2025 7th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS) | 979-8-3315-5982-3/25/$31.00 ©2025 IEEE | DOI: 10.1109/BRAINS67003.2025.11302941

Demo: MedTech Chain 2.0 Integration of Homomorphic Encryption into a Decentralized Platform for Medical Device Data Research*

Yavor Pachedzhiev†, Andaç Durmaz†, Ege Yarar,
Ivan Markov, Sinan Önen, Roland Kromes and Zekeriya Erkin
Delft University of Technology Departments of Intelligent Systems and Software Technology
Delft, The Netherlands
{y.s.pachedzhiev-1, a.durmaz, e.yarar, i.k.markov, s.onen.}@student.tudelft.nl, {r.g.kromes, z.erkin}@tudelft.nl

***Abstract*—Hospitals produce vast amounts of medical device data, making their protection and analysis crucial in Cyber Threat Intelligence (CTI) settings. MedTech Chain 1.0 allowed cybersecurity researchers to run queries for data analytics. Even if the platform applied differential privacy, data storage was in plaintext and certain analysis capabilities were still missing. To address these limitations, we propose MedTech Chain 2.0, a platform that integrates homomorphic encryption, and enhanced mechanisms for encryption key management, and expanded query support. These improvements strengthen data protection while enabling deeper insights, advancing in cybersecurity and CTI research.**

***Index Terms*—blockchain, homomorphic encryption, healthcare, security, privacy, networked medical devices, Hyperledger Fabric**

I. INTRODUCTION

Modern healthcare providers generate large volumes of medical device data through their digital systems. This data includes information such as firmware versions, device types, operational status, and more. Although these data do not contain direct patient information, they remain highly sensitive because they can reveal vulnerabilities within hospital infrastructures, such as outdated devices that allow potential cyberattack targets [1]. Due to the crucial nature of this information, a secure and privacy-preserving platform is necessary to protect medical device data while still enabling analysis to identify weaknesses.

To address this challenge, an initial version of such a platform, MedTech Chain 1.0, was developed as part of the EU Septon project. It was a platform on which cybersecurity researchers could run queries on medical device data via a blockchain-based architecture to perform data analytics within hospital systems. [2] However, this first version had key limitations, especially in terms of privacy and encryption. The medical device data stored on the blockchain was plaintext, making potentially sensitive information accessible to peers on the blockchain. In addition, only three query operations were available: count, grouped count, and average. These operations were limited in scope and provided minimal insight to cybersecurity researchers. To mitigate these problems, the development of MedTech Chain 2.0 was essential.

MedTech Chain 2.0 is a blockchain-based platform designed to enable secure, decentralized, and privacy-preserving analysis of medical device data. Building upon the first version of the platform, [2] MedTech Chain 2.0 aims to solve the query capability problem and privacy concerns regarding how networked medical device data are stored in the system. To address these problems, MedTech Chain 2.0 integrates homomorphic encryption, allowing computations such as summation and multiplication to be performed directly on encrypted medical device data. [3] This eliminates the need to decrypt sensitive medical device data during specific data analytics operations, preventing their exposure to any party in the system. As a result, cybersecurity researchers can continue to obtain insightful information on hospital systems without accessing identifiable device-level data. The platform now also supports new queries—including sum, standard deviation, and linear regression—that leverage homomorphic properties. In addition, a histogram and a unique count query were introduced to provide more information to researchers, but they require decryption due to their calculations. To ensure that homomorphic encryption of medical device data, secure management of keys, and decryption of query results are performed reliably, the system requires a component trusted by all entities. Therefore, a Trusted Third Party (TTP) is introduced, which is responsible for executing encryption and decryption operations, managing key rotation, and maintaining secure key versioning. Additionally, a controlled noise is added to the decrypted query results with ϵ-differential privacy to protect sensitive information while allowing analysis on aggregate results. [4]

II. FEATURES OF MEDTECH CHAIN 2.0

A. Participants and responsibilities

Hospitals. Hospitals collect device metadata (e.g., model, firmware, usage) and submit it for encryption. They never

† These authors contributed equally to this work.

* This work is supported by the European Union's Horizon Europe research and innovation programme under grant agreement No. 101094901 (SEPTON).

write plaintext to the ledger.

Trusted Third Party (TTP). The TTP manages the cryptography. It generates and stores all secret keys, encrypts hospital submissions, and decrypts the final aggregate result. In this paper, an *aggregate result* means a summary value computed over many device records (e.g., a total count or an average), rather than any individual record of a query before it is returned. Keys never leave the TTP.

MedTech Chain Organization. This party operates the Hyperledger Fabric network and chaincode, manages access control, and applies ϵ-differential privacy to each result before release. It also provides the web API that hospitals use to upload data and that researchers use to submit queries. Administrators may initiate key rotations, which are then handled by the TTP.

Researchers. Cybersecurity researchers use the web application to submit aggregate queries with optional filters. They do not see device-level records. they only receive differentially private aggregates.

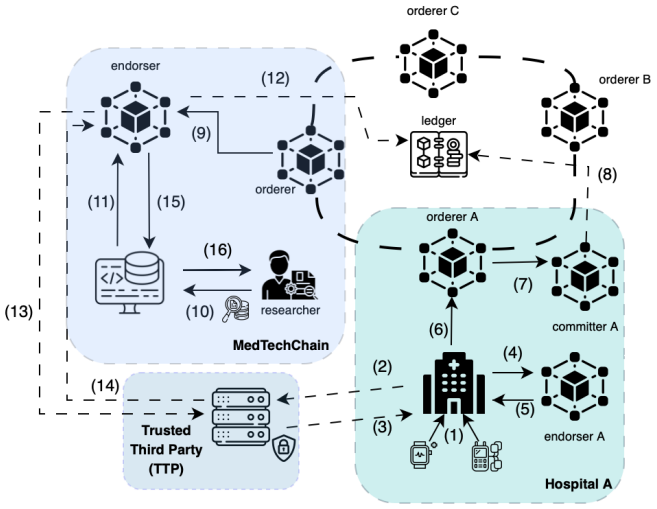

Fig. 1. Hyperledger Fabric–based MedTech Chain Architechture

B. Data Lifecycle

1) Ingestion: Hospital → Ledger (Steps 1–9): Hospitals collect metadata from their medical devices. In the current setup, two device categories are considered: portable and wearable, but the model can be extended to other types. Once data is gathered (step 1), the hospital requests encryption of the relevant fields from the TTP (step 2). The TTP encrypts the values using the active scheme (Paillier or BFV) and attaches metadata identifying the key in use (step 3). The hospital service then prepares a blockchain transaction with the encrypted data and submits it to the network. Endorsing peers validate the transaction (steps 4–5), after which it is sent to the ordering service (step 6). The ordered transaction is packaged into a block and distributed to all committing peers, which update the world state (steps 7–8). As a result, the final ledger state contains only ciphertexts and metadata, replicated across the network (step 9). At no point during this process is raw device data written to the blockchain.

2) Analysis: Researcher → Result (Steps 10–16): A researcher submits a query through the web application (step 10). The request is turned into a read-only transaction and sent to endorsing peers (step 11), which invoke the chaincode to read the ledger state and evaluate the query (step 12).

Since device records on the ledger are encrypted, peers evaluate queries over ciphertexts according to the active scheme selected during system setup (Paillier or BFV; see II-C). Peers do not hold secret keys.

Some queries, for example *grouped count*, *unique values*, or *histogram*, require comparisons or ordering of values. Since this functionality is not fully supported under encryption in the current implementation, the chaincode produces intermediate aggregates and forwards them to the TTP (step 13). *Intermediate aggregates* are partial summaries (for example, per-bin counts in a histogram) that are later combined into the final aggregate result. The TTP decrypts only what is necessary for the query to continue and returns the intermediate aggregate results to the peer (step 14). Regardless of whether a query required intermediate decryption, every query follows the same final path: once computation is finished, the resulting aggregate is sent to the TTP for final decryption (steps 13–14) before being returned to the researcher. Individual device records are never decrypted in this process.

Before returning the value to the user, ϵ-differential privacy is applied to the final aggregate. The result is then returned to the application and shown to the researcher (steps 15–16).

C. Cryptography and Secret Keys

All medical device data on the ledger is stored in encrypted form. The platform supports two homomorphic encryption schemes: Paillier, which allows addition of ciphertexts and multiplication by constants, and BFV, which supports both addition and multiplication. The choice of encryption scheme (Paillier or BFV) is fixed at deployment, but keys for the chosen scheme can be rotated during operation. Regardless of the scheme, no peer in the network ever sees secret keys, and only ciphertexts and metadata are written to the blockchain.

The Trusted Third Party (TTP) is responsible for key management. It generates and stores all secret keys and uses the public key to encrypt incoming hospital data. When administrators initiate a key rotation, the TTP issues a new key pair. New records are encrypted under the fresh public key, while metadata stored with each record ensures that older ciphertexts can still be decrypted with the correct key.

D. Queries

MedTech Chain 2.0 extends the set of queries available to cybersecurity researchers for analyzing medical device data. These queries allow cybersecurity researchers to look for patterns that may reveal vulnerabilities or weak points in hospital infrastructures, without exposing individual device records.

Count, *sum*, and *average* are executed entirely under encryption. They can be used, for example, to measure how many devices of a given type are present in a hospital, or to check the average usage hours across a fleet of devices. *Standard deviation* and *linear regression* provide insight into variability and correlations, helping to detect unusual usage patterns or relationships between device properties, such as production date and battery performance. These are computed on ciphertexts, except for the final decryption of the aggregate result. Queries that require comparing values, such as *grouped count*, *histogram*, or *unique values*, cannot be performed fully homomorphically in the current setup. For these, the chaincode produces partial aggregates, which are briefly decrypted by the TTP before the query performing instances continue execution on decrypted aggregate data.

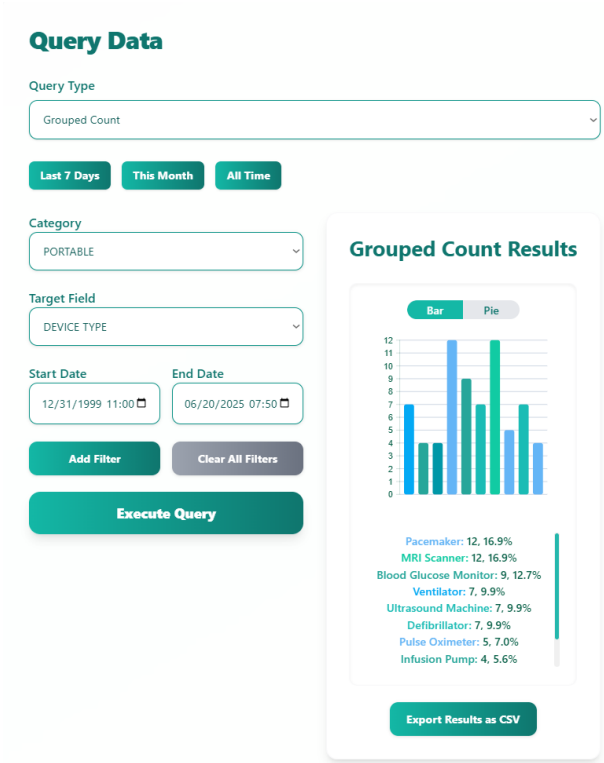

Fig. 2. Graphical interface for the grouped count query

E. Implementation

The platform is deployed on a Hyperledger Fabric network, where hospitals and the MedTech Chain Organization each run peers and orderers. Homomorphic operations are not implemented directly inside the chaincode but delegated to external binaries. This ensures that the complex arithmetic required by Paillier and BFV can be handled in efficient native code, while the chaincode remains simple and only forwards inputs and collects outputs. For Paillier, the binary is a Rust module supporting addition and multiplication by constants. For BFV, a C++ binary based on the OpenFHE library provides ciphertext addition and multiplication. All peers use the same binaries packaged with the chaincode container, so endorsement results remain deterministic across the network.

III. Maturity of MedTech Chain 2.0

Experiments were conducted to measure the computation overhead introduced by homomorphic encryption in MedTech Chain 2.0. Only the standard deviation query was tested, as it is the most computationally intensive and a good indicator of encryption overhead.

The experiment setup simulated the MedTech Chain Organization and two hospitals. All services were deployed with Docker, and each hospital uploaded randomized medical device data until the blockchain contained around 10 000 assets. The experiments were run on an AMD Ryzen Threadripper 7970X processor with 32 cores and 64 threads, running at a base frequency of 1.5 GHz.

TABLE I
AVERAGE EXECUTION TIME OF STANDARD DEVIATION QUERY OVER 1000 RUNS

Scheme	Avg. Time
None	2076.9 ms
Paillier	2346.4 ms
BFV	4142.3 ms

The results show that Paillier adds about 13% overhead compared to plaintext execution, while BFV is roughly twice as slow. Although BFV is less efficient in the current setup, queries still finish within a few seconds. The BFV implementation can also be considerably optimized further, as the current setup is not tuned for efficiency, so these numbers should be considered an upper bound rather than a limitation of the system.

IV. Additional Materials

In this paper, we presented MedTech alongside the improvements made in its second version. A video illustrating the main features and workflow of MedTech Chain 2.0 is available at https://github.com/MedTech-Chain-2-0/.github/blob/main/demo.mp4. In addition, the source code, and documentation are hosted at https://github.com/MedTech-Chain-2-0.

References

[1] Medical Device Cybersecurity Working Group, "Principles and practices for medical device cybersecurity," International Medical Device Regulators Forum, Tech. Rep., 2020.

[2] A. Petru-Rosu, T. Tataru, J. Zelenjak, R. Kromes, and Z. Erkin, "Medtech chain: Decentralised, secure and privacy-preserving platform for medical device data research," in *2024 6th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*, 2024, pp. 1–8.

[3] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Comput. Surv.*, vol. 51, no. 4, Jul. 2018. [Online]. Available: https://doi.org/10.1145/3214303

[4] A. Pankova and P. Laud, "Interpreting epsilon of differential privacy in terms of advantage in guessing or approximating sensitive attributes," in *2022 IEEE 35th Computer Security Foundations Symposium (CSF)*, 2022, pp. 96–111.