

CONFIDENTIAL COMPUTING AND PRIVACY

Policy Implications of Trusted Execution Environments

JULY 2024



**FUTURE OF
PRIVACY
FORUM**

AI

AUTHORS

Samuel Adams

Policy Counsel, Future of Privacy Forum

Stacey Gray

Senior Director for U.S. Policy, Future of Privacy Forum

Dr. Aaron Massey

Senior Policy Analyst and Technologist, Future of Privacy Forum

Dr. Rob van Eijk

Managing Director for Europe, Future of Privacy Forum

ACKNOWLEDGEMENTS

The authors would like to thank Lee Matheson, Senior Counsel for Global Privacy (Future of Privacy Forum), Amie Stepanovich, Vice President for U.S. Policy (Future of Privacy Forum), Yasha Doddabele, 2023 U.S. Policy Intern (Future of Privacy Forum), Ryan Campbell, 2024 U.S. Policy Intern (Future of Privacy Forum), Sonia Saini, 2024 U.S. Policy Intern (Future of Privacy Forum), as well as the many expert stakeholders from the FPF community who provided invaluable feedback on early drafts of this paper.



About Future of Privacy Forum (FPF)

The **Future of Privacy Forum (FPF)** is a non-profit organization that serves as a catalyst for privacy leadership and scholarship, advancing principled data practices in support of emerging technologies. Learn more about FPF by visiting fpf.org.

About FPF's Center for Artificial Intelligence

The **Center for Artificial Intelligence** at the Future of Privacy Forum is dedicated to navigating the complex landscape of AI governance and its intersection with privacy and data protection law.

Drawing on expertise from a global Leadership Council comprising industry leaders, academics, civil society, and policymakers, the Center provides sophisticated, practical policy analysis to help organizations align innovation with responsible implementation while meeting evolving regulatory requirements. Learn more about the FPF Center for AI at fpf.org/ai.



All FPF materials that are released publicly are free to share and adapt with appropriate attribution. Learn more at creativecommons.org.

Executive Summary

Confidential computing promises a significant shift in trustworthiness and verifiability of data processing for the use cases it supports, including training and use of artificial intelligence (AI) models. The technology allows organizations to restrict access to information (such as personal information, intellectual property, or other sensitive or high-risk data) only to specific parties and for specific purposes, through the use of a secure, hardware-based enclave, or “trusted execution environment” (TEE). A TEE allows for isolated data processing with the use of a separate attestation process that governs access and use of data **(Part A)**. Especially for cloud service providers, the use of hardware isolation and cryptographic attestation can reduce the need for trust between entities that otherwise rely primarily on contractual agreements.

Early adoption of confidential computing is particularly prominent in economic sectors that are facing privacy and security challenges due to regulatory pressures, the sensitivity of data, and the use of neural network-based AI models, including financial services, healthcare, and advertising **(Part B)**. As manufacturers develop and adopt confidential computing technologies, policymakers and practitioners should consider a range of potential data protection implications — including transparency and accountability, risks of re-identification, regulatory restrictions related to access and “sale,” law enforcement access, cross-border data transfers, and data localization **(Part C)**.

A. The Basics: What is Confidential Computing?

Confidential computing is an emerging term for a **privacy-enhancing technology (“PET”)** that has the novel capability to isolate data processing from the rest of a computer and prevent unauthorized access, even from administrators with elevated system privileges. Although organizations define the precise contours of the term differently (**Table 1**), confidential computing generally involves a range of hardware-based technologies that isolate data processing and provide technically-enforceable administrative safeguards for data while it is actively in use.

Source	Leading Descriptions and Definitions ¹
IBM	“Confidential computing technology isolates sensitive data in a protected CPU enclave during processing. The contents of the enclave, which include the data being processed and the techniques that are used to process it, are accessible only to authorized programming codes. They are invisible and unknowable to anything or anyone else, including the cloud provider.”
Fortinet	“Confidential computing refers to cloud computing technology that can isolate data within a protected central processing unit (CPU) while it is being processed.”
National Institute of Standards and Technology (NIST)	“Hardware-enabled features that isolate and process encrypted data in memory so that the data is at less risk of exposure and compromise from concurrent workloads or the underlying system and platform.”
Intel	“Confidential Computing offers a hardware-based security solution designed to help protect data in use via unique application-isolation technology called a Trusted Execution Environment (TEE).”
Confidential Computing Consortium	“Confidential Computing is the protection of data in use by performing computation in a hardware-based, attested Trusted Execution Environment.”
Information Commissioner's Office (ICO)	“A trusted execution environment (TEE) is a secure area inside a computing device’s central processing unit (CPU). It runs code and accesses information in a way that is isolated from the rest of the system.”

Table 1: Comparison of leading descriptions of confidential computing. Confidential Computing Consortium (CCC) and NIST provide industry-wide definitions that many organizations follow, while other descriptions may reflect competing architectural implementations or marketing language.

Confidential computing leverages two key technologies: trusted execution environments and attestation services. **Trusted execution environments (“TEEs”)** are separate and secure areas within a central processing unit (“CPU”) that isolate data processing from the rest of a computer system.² Such isolated processing can offer a spectrum of security and privacy protections for highly sensitive data, the processing of which warrants enhanced safeguards against unauthorized access.³

A core function of a TEE is the use of a separate **attestation process** to verify the proper configuration of the TEE. The process of attestation provides **cryptographic verification** of the hardware and configuration elements of the TEE that govern access to the isolated processing.⁴ The configuration of the TEE provides the “rules,” insofar as it governs which entities are authorized to input data, the kinds of datasets and programs that can be processed, and the permitted outputs. Meanwhile, attestation provides a layer of cryptographic verification that those “rules” have been followed, i.e., verification of the TEE’s configuration. The organization that provides a TEE is often, but not necessarily, the same entity that owns or hosts the attestation services for that TEE. Similarly, the physical infrastructure for TEEs and their accompanying attestation services do not need to be located in the same place.⁵

Confidential computing differs in key ways from other Privacy Enhancing Technologies (PETs), including **homomorphic encryption (HE)** and **secure multiparty computation (SMPC)**. All are cryptographic techniques used to ensure privacy and security in computations involving sensitive data.⁶ However, while both HE and SMPC offer the means of **protecting data in use** for limited forms of mathematical comparison between datasets, trusted execution environments provide greater flexibility by establishing an isolated environment for a range of general-purpose computing.⁷ By offering the ability to isolate data processing in an attestable environment, confidential computing can alleviate a fundamental lack of trust that many organizations have in **cloud computing**, the on-demand service for computing resources, including applications, servers, data storage, and networking.

B. Emerging Sectoral Applications

Early adopters of confidential computing include organizations in regulated industries, such as healthcare and financial services, because these organizations commonly process highly sensitive personal data in collaboration with other organizations.⁸ Meanwhile, the regulatory and platform pressures on the online advertising industry are creating incentives for advertising providers to implement stricter safeguards and protections around personal information. At least some large platforms, including Google and Microsoft, are beginning to implement trusted execution environments in browser-based advertising as an alternative to on-device processing.

1. Health Care

Health care data, such as that located within patient records, can reflect inherently sensitive details about individuals’ lives, including medical conditions, treatments, and medications. Many laws, in the U.S. and globally, regulate the collection, use, and disclosure of health records and codify security and data protection requirements to safeguard against unauthorized or unlawful access.⁹

At the same time, large quantities of representative and detailed real-world health data are a prerequisite for training reliable and generalizable AI models.¹⁰ Across the healthcare sector, AI technologies are being applied for a wide variety of uses, including administrative tasks, patient engagement, and diagnosis and treatment recommendations.¹¹ Given the sensitivity and risk associated with processing such data, medical institutions have historically been limited in their ability to share or pool data across datasets to achieve these ends, relying instead on legal exemptions for de-identified data or the use of “limited data sets,” the latter of which refer to health information stripped of certain patient identifiers and subject to a data use agreement.¹² While effective for many purposes, HIPAA-compliant de-identification typically reduces the utility of datasets, often making it more challenging to detect and mitigate bias or study rare diseases or conditions.¹³

In response to these challenges, some health research institutions are adopting confidential computing to safeguard patient health records used to train AI models from unauthorized access. For example, the University of California at San Francisco's Center for Digital Health Innovations (CDHI) hosts a confidential computing platform to enable AI developers' algorithms to interact with clinical data for training purposes.¹⁴ According to CDHI, the clinical datasets, the developers' AI models, and the processing environment (where the algorithm interacts with the data) are stored in separate secure TEEs.

2. Banking and Financial Services

Banking and financial institutions are heavily regulated across several legal regimes, including in laws on consumer protection, financial stability, and fraud prevention.¹⁵ In particular, anti-money laundering (AML) laws require financial institutions to record, analyze, and proactively report certain financial transactions suspected of being associated with money laundering or other illegal activity.¹⁶ As a result, large institutions may analyze billions of transactions per day per institution, in furtherance of compliance programs that cost U.S. banks an estimated \$25 billion annually and costs the average bank around \$48 million per year.¹⁷

In response to these compliance challenges, many institutions implement automated AML programs using artificial intelligence models to detect money laundering and other forms of financial fraud.¹⁸ In many cases, AML and related fraud detection services are procured from vendors that train AI models on large datasets of financial information, including across multiple institutions.¹⁹ In sharing customer data with AML vendors, financial institutions must navigate restrictions on data sale and transfer (see below, Part C) and typically rely on contractual agreements that the vendor will not re-use or fail to adequately protect the data. Meanwhile, vendors may seek to add additional security measures to ensure they meet regulatory requirements and contractual obligations without impacting their ability to effectively and efficiently analyze the data.

In this context, confidential computing can allow financial institutions to transfer encrypted customer financial information with AML vendors with technical safeguards that reduce the risk

that the data can be accessed by unauthorized individuals or for unauthorized purposes, decreasing the reliance on trust through contractual agreements. Although banks and financial institutions are often slow to migrate from legacy technologies, they are also driven by the need for advanced data analytics while complying with regulatory standards.²⁰ As a result, banks and other financial institutions, such as Swift (Society for Worldwide Interbank Financial Telecommunication), are beginning to use confidential computing to enable AI models that can detect potential fraud or other anomalies in financial data without violating privacy.²¹

3. Online Advertising

For several decades, digital advertising has relied on the relatively unrestricted collection and use of individuals' personal information within a complex ecosystem of advertising intermediaries designed to broker sales between advertisers and publishers, deliver targeted ads, and measure ad effectiveness. These practices have increasingly been the subject of regulatory pressure from U.S. and global lawmakers focused on strengthening individual privacy protections and preventing data-related harms. Meanwhile, large platforms and web browsers have taken significant steps to limit access to advertising-related data about their users, through, for instance, Apple's Intelligent Tracking Prevention (ITP), Google's ongoing deprecation of third party cookies, and similar efforts.²²

As a result of these pressures, advertising technology providers are actively testing a range of privacy-enhancing technologies (PETs), as part of an overall trend toward limiting the unrestricted sharing of personal identifiers.²³ Both Google and Microsoft have begun implementing trusted execution environments in browser-based advertising, in order to leverage the speed, energy efficiency, and scalability that remains challenging to achieve with on-device processing.²⁴ Meanwhile, Apple is also deploying confidential computing as part of their artificial intelligence infrastructure.²⁵

In comparison to typical cloud infrastructure, a TEE can offer cryptographic restrictions that limit access to individual browsing data used for online advertising, as well as potential benefits related to auditability that cannot exist

with on-device processing (see below, Policy Implications, Section C-I). For example, Google currently uses TEEs as part of Aggregation Service, a tool that generates an aggregated report of the effectiveness of an ad campaign conducted using Privacy Sandbox.²⁶ In order to mitigate risks of re-identification within any generated report, Aggregation Service injects statistical noise into the raw datasets prior to creating the aggregated report. In doing so, it processes browsing data within a TEE, allowing for the data to be decrypted in an isolated hardware environment, obfuscated with statistical noise, and released in the form of an aggregated report. In other words, the service offers a cryptographic safeguard against an unauthorized user (including the TEE provider) accessing or re-using the underlying data. Notably, Microsoft recently announced a similar confidential computing use case as part of the company's Edge browser.²⁷

While such uses of TEEs remain nascent in online advertising, the combination of benefits and real-world practicality suggests that it may become more prominent over time. For example, Microsoft has recently highlighted the role of Azure confidential computing in enabling “data clean rooms,” or secure environments in which multiple partners access and leverage shared data.²⁸ Although applicable across a range of industries, data clean rooms, many of which leverage TEEs, are a growing trend in online advertising that will likely continue in light of growing regulatory pressures.²⁹

C. Policy Considerations

Confidential computing remains a relatively new computing paradigm, and security researchers are actively researching its use in practice, discovering potential new threats, and improving the protocol.³⁰ For data protection practitioners and privacy professionals evaluating new tools, confidential computing offers potential benefits for accountability and transparency, and overall risk mitigation. At the same time, its potential roles with respect to restrictions on the “sale” or transfer of data, law enforcement access, cross-border data transfers, and data localization all remain largely untested and contingent on the specific details of management and configuration.

1. Accountability and Transparency

A unique promise of confidential computing is the separation of the attestation service from the underlying computation processes, which can provide certain **accountability and transparency benefits** to users of sensitive datasets, including through cryptographically guaranteed verification of the configuration of the TEE. At least in theory, this could include the ability to offer external monitoring or control of a dataset to a third-party entity, such as an auditor or regulator.

Specifically, **attestation server logs** can verify that server activities are aligned with established policies, including verifying that only acceptable metadata, data, and programs are provided entry into the TEE.³¹ In addition, attestation server logs can cryptographically authenticate the entities that request access to the TEE, including the authentication of each program or data point provided to the system. Particularly in the context of high-risk or sensitive data, the use of attestation server logs allows users of a TEE to rely on cryptographically guaranteed verification that the data processing is in fact occurring in an isolated environment and subject to the intended policies.

More recent iterations of confidential computing offer the ability to separate the attestation service from the TEE even further, for example by operating the attestation service independently by a trusted third party.³² Attestation services can be set up on independent machines, maintained, and operated independently by companies, government agencies, and research institutes, or any other trusted third party, even when the TEE or data storage infrastructure are run by a traditional third-party cloud provider. At least in theory, such arrangements could provide even greater monitoring and accountability capabilities to the independent third party.³³

2. Reducing Risks of Identifiability (Pseudonymization or De-Identification)

Most data protection laws either require or incentivize an organization processing personal data to take steps to reduce the data's identifiability or the likelihood that data could be linked, directly or indirectly, to a specific identified or identifiable person.³⁴ An analysis of legal re-identification risk typically includes the existence

and strength of “administrative safeguards,” such as internal controls over who can access data.³⁵ In aiming to establish greater administrative safeguards, an owner of a TEE could architect the attestation process to create greater controls and limitations around access to the data within those protected datasets.

However, while processing data within a TEE may reduce risks of unauthorized access, it does *not* lead to any technical reduction in the underlying data’s identifiability. In fact, processing data in a TEE can enable the exact opposite – i.e., the ability to process personally identifiable data in more secure and limited ways, rather than taking steps to de-identify it at a technical level. As a result, while an organization processing personal data in a TEE can have greater control and auditing capacity over the administrative access that is granted, to whom it is granted, and the kind of output that the system allows, the underlying data may still be highly identifiable. Any evaluation of whether information processed in a TEE is sufficiently legally de-identified would require a holistic evaluation of all relevant factors, including both the technical and administrative risks of re-identification and linkability with external data.

3. Access, Sharing, and Sale (“Do Not Sell”)

Emerging U.S. state laws establish a variety of obligations for respecting individual rights to control their personal information, most often in the form of the right to “opt-out” of a company’s “sale” or “sharing” (i.e., transfer) of the information. In most cases, “sale” is a legal term defined broadly enough to include any transfer of personal information to another entity for any accompanying benefit, whether financial or otherwise, outside of the context of a service provider or processor relationship.³⁶ As a result, a key legal and policy question for many US entities is whether information placed in a TEE may be considered to have been transferred or sold to another entity.

As a threshold matter, the transfer of data to **service providers** that process data *solely* on an entity’s behalf (for example, cloud storage providers) is typically *not* considered a sale or transfer.³⁷ Businesses generally rely on legally enforceable agreements (contracts) to ensure

that their service providers adhere to their data protection promises and obligations. In this context, confidential computing can offer additional layers of technical protections for a company to ensure that a recipient service provider is unable to access data in any functional manner and, therefore, unable to re-use or re-sell it — not solely as a matter of law, but also as a matter of technology. As a result, confidential computing can allow a business to have greater confidence in the relationship and in some cases demonstrate compliance to regulators.

In non-service provider relationships, such as shared data pools or co-ops in which multiple companies place information into a TEE and share insights derived from it, the question of whether the activity constitutes a legal sale or transfer is more complex. Generally speaking, the provision of access to personal data will be considered a sale when the transferring entity receives some form of benefit and the recipient is not restricted in its subsequent uses. For example, in a recent settlement, the California Attorney General alleged that the participation in (and benefiting from) a marketing co-op designed to support the contribution of personal information from unrelated businesses constituted a sale.³⁸

In contrast, a TEE could be engineered to allow for a similar co-op relationship, in which unrelated businesses pool information, but with greater limits established on the “output” to avoid any party having direct visibility into the information. If this were the case, the question of whether data has been “shared” would likely depend on the details of the system. For example, relevant factors would include: whether it allows for individualized information or 1:1 matching (e.g., “enhancing” first-party data); whether it allows for only aggregated inferences (similar to differential privacy solutions³⁹); and whether businesses are limited in the number of queries they can make. All of these questions would benefit from greater guidance from regulators and are sure to be a key topic of interest in years to come.

4. Law Enforcement Access

In general, the use of a trusted execution environment to protect sensitive or personal information can provide technical safeguards for the data and help ensure that attempts to access it are conducted through appropriate channels. For example, a service provider that is providing a TEE would likely not be capable of responding to a law enforcement request for data residing in the TEE without providing access through the designated attestation process. This can help prevent, for example, threats from a “rogue employee,” or other ways of granting access to unencrypted data in informal or unaccountable ways. Notably, however, the use of a TEE does not prevent the primary controller or owner of the data from responding to legitimate requests, either voluntarily or in response to a legal order.

In some cases, a TEE could also benefit a company’s internal processes for responding to law enforcement requests, insofar as it could be used to centralize the processing of information and limit access to a controlled number of entities and service providers.⁴⁰ For example, a non-U.S. company using a U.S.-based cloud service provider could use TEEs, among other safeguards, to ensure that law enforcement requests are made using appropriate legal channels, such as the CLOUD Act, and subject to review using established internal standards and processes.⁴¹

5. Risk Mitigation in Cross-Border Data Transfers

Organizations that operate globally must navigate legal restrictions that many jurisdictions place on transferring personal information to other jurisdictions (i.e., “cross-border data transfers”). This includes the European Union (EU), which requires that the personal data of Europeans only be processed in a jurisdiction with legal protections that meet or exceed EU data protection standards.⁴² In this context, the use of confidential computing would likely be a relevant, if not dispositive, factor in a transfer impact assessment (TIA) that accounts for security risks, especially risks related to insider threats and information-gathering conducted through informal or extra-judicial means.

In general, the GDPR authorizes cross-border transfers of personal data, in the absence of an adequacy decision from the European Commission, if the data exporter “has provided appropriate safeguards, and on the condition that enforceable data subject rights and effective legal remedies for data subjects are available.”⁴³ Among the options specified in Article 46, most organizations rely on standard contractual clauses (SCCs), or model contracts that serve to raise the level of data protection to GDPR-like standards.⁴⁴ However, in recent years, the European Court of Justice (CJEU) has added complexity to this process through its 2020 judgment in *Schrems II* that SCCs are not *alone* sufficient to address risks of access by public authorities in an importing country.⁴⁵ Following *Schrems II*, organizations that rely on SCCs to transfer data to a non-adequate jurisdiction must engage in additional analysis by conducting a **transfer impact assessment (TIA)** and considering “**supplementary measures**” to ensure sufficient data protection, depending on the laws of the importer country.

A **transfer impact assessment (TIA)** involves a fact-specific analysis of the circumstances of each data transfer, the legal obligations based on the data subject and origin jurisdictions, the safeguards the organizations put in place to protect the personal data, and the laws in the destination country.⁴⁶ This includes the practical risks related to a public authority in the importer country being capable of, and likely to succeed in, an attempt to gain access to the data through formal or informal means, including through backdoors, rogue employees, or mass surveillance mechanisms.⁴⁷ If the outcome of a TIA suggests that the organization’s safeguards for the data are not sufficient in light of the legal framework in the recipient country, an organization can adopt supplementary measures, including “technical, contractual and organizational measures,” to fill in the gaps and bring the level of data protection up to GDPR standards.⁴⁸

In this context, processing data within a TEE can potentially provide greater security safeguards for unencrypted data than a traditional cloud processing arrangement. In the latter, owners of data must trust that their cloud provider can sufficiently protect unencrypted data against

unauthorized access through side-channel attacks, hardware-based vulnerabilities, and insider threats.⁴⁹ In contrast, confidential computing can provide evidence in the attestation process that the cloud provider is complying with certain contractual obligations (e.g., not to inspect or use private data), that the technical safeguards are in place, and that they have not changed from the data owner's intended configuration. In other words, the hardware-based separation of data processing into a TEE creates an additional cryptographic safeguard against external hardware-based threats and hacking. In a verifiable way, it can also prevent the cloud provider itself, or any other insider or employee, from accessing the data without permission.⁵⁰

These factors, which may be relevant both to the outcome of a TIA and to the potential use of confidential computing solutions as supplementary measures, must be considered on a case-by-case basis, and their relevance will depend on the specific details of its implementation. For example, in addition to the physical location of data, a TIA might consider: the locations, number, and jurisdictions of any additional entities authorized to access data; the location and management of (i.e., keys to) the attestation service itself; and whether the organization is within scope of the CLOUD Act.⁵¹ Each of these factors would influence both the potential security risks, as well as the legal and technical ability for a cloud provider, the owner of the attestation service (if different), or an entity with permitted access, to comply with orders from public authorities that can compel access through the nation's legal system.

6. Data Localization and Other Restrictions

In addition to the limitations on cross-border data transfers discussed above, a growing number of jurisdictions have laws requiring that certain types of personal information be processed solely within the geographical boundaries of that nation.⁵² In some jurisdictions, including the United States, restrictions exist to prohibit or restrict the transfer of certain personal information, such as sensitive personal information, to a particular subset of adversarial foreign nations.⁵³

In most cases, if there is a role for confidential computing in complying with these requirements, it will depend on the specific legal terms and definitions of the relevant law, as well as the configuration of the TEE and the configuration, location, and management of the accompanying attestation service. Similar to the evolving legal questions on the "sale" of data (described above), the processing of protected information in a TEE with limited access could provide additional assurances that the data could not be easily exfiltrated to another location.

D. What's Ahead

By processing data in a hardware-based isolated environment and with cryptographic attestation processes, confidential computing solutions have the potential to provide unique data protection benefits for organizations that rely on processing sensitive or high-risk data. As with many privacy enhancing technologies, the usefulness, scale of impact, and potential regulatory compliance benefits of confidential computing will depend on the configuration and management of the TEE and its accompanying attestation service. As a result, privacy professionals should consider the full range of implementation details for any confidential computing system in balancing the relevant tradeoffs of cost, scale, and value for the relevant purposes.

NOTES

[illegible]

ENDNOTES

- 1 *What is confidential computing?* IBM, Jun. 4, 2024, <https://www.ibm.com/topics/confidential-computing> (last accessed Jun. 21, 2024); *Confidential Computing*, Fortinet, <https://www.fortinet.com/resources/cyberglossary/confidential-computing> (last accessed Jun. 21, 2024); *Confidential Computing*, Nat'l Inst. Standards and Tech. Comput. Sec. Res. Ctr., https://csrc.nist.gov/glossary/term/confidential_computing (last accessed Jun. 21, 2024); Frederick de Ryck, et al., *Enabling Sovereign Landing Zones with Confidential Computing*, Intel, available at https://cdrdv2-public.intel.com/783855/enabling-sovereign-landing-zones-with-confidential-computing-white-paper-1_1.pdf (last accessed Jun. 21, 2024); *About the Confidential Computing Consortium*, Confidential Computing Consortium, <https://confidentialcomputing.io/about/> (last accessed Jun. 21, 2024); *Trusted execution environments*, Information Commissioner's Office, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies/what-pets-are-there/trusted-execution-environments/> (last accessed Jun. 21, 2024).
- 2 The **central processing unit (CPU)** is the primary component of a computer responsible for interpreting and executing commands from the computer's other hardware and software. Strictly speaking, data in a TEE may reside in any memory address completely controlled and protected by the memory controller of the CPU. For purposes of policy, this paper does not distinguish between the CPU and computer's memory.
- 3 This can include, but is not limited to, the legal category of "sensitive personal data," referring to personal data under comprehensive data protection laws that require heightened protections due to the significant risks of financial, reputational, or other harm to individuals if misused or accessed by unauthorized entities. See, e.g., California Consumer Privacy Act, Cal Civ. Code § 1798.140(ae); Colorado Privacy Act, C.R.S. § 6-1-1303(24); Regulation (EU) 2016/679 Art. 9.
- 4 Cryptographic verification refers to the mathematical confirmation that the configuration of hardware, data, and software is within expected parameters based on the original specification of the TEE. See *Consumer Security Resource Center*, NIST, <https://csrc.nist.gov/glossary/term/verification> (last accessed Jun. 13, 2024); and Rahul Awati, *Cryptographic checksum*, TechTarget, <https://www.techtarget.com/searchsecurity/definition/cryptographic-checksum>, Dec 2021, (last accessed Apr. 19, 2024).
- 5 At least some TEE architectures have support for physical separation of the attestation services. See Michael Bartock, et al., *Hardware-Enabled Security: Enabling a Layered Approach to Platform Security for Cloud and Edge Computing Use Cases* (NIST IR 8320), NIST, May 2022, <https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8320.pdf> (last accessed Jun. 21, 2024). See also, e.g., Muhammed Usama Sardar et al., *Formal Specification and Verification of Architecturally-defined Attestation Mechanisms in Arm CCA and Intel TDX* (Nov 2023), Institute of Electrical and Electronics Engineers, available at https://www.researchgate.net/publication/375592777_Formal_Specification_and_Verification_of_Architecturally-defined_Attestation_Mechanisms_in_Arm_CCA_and_Intel_TDX.
- 6 Homomorphic encryption (HE) allows for computation on an encrypted dataset without intermediate decryption or knowledge of the decryption key, while secure multiparty computation (SMPC) is a protocol that allows multiple parties to jointly compute a function over their (encrypted) inputs while keeping those inputs private. See *Privacy-Enhancing Technologies (PETs)*, Information Commissioner's Office, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies> (last accessed Jun. 13, 2024).
- 7 On conventional computers, data can be encrypted when stored on a hard drive or in cloud storage ("at rest") and when sent across a network ("in transit"). But encryption's benefits generally end once an entity needs to process the data ("in use"). See *Strong Data Encryption Protects Everyone*, Future of Privacy Forum (Jul. 9, 2020), <https://fpf.org/encryption-infographic/>.
- 8 See *Confidential Computing – The Next Frontier in Data Security*, Confidential Computing Consortium, Oct. 2021, p. 13, https://confidentialcomputing.io/wp-content/uploads/sites/10/2023/03/Everest_Group_-_Confidential_Computing_-_The_Next_Frontier_in_Data_Security_-_2021-10-19.pdf (last accessed Jun. 21, 2024) ("[o]ver 75% of demand is driven by regulated industries like banking, finance, insurance, healthcare, life sciences, public sector, and defense in 2021." (emphasis added)).
- 9 In the United States, the Health Insurance Portability and Accountability Act ("HIPAA") governs the privacy and security of patient health records held by health providers and certain affiliates. HIPAA, 110 Stat. 1936; see also the HIPAA Privacy and Security Rules, 45 CFR Part 160; *Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule: A Cybersecurity Resource Guide*, Nat'l Inst. of Standards and Tech., <https://csrc.nist.gov/pubs/sp/800/66/r2/final> (last accessed Apr. 18, 2024).
- 10 Tajabadi M, Grabenhenrich L, Ribeiro A, Leyer M, Heider D. Sharing Data With Shared Benefits: Artificial Intelligence Perspective. *J Med Internet Res*. 2023 Aug 29;25:e47540. doi: 10.2196/47540. PMID: 37642995; PMCID: PMC10498316, available at <https://www.jmir.org/2023/1/e47540> (last accessed Jun. 20, 2024).
- 11 See Davenport T, Kalakota R., *The potential for artificial intelligence in healthcare*. *Future Healthc J*. 2019 Jun;6(2):94-98. doi: 10.7861/futurehosp.6-2-94. PMID: 31363513; PMCID: PMC6616181, available at <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6616181/> (last accessed Jun. 20, 2024).
- 12 See 45 C.F.R. §§ 164.502(d)(2) ("Uses and disclosures of protected health information"); 45 C.F.R. § 164.512(i) ("Uses and disclosures for which an authorization or opportunity to agree or object is not required"); How Can Covered Entities Use and Disclose Protected Health Information for Research and Comply with the Privacy Rule?, US Dep't of Health & Human Svcs, https://privacyruleandresearch.nih.gov/pr_08.asp (last accessed Jun. 17, 2024).
- 13 See, e.g., Im E, Kim H, Lee H, Jiang X, Kim JH., *Exploring the tradeoff between data privacy and utility with a clinical data analysis use case*. *BMC Med Inform Decis Mak*. 2024 May 30;24(1):147. doi: 10.1186/s12911-024-02545-9. PMID: 38816848; PMCID: PMC11137882., available at <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11137882/> (last accessed Jun. 20, 2024).
- 14 Press Release, UCSF, Fortanix, Intel, and Microsoft Azure Utilize Privacy-Preserving Analytics to Accelerate AI in Health Care (Oct. 8, 2020), available at <https://www.ucsf.edu/news/2020/10/418736/ucsf-fortanix-intel-and-microsoft-azure-utilize-privacy-preserving-analytics> (last accessed Jun. 21, 2024).
- 15 For example, the Gramm Leach Bliley Act (GLBA) requires that financial services organizations meet transparency obligations with respect to the privacy of customer financial information, usually in the forms of annual privacy notices or providing individuals with the ability to object to sharing of personal information with third parties. 15 U.S.C. §§ 6801-809.
- 16 See, e.g., Bank Secrecy Act ("BSA") of 1970. 31 U.S.C. § 5311; See also 31 CFR. § 1010.311 (Currency Transaction Report); 31 C.F.R. § 1010.320 (Suspicious Activity Report); 31 C.F.R. § 1010.410 (BSA recordkeeping requirements).
- 17 A recent KPMG case study found that US banks spend \$25 billion on AML programs annually, with the average bank spending \$48 million. *Combating Financial Crime*, KPMG, <https://kpmg.com/mc/en/home/insights/2019/03/combating-financial-crime-fs.html> (last accessed Jun. 21, 2024).

- 18 A recent study found that roughly 70 percent of financial services institutions use AI models to combat (among other forms of fraud) money laundering. *Seven in 10 Financial Institutions Use AI and ML to Combat Fraud*, PYMTS, May 26, 2024, <https://www.pymnts.com/news/security-and-risk/2024/seven-in-10-financial-institutions-use-ai-and-ml-to-combat-fraud/> (last accessed Jun. 20, 2024). The US Department of Treasury is aware of this trend and is actively seeking information on how institutions are using AI as part of compliance programs that include, among other things, AML. See *Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector*, Dep't of Treasury, Jun. 14, 2024, <https://www.govinfo.gov/content/pkg/FR-2024-06-12/pdf/2024-12336.pdf> (last accessed Jun. 21, 2024).
- 19 For examples of AML vendors currently offering confidential computing solutions, see, e.g., Anubhav Gupta, Kiran Kannur, *Preventing Money Laundering in a Confidential Computing Way*, Fortanix, Feb. 1, 2023, <https://www.fortanix.com/blog/preventing-money-laundering-in-a-confidential-computing-way> (last accessed May 8, 2024); and *Improved Money Laundering Detection with Predictive Analytics*, C3.ai, <https://c3.ai/products/c3-ai-anti-money-laundering/> (last accessed Mar. 18, 2024).
- 20 See *supra* note 8, *Confidential Computing – The Next Frontier in Data Security*, Confidential Computing Consortium.
- 21 See, e.g., Tyler Pichach, et al., *Combat financial crime with AI and advanced technology from Microsoft*, Jun. 5, 2023, Microsoft, <https://www.microsoft.com/en-us/industry/blog/financial-services/2023/06/05/combating-financial-crime-with-ai-and-advanced-technology-from-microsoft/> (last accessed Jun. 21, 2024); see also *Harnessing AI in the fight against payments fraud*, Swift, May 30, 2024, <https://www.swift.com/news-events/news/harnessing-ai-fight-against-payments-fraud> (last accessed Jun. 21, 2024). Notably, confidential computing can also facilitate the generation of synthetic datasets, which are particularly important in the financial industry. See R. Searle, et al., "Secure Implementation of Artificial Intelligence Applications for Anti-Money Laundering using Confidential Computing," 2022 IEEE International Conference on Big Data (Big Data), Osaka, Japan, 2022, pp. 3092-3098, doi: 10.1109/BigData55660.2022.10021108, available at <https://ieeexplore.ieee.org/abstract/document/10021108> (last accessed Jun. 20, 2024).
- 22 S. Gray, "Advertising in the Age of Data Protection: Background for a Proposed Risk-Utility Framework for Novel Advertising Solutions," Future of Privacy Forum, Apr. 2024. Available: <https://fpf.org/wp-content/uploads/2024/04/FPF-Proposed-Advertising-Risk-Utility-Framework-Apr-2024-v1.0.pdf> (last accessed Jun. 20, 2024).
- 23 See, e.g., Priyanka Chatterjee, et al., *Overview of FLEDGE Services*, GitHub, Apr. 22, 2024, https://github.com/privacysandbox/protected-auction-services-docs/blob/main/trusted_services_overview.md (last accessed Jun. 21, 2024) (Google's proposed specification); Erik Anderson, *Ad Selection API details*, GitHub, Mar. 5, 2024, <https://github.com/WICG/privacy-preserving-ads/blob/main/API%20Details.md> (last accessed Jun. 21, 2024) (Microsoft's proposed specification).
- 24 In many contexts, on-device processing can offer greater computational speeds insofar as it eliminates the transfer of data from the user to an external server, particularly for people with powerful devices and strong network connections. However, on-device processing remains challenging for auction-based advertising, which is typically performed at scale by ad exchanges with low latency, high throughput network connections to deliver ads within the first 100 milliseconds of a page load. See Maciej Zawadziński, Mike Sweeney, *How Does Real-Time Bidding (RTB) Work?*, Clearcode, Apr. 16, 2024, <https://clearcode.cc/blog/real-time-bidding/> (last visited Jun. 21, 2024); see also Amanda Martin, *Privacy Sandbox's Latency Issues Will Cost Publishers*, Ad Exchanger, Mar. 4, 2024, <https://www.adexchanger.com/the-sell-side/privacy-sandbox-latency-issues-will-cost-publishers/> (last accessed May 9, 2024).
- 25 See Apple, *Apple Intelligence | Privacy*, YouTube (Jun. 10, 2024), <https://youtu.be/546ufMY7488?si=kQ8sCG4e4vZeO40d> (last accessed Jun. 18, 2024) (Apple's 2024 Worldwide Developer Conference (WWDC) Keynote). See also, Apple Security Research, *Private Cloud Compute: A New Frontier for AI Privacy in the Cloud*, Apple Security Research Blog (Jun. 10, 2024), <https://security.apple.com/blog/private-cloud-compute/> (last accessed Jun. 18, 2024).
- 26 Aggregation Service overview, Google, <https://developers.google.com/privacy-sandbox/relevance/aggregation-service> (last accessed Apr. 18, 2024).
- 27 *New Privacy-Preserving Ads API coming to Microsoft Edge*, Microsoft Windows Blogs, Mar. 5, 2024, <https://blogs.windows.com/msedgedev/2024/03/05/new-privacy-preserving-ads-api/> (last accessed Mar. 20, 2024). Microsoft has posted further details on Github, <https://github.com/WICG/privacy-preserving-ads> (last accessed Apr. 18, 2024).
- 28 Ananya Garg, Microsoft Azure Architecture Blog, *Enabling data clean rooms with confidential computing* (Jan. 3, 2024), <https://techcommunity.microsoft.com/t5/azure-architecture-blog/enabling-data-clean-rooms-with-confidential-computing/ba-p/4020538> (last visited Jun. 21, 2024).
- 29 E.g., Habu, *Microsoft Azure And Habu*, <https://habu.com/partners/microsoft-azure/> (last accessed Jun. 18, 2024). See generally, Apoorv Durga, MarTech, *Data clean rooms: A beginner's guide* (July 12, 2023), <https://martech.org/data-clean-rooms-a-beginners-guide/> (last accessed Jun. 21, 2024); Ryan Barwick, MarketingBrew, *Clean rooms, explained: How they became the buzziest tool in ad tech* (Nov. 22, 2021), <https://www.marketingbrew.com/stories/2021/11/22/clean-rooms-explained-how-they-became-the-buzziest-tool-in-ad-tech> (last accessed Jun. 18, 2024).
- 30 See, e.g., L. H. Newman, WIRED, *Intel Let Google Cloud Hack Its New Secure Chips and Found 10 Bugs* (Apr. 24, 2023), <https://www.wired.com/story/intel-google-cloud-chip-security/> (last accessed Feb. 20, 2024). For a survey of security properties and challenges, see Muñoz, Antonio, Ruben Ríos, Rodrigo Román, and Javier López, *A Survey on the (in)Security of Trusted Execution Environments*, Computers & Security 129 (Jun. 1, 2023): 103180. <https://doi.org/10.1016/j.cose.2023.103180>. For a more practitioner-oriented overview, see Jauernig, Patrick, Ahmad-Reza Sadeghi, and Emmanuel Stäpf, *Trusted Execution Environments: Properties, Applications, and Challenges*, IEEE Security & Privacy 18, no. 2 (Mar. 2020): 56–60. <https://doi.org/10.1109/MSEC.2019.2947124>.
- 31 Server logs are automatically-generated text documents that contain lists of all activities of a specific server in a given period of time. See also *Server Log Files in a Nutshell*, graylog (Feb. 26, 2021), <https://graylog.org/post/server-log-files-in-a-nutshell/> (last accessed Mar. 18, 2024).
- 32 Intel's SGX, for example, includes support for remote attestation services. See Strengthen Enclave Trust with Attestation, Intel, <https://www.intel.com/content/www/us/en/developer/tools/software-guard/extensions/attestation-services.html> (last accessed Jun. 21, 2024). For an annotated use of remote attestation on Intel's SGX, see Attestation, <https://ssl-gatech.github.io/sgx101/pages/attestation.html> (last accessed Jun. 21, 2024).
- 33 For example, a government agency concerned about the confidentiality and integrity of personnel files could set acceptable policies in a TEE, be able to rely on cryptographic verification that these policies were used by the hardware to complete the computation, and potentially even be able to comprehensively monitor every event within the server and TEE. See, e.g., Devlin Barrett et al., *U.S. Suspects Hackers in China Breached About 4 Million People's Records, Officials Say*, Wall St. J., Jun. 5, 2015, <https://www.wsj.com/articles/u-s-suspects-hackers-in-china-behind-government-data-breach-sources-say-1433451888> (last accessed Jun. 21, 2024).

- 34 Most if not all state privacy laws exempt de-identified data from the definitions of personal data. See, e.g., California Consumer Privacy Act (as amended), Cal. Civ. Code § 1798.140(v)(3); Colorado Privacy Act, CRS § 6-1-1303(17)(b); Connecticut Data Privacy Act, Conn. Gen. Stat. § 42-515.
- 35 Federal Trade Commission, Protecting Consumer Privacy in an Era of Rapid Change (Mar. 2012), <https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf>. See also Lore Leitner, et al., *Anonymisation through separation: what recent cases teach us about the EU'S Anonymisation standards*, https://www.linkedin.com/posts/gabe-maldoff-74755646_anonymisation-through-separation-activity-7153432687096029184-yFxD/ (last accessed Jun. 21, 2024).
- 36 California Consumer Privacy Act (“CCPA”). Cal. Civ. Code 1798.140(ad), 1798.140(ah).
- 37 See, e.g., Cal. Civ. Code 1798.140(ag)(1) (requiring “service providers” and the businesses for whom they process personal information to have a written contract that prohibits service providers from selling or sharing the personal information).
- 38 *California v. DoorDash, Inc.*, Compl. at 2, available at <https://oag.ca.gov/system/files/attachments/press-docs/DoorDash%20Complaint.pdf> (last accessed Apr. 23, 2024), see also Press Release, Att’y Gen. Bonta Announces Settlement with DoorDash, Investigation Finds Company Violated Multiple Consumer Privacy Laws (Feb. 21, 2024), <https://oag.ca.gov/news/press-releases/attorney-general-bonta-announces-settlement-doordash-investigation-finds-company> (last accessed Jun. 21, 2024).
- 39 See *What Is Differential Privacy?*, Inst. of Elec. and Elec. Eng’r, <https://digitalprivacy.ieee.org/publications/topics/what-is-differential-privacy> (last accessed Jun. 21, 2024).
- 40 For example, following the United States Supreme Court’s decision in *Dobbs v. Jackson Women’s Health Organization*, 597 U.S. 215 (2022), overturning *Roe v. Wade*, 410 U.S. 113 (1973), companies have faced heightened scrutiny over their internal management of law enforcement requests for health-related information. See Press Release, Wyden, Jayapal and Jacobs Inquiry Finds Pharmacies Fail to Protect the Privacy of Americans’ Medical Records; HHS Must Update Health Privacy Rules (Dec. 12, 2024), <https://www.finance.senate.gov/chairmans-news/wyden-jayapal-and-jacobs-inquiry-finds-pharmacies-fail-to-protect-the-privacy-of-americans-medical-records-hhs-must-update-health-privacy-rules> (last accessed Jun. 21, 2024).
- 41 The CLOUD Act both authorizes extraterritorial reach for judicial process issued under certain U.S. legal authorities and provides a framework for other countries to be approved to issue legal process directly to companies in the United States’ jurisdiction. Pub. L. No. 115-141 (Mar. 23, 2018), available at <https://www.congress.gov/115/plaws/publ141/PLAW-115publ141.pdf> (last accessed June 21, 2024).
- 42 The GDPR creates obligations for transfers of information to countries not deemed “adequate,” a term that refers to a determination by the European Commission that non-European Economic Area (EEA) nations provide levels of data protection, rights, and access restrictions commensurate with EU standards. Art 45; see also GDPR Recital 108. At the time of writing, the European Commission has adopted adequacy decisions for 15 countries. *Adequacy Decisions*, European Commission, https://commission.europa.eu/law/law-topic/data-protection/international-dimension-data-protection/adequacy-decisions_en (last accessed Jun. 21, 2024).
- 43 GDPR Art 46(1); see also Recital 108.
- 44 See Nigel Cory, Ellysse Dick and Daniel Castro, *The Role and Value of Standard Contractual Clauses in EU-U.S. Digital Trade*, Information Technology & Innovation Foundation, Dec. 2020, <https://itif.org/publications/2020/12/17/role-and-value-standard-contractual-clauses-eu-us-digital-trade/>; Lee Matheson, *Not-So-Standard Clauses: Examining Three Regional Contractual Frameworks for International Data Transfers*, Future of Privacy Forum (Mar. 2023), <https://fpf.org/blog/fpf-report-not-so-standard-clauses-an-examination-of-three-regional-contractual-frameworks-for-international-data-transfers/> (last accessed June 21, 2024); *EU Standard Contractual Clauses*, European Commission, https://commission.europa.eu/law/law-topic/data-protection/international-dimension-data-protection/standard-contractual-clauses-scc_en (last accessed Mar. 5, 2024).
- 45 *Data Protection Commissioner v. Facebook Ireland LTD, Maximilian Schrems*, C-311-18 (holding that the EU-US Privacy Shield Framework failed to provide for an “essentially equivalent” level of protection for outbound EU personal data to the US) (“*Schrems II*”). The successor to these transfer agreements is the EU-US Data Privacy Framework, which the European Commission has found to offer adequate protection.
- 46 See *International Data Transfers*, European Data Protection Board, https://www.edpb.europa.eu/sme-data-protection-guide/international-data-transfers_en (last accessed Jun. 21, 2024).
- 47 See David Rosenthal, *Transfer Impact Assessment Templates*, Int’l Ass’n of Privacy Prof’l (IAPP), Sept. 1, 2021, <https://iapp.org/resources/article/transfer-impact-assessment-templates/> (last accessed Jun. 21, 2024).
- 48 *Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data (Version 2.0)*, European Data Protection Board, adopted on 18 Jun. 2021, https://www.edpb.europa.eu/our-work-tools/our-documents/recommendations/recommendations-012020-measures-supplement-transfer_en (last accessed Mar. 18, 2024).
- 49 Recent hardware vulnerabilities have shed a light on this particular issue. See, e.g., *Unpatchable vulnerability in Apple chip leaks secret encryption keys*, ArsTechnica, Mar. 21, 2024, <https://arstechnica.com/security/2024/03/hackers-can-extract-secret-encryption-keys-from-apples-mac-chips/> (last accessed Apr. 23, 2024) (discussing critical security vulnerabilities with Apple’s M1 chips); “*Meltdown*” and “*Spectre*”: Every modern processor has unfixable security flaws, ArsTechnica, Jan. 1, 2018, <https://arstechnica.com/gadgets/2018/01/meltdown-and-spectre-every-modern-processor-has-unfixable-security-flaws/> (last accessed Apr. 23, 2024) (discussing vulnerabilities with Intel chips).
- 50 For these reasons, confidential computing is sometimes referred to as “zero trust,” although the phrase is not always strictly accurate. A true “zero trust” architecture involves a wider set of considerations, including network management, ID requirements, and guarantees of data availability. Nat’l Inst. of Standards & Tech., Zero Trust Architecture, Aug. 2020, <https://csrc.nist.gov/pubs/sp/800/207/final> (last accessed Apr. 23, 2024).
- 51 See *supra*, note 47 (Transfer Impact Assessment templates).
- 52 See *Financial Data Localization: Conflicts and Consequences*, Future of Privacy Forum, Dec. 7, 2017, <https://fpf.org/blog/financial-data-localization-info-graphic-conflicts-and-consequences/> (last accessed Jun. 21, 2024); Hunter Dorwart, *Demystifying Data Localization in China: A Practical Guide*, Future of Privacy Forum, Feb. 2022, <https://fpf.org/wp-content/uploads/2022/02/Demystifying-Data-Localization-Report.pdf> (last accessed Jun. 21, 2024).
- 53 See Protecting Americans’ Data from Foreign Adversaries Act of 2024, H.R. 815, 118th Cong. § 1(a)(9) (2024), available at <https://www.congress.gov/bill/118th-congress/house-bill/815/text> (last accessed Jun. 20, 2024); Exec. Order No. 14117, 28 C.F.R. § 202 (2024), available at <https://www.federalregister.gov/documents/2024/03/01/2024-04573/preventing-access-to-americans-bulk-sensitive-personal-data-and-united-states-government-related> (last accessed Jun. 20, 2024).

