CONFIDENTIAL COMPUTING CONSORTIUM
Governance Risk
& Compliance
Special Interest Group

GDPR Compliance with Confidential Computing in the AI Era

August 2025

Abstract

This article explores the pivotal role of Confidential Computing in achieving GDPR compliance, particularly in the context of AI data processing. GDPR, a comprehensive data protection regulation, mandates stringent measures to ensure the secure handling of personal data. While traditional encryption methods secure data at rest and in transit, they do not protect data during active processing—a critical phase for AI operations. Confidential Computing addresses this gap by utilizing hardware-based Trusted Execution Environments (TEEs) to create secure enclaves, ensuring data remains encrypted and shielded from unauthorized access during computation. This technology not only enhances GDPR compliance by safeguarding data in use but also aligns with the regulation's principles of data minimization and security by design.

Special considerations for AI are highlighted, as AI models require access to large datasets, often containing sensitive information. Confidential Computing enables secure AI model training, inference, and multi-party collaboration, ensuring data privacy and compliance with GDPR standards. By protecting data throughout the AI processing workflow, organizations can demonstrate due diligence, build trust, and reduce breaches, their impacts, and consequent notification obligations. Furthermore, Confidential Computing supports GDPR-compliant cross-border data transfers, providing robust safeguards against unauthorized access during international processing. This article illustrates how Confidential Computing not only meets GDPR's legal and ethical imperatives but also empowers organizations to responsibly leverage AI technologies while maintaining high data protection

standards.

Table of Contents

Introduction

Data privacy and security are paramount for organizations globally, particularly in the digital and artificial intelligence (AI) era. The General Data Protection Regulation (GDPR), enacted by the European Union, stands as one of the most comprehensive data protection frameworks. It grants individuals enhanced control over their personal data and obligates organizations to implement rigorous data privacy measures. The regulation's key principles include data minimization, purpose limitation, and integrity and confidentiality. Non-compliance can lead to severe financial and reputational repercussions.

Central to GDPR compliance is the necessity for robust data security practices. Encryption serves as a fundamental mechanism, safeguarding data from unauthorized access by converting it into an unreadable format. Traditionally, encryption has effectively secured data at rest (stored data) and data in transit (data moving across networks). However, securing data during active processing or use is often neglected. During computation, data that was encrypted at rest or in transit typically requires decryption, exposing it to unauthorized access and modification by various actors, including device owners and administrators.

These vulnerabilities are particularly critical in the context of AI. AI models necessitate access to substantial data volumes for training and inference, often encompassing personal or sensitive information. Processing such data without adequate protections poses significant risks, as exposure during this phase can result in data breaches, GDPR non-compliance, and a loss of trust from users and stakeholders. Traditional encryption methods, while essential, are insufficient for protecting data in use during AI operations, underscoring a gap in current data protection strategies.

Confidential computing emerges as a transformative technology to address these challenges. It involves utilizing hardware-based Trusted Execution Environments (TEEs), sometimes called secure enclaves, to create isolated and protected spaces within a computing system. These environments ensure that data remains encrypted at rest and in transit, while also shielding it against unauthorized use during processing. This approach significantly reduces the risk of data exposure, even if the underlying system is compromised. By extending isolation safeguards to data in use, confidential computing meets the advanced data protection needs outlined by GDPR and other global regulations, providing a critical safeguard for AI data processing.

In this paper, we explore how confidential computing strengthens compliance with GDPR's requirements by addressing the inherent vulnerabilities associated with data actively in use. Confidential computing acts as a natural extension that bolsters traditional data protection methods, playing a critical role in specific AI use cases to secure data processing. This ensures that AI models can be trained and deployed without compromising the privacy and rights of individuals.

Through this analysis, we aim to illustrate that confidential computing not only provides a technological advantage but also aligns with the ethical and legal imperatives mandated by GDPR and other data privacy regulations. By embracing this advanced technology, organizations can enhance their data protection strategies, foster trust with their users, and responsibly harness the power of AI without fear of non-compliance or data exposure.

Confidential Computing & GDPR Compliance

GDPR mandates stringent data protection standards that require organizations to implement robust measures ensuring the confidentiality, integrity, and availability of personal data. While traditional security methods, such as encryption, are essential, they can fall short when it comes to protecting data during active processing. This is where confidential computing plays a critical role, ensuring that decrypted data in use remains shielded and thereby enhancing GDPR compliance in significant ways.

CORE PRINCIPLES OF CONFIDENTIAL COMPUTING

Secure Data Processing: Confidential computing leverages hardware-based Trusted Execution Environments (TEEs) to create secure, isolated environments where data can be processed without exposing it to unauthorized parties. These TEEs ensure that data is only decrypted within a protected enclave, making it unreadable to unauthorized entities, even if the system is compromised.

Confidentiality and Integrity: The design of TEEs guarantees that data processed within the enclave is not exposed to other applications, users, or in some constructions even the operating system. This ensures that even highly privileged users or malware with administrative access cannot intercept or tamper with the data.

Auditability: Confidential computing includes Remote Attestation in which the hardware digitally signs evidence about the security state of the system. These secure digital messages can be logged and used for audit to show that security measures were in place during given events.

Performance: Typically confidential computing products implement memory encryption in hardware. This hardware accelerated encryption minimizes performance impacts to workloads. Moreover, hardware based Trusted Execution Environments generally offer the best performance to organisations compared to other Privacy Enhancing Technologies such as homomorphic encryption.

BENEFITS FOR ORGANIZATIONS IMPLEMENTING CONFIDENTIAL COMPUTING

Minimized Breach Notification Obligations: Article 34 of GDPR outlines the conditions under which data controllers must notify data subjects of a breach. If data is encrypted and unintelligible to unauthorized users, organizations may be exempt from this requirement. By using confidential computing to maintain data encryption during processing, the likelihood of breach notifications and associated reputational damage is significantly reduced.

Building Trust and Demonstrating Due Diligence: Implementing confidential computing can help organizations showcase their commitment to data protection and compliance. This proactive approach can enhance trust among customers and stakeholders, demonstrating that the organization goes beyond the minimum legal requirements to protect personal data.

Future-Proofing Data Security: As GDPR evolves and other data privacy regulations are introduced or updated, the foundational security provided by confidential computing positions organizations to adapt more seamlessly to new standards. The capability to secure data in use aligns with global trends in privacy legislation, which increasingly emphasize comprehensive data protection.

Data in Use Protection: By protecting data during the processing phase, confidential computing meets these criteria in Article 32 which mandates the need for measures that ensure the security of processing activities, including encryption.

Reducing the Attack Surface: With data shielded during processing, confidential computing significantly reduces the potential attack surface. This added security measure prevents data exposure even in the presence of vulnerabilities such as memory scraping attacks, which could otherwise compromise sensitive data.

Data Minimization and Security by Design: GDPR advocates for data minimization and the concept of security by design. Confidential computing supports these principles by only processing the necessary data in a secure manner, inherently aligning with a minimal and protected approach to data handling.

IMPLEMENTATION CHALLENGES AND CONSIDERATIONS

Infrastructure and Integration: While confidential computing offers significant benefits, organizations need to assess their existing infrastructure to determine compatibility and integration requirements. Deploying TEEs may involve updating hardware, software, or both.

Cost and Resource Allocation: The implementation of confidential computing technologies can require upfront investment in enabled hardware, attestation infrastructure and processes, and training for IT teams. However, the long-term benefits in terms of enhanced compliance and reduced risk of breaches often outweigh these initial costs.

Adoption and Skill Gaps: Organizations may face challenges related to expertise, as architects must understand confidential computing principles. Deployment, however, generally does not require expertise beyond traditional systems administration. Investments in training and partnerships with vendors who provide confidential computing solutions can help bridge these gaps.

GDPR COMPLIANCE STRATEGY

Confidential computing should be viewed as part of a broader data protection strategy resulting in a Comprehensive Security Posture which includes traditional encryption, access controls, and data governance policies. By addressing the limitations of traditional encryption and providing security during data processing, confidential computing complements existing measures to create a comprehensive security framework.

Confidential computing can be used alongside other privacy-enhancing technologies (PETs), such as homomorphic encryption and secure multi-party computation, to further strengthen GDPR compliance and create robust solutions for data processing in complex environments.

GDPR is built around several key articles that establish the expectations and requirements for data protection. Confidential computing addresses gaps in traditional data security practices, aligning with these articles to enhance compliance. Below, we highlight specific GDPR articles and illustrate how confidential computing can support compliance efforts:

ARTICLE 5: PRINCIPLES RELATING TO PROCESSING OF PERSONAL DATA

What It States: Article 5 outlines key principles for processing personal data, including lawfulness, fairness, transparency, data minimization, accuracy, storage limitation, integrity, and confidentiality.

How Confidential Computing Helps: Confidential computing enhances the principle of integrity and confidentiality by ensuring that personal data is protected not only at rest and in transit but also during active processing. By leveraging secure enclaves, organizations can process data in an encrypted state, reducing the risk of unauthorized access and ensuring that data handling aligns with GDPR's overarching data protection principles.

ARTICLE 24: RESPONSIBILITY OF THE CONTROLLER

What It states: Article 24 mandates implementation of appropriate technical measures to ensure and to be able to demonstrate that personal data processing is performed sufficiently, assessed, for example, by external audit or respective certification procedure.

How Confidential Computing Helps: Remote attestations provide hardware signed evidence about the security state of a system. These digitally signed messages should be logged for audit. Given their cryptographic nature they provide strong evidence.

ARTICLE 25: DATA PROTECTION BY DESIGN AND BY DEFAULT

What It States: Article 25 mandates that data protection measures be integrated into processing activities from the outset. Organizations must implement appropriate technical and organizational measures to uphold data protection principles and safeguard data subjects' rights.

How Confidential Computing Helps: Confidential computing embodies the principle of security by design, as it allows organizations to define isolation at design time by for example indicating which programs must run in TEEs. Additionally, specifying confidential computing in infrastructure design satisfies secure by default goals. For example, by using confidential virtual machines rather than conventional virtual machines, memory is automatically protected without need for software implementation or runtime configuration for individual workloads.

ARTICLE 32: SECURITY OF PROCESSING

What It States: Article 32 requires organizations to implement appropriate technical and organizational measures to ensure a level of security appropriate to the risk, such as encryption, the ability to ensure ongoing confidentiality, and protection against accidental or unlawful destruction or loss.

How Confidential Computing Helps: While traditional encryption secures data at rest and in transit, Article 32 implicitly calls for protecting data during active processing as well. Confidential computing fills this gap by ensuring that data remains encrypted within secure enclaves during computation. This approach enhances the security of processing activities and aligns with GDPR's expectations for comprehensive data protection measures.

ARTICLE 34: COMMUNICATION OF A PERSONAL DATA BREACH TO THE DATA SUBJECT

What It States: This article mandates that data controllers inform data subjects of a data breach if it is likely to result in a high risk to their rights and freedoms. However, if the breached data is encrypted and rendered unintelligible, the notification requirement may not apply.

How Confidential Computing Helps: By ensuring that data is encrypted and protected during processing, confidential computing minimizes the risk of data breaches that could lead to exposure. If a breach were to occur, the use of confidential computing might eliminate or lessen the requirement to notify data subjects, as the data would be encrypted and unreadable to unauthorized parties.

ARTICLE 35: DATA PROTECTION IMPACT ASSESSMENT (DPIA)

What It States: Article 35 requires organizations to conduct DPIAs for data processing activities that are likely to result in high risks to the rights and freedoms of individuals. The DPIA must demonstrate that data protection measures are in place to mitigate these risks.

How Confidential Computing Helps: Implementing confidential computing can be a key part of a DPIA strategy, demonstrating that advanced data protection measures are in place to safeguard data during processing. Organizations can showcase that data is processed within secure enclaves, significantly reducing risk exposure and aligning with GDPR's emphasis on high standards of protection.

ARTICLE 89: SAFEGUARDS AND DEROGATIONS RELATING TO PROCESSING FOR ARCHIVING PURPOSES IN THE PUBLIC INTEREST, SCIENTIFIC OR HISTORICAL RESEARCH PURPOSES, OR STATISTICAL PURPOSES

What It States: Article 89 allows for the processing of personal data for research, archiving, or statistical purposes under specific conditions, including the implementation of appropriate safeguards to protect data subjects' rights.

How Confidential Computing Helps: Confidential computing supports research and statistical data processing by enabling secure data handling without exposing raw data. This technology can ensure that even sensitive data used for AI model training or large-scale research is protected, fulfilling GDPR's requirements for strong safeguards in specialized data processing scenarios. Additionally remote attestation provides evidence about the security measures in place when data was processed.

Cross Border Transfers

In addition to the aforementioned articles, Cross Border Transfers impose additional requirements which are also satisfiable with confidential computing. Confidential computing offers significant advantages in this context by minimizing the risk of data exposure. Data processed within a Trusted Execution Environment (TEE) remains encrypted, making it unreadable to unauthorized parties and reducing the risk of exposure during transfers. This technology provides a stronger demonstration of compliance, as it shows that an organization is taking advanced steps to protect data, which can be critical when dealing with Standard Contractual Clauses (SCCs), Binding Corporate Rules (BCRs), or when demonstrating due diligence to regulatory bodies. Furthermore, confidential computing enhances trust in data processing partners. When engaging third-party processors or collaborating internationally, organizations can use remote attestation to ensure that partners uphold the same data protection standards.

DATA PROTECTION DURING TRANSFERS (ARTICLE 44)

Article 44 and General Principles: Confidential computing ensures that data is protected not just at rest and in transit, but also during active processing in third countries. This aligns with Article 44's requirement that data transferred outside the EU must still adhere to GDPR's level of data protection. By keeping data encrypted within secure enclaves during processing, confidential computing reduces the risk of unauthorized access, even when the data is processed in less regulated regions.

SUPPORT FOR APPROPRIATE SAFEGUARDS (ARTICLE 46)

Standard Contractual Clauses (SCCs) and Binding Corporate Rules (BCRs): When organizations use SCCs or BCRs as safeguards for cross-border transfers, they need to demonstrate that data is protected at all stages. Confidential computing strengthens these safeguards by providing an additional layer of protection during processing, ensuring that personal data remains encrypted and secure within a Trusted Execution Environment (TEE). This makes it easier for organizations to meet GDPR's requirements for demonstrating adequate protection under SCCs or BCRs.

PROTECTING AGAINST UNAUTHORIZED ACCESS (ARTICLE 48)

Limitations on Non-EU Authority Requests: Confidential computing helps ensure that data processed outside the EU remains secure, even if external authorities attempt to access it. With data protected within TEEs, unauthorized entities cannot access or decrypt the data without proper authorization. This compliance measure aligns with Article 48, which restricts data transfers in response to non-EU authorities' requests unless based on an international agreement.

ENABLING COMPLIANT CROSS-BORDER PROCESSING (ARTICLE 49)

Derogations for Specific Situations: For cases where data needs to be transferred under specific derogations (e.g., explicit consent or public interest), confidential computing provides peace of mind that data is secure even when processed in potentially less secure jurisdictions. This technology ensures that personal data remains encrypted and protected, reducing risks and enhancing compliance with GDPR during cross-border processing.

Special Considerations for AI

The rapid integration of artificial intelligence (AI) across various sectors has introduced significant advancements in data analysis, decision-making, and automation. However, the use of AI often requires access to large datasets that can include sensitive personal information. This necessity for data processing raises significant concerns about privacy and compliance with regulations like the GDPR, which mandates strict data protection measures. Confidential computing has emerged as a vital solution to address these challenges, enabling AI to process sensitive data securely and in compliance with regulatory requirements.

Enhanced Data Security: Data remains shielded throughout the AI processing workflow, reducing the risk of exposure and aligning with GDPR's data protection standards.

Greater Trust and Compliance: Organizations can demonstrate that they are taking advanced steps to protect personal data, which fosters trust among clients and regulators and supports overall GDPR compliance.

Scalable Privacy Solutions: Confidential computing provides a scalable way to integrate data protection into AI operations without significantly affecting performance, making it suitable for a wide range of industries and applications.

PROTECTING DATA DURING AI MODEL TRAINING

The Challenge: AI model training often involves extensive datasets, which can include personal data such as customer records, medical information, or financial transactions. Traditional data protection measures, like encryption at rest and in transit, do not protect data when it is being processed during training. This leaves a critical gap where sensitive data is exposed, potentially violating GDPR's standards for data security.

How Confidential Computing Helps: Confidential computing ensures that data remains shielded during the training process by using Trusted Execution Environments (TEEs). These secure enclaves isolate and protect data, preventing it from being exposed even if the system's other components are compromised. This allows organizations to train AI models using sensitive data while maintaining GDPR compliance and protecting data integrity and confidentiality.

SECURE AI INFERENCE AND REAL-TIME PROCESSING

The Challenge: AI-powered applications often require real-time data processing, such as predictive analytics in financial services or diagnostic tools in healthcare. During inference, AI models process new inputs to generate predictions or insights. If these inputs contain personal data, traditional methods may expose them during processing.

How Confidential Computing Helps: By leveraging confidential computing, AI inference can be conducted in a protected environment where input data remains shielded throughout the computation process. This ensures that sensitive data is never exposed, aligning with GDPR's requirement for robust data security. For instance, in healthcare, patient data used in diagnostic tools can be processed securely, maintaining patient privacy and reducing compliance risks.

MULTI-PARTY AI COLLABORATION

The Challenge: Collaborative AI projects often require data sharing among different organizations. For example, research institutions, healthcare providers, or financial entities may need to collaborate on shared models or data sets. GDPR requires that data sharing be handled with strict controls to prevent unauthorized access or breaches.

How Confidential Computing Helps: Confidential computing enables secure multi-party computation by allowing multiple parties to process shared data collaboratively within a TEE, without revealing the raw data to any of the participants. This facilitates privacy-preserving AI collaboration and ensures that each party's data is protected throughout the computation process. Organizations can train shared models on collective datasets without compromising individual data privacy, achieving both innovation and GDPR compliance.

FEDERATED LEARNING WITH ENHANCED PRIVACY

The Challenge: Federated learning is an AI training technique where models are trained across decentralized devices or servers holding local data samples, without exchanging the data itself. While federated learning enhances privacy by keeping data local, the aggregation phase can still present security risks. Each one of these devices is a trusted portion of the overall system and it can be difficult to enforce and audit security measures across such a distributed system.

How Confidential Computing Helps: By integrating confidential computing, the aggregation of the locally trained models can occur within a TEE, ensuring that intermediate and final results are protected during processing. This provides an added layer of security, ensuring that data privacy is maintained throughout the entire federated learning process. While this approach may not directly minimize the data the isolation of the aggregation supports GDPR's requirements for data minimization and secure processing, enabling the use of distributed AI techniques in a compliant manner. Additionally, if all of the distributed devices implement confidential computing, then there is added assurance that each device is processing its portion of the data in a compliant manner. Moveover, with remote attestation, evidence of that compliance is auditable.

PRIVACY-PRESERVING AI FOR SENSITIVE INDUSTRIES

Healthcare: AI applications in healthcare often require the processing of highly sensitive patient data. Confidential computing enables secure data analysis for diagnostic models, predictive tools, and personalized treatment plans while maintaining patient privacy. This ensures that healthcare organizations can use AI to improve patient outcomes without breaching GDPR guidelines.

Financial Services: Financial institutions use AI for fraud detection, risk assessment, and customer insights, all of which involve processing personal financial information. Confidential computing ensures that this data is protected during analysis, preventing unauthorized access and aligning with GDPR's data protection principles.

Government and Public Sector: AI can support data-driven policymaking and public services but often requires access to citizens' data. Confidential computing helps protect this data during analysis, enabling governments to harness AI capabilities while maintaining public trust and complying with stringent privacy regulations.

EXAMPLE APPLICATIONS

Healthcare AI Diagnostics: A medical research institution develops an AI model to detect diseases using patient imaging data. By utilizing confidential computing, the institution can process sensitive patient data within TEEs, ensuring that data is protected throughout the training and inference processes. This approach meets GDPR requirements for data security and patient confidentiality.

Financial AI Analytics: A financial firm uses confidential computing to analyse transaction data in real-time for fraud detection without exposing sensitive customer information. The firm processes data within secure enclaves, ensuring that data privacy is preserved and demonstrating compliance with GDPR.

Conclusion

Confidential computing is a transformative technology that significantly enhances data protection by securing data in use. This advancement is crucial for organizations aiming to comply with stringent data privacy regulations like GDPR, especially in AI and data-intensive operations. By adopting confidential computing, organizations can ensure comprehensive data protection throughout its lifecycle, thereby reducing compliance risks and fostering trust with stakeholders. Furthermore, confidential computing facilitates GDPR-compliant cross-border data transfers, enabling secure international data processing while maintaining high privacy standards. This technology not only supports innovation in AI but also strengthens the overall data protection framework, making it an indispensable tool for modern enterprises.

Acknowledgements

The Confidential Computing Consortium (CCC) extends its gratitude to Paul O’Neill and Intel Corporation for his insightful analysis and for authoring the initial draft of this document, which served as the foundation for subsequent contributions and refinements. The CCC also thanks Dan Middleton and Johnita Denson for their extensive work incorporating community feedback and preparing the final draft.