applied sciences

MDPI

Article

SIBERIA: A Self-Sovereign Identity and Multi-Factor Authentication Framework for Industrial Access

Daniel Paredes-García [1,*], José Álvaro Fernández-Carrasco [1], Jon Ander Medina López [1], Juan Camilo Vasquez-Correa [1], Imanol Jericó Yoldi [1], Santiago Andrés Moreno-Acevedo [1], Ander González-Docasal [1,2], Haritz Arzelus Irazusta [1], Aitor Álvarez Muniain [1] and Yeray de Diego Loinaz [1]

[1] Digital Security Department, Fundación Vicomtech, Basque Research and Technology Alliance (BRTA), Mikeletegi 57, 20009 Donostia-San Sebastian, Spain; jafernandez@vicomtech.org (J.Á.F.-C.); jamedina@vicomtech.org (J.A.M.L.); jcvasquez@vicomtech.org (J.C.V.-C.); ijerico@vicomtech.org (I.J.Y.); samoreno@vicomtech.org (S.A.M.-A.); agonzalezd@vicomtech.org (A.G.-D.); harzelus@vicomtech.org (H.A.I.); aalvarez@vicomtech.org (A.Á.M.); ydediego@vicomtech.org (Y.d.D.L.)

[2] Aragon Institute for Engineering Research, University of Zaragoza, Mariano Esquillor, 50018 Zaragoza, Spain

* Correspondence: dparedes@vicomtech.org

Featured Application

The potential applications of this work span multiple sectors that require secure, user-centric, and privacy-preserving identity management. With a modular design that integrates Self-Sovereign Identity, verifiable credentials, and advanced biometrics, the system is adaptable beyond its primary focus on industrial access control. In critical infrastructure sectors such as energy, water treatment, and transportation, it can ensure that only authorized personnel access sensitive systems by verifying their identity through decentralized IDs and voice or behavioral biometrics while maintaining a verifiable audit trail on the blockchain. In healthcare, doctors can use their SIBERIA wallet to access electronic health records (EHRs), with credential verification (e.g., medical license, hospital affiliation) occurring without storing sensitive information directly, while patients retain control over access to their data in compliance with regulations like the GDPR. In high-security corporate and financial environments, SIBERIA could protect access to sensitive financial data, intellectual property, or corporate intranets. An employee would use their corporate-issued VCs for access, with continuous authentication via behavioral biometrics monitoring their keyboard and mouse patterns to detect potential session hijacking in real-time.

Abstract

The growing need for secure and privacy-preserving identity management in industrial environments has exposed the limitations of traditional, centralized authentication systems. In this context, SIBERIA was developed as a modular solution that empowers users to control their own digital identities, while ensuring robust protection of critical services. The system is designed in alignment with European standards and regulations, including EBSI, eIDAS 2.0, and the GDPR. SIBERIA integrates a Self-Sovereign Identity (SSI) framework with a decentralized blockchain-based infrastructure for the issuance and verification of Verifiable Credentials (VCs). It incorporates multi-factor authentication by combining a voice biometric module, enhanced with spoofing-aware techniques to detect synthetic or replayed audio, and a behavioral biometrics module that provides continuous authentication by monitoring user interaction patterns. The system enables secure and user-centric identity management in industrial contexts, ensuring high resistance to impersonation and credential theft while maintaining regulatory compliance. SIBERIA demonstrates that it is

check for updates

Academic Editor: Gianluca Lax

Received: 12 June 2025
Revised: 22 July 2025
Accepted: 25 July 2025
Published: 2 August 2025

Citation: Paredes-García, D.; Fernández-Carrasco, J.Á.; López, J.A.M.; Vasquez-Correa, J.C.; Yoldi, I.J.; Moreno-Acevedo, S.A.; González-Docasal, A.; Irazusta, H.A.; Muniain, A.A.; Loinaz, Y.d.D. SIBERIA: A Self-Sovereign Identity and Multi-Factor Authentication Framework for Industrial Access. *Appl. Sci.* **2025**, *15*, 8589. https://doi.org/10.3390/app15158589

possible to achieve both strong security and user autonomy in digital identity systems by leveraging decentralized technologies and advanced biometric verification methods.

Keywords: SSI; verifiable credentials; wallet; digital identity; GDPR; access control; biometrics; blockchain; spoofing aware speaker verification

1. Introduction

Traditional authentication mechanisms—such as username–password schemes and centralized identity providers—pose critical challenges in industrial settings, including single points of failure, vulnerability to credential theft, and limited user control over identity data [1].

To address these issues, biometric authentication has become increasingly relevant in modern access control systems, as it enables identity verification based on inherent user characteristics. Among these, voice biometrics has emerged as a particularly effective alternative due to its non-invasive nature and ease of integration in industrial environments, where the deployment of specialized sensors may be constrained. Similarly, behavioral biometrics supports continuous authentication (CA) through the dynamic analysis of user interaction patterns, offering an additional layer of protection against anomalous or unauthorized access [2]. Nevertheless, the intrinsic characteristics of biometric systems, combined with various external factors that may affect their performance, suggest that authentication relying solely on biometrics might not achieve the levels of reliability required in industrial environments. As a result, it is essential to complement these systems with additional mechanisms to strengthen the overall security and robustness of the authentication process.

To address these challenges, SIBERIA is introduced as a novel identity management system designed for secure authentication and authorization in private industrial services. SIBERIA integrates Self-Sovereign Identity (SSI) principles [3], Spoofing Aware Speaker Verification (SASV) authentication, and behavioral biometric monitoring to enhance security and user control. The system consists of three main modules:

- **SSI Module:** Implements a decentralized identity model where users control their information via a mobile wallet. Identity credentials, stored as Verifiable Credentials (VCs), are issued and signed by trusted entities. A private blockchain based on EOSIO manages identity-related smart contracts, ensuring compliance with the European Blockchain Services Infrastructure (EBSI) standards.
- **Secure SASV Module:** Requires users to authenticate using their voice in addition to traditional username and password credentials. The system generates a voiceprint, issues a corresponding VC, and stores it in the user's wallet. During authentication, a similarity score is computed to grant or deny access. This module incorporates anti-spoofing measures and protection against deepfake audio attacks.
- **Behavioral Biometric Authentication Module:** Monitors user interaction with the protected service by analyzing events such as mouse movements, visited options, and session durations. This module detects anomalies in user behavior and triggers alerts or terminates the session in case of suspicious activity.

By integrating these modules, SIBERIA provides a robust and scalable identity management solution that enhances security in industrial environments. This paper presents the architecture, implementation, and evaluation of SIBERIA, demonstrating its effectiveness in securing access to critical industrial services.

1.1. Self-Sovereign Identities

SSI is a user-centric digital identity model that shifts authority from centralized issuers and identity providers to the individual [4,5]. Based on principles of autonomy, privacy, and security, SSI enables individuals to create, hold, and present VCs without relying on a single intermediary. In practice, this decentralization streamlines identity workflows and reduces single points of failure. It also limits large-scale data collection and lowers exposure to breaches and identity theft. Users decide what data to disclose and can revoke compromised credentials. Meanwhile, issuers and verifiers maintain transparency through signed issuance and revocation records. By combining cryptographic assurances with decentralized network infrastructure, SSI offers a resilient, privacy-enhancing alternative to conventional identity management approaches. It aligns with evolving expectations for individual control in the digital age [6].

A core element of SSI architectures is the Decentralized Identifier (DID), a unique and persistent identifier that represents entities—such as individuals, organizations, or devices—without reliance on a central authority. The World Wide Web Consortium (W3C) defines the syntax, structure, and usage of DIDs to promote interoperability and trust across platforms [7]. A DID has the following standardized format:

$$\text{did} : \langle \textit{method} \rangle : \langle \textit{identifier} \rangle \tag{1}$$

The DID contains metadata necessary to verify control over the identifier. This includes public cryptographic keys, authentication methods, and service endpoints. These records are stored in a Verifiable Data Registry (VDR), typically implemented using blockchain or distributed ledger technologies. This allows verifiers to resolve a DID, retrieve its associated document, and authenticate the entity it represents [8].

Several key actors support the issuance, storage, and verification of VCs. Each plays a distinct role within the SSI architecture:

- Holders: Individuals or devices that manage their digital identities within the ecosystem. They store credentials in a secure digital wallet, which acts as the container for their VCs. Holders retain full control over their data and decide when and with whom to share their credentials.
- Issuers: Authorized entities that generate VCs based on the holder's identity. Issuers validate the information and issue credentials, which are registered on the blockchain to ensure authenticity.
- Verifiers: Entities that must validate credentials before granting access to a service. They ensure that the credentials presented by holders are legitimate and have not been revoked, using the blockchain as a source of verification.

In addition to these three actors, the VDR provides the trusted infrastructure for resolving identifiers, retrieving issuer keys, and verifying credential status. All identity-related operations—including registration, issuance, storage, and verification of VCs—are immutably recorded on a blockchain or distributed ledger. This guarantees data integrity, transparency, and resistance to tampering, enabling verifiers to detect any unauthorized changes to the trust framework.

Figure 1 provides a visual overview of the actors involved in the SSI ecosystem and the interactions between them.

Within SSI systems, digital identity data is exchanged through two core structures: VCs and Verifiable Presentations (VPs) [9]. A VC is a digitally signed assertion issued by a trusted entity, containing claims about an individual or organization. These credentials may represent identity documents, diplomas, licenses, or attestations of attributes. Unlike traditional credentials, VCs can be cryptographically verified and selectively disclosed by the holder. When information needs to be shared with a third party, the holder generates a VP, which packages one or more VCs—or selected claims from them—and attaches a cryptographic proof demonstrating control over the disclosed data.

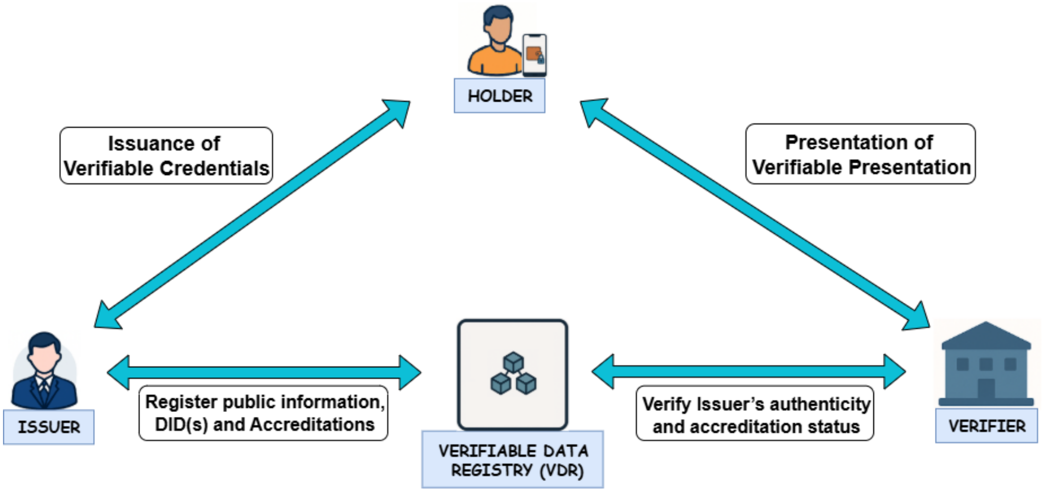
HOLDER
Issuance of Verifiable Credentials
Presentation of Verifiable Presentation
ISSUER
Register public information, DID(s) and Accreditations
VERIFIABLE DATA REGISTRY (VDR)
Verify Issuer's authenticity and accreditation status
VERIFIER

Figure 1. Architecture of the SSI system.

In the European context, EBSI represents an official implementation of SSI principles and standards. Developed by the European Commission and the European Blockchain Partnership, EBSI operates as a decentralized, cross-border network for digital services aligned with EU regulations and the core values of SSI. Its goal is to provide a trusted and interoperable digital identity framework for EU citizens, organizations, and institutions, supporting use cases such as diploma issuance, identity verification, and access to public services [10]. In the EBSI ecosystem, DIDs are alphanumeric strings that uniquely identify a subject without disclosing any personal information. A distinction is made between DIDs assigned to natural persons and those associated with legal entities. The guidelines and components defined by EBSI serve as a key technical reference for developing solutions aligned with SSI principles, a foundation upon which SIBERIA is built.

For natural persons, the decentralized identifier follows the syntax did:key:<method-specific-identifier>. The method-specific identifier (MSI) must be unique and case-sensitive. It is derived from the subject's public key, encrypted and encoded using Base58BTC, and always preceded by the letter "z". In the case of legal entities, the syntax used by EBSI is did:<network>:<method-specific-identifier>, where <network> is set to "ebsi". In this case, the MSI must also be unique and case-sensitive, similarly encoded in Base58BTC, and starting with the character "z". However, unlike in the case of natural persons, the MSI is not derived from an encrypted public key but consists of 16 random bytes.

To ensure trust, interoperability, and security within its SSI network, EBSI relies on several core registries that form the backbone of its decentralized identity system. Each registry plays a vital role in maintaining the integrity, functionality, and immutable management of decentralized identities, trusted issuers, and credential schemes within the ecosystem.

- DID Registry: This is the registry where DIDs are created, managed, and resolved. It ensures that entities within the EBSI ecosystem, such as individuals or organizations, can be uniquely identified in a decentralized manner. This registry is crucial for the resolution of DID documents, which contain the necessary information for authentication and verification.
- Trusted Issuers Registry (TIR): This registry lists the trusted issuers authorized to issue VCs within the EBSI network. By maintaining a registry of trusted issuers,

EBSI ensures that only authorized entities can participate in the issuance of VCs, thereby establishing trust and authenticity in the credentials that are shared across the ecosystem. Issuers in the TIR are categorized into three types:

  - Root Trusted Accreditation Organisation (Root TAO): A root entity capable of self-accrediting and accrediting other organisations.
  - Trusted Accreditation Organisation (TAO): Organisations that accredit third parties to issue VCs.
  - Trusted Issuer (TI): Issuers authorized to create and transmit credentials linked to a specific subject.

- Trusted Schema Registry (TSR): This registry is responsible for defining and maintaining the trusted schemas for the VCs issued within the system. These schemas establish standardized formats and structures for different types of credentials, ensuring consistency and compatibility across different issuers and verifiers.

A key component for user interaction within SSI systems is the Digital Wallet. This concept refers to software applications that enable holders to securely store, manage, and present their VCs. Serving as the primary interface for identity control, the wallet allows users to decide when, how, and with whom to share personal information. Depending on the architecture, digital wallets can be implemented as mobile apps, browser extensions, desktop clients, or cloud-based services. Many wallets also include backup and recovery options, multi-device synchronization, or integration with trusted hardware for enhanced security [11]. Aligned with these principles, SIBERIA includes its own digital wallet component, designed to deliver secure and autonomous identity management.

1.1.1. Related Work on SSI

SSI represents a paradigm shift in digital identity management, granting individuals full control over their personal data while eliminating the need for centralized authorities. This subsection reviews the most relevant surveys and studies on SSI, focusing on the core elements of SSI frameworks and the underlying technologies commonly employed. The selected works highlight the most representative and influential contributions in the field, providing an overview of the current state of technological maturity.

The IdM system for public transportation was introduced in [12] as an SSI-based identity manager leveraging blockchain technology for a European public transport system. This system allowed students to access transport discounts using VCs issued and signed by their respective universities, which certify their student status. The proposed solution ensured secure and decentralized verification of student identities, enhancing privacy and autonomy. However, the design remained theoretical, with no practical implementation. The authors of [13] explore the application of SSI in IoT networks, comparing its implementation with Pretty Good Privacy (PGP) and X.509 standards. Their analysis examined core components and functionalities across these technologies. The study concluded that SSI was particularly advantageous in IoT environments, as it granted users full control over their identities, eliminated dependence on third parties, and enhanced privacy. The SSIBAC system was presented in [14] as an access control system based on SSI. The tool combined conventional access control elements with blockchain technology (using Sovrin) to provide decentralized authentication and centralized authorization. The system was validated in an academic context, where VCs represent academic information such as degrees or qualifications. In this way, students are the holders, educational institutions are the issuers, and verifiers are the potential employers. Finally, the authors of [15] proposed a theoretical SSI-based identity model for healthcare, a sector characterized by highly sensitive data. The study highlighted the potential of SSI to enhance patient confidentiality and ensure compliance with existing regulations.

Recognizing the potential of such identity models and anticipating emerging trends and regulations related to identity and rights protection, several commercial solutions and products have been developed in recent years. These tools have demonstrated, to varying degrees, successful use cases and real-world applications, and therefore, they warrant discussion and analysis. Table 1 provides a comparison of their general characteristics. The tools are outlined below:

- SelfKey [16] is a blockchain-based SSI management system based on Ethereum blockchain technology to ensure secure, private, and efficient identity verification processes. The SelfKey Marketplace offers access to various financial, immigration, and cryptocurrency services. Users can find and apply for products like bank accounts, residency, and company incorporation services directly through the SelfKey platform. The SelfKey ecosystem is powered by the KEY token (an ERC-20 token on the Ethereum blockchain), which is used for transactions within the network. KEY tokens facilitate identity verification services and access to the marketplace, and they incentivize participation in the SelfKey ecosystem.
- Sovrin [17] is a decentralized, global public utility designed specifically for managing digital identities. It operates as an open-source, blockchain-based platform that enables individuals and organizations to create, manage, and verify digital identities. At the heart of the Sovrin Network is distributed ledger technology, which provides a tamper-proof record of identity transactions. This ledger is maintained by a network of independent stewards, which are organizations that operate the nodes and uphold the integrity of the system. These stewards include reputable institutions from various sectors. Sovrin's architecture supports a wide range of use cases, from verifying academic credentials and professional qualifications to enabling secure access to services and compliance with regulatory requirements. Organizations can issue VCs that users store in their digital wallets and present when needed.
- LifeID [18] is an SSI platform with zero-knowledge proof technology to enable users to only present the needed information without revealing other attributes. In addition, LifeID provides an app with biometric authentication to protect against identity theft. Finally, users can recover their identity through backup, with close family/friends, and with a trusted organization.
- Evernym [19] is a company that provides SSI technology, developing the Sovrin Network and empowering individuals and organizations to manage digital identities securely and privately. One of Evernym's flagship products is Verity, a comprehensive platform for building and deploying SSI solutions. Verity provides tools for creating and managing digital credentials and facilitating secure interactions among issuers, holders, and verifiers.
- Hyperledger Indy [20] is a distributed ledger purpose-built for decentralized identity management, providing the tools and libraries necessary for creating and using independent SSI. Developed under the Hyperledger umbrella, which is hosted by the Linux Foundation, Indy aims to enable identity owners to control their own identity and verifiable claims. Indy enables zero-knowledge proofs and selective disclosure, which means users can prove certain attributes about themselves without revealing their full identity or other personal information. Indy also supports interoperability and adherence to global identity standards. It is designed to work seamlessly with other decentralized identity solutions and technologies, promoting a cohesive and unified approach to digital identity management. The platform is compatible with the W3C standards for DIDs and VCs.
- EverID [21] aims to solve issues related to identity verification with blockchain technology. The platform supports biometric data such as fingerprints, facial recognition,

and iris scans. EverID is a user-friendly identity account system for managing crypto assets on the everPay network. It simplifies the process for traditional industry users to create and handle non-custodial accounts, thus easing their transition into the Web3 environment. EverID provides two main types of account management. The first is through wallet addresses, which aligns with the standard method employed by most blockchain projects. The second type uses the FIDO authentication system, enabling account creation via email addresses, biometrics, and other secure methods. This approach significantly reduces the barrier for users, making it easier and more accessible to manage Web3 crypto assets.

Table 1. Comparison between SSI-based commercial products.

	Blockchain Based	Wallet	Web3 Standard	Open-Source	DID Based
Selfkey	✓	✓	✓	✗	✗
Sovrin	✓	✗	✗	✓	✗
LifeID	✓	✓	✓	✓	✓
Evernym	✓	✓	✗	✓	✓
Hyperledge Indy	✓	✓	✓	✓	✓
EverID	✓	✗	✓	✗	✓

✗ No. ✓ Yes.

The described commercial solutions reflect the increasing maturity and adoption of SSI technologies across sectors such as public transportation, IoT, and healthcare. However, most remain either conceptual or limited to narrow domains, with little attention paid to the stringent requirements of industrial environments. Specifically, aspects such as CA, biometric privacy, or real-time credential verification in constrained Operational Technology (OT) networks are poorly addressed. This gap highlights the novelty of SIBERIA, which applies SSI principles to high-security industrial contexts and complements them with adaptive, privacy-preserving multi-factor authentication.

1.1.2. IAM Frameworks and Industrial Standards in OT Environments

While the adoption of SSI technologies continues to evolve, industrial environments, especially those governed by OT, are subject to well-established standards and frameworks for identity and access management (IAM). These frameworks are essential for ensuring secure authentication, authorization, and accountability within critical infrastructure systems.

A primary reference in the industrial cybersecurity domain is the IEC 62443 series [22], which outlines comprehensive security requirements for control system components. It emphasizes role-based access control, secure credential management, and multi-factor authentication tailored to OT architectures. Similarly, the NIST SP 800-82 Revision 3 [23] offers detailed guidance on applying the NIST Cybersecurity Framework to ICS networks, reinforcing the relevance of layered IAM mechanisms in sensitive industrial contexts.

In the corporate IAM landscape, platforms such as Azure Active Directory increasingly integrate with OT networks via gateway interfaces or edge connectors [24]. While these solutions enable centralized policy enforcement, they often require tight infrastructure alignment and may conflict with the decentralization principles championed by SSI [25].

Additionally, OAuth 2.0 and OpenID Connect (OIDC) [26] are widely adopted in industrial IoT (IIoT) platforms and cloud-native environments. These protocols provide standardized mechanisms for token-based delegation and federated identity. However, they are inherently centralized and lack native support for VCs or DIDs. Limitation underscores the distinct positioning of SIBERIA, which combines privacy-preserving identity models with industrial-grade authentication flows.

Recent studies [27,28] have explored blockchain-enabled IAM in OT environments, addressing challenges related to access decentralization, resilience, and traceability. While conceptually aligned with SIBERIA, these works do not incorporate advanced biometric authentication or continuous behavioral monitoring. These limitations open the door to novel models that can provide stronger guarantees of autonomy, adaptability, and privacy in OT environments. SIBERIA proposes a hybrid approach that combines compliance with industrial standards and regulatory frameworks (e.g., GDPR and EBSI) with innovative features such as local biometric processing, CA, and decentralized credential management. This strategy bridges existing gaps and aligns with the evolving security demands of OT ecosystems.

1.2. Secure Voice Biometric Authentication

One of the most critical aspects in the design of SSI systems is ensuring that VCs are presented exclusively by their rightful holder. Biometric authentication serves as a strong complement to traditional username–password mechanisms, enhancing protection against unauthorized access. Various human traits, such as the retina, face, and fingerprint, can be used to verify the identity of a VC holder. Among these, voice offers several advantages for biometric applications. It can be captured remotely, requires no physical contact, and does not necessitate specialized hardware, as microphones are already embedded in most consumer devices [29]. Automatic Speaker Verification (ASV) systems have gained popularity in recent years, largely due to advances in deep learning [30]. Modern ASV systems typically rely on large-scale Deep Neural Networks (DNNs) trained to extract speaker embeddings that represent an individual's vocal identity [31]. The most well-known DNN architectures used for this purpose include X-vectors [32], ECAPA-TDNN [33], and, more recently, TiTANet [34].

A typical ASV system comprises two main stages: (1) enrollment, during which a set of utterances from the holder is used to extract a voiceprint (referred as a speaker embedding) using one of the previously mentioned DNN models; and (2) authentication, where a new speech sample from the speaker is processed to extract a second voiceprint. This speaker representation is then compared to the one obtained during enrollment in order to retrieve a similarity score used for final authentication when it surpasses a predefined threshold. Common techniques for comparing speaker voiceprints include simple metrics such as cosine similarity, as well as more sophisticated and adaptive methods like Probabilistic Linear Discriminant Analysis (PLDA) [35].

Despite the multiple benefits of ASV systems, their performance is known to degrade significantly in the presence of fabricated or manipulated speech inputs [36]. These so-called spoofing attacks can be categorized as either logical or physical. Logical attacks involve the use of voice cloning or text-to-speech (TTS) systems to synthetically generate a speaker's voice, whereas physical attacks rely on replaying pre-recorded speech samples of the legitimate user during the authentication phase [36]. Among these, logical attacks currently pose the greatest challenge due to recent advancements in generative AI technologies and modern speech synthesis frameworks [37].

With the aim of improving the robustness of ASV systems against spoofing attacks, it is essential to incorporate anti-spoofing mechanisms capable of detecting machine-generated or replayed speech samples [38]. Given the importance of creating these mechanisms, the research community has actively developed and benchmarked a wide range of models and detection engines [39]. The most common approach to integrate an anti-spoofing mechanism in an ASV system is to design separate modules that can be combined in a cascade or parallel fashion. This integration can be performed at the score, decision, or embedding levels. This fusion of an ASV with an anti-spoofing mechanism is what is known in the research community as SASV systems [40].

When implementing these biometric systems, it is crucial to ensure the security of all data, including the voiceprints derived from the VC. According to the European Union's data privacy regulations [41], biometric data are classified as personal and highly sensitive information, thereby warranting strong privacy protections. In line with this, the ISO/IEC IS 24745 standard [42] specifies requirements for the protection of biometric information, addressing confidentiality, integrity, and renewability/revocability during both storage and transmission. To meet these requirements, voiceprint protection strategies typically focus on three main principles: (1) unlinkability, ensuring that voiceprints cannot be correlated across different applications or databases; (2) irreversibility, guaranteeing that biometric samples or other personal traits cannot be reconstructed from the stored voiceprints; and (3) renewability, which mandates that multiple voiceprints generated for the same user are treated as independent and non-interchangeable [43].

To satisfy these principles, voiceprints must be stored and processed using robust encryption mechanisms. Although these voiceprints are usually stored as feature embeddings generated by neural networks, recent studies have shown that sensitive information such as gender, age, or even health-related biomarkers can still be inferred from these embeddings [43,44]. Traditionally, voiceprints are encrypted only during storage (i.e., in the enrollment phase) but are decrypted during the authentication stage [45]. This decryption exposes the raw embeddings to the server and, potentially, to adversaries who may compromise the server, thereby creating privacy risks and enabling user tracking. To address this vulnerability, the processing of voiceprints should be carried out without decryption. This can be achieved using homomorphic encryption (HE), which allows mathematical operations to be performed directly on encrypted data. In this way, both verification and anti-spoofing scores can be computed in the encrypted domain, preserving the privacy of the holder's biometric data throughout the entire authentication process.

SIBERIA implements a secure state-of-the-art SASV system for holder biometric authentication, composed of two independent modules: (1) a combined, cascaded ASV and anti-spoofing component that generates encrypted voiceprints and validates the user's identity, producing independent encrypted verification and anti-spoofing scores; (2) a key management component responsible for generating the cryptographic keys used to encrypt the holder's voiceprint and to decrypt the ASV and anti-spoofing scores. This module subsequently produces a final combined score that determines the holder's identity when accessing a specific service. Additional implementation details of these modules are provided in Section 2.2.

1.3. Behavioral Biometrics

CA has emerged as a promising approach to enhancing digital security by verifying user identity throughout a session, rather than relying on a single login event. This method enables the detection of anomalous behavior and helps prevent unauthorized access or account takeover during active sessions.

A core component of CA is behavioral biometrics, which models distinctive user habits such as typing patterns, touch dynamics, or device usage that are difficult to mimic or steal. Unlike physiological biometrics, behavioral data can be collected passively and unobtrusively through everyday interactions with smartphones, computers, or other devices [46,47].

Behavioral biometrics rely on embedded sensors (e.g., accelerometers, gyroscopes, touchscreens, and microphones) to capture temporal signals that reflect unique user patterns. These signals are processed and transformed into features used to train supervised machine learning models that distinguish between legitimate users and impostors. In industrial settings aligned with Industry 4.0, high accuracy has been achieved by analyzing keyboard and mouse usage using models like Random Forest and Gradient Boosting [48]. Similarly, smartphones and smartwatches use motion data with algorithms such as Support Vector Machine (SVM), K-Nearest Neighbors (KNN), or Convolutional Neural Network (CNN) to enable efficient and low-latency CA [47].

Although a wide range of techniques is employed in CA systems, balancing security, usability, and privacy remains a challenge [46]. Ensuring robustness against attacks and adaptability to behavioral changes often requires online learning or periodic model retraining. Additionally, the incorporation of contextual information (e.g., time of day and device location) is being explored to improve classification reliability [49]. Ultimately, the synergy between behavioral biometrics and machine learning enhances the precision of identity verification while enabling a seamless user experience, making CA a compelling solution for securing modern digital ecosystems.

The effectiveness of CA systems based on behavioral biometrics depends heavily on the quality, richness, and representativeness of the data used to train and evaluate the models. Among the most relevant behavioral signals for this purpose are touch dynamics, keystroke patterns, inertial sensors (e.g., accelerometers, gyroscopes, and magnetometers), and contextual usage patterns. These signals enable detailed modeling of how users physically interact with their devices. The Hand Movement, Orientation, and Grasp (HMOG) dataset is one of the main references in the field of behavioral biometrics. It comprises a set of features obtained from hand micro-movements recorded by smartphone motion sensors while users perform everyday activities such as reading, writing, or walking [50]. This set includes data from more than 100 participants in different physical conditions, both at rest and in motion, making it a key tool for evaluating the robustness of CA systems. Accelerometer, gyroscope, and magnetometer sensors record high-frequency synchronized data, enabling the extraction of fine-grained temporal characteristics.

Using datasets such as HMOG, recent studies have explored deep learning techniques to automatically extract representations from raw behavioral signals. A study highlighted in [51] evaluated models such as CNNs and Long Short-Term Memory (LSTM) networks, demonstrating their effectiveness in capturing both temporal dependencies and spatial correlations in data obtained from smartphone sensors. These architectures were tested with data gathered during real-world usage sessions and showed strong generalization across different users and scenarios. The study also emphasized the importance of using signals with high temporal resolution and sufficient duration to train robust models and prevent overfitting.

The adoption of deep architectures such as CNNs and LSTMs has advanced the field of CA based on behavioral biometrics. These models are particularly well-suited to handle the complex, high-dimensional, and temporally correlated data obtained from motion sensors and user interaction patterns on mobile devices. The HMOG corpus, presented in [50], serves as an ideal benchmark for evaluating these models. The dataset enables the training of models capable of distinguishing users based on precise motion signatures across various physical conditions.

Hybrid architectures combine the spatial feature extraction capabilities of CNNs with the temporal modeling power of LSTMs. For instance, in [52], a CNN-BiLSTM architecture was introduced for processing mobile sensor data in the context of CA. In this approach, convolutional layers identify local patterns and structures in raw signals, while bidirectional LSTM layers capture temporal relationships in both the forward and backward directions, thereby improving authentication accuracy. Building upon this work, the model proposed in [53] incorporated an attention mechanism on top of the CNN-BiLSTM architecture, allowing it to focus on the most relevant segments of the behavioral sequence and increasing robustness to noisy or uninformative data. On the other hand, the DeepConvLSTM [54] sequentially combined CNN and LSTM layers to jointly learn spatial and temporal dynamics from activity data and has demonstrated outstanding performance in mobile user authentication scenarios.

SIBERIA proposes a behavioral biometric authentication system based on the continuous monitoring of user interactions with the keyboard and mouse, processed through a locally deployed deep learning model. Unlike previous studies, which typically focus on static datasets and centralized architectures, SIBERIA emphasizes real-time, on-device evaluation, preserving user privacy and enabling seamless integration into everyday use.

1.4. Practical Relevance and Application Scenarios

SIBERIA offers a secure and privacy-oriented identity management solution applicable across a wide range of sensitive sectors. Below are several representative scenarios where its implementation would be particularly beneficial:

- Critical Infrastructure Management: In sectors such as energy, water treatment, and transportation, SIBERIA could restrict access to control rooms strictly to authorized personnel through the use of decentralized credentials and biometric verification (voice or behavioral), ensuring operational traceability and system integrity.
- Healthcare: SIBERIA would enable secure access to medical records and hospital systems without directly storing user credentials. Furthermore, CA via behavioral biometrics would ensure that the medical professional accessing a patient's data remains verified throughout the session, preventing unauthorized use of shared terminals. Patients could also grant temporary, verifiable access to other specialists while maintaining full control over their data, in compliance with the GDPR.
- Financial and Corporate Services: In corporate environments, SIBERIA could safeguard access to sensitive financial data or intellectual property. Additionally, in high-value or high-risk transactions, initial voice-based authentication combined with continuous behavioral monitoring would ensure that the individual performing the operation remains consistently verified, significantly reducing the risk of internal or external fraud.
- Public Services: SIBERIA could provide citizens with a self-managed digital identity to securely access public services such as transportation, social programs, or electronic voting. Using VCs, users could disclose only the attributes strictly necessary, in accordance with data minimization and SSI principles.

2. Materials and Methods

This section outlines the technical foundations, software components, and implementation strategies employed in the development of SIBERIA. The system comprises three independent but interoperable modules: an SSI module, a secure SASV system, and a behavioral biometric monitoring module. Each module was implemented using a combination of custom-developed and established technologies, aiming to deliver a secure and scalable identity management solution for industrial environments. Novel methods and protocols are described in detail to ensure replicability, while existing standards and technologies are referenced appropriately. Software versions and the availability of relevant codebases are specified where applicable.

2.1. SSI Module

SIBERIA adopts an SSI model tailored for private industrial environments, where decentralized and secure identity management is essential. The architecture implements a modular and decentralized identity management system, offering individuals complete control over their identity data while ensuring robust authentication and integrity mechanisms.

At the core of the identity infrastructure is the VDR, implemented on a private EOSIO-based blockchain, which manages digital identities and VCs within the system. The choice of EOSIO is based on its implementation of a Delegated Proof-of-Stake (DPoS) consensus mechanism, which enables efficient transaction validation without relying on the costly and slow mining processes typical of other public blockchains.

Unlike the public EOSIO blockchain, where the 21 most-voted nodes among all participants act as producer nodes, this private EOSIO blockchain consists of only three nodes, which elect each other to become active producer nodes. In this network, a single administrative account on the EOSIO chain manages the deployment of the smart contract responsible for defining and maintaining three essential SSI registries inspired by the EBSI ecosystem: DID Registry, TIR, and TSR. These registries, implemented on the blockchain, enable the secure and immutable management of decentralized identities, trusted issuers, and credential schemas.

Regarding DIDs, SIBERIA closely aligns with the identity model defined by EBSI. For natural persons, SIBERIA adopts the did:key method, following the structure established by EBSI, where identifiers are derived from a public key and encoded using the Base58BTC format. However, within the SIBERIA context, this type of DID is reserved exclusively for identity holders. Conversely, DIDs associated with legal entities are designated for issuers, who utilize a network-specific DID structure in the format did:siberia:<method-specific-identifier>. This approach ensures compatibility with existing DID resolution mechanisms while enabling the contextualization and differentiation of the service within the SIBERIA ecosystem.

In SIBERIA, the DID Registry is specifically designed to record only the DIDs and DID documents associated with legal entities, namely issuers. Consistent with EBSI recommendations, registering a DID document requires the issuer to first obtain a verifiable authorization to onboard, ensuring a trusted entry into the system.

Regarding the TIR, SIBERIA adopts a streamlined approach by reducing the issuer categories from three to two: TAO and TI. The TAO designation is exclusively reserved for SIBERIA’s issuer, while all other entities are classified as TI. To be listed as a trusted issuer in the VDR, issuers must present a valid trusted issuer credential. Furthermore, to enhance operational efficiency and privacy, issuers have the option to register a proxy service within the TSR. This proxy facilitates verifiers in remotely querying the status of credentials issued, such as revocation status, without necessitating direct communication with the issuer.

In order to register a schema within the TSR, an issuer must first obtain an authorization token granting permission to perform the registration. This serves as a form of access control, ensuring that only authorized issuers can contribute schemas to the registry. Once the issuer holds this token, they may proceed to register one or multiple schemas.

Each registered schema represents a formal definition of the structure, attributes, and data formats that VCs issued under that schema must comply with. This mechanism facilitates validation processes and promotes a standardized framework for credential issuance within SIBERIA.

To support interaction with the SIBERIA SSI ecosystem, three specialized API interfaces have been developed, each targeting distinct roles within the infrastructure:

- TAO API: This API is managed by the authorized TAO within the SIBERIA environment and exposes the necessary endpoints for managing onboarding requests to the ecosystem. It allows external entities to initiate the credential request process required to become recognized issuers.

- Token Issuance API: This API handles the secure issuance of authorization tokens required to perform sensitive actions within the system, such as schema registration or issuer onboarding. Tokens issued through this API serve as access control mechanisms, ensuring that only authenticated and authorized entities can perform critical operations.
- SIBERIA API: This API enables direct interaction with the private blockchain and the VDRs. It supports a wide range of functions, including issuer registration, schema registration, querying trusted issuers and schemas, and credential sharing, among others.

These API layers are fundamental for ensuring modularity, security, and scalability within the SIBERIA architecture. By separating roles and scopes, the system enforces principle-based access control, ensuring that each actor interacts only with components relevant to their responsibilities.

The identity wallet plays a pivotal role in managing digital identities for natural persons, serving as the primary interface for holders within the SSI framework. Designed exclusively for Android devices and strictly scoped to the private SIBERIA ecosystem, the wallet adheres to the core principles of SSI: user control, privacy, security, and portability.

After completing the registration process with a username and password, the holder is automatically provisioned with a did:key identifier, a corresponding cryptographic key pair, and an associated DID document. In line with SSI principles, all identity-related data—keys, identifiers, credentials, and metadata—are generated and stored locally on the user's device, ensuring full user control without any form of external synchronization or centralized storage. Additionally, the wallet integrates biometric authentication, allowing users to link fingerprint-based access for enhanced security and convenience.

Regarding credential management, the wallet provides functionality to receive VCs and to store and manage them locally on the device. The wallet also supports the generation of VPs, allowing the holder to selectively disclose one or more credentials. Each VP is constructed in compliance with the appropriate schema registered in the blockchain-based TSR, ensuring both structural and semantic validity. Once generated, the presentation is signed with the holder's private key, providing cryptographic proof of authenticity and integrity before it is securely shared with the verifier.

To support device migration, the wallet allows exporting all user identity data (DID, keys, and credentials) in a secure JavaScript Object Notation (JSON) format. These data can then be imported into a new device, ensuring full user data portability.

In addition to the components described above, the SIBERIA SSI architecture incorporates two databases, each serving distinct critical functions:

- A database managed by the verifier that stores users' email addresses and passwords. This enables secure authentication and efficient account management.
- A database managed by the issuer that maintains the status of credentials issued to holders, indicating whether a credential is active, suspended, or revoked. Importantly, this database stores only credential status, not the credential data itself.

To facilitate efficient and privacy-preserving status verification, the issuer registers a proxy service within the TSR. This proxy acts as an intermediary, enabling verifiers to remotely query credential status stored in the issuer's database without communicating directly with the issuer. This design optimizes operational efficiency and enhances privacy by decoupling credential status checks from direct issuer interactions.

This module—including the API and databases—integrates seamlessly with the overall SIBERIA SSI framework, complementing the blockchain-based verifiable data registries and the identity wallet to deliver a comprehensive, secure, and scalable SSI system.

Figure 2 provides an architectural overview of the SIBERIA SSI module, illustrating the interaction between its main components.

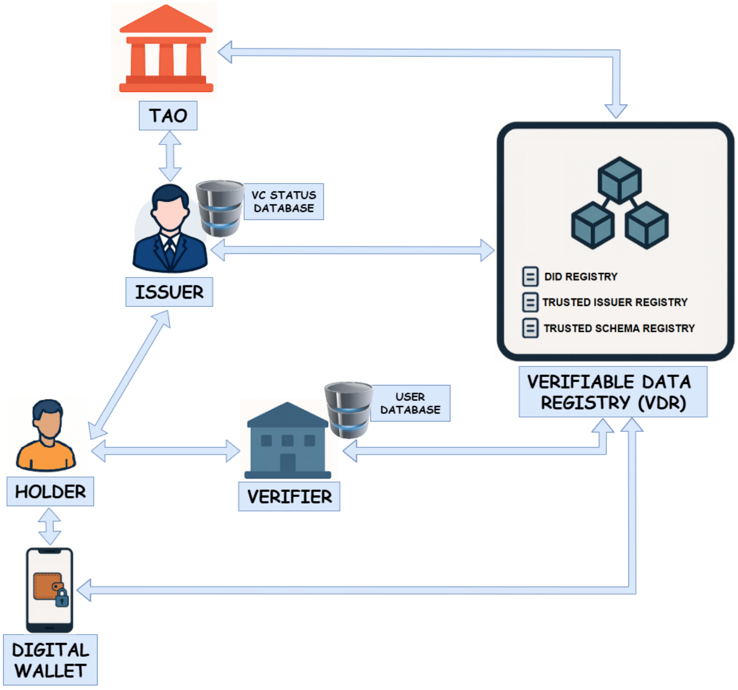

Figure 2. High-level architecture of the SIBERIA SSI module.

2.2. *Secure SASV Module*

The overall architecture of the secure SASV module implemented in SIBERIA is shown in Figure 3. It comprises two independent components interacting with the SSI module: (1) the Keys Manager, responsible for the generation and management of keys used in the HE of voiceprints, and (2) the biometric processor, which handles voiceprint generation and the identity verification of the holder. These two components are deployed on separate servers to enhance security and privacy, ensuring that the private keys required for decryption are stored independently from the encrypted voiceprints. The process of biometric authentication, considering the SIBERIA SASV module, consists of five steps, detailed as follows:

0. The issuer in the SSI module sends a request to the Keys Manager to generate a set of cryptographic keys for secure SASV processing.
1. The Keys Manager returns the set of public and cryptographic keys to the holder. The private key is securely retained by the Keys Manager and is later used to decrypt the biometric scores.
2. During enrollment, the issuer submits an audio sample from the holder along with the public key to the biometric processor, which generates the encrypted voiceprint that is stored in a voice VC. During authentication, the verifier provides a verification audio sample from the holder, the encrypted enrollment voiceprint stored within the VC, and the set of public, Galois, and relin keys. These are used to perform identity verification entirely within the HE domain.
3. The biometric processor returns either the encrypted voiceprint to the issuer during the enrollment phase or a pair of encrypted biometric and anti-spoofing scores during authentication of the holder’s identity.
4. The pair of biometric and anti-spoofing scores is sent by the SSI verifier to the Keys Manager, which decrypts them and performs a score-level fusion to produce the final SASV authentication score.

5. Finally, the SASV score is sent back to the SSI verifier, where it is evaluated against a predefined threshold to determine whether the identity should be accepted or rejected.

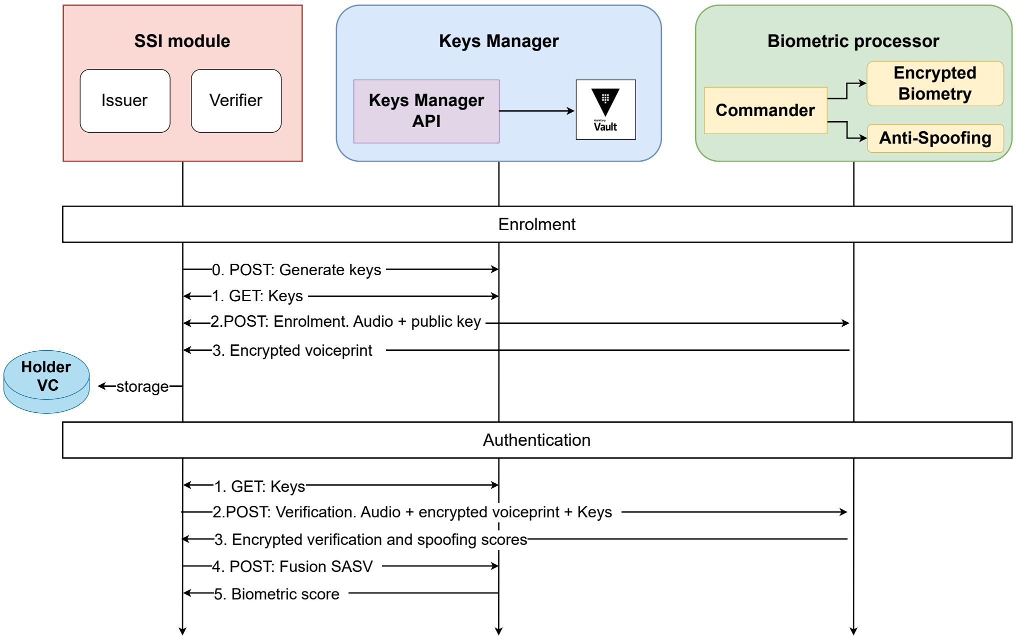

Figure 3. Overall architecture and flow chart of the secure SASV module implemented in SIBERIA. VC: voice credential.

To preserve the holder's privacy and adhere to SSI principles, no identifiable information is stored on either the Keys Manager or the biometric processing servers. The encrypted voiceprint is stored locally on the holder's digital wallet, ensuring that biometric data remains under the holder's control. The Keys Manager retains a copy of the cryptographic keys, which are required for encrypting voiceprints and decrypting the biometric and anti-spoofing scores. The public keys are made accessible to the holder through authenticated API calls to the Keys Manager. The biometric module is stateless with respect to the holder. It does not retain any personal or biometric information. All data required for enrollment or authentication are transmitted through secure API requests initiated by the SSI module.

The following subsections describe each of the individual components included in both the Keys Manager and the biometric authentication module.

2.2.1. Keys Manager

The Keys Manager is responsible for generating and managing the set of keys required by the SASV module. The public key is used to encrypt the voiceprint generated during enrollment, while the private key is employed to decrypt the biometric and anti-spoofing scores computed during the authentication phase. In addition to the public and private keys, the Keys Manager also generates Galois and relinearization (relin) keys to support HE operations. These two auxiliary keys enable specific encrypted computations that cannot be performed using standard keys alone. In particular, the Galois key allows for encrypted vector rotations and other Galois automorphisms on ciphertexts, which are essential for implementing matrix–vector multiplications during the authentication phase. Conversely, the relin key is used to reduce the size and computational complexity of ciphertexts after homomorphic multiplications, thereby improving efficiency and maintaining manageable noise levels [55].

The set of cryptographic keys is generated using Microsoft SEAL [56], a lattice-based HE library that supports secure computation by enabling addition and multiplication operations directly on encrypted integers or real numbers. To ensure secure and auditable key storage, the Keys Manager integrates an instance of Vault [57], which provides robust, policy-based access control and secure storage for the generated encryption keys.

Beyond key generation and storage, the Keys Manager is also responsible for computing the final SASV score (SS) returned to the SSI module. It receives a pair of encrypted biometric and anti-spoofing scores, which are individually decrypted using the private key. These decrypted scores are then combined using Equation (2), where b_s and a_s represent the decrypted biometric and anti-spoofing probability scores, respectively, while σ denotes the sigmoid function used to normalize the final score [58]. α_a and α_b are learned weights used for the fusion.

$$\mathrm{SS} = \sigma\left(-\log\left|\alpha_a e^{-a_s} + \alpha_b e^{-b_s}\right|\right) \tag{2}$$

2.2.2. Commander

The Commander serves as the orchestrator for the biometric and anti-spoofing pipelines, handling user HTTP requests and managing data flow between the different deployed modules. Supplementary data needed to complete each task are provided by any of the instances of the SSI module in JSON format, which may include audio files and the public keys needed by the technological modules. Upon receiving a request, tasks are stored in a processing queue, and a unique string identifier is returned to the SSI components. The steps of the task are executed sequentially when resources are available, and the result is stored in memory for retrieval via a subsequent HTTP request, also in JSON format.

To improve performance, this component supports the deployment of multiple instances of each technological module and distributes tasks among them, reducing processing times when sufficient computational resources are available. Instance management can be triggered via a single HTTP request, without requiring a service restart.

In the specific case of SIBERIA, two different pipelines have been developed: enrollment and verification. The first one only parses an input request and sends its content to the biometry component. In the second pipeline, the commander parses the input request from the SSI module and sends two different requests to each technological module. Once the responses are received, both outputs are merged and saved as the result of the verification task.

2.2.3. Encrypted Biometry

This component is responsible for generating the biometric voiceprint of the wallet holder during the enrollment phase and for producing the encrypted biometric score during the authentication stage. The voiceprint is generated as an embedding vector extracted from a large DNN model. Specifically, the model consists of a 101-layer Residual Network (ResNet) [59] trained on audio samples at 16 kHz from various media domains, including data from Voxceleb1 (323 h of speech from 1211 speakers) [60], Voxceleb2 (2290 h from 5994 speakers) [61], and CN-CELEB (264 h from 973 speakers) [62]. The pre-trained DNN model is publicly available at [63]. During enrollment, the resulting voiceprint is encrypted using the public key and stored in the VC of the wallet holder for secure storage, in alignment with the SSI paradigm.

In the authentication phase, the module receives the input audio used for the holder's verification, the voiceprint securely stored in the VC, and the set of public, Galois, and relin keys. A new verification voiceprint is computed from the input audio using the same process as in the enrollment phase. To verify the holder's identity, the similarity between the encrypted enrollment and verification voiceprints is calculated using cosine distance. This operation is performed entirely in the HE domain, leveraging the Galois and relin keys. The result is an encrypted biometric score, which is returned to the SSI module for further processing.

2.2.4. Anti-Spoofing

The core objective of this module is to detect fraudulent identity verification attempts resulting from logical spoofing attacks, such as voice cloning or other audio deepfake techniques. The anti-spoofing module also relies on a DNN model trained to determine whether the input audio used during the authentication phase was produced by a genuine speaker or synthesized using voice generation technologies. The DNN model considered for anti-spoofing is based on a Self-Supervised Learning (SSL) upstream model [64], focused on computing an embedding representation of the input audio, and a downstream classifier, trained for the anti-spoofing task. Specific details about the neural architecture can be found in [65]. The trained model produces a score representing the similarity between the input audio and a reference vector representing a set of genuine samples. The computed score is finally encrypted using the public key and returned to the SSI module.

The anti-spoofing model was trained using the combination of ASVSpoof2019 [66] (111 h of speech from 107 speakers), ASVSpoof2021 [67] (133 h of speech from 107 speakers), and a subset of MLAAD corpus [68] (24 h of selected multilingual spoofed samples). This MLAAD subset was selected to include multilingual information in four languages: Swedish, Greek, Spanish, and French, in order to adapt the model to those specific languages. The training corpus is completed with an additional 24 h of synthetic speech generated using the VITS TTS model [69]. The training process employed the one-class softmax loss function [70], a dropout rate of 0.2 where applicable, and the ADAM optimizer with a learning rate of 3×10^{-4}.

2.2.5. SASV Validation

The technological components of the secure voice-based biometric module are evaluated using the ASVspoof 2024 Challenge corpus (ASVspoof5) [71]. This dataset comprises speech data from over 4000 speakers, containing more than 100,000 genuine utterances and over 500,000 spoofed samples generated using 16 different TTS and voice cloning algorithms. In particular, we consider the dev partition, which includes 140,950 audio samples, covering 8 types of spoofing attacks and more than 700 speakers. Each component of the SASV module is assessed individually using the Equal Error Rate (EER) metric, which corresponds to the operating point where the False Acceptance Rate (FAR) equals the False Rejection Rate (FRR). The evaluation is also conducted in terms of the architecture-agnostic detection cost function (a-DCF), a recently introduced metric for the assessment of SASV systems [72]. The a-DCF reflects the cost of decisions in a Bayes risk sense, with explicitly defined class priors and a detection cost model.

Evaluation results are presented in Table 2, reporting EER values for both the ASV and anti-spoofing modules, as well as the combined SASV score. The table also includes the corresponding acceptance probability thresholds used to classify the holder as legitimate at the EER point, in addition to the a-DCF cost. Our results are also compared with those obtained in previous studies.

Table 2. Error of the SASV module in the ASVSpoof5 corpus in terms of the EER for the ASV, anti-spoofing, and overall SASV scores, in addition to the a-DCF cost function.

Model	EER ASV (%)	EER Anti-Spoofing (%)	EER SASV (%)	EER Threshold	a-DCF
SIBERIA-SASV	1.30	0.29	1.17	0.21	0.021
[73]	-	10.8	6.53	-	0.134
[74]	-	-	-	-	0.125
[75]	5.10	2.97	2.51	-	0.048
[76]	-	-	1.28	-	0.028
[77]	1.28	0.74	0.85	-	0.021
[78]	-	0.76	0.58	-	0.008

The obtained results are competitive with those reported in the current state of the art. The considered approach achieved the third-best results in terms of a-DCF, after [78]. In addition, note especially that our system is the best model to detect spoofed samples (EER = 0.29%), compared with the rest of the studies. The EER for the overall SASV module in our system is 1.17%, achieved at a decision threshold of 0.21, which defines the acceptance criterion for identifying a legitimate holder. However, this threshold may not represent the optimal operating point for SIBERIA, particularly in applications where minimizing the FAR is crucial for enhancing system security. To maintain both robust authentication and usability, it is essential to select a threshold that strikes an appropriate balance between FAR and FRR. Consequently, the decision threshold should be carefully tuned based on the deployment context and security requirements.

To illustrate how the decision threshold can be selected, Figure 4 presents the log-likelihood ratio scores produced by the SASV module over the evaluation dataset. These scores were computed using Equation 2, omitting the sigmoid function to preserve the raw scale of the log-likelihood ratios. This representation facilitates the analysis of score distributions for both genuine, fraudulent, and spoofed trials, enabling the selection of an optimal threshold that best separates the three classes.

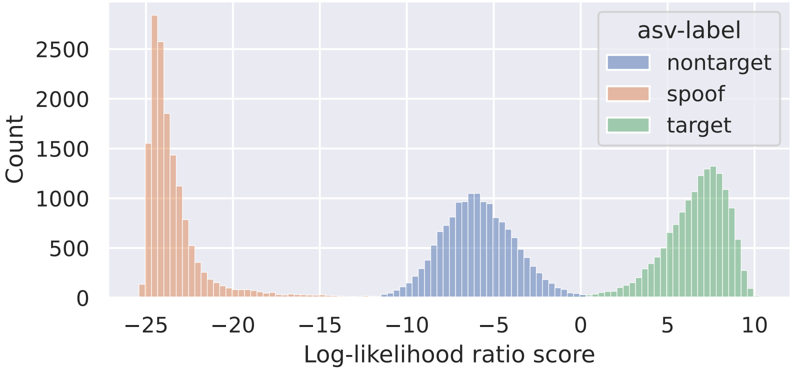

Figure 4. SASV score distribution for genuine (target), fraudulent (nontarget), and spoof samples.

When assessing a biometric solution such as the one proposed in SIBERIA, it is crucial to evaluate not only its accuracy but also the run-time performance of its different processing stages. Table 3 presents the execution time for each component involved in both enrollment and authentication workflows. The Keys Manager operates on a CPU server equipped with a 24-core Intel (R) Xeon (R) E5-2620 processor (Intel Corporation, Santa Clara, CA, USA) and 64 GB of RAM. The Encrypted Biometry module runs on a separate machine featuring a 16-core Intel (R) Xeon (R) Gold 6130 CPU, 346 GB of RAM, and an NVIDIA (NVIDIA Corporation, Santa Clara, CA, USA) GeForce GTX 1080 Ti GPU with 12 GB of VRAM. Overall, the processing times are consistent and remain low across all stages. The most time-consuming step is the SASV score fusion, performed on encrypted data within the Keys Manager module, due to the computational overhead introduced by the decryption operations of the scores.

Table 3. Run time performance of the different stages of the SASV Biometric module in SIBERIA. Time (in seconds) is presented in terms of the mean ± standard deviation of a set of 10,873 enrollment and authentication requests.

Process	Module	Time (s)
	Processing Times	
GET keys	Keys manager	1.00 ± 0.25
Enrollment	Encrypted Biometry	1.06 ± 0.19
Authentication	Encrypted Biometry	1.76 ± 0.17
SASV Fusion	Keys manager	3.59 ± 0.38
	Audio Duration	
Enrollment		9.40 ± 0.03
Authentication		9.40 ± 0.03

2.3. Behavioural Biometric Authentication Module

One of the core components of SIBERIA is the behavioral biometric authentication module, deployed directly on the protected service. It centrally collects and analyzes features derived from the user's keyboard and mouse interaction patterns. This design enables tighter control over the authentication process, streamlines integration with existing security mechanisms, and allows centralized updates to the model. By leveraging human–machine interaction signals already present in typical industrial software environments, the module eliminates the need for additional biometric hardware, representing a significant advantage in regulated operational contexts.

The real-time collection of these features enables the system to capture unique interaction signals that are difficult for attackers to imitate or replicate accurately [47,48]. Studies have shown that attributes such as cursor speed, keystroke intervals, and the use of modifier keys (e.g., Shift or Ctrl) exhibit stable, user-specific patterns over time [49]. This stability makes these signals a reliable source for adaptive verification and CA systems. In our implementation, feature snapshots are sampled every 1 s from the protected application's backend telemetry, which is linked to the authenticated user session. The inference window size is administratively configurable; a 10 s sliding window with a 1 s stride is recommended as a practical balance between responsiveness and statistical stability. Longer windows reduce noise but slow detection, whereas shorter windows improve reactivity but may result in sparse data during low-activity periods.

Tables 4 and 5 show the features extracted from the user´s mouse and keyboard behavior, respectively, and used later to train the DNN models for behavioral biometrics.

Table 4. Mouse-based behavioral biometric features.

Feature	Description
cursor_speed	Average speed of cursor movement during usage
left_clicks	Total number of left mouse button clicks
right_clicks	Total number of right mouse button clicks
left_double_clicks	Total number of left double-click actions
right_double_clicks	Total number of right double-click actions
movement_to_click_time	Average time from cursor movement to a click

Table 5. Keyboard-based behavioral biometric features.

Feature	Description
avg_key_interval	Average time interval between two key presses
avg_key_duration	Average duration a key remains pressed
backspace	Number of times the backspace key is used
enter	Number of enter key presses
shift	Number of times the shift key is activated
ctrl	Number of control key presses
alt	Number of Alt key usage
caps_lock	Number of caps lock key activations
tab	Number of tab key presses
esc	Number of escape key usage
space	Number of spacebar presses

The variables listed in Table 4 extracted from mouse interactions capture the user's motor behavior during system use. For example, cursor speed and the latency between movement and click reflect individual differences in movement precision and decisiveness. Likewise, the frequency of single and double clicks helps characterize habitual interaction patterns with the interface, among other aspects.

In the case of the keyboard features shown in Table 5, the selected variables describe key aspects of typing dynamics. Metrics such as the time between keystrokes and key press durations are well-established in keystroke dynamics research. Special keys—such as Backspace, Shift, or Ctrl—provide further discriminatory power, revealing individual habits in error correction and the execution of key combinations, thereby enriching the user's biometric profile.

This module, integrated within the CA architecture, acts as an additional security layer automatically activated during regular application usage. It captures and analyzes behavioral biometric patterns derived from mouse and keyboard activity without impacting the user experience. Data capture occurs at the backend associated with the authenticated user session. Raw event streams remain local to the protected service, while only derived anomaly scores or risk signals need to be shared with higher-level orchestration components. This approach reinforces the privacy-preserving design of the broader SIBERIA framework.

To effectively perform this analysis, the system employs a DNN model designed to process the extracted features during application usage. The model is optimized to handle continuously collected temporal data, enabling it to identify unique behavioral patterns over time. Its architecture is tailored to integrate multiple information sources and to deliver continuous identity predictions throughout the session.

Table 6 summarizes the DNN model designed to process behavioral biometric signals. The model adopts a dual-branch autoencoder structure, with distinct convolutional and recurrent pathways independently processing mouse and keyboard signals. These temporal sequences are analyzed separately in each branch, and their representations are fused into a shared latent layer. From this combined representation, the decoders reconstruct the original signals, enabling both user identity verification and anomaly detection.

The input sequence length is set to 10, allowing the model to capture temporal windows that are representative of the user's behavior. The latent space dimension is fixed to 32, providing a compact yet expressive encoding of biometric features. Training is performed in batches of 16 samples, balancing computational efficiency and learning stability, over 10 epochs. Mean squared error is employed as the loss function, as it is well-suited for measuring the difference between the original sequence and its reconstruction by the autoencoder. After training, a decision threshold is calculated to detect significant deviations from the user's normal behavior pattern. This threshold is defined as the mean of the combined reconstruction error plus three times its standard deviation, following a statistical criterion that promotes high sensitivity to anomalies without compromising the false-positive rate.

Table 6. Summary of the dual autoencoder model architecture.

Layer Mouse (M)/Keyboard (K)	Description	Output Shape	Parameters
InputLayer × 2 (M input/K input)		(None, 10, 6)/(None, 10, 11)	0/0
Conv1D × 2 (M/K)	Extract patterns from cursor	(None, 10, 64)/(None, 10, 64)	1216/12,352
Dropout	Helps prevent overfitting	(None, 10, 64)	0
MaxPooling1D	Downsamples sequence length	(None, 5, 64)	0
LSTM × 2 (M/K)	Summarize temporal information	(None, 5, 128)/(None, 64)	98,816/49,408
Conv1D × 2 (M/K)	Extract patterns from keyboard	(None, 10, 64)/(None, 10, 64)	2176/12,352
Dropout	Helps prevent overfitting	(None, 10, 64)	0
MaxPooling1D	Downsamples sequence length	(None, 5, 64)	0
LSTM × 2 (M/K)	Summarize temporal information	(None, 5, 128)/(None, 64)	98,816/49,408
Concatenate	128-dimensional vector	(None, 128)	0
Dense (latent)	ReLu activation	(None, 32)	4128
RepeatVector (mouse)	Match half original sequence length	(None, 5, 32)	0
LSTM	Decodes temporal patterns	(None, 5, 64)	24,832
UpSampling1D	Restores original sequence length	(None, 10, 64)	0
Conv1D	Refines output	(None, 10, 64)	12,352
TimeDistributed (mouse output)	Reconstructed cursor signal	(None, 10, 6)	390
RepeatVector (keyboard)	Match half original sequence length	(None, 5, 32)	0
LSTM	Decodes temporal patterns	(None, 5, 64)	24,832
UpSampling1D	Restores original sequence length	(None, 10, 64)	0
Conv1D	Refines output	(None, 10, 64)	12,352
TimeDistributed (keyboard output)	Reconstructed keyboard signal	(None, 10, 11)	715
Total parameters			404,145
Trainable parameters			404,145
Non-trainable parameters			0

The model is trained using samples collected over a 5 min period, with a sampling rate of one sample per second on a single computer. To preserve the temporal coherence of the sequence, the samples are kept contiguous without any random shuffling. The full dataset is initially split by reserving the last 20% of the samples, which is further divided equally: 50% of this subset is used to generate the model's threshold, while the remaining 50% serves for validation. The initial 80% is used as the training set. For operational deployments, however, we recommend collecting data across all application sections where the user performs work tasks (e.g., monitoring, input forms, and maintenance tools) to ensure that the learned profile captures context-dependent interaction styles and reduces false alarms when workflows change.

In real environments, several practical factors can influence CA performance. These include (1) behavioral drift over time as users change habits, (2) hardware variability or latency across thin clients, (3) sparse interaction intervals when operators monitor rather than interact, (4) automated macro or scripted inputs that may attempt to mimic user activity, and (5) privacy constraints regarding the storage of raw interaction traces. The current implementation addresses these by supporting periodic re-calibration of user thresholds using recent accepted sessions, normalizing timing to device timestamps with optional per-endpoint calibration, adapting window size or confidence weighting in low-activity periods, detecting low-entropy timing signatures indicative of scripted replay, and retaining raw events locally while exporting only aggregated scores to upstream decision logic.

Table 7 presents the model validation results on the HMOG dataset, which contains biometric data from more than 100 users applied in a similar CA context.

Table 7. HMOG's results for the validation of 1 user on all other users of the dataset.

Accuracy	F1-Score	Accuracy (Std)	F1-Score (Std)
0.8592	0.9132	0.1795	0.1217

3. Results

This section presents the operational flow of the SIBERIA system, illustrating how its components interact to provide secure and self-sovereign authentication and authorization within industrial environments. The complete process is detailed through the typical user lifecycle, including identity creation, credential issuance, service access, and ongoing monitoring. Each step highlights the role of the SSI, secure SASV, and behavioral biometric modules. Beyond the user-centric perspective, the system also includes preparatory mechanisms that establish the trust infrastructure supporting identity and credential verification. The overall operational flow of the SIBERIA system is outlined as follows:

1. Initial System Configuration: The trust infrastructure is established by registering foundational elements necessary for credential issuance and verification.
2. User Identity Creation: Users generate decentralized identities via a mobile identity wallet, enabling secure key management and system interaction.
3. Issuer Onboarding: Credential issuers are integrated by registering their identities and setting up the components required to issue credentials in a trusted and verifiable manner.
4. Credential Issuance: Users request biometric credentials by providing requisite information or biometric data. The issuer evaluates these requests and issues VCs, optionally storing associated metadata or status information.
5. Credential Storage: The issued credential is securely stored in the holder's wallet for later use.
6. Service Access Request: The user presents the required credentials when attempting to access a protected service, along with fresh input to support authentication.
7. Verification and Access Grant: The verifier (the service owner) validates the authenticity of the presentation and compares it with the provided evidence to determine whether access should be granted.
8. Ongoing Monitoring: Once access is granted, SIBERIA performs CA to ensure that user behavior remains consistent with expectations, triggering actions if anomalies are detected.

3.1. Initial System Configuration

To initially set up the SIBERIA environment, it is necessary to register the DID and DID document of the TAO, the trusted issuer of SIBERIA. Through the SIBERIA API, the TAO's DID is created along with its corresponding public and private keys. Subsequently, this identity is registered in the DID Registry, thereby establishing the TAO as a recognized entity within the system.

Next, the general schema called SIBERIA Verifiable Attestation is defined and registered. This schema forms the basis for all credentials issued by the various issuers and incorporates the minimum required elements for credentials, following recommendations from EBSI and W3C. These elements include the following:

- @context: Semantic context defining the terms used within the document.
- id: Globally unique identifier for the verifiable credential.
- type: Declaration of the credential type.
- issuer: Entity issuing the credential.
- issuanceDate: Official issuance date.

- issued: Similar to issuanceDate, indicating the issuance date.
- validFrom: Date from which the credential is valid.
- credentialSubject: Information and claims related to the subject of the credential.
- credentialSchema: Reference to the structure that the credential adheres to.

Additionally, other fundamental schemas are registered for the system's operation, including the following:

- The "SIBERIA Verifiable Presentation" schema, defining the structure of VPs created by holders when sharing their credentials with verifiers.
- The "SIBERIA StatusList2021 Credential" schema, representing status lists indicating the validity or revocation of other credentials; its use is optional if the credential's validity is specified directly in the issued credential.
- The "SIBERIA Voice ID" schema, used for credentials issued through the secure SASV module.

Finally, the TAO is registered as a TI in the TIR, authorizing it to issue onboarding credentials to other issuers and to issue credentials conforming to the SIBERIA Verifiable Attestation schema. These steps establish the initial deployment of the SIBERIA environment, enabling secure and controlled participation of SIBERIA stakeholders: regular users (holders), owners of the service to be protected (verifiers), and SASV module owners (issuers).

3.2. User Identity Creation

The next phase involves onboarding a user (holder) into the SIBERIA ecosystem. This process is carried out via the SIBERIA digital wallet, where the user initiates registration by providing an email and password. Upon completion, the system automatically assigns a DID along with an associated public–private key pair and DID document, thereby establishing a secure and verifiable digital identity. The wallet further enables users to manage their VCs, export their identity data in JSON format, and optionally enroll biometric authentication—such as fingerprint recognition—to enhance account security and facilitate user authentication.

3.3. Issuer Onboarding

The onboarding process for a new TI within SIBERIA involves a structured sequence of steps designed to ensure the authenticity and authorization of the entity to issue VCs. In this use case, typically, the entity in charge of issuing credentials is the Secure SASV Module, but the registration of additional issuers that can issue credentials within the SIBERIA ecosystem is supported.

First, the entity must create its identity as an issuer. This is performed via the SIBERIA API, which generates a DID of type legal entity, a cryptographic key pair (public and private), and a DID document that defines the decentralized identity of the entity.

Once the identity is created, it must be registered in the DID Registry. Following EBSI guidelines, the entity must request a credential from the TAO known as the Verifiable Authentication Onboard via the TAO API. This credential serves as proof that the entity has been evaluated and accredited by the TAO to operate as an issuer within the SIBERIA ecosystem.

With this credential, the entity requests a token from the Token Issuance API to register its identity in the DID Registry. Upon validation of the credential, a token is issued, permitting the entity to register its DID document using the SIBERIA API.

This registration process within the DID Registry is illustrated in Figure 5.

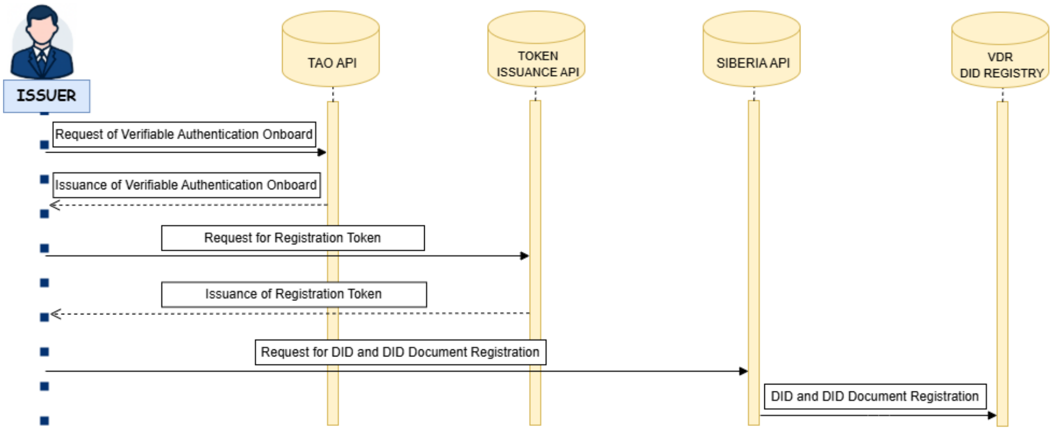
ISSUER
TAO API
TOKEN ISSUANCE API
SIBERIA API
VDR DID REGISTRY
Request of Verifiable Authentication Onboard
Issuance of Verifiable Authentication Onboard
Request for Registration Token
Issuance of Registration Token
Request for DID and DID Document Registration
DID and DID Document Registration

Figure 5. Registration in DID registry.

Following successful registration in the DID Registry, the entity may proceed to be recognized as a TI. At this stage, two scenarios are considered: the entity may either adhere to the existing schemas registered in the system or it may require additional or custom schemas. It is important to emphasize that adherence to the general SIBERIA Verifiable Attestation schema is mandatory for all TIs, as it serves as the foundational structure for all issued credentials.

If new schemas are necessary, they must first be registered in the TSR. This process involves the entity requesting a token via the Token Issuance API, followed by submitting the schema registration through the SIBERIA API.

Once all required schemas are registered, the entity can proceed with its registration in the TIR. As illustrated in Figure 6, the entity must request a new credential from the TAO, known as Verifiable Accreditation to Attest, again through the TAO API. This credential certifies the entity's authorization to act as a TI and specifies the schemas it is permitted to use. Concurrently, the TAO reserves a corresponding entry in the TIR for the entity.

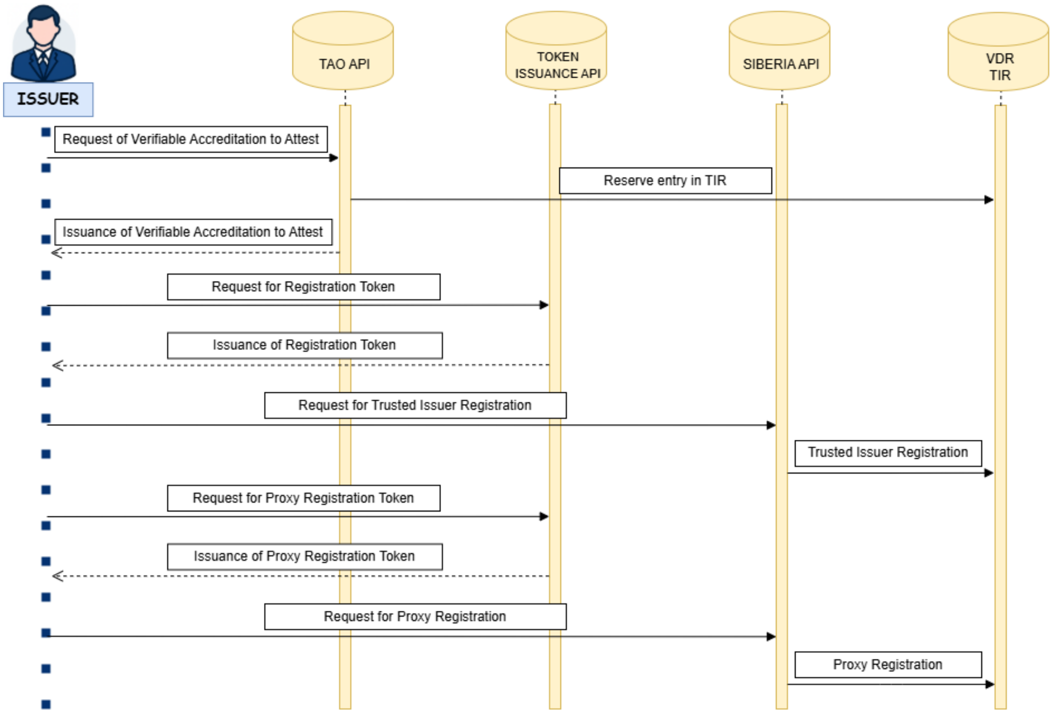
ISSUER
TAO API
TOKEN ISSUANCE API
SIBERIA API
VDR TIR
Request of Verifiable Accreditation to Attest
Reserve entry in TIR
Issuance of Verifiable Accreditation to Attest
Request for Registration Token
Issuance of Registration Token
Request for Trusted Issuer Registration
Trusted Issuer Registration
Request for Proxy Registration Token
Issuance of Proxy Registration Token
Request for Proxy Registration
Proxy Registration

Figure 6. Registration in TIR.

Subsequently, the entity presents this credential to the Token Issuance API, which, upon successful validation, issues a token for TIR registration. Using this token, the entity completes the onboarding process by submitting its registration request through the SIBERIA API. The TIR is updated in the reserved entry with the TAO-issued credential, which is made publicly accessible to ensure transparency within the ecosystem.

Optionally, if the issuer supports the SIBERIA StatusList2021 Credential schema, it may register a status list proxy that allows verifiers to query the real-time status of credentials it has issued, e.g., whether they are revoked, suspended, or valid. This involves requesting a token from the Token Issuance API, followed by proxy registration via the SIBERIA API, specifying the endpoint for accessing the status list. This proxy is then linked to the issuer's public profile in the TIR. Upon completion of these procedures, the entity is fully onboarded as a TI and is authorized to issue VCs under the approved schemas within the SIBERIA ecosystem.

3.4. Credential Issuance

The next step in the process involves the issuance of a VC that enables users to access the various services within the SIBERIA ecosystem. To allow users to participate in the SIBERIA ecosystem and access its services, a credential issuance process is carried out. This involves creating a digital proof of identity, such as the voice credential, which users will store and control in their digital wallets.

As part of the onboarding procedure, participating entities must issue holders a voice-based credential, utilizing the SASV module developed within SIBERIA.

To generate this credential, the holder must first provide their DID along with a set of audio samples. Based on this input, a dedicated public–private key pair is generated for the user's voice-based authentication. The audio samples are then converted to Base64 format and processed during the enrollment phase, resulting in the creation of a unique voiceprint. This voiceprint, encrypted with the user's public key, serves as the core biometric identifier of the credential.

The issued credential conforms to the mandatory general schema defined by SIBERIA, which ensures consistency and interoperability across the ecosystem. According to the SIBERIA Voice ID schema, the voiceprint is embedded as the primary attribute within the credentialSubject field, alongside the holder's DID. This structure is compatible with both the general schema and the specific requirements of voice-based credentials.

Additionally, issuers may choose to manage the validity of the credential dynamically. Instead of specifying a static expiration date via the validUntil field, he issuer can leverage a status list mechanism based on the SIBERIA StatusList2021 Credential schema. This mechanism allows the issuer to track the real-time status of issued credentials—such as active, suspended, or revoked—within an authoritative registry.

This credential enables users to authenticate themselves across the service within the ecosystem by recognizing their voice, ensuring a high level of biometric security and a seamless user experience.

3.5. Credential Storage

Once the credential is issued and cryptographically signed by the issuer, it is delivered to the user in the form of a QR code. The user retrieves and scans this code using the SIBERIA mobile wallet application, which integrates a QR scanner for this purpose. Upon successful scanning, the credential is decoded and securely stored in the wallet, making it readily available for future use and disclosure when interacting with services.

It is important to note that the credential is not stored by the system. Aside from maintaining a reference to its current status, no personal data or content of the credential is retained on the issuer's side. This approach ensures that all sensitive information remains exclusively with the holder, stored locally on their device, in alignment with the principles of SSI.

Within the wallet, the credential can be locally managed by the user, including actions such as storing, deleting, or exporting it as a JSON file. At all times, the data remain under the sole control of the holder, reinforcing privacy, autonomy, and user-centric identity management.

The sequence of credential issuance and storage described above is illustrated in Figure 7.

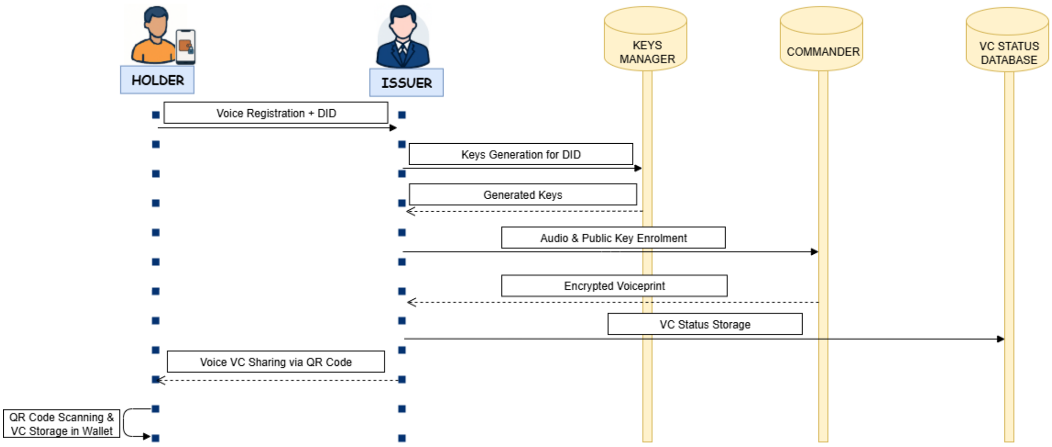

Figure 7. Credential issuance and storage.

3.6. Service Access Request

In order to initiate access to a protected service within the SIBERIA ecosystem, the holder must undergo a multi-factor authentication process that combines traditional and decentralized identity mechanisms. This step is essential to ensure that the individual attempting to access the service is indeed the legitimate owner of the VCs presented.

The access request process begins with a conventional login procedure, wherein the holder must provide an email (or alternatively, their DID) and a password. These credentials must have been previously registered and stored in the verifier's internal database. This first factor establishes a baseline for identity verification using traditional authentication methods.

Concurrently, the verifier displays a QR code, which the holder scans using their wallet application. Upon scanning, the holder is prompted to select and share the required VCs, including the mandatory voice-based credential. This sharing is executed through the creation of a Verifiable Presentation (VP), transmitted via a secure HTTP request between the wallet and the SIBERIA API. The VP includes the holder's DID, the selected VCs, and the holder's digital signature, ensuring data integrity and provenance. The structure of the VP conforms to the SIBERIA Verifiable Presentation schema, guaranteeing compatibility with the verification components of the system.

As a third factor, in accordance with the SASV module's requirements, the holder must record a live voice sample during the access attempt. This sample, captured in real time, is used as a biometric proof and is compared against the encrypted voiceprint included in the voice credential previously shared via the VP.

Only after successfully completing all three steps—traditional login, verifiable presentation submission, and live voice sample capture—does the access request proceed to the final verification phase. This multi-layered approach significantly enhances the security of the system, ensuring that identity assertions are trustworthy, private, and resistant to impersonation.

3.7. Verification and Access Grant

This phase is responsible for determining whether the holder is indeed the legitimate subject of the shared VCs and whether access to the requested service should be granted. To achieve this, a multi-step verification process is carried out, combining traditional authentication methods, cryptographic credential validation, and secure SASV.

As illustrated in Figure 8, the process begins with a classical authentication mechanism, based on the verification of an identifier (email address or DID) and a password previously registered in the service provider's database. If the provided credentials do not match the stored records, access is immediately denied.

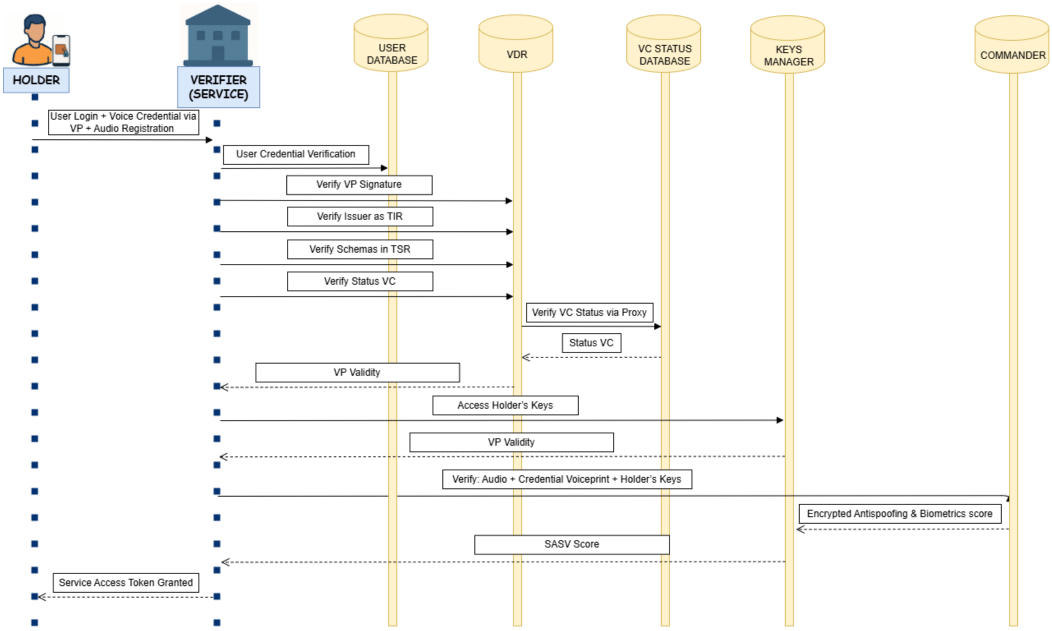
HOLDER
VERIFIER (SERVICE)
USER DATABASE
VDR
VC STATUS DATABASE
KEYS MANAGER
COMMANDER
User Login + Voice Credential via VP + Audio Registration
User Credential Verification
Verify VP Signature
Verify Issuer as TIR
Verify Schemas in TSR
Verify Status VC
Verify VC Status via Proxy
Status VC
VP Validity
Access Holder's Keys
VP Validity
Verify: Audio + Credential Voiceprint + Holder's Keys
Encrypted Antispoofing & Biometrics score
SASV Score
Service Access Token Granted

Figure 8. Verification and access grant.

Following successful traditional authentication, the system proceeds to validate the VP received from the holder. The verifier submits an HTTP request to the SIBERIA API, which handles the VP validation. This involves verifying the cryptographic signature of the holder using their registered public key, ensuring the presentation has not been altered and that the identity of the sender is authentic. Additionally, the VP must comply with the SIBERIA Verifiable Presentation schema, confirming structural and semantic integrity.

If the VP passes this validation stage, each of the credentials included within the presentation is individually assessed. The issuer's digital signature is first verified to ensure the credential has been issued by a TI registered in the TIR and that its contents have not been tampered with. The system also checks that the credential conforms to the expected schema and that the issuer is authorized—according to their TIR record—to issue credentials of that specific type.

Subsequently, the validity period of the credential is assessed. If the credential includes an explicit expiration date, it is checked against the current timestamp. However, if the credential does not define an expiration date, it is assumed that the issuer maintains an external status registry. In such cases, an HTTP request is sent to the issuer's status proxy—whose endpoint must be registered in the SIBERIA proxy registry—to obtain a status token.

This token contains a VC that represents the current status of the original credential. It must pass the same verification checks as any other credential: issuer authentication via the TIR, compliance with the SIBERIA Status Credential schema, and temporal validity.

If the status credential is valid, its content is then examined. In accordance with the W3C StatusList2021 specification, it includes fields such as the following:

- id: DID of the credential subject.
- type: Expected to be StatusList2021.
- statusPurpose: The purpose of the status entry (revocation or suspension)
- encodedList: A compressed bitstring that encodes the status of multiple credentials.
- statusListCredential: A reference to the credential containing the status list.

The encodedList is used to determine the status of the specific credential based on its index in the list. Following W3C recommendations, the list is GZIP-compressed to optimize storage and transmission. Each credential is assigned a bit: 1 indicates a revoked or suspended credential, while 0 indicates a valid one. Upon decompressing the list and identifying the relevant bit, the system determines whether the credential remains valid.

If the credential is confirmed as valid, the verifier then extracts the actual subject data, i.e., the holder's DID and encrypted voiceprint.

The final step consists of biometric verification through voice matching. The voice sample recorded by the holder during the service access request is converted to Base64. An HTTP call is then made to the SIBERIA Keys Manager module API in the SASV module to retrieve the voice authentication keys associated with the holder's DID. Using this information, a second HTTP request is sent to perform biometric verification.

This verification compares the encrypted voiceprint from the credential with the newly recorded audio sample. The process includes anti-spoofing analysis and returns two scores: one indicating the similarity between the voices, and another assessing the likelihood of the sample being genuine (i.e., not replayed or synthetically generated). These scores are used to compute a final biometric confidence score.

If this score exceeds a predefined threshold, the system considers the voices to match, confirming that the holder is the legitimate subject of the credential. At this point, an authentication token is issued to the holder, granting access to the requested service.

3.8. Ongoing Monitoring

Once access to the service has been granted, continuous user verification is maintained through the integration of the behavioral biometrics module implemented within the service infrastructure. This module enables real-time monitoring of user behavior to confirm the identity of the authenticated holder throughout the session.

The system captures and analyzes various behavioral parameters associated with user interaction, including mouse movement dynamics, click patterns, and keyboard input behavior. These parameters were detailed in Tables 4 and 5, which summarize the extracted features used to model user behavior through mouse and keyboard activity, respectively. The collected features are processed by a behavioral model specifically trained for individual user profiles.

This personalized training enables the model to establish a precise behavioral baseline for each subject. During application usage, live input is continuously compared against this baseline, allowing for dynamic and adaptive identity verification. If the model detects significant deviations from the holder's expected behavior—indicating a potential account takeover—the system raises an alert within the service environment. If these anomalies persist over time, the system will automatically terminate the user's session, requiring a new authentication process to regain access.

The classification and severity of these alerts are determined by a systematic evaluation process. This process involves comparing the prediction values of user-extracted samples against a predefined threshold. To ensure a consistent and uniform scale, these samples are first normalized to the range [0, 1]. Based on the resulting value's deviation from the threshold, the system classifies the alert into one of three escalating levels of severity: (1) basic, for values between the threshold and the threshold plus twice the standard deviation of the error; (2) medium, for values exceeding the basic tier's upper limit but remaining below the midpoint of the remaining interval towards 1; and (3) high, for any value exceeding this midpoint up to the maximum scale value of 1.

$$\text{Basic level:} \quad x \in [T,\ T + 2\sigma)$$

$$\text{Medium level:} \quad x \in \left[T + 2\sigma,\ T + 2\sigma + \frac{1 - (T + 2\sigma)}{2}\right)$$

$$\text{High level:} \quad x \in \left[T + 2\sigma + \frac{1 - (T + 2\sigma)}{2},\ 1\right]$$

where:

- x denotes a normalized prediction sample, with $x \in [0, 1]$;
- T is the predefined threshold;
- σ is the standard deviation of the prediction error.

The incorporation of this alert-level system enables the activation of response mechanisms associated with each detected risk level.

4. Discussion

SIBERIA represents a significant advancement in the implementation of authentication mechanisms for industrial environments by integrating voice and behavioral biometrics, combined with SSI principles. This innovative approach overcomes the limitations of traditional models based solely on passwords or physical tokens, offering an architecture that ensures not only robust identity verification but also CA throughout the user's active session.

One of the system's most relevant contributions is the incorporation of voice-based biometric credentials as another authentication factor, eliminating the need for additional hardware. This approach aligns with data minimization and user-control principles fundamental to SSI frameworks, while also ensuring the confidentiality of biometric information through the use of HE techniques. This mechanism guarantees that voiceprints are processed in their encrypted form and are never exposed in plaintext.

The decentralized management of VCs, stored locally in user-held wallets, provides granular control over the entire identity lifecycle, including issuance, revocation, and recovery. Furthermore, the technical structure of the behavioral biometric model, based on a dual architecture that separately processes mouse and keyboard signals, enables the precise and complementary characterization of temporal interaction patterns. This enhances the system's robustness against individual variability and improves the reliability of CA.

Operationally, SIBERIA balances security and usability, integrating mechanisms that ensure secure credential reissuance in case of wallet loss while maintaining a smooth user experience. The adoption of advanced cryptographic methods, combined with decentralized identity management, reinforces system integrity and aligns with privacy-by-design principles. As a result, SIBERIA is positioned as a reliable and adaptable solution for modern industrial access management challenges.

5. Conclusions

SIBERIA demonstrates a novel and effective approach to identity management in industrial environments, enhancing both security and user autonomy [5,33]. By integrating an SSI framework with advanced biometric technologies, it addresses critical vulnerabilities found in conventional centralized authentication systems [4].

The SIBERIA architecture demonstrates a successful fusion of three key technologies: a decentralized identity model based on SSI principles, multi-factor authentication through SASV, and CA through behavioral biometrics. The use of an SSI framework, aligned with European standards such as EBSI and GDPR, empowers users by giving them direct control over their digital credentials, which are stored securely in a personal digital wallet [6].

A significant contribution of this work is the enhancement of biometric security and privacy. The SASV module not only provides robust protection against sophisticated spoofing and deepfake attacks but also guarantees the confidentiality of sensitive voiceprint data by using HE. This allows verification computations to be performed directly on encrypted data, ensuring that raw biometric templates are never exposed. Furthermore, the behavioral biometrics module offers a non-intrusive layer of continuous security by monitoring user interaction patterns post-authentication, enabling the detection of potential session hijacking in real-time.

Moreover, this work demonstrates the validity of the integrated behavioral biometric model, whose effectiveness lies in a robust structural design based on a dual-head architecture. This configuration, which separates and processes user behavior signals from mouse and keyboard, enables the precise and complementary capture of temporal interaction patterns. This structure not only facilitates the efficient reconstruction of the original signals but also enables a robust characterization of user behavior, an essential feature for a reliable CA system. Its integration within SIBERIA enhances the detection of behavioral deviations accurately and non-intrusively, consolidating its applicability in industrial environments with high security demands.

Ultimately, SIBERIA proves that it is feasible to build a secure, scalable, and user-centric identity management system for critical industrial services. It establishes a strong balance between high-level security and user autonomy, mitigating risks such as credential theft and impersonation while ensuring regulatory compliance.

6. Limitations and Future Work

While SIBERIA presents a robust and innovative approach to industrial identity management, it is important to acknowledge certain technical and operational limitations inherent to its current implementation. Recognizing these constraints not only provides a more realistic assessment of the system's applicability but also helps to outline clear directions for future research. This section discusses the most relevant limitations identified during the development and validation phases, as well as prospective improvements and architectural extensions aimed at enhancing the system's scalability, interoperability, and resilience in real-world industrial contexts.

6.1. Limitations

One of the main operational constraints of SIBERIA lies in its reliance on persistent Internet connectivity. Core operations such as DID creation, credential issuance, and credential verification require continuous network availability. In industrial environments with intermittent or constrained connectivity, such as remote facilities, this dependency may hinder operational continuity.

Additionally, while the mobile wallet adheres to privacy-by-design principles, its current implementation is limited to Android-based platforms. This restricts its applicability in contexts where other platforms such as iOS are prevalent and highlights the need for cross-platform alternatives.

From an infrastructure perspective, the large number of operations involved in the authentication flow, such as multiple HTTP requests for presentation verification and credential status retrieval, can introduce significant latency under heavy system load. This challenge is further exacerbated by the reliance on blockchain interactions, where slow or unstable connectivity may degrade system responsiveness and user experience.

Finally, the behavioral biometrics module is designed for deployment within the backend of the protected service, ensuring that sensitive data remains local. While this approach supports privacy and low-latency operation, it presumes the presence of a compatible backend infrastructure, which may limit integration in legacy or resource-constrained environments.

6.2. Future Work

To address these limitations and further enhance system capabilities, several research directions have been identified:

- Performance and Scalability Optimization: Future efforts will focus on improving the efficiency of the authentication process by adopting optimized communication protocols and resource management techniques to reduce latency and increase responsiveness under high load.
- Enhanced Wallet Recovery Mechanisms: More seamless and secure mechanisms for wallet recovery will be explored, including social recovery schemes and multi-device synchronization approaches, in alignment with SSI principles.
- Adaptability of Behavioral Biometrics: The behavioral biometrics module will be refined to support scenarios with minimal interaction or changing user workflows, potentially through online learning strategies and the incorporation of contextual data.
- Adaptive Decision Thresholds: Dynamic thresholding mechanisms will be studied, allowing the system to adjust sensitivity based on contextual risk factors or historical authentication patterns.
- Integration with Edge Computing and IIoT Devices: Future versions of SIBERIA may delegate parts of the authentication process, such as preliminary verification or anomaly detection, to edge computing nodes or IIoT gateways. This could increase autonomy and resilience in environments with limited network access.
- Validation in Real-World Scenarios: The system will be validated in representative industrial contexts that replicate actual operating conditions and involve a diverse user base to improve model generalizability and robustness.
- Multi-Level Alert System: A hierarchical alert mechanism is envisioned, capable of responding to anomalous behavior with graduated actions, ranging from warnings to re-authentication requests. This would strengthen security without degrading user experience.

Author Contributions: Methodology, D.P.-G., J.Á.F.-C., J.C.V.-C., and A.Á.M.; Software, D.P.-G., J.C.V.-C., S.A.M.-A., A.G.-D., and Y.d.D.L.; Validation, D.P.-G., J.Á.F.-C., J.A.M.L., J.C.V.-C., and Y.d.D.L.; Formal analysis, J.C.V.-C., A.Á.M.; Investigation, D.P.-G., J.Á.F.-C., J.C.V.-C., and A.Á.M.; Resources, A.Á.M.; Data curation, J.A.M.L., I.J.Y., S.A.M.-A. and H.A.I.; Visualization, D.P.-G., J.C.V.-C., and Y.d.D.L.; Supervision, D.P.-G., J.Á.F.-C., and A.Á.M.; Project administration, D.P.-G., J.Á.F.-C., and A.Á.M.; Writing—original draft: D.P.-G., J.Á.F.-C., J.A.M.L., J.C.V.-C., and I.J.Y.; Writing—review and editing, D.P.-G., J.Á.F.-C., and J.C.V.-C. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Vicomtech grant number E22109/2024. The APC was funded by Vicomtech.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ASV	Automatic Speaker Verification
CA	Continuous Authentication
CNN	Convolutional Neural Network
DID	Decentralized Identifier
DNN	Deep Neural Network
DPoS	Delegated Proof-of-Stake
EBSI	European Blockchain Services Infrastructure
EER	Equal Error Rate
FAR	False Acceptance Rate
FRR	False Rejection Rate
HE	Homomorphic Encryption
HMOG	Hand Movement, Orientation, and Grasp
IAM	Identity and Access Management
IIoT	Industrial IoT
IoT	Internet of Things
JSON	JavaScript Object Notation
KNN	K-Nearest Neighbors
LSTM	Long Short-Term Memory
MSI	Method-Specific Identifier
OIDC	OpenID Connect
OT	Operational Technology
PLDA	Probabilistic Linear Discriminant Analysis
ResNet	Residual Network
SASV	Spoofing Aware Speaker Verification
SSI	Self-Sovereign Identity
SSL	Self-Supervised Learning
SVM	Support Vector Machine
TAO	Trusted Accreditation Organization
TI	Trusted Issuer
TIR	Trusted Issuers Registry
TSR	Trusted Schema Registry
TTS	Text-to-Speech
VC	Verifiable Credential
VDR	Verifiable Data Registry
VP	Verifiable Presentation
W3C	World Wide Web Consortium

References

1. Preukschat, A.; Reed, D. *Self-Sovereign Identity*; Manning Publications: Shelter Island, NY, USA, 2021.
2. Bhattacharyya, D.; Ranjan, R.; Alisherov, F.; Choi, M. Biometric authentication: A review. *Int. J. Serv. Sci. Technol.* **2009**, *2*, 13–28.
3. Wagner, K.; Némethi, B.; Renieris, E.; Lang, P.; Brunet, E.; Holst, E. *Self-Sovereign Identity: A Position Paper on Blockchain Enabled Identity and the Road Ahead*; Blockchain Bundesverband: Berlin, Germany, 2018; pp. 1–56.
4. Mühle, A.; Grüner, A.; Gayvoronskaya, T.; Meinel, C. A survey on essential components of a self-sovereign identity. *Comput. Sci. Rev.* **2018**, *30*, 80–86.
5. Soltani, R.; Nguyen, U.T.; An, A. A survey of self-sovereign identity ecosystem. *Secur. Commun. Netw.* **2021**, *2021*, 1–26.
6. Allende López, M. Self-Sovereign Identity: The Future of Identity: Self-Sovereignity, Digital Wallets, and Blockchain. 2020. Available online: https://publications.iadb.org/en/self-sovereign-identity-future-identity-self-sovereignity-digital-wallets-and-blockchain (accessed on 14 May 2025).
7. World Wide Web Consortium (W3C). Decentralized Identifiers (DIDs) v1.0. 2022. Available online: https://www.w3.org/TR/did-core/ (accessed on 14 May 2025).
8. Michalopoulos, F.; Misiakoulis, G.; Vavilis, S.; Niavis, H.; Loupos, K. A Verifiable Data Registry for Secure and Scalable Decentralised Identity Management in the IoT Context. In Proceedings of the International Symposium on Distributed Computing and Artificial Intelligence Salamanca, Spain, 25–27 June 2024; Springer: Berlin/Heidelberg, Germany, 2024; pp. 271–280.
9. Sporny, M.; Longley, D.; Chadwick, D.; Kellogg, G.; Herman, I. Verifiable Credentials Data Model 2.0. W3C Recommendation. 2025 Available online: https://www.w3.org/TR/vc-data-model-2.0/ (accessed on 17 May 2025).
10. European Commission. European Blockchain Services Infrastructure (EBSI). 2025. Available online: https://ec.europa.eu/digital-building-blocks/sites/display/EBSI (accessed on 14 May 2025).
11. Čučko, Š.; Keršič, V.; Turkanović, M. Towards a Catalogue of Self-Sovereign Identity Design Patterns. *Appl. Sci.* **2023**, *13*, 5395.
12. Stockburger, L.; Kokosioulis, G.; Mukkamala, A.; Mukkamala, R.R.; Avital, M. Blockchain-enabled decentralized identity management: The case of self-sovereign identity in public transportation. *Blockchain Res. Appl.* **2021**, *2*, 100014.
13. Fedrecheski, G.; Rabaey, J.M.; Costa, L.C.; Ccori, P.C.C.; Pereira, W.T.; Zuffo, M.K. Self-sovereign identity for IoT environments: A perspective. In Proceedings of the 2020 Global Internet of Things Summit (GIoTS), Dublin, Ireland, 3 June 2020; pp. 1–6.
14. Belchior, R.; Putz, B.; Pernul, G.; Correia, M.; Vasconcelos, A.; Guerreiro, S. SSIBAC: Self-sovereign identity based access control. In Proceedings of the 2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), Guangzhou, China, 29 December 2020–1 January 2021; pp. 1935–1943.
15. Shuaib, M.; Alam, S.; Alam, M.S.; Nasir, M.S. Self-sovereign identity for healthcare using blockchain. *Mater. Today Proc.* **2023**, *81*, 203–207.
16. SelfKey Foundation. SelfKey. Available online: https://selfkey.org/ (accessed on 15 July 2025).
17. Tobin, A.; Reed, D. The inevitable rise of self-sovereign identity. *Sovrin Found.* **2016**, *29*, 18.
18. lifeID. Available online: https://lifeid.io/ (accessed on 15 July 2025).
19. Evernym. https://github.com/evernym (accessed on 14 July 2025).
20. Hyperledger Foundation. Hyperledger Indy. 2024. Available online: https://www.hyperledger.org/projects/hyperledger-indy (accessed on 16 July 2025).
21. EverID Project. EverID. Available online: https://app.everpay.io/ (accessed on 14 July 2025).
22. ISA Global Cybersecurity Alliance (ISAGCA). ISA/IEC 62443 Series of Standards. 2025. Available online: https://isagca.org/isa-iec-62443-standards (accessed on 21 July 2025).
23. Stouffer, K.; Pease, M.; Tang, C.; Zimmerman, T.; Pillitteri, V.; Lightman, S.; Hahn, A.; Saravia, S.; Sherule, A.; Thompson, M. *Guide to Operational Technology (OT) Security*; Special Publication NIST SP 800-82 Rev. 3; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2022.
24. Microsoft. OT Cloud Enablement—Azure Active Directory Tenant. 2023. Available online: https://techcommunity.microsoft.com/blog/azurearchitectureblog/ot-cloud-enablement-%E2%80%93-azure-active-directory-tenant/3707898 (accessed on 18 April 2025).
25. Goel, A.; Rahulamathavan, Y. A comparative survey of centralised and decentralised identity management systems: Analysing scalability, security, and feasibility. *Future Internet* **2024**, *17*, 1.
26. Hardt, D. The OAuth 2.0 Authorization Framework. RFC 6749, 2012. Available online: https://www.rfc-editor.org/info/rfc6749 (accessed on 11 May 2025).
27. Sekar, R.R.; Masna, A.; Sharma, S.; Abraham, A.; Pagilla, P.R. Decentralized identity and access management (IAM) using blockchain. In Proceedings of the 2024 International Conference on Intelligent Systems for Cybersecurity (ISCS), Gurugram, India, 3–4 May 2024; pp. 1–6.
28. Ghadge, N. Use of blockchain technology to strengthen identity and access management (IAM). *Int. J. Inf. Technol.* **2024**, *1*, 1–17.

29. Mittal, A.; Dua, M. Automatic speaker verification systems and spoof detection techniques: Review and analysis. *Int. J. Speech Technol.* **2022**, *25*, 105–134.
30. Bai, Z.; Zhang, X.L. Speaker recognition based on deep learning: An overview. *Neural Netw.* **2021**, *140*, 65–99.
31. Jakubec, M.; Jarina, R.; Lieskovska, E.; Kasak, P. Deep speaker embeddings for speaker verification: Review and experimental comparison. *Eng. Appl. Artif. Intell.* **2024**, *127*, 107232.
32. Snyder, D.; Garcia-Romero, D.; Sell, G.; Povey, D.; Khudanpur, S. X-vectors: Robust dnn embeddings for speaker recognition. In Proceedings of the 2018 IEEE international conference on acoustics, speech and signal processing (ICASSP), Calgary, AB, Canada, 15–20 April 2018; pp. 5329–5333.
33. Desplanques, B.; Thienpondt, J.; Demuynck, K. Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification. *arXiv* **2020**, arXiv:2005.07143.
34. Koluguri, N.R.; Park, T.; Ginsburg, B. Titanet: Neural model for speaker representation with 1d depth-wise separable convolutions and global context. In Proceedings of the ICASSP 2022–2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Singapore, 22–27 May 2022; pp. 8102–8106.
35. Kanagasundaram, A.; Vogt, R.; Dean, D.; Sridharan, S. PLDA based speaker recognition on short utterances. In Proceedings of the Speaker and Language Recognition Workshop: Odyssey 2012, Singapore, 25–28 June 2012; pp. 28–33.
36. Wu, Z.; Evans, N.; Kinnunen, T.; Yamagishi, J.; Alegre, F.; Li, H. Spoofing and countermeasures for speaker verification: A survey. *Speech Commun.* **2015**, *66*, 130–153.
37. Guo, H.H.; Hu, Y.; Liu, K.; Shen, F.Y.; Tang, X.; Wu, Y.C.; Xie, F.L.; Xie, K.; Xu, K.T. Fireredtts: A foundation text-to-speech framework for industry-level generative speech applications. *arXiv* **2024**, arXiv:2409.03283.
38. Sahidullah, M.; Delgado, H.; Todisco, M.; Nautsch, A.; Wang, X.; Kinnunen, T.; Evans, N.; Yamagishi, J.; Lee, K.A. Introduction to voice presentation attack detection and recent advances. In *Handbook of Biometric Anti-Spoofing: Presentation Attack Detection and Vulnerability Assessment*; Springer: Berlin/Heidelberg, Germany, 2023; pp. 339–385.
39. Speech Arena. Speech DeepFake Leaderboard. 2025. Available online: https://huggingface.co/spaces/Speech-Arena-2025/Speech-DF-Arena (accessed on 3 May 2025).
40. Jung, J.w.; Tak, H.; Shim, H.j.; Heo, H.S.; Lee, B.J.; Chung, S.W.; Yu, H.J.; Evans, N.; Kinnunen, T. SASV 2022: The first spoofing-aware speaker verification challenge. *arXiv* **2022**, arXiv:2203.14732.
41. Council, E. Directive 2016/680 of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data by Competent Authorities for the Purposes of the Prevention, Investigation, Detection or Prosecution of Criminal Offences or the Execution of Criminal Penalties, and on the Free Movement of Such Data, and Repealing Council Framework Decision 2008/977/JHA. 2016. Available online: https://eur-lex.europa.eu/eli/dir/2016/680/oj (accessed on 2 May 2025).
42. *ISO/IEC 24745:2022*; Information Security, Cybersecurity and Privacy Protection—Biometric Information Protection. ISO: Geneva, Switzerland, 2022. Available online: https://www.iso.org/standard/75302.html (accessed on 30 July 2025).
43. Nautsch, A.; Isadskiy, S.; Kolberg, J.; Gomez-Barrero, M.; Busch, C. Homomorphic encryption for speaker recognition: Protection of biometric templates and vendor model parameters. *arXiv* **2018**, arXiv:1803.03559.
44. Botelho, C.; Teixeira, F.; Rolland, T.; Abad, A.; Trancoso, I. Pathological speech detection using x-vector embeddings. *arXiv* **2020**, arXiv:2003.00864.
45. Rahulamathavan, Y. Privacy-preserving similarity calculation of speaker features using fully homomorphic encryption. *arXiv* **2022**, arXiv:2202.07994.
46. Baig, A.F.; Eskeland, S. Security, Privacy, and Usability in Continuous Authentication: A Survey. *Sensors* **2021**, *21*, 5967. [CrossRef]
47. Verma, A.; Moghaddam, V.; Anwar, A. Data-Driven Behavioural Biometrics for Continuous and Adaptive User Verification Using Smartphone and Smartwatch. *Sustainability* **2022**, *14*, 7362. [CrossRef]
48. López, J.M.E.; Celdrán, A.H.; Esquembre, F.; Pérez, G.M.; Marín-Blázquez, J.G. A Supervised ML Biometric Continuous Authentication System for Industry 4.0. *IEEE Trans. Ind. Inform.* **2022**, *18*, 9132–9140. [CrossRef]
49. Madavarapu, J.B.; Mittal, M.; Salagrama, S.; Adnan, M.M.; Rana, A.; Yadav, K. Behavioral Biometrics Authentication Systems: Leveraging Machine Learning for Enhanced Cybersecurity. In Proceedings of the 2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE), Gautam Buddha Nagar, India, 9–11 May 2024; pp. 1478–1483. [CrossRef]
50. Sitová, Z.; Šeděnka, J.; Yang, Q.; Peng, G.; Zhou, G.; Gasti, P.; Balagani, K.S. HMOG: New Behavioral Biometric Features for Continuous Authentication of Smartphone Users. *IEEE Trans. Inf. Forensics Secur.* **2016**, *11*, 877–892. [CrossRef]
51. Uslu, U.; Özlem Durmaz İncel.; Alptekin, G.I. Evaluation of Deep Learning Models for Continuous Authentication Using Behavioral Biometrics. *Procedia Comput. Sci.* **2023**, *225*, 1272–1281. [CrossRef]
52. Alzahrani, S.; Alderaan, J.; Alatawi, D.; Alotaibi, B. Continuous Mobile User Authentication Using a Hybrid CNN-Bi-LSTM Approach. *Comput. Mater. Contin.* **2023**, *75*, 651–667. [CrossRef]

53. Mao, R.; Wang, X.; Ji, H. ACBM: Attention-based CNN and Bi-LSTM model for continuous identity authentication. *J. Physics Conf. Ser.* **2022**, *2352*, 012005. [CrossRef]
54. Mekruksavanich, S.; Jitpattanakul, A. Deep Learning Approaches for Continuous Authentication Based on Activity Patterns Using Mobile Sensing. *Sensors* **2021**, *21*, 7519. [CrossRef]
55. Cheon, J.H.; Costache, A.; Moreno, R.C.; Dai, W.; Gama, N.; Georgieva, M.; Halevi, S.; Kim, M.; Kim, S.; Laine, K.; et al. Introduction to homomorphic encryption and schemes. In *Protecting Privacy through Homomorphic Encryption*; Springer: Berlin/Heidelberg, Germany, 2021; pp. 3–28.
56. Microsoft SEAL (Release 4.1). Microsoft Research: Redmond, WA, USA, 2023. Available online: https://github.com/Microsoft/SEAL (accessed on 6 April 2025).
57. Vault. Hashicorp. 2025. Available online: https://developer.hashicorp.com/vault (accessed on 16 April 2025).
58. Martín-Doñas, J.M.; Rosello, E.; Gomez, A.M.; Álvarez, A.; López-Espejo, I.; Peinado, A.M. ASASVIcomtech: The Vicomtech-UGR speech deepfake detection and SASV systems for the ASVspoof5 Challenge. In Proceedings of the ASVspoof, Kos, Greece, 31 August 2024; pp. 144–151.
59. Landini, F.; Profant, J.; Diez, M.; Burget, L. Bayesian hmm clustering of x-vector sequences (vbx) in speaker diarization: Theory, implementation and analysis on standard tasks. *Comput. Speech Lang.* **2022**, *71*, 101254.
60. Nagrani, A.; Chung, J.S.; Zisserman, A. VoxCeleb: A Large-Scale Speaker Identification Dataset. In Proceedings of the INTERSPEECH, Stockholm, Sweden, 20–24 August 2017; p. 2616–2620.
61. Chung, J.; Nagrani, A.; Zisserman, A. VoxCeleb2: Deep speaker recognition. In Proceedings of the INTERSPEECH, Hyderabad, India, 2–6 September 2018; p. 1086–1090
62. Fan, Y.; Kang, J.; Li, L.; Li, K.; Chen, H.; Cheng, S.; Zhang, P.; Zhou, Z.; Cai, Y.; Wang, D. Cn-celeb: A challenging chinese speaker recognition dataset. In Proceedings of the ICASSP, Virtual, 4–8 May 2020; pp. 7604–7608.
63. VBHMM: X-Vectors Diarization. 2020. Available online: https://github.com/BUTSpeechFIT/VBx/tree/master/VBx/models/ResNet101_16kHz (accessed on 4 May 2025).
64. Babu, A.; Wang, C.; Tjandra, A.; Lakhotia, K.; Xu, Q.; Goyal, N.; Singh, K.; Von Platen, P.; Saraf, Y.; Pino, J.; et al. XLS-R: Self-supervised cross-lingual speech representation learning at scale. *arXiv* **2021**, arXiv:2111.09296.
65. Martín-Doñas, J.M.; Álvarez, A.; Rosello , E.; Gómez, Á.M.; Peinado, A.M. Exploring Self-supervised Embeddings and Synthetic Data Augmentation for Robust Audio Deepfake Detection. In Proceedings of the INTERSPEECH, Kos, Greece, 1–5 September 2024.
66. Nautsch, A.; Wang, X.; Evans, N.; Kinnunen, T.H.; Vestman, V.; Todisco, M.; Delgado, H.; Sahidullah, M.; Yamagishi, J.; Lee, K.A. ASVspoof 2019: Spoofing countermeasures for the detection of synthesized, converted and replayed speech. *IEEE Trans. Biom. Behav. Identity Sci.* **2021**, *3*, 252–265.
67. Liu, X.; Wang, X.; Sahidullah, M.; Patino, J.; Delgado, H.; Kinnunen, T.; Todisco, M.; Yamagishi, J.; Evans, N.; Nautsch, A.; et al. Asvspoof 2021: Towards spoofed and deepfake speech detection in the wild. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2023**, *31*, 2507–2522.
68. Müller, N.M.; Kawa, P.; Choong, W.H.; Casanova, E.; Gölge, E.; Müller, T.; Syga, P.; Sperl, P.; Böttinger, K. Mlaad: The multi-language audio anti-spoofing dataset. In Proceedings of the 2024 International Joint Conference on Neural Networks (IJCNN), Yokohama, Japan, 30 June–5 July 2024; pp. 1–7.
69. Casanova, E.; Weber, J.; Shulby, C.D.; Junior, A.C.; Gölge, E.; Ponti, M.A. Yourtts: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone. In Proceedings of the International Conference on Machine Learning, Baltimore, ML, USA, 17–23 July 2022; pp. 2709–2720.
70. Zhang, Y.; Jiang, F.; Duan, Z. One-class learning towards synthetic voice spoofing detection. *IEEE Signal Process. Lett.* **2021**, *28*, 937–941.
71. Wang, X.; Delgado, H.; Tak, H.; Jung, J.w.; Shim, H.j.; Todisco, M.; Kukanov, I.; Liu, X.; Sahidullah, M.; Kinnunen, T.H.; et al. ASVspoof 5: Crowdsourced speech data, deepfakes, and adversarial attacks at scale. In Proceedings of the ASVspoof, Kos, Greece, 31 August 2024; pp. 1–8.
72. Shim, H.J.; Jung, J.W.; Kinnunen, T.; Evans, N.; Bonastre, J.F.; Lapidot, I. a-DCF: An architecture agnostic metric with application to spoofing-robust speaker verification. In Proceedings of the ODYSSEY 2024, the Speaker and Language Recognition Workshop, Quebec City, QC, Canada, 18–21 June 2024.
73. Villalba, J.A.; Feng, T.; Thebaud, T.; Lee, J.; Narayanan, S.; Dehak, N. The SHADOW team submission to the ASVSpoof 2024 Challenge. In Proceedings of the Automatic Speaker Verification Spoofing Countermeasures Workshop (ASVspoof 2024), Kos, Greece, 31 August 2024; pp. 36–42. [CrossRef]
74. Kurnaz, O.; Demirtaş, S.C.; Büker, A.; Mishra, J.; Hanilçi, C. Spoofing-Robust Speaker Verification Using Parallel Embedding Fusion: BTU Speech Group's Approach for ASVspoof5 Challenge. *arXiv* **2024**, arXiv:2408.15877.

75. Rohdin, J.; Zhang, L.; Oldřich, P.; Staněk, V.; Mihola, D.; Peng, J.; Stafylakis, T.; Beveraki, D.; Silnova, A.; Brukner, J.; et al. BUT systems and analyses for the ASVspoof 5 Challenge. In Proceedings of the ASVspoof 2024, Kos, Greece, 31 August 2024; pp. 24–31.
76. Aliyev, A.; Kondratev, A. Intema system description for the ASVspoof5 Challenge: Power weighted score fusion. In Proceedings of the ASVspoof 2024, Kos, Greece, 31 August 2024; pp. 152–157.
77. Duroselle, R.; Boeffard, O.; Courtois, A.; Nourtel, H.; Champion, P.; Agnoli, H.; Bonastre, J.F. Data augmentations for audio deepfake detection for the ASVspoof5 closed condition. In Proceedings of the ASVspoof Workshop 2024, Kos, Greece, 31 August 2024; pp. 16–23.
78. Falez, P.; Marteau, T. Whispeak speech deepfake detection systems for the ASVspoof5 Challenge. In Proceedings of the ASVspoof 2024, Kos, Greece, 31 August 2024; pp. 32–35.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.