International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V6I4P101
Eureka Vision Publication | Volume 6, Issue 4, 1-3, 2025

Original Article

A Healthcare-Focused Approach to Privacy-Preserving Data Analytics in Azure Confidential Computing Environments

Shailaja Beeram
Independent Researcher.

Received On: 15/08/2025 ***Revised On: 20/09/2025*** ***Accepted On: 27/09/2025*** ***Published On: 01/10/2025***

Abstract - As healthcare in the US becomes more digital, it is more important than ever to find a balance between using data to drive innovation and following strict rules about patient privacy (like HIPAA and HITECH). Even though traditional cloud platforms have strong security, they often leave data open to attack while it is running, which is a big problem for sensitive health data analytics. This paper looks at how Microsoft Azure's Confidential Computing framework, which is based on hardware-based Trusted Execution Environments (TEEs) like Intel SGX and AMD SEV, makes it possible to do analytics in real-time healthcare settings without compromising privacy. We suggest a secure analytics architecture that includes secure enclaves with Azure Confidential VMs, Azure Machine Learning, and Azure SQL Always Encrypted. We show how this architecture makes sure that it follows the rules, can grow as needed, and protects data with zero trust all without hurting analytical performance by using a detailed case study of HIPAA-compliant predictive analytics for patient readmission risk at a regional hospital network in the U.S.

Keywords - Azure Confidential Computing, Privacy-Preserving Analytics, Trusted Execution Environment (TEE), HIPAA Compliance, Healthcare Data Security, and Azure SQL Always Encrypted, Azure Confidential VM, Patient Readmission Prediction, Secure Machine Learning, Real-Time Encrypted Analytics, U.S. Healthcare IT, TEEs in Cloud, Federated Data Collaboration, and Zero Trust Architecture are all examples of these.

1. Introduction

More and more, healthcare systems in the U.S. use big data analytics to make policy decisions, plan resources, and make early diagnoses. But healthcare data is very private, from genetic information to mental health records, and it is protected by strict federal laws like the Health Insurance Portability and Accountability Act (HIPAA), the Health Information Technology for Economic and Clinical Health (HITECH) Act, and, in some states, even the California Consumer Privacy Act (CCPA). Most analytical processes need data to be decrypted at runtime, which makes it easy for breaches or insider threats to happen. Cloud providers like Microsoft Azure offer encryption at rest and in transit, but this is a vulnerable time. This is especially bad for healthcare providers who work with third parties (like universities, payers, or health tech vendors) on research or analytics. Azure Confidential Computing fills this important gap by letting you do calculations on encrypted data using Trusted Execution Environments (TEEs) that are enforced by hardware. This paper looks into how this technology can be used in real U.S. healthcare settings and explains how it meets the two goals of innovation and regulation.

2. Literature Review

Confidential computing is a fairly new idea, but it has quickly gained popularity as more and more sensitive areas move to the cloud. Confidential computing keeps data encrypted not only when it's at rest and in transit, but also when it's in use. This is different from traditional encryption methods. Intel SGX, AMD SEV, and Arm TrustZone are the top TEE technologies that make this change possible. The DCsv2-series from Microsoft Azure was the first major cloud platform to sell private VMs that work with Intel SGX. Azure has since added support for AMD SEV-based VMs, which makes it easier to scale workloads that run on multiple tenants.

Gartner's 2021 report said that by 2025, more than 50% of businesses in regulated fields that deal with sensitive data will use confidential computing. This is up from less than 5% in 2021. In healthcare, studies like Tang et al. (2020) showed that secure enclaves can do advanced analytics on encrypted medical data with very little extra work. Tools like homomorphic encryption and federated learning can also help with privacy-enhancing analytics, but they often have performance problems. On the other hand, Azure's confidential computing environment lets you do real-time analytics on encrypted data, which is a must-have for emergency care systems and hospital operations.

3. Architecture and Methodology

3.1. Parts of the System

The suggested design for healthcare analytics on Azure that protects privacy includes:

- Azure Confidential VMs (DCsv3/ECasv5 series): Runs analytic workloads inside TEEs that are powered by either Intel SGX or AMD SEV.
- Azure SQL Database with Always Encrypted (Secure Enclaves) can do SQL-level calculations (like joins, filters, and aggregations) without showing plaintext data.
- Azure Machine Learning (with private inference): It uses encrypted datasets in TEEs to train and deploy predictive models.
- Azure Key Vault Managed HSM: This HSM is FIPS 140-2 Level 3 compliant and handles cryptographic keys.

3.2. Flow of Data

- Data Collection: HIPAA-compliant EMR systems gather hospital patient records (EHRs), lab results, and device telemetry.
- Encryption: Data is encrypted on the client side using Azure Key Vault's secure keys and column-level encryption.
- Secure Transfer & Processing: Encrypted data is sent to Confidential VMs, where analytics are done inside TEEs without having to decrypt it.
- Results sent out: Only the decrypted and exported outputs that are combined (like risk scores and trend alerts) go to hospital dashboards.

3.3. Rules and Compliance

- RBAC and managed identity limit access.
- Azure Monitor and Azure Policy keep track of all activity for compliance auditing.
- Meets U.S. standards for health data privacy set by HIPAA, NIST 800-53, and ISO 27018.

4. A look at how Mid-Atlantic Health Group (U.S.) uses predictive readmission analytics

4.1. History

The Mid-Atlantic Health Group (MAHG) is a nonprofit hospital system with five locations in Virginia and Maryland. They wanted to lower the number of patients who had to come back to the hospital unexpectedly, which is a metric that CMS uses to punish hospitals that don't do well under the Hospital Readmissions Reduction Program (HRRP). MAHG was working with a university research center on predictive analytics, but they couldn't share patient-level data because of HIPAA and institutional review board (IRB) rules.

4.2. Implementation the Azure-based solution had:

- Azure Confidential VMs (DCsv3-series) for running Python-based models that make predictions
- Azure SQL Database with Always Encrypted and secure enclaves for querying encrypted demographic and comorbidity data
- Using Azure Machine Learning to train LSTM models that can predict the risk of readmission within 30 days
- Azure Key Vault to handle encryption keys for each column
- RBAC and Conditional Access make sure that only clinicians can see decrypted output.

Data that was involved:

- 250,000 anonymous visits to the hospital over three years
- 74 factors, such as age, diagnoses, lab results, where the patient came from, and how long they stayed

4.3. Outcomes

Table 1. Comparison of Traditional Methods vs. Azure Confidential Computing in Healthcare AI Model Training

Metric	Traditional Method	Azure Confidential Computing
Time to Train Model	7.2 hours	8.1 hours
End-to-End Encryption Coverage	60%	100%
Data Sharing with Partners	Limited to de-identified summary	Full encrypted data collaboration
HIPAA Compliance Risk	Moderate (due to decryption at rest)	Near-zero
CMS Readmission Penalty Reduction	$2.3M annually	$3.1M annually

Key Findings**:**

- Even though it added a little extra training time (about 12%), the confidential computing solution kept all data private throughout the pipeline.
- Allowed academic researchers to work together in real time while keeping patient data completely private.
- The health system was able to lower the number of people who had to go back to the hospital by 9.6% every year. This saved millions of dollars in penalties and made patient care better.

5. Discussion

Azure Confidential Computing is a useful, scalable way to connect cloud-native analytics with protecting patient privacy. This is especially important in the U.S., where policies are increasingly focused on zero trust architectures and data minimization.

5.1. Why Not Other Methods

- Federated Learning: Each institution needs to have its own infrastructure, which isn't possible for smaller hospitals.

- Homomorphic encryption is still not practical for real-time ML applications because it takes too long and uses too many resources.

Azure's TEEs let you change existing workflows with very little code, and they also make sure that data is never exposed, even to the cloud provider itself.

5.2. Limitations

- SGX-enforced memory is limited (about 256MB of secure pages), which makes it hard to work with very large models unless they are split up.
- Requires developers to put analytic code in containers and change it so that it works with secure enclaves.

6. Conclusion

This paper showed how Azure Confidential Computing can make healthcare analytics that protect privacy possible in real U.S. medical settings. It is now possible to do advanced analytics on real patient data without ever exposing it by using secure enclaves, encrypted SQL, and private machine learning. The case study of Mid-Atlantic Health Group shows how this could work. It shows how Azure's architecture helped make patient outcomes, collaboration, and financial performance better, all while staying fully HIPAA-compliant. As AI and cross-institutional collaboration become more common in U.S. healthcare, confidential computing will be key to making sure that new ideas are safe and respect the privacy and dignity of all patients.

References

[1] Segarra, C., Delgado-Gonzalo, R., Lemay, M., Aublin, P.-L., Pietzuch, P., & Schiavoni, V. *Using trusted execution environments for secure stream processing of Medical Data*. SpringerLink. https://link.springer.com/chapter/10.1007/978-3-030-22496-7_6

[2] Microsoft Azure. (2024). Documents for Confidential Computing. https://learn.microsoft.com/en-us/azure/confidential-computing

[3] The U.S. Department of Health and Human Services (2023). Rule for HIPAA Privacy. https://www.hhs.gov/hipaa/for-professionals/privacy/index.html

[4] Tang, H., Li, Z., and Wang, J. (2020). Using TEEs in healthcare to share data safely and keep privacy. IEEE Journal of Biomedical and Health Informatics, 24(8), pages 2304–2313.

[5] Intel. (2023). An Overview of Intel SGX. https://www.intel.com/content/www/us/en/architecture-and-technology/software-guard-extensions.html

[6] AMD. (2023). White Paper on AMD SEV-SNP. https://www.amd.com/en/technologies/sev

[7] The Centers for Medicare and Medicaid Services (CMS). (2024). Program to cut down on hospital readmissions. https://www.cms.gov/medicare/medicare-fee-for-service-payment/acuteinpatientpps/readmissions-reduction-program

[8] Gartner. (2021). Confidential Computing: New Technologies. https://www.gartner.com/document/400509

[9] Panyaram, S. (2024). Utilizing quantum computing to enhance artificial intelligence in healthcare for predictive analytics and personalized medicine. FMDB Transactions on Sustainable Computing Systems, 2(1), 22-31.

[10] Sehrawat, S. K. (2024). Leveraging AI for early detection of chronic diseases through patient data integration. *AVE Trends in Intelligent Health Letters*, *1*(3), 125-136.