

D6.2 GDPR guidelines for FL implementations

Lead Author: TLX

With contributions from: VHIR, CHUL, GRAD, QUIBIM, INRIA, TVS

Reviewer: Luc Chatty (HL7), Patrick Dufлот (CHUL)

Deliverable nature	Other
Dissemination level	PUB
Delivery date	04-05-2026 [M36]
Version	1.0
Total number of pages	59
Keywords	Federated learning; GDPR; health data; privacy-enhancing technologies; data hubs; secure processing environments; European Health Data Space; Data Governance Act; AI Act; Medical Device Regulation; differential privacy; trusted execution environments; secure multi-party computation.

EXECUTIVE SUMMARY

This deliverable provides GDPR guidelines for federated learning implementations in healthcare, with a focus on the FLUTE platform. It explains how federated learning (FL) can support privacy-preserving collaboration by allowing model training across distributed datasets without centralising raw health data.

The report analyses the application of the GDPR to FL, including the concepts of personal data, anonymisation, pseudonymisation, controllership, joint controllership, processing, and the application of key GDPR principles, such as transparency, purpose limitation, data minimisation, storage limitation, confidentiality, integrity and accountability. It emphasises that data stored at local nodes, model updates, trained models, and model outputs may, depending on the circumstances, qualify as personal data. As a result, FL participants should not assume that data protection obligations fall away merely because raw datasets do not leave the data holder's environment.

The deliverable further examines the interaction between FL and the EU AI Act, the Medical Device Regulation, the Data Governance Act, and the European Health Data Space Regulation. In particular, it focuses on the requirements for AI systems developed through federated learning that qualify as high-risk AI systems where they are used in or as medical devices subject to third-party conformity assessment.

The report also assesses the relevance of the Data Governance Act and the EHDS for federated learning platforms. It concludes that a closed consortium platform is unlikely to qualify as a regulated data intermediation service under the DGA, but this assessment may change if the platform is later opened to an indeterminate group of external data holders and users. Under the EHDS, FL may support secondary use of electronic health data, especially where implemented within secure processing environments or a network of trusted secure processing environments. Still, the exchange of model updates and outputs must be carefully assessed to ensure that personal data are not extracted, reconstructed, or unlawfully disclosed.

From a technical perspective, the deliverable identifies privacy-enhancing technologies and governance processes that can support compliant federated learning. These include Trusted Execution Environments, secure multi-party computation, homomorphic encryption, secret sharing, differential privacy, secure aggregation, robust access controls, logging, local data governance, dataset quality checks, and study agreements that define security and privacy requirements. The report highlights the need to align legal obligations with technical guarantees, so that contractual commitments, system architecture, and algorithmic safeguards operate together.

The deliverable concludes that GDPR-compliant federated learning in healthcare requires a combined legal, technical and organisational framework. Key practical measures include a table with a clear allocation of GDPR obligations to specific FL stakeholders. Finally, the report identifies a need for authoritative EU-level guidance, ideally involving ENISA and data protection authorities, to support practical assessment and certification of federated learning implementations under the GDPR.

DOCUMENT INFORMATION

Grant agreement No.	101095382	Acronym	FLUTE
Full title	Federate Learning and mUlti-party computation Techniques for prostatE cancer		
Call	HORIZON-HLTH-2022-IND-13-02		
Project URL	https://cordis.europa.eu/project/id/101095382 https://www.fluteproject.eu/		
EU project officer	Nihal Yildirim		

Deliverable	Number	D6.2	Title	GDPR guidelines for FL implementations
Work package	Number	WP6	Title	Ethics, regulatory acceptability, GDPR guidelines and standards
Task	Number	T6.2	Title	GDPR guidelines and regulatory liaising

Date of delivery	Contractual	M36	Actual	M36
Status	version 1.0 <input checked="" type="checkbox"/> Final version			
Nature	<input type="checkbox"/> R <input type="checkbox"/> DEM <input type="checkbox"/> DMP <input type="checkbox"/> DEC <input type="checkbox"/> ETHICS <input checked="" type="checkbox"/> OTHER			
Dissemination level	<input checked="" type="checkbox"/> Public <input type="checkbox"/> Sensitive			

Authors (partners)	TLX
Responsible author	Magdalena Kogut-Czarkowska
	magdalena.kogut@timelex.eu

Summary (for dissemination)	<p>This deliverable provides legal, technical and organisational guidelines for GDPR-compliant federated learning in healthcare. It analyses the roles of data holders, data users and platform operators in federated learning ecosystems. The report explains how federated learning can support privacy-preserving research by reducing the need to centralise sensitive data, while identifying residual risks such as inference attacks, model leakage, unclear role allocation and international data access. It also assesses relevant obligations under the AI Act, the Medical Device Regulation, the Data Governance Act and the European Health Data Space Regulation. The deliverable concludes with practical recommendations on GDPR compliance and data governance.</p>
Keywords	<p>Federated learning; GDPR; health data; privacy-enhancing technologies; data hubs; secure processing environments; EHDS; DGA; AI Act; MDR; differential privacy; trusted execution environments; secure multi-party computation.</p>

VERSION LONG			
Issue Date	Rev. No.	Author	Change
08-07-2025	0.1	Magdalena Kogut-Czarkowska	Index
22-12-2025	0.2	Magdalena Kogut-Czarkowska	First draft of the legal part
12-12-2025	0.31	Carlota Cañamero Herrero, Mario López Feijóo	Introduction to FL & TEEs
15-12-2025	0.32	Carlota Cañamero Herrero	TEEs
19-02-2026	0.4	Magdalena Kogut-Czarkowska	Updated versions of sections on GDPR, EHDS and DGA
27-02-2026	0.5	Magdalena Kogut-Czarkowska, Marta Wilinska	Internal review
09-04-2026	0.6	Magdalena Kogut-Czarkowska	Version with input from VHIR, CHUL, GRAD, QUIBIM and INRIA
28-04-2026	0.7	Magdalena Kogut-Czarkowska	Final version for internal review
29-04-2026	0.8	Magdalena Kogut-Czarkowska	Internal peer-review comments addressed
29-04-2026	0.9	Magdalena Kogut-Czarkowska	Clean pre-final version
4-05-2026	1.0	Magdalena Kogut-Czarkowska	Final version

Table of contents

1. Introduction.....	9
2. Legal recommendations for FL under GDPR.....	10
2.1. Literature on federated learning on health data under GDPR	10
2.2. FL stakeholders in the context of FLUTE platform	11
2.3. Applicability of GDPR to FL	11
2.4. Use of personal data in FL in health care context	12
2.5. Roles of the FL participants under GDPR	15
2.6. Legal basis for processing of health data and transparency toward the data subjects	17
2.7. Purpose limitation and data minimisation in FL	18
2.8. Data accuracy FL under GDPR	19
2.9. Storage limitation in FL under GDPR	20
2.10. Confidentiality and integrity in FL under GDPR	20
2.11. Accountability: data sharing agreements and DPIAs	21
2.12. Transfer of personal data to and from EEA	23
3. Legal recommendations for FL under AI Act, MDR, Data Governance Act and EHDS	24
3.1. Artificial Intelligence Act.....	24
3.2. The intersection between the AI Act and the MDR, and its implications for the use of Federated Learning in medical devices	29
3.3. Data Governance Act (DGA)	35
3.4. European Health Data Space Regulation (EHDS)	38
4. Recommendations on appropriate privacy-preserving techniques and processes for FL implementations in healthcare.....	42
4.1. Introduction	42
4.2. Modelling security and privacy requirements and algorithm guarantees	42
4.3. Trusted Execution Environments (TEEs).....	43
4.4. Adaptive encryption techniques	44
4.5. Statistical privacy	44
4.6. Role of data holders and data hubs.....	45
5. Conclusions	47
6. References	50
Annex 1.....	52
Table 1: Data Protection Requirements in FL in Healthcare	52

Table of Figures

Figure 1: ENISA compliance building blocks.....	23
Figure 2: EHDS secondary use framework	38

List of Tables

Table 1: CJEU case law on the notion of personal data.....	14
Table 2: Categories of AI systems based on AI Act risk based approach	25
Table 3: Classification rules for AI high-risk systems.....	26
Table 4: Responsibilities of health data holders, users and HDAB under EHDSR	40
Table 5: Brief description of EHDS bodies.....	41

ABBREVIATIONS AND ACRONYMS

AEPD: Agencia Española de Protección de Datos (the Spanish Data Protection Authority)

AI Act: Artificial Intelligence Act

AI: Artificial Intelligence

API: Application Programming Interface

AWS: Amazon Web Services

B2B: Business-to-business

CE: Conformité Européenne / European conformity marking

CJEU: Court of Justice of European Union

CMP: Consent Management Platform

CT Scans: Computed Tomography Scans

DGA: Data Governance Act

DIS: Data Intermediary Service

DISP: Data Intermediation Service Provider

DP: Differential Privacy

DPA: Data Protection Authority

DPIA: Data Protection Impact Assessment

DPO: Data Protection Officer

EDPB: European Data Protection Board

EDPS: European Data Protection Supervisor

EEA: European Economic Area

EHD: electronic health data

EHDS: European Health Data Space Regulation

EHR: Electronic Health Record

ELSI: Ethical, Legal and Social Implications

ENISA: European Union Agency for Cybersecurity

EU: European Union

FEDAVG: Federated Averaging

FHIR: Fast Healthcare Interoperability Resources

FL: Federated Learning

GA4GH: Global Alliance for Genomics and Health

GDI: European Genomic Data Infrastructure

GDPR: General Data Protection Regulation

HDAB: Health Data Access Body

HRAIS: High-risk AI system(s)

ICO: Information Commissioner's Office (the UK Data Protection Authority)

IEC: Independent Ethics Committee

IMCO: European Parliament Committee on the Internal Market and Consumer Protection

IRB: Institutional Review Board

IT: Information Technology

LIBE: European Parliament Committee on Civil Liberties, Justice and Home Affairs

MDR: Medical Device Regulation

ML: Machine Learning

MPC: Multi-party Computation

NON-IID: Non-Independent And Identically Distributed

OECD: Organisation for Economic Co-operation and Development

PETs: Privacy Enhancing Technologies

PSUR: Periodic Safety Update Report

QMS: Quality Management System

SGX: Software Guard Extensions

SMPC: Secure Multi-party Computation

SNOMED-CT: Systematized Nomenclature of Medicine Clinical Terms

SPE: Secure Processing Environment

SRB: Single Resolution Board

TCF: Transparency and Consent Framework

TEE: Trusted Execution Environment

TEHDAS2: Second Joint Action Towards the European Health Data Space

THDH: trusted health data holder

TIA: Transfer Impact Assessment

TOMs: Technical and Organisational Measures

VIN: vehicles' unique identifying numbers

WP: Work Package

1. Introduction

Federated Learning (FL) is a distributed machine-learning paradigm designed to enable collaborative model development without requiring the centralization of raw data. Instead of aggregating sensitive datasets in a single repository, FL sends a global model to each participating node (such as hospitals, research centres, or edge devices) where local training occurs. The updated model parameters (e.g., gradients or weights) are subsequently returned to a coordinating server, which aggregates them to produce an improved global model. Crucially, raw data never leaves the local environment, making FL an attractive solution for domains governed by strict data-protection regulations, including healthcare [1].

In the context of health data sharing, FL has been applied to multicentre learning on medical imaging, electronic health records (EHRs), genomic data, and intensive-care monitoring. For example, FL has been used to enable collaborative development of deep neural networks for COVID-19 diagnosis from chest CT scans, oncology prognosis prediction, and clinical outcome modelling across international hospital networks [2] [3]. These applications demonstrate that FL can unlock large-scale learning opportunities while respecting institutional boundaries that traditionally impede data sharing.

From a technical perspective, the benefits of FL stem from its privacy-preserving data locality and its capacity to leverage data heterogeneity across distributed nodes. By maintaining data at its source, institutions can adhere more easily to requirements such as General Data Protection Regulation (GDPR) principles of data minimization and purpose limitation. In addition, FL reduces selection and sampling bias by incorporating data from diverse populations, thereby improving model generalizability. Methods such as federated averaging (FedAvg) also reduce communication costs by transmitting only aggregated updates rather than full models or datasets, facilitating large-scale collaborations even among institutions with limited bandwidth [4].

Despite these advantages, FL is subject to several limitations and risks. Technically, FL is vulnerable to privacy-related attacks such as model inversion, which attempts to reconstruct input data from gradients, and membership inference, which seeks to determine whether specific patient records contributed to model updates. Furthermore, the non-independent and identically distributed (non-IID) nature of healthcare data (arising from institutional differences in patient demographics, imaging protocols, or clinical practices) poses challenges to model convergence and robustness. FL systems are also susceptible to poisoning attacks, where malicious participants may introduce corrupted updates to bias the global model. Addressing these vulnerabilities requires complementary techniques such as differential privacy, secure aggregation, homomorphic encryption, and robust optimization strategies [5] [4].

Overall, federated learning provides a technically rigorous and privacy-conscious framework for enabling large-scale collaboration in healthcare AI. While the approach presents clear advantages in terms of privacy preservation and data-sharing compliance, its safe and reliable deployment requires careful consideration of algorithmic robustness, communication security, and governance mechanisms to mitigate residual risks.

This deliverable sets out the FLUTE guidelines for secure and trustworthy federated learning in healthcare. It provides a structured legal and governance framework tailored to health data sharing through federated architectures, with a primary focus on the GDPR and its application in data hubs and federated processing environments. The analysis also integrates relevant provisions of the Data Governance Act, the European Health Data Space and the AI Act. Building on the technical outcomes of the project, the document identifies concrete privacy-preserving techniques (PETs) and organisational measures required to ensure

and demonstrate compliance, including safeguards against inference attacks and data leakage. It addresses cross-border challenges, clarifies legal grey areas, and formulates practical recommendations for developers, operators and supervisory authorities.

2. Legal recommendations for FL under GDPR

2.1. Literature on federated learning on health data under GDPR

The topic of federated learning and its implications under Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation, GDPR) has been critically discussed in scientific papers, raising significant questions about ‘overpromise’ of FL as a silver-bullet for legal and ethical compliance. For instance, Bak et al. [6] argue that federated learning should not be equated with GDPR compliance or privacy preservation. Their core claim is that FL merely reduces data centralisation but does not, by itself, provide legal or technical privacy guarantees. Under GDPR, ethical and legal justifications for processing remain necessary, including a valid legal basis, purpose limitation, and security obligations. A key challenge identified is accountability: because data remain distributed across multiple participants, responsibility for compliance may become fragmented, increasing the risk that no single actor ensures adequate governance. The authors also stress that model updates and gradients can leak personal data, meaning that FL models themselves may qualify as personal data. This undermines the common assumption that FL is ‘privacy-preserving by design’. Further challenges include vulnerability to model poisoning, difficulties in auditing and demonstrating compliance due to the black-box nature of models, and reduced transparency, all of which complicate GDPR obligations relating to accountability, integrity, and fairness.

Brauneck et al. [1] adopt a more operational and compliance-oriented perspective. They accept that FL can support GDPR compliance, but only if accompanied by additional technical and organisational measures under Articles 24, 25, and 32 GDPR. Hence, they recommend actions in three areas: privacy-enhancing measures, documentation, and trust and certification. However, the authors highlight a central challenge for researchers: the vagueness of the GDPR standard of ‘appropriate measures’, which creates legal uncertainty when designing FL systems. They recommend concrete safeguards such as secure multi-party computation, differential privacy, and homomorphic encryption, while recognising the inevitable trade-off between data protection and model performance. Another key challenge is documentation and demonstrability. Researchers must be able to evidence compliance through detailed records of lawful basis, data flows, retention, and model design. The article also points to unresolved blind spots, notably transparency, fairness, and bias detection in federated settings, which remain difficult to assess when data and models are distributed across multiple actors.

In reports published by EDPS [7] and OECD [8], FL is seen as one of the privacy enhancing technologies (PETs), which may help safeguard privacy, however its implementation should be considered on a case-by-case basis. Federated networks have been seen as enablers of collaborative research across distributed datasets without requiring the data to leave its source [1]. At the same time, it has been noted that, by design, they involve multiple stakeholders collaborating across institutional and sometimes national boundaries without centralizing data [9]. Consequently, this decentralization leads to unique legal and organizational questions [10] under the GDPR and other applicable regulations.

2.2. FL stakeholders in the context of FLUTE platform

In this part of the report, we examine in detail the obligations arising from GDPR in the context of FL scenarios and link to the proposed appropriate measures as safeguards to privacy preserving FL to its stakeholders. In the context of FLUTE platform, the following stakeholders need to be distinguished:

- **Platform operator:** The entity that designs, deploys and manages the federated learning infrastructure. It provides and maintains the technical environment enabling model training across distributed datasets without centralising raw data. The platform operator typically allows ‘matching’ between data holder and data user, implements security and privacy safeguards, and ensures the integrity, confidentiality and traceability of the process. Depending on its role and access, it may act as a processor or, in certain configurations, as a controller under data protection law.
- **Data holders (also known as data owners or data providers), including Data Hubs:** Organisations that lawfully control and store the underlying datasets used for federated learning. In healthcare, these are typically hospitals, research institutions or biobanks. They retain custody of the data and perform local model training within their secure environments. Data Hubs are structured intermediaries that facilitate access, governance and coordination of datasets from multiple sources, often providing standardised access procedures, secure processing environments and compliance oversight. Data holders usually determine the purposes and means of the original data processing and are commonly controllers for the underlying health data.
- **Data users:** Entities that seek to train, validate or benefit from federated learning models. They define the research or development objective, submit the model for distributed training, and receive the aggregated model outputs. Data users may be other Data holders, research organisations, pharmaceutical companies, public authorities or other innovators. Depending on their level of influence over the purposes and means of processing, they may qualify as controllers, joint controllers or processors in relation to the federated learning operation.

While the detailed analysis of the requirements is presented below, Annex I provides a summary table that assigns specific responsibilities under GDPR obligations to data holders (including data hubs), platform operator and data users.

2.3. Applicability of GDPR to FL

The GDPR has a broad territorial and material scope. Under Article 3(1), it applies to any controller or processor established in the EU or EEA, irrespective of whether the processing of personal data (see below) itself takes place within the Union. In addition, Article 3(2) extends the GDPR to non-EU organisations that offer goods or services, whether free or paid, to individuals in the EU or EEA, or that monitor their behaviour, including through online tracking or profiling. The GDPR applies across all EU Member States and the wider EEA, including Iceland, Liechtenstein, and Norway, while the United Kingdom operates a largely equivalent framework under the UK GDPR. For small and medium sized enterprises, certain obligations may not apply where processing is not a core activity and does not pose a high risk, such as the obligation to appoint a data protection officer under Article 37, but the general principles of processing set out in Article 5 and the core compliance obligations remain fully applicable.

In the context of FL it’s important to note that there are several stakeholders that are involved, namely: (i) platform operator, (ii) data holders (including data hubs), (iii) data users. Each of these parties will need to assess the application of GDPR to its processing activities. In particular, the fact that the processing itself may be conducted outside of EEA, does not necessarily imply that it is not subject to the requirements of

GDPR. For instance, if a data holder is a controller established in the EU, it would still need to ensure that it has a legal basis for making available the personal data to a data user which is outside of EU. Furthermore, potential establishment of the platform operator outside of EU does not necessarily imply that its activities concerning personal data in the FL infrastructure are outside of the GDPR.

2.4. Use of personal data in FL in health care context

The material scope of the GDPR covers any information relating to an identified or identifiable natural person within the meaning of Article 4(1), including but not limited to names, contact details, IP addresses, location data, and special categories of personal data such as health or biometric data as referred to in Article 9. By contrast, GDPR does not apply to anonymous information i.e. information which does not relate to an identified or identifiable natural person or relates to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable. The concept of personal data has been the subject of much debate, owing it to the interpretations of the term ‘identified’ or ‘identifiable’ by the Court of Justice of the European Union (CJEU or Court). Some of the relevant cases are briefly discussed below.

Name, case reference, date	Facts and findings
Patrick Breyer v. Federal Republic of Germany (C-582/14)	<p>Facts: Mr. Breyer sought to prevent the Federal German institutions running websites that he visited from registering and storing his IP address.</p> <p>Main questions: Can a dynamic IP address can be considered as personal data?</p> <p>Findings:</p> <ul style="list-style-type: none"> • It must be determined whether the possibility of combining a dynamic IP address with the additional data held by the internet service provider constitutes a means likely reasonably to be used to identify the data subject • Identification would not be considered reasonably likely where it is prohibited by law or practically impossible on account of the fact that it requires a disproportionate effort in terms of time, cost and man-power, so that the risk of identification appears in reality to be insignificant • It is not necessary that all information enabling identification of a data subject must be in the hands of one person.
Gesamtverband Autoteile-Handel eV v. Scania CV AB (C-319/22)	<p>Facts: A German trade association of independent wholesalers of vehicle parts, asked Scania to provide it with access to vehicles’ unique identifying numbers (VINs) which it deemed necessary to ensure competition on the motor vehicle aftermarket.</p> <p>Main questions: Can Scania rely on Article 6(1)(c) as a legal basis for disclosing the information about VIN?</p> <p>Findings:</p> <ul style="list-style-type: none"> • Information is to be regarded as personal when, ‘by reason of its content, purpose and effect’ it is linked to a particular person.

	<ul style="list-style-type: none"> • In order to determine whether a natural person is identifiable, directly or indirectly, account should be taken of all the means likely reasonably to be used either by the controller [...] or by any other person, to identify that person, without, however, requiring that all the information enabling that person to be identified should be in the hands of a single entity. • VINs are assigned to a vehicle to ensure the latter's proper identification, they are not – as such – personal data. They may, however, become personal data as regards someone who reasonably has the means to link the VIN to a specific person.
<p>IAB Europe v. Gegevensbeschermingsautoriteit (C-604/22)</p>	<p>Facts: IAB Europe developed Transparency and Consent Framework (TCF) which intended to enable online publishers, data brokers and others, to obtain user's consent and lawfully process their personal data through a Consent Management Platform (CMP). The TCF facilitates capture, through the CMP, of the user's preferences. These preferences are then coded and stored as 'TC String'.</p> <p>Main questions: Are TC Strings personal data?</p> <p>Findings:</p> <ul style="list-style-type: none"> • TC strings contain individual preferences of an individual user in relation to the processing of their personal data. If a combination of a TC string with additional data such as IP address, allows identification of a user, then the TC string contains information concerning an identifiable user and constitutes personal data within the meaning of Article 4(1) GDPR.
<p>EDPS v Single Resolution Board (SRB) (C-413/23)</p>	<p>Facts: The SRB gathered stakeholder's comments regarding proposed after Banco Popular's resolution. The comments were pseudonymised, and then sent to Deloitte for assessment. The SRB alone retained the alphanumeric codes that could link the data back to individuals.</p> <p>Main questions: Should SRB's privacy notice list Deloitte as a recipient of stakeholder's personal data?</p> <p>Findings:</p> <ul style="list-style-type: none"> • Information may relate to an identifiable person by virtue of its content, purpose or effect – and that these criteria operate independently. The content of the data (stakeholder comments) reflected their personal views, thus it was inherently linked to the individual who authored it. • Pseudonymized data are not always personal for every recipient. If adequate safeguards ensure that the recipient cannot identify the data subject using reasonably available means, the information ceases to qualify as personal data for that recipient. • In case of controller's obligation to inform data subjects about data recipients, the identifiable nature of a data subject must be assessed at the time of data collection and from the controller's perspective.

Table 1: CJEU case law on the notion of personal data

Given the above judgements of the CJEU, the following main take-aways need to be taken into account:

- Notion of personal data has boundaries: identification must be realistically possible.
- Identifiability depends on whether the controller or any other recipient is reasonably likely to use available means to identify the individual, taking into account all objective factors: such as cost, time, and technological resources, but not excluding means simply because their use may be restricted by law.
- Data considered anonymous for one entity may become personal if shared with another capable of reidentification through feasible or lawful means.

In the context of FL it needs to be considered *which information* processed in the network could be considered personal. The following could potentially be considered as ‘personal data’:

- Data stored in the nodes, for example, training, testing or validation data shared by the data holders (data hubs),
- Any data that is exchanged between the nodes, including updates/gradients of the models,
- Models trained on the data, as personal data could potentially partially memorised in the resulting local model,
- Future outputs of the models trained on the data through FL.

While each type of data requires its own assessment, taking into account the ‘means reasonably likely to be used’, recent CJEU case law suggests that this assessment may differ for each stakeholder involved in FL. In particular, data held at the nodes may constitute personal data for the data holder, as they may have the means to re-identify data subjects. The same information, however, may be considered non-personal for other participants, where the organisational and technical measures of the FL infrastructure effectively prevent any possibility of re-identification by data users or the operator. From this perspective, FL infrastructures, which by design limit data sharing, may function as mechanisms that prevent re-identification. As noted by EDPS in TechDispatch on Federated Learning [7], ‘*Given that only weights and/or gradients are shared, FL has lower risks from a personal data protection point of view than exchanging the full training datasets. Reconstructing training data from data exchanged in a FL setting is complicated and will only work for a fraction of the training data (the difficulty level and rate of success depend on the FL setup).*’ Still the assessment needs to be performed on a case by case basis.

Scheibner et al. [11] note that although for federated data-sharing approach ‘*the risk or reidentification is reduced compared to the centralized approach, the federated approach remains vulnerable to the same inference attacks of the meta-analysis approach*’. They add that in legal terms the aggregate-level data can potentially be considered personal data and that obfuscation techniques can be used to anonymize the model’s updates at each iteration, however this can negatively affect the performance of the final model (Darzidehkalani [12]).

At the same time, the risk of reconstructions and extraction of personal data from models trained on the data, even in a federated way, still remain. As noted by EDPS [7], ‘*it can be concluded there is a risk that part of the personal data used in the training could be extracted from the resulting ML models*’, in particular extraction attacks can lead to reconstruction of personal data used in the training phase, while membership inference attack can infer if specific data samples were present in the training dataset. The risk has been

analysed in EDPB Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models [13].

2.5. Roles of the FL participants under GDPR

The processing of personal data performed by FL stakeholders (platform operator, data holders and data users) needs to be assessed from the perspective of possible roles under GDPR:

- **Controller:** A natural or legal person, public authority, agency, or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data.
- **Joint controller:** Two or more controllers which jointly determine the purposes and means of the processing of personal data.
- **Processor:** A natural or legal person, public authority, agency, or other body which processes personal data on behalf of the controller.

While, it has been argued that FL improves accountability and auditability, as controllers have clearer oversight of how personal data is processed (EDPS in TechDispatch on Federated Learning [7]), there have been much debate about the roles of each of the stakeholders involved in the processing.

Casaletto et al. [10] argue that GDPR roles in federated analysis are determined by factual circumstances and technical design, not by contractual labels. They distinguish between different types of nodes in federated networks and emphasise that actors should first assess whether they process identifiable personal data and, if so, whether they determine purposes and means. Their preferred model treats each data-contributing node as a controller for its own local processing, while third-party compute providers act as processors. Central nodes and downstream users should not be considered controllers or processors if they only receive non-identifiable outputs and lack ‘means reasonably likely to be used’ for re-identification. They stress that careful technical and organisational design can significantly reduce compliance burdens and the risk of joint controllership.

Becker et al. [14] focus on secure processing environments and the role of data providers (data holders). They argue that when a data provider hosts data in a secure environment and actively enables research by providing infrastructure, tools, and oversight, it is no longer merely disclosing data but is processing it as part of the research project. In such cases, the provider is at least a processor. However, they caution that restricting access, tools, or security measures alone does not automatically lead to joint controllership, which instead depends on factors such as active collaboration in study design or benefit sharing.

Gedeborg et al. [15] emphasise that federated analysis changes traditional assumptions that control follows data transfer. In federated settings, responsibility may be distributed among collaborating research principals even when personal data remain local. They underline that the allocation of controller and processor roles depends on how responsibility for purposes and means is divided and on each party’s actual influence over the analysis. Clear procedures and role definitions are essential to prevent unintended shifts in control and to ensure accountability both towards data subjects and between collaborating parties.

Bernier et al. [16] argue that federated data analysis can reconcile large-scale biomedical research with GDPR, provided that identifiable personal data are not disclosed between nodes. They recommend that each node owner be considered a controller only for its own local processing, without joint controllership arising merely from participation in a federated network. Nodes that receive or combine non-identifiable outputs should not be treated as controllers or processors. They stress the importance of contractual

allocation of responsibilities, technical safeguards, and regulatory guidance to ensure legal certainty and to avoid treating federated analysis as international data transfers or joint controllership by default.

Scheibner et al. [11] take a technology-driven perspective, focusing on privacy-enhancing technologies such as homomorphic encryption and secure multiparty computation. They argue that these techniques can enable joint computation while complying with GDPR and that data processed in this way may be considered anonymous under EU law. Unlike other authors, they accept that joint controllership may still arise between participating entities, even when data are effectively anonymised, although PETs can reduce administrative burdens. They also note that federated learning can enhance accountability and auditability by giving controllers clearer oversight of processing.

Bradshaw et al. [17] approach the issue from a governance and implementation perspective. They highlight that unclear allocation of roles in data sharing agreements creates organisational and legal challenges. They stress that controllers and processors must be clearly identified, and that additional roles such as data custodians, data stewards, and data recipients should also be defined. Clear role definition is presented as a practical necessity for compliant and functional data sharing in complex, multi-party research collaborations.

ENISA in Engineering Personal Data Protection in EU Data Spaces Final report [18], explains that, under the generic data space model, GDPR roles cannot be reliably identified. It is not possible to determine with certainty who acts as controller, whether multiple entities qualify as joint controllers, whether a processor is involved, or whether data users are merely recipients. Although the GDPR attaches specific obligations to each role, the model does not clarify which entity, alone or jointly, determines the purposes and means of processing, which entity acts on behalf of a controller, or which entity receives personal data. ENISA recalls the AEPD's view that the defining element of any processing operation is its purpose (AEPD, Approach To Data Spaces From GDPR Perspective [19]).

In summary, across the literature, there is broad agreement that GDPR roles in federated data sharing depend on factual influence over purposes and means and on access to identifiable personal data, rather than on formal labels. Casaletto et al. [10] and Bernier et al. [16] strongly converge on a design-driven approach that seeks to avoid joint controllership by ensuring that only local nodes process personal data and that all shared outputs are non-identifiable. Becker et al. [14] and Gedeborg et al. [15] introduce more nuance, warning that operating secure environments or influencing research design can still shift roles towards processor or controller status. ENISA [18] states that even though the operator of the infrastructure (data intermediary) may or may not be part of the group of decision-makers (based on being either a controller or a processor), but definitely is one of the entities that need to actually implement the set of PETs chosen. Scheibner et al. [11] depart slightly by accepting joint controllership even in highly privacy-preserving technical setups, while arguing that the data may nevertheless be legally anonymous.

Last but not least, the recent EDPB Guidelines 1/2026 on processing of personal data for scientific research purposes [44], EDPB notes that 'active participation in the determination and the definition of a scientific research protocol, by clearly determining the purpose and essential elements of the means to achieve the objectives pursued, would normally qualify an entity as a controller'. However, this does not imply that the research funder or entity or person consulted in the process of drafting of the protocol (e.g. ethics adviser) should be attributed a controller role.

The practical advice emerging from this literature is to align technical architecture, governance, and contracts: minimise access to identifiable data, clearly delimit decision-making power, and explicitly

allocate responsibilities, while being cautious that increased influence over analysis design or infrastructure can alter GDPR role attribution.

2.6. Legal basis for processing of health data and transparency toward the data subjects

According to Article 5(1)(a) GDPR personal data must be processed lawfully, fairly and in a transparent manner in relation to the data subject. Personal data related to health (which constitutes special category of data), unlike 'regular' personal data, which can be processed if there is a lawful basis under Article 6 GDPR, cannot be processed, unless it meets one of the exceptions outlined in Article 9(2) GDPR. This means that processing health data requires both a lawful basis under Article 6 GDPR and one of the exceptions in Article 9(2) GDPR, such as explicit consent from the data subject.

Processing of personal data is defined broadly in the GDPR as any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction (Article 4(2) GDPR).

In the context of federated learning, several distinct processing operations take place.

First, data holders are typically responsible for **collecting and preparing the data**, for example by pseudonymising it and **adapting it** to agreed standards, and for storing it within the local node. Collection, adaptation and storage qualify as processing under Article 4(2) GDPR. These activities therefore require the data holder to identify and document a valid legal basis.

Second, once a data user identifies a suitable dataset and is granted access, further processing occurs. The **use of data** is expressly included in the definition of processing under Article 4(2) GDPR. Federated learning enables users to query or process data within the local node (Hallock [20]), without granting direct access to the underlying dataset. In practice, a user's AI model can visit the data during federated learning (Thorogood [21]). The user initiates this activity in order to advance specific scientific objectives. Given the broad scope of the GDPR definition, which covers any operation performed on personal data, the analysis carried out at the local node constitutes processing (Gedeborg [15]). Likewise, conducting federated AI training on personal data falls within the scope of the GDPR (Rossello [22]). Although the federated architecture may prevent direct user interaction with the dataset, it is ultimately the user who determines the research purpose and the essential means, such as the specific query or model configuration. In that role, the data user will typically qualify as controller for the research use of health data and – if that data is personal - must demonstrate an appropriate legal basis for that processing.

The storage of data in the federated node by the data holder and the subsequent access and research use by one or more data users therefore pursue different purposes. They constitute separate processing operations, potentially carried out by different controllers. To ensure compliance with the GDPR, each controller involved in the federated network must establish a valid legal basis for its own processing activities.

In practice, a single consent from the data subject may not be sufficient to cover multiple operations performed by different actors within a federated network. Alternative approaches, such as dynamic consent (Kaye et al. [45]), may therefore be required to ensure that all controllers can demonstrate a valid legal basis for their respective processing activities. In its recent guidelines [44], EDPB indicated that: *'Where the purposes of research are not fully known at the time of collecting the personal data, it is possible for controllers to rely on the consent of a data subject to collect and process personal data in a certain*

area of scientific research (so-called broad consent). To rely on broad consent, the controller should process personal data in accordance with ethical standards for scientific research and put additional safeguards in place to compensate for the lack of purpose specification. Controllers can also ask data subjects to consent to different individual research projects, or parts thereof, separately, as soon as the purposes of those projects become known (so-called dynamic consent). A combination of both approaches to consent is also possible.'

Moreover, under Article 13 and 14 GDPR, data subjects have the right to be informed about the processing of their personal data. This information can be communicated orally during data collection or provided in written form either before or during the collection process with the latter being preferable from an accountability perspective. To meet this requirement, information notices must be prepared to clarify various aspects such as the purpose and methods of data processing, the responsible parties involved, and the rights of the data subjects. They should also explain how data subjects may exercise their rights, such as the right to withdraw consent or object to processing, in the context of processing through the federation. In practice, patients would typically need to contact the data holder hosting the dataset that contains their data. When relevant, this information can be included into more general materials given to pilot or study participants. In the context of FL, following the CJEU judgement in the SRB case, the notice should include also information about the potential data users, even if from their perspective the data will not be identifiable.

For transparency and to help potential data users assess the legal bases available to them, the dataset metadata should include the template consent form used by the data holder when collecting the data in the node, together with the related information notice.

2.7. Purpose limitation and data minimisation in FL

Principle of purpose limitation under Article 5(1)(b) GDPR imposes that personal data should only be collected for specified, explicit, and legitimate purposes and not further processed in a manner that is incompatible with those purposes. In particular, the specific purposes for which personal data are processed should be explicit and legitimate and determined at the time of the collection of the personal data. However, further processing for archiving purposes in the public interest, scientific, or historical research purposes or statistical purposes (in accordance with Article 89(1) GDPR) is not considered to be incompatible with the initial purposes. Thus, it is not necessary to undertake the compatibility test, pursuant to Article 6(4) GDPR. Still, the controller must assess whether for the further processing of personal data for scientific research purposes, they can rely on the same legal basis that the initial processing was based [44]. The same guidelines also mention that *'When providing personal data to another controller for scientific research purposes, neither the providing (controller A) nor receiving controller (controller B) needs to undertake a compatibility assessment, pursuant to Article 6(4) GDPR'*.

Connected principle of data minimisation (Article 5(1)(c) GDPR) states that the processing of personal data must be adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed. Personal data should be processed only if the purpose of the processing could not reasonably be fulfilled by other means. This requires, in particular, ensuring that the period for which the personal data are stored is limited to a strict minimum.

FL offers significant benefits by minimising personal data sharing. In particular, the decentralised approach aligns with the core principles of data protection, such as data minimisation and purpose limitation, by ensuring that personal data remains under the control of the data holder and is not exposed to external parties. In FL, instead of sending the whole dataset to other parties only the model parameters or their

updates are transmitted. Additionally, by keeping raw data on local devices/servers and only sharing models or model updates (gradients or weights), FL can enhance the confidentiality of personal data limiting the need for its centralisation and reducing the impact of large-scale data breaches (EHDS, TechDispatch on Federated Learning [7]).

To further uphold the principles of data minimization and purpose limitation, data stored in the nodes should be pseudonymized or anonymized by the data holders. They should agree on the specific anonymization or pseudonymization techniques to be applied, often referred to as data sanitization, while recognising the practical limits of achieving full anonymization of personal data. The objective of data sanitization is to lower the risk of re-identification by removing or modifying both direct identifiers, such as names or identification numbers, and quasi-identifiers. The widely recognized methods for data sanitisation are: generalisation, suppression, perturbation, anatomisation [23]. The fundamental trade-off between data utility and privacy (C4DT [24]) must be acknowledged, thus decision on the anonymization of the data must be balanced with ensuring the utility of the data for given research purpose. Ensuring full anonymisation of data, and health data in particular, is a challenging task, if even possible and attempt to achieve it can destroy important features and reduce the value of the original data. As noted by EDPB, *'training models on data that has been anonymized or heavily minimized may result in decreased performance. This can lead to frustration, as the data scientist feels constrained by limitations that prevent them from achieving optimal results.'*[23]

2.8. Data accuracy FL under GDPR

Pursuant to Article 5(1)(d) GDPR, controllers are required to ensure that personal data are accurate and, where necessary, kept up to date. They must take all reasonable steps, in light of the purposes of the processing, to ensure that inaccurate personal data are rectified or erased without undue delay. This entails, in particular, that controllers accurately record the personal data they collect or receive, as well as the source of those data.

It is important not to confuse the data quality of the training data set (in particular, the characteristics of accuracy and precision) with the GDPR accuracy principle. In an FL setting, checking data quality is more difficult as the data sources are not centralised and not transmitted. Thus, each source of data cannot be compared against the other data sources (no cross-source data quality checks), there are no possibilities to check the data quality of all the training data as a whole and it might be difficult to check the credibility of each data source. Hence, it has been suggested (EDPS, TechDispatch on Federated Learning [7]) that the central platform operator or intermediary plays a role in facilitating exchange of the data in the format in which it receives it from the data holder. It can convert the data into specific formats only to enhance interoperability within and across sectors or if requested by the data user or where mandated by Union law or to ensure harmonisation with international or European data standards. Other potential solutions include establishing: (i) data verification protocols and tools and (ii) safeguards needed to detect and mitigate bias present in the source data. Robust descriptions of the datasets may also offset the challenges with data accuracy.

Some of the solutions recommended [25] include:

- **Data distribution–based evaluation:** The central server assesses statistical characteristics of local datasets to estimate their relevance and assigns higher weight to model updates derived from more representative or valuable data.

- Model utility–based evaluation: Local models are evaluated based on their contribution to improving the performance of the aggregated model in the next training round.
- Statistical metrics–based evaluation: Local models are assessed using predefined statistical performance metrics, such as accuracy or loss, without direct access to the underlying data.

As examined in the context of genomic data, there can be various modes of access to the data in the federation approach in order to limit the amount of the exposed or shared data, while retaining the possibility of the users to inspect the quality of the data and its fit-for-purpose, for example: access to parts of the dataset or returning just aggregated results. However, even with those modes, there are specific challenges for the data user (Rambla [26]).

2.9. Storage limitation in FL under GDPR

Under the principle of storage limitation (Article 5(1)(e) GDPR), personal data should only be kept in a form which permits identification of data subjects for as long as is necessary for the purposes for which the personal data are processed. In order to ensure that the personal data are not kept longer than necessary, time limits should be established by the controller for erasure or for a periodic review.

In the context of federated learning, storage limitation should be enforced through retention rules for (i) local training data on data in the nodes, (ii) intermediate artifacts such as gradients, feature caches, and model checkpoints, and (iii) server-side logs and metadata. Data holders should set a maximum retention period for local data aligned with the training purpose. Data users should delete or anonymize data after use or after a set number of rounds. Model updates may be personal data and should be stored only as long as necessary. Logs should be pseudonymized, with backups following the same limits.

2.10. Confidentiality and integrity in FL under GDPR

Another principle of GDPR is the principle of integrity and confidentiality (Article 5(1)(f) GDPR). Personal data should be processed in a manner that ensures appropriate security and confidentiality of the personal data, including protection against unauthorised or unlawful access to or use of personal data and the equipment used for the processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures. In parallel, the GDPR requires data protection by design and by default and the implementation of appropriate technical and organisational measures, having regard to the state of the art, the cost of implementation, the nature, scope, context and purposes of the processing, and the risks to the rights and freedoms of natural persons. Article 89(1) GDPR requires controllers to adopt of appropriate safeguards to ensure the rights and freedoms of data subjects.

In the context of FL in healthcare, that requires genuine data minimisation, appropriate access controls, pseudonymisation where suitable, secure communications, and adequate traceability. As EDPB clarified *'Appropriate safeguards that may be particularly important when processing genetic or biometric data include, but are not limited to, pseudonymisation, ethical approval, particularly restrictive purpose limitations, federated storage with access through secure processing environments, as well as rigid and role-based access controls'* [44]. Still, the mere use of a distributed architecture is not, in itself, evidence of compliance.

FL setting offers the opportunity to attack the local models as they are trained by the devices and attacking the weakest link can put at risk the whole structure. Thus, the system architecture should be implemented in a way that prioritises data protection by design and by default, ensuring that data access among

federated parties is carried out balancing the level of risk of the processing, accuracy and usefulness of the resulting model (EDPS TechDispatch on Federated Learning [7]).

Another concern for FL is the potential for data leakage through model updates, as even without direct access to raw data, an attacker could infer sensitive information by analysing the gradients or weights shared between devices (and the central server where there is one). This vulnerability opens the door to membership inference attacks, where adversaries can determine whether specific data points were part of the training set.

Moreover, FL must put in place specific distributed training data quality assurance measures, and be free of bias, when data is processed for an intended purpose. (EDPS TechDispatch on Federated Learning [7]). Hence, in the context of FL, safeguards should be implemented on: the data at rest and the models on the nodes; the communication between nodes devices and the central server and the central server itself containing the intermediary and final models. Discussion on the possible confidentiality measures is provided in Section 4.

2.11. Accountability: data sharing agreements and DPIAs

The accountability principle requires the controller to document compliance with data protection obligations and to be able to demonstrate such compliance at any time. In the context of federated learning, a key characteristic of the processing is the involvement of multiple parties acting in different roles, including controllers, joint controllers and processors. A central issue is therefore the clear allocation of responsibilities among all parties participating in the federated set up. Depending on how decisions on the purposes and means of processing are taken, the relationships between the parties may qualify as controller to processor, controller to controller, or joint controllership.

In the context of FL, to ensure compliance with the GDPR, it is necessary to put in place formal agreements governing the exchange of personal data between the parties. These arrangements should be incorporated into a broader study agreement that regulates personal data exchanges within the federated learning platform. The specific type of agreement required depends on the roles assumed by the parties in the processing activities.

- Where one party processes personal data on behalf of another, a **data processing agreement** is mandatory under the Article 28 GDPR. This agreement must clearly define the roles and responsibilities of the controller and the processor and ensure that processing is carried out in compliance with the GDPR. The controller must verify that the processor provides sufficient guarantees to implement appropriate technical and organisational measures to protect the rights of data subjects. The data processing agreement should specify the subject matter and duration of the processing, its nature and purpose, and the types of personal data and categories of data subjects involved. It should also set out the obligations and rights of the controller and the processor, including the requirement that the processor acts solely on documented instructions from the controller. Further elements include conditions for engaging sub processors, confidentiality obligations, security measures, rules on deletion or return of personal data at the end of the service, provisions on demonstrating compliance and supporting audits, and assistance to the controller in fulfilling its GDPR obligations. Template of data processing agreement is available by the European Commission [[link](#)].
- In situations where two or more entities jointly determine the purposes and means of processing, they qualify as joint controllers and must enter into a **joint controller arrangement**. In the context of scientific research projects, this often results in most or all consortium partners being considered joint

controllers. A joint controller arrangement is therefore essential where project partners collectively decide on how and why personal data are processed. This arrangement must transparently allocate responsibilities among the joint controllers, in particular with regard to compliance with data subject rights and the provision of information under Articles 13 and 14 of the GDPR. The essence of the arrangement must be made available to data subjects, and the parties may also designate a single contact point for data subjects. A template of a joint controller agreement can be found here [\[link\]](#).

Brauneck [1] points out that *‘federated learning collaborations must be divided into trustful (i.e., all participants trust each other) and untrustful (i.e., participants do not necessarily trust each other) collaborations.’*, indicating that in case of trustful collaboration *‘enforceable collaboration agreement and the usual principles of collaboration constitute a sufficient safeguard against attacks by participants’*. In case of untrustful collaborations, she suggests the use of trusted intermediates, such as certification of data processing practices and implementation of ISO standards, combined with TEE and third party intermediaries. Kogut-Czarkowska and Shabani [46] provide a more detailed analysis of the challenges and approaches to establishing contractual frameworks in federated repositories.

Another aspect closely connected to the principle of accountability, is conducting a data protection impact assessment. A data protection impact assessment, or DPIA, is a structured assessment required under Article 35 GDPR where processing is likely to result in a high risk to the rights and freedoms of natural persons. It serves to describe the intended processing, assess its necessity and proportionality, identify and evaluate risks to data subjects, and define measures to address those risks. The DPIA is a concrete expression of the accountability principle, as it obliges controllers to anticipate, document, and mitigate data protection risks before the processing begins.

In the context of federated learning, a DPIA must take into account that risks may arise not only from the underlying local datasets but also from the architecture and functioning of the federated system itself. Even where raw data remain locally stored, the exchange of model updates, gradients, or parameters can create risks such as reconstruction attacks, unintended inference, or re-identification. Security vulnerabilities in aggregation servers or in the broader data space infrastructure must also be assessed. Each controller participating in the federated learning setup is required to evaluate these risks within its own DPIA.

Where a data intermediary (platform operator) provides infrastructure or services, it may conduct a risk assessment of its own systems and services and share the identified risks and relevant information with participating controllers. Controllers may rely on this information when preparing their own DPIAs and incorporate the intermediary’s risk analysis into their assessment. However, the responsibility to conduct and complete a DPIA remains with each controller, and reliance on an intermediary’s assessment does not transfer accountability.

Supporting a holistic DPIA in a federated learning environment that involves multiple controllers and possibly several intermediaries is a complex undertaking that requires careful coordination among the parties involved. Clear allocation of roles and responsibilities, detailed mapping of data and model update flows, and alignment of technical and organisational safeguards are necessary. These aspects must be addressed during the engineering and deployment of the overall processing operation, including within EU data spaces, rather than treated as a purely formal or ex post compliance exercise.

ENISA [18] proposes the following main building blocks towards achieving accountability in EU data spaces. While they are addressed at policy makers designing the EU data spaces, similar ‘building blocks’ can be provided for FL scenarios.

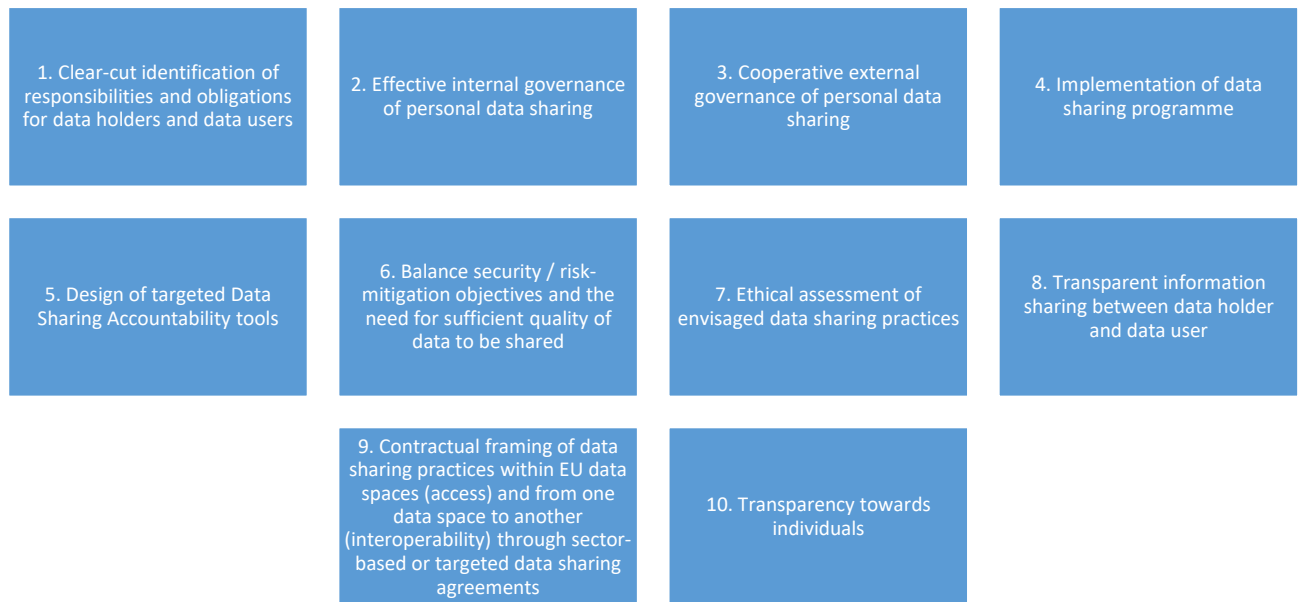


Figure 1: ENISA compliance building blocks

2.12. Transfer of personal data to and from EEA

GDPR puts in place strict conditions (Chapter V GDPR) for transferring personal data to countries outside the European Economic Area which do not provide for adequate level of data protection (non-adequate third countries). Data transfers to such countries require putting in place safeguards listed in Art. 46(1) GDPR (such as standard contractual clauses) and are possible on condition that enforceable data subject rights and effective legal remedies for data subjects are available. To determine those safeguards, the data exporter, together with the data importer, must conduct a transfer impact assessment. Transfer impact assessment (TIA) under the GDPR is a documented evaluation to assess whether personal data transferred to a third country will receive a level of protection essentially equivalent to that guaranteed in the EU. It examines the legal framework and practices of the destination country, particularly access by public authorities, and determines whether supplementary measures are needed to ensure compliance with EU data protection standards.

These strict legal conditions drive the question whether accessing data through the federated network constitutes a data transfer. Answer to this question is an important factor in determining the GDPR compliance of the network, in particular if the network allows data users from non-adequate third countries to use the data in the federation. The term ‘data transfer’ is not defined by GDPR. In its guidelines EDPB [27] provided exhaustive analysis of the criteria of the data transfer, and among them explained that transfer should be understood as an act by which the exporter discloses data by transmission or otherwise making personal data available to another entity (importer). The additional criteria of transfer being that the exporter is subject to the GDPR for the given processing and that the importer is in a third country, irrespective of whether or not this importer is subject to the GDPR for the given processing.

When the data user is using the data in federated networks, the data typically would not ‘leave’ the node [28] [29] [30]. In particular, when federated data is used for federated learning or federated analysis, the

data is not ‘moved’, but the node is ‘visited’ by the AI model (or a user query) and only the computed answers or updated local model parameters [31] are transmitted back to the user. Following the broad EDPB interpretation of the concept of data transfer, even in the first scenario, i.e. when the data in the federated node are made available for research use (processing) by researcher in a non-adequate third country, even without this researcher downloading or copying the data, such use could be considered a data transfer. As result, the broad understanding of data transfers from the EU impose additional compliance obligations of Chapter V GDPR on federated networks which allow data users from third countries. Bernier [16] invited the European Commission to ‘*elaborate the legal treatment of cross-border <<data visitation>> arrangements. Data visitation is the provision of access to data through a controlled, secure processing environment that enables data analyses in the cloud, without enabling data download.*’

3. Legal recommendations for FL under AI Act, MDR, Data Governance Act and EHDS

3.1. Artificial Intelligence Act

3.1.1. Introduction

The European Parliament and the Council Regulation 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (**AI Act**) was published on 12 July 2024 and came into force on 1 August 2024. AI Act can be viewed as product safety law, with its aim is to promote innovation in and the uptake of AI, while ensuring a high level of protection of health, safety and fundamental rights in the Union, including democracy and the rule of law.

The AI Act applies mostly to **AI systems**, which, under Article 3(1) AI Act are defined as *a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate output such as predictions, content, recommendations, or decisions that can influence physical or virtual environments..*

The legal obligations under the AI Act are not applicable equally to all AI systems. The risk classification is **dependent on the purpose for which the AI system is intended to be used**, in line with existing EU product safety legislation¹. This approach is referred to as “risk based approach” and the different categories of AI systems are often described in the following way; the resulting obligations are proportionate to the level of risk the AI system poses.

Category	Description	Example	Legal obligations
Unacceptable risk²	AI systems posing unacceptable risks to fundamental rights and Union values are prohibited	AI system that deploys subliminal techniques beyond a person’s consciousness with the objective of materially distorting the behaviour of a person causing	Not permitted to be placed or deployed in the EU market under Article 5 AI Act

¹ European Commission, Artificial Intelligence – Questions and Answers < [Artificial Intelligence – Q&As](#)>

² https://ai-act-service-desk.ec.europa.eu/sites/default/files/2026-01/guide-prohibited_en.pdf#nameddest=para_44

		them to take a harmful decision	
High-risk	AI systems posing high risks to health, safety and fundamental rights are subject to a set of requirements and obligations	AI-driven medical device software	AI systems that can be placed or deployed in the EU market subject to compliance requirements. See Section 3.1.2 below for details.
Transparency risk	AI systems posing limited transparency risk are subject to transparency obligations under Article 50 AI Act	AI system that generates synthetic audio, image, video or text content	Subject to transparency and information requirements
Minimal or no risk	AI systems posing minimal to no risk	Spam filter	Not regulated by the AI Act; self-regulation is encouraged through use of voluntary codes of conduct ³ .

Table 2: Categories of AI systems based on AI Act risk based approach

3.1.2. High-risk AI Systems (HRAIS)

The classification rules for AI high-risk systems (**HRAIS**) are set out under Article 6 AI Act, and they recognize two different types of HRAIS: a) first one, defined by Article 6(1) AI Act and Annex I, are **harmonised sector specific**, and the b) second one, defined by Article 6(2) AI Act and Annex III, are **context specific**.

Harmonised sector specific (Annex I AI Act)	Context specific (Annex III AI Act)
<p>Two conditions are jointly fulfilled:</p> <ul style="list-style-type: none"> The AI system is intended to be used as a safety component of a product, or the AI system itself is a product under legislation listed in Annex I AI Act; <u>AND</u> The product whose safety component is the AI system, or the AI system itself as a product is required to undergo a third-party conformity assessment, with a view to the placing on the market or putting into service of the product 	<p>Standalone AI systems intended to be used in eight specific areas for very specific purposes which are listed in Annex III AI Act.</p> <p>Classified as high risk due to their significant potential harm to health, safety, fundamental rights, environment, democracy and rule of law unless there are subject to one of the exceptions provided under Article 6(3) AI Act.</p>

³ In addition to the above four categories, the AI Act also regulates general purpose AI models and systems and requires specific obligations to be fulfilled for such AI models and systems.

<p>subject to sectoral legislation listed under Annex I.</p>	
<p>Example:</p> <ul style="list-style-type: none"> • Annex I AI Act includes reference to the Medical Device Regulation (MDR). • Therefore, an AI system that qualifies as a medical device software that is required to undergo a third-party conformity assessment would fall within the definition of a high-risk AI system. • Under the MDR, a medical device of Class IIa, IIb or III are required to complete conformity assessments and therefore fall within the meaning of a high-risk AI system (see Section 3.2 below). The European Commission in its recent study on the deployment of AI in healthcare <i>inter alia</i> confirmed that most healthcare AI applications such as clinical decision support systems largely fall under high-risk category.⁴ 	<p>Example:</p> <ul style="list-style-type: none"> • AI system for creditworthiness assessment used by a retail bank in the EU

Table 3: Classification rules for AI high-risk systems

For a high-risk AI system to be placed on the market, put into service or used in the EU, a number of requirements listed in AI Act will apply to the provider of these systems. Under these requirements provider of the HRAIS must:

- establish a risk management system to identify, evaluate, and mitigate risks throughout entire life-cycle of AI system (Article 9);
- ensure adequate data governance and management practices i.e. training, validation, and testing data must be relevant, representative, free of errors, and complete to the best extent possible (Article 10);
- draw up technical documentation before the system is placed on the market (Article 11);
- ensure traceability (automatic logging functionalities that enable the traceability of results, Article 12);
- enable deployers to interpret the AI system's output by providing clear instructions for use, including system capabilities and limitations (Article 13);
- enable human oversight i.e. design the system to allow effective oversight by natural persons during use (Article 14);

⁴ European Commission: Directorate-General for Health and Food Safety, EEIG, Open Evidence and PwC, Study on the deployment of AI in healthcare – Final report, Publications Office of the European Union, 2025, <https://data.europa.eu/doi/10.2875/2169577>

- ensure an appropriate level of accuracy, robustness and cybersecurity, making sure that system is resilient to errors and attacks (Article 15);
- label the AI system, i.e. indicate on the high-risk AI system (or, where that is not possible, on its packaging or its accompanying documentation, as applicable), their name, registered trade name or registered trademark, the address at which they can be contacted;
- draw up an EU declaration of conformity (Article 47);
- affix the CE marking to the high-risk AI system (Article 48);
- have a quality management system in place (Article 17);
- keep the documentation referred to in Article 18;
- when under their control, keep the logs automatically generated by their high-risk AI systems as referred to in Article 19;
- ensure that the high-risk AI system undergoes the relevant conformity assessment procedure as referred to in Article 43. This conformity should be demonstrated upon a reasonable request by a competent national authority;
- ensure that the high-risk AI system complies with accessibility requirements in accordance with Directives (EU) 2016/2102 and (EU) 2019/882.

Providers must establish a post-market monitoring system to collect and analyse performance data and incidents (Article 61), and must notify serious incidents and malfunctioning to the competent national authority (Article 62), take the necessary corrective actions and provide information as required in Article 20 and cooperate with national competent authorities as required in Article 21. High-risk systems must be registered in the EU-wide database for high-risk systems prior to being placed on the market (Article 60).

3.1.3. Data governance and quality requirements

As set out under Article 10 AI Act, training, validation and testing datasets must meet quality criteria and be subject to appropriate data governance and management. This provision applies to high-risk AI systems trained with data. For systems not using training techniques, requirements apply only to testing data.

Under the **data governance and management** requirements, providers must define and document practices covering:

- Design choices relevant to data use.
- Data collection processes and data origin, including original purpose where personal data are used.
- Data preparation steps, such as annotation, labelling, cleaning, updating, enrichment and aggregation.
- Assumptions about what the data measure and represent.
- Assessment of data availability, quantity and suitability.
- Examination of potential biases affecting health and safety, fundamental rights or leading to unlawful discrimination, including feedback loop risks.
- Measures to detect, prevent and mitigate identified biases.
- Identification of data gaps or shortcomings and how they are addressed.

Under the **contextual relevance and data quality requirements**, the datasets must:

- Reflect, where required by purpose, the geographical, contextual, behavioural or functional setting of use.
 - Be relevant and sufficiently representative for the intended purpose.
 - Be as error-free and complete as reasonably possible.
 - Be statistically appropriate, including in relation to affected persons or groups.
- These requirements may be met through individual datasets or combined datasets.

Use of **special categories of personal data** is exceptionally permitted only where strictly necessary for bias detection and correction, and subject to GDPR, and additionally:

- No effective alternative using other data, including synthetic or anonymised data.
- Technical limits on reuse and state-of-the-art security and privacy measures, including pseudonymisation.
- Strict access controls, safeguards, documentation and confidentiality obligations.
- No sharing or access by other parties.
- Deletion once bias is corrected or retention period ends, whichever is earlier.
- Records of processing must justify strict necessity and lack of alternatives.

3.1.4. AI Act and scientific exceptions

The AI Act sets out a number of general exemptions from its scope, including:

- AI systems or AI models, including their output, specifically developed and put into service for the sole purposes of scientific research and development⁵;
- Any research, testing or development regarding AI systems or AI models prior to their being placed on the market or put into service. Such activities should be conducted in accordance with applicable Union law. Testing in real world conditions should not be covered by that exclusion.⁶

3.1.5. Timeline and proposal for a change of AI Act

Under the AI Act's original schedule, harmonised sector specific HRAIS listed in Annex I, such as medical devices, would have needed to comply with the Act by 2 August 2027.

As part of the Digital Omnibus package presented on 19 November 2025, the Commission has proposed to simplify existing rules on Artificial Intelligence, cybersecurity, and data (Digital Package). Regarding the AI Act, the Commission proposes linking the entry into application of the rules governing high-risk AI systems to the availability of support tools, including the necessary standards. This would mean a delay for the high-risk requirements. Specifically, for AI systems listed in Annex I (Harmonised sector specific), the proposal introduces a two-tier timeline:

- **Tier 1:** an acceleration mechanism: Annex I systems would need to comply 12 months after the Commission issues a decision confirming that sufficient compliance tools—such as harmonised

⁵ Article 2(6) AI Act.

⁶ Article 2(8) AI Act.

standards, common specifications, or administrative guidance—are available. This gives the Commission the option to bring the application date forward when it considers the supporting framework complete.

- **Tier 2**, if no such decision is adopted, a fixed deadline still applies: Annex I systems must comply with the high-risk obligations by 2 August 2028.

On 26 March 2026, the European Parliament, approved amendments to the Digital Omnibus on AI proposal drafted by the IMCO and LIBE committees. The adopted text⁷ sets out Parliament’s final position ahead of negotiations with the Council of the EU, following months of committee-level discussions.

3.2. The intersection between the AI Act and the MDR, and its implications for the use of Federated Learning in medical devices

Where a federated learning system is used for development or as a medical device, the analysis begins with the legal characterisation of the product and its intended purpose. Under Regulation (EU) 2017/745 (Medical Device Regulation or MDR), software may qualify as a medical device where the manufacturer intends it to be used for one or more specific medical purposes, including diagnosis, prediction, prognosis, treatment or alleviation of disease. Software intended to provide information used for diagnostic or therapeutic decision-making is, as a general rule, classified at least as Class IIa, unless its impact places it in Class IIb or III⁸.

In the medical-device context, the principal route to high-risk classification under the AI Act is Article 6(1). An AI system is high-risk, as previously mentioned, where it is itself a product, or a safety component of a product, covered by the Union harmonisation legislation listed in Annex I, and that product is subject to third-party conformity assessment. The MDR is expressly listed in Annex I, Section A, of the AI Act.

It does not follow that every FL system used in healthcare is automatically high-risk. The answer depends on whether the system qualifies as a medical device or safety component, on its intended purpose, on its classification under the MDR, and on whether notified body involvement is required under Article 52 MDR. In practice, medical software will often fall within the high-risk regime because Class IIa, IIb and III devices are subject to third-party conformity assessment, whereas Class I devices generally remain within the scope of self-assessment⁹.

The sequence of analysis is therefore clear. The software or FL system should first be assessed under the MDR. Once that exercise has been completed, it should then be determined whether the conditions of Article 6(1) AI Act are met.

The AI Act and the MDR apply cumulatively, but they should not be treated as separate compliance silos. For high-risk AI systems that are safety components of products covered by Annex I, Section A, the AI Act provides that the requirements in Chapter III, Section 2 are to be addressed through the applicable sectoral conformity-assessment framework. It also allows testing, documentation and reporting processes to be incorporated into documentation already required under sector-specific legislation, so as to avoid duplication. For manufacturers of medical devices incorporating AI systems trained through FL, the practical consequence is the need for a single quality and documentation framework capable of demonstrating compliance with both regimes.

⁷ https://www.europarl.europa.eu/doceo/document/TA-10-2026-0098_EN.html

⁸ Under Rule 11 of Annex VIII MDR.

⁹ Subject to the exceptions set out in Article 52(7) MDR.

As regards regulatory responsibility, the manufacturer of the product is treated as the provider of a HRAIS where that system is a safety component of a product covered by Annex I, Section A, and is placed on the market or put into service under that manufacturer's name or trademark. In the case of an FL-trained medical device, responsibility does not pass to a participating hospital merely because training takes place locally. A different position arises only where the hospital, or another third party, introduces a substantial modification or changes the intended purpose¹⁰.

3.2.1. Data governance, privacy and operational obligations in FL architectures

The fact that clinical data remain at each participating site does not displace the application of the GDPR, nor does it reduce the obligations imposed by the AI Act. Article 2(7) AI Act makes clear that Union law on the protection of personal data, privacy and the confidentiality of communications remains fully applicable. The principles set out in Article 5 GDPR therefore continue to govern processing in FL environments, including lawfulness, fairness and transparency, purpose limitation, data minimisation, accuracy, storage limitation, integrity and confidentiality, and accountability, as further discussed in Section 2 above

Each FL deployment requires a case-specific assessment of who acts as controller, joint controller or processor, since that classification depends on who determines the purposes and means of the processing. Where several participants jointly determine those matters, Article 26 GDPR requires an arrangement setting out their respective responsibilities in a transparent manner. There is no general rule that resolves the point in the abstract. The answer depends on the structure of the project, the governance of the federated network, the allocation of decision-making power, the degree of operational control, and the contractual division of roles.

Nor does FL automatically render information anonymous. The GDPR draws a clear distinction between anonymisation and pseudonymisation and makes equally clear that pseudonymised data remain personal data where attribution to an individual remains reasonably possible by use of additional information. The prudent legal approach is therefore to proceed on the basis that, wherever direct or indirect identification remains reasonably possible, the GDPR continues to apply.

Any processing carried out through FL must rest on a lawful basis under Article 6 GDPR and, where health data or other special categories of data are involved, on a condition under Article 9(2) GDPR. In healthcare settings, the most relevant grounds will often be Article 9(2)(h), concerning medical diagnosis and the provision or management of care, Article 9(2)(i), concerning public interest in the area of public health, or Article 9(2)(j), concerning scientific research or statistical purposes, in each case subject, where necessary, to a supporting basis in Union or Member State law. Article 9(2)(a), explicit consent, may be available in some circumstances, but it should not be assumed to be the default basis in clinical-care settings.

Where the FL system qualifies as a HRAIS, Article 10 AI Act imposes substantive data-governance obligations even where data remain at source. The provider must ensure that training, validation and testing datasets are subject to governance and management practices appropriate to the intended purpose. Those requirements extend to design choices, data collection and origin, relevant preparation operations, the formulation of assumptions, assessment of availability and suitability, the examination of bias, bias detection and mitigation measures, and the identification of data gaps. The datasets must also be relevant, sufficiently representative, and, to the best extent possible, free of errors and complete, having

¹⁰ Within the meaning of Article 25 AI Act.

regard to the intended purpose, and must display appropriate statistical characteristics in relation to the persons or groups on whom the system is intended to be used.

The Regulation expressly recognises that those requirements may be met either at the level of each individual dataset or at the level of the datasets taken together. That is particularly important in multicentre FL networks. The distributed nature of training does not relieve the provider of the burden of showing that the combined dataset architecture is suitable for the relevant clinical purpose. In practice, that requires documented site-onboarding criteria, data-quality standards, local consistency checks, bias-management procedures, and a reasoned justification for the use of the federated dataset as a whole. Those obligations are reinforced by Articles 16 and 17 AI Act, which require the provider to ensure conformity, maintain documentation, operate a quality-management system, and establish appropriate data-governance arrangements.

Article 10(5) AI Act permits, on an exceptional basis, the processing of special categories of personal data for the detection and correction of bias, but only where this is strictly necessary and where that objective cannot effectively be achieved by other means, including synthetic or anonymised data. The provision is narrowly drawn. It is subject to technical limitations on re-use, state-of-the-art security and privacy-preserving measures, including pseudonymisation, strict access controls, a prohibition on onward transmission to third parties, and erasure once the bias has been corrected or the applicable retention period has expired. It should therefore be treated as a limited exception, not as a general licence for broader re-use of sensitive clinical data.

3.2.2. Technical and organisational requirements: accuracy, robustness, cybersecurity and software lifecycle

Article 15 AI Act requires HRAIS to achieve appropriate levels of accuracy, robustness and cybersecurity throughout their lifecycle. The system must be as resilient as possible against errors, faults and inconsistencies, and against attempts by unauthorised third parties to alter its use, outputs or performance by exploiting vulnerabilities. The AI Act expressly refers, where relevant, to data poisoning, model poisoning, adversarial examples, attacks on confidentiality, and model flaws.

Those issues are particularly acute in FL, because the risk surface extends beyond the final device to the interactions between participating sites, coordination layers and update mechanisms. That said, neither the AI Act nor the MDR prescribes a single technical architecture. The legislation does not mandate the use of a secure aggregation server, a particular parameter-sharing method, or any specific encryption technique. The legal requirement is that the measures adopted be appropriate to the risks presented and to the state of the art, and that they be properly selected, justified, documented and maintained.

In the case of medical software, Annex I MDR reinforces that position by requiring software to be developed in accordance with the state of the art, taking account of the software lifecycle, risk management, including information security, and verification and validation. The manufacturer must also specify minimum requirements relating to hardware, IT networks and IT security measures, including protection against unauthorised access.

For an FL-trained medical device, the central legal issue is whether the manufacturer can demonstrate, through its technical documentation and quality-management system, that the measures adopted are sufficient to preserve safety, reliability, training integrity and the confidentiality of the information processed, having regard to the intended purpose of the system and the reasonably foreseeable risks.

3.2.3. Transparency, human oversight and allocation of obligations between manufacturer and hospital

Article 13 AI Act does not require every model to be fully interpretable in a strict technical sense. What it requires is a degree of transparency sufficient to enable the final user (referred to as deployer under the AI Act¹¹) to interpret the system's output and use it properly. The instructions for use must include the intended purpose, the capabilities and limitations of the system, the expected levels of accuracy, robustness and cybersecurity, the known or foreseeable circumstances that may affect performance, and, where appropriate, the system's ability to provide information relevant to explaining its output.

In FL settings, that requires documentation not only of the model itself, but also of the limitations arising from data heterogeneity and from the clinical settings in which the system was trained and validated. If the model is still operating with the use of federated infrastructure, a hospital or clinician must be in a position to understand the boundaries of safe use, not merely the existence of the model.

Article 14 AI Act requires high-risk systems to be capable of being effectively overseen by natural persons during use. In addition, Article 26 places specific duties on the deployer. These include the duty to use the system in accordance with the instructions for use, to assign human oversight to persons with the necessary competence, training and authority, to ensure that input data are relevant and sufficiently representative where those data are under its control, to monitor the operation of the system, to inform the provider and the competent authority where risks are identified, and to retain logs where they are under their control for an appropriate period and, at a minimum, six (6) months, unless sector-specific law provides otherwise.

In a hospital setting, this requires more than technical implementation. It requires an effective internal structure for oversight, monitoring and incident management. The operational use of the system must remain capable of meaningful human review.

A hospital or another third party may itself become a provider where it introduces a substantial modification to a system already placed on the market or changes its intended purpose in such a way that the system remains, or becomes, high-risk. That issue is particularly relevant in FL projects involving local retraining, parameter adjustment or centre-specific configuration. The manufacturer should therefore define, both contractually and by design, which local changes are permitted, which remain within the validated perimeter, and which would trigger a fresh allocation of regulatory responsibility under Article 25 AI Act.

3.2.4. Post-market monitoring, incidents and adaptive FL systems

Both the AI Act and the MDR require post-market monitoring. Under Article 72 AI Act, the provider must establish and document a post-market monitoring system proportionate to the nature of the technology and the risks of the system, capable of actively and systematically collecting, documenting and analysing relevant data on the system's performance throughout its lifecycle. Under the MDR, the manufacturer must plan, establish, document, implement, maintain and update a post-market surveillance system integrated within its quality-management system. That system must serve to update the benefit-risk determination, risk management, clinical evaluation, design, manufacture, instructions for use and, where appropriate, to trigger preventive or corrective action.

¹¹ Under Article 3(4) AI Act, 'deployer' means a natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity.

For medical devices falling within the scope of the MDR, the AI Act allows post-market monitoring requirements to be integrated into the systems already required under sector-specific legislation, provided that an equivalent level of protection is ensured. The object is not duplication of records, but a coherent compliance structure capable of satisfying both regimes.

Where an AI system continues to learn after being put into service, the key legal issue is whether a substantial modification has taken place. The AI Act requires a new conformity assessment where a HRAIS undergoes such a modification. At the same time, changes to the performance or algorithm of a system that continues to learn do not amount to a substantial modification where those changes have been predetermined by the provider in the initial conformity assessment and are reflected in the technical documentation.

For an AI-based medical device, that means the permitted change perimeter must be defined in advance, together with validation criteria, rollback conditions, and the circumstances in which further regulatory assessment becomes necessary. Continuous learning cannot be treated as an open-ended operational discretion. It must remain within a pre-defined and documented regulatory envelope.

The position on impact assessments should also be stated precisely. The Fundamental Rights Impact Assessment under Article 27 AI Act is not required merely because an AI system in a medical device qualifies as high-risk under Article 6(1). Article 27 applies to certain high-risk systems under Article 6(2), namely the Annex III route, and to the *deployers* expressly identified in the Regulation. That should be distinguished from the DPIA under Article 35 GDPR, which will commonly be required where the processing, particularly through the use of new technologies, is likely to result in a high risk to the rights and freedoms of natural persons. The GDPR expressly refers, among other situations, to systematic and extensive evaluation based on automated processing and to the large-scale processing of special categories of data, including health data.

3.2.5. Recommendations for the FL stakeholders under AI Act and MDR

In light of the foregoing, any deployment of AI in or as a medical device, including one that was developed or tested with the use of FL, should be supported by a compliance framework that addresses, in an integrated manner, the requirements of the AI Act, the MDR and, where personal data are processed, the GDPR.

The following measures should be treated as minimum practical steps:

- The manufacturer should undertake a formal classification exercise at the outset of the project. That exercise should determine whether the AI-enabled software qualifies as a medical device under the MDR, including by reference to Rule 11 of Annex VIII, and whether the system falls within the high-risk regime under Article 6(1) AI Act. That analysis should be recorded in writing and retained as part of the technical documentation.
- The intended purpose of the system should be defined with precision and kept under strict change control. In particular, the documentation should identify the medical function of the system, the intended users, the relevant clinical setting, the nature of the input data, and the limits of use. That description is central to qualification and classification under the MDR and also determines the scope of the provider's obligations under the AI Act.
- The parties participating in the federated architecture should adopt a clear governance model for data protection purposes. Each project should identify, on a case-specific basis, who acts as controller, joint

controller or processor. Where the purposes and means of processing are jointly determined, the parties should enter into an arrangement compliant with Article 26 GDPR, allocating responsibilities in a transparent and operationally effective manner.

- The lawful basis for all personal-data processing should be identified and documented before deployment. Reliance on a federated architecture does not remove the need for a proper legal basis under Article 6 GDPR and, where health data or other special categories of data are involved, a valid condition under Article 9(2) GDPR. Where Union or Member State law is required to support the relevant condition, that supporting basis should be identified expressly in the compliance file.
- The provider should establish and maintain a dedicated Article 10 AI Act data-governance file. That file should address the origin of the data, collection logic, preparation steps, assumptions, assessment of suitability and availability, bias risks, mitigation measures, and data gaps. In multicentre FL deployments, the provider should be in a position to justify not only the adequacy of each local dataset, but also why the datasets, taken together, are sufficiently representative for the intended clinical purpose.
- The FL architecture should be designed and operated in accordance with the principles of data protection by design and by default. In practice, this requires genuine data minimisation, appropriate access controls, pseudonymisation where suitable, secure communications, and traceability measures capable of supporting both accountability and incident analysis. The fact that data remain at source should not be treated as sufficient, in itself, to establish compliance with the GDPR.
- Cybersecurity and model integrity should be treated as regulated design requirements rather than purely technical matters. The manufacturer should conduct a risk-based assessment of threats relevant to FL, including poisoning attacks, confidentiality attacks, unauthorised alteration of outputs or performance, and vulnerabilities affecting the training process. The technical and organisational measures adopted in response should then be documented so as to demonstrate compliance with Article 15 AI Act and with the MDR requirements relating to software lifecycle, risk management and information security.
- The instructions for use should be drafted so as to be meaningful in practice. They should set out, in clear terms, the intended purpose of the system, its performance limitations, known constraints, foreseeable failure conditions, and the circumstances in which human review, non-use or escalation is required. In a clinical setting, compliance with Article 13 AI Act depends not on abstract transparency, but on whether the user is given sufficient information to use the system safely and appropriately.
- The manufacturer should define formal limits on local retraining, parameter adjustment and site-specific configuration. The documentation and contractual framework should distinguish between changes that are permitted within the validated perimeter and changes that would amount to a substantial modification or a change in intended purpose. This is necessary both to preserve regulatory control over the system and to avoid an unintended transfer of responsibility under Article 25 AI Act.
- Post-market monitoring should be integrated into the manufacturer's existing surveillance framework under the MDR. A single process should be established to capture incidents, performance degradation, drift indicators, corrective actions and update decisions affecting the AI system. Where the system falls within the MDR, the provider should align the requirements of Article 72 AI Act with the manufacturer's post-market surveillance obligations under the MDR, including Article 83 and, where applicable, the preparation of a Periodic Safety Update Report under Article 86.

3.3. Data Governance Act (DGA)

3.3.1. Introduction

Data Governance Act Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act, DGA) [33] is a EU regulation aimed to increase trust in data sharing, boost data availability, and foster a single market for data. It facilitates the reuse of protected public sector data, regulates data intermediation services, and encourages data altruism.

DGA introduced data intermediation service (DIS), a new business model for entities acting as intermediaries between data holders and data users. DIS means a service which aims to **establish commercial relationships for the purposes of data sharing** between an undetermined number of data subjects and data holders on the one hand and data users on the other, through technical, legal or other means, including for the purpose of exercising the rights of data subjects in relation to personal data.

As pointed out by scholars, *'the regulation under the DGA is aimed at both strengthening these data intermediaries by establishing trust in them and, simultaneously, at pre-emptively curbing potential abuse of their intermediary position'* [34]. There are different types of data intermediaries under DGA, however in the context of FLUTE, this report focuses only on business-to-business (B2B) data intermediaries: data marketplaces and industrial data platforms.

- **Data marketplace** is a typical model of a two-sided matching platforms [35], where data holders can offer their data to potential data users, while users can browse different data offerings to find the purpose-specific data they require.
- **Industrial data platforms** do not function as matchmakers between data holders and users. Rather, they primarily aim at providing the technical infrastructure for companies to share data with each other as part of their broader collaboration [34].

However, the following services are excluded from the regulation of a DIS under DGA:

- services that obtain data from data holders and aggregate, enrich or transform the data for the purpose of adding substantial value to it and license the use of the resulting data to data users, without establishing a commercial relationship between data holders and data users (example: data brokers)
- services that focus on the intermediation of copyright-protected content;
- services that are exclusively used by one data holder in order to enable the use of the data held by that data holder, or that are used by multiple legal persons in a closed group, including supplier or customer relationships or collaborations established by contract, in particular those that have as a main objective to ensure the functionalities of objects and devices connected to the Internet of Things;
- data sharing services offered by public sector bodies that do not aim to establish commercial relationships.

3.3.2. Requirements for data intermediation service providers under DGA

Entity which intends to provide the DIS must submit a notification to the competent authority in the Member State where it is established. Acceptance of the notification is not required. DIS providers (DISPs) located outside the EU have to appoint a representative in the EU. The services are to be supervised by competent public authorities that the member states are obliged to designate (Article 13 DGA). DIS providers must observe a number of requirements when providing the service.

Under the DGA, DISPs must act as neutral facilitators. They may not use the data they intermediate for any purpose other than making it available to users. This excludes analytics, profiling, or any secondary exploitation for their own benefit. Any activity related data, such as time, location, or duration of use, may only be processed for service development purposes like fraud detection or cybersecurity and must be shared with data holders upon request.

Structural separation is a core requirement. Data intermediation services must be provided through a separate legal entity and be operationally isolated from other services within the same corporate group. Users may not receive more favourable conditions because they also use other group services. Commercial terms, including pricing, must be independent and not bundled or cross subsidised. Only auxiliary services strictly supporting data exchange, such as anonymisation or formatting, may be offered without breaching this separation.

DISPs must facilitate data exchange in the original format. Data conversion is permitted only where necessary for interoperability, at the user's request, to comply with the EU law, or for harmonisation purposes. Where conversion is not legally mandated, users must be offered an opt out. Additional tools such as storage, curation, anonymisation, or pseudonymisation may be provided only upon explicit request or approval, and providers must ensure that third party tools are not misused.

Access to the service must be fair, transparent, and non-discriminatory for all parties. This applies equally to pricing, contractual terms, and technical conditions. DISPs are required to implement procedures to prevent fraudulent or abusive access and to maintain detailed logs of all data intermediation activities.

Strong operational safeguards are required. DISPs must ensure appropriate security for the storage, processing, and transmission of non-personal data, with the highest level of protection for competitively sensitive information. They must also put in place measures to prevent unlawful access or transfers under Union or national law and must immediately inform data holders of any unauthorised transfer, access, or use of non-personal data.

Interoperability is another obligation. DISPs must take measures to ensure interoperability with other data intermediation services, relying on open standards where possible. In the event of insolvency, they must ensure continuity of access and enable data holders and users to access, retrieve, or transfer their data, while also ensuring that data subjects can continue to exercise their rights.

Where personal data and data subject rights are involved, DISPs must act in the best interests of data subjects. They must provide clear, transparent, and accessible information on intended data uses and applicable terms before consent is given. If consent or permission management tools are offered, these must specify any intended third country jurisdiction for data use and allow data subjects to easily give and withdraw consent or permissions.

Currently, a number of data intermediation services have been established in the EU [\[link\]](#). The providers of those services function as neutral third parties that connect data holders with data users. While they may charge for facilitating the data sharing between the parties, they cannot directly use the data that they intermediate for financial profit (e.g. by selling it to another company or using it to develop their own product based on this data).

3.3.3. Challenges of establishing platforms under DGA

According to recital 33 in the DGA, data intermediation services should be provided through a legal person that is separate from other activities of a given data intermediation service provider. The goal of this structural separation is to avoid conflicts of interests [36]. Consequently, however the strict measures

imposed on DISPs in order to protect competition could severely limit the potential value that could be created by data intermediation services, as they would not be able to provide high value analytic services to users [34]. Accordingly, in practice, only a few DISPs have been established with success.

Furthermore, while the regulatory framework for data intermediaries is strict, DGA only targets certain types of such data intermediaries and *'leaves some leeway to avoid the applicability of the regulation'* [34]. Moreover, the DGA definition of the data intermediation services is quite vague as DGA does not provide a precise definition of these services with specified criteria, but rather examples [37].

In order to be classified as a DISP, providers must play an **active role in establishing direct commercial relationships between businesses**. Thus, it is not sufficient to merely provide the technical tools for data sharing without the aim to establish or gather information on commercial relationships between data holders and users. Hence, providers of Application Programming Interfaces (API) for sharing data, cloud storage, web browsing or email services are not DISPs. Also, the DGA does not cover closed-membership services, i.e. those only available to a single data holder or to a pre-selected group of entities.

Further to the Digital Omnibus Regulation Proposal, significant changes are expected to affect the DGA and the overall requirements for data intermediation services. If the proposal is accepted, the DGA will be repealed, and some of its provisions will be incorporated into the Data Act. In addition, the definition of 'data intermediation service' will be revised and the related conditions will become less strict. For example, a legal separation of activities will no longer be required; a functional separation will suffice. The regime will also become voluntary, meaning that adopting the 'data intermediation service' label will be an option for providers that wish to distinguish themselves in the market.

3.3.4. Recommendations for the FL stakeholders under DGA

- **Platform Operator:** If the platform used for FL is used exclusively within a defined group, such as research consortium, it does not qualify as a data intermediation service under the DGA. The closed and contractually defined nature of the consortium falls within the DGA exclusion for services provided within a closed group. In this scenario, the operator is not subject to the DGA obligations applicable to DISPs. If, after the end of the Project, the platform is opened to an indeterminate group of external users, its qualification may change. In that case, the operator could fall within the scope of a data intermediation service. The operator would then need to comply with the DGA requirements, including neutrality, separation of roles, and the prohibition on using intermediated data for its own commercial purposes, unless the planned amendments under the Digital Omnibus alter the legal framework. The commercial plan should therefore clearly define the access model and assess regulatory implications before any expansion beyond the consortium.
- **Data holders:** When data sharing takes place solely within the consortium, data holders do not engage with data users through a regulated data intermediation service under the DGA. Their relationship remains contractual and confined to a closed group. If the platform becomes accessible to external users and qualifies as a data intermediation service, data holders would interact with a regulated intermediary. They should then verify that the operator complies with the DGA framework and reflect this status in their contractual arrangements and risk assessments.
- **Data users:** Within the consortium model, data users access data in a closed, contract-based setting. The DGA regime for data intermediation services does not apply. If the platform evolves into an open model and qualifies as a data intermediation service, data users would obtain data through a regulated

intermediary. They should ensure that the intermediary’s DGA status is properly documented and that contractual terms align with the applicable regulatory obligations.

3.4. European Health Data Space Regulation (EHDS)

3.4.1. Introduction

The European Health Data Space Regulation (**EHDS**) [38] was adopted in 2025 to govern portability and re-use of electronic health data (EHD) in the EU, with two main aims: facilitating natural persons’ access to and control over their personal electronic health data, in the context of healthcare (primary use) and enabling secure and lawful sharing and reuse of health data for research, innovation, and policy-making purposes (secondary use). The EHDS framework on secondary use is of particular relevance to platforms that enable FL, as well as data holders and data users interested in using such platforms during their scientific research activities.

3.4.2. Secondary use of electronic health data under EHDS

Once EHDS is implemented, the secondary use framework will streamline access and re-use of EHD for such purposes as:

- Public interest in the areas of public or occupational health;
- Education or teaching activities in health;
- Scientific research related to health or care sectors that contributes to public health or health technology assessments including development and innovation activities for products or services and training, testing and evaluation of algorithms, including in medical devices, *in vitro* diagnostic medical devices, AI systems and digital health applications;
- Policy making or regulatory activities (Article 53 EHDS).

The schema (adapted from the works of TEHDAS2, for example in D7.1 Guideline on how to use data in a secure processing environment) of the secondary use may be shown as follows:

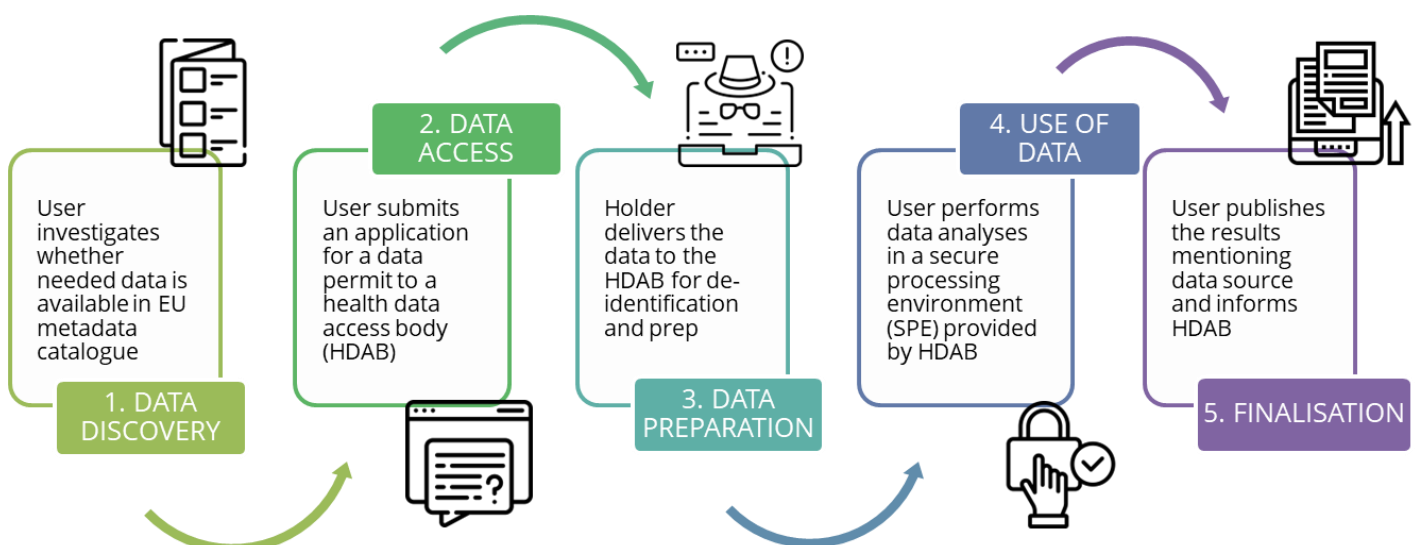


Figure 2: EHDS secondary use framework

The steps shown can be detailed as follows:

1. EHD datasets are discoverable through catalogues of national datasets available for re-use maintained by Health Data Access Bodies (HDAB). These catalogues are interconnected through a central EU-level system, the HealthData@EU infrastructure, to enhance dataset visibility across Member States. HDAB are public authorities designated by Member States to serve as central coordinators for granting access to EHD for secondary purposes. Health data holders must notify their HDAB of the datasets they manage and provide a detailed description of the dataset.
2. Organisations seeking to re-use EHD as health data users – such as research institutions, healthcare innovators or public bodies – can apply for a data permit from their national HDAB. Applications must clearly define the intended purpose, scope, and anticipated outcomes of data processing. Based on evaluation, the HDAB must either approve the application and issue a data permit or reject the application.
3. If the data user is granted a data permit, the HDAB retrieves the requested data from the health data holder and prepares the data received from the holder into an anonymised format, or into a pseudonymised format where the health data user sufficiently demonstrates that the intended purpose of the processing cannot be achieved using anonymised data. Once de-identified, the data is placed in a secure processing environment (SPE).
4. Data user can only access the individual-level health data in the SPE; it cannot be downloaded or copied from that environment.
5. Once the data user finalizes their purpose, they must also inform the HDAB of the results. They must also publish the results (in anonymous format) of secondary use of health data within 18 months of the completion of the data processing that the results have been obtained by using data in the framework of the EHDS.

A more detailed overview of the respective roles of EHDS stakeholders is provided in table below:

Health data holders	<ul style="list-style-type: none"> • provide the HDAB with detailed information about the datasets they hold and their characteristics (Article 60.3) • upon request by a HDAB, make the requested electronic health data available within a reasonable period and, in any event, no later than three months from receipt of the request (Article 60.1 & 2)
Health data users	<ul style="list-style-type: none"> • apply for a data permit from the competent HDAB, by submission of a formal health data access application • may process the data solely for the secondary use purposes specified in the permit and subject to the conditions laid down in Chapter IV of the EHDS. • may submit to HDAB a health data request (Article 69) to obtain anonymised statistical outputs from analysis of the EHD. In this case, the user does not obtain access to the underlying electronic health data. • must publish the anonymous results of their secondary use within 18 months after completing the processing or receiving a response to health data request • reported their results to the HDAB

	<ul style="list-style-type: none"> • must acknowledge in publications that the results were obtained using data accessed under the EHDS framework
HDAB	<ul style="list-style-type: none"> • maintain a national dataset catalogue listing the available datasets which are held by health data holders • inform the public about the conditions under which the data may be accessed • issue data permits (Article 68) and decide on health data requests (Article 69) • after a permit is issued, retrieve the requested EHD from the relevant health data holder (Article 68.7 EHDS), • carry out any necessary pre-processing (e.g., pseudonymization or dataset integration) • provide data users with access to EHD through a secure processing environment (Article 57.1(b) EHDS) • act as enforcement authorities (Article 63) • carry out tasks toward data subjects (Article 58)

Table 4: Responsibilities of health data holders, users and HDAB under EHDSR

3.4.3. Existing mechanisms for sharing health data under EHDS

EHDS explicitly clarifies that it is not intended to replace existing data-sharing solutions. Article 1(8) provides that the regulation shall not affect access to electronic health data for secondary use agreed in the framework of contractual or administrative arrangements between public or private entities. Similarly, recital 52 emphasizes that *'Without hindering or replacing contractual arrangements or other mechanisms in place, this Regulation is aimed at establishing a common mechanism to access electronic health data for secondary use across the Union.'*

This suggests that the EHDS does not seek to completely replace existing governance models or national mechanisms. Moreover, it creates new categories of data governance bodies, beyond those mentioned in Section 3.4.2 above. The table below provides a brief description of those bodies.

Trusted health data holders (Article 72) designated by Member States
<ul style="list-style-type: none"> • Health data holders may apply to obtain a status of a trusted health data holder, if they demonstrate: (i) ability to provide access to health data through a SPE, (ii) necessary expertise to assess health data applications and health data requests; (iii) necessary guarantees to ensure compliance with the EHDS. The specific procedures for designation as a trusted health data holder will be specified in national legislation of the Member States. • In case the application or request refer to the dataset held by the trusted health data holder (THDH), HDAB may forward data request to a trusted health data holder (THDH). THDH are required to assess the health data access applications submitted under this simplified procedure, based on their expertise in dealing with the type of health data they are processing, and issue a recommendation regarding a data permit. HDAB will however make the final decision on the application.

<ul style="list-style-type: none"> • HDAB issues data permit, THDH prepares and makes the data available to the user via a secured processing environment.
Authorized participants in HealthData@EU (Article 75.4)
<ul style="list-style-type: none"> • Health-related research infrastructures or similar infrastructures whose functioning is based on EU law and which provide support for the use of electronic health data for research, policy making, statistical, patient safety or regulatory purposes may become authorised participants in HealthData@EU and connect to it.
Health data intermediation entities (Article 50.3)
<ul style="list-style-type: none"> • Member states may provide in their national law that the duties of certain categories of health data holders are to be fulfilled by health data intermediation entities.
Cross-border registries or databases from a number of Member States (Article 76)
<ul style="list-style-type: none"> • They may designate a coordinator to ensure the provision of data from the registries' or databases' network for secondary use.

Table 5: Brief description of EHDS bodies

3.4.4. Role of secure processing environments (SPE)

EHDS permits health data users to access electronic health data through secure processing environments. This regulation refers to the definition of SPE contained in the Data Governance Act and further elaborates on the requirements applicable to such environments for the purposes of secondary use of electronic health data (Article 73(1) EHDS). A SPE constitutes a physical or virtual environment combined with organisational arrangements aimed at ensuring compliance with the European Union and national law, including the GDPR, notably in relation to data subjects' rights, intellectual property, and commercial and statistical confidentiality, as well as data integrity and accessibility. It also allows the entity operating the SPE to determine and oversee all data processing operations, including data display, storage, download, export, and the production of derivative data through computational means.

Health Data Access Bodies may provide access to electronic health data pursuant to a data permit only via an SPE. These environments must implement state of the art technical and organisational measures, including restricting access to authorised natural persons identified in the data permit, reducing the risk of unauthorised access or alteration of data, limiting data handling operations to a small number of identifiable authorised individuals, and ensuring that users can access only the data covered by their permit through secure and individual access credentials. SPEs must also maintain identifiable logs of access and activities for audit and verification purposes, retained for at least one year, and must ensure ongoing compliance with security requirements and monitoring of potential risks.

The European Commission will further specify the technical, security, and interoperability requirements for SPEs through implementing acts. ENISA [18] has highlighted that challenges remain in defining and operationalising appropriate technical and organisational measures, particularly in view of novel processing activities and the lack of clarity surrounding roles and responsibilities.

3.4.5. Recommendations for the FL stakeholders under EHDS

Under the EHDS, federated learning can fit the secondary-use regime because training and evaluating algorithms for healthcare is an allowed purpose, access to individual-level data requires a data permit, and

the default logic is that data are accessed in pseudonymised form only where anonymised data do not suffice. If FL platform is organised as a secure processing environment, or as a federation of trusted SPEs, it must still satisfy Article 73 requirements, such as: access limited to authorised users, state-of-the-art measures against unauthorised reading, copying, modification or removal, logging and monitoring, and a strict bar on downloading personal data, with only reviewed non-personal outputs allowed to leave the environment. That is precisely why FL is attractive here: raw data remain local and only model updates or weights are exchanged, but the compliance focus then shifts to whether those intermediate outputs, or the final model, can still be regarded as non-personal rather than containing or reproducing underlying health data. TEHDAS2's draft SPE specifications address federated learning expressly on that basis, noting both that raw data are not exchanged and that, because the anonymity of interim outputs is difficult to assess, federated learning should run within a network of trusted, EHDS-compliant SPEs. For SPEs the EHDS points to a future Commission implementing act, as under Article 73(5), by 26 March 2027 the Commission must lay down the detailed technical, organisational, information-security, confidentiality, data-protection and interoperability requirements for SPEs, including the tools available to the health data user.

4. Recommendations on appropriate privacy-preserving techniques and processes for FL implementations in healthcare

4.1. Introduction

Data holders, such as hospitals or data hubs, typically already have a strong protection of the sensitive data which is stored locally. These data protection measures are often already approved for the purposes of treating patients and performing research locally on the data holder's premises. Several steps in a FL process can benefit from these existing local protection and approval processes.

A typical FL scenario consists of multiple steps:

- Collecting and preparing data: Here, as in the FLUTE approach, typically every data holder locally prepares data (anonymization, normalization, minimization ...) according to their own local procedures and ethical committee & DPO approved processes.
- The training of the FL model, which will be discussed in depth below.
- Sharing of the model. Models can only be shared to the extent they do not contain sensitive data which can be extracted. A common approach is to use differential privacy, i.e., to guarantee that not more than a certain amount of sensitive information (considered negligible) can be extracted from the shared model.
- Possibly fine-tuning of the model towards a specific data holder (with its specific patient population and environment). This process can be performed locally, so the data protection risks are more limited. The results for more details on fine-tuning were reported under WP9.
- Applying the model to make predictions for individual new patients. This too can be performed locally at the location where the patient is treated. As the data does not need to leave that location, the local data holder protection is sufficient.

4.2. Modelling security and privacy requirements and algorithm guarantees

When multiple Data holders collaborate for performing FL, it is important to agree on the parameters of this collaboration. Next to a legal, e.g., joint controller, agreement specifying the legal responsibilities of

the involved parties, a technical agreement, e.g., as an annex, can specify the several technical parameters of the collaboration. These can include agreements on whether the parties are trusted to perform operations correctly (or else a verification mechanism such as a TEE is desired), what level of statistical privacy will be required on the output (e.g., the shared models), or what level of security is required (e.g., should the system be robust against colluding participants, against adversaries with unlimited computational power etc.).

FLUTE proposes a technical implementation for this idea: first deliverable D4.1 describes the concept of a “study agreement” which is a data structure shared between the several machines of the participating stakeholders to agree on these technical parameters. Then, deliverable D5.2 describes in more detail an ontology to represent these parameters and security requirements. Next to specifying security requirements, the FLUTE system also allows one to annotate algorithms with security/privacy guarantees. This can then be used to check whether queries and algorithm asked by the researcher comply with the requirements.

While the modelling of security and privacy requirements provides a structured way to specify the guarantees expected from a federated learning collaboration, these requirements must ultimately be enforced through concrete technical mechanisms. privacy-enhancing technologies (PETs) play a central role in this regard, as they enable the implementation of verifiable security and privacy guarantees at different stages of the FL pipeline. Depending on the threat model and trust assumptions defined in the study agreement, different PETs can be employed to protect data, computations, and communications.

4.3. Trusted Execution Environments (TEEs)

Deploying FL in healthcare environments requires privacy guarantees that extend beyond data locality and model-level protections. While FL reduces the need for centralized data aggregation, sensitive information may still be exposed during training, aggregation, or coordination phases. Therefore, robust privacy-preserving mechanisms must be integrated across the entire FL lifecycle to ensure compliance with regulatory frameworks and to maintain institutional trust.

In this context, trusted execution environments (TEEs) have emerged as a particularly effective technical safeguard. TEEs provide hardware-enforced isolated execution environments in which sensitive computations (such as local model training, gradient computation, or secure aggregation) can be executed without exposure to the host operating system, hypervisor, or other potentially compromised system components. By ensuring that raw patient data and intermediate model updates remain protected at the hardware level, TEEs enable healthcare institutions to participate in collaborative FL workflows while retaining strict control over locally stored data (Costan and Devadas [39]).

A critical component of TEE-based FL deployments is the implementation of a robust remote attestation process. Remote attestation allows each participant in the federated network to cryptographically verify that other nodes are running trusted code within genuine TEEs before exchanging any model updates. This verification step is essential in healthcare collaborations, where institutions must trust not only their own execution environment but also the integrity and security posture of external partners. Without reliable attestation, the risk of malicious or misconfigured nodes undermining the FL process remains significant (Schuster et al [40]; Knauth et al. [41]).

From a system-design perspective, TEEs are particularly well suited to complement other privacy-enhancing technologies (PETs) commonly used in FL. Unlike techniques such as homomorphic encryption or secure multi-party computation, which may introduce substantial computational and communication

overhead, TEEs rely on hardware-based isolation and therefore enable efficient execution of complex machine-learning workloads. As a result, they can be deployed in real-world clinical environments without significantly degrading training performance or scalability (Sabt et al. [42]).

However, the effectiveness of TEEs depends not only on technical implementation but also on sound organizational and governance practices. Secure key management, strict access-control policies, regular security audits, and clear delineation of roles and responsibilities are necessary to prevent operational vulnerabilities. Furthermore, TEE-based FL deployments must be aligned with applicable legal and ethical frameworks, such as GDPR, ensuring that hardware-level protections are complemented by transparent data-governance policies and accountability mechanisms (Brauneck [1]).

In summary, the integration of trusted execution environments, supported by remote attestation and strong organizational processes, constitutes a recommended best practice for privacy-preserving FL in healthcare. When combined with complementary PETs and appropriate governance structures, TEEs can significantly enhance trust, security, and regulatory compliance in distributed medical AI systems.

4.4. Adaptive encryption techniques

TEEs described above have however some limitations. They rely on a private key of a hardware manufacturer. In that sense, corruption of a hardware manufacturer could have severe impact on all users of the hardware. Using TEEs to attest correct computation provides assurance to other participants, provided that the TEE administrator is trusted not to manipulate the environment. Even in case of corruption of the hardware manufacturer, all attestations provided before that point can be considered truthful. The risk for using TEEs for protecting privacy in the collaboration between data holders, e.g., when transferring encrypted information between data holders, is higher. Any information encrypted only using a TEE may be stored and revealed after the corruption of a TEE hardware manufacturer. Therefore, it is recommended to use for the collaboration between multiple parties cryptographic techniques such as multi-party computation (MPC) which cannot be affected by such attacks, e.g., secret sharing can be information-theoretically secure, so data holders are guaranteed that without their consent no operation can be performed on anything that depends on (unrevealed parts of) their data.

While MPC can ensure that no intermediate results are revealed, its costs strongly depends on the security requirements, e.g., if one requires the computation that is robust against all but one parties colluding with each other, then the cost per party is likely linear in the number of parties. In contrast if one only requires that the computations are secure against a fraction of the parties colluding (e.g., the honest majority setting), then costs per party may be only logarithmically in the number of parties, while if one assumes all parties are honest and do not collude, the cost per party may be independent of the number of parties. To make the most appropriate choice, it is important to have made explicit the security requirements.

In conclusion, the use of MPC (encryption based on secret sharing, homomorphic encryption or a similar technique) is recommended to ensure no intermediate results can be revealed when exchanging information between parties. Having made explicit the security requirements allow one to choose the most appropriate technique which both can guarantee the security requirement and can keep computation and communication costs limited.

4.5. Statistical privacy

To be useful, every algorithm eventually has an output, e.g., a statistic which has been computed or a prediction or recommendation for an individual patient. It is important to avoid that from this output one can infer sensitive information (e.g., can identify an individual or can infer attributes of an individual). This

applies even if all intermediate results of a computation have been hidden by appropriate use of software or hardware based encryption techniques.

To make the output privacy-preserving, one can apply statistical privacy. The most popular form is differential privacy (DP). The effect is that the model output is (preferably only slightly) perturbed so that exact inference of sensitive information becomes impossible. When the confidence in identifying an individual becomes negligibly small, the output is considered privacy preserving.

In the context of application in medicine, where output of models may affect patients, it is ethical to maximize the precision of such outputs. Applying the well-known, generic formulas for differential privacy may guarantee privacy but may also make the output less precise than what is achievable in a privacy-preserving way.

We therefore recommend to perform a case-by-case analysis to determine the actual amount of noise needed to guarantee privacy. Factors that can be taken into account are (a) the amount of noise already present in the data, e.g., measurement noise, (b) the security requirements on which the parties agreed, and (c) a refined analysis of the sensitivity of the different data used.

4.6. Role of data holders and data hubs

Healthcare institutions participating as data holders play a central role in privacy-preserving FL deployments. Within the project framework of the FLUTE project, data hub partners such as VHIR, IRST, CHUL, and SR are responsible for maintaining secure local environments where sensitive clinical and imaging data remain stored and processed in accordance with institutional, national, and European data-protection regulations.

A fundamental principle of federated learning is that patient-level data never leave the originating institution. Instead, model training is performed locally within each data holders' secure infrastructure (node), and only model parameters or aggregated updates are shared with the federated aggregation service. This architecture significantly reduces the risks associated with centralized data storage while enabling collaborative model development across multiple institutions. Data holders therefore act as trusted custodians of the original datasets, ensuring that all local data processing complies with applicable legal frameworks such as the GDPR as well as with institutional ethical approvals.

To support secure participation in the federated network, each data hub or data holder must implement appropriate technical and organizational safeguards. These include secure computing environments capable of hosting the FL node, controlled access to clinical and imaging datasets, encryption of communications during model-update exchanges, and strict authentication mechanisms for participating nodes. In addition, local governance procedures and their translation into the requirements stated in the study agreement must ensure that data use remains consistent with the scope of informed consent and ethical approvals under which the data were originally collected.

Data holders also contribute to the quality, interoperability, and reliability of the federated training process. This includes the curation and harmonization of clinical and imaging datasets, the implementation of standardized preprocessing workflows where applicable, and the validation of locally generated model updates before they are shared with the federated system. Such procedures help ensure that collaborative models developed through FL are robust, reproducible, and clinically meaningful across different healthcare settings.

Still, despite minimizing data sharing challenges, FL in healthcare involves the secondary use of sensitive health data. This secondary use of data raises challenges about governance, ethical concerns, privacy, hospital cybersecurity and data quality [43].

4.6.1. Data governance

Governance agreements (such as data sharing agreements or study agreements discussed above) should be established with external FL stakeholders before data is accessed in federated studies to ensure data protection and intended use purposes are upheld. In FL, contractual agreements are complex due to the sheer number of parties involved and the need for any amendments to be synchronized. Clear and formalized policies, procedures, and standards need to be developed and rolled out. Governance agreements should outline the explicit roles of the parties (coordinating server operator, infrastructure, model developer, data holder, model validator), define responsibilities for incident response, model misuse, or data breach and the procedures for onboarding and offboarding participating sites. Governance agreements should also ensure a fair participation of every data node in the network and a mitigation plan to address such imbalance. In a project such as FLUTE, where the parties to the collaboration are clearly defined, governance arrangements were established through a joint controller agreement between the consortium partners. By contrast, where the platform operates as a marketplace, and depending on the specific setup of the FL platform, governance may rely on a combination of agreements, including separate agreements with the platform operator and study-specific agreements concluded with other research participants, such as data holders or researchers, for each individual study.

4.6.2. Ethical concerns

Health research projects often use retrospective data (i.e. data obtained for other purposes or studies), for the training of AI systems. Even though retrospective studies use existing data, they are often still considered human-subject research and may require ethics review. Institutional Review Boards / Independent Ethics Committees (IRB/IEC) evaluate, in the local regulatory landscape, whether data are truly existing (no prospective collection allowed for “retrospective” status), whether data minimization principles are applied, whether waivers of informed consent and GDPR obligations are justified. Federated learning introduces new ethical, governance, security, and transparency challenges compared with traditional centralized clinical research. IRBs can evaluate such projects more effectively through a clear data governance plan describing:

- data privacy controls, monitoring processes, formalized data-sharing or non-sharing agreements, transparent evaluation steps.
- stakeholder engagement and clear communication channels.
- designated data stewards, oversight bodies, and clearly assigned roles across institutions.

4.6.3. Privacy

Federated learning does not automatically guarantee privacy. Data holders should check the privacy tools available in the FL platform and implemented in the study, such as differential privacy, secure multiparty computation, encryption, or hybrid approaches. The assessment of the actual privacy guarantees may be out of the expertise of data holders. Therefore, data holders should seek access to certification labels, penetration testing reports or threat-model analysis.

4.6.4. Hospital cybersecurity

Healthcare systems are high-value cyberattack targets, and security weaknesses in servers, networks, or endpoints can compromise the entire data ecosystem. Participation in FL network mechanically increases the surface of attacks. Data holders must protect systems from external attackers and malicious insiders exploiting vulnerabilities for hacking, theft, espionage. Cybersecurity measures such as segmentation, 2-factor authentication, encryption, monitoring, threat detection applies to FL nodes as to other clinical IT systems, with a special focus point on denial-of-service attacks which may compromise a data holder node, hence disrupting model updates.

4.6.5. Data quality

In traditional centralized ML, a central data science team receives all datasets and performs: data cleaning, formatting and normalization, feature construction, missing-value handling, quality checks. FL shifts these responsibilities onto data holders. This shift is not a side-effect but a structural property of how FL works.

Therefore, FL also places a greater onus on capability development within the nodes to ensure the data holders have the skills necessary to locally preprocess data. Data holders must not only clean their data — they must do it in a way that is consistent with other nodes, even though they cannot see each other's data. On an organizational level, this requires trust, cooperation and governance. On the technical level, this requires knowledge in international data format such as FHIR and standard ontologies such as SNOMED-CT.

5. Conclusions

Federated learning offers a strong model for enabling collaborative health research while reducing the need to centralise sensitive clinical data. Its main legal and technical advantage is that raw patient-level data can remain within the infrastructure of the data holder, while only model updates, parameters, or aggregated outputs are exchanged. This design can support the GDPR principles of data minimisation, purpose limitation, integrity and confidentiality. It may also align well with the EHDS model of secure processing environments for secondary use of electronic health data.

However, FL should not be treated as automatically compliant with the GDPR. The fact that raw data remain local does not remove the need to assess whether personal data are processed. In particular, local datasets, model parameters, trained models, and outputs may still qualify as personal data if individuals can reasonably be identified or inferred. This assessment must be performed case by case and from the perspective of each relevant actor.

A compliant federated learning deployment therefore requires a clear allocation of roles and responsibilities. Data holders, data users and platform operators must determine whether they act as controllers, joint controllers, processors or recipients. These roles should not be assigned only by contractual labels. They must reflect factual influence over the purposes and means of processing, access to identifiable information, and control over the federated infrastructure.

The legal basis for processing health data must be identified and documented before any federated study begins. In healthcare and research settings, this requires both an Article 6 GDPR legal basis and an Article 9(2) GDPR condition. Transparency obligations must also be addressed. Data subjects should be informed about the relevant processing operations, the categories of recipients, including the possible use of their data for AI training through FL, unless a valid exemption applies under EU or national law.

Accountability is central. Federated learning collaborations should be supported by data sharing agreements, joint controller arrangements or data processing agreements, depending on the role allocation. These agreements should be complemented by technical annexes or, if the platform allows, case specific study agreements that define security requirements, privacy guarantees or permitted processing operations. The platform should also make clear, through user manuals and terms of use, the foreseen model update flows, data retention rules, logging and audit rights, onboarding and offboarding procedures, and incident response mechanisms.

DPIAs will often be required in federated learning deployments involving health data. They should cover not only local datasets, but also the federated architecture itself. Relevant risks include inference attacks, reconstruction attacks, membership inference, model poisoning, data poisoning, unauthorised access to aggregation servers, weak node security, leakage through model outputs, and international access to personal data. Where several controllers participate in the same federated network, a coordinated DPIA approach is recommended, while each controller remains accountable for its own assessment.

Federated learning used in development of or as a medical device may fall within both the MDR and the AI Act. If the AI system qualifies as a high-risk AI system, the provider must demonstrate compliance with requirements on data governance, technical documentation, accuracy, robustness, cybersecurity, transparency, human oversight, logging, post-market monitoring, and quality management, among others. These obligations should be integrated with the MDR compliance framework rather than managed in isolation.

The DGA and EHDS may affect federated learning platforms depending on their access model and intended use. A platform intended to be used by a closed research consortium will usually raise different regulatory issues from one operated as an open data marketplace. If a federated platform is made available to an indeterminate group of external data holders and data users, its status under the DGA, and potentially under future amended data legislation, should be reassessed. Under the EHDS, federated learning can be a useful mechanism for secondary use and deployed in secure processing environments (SPEs), but only if implemented with sufficient safeguards, secure processing controls, and output review procedures.

In parallel, the technical safeguards supporting the FL platform security must be selected according to the threat model and the trust assumptions of the collaboration. Trusted execution environments, multi-party computation, homomorphic encryption, secret sharing, secure aggregation and differential privacy may each play a role. No single technique is sufficient in all cases. The appropriate combination depends on the sensitivity of the data, the number and trustworthiness of participants, the risk of collusion, the need for model utility, and the legal requirements applicable to the use case.

Data holders and data hubs remain central to compliant federated learning. They must maintain secure local environments (nodes), apply appropriate data preparation and pseudonymisation measures, ensure consistency with ethical approvals and consent or other legal bases, and support data quality and interoperability, and validate study agreement conditions before participation in federated training. They also need sufficient technical capacity to perform local preprocessing in a consistent manner, using relevant standards and ontologies where appropriate.

The project results indicate further need for practical and authoritative guidance on GDPR-compliant federated learning. Potentially, such authority as ENISA, in cooperation with data protection authorities and other relevant EU bodies, would be well placed to develop a structured framework defining minimum legal, organisational and technical requirements for federated learning systems. Such guidance would

support legal certainty, reduce reliance on ad hoc legal opinions, and help organisations deploy federated learning in a trustworthy and scalable manner.

Overall, the safe deployment of federated learning in healthcare requires more than distributed architecture. It requires an integrated compliance framework in which legal analysis, governance arrangements, privacy-enhancing technologies, cybersecurity measures, documentation, and operational controls are designed together from the outset.

6. References

1. Brauneck, A., et al., *Federated machine learning, privacy-enhancing technologies, and data protection laws in medical research: scoping review*. Journal of medical Internet research, 2023. **25**: p. e41588.
2. Dayan, I., et al., *Federated learning for predicting clinical outcomes in patients with COVID-19*. Nat Med, 2021. **27**(10): p. 1735–1743.
3. Rieke, N., et al., *The future of digital health with federated learning*. NPJ Digit Med, 2020. **3**: p. 119.
4. Kairouz, P. and H.B. McMahan, *Advances and open problems in federated learning*. Foundations and trends in machine learning, 2021. **14**(1-2): p. 1–210.
5. Xu, J., et al., *Federated learning for healthcare informatics*. Journal of healthcare informatics research, 2021. **5**(1): p. 1–19.
6. Bak, M., et al., *Federated learning is not a cure-all for data ethics*. Nature Machine Intelligence, 2024. **6**(4): p. 370–372.
7. Supervisor, E.D.P., in *Federated Learning*. https://www.edps.europa.eu/press-publications/publications/techsonar/federated-learning_en.
8. OECD, *Sharing Trustworthy AI Models With Privacy-Enhancing Technologies*, in *OECD Artificial Intelligence Papers, no. 38 (2025)*. 2025: https://www.oecd.org/en/publications/sharing-trustworthy-ai-models-with-privacy-enhancing-technologies_a266160b-en.html
9. Kogut-Czarkowska, M. and M. Shabani, *Federated networks and secondary uses of health data: Challenges in ensuring appropriate safeguards for sharing health data under the GDPR and EHDS*, in *The European Health Data Space*. 2025, Taylor & Francis.
10. Casaletto, J., et al., *Federated Analysis for Privacy-Preserving Data Sharing: A Technical and Legal Primer*. Annu Rev Genomics Hum Genet, 2023. **24**: p. 347–368.
11. Scheibner, J., et al., *Revolutionizing Medical Data Sharing Using Advanced Privacy-Enhancing Technologies: Technical, Legal, and Ethical Synthesis*. J Med Internet Res, 2021. **23**(2): p. e25120.
12. Darzidehkalani, E., M. Ghasemi-Rad, and P.M.A. van Ooijen, *Federated Learning in Medical Imaging: Part II: Methods, Challenges, and Considerations*. J Am Coll Radiol, 2022. **19**(8): p. 975–982.
13. EDPB, *Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models*. 2024: https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en.
14. Becker, R., D. Chokoshvili, and E.S. Dove, *Legal bases for effective secondary use of health and genetic data in the EU: time for new legislative solutions to better harmonize data for cross-border sharing?* International Data Privacy Law, 2024. **14**(3): p. 223–246.
15. Gedeberg, R., et al., *Federated analyses of multiple data sources in drug safety studies*. Pharmacoepidemiology and Drug Safety, 2023. **32**(3): p. 279–286.
16. Bernier, A., et al., *Reconciling the biomedical data commons and the GDPR: three lessons from the EUCAN ELSI collaboratory*. European Journal of Human Genetics, 2024. **32**(1): p. 69–76.
17. Bradshaw, A., et al., *Data sharing in neurodegenerative disease research: challenges and learnings from the innovative medicines initiative public-private partnership model*. Frontiers in neurology, 2023. **14**: p. 1187095.
18. ENISA, *Engineering Personal Data Protection in EU Data Spaces Final report 2024*: <https://www.enisa.europa.eu/publications/engineering-personal-data-protection-in-eu-data-spaces>.
19. AEPD, *Approach To Data Spaces From GDPR Perspective*,. <https://www.aepd.es/documento/approach-to-data-spaces-from-gdpr-perspective.pdf>.
20. Hallock, H., et al., *Federated Networks for Distributed Analysis of Health Data*. Front Public Health, 2021. **9**: p. 712569.
21. Thorogood, A., et al., *International federation of genomic medicine databases using GA4GH standards*. Cell Genomics, 2021. **1**(2).
22. Rossello, S., *Data protection by design in AI? The case of federated learning*. The case of federated learning (May 30, 2021). S. Rossello, R. Díaz Morales, L. Muñoz-González, 'Data Protection by design in AI, 2021: p. 1–11.
23. EDPB, *Fundamentals of Secure AI Systems with Personal Data*. https://www.edpb.europa.eu/our-work-tools/our-documents/support-pool-experts-projects/fundamentals-secure-ai-systems-personal_en.

24. Theresa Stadler, C.T., Melanie Kolbe-Guyot, *Purpose First: The Need for a Paradigm Shift in Privacy-Preserving*, in *C4DT Insight Paper 2025*: <https://drive.switch.ch/index.php/s/nrMAQWPTTImxMGp>.
25. Soltani, B., et al., *A survey of federated evaluation in federated learning*. arXiv preprint arXiv:2305.08070, 2023.
26. Rambla, J., Beltran, S., & D'Altri, T. , *European Genomic Data Infrastructure project (GDI) D8.4 Report on federated data access scenarios*. 2023: Zenodo. <https://doi.org/10.5281/zenodo.8208439>.
27. EDPB, *Guidelines 05/2021 on the Interplay between the application of Art. 3 and the provisions on international transfers as per Chapter V of the GDPR*. 2023.
28. Alvarellos, M., et al., *Democratizing clinical-genomic data: How federated platforms can promote benefits sharing in genomics*. *Frontiers in Genetics*, 2023. **13**: p. 1045450.
29. Townend, D., *Conclusion: harmonisation in genomic and health data sharing for research: an impossible dream?* *Human genetics*, 2018. **137**(8): p. 657–664.
30. *Federated Data Systems: Balancing Innovation and Trust in the Use of Sensitive Data*. 2019, World Economic Forum: http://www3.weforum.org/docs/WEF_Federated_Data_Systems_2019.pdf
31. Narmadha, K. and P. Varalakshmi, *Federated Learning in Healthcare: A Privacy Preserving Approach*. *Stud Health Technol Inform*, 2022. **294**: p. 194–198.
32. Mandl, K.D., et al., *The Genomics Research and Innovation Network: creating an interoperable, federated, genomics learning system*. *Genet Med*, 2020. **22**(2): p. 371–380.
33. *Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act)*, E. Union, Editor.: OJ L 152, 3.6.2022.
34. Von Ditzfurth, L. and G. Lienemann, *The Data Governance Act:—Promoting or Restricting Data Intermediaries?* *Competition and Regulation in Network Industries*, 2022. **23**(4): p. 270–295.
35. Koutroumpis, P., A. Leiponen, and L.D. Thomas, *Markets for data*. *Industrial and Corporate Change*, 2020. **29**(3): p. 645–660.
36. Ruohonen, J. and S. Mickelsson, *Reflections on the data governance act*. *Digital Society*, 2023. **2**(1): p. 10.
37. Baloup, J., et al., *White paper on the data governance act*. 2021.
38. *Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space and amending Directive 2011/24/EU and Regulation (EU) 2024/2847*, E. Commission, Editor. 2025: OJ L, 2025/327.
39. Costan, V. and S. Devadas, *Intel SGX explained*. *Cryptology ePrint Archive*, 2016.
40. Schuster, F., et al. *VC3: Trustworthy data analytics in the cloud using SGX*. in *2015 IEEE symposium on security and privacy*. 2015. IEEE.
41. Knauth, T., et al., *Integrating remote attestation with transport layer security*. arXiv preprint arXiv:1801.05863, 2018.
42. Sabt, M., M. Achemlal, and A. Bouabdallah. *Trusted execution environment: What it is, and what it is not*. in *2015 IEEE Trustcom/BigDataSE/ISPA*. 2015. IEEE.
43. Eden R, Chukwudi I, Bain C et al; *A scoping review of the governance of federated learning in healthcare*; *NPJ Medicine* (2025)
44. EDPB, *Guidelines 1/2026 on processing of personal data for scientific research purposes*, https://www.edpb.europa.eu/our-work-tools/documents/public-consultations/2026/guidelines-12026-processing-personal-data_en
45. Kaye, J., Whitley, E., Lund, D. et al. *Dynamic consent: a patient interface for twenty-first century research networks*. *Eur J Hum Genet* **23**, 141–146 (2015).
46. Kogut-Czarkowska M, Shabani M. *Anonymization, accountability, and access : legal dimensions of health data sharing in federated networks : perspectives from empirical study*. *Frontiers In Digital Health*. 2026;8

Annex 1

Table 1: Data Protection Requirements in FL in Healthcare

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
1.	<p>Legal basis</p> <p><i>The legal bases for the processing of personal data are set out in Article 6 of the GDPR. Pursuant to that provision, consent constitutes one of the lawful bases for processing. Another legal basis may be, for example, compliance with a legal obligation to which the controller is subject.</i></p> <p><i>In order to process special categories of personal data (such as health data) additionally one of the exceptions provided for in Article 9(2) GDPR must be met.</i></p>	<p>Data Holders as controllers should indicate which legal basis and, in case of special categories of data, which additional exception they will rely on for personal data processing.</p> <p>If this will be consent, it must be voluntary, free, informed and specific (i.e. list the purpose for which the data will be used).</p> <p>*In its recent guidelines, EDPB confirmed that broad consent for future research projects within certain areas of scientific research may be used under specific safeguards¹².</p> <p>Data Hubs if they are not controllers, they need to</p>	<p>The Platform Operator should ensure that the platform enables Data Holders and Data Users to:</p> <ul style="list-style-type: none"> record the legal basis for each processing activity, document and audit all processing operations, including data retention and deletion, and regularly update this documentation to reflect the categories of data and the processing activities carried out. <p>The Platform should enable Data Holders to publish and the Data Users to view the dataset metadata, including the legal limitations for its use and basis for personal data processing (in case of</p>	<p>If Data Users process personal data as controllers or joint controllers, they need to demonstrate legal basis for use of the health data for their research through FL.</p> <p>*If the legal basis is consent, potentially Data Users could also rely on the same broad consent as obtained by the Data Holders¹³</p> <p>If it can be demonstrated that the data is not personal from the perspective of the Data Users, legal basis is not needed.</p>

¹² See Guidelines 1/2026 on processing of personal data for scientific research purposes

¹³ *Ibidem*.

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
		ensure that the controller (provider of the data to the hub) can demonstrate legal basis.	consent based processing, the remits of the consent, including template consent form, should be available to the Data Users). The Platform Operator when acting in a controller or join controller role and processing personal data, needs to demonstrate legal basis for use of the health data for their research through FL.	
2.	<p>Transparency obligations</p> <p><i>Data subjects must be informed about the processing of their personal data.</i></p> <p><i>Respect their rights and make sure they can exercise them effectively unless special derogations exist under national laws.</i></p>	<p>Data Holders should provide the data subjects (research participants) with necessary information (as per Article 13 and 14 of the GDPR), unless they can demonstrate that an exemption applies.</p> <p>The information should include information about data recipients (such as Platform Operator and Data Users).</p> <p>Data Hubs if they are not controllers, they need to ensure that the controller</p>	<p>Platform Operator should ensure that the data subjects have been informed about the processing of their personal data, unless they can demonstrate that an exemption applies.</p>	<p>Data Users should ensure that the data subjects have been informed about the processing of their personal data, unless they can demonstrate that an exemption applies.</p> <p>*Further to SRB case (see Section 2.4 above), even if the Data Users cannot identify the data subjects, the Data Holders need to still fulfil their obligation to inform the data subjects about the recipients of personal data.</p>

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
		(provider of the data to the hub) can demonstrate fulfilment of transparency obligations.		
3.	<p>Purpose limitation and data minimization</p> <p><i>Personal data must be collected for specific, explicit purposes and not used later in ways that are incompatible with those purposes.</i></p> <p><i>Only the data that is necessary and proportionate for achieving those purposes should be collected and processed.</i></p>	<p>Data Holders should:</p> <ul style="list-style-type: none"> • minimise the collection and use of personal data, bearing in mind the purpose of the data sharing • provide and agree upon the anonymization or pseudonymization measures to be followed • make sure that the implemented measures are documented for reference purposes. <p>The teams which pseudonymize the data should not be the same which then use the data for research purposes (factual separation).</p> <p>Ultimately, out of caution, FL platforms may need to treat</p>	<p>Platform Operator should ensure that there is no possibility for the Data Users to copy or extract any personal data from the Platform. Air locks (previous checks) can be considered for any data to be downloaded. Measures need to be put in place so that risk that personal data is reconstructed using the gradients or weights, or the local/central models is mitigated.</p>	<p>Data Users should design AI models to operate on minimal personal data, reducing unnecessary collection.</p> <p>Decisions regarding which data will be used should be documented, only minimum data categories should be used.</p>

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
		all datasets as sensitive (personal data).		
4.	<p>Data accuracy</p> <p><i>Personal data must be accurate and, where necessary, kept up to date. Inaccurate data must be erased or corrected without delay, taking into account the purposes for which the data is processed.</i></p>	<p>Data Holders should:</p> <ul style="list-style-type: none"> • verify data, allow corrections, and periodically review whether stored data remains correct and relevant; • implement distributed data quality management procedures; • follow instructions and guidelines on the data upload and promote interoperability. 	<p>Platform Operator:</p> <ul style="list-style-type: none"> • should issue guidelines and/or instructions for the Data Holders on data management and accuracy; • may provide tools that require data checks before the data is made available in the nodes; • may facilitate interoperability and the sharing of personal data, assisting Data Holders with anonymising or pseudonymising personal data. 	<p>Data Users should inspect the dataset metadata which is made available by the Data Holders and document which datasets were accessed during FL. If available, they should verify the accuracy of data and suitability for their research through the recommended tools.</p>
5.	<p>Storage limitation</p> <p><i>Storage periods for the data must be determined.</i></p> <p><i>No personal data should be available beyond the data retention period – any stored personal data must be anonymized or deleted</i></p>	<p>Data Holders should:</p> <ul style="list-style-type: none"> • Define storage periods and ensure they are maintained. • Storage period must be communicated to data subjects, when collecting data from them. 	<p>Platform Operator should set automated retention periods to delete or anonymize data used by after its purpose is fulfilled or Data Holders so instructs.</p>	<p>Data Users need to be aware of the storage limitations and take this into account when planning research and documenting data use.</p>

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
	<i>following the data deletion policy.</i>			
6.	<p>Confidentiality and integrity (data security)</p> <p><i>Personal data must be protected from unlawful access, deletion and modification by implementation of adequate security measures. Determination should take into account: state of the art, implementation costs, the nature, scope, context and purposes of processing, and the type and volume of data involved. Measures are appropriate if they demonstrably reduce identified risks to an acceptable level and are regularly tested, reviewed and updated.</i></p>	<p>Data Holders is responsible for implementing appropriate safeguards to protect personal data in their node.</p> <p>It is recommended to use encryption on the data at rest on the nodes in order to mitigate attacks that could directly compromise those nodes.</p> <p>They must ensure protections against data poisoning (where false data is injected into the training process) including mechanisms for anomaly detection and secure update verification.</p>	<p>Platform Operator should issue guidelines for the Node Holders on the minimum requirements to establish a node.</p> <p>Platform Operator must establish identity and access management policy (with authentication, authorisation and logging processes).</p> <p>Platform Operator should include implementation of appropriate PETs, as robust security measures to protect the personal data from unauthorised access or modifications. Measures that can be available on a FL platform include: Secure aggregation and SMPC, Trusted Execution Environments (TEE), Differential privacy, ensuring protection of data not only at rest and in transit, but also</p>	<p>Data Users should verify the security of the model, including assessing potential privacy risks such as information leakage through model outputs, susceptibility to inference attacks, and robustness against adversarial manipulation.</p> <p>Where applicable, they should ensure that models are trained and accessed within secure environments (e.g., TEEs) or using PET-supported mechanisms provided by the platform.</p>

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
			during processing (data in use).	
7.	<p>Accountability</p> <p><i>Controllers and processors keep on file documentation ready to show to both data subjects and supervisory authorities the measures that have been taken to achieve compliance with the data protection principles.</i></p>	<p>Data Holders should document, in particular:</p> <ul style="list-style-type: none"> • Description of security measures (TOMs) applied in the node; • Record of processing activities, • Record of incidents related to data, • Data subjects' requests. 	<p>Platform Operator:</p> <ul style="list-style-type: none"> • Description of security measures (TOMs) applied in the central platform • defined data processing activities undertaken in the platform. 	<p>Data Users should produce a technical documentation of the model.</p>
8.	<p>Data sharing agreements</p> <p><i>Define the roles of the consortium partners (controller, joint controller, processor) and prepare data sharing agreements.</i></p>	<p>Data Holders should enter into collaboration agreement with Data Users, including provisions for data sharing, depending on the case.</p>	<p>Platform Operator can facilitate drafting and executing personal data sharing agreements by providing template agreements to be entered between Data Holders and Data Users.</p> <p>Platform Operator, if acting as an intermediary, and not using the data exchanged for its purposes, would act as processor.</p>	<p>Data Users should enter into collaboration agreement with Data Holders (e.g. study agreement), including provisions for data sharing and applicable security measures, depending on the case.</p>

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
			The terms of use of the Platform and its privacy policy should be available to Data Holders and Data Users.	
9.	<p>DPIA</p> <p><i>Evaluate the risks related to the data processing activities of the FL.</i></p>	<p>Data Holder as controller is responsible for assessing if DPIA is needed and carrying out a DPIA.</p>	<p>Platform Operator to support the controllers, consider preparing a common DPIA template with platform description and available risk mitigation measures, such as PETs. information.</p>	<p>If Data User acts as controller, it is responsible for assessing if DPIA is needed and carrying out a DPIA.</p>
10.	<p>Transfer of data to and from EEA</p> <p><i>Transfers of data to countries beyond EEA (EU + Iceland, Liechtenstein, and Norway) are subject to strict rules.</i></p> <p><i>This applies also to storage of personal data and access to it from outside of EEA.</i></p> <p><i>If the country is not adequate, additional security measures are required.</i></p>	<p>Node Holders need to be careful with using cloud storage solutions (AWS, Google etc) for their node, as additional restrictions can apply.</p>	<p>Platform Operator should assess whether the federated infrastructure enables access to personal data, model updates, logs, metadata, or other potentially identifiable information from outside the EEA.</p> <p>This assessment should include remote administration, cloud hosting, support access, aggregation services, and access by non-EEA data users.</p>	<p>Data Users located outside the EEA, or accessing the federated infrastructure from outside the EEA, should assess whether their participation involves a transfer (including access) of personal data under Chapter V GDPR. This assessment should cover not only raw data, but also access to pseudonymised data which may be reidentified, model updates, gradients, AI</p>

	Title and Description	Data Holders / Data Hubs	Platform Operator	Data Users
			<p>Platform should include technical safeguards enforcing the data use restrictions in the study agreements and guard against downloading, copying or attempting to reconstruct personal data unless expressly authorised and legally covered.</p> <p>Where a transfer under Chapter V GDPR may occur, the Platform Operator should support controllers by documenting relevant data flows, identifying third-country access points, implementing supplementary technical measures, and providing information needed for TIA. Where appropriate, the Platform Operator should ensure that standard contractual clauses or other transfer mechanisms are in place.</p>	<p>models and their outputs that may allow re-identification.</p> <p>Where required, Data Users should enter into appropriate transfer arrangements, support TIAs, comply with supplementary safeguards, and avoid downloading, copying or attempting to reconstruct personal data unless expressly authorised and legally covered.</p>