




Zero-Knowledge Cryptographic Proofs as a Trust Mechanism For Financial and Government Digital Services

Aigumov T.G., ^a, Abdulmukminova F.M., ^b, Sakhbieva A.I. ^c

¹PhD in Economics, Associate Professor, Associate Professor of the Department of Information Security and Software Engineering FSBEI HE "Dagestan State Technical University", Makhachkala, Russian Federation

²Postgraduate student of the Department of Information Security and Software Engineering at the Federal State Budgetary Educational Institution of Higher Education Dagestan State Technical University, Makhachkala, Russian Federation

³associate professor of Department of Financial Markets and Financial Institutions, Kazan Federal University, Russian Federation

915533@mail.ru, eguri@inbox.ru, aminasmile@mail.ru,


Keywords: zero-knowledge proofs, cryptographic protocols, trust infrastructure, digital financial services.

Abstract: This article explores the role of zero-knowledge cryptographic proofs as a basic trust mechanism for financial and government digital services. The theoretical section explains how such proofs shift compliance verification from a "reveal data and verify" mode to a "prove property and admit access to service" mode, thereby reducing transaction costs, mitigating information asymmetries, and limiting agency costs. Drawing on insights from new institutional economics and mechanism theory, it demonstrates how formalized verifiability improves rule robustness, makes truthful reporting individually rational, and reduces the negative externalities of leaks. The analytical section systematizes classes of constructions, highlighting tradeoffs between proof size, latency, the need for trusted configuration, and operational risks. The practical section describes applications in payment infrastructure, lending, insurance, cross-border settlements, digital identity, and ledger management: range constraint verification, threshold and attribute verification, verifiable computation, and selective disclosure. It is demonstrated that with proper design, it is possible to combine data minimization with targeted transparency and effective enforcement, including in projects involving central bank digital money and pan-European identity wallets. Implementation metrics are discussed: marginal cost of proof and verification, average latency, fault tolerance, and the proportion of cases where primary documents are replaced with proof of property.

1 INTRODUCTION

At the beginning of the third decade, the digital economy faces a dual challenge: on the one hand, government and financial platforms strive for maximum automation, interoperability, and reduced service costs; on the other hand, society's aggregate demand for privacy and control over personal data is growing. In economic theory, trust is viewed as a factor directly influencing transaction costs and the effectiveness of institutions. Oliver Williamson demonstrated that the costs of information search, contract formation, and contract enforcement

increase with the increase in opportunistic behavior, and therefore, institutions that reduce uncertainty and protect rights enhance the efficiency of exchange. Douglass North linked long-term growth to the quality of institutions and law enforcement mechanisms. Under conditions of information asymmetry, as described by George Akerlof, adverse selection drives out "good" goods, and therefore quality and reliability checks became a key public good. Finally, Francis Fukuyama linked sustainable prosperity to high levels of interpersonal and institutional trust. Zero-knowledge cryptographic proofs arise as a response to this dilemma: they allow

^a <https://orcid.org/0000-0002-8737-0228>

^b

^c

^h

^h

^t

^p

^ø

^s

^l

[/]

^ó

the verifier to verify the truth of a statement without receiving any additional information about the hidden data. The concept was formulated in the work of Shafi Goldwasser, Silvio Micali, and Charles Rackoff within the framework of interactive proof theory; such protocols were later shown to exist for a wide range of problems. Intuitively, this is a "proof without disclosure": a participant can confirm knowledge of a secret without revealing either the secret itself or any side information. From a trust perspective, this means a shift from verification through data disclosure to verification of the properties of the data. This shifts the economic tradeoff between security and privacy, reducing verification costs and the risk of leaks.

The next step is to eliminate interaction rounds and develop non-interactive schemes on a shared random string: Manuel Blum, Paul Feldman, and Silvio Micali demonstrated that interactivity can be replaced by shared randomness, paving the way to compact proofs suitable for mass digital services. This is critical for financial and government platforms with tight latency and bandwidth constraints: minimizing communication curbs transaction costs and improves user experience. «More recent work has substantiated the feasibility of terse knowledge arguments, where proof size and verification time grow slowly relative to the complexity of the computation being verified, laying the foundation for practical protocols in payment infrastructure and identity services» [3, p. 81].

The current landscape includes a whole family of constructions. Jens Groth's construction ensures extremely small proof size and fast verification—properties that directly impact the marginal costs of verification on nodes in mass-market systems. Scalable "transparent" knowledge arguments developed by Eli Ben-Sasson et al. eliminate the need for trusted pre-configuration and are focused on post-quantum security, reducing institutional risks and the costs of managing key materials. Finally, "short proofs for confidential transactions" proposed by Benedikt Bünz et al. enable efficient proving of range constraints without disclosing amounts, which is particularly important for financial messaging and reporting. These results essentially transform trust from an organizational resource into a computational procedure with a predictable cost.

The economic significance of this shift is particularly noticeable in payment systems and in the design of central bank digital currencies. The Bank for International Settlements notes that central banks are widely studying digital currency architectures, with the key issue being how to reconcile user privacy requirements with anti-money laundering and

counter-terrorism financing objectives and financial oversight tasks. Research shows that user trust in central bank digital currencies is significantly dependent on data protection regimes; privacy-enhancing technologies, including zero-knowledge proofs, are viewed as a tool for implementing selective compliance checks without blanket disclosure. From an economic perspective, this means reducing agency costs and moral hazard risks while simultaneously meeting regulatory constraints.

Government digital services face similar challenges: identity verification, attribute verification, access to registries, and filing declarations—all of which require reliable verification while minimizing the volume of personal data transferred. The European regulatory framework for digital identity envisages the introduction of pan-European identity wallets; official documents explicitly describe the possibility of selective disclosure and the use of zero-knowledge proof to confirm facts (for example, reaching a certain age) without disclosing unnecessary information. Project documents on verifying the revocation of attestations and concealing publisher data using such proof are already being published; a practical roadmap for age verification while preserving privacy is also being developed. This reduces the regulatory and social costs of data processing and increases trust in public services. From the perspective of economic mechanism design, pioneered by Leonid Gurvich and developed by Eric Maskin and Roger Myerson, zero-knowledge technologies enable the construction of rules of the game in which truthful disclosure becomes individually rational, and the social choice function is achievable without excessive disclosure. Verifying compliance with constraints (e.g., capital adequacy, geographic residency, or age criteria) through proof of property, rather than data disclosure, reduces the likelihood of strategic manipulation while simultaneously protecting commercial and personal privacy. In terms of new institutional economics, this reduces reliance on centralized intermediaries for trust functions, transferring some enforcement to the cryptographic layer of the infrastructure, which ultimately reduces transaction costs and expands the scope of digital services markets.

Of course, implementation limitations are important in practice. Some schemes require trusted parameter configuration; its compromise creates systemic risks, and operational procedures for securely conducting ceremonies generate fixed costs. Modern transparency-focused designs seek to eliminate this need, but are offset by higher computational costs and requirements for provable

computational paths. An economically feasible analysis of the marginal cost of proof and verification per unit of service, taking into account scalability, latency, and integration with existing certificate authority and registry protocols, is appropriate. Taken together, the maturity of the theory and the availability of practical regulatory guidance mean that zero-knowledge proofs have moved from the realm of abstract cryptography to the toolkit of data policy and trust architecture for public and financial platforms.

Zero-knowledge cryptographic proofs can be viewed as a trust-saving technology: they replace data disclosure with verification of data properties, transforming trust from fuzzy social capital into a reproducible computational procedure with a predictable cost. For digital platforms, this means lower transaction costs of surveillance and control, reduced agency costs between service providers, users, and regulators, and mitigation of information asymmetries, which, as George Akerlof noted, lead to adverse selection and destroy markets. In the terms of new institutional economics (Oliver Williamson, Douglass North), this technology allows for the design of rules and institutions to minimize opportunism while preserving privacy—that is, improving the "quality" of institutions without increasing the volume of personal data in circulation. This achieves a reduction in the costs of searching, concluding, and monitoring contracts while simultaneously increasing the predictability of enforcement.

The idea of zero-knowledge proofs arose in the context of interactive proofs: it was shown that it is possible to convince a verifier of the truth of a statement without revealing anything beyond its truth value. This shift—from content verification to correctness verification—created the possibility of information-minimal trust mechanisms. Further development led to non-interactive schemes in which interaction is replaced by shared randomness or other sources of shared parameterization. This is critical for the economy of digital services: reducing communication rounds reduces latency, communication costs, and increases the throughput of queueing systems. This has given rise to a whole class of laconic knowledge arguments, in which the proof size and verification costs grow slowly relative to the complexity of the original computation, directly impacting the marginal cost of verification in payment gateways, identification systems, and registries.

«A key contribution to practical applicability was made by results demonstrating the feasibility of ultra-

small proofs and fast verification. Jens Groth's 2016 construction has become the starting point for numerous production systems, as it enables extremely compact proofs and a small amount of computation on the verifier's side. In economic terms, this translates into reduced load on infrastructure nodes and, consequently, a lower total cost of ownership for platform operators, especially with large request flows. However, this efficiency is often achieved at the expense of pre-trusted parameter configuration, which entails organizational risks and fixed costs of managing key assets. These tradeoffs are a classic subject of contract theory analysis: we gain benefits in ongoing operating expenses but incur the risks and costs of a one-time configuration "ceremony," which should be factored into the service lifecycle calculation» [1, p. 193].

An alternative development path is represented by "transparent" and scalable knowledge arguments, which eliminate the need for trusted configuration. The work of Eli Ben-Sasson and co-authors has shown that it is possible to construct proofs resistant to a quantum adversary, with verification growing significantly slower than the volume of initial data. From an economic analysis perspective, this reduces institutional risks (there is no trusted party, no threat of capture or loss of master parameters) and the costs of compliance control over parameter generation procedures. The cost is increased proof size and generation time, meaning increased computational costs for the prover. Optimizing the balance between operational transparency and computational burden is the subject of design decisions for specific services: financial payment systems, attribute certification systems, registries, etc.

Of particular importance for financial services are "short proofs" of range constraints, which allow confirmation that a hidden value lies within an acceptable interval without revealing the value itself. Such proofs enable compliance with rules (e.g., the correctness of amounts, limits, and reserve ratios) to be verified without disclosing transaction details or commercial secrets. In terms of risk management, this reduces the leakage surface, reduces the likelihood of price discrimination based on identified volumes, and enables selective disclosure of only what is strictly necessary for regulatory review. The economic benefits are felt on both the consumer and provider sides: privacy costs are reduced, the perceived quality of service and, consequently, its acceptance are increased.

From a practical perspective, zero-knowledge proofs act as a "transmission belt" between privacy objectives and law enforcement requirements. In a

recent study on privacy-enhancing technologies for digital payments, the Bank for International Settlements notes the potential of modern cryptographic proofs, while simultaneously highlighting their current limitations in terms of security and computational capacity. The conclusion is that further engineering refinement and integration with data minimization architectural techniques are required. For central bank digital currency projects and related payment solutions, this means the ability to build selective compliance verification modes: the user proves the correctness of transaction parameters (e.g., compliance with limits or the absence of "black" connections) without disclosing full data sets. «This approach reduces agency costs between the issuer, intermediaries, and users and, according to international institutions, can increase trust in systems while maintaining supervisory effectiveness» [7, p. 72].

In the government digital identity sector, zero-knowledge proofs are closely linked to the concept of selective attribute disclosure. The European initiative for a common European identity wallet explicitly requires proof of facts (e.g., "age of majority" or "resident of a given jurisdiction") without sharing unnecessary data. Official materials and accompanying documents describe the role of such proofs as a way to maximize compliance with the data minimization principle and reduce the risk of linkability of user actions across different domains. For the public service economy, this means reduced regulatory and social costs of personal data processing, as well as reduced legal exposure for service providers, since the volume of stored and processed data is minimal.

A current example of a practical trust mechanism is age verification for online services. In 2025, the European Commission presented a roadmap and pilot solutions proposing to implement verification of "over eighteen" using cryptographic proofs, excluding the transmission of passport data and other identifiers. This design reconciles three goals: compliance with child protection requirements, data minimization, and increased user adoption. Economically, this reduces compliance costs for platforms (less risky storage, less handling of sensitive data) and mitigates potential externalities from breaches. The "mini-wallet" option allows for the mechanism to be introduced before the widespread deployment of a pan-European identity wallet, while remaining compatible with the subsequent transition. Equally important is the "verifiable computation" approach, where the correctness of a complex algorithm's execution is

proven without requiring a verifier to rerun it. This is relevant for government registries, tax services, and all sorts of decision-making subsystems that outsource computations. The economic rationale is a radical reduction in the costs of monitoring the quality of outsourced services and a lowering of barriers to specialization: the verifier receives a brief proof of correctness without the expense of a full recalculation, and the provider demonstrates compliance with the specification without disclosing internal data and code. Moreover, theoretical and engineering results in recent years indicate the practical feasibility of such schemes, expanding the class of problems where "trust through proof" is displacing "trust through disclosure."

From a mechanism theory perspective (Leonid Gurvich, Eric Maskin, Roger Myerson), zero-knowledge cryptographic proofs can be interpreted as a tool for enhancing the feasibility of desired rules under rationality and privacy constraints. While the social choice function requires verifying compliance with constraints (age, geography, wealth, absence of conflicts of interest), proofs of properties without disclosing the original data enable truthfulness strategies to be implemented while preserving individual rationality and incentive compatibility. In other words, "properly designed" rules combined with such cryptography make truth-telling the dominant strategy without imposing full disclosures—which reduces the risks of manipulation and minimizes the welfare losses from forced data exchange.

Of course, any trust mechanism has operational and technological limitations. Modern design approaches distribute costs between prover and verifier differently, differ in proof size, the need for trusted configuration, resistance to a quantum adversary, and the difficulty of integrating with existing attestation formats. Some schemes, while conserving computing resources for the verifier, impose a high workload on the prover; others, being "transparent," reduce institutional risks but increase traffic and verification time. International organizations emphasize that the technology is promising but requires development and must be combined with an architecture that minimizes data, localizes responsibility, and ensures procedural transparency—only in this way can a sustainable balance be achieved between privacy, oversight, and efficiency. For service designers, this means calculating the total cost of ownership, taking into account the workload profile, risk scenarios, and regulatory requirements.

Returning to the broader economics of trust, zero-knowledge cryptographic proofs act as a technological "bridge" between private and public efficiency. They reduce information asymmetry by avoiding adverse selection; reduce transaction costs of monitoring without shifting the risk of leaks to users; and strengthen institutional expectations without requiring expanded control. Moreover, in areas where society places heightened demands on rights and freedoms—from payment infrastructure to authentication services—such technology supports trust as a public good, enhancing the legitimacy and acceptability of digital platforms. This is why regulators and research centers are incorporating zero-knowledge proofs into the "toolkit" of privacy-enhancing technologies, seeing them as a way to reconcile the goals of privacy, law enforcement, and innovative growth [4, p. 44].

Finally, it's worth emphasizing that this is not about replacing institutions with technology, but rather about redistributing trust functions between legal norms, organizational architecture, and cryptographic procedures. As institutional researchers demonstrate, sustainable growth is ensured not only by property rights and independent courts, but also by adaptive mechanisms that reduce the uncertainty of interactions. In this logic, zero-knowledge proofs become the "hidden" component of the trust infrastructure that allows government and financial services to scale without becoming repositories of redundant personal data. The economic result is more transactions with less risk and lower control costs, i.e. a shift in the production possibilities curve of the digital state and digital market to the right.

The use of zero-knowledge proofs in financial digital services and the public sector is rationally viewed through the lens of reducing transaction costs, correcting information asymmetries, and streamlining enforcement mechanisms while maintaining privacy. In the financial sector, this reduces monitoring costs, simplifies compliance, and mitigates the likelihood of adverse selection. In public services, it combines data minimization with verifiability of citizens' rights and attributes. International organizations emphasize that privacy-enhancing technologies are already capable of combining confidentiality and granular transparency: instead of mass data disclosure, rules are being developed that only verify compliance with requirements, while sensitive information remains hidden. This shift in trust from "reveal and verify" to "prove property and allow access to the service,"

thereby adjusting participant incentives and reducing the overall cost of compliance.

In the payment infrastructure, zero-knowledge proofs allow the verification of regulatory constraints to be embedded into the very logic of transactions. The concept of "built-in compliance" is that rules verification is performed automatically and does not require the transfer of unnecessary data between intermediaries and oversight. Project work at international platforms demonstrates that cross-border compliance requirements can be encoded into formal conditions that are verified before or during settlement: participants prove compliance with thresholds, sanctions restrictions, or origin-of-funds rules without disclosing unnecessary transaction parameters. This reduces information flows, reduces the risk of leaks, and simultaneously speeds up settlements—a crucial factor for the network effects of payment systems, where latency and friction directly impact utility for the marginal user.

Central bank digital currencies are another area where economic logic requires a tradeoff between household privacy and the ability to achieve targeted oversight. Pilot and research projects demonstrate the feasibility of "targeted verifiability" regimes, in which individual transaction attributes are verified without revealing the payer's identity or other details. Thus, studies highlighting the results of Swiss research conducted in collaboration with international partners demonstrate that it is possible to achieve "payer anonymity" while maintaining the visibility of the recipient, as well as verify the authenticity of funds and the absence of abuse without disclosing excessive information with each payment. From an economic perspective, this reduces agency costs between the issuer, intermediaries, and users, maintains trust in the unit of account, and mitigates the negative externalities of centralized personal data storage.

The findings of industry reviews on confidential payments are also of practical interest: they classify approaches along the "privacy-verifiability" axis and emphasize that strong technological privacy, combined with limited target transparency, is achievable through cryptographic procedures, including zero-knowledge proofs. A typical tradeoff is the redistribution of the computational load between the prover and the verifier, which is important to consider when calculating marginal costs and choosing the architecture of the payment platform. For retail systems with mass micropayments, low latency and low verification costs at the network node are critical; for cross-border settlements, the ability to encode complex regulatory

rules into formal terms that are verified automatically and locally is crucial.

$$C_i = \text{gmihri}, C = \text{gmhr} = i = 0 \prod [n - 1 C_i 2i \cdot \text{hr} - \sum 2i r_i]$$

In the insurance market, zero-knowledge proofs help reduce information rent and the problem of moral hazard. Policyholders can confirm compliance with rate conditions (age threshold, absence of certain events in the policyholder's history, medical test results within a specified range) without disclosing their entire history or underlying indicators. For the insurer, this means more accurate pricing with lower data processing costs and reduced legal liability for data storage; for the client, it reduces privacy costs and increases product acceptance. Industry analyses highlight emerging opportunities, particularly in segments where sensitive data traditionally hinders demand—health and personal risks—and point to the need to standardize interfaces and verification procedures. «The overall effect is increased reach and lower barriers to entry with sustainable risk management» [6, p. 52].

In lending, proof of property without disclosing original documents helps mitigate information asymmetries in creditworthiness assessments. Borrowers can confirm income above a threshold or the adequacy of collateral relative to the loan amount without providing a full set of documents. At the incentive level, this reduces the likelihood of strategic behavior (concealing unfavorable facts or, conversely, excessive disclosure with the risk of leaks), and at the cost level, it reduces the burden on document exchange channels and compliance departments. Modern reviews of privacy-enhancing technologies recommend combining such evidence with a data minimization architecture, limiting the storage and replication of copies of personal information. This improves the quality of selection and reduces adverse selection, which, taken together, expands the supply of credit without worsening the risk profile.

For cross-border payments and trade settlements, the added value of this technology lies in reducing the costs of coordinating requirements across different jurisdictions. "Rules-in-code" mechanisms, where compliance with regulations is automatically verified before settlement, allow banks and payment institutions to reduce delays caused by regulatory verification. Analysis of international platforms emphasizes that linking fast payments across borders requires not only technical connectivity but also coordinated oversight regimes: zero-knowledge proofs act as a "universal adapter" to enable verification without the cross-border transfer of large amounts of personal data. This reduces operational

risks and facilitates the participation of new providers, creating a more competitive environment.

In the public sector, identity verification and attribute verification are central. The European regulatory framework for a pan-European identity wallet explicitly provides for the ability to verify facts (such as age, residency, or license validity) without disclosing unnecessary information. Official documents emphasize the role of zero-knowledge proofs as a key security and privacy mechanism; citizens control what data is disclosed and to whom, while the verifying party receives only confirmation that conditions are met. Economically, this shifts the trust function from centralized databases to distributed cryptographic procedures, reducing the regulatory and social costs of processing personal data and increasing the acceptance of digital authentication services. A separate, already formalized scenario is proof of age for access to restricted resources. European authorities have presented a "standard blueprint" for a solution whereby the user proves they have passed the age threshold without disclosing their date of birth or supporting documents. For platforms, this means reduced legal risks and document handling costs, and for citizens, no need to provide operators with copies of their IDs. Pilot projects and guidelines link such checks with future identity wallets, ensuring compatibility and the expected interoperability. At the public good level, this supports the goals of protecting minors without imposing ubiquitous identification or creating redundant personal data storage.

$$R \subseteq \{0,1\}^{**} \times \{0,1\}^*$$

$$LR = \{x \mid \exists w: R(x,w) = 1\}.$$

Standardization is critical for reducing implementation costs and network interoperability. Electronic attribute attestation profiles and relevant organizations' materials support selective disclosure and cryptographic proof of properties as a core function of certification systems. This establishes a "common language" for agencies, banks, and private service providers, enabling the reuse of proven building blocks and avoiding the creation of incompatible solutions. From an economic perspective, standards reduce fixed integration costs, create predictable network externalities, and stimulate competition through component interchangeability.

$$\Pr(x, \pi^*) = 1 \wedge \neg \exists w: R(x,w) = 1 \leq \epsilon(\lambda)$$

In the management of public registries and permitting procedures, zero-knowledge proofs solve the problem of "attribute-based access." Instead of submitting a complete document to the authority, the applicant confirms only the required property:

compliance with a professional license, the absence of a ban, or the property's required legal status. This reduces the burden on registries, reduces the volume of data that must be stored and protected, and speeds up permitting processes. Overall, this reduces queuing and monitoring costs, increases the predictability of deadlines, and increases applicant satisfaction—metrics that, in the digital economy, are directly linked to the investment attractiveness of jurisdictions and the efficiency of administrative procedures. Official reports on the digital infrastructure of public services emphasize the need to link such mechanisms with common platforms and partnerships between countries to ensure cross-border recognition of certifications.

Tax administration and regulatory reporting also benefit from proof of properties without disclosing source documents. A taxpayer or reporting organization can prove compliance with thresholds, the absence of prohibited links, and the accuracy of aggregates without submitting the full data set. For the government, this reduces verification and storage costs, decreases vulnerability to leaks, and improves the accuracy of sample checks. For businesses, it reduces the risk of disclosure of trade secrets and compliance costs. Academic and interagency materials point to the particular value of such procedures for "regulatory on-site audits," where it is important to confirm a rule but it is undesirable to transfer the entire array of primary data across agency and national borders.

Sectoral examples can be expanded to include social support and education. When applying for benefits, a citizen can confirm eligibility without disclosing excessive information about family composition or health status; when applying for admission to an educational institution, they can confirm the possession of a high school diploma and grades within the specified range without submitting the entire file. Such "pinpoint evidence" reduces stigma, reduces the likelihood of discrimination based on excessively disclosed information, and simultaneously facilitates interagency exchange. As a result, public trust in digital channels grows, while budget costs are reduced by automating verification and reducing manual approvals. European roadmaps for digital identity and public infrastructure emphasize precisely this combination: data minimization, unified specifications, and verifiability of the result.

It is important to consider the marginal costs of computation and communication. Some designs require significant resources on the prover side; in mass services, this can impact throughput and latency.

In platform economics, this means that architectural decisions must be made taking into account the workload profile: in payment systems, cheap and fast verification is critical; in certifying systems, resilience to attestation collisions and the ability to perform offline verification; in cross-border settlements, flexible implementation of complex rules without centralizing databases. Design reviews of new settlement and certifying systems provide detailed lists of tradeoffs: privacy versus observability, resistance to quantum threats versus performance, local verifiability versus message size, and so on. The rational choice of balance here is the task of comparing alternatives in terms of total cost of ownership and resilience to operational and legal risks.

The overall regulatory focus in the economy is emphasizing technologies that simultaneously reduce risks and avoid barriers to innovation. Reports from international institutions emphasize that zero-knowledge proofs are particularly useful for demonstrating compliance without sharing sensitive data, thereby reducing the costs of interaction between regulated entities and oversight bodies. This creates the foundation for new market models—from tokenized assets with programmable compliance to "smart" reports, where verification is performed by the service provider and the regulator receives confirmation of the result. Against this backdrop, entry barriers for new participants are lowered, and competition for quality and convenience improves overall user well-being.

Finally, a coherent ecosystem is being built at the level of pan-European initiatives: digital identity wallets, unified attribute attestation profiles, age verification guidelines, and platforms for experimentation and the exchange of best practices. This allows governments and private providers to implement zero-knowledge proofs not as an exotic feature, but as a standard element of the trust architecture. In institutional terms, this corresponds to a transition from individual agreements to formalized rules and interfaces, which reduces uncertainty and predictably lowers transaction costs as digital services scale. As components mature and interfaces become standardized, we can expect an expanding range of services where property verification replaces the transfer of raw data—from social programs and education to taxation and permitting procedures—with tangible benefits for the budget, citizens, and businesses alike.

As a result, zero-knowledge proofs are being integrated into the economy of financial and

government digital services as a universal trust mechanism that reduces information redundancy, protects privacy, and ensures the verifiability of rules. Their use reduces monitoring costs, curbs the negative externalities of leaks, mitigates the problem of adverse selection, and shifts government-business-citizen interactions to a more efficient technological foundation. In the context of increasing digitalization and interoperability, it is precisely these solutions that ensure sustainable alignment of goals: protection of rights, effective law enforcement, and resource savings through the formalization of trust in the form of cryptographic procedures.

2 CONCLUSION

Zero-knowledge proofs thus transform trust from an informal public resource into a reproducible cryptographic procedure embedded in the rules of interaction and measurable in terms of costs and benefits. For financial and government digital services, this means a shift from verification through the disclosure of data sets to verification of data properties, which simultaneously reduces transaction costs, mitigates information asymmetries, reduces agency costs, and limits the negative externalities of data breaches. Within such an architecture, participants prove compliance rather than revealing raw data; enforcement relies on formally verifiable conditions rather than centralized accumulations of personal data.

In the financial sector, this provides practical mechanisms for settlements, lending, and insurance: thresholds and limits are verified without the transmission of sensitive parameters, reducing the risk of adverse selection, reducing surveillance costs, and speeding up transactions. In cross-border scenarios, zero-knowledge proofs allow for the coordination of requirements across different jurisdictions without the constant exchange of personal information, thereby reducing latency and increasing resilience to operational risks. In payment systems and central bank digital money projects, the technology supports a "selective verifiability" regime, combining privacy protection with targeted oversight and improving the perceived quality of the settlement infrastructure. Network effects are enhanced by reducing friction when onboarding new participants and simplifying compliance.

In the public sector, zero-knowledge proofs enable selective verification of attributes in identity verification, registry access, benefit allocation, and tax administration. This reduces the volume of data

collected and stored, speeds up procedures, and increases citizen trust in digital channels. This creates a public good effect: security and privacy are simultaneously enhanced, while predictability and transparency of procedures strengthen the legitimacy of services.

However, economic tradeoffs remain. Different designs distribute computational costs between prover and verifier in different ways; some solutions impose pre-configuration requirements, which create fixed costs and institutional risks; and latency and scalability limitations are noticeable. Coordination costs are no less important: interface standardization, certification procedures, legal certainty regarding responsible parties, and the ability to challenge results. Therefore, rational implementation requires calculating the total cost of ownership, taking into account the workload profile, risk scenarios, and regulatory requirements, as well as a mechanism for distributing benefits and costs among participants.

The political and economic challenge for the near future is to enshrine the principle of data minimization "by default," develop open specifications and compatible attestation formats, implement independent verification of correctness, and develop performance indicators: marginal cost of proof and verification, average latency, fault tolerance, and the proportion of cases in which the transfer of primary documents can be replaced by proof of property. The scientific agenda includes an analysis of incentive mechanisms, an assessment of distributional effects for households and businesses, a study of resilience to future computational threats, and modeling intertemporal tradeoffs during infrastructure modernization.

Thus, zero-knowledge cryptographic proofs are not a highly specialized add-on, but a fundamental trust mechanism for digital government and digital finance. They allow for the reconciliation of three previously seemingly contradictory goals: data protection, effective law enforcement, and resource conservation. By embedding verifiability into the very fabric of processes, these proofs expand the boundaries of what is possible: more transactions with lower risk, higher quality services with lower costs, greater institutional resilience with less reliance on centralized data repositories.

REFERENCES

- Cheremushkin, A.V., 2009. Cryptographic Protocols. Basic Properties and Vulnerabilities. Moscow: Academy, 272 p.

- Danchenko, A.V., 2012. Theory and Technology of Information Transmission. Textbook. Moscow, 2012. 304 p.
- Konkin, A.Yu., Zapechnikov, S.V., 2024. Ensuring Information Confidentiality in Distributed Ledger Systems Using Zero-Knowledge Proofs. Information Technology Security, 31 (1), 75–85.
- Martynenkov, I.V., 2023. Brief Non-Interactive Zero-Knowledge Arguments Based on Sets of Polynomials. Applied Discrete Mathematics, 2023 (1), 20–57.
- Martynenkov, I.V., 2023. Methods for Improving the Performance of Short Non-Interactive Zero-Knowledge Arguments and Analysis of the Achieved Results. Applied Discrete Mathematics, 2023 (60), 40–58.
- Martynenkov, I.V., 2023. Secure Generation of Public Parameters and Elimination of Vulnerabilities in Short Non-Interactive Zero-Knowledge Arguments. Applied Discrete Mathematics, 2023 (61), 28–43.
- Muzykantskiy, A.I., Furin, V.V., 2011. Lectures on Cryptography. Moscow: MCNO, 232 p.
- Shlyakhtina, E.A., Gamayunov, D.Yu., 2021. A Group Authentication Scheme Based on Zero-Knowledge Proof. Applied Discrete Mathematics, 51, 68–84.
- Voloshin, S.K., Romankov, V.A., 2023. Discrete differentiations and integrations and their possible applications to algebra and cryptography. Applied Discrete Mathematics, 2023 (3), 5–14.
- Zubov, A.Yu., 2011. Almost perfect ciphers and
a
u
t
h
e
n
t
i
c
a
t
i
o
n

c
o
d
e
s
.

rete Mathematics, 4 (14), 28–33.