



The Case for Confidential Computing

Delivering Business Value
Through Protected, Confidential
Data Processing

Suzanne Ambiel

July 2024

The Case for Confidential Computing

CONFIDENTIAL COMPUTING CONSORTIUM

Unites vendors, cloud providers, and developers to accelerate Trusted Execution Environment (TEE) technology and standards adoption.




CONFIDENTIAL COMPUTING

Enhances data security during use by performing computations within a hardware-based, attested Trusted Execution Environment.




CONFIDENTIAL COMPUTING

Secures data in use by performing computations in a hardware-based, attested Trusted Execution Environment.



CONFIDENTIAL COMPUTING USE CASE



Deliver Business Value: improve lives, discover new drugs, catch thieves, & protect data from tampering.

CONFIDENTIAL COMPUTING USE CASE

Enable Public Cloud Adoption: Securely deploy sensitive workloads in the public cloud, addressing data privacy & security concerns.



CONFIDENTIAL COMPUTING USE CASE



Expand into new markets: Enhance GDPR & HIPAA compliance by encrypting sensitive data and preserving border-restricted access.

CONFIDENTIAL COMPUTING USE CASE

Deploy Confidential AI: Secure models and data throughout their lifecycle, ensuring safe training, inferencing, & data privacy.



CONFIDENTIAL COMPUTING USE CASE



Protect AI Investments: Secure models and data, & control access "by third parties" or public cloud providers.

CONFIDENTIAL COMPUTING USE CASE

Build Better Campaigns, Engage New Customers: Deploy targeted campaigns, protect personal information, and **securely** share data to understand customer behavior.



CONFIDENTIAL COMPUTING USE CASE

Identify Suspicious Transactions: Meet & global banking regulatory requirements for customer privacy while pooling data to identify money laundering.



CONFIDENTIAL COMPUTING USE CASE

Save Lives, Improve Patient Outcomes: Detect disease faster, develop innovative treatment plans by securely analyzing patient data, while preserving privacy & meeting regulatory requirements.



CONFIDENTIAL COMPUTING USE CASE

Develop New Drugs: Protect patient data, ensure regulatory compliance, & maintain ethical standards while collaborating with research institutions.



Contents

- Introduction 4
- Confidential data clean rooms deliver efficient marketing outcomes..... 6
- Protected data processing enables public cloud adoption..... 8
- Confidential computing powers confidential AI for enterprises 10
- Attestation and secure enclaves help deliver global data protection and compliance..... 12
- Industry-specific use cases 13
- Healthcare and medical research: privacy-preserving data aggregation leads to better outcomes 15
- Conclusion..... 17
- Methodology 18
- Acknowledgments..... 18
- About the author 18
- Appendix..... 18

Introduction

Regardless of the industry, business, or company size, today's business leaders must confront the data challenge: how to safely and securely collect, curate, and use company and customer data without compromising privacy, all the while maintaining strict compliance with regulations and industry-specific laws. And with the rise of artificial intelligence (AI) and AI / machine learning (ML) and the immense data requirements to fuel those initiatives, the challenge grows even more urgent.

Storage and network encryption technologies provide means to protect data at rest and in transit, but data in use remains vulnerable. When data is unencrypted for processing, it's susceptible to taint, tampering, and other forms of compromise.

Missed opportunities: Cloud computing and multiparty data collaboration

For many enterprises, their data is simply too valuable to introduce new risks. The specter of data breaches or tampering leads these companies to limit the distribution and use of their data. For some, it's forgoing the opportunity to leverage cloud computing; for others, it means retaining the data in company-designated silos to reduce risk. These data isolation strategies limit a company's ability to put their data to work to help them grow their business, improve operational efficiency, or expand their customer base.

Cross-industry or third-party collaboration is another potential missed opportunity. The ability to securely access third-party data to augment internal proprietary data can help a company detect unseen patterns or develop new products. Third-party data is also essential to train large AI models, validate research, or discover market opportunities. For companies in industries

such as healthcare or financial services, stringent data protection and privacy regulations coupled with potential penalties for noncompliance force many organizations to opt out or limit participation in these new opportunities.

Protect data and code with confidential computing

Lacking a method to share and process data while keeping it protected impacts an enterprise's ability to achieve its goals. Plenty is at stake: competitive advantage, new product development, reduced operational costs, increased market share, and better return on investment (ROI), among others. However, embracing a strategy with confidential computing can unlock these possibilities.


Confidential computing provides a protected Trusted Execution Environment (TEE) to process data out of view of unauthorized software, administrators, or other tenants on shared infrastructure. The data in use is secure; unauthorized parties can't change it and can't look at it. That protection extends to the application as well. The TEE helps ensure the integrity of the application, protecting it from tampering or theft.

The power of confidential computing

With confidential computing, your data remains your data—no unauthorized users can see it, taint it, or steal it. There is data protection for all three modes: at rest, in transit, and in use. A confidential computing platform offers a protected foundation for innovation and collaboration across industries. Organizations use confidential computing to unlock data silos, enable public cloud adoption, and meet global data privacy compliance requirements.

- **CONFIDENTIAL COMPUTING:** a hardware-based technology solution that protects data and code “in use” through the creation of hardware-based, attested Trusted Execution Environments. Confidential computing aims to provide end-to-end security for sensitive data, ensuring it remains protected from unauthorized entities throughout its lifecycle
- **TRUSTED EXECUTION ENVIRONMENT (TEE):** an environment that provides data integrity, data confidentiality, and code integrity such that unauthorized entities cannot alter, remove, or view the data or code while in use in the TEE.
- **ENCLAVES:** An enclave is a secure, isolated portion of memory within a TEE where sensitive data and code can be processed securely. TEEs are implemented at the processor level to provide a comprehensive secure environment. Enclaves are specific implementations within a TEE, created to provide additional isolation for sensitive code and data.
- **ATTESTATION:** Attestation verifies, through a cryptographically signed proof, the integrity and authenticity of the software and hardware components within a Trusted Execution Environment (TEE). This process ensures that the TEE is running genuine, unaltered code before it is trusted to handle sensitive data.
- **DATA LIFECYCLE:** Data exists in one of three modes: it’s either at rest (in storage), in motion (on the network), or in use (being processed). Today’s encryption technologies tackle the first two modes; confidential computing provides hardware-backed protection for the last mode, data in use, by creating a secure processing environment.

For more details about confidential computing, including architectures and technical implementations, read “**A Technical Analysis of Confidential Computing**” and “**Confidential Computing: Hardware-Based Trusted Execution for Applications and Data**”.



In healthcare and pharmaceutical enterprises, the ability to engage in multiparty data collaboration safely and securely enables them to explore new treatments, develop new drugs, and improve patient outcomes more effectively than before. For financial services enterprises, confidential computing not only fuels innovation but also permits secure cross-industry collaboration to detect money laundering and other forms of financial crime.

In each of the following use case briefs, you’ll learn how confidential computing can help organizations tap into the value of enterprise data safely, securely, and with full confidence. By choosing a confidential computing solution and putting data to work, organizations can grow their business; deliver better service to clients, patients, and customers; tackle tough data challenges; and meet regulatory compliance demands head-on.

- Confidential Data Clean Rooms Uncover Efficient Marketing Outcomes
- Secure Data Processing Enables Public Cloud Adoption
- Attestation and Secure Enclaves Strengthens Global Data Protection and Compliance
- Confidential Computing Boosts Confidential AI for Enterprises
- Multiparty Data Collaboration Helps Combat Money Laundering
- Privacy Preserving Data Aggregation Leads to Better Healthcare Outcomes

To learn more about confidential computing, visit the Confidential Computing Consortium (CCC), and take advantage of their resources, white papers, technical documentation, and more.

Confidential data clean rooms deliver efficient marketing outcomes

INDUSTRY: All; Marketing Strategists; Digital Media and Advertising

CHALLENGE: Insufficient first-party data and elimination of third-party cookies render marketing and ad targeting less accurate, reducing effectiveness and lowering the return on ad spending.

SOLUTION: Confidential computing-enabled data clean rooms can securely enrich first-party data with third-party data.

BENEFITS:

1. Improve marketing and advertising program results to deliver better ROI.
2. Expand market share with enhanced prospect and customer profiles.
3. Reduce the risk of data compromise and subsequent regulatory and legal penalties.

“Third-party cookies have underpinned Internet advertising for decades. With its deprecation, publishers and brands must rethink entirely new ways to reach audiences in a privacy-first world. With Decentriq and confidential computing, Goldbach enables a trust layer, allowing brands to onboard their first-party data without compliance risks.”¹

—JOCHEN WITTE, CTO, GOLDBACH GROUP AG

¹ <https://www.decentriq.com/products/media-advertising>

Overview

John Wanamaker, a retail and marketing pioneer, reportedly declared: “Half the money I spend on advertising is wasted; the trouble is I don’t know which half.” And while nearly 125 years have passed since he uttered those words, they still ring true. Accurately measuring marketing and advertising ROI across digital and traditional media challenges even the best marketers and agencies.

First-party data collected from interactions and transactions with actual customers provides organizations with some insights into behavior and buying patterns. However, it only covers one relationship and lacks insights into how the customer interacts with other brands, whether market competitors or merely complementary product providers.

In the past, third-party cookies captured that additional data and helped to paint a more complete picture of the customer. This combined data—first and third-party data—served as an essential foundation for crafting personalized experiences, driving targeted marketing efforts and advertising campaigns, and creating new product and service offerings.

Today various regulations protect first-party data, often called personally identifiable information (PII), elevating the need for stringent data security and protection. It also presents barriers to sharing and third-party collaboration. The E.U.’s hallmark GDPR is the most well-known and widely cited. Adding to a marketer’s challenge, the curtailing of traditional third-party cookies limits the amount of customer persona data available to organizations.

This leaves organizations with a potentially reduced dataset and limited view of customer behavior, once again wondering which half of their advertising and digital marketing efforts are most effective.

Solution: Data clean rooms enabled by confidential computing

Confidential computing provides a solution to this challenge. By creating a secure environment where data can be connected and processed without revealing its contents, organizations can enrich their first-party data with additional third-party data.

These environments, called data clean rooms, allow multiple data providers to securely connect first-party data to gain insights from a broader, previously unavailable set of data. Users can create “lookalike” audience profiles without the risk of data compromise or regulatory challenges. With access to a larger data pool, companies can use these insights to design more efficient marketing and advertising programs, helping to fill the prospect pipeline with ready buyers. Media companies can also monetize their data without compromising privacy or running afoul of international and regional privacy laws.

“Confidential computing and data clean rooms present us with a great chance to exchange information in a legal, efficient way. This allows us to do better media planning and advertising, targeting more accurately, and, in the end, driving more revenue while complying with all data protection regulations.”²

—ZHAO WANG, HEAD OF DATA TECHNOLOGY, RINGIER

Learn more

- **Enhancing Privacy and Data Protection With Confidential Computing** (Zonar)
- **Is Your Organization Ready for a “Cookie-Less” World?** (Decentriq)
- **Improving Ad Efficiency Through First-Party Data Collaboration** (Decentriq)
- **Reducing Marketing Costs With Tailored Lookalike Audiences** (Decentriq)
- **Sharing Data Across Organizations Without Sharing** (Enclave)

Protected data processing enables public cloud adoption

INDUSTRY: All; Business leaders: CIO, CISO, CTO

CHALLENGE: Cloud computing promises scale and favorable economics, but the risk of exposing proprietary data to cloud providers or other tenants limits cloud adoption.

SOLUTION: Deploy your applications and data in the public cloud protected by confidential VMs and enclaves.

BENEFITS:

1. Cloud economics, scalability, and flexibility
2. Hardware-enforced data and application security, privacy, and control

Overview

In nearly every industry, you'll encounter widespread adoption of public cloud. Its advantages are widely acknowledged: flexibility, scalability, cost management, portability, and beyond. However, despite its popularity, it's essential to recognize that it's not without risk. Processing proprietary or highly regulated data in the cloud, on shared infrastructure, introduces the risk of data or applications being compromised, stolen, or breached.

According to Sebastian Gajek, CTO, Enclave, "Public clouds, while offering many benefits, often raise security questions due to shared infrastructure and potential vulnerabilities in a multi-tenant environment."³ Your protection is one of social trust coupled with contractual penalties for cloud provider failure—

³ Five Reasons Why a Confidential Cloud Trumps a Public Cloud

⁴ Five Reasons Why a Confidential Cloud Trumps a Public Cloud

not enough to convince risk-averse CIOs or CTOs to move their most sensitive workloads to the cloud. For many enterprises, the only reliable solution is to keep the workload and the data on-premise, which is, in many instances, a more costly and unwieldy choice.

*"The enclaves are kept separate from other applications and services running on the same system, making sure that the attack surface is kept to a minimum. In today's world, where businesses are thinking "cloud-first," this technology is of great importance. Any workloads that were previously not considered to be uploaded to the cloud, because of security and compliance concerns, can now take advantage of these services."*⁴

—SEBASTIAN GAJEK, CTO, ENCLAVE

Protect applications and data with confidential computing

Today's major cloud providers address this challenge by offering confidential computing solutions. Industry analysts consider

confidential computing a viable solution to making cloud computing safer for even the most sensitive applications and their data. Confidential computing provides a secure hardware-based processing environment, called a trusted execution environment (TEE), or enclave, that provides additional protection for sensitive data and applications.

In a TEE, applications and their data are isolated from all other users and operators. No unauthorized party, not even the cloud provider, can access the data or the code—they can't add to it, remove it, or alter it in any way. This solution provides hardware-enforced data confidentiality and application integrity, strengthening the security and privacy of the data, and the applications as well. Nosy neighbors, hackers, and other potential security threats are significantly altered and minimized. For organizations with highly sensitive or regulated data, this additional protection and level of assurance can open the door to a new operating model in the cloud.

By deploying applications and data to the cloud in a confidential computing VM or enclave, organizations can enjoy all the benefits of cloud computing while safeguarding data and applications from third-party access in a secure environment, protected against compromise.

By adopting confidential computing-based public cloud solutions, enterprises can enjoy the many benefits of cloud computing even with their most sensitive or regulated data.

Learn more

- **Five Reasons Why a Confidential Cloud Trumps a Public Cloud** (Enclave)
- **Confidential Computing** (Google)
- **Azure Confidential Computing—Protect Data in Use** (Azure)
- **O.C. Tanner Protects Customer Data Across Hybrid, MultiCloud, and Multi-Site Environments** (Fortanix)

Confidential computing powers confidential AI for enterprises

INDUSTRY: All; Business leaders: CTO, CISO, CPO, CIOs

CHALLENGE: AI models and associated data are at risk of theft, tampering, or compliance violations if left unprotected.

SOLUTION: Deploy and operate AI models using confidential computing solutions.

BENEFITS:

1. Safeguards investment in proprietary data, algorithms, and models
2. Leverages the economics and scale of cloud computing
3. Lowers costs, faster time to market

Overview

The transformational potential of AI is top-of-mind for every organization today, and the desire to invest, invent, and innovate spans all industries. According to Menlo Ventures, today's enterprises are investing more than \$75 billion annually in AI-related technologies.⁵ The bulk of that spending occurs in product development and engineering departments—the heartbeat of an organization's innovation and revenue machine. To protect their business and their investment, the proprietary data and algorithms that form the foundation for these initiatives must be kept secure from competitors and potential tampering.

⁵ Source: Menlo Ventures, *The State of Generative AI in the Enterprise, 2023*

Striking a balance: Security versus access

Proper data governance and security are critical to enterprise success and the deployment of AI solutions. Securely connecting internal data stores to either third-party foundation models such as OpenAI or Anthropic or cloud provider-hosted models poses barriers to adoption; it's a risky move for many business leaders. Operating AI / ML models exclusively on-premise not only consumes extraordinary amounts of internal compute resources but also constrains the outcomes. With access limited to only first-party data, the results may be incomplete or suboptimal.

But using cloud computing introduces significant risks. Building an AI-enabled product or service is expensive, time-consuming, and relies on proprietary data and custom algorithms, making it a tempting target for compromise. The global nature of cloud providers also introduces a complex web of compliance challenges, as data protection laws differ across regions. Cloud services require the transfer of sensitive information to external servers, which can raise questions about data ownership, jurisdiction, and control—still more risk.

More than just AI—it's confidential AI

With AI, it's not only the data that needs protection. The model, the weights, the algorithms, and the outcomes—they all need to be kept private and secure. By deploying enterprise AI in a confidential computing environment, protection extends from the data to the model and application. Only authorized users have access; confidential computing prevents the cloud provider from viewing, tampering, or stealing the data. The TEE safeguards the model as well, preventing tampering or taint, helping to assure the integrity of the model outcomes.

According to Glen Otero, VP of Scientific Computing, Translational Genomics Research Institute, “Confidential computing that can provide protection of algorithms as well as the data while computing will be the default requirement for data privacy and the future of AI modeling over the next five to ten years.”⁶

For healthcare providers or financial services enterprises, AI poses unique challenges due to the amount of heavily regulated data consumed by the model or application. Since confidential computing offers a trusted execution environment that protects the model and data from unauthorized access, these new AI-enabled initiatives can deliver results while meeting privacy, security, and ethical concerns.

In addition, a confidential AI environment makes secure, compliant use of third-party data possible, preserving confidentiality at all times. Accessing and leveraging third-party LLMs while preserving company and customer private data is possible with confidential computing.

With these assurances of privacy and security, organizations can more confidently adopt and invest in AI, knowing that their proprietary information remains confidential, intact, and immune to breaches and tampering. Confidential computing environments safeguard their investments.

Learn more

- **How Jamworks Protects Confidentiality While Integrating AI Advantages** (IBM)
- **TGen Secures Genome Data for Richer Healthcare AI Models With Fortanix Confidential Computing** (Fortanix)
- **Privacy-Preserving Data-Collaboration Methods That Accelerate Healthcare Innovation** (Intel)

“Confidential computing platforms—[CCP] allow us to reduce the cycle time to validate an algorithm in half. It also cuts the costs almost in half. Those kinds of savings allow us to train, validate, and bring to market generalizable algorithms much faster. And, it will only get faster and less costly as the technology and processes underlying CCP mature.”⁷

—MARYBETH CHALK, CO-FOUNDER AND CHIEF COMMERCIAL OFFICER, BEEKEEPERAI, INC.

⁶ TGen Secures Genome Data for Richer Healthcare AI Models With Fortanix Confidential Computing

⁷ Accelerating Development of Clinical AI Algorithms

Confidential computing helps deliver global data protection

“Confidential computing ... was the clear choice to secure our customers’ data and ensure GDPR and Schrems II compliance for European needs.”⁸

—GORDON WADDELL, SENIOR VICE PRESIDENT OF SOFTWARE DEVELOPMENT, ZONAR

INDUSTRY: All; Business Leaders: CTO, CIO, CISO/CPO

CHALLENGE: Compliance requirements for all types of data continue to grow; costs for security and privacy technologies, audits, and reporting consume ever more of the IT budget.

SOLUTION: Adopt confidential computing solutions to deliver assured, attested, and auditable results cost-effectively and consistently.

BENEFITS:

1. Lower costs through streamlined compliance
2. Increased flexibility
3. Attested and auditable data security and privacy

Overview

Data compliance regulations designed to protect personal and private information continue to become more prevalent, complex, and difficult to address. With local, regional, and national versions of regulation applying across all types of consumer data, it’s a dizzying matrix of compliance challenges.

From GDPR to HIPAA, if your business handles data from patients, customers, prospects, and third-party providers—even those beyond your geographic borders—these rules apply. In addition, these regulations are generally applicable regardless of enterprise size or location. It’s also a safe bet to always default to the most stringent set of rules for ALL your data.

All this adds up to a very challenging data management and processing landscape, especially with cloud computing strategies in place. Some data requires geo-fenced cloud environments and datasets to ensure compliance. For some enterprises, the problem is too complex, the penalties too steep, and the compliance costs too high, forcing them to limit their business, or forgo promising partnerships.

Trusted execution and attestation: A powerful combination

Placing regulated data in a confidential computing environment addresses many of the challenges posed by regulatory agencies. In a confidential computing solution, the TEE or enclave isolates and protects data from software and individuals not explicitly authorized to access the TEE. Further, attestation provides a cryptographic confirmation that the TEE is genuine and operating correctly. In this scenario, the TEE provides protection and isolation, while the attestation provides assurance.

As organizations pursue new business partnerships or expand into new markets, safeguarding data and providing a secure means to process the data becomes a key to success. With the protections provided by a confidential computing platform, organizations can meet the growing regulatory challenge of data management while helping CIOs and CISOs rest a little easier.

Learn more

- [Zonar Helps Ensure GDPR and Schrems II Compliance by Enhancing Privacy and Data Protection](#) (Google, Zonar)

⁸ Zonar Helps Ensure GDPR and Schrems II Compliance by Enhancing Privacy and Data Protection

FINANCIAL SERVICES

Multiparty data collaboration helps combat money laundering

INDUSTRY: Financial Services; Business Leaders: CTO, CLO, CIO, CISO/CPO, GM, Product Manager

CHALLENGE: Financial institutions must safeguard customer data while also complying with strict anti-money laundering regulations to prevent the receipt of proceeds of crime.

SOLUTION: Securely pool transaction data from multiple institutions enabled by confidential computing solutions.

BENEFITS:

1. Save time and money with faster detection.
2. Support compliance demands with attested results.
3. Reduce cost with efficient multiparty collaboration while increasing success.
4. Accelerate fraud detection through pattern recognition, algorithm development, and predictive modeling.

“Fraud in payments is a major concern for our clients—in the order of magnitude of billions of dollars every year. Preventing financial crime is a problem that we cannot solve individually as a bank; we need to solve it collectively.”¹⁰

—ISABEL SCHMIDT, CO-HEAD OF GLOBAL PAYMENTS PRODUCTS AT BNY MELLON, A MEMBER OF THE SWIFT NETWORK

Overview

Multinational smugglers. Drug syndicates. Sanctioned companies and countries. They all conduct financial transactions and rely on public banks and financial institutions in their operations. Regardless of the customer, financial institutions must contend with dueling mandates. On the one hand, they must abide by strict privacy rules, safeguarding customer data, while at the same time, the growing anti-money laundering and “know your customer” laws force institutions to examine their customer data very cautiously. This tension renders swift detection of transaction anomalies, especially between institutions, difficult. It’s this siloed nature of financial data that enables bad actors to mask their illicit activities.

By moving their funds swiftly and covertly between institutions, criminals can exploit the global financial network and slip through the cracks unnoticed. According to the United Nations Office on Drugs and Crime, up to \$2 trillion of illicit funds are laundered through global financial networks every year.⁹

To stem the tide, financial institutions face a growing list of regulations, reporting requirements, and laws, and failure to comply invites massive penalties. Today, financial institutions must implement strict due diligence and operating procedures, with auditable records to prove compliance, to ensure they are not unwittingly processing payments in restricted sectors, for sanctioned individuals or entities, or for criminal means.

⁹ <https://www.unodc.org/unodc/en/money-laundering/overview.html>

¹⁰ Swift Innovates With Azure Confidential Computing To Help Secure Global Financial Transactions

The price for noncompliance is steep, with penalties frequently reaching into the hundreds of millions to billions of dollars. However, money laundering is tough to catch, especially if you're fighting the battle alone. A single company's dataset is not sufficient for anomalies that might indicate money laundering, but aggregating that data across institutions might reveal illicit transactions. However, sharing this information, especially with third parties and competitors, also introduces risk. A breach could mean not only stiff financial penalties but damaging front-page reputational harm as well.

Leveraging confidential computing for enhanced data sharing and risk mitigation

But just as bad actors exploit the global network, financial institutions need to leverage that same network to combat money laundering. Pattern recognition improves exponentially when the dataset increases—the richer the data, the better the results. By pooling data and collaborating with other financial institutions or stakeholders securely, detection can happen sooner, saving everyone time and money. It's a task especially suited for confidential computing. The two methods used for collaboration are multiparty data sharing and a federated learning model.

Multiparty data sharing in a confidential computing environment ensures that all participant data is kept private and secure, eliminating the threat of exposure or contamination. With confidential computing, financial institutions can access data clean rooms—a controlled, secure, centralized environment that pools data from multiple parties without exposing raw identifiable information to ensure confidentiality and regulatory compliance.

The federated method is often used when the datasets are too big, too sensitive, or too regulated to move off-premise or share beyond company-private cloud instances. In this scenario, the institutions' internal systems preserve all data locally. Institutions

securely share models and applications for local processing. All parties benefit from the combined results that come from the shared common models or applications.

In both scenarios, the protections offered by a confidential computing environment give participating institutions the confidence that their data will be safeguarded against exposure or compromise.

Learn More

- **Swift Innovates With Azure Confidential Computing** To Help Secure Global Financial Transactions (Microsoft)
- **MonetaGO Detects Duplicate Financing Fraud** (Google / AMD)
- **Applications for Digital Fraud Defense** (FiVerity)

Privacy-preserving data aggregation leads to better outcomes

INDUSTRY: Healthcare Providers, Medical Research, Pharmaceuticals

CHALLENGE: Patient data is one of the most highly regulated data types; collaboration and data sharing are largely prohibited, yet offer the promise of better outcomes and lower costs.

SOLUTION: Data aggregation enabled by confidential computing gives healthcare providers the ability to improve patient or research outcomes while safeguarding patient privacy.

BENEFITS:

1. Improved patient care and outcomes
2. Lower costs through more efficient diagnosis and treatment
3. Enhanced clinician and physician learning
4. Larger, more diverse clinical trials

Overview

Patient data is one of the trickiest types of data around. The collection, use, storage, access, and processing are highly controlled, with good reason, and regulated by a variety of laws, regulations, and regulators, with HIPAA in the U.S. being the most well-known. Penalties for violating patient privacy are steep—reputational and financial—rendering stewards of healthcare data reluctant to share with anyone for any reason. And yet, sharing this information can yield benefits that accrue directly to the patient and their providers.

To deliver better healthcare and patient outcomes, aggregating patient data across providers is key. Aggregated patient data can enable more efficient, coordinated care across institutions and hospitals or help facilitate specialized medical care. It can also help physicians, especially those in rural areas with a smaller patient population, to deliver better patient outcomes. With a larger pool of data at their fingertips, physicians are better able to quickly compare symptoms, diagnose disease, and recommend advanced treatment plans.

For clinical researchers and pharmaceutical companies, the development and approval of new novel drugs and treatments require data. Access to high-quality, large-volume datasets is essential to success. Clinical trials and advanced medical research require real-world data—including electronic health records—that's often distributed across multiple healthcare providers, clinicians, and countries. These groups and jurisdictions severely restrict or even prohibit access, making the barriers to access oftentimes appear insurmountable. Lacking the means to safeguard patient data while sharing information makes clinical trials even more challenging.

Alternative methods

Alternatives to data aggregation and collaboration include anonymization and cleansing—eliminating any PII. But these alternatives are cumbersome and inefficient; with each institution hewing to different data naming, storage, and curation standards, the barriers to normalized data are high. Once anonymized, that data often can't yield results in real time and may deliver ambiguous results.

Establishing multiparty data sharing and collaboration agreements is another option. Yet establishing and maintaining compliant,

secure raw patient data-sharing agreements between academic institutions, hospitals, care providers, or individual physicians is time-consuming and riddled with regulatory and legal challenges.

As Chris Gough, General Manager of Health and Life Science at Intel, notes, “Collection and analysis of protected medical data is fraught with challenges, including the sensitivity of and distributed nature of the data across multiple healthcare entities and systems, which requires strong protections for patient privacy.”¹¹

Sharing is caring

By sharing data, medical professionals can shift healthcare from a one-size-fits-all model to a more precise, targeted intervention, improving outcomes and quality of life for patients in a cost-effective way. For patients with specialized or complex needs, aggregated data enables more efficient cross-institution, multi-provider treatment programs. As clinical trials grow more complex, building systems that can address patient privacy and data security while enabling collaboration is essential to discovering new treatments.

“Leidos understands the technological challenges associated with clinical information systems and the need to create trusted computing environments to securely share information. [Confidential computing] ... gives us the foundation needed to build an ecosystem of partners that can confidently share data privately and securely while still meeting the stringent compliance regulations in the space.”¹²

—ERIKA KILLIAN, FDA PORTFOLIO DIRECTOR, LEIDOS

¹¹ Intel Teams With Leidos, Fortanix To Accelerate Clinical Trials

¹² Ibid, Intel Teams With Leidos, Fortanix To Accelerate Clinical Trials

Despite the many potential benefits of patient data aggregation and collaboration, the high threshold of patient privacy and data protection means these benefits remain out of reach. An inefficient and non-compliant method to share data restricts healthcare providers to a dataset that yields limited information to aid in diagnosis, treatment, and care. It’s certainly not enough to develop new drugs.

Confidential computing can deliver a secure multiparty collaboration or federated learning solution ideally suited for healthcare providers to share and aggregate patient data efficiently and safely across institutions and providers.

Medical professionals can deliver better patient outcomes, knowing that their patient’s privacy is protected and their data is safe. Patients get healthier sooner, with less time in treatment and office visits. Coordinated care plans across providers enabled by confidential computing reduce friction in all types of healthcare settings. Physicians and clinicians benefit as well through exposure to additional cases and alternative treatment plans. For researchers, pooled data and cross-institution collaboration can lead to innovative treatments and accelerate therapies to market.

Learn more

- **Intel Teams With Leidos, Fortanix To Accelerate Clinical Trials** (Intel, Leidos, Fortanix)
- **Privacy-Preserving Data-Collaboration Methods That Accelerate Healthcare Innovation** (Intel, BeekeeperAI)
- **Healthcare Data Collaboration by the Numbers: Balancing Progress and Privacy** (Decentriq)
- **Decentriq To Facilitate Analysis of Data From Over 1M Cardiovascular Disease Patients** (Decentriq)

Conclusion

Confidential computing leads to better outcomes

Improve lives, discover new drugs, catch thieves, and protect data from tampering? It's all possible with confidential computing.

When organizations adopt confidential computing, they demonstrate a strong, comprehensive commitment to data security—in any state. Implementing solutions with hardware-based TEEs alongside traditional data encryption technologies protects an organization's data not only when it's at rest and in transit but also when it's in use. These solutions also assure code integrity. When a TEE protects the application, tampering, tainting, or stealing the code is not possible.

The benefits span industries, enabling new business, increasing compliance, and streamlining new product development and deployment.

- For enterprise marketers, securely enriching first-party data with third-party data can increase marketing campaign efficiency, reduce unwanted and unproductive campaign activities, and improve customer satisfaction.
- For CIOs and CISOs, the possibility of custom and confidential AI becomes a reality with an assured, private, and secure environment to build, operate, and tune their proprietary models.
- For global financial services institutions, the ability to pool their customer and transaction data with other potentially competitor financial institutions and third parties could help prevent fraud, improve lending practices, and accelerate innovation.

- For healthcare providers, access to large, diverse datasets well beyond their own collection could yield faster and more precise disease diagnosis and optimal treatment plans, reducing costs while improving patient outcomes and preserving patient data privacy.
- For pharmaceutical companies and medical researchers, data sharing opens the door to larger and potentially more effective clinical trials, enabling the delivery of new, innovative treatments sooner.

The stories included here draw upon the experience of industry leaders and innovators such as Intel, Google, Decentriq, Microsoft, and BeeKeeperAI. Organizations such as TikTok continue to experiment with new uses for confidential computing, all in the pursuit of building trust with their users and customers. With four different use cases in the works, including privacy-preserving data sharing, TikTok looks to confidential computing to support its goals. Both Vini Jaiswal, Manager, Developer Advocacy & Open Source at TikTok, and Mingshen Sun, Research Scientist at TikTok, agree that "Privacy is one of our [TikTok's] top initiatives because we work with 170 million users, and we want to make sure that our platform is reliable and trustworthy for those who consume the platform." Confidential computing is one of the many ways TikTok continues to deliver on its promise of trust, privacy, and security to all its users and customers.

But the adoption of confidential computing doesn't need to stop here. To hear Hushmesh, a U.S.-based software firm, describe the potential of confidential computing is to learn about a future Internet that's based on confidential computing to implement a universal "zero trust" model that starts at the chip level. In that world, privacy and security would cease to be an added feature, something implemented through a diverse array of software and

layered on top of the infrastructure. Instead, data security and privacy would be innate to the infrastructure— automatic and built-in, at the chip—not added or managed by a human. According to Manu Fontaine, Hushmesh CEO, “We believe that confidential computing enables the creation of a global infrastructure and an information space that completely automates end-to-end cryptographic security across everything and everyone.” A big vision for confidential computing indeed.

Methodology

This research was conducted between February and May 2024 by the Linux Foundation Research Team and sponsored by the Confidential Computing Consortium. The project team conducted a series of primary research interviews, with secondary research supplemented by Consortia member contributions. Additional insights were gathered from in-depth literature reviews, market research reports, analyst reports, vendor websites, media articles, and vendor-sponsored whitepapers.

Acknowledgments

Thanks to Mike Ferron-Jones for his invaluable contributions, from guiding outline development and identifying study participants to his thoughtful and thorough review and commentary during the report development.

About the author

Suzanne Ambiel is a technology industry veteran; her experience ranges from fault-tolerant computing to VMs, containers, and open source software. During her tenure at VMware (now Broadcom), Suzanne led OSPO strategy and communications in addition to managing brand research and insights. As a former Linux Foundation Board member and LF Research Advisory

Board member, Suzanne personally observed the essential and expanding role open source plays in today’s digital world.

Appendix

Interview participants and contributors

Vini Jaiswal, *TikTok*

Mingshen Sun, *TikTok*

Dayeol Lee, *TikTok*

Manu Fontaine, *Hushmesh*

Marcus Hartwig, *Google*

Stanislav Nikolskiy, *GenoMex*

Malini Bhandaru, *Intel*

Mike Ferron-Jones, *Intel*

Mona Vij, *Intel*

Paul O’Neill, *Intel*

Nikolas Molyndris, *Decentriq*

Andrew Knox, *Decentriq*

Emily Fox, *Red Hat*

Mike Bursell, *Confidential Computing Consortium*

Additional resources published by the Confidential Computing Consortium

Common Terminology for Confidential Computing

Confidential Computing—The Next Frontier in Data Security

A Technical Analysis of Confidential Computing v1.2

Confidential Computing: Hardware-Based Trusted Execution for Applications and Data

Confidential Computing in Financial Services: Use Cases for Data Security

Confidential Computing Consortium YouTube Channel



Sponsored by the Linux Foundation, the CCC is a community focused on projects securing data in use using hardware-based TEEs and accelerating the adoption of confidential computing through open collaboration. The CCC brings together hardware vendors, cloud providers, and software developers to foster the adoption of TEE technologies and standards.



Founded in 2021, **Linux Foundation Research** explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks, and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations the world over.



Copyright © 2024 **The Linux Foundation**

This report is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International Public License**.

To reference this work, please cite as follows: Suzanne Ambiel, "The Case for Confidential Computing: Delivering Business Value Through Protected, Confidential Data Processing," The Linux Foundation, July 2024.