# Multi-Party Computation in Corporate Data Processing: Legal and Technical Insights*

Sebastian Becker[1], Christoph Bösch[1], Benjamin Hettwer[1], Thomas Hoeren[2], Merlin Rombach[2], Sven Trieflinger[1], and Hossein Yalame[1]

[1] Bosch Research, Renningen, Germany
{first.last}@de.bosch.com
[2] Institute for Information, Telecommunications and Media Law, Münster, Germany
{hoeren,merlin.rombach}@uni-muenster.de

**Abstract.** This paper examines the deployment of Multi-Party Computation (MPC) in corporate data processing environments, focusing on its legal and technical implications under the European Union's General Data Protection Regulation (GDPR). By combining expertise in cryptography and legal analysis, we address critical questions necessary for assessing the suitability of MPC for real-world applications.
Our legal evaluation explores the conditions under which MPC qualifies as an anonymizing approach under GDPR, emphasizing the architectural requirements, such as the distribution of control among compute parties, to minimize re-identification risks effectively. The assertions put forth in the legal opinion are validated by two distinct assessments conducted independently.
We systematically answer key regulatory questions, demonstrating that a structured legal assessment is indispensable for organizations aiming to adopt MPC while ensuring compliance with privacy laws. In addition, we complement this analysis with a practical implementation of privacy-preserving analytics using Carbyne Stack, a cloud-native open-source platform for scalable MPC applications, which integrates the MP-SPDZ framework as its backend. We benchmark SQL queries under various security models to evaluate scalability and efficiency.

**Keywords:** Multi-Party Computation · GDPR Compliance · Data Anonymization · Privacy-Preserving Analytics.

## 1 Introduction

The increasing complexity of privacy regulations, such as the European Union's General Data Protection Regulation (GDPR), presents significant challenges for organizations seeking to process sensitive data across borders. While GDPR emphasizes both data protection and free data flow, achieving this balance often

* This paper presents the results of a publicly funded research project in which Bosch contributed the technical use case and Prof. Hoeren from ITM, Münster contributed the legal assessment. Note that the presented technical solution is not implemented at Bosch and that the legal assessment is an independent academic study and shall not be construed as an official position or endorsement by Robert Bosch GmbH.

requires innovative technical solutions. Among these, Secure Multi-Party Computation (MPC) [19, 36] emerges as a transformative Privacy-Enhancing Technology (PET), enabling collaborative data processing without compromising individual privacy [9,21,22]. By leveraging cryptographic techniques, MPC ensures that sensitive data remains secret-shared and inaccessible during computations, aligning with GDPR's principles of privacy by design and by default.

MPC has gained traction for its potential to enable privacy-preserving computations in a variety of fields, such as financial fraud detection, genomics research, and privacy-aware advertising [10]. Moreover, the MPC Alliance—a collaborative effort among industry leaders—has further advanced MPC adoption, emphasizing its utility in addressing real-world privacy challenges [28]. Despite these successes, the broader adoption of MPC remains hindered by legal ambiguities, particularly regarding its compliance with data privacy regulations like GDPR. To address these challenges, our assessment leverages a real-world use case utilizing Carbyne Stack [4], an open-source cloud stack designed for scalable MPC applications, with MP-SPDZ serving as the MPC backend. Carbyne Stack's cloud-native architecture enables secure and efficient implementation of MPC protocols in real-world settings.

**The Technological Promise and Legal Challenge of MPC.** MPC provides a framework for secure computations in multi-entity scenarios, such as cross-border human resource (HR) analytics or collaborative benchmarking across corporate subsidiaries [3, 5, 18, 27, 29, 30]. By enabling computations on encrypted data without revealing the underlying plaintext, MPC safeguards individual privacy. This capability is particularly valuable in corporate environments constrained by stringent data-sharing regulations like the GDPR, which emphasizes principles such as data minimization, anonymization, and lawful processing.

However, despite its robust privacy guarantees, deploying MPC within existing legal frameworks poses significant challenges. A fundamental question is whether data processed via MPC qualifies as *anonymized* under GDPR or remains classified as personal data. This distinction is critical because it determines the regulatory obligations attached to MPC implementations. Furthermore, the lack of standardized legal guidelines for evaluating MPC creates uncertainty for organizations, especially in multi-jurisdictional scenarios. Bridging this gap requires a nuanced understanding of cryptographic principles alongside the evolving regulatory landscape to ensure compliance and build trust in MPC deployments. This interplay between technical innovation and legal feasibility underscores the importance of interdisciplinary efforts to advance the practical adoption of MPC. By addressing these challenges, MPC can unlock its potential to enable secure, privacy-preserving collaborations across regulated industries while adhering to complex legal requirements.

**The Aim of This Work.** This paper bridges the gap between MPC's technical promise and its legal feasibility by conducting an interdisciplinary analysis combining expertise from cryptographers and legal scholars. Focusing on HR analytics as a case study, we evaluate how MPC can facilitate privacy-preserving collaboration across subsidiaries without violating data protection laws. By systemat-

ically exploring the conditions under which MPC qualifies as an anonymization method under the GDPR, this work provides practical insights for deploying MPC in privacy-sensitive contexts and offers a pathway for its broader adoption in real-world corporate environments.

**Scope and Related Work.** PETs have been extensively studied to evaluate their compliance with data protection regulations. For instance, differential privacy has been analyzed for its ability to mitigate re-identification risks and produce outputs compliant with frameworks such as GDPR and FERPA [2, 31]. In the U.S. context, works like [35] focus on data privacy laws affecting private entities. Research into MPC as a PET has primarily centered on its alignment with GDPR. Studies like [23] discuss MPC's ability to meet Article 25 requirements for privacy by design and data minimization, while [32] examine its role in securely exchanging health data under GDPR, emphasizing anonymity. Similarly, [34] highlights MPC's advantages for data exchanges between legal entities but do not comprehensively address whether MPC achieves *anonymization* or its applicability in multi-jurisdictional scenarios. Additionally, works such as [16] and [33] discuss the broader tension between cryptographic techniques and privacy regulations but lack a detailed exploration of MPC's compatibility with GDPR.

Despite these contributions, significant gaps remain in understanding MPC's role in achieving GDPR compliance, particularly in assessing whether its processed data qualifies as anonymized and how it can be implemented in real-world scenarios. Our work addresses these gaps by combining expertise from cryptography and law to systematically evaluate MPC's compliance with GDPR. Specifically, we analyze its potential as an anonymization method and its deployment in privacy-sensitive applications like HR analytics across corporate subsidiaries.

**Citations in this Paper.** This paper employs two citation styles to align with its interdisciplinary nature. Legal citations are provided as footnotes to ensure accurate identification and retrieval, including non-English sources specific to individual jurisdictions. Technical references are cited using standard academic conventions.

## 2 Background and Context

**Secure Multi-Party Computation (MPC).** MPC enables multiple entities to collaboratively compute a function over their private data without revealing the inputs. The setup involves input parties that provide secret-shared data, compute parties that perform cryptographic computations without accessing raw data, and output parties that receive the computed result. MPC has been applied to domains such as financial fraud detection [14, 20], privacy-preserving health analytics [6, 12], and secure data aggregation [7, 17], offering robust privacy-preserving collaboration.

**Comparisons with other Privacy-Enhancing Technologies.** MPC is one of several PETs with distinct advantages and trade-offs. Confidential Computing (CC) uses Trusted Execution Environments (TEEs) to secure computations

within isolated hardware zones, though TEEs can be vulnerable to side-channel attacks [8]. Homomorphic Encryption (HE) enables computations directly on encrypted data but suffers from high computational overhead, limiting its practicality for large datasets [1]. Differential Privacy (DP) [13] protects individual privacy by adding controlled noise to outputs, making it less suitable for precise collaborative computations where MPC excels.

**Focus of this Work.** This paper focuses on MPC's unique capabilities for privacy-preserving computations and its ability to address legal and organizational challenges under GDPR.

**Legal Landscape and Regulatory Considerations.** MPC aims to ensure data security throughout the computational process, making it a pivotal tool for scenarios involving personal data. Its application raises critical questions of admissibility under data protection laws, particularly the GDPR. These considerations can be grouped into two categories: First, MPC enables previously restricted data processing by providing strong protective measures. Second, its deployment must comply with stringent data protection requirements. As one of the most comprehensive frameworks globally, the GDPR establishes high standards for safeguarding data subjects' rights while imposing extensive obligations on data controllers and processors.

- **Personal Data Processing and Anonymization:** For MPC applications, the primary legal consideration is the processing of personal data and its potential anonymization. According to Article 4(1) GDPR, personal data encompasses any information related to an identified or identifiable individual. The European courts adopt a risk-based approach to assess re-identification possibilities, requiring case-specific evaluations of individual risks, the nature of the data, and the specific processing operations involved.[1]
- **Participants and their Roles:** Under GDPR (Art. 4(7), 4(8), 26, 28), the distinction between joint controllership and data processing arrangements hinges on the decision-making authority over processing purposes and means. Processors act in a supporting role under the controller's instructions,[2] whereas controllers retain the authority to determine these purposes and means.[3] MPC implementations necessitate a case-specific evaluation of these roles to ensure compliance with GDPR criteria.
- **International Data Transfers:** Chapter V of the GDPR (Art. 44-50) imposes strict requirements for international data transfers, requiring an adequacy decision, appropriate safeguards, or specific derogations to ensure equivalent EU data protection standards.[4] By protecting data from unauthorized third-party access, MPC offers a promising solution to simplify third-country transfers while maintaining compliance with GDPR standards.
- **Technical and Organizational Measures:** Article 32 of the GDPR requires entities processing personal data to ensure its security by considering potential risks to individuals' rights and freedoms, taking into account the state of the art, implementation costs, and processing risks. In this regard, MPC technology holds significant promise as a robust technical measure to enhance data security and mitigate associated risks.[5]

## 3 Technical Requirements

Deploying Multi-Party Computation (MPC) requires a comprehensive understanding of its technical, security, and organizational requirements, particularly for privacy-sensitive applications such as human resource analytics.

**Infrastructure Requirements.** MPC systems are inherently distributed and involve multiple parties collaborating securely. The key components of an MPC system include:

- **Input Parties (Data Owners):** These are corporate entities or organizational units that provide data in a secret-shared form to ensure privacy throughout the computation process.
- **Compute Parties (CPs):** Entities, such as Virtual Cloud Providers (VCPs) that execute the cryptographic protocols required for secure computation, ensuring that no plaintext data is accessed at any stage.
- **Output Parties (Analysts):** Entities or individuals define computational queries and reconstruct results while adhering to privacy guarantees.

**Security Protocols.** The security of MPC systems is maintained through advanced cryptographic techniques designed to safeguard data privacy and integrity during computations. Key security protocols include:

- **Secret Sharing:** This technique involves splitting input data (each data item separately) into multiple secret shares distributed across compute parties. This ensures that no single party can reconstruct the original dataset, providing robust confidentiality.
- **Adversarial Models:** MPC deployments may operate under different security assumptions, the most significant of which are:
  - *Semi-Honest Model:* Assumes that parties follow the protocol but may attempt to infer information from received data.
  - *Malicious Model:* Provides stronger security by protecting against parties that may deviate from the protocol or collude with others.

**Integration with Privacy Regulations.** Adhering to privacy regulations like the General Data Protection Regulation (GDPR) is a fundamental requirement for MPC deployment. The ability of MPC to process encrypted data aligns with GDPR principles, particularly data minimization and purpose limitation. By enabling computations on secret-shared data, MPC facilitates cross-border data sharing while ensuring compliance with stringent legal frameworks.

**Case Study: Privacy-Preserving People Analytics.** *Collaborative benchmarking* represents a distinctive form of benchmarking, whereby independent legal entities within a company engage in joint comparison and assessment of their processes and metrics. Given the typically sensitive nature of the data involved in such analysis, compliance with relevant legal regulations is of paramount importance. This is especially the case when the data analysis is conducted across national borders and subject to various regulatory frameworks.

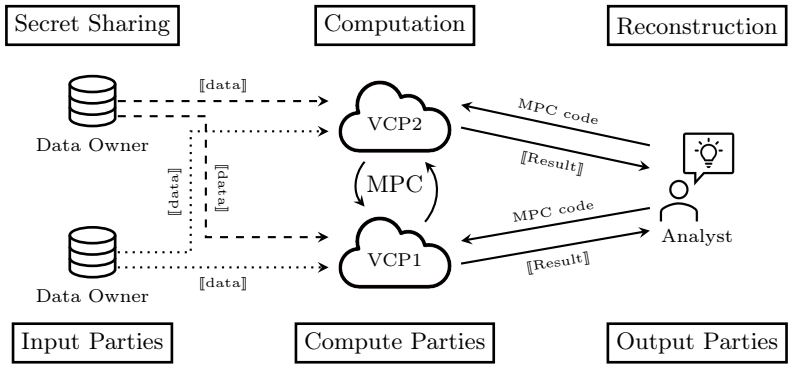


**Fig. 1:** System model for privacy-preserving collaborative data analysis. [[data]] and [[Result]] are secret shared data items.

Our Privacy Preserving People Analytics (PP-PA) use case focuses on pooling and analyzing human resource (HR) data spanning the entire life cycle of an employee across various legal entities in different countries within a company in a legally compliant way. Ideally, this should be possible with minimal risk to, e.g., understand and evaluate the organizational structure or to gain insights into the distribution of metrics like the span of control throughout the company. Therefore, we are analyzing a PP-PA system for secure data pooling and privacy-preserving multi-party analysis using MPC.

Our PP-PA system consists of several hundred legal entities within a company in multiple countries that plan to better understand their organizational structure to foster diversity and to derive comparisons on organizational metrics across the company. These are, for example, span of control, average age, and gender and headcount distribution. By pooling data from all of the participating entities, the scale of the database can be significantly increased. In this way, planning and forecasting can be made with higher confidence.

As seen in Figure 1, the system consists of data owners that act as input parties, several virtual cloud providers (VCPs) that perform the actual multi-party computation, and an analyst who provides the queries—in form of an MPC program—and acts as an output party. After agreeing on a common query/function, the data owners secret share the required data and distribute the shares to the VCPs that act as the computing parties for the MPC protocol and execute the computation on the secret-shared data without having access to the plain text data. The still secret-shared result is subsequently transmitted to the analyst, who then reconstructs the final result by combining the received secret shares.

## 4 Question and Answer Catalogue

This section provides a structured question-and-answer catalog addressing critical legal, organizational, and technical aspects of MPC deployments. These questions are categorized to guide organizations in assessing MPC's suitability for privacy-compliant data processing.

**Organizational and Deployment Considerations.** MPC systems consist of *input parties*, *compute parties* (CP), and *output parties*. The following questions relate to the legal implications of the composition and hosting of compute parties.

**Question 4.1: What are the implications of who controls the CPs? Consider the following scenarios (subsidiaries):**

1. **All same legal entity (LE)**
2. **Different LEs within the Bosch Group**
3. **Mix of LE(s) and external providers**
4. **All external providers**

The control architecture of CPs within MPC protocols has significant implications for data protection law assessment, as a thorough analysis of the relevant control scenarios reveals:

**All same legal entity (LE)**: In scenarios where a single LE maintains control over all CPs, the implementation of MPC protocols cannot achieve an anonymizing effect. This conclusion stems from the fundamental observation that the secret-sharing mechanisms becomes functionally ineffective, as all shares remain within the sphere of influence of a single entity. Despite the technical distribution across multiple software instances, this arrangement creates a critical single point of failure, as the controlling entity maintains comprehensive access to all information necessary for re-identification. In this context, the European Court of Justice's doctrine regarding non-consideration of unlawful actions is inapplicable due to the controlling entity's direct access to all relevant info.[6]

**Different LEs within one corporate group**: Conversely, when control is distributed among multiple legal entities within a corporate group, the potential for effective anonymization emerges. This assessment is predicated upon: i) the legal autonomy of individual group entities in matters of data protection,[7] ii) the automatic nullity of unlawful directives from the parent Company, and iii) the fundamental obligation of corporate management to ensure legal compliance.[8] These legal safeguards effectively neutralize the relevance of corporate affiliation in the context of unauthorized re-identification attempts.

**Mix of LEs within one corporate group and external providers:** The introduction of external providers or mixed control scenarios can further enhance data protection independence and provide anonymization. However, this enhancement is contingent upon maintaining genuine autonomy among multiple

independent entities and implementing robust legal and technical safeguards that prevent unilateral re-identification capabilities. The effectiveness of such arrangements depends crucially on the practical implementation of independence mechanisms and the maintenance of proper governance structures.

**Takeaway 4.1:** MPC's anonymization efficacy fundamentally depends on its control architecture. While single entity control negates anonymization by enabling comprehensive re-identification, distribution among legally autonomous entities or external providers establishes effective anonymization, contingent upon maintained independence and robust safeguards.

### Question 4.2: What are the implications of where the CPs are hosted, e.g., on-premise vs. public cloud infrastructure (in- or outside EU)?

The specific hosting location of the CPs is largely immaterial to the legal assessment. The determinative factor remains the nature of control over these CPs, as addressed in Question 4.1. If one entity maintains control, MPC cannot be considered to provide anonymization, regardless of the hosting arrangements. Conversely, when control is properly distributed, public cloud infrastructure may be utilized without legal impediment, provided there are adequate safeguards—particularly contractual provisions—against re-identification attempts.

This principle extends to cloud services operating outside the EU. Where MPC implementation, following the initial secret shares fragmentation, achieves anonymization (see Question 4.6), it falls outside the scope of third-country transfer regulations under GDPR, thus exempting it from Chapter Five compliance requirements. For scenarios where MPC's anonymizing effect is not established, please refer to Question 4.7 regarding extra-EU, esp. US-transfers.

**Takeaway 4.2:** The legal assessment of MPC focuses on control architecture rather than geographical hosting location. Where proper distributed control enables anonymization, even non-EU cloud services may be utilized without GDPR third-country transfer requirements; conversely, single-entity control negates anonymization regardless of hosting arrangements.

### Question 4.3: What are the implications on liability and service level agreements if CPs jointly offer a service to data owners?

The fundamental premise is that the processing is considered anonymous and thus outside the GDPR's scope until a potential data breach occurs. In the event of a data breach resulting in de-anonymization, the legal ramifications substantially depend on the specific circumstances of the breach, including its scope, duration, and other factors. Such de-anonymization automatically triggers the applicability of data protection law under Article 2(1) GDPR, necessitating compliance with all relevant data protection obligations where processing falls within Article 4(2) GDPR's scope.

The legal consequences differ significantly depending on whether the de-anonymization is characterized as active or passive.[9] Active de-anonymization, constituting an intentional process aimed at restoring personal references in anonymized data, typically results in more severe legal consequences. Such deliberate actions generally trigger comprehensive GDPR compliance obligations, including breach notification requirements under Article 33, data subject notifications per Article 34, potential administrative fines pursuant to Article 83(4),(5), and compensation obligations under Article 82 GDPR in conjunction with national law of the EU member states.[10]

Conversely, passive de-anonymization—occurring through exogenous factors such as technological advancement or the emergence of new sources of information—calls for a more nuanced assessment. Recital 26 GDPR acknowledges inherent de-anonymization risks even in properly implemented anonymization measures. Consequently, administrative fines may be avoided under Article 83(2)(b), (c), (d), (k) GDPR if: i) the initial anonymization adhered to GDPR standards, ii) documentation and due diligence obligations were fulfilled, and iii) immediate risk-mitigating measures were implemented upon discovery.[11]

Service-level agreements (SLAs) between the involved CPs must establish clear parameters for the distribution of responsibilities among participating CPs and define specific obligations regarding breach prevention and response mechanisms. The agreements should explicitly delineate liability allocation among the parties and specify the technical and organizational measures required to maintain effective anonymization. Furthermore, they must include detailed response protocols for potential de-anonymization scenarios to ensure prompt and appropriate action when necessary. A crucial element of these agreements must be the explicit prohibition of any cooperation between parties aimed at re-identification, reinforced through contractual penalties. This prohibition serves as an essential safeguard against intentional circumvention of anonymization measures and strengthens the overall data protection framework.

**Takeaway 4.3:** Data processing remains outside GDPR scope until de-anonymization occurs, whereupon legal consequences vary by breach circumstances. SLAs between CPs must delineate responsibilities, prevention measures, and response protocols, explicitly prohibiting intentional re-identification attempts.

**GDPR-Specific Considerations.** The following questions address GDPR implications for MPC systems processing EU citizens' personal identifiable information (PII.)

### Question 4.4: Does MPC imply joint controllership? What obligations arise for data controllers?

In accordance with Art. 26 GDPR, joint control regularly arises between the input parties and the parent company issuing the directive to use MPC. This arrangement stems from their mutual determination of purposes and means of

processing: the parent company provides technical infrastructure and usage directives, while subsidiary companies, as independent Legal Entities (LEs), execute the actual processing operations. Despite diverging operational purposes, a superordinate common goal of anonymized processing of employee data in the group and a common economic interest can be established, which can be seen as a converging purpose in any case. This conclusion stems in particular from the ECJ's extensive interpretation of the scope of joint controllership, according to which a small degree of cooperation and influence—even without access to the data—can establish joint controllership.[12] However, non-EU subsidiaries are exempt from the aforementioned provisions, and joint responsibility cannot be established based solely on the parent company's instructions.

If joint controllership is assumed in the above-mentioned cases, the parties involved must, in accordance with Article 26(1), (2) of the GDPR, define in a transparent obligation which entity fulfils which obligations of the GDPR. In doing so, it should be stipulated that only the subsidiaries assume these obligations in order to ensure that data is handled in the most data-minimizing way possible. Otherwise, the parent company would have to be given access to the personal data in order to fulfill this obligation. With regard to the other obligations, there are no special features compared to a conventional joint responsibility.

**Takeaway 4.4:** Joint controllership under GDPR Article 26 typically exists between EU subsidiaries and parent companies implementing MPC, requiring transparent obligation allocation, while non-EU subsidiaries remain outside this framework despite parental directives.

### Question 4.5: What is the impact of commissioned data processing regarding CPs? What contractual requirements exist between involved parties?

The individual CPs do not have access to personal data and are therefore not processors within the meaning of the GDPR, despite the joint processing of data with other CPs. Neither with respect to the parent company nor with respect to other CPs or the input parties. As an additional safeguard, the CPs should nevertheless conclude an agreement with other CPs and the other entities involved, which prohibits the cooperation for the purpose of carrying out a re-identification and imposes a contractual penalty.

**Takeaway 4.5:** Since the individual CPs have no access to personal data, they are not processors within the meaning of the GDPR.

### Question 4.6: How does end-to-end encryption or cryptographic methods affect data classification under GDPR?

According to Article 2(1) GDPR, the regulation applies exclusively to personal data processing, implicitly excluding anonymous data from its scope, as confirmed by Rec. 26. The determination of anonymity employs a dual analyti-

cal framework combining risk-based and relative approaches. The prevailing interpretation, endorsed by European courts, requires neither absolute nor irreversible anonymization. Instead, it demands a reduction of re-identification risk to a negligible level, considering "all means reasonably likely to be used" for identification.[13] This assessment examines practical factors such as cost, time, and technical feasibility of re-identification within specific processing contexts. Significantly, the Court of Justice of the European Union (CJEU) and recently the General Court (T-557/20) have established that anonymity must be evaluated from the perspective of the specific data processor.[14] Thus, data may qualify as anonymous for one entity while remaining personal data for another, depending on their reasonable means of re-identification. The mere theoretical possibility of third-party re-identification does not preclude anonymization if the processing entity cannot reasonably achieve identification through legally and practically available means.[15]

In accordance with these case law guidelines, the implementation of technical or cryptographic procedures can achieve an anonymizing effect, provided that robust and uninterrupted encryption is assured, irrespective of the specific technology employed. Following a relative, risk-based approach, the perspective of the respective data processing entity is decisive. Cryptographic-based procedures offer a distinct advantage over hardware-based PETs, like confidential computing, as they avoid the inherent risks associated with hardware-specific attack vectors and failures.

Considering the aforementioned approach to anonymizing personal data, which hinges on the residual risk of re-identification due to the unavailability of personal data, the deployment of MPC on the basis of the presumed prerequisites (correct implementation, contractual obligations, etc.) has the potential to result in the anonymization of the data subjected to processing in the examined use case. The personal reference is not eliminated when using MPC by removing individualizing features, as is the case with conventional anonymization methods. Rather, it is eliminated by encrypting and splitting the data into secret shares, which are distributed to the MPC instances. Through secret-sharing and supplementary measures, the re-identification risk can be reduced to such a negligible level that only the input party retains personal data, while other participating entities process effectively anonymized information.[16]

**Takeaway 4.6:** Data qualifies as anonymous when re-identification risk becomes negligible from the specific processor's perspective, regardless of theoretical third-party capabilities. Properly implemented MPC, through secret-sharing and supplementary measures, can achieve anonymization by reducing re-identification risk to negligible levels for all parties except the input entity.

### Question 4.7: What is the impact of using MPC for data transfers to non-EU countries?

GDPR Chapter Five establishes a complex framework for international data transfers to ensure the maintenance of Union-level data protection standards

while enabling proportionate cross-border data flows. The regulatory scheme necessitates a dual examination: First, meet the general GDPR processing requirements. Then, satisfaction of the specific transfer mechanisms in Articles 45-50.[17] These mechanisms include European Commission adequacy decisions and appropriate safeguards like Standard Contractual Clauses or Binding Corporate Rules. The ECJs Schrems II decision has enhanced these requirements, mandating comprehensive Transfer Impact Assessments of recipient jurisdictions' legal frameworks and practices, even where Article 46 safeguards exist.[18]

According to the position advanced herein, proper implementation of a MPC environment can achieve data anonymization, thereby generally precluding the application of aforementioned GDPR requirements for data transfers. However, the increased risk associated with transfer to third countries must be included in the assessment of the re-identification risk, and anonymization may only be considered if the previously discussed conditions are met (in particular, the effective distribution of secret shares outside the control of a single entity). Should this legal interpretation be contested and the anonymizing effect of MPC implementation be rejected, MPC deployment would nevertheless maintain significant relevance as a supplementary technical protective measure in accordance with the Schrems II jurisprudence.

The distribution of secret shares across multiple CPs can effectively prevent the concentration of re-identification-enabling data within jurisdictions characterized by extensive intelligence service access rights. This architectural approach substantially ensures the integrity of processed data. Such distribution mechanisms, when properly implemented, serve as robust technical measures aligning with post-Schrems II requirements for international data transfers, even in scenarios where MPC's anonymizing effect is not recognized.

**Takeaway 4.7:** GDPR's international transfer requirements become irrelevant when MPC achieves proper anonymization through distributed secret-sharing. However, if anonymization is contested, MPC's architecture still serves as a valuable technical safeguard under Schrems II by preventing data concentration in jurisdictions with extensive surveillance powers.

### Question 4.8: What defines a state-of-the-art technology under GDPR, and do MPC solutions meet this? What are the consequences of not using SOTA technology in case of a data breach?

The GDPR places significant emphasis on the implementation of technical and organizational measures (TOMs) across multiple provisions. Article 32(1) GDPR establishes a fundamental framework requiring controllers and processors to implement appropriate TOMs while considering the state of the art, implementation costs, and processing circumstances, alongside the probability and severity of risks to natural persons' rights and freedoms.[19] The concept of "state of the art" appears throughout the GDPR (notably in Articles 24, 25, 28, 32, and 33), though the regulation notably lacks a legal definition of this term.[20] Established legal interpretation defines it not merely as novel technological developments but

rather as proven methodologies integrated into technological practice. This encompasses known, tested, and effective measures currently available in the market. This concept requires economically feasible implementation but does not mandate the adoption of highly advanced technological developments that are still emerging.[21]

The classification of MPC as a state-of-the-art measure depends on its specific application context—in simpler scenarios, it may qualify as an established technology already, while in more complex implementations, it may represent cutting-edge scientific research beyond current technical standards.[22]

Importantly, Article 32(1) GDPR requires only the consideration of state-of-the-art measures, not their mandatory implementation. The failure to employ such measures does not automatically trigger legal consequences; rather, it requires an assessment of all circumstances of the individual case. However, disregarding this obligation may independently result in fines and factor into penalty calculations in case of data breaches. Regarding advanced cryptographic protocols not yet considered "standardized," similar principles apply. The deployment of more progressive methods, while potentially reflecting current scientific research rather than established technical standards, remains permissible provided their security can be effectively demonstrated. The absence of standardization may necessitate thorough preliminary security validation, potentially through external assessment. In this context, open-source approaches, such as those employed by Carbyne Stack, significantly contribute to external validation and security assurance.

**Takeaway 4.8:** GDPR mandates appropriate TOMs considering state-of-the-art technology, defined as proven, market-available methodologies rather than emerging developments. While MPC's classification as state-of-the-art varies by application context, its implementation remains permissible when security can be demonstrated. Please refer to Appendix A1 for a legal comparison to other PETs.

### Question 4.9: Can secret sharing be considered an encryption mechanism where parties hold decryption keys via their secret shares?

Article 32(1)(a) GDPR explicitly mentions encryption as an exemplary security measure, though without providing a legal definition. However, GDPR adopts a functional rather than purely technical approach to encryption, as evidenced by Article 34(3)(a)'s focus on rendering personal data "unintelligible to any person who is not authorized to access it, such as encryption".[23]

Under this technology-neutral framework, Multi-Party Computation (MPC) qualifies as encryption within GDPR's meaning despite its technical differences from traditional encryption methods. While conventional encryption transforms plaintext into ciphertext, MPC distributes data fragments among parties. However, both approaches achieve the fundamental objective of protecting data from unauthorized access. This functional achievement of security objectives,

rather than adherence to specific technical implementations, aligns with GDPR's technology-neutral approach to data protection.

**Takeaway 4.9:** GDPR adopts a functional, technology-neutral approach to encryption, focusing on data unintelligibility rather than specific technical implementations. Consequently, MPC qualifies as encryption by achieving the objective of protecting data from unauthorized access.

### Question 4.10: How does the use of non-standard cryptographic methods affect the legal assessment of MPC?

The principles governing "state of the art" technology equally apply to the implementation of advanced cryptographic protocols that have not yet achieved standardization status. Advanced methods, while potentially reflecting current scientific research rather than established industry standards, may be deployed provided their security effectiveness can be conclusively demonstrated.

This approach aligns with GDPR's technology-neutral framework,[24] though the absence of standardization necessitates heightened scrutiny. Specifically, such implementations require thorough security validation, potentially through independent third-party assessment, prior to deployment. In this context, open-source approaches, such as those adopted by Carbyne Stack, significantly facilitate external validation by enabling comprehensive security review by the broader technical community.

**Takeaway 4.10:** Non-standard cryptographic methods may be deployed provided their security effectiveness can be conclusively demonstrated.

### Question 4.11: How can non-collusion between CPs be legally enforced? Is a contractual prohibition sufficient?

The integrity of MPC fundamentally relies on distributed calculation of data fragments (secret shares). This architecture, while robust, presents potential confidentiality risks in scenarios where multiple CPs engage in malevolent cooperation, depending on the chosen security model. The assessment of collusion risks among CPs must be evaluated according to established jurisprudential principles regarding the attribution of third-party knowledge that could enable re-identification. ECJs precedent establishes that third-party knowledge should be considered irrelevant where access would be legally impermissible or re-identification practically unfeasible (see above).

In the analyzed use-case, re-identification is already prohibited by law, theoretically eliminating the need for additional contractual measures under strict application of ECJs jurisprudence. However, prudent practice suggests implementing explicit contractual prohibitions against re-identification, potentially reinforced by penalty clauses, particularly when engaging external service providers.[25] Importantly, such agreements should not be structured as data processing agreements under Article 28 GDPR, as the processing does not involve personal

data, and the instruction rights inherent in such agreements would fundamentally conflict with the required autonomy of individual MPC instances.

**Takeaway 4.11:** While MPC's distributed architecture inherently protects against re-identification, explicit contractual prohibitions against CP collusion and re-identification are recommended.

### Question 4.12: If computation is performed securely, what is GDPR's stance if input data can be derived from computation results?

While MPC protocols effectively ensure input data confidentiality during processing, they do not inherently address potential privacy concerns in output data.[26] The evaluation of whether analysis results constitute personal data requires—once again—an assessment under the relative approach to personal data qualification, considering each receiving entity's perspective. Analysis results may reveal granular or statistically significant information about specific subpopulations or individuals, particularly when examining outliers or rare attribute combinations. In the context of personnel analytics within corporate groups, aggregated statistics, correlation analyses, or pattern recognition results could potentially enable re-identification of individual employees or inference of sensitive attributes.

The actual re-identification risk depends significantly on contextual factors such as organizational size, regional distribution, and information asymmetry between group entities. For instance, local entities with direct knowledge of their workforce face different risk profiles than parent companies lacking detailed employee information.

To mitigate these risks, implementation of Differential Privacy (DP) systems could complement MPC processing by introducing controlled noise into output data.[27] However, DP implementation requires careful calibration of privacy budgets to balance data utility with privacy protection. The ultimate determination of whether analysis results contain personal data remains context-dependent, with access control serving as a crucial parameter for maintaining output anonymity.

**Takeaway 4.12:** While MPC ensures input confidentiality, output privacy requires separate evaluation of potential re-identification risks.

### Question 4.13: In the context of employee PII, what consent is required for data processing, particularly when automation is a goal?

The initial generation of secret shares constitutes a processing operation subject to GDPR requirements. While this preliminary step requires legal justification, the subsequent MPC processing falls outside GDPR's scope due to its anonymizing effect.

Article 6(1)(f) GDPR emerges as the primary legal basis for the initial anonymization process, requiring a balancing of controller interests against data

subject rights.[28] This balance typically favors processing, as anonymization generally aligns with data subjects' interests in privacy protection. The legitimate interest assessment becomes particularly favorable when controllers implement sufficient risk controls in the anonymization process.[29] MPC's encryption through share generation significantly minimizes re-identification risks, strengthening the controller's position in the balancing test. Furthermore, data subjects' rights to informational self-determination remain protected rather than compromised by such processing. It should be noted, however, that the precise requirements may vary from one jurisdiction to another, given that the GDPR permits member states to deviate from the requirements for the processing of employee data under national law.

Notably, automated decision-making provisions under Article 22 GDPR become inapplicable to the subsequent MPC processing, as the data no longer qualifies as personal data. This example illustrates how MPC's anonymizing effect fundamentally transforms the legal framework applicable to data processing operations, effectively rendering traditional consent requirements irrelevant for the primary processing phase while focusing compliance obligations on the initial anonymization step.

**Takeaway 4.13:** The generation of secret shares in MPC requires GDPR compliance, esp. through Article 6(1)(f) GDPR. Once anonymization is achieved, subsequent processing falls outside GDPR scope, including Article 22's automated decision-making provisions, though initial anonymization requirements may vary by member state jurisdiction.

#### Independent Legal Verification

All arguments and assessments presented in this section have been reviewed and verified by two independent external legal experts specializing in data protection law. Both experts concur with our findings and confirm that the outlined legal interpretations align with established GDPR principles.

## 5 Technical and Experimental Analysis

In addition to our comprehensive legal assessment, this section presents the implementation and performance evaluation of our PP-PA system. The practical implementation complements the legal analysis by providing concrete performance metrics and feasibility insights. The system executes SQL queries using MP-SPDZ [24]. We evaluated the system under different security models to ensure robustness and compliance in real-world scenarios.

**Implementation Details.** Our implementation leverages MP-SPDZ to execute SQL queries that are pivotal to the PP-PA use case. This includes SELECT queries with aggregate functions (e.g., COUNT, SUM, AVG) and age distribution queries that involve sorting. The implementation aligns with the legal requirements, thereby ensuring secure and privacy-preserving computation across

| | | SELECT, incl. SUM, COUNT, AVG | | | | Age distribution incl. SORT | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Attribute | MASCOT [25] | SPDZ2K [11] | Semi [26] | Semi2k [19] | MASCOT [25] | SPDZ2K [11] | Semi [26] | Semi2k [19] |
| w/o edaBit [15] | Time [s] | 698.47 | 600.69 | 372.64 | 382.66 | 7650.28 | 8253.3 | 3062.16 | 2841.2 |
| | Rounds | 531,345 | 531,347 | 329,283 | 329,283 | 4,109,887 | 5,638,755 | 2,063,107 | 2,063,107 |
| | Data [MB] | 1727.4 | 1701.06 | 1699.08 | 418.19 | 111,538 | 111,553 | 45,602 | 22,774 |
| w/ edaBit [15] | Time [s] | 955.05 | 842.94 | 533.89 | 481.59 | 8262.29 | 8901.49 | 3544.87 | 3119.37 |
| | Rounds | 654,829 | 654,831 | 370,443 | 370,443 | 4,168,703 | 5,697,571 | 2,121,913 | 2,121,913 |
| | Data [MB] | 84.87 | 84.67 | 55.40 | 35.40 | 107,801 | 107,868 | 41,868 | 20,951 |

**Table 1:** Online phase benchmarking results for SELECT and Age Distribution queries with and without *edaBits* optimization under various security models. The test data set utilized in this study encompasses a total of 411,000 rows.

multiple entities. Our test data set utilized in this study encompasses a total of more than 411,000 database rows and approximately 50 columns.

**Experimental Setup.** To simulate a realistic environment, we deployed our system across two cloud providers using `x86_64` architecture with 8 vCPUs and 32 GB RAM. The instances used were Azure `Standard D8ads_v5` and AWS `t2.2xlarge`, both running `Ubuntu 22.04` as the operating system. The network setup included an average round-trip time (RTT) of 2.53 ms and a bandwidth of 100 Mbps. This distributed cloud configuration reflects practical deployment scenarios for privacy-preserving analytics.

**Performance Results (Online Only).** The performance results for SELECT and age distribution queries are summarized in Table 1. Key metrics include computation time, communication rounds, and data exchanged under various configurations, both with and without *edaBits* [15] optimization. Notable observations include the substantial impact of *edaBits*, which significantly reduced data exchange, particularly in malicious settings, but increased computation time due to a higher round complexity. Additionally, a clear security-efficiency trade-off was observed: the malicious model, while incurring higher overhead, offers enhanced security, making it crucial for sensitive applications.

## 6 Conclusion

This work provides a comprehensive analysis of the deployment of Multi-Party Computation (MPC) from both legal and practical perspectives. It addresses regulatory challenges under the GDPR alongside a practical implementation and evaluation of MPC's performance in real-world scenarios. The legal assessment systematically examines key questions, including the conditions under which MPC qualifies as an anonymizing approach under GDPR. The analysis emphasizes that the effectiveness of MPC as an anonymization method relies heavily on its architectural design, particularly the distribution of control among compute parties. The practical implementation evaluates SQL queries for privacy-preserving people analytics using MP-SPDZ, testing performance under various

security models. The findings confirm that MPC facilitates secure processing of sensitive organizational data while ensuring compliance with privacy regulations.

This work bridges regulatory and technical dimensions, offering organizations a clear roadmap for adopting MPC as a privacy-compliant data processing solution. By clarifying its legal implications and demonstrating its practical benefits, the study promotes the broader adoption of MPC for secure and efficient data analytics.

**Acknowledgments.** This work has been done partially in the CRYPTECS project that received funding from the Federal Ministry of Education and Research under Grant Agreement No. 16KIS1441.

## References

1. Acar, A., Aksu, H., Uluagac, A.S., Conti, M.: A Survey on Homomorphic Encryption Schemes: Theory and Implementation. ACM Computing Surveys (2018)
2. Altman, M., Cohen, A., Nissim, K., Wood, A.: What a Hybrid Legal-Technical Analysis Teaches Us About Privacy Regulation: The Case of Singling Out. BUJ Sci. & Tech. L. (2021)
3. Araki, T., Furukawa, J., Ohara, K., Pinkas, B., Rosemarin, H., Tsuchida, H.: Secure Graph Analysis at Scale. In: CCS (2021)
4. Becker, S., Duplys, P., Graf, J., Graffi, K., Grassi, A., Greven, D., Grewe, J., Jain, S., Klenk, T., Matyunin, N., Modica, H., Raskin, V., Scherer, P., Suschke, V., Trieflinger, S., Vlasakiev, V., Weinfurtner, J.: Carbyne stack, `https://carbynestack.io`
5. Ben-Itzhak, Y., Möllering, H., Pinkas, B., Schneider, T., Suresh, A., Tkachenko, O., Vargaftik, S., Weinert, C., Yalame, H., Yanai, A.: ScionFL: Efficient and Robust Secure Quantized Aggregation. In: IEEE SaTML. IEEE (2024)
6. Boemer, F., Cammarota, R., Demmler, D., Schneider, T., Yalame, H.: MP2ML: A Mixed-Protocol Machine Learning Framework for Private Inference. In: ARES (2020)
7. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A., Seth, K.: Practical Secure Aggregation for Privacy-Preserving Machine Learning. In: CCS (2017)
8. Brasser, F., Müller, U., Dmitrienko, A., Kostiainen, K., Capkun, S., Sadeghi, A.R.: Software Grand Exposure:SGX Cache Attacks Are Practical. In: WOOT (2017)
9. Brüggemann, A., Schick, O., Schneider, T., Suresh, A., Yalame, H.: Don't Eject the Impostor: Fast Three-Party Computation With a Known Cheater. In: IEEE S&P (2024)
10. University of California, B.: MPC Deployments (2024), `https://mpc.cs.berkeley.edu/`
11. Cramer, R., Damgård, I., Escudero, D., Scholl, P., Xing, C.: SPD $\mathbb{Z}_{2^k}$: Efficient MPC mod $2^k$ for Dishonest Majority. In: CRYPTO (2018)
12. Duddu, V., Das, A., Khayata, N., Yalame, H., Schneider, T., Asokan, N.: Attesting Distributional Properties of Training Data for Machine Learning. In: ESORICS (2024)
13. Dwork, C., Roth, A., et al.: The Algorithmic Foundations of Differential Privacy. Foundations and Trends® in Theoretical Computer Science (2014)

14. van Egmond, M.B., Dunning, V., van den Berg, S., Rooijakkers, T., Sangers, A., Poppe, T., Veldsink, J.: Privacy-preserving Anti-Money Laundering using Secure Multi-Party Computation. In: FC (2024)
15. Escudero, D., Ghosh, S., Keller, M., Rachuri, R., Scholl, P.: Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits. In: CRYPTO (2020)
16. Feigenbaum, J., Weitzner, D.J.: On The Incommensurability of Laws and Technical Mechanisms: Or, What Cryptography Can't Do. In: Cambridge International Workshop on Security Protocols (2018)
17. Fereidooni, H., Marchal, S., Miettinen, M., Mirhoseini, A., Möllering, H., Nguyen, T.D., Rieger, P., Sadeghi, A., Schneider, T., Yalame, H., Zeitouni, S.: SAFELearn: Secure Aggregation for Private Federated Learning. In: DLS@S&P (2021)
18. Gehlhar, T., Marx, F., Schneider, T., Suresh, A., Wehrle, T., Yalame, H.: SAFEFL: MPC-friendly Framework for Private and Robust Federated Learning. In: DLS@S&P (2023)
19. Goldreich, O., Micali, S., Wigderson, A.: How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In: STOC (1987)
20. Günther, D., Schmidt, J., Schneider, T., Yalame, H.: FLUENT: A Tool for Efficient Mixed-Protocol Semi-Private Function Evaluation. In: ACSAC (2024)
21. Harth-Kitzerow, C., Suresh, A., Wang, Y., Yalame, H., Carle, G., Annavaram, M.: High-Throughput Secure Multiparty Computation with an Honest Majority in Various Network Settings. PoPETs (2025)
22. Hegde, A., Möllering, H., Schneider, T., Yalame, H.: SoK: Efficient Privacy-Preserving Clustering. PoPETs (2021)
23. Helminger, L., Rechberger, C.: Multi-Party Computation in The GDPR. In: Privacy Symposium (2022)
24. Keller, M.: MP-SPDZ: A Versatile Framework for Multi-Party Computation. In: ACM CCS (2020)
25. Keller, M., Orsini, E., Scholl, P.: MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In: CCS (2016)
26. Keller, M., Sun, K.: Secure Quantized Training for Deep Learning. In: ICML (2022)
27. Knott, B., Venkataraman, S., Hannun, A.Y., Sengupta, S., Ibrahim, M., van der Maaten, L.: CrypTen: Secure Multi-Party Computation Meets Machine Learning. In: NeurIPS (2021)
28. MPCALLIANCE: Powering Secure Computation, Together (2024), `https://www.mpcalliance.org/`
29. Nguyen, T.D., Rieger, P., Chen, H., Yalame, H., Möllering, H., Fereidooni, H., Marchal, S., Miettinen, M., Mirhoseini, A., Zeitouni, S., Koushanfar, F., Sadeghi, A., Schneider, T.: FLAME: Taming Backdoors in Federated Learning. In: USENIX Security (2022)
30. Patra, A., Schneider, T., Suresh, A., Yalame, H.: ABY2.0: Improved Mixed-Protocol Secure Two-Party Computation. In: USENIX Security (2021)
31. Rubinstein, I.S., Hartzog, W.: Anonymization and Risk. Wash. L. Rev. (2016)
32. Scheibner, J., Raisaro, J.L., Troncoso-Pastoriza, J.R., Ienca, M., Fellay, J., Vayena, E., Hubaux, J.P.: Revolutionizing medical data sharing using advanced privacy-enhancing technologies: technical, legal, and ethical synthesis. JMIR (2021)
33. Spindler, G., Schmechel, P.: Personal data and encryption in the european general data protection regulation. J. Intell. Prop. Info. Tech. & Elec. Com. L. (2016)
34. Treiber, A., Müllmann, D., Schneider, T., genannt Döhmann, I.S.: Data protection law and multi-party computation: Applications to information exchange between law enforcement agencies. In: WPES. (2022)

35. Walsh, J.M., Varia, M., Cohen, A., Sellars, A., Bestavros, A.: Multi-regulation computing: Examining the legal and policy questions that arise from secure multiparty computation. In: Symposium on Computer Science and Law, CSLAW. (2022)
36. Yao, A.C.C.: Protocols for Secure Computations (Extended Abstract). In: FOCS (1982)

## Legal References (Footnotes)

[1]EuGH, Urt. vom 19.10.2016, Breyer – C-582/14, EU:C:2016:779, para. 31 ff., 46; most recently, in addition: EuG, judgment of 26 April 2023 – T-557/20, CR 2023 532, para. 103; on the decision-making history of European courts, see also: Arning/Rothkegel in: Taeger/Gabel, DS-GVO BDSG TTDSG, Art. 4 para. 48; Gierschmann, ZD 2021, 482 (483).

[2]Petri in: Simitis/Hornung/Spiecker gen. Döhmann, DSGVO/BDSG, Art. 28 para. 33; Martini in: Paal/Pauly, DS-GVO, BDSG, Art. 28 para. 8a; Bertermann/Peintinger in: Ehmann/Selmayr, DS-GVO, Art. 28 para. 8.

[3]Piltz in: Gola/Heckmann/Brand et al., DS-GVO BDSG, Art. 26 para. 4; Hartung in: Kühling/Buchner, DS-GVO, BDSG, Art. 26 para. 11.

[4]ECJ, judgment of October 6, 2015, Schrems – C-362/14, EU:C:2015:650, paras. 72 et seq.; on the purpose of the third-country transfer rules: Gabel in: Taeger/Gabel, DS-GVO BDSG TTDSG, Art. 44 para. 1; Thalhofer in: Bräutigam, IT-Outsourcing und Cloud Computing, Part 14 B. para. 68; Schantz in: Simitis/Hornung/Spiecker gen. Döhmann, DSGVO/BDSG, Art. 44 para. 6; Bussche/Raguse in: Plath, DSGVO/BDSG/TTDSG, Art. 44 para. 1.

[5]Piltz/Zwerschke in: Kipker/Reusch/Ritter et al, DS-GVO Art. 32 para. 6 ff; Hansen in: Simitis/Hornung/Spiecker gen. Döhmann, DS-GVO/BDSG, Art. 32 para. 18 ff.

[6]In detail: Schweinoch/Peintinger, CR 2020, 643 (644); Nink/Pohle, MMR 2015, 563 (565); Bierbauer/ Helminger, ALJ 2023, 1 (20); Kroschwald, ZD 2014, 75 (76).

[7]Heinson in: Weth/Herberger/Wächter, Datenschutz und Persönlichkeitsschutz im Arbeitsverhältnis, B. XII. para. 2; Knappertsbusch in: Braun/Wisskirchen, Braun/Wisskirchen/KonzernArbR, Part II, Chap. 8 A. para. 26; Schulz in: Gola/Heckmann/Brand et al, DS-GVO BDSG, Art. 6 para. 131.

[8]Regarding German law: Emmerich in: Emmerich/Habersack, Konzernrecht, AktG § 308 para. 52a.

[9]On this pair of terms: Stummer, DuD, 368 (369).

[10]Stummer, DuD, 368 (371).

[11]Stummer, DuD, 368 (371).

[12]ECJ, judgment of 5.6.2018, Wirtschaftsakademie SH - C-210/16, K&R 2018 475, para. 36; judgment of 29.7.2019, Fashion ID - C-40/17, NJW 2019 2755, para. 69, 79 ff.; however, this broad understanding of the term is not viewed uncritically in the literature and it is argued that the excessive extension of the scope of application contradicts the intention of the legislator. With further references: Hartung in: Kühling/Buchner, DS-GVO, BDSG, Art. 26 para. 43 ff; Spoerr in: Wolff/Brink/v. Ungern-Sternberg, BeckOK DatenschutzR, Art. 26 para. 26 describes this wording of the ECJ as misleading and states - without further justification - that a party cannot be (jointly) responsible if that party has permanent neither legal nor de facto access to the personal data. However, it is doubtful whether this is in line with the ECJ's interpretation.

[13]EuGH, Urt. vom 19.10.2016, Breyer – C-582/14, EU:C:2016:779, para. 31 ff., 46; most recently, in addition: EuG, judgment of 26 April 2023 – T-557/20, CR 2023 532, para. 103.

[14]EuG, Urt. vom 26.4.2023 – T-557/20, CR 2023 532, Rn. 103.

[15]Baumgartner, ZD 2023, 399 (403); Wildberg/Lee-Wunderlich, CCZ 2023, 281 (283); Bierbauer/ Helminger, ALJ 2023, 1 (22); Conrad/Folkerts, K&R 2023, 89 (91).

[16]Also coming to this conclusion: Treiber et al. in: Hong, WPES'22, Data Protection Law and Multi-Party Computation p. 74; also: Bierbauer/Helminger, ALJ 2023, 1 (22); briefly discussing the possibility of anonymization by MPC within the scope of the GDPR: Walsh et al. in: Weitzner/Feigenbaum/Yoo, Multi-Regulation Computing 58, whereby the authors come to the conclusion with regard to US anonymization requirements - which are comparable in parts to the European ones - that anonymization may be possible.

[17]Schumacher in: Bräutigam, IT-Outsourcing und Cloud Computing, Part 5 B. para. 99; Zerdick in: Ehmann/Selmayr, DS-GVO, Art. 44 para. 17; Gabel in: Taeger/Gabel, DS-GVO BDSG TTDSG, Art. 44 para. 15; Dovas/Grapentin in: Auer-Reinsdorff/Conrad, Hdb. IT- u. DsR., § 35 Grenzüberschreitende Datenverarbeitung para. 15.

[18]ECJ, judgment of 16.7.2020, Schrems-II - C-311/18, EU:C:2020:559, para. 128 et seq.

[19]Piltz/Zwerschke in: Kipker/Reusch/Ritter et al, DS-GVO Art. 32 para. 6 ff; Hansen in: Simitis/Hornung/Spiecker gen. Döhmann, DS-GVO/BDSG, Art. 32 para. 18 et seq.

[20]Bartels/Backer, DuD 2018, 214 (215); Piltz/Zwerschke in: Kipker/Reusch/Ritter et al, DS-GVO Art. 32 para. 14; Hansen in: Simitis/Hornung/Spiecker gen. Döhmann, DS-GVO/BDSG, Art. 32 para. 22.

[21]Bartels/Backer, DuD 2018, 214 (216); Hansen in: Simitis/Hornung/Spiecker gen. Döhmann, DS-GVO/BDSG, Art. 32 para. 23; Piltz in: Gola/Heckmann/Brand et al, DS-GVO, BDSG, Art. 32 para. 19.

[22]Classifying MPC as a sensible measure within the meaning of Art. 25 GDPR: Bierbauer/Helminger, ALJ 2023, 1 (26).

[23]On this functional approach to the definition: Piltz/Zwerschke in: Kipker/Reusch/Ritter et al, DS-GVO Art. 32 para. 48.

[24]On the technology-neutral approach of the GDPR in the context of technical and organizational protection measures: Piltz in: Gola/Heckmann/Brand et al, DS-GVO BDSG, Art. 32 para. 23; Hansen in: Simitis/Hornung/Spiecker gen. Döhmann, DS-GVO/BDSG, Art. 32 para. 20; Piltz/Zwerschke in: Kipker/Reusch/Ritter et al, DS-GVO Art. 32 para. 38.

[25]For details on the risk-mitigating effect of a contractual prohibition of re-identification: Conrad/Folkerts, K&R 2023, 89 (92, 94); also assuming a risk-mitigating effect: Johannes/Geminn, MedR 2023, 368 (371); Bierbauer/Helminger, ALJ 2023, 1 (9) expressly in the context of MPC instances: "In this consideration, organizational measures, such as explicitly placing the cooperation of two parties (contrary to protocol) under contractual penalty, can also be taken into account."

[26]The question of whether the results of a calculation (output) are anonymous is summarized under the term "output privacy", see: Böhler, Input Secrecy & Output Privacy: Efficient Secure Computation of Differential Privacy Mechanisms, p. i; explicitly referring to this problem in the context of MPC: Lindell, Secure Multiparty Computation (MPC), p. 6.

[27]On the potential benefits of differential privacy in the context of an MPC system: Liagouris et al, Secrecy: Secure collaborative analytics on secret-shared data, p. 15.

[28]Schulz in: Gola/Heckmann, DSGVO, Art. 6 para. 59 et seq.

[29]Hornung/Wagner, ZD 2020, 223 (225).

# Appendix

## A1. Comparison to Other Technologies

MPC is compared to other PETs to assess differences in legal compliance and technical security.

1. How does Computing on Encrypted Data (CoED), specifically MPC, compare with Confidential Computing (CC) regarding data protection? Would a data breach differ in severity depending on whether CC or MPC was used?
2. What are the implications of ongoing practical attacks in CC implementations versus formal security proofs in CoED methods?

The comparison of Computing on Encrypted Data (CoED), specifically MPC, and hardware-based PETs like Confidential Computing (CC) reveals differences in their approaches to data protection and security. MPC, as a CoED technology, enables computations on encrypted data without requiring plaintext exposure after initial secret share fragmentation. The fundamental distinction lies in their processing methodology: while CC requires data decryption within protected hardware environments (Trusted Execution Environments, TEEs),[30] MPC maintains continuous encryption through distributed computation. This difference becomes particularly relevant in breach scenarios, where CC vulnerabilities might expose plaintext data, while attacks on one MPC-compute party would only reveal encrypted fragments, requiring coordinated compromise across multiple parties for data reconstruction.

CC's hardware-based approach has demonstrated vulnerabilities through exploits such as Meltdown and Spectre. However, these vulnerabilities' practical

impact is significantly mitigated by the complexity of exploitation, which requires extraordinary technical expertise and substantial resources, often rendering attacks impractical in real-world scenarios. Moreover, rapid security patch deployment and the limited utility of potentially extractable data further diminish these risks.[31] Conversely, CoED methods like MPC offer formally proven security guarantees based on mathematical foundations. While this provides robust theoretical security, practical implementation must consider the chosen security models' implications and potential collusion risks among participating parties. The effectiveness of these security guarantees depends heavily on proper protocol implementation and participant behavior.

The selection between these technologies ultimately depends on specific use case requirements rather than absolute security superiority. For distributed analytics across multiple entities, MPC's cryptographic approach typically offers more appropriate security guarantees, particularly when continuous data encryption is paramount. Conversely, CC might better serve scenarios requiring secure single-entity processing with cloud protection, where hardware-based isolation provides sufficient security assurance in a scenario where the CSP would otherwise not be trusted.

**Takeaway 6:** CoED technologies and hardware-based PETs face different vulnerability types—e.g. hardware exploits for CC versus potential party collusion for MPC—their selection should primarily depend on specific use-case requirements rather than absolute security considerations.