

Technology Program Edge Data Economy
German Federal Ministry of Research, Technology and Space



PRIVACY-ENHANCING TECHNOLOGIES FOR INFORMATION SECURITY IN EDGE-CLOUD APPLICATIONS

Imprint

Authors

Nils Jahnke
Sarah Schimankowitz

Fraunhofer Institute for Software and Systems Engineering ISST
Dortmund

Publishers

Peter Gabriel
Dr. Nicole Wittenbrink

Begleitforschung Edge Datenwirtschaft
Institute for Innovation and Technik (iit)
within the VDI / VDE Innovation + Technik GmbH
Berlin

Date

February 2026

Layout

PRpetuum GmbH

This study was commissioned by the German Federal Ministry of Research, Technology and Space as part of the coordination and support action (Begleitforschung) for the technology program Edge Data Economy (Edge Datenwirtschaft)

This English version has been produced by a machine translation tool with minimal human intervention. We do not guarantee the accuracy of this translated version and accept no liability for any errors. Please visit https://www.digitale-technologien.de/DT/Redaktion/DE/Downloads/Publikation/EDGE-Datenwirtschaft/Orientierungshilfe_Privacy-Enhancing-Technologies.html for the original publication in German language.

Executive Summary

Edge-cloud systems enable applications that create economic value and address societal challenges based on data from intelligent objects and infrastructures. This often requires the sharing of data with partners in established value networks or along the edge-cloud continuum. A fundamental requirement in data sharing is to ensure the protection of sensitive company and personal information. While local data processing at the edge enables a basic level of data protection and information security, relying exclusively on this measure is often not sufficient to meet these requirements while simultaneously achieving the envisioned value of data-driven applications. For example, there is often a need to aggregate sensitive data in a central location, such as the cloud, to gain rich insights or ensure the integrity of the data used. This is where privacy-enhancing technologies (PETs) come into focus. PETs include mechanisms for integrating data protection, information security, and data sovereignty "by design" into system architectures. PETs are a class of individual tools that can address specific information security requirements and risks in edge-cloud systems. Practitioners face the challenge of developing suitable PET strategies based on the specific needs of their application domain and the available PET tools to enable edge-cloud applications that comply to existing requirements for information security and deliver business value alike.

This study supports practitioners in developing their own PET strategies for edge-cloud Applications. It provides assistance in identifying information security requirements and risks, selecting suitable PET tools, and seamlessly integrating them into the application design. A core element of this study is the analysis of PET tools in real-world edge-cloud applications. The study shows how PET tools can contribute to the implementation of information security, what prerequisites must be created for their use in specific scenarios, and what implications arise from their practical implementation. To this end, the guidance draws on the findings of early adopters of edge-cloud systems and PETs. The early adopters stem from projects part of the technology program "Edge Data Economy" commissioned by the German Federal Ministry of Research, Technology, and Space (BMFTR). This study's guidance is particularly aimed at system architects and data protection officers who are required to design data processing processes in edge-cloud systems in compliance with data protection regulations.

Based on the presentation of possible risks such as physical and cyber attacks, uncertain data sovereignty, insider threats, and misconfigurations, as well as requirements such as data minimization, data integrity, purpose limitation, and the prevention of data leaks "by design" in edge-cloud applications, this study analyzes five PET-tools in practical application scenarios:

- Hardware keys for secure authentication without personal data disclosure in the food industry.
- Federated learning for collaborative AI training without raw data transfer in industrial manufacturing.
- Compute-to-data for performing analyses in the data owner's environment in industrial manufacturing.
- Zero-knowledge proofs for data-based verification without disclosure of sensitive data in the energy industry.

- Trusted execution environments for confidential calculations in isolated hardware environments in the energy industry.

The study additionally presents four areas of action and associated recommendations for the successful use of PETs in edge-cloud applications:

- 1) Establishing a trustworthy partner ecosystem and creating necessary incentive mechanisms.
- 2) Creating the operational prerequisites needed for PET use, including training and awareness.
- 3) Ensuring the technical validity and integrability of PETs in the application context.
- 4) Ensuring the regulatory compliance of the PET-supported edge-cloud application.

Edge-cloud systems and data sharing become increasingly relevant for data value creation. At the same time, the requirements for the protection of sensitive company and personal data remain challenging. In this context, implementing PET-based data processing becomes an important success factor. PETs not only enable compliance with regulatory requirements but also create the basis for trust-based cooperation in complex edge cloud ecosystems. Companies that want to jointly pursue data-driven value creation in the future should actively engage with PET.

Contents

1	Introduction	5
2	Background.....	7
2.1	Edge-cloud systems.....	7
2.2	Data protection and information security through Privacy-Enhancing Technologies	9
3	The Problem Space – Requirements and Risks for Data Protection and Information Security in Edge-Cloud Applications.....	15
3.1	Sensitive data and information in edge-cloud applications	15
3.2	Specific requirements and risks for data protection and information security in edge-cloud systems.....	17
4	The Solution Space – PET Tools for Implementing Data Protection and Information Security in Edge-Cloud Applications	22
4.1	PET tools	22
4.2	Summary of existing guidance on PETs.....	25
4.3	Scenarios for applying PET tools in edge-cloud applications.....	26
4.3.1	Hardware keys in the context of quality control in the agrifood industry.....	28
4.3.2	Federated learning in the context of industrial manufacturing.....	30
4.3.3	Compute-to-data in the context of industrial manufacturing.....	33
4.3.4	Zero-knowledge proofs in the energy industry.....	36
4.3.5	Trusted execution environments in the energy industry.....	39
5	Recommendations for Action and Outlook.....	42
	Appendix A – Detailed Overview of PET	50

1 Introduction

Edge-Computing moves compute and storage closer to where data is generated and enables data processing outside of centralized systems. Edge-cloud systems synergistically combine edge and cloud capabilities. The edge and cloud play complementary roles: Raw data is collected and processed at the edge, for example using AI methods, while the cloud is responsible for orchestration, long-term data storage and the provision of aggregated results. Edge cloud systems open new opportunities for processing, analyzing, and sharing data, thereby enabling application scenarios with economic and social added value. By processing data in local environments rather than on centralized (cloud) resources, edge-cloud systems also fundamentally promote data protection and confidentiality.

However, in many cases, the use of edge-cloud systems alone is not sufficient for protecting data. Rather, distributed data processing and the use of AI in edge-cloud systems give rise to a number of novel challenges relating to data protection and information security (Sheikh et al., 2025). For example, potentially sensitive environmental data is often recorded and analyzed on a large scale. Furthermore, there is an increased risk of physical or cyberattacks on decentralized edge devices. In addition, misconfigurations or human errors while handling cloud services can lead to data leaks. This results in an increased need for the protection of personal or business-critical data in edge-cloud systems.

This is where privacy-enhancing technologies (PET) (Hes & Borking, 1995) come into play. With the help of PETs, aspects of data protection and information security can already be ensured through system design. PETs enable the usage and processing of data while promoting privacy and confidentiality, strengthening trust between partners, and supporting compliance with legal requirements such as the General Data Protection Regulation (GDPR). The term PET is to be understood as an umbrella term for a wide variety of technical tools for ensuring privacy. It includes approaches such as confidential computing or federated learning, which in turn can be used for solving specific problems.

To use PET tools, technical, organizational (e.g., roles, understanding of technology, acceptance) and legal challenges must be clarified (Klymenko et al., 2025). Data processing companies must therefore develop a suitable PET strategy that enables them to implement PET-based information security in edge-cloud systems as efficiently and effectively as possible. A PET strategy answers the following questions:

- Which PET tool is most suitable for protecting data in my use case, taking into account the type, scope, context, and purpose of data processing in the designed edge-cloud system?
- What are the potential challenges and limitations of using specific PET tools in edge-cloud systems that need to be addressed in the application design?
- How can the PET tool be integrated into data processing workflows in edge-cloud systems in order to fully exploit the potential for information security?

This study helps practitioners develop their own PET strategies for ensuring information security in their edge-cloud applications. It characterizes information security requirements and risks specific to edge-cloud applications and identifies suitable PET tools that can be used to meet these challenges. The study also highlights the requirements that must be met for the use of PET tools

in real application contexts and discusses the implications of their use. To this end, case studies of the use of PETs in real-world edge-cloud applications are presented, and the experiences of early adopters are shared. In addition, the study provides recommendations for the implementation of PET-based edge-cloud applications. The study aims to support system architects and data protection officers who are responsible for information security in edge-cloud applications.

The remainder of the study is structured as follows: Section 2 explains the conceptual basis of this study: edge-cloud systems and PETs as tools for implementing information security. Section 3 presents the problem space by identifying the types of data needing protection and possible requirements and risks for information security in edge-cloud systems. The corresponding solution space is presented in Section 4. This section provides an overview of possible PET tools, identifies existing solutions that support the design and implementation of PET-based data processing architectures, and analyzes practical scenarios for the use of PET tools in real-world edge-cloud applications. The concluding outlook provides recommendations for practitioners who want to ensure information security using PETs in their own edge-cloud endeavors (Section 5).

The collaboration with experts from the early adopter projects of the technology program “Edge Data Economy” has contributed significantly to analyzing and preparing practical insights into the integration of PET tools in edge-cloud systems. The team of authors would therefore like to take the opportunity to thank all the interviewees once again for their support in the creation of this study:

- Sabine Haag, Robert Bosch GmbH, project *EASY*
- Tobias Schlagenhaut, Robert Bosch GmbH, project *EASY*
- Alexander Tessmer, University of Osnabrück, project *FRED*
- Marvin Ehaus, Fraunhofer Institute for Applied Information Technology FIT, project *DEER*
- Fabian Gast, Institute for Production Management, Technology and Machine Tools (PTW) | TU Darmstadt, project *ESCOM*
- Felix Förster, OLI Systems, project *DEER*

The responsibility for the content of this study lies exclusively with the authors.

The study was produced as part of the coordination and support action (Begleitforschung) for the technology program "Edge Data Economy" of the Federal Ministry of Research, Technology and Space (BMFTR). The program comprises ten projects that are developing new forms of edge-cloud systems for important sectors of the German economy.

2 Background

To establish data protection and information security in edge-cloud systems using privacy-enhancing technologies (PETs), a basic understanding of the core concepts is needed. To this end, this section first describes the conceptual fundamentals of edge-cloud systems and the related data processing activities and presents why the topology of edge-cloud systems necessitates a comprehensive consideration of data protection aspects (Section 2.1). This is followed by a presentation of the fundamentals of data protection, information security, and their risk management, as well as an explanation of PET as a tool for reducing information risks (Section 2.2).

2.1 Edge-cloud systems

Edge computing is a concept for executing data processing close to the location of data generation to reduce the amount of data that needs to be distributed to the cloud or other centralized systems. The use of edge computing results in lower latency in data processing, more efficient use of resources, and advantages in terms of the confidentiality of critical data. However, centralized systems remain important as these are needed for tasks with increased requirements for computing power, long-term availability of aggregated data, or resource orchestration.

To reap the benefits of cloud and edge computing resources at the same time, both components are combined in edge-cloud systems. Edge-cloud systems are hybrid computing architectures that draw on various resources—typically computing power, storage, and network resources—from the so-called edge-cloud continuum. Such resources range from IoT devices to public cloud infrastructure. Various attempts have been made in both scientific and practical literature to classify the edge-cloud continuum.

In this publication, the edge-cloud continuum is divided into four levels, which are arranged according to their distance from the location where the data is generated, starting with the closest level on the left (see Figure 1). The levels also differ in two key respects: the available computing and storage capacities, and the actors typically responsible for providing and governing resources at each level.

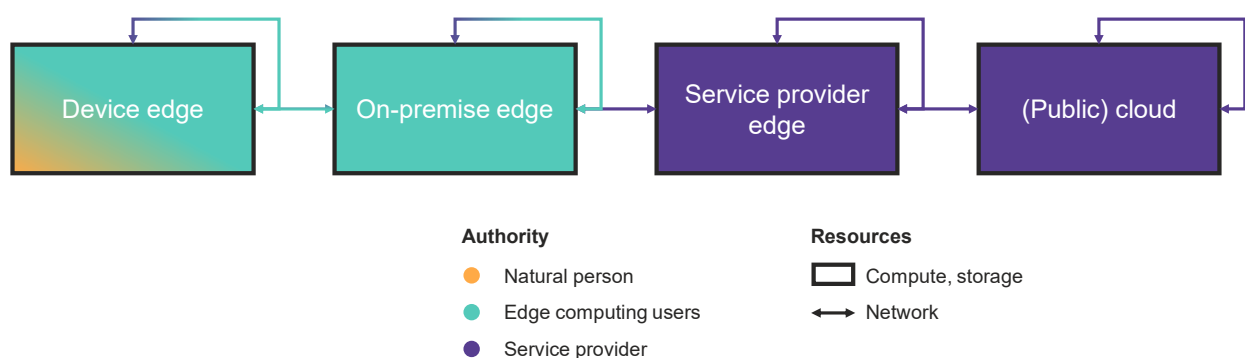


Figure 1: Layers of the edge-cloud continuum for edge-cloud systems

On the one hand, the *device edge* includes so-called "constrained devices," i.e., microcontroller-based devices with limited memory capacity, computing power, and power consumption, which are often developed for specific application scenarios. On the other hand, it also includes "smarter" devices that have greater technical capabilities and provide application-agnostic computing power and storage. Examples of such devices are smartphones or tablet computers. The responsibility for operating the device edge lies either with natural persons, in the case of operating applications on their smartphones, or, in the case of industrial applications, with the respective companies. The *on-premise edge* comprises server hardware housed in conventional, physically secure data centers or in modular data center units located near or in office buildings and factories. These resources usually belong to and are operated by the user company. With the *service provider edge*, cloud service providers or telecommunications providers bring cloud computing capabilities closer to potential users. The resources are located in radio masts, network nodes, or urban agglomerations, for example, and thus remain in close proximity to the location of data generation. The services offered at the service provider edge are based on standardized hardware and software and can be consumed by different customers. *The (public) cloud* provides nearly infinitely scalable computing and storage resources in centralized locations that can be flexibly consumed by its customers.

The connection between the individual storage and computing resources in the edge-cloud continuum is established via network connections, which are operated by one or more players. In the levels of device edge and on-premise edge, networking services can be provided by the edge computing users in the form of WiFi, 5G campus networks, or Ethernet connections. In case operations take place in more remote regions or by moving objects, services from (mobile) network operators (4G/5G) and Internet service providers are used. Such services are also needed to transport data over longer distances to the service provider edge or public cloud. So-called interconnection services can be used for direct network connections that bypass the public Internet.

Data processing in edge-cloud systems

When implementing data-driven applications, the individual levels of an edge-cloud system typically perform specific tasks along the data processing chain. For data acquisition, device edge and on-premise edge can collect large amounts of data at high speed and with low probability of failure from sensors deployed in the field. Additionally, at both levels, data preparation functions are executed. This can, on the one hand, ensure local data anonymization or pseudonymization. On the other hand, applying data minimization techniques leads to reduced data transfers into the cloud, thereby reducing the needed bandwidth and network costs.

To obtain actionable information based on the available data, it must be combined and evaluated using (AI-based) data analysis techniques. For applications that require real-time information provision (e.g., machine control or autonomous driving), data analysis can be performed at the edge level, taking given resource constraints into account. With increasing storage and computing requirements and increasing numbers of data sources, more centralized data processing levels such as the service provider edge or the public cloud are used for information generation.

The information obtained must be made available for use by humans and machines along the edge-cloud continuum. The edge levels can serve as a relay point for information delivery, distributing selected information from cloud services to specific devices and users. Conversely, information can be passed from the edge to cloud services to make this information available and usable on a global scale.

Edge-cloud systems in the data economy

While many companies already use edge-cloud systems for internal purposes, local data processing is no longer sufficient to generate value in the data economy. Instead, collaborative business models are increasingly being pursued in which data or information must be shared across different parties. This includes applications in which data is analyzed by a service provider with the appropriate expertise, or applications in which data is made available at a central location for better monitoring and control. One example of this is the tracking of goods throughout the entire supply chain, to ensure that they are disposed of if they exceed certain temperatures or forces during transport. Furthermore, requirements regarding regulatory compliance or the presentation of verifiable information to third parties may necessitate the sharing of data. Finally, data transfer is conceivable for product development, the generation of insights, or for the purpose of knowledge transfer, for example, through the collaborative training of AI models. In these cases, edge-cloud systems continue to cater for tasks such as data acquisition, data preparation and data provisioning. However, to fully reap the data-based benefits, data shared with other parties or aggregated at a central point must often continue to contain commercially sensitive or personally identifiable content. Yet, this conflicts with protection needs of the organizations or data subjects involved as well as with legal requirements.

2.2 Data protection and information security through Privacy-Enhancing Technologies

Based on the general rights of individuals to privacy and of companies to protect trade secrets, claims arise in the digital space regarding data protection and the secure processing and storage of information. Data protection means complying with and ensuring the existing rights, obligations, and freedoms of the data subjects with regard to the collection, storage, use, publication, disclosure, and deletion of sensitive information. The main goal of data protection is to ensure the legitimate use of data throughout its entire life cycle. In addition to data that is sensitive from the perspective of the privacy of natural persons, this study also takes into account information that is sensitive and should be protected from a business perspective.

Risks to data protection and information security generally refer to the possible effects resulting from a lack of information regarding the confidentiality, integrity, and availability of data, as well as compliance with data protection regulations (ISO, 2024). To identify, analyze, and address these uncertainties, data protection and information security risk management must be implemented. Risk management is usually embedded in an organization's broader risk management framework. Managing data protection risks primarily involves the following activities (ISO, 2024):

- Developing the necessary understanding of the application context and use case, including the analysis of data protection requirements
- Identifying, analyzing, and evaluating the risks with regard to the data and information to be protected
- Developing and implementing appropriate countermeasures to avoid or reduce risks
- Communicating and consulting with affected stakeholders on risks and countermeasures
- Continuous monitoring and review of risks and countermeasures

This work focuses on the first three aspects of the risk management process. The following section begins by describing approaches that support the identification of data protection and information security requirements as well as potential risks. This is followed by a conceptual classification of privacy-enhancing technologies (PET) as a tool for meeting these requirements and minimizing potential risks by design.

Identification of data protection and information security requirements

The first necessary aspect of risk management is the creation of sufficient knowledge about the application, the sensitive data and information used or generated, and possible data protection requirements. While the sensitive data and information result from the specific application context, protection and security requirements can generally be derived from the perspectives described in Table 1 and thus be motivated by a number of factors. The factors are arranged according to the granularity of data processing requirements imposed. Legal and regulatory factors set the general framework within which data processing activities may take place. Contractual factors define additional, more extensive claims and obligations for the parties involved. Application factors refer to specific technical and organizational elements to ensure data protection and information security in practice.

Legal and regulatory factors	Contractual factors	Application factors	Other factors
- International, national, or local laws	- Agreements between stakeholders	- Characteristics of the targeted application	- Individual privacy requirements of those affected
- Court decisions	- Existing company rules and guidelines	- Industry-specific guidelines or best practices	- Internal control systems
- Agreements with trade unions or works councils		- Reputation factors	

Table 1: Origin of data protection requirements according to ISO (2024)

Legal and regulatory requirements result primarily from data legislation. As a pioneer in data legislation, the European Union formulates comprehensive requirements for the protection of (personal) data. The central instrument is the General Data Protection Regulation (GDPR, European Parliament and Council of the European Union (2016)). The GDPR establishes seven principles for the processing of personal data (Art. 5) and, based on these principles, formulates obligations for the implementation of data protection through technical solution design (Art. 25), amongst others. In addition to the GDPR, the AI Act (European Parliament & Council of the European Union, 2024) and the Data Act (European Parliament & Council of the European Union, 2023) describe potential additional requirements for the protection of data and information. On the one hand, the AI Act refers to the principles of the GDPR. Additionally, it also explicitly mentions the use of PETs to train AI systems without the need to transfer or copy data between parties (Rec. 39). The Data Act is intended to simplify access to raw data and pre-processed data from connected products and services. It provokes a tension between data transfer obligations and the requirements for protection against the unlawful disclosure of information and associated trade secrets. The Data Act also refers to so-called technical protection measures as a means for restricting unauthorized access to data and metadata while at the same time complying with the data provisioning obligations. In addition, industry-specific and international legislation that defines certain requirements for the cross-border flow of data must be taken into account.

Court rulings that clarify the interpretation of the legal provisions help in the implementation of the given requirements. Comprehensive court rulings are already available in case of the GDPR. Since the new data legislation (AI Act and Data Act) is not fully applicable yet, court rulings on possible disputes can only be expected in the coming years.

Other noteworthy requirements may arise from existing agreements or the immediate interests of workers' councils. Examples of applications usually viewed critically include video surveillance or other means of recording workers' activities. Even if the collected data in these applications is not being used for the purpose of employee monitoring, it is in the interest of employees to meet privacy protection requirements.

Contractual agreements are the second area from which data processing protection claims are formulated. Contractual agreements can be defined between a wide variety of parties. Examples include agreements between the data subjects or companies and the service providers or data processors, such as cloud service providers. In addition, agreements on bilateral data flows between companies (so-called data usage agreements) are documents that specify requirements for data processing and information use.

Further privacy requirements may arise from the organization's existing guidelines and rules. Organizations formulate internal requirements for data processing, information handling and the development of data architectures. This may also include requirements for data sovereignty, i.e., the possibility to exercise control over data even after it has been transferred to another party.

Application factors are considered the third source of requirements for data protection and information security. They describe conditions that result from the specific application context, the application objectives, the application features, and the expectations of the application stakeholders. Application factors can vary greatly depending on the application, industry, and company. For example, different standards and guidelines exist depending on the industry, which must be met by the corresponding applications. The requirements arising from these factors often describe "how" protection and security can be realized while the previous categories formulate "what" should be guaranteed. It should also be considered what negative implications non-compliance with given standards or technical measures could have, for example, with regard to the reputation of the company and the use of the provided application (Janakiraman et al., 2018).

Fourth, there are other general factors that play a role in deriving data protection and information security requirements. Consideration should be given to the individual privacy expectations of the data subjects. Data subjects may have additional expectations that are not necessarily agreed upon in the terms and conditions. Furthermore, only about a quarter of data subjects actually read the terms and conditions in their entirety (Gerber et al., 2018). Personal attitudes toward privacy and the assessment of perceived risks depend on many factors, including the purpose of the transaction and its subjective benefits, technical understanding, trust in the data processor, previous experience, and the self-confidence of the data subjects (Gerber et al., 2018). In addition, internal control systems can influence specific data protection measures. Components of such control systems include, for example, data governance, risk management, or other protective measures agreed upon by the company in the context of auditing and certification.

Identification of risks to data protection and information security

After the requirements for data protection and information security have been identified, an analysis of the potential risks to fulfilling these requirements is needed. Several approaches can be used to identify risks to data protection and information security. Two popular approaches are the asset-based approach and the event-based approach (ISO, 2022). In the asset-based approach, organizational assets are the starting point for consideration, and the threats and vulnerabilities that could damage these assets are examined. In addition to data and information worthy of protection, assets also include systems and software applications, the hardware used, and other factors such as reputation, which can be damaged through data breaches. Once the assets worthy of protection have been identified, relevant threats to data protection and information security must be identified. Threats can arise, for example, from the organization, processes, system configuration, hardware and software used, or from third parties.

The event-based approach to risk identification focuses on the analysis of (previously defined) events that have an impact on data protection or information security. Such events include, for example, the use of sensitive data beyond the intended purpose, the comprehensive monitoring of individuals, or a lack of transparency regarding data use. Scenarios are created for these events in order to better assess the consequences. While the asset-based approach tends to take a technical perspective, the event-based approach is particularly suitable for considerations at the management level. Both approaches can also be combined.

Possible general risks for data protection and information security include unauthorized access, unauthorized modification, or loss or theft of sensitive data. In addition, further risks may arise from data processing. These include the collection of sensitive data beyond its intended use, the inappropriate linking of personal data, the disregard of legal regulations, limited transparency in data processing, or the sharing of sensitive data with third parties without the consent of the persons concerned.

Privacy-Enhancing Technologies

As the data economy continues to grow, there is an ever-increasing demand for tools that enable data processing in compliance with data protection regulations and information security requirements. Although market research organization Gartner expects that by 2025, around 60 percent of large companies will be using at least one Privacy-Enhancing Technology (PET), the PET market is still relatively young (Noble, 2023). As a result, there is currently no commonly agreed definition of the concept "PET" and no consensus on which technical measures should be actually be considered a PET. This study understands PET tools in the sense of the European Union Agency for Cybersecurity as "software and hardware solutions, i.e. systems encompassing technical processes, methods, or knowledge to achieve specific privacy or data protection functionality or to protect against risks to privacy of an individual or a group of natural persons" (Schiffner, 2015b, p. 9). While this definition emphasizes the protection of personal data, PETs are also capable of improving the security of information that is sensitive from a business perspective. In the context of PET, the term privacy-preserving technology (PPT) is often used synonymously. PETs also include concepts of privacy-preserving machine learning (Xu et al., 2021), in which privacy protection techniques are integrated into machine learning processes or special AI algorithms are used.

PETs are used to ensure the protection of sensitive data and information from a certain point in the data processing chain. They can be integrated in the phases of data collection, processing,

analysis, and transfer. PETs can be integrated into data processing architectures as individual technical solutions or in combination with each other. For many data-driven cooperation models, such a combination of different PETs is essential to meet complex requirements. Accordingly, there is no universally applicable PET. Instead, optimal protection requires a context-specific selection and, if necessary, integration of several technologies, depending on the objectives to be achieved, possible challenges, and the technical and organizational environment (ISACA, 2024) – in other words, a suitable PET strategy.

In addition to increased protection of sensitive data and information, the use of PETs offers further advantages: they strengthen the trust of users and partners in data processing, enable more legally secure collaborations, and allow data-driven innovations while maintaining data sovereignty. In addition, selected PETs reduce dependence on trusted intermediaries and make the type and scope of data protection measurable. However, PETs increase the technical implementation effort of data architectures, often require specialized expertise, and demand a high level of technical understanding. Furthermore, they sometimes induce performance losses, higher costs, and expenses that users are rarely willing to pay. In addition, while the use of PETs strengthens data privacy, it needs to be accompanied by additional organizational and technical data protection measures.

The characteristics of PETs can be distinguished on the basis of several functional criteria (Heurix et al., 2015). An overview of these criteria is shown in Figure 2.

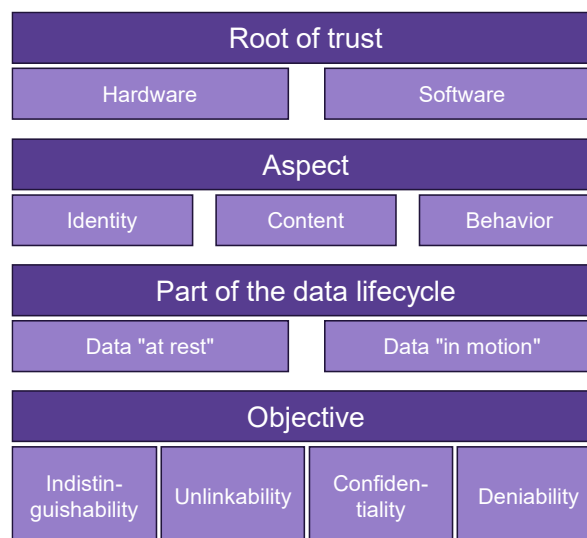


Figure 2: Taxonomy of PET based on Heurix et al. (2015)

Firstly, PETs can be differentiated according to their root of trust, or rather the protection mechanism used. Hardware-based PETs primarily achieve their protective function through specialized physical architectures (supplemented by software libraries) that isolate and protect data and code from the environment during data processing. Software-based PETs primarily rely on mathematical or cryptographic methods to ensure data protection.

PETs also focus on one or more aspects of data protection. The aspect dimension distinguishes between three main categories: identity, content, and behavior. Identity protection refers to the concealment or masking of actively or passively involved entities. Active entities provide data about themselves, while passive entities are included in third-party data. The aspect of content

protection refers to the data or information that is created during the use of a service and subsequently managed by the service provider. Behavioral data includes the actions of the individuals concerned when generating data. The protection of this data is relevant when the data content is already hidden, but access patterns or metadata can reveal sensitive information about the user.

The "data lifecycle" dimension describes which data is addressed and protected by a PET. PETs can generally be used throughout the entire data lifecycle. A distinction can be made between two basic data statuses: data "at rest" and data "in motion." Data "at rest" refers to data or information stored in repositories such as databases or file systems. Procedures aimed at protecting this data ensure that unauthorized persons who gain access to it cannot violate privacy. In the context of this publication, data "in motion" includes both data and information during transmission from one party to another and during processing.

The objective of a PET defines what the PET should do and specifies the means to be used to protect privacy and ensure information security. Unlinkability describes the objective of not being able to link individual entities to other data, such as transactions or data entries. This prevents third parties from determining whether certain entities belong to a larger group. Another objective is indistinguishability. This makes it possible to hide an entity within a larger set of entities. For example, this is achieved by changing or clustering the attributes in a table in such a way that it is not possible to trace them back to a single entity. Maintaining secrecy or ensuring confidentiality refers to the requirement to protect data from unintentional disclosure. Finally, deniability describes the ability of a party involved to credibly deny to third parties that they participated in a particular communication or transaction, even if the immediate counterpart in the situation was convinced of its authenticity. Examples of such solutions are group mechanisms in which "someone from the group" performed a corresponding activity.

The explanations in this section highlight the types of requirements that generally arise for data protection and information security and how associated risks can be identified. In addition, this section explains how PETs can generally be used as a tool to support appropriate measures for reducing privacy and information security risks. In the context of Edge-Cloud-Applications, however, there are several specific conditions, such as the comprehensive generation of data from the environment, limited local computing power, and distributed data processing, which are associated with specific requirements and risks for data protection and information security. Further instruments are currently needed to address the challenges in edge-cloud applications by means of PETs.

3 The Problem Space – Requirements and Risks for Data Protection and Information Security in Edge-Cloud Applications

Protecting data and ensuring information security in edge-cloud systems involves specific challenges that go beyond the general problems of data-driven applications. This section provides an understanding of the specific challenges involved, addressing their context and root causes. To this end, it first outlines the types of sensitive data and information generated in edge-cloud systems (Section 3.1). It then presents key requirements for data protection and information security in edge-cloud systems and describes corresponding risks (Section 3.2). The diversity of data and information needing protection, the associated requirements and pertinent risks result in a complex and high-dimensional problem space. A comprehensive analysis of this problem space during the design and development of edge-cloud systems forms the basis for identifying suitable privacy-enhancing technology (PET) tools for implementing data protection and information security.

3.1 Sensitive data and information in edge-cloud applications

The direct collection and integration of data from physical devices, applications, and their environment generates dedicated types of sensitive data and information in edge-cloud systems. These can differ greatly from the personal data collected in, e.g., form-based applications. Table 2 provides an overview of sensitive personal data and information. In the following, these types of sensitive information are analyzed in the context of edge-cloud systems.

Sensitive personal data	Sensitive company data and information
- Identifiers	- Intellectual property
- Other characteristics that can distinguish individuals	- Data from customer relationships
- Pseudonymized data that can be linked to an individual	- Non-public financial figures and statistical data
- Metadata and behavioral data	- Data that must be protected based on other industry-specific regulations
- Unsolicited, randomly generated personal data	

Table 2: Overview of data and information requiring protection

Sensitive personal data

The GDPR requires comprehensive protection of personal data. Personal data is defined as any data relating to natural persons or data that can be linked to such persons with reasonable effort. First and foremost, this includes identifiers that can uniquely identify a natural person. Examples of identifiers are passport numbers or mobile phone numbers. In edge-cloud applications, further types of identifiers emerge. These include device identifiers that can be uniquely assigned to natural persons. Examples of such identifiers are IMEI/IMSI, MAC and Bluetooth IDs, IP addresses, license plate numbers, or badge numbers. In addition, biometric identifiers are generated

and used in a number of edge-cloud applications. Such identifiers arise primarily in scenarios such as camera surveillance, access control systems, time recording terminals, or wearables.

There are also other data attributes that are not directly assigned to a person but can provide information about a person in certain contexts and based on other available data or a combination of data attributes. These include processing operations on machines that can be assigned to individual persons based on the knowledge of the shift schedule.

Various pseudonymization techniques are used in edge-cloud systems to work with personal data. Data is considered pseudonymized if the attributes associated with the pseudonym are not sufficient to identify the person behind it and the pseudonymization cannot be reversed without disproportionate effort. However, if the data pseudonymized at the edge is merged with other data at a central location, such as the cloud, identification may be possible due to the larger number of available attributes. An increasing number of attributes associated with a pseudonym increases the likelihood of re-identification.

Edge-cloud systems often generate large amounts of metadata. Metadata is created during the processing, transmission, and storage of data between end devices (edge), intermediate nodes (fog/edge servers), and the cloud, and provides contextual information that is crucial for control, evaluation, and security in edge-cloud systems. It includes, for example, the location of data generation for mobile objects such as smartphones or autonomous vehicles, timestamps for data transmission, or access data for data retrieval in the cloud. Metadata can therefore also contain sensitive information.

Finally, the possibility of generating unsolicited personal data in edge-cloud applications should also be considered. Unsolicited personal data refers to personal data that is generated during the runtime of the application without the intention of recording this data in the application design. For example, audio and video recordings may capture unintended sensitive content if individuals are present in the recording areas without authorization or by accident. Further, if error logs created for machine statuses are filled with personal data by operators, personal data can be generated in an unforeseen way.

Sensitive company data and information

Sensitive company data and information should also be protected in edge-cloud systems. Sensitive data or information is defined as any data or information whose disclosure could cause economic, legal, or strategic damage to the organization or its business partners. Possible damages include loss of revenue, damage to reputation, regulatory penalties, the creation of vulnerability to blackmail, the interruption of value-added processes, and the investigation costs to reconcile the disclosure.

Data describing *intellectual property* is a category that is particularly worthy of protection. Intellectual property primarily includes information relating to a company's products. This includes technical documentation, design plans, and software source code. In addition, trade secrets such as production processes, pricing mechanisms, market research, and unpublished research results are also considered intellectual property. Trade secrets are protected under the EU Trade Secrets Directive and comparable international laws, provided that appropriate confidentiality measures are taken. In edge-cloud applications, data subject to trade secrets is created in domains such as industrial production. For example, machine parameters are collected for condition monitoring, which can provide information about production processes or part geometries.

Furthermore, data containing information about corporate customers represents another category worthy of protection. This applies not only to customers' intellectual property, but also to elements such as contract details, service histories, communication content, and usage data. Such data is often protected by bilateral agreements such as confidentiality agreements. Third parties could use such data to generate economic advantages or uncover possible misconduct on the customer side. Here, too, data from machines and systems is a relevant example. For example, job shops could potentially disclose their customers' part geometry to providers of machine-related services.

Internal financial information and other statistical information, such as sales forecasts, cost structures, investment plans, or data on mergers and acquisitions, are also of strategic importance to companies. Their disclosure can lead to market manipulation, insider trading, or competitive disadvantages. In edge-cloud systems, such data is potentially at risk if recorded production numbers or energy consumption data allow to draw conclusions about the economic situation of a company.

Finally, there are industry-specific regulations that impose additional requirements on the protection of certain company data and information and prohibit their aggregation. This applies in particular to critical infrastructures, the financial sector, and the medical sector. Also these regulations can impose certain restrictions and requirements on how data is being handled in edge-cloud applications.

3.2 Specific requirements and risks for data protection and information security in edge-cloud systems

In addition to the general requirements and risks for data protection and information security that arise from the application context, the implementation of an application through edge-cloud systems results in further challenges. For example, the combination of different components along the edge-cloud continuum or the distribution of data processing across different partners along the data value chain results in specific risks and requirements for information security. Possible risks and requirements for data protection and information security in edge-cloud systems are shown in Figure 3 and explained below.

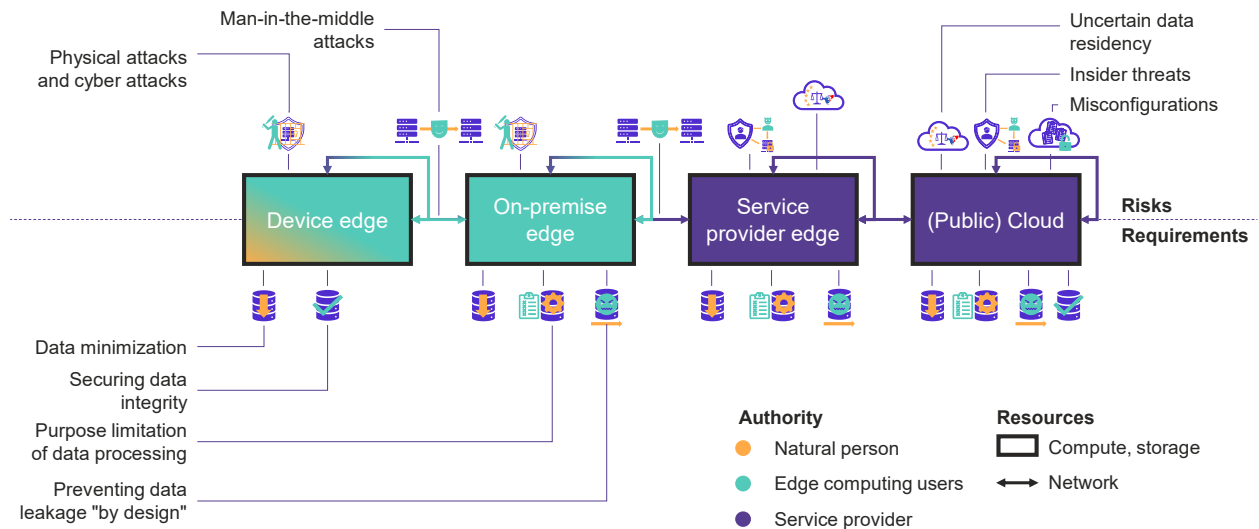


Figure 3: Risks and requirements for data protection and information security in edge-cloud systems

Increased vulnerability to physical attacks and cyberattacks

Operators of large data centers, most of which also support cloud service providers as part of co-location offerings, establish comprehensive security concepts that regulate and restrict access to physical computing resources. The utilized security mechanisms include access controls, video surveillance, and structural protective measures. With the use of edge devices, computing resources are placed in decentralized locations that are often easier to access and more difficult to monitor (Chang & Wu, 2021). This limited physical security gives rise to various risks. On the one hand, edge devices are vulnerable to theft and intentional and unintentional damage, for example due to accidents, weather events, or operator error. If an edge node is physically destroyed, this can lead to service interruptions. Furthermore, there is a risk of physical manipulation, whereby attackers can modify the hardware, install malware, or extract local data carriers and read data to gain access to sensitive information. In addition, edge devices can become a gateway for attackers if attacks on an edge device are extended to other connected networks and the perpetrators spread laterally until critical systems are compromised (Sheikh et al., 2025). Attackers can, for example, place backdoors and then activate them later from outside. These possibilities are facilitated by the fact that many edge devices do not have secure hardware roots of trust and are rarely equipped with tamper-proof housing or self-destruction mechanisms.

Danger from man-in-the-middle attacks

Another risk is the vulnerability of edge-cloud systems to network attacks. Edge-cloud systems feature a wide variety of combinations of local network components (LAN: Ethernet, WLAN, 5G campus networks) and wide area networks (WAN) connecting a large number of devices, storage and network components, and the corresponding services in both horizontal and vertical directions. Especially the local network level is susceptible to attacks. A relevant attack pattern is the man-in-the-middle attack. In a man-in-the-middle attack, the attacker positions themselves between two communicating parties in order to intercept, alter, or manipulate data without the parties noticing. In addition to the data content itself, also metadata that can reveal relevant information about users might be subject to such attacks. If data is manipulated by an attacker, the system behavior can be compromised. In edge-cloud systems, these risks are amplified by the distributed locations of devices and different network types, which increase the attack surface. Additionally,

the use of resource-limited IoT devices can hamper security. These devices often have only limited security mechanisms and are therefore particularly vulnerable to attacks if configured incorrectly (Ali & Al-Sharafi, 2025).

Challenges posed by insecure data residency

The use of cloud services and other centralized third-party services in edge-cloud systems offers companies high flexibility and scalability as well as low data processing costs, but also entails specific data protection and security risks. One of these risks is uncertain data residency (data sovereignty). Data residency encompasses the physical location of data transferred to the cloud and the legal possibilities for accessing it. European companies must ensure that personal data meets the requirements of the GDPR and is not transferred to other jurisdictions without permission (European Commission, 2023). When leveraging non-European service providers, it must be ensured that personal or other sensitive data is only processed and stored on the agreed servers in Europe. This is in potential conflict with legislation in other regions, first and foremost, the US CLOUD Act (115th Congress [2017-2018], 2018). The CLOUD Act allows US law enforcement agencies access to data stored in the clouds of US-based service providers, even if this data is processed in Europe (e.g., by non-US subsidiaries). Access is possible, for example, for law enforcement purposes or in cases of national security concerns. As a result, US cloud providers are currently promoting new offerings in Europe with a stronger focus on data sovereignty. However, even when relying on these offerings, data transfer under legal pressure cannot be ruled out.

Danger from insider threats

Another risk category on the part of cloud service providers or other third parties is insider threats (Chang & Wu, 2021). Insiders are defined as individuals who hold or have held a position within the cloud service provider. These can be either internal employees or contractors who, due to their insider status, have legitimate access to sensitive or critical data. Such individuals are therefore able to intentionally or unintentionally access and manipulate sensitive data or information. Often, there are no direct detection mechanisms for such activities.

Increased risk of misconfigurations

Also, misconfigurations are a relevant cause of data leaks on the side of cloud service providers. Misconfigurations can affect the entire cloud tenant, i.e., the entire customer account, or individual cloud services, such as analysis or storage services. Among the most common misconfigurations are incorrectly set access rights, openly accessible storage services, or insufficiently secured secrets (digital access data for non-human users) and user accounts. In previously known incidents, personal data, financial information, and health data, among others, were made publicly accessible (Chang & Wu, 2021). For example, Verizon placed data from approximately 6 million customers in a public AWS S3 storage¹, while Decathlon disclosed approximately 9 GB of data, including sensitive personal data, through a misconfigured Elasticsearch server². While the companies affected mostly describe these situations as the misconduct of individual employees, the complexity

¹ <https://www.cbsnews.com/philadelphia/news/verizon-data-leaked-online/>

² <https://www.computerweekly.com/news/252479101/Sports-retailer-Decathlon-left-employee-data-exposed>

of modern cloud platforms and the lack of uniform, consistent configuration standards generally increase the risk of such incidents.

In addition to the specific data protection risks relevant to edge-cloud systems, organizations also express requirements for their information security, which may arise from legal, business, and social conditions, as well as from the individual characteristics of the business model and the underlying architecture. These requirements are elaborated below.

Requirement for data and information minimization

Data minimization of personal data is formulated as a fundamental principle in the GDPR (Art. 5). On the other hand, minimization can also refer to the processing of company-related information. Data minimization generally means only collecting, processing, and sharing data that is appropriate and necessary for the purpose of an application (ISACA, 2024). This principle is becoming increasingly important in edge-cloud systems, as data is typically collected and processed in a distributed manner and must be transferred across various nodes along the edge-cloud continuum. In addition to ensuring regulatory compliance, data minimization reduces the attack surface, as less information can be disclosed in the event of security incidents. Furthermore, it reduces the scope for potential data misuse by partners along the data processing chain. Accordingly, once data minimization has been implemented, processing by less trustworthy third parties may also be possible. A positive side effect of data minimization is that it reduces technical complexity by reducing storage requirements or network load, thereby saving energy. Ultimately, the minimization of data and information may not only be a requirement of the data subjects or data-providing organizations, but also of the data processors to simplify data processing. For example, this can reduce the necessary protective measures.

Need to ensure data and information integrity

Ensuring integrity and confidentiality is another principle for the processing of personal data formulated in the GDPR. Data integrity refers to consistent accuracy, completeness, and quality of data throughout the entire data processing chain. Data integrity is an important criterion for the processing of personal data. If data does not accurately reflect the characteristics of a natural person, this can result in significant disadvantages for that person. Examples include inappropriate treatment due to incorrect health data or financial disadvantages when calculating creditworthiness. Integrity is also important when sharing company-related information, for example, when billing in subscription models is to be based on IoT data or orders are to be distributed based on the availability and cost of resources. In edge-cloud systems, integrity plays an important role as data is collected and used directly from sensors or actuators (Adam et al., 2024). Further, the correct functionality of autonomous systems can only be ensured when the provided input data is accurate and complete. Like data minimization, data integrity is often a requirement of data subjects, data-providing organizations, and data processors.

Necessity of purpose limitation in data processing and information use

The purpose limitation of data processing means that data may only be collected and processed for the specified and legitimate purpose (ISACA, 2024). Purpose limitation is also a principle formulated in the GDPR. The GDPR states that personal data must not be processed in a manner that is incompatible with the predefined purpose. With regard to business information, it must be ensured that sensitive information is used by data processors exclusively for the agreed purposes.

Using data for previously undefined tasks may result in negative consequences for the data subjects. For example, health data could be misused for insurance risk assessments or marketing purposes. When processing sensitive business information, purpose limitation is important, among other things, to protect trade secrets and avoid competitive disadvantages. For example, order and utilization data, which is intended to optimize the supply chain performance, could be used as an unfair competitive disadvantage in price negotiations. Examples of the need for purpose limitation in data processing in the context of edge-cloud applications include services that collect personal vehicle data and should not make it available for insurance services, or pay-per-use business models in which consumption data collected at the edge should be used exclusively for billing purposes. Traditionally, the purpose limitation of data processing is achieved through organizational and contractual concepts such as data classifications, authorizations, and retention periods. Such design measures are closely related to the concept of data sovereignty, i.e., the ability of data-providing organizations or persons to decide on the use of data and information throughout the entire data lifecycle (Scherenberg et al., 2024).

Need to prevent data leaks by design

Data usage agreements and supplementary organizational measures are often sufficient to convince data-providing organizations and data subjects that the purpose limitation of data processing will be complied with. However, if there is little trust in the data-processing party or if this party has strong incentives to use the data obtained for additional purposes, such measures are not sufficient. Instead, data-providing organizations are required to prevent data leaks and the use of information for undefined purposes through the application architecture. Appropriate "privacy by design" measures use technical means to ensure that data can only take predefined paths (Schiffner, 2015a). The implementation of such technical design mechanisms is particularly useful in edge-cloud systems, where there is greater freedom in terms of the location of data processing, the selection of components and partners, and the data to be collected and processed. For example, a large proportion of the data can already be processed locally, i.e., in the environment controlled by the data provider. If data leaks can be prevented "by design," there is a greater chance of success in implementing applications that access critical data from external parties.

4 The Solution Space – PET Tools for Implementing Data Protection and Information Security in Edge-Cloud Applications

Based on the requirements and risks for data protection and information security in edge-cloud applications described in Section 3 (problem space), this section supports the design of Privacy-enhancing technology (PET)-based solutions to address these challenges. The section first provides an overview of generally available PET tools for the technical implementation of data protection and information security (Section 4.1). Afterwards, it presents existing solution aids that support the requirements analysis and system design, piloting and implementation, and evaluation and testing of PET tools in real application scenarios (Section 4.2). Finally, this section presents practical application scenarios of PET tools to meet data protection and information security requirements in edge-cloud applications in various industries (Section 4.3). The presented application scenarios are based on insights from the practice projects that are part of the technology program “Edge Data Economy,” funded by the German Federal Ministry of Research, Technology and Space (BMFTR). The application of PETs is presented in the context of real-world edge-cloud applications and their specific implementation requisites and implications are described from the perspective of the early adopters involved. These insights provide new knowledge on how PETs can be used in the practical context of edge-cloud systems and what contextual factors need to be managed, thereby providing valuable information for companies developing their own PET strategies.

4.1 PET tools

Various PETs could potentially be used to technically ensure the protection of data and information in edge-cloud systems. Figure 4 provides an overview of the PETs currently in widespread use, sorted according to the root of trust through which data protection and information security are generated, as well as the typical area of application of the PET along the data lifecycle. The types of PET tools are briefly described below. A more detailed explanation of the listed PETs can be found in the appendix to this study.

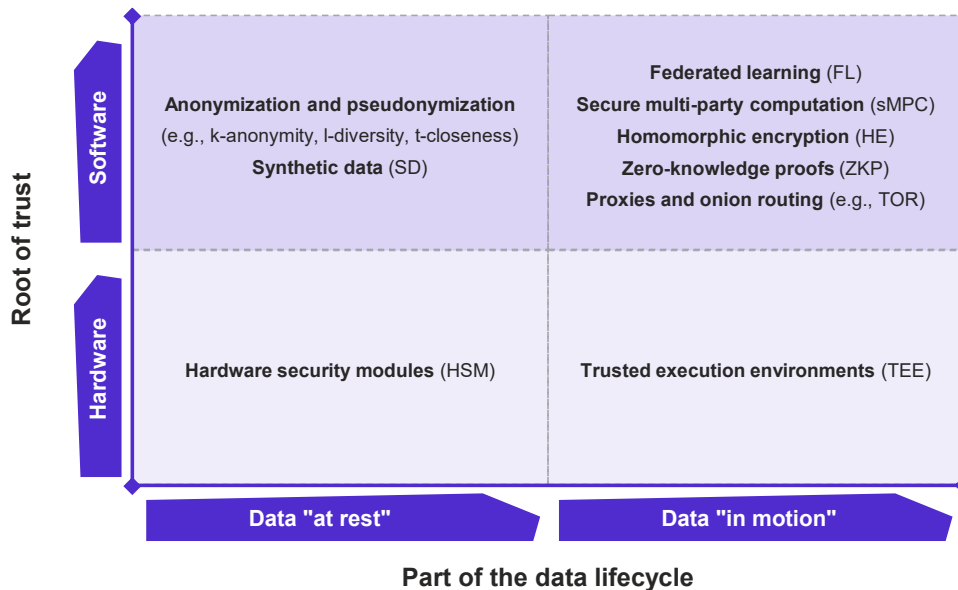


Figure 4: Overview of popular PETs

Anonymization and pseudonymization

Anonymization involves removing or generalizing attributes from data sets to prevent identification and make entities indistinguishable. Effectively anonymized data is no longer considered personal data. In pseudonymization, for example, identifiers in data records are replaced by pseudonyms. Typical approaches to pseudonymization include k-anonymity, l-diversity, and t-closeness. Due to the possibility of re-identification, pseudonymized data is still considered personal data. Both approaches are usually applied to data "at rest."

Synthetic data

Synthetic data is machine-generated data that mimics the statistical properties of the population and protects the privacy of the data generators. Today, machine learning methods are increasingly being used for this purpose. They learn the underlying distribution and generate new data from it. Synthetic data can be completely synthetic (all variables are generated by a model), partially synthetic (only some variables are synthesized), or hybrid (generated from the real set of data and a completely synthetic set).

Federated learning

Federated learning is a collaborative approach to machine learning. A global AI model is initially provided by a central server and distributed to various nodes/clients (e.g., servers, edge devices, smartphones). The nodes train the model locally with their own data and only transfer updated model parameters to the central server. The server aggregates these values, updates the global model, and returns the current global model to the clients.

Secure multi-party computation

Secure multi-party computation is a cryptographic approach that enables multiple parties who distrust each other to perform calculations based on pooled data without revealing the input data. To implement secure multi-party computation, the inputs of the participants are broken down and

distributed among the other parties so that no single party has access to all the information. The parties then perform their individual calculations and combine the calculation results at the end. This process requires a minimum number of participating parties. Other current challenges include low performance and the restriction to dedicated computing methods.

Homomorphic encryption

Homomorphic encryption is a cryptographic method that allows calculations to be performed on encrypted data without having to decrypt it. Fully homomorphic encryption allows any mathematical operations to be performed directly on the ciphertext. After the calculations are complete, the encrypted result can be decrypted to obtain the real result. Homomorphic encryption can be used in particular for calculations that can be expressed as polynomial functions, as well as for selected AI applications. Current limitations include increased computational effort and calculation inaccuracies (Stock et al., 2022).

Zero-knowledge proofs

Zero-knowledge proofs are a cryptographic concept that enables one party (prover) to prove the validity of a statement to another party (verifier) without having to provide the underlying data. Two types of methods for creating zero-knowledge proofs exist. Interactive methods create zero-knowledge proofs that require multiple communications between the prover and verifier. Non-interactive methods allow a one-time proof to be provided (Lavin et al., 2024). Established areas of application for zero-knowledge proofs include identity management and finance.

Proxies and onion routing

Proxy servers and onion routers control requests and communication between different parties. A proxy server mediates between the client and the target system, forwards requests and returns responses, and can conceal the identity of the source (e.g., IP address). Anonymous proxies and highly anonymous proxies are particularly relevant for privacy protection. Anonymous proxies hide the user's IP address, while highly anonymous proxies go a step further and conceal the fact that they themselves are acting as proxies.

Onion routing is a network technology that routes bidirectional communication via a chain of relays (onion routers) in order to decouple the sender and destination from each other. The aim is to protect both the content of the communication and the communication partners from surveillance. To do this, the messages are encrypted in layers by the client, similar to the layers of an onion. Each onion router in the connection chain removes only its own layer and forwards the message accordingly. A single onion router thus only knows its immediate predecessor and successor and has no access to the complete content or the entire route.

Hardware security modules

Hardware security modules generate, store, and manage cryptographic keys in tamper-proof hardware and offer functions for encrypting and decrypting data and generating digital signatures. They thus help to protect intellectual property and sensitive data from unauthorized access. The cryptographic keys are created and remain in the crypto processor. Hardware security modules provide access to these keys via defined APIs. In case of physical attacks, data protection is

ensured by deleting sensitive data, amongst other things. Hardware security modules are primarily used in areas with high compliance requirements and are often combined with other PETs.

Trusted execution environments

Trusted execution environments are hardware architectures that enable applications to be executed in isolation from the underlying operating system. To do this, they use specialized CPU solutions that prevent access to the memory area. They create a trustworthy processing environment in which data and processes are protected from unauthorized access, manipulation, or monitoring, even if the host or underlying operating system has been compromised. Even operators of these environments cannot view the data. The use of remote attestation also allows users to verify the integrity and authenticity of the code and data.

4.2 Summary of existing guidance on PETs

There are already a number of publications that can help practitioners select, design, and implement PETs to ensure data protection and information security. These publications mainly take an application-agnostic approach. An overview of the available resources is provided in Table 3. The publications provide guidance that can be valuable in the areas of requirements analysis and system design ("When can which PET provide support?"), piloting and implementation ("How can a PET-based application be implemented?"), and the evaluation of PET ("What criteria can be used to evaluate a PET?").

A handful of resources provide a detailed overview of individual PETs that go beyond the descriptions in this publication. These overviews stem from different perspectives, such as privacy by design, data protection engineering, or compliance with data protection regulations. In addition, ENISA (2022) shows which technologies can be used complementary to PETs. The Centre for Data Ethics and Innovation (2021) provides a decision tree that can be used to identify suitable PETs for individual data protection needs.

Assistance with piloting and implementing PETs is primarily provided by the presentation of PET use cases. Examples of this include the work of the Centre for Data Ethics and Innovation (2021), the United Nations (2023) and Noble (2023), with the latter focusing on the use of PET for public authorities and in the public sector. Noble (2023) also provides an overview of relevant standards in the context of PET implementations.

For the evaluation and testing of PET tools, the Future of Privacy Forum (2024) presents a summary of existing test environments. Noble (2023) provides information on individual evaluation aspects of PET as part of the presentation of PET use cases. To support the development of a catalog of criteria for evaluating PET technologies, ENISA (2016) has formulated a framework that includes several criteria and indicators.

Resource	Content	Requirements analysis and system design	Piloting and implementation	Evaluation and testing
Klymenko et al. (2025)	Detailed overview of PETs; recommendations for implementing privacy by design with PETs	x		
OECD (2023)	Overview of PETs sorted by their intended use (data obfuscation, encrypted data processing, federated analytics, data control tools)	x		
ENISA (2022)	Overview of PETs in the context of data protection engineering; explanation of complementary technologies for the technical implementation of data protection.	x		
ISACA (2024)	Overview of PETs with a focus on standards and compliance with regulatory data protection requirements	x		
Centre for Data Ethics and Innovation (2021)	Decision tree for selecting suitable PETs; overview of PET use cases	x	x	
United Nations (2023)	Overview of PETs; PET use cases for public institutions	x	x	
Future of Privacy Forum (2024)	Overview of regulatory activities, studies and test environments related to PETs	x	x	x
Noble (2023)	Overview of PETs and their comparison; overview of standards related to PETs; PET use cases in the public sector	x	x	x
ENISA (2016)	Framework for analyzing PETs based on various criteria and indicators			x

Table 3: Overview of solution aids for the development of PET-supported applications

4.3 Scenarios for applying PET tools in edge-cloud applications

Going beyond the general guidance presented above, this section analyzes PET tools in the context of edge-cloud systems. It describes how PETs in various application domains minimize the risks to data protection and information security in edge-cloud systems described in Section 3. The presented findings are based on research and development work carried out within the research projects in the technology program "Edge Data Economy," funded by the German Federal Ministry of Research, Technology and Space (BMFTR), and cover relevant solutions within the program. The presentation of the PET tools comprises their technical implementation as well as the integration of the technology into the application context, including the associated organizational, structural, and cultural consequences. Further, the feasibility of implementing each PET tool in its application context is evaluated by the early adopters from the research projects and further implications of their implementation are identified. Table 4 summarizes the PET tools presented in this section and the risks and requirements motivating their implementation in the respective application domain.

Addressed requirements and risks for data protection and information security	PET tool	Application domain within the technology program "Edge Data Economy"
Data minimization, physical and cyber attacks	Hardware keys (4.1)	Agrifood industry
Data minimization, physical and cyber attacks	Federated learning (4.2)	Industrial manufacturing
Purpose limitation, preventing data leakage by design, uncertain data residency, insider threats, misconfigurations	Compute-to-data (4.3)	Industrial manufacturing
Data minimization, data integrity	Zero-knowledge proofs (4.4)	Energy industry
Purpose limitation, data integrity, preventing data leakage by design, insider threats	Trusted execution environments (4.5)	Energy industry

Table 4: PET tools and addressed requirements and risks for data protection

In the following sections, the individual PET tools are described according to a uniform scheme. First, the PET tool and its motivation are presented in the context of its application area. Specifically, the challenges for data protection and information security as well as other requirements arising from the concerns of stakeholders in the application context are explained. In addition, the technical functioning of the PET tool is explained and the roles required for the implementation of PET in the application context are described. Subsequently, the prerequisites for implementing the PET tool in the application domain are analyzed and the effort required to create these prerequisites is assessed from the perspective of the consulted practitioners³. The possible requisites are divided into the following categories:

- *Organizational requisites* include, for example, setting up organizational structures, creating incentives, defining roles and responsibilities, and establishing partnerships.
- *Technical requisites* include, for example, the technical specification or the development, integration, and maintenance of the application.
- *Requisites in terms of skills and understanding* include, for example, acquiring knowledge and expertise about the PET tool or building trust in the technology used.
- *Requisites for operation and integration* include, for example, adapting ongoing processes or ensuring compatibility with existing protocols and services.
- *Regulatory requisites* include, for example, ensuring legal compliance in the use of the technology in the application context, and documenting this appropriately.

Finally, the intended and unintended implications caused by the use of the PET tool in its application context are analyzed. Implications may arise, for example, about future cooperation between different actors, long-term costs and expenses, or resource consumption. In this context, *enablers* are understood as positive implications that can provide additional motivation for the use of a PET tool. Implications that could hinder the widespread adoption of the technology are referred to as *constraints*.

³ The early adopters' assessment was based on a six-point Likert scale ranging from "very easy to implement" to "very difficult to implement." The assessment only provides an impression of the relative complexity of the implementation requisites for each application context from the users' perspective and cannot be generalized across the various applications presented.

4.3.1 Hardware keys in the context of quality control in the agrifood industry

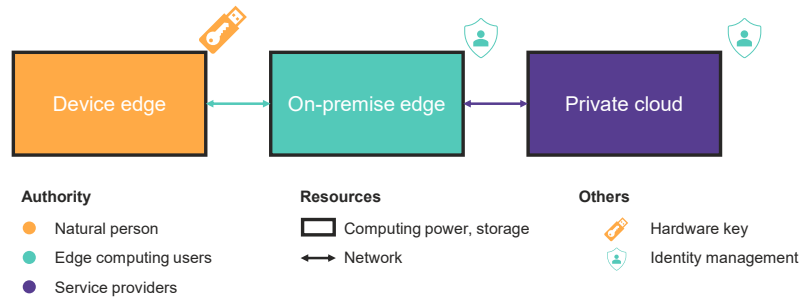


Figure 5: Schematic representation – use of hardware keys in the agrifood industry

In edge-cloud systems, hardware keys are used to achieve secure authentication of edge device users without having to record personal data that is generated during standard login procedures using a username and password. Hardware keys are a form of hardware security modules. In this agrifood application (see Figure 5), quality inspectors from a service company use edge computing-based scanner systems (device edge) to check the quality of fruit and vegetables from a food producer at certain stages of the food supply chain, for example after the harvest in food packaging. The quality data is then transferred to the food company's software systems for long-term documentation. The software systems are provided either in a private cloud or via on-premise edge systems. From there, the collected data can be made available to other authorized parties along the supply chain. The food producer requires personal data to be removed to facilitate the subsequent data handling processes. Additionally, it is important that the quality service company's employees can be easily included in the authentication method. In addition, the authentication method should be integrable into the food manufacturer's existing identity management system and the authentication effort for users should remain low. Furthermore, authentication should also work in areas with limited network access, as this is often the case in agricultural settings.

During implementation, application administrators first configure the hardware keys. The WebAuthn standard⁴ is used to integrate the hardware keys with the existing identity management infrastructure. A cryptographic procedure (challenge-response procedure) ensures that each key can be uniquely assigned and that authentication credentials cannot be extracted from the hardware key and are therefore inaccessible to third parties. The configured hardware key can be used to enforce robust access control at various levels, from services at the device edge to the on-premises edge to the cloud, without relying on personal data. Finally, the hardware keys are distributed to the quality service provider, who makes them available to the relevant employees. The employees use the hardware key to authenticate themselves to the scanner system through the associated quality management system's web interface. Depending on the application context, the quality management system can be operated at the edge, in the cloud, or in a distributed manner.

⁴ <https://www.w3.org/TR/webauthn-3/>

Implementation requisites



Figure 6: Evaluation of implementation requisites – hardware keys in the food industry | Subjective assessment of the relative complexity of implementation requisites by early adopters of PETs in the presented use case, measured on a six-point Likert scale ranging from "very easy to implement" to "very difficult to implement."

A key aspect of successful implementation is the integration of the hardware key-supported system into the existing processes. According to the early adopters, the complexity of this integration depends largely on the existing degree of digitization of the processes under concern. In environments where digital user management systems are already in place, integration is relatively straightforward and usually only requires adjustments to existing authentication modules. However, integration into analog structures, such as paper-based user registrations, is much more complex. In such cases, the introduction of the solution goes hand in hand with a comprehensive redesign of processes and becomes part of a fundamental digitization project that entails challenges beyond the specific technology.

Additionally, according to the early adopters, there is little need to build up special skills and understanding on the user side. Hardware keys are used in the same way as other physical keys, and modern operating systems automatically support authentication processes using hardware keys. A certain level of web development expertise and knowledge of client-server architectures is required for technical solution development. However, the need for knowledge can be greatly reduced by using existing frameworks and libraries. From a legal perspective, the use of hardware does not require any special considerations according to the early adopters.

Implications

The use of hardware keys in the quality control of fruit and vegetables in the food industry has the following implications:

Enablers	Neutral	Constraints
<p>Future viability: Due to the low hardware costs and the support of the solution by companies such as Apple and Google, it is foreseeable that security chips will be installed as standard on many end-user devices such as smartphones and tablets, which will then serve as hardware keys themselves.</p>	<p>Costs: Initial investments in hardware (currently approx. €100 per key) are necessary. On the other hand, IT costs for user administration and additional expenses for employee hardware can be avoided.</p>	<p>Wear and tear: As with conventional keys, hardware keys can also be lost or damaged. This results in additional costs for blocking or reissuing authorizations and may temporarily restrict the usability of the system.</p>
<p>Usability: Users do not have to remember passwords, and login is quick and easy. Authentication is consistent, platform-independent, and reduces human error during login.</p>		<p>Culture and collaboration: Low acceptance can result if the solution is introduced "from above" without clearly communicating the benefits. In addition, uncertainties about actual security can lead to mistrust among many users.</p>
<p>Areas of application: The authentication supported by hardware keys can be used in a wide variety of contexts and is not limited to the application of quality testing for fruit and vegetables.</p>		
<p>Economic efficiency: Currently, the number of hardware key providers is still small. In the long term, it can be assumed that competition in this area will intensify, potentially reducing costs.</p>		

Table 5: Implications – Hardware keys in the food industry

4.3.2 Federated learning in the context of industrial manufacturing

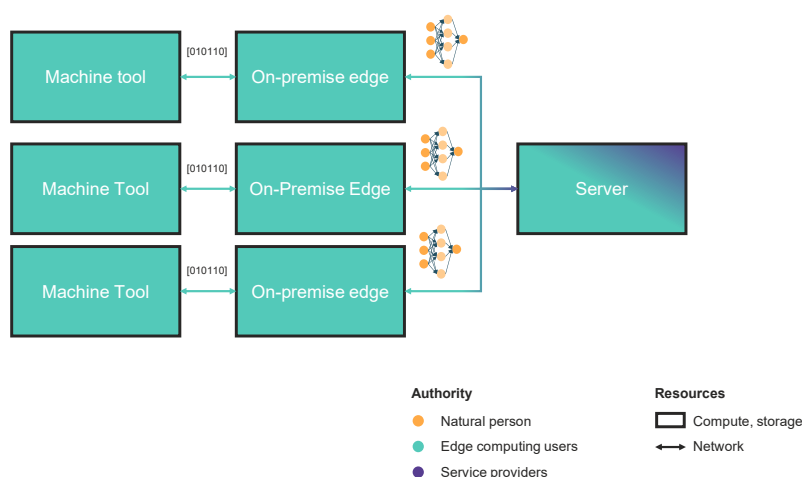


Figure 7: Schematic representation – Federated learning in industrial manufacturing

Federated learning allows for the training of AI models using data from distributed clients without having to transfer raw data to a central location. Instead, only model updates or gradients are sent to an aggregator, which combines them into a global model. The updated global model can then be distributed to all participating clients so that they can learn from each other. In the given

context, federated learning is being tested for an internal use case in a multinational industrial enterprise. The goal is to use data sources distributed across different production sites for model training without having to centralize or directly exchange sensitive machine data. This is particularly relevant for applications where there are significant technical and organizational barriers to accessing production data. Specifically, this can be the case when security zones within the plants restrict access or when factory managers resist making production data available. Furthermore, federated learning allows individual clients to learn from each other. This enables, for example, the detection of errors that previously only occurred on individual machines for all machines participating in the federated learning system. In the future, it is conceivable that federated learning could also be used directly in product service systems to create continuous product improvement without having to disclose sensitive information. To the satisfaction of users, the federated learning process should integrate seamlessly into the application and be fully automated.

In the given federated learning scenario (see Figure 7), the process data (raw data) is collected from the machine tool and loaded into a more powerful edge environment (on-premise edge) to satisfy, amongst others, the requirements for computing power. At the edge, the data is processed by the initial AI model to perform AI inference and model training. The model updates or gradients are then transferred to a central orchestrator that merges the gradients of the individual machine tools. The orchestrator can be implemented, for example, in the private or public cloud or on a local server computer. Finally, the updated model is distributed back to the edge environments, ready for the learning cycle to begin again.

Implementation requisites



Figure 8: Evaluation of implementation requisites – Federated learning in industrial manufacturing | Subjective assessment of the relative complexity of implementation requisites by early adopters of PETs in the presented use case, measured on a six-point Likert scale ranging from "very easy to implement" to "very difficult to implement."

In particular, the technical implementation of federated learning in industrial production is considered challenging by early adopters (see Figure 8). A key technical obstacle lies in the heterogeneity of the machines, products, and processes in industrial manufacturing. This leads to strongly heterogeneous (non-IID) data distributions across clients, which substantially complicates the aggregation step and hinders the training of a robust, well-generalizing global model. At the same

time, frameworks and technology stacks that go beyond the actual federated learning technologies are necessary to ensure seamless communication and continuous updates of the federated learning network, for example, to generate a good user experience. Accordingly, comprehensive development efforts are required that go beyond solving the actual problem of privacy-preserving machine learning.

According to early adopters, this also means that integrating and operating the solution is a rather complex undertaking. On the one hand, the implementation of federated learning applications is often not an immediate step, but rather a subsequent stage in the data-driven transformation of industrial production. Before starting federated learning endeavors, data has to be made available and potentially analyzed by more mundane technologies. A particular challenge for the successful operation of federated learning solutions is to attract the necessary number of participating entities for the federated learning use case. The practical benefits of a federated learning applications with regard to mutual learning only emerge with a broad participation of numerous machines or locations. If only few clients participate in such a system, it is hard to demonstrate the benefit to further interested parties. Once a critical mass has been reached and the added value of federated learning becomes apparent, it will be easier to convince other parties to participate. This challenge resembles the chicken-and-egg problem present in other platform business models.

According to the early adopters, the organizational complexity of introducing federated learning within industrial companies depends primarily on existing organizational structures. As discussed earlier, the sustainable success of federated learning requires the involvement of a large number of internal and external stakeholders. As the size and complexity of an organization increases, so does the number of stakeholders and, with it, the coordination effort required. However, federated learning does not differ significantly from conventional technology initiatives in this respect. Accordingly, provided there is strong management support, federated learning technologies can also be introduced quickly and successfully within companies.

Large industrial enterprises generally employ enough qualified personnel to implement federated learning applications. While initial approval for pilot projects (e.g., by individual innovation drivers) is often easy to obtain, building broader support for an organization-wide roll-out beyond the pilot projects proves to be much more difficult due to the previously elaborated chicken-egg problem. Therefore, good project management skills and powers of persuasion are needed.

There are currently no fundamental restrictions on the regulatory-compliant implementation of federated learning in industrial scenarios. However, depending on the design of the application, additional industry-specific compliance requirements and tax and liability issues must be considered, such as those relating to the use and processing of data across national borders. For instance, model weights may be considered intangible assets. Sharing these assets across national borders constitutes a cross-border exchange of services that must be priced accordingly even when occurring within an enterprise.

Implications

The use of federated learning for internal analysis of machine data in industrial production has the following implications:

Enablers	Neutral	Constraints
<p>Economic</p> <p>Federated learning enables the development of better products, provided that a sufficiently large number of participants (machines) take part in the federated learning ecosystem. The aggregated use of decentralized data creates a broader knowledge base, which can lead to higher-quality models and benefits for all parties involved.</p>	<p>Usability:</p> <p>Ideally, federated learning should be invisible to end users, causing no noticeable restrictions and having no implications for the user experience.</p>	<p>Legal implications:</p> <p>In cross-company federated learning projects, uncertainties regarding data sovereignty, responsibility, value distribution, and industry-specific regulations may exist, limiting its application.</p>
<p>Future viability:</p> <p>The integration of federated learning services in product-service systems can create novel benefits and thus serve as a differentiating feature, as well as enable novel and sustainable business models.</p>	<p>Resource consumption:</p> <p>The actual efficiency, resource consumption, and performance of the application must be evaluated on a case-by-case basis. There is currently still uncertainty about how resource-efficient federated learning can be under real-world conditions, especially when many clients are involved.</p>	
<p>Culture and collaboration:</p> <p>The economies of scale of federated learning make collaboration between companies or business units a prerequisite. Parties involved are encouraged to recruit additional trustworthy partners to increase the accuracy and robustness of the application.</p>		

Table 6: Implications – Federated learning in industrial manufacturing

4.3.3 Compute-to-data in the context of industrial manufacturing

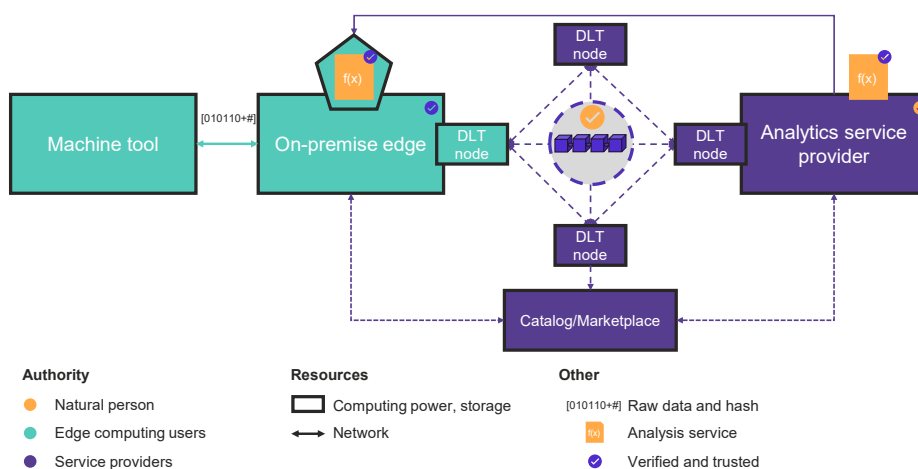


Figure 9: Schematic representation – Compute-to-Data in industrial manufacturing

The core concept of the compute-to-data approach is to reverse the conventional practice of transferring data to an external analysis service for data processing or insight generation. Instead,

analysis services are transferred to and executed within an environment under the control of the data-providing organization. In the manufacturing industry, this means that services based on machine or machine component data can be executed without having to disclose sensitive information such as geometries or process data to other parties. In addition, no personal data leaves the data provider's environment, which strengthens legal feasibility and augments the data provider's data sovereignty. In the presented use case, a component manufacturer's analysis applications are made available to the machine operator (who is also the component user) via a compute-to-data framework.

Beyond data protection and security, the application of compute-to-data paradigms in industrial manufacturing introduces additional requirements. First, mechanisms must be established to ensure the trustworthiness of all participating parties, including both service providers and service consumers. Second, it is often critical to verify the integrity and correctness of input data collected at the machine level. This requirement becomes particularly salient in scenarios such as subscription-based business models, where equipment manufacturers depend on accurate condition monitoring and usage data to deliver their services effectively.

To implement compute-to-data approaches, the given use case (see Figure 9) leverages the Pontus-X ecosystem⁵. Pontus-X builds upon the Ocean Protocol and leverages Web3 technologies, employing smart contracts to technically enforce rules, guidelines, and interaction protocols for data sharing. The underlying distributed ledger technology (DLT) network is operated collaboratively by a consortium of public and private institutions adhering to jointly agreed-upon governance principles. Notably, this federated governance model allows manufacturing companies to actively participate in network operation, although such participation is not a prerequisite for utilizing the network's services. To access compute-to-data capabilities, prospective service providers and consumers must first complete an onboarding process into the Pontus-X ecosystem. This requires a Gaia-X credential, which contains verifiable statements about the identity, characteristics, and conformity of participants in the Gaia-X ecosystem⁶. These credentials serve to establish organizational trustworthiness. Following onboarding, data and service offerings can be made available and their terms of access and use described via the Pontus-X ecosystem.

Once onboarded, a manufacturing company seeking to procure compute-to-data services must first discover and select an appropriate analysis service through the network's decentralized marketplace, submitting a usage request if required. The service metadata accessible through the marketplace specifies both technical and organizational prerequisites such as required input data formats and execution environment specifications that must be satisfied prior to service invocation. To execute the service, the manufacturing company (or an authorized third party) provisions a secure execution environment within a Kubernetes cluster⁷ at the on-premise edge, utilizing Ocean Protocol libraries⁸. Subsequently, an asset entry is created in the Web3 network, containing a reference to the production data designated for processing within the corresponding cluster. After the service provider and machine operator have automatically concluded an agreement on service use, the analysis application is transferred to the secure environment and executed with the provided production data. To ensure data integrity, a cryptographic hash may be generated at the machine level and verified prior to service execution, thereby detecting any unauthorized

⁵ <https://www.pontus-x.eu/>

⁶ <https://gaia-x.eu/>

⁷ <https://kubernetes.io/>

⁸ <https://oceanprotocol.com/>

modifications during data transfer from the machine to the secure execution environment. Throughout execution, the service provider can monitor application status via automatically generated logs. Upon completion, analysis results are made available to authorized parties according to predefined access policies. For standard analytical procedures, results may be restricted exclusively to the manufacturing company. However, in scenarios involving AI model training, derived artifacts such as model weights can additionally be shared with the algorithm provider, enabling collaborative model improvement while preserving data sovereignty.

Implementation requisites



Figure 10: Evaluation of implementation requisites – Compute-to-Data in industrial manufacturing | Subjective assessment of the relative complexity of implementation requisites by early adopters of PETs in the presented use case, measured on a six-point Likert scale ranging from "very easy to implement" to "very difficult to implement."

According to the early adopters, the implementation of the compute-to-data application poses challenges, particularly regarding the skills and understanding required for the technical realization of the solution (see Figure 10). For traditional companies in the mechanical engineering sector, it is necessary to build up expertise in the technologies used, such as DLTs and trusted environments. Especially if there is no trust in third parties, the trusted environments must be provided by the company itself, which requires technical effort. On the other hand, the use of available open-source software simplifies the development of the technical solution. In general, the solution complements existing processes in the area of data and analytics. Therefore, only minor hurdles in terms of operation and integration exist. However, after the analysis process, further steps are needed to integrate the resulting data into the existing data infrastructure. From an organizational perspective, the roles of the parties are clearly defined by the Pontus-X and Gaia-X frameworks. Also, the incentive mechanisms for individual parties, which also include the support of the decentralized ecosystem, are clarified. Likewise, the solution provides security about legal compliance through standardized terms and conditions and the enforcement of governance through smart contracts.

Implications

The following implications arise from the use of the compute-to-data approach to analyze machine data in the field of industrial production:

Enablers	Neutral	Constraints
<p>Business models</p> <p>The decentralized autonomous organization approach gives rise to novel business models for organizations. Examples include partial ownership of the data infrastructure, provision of value-added solutions, or the establishment of proprietary data spaces within the ecosystem.</p>	<p>Resource consumption</p> <p>Despite the use of DLT, the chosen consensus mechanism (proof-of-stake) results in only minor additional computing costs and environmental impacts compared to centralized solutions.</p>	<p>User experience</p> <p>The initial setup of a compute-to-data application requires additional effort, such as creating the necessary credentials or describing the data and service offerings. Consequently, more time elapses before an application can be executed for the first time.</p>
<p>Culture and collaboration</p> <p>A trustworthy ecosystem is created that makes partners, data, and service offerings visible and promotes collaboration between data owners and service providers.</p>		<p>Usability</p> <p>Depending on the application's computational requirements, users may face additional infrastructure costs to perform the necessary operations.</p>
<p>Future viability</p> <p>The decentralized autonomous organization and the open-source software foundation enable long-term availability of solutions without restrictions.</p>		

Table 7: Implications – Compute-to-Data in Industrial Manufacturing

4.3.4 Zero-knowledge proofs in the energy industry

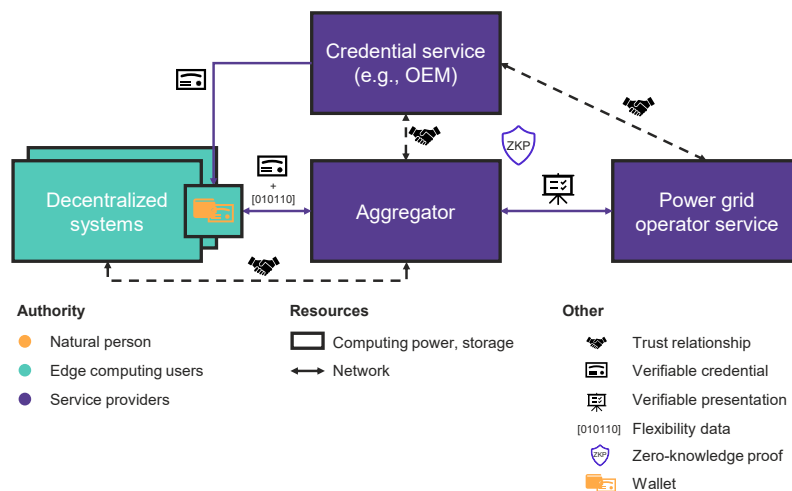


Figure 11: Schematic representation – Zero-knowledge proofs in the energy industry

Zero-knowledge proofs are mathematical procedures that enable the provisioning of data-based proof of a specific fact without disclosing the underlying data itself. For example, age verification can be performed without revealing the actual age. In the energy industry, zero-knowledge proofs can be combined with self-sovereign identity mechanisms to aggregate the available capacities

of decentralized energy generation plants and report them to a grid operator. This avoids exposing potentially sensitive personal or corporate data, such as the location or utilization of a specific plant, while generating certainty about the capabilities of such systems. In the application under consideration (see Figure 11), an aggregator uses this principle to bundle redispatch⁹ capacities from decentralized plants and make these capacities available to a grid operator.

In this scenario, the grid operator wants to verify the data aggregated by the aggregator without accessing it directly. For example, the grid operator must ensure that the capacities offered by the decentralized power plants are actually available and that the plants are located in the relevant grid area. Additionally, the grid operator should be able to perform this verification as efficiently as possible.

The concept is based on self-sovereign identities and public/private key infrastructures. The starting point is an existing chain of trust between network operators, plant OEMs, aggregators, and plant operators. During the setup phase, the decentralized plants are equipped with a trustworthy hardware module. Initially, a verifiable digital proof (verifiable credential)¹⁰ relating to the master data of the plant is transferred to the wallet of this hardware module from the credential service of a trustworthy party (e.g., OEM or network operator). The plant owner can then control this proof. One piece of information contained in this verifiable credential is, for example, proof of eligibility to participate in redispatch. The owner delegates the verifiable credential to the aggregator and authorizes it to create aggregated flexibility offers. The plants then automatically make their individual flexibility offers available to the aggregator. The aggregator can then consolidate the individual flexibility potentials and make them available to the grid operator as a verifiable presentation. The signatures of the plants prove that the offers originate from certified plants. At this point, zero-knowledge proofs are also used to create further necessary evidence of aggregated offers without disclosing the data of individual plants. For example, it can be proven that the plants are located in a specific grid area or that the capacities are actually available. In the future, the aim is to supplement the approach with hardware security modules in decentralized plants to rule out further manipulation.

⁹ In this context, redispatch means that the grid operator, in return for payment, shuts down decentralised generation plants or utilises battery storage capacities to stabilise the power grid.

¹⁰ <https://www.w3.org/TR/vc-overview/>

Implementation requisites



Figure 12: Evaluation of implementation requisites – zero-knowledge proofs in the energy industry | Subjective assessment of the relative complexity of implementation requisites by early adopters of PETs in the presented use case, measured on a six-point Likert scale ranging from "very easy to implement" to "very difficult to implement."

According to the early adopters, building expertise and understanding is a key challenge in implementing zero-knowledge proof applications (see Figure 12). Zero-knowledge proofs are still a young technology with a relatively small developer community. Building expertise on the part of application developers therefore requires considerable effort. Furthermore, acceptance among users is not immediately guaranteed, as the technology is difficult to understand due to its complex mathematics and is often perceived as opaque without additional explanations.

Comprehensive work is also necessary to create the organizational prerequisites. The use of zero-knowledge proofs is particularly useful when data protection needs to be guaranteed across different parties. However, realizing cooperation and exchange across parties leads to considerable additional organizational effort. For example, manufacturers, network operators, and aggregators must cooperate to define common standards for certificates and trustworthy data and to establish the necessary organizational trust mechanisms. For the general implementation of the use case, the approach must also attract a critical mass of participants to generate the necessary network effects.

According to the early adopters, there are also major challenges at the legal level, although these result primarily from the application of zero-knowledge proofs in the strictly regulated electricity market. For example, certain cryptographic methods and programming languages for zero-knowledge proofs are not yet approved for use in critical infrastructures. In addition, zero-knowledge proofs need to be evaluated from a liability perspective to build trust among the necessary stakeholders.

Once the necessary knowledge has been established among those responsible, the actual technical implementation is relatively easy to achieve. From the perspective of the operation and integration of zero-knowledge proofs for the given application, the chosen greenfield approach does not present any hurdles. In the future, however, additional efforts may be required to harmonize the effort with already established standards and protocols in the energy market.

Implications

The use of zero-knowledge proofs for decentralized redispatch in the energy industry has the following implications:

Enablers	Neutral	Constraints
<p>Culture and cooperation: The use of zero-knowledge proofs in the energy sector promotes the development of an ecosystem in which producers, aggregators, and grid operators work together, resulting in benefits for all parties. Zero-knowledge proofs also promote the introduction of aggregators as a new form of intermediary that enables novel applications in the energy industry.</p>	<p>User experience: The creation of zero-knowledge proofs runs in the background and is not visible to flexibility providers. Full automation of the processes prevents user intervention.</p>	<p>Black box process: The process of creating zero-knowledge proofs is difficult for users to understand and can create a black-box impression that undermines confidence in the technology.</p>
<p>Economic potential: Economic potential can be leveraged by creating new roles and implementing previously impossible applications. For example, the application presented can help reduce redispatch costs, which in the long term reduces system costs in the electricity market and leads to potentially lower grid fees.</p>		<p>Resource consumption: The calculations involved in zero-knowledge proofs result in higher computing costs, which limits real-time capability and restricts operational use in electricity grids with requirements in the millisecond range.</p>
<p>Sustainability: Zero-knowledge proofs integrate seamlessly into concepts such as self-sovereign identities and data spaces, thereby promoting further innovation and combining governance aspects with the technical guarantee of data sovereignty.</p>		<p>Sustainability: Further regulatory work is needed to enable the widespread adoption of zero-knowledge proofs in critical infrastructures.</p>

Table 8: Implications – Zero-knowledge proofs in the energy industry

4.3.5 Trusted execution environments in the energy industry

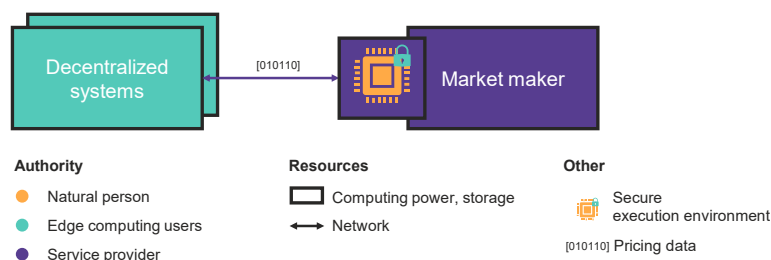


Figure 13: Schematic representation – Trusted execution environments in the energy industry

Trusted execution environments are hardware-supported, isolated execution environments that protect code and data from being accessed or manipulated by other elements of the system during data processing. This allows, for example, confidential data to be processed in the cloud or model weights to be protected when using artificial intelligence. In the energy sector, trusted execution environments can be applied to pricing processes in local energy markets. For pricing,

both the pricing algorithms and the electricity offers must be brought together at a central location. Trusted execution environments allow price calculations to take place in a protected environment to prevent individual parties from gaining an advantage and to protect potentially sensitive electricity offers from being viewed by other providers.

In addition to protecting the data necessary for price determination, market participants attach particular importance to fairness, transparency, and traceability in the execution of the price determination algorithm. It is also important that no central body acts as the sole supervisory authority. Furthermore, the pricing process should be as cost-effective as possible. Since calculation costs increase with the frequency of price calculations, cost efficiency enables shorter, more granular pricing intervals. Furthermore, the solution should be scalable to allow the participation of as many parties as possible.

In the examined case, the trusted execution environment is realized through a cloud service provider's offering, configured by the market maker (see Figure 13). The pricing algorithm is available to plant operators as open-source code. The trusted execution environment can use cryptographic methods to verify and transparently demonstrate that exactly this code is executed. The data required for pricing is fed directly from the plant into the trusted execution environment. End-to-end encryption ensures that the data is protected during transfer.

Implementation requisites



Figure 14: Evaluation of implementation requisites – Trusted Execution Environments in the energy industry | Subjective assessment of the relative complexity of implementation requisites by early adopters of PETs in the presented use case, measured on a six-point Likert scale ranging from "very easy to implement" to "very difficult to implement."

According to early adopters, the biggest challenge in implementing the described application lies in creating the necessary skills among application developers and technical understanding among plant operators (see Figure 14). Trusted execution environments are a rather "exotic" technology that requires specialized expertise. Companies must be prepared to allocate internal resources to build up the necessary skills. At the same time, acceptance among plant operators remains a challenge. While technically savvy users can understand the benefits, there is often a lack of understanding of the underlying logic, which, combined with the higher costs of the infrastructure, can lead to reduced technology acceptance.

The ability to use specialized cloud services for trusted execution environments abstracts much of the technical complexity. Nevertheless, setup and operation require more effort than conventional cloud services. In addition, trusted execution environments are just one part of an overall system, necessitating additional safeguards, for example, in data transmission. Vulnerabilities in individual system components can increase the effort required for continuous adjustments. Furthermore, new attack vectors on trusted execution environments have been identified repeatedly over the years; therefore, potential risks should be monitored continuously.

From an organizational, operational, and legal perspective, early adopters generally see no major obstacles to the use of trusted execution environments. They can be purchased like conventional software services and can generally be implemented without extensive partner involvement. For operational integration, trusted execution environments typically require only adaptation to operational processes. As infrastructure components, they can be integrated in the same way as other software modules. In contrast to other cryptographic approaches, there are comparatively few legal hurdles. Regulatory requirements can be met by using European cloud providers that guarantee compliance with the necessary technical and organizational measures. In certain contexts, the use of trusted execution environments can also be viewed positively from a legal perspective, as it provides additional integrity and security.

Implications

The use of trusted execution environments in the energy industry has the following implications:

Enablers	Neutral	Constraints
<p>Integrity: In addition to data protection aspects, the implementation of trusted execution environments also increases the integrity and trustworthiness of data processing operations.</p>	<p>Future viability: Although comprehensive service offerings already exist, trusted execution environments remain more of a niche technology. However, they may become more widespread in the future due to the increasing demands for transparency and traceability in digital processes.</p>	<p>System integration: Trusted execution environments are always part of a larger system that must also meet data protection requirements. The use of trusted execution environments alone cannot guarantee comprehensive protection against data leakage.</p>
<p>Usability: Trusted execution environments support the reliability and stability of processes. End users benefit from a system that works reliably in the background. At the same time, the technology should be largely invisible to users.</p>		<p>Culture and collaboration: The "black box" nature of the technology can increase skepticism among less tech-savvy stakeholders. This can become a problem, especially when interacting with end users.</p>
<p>Trustworthiness Trusted execution environments enable the implementation of applications in which participants do not need to trust each other (trustless applications).</p>		<p>Cost-effectiveness: Initial and ongoing costs for hardware, integration, and operation are higher compared to conventional cloud solutions and could challenge the cost-effectiveness of a solution.</p>

Table 9: Implications – Trusted execution environments in the energy industry

5 Recommendations for Action and Outlook

In addition to the specific requirements for the use of privacy-enhancing technologies (PETs) in dedicated edge-cloud applications, the findings from the examined practical implementations can be used to generalize further recommendations. The recommendations presented in this section are based on statements from early adopters of PET solutions. Qualitative, inductive content analysis was used to analyze transcripts and thought protocols from interviews with early adopters to identify common prerequisites for the success of PETs and derive actionable recommendations.

Specifically, four fields of action (see Figure 15) have been identified for the effective use of PETs to ensure data protection and information security in edge-cloud systems. These fields are interdependent and should therefore be addressed together. Each field and its associated design objects are described below.

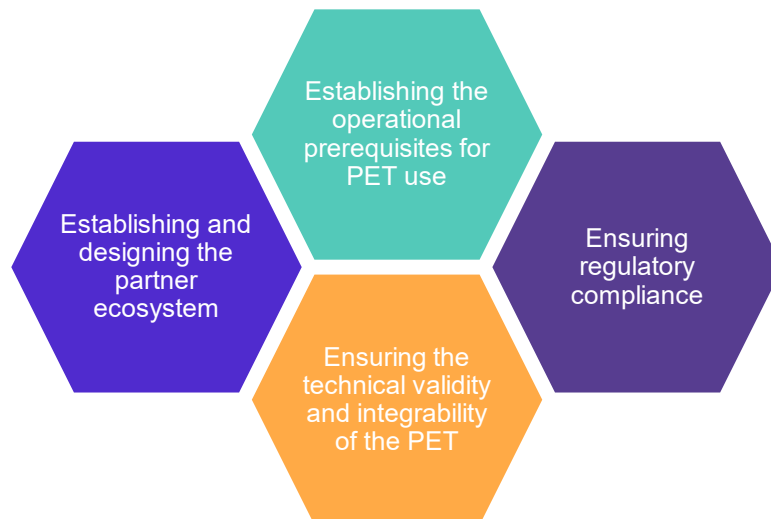


Figure 15: Fields of action for the effective use of PETs in edge-cloud Applications

Establishing and designing the partner ecosystem



Due to the high technological complexity of PETs such as zero-knowledge proofs, trusted execution environments, and federated learning, medium-sized companies in particular are often unable to implement and scale PET-based architectures on their own. At the same time, there is currently no technology provider that can fully cover all PETs in their service offering. Conversely, the use cases enabled by PET and their underlying data-sharing relationships inevitably result in constellations in which individual partners can both cooperate and compete with each other. This is the case, for example, when individual machine operators participate in a federated learning-based quality management system but at the same time produce competing products or compete for available orders. Accordingly, it is necessary to consider the structure and sustainable operation of the partner ecosystem. This is the task of the implementing organization, which thus automatically becomes an orchestrator of this ecosystem, regardless of intent.

Identify and integrate trustworthy partners:

In the edge-cloud applications to be realized, potential users—who are simultaneously data providers—either originate from existing business relationships or are targeted as new customers. The implementing organizations must select potential technology partners, implement the business application together with them, and integrate the PET solutions with already established software systems. Since these specialized software providers often lack the necessary domain knowledge, it is essential to integrate them closely into the business case at hand to strengthen cooperation and understanding. The close involvement of PET partners often results in interactions that are visible to other parties in the ecosystem.

Accordingly, it must be ensured that the PET partner is perceived as a trustworthy organization by the other parties involved. In the context of increasing geopolitical tensions, digital sovereignty should also be considered. For example, trusted execution environments are typically provided by cloud service providers and require data to be transferred to this environment. In highly regulated areas or those with heightened concerns regarding information security, selecting a European cloud service provider could mitigate concerns about digital sovereignty. Furthermore, PET services should be provided by neutral intermediaries with no economic stake in the business models of the partners involved, as exploitable vulnerabilities may persist despite the use of PETs.

Design necessary incentive mechanisms and business models:

In PET-based applications, it is necessary to create sufficient commercial incentives for all partners involved. The use of PET solutions often involves additional organizational and commercial effort, which affects all parties in the application ecosystem and must be managed accordingly. The added value generated by the implemented application must be distributed fairly amongst the participating parties. This applies to both the distribution of value between the various actor groups and within individual actor groups. For example, in applications based on federated learning, individual parties may contribute disproportionately to joint model training, while others may benefit disproportionately from these trained models. Furthermore, new incentive and reward systems can be explicitly created in PET partner ecosystems. This is demonstrated, for example, by the establishment of the DLT ecosystem in the compute-to-data approach (see Section 4.3.3), which enables transaction-based remuneration for infrastructure providers and thus makes it possible to scale the number of partners involved.

Establishing the operational prerequisites for PET use

At the operational level, it is also important to generate sufficient incentives for users of PET applications. For users, many PETs have a "black box" nature. They are incomprehensible due to their high technical or mathematical complexity. Similarly, PETs can sometimes involve significant additional procedural effort, which reduces user acceptance.

Generate knowledge, acceptance, and trust in the solution:

Accordingly, it is important to communicate the solution and provide appropriate training to decision-makers, users, and other stakeholders to build understanding, promote trust, and ultimately increase technology acceptance. On the one hand, existing training materials from partners can often be used. On the other hand, it is also conceivable to provide test environments in which users can access the views of different roles and thus test the technology "hands-on" and gain clarity about the design of processes and data access permissions.

Abstract PETs:

Although building trust in the technologies is a key success factor, the use of PETs should not significantly change the user experience. Instead, like other security technologies, PETs should work behind the scenes, so that the technical complexity recedes into the background and the benefits of the application come to the fore. Simple, targeted information for users may, however, be useful. For example, it is conceivable to visualise the system's current status during operation using simple displays. One example is a "traffic light view" that indicates correct functionality and thus builds confidence in continuous availability.

Rethink process and system design:

Furthermore, organizations should follow an approach of ensuring data protection and information security "by design." Compliance with data protection and information security requirements should be fully considered in the design phase of systems, rather than attempting to retrofit an existing system to comply with data protection requirements using PET tools. Similarly, processes should also be designed to ensure the most efficient use of PET tools. On the one hand, this requires the digitization of analog processes, such as manual data preprocessing or organizational processes that precede integration with PETs. On the other hand, digital processes should also be designed in such a way that the use of PET tools generates as little additional effort as possible, for example through data minimization.

Ensuring the technical validity and integrability of the PET

From a technical perspective, the use of PET tools in edge-cloud applications must be designed in such a way that both technical validity (proof of correct and secure functioning) and seamless integration into heterogeneous edge-cloud architectures are guaranteed. On the one hand, it is important to consider the challenge that individual vulnerabilities in the overall architecture can compromise the protective functions of the entire application. On the other hand, the integration of PET tools into existing system architectures is costly and complex due to the high degree of technical heterogeneity. Furthermore, the use of PETs often results in significant additional technical effort, which leads to higher latency in the edge cloud application, especially when cryptographic methods are used.

Establish the end-to-end trust chain:

For PETs to function reliably across organizational and system boundaries, a consistent end-to-end chain of trust with clear trust anchors must be established. Trust refers to the partners, the technical infrastructure and the data in a given application scenario. An important aspect is the technical implementation of zero-trust concepts and the pursuit of a "privacy by design" approach. One possible measure is the use of trusted issuers for identities and characteristics of the organizations, individuals, systems, and physical objects involved in value creation. Furthermore, the integrity of actors, code, and data should be made verifiable through cryptographic procedures without disclosing unnecessary information. Mechanisms such as remote attestation or technologies originating from the field of decentralized identities can support such processes.

In addition, onboarding processes for devices, software, and data should be clearly defined. In such processes, devices receive verified identities and certificates, software applications and configuration states are signed, and necessary roles are assigned via verifiable credentials. Data, models, and model weights should also be provided with hashes and signatures and made traceable via audit trails so that only artifacts with a valid attestation are leveraged.

Maintain an edge-first approach:

Despite incorporating PET solutions, organizations should perform as many operations as possible at the edge. This approach offers three key benefits: it reduces computing overhead, minimizes latency, and shrinks potential attack surfaces. For example, data can be pre-processed at the edge before being transmitted to a central location. This reduces network load and lowers the computational burden associated with cryptographic operations. All data transfers should be secured through encryption, authorization, and authentication procedures. By following this edge-first principle, the use of PET solutions can be limited to essential functions, keeping the overall system lightweight.

Use open interfaces, standards, and open-source software:

When expanding edge-cloud architectures with PETs, organizations should prioritize open interfaces, standards, and open-source software. This approach ensures that PET solutions can integrate smoothly into existing system landscapes and processes while maintaining interoperability, expandability, and flexibility. Open standards and interfaces make it easier to add new functions or components, allowing PET implementations to evolve alongside future technologies. Proprietary solutions, by contrast, often create vendor lock-in and complicate future adjustments or migrations. This risk is particularly acute in the PET market, where a limited number of providers exist—potentially giving them significant leverage over their users. Open-source software allows public review of the source code and cross-organizational collaboration. Especially in the case of PETs, which are designed to protect sensitive data and information, transparency is crucial for trust in the security and functionality of the systems. If open-source software is used by many companies in different contexts, security vulnerabilities are more likely to be detected and fixed more quickly. Furthermore, existing open-source PET implementations can be further extended and adapted to specific requirements, which accelerates innovation processes and reduces parallel, isolated new developments. Examples of existing open-source frameworks that meet these conditions are Flower and PySyft in the field of federated learning (Riedel et al., 2024) .

Ensuring regulatory compliance

Beyond organizational feasibility and technical functionality, PET-based edge-cloud architectures must also meet regulatory requirements. Practitioners face several challenges in this regard. Some PETs, such as federated learning, operate non-deterministically and represent a "black box" for user companies. As a result, certain risks remain regarding the potential disclosure of personal data by malicious actors, and complete privacy cannot be guaranteed. Furthermore, there are areas in which the use of PETs is not yet approved. The energy industry in Germany, for example, is considered critical infrastructure and subject to special IT security measures. Lengthy approval procedures are expected, particularly for technologies such as zero-knowledge proofs, where each individual implementation may require separate review and approval. There are also additional challenges when data or model weights are to be transferred across national borders. The General Data Protection Regulation and the Data Act set out specific requirements for cross-border data transfers, which should be reviewed and complied with when using non-European service providers, for example. Models and model weights may also constitute intangible assets that must be capitalized under international reporting standards when transferred across borders, requiring corresponding monetary consideration. This applies even to internal company data transfers—for example, when a foreign subsidiary transfers a locally trained model

to an entity registered in Germany. In such cases, the arm's length principle applies: affiliated companies must charge for the model transfer as if sharing with independent third parties.

These challenges demonstrate that legal safeguards for PET-based edge-cloud applications remain necessary, even when sensitive data has been technically anonymized. Organizations should consider regulatory compliance early in the project to identify potential obstacles. Decisions and considerations should be documented, for example, through data protection impact assessments, which are mandatory when processing personal data. Potential challenges related to cross-border data transfers should also be anticipated early. In research projects or initial trials, experimentation clauses may provide additional flexibility.

References

- 115th Congress. (2018). *H.R.4943 - CLOUD Act*. <https://www.congress.gov/bill/115th-congress/house-bill/4943>
- Adam, M., Hammoudeh, M., Alrawashdeh, R., & Alsulaimy, B. (2024). A Survey on Security, Privacy, Trust, and Architectural Challenges in IoT Systems. *IEEE Access*, 12, 57128–57149. <https://doi.org/10.1109/ACCESS.2024.3382709>
- Ali, M. A., & Al-Sharafi, S. A. H. (2025). Intrusion detection in IoT networks using machine learning and deep learning approaches for MitM attack mitigation. *Discover Internet of Things*, 5(1). <https://doi.org/10.1007/s43926-025-00104-w>
- Centre for Data Ethics and Innovation. (2021). *Privacy Enhancing Technologies Adoption Guide*. <https://cdeiuk.github.io/pets-adoption-guide/>
- Chang, W., & Wu, J. (Eds.). (2021). *Springer eBook Collection: Vol. 83. Fog/edge Computing For Security, Privacy, and Applications* (1st ed. 2021). Springer International Publishing; Imprint Springer. <https://doi.org/10.1007/978-3-030-57328-7>
- DIN. (2020). *DIN EN ISO/IEC 29134*.
- ENISA (2016). PETs controls matrix: A systematic approach for assessing online and mobile privacy tools. <https://www.enisa.europa.eu/publications/pets-controls-matrix/pets-controls-matrix-a-systematic-approach-for-assessing-online-and-mobile-privacy-tools>
- ENISA (2022). DATA PROTECTION ENGINEERING: From Theory to Practice. <https://www.enisa.europa.eu/publications/data-protection-engineering>
- European Parliament, & Council of the European Union. (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance)*. <http://data.europa.eu/eli/reg/2016/679/oj>
- European Parliament, & Council of the European Union. (2023). *Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data and amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act) (Text with EEA relevance)*. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32023R2854>
- European Parliament, & Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance)*. <http://data.europa.eu/eli/reg/2024/1689/oj>
- European Commission. (2023). *Study on the economic potential of far edge computing in the future smart Internet of Things*. <https://op.europa.eu/en/publication-detail/-/publication/ff35c457-8f3b-11ee-8aa6-01aa75ed71a1>

- Future of Privacy Forum. (2024). *Repository for Privacy Enhancing Technologies (PETs) - Future of Privacy Forum*. <https://fpf.org/global/repository-for-privacy-enhancing-technologies-pets/>
- Gerber, N., Gerber, P., & Volkamer, M. (2018). Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & Security*, 77, 226–261. <https://doi.org/10.1016/j.cose.2018.04.002>
- Heurix, J., Zimmermann, P., Neubauer, T., & Fenz, S. (2015). A taxonomy for privacy enhancing technologies. *Computers & Security*, 53, 1–17. <https://doi.org/10.1016/j.cose.2015.05.002>
- ISACA. (2024). *Exploring Practical Considerations and Applications for Privacy Enhancing Technologies*. <https://www.isaca.org/resources/white-papers/2024/exploring-practical-considerations-and-applications-for-privacy-enhancing-technologies>
- ISO. (2022). *ISO/IEC 27557:2022*. ISO.
- ISO. (2024). *ISO/IEC 29100:2024*. ISO.
- Janakiraman, R., Lim, J. H., & Rishika, R. (2018). The Effect of a Data Breach Announcement on Customer Behavior: Evidence from a Multichannel Retailer. *Journal of Marketing*, 82(2), 85–105. <https://doi.org/10.1509/jm.16.0124>
- Klymenko, A., Meisenbacher, S., & Matthes, F. (2025). *Privacy-Enhancing Technologies: A Comprehensive Guide for Non-technical Practitioners*.
- Lavin, R., Liu, X., Mohanty, H., Norman, L., Zaarour, G., & Krishnamachari, B. (2024). *A Survey on the Applications of Zero-Knowledge Proofs*. <https://doi.org/10.48550/arXiv.2408.00243>
- Noble, A. (2023). *From privacy to partnership*. The Royal Society. <https://royalsociety.org/-/media/policy/projects/privacy-enhancing-technologies/From-Privacy-to-Partnership.pdf>
- OECD (2023). *Emerging privacy enhancing technologies: Maturity, opportunities and challenges*. https://www.oecd.org/en/publications/emerging-privacy-enhancing-technologies_bf121be4-en.html
- Riedel, P., Schick, L., Schwerin, R. von, Reichert, M., Schaudt, D., & Hafner, A. (2024). Comparative analysis of open-source federated learning frameworks - a literature-based survey and review. *International Journal of Machine Learning and Cybernetics*, 15(11), 5257–5278. <https://doi.org/10.1007/s13042-024-02234-z>
- Scherenberg, F. von, Hellmeier, M., & Otto, B. (2024). Data Sovereignty in Information Systems. *Electronic Markets*, 34(1). <https://doi.org/10.1007/s12525-024-00693-4>
- Schiffner, S. (2015a). *Privacy and Data Protection by Design*. <https://www.enisa.europa.eu/publications/privacy-and-data-protection-by-design>
- Schiffner, S. (2015b). *Readiness analysis for the adoption and evolution of privacy enhancing technologies: Methodology, pilot assessment, and continuity plan : Approved, version 1.0, public*. ENISA. <https://doi.org/10.2824/614444>
- Sheikh, A. M., Islam, M. R., Habaebi, M. H., Zabidi, S. A., Bin Najeeb, A. R., & Kabbani, A. (2025). A Survey on Edge Computing (EC) Security Challenges: Classification, Threats, and Mitigation Strategies. *Future Internet*, 17(4), 175. <https://doi.org/10.3390/fi17040175>

- Stock, J., Petersen, T., Behrendt, C.-A., Federrath, H., & Kreuzburg, T. (2022). Privatsphärefreundliches maschinelles Lernen. *Informatik Spektrum*, 45(3), 137–145. <https://doi.org/10.1007/s00287-022-01440-9>
- United Nations. (2023). *United Nations Guide on Privacy-Enhancing Technologies for Official Statistics*. https://unstats.un.org/bigdata/task-teams/privacy/guide/2023_UN%20PET%20Guide.pdf
- Xu, R., Baracaldo, N., & Joshi, J. (2021, August 10). *Privacy-Preserving Machine Learning: Methods, Challenges and Directions*. <http://arxiv.org/pdf/2108.04417v2>
<https://doi.org/10.48550/arXiv.2108.04417>

Appendix A – Detailed Overview of PETs

Anonymization and Pseudonymization

Anonymization involves removing attributes from data sets to prevent the identification of data subjects. Anonymization ensures that entities within a data set are indistinguishable. In pseudonymization, characteristics such as identifiers are replaced by pseudonyms. While effectively anonymized data is no longer considered personal data, pseudonymized data can often be linked to a natural person using additional, separate information and therefore continues to be considered personal data. Anonymization often results in a tension between the usefulness of the data and the extent of protection. Anonymization and pseudonymization techniques are typically applied to data "at rest" during the data preparation phase. Anonymization techniques include the removal or generalization of certain attributes of a data set. Anonymization techniques include:

- k-anonymity: Information about each entity concerned is indistinguishable from at least $k-1$ other entities within the data set, for example through clustering. This primarily protects the identity of the entity.
- l-diversity: Extends k-anonymity by forming clusters in such a way that the sensitive attribute differs l times within the class. This prevents the characteristics of the person in the cluster from being deduced from their membership in the cluster.
- t-closeness: Extends l-diversity by ensuring that the value distribution within a cluster is close to the value distribution of the population (t). This makes identification even more difficult.

The following procedures, among others, are conceivable for pseudonymization:

- Digital pseudonyms: Users choose their own pseudonyms or are assigned randomly generated character strings as pseudonyms.
- Cryptographic methods: Cryptographic methods such as hash functions are used to generate pseudonyms.
- Trusted third parties: A third party is entrusted with keeping the key secret that links digital pseudonyms to the true identities of their users.

Synthetic data

Synthetic data is (machine-generated) data that mimics the statistical properties of the population without violating the privacy of data subjects. Historically, statistical models such as Monte Carlo simulations have been used to generate synthetic data. Nowadays, machine learning methods are increasingly employed. Machine learning methods learn the underlying distribution and generate new data with similar properties. Synthetic data can be completely synthetic (all variables are generated by a model), partially synthetic (only some variables are synthesized), or hybrid (generated from the real set and a completely synthetic set).

Federated learning

Federated learning is a collaborative approach to machine learning. Typically, a global AI model is initially provided by a central server and distributed to various nodes/clients (e.g., servers, edge

devices, smartphones). These train the model locally with their own data and only transfer updated model parameters to the central server. The server aggregates the individual gradients, for example by calculating weighted averages, and updates the global model. This cycle can be continued indefinitely, but ideally reaches convergence. Federated learning exists in variants such as cross-device learning with many end devices or cross-silo learning with a few powerful nodes that already aggregate data in their own environment (e.g., clinics, banks). In addition, there is a difference between "horizontal federated learning," in which the clients' data has the same attributes but different entities, and "vertical federated learning," in which the clients have data on the same entities but with different attributes. The best-known examples of federated learning are next word prediction, auto-correction, and emoji suggestions, which are trained and executed on millions of smartphones.

Secure multi-party computation

Secure multi-party computation is a cryptographic approach that enables multiple parties who distrust each other to perform calculations together without revealing their individual input data. All parties receive only the result of the calculations. The individual inputs of the participants remain confidential throughout the entire process. Specifically, protocols such as secret sharing are used to implement secure multi-party computation. The input of those involved in the joint calculation is divided into parts and distributed among the other parties so that no single party has access to all the information. The parties then perform their individual calculations and combine the calculation results at the end. The process therefore requires a minimum number of participating parties. The use of secure multi-party computation is currently still a field of research. Current challenges include low performance and the restriction to dedicated computing methods.

Homomorphic encryption

Homomorphic encryption is a cryptographic method that allows calculations to be performed on encrypted data without having to decrypt it first. While conventional encryption methods secure the storage and transmission of sensitive information, they require the data to be decrypted for further processing, analysis, or machine learning applications. This poses a risk to data protection and confidentiality, as unauthorized persons can access the raw data during processing. Full homomorphic encryption allows any mathematical operation to be performed directly on the ciphertext. Once the calculations are complete, the encrypted result can be decrypted to obtain the real result. This corresponds to the same output that would be generated by a calculation with the decrypted data. Homomorphic encryption can be used for a wide range of analyses, especially those that can be expressed as polynomial functions, as well as for selected AI applications. However, there are also challenges, such as increased computational effort and limited calculation accuracy (Stock et al., 2022).

Zero-knowledge proofs

Zero-knowledge proofs are a cryptographic concept that enables one party (prover) to prove the validity of a statement to another party (verifier) without revealing information about the underlying data itself. Zero-knowledge proofs primarily protect data during processing and transmission by minimizing the disclosure of information. There are various mathematical methods for creating zero-knowledge proofs, which can be divided into two basic categories. Interactive zero-

knowledge proofs require multiple communications between the prover and the verifier. The verifier makes several rounds of inquiries to the prover, who must answer them to ensure the credibility of the proof. In contrast, non-interactive zero-knowledge proofs do not require repeated communication between the parties. Instead, the prover creates a one-time proof that can be verified directly and independently by the verifier (Lavin et al., 2024). Established areas of application for zero-knowledge proofs include identity management and finance.

Proxies and onion routing

A proxy server is an intermediary in a computer network that receives requests from clients and forwards them to the target system. The proxy acts as an intermediary that controls communication and can conceal the identity of the requesting client. The proxy receives the request, independently establishes the connection to the destination, and forwards the response back to the client. The proxy can thus be used, for example, to hide the IP address and thus obscure the client's identity. There are different types of proxies. Anonymous proxies and highly anonymous proxies are particularly relevant for protecting privacy. Anonymous proxies hide the user's IP address, while highly anonymous proxies also conceal the fact that they themselves are acting as proxies.

Onion routing is a network technology that routes connection-oriented, bidirectional communication through a chain of relays (onion routers) in order to decouple the sender and destination from each other and thus make many forms of network traffic analysis more difficult. Essentially, onion routing replaces direct connections with a series of entry nodes (guard/entry), intermediate nodes, and exit nodes (exit), known as onion routers. Messages are encrypted in layers by the client, similar to the layers of an onion. Each node removes only its own layer, knows only its immediate predecessor and successor, and has no access to the complete content or the entire route. Therefore, each router only knows from whom it received the data and where it is sending it to. This protects both the content of the communication and the communication partner from surveillance. Unlike conventional encryption methods, onion routing ensures that compromising a single router does not expose the entire communication path. Existing Internet services and applications do not need to be adapted to use onion routing.

Hardware security modules

Hardware security modules are specialized hardware devices that can securely generate, store, and manage cryptographic keys. These keys are used to encrypt and decrypt data and to create digital signatures and certificates. Hardware security modules can therefore make a significant contribution to protecting intellectual property and other sensitive data from unauthorized access. Specifically, hardware security modules generate keys directly within their cryptographic processor. Applications can access the cryptographic functions via dedicated APIs and use them to encrypt or decrypt data within the secure environment. The cryptographic functions and algorithms are known only to the manufacturer. Hardware security modules are designed to be tamper-proof. Therefore, physical access to the devices can either be prevented or lead to the removal of sensitive data. Hardware security modules are primarily used in areas with high compliance requirements, such as finance and telecommunications. In addition, they are often combined with other PETs to reinforce existing security concepts.

Trusted execution environments

Trusted execution environments are specialized hardware architectures that enable applications to run in isolation from other parts of the system. Their fundamental goal is to create a trustworthy environment in which data and processes are protected from unauthorized access, manipulation, or monitoring—even if the host system or underlying operating system has been compromised. Trusted execution environments enable what is known as "confidential computing".

Trusted execution environments require specialized CPU solutions, currently available from only a few manufacturers. They use dedicated on-chip memory and virtual memory management functions to prevent external access to their memory area. Hardware encryption secures data that must be moved between the CPU and system memory. Remote attestation allows users to verify the integrity and authenticity of code and data executed within a trusted execution environment. The hardware forms the trust anchor, while software components control security-related processes such as encryption, authentication, and attestation. In addition to ensuring data protection by preventing unauthorized parties—including the operators of the environment—from viewing the data, TEEs also guarantee the integrity of both data and code.