Computer Communications 239 (2025) 108180

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

VESPACE: A verifiable blockchain-based data space solution to empower the data economy

Andrea Roberta Costagliola [a], Carlo Mazzocca [b], Armir Bujari [a,*], Rebecca Montanari [a], Paolo Bellavista [a]

[a] Department of Computer Science and Engineering, University of Bologna, Bologna, Italy
[b] Department of Information and Electrical Engineering and Applied Mathematics, University of Salerno, Fisciano, Italy

ARTICLE INFO

Dataset link: https://github.com/AndreaCostagliola/VSPACE

Keywords:
Data space
Blockchain
Decentralized Identifier
Verifiable Credential
Decentralized storage
Data economy

ABSTRACT

In the rapidly evolving data economy, the ability to securely and efficiently share data between organizations has become paramount, unlocking new opportunities for innovation and growth. In this context, different initiatives have worked on conceptual proposals and enabling technological building blocks, addressing design aspects of data spaces. However, the current landscape lacks practical implementations and integration of secure data-sharing primitives supporting a decentralized data ecosystem. To this end, we conduct an analysis of previous efforts and initiatives, identifying gaps. We then introduce VESPACE, a blockchain-based platform for data spaces that enables participants to selectively and securely share verifiable data with authorized users while maintaining control over their access. Our framework incorporates data sovereignty principles implemented through Decentralized Identifiers (DIDs), Verifiable Credentials (VCs), and blockchain technology, qualifying decentralized identity and access control as key features to establish user trust in decentralized data ecosystems. We present a prototype system implementation of VESPACE, evaluating the design choices, showcasing the feasibility of our proposal.

1. Introduction

The data economy is paramount, as data-driven innovation accelerates economic growth and competitiveness. A study by the European Commission reports that in 2019, the data economy contributed €300 billion in value added, representing 2% of the European Gross Domestic Product [1]. This figure is projected to more than double, reaching €739 billion by 2025. The use of data-informed approaches has already transformed key sectors such as manufacturing, transportation, and healthcare.

One of the primary drivers behind the explosion of data is the Internet of Things (IoT), which encompasses everything from smart home devices to industrial machines. In addition to this, social media platforms contribute significantly, as millions of users continuously generate and share content. Although this vast amount of data has no intrinsic value, its effective utilization can unlock unprecedented insights and drive innovative applications across industries [2].

Artificial Intelligence (AI) is a key enabler in transforming raw data into a powerful resource for decision-making and automation [3]. In the context of smart cities, Digital Twins (DTs) [4], virtual representation of physical systems, can leverage Machine Learning (ML) and IoT data to optimize city operations [5]. For example, DTs help reduce congestion and improve mobility [6]. Achieving such an optimization requires cooperation among various stakeholders in the smart city, as addressing congestion requires integrating data from traffic cameras, GPS signals, public transport schedules, and weather conditions. By breaking down data silos across these domains, cities can fine-tune traffic light timings, recommend alternative routes, and dynamically adjust public transportation schedules to alleviate bottlenecks.

This example emphasizes the critical need for a specialized data platform that enables secure and seamless data sharing among different stakeholders, including both data producers and consumers operating in various domains. Such a platform must empower data owners with granular control on how, when, and where their data can be accessed within the network [7]. System users should be able to track and audit operations performed on the data to detect or prevent privacy violations.

Equally important is the assurance for data consumers that the information originates from trusted and verified sources, such as municipalities, public agencies, or reputable companies. To achieve this, data management within the platform should be designed to operate

* Corresponding author.
E-mail addresses: andrea.costagliola@unibo.it (A.R. Costagliola), cmazzocca@unisa.it (C. Mazzocca), armir.bujari@unibo.it (A. Bujari), rebecca.montanari@unibo.it (R. Montanari), paolo.bellavista@unibo.it (P. Bellavista).

https://doi.org/10.1016/j.comcom.2025.108180
Received 14 July 2024; Received in revised form 18 February 2025; Accepted 14 April 2025
Available online 29 April 2025

within a decentralized ecosystem, adhering to the principles of Findability, Accessibility, Interoperability, and Reusability (FAIR) [8]. This approach not only fosters trust and transparency, but also improves the overall efficiency and effectiveness of the use of smart city data.

Although previous efforts have been dedicated to developing novel data-sharing platforms, they primarily focus on supporting data streams, overlooking real-world scenarios where users need access to previously generated data. These solutions typically rely on centralized trust and do not address the challenges related to the verifiability of shared data [9,10], which is one of the key contributions of our work. Furthermore, they are not aligned with Self-Sovereign Identity (SSI) [11] principles, as data owners do not retain direct control over their data and data access is regulated through intermediaries.

To fill these gaps, this article first outlines the requirements drawn from the reference data space initiatives. Building on these foundations, we present a Verifiable Blockchain-based Data Space Solution to Empower the Data Economy (VESPACE), a decentralized data sharing platform that embodies the principles of SSI by leveraging Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) [12]. These key standardized technologies empower users in the data ecosystem with full control over their digital identities and data. Each participant in the platform is uniquely identified via a DID. Data owners use VCs to verify the authenticity of the dataset and to manage consumer access rights. Since these credentials are issued directly by data producers, they retain the ability to revoke access at any time, providing dynamic and flexible control over data sharing. Our platform also archives auditing of data access, as information related to dataset certification and access rights is recorded on the blockchain. Datasets are stored in a decentralized storage system managed by the service/participants of the data platform. To validate the design choices, we implemented a prototype of VESPACE and comprehensively evaluated its performance with different open datasets from the municipality of Bologna, Italy. The main contributions of this work can be summarized as follows:

- We analyze existing standards and reference projects for data spaces to identify guidelines for designing a decentralized data sharing ecosystem that ensures verifiable data, audits operations, and enables fine-grained access control;
- We propose VESPACE, a verifiable, blockchain-based data space that adheres to the identified design guidelines while aligning with FAIR and SSI principles;
- We implement a prototype of the proposed framework and evaluate its performance, demonstrating its feasibility using actual datasets and synthetic triggered actions to assess the scalability of the security primitives.

The remainder of this paper is structured as follows. Section 2 presents a concise survey of relevant European projects on data platforms and data spaces, followed by a discussion of the key technological building blocks and standards that support our proposal. Section 3 draws the identified design guidelines for implementing a secure and verifiable platform for sharing data, while in Section 4, we present VESPACE, a verifiable blockchain-based data space that meets the identified requirements. Section 5 describes how VESPACE achieves secure and verifiable data sharing, followed by Section 6 where we assess aspects of VESPACE employing real datasets. Section 7 discusses the main related works and compares them according to their compliance with the identified design requirements. Finally, Section 8 draws the conclusions and identifies future work.

2. Background

In the following, we present a concise review of the main European initiatives from which we draw the design guidelines for VESPACE. We begin by examining projects that advocate for a federated ecosystem of data platforms, moving to data spaces, identified as decentralized data ecosystems where trust is paramount. Finally, we explore the technological building blocks and reference security standards essential for establishing trust and enforcing data sovereignty principles within a decentralized data ecosystem.

2.1. Towards a decentralized ecosystem of data

Different European initiatives promote the concept of open data by adapting functional building blocks and standards to specific operational domains, leveraging federated software systems to store, manage, process and analyze data. As an example, the European Open Science Cloud (EOSC) [13] is a federation of data platforms promoted by the European Commission to provide a scalable, interoperable, and secure digital ecosystem for research, innovation, and education. EOSC integrates distributed data repositories, computing infrastructures, and research tools, enabling seamless access to scientific resources across following a federated identify provider model. By adhering to the FAIR principles, it ensures metadata harmonization, cross-platform interoperability, offering a robust framework for managing and processing research data at scale. In parallel, OpenAIRE [14] acts as a decentralized Open Science platform that aggregates, enriches, and interlinks research output, including publications, datasets, and project metadata, from multiple repositories and data providers. Its technical backbone consists of metadata aggregation services, knowledge graphs, and machine-readable interoperability frameworks that facilitate automated content linking, semantic enrichment, and research impact tracking.

Moving beyond the notion of centralized and/or federated data platforms, data spaces are recognized as decentralized environments based on standards and structured data that enable the secure sharing of data between different organizations, systems, and users [15]. A data space encompasses a broader ecosystem, integrating diverse data types, including proprietary data, across multiple domains, fostering a more flexible and expansive data-sharing environment.

In this context, the International Data Space (IDS) is a key standard for data spaces [15]. With its integration into Data Spaces Business Alliance (DSBA), founded in 2021 by IDS Alliance, FIWARE, etc., DSBA is emerging as the leading framework for data space creation and interoperability. Currently under specification, DSBA advocates for a secure data sharing architecture that enables cross-company data exchange in a trusted and sovereign environment. The data sharing environment is supported by a decentralized architecture and relies on standardized interfaces for data exchange, with the aim of promoting secure and transparent data sharing between different organizations, sectors, and countries [16].

VESPACE draws from the initiatives mentioned above and proposes an architecture-agnostic approach that can integrate new dataspaces and data platforms that emerge, allowing, for instance, to introduce new standards, data structures, and ontologies, while enabling secure and trustworthy data exchange through a fine-grained access control mechanism built via decentralized technologies.

2.2. Decentralized technologies for data sovereignty

The key objective of Self Sovereign Identity is to empower data owners with full control over their personal information, allowing them to decide what data to share, with whom, and when. The concept of data sovereignty, pictorially depicted in Fig. 1, is typically achieved by using Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) [12], both recently standardized by the World Wide Web Consortium (W3C). A DID is a unique identifier that distinguishes entities within a decentralized system. Each DID is linked to a cryptographic key pair and is recorded in a DID Document, which is stored on a verifiable data registry (Step 1 in Fig. 1), such as a blockchain. The DID Document contains publicly accessible information, including the public key, which facilitates decentralized identity verification.

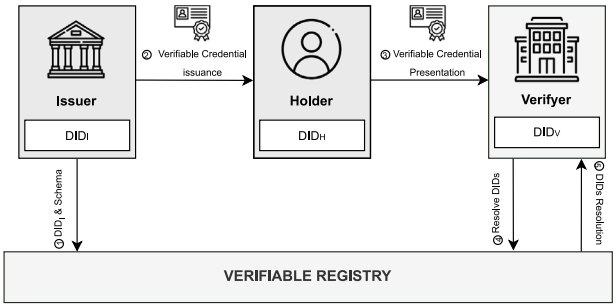

Fig. 1. SSI overview.

While DIDs handle identification, VCs enable the secure and verifiable sharing of information. A VC is digitally signed by a trusted authority, such as a government, to attest to specific claims or attributes about an entity (Step 2 in Fig. 1). It includes references to both the credential holder and issuer (via their DIDs). When sharing a credential, the holder generates a Verifiable Presentation (VP) by signing the VC with their private key (Step 3 in Fig. 1). The verifier can then resolve the referenced DIDs (Step 4 in Fig. 1) to retrieve the associated DID Documents (Step 5 in Fig. 1), which provide the public keys of both the issuer and the holder. Using the issuer's public key, the verifier confirms the credential's authenticity, while the holder's public key ensures ownership of the credential. Notably, because DID Documents are stored on a verifiable registry, verifiers can authenticate credentials without direct interaction with the issuer. Since VCs may contain sensitive information, they are typically shared off-chain.

In VESPACE, DIDs are used to identify participants in the system. Unlike traditional SSI implementations, where trusted organizations issue VCs to certify attributes of an entity, VESPACE allows data owners to issue credentials directly. These VCs certify the datasets that owners wish to share and grant access permissions to data consumers. By issuing VCs themselves, data owners maintain full control, including the ability to revoke credentials when necessary. To ensure transparency, VCs that certify the authenticity of the data are shared on the blockchain, while datasets themselves are stored using distributed content-addressable storage systems such as the InterPlanetary File System (IPFS) [17].

3. Design guidelines

Designing a data space environment where customers can directly share and verify the authenticity of data is paramount. In this section, we outline some of the principles for designing a verifiable data space that facilitates cross-domain data sharing.

3.1. Functional requirements

This subsection highlights some fundamental functional requirements of a dataspace environment in accordance with the EU initiatives discussed previously.

FR-1 Data Sovereignty. In data spaces, data sovereignty refers to the ability of organizations, governments, and individuals to maintain full control over their data [18]. This capability encompasses how data are shared and used by others, as users should have the ability to dynamically update access to their information at any time. IDSA [19] offers guidelines and a framework for ensuring data sovereignty in data spaces, emphasizing the importance of clearly defined data usage policies and contracts. This approach ensures that entities maintain control over their data, sharing them based on their specific preferences and conditions.

FR-2 User Management. A data space must protect all participants and systems within the ecosystem. All entities involved should be registered, and mutual authentication must be established before accepting any requests. Users can interact with the system and perform actions based on their authorization privileges, which determine their permitted operations. These include access to specific datasets and authorization to share data. Additionally, since access conditions may change, the system must include mechanisms to dynamically update user permissions.

FR-3 Discovery and Selection. Given the potentially large amount of heterogeneous data that populate data spaces, product discovery and selection are crucial capabilities of dataspaces. The system should empower users to search for information within the data space, supporting queries that match their specific requirements with relevant products. In EOSC [13], these features are implemented through catalogs (e.g., for datasets, services, standards), which allows identifying information through machine-readable metadata. Thus, each dataset should be associated with a comprehensive set of metadata detailing its attributes and characteristics, enabling users to conduct targeted searches based on specific criteria. This functional requirement addresses the first FAIR principle of findability, ensuring that users can locate and access data products efficiently.

FR-4 Data Access. This requirement aligns with the second FAIR principle of accessibility. It is essential to have well-defined data access policies that allow data owners to decide with whom to share information and when to deny access to specific users. These policies should ensure that only authorized users can access sensitive data while also making public data readily available to the research community. This granular control over data access helps protect the privacy and security of information while maintaining a high level of accessibility. Authorized users must therefore be able to access and share data through trusted data repositories, ensuring the long-term sustainability of research data. OpenAIRE guidelines for repository managers highlight the importance of ensuring data accessibility through standard protocols and facilitating data reuse by complying with community standards [14].

FR-5 Data Transfer and Exchange. A key functionality of data spaces is the transfer of data from one participant to another. Data owners should be able to publish verifiable data, while consumers can easily retrieve them. Data space platforms must ensure that consumers can always detect unauthorized alterations of requested data and that data delivery is efficient, minimizing latency between data access and delivery. For example, the DBSA reference architecture incorporates mechanisms to detect unauthorized alterations, thus supporting efficient and secure data transfers.

FR-6 Data Interoperability and Portability.

Secure and efficient data transmission across platforms is vital for collaboration and maximizing data utility. Projects and initiatives like EOSC and IDSA leverage interoperability to enable seamless, mutually beneficial data exchange. Standardized protocols and formats ensure dataset integration while preserving integrity and context. Embracing interoperability fosters dynamic ecosystems where information flows freely, driving data-driven insights, innovation, and cross-border collaboration.

FR-7 Auditing and Compliance. Data spaces must monitor and record the entire lifecycle of data, with particular emphasis on data accesses. This capability is crucial for detecting unauthorized use or potential data breaches. Monitoring operations are also necessary to ensure compliance with existing data protection regulations, such as GDPR. OpenAiRE infrastructure supports the tracking of data usage and compliance with legal and ethical standards, which is crucial for maintaining data integrity and ensuring that data reuse adheres to established norms.

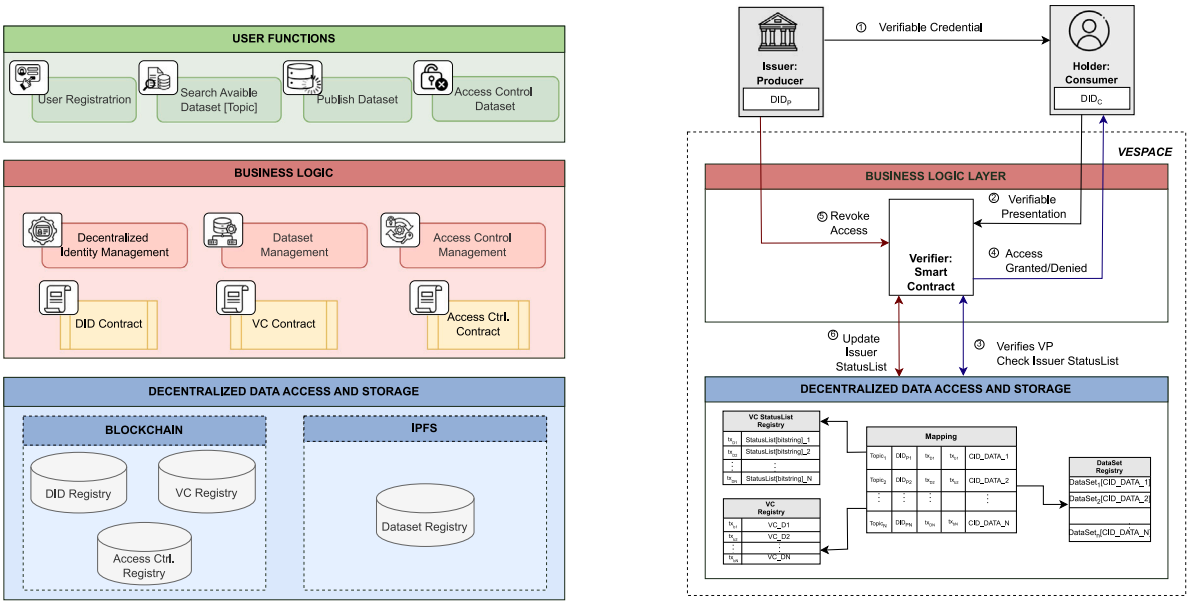

(a) 3-tier architecture.

(b) Data access management.

Fig. 2. On the left is shown VESPACE architecture comprising user functions, which act as abstractions built on the business layer's functionalities. The business layer manages data transactions through on-chain interactions while facilitating storage and retrieval via decentralized storage systems. The right side illustrates the flow of interactions for granting and revoking data access to consumers.

3.2. Non-functional requirements

Non-functional requirements define the qualitative characteristics and constraints of a system, influencing how it meets functional requirements. In this subsection, we outline the primary non-functional requirements for verifiable data spaces.

NFR-1 Security and Privacy. Data spaces must guarantee secure storage and transmission of data. To promote participation, data owners should be confident that their data are managed responsibly and ethically following existing regulations and using secure technologies. Unauthorized entities that are not directly involved in a specific data exchange must be prevented from accessing the data. This requires robust encryption protocols for both data at rest and data in transit, ensuring that data remain confidential and tamper-proof throughout its lifecycle.

NFR-2 Reliability and Usability. The data space must ensure continuous availability and functionality, minimizing downtime and errors. This requirement requires implementing redundancy and fault tolerance measures, as well as robust backup and recovery processes to maintain seamless operations. In addition, the system should be user-friendly and offer intuitive interfaces that simplify the process of identifying and accessing the desired information.

NFR-3 Verifiability and Trustworthiness. Users of data spaces should feel confident that the collected data have been released by reputable organizations and individuals. Verifiability ensures that the data are neither manipulated nor altered during transmission or storage. This is essential to maintain the trustworthiness of the data and to ensure the validity of any evaluations or analyses based on them.

NFR-4 Scalability and Performance. Data spaces should be able to adapt to varying numbers of users and volume of data. In addition, they should support high-performance data analytics use cases and facilitate rapid data transfers between the platform and its participants.

NFR-5 Auditability and Transparency. Users, governing bodies, and regulators should be allowed to conduct comprehensive analyses of operations within data spaces. This involves accessing information on the types of data exchanged, the purposes for which they are used, and the duration of their use. Individuals should be fully informed about their ability to contribute to the data space, the entities with access to their data, the storage locations, the security measures to safeguard their data, and the methods available for interacting with it.

4. VESPACE architecture

VESPACE provides a platform for sharing data between different stakeholders according to the principles of SSI. Each participant in the ecosystem is identified through a DID, which enables accessing corresponding public keys from DID Documents. Data owners, referred to as producers, retain full control over their data, since access is granted through self-issued VCs. Thus, producers can eventually revoke access permission to their generated data. The credentials contain the hash of the corresponding dataset, which guarantees the integrity of the dataset, and the necessary information used to revoke access to the data.

As shown in Fig. 2(a), VESPACE is logically organized in a 3-tier model consisting of a Presentation (User Functions), a Business Logic, and a Decentralized Access & Storage Layer. The Presentation Layer is the entry point for users, offering high level functions to securely share and access data. The Business Logic Layer encapsulates the core functions, handling data processing, access control management, and decision-making operations. Finally, the Data Access Layer manages all the information for the storage, retrieval, and access control rights of data. In the following, we provide a detailed description of the main components of our platform.

4.1. System users

In a data space, users refer to individuals or organizations that derive value from sharing or collecting data. A user may assume the role of a data producer, a data consumer, or both, as detailed below.

Producer. The producer p is any entity that shares data within the data space. Each producer is uniquely identified by a DID (did_p), linked to a pair of cryptographic keys (sk_p, pk_p). Producers act as issuers as

they release VCs to authorize data access and certify data. Given a dataset D, the producer p certifies the authenticity of the dataset with a credential vc_D and provides an authorized consumer $c \in C$ with a credential vc_D^c. These credentials are signed using sk_p (Step 1 of Fig. 2(b)), ensuring decentralized verification of access rights and data authenticity. Furthermore, producers retain fine-grained control over data access and revocation through the W3C Bitstring Revocation List vb_t [20], comprised of a bitstring b_t, where each bit indicates the state of vc_D^c: if the bit is set to 1, the credential is revoked; if it is set to 0, the credential is still valid.

Consumer. Consumers are entities that seek to use datasets shared via VESPACE. Each consumer c is uniquely identified by a DID (did_c), associated with a pair of cryptographic keys (sk_c, pk_c). Only authorized consumers can query the platform and access the data using VCs issued by the producers. Given a credential vc_D^c, VESPACE verifies its authenticity using the producer's public key pk_p. To access data, consumers generate VPs signed with their secret key sk_c as shown in Step 2 of Fig. 2(b). This mechanism prevents replay attacks, as only the legitimate holder has sk_c corresponding to did_c referenced in vc_D^c.

4.2. User functions

Part of a presentation layer, they offer an interface to producers and consumers to use VESPACE's high-level abstractions. The platform is implemented as a Decentralized Application (DApp), which interacts with the back-end of the blockchain through smart contracts. The essential functionalities are:

- *User Registration*: Register users by presenting their DID and necessary information as VC issued by a trusted authority such as a government or any accredited institution;
- *Publish Dataset* and correspondingvc_D acting as dataset certification;
- *Search Available Dataset*: View available dataset metadata-based research such as topic t_D;
- *Access Control Dataset*: This function allows data producers to grant or revoke access permission by updating the access control list associated with that dataset.

4.3. Business logic

The business logic is the core of the VESPACE platform as it encapsulates the main functionalities of our system and the processes needed to manage and regulate access to data. This layer relies on blockchain technology and smart contracts, which ensure the secure, transparent, and automated execution of key operations. In VESPACE, the blockchain stores the DID Document of each user, vb_t and vc_D for each dataset D, and the binding with the information to manage data access. This eliminates the risk of tampering and establishes a transparent audit trail accessible to all authorized parties.

User functions abstract several management flows enacted at the business logic level used to, e.g., control, grant or revoke access to a dataset. Most of these operations are supported by smart contracts that require interactions with the blockchain:

- *Decentralized Identity Management*: Support DID and VC operations, such as their verification during the registration phase, uploading both vc_D to verify data set authenticity and vc_D^c for operations related to access control rights.
- *Dataset Management*: Manages operations related to storing, retrieving, and modifying datasets.
- *Access Control Management*: Tracks every dataset uploaded by each producer p, its Topic t_D, did_p, and the associated certification vc_D and revocation list vb_t. It checks if a consumer is allowed to access a specific dataset.

Indeed, smart contacts play a key role in the VESPACE architecture, used to connect applications that implement business logic with the data access layer. To deliver VESPACE services, we rely on a collection of smart contracts deployed on a blockchain as follows:

- *DID Contract*: Manages the DID lifecycle, including registration and resolution.
- *VC Contract*: Manages all credentials involved in the registration, certification, access, and revocation processes.
- *Access Control Contract*: Handles all the operations on the Access Control Registry (ACR), which maintains a binding among the dataset topic t_D, did_p, the transaction identifier tx_b corresponding to the dataset's revocation list, the transaction identifier tx_D linked to the dataset's certification, and the content identifier cid_D used to reference the dataset itself.

4.4. Decentralized data access and storage

The Data Access Layer is essential for efficient and secure data access management within VESPACE. This layer comprises a blockchain and a storage component, implemented through IPFS.

Decentralized Storage. In VESPACE, datasets are stored in an IPFS cluster [21]. Each dataset D is assigned a content identifier cid_D which acts as a unique identifier used to denote and retrieve the dataset upon authorization. Distributing data across multiple nodes reduces server load and improves scalability, allowing our architecture to handle increased user traffic without performance degradation. Moreover, IPFS improves resilience by seamlessly retrieving data from alternative nodes if primary access points fail, thus increasing VESPACE availability and fault tolerance. Finally, IPFS uses highly connected pinning services and parallelization, which makes it particularly beneficial for large or frequently accessed datasets.

Verifiable Registry. We use the blockchain to record DID documents, VCs, and the Access Control Registry. Whenever a dataset is saved to IPFS, a vc_D is issued and stored on the blockchain to certify its authenticity. As a result, the blockchain returns a confirmation transaction tx_D. Simultaneously, the producer p generates a status list vb_t, saves it on the blockchain, and stores the returned tx_b along with the t_D, the did_p, the tx_D and the dataset cid_D in the ACR. This data structure allows the efficient retrieval of all information to prove the dataset authenticity, also regulates access.

The transparency and immutability of the blockchain ensure that the information is securely stored, preventing any alteration and boosting the trustworthiness of participants in using the data space. Our blockchain is implemented through the Sepolia testnet, which allows testing smart contracts and DApps in a real-world Ethereum environment. To access the blockchain and interact with Sepolia's nodes, we use Alchemy, which provides reliable, scalable, and seamless communication with the blockchain.

5. Secure and verifiable data sharing

This section describes how VESPACE enables secure and verifiable data sharing, meeting the requirements identified in Section 3. It should be noted that all participants must undergo a registration process, in which their identity is verified based on DIDs and VCs issued by a trusted organization, such as a government. In Data Sharing and Certification, producers make their data available and certify them, leveraging vc_D. In the data retrieval and verification phase, consumers interact with VESPACE to retrieve datasets related to a given topic. The VESPACE checks against the access policies for the requested data, and if the user can access the information, it grants the requested data. Finally, in the access control management phase, producers can interact with VESPACE to check and update the list of users who can access the information.

5.1. Registration

To interact with VESPACE, users must be registered within the system. We operate under the assumption that the entities involved adhere to the SSI framework, where participants have DID associated with a VC issued by trusted entities such as governments or other authorized third parties, ensuring the reliability of producers. This approach aligns with the principles outlined in the eIDAS 2.0 regulation, enacted in May 2024, which aims to enhance trust and security in electronic transactions throughout the European Union by promoting the use of digital identities and trust services [22].

In VESPACE, both producers and consumers have a key pair associated with their DID, generated by the Ed25519 algorithm $(1\lambda) \rightarrow (sk, pk)$. Each user also has a VC issued by a trusted authority, which is used to attest some of their properties to the User Agent for registration purposes. As users must already be registered with the entity issuing the VC, verification of the data provisioned is realized by interacting with the blockchain storing DID Documents, which we refer to as the Identity Blockchain. In case the verification process succeeds, VESPACE registers the user information on the blockchain and issues a new credential, namely, VC Registration vc_r^c, serving as an authorization token to interact with the platform.

This registration procedure meets the requirements of FR-1 (Data Sovereignty), FR-2 (User Management), and NFR-1 (Security and Privacy). Moreover, using the VC standard also satisfies the FR-6 requirement of Data Interoperability and Portability.

Algorithm 1 VESPACE data sharing and certification.

```
Input Initialization
h_D ← Hash Dataset using SHA256 Algorithm
vc_D ← issue VC to certify Dataset
vb_t ← issue VC status list containing b_t managing access to the dataset

function PublishDataset(
      did_p, D,
      vc_D[t_D, h_D, MD], vb_t )
result_vc ← Verify vc_D using pk_p on VESPACE
if result_vc then
    result_id ← Verify vb_t using pk_p on VESPACE
    if result_id then
        tx_D ← store vc_D
        tx_b ← store vb_t
        cid_D ← upload D to IPFS
        newEntry ← Update ACR
      t_D, did_p, tx_D, tx_b,
      cid_D )
    end if
end if
```

5.2. Data sharing and certification

Authorized producers can share data in VESPACE following the data sharing and certification process reported in Alg. 1. First, the producer produces a hash of the dataset D using the SHA256 hash algorithm $h_d \leftarrow h_{sha256}(D)$. Then, it issues a self-signed VC:

$$vc_D = [did_p \| t_D \parallel h_D \| m_D] \tag{1}$$

where t_D is the topic associated with the dataset, h_D is the dataset hash and m_D contains metadata, including the title, detailed description, quality certifications, provenance information, and other pertinent details that improve the reliability and usability of the dataset. Metadata offers consumers a clear view of the data available in the data space, in compliance with FR-3 (Discovery and Selection). Furthermore, for each consumer c, the producer issues a vc_D^c containing its DID did_p, the DID of the consumer did_c, the topic associated with that dataset t_D, and the index i of the credential in the bitstring b_t:

$$vc_D^c = [did_p \| did_c \parallel t_D \| i] \tag{2}$$

Then, a VC status list vb_t is generated, used to regulate access to D. The dataset D, the corresponding certification vc_D, and the revocation list vb_t (FR-1) are shared with VESPACE, while consumers are provided with vc_d^c.

VESPACE resolves the DID of the producer did_p contained in the credential, and collects the corresponding public key pk_p from its DID Document. The public key is used to verify the authenticity of vc_D and vb_t (FR-5). Once the verifications are confirmed, VESPACE saves the vc_D certifying the dataset on the blockchain, along with its vb_t. This produces two transaction identifiers, respectively, tx_D and tx_b, which will be used to retrieve information in the following operations. The dataset is then stored on IPFS (FR-3), which binds it to a content identifier cid_D. Finally, all this information is included in a new entry of the ACR, which is structured as:

$$acr \leftarrow [t_D \| did_p \| tx_D \| tx_b \| cid_D] \tag{3}$$

Storing the vc_D on the blockchain, the dataset on decentralized storage, and the subsequent ACR mapping on the blockchain fulfills the functional requirements FR-7 and FR-8. It is worth noting that a producer can associate more data to the same topic.

Regarding non-functional requirements, our proposal adheres to the NFR-2, NFR-4, and NFR-5 requirements. The use of blockchain and decentralized storage improves NFR-2 by providing fault tolerance, data redundancy, and continuous availability, reducing downtime risks. NFR-4 is addressed by leveraging decentralized architectures that support increasing data volumes and user demands while ensuring efficient data discovery and rapid retrieval through content identifiers. Finally, NFR-5 is inherently supported by the immutability of blockchain records, which store verifiable credentials and revocation lists, allowing traceability, verification of data provenance, and comprehensive oversight of data exchange operations within VESPACE.

Algorithm 2 VESPACE data retrieval and verification.

```
Function SearchAvaibleDataset(Topic)
vc_D^c[did_p, did_c, h_D, i] ← c request(p,Dataset[t_D])
vp_D^c ← c issue VP from vc_D^c
result ← Verify vp_D^c on VESPACE
if result then
    for (topic, did, tx_D, tx_b, cid) ∈ ACR do
        if topic == t_D then
            if did == did_p then
                tx_D ← tx_D
                tx_b ← tx_b
                break
            end if
        end if
    end for
    vc_D ← recover VC Dataset using tx_D
    vb_t ← recover VC status list using tx_b
    b_t ← extract bitstring from vb_t
    i ← extract index from vp_D^c
    if b_t[i]! = 1 then
        cid_D ← from ACR
        D ← IPFS[cid_D]
        return D, vc_D
    else
        Access Denied
    end if
end if
```

5.3. Data retrieval and verification

Consumers can interact with the platform to browse and access datasets related to topics of interest. Algorithm 2 details the interactions among consumers and VESPACE to retrieve datasets for a

specific topic. Given a valid credential vc^c_D, the consumer c generates a verifiable presentation vp^c_D by signing it with its own private key sk_c. VESPACE verifies the signature by resolving $did_c \in vc^c_D$, and obtaining pk_c.

Before accessing the dataset, VESPACE must ensure that the consumer's access has not been revoked (FR-4). Using t_d and did_p contained in vc^c_D, VESPACE retrieves tx_D and tx_b from acr for each D associated with t_d. These transaction identifiers are used to retrieve vc_D and vb_t, respectively. Our platform verifies whether the following equation holds $b_t[i] \neq 1$ where $i \in vc^c_D$ (Step 3 of Fig. 2(b)). If the equation is verified, VESPACE retrieves D through cid_D (Step 4 of Fig. 2(b)).

Consumers receive the datasets along with associated vc_D. For each retrieved dataset D_r, the consumer verifies if $h_{sha256}(D_r) = h_D \in vc_D$ (FR-5, NFR-3). This allows each consumer to directly verify the authenticity of the collected data in a fully decentralized manner.

5.4. Access control management

In VESPACE producers have direct control over who can access their data, including the right to decide when a consumer can no longer access them (FR-1). This property is aligned with the SSI principles. Algorithm 3 shows operations performed by VESPACE to revoke access for a topic t_D to a specific consumer c. Specifically, the producer is required to update b_t setting to 1 the bits in the ith position where $i \in vc^c_D$ of the consumer whose access rights must be revoked (Step 5 of Fig. 2(b)). Consequently, the updated version of b_t is stored on the blockchain, generating a new tx_b. Finally, the ACR is updated by adding a new entry that contains the transaction identifier of the latest computed status list (Step 6 of Fig. 2(b)).

Algorithm 3 VESPACE Data access management.

```
1: Function Revoke Access Dataset(vc^c_D, vb_t)
2: i ← extract index from vc^c_D
3: extract b_t from vb_t
4: search i in b_t
5: b_t[i] ← 1
6: Update b_t in vb_t
7: issue vb_t
8: tx_b ← store vb_t
9: Update ACR with new tx_b
```

6. VESPACE evaluation

This section provides a comprehensive evaluation of VESPACE. In particular, we designed a series of experiments to evaluate the dataset certification and access control functionalities, assessing the scalability of the mechanisms when varying the number of participants and datasets involved in the various operations. Moreover, our experiments also aim to assess the usability of the system for added latency. To this end, we use the Response Animation Idle Load (RAIL) model, a performance model proposed by Google that provides metrics to assess the usability of a platform from a user perspective [23]. All the experiments were executed 50 times and results were averaged.

6.1. Implementation settings

We assess a real implementation of the VESPACE available in [24]. The generation and verification of DIDs, DID documents, and VCs are performed using the Digital Bazaar library [25], a widely adopted solution for managing digital identities and credentials in compliance with W3C standards. For blockchain-based operations, we utilize the Sepolia Ethereum testnet. The smart contracts, written in Solidity, primarily handle mappings between topics, producer DIDs, transaction IDs, CIDs, and bitstring-based status lists. Dataset storage is implemented through a decentralized IPFS cluster consisting of three nodes. This cluster employs a Conflict-Free Replicated Datatype (CRDT)-based consensus mechanism to maintain a global pinset, ensuring data availability and redundancy. Furthermore, to evaluate our system in a realistic scenario, we used urban datasets from the Municipality of Bologna, Italy, sourced from the Open Data Bologna Platform [26] under the CC BY 4.0 license. Datasets include diverse urban data, such as parking availability, traffic flow, and environmental metrics, providing a comprehensive, real-world test environment.

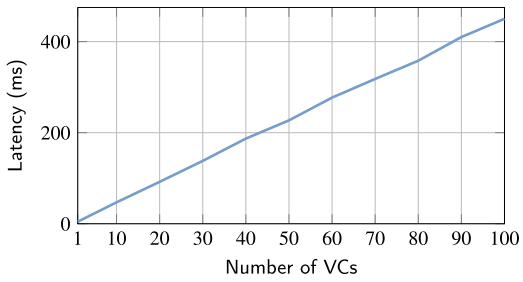

Fig. 3. Latency (in milliseconds) for the issuance of verifiable credentials.

6.2. Dataset sharing and certification

This set of experiments aims to analyze the time required for data owners to certify the authenticity of the dataset and to grant access to consumers. In particular, we measure the latency associated with the process to issue vc^c_d and the corresponding status list contained in vb_t, regulating access in VESPACE. As discussed in Section 5, each vc^c_d also contains an index i, referring to a ith position in the status list, and the dataset topic t_D, providing c with integrity and authenticity. We varied the number of consumers from 1 to 100. The time to generate vc_D can be neglected as each dataset is only certified once. We did not observe remarkable differences while considering different datasets.

Fig. 3 reports the collected results. As expected, the latency grows linearly with the number of credentials issued, ranging from 4,5 ms when issuing a single VC to 449 ms in scenarios with 100 consumers. According to the RAIL model, latency between 100 ms and 1000 ms is perceived as the natural and continuous progression of tasks, allowing users to maintain focus and flow without feeling interrupted or delayed. The results suggest that the issuance process has relatively low overhead as more credentials are generated. The overall latency remains within an acceptable range for real-world applications, ensuring that even a large batch of credentials can be generated rapidly.

6.3. Access right verification

Herein, we evaluate the performance of the system in handling the authorization verification process for a consumer while varying the number of datasets that belong to the same topic. In particular, we consider all the phases needed for varying access rights, which include:

1. Resolving the DID Document associated with did_c to retrieve the consumer's public key and consequently verify the vp^c_D issued (line 4 of Alg. 2);
2. Querying acr to obtain tx_b and tx_D and retrieve the corresponding status list and credentials vc_D that certify the data from the blockchain (line 6–16 of Alg. 2);
3. Assessing whether the ith position of the status list, which corresponds to the index included in the consumer credential is set to 0 (authorized) or 1 (revoked) (line 17–18 of Alg. 2).

We measure the latency of the entire process, reporting the results in the second column of Table 1 and in Fig. 4, which shows the authorization verification time across different dataset counts. Experimental results demonstrate that VESPACE can scale effectively while increasing the number of datasets. This underscores the system's validity in

Table 1
Comparison of Authorization and Revocation Times.

Number of Datasets	Authorization Time (ms)	Revocation Time (ms)
1	668	528
2	1337	1056
4	2648	2112
6	4100	3182
8	5448	4233
10	6690	5284

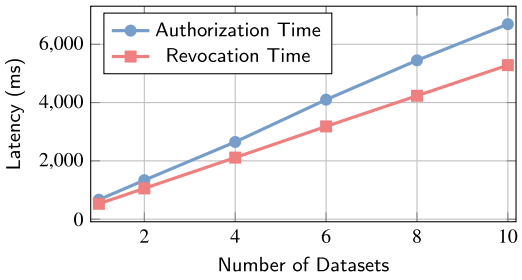

Fig. 4. Comparison of authorization and revocation times.

handling access to multiple datasets belonging to a single topic while maintaining an acceptable response time. It should be noted that we did not measure the time required to download the datasets, as this depends on network connectivity and various cluster settings, which affects all platforms used to share data.

6.4. Access right revocation

Finally, we evaluate the efficacy of VESPACE in revoking access rights, a critical aspect of the SSI paradigm, as it ensures that data owners maintain control over their data over time. We report that the time required to revoke access rights by generating the bitstring status list remains constant when varying the number of consumers from 1 to 100. This is because revocation involves setting to 1 the bits corresponding to the indexes contained in the consumer credentials used to access the data. This introduces an average latency of approximately 27 ms, with no significant variations as the number of revoked users increases.

Moreover, we evaluate the time for managing access revocation when the increasing number of datasets are issued by the same DID producer on the same topic. This involves the generation of an updated bitstring and the update of the access control registry on the blockchain. Table 1 presents the revocation times, which are illustrated in Fig. 4, showing the time required for revoking access while varying the number of datasets up to 10. For a single dataset, the entire process takes approximately 528 ms. As the number of datasets increases, the revocation time scales accordingly. This linear growth suggests that VESPACE exhibits predictable behavior, which is critical when a producer needs to revoke a consumer's access to multiple datasets simultaneously, ensuring a timely and secure update of credential status.

The results indicate that revocation is generally faster than authorization, with an average ratio of approximately 1.26, computed as the authorization time divided by the corresponding revocation time across different dataset sizes. Conversely, the inverse ratio represents the number of authorizations per revocation, highlighting that revocation has a limited impact: its execution time is comparable to that of authorization. This ensures that the system can continue processing new access requests without significant delays.

The primary factor influencing revocation time is the upload of the bitstring to the blockchain. However, since revocations occur less frequently than authorizations, their overall impact on system performance remains minimal. Additionally, the system exhibits scalability

Table 2
Analysis of Relationships and Timing between Authorization and Revocation.

Number of Datasets	Auth/Revoke Ratio
1	1.26
2	1.26
4	1.25
6	1.28
8	1.28
10	1.26

as dataset size increases. As shown in Fig. 4, authorization times scale proportionally with dataset size, while revocation times follow a similar trend but remain consistently lower. This demonstrates that the revocation mechanism remains efficient and does not introduce usability concerns (see Table 2).

7. Related work

In this section, we offer a comprehensive literature review of the main blockchain-based frameworks that favor data sharing among different stakeholders. Table 3 compares the solutions on the functional and non-functional requirements introduced in Section 3.

7.1. Data stream sharing platform

A data stream sharing platform is a system designed to facilitate the real-time exchange of continuous data streams between multiple producers and consumers. FAST [9] is an IoT data marketplace where users are authenticated through DIDs. The blockchain is used to list the available data streams, verify identities, and settle payments. Similarly, Sober et al. [10] propose another blockchain-based IoT data marketplace based on a broker, facilitating data exchange between producers and sellers. The marketplace leverages smart contracts to manage participation and information on data products. Missier et al. [29] proposed a decentralized infrastructure for fair and trusted IoT data trading. Blockchain tracks information on data streams that flow from IoT devices to value-added services deployed on cloud resources.

Bajoudah et al. [30] also present a decentralized structure for IoT data streams that leverages blockchain technology and smart contracts to facilitate secure and trustworthy data exchange without prior trust between participants. Although their approach ensures data integrity and payment settlement through Ethereum-based smart contracts, it focuses mainly on providing data streams and does not address the challenges related to secure storage and verifiability of shared data.

7.2. Decentralized and distributed data storage

Ramachandran et al. [27] lay the foundation for implementing a data space for smart cities based on distributed ledger technologies. As for VESPACE, the blockchain only maintains metadata associated with valuable data, which must be stored off-chain (e.g., through an IPFS Cluster). DEON [28] offers a marketplace for different contexts, including IoT applications such as environmental monitoring and anomaly detection. Similarly to VESPACE, it employs IPFS to store data whose corresponding identifier is maintained on a blockchain. User participation is regulated using DIDs and VCs. However, DEON is mainly proposed as a blockchain-based platform to deploy decentralized off-grid networks, and the framework is only presented from a very high-level perspective.

Bernabé-Rodríguez et al. [31] propose a framework that leverages blockchain and secure multi-party computation (SMPC). The system employs blockchain to guarantee transparency and trust, while SMPC enables privacy-preserving computations, ensuring that data are processed securely without revealing private information. Although their

Table 3
Comparison of data spaces solutions based on function and non-functional requirements. We use "✓" if the requirement is guaranteed, "~" if it is partially met, "–" when it is not addressed, and "✗" if it is not guaranteed.

Requirement	[27]	[28]	[9]	[10]	[29]	[30]	[31]	[32]	Ours
User Management	~	~	✓	✓	✗	~	✗	✓	✓
Data Sovereignty	✗	✓	✓	–	–	–	–	✓	✓
Discovery and Selection	~	✓	✓	✓	~	~	~	~	✓
Data Access	✓	✓	~	~	~	~	✓	–	✓
Data Transfer and Exchange	✓	~	✓	~	✓	~	~	~	✓
Data Interoperability and Portability	✓	✓	✓	✗	✗	✗	–	✓	✓
Auditing and Compliance	✗	✗	✓	✓	✓	✗	✓	✓	✓
Security and Privacy	✓	✓	✓	✓	–	~	✓	✓	✓
Reliability and Usability	–	✓	–	✓	✓	–	✗	~	✓
Verifiability and Trustworthiness	✗	✗	✓	✗	✗	–	✗	✓	✓
Scalability and Performance	✓	✓	✓	✓	✓	~	~	✓	✓
Auditability and Transparency	✓	✓	✓	✓	~	✓	✓	✓	✓

approach emphasizes computational privacy and secure data processing, it does not specifically address the verifiability of shared data and the use of SSI.

Instead, Yoon et al. [32] introduce a blockchain-based personal data exchange system using DID and VC. Users authenticate their identity and claim ownership of the data without relying on centralized systems. Data integrity and transaction history are maintained through a blockchain implemented using Hyperledger Fabric. However, the system mainly focuses on personal data exchange and does not fully address the IoT data space.

7.3. Comparison with previous works

Existing solutions typically use DIDs and VCs for authentication, where VCs contain verifiable claims about an entity that are leveraged to grant access to data sharing platforms. However, these frameworks fail to fully embrace SSI principles, particularly when it comes to granting producers full control over their data. In most cases, credentials are issued by trusted third-party organizations, rather than the data owners themselves, thereby limiting the level of control producers have over their shared datasets. Moreover, many existing proposals overlook revocation entirely, and those that address it often rely on third-party components for this purpose.

In contrast, VESPACE directly addresses these gaps by enabling data owners to revoke access rights to their data in a fully decentralized manner. This is achieved through the W3C bitstring structure, which allows for efficient and direct access control updates by the producers themselves, ensuring full autonomy. To our knowledge, VESPACE is the first solution to regulate access rights using this data structure, providing a groundbreaking approach to decentralized access control and revocation.

Furthermore, while VCs are traditionally used by trusted issuers to certify claims about entities, VESPACE innovatively repurposes VCs to allow producers to self-certify the authenticity of their datasets in a decentralized manner by issuing self-signed credentials. This approach not only empowers producers with greater control over their data but also addresses data provenance, a critical concern that is often overlooked in many related works. By enabling producers to certify the authenticity of datasets, VESPACE enhances trust and transparency in data sharing within dataspaces, which is particularly valuable in contexts where the integrity and origin of data are crucial.

Finally, VESPACE overcomes a significant limitation in existing data stream platforms, which often fail to account for real-world scenarios where consumers may need access to both real-time data and historical datasets. Unlike traditional platforms that focus solely on live data streams, VESPACE ensures that previously shared data can be accessed in a secure, decentralized manner, offering a more flexible and scalable solution for data consumers and producers alike.

8. Conclusions

Data have become an invaluable asset in the modern economy, providing unprecedented value when used to develop new applications and services. This trend is also underscored by substantial investments from organizations such as the European Union, which has recently funded several projects to promote effective data sharing. However, practical solutions are often hindered by obstacles such as the lack of data verifiability and centralized trust model.

To address this gap, this article proposes VESPACE, a verifiable blockchain-based data space designed to empower the next-generation data economy. We developed this data platform by following functional and non-functional requirements identified by reviewing the major European data space projects and initiatives. VESPACE leverages DIDs, VCs and blockchain to establish a decentralized governance model, allowing users to validate data authenticity. We implemented a prototype and evaluated its performance using urban datasets, showing that our framework can efficiently enable secure and verifiable data sharing.

CRediT authorship contribution statement

Andrea Roberta Costagliola: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Data curation. **Carlo Mazzocca:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Data curation, Conceptualization. **Armir Bujari:** Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Data curation, Conceptualization. **Rebecca Montanari:** Validation, Supervision, Methodology. **Paolo Bellavista:** Writing – review & editing, Validation, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work is supported by the PRIN PNRR Digit4Circle project, cod. P2022788KK under the Italian National Recovery and Resilience Plan.

Data availability

In the spirit of open science and reproducibility, we have made available the source code of the proposed solution, which can be consulted in https://github.com/AndreaCostagliola/VSPACE.

References

[1] European Union, A European strategy for data, 2024, URL: https://digital-strategy.ec.europa.eu/en/policies/data-spaces. Accessed on July 2024.

[2] C. Yang, Q. Huang, Z. Li, K. Liu, F. Hu, Big Data and cloud computing: innovation opportunities and challenges, Int. J. Digit. Earth 10 (1) (2017) 13–53.

[3] Q. Qi, F. Tao, Digital Twin and Big Data Towards Smart Manufacturing and Industry 4.0: 360 Degree Comparison, IEEE Access 6 (2018) 3585–3593.

[4] M.J. Kaur, V.P. Mishra, P. Maheshwari, The Convergence of Digital Twin, IoT, and Machine Learning: Transforming Data into Action, in: M. Farsi, A. Daneshkhah, A. Hosseinian-Far, H. Jahankhani (Eds.), Digital Twin Technologies and Smart Cities, Springer International Publishing, Cham, 2020, pp. 3–17.

[5] M.M. Rathore, S.A. Shah, D. Shukla, E. Bentafat, S. Bakiras, The Role of AI, Machine Learning, and Big Data in Digital Twinning: A Systematic Literature Review, Challenges, and Opportunities, IEEE Access 9 (2021) 32030–32052.

[6] B. Lei, P. Janssen, J. Stoter, F. Biljecki, Challenges of urban digital twins: A systematic review and a Delphi expert survey, Autom. Constr. 147 (2023) 104716.

[7] J. Ernstberger, et al., SoK: Data Sovereignty, in: Proc. of IEEE European Symposium on Security and Privacy (EuroS&P), 2023, pp. 122–143, http://dx.doi.org/10.1109/EuroSP57164.2023.00017.

[8] M. Wilkinson, et al., The FAIR guiding principles for scientific data management and stewardship, Sci. Data 3 (1) (2016) 160018.

[9] A. Dixit, A. Singh, Y. Rahulamathavan, M. Rajarajan, FAST DATA: A Fair, Secure, and Trusted Decentralized IIoT Data Marketplace Enabled by Blockchain, IEEE Internet Things J. 10 (4) (2023) 2934–2944.

[10] M. Sober, G. Scaffino, S. Schulte, S.S. Kanhere, A blockchain-based IoT data marketplace, Clust. Comput. 26 (6) (2023) 3523–3545.

[11] A. Mühle, A. Grüner, T. Gayvoronskaya, C. Meinel, A survey on essential components of a self-sovereign identity, Comput. Sci. Rev. 30 (2018) 80–86.

[12] C. Mazzocca, A. Acar, S. Uluagac, R. Montanari, P. Bellavista, M. Conti, A Survey on Decentralized Identifiers and Verifiable Credentials, 2024, arXiv preprint arXiv:2402.02455.

[13] EOSC, EOSC Future Results, 2024, URL: https://eoscfuture.eu/results/. Online; accessed June 2024.

[14] OpenAIRE, Openaire guidelines, 2024, URL: https://guidelines.openaire.eu/en/latest. Accessed on July 2024.

[15] International Data Spaces Association, Why data spaces?, 2025, URL: https://internationaldataspaces.org/why/data-spaces/. (Accessed 6 February 2025).

[16] Data Spaces Business Alliance, Technical convergence, 2025, URL: https://data-spaces-business-alliance.eu/wp-content/uploads/dlm_uploads/DSBA-Technical-Convergence.pdf. (Accessed 6 February 2025).

[17] IPFS, IPFS: The InterPlanetary File System, 2025, URL: https://ipfs.tech. [Accessed 6 February 2025].

[18] P. Hummel, M. Braun, M. Tretter, P. Dabrock, Data sovereignty: A review, Big Data & Soc. 8 (1) (2021) 2053951720982012.

[19] IDS, International Data Spaces, 2024, URL: https://internationaldataspaces.org/. Online; Accessed June 2024.

[20] W3C, Verifiable Credentials Bitstring Status List, 2025, URL: https://www.w3.org/TR/vc-bitstring-status-list/. (Accessed 16 February 2025).

[21] IPFS Community, IPFS Cluster, 2024, URL: https://ipfscluster.io/. Online; Accessed June 2024.

[22] EUR-Lex, EUR-lex, 2024, URL: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202401183. Accessed on July 2024.

[23] RAIL, RAIL, 2024, URL: https://developer.mozilla.org/en-US/docs/Glossary/RAIL. Accessed on July 2024.

[24] A.R. Costagliola, et al., Verifiable Blockchain-based Data Space Solution to Empower the Data Economy: Source Code, 2024, URL: https://github.com/AndreaCostagliola/VSPACE. Accessed on July 2024.

[25] Digital Bazaar, Digital Bazaar GitHub Repository, 2025, URL: https://github.com/digitalbazaar. (Accessed 16 February 2025).

[26] Bologna Municipality, Open Data - Bologna Municipality, 2025, (Accessed 16 February 2025). URL: https://opendata.comune.bologna.it/.

[27] G.S. Ramachandran, R. Radhakrishnan, B. Krishnamachari, Towards a Decentralized Data Marketplace for Smart Cities, in: Proc. of IEEE International Smart Cities Conference, 2018, pp. 1–8.

[28] H. Niavis, N. Papadis, V. Reddy, H. Rao, L. Tassiulas, A Blockchain-based Decentralized Data Sharing Infrastructure for Off-grid Networking, in: Proc. of IEEE International Conference on Blockchain and Cryptocurrency, 2020, pp. 1–5.

[29] P. Missier, S. Bajoudah, A. Capossele, A. Gaglione, M. Nati, Mind my value: a decentralized infrastructure for fair and trusted IoT data trading, in: Proc. of the Seventh International Conference on the Internet of Things, IoT '17, Association for Computing Machinery, New York, NY, USA, 2017.

[30] S. Bajoudah, C. Dong, P. Missier, Toward a Decentralized, Trust-Less Marketplace for Brokered IoT Data Trading Using Blockchain, in: Proc. of IEEE International Conference on Blockchain, 2019, pp. 339–346.

[31] J. Bernabé-Rodríguez, A. Garreta, O. Lage, A Decentralized Private Data Marketplace using Blockchain and Secure Multi-Party Computation, ACM Trans. Priv. Secur. (2024).

[32] D. Yoon, S. Moon, K. Park, S. Noh, Blockchain-based Personal Data Trading System using Decentralized Identifiers and Verifiable Credentials, in: Proc. of International Conference on Information and Communication Technology Convergence, 2021, pp. 150–154.