bondex™

When Agents Hire Humans

Verified Reputation Infrastructure for Hiring.

● PUBLISHED

01 June, 2026

● AUTHORS

Kyle Burke and Paul Martin.

A TECHNICAL WHITEPAPER ON THE BONDEX REPUTATION SCORE

ACKNOWLEDGEMENTS

A Paslar, Samuel Rowlands, Ignacio Palomera, Etienne Michel.

Summary

AI has overwhelmed the infrastructure that sits under hiring. The average job application response rate has collapsed to around 2%, not because qualified people are scarce, but because hiring managers now face a tsunami of applications and are forced to use AI to filter them. Every new job post draws more automated applications, and each wave pushes recruiters to rely more heavily on machine filtering; the arms race accelerates in both directions.

Resumes are naturally lossy and, by nature of just being words on a page, may not represent truth. This leads employer AI agents to transact over signals they cannot trust: resumes they cannot verify, endorsements with no cryptographic backing, and claims that may not originate from a human at all. The consequences are exclusionary at scale. High-signal candidates get buried beneath optimized noise, entire geographies get filtered out at the recognition step, and the market rewards prestige over proof.

Verified reputation inverts that dynamic. It elevates honesty over brand recognition, surfaces talent regardless of geography, and gives AI agents a machine-readable signal layer they can actually act on. This paper outlines the requirements any such system must satisfy, evaluates why the current industry norms fall short, and specifies an architecture that meets them. The Bondex Reputation Score is a step towards realizing this future.

1. The Signal-to-Noise Crisis in Professional Hiring

The integrity of the global labor market depends on a single assumption: that the people claiming to be qualified actually are. That assumption is breaking. When AI agents begin making hiring decisions for employers, and they will, the signals they act on will determine who gets hired, who gets seen, and who gets left behind. If those signals are fabricated, the consequences are exclusionary at scale: the most honest professionals become the least visible, and entire populations get locked out by systems that reward noise over proof.

1.1 The Collapse of Trust in Self-Reported Data

The hiring ecosystem has long operated on self-reported claims: candidates document employment history and achievements in textual resumes, employers accept them in good faith. That paradigm was already fragile. Generative AI has broken it.

Synthesizing a compelling CV, portfolio, and social proof now takes seconds. 78% of U.S. Millennials now use AI when applying for a job,[1] and 95% of executives are concerned about the accuracy of data gathered on candidates.[2]

This is a textbook instance of the "Market for Lemons":[3] information asymmetry between buyers (employers) and sellers (candidates) leads to adverse selection, and high-signal candidates get buried beneath optimized noise. Organizations report losses exceeding $50,000 per fraudulent hire when accounting for recruitment, onboarding, and productivity loss.[4]

Response rates have collapsed accordingly. The average job application response rate now sits near 2%, while submission volumes per role keep climbing. In high-throughput ecosystems the dynamic is already mature: over 1.5 million job seekers compete monthly for Web3 roles while companies post between 30 and 50 positions per day.[12] Existing matching infrastructure is functionally obsolete.

1.2 Structural Bias and Credential Portability

A high-growth startup in Eastern Europe or a leading technical university in Southeast Asia frequently lacks the reputational weight to trigger interest from North American hiring managers, not because the work is weaker but because the institution is unfamiliar. This is structural bias, not friction.

The mechanism is well-documented. People extend trust most easily to people who look like them and come from institutions they recognise.[5] Homophily systematically disadvantages high-performing profiles that fall outside the evaluator's local frame of reference. The effect multiplies at every filter: recruiters who triage a stack of resumes by recognition signals drop unfamiliar candidates first, AI screeners trained on historical hiring data replicate those same biases at higher throughput, and the outcome is a labor market in which a professional's geography silently determines how visible their work can be. Quantitative reputation systems with

verified reviews and standardised scores are among the most effective tools for breaking this pattern, enabling managers to extend trust to unfamiliar high-performing candidates.[6]

1.3 The Transition to Agent-Mediated Hiring

Median U.S. employee tenure has declined from 4.6 years in 2014 to 3.9 years in 2024.[7] As professionals accumulate denser, shorter-duration experience across a wider range of institutions and geographies, static self-reported credentials become unworkable. A resume captures what someone did once; it does not reflect who they are now. The problem doubles at exactly the moment the evaluator also changes.

Within the next few years, the dominant hiring flow will be an employer's AI agent querying a talent pool, evaluating candidates, and initiating engagement before a human is involved. On the candidate side, personal AI agents may respond and negotiate across multiple opportunities simultaneously.

In this environment, existing reputation infrastructure is not merely inefficient. It is unusable. A PDF resume is unstructured text; the agent that reads it has no way to verify whether a single line in it is true, or whether a human even authored it. A LinkedIn profile is a self-reported claim locked inside a proprietary platform. An endorsement is a social gesture with no cryptographic backing. None can be programmatically verified, and none carry the anchoring an agent needs to act without a human reviewing first.

What AI agents need is a machine-readable, cryptographically anchored signal layer where every credential is proven and every score has been evaluated without human judgment. A protocol hiring a Solidity engineer and a startup hiring a growth marketer should be able to query the same data with different weightings and get back a ranked, trustworthy shortlist in seconds, not weeks.

Building this infrastructure allows a future in which AI agents could fully mediate hiring.

2. What Verified Professional Reputation Enables

Before specifying the architecture, five outcomes are worth naming. Each is currently blocked by the absence of a machine-readable, cryptographically anchored signal standard. Each unlocks once the standard exists.

Borderless Talent Discovery. A verified professional anywhere becomes visible on the same terms as a candidate in a major hiring hub; evaluation is based on proven credentials, not familiarity. The evidence is legible regardless of where the work was done, where the institution was founded, or where the candidate happened to be born. For the first time, the global labor market becomes a single legible pool rather than a collection of regional ones that rarely see each other. The most-qualified candidate for any given role can be surfaced without the recruiter's geography doing the filtering.

Reputation that updates in real time. Today, professional reputation gets reconstructed from memory every time someone enters a job search. Verified reputation replaces this with a living signal that updates continuously as credentials are earned, contributions made, and peers attest. A professional's standing reflects what they are doing now, not what they wrote down last time they were looking.

Agent-Ready Talent Matching. AI agents receive structured, multi-dimensional, independently verifiable signals that can be queried, weighted for role-specific criteria, and evaluated without human interpretation. A protocol hiring a smart contract engineer queries the same data as a research institution hiring an economist, with different weights. The matching work collapses from weeks of recruiter review into seconds of programmatic evaluation, without sacrificing rigor, because every signal underneath can be cryptographically anchored.

Verified honesty can outperform prestige. A professional who proves their credentials may produce a stronger signal than someone who claims a prestigious background without offering proof. The system is designed to reward integrity alongside brand recognition, which matters most for professionals outside the usual hiring corridors — those whose employers and universities lack legacy name recognition but whose work stands up to verification.

Reputation belongs to the individual, not the intermediary. Reputation moves with the person across platforms, ecosystems, and geographies as a single ownable identity layer that no platform can revoke or gatekeep. When a professional's reputation is portable and cryptographically anchored, no single platform holds their career hostage.

3. Design Principles for Verified Professional Reputation

An AI-mediated labor market trusts only what it can verify. These nine principles define what any viable reputation system must satisfy.

1. **Adversarial behavior is assumed** → Any mechanism that can be exploited will be. Countermeasures against manipulation are embedded at every layer of the architecture.
2. **AI-native and agent-ready** → Credentials and reputation signals must be verifiable and actionable by AI agents, enabling programmatic discovery, evaluation, and shortlisting. When a human makes the final hiring decision, the signal they are acting on is already proven.
3. **Reputation is dynamic** → A user's standing is a living metric. Actions taken on or off-platform continuously create opportunities to gain or lose reputation, ensuring the score reflects current behavior rather than historical achievement.
4. **Trust is asymmetric** → Consistent, verified good behavior is required to build a high score, whereas a single fraudulent action results in significant and durable penalties. This asymmetry reflects the structure of real-world trust.
5. **Reputation is earned by association** → The system leverages transitive trust. Verified association with reputable institutions and peer-to-peer attestations materially amplify individual signal strength, analogous to letters of recommendation in traditional hiring.
6. **Activity decay is time-relative** → To prevent free-riding on legacy reputations, recent events are weighted more heavily than historical ones. The score is a predictive indicator of current professional capacity, not a permanent record of past achievement.
7. **Foundational credentials are evergreen** → While most signals decay, certain bedrock credentials (such as a verified undergraduate degree) maintain a permanent floor weight, so durable qualifications retain proportionate value indefinitely.
8. **Relative scoring takes precedence over absolute scoring** → The BRS represents a user's standing relative to the entire ecosystem. A score of 100 reflects the current best-in-class signal within the network, not a theoretical ceiling.
9. **Scoring is transparent, not exploitable** → The general framework is publicly documented so users can understand and improve their standing. Precise algorithmic weightings and scoring formulas are intentionally withheld, functioning analogously to a credit score: categories and mechanisms are known, the exact formula is not.

Principle Evaluation Matrix: Existing Solutions vs. BRS Design Standards

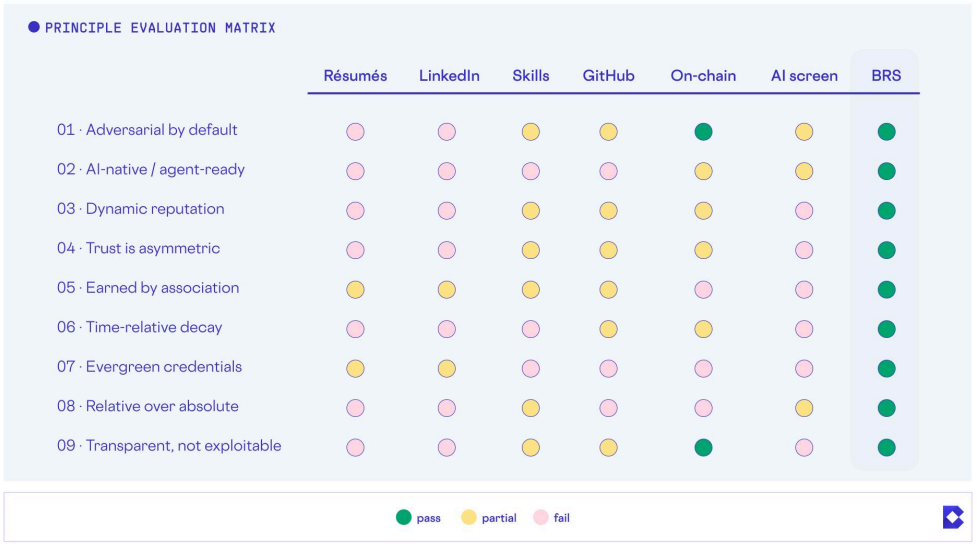

No existing solution satisfies all nine principles. Resumes and LinkedIn fail on every objective dimension: static, self-reported, trivially gameable. Skills-test platforms and GitHub profiles produce genuine signal in narrow domains but lack breadth, don't decay appropriately, and aren't structurally resistant to adversarial behavior. On-chain activity is tamper-resistant but applies only to Web3 and captures only a thin slice of professional identity. AI screening tools filter at speed against the same unverified inputs. The BRS addresses all nine dimensions simultaneously

4. Algorithmic Architecture

4.1 Overview

The Bondex Reputation Score aggregates verified evidence from five proof categories into a single normalized value. Unlike traditional professional profiles, which are updated episodically during job searches, BRS scores are computed and re-ranked continuously. A new OAuth-authenticated skill platform connection, a verified employment transition, a governance vote on-chain, or a peer attestation each updates the professional signal in near real-time.

Each claim is a hypothesis that accumulates evidential support through multiple independent verification pathways. No single verification is presumed to authenticate a complete claim; the verification coefficient increases as additional independent evidence sources corroborate the underlying assertion.

4.2 The Five Proof Categories & Their Default Weight

The BRS is structured around five categories of verifiable professional proof. Each represents a distinct dimension of professional identity: not what a candidate claims about themselves, but what they can demonstrate through evidence.

Each category is independently scored on a [0, 1] interval (a vector representing the strength of signal in that dimension) and weighted by its empirically assessed contribution to professional signal quality. The default distribution is in Table 1. Precise weighting methodology and inter-category scoring formulas are withheld per Principle 9.

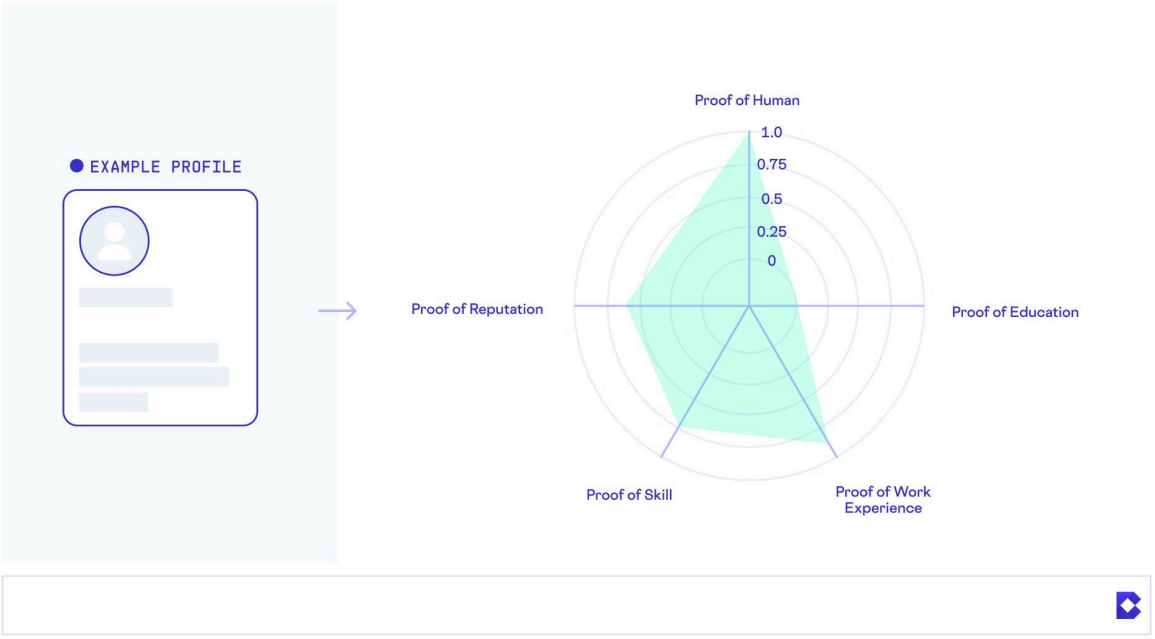

● TABLE 1 · BRS PROOF CATEGORIES, DEFAULT WEIGHTS, AND SCOPE

CATEGORY	DEFAULT WEIGHT	WHAT IT MEASURES
Proof of Human	0.10	Confirmation that the profile belongs to a real, unique individual rather than an AI agent, bot, or duplicate account.
Proof of Education	0.30	Verified evidence of formal academic credentials and institutional affiliations.
Proof of Work Experience	0.30	Verified evidence of professional employment history, organizational affiliations, and role-level claims.
Proof of Skill	0.20	Demonstrated competency through authenticated third-party platform activity: the verifiable receipts of professional output.
Proof of Reputation	0.10	The accumulated trust signal from peer attestations, behavioral integrity, and the absence of fraudulent submissions.

Default weights represent a generalist hiring signal. Operators and platform integrators can configure custom distributions for role-specific evaluation. A protocol hiring smart contract engineers may elevate Proof of Skill significantly; a research institution may heavily weight Proof of Education. Proof of Human weight may be adjusted toward zero in AI-agent evaluation contexts, anticipating a future in which autonomous agents legitimately hold professional profiles within the Bondex ecosystem (Principle 2).

4.3 Proof of Human: Identity Verification

Proof of Human is a measure of identity authenticity; confirmation that a profile represents a real, unique individual. In an ecosystem where synthetic profiles, bot accounts, and AI-generated personas are increasingly plausible actors, the ability to confirm human presence carries meaningful signal weight.

The Proof of Human score is determined by the highest-tier verification a user has completed. Verification tiers progress from basic email confirmation to cryptographic biometric uniqueness proof:

● TABLE 2 · PROOF OF HUMAN VERIFICATION TIERS

VERIFICATION METHOD	APPROX. SCORE	DESCRIPTION
Biometric uniqueness proof	1.0	Cryptographic confirmation of unique human identity. The gold standard, proving personhood without disclosing identifying information (e.g., World ID).
Institutional KYC	~0.7	Government-grade identity verification through an accredited KYC or identity provider (e.g., Didit, or comparable service).
Government document check	~0.5	Document-based identity verification without biometric uniqueness proof.
Verified institutional email	~0.3	Domain-verified email confirming organizational affiliation. Confirms association; does not authenticate the specific claim.
Unverified	0.0	No identity verification completed.

Biometric uniqueness proof, which confirms that a profile is bound to a single unique human without disclosing identifying information, represents the gold standard for human verification and returns a coefficient of 1.0. Institutional KYC through accredited providers (such as Didit or comparable services) provides a strong but non-biometric alternative. Domain-verified institutional email confirms affiliation but does not establish unique human identity.

This category's configurable weight is a deliberate architectural choice: operators who prioritize human-verified talent may increase this weight for roles where identity certainty is critical. Those evaluating talent in capabilities-first or autonomous-agent contexts may reduce it accordingly.

4.4 Proof of Education: Credential Authentication

Proof of Education captures verified evidence of formal academic credentials, assessing each claim component independently. A domain-verified email confirms institutional affiliation, not degree details or status; those require separate verification. The coefficient increases as independent evidence (email verification, institutional API confirmation, peer attestation) accumulates, requiring candidates to build evidential depth rather than completing one step.

Within-category scores aggregate via additive diminishing returns: each additional verified credential helps, but less than the previous one. Credentials in different fields or at different levels add proportionally greater weight, signaling interdisciplinary breadth. Adding any legitimate credential can never reduce the score. Credential scores decay over time, with academic degrees maintaining higher value and a longer half-life than professional certifications (Section 4.10).

4.5 Proof of Work Experience: Employment Corroboration

Proof of Work Experience verifies professional history using a similar multi-layer model. Initial verification, typically an occupational email, confirms affiliation but not role specifics or tenure; those require independent confirmation. Cross-category corroboration is vital: on-chain activity, developer contributions, or authenticated third-party data during the employment period add weight even if the primary email verification fails, for example when a domain expires after the employer ceases operations.

Within-category aggregation follows the Proof of Education model: each verified entry boosts the score with diminishing returns. The score rewards a consistent, broad career history and includes a quality floor to limit the impact of very brief or junior roles. Decay rates are faster than for academic credentials, reflecting rapid industry change. The core principle holds throughout: verified honesty beats unverified prestige. A fully corroborated history at a lesser-known company may signal more strongly than an unverified claim at a famous one.

4.6 Proof of Skill: Demonstrated Competency

Proof of Skill collects authenticated evidence of professional output across a broad range of skill domains. Rather than inferring competency from activity metrics alone, the system requires platform-authenticated evidence. A candidate who states "I am a Solidity developer" is making a claim. A candidate who authenticates a GitHub account via OAuth thereby exposing a history of Solidity repository contributions, merged pull requests, and recognized open-source participation, is providing receipts. The former is self-report. The latter is evidence.

The framework operates on a Claim → Evidence → Score logic:

- **Claim:** The candidate asserts competency in a professional domain (software development, content creation, community management, financial product development, research output).
- **Evidence:** The candidate authenticates one or more third-party platform accounts via OAuth or equivalent secure data-sharing protocols. Platform data is obtained directly from the source without the candidate as intermediary, carrying a high verification coefficient.

- **Score:** Platform data is normalized within its domain to a [0, 1] interval relative to the Bondex population of users with that platform connected, then aggregated across domains using a diminishing returns model. The first skill domain provides full weight; each additional domain contributes incrementally less, preventing platform-stacking from producing disproportionate returns.

Skill domains currently integrated or in development span:

- **Developer and Technical:** Version-controlled code repositories, competitive programming platforms, and technical contribution records.
- **Design and Creative:** Portfolio platforms and professional design community profiles.
- **Writing and Content Creation:** Publishing and audience-building platforms with verifiable production and engagement metrics.
- **Social and Audience Building:** Platforms with verifiable engagement and audience metrics, filtered for authentic engagement over manufactured following.
- **Business and Revenue:** Revenue platform integrations providing verifiable evidence of commercial activity, such as payment processing records and published application revenue data.
- **Analytics and Growth:** Marketing and analytics platforms demonstrating verifiable audience development and traffic generation.
- **Community and Open Source:** Community platform roles, open-source maintainership records, on-chain governance participation, and DAO contribution histories.
- **Freelance and Client Work:** Marketplace platforms with verified client ratings, project completion rates, and earnings records.
- **Research and Academic Output:** Scholarly indexing and citation databases with verifiable publication and citation records.
- **On-Chain and Web3 Activity:** Protocol interaction histories, governance participation records, DeFi engagement, and on-chain deployment records authenticated directly from immutable ledger data. This domain explicitly recognizes Web3 professional readiness as a distinct and verifiable skill signal.

The on-chain domain explicitly treats Web3 professional readiness as a distinct and verifiable skill signal.

The primary challenge is normalizing heterogeneous inputs. GitHub commit frequency, Stripe revenue, and Google Analytics traffic are measured in different units and reflect different professional contexts. The BRS uses population-relative normalization: each metric is scored relative to the distribution of Bondex users with that specific platform connected, so a strong GitHub profile and a strong Stripe revenue profile represent equivalent levels of demonstrated output relative to peers. The architecture is intentionally incremental, expanding the set of authenticated integrations per skill domain as normalization benchmarks mature across a growing population.

4.7 Proof of Reputation: Social Standing and Network Quality

Proof of Reputation is a composite signal drawn from three sources: automated social graph analysis, peer attestations from verified professionals, and institutional endorsements from employers or academic bodies. This category measures the quality of a professional's standing within their field as perceived and attested to by others. Questions of behavioral integrity and

honest dealing are handled by a dedicated post-computation layer, the Trust Multiplier, described in Section 4.8.

- **Social Graph Analysis:** An automated, passive signal that evaluates network characteristics rather than network size: the verified BRS standing of first-degree connections, the density of cross-platform relationship overlaps, and the consistency of network topology with the candidate's stated professional trajectory. A network heavily populated by verified, high-BRS professionals in the candidate's declared field is a stronger signal than a large but unverified or diffuse connection graph.
- **Peer Attestations:** A direct signal from other verified BRS profiles vouching that a candidate's professional conduct or demonstrated skill is high quality or factual. Attestations carry weight proportional to the attesting profile's own BRS standing, and the total attestation bonus is bounded to prevent reciprocal attestation rings from artificially inflating scores.
- **Institutional Endorsements:** Formally verified organizations (employers, universities, professional bodies whose identity has been authenticated within the BRS ecosystem) can endorse current or former affiliates. Unlike peer attestations, institutional endorsements carry the authority of the verified organization, independent of the endorsing individual's BRS. A verified former employer's endorsement provides a strong corroborating signal, calibrated by the institution's standing within the network.

4.8 The Trust Multiplier: Integrity as a Score Modifier

The five proof categories measure what a professional has done. The Trust Multiplier addresses a different question: *given what we know about this person's honesty and professional conduct, how much should we trust those scores?* It is a post-computation layer that modulates the weighted category aggregate before final ranking, ensuring that a high BRS cannot be sustained by someone whose integrity has been materially called into question.

The architecture mirrors the staking and slashing model of proof-of-stake blockchain protocols. Consistent compliant behavior earns standing over time; a single act of serious misconduct results in an asymmetric, durable penalty.[9] The economic logic is identical: the expected cost of dishonesty must exceed the expected benefit to make honest participation the dominant strategy.

Two input streams can negatively impact the Trust Multiplier:

- **Submission Integrity:** The BRS maintains a submission log for all profile claims. When a newly submitted claim differs materially from what is subsequently verified or previously claimed, the discrepancy is flagged and assessed as a slash event. This mechanism penalizes users who submit inflated or fabricated profiles expecting to correct them once they realize verification is required.
- **External Conduct Events:** Independently-sourced information about serious professional misconduct that falls outside what a candidate has submitted: a fraud conviction, a regulatory sanction from a named supervisory authority, a documented association with a project subsequently confirmed to be a scam. A professional can submit a fully truthful profile and still have a public record of conduct that is material to their professional trustworthiness.

4.8.1 Slash Events, Recovery, and the Permanent Scar

The Trust Multiplier begins at 1.0 for all profiles. Each recorded slash event (submission integrity failure or external conduct event) reduces the base value and potentially lowers the maximum achievable multiplier. The magnitude scales with severity:

TABLE 3 · TRUST MULTIPLIER: SLASH EVENT SEVERITY TIERS

DISCREPANCY SEVERITY	EXAMPLE	IMPACT ON MULTIPLIER
Minor	Date correction, role title variation within reasonable margin.	Small recoverable penalty. Full recovery achievable through continued verified positive behavior.
Significant	Material role title inflation, organizational scale misrepresentation, tenure extension.	Moderate penalty. Maximum achievable Proof of Reputation score is temporarily capped; recovery trajectory is extended.
Major	Fabricated employer, fabricated degree, identity-level fraud.	Substantial and durable reduction with a persistent floor below 1.0. Profile is flagged; a material integrity event is permanently recorded.

Once a slash event is recorded, the Trust Multiplier recovers over time through a linear function of days elapsed since the most recent event and continued absence of further violations:

```
TrustMultiplier = min(MaxPossible, BaseValue + RecoveryRate × DaysSinceLastEvent)
```

The MaxPossible ceiling is set at the time of the slash and reflects the severity of the event. For Minor events, MaxPossible sits close to 1.0 and full effective recovery is achievable with time. For Significant events, the ceiling is lower and recovery is extended. For Major events (fabricated identity, fraud conviction, scam association), MaxPossible sits at a persistent sub-1.0 value that does not approach 1.0 regardless of time elapsed.

Wounds heal, but they leave a scar. A professional who committed major fraud ten years ago and has since maintained a clean record recovers substantially, but the Trust Multiplier never returns to 1.0. That signal is permanent. Future protocol versions may introduce a peer vouching mechanism whereby verified high-BRS professionals formally attest to rehabilitated integrity, providing a social recovery pathway alongside the temporal one.[10]

4.8.2 Application to the Final Score

The Trust Multiplier is applied after the weighted category aggregate is computed and before ecosystem-relative normalization. For a candidate whose weighted category sum yields an effective score of 80 out of 100, a Trust Multiplier of 0.75 (reflecting a Significant event with partial recovery) produces a trust-adjusted score of 60. The five proof category scores remain what they are; the multiplier reflects reduced confidence in the overall profile's reliability. The formal computation follows in Section 5.1.

4.9 Multi-Layer Verification and the Coefficient

Consider two candidates claiming the same role at the same company. The first has connected a corporate email domain. The second has the same email domain plus an OAuth-authenticated GitHub account showing contribution history to the company's repositories during that period, plus an on-chain record of governance votes from a company-affiliated wallet. Both assert the same fact. Only the second has substantiated it through multiple independent sources. The verification coefficient quantifies this difference.

The BRS coefficient increases continuously as independent evidence sources are added, reflecting the information-theoretic principle that each independent pathway reduces epistemic uncertainty by a measurable, non-overlapping amount. Confidence accumulates as independent sources converge. The coefficient for any given claim reflects three properties:

- **Strength:** The evidential value of the verification method. OAuth-authenticated third-party platform data originates directly from the source without the candidate as intermediary, a substantially stronger signal than a self-reported claim or a domain-verified email.
- **Freshness:** How recently the verification was performed. A domain check completed two years ago for an organization the candidate no longer works for carries less weight than one done in the past six months.
- **Source Independence:** Whether evidence pathways are genuinely independent. Two confirmations from the same organizational domain do not constitute two independent verifications; they originate from a single source and provide correlated rather than additive evidence. True independence (confirmation through an unrelated third party or a different verification modality) carries substantially more evidential weight.

The overall coefficient moves from a base value for self-reported claims through progressively higher tiers as independent evidence accumulates: domain verification, platform OAuth authentication, institutional API confirmation, and multi-source cryptographic attestation. Precise coefficient values at each tier are not disclosed (Principle 9); the framework and its directional logic are. The practical outcome: candidates who substantiate claims through multiple independent pathways carry a meaningfully higher coefficient, making verifiable honesty the dominant strategy.

4.10 Credential Decay: The Half-Life Model

The BRS applies time-relative decay to all credential and activity scores (Principle 6). Rather than applying uniform exponential decay across every credential type, which would unfairly erode genuinely durable qualifications, each credential type is assigned parameters that reflect how quickly its relevance diminishes in practice:

- **Floor:** The minimum permanent weight a credential retains regardless of age. Foundational academic credentials maintain a higher floor. Short-cycle technical certifications have a lower floor, reflecting dependence on current technology.
- **Grace period:** The period during which a newly earned credential retains full weight before decay begins.
- **Half-life:** The time after which the credential's decayable value reduces by half. Rapidly-evolving technical certifications have short half-lives; doctoral degrees and professional licenses have long ones.

A doctoral degree from fifteen years ago retains significantly more value than a technical certification from the same period. A cutting-edge developer certification earned two years ago decays faster than an undergraduate computer science degree of the same age. A continuously active professional scores meaningfully higher than one coasting on a strong credential history from a decade prior. Precise parameters are calibrated internally and not publicly disclosed, consistent with Principle 9.

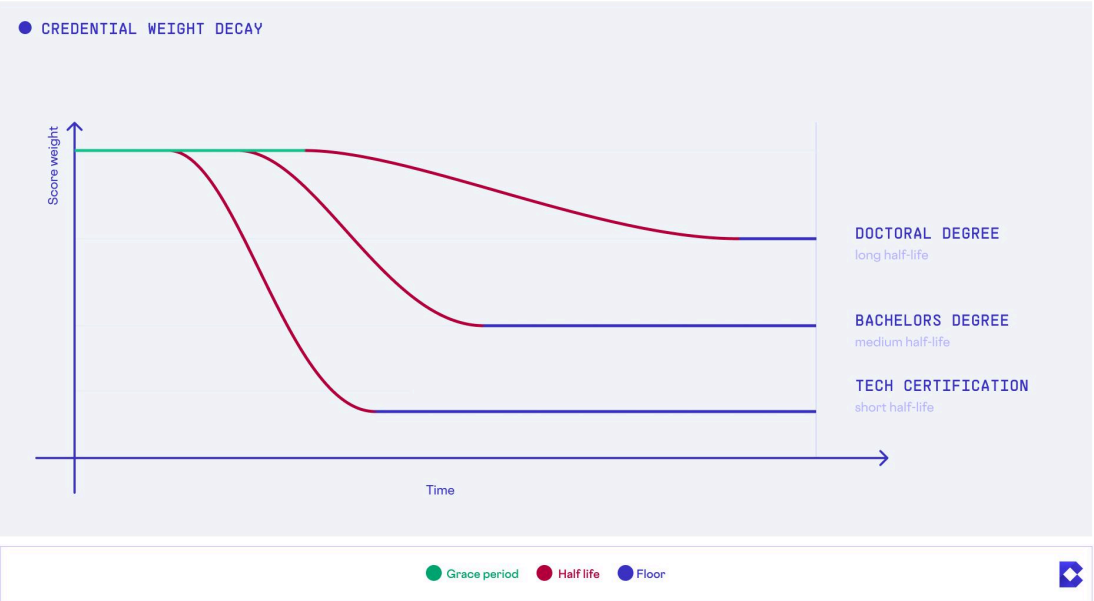

4.11 Category-Level Aggregation and Missing-Category Redistribution

Instance scores within each category aggregate into a single categorical score bounded on [0, 1]. The methodology is consistent across categories: additive diminishing returns ensure each additional verified instance contributes positively with decreasing marginal effect. A critical property is monotonicity: adding any legitimate credential can never reduce the score. A candidate who adds an early-career role, a second degree, or a third skill platform never sees their score decrease as a result.

Where a profile lacks a category entirely, for example, a professional who has not yet connected any Proof of Skill platforms, the missing category's default weight is redistributed proportionally across all active categories, so a full [0, 100] score remains achievable for profiles that do not span all domains. A professional with verified academic credentials and a strong work history but no third-party skill connections is not penalized; the Proof of Skill weight is redistributed. The formula:

$$\text{AdjustedCategoryWeight}_i = \text{CategoryWeight}_i \,/\, \Sigma[\text{AllActiveCategoryWeights}]$$

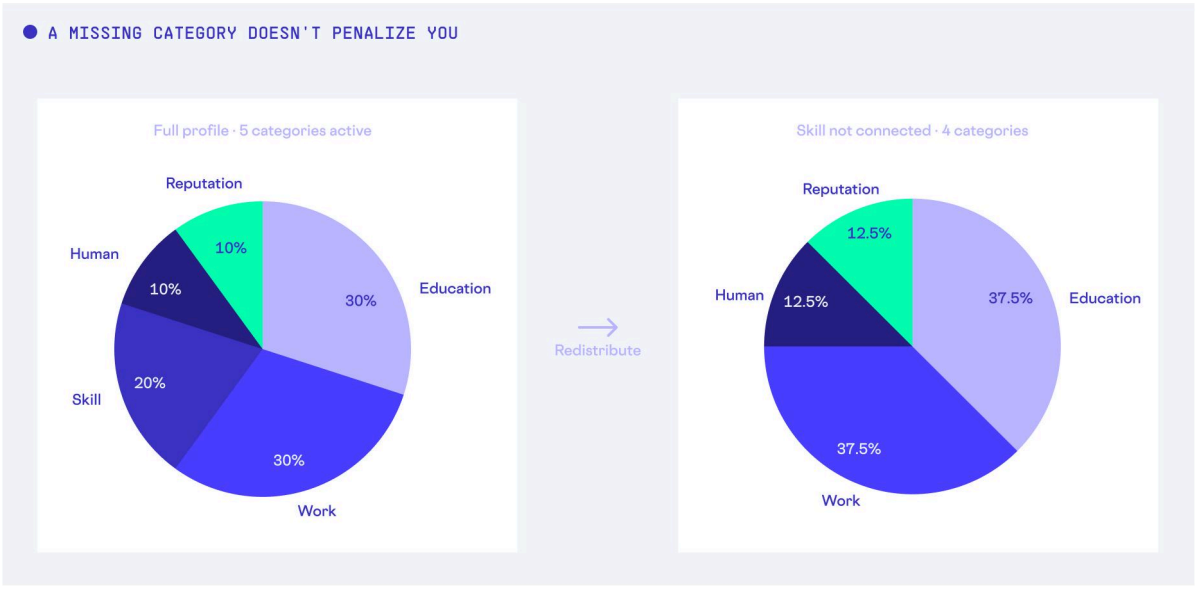

5. Final Score Computation

5.1 Weighted Aggregation and Trust Adjustment

The final BRS is computed in three steps. First, the weighted sum of all active category scores is computed and scaled to a [0, 100] interval. Second, the Trust Multiplier is applied to produce the trust-adjusted score. Third, the trust-adjusted score is used as the input to ecosystem ranking:

$$\text{RawScore} = \Sigma[\text{CategoryWeight}_i \times \text{CategoryScore}_i] \times 100$$

$$\text{FinalScore} = \text{TrustMultiplier} \times \text{RawScore}$$

CategoryScore_i is the normalized [0, 1] score for each of the five proof categories. CategoryWeight_i is the corresponding weight, summing to 1.0 across all active categories. TrustMultiplier is the post-computation integrity modifier from Section 4.8. For profiles with no recorded slash events, TrustMultiplier = 1.0 and FinalScore equals RawScore. For profiles with recorded integrity events, the TrustMultiplier reduces the effective score proportionally to severity and recency, ensuring that a strong credential profile cannot insulate a candidate whose honesty has been materially called into question.

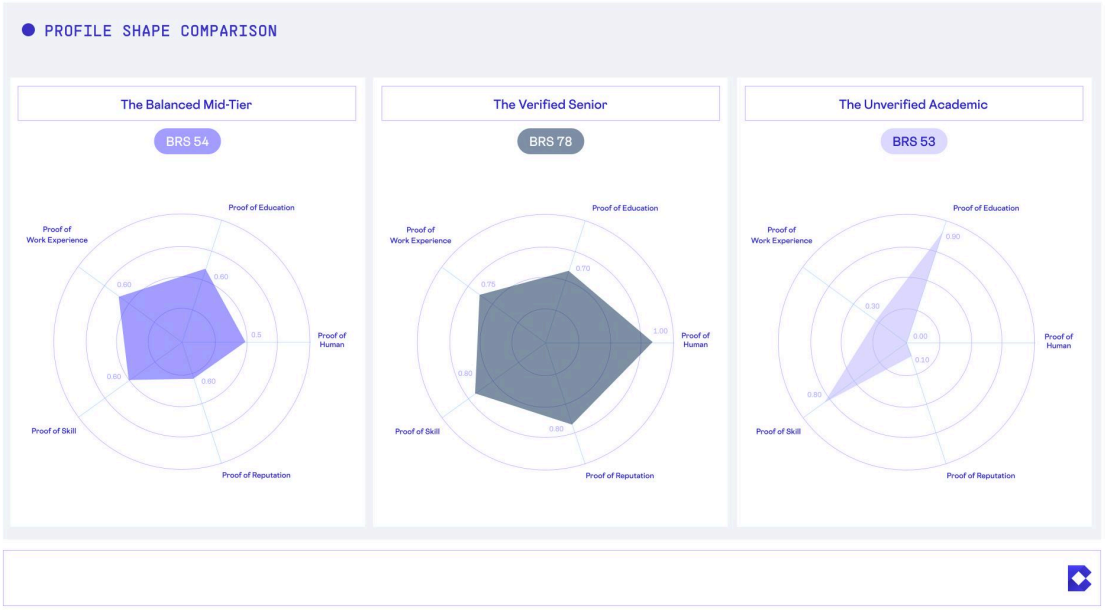

5.2 Ecosystem-Relative Ranking

The absolute FinalScore is then normalized against the distribution of scores across the entire Bondex user base, so the BRS functions as a relative percentile indicator reflecting where a professional stands within the contemporary talent ecosystem:

```
RankedScore = [[FinalScore - GlobalMin] / [GlobalMax - GlobalMin]] × 100
```

A score of 100 denotes the current best-in-class verified signal in the network. As the global score distribution evolves with platform growth, new data integrations, and shifts in user behavior, the normalization anchors are recalibrated (Principle 8). With the verified population growing, the distribution remains meaningful and competitive, rewarding professionals who invest in verification.

The score displayed to a candidate on their own profile reflects their precise position. When profiles are surfaced to employers in a talent pool context, scores may be presented within a calibrated range rather than as a single exact figure, preserving signal quality for hiring operators while maintaining the opacity needed to prevent adversarial formula extraction (Principle 9). The five proof categories are independently queryable by employer AI agents, enabling programmatic evaluation against role-specific criteria without human review of unstructured profile text (Principle 2).

6. Future Directions and System Evolution

The research program around BRS extends across two horizons: deeper integration with the agent economy, and ongoing improvements to verification integrity.

Agentic web and machine-readable signal infrastructure

- **x402 protocol integration.** The x402 protocol, an open internet-native payment standard developed by Coinbase and the Linux Foundation that revives the HTTP 402 "Payment Required" status code, enables autonomous on-chain stablecoin payments for API access without accounts or sessions.[11] Integrating x402 into the BRS query infrastructure lets employer AI agents pay for verified signal access on a per-query basis, aligning the economics of talent discovery with the agentic web.
- **AI-native talent discovery.** As employer-side AI agents mediate more of candidate evaluation, the BRS serves as the primary machine-readable signal layer for autonomous hiring workflows. The five proof categories can be queried, filtered, and weighted programmatically, removing the human-review bottleneck that defines high-volume hiring.
- **Zero-knowledge proof infrastructure.** ZK systems let candidates prove a claim meets a specified threshold without revealing the underlying data, enabling privacy-preserving talent matching aligned with the W3C Verifiable Credentials standard.
- **Machine learning-based categorical scoring.** Proof of Skill domain scoring and Proof of Reputation integrity assessment are well-suited to supervised models trained on post-hire performance metrics and dispute resolution outcomes.
- **Operator-configurable professional archetypes.** Pre-configured weight profiles for common archetypes (engineering, marketing, founder, research) let operators surface role-optimized candidates without manual weight configuration.

Verification integrity and coverage

- **Expanded Proof of Skill integrations.** The system begins with a limited set of authenticated platforms per domain and expands as normalization benchmarks mature.
- **Institutional credential verification APIs.** Direct integration with National Student Clearinghouse or equivalent national databases lets the system confirm with institutions whether a specific degree was awarded to a given individual in a given year.
- **Credential mill resistance.** Type-specific ceilings on base values ensure short-cycle online certificates from prestige institutions are scored appropriately relative to full degree programs.
- **Ghost employer domain defense.** Mitigating the expired-domain attack via cross-referencing registration dates against claimed employment dates and requiring corroboration across independent addresses at the same domain.
- **Career break accommodation.** Differentiated decay rates for documented parental leave, medical leave, and sabbatical periods so professionals are not penalized for time away.
- **Peer vouching for trust recovery.** Verified high-BRS professionals attesting to rehabilitated integrity, introducing a social rehabilitation pathway alongside the temporal one.

- **Corroboration quality function.** Weighting peer attestations by the attesting profile's BRS, penalizing reciprocal attestation patterns, and applying cycle detection to prevent multi-hop fraud.
- **Negative verification.** Claims affirmatively disproven through dispute resolution yield a negative verification coefficient, actively reducing the category score and potentially triggering a Proof of Reputation or Trust Multiplier penalty.
- **Sybil resistance enhancement.** Account activity patterns, temporal synchronization across actions, and network graph structure incorporated to detect likely Sybil or duplicate profiles.
- **Challenger staking and decentralized dispute resolution.** A token-based dispute economy in which verified users stake against specific claims, with successful challenges rewarded and failed challenges penalized.

7. Applications Beyond Hiring

Hiring is the first application of verified professional reputation, but the primitive the BRS defines (portable, cryptographically anchored, multi-dimensional) could extend wherever trust currently runs on unverified claims or high-friction paperwork. The applications below are directional: potential future extensions of the same primitive, not commitments of what the BRS does today.

- **Verified talent infrastructure for global migration.** More than 300 million people live outside their country of birth,[13] flowing through H-1B, EU Blue Card, UK Skilled Worker, Canada Express Entry, and equivalents on a paper-based, months-long, fraud-prone verification layer. A future BRS could collapse that verification from months to a single cryptographic query.
- **Credit scores for DeFi.** DeFi lending crossed $55 billion in total value locked in mid-2025[14] but remains stuck at typically 130–200% over-collateralization[15] because protocols cannot verify their borrowers. The $6.4 trillion global unsecured loan market[16] excludes the billion-plus remote, non-W2 workers traditional credit bureaus cannot reach. With sufficient adoption, BRS could become a credit rail for a workforce that earns everywhere and lives nowhere.
- **Insurance for a workforce that is structurally uninsurable.** More than 400 million online gig workers operate globally,[17] with fewer than half covered by basic health insurance and under one in five covered for employment injury, disability, or pension.[18] Against a USD 9 trillion global insurance protection gap,[19] a future BRS could unlock entire product categories that do not yet exist.
- **A potential replacement for the global background-check industry.** The global employment-screening market sits at $7.2 billion in 2025 and is projected to nearly double to $13.7 billion by 2034,[20] still running on manual phone, paper, and fax verification. A future BRS could collapse every check into one cryptographic query and expand the category by making verification live and continuous rather than a one-time snapshot at hire.
- **A portable license standard for the world's regulated professions.** Tens of millions of licensed professionals (doctors, lawyers, engineers, accountants, nurses) re-credential from scratch each time they cross borders.[21] BRS could become the license-passport standard that compresses cross-border mobility from months to minutes.
- **Identity and reputation rails for the AI agent economy.** The agentic economy has no credit bureau. McKinsey projects USD 3–5 trillion in AI-agent-orchestrated consumer commerce by 2030.[22] Every agent transacting with another agent will need an identity it can prove and a reputation it can be judged on. BRS could become that layer.

8. Addressing Common Concerns

A system that assigns portable, verifiable scores to professionals raises predictable objections. The most common are addressed below.

This is a social credit score. No. Social credit systems are government-administered, involuntary, and evaluate personal and civic behavior, including political activity. The BRS is voluntary, professional-context-only, and self-sovereign. Nothing enters the profile without the user's explicit action, and only verified professional signals are scored.

Privacy and GDPR. Privacy is a design constraint, not an afterthought. Cryptographic email verification confirms institutional affiliation without retaining the email address itself. OAuth-authenticated platform connections are permissioned by the user and scoped to the data fields required for scoring. Zero-knowledge proof infrastructure, in active development, will allow claims to be proven without the underlying personal data being disclosed. The BRS is designed for GDPR compliance, and users retain the right to request deletion of their profile data.

Who controls the institutional rankings? Bondex maintains none. Where institutional reputation contributes to a score (as with university affiliation), the system draws on publicly available, independently governed databases such as QS World University Rankings and Times Higher Education. The verification coefficient ensures institutional ranking is only one input: a verified credential from a lower-ranked institution produces a stronger signal than an unverified credential from a higher-ranked one.

What happens when the algorithm is wrong? No scoring system is infallible. A formal dispute process lets users contest flagged discrepancies. Cross-category signal reinforcement ensures a single mis-scored category cannot dominate the final result. Operators are advised to treat the BRS as a high-quality input to a hiring decision, not a sole criterion.

Collusion, gaming, and adversarial behavior. The system is designed on the premise that adversarial behavior is inevitable (Principle 1). Countermeasures are embedded at every layer: the Trust Multiplier (Section 4.8) applies a post-computation penalty when integrity events are recorded; the verification coefficient rewards multi-source corroboration; cross-category analysis flags inconsistencies; and peer-attestation weighting defeats reciprocal attestation rings.

9. References

[1]: Statista+. (2026, February 10). *Is AI the Future of Job Applications? / How and Why Americans Are Using AI to Apply for Jobs*. Statista. https://www.statista.com/chart/35806/use-ai-in-job-application/
[2]: Deloitte. (2026, March 3). *Managing Disinformation at Scale — 2026 Global Human Capital Trends*. Deloitte Insights. https://deloitte.com/us/en/insights/topics/talent/human-capital-trends/2026/managing-disinformation-at-scale.html
[3]: Akerlof, G. A. (1970). The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500.
[4]: CareerBuilder. (2017). *The high cost of a bad hire*. CareerBuilder Research Report. https://www.careerbuilder.com
[5]: McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1), 415–444.
[6]: Pallais, A., & Sands, E. G. (2016). Why the referential treatment? Evidence from field experiments on referrals. *Journal of Political Economy*, 124(6), 1793–1828.
[7]: U.S. Bureau of Labor Statistics. (2024). *Employee tenure in 2024*. News Release USDL-24-1810. https://www.bls.gov/news.release/archives/tenure_09262024.htm
[9]: Buterin, V., et al. (2020). *Ethereum proof-of-stake: The beacon chain*. Ethereum Foundation. https://ethereum.org/en/upgrades/beacon-chain/
[10]: Resnick, P., Kuwabara, K., Zeckhauser, R., & Friedman, E. (2000). Reputation systems. *Communications of the ACM*, 43(12), 45–48.
[11]: Coinbase & Linux Foundation. (2025). *x402: The HTTP Payment Protocol*. https://x402.org
[12]: web3.career. (2026). Job seeker and employer traffic data. https://web3.career
[13]: United Nations Department of Economic and Social Affairs, Population Division. (2025, January). *International Migrant Stock 2024: Key Facts and Figures*. https://www.un.org/development/desa/pd/sites/www.un.org.development.desa.pd/files/undesa_pd_2025_intlmigstock_2024_key_facts_and_figures_advance-unedited.pdf
[14]: The Block (citing DeFiLlama). (2025, June). *DeFi lending hits record $55 billion TVL as Aave, Maple, and Morpho lead the charge*. https://www.theblock.co/post/358368/defi-lending-hits-record-55-billion-tvl-as-aave-maple-and-morpho-lead-the-charge
[15]: Cornelli, G., & Gambacorta, L. (2024, June). *Why DeFi lending? Evidence from Aave V2*. BIS Working Papers No. 1183. Bank for International Settlements. https://www.bis.org/publ/work1183.pdf
[16]: TechSci Research. (2025). *Unsecured Loan Market – Global Industry Size, Share, Trends, Opportunity, and Forecast, 2026–2031*. https://www.techsciresearch.com/report/unsecured-loan-market/14544.html
[17]: Datta, N., Rong, C., Singh, S., et al. (2023, September). *Working Without Borders: The Promise and Peril of Online Gig Work*. World Bank Group. https://documents1.worldbank.org/curated/en/099071923113511279/pdf/P17730205fbe2002709043043e4d4f7efee.pdf
[18]: International Labour Organization. (2021, February). *World Employment and Social Outlook 2021: The Role of Digital Labour Platforms in Transforming the World of Work*. ILO. https://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_771749.pdf
[19]: MAPFRE Economics. (2025). *The Global Insurance Potential Index 2025*. MAPFRE. https://www.mapfre.com/en/insights/economy/the-global-insurance-potential-index-2025/
[20]: Fortune Business Insights. (2025). *Employment Screening Services Market Size, Share & Industry Analysis, 2026–2034* (Report ID: FBI115476). https://www.fortunebusinessinsights.com/employment-screening-services-market-115476
[21]: Boniol, M., Kunjumen, T., Nair, T. S., Siyam, A., Campbell, J., & Diallo, K. (2022). The global health workforce stock and distribution in 2020 and 2030: A threat to equity and 'universal' health coverage? *BMJ Global Health*, 7(6). https://pmc.ncbi.nlm.nih.gov/articles/PMC9237893/
[22]: McKinsey & Company (QuantumBlack). (2025, October). *The Agentic Commerce Opportunity: How AI Agents Are Ushering in a New Era for Consumers and Merchants*. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-commerce-opportunity-how-ai-agents-are-ushering-in-a-new-era-for-consumers-and-merchants