**, on the other hand, creates artificial datasets that mimic statistical properties of the original data. Depending on the methodology used and the residual risk of re-identification, such data may or may not constitute personal data under the Regulation (EU) 2016/679 (GDPR). While distinct, both methods require the HDAB to establish similar processes for evaluating **data quality**, performing **privacy risk assessments** (looking at risks of re-identification and inferring sensitive information), and implementing **disclosure controls**. All such activities must be thoroughly **documented** to ensure transparency and accountability. It should be noted that the EHDS does not impose legal obligations regarding synthetic data generation, but HDABs may support its use via evaluation frameworks, as part of enabling responsible data access.

1.1 Abbreviations

Term	Abbreviation
Computed tomography	CT
Digital imaging and communications in medicine	DICOM
Differential privacy	DP
Electrocardiograms	ECG
Electrodermal activity	EDA
European data protection board	EDPB
Electroencephalograms	EEG
Regulation (EU) 2025/327	EHDS
Electronic health record	EHR
Electromyograms	EMG
European Union	EU
Generative adversarial network	GAN
Regulation (EU) 2016/679	GDPR
Health data access body	HDAB
Kidney transplantation	KT
Large language models	LLM
Machine learning	ML
Magnetic resonance imaging	MRI
Named entity recognition	NER
Positron emission tomography	PET
Secure processing environment	SPE
Trusted health data holder	TDH
Second joint action towards the European health data space	TEHDAS2
Trusted third party	TTP
Variational autoencoders	VAE
Work package 7	WP7
Whole Slide Image	WSI
Zone improvement plan	ZIP code

2 Introduction

Advancing health data use in the European Health Union

As part of its work on the European Health Union, the European Union (EU) is advancing the use of health data for secondary purposes, including research, innovation and policymaking. Smooth and secure access to data will drive the development of new treatments and medicines and optimise resource utilisation – all with the overarching goal of improving the health of citizens across Europe.

TEHDAS2, the second joint action Towards the European Health Data Space, represents a significant step forward in this vision. The project will develop guidelines and technical specifications to facilitate smooth cross-border use of health data, and support data holders, data users and the new health data access bodies in fulfilling their responsibilities and obligations outlined in the European Health Data Space (EHDS) Regulation.

TEHDAS2 focuses on several critical aspects of health data use.

- Data discovery: findability and availability of health data, ensuring it is accessible for secondary purposes.
- Data access: developing harmonised access procedures and establishing standardised approaches for granting data access across Member States.
- Secure processing environment (SPE): defining technical specifications for environments where sensitive health data can be processed safely.
- Citizen-centric obligations: providing guidance on fulfilling obligations to citizens, such as communicating significant research findings that impact their health, informing them about research outcomes and ensuring transparency in how their data is used.
- Collaboration models: developing guidance on collaboration and guidelines on fees and penalties as well as third country and international access to data.

TEHDAS2 will contribute to harmonised implementation of the Regulation (EU) 2025/327 (EHDS) through the concrete guidelines and technical specifications. Some of these documents and resources will also provide input to implementing acts of the regulation. Hence, the joint action will increase the preparedness for the EHDS implementation and lead to better coordination of Member States' joint efforts towards the secondary use of health data, while also reducing fragmentation in policies and practices related to secondary use.

This document should be understood as an expert opinion and guidance document developed within the TEHDAS2 framework, reflecting technical and expert input from the project partners. It is not legally binding and does not constitute a formal guideline or technical specification under the European Health Data Space.

This document does not represent the position of the European Commission.

Legally binding and enforceable requirements under the European Health Data Space are laid down in Regulation (EU) 2025/327 and, where applicable, in Implementing Acts adopted by the European Commission, within the limits of the empowerments provided by the Regulation.

The work performed in work package 7 (WP7) addresses “Safe and secure processing” of electronic health data within the HealthData@EU infrastructure. The goal is to enable secure processing of EU citizens’ electronic health data for secondary purposes while fostering a secure, interoperable, and efficient health data ecosystem. The output of this work package consists of guidelines and technical specifications that shall further inform decisions and technical frameworks to set up the EHDS.

The results of WP7 are distributed across five tasks. Task 7.1 provides guidance to users about their duties and responsibilities when analysing data in a SPE. Next, guidelines for data minimisation, pseudonymisation, anonymisation and synthetic data generation give guidance on how to address these topics and their challenges (task 7.2 includes sub-tasks: 7.2.1, 7.2.2, 7.2.3 & 7.2.4). Specifications for the implementation of a common IT infrastructure (task 7.3) shall help member states to connect to the EHDS ecosystem. To ensure interoperability, common security requirements applicable to all SPEs are defined in addition to functional and technical services that should be part of all SPEs (task 7.4). Lastly, information about data linkage techniques and possibilities of quality control of linked data are collected (task 7.5).

Here is an overview of the documents that are part of WP7:

- Guidelines for data users on how to use data in a secure processing environment (task 7.1);
- Guidelines for Health Data Access Bodies on data minimisation, pseudonymisation, anonymisation and synthetic data (task 7.2);
- Technical specifications for Health Data Access Bodies on the implementation of the common IT infrastructure (task 7.3);
- Technical specifications for Health Data Access Bodies on the implementation of secure processing environments (task 7.4);
- Guidelines for Health Data Access Bodies on linkage of health datasets (task 7.5).

2.1 Overview of the content and scope of this guideline

Figure 1. The EHDS user journey depicting data minimisation, pseudonymisation and anonymisation.

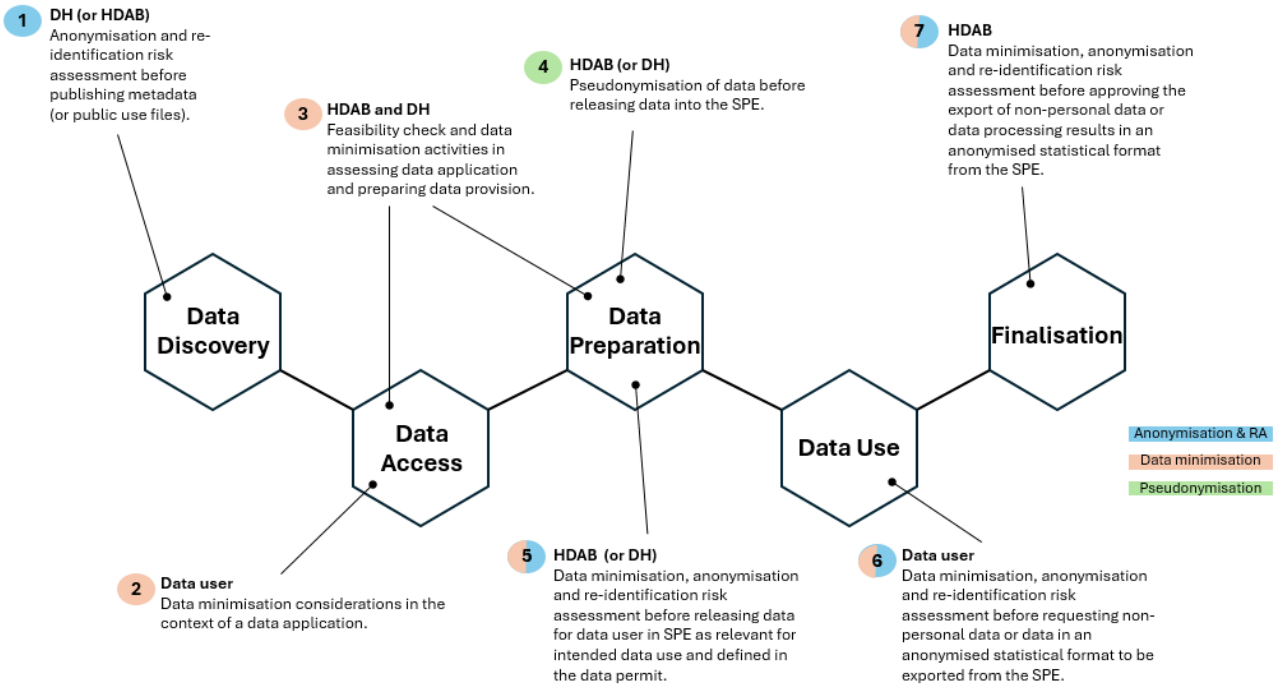

Abbreviation: DH: Data holder; HDAB: Health Data Access Body; RA: re-identification risk assessment; SPE: secure processing environment.

Figure 1 provides an overview of the key phases of the EHDS user journey where data minimisation, anonymisation and re-identification risk assessment, and pseudonymisation are relevant. Several phases involve multiple of these aspects simultaneously. For simplicity, trusted health data holders, intermediation entities, and trusted third parties have been excluded from the figure.

The scope of this guideline focuses on the phases after the data discovery phase (see Figure 1 and Annex 3: User journey) and covers various aspects of health data minimisation and de-identification.

Figure 2. The EHDS user journey divided into before and after the actual data processing takes place.

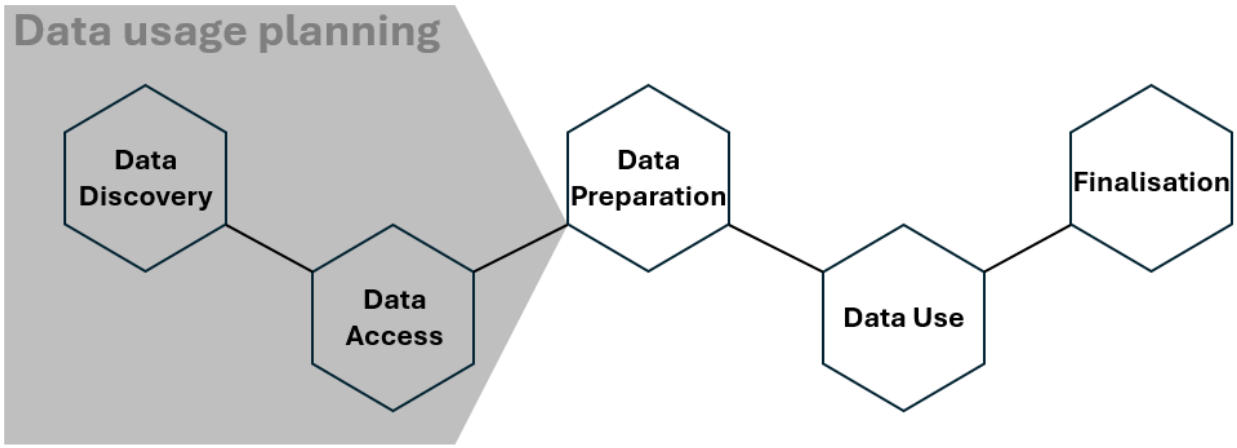

Along the user journey, the first phases include *data usage planning* (see Figure 2). This planning involves specifying an approach on how to perform data minimisation, specific requirements for pseudonymisation, anonymisation or synthetic data generation before the data is actually being processed. The data usage planning ends when the health data request (Article 69, EHDS) or data access application (Article 67, EHDS) is approved and data preparation begins.

Regarding data minimisation, this guideline elaborates on the dimensions of data minimisation for the use of secondary health data. This includes limiting the amount, type, and granularity of data during data preparation. While a first assessment must occur prior to the approval of the data access application or data request (Article 66(1), EHDS), the obligation to minimise data continues throughout all processing activities (i.e., data access, data preparation, data use and finalisation phases). It applies equally to data holders, HDABs, and data users.

Pseudonymisation should be performed as early as possible where appropriate, and additional safeguards – including re-pseudonymisation where necessary – may be required before data provision in a SPE, depending on the specific context and risk assessment.

Furthermore, it should be noted that anonymisation is not a binary nor permanent status. Previously anonymised data may, in the future, cease to meet the conditions for anonymity due to technological advances or the ability to combine multiple datasets. Under Recital 26 of the Regulation (EU) 2016/679 (GDPR), data are considered anonymised only if the data subject is not identifiable by any means reasonably likely to be used, taking into account objective factors such as cost, time, and available technologies. Therefore, the concept of anonymised data in this guideline refers to data that has been processed through anonymisation techniques that meet this “reasonableness”.

Anonymisation and synthetic data generation are related but distinct concepts. While they may involve similar techniques, anonymisation implies a transformation of real data to prevent re-identification, whereas synthetic data generation creates artificial data based on models or distributions. Synthetic data is not necessarily anonymised. The EDPB stated that the documentation of synthetic data generation should include the model’s theoretical resistance to re-identification techniques (§58e) and meet the purpose and data minimisation principles (§64) (EDPB Opinion of the Board (Art. 64), Opinion 28/2024 on certain data

protection aspects related to the processing of personal data in the context of AI models[1]). In other words, synthetic data may fall under GDPR if individuals can still be re-identified with reasonable effort. Therefore, it is necessary to demonstrate the resistance to re-identification.

Next, it should be noted that oftentimes, this guideline left out trusted data holders for simplicity reasons.

It is recognised that different types of SPE architectures may be implemented across Member States, depending on technical and organisational choices. For example, there may be isolated components dedicated to data preprocessing by the HDAB, and other components enabling data processing by authorised users. In this guideline, the term SPE (Article 2(1)(c), EHDS) refers to the complete set of environments or components that, taken together, meet the requirements set out in Article 73 of the EHDS Regulation. These include strict control over data access, processing, export, logging, and compliance with data protection, intellectual property, and confidentiality obligations.

Lastly, it should be mentioned that this document does not take into account the EC's Digital Omnibus proposal, published on 19 November 2025, as it still is at proposal stage and is currently subject to the legislative procedures and interinstitutional negotiations. As such, its final content, scope and legal effects are not yet determined and may be subject to amendment.

2.1.1 Data types

Under the EHDS Regulation, all categories of data processed must be subject to appropriate minimisation, pseudonymisation or anonymisation measures depending on purpose and context (Article 66(3), EHDS). Similarly, the generation of synthetic data may be relevant across all data types. In the following, a high-level categorisation of relevant datatypes is presented below (please see D5.1 Guideline on data description, section 5.3 for details on classification on EHDS categories):

- **Structured data** refers to healthcare or health-related information that is stored according to a predefined schema typically in tabular form, where each row represents an individual person or an observation, and columns correspond to specific attributes such as demographics, diagnoses, treatments, lab results, and service usage. Structured health data can be cross-sectional or longitudinal. It may comply with a healthcare standard, such as HL7 FHIR, but can also exist in non-standardised formats.
- **Medical imaging data** refers to digital representations of visual information captured for clinical purposes. This includes traditional scans such as X-rays, CT scans, MRI, ultrasound, and PET, as well as clinical photographs – such as images of skin, wounds, or surgical sites. These data are typically stored in standardised formats like DICOM (Digital Imaging and Communications in Medicine), and may include raw image files, metadata (e.g., patient identifiers, imaging parameters), and annotations used for diagnosis, monitoring, or research. Image metadata may include direct identifiers (e.g., patient names or IDs in DICOM headers).
- **Bio-signal data** refers to measurable signals derived from a biological source, reflecting underlying physiological or biological processes recorded from the human

[1] https://www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf (Adopted 17 December 2024).

body. These signals include electrocardiograms (ECG) for heart activity, electroencephalograms (EEG) for brain activity, electromyograms (EMG) for muscle function, electrodermal activity (EDA) for skin conductance, audiological signals (e.g., tympanometry), and other measurements such as temperature, blood pressure, respiration rate, and oxygen saturation.

- **Genetic data** refers to information derived from an individual's DNA, representing their genetic makeup. It includes sequences of nucleotides (A, T, C, G), genetic variants, mutations, and structural variations that influence traits, disease susceptibility, and responses to treatments. Next to genetic data, genomic data refers to data on all genes, while genetic data refers to data on some genes.
- **Textual data** refers to unstructured or semi-structured information recorded in free-text form within electronic health records (EHRs) or other information systems that include clinical notes, discharge summaries, pathology reports, patient histories and other relevant information. This data often includes physician observations, diagnostic assessments, treatment plans, and patient-reported symptoms, but can also be non-clinical data such as interviews and questionnaire responses.
- **Multimodal data** refers to information that combines different types of media, such as text, images, graphs, audio, and video, in defined formats. This data is typically provided on forms that are filled out by devices used in testing, or manually by people (usually patients) undergoing testing for diagnostic or screening purposes, and the completed forms are digitised. This data includes audiograms, dementia memory tests, cognitive tests, and other tests whose results include media other than text.

2.1.2 Terminological notes

- In this guideline, references to activities conducted by the HDAB may also encompass activities carried out by a processor acting on behalf of the HDAB, such as the operator of a SPE. This is consistent with Article 28 (GDPR) and Article 73 of the EHDS Regulation. In such cases, the HDAB remains responsible for ensuring full compliance with the applicable legal framework and must implement appropriate contractual and technical safeguards to govern the processing by the processor.
- The term *health data user* refers to an organisation or natural person who has been lawfully granted access to electronic health data for secondary purposes in the context of the EHDS Regulation (in line with Article 2(2)(u), EHDS). In this deliverable, *health data user* may also refer to the individual staff member acting on behalf of the organisation processing the data as employee of the data user organisation, provided there is no ambiguity. Please note that instead of the defined terms *health data user* or *health data holder*, shortened versions are also used, i.e., *data user* or *data holder*.

3 Data minimisation

Data minimisation is a fundamental principle under Article 5(1)(c) of the GDPR, requiring that personal data be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed"[2].

In the context of secondary use of personal data as governed by Chapter IV of the EHDS Regulation, data minimisation and purpose limitation must be applied throughout the entire data lifecycle including:

- during data collection and preparation by the data holder,
- when assessing the data request or data access application (by the HDAB and potentially the trusted data holder),
- during the use and processing of data by the data user, including export of results from the SPE.

Purpose limitation in the context of data minimisation implies that HDABs should assess whether the purposes stated by the applicant are sufficiently specific and compatible with the EHDS framework, and where needed, request clarifications or a narrowing of scope. This should be documented clearly and reflected in the data permit conditions under Article 68(3) EHDS.

Data minimisation may include:

- reducing the **volume of data** made available,
- limiting the **granularity** and **sensitivity** of variables,
- restricting **temporal or geographical scopes** where appropriate.

It forms part of the broader risk mitigation strategy that includes pseudonymisation, anonymisation, and organisational safeguards. It also supports compliance with other regulatory requirements, such as:

- the protection of intellectual property and trade secrets (Article 52, EHDS), when relevant.
- the exclusion of public interest risks from permitted uses (Articles 68(2) & 69(3), EHDS).

Within their risks evaluation strategy towards data provision, HDABs and data users must consider not only re-identification risk but also inference attacks. Attribute inference is described as the process of inferring unknown attributes of known individuals from the data. Membership inference refers to inferring the presence of known individuals in the data (see section 5). Attacks on attribute inference, group membership inference, linkage attacks should be considered during the risk evaluation before data provision. While data minimisation is primarily a data protection principle (Article 5(1)(c) GDPR), it can also reduce the attack surface by limiting the volume and granularity of data processed. Security of processing remains governed by Article 32 GDPR and the EHDS requirements for secure processing environments.

[2] See Article 5(1)(c) (GDPR).

While minimisation is not an active measure of enhancing data security control, it significantly reduces exposure related to:

- confidentiality (unauthorised access),
- integrity (unauthorised alterations), and
- availability (data not accessible to authorised parties).

Hence, data minimisation indirectly supports data confidentiality, integrity and availability, by limiting the amount of the to be exposed data. However, the GDPR triad (confidentiality, integrity, availability) must be clearly distinguished from broader data protection risks and traditional IT security domains (e.g., authenticity in Recital 49, GDPR).

Under Article 68 (EHDS), HDABs shall assess whether there is sufficient justification that pseudonymised data should be used instead of anonymised data (see Article 68(1)(c)). Regardless of the legal status of the data, data minimisation applies to all phases of data user's journey and the actors involved (data holders, HDABs, data users), including within and beyond the SPE. In fact, data minimisation also applies after the data is made available, including during processing by the data user, during analysis and results export.

If the HDAB provides a response to a data request, the data will be processed and the data made available to the data user will be in an anonymised statistical format. The detailed procedures should be aligned with the user and documented in the data request approval. The HDABs must pay attention to which relevant data elements should be provided to data users, based on their stated purposes, that should not be exceeded. More information and insights on the procedures and techniques to produce aggregated data within data requests are offered in section 5, i.e., Anonymisation chapter.

Regarding data permits, data minimisation procedures and techniques that can be used by HDABs and data holders will be further described in this paragraph and should be considered by applicants in the data access phase of the EHDS user journey (Figure 1).

3.1 When should data minimisation be performed?

Controllership. In general, a data controller is a person or organisation that determines the purposes and means of processing personal data, as regulated by the GDPR. According to Article 4(7) (GDPR), the data controller is responsible for ensuring that data processing complies with the principles set out in the regulation, including data minimisation. The EHDS Regulation specifies (Article 74) the controllership rules applicable to data processing workflows within the EHDS ecosystem. Generally, data holders act as controllers for the initial processing and provision of data to the HDABs. This condition may occur when pseudonymised or anonymised data are included in the data permit, or when HDAB will perform data aggregation of personal data to be offered in an anonymous statistical format. HDABs act as controllers for processing personal data within the scope of their tasks indicated by the EHDS Regulation (Article 57, EHDS), such tasks include receiving data from data holders, for activities like data linkage, additional pseudonymisation that may be needed, additional risk reduction or anonymisation, upload and provision to the SPE. Data users are deemed controllers within the limits of the data permit and the regulatory framework for the data that has been made available to them in a SPE following an approved permit (see also D7.1, TEHDAS2).

Communication channels. Specific considerations on data minimisation should always be made by data applicants, because the EHDS Regulation requires the HDAB to ensure that the requested data are limited to what is necessary for the permitted purposes. HDABs will have to assess compliance with data minimisation principles of the data access application or the data request received, before issuing the data permit (see Article 68, EHDS) or approving the data request (see Article 69, EHDS). If compliance is not sufficiently demonstrated by the data applicant, some delays in the process, a changeover from a data permit to a data request (Article 68(3), EHDS) or even a rejection may occur. For that reason, it is important to establish an effective communication channel between the data applicant and the HDAB, that could resolve doubts or controversial points specific to data minimisation early in the process (see also D6.2, TEHDAS2).

Another communication channel may be needed between HDAB and data holders, to clarify some aspects of data minimisation, especially when the application is incomplete (see D6.2 Guideline for data users on good applications and access practice, section 8.5) to manage the amount and the granularity of information requested by the data applicant.

Further ahead in the user journey, data users shall take data minimisation principles into consideration during the use of data in the SPE and when requesting data exports from the SPE (see more in section 5 of this document).

Summing up all those considerations, data minimisation shall be applied by different actors and most importantly during the phases presented in Table 1. See also Figure 1 for a more general overview.

Table 1. Role-based involvement for data minimisation in the EHDS. The list below addresses data access applications, while similarities for data requests hold. For simplicity reasons, trusted data holders and trusted third parties are left out from this table. For each data-minimisation related activity, actors' roles are distinguished by the following notation: A/R = Accountable & Responsible; C = Consulted; I = Informed.

User's journey phase/activity	Data User	HDAB	DH	Relevant Data Minimisation activities
Data Discovery				
Publishing Datasets Metadata		**I**	**A/R**	Specifying relevant information on published datasets that may be useful for the next phases of data user's journey. (please find more information in deliverable D5.1)
Exploring Useful Datasets	**A/R**	**C**	**C**	Communicating with the HDAB and the DHs about the characteristics, availability, and granularity of information, or for any clarification issues.
Data Access				

Data Application Preparation	**A/R**	**C**	**C**	Considering relevant data minimisation actions on tables, variables, granularity levels, quasi-identifiers. HDAB and DHs may be involved in helping Data user refining their applications.
Data Application Submission	**A/R**	**C**		Documenting relevant data minimisation actions on tables, variables, granularity levels, and quasi-identifiers.
Data Preparation				
Data access application evaluation	**C**	**A/R**	**I**	Engaging with the data user to assess whether data minimisation has been adequately demonstrated, and, if needed, formulating adjustments to ensure compliance with Articles 66–68 EHDS.
Data Permit / Data Request Issuing	**I**	**A/R**	**I**	Deciding on Data Request / Data Permit issuing, also considering data minimisation aspects.
Data Extraction pursuant to a Data Permit		**C**	**A/R**	Verifying the availability and accuracy of information specified in the data permit issued. Carrying out data extraction activities from the relevant datasets to match the requirements stated in the Data Permit. Whenever agreed with the HDAB, some formatting, data minimisation and/or re-identification risk reduction activities may be performed by the DH during this phase.
Data Provision towards the SPE		**I**	**A/R**	Providing datasets to the SPE or the specific data repository agreed with the HDAB.
Data Preparation before granting access to data in the SPE	**I/C**	**A/R**	**I**	Verifying what is received from the DHs and carrying out more data preparation activities that may be needed (i.e., anonymisation and/or further pseudonymisation and minimisation activities, data linkage etc...) Any deviations, inconsistencies, and data quality problems identified during the data preparation shall be communicated by the HDAB to the data user.
Data Access Grant and Data Ingestion to the SPE	**I**	**A/R**		Preparing the SPE and granting access to Data Users.
Data Use				
Data Cleaning & Transformation	**A/R**			Applying further data transformations that may be needed for the project, including minimisation, cleaning, and validation activities.
Data Project Execution	**A/R**			Carrying on the project activities.

Finalisation				
Preparation and request for result export from the SPE	**A/R**	**C**		Proposing the structure of result export to the HDAB. Data minimisation and study purposes should guide the output attribute selection.
Approval of result export from SPE	**I**	**A/R**		Evaluating the result export proposal.
Data export		**A/R**		Downloading a data export compliant to what is approved by the HDAB.

- The GDPR emphasises the necessity to manage data minimisation activities also at the data holder's side before any data permit has been issued (i.e., at the data collection phase and when pre-processing data to be subsequently shared). However, those activities are out of scope for this document. Nonetheless, one point to be cited here on data minimisation activities during data collection and data curation is the following: Data holders should make publicly available relevant details on datasets for secondary use as required by the regulation (see art. 60(3) and art. 77(2), EHDS). If there are default data minimisation measures, it is recommended that data holders specify any standard data minimisation rules that are applicable by default to the datasets offered. For instance, a data holder may specify that the information on healthcare operators, e.g., *clinicians*, *general practitioners* etc. is normally encrypted; that *birthdate* is normally available in terms of year (*Year of birthdate*). Any deviation from the default standard offered by the data holder may be checked before data access application submission and should then find clear justification in respect to the study objectives in the data permit (see also D5.1 Guideline for data holders on data description).

3.2 Direct Identifiers and quasi-identifiers

Within the context of human subjects' studies, **direct identifiers** are variables that point explicitly to specific individuals and are sufficient to identify the data subject either alone or in their mutual combination (e.g., names, national id numbers, social security numbers, telephone numbers, email addresses, fingerprints). As widely discussed in the EDPB guideline on Pseudonymisation[3], effectively pseudonymised data do not contain direct identifiers as they are stored separately as part of the "additional information". Pseudonymised data provisions, however, in some cases may contain variables whose combination is sufficient to attribute at least part of the pseudonymised data to identifiable data subjects. Those attributes are called quasi-identifiers (or indirect identifiers[4]). Some examples of quasi-identifiers are age, date of birth, sex, ethnicity, education, employment status, marital status, income, place of residence or work/study, or a sequence of hospital visit dates. In the context of employee data, relevant quasi-identifiers may include structural roles, number of working hours, and length of service.

Quasi-identifiers may be considered critical variables that increase the assessed privacy risk (including the re-identification risk) and should be removed from the data provision, unless they play a crucial role in data usage. When they play a crucial role, only strictly necessary quasi-identifiers should be retained, and their risk mitigated, in line with the

[3] See EDPB Guidelines 01/2025.
[4] See EDPB Guidelines 01/2025.

GDPR's data minimisation principle. Appropriate risk mitigation techniques may be applied to reduce the likelihood of re-identification, such as generalisation, suppression, or randomisation. Quasi identifiers may sometimes emerge from linking different datasets from different DHs. This circumstance should be verified by the data user before the submission of the access application and managed consequently.

In the GDPR (Recital 26), directly or indirectly identifiable data is data that can identify a natural person by reference to an identifier or more factors (Article 4(1), GDPR).

Reducing re-identification risk typically requires some kind of transformation in the original datasets. Generalisation is the most used method and involves reducing the level of detail (e.g., using age brackets instead of exact age). Suppression involves removing or hiding specific cell values. Randomisation[5], which introduces noise to diminish predictability like any other forms of non-deterministic methods should be carefully calibrated because they may significantly reduce data utility. Non-deterministic techniques like noise injection are more commonly used in anonymisation than in pseudonymisation contexts, especially under EHDS which prioritises data fidelity. See more in the next chapters.
Some examples are as follows:

- An example of generalisation is when income data are grouped into ranges (e.g., €20,000-€30,000) to prevent identification.
- An example of suppression is when the values of cells that represent a small number of contributing entities are suppressed, for instance the extreme values in a lab test are removed. For aggregated data, an example of suppression is removing the number of patients discharged from a small hospital for a low-incidence disease.
- An example of randomisation is altering the number of working hours by adding or subtracting a small random value to/from each entry.

Although a case-by-case approach may be adequate, some recommendations and general principles to be respected will be further discussed in the following paragraphs.

In the EHDS, when data access to a pseudonymised or anonymised dataset is provided, the HDAB should implement appropriate measures to reduce re-identification and inference risks, taking into account the means reasonably likely to be used (Recital 26, GDPR) and the EHDS safeguards. The data user must refrain from any attempt to re-identify individuals (Recital 72, EHDS).Therefore, it should be clear that ensuring adequate protection against re-identification is a shared responsibility between the HDAB and the data user.

The EHDS also states that HDABs should apply tested state-of-the-art techniques that ensure that data processing preserves the privacy of the information contained in the data, including methods such as generalisation, suppression or randomisation (Recital 65, EHDS).

3.3 The dimensions of data provision

Whenever a data permit is issued by an HDAB, meaning the HDAB has determined that access to individual-level data (pseudonymised or anonymised) is justified and appropriate based on the purpose and privacy risk assessment, a data provision with individual-level data will be offered to the users in an SPE. It is possible to identify some specific dimensions of

[5] See EDPB Guidelines 01/2025.

such data provision, as described in Figure 3. It should be noted that access to data via a data permit is not necessarily equivalent to granular pseudonymised data, but can be restricted, depending on proportionality and risk mitigation principles.

Figure 3. Data provision dimensions.

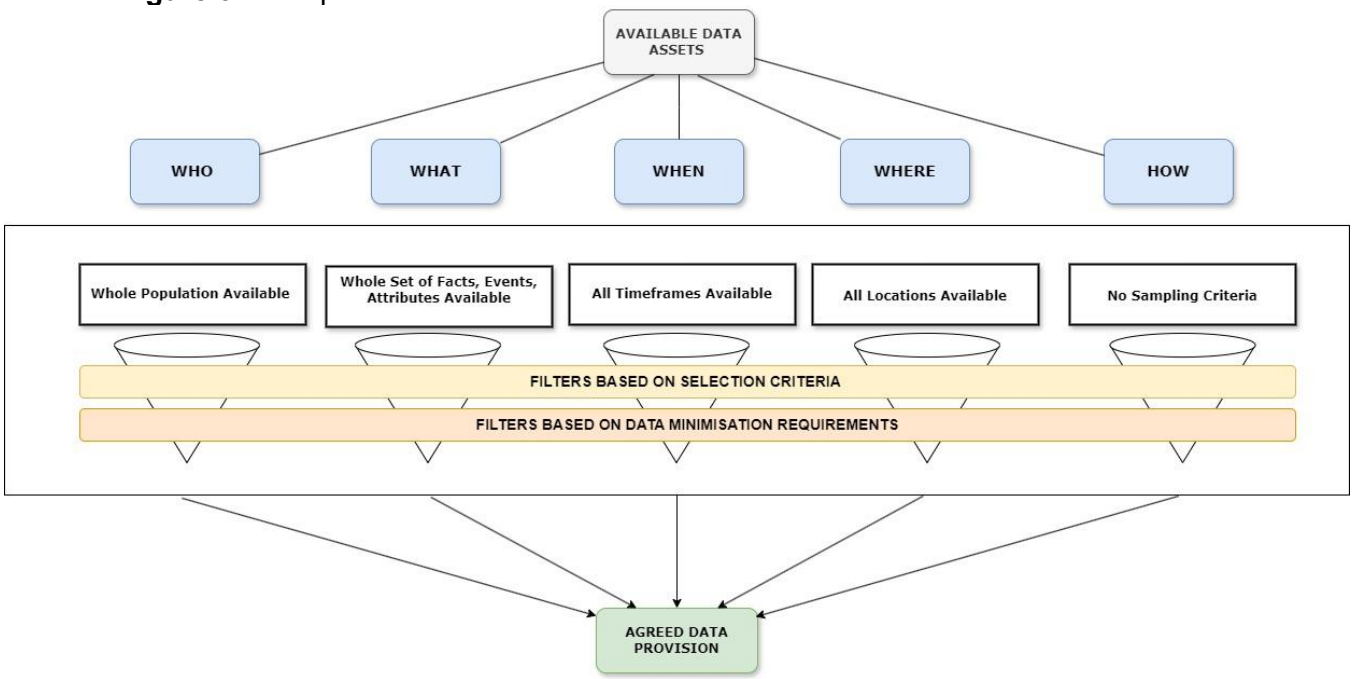

During the data discovery phase (Figure 1), the data applicants would check data assets that are available from specific data holders of their interest and would select some of them that can work well in addressing a limited number of predefined secondary use objectives. Data applicant purposes and objectives for secondary use studies can be very different between each-other; they may be focused on hypothesis-driven, descriptive, predictive, exploratory or explanatory endeavours. In any of these cases, **information from available data assets** can be distinguished for practical reasons into the following **five dimensions**:

- **Who** - Which individuals are relevant for the study purposes (Study population), from the population;
- **What** - What information is relevant for the study purposes (Study variables), from the whole set of facts, events, attributes available;
- **When** - Which timeframes are relevant for the endeavour (Study timeframes), from the whole available periods of time;
- **Where** - Which specific locations, e.g., place of residence, work or health assistance, are relevant for the study purposes (Study geographical perimeter), from the whole geographical distribution available;
- **How** - What extraction methods are used to get samples of interest (Study extraction methods), from the entire available sets of information.

Please also see D6.3 Guideline for Health Data Access Bodies on the procedures and formats for a data access application template (Annex 5), where the five dimensions are also used to help applicants adhering to the data minimisation principle.

An overarching further dimension that should guide the data user's inclusion/exclusion decisions on selected data is the **"Why"** dimension – Why a single piece of information is

selected to be included in the study endeavour. This dimension acts in terms of justifying the appropriateness of data minimisation decisions for the sake of the study endeavour. It is not further directly described in this document but has some explicit reference in all the other five data dimensions, and finds direct use in the data access application, because it provides the HDAB with crucial elements to evaluate the appropriateness of choices made by the data user (see more in D6.2 - Draft guideline for data users on good application and access practice – specifically in the sections of the access application form).

When selecting multiple pieces of information for the study endeavour, the cross-classification of multiple variables is the real issue to be tackled. In a dataset containing a medium to high number of variables, a large proportion of records will be unique, due to the cross-classification of their values. Moreover, in the case of quasi-identifiers, it could also occur a risk of singling out (which corresponds to the possibility to isolate some or all records which identify an individual in a dataset) due to combinations of different personal information (see Opinion 05/2014 on Anonymisation Techniques).

Based on the study objectives, applicants should identify the necessary information and apply inclusion and exclusion criteria accordingly. This process may reduce the data requested across the five data provision dimensions. Data minimisation principles and data protection risk management should inform the exclusion of information not directly relevant to the study or that presents additional privacy risks. Applicants must justify their data selection decisions ("Why") across all dimensions.

The result of those activities, that imply a mediation between the study objectives and the data minimisation requirements, which may require negotiation between the applicant and the HDAB (or even the data holder in particular cases, see more in the next paragraph), will end up defining data provision characteristics, and the specificity of each piece of information required by the applicant, in all five dimensions of data provision.

A general consideration that applies to all those dimensions is that it must be stated in each data permit (Art. 68(10), EHDS) which are "the *categories,* ***specification*** and *format* of the electronic health data to be accessed". Some specifications are consequently considered mandatory in the data access application form and need to be justified by the applicants as required for reaching the study objectives and yet compliant with the data minimisation principles. The access application form, as well as the data request template can be inspected in Annexes 5 and 6 of the M6.3 guideline.

Practical strategies to enhance the data provision may include the following where feasible:

- It is possible to provide a certain percentage of the data during data provision (e.g., 1 million of data subjects, or 10 years of data) that allow data users to build and test their analysis scripts and elaborate on intermediate results, then executing the final scripts on the complete dataset applied for.
- It is possible to provide pre-defined granularity levels by the data holder that the applicant can consider.

Those ways of working are not generally required and should be made clear as early as possible, so that applicants may know this possible chance in advance and be aware of this already during the data discovery phase.

In the following paragraphs, special attention will be given to each data provision dimension.

3.3.1 The “Who” dimension

Explicit study population definition support both data minimisation and purpose limitation. Specification of data subjects included in the study population is mandatory when individual-level pseudonymised or anonymised data are requested by the data applicants. Applicants should clearly define their study population, and the relevant inclusion and exclusion criteria from the whole population included in the datasets offered by the data holders[6]. This decision will affect which subjects will be included in the data provision. Moreover, eligibility criteria are also critical for enabling the HDAB to assess necessity and proportionality of the requested sample. Common inclusion and exclusion criteria normally act jointly on demographic characteristics of subjects and on study-specific variables of interest like diseases, exposures, treatments, and comorbidities. All eligibility criteria specified by the data applicant would define the study population.

An estimate of the expected cohort size should be provided, whenever possible, by the applicant in the application phase, along with the inclusion and exclusion criteria used to derive this estimate from the data assets of interest. This information serves to support the HDAB in assessing the proportionality and feasibility of the requested data provision. For instance, the applicant may state: 'We apply inclusion criteria A, B, and C to dataset X, which contains approximately 120,000 individuals as reported by the Data Holder. Based on prior studies, we estimate that 5,000 individuals will match our criteria.' The HDAB may then verify this estimation against the structure of the dataset and, if needed, request clarification or propose adjustments (e.g., stratified sampling, changes in inclusion logic). Such a check does not imply that the HDAB re-derives the full cohort itself but helps identify inconsistencies or misalignments between the stated study population and the available data.

It should be noted that most demographic variables typically fall into the category of quasi-identifiers and hence require some remarks as follows:

- ***Date of birth*** and ***Date of death*** are strong quasi-identifiers and should be avoided whenever possible. They can be substituted with other relevant variables like year of birth, age group or age at death. For age-sensitive studies year/month of birth, year/month of death, or even days at hospital discharge for children might be relevant.
- ***Age*** is a quasi-identifier that requires attention. Age should only be used as an exact value when it is indispensable for the research question and cannot be replaced by age categories, i.e., generalised, without loss of analytical validity. In some cases, where knowing the exact age is fundamental for the analysis, it may, from a data minimisation perspective (to reduce privacy risk), be relevant to categorise the extreme values where the number of subjects reduces considerably (e.g., ages from 0 to 5 and above 85 are grouped together).
- ***Nationality, ethnicity, education, profession, employment status, marital status, income, state of vulnerability*** (persons under guardianship, trusteeship) are quasi-identifiers that require attention and should be used only when strictly relevant, taking the precaution of standardising, classifying and/or grouping their possible values into the lowest level of detail that is useful to meet the study objectives.

Please note that, beside variables concerning health, genetic or biometric data, if any of the other requested demographic variables can reveal a natural person’s racial or ethnic origin,

[6] see also TEHDAS2 D6.2 Guideline for data users on good application and access practice.

political opinions, religious or philosophical beliefs, trade union membership, sex life or sexual orientation, they may be classified as special category data under Article 9(1) (GDPR).

Moreover, **combinations of demographic variables**, like “date of birth plus sex plus postcode” should be avoided because they significantly increase the re-identification and privacy risk of the data subject.

3.3.2 The “What” dimension

Specification of datasets, tables, objects and data elements needed to pursue the aim and topic of the application is mandatory. During the data discovery phase, applicants should invest some effort to understand what is available from a specific data holder of their interest, and which information is relevant for their study objectives. In accordance with Article 5(1)(b) of the GDPR, personal data must not be made available for undefined or overly broad secondary use purposes. Instead, applicants should request only the minimum data that is necessary to accomplish the clearly defined secondary use objectives specified in their data access application. Variables to be requested can be divided in two broad categories:

- **Study-specific variables:** the ones that are functional to answer research questions and address study objectives, and that may be linked with patients, diseases, characteristics of diseases, treatments, exposures, comorbidities, disabilities and so on.
- **Control & confounding variables**: the ones that are needed not to address the study objectives, but to enhance the validity of the study, giving the researcher the ability to control external and extraneous influences on the observed outcomes.

Data applicants are strongly encouraged to explicitly mention, during the application phase, to which of these categories each requested variable refers to. Separating variables in these two categories may help HDABs assess the relevance of the data application and its data minimisation compliance. In some cases, a transformation in the type, format, level of detail or coding of a variable made available by a data holder may be asked by the applicant in the data access application form or suggested by the HDAB for data minimisation compliance. In such cases it is important to define who will be in charge of applying those transformations, that virtually can be done both at data holder’s or HDAB’s side and could influence the production and costs of data provision (see more at paragraph 5.4 and in the Annex 6).

The ”What” data provision dimension consists of a list of objects and data elements that are strictly useful for the expected data usage, and do not fall in any other dimension like “who”, “when”, “where”, or “how”. The “what” dimension can sometimes manage direct and/or indirect identifiers. Overall, the higher the number of indirect identifiers requested by applicants, the higher the risk of re-identification. For more clarification on the EHDS categories of data for secondary use see section 5.3 in D5.1 Guideline on data description. They require some general remarks:

- ***Medical images*** produced by healthcare procedures may reveal unique body features of individuals with specific diseases (e.g., head and neck cancers). Alternatively, histopathological imaging which is used in cancer diagnostics (so called Whole Slide Images or WSIs) contain detailed cellular structures in order of Gigapixels, which are also unique (but not well reproducible as the biological material has been removed from patient body). This uniqueness may contribute to increase the risk of re-identification of data subject’s identity. Depending on the case, those images could be considered as direct-identifiers or quasi-identifiers and should be treated consequently – they can either identify the patient (facial imaging) or can act

as linkage across data sets if the same image is used in different datasets. It is important to stress that evaluations depend on the specific case and the re-identification risk.

- ***Genetic data*** require special attention because depending on the study endeavour may represent direct identifiers or quasi-identifiers. Genetic data are classified as sensitive personal data under Article 9(1) (GDPR), and special safeguards may apply to them.
- ***Rare disease codes*** can be considered quasi-identifiers. Therefore, when researching rare diseases, attention must be posed to remove or generalise from data provision other quasi-identifiers. Practically, the higher the number of indirect identifiers requested by applicants, the higher the risk of re-identification, and consequently, the more importance should be given to risk reduction (e.g., anonymisation, differential privacy techniques) and risk management procedures.
- ***Unstructured data*** always needs particular attention, because attributes selection is difficult to manage. Thus, unstructured data potentially includes more background information on facts and people being investigated than what is strictly necessary to data usage purposes. Also, the activities that may be performed to remove direct and quasi-identifiers (e.g., searching and removing or obfuscating them) are less deterministic compared to the ones performed for structured data.

3.3.3 The "When" dimension

Specification of the data provision timeframe is mandatory. It may regard years, or months, or days, and both *absolute* and *relative* timeframes are possible. An absolute timeframe specification may specify a given calendar period, e.g., from the year 2020 to 2025. A relative timeframe specification may be relevant for longitudinal studies, where the timeframe can be aligned with the patient inclusion period (e.g., +/- 1 year after/before hospitalisation, +-30 days from intensive care unit admission, etc..). Specifying the timeframe does not mean that it cannot be the maximum available, it only stresses out the necessity to align it to the study purposes/objectives, and so a clear justification of the timespan is required. Applicants are expected to provide considerations on the granularity of time-specific variables also, because the granularity of temporal data (especially when requesting day-level information or even shorter timespans) is directly linked to privacy risk and must be justified accordingly. For instance, if the data applicant is requesting information on the incidence of a disease for a selected population, it should be clear enough if the year of incidence (instead of the month or the exact date) is relevant for the study objectives.

Quasi-identifiers that fall under the "When" category of data provision dimension require some remarks:

- ***Dates*** are often critical variables, especially when combined with other information. Some consideration on dates linked to patient's demographics have already been discussed in paragraph 5.3.1. Other relevant dates available in a data provision may be time-specific ones. For instance, hospital admission and discharge dates or treatments dates. Those dates may point at a specific episode of care in the life of a person, that linked with other information may reveal data subjects' personal information or their identity. Therefore, absolute dates are variables that should be avoided when not essential and replaced by relative timespans (e.g., time since diagnosis, age at event), which may often be safer and sufficient for most analyses. Possibly, year/month of hospital admission and length of stay may be considered valid substitutes of hospital admission and discharge dates.

These considerations align with the GDPR's overarching principles of confidentiality (Article 5(1)(f)) and with recommendations under EDPB Guidelines 01/2025 for controlling re-identification risk in datasets containing high-dimensional or temporally specific variables.

3.3.4 The "Where" dimension

Specification of geographic constraints applicable to the study population should be indicated by the applicant in the data access application form submitted to the HDAB (see Annexes 5 and 6 of Guideline M6.3, TEHDAS2). Geographic attributes may refer directly to patients or to other relevant entities like healthcare facilities, local health authorities, multicentre clinical studies (i.e., site-specific vs pooled analysis) and so on.

The geographical dimension can directly affect the amplitude of data provision.

Selecting a large study population may be beneficial, especially when machine learning models are going to be used but may also increase the need to control for confounding variables when geographical differences imply differences in the health determinants and/or in the healthcare service characteristics and consumptions. Choosing a very large population may also increase the data minimisation burden, as it increases the probability of providing records with unique characteristics that may be correlated with identifiable data subjects and/or that may be linked with external available sources of information (i.e., using geotags or geocoding[7]). A very small population, on the other hand, may increase the privacy risks because their demographic and geographic variables may easily point out toward a data subject that is known to be included in the dataset. In both cases, a general guidance can be to pay close attention on the variable selection. Next to population size, the granularity of the location data is of great importance and should always be generalised by default unless granular data is clearly justified and proportionate.

Geographic attributes are generally easily generalisable without reducing the utility of data provided. Examples include:

- ***Place of residence, of work, of health assistance***, being critical variables, especially when combined with other information, should be provided at the highest possible level, like a ZIP code with only 2 digits instead of 5, or the region, or the local health authority. On the other end, data projects that are very specific in terms of geographic scope (for instance, comparing the catchment area of the city hospitals for specific procedures) should be critically assessed and should offer very controlled variables (for instance, ZIP codes could be provided with 4 or 5 digits but most of the other demographic variables would be omitted or generalised, or otherwise specific variables like "distance from workplace" or "distance from residence" could be calculated and offered to data user).
- ***General Practitioner*** associated with the data subject should always be pseudonymised and may give a strong local attribution if linked with other information included in the data provision like general practitioner's visits.
- ***Health Facilities*** linked to data subjects' episodes of care, having a known geographic location, may sometimes require attention and may suggest using some form of obfuscation (i.e., substituting identifying facilities codes with sequential numbers).

[7] See DOI: 10.3389/fsoc.2022.910111.

3.3.5 The “How” dimension

Specification of how the general population available from a data holder should be sampled and/or managed to obtain the final data provision may be included in the data access application form.
In some cases, data applicant may ask to apply specific data cleansing methods to the identified target population to increase data quality or data analysis (remove redundant records, apply standardisation procedures to specific columns, apply specific rules to treat missing information etc..). However, these methods may conveniently be applied once data provision has been issued by data users themselves.

In other circumstances, modifying the consistency of the whole population available at data holder side may be needed for data minimisation necessities (i.e., removing outliers, aggregating information for low-frequency groups when they do not influence the study objectives).

Sampling procedures that produce statistical samples from the identified target population are also possible (i.e., the applicant want to conduct a quicker preliminary study on available data before starting the main study).

In any of those situations it is required that:

- The principle of data protection by design and by default (Article 25, GDPR) should guide sampling decisions, particularly in how subsets of the population are chosen and filtered.
- Sampling methods must preserve representativeness and avoid introducing bias.
- Data applicants should be transparent on methods to be applied in the data application form.

3.4 Data minimisation steps towards issuing a data permit

The data minimisation principle applies at all phases of health data access application.

One of the primary responsibilities of HDABs, as illustrated by Article 57 (EHDS), is to issue data permits, based on specific criteria, as expressed in Article 68 (EHDS). During the data access application assessment phase, HDABs are legally responsible for ensuring that the scope and granularity of requested data are limited to what is necessary. Note that trusted data holders may also perform the first application assessment.

This section outlines the main procedural steps HDABs should implement to ensure that only data that are necessary for the stated study purposes are made available, and that personal data are minimised at every phase of the permit process.

The main steps that HDABs should put into effect towards the goal of **issuing data permits and requests** are as follows[8]:

[8] TEHDAS2 D6.3 Guideline for Health Data Access Bodies on the procedures and formats for data access.

1. **Completeness and relevance check**: Review the data access application to ensure that all required information is present, and that each requested element is clearly justified, relevant to the secondary use objectives, and aligned with the data minimisation principle to ensure that all required information is present, and that each requested element is clearly justified, relevant to the study objectives, and aligned with the data minimisation principle.
2. **Data access application assessment**: The HDAB should examine the content of the application in detail; identify critical elements that may pose minimisation challenges; clarify applicant choices and goals where needed and initiate communication with the data applicant and/or the data holder to resolve ambiguities and reduce unnecessary data demands.
3. **Decision and outcome**: Based on the assessment, the HDAB should either issue the data permit, request specific revisions to ensure compliance with the minimisation principle, or propose a changeover to a data request (e.g., if only aggregated data in a statistical format are appropriate).

3.4.1 Data access application assessment

In evaluating a data access application (after a successful data access application completeness check, see Annex 7 in M6.3), HDABs must ensure that all requested data are strictly necessary for the stated data usage objectives and comply with the principles of data minimisation and purpose limitation (Articles 66 & 68, EHDS; Article 5(1)(b–c) GDPR). This assessment consists of two layers:

Layer 1. HDABs verify whether the data access application justifies access to the data applied for. The HDAB examines whether the applicant can get access to anonymised or pseudonymised data (Article 66). They also examine (possibly with the help of the data holder) whether risks related to public interest, intellectual property, or trade secrets are implicated (Articles 52, 68(2), and 69(3), EHDS).

Layer 2. Once these general conditions are considered, a more **detailed examination** is conducted. This includes reviewing the proposed data provision across all five dimensions ("Who", "What", "When", "Where", and "How") to ensure that:

- each data element is relevant and justified for the study purpose;
- the granularity and sensitivity of the data are proportionate;
- unnecessary or overly detailed data are excluded or transformed.

Applicants should take a proactive role in anticipating potential data minimisation concerns, and structure their requests accordingly. HDABs may engage in dialogue with applicants and/or data holders to refine the request and ensure regulatory compliance.

The assessment of a data access application involves both **general** and **domain-specific considerations**. The following section first outlines cross-cutting criteria that apply to all applications, regardless of data type or research domain. It then provides tailored recommendations for specific types of data, such as structured datasets, images, genomic datasets, and unstructured data, which require additional attention due to their unique privacy or minimisation implications.

The following considerations also apply in the case of data permit amendments (for instance upon the user's request for additional variables or changes in scope) which may occur after

the initial permit is granted. The HDAB must reassess the relevance and proportionality of any such modifications before approval and remains responsible for data minimisation and necessity assessment when such modifications are requested.

General considerations specific to data minimisation and purpose limitation that should be made during the detailed examination of a data application assessment include:

- **Revising** the specificity of data usage objectives. Significantly different objectives should be managed with different data applications and data provisions.
- **Revising** the appropriateness of the study population to answer data usage objectives, in terms of number and characteristics of the data subjects for which personal data are expected to be included in the data provision.
- **Verifying** that no direct identifiers are asked to be provided, and that quasi-identifiers' inclusion is based on solid motivations and pertinent risk reduction safeguards. (Please note that in the case of data linkage, direct identifiers used for linkage should not be disclosed to the user and should be handled by HDAB/DH/TTP under strict safeguards).
- **Evaluating** the legitimacy, proportionality, and granularity of quasi-identifiers and the potential increases in re-identification risk that may arise from their combination.
- **Exploring** all five dimensions of the expected data provision to check if justification for inclusion/exclusion of all variables is available and appropriate.
- **Evaluating** the impact of cross-classification of multiple variables: for instance, in a dataset containing a high number of variables, a large proportion of records can result to be unique, due to the cross-classification of their values. In these cases, there is a trade-off to be managed between the limitation of re-identification risk and the retention of utility for the expected data provision.

Domain-specific considerations on data minimisation and purpose limitation that should be made during the detailed assessment of a data access application may include:

For **datasets** and **data elements**:

- **Limiting** the **tables** requested to what is strictly necessary to the data usage objectives, in terms of tables that contain study-specific and/or control/confounder variables only.
- **Limiting** the number of **columns (variables)** to what is necessary to data usage objectives. For each variable that is not directly justified as related to data usage objectives, HDAB will ask the applicant for a revision of the application.
- **Confirming or changing the provision of columns (variables).** The use of variables should be assessed considering data minimisation and purpose limitation issues as well as data usage objectives. Variables that are not linked to study objectives and cannot be considered as control/confounder variables should be excluded. If different variables that point to similar kind of information is available at data holder’s side, these variables should be evaluated as alternatives in terms of utility and risks. HDABs could propose data transformations for the purpose of minimisation (e.g., replacing detailed dates with derived indicators like hospital length of stay). This activity is typically limited to predefined variables and options already made available through data catalogues. Unless otherwise agreed with the data holder, custom derivations are not required.
- **Confirming or changing the specificity of information** (i.e., the structure and granularity of values in the columns). HDABs will confirm data applicant’s choices or offer provision of more aggregated information (e.g., variables with categorised data

or at a higher geographical level) or diluted information (e.g., suppression of row or, randomisation). In publishing data catalogues, data holders may pro-actively anticipate the available levels of specificity for each variable provided, to allow data minimisation considerations during the data discovery phase.
- **Confirming or changing cross-linkage between tables**. Only tables that are explicitly intended to be linked should be cross-linked within the data provision, to keep re-identification risk controlled.

For **images**:
- **Limiting metadata** to what is strictly necessary to the secondary use objectives. Metadata should not include direct identifiers. Technical metadata should be removed too, if not explicitly required.
- **Digital blurring** of specific image areas (i.e., where personal identifiers or quasi-identifiers are present).
- **Digital masking** of specific image areas (i.e., when only part of the image is relevant or must be highlighted for the intended purposes).
- **Reducing** the **resolution** of the image whenever is relevant, to what is pertinent to data usage objectives.

For **genomic datasets**:
- **Limiting metadata** to what is strictly necessary to data usage objectives. Technical metadata should be removed too, if not explicitly required.
- **Filtering** for **specific regions** of the genome (i.e., give access only to a limited set of chromosomes).
- **Filtering** for relevant **variables** (i.e., if the analysis concerns only some mutations - like BRCA1 for breast cancer - only those variants can be preserved from the entire genome).
- **Preparing files** in the format most suitable for analysis. Raw genomic data is generally avoided unless clearly necessary and explicitly justified.

For **unstructured data** and **qualitative research**:
- **Evaluating** the data collection approach and its sensitivity to minimise personal data. It might be helpful to use sample data from data holders.
- **Pre-processing** text files to remove direct identifiers, quasi-identifiers and data attributes that are described in the application as not linked to data usage purposes. For example, this may mean using regular expressions ("regex") to locate identifiers in free text, and to base these regular expressions on public databases such as a list of common surnames, or the exhaustive list of European communes or ZIP codes. The length and format of the text snippet to be extracted should be specified (ideally per variable) (e.g., detecting sequences of 13 digits in free text to identify a social security number). Automated open-source text minimisation applications are emerging and may help in making available for text mining only relevant information.

Having assessed in detail the structure and the content of the application, HDABs will end up with different possibilities:
1. to issue a data permit;
2. to refuse the data access application;
3. to propose a revision of some critical aspects or gaps that need to be addressed to accept the application (see Article 68(3));

The first two cases are straightforward: HDAB will take its decision and document it. Refusing the data access application may lead to a new data request submission by the data applicant, when HDAB propose aggregated data instead of individual-level data. The third case is indeed the more complex scenario to be managed by the HDAB and may require communication with the data applicant and data holders. After the detailed analysis of the application, the HDAB might have a list of the gaps that need to be addressed to accept the application, for instance:

- requested objects, tables and variables match the study objectives, quasi-identifiers are justified and proportionate, and safeguards to reduce re-identification risk are in place, yet some of them or their cross-combination is considered critical and need to be managed further (**Gap 1**).
- some of the requested objects, tables or variables do not seem to match the study objectives or require clarification from the applicant to understand if they are strictly necessary and/or they can be offered with a lower level of detail (**Gap 2**).

These gaps may be resolved by direct communication with the applicant and/or data holders. Whenever clarification from the applicant is needed (i.e., Gap 2 of the previous bullet point), the HDAB will contact the data applicant and start a discussion aimed at closing the gap towards an acceptable data access application.

In any case, before issuing a data permit, the HDAB must always assess whether the provision of pseudonymised or anonymised individual-level data is justified, taking into account the purpose of processing and the risk of re-identification (Articles 68(6) & 70(2), EHDS). In particularly sensitive or complex cases, (i.e., Gap 1), the HDAB may ask the data holder to apply quality metrics (e.g., k-anonymity, uniqueness analysis) to the expected datasets, to support this assessment, or ask the data holder to provide a sample of the expected datasets. Nevertheless, the responsibility to determine whether the re-identification risk is acceptably low and to decide whether data may be made available at individual level remains with the HDAB.

The quality of the data access application may also affect the level of interaction needed between the HDAB, the applicant and the data holder:

Scenario A – Detailed application: The applicant specifies precise variables, granularity levels, and richly justifies the use of quasi-identifiers. The HDAB validates alignment with minimisation principles, and the data holder can provide a cost estimate (see TEHDAS2 M4.1 for guidelines on fees related to the EHDS) based on well-defined parameters.

Scenario B – Vague application: The applicant requests broad tables or domains with insufficient detail. The HDAB must act as a facilitator, requesting clarification from the applicant and inviting the data holder to suggest a concrete variable set and associated minimisation techniques. The data holder may also propose specific transformations and provide a cost estimate based on the resulting dataset.

In each of the two scenarios, pursuing a collaborative approach allows the HDAB to make a legally sound, proportionate and documented data permit decision. It also reduces the likelihood that the data holder would later need to request modifications due to unmitigated re-identification risk or unresolved minimisation issues.

When exploratory or data-driven studies are envisaged, and the value of variables may only become apparent after the initial analysis, several approaches are possible:

- Focusing on the main purpose of the study will reduce potential variable selection even for this kind of studies.
- Demographic information on patients may be reduced to a minimum, strongly minimised, strongly encrypted (i.e. with advanced techniques that preserve machine readability) or even excluded.
- Studies may be divided into a preliminary phase where a reduced number of individuals are targeted, and/or where stronger minimisation techniques are used, and then a second fine-tuned analysis extended to more individual which includes only the more promising variables with lighter minimisation interventions.

Proactively anticipating possible issues with excessive or unprecise data access application forms is a shared responsibility between actors that may boost costs containment and processing times.

3.5 Closing remarks

Data minimisation is a foundational principle of lawful and proportionate secondary use under the EHDS Regulation. This section has outlined how HDABs, data holders and data applicants/users can operationalise this principle across the data lifecycle – from application, through assessment, to permit issuance and finally to data usage. Each dimension of data provision has been reviewed in terms of minimisation logic, risk management, and justification requirements.

With the procedural clarifications and practical examples provided, this section supports HDABs in implementing a robust, transparent and compliant data minimisation process. It also provides a clear framework for dialogue between actors involved in the data access process.

Further relevant methods and tools to improve process optimisation and harmonisation between actors are expected to emerge once HDABs will widely manage data requests and data permits.

The principles and methods described here should remain applicable not only to initial applications, but also to any subsequent amendments, ensuring continuity in the application of the minimisation principle.

4 Pseudonymisation

Pseudonymisation is defined in Art. 4(5) (GDPR) as “the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person.” (EDPB-G, §16).

Figure 4. Depiction of the pseudonymising transformation taken from the EDPB Guideline 01/2025.

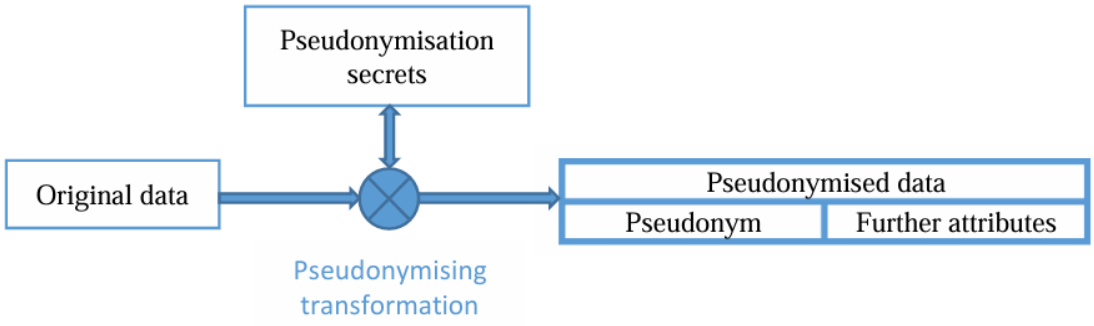

The purpose of pseudonymisation is balancing between privacy protection of data subjects, retaining data fidelity and quality, and the ability of data subjects to exercise their rights. Pseudonymised data is still personal data. While pseudonymisation allows for preserving value and utility of data (compared to anonymisation), pseudonymised data remain personal data but reduce direct identifiability. Processing pseudonymised data hence needs to be complemented with organisational (e.g., contractual limitations of use, contractual requirements) and technical measures (e.g., access control, encryption) to ensure adequate data protection.

In the context of this section, we will use terminology depicted in Figure 4, consistent with EDBP Guideline 01/2025. *Pseudonymisation* is a transformation, in which *direct identifiers are replaced by new identifiers called pseudonyms*.[9] The mapping of pseudonyms to the direct identifiers they replace is performed using *pseudonymisation secrets* which have to be kept separately from the *pseudonymised data* with appropriate technical and organisational measures. Pseudonymisation secrets constitute the *additional information* described by GDPR (Article 4(5)), which is necessary to map the data to data subjects. Note that we use the term *additional information* in this specific meaning in this section.

4.1 The purpose of processing pseudonymised data within the EHDS

The overarching purpose of pseudonymised data within the EHDS is to enable the provision of personal data for secondary use by data users, pursuant to a data permit. Processing of pseudonymised data within the EHDS must comply with both the EHDS Regulation and the GDPR. Pseudonymisation enables lawful secondary use in a SPE as part of a data permit granted by an HDAB, under Article 68 (EHDS).

[9] See §83 EDPB Guidelines 01/2025.

The following goals of processing pseudonymised data are considered within the EHDS, when a data applicant files a data access application (i.e., applies to get a data permit):

- Ensuring high data fidelity while applying strict data minimisation according to the approved research purpose (Article 66(1), EHDS; Article 5(1)(c), GDPR);
- Enabling linkage of data records across datasets from one or more data holders, under the supervision of the HDAB, while managing re-identification risks (Article 70(2), EHDS);
- Supporting the implementation of data subject rights, such as opt-out processes (see e.g., TEHDAS2 Deliverable D8.1 on opt-out and Article 71, EHDS), and significant findings (see e.g., TEHDAS2 Deliverable D8.2 on significant findings) (Recital 67 & Article 58(3), EHDS) where reversible pseudonymisation allows HDAB link back to the patient to report significant finding.

Justification of pseudonymised data access under data permit in SPEs. Under the EHDS Regulation, an HDAB shall provide pseudonymised data where there is sufficient justification that the purpose cannot be achieved with anonymised data (Article 68(1)(c), EHDS). The EHDS Regulation embeds the decision-making process within the framework of the GDPR, in particular the principles of necessity, proportionality, and data minimisation. The HDAB is required to assess each application individually and to issue a data permit that is proportionate to the purposes of the application and specifying the necessity of the data to be made available.

For further reference, a comprehensive definition of required justification can be found in the TEHDAS2 Deliverable D5.2.

When the subject of the data access application is access to pseudonymised data, the HDAB should assess this primarily. Providing anonymised data instead, where the user disagrees, would imply a rejection of the original application. Since the data permit constitutes an administrative decision, it is subject to judicial review and supervisory oversight and the HDAB is obliged to provide adequate reasoning and to apply a case-by-case necessity test. The instruments of judicial oversight operate as safeguards against a HDAB policy favouring anonymised data.

4.2 The concept of pseudonymisation in the context of the EHDS

Pseudonymised data can often be usefully analysed since, in large part, the information content of the original data can still be evaluated (i.e., fidelity of data is preserved). Moreover, if consistent pseudonymisation is used, the insertion of pseudonyms enables the linkage of various records of pseudonymised data relating to the same person without the need to use additional information.[10]

Processing pseudonymised data preserves fidelity of data, as only direct identifiers are replaced by pseudonyms and the data is only subject to data minimisation for the purpose defined by the data applicant in the data access application. To make data available for secondary use in the EHDS, pseudonymisation should be performed as early as possible, taking into account the purposes of the processing (Recital 72, EHDS). Data minimisation pays particular attention to quasi-identifiers, which in conjunction with other attributes imposes increased risk of indirect identification of data subject in the dataset.

[10] See §31 EDPB Guidelines 01/2025.

Pseudonymisation supports linkage of data across different datasets available at a single data holder, or, when consistent pseudonymisation is used across different data holders directly or via trusted third parties is used, at multiple data holders within one or even multiple countries (i.e., when data subjects have health records in several Member States). As dataset linkage increases privacy risks (namely re-identification risk and attribute inference), its consequences need to be carefully assessed by implementing data minimisation and by implementing adequate technical and organisational data protection measures.

Moreover, it may also be possible to use additional information to link different sets of pseudonymised data whose linkage has not been planned at the outset, i.e., at the time the purposes and means for the processing of secondary use have been determined by the controller or controllers involved. Implementing such linkage should be performed only by persons specifically authorised for this purpose[11], such as the HDAB, for example. See TEHDAS2 M7.5 on data linkage for more details.

Linkage is very relevant in the context of the EHDS, as health data is often split across different data holders (e.g., different hospitals, registers, and biobanks) and different datasets (e.g., different sub-systems comprising the overall hospital information system, which may use different identifiers for the same patient – acting as pseudonyms).

- Within the EHDS framework, linkage is performed under the supervision or responsibility of the HDAB, either directly or through authorised entities (e.g., (trusted) data holders or TTPs). [12].

For pseudonymisation to be effective, pseudonymised data must not contain direct identifiers (e.g., personal or social security identifiers) whenever those direct identifiers could be used in the pseudonymisation domain (i.e., the environment where the controller or processor wishes to prelude attribution of data to specific data subjects) to easily attribute the data to the data subjects. To this end, those direct identifiers are removed in the course of the pseudonymising transformation.

Table 2. Role-based involvement for pseudonymisation in the EHDS. The list below addresses data access applications, while similarities for data requests hold. For simplicity reasons, trusted data holders are left out from this table.

Actors	Use journey phase	Task	Activities
Data Holder	Data discovery	Dataset description	Ensuring sufficient information about the dataset (i.e., preferably including information about the dataset, sample distribution and personal data).
HDAB & Data Applicant	Data access	Data access application preparation	Agreement on properties and parameters of the pseudonymisation. Ensuring adequate utility of the dataset for the purpose defined by the applicant. This includes identification of direct identifiers to be removed from

[11] See §33 EDPB Guidelines 01/2025.

[12] See TEHDAS2 M7.5 Guideline for Health Data Access Bodies on linkage of health datasets.

			the dataset and replaced by the pseudonym.
HDAB (involving data applicant(s) and data holder(s) whenever relevant)	Data access	Data access application evaluation	The process of defining data minimisation parameters is informed by the risks of determined quasi-identifiers (indirect identifiers) during the data access application process.
HDAB	Data preparation	Pseudonymisation transformation pursuant to a data permit	Ensuring adequacy of the process, including: - decision of responsible entity/entities for pseudonymisation (HDAB vs. data holder(s) vs. trusted third parties) - approval of techniques used for pseudonymisation - approval of technical and organisational measures for storing pseudonymisation secrets (including separation of roles and adoption of measures such as use of encryption and strong access management systems, key rotation, escrow, incident response, audit logging) - design and approval of technical and organisational measures for processing pseudonymised data

Insofar as necessary for pseudonymisation to have the intended effect, it is complemented by the data minimisation process. In the context of pseudonymisation in the EHDS, which comes with complex combination of technical and organisational measures to minimise risks. The pseudonymisation transformation of data (except for generation of pseudonyms) should only use well-defined documented deterministic operations such as suppression or generalisation.[13] Where linkage and data fidelity are required, deterministic methods are preferable. Non-deterministic techniques are typically more suitable in anonymisation contexts.

- Note that this suggested range of techniques is narrower than the list in §84 of the EDPB Guidelines 01/2025. Non-deterministic operations disturb the original data to the point that the user can only have guarantees on statistical properties on the dataset level and not on the individual data values. Noise modulation techniques and other non-deterministic processes are suitable as a part of anonymisation techniques portfolio. Therefore, the choice of the pseudonymisation method may depend on whether data linkage is required for the purpose of processing and should be aligned between relevant involved parties before data access approval.

[13] Based on modified §84 EDPB Guidelines 01/2025.

- To prevent unauthorised attribution of pseudonymised data, the pseudonymising transformation regularly involves secret data. The controller may choose these data prior or during the execution of the transformation. These data are often either cryptographic keys (for encryption), salt values for hashing algorithms (for one-way functions) or tables matching pseudonyms with the personal data they replace. This secret data will be stored as a part of pseudonymisation secrets.[14]
- When implementing consistent pseudonymisation, controllers need to define which sets of personal data will be pseudonymised consistently based on objectives of the pseudonymisation. For example, they may decide to pseudonymise all data they collect on the same day consistently allowing for the linkage of two data records pertaining to the same data subject and collected on the same day, but preventing linkage of records of data collected on different days.
    - In particular, three ways to arrange for controlled linkage of pseudonymised data are widely used: person, relationship, and transaction pseudonyms. Note, however, that other ways to segment the pseudonymised data are available and may be appropriate for the respective use case[15].

Generally, pseudonymisation can be reversible or irreversible, depending on whether the pseudonymisation secrets are kept (reversible) or discarded (irreversible). Irreversible pseudonymisation can be used in cases where re-identification is not legally or ethically required or desired. EHDS should rely on reversible pseudonymisation to support implementation of data subject rights such as return of information on significant findings. Also, if pseudonymised and not identifiable data is stored at the data holder, reversible pseudonymisation is needed to implement opt-out for future data access applications. Reversal of pseudonymisation requires cooperation of the entity storing pseudonymisation secrets (e.g., trusted third parties).

4.3 Pseudonymisation with respect to the different phases of the EHDS user journey: from data discovery, access, to data processing

Data discovery phase: Pseudonymisation does not come directly into the discovery phase of the user journey. There are, however, certain aspects related to pseudonymisation, which are recommended to be advertised as a part of the data catalogue using HealthDCAT-AP (please consult D5.1, section 7.3.4 for a detailed overview of HealthDCAT-AP properties. For each dataset published by the data holder, it should be indicated:

- if the described original dataset dcat:hasPersonalData (see below);
- if the dataset is categorised as *personal electronic health data* (i.e., genetic data or data concerning health (Article 4 (13), (15), Directive 95/46/EC (General Data Protection Regulation) an open access subset should be provided to ensure meaningful use and interpretation of non-public datasets. Examples are anonymised subsets of an original dataset, or synthetic subsets. Information can be included as sample distribution in HealthDCAT-AP, using the dcat:sampledistribution property and accompanied by appropriate documentation (e.g., data dictionaries, codebooks, or a description of the anonymisation method). This approach ensures clarity and

[14] See §85 EDPB Guidelines 01/2025.
[15] See §115 EDPB Guidelines 01/2025.

avoids mixing metadata related to the original dataset with that of its processed versions.
- if the described dataset is suitable for linkage activities by indicating the properties:
  - hasPersonalData: Indicates whether the dataset includes personal data, and if so, what type. This can support evaluation of linkage feasibility by identifying the presence of shared identifiers or quasi-identifiers (e.g., patient IDs, dates)
  - sampleDistribution: Can provide a structural view of the dataset through a data dictionary (e.g., CSVW) or an anonymised/synthetic dataset. This supports the evaluation of shared variables (e.g., time periods, population segments) that are essential for assessing linkage potential.
  - HealthDCAT-AP does not directly enable linkage but it helps data users and HDABs to evaluate compatibility between datasets for linkage planning. HealthDCAT-AP provides data holders the possibility to inform data users that a dataset has already been successfully linked in the past with other datasets. This can be achieved by advertising the linked data set via HealthDCAT-AP record and use DCAT-AP “dct:source” property to reference related datasets from which the linked dataset has been derived.

Next to the information already specified in HealthDCAT-AP, it might be helpful to integrate additional information on the dataset regarding:

- pseudonymisation method: reversible vs. irreversible pseudonymisation used at data holder;
- linkage information: Can it be deterministically linked with other datasets from the same data holder, and if yes, with which datasets?
  - Can it be deterministically linked with other datasets from other data holders either specifically (e.g., oncology registries) or generally (e.g., any identifying patient records using the same citizen identifier)?;
  - Can be deterministically linked with other datasets from other data holders in other countries, and if yes, with which datasets either specifically (e.g., hospital records on rare disease patients with patient registers of European Reference Networks) or generally?;

Providing these descriptors will help manage expectations of data applicants during data discovery and streamlines the communication between data applicants and HDABs in the subsequent data access application phase.

Data access application phase: in this phase, an agreement needs to be reached between the data applicant and the HDAB (or TDH) on application of pseudonymisation and resulting data quality and fidelity. This is to ensure that the dataset to be released and paid for by the data applicant meets their requirements specified in the data access application and subsequently stated in the data permit (Article 68(10), EHDS) to ensure traceability and auditability of the agreement. This pertains to:
- When pseudonymised data can be provided to the data user into the SPE. If resources allow, the pseudonymising entity could provide more information on:
  - direct identifiers to be removed from the dataset and replaced by the pseudonym;
  - determined quasi-identifiers and their planned handling as a part of data minimisation.
  - pseudonymisation policy to be applied: deterministic pseudonymisation, or document-randomised pseudonymisation (i.e., randomly generated identifiers

of each appearance), or fully randomised pseudonymisation (i.e., for any occurrence)[16]

- When data linking is to be implemented, the determination of direct identifiers and quasi-identifiers needs to consider the whole resulting linked dataset, since some combined data variables may become quasi-identifiers only after the linkage (e.g., in a simple hypothetical scenario, where dataset 1 contains day of birth, dataset 2 contains month of birth, dataset 3 contains year of birth). This may refer to data combination (i.e., when bringing together data from multiple datasets based on one or multiple data permits, or legal basis), but also data linkage (i.e., when bringing together datasets from several sources based on one topic or subject) (see Annex 4 - Glossary).

Data preparation phase:

- Scenarios with or without linking data:
  - Identifying data arrives from data holder(s), pseudonymisation is implemented at HDAB(s) or at data holders when HDAB approved processes are in place.
  - Pseudonymised data arrives from data holder(s) and is passed on by the HDAB(s) to the data user without re-pseudonymisation. This is only possible if the data holder can implement per project pseudonym generation and uses a pseudonymisation algorithm approved by the HDAB.
  - Pseudonymised data arrives from data holder(s), data is re-pseudonymised by the HDAB(s) and passed on to the data user.
- Linking-specific scenarios:
  - Linkable data *in a single data holder over time* (i.e., pseudonyms stable in time for linking data which may be coming continuously or in batches over time) to allow for longitudinal research. Pseudonyms can be generated by data source or by the HDAB (or by a trusted third party – but not any specific advantage in this scenario).
  - Linkable pseudonymisation requested *across multiple data holder(s) in a single country*. Linkable pseudonyms are generated by the HDAB, or by a trusted third party (TTP) contracted by the HDAB or defined by law and imposed on data holders.
  - Linkable pseudonymisation *across multiple data holder(s) in multiple countries* (including rare diseases, or when cross-border registers are to be linked with other national data holder(s) data). Linkable pseudonyms are generated by a TTP contracted by all the involved HDAB and agreed by each HDAB on the data holder(s) in the specific country. An example of such a TTP is the SPIDER pseudonymisation tool operated by the European Commission[17].
  - While there is no formal requirement under the EHDS Regulation for such TTPs to be certified or audited under a specific scheme, they must implement appropriate technical and organisational measures in accordance with Article 32 of the GDPR and be subject to contractual oversight by the controllers (i.e., HDABs). The use of a common TTP should be based on mutual agreement between the HDABs and supported by documented responsibilities, safeguards, and auditability provisions.

[16] See p.13 Data Pseudonymisation: Advanced Techniques and Use Cases report by ENISA for more information.

[17] See https://eu-rd-platform.jrc.ec.europa.eu/spider/.

- Instead of contracting with a TTP, linkable pseudonyms can be generated by an algorithm based on secure multi-party computing[18] when agreed by all the parties involved.
- Linkable pseudonymisation across one or more data holder(s) *and authorised participants*. Linkable pseudonyms are generated by a TTP contracted by all the involved HDABs and authorised participants, or by algorithms based on secure multi-party computing agreed by all the involved HDABs and authorised participants.

Data use phase:

- HDAB (and Trusted data holder where appropriate) is responsible for ensuring that pseudonymised or anonymised data is delivered into the SPE, possibly in collaboration with the data holder(s) (Articles 72(2), 73(2), EHDS);
- HDAB is responsible for imposing adequate technical measures for processing pseudonymised data in the SPE, including requirements on the SPE provider and the data user (requirements on the data user will be formulated in the data permit and become legally binding for the data user) – these include strong access control to data including strong authentication and authorisation, encryption of data at rest, technical means to limit possibility of data extraction from the SPE (these requirements are in detailed specified in D7.4);
- HDAB is responsible for imposing adequate data protection-related organisational measures for processing pseudonymised data in the SPE, including requirements on the SPE provider (D7.4) and the data user requirements formulated in the data permit (thus becoming legally binding for the data user) – these include requirements on the data user to avoid any reidentification of person, to avoid any tampering of SPE security mechanisms, to avoid any attempts to extract data from the SPE, to impose requirements on training of data user, to require reporting of any incidents to the HDAB, etc.

Finalisation phase:

- Archival of pseudonymised data is possible based on the data permit (subject to Article 68(12), EHDS). Archival refers here to time-limited storage within the SPE for reproducibility purposes, aligned with the validity of the data permit (as described in TEHDAS2 Deliverable 7.1) - possibly on a “cold storage”. Archival should not be interpreted as long-term or open-ended storage.

4.4 Pseudonymisation requirements

Pseudonymisation, risk assessment, and risk management needs to be considered in the overall context of the EHDS. The data access mechanism of the EHDS only allows processing of data in SPEs, which have their own technical and organisational safeguards and must have guarantees on the level of anonymisation of the results or any other output research object (e.g., AI model, software). These safeguards can be taken into consideration with managing risks of quasi-identifiers handling the data minimisation (i.e., if a pseudonymisation domain is in a SPE, data minimisation may accept higher risks in quasi-identifiers taking into account the specific purposes of the processing compared to a completely generic case of pseudonymisation in the pseudonymisation domain).

[18] See https://doi.org/10.1093/bioinformatics/btaa764, https://doi.org/10.1109/TIFS.2021.3114026.

Recommended pseudonymisation techniques and best practices to achieve reversible pseudonymisation:

- Use of *lookup tables*:
  - Need to save the complete mapping to implement reversibility.
  - Examples of possible techniques: counters, random number generators (typical for so called tokenisation mechanisms).
- Use of cryptographical functions to generate pseudonyms:
  - Need to save function used, exact specification of inputs, value of salt (i.e., a randomly generated list of characters added to the data before creation of the pseudonym, the salt must be kept separate and secure).
  - Recommended to save also complete mappings between unique identifiers in the dataset and pseudonyms if technically feasible (e.g., due to capacity constraints) – to handle situations that input data used to generate the pseudonym is changed at source subsequently after pseudonymisation (e.g., correcting errors in inputs such as names or birth dates or national identifiers), where required.
  - Examples of possible techniques: *message authentication code (MAC)* (basically hash functions with additional secret key input – so called salt), *keyed hash functions, symmetric encryption*[19].
    - chosen schemes and their configuration (e.g., size of keyed hash functions) should consider post-quantum security requirements;
    - if the encryption has homomorphic properties, the impact needs to be specifically assessed in terms of risk-to-benefit ratio.
  - Use of unsalted hash functions used on direct identifiers, or any other similar method where the knowledge of the method and knowledge of identifiers of one or more persons allows deterministic generation of pseudonyms, does not qualify as pseudonymisation.
  - Given the architecture of the EHDS, where the process is controlled by the HDAB, and data holders are legally required to cooperate with the HDAB, advanced scenarios such as chained pseudonym generation or pseudonyms with the proof of ownership are usually not necessary.

Any pseudonymisation secrets must be stored securely with adequate technical and organisational safeguards in accordance with Article 32 and Article 4(5) (GDPR) in a way, that reversibility of the pseudonymisation can only be implemented by the HDAB or a designated TTP and *not* by the data user (Article 66(3).

The recommended pseudonymisation policy, whether deterministic pseudonymisation, or document-randomised pseudonymisation, or fully randomised pseudonymisation,,[20] depends on the purpose and requirements of data processing. ENISA suggests choosing the pseudonymisation technique based on the identified risk and the identified or expected utilisation of the pseudonymised dataset. Random number generators and MACs are stronger encryption methods as they prevent exhaustive search, dictionary search and random search. However, the pseudonymisation entity may lean towards a combination of different methods due to practicality reasons. Regarding the pseudonymisation policies, fully randomised pseudonymisation offers the best protection, while hindering linkage. Therefore, document-randomised and deterministic methods are recommended, as they allow for data

[19] See Data Pseudonymisation: Advanced Techniques and Use Cases report by ENISA for more information.

[20] See p.13 Data Pseudonymisation: Advanced Techniques and Use Cases.

linkage[21]. To enhance data utility for the user, the pseudonymisation policy could be agreed upon between the data user and the HDAB in the data access application phase and stated in the data permit in accordance with (Article 68(10)(a), EHDS) to ensure traceability and auditability. In the context of the EHDS with all other technical and organisational safeguards in place, it is advisable to consider deterministic pseudonymisation, which gives the data user ability to recognise the same patient in the dataset. Only when this is not needed by the data user, one of the randomised policies can be used.

The use of irreversible pseudonymisation is discouraged within the context of the EHDS and should only be used for legacy situations where data is already irreversibly pseudonymised at the data holder. Irreversible pseudonymisation may limit the ability to exercise certain data subject rights and does not automatically meet the threshold for anonymisation under Recital 26 GDPR.

- Pseudonyms MUST NOT be reused across different data permits, to minimise risks of intentional or unintentional linkage of data by the data user.
- For data requests, HDABs can repeatedly process datasets with the same pseudonyms, as the data does not leave the HDAB and only a response in an anonymised statistical format is shared with the data user.

4.5 Safeguarding pseudonymised data in the EHDS

- Separation must be implemented between pseudonymisation secrets and pseudonymised data. This includes implementing technical and organisational measures at the entity responsible for pseudonymisation (HDAB or TTP or data holder, depending on the scenario).
- Identification of quasi-identifiers and handing over information on them to the data minimisation process (see section 3) such as:
    - Minimisation of quasi-identifiers in the dataset as a part of data minimisation process – because during data minimisation, the purpose of processing and requirements on data quality are known (fitness for purpose), which helps to consider which quasi-identifiers need to be released and if any generalisation/suppression/noise generation can be applied to them
    - Provide health-data specific quasi-identifiers beyond what is in the EDPB guidelines (e.g., dates of visits in the hospital also become quickly identifying – possibly from the public health databases and health care registers)

4.6 Examples of open-source/reference tools for pseudonymisation

- **Mainzelliste**[22] (and its extensions)
    - Open-source first-level pseudonymisation service with record linkage (often used as a TTP component).
    - Extensions for *privacy-preserving record linkage* using secure MPC exist (MainSEL / SecureEpiLinker / MainzellisteSEL).
- **MOSAIC suite**[23]

[21] See p.13 Data Pseudonymisation: Advanced Techniques and Use Cases.
[22] https://bitbucket.org/medicalinformatics/mainzelliste/
[23] https://github.com/mosaic-hgw/

- **E-PIX** (ID management + probabilistic record linkage; published as open source, AGPLv3 is explicitly stated in the literature).
- **gPAS** (generic pseudonym administration service).
- **gICS** (consent management; often paired with pseudonym services).

4.7 Data subject rights within the EHDS

- Processing pseudonymised data means processing personal data – hence all the data subject rights apply (but see also Article 11, GDPR).
- Opt-out and significant findings propagation, as these scenarios require reversible pseudonymisation methods to be implemented. For more information on the notifications of natural persons of significant findings can be found in D8.2.
- Opt-out only applies to the new projects. For ongoing project, the opt-out does not apply (see Article 71(3), EHDS). For more information on opt-out, please refer to D8.1.

4.8 Relative anonymity

Rulings of the European Court of Justice[24] have pointed out that pseudonymised data may be understood as *relatively anonymous* where the data recipient has no legal means to re-identify data subjects. The assessment of relative anonymity depends on the legal possibilities and technical capabilities of the data recipient or processor and therefore can only be made in the context of a specific recipient or processor. Even if the data is not identifiable for a specific recipient taking into account the means reasonably likely to be used, hence no longer considered personal for that particular recipient of data, it still needs to be accompanied by technical and, in particular, organisational measures to prevent situations where the relative anonymity status could change and the data could become identifying or be re-identified, or where other privacy risks could emerge (note that attribute inference or group-membership attacks are also privacy risks, not only re-identification risk). This means that even if the data is no longer personal, the recipient or processor must be limited/prevented from further sharing the data or changing the context in which it is processed.

In the context of EHDS, the status of relative anonymity could be considered in three scenarios:
(a) Retrieving responses to data requests to HDAB (Article 69, EHDS);
(b) Retrieving data from the SPE under data permit (Article 68, EHDS);
(c) Processing anonymised data in the SPE.

For *data requests*, scenario (a), the regulation stipulates that the *response can only be provided in an anonymised **statistical format***. Relatively anonymous pseudonymised data is typically more granular than statistical-format data (i.e., aggregated data), and therefore the concept of relative anonymity is not relevant in this case.

For *data permit*, scenario (b), the regulation stipulates in the Article 73 that the *HDAB shall review the electronic health data included in a download request to ensure that health data users are only able to download non-personal electronic health data, including electronic health data in an anonymised statistical format, from the secure processing environment.* If relative anonymity is used as a basis for exporting individual-level data from the SPE, the assessment of all risks (as defined in the generic anonymisation section below) must be carried out in the context of the particular data user. The permit must also prevent the data user from passing on the data to other entities and may

[24] Breyer case (C-582/14) and more recently EDPS v. SRB case (C-413/23 P).

need to restrict processing to a specific context, since the relative anonymity status could otherwise change.

For processing data in within an SPE, scenario (c), SPEs must in any case support the processing of personal data, and anonymisation within the SPE may only be useful as an additional risk-mitigation measure where the data user does not need personal (pseudonymised) data. Relative anonymity could be used to justify less stringent technical (e.g., lower requirements for multi-factor authentication) and organisational measures (e.g., less extensive user training and less demanding permit requirements) for a specific processing activity in the SPE, in the context of a specific user and processing purpose.

5 Anonymisation and synthetic data generation

Anonymisation and synthetic data generation techniques can be applied throughout the EHDS data lifecycle to provide stronger protection of personal data than pseudonymisation alone. These approaches are distinct privacy-preserving methods and support the fulfilment of obligations of both the Health Data Access Body (HDAB) (Article 73(2), EHDS) and the health data user (Article 61(4), EHDS) regarding the export of non-personal electronic health data (including anonymised statistical format), as applicable. Such techniques may be applied by HDABs, data holders, and data users at the stage where data or data processing results produced under a data permit are exported from secure processing environments (SPEs), or when data resources are disclosed to a data user following a data request. In addition, in specific cases, HDABs and data holders may, in agreement with the data user, apply these methods to data resources processed within the SPE itself in order to achieve enhanced data protection.

Anonymisation involves modifying or aggregating personal data so that it can no longer be linked to a specific individual, in compliance with data protection regulations. In contrast, **synthetic data generation** creates entirely new, artificial data that reflects the statistical properties of the original dataset. Synthetic data is typically produced using statistical or AI-based models trained on real pseudonymised data. This guideline addresses synthetic data generation primarily from a data protection perspective, noting that synthetic data is not inherently anonymous and requires appropriate safeguards. Beyond data protection, synthetic data may also be generated within the SPE in order to augment datasets, for instance to enhance the performance of deep learning models.

Although the methodologies for generating anonymised and synthetic data differ, both require the HDAB to establish similar processes and functional components. Producing either type of data requires evaluation of data quality metrics, conducting privacy risk assessments (including assessment of re-identification and inference risks), and enforcing disclosure control measures. Therefore, this chapter considers anonymisation and synthetic data generation in parallel. The main differences between the two approaches lie in the specific technical methods and tools used, which will be addressed separately.

5.1 Objectives

Data protection by using anonymisation and synthetic data generation techniques is especially needed when allowing the data user to export data processing results or data contents from SPEs. These techniques are also valuable earlier in the workflow, for example to enable data users to develop and validate analysis scripts within the SPE before access to real data is granted. The HDAB may also generate anonymised or synthesised data to enable testing with *public use files* outside SPEs. Anonymisation is also essential when data is prepared for use under *data request*.

It is essential that the EHDS infrastructure supports the assessment of anonymised and synthetic data quality, taking into account the three key dimensions: fidelity, utility, and privacy. From the HDAB's perspective, particular emphasis should be placed on the assessment of privacy risks. Importantly, privacy risk in this context extends beyond the data subject to include other individuals, such as relatives and health care professionals.

The objectives of this chapter are to:

- Define the contexts ("use cases") and high-level architecture where anonymisation, synthetic data generation and privacy risk assessment are carried out in the EHDS.
- Provide guidelines concerning the methodology and tools which should be deployed to ensure safe and efficient anonymisation, synthetic data generation and privacy risk assessment implementation.

5.2 Scope and assumptions

5.2.1 Assumptions about data

Foundation in real, individual-level data. It is assumed that anonymisation and synthetic data generation are performed using real, individual-level data or results derived from such data. For instance, synthetic data generated solely from public aggregate statistics is anonymous by nature and, therefore, falls outside the scope of this deliverable.

Data utility and fidelity. While privacy protection is the primary concern from the HDAB's perspective, assessing the utility and fidelity of anonymised or synthetic data is essential to ensure the data remains suitable for its intended use.

Evolving privacy risk. It is recognised that anonymisation and synthetic data generation almost always entail some residual privacy risk. These risks may increase over time as technologies and threat landscapes evolve. Therefore, anonymisation and synthetic data generation and the related privacy risk assessment should be understood as ongoing, adaptive processes that require regular review and maintenance.

Context-dependent risk tolerance. It is recognised that acceptable levels of privacy risk may vary case-by-case, with higher risk potentially tolerated for less sensitive data.

Data subjects. In addition to primary data subjects, personal data of healthcare professionals or other contributors to care may be present in the original data. The processes and methods referred to in this chapter should be applied to ensure the protection of such personal data wherever applicable.

5.2.2 EHDS scope

Purpose of data use. It is assumed that data usage is driven by the permitted purposes defined in the EHDS Regulation. Anonymisation or synthetic data generation are not standalone permitted purposes for data use.

Responsible actors. Anonymisation, synthetic data generation, and related risk assessments may be carried out by HDABs, data holders, or data users, depending on the use case. HDABs and data holders may also assign these tasks to other entities, such as SPE operators, on a contractual basis.

Data holders' role. The EHDS Regulation defines through Articles 57 and 72 that anonymisation can be carried out by *trusted data holders*. In line with Recital 72 of EHDS this deliverable assumes that any data user may be allowed to carry out anonymisation and synthetic data generation activities.

Data processing infrastructure. This deliverable focuses on anonymisation, synthetic data generation, and privacy risk assessment activities that support the responsibilities defined in the EHDS Regulation. It is assumed that anonymisation and synthetic data generation takes place in a single SPE as outlined by (Article 73 EHDS). Activities conducted by health data holders outside the EHDS infrastructure – such as the sharing of non-personal data via open databases in accordance with Article 60 of the EHDS – are not in the scope of this deliverable.

Synthetic data. The EHDS Regulation does not explicitly require synthetic data generation to be supported by the EHDS infrastructure. However, it is included in this guidance because synthetic data generation is widely recognised as a valuable approach for enabling research and innovation in the EHDS context.

5.2.3 Limitations

Criteria for disclosure decision. Definition of the exact qualitative and quantitative criteria for disclosing anonymised or synthetic data for the data user is not in the scope of this deliverable as such definitions are heavily case-dependent.

Policy considerations. Policies related to granting permissions to data users to create and use anonymised or synthetic data – such as exporting data from the SPE – are outside the scope of this deliverable.

Coverage of costs. This guideline does not address the allocation of costs for anonymisation and synthetic data generation activities among stakeholders.

HDAB collaboration. To promote efficiency in implementation and maintenance, many of the described tasks and processes should be reused, and mutual recognition adopted where appropriate. This deliverable, however, does not provide guidance on the organisation of such collaborative arrangements.

5.3 Use cases

The use cases for data anonymisation, synthetic data generation, and privacy risk assessment are listed in Table 3 and mapped to the EHDS user journey in Figure 5. Use case 1 covers EHDS activities that are not tied to a specific data permit but are part of regular activities of HDAB or data holder in enabling secondary use of data. Use cases 2–4 are related to a specific data permit (or data request in case 2), and the associated anonymisation and synthetic data generation measures must be defined and agreed during the data access application phase. This ensures that both the data applicant and the HDAB are aligned on the conditions for data access before the permit is granted and the data applicant commits to any data retrieval costs.

In use cases 3 and 4, the data user has been granted access to pseudonymised data within the secure processing environment (SPE) and seeks to export processing results for reporting or publication. In addition to such results, the data user may wish to export anonymised or synthetic data in order to continue research or testing activities outside the SPE. The anonymisation or synthetic data generation is performed either by the data user (use case 3) or by the HDAB or data holder (use case 4). In all cases, the HDAB retains final responsibility for disclosure approval. Use case 2 may be adopted by the HDAB in cases where pseudonymised data cannot be released due to high privacy risks, but anonymised or synthetic data set can be released instead.

Annex 5 provides example scenarios for each use case.

Table 3. Use cases for data anonymisation, synthetic data generation and privacy risk assessment.

Use case	Roles	Rationale
1. HDAB (or data holder) wants to create public use files. (Article 57(1l) EHDS)	• HDAB (or data holder) anonymises data or generates synthetic data • HDAB (or data holder) assesses privacy risk of data intended to be published.	• Public use files are provided for testing purposes prior to data access application
2. HDAB wants to anonymise data or generate synthetic data for the data user under approved data request (Article 69 EHDS) or data permit (Article 68 EHDS).	• HDAB (or data holder) anonymises data or generates synthetic data • HDAB assesses the privacy risk • HDAB optionally assesses the utility and/or fidelity of the anonymised or synthetic dataset.	• Data access is granted based on data request, which allows data disclosure only in anonymous statistical format • HDAB considers that data can only be released for processing by the data user in the SPE in an anonymised form due to the nature of the project (see Section 4.1).
3. Data user wants to anonymise data or data processing results or to generate synthetic data to enable their export from the SPE (Articles 61(4) and 73(2) EHDS).	• Data user anonymises data or data processing results or generates synthetic data. • Data user and HDAB assesses the privacy risk of the data or data processing results • HDAB approves the export of data or data processing results.	• Data user needs to export data or data processing results for scientific publication and other secondary use purposes.
4. HDAB wants to anonymise data or data processing results or to generate synthetic data to enable their export from the SPE(Articles 61(4) and 73(2) EHDS).	• HDAB anonymises data or data processing results or generates synthetic for the data user. • HDAB assesses privacy risk • HDAB approves the export of data or data processing results..	• Data user needs to export data or data processing results for scientific publication and other secondary use purposes.

Figure 5. Anonymisation, synthetic data generation and privacy risks assessment use cases (1-4) mapped on the EHDS user journey. Data minimisation and pseudonymisation activities are covered in section 3 and 4, respectively, of this deliverable.

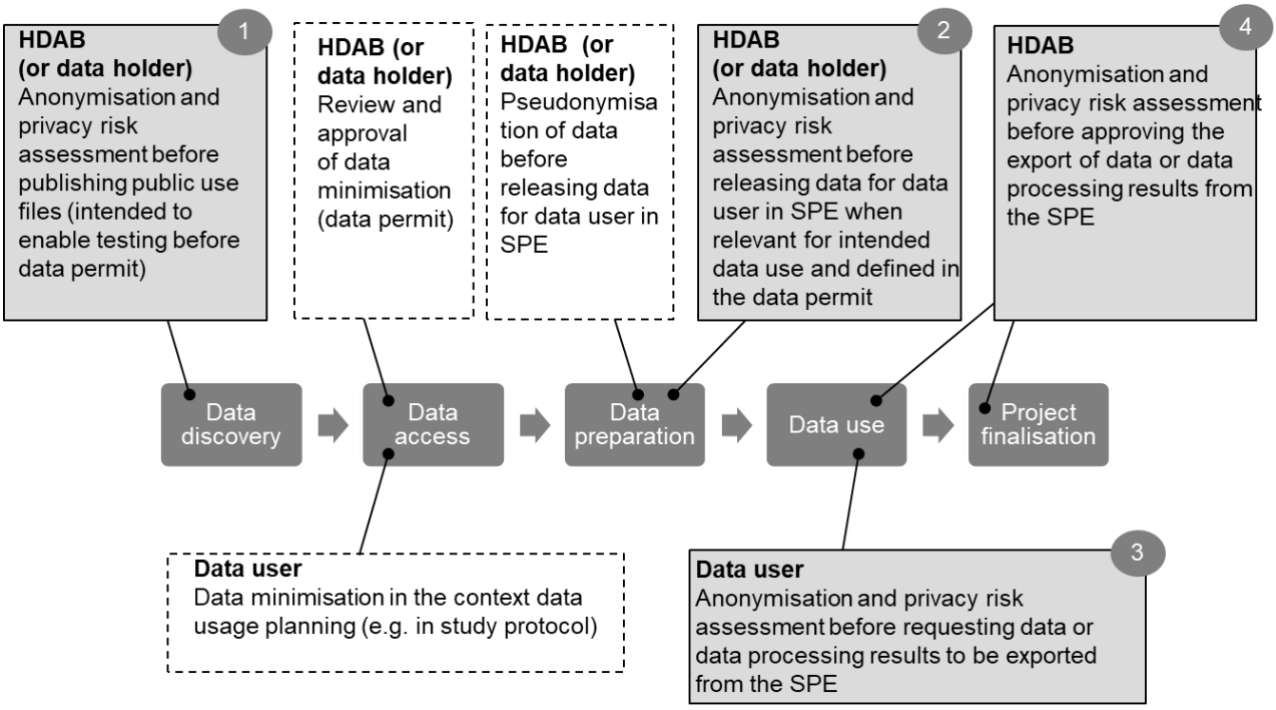

5.4 Architecture

The high-level architecture presented in Figure 6 defines the overall framework within which anonymisation, synthetic data generation, privacy risk assessment, and related activities are performed. These activities are mostly carried out by the HDAB and the data user as outlined in Table 3. The high-level architecture leaves freedom for the HDABs in implementing the detailed processes following existing models for disclosure control. The functions referred in the architecture are designated to take place within a SPE[25]. In this context, the term SPE is understood broadly, encompassing both the HDAB's infrastructure and any SPE provider infrastructure where data is processed in accordance with the EHDS Regulation throughout the data lifecycle.

[25] TEHDAS2 Deliverable document D7.4

Figure 6. Overall architecture for secure disclosure of anonymised data, processing results, and synthetic data. This setup applies to all Table 3 use cases, with role activities (HDAB, data user, health data user) varying by use case.

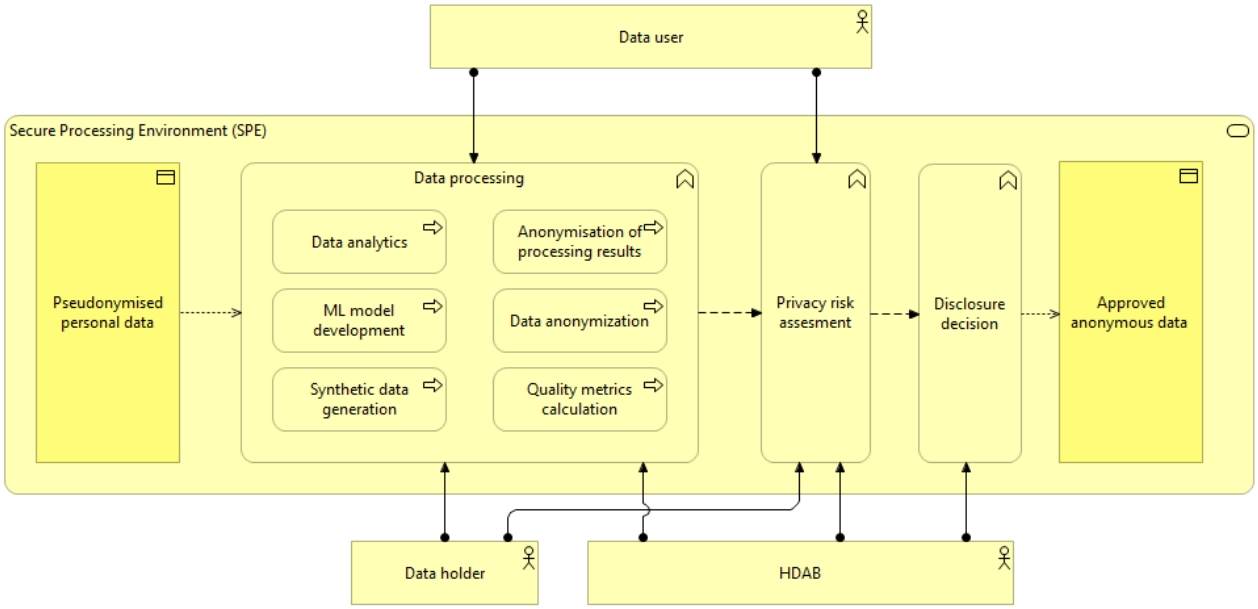

The high-level architecture is intended to be agnostic to the use cases and to the methods used for anonymisation or synthetic data generation. Therefore, it is applicable across all use cases and it allows the most appropriate methods to be selected for anonymisation and synthetic data generation case-by-case.

The components of the architecture are:

Pseudonymised personal data can be:

- Use case 1: any individual-level data which the HDAB or data holder uses for generating public use files.
- Use case 2: data permit-specific data which the HDAB uses as basis for generating an anonymised or synthetic dataset to be released for the data user in the SPE.
- Use cases 3-4: data permit-specific data which the data user processes to create data processing results and which data user or HDAB uses as a basis for generating an anonymised or synthetic dataset.

Data processing refers to the data processing activities carried out by the data user, HDAB or data holder in the scope of use cases 1-4:

- Data analytics and machine learning (ML) model development are data user's activities in utilising pseudonymised data as enabled by an approved data data permit
- Anonymisation of processing results is carried out by the data user, HDAB or data holder in order to enable data processing outputs to be exported from the SPE for reporting and publishing purposes.
- Data anonymisation or synthetic data generation are carried out by (1) data user as an activity enabled by approved data data permit, or (2) by HDAB or data holder for

generating data for public use or specificly for a data user as enabled by the data permit.
- Quality metrics calculation covers privacy, fidelity, and utility dimensions and is carried out in the context of data anonymisation or synthetic data generation by the data user, HDAB, or data holder.

Privacy risk assessment refers to the processes for assessing the privacy risks of anonymised data, data processing results, or synthetic data. The activity involves comparing similarity-based privacy metrics against pre-established criteria, as well as assessing attribute and linkage inference risks. These activities are carried out by the HDAB, data holder, and data user, depending on the use case.

Disclosure decision refers to the process by which HDAB decides whether anonymised data, data processing results, or synthetic data may be disclosed and exported from the secure processing environment (SPE) to the data user. In the context of use case 1, the decision concerns approval of releasing a public use file for open use.

Approved anonymous data refers to anonymised or synthetic data or data processing result which has been deemed to be anonymous and approved to be used within the SPE (use case 2) or to be exported from the SPE (use cases 1, 3 and 4).

5.5 Guidelines

The high-level architecture depicted in Figure 6 outlines a disclosure control framework for the HDAB while allowing flexibility in the choice of methodologies, state-of-the-art technologies and tools. The following guidelines are intended to support the HDAB in establishing the capabilities needed for efficient generation of anonymised and synthetic data as well as for implementing controls that ensure the safe disclosure of anonymous data and processing results. Strategies, processes and methods applied for anonymisation, synthetic data generation, and associated privacy risk assessments shall be periodically reviewed to verify their effectiveness and continued validity.

The evaluation of quality metrics and assessment of privacy risk fall primarily under the HDAB's responsibility. However, data users may only request the export of anonymous data from the SPE and must therefore be able to assess the associated privacy risks. Data users are also strongly motivated to evaluate the fidelity and utility of the exported data.

5.5.1 Documentation of anonymisation or synthetic data generation

Metadata and traceability. All anonymisation and synthetic data generation activities must be thoroughly documented. This documentation should be linked as metadata to the resulting anonymised or synthetic dataset, ensuring that essential information remains available to any future user of the data**.** Most essentially, the documentation should support the process of evaluating the privacy risk and in making the disclosure decision. The documentation will also be highly important when a future user of the data assesses the applicability of the data for a particular purpose.

Responsibility of documentation. All parties involved in the processes indicated in Figure 6 must produce relevant documentation for the actions they perform. The HDAB is responsible for ensuring that such documentation is created and attached to the disclosed data or data processing results. This includes establishing suitable control functions and

providing documentation guidance to HDAB staff, data users, data holders and all downstream users of the anonymised or synthetic data or data processing results.

Documentation content. Where applicable, the following elements are recommended for inclusion in the documentation, together with timestamps[26]:

- **Identification information**
  - Data creator identifier
  - Data permit identifier
- **Metadata of the original dataset** that was anonymised or used to generate synthetic data, including provenance information, dataset size, descriptive statistics and data quality and utility label (EHDS Regulation, Article 78).
- **Data processing documentation**
  - Dataset size and descriptive statistics of the used subset of data
  - Data processing steps, methodology, transformations and parameters used;
  - Tools, software and platforms used (with versions)
  - Specific measures applied to ensure privacy protection during ML model training or synthetic data generation (e.g., differential privacy parameters).
  - Quality metrics of anonymised or synthetic data
  - ML model validation and testing results
- **Privacy risk assessment documentation**
  - Description of methods, standardised metrics, criteria and thresholds used for risk assessment
  - Qualitative or quantitative risk assessment results
- **Disclosure documentation**
  - Categorisation of the data contents to be disclosed (anonymised data, data processing results, synthetic data)
  - Disclosure decision, including compliance statements and any remarks from the decision process
  - Data recipient and intended use
  - Any restrictions or conditions attached to the disclosed data

Anonymity of documentation. The documentation shall not contain personal or identifying information relating to the data subjects whose data have been processed under the data permit. This allows the documentation to be shared with the data user alongside the anonymisation results or synthetic data.

Documentation structure. Documentation shall be provided in a structured, standard format. Preferably, it should be machine-readable to support the development of both fully automated and assisted approval processes. Preferably, the documentation should be machine-readable to support the development of fully or partly automatic approval processes. Standard structure for anonymised or synthetic data documentation is currently not existing and should be developed.

Documentation provenance and traceability. Documentation shall include clear information on the origin, authorship, and time of creation or update for each documentation entry. Appropriate mechanisms shall be in place to maintain traceability across documentation updates, enabling verification of how the documentation has evolved over time.

[26] See also De-identification Guidelines for Structured Data (Annex J)

5.5.2 Ensuring anonymity of statistical data processing results

Ensuring anonymity of results. HDABs are responsible for ensuring the anonymity of data processing results intended to be exported from the SPE in line with Article 73(2) EHDS. The HDAB controls and approves the export, while the data user must comply with permit and not attempt disclosure of personal data.. Besides the HDAB, the data user has a responsibility[27] in ensuring that the data processing results and outputs to be exported from the SPE are anonymous. Thus, also the data user benefits from the privacy risk assessment guidance.

HDAB shall establish well-defined criteria for privacy risk assessment, along with clear guidance for both its staff and data users to support them in fulfilling their respective responsibilities[28]. Findata's criteria and guidance[29] for producing anonymous results are a good example covering the following result types:

- descriptive analysis and indicators
- correlations and regression-type analysis
- graphs of different types
- images and other imaging materials
- results based on genome data
- machine learning models
- individual-level result materials
- synthetic data
- results of qualitative research.

5.5.3 Controlling privacy impacts of machine learning models

Data users are responsible for mitigating privacy risks in ML models and their outputs, whereas HDABs retain the final responsibility for ensuring that any ML models or outputs exported from the SPE do not leak personal information. Privacy impacts of two threat models shall be considered[30]: the ML model itself and the outputs of generative models (e.g. synthetic data). Large models can be difficult to assess directly for privacy risks[31]. When such assessment is not feasible, the data user shall mitigate privacy risks by training the model on an anonymised dataset or by applying relevant methods, in particular differential privacy (DP) during the model development[32, 33].

Applying DP to ML model training is effective for addressing both threat models referred to above. Controlled noise is introduced into the training process, ensuring that the contribution of any single individual in the dataset cannot be distinguished. This can be applied, for example, to the gradients computed during training, thereby limiting the information that the

---

[27] EHDS Regulation Article 61(4).
[28] Guidelines for output checking.
[29] Producing anonymous results - Findata.
[30] https://jair.org/index.php/jair/article/view/14649/26952
[31] edpb_opinion_202428_ai-models_en.pdf.
[32] https://findata.fi/en/services-and-instructions/producing-anonymous-results/#other-result-types .
[33] https://www.sciencedirect.com/science/article/pii/S0167739X23002315

model can memorise about specific data points. By carefully calibrating the noise and the privacy budget, DP provides a quantifiable and formally defined guarantee of privacy protection. Models trained with DP aim to preserve overall utility while bounding the risk of re-identification. Implementing DP requires the selection of appropriate algorithms and parameters, as well as accounting for cumulative privacy loss over multiple training iterations.

While introducing noise into the training process, DP can affect the quality of the resulting model. It is therefore good practice to assess the utility of a DP-trained model by comparing it with a model trained on real data. This assessment can be safely carried out within the SPE prior to exporting the DP-trained model. The use of DP is not required in cases where neither the ML model nor individual-level model outputs are intended to be exported from the SPE, but used only for generating aggregated results within the SPE.

Detailed documentation of the ML development process, including quality metrics and the description of the measures taken to prevent privacy leakage from the ML model shall be included in the anonymisation documentation to support privacy risk assessment and disclosure decision-making.

5.5.4 Anonymisation of individual-level data

Anonymisation methods. HDABs should be prepared to apply various anonymisation techniques and to assess the privacy risk of the resulting anonymised data. A variety of methods are used in healthcare data anonymisation, and the optimal approach depends on both the type of data, the relevant use case (Table 3) and its intended use outside the SPE. As a baseline, all direct personal identifiers and pseudonyms must be removed from the dataset (Chapter 4). In addition, the following anonymisation methods are commonly applied:

- Common anonymisation approaches for tabular data are[34,35,36,37,38, 39, 40]:
  - Perturbation-based methods: noise addition, data shuffling, micro-aggregation, data swapping.
  - Generalisation- and suppression-based methods: generalisation of values (e.g., age ranges), suppression or removing of identifiers, top- and bottom-coding, masking.
  - Aggregation-based methods: grouping, binning, statistical summarisation
- Common anonymisation approaches for imaging and bio-signals data are[41]:
  - Metadata anonymisation, pixel redaction.
  - Image de-identification: face removal, surface rendering anonymisation, skull stripping.
  - Perturbation-based methods.

[34] guide-to-basic-anonymisation-(updated-24-july-2024).pdf.
[35] Perturbation Methods for Protecting Data Privacy: A Review of Techniques and Applications.
[36] Opinion 05/2014 on Anonymisation Techniques.
[37] Anonymisation and Personal Data - Finnish Social Science Data Archive (FSD).
[38] ENISA, Data protection engineering, January 2022.
[39] Techniques - De-Identification Profile v0.0.1-current
[40] Statistical Disclosure Control (SDC) methodology | Eurostat CROS
[41] Ministry of Social Affairs and Health (Finland), VN/23353/2022.

- Common anonymisation approaches for free text are[42]:
  - Named entity recognition (NER) and rule-based filtering;
  - Text generalisation and masking.
  - Text perturbation and data synthesis.
- Common anonymisation methods for genome data[43]:
  - Suppressing or removing identifiable variants.
  - Reducing resolution of genetic data.

Privacy properties targeted in tabular data anonymisation are:

- k-anonymity: each record is indistinguishable from at least k−1 others with respect to quasi-identifiers.
- l-diversity: each equivalence class has at least l "well-represented" sensitive values.
- t-closeness: the distribution of sensitive attributes in each equivalence class is within a threshold t of the global distribution.

Anonymisation methods should be selected based on a contextual risk assessment and not applied as fixed recipes. The risk of re-identification must be assessed using state-of-the-art techniques appropriate to the data type and use case. Specific attention shall be given to omics data, such as whole genome sequences, rare variants, methylation profiles, and proteomic signatures, which may be inherently identifying.

In addition to privacy considerations, it is important to assess the utility of the data after anonymisation to ensure that it meets the quality requirements of the intended use.

Aggregated data. When a data user is granted access to data under a data request, the individual-level data shall be converted into anonymous statistical format in accordance with the EHDS Regulation. The data must be aggregated over a sufficient number of individuals ensuring compliance with relevant k-anonymity criteria.

5.5.5 Synthetic data generation

Methods for synthetic data generation. HDABs should be prepared to apply commonly used synthetic data generation techniques and to assess the privacy risk of the resulting data. A variety of methods are used in generating synthetic data in healthcare, and the optimal approach depends on both the type of data and its intended use[44]. Synthetic data generation methods can be roughly divided into three groups:

- Statistical methods (e.g., multivariate Gaussian models, copula-based methods, generalised linear models, Bayesian networks, Markov models, Markov random fields, multiple imputation). These methods learn multivariate distributions and dependency structures from real data, which are then sampled to generate synthetic data.

[42] De-identification of Free Text Data containing Personal Health Information: A Scoping Review of Reviews | International Journal of Population Data Science.
[43] Computational tools for genomic data de-identification: facilitating data protection law compliance | Nature Communications.
[44] Synthetic data generation methods in healthcare: A review on open-source tools and methods - ScienceDirect.

- Classical ML-based methods (e.g., decision trees, random forests, gradient-boosted trees such as XGBoost). These approaches learn conditional relationships between variables and can be used for synthetic data generation via sequential modeling, conditional sampling, or model chaining.
- Deep generative models (e.g., generative adversarial networks (GANs), variational autoencoders (VAEs), diffusion models, transformer-based models, autoregressive models and large language models (LLMs)). These models leverage neural networks to learn complex data distributions and generate synthetic samples, including images, text, and other high-dimensional data.

Reducing privacy risk. Synthetic data may still qualify as personal data under GDPR if it can be linked back to individuals with reasonable effort. This must be considered when evaluating privacy risk.

Privacy risks can be controlled by incorporating privacy-enhancing techniques into the synthetic data generation process as referred in Section 5.5.3. Applying differential privacy during model training can limit the influence of any individual record, thereby improving privacy guarantees without requiring extensive post-processing.

Also various post-processing methods can be used to assess and reduce the privacy risk of synthetic data. Their applicability depends on the specific context and type of data:

- *Top and bottom coding* of continuous variables to reduce the influence of extreme outliers that may resemble real individuals too closely.
- *Applying distance metrics* (e.g., record-level closeness measures) to identify and remove synthetic data records that are too similar to real data records. A typical approach is to compare the distribution of synthetic samples with holdout records that were not used in model training. While such assessments are typically performed as part of ex post evaluation, incorporating them earlier in the data generation pipeline can reduce the number of required iterations.
- *Identifying and mitigating rare attribute combinations* that may pose disclosure risks. This includes techniques such as:
  - application of *k*-anonymity-like checks to detect and manage unique low-frequency combinations in synthetic datasets;
  - detecting and removing decision tree leaves or other model structures that capture subgroups with very small counts

5.5.6 Quality metrics

Data quality metrics shall be generated and documented during the anonymisation or synthetic data generation process by the relevant actor, as defined by the applicable use case (data user, HDAB, or Data Holder). Specific tools, as outlined in Section 5.5.8, may be used for this purpose. Privacy metrics are essential inputs for the HDAB's privacy risk assessment process. Fidelity and utility metrics are important for evaluating the applicability of anonymised or synthesised data or an ML model for a particular purpose. In most cases, increasing privacy will decrease utility. A privacy-utility trade-off, which enables the achievement of the data processing activity with acceptable privacy risk should be achieved.

This may require an iterative process in which data processing, metric calculation, and privacy risk assessment are repeated.

5.5.7 Privacy risk assessment

Privacy risk assessment is a shared responsibility between data users and HDABs, with HDABs retaining primary responsibility for oversight and final decisions. The assessment should be based on a risk-based approach, with methods and criteria adapted to the intended use of the exported data or model and the sensitivity and size of the original data.

It evaluates the potential for disclosure of personal or sensitive information and applies to anonymised datasets, synthetic datasets, and machine learning models trained on such data. The assessment is based on analyses of privacy metrics calculated at earlier stages, including record-level closeness measures, distances to closest records, and membership and attribute inference risk indicators. Where applicable, the assessment shall also include analysis of differential privacy implementation, including the related privacy parameters and their impact on privacy risk.

The results of the privacy risk assessment guide decisions on data release, access controls, and the need for further mitigation measures to ensure that residual risks remain within acceptable levels.

Key components of the assessment include:

- **Re-identification risk:** Evaluation of the likelihood that an individual could be re-identified from the data or model outputs, considering both direct matches and probabilistic inference.
- **Membership inference risk:** Assessment of whether it is possible to infer whether a specific individual or record was included in the original dataset used to generate the anonymised or synthetic data or to train a machine learning model.
- **Attribute inference risk:** Assessment of whether sensitive or otherwise protected attributes of an individual could be inferred from the dataset or model outputs, including through generative outputs, model queries, or linkage with auxiliary information.
- **Other relevant considerations:** Additional factors that may affect privacy, such as outlier or rare records, linkage with auxiliary data, model memorisation, and compliance with regulatory requirements.

Most comprehensive approaches for inference risk assessment involve explicit **attack simulations.** Membership inference may be based on generating shadow models that mimic the target model's behavior, while attribute inference attacks aim to predict sensitive attributes from model outputs or synthetic data[45]. These simulations provide a detailed understanding of potential privacy leakage but can be computationally intensive, require specialised expertise, and depend on assumptions about the attacker's knowledge, which can significantly affect the results[46].

As a complement to simulated attacks, **lighter-weight methods** can be used to approximate membership and attribute inference risks based on observable properties of the data or

[45] https://link.springer.com/chapter/10.1007/978-3-032-07884-1_24

[46] https://arxiv.org/pdf/2508.08353

model outputs. These methods typically rely on distance-based metrics, distributional comparisons, and exposure analyses.

Examples of lighter-weight methods include:

- **Record-level distance metrics** (e.g. Distance to Closest Record (DCR), Nearest Neighbour Distance Ratio (NNDR)) to identify synthetic records that are excessively similar to real individuals, which can indicate elevated membership inference risk.
- **Threshold-based similarity tests** (e.g. τ-DCR tests) assessing the proportion of synthetic records that are within a predefined similarity threshold of real records.
- **Outlier-focused analysis**, evaluating privacy metrics specifically for rare or extreme records, which are more vulnerable to inference attacks.
- **Attribute disclosure tests**, measuring whether sensitive attributes can be predicted with higher accuracy from anonymised or synthetic data than from baseline population statistics.
- **Distributional stability checks**, comparing marginal and joint distributions between real and synthetic data to detect overfitting or memorisation.
- **Cumulative exposure analysis**, considering the increased privacy risk caused by prior data releases involving the same individuals, which can amplify the likelihood of re-identification or inference when combined with new anonymised or synthetic datasets.

Evidence from recent research suggests that privacy risk assessments based on inference attacks are more informative and should be prioritised over methods relying solely on record-level similarity metrics[47,48]. Inference attack simulations can be valuable in high-risk or research settings, but may not be feasible in most EHDS scenarios. The appropriate methodology should be selected on a case-by-case basis, taking into account the synthetic data generation or anonymisation scenario, sensitivity and size of used data set, and HDAB capabilities. The requirements for the assessment may be relaxed when differential privacy has been applied during model training to protect ML models and synthetic data, as described in Section 5.5.3.

5.5.8 Tooling

The SPE should provide built-in and controllable tooling to support HDABs in fulfilling their regulatory obligations under the EHDS Regulation. Such support may include anonymisation of data and data processing results, generating synthetic data[49], assessing anonymisation and privacy risks, validating synthetic data generation, and ensuring the traceability and auditability of all processing actions. These tools shall be customised as needed to assist HDABs in fulfilling their responsibilities under the EHDS Regulation. Where appropriate, the tools should also be made available to data users and data holders.

All tools used in the SPE for anonymisation, or synthetic data generation must be either provided or explicitly validated and approved by the HDAB to ensure regulatory compliance

[47] https://link.springer.com/chapter/10.1007/978-3-032-07884-1_24

[48] https://www.sciencedirect.com/science/article/pii/S2666389925001680

[49] Synthetic data generation methods in healthcare: A review on open-source tools and methods - ScienceDirect.

and auditability. The HDAB shall establish processes to maintain a list of approved tools and the conditions for their use. The HDAB shall also provide internal guidance and training to ensure that staff have sufficient competences to use and maintain these tools.

Data user-provided tools. The HDAB may provide mechanisms for approving, installing, or uploading additional tools requested by data users, where these tools support anonymisation, synthetic data generation, or privacy risk assessment.

Documentation generation. Where applicable the tools should create information about the anonymisation or synthetic data generation operations as needed for the documentation referred to in Section 5.5.1.

Privacy metrics and risk assessment. The functionalities below are considered useful in tools supporting privacy and privacy risk assessment. In particular, the functionalities will support the HDAB in carrying out its duties related to disclosure control[50]. Their applicability may depend on the specific context and type of data:

- Identifier and risk factor detection
    - Detect direct identifiers of data subjects, care personnel, and care organisations.
    - Detect quasi-identifiers or outliers that may lead to re-identification.
    - Classify variables and groups of data subjects according to their sensitivity levels enabling tailored protection strategies.
- Privacy metrics
    - Perform similarity analysis between real and anonymised/synthetic datasets using standard distance metrics to quantify differences.
- Re-identification and inference risk assessment
    - Evaluate the overall re-identification risk of datasets and data processing results.
    - Support assessment of re-identification risks associated to membership inference attacks and attribute inference attacks.
    - Identify risks where known individuals may have their presence in the data revealed (membership inference).
    - Detect cases where unknown attributes of known individuals could be inferred from released data (attribute inference).
- Reporting and visualisation
    - Generate structured reports or visual summaries of privacy risk assessments and privacy metric outputs to support decision-making.
    - Parse anonymisation or synthetic data generation metadata to retrieve and visualise parameters used and metrics computed in the earlier processing steps.

Fidelity metrics. The following functionalities are considered useful in tools supporting anonymised and synthetic data fidelity assessment. Their applicability may depend on the specific context and type of data:

- Comparison of dataset variables using standard univariate and multivariate statistical analyses to quantify differences between original and anonymised datasets.

[50] EHDS Regulation Article 73(2).

- Analysis of multivariate statistical relationships, such as comparing feature correlation properties between original and anonymised datasets.
- Visual comparison of original and anonymised datasets or related statistical properties such as heatmaps for comparing correlation matrices.
- ML based methods such as data labelling analysis.

Utility of anonymised or synthetic data. Utility assessment tools should support comparison of machine learning models trained on real versus anonymised or synthetic data[51]. This includes evaluating both the similarity of model outputs (e.g., predicted labels, probability distributions) and the model performance metrics (e.g., accuracy, precision, recall, F1-score).

Utility of ML model. Utility assessment[52] tools should enable evaluation of how the use of anonymised data or differential privacy during training affects model utility. This is typically done by comparing the resulting models to reference models trained on the original, non-anonymised dataset (see also Section 5.5.3). It is important to note that utility tools are not generic, but shall be adapted for the specific task of interest.

Examples of applicable tools. The following list is an exemplary, non-exhaustive list of open-source libraries and tools.

For synthetic data generation:

- Synthcity (Python library that makes many existing synthesis methods for different data types easily accessible. Also implements utility and privacy metrics.)
- Synthpop (R package for fast synthesis for single tables. Showed high utility in several benchmark papers.)
- Synthetic Data Vault (Python library that includes several methods for synthetic data generation. Supports single tables, relational datasets including multiple tables and sequential or time series data. Also offers a quality report including several quality metrics. Please note that there is no real open-source license.)
- REaLTabFormer (Python library, GPT-2-based transformer model that can synthesise single tables and relational datasets including multiple tables.)
- CTGAN (collection of Deep Learning based synthetic data generators for single table data, which are able to learn from real data and generate synthetic data)
- Synthea (synthetic patient population simulator to output synthetic patient data and associated health records in a variety of formats).

For synthetic data privacy evaluation:

- Anonymeter (Python library that evaluates different types of privacy risks (singling out, linkability, inference risks) in synthetic tabular data.)
- Shadow model attacks (Python library that can evaluate the privacy-utility trade-off of synthetic data publishing.)
- TAPAS (Python library that can evaluate the privacy of synthetic data within various attack scenarios.)

[51] Can I trust my fake data – A comprehensive quality assessment framework for synthetic tabular data in healthcare - ScienceDirect.

[52] Frontiers | Comprehensive evaluation framework for synthetic tabular data in health: fidelity, utility and privacy analysis of generative models with and without privacy guarantees.

- SDmetrics (Python library that compares the original data with the synthetic data. Itis model agnostic and provides diagnostic, quality and privacy metrics.)
- Synthcity (Python library that makes many existing synthesis methods for different data types easily accessible. Also implements utility and privacy metrics.)

For synthetic data quality evaluation:

- SDmetrics (Python library that compares the original data with the synthetic data. It is model agnostic and provides diagnostic, quality and privacy metrics.)
- Synthcity (Python library that makes many existing synthesis methods for different data types easily accessible. Also implements utility and privacy metrics.)
- STDG evaluation metrics (Python library to evaluate resemblance, utility and privacy for synthetic tabular data.)

For anonymisation and residual privacy risk assessment

- ARX Data Anonymization Tool (Java software for anonymising sensitive personal data with various privacy models – including k-anonymity, l-diversity, and differential privacy – and providing tools for risk analysis, data transformation, and utility evaluation).
  sdcMicro (R package for applying statistical disclosure control (SDC) to tabular data).
- Amnesia (on-premise anonymisation tool, including pseudonymisation, k-anonymity, masking)
- Greenmask (utility for logical database backup dumping, anonymisation, synthetic data generation and restoration)
- Privacy Meter (library for auditing data privacy specifically within statistical and machine learning algorithms. It helps with data protection impact assessments by using state-of-the-art membership inference attacks)

6 Open questions and recommendations

To support harmonised and risk-based implementation across Member States, future EU-level work could focus on the development of:

- Improved data minimisation methods (e.g., semi-automated processes or tools that help to adequately define the appropriate data granularity)
- Quantitative privacy criteria and recommended parameter ranges for anonymised and synthetic data as applied for various data types and use cases, and to be applied alongside context-based anonymity assessment.
- A standardised, machine-readable metadata format for documenting anonymised and synthetic datasets.

These initiatives would support Health Data Access Bodies (HDABs) in fulfilling their responsibilities under Articles 73, 78 and 79 of the EHDS Regulation, in particular regarding disclosure control, auditability, and quality labelling.

Data minimisation

This paragraph presents some open issues related to data minimisation.
Despite the centrality of data minimisation under Articles 5(1)(c) GDPR and 66(1) EHDS, several implementation challenges remain to be addressed:

- Estimating resources and ensuring process efficiency, particularly for large-scale datasets and time-constrained processing in SPEs;
- Introducing semi-automatised approaches to data minimisation activities that are proposed by data users and then assessed by HDABs (i.e. suggestions given by the data access application management system comparing the single access application form to similar ones; defining harmonised criteria of combinations of quasi identifiers to be categorised into a scale of low-medium-high risk of re-identification, and so on).
- Developing tools that help determining the appropriate level of data granularity to meet the research purpose without exceeding necessity – especially in complex settings such as cohort selection or decision-tree-based models; future work on structured variable selection from standardised ontologies and dictionaries (i.e., OMOP concepts, FHIR resources, LOINC, SNOMED, ICD10 codes) may be helpful in this context.
- Adapting variable selection strategies depending on the type of analysis – i.e, ensuring lawful minimisation while supporting feature-rich datasets in machine learning as opposed to traditional statistical models: while the use of machine learning methods do not relax the requirements of data minimisation, staged access and other process precautions may help reaching the endeavour objects while preserving data protection.
- Offering detailed guidance for data that show high granularity, high sensitivity or strong uniqueness (i.e., genomic data, multi-omics integration, rare disease cohorts, and precision medicine contexts).

Privacy criteria for anonymised and synthetic data
Definition of quantitative privacy criteria for anonymised and synthesised datasets would be highly important for HDAB's but remains an open issue. The relevance of privacy assessment methods and related criteria varies on a case-by-case basis.

It is unlikely that single fixed values for privacy parameters – such as those used in *k*-anonymity, *l*-diversity, or *t*-closeness, or ε for privacy budget in differential privacy – can be defined to cover all use cases. For example, according to responses to the HDAB questionnaire conducted under WP7.2 (see Annex 1 – Methodology), selected k-values ranged between 3 and 100, highlighting the diversity of acceptable thresholds across different cases. Instead of fixed quantitative values, it may be more practical to specify acceptable parameter ranges and examples depending on the data type, intended use, and context.

Anonymisation and synthetic data documentation (metadata) structure

- Anonymised and synthetic data should be documented in a harmonised format to support the HDAB's disclosure control process and any later usage of a dataset exported from an SPE. Thus, a common, machine-readable documentation definition (metadata structure) should be specified as outlined in Section 5.5.1. Such definition

could follow the example of the dataset description template defined for data linkage[53].

The documentation template should be standardised as part of the EHDS implementation and continuously updated to reflect evolving techniques and regulatory expectations.

[53] TEHDAS2 deliverable D7.5

7 Annexes

Annex number	Annex title
1	Methodology
2	Public consultation summary
3	User journey
4	Glossary
5	Anonymisation and synthetic data example scenarios
6	Data minimisation example scenarios

Annex 1 – Methodology

The first input to this guideline was based on a survey we developed over the summer of 2024, which served as scoping aid and starting point. During weekly task meetings the results of the survey were discussed and additional external sources incorporated (such as the EDPB Guideline 01/2025 or scientific literature, for example).

Survey development:

In the preparatory phase, thematic brainstorming sessions were performed internally, followed by drafting the first version of the survey questions. After an internal feedback loop, feedback was provided by the EC, and the result was further distributed to the major and minor contributors to provide comments. The final version was implemented in an online survey tool (i.e., LimeSurvey), published and distributed to the whole TEHDAS2 consortium and further to maximise outreach.

Demographic information from the surveys:

Data minimisation: A total of 122 responses were recorded. Incomplete responses and multiple entries by one Institution were cleaned, which resulted in 29 (full: 25, partial: 4) contributions. The top three personal fields of expertise were: Data management (12), Project management (11) and Data science (9). The most frequent represented countries were Italy (4), Spain (3), Germany (3), Finland (2) and Sweden (2). The top three roles that were represented were: Project coordination (6), Project lead (6) and Director (6). Among the respondents, there were data holders (15), HDABs (12), data users (6), a trusted third party (1) and others (5). The data these respondents are (planning) to make accessible are tables (23), relational databases (19), unstructured data (11), imaging data (9), genomic data (7), bio-sample data (4), and other (3).

Pseudonymisation: A total of 65 responses were recorded. Incomplete responses and multiple entries by one Institution were cleaned, which resulted in 31 (full: 23, partial: 8) contributions. The top three personal fields of expertise were: Data science (14), Project management (13) and Data management (11). The most frequent represented countries were Italy (4), Belgium (3), Germany (3), Finland (2), Sweden (2), Cyprus (2), Luxembourg (2) and Spain (2). The top three roles that were represented were: Project lead (7), Team lead (6), Project coordination (6) and other (6). Among the respondents, there were data holders (15), HDABs (11), data users (7), trusted third parties (3) and others (5). The data these respondents are (planning) to make accessible are tables (22), relational databases (20), unstructured data (15), imaging data (11), genomic data (7), bio-sample data (6), and other (4).

Anonymisation: A total of 43 responses were recorded. Incomplete responses and multiple entries by one Institution were cleaned, which resulted in 27 (full: 23, partial: 4) contributions. The top three personal fields of expertise were: Data science (11), Data management (11) and Project management (10). The most frequent represented countries were Finland (4), Italy (3), Belgium (2), Germany (2), Sweden (2), and Spain (2). The top three roles that were represented were: Project lead (9), Project coordination (7) and other (5). Among the respondents, there were data holders (15), HDABs (8), data users (5), trusted third parties (2) and others (4). The data these respondents are (planning) to make accessible are tables (21), relational databases (17), unstructured data (11), imaging data (11), genomic data (5), bio-sample data (3), and other (3).

Synthetic data: A total of 54 responses were recorded. Incomplete responses and multiple entries by one Institution were cleaned, which resulted in 23 (full: 21, partial: 2) contributions. The top three personal fields of expertise were: Data science (12), Data management (12) and Project management (11). The most frequent represented countries were Italy (4), Germany (3), Finland (2), Belgium (2) Sweden (2), and Spain (2). The top three roles that were represented were: Project lead (7), Project coordination (6) and other (5). Among the respondents, there were data holders (10), HDABs (8), data users (6), trusted third parties (2) and others (5). The data these respondents are (planning) to make accessible are tables (19), relational databases (16), unstructured data (9), imaging data (7), genomic data (6), bio-sample data (4), and other (2).

Annex 2 – Public consultation summary

A draft version of this document was in public consultation between the 30th of September and 30th of November 2025. This document was commented in total 99 times. The number of responses may contain some duplicates as there was no individual identification and verification required to respond to the surveys. Some respondents have also responded both from data holder's and data user's perspective. The responses came from 18 different countries from the EU countries and the European Economic Area countries. Responses from Eastern and Southern European countries and international organisations were largely missing. The respondents were primarily from three main types of organisations, listed in order of prevalence: public organisations, academic/research organisations, and private organisations.

The comments were screened and labelled (labels: relevant and implement; relevant but will not be implemented; relevant but not feasible/postponed; required discussion; irrelevant/out of scope). The comments from the Generic Feedback section were sorted to the chapters they pertained to. Based on the comments, each chapter was adapted, which is described in more detail below. Some of the most frequent comments concerned the concept of relative anonymity and the scarce inclusion of use cases or concrete examples. Now, you can find scenarios for data anonymisation and synthetic data generation in Annex 5 and scenarios for data minimisation in Annex 6 and a paragraph on relative anonymity in the pseudonymisation section.

Data minimisation

A total number of 239 comments specific to data minimisation were drawn from the 99 public consultation responses to D7.2. Among these, around two/thirds come directly from the data minimisation sections of the questionnaire; the rest were mentioned in the Generic Feedback section and yet were linked to data minimisation.

Most of the comments were well articulated and asked for more clarification, more elaboration, the introduction of specific subsections, examples, harmonisation or linkage improvements between TEHDAS2 documents.

More than half of the comments were classified as *Relevant* (61%) to the data minimisation section and treated consequently. Some of them were considered "Relevant but not feasible" within the target of the guideline, that cannot offer specific guidance to each possible study context or domain. The span of the study objectives, methods, and their attributes (i.e., rarity, sensitivity, volume, granularity) can be extremely wide, and this make it difficult to define deterministic sequences of activities and/or precise quantitative risk-utility evaluation to be widely applied. Rather, a high-level framework replicable in all contexts, as the one provided in the data minimisation section, and some domain-specific, expertise-driven suggestions appear to be more realistic to be pursued.

In response to other relevant comments:

- Table 1 has been restructured offering a RACI perspective (Responsible-Accountable-Consulted-Informed) and some clarification on the process.

- Annex 6 has been included, adding some worked examples of significant use cases in which potential process ambiguities and/or data minimisation issues have been elaborated.
- The “Why” dimension has been added to the data dimensions framework, for which some specification and clarification of data features have been provided.
- The “Open Questions and recommendations” section has been integrated with relevant suggestions from the comments received.

Pseudonymisation

A total of 244 comments were received pertaining to pseudonymisation. Public consultation feedback on the pseudonymisation chapter repeatedly asked for (i) clearer articulation of why pseudonymised data are needed in EHDS workflows (incl. federated/longitudinal use cases), (ii) clearer positioning versus anonymisation and especially relation to relative anonymity, and (iii) clearer governance around reversibility, linkage and safeguards.

In response, the pseudonymisation section was revised and expanded as follows:

- Clarified the nature and limits of pseudonymisation (GDPR framing). The text now states explicitly that pseudonymised data remain personal data and that pseudonymisation must be complemented by organisational and technical controls (e.g., contractual restrictions, access controls, encryption) to manage re-identification and broader privacy risks.
- Strengthened EHDS-specific purpose and decision logic (avoid “default-to-anonymised” where unjustified). A new subsection explains that HDABs should not systematically substitute anonymised data where the approved purpose demonstrably requires pseudonymised data, and it links the decision to GDPR necessity/proportionality and to the requirement for reasoned permits.
- Expanded/clarified linkage scenarios and longitudinal needs. The EHDS “user journey” part was refined to better reflect linkage realities, including an explicit scenario for stable pseudonyms over time within a data holder to support longitudinal research, and clearer cross-actor arrangements (HDAB / data holder / TTP / cross-border).
- Tightened requirements and “what counts” as pseudonymisation. The requirements now explicitly state that unsalted hashing of direct identifiers (or similar deterministic schemes enabling regeneration) does not qualify as pseudonymisation, and the requirements text was strengthened around secrets handling, policy choices (deterministic vs randomised), and preventing cross-permit linkage (no reuse of pseudonyms across permits).
- Improved treatment of data subject rights and reversibility. The chapter now more explicitly connects reversible pseudonymisation to EHDS-relevant processes (opt-out; significant findings) and clarifies applicability constraints (e.g., opt-out for new projects).
- Added a dedicated subsection on “relative anonymity”. Given that this was among the most frequent consultation themes, the chapter now includes a paragraph explaining “relative anonymity” (contextual assessment; risk of status change; need for organisational restrictions) and how it may apply in EHDS processing scenarios.

Anonymisation and synthetic data generation

A total of 194 comments were received pertaining to anonymisation and synthetic data generation. A substantial part of the comments received for Section 5 on anonymisation and synthetic data generation where related to the needs for more detailed specifications on the anonymisation and synthetic data generation methods as well as the methodology to assess the related privacy risks. In response to these comments, we have included more detailed texts on the various methods available. For example, we have extended the description concerning the role and use of differential privacy in machine learning model development and we have elaborated on the usage of methods used for privacy risk assessment. We have also added concrete examples concerning different types of use cases. However, due to the extremely wide spectrum of different data types and data usage scenarios it has not been possible to provide detailed specifications, such as definitive criteria for disclosure approval. The role of synthetic data in EHDS context has caused confusion as it is not referred by the regulation. We have strengthened the motivation for covering synthetic data as an approach to privacy protection, particularly for cases where individual-level data would need to be exported from SPEs. The text also clarifies that synthetic data is not inherently anonymous and requires appropriate safeguards and assessment.

Annex 3 – User journey

When a data applicant[54] applies for electronic health data for secondary use purposes, such as research and innovation activities, education, and policy-making, within the EHDS, the user journey consists of several phases (see Figure A1-1). Access for certain purposes (public or occupational health, policy-making and regulatory activities, and statistics) is reserved for public sector bodies and Union institutions (see Chapter IV, Art. 53(1) and 53(2)).

Figure A1-1: EHDS user journey consists of five main phases: data discovery, data access, data preparation, use of data and finalisation.

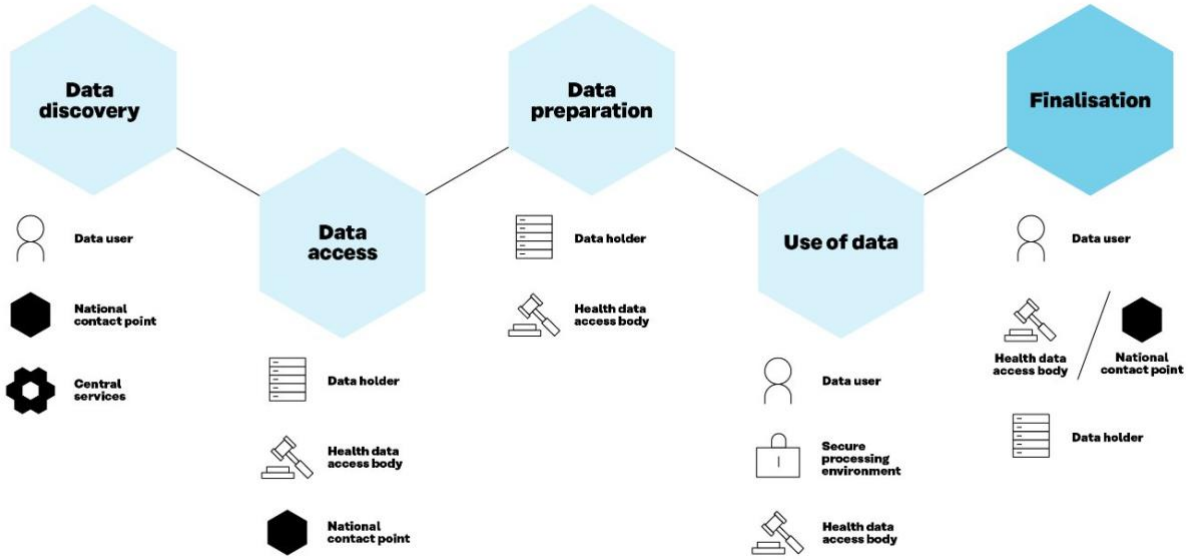

Data discovery

Before being able to use the data, the user needs to investigate whether the data needed is available, and whether it is available in the necessary format for the secondary use purpose. This phase is called data discovery. Datasets available in the EU can be found in a metadata catalogue at https://qa.data.health.europa.eu/. Once the data discovery is completed, the user can begin the process of applying for the data.

Data access

In the data access phase, the user fills in and submits a dedicated and standardised data access application or data request form to a HDAB[55]. The user must complete the information required in the form, upload necessary documents, and provide justifications as needed. The data minimisation principle (as per the GDPR) must be respected to ensure privacy.

[54] Data applicant = a person applying to use electronic health data for a secondary use purpose

[55] Health data access body (HDAB) = the authority responsible for assessing the information provided by the data user who applies for electronic health data for a secondary use purpose

Data access application is used when the user seeks to use personal level data. **Data request** is for cases when the user wants to apply for anonymised statistical data.

Data preparation

During this phase, the data holder(s)[56] deliver(s) the necessary data to the HDAB, which starts to prepare the data for secondary use. Techniques for pseudonymisation, anonymisation, generalisation, suppression, and randomisation of personal data are employed. The data minimisation principle (as per the GDPR) must be respected to ensure privacy.

Use of data

In this phase, the user performs analyses based on the received data for the purpose defined in the application phase. Analysing personal level data must be performed in a secure processing environment (SPE)[57]. The duration of this phase is specified in the Regulation (Art 68(12)).

Finalisation

This last phase of the user journey concerns data user's duties regarding analysis outcomes derived from secondary use of data. Data user must publish the results of secondary use of health data within 18 months of the completion of the data processing in a SPE or of receiving the requested health data. The results should be provided in an anonymous format. The data user must inform the HDAB of the results. In addition, the data user must mention in the output that the results have been obtained by using data in the framework of the EHDS.

[56] Data holder = Any natural or legal person, public authority or other body in the healthcare or the care sectors that has the right or obligation to provide electronic health data for secondary use purposes or the ability to make such data available (see more EHDS Regulation Art. 2 (1t)).

[57] Secure processing environment = an environment with strong technical and security safeguards in which the data user can process personal level electronic health data

Annex 4 – Glossary

Please note that we inserted the definitions below also in the master glossary.

Term	Description
Additional information (related to pseudonymisation)	Additional information is information whose use enables the attribution of ***pseudonymised data*** to identified or identifiable persons (EDPB Guideline 01/2025, Glossary). This term is specific to ***pseudonymisation*** and part of the "additional information" referred to in Regulation (EU) 2016/679 Article 4(5) (GDPR).
Anonymisation	Anonymisation means the process by which personal data are transformed into data that do not relate to an identified or identifiable natural person, taking into account all means reasonably likely to be used, in accordance with Recital 26 of Regulation (EU) 2016/679 (GDPR).
Anonymisation metadata	Where applicable, anonymisation metadata refers to a structured set of detailed information describing (a) the methods and parameters used to anonymise a dataset, and (b) the resulting ***quality metrics*** used to anonymise a dataset or data processing result, or to assess their anonymisation. It includes details e.g., on applied techniques and transformation logs. This metadata helps assess data protection, track modifications, and ensure compliance with anonymisation criteria.
Anonymisation result	The output of anonymisation, which can be an anonymised dataset or a data processing result including ***anonymisation metadata***.
Anonymised statistical format	An anonymised statistical format refers to aggregated data that does not include information on individual data subjects or entities. Aggregation is one possible anonymisation technique.
Attribution of pseudonymised data to data subjects	Process that establishes that ***pseudonymised data*** relate to an already identified person, or links the data to other information with reference to which the data subjects could be identified. (EDPB Guideline 01/2025, Glossary, version adopted for public consultation)
Consistent pseudonymisation	Two sets of data are considered to be pseudonymised consistently if data contained in those sets and relating to the same person can be linked on the basis of the ***pseudonyms*** they contain (EDPB Guideline 01/2025, Glossary). Consistency is context-specific and may be limited to a ***pseudonymisation domain***.

Data aggregation	Process by which information is collected, manipulated and expressed in summary form (ISO/TR 12300:2014(en), 2.1.4)
Data anonymisation framework	A set of processes and practices designed to ensure data privacy through anonymisation and ***privacy risk assessment***.
Data combination	The process of bringing together data from multiple ***datasets*** that can be processed pursuant to one or multiple data permit(s) or data request(s) (Regulation (EU) 2015/327 (EHDS) Articles 57, 68, 69) or other legal basis (such as consent or permits based on other legislation than EHDS). ***Data linkage*** can be part of this process.
Data linkage	The process of combining ***datasets*** "from several sources on one topic or data subject" (ISO 5127:2017, 3.1.11.12). This can be done using unique identifiers, probabilistic methods, or a combination of techniques.
Data minimisation	A principle mandating to only collect, store and process personal data in a manner that is adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed. (GDPR Article 5(1)(c)) Access is only provided to electronic health data that is "adequate, relevant and limited to what is necessary in relation to the purpose of processing indicated in the health data access application by the health data user and in line with the data permit issues pursuant to Article 68." (EHDS Regulation, Article 66(1)) Data minimisation applies to all stages of the data lifecycle.
Data permit	An "administrative decision issued by a health data access body to process certain electronic health data specified in the data permit for specific secondary use purposes, based on conditions laid down in Chapter IV of this Regulation"; (Regulation (EU) 2025/327 (EHDS), Article 2(2)(v))
Data processing result	Refers to outputs from data processing activities carried out by the health data user. It may be generated from statistical analysis or machine learning algorithms, including descriptive statistics, model coefficients, performance indicators, visualisations.
Data protection	Processing data respecting the principles laid down in GDPR Article 5(1). The "implementation of appropriate administrative, technical or physical means to guard against unauthorised intentional or accidental disclosure,

	modification, or destruction of data (ISO/IEC 20944-1:2013(en), 3.6.5.1).
Dataset	"'Dataset' means a structured collection of electronic health data" Regulation (EU) 2025/327 (EHDS), Article 2(2)(w)).
Dataset provenance	Data provenance means a **description of the source** of the data, including context, purpose, method and technology of data generation, documenting agents involved in the provenance of data, data validation routines, source data verification, traceability of changes, and quality control of data (QUANTUM, D1.1).
Direct identifier	A direct identifier is a data element (or set thereof) that has been assigned or is being used to distinguish the data subject it refers to from all others in the given context without requiring the use of ***additional information***. Examples are passport or social security numbers, or the set consisting of first and last name as well as date of birth (EDPB Guideline 01/2025, Glossary, versopn adopted for public consultation).
Fidelity	Fidelity (or resemblance**)** refers to the extent to which processed data – such as anonymised data – retains the statistical properties, relationships, and structural characteristics of the **original/source data**. High fidelity means that distributions, correlations, and key patterns remain unchanged.
Health data access body (HDAB)	Member state-designated authority that facilitates the secondary use of electronic health data. HDABs assess the information provided by the health data applicant and decide on health data requests and access applications, authorise and issue data permits, obtain data from data holders and make data available in secure processing environments. HDABs systematically track the data request and data access applications received and the data permits issued. (EHDS Article 55 and Recital 52)
Health data holder	Any person, organisation or public body involved in healthcare, care services, health-related products, wellness apps or health(care) research, that has the right to process data for health care provision or for public health purposes, reimbursement, research, policy making, official statistics or patient safety. This includes, for example, hospitals, insurers, research institutes and EU institutions. For a more detailed definition: EHDS Regulation, Article 2(2) point (t))

Health data user	A "'natural or legal person, including Union institutions, bodies, offices or agencies, which has been granted lawful access to electronic health data for secondary use pursuant to a data permit, a health data request approval or an access approval by an authorised participant in HealthData@EU;" (Regulation (EU) 2025/327 (EHDS), Article 2(2)(u)).
Irreversible pseudonymisation	A pseudonymisation method where the ***pseudonymising transformation*** cannot be reversed. The information necessary to re-establish the link between the ***pseudonym*** and the ***original data*** has been permanently destroyed or is otherwise unavailable.
Original/source data	Individual-level health data prior to any application of ***pseudonymisation***, ***anonymisation***, or ***synthetic data generation***. It consists of raw data that directly represent real-world individuals.
Privacy (of synthetic or anonymised data)	Privacy measures the extent to which anonymised or synthetic data protects individuals from re-identification, membership inference, or sensitive information leakage. High privacy ensures that no single individual can be traced back to the real dataset, nor can their participation in the ***dataset*** be inferred.
Privacy risk assessment	"Overall process of identifying, analysing, evaluating, consulting, communicating and planning the treatment of potential ***privacy*** impacts with regard to the processing of personally identifiable information (3.7), framed within an organisation's broader risk management framework" (ISO/IEC 29100:2024(en), 3.18). ***Re-identification risk assessment*** falls under privacy risk assessment, together with attribute inference and group membership, for example.
Pseudonym	Identifier that is added to data in the course of the ***pseudonymising transformation*** and set in such a way that it can be attributed to data subjects only using ***additional information***. (EDPB Guideline 01/2025, Glossary, version adopted for public consultation)
Pseudonymisation	The processing of personal data in such a way that the "data can no longer be attributed to a specific data subject without the use of ***additional information***, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person". (Regulation (EU) 2016/679 (GDPR) Article 4(5))

Pseudonymisation domain	Environment in which the controller or processor wishes to preclude attribution of data to specific data subjects. May incorporate persons acting under the authority of the controller or processor, respectively, other natural or legal persons, public authorities, agencies or other bodies, and their respective technological and informational resources. Does not include persons authorised to process additional data allowing the attribution of the ***pseudonymised data*** to data subjects. (EDPB Guideline 01/2025, Glossary, version adopted for public consultation)
Pseudonymisation entity	The entity responsible of processing identifiers into pseudonyms using the pseudonymisation function. It can be a data controller, a data processor (performing pseudonymisation on behalf of a controller), a ***trusted third party*** or a data subject, depending on the pseudonymisation scenario. It should be stressed that, following this definition, the role of the pseudonymisation entity is strictly relevant to the practical implementation of pseudonymisation under a specific scenario. (ENISA, Pseudonymisation techniques and best practices, p. 10 - *modified*)
Pseudonymisation secrets	Data that is used in the application of the ***pseudonymising transformation*** or is created during that process, for example, cryptographic keys or salts, for example. Allows the computation of pseudonyms from certain identifying attributes. Part of ***additional information***. (EDPB Guideline 01/2025, Glossary, version adopted for public consultation)
Pseudonymised data	Result of applying the ***pseudonymising transformation*** to some personal data. Cannot be attributed to a specific data subject without ***additional information***. (EDPB Guideline 01/2025, Glossary, adopted for public consultation)
Pseudonymising controller or processor	Controller or processor that uses pseudonymisation as a safeguard and modifies ***original data*** according to Regulation (EU) 2016/679 (GDPR) Article 4(5). (EDPB Guideline 01/2025, Glossary, version adopted for public consultation)
Pseudonymising transformation	Procedure that modifies ***original data*** in a way that the result cannot be attributed to a specific data subject without ***additional information***. (EDPB Guideline 01/2025, Glossary, version adopted for public consultation)
Public use file	A ***dataset*** made available to the public, typically containing anonymised, synthetic or aggregated

	data to protect individual privacy. These files can be released to data users for information and testing purposes before they apply for a data permit. It is based on **original data**. (Eurostat CROS).
Quality metrics	Refer to qualitative and quantitative indicators used to assess the fitness for purpose of a dataset. In the context of synthetic and anonymised data, quality metrics are particularly relevant to evaluate how transformations affect the data's ***utility***, ***fidelity***, and ***privacy***. Quality metrics may also be used to assess pseudonymised or original datasets, particularly when serving as a benchmark or when evaluating fitness for specific secondary use purposes. (Adapted from ISO and EHDS principles, see Regulation (EU) 2015/327 (EHDS) Article 66 and Recital 58).
Quality metrics evaluation	Refers to the calculation or derivation of the ***quality metrics***.
Quality metrics tool	Quality metrics tool (or "metrics tool") refers to a software, an algorithm, a processing pipeline, a documented manual process, or a combination of these, designed to perform ***quality metrics evaluation***.
Quasi-identifier	A ***dataset*** attribute that, when considered in conjunction with other attributes are sufficient to attribute at least part of the pseudonymised data to data subjects (EDPB Guideline 01/2025, §101, adopted for public consultation).
Re-identification	The “process of associating data in a de-identified ***dataset*** with the original data principal” (i.e., data subject) (ISO/IEC 20889:2018(en), 3.31).
Re-identification risk	The “risk of a successful re-identification attack” (ISO/IEC 20889:2018(en), 3.33), which describes an “action performed on de-identified data by an attacker with the purpose of ***re-identification”*** (ISO/IEC 20889:2018(en), 3.32).
Re-pseudonymisation	The processing of ***pseudonymised data***, where project ***pseudonyms*** are generated using a pseudonymisation algorithm, replacing previously generated pseudonyms. Re-pseudonymisation should not be confused with attempts to reverse the pseudonymisation, which is not meant here.
Reversible pseudonymisation	The ***pseudonymisation entity*** uses a ***pseudonymising transformation*** process that allows the pseudonymisation entity to reverse the ***pseudonym***, if necessary. For example, by using separately kept matching tables of pseudonyms and identifying data, or computable secrets allowing for calculating back to the original input.

Secure Processing Environment (SPE)	A secure processing environment means a controlled environment in which electronic health data can be processed, including for data preparation and access, under the responsibility of a health data access body and in compliance with a data permit, subject to technical and organisational measures ensuring security, confidentiality and traceability, in accordance with Article 73 of Regulation (EU) 2025/327 (EHDS).
Sensitive data	Data with potentially harmful effects in the event of disclosure (i.e., providing access to data to a third party) or misuse (ISO 5127:2017(en), 3.1.10.16)).
Statistical disclosure control	Statistical disclosure control can be defined as the set of "methods to reduce the risk of disclosing information on the statistical units (natural persons, households, economic operators and other undertakings, referred to by the data), usually based on restricting the amount of, or modifying, the data released" (Eurostat CROS).
Synthetic data	Synthetic data means artificially generated data created from an original dataset to reproduce its statistical properties, while not directly corresponding to real individuals. Synthetic data may constitute personal data where individuals remain identifiable, in accordance with Regulation (EU) 2016/679 (GDPR).
Synthetic data documentation	Documentation of a synthetic dataset generated automatically or semi-automatically by the ***synthetic data generator***. The documentation shall be anonymised so that it can be accompanied by the synthetic dataset when released for the data user or for public use.
Synthetic data generator	A synthetic data generator is a software application, model or algorithm designed to generate ***synthetic data***. It uses real-world data as input and generates a synthetic dataset. It is also possible to use parameters derived from the ***original data*** as input and/or modify additional parameters entered by the user.
Trusted health data holder (TDH)	A member state designated health data holder for whom a simplified procedure can be followed for the issuance of data permits. Trusted data holders leverage their expertise on the data they hold to assist the HDAB by providing assessments of data requests or access applications. Once data permits are authorised, these trusted data holders provide the data within an SPE that they manage (Regulation (EU) 2025/327 (EHDS), Article 72 and Recital 72).

Trusted third party (TTP)	A ***pseudonymisation entity*** which is independent from the data user and data holder that processes identifiers into pseudonyms. (ENISA, Pseudonymisation techniques and best practices, p. 10, *modified*). The TTP needs only to know the identifiers of the data subjects on the basis of which it will compute the ***pseudonyms***, and no other data (EDPB Guideline 01/2025, §126, version adopted for public consultation).
Utility	Utility refers to how well the data supports its intended use, such as syntactical testing, analytical tasks, decision-making, or machine learning model performance. In the context of anonymised and synthetic data high utility means that insights, predictions, or outcomes derived from the data closely match those obtained using the ***original data***.

Annex 5– Anonymisation and synthetic data example scenarios

The table below presents anonymisation and synthetic data example scenarios related to the use cases listed in Table 3. Each scenario is characterised by the main activities, which are aligned with the functional groups of the overall architecture shown in Figure 6. The examples are loosely based on real-world cases and have been adapted to the EHDS context.

Example scenario	Original pseudonymised personal health data set	Data processing	Privacy risk assessment	Disclosure decision
1. HDAB creates an anonymised Public Use File replicating the structure and overall characteristics of a real data set to be made available for data users. The file can be used by data users for technical preparation of research projects before data permit approval.[58] **(use case 1, Table 3)**	Health insurance claims data	**Data anonymisation (HDAB)** • Breaking down the correlations between variables. • Coarsening of variables (k-anonymity). • Replacement of pseudonyms with random values. • Picking a simple random 1% sample from the original data. • Publication of only one reporting year per data model.	**Content to be exported** • Anonymised data set released to be freely used and published. **Privacy risk assessment (HDAB)** • Assessment of re-identification and residual disclosure risk, including record uniqueness, small subgroups, and rare attributes. • Assessment of attribute inference and correlation risk, including sensitive attributes and unusual variable combinations.	• Decision on whether the generated data set can be approved for public use, based on anonymity and acceptable privacy risk.
2. HDAB creates a synthetic Public Use File replicating the structure and overall characteristics of real birth registry data. The file is intended to be published for research and policy-making purposes.[59] **(use case 1, Table 3)**	Birth registry	**Synthetic data generation (HDAB)** • Selection of a subset of registry variables and time limits from the original data (data minimisation). • Configuration of a differentially private synthetic data generation pipeline. • Definition of data quality and privacy parameters (e.g., privacy loss budget versus fidelity trade-off). • Hyperparameter configuration, model training, and synthetic data generation using a Bayesian network–based generative model (PrivBayes). • Post-processing of synthetic data to enforce record-level plausibility and consistency constraints (e.g., removal of extreme, implausible, or low-utility synthetic records). • Computation of data fidelity and utility metrics to	**Content to be exported** • Synthetic data set enabled to be freely used and published. **Privacy risk assessment (HDAB)** • Evaluate ε, δ, and other DP settings used during synthetic data generation to ensure they meet organisational privacy policies and acceptable risk levels. • Check for record-level similarity to the original data, attribute inference risks, uniqueness or sparsity of records, and potential privacy risks in small subgroups	• Decision on whether the generated data set can be approved for public use, based on anonymity and acceptable privacy risk.

[58] https://zenodo.org/records/15057924
[59] https://arxiv.org/pdf/2405.00267v1

		be released as metadata together with the Public Use File		
3. HDAB generates a synthetic data file that replicates the structure and overall characteristics of a real dataset. The file maintains a high level of usability while greatly reducing the risk of disclosing personal data. It is made available to data users only in a secure processing environment, but with less stringent data permit conditions than those applied for pseudonymised data[60]. **(use case 2, Table 3)**	Health insurance claims data	**Synthetic data generation (HDAB)** • A dependency graph is trained on the health data to learn probabilities for variable distributions and relationships. • A synthetic population is generated, with demographics, diagnoses, and other variables sampled according to the dependency graph. • Learned transition probabilities between consecutive years are applied to maintain temporal consistency in the synthetic data.	**Content to be exported** • The data set is not allowed to be exported, but intended only for use in the secure processing environment **Privacy risk assessment (HDAB)** • Assessment of residual disclosure and record-level risk, including similarity to original data, record uniqueness, and small or rare subgroups. • Assessment of attribute inference and statistical consistency, including sensitive attributes and marginal/joint distributions. • Relaxed criteria to be applied as the data is not allowed to be exported from the secure processing environment.	• Decision on whether the generated data set can be approved for use within the secure processing environment, based on relaxed criteria for anonymity and acceptable privacy risk.
4. Data user develops predictive machine-learning models to identify population groups with elevated health and social service needs[61]. The data user seeks to export the resulting models and analytical outputs for publication. [62] **(use case 3, Table 3)**.	EHR data	**Data analytics and ML model development (data user)** • Data is pre-processed to ensure quality and usability for analysis and modelling • Features are identified, including demographics, existing diseases, medication and healtcare services usage • Baseline statistics are calculated from the data set • Training and validation of logistic regression and regression tree (XGBoost) models is carried out • Model performance and predictor importance results are collected and presented in graphs and tables.	**Content to be exported** • Descriptive statistics of features and endpoints. • Modelling results, including regression coefficients, feature importance, contingency tables, ROC curves, and performance indicators. **Privacy risk assessment (data user)** • Verify that no individual-level data is exported, including raw records or unaggregated outputs. • Verify that only variables relevant for the study results are exported. •Verify that descriptive statistics and model features meet k-anonymity requirements ($k \geq 5$), avoiding rare individuals in aggregates or feature usage. • Assess potential inference from model outputs, including predictions, feature importance, and interactions that could reveal individual information. **Privacy risk assessment and disclosure decision (HDAB)**	• Decision on whether the requested export of data processing results can be approved based on anonymity and acceptable privacy risk.

[60] https://www.forschungsdatenzentrum-gesundheit.de/infoportal/datenbereitstellungsformen/scientific-use-file

[61] https://www.tandfonline.com/doi/full/10.1080/02813432.2024.2372297

[62] https://findata.fi/en/services-and-instructions/producing-anonymous-results/

			• Assessment of privacy risks based on the documentation submitted by the data user.	
5. Data user has conducted a retrospective study using a pseudonymised dataset[63]. The data user wants to use the original data to create a cross-sectional, synthetic dataset that can be exported from the secure processing environment and used in a real-world clinical setting for demonstration and testing purposes. **(use case 3, Table 3)**.	EHR data	**Synthetic data generation (data user)** • The longitudinal EHR data are transformed into a cross-sectional dataset (one record per patient at a defined time point). • Synthetic data is generated using CART (Classification and Regression Trees) and random sampling methods • Factors with >60 levels are binned or dropped • Dates are aggregated to the month level • No rare or individual-identifying codes are retained. • Computation of data fidelity and utility metrics to be released as metadata together with the Public Use File.	**Content be exported** • Synthetic data set released to be freely used and published. **Privacy risk assessment (data user)** • Propensity Score Mean Squared Error analysis is used to assess if any variable alone carries meaningful identifying signal that could enable record-level re-identification. • Assessment of attribute inference and statistical consistency, including sensitive attributes and marginal/joint distributions. **Privacy risk assessment (HDAB)** • Assessment of privacy risks based on the documentation submitted by the data user.	• Decision on whether the requested synthetic data export can be approved based on anonymity and acceptable privacy risk.
6. Data user develops computed tomography (CT) image analysis algorithms to support early lung cancer diagnosis. The data user seeks to export anonymised CT images and a set of linked EHR data (demography and main diagnoses) from the secure processing environment for the purpose of testing and demonstrating the algorithm performance in operational clinical information systems[64]. **(use case 4, Table 3).**	Image and EHR data	**Image anonymisation (HDAB)** • All identifying information (e.g., name, exact birthdate, address, exposure date) are removed from the image metadata (DICOM header). • All visual identifying features (such as facial features, extremely rare findings or tattoos) are removed from the images by masking or deformation operations. **EHR data anonymisation (HDAB)** • Demography data to be attached to the images is anonymised by appropriate coarsening operations, such as converting birth dates to birth years. • Rare diagnoses or rare diagnosis combinations are removed from the list of diagnoses to be linked with the images. • Exact diagnosis dates are removed and replaced by coarsened dates.	**Content to be exported** • Anonymised images with linked EHR data. **Privacy risk assessment (HDAB)** • Assess whether the anonymisation measures applied are sufficient and meet the established criteria. • Assessment of re-identification and residual disclosure risk, including uniqueness of patient records, rare diagnoses, or small subgroups. • Assessment of attribute inference and correlation risk, including sensitive features in EHR data and identifiable patterns in medical images. • Consider other existing releases of patients' medical images and EHR data and assess their impact on privacy risk.	• Decision on whether the export of the anonymised images and EHR data can be approved based on anonymity and acceptable privacy risk.
7. Data user develops an AI model detecting tissue features in digital pathology imaging[65].	Digital pathology imaging (whole-	**Image anonymisation (data holder, HDAB)** • All identifying information (e.g., name, exact birthdate, address, exposure date) are removed from the image metadata (DICOM header).	**Content to be exported** • Trained AI model • Synthetic WSIs released to be freely used and published for demonstration purposes.	• Decision on whether the trained AI model and requested synthetic data export

[63] https://www.sciencedirect.com/science/article/pii/S2352914825000929

[64] Ministry of Social Affairs and Health (Finland) – guidance on image and signal data processing

[65] https://www.nature.com/articles/s41467-023-37991-y

	slide images, WSI)	• All visual identifying features (such as barcodes or names and dates of birth) are removed from the images by masking. • Rare diagnoses or rare diagnosis combinations are removed from the image data set as they can be themselves revealing. • Before the data is released for the data user: Data holder considers that the same image was not released as a part of other data sets (or under another data permit) possibly accessible to the data user and assess the related linkability risk, as the image could be used to link patients to those data sets and enrich data users knowledge.	**Privacy risk assessment (data user)** • Assessment of the privacy risks associated with the developed AI model (e.g., susceptibility to group membership attacks, or ability of the model to reproduce source data). **Privacy risk assessment (HDAB)** • Assessment of privacy risks based on the documentation submitted by the data user. • Additional privacy risk assessments might be required by the HDAB.	can be approved based on anonymity and acceptable privacy risk.

Annex 6 – Data minimisation example scenarios

Example scenario	Original pseudonymised personal health data set	Example of Data user's journey activities
1. Data user requires information on nephropathic patients both from administrative data & lab test results to evaluate usage & outcomes of immunosuppressive therapies for kidney transplantation (KT).	Administrative data + Electronic Health Records	**Data discovery (data user):** • Data user checks for the availability of relevant information for the study endeavour, select their geographic perimeter of interest and the regional institution offering most of the information needed; Checks from metadata descriptions that information from lab test reports is available from many data holders in the same region. **Data access (data user):** • Data user fills out the Access Application form, specifying detailed information on data minimisation according to the 5 dimensions (Who, What, Where, When, How). Pseudonymised data from multiple datasets and multiple data holders are needed; a nationally agreed procedure for pseudonymisation is available for all data holders. The user is insecure on the level of detail and what information on lab results may be needed. Based on the literature research, there are clear threshold values to categorise values into relevant groups. After a brief consultation with the HDAB, the user agrees that having

		the values of the blood test aggregated into groups may be beneficial, but asks for an explicit inclusion of an attribute specifying if the reported value is in- or outside the biological scale, that may hinder a data-quality problem on the data. Please note that only the blood test that specifically addresses the study object has been selected. **Data access application evaluation (HDAB):** • HDAB performs all the relevant checks needed; variables to be provided by DHs are described in detail and linked to the study objectives or the control of confounders; there are no gaps toward a sound, proportionated and well documented data Permit issuing. **Data pseudonymisation (DH):** • Data comes from 9 data holders: a regional institution owning administrative data, and 8 different providers of lab tests within the regional territory. National ID Numbers are pseudonymised following a national agreed upon reversible algorithm. **Data preparation (DH):** • Each DHs extract data from their datasets, pertaining patients that for the agreed timespan have done a specific blood test. **Data minimisation (DH, HDAB):** • The regional central institution uploads the information of people who did the specific test during the agreed period into the SPE, the immunosuppressive therapies received before, during & after hospital stay, formatted as dates, main active ingredients and pharmaceutical forms, date of KT, length of stay, hospital facility and encrypted information on KT Team, and some demographic information on patients (Year of death, age at KT, sex, region/country of origin of the patient, information regarding previous KT & dialyses). • The lab test providers upload information on agreed pseudonymised patients into the SPE, the specific procedure code, the unit of measurement, the date of the agreed procedure performed and its values.

		• Different Member-states specific alternative data flows are possible (i.e., if lab tests result are available by the regional central institution, DH may become just one). • The HDAB links the data received from DHs; performs quality checks; removes patients who did not undergo KT from the lab test providers' datasets; asks to a national point of contact if there are people who opt-out for secondary use data treatments and removes them from the datasets; encrypts information on KT health facility; transforms procedures and drug dispensation dates to a Year/Month or Days Before/After KT format; groups extreme age values into groups whenever relevant, i.e., children under the age of 10 in this specific context where we have 2/3 patients per year; groups lab test values into agreed ranges and use a specific attribute to indicate extreme values that goes out of the biological scale. At that point, it turns out that one lab test provider uses a different analytical procedure, and results cannot be directly transformed into agreed ranges. HDAB informs the data user, and they find the right choice to be put in place. Different analytical procedures may be differentiated in two fields, or standardised before grouping their values into ranges, or, if not feasible, patients coming from the different lab test provider may be excluded.
2. Data user requires information on retroperitoneal sarcoma patients both from clinical & population-based registries across Europe to investigate patients, tumour and treatment characteristics of retroperitoneal sarcoma (RPS), and to confirm prognostic factors, including hospital case volumes.	Clinical + population-based registries.	**Data discovery (Data user):** • Data user checks for the availability of relevant information for the study endeavour. The main issue is finding relevant clinical registries that have enough specific variables to study this rare disease. Data user explores the EU central platform catalogue. For two hospitals and one international multi-site cancer registry she finds full or nearly full coverage of the needed information. Details are missing for many other hospitals, yet she does not know exactly how many retroperitoneal sarcoma patients have been treated at most of them. She contacts the national main HDAB and receives some suggestion: the endeavour may in principle be divided in two different analyses, one asking aggregated data for hospital case volumes and the second one asking for pseudonymised data on the characteristics of patients and tumours from a selection of the hospitals and from the EURACAN registry. Otherwise, one single data access proposal may fit the use case, reducing the burden of activities for the data user and providing complete vision of the objectives pursued to the HDAB. In any case, if two different proposals are submitted by the data user, the correlations between them shall be taken into consideration and may directly influence choices on data minimisation.

		Data access (Data user): • Data user chooses her national HDAB as the recipient for the data access application form, or otherwise accordingly to art.67, EHDS. She decides to fill out one access application form, yet she specifies that needs aggregated data on case volumes pertaining both RPS and sarcomas in general to investigate the prognostic factor of case volumes, and pseudonymised information specific to RPS from EURACAN registry and two selected hospitals, who can provide detailed information on the tumour characteristics (i.e., size, grade, histology, multifocality), the patient characteristics (i.e., age, sex, country of residence), the treatments received (i.e., surgical resection, chemotherapy, radiotherapy). She specifies also that EURACAN registry, being a population-based registry, should be contacted for the detailed information but not for the information on grading that may be inadequate for sarcomas. In addressing information on the aggregate case volumes, she asks to provide information only on facilities with more than 10 cases treated during the period of interest (i.e., 2010-2019, with/without surgery). For the pseudonymised subset of information, she proposes to convey ICD-O codes of the specific tumour into six major histological groups. **Data access application evaluation (HDAB):** • Once submitted, the data access application form is forwarded automatically to all relevant HDABs: to the HDAB of the member state where EURACAN registry is established, to the HDABs of the two member states where selected hospitals are located. The HDAB responsible to issue the Data Permit will remain the one to which data user have submitted the application. This HDAB will examine the proposal in depth and consult the HDAB of the member state in which the coordinator of the EURACAN registry is established (Art 76, EHDS). Being a rare disease involved in the study, particular attention is posed to the characteristics of patients to be provided and the granularity of variables. Specific treatment data have not been requested by the user. Besides age and sex, that act as quasi-identifiers but are extremely relevant for the study objectives, the user applied for the information on the patients' countries of residence, that at first sight seem a good generalisation of the information regarding the place of residence. She asked also for the information on the treatment's health facilities. Being identity disclosure an

		issue to be carefully tackled in this context, HDAB might encrypt the information on facilities, but this measure would be mostly useless, because based on the numbers of the rare disease and the given nationality of patients, it would be easy to infer the facility name (attribute disclosure). These considerations lead the HDAB to contact the data user, make sure that removing the country of residence from the data provision would not invalidate the study objectives, and then propose a revision of the proposal. Data user will re-submit the data access application proposal, using the previous form compiled, removing information on the country of residence and asking to encrypt the information on health facilities. **Data preparation (DHs):** • EURACAN is informed on the availability of more granular data for grading at the two hospitals involved in the study endeavour and prepares its datasets excluding these two hospitals from the data provision; calculate aggregate measures on case volumes requested. • The two selected hospital that maintain clinical registries extract data from their datasets, pertaining patients that for the agreed diagnoses and timespan have been treated, and calculate aggregate measures on case volumes requested. • Pseudonymised data from the different DHs are conveyed to the SPE specified by the HDAB. To exclude potential bias for patients receiving treatments in more than one facility included in the study, both hospitals who participate to the study, being part of the EURACAN registry, are asked to use EURACAN pseudonymisation technique. **Data minimisation (HDAB):** • The HDAB performs quality checks on received data; asks to the relevant national point of contacts if there are people who opt-out for secondary use data treatments and removes them from the datasets; encrypts information on healthcare facilities; groups diagnosis codes into the agreed six major histological groups.
3. User aims to develop an AI solution for diagnostics in digital	Imaging + Clinical data sets	**Data discovery (Data user):**

pathology, e.g., (a) carcinoma detection in histopathological imaging, and (b) optimum treatment predictor.		• Data user searches for high-resolution scans of digital pathology imaging for a specific disease (e.g., ICD-10 C61 prostate cancer, digital scans of hematoxylin-eosine stained histopathological slides with resolution >=0.25um/px), with slide-level annotations (carcinoma/non-carcinoma). For treatment predictoe [case (b)] looks for cases where besides the digital pathology imaging, data is available about diagnostics (pTNM, UICC stage, Gleason score), and treatments, outcomes of treatments and survival information (survival not only as years but also cause of death available). The data user wants to use data coming from different healthcare institutions to make sure that resulting models are multi-institutionally applicable thanks to being trained and tested on data from different institutions. **Data requesting (Data user, optional):** • The data user asks the HDAB about the total number of cases (patients) per data holder which are fulfilling inclusion criteria and requirements on data availability. **Data access application and permit (Data user, HDAB):** • The data user applies for access to the data, specifying purpose of processing, description of methods used, inclusion criteria, and all known requirements on data (defining fitness for purpose): including requirements on preparation and staining of histopathological slides, requirements on cleaning and scanning of the slides, acceptable file formats, requirements on slide-level annotations [case (a)]. Through the description of the method, the application contains justification for the treatment predictor [case (b)], also requirements on clinical data, namely required semantics of data. Data User justify access to pseudonymised data indicating the need for longitudinal information across a specific period and specifying that any anonymisation attempt would not be in relation to the benefit gained from performing the research, time consuming and costly [66]. The HDAB contacts the DHs and verifies the availability of the data at source. DHs make estimates related to costs of data preparation and data transfers (digital pathology data can be in range of petabytes, where even data transfer can be significant cost) and HDAB composes fees for the permit. **Data preparation (DH, HDAB):** • DHs transform digital pathology imaging to a format acceptable for the data user (e.g., from a proprietary format to DICOM) and validates that all the provided data

[66] https://www.nature.com/articles/s41467-023-37991-y

		is conform to the requirements. In order to enable data minimisation and pseudonymisation, each DH checks if the data contains identifiers as a part of image data (e.g., barcodes, QR codes, names and birth dates of patients). This check may be performed by HDAB too. Slide-level annotations [case (a)] are linked to the images. For the treatment predictors [case (b)], the required clinical data is also extracted and linked to imaging data. **Data minimisation (HDAB):** • The extracted data set is analysed for presence of any data which was not explicitly requested and this information is removed. **Data pseudonymisation (HDAB):** • Data and its related metadata are analysed for presence of any directly identifying information, including identifiers that are used for health care purposes. Pseudonyms are generated, linked to the primary identifiers used in the original data sets, and stored separately and securely as pseudonymisation secrets. All these identifiers are removed – including identifiers that might be present in the image data (where it can be deleted on the pixel-level) and replaced with pseudonyms.