http://dx.doi.org/10.62762/TISC.2026.917693

RESEARCH ARTICLE

A Secure and Auditable Distributed Data-Sharing Framework for Multi-Domain Systems

Sarwar Sayeed[1,*], **Pavlos Papadopoulos**[1], **Dimitrios Kasimatis**[1], **Nikolaos Pitropakis**[2], **William J. Buchanan**[1], **Evangelos K. Markakis**[3], **Dimitra Papatsaroucha**[3], **Konstantina Pityanou**[3] **and Ilias Politis**[4]

[1] Blockpass ID Lab, School of Computing, Engineering and the Built Environment, Edinburgh Napier University, Edinburgh EH10 5DT, United Kingdom
[2] Department of Information Technology, The American College of Greece, Athens 15342, Greece
[3] Department of Electrical and Computer Engineering, Hellenic Mediterranean University, Heraklion 71410, Greece
[4] Industrial Systems Institute, Athena Research Center, Patras 26504, Greece

Abstract

Data is one of the essential assets of every organisation, and its security has frequently been highlighted as a concern due to unauthorised access, breaches, and tampering. Conventional data repositories are often based on centralised setups, which are prone to problems such as single points of failure, privacy breaches, and unpredictable system behaviour. While an auditing mechanism can strengthen transparency by tracking how data is accessed, such tools are often managed by external third parties or rely on manual methods. Additionally, due to complexities such as duplicate user credentials across multiple services and systems, it becomes even more difficult for users to control their content. Furthermore, without an access-monitoring system in place, attackers can usually evade detection, leaving the system compromised while they manipulate sensitive data. To address these shortcomings, we propose a distributed data-sharing framework, developed as part of a Trusted Mediator Platform, that enhances data privacy and security by employing a distributed data repository for secure storage and retrieval, integrated with an audit framework implemented on a permissioned blockchain. Our method is designed to mitigate unauthorised access, data tampering, and single points of failure (SPOFs) by distributing data across trusted nodes and providing a tamper-resistant ledger, ensuring traceability and integrity. Our experiment demonstrates strong performance, with average blockchain transaction latency remaining below 1 s for batches of up to 10,000 transactions (tx) and throughput attaining over 115 tx/s, while IPFS achieved an upload throughput of approximately 35 one-megabyte files per second at a batch size of 300, thus indicating their feasibility for the exchange and secure storage of essential data.

Keywords: distributed storage, audit trail, data sharing, privacy, transparency.

Submitted: 11 August 2026
Revised: 10 September 2026
Accepted: 16 September 2026
Published: 23 September 2026

Vol. 2, **No.** 3, 2026.
10.62762/TISC.2026.917693

***Corresponding author:**
✉Sarwar Sayeed
S.Sayeed2@napier.ac.uk

Citation
Sayeed, S., Papadopoulos, P., Kasimatis, D., Pitropakis, N., Buchanan, W. J., Markakis, E. K., Papatsaroucha, D., Pityanou, K., & Politis, I. (2026). A Secure and Auditable Distributed Data-Sharing Framework for Multi-Domain Systems. *ICCK Transactions on Information Security and Cryptography*, 2(3), 172–191.

1 Introduction

The current era is often labelled the age of digital innovation, where data is becoming the driving force, transforming every aspect of life. This transformation spans healthcare, automotive, education, energy, space, food tracking, and more, generating vast volumes of transactional records that must be handled with transparency and accountability — a requirement that blockchain-driven audit trail systems have been specifically designed to address [1]. Given the data-intensive nature of these domains, rapid technological advancement influences how data integrity should be guaranteed. Data integrity is indeed crucial for ensuring that a storage method is secure, reliable, and protected from unauthorised access, modification, or fabrication [2]. Although data integrity can be achieved with a single database, protecting data against adversaries becomes challenging. Therefore, while industries are concerned about the constantly changing environment, they tend to seek more secure strategies to stay aligned with rapid technological shifts.

IBM's 2026 Cost of a Data Breach Report [3, 4] reveals that the global average cost of a data breach reached $4.99 million, an increase of 12% over the last year, with the Healthcare sector being the most expensive for the past 13 years. Organisations also faced severe consequences due to inadequate security, and it took an average of 247 days to identify and contain a data breach globally, highlighting the necessity of data security. Additionally, the National Public Data (NPD), which provides background checks, became the subject of a class-action lawsuit after attackers stole sensitive user information. The incident in 2024 is considered one of the largest data breaches, compromising nearly 3 billion personal records, including full names, addresses, and Social Security numbers [5]. Similarly, the AT&T data breach has affected nearly all of its customers and is considered one of the largest in 2024 [6]. The telecom giant has, in fact, suffered two separate data breaches in 2024, months apart, the larger of which affected nearly 110 million people.

Conventional data security approaches have predominantly focused on shielding with "harder walls", either by adding additional layers of authentication for access control or by employing robust encryption techniques. However, these techniques ultimately rely on the same principle: once communication with the system is established, the resources can be tampered with and eventually become compromised. Additionally, the access control in traditional systems is often managed through centralised databases and identity management infrastructures, thus creating a single entry point for attackers to exploit; comparative analyses of centralised and decentralised identity management systems have highlighted these scalability and security shortcomings [7]. Centralised models are also often vulnerable to insider attacks, as they enable malicious administrators to manipulate critical data. Because data is stored in a central location, it repeatedly becomes a prime target for attackers, resulting in various attacks, including Structured Query Language (SQL) injection, ransomware, and Distributed Denial of Service (DDoS).

As the technology advances, the significance of the distributed technologies continues to grow due to increasing demands for security, scalability, interoperability, immutability, and transparency. Although numerous peer-to-peer file-sharing applications have emerged over the past years, most have failed due to a lack of accountability and content verification [8]. BitTorrent, however, succeeded in distributing files using networks of untrusted peers, yet such applications failed to meet the requirements of sensitive use cases. Amid recent challenges, cutting-edge technologies such as the InterPlanetary File System (IPFS) and blockchain offer promising solutions for modern-day data storage and sharing [9]. IPFS is a peer-to-peer distributed data storage and sharing protocol with a content-addressing mechanism that provides a tamper-evident hash utilising cryptographic techniques [10]. Its censorship-resistant design is intended to prevent any single entity from restricting or removing data from the public network. This helps sustain the availability of the content for as long as at least one node retains (pins) it, while the content-addressed hash makes unauthorised modification detectable. Blockchain, in contrast, is a Distributed Ledger Technology (DLT) that does not require a central administrative authority, reducing exposure to attacks that exploit a central authority, such as insider threats, Man-in-the-Middle (MITM) attacks, Denial-of-Service (DoS), and more, although blockchain systems also have their own attack surface [11, 12]. Its ability to offer a tamper-resistant audit mechanism to protect against challenges such as data breaches and tampering while ensuring data integrity, transparency, and security is critical.

This work presents a distributed data-sharing

framework as part of the Trusted Mediator Platform, enabling users to store data in a distributed data repository. Each piece of data stored is identified by a unique hash, ensuring data integrity and immutability while making tampering attempts easily detectable. Subsequently, all activities are recorded in a permissioned ledger implemented in Hyperledger Fabric. The broader platform addresses the requirements of its end users, such as consumers, organisations, and industrial partners, for federated and Findable, Accessible, Interoperable, and Reusable (FAIR) data access through a private framework of public and private organisations [13]. The architecture is designed in line with the Horizon Europe guidelines for FAIR and open access to data.

The main contributions of this work can be summarised as follows:

- Design and implementation of a private distributed data repository, empowering a scalable and secure storage solution for managing digital content.
- Development of a robust and auditable Transaction Recorder utilising a permissioned blockchain to record all activities for accountability and transparency.
- Seamless integration of the components into a unified framework in tackling complex data-sharing requirements.

The rest of the paper is organised as follows: Section 2 discusses related technological components, providing a detailed overview of their technical capabilities, including a comparative analysis of contemporary distributed data repositories, current permissioned blockchain frameworks and associated monitoring and visualisation tools that support them; Section 3 discusses the Trusted Mediator Platform and the architectural components developed in this work; Section 4 conducts a comprehensive performance assessment to validate the efficiency of the proposed system; and finally, Section 5 discusses the overall findings and provides avenues for future work.

2 Background and Literature Review

The growing reliance on data-driven systems has underscored the need for a secure, transparent, and reliable framework that goes beyond traditional centralised systems. Alongside outlining the limitations of centralised models, this section delves into the foundational concepts and technologies that shape a distributed data-sharing platform. Moreover, it reviews related studies, emphasising the current advancement in the area.

Figure 1 depicts an example of a centralised data-sharing platform, which is often fragile due to its full reliance on a single trusted authority, making it prone to tampering, unauthorised access, SPOF, and insider threats. In this scenario, a data owner (e.g. a hospital) submits sensitive information to the central admin, which serves as the primary controller and is responsible for managing centralised storage and access control. As identified in prior work on blockchain-based hierarchical data sharing, this fully centralised architecture exposes inherent security vulnerabilities and single points of failure [14]. The admin uploads the submitted data to a cloud storage service operated solely by the admin. The regulatory body is the central authority that governs data access, use, and compliance, while domain authorities, such as healthcare, ensure data meets sector-specific compliance requirements. Any data access request made by a user is assessed by the admin in accordance with sector regulations prior to approving or denying access. The centralised architecture and full dependence on a single trusted authority expose security vulnerabilities and architectural fragility.

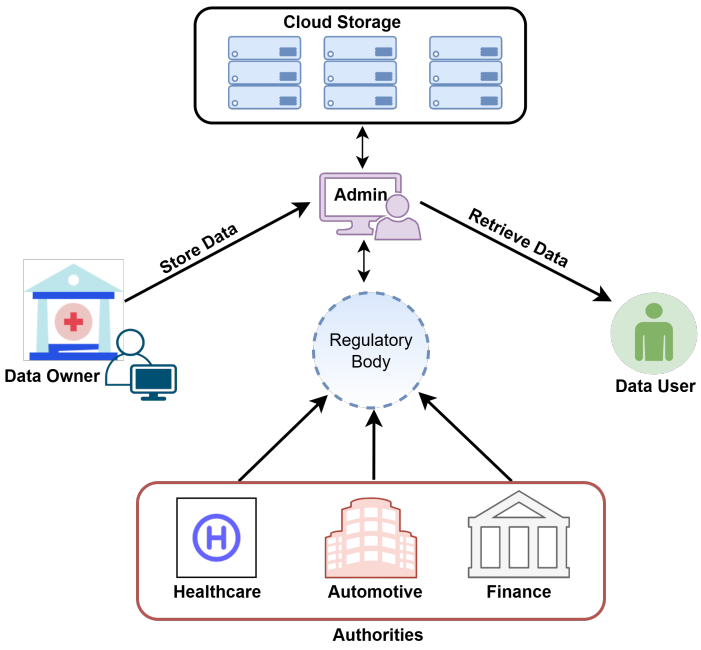

Figure 1. Centralised data sharing model with domain-specific data management.

2.1 Distributed Ledger Technology (DLT)

DLT is a record-keeping system that replicates data across multiple nodes, rather than being handled by a

Table 1. Comparison of permissioned blockchains.

Permissioned Frameworks	Consensus Mechanism	Transactions per Second (TPS)	Permissioned Control	Token / Crypto	Primary Domains	Limitations
Hyperledger Fabric	Raft, SmartBFT (v3.x); Kafka, Solo (deprecated)	3,000–20,000 (20,000 with FastFabric) [15]	Membership Service Provider	✗	Supply Chain, Healthcare, Automotive	Highly modular but requires expert configuration
R3 Corda	Notarisation (Validity & Uniqueness)	∼600–1,678 [16]	Node Identity Service	✗	Financial service, Trade	Process single flow at a time unless optimised
MultiChain	Round-Robin	2,500 (peak) [17]	Permissions Management	✗	Data sharing, IoT	Inflexible and not designed to align with smart contracts
Quorum	RAFT, Istanbul BFT	∼300 (peaks >600) [18]	Quorum Permissioning	✗	Financial tx, Supply Chain, Voting	Limited adoption & costly maintenance
Ripple	RPCA	Lower than theoretical capacity in practice [19]	Permissions Management	✓	Banking, Remittances	Limited privacy features

centralised authority [20]. The foundation of DLT lies in a distributed, decentralised peer-to-peer network, supported by core security mechanisms. DLT ensures data integrity, transparency, and trust in business settings.

In recent times, there has been significant attention to trusted distributed systems due to their robust security approach and their ability to address the challenges faced by existing systems [21, 22]. Blockchain is a DLT that provides a secure method for recording transactions. The distributed nature of blockchain ensures that transactions are stored across multiple locations rather than a centralised database, thereby reducing the risk of SPOF [23]. Although public and permissioned blockchains are the two primary types of blockchain networks [24], permissioned blockchains are increasingly adopted for enhanced privacy, confidentiality, faster performance, and scalability across domains, including healthcare, finance, automotive, farming, and more.

2.2 Decentralised Data Repository

A data repository can be implemented as a managed storage system combining database infrastructure with distributed file systems to support scalable data storage and access [25]. Data repositories often play a crucial role in enhancing businesses' decision-making processes. Several methods can be adopted to store data, ranging from centralised approaches such as data warehouses and data lakes to decentralised alternatives that leverage blockchain and distributed file systems for secure storage and access control. The latter approach has demonstrated particular promise, as exemplified by blockchain-integrated IPFS storage schemes that enforce fine-grained access control while ensuring data integrity [26]. A data warehouse is a relatively large-scale repository that collects data from various sources, while a data lake stores unstructured data and categorises and tags it with metadata. Data marts can be subsets of a data repository that provide a user-friendly approach focusing on the user's requirements. Beyond these, metadata repositories and data cubes offer further options for collecting and storing data in different ways.

Despite the simplicity of the centralised data repository, its limited transparency and inherent challenges motivate a shift towards decentralised alternatives. Moreover, centralised architectures weaken the Confidentiality, Integrity, and Availability (CIA) triad, where insider threats can breach confidentiality, alteration to records can compromise integrity, and a SPOF can disrupt the service for the entire system [27]. A decentralised and distributed repository mitigates these risks through replication, content-addressed integrity verification, fault tolerance, and resilience to failures [10]; consensus, end-to-end encryption, and fine-grained access control can be layered on top of it. Its distributed architecture prevents performance bottlenecks when processing large datasets and enhances data retrieval and analysis speed.

2.3 Network Visualisation and Analysis

Network visualisation displays complex data relationships in graphs or models to understand how elements are connected. It helps to understand a system's structure, behaviour, and interactions. In the context of blockchain, it facilitates the identification of various attacks that potentially threaten the integrity and confidentiality of distributed records [28]. When

mapped with network nodes, transaction flows, and access patterns, it visualises real-time monitoring, enabling the detection of anomalies and unusual behaviour and providing complete insights into the network and its associated nodes. Network visualisation allows security analysts to monitor, detect, and respond to internal and external threats. Likewise, it helps detect suspicious activities easily by providing a complete forensic analysis of privilege misuse through visual audit trails. Besides that, it makes it easier to assess how nodes behave in a distributed environment, as per the access rules and imposed policies.

2.4 Comparative Assessment of Permissioned Blockchains

Modern systems offer various methods for sharing, storing, and managing data, often making it challenging to choose the right mix. This part provides a detailed comparative analysis of key technologies, focusing on their core functionalities, strengths, and limitations. The permissioned blockchains are based on trusted networks where access to transactions is restricted unless authorised. Unlike the full transparency of public blockchains, permissioned blockchains offer controlled or selective transparency based on access controls. Some of the major advantages of permissioned blockchains are robust privacy, modularity, faster transactions, and high scalability. Hyperledger Fabric, Ethereum Geth, Quorum, MultiChain, and R3 Corda are some of the permissioned blockchain frameworks with distinct features [24].

Table 1 provides an overview of permissioned blockchains, highlighting their unique attributes, consensus mechanisms, capabilities to record transactions, and use-cases, as well as application areas they are suitable for. The comparison indicates that Hyperledger Fabric offers the strongest combination of reported throughput, modularity, and ecosystem maturity, particularly when compared to R3 Corda, MultiChain, Quorum, and Ripple, albeit at the cost of a more demanding configuration. While R3 Corda facilitates a wide range of use-cases [33, 34], its reported throughput is lower than that reported for Hyperledger Fabric [15, 35]: vendor-reported figures indicate roughly 600 TPS for a realistic two-party payment flow and up to ∼1678 TPS for single-node issuance transactions [16]. Additionally, MultiChain and Quorum are domain-specific, while Hyperledger, in contrast, provides a modular architecture that is highly customisable depending on the use case and participating entities, making it suitable for various sectors including healthcare, cyber security, automotive, and finance [36, 37]. MultiChain lacks full smart contract support, as evidenced by reported challenges in applying it to real-world use cases such as real estate, and is therefore not well-suited for general-purpose applications requiring programmable business logic [38]. Quorum requires costly maintenance for enhanced performance; thus, it is often not widely adopted [39, 40]. Similarly, Ripple also has limitations, being limited to tokenised domains with limited privacy features [41, 42] and a significantly lower observed transaction rate in practice compared to its reported theoretical capacity, as demonstrated by empirical testing of the Ripple consensus protocol [19]. Therefore, when considering performance, scalability, and operational flexibility, Hyperledger Fabric demonstrates distinct advantages over other alternative frameworks.

2.5 Comparative Assessment of Distributed Repositories

Distributed data repositories have been shown to provide high availability, improve throughput, and mitigate SPOFs [43]. Table 2 provides a comparison of distributed storage systems, highlighting their storage models, key application domains, limitations, tokenisation mechanisms, and data access times. While the reported access times vary in format and different testing conditions, ranging from latency to full retrieval, they provide essential insights into how those systems perform under different use cases and scenarios. IPFS is a distributed file storage protocol based on a peer-to-peer network that provides fast content retrieval in under a second and applies to multi-domain applications, including healthcare, finance, automotive, space, and more. Other storage alternatives, including Arweave and Filecoin, focus on permanent storage but fall short in data retrieval speed [44]. Compared to IPFS, Arweave generally involves higher storage costs as it requires payment with AR tokens. The reported latency reflects block propagation rather than data access and rises from 200 ms to 1200 ms with increasing block size [30]. Likewise, Filecoin has longer retrieval times with complex setup requirements. Despite offering similar functionalities, both Storj and Sia are limited to specific use cases. They pose challenges due to scalability limitations, storage cost variances, and full dependence on host availability [32, 45]. Therefore, IPFS widely stands out as a cost-effective and reliable distributed

Table 2. Comparison of distributed data repositories.

Frameworks	Storage Model	Access Time	Token	Primary Domains	Limitations
IPFS	P2P distributed file system	<1 s (private deployment; Sect. 4)	✗	Healthcare, Finance, Agriculture, Automotive	Complex setup for non-developers
Storj	Decentralised P2P cloud storage	High throughput; performance varies by network configuration [29]	✓	Cloud-based backups for governments, banks and enterprises	Moderate costs, Archival issue
Arweave	Blockchain-based permanent storage	~200–1200 ms (block propagation) [30]	✓	Long-term archiving for media	Higher costs, Slower retrieval
Filecoin	Decentralised storage network	<1 s (if cached); 5–10 min if unsealing needed [31]	✓	Multidomain: Entertainment Finance, Health, Energy	Higher Costs, Price fluctuations
Sia	Decentralised cloud storage	2–3× faster than Google Drive (512 MB–10 GB) [32]	✓	Backup solutions for small-medium scale businesses	Moderate scalability, Host reliability

Table 3. Comparison of blockchain monitoring and visualisation tools.

Frameworks	Scope	Blockchain Integration	Capabilities	Limitations
Hyperledger Explorer	Network monitoring & visualisation	✓	Dashboard for visualising transactions, nodes, blocks and network metrics	Limited to Hyperledger Fabric blockchain
Kibana	Data visualisation & analytics	✓	Browser-based visualisation tool for exploring large log volumes using line graphs, bar charts and pie charts	Limited transaction data, Inefficient to identify threats
Elasticsearch	Data storage and search	✓	Supports complex queries, real-time analytics, and systematic support for filtering, sorting, pagination, and aggregations	Requires additional setup tools, such as Kibana
D3.js	Data-driven visualisations	✓	Open-source JavaScript library for creating dynamic and flexible web-based graphics	Requires custom setup & advanced programming skills
Caliper	Performance benchmarking	✓	Comparative performance evaluation tool for different blockchain solutions	Inefficient as visualisation tool

file storage system across various domains, offering quick data retrieval.

2.6 Comparative Assessment of Monitoring Tools

Monitoring tools integrated with data systems support tracking potential vulnerabilities and maintaining the integrity of transactional records [46]. Real-time monitoring of network activities helps to track the flow of the recorded transactions, access controls, and all user interactions.

Table 3 provides a comparative summary of tools that can be integrated with blockchain for monitoring network activity and benchmarking. It highlights their primary purpose, capabilities, and limitations, emphasising how they differ in their approach to real-time network analytics. Hyperledger Explorer was a monitoring tool from the Hyperledger project that connected to blockchain nodes to provide network visualisation. Explorer allowed authorised users to audit transactions, examine smart contract execution, and inspect consensus mechanisms. However, Hyperledger Explorer has since reached end-of-life, and its repository was archived in February 2026 [47]. Besides Explorer, Hyperledger Labs has other community projects that integrate Fabric with Elasticsearch and Kibana for log visualisation and analytics. In parallel, research efforts have explored tighter integration of Hyperledger Fabric with IPFS-based storage to enhance secure multi-group data sharing [48], further demonstrating the complementarity of these two distributed technologies. However, integrating such monitoring tools with blockchain systems introduces additional complexity and operational challenges that are not yet fully resolved — a reflection of the broader unsolved interoperability and tooling challenges that remain open in the blockchain domain [49]. Both Elasticsearch and Kibana display limited transaction data and often fail to identify potential threats in Hyperledger Fabric. Moreover, experiments reveal that aggregating additional data sources is often challenging when using standard tools such as the Elastic Stack [28]. D3.js is a flexible, low-level visualisation library that can also be integrated with Hyperledger. It typically requires a custom setup and

users with advanced programming proficiency [50]. Although Hyperledger Caliper can be a valuable tool for visualising performance metrics, including transactions per second, latency, and resource usage, it is inefficient for monitoring live transactions and resource metrics [51]. Prometheus and Grafana, in contrast, are effective for resource monitoring but do not provide functionality for blockchain monitoring.

In summary, the comparative performance specifications of the permissioned blockchains, distributed data repositories and monitoring and visualisation tools originate from heterogeneous platforms and experimental environments with differences in hardware, network configuration and implementation settings. The assessment hence provides an indicative comparison only and does not provide fully comparable benchmarks, a limitation that is common to empirical performance studies of permissioned platforms [52].

2.7 Related Works

Distributed systems continue to reshape multiple sectors through robust consensus mechanisms along with effective security solutions. The following reviews existing studies on distributed technologies and related approaches for secure data management and sharing.

Chen et al. [8] propose an improved peer-to-peer (P2P) file storage scheme combining IPFS and blockchain technology. Their method primarily addresses the high-throughput problem faced by IPFS users by introducing the role of content service providers. To mitigate data reliability, availability, storage overhead, and other issues for service providers, they have introduced a zigzag-based storage model to improve the IPFS block storage model, integrated with blockchain.

Meanwhile, Kang et al.'s [53] method combines Named Data Networking (NDN) technology with blockchain and IPFS. The main idea is to develop a blockchain-based knowledge storage and sharing technique based on NDN.

Another study by Wang et al. [54] utilises a blockchain-based distributed storage system to enhance Knowledge Graph security. The key concept for their experiment is to investigate the efficacy of a knowledge graph. The proposed scheme processes file storage and traceability of the knowledge graph, leveraging blockchain and a distributed file storage system.

Azbeg et al. [55] propose BlockMedCare, a secure healthcare system for remotely monitoring patients. Their method integrates the Internet of Things (IoT) with blockchain and an off-chain database implemented with IPFS. Security, scalability, and processing time have been the primary considerations for BlockMedCare, and therefore, the security methods are established using a re-encryption proxy, combined with Blockchain to store hash data.

In contrast, Ye and Park [56] highlight the importance of vehicle data, noting that security and capacity remain challenges when storing data securely and efficiently. They have proposed a system capable of securely storing vehicle data, utilising blockchain and IPFS, enabling users to easily manage vehicle data using a decentralised application (DApp) implemented on Ethereum. The experiments conducted show the performance of the proposed method, thus establishing the advantages in terms of data-processing speed and cost.

Pawar et al. [57] propose an agriculture supply chain management system combining IPFS, blockchain-based Ethereum smart contract, and a decentralised repository to automate the process and data exchange. Instead of relying on third parties, only the file's hash is stored on the blockchain. Their proposed approach ensures full decentralisation, thereby reducing the various challenges often introduced by centralised systems. Besides that, their evaluation demonstrates a significant reduction in file upload and download times.

Wang and Susilo [58] have designed a data security storage model for IoT systems. It is based on a permissioned blockchain and ensures integrity through a non-tampering approach and a verification layer to prevent the storage of malicious data. The security experiments demonstrate that the proposed approach effectively enhances the protection of IoT data.

To address challenges with the traditional healthcare architecture, Jayabalan and Jeyanthi [59] proposed a scalable blockchain-based model that leverages off-chain storage. The authors utilised IPFS to address storage challenges. Their method comprises two-factor and multi-factor authentication for mitigating fake node attacks. The Patient-centric access model enables patients to serve as digital stewards of their data, granting doctors and hospitals on-demand access and revoking it after a specific period.

Likewise, Tovanich et al. [60] systematically reviewed visualisation approaches for blockchain data, highlighting the role of visual analytics in exploring transaction and network activity.

Despite many prior studies advancing the interdisciplinary research, they are primarily research prototypes that either explored distributed solutions with application-specific blockchain architectures or depend on public blockchains such as Ethereum, which can often be less suitable due to privacy issues, transaction fees, and governance complexity. In contrast, the novelty of the proposed method concentrates on the integrated use of cutting-edge technologies by combining on-chain immutable transaction records with distributed off-chain data storage. The combined design of Fabric-based auditability and IPFS-based secure storage solution is further intended to go beyond a single-domain prototype through a multi-domain and scalable architecture. The proposed solution therefore aims to address transaction cost, storage scalability and organisational challenges while supporting common data sharing.

3 Trusted Mediator Platform and Implementation

3.1 Architecture

The Trusted Mediator Platform is designed, utilising a collection of modular components centred on data sovereignty, Homomorphic Encryption (HE), and agreement mechanisms to facilitate compliance with data protection rules, including the General Data Protection Regulation (GDPR) [61–63]. Its multi-sided architecture further draws on established principles for designing federated data platforms and data spaces [64], acting as a federated, privacy-preserving intermediary between distributed Data Providers and Data Consumers.

Figure 2 outlines the broad perspective of the Trusted Mediator Platform, which incorporates Data Privacy Assurance Services, Trust Assurance and Data Protection Services, Artificial Intelligence (AI) Services, and the Storage Services. The Data Privacy Assurance Services block consists of the Homomorphic Enabled Data Fusion components, the Self-Sovereign Identity-Homomorphic Framework modules, the Data Protection Impact Assessment component, and the GDPR Compliance Checker. The AI Services block includes the AI Models-as-a-Service components and the Trustworthiness AI Support Design Framework, which is also part of the Trust Assurance and Data Protection Services block that further includes the Security and Trust Manager, the Authentication Manager, and the Accountable Transaction Recorder. In addition, the Storage Services block comprises the Data Lake, Knowledge Repository, and One-Stop-Shop that additionally contains the External Connectors modules, which facilitate the communication of the Trusted Mediator Platform with external Data Spaces, such as Gaia-X. Certain services of the platform are deployed in a containerised form at the premises of the end users to facilitate the operation of encrypted processes on their side (encrypted search and encrypted computation), which highlights that the Trusted Mediator platform (its centralised part) never sees or accesses either raw or encrypted data.

Data Providers provide descriptions of their datasets in order to register them as part of the federated data sources of the Trusted Mediator, while the raw data remains at their premises, ensuring data sovereignty and total control are retained. They are required to perform a GDPR compliance check, data privacy impact assessment, and agreement signing, indicating which computations they allow to be performed over their datasets in the encrypted domain, as well as if they allow knowledge stemming from the use of their datasets to be published in the platform. In contrast, Data Consumers utilise the Dashboard in order to construct and submit search queries that are executed in the encrypted domain. In addition to that, they are allowed to access cached results from previously conducted and published encrypted search and computation processes without the need to initiate new queries, while the underlying information is not disclosed at any point in the process. Before executing the computation in the encrypted domain, Data Consumers are required to follow the necessary agreement signing, GDPR compliance check, and data privacy impact assessment procedures enforced by the Trusted Mediator. When they obtain the result, they may decide whether they wish to publish it on the platform, making it available for other users to explore. As a result, the storing and publication of knowledge are governed by the bilateral consent mechanism. The Data Providers consent for knowledge storage is defined in the dataset agreement. Likewise, the Data Consumers consent is required for publication of knowledge derived from the encrypted process they perform. If such consent is not granted from both ends, search query results and computation results are not published on the platform.

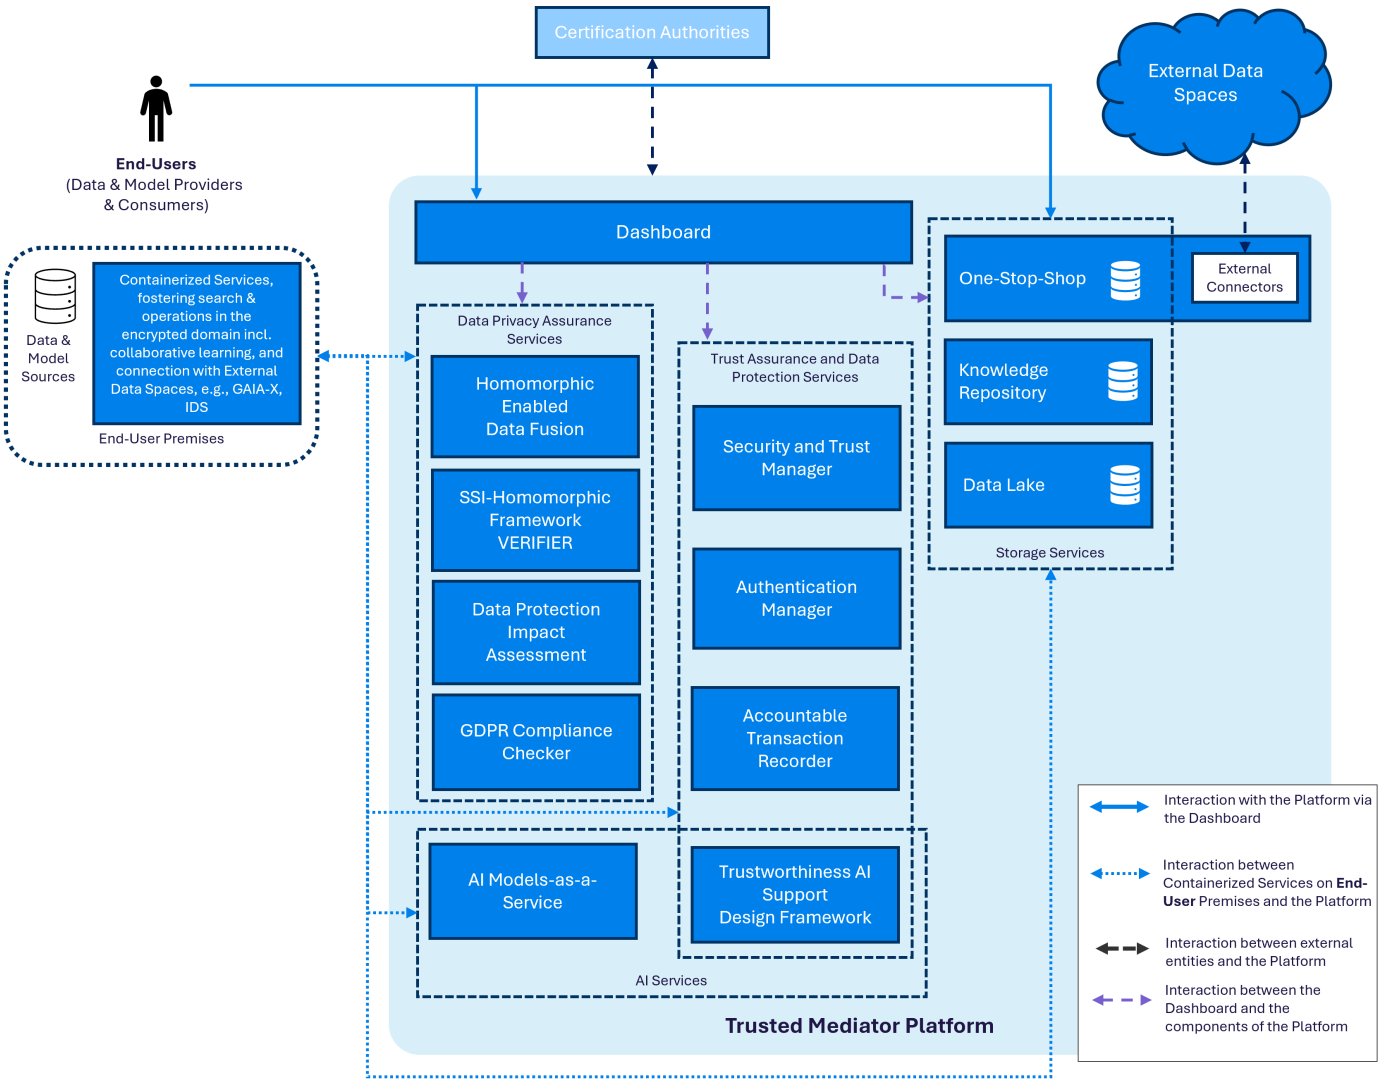

Figure 2. Trusted Mediator Platform architecture.

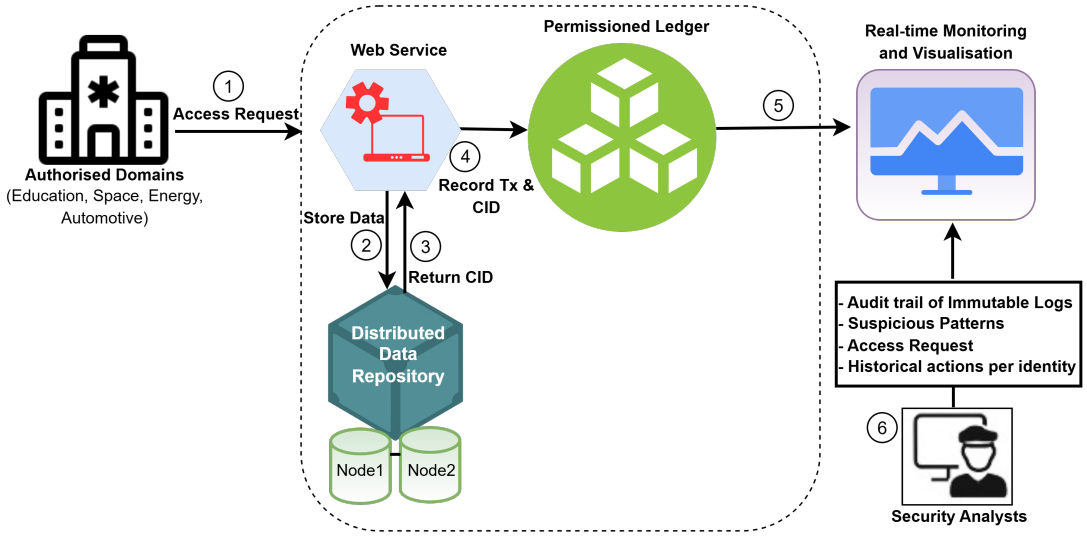

Figure 3. Distributed data sharing and auditable transaction recording architecture.

In this manner, the platform enables secure access, fusion, and analysis of sensitive datasets from multiple domains (e.g. health, education, energy). Data always remains encrypted, and the platform never

requires access to the raw data itself, thereby helping to preserve privacy and maintain compliance with data regulations.

3.2 Implementation

The implementation emphasises two core components of the Trusted Mediator Platform: the Accountable Transaction Recorder and the Knowledge Repository. The experimental environment uses Hyperledger Fabric v2.5 [37] for ledger and transaction management, Fabric CA v1.5 for managing identity and certificates, IPFS/Kubo v0.29.0 for distributed data storage and Docker client v27.4.0 for containerisation. Figure 3 presents the high-level architecture of the proposed distributed data sharing framework, where an authorised entity (e.g. the education domain) submits a request to store both data and corresponding transactions. The *Web Service* handles the dual request by first storing the data contents in the IPFS-based distributed repository, retrieving the resulting Content Identifier (CID) and associating it with the corresponding metadata. This metadata, along with the related transactional information, is then recorded on the blockchain ledger to ensure auditability. Algorithm 1 provides the corresponding pseudocode for storing and retrieving using blockchain and IPFS.

3.2.1 Blockchain Transaction Management

To facilitate secure auditability and immutability, the implementation leverages Hyperledger Fabric permissioned blockchain. The dedicated chaincode, developed in JavaScript, executes the core business logic to record every transaction of the platform within the immutable ledger. The chaincode records CIDs generated by IPFS, effectively linking them to domains, domain-specific metadata, and the transaction identifier (TxID) for fast retrieval from the ledger. Figure 4 illustrates the process flow for the single-channel permissioned blockchain. The architecture comprises two generic Organisations, *Org1* and *Org2*, each equipped with a peer node and a network-wide orderer using Fabric's Raft-based ordering mechanism. Each transaction is cryptographically signed and validated by the endorsing peers before being included in the ledger. The Organisations are equipped with Certificate Authorities (CAs) and Membership Service Providers (MSPs) to enforce strict identity management, enabling only authenticated and authorised users to participate in network activity. All communications among network components are secured using Transport Layer Security (TLS) certificates, ensuring encrypted information exchange and mitigating potential network attacks. All the blockchain components, including the peers, orderers, and CAs, have been orchestrated using Docker Compose, as containerisation streamlines infrastructure management, ensuring full consistency without disruption when additional components are integrated.

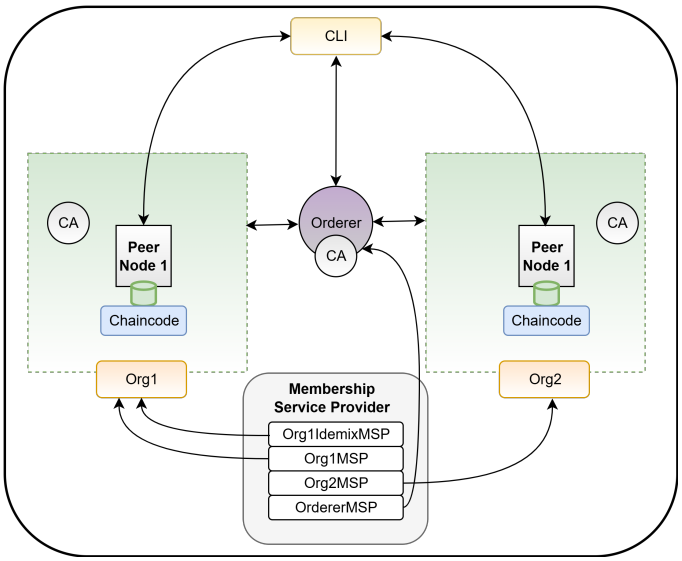

Figure 4. Topology of single-channel permissioned blockchain.

Access Validation and Policy Enforcement: The access control policy is imposed based on endorsement policies defined in the channel configurations. It therefore impedes any request occurring from unauthorised sources that do not comply with the policies. Each transaction submitted by authorised domains (e.g. space, automotive, healthcare) follows a structured JSON payload format containing 20 transaction fields, including the submitting provider's ID (ProviderID), target domain, organisation unit, and identifiers to data generated from the distributed data repository. The schema is designed to be domain-agnostic, allowing one to follow the same structure for seamless communications across multiple domains.

Tamper-resistant Transaction Recording: Each transaction submitted is recorded immutably comprising domain-specific metadata along with the CIDs from the data repository. The transaction data and the contextual metadata are logged as an auditable record referencing to the underlying data stored in the off-chain repository. Every transaction is assigned a globally unique TxID, which serves as an immutable reference for post-retrieval and audit operations. The

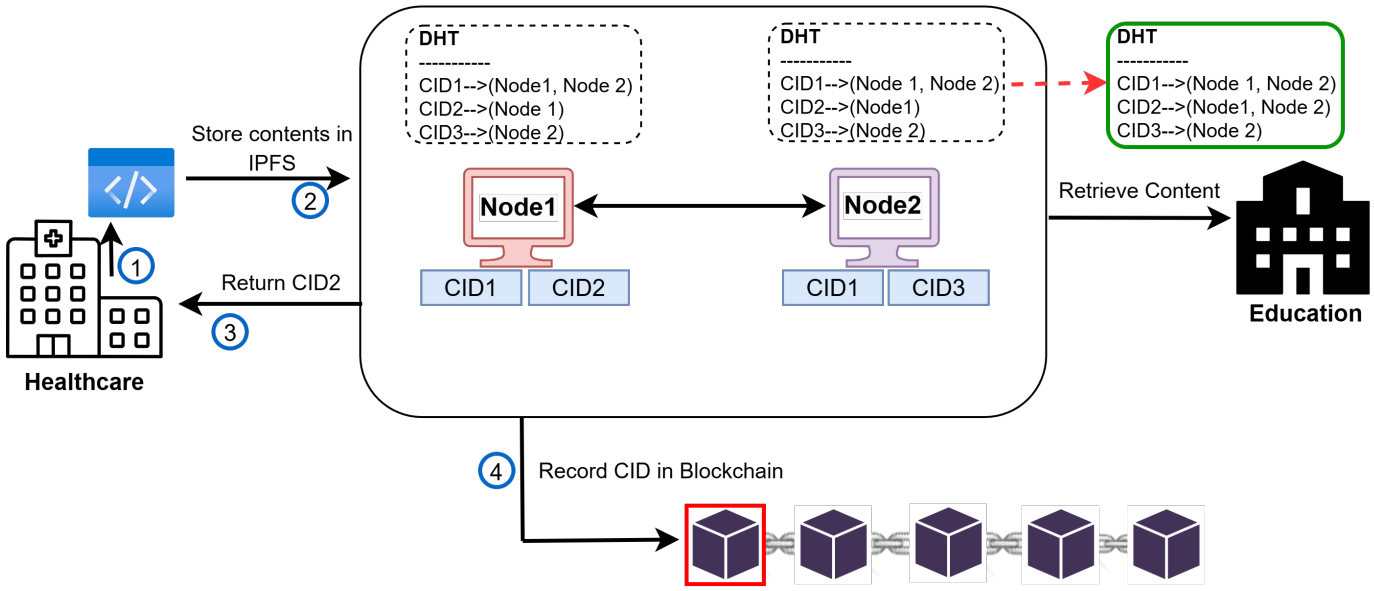

Figure 5. Distributed data repository based on IPFS.

Algorithm 1: Data content storage and retrieval

Input: Data content R, metadata M
Output: Retrieved data content R'
Receive data content R and associated metadata M;
Submit R to the distributed data repository (IPFS);
$CID \leftarrow \text{IPFSAdd}(R)$;
Create a Hyperledger Fabric record containing CID and metadata M;
Submit the record to Hyperledger Fabric;
Validate and commit the transaction to the Hyperledger Fabric ledger;

When an access request is received:;
Verify the requester's access permission;
Query the Hyperledger Fabric ledger using the relevant transaction ID;
$TX' \leftarrow \text{QueryTransaction}(\text{transactionID})$;
$CID' \leftarrow \text{ExtractCID}(TX')$;
Use CID' to request the data content from the distributed data repository;
$R' \leftarrow \text{IPFSGet}(CID')$;
return R';

composite key functionality implemented within the chaincode (e.g. ProviderID) enables partial filtering of domain-specific transactions without scanning the full ledger. This method enhances retrieval performance, giving each domain the ability to sort transactions independently when performing domain-specific audits. The comprehensive ledger query facilitates sequential navigation from the genesis block to the latest committed block, enabling navigation of historical audits of transactions and providing security analysts' compliance verification and reporting leveraging the same ledger platform.

3.2.2 *Distributed Data Repository*

The distributed data repository is implemented utilising IPFS, which is configured to provide enhanced data privacy, controlled access, performance optimisation, and secure data purge capabilities. To guarantee isolation and enhanced security, a private IPFS network with two nodes is customised in a way that remains fully isolated from the public IPFS network. This ensures a controlled environment for content storage, facilitating only designated nodes to take part in exchanging digital content.

The bootstrap nodes, which typically act as entry points for IPFS connections, have been removed from the configuration of both nodes so that these private instances could never be discovered by any public nodes. The elimination reduces the possibility of any malicious connections or potential threats that could lead to data leaks.

The implementation is fully containerised using Docker for easier management. Dockerisation facilitates consistency across diverse deployment environments and simplifies the scaling process by permitting new nodes without disrupting the flow of the network. Figure 5 defines the structural overview within the private IPFS network. An authorised domain (e.g. Healthcare) stores its content on *Node1* and receives the content hash (CID2) in response. The

Distributed Hash Table (DHT) of each node maintains metadata providing insights on which nodes store a particular piece of content. Afterwards, when another domain (e.g. Education) connects to *Node2* to retrieve the same content, *Node2* first checks its local DHT. As the DHT indicates *Node1* is the holder of the content, a query is sent for retrieval, and the DHT table is updated for *Node2*, optimising future references.

Data Integrity and Content Addressing through DHT: The content-addressing mechanism ensures the integrity of any data stored by leveraging cryptographic hashing to determine that it has not been tampered with. Since each submitted data generates a unique hash, it guarantees strong data integrity and acts as a tamper-evident pointer, preventing any tampering attempts. The dual-node implementation enhances load balancing and optimises performance. Besides, the DHT records which peers provide each CID, so that content can be located and fetched from any node holding a copy; replicas are created when a node retrieves or pins the content.

Secure Data Storage and Distribution: Following API key authentication, authorised entities submit their data to be stored in the distributed repository. Stored data can be replicated to the other node(s) on demand (through retrieval) or by explicit pinning, which allows retrieval from any node holding a copy and improves redundancy, availability, and fault tolerance. The original content or data is never stored on-chain (e.g. on the blockchain); instead, only the CID and relevant metadata are recorded on-chain. In this way, it creates a linkage between the stored contents/data within IPFS and tamper-resistant transactions on-chain, reducing attack surface by enhancing data confidentiality and privacy.

3.2.3 Custom API and Web Service

All communication within the framework occurs through the implemented custom API that enables seamless communication with the permissioned ledger. The API provides clearly defined endpoints for multiple ledger operations, including submitting transactions, querying domain-specific records, querying records by TxID, and listing all historically recorded transactions. Besides the custom blockchain API, the IPFS API is used to streamline the interactions with the distributed data repository. It supports storing, retrieving, and managing data via CID, ensuring secure data referencing within the distributed environment.

In addition, a Web Service is designed and developed to serve as an intermediary between the IPFS-based data storage and the permissioned blockchain. It serves as a unified method of secure communication, streamlining communication by handling two client interactions. The Web service reduces integration complexity by accepting data submissions and storing them directly on IPFS nodes. Subsequently, it packages the acquired CIDs, along with domain-specific metadata and other structured transaction payloads, and transmits all to the blockchain ledger.

4 Evaluation and Discussion

This section presents the test methodology, performance indicators, and a full breakdown of the experimental results, assessing the scalability, reliability, and efficiency of the proposed architecture. The evaluation has been built upon the two core distributed components, permissioned blockchain and IPFS, of the distributed data sharing framework to assess their unique strengths in providing tamper-resistant auditable functionality and tamper-evident storage. It therefore simulates real operational loads and responses across various scenarios to identify potential bottlenecks when deployed in a multi-domain data-sharing environment. The proof-of-concept has been configured with the following technical specification: a 3.40 GHz quad-core Intel Core i7 4th-generation processor, 16 GB of RAM and a 512 GB SSD. The evaluation aims to determine the latency, throughput, and resource utilisation of the blockchain component, and the upload and retrieval throughput and resource utilisation of the IPFS component.

4.1 Blockchain Performance Analysis

To evaluate the performance of the proposed blockchain audit tool, systematic experiments were conducted to assess responsiveness, transaction-handling capacity, and resource utilisation under varying loads. The evaluation assesses how the system performs when it receives transactions in batches of 1,000, 3,000, 5,000, and 10,000. These transactions were generated across 10 representative domains, including Energy, Health, Automotive, Space, Education, Finance, Transport, Supply Chain, Agriculture and Defence. For each batch, transactions were randomly assigned to these domains and and submitted using multithreaded execution with 80 concurrent threads to simulate concurrent transaction activity. Every transaction payload was synthetically generated using a 20-field schema, including unique

metadata fields (such as date/time, domain tags, content hashes, and text fields).

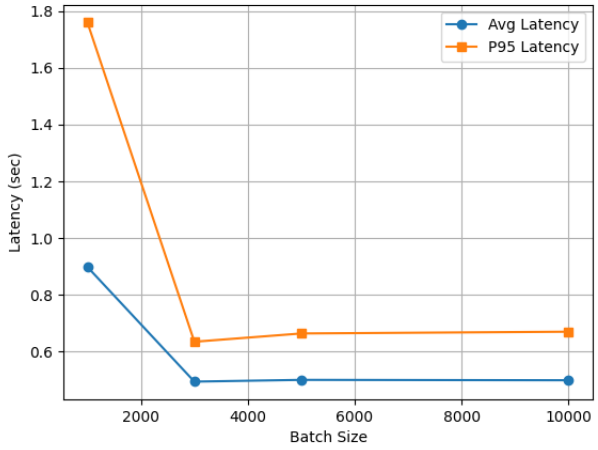

Figure 6. Transaction latency across batch sizes.

Figure 6 measures the average latency and the P95 latency for each batch of transactions recorded to the ledger. As the batch size increases, the average latency stays below ∼1 s, and the 95th-percentile (P95) latency also remains within a bounded range. This stability indicates its ability to process large, complex transaction sets, along with greater responsiveness and scalability.

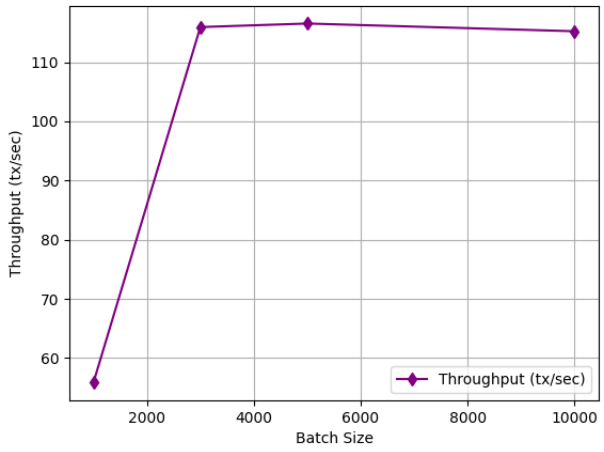

Figure 7. Transaction throughput across batch sizes.

Figure 7 shows the number of transactions the system can handle per second. It demonstrates that throughput improves as batch size increases, reaching approximately 116 tx/s for larger batches. This indicates its capacity to record larger transaction volumes more efficiently when deployed for large-scale applications.

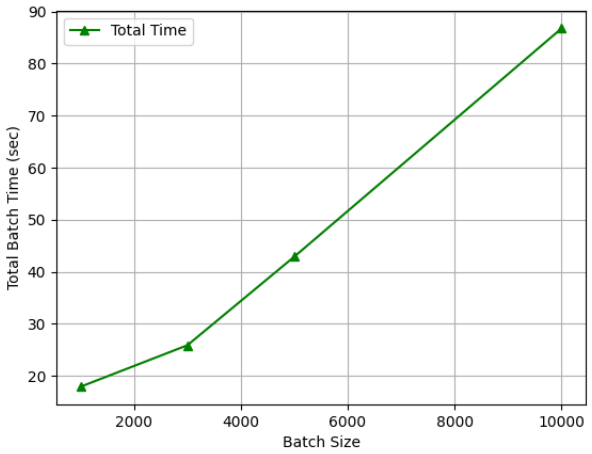

Figure 8. Total processing time across batch sizes.

Figure 8 represents the total time taken to finish recording all transactions for each batch. As the batch size increases, the total time increases linearly and in a controlled manner, demonstrating that the system functions as predicted and avoids performance degradation under higher loads.

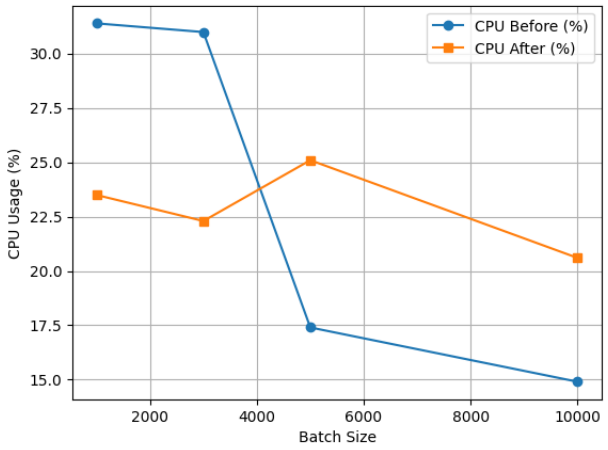

Figure 9. CPU usage before and after each batch.

The CPU utilisation in Figure 9 is recorded immediately before and after each batch execution. The measurement provides a comparison of pre- and post execution CPU usage across different batch sizes, showing resource utilisation changes around each workload. The objective of this experiment was to address the immediate resource impact across increasing workload sizes rather than to profile transient CPU peaks during execution.

Figure 10 demonstrates the memory load recorded before and after each transaction batch is submitted to the ledger. In a manner similar to CPU usage, memory

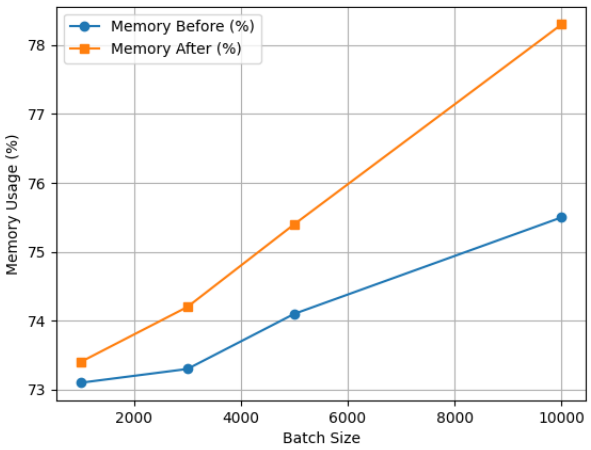

Figure 10. Memory usage (%) before and after each batch.

usage also increases when executing larger batches, but remains within acceptable bounds under the tested load.

4.2 Distributed Repository Evaluation

To evaluate the scalability and efficiency of the IPFS-based distributed data repository, a series of upload and retrieval experiments were carried out on *'Node 1'* using 1 MB randomly generated JavaScript Object Notation (JSON) files, executed in different batch sizes (10, 20, 50, 100, 200, 300, 400, and 500). Each batch size corresponds to the number of distinct files uploaded (or retrieved) simultaneously. For instance, a batch size of 500 refers to 500 parallel uploads or retrievals at the same time, reflecting a real-world scenario in which multiple domains access the system to store their datasets or query records.

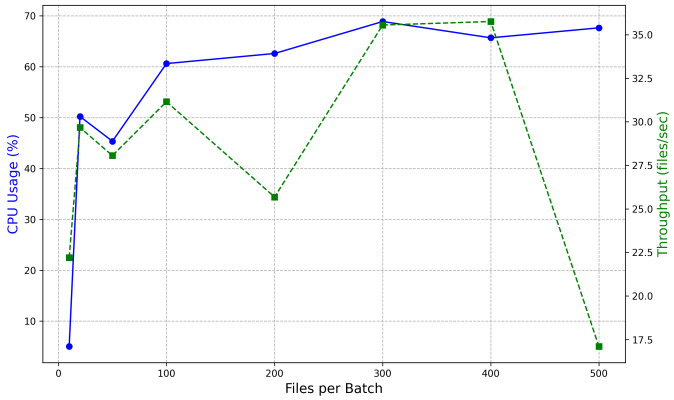

Figure 11. CPU usage and upload throughput for different batch sizes.

Figure 11 depicts the correlation between CPU usage and content upload throughput while batch size increases. The CPU usage rises steadily with batch size from 5% at batch size 10 to nearly ∼68% for 500 parallel uploads. The system achieves a peak upload throughput of approximately 35 files/s at a batch size of 300, performing efficiently at low to medium (100–400) batch sizes, while the drop to ∼17 files/s at batch size 500 may be due to resource saturation (e.g. disk I/O) or operating system-level throttling, since CPU usage alone is not yet saturated.

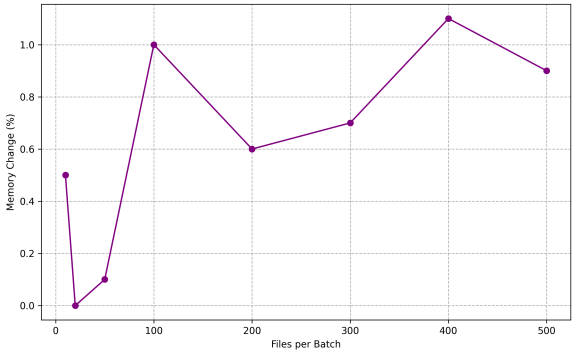

Figure 12. Memory usage during uploads for different batch sizes.

Figure 12 marks the memory change (%) during uploads with increasing file batch sizes. The stability of memory usage without memory spikes across all test cases (with a maximum increase of approximately 1.1%) indicates that the system is memory-efficient during concurrent uploads under the tested workloads.

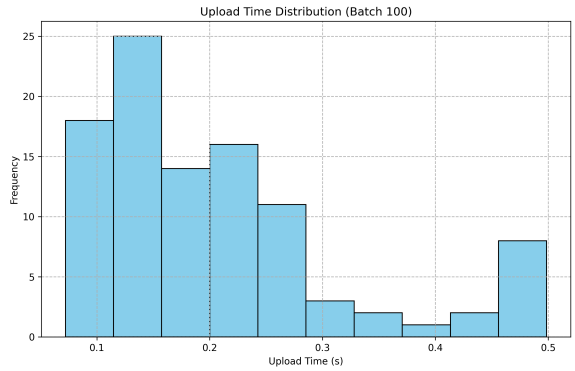

Figure 13. Upload time distribution for a batch size of 100

Figure 13 presents the distribution of upload times in seconds. The experiment shows how long each file, in a batch of 100, took to upload. Most uploads took 0.2 seconds, indicating that the upload process is consistently fast and well optimised for environments where concurrent uploads are necessary.

Figure 14(a) assesses the efficiency in fetching batches

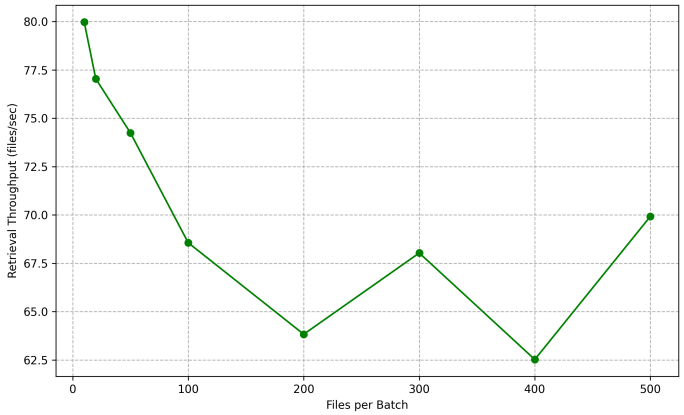

(a) Retrieval throughput across different batch sizes of files

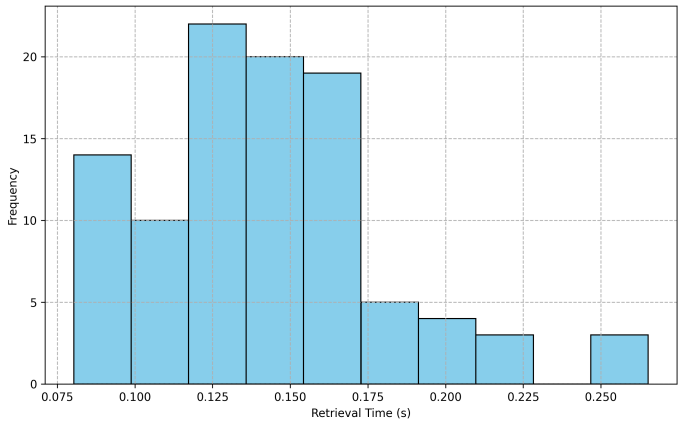

(b) Distribution of retrieval times for a batch size of 100

Figure 14. Retrieval performance of the distributed repository.

of previously uploaded contents in parallel. As the retrieval throughput remains consistently high, varying between ∼63 and 79 files per second across all batch sizes, this indicates that the system processes concurrent downloads efficiently and is thus effective in content distribution across domains.

In Figure 14(b), the individual retrieval times for a batch size of 100 have been measured. The batch size of 100 indicates that 100 content retrieval requests were executed simultaneously, reflecting a real-world use case in which multiple queries will be made. The experiment shows good responsiveness, with a large portion of files fetched in under 0.15 seconds. The tight clustering indicates low variance in retrieval time. This suggests that the framework is responsive and consistent under this workload.

Taken together, the experimental evaluation conducted provides an extensive view of the proposed framework under various loads. These results together indicate the efficiency and scalability, within the tested range, that are necessary for deployment in multi-domain settings (reliability under fault conditions was not evaluated) where real-time, concurrent, and secure transactions, as well as data handling, are essential.

4.3 Discussion

This work provides a comparative, empirical analysis to underpin the implementation of a secure, distributed data-sharing architecture that connects multiple domains. The comparative analysis supported the selection of the technologies best suited for achieving scalable, access-controlled, and verifiable data exchange. The integration of Hyperledger Fabric with IPFS, supported by a web service and a custom API layer, has demonstrated the efficiency of blockchain-driven auditability and content-addressed distributed storage without compromising performance. The experimental results have shown minimal latency when the batch size increased during recording a significant volume of transactions. Besides, the IPFS-based repository has shown consistent CPU and memory usage movements with high-volume upload and retrieval options. These results indicate that the framework introduces only moderate computational overhead under the tested loads (no baseline without the framework was measured), and support the feasibility of combining the permissioned blockchain with a P2P distributed data repository for scalable data storage with enhanced transparency while maintaining operational efficiency.

The key strength of this work has been its comprehensive comparative foundation, followed by a validated implementation that bridges emerging technologies. Our method offers modularity, transparency, and interoperability across domains, enabling cross-domain data exchange to be handled securely and efficiently. While the experiments indicate reliable operation under evaluated workloads, they did not include controlled fault testing or adversarial scenarios, as the broader security considerations of the framework were explicitly discussed in our earlier work [13]. These results therefore concentrate on the scope of robustness and reliability. Additionally, the evaluation was conducted with limited network peers and in a controlled testing environment, and thus broader distributed testing would enable further evaluation of resilience under diverse network conditions.

5 Conclusion

This work presents the architectural design and implementation of a distributed audit and data-sharing framework. It has been developed as part of the Trusted Mediator Platform, providing a secure, privacy-preserving data infrastructure across multiple sectors. Through our comparative analysis, we have highlighted the most competent technologies essential for eliminating traditional security breaches, enhancing security, and ensuring compliance. The development of the proposed method through the effective integration of blockchain and IPFS has ensured tamper-resistant transaction recording and efficient data storage. The experimental evaluation provided insights into the proposed system and its operational elements across different scenarios. Despite increasing parallelism, the system demonstrated consistent behaviour, indicating its efficiency across multiple domains. The experimental observations have indicated its performance in scalability and responsiveness under the tested workloads, as well as its potential to support a multi-user environment in real-world deployment.

Future directions could focus on stronger privacy features, such as anonymous credentials. In particular, emphasis should be given to integrating privacy-preserving identity mechanisms and attribute-based authentication. Future work should also extend the experimental evaluation to larger loads beyond the throughput drop observed at 500 parallel uploads, to multi-host and multi-orderer deployments, and to fault-injection and adversarial scenarios.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported by the Horizon Europe Project Trust & Privacy Preserving Computing Platform for Cross-Border Federation of Data (TRUSTEE) under Grant 101070214.

Conflicts of Interest

Pavlos Papadopoulos and Nikolaos Pitropakis served as Associate Editors of *ICCK Transactions on Information Security and Cryptography* at the time of manuscript submission. To ensure the integrity of the peer-review process, neither Pavlos Papadopoulos nor Nikolaos Pitropakis was involved in the editorial handling, peer review, or decision-making process for this manuscript. The manuscript was handled independently by another editor. The remaining authors declare no conflicts of interest.

AI Use Statement

The authors declare that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] Ahmad, A., Saad, M., Al Ghamdi, M., Nyang, D., & Mohaisen, D. (2021). Blocktrail: A service for secure and transparent blockchain-driven audit trails. *IEEE Systems Journal, 16*(1), 1367-1378. [CrossRef]

[2] Sun, Y., Zhang, J., Xiong, Y., & Zhu, G. (2014). Data Security and Privacy in Cloud Computing. *International Journal of Distributed Sensor Networks, 10*(7), 190903. [CrossRef]

[3] IBM. (2026). *Cost of a Data Breach Report 2026*. Retrieved from https://www.ibm.com/reports/data-breach

[4] Sobers, R. (2025). *82 Must-Know Data Breach Statistics*. Varonis. Retrieved from https://www.varonis.com/blog/data-breach-statistics

[5] Dhaliwal, J. (2025). *Data Breach Exposes 3 Billion Personal Information Records*. McAfee Blog. Retrieved from https://www.mcafee.com/blogs/security-news/data-breach-exposes-3-billion-personal-information-records/

[6] Whittaker, Z. (2024). *The biggest data breaches in 2024: 1 billion stolen records and rising*. TechCrunch. Retrieved from https://techcrunch.com/2024/10/14/2024-in-data-breaches-1-billion-stolen-records-and-rising/

[7] Goel, A., & Rahulamathavan, Y. (2025). A comparative survey of centralised and decentralised identity management systems: analysing scalability, security, and feasibility. *Future Internet, 17*(1), 1. [CrossRef]

[8] Chen, Y., Li, H., Li, K., & Zhang, J. (2017). An improved P2P file system scheme based on IPFS and Blockchain. In *2017 IEEE International Conference on Big Data (Big Data)* (pp. 2652-2657). [CrossRef]

[9] Lo, O., Buchanan, W. J., Sayeed, S., Papadopoulos, P., Pitropakis, N., & Chrysoulas, C. (2022). GLASS: A Citizen-Centric Distributed Data-Sharing Model within an e-Governance Architecture. *Sensors, 22*(6), 2291. [CrossRef]

[10] Benet, J. (2014). IPFS - Content Addressed, Versioned, P2P File System. *arXiv preprint arXiv:1407.3561*. [CrossRef]

[11] Li, X., Jiang, P., Chen, T., Luo, X., & Wen, Q. (2020). A

survey on the security of blockchain systems. *Future Generation Computer Systems, 107*, 841-853. [CrossRef]

[12] Saad, M., Spaulding, J., Njilla, L., Kamhoua, C., Shetty, S., Nyang, D., & Mohaisen, D. (2020). Exploring the Attack Surface of Blockchain: A Comprehensive Survey. *IEEE Communications Surveys & Tutorials, 22*(3), 1977-2008. [CrossRef]

[13] Sayeed, S., Pitropakis, N., Buchanan, W. J., Markakis, E., Papatsaroucha, D., & Politis, I. (2023, August). TRUSTEE: Towards the creation of secure, trustworthy and privacy-preserving framework. In *Proceedings of the 18th International Conference on Availability, Reliability and Security* (pp. 1-10). [CrossRef]

[14] Athanere, S., & Thakur, R. (2022). Blockchain based hierarchical semi-decentralized approach using IPFS for secure and efficient data sharing. *Journal of King Saud University - Computer and Information Sciences, 34*(4), 1523-1534. [CrossRef]

[15] Gorenflo, C., Lee, S., Golab, L., & Keshav, S. (2020). FastFabric: Scaling hyperledger fabric to 20 000 transactions per second. *International Journal of Network Management, 30*(5), e2099. [CrossRef]

[16] Ward, M. (2018). *Transactions Per Second (TPS)*. Corda blog (Medium). Retrieved from https://medium.com/corda/transactions-per-second-tps-de3fb55d60e3

[17] MultiChain. (2025). *Tips for performance optimization*. Retrieved from https://www.multichain.com/developers/performance-optimization

[18] Søgaard, J. S., Eklund, P. W., Herskind, L., & Spasovski, J. (2023). Improving trust in a (trans) national invoicing system: the performance of crash vs. byzantine fault tolerance at scale. *Applied Sciences, 13*(12), 6941. [CrossRef]

[19] Christodoulou, K., Iosif, E., Inglezakis, A., & Themistocleous, M. (2020). Consensus crash testing: Exploring ripple's decentralization degree in adversarial environments. *Future Internet, 12*(3), 53. [CrossRef]

[20] Alt, R., & Gräser, M. (2025). Distributed ledger technology. *Electronic Markets, 35*(1), 53. [CrossRef]

[21] El Ioini, N., & Pahl, C. (2018, October). A review of distributed ledger technologies. In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"* (pp. 277-288). Cham: Springer International Publishing. [CrossRef]

[22] Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and informatics, 36*, 55-81. [CrossRef]

[23] Yaga, D., Mell, P., Roby, N., & Scarfone, K. (2018). *Blockchain Technology Overview* (NISTIR 8202). National Institute of Standards and Technology. [CrossRef]

[24] Polge, J., Robert, J., & Le Traon, Y. (2021). Permissioned blockchain frameworks in the industry: A comparison. *ICT Express, 7*(2), 229-233. [CrossRef]

[25] Huang, H., Lin, J., Zheng, B., Zheng, Z., & Bian, J. (2020). When blockchain meets distributed file systems: An overview, challenges, and open issues. *IEEE Access, 8*, 50574-50586. [CrossRef]

[26] Sun, J., Yao, X., Wang, S., & Wu, Y. (2020). Blockchain-based secure storage and access scheme for electronic medical records in IPFS. *IEEE access, 8*, 59389-59401. [CrossRef]

[27] Samonas, S., & Coss, D. (2014). The CIA Strikes Back: Redefining Confidentiality, Integrity and Availability in Security. *Journal of Information System Security, 10*(3), 21-45. https://www.proso.com/dl/Samonas.pdf

[28] Putz, B., Böhm, F., & Pernul, G. (2021, June). HyperSec: Visual Analytics for blockchain security monitoring. In *IFIP International Conference on ICT Systems Security and Privacy Protection* (pp. 165-180). Cham: Springer International Publishing. [CrossRef]

[29] Li, H., Mi, X., Dou, Y., & Guo, S. (2023, June). An empirical study of storj dcs: Ecosystem, performance, and security. In *2023 IEEE/ACM 31st International Symposium on Quality of Service (IWQoS)* (pp. 1-10). IEEE. [CrossRef]

[30] Reno, S., Roy, K., & Tabassum, T. (2025). Blockweave: An Arweave-Based Decentralized Storage Solution to Tackle the Blockchain Trilemma. *Engineering Reports, 7*(7), e70259. [CrossRef]

[31] Filecoin. (2024). *Network Performance*. Filecoin Docs. Retrieved from https://docs.filecoin.io/networks/mainnet/network-performance

[32] Austria, P., Park, C. H., Hoffman, A., & Kim, Y. (2021). Performance and Cost Analysis of Sia, a Blockchain-Based Storage Platform. In *2021 IEEE/ACIS 6th International Conference on Big Data, Cloud Computing, and Data Science (BCD)* (pp. 98-103). [CrossRef]

[33] Kumari, K. A., Sangeetha, S., Rajeevan, V., Dharshini, M. D., & Haritha, T. (2022, December). Trade management system using R3 Corda blockchain. In *International Conference on Intelligent Systems Design and Applications* (pp. 257-275). Cham: Springer Nature Switzerland. [CrossRef]

[34] Brown, R. G., Carlyle, J., Grigg, I., & Hearn, M. (2016). Corda: an introduction. *R3 CEV, August, 1*(15), 14. https://smallake.kr/wp-content/uploads/2016/10/corda-introductory-whitepaper-final.pdf

[35] Kuzlu, M., Pipattanasomporn, M., Gurses, L., & Rahman, S. (2019, July). Performance analysis of a hyperledger fabric blockchain framework: throughput, latency and scalability. In *2019 IEEE international conference on blockchain (Blockchain)* (pp. 536-540). IEEE. [CrossRef]

[36] Stamatellis, C., Papadopoulos, P., Pitropakis, N., Katsikas, S., & Buchanan, W. J. (2020). A privacy-preserving healthcare framework using hyperledger fabric. *Sensors, 20*(22), 6587. [CrossRef]

[37] Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., ... & Yellick, J. (2018, April). Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the thirteenth EuroSys conference* (pp. 1-15). [CrossRef]

[38] Avantaggiato, M., & Gallo, P. (2019). Challenges and Opportunities using MultiChain for Real Estate. In *2019 IEEE International Black Sea Conference on Communications and Networking* (*BlackSeaCom*) (pp. 1-5). [CrossRef]

[39] Mazzoni, M., Corradi, A., & Di Nicola, V. (2022). Performance evaluation of permissioned blockchains for financial applications: The ConsenSys Quorum case study. *Blockchain: Research and applications, 3*(1), 100026. [CrossRef]

[40] Baliga, A., Subhod, I., Kamat, P., & Chatterjee, S. (2018). Performance evaluation of the quorum blockchain platform. *arXiv preprint arXiv:1809.03421*. [CrossRef]

[41] Moreno-Sanchez, P., Zafar, M. B., & Kate, A. (2016). Listening to Whispers of Ripple: Linking Wallets and Deanonymizing Transactions in the Ripple Network. *Proceedings on Privacy Enhancing Technologies, 2016*(4), 436-453. [CrossRef]

[42] Armknecht, F., Karame, G. O., Mandal, A., Youssef, F., & Zenner, E. (2015, August). Ripple: Overview and outlook. In *International conference on trust and trustworthy computing* (pp. 163-180). Cham: Springer International Publishing. [CrossRef]

[43] Andrian, Y., Kim, H., & Ju, H. (2019). A Distributed File-Based Storage System for Improving High Availability of Space Weather Data. *Applied Sciences, 9*(23), 5024. [CrossRef]

[44] Daniel, E., & Tschorsch, F. (2022). IPFS and friends: A qualitative comparison of next generation peer-to-peer data networks. *IEEE Communications Surveys & Tutorials, 24*(1), 31-52. [CrossRef]

[45] Merlec, M. M., & In, H. P. (2024). Blockchain-Based Decentralized Storage Systems for Sustainable Data Self-Sovereignty: A Comparative Study. *Sustainability, 16*(17), 7671. [CrossRef]

[46] Zheng, P., Zheng, Z., Luo, X., Chen, X., & Liu, X. (2018). A detailed and real-time performance monitoring framework for blockchain systems. In *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice* (pp. 134-143). [CrossRef]

[47] Hyperledger Labs. (2026). *Hyperledger Explorer* (*blockchain-explorer*) [Software repository, archived 5 February 2026]. Retrieved from https://github.com/hyperledger-labs/blockchain-explorer

[48] Wen, F., Wang, Z., Qu, L., Huang, H., & Hu, X. (2024). Enhancing secure multi-group data sharing through integration of IPFS and hyperledger fabric. *PeerJ Computer Science, 10*, e1962. [CrossRef]

[49] Zheng, Z., Xie, S., Dai, H. N., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International journal of web and grid services, 14*(4), 352-375. [CrossRef]

[50] Meeks, E., & Dufour, A.-M. (2024). *D3.js in Action* (3rd ed.). Manning Publications. https://www.manning.com/books/d3js-in-action-third-edition

[51] Kaushal, R. K., & Kumar, N. (2024, March). Exploring hyperledger caliper benchmarking tool to measure the performance of blockchain based solutions. In *2024 11th international conference on reliability, infocom technologies and optimization* (*trends and future directions*)(*ICRITO*) (pp. 1-6). IEEE. [CrossRef]

[52] Dabbagh, M., Choo, K.-K. R., Beheshti, A., Tahir, M., & Safa, N. S. (2021). A survey of empirical performance evaluation of permissioned blockchain platforms: Challenges and opportunities. *Computers & Security, 100*, 102078. [CrossRef]

[53] Kang, P., Yang, W., & Zheng, J. (2022). Blockchain private file storage-sharing method based on IPFS. *Sensors, 22*(14), 5100. [CrossRef]

[54] Wang, Y., Yin, X., Zhu, H., & Hei, X. (2020, July). A blockchain based distributed storage system for knowledge graph security. In *International Conference on Artificial Intelligence and Security* (pp. 318-327). Cham: Springer International Publishing. [CrossRef]

[55] Azbeg, K., Ouchetto, O., & Andaloussi, S. J. (2022). BlockMedCare: A healthcare system based on IoT, Blockchain and IPFS for data management security. *Egyptian informatics journal, 23*(2), 329-343. [CrossRef]

[56] Ye, H., & Park, S. (2021). Reliable vehicle data storage using blockchain and IPFS. *Electronics, 10*(10), 1130. [CrossRef]

[57] Pawar, M. K., Patil, P., Sharma, M., & Chalageri, M. (2021, June). Secure and scalable decentralized supply chain management using ethereum and ipfs platform. In *2021 International Conference on Intelligent Technologies* (*CONIT*) (pp. 1-5). IEEE. [CrossRef]

[58] Wang, P., & Susilo, W. (2021). Data Security Storage Model of the Internet of Things Based on Blockchain. *Computer Systems Science and Engineering, 36*(1), 213–224. [CrossRef]

[59] Jayabalan, J., & Jeyanthi, N. (2022). Scalable blockchain model using off-chain IPFS storage for healthcare data security and privacy. *Journal of Parallel and distributed computing, 164*, 152-167. [CrossRef]

[60] Tovanich, N., Heulot, N., Fekete, J. D., & Isenberg, P. (2019). Visualization of blockchain data: A systematic review. *IEEE transactions on visualization and computer graphics, 27*(7), 3135-3152. [CrossRef]

[61] Papatsaroucha, D., Pityanou, K., Sayeed, S., Papadopoulos, P., Kasimatis, D., Zotou, M., ... & Markakis, E. K. (2025, October). A Greener Approach to Privacy-Aware Cross-Border Data Federation Through Knowledge Reuse. In I*nternational Conference on Digital Sovereignty* (pp. 407-419). Cham: Springer Nature Switzerland. [CrossRef]

[62] Papatsaroucha, D., Regueiro, C., Escalante, M., Martignano, M., Hillman, V., & Markakis, E. K. (2025, December). Secure and Privacy-Preserving Secondary Data Use: A Framework for Cross-Domain Computation in the Encrypted Domain. In *2025 Annual Computer Security Applications Conference Workshops* (ACSAC Workshops) (pp. 212-221). IEEE. [CrossRef]

[63] Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys (Csur)*, *51*(4), 1-35. [CrossRef]

[64] Otto, B., & Jarke, M. (2019). Designing a multi-sided data platform: findings from the International Data Spaces case: B. Otto, M. Jarke. *Electronic markets*, *29*(4), 561-580. [CrossRef]

Dr Sarwar Sayeed is a researcher at Ulster University and a Visiting Fellow at Edinburgh Napier University. He holds a Ph.D. in Computer Science from the University of the West of Scotland and an M.Sc. in Information Technology. His doctoral research focused on blockchain security and received the IEEE Spain Blockchain Initiative Award for Best Doctoral Thesis in 2022. His research interests include blockchain, cybersecurity, distributed systems, and secure data sharing. Sarwar's current research focuses on developing and evaluating secure data-driven solutions for supply chain traceability and digital transformation. He has contributed to several major UK- and EU-funded research and innovation projects, including TRUSTEE, GLASS, and the Safety Tech Challenge Fund, where he worked on distributed architectures for secure data sharing and digital identity. He has authored numerous peer-reviewed journal and conference publications in blockchain and cybersecurity. He has been recognised as a Global Talent by UK Research and Innovation (UKRI) and also received the First Prize in Academic Excellence. (Email: S.Sayeed2@napier.ac.uk)

Dr Pavlos Papadopoulos is a Lecturer in Cyber Security at Edinburgh Napier University and the Recruitment Lead of the Cyber Security and Systems Engineering subject group. He received his PhD in Privacy-Preserving Systems around Security, Trust and Identity in 2022, and an MSc in Cyber Security (Advanced Security and Digital Forensics) in 2019 from Edinburgh Napier University. He also holds a Bachelor's Degree in Computer Science (Department of Digital Systems) from the University of Piraeus in Greece. Pavlos is an Enterprise Fellow of the Royal Academy of Engineering, an Associate Fellow of the Higher Education Academy (AFHEA), a Senior Member of the Institute of Electrical and Electronics Engineers (SMIEEE), a Member of the Institution of Engineering and Technology (MIET), Certified in Cybersecurity and a Member of the International Information System Security Certification Consortium (ISC2), and a Member of the Blockpass Identity Lab (BIL). Pavlos' research interests focus on cybersecurity, distributed ledger technology, privacy-preserving machine learning, AI security, software security, and related topics. He has over 30 peer-reviewed publications and has been the Principal Investigator on projects exceeding £850,000 in competitive research and innovation funding from EU Horizon, Innovate UK, Scottish Enterprise, and the Royal Academy of Engineering. Additionally, Pavlos received the "Outstanding Young Person in Cyber" award at the Scottish Cyber Awards 2022, and his expertise led him to found TrueDeploy, which received the "Start Up of the Year" award at the Scottish Cyber Awards 2025. (Email: P.Papadopoulos@napier.ac.uk)

Dimitrios Kasimatis received the B.Sc. degree in digital systems from the University of Piraeus, Greece, and the M.Sc degree with distinction in advanced security and digital forensics in 2023 from Edinburgh Napier University, U.K. He is currently completing his Ph.D. at Edinburgh Napier University. He received a conditional pass with minor corrections following the successful defense of his doctoral thesis in July 2026. He has 10 publications in international journals, conferences, and books. His research focuses on cybersecurity, digital trust, digital identities, AI security, and related topics. He is a Fellow of Advance HE (FHEA) (formerly the Higher Education Academy) and a member of the Blockpass Identity Lab (BIL). (Email: d.kasimatis@napier.ac.uk)

Dr Nikolaos Pitropakis received the degree in computer science from the University of Athens, Greece, in 2009, and the M.Sc. degree in advanced information systems from the Athens University of Economics and Business in 2011. He is currently an Associate Professor with the Hellenic American College, where he co-leads the Cybersecurity Research Laboratory and also serves as a Visiting Associate Professor at Edinburgh Napier University. His research interests include adversarial machine learning, trust and privacy in distributed ledger technologies, cyber attack attribution, and the application of data science to cybersecurity and the IoT device security. He was a Lead External Examiner for the M.Sc. Cyber Security Program, Northumbria University. With more than 90 peer-reviewed publications and several prestigious awards, he has made substantial contributions to advancing cybersecurity research and practice. He is a fellow of the Higher Education Academy. He is actively engaged in organising major international conferences and contributes to the field as an editor for leading academic journals. (Email: npitropakis@acg.edu)

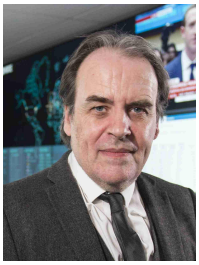

Prof William (Bill) J. Buchanan is a Professor of applied cryptography in the School of Computing, Engineering and the Built Environment, at Edinburgh Napier University. He was appointed as an Officer of the Order of the British Empire (OBE) in the 2017 Birthday Honours for services to cybersecurity. His true academic passion is building, analyzing, and breaking cryptographic methods. He also has a strong belief in the power of education

and in supporting innovation from every angle and currently leads the Scottish Centre of Excellence in Digital Trust and DLT, the Blockpass ID Lab and the Centre for Cybersecurity, IoT and Cyberphysical. He works in the areas of cryptography, blockchain, trust, and digital identity. He has one of the most extensive cryptography sites in the World (asecuritysite.com), and is involved in many areas of novel research and teaching. He has published more than 30 academic books and more than 450 academic research articles. Along with this, his work has led to many areas of impact, including a number of highly successful spin-out companies (including Zonefox, Symphonic Software and Cyan Forensics), along with awards for excellence in knowledge transfer and for teaching. He was appointed as a fellow of the Royal Society of Edinburgh (FRSE), in 2024. He recently received an "Outstanding Contribution to Knowledge Exchange" award and was included in the FutureScot "50 Scottish Tech People Who Are Changing The World", along with being a regular keynote speaker at a range of conferences. (Email: b.buchanan@napier.ac.uk)

Dr Evangelos K. Markakis is an Assistant Professor in the Department of Electrical and Computer Engineering at the Hellenic Mediterranean University and Principal Investigator and Director of the Pasiphae R&D Laboratory. He holds a Ph.D. in peer-to-peer constellations in broadband networks from the University of the Aegean. His research focuses on cybersecurity, secure communications, post-quantum cryptography, secure software engineering, cyber-physical systems, and critical infrastructure protection. He has authored over 150 scientific publications and participated in more than 60 European R&I projects across FP5, FP6, FP7, Horizon 2020, and Horizon Europe, serving in several as Coordinator, Technical Coordinator, or key scientific contributor. His work combines academic research and applied cybersecurity engineering, with particular emphasis on practical, interoperable, and standards-aware security solutions. He is also the Chairman of the Pan-European Emergency Application Consortium, contributing to major initiatives in emergency communications. (Email: emarkakis@hmu.gr)

Dimitra Papatsaroucha is a Cybersecurity Associate Researcher in the Department of Electrical and Computer Engineering at the Hellenic Mediterranean University and Senior Research Project Manager and Lab Team Leader at the Pasiphae R&D Laboratory. She is currently pursuing a Ph.D. focused on Human Vulnerability Assessment and human-centered cybersecurity. She has contributed to more than 7 European R&I projects across Horizon 2020 and Horizon Europe, serving in roles including Deputy Technical Coordinator and key scientific contributor in activities related to cybersecurity architectures, trusted digital ecosystems, privacy-preserving systems, and cyber resilience. She has co-authored more than 20 scientific publications and presented the findings of her work at international conferences and scientific events. Her work combines applied cybersecurity research, secure system design, privacy-preserving technologies, and human-centered approaches to cybersecurity and resilient digital infrastructures. (Email: d.papatsaroucha@pasiphae.eu)

Konstantina Pityanou received the B.Sc. degree in Informatics Engineering–Software Engineering from the Technological Educational Institute of Crete and the M.Sc. degree in Informatics Engineering from the Hellenic Mediterranean University. She has served as a key scientific contributor across multiple EU-funded research projects, including Horizon TRUSTEE, EMERGE, SUNRISE, and CSSBoost. Her research interests include cybersecurity best practices, browser technologies, server-side technologies, database management, API development, and User Experience and User Interface design. Her skills encompass both front-end and back-end development, reflecting her commitment to supporting innovative research efforts. (Email: k.pityanou@pasiphae.eu)

Dr Ilias Politis is a Research Associate at the Industrial Systems Institute of the ATHENA Research and Innovation Center, Greece. His research focuses on cybersecurity, privacy, trust, secure communications, and emerging technologies, including quantum-safe security and privacy-preserving artificial intelligence. He has authored more than 90 peer-reviewed journal and conference publications and has contributed to several national and European research projects. (Email: politis@athenarc.gr)