Description

Title

INTEGRATED CREDENTIALING SYSTEM WITH VERIFIABLE MICRO-CREDENTIALS, PREDICTIVE PRIVILEGING AI, BIOMETRIC ATTESTATION, AND SMART-CONTRACT GOVERNANCE

Technical Field

This invention relates to digital credentialing and privileging systems for professionals, particularly healthcare providers. More specifically, it addresses an integrated credential management platform that leverages distributed ledger technology, machine learning risk prediction, biometric authentication, smart contracts, and regulatory automation to securely manage and verify professional credentials and privileges across institutions.

Background of the Invention

Healthcare provider credentialing and privileging is a complex, multi-step process involving verification of qualifications, tracking of ongoing competencies, and ensuring compliance with regulatory standards. Conventional credentialing is often slow and manual, leading to onboarding delays, administrative burden, and compliance risks. Failures in credentialing can have severe consequences; for example, inadequate physician privileging has resulted in facility penalties and patient harm (e.g., the Joan Rivers case where a clinic lost accreditation due to a lapse in credentialing oversight). A need exists for a comprehensive solution that not only digitizes and automates credentialing, but also proactively enhances its reliability and scope.

Existing solutions only partially address this need. For instance, Silversheet launched an intelligent platform to digitize credential documents, automate primary source verification, and remind physicians of expirations . This reduces paperwork and errors, helping facilities maintain

compliance by replacing paper files and faxes with easy-to-use software . Medallion, another provider, offers an AI-powered platform that automates end-to-end onboarding, credentialing, and privileging workflows in line with Joint Commission standards. Medallion's system expedites privileging applications and uses built-in business rules to ensure compliance, thereby mitigating risk and reducing delays . Evercred takes a decentralized approach: it issues permanent digital credentials and creates a "perpetual, automated primary source of verification", built on self-sovereign identity standards (W3C and Decentralized Identity Foundation) . Evercred's solution empowers clinicians to own and control portable credentials, eliminating redundant reverification and leveraging open standards for trust . CertifyOS similarly advertises a "one accurate, AI-enabled source of truth" for provider data, using APIs to integrate credentialing, licensing, and network monitoring into healthcare workflows (press release, 2024). Despite these advances, current platforms lack certain forward-looking features such as predictive risk analytics, in-person biometric attestation of supervision, dynamic smart-contract controls on privileges, and privacy-preserving sanction checks.

In terms of prior art in the patent literature, some components of credentialing systems have been disclosed, but no single reference integrates all the functionalities of the present invention. Axuall's U.S. Patent No. 12,079,891 (issued Sept. 3, 2024) discloses "systems and methods for verifying and managing digital credentials" in a healthcare context. Axuall's patent teaches tracking the lifecycle of digitally verified credentials (identity, education, licenses, work history, outcomes, etc.) through a network of certified primary-source issuers, using a verifiable data registry for credential validity and status. It also describes credential presentation and subscription processes via agent software for various actors. However, while it introduces a blockchain-like data network for portable credentials, it does not address AI-based privileging risk prediction or biometric co-signing of procedures.

Another relevant thread of prior art involves blockchain and smart contracts for verifying practitioner qualifications. For example, WIPO Publication WO2020/192342A1 (Shenzhen Launch Tech.) describes a medical service platform wherein a provider's node on a blockchain must be authenticated by an authorized authority node before the provider can take on appointments . This ensures that only practitioners with verified qualifications (e.g., valid licenses and hospital employment) can participate, preventing "ghost" providers . While this approach uses blockchain-based identity verification of medical credentials to guard system access, it is focused on appointment scheduling and does not extend to granular skill tracking or AI risk analysis.

With respect to predictive analytics, services like Preverity have demonstrated that aggregating malpractice data and provider performance metrics can yield predictive models of malpractice risk (as of 2025, Preverity maintains one of the largest databases of malpractice events to score likelihood of future claims). Yet, such analytics have not been tightly integrated into credentialing or privileging workflows in prior solutions. No known patent specifically teaches using machine learning to forecast a provider's risk of future sanctions or malpractice as an input to credentialing decisions.

Concerning biometric attestation and privacy, cryptographic methods for identity and credential verification exist in other domains. Microsoft's US Patent 9,768,962 B2 (2017) discloses a "minimal disclosure credential verification" system using cryptographic accumulators and zero-knowledge proofs to verify that a user's identifier is not on a revocation list (or, conversely, is on a whitelist) without revealing the identifier itself . The system issues each user a credential with an embedded unique identifier and uses a revocation authority to update an accumulator of revoked IDs; the user can later produce a non-membership proof that their ID is not in the revoked set, thereby proving their credential is still valid . This technique, while not originally aimed at medical credentialing, is highly relevant to zero-knowledge sanction checks – confirming a provider has no disciplinary record without exposing sensitive data. Prior art does not show this being applied to sanction or license verification in healthcare, which the present invention adapts for confirming the absence of negative events in a privacy-preserving manner.

In summary, there is a clear need for an omnibus credentialing system that unifies these disparate innovations: a ledger of micro-credentials earned via real work, AI-driven risk prediction, biometric confirmation of training/supervision events, smart-contract-based privilege management, automated regulatory tracking, integration with scheduling systems, device-linked competency tracking, and zero-knowledge proofs for clean reputations. No single prior art reference or product addresses all these facets in one interoperable framework. The present invention fills this gap by providing a comprehensive platform that not only streamlines and secures credentialing, but also adds intelligence and future-proof adaptability to the credentialing and privileging process.

Summary of the Invention

Embodiments of the present invention provide a novel credentialing and privileging system combining multiple innovative components that collectively improve the verification, maintenance, and governance of professional credentials. In one aspect, the system comprises a "Skills-Currency" ledger module that translates live activity logs (from Electronic Health Records (EHRs), procedure documentation systems, and connected devices) into verifiable micro-credentials. Each procedure performed, device used, or skill demonstrated by a provider can be automatically captured and logged as a credential "token" or credit on an immutable ledger. These micro-credentials function like a currency of skills, accumulating over time to evidence experience and competencies (e.g., number of central line placements, or hours of ultrasound use). They are digitally signed and verifiable using, for example, W3C Verifiable Credential standards, making them portable across institutions and audit-ready . The ledger may be underpinned by blockchain technology to ensure tamper-proof recording and sharing of these credentials.

Another aspect is a Predictive Privileging AI component. This uses machine learning models trained on historical provider data (such as malpractice claims, patient outcomes, peer reviews, and sanction records) to forecast the risk of future adverse events for a given provider. For example, the AI can output a risk score or flag if a surgeon's pattern of outcomes and practice history statistically correlates with higher malpractice risk in the next year. This risk prediction

can be used by credentialing committees to make data-driven privileging decisions – for instance, identifying providers who may require additional oversight or denying certain privileges if predicted risk is above a threshold. The model continuously updates as new data (including the Skills-Currency ledger entries and external data like national practitioner databank reports) become available, effectively providing an early-warning system for potential quality or compliance issues. This is a proactive complement to traditional credentialing, which typically reacts to past sanctions rather than anticipating them.

The system further includes a Bio-Signed Procedure Attestation mechanism for in-person supervision and training verification. This feature enables cryptographic biometric confirmation that a given procedure or task was directly supervised or performed by specific individuals. In practice, when a trainee (e.g., a resident physician or a nurse on orientation) performs a procedure under supervision, both the trainee and the supervising practitioner can attest to the event in real-time using biometric authentication (such as a fingerprint scan, facial recognition, or FIDO2 security key) on a mobile device or workstation. The biometric factors are linked to cryptographic keys (for example, via the FIDO2/WebAuthn standard, which ensures a strong binding between the person's biometric and their digital credential). The attestation is then recorded on the system – e.g., as a signed verifiable credential stating "Dr. X supervised Procedure Y performed by Provider Z on date/time," with both parties' digital signatures. This ensures non-repudiation of supervision: it's verifiably proven that the supervising professional was physically present and confirmed the trainee's competency at that moment. Such bio-signed attestations can be required for high-stakes procedures or for initial granting of privileges, creating a trusted record of hands-on competency demonstration.

Additionally, the invention provides Smart-Contract Conditional Privileges. In this model, the assignment of clinical privileges (e.g., the ability to perform a certain surgery or operate a type of equipment) is governed by smart contracts that encode time- and scope-limited conditions. A privilege credential can be represented as a smart contract on a blockchain network (for example, a Hyperledger Fabric chain or Ethereum-based network), which automatically expires or revokes the privilege when conditions are not met. Conditions might include time limits (e.g., privileges auto-expire after 2 years unless renewed), continuing education requirements (e.g., completion of specific CME credits by a deadline), volume thresholds (e.g., at least 5 procedures performed per year to maintain proficiency), or compliance status (no new sanctions). The smart contract can query the Skills-Currency ledger and external data feeds to determine compliance with conditions. If a condition is violated or expiration date reached, the contract self-executes to suspend or revoke the credential, and can even notify relevant parties. This conditional privileging via smart contracts ensures that credentials are dynamically kept up-to-date and that lapsed requirements automatically result in loss of privilege until remedied, without manual intervention. The smart contracts also enable fine-grained delegation – for example, issuing a locum tenens doctor a temporary hospital privilege that activates only for the week they're on rotation and then auto-revokes.

Another innovative feature is a Regulatory Horizon-Scanner, which is essentially an AI (e.g., an NLP agent or Large Language Model) configured to continuously monitor emerging regulations, guidelines, and requirements from sources such as medical boards, accreditation bodies, and

government agencies (CMS, FDA, etc.). This agent scans bulletins, websites, and publications for changes in rules that could affect credentialing (for instance, a new state requirement for telehealth certification, or a new CMS rule on nurse anesthetist supervision). When it detects a relevant change, the system automatically generates new tasks or credential requirements in the platform's workflow. For example, if a state board announces that all providers must undergo training on a new opioid prescribing protocol, the system's scanner agent identifies this and creates a "to-do" credential item for all providers in that jurisdiction, perhaps even triggering issuance of a new conditional privilege that must be fulfilled (in this case, completion of the opioid training) by a deadline. This horizon-scanning ensures the credentialing office stays ahead of regulatory changes without needing to manually research them, effectively automating compliance updates. The LLM-based agent can also summarize the new requirements and suggest what credentials or documents are needed, reducing the load on human administrators.

Moreover, the system ties into operational enforcement through a Workforce Scheduler Guardrails integration. This means the credentialing platform provides APIs or plugins to hospital workforce management and scheduling systems. Before a schedule is finalized or a staff member is assigned to a shift/role, the scheduling system will automatically check via the API whether that individual's credentials and privileges are valid for the intended role at that future time. For instance, if a nurse is being scheduled for an ICU shift next month, the system will confirm that the nurse will still have an active ACLS certification and ICU privilege at that time; if not (e.g., certification would expire beforehand), the scheduling system can flag it or prevent the assignment. Similarly, for a surgeon scheduled to perform a specialized procedure, the system ensures they currently hold that privilege and have no active compliance issues. This acts as a real-time enforcement guardrail, ensuring that only properly credentialed providers are assigned to patient care duties. In particular, locum tenens workflows benefit greatly: when bringing in temporary staff, their credential verification (licenses, temp privileges, etc.) can be programmatically confirmed via this system before they are slotted into the roster, avoiding last-minute surprises. The guardrail integration essentially links administrative actions (like scheduling or granting system access) with credential status, enforcing compliance at the point of decision.

The invention also covers Device-Specific Competency Tokens. Modern medical devices (such as smart infusion pumps, ventilators, diagnostic analyzers, or robotic surgery systems) often have digital interfaces and user authentication. In this system, such devices are connected to feed usage and training data back into the credentialing platform. For example, an infusion pump might require users to log in with their ID before use; the pump's software can report how many times a user has successfully set it up and run infusions. Similarly, a blood analyzer machine can log when a lab technician completes a maintenance or calibration procedure. These logs are converted into competency tokens – essentially micro-credentials tied to specific device models or procedural skills. A token might represent that "Nurse A has performed 50 successful infusions on Pump Model X" or "Technician B is certified on Analyzer Y as of [date]". The device itself or its accompanying software could act as an issuer of a verifiable credential to the user's digital wallet, or the system can pull the log and issue the token on behalf of the facility. In either case, the integration ensures ground-truth data from equipment use is captured as proof of competency. This is especially important for high-tech or high-risk equipment where

training is required – the system can track not just that a one-time training was done, but that the provider maintains proficiency through actual usage. If a device manufacturer has a training program, their system could directly issue a credential upon completion (for example, a manufacturer of a surgical robot issues a digital certificate to a surgeon after training, which the hospital's credentialing system then incorporates as a requirement for granting privileges on that robot).

Finally, the invention implements a Zero-Knowledge Sanction Check capability to preserve privacy while verifying the absence of negative credentials. Typically, credentialing includes checking databases for any sanctions, disciplinary actions, or license revocations related to a provider. Using zero-knowledge cryptography, the system enables a proof whereby a provider (or an authoritative source) can prove to the hospital's credential system that no sanctions exist for that provider, without revealing any additional personal data. One embodiment uses a cryptographic accumulator updated with identifiers of all sanctioned providers (e.g., an accumulator of license numbers under discipline). The provider's credential wallet can generate a zero-knowledge non-membership proof that their identifier is not in that accumulator . The verifier (the credentialing system) checks this proof using the corresponding public parameters, and if valid, it is cryptographic assurance that the provider has a clean record up to date. Alternatively, a zero-knowledge proof could be constructed with a trusted third-party (like a board or the NPDB – National Practitioner Data Bank) where the third-party vouches in zero-knowledge form that "Dr. X has no adverse records as of today." The result is a privacy-preserving background check: sensitive data about other providers or details of the check are never exposed, only a pass/fail cryptographic attestation. This encourages frequent checks (even continuous monitoring) without burdening privacy or requiring full raw data exchange.

In combination, these features form a comprehensive credentialing ecosystem that is more secure, efficient, and intelligent than prior approaches. The system leverages open standards for identity (DID – Decentralized Identifiers and W3C Verifiable Credentials) to ensure interoperability and future-proofing . Data integrity and auditability are enhanced by using blockchain or distributed ledger anchoring for credential records (for example, writing a hash of each credential issuance or revocation to a consortium ledger shared by stakeholders). Identity verification of users is strengthened via protocols like FIDO2 for passwordless, phishing-resistant authentication – not only does this secure access to the credentialing system, but it enables the biometric co-signing features described . By integrating with enterprise blockchain frameworks (e.g., Hyperledger Fabric/FireFly) and identity frameworks, the system can operate as a trust network among hospitals, clinics, insurers, and regulators, where credentials and privileges are accepted across organizations with cryptographic trust instead of redundant re-verification.

Overall, the invention provides an omnibus solution for managing credentials and privileges across a provider's career lifecycle. It ring-fences various functionalities – from initial credential issuance and verification, through ongoing monitoring of skills and compliance, to automated enforcement and renewal – in one interoperable platform. Embodiments of the system can be realized as a web-based SaaS platform, an on-premise module integrated into hospital EHR

and HR systems, or a hybrid decentralized network of credential issuers and verifiers. By covering all aspects (ledger of skills, AI risk scoring, biometric attestation, smart contracts, regulatory AI, scheduling integration, device data integration, and privacy-preserving checks), the invention ensures that healthcare organizations (and other industries with similar needs) can credential and privilege professionals faster, safer, and smarter than ever before.

Brief Description of the Drawings

FIG. 1A illustrates the three-party model of decentralized identity credentials, showing how an Issuer entity issues verifiable credentials to a Holder (the professional), who can later present them to a Verifier (e.g., a hospital). This model underpins the Skills-Currency Ledger and portable credential aspect of the invention, using standards from W3C.

FIG. 1B is a schematic diagram of an example credentialing network architecture, including components for issuers, holders, and verifiers as well as supporting services (storage, status registry, etc.) to manage verifiable credentials in a distributed ecosystem. This architecture highlights possible software components (e.g., agent services and coordinators) that implement the credential issuance and verification flows.

FIG. 2 is an overview of a blockchain-enabled credentialing platform architecture, such as a Hyperledger FireFly-based system, which can be used to implement the invention. In this figure, a core node orchestrates various runtimes and plugins: a blockchain ledger for transactions (credential issuance/revocation events), an off-chain data exchange (for documents and biometric proofs), an identity registry (for DIDs and keys), and connectors to enterprise systems (EHR, scheduling software, device interfaces). The layered architecture ensures pluggability – for example, different blockchain technologies or identity providers can be swapped in.

FIG. 3 illustrates an example of biometric authentication integration using the FIDO2 WebAuthn flow. The diagram shows how a user's device (authenticator) registers and later produces an assertion to authenticate, communicating with a server. In the context of this invention, this flow is employed for Bio-Signed Procedure Attestations: both supervisor and performer use a similar flow to generate cryptographic signatures tied to their biometrics, which are then recorded with the procedure log.

(FIG. 4 and subsequent figures are reserved for additional flowcharts or use-case illustrations as needed – for example, a flow diagram of the predictive privileging AI process or the regulatory horizon-scanner pipeline. They are omitted here for brevity.)

Detailed Description of Embodiments

System Architecture Overview

An embodiment of the credentialing system can be implemented on a multi-tier architecture combining client applications, server-side logic, and distributed ledger/network components. At

the highest level, the system comprises: (1) Client Interfaces – such as a web portal and mobile app for users (credentialing staff, providers, supervisors) to interact with the system (submitting credentials, approving tasks, providing biometric auth, etc.); (2) Application Server – which houses the business logic for credential issuance, verification, rule evaluation (AI and smart contracts), and integrations; and (3) Distributed Network Components – including a ledger/blockchain network for recording credential transactions, a decentralized identity infrastructure (DID and verifiable credential registry), and possibly secure enclaves or cryptographic services for the zero-knowledge proofs.

In a web-based embodiment, the system is offered as a cloud service where the Application Server is hosted centrally. Hospitals and providers access a web dashboard to manage credentials. The central server connects to a consortium blockchain (or an internal ledger) where each credential or privilege change is anchored. External data sources (like regulatory feeds or NPDB) are accessed via API by the server. The AI model for risk prediction can be hosted in the cloud, and biometric attestation might use the user's smartphone (with a companion mobile app) for fingerprint or face recognition.

In an EHR-integrated embodiment, key functionalities of the system are embedded within a hospital's existing Electronic Health Record or credentialing software. For instance, the EHR could have a credentialing module powered by this system's software development kit (SDK). The Skills-Currency ledger updates could occur in real-time as clinicians document cases in the EHR: each procedure note signed in the EHR triggers a call to the credentialing system to issue a micro-credential token. The blockchain or ledger might be hosted on-premises or in a private network among affiliated hospitals for data locality. Biometric sign-offs could be done via hospital single-sign-on systems that support FIDO2 (e.g., a badge reader or biometric sensor integrated with the workstation login). The AI risk model could run on the hospital's analytics servers, possibly using de-identified internal data combined with national benchmarks for training.

Crucially, all embodiments utilize standardized data formats and protocols for interoperability. Each credential or micro-credential is represented in a format compliant with the W3C Verifiable Credentials Data Model, containing claims (e.g., "Provider X has skill Y level Z") and a digital signature by the issuer. The use of DIDs means each entity (practitioner, hospital, device, regulator) has a unique decentralized identifier for signing and verifying credentials. As shown in Fig. 1A, the three roles of issuer, holder, verifier can be distinct or played by the same organization in different contexts. For example, a medical board is an issuer for licenses, the provider is the holder, and a hospital is the verifier; whereas the hospital can also be an issuer for hospital-specific privileges, etc. The system's architecture (see Fig. 1B) includes issuer agents for various credential issuers (state boards, training programs, device systems), a holder wallet for each provider (which could be an app or cloud wallet where their credentials are stored under their control), and verifier services used by organizations to validate incoming credentials. By adhering to this architecture, the system ensures that credentials can be ported globally – a "digital passport" of qualifications for the provider – which aligns with the applicant's intent to file globally and cover interoperability across jurisdictions.

The blockchain ledger (or alternative tamper-evident log) in the system can be implemented with enterprise blockchain frameworks. Hyperledger Fabric or Hyperledger FireFly is a suitable choice in one embodiment, allowing permissioned sharing of data among authorized parties (hospitals, certifying bodies, etc.). The FireFly “supernode” architecture (see Fig. 2) provides an orchestration layer that combines on-chain transactions (for asset transfers or smart contract invocations) with off-chain messaging and data storage for heavier data. In this system, the Skills-Currency Ledger entries might be recorded as token transfers or events on-chain (if using a token model for skills) or simply logged in a secure database with a hash anchored to the chain. Smart contracts are deployed on the chain to represent conditional privileges; these contracts encode the logic for expiration and conditions. For example, a smart contract for “ICU Privilege for Dr. Y” would have state variables for expiry date and required CME credits, and functions that automatically revoke or renew based on inputs (like receiving a transaction indicating CME completion). Multiple smart contracts can exist per provider, or a single contract per provider with different privilege flags.

The Predictive AI module can be implemented as a microservice in the architecture. It may have access to a data warehouse that aggregates relevant data: the provider’s own history (from HR and credential files), industry data (de-identified peer benchmarks, malpractice databases), and perhaps even real-time performance indicators (like patient satisfaction scores, complication rates). A machine learning pipeline is used: data preprocessing, feature extraction (possibly using NLP on clinical notes or clustering procedure types, etc.), a prediction model (which could be a gradient boosting machine, neural network, or an ensemble), and an output interface. The output might be a risk score (0 to 100), or classification (low/medium/high risk), or predicted probability of an event (e.g., “5% chance of a malpractice claim in next 1 year”). This output is then integrated into the credentialing workflow. In practice, when a credentialing analyst or committee reviews a provider’s file for appointment or reappointment, the system displays this AI-derived insight prominently (with explanations if possible). In some embodiments, the AI is used in an automated manner: for instance, the system could automatically restrict granting of additional privileges to a provider flagged as high-risk until a human review is done, or conversely fast-track low-risk candidates.

The Bio-Signed Procedure Attestation feature requires both hardware and software integration for biometrics. In one embodiment, each supervising physician and each practitioner has a FIDO2-compliant security key or device (this could be as simple as their smartphone with biometric unlock, registered as an authenticator). When supervision is required, the trainee initiates an attestation request via the system’s mobile app or workstation. The supervisor is prompted on their device to confirm supervision; using their fingerprint or Face ID (which never leaves the device), the device’s secure element signs a challenge from the server, returning a signed assertion that can be verified by the server . The server then packages the evidence (who, what procedure, when, plus the signatures) into a verifiable credential or ledger entry. Fig. 3 illustrates a generic WebAuthn flow as used for registration and authentication – the system utilizes this standard flow so that biometric attestations are strongly bound to identities without the server ever seeing raw biometric data. In a variation of this embodiment, if hardware like a biometric scanning pad or an iris scanner is available in a procedure room, that hardware can directly interface with the system’s app to capture both parties’ biometrics on-site. The key point

is that the attestation is cryptographically signed by each party, preventing later denial. In an alternate embodiment for lower-resource settings, a simpler approach uses one-time passwords or QR codes co-signed by both parties – not as strong as biometrics, but still providing a digital signature trail (e.g., both scan a QR code that generates a signed attestation via an app).

Now, the Smart-Contract Conditional Privileges aspect will be described with an example. Suppose a hospital grants a surgeon "Robotic Surgery Privileges" on a da Vinci™ surgical system, conditional on (a) completing manufacturer training, (b) performing at least 10 robotic surgeries per year, and (c) maintaining general surgery board certification. In the traditional process, tracking these conditions is manual – the credentialing office would have to remind the surgeon of these and periodically verify compliance. With this invention, when the privilege is granted, the system deploys a smart contract (or creates a conditional credential entry) that encodes these rules. The manufacturer training completion can be automatically confirmed if the device-specific token for that training is present in the ledger (as issued by the manufacturer's training system). The volume of surgeries is tracked via the Skills-Currency ledger (which logs each robotic case); the smart contract can read those counts from an oracle or from the ledger state. Board certification status could be monitored via integration with the board's credential (perhaps the board issues a verifiable credential for certification that expires if not renewed – the system would know if it's active or not). The smart contract continuously evaluates whether all conditions hold. If, after a year, only 8 robotic cases were performed, the contract could mark the privilege as "Suspended" and trigger an alert to the surgeon and hospital. The surgeon might then need to undergo re-proctoring or additional training to restore the privilege. The contract can also handle time-bound logic (expiring after 2 years regardless, requiring renewal application which could be another on-chain transaction to extend it). This approach is like a digital rights management for clinical privileges – it ensures compliance in a transparent yet automated way. Importantly, the use of smart contracts means that if multiple hospitals share a blockchain network, a provider's privileges can be recognized across them or at least uniformly governed (subject to inter-organizational agreements). For example, a locum tenens doctor could carry a smart contract credential that multiple hospitals' systems read and enforce (perhaps a token that says "Active privileges at Hospital A until date X, automatically void if not used in 6 months").

Turning to the Regulatory Horizon-Scanner, an embodiment uses a combination of web crawling, text processing, and natural language understanding. The system may subscribe to RSS feeds of state medical boards, the Joint Commission news, CMS transmittals, etc. Additionally, a web crawler could regularly scrape known websites (like government gazettes or medical licensing websites). An NLP pipeline then classifies these text updates to detect those relevant to credentialing. For example, if a new law is passed mandating racial bias training for all physicians, the text of that law or a news update would contain keywords like "require" "all physicians" "training" etc. The system's LLM, possibly fine-tuned on regulatory texts, can parse the obligations and identify who must do what by when. The output might be a structured data: profession = physician; requirement = "BiasTrainingCourse"; due_date = Jan 1, 2026. The credentialing system, upon getting this, will create a new credential requirement entry called "Bias Training Course Completion" for all physician users, and perhaps link to an available course if known. It might also send notifications to admins that "New requirement added per

[Source]". This component essentially automates continuous compliance. It can also scan for changes in existing rules (e.g., if a state extends the license renewal period from 2 to 3 years, the system could adjust its reminders and expirations accordingly). One embodiment might leverage an LLM agent that can query a site and have a conversation-like parsing (some regulators provide Q&A or guidelines which the agent can interpret). Another simpler embodiment might rely on a regularly updated knowledge base (maintained by a service or consortium) that the system pulls updates from (like subscribing to a standards feed). In all cases, the end goal is to minimize the lag between a new rule and the credentialing office's reaction.

The Workforce Scheduler Guardrails integration is implemented via APIs or webhooks. In one scenario, the scheduling software queries the credentialing system's API with a question like: "Can provider X work role Y on date Z?" The credentialing system then evaluates if all necessary credentials for that role (and date) are in order. This includes checking: is the license valid through that date? Are privileges for that department active? Does the provider meet any specialty-specific requirements (like vaccination status if relevant, or special training if the shift is in a stroke center, etc.)? The API responds with a boolean or a detailed report of any deficiencies. The scheduling system can then act – e.g., warn the scheduler "Provider X lacks requirement Q for that assignment" or automatically replace the provider. For a more automated approach, the scheduling system might continuously feed the final schedule to the credentialing system, which then cross-validates the entire roster and flags any compliance issues. The guardrails could extend to other systems too: for example, an OR scheduling system might check surgeon credentials before allowing a surgery booking; a telehealth platform might check if a doctor is licensed in the state of the patient before connecting the session (via an API call to the credentialing system which has up-to-date multi-state license info). This integration ensures real-time compliance enforcement, essentially operationalizing the credentials data rather than keeping it siloed in an HR file.

Now focusing on Device-Specific Competency Tokens, consider two concrete embodiments:

- Infusion Pump Competency: A smart infusion pump system is configured such that each user must tap their ID card or enter a code to unlock the pump for use (user authentication to the device). The pump's internal software, integrated with the hospital network, logs each use including the user ID and any noteworthy event (e.g., completed infusion, any infusion errors). The credentialing system periodically receives or pulls these logs. When a user hits predefined milestones (e.g., 10 error-free setups of the pump, or completion of an interactive training mode on the pump's interface), the system can automatically issue a credential "Certified in Pump Model ABC". This credential may be valid for, say, 1 year, after which revalidation is needed (which can also be triggered via device usage – e.g., if the user hasn't used the device in 6 months, the credential could lapse pending a refresher). The pump manufacturer could also be involved by defining the criteria or by providing a validation signal. An alternate approach is that the pump itself, being IoT-enabled, signs a message each time a user completes a session and sends it to the ledger – thereby the device acts as an issuer of a verifiable usage

record.

- Laboratory Analyzer Training: A laboratory analyzer (say a new blood test machine) often requires technicians to be trained by the vendor. In this scenario, the vendor's training application issues a verifiable completion certificate to each tech who passes the training. Using decentralized identity, the vendor can issue this credential directly to the technician's digital wallet. The hospital's credentialing system is subscribed (with the technician's consent or via an organizational wallet) to receive a proof of that credential. Once received, the system updates the technician's profile to mark them as authorized on that analyzer. The analyzer itself might enforce that only authorized users (with that token in their wallet) can operate it – for example, the technician scans a QR code from their credential wallet at the machine to unlock it. This ensures both that training is verified and that the device is used only by trained personnel.

Finally, the Zero-Knowledge Sanction Check can be technically realized as follows. A third-party service (or a consortium of hospitals) maintains an accumulator of "bad identifiers" (e.g., sanctioned provider license numbers). This could be built using known cryptographic accumulator techniques (such as RSA accumulators or Merkle trees for a simpler approach). The credentialing system, whenever it needs to check a provider, retrieves from the provider a proof that their identifier is not in the set. In practice, the provider's wallet or account would be linked to their license number or NPI (National Provider Identifier). The system might daily update a public accumulator value and trapdoor (private key) for revocations. The provider's side (which could just be the credentialing server acting on their behalf in this implementation) computes a witness value proving non-membership and sends it with the application file. The hospital system verifies it quickly (this can be done in milliseconds even for large sets, with the right algorithms). If the proof fails, it means either the provider is on the list or something is wrong; in either case, the system would then fall back to a direct check for details. If the proof succeeds, the system can confidently skip showing or storing any sanction data because the proof assures none exist. This approach can be supplemented with selective disclosure: if a provider does have a past issue that's minor or expunged, they could present a proof that no current sanctions exist, without revealing the past resolved issue – thus balancing transparency and privacy. In another embodiment, instead of a pure cryptographic solution, a secure federated query could be used: e.g., using a secure enclave or homomorphic encryption, the system queries a national database in an encrypted manner to get a yes/no answer on sanctions. These variations achieve similar goals: ensuring that in the sharing of credential data, sensitive negative information is tightly controlled and only the necessary boolean outcome is shared.

Use Case Examples

To illustrate how all these components interact, consider the following scenario involving a locum tenens physician (a traveling doctor) onboarding at a hospital:

- Credential Wallet Import: Dr. Alice, a locum tenens physician, has a digital wallet containing her core credentials: medical license verifiable credentials from two states, board certification credential, several hospital privilege credentials from past assignments, and various Skills-Currency tokens (e.g., 500+ hours ICU experience, 50 central line insertions logged, etc.). She applies to a new hospital. Instead of filling out lengthy forms, she shares her verifiable credentials from her wallet. The hospital's system (verifier) checks the signatures and validity of each – since many are issued by trusted authorities (state boards, ABMS, prior JCAHO-accredited hospitals), much of the primary source verification is instantaneous and cryptographically assured. This dramatically reduces manual verification time (consistent with claims by digital credential networks like Axuall that direct connection to issuers cuts delays).

- Gaps Identified & Filled: The system's AI and rules engine compare Alice's credentials against the hospital's requirements for her role (say, an ICU locum). It finds she is missing a new requirement: the hospital, prompted by the Regulatory Scanner, recently requires all ICU staff to complete a COVID-19 emergency preparedness training introduced by CMS last month. The system automatically assigns this as a task to Alice, with a link to the online training. She completes it, and the training platform issues a completion credential which the system records. Meanwhile, the system's predictive privileging AI runs on Alice's profile and flags no concerns (perhaps even noting her performance metrics from prior hospital data put her in a low-risk category).

- Biometric Agreement: When Alice arrives for her first shift, the hospital requires a proctoring of two procedures to finalize her privileges (common for locums on high-risk procedures). A senior physician supervises her in a central line insertion and a ventilator management case. Using a tablet, the supervisor and Alice both use fingerprint login to attest these were successfully done under supervision. Those attestations are logged, and the smart contract governing Alice's "Full ICU Privileges" sees that the condition "2 proctored procedures" is now satisfied, thereby activating her privilege fully.

- Active Monitoring and Scheduling: As Alice is scheduled for shifts, the integrated scheduling guardrail automatically checks her credentials. One day, her state license is due to expire in 10 days. The system already alerted her of this, and she's in process of renewal. The scheduling system sees that an upcoming shift is beyond the expiry; it warns the credentialing office. Because Alice has provided proof of renewal application and the state's system (via a verifiable interim credential or API) indicates it's in progress, the hospital may choose to allow scheduling under a conditional grace period. The smart contract on her privileges could be temporarily extended based on a signed attestation from the state board of a pending renewal – all handled within the platform.

- Outcome Logging and Feedback: Over the assignment, Alice logs many procedures. Each is added to her Skills-Currency ledger, not only enriching her own professional portfolio but also feeding back data to the hospital. The hospital can analyze aggregate data (de-identified) to see, for instance, how experience correlates with outcomes among

its staff, feeding the AI model. At the end of her tenure, the hospital issues her a verifiable credential of service (like a letter of reference with embedded metrics). If anything adverse had occurred, they could also issue – or at least record – a flag in her data (which could be coded in a privacy-preserving way if needed).

This scenario shows the end-to-end cycle: onboarding expedited by credential portability, compliance assured by smart rules and AI, competency demonstrated and recorded in real-time, and portability ensured for the next engagement. All along, security and trust are maintained via cryptographic methods (signatures, blockchain, zero-knowledge proofs) so that each participant (the provider, hospital, regulators) can trust the data without heavy manual processes.

Alternative and Additional Embodiments

While the description thus far has focused on healthcare provider credentialing, the inventive system is broadly applicable to other fields that require rigorous credentialing and privileging. For example, in aviation, pilots and mechanics require certifications, logged hours, checkrides, etc. A skills ledger could log flight hours automatically from aircraft systems, and AI could predict risk of incidents based on training data. Biometric attestation could confirm that a safety check was performed by a certified mechanic (with a fingerprint scan on the aircraft's maintenance computer). Similarly, in education, teacher certifications and continuing education credits could be managed with smart contracts (auto-expiring if not kept current), and classroom performance could feed into micro-credentials (e.g., successful teaching of special-needs students logged and credited). The integration with scheduling in education (for assigning teachers) or in finance (for assigning traders with certain licenses to certain roles) follows a similar pattern as the healthcare scheduling guardrails described.

Another embodiment could incorporate Hyperledger Indy/Aries for the identity aspect, which is tailored for self-sovereign identity with privacy-preserving credential presentations (e.g., only disclosing the necessary attributes to verifiers). This could amplify the privacy features of the system – for instance, using predicate proofs to prove a value is above a threshold without revealing the exact value (useful if, say, verifying a provider has done at least X procedures without revealing the total number). This aligns with the zero-knowledge approach and can be seen as an extension of it.

Additionally, the machine learning component can leverage not just classical features but also modern approaches like graph neural networks (modeling the network of professional relationships or referrals) or large language models on unstructured text (to, for example, read de-identified patient feedback). The system's predictive insights could be explained to users in natural language, bridging the gap between raw data and actionable knowledge (e.g., "The provider's malpractice risk is in the top 10%, primarily due to a higher-than-average surgical complication rate and fewer continuing education hours . Recommendation: assign a mentor or require refresher training.").

The smart contracts for privileges could be implemented on different technologies: a public blockchain (if broadest trust is needed) or a private one. In some jurisdictions, FireFly could be used as in FIG. 2 to coordinate off-chain data (like storing the actual credential files or large documents in IPFS, with only their hash on chain). FireFly's event engine can sequence multi-party workflows, which is useful if, say, multiple signatures are needed on a credential (a collaborative credential issuance by two institutions).

The Regulatory Scanner could incorporate a collaborative element as well – multiple hospitals could share data about regulatory changes they've detected, perhaps via a decentralized knowledge graph. In a PCT context, the system could adapt to various countries' regulatory bodies (for example, scanning GMC updates in the UK, or provincial college rules in Canada, etc.), making it a globally aware system that still provides local specificity.

Overall, the present invention admits many variations in implementation while still falling within the scope of the inventive concepts. The specific technologies (blockchain type, AI model type, biometric device, etc.) can be interchanged or updated as new standards emerge (for instance, if a new FIDO3 standard comes out, it could be used similarly). The key novelty lies in the integration and interplay of all these components to create a credentialing system that is secure, automated, intelligent, and interoperable.