D7.4 Technical specification for Health Data Access Bodies on the implementation of secure processing environments

Technical, functional and security specifications of secure processing environments

TEHDAS2 – Second Joint Action Towards the European Health Data Space

24 February 2026

Co-funded by the European Union

0 Document info

0.1 Authors

Lead Author(s)	Lead organisation
Heikki Lehväslaiho	CSC – IT Center for Science, Finland
Helena Lodenius	CSC – IT Center for Science, Finland
Beatriz Barros	Sciensano, Belgium
Alexandre Berna	Health Data Hub, France
Lucas Bréchot	Health Data Hub, France
Krisztina Fekete-Molnar	Sciensano, Belgium
Zdenek Gütter	Ministry of Health of the Czech Republic
Hans Aage Huru	Norwegian Institute of Public Health, Norway
Yohan Jarosz	Luxembourg National Data Service, Luxembourg
Todor Kondić	Luxembourg National Data Service, Luxembourg
Jaakko Lähteenmäki	VTT Technical Research Centre of Finland Ltd, Finland
Max Martens	BfArM - Federal Institute for Drugs and Medical Devices, Germany
Minerva Alvarez	Spanish Ministry of Health
Juha Pajula	VTT Technical Research Centre of Finland Ltd, Finland
Thomas Sondag	Luxembourg National Data Service, Luxembourg
Christophe Trefois	Luxembourg National Data Service, Luxembourg
Emmi Turunen	HUS Group, the joint authority for Helsinki and Uusimaa, Finland

0.2 Keywords

Keywords	TEHDAS2, Joint Action, Health Data, Health Data Space, Secure Processing Environments, SPE Federation, Federated computing

0.3 Document history

Date	Version	Editor	Change	Status
11/10/2024	0.1	Helena Lodenius, Heikki Lehväslaiho	Table of Contents	Draft
27/06/2025	0.2	Heikki Lehväslaiho, Helena Lodenius, Beatriz Barros, Alexandre Berna, Todor Kondić, Jaakko Lähteenmäki, Juha Pajula	Draft to be reviewed by the Consortium	Draft
12/09/2025	1.0	Heikki Lehväslaiho, Helena Lodenius, Beatriz Barros, Jaakko Lähteenmäki	Document to be submitted for public consultation	Final
02/02/2026	1.1	Heikki Lehväslaiho, Helena Lodenius, Beatriz Barros, Krisztina Fekete-Molnar, Jaakko Lähteenmäki, Juha Pajula, Anne Heidi Skogholt	Draft D7.4	Draft
24/02/2026	1.2	Heikki Lehväslaiho, Helena Lodenius, Beatriz Barros, Krisztina Fekete-Molnar	D7.4	Final

Accepted in Project Steering Group on 24 February 2026.

Table of contents

1 Executive summary

This report presents the technical, functional, and security specifications of **Secure Processing Environments (SPEs), a central component of the European Health Data Space (EHDS) as required under Article 73 of Regulation (EU) 2025/327**. SPEs are designed to enable the safe secondary use of electronic health data while ensuring compliance with data protection, confidentiality, and information security obligations.

Based on a thorough analysis, this report defines **a structured set of minimum requirements for SPEs**. It covers core capabilities needed to safeguard any sensitive data, as well as enabling requirements arising from the needs of scientific research principles. **A generic SPE specification** is developed that is flexible enough to fulfil current and future functional requirements. Functional and operational needs of **two main SPE use cases identified to be needed by EHDS** are shown to be derivable from this model.

The report goes beyond the obligatory demands of EHDS to define minimal requirements of interoperability between compatible services that form an **SPE-based federation** that is required when data needs to be transferred between organisations. The SPE federation model is further expanded with a set of tentative **practical implementation requirements for federated computing** that is still an active research area.

The primary purpose of this report is to support policy alignment, harmonised interpretation of legal obligations, and consistent high-level design choices across Member States. It does not aim to serve as a complete practical implementation manual or prescribe specific technologies or architectural solutions. Detailed implementation guidance, reference architectures, and operational playbooks will be developed in subsequent work, including implementing acts, pilot activities and follow-up guidance produced beyond TEHDAS2, for instance by the SPE Community of Practice subgroup. The present document should therefore be understood as a foundation for further specification, standardisation and implementation activities.

Justification of the report focus on high-level functional requirements over technical ones is given in the appendices that cover existing solutions, pitfalls of too narrow approaches, and crucial interplay of SPEs with other services in the ecosystem. Attention is given to how various approaches and technical implementations affect the **trade-offs between data security and usability, how they affect the health data users, and the responsibilities of actors accessing sensitive data.**

These specifications are intended to support Member States and stakeholders in the design and operation of SPEs and to inform the work of the European Commission in the preparation of the implementing act under Article 73(5) of the EHDS Regulation.

2 Introduction

Advancing health data use in the European Health Union

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

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

TEHDAS2 focuses on several critical aspects of health data use.

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

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

The aim of TEHDAS2 T7.4 deliverable is to develop technical specifications for the lawful and effective operation of Secure Processing Environments (SPEs) in accordance with Article 73 of the European Health Data Space (EHDS) Regulation. These requirements are used to describe functional, security and operational capabilities needed to build SPEs and supporting services. This report defines common concepts, minimum requirements and high-level specifications for SPEs. The document is intended to support policy alignment, strategic planning and high-level design across Member States and stakeholders. This report is not intended to be a standalone implementation manual. It does not prescribe specific products, detailed configurations, or deployment procedures, and it cannot by itself guarantee

operational compliance or certification of SPEs. Implementation requires additional technical documentation, organisational procedures, risk assessments and national choices that are outside the scope of this deliverable. The content presented here should be understood as a framework for future work. More detailed implementation guidance, including reference implementations, interoperability profiles and operational guidelines, will be developed beyond TEHDAS2, in implementing acts of the EHDS Regulation and in national and European initiatives building and operating SPEs in practice. Further details on the methodological approach used to develop these specifications are provided in Annex 1: Methodology. A summary of the stakeholder feedback collected during the public consultation is available in Annex 2: Public consultation summary.

SPEs must support multiple use cases for secondary data use. Design and implementation choices will significantly influence their ability to accommodate these use cases in a compliant and effective manner. The impact of a precise definition of an SPE will be critical for the context of secondary use of health data. It will guide the use of technologies and capabilities of sensitive data environments across many domains within the European Union and potentially beyond. It is therefore important that the definition of SPE focuses on essentials and enabling the full data management life cycle.

Health data is inherently dynamic and evolves over time as individuals generate and share information throughout their lives and during medical care. Initially used for direct patient treatment (primary use), this data can later be repurposed for broader goals like public health, research, and policymaking (secondary use). As data moves through various stages – collection, processing, analysis, and storage – it is transformed and enriched, often combined with other datasets. This dynamic flow of information across systems and organisations requires strong collaboration and clear governance. The EHDS aims to give citizens more control over their data and ensure transparency about its use. To achieve data life cycle support, interdisciplinary understanding across legal, organisational, semantic, and technical layers, as outlined in the European Interoperability Framework (EIF), is required.

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

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

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

3 Scope

This report takes the concept of secure processing environment (SPE) and analyses to ensure the data privacy of processed sensitive data, as well as the privacy and accountability of the user of sensitive data.

While the context of this report is the EHDS Regulation, the use and applicability of SPE is wider. This has led to several distinct but interdependent sets of requirements where the terminology reflects interpretation of EHDS obligations and recommended minimum standards but does not constitute binding legal requirements beyond the EHDS Regulation. This report is divided into two main parts to separate these sets clearly (Figure 3.1). In addition, requirements dealing with federated aspects of SPE are not mandatory under EHDS Regulation.

Chapter 4 analyses the general SPE concept first as a stand-alone environment that offers a uniform approach to GDPR requirements. Then it expands its view to a collection of distributed SPEs that need varying levels of interoperability to accomplish increasingly complex computations. The results are presented as a series of requirements aimed at defining high-level design for secure processing that enables scientific research that sets guidelines for policy alignment that is independent of technical implementation details.

Chapter 5 goes into details analysing how the EHDS Regulation fits into this framework. Specific functional and operational SPE requirements of EHDS legislation are presented and EHDS-specific roles responsible for them are named. The stand-alone SPE, that is a strict requirement of the EHDS Regulation, is analysed separately from the needs of an SPE federation. Technical interoperability requirements for an EHDS-compliant federated learning setup are presented for future reference.

Annexes give more details of the regulatory analysis, existing solutions of SPEs, and challenges in implementing SPE functionalities.

Readers are encouraged to familiarise themselves with ArchiMate Enterprice Architecture notation used throughout in figures and the standard modality terminology for requirements that are defined in Annex 1.

SPEs are critical to enabling lawful secondary use of health data while safeguarding individual privacy and ensuring compliance with GDPR and EHDS Regulations (EHDS SPEs). To clarify the distinction between data processed within the EHDS regulated SPEs and data processed in generic SPEs, it is essential to underline that only EHDS SPEs are subject to legally mandated requirements for access control, auditability, interoperability, controlled outputs, and continuous security monitoring. While generic SPEs may offer technical safeguards, they are not bound by the statutory obligations defined in Article 73 of the EHDS Regulation and therefore do not provide the same level of governance oversight, and protection for secondary use of electronic health data.

Figure 3.1. Structure of the report. The ArchiMate graph shows as a work package with the drivers sensitive data protection and scientific research and goals with requirements are numbered with report chapters. Colours are for illustrative purposes.

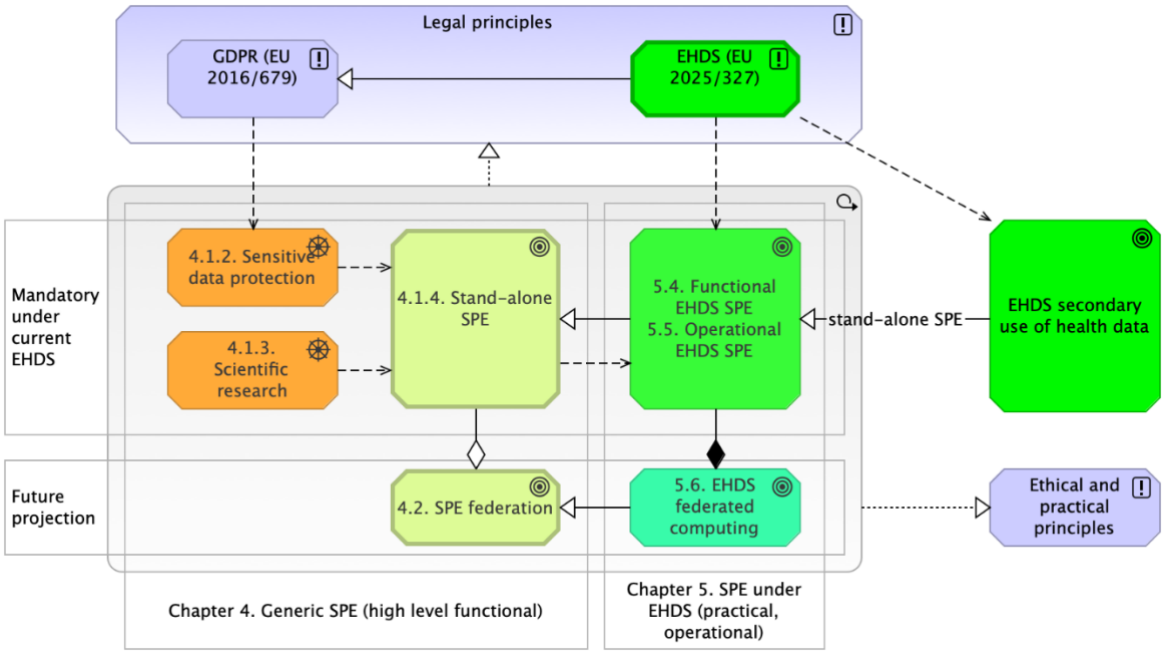

4 Generic SPE

4.1 SPE as service

While cybersecurity and data confidentiality has long history, the European General Data Protection Regulation (GDPR EU 2016/679) precisely defined the defined means and sensitivity of the data. See chapter Key GDPR data and processing requirements in Annex 9: Overview of relevant EU regulations and Annex 5: Historical context and legacy models for earlier attempts and consideration to restrict access to data.

The scope of EU regulation was later expanded to more explicitly mention non-commercial and academic activities. The European Union Data Governance Act (DGA, EU 2022/868) provides the currently valid definition of SPE:

> ***'secure processing environment' means the physical or virtual environment and organisational means to ensure compliance with Union law, such as Regulation (EU) 2016/679, in particular with regard to data subjects' rights, intellectual property rights, and commercial and statistical confidentiality, integrity and accessibility, as well as with applicable national law, and to allow the entity providing the secure processing environment to determine and supervise all data processing actions, including the display, storage, download and export of data and the calculation of derivative data through computational algorithms****;* (DGA, Article 2)

We will proceed analysing **SPE as a service** that can be used in multiple use cases, domains of knowledge, and governance structures. We will define high-level requirements of SPEs based on legal drivers and system-level goals. Each set of requirements will need to be defined separately as their scope is wider than following ones that depend on them.

4.1.1 Principles

Principles provide justification for service design. Here**, legal principles** are EU-wide laws that guide or define legal requirements. They can be referred to by name or down to a specific paragraph of an article that spell out the specific requirement. The most important laws for SPEs and sensitive data are covered in Annex 9: Overview of relevant EU regulations.

With respect to service requirements, practical and ethical design principles are best categorised as either **enabling or trust-dependent** ones. Trust covers traditional cybersecurity principles but also wider aspects like privacy and transparency (Figure 4.1).

Figure 4.1. Principles affecting the SPE service provision that separates enabling and trust-dependent principles. GDPR is the main European legislation determining the application of data confidentiality requirements that balances a wide selection of these principles.

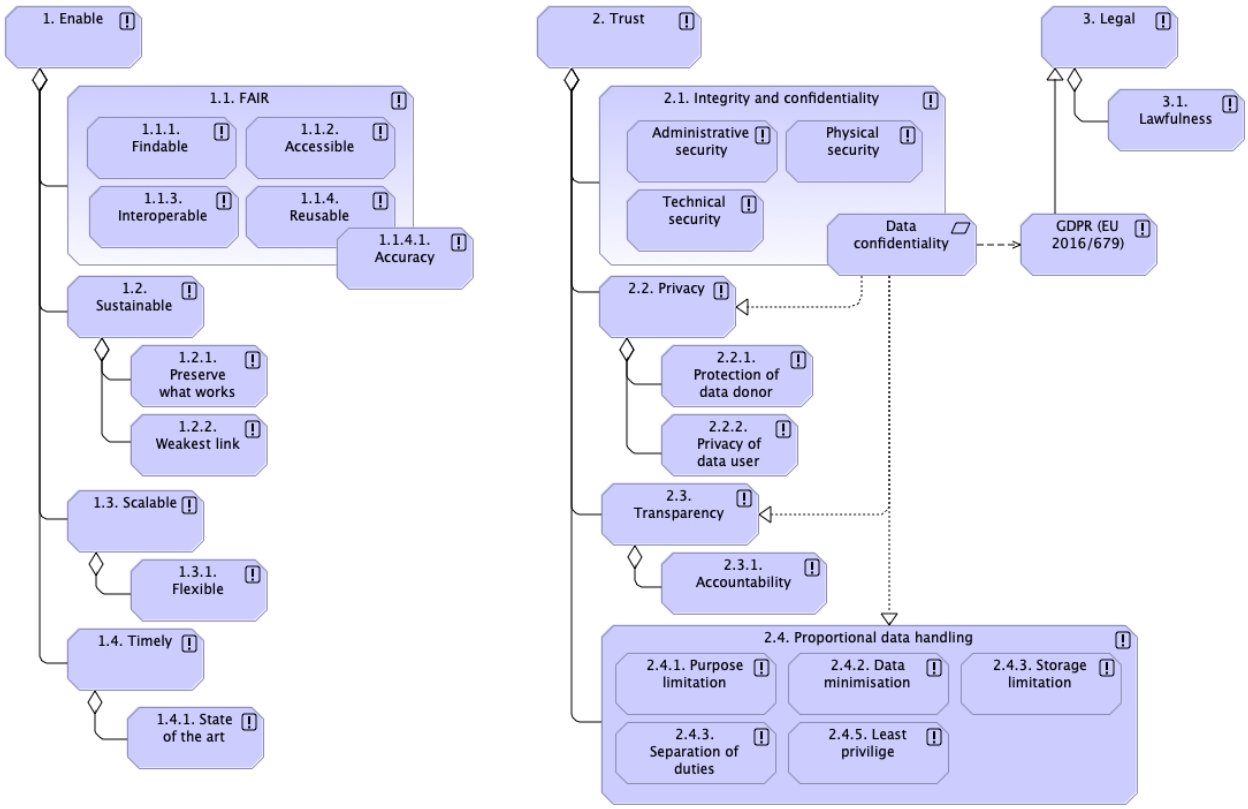

Enabling principles relate to the benefits and functionalities that SPEs are expected to deliver. These may vary depending on the perspective – for example, from a societal viewpoint (e.g. scientific progress), a service provider's viewpoint (e.g. performance, scalability, maintainability), or a data user's viewpoint (e.g. usability, analytical capabilities, cost).

The most important set of enabling principles data management are FAIR. The combination of findability, accessibility, interoperability and reusability determine how well data and services support scientific research.

Other enabling aspects gauge if the design is reasonable for current needs (state of the art), it can be maintained with the resources available (sustainable), will be able to respond to ever-changing needs (flexible), and provide services in a timely manner (timely).

Data confidentiality is a key security principle applied in data management to ensure trust. The GDPR covers all these aspects and requires appropriate measures for confidentiality and integrity based on the level of perceived risk.

4.1.2 Sensitive data protection requirements

Sensitive data differs from classified data in its requirement to be not only protected but also be available for secure processing

The sensitive data protection requirements listed in table 4.1 and illustrated in figure 4.2 are sequentially numbered and preceded by the namespace 'Sensitive Data' acronym 'SD' followed by the letter 'R'. SDR requirements fall under EHDS Regulation.

Table 4.1: Sensitive data requirements.

#	Requirement	Source	Importance
SDR-1	Unauthorised persons MUST NOT be able to access sensitive data	GDPR Art. 25, GDPR Art. 32	Mandatory
SDR-2	Service administrators SHOULD NOT access sensitive data	GDPR Art. 25, Principle 2.4.1 Purpose limitation	Recommended
SDR-3	Sensitive data MUST be in a protected format at rest and in transit	GDPR Art. 32(1)(c)	Mandatory
SDR-4	Sensitive data protection MUST be done with widely accepted, secure algorithms combined with effective isolation measures	GDPR Art. 32(1)(c)	Mandatory

For a service to process sensitive data, it needs legal basis (GDPR Art. 6 and 9), follow core data protection principles (GDPR Art. 32) and implement appropriate safeguards (GDPR Art. 89).

All processing of sensitive data must be based on the **sensitive data privacy requirements**. The most important concept is the **authorised user** that identifies persons who have access to the sensitive information contained in the data (**SDR-1**). This may be defined as the collector of the information (primary use) or through an explicit permit (as in secondary use). It implements the principles of segregation of duties and least privilege.

Everyone else is unauthorised, including possible administrators of the service used (**SDR-2**). In some cases, this may be difficult to maintain, hence this requirement is only a recommendation to allow administrators access when their help is needed in exceptional situations. Security and privacy measures to protect sensitive data should always mention if they are directed to support or control authorised users, deter unauthorised users, or both.

Sensitive data security is fundamentally based on user identification and the processes used for ID authentication, validation, and verification. Authorisation to access to that is internal to the project that may be defined by an explicit permit. User identification and measures to authenticate users are the foundation of sensitive data security. A high level of identity assurance should be required, and further measures of authentication like multifactor authentication should be adopted.

To narrow the possibility that sensitive data is available to unauthorised users, it needs to be protected by default (**SDR-3**). Encryption is the default way to protect digital information, and it lessens needs to isolate and guard the data. However, this creates the problem of handling encryption keys (secrets management) that need to be present when data will be used.

Encryption algorithms are constantly improved to answer to new vulnerabilities that technical development brings up. Encryption should rely on peer-reviewed, widely accepted algorithms that are aligned with recognised international standards. Only open encryption algorithms are

generally considered safe. Encryption protection is complemented with isolation measures that are either physical, technical or operational (**SDR-4**).

Additionally, pseudonymisation is a widely used method to protect the privacy of individuals from authorised users when they do not need all information and from accidental exposure to unauthorised users. Pseudonymisation and non-sensitive data types such as anonymised and synthetic data are covered in TEHDAS2 *D7.2 Guideline on data minimisation, pseudonymisation, anonymisation and synthetic data.*

Figure 4.2. General sensitive data protection requirements ArchiMate graph.

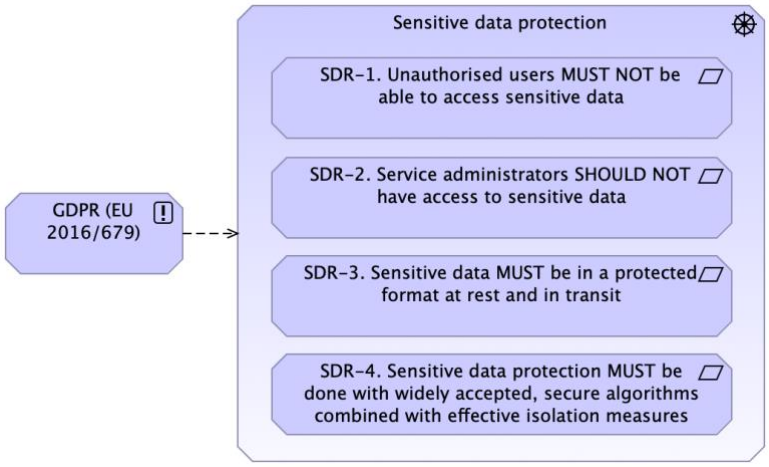

4.1.3 Enabling needs of scientific research

In this report, we are using scientific research as the ultimate yardstick for the enabling needs of sensitive data processing. In its purest form, scientific research is open-ended. Its results cannot be fully predicted in advance. Unexpected results that give new light to the study area are a real goal of scientific research. Same applies to applications needed to gain these insights: only the researchers themselves who are developing them will be able to evaluate their security and install them. To advance scientific research, users need to be able to perform all these actions.

Scientific research works as a good representative for controlled, IP-regulated (proprietary, innovation-driven) data processing. However, not all use cases need wide freedom. A more limited scope for data user makes it possible to simplify and streamline services created for sensitive data processing. Policy making and quality assessment, for example, can be served with services with less user choice.

The scientific research requirements listed in table 4.2 and illustrated in figure 4.3 are sequentially numbered and preceded by the namespace ‘Scientific Research’ acronym ‘SR’ followed by the letter ‘R’. SRR requirements fall under EHDS Regulation.

Table 4.2: Scientific research requirements

#	Requirement	Source	Importance
SRR-1	Sensitive data protection SHOULD NOT prevent processing for scientific research	GDPR Art. 32(1)(a), Art. 32(2), Art. 89	Recommended
SRR-2	Data protection needs MUST minimise impact on application of FAIR principles to sensitive data		Mandatory

It is important to see that the requirements and functionalities of SPEs depend equally on the restrictive pressure to **protect sensitive data** and the **enabling needs of scientific research** that is a is one of the favoured uses in GDPR restrictions (**SRR-1**).

The need to protect the privacy of the data donor and data processing as expressed in GDPR constrains how the FAIR principles that demand open science principles of findable, accessible, interoperable and reusable are applied to sensitive data (**SRR-2**). GDPR primarily affects the accessibility and reusability aspects of FAIR by needing accountability for data access.

The connection to scientific research helps us to understand the important key difference between sensitive and classified data. Similar methods are used to secure both types of data, but these data types have opposing aims. Secrets are meant to be dangerous facts that society wants to protect and make available for a very limited group of people only when necessary. Protection of sensitive data aims to enable lawful and open-ended research on sensitive data as widely as possible to benefit society.

Scientific research is a curious combination of society's need to promote the academic curiosity of individuals to reap benefits from their explorations of unknown in a way that unexpected results are common (Figure 4.3).

Figure 4.3. Scientific research requirements. Requirement identifiers follow the convention 'SRR–n', where 'SR' indicates the scientific research requirements namespace and 'R' denotes requirement.

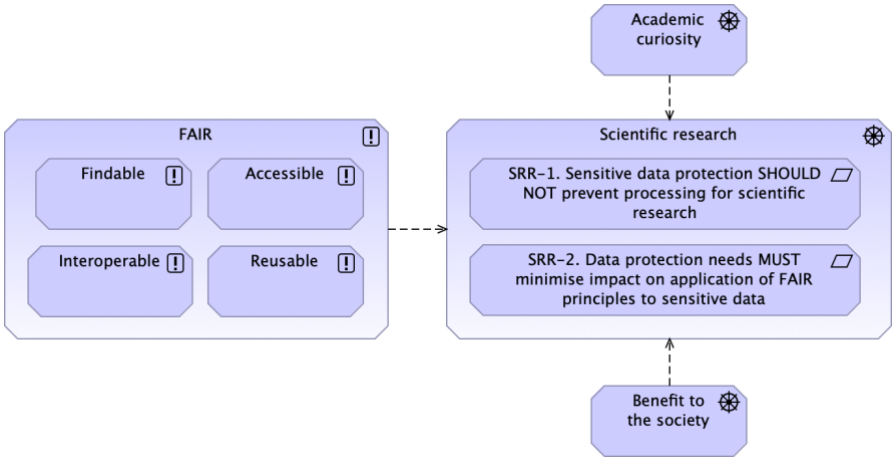

4.1.4 Stand-alone SPE

When applying the wide and risk-based requirements from general sensitive data protection to GDPR-based standalone SPE according to its DGA definition, they can be reduced to two main security conditions: monitoring and isolation (Figure 4.4).

Figure 4.4. Legal conditions of SPE

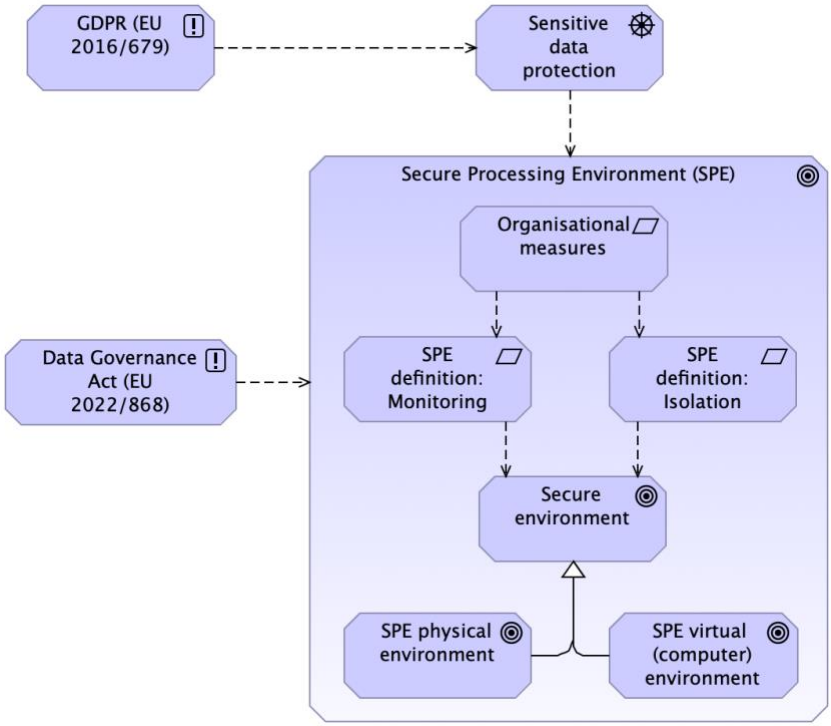

In this context, we can ignore the physical environment and focus on SPE as a virtual computer environment. Then we can say that **SPE operators** provide a secure service that isolates sensitive data processing within its confines. Isolation is provided based on projects that are separated from each other. The monitoring covers service access, data transfers, and of course the success of the isolation (Figure 4.5).

Users have access to their **project space** that contains their sensitive data. The project space has the ability store and process sensitive data. For **data users**, this project space is their “SPE” that is different from the SPE service. This difference between the two SPE concepts has been a source of much confusion.

Figure 4.5. Strategic view of a generic SPE. SPE is represented as abstract ArchiMate capability that has resources it needs to have or react to.

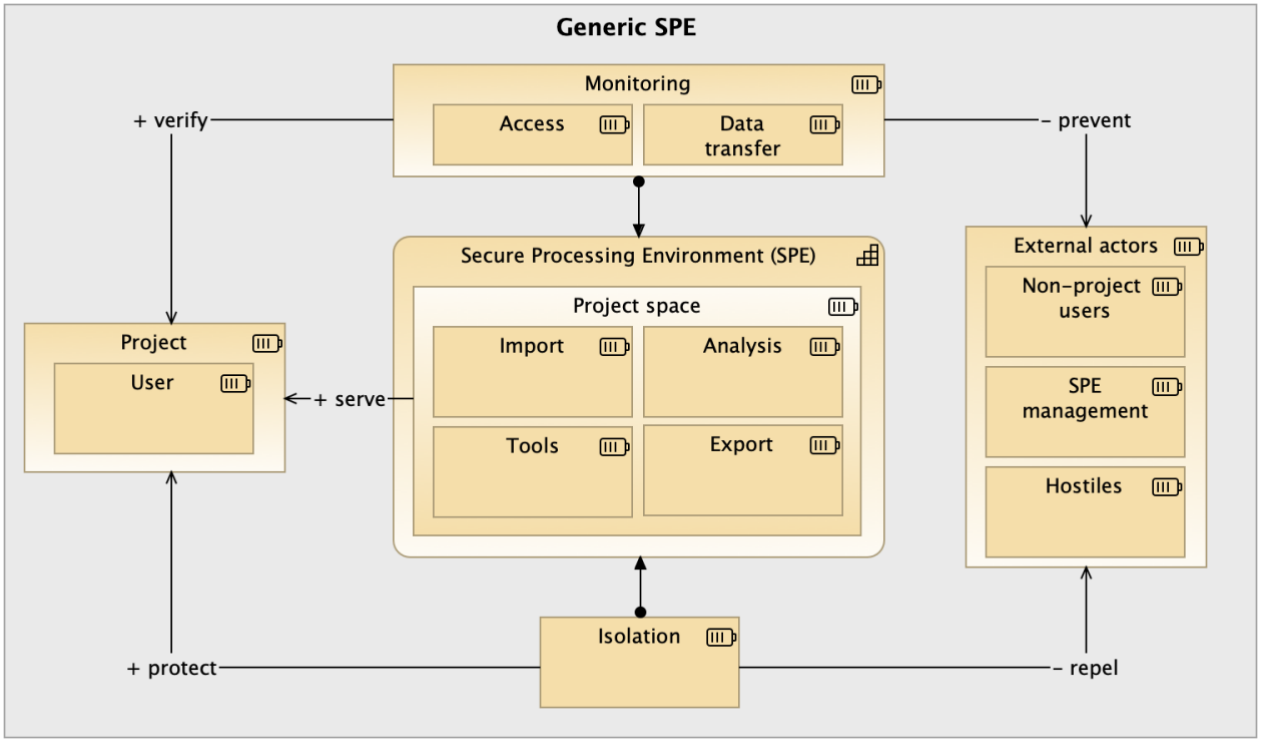

The general purpose SPE requirements listed in table 4.3 and diagrammed in figure 4.6 are sequentially numbered and preceded by the namespace acronym ‘SPE’ followed by the letter ‘R’. SDR requirements fall under EHDS Regulation.

Table 4.3: Stand-alone minimum SPE requirements

#	Requirement	Source	Responsible role	Importance
SPER-1	SPE MUST enable scientific research on sensitive data	SSR-1	SPE Operator	Mandatory
SPER-2	There SHOULD be a diverse selection of SPEs for the varied needs of sensitive data research	SSR-1	SPE Operator	Recommended
SPER-3	It MUST be possible to transfer sensitive data between, in and out of SPEs	SSR-1, SDR-3	SPE Operator	Mandatory
SPER-4	SPE MUST provide adequate protection against exposing sensitive data to unauthorised users	SDR-1, SDR-2	SPE Operator	Mandatory
SPER-5	SPE design SHOULD promote collaboration	SRR-2	SPE Operator	Recommended

	among authorised users			
SPER-6	Authorised users MUST protect sensitive data they display	SDR-3	Data User	Mandatory
SPER-7	Project-based user environments of SPE MUST NOT have open connections	SDR-3, SDR-4, SPE definition: Isolation	SPE Operator	Mandatory
SPER-8	SPE user accesses and data transfers MUST be logged and monitored	SDR-4, SPE definition: Monitoring	SPE Operator	Mandatory

SPE has been defined to meet the demands of scientific sensitive data processing and, by extension, other uses that need high level of protection (**SPER-1**)[1].

Openness and interoperability as expressed in the FAIR principles are essential elements of scientific research. In the context of sensitive data processing, they also embody the unpredictability of scientific research results that must be considered in the SPE design. Since it is impossible to determine beforehand what twists and turns empirical scientific research will take, the functional definition of an SPE must be open-ended. In other words, we cannot limit the SPE capabilities *a priori*.

Hardware and software capabilities of SPE instances is and will remain to be variable. It is impossible for any single SPE to fulfil all present and future requirements. Data capacity and analysis capabilities will differ widely and create a market for different SPEs (**SPER-2**).

For users, the primary advantage of SPE is the compliance to GDPR requirement that allows them to focus on data analysis. In the default case, the users are primary controllers of sensitive data and they need the ability to move code, applications and sensitive data as they please for processing in the safe environment that SPE provides (**SPER-3**). These abilities may be limited by data permit if the sensitive data is under controlled access secondary use.

Activities that create sensitive data are out of scope but to make datasets available for analysis they need to be transferred into an SPE. Also, the diversity and changing needs of research might make it necessary that sensitive data must be transferred between SPEs during the analysis. This does not mean that that transfer is necessary direct.

Users should be able trust the technical functionalities the SPE is providing (**SPER-4**).

SPE users have the right to expect that basic tools needed for their analysis, and more generally conducting their research on sensitive data, are supported by the SPE service (**SPER-5**). This could include a collection of preinstalled common libraries and tools but also more comprehensive collaboration among project members (See Annex 7: Design considerations and expert commentary – SPE as collaboration area).

A related and resolved issue is the SPE support to anonymous reviewing of results in cases where sensitive data is prohibited to be exported out of the SPE. SPE service might respond to this by making it possible to have unidentifiable aliases to users.

[1] The 2025 ruling by the European Court of Justice on lesser sensitivity of pseudonymised personal data is included in the Digital Omnibus legal package under preparation. It could have impact on SPE use.

Authorised users must bear the responsibility of their actions. For most of the expected SPE use, authorised users display plain-text sensitive data on their computer screens, and they must be aware that it is their responsibility to protect it (**SPER-6**).

Monitoring displayed sensitive information in modern distributed and virtualised environments is hard. If the use of SPE is made too cumbersome for technical or regulatory reasons, users are easily tempted to take note of the displayed information from the screen rather than follow official ways. In practice, this means that any addition of any security measures that limit the ability of authorised users to process sensitive data need to be evaluated against this danger. Ultimately there are no foolproof technical means to ensure that users do the right thing, or even that they know what that is. It is therefore necessary to ensure that SPE users have the proper understanding of their responsibilities. See Annex 7: Design considerations and expert commentary – Priority of user training.

Stand-alone SPEs must be by default designed in such a way that user environments within it are clearly and completely isolated from each other and from the open Internet (**SPER-7**). Any explicitly opened connections must be secure. This is needed for both control and protection of the environment. We will see that there are many practical reasons to open connections in controlled way to meet real world needs.

The first of them applies to use cases where users are primary controllers of the sensitive data. They might want to give access to other projects within the same SPE. Secure sharing of datasets can then be part of the SPE interface, and projects can focus on the contractual side of the transaction.

Monitoring and logging are requirements that include transparency and data minimisation under GDPR and are also mentioned in the SPE definition. The requirement **SPER-8** takes the stand balance to risk and feasibility as required by GDPR Article 5(1)(c) and Article 32 that access logs and data transfer logs are crucial to monitor and contain useful information. User action logs are seen to be coved by user privacy and unlikely to contain useful information in its impractically massive volumes. See the reasoning in more detail in Annex 7: Design considerations and expert commentary – Monitoring of SPE use

Figure 4.6. SPE minimum requirements and their derivation. Requirement identifiers follow the convention 'SPER–n', where 'SPE' indicates the general purpose SPE requirements namespace and 'R' denotes requirement. SPER requirements fall under EHDS Regulation.

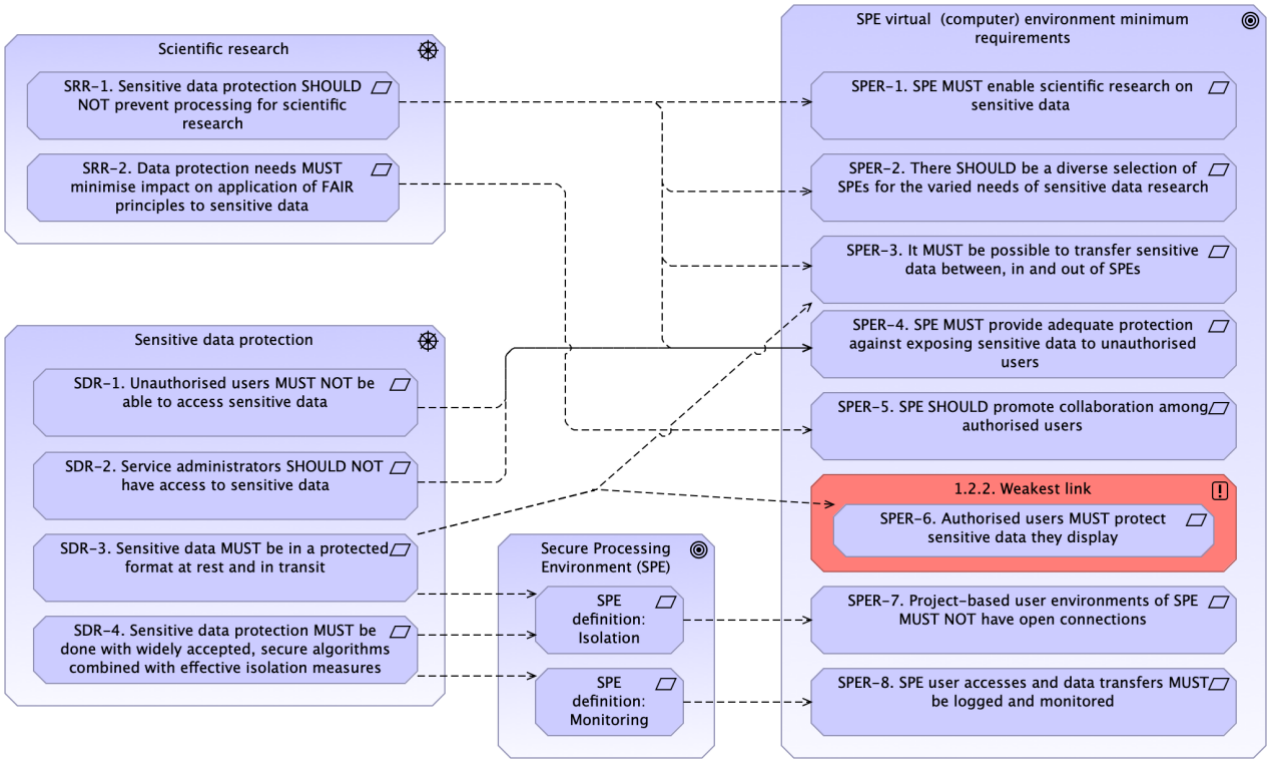

4.2 SPE federation

The preceding analysis of SPE requirements described a service that is strictly an independent service. Its legal definition stresses the legality, security and accountability of sensitive data processing as if there were only one of SPE. Having all data needed for a project in one SPE remains the main use case, but there are situations where it is not possible. This chapter expands those requirements to many SPEs and interoperability between them. Any means to enable data processing in more than one environment is called distributed processing.

The first stage of interoperability needs to come from the set of services that fall under the same category. From both policy and user point of view it is necessary to know how many there are, and in what respects they differ from each other. There is a need for a findability service that in turn needs a shared vocabulary that is used to build a registry of services.

Services like SPE that need data must have means to transfer data in and out, as well as programmes implementing algorithms to process the data. Convenience together with security concerns pushes these to follow common practices. Users and data managers have the need to utilise different capabilities of SPE services that change over time.

SPE was envisioned as a convenient place that needs a set of safety rules to perform calculations on sensitive data. The need to provide continuous protection to sensitive data

in a distributed world has led to ecosystem thinking like in European data spaces. Within them, SPE functionalities and safety requirements need to cover all aspects of data processing. Under European laws, this processing is best organised using SPEs through which users interact with sensitive data.

4.2.1 SPE federation requirements

Interoperability works best when it is defined in hierarchies with a tightening set of rules and their applications. A general federation covers the most fundamental interoperability rules, next one defines a domain, next a use case, always building on top of previous rules. This way a project involving SPEs in different countries is a federation defined by the data permit that provides the contractual basis.

We sketch this out by defining base requirements for SPE federations in table 4.4.

Table 4.4: Minimum SPE federation requirements. The requirements of general purpose SPE federation have the namespace acronym 'FSPE' for Federated SPE followed by the letter 'R'. FSPE requirements do not fall under EHDS Regulation. Importance levels of its requirements are valid only within its namespace.

#	Requirement	Importance
FSPER-1	Legal or contractual agreement MUST cover the SPE federation across organisations	Mandatory
FSPER-2	Federation user identities MUST match over services	Mandatory
FSPER-3	SPE federation environment MUST fulfil sensitive data processing requirements of stand-alone SPEs	Mandatory
FSPER-4	All interactive user actions on sensitive data in the federation MUST be through SPE	Mandatory
FSPER-5	Federation governance structure MUST cover secure, shared data access and export from SPE	Mandatory
FSPER-6	Federation user identities SHOULD be shared	Recommended
FSPER-7	Federation SPEs MUST share technical and semantic interoperability needed for shared processing	Mandatory
FSPER-8	Federation SPEs MUST use shared secure communication and data transfer protocols	Mandatory
FSPER-9	The federation MUST support distributed processing through authorisation and accounting services	Mandatory
FSPER-10	A federation SPE MAY fulfil federated computing requirements	Optional

A formal agreement between SPEs is the foundation of an SPE federation (**FSPER-1**). The rules of federation must be based on writing for transparency and enable shared sensitive data processing across organisations.

An SPE federation needs ways to identify and authorise its users in all federation services (**FSPER-2**). This can be done by matching local identities over federated identities.

An SPE federation is based on SPEs providing the main sensitive data protection. None of the stand-alone SPE rules should be violated in an SPE federation (**FSPER-3**).

While the aim of the SPE federation is to widen the scope and quantity of sensitive data processing via automated, remote processes, the users must log in to an SPE to initiate the

processing and export the results (**FSPER-4**). The customary external programmes that control the processing in general federated computing, are not allowed.

The federation will need to provide the governance structure that defines operational procedures for secure, shared processes for data transfer (**FSPER-5**).

While the federation can function with matching user identities as required in FSPER-2, its authorisation services are significantly strengthened by a federation-wide identity system **FSPER-6**. This is a nod of approval towards the shared European eiDAS system. A shared user identity system may enhance service authorisation and auditability to better comply with GDPR principles.

Shared processing needs agreements on technical and semantic standards that will enable it (**FSPER-7**). In a European context, the most common health data model for remote data queries is OMOP CDM.

Security of communication between federation members, especially when transferring sensitive data, is the fundamental requirement (**FSPER-8**). EU eDelivery and GA4GH crypt4GH are examples of these (see chapter Data access management and SPE interoperability).

A federation needs – usually centrally managed – registries that secure non-interactive, distributed processing. While authorisation has been traditionally been an integral part of grid computing, accounting and billing has been ignored to the detriment of long-term sustainability of the service (**FSPER-9**). Federation registries need to have infrastructure metadata about its authorised services and service providers, content metadata for dataset discoverability, and governance metadata about data permits that link project and users to datasets and resource usage[2].

Specialised distributed computing methods are usually called federated computing that is a separate set of requirements that are based on federation principles (**FSPER-10**). Federated processing is its own branch of distributed processing that some members of the SPE federation may enable (See chapter Federated computing).

[2] DARE UK Federated Architecture Blueprint https://doi.org/10.5281/zenodo.14192786

Figure 4.7. SPE federation requirements.

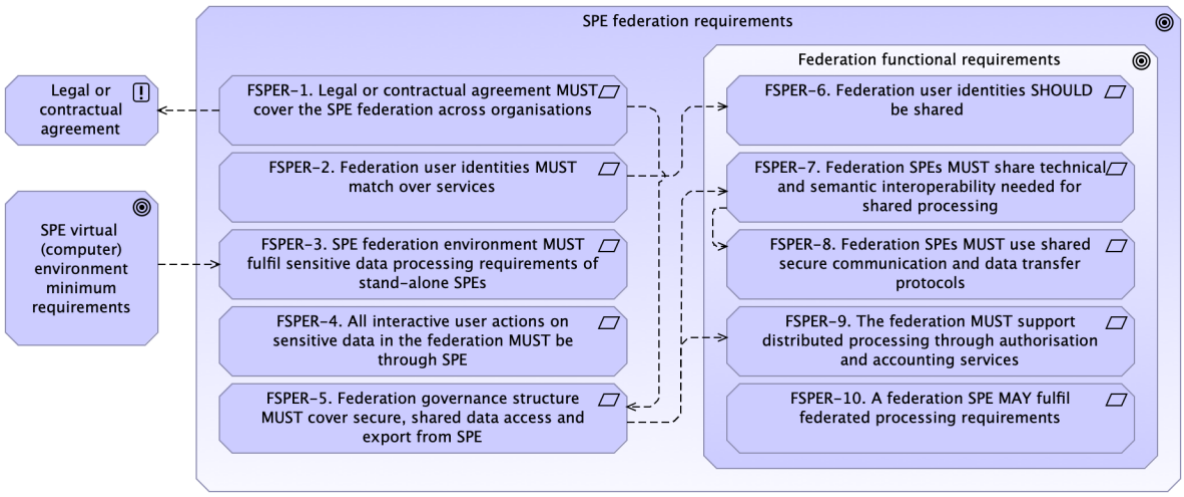

4.2.2 Federated computing requirements

Scope and definitions

This section provides a proposal for supporting federated computing within the SPE federation defined in the previous chapter. A high-level conceptual architecture is presented. The impact of federated computing on the data permit application phase, including how the involved SPEs are selected by the health data user and approved by the HDAB is out of scope of the deliverable.

A wide variety of definitions related to federated computing exist. For the EHDS implementation purposes, it is important to distinguish between two scenarios: (1) cases where node outputs are anonymous and can be directly returned to data users, and (2) cases where the anonymity of node outputs cannot be fully guaranteed. The terms **"federated analysis"** and **"federated learning"** are used to describe these scenarios, respectively. This categorisation aligns with the different needs of data users and supports the possibility of gradual implementation of federated computing, starting from simpler and less resource-intensive setups.

It is also important to note that, the EHDS Regulation does not mandate that HDABs to provide federated computing services.

Federated computing[3] is a decentralised data processing approach where computations occur locally on distributed SPEs rather than being centralised into a single SPE. Such approach is encouraged by EHDS Regulation (Recital 80) which promotes the principle of "bringing algorithms to the data" to enhance privacy-preserving computation. Federated computing methods enable data to remain closer to their original location while only aggregated results or model updates are shared, enhancing privacy and security.

[3] Federated computing is not mandated under the EHDS Regulation, but may be supported under future implementing acts (Article 73(5))

Federated analysis is a form of federated computing in which each participating SPE performs local computations and returns results that can be considered anonymous by design (e.g., aggregated or statistical outputs). These results can be directly combined by a coordinating entity (the data user application, see Section 5.6) without additional anonymisation steps. This method enables benchmarking, multi-site research, and collaborative analytics while preserving data privacy and security.

Federated learning is a form of federated computing in which participating SPEs collaboratively train models by exchanging intermediate results (such as model parameters or gradients). Individual-level data are not exchanged between SPEs. Instead, only model updates are communicated thereby enhancing data privacy and security. In this case, the anonymity of exchanged information cannot be fully guaranteed, and additional privacy-preserving and trust mechanisms may be required. Federated learning may involve information exchange between computing nodes under defined mutual trust conditions.

Due to the difficulty of assessing the anonymity of intermediate outputs, it is essential that federated learning computation and information exchange takes place within a network of trusted SPEs (the SPE federation).

High-level federated architecture

The proposed federated computing approach aligns with Recital 80 of the EHDS Regulation, which encourages bringing questions to the data instead of moving the data. However, the EHDS Regulation does not provide further direction on how to support the implementation of federated computing. Further guidance is expected in the context of implementing acts concerning secure processing environments.

Two settings can be identified. In the first (Figure 4.8.a), the health data user's application interacts independently with multiple SPEs via API interface, retrieving anonymised outputs from each and combining them to compute statistical results. In this case, no communication occurs between SPEs. A prerequisite for this setting is that the SPE operator provides an API interface and related governance processes as outlined in Section 5.6.

Figure 4.8.b illustrates a second setting where a master SPE (SPE#1) communicates with other SPEs to perform a federated learning task. The master SPE is responsible for orchestrating the federated learning process, collecting computation results from the other SPEs to iteratively train or validate a machine learning model. The health data user retrieves the final output (e.g., the trained model) from the master SPE after privacy risk assessment and disclosure approval by the HDAB (see TEHDAS2 deliverable D7.2).

A prerequisite for this setting is that the SPE operator provides an open and standardised interface needed to exchange information with other trusted SPEs as outlined in Section 5.6. Furthermore, an SPE federation (see SPE federation) is a prerequisite for the setting of Figure 4.8.b to establish the framework for communication between trusted SPEs. In order to further minimise and control risks of personal data leaks the SPE operator may support and require the use of privacy protection methods (such as differential privacy).

Figure 4.8. (a) Data user application collects and further processes anonymous computation results from multiple SPEs. (b) Data user application is connected with one master SPE, which in turn communicates with other SPEs under trust conditions to accomplish a federated computing task.

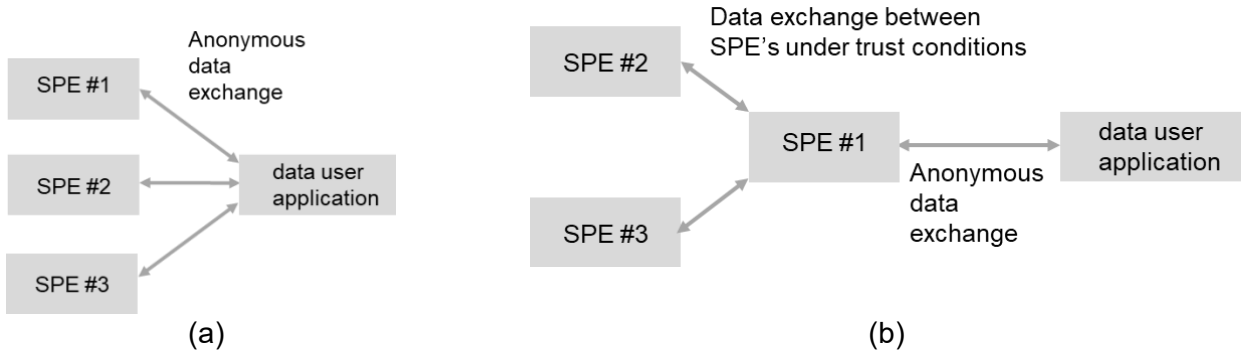

The scenario illustrated in Figure 4.8.a is relevant for federated analysis use cases, where only anonymous aggregated results need to be transferred to the health data user by each SPE. Figure 4.8.b is relevant for federated learning use cases, where it is difficult to ensure the anonymity of the interim outputs of the SPEs and which thus need to be kept within the network of federated SPEs.

Assumptions

Proposed functional requirements for supporting federated computing are based on the following assumptions:

- Federated computing is applied in the framework of the EHDS Regulation complemented by more detailed specifications (e.g. implementing acts)
- All computations on sensitive personal data are carried out in an SPE environment aligned with EHDS requirements (e.g. Article 73 of the EHDS Regulation)
- All SPE services involved in federated computing are EHDS compliant, trusted, audited and under control of HDABs or trusted data holders (directly or through data processing agreements with SPE operators).
- SPEs involved in federated computing, i.e. **SPE federation,** may be located in one or more countries.
- Federated computing support is a recommended but optional feature of an SPE ("MUST" in a requirement means that the requirement is applicable if federated computing is supported by the SPE).
- An SPE may also support only a subset of federated computing requirements, e.g. limiting to federated analysis.

- All data resources used in a federated computing setup shall be covered by the same data permit, with approval by the relevant HDABs or trusted health data holders.

Functional requirements

The requirements of computing performed by a federation of SPEs are listed in Table 4.5. These requirements have the namespace acronym 'FC' followed by the letter 'R'. FC requirements cannot be derived directly from the EHDS Regulation but will be subject to it if these processes are used for secondary use of health data. Importance levels of its requirements are valid only within its namespace.

Table 4.5: Functional requirements of computing performed by a federation of SPEs.

#	Requirement	Importance	Role
FCR-1	SPE project space MUST support a shared data model	Mandatory	SPE operator
FCR-2	Controlled and secure data transfer supporting federated computing MUST be enabled between the participating SPE project spaces	Mandatory	SPE operator
FCR-3	Participating SPE project space MUST support the deployment and execution of software components needed	Mandatory	Data user
FCR-4	The project MUST be authorised to use federated computing	Mandatory	Data controller

The effect of these functional requirements is illustrated Figure 4.9. to project a vision of a federated SPE that communicates to other members of the federation through isolated secure gateways (security server). The main practical implications are that 1) the data accessible to users must be independently accessible to (permitted) federation services and that 2) security of the encrypted data needs to be maximised, preferably by dynamic, one-time keys managed by a dedicated secrets management component of SPE.

Figure 4.9. Projection of functional requirements for SPE in a federation. Motivation elements (violet) show the justification and strategy elements show how all communications to SPE and the user's project environment must go through isolated gateways within the security server that has its counterparts in all federation resources.

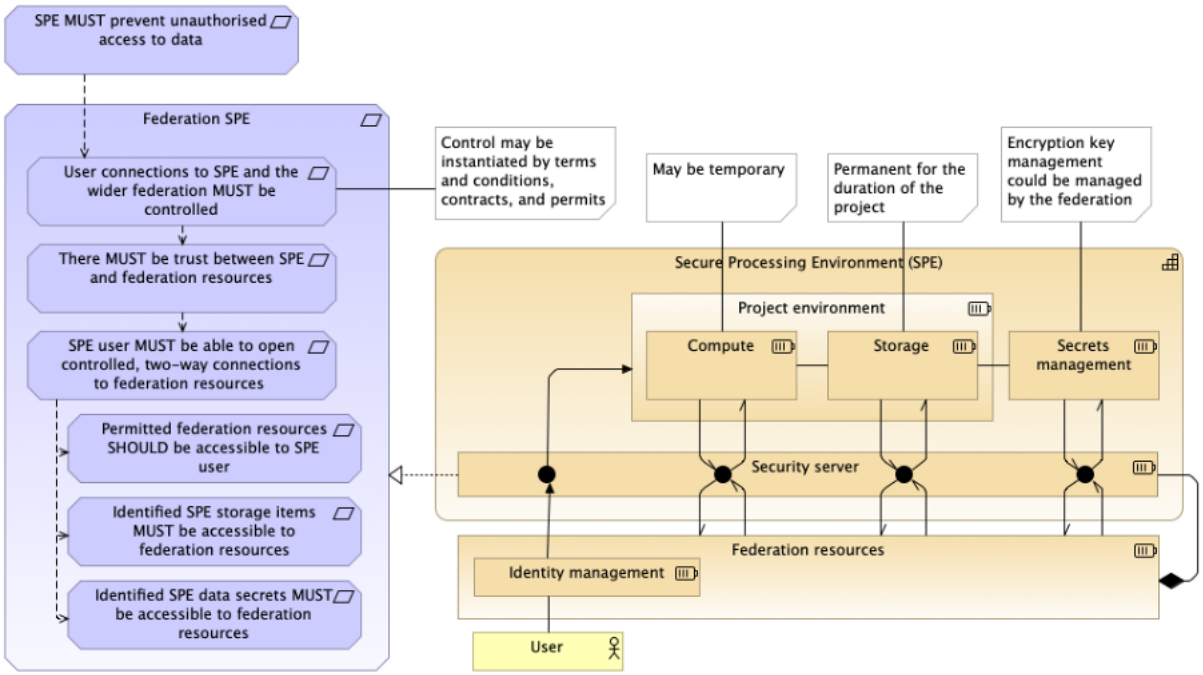

5 SPE under EHDS

5.1 Preliminary life cycle components of EHDS SPE

The use of an SPE within the HealthData@EU infrastructure occurs after a health data user (in distinction from the generic SPE user that is called *data user*) has been granted access to a specific dataset and a data permit has been issued by the HDAB[4]. The general life cycle of an SPE, as outlined in TEHDAS1[5], consists of the following steps:

Environment Creation: Once the data permit is issued, the designated SPE operator sets up an isolated environment instance, tailored to the health data access application form. In the current IT landscape, this may be done by deploying a virtual machine or container, either in a dedicated cluster or a cloud computing environment.

Data Reception: Depending on the interface arrangements between the SPE and the data holder(s), data can either be pulled by the SPE operator from the data holder(s) or pushed by the data holder(s) to the assigned SPE storage.

Data Upload: HDABs may use an *intermediate* SPE to prepare the data. This includes anonymisation or pseudonymisation, as well as data linkage (combining the datasets requested by the applicant at individual level). Pseudonymisation ensures that data cannot be linked to an individual without access to a separate key, which only the HDAB holds. Once the data, sometimes from several data holders, have been prepared, all of it is securely transferred from the intermediate SPE to a user-facing SPE.

Data Analysis: The health data user then processes the uploaded data within the environment using the tools provided in that SPE instance to derive the insights they are seeking.

Results Extraction: Once the analysis is complete, the health data user must request permission to download the fully anonymised aggregated results (whether partial or final) from the SPE. This step involves controlling and monitoring the data that is allowed to leave the SPE by the HDAB.

Environment Decommissioning/Archival: After the project concludes or the data permit expires, the environment may either be decommissioned (with all contents destroyed) or archived for potential future use, under new conditions, such as for reproducibility of results or new project permits.

[4] D6.2 Guideline for data users on good application and access practice https://tehdas.eu/wp-content/uploads/2025/10/d6.2-guideline-for-data-users-on-good-application-and-access-practice.pdf
[5] D7.2. Options for the services and services architecture and infrastructure for secondary use of data in the EHDS https://tehdas.eu/tehdas1/results/tehdas-proposals-for-the-implementation-of-ehds-technical-infrastructure/

Figure 5.1. The SPE functionalities as specified in the TEHDAS1 project.

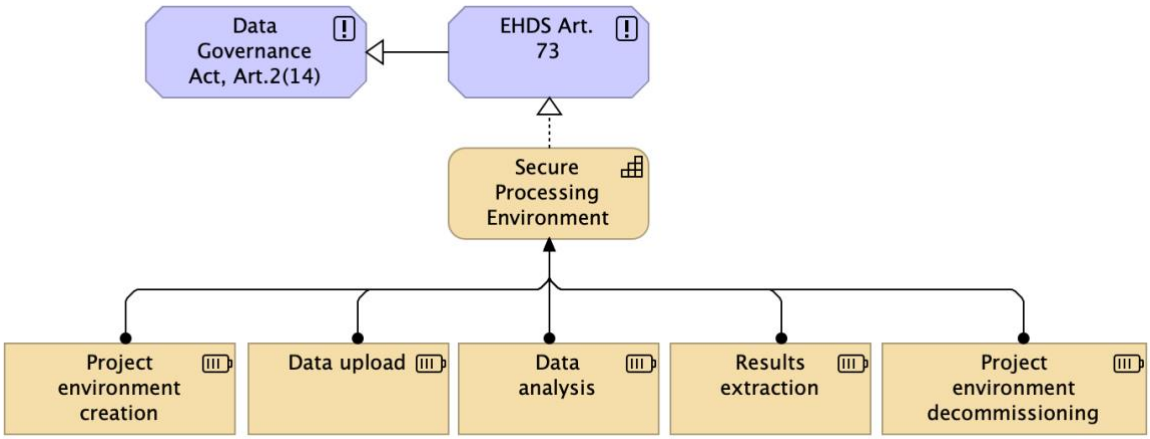

5.2 User stories

In the context of the EHDS framework, two typical user profiles can be distinguished in relation to the use of SPEs, HDAB and health data user, each associated with different roles, restrictions, and levels of responsibility.

We must differentiate the two uses of the term SPE. For the SPE operator, it is a service that must be maintained to its users. SPE is also used to refer to collective functionalities of SPE user accessible instances. For users, SPE is a project-based environment that is fully separated from other project environments within the service. To distinguish these, the full term for the latter would be **SPE project-based user space**. Any shorter versions of this should include either ‘user’ or ‘project’ to be clear what environment is meant.

The main use case is the one where the user is the health data user processing health data for secondary use according to a data permit. The whole process is overseen by an HDAB that provides the prepared sensitive data to the user, and the user is not allowed to export anything but anonymous results out of the SPE.

The second use case defines the HDAB itself as the user. An HDAB employee must do the final processing of the permitted health data that in the next stage will be handed over to the health data user’s SPE environment. The HDAB user is fully responsible for both managing the import of datasets, possibly from multiple health data holders, and export of the final permitted dataset with personal details. This use case utilises SPE solely to protect the health data and its processing from non-users. It is subject to general access control monitoring of services, but nothing more. In wider context, this matches the requirements of generic scientific research SPE protecting any sensitive data processing shown in Figure 4.5.

These two use cases reflect different trust and control models, which have implications for design of SPEs and how technical and operations measures (TOMs) are implemented. To emphasise the importance of trust, we name these **distrusting and trusting use cases**

(Figure 5.2). The challenge is to understand the implications of both use cases to the functional and technical design of SPE.

Figure 5.2. SPE trusting and distrusting use cases.

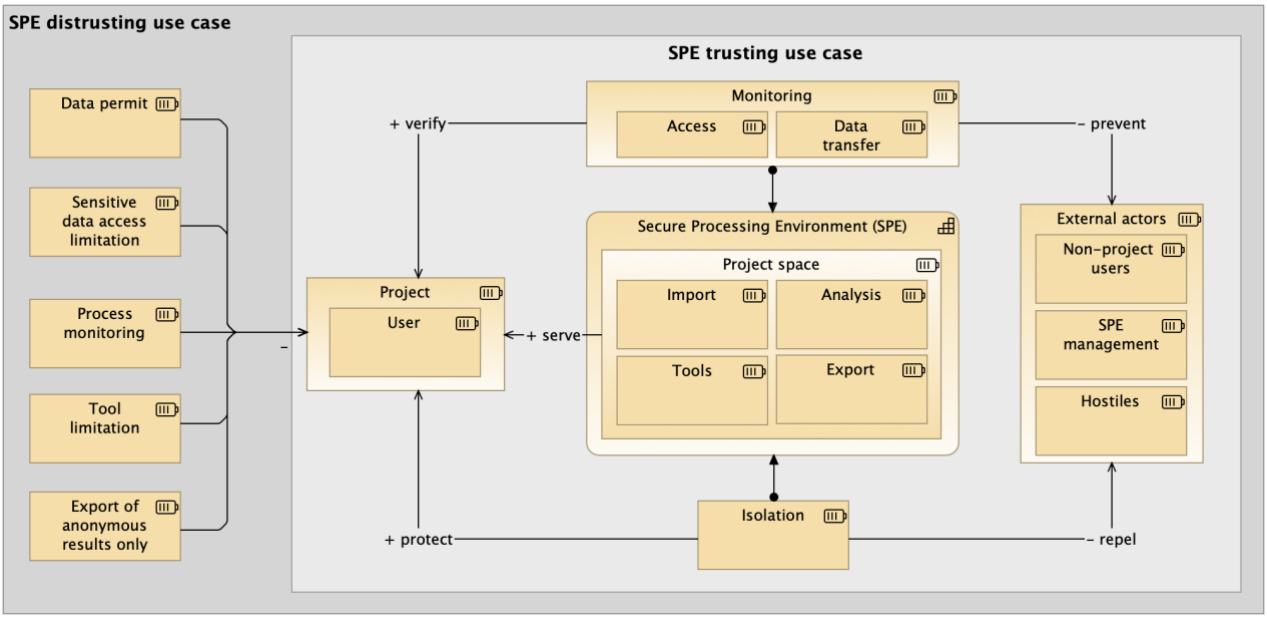

5.3 SPE requirements from EHDS Regulation

Article 73 of the EHDS Regulation lists baseline legal requirements for SPEs, serving as a foundation and minimum standard for the guidelines. Article 73(5) provides the empowerment for further implementing acts that this report aims to advise.

Article 73 requires that all processing of electronic health data for secondary use purposes must take place in a secure processing environment (SPE). These environments must have robust data protection processes, restricted unauthorised access, and prevent data from being copied or transferred unlawfully. SPE providers must comply with relevant EU laws, ensuring transparency and oversight to maintain trust and privacy.

Requirements derived from Article 73 are listed in table 5.1 and in Figure 5.3. The detailed analysis of Article 73 and derivation of its requirements is provided in Annex 9: EHDS article 73 analysis to deduce SPE requirements.

The functional and operational requirements of SPEs derived from the EHDS Regulation have by the namespace acronym ‘EHDS’ followed by the letter ‘R’. These requirements and their importance are interpretations of the EHDS Regulation.

Table 5.1: Functional and operational requirements of SPEs derived from the EHDS Regulation.

#	Requirement	Source	Responsible role	Importance
EHDSR-1	HDAB MUST grant access to EHD using a data permit	EHDS Art. 73(1)	HDAB	Mandatory
EHDSR-2	EHD MUST be accessed using an SPE	EHDS Art. 73(1)	HDAB	Mandatory
EHDSR-3	Natural persons listed in the data permit MAY access the identified EHD in SPE	EHDS Art. 73(1)(a) and (c)	HDAB, SPE Operator	Mandatory
EHDSR-4	TOMs MUST minimise the risk of unauthorised EHD access in SPEs	EHDS Art. 73(1)(b)	HDAB, SPE Operator	Mandatory
EHDSR-5	Authorised health data users MUST be strongly identified	EHDS Art. 73(1)(c) and (d)	HDAB, SPE Operator	Mandatory
EHDSR-6	All access and operation logs of SPE MUST be available for verification and auditing	EHDS Art. 73(1)(d) and (e)	SPE Operator	Mandatory
EHDSR-7	All SPE logs MUST identify the actor	EHDS Art. 73(1)(e)	SPE Operator	Mandatory
EHDSR-8	All SPE logs MUST be kept at least for one year	EHDS Art. 73(1)(e)	SPE Operator	Mandatory
EHDSR-9	TOMs of SPEs MUST be monitored for security	EHDS Art. 73(1)(f)	SPE Operator	Mandatory
EHDSR-10	EHD MUST be identified in the data permit	EHDS Art. 73(2)	HDAB	Mandatory
EHDSR-11	Health data holder MUST upload the permitted EHD to be available in an SPE for the health data user	EHDS Art. 73(2)	Data Holder	Mandatory
EHDSR-12	Health data user MAY download only non-personal EHD from SPE. Anonymised personal data is non-personal	EHDS Art. 73(2)	Data User	Mandatory
EHDSR-13	HDAB MUST ensure by reviewing that no personal data is taken out of the SPE by the health data user	EHDS Art. 73(2)	HDAB	Mandatory
EHDSR-14	Regular internal and external security audits MUST be done on SPE TOMs	EHDS Art. 73(3)	SPE Operator	Mandatory
EHDSR-15	SPE TOMs MUST undergo risk assessments	EHDS Art. 73(3)	SPE Operator	Mandatory
EHDSR-16	HDABs MUST ensure that SPE TOMs audits are carried out and that risk assessments lead to risk mitigations	EHDS Art. 73(3)	HDAB	Mandatory
EHDSR-17	When SPEs mentioned in the EU	EHDS Art. 73(4)	SPE Operator	Mandatory

	Data Act are used for EHD, EHDS rules and requirements MUST be followed			
EHDSR-18	SPEs MUST adapt to TOMs that the Commission will write into EHDS implementing acts	EHDS Art. 73(5)	SPE Operator	Mandatory

Figure 5.3. SPE requirements derived from EHDS Article 73. Requirement identifiers follow the convention 'EHDSR–n', where 'EHDS' indicates the EHDS requirements namespace and 'R' denotes requirement.

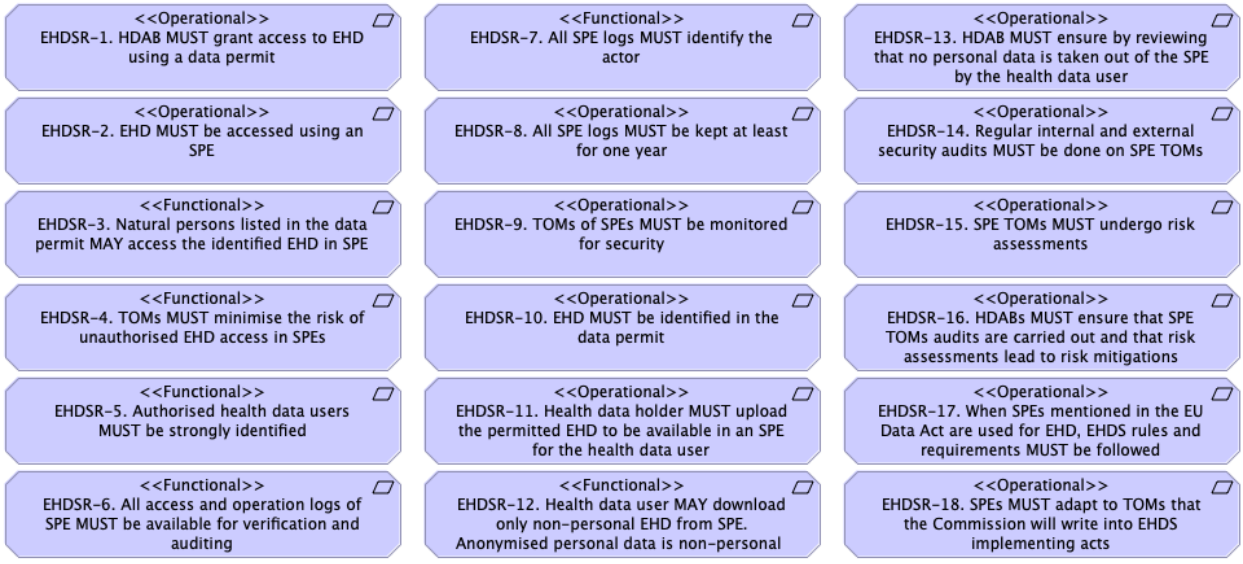

Seven of these 18 requirements are functional in a way directly affecting the functionality of SPEs under EHDS (Figure 5.4). These requirements touch either security of data processing and user identity giving specific demands (SPER-1), or logging requirements (SPER-9) already covered in more general terms in generic SPE requirements. The problematics about logging details in EHDSR-6 already seen coming from the SPE definition is covered in chapter Monitoring of SPE use in Annex 7: Design considerations and expert commentary.

Figure 5.4. EHDS SPE functional requirements. Requirement identifiers follow the convention: 'SPER–n' for general purpose SPE requirements ('SPE' namespace + 'R' for

requirement) and 'EHDSR–n' for EHDS requirements ('EHDS' namespace + 'R' for requirement).

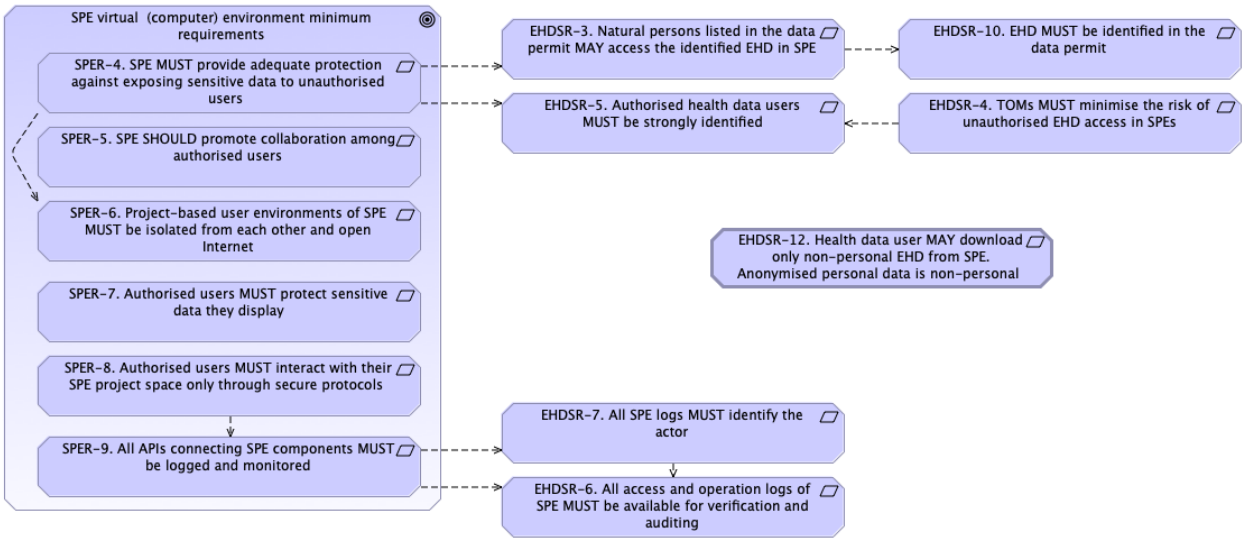

5.4 Functional requirements for EHDS SPE

5.4.1 Export control

The only functional requirement that separates SPE under EHDS from generic SPE concept is its legal obligation to health data users to only export anonymised results (EHDSR-12). This topic will be discussed in chapter Data export from SPE of Annex 7: Design considerations and expert commentary. Export control is there further separated into three main areas affected:

1. Export of anonymous results
2. Returning of clinically significant results to the original data holders
3. Creation of new datasets of enriched research data

The anonymisation of results is further discussed in TEHDAS2 *D7.2 Guideline on data minimisation, pseudonymisation, anonymisation and synthetic data.*

5.5 Operational requirements for EHDS SPE

Context: As previously explained, the goal of this deliverable is to support policy alignment, strategic planning, and high-level design of SPEs across Member States and stakeholders. As stipulated in Article 73 of the EHDS Regulation, the implementing acts will legally define the technical, organisational, information security, confidentiality, data protection, and interoperability requirements for SPEs. This chapter therefore delves into operational requirements to provide a robust basis for discussion and decision-making during the development and adoption of the EHDS implementing acts. That said, it is important to note that none of the requirements presented below are legally prescribed by this work. The exercise undertaken was to review established standards, certifications, directives, and best practices to formulate a list of recommended and mandatory operational requirements for

SPEs. In the following sections, a list of recommendations is presented for key operational areas, including the suggested level of importance (mandatory or recommended) and the rationale for each requirement. This rationale may reference the conditions set in Article 73 of the EHDS Regulation and may also be complemented by guidance from other standards and frameworks relevant to information security, compliance, and operational management.

Operational requirements in SPEs refer to the necessary processes, controls and capabilities that ensure systems handling sensitive data operate securely, reliably and in line with the functional obligations defined in Article 73 of the EHDS Regulation. These requirements cover both technical and procedural aspects, such as access management, system configuration, monitoring, incident handling and service continuity. They are primarily derived from the functional requirements laid down in the EHDS Regulation, which establishes the mandatory framework for SPE setup and operation.

However, it is important to consider that the EHDS Regulation enters a broader ecosystem of established standards, directives, and best practices that inform the design and implementation of operational controls. In developing the operational requirements presented in this section, we have analysed several state-of-the-art references to complement the EHDS provisions and provide practical guidance for SPE operators:

- **ISO/IEC 27000 series**[6]:Internationally recognised standards for information security management, providing a structured approach for establishing, implementing, and maintaining an Information Security Management System (ISMS). While ISO 27001/27002 are not explicitly mandated by the EHDS Regulation, they offer widely accepted guidance for implementing robust security processes and controls.

- **NIS2 Directive (Directive (EU) 2022/2555):** A European directive establishing cybersecurity risk management obligations. Importantly, NIS2 obligations may apply to SPE operators depending on their organisational context and national transposition. NIS2 provides guidance on areas such as incident response, access control, supply chain security, and business continuity. Its obligations serve as a reference for state-of-the-art operational practices, but compliance depends on the scope defined by national legislation.

- **FitSM:** A practical framework for IT service management that supports structured service governance, clear role definitions, and disciplined operational processes. FitSM is not a mandatory requirement; rather it can be adapted to an organisation's needs and resources. Implementation can be approached in phases, starting with core processes for operational stability, followed by scalability-focused processes, and culminating in continual improvement to enhance service performance and maturity. FitSM is proposed here as a helpful approach to unify governance and operational practices across SPE providers.

- **General Data Protection Regulation (GDPR):** Establishes the legal framework for data protection and privacy of natural persons, including requirements for data processing accountability, security, confidentiality, and breach reporting. In the SPE context, GDPR provides crucial guidance for handling personal health data, defining principles for lawful

[6] ISO27001 https://www.iso.org/standard/27001

processing, auditability, and data minimisation. Compliance with GDPR ensures that SPE operations protect the privacy rights of individuals while supporting lawful secondary use of health data.

By drawing on these references, SPE operators can better understand the practices, protocols, and controls that support secure and resilient operation. Note that this report is **not intended to be a standalone implementation manual**. It does not prescribe specific products, detailed configurations, or deployment procedures, and it cannot by itself guarantee operational compliance or certification of SPEs.

We will formulate SPE operational requirements under six categories (Figure 5.5):

- Main SPE roles
- SPE setup and access management
- SPE auditing, compliance and reporting
- Monitoring and incident management
- Risk management and mitigation
- Maintenance and support

Figure 5.5. SPE operational requirement categories

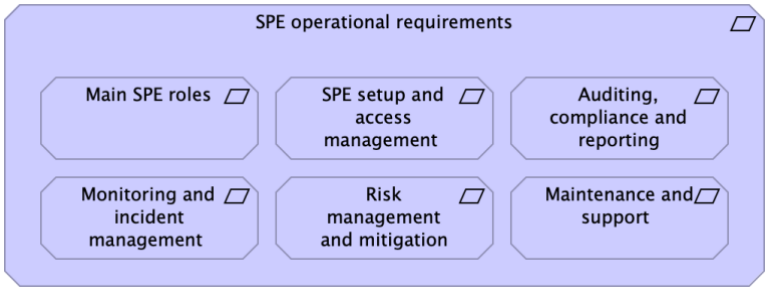

We will mark these requirements in relation to the EHDS Regulation, its identified functional requirements, and to the four aforementioned regulations.

Main SPE roles

In the context of the EHDS, operational requirements are shaped by a framework of clearly defined roles that collectively ensure the secure, lawful, and ethical secondary use of electronic health data.

The core roles - Health Data Access Body, SPE Operator, Data Holder, and Health Data User - each carry specific responsibilities that influence or depend on robust operational processes. These roles do not operate in isolation but rather in close interaction through structured processes, mutual dependencies and shared responsibilities that ensure compliance with legal standards, while also enabling technical interoperability and data protection. Their relation to EHDS SPE requirements is shown in Figure 5.6.

Figure 5.6. The main SPE roles and their relation to identified EHDS SPE requirements. Purely functional requirements have been greyed down. Requirement identifiers follow the

convention ‘EHDSR–n’, where ‘EHDS’ indicates the EHDS requirements namespace and ‘R’ denotes requirement.

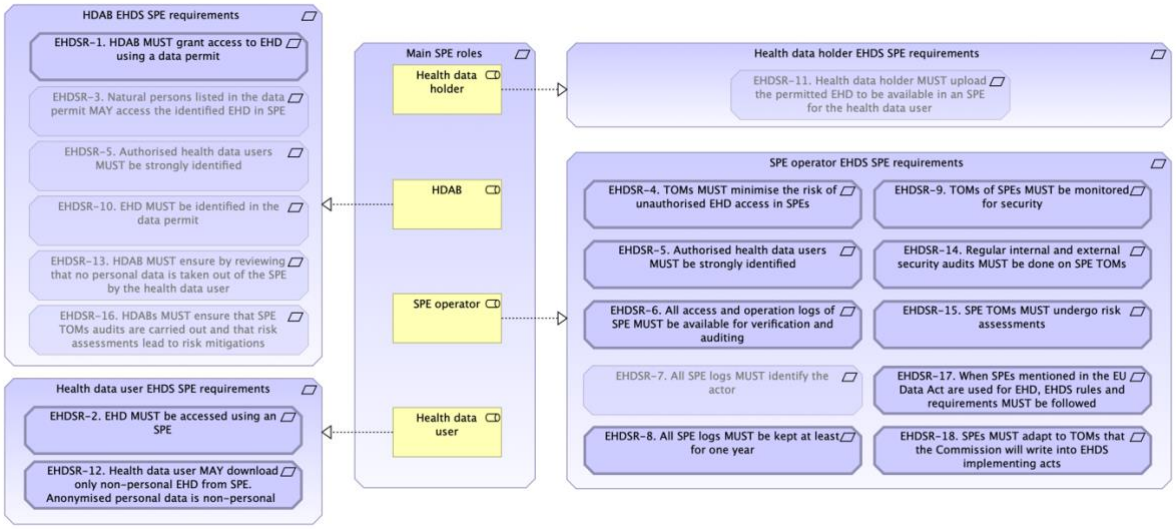

HDAB: The HDAB is the national authority designated under the EHDS Regulation to govern and oversee access to electronic health data for secondary use. In the context of SPEs, the HDAB plays a central regulatory and supervisory role to ensure that access to, processing of, and outputs derived from electronic health data comply with the conditions set out in the data permit and with applicable legal, security, and confidentiality requirements. In particular, the HDAB is responsible for ensuring that electronic health data made available by health data holders can be uploaded to, and accessed by health data users within a SPE, in the format and under the conditions specified in the data permit. The HDAB oversees that access within the SPE is limited to health data users and aligned with the approved purposes and safeguards defined in the permit. The HDAB is also responsible for reviewing electronic health data included in download requests from the SPE, in order to ensure that health data users are only able to extract non-personal electronic health data, including data in anonymised statistical form, in accordance with the EHDS Regulation and the conditions of the data permit. Furthermore, the HDAB must ensure that audits of SPEs are carried out on a regular basis, including audits performed by third parties. Where audits identify shortcomings, risks, or vulnerabilities in the SPE, the HDAB is required to ensure that appropriate corrective actions are taken. As part of its transparency obligations, each HDAB must publish a publicly available activity report every two years, including information on audits conducted on health data users and on the compliance of secure processing environments with the applicable standards, specifications, and requirements. Beyond these core regulatory responsibilities, HDABs may in practice act as a central repository of expertise on the secondary use of health data and serve as a primary point of contact for stakeholders. As a matter of good practice, HDABs may provide guidance on lawful and appropriate data use, deliver training on the handling of sensitive health data, and operate helpdesk or support functions to address enquiries related to data access, use, and compliance. These supporting activities, while not mandated by the EHDS Regulation, can contribute to consistent interpretation of requirements and to the effective and compliant use of secure processing environments.

SPE Operator: The SPE Operator refers to the organisational function responsible for the day-to-day operation, maintenance, and technical management of the SPE. While the EHDS Regulation does not explicitly define "SPE Operator" as a distinct legal role, this concept is introduced in this work to address the practical need for a clearly identified entity or function accountable for operating and maintaining the SPE in accordance with the requirements set out in Article 73 of the EHDS Regulation. In practice, the SPE Operator function may be fulfilled by the HDAB itself or by another organisation entrusted with operating the SPE. Regardless of the organisational setup, the allocation of tasks, responsibilities, and accountability related to SPE operations should be formally defined and documented to ensure clarity, traceability, and effective governance. The SPE Operator should be responsible for ensuring that the SPE operates securely, reliably, and continuously, that technical and organisational measures for access control, logging, monitoring, incident management, and service continuity are implemented and enforced, and that processing conditions defined in data permits, such as access restrictions, separation of processing contexts, and termination of processing upon permit expiry, are supported. Furthermore, the SPE Operator must ensure that the SPE remains compliant with applicable security, confidentiality, and data protection requirements relevant to its operation. It is important to note that, in the context of this work, the responsibilities of the SPE Operator are related and limited to the operation and management of the SPE itself and do not replace or override the regulatory responsibilities of the HDAB or the legal obligations of health data users. A clear separation and coordination of responsibilities between the SPE Operator, the HDAB, and health data users, including research leads, is essential to avoid ambiguity, duplication, or gaps in accountability. This work emphasises the importance of explicitly defining responsibility boundaries between the HDAB, the SPE Operator, and health data users, formally assigning and documenting operational responsibilities and delegated tasks related to SPE operation, and providing further guidance, at the implementation level, on governance models for establishing and operating SPEs, including scenarios where SPE services are provided by third parties. By introducing the SPE Operator as an operational role, this document aims to support clarity in responsibility allocation and to facilitate consistent implementation of the EHDS requirements across Member States, without prescribing a specific organisational or contractual model.

Data Holder: The Data Holder is the organisation or entity that possesses the original health datasets, such as hospitals, research institutions or health registries. Health Data Holders are responsible for providing electronic health data in accordance with the conditions set out in the data permit and under appropriate legal, technical, and security safeguards. In the context of the SPEs, Health Data Holders make the requested data available for analysis within the SPE, ensuring that the data is provided in the format specified by the data permit and can be accessed securely by authorised health data users.

Health Data user: The Health Data User is a natural or legal person, including Union institutions, bodies, offices, or agencies, that has been granted lawful access to electronic health data for secondary use pursuant to a data permit or a health data request approval. Health Data Users are authorised to access and process electronic health data only in accordance with the conditions set out in the data permit. When working within an SPE, Health Data Users must not provide access to the data, nor make it available, to any third party not explicitly listed in the data permit. Access rights are strictly tied to the authorised

users identified in the data permit, ensuring that only those individuals can interact with the health data in the SPE.

Health data user is here defined from the EHDS perspective as the entity that carries the responsibility for the approved health data. Natural persons mentioned in the health data permit are authorised users. For researchers, the situation is more complicated. To clarify it, we need to specify the correct **project context for their roles (**Figure 5.7).

A **research project** may consist of many members with different roles, and not all of them are accessing sensitive data specified by the data permit. To get the data permit, the principal investigator (or one of the many) acts as the health data applicant to submit a data application. A positive decision creates the data permit that includes identifiers to approved **EHDS project** authorised users, of which the principal investigator is one. The data permit gives access to applied data and to at least one SPE service. Based on the data permit, the principal investigator establishes a project in the SPE to start analysing the data and includes other authorised users to it.

Figure 5.7. Illustration of the use of the role name principal investigator in three different project contexts (work package ArchiMate elements). The principal investigators mean different things to the research group, the HDAB and SPE service.

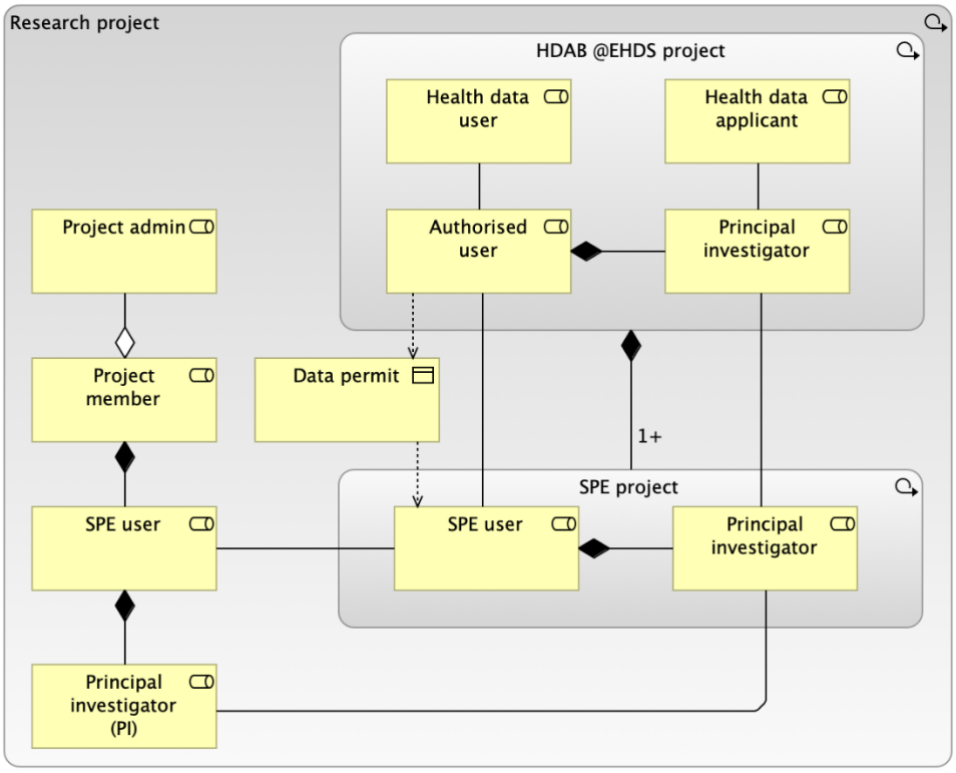

SPE setup and access management

An SPE that is configured in accordance with the conditions set out in a data permit (Article 73, EHDSR-1). The data permit defines, among other elements, the purpose of processing (EHDSR-2), the data that will be accessed (EHDSR-10), the authorised health data users (EHDSR-3), and the specific SPE in which the processing is allowed. The Regulation further requires that SPEs be set up in a timely manner, so that HDABs are able to make health data available to the health data user within two months of receiving the data from the health data holder. This implies that the SPE must be operational within this timeframe and configured in line with the terms of the data permit. In addition, the Regulation requires that access to the SPE be restricted to authorised natural persons listed in the data permit, that appropriate technical and organisational measures be applied to prevent unauthorised access, and that the SPE be terminated after the expiry of the data permit, with the corresponding deletion or rendering unrecoverable of electronic health data within the legally defined timeframe.

To give effect to these regulatory obligations, SPEs must rely on a set of core operational capabilities related to environment setup and access management. These include identity and access management mechanisms that ensure unique, auditable user identities, procedures for granting, modifying, and revoking access rights in line with data permit conditions, credential lifecycle management processes and controls for managing privileged access to the underlying infrastructure. In addition, SPE operators require service and configuration management practices that allow them to maintain an accurate overview of the services, systems, and components that form part of the SPE, and to ensure that each environment is configured and operated consistently with its authorised purpose.

Defining operational requirements in this area presents several challenges. Access rights and system configurations are not static but are tightly coupled to administrative decisions taken by HDABs through data permits, which may change over time. SPE operators must therefore be able to reflect permit updates promptly and accurately in access controls and environment configurations. At the same time, the need to enforce strict access restrictions must be balanced with operational requirements for logging, auditability, incident response, and, where applicable, legal obligations related to traceability or scientific reproducibility. The termination and deletion of SPE environments after permit expiry further requires carefully designed procedures to ensure confidentiality and compliance, while allowing for exceptions mandated by applicable law.

The operational requirements set out below provide a structured and practical basis for addressing these challenges. They describe recommended operational requirements for SPE setup and access management, indicating where requirements are considered mandatory or recommended, and explaining how they support compliance with the EHDS Regulation while enabling secure and effective operation of SPEs.

Table 5.2. lists these requirements and their derivation that is shown in figure 5.8.

Figure 5.8. SPE setup and access management. Requirement identifiers follow the convention: ‘SPER–n’ for general purpose SPE requirements (‘SPE’ namespace + ‘R’ for requirement), ‘EHDSR–n’ for EHDS requirements (‘EHDS’ namespace + ‘R’ for

requirement) and ‘OPR–n’ for operational requirements (‘OP’ namespace + ‘R’ for requirement). OP requirements and their importance are interpretations of the EHDS Regulation.

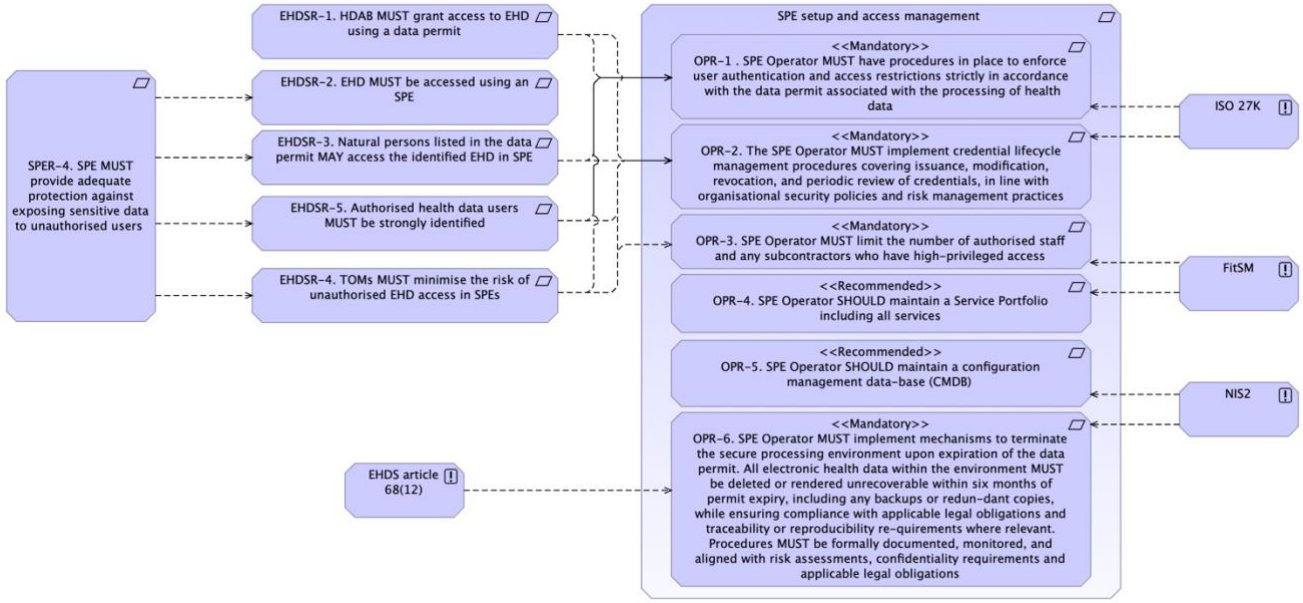

Table 5.2: SPE setup and access management

#	Link	Area	Requirement	Rationale	Importance
OPR-1	EHDSR-1 EHDSR-3 EHDSR-4 EHDSR-5	Access management	**SPE Operator MUST have procedures in place to enforce user authentication and access restrictions** strictly in accordance with **the data permit associated with the processing of health data.** These procedures shall ensure that only authorised natural persons listed explicitly in the data permit can access the SPE. Access rights shall be derived from, limited to, and continuously aligned with the conditions of the data permit, and shall be reviewed and adjusted as required to reflect permit changes. User identities shall be unique, persistent, auditable, and non-	Article 73(1)(a) - Restriction of SPE access to authorised natural persons listed in the data permit Article 73 (1)(d) - Use of individual and unique user identities and confidential access modes to assure health data users have access only to the electronic health data covered by their data permit. ISO/IEC 27001 and ISO/IEC 27002 - Best practices on authentication and access control	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			transferrable. Access control mechanisms, including role-based access control, shall be implemented to ensure that users have only the minimum level of access necessary to perform their authorised tasks under the data permit. Where the same natural person is authorised under multiple concurrent data permits, the SPE Operator MUST ensure logical separation of access rights and processing contexts per data permit, in order to prevent unauthorised cross-access between permits.		
OPR-2	EHDSR-1 EHDSR-3 EHDSR-4 EHDSR-5		**The SPE Operator MUST implement credential lifecycle management procedures covering issuance, modification, revocation, and periodic review of credentials, in line with organisational security policies and risk management practices.** If an alias or user ID is used, it must be uniquely attributable to a single individual within the organisation's identity management system. The true identity behind any alias must be known and auditable at all times.	ISO/IEC 27001 and ISO/IEC 27002 and EHDSR principles require secure management of credentials to protect sensitive health data, ensuring accountability and traceability of all actions within the SPE. Proper identity management prevents unauthorised access and supports audit readiness.	Mandatory
OPR-3	EHDSR-4	Access management	**SPE Operator MUST limit the number of authorised staff and any subcontractors who have high-privileged access enabling them to access or process health data and MUST implement effective procedures for managing and monitoring such**	Article 73(1)(c) - Restrict access to electronic health data in the SPE (input, inspection, modification, and deletion) to a limited number of authorised and identifiable individuals.	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			access within the SPE infrastructure. Privileged access rights shall be reviewed at regular intervals and immediately revoked when no longer necessary. Exceptionally, administrators may obtain privileged access for technical troubleshooting or incident response, provided this is explicitly authorised, time-limited, and subject to enhanced monitoring and documentation.	ISO/IEC 27002 recommendation - organisations must strictly limit and manage privileged access by assigning it only when necessary, regularly reviewing permissions, and logging all privileged actions. Elevated access must be controlled: authorised temporarily, detailed in logs, and promptly reviewed for necessity and accountability.	
OPR-4		Set up	**SPE Operator SHOULD maintain a Service Portfolio including all services** (e.g., data ingestion, analysis platforms, audit tools). Include: Descriptions Availability Security levels Eligible user groups (e.g., researchers, data owners) Publish the catalogue internally and regularly update it.	FitSM Service Portfolio recommendation	Recommended
OPR-5		Set up	**SPE Operator SHOULD maintain a configuration management database (CMDB)** that includes: - All virtual machines, services, APIs, and databases - Tag systems based on data classification (e.g., sensitive, anonymised) The CMDB SHOULD be kept up to date and reflect the current operational state of the SPE.	FitSM Configuration Management recommendation	Recommended
OPR--6	EHDSR-1	Set up	**SPE Operator MUST implement**	Article 68(12) - electronic health	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			mechanisms to terminate the secure processing environment upon expiration of the data permit. All electronic health data within the environment MUST be deleted or rendered unrecoverable within six months of permit expiry, including any backups or redundant copies, while ensuring compliance with applicable legal obligations and traceability or reproducibility requirements where relevant. Procedures MUST be formally documented, monitored, and aligned with risk assessments, confidentiality requirements and applicable legal obligations.	data in the SPE must be deleted within six months of the data permit's expiry. NIS2 emphasises secure data disposal as part of business continuity and disaster recovery planning, including managing backups and redundant data to prevent unauthorised access or data leaks. Traceability, reproducibility, and other legal obligations (AI Act, Data Act) may require certain processed results or outputs to be retained in accordance with applicable laws. The deletion and retention policy should balance these obligations with the mandate to securely terminate the SPE.	

SPE auditing, compliance and reporting

Technical and organisational security measures

The EHDS Regulation establishes that SPEs must operate under robust oversight mechanisms to ensure compliance with legal, organisational, and security obligations. In particular, Article 73 requires that SPEs have in place appropriate technical and organisational measures to minimise security risks, ensure confidentiality and integrity of electronic health data, and enable verification and auditing of processing operations. HDABs are required to ensure that audits of SPEs are carried out on a regular basis, including by third parties, and that any identified shortcomings, risks, or vulnerabilities are addressed through corrective actions. The Regulation further requires the keeping of identifiable logs of access to and activities within the SPEs for the period necessary to verify and audit all processing operations.

To support these obligations in practice, SPEs require a coherent set of operational capabilities related to auditing, compliance, and reporting. These include mechanisms for conducting regular internal and external audits, processes for retaining and managing logs and access records that enable traceability of key activities and the maintenance of accurate and up-to-date documentation covering technical configurations, organisational procedures, and security controls. In addition, effective compliance relies on clearly assigned responsibilities, including dedicated and appropriately trained personnel who oversee adherence to legal, organisational, and technical obligations, act as points of contact for audits, and coordinate responses to identified risks or non-compliance. Structured management systems, such as information security and service management systems, further provide a systematic framework for implementing, monitoring, and continuously improving these controls, including supplier and change management practices.

Defining operational requirements in this area also presents diverse challenges. Audit and compliance mechanisms must be risk-based and proportionate, considering the size, complexity, and risk profile of the SPE, while remaining sufficiently robust to meet regulatory expectations. Governance arrangements between HDABs, SPE operators, and other involved actors may differ across Member States and organisational contexts, requiring flexibility in how responsibilities are assigned and exercised. A particularly significant challenge lies in balancing operational feasibility with privacy and data protection considerations, especially with regard to logging and traceability. Logging practices must be detailed enough to enable effective audits and investigations, yet carefully designed to avoid excessive data collection, unnecessary intrusion into user activities, or the creation of unmanageable volumes of log data. Ensuring practical, proportionate, and privacy-respecting traceability is therefore a central consideration when defining operational controls in this domain.

The operational requirements set out below provide a structured description of recommended practices for SPE auditing, compliance, and reporting. They specify mandatory and recommended requirements aimed at supporting compliance with the EHDS Regulation, while aligning with state-of-the-art standards and good practices for secure and accountable operation of SPEs.

To ensure that health data is handled securely, SPEs must follow strict technical and organisational security measures. These measures help protect sensitive data, prevent unauthorised access, and comply with legal obligations under the EHDS Regulation and the NIS2 Directive.

The EHDS Regulation, particularly Article 73, requires that SPEs minimise security risks, monitor compliance, and undergo regular audits to identify and fix vulnerabilities. At the same time, the NIS2 Directive sets additional cybersecurity rules, including risk assessments, encryption policies, incident response plans, and security checks for suppliers.

One key requirement is that SPE Operators implement an ISMS to establish and maintain security policies aligned with applicable laws, such as the GDPR. The ISMS should also be complemented with backup and disaster recovery plans, employee cybersecurity training and use of encryption to protect data and security assessment of suppliers and external systems. By following these security measures, SPEs can provide a safe and reliable environment for processing health data while ensuring compliance with EU regulations.

Table 5.3 lists these requirements and their derivation that is shown in Figure 5.9.

Figure 5.9. SPE auditing, compliance and reporting. Requirement identifiers follow the convention: 'EHDSR–n' for EHDS requirements ('EHDS' namespace + 'R' for requirement) and 'OPR–n' for operational requirements ('OP' namespace + 'R' for requirement). OP requirements and their importance are interpretations of the EHDS Regulation.

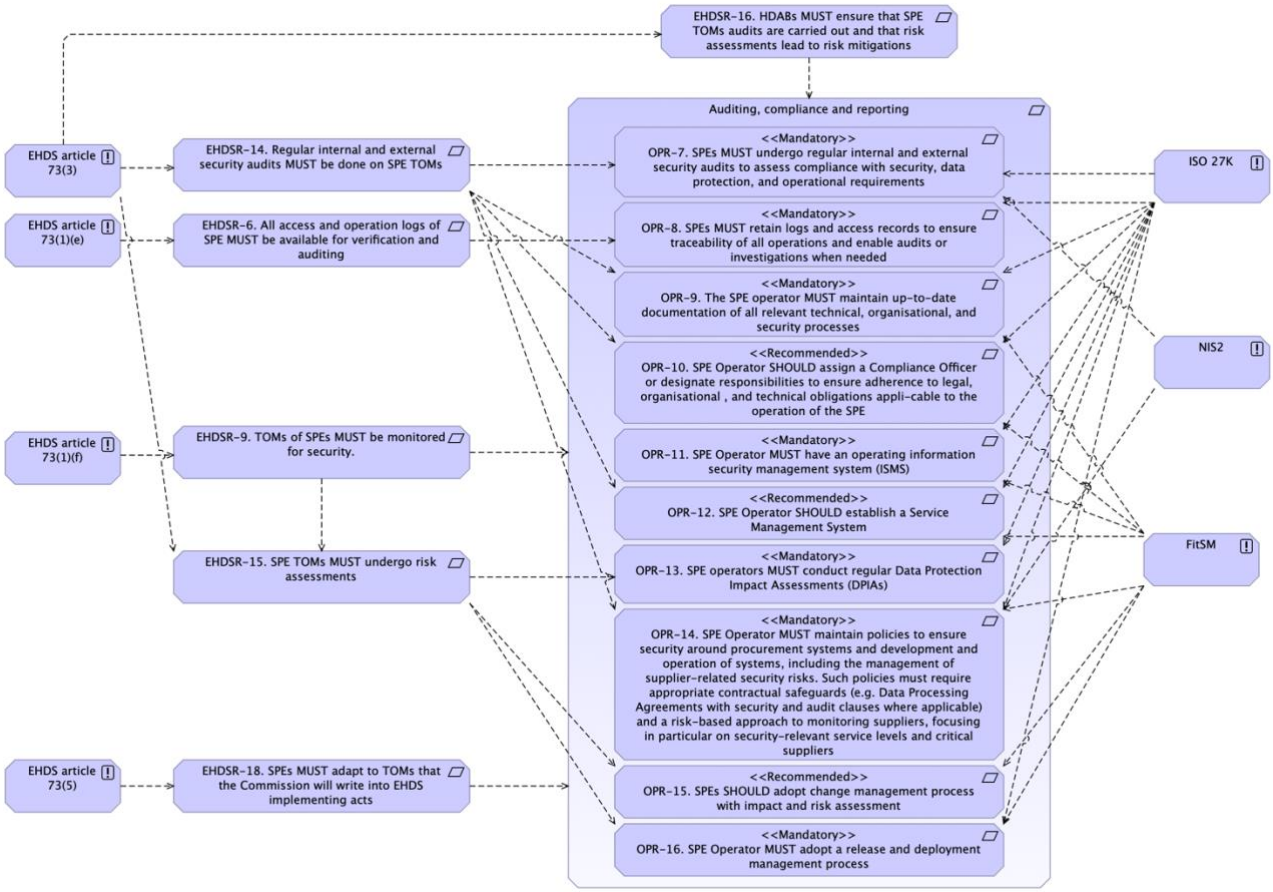

Table 5.3: SPE auditing, compliance and reporting

#	Link	Area	Requirement	Rationale	Importance
OPR-7	EHDSR-14	Auditing	**SPE Operator MUST undergo regular internal and external audits to assess compliance with security, data protection,**	EHDS Regulation (Article 73(3)) explicitly requires regular audits by internal or third-party entities.	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			and operational requirements. Audits must cover essential aspects including risk assessment, planning, execution, corrective actions, and documentation, and must be conducted in a risk-based and proportionate manner, considering the SPE's size, complexity, and risk profile. Audit arrangements, including scope and frequency, must be defined in agreement with the relevant HDAB. At minimum: • **Internal audits** shall include a full audit of the SPE's ISMS and operational controls at least once per year. Additional internal audits must be conducted when significant changes occur to the SPE's services, infrastructure, or risk profile • **External audits** must be conducted at least once every three years. Additional external audits must be carried out following significant structural or security-relevant changes to the SPE, or after a serious security incident. Internal audits must be conducted by personnel independent from the audited activities, and external audits must involve third parties with appropriate competence or accreditation.	NIS2 Directive mandates the evaluation of risk-management measures and corrective actions. ISO/IEC 27001 and ISO/IEC 27002 require planned internal audits to verify the ISMS and other controls. In common practice, organisations aligned with ISO/IEC 27001 conduct internal audits at least annually. External certification under ISO/IEC 27001 follows a three-year cycle, with a full recertification audit every three years. To maintain certification, accredited bodies conduct surveillance audits at least once per year during the first and second years of the certification cycle. These surveillance audits are external, risk-based, and narrower in scope than the three-year recertification audit. A risk-based and proportionate audit approach, aligned with governance arrangements agreed with the HDAB, supports consistent implementation across SPEs of different sizes and complexity while ensuring audit independence, accountability, and operational feasibility.	
OPR-8	EHDSR-6	Auditing	**SPE Operator MUST retain logs and access records to ensure traceability of all operations and enable audits or investigations when needed.** The logging should be focused on key actions, including data access, processing, and administrative activities, while balancing operational feasibility and privacy considerations.	EHDS Regulation (Article 73(1)(e)) obliges SPEs to ensure logs of data access and operations/activities in the SPE. ISO/IEC 27002 defines how logs and access records should be retained and reviewed: • Logs should capture key details such as user IDs, timestamps, and the type of activity performed. They must be protected against tampering, stored in a time-synchronised environment, and regularly reviewed to detect anomalies or	Mandatory

#	Link	Area	Requirement	Rationale	Importance
				unauthorised access. Logs access should be restricted to authorised personnel. GDPR Article32(1)(b) requires measures ensuring the ongoing confidentiality, integrity, availability, and resilience of processing systems, for which audit trails are a key control. Article 5(2) establishes the accountability principle, requiring controllers and processors to demonstrate compliance with data protection principles, which operationally relies on reliable logs. Article 30 requires records of processing activities, which are supported by structured logging of system usage and access to personal data.	
OPR-9	EHDSR-14	Compliance/ Reporting	**SPE Operator MUST maintain up-to-date documentation of all relevant technical, organisational, and security processes.**	ISO/IEC 27001 and ISO/IEC 27002 mandate control documentation and maintenance of procedures and records. FitSM supports structured documentation as a basis for consistent compliance and audit-readiness.	Mandatory
OPR-10	EHDSR-14	Compliance	**SPE Operator SHOULD assign a Compliance Officer or designate responsibilities to ensure adherence to legal, organisational, and technical obligations applicable to the operation of the SPE.** This role may be supported by other functions, such as Privacy Officer or, where applicable under the GDPR, a DPO (Data Protection Officer), and servesas a point of contact for audits, compliance, and regulatory matters related to SPE operations.	ISO/IEC 27002 recommends that roles and responsibilities related to security and compliance be clearly assigned and communicated. FitSM recommends the assignment of roles such as Information Security Manager or Compliance Officer to oversee governance. This ensures accountability and provides a single point of contact for audit, compliance, and regulatory matters.	Recommended
OPR-11	EHDSR-15 EHDSR-16	Compliance	**SPE Operator MUST have an operating information security management system (ISMS)** in line with the state of the art, for example following the requirements of ISO/IEC 27001. Following the recommendations from the FitSM standard, this should include: - End-to-end encryption	To have an ISMS in place is already a requirement of other laws (including GDPR). FitSM - Information Security Management requirement ISO/IEC 27001 and ISO/IEC 27002 requirement	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			- Role based access control - Multifactor authentication - Vulnerability assessments and penetration tracking - Document and update security controls		
OPR-12	EHDSR-14	Compliance	**SPE Operator SHOULD establish a Service Management System**: Define the scope of the SMS to include all services and components involved in processing health data. Appoint a Service Management Officer (SMO) responsible for maintaining SMS compliance.	FitSM SMS requirement ISO/IEC 27002: Maintain a real-time inventory of assets used to process health data	Recommended
OPR-13	EHDSR-14	Auditing	**SPE Operator MUST conduct regular Data Protection Impact Assessments (DPIAs)**	ISO/IEC 27002	Mandatory
OPR-14	EHDSR-14	Compliance	**SPE Operator MUST maintain policies to ensure security around procurement systems and development and operation of systems, including the management of supplier-related security risks. Such policies must require appropriate contractual safeguards (e.g. Data Processing Agreements with security and audit clauses where applicable) and a risk-based approach to monitoring suppliers, focusing in particular on security-relevant service levels and critical suppliers.**	This is a NIS2 requirement to ensure security around supply chains and the relationship between the company and the direct supplier. Companies must choose security measures that fit the vulnerabilities of each direct supplier. And then companies must assess the overall security level for all suppliers. FitSM requirements - Require Data Processing Agreements (DPAs) with security and audit clauses. - Monitor SLA compliance of suppliers, particularly for critical infrastructure - ISO/IEC 27002 – security clause within procurement contracts	Mandatory
	EHDSR-18	Compliance	SPE Operator MUST adapt to TOMs that the Commission will write into EHDS implementing acts.	Article 73(5)	Mandatory
OPR-15	EHDSR-15	TOM	**SPE Operator SHOULD adopt change management process with impact and risk assessment**	FitSM requirements	Recommended
OPR-16	EHDSR-15	TOM	**SPE Operator MUST adopt a release and deployment management process.** These may include staging environments, rollback plans and automated deployment pipelines (CI/CD). The	FitSM requires that changes to services and service components are planned, tested, approved, and documented in order to minimise the risk of service disruption and security incidents.	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			release and deployment process MUST include controls to ensure software integrity and secure development practices throughout the software lifecycle, including measures such as controlled code changes, verification prior to deployment, and prevention of unauthorised or unreviewed modifications to production systems.	ISO/IEC 27001 and ISO/IEC 27002 require organisations to control changes to information systems and to ensure that software development, release, and deployment processes protect the confidentiality, integrity, and availability of information assets. This includes safeguarding software integrity and preventing the introduction of vulnerabilities through uncontrolled or insecure development and deployment practices. Given the sensitivity of electronic health data processed within an SPE, robust release, deployment, and secure development lifecycle controls are necessary to reduce operational and security risks and to support compliance with EHDS requirements for SPEs.	

Monitoring and incident management

The EHDS Regulation requires that SPEs enable effective monitoring of processing activities and provide mechanisms to detect, manage, and respond to security incidents. Article 73 establishes that SPEs must retain, for a minimum of one-year, identifiable logs of access to and processing of electronic health data, while ensuring confidentiality and integrity of those logs (Article 73(1)(e)). The Regulation further implies the need for prompt intervention in cases of misuse or security breaches, including the ability to restrict or halt access to the environment when necessary to protect data and prevent further harm. These obligations form the legal basis for operational monitoring, incident detection, and response capabilities within SPEs.

To fulfil these requirements in practice, SPEs must be equipped with operational functionalities that support continuous monitoring and effective incident management. This includes the systematic recording and retention of security- and compliance-relevant logs, covering user authentication, data access and movement, and administrative actions. Secure storage mechanisms are required to protect logs against unauthorised access or tampering and to ensure their availability for audits and investigations. In addition, SPEs must establish clear incident reporting processes to notify HDABs and, where applicable, other relevant authorities of security incidents or non-compliance findings, including data breaches or misuse. Effective incident management further requires the technical and organisational capability to promptly suspend or restrict access and processing activities within the SPE, as well as robust backup and disaster recovery procedures to restore service availability and data integrity following an incident.

Defining operational requirements in this area involves addressing several important challenges. A key consideration is ensuring that logging practices are proportionate, risk-based, and technically feasible. Excessive or overly detailed logging may be operationally burdensome, privacy-invasive, and disproportionate to the risk reduction achieved, while insufficient logging may undermine auditability and incident detection. Therefore, logs must be carefully scoped to capture information necessary for accountability and misuse detection, without collecting unnecessary detail on user behaviour. At the same time, logged information must be adequately protected against unauthorised access or manipulation. Incident reporting timelines and response procedures must also balance urgency with accuracy, ensuring timely notification without compromising the quality of incident analysis. These challenges have been explicitly taken into account in the formulation of the operational requirements below, which draw on established standards and best practices, including those reflected in ISO/IEC standards and incident reporting frameworks referenced in instruments such as the NIS2 Directive, to define proportionate, privacy-respecting, and operationally realistic requirements.

The following operational requirements provide a structured description of recommended practices for monitoring and incident management in SPEs. They specify mandatory and recommended measures intended to support effective oversight, timely incident response, and resilient operation of SPEs in line with the EHDS Regulation. Table 5.4 lists these requirements and their derivation that is shown in Figure 5.10.

Figure 5.10. Monitoring and incident management. Requirement identifiers follow the convention: 'EHDSR–n' for EHDS requirements ('EHDS' namespace + 'R' for requirement) and 'OPR–n' for operational requirements ('OP' namespace + 'R' for requirement). OP requirements and their importance are interpretations of the EHDS Regulation.

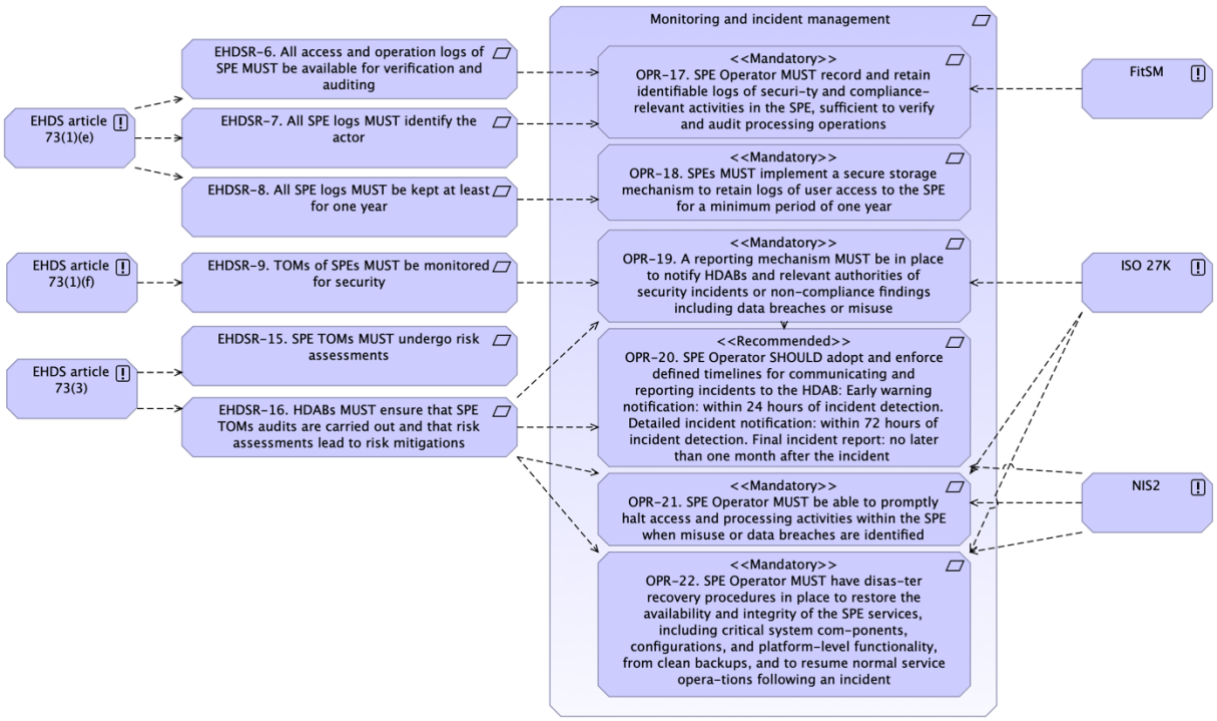

Table 5.4: Detailed operational requirements for monitoring and incident management

#	Link	Area	Requirement	Rationale	Importance
OPR-17	EHDSR-6 EHDSR-7	Monitoring	**SPE Operator MUST record and retain identifiable logs of security and compliance-relevant activities in the SPE, sufficient to verify and audit processing operations. Logging must, at a minimum, cover:** • **User authentication and access to electronic health data** • **Data import, export, or release of outputs** • **Administrative actions affecting access rights or data availability.** Logging MUST be proportionate, risk-based, and technically feasible, and MUST avoid unnecessary collection of detailed user behaviour not required for auditability or misuse detection. Logged information MUST be protected against unauthorised access or tampering and retained for the period necessary to support audits and investigations in accordance with applicable law.	Article 73(1)(e) - the keeping of identifiable logs of access to and activities in the secure processing environment for the period necessary to verify and audit all processing operations in that environment. Proportionate, risk-based logging of security- and compliance-relevant events is consistent with recognised information security practices and avoids unnecessary operational burden, excessive data generation, or privacy risks. FitSM guidance recommends that recommends that track incidents and requests should be logged using ITSM tools (e.g., Jira Service Management, Freshservice), classify and prioritise events based on data sensitivity and service impact, provide a clear user-facing incident reporting process, and maintain a 24/7 on-call rota for critical service issues. These practices support effective incident handling and oversight without prescribing specific technical implementations.	Mandatory
OPR-18	EHDSR-8	Monitoring	**SPE Operator MUST implement a secure storage process to retain logs of user access to the SPE for a minimum period of one year**. Log storage must ensure integrity, confidentiality, and protection against unauthorised access or tampering. While one year is the minimum retention period, longer retention (5 years) is recommended to support effective auditing procedures.	Article 73(1)(e) - logs of access shall be kept for at least one year	Mandatory
OPR-19	EHDSR-9 EHDSR-16	Incident management	**A reporting process MUST be in place to notify HDABs and relevant authorities of security incidents or non-compliance findings including data breaches or misuse. The reporting process must define reporting responsibilities, notification timelines,**	EHDS Regulation (Art. 73(3)) requires that SPEs allow for compliance checks and audits by competent authorities to notify the HDAB of any incidents affecting the integrity or confidentiality of data.	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			escalation paths, and required documentation	ISO/IEC 27002 requirement: Incident reporting procedures should define what must be reported, by whom, how quickly, and through which communication channels. This includes the escalation path, responsibilities, and necessary documentation to ensure timely and effective response.	
OPR-20	EHDSR-16	Incident management	**SPE Operator SHOULD adopt and enforce defined timelines for communicating and reporting incidents to the relevant HDAB(s):** • Early warning notification: within 24 hours of incident detection • Detailed incident notification: within 72 hours of incident detection • Final incident report: no later than one month after the incident Incident reports must include minimum fields such as incident ID, timestamp, severity, affected data, and mitigation actions.	The NIS2 Directive mandates incident notification to the relevant authorities within 24 hours of awareness, with a more detailed incident notification to be submitted within 72 hours (or within 24 hours for certain entities, such as trusted service providers). Additionally, a final incident report must be provided no later than one month after the incident.	Recommended
OPR-21	EHDSR-16	Incident Management	**SPE Operator MUST be able to promptly halt access and processing activities within the SPE when misuse or data breaches are identified.**	Article 73(3) - HDABs shall take corrective action for any shortcomings, risks or vulnerabilities identified. ISO 27001 and ISO/IEC 27002 and NIS2 requirements identify the need for organisations to have procedures to stop or restrict access and processing in response to security incidents or data breaches.	Mandatory
OPR-22	EHDSR-16	Incident management	**SPE Operator MUST have disaster recovery procedures in place to restore the availability and integrity of the SPE services, including critical system components, configurations, and platform-level functionality, from clean backups, and to resume normal service operations following an incident.** Project-specific data recovery may be limited to what is feasible within backup and retention policies. Disaster recovery measures should be proportionate to the SPE's size,	ISO 27001 and ISO/IEC 27002 - Organisations must have a disaster recovery process that includes restoring systems and service functionality from backups and return to normal operations. NIS2 requirement - Organisations must have plans for business continuity and disaster recovery, including the ability to restore critical systems and services after an incident.	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			complexity, risk profile, and service criticality.		

Risk management and mitigation

Risk management and mitigation are fundamental to ensuring the security, reliability, and resilience of SPEs. The EHDS Regulation establishes clear obligations for this purpose: Article 73(1)(b) requires that technical and organisational measures are in place to minimise the risk of unauthorised reading, copying, modification, or removal of electronic health data hosted in the SPE. Beyond the regulatory requirements, proactive risk management and mitigation represent core functionalities of secure information systems, providing a structured approach to identifying, assessing, and addressing potential threats to system operations, data confidentiality, and user compliance. Based on the identified threats to SPEs (see Annex 10: Classification of risks and threats against SPEs for examples), the associated risks can be assessed and appropriate mitigation strategies planned. By implementing a comprehensive risk management framework, SPE providers can protect sensitive data and maintain system integrity.

To implement effective risk management in SPEs, a combination of technical, organisational, and operational functionalities is required. These include robust security measures, such as firewalls, encryption, intrusion detection systems, and controlled access mechanisms, to prevent unauthorised access or manipulation of sensitive data. Personnel interacting with the SPE, including HDAB staff, SPE operator staff, and health data users, must receive role-specific training covering EHDS compliance, secure data handling, and relevant GDPR principles. Operational measures must also incorporate strategies for backup management, disaster recovery, and crisis management to ensure business continuity and resilience in case of incidents. Complementary operational practices, such as service-level agreements (SLAs) and regular restoration of backups, support preparedness and accountability across the entire SPE workflow.

Defining operational requirements for risk management and mitigation involves addressing several important considerations. Effective risk mitigation requires coordination across all actors involved in the data provision and processing workflow, from the HDAB to the SPE operator and down to health data users. Additionally, preparedness for recovery and incident response is critical in a live processing environment, where ongoing data analysis must be safeguarded against loss, disruption, or compromise. Effective measures must therefore ensure that work-in-progress analyses can continue safely while maintaining protection of data integrity and confidentiality. These considerations have guided the formulation of the operational requirements outlined below.

The following operational requirements provide a structured description of recommended practices for risk management and mitigation in SPEs. They specify mandatory and recommended measures designed to safeguard sensitive data, maintain system integrity, and ensure continuity of SPE operations in line with the EHDS Regulation.

Table 5.5. lists additional requirements and their derivation that is shown in Figure 5.11.

Figure 5.11. Risk management and mitigation. Requirement identifiers follow the convention: 'SPER–n' for general purpose SPE requirements ('SPE' namespace + 'R' for requirement), 'EHDSR–n' for EHDS requirements ('EHDS' namespace + 'R' for requirement) and 'OPR–n' for operational requirements ('OP' namespace + 'R' for requirement). OP requirements and their importance are interpretations of the EHDS Regulation.

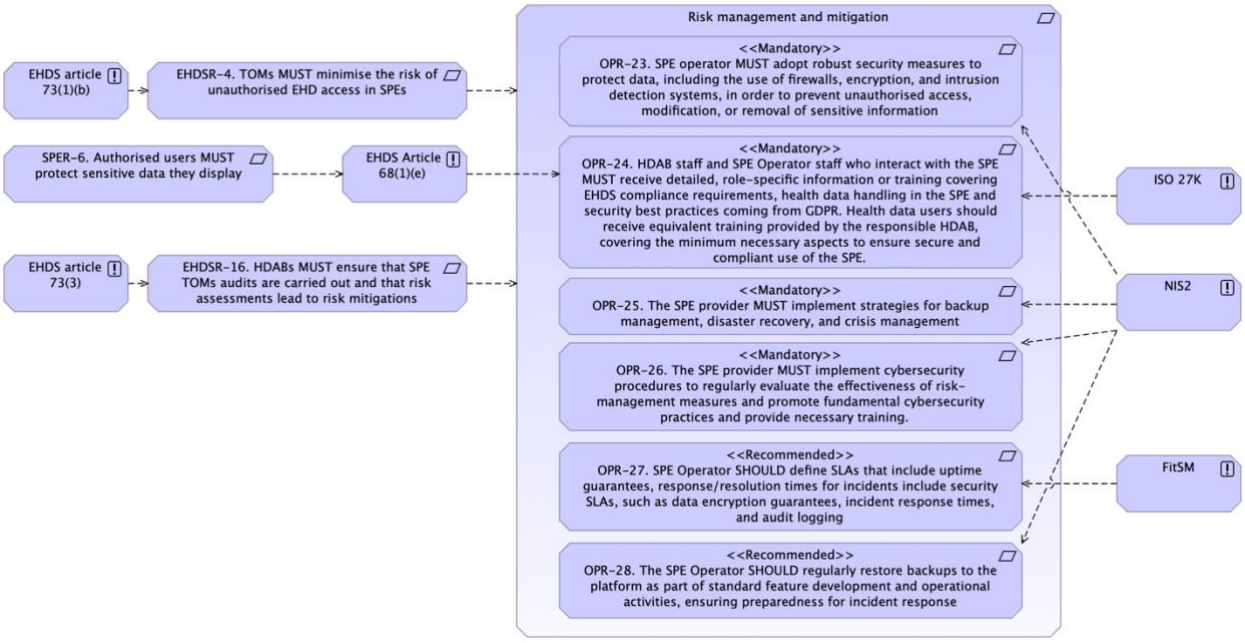

Table 5.5: Risk management and mitigation

#	Link	Area	Requirement	Rationale	Importance
OPR-23	EHDSR-16	Risk mitigation	**SPE Operator MUST adopt robust security measures to protect data, including the use of firewalls, encryption, and intrusion detection systems, to prevent unauthorised access, modification, or removal of sensitive information.**	Article 73(1)(b) - Technical and organisational measures must be in place to minimise the risk of the unauthorised reading, copying, modification or removal of EHD hosted in the SPE. NIS2 Directive explicitly references the implementation of appropriate security measures, such as incident detection and regular security assessments, to protect data and systems. Measures like encryption and multi-factor authentication are considered state-of-the-art practices that can support compliance with these requirements.	Mandatory
OPR-24	EHDSR-16	Risk mitigation	**HDAB staff and SPE Operator staff who interact with the SPE**	Article 68(1)(e) - The health data applicant must demonstrate sufficient	Mandatory

#	Link	Area	Requirement	Rationale	Importance
			MUST receive detailed, role-specific information or training covering EHDS compliance requirements, health data handling in the SPE and security best practices coming from GDPR. Health data users should receive equivalent training provided by the responsible HDAB, covering the minimum necessary aspects to ensure secure and compliant use of the SPE.	technical and organisational measures to prevent data misuse and protect the rights of data holders and individuals. Article 73(1)(b) - Technical and organisational measures must be in place to minimise the risk of the unauthorised reading, copying, modification or removal of EHD hosted in the SPE. ISO 27001 and ISO/IEC 27002 require role-appropriate security training for all staff and relevant contractors, covering policies, responsibilities, and basic security practices.	
OPR-25	EHDSR-16	Risk mitigation	**The SPE Operator MUST implement strategies for backup management, disaster recovery, and crisis management.**	NIS2 directive stipulates that organisations should have a plan for managing business operations during and after a security incident. This means that backups must be up to date. There must also be a plan for ensuring access to IT systems and their operating functions during and after a security incident.	Mandatory
OPR-26	EHDSR-16	Risk mitigation	**The SPE Operator MUST implement cybersecurity procedures to regularly evaluate the effectiveness of risk-management measures and promote fundamental cybersecurity practices and provide necessary training.**	NIS2 directive requires that employees receive cybersecurity training and practice for basic computer hygiene.	Mandatory
OPR-27	EHDSR-16	Risk mitigation	**SPE Operator SHOULD define SLAs that include uptime guarantees, response/resolution times for incidents include security SLAs, such as data encryption guarantees, incident response times, and audit logging.**	FitSM Service Level Management recommendation – a service catalogue must be maintained, and SLAs should clearly define service targets such as uptime, response times, and resolution times.	Recommended
OPR-28	EHDSR-16	Risk mitigation	**The SPE Operator SHOULD regularly restore backups to the platform as part of standard feature development and operational activities, ensuring preparedness for incident response.**	NIS2 directive - organisations should have a plan for managing business operations during and after a security incident. This means that backups must be up to date. There must also be a plan for ensuring access to IT systems and their operating functions during and after a security incident.	Recommended

Maintenance and support

Maintenance and support mechanisms are essential to ensure that SPEs remain secure, reliable, and fully operational over time. While the EHDS Regulation does not explicitly define detailed operational requirements for these activities, Article 73(3) establishes that corrective actions must be taken for any shortcomings, risks, or vulnerabilities identified during audits of SPEs. Furthermore, the regulation implies that SPEs must be resilient and ensure service continuity, which necessitates responsive support and structured maintenance processes. Beyond regulatory obligations, maintenance and support are fundamental practices in information systems management, providing the mechanisms to manage updates, resolve incidents, and preserve operational integrity while sustaining user confidence.

To achieve effective maintenance and support in SPEs, several high-level functionalities are needed. These include robust patch management and system update processes to promptly identify, test, and deploy security patches, updates to operating systems, firmware, and application software, and controlled change management procedures to ensure that updates do not compromise ongoing analyses or system stability. Dedicated technical support with clearly defined SLAs and escalation paths is required to assist users and address operational issues promptly. Support staff must be trained in information security, privacy procedures, and specific SPE operational practices. Additionally, maintaining a knowledge base and support documentation helps streamline issue resolution, reduces response times, and mitigates the risk of user errors or repeated support requests.

Defining operational requirements for maintenance and support involves careful consideration of several challenges. SPEs must balance the need to keep systems up-to-date and secure with the need to preserve reproducibility, operational continuity, and compliance with active data permits. Changes, updates, or patches can have unintended consequences on ongoing analyses, so a controlled change management process is essential to assess backward compatibility, identify potential risks, and coordinate updates with all relevant stakeholders. Furthermore, as SPEs will be new systems used by multiple actors, clear support structures, comprehensive documentation, and accessible guidance are critical to prevent overloading technical support teams and to ensure smooth onboarding and user adoption. These considerations underpin the operational requirements elaborated below, providing guidance to maintain SPE functionality, reliability, and regulatory compliance.

The following operational requirements provide a structured description of recommended practices for maintenance and support in SPEs, focusing on system updates, patch management, change control, and user assistance to ensure resilient and secure operation. Table 5.6 lists these requirements and their derivation that is shown in Figure 5.12.

Figure 5.12. Maintenance and support. Requirement identifiers follow the convention: ‘EHDSR–n’ for EHDS requirements (‘EHDS’ namespace + ‘R’ for requirement) and ‘OPR–

n' for operational requirements ('OP' namespace + 'R' for requirement). OP requirements and their importance are interpretations of the EHDS Regulation.

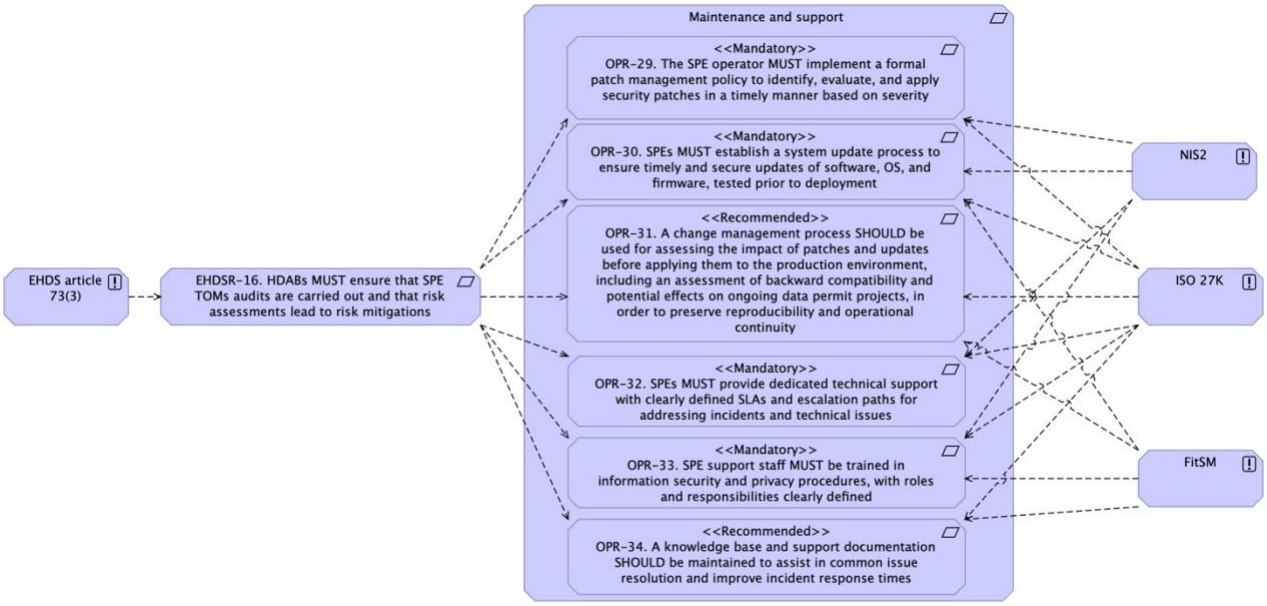

Table 5.6: Maintenance and support

#	Link	Area	Requirement	Rationale	Importance
OPR-29	EHDSR-16	Maintenance	**The SPE Operator MUST implement a formal patch management policy to identify, evaluate, and apply security patches in a timely manner based on severity.**	EHDS Regulation (Article 73(3) requires corrective actions for any shortcomings, risks or vulnerabilities identified un the process of SPE audits. NIS2 Directive mandates organisations to implement risk management measures that include updating systems to protect against known vulnerabilities. ISO/IEC 27001 and ISO/IEC 27002 explicitly require a patch management process.	Mandatory
OPR-30	EHDSR-16	Maintenance	**SPE Operator MUST establish a system update process to ensure timely and secure updates of software, OS, and firmware, tested prior to deployment.**	NIS2 Directive stipulates regular system maintenance and software updates to reduce risk exposure. ISO/IEC 27001 and ISO/IEC 27002 require that information systems are regularly updated with new versions and patches. FitSM requirement - structured and traceable system update processes must be in place	Mandatory
OPR-31	EHDSR-16	Maintenance	**A change management process SHOULD be used for assessing the impact of patches and updates**	ISO/IEC 27001 and ISO/IEC 27002 requirements recommend that changes to information systems are controlled and authorised to avoid	Recommended

#	Link	Area	Requirement	Rationale	Importance
			before applying them to the production environment, including an assessment of backward compatibility and potential effects on ongoing data permit projects, in order to preserve reproducibility and operational continuity.	disruptions or vulnerabilities. FitSM requirement - all changes to services or infrastructure must be reviewed, approved, and documented, ensuring continuity and stability during patching and updates. This reduces the risk of unplanned outages or system compromise.	
OPR-32	EHDSR-16	Support	**SPE Operator MUST provide dedicated technical support with clearly defined SLAs and escalation paths for addressing incidents and technical issues.**	EHDS Regulation (Article 73(3)) requires SPEs to be resilient and ensure service continuity, implying that support services must be available and responsive. NIS2 Directive: Need for incident handling and response capabilities. ISO/IEC 27001 and ISO/IEC 27002 requires that technical support be available to ensure effective incident detection, response, and recovery.	Mandatory
OPR-33	EHDSR-16	Support	**SPE Operator support staff MUST be trained in information security and privacy procedures, with roles and responsibilities clearly defined.**	ISO/IEC 27002 highlights the importance of assigning roles and responsibilities related to security and ensuring adequate training. NIS2 Directive requires that personnel involved in system operation receive regular cybersecurity awareness training. FitSM requirement **-** ongoing education and role-based training should be provided to ensure staff can respond effectively to incidents and manage secure systems.	Mandatory
OPR-34	EHDSR-16	Support	**A knowledge base and support documentation SHOULD be maintained to assist in common issue resolution and improve incident response times.**	ISO/IEC 27002 - Documentation of known issues and responses shall be kept to improve incident handling and system stability**.** FitSM requirement - Building and maintaining a knowledge base as part of support operations to enable faster resolution and avoid repetition of known problems. This contributes to efficiency and organisational learning in support environments.	Recommended

5.6 Technical interoperability requirements

General

Existing EHDS technical requirements (Specific Contract no.22 documents: D02.03 Requirements Catalogue[7], D03.03 Architecture Artefacts[8], D03.03 System Specifications[9]) are not directly addressing technical interoperability, but we have aligned with these documents where applicable, adopting consistent terminology and using the architecture described in them as a reference for interoperability requirements.

The high-level technical architecture of the EHDS SPE is depicted in Figure 5.13. Interfaces towards the following entities are identified:

- Data user applications
- Data holder applications
- External SPEs
- HDAB coordinator portal

This section provides high-level technical specifications for the data user and data holder application interfaces as well as the interface for connecting external SPEs. The specifications are applicable both for federated computing and traditional data processing. The governance and safeguarding processes needed for setting up and managing the interfaces are beyond the scope of this deliverable.

[7] D02.03 Requirements Catalogue for Scale-Up version. Specific Contract no. 22 under Framework Contract SLG.AVT.DI07926 - BEACON – Lot2, Date: 17.12.2024, Version 3.1.
[8] D03.03 Architecture Artefacts – Scale-Up version. Specific Contract no. 22 under Framework Contract DI7925-DI7932 – BEACON – Lot2: Analysis and design of the European Health Data Space infrastructure for secondary use of health data (HealthData@EU), Date: 4.12.2024, Version 1.0.
[9] D03.03 System specifications – Scale-Up version. Specific Contract no. 22 under Framework Contract DI7925-DI7932 – BEACON – Lot2: Analysis and design of the European Health Data Space infrastructure for secondary use of health data (HealthData@EU), Date: 10.12.2024, Version 1.0.

Figure 5.13. High-level technical architecture of EHDS SPE (modified from D03.03 Architecture Artefacts).

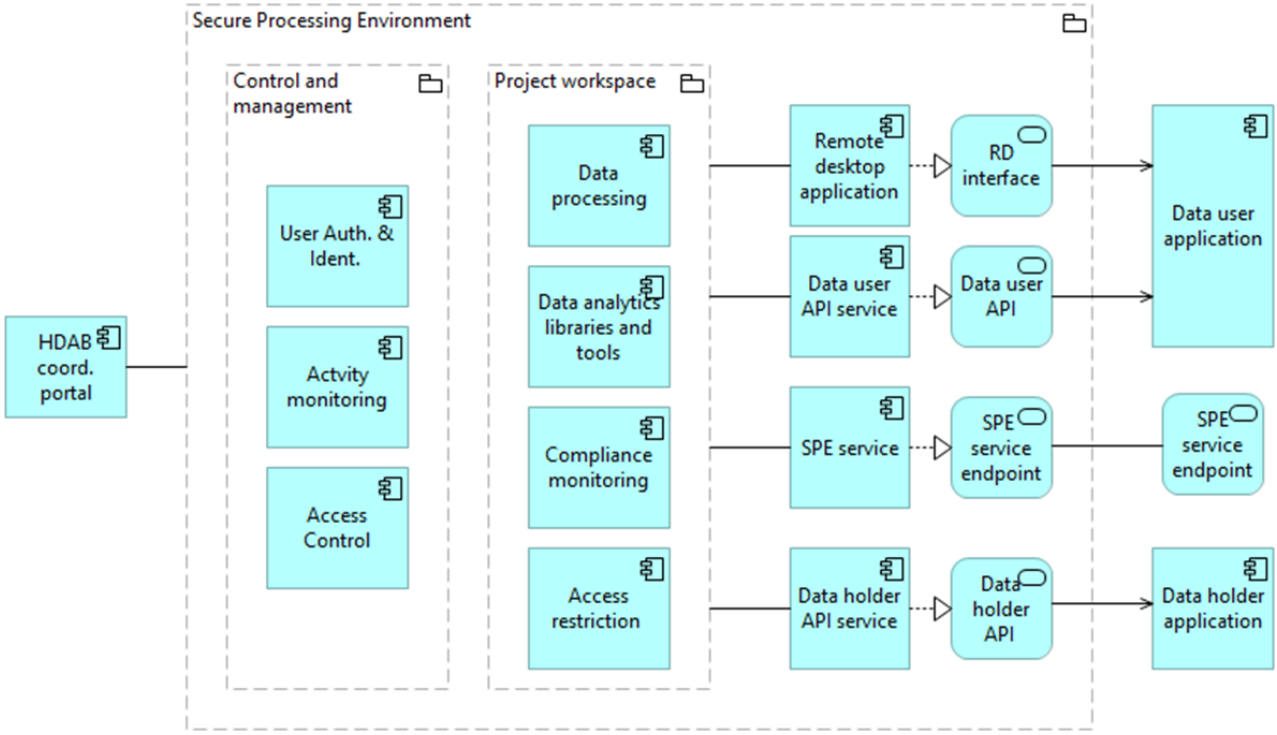

Data user application is a user-facing interface supported by the SPE infrastructure. It allows users to securely connect to their designated project workspace, which is isolated from the other workspaces. Two technical approaches should be supported by the SPE: (1) remote desktop connection and 2) API based connection.

The remote desktop (RD) interface allows the data user to interactively process data in the project workspace. Remote desktop is the legacy approach widely used in the currently existing SPE installations. The data user application (remote desktop client) allows the user to access the project workspace as if working directly on its host machine, supporting interactive tasks such as file management, software execution, and visualisation.

The API interface enables more versatile user applications, such as web portals and standalone applications provided or authorised by the HDAB. The API based data user application may enable to trigger execution of functions (e.g. data processing scripts), monitor progress of function execution, retrieve anonymous data analysis results, download anonymised data sets, communicate with authorised organisation (HDAB, authorised participant) and upload custom tools, supporting contents (e.g. vocabularies) and additional data sets. The API interface also is a prerequisite for federated analysis as outlined in Section 4.2.2.

The SPE operator is responsible for setting up the required processes and controls to ensure that the exposed API services ensure privacy protection and compliance with EHDS Regulations. For instance, it should be possible for the SPE operator to manually or automatically verify that the downloaded contents meet anonymity requirements.

Data holder application is the data holder's interface supported by the SPE infrastructure. The application enables the data holder to upload data to the SPE based on an accepted data permit. The data holder application uses an API based connection mechanism.

SPE service endpoint enables connections to external SPEs. Such connections enable federated networks of trusted SPEs to be formed. Options for setting up such networks are discussed in the SPE federation chapter.

HDAB coordinator portal connection enables the relevant organisation (HDAB, authorised participant or SPE operator) user to carry out overall management of the SPE and to execute specific tasks needed to support data usage and ensure regulatory compliance. Technical specifications for connecting the HDAB coordinator portal with SPEs may vary between implementations and are not in the scope of this deliverable.

To fulfil the regulatory requirements (EHDS, Article 73) the SPE operator shall provide the RD interface and the data holder API shown in Figure 5.13. Implementation of the data user API and SPE service endpoint is recommended, but not obligatory.

Technical requirements for the interfaces are listed in the following subsections.

Scope

The following are **out of scope** for this deliverable:

- Implications for EHDS governance, such as the authorisation of data transfers between SPEs and the deployment or execution of required software components (see also SPE federation)
- Manual implementation of federated analysis, for example, cases where the user manually combines results from separate analyses, as permitted by their approved data permit(s).

5.6.1 Data user and data holder API requirements

Technical interoperability requirements of federated SPE under EHDS Regulation have the namespace acronym 'TI' for **Technical Interoperability** followed by the letter 'R'. TI requirements do not fall under EHDS Regulation. Importance levels of its requirements are valid only within its namespace.

Table 5.7: General API requirements

#	Requirement	Importance
TIR-1	**Client application.** The API services MUST be accessed via data user and data holder applications, which can be dedicated client applications or browser-based applications.	Mandatory
TIR-2	**Architectural style.** The API MUST follow a web services architecture (e.g., RESTful, GraphQL), enabling stateless request/response interactions and supporting secure file transfer protocols. The implementation MUST adhere to common	Mandatory

	architectural and security practices, including resource grouping and, where appropriate, the use of microservices to isolate sensitive services.	
TIR-3	**Communication protocol.** The API MUST use industry-standard security protocols to en-sure compatibility with standard web clients and server implementations. For API requests, HTTPS MUST be supported, and for file operations, secure file transfer protocols such as HTTPS, SFTP, or FTPS MUST be supported, where applicable.	Mandatory
TIR-4	**Data exchange format.** The API MUST support industry-standard data formats, ensuring at least JSON is available for requests and responses, while allowing additional formats as needed.	Mandatory
TIR-5	**File transfers.** The API MUST be capable of uploading and downloading files using stand-ard file types (e.g. CSV and parquet). It MUST provide configurable restrictions on allowed file types and transfer directions (upload/download) separately for each API and workspace. For large files, the API MUST support efficient handling, including resumable transfers, asynchronous processing where appropriate, and configurable timeout settings to accom-modate long-duration operations.	Mandatory
TIR-6	**API documentation.** The API MUST provide clear, machine-readable documentation (e.g., OpenAPI/Swagger) in accordance with ISO 27001 information security and change management controls to ensure proper governance, traceability, and compliance.	Mandatory

Table 5.8: API security and access control requirements

#	Requirement	Importance
TIR-7	**Authentication and authorisation.** All requests to data user and data holder API functions MUST be authenticated and authorised using industry-standard methods (e.g., OpenID Connect, OAuth 2.1 with JWT, API keys, and JWS-signed requests). The access control mechanism MUST enforce attribute-based access control (ABAC), limiting access based on attributes such as user role, organisation, clearance level, and data permit approval sta-tus.	Mandatory
TIR-8	**Data encryption.** The API MUST ensure data in transit is encrypted using industry-standard protocols (e.g., TLS). For added security, the API MUST support content encryption (e.g., AES-256) prior to transmission, enforced by the HDAB, to protect data even if stored by the API. When encryption is applied, encryption keys MUST be managed and stored securely in accordance with industry best practices, such as via a dedicated key management system (KMS).	Mandatory
TIR-9	**Input validation and sanitisation.** The API MUST validate and sanitise all incoming data using industry-standard practices to prevent injection attacks (e.g., SQLi, XSS) and mal-formed requests. Encrypted payloads MUST be decrypted before validation, and only pre-defined, approved request types are allowed to be processed.	Mandatory

TIR-10	**Error handling and security.** Error responses MUST avoid exposing sensitive details (e.g., stack traces, internal error codes) while providing meaningful messages for debug-ging.	Mandatory
TIR-11	**Network access control.** API access MAY be restricted based on network-level controls, including firewall rules, IP whitelisting, or Virtual Private Network (VPN) restrictions.	Recommended

Table 5.9: API service management and monitoring

#	Requirement	Importance
TIR-12	**Logging and auditing.** The API MUST log all requests, responses, and errors following industry-standard practices to support monitoring and compliance. Logging MUST adhere to a predefined purpose and retention period, ensuring that data is kept only as long as necessary and in accordance with applicable policies.	Mandatory
TIR-13	**Monitoring and alerts.** The API MUST monitor usage, failures, performance metrics, and service availability in accordance with industry-standard practices, and alerts MUST be generated for any detected anomalies to support operational reliability and compliance.	Mandatory
TIR-14	**API version management.** The API MUST have mechanisms to approve, deploy, and retire API versions, following industry-standard practices and aligning with ISO 27001 change management controls to ensure secure and controlled updates.	Mandatory

Table 5.10: API performance and scalability

#	Requirement	Importance
TIR-15	**Response time.** The API SHOULD meet defined latency and response time SLAs.	Recommended
TIR-16	**Load handling.** The system SHOULD handle expected and peak loads efficiently, with rate limiting and throttling mechanisms in place if necessary.	Recommended
TIR-17	**Data Validity.** The API SHOULD ensure that payloads of requests and responses comply with the specified data formats and content standards, applying appropriate semantic valida-tion mechanisms. This ensures consistency, accuracy, and interoperability of exchanged data, preventing the delivery of erroneous or inconsistent content.	Recommended

5.6.2 Data user remote desktop interface

Table 5.11: General remote desktop interface requirements

#	Requirement	Importance
TIR-18	**Functionality.** The remote desktop interface MUST enable the data user to interactively access and process data with a	Mandatory

	remote desktop application ("SPE UI" in D03.03 Architecture Artefacts).	
TIR-19	**Data user application.** The data user MUST be able to use a standard browser or a dedicated client to access the remote desktop environment.	Mandatory
TIR-20	**Platform compatibility.** The remote desktop solution MUST support remote desk-top clients running on standard operating systems including (e.g., Windows, macOS, Linux).	Mandatory
TIR-21	**Protocol support.** The remote desktop solution MUST use industry-standard, purpose-built remote access protocols (e.g., RDP or VNC) to establish connections between the client and server. All communication MUST be encrypted in transit using industry-standard secure protocols. If SSH is used, it SHALL be restricted to transport or tunneling purposes only and SHALL NOT provide interactive shell or file transfer access to end users.	Mandatory

Table 5.12: Remote desktop security and access control

#	Requirement	Importance
TIR-22	**Authentication and authorisation.** Access to the remote desktop interface MUST require multifactor authentication (MFA) in accordance with industry-standard security practices. The system MUST enforce attribute-based access control (ABAC), limiting access based on attributes such as user role, organisation, clearance level, and data permit approval sta-tus, ensuring that only authorised users can access the system.	Mandatory
TIR-23	**Network Access Control.** Access MAY be restricted based on network-level controls, including firewall rules, IP whitelisting, or VPN restrictions	Recommended
TIR-24	**Session management.** Idle sessions MUST be automatically terminated after a predefined time to prevent unauthorised access. Users MUST be logged out or locked after a period of inactivity.	Mandatory
TIR-25	**Data transfer restrictions.** Clipboard sharing and file transfers MUST be restricted by default to prevent unauthorised export of data from the SPE. Changes to or removal of these restrictions MUST be configurable at the discretion of the HDAB.	Mandatory

Table 5.13: Remote desktop service management and monitoring

#	Requirement	Importance
TIR-26	**Logging and auditing.** All remote sessions MUST be logged in accordance with industry-standard practices, including authentication attempts, connection times, and actions per-formed during the session. Logging MUST adhere to a predefined purpose and retention period, ensuring that data is kept only as long as necessary and in accordance with applicable policies.	Mandatory

TIR-27	**Monitoring and alerts.** Service usage, failures, performance metrics and service availability MUST be monitored, with alerts for anomalies.	Mandatory

5.6.3 SPE service endpoint

Table 5.14: General SPE service endpoint requirements

#	Requirement	Importance
TIR-28	**Functionality.** The service endpoint MAY enable communication between trusted SPEs.	Recommended
TIR-29	**Communication Protocol.** The interface MUST support protocols required for client-server and bidirectional streaming communication (e.g., gRPC with HTTP/2 or an equivalent). All such connections MUST use industry-standard secure protocols (e.g., TLS) to protect data in transit.	Mandatory

Table 5.15: SPE service endpoint security and access control

#	Requirement	Importance
TIR-30	**Authentication and authorisation.** All connections MUST be authenticated using industry-standard mechanisms (e.g., mutual TLS (mTLS) or token-based authentication), and all data MUST be transmitted over industry-standard secure channels to ensure confidentiality and integrity.	Mandatory
TIR-31	**Client whitelisting.** Connecting clients MUST be preapproved before establishing a connection.	Mandatory
TIR-32	**Network access control.** The interface MAY be configured to restrict access using firewall rules, IP whitelisting, VPN access, other network-level policies and holistic approaches such as virtual closed networks.	Recommended

Table 5.16: SPE service endpoint service management and monitoring

#	Requirement	Importance
TIR-33	**Logging and auditing.** All connection sessions MUST be logged in accordance with indus-try-standard practices, including requests, responses, streaming events and connection times. Logging MUST adhere to a predefined purpose and retention period, ensuring that da-ta is kept only as long as necessary and in accordance with applicable policies.	Mandatory
TIR-34	**Monitoring and alerts.** Real-time monitoring of streaming sessions, errors, latency, and security threats MUST be implemented. Alerts MUST be generated based on anomalous activity (e.g., failed authentication attempts, unusual data patterns).	Mandatory

Table 5.17: SPE service endpoint performance and scalability

#	Requirement	Importance
TIR-35	**Response time and streaming performance.** The interface MUST meet defined latency and response time SLAs.	Mandatory
TIR-36	**Load handling.** The system MUST implement flow control to manage varying loads and prevent overloads, using rate limiting, throttling, and adaptive resource allocation as needed.	Mandatory

5.6.4 Relation to existing specifications

Existing specifications (D03.03 Architecture Artefacts) use the term “SPE UI”, which seems to refer to a remote desktop type interface used in legacy SPE implementations. In this deliverable, we have proposed two types of data user applications: (1) the legacy remote desktop approach and (2) a new approach where the data user application interacts with the SPE via a web API interface. Approach (1) provides a familiar, typically browser-based, method for data users. Approach (2) enables new ways to enable more efficient ways to exchange information between the health data user and the SPE environment, including the possibility to execute federated analysis in multiple SPEs.

Similarly, we propose the API approach to be used also for the data holder application, where such approach would enable a more efficient process with less manual interaction needed.

We have included the interface for connecting external SPEs, which is not covered in D03.03 Architecture Artefacts. This interface enables advanced federated learning scenarios requiring direct information exchange between SPEs. These scenarios need careful governance and implementation to comply with EHDS Regulation which are not in the scope of this deliverable. The development of more detailed specifications for the API and the SPE service point, together with the associated governance models, shall be key objectives of subsequent development activities.

5.7 Challenges for EHDS SPE

5.7.1 Operational requirements for EHDS SPE federation

The paragraphs of Article 73 of the EHDS Regulation contain very specific legal demands for SPEs that function primarily in isolation. We have interpreted these to form 18 practical requirements out of them (See chapter EHDS SPE requirements and Annex 9: EHDS article 73 analysis to deduce SPE requirements). Many of them are purely operational, meaning their successful and secure application depend on the interplay of the personnel maintaining the services and the processes they implement, rather than being purely technical. Extending these requirements to fully support federated processing within EHDS, these requirements need to be interpreted and deployed in a uniform manner throughout the federation to maintain the functioning and security of services.

In addition to EHDS, the catch-all security standards applied to high-end security services and organisations are the ISO/IEC 27000 -family of standards (ISO 27K) where ISO/IEC 27002 deals with information security controls. These are not very helpful for determining

specific technical requirements for SPEs that go beyond it. The bulk of ISO27K requirements are operational, affecting the organisation maintaining the service.

The NIS2 Directive builds on ISO 27K security demands and requires reporting within member countries and collection of all serious incidents to an EU-wide registry.

We have organised and tabulated various operational demands to the SPE from EHDS, ISO 27K and NIS2 (see Operational EHDS SPE requirements), but have excluded operational requirements that direct HDAB and its reporting requirements to the EU central platform.

Our evaluation also includes the FitSM standard, which we suggest as a lightweight framework to support cohesion of operation and reporting for EHDS-compliant SPEs because we feel it offers significant advantages to the EHDS over ISO 27K that is closed, heavyweight and expensive to certify against.

FitSM is a lightweight, open-source standard for IT Service Management, developed through an EC funded project (FedSM, 2012-2015) and still actively maintained today. It supports any service delivery scenario but is unusual in supporting federated service provision and being much easier to implement in academic and public sector than other heavier traditional frameworks. FitSM has a wider scope than ISO 27K but covers information security management and can be easily integrated with an ISO27K management system. It offers strong advantages in that it can then connect security requirements more thoroughly to service design and delivery, service level management, capacity and availability management, customer relationship management and other IT Service Management processes (e.g. Service Portfolio Management, Incident & Problem Management, Change Management, Service Level Management).

The core FitSM documentation is freely available under a Creative Commons licence, so it can be easily embedded and referenced in organisational and national documentation. This lets it be used as a free and lightweight reference model for managing services and for connecting security and operational requirements across EHDS.

FitSM is proposed here as a practical framework to support coordination and documentation but does not replace the legal obligations established under the EHDS Regulation or related EU legislation.

To ensure mutual trust across the EHDS SPEs, the documentation produced by each MS should demonstrate that auditing processes are implemented consistently and meet the common safeguards required under Article 73. This documentation should enable Member States to recognise each other's SPEs as compliant with the Regulation. Without such transparency and alignment, the SPE federation cannot guarantee a uniform level of protection for sensitive data across borders, thereby undermining its core objective

FitSM should offer a lightweight and cost-effective way for EHDS SPE federation to coordinate its operational requirements to fulfil that goal.

We recommend that EHDS should build its guidance for EHDS SPE federation compliance reporting based on FitSM.

5.7.2 Cybersecurity of SPE infrastructure

This chapter of the SPE security considerations deals with challenges arising from virtualisation that are wider than SPE, e.g. the 18 security functions of the CIS framework[10].

SPE implementation is typically run as part of a virtualisation service where the aim is to isolate the user's project environment from other user environments and fully control the access to additional services. This isolation is built using nested security structures, where user applications run within virtualised environments that simulate their own operating systems, hosted on a local network. Eventually, the virtualisation layers interface with the physical hardware, where all processes are ultimately executed.

The main aim of cybersecurity guidelines is to determine what is the adequate protection for a given setup balancing both enabling and limiting aspects. Given the aim to tap into creativity of researchers, we have already highlighted the need for openness and information sharing among the project members that requires them to understand their responsibilities and implications of their actions because too onerous technical actions are counterproductive.

In our minimum requirements for SPEs these are expressed by **SPER-4** and **SPER-6**. We may rephrase these sentiments to: SPE isolation measures SHOULD protect everything else except user's own immediate environment. Unfortunate users of an SPE should be allowed to find themselves in a situation where they exercise their freedom of actions to the extent that they lose the most recent results and their environment must be recreated to re-allow them access to the permitted sensitive data.

Following the same logic, it should not matter what applications users have access to or what they have installed themselves. The operating system setup should prevent the harm from spreading outside the immediate user environment. If the users or their applications gain elevated access privileges, the network isolation is there for the next layer of protection. They should never reach administrative control of the system.

However, it is not elephants all the way down. Networks and processes run on physical hardware. Secure processing running on the hardware maintained by the SPE service provider is seen as private cloud. In contrast, services running on public cloud necessarily involve trust to the cloud provider.

The figure 5.14 illustrates the nested structure of computer cybersecurity layers and the full chain of trust inside a private cloud should enable the provider to control the operating system layer and offer its users SPE as a Platform as a Service (PaaS) where users have wide freedoms to import and install their own software. If services are built on public cloud, SPE operator might not have enough trust to the underlying services to provide high enough security, so they might have to restrict user's abilities to install software. Then SPE is essentially providing Software as a Service (SaaS).

[10] CIS Controls https://www.cisecurity.org/controls/cis-controls-list

Figure 5.14. Nested isolation of SPE processing.

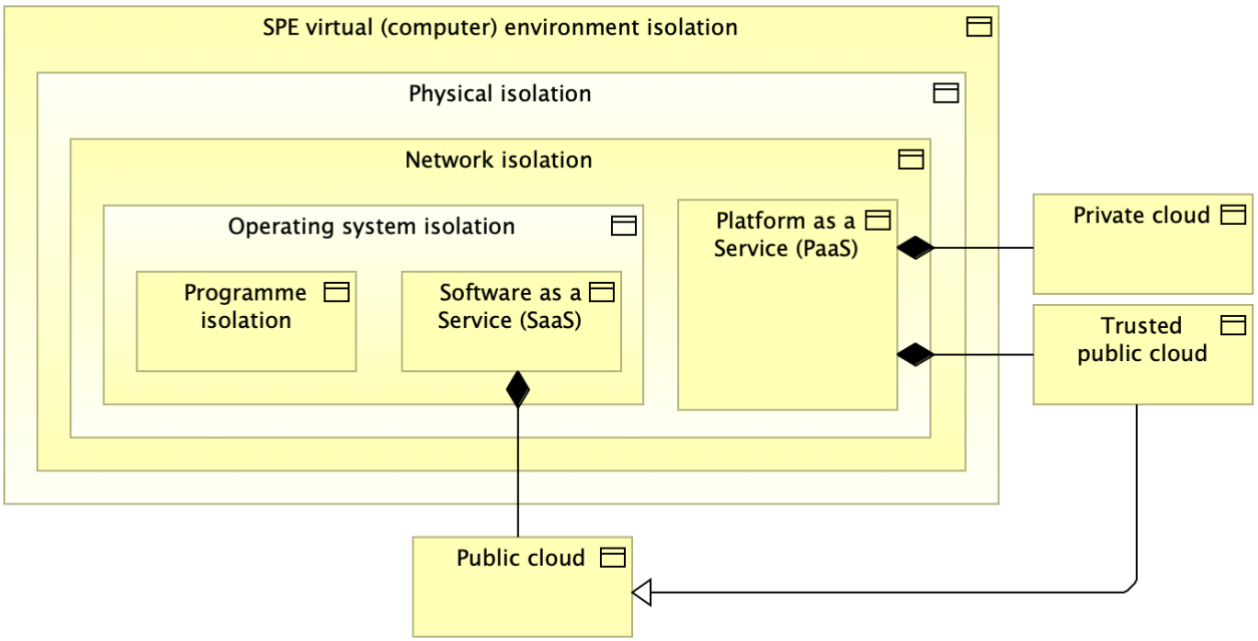

A key outcome of the UK community-driven SATRE project is a standard architecture specification for TREs (See Annex 8: SATRE). Closely aligned with TEHDAS2, it defines core TRE capabilities and criteria for assessing compliance. Regarding network security and nested isolation, the SATRE specifications include the following mandatory requirement: SATRE 2.2.9. "Your TRE must control and manage all of its network infrastructure". However, globally available public cloud services are common use in UK.

The recent proliferation of public clouds is based on reliance on laws and regulations providing foundations for usual business subcontracts that have been seen enough to elevate public clouds to the same level as private clouds.

Public cloud providers may be subject to extraterritorial legislation (e.g. US CLOUD Act[11]). The suitability of such providers for SPE hosting must therefore be reassessed under EU data protection, cybersecurity, and sovereignty requirements. The recent American governmental policy shift to emphasise US presidential prerogative power has already led to calls for purely European cyber infrastructure to safeguarding sensitive European data[12].

The current political climate has made it clear that cybersecurity requirements for SPEs and their federation will need to be reconsidered to redefine the adequate level of trust for sensitive data computer infrastructure is run on.

5.7.3 Projection to EHDS SPE Federation

Secondary use in EHDS is a federation of interconnected services sharing information about available health data, as detailed in the EHDS Regulation. It provides legal requirements for

[11] CLOUD Act https://en.wikipedia.org/wiki/CLOUD_Act
[12] e.g. EuroStack https://euro-stack.eu/

member states to produce metadata of their health datasets and requires those to be findable centrally though the HealthData@EU infrastructure.

EHDS has a legal obligation for moving sensitive data only within jurisdiction, demanding at least one SPE service for processing the permitted data in each member country and one for the European Commission. The latter will primarily serve European Union's institutions, bodies, offices and agencies (Article 75(2)). Additionally, the member state EHDS official, HDAB, may allow their data to be processed alone or in together with other permitted datasets in any of the approved European SPEs.

Moreover, EHDS calls other European infrastructures like ERICs and EDICs Authorised Participant (AP) (Article 75(4)) that may participate in the HealthData@EU. EHDS data permits will then allow users to access the data based on the legal framework of that AP and get access to the data in the EHDS 73 compliant SPE provided by that AP.

The initial design of EHDS was based on the assumption of a stand-alone SPE maintained by HDAB. Data transfer would have been solely within each jurisdiction and would not need to be included into EHDS. However, the need to combine datasets over national borders and especially the rise of federated computing (see Federated computing) makes it necessary to bridge the gap from stand-alone SPE that works under one organisation by defining general requirements of an SPE federation.

EHDS is already an SPE federation in almost all but in name and technical support. The defined general SPE federation requirements are regardless of federation structure (Figure 4.7). Their main limitations arise from sensitivity of data. They enable sensitive data processing beyond a single service to cover processing over multiple separate but linked services within one organisation (like database or HPC), as well as enabling processing over organisational borders. These requirements should work for *ad hoc* proprietary data as well as EHDS-based processing.

Federation rules will need to direct the collaboration and establish federation services for management, mutual trust, findability and information sharing. An SPE federation will need findability services to locate datasets and SPEs. EHDS already includes detailed instructions for data findability services, but a similar approach will be needed to SPEs, too, depending on the national setting. Heath data users will need to search and compare a catalogue to find EHDS compliant SPEs nationally across EU.

5.7.4 Data access management and SPE interoperability

This section outlines the key requirements for interoperable and secure data access in an SPE federation, with emphasis on automation, standardisation, and support for scalable cross-border secondary use of health data under the EHDS Regulation.

The interoperability and scalability of an SPE federation largely depends on the efficiency of data access management. The security of large-scale dataset management needs automated procedures that manual data management cannot match. It has been estimated that to reach similar throughput of data permits in the future European EHDS framework serving the whole EU as the Findata authority did manually in 2024 for Finland, would need European member states altogether hire 10000 persons to handle all the applications (Jaakko Leinonen, CSC, FI, personal communication).

The lack of scalability in manual management of sensitive data life cycles (see Annex 6: Sensitive data lifecycles) has been the driver for designing the automated "state-of-art" approaches (see Annex 7: Scenarios). It is also the reason why this report defining the requirements and capabilities of SPE has been unable to do it without considering all the sensitive data lifecycle components. They are totally interdependent on each other. Enabling better functionality in one will not be possible without the other components supporting that.

The core functionality of the SPE is to process sensitive data. We already have a good grasp on how it can be done securely and ample approaches to make it happen more efficiently (see Annex 7: Scenarios), but they all depend on the capabilities of other components of the federation.

Cornerstones of this scalable infrastructure are:

- Shared identity and authorisation
- Machine-actionable access permits
- Data streaming
- Structured data warehouses

All security and accountability of sensitive data processing is based on identifying users with certainty (see Annex 7: Identity and authorisation). The data applicants and health data users should be identified in the beginning of the application process and that information needs to be maintained unbroken throughout. This means the official, human-readable document that is the data permit needs to be converted to a machine-readable access permit. This access permit combines the identifiers of the data permit to users' identities and to the newly created dataset that users will have access to.

An access permit cannot exist before the dataset promised in the data permit has been created and stored with a unique identifier. This new dataset is usually a combination of several source datasets with unnecessary items removed (according to GDPR minimisation principle) and pseudonymised.

Manual data transfer increases operational complexity and potential risk of human error. Automated streaming based on machine-readable access permits strengthens traceability and compliance with Article 73. The data manager will continue to have full control of the dataset and access to it. They can also make user-requested amendments to data if needed.

The primary technical requirement to start an SPE federation is to agree on at least one data transfer protocol (**FSPER-8**). While EHDS has already committed itself to eDelivery as the secure communication protocol for sending metadata records, data access application forms and messages between national contact point and the EU central platform, the data transfer details have not been agreed on. In practice, there is only one set of implemented community standards that cover the needs of sensitive data transfer and data access management for research: the Global Alliance for Genomics and Health (GA4GH) standards[13].

GA4GH crypt4GH[14] is an open source, secure, streaming capable encryption algorithm. It is created for huge genomic sequences but applicable to any file type. The encrypted file is

[13] GA4GH implementations https://www.ga4gh.org/our-products/implementations/
[14] Crypt4GH https://www.ga4gh.org/product/genetic-data-encryption-crypt4gh/

separated to header and payload. The payload is encrypted with a key that is stored in the header and never exposed. New user keys can be added and at will. Header can be stored and updated independently of the payload that cannot be decrypted with brute force methods. Sending header and payload separately creates a quantum safe data transfer. Storing crypt4gh headers in a secure database enables automatic user key management for secure data projects. The streaming capability used by GA4GH htsget API[15] can be linked to token-based data access control implementing the GA4GH Passport standard[16].

These standards are not mandated under the EHDS Regulation but could be considered as candidates for future implementing acts or common specifications under Article 73(5).

The use of these standards already forms the basis of plans for the framework for the EU One Million Genomes Initiative[17]. Its framework[18] gives detailed information about the initiative that aims to establish itself as a genomics Authorised Participant for EHDS.

Finally, one of the challenges in manual data management is that many European health data is either not stored in an electronic format at all or not in a standardised format. This makes determining the data applicants' needs and the creation of the needed new dataset difficult. The secondary use of health data should be one more driver towards rapid digitalisation and standardisation of health data at source.

[15] GA4GH htsget https://www.ga4gh.org/product/htsget/)
[16] GA4GH passport https://www.ga4gh.org/product/ga4gh-passports/
[17] 1+MG https://digital-strategy.ec.europa.eu/en/policies/1-million-genomes
[18] 1+MG framework https://framework.onemilliongenomes.eu/

6 Annexes

Annex number	Annex title
1	Methodology
2	Public consultation summary
3	User journey
4	Glossary
5	Historical context and legacy models
6	Sensitive data life cycles
7	Design considerations and expert commentary
8	Existing solutions for secure processing
9	Overview of relevant EU regulations
10	Classification of risks and threats against SPEs

Annex 1: Methodology

The specifications outlined in this report are the result of an in-depth analysis of the EHDS regulatory text, as well as findings from the PwC–Sopra Steria analysis and design of the European Health Data Space infrastructure for secondary use of health data, specific contract under FRAMEWORK CONTRACT N° DI7925-DI7932 European Commission 2022-2024, existing solutions and project outcomes, and the discussions and feedback received from the Subgroup on SPEs of the Community of Practice. This foundation ensures the specifications are aligned with both regulatory and practical implementation needs.

A key starting point was the thorough examination of the EHDS Regulation, especially its Article 73 on secure processing environment (SPE) which underpins Task 7.4 and provides essential guidance for the work. This analysis was further enriched by a detailed review of existing SPE solutions and related project outcomes, ensuring a holistic approach to requirements and capabilities.

Several past and ongoing projects are addressing various aspects of SPEs, providing valuable insights that were analysed to extract relevant requirements and identify reusable approaches. Task partners were requested to suggest relevant projects. The examined projects were primarily selected for their assessment of existing SPEs and their surveys of current infrastructure providers, which eliminated the need for a new survey in this task. Additionally, some projects focus on developing a blueprint for SPEs or are piloting the infrastructure. Including these projects in the evaluation was crucial to ensure alignment among them. These projects are presented in Annex 8: Existing solutions for secure processing.

The major contributors of Task 7.4 analysed the collected material and conducted in-depth research on various topics according to their specific areas of expertise and interest. The findings were then presented and discussed with all task partners.

Following the analysis of project results and current requirements for existing SPEs, the next step was to identify gaps and areas for improvement within the current landscape. These results were utilised to establish the minimum functional, operational, security and interoperability requirements for future SPEs.

A draft version of this document was in public consultation in September 2025. A summary of the stakeholder feedback collected during the public consultation is available in Annex 2: Public consultation summary.

Enterprise Architecture

The analysis of SPE environment in this report is presented using The ArchiMate® Enterprise Architecture Modeling Language[19] version 3.2 as implemented by open-source Archi[20] application version 5.6.0.

[19] The ArchiMate® Enterprise Architecture Modeling Language https://www.opengroup.org/archimate-forum/archimate-overview
[20] Archi https://www.archimatetool.com/

ArchiMate® has not been widely used in biomedical context despite its advantages over other graph generating approaches. ArchiMate® offers a rich selection of conceptual elements that cover from high abstraction level motivation elements through strategy, business, application levels down to technology. The richness of this vocabulary and restricted relations linking elements together form an expressive graphical framework enforcing the separation of abstraction levels.

Annex 1: Figure 1 presents a simplified metamodel highlighting the business-level elements (in yellow), which represent real-world user-facing functions. These elements are linked to their strategic capacities (orange) and motivational drivers (purple), reflecting how real-world services relate to abstract capabilities and regulatory motivations.

In architecture terminology, this report focuses on determining the strategic level minimum requirements and describing the capabilities needed for SPEs in different implantation scenarios.

Annex 1: Figure 1. A simplified ArchiMate® metamodel showing the elements used to show how motivational elements (purple) influence strategic elements (orange), as abstract representations of services and their component roles, processes, and objects (yellow).

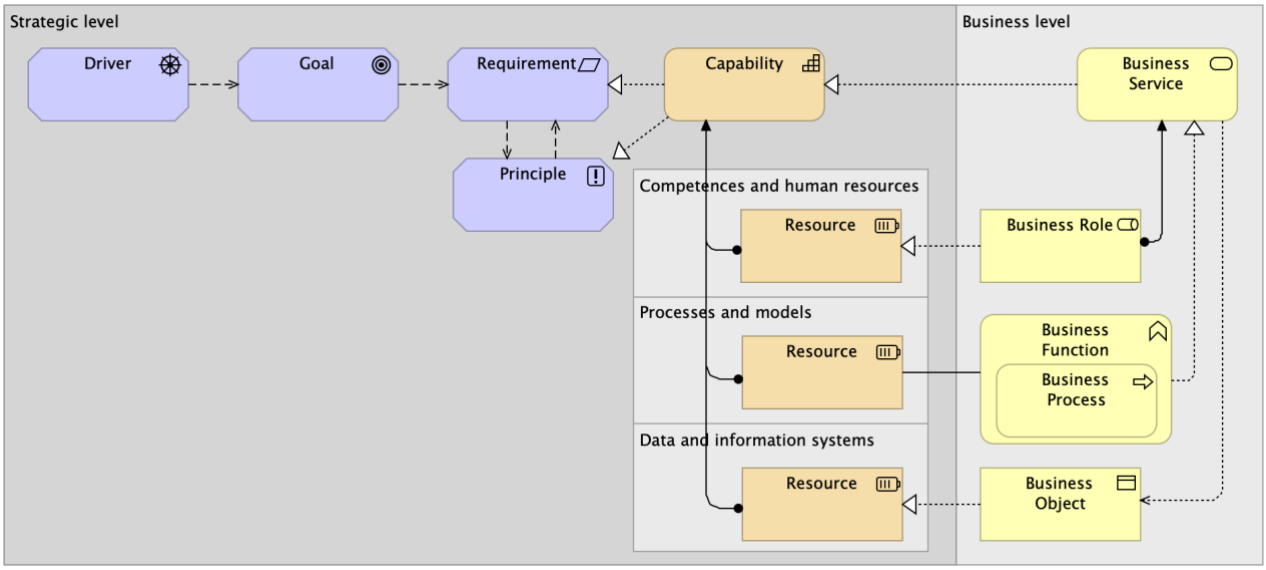

Modality of requirements

Requirements are the a priori conditions that capabilities fulfil. The modality of the requirement is given according to the Internet convention IETF BCP 14[21] as summarised in Table 1.1.

[21] S. Bradner, B. Leiba; BCP14; The Internet Engineering Taskforce Best Current Practice; https://www.ietf.org/rfc/bcp/bcp14.html

Table 1.1. Words used in capital letters to indicate the modality of the requirement.

Adjective	Modal verbs positive	Modal verbs negative
REQUIRED	MUST, SHALL	MUST NOT, SHALL NOT
RECOMMENDED	SHOULD	SHOULD NOT
OPTIONAL	MAY	

Annex 2: Public consultation summary

A draft version of this document was in public consultation in September 2025. This document was commented in total for 80 times. The number of responses may contain some duplicates as there was no individual identification and verification required to respond to the surveys. Some respondents have also responded both from data holder's and data user's perspective. The responses came from 13 different countries from the EU countries and the European Economic Area countries. Responses from Eastern European countries and international organisations were largely missing. The respondents were primarily from three main types of organisations, listed in order of prevalence: public organisations (41%), academic/research organisations (29%) and private organisations (13%).

Summary of comments

The public consultation attracted extensive and constructive feedback from a broad range of stakeholders. Overall, respondents acknowledged the value of the document while also identifying areas where further clarification and development would be beneficial. Stakeholders appreciated the concreteness of the content and recognised that, while further technical specification will ultimately be needed for implementation, the current level of detail is appropriate for a policy guidance document. The document was also commended for clearly identifying key challenges for SPEs and for including annex examples that support knowledge exchange. Terminology was generally considered clear and appropriate.

Several stakeholders wished for more detailed technical requirements to directly support practical implementation and noted that additional specifications would be required. Concerns were also expressed regarding the clarity of compliance and audit expectations.

Respondents further observed that operational guidance is limited, particularly with respect to reference architectures, deployment scenarios, and concrete lessons learned. Some EHDS-related requirements were viewed as infeasible or overly costly (especially the exhaustive logging), with concerns also raised about potential privacy implications. Although these issues were addressed in the annexes, stakeholders recommended that they should be explicitly acknowledged in the main text.

The feedback also pointed to challenges in readability and structure. Requirements were perceived as scattered, repetitive, and at times mixed with recommendations, with insufficient distinction between mandatory and optional elements and unclear allocation of responsibilities. Some terminology was considered misleading, and concepts related to federation were felt to be insufficiently integrated.

Several comments addressed topics beyond the intended scope of the document, including multi-dataset linkage workflows, privacy-enhancing technologies, synthetic data generation, multi-authority coordination, and payment models. These topics are addressed in other TEHDAS2 deliverables.

Report revisions

Based on the feedback received through public consultation, the document has been substantially revised and strengthened in several key areas. The revisions focused both on improving clarity and structure, and on ensuring that stakeholder comments were meaningfully incorporated into the final specifications. These have not changed the main outcomes of the report.

The introduction has been expanded and refined to better express the aims of the report and its scope. Additional context has been provided to explain how the document should be used and how it relates to other deliverables.

The progression and interdependence of requirements have been clarified. Fundamental drivers are now clearly discussed in relation to GDPR, and the generic SPE section does not refer to aspects related to the EHDS Regulation. A new explanation has been added describing the logical progression of requirements and how they are interconnected. Requirements for export control and reporting of clinically significant findings that are unique to EHDS SPE are now mentioned in the main report.

All requirements are presented in tables that follow the same overall structure. All requirements deriving from EHDS Regulation now have explicit role allocation, directly addressing requests to clarify responsibility for each requirement.

In response to specific comments, several clarifications have been introduced across the document, focusing mainly on improving the interpretation of requirements and refining terminology.

The comments received during the public consultation on operational requirements primarily highlighted areas where additional clarification was needed, ranging from the general description of roles to the specifics of individual requirements. As a result, the operational requirements were refined primarily to provide added clarity, ensure proportionality and feasibility, and explicitly address the regulatory, legal, and security considerations highlighted by the consultation.

Adjustments included reinforcing connections to the data permit, clarifying credential and access management, specifying the minimum scope of auditable and compliance-relevant events, and strengthening the rationale for security, risk management, and software integrity practices. These revisions aim to ensure that the operational requirements offer clear guidance while remaining practical, proportionate, and aligned with EHDS objectives and best practices in information security.

Several adjustments have been made to the federated computing requirements (FCR). Requirements relating to federated computing are now clearly indicated to be non-mandatory, reflecting the feedback received. While several comments recommended providing more detailed definitions for the implementation of federated learning, it was determined that this level of detail falls outside the scope of the current deliverable. Additionally, feedback regarding the definitions of key terms has been addressed; the definitions of federated computing, federated analysis, and federated learning have been clarified, and the rationale for these concepts from the EHDS perspective has been strengthened.

Furthermore, various changes have been applied to the technical interoperability requirements (TIR). These requirements now reference industry standards and make clear that specific technologies mentioned are intended solely as examples. Importantly, the technical interoperability requirements have been revised so that they are applicable to all SPEs, rather than being limited to those involved in SPE federation.

Several changes have been made to improve the overall structure and coherence of the report. The technical interoperability requirements (TI) have been repositioned to follow the EHDS SPE operational requirements, while the federated computing (FC) requirements are now presented as a subsection of the SPE federation chapter (FSPE) to better reflect their close relationship and avoid fragmentation of related content.

Annex 3: User journey

When a data user applies for electronic health data for secondary use purposes, such as research and innovation activities, education, and policy-making, within the European Health Data Space (EHDS), the user journey consists of several stages (see Annex 3: Figure 1). Access for certain purposes (public or occupational health, policy-making and regulatory activities, and statistics) is reserved for public sector bodies and Union institutions (see Chapter IV, Art. 53(1) and 53(2)).

Annex 3: Figure 1. EHDS user journey consists of five main phases: data discovery, data access, data preparation, use of data and finalisation.

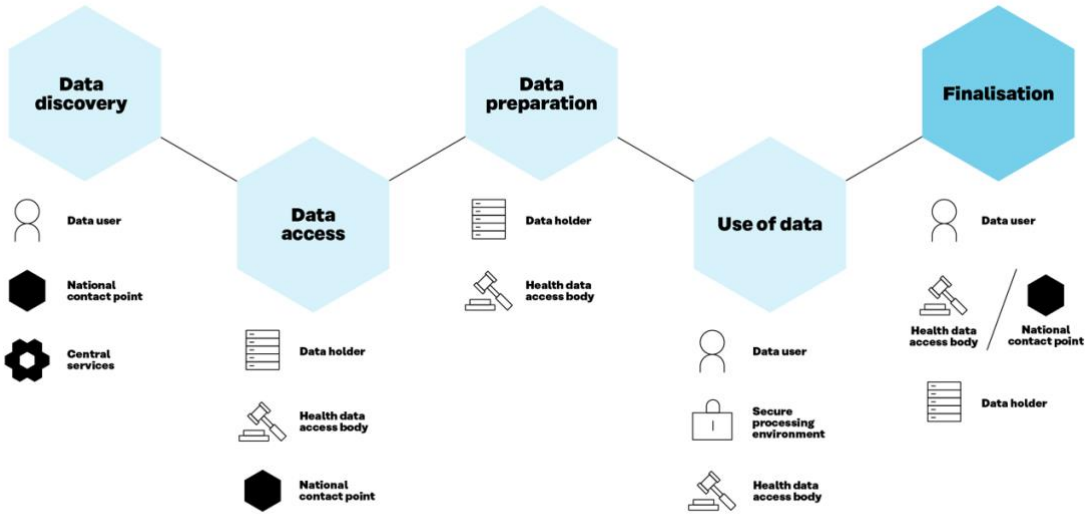

Data discovery

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

Data access

In the data access phase, the user fills in and submits a dedicated and standardised data access application form or a data request to a health data access body (HDAB). The user must complete the information required in the form, upload necessary documents, and provide justifications as needed.

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

Data preparation

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

Use of data

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

Finalisation

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

Annex 4: Glossary

Table 4.1. Key terminology

Term	Description
Access permit	Machine-actionable data structure that contains the information from data permit in a standardised format that can be securely transferred and acted on by computer services
Anonymisation	The process by which personal data is altered in such a way that a data subject can no longer be identified directly or indirectly. (Directive (EU) 2019/1024 of the European Parliament and of the Council of 20 June 2019 on open data and the re-use of public sector information, Recital 52; EHDS Regulation, Recital 92)
Authorised user	An authorised natural person or legal person listed in the data permit, giving them the rights to process sensitive data inside a secure processing environment.
API	Application Programming Interface
Data permit	An administrative decision issued to a health data user by a health data access body to process certain electronic health data specified in the data permit for specific secondary use purposes based on conditions laid down in Chapter IV of EHDS Regulation. (EHDS Regulation, Article 2(2) point (v))
Data holder application (a software linked to a secure processing environment)	A software application that provides the data holder with secure digital access to the Secure Processing Environment (SPE). Its core functions include facilitating the upload and download of data in accordance with the data holder's responsibilities under the EHDS Regulation.
Data processor	The data processor should handle data exclusively in the manner prescribed by the controller. A data processor acts under the detailed instructions of the data controller only, by processing personal data on his behalf. (GDPR, Article 4(1)(8))

Term	Description
Data user application (a software linked to a secure processing environment)	A software application that provides the data user with secure, computerised access to their workspace within the SPE. Its primary functions include facilitating the upload and download of data while ensuring robust authentication and authorisation mechanisms to prevent unauthorised access.
EDIC	European Digital Infrastructure Consortium
EHD	Personal or non-personal electronic health data (EHDS Article 2(2) point (c)).
EID	European Interoperability Framework
ERIC	European Research Infrastructure Consortium
Federated analysis	A decentralised approach to data analysis where statistical results are computed locally on distributed data resources rather than aggregating raw data centrally. This method enables benchmarking, multi-site research, and collaborative analytics while preserving data privacy and security. Only aggregated insights or summary statistics are shared between nodes, ensuring compliance with data protection regulations.
Federated learning	A decentralised machine learning approach where models are trained and validated on distributed data resources without transferring raw data. Instead, only model updates or gradients are exchanged between nodes, enhancing data privacy and security. This method enables collaborative model development across multiple organisations or devices while maintaining local data sovereignty and regulatory compliance.
Federated processing	A decentralised data processing approach where computations occur locally on distributed nodes rather than being centralised. This method enables data to remain on local devices or servers while

Term	Description
	only aggregated results or model updates are shared, enhancing privacy and security. It is commonly used in machine learning ("federated learning"), analytics ("federated analysis"), and secure data collaborations across multiple organisations.
Federation SPE	Secure processing environment (SPE) that is engaged in federated computing.
FTPS	Secure File Transfer Protocol
gRPC	Google Remote Procedure Call
Health data access body (HDAB)	Member state-designated authority that facilitates the secondary use of electronic health data. HDABs assess the information provided by the health data applicant and decide on health data requests and access applications, authorise and issue data permits, obtain data from data holders and make data available in secure processing environments. HDABs systematically track the data request and data access applications received and the data permits issued. (EHDS Article 55 and Recital 52)
Health data user	A natural or legal person, including Union institutions, bodies, offices or agencies, which has been granted lawful access to electronic health data for secondary use pursuant to a data permit, a health data request approval or an access approval by an authorised participant in HealthData@EU. (EHDS Regulation, Article 2(2) point (u))
High Performance Computing (HPC)	The use of advanced and not commonly available computational infrastructure – such as supercomputers or compute clusters – to solve highly complex and resource intensive computational problems.
HTTP/2	Hypertext Transfer Protocol Version 2

Term	Description
HTTPS	Hypertext Transfer Protocol Secure
JSON	JavaScript Object Notation
JWT	JSON Web Token
OAuth	Open Authorisation
Observational Medical Outcomes Partnership (OMOP) common data model (CDM)	A standardised, common data model (CDM) specification originally developed by the Observational Medical Outcomes Partnership (OMOP) and now maintained by the Observational Health Data Sciences and Informatics (OHDSI) community. It defines a consistent structure and a set of standardised vocabularies for observational health data, enabling researchers to perform large-scale, reproducible analyses across diverse databases.
Pseudonymisation	The processing of personal data in such a way that the data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure non-attribution to an identified or identifiable person. (GDPR Article 4(5))
RDP	Remote Desktop Protocol
Representational State Transfer Application Programming Interface (RESTful API)	An application programming interface used for building scalable and interoperable web services. RESTful API follows the principles of Representational State Transfer (REST), using standard HTTP methods to perform operations on resources identified by URLs. It emphasises stateless interactions, meaning each request contains all necessary information without relying on server-side sessions.
Secure Processing Environment (SPE)	An environment in which access to electronic health data can be provided in following a data permit. A secure

Term	Description
	processing environment is subject to technical and organisational measures and security and interoperability requirements. Specifically allowing access to only those persons listed in the permit, as well as user authentication, authorisation, restricted data handling, logging and the compliance monitoring of respective security measures. (EHDS Regulation, Article 73) Used to mean the SPE service run be the operator and to refer to collective functionalities of SPE user accessible instances.
sFTP	Secure File Transfer Protocol
SLA	Service Level Agreement
SPE project-based user space	Project-based environment within an SPE service that users access. Any shorter versions of this should include either 'user' or 'project' to be clear what environment is meant.
SSH	Secure Shell
SQLi	SQL Injection
TLS	Transport Layer Security
TOMs	Technical and Operational Measures
Trusted Research Environment (TRE)	A research environment that aims to create trusted, auditable access to sensitive data, often under national governance frameworks. TREs are not the same as secure processing environments, which are legally defined in the EHDS Regulation. TREs emerged from the UK health research sector, shaped by community-led principles and structured around flexible, function-based zones.
UDAS	User and Data Access Services
URL	Uniform Resource Locator
Virtual environment	Virtual environment is a networked application that gives users of simulated experience of something that doesn't directly correspond to underlying hardware.

Term	Description
	Commonly used for scalability reasons to maintain simpler environments using vastly more complex hardware.
VNC	Virtual Network Computing
VPN	Virtual Private Network
XSS	Cross-Site Scripting

Annex 5: Historical context and legacy models

The introduction of GDPR at the European Union a decade ago created the need protect sensitive data in way touched all citizens. We are still reeling from the effort to interpret and adapt to it. Sensitive does not mean secret and information is no longer contained in physical documents. The first technical solutions designed for analysing small, structured datasets like statistical information are no longer sufficient. Opening up sensitive data to distributed use brings to the fore the human aspects of these activities that are difficult to contain.

Legacy approach: Secret

The topic of securing sensitive data immediately evokes popular images of clandestine operations or military secrets handled behind closed doors. If that sounds familiar, you already have got the wrong end of the stick. Confidential information is where a huge amount of effort is put to keep information as tightly limited as possible to as few people as possible. In contrast, health data for secondary use is meant to be utilised as efficiently as possible for the benefit of the society at large and specifically for future health care choices. The ideas contained and the implications of findings and decisions are meant to be discussed, debated, and published as widely as possible.

Unfortunately, the technology and therefore vocabulary that we use obscures the distinction between the two. It must be remembered that research on health data has always been done, and it is being done at many different approaches, situations, and purposes. A well-intended regulation that is guided by oversimplification can lead to serious loss of institutional knowledge.

Legacy approach: Physical isolation

Confidential legacy is also strongly linked to paper-based approach in information handling. The powerful imagery of red "TOP SECRET" stamped to the corner of the paper permeates popular imagination. Most European national legislations only started changing from data-based definition of official information in the 1990s and the effects of that are still seen in practice. The idea of a physical entity holding specific information sits fast.

In the brave new world of electronic data storage and communication, the realities are quite different. Information can be searched and accessed much faster, flexibly and scalably. Secure storage requirements are conceptually different and need abilities to back up and recover. Sharing, collaboration, and tracking of changes can be automated. Data security becomes paramount to preserve data integrity, prevent unauthorised access, allow auditing and tracking.

These are the very challenges that carefully designed SPEs and supporting services face when they strive to combine security for improved efficiency, productivity, and collaboration.

Legacy approach: Statistical analysis of registries

The first implementations of electronic environments tended to focus on statistical analysis of large amounts of social and demographic information from national registries. These analyses have been typically done using standard statistical packages like SAS along with R. The source data are mostly ordered columns of textual and numeric values that require modifications to conform to the needs of the analysis appropriate for the question.

In other kinds of research, data types can be more varied, datasets significantly larger, the duration of the study much longer, and types of analysis more complex with a need for the development of completely new algorithms and software implementations. For these, the single closed environment created for the purpose of one analysis is a poor fit. A long-term clinical study or genomic study of a rare disease does challenge the boundaries of this approach. They need a more dynamic and distributed environment. That is the current problem of defining how to allow for a controlled and secure ecosystem that combines secure and precise communication to sensitive data services. We will have to draw from the research and practices of large-scale federated data management and combine these to the special needs of secondary use of health data in the context of EHDS.

Technical solutions can never be completely secure on their own

Development of sensitive data processing computer environments is driven by increasing security requirements. Although proper risk assessment is supposed to balance the impact of a threat with its probability, there is a clear pressure to keep increasing limiting security features to services. Adding them have the tendency to increase the cost-of-service maintenance and reduce the user experience. The usefulness of new security measures must be considered against the least secure existing aspects of the service.

In SPE, the graphical interface of the computer that allows the authorised health user to see unprotected health data is that element of least security. It is also the most important for users. Any technical security measures that users do not see the benefit of or make the routine use of the service harder for the user discourages them to use the service at all or prompts them to invent new ways to circumvent these obstacles.

This implies correctly that non-interactive means of accessing sensitive data can be more secure than interactive ones. This builds a case for query-based sensitive data services where queries and replies can be unequivocally recorded.

Together, these mean that regulation of authorised users will have to build largely on trust and accountability, i.e. logs. Trust is based on the health data users' willingness and ability in their professional capacity to follow the data access permit. Clear separation of internal and external threat issues and their countermeasures should solve many governance and scalability issues these services face.

Interoperability between SPEs will have a major unifying impact on services

The only way secure sensitive data interoperability can be ensured is through shared user and project identity. Only with these, can usage and the flow of information be tracked across a network of services.

The establishment of secure and trusted networks of communication between services will make evident the benefits of sharing resources. Many of the resources that SPEs need are available from the Internet but such resources need additional assurances to clear them for sensitive data and a secure communication layer. These in place, these resources will start serving a larger community of services. A prime example is software libraries for commonly used programming languages.

Annex 6: Sensitive data life cycles

The security of data processing inside SPE is tightly dependent on data transfer and supporting services. **Sensitive data access management and interoperability requirements** link SPEs to wider sensitive data cycles (Annex 6: Figure 1) that are common to all sensitive data types and uses, including sensitive intellectual property and consent-based data processing.

This means that the whole ecosystem where SPEs function must be designed to support modularity and flexibility in order to accommodate different governance models and data flows (Principle 1.3.1. Flexible).

Both sensitive data management and interoperability depend on reliable and universally applicable **identity management**, which must extend to all participants and services in an interconnected environment. All these services also rely on stable and secure **network connectivity** to ensure availability and communication.

The generic, high-level use of SPEs (Annex 6: Figure 1) has three use cases for sensitive data processing:

1. Primary use
2. Secondary use
3. Preparation of transient datasets for secondary use

The primary use might use SPEs for protecting the processing, even in the absence of regulatory obligation. SPE use for safeguarding sensitive data for general use is a major driver for SPE design. SPE should be a convenient way to fulfil the general GDPR and national legislation requirements of due diligence when wanting to protect any kind of sensitive data.

In the secondary use of sensitive data, the datasets must have descriptive and governance metadata that make them findable and enable the decision-making process of granting access to an identifiable dataset using a data permit.

When the permitted data is not handed over to the user exactly as it is in the data repository, like when data is first pseudonymised, combined or minimised, the dataset preparation adds another processing step (Transient sensitive data management) before it is ready for the user (TEHDAS2 D7.2 Draft guideline on data minimisation, pseudonymisation, anonymisation and synthetic data).

Annex 6: Figure 1. Key capabilities of sensitive data life cycles.

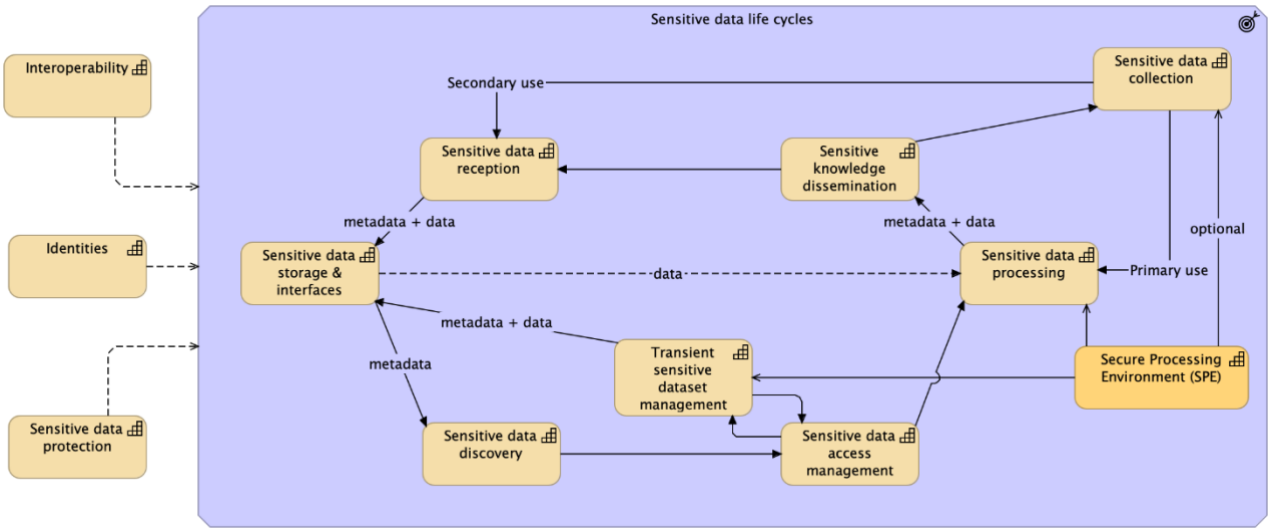

To be maintainable and economically feasible, the specific requirements in EHDS for the secondary use of health data need to fit into these wider SPE ecosystem requirements.

Annex 6: Figure 2 presents these sensitive data life cycles in a more commonly used linear use case format for easier reading. Capabilities are here directly linked to actors with role names and the functions they perform in EHDS context.

Annex 6: Figure 2. EHDS sensitive data life cycles in linear format with actors.

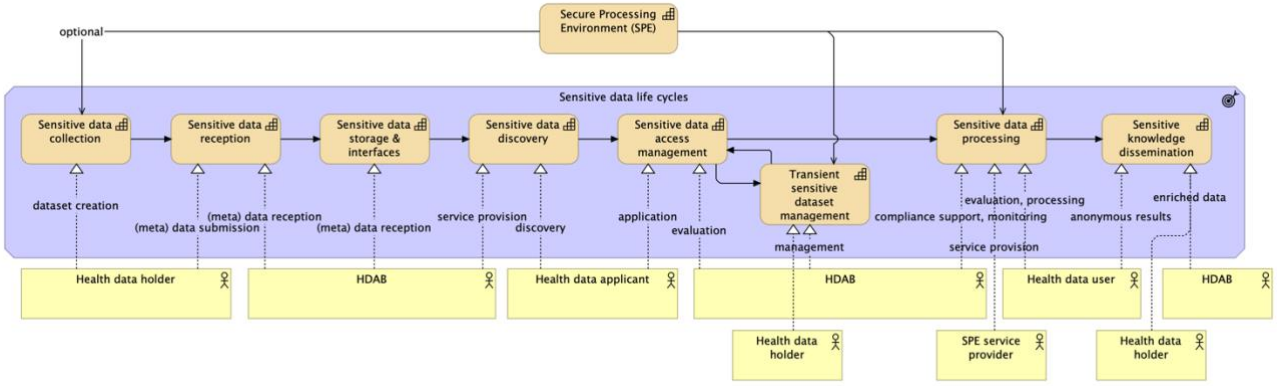

Annex 7: Design considerations and expert commentary

Functional requirements of the SPE cannot be fully separated from the sensitive data life cycles ecosystem (Annex 7: Figure 1). Some of its functional requirements must be shared with the whole ecosystem. This becomes clearer when we consider the needs of data transfer and interoperability that arise from the needs of distributed and federated computing.

In architecture, a stand-alone SPE concept is abstracted to a capacity that has resources (Annex 7: Figure 1). We will follow the order of active resources to discuss the various aspects of SPE functionalities.

Annex 7: Figure 1. Resource map of SPE capability.

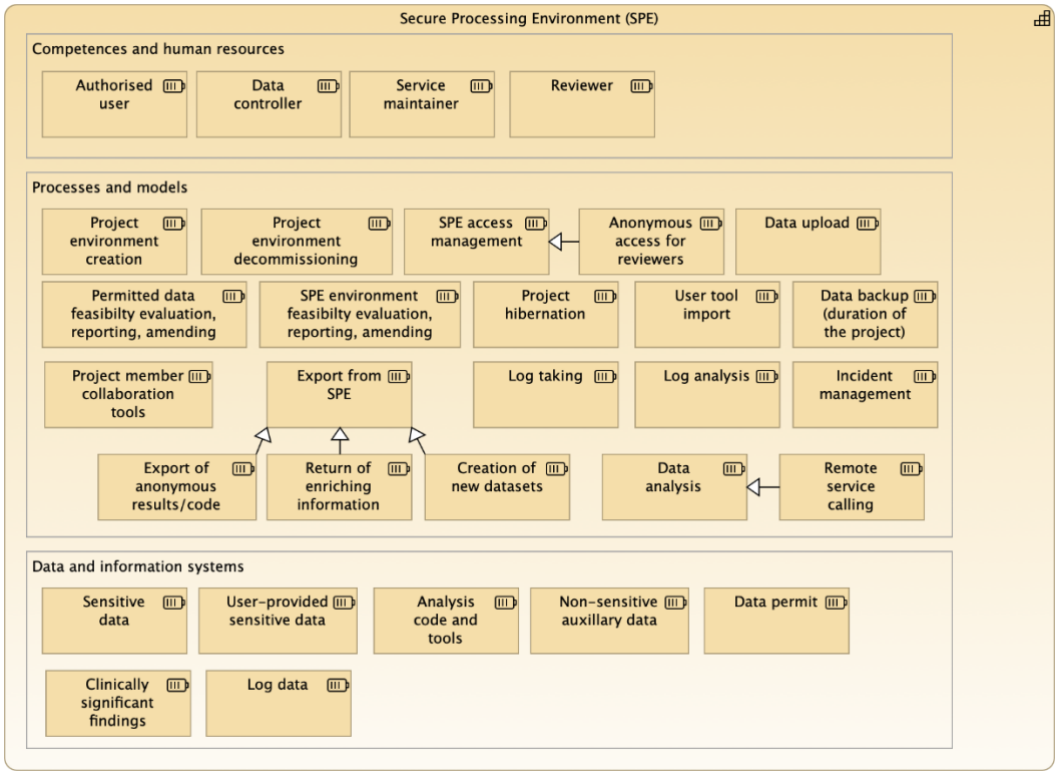

Identity and authorisation

The absolute first requirement of any secure system is that all its users are uniquely and reliably identified (see **SDR-1**) as stated in the EHDS Article 16. That requirement includes that the chain of trust created by this identification is carried uninterrupted throughout the data application and use.

EHDS needs national and cross-border authorisation of users on a scale not yet realised within the EU. Therefore, its success will depend heavily on the widespread adoption of European electronic identification system with its cross-border eIDAS network connecting national electronic IDs and their efficient use in the European Digital Identity (EUDI) Wallet.

The use of eID for sensitive data processing should require high level of assurance that is equivalent to initial registration of turning up in person and authenticating with an official document[22].

Implementing acts and secure protocols underlying the EU Wallets are still under development, raising a possibility that EHDS services will initially need to allow the use of alternative authentication methods that should not be lower than the aforementioned assurance level. Also, the aim of the EUDI Wallet system is to provide identity services first for official uses and the needs of education and research are only fourth in priority list.

In addition to authorisation, EHDS needs to enable authentication of access to sensitive data with equally high level of assurance. This would allow converting the official, human-readable data permit to machine-actionable access permit that can be used to channel permitted health data to the health data user's project space inside SPE.

In Europe, automated research data access authentication services are offered only by Life Science Login (LS LOGIN) by EOSC and ELIXIR that implement GA4GH passports and visas[23] that are secure protocols able to give remote access to electronic resources.

GA4GH visas allow for additional attributes to be added to the person and the permit (as required in EHDS article 68(1)(d) about needed professional qualifications). Users' ability to access the data could be linked to their affiliation (organisation, position) as well as their abilities to handle sensitive data (see the next chapter Priority of user training). The LS Login currently has "de facto researcher" attribute that ensures that the person has the research status in their current organisation that is not available to EUDI Wallet. The word is out that data spaces will be responsible for implementing the attributes and services they will need to as Wallet verifiable credentials.

It should be a high priority for EHDS to implement these approaches. Luckily, EU Digital Wallet and GA4GH passport protocols are both based on OAuth and OpenID Connect specifications, making the work easier.

Priority of user training

There are three main principles governing the implementation of SPE:

1. Enabling (Principle 1.)
2. Privacy (Principle 2.2.)
3. Accountability (Principle 2.3.1)

Various technical and operational ways ensure the privacy of data and the accountability of its processing. The enabling SPE requirement **SPER-7** separates the sensitive data from classified information and creates the weakest link in the SPE usage: Users in all practical cases are required to see, learn from its details, and reorganise the sensitive data before it can be effectively analysed.

[22] eIDAS https://ec.europa.eu/digital-building-blocks/sites/display/DIGITAL/eIDAS+Levels+of+Assurance
[23] GA4GH passports https://www.ga4gh.org/product/ga4gh-passports/

This has two major consequences. Firstly, users need to understand the responsibility of accessing sensitive data and conversely that understanding needs to be demonstrated to representatives of data controllers before they are given access to the data. Secondly, all security measures directed to data users need to take this weakest link into account. Overdoing measures against authorised users are waste of effort and money, as well as increase the security risk due to user irritation and frustration.

In other words, if data processing is based on users learning new insights, the system must have trust to the users. This needs to be earned by demonstrating ability to shoulder the responsibility based on clearly defined curriculum. The modular SPE certification framework could be formalised to have several modules covering different levels of necessary knowledge, e.g. sensitive data protection, user responsibilities and penalties under EHDS, specifics of the SPE user. The licence could answer part of the requirements from EHDS Article 68(1) points (d) and (e) for health data applicant to demonstrate their qualifications to perform the intended research (see also previous chapter Identity and authorisation).

SPE as collaboration area

The other aspect of the enabling principle is that SPE needs to promote communication between project members to do the data analysis (**SPER-5**). The requirements of sensitive data demand that all that communication should happen within the SPE. A prudent principal investigator would immediately set up a shared file that project members will be using as the project logbook. That would allow project members to know exactly what is happening in the project without the risk of revealing sensitive project data to outsiders.

We need to take that approach even further: SPE should be designed to enable information transfer and communication among project members. While each researcher needs their own private space within the SPE, the default should be that actions and files need to be fully shared. Ideally, all record keeping tools should allow for concurrent editing. This thinking should be extended to ongoing analysis runs so that project members working in different physical locations and time zones stay up to date with the project. The SPE should offer these tools. Similarly, the immediate communication in the form of chats, and phone and video calls should happen within the SPE when the technical capabilities make that possible, like in generic virtual desktops.

From this point of view, the SPE is a new kind of social experiment in computing. The isolation of the SPE recreates the open and protected environment of isolated mainframes of the late 1960s, when file and user project space protection was in infancy, but enthusiastic users explored the potentials new technology in. This era also saw an explosion of new ideas getting implemented and information was freely exchanged. The roots of the open-source movement are in those times.

We are still so early in developing the SPE concept that these ideas have not yet turned into guidelines or implementations[24].

Finally, the project logbook is a good example of the privacy problems of sensitive data processing. It allows project members to focus on data analysis and use freely identifying

[24] Lehväslaiho, H. 2025-06-09 Secure Processing Environment in search of a metaphor. https://research.csc.fi/2025/06/09/secure-processing-environment/

pseudonyms in their data. The accountability of their work depends on the accuracy of their log keeping. This raises open questions about how accountability and anonymisation obligations intersect. Further guidance will be needed on retention of project-level collaborative metadata.

Data analysis

Data processing involves the software and frameworks provided within the SPE that enable data users to manipulate, transform and extract actionable insights from data. These tools enable the SPE to fulfil its core function, ensuring compliance to secure sensitive data processing rules and protecting against threats.

By default, in an isolated multipurpose SPE that is based on the secure virtual desktop concept, the data analysis itself should not be that different from normal desktop computing -- apart from lack of direct and open network access.

SPE environments that are built for specific analysis using pre-installed software for specific purposes, based for example on Jupyter Notebook, can be more straightforward to manage and run.

Tilting the balance towards more advanced technologies is the reality that building SPEs and their supporting services is an expensive and time-consuming effort. The cost is so high that it would be unrealistic to build a low technology solution for EHDS and develop and maintain simultaneously a higher-end SPE for other purposes. Low-end solutions are quick to set up but have higher running costs with more error-prone human steps. Deploying low-end SPE solutions may lead to higher long-term costs and reduced scalability, making upgrade paths a critical consideration for EHDS infrastructure planning.

This is where we move the focus of this treatise from abstract to more concrete. Different use cases and implementation choices will strongly affect the flexibility, scalability and capabilities of SPE-based sensitive data processing.

SPE environment management

The definition of an SPE regardless of the use case is based on isolation and accountability. The isolation extends to the separation of roles to manage the SPE environment separately from data access and other services maintained within the same organisation. Additionally, it requires that data from different projects are not mixed at any point and a new SPE instance is launched for each of them.

The trusted use case: HDAB

The trusted use case (see chapter User stories) of SPE gives the user the freedom to select the environment and modify its details to meet their needs. For the HDAB, this means that the tools that the HDAB data managers need can be largely predicted and pre-installed to the environment where the HDAB data managers bring in the source datasets provided by the health data holders, process them, bring them out of the SPE, and place them for

importing or streaming into the distrusted use case SPE for processing according to the data permit.

The distrusted use case: Health data users

In the distrusted use case of EHDS, various limitations and restrictions are imposed on the health data user. Under EHDS (Article 61), the HDAB designates the SPE for the health data user based on availability and user needs expressed in the data access application form. The health data user is required to pre-declare their own datasets containing personal data needed in the environment. Depending on the SPE setup, some needed resources (applications, libraries, datasets) will be provided into the SPE environment by the SPE provider (Article 67(2)(i)). In any case, users will need to import various kinds of files into their SPE project space. In an ideal case, there should be no limitations for user imports, but this can depend on the SPE setup (see chapter Cybersecurity of SPE infrastructure).

The SPE operators might want to agree on minimal security procedures that will allow them to share the burden of preparing them. Annex 7: Figure 2 gives an overview of a proposal under preparation.

To summarise, security measures should strike a balance between protection and usability. They must not impose unnecessary restrictions or unnecessarily impede researchers' work. A streamlined process is needed to ensure tools can meet users' needs in a fast and efficient way.

Annex 7: Figure 2. Resource-checking procedure for SPE operators, currently a draft best practice under development by the SPE Community of Practice subgroup (personal communication, Miikka Kallberg, CSC, FI).

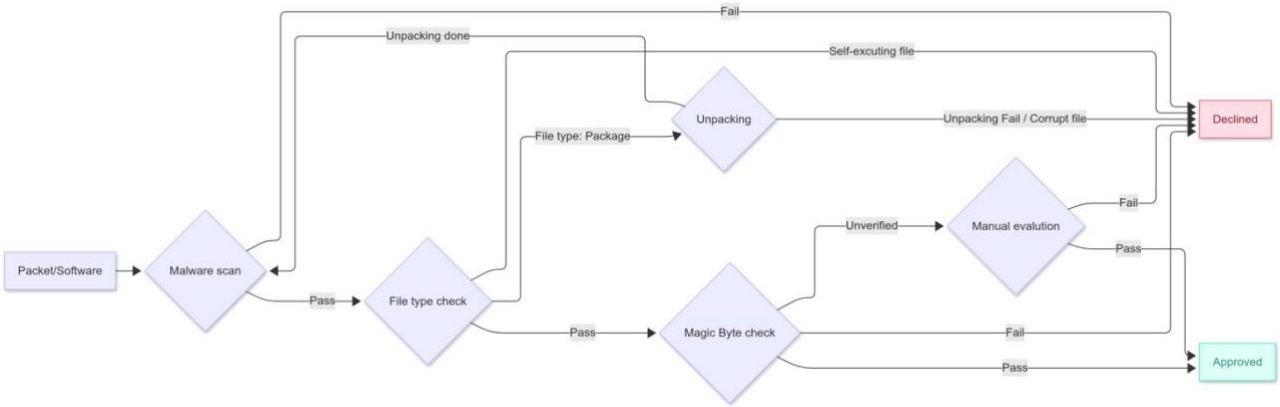

The first task for users getting access to a new SPE is to evaluate its appropriateness to its intended use. Data controllers, represented by the HDAB, are responsible that the provided data is according to the data permit. The SPE provider should be approached in technical questions about the environment. There should be written instructions for most common problems (as detailed in the chapter Priority of user training).

The goals, approaches, and timetables of a project tend to change. Some of these changes will need changes to the data permit. Managing the timely onboarding and offboarding of researchers presents a significant security and operational challenge for projects, which can be addressed effectively with automated identity and affiliation management. It would be convenient if the management of the list of authorised participants could be amended directly by the principal investigator.

The specifics that are mentioned in the data permit will need to be passed on to the HDAB for recording or approval. The principal investigator of the project carries the main onus that official requirements are fully followed (EHDS recital 62). The requirement of destroying user enriched data from SPE at the end of the project (after 6-month grace period) will cause requests for long extensions to keep the project alive at least until its main findings are published (see further discussion in the chapter Data export from SPE).

Ending and pausing the user SPE

At the end of the project, or in case of a detected security incident, the HDAB and the SPE Operator will need to be able to cut access to data or SPE.

The health data users also need to be able to control the data and processing for security and cost reasons. The project area of SPE always has two main components: storage and analysis. Depending on the SPE implementation, these can be tightly or loosely coupled. A very loose coupling is exemplified by the system where the secure storage is fully independent of the processing component of which there can be multiple instances accessing data (see Operational SPEs in Europe: CSC SD Services)

Users might need an *SPE backup facility* to safeguard some intermediate results from their temporary storage area into the secure storage.

An extended pause in the activities might put a severe strain on the budget of a project. SPE operators should offer an *SPE hibernation* service to maintain the project in inactive state with significantly lower cost. This same sentiment is expressed in EHDS article 68(12) suggesting that HDAB could store permitted datasets in a system with lower cost to the user.

Monitoring of SPE use

Accountability (Principle 2.3.1) is the way to ascertain that the permitted health data user's actions correspond to the rules for sensitive data processing and specifically to their granted data permit.

Monitoring is the way to implement accountability. User actions are recorded by the SPE operator and monitored for irregularities. However, the legal responsibility of investigating them is on the HDAB. The sharing of executive actions on incidents between the SPE operator and the HDAB needs more clarification, especially when they belong to different organisations.

The DGA's definition of an SPE and again the EHDS legislation stress that all user actions must be recorded and be available for auditing. They are usually seen to be the **access log, the transfer log, and the activity log**. The first two record who accesses what service and

when, and what is being transferred in or out of the service. They are seen as standard procedures for all services and make perfect sense to monitor those for sensitive data processing SPE.

The inclusion of obligation to record user activity within the service is highly contentious. It means that each and every action of the user within a service needs to be recorded, regardless if it has to do with sensitive data processing. Very few existing services are known to implement this. We feel this severely violates the user privacy. Their credentials for accessing sensitive data have already been vetted and approved. They have been given access to sensitive data they applied for, making them the data controllers of the approved data limited by conditions in the data permit in an environment that should protect them and sensitive data from outside perusal in addition to promote collaboration among project members. Collecting unspecified personal data on user actions inside the SPE is the kind of activity that GDPR was written to prevent.

Not only is collecting all user actions a violation of privacy, but it also generates huge amounts of data that is very difficult to analyse. Further, graphical user programmes may do whatever with the data in their memory without leaving any sensible trace in the logs, undermining the usefulness of action logs. The increase in log volume is also significant cost issue.

The task of logging and monitoring needs to balance scalability with accountability requirements. Continuous monitoring tasks can see the most obvious and often unintentional breaches of requirements. The other end of the spectrum arises from the novelty of research results and the changing sensitivity of information over time as reguired by GDPR. The latter is near impossible and time consuming to define. The cost is that arises either from human labour or automated tools is prohibiting.

Rather than focusing on action logs of dubious legality and utility, its efforts should be spent on ensuring that exported information is properly identified and anonymised. The chapter 'Data export from SPE' below proposes that all exported data from the SPE should be kept for an extensive period in case suspicions of their appropriateness or accuracy arises later.

The division of monitoring responsibilities between SPE operator and HDAB should clearly be stated in the implementation act. These cannot be left to be determined separately between each HDAB and SPE nor should they depend and vary on every data permit. SPE provider should do the base monitoring of logs. When an incident is detected, SPE operator should evaluate the severity and pass their significant findings to the HDAB that will take responsibility of the proceedings and ask for more details when needed.

Data export from SPE

Data export from the SPE in the trusting use case is role-based without explicit controls required by the regulation. HDAB staff are trusted to keep sensitive data secure outside SPE. They can be trusted to store and distribute the sensitive data responsibly in encrypted format and pass it on to a secure processing environment for further analysis by health data users.

The untrusted SPE use case under EHDS have three separate cases for data export that have distinct functional requirements for SPE (Annex 7: Figure 3):

1. Export of anonymous results
2. Returning of clinically significant results to the original data holders

3. Creation of new datasets of enriched research data

Annex 7: Figure 3. Export from SPE.

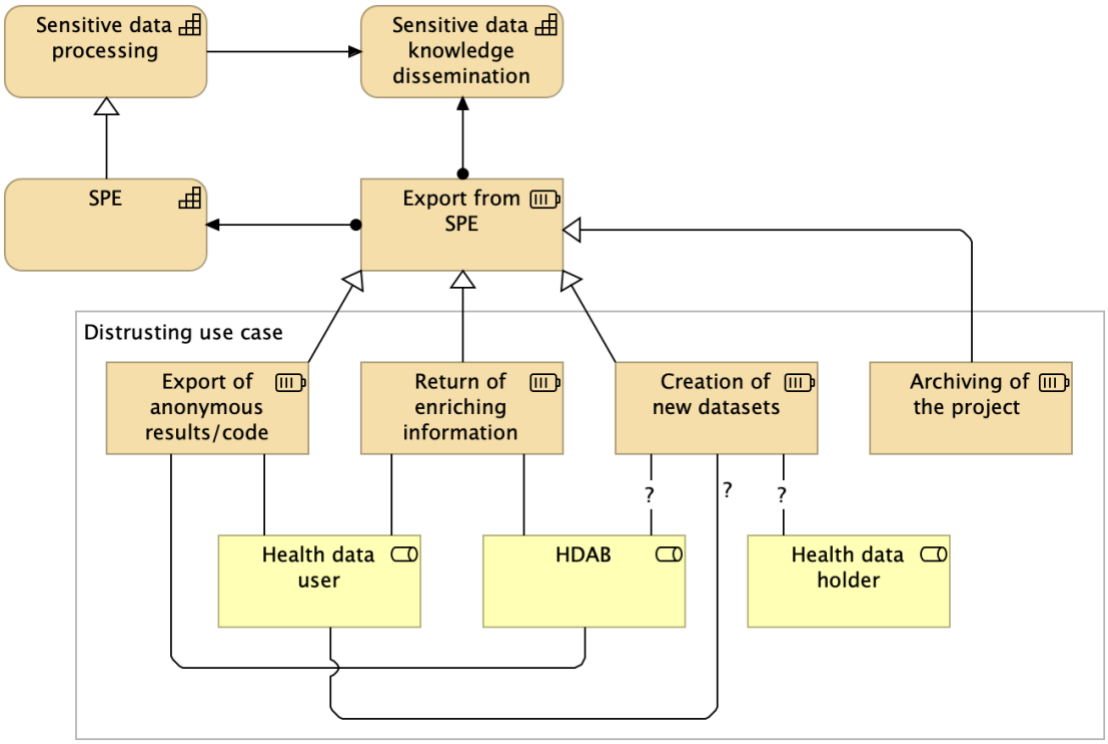

Export of anonymous results

EHDS requires HDABs to verify that exported results are anonymised, but it does not specify the method for performing such verification (Article 73(1)(f)). Member States or HDABs may develop tools or guidelines to support this obligation. The implementing acts under Article 73 may also provide common EU-level requirements or procedures to ensure consistency. Pseudonymisation and non-sensitive data types such as anonymised and synthetic data are covered in TEHDAS2 D7.2 Draft guideline on data minimisation, pseudonymisation, anonymisation and synthetic data.

Export of anonymous results may be divided to two major steps: (1) generation of anonymous results from analysed data and export of it by the health data user and (2) monitoring of the exported files for anonymity by the HDAB (EHDSR-13).

The feasibility of these tasks ranges from straightforward (exported information in code or large numbers of register data in aggregated format) to outright impossible (complex mix of data types).

It should be noted that health data users can be expected to export not only final analysis results for publication, but also intermediary results and the frequency and volume of these ones can be high.

Most obvious mistakes in anonymisation can and should be semi-automatically detected before the export is allowed to happen. Tools and guidelines are currently being developed (see e.g. Anonymity verification tool (Finland), Handbook on Statistical Disclosure Control for Outputs[25], and the PHASE IV AI project deliverable 5.3[26]). Related efforts include the SACRO-project[27], which explores semi-automated checking of research outputs to identify potentially sensitive content.

Ultimately, the requirement of checking all output from open-ended scientific research, even by reviewing, can be extremely difficult. HDABs would need personnel who are experts in all health research domains who would spend long time trying to understand the risks of exports.

A more scalable approach would be to focus on helping researchers to anonymise their results effectively by training and automated feedback tools and store all output results as if they were logs of activities. Analysis of results anonymity would be a part of the required continuous activity of HDABs that can be subjected to more thorough reviewing in case of a security concern.

Returning of clinically significant results

Clinically significant findings (EHDS Article 65(5)) are an exception to the rule that health data users can export only anonymous data from the SPE (Annex 7: Figure 4). While the process is to export out of the SPE, sensitive data is not seen to leave the EHDS secondary use domain as it is to be transferred securely to the HDAB who has the responsibility to inform the original data holders for the ultimate purpose of letting the citizen to know. The details of this process will be decided by member states.

Still, from the SPE service implementation perspective, this is an export of personal information from the closed environment that differs from other exports only in having a different recipient.

[25] Handbook on Statistical Disclosure Control for Outputs https://securedatagroup.org/guides-and-resources/sdc-handbook/
[26] The Phase IV AI project D5.3 https://www.phase4ai-project.eu/wp-content/uploads/2025/09/D5.3_-Data-Hub-Design-and-Data-Market-v1_PU.pdf
[27] SACRO https://dareuk.org.uk/how-we-work/previous-activities/dare-uk-phase-1-driver-projects/sacro-semi-automated-checking-of-research-outputs/

Annex 7: Figure 4. The EHDS outline for the procedure for clinically significant findings from SPE.

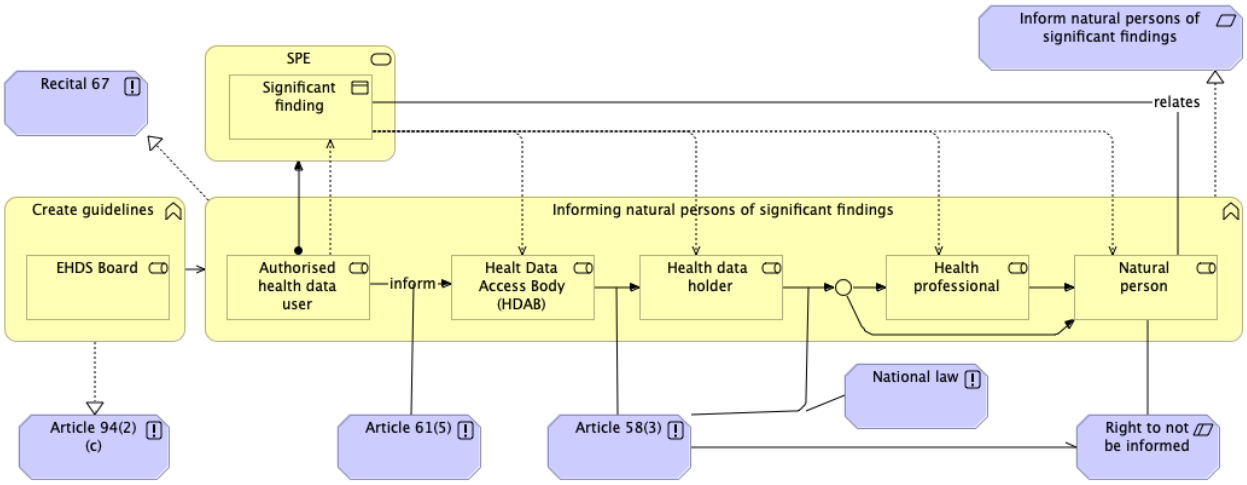

Creation of new datasets of enriched research data

EHDS does not give any guidance on reusing datasets cleaned and enriched by health data users, and this is generally seen to be outside the basic EHDS use case. Enrichment is mentioned in the EHDS recitals.

Two main arguments may be raised in favour of reusing of datasets enriched in research. The effort of cleaning, harmonising and combining of research data for analysis is often the most time-consuming part of the work. From the practical point of view, throwing this away is this massive waste of effort will have to be again for any updated analysis. At a more philosophical level, this goes against the principle of sharing and the results of enriched data for future studies (Principle 1.1.4. Reusable in FAIR).

The **TEHDAS2 Deliverable D5.4: Short Guide for Data Enrichment for HDABs, Data Holders, and Data Users** will have guidance for returning enriched data to data holders. In addition, member states are recommended to consider enabling national legislations for storing and reusing of their health datasets.

Enriched datasets can be seen to cover two distinct situations:

Enhanced versions of original datasets – This refers to datasets whose quality or structure has been improved (e.g. through cleaning, deduplication, harmonisation), without combining personal data from multiple sources. These may, under certain conditions, be returned to the original data holder within the EHDS framework, especially if no new personal data are added or created.

Newly created datasets resulting from linkage or analysis – These are research outputs combining data at the individual level from multiple sources. Distributing would raise significant concerns about data controllership, re-use conditions, and compliance with the

EHDS secondary use rules. These datasets cannot be exported or reused unless they are fully anonymised or covered by a clearly defined legal exception.

If these ELSI aspects can be solved, the practical solution for restricted access datasets are offered by the European Federated EGA services[28]. Federated EGA requires that researchers depositing datasets are members of research organisations that guarantee their identity and trustworthiness and will act as the broker for re-use requests of datasets.

Scenarios

We will somewhat arbitrarily divide available implementation options to three that we call local, state-of-the-art, and distributed (Table 7.1.) that will enable progressively more capabilities and choices for data processing. Most existing SPE environment solutions already combine their features across these simplistic categories.

Table17.1. Implementation options

SPE Capabilites	Local	State-of-the-art	Distributed
User identity	Local	Federated	Federated
Data protection	Full	Full	Full
Key management	Manual	Automatic	Automatic
Data access	Local copy	Managed, streamed	Managed, streamed
Interface	Virtual machine or fit for purpose restricted access	Virtual machine	Federation services for distributed service provision across organisations
Service type	SaaS	PaaS	

The local scenario represents a situation where the hosting organisation starts from scratch building an SPE from a virtual desktop by securing its operating system and deploys an automation tool inside a secured local network. Users need to be initially authenticated manually and assigned a local identity that will be enabled to a designated virtual desktop. The desktop image will need to be installed with most or all user tools and data need to be copied to the desktop once service providers and possibly users have installed their software products, and it has been secured. The requirement to keep sensitive data encrypted in storage and during transfers causes complications in encryption key management.

[28] FEGA https://ega-archive.org/about/projects-and-funders/federated-ega/

The local scenario can reach a high level of data security but at higher maintenance cost arising from the cost of manual labour. More people involved always increases the administrative cost of minimising human error. It is likely that a manual solution cannot effectively enough secure the SPE user environment to allow users to manage their own runtime software installations, making them more dependent on the service provider. This classifies the service as Software as a Service (SaaS).

The state-of-the art SPE system uses more effectively supporting services. User identification to high level of assurance is outsourced to their host organisation or country of origin (See Identity and authorisation). Access control to the SPE and project data is controlled directly by the data controller. The human-readable data permit is converted to a machine-readable access permit that automatically links permitted users and data to the project. Data is stored independently of its use and streamed by demand to the SPE project space. The lack of error-prone human intervention to SPE management and high degree of data protection by design and by default (GDPR Article 25) allow for Platform as a Service (PaaS), where users are able to manage their application environment themselves (See Cybersecurity of SPE infrastructure).

To future-proof sensitive data processing, we need to carefully think what current sensitive data services can do, what is just becoming reachable based on research activities on public data computation, and what needs to be added to them to make sensitive data processing possible in them.

The first step is clear: We need to widen our thinking of secure processing environment (SPE) beyond a single predetermined virtual desktop. **The concept needs to include all types of processing of special categories of sensitive data** that follows the defined principles, regardless of place or time:

SPE implements sensitive data processing

According to this, the desktop is the place where the user interacts with sensitive data, either directly or indirectly.

For the EHDS health data user experience, this means that the user is allowed to initiate processing tasks that have been audited to be part of the EHDS federation or explicitly mentioned in the data permit.

What this means in practice will need a more detailed look into various processing options and modern cryptographic abilities.

Annex 7: Figure 5 provides one way to classify different processing approaches. The current SPE concept can be detached from physical location by implementing the desktop into ephemeral containers that can be executed in multiple compatible platforms (Serverless compute).

Annex 7: Figure 5. Sensitive data processing approaches.

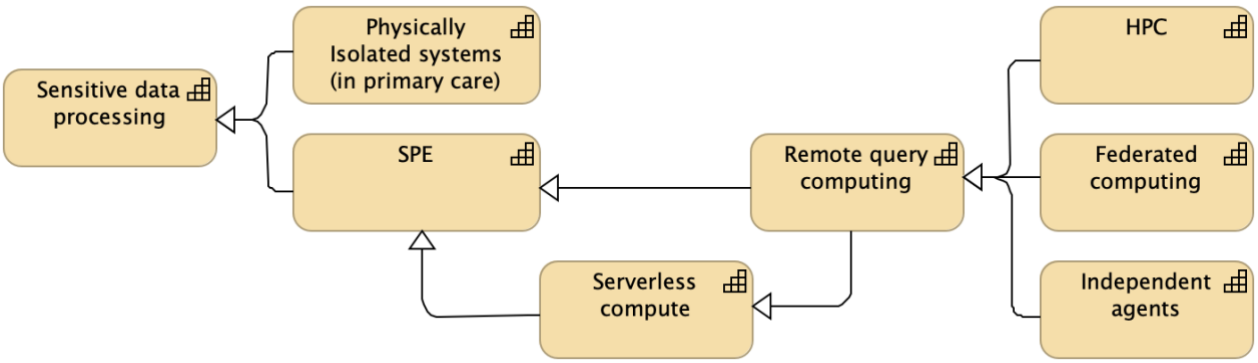

Allowing a user to send an interactive query to a remote server (Remote query computing), can dissociate the user from the externally managed sensitive data. This increases the security but may distance the user from immediate experience of data. The main questions to ask from a remote query are 1) the permissibility of the query and 2) are returned answers to query sensitive.

High Performance Computing (HPC) utilises massively parallel and interconnected computer systems that by design have been built to be shared simultaneously with multiple users. The challenge is to allow for isolated queuing and executing of sensitive data jobs without leaking any personal data. The implementation must ensure that the job submitter has permission to the processed data (see an example of HPC sensitive data processing).

Queries can be automatically sent to more than one target processor and results returned to the sender or the query itself can be the model to be trained, opening increasingly complex possibilities (see Implementing federated computing) with their own requirements.

Finally, while the processing in an SPE may be anything from simple command line tool execution to complex black box of a large language model (LLM) implementing artificial intelligence (AI), the demand for ever more complex reasoning is driving the development for independent and interacting agents that look for and evaluate information (Independent agents).

Annex 8: Existing solutions for secure processing

The Annex 8 gives a listing of SPEs in Europe and brief explanations of some past and current projects. They provide a solid groundwork for SPEs in the EHDS, emphasising a balance between security, computational performance, and interoperability. These insights will guide the development of SPE specifications in TEHDAS2 and EHDS, ensuring that health data can be processed securely and efficiently while complying with European legal standards.

Contents:

Operational SPEs in Europe

Assessing existing SPEs provides a basis for developing technical specifications. Examining national implementations helps identify best practices for EU-wide adoption while also uncovering gaps, inconsistencies, and areas for improvement. This ensures that the specifications address emerging challenges and strengthen interoperability between SPEs.

15 operational SPEs across Europe that process individual-level health data for secondary use have been identified and are listed below (Table 8.1). Environments limited to internal organisational use or specific projects have been excluded. Other relevant SPEs may exist but were not identified within the scope of this analysis.

Table 8.1. Examples of operational SPEs in Europe

Organisation	Country	Name	Access type	Website
The Danish Health Data Authority	Denmark	The Secure Research Platform	Virtual Desktop	https://sundhedsdatastyrelsen.dk/da/english/hea
Social and Health Data Permit	Finland	Kapseli	Virtual desktop	https://findata.fi/en/kapseli/

Authority (Findata)				
The wellbeing services county of Southwest Finland (Varha)	Finland	Atolli	Virtual desktop	https://www.auria.fi/tietopalvelu/en/atolli/index.ht
HUS Helsinki University Hospital	Finland	HUS Acamedic	Virtual desktop	https://www.hus.fi/en/research-and-education/hu
CSC – IT Center for Science	Finland	SD Desktop	Virtual desktop	https://research.csc.fi/-/sd-desktop
Esior Ltd.	Finland	SPESiOR	Virtual Desktop	https://esior.fi/en/spesior/
Statistics Finland	Finland	Fiona	Virtual desktop	https://stat.fi/tup/tutkijapalvelut/fiona-etakayttoja
Health Data Hub	France	HDH Technological platform	Virtual desktop	https://www.health-data-hub.fr/page/faq-english
Central Statistics Office	Ireland	Researcher Online System for Applications (ROSA)	Virtual desktop	https://www.cso.ie/en/media/csoie/aboutus-new/
Statistics Netherlands (CBS)	Nether-lands	Remote access environment	Virtual desktop	https://www.cbs.nl/en-gb/our-services/customise

The Statistical Office of the Republic of Slovenia (SURS)	Slovenia	Remote access environment	Virtual desktop	https://www.stat.si/StatWeb/en/StaticPages/Inde
NHS Digital	UK	Secure Data Environment (SDE)	Virtual desktop	https://digital.nhs.uk/services/trusted-research-e
NTNU – Norwegian University of Science and Technology	Norway	HUNT Cloud	Virtual desktop	https://about.hdc.ntnu.no/
University of Oslo (UiO)	Norway	Services for sensitive data (TSD)	Virtual desktop	https://www.uio.no/english/services/it/research/s
University of Bergen (UiB)	Norway	SAFE (secure access to research data and e-infrastructure)	Virtual desktop	https://www.uib.no/safe

.

TEHDAS1

TEHDAS1 deliverable D7.2[29] provides an in-depth analysis of SPEs as defined in the EHDS legislative proposal, along with guidelines covering technical, information security and interoperability requirements.

The report indicates that the preferred architectural model for SPEs is a decentralised system where multiple SPE providers across different EU Member States can offer compliant environments. This model also supports a federated learning approach, where data stays within national boundaries, but models and insights can be shared across borders. Each SPE would adhere to common EU-wide standards, ensuring interoperability and trust among

[29] D7.2. Options for the services and services architecture and infrastructure for secondary use of data in the EHDS https://tehdas.eu/tehdas1/results/tehdas-proposals-for-the-implementation-of-ehds-technical-infrastructure/

Member States. This approach also facilitates cross-border research while maintaining national control over data.

Additionally, the report emphasises that SPEs are not only security-focused but must also be flexible enough to handle a wide range of computational tasks. This includes everything from basic data analysis to more complex tasks like deep learning and artificial intelligence (AI). Given the increasing reliance on advanced computational techniques in health research, SPEs should be equipped to utilise high-performance computing (HPC), GPUs, and other advanced computing resources. Systems must be scalable to accommodate large datasets and resource-intensive processes, without compromising security.

To fulfil the requirements of the services outlined in Article 73 of the EHDS Regulation, the key functional capabilities for SPEs identified in TEHDAS1 are as follows:

- **Data Processing Capabilities**: Advanced analysis tools to handle sensitive data, including statistical software, AI libraries, and version control systems for code management.
- **Interactive Access:** Secure, remote access options such as remote desktop and secure shell connections. Some SPEs may also offer API-based access to support federated analysis.
- **Strong Access Control**: Comprehensive access management (data holders for data deposition, data users for data analysis, and system administrators for SPE operations)
- **Controlled Communications**: Data imports, exports, and other outbound communications.
- **High Security Standards**: Adherence to stringent security measures to protect data integrity and confidentiality.
- **Defined Operational Protocols:** Clear, well-documented procedures governing the operation and management of the system.

SPEs must also integrate harmonised security measures across the EU, allowing for consistent and reliable processing environments. This approach promotes cross-border data collaboration, with interoperability as a key consideration. To this end, SPEs must support standardised interfaces and protocols for seamless data sharing and access management. This ensures that data can be processed and analysed across different countries while maintaining a high level of security and legal compliance. Aligning security standards with recognised frameworks such as ISO/IEC 27001 and ENISA guidelines is recommended.

Five Safes

The concept of Five Safes – a framework for planning confidential data governance and management solutions has emerged and evolved in the United Kingdom over the past two decades.

The Five Safes framework has been initially conceived to help establish a virtual microdata research data centre at the UK Office for National Statistics in 2003[30]. Since then, it rose to international prominence and has been adopted as governance/data management standard

[30] Green, E., & Ritchie, F. (2023). The Present and Future of the Five Safes Framework. Journal of Privacy and Confidentiality. doi: https://doi.org/10.29012/jpc.831

for introducing (or retrofitting) various confidential data access solutions in the UK, the statistical offices in Canada, Australia, New Zealand, Norway, as well as Eurostat. Moreover, the framework has been built into or has influenced legislation such as Digital Economy Act in the UK and several state laws in Australia.

The Five Safes defines five dimensions of confidential data management: safe projects, safe people, safe settings, safe data and safe outputs. These are often formulated as "key questions" (Table 8.2).

Table 8.2. The Five Safes and key questions (Green & Ritchie, 2023)

Element	Typical question	Example of problems being addressed
Safe projects	Is this appropriate use and management of the data?	- What is the purpose of the access request? - Is this an ethical and lawful use of the data? - What is the benefit to society or to the organisations sharing data? - Is there a data management plan in place? - What happens to the data at the end of the project?
Safe people	How much can I trust the users to use the data appropriately?	- Do the users have the necessary technical skills? - Do the users need training inhandling confidential data? - Are users likely to follow procedures?
Safe settings	How much protection does the physical environment afford to the data?	- How are data stored? - Are there physical restrictions on the users? - Does the IT prevent unauthorised use? - Are mistakes by authorised users likely to be detected?
Safe outputs	How much risk is there in the outputs of the access breaching confidentiality?	- If the aim of access is to produce statistics, is there any residual risk by, for example, showing outliers? - If the aim of the access is to produce data for onward transmission, how do we make sure that the released data are appropriate for the next use?
Safe data	Is the level of detail in the data appropriate?	- Is there sufficient detail to allow the project to go ahead? - Is there excessive data not necessary for the project?

As Green & Ritchie underline, the dimensions are not limits, they are scales. What "safe" means, or how restrictive the safety requirements of each dimension are, is dependent on the context. For example, for open data, only the Safe Data dimension must be controlled. In

a secure compute environment, data is safer, perhaps it needs to be deidentified only, while there is a higher degree of control across other dimensions. This flexibility goes hand in hand with "principle-based" approach to designing governance and legal measures.

In modern understanding, the Five Safes are an aid in identifying structures and goals of a particular confidentiality solution, not a rigid set of guidelines or rules. The right way to approach a design of a system according to the Five Safes is an analysis of the approach to the design problem, then specification of broad principles and aims of the solution, followed by using the Five Safes to provide the structure and, finally, identifying the good/best practices in each area of Five Safes.

Trusted Research Environments

Based on the Five Safes framework, the past two decades have seen the rise and adoption of the Trusted Research Environments (TREs), particularly in the United Kingdom, as a safer computer environment to protect the identity of human genomic sequences and then expanded to cover health and social data. The TREs are now seen as the preferred approach (DARE UK, 2024) to ensure the safety of processing of any sensitive data. While there exists a set of community-defined requirements and a derived high-level architecture, interestingly, TRE does not have an agreed definition. It has been called by many alternative names like Secure Data Environment (SDE by the England National Health Service) or Data Safe Heaven (Scotland), and the TRE specification allows it to be organised in many ways.

The DARE UK Federated TRE Blueprint[31] clarifies the concept by specifying three functional zones that any TRE may contain in any combination:

- **Research Analytics Zone (RAZ)** for project-specific data processing by users
- **Secure Data Zone (SDZ)** for active data management roles of data governance
- **Query Management Zone (QMZ)** to provide secure remote data access services

In an extreme case, a TRE can be composed only of RAZ making it impossible to draw any functional conclusions from the name. Any practical discussion about TRE functionalities should now indicate the zone.

Both Research Analytics Zone of TRE and SPE are based on securing sensitive data privacy within a secure computer environment for the exclusive use by data users. Both isolate users by project and demand clearly defined, limited, and secure interfaces out of them[32] (see Annex 8: Figure 1 and Figure 2).

[31] DARE UK Federated TRE Blueprint https://zenodo.org/records/14192786
[32] Lehvaslaiho H. (2025) SPE and TRE terminology for sensitive data processing https://zenodo.org/records/15696511

Annex 8: Figure 1. Equivalence of SPE and TRE-RAZ.

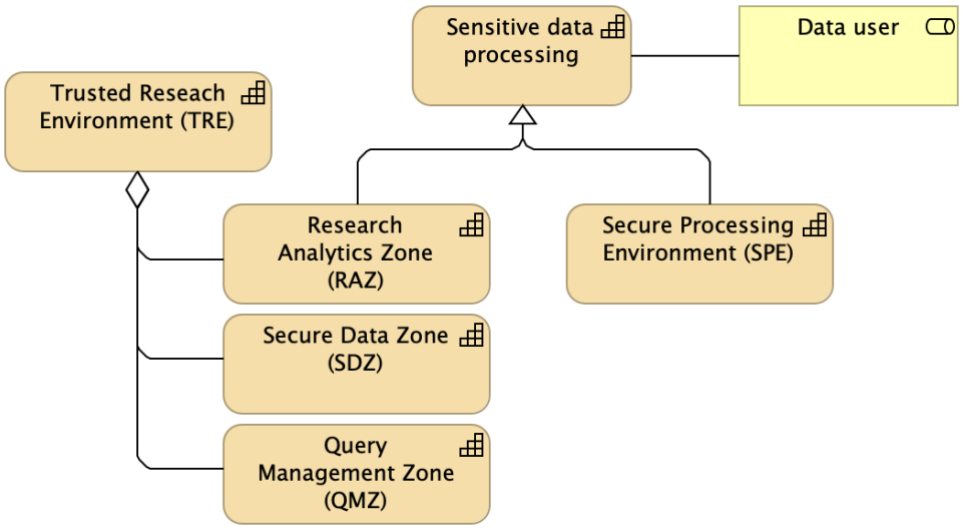

Annex 8: Figure 2. SPE and TRE-RAZ may contain multiple isolated project environments that have only secure and limited access to the outside.

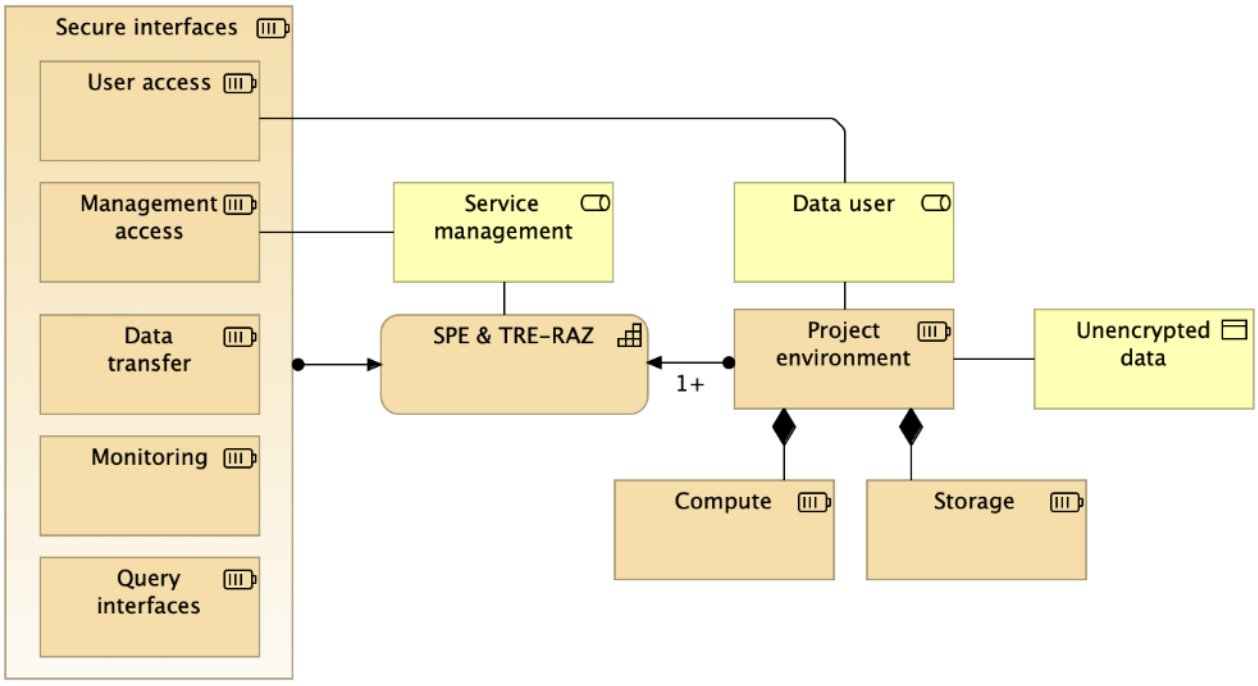

DARE UK

The DARE UK (Data and Analytics Research Environments UK) project is a UK-wide initiative aimed at developing a coordinated, trustworthy, and secure national data research infrastructure. It is designed to support research using sensitive data (like health and administrative data) in a way that is safe, ethical, and for public benefit. In its first phase (2021-2024), it focused on developing the technical and governance foundations. The outputs of the projects[33] were categorised according to four of the Five Safes:

- Safe data: Semi-automated Risk Assessment of Data Provenance and Clinical Free-text in TREs (SARA) aimed at using machine learning to better understand privacy risks in the free-text data.
- Safe outputs: Semi-automated Checking of Research Outputs (SACRO) focused on introducing efficiencies into checking research results for disclosure risk before they leave TRE.
- Safe projects: Maintaining the safety of projects spanning multiple TREs (TRE-FX and TELEPORT projects).
- Safe settings: the SATRE project set out to assimilate the essential features of TREs into a common specification and provide a first blueprint for new TRE builders.

  The currently ongoing phase two of the project has yielded a blueprint for a federated TRE architecture[34]. For this work, of particular interest are the SATRE specification and the federated architecture blueprint.

SATRE

> *SATRE (Standardised Architecture for Trusted Research Environments) started as a DARE UK Driver Project working to standardise access to secure data in trusted research environments, and included University of Dundee, Alan Turing Institute, UCL, Ulster University, and Research Data Scotland. (*SATRE website[35]*)*

A major outcome of the project has been a standard architecture specification for Trusted Research Environments. The specification is pertinent to TEHDAS2 specification as it defined key capabilities of a general TRE and a comprehensive list of criteria that can be used to evaluate compliance against the framework of any given TRE solution.

The four key capabilities – or pillars in SATRE parlance – are[36]:

- Information Governance
- Computing Technology and Information Security
- Data Management

[33] DARE UK (Data and Analytics Research Environments UK). (2024). The 2023 DARE UK Driver Projects: Summaries and lessons learned. Zenodo. https://doi.org/10.5281/zenodo.11443328
[34] DARE UK. (2024). DARE UK Federated Architecture Blueprint (2.2).
[35] SATRE https://satre.uktre.org/en/page/about/
[36] SATRE Specification. https://satre-specification.readthedocs.io/en/stable/

- Supporting Capabilities

The *Information Governance* pillar consists of requirements ensuring information risk is measured and managed to an acceptable level (EHDS Article 73).

The *Computing Technology and Information Security* pillar lays out a set of technical requirements related to the systems used to secure, manage and provide compute capabilities to *Data Consumers*.

The *Data Management* pillar is concerned with managing data assets while they exist (temporarily or permanently) within a TRE. Note that some of the requirements, particularly those that concern data management from the point of view of a data holder are beyond the scope of an EHDS compliant SPE.

The *Supporting Capabilities* pillar is only indirectly associated with a SPE provider. Most of the requirements are related to HDAB itself.

SATRE Roles and their Relation to the EHDS SPE

SATRE Project also defined roles belonging to different stakeholders involved with the data lifecycle[37]. The roles may fit well the *health data permit-based* project concept as data processing roles are tied to natural persons, as well as the kind of expertise needed to run a SPE provider.

The role of a *Data Consumer* aligns with *health data users* as defined in the EHDS legislation[38]. In particular, *Project Manager* maps to the role of Principal Investigator under EHDS (Article 68). The role of *Data Analyst* corresponds to any authorised natural person who accesses the data with the purpose of processing.

The Data Management roles are broadly applicable to organisations which curate the data in addition to providing it in a SPE. In this sense, they may be beyond the scope of an EHDS SPE directly --- except for *Output Checker*, who may need to be listed on a health data permit with the responsibility of ensuring that no confidential data exists in the outputs of SPE project, the data which is supposed to be released out of the secure environment.

The SATRE roles concerned with infrastructure management describe the governance needs of an *SPE Operator*, in order to be applicable to the broadest community of existing SPE providers.

[37] SATRE specification role https://satre-specification.readthedocs.io/en/stable/roles.html#roles
[38] Directorate-General for Health and Food Safety. (2025, 03 05). 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. Retrieved from European Commission: https://health.ec.europa.eu/publications/regulation-eu-2025327-european-health-data-space-and-amending-directive-201124eu-and-regulation-eu_en

EOSC-ENTRUST

Launched in March 2024, EOSC-ENTRUST[39] is a three-year initiative that brings together partners from 15 European countries to advance the development and interoperability of TREs. These environments play a critical role in enabling secure and compliant access to sensitive research data.

At the heart of the project is the development of an Interoperability Blueprint[40], a foundational framework that outlines how sensitive data can be accessed and analysed across a federated network of composable TREs. By addressing both technical and policy-level interoperability, EOSC-ENTRUST supports the broader goals of the European Open Science Cloud (EOSC) and is working under both the EOSC and ELIXIR umbrellas. The project also maintains strong ties with other ELIXIR and European Commission-funded initiatives and is aligned with the EuroHPC infrastructure to explore high-performance computing for secure data analysis.

A key objective of EOSC-ENTRUST is the establishment of a Provider Forum, a community platform for SPE/TRE providers to exchange knowledge, share best practices, and collaborate. This forum is designed to continue beyond the project's lifetime, ensuring sustained support and onboarding of new providers into the ecosystem.

The project is also developing complementary tools and resources, including:

- A set of Training Packages for TRE providers and users,
- A machine-readable TRE Provider Catalogue that maps available secure environments within EOSC, along with their capabilities.

EOSC-ENTRUST's work is driven by diverse use cases, or "Drivers," which include:

- Federated human genomics,
- Social science data sharing,
- Clinical research data interoperability,
- Public-private collaboration involving health and environmental data.

To maximise relevance and impact, EOSC-ENTRUST closely aligns with related European initiatives. In particular, the project collaborates with TEHDAS2 in support of the EHDS, and with the Genomic Data Infrastructure (GDI) for genomics-related use cases. These connections help ensure that EOSC-ENTRUST complements ongoing efforts and contributes meaningfully to a cohesive European data ecosystem.

Throughout the project, EOSC-ENTRUST will also follow the project outcomes of its sister projects SIESTA (Secure Interactive Environments for SensiTive data Analytics) and TITAN (Trusted envIronments for confidenTiAl computiNg and secure data sharing), which are

[39] EOSC-ENTRUST https://eosc-entrust.eu/
[40] Sætrom, P., Lehväslaiho, H., Stansberg, C., Awan, H., & Hesam, A. (2024). EOSC-ENTRUST D13.4 Year one version of EOSC-ENTRUST Blueprint & Interoperability Framework. https://doi.org/10.5281/zenodo.14362388

funded under the same HORIZON 1.3 call on Trusted environments for sensitive data management in EOSC.

HealthyCloud

The HealthyCloud project aimed to lay the foundation for a European Health Research and Innovation Cloud (HRIC) by promoting the secure, ethical, and efficient sharing and reuse of health data across Europe. It brought together a wide range of stakeholders to define a strategic agenda and develop a practical framework to support cross-border health research and innovation, while ensuring compliance with legal and ethical standards.

Within this context, the objective of Deliverable D5.1 (confidential and unpublished) was to identify common patterns in the design and operation of SPEs, focusing on key requirements that could inform the development of reference guidelines for the planned HRIC. The report draws on in-depth assessments of over 13 representative SPEs, supported by interviews with experts familiar with their implementation and operation. These examples were selected from a broader inventory of SPEs across Europe.

The report concentrates on identifying typical models and recurring patterns, along with the associated technical and governance requirements.

The report highlights that the fragmented policy landscape has resulted in a similarly fragmented architectural landscape, leading to the emergence of various SPE designs, each with distinct objectives, including:

- Secure access to controlled data
- Secure collaboration
- Distributed computational approaches
- Data lakes

Among the 13 infrastructures examined in detail, most focus on enabling secure collaboration, while a few focus on secure access to controlled data, and one each provides “compute-to-data” capabilities and data lake functions.

HealthyCloud deliverable D7.4[41], on the other hand, gathered information on security policies and breach response protocols related to the same infrastructures examined in D5.1. This report presents an overview of current practices related to identification of users and access control, data processing, managing and monitoring the environment, as well as organisational policies and procedures. Given that these environments are built and operated in various ways, they each have distinct privacy requirements. Nevertheless, there is a strong consensus on the establishment of SPEs. The key findings are summarised as follows:

- Most infrastructures employ federated authentication and view multi-factor authentication (MFA) as crucial for effective identity and access management.
- Each project should operate within a dedicated environment that is technically and logically isolated from others, grouped by project rather than user.
- Infrastructures managing permits automatically lock user environments once data permits expire and regularly verify the validity of access.

[41] HealthyCloud deliverable D7.4 https://zenodo.org/records/10225422

- For health data processing, pseudonymised data with minimised variables is used, transferred in encrypted form, and must be audited before leaving the infrastructure.
- All infrastructures implement various technical and organisational measures for security management, guided by institutional policies for federated infrastructures.
- Most respondents have internal and external policies for users and staff, with certified sites offering more comprehensive documentation and staff training as a common best practice.

Global Alliance for Genomics and Health (GA4GH)

The Global Alliance for Genomics and Health (GA4GH)[42] develops open, interoperable standards to facilitate secure, ethical, and scalable genomic and health data sharing. These standards provide technical frameworks for authentication, encryption, access control, federated analysis and data governance, which are also key elements for SPEs.

Among these, **Crypt4GH** is a secure encryption standard specifically designed for genomic data. It ensures data remains protected both at rest and in transit through multi-layer encryption, allowing authorised users to decrypt only the portions they are permitted to access. Another standard is **GA4GH Passports**, which streamline authentication and authorisation by providing a federated identity management system. This allows users to securely access SPEs using credentials from their home institutions while enforcing fine-grained access controls.

SPE providers benefit from a user-friendly, secure, interoperable, and privacy-preserving toolkit designed to support the management and processing of sensitive biomedical data.

GDI and 1+MG

The *Genomic Data Infrastructure* (GDI)[43] project is enabling access to genomic and related phenotypic and clinical data across Europe. It is doing this by establishing a federated, sustainable and secure infrastructure to access the data according to the 1+ Million Genomes Initiative[44]

The GDI project places emphasis on federated processing of genomic data which – in case the data in question falls under EHDS – provides a strong driver to solve SPE interoperability and federation in a standardised manner. Additionally, the scale of compute and storage resources needed for genomic data will impact the high-end SPE requirements, for example, driving the high-performance and cloud compute systems used for GDI data to comply with EHDS SPE requirements.

Access to categories of data in focus of the GDI project can be split in three tiers: *open, registered* and *controlled*. These permissions are designed to be machine readable (Ga4GH passports).

[42] GA4GH https://www.ga4gh.org/
[43] GDI https://gdi.onemilliongenomes.eu/
[44] 1+ Million Genomes Initiative https://digital-strategy.ec.europa.eu/en/policies/1-million-genomes

The 1+MG Framework[45] is a series of components based on the output of the 1+MG projects that provide guidance on ELSI, data quality, data standards, and technical infrastructure standards and APIs.

Sensitive Data HPC strategy (CSC, Finland)

SD HPC is the Sensitive Data High Performance Computing solution from CSC – IT Center for Science Ltd., Finland. It is based on the secure transfer of batch job requests from its Secure Processing Environment (SPE) SD Desktop. It contains a batch job management process not visible to other users, secure transfer of user data into a dedicated compute node when the job enters the execution, isolation of the compute node during the execution, and return encrypted results to the SPE user storage space.

The SD HPC service is currently deployed in the CSC Puhti supercomputer. It is in alpha user testing phase.

SD HPC receives job description files from users in the SD Desktop SPE project environment (Annex 8: Figure 3). The HPC system must recognise the user and their project. In writing the job description file, the user must consider the capabilities and limitations of the HPC. The dedicated sensitive data execution manager queues and launches the job in a dedicated node. It also copies user resources from SPE storage before isolating the node for the duration of the job execution. It returns the execution results to the user project environment and cleans any trace of the execution event from the node.

Annex 8: Figure 3. CSC's sensitive data HPC strategy view

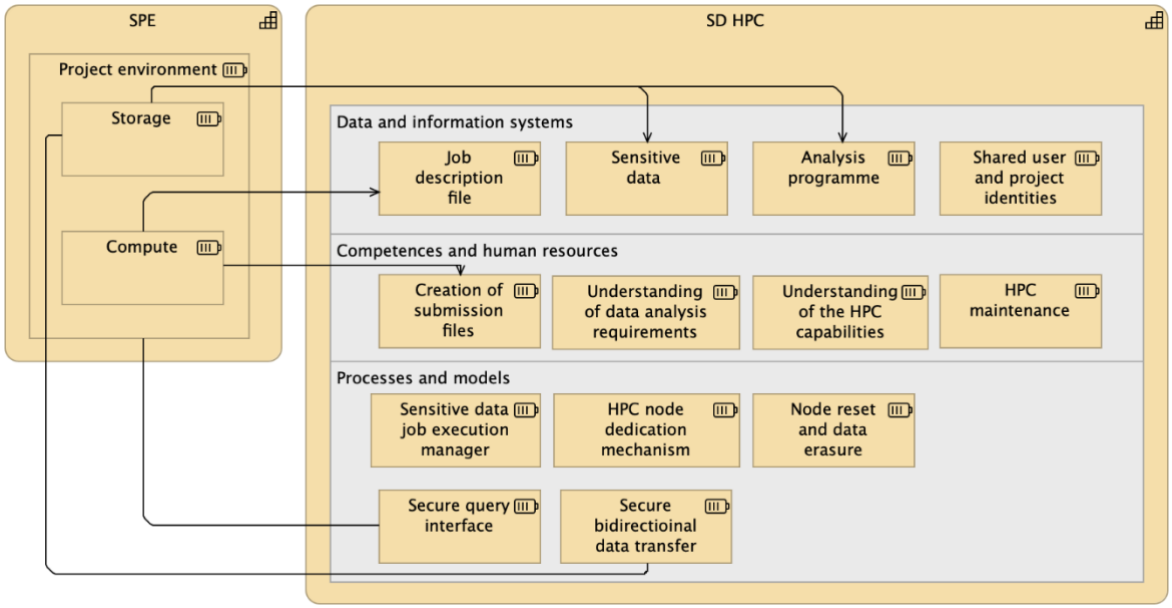

[45] 1+MG framework https://framework.onemilliongenomes.eu/

The job request message from the user is passed through the backend gateways to the main HPC gateway called SDSI server (Annex 8: Figure 4). The acronym comes from Sensitive Data Slurm Interface indicating the used batch job controller. The SDSI server has complete control of the job management once the job request has been submitted to the HPC system. This isolates the client side cleanly from the HPC execution. Expansion of SD HPC to other HPC systems requires only the adaptation of the SDSI server to their specific capabilities and architecture that needs extensive involvement of their system maintainers.

Inside the SD HPC service, the SDSI server formats the job submission to the requirements of the batch job controller (Annex 8: Figure 4). Once the job enters the execution, it fetches the header parts of the user input files from the internal vault component of SD Connect. Before dispatching the headers, the vault re-encrypts the headers with job-specific encryption key according to the SDSI server instructions. These headers are combined with the encrypted payload from SD Connect before passing them into the HPC compute node reserved for the job. The separation of headers and their separate encryption from the payload is an envelope encryption capability of the GA4GH Crypt4GH[46] file container format.

Once all files have been copied to the job node, all access to the node is cut off. Communication to the node is resumed only after the execution is finished, decrypted input files have been removed, and the result files are encrypted with the project key. The SDSI server then copies all result files to a new folder in the SD Connect project space, removes all user files, resets the HPC compute node, and returns it to the pool.

[46] Global Alliance for Genomics and Health (GA4GH) Cryp4GH https://samtools.github.io/hts-specs/crypt4gh.pdf

Annex 8: Figure 4. Service diagram showing the user initiating the SD HPC call. Solid arrows represent message directions and dashed lines file transfers.

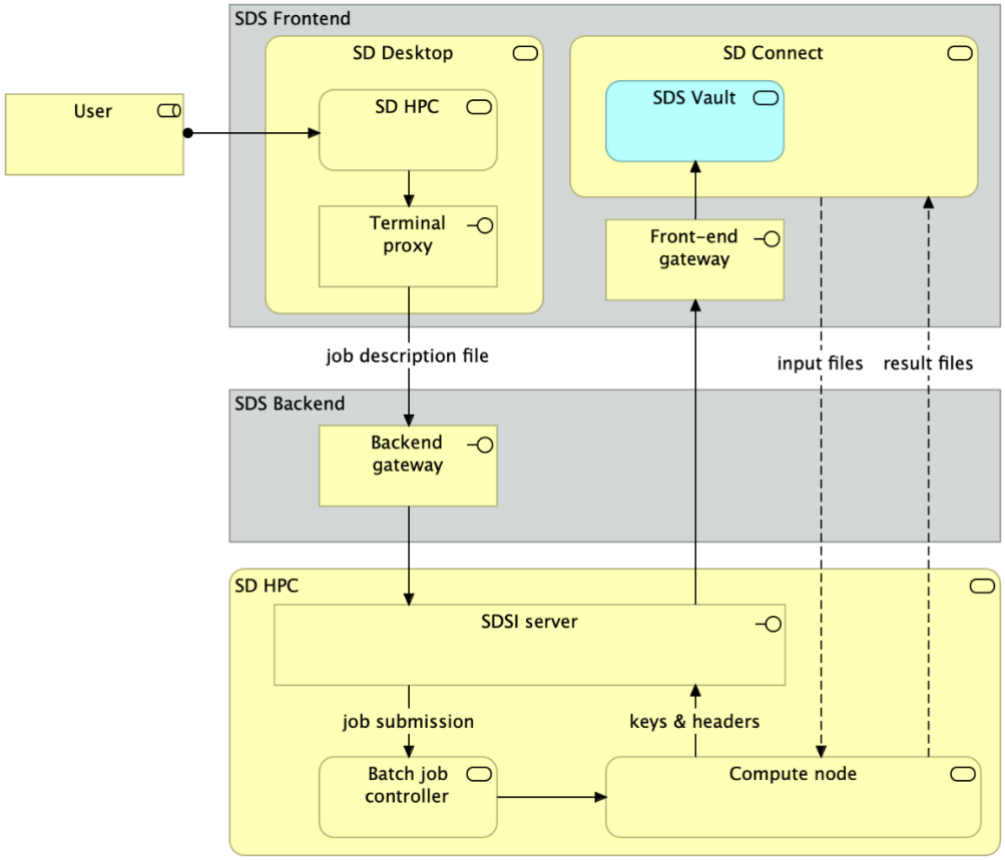

Using the language of the DARE UK Federated Architecture Blueprint[47], the SD HPC is implemented as an indirect query. It also fulfils other requirements of data privacy and accountability for secure federated analysis: Identity, source and target of all transfers are known, validated, and accounted for, and all data and communication are encrypted in transit. The guaranteed ephemerality of personal data and isolation of processing within the HPC system makes SD HPC a special case of federated analysis with exceptionally strong data privacy guaranties.

NORTRE, infrastructures for sensitive data in Norway

NORTRE[48] (Norwegian Trusted Research Environments) is a collaboration between the three main institutional research infrastructures for sensitive data in Norway, TSD[49] (services for sensitive data) at University of Oslo (UiO), HUNT Cloud[50] at the Norwegian University of

[47]DARE UK Federated Architecture Blueprint https://zenodo.org/records/14192786)
[48] NORTRE https://nortre.no/
[49] Tjenester for Sensitive Data (TSD) https://www.uio.no/tjenester/it/forskning/sensitiv/
[50] HUNT Cloud https://about.hdc.ntnu.no/

Science and Technology (NTNU) and SAFE[51] (secure access to research data and e-infrastructure) at University of Bergen (UiB). The three partners share knowledge and expertise so scientists and data controllers from Norway and around the world can collect, analyse, store, share and collaborate on sensitive data in an optimised and trustworthy manner.

In accordance with EHDS2, Norway has an ongoing project SPUHiN[52] which, among others, aims to prepare NORTRE for EHDS Regulation. This includes work such as drafting a requirements list, closely aligned with the requirements of the EHDS Regulation, which will be a goal for the three SPEs to achieve. SPUHiN is also creating a GAP analysis/list of points to be followed up with NORTRE to close these gaps.

Secure data transfer solutions (Finland)

In Finland, CSC has developed Supertunneli ("Super Tunnel")[53], an advanced solution designed to address limitations in current data transfer services, particularly for large datasets. Based on the widely adopted S3 interface, Supertunneli significantly improves capacity compared to earlier systems, which were limited to a 4 GB maximum file size per transfer. This enhancement reduces manual workload and minimises the risk of errors typically associated with managing large datasets, particularly for projects utilising AI or machine learning.

Supertunneli enables the secure transfer of large datasets in a single operation and initially supports transfers of at least 1 terabyte (1 TB). Transfers exceeding this threshold are also supported but require prior notification to Findata, Finland's data permit authority. Additionally, the process can be optimised by automating data encryption and decryption.

The connection used in the transfer of the files is encrypted and requires private/public keys exchange to allow the transfer. The transferred data is encrypted with Crypt4GH -encrytion tools provided by Global Alliance for Genomics and Health.

In compliance with Finland's Secondary Use Act, Supertunneli will serve as the standard interface for transferring datasets to approved secure processing environments. This ensures both regulatory alignment and improved security during data exchange.

Further developments include leveraging the S3 interface to support the anonymisation and submission of research results. A complementary system, Tulostunneli ("Results Tunnel"), is under development and will enable secure environments to automate the delivery of published outputs directly to Findata. This replaces the current manual process, traditionally the responsibility of researchers, thereby enhancing security and ensuring consistent anonymisation of disseminated results.

[51] SAFE https://www.uib.no/safe
[52] SPUHiN https://www.helsedirektoratet.no/om-oss/forsoksordninger-og-prosjekter/fair-secure-provision-and-use-of-health-data-in-norway-spuhin
[53] Supertunneli https://findata.fi/en/news/supertunneli-launching-in-may-2025-as-part-of-transfer-service-update/

For smaller or individual data transfers, the existing Tunneli service will remain in operation, providing a flexible, tiered approach to data exchange infrastructure in Finland.

Anonymity verification tool (Finland)

In Finland, Findata is launching a new tool called Portti[54] in Kapseli to simplify and speed up the process of verifying the anonymity of research results. With this update, users no longer need to complete separate summary forms.

Portti – Finnish for *Gateway* – allows users to send results directly for anonymity verification. It enhances data security and streamlines the process for both users and Findata staff.

Portti is strictly for transferring anonymous results. Personal data or non-anonymised information must not be submitted via the tool. To use Portti, users create a new transfer, upload the necessary files, fill in the required details, and select the result type. The results are then submitted to Findata for verification.

Findata processes submissions as quickly as possible, and within a maximum of five working days. Once approved, results are automatically delivered to the user's workspace, where they can be downloaded. Approved results remain available in the workspace for six months.

If anonymity issues are found, Findata will send instructions for corrections via email and may request additional information if needed before final approval.

The tool is being developed as part of the FinHITS project, co-funded by the European Union.

Building a secure health data network (Norway)

In Norway, Norwegian Directorate of Health (NDH), National Institute of Public Health (FHI) and NORTRE[55] are working together in the SPUHiN[56] project to ensure that sensitive data transfers within Norwegian data holders, SPEs and Authorities (Coming HDABs) are secured and guaranteed.

This will be done by building a HealthData@NO network based on the 4-corner model for the eDelivery model[57]. The 4-corner model ensures that data are encrypted, signed and addressed before leaving their safe environments. The 4 corner models addressing system, implemented by among other EU central Servies, guarantees that the data can only be delivered to an approved organisation registered at a central exchange register.

The network aims also to be HealthData@EU compatible by following the eDelivery standard when there are cross border data transfers.

[54] Portti https://findata.fi/en/news/new-tool-for-kapseli-result-submission-will-launch-in-september/
[55] NORTRE
[56] FAIR Secure Provision and use of health data in Norway (SPUHiN) - Helsedirektoratet
[57] How does eDelivery work

Annex 9: Overview of relevant EU regulations

The EHDS will provide a trustworthy setting for secure access to and processing of a wide range of health data. It builds further on the General Data Protection Regulation (GDPR), Data Governance Act (DGA) and Network and Information Systems Directive (NIS2).

These regulations collectively shape the landscape for SPEs, influencing their design, implementation, and operation in several critical ways. The following sections explore in greater detail how EHDS, GDPR and NIS2 shape the architecture and operational principles of SPEs, ensuring that they can securely and efficiently support the goals of the EHDS and GDPR. Notably, EHDS Article 73 serves as the primary reference for the general requirements of SPEs within the EHDS framework.

EHDS article 73 analysis to deduce SPE requirements

The logic of creating EHDS-specific requirements is presented as architecture graphs with the source paragraph as the figure legend.

Annex 9: Figure 1. EHDS article 73(1). Health data access bodies shall provide access to electronic health data pursuant to a data permit only through a secure processing environment which is subject to technical and organisational measures and security and interoperability requirements. In particular, the secure processing environment shall comply with the following security measures:

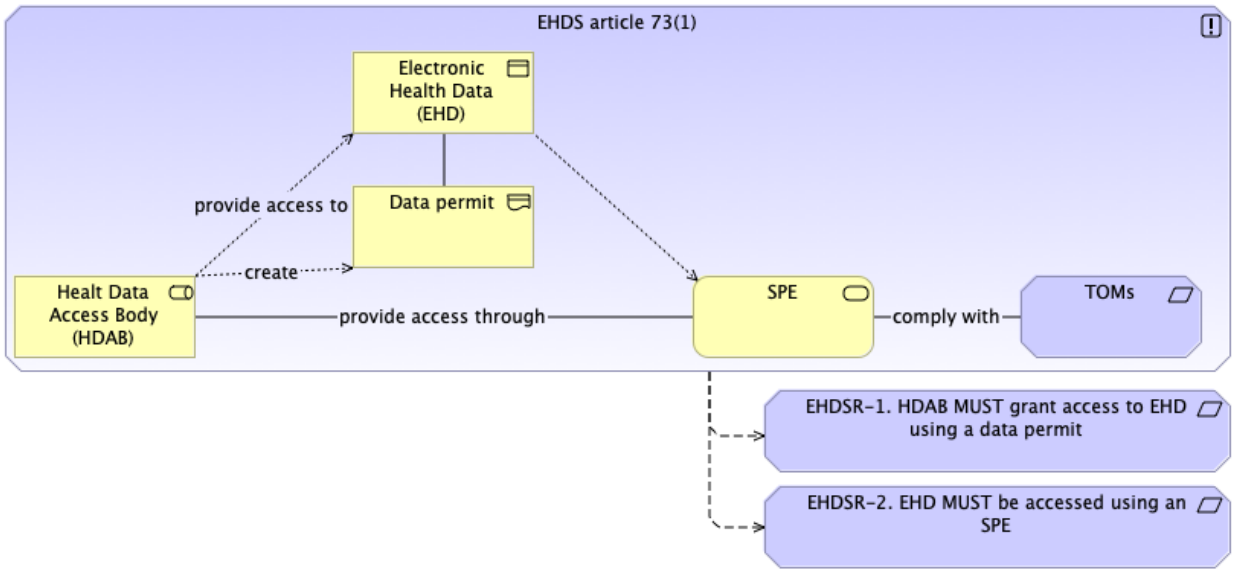

Annex 9: Figure 2. EHDS article 73(1)(a). *The restriction of access to the secure processing environment to authorised natural persons listed in the data permit issued pursuant to Article 68;*

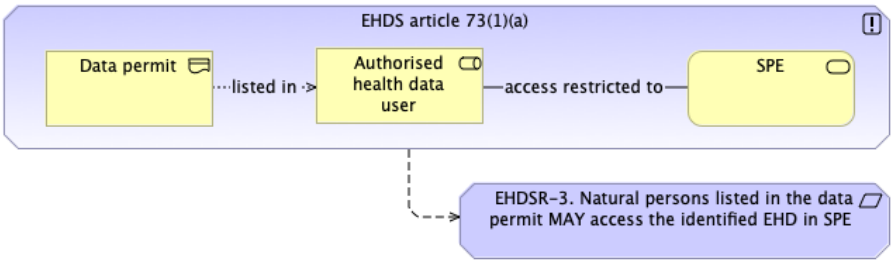

Annex 9: Figure 3. EHDS article 73(1)(b). *The minimisation of the risk of the unauthorised reading, copying, modification or removal of electronic health data hosted in the secure processing environment through state-of-the-art technical and organisational measures;*

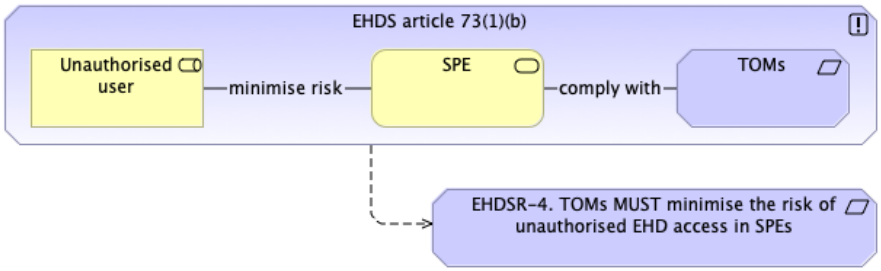

Annex 9: Figure 4. EHDS article 73(1)(c). *The limitation of the input of electronic health data and the inspection, modification or deletion of electronic health data hosted in the secure processing environment to a limited number of authorised identifiable individuals;*

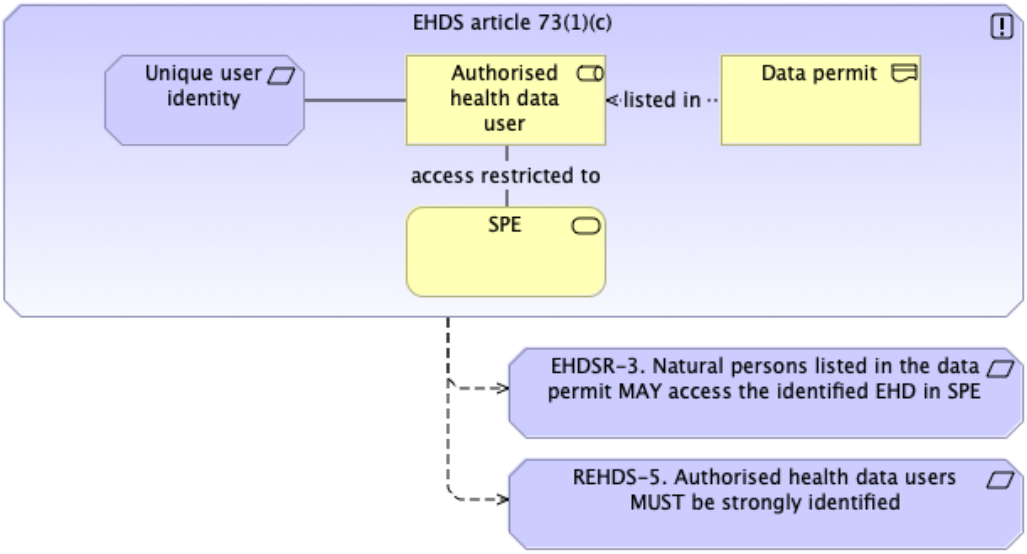

Annex 9: Figure 5. EHDS article 73(1)(d). *Ensuring that health data users have access only to the electronic health data covered by their data permit, by means of individual and unique user identities and confidential access modes only;*

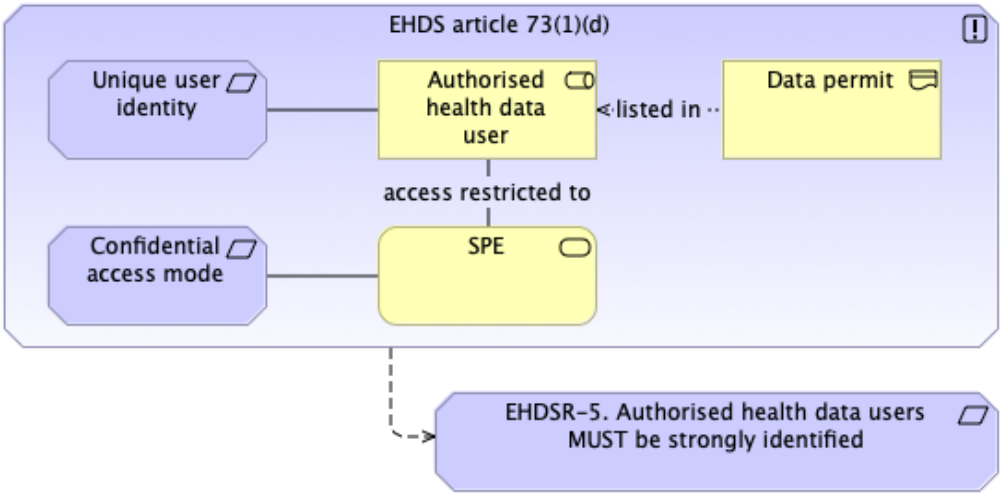

Annex 9: Figure 6. EHDS article 73(1)(e). *The keeping of identifiable logs of access to and activities in the secure processing environment for the period necessary to verify and*

audit all processing operations in that environment; logs of access shall be kept for at least one year;

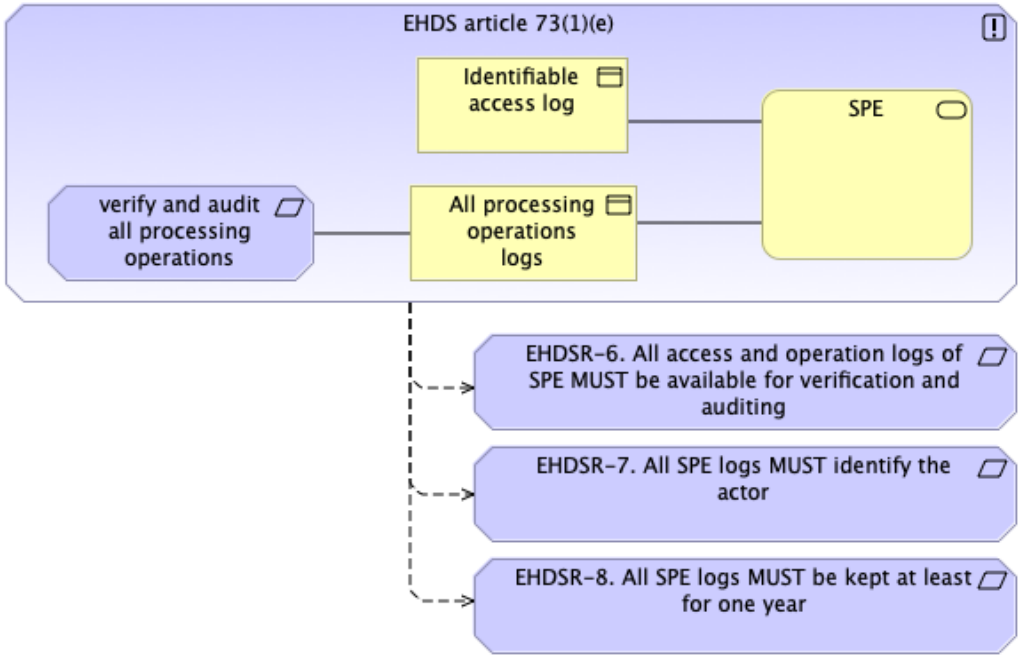

Annex 9: Figure 7. EHDS article 73(1)(f). *Ensuring compliance and monitoring the security measures referred to in this paragraph to mitigate potential security threats.*

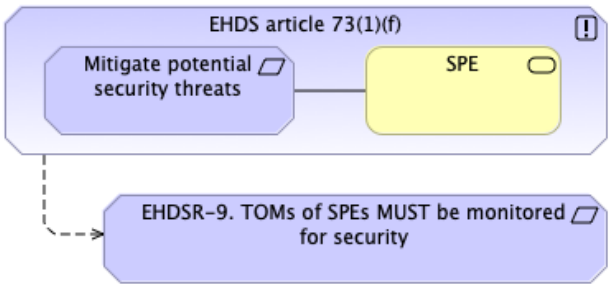

Annex 9: Figure 8. EHDS article 73(2). *Health data access bodies shall ensure that electronic health data from health data holders in the format specified in the data permit can be uploaded by those health data holders and can be accessed by the health data user in a secure processing environment.*

Health data access bodies shall review the electronic health data included in a download request to ensure that health data users are only able to download non-personal electronic health data, including electronic health data in an anonymised statistical format, from the secure processing environment.

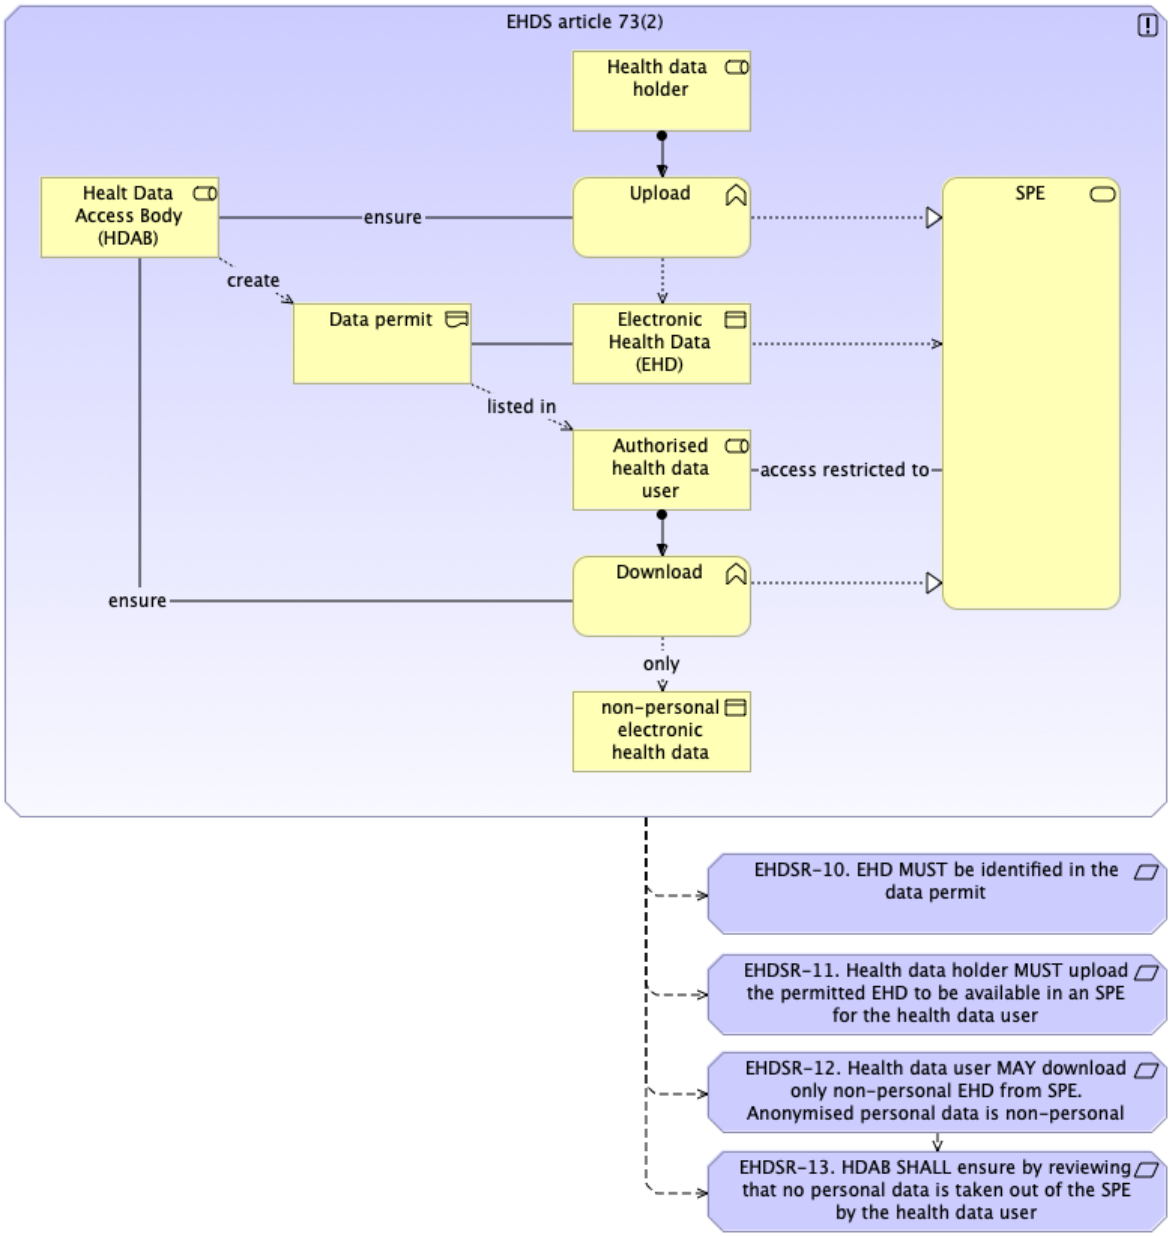

Annex 9: Figure 9. EHDS article 73(3). *Health data access bodies shall ensure that audits of the secure processing environments are carried out on a regular basis, including by third*

parties, and shall take corrective action for any shortcomings, risks or vulnerabilities identified by those audits in the secure processing environments.

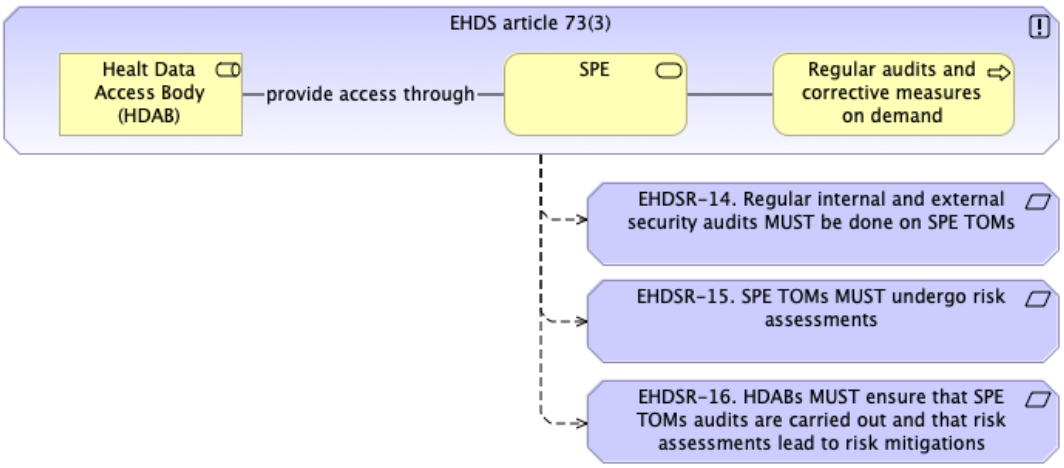

Annex 9: Figure 10. EHDS article 73(4). *Where recognised data altruism organisations under Chapter IV of Regulation (EU) 2022/868 process personal electronic health data using a secure processing environment, those environments shall also comply with the security measures set out in paragraph 1, points (a) to (f), of this Article.*

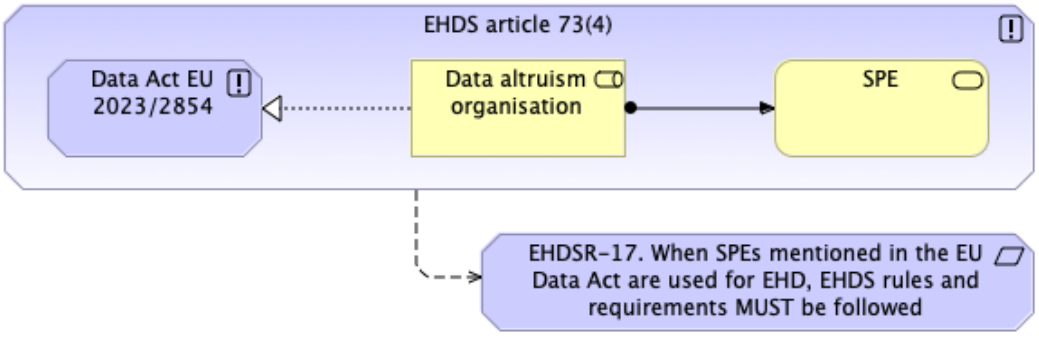

Annex 9: Figure 11. EHDS article 73(5). *By ... [two years from the date of entry into force of this Regulation], the Commission shall, by means of implementing acts, lay down the technical, organisational, information security, confidentiality, data protection and interoperability requirements for the secure processing environments, including with regard to the technical characteristics and tools available to the health data user within the secure*

processing environments. Those implementing acts shall be adopted in accordance with the examination procedure referred to in Article 98(2).

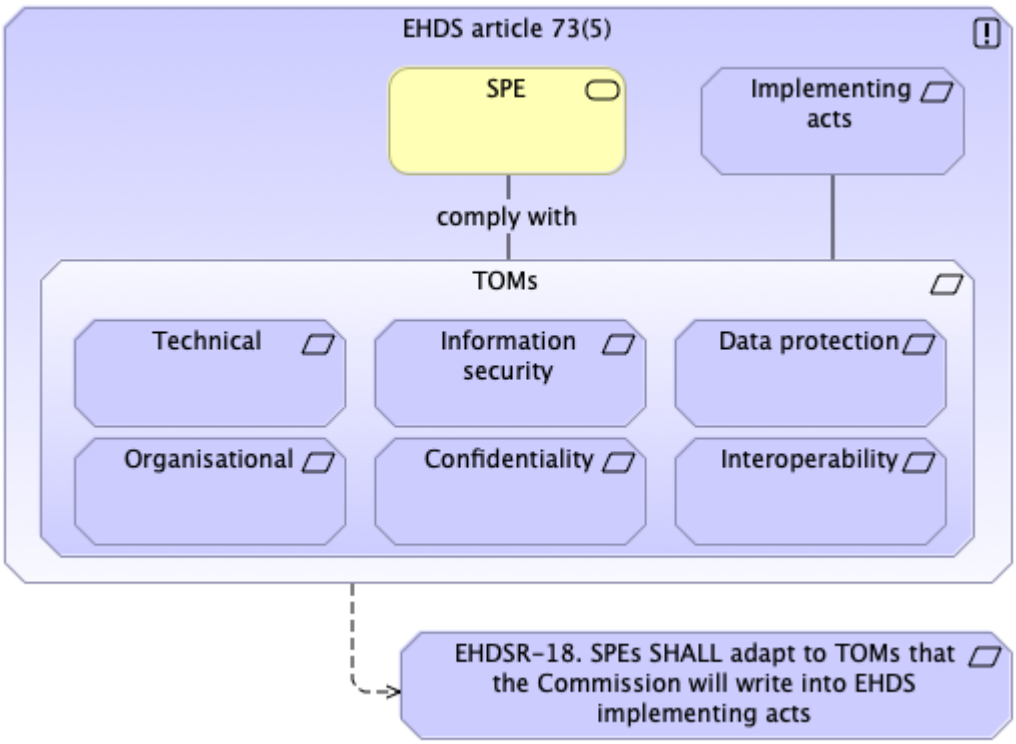

Key GDPR data and processing requirements

When developing technical specifications for SPEs, several GDPR articles must be considered to ensure compliance with data protection requirements. These articles include:

Article 5 – Principles relating to processing of personal data: This article outlines the core principles of data processing, including data minimisation, purpose limitation, accuracy and storage limitation. SPEs must be designed to process only the necessary data for defined purposes and ensure data is accurate and not retained longer than necessary. (Annex 9: Figure 12)

Annex 9: Figure 12. GDPR Article 5 – Principles relating to processing of personal data

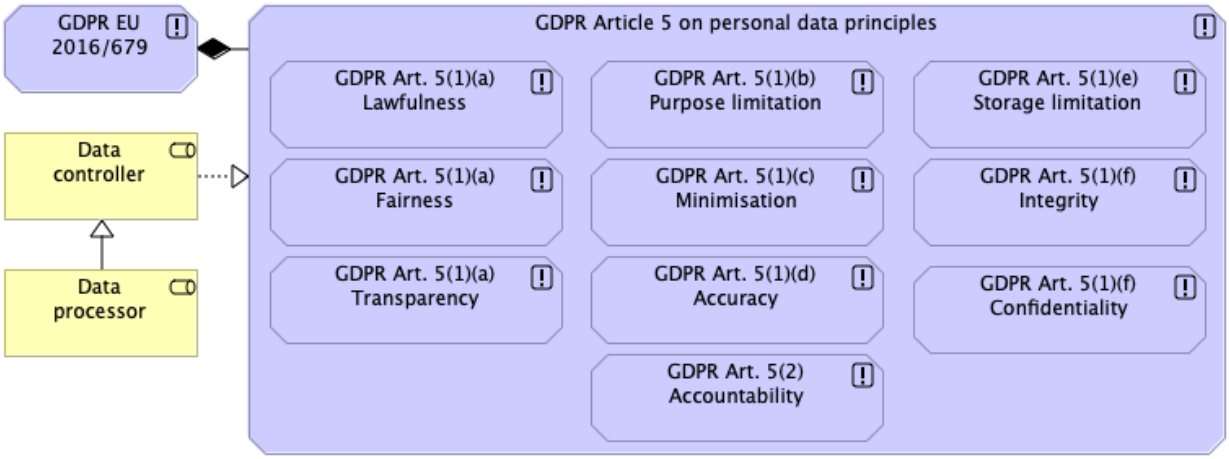

Article 25 – Data protection by design and by default: This article requires incorporating data protection measures into the design of systems, ensuring that privacy is considered from the outset and that the SPE is configured to minimise data exposure by default (Annex 9: Figure 13).

Annex 9: Figure 13. GDPR Article 25 – Data protection by design and by default

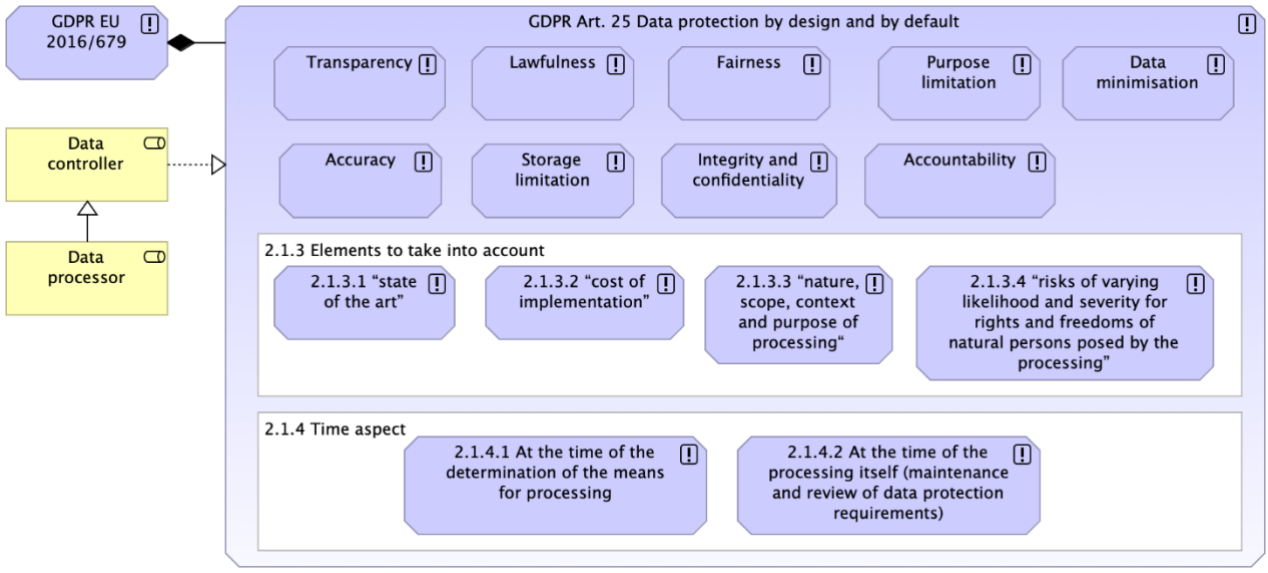

Article 32 – Security of processing: This article mandates that appropriate technical and organisational measures, such as encryption, access control, and regular security

assessments, are implemented to ensure the security of personal data processed within SPEs (Annex 9: Figure 14).

Annex 9: Figure 14. GDPR Article 32 – Security of processing

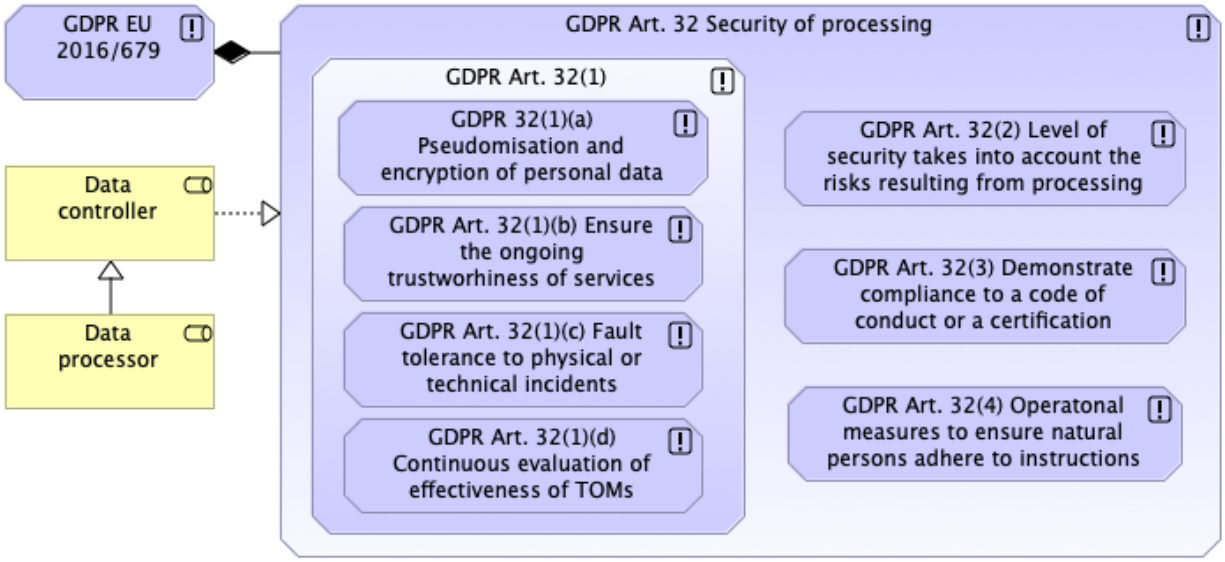

The Article 32 requires specific measures to protect sensitive health data and ensure that it is processed securely.

Firstly, data protection by design and by default, as specified in Article 25, is essential. SPEs should isolate health data from other environments to prevent unauthorised access. By separating sensitive data, SPEs limit exposure to only those with verified access permissions. Additionally, role-based access control systems must be implemented to ensure that only authorised individuals, based on their specific roles and responsibilities, can access the data. This ensures that access is tightly controlled and compliant with GDPR's requirements for data security.

In addition to isolation and access control, encryption is a cornerstone of the SPE security protocols, directly supporting Article 32(1)(a), which mandates appropriate security measures. All health data within the environment should be encrypted both in transit and at rest, to prevent unauthorised access and helping to protect data integrity and confidentiality.

The integration of privacy-preserving technologies, such as anonymisation and pseudonymisation, aligns with Article 32(1)(a) and Article 25(1). These techniques reduce the risk of re-identification and protect data subjects' privacy while enabling secure data processing. The guidelines for implementing these privacy-preserving technologies will be outlined as part of T7.2 in TEHDAS2.

SPEs should also undergo regular security assessments and penetration testing to identify any vulnerabilities and maintain compliance with Article 32(1)(d), which requires ongoing evaluation of the effectiveness of security measures. Robust incident detection and response protocols must be in place to swiftly manage any potential data breaches.

Finally, accountability and documentation are crucial to maintaining GDPR compliance. In accordance with Article 5(2), which emphasises the accountability principle, and Article 30, which requires detailed records of processing activities, SPEs should maintain comprehensive records of all security measures and access logs. Regular internal and external audits should be conducted to verify that the SPE is consistently meeting GDPR's privacy and security standards.

NIS2 Directive

The NIS2 Directive (Directive (EU) 2022/2555) is the EU's key cybersecurity legislation, aimed at enhancing the cybersecurity posture across the Union by improving the resilience and incident response capacities of essential and important entities in critical sectors.

The NIS2 Directive establishes stronger cybersecurity requirements for critical infrastructure across the EU, with the specific obligations defined through national transpositions. SPEs handling sensitive health data will generally fall within scope, requiring compliance with stringent security measures such as endpoint protection, intrusion detection systems, and incident response protocols. They must also demonstrate resilience against cyberattacks and operational risks, ensuring continuity of service and the secure handling of data.

Role of CSIRTs under NIS2

Under the NIS2 Directive, each EU Member State is required to establish its own Computer Security Incident Response Team (CSIRT). These national CSIRTs act as the single point of contact for receiving notifications related to cybersecurity incidents, threats, and near misses within their respective countries.

Cybersecurity risk-management measures under NIS2

According to NIS2, entities must implement cybersecurity risk-management measures based on an all-hazards approach, designed to protect both network and information systems and their physical environment from incidents. At a minimum, these measures must include:

(a) policies on risk analysis and information system security

(b) incident handling

(c) business continuity, such as backup management and disaster recovery, and crisis management

(d) supply chain security, including security-related aspects concerning the relationships between each entity and its direct suppliers or service providers

(e) security in network and information systems acquisition, development and maintenance, including vulnerability handling and disclosure

(f) policies and procedures to assess the effectiveness of cybersecurity risk-management measures

(g) basic cyber hygiene practices and cybersecurity training

(h) policies and procedures regarding the use of cryptography and, where appropriate, encryption

(i) human resources security, access control policies and asset management

(j) the use of multi-factor authentication or continuous authentication solutions, secured voice, video and text communications and secured emergency communication systems within the entity, where appropriate

Annex 10: Classification of risks and threats against SPEs

The table below (Table 10.1) is sourced from the French Digital Health Agency and is part of a report designed to help healthcare facilities implement the Politique Générale de Sécurité des Systèmes d'Information de Santé[58], the legislative framework governing IT security in the health sector.

Table 10.1. An example classification of risks and threats against general cybersecurity threats.

Event	Threat scenarios	Reference	Level of risk
Unavailability of functions (medical application, connected device, etc.) or information (personal health data or personal data of a user) that may lead to service disruption, a negative impact on the image among users, and/or a user risk	Malicious code (virus, Trojan horse); Degradation or interruption of the network (local network, WAN and Internet access, WiFi network); Power supply failure; Failure of server room air conditioning; Handling error by IT staff; Lack of maintenance;	R-01	High
	Application saturation; Uncontrolled modification of software.	R-02	Medium
	Damage to computer equipment (fire, water damage, etc.); Unavailability of personnel (pandemic, health crisis, difficulties accessing buildings, etc.).	R-03	High
Alteration of functions (medical application, connected device, etc.) or information (for example, personal health data or personal data of a user) that may lead to service disruption, a negative impact on the image among users, and/or a patient risk.	Malicious code (virus, Trojan horse); Input or command error by an IS user; Lack of maintenance; Equipment failure.	R-04	High
	Uncontrolled modification of software (software update or configuration/parameter changes); Misuse of software's intended purpose (abuse of system or application rights, direct access to application data, etc.).	R-05	Medium
	Computer intrusion.	R-06	Medium
Alteration of evidence elements generated and stored by the IS (e.g., application logs, etc.) that may increase the legal risk for the organisation in case of litigation.	Misuse of a software's intended purpose (abuse of system or application rights, direct access to application data, etc.).	R-07	Medium
	Computer intrusion.	R-08	Medium
Access to personal health data or personal data of a patient by an unauthorised third party, constituting a violation of privacy and/or professional secrecy.	Loss or uncontrolled removal of equipment (laptop, removable storage media, etc.).	R-09	High
	Misuse of a software's intended purpose (abuse of system or application rights, direct access to software data).	R-10	Medium

[58] PGSSI-S https://esante.gouv.fr/sites/default/files/media_entity/documents/pgssi-s_guide_pssi_non-expert-ssi-v-1.0.pdf

	Computer intrusion.	R-11	Medium

PerMedCoE[59], the HPC/Exascale Centre of Excellence for Personalised Medicine in Europe project, published a report[60] 2022 on about guidelines for HPC systems on cybersecurity. The architecture graph of their main findings is in Annex 10: Figure 1.

Annex 10: Figure 1. Personalised medicine applications HPC security breach risks and impacts.

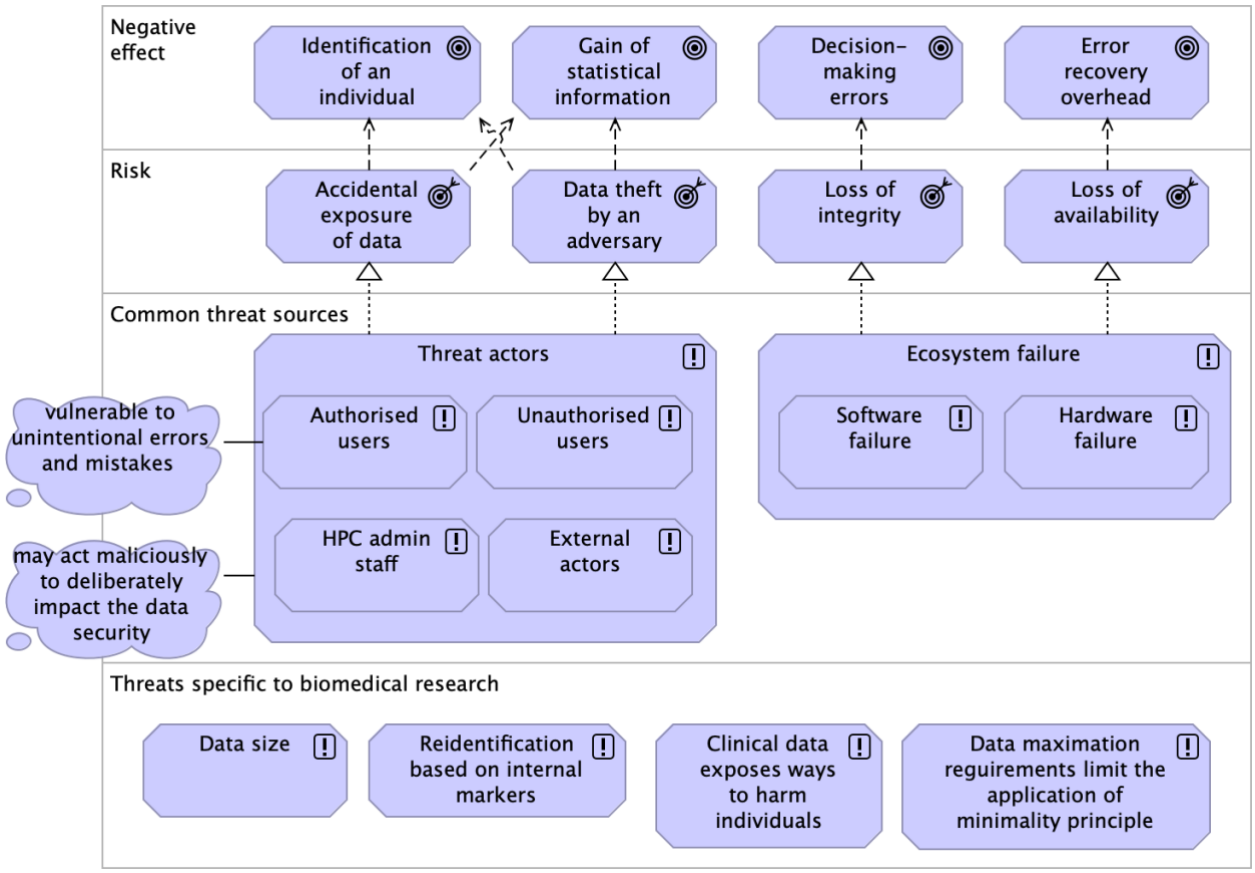

[59] PerMedCoE https://permedcoe.eu/
[60] D5.3 Derivation of general guidelines on data protection and privacy preservation of a use-case independent method/software development that will be exascale ready (September 2022) https://permedcoe.eu/wp-content/uploads/2024/10/PMC-D5.3.pdf