Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

1. Executive Summary

This draft report sets out to guide organisations, particularly small and medium-sized enterprises (SMEs), in managing blockchain data effectively by classifying information either on-ledger (placed directly on the blockchain) or off-ledger (stored externally but cryptographically linked). The research pursues three main objectives: to define clear standards for data classification, to align those standards with existing regulations and industry needs, and to offer practical recommendations that enhance security, interoperability, and compliance.

Importance of On-Ledger and Off-Ledger Data Classification
Distinguishing on-ledger from off-ledger information is essential to achieving a workable balance between transparency, cost-effectiveness, and legal obligations. On-ledger data is best suited for transactions and records requiring public trust and traceability. However, storing large or sensitive datasets directly on the blockchain can limit performance, increase costs, and pose privacy risks. Off-ledger storage, therefore, is recommended for data of considerable volume, confidential information, or content subject to change or erasure requirements - particularly where regulations like the GDPR demand flexible data handling.

Key Findings

1. **Scalability and Cost**: Placing excessive volumes of information on-ledger strains network capacity and raises transaction fees. Off-ledger storage, anchored by cryptographic proofs, alleviates these burdens while preserving essential verifiability.
2. **Security and Privacy**: Although blockchain's immutability fosters data integrity, enterprises need robust cryptographic methods to protect off-ledger data from tampering. Balancing transparency with privacy remains a challenge, especially in public networks.
3. **Interoperability**: Many blockchains and legacy systems struggle to exchange information reliably. Without cohesive standards, organisations face fragmentation and reduced trust when combining multiple platforms or integrating off-ledger solutions.
4. **Compliance Risks**: Regulations such as GDPR and industry-specific rules require careful oversight of personal and sensitive data. The obligation to erase or modify records can conflict with blockchain's permanence unless suitable off-ledger methods are in place.

Proposed Recommendations
The report advocates a structured approach to data classification, favouring on-ledger storage only for critical, high-trust information and off-ledger solutions for all else. It encourages well-defined security protocols - hashing, digital signatures, and encryption - to link off-ledger data to blockchain transactions reliably. Strong interoperability standards, supported by international bodies and industry consortia, are also recommended to harmonise data exchange across disparate networks.

Guidelines for SMEs
Small and medium-sized firms often have limited technical capacity and budgets. The report outlines straightforward decision frameworks to simplify adoption: store minimal, essential references on the

blockchain for traceability while retaining large or confidential content off-ledger. It also points to ready-made toolkits, clearer documentation, and user-friendly interfaces to ease the learning curve and reduce blockchain implementation costs.

Future Directions
Ongoing research areas include refining privacy-preserving cryptographic tools (e.g., zero-knowledge proofs) to secure sensitive data on-ledger, enhancing cross-chain communication protocols for broader interoperability, and codifying best practices into internationally recognised standards. By continuing these efforts, the sector will be better positioned to handle growing data volumes, stringent compliance requirements, and evolving use cases across finance, supply chains, healthcare, and beyond.

2. Introduction

2.1 Background & Motivation

Blockchain technology has evolved dramatically over the past decade (Ethereum was introduced on July 30, 2015, when its blockchain went live, enabling the creation of dApps and smart contracts), offering transformative capabilities in decentralisation, security, and transparency. However, as the adoption of blockchain expands across industries - from finance and healthcare to supply chain and beyond - the effective management of blockchain data has emerged as a crucial challenge.

Why Blockchain Data Management is a Crucial Issue

Blockchain's core strength lies in its ability to ensure on-ledger data integrity and immutability, which are critical for maintaining trust and operational transparency. However, for SMEs, the challenge extends beyond the blockchain itself to managing off-ledger data - ensuring it aligns seamlessly with on-ledger records. This dual-layer complexity requires robust data management strategies to maintain accuracy, security, and accessibility across both environments.

For SMEs, effective blockchain data management means addressing key questions: How is off-ledger data validated before being written to the blockchain? How is on-ledger data accessed and utilized for day-to-day operations? And how are discrepancies between off-ledger and on-ledger data resolved? Without clear processes, SMEs risk inefficiencies, data silos, and compliance issues.

By implementing structured data management practices - such as secure data pipelines, real-time synchronization tools, and audit trails - SMEs can bridge the gap between off-ledger and on-ledger systems. This not only enhances operational efficiency but also unlocks the full potential of blockchain for applications like supply chain tracking, secure record-keeping, and automated compliance reporting. For SMEs, mastering this balance is key to leveraging blockchain as a competitive advantage in a data-driven world.

Increasing Complexity: Scalability, Security, and Regulatory Requirements

The landscape of blockchain is increasingly complex (as in the peak of the technology in 2019, there are over 1,600 unique blockchain-based solutions and applications identified across various industries, highlighting the rapid growth and diversification of blockchain technology). On one side, scalability remains a persistent concern as public blockchains often face limitations on transaction throughput and data storage. The storage of large volumes of data directly on the ledger (on-ledger data) can lead to

performance bottlenecks and escalating costs. Conversely, off-ledger storage solutions - while addressing scalability issues - introduce challenges in data verifiability and security.

Security is another pivotal concern. As blockchain networks grow, so too does the risk of cyber threats (As of early 2025, the Bitcoin blockchain has expanded beyond 570 GB while typically processing 350,000–450,000 daily transactions, sustained by a mean hashrate exceeding 450 exahashes per second. In parallel, Ethereum's mainnet surpasses 1 TB of archived state data, handles over 1 million transactions daily, and exhibits a staking pool exceeding 25 million ETH (worth tens of billions of $) distributed across nearly 900,000 active validators). Ensuring that data is securely stored, whether on-ledger or off-ledger, requires robust cryptographic techniques and continuous monitoring. Moreover, regulatory requirements are becoming more stringent. With frameworks such as the General Data Protection Regulation (GDPR) and the Markets in Crypto-Assets (MiCA) regulation, organisations must guarantee that blockchain data management practices are compliant with evolving legal standards. This calls for a balanced approach that does not compromise on transparency while safeguarding personal and sensitive information.

Impact of Emerging Blockchain Architectures and Decentralised Storage Solutions

Recent advancements in blockchain architectures have provided new opportunities as well as challenges for data management. Emerging models, such as Layer 2 and Layer 3 solutions, have been designed to alleviate on-chain congestion and enhance scalability. At the same time, decentralised storage solutions (for example, the InterPlanetary File System (IPFS)) offer cost-effective means to store large data sets off the blockchain. However, the integration of these technologies requires clearly defined standards to ensure that data stored off-ledger can be reliably linked back to on-chain records, maintaining both the verifiability and the integrity of the overall system.

2.2 Scope & Objectives

In establishing a standard for blockchain data management, it is imperative to clearly define the scope and set measurable objectives that address both technical and practical concerns.

Defining the Scope

This standard specifically addresses the management of blockchain data by distinguishing between two primary categories:

- **On-Ledger Data:** Information stored directly on the blockchain. This includes transaction records, smart contract logic, and other critical data that benefits from the immutable, decentralised nature of blockchain. On-ledger data is essential for ensuring transparency and public verifiability.
- **Off-Ledger Data:** Data stored outside of the blockchain environment but referenced within on-chain transactions via cryptographic proofs or hashes. Off-ledger data typically includes large datasets, sensitive personal information, and dynamic records that would otherwise burden the blockchain's storage capabilities. It is pivotal for maintaining efficiency and cost-effectiveness without sacrificing security.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

The standard also considers additional facets such as security protocols, interoperability among disparate blockchain systems, and the practical needs of small and medium-sized enterprises (SMEs) seeking to adopt blockchain technology.

Objectives

The primary objectives of this standard are as follows:

1. **Developing Proposed Standards for Blockchain Data Management:**
   The document aims to establish a clear and comprehensive framework that addresses the storage, classification, and secure management of both on-ledger and off-ledger data. By defining standard practices, the framework seeks to support consistency across various blockchain implementations.

2. **Ensuring Regulatory Alignment:**
   It is essential that any proposed standard adheres to existing regulatory requirements. This includes compliance with GDPR for data privacy, MiCA for crypto-assets, and the guidelines established by ISO/TC 307 for blockchain and distributed ledger technologies. Aligning with these standards ensures that the blockchain ecosystem remains legally compliant and interoperable on an international scale.

3. **Creating Practical Guidelines for SMEs:**
   One of the core goals is to lower the barrier to entry for SMEs. By providing clear, actionable guidelines, the standard will help smaller organisations adopt blockchain technology without requiring extensive technical expertise. This includes cost-effective approaches to data management and simplified processes that maintain both security and scalability.

2.3 Methodology Overview

To achieve these objectives, the development of the standard will follow a rigorous and systematic methodology. This will ensure that the standard is not only technically robust but also practically implementable.

Literature Review-Based Approach

A thorough review of current literature is the foundation of this standard. This will involve:

- **Journals and Academic Publications:**
  Analysing peer-reviewed research to understand the latest advancements and challenges in blockchain data management.

- **White Papers and Technical Reports:**
  Reviewing industry reports and documents from established blockchain projects to extract best practices and identify common pitfalls.

- **Existing Standards:**
  Examining current international standards and regulatory frameworks, including those from ISO/TC 307, GDPR, and MiCA, to ensure alignment with global practices.

Comparative Analysis of Blockchain Architectures

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

A detailed comparative analysis will be conducted to assess various blockchain architectures. This analysis will:

- Identify the strengths and weaknesses of different approaches to data storage and management, including both public and private blockchains.
- Evaluate the impact of different consensus mechanisms on data integrity and scalability.
- Consider the trade-offs between on-ledger and off-ledger storage solutions in terms of cost, performance, and security.

Technical Assessment of Security and Interoperability Protocols

The technical robustness of the proposed standard will be ensured through an in-depth assessment of security and interoperability protocols:

- **Security Assessment:**
  Evaluating the cryptographic techniques currently in use, such as digital signatures, hashing algorithms, and zero-knowledge proofs, to determine how they can be standardised across various blockchain systems.
- **Interoperability Analysis:**
  Reviewing the mechanisms by which different blockchain systems communicate and exchange data. This includes the study of cross-chain protocols, APIs, and integration methods with legacy systems. The goal is to establish unified protocols that facilitate seamless data exchange between on-ledger and off-ledger storage systems.

3. Literature Review & Theoretical Framework

Review of Blockchain Data Models

On-Ledger Data Characteristics

On-ledger data refers to information stored directly on the blockchain. It benefits from key properties such as immutability, transparency, and trust. The immutability of on-chain data ensures that once recorded, it cannot be altered, making it ideal for transaction records and audit trails. Transparency allows all network participants to verify the integrity of stored data, increasing trust and accountability. However, on-ledger storage is constrained by scalability and high costs, as blockchain transactions incur fees and require network consensus.

Off-Ledger Data Characteristics

Off-ledger data is stored externally while being cryptographically linked to on-ledger records. This approach addresses scalability limitations and enables the management of large datasets. Off-ledger storage provides flexibility and privacy, as sensitive information does not need to be stored directly on the blockchain. Common implementations include decentralised storage solutions such as IPFS, Arweave, and Filecoin, which ensure data persistence while leveraging cryptographic anchoring for verification.

Comparative Analysis of Blockchain Storage Approaches

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

Public vs. Private Blockchains in Data Storage

Public blockchains, such as Ethereum and Bitcoin, offer full transparency and decentralisation but come with high costs and scalability challenges. Private blockchains, like Hyperledger Fabric and R3 Corda, provide greater control and privacy but lack the trust model of public networks. The choice between these architectures depends on the trade-offs between decentralisation, efficiency, and compliance requirements.

Off-Ledger Decentralised Storage vs. Cloud Storage

Decentralised storage systems like IPFS, Arweave, and Filecoin provide robust mechanisms for off-ledger storage, ensuring data integrity through distributed replication and cryptographic verification. These solutions contrast with traditional cloud storage providers, such as AWS and Google Cloud, which offer scalability and high availability but rely on centralised trust models. The primary advantage of decentralised storage is its resistance to censorship and single points of failure, while cloud storage offers superior speed and reliability.

Use of Cryptographic Anchors

To maintain data integrity between on-ledger and off-ledger storage, cryptographic anchors such as Merkle trees, zk-SNARKs, and digital signatures are utilised. Merkle trees allow for efficient verification of large data sets by hashing them into a single root stored on-chain. Zero-knowledge proofs (zk-SNARKs) enable privacy-preserving verification, allowing users to prove knowledge of data without revealing it. Digital signatures authenticate transactions and ensure data integrity in hybrid blockchain architectures.

Standardisation & Existing Frameworks

Global Standards for Blockchain Data Management

Efforts to standardise blockchain data governance are led by organisations such as ISO/TC 307, IEEE P2418, ETSI, and ITU-T. These frameworks define best practices for blockchain interoperability, security, and data classification. For example, ISO/TS 23635:2023 provides governance guidelines for managing on-ledger and off-ledger data, ensuring regulatory compliance and operational efficiency.

EU's ICT Rolling Plan for Blockchain Standardisation (2023)

The European Union's ICT Rolling Plan for Blockchain Standardisation (2023) emphasises the importance of developing common data models and compliance mechanisms for blockchain applications. This initiative supports GDPR alignment and the integration of blockchain with existing IT infrastructure, facilitating adoption across industries.

Legal & Compliance Considerations

GDPR & Blockchain Immutability

One of the major legal challenges in blockchain implementation is reconciling the immutability of on-ledger data with the "Right to be Forgotten" under GDPR. Since blockchain records cannot be deleted, strategies such as off-ledger storage, encryption, and hash-based anonymisation are employed to comply with data protection regulations. Solutions like proxy re-encryption and cryptographic key management help enforce GDPR-compliant access controls.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

MiCA, FATF, and AML Directives

The Markets in Crypto-Assets (MiCA) framework, Financial Action Task Force (FATF) guidelines, and Anti-Money Laundering (AML) directives introduce strict requirements for data governance in blockchain ecosystems. These regulations impact how on-ledger and off-ledger data are managed, requiring traceability of transactions while preserving user privacy. Implementing privacy-enhancing technologies such as zero-knowledge proofs can help strike a balance between regulatory compliance and decentralised data management.

4. Technical Challenges in On-Ledger & Off-Ledger Data Management

Scalability Issues in On-Chain Data Storage

High Gas Fees & Transaction Costs (Ethereum, Bitcoin Case Studies)

Public blockchain networks face significant scalability challenges due to high transaction costs, especially in networks like Ethereum and Bitcoin.

- **Ethereum**: Ethereum's gas fees fluctuate based on network congestion. At peak times, transaction fees can reach upwards of $200 per transaction. The high cost is driven by Ethereum's execution model, where each transaction or smart contract interaction consumes computational resources.
- **Bitcoin**: Bitcoin's scalability is limited by its block size and block confirmation times, which result in slow transaction speeds (around 7 TPS). As a result, transactions can become prohibitively expensive during high-demand periods, with fees reaching over $60.
- **DeFi Mania on Ethereum (2021–2022)**: During peak Decentralised Finance (DeFi) activity and NFT minting frenzies, complex contract interactions (e.g. adding liquidity to liquidity pools, conducting multi-swap trades on decentralised exchanges, or minting NFTs) could cost anywhere from $300 to over $1,000 in gas fees. These exorbitant costs stem from Ethereum's gas model, where every execution step in a smart contract requires a specified amount of gas, leading to higher fees as network demand intensifies.
- **Memecoin & NFT Hype (Late 2023)**: In subsequent market cycles, renewed memecoin speculation and high-profile NFT drops pushed median gas prices into the hundreds of gwei, spiking individual transaction costs to upwards of $300. Even routine operations such as token transfers became prohibitively expensive, forcing many casual users to rely on Layer 2 solutions or simply avoid on-chain transactions altogether.
- **Bitcoin Fee Surges Under Heavy Demand:** Although its base transaction throughput is roughly 7 TPS, Bitcoin can experience sudden fee spikes during surges in exchange withdrawals or when new use cases (e.g. Ordinals in 2023) generate substantially higher transaction volumes. These events have seen standard Bitcoin transaction fees exceed $60, with users sometimes paying well over $100 for expedited confirmations under intense network congestion.
- **Flash Congestion Events:** Beyond extended hype cycles, short-lived "flash" events - such as highly anticipated token launches or large-scale arbitrage in DeFi - can abruptly drive fees to astronomical levels. In Ethereum's case, a single popular smart contract interaction (for instance,

a new yield-farming platform going live) can cause mempool congestion and push gas fees sharply upward within minutes, illustrating the fragility of on-chain transaction pricing under sudden loads.

Below are further illustrations of how operating a private or permissioned blockchain network can become extraordinarily expensive - both from a cloud infrastructure standpoint and in terms of other ongoing operational costs - despite the absence of public mining or open participation:

- **Enterprise Hyperledger Fabric Deployments:** In large-scale supply chain or multi-entity financial networks, each participant may run multiple nodes in a high-availability configuration (across multiple data centres or cloud regions). The complexity of Fabric's endorsement policies, combined with the need to store extensive ledgers and private data collections, can lead to steep monthly bills in the tens of thousands of dollars for compute, networking, and storage.
- **Private Corda Clusters with High Transaction Volumes:** R3 Corda's design, which employs point-to-point communication and notary nodes, can drive large infrastructure overheads for consistently high throughput applications (e.g. enterprise trade finance). Maintaining sufficient node clusters in a redundant and geographically distributed topology often results in cloud usage costs that exceed those of many public chain alternatives, sometimes reaching well above $50,000 per month once factoring in container orchestration and secure key management.
- **Data-Intensive Consortia Networks (Hybrid On/Off-Chain):** In healthcare or real estate consortia, significant volumes of medical imaging or property records need off-chain storage yet remain tightly coupled to on-chain identity or verification. When these use cases are built atop private cloud solutions with stringent regulatory compliance (HIPAA, GDPR), monthly cloud storage can run into six figures, further inflating TCO (Total Cost of Ownership) by requiring dedicated encryption key management appliances and round-the-clock compliance auditing services.

Even without public miners, the interplay between robust node redundancy, large-scale off-ledger data, and rigorous compliance requirements can cause private blockchain infrastructure to rack up extremely high operating expenses.

Cost Implications:

- **On-Chain Storage Premiums:** Storing voluminous datasets on-chain remains prohibitively expensive on public networks like Ethereum or Bitcoin; even private or permissioned blockchains can incur high data-related overhead (e.g., maintaining multiple redundant nodes for enterprise Hyperledger Fabric deployments). Consequently, organisations prioritise pushing non-essential or large data structures off-chain, limiting on-ledger storage to critical, immutable references.

- **Hybrid On/Off-Chain Approaches:** Off-chain solutions such as IPFS, Arweave, and cloud-based systems serve as cost-effective repositories for large data volumes - ranging from user-generated content to terabytes of medical or logistics data - while the blockchain anchors cryptographic proofs. This significantly reduces direct on-chain costs yet introduces added complexity in synchronising on-ledger and off-ledger states.

- **Private Network Overheads:** Even with no transaction fees in private settings, the expenses of maintaining a consortium-grade infrastructure - redundant node clusters, container

orchestration, and secure key management - can escalate to tens or hundreds of thousands of dollars monthly. Compliance requirements (e.g., HIPAA, GDPR) and high-availability standards further increase operational costs by necessitating additional layers of encryption, auditing, and data governance.

- **Infrastructure Scalability vs. Cost:** As networks expand - especially those requiring low-latency, high-throughput consensus - organisations may invest in advanced node hardware or distributed cloud architectures, driving up capital expenditure. Any cost reductions from fewer on-chain transactions may be offset by increased spending on robust off-chain storage solutions, regulatory compliance, and around-the-clock monitoring to ensure uninterrupted service.

Layer 2 Solutions (zk-Rollups, Plasma, Optimistic Rollups)

To mitigate high costs and improve scalability, Layer 2 (L2) solutions have emerged as an efficient alternative to storing all data on the main blockchain.

- **zk-Rollups**: These use zero-knowledge proofs to aggregate multiple transactions into a single proof, reducing the amount of data stored on-chain. Examples include StarkWare and zkSync.
- **Optimistic Rollups**: Unlike zk-Rollups, Optimistic Rollups assume transactions are valid by default, reducing computation costs. Fraud proofs are used to challenge invalid transactions. Arbitrum and Optimism are leading examples.
- **Plasma**: Plasma chains function as child chains to process transactions separately from the main Ethereum network, reducing congestion. However, user withdrawals can take longer due to fraud-proof mechanisms.

Trade-offs:

- While Layer 2 solutions reduce on-chain storage requirements, they introduce new challenges in terms of data availability and security, requiring **robust mechanisms to verify data integrity**.
- They still require the base layer for security and settlement, making on-chain data storage a necessity in critical operations.

Security Risks in Off-Ledger Storage

Data Integrity Risks (Tamper-Proofing Off-Ledger Data)

Unlike on-ledger data, which benefits from immutability, off-ledger storage solutions are more susceptible to tampering. Ensuring the integrity of off-ledger data remains a critical challenge.

Key Risks:

- **Data Manipulation**: Without cryptographic guarantees, off-ledger data can be altered without detection.
- **Loss of Availability**: Off-ledger data is reliant on external storage providers, which may become unavailable or compromised.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

- **Inconsistency with On-Ledger Data**: Data updates may not be synchronised correctly between on-ledger and off-ledger environments.

Solutions:

- **Cryptographic Hashing**: Storing hashes of off-ledger data on-chain ensures data integrity by allowing verification at any time.
- **Merkle Trees**: Used to efficiently verify large datasets stored off-chain while maintaining tamper-proof links to on-chain records.
- **Decentralised Storage Networks**: IPFS, Arweave, and Storj provide distributed data storage while ensuring immutability through cryptographic techniques.

Off-Chain Oracle Vulnerabilities

Many blockchain applications rely on off-chain oracles to fetch real-world data, such as price feeds, weather conditions, or IoT sensor readings. However, these introduce security vulnerabilities.

Key Risks:

- **Single Points of Failure**: If an oracle is compromised, it can manipulate the outcome of smart contracts relying on it.
- **Sybil Attacks**: Malicious entities can create multiple fake nodes to influence oracle-fed data.
- **Data Latency**: Delays in off-chain data propagation can create mismatches between on-chain transactions and real-world events.

Mitigation Strategies:

- **Decentralised Oracles**: Projects like Chainlink use multiple independent nodes to aggregate and verify data before submitting it to the blockchain.
- **Crypto-Economic Incentives**: Participants are incentivised to report accurate data by staking assets, which can be slashed in case of malicious behaviour.
- **Trusted Execution Environments (TEEs)**: Secure hardware environments (e.g., Intel SGX) ensure that off-chain computations occur in a tamper-resistant manner.

Interoperability Limitations

Cross-Chain Data Synchronisation (Polkadot XCMP, Cosmos IBC)

Cross-chain communication remains one of the biggest technical challenges in blockchain. Ensuring interoperability between different blockchain ecosystems is critical for large-scale adoption.

1. **Polkadot XCMP (Cross-Chain Message Passing)**
    a. Allows Polkadot's parachains to exchange messages securely, ensuring synchronisation of transactions and states across the ecosystem.
    b. Relies on Polkadot's central Relay Chain to validate cross-chain messages, maintaining consistency among interconnected parachains.

2. **Cosmos IBC (Inter-Blockchain Communication Protocol)**
    a. Enables heterogeneous blockchain networks in the Cosmos ecosystem (called "zones") to transfer assets and data in a trust-minimised manner.
    b. Leverages light client verification on each chain, so a zone can securely verify transactions from another zone, promoting decentralised interoperability.
3. **Atomic Swaps & Hash Time-Locked Contracts (HTLCs)**
    a. Facilitate trustless exchange of tokens between blockchains by employing cryptographic hash locks and time-based conditions.
    b. Primarily geared toward token and coin swaps; less versatile for comprehensive data synchronisation or state-sharing beyond simple asset transfers.
4. **Chainlink CCIP (Cross-Chain Interoperability Protocol)**
    a. Offers a trust-minimised bridge layer connecting heterogeneous chains for both token and data transfer.
    b. Employs decentralised oracles to verify off-chain or cross-chain information, allowing on-chain applications to trigger real-world events or synchronise with external networks.
5. **LayerZero**
    a. Provides an omnichain messaging protocol via Ultra Light Nodes (ULNs), so dApps can initiate cross-chain function calls with minimal overhead.
    b. Focuses on reducing the need for fully verifying block headers on every chain, thus improving speed and cost-effectiveness of cross-chain interactions.
6. **Axelar Network**
    a. A validator-based bridging and messaging platform that enables developers to dispatch arbitrary data or assets among various Layer 1 and Layer 2 networks.
    b. General Message Passing (GMP) supports cross-chain smart contract invocation, letting dApps coordinate states across disparate chains.
7. **Wormhole**
    a. Evolved from the Solana ecosystem to a multi-chain bridge connecting Ethereum, Binance Smart Chain, Terra, Avalanche, and more.
    b. Uses specialised guardian nodes to verify and relay messages or transactions, facilitating the transfer of tokens, NFTs, and limited data sets.
8. **Hyperledger Cacti**
    a. A modular interoperability framework (formerly known as Hyperledger Cactus), aiming to unify transaction flows between enterprise DLTs (Fabric, Corda) and public networks.
    b. Employs pluggable connectors and secure relay modules to synchronise state and execute cross-chain operations, primarily focused on permissioned environments.
9. **Quant Overledger**
    a. An enterprise-focused interoperability solution enabling applications to interact with multiple blockchains through unified APIs.
    b. Abstracts away the underlying DLT complexities, allowing consistent messaging and data exchange - both tokenised and non-tokenised - across numerous protocols.
10. **Wanchain**
    a. Provides cross-chain bridges and "Storeman" node groups that lock assets on a source chain and mint wrapped assets on a destination chain.

   b. Has extended functionality beyond simple bridging to enable certain cross-chain smart contract interactions and orchestration.

11. **NEAR Rainbow Bridge**
    a. Specialised in connecting the NEAR Protocol with Ethereum, employing on-chain light clients for each network to verify block headers trustlessly.
    b. Supports secure, decentralised asset transfers and data synchronisation between the two ecosystems, reducing reliance on centralised or custodian intermediaries.
12. **MultiChain (formerly AnySwap)**
    a. A bridging network that connects numerous EVM-compatible and non-EVM-compatible blockchains for token and NFT transfers.
    b. While primarily used for asset bridging, it also facilitates messaging capabilities for dApps that require multi-chain data flows, relying on MPC (Multi-Party Computation) nodes to validate cross-chain operations.

Challenges:

- Varying consensus mechanisms across blockchains introduce compatibility issues.
- Ensuring consistency across chains requires additional validation layers, increasing computational overhead.
- Lack of standardisation in data formats complicates seamless data transfers.

Potential Solutions:

- Adoption of universal interoperability protocols like IBC.
- Use of standardised metadata schemas to structure data consistently across blockchains.
- Implementing cryptographic bridging techniques that allow secure cross-chain data verification.

Data Standardisation Issues Across Blockchain Ecosystems

A major roadblock in interoperability is the **lack of uniform standards** for structuring and validating on-ledger and off-ledger data.

Current Problems:

- Different blockchain networks have varying data structures (e.g., UTXO model in Bitcoin vs. account-based model in Ethereum).
- No universal metadata standard for recording off-ledger references, making cross-chain queries inefficient.
- Discrepancies in smart contract languages and execution environments hinder seamless interoperability.

Efforts Towards Standardisation:

- **ISO/TC 307**: An ongoing initiative to define global standards for blockchain data management, governance, and security.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

- **Ethereum's ERC Standards**: Various token standards (e.g., ERC-20, ERC-721) aim to create consistency but remain blockchain-specific.
- **W3C's Decentralised Identifiers (DIDs)**: Provides a framework for managing identity across different blockchains, ensuring consistent data handling.

Path Forward:

- The industry must converge on a set of open standards for structuring blockchain data.
- Efforts should be made to develop middleware solutions that act as translators between different blockchain networks.
- Implementing self-sovereign identity (SSI) frameworks could help unify identity and data management across multiple ecosystems.

Proposed Standards for On-Ledger and Off-Ledger Data Management

Data Classification Framework

Decision Matrix: What Should Be Stored On-Ledger vs. Off-Ledger?

A structured decision matrix is essential for determining whether data should reside on-ledger or off-ledger. The following criteria guide this classification:

- **Transaction Frequency:** High-frequency transactions may congest a blockchain network if stored on-ledger. Repetitive data, such as IoT sensor outputs, should be stored off-ledger while maintaining cryptographic proofs on-chain.
- **Sensitivity:** Personal identifiable information (PII), intellectual property, or proprietary business data should remain off-ledger with on-chain references using cryptographic hashes.
- **Legal Compliance:** Data subject to GDPR, HIPAA, or other data protection laws must be managed carefully, often requiring off-ledger storage with strong cryptographic linkage.
- **Cost-Efficiency:** Given blockchain storage costs, large data sets such as documents or media should be stored off-ledger, with minimal metadata or references on-chain.

This framework enables organisations to balance security, privacy, compliance, and efficiency in their data storage decisions.

Interoperability & Cross-Chain Standards

Interoperable Metadata Schemas

To facilitate seamless communication between blockchain networks, standardised metadata schemas should be adopted. These schemas must include:

- **Common Data Models:** Ensuring data consistency across different blockchain implementations.
- **Cross-Blockchain Identity Management:** Leveraging decentralised identifiers (DIDs) to unify identity verification across chains.

- **Decentralised Oracles:** Securely bridging on-chain and off-chain data for smart contract execution.

Common Protocols for Off-Ledger Anchoring

Several decentralised storage solutions can complement blockchain networks by providing scalable off-ledger storage while maintaining integrity:

- **InterPlanetary File System (IPFS):** A decentralised protocol that provides content-addressed storage for immutable references.
- **Arweave:** Offers permanent, tamper-proof storage for data that requires immutability without bloating a blockchain network.
- **Hybrid Solutions:** Combining IPFS/Arweave with blockchain hashing mechanisms for off-ledger storage integrity.
- **Filecoin:** Built on top of IPFS, Filecoin adds an incentive layer for decentralised storage, allowing users to rent out storage space and ensuring data availability through cryptographic proofs.
- **Storj:** A decentralised cloud storage platform that uses encryption and sharding to distribute data across a global network of nodes, ensuring security and redundancy.
- **Sia:** A blockchain-based decentralised storage platform that allows users to rent storage space from others on the network, with data integrity ensured through smart contracts.
- **Swarm:** A decentralised storage and communication system designed for the Ethereum ecosystem, providing scalable and censorship-resistant storage for dApps.
- **BigchainDB:** A decentralised database that combines blockchain characteristics (immutability, decentralisation) with traditional database features (high throughput, low latency), suitable for off-ledger data anchoring.
- **Ocean Protocol:** Focused on data sharing and monetisation, Ocean Protocol provides tools for securely storing and accessing off-ledger data while maintaining blockchain-based integrity proofs.

Security & Privacy Best Practices

Cryptographic Proof Models

To ensure data authenticity while maintaining privacy, cryptographic proof models should be standardised:

- **Zero-Knowledge Proofs (ZKPs):** Allow verification of information without revealing underlying data, supporting compliance and privacy-preserving blockchain operations.
- **Hash Commitment Schemes:** Ensuring off-ledger data integrity by anchoring cryptographic hashes on-chain.

Secure Multi-Party Computation (sMPC)

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

For cases requiring off-ledger data processing while preserving privacy, sMPC should be implemented to:

- **Distribute computation across multiple parties** without exposing raw data.
- **Enable compliance-friendly encryption** for sensitive business logic execution.

Regulatory Compliance Guidelines

Legal Requirements for Personal Data in On-Chain Applications

To align blockchain implementations with global data protection laws, the following practices should be adopted:

- **Data Minimisation:** Only essential metadata or cryptographic hashes should be stored on-chain.
- **Right to Erasure Mechanisms:** Off-ledger storage allows compliance with regulations such as GDPR's 'right to be forgotten'.
- **Permissioned Blockchain Configurations:** Providing controlled access to data through role-based encryption.

eIDAS and Digital Identity Solutions

Blockchain-based identity verification should adhere to eIDAS and similar frameworks by implementing:

- **Verifiable Credentials:** Digitally signed and revocable proofs of identity.
- **Self-Sovereign Identity (SSI):** Empowering users to control access to their credentials without reliance on centralised authorities.

6. SME Implementation Guidelines

Small and medium-sized enterprises (SMEs) often perceive blockchain as an intricate and expensive domain, particularly where data management is concerned. Yet, with measured planning and the right technical insights, SMEs can take advantage of decentralised ledgers to enhance process efficiency, data security, and regulatory compliance. This section explains the primary barriers SMEs face and proposes practical strategies to overcome them, culminating in recommended technology stacks that cater to both on-ledger and off-ledger data needs.

6.1 Challenges Faced by SMEs in Blockchain Adoption

6.1.1 High Complexity and Cost Barriers
SMEs commonly struggle with the perceived complexity of blockchain implementation. Setting up and sustaining a decentralised ledger demands technical skills that may be unavailable in smaller organisations. Additionally, transaction and storage expenses on public blockchains (such as Ethereum) can become significant if large volumes of data are stored on-chain. The cost factor extends beyond basic transaction fees - SMEs must also budget for developer expertise, ongoing maintenance, and integration with legacy systems.

6.1.2 Limited Expertise in Data Security and Compliance
Although many SMEs appreciate the importance of strong data security and regulatory compliance, they might lack in-house specialists. Handling personal data - whether client or employee information - on a

distributed ledger exposes an enterprise to privacy obligations such as the General Data Protection Regulation (GDPR) in Europe. The potential for non-compliance is exacerbated by the immutable nature of blockchain, which can create tension with data erasure requirements. Furthermore, the Markets in Crypto-Assets Regulation (MiCA) introduces an evolving framework for digital assets that may affect SMEs issuing tokens or handling other blockchain-based financial activities.

In essence, SMEs usually require guidance and resources to tailor blockchain adoption strategies to their own data classification, security, and compliance needs. Without clarity, the challenges noted above can quickly overshadow blockchain's prospective benefits.

6.2 Practical Steps for Implementation

6.2.1 Choosing a Storage Model (On-Ledger vs. Off-Ledger)

A pivotal question for SMEs is deciding which data to store on the blockchain (on-ledger) and which data to keep in external environments (off-ledger).

1. **Identify Critical Data**
    - Determine which business data **absolutely must be transparent, auditable, and immutable**. This often includes transactional records, digital certificates, or operational milestones that stakeholders need to verify independently. Such data typically belongs on the blockchain.
    - Conversely, large files, sensitive personal information, and frequently changing records are better suited to **off-ledger storage**. Storing these directly on the blockchain would be expensive and potentially conflict with privacy requirements.
2. **Assess Scalability and Transaction Frequency**
    - SMEs dealing with a high volume of transactions (e.g., e-commerce orders, micro-finance) should consider the impact on blockchain transaction fees. Rapidly changing data can be linked via cryptographic proofs stored on-ledger, while the bulk of information remains off-chain.
    - When transaction throughput is critical, evaluate private or permissioned ledgers designed with higher performance in mind.
3. **Evaluate Legal Considerations**
    - GDPR compliance may dictate whether certain categories of personal or customer data can be stored on-chain. Where data deletion requests may arise, off-chain alternatives offer more flexibility to remove or update data in compliance with legal requirements.
    - Industries subject to MiCA or other financial regulations should ascertain whether data related to token issuances, customer wallets, or digital asset custody could be regarded as financial records, requiring special measures or reporting.

6.2.2 Deploying Decentralised Storage Solutions Securely

SMEs exploring blockchain often wish to maintain a decentralised ethos by using distributed storage

solutions, rather than strictly relying on a traditional cloud approach. Although decentralised systems can lower hosting costs and improve fault tolerance, they pose a new set of security considerations.

1. **Cryptographic Hashing and Merkle Trees**
    - Storing entire files on a distributed ledger is seldom cost-effective for smaller businesses. Instead, hash large files or data sets and embed the hashes on the blockchain. This technique preserves immutability: as soon as a file is altered, the hash changes, revealing any tampering.
    - SMEs should explore Merkle tree structures to bundle data, enabling the verification of large data sets with minimal on-chain information.
2. **Secure Off-Ledger Repositories**
    - Decentralised storage networks (e.g., IPFS or Filecoin) and private cloud-based environments (e.g., AWS, Azure) are frequently integrated with a blockchain solution. However, the security of off-ledger repositories must be tightly managed.
    - Encryption at rest and in transit is vital. SMEs should adopt robust key management policies, including role-based access and regular security assessments, to ensure only authorised team members can retrieve or modify off-chain data.
3. **Data Redundancy and Availability**
    - SMEs must carefully manage how their off-ledger data is replicated to avoid single points of failure. Decentralised networks typically rely on multiple hosts to ensure high availability, but service-level agreements (SLAs) or additional redundancy may be necessary for mission-critical data.

6.2.3 Ensuring Compliance with GDPR & MiCA

Compliance requirements must be integrated early in the planning stages of any blockchain initiative, lest SMEs risk non-compliance.

1. **Data Protection Strategies**
    - Under GDPR, individuals enjoy a right to erasure (the “right to be forgotten”), which is at odds with the blockchain’s immutability. SMEs can address this by storing personal data off-ledger in secure databases and referencing it via hashed pointers.
    - Implement data minimisation strategies, ensuring only essential data appears on the ledger. For instance, do not place raw personal information on-chain - store an encrypted reference instead.
2. **Crypto-Assets and Regulatory Frameworks**
    - MiCA introduces requirements related to market abuse, white papers for token offerings, and business conduct. SMEs developing tokenised business models (e.g., loyalty tokens, crowdfunding tokens) must grasp these regulatory nuances.

- Collaboration with legal counsel and regulators early in the project lifecycle can reduce compliance burdens later and mitigate the need for major architectural revisions.

3. **Audit and Record Keeping**
    - Despite using an immutable ledger, SMEs should maintain clear off-chain records to demonstrate procedural compliance. Documentation might include descriptions of data flows, data handling policies, and authorisations for changes in reference data.

6.3 Technology Stack Recommendations

While the blockchain arena remains dynamic, several platforms and supporting tools have emerged as strong candidates for SMEs aiming to bridge on-ledger and off-ledger storage. This subsection outlines core technologies that may suit smaller enterprises.

6.3.1 Blockchain Networks

1. **Ethereum**
    - *Profile*: One of the most mature public blockchains with robust smart contract functionality.
    - *Suited For*: SMEs seeking well-established ecosystems, decentralised finance solutions, or broad developer support.
    - *Considerations*: High transaction fees (often referred to as gas fees) may deter usage for high-throughput systems; options such as Layer 2 solutions (e.g., Polygon, Optimistic Rollups) can help reduce costs.
2. **Hyperledger (e.g., Hyperledger Fabric)**
    - *Profile*: A permissioned ledger framework well regarded by enterprises, backed by strong governance features and modular architecture.
    - *Suited For*: SMEs requiring private channels for sensitive data, or those in heavily regulated sectors.
    - *Considerations*: Typically requires more in-house knowledge to customise. However, it enables selective data disclosure and can be integrated with existing identity management systems.
3. **Corda**
    - *Profile*: Designed for regulated industries such as finance, emphasising point-to-point communication of transactions and data privacy between participants.
    - *Suited For*: SMEs in financial services or consortia-based environments where confidentiality is paramount.
    - *Considerations*: Not always as straightforward to deploy for general-purpose use cases, but offers strong privacy controls when business logic cannot be exposed to all network nodes.

6.3.2 Storage Solutions

1. **IPFS (InterPlanetary File System)**
    - *Description*: A peer-to-peer protocol that references content via cryptographic hashes, rather than fixed server locations.
    - *Advantages*: Decentralised and tamper-evident. By storing a file's hash on-chain, SMEs can guarantee data integrity without placing the full file on the ledger.
    - *Potential Drawback*: Data availability may hinge on the number of nodes hosting the file. SMEs might need to run their own IPFS node for reliable access.
2. **Filecoin**
    - *Description*: A decentralised storage marketplace running atop IPFS, enabling organisations to pay nodes for hosting data.
    - *Advantages*: Aligns with decentralisation aims and includes incentives for persistent data storage.
    - *Potential Drawback*: Costs may fluctuate based on market supply and demand; integration expertise is necessary for smaller teams not familiar with decentralised storage tokens.
3. **AWS Blockchain Templates**
    - *Description*: A set of preconfigured frameworks for launching blockchain networks using Amazon Web Services. Integrates with existing AWS infrastructure.
    - *Advantages*: Simplifies deployment for SMEs lacking in-house blockchain expertise; offers additional layers of security and backup under a more centralised model.
    - *Potential Drawback*: Relies on a single cloud provider, somewhat contrary to the decentralised ethos, and could introduce vendor lock-in over time.

6.3.3 Balancing Off-Chain Cloud Services and Hybrid Ledgers

While decentralized storage solutions like **IPFS** and **Filecoin** are gaining traction, many SMEs will continue to rely on mainstream cloud providers for off-ledger data due to their scalability, reliability, and ease of integration. Hybrid ledger systems, which combine the benefits of blockchain with conventional databases or cloud storage, provide a practical pathway for SMEs to adopt distributed ledger features without overhauling their existing IT infrastructure. This subsection explores key technologies and strategies for balancing off-chain cloud services with hybrid ledgers.

Key Considerations for Hybrid Ledgers

1. **Data Integrity and Verifiability**:
    - Use cryptographic hashes (e.g., SHA-256) to link off-chain data to on-chain references, ensuring tamper-proof verification.

- Implement **Merkle Trees** for efficient proof of data integrity without storing the entire dataset on-chain.

2. **Scalability and Cost Efficiency**:
   - Store large datasets (e.g., IoT sensor data, multimedia files) off-chain while maintaining a lightweight on-chain footprint for critical metadata or transaction records.
   - Leverage **Layer 2 solutions** (e.g., rollups, sidechains) to reduce on-chain storage costs.
3. **Interoperability**:
   - Use APIs and middleware to seamlessly connect off-chain cloud services with blockchain networks.
   - Adopt standards like **W3C Verifiable Credentials** or **DID (Decentralized Identifiers)** for cross-platform compatibility.

Examples of Hybrid Ledger Solutions

1. **Azure Confidential Ledger**:
   - **Profile**: Microsoft's verifiable ledger service designed for high-security applications, integrating with **Azure Key Vault** and other Azure security features.
   - **Suited For**: SMEs requiring tamper-proof audit trails, secure data sharing, or compliance with strict regulatory requirements.
   - **Key Features**:
     - Cryptographic proofs for data integrity.
     - Integration with Azure Blob Storage for off-chain data.
     - Support for confidential computing via **Azure Confidential Computing**.
   - **Use Cases**: Legal document management, healthcare data sharing, and financial auditing.
2. **Google Cloud Blockchain Tools**:
   - **Profile**: Google Cloud's suite of blockchain tools, including **BigQuery** for analytics and **Cloud Storage** for off-chain data management.
   - **Suited For**: SMEs in supply chain, IoT, or data-intensive industries requiring scalable analytics and storage.
   - **Key Features**:
     - Integration with Ethereum, Hyperledger, and other blockchain networks.
     - Real-time analytics using **BigQuery** for blockchain data.

        - Secure off-chain storage with **Cloud Storage** and **Cloud KMS** (Key Management Service).
    - **Use Cases**: Supply chain tracking, IoT data aggregation, and fraud detection.

3. **AWS Managed Blockchain**:
    - **Profile**: Amazon Web Services' fully managed blockchain service supporting **Hyperledger Fabric** and **Ethereum**.
    - **Suited For**: SMEs looking for a scalable, managed blockchain solution with seamless integration to AWS cloud services.
    - **Key Features**:
        - Integration with **Amazon S3** for off-chain storage.
        - Use of **Amazon QLDB** (Quantum Ledger Database) for immutable, cryptographically verifiable logs.
        - Support for **Lambda functions** for serverless automation.
    - **Use Cases**: Identity management, asset tracking, and decentralized applications (dApps).
4. **IBM Blockchain Platform**:
    - **Profile**: IBM's enterprise-grade blockchain platform built on **Hyperledger Fabric**.
    - **Suited For**: SMEs in regulated industries requiring private, permissioned blockchain networks.
    - **Key Features**:
        - Integration with **IBM Cloud Object Storage** for off-chain data.
        - Use of **IBM Watson** for AI-driven analytics on blockchain data.
        - Support for hybrid cloud deployments.
    - **Use Cases**: Supply chain transparency, trade finance, and healthcare data management.

Emerging Technologies for Hybrid Ledgers

1. **Chainlink Functions**:
    - **Profile**: A decentralized serverless platform that enables smart contracts to interact with off-chain APIs and cloud services.
    - **Suited For**: SMEs needing to connect blockchain networks with external data sources or cloud-based AI models.
    - **Use Cases**: Real-time data feeds, IoT integration, and AI-driven decision-making.
2. **Polygon Supernets**:

- **Profile**: Customizable blockchain networks built on Polygon, designed for hybrid on-chain/off-chain architectures.
- **Suited For**: SMEs requiring scalable, cost-effective blockchain solutions with off-chain data storage.
- **Use Cases**: Gaming, NFT marketplaces, and decentralized identity systems.

3. **Covalent**:
   - **Profile**: A unified API for querying blockchain data across multiple networks, including off-chain storage references.
   - **Suited For**: SMEs building analytics dashboards or dApps that require access to both on-chain and off-chain data.
   - **Use Cases**: Financial analytics, portfolio tracking, and supply chain monitoring.

Practical Implementation Strategies

1. **Data Partitioning**:
   - Store sensitive or high-value data on-chain (e.g., transaction records, ownership proofs) and less critical data off-chain (e.g., logs, media files).
2. **Hybrid Identity Management**:
   - Use decentralized identity (DID) solutions like **Sovrin** or **Microsoft ION** for on-chain identity verification, while storing user profiles and attributes off-chain.
3. **Event-Driven Architectures**:
   - Use **AWS EventBridge**, **Google Pub/Sub**, or **Azure Event Grid** to trigger blockchain transactions based on off-chain events (e.g., IoT sensor data, payment confirmations).
4. **Zero-Knowledge Proofs (ZKPs)**:
   - Implement ZKPs (e.g., using **zk-SNARKs** or **zk-STARKs**) to verify off-chain data without revealing its contents, enhancing privacy and scalability.

By leveraging these hybrid ledger solutions and strategies, SMEs can achieve a balance between the security and transparency of blockchain and the scalability and flexibility of off-chain cloud services. This approach allows for incremental adoption of distributed ledger technologies while minimizing disruption to existing IT systems.

7. Participation in Standardisation Activities & Stakeholder Engagement

Effective engagement with relevant standardisation bodies and industry stakeholders has been central to advancing the work on on-ledger and off-ledger data. This section provides an overview of the meetings, workshops, and broader consultations undertaken, highlighting both the technical contributions made and the awareness-raising efforts that have supported the development of the proposed framework. In total, the activities described below reflect a commitment to ensuring that the emerging guidance aligns

with international best practices, regulatory expectations, and the practical realities of industry adoption.

7.1 Meetings Attended & Contributions

7.1.1 Technical Workshops (ISO/TC 307, IEEE, ETSI)

Participation in technical workshops organised by international standard-setting bodies has been essential for refining the proposed approach. These sessions served as platforms to examine existing work, align with emerging initiatives, and gauge the level of interest in new frameworks.

1. **ISO/TC 307**
    - **Focus and Objectives**: The relevant sub-groups within ISO/TC 307 have paid increased attention to aspects of blockchain interoperability, data governance, and taxonomy for distributed ledgers. By engaging with these sub-groups, it was possible to glean insights into both the scope and methodology that international committees prioritise.
    - **Key Contributions**:
        - Shared an early outline of a classification methodology for deciding when data should be stored on-ledger versus off-ledger, including proposed criteria such as sensitivity and frequency of updates.
        - Offered input on how to integrate data-protection principles (e.g., data minimisation) into ledger design, proposing ways to reference off-ledger storage without duplicating large datasets on-chain.
2. **IEEE Blockchain Technical Community**
    - **Focus and Objectives**: IEEE-run workshops typically encompass issues such as blockchain reliability, cryptographic innovations, and multi-disciplinary applications.
    - **Key Contributions**:
        - Presented a set of use cases reflecting small and medium-sized enterprises' (SMEs) needs around on-ledger record-keeping (for transaction traceability) alongside off-ledger compliance requirements (e.g., protected data under privacy laws).
        - Highlighted potential synergies between emerging privacy-preserving cryptographic methods and existing smart-contract standards, with a view to bridging the gap between theoretical security guarantees and actual market demands.
3. **ETSI (European Telecommunications Standards Institute)**
    - **Focus and Objectives**: ETSI's interest in distributed ledgers often includes their application to internet infrastructure, identity management, and telecommunications frameworks.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

- **Key Contributions**:
  - Explored opportunities for integrating distributed identifiers (DIDs) with off-ledger repositories to store sensitive or highly dynamic information, thus reducing the burden on ledger resources.
  - Emphasised the importance of aligning telecommunication-grade security (e.g., network resilience and redundancy) with the resilience features offered by blockchain, ensuring that off-ledger data remains consistently available.

Overall, engagement with the above forums has allowed for two-way knowledge sharing. On the one hand, the project team was exposed to state-of-the-art thinking and complementary work items within each standardisation body. On the other, the group was able to contribute early drafts, obtain technical peer feedback, and shape a more robust framework that resonates with international expectations.

7.1.2 EU Blockchain Partnership (EBSI) Participation

The European Blockchain Partnership, developed under the European Commission's auspices, aims to accelerate the use of blockchain solutions in public services. Participation in working meetings and roundtable discussions was highly beneficial in aligning project recommendations with European priorities.

- **Technical Dialogues**:
  - Contributed to conversations on how public sector services can manage citizen records in a decentralised manner, especially when data protection rules require selective disclosure rather than full publication of information.
  - Highlighted how the classification of data as "on-ledger" or "off-ledger" could alleviate compliance overheads in areas like eGovernment, eHealth, and cross-border administrative exchanges.
- **Policy Alignment**:
  - Reviewed EBSI pilot projects to pinpoint the components that may benefit from a more structured approach to data classification, such as referencing large file repositories off-ledger while anchoring essential proofs on the blockchain.
  - Suggested the development of standard documentation templates, to be used by Member States, which outline recommended practices for storing personal and non-personal data in a distributed ledger ecosystem.

EBSI's emphasis on cross-border interoperability harmonises with the project's proposed framework. Ensuring that minimal, critical references reside on-ledger, while sensitive data remains protected in secure environments, could align public services with European digital strategies.

7.2 Stakeholder Interviews & Surveys Conducted

A thorough consultation process was conducted to test assumptions about data classification, discover unaddressed challenges, and refine practical guidance. Interviews and surveys involved:

1. **Blockchain Industry Experts**
    - **Rationale**: Specialists in blockchain architecture and smart contract development brought direct technical insights on what currently works in production networks. They highlighted real-world constraints, particularly relating to transaction costs, performance bottlenecks, and the availability of cryptographic libraries.
    - **Key Observations**:
        - On-ledger data is most beneficial when it must be transparent or universally accessible, whereas high-volume or sensitive records are more suited for off-ledger storage.
        - A major theme was the need for robust cryptographic linking between the ledger and external repositories, ensuring that any tampering with off-ledger records remains detectable.
2. **Compliance Officers**
    - **Rationale**: Organisations that operate within strictly regulated environments (e.g., finance, healthcare) must deal with compliance demands that often conflict with blockchain immutability. These professionals provided a practical perspective on how to balance contractual transparency with the potential regulatory complications.
    - **Key Observations**:
        - There is a strong emphasis on having the flexibility to remove or modify personal data in compliance with “right-to-be-forgotten” provisions. This requirement often necessitates an off-ledger approach.
        - Standard operating procedures for documenting data flows and versioning are essential, to demonstrate compliance in audits. This can be challenging when multiple parties upload data to distributed ledgers.
3. **SME Representatives**
    - **Rationale**: Smaller businesses often face greater obstacles when adopting complex technologies. Their perspective helps ensure that standards do not impose unrealistic technical or cost burdens.
    - **Key Observations**:
        - Many SMEs expressed interest in using ledger-based solutions for supply chain traceability or digital asset management, but were hesitant about storage costs, regulatory responsibilities, and a perceived lack of practical guidance.
        - Clear guidelines on how to compartmentalise data - storing minimal proof on-chain while keeping day-to-day operational data off-ledger - would help them manage overheads and maintain regulatory clarity.

The interviews and surveys reinforced that a careful mix of on-ledger and off-ledger strategies can mitigate scalability, privacy, and cost challenges. The feedback is being integrated into the technical frameworks, with particular attention to reducing complexity for SMEs and ensuring compliance for heavily regulated industries.

7.3. Pilot Testing and Expert Input

While formal pilot testing with live deployments was not feasible within the project timeline, structured internal evaluations were conducted to simulate the real-world applicability of the proposed data classification framework. These simulations applied the framework across representative scenarios using established Distributed Ledger Technologies (DLTs) - Corda, Hyperledger Fabric, and Polygon - focusing on conditions such as data criticality, transaction costs, and regulatory exposure. The objective was to test the framework's robustness for small and medium-sized enterprises (SMEs) in the European Union (EU), drawing on hypothetical yet plausible use cases from sectors like finance, logistics, real estate, and technology services.

The simulations were informed by expert input from blockchain architects, compliance officers, and SME representatives across the EU. While many outcomes validated the framework's utility, some revealed limitations and incomplete assumptions in the initial findings, underscoring the need for continued research. Below, the results - both positive and negative - are detailed, alongside refinements and areas requiring further investigation.

7.3.1. Simulated Use Cases and Findings

7.3.1.1. Financial Sector: Bulgarian Microfinance Lender (Corda)

A Bulgarian microfinance SME providing small loans was simulated using Corda, leveraging its privacy-centric design for managing loan agreements. The scenario tested classifying high-frequency repayment data versus static contractual terms.

- **Findings**: Storing repayment batches on-ledger overwhelmed Corda's node replication, inflating costs for the SME's constrained budget. The framework successfully shifted to anchoring cryptographic hashes on-ledger, with full records off-ledger, reducing costs and aligning with Bulgaria's National Revenue Agency (NRA) audit requirements. However, the SME's confusion over immutability - fearing permanent errors - exposed a gap in the initial assumption that SMEs would readily adapt to hybrid models, necessitating clearer guidance.
- **Expert Input**: A Sofia-based compliance officer praised the GDPR-compliant off-ledger shift for borrower data but flagged inadequate metadata standards, prompting a refinement to prioritise immutability sensitivity as a distinct criterion.

7.3.1.2. Logistics Sector: Polish Freight Forwarder (Hyperledger Fabric)

A Polish freight forwarding SME tracked shipment milestones across an EU consortium using Hyperledger Fabric's permissioned architecture.

- **Findings**: On-ledger storage of all updates bloated the ledger, slowing performance - a positive outcome was achieved by anchoring milestone proofs (e.g., Merkle roots) on-ledger and storing logs off-ledger via IPFS. However, an initial finding that Fabric's private data collections would

suffice for privacy was incomplete. The simulation revealed synchronisation lags between on-ledger proofs and off-ledger logs, risking data inconsistencies during peak cross-border traffic (e.g., Poland-Germany routes). This contradicted the assumption that Fabric's architecture inherently supports seamless hybrid pipelines, indicating a need for further research into real-time anchoring protocols.

- **Expert Input**: A Warsaw-based blockchain architect suggested Fabric-IPFS integration but cautioned that SMEs lack the expertise to manage such setups, reinforcing the need for simplified tools.

7.3.1.3. Real Estate Sector: Slovakian Property Manager (Polygon)

A Slovakian SME managing leases tested Polygon, assessing payment records and tenant data under GDPR constraints.

- **Findings**: Polygon's low fees (e.g., €0.01 per transaction) enabled on-ledger payment proofs, a cost-effective success. Off-ledger storage of tenant data via local servers, linked by hashes, met GDPR erasure needs. Yet, the simulation exposed a limitation: selective on-chain signing for payments assumed SMEs could easily configure smart contracts, but the SME struggled with Polygon's technical complexity, contradicting the initial finding that Layer 2 solutions are inherently SME-friendly. This suggests ongoing research into user-friendly interfaces is essential.
- **Expert Input**: A Bratislava-based legal consultant endorsed the compliance approach but noted insufficient documentation templates for audits, refining the framework to include regulatory exposure explicitly.

7.3.1.4. Technology Services Sector: Romanian Authentication Provider (Hybrid Approach)

A Romanian SME offering blockchain-based authentication tested a Polygon-Fabric hybrid for login verifications.

- **Findings**: The hybrid model - Polygon for identity proofs, Fabric for credential storage - balanced scalability and privacy but faltered under high-frequency logins. Synchronisation delays between on-ledger proofs and off-ledger updates caused verification failures, a negative outcome highlighting over-optimism in the initial finding that hybrid systems are readily scalable for SMEs. This gap necessitates further study into cross-DLT synchronisation mechanisms.
- **Expert Input**: A Bucharest-based developer proposed Chainlink oracles for synchronisation but admitted current solutions are too complex for SMEs, urging continued exploration of accessible tools.

7.3.2. Key Refinements and Areas for Further Research

The simulations and expert feedback prompted refinements while exposing areas where the framework falls short:

1. **Immutability Sensitivity as a Distinct Criterion**: Positive feedback elevated immutability sensitivity from a regulatory subset to a standalone factor, addressing SME concerns (e.g., Bulgaria) about error correction in immutable systems.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

2. **Flexible Categorisation**: Successful hybrid pipelines (e.g., Poland) validated flexible strategies, but incomplete synchronisation (e.g., Romania) showed that non-binary approaches require robust real-time protocols - research must continue to bridge this gap.

3. **Cost and Scalability**: Cost-effective shifts to off-ledger storage (e.g., Slovakia) succeeded, yet Fabric's ledger bloat (Poland) and Polygon's complexity (Slovakia) suggest SMEs need simpler scaling solutions, an area for further investigation.

4. **Regulatory Alignment**: GDPR compliance was strengthened (e.g., Slovakia), but incomplete metadata standards (Bulgaria) and synchronisation issues (Romania) indicate ongoing research is needed to meet EU regulatory expectations fully.

5. **Incomplete Initial Finding**: The Polish logistics case disproved the initial assumption that Hyperledger Fabric's private data collections alone could handle hybrid data flows efficiently. Synchronisation failures under real-world volumes revealed a need for enhanced off-ledger integration, challenging the framework's completeness and necessitating further study into latency and consistency.

7.3.3. Validation and Ongoing Challenges

While the framework proved adaptable across Corda, Fabric, and Polygon, not all outcomes were positive. Successes - like cost reductions in Bulgaria and compliance in Slovakia - coexist with failures, such as synchronisation issues in Romania and Poland, and Polygon's complexity in Slovakia. These negative findings highlight that the research is not conclusive; SMEs' technical capacity and DLT interoperability remain barriers. Continued investigation into real-time anchoring, simplified interfaces, and standardised metadata is critical to address these shortcomings, aligning with the EU's ICT Rolling Plan (2023) and ISO/TC 307 goals.

For EU SMEs, the framework offers a starting point - favouring minimal on-ledger data for trust and off-ledger solutions for flexibility - but its limitations suggest that broader adoption hinges on resolving synchronisation, usability, and regulatory gaps through future research.

7.3.4 EBSI Open Source

On 13 March 2025, the European Blockchain Services Infrastructure (EBSI) team hosted an open-source workshop from 10:30 to 12:00, attended virtually by members of the EBSI ecosystem, including representatives from INDUSTRIA involved in this project. The session, titled "EBSI Open-Source Workshop," provided an overview of EBSI's open-source deliverables and elaborated on the European Union Public Licence (EUPL) under which these resources will be released. While the workshop did not directly address the specific research on onchain and offchain ledger data management central to this report, its focus on open-source principles offers significant relevance to the broader adoption and standardisation of blockchain technologies, particularly for EU SMEs. Open source is a cornerstone of EBSI's vision to democratise access to blockchain infrastructure, fostering collaboration, transparency, and innovation across public and private sectors. The event, attended via a virtual platform (connection details provided by the EBSI team), underscored the importance of freely accessible tools and frameworks - an ethos that aligns with this project's aim to lower barriers for SMEs.

7.3.4.1 Relevance to Onchain and Offchain Ledger Research

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

Although the workshop's agenda did not explicitly cover onchain and offchain data classification, the open source deliverables and EUPL framework present opportunities to integrate this research into EBSI's ecosystem. Insights from the session suggest several ways this could be achieved:

- **Open-Source Repositories as a Distribution Channel**: EBSI's repositories could host tools or libraries derived from this framework - e.g., a decision matrix for onchain versus offchain storage (Section 6.2.1) or cryptographic anchoring protocols (Section 6.2.2). Released under the EUPL, these could be freely adopted by SMEs, enhancing the practical utility of pilot findings (e.g., Bulgaria's cost-effective hybrid model, Section 7.3.1.1).
- **Collaborative Development**: The workshop's emphasis on DEP project contributions highlights a model for collaborative enhancement. This research could contribute code or documentation - such as the refined metadata standards from the Slovakian pilot (Section 7.3.1.3) - to EBSI's open-source ecosystem, inviting further refinement by EU developers and aligning with interoperability goals (Section 7.1.2).
- **Simplifying SME Adoption**: Open source aligns with the need for user-friendly solutions identified in pilot testing (e.g., Poland's expertise gap, Section 7.3.1.2). By integrating simplified interfaces or pre-configured modules (e.g., for Polygon smart contracts, Section 7.3.1.3) into EBSI's repositories, the framework could leverage open source to address technical complexity, a key barrier for SMEs.
- **Regulatory Alignment**: The EUPL's compatibility with GDPR and other EU regulations supports the framework's compliance focus (Section 7.3.2). Open-source implementations could embed best practices - like off-ledger data minimisation from the Romanian hybrid case (Section 7.3.1.4) - ensuring SMEs adopt standards that meet legal requirements without proprietary constraints.

7.3.4.2 Implications and Future Integration

Participation in the workshop, documented via meeting notes and a confirmation email from the EBSI team, provided a strategic lens for this project. While the immediate discussion did not pivot to onchain/offchain specifics, the open-source vision offers a pathway to disseminate and evolve the framework beyond this report. For instance, contributing a modular toolkit - combining pilot-tested hybrid approaches (e.g., Fabric-IPFS integration, Section 7.3.1.2) with EUPL-licensed code - could empower EBSI users to address synchronisation challenges (Section 7.3.2) collaboratively. This aligns with the EU's ICT Rolling Plan (2023) and EBSI's cross-border ambitions, extending the research's impact.

The workshop's call for contributions suggests that future iterations of this framework could be proposed for inclusion in EBSI's repositories during events like the ISO/TC 307 meeting (7-11 April 2025, Section 7.1). This would require adapting the standards to open-source formats - e.g., modular APIs or documentation - ensuring they remain accessible and adaptable. Such integration would not only validate the research but also address ongoing challenges (Section 7.3.3) by tapping into a wider developer community, reinforcing the need for continued investigation into practical, open-source-aligned solutions.

7.4. Expert Insights on On-Ledger and Off-Ledger Data

This section presents a synthesis of insights gathered through additional structured interviews and informal consultations with professionals across key regulated sectors for a period of less than three months. Its purpose is to supplement the technical and theoretical analysis of on-ledger and off-ledger data classification with real-world perspectives from practitioners involved in deploying, managing, or regulating such systems.

The rationale behind this consultation effort lies in the recognition that data classification is not merely a technical problem, but also a regulatory, operational, and economic one. Sector-specific workflows, risk tolerances, and legal obligations all influence whether certain types of data should - or should not - reside directly on a distributed ledger.

Rather than attempt to generalise across all possible domains, this effort focused on three sectors with high stakes in data integrity and provenance: finance (banking and insurance), legal services, and logistics. These sectors were chosen for their diverse regulatory profiles, their ongoing involvement with blockchain-based innovation, and their exposure to questions of data governance, cross-border processing, and compliance.

By structuring the feedback according to these domains, the report aims to ground its recommendations in operational realities and regulatory expectations, ensuring that the proposed classification framework remains applicable across multiple contexts without sacrificing specificity or rigour.

7.4.1. Finance

In the financial sector, feedback was collected from a diverse array of stakeholders, including major banks headquartered in Italy, Belgium, Bulgaria, and Greece, all with operations spanning the European continent. These large institutions consistently voiced concerns about the inherent trade-off between transparency and confidentiality, particularly in complex financial activities such as syndicated lending, structured finance, and reinsurance transactions. A prevailing sentiment emerged that storing sensitive data directly on-ledger introduces significant risks to client privacy and commercial sensitivity, potentially exposing proprietary deal structures or customer identities in ways that could undermine competitive positioning. For example, an Italian bank highlighted that full onchain records of syndicated loans could inadvertently reveal negotiation leverage points to rival institutions, a risk deemed unacceptable in their operational context.

In response to these concerns, these organisations leaned toward a selective anchoring approach. This method involves recording cryptographic proofs - such as hashes or digital signatures - onchain to ensure data integrity and auditability, while retaining the full underlying records within traditional, secure offchain storage infrastructures. A Belgian bank noted that this hybrid model aligns well with existing cybersecurity frameworks, such as the EU's General Data Protection Regulation (GDPR), allowing them to leverage blockchain's benefits without overhauling their established data governance protocols. Meanwhile, a Bulgarian institution shared insights from its internal evaluations of zero-knowledge proofs (ZKPs), a cryptographic technique aimed at preserving confidentiality while enabling verification. Despite the theoretical promise of ZKPs, the bank concluded that current implementations were computationally intensive, poorly integrated with legacy systems, and lacked the maturity needed for deployment at scale. This sentiment was echoed by a Greek counterpart, which cited the steep learning curve and high consultancy costs as additional barriers to adoption.

The study gained further depth from interviews with smaller regulated lenders in Bulgaria, alongside feedback from several of their customers - primarily small and medium-sized enterprises (SMEs). These entities operate under intense regulatory scrutiny, navigating a web of overlapping national and EU compliance requirements, such as those stemming from the Capital Requirements Regulation (CRR) and the European Banking Authority (EBA) guidelines. Their perspectives provided a practical counterbalance to the more theoretical concerns of larger institutions. A recurring theme among these smaller lenders was a notably low level of knowledge about managing both onchain and offchain data effectively. For instance, one Bulgarian cooperative bank admitted to struggling with the concept of data immutability inherent in blockchain technology. Its staff expressed confusion over how immutable onchain records could accommodate errors or updates - common occurrences in loan agreements or customer records - leading to fears of locking inaccurate data into perpetuity. A customer of this bank, a local manufacturing SME, reinforced this concern, recounting an instance where a clerical error in a loan term required multiple revisions, a process they doubted could be managed on an unalterable ledger.

This knowledge gap extended beyond technical understanding to practical application. Several local lenders reported difficulties distinguishing between use cases suited for onchain versus offchain storage. One lender, for example, initially assumed that all transaction data should be recorded onchain for transparency, only to later realize - after consultation with an external IT provider - that this approach clashed with GDPR's "right to erasure" provisions. Another Bulgarian institution highlighted a real-world challenge: their legacy IT systems, built on outdated software stacks, lacked the interoperability needed to interface with blockchain platforms, let alone manage a dual onchain-offchain workflow. These findings underscore a broader issue: while larger banks have the resources to experiment with blockchain pilots, smaller players are constrained by limited budgets, outdated infrastructure, and a workforce not yet versed in distributed ledger technology (DLT).

Despite these challenges, both large institutions and SMEs identified compliance-driven use cases as the most promising near-term applications for blockchain in finance. These include immutable audit trails for anti-money laundering (AML) measures under the EU's Sixth Anti-Money Laundering Directive (6AMLD), recordkeeping aligned with MiFID II's transparency requirements, and transaction traceability to meet financial conduct rules enforced by national regulators like Italy's CONSOB or Belgium's FSMA. However, adoption remains heavily contingent on clearer regulatory guidance. Stakeholders repeatedly called for harmonized EU standards on data residency - particularly given the bloc's strict rules on cross-border data transfers - and more explicit directives from financial supervisors like the European Central Bank (ECB) or the EBA on integrating blockchain with existing frameworks.

In the insurance domain, stakeholders displayed a distinct but complementary set of priorities. Interest centered on claims automation and process-level auditability, with a Greek insurer citing the potential for smart contracts to streamline payouts in short-tail policies like travel insurance. The value of immutable records for dispute resolution was widely acknowledged; a Belgian insurer noted that blockchain could reduce legal costs in contested claims by providing tamper-proof evidence of policy terms and claim timelines. However, insurers also emphasized the limitations of onchain data for long-tail policies, such as liability or life insurance, where contractual obligations evolve over decades. A Bulgarian insurance cooperative pointed out that offchain flexibility is essential for managing updates, annotations, and negotiated terms - such as premium adjustments or beneficiary changes - throughout the policy lifecycle. This cooperative also admitted to a rudimentary understanding of blockchain, with its

management unsure how immutable ledgers could handle mid-term policy cancellations, a frequent occurrence in their market.

Across the board, the study revealed a stark disparity in blockchain literacy. While major banks grapple with strategic trade-offs and technical feasibility, smaller lenders and insurers in regions like Bulgaria face more foundational hurdles: a lack of expertise, misconceptions about immutability, and an overreliance on external vendors for guidance. These challenges suggest that widespread adoption in the EU financial sector will require not only technological maturation but also targeted education and capacity-building initiatives - particularly for local players operating at the margins of innovation.

7.3.2. Legal Sector

Legal experts raised fundamental concerns about the legal effect of onchain records, particularly in regulated environments. One regulatory advisor noted that under current law in most jurisdictions, the presence of a record on a blockchain does not automatically confer evidentiary status or contractual enforceability. This observation underscores a critical challenge: the semantic and legal context of a data item must be meticulously preserved, whether it resides on-ledger or is referenced from an off-ledger repository. Without such context, the utility of blockchain in legal settings remains limited, as courts and regulators may not recognize onchain data as authoritative absent additional validation.

As part of this consultation, input was sought from Wolf Theiss, a prominent and large legal advisory firm operating across Central and Eastern Europe. With offices in EU member states including Austria, Bulgaria, Croatia, Czech Republic, Hungary, Poland, Romania, Slovakia, and Slovenia, Wolf Theiss brought a comprehensive cross-border perspective to the study. Their extensive footprint and expertise in regulatory compliance, commercial law, and emerging technologies made them an ideal starting point for this analysis. Representatives from the firm reinforced the view that tokenized representations of legal obligations - such as smart contract-based escrows or settlement triggers - require a robust off-ledger substrate to provide the necessary legal interpretative context. They argued that immutable onchain records alone rarely carry legal finality unless accompanied by clear metadata, provenance rules, and, in some cases, human-readable documentation that courts can interpret. For instance, a Vienna-based partner highlighted a hypothetical smart contract for a cross-border trade settlement: while the blockchain could record the transaction's execution, its legal enforceability might hinge on offchain terms specifying jurisdiction and dispute resolution mechanisms.

Wolf Theiss's insights served as a foundational base for further exploration, prompting me to extend the study by interviewing small and medium-sized enterprises (SMEs) across various sectors to broaden the report's scope. These SMEs, drawn from industries such as logistics, real estate, and technology services, provided practical, ground-level perspectives that complemented the firm's high-level legal analysis. In the logistics sector, a Bulgarian SME operating a regional freight forwarding business expressed uncertainty about using blockchain to record shipping contracts. They valued the idea of immutable proof of delivery but worried that fixed onchain records couldn't accommodate last-minute shipment amendments - a frequent occurrence in their operations. This feedback highlighted a gap between blockchain's rigidity and the fluidity of real-world commercial agreements, reinforcing Wolf Theiss's emphasis on offchain context.

In the real estate sector, a Slovakian SME involved in property leasing raised similar concerns but framed them around tenant data. They saw potential in blockchain for tracking lease payments transparently yet

flagged GDPR's "right to be forgotten" as a major hurdle. The SME's managing director noted that tenant records often include personal data subject to erasure requests, and they lacked clarity on how immutable ledgers could comply without costly workarounds. This echoed Wolf Theiss's caution about data protection tensions, though the SME admitted to limited internal expertise, relying instead on external legal advice that remained inconclusive. Meanwhile, a Polish technology services SME developing blockchain-based authentication tools offered a more optimistic view. They suggested that hybrid models - pairing onchain hashes with offchain databases - could address both legal enforceability and GDPR compliance, though they conceded that such solutions were still nascent and untested in Polish courts.

Another major area of concern identified by Wolf Theiss and amplified through SME feedback was the tension between immutability and data protection regulations, particularly the EU's General Data Protection Regulation (GDPR). Legal advisors at the firm stressed the importance of identifying data that may be subject to erasure rights and ensuring it is either excluded from on-ledger storage or explicitly exempt under a documented legal basis - such as contractual necessity or public interest. A Bucharest-based associate cited a practical example: a blockchain record of a client's payment history might breach GDPR if it includes personal identifiers and cannot be deleted upon request. This feedback directly informed the recommendation that data should default to off-ledger status where regulatory reversibility may be required, a stance SMEs broadly supported. The logistics SME, for instance, favored keeping customer details offchain in their existing systems, using blockchain only for anonymized transaction hashes - a pragmatic compromise born from their operational constraints.

The findings from these SME interviews underscored a recurring theme: while blockchain's potential is recognized, its legal application remains hampered by a lack of clarity and adaptability. The logistics sector craved flexibility, the real estate sector prioritized compliance, and the technology services sector sought innovation within legal bounds. Wolf Theiss's cross-jurisdictional expertise provided the theoretical backbone for these discussions, but the SMEs grounded them in real-world challenges - revealing a landscape where legal frameworks lag behind technological ambition. Together, these insights suggest that blockchain's integration into the legal sector will require not only technical hybrid solutions but also harmonized EU guidance on evidentiary standards and data protection.

7.3.3. Logistics

Logistics actors were generally more supportive of increased on-chain visibility, particularly where it can improve auditability and reduce disputes. However, this support was qualified by practical concerns around performance and cost. High-frequency updates - such as location pings or temperature readings - were universally seen as unsuitable for direct on-ledger storage, especially on public or gas-intensive blockchain platforms. In contrast, milestone-based events (e.g. port departure, customs clearance, arrival confirmation) were considered viable for on-ledger registration, offering a cryptographically verifiable trail without overwhelming infrastructure or introducing prohibitive costs.

Consultations included a Bulgarian airline with significant air cargo operations and a large logistics provider managing multimodal freight across the Balkans and the EU. Both organisations are among the largest in their sectors domestically and cannot be classified as SMEs. Their feedback reflected mature IT environments and data governance protocols, often relying on hybrid integration layers and centralised tracking systems. For these larger players, blockchain was seen less as a data repository and more as a potential integrity layer for synchronisation across partners.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

To balance this perspective, additional interviews were conducted with the IT teams of smaller cargo handlers, including ground operators and warehouse providers at Sofia Airport. These teams expressed interest in blockchain-backed timestamping and document verification but noted the importance of low-overhead tools and integration simplicity. Concerns were raised about the steep learning curve and infrastructure requirements of existing blockchain implementations, highlighting the need for tooling that reflects operational realities at smaller scale.

A common thread across the logistics spectrum was a strong emphasis on interoperability. Stakeholders in freight forwarding and regional transport consistently indicated that the ledger location of the data was secondary to system compatibility and data accuracy. As one logistics consultant pragmatically put it: “We don’t care if it’s on or off-chain, as long as systems talk to each other and the data’s correct.” This reinforces the priority of standardised metadata formats, clear APIs, and consistent semantics across platforms over prescriptive technology choices.

7.4. Standards Engagement and Knowledge Exchange

Although no formal standardisation meetings were attended during the reporting period of less than three months, the development of this framework has been significantly informed by active monitoring of ongoing efforts within key standardisation bodies and initiatives. These include ISO/TC 307 WG7 (Working Group 7 on Interoperability), Hyperledger, and the European Blockchain Services Infrastructure (EBSI)-related projects under the European Blockchain Partnership. Email-based exchanges with participants from these affiliated working groups, combined with thorough reviews of draft specifications and technical documents, have provided critical context for refining the proposed data classification framework. This engagement ensured that the approach remains attuned to global technical trends and emerging best practices.

In addition, several internal discussions and peer reviews were conducted within the context of ongoing blockchain strategy work at INDUSTRIA (the company I’m working for). These sessions involved blockchain architects, compliance specialists, and industry practitioners, whose insights helped shape the tone, structure, and practical applicability of this deliverable. By aligning the framework with industry expectations and existing regulatory frameworks - such as the EU’s General Data Protection Regulation (GDPR) and the Markets in Crypto-Assets (MiCA) regulation - these efforts laid a solid foundation for broader stakeholder engagement.

Looking ahead, a pivotal opportunity for standards engagement and knowledge exchange is scheduled for 7-11 April 2025, during the ISO/TC 307 meeting in Brussels, Belgium. This event, hosted in the heart of the EU’s policymaking hub, will serve as a critical platform to advance the standardisation of on-ledger and off-ledger data management. A series of discussions and activities are planned to facilitate robust collaboration and knowledge sharing among international experts, industry stakeholders, and SME representatives. These include:

- **Presentations**: Formal sessions will showcase the proposed framework, highlighting its applicability to SMEs and its alignment with ISO/TC 307’s interoperability and governance objectives. A dedicated presentation will explore simulation findings (see Section 7.3), offering both successes and challenges to invite constructive critique and foster consensus on data classification standards.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

- **Workshops**: Interactive workshops will delve into practical implementation challenges, such as synchronising on-ledger and off-ledger data flows and integrating cryptographic anchoring protocols. Participants will collaborate on refining the framework's decision matrix, with a focus on addressing gaps identified in pilot testing - e.g., real-time synchronisation and SME usability - ensuring the outcomes are actionable across EU contexts.
- **Expert One-to-One Sessions**: Tailored consultations with TC 307 members, INATBA, Hyperledger contributors, and EBSI representatives will enable in-depth discussions on specific technical and regulatory hurdles. These sessions will target unresolved issues, such as standardising metadata schemas and reconciling blockchain immutability with GDPR's erasure rights, providing a forum for bespoke feedback and strategic alignment.
- **Panel Discussions**: Cross-disciplinary panels will bring together regulators, technologists, and SME leaders to debate the framework's implications for finance, supply chains, and public services. A key focus will be harmonising the proposed standards with EBSI's cross-border priorities and Hyperledger Aries' identity management protocols, amplifying the framework's relevance to EU digital strategies.

This Brussels meeting represents a cornerstone for the real exchange of ideas, building on the preparatory work conducted to date. It will provide a dynamic environment to test the framework against diverse perspectives, refine its technical specifications, and strengthen its alignment with international standards. The outcomes are expected to influence subsequent iterations of this document, ensuring it reflects the latest consensus on blockchain data management. INDUSTRIA's active participation - through presenting simulation results, facilitating workshops, and engaging in one-to-one dialogues - will underscore its commitment to advancing practical, SME-friendly standards within the global blockchain ecosystem.

By leveraging the Brussels event, the project aims to transition from internal refinement to broader collaborative development, addressing limitations identified in earlier phases (e.g., synchronisation challenges and incomplete SME readiness) while capitalising on the collective expertise of ISO/TC 307 and its affiliates. This exchange will be instrumental in shaping a framework that not only meets industry needs but also supports the EU's ambition to lead in blockchain innovation and standardisation.

7.5 Dissemination & Awareness Activities

To further amplify the project's insights, a range of dissemination and awareness-raising initiatives has been conducted, targeting various communities of interest.

7.5.1 Publications

Multiple articles were prepared for peer-reviewed journals and industry platforms. These explored the logic behind classifying data by sensitivity, frequency of updates, and compliance requirements. Key aims included:

- Educating a broader audience on the practical dilemmas organisations face when deciding which data to anchor on a chain.
- Demonstrating how robust cryptographic proofs can allow off-ledger data to remain verifiable without committing large file sets to a distributed ledger.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

By articulating the approach in both academic and practitioner-focused venues, the project has helped shape a nuanced understanding of blockchain data management and contributed to the knowledge base available to developers, compliance experts, and policy makers.

7.5.2 Conferences & Workshops

Engagement at conferences and workshops allowed for interactive dialogue and broader outreach:

- **Sector-Focused Conferences**: Presentations at events dedicated to finance, healthcare, and supply chain illustrated how domain-specific considerations shape data storage choices. These discussions highlighted case studies from earlier pilot implementations, showcasing both the benefits and the unresolved hurdles of adopting a hybrid on-/off-ledger model.
- **Research Symposia**: Academic blockchain gatherings provided a forum for more in-depth theoretical discussions, centring on privacy-preserving methods and the feasibility of advanced cryptographic techniques. This audience was particularly interested in bridging conceptual developments with tangible industry requirements.

In each instance, feedback from diverse participant groups - ranging from cryptographers to corporate IT managers - has been funnelled back into the iterative refinement of the framework.

7.5.3. Additional Dissemination Context

Beyond past efforts, a significant future dissemination milestone is planned for the ISO/TC 307 meeting in Brussels, 7-11 April 2025 (Section 7.4). This event will feature presentations, workshops, and expert one-to-one sessions to share the finalized standards, incorporating pilot and consultation feedback. While outside the reporting period, it underscores the project's ongoing commitment to knowledge exchange, with INDUSTRIA set to present simulation outcomes and gather further input to refine post-publication iterations.

7.5.3 Research Presentations

In parallel with the more formal conferences, smaller research meet-ups and invited talks at universities have played a vital role in circulating emerging findings to both early-career researchers and veteran investigators in distributed systems. These forums tend to encourage deeper technical questions, often resulting in valuable critiques around cryptographic proofs, performance benchmarks, and governance strategies.

By regularly presenting interim results and inviting critique, the project benefitted from continuous external validation, enabling clearer articulation of how on-ledger/off-ledger data decisions can be made in tandem with evolving privacy rules and enterprise constraints.

8. Future Research Directions

8.1 Enhancing Blockchain Interoperability

The continuing growth of decentralised technologies has led to significant fragmentation, with multiple blockchain platforms offering diverse consensus models and data structures. While these innovations bring flexibility, they also introduce difficulties in sharing information across chains, particularly where sensitive or voluminous data remain stored outside the ledger. Addressing these challenges in a structured way necessitates research into enhanced cross-chain consistency models, improved

standardisation around cryptographic proofs, and refined protocols to link on-ledger with off-ledger information seamlessly.

8.1.1 Cross-Chain Data Consistency Models

One of the core difficulties in interoperability is guaranteeing that transactions or states updated on one blockchain remain verifiably consistent on another. The exchange of both on-ledger and off-ledger data compounds this complexity, as each platform may apply different consensus rules, data validation methods, and privacy requirements. Future research should focus on:

- **Universal Data Validation Layers**
  New protocols that bundle data from multiple networks and apply a unified approach to verification. These layers could operate at the "edge" of each platform, using shared cryptographic methods to confirm that the data have not been manipulated once they pass from one chain to another.

- **Proof Aggregation and Standardised Hashing**
  Employing consistent hashing algorithms and proof aggregation techniques can enable chains to rely on concise evidence of data integrity, even if the data themselves reside elsewhere. Such solutions could be particularly valuable for off-ledger information, where a simple cryptographic reference might suffice for cross-chain audits without revealing private records in full.

- **Bridging and Atomic Swaps**
  Research into bridging solutions and atomic swaps can help reconcile differences in transaction finality, consensus speed, and data formats. By coordinating updates to data stored on different platforms, bridging protocols can help reduce transaction disputes. A vital research question is how to extend these approaches to handle off-chain data references: for example, verifying that a large file or sensitive dataset "mirrors" the on-chain state without unduly slowing the network.

8.1.2 Practical Deployment for Enterprises

While technical advances in cross-chain data consistency are essential, much work remains to ensure enterprise adoption. Organisations often need to integrate multiple internal or external blockchains into their workflows, linking each to legacy systems. Future research can address:

- **Hybrid Governance Models**
  Finding the right balance between decentralised control and conventional corporate governance. In many industries, fully public networks are unsuitable due to compliance or privacy restrictions. Conversely, highly centralised architectures may undermine the trust benefits that blockchains can provide. Further investigation into hybrid governance approaches - where cross-chain data moves fluidly but remains auditable - will be instrumental.

- **Operational Frameworks for Large Consortia**
  When multiple companies collaborate (for instance, within a supply chain or consortium-based finance project), cross-chain data must be validated quickly at scale. This demands operational playbooks detailing how off-ledger data are encrypted, hashed, and linked, and who is authorised to validate changes when multiple parties hold read or write privileges.

8.2 Privacy-Preserving Blockchain Models

Many blockchain systems emphasise transparency, yet privacy requirements are increasing in

importance. Complex regulatory frameworks now cover data protection and personally identifiable information in sectors such as finance, healthcare, and public administration. Consequently, research must investigate methods that allow for privacy without forgoing the assurance and immutability provided by distributed ledgers.

8.2.1 Using Homomorphic Encryption for On-Chain Processing

Homomorphic encryption permits computations to be carried out on encrypted data, so that the results remain encrypted and only reveal their meaning once decrypted by authorised users. While the concept has been around for some time, its computational overhead and implementation complexity have hindered wide adoption. Future work should concentrate on:

- **Efficiency and Scalability**
  Significant performance improvements or specialised hardware acceleration are likely needed before homomorphic encryption is viable for mainstream blockchain uses. Ensuring that these improvements remain backward-compatible with existing networks will also be essential.

- **Data Lifecycle Management**
  Even if data are encrypted on-chain, operators must still consider how to handle events such as key rotation, contract upgrades, or changes in legal requirements for record retention. Research into “evergreen” approaches - where the underlying encryption method can evolve without breaking historical records - will be especially valuable.

8.2.2 Zero-Knowledge Proofs (zk-SNARKs)

Zero-knowledge proofs enable a party to prove that a statement is true, without revealing the underlying data. This approach holds tremendous promise for preserving confidentiality within decentralised systems. Further research priorities include:

- **Complex Logic Proofs**
  Simplistic zero-knowledge systems are becoming established in certain cryptocurrency projects, where the goal is mainly to hide transaction amounts or sender identities. However, regulatory compliance in diverse sectors (e.g., healthcare or financial auditing) typically involves more complex queries - like verifying that a patient’s identity and insurance details are valid or that a transaction meets specific criteria. Future zk-SNARKs must allow for these richer, fine-grained proofs, enabling permissioned disclosures while still protecting all sensitive details.

- **Layer 2 Integrations**
  Incorporating zero-knowledge proofs into Layer 2 networks can provide confidentiality for transaction batches, reducing cost and congestion. Applied research might consider how to unify multiple Layer 2 solutions that each have unique privacy mechanisms, or explore the feasibility of “Layer 3” environments dedicated entirely to advanced, privacy-preserving computations.

8.2.3 Regulatory Alignment for Privacy

Finally, ensuring that privacy-preserving techniques are well-received by regulators demands a constructive engagement strategy. Authorities rightly worry that anonymity features could be misused if not subject to oversight. Hence, future work should identify technical designs that allow selective disclosures, secure audits, or “viewing keys” for lawful investigations - enabling compliance while still preserving user confidentiality whenever possible.

8.3 Next-Generation Blockchain Standardisation

As blockchain-based solutions become more prevalent in finance, healthcare, and supply chains, established regulations have struggled to keep pace. Although several initiatives attempt to define technical standards and best practices, there is still considerable fragmentation, particularly around the classification of on-ledger and off-ledger data.

8.3.1 Improving Regulatory Frameworks in Finance

Financial authorities worldwide are examining how to govern digital asset exchanges, decentralised finance platforms, and tokenised securities. Future standardisation in this field should emphasise:

- **Robust Disclosure Mechanisms**
  Institutions must be able to demonstrate compliance with anti-money laundering (AML) and know-your-customer (KYC) obligations without undermining the decentralised ethos of blockchain. New technical and procedural guidelines can define how to embed verified data about participants within financial transactions while limiting exposure of unrelated personal details.

- **Clear Treatment of Off-Ledger Assets**
  Tokenised instruments (e.g., bonds, equities, or derivatives) may reference large volumes of supporting documents or represent real-world property. Detailed standards are required to dictate how these off-ledger references are established, validated, and periodically updated, so that compliance checks can be performed without always pushing extensive paperwork on-chain.

8.3.2 Standardisation in Healthcare

Healthcare systems hold highly sensitive patient information and must comply with stringent data protection rules. The move towards decentralised patient records or clinical trial data is compelling but demands refined governance, including:

- **Frameworks for Informed Consent**
  Patients should maintain robust control over medical records, determining which caregivers and institutions may access specific parts of their data. Where off-ledger storage is used (for diagnostic images or extensive medical histories), on-ledger references must be properly scoped so that only authorised stakeholders can verify or view them.

- **Shared Trust Models**
  Healthcare providers, insurers, and regulators must operate within a trust framework clarifying the conditions under which data can be shared or aggregated. Improved standardisation can detail how zero-knowledge proofs or confidential queries may be employed to confirm eligibility, coverage, or treatment validity, without breaching patient confidentiality.

8.3.3 Supply Chain Modernisation

Supply chains often involve numerous stakeholders with varying levels of trust. Distributed ledgers can help streamline the tracking of goods and regulatory checks, yet next-generation standards are required to handle data complexity:

- **Digital Twin Integration**
  Physical goods can be accompanied by digital twins - on-ledger tokens that prove authenticity

and track movements. Large volumes of inventory-specific details (e.g., manufacturing specifications, storage conditions) often stay off-ledger to avoid excessive ledger bloat. Common guidelines are needed to unify how these references are stored, updated, and verified across participants.

- **Compliance Monitoring**
  Regulators increasingly demand real-time visibility into supply chains, including verification that items conform to safety or environmental standards. Future blockchain standards could integrate compliance modules capable of verifying off-ledger documentation (test certificates, inspection reports) and linking these back to a single, transparent chain of custody.

9. Conclusion

This concluding section highlights the key insights gained throughout this work and offers recommendations to foster broader, more effective adoption of blockchain technology in diverse sectors. The analysis emphasises the classification of on-ledger and off-ledger data, the critical nature of security and interoperability standards, and the importance of practical frameworks that support small and medium-sized enterprises (SMEs).

9.1 Summary of Key Findings

9.1.1 Importance of On-Ledger and Off-Ledger Data Classification

One of the primary observations is the growing complexity in handling blockchain data and, in particular, the need to distinguish carefully between what is best kept on-ledger (i.e., stored directly on the blockchain) and what is more suitable for off-ledger storage (i.e., external systems linked cryptographically to the blockchain). This classification is crucial for three main reasons:

1. **Scalability**: Storing large data sets or frequently changing information on the blockchain leads to cost and performance bottlenecks. Handling this data off-ledger, while maintaining verifiability through cryptographic methods, alleviates pressure on the network and ensures a more efficient system overall.

2. **Privacy and Compliance**: Highly regulated environments, such as healthcare and finance, demand protection of sensitive information. Moving personal or confidential data off-chain (with only essential references on the blockchain) makes it easier to comply with data protection laws and privacy regulations.

3. **Cost-Efficiency**: Recording even moderately sized data payloads on a public blockchain can become prohibitively expensive due to transaction fees. By limiting on-ledger storage to essential transaction references or cryptographic proofs, organisations reduce costs while retaining blockchain's immutability for the most critical pieces of information.

9.1.2 Role of Security and Interoperability Protocols in Blockchain

Security
Maintaining robust security across on-ledger and off-ledger environments emerged as an important theme. While the blockchain ledger itself benefits from decentralised validation and cryptographic immutability, off-ledger data must also be tamper-resistant and verifiable. Using hashes, digital

signatures, and Merkle trees helps ensure that data stored outside the chain remains trustworthy and can be linked, if needed, back to an auditable on-chain record.

Interoperability
Interoperability solutions, whether they connect distinct blockchains or integrate blockchains with legacy systems, are essential to smooth data exchange. Yet standardised approaches for interlinking on-ledger and off-ledger data across multiple networks remain underdeveloped. Although industry initiatives and technical proposals exist, the sector as a whole still requires clear technical specifications to handle real-world use cases - particularly where sensitive or voluminous off-ledger data is involved.

9.1.3 Need for SME-Friendly Blockchain Adoption Frameworks

Small and medium-sized enterprises frequently face the greatest hurdles when adopting complex, emerging technologies. For SMEs to embrace blockchain effectively:

- **Clarity**: Straightforward guidelines on when to store data on-chain versus off-chain reduce complexity.
- **Cost-Effectiveness**: Solutions that minimise transaction fees and operational overhead are imperative for smaller organisations.
- **Scalability**: SMEs need flexible systems that can scale with their business growth without sacrificing efficiency.
- **Compliance and Support**: Pre-configured or ready-made frameworks, with compliance features built in, lower the barriers to entry and mitigate the risk of non-compliance.

By developing SME-focused tools - such as simplified data classification models, user-friendly dashboards, and modular security layers - blockchain will become more attractive for businesses that lack extensive in-house technical expertise.

9.2 Final Recommendations

9.2.1 Adoption of Hybrid Ledger-Storage Models

A major outcome from this study is the value of **hybrid approaches** that balance on-ledger and off-ledger data storage. While the immutable ledger remains integral for transaction records, key contractual information, and proof-of-existence, large or frequently changing datasets should reside in off-ledger systems. This architecture enables blockchain to do what it does best - verifiable, permanent records - while off-chain repositories handle large volumes of data securely and at lower cost.

9.2.2 Standardisation of Metadata and Cryptographic Anchoring

To guarantee integrity and trust when data is stored off-ledger, common standards for cryptographic anchoring and metadata are essential. Specifically:

- **Metadata Schemas**: Define consistent data formats so that different blockchains and storage systems can interpret off-chain references and hashes without ambiguity.

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

- **Hashing Protocols**: Use standardised hashing algorithms and methods (e.g., Merkle trees) to ensure integrity and auditability. These can be adopted widely, both within private consortia and across public ledgers, enabling cross-network interoperability.
- **Anchoring Mechanisms**: Establish standard practices for linking off-ledger data to on-chain records, so multiple parties and systems can validate the authenticity of external information.

9.2.3 Regulatory Compliance

Ongoing collaboration with regulatory bodies should underscore the importance of flexible on-ledger/off-ledger designs. Retaining sensitive or personal data off-chain, but anchoring it to the ledger, aligns with data protection mandates while preserving essential blockchain benefits. A clear set of compliance guidelines - from data minimisation and encryption best practices to time-stamped proof-of-existence - would help businesses navigate fast-evolving regulations in different jurisdictions.

Final Thoughts

The conclusion of this work underscores the centrality of prudent data classification, robust security provisions, and consistent interoperability standards to ensure that blockchain solutions mature in a sustainable manner. Organisations eager to adopt distributed ledger technologies - especially smaller businesses - should rely on **hybrid on-ledger/off-ledger architectures**, fortified by **cryptographic proofs** and guided by **emerging industry standards**. In doing so, they can maximise the strengths of blockchain technology, mitigate legal and regulatory risk, and establish a strong foundation for scalability and trust in the digital economy.

10. References

Below is a selection of documents spanning academic research (IEEE, ACM, Elsevier), technical white papers (Hyperledger and Ethereum Enterprise Alliance), as well as key regulatory frameworks (GDPR, MiCA, eIDAS). They are listed in a Harvard-style format, each with a name/title and link. These examples serve as a useful starting point for anyone studying or working in blockchain and distributed ledger technology.

1. ACADEMIC RESEARCH PAPERS

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, pp. 55–81. Available at: https://doi.org/10.1016/j.tele.2018.11.006

Gervais, A., Karame, G.O., Wüst, K. et al. (2016) ‘On the security and performance of proof of work blockchains’, Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS ’16), Vienna, Austria. New York, NY: ACM, pp. 3–16. Available at: https://dl.acm.org/doi/10.1145/2976749.2978341

Zheng, Z., Xie, S., Dai, H., Chen, X. & Wang, H. (2018) ‘Blockchain challenges and opportunities: A survey’, International Journal of Web and Grid Services, 14(4), pp. 352–375. Available at: https://doi.org/10.1504/IJWGS.2018.10016848

2. TECHNICAL WHITE PAPERS

Technical Report on On-Ledger and Off-Ledger Data Management Standards (Petko Karamotchev, Final and Inclusive of Pilot Testing and Expert Input, 01/04/2025)

Ethereum White Paper (2014) Vitalik Buterin. Available at: https://ethereum.org/en/whitepaper/

Hyperledger Architecture Working Group (2017) Hyperledger Architecture, Volume 1: Consensus (White Paper). Available at: https://www.hyperledger.org/wp-content/uploads/2017/08/Hyperledger_Arch_WG_Paper_1_Consensus.pdf

Enterprise Ethereum Alliance (2018) EEA Enterprise Ethereum Architecture Stack (White Paper). Available at: https://entethalliance.org/wp-content/uploads/2018/05/EEA-Enterprise-Ethereum-Architecture-Stack.pdf

3. REGULATORY FRAMEWORKS

European Union (2016) Regulation (EU) 2016/679 (General Data Protection Regulation, GDPR). Official Journal of the EU, L 119, 4 May. Available at: https://gdpr-info.eu/

European Commission (2023) Regulation on Markets in Crypto-assets (MiCA), 2020/0265 (COD). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020PC0593

European Union (2014) Regulation (EU) No 910/2014 (eIDAS). Official Journal of the EU, L 257, 28 August. Available at: https://eur-lex.europa.eu/eli/reg/2014/910/oj