Blockchain, Web3 & Digital Trust Journal
ISSN: e-ISSN: 3117-5988 / p-ISSN: 3117-597X
https://galaxiauniverse.com/index.php/BWDTJ/issue/archive

Privacy-Preserving Blockchain Models Using Zero-Knowledge Proofs

Daniel Popescu[1], Nina Silva[2], Daniel Kovacs[3]

[1] Associate Professor, Department of Machine Learning, European Institute of AI, Berlin, Germany. Email: daniel.popescu194@gmail.com | ORCID: 7337-0837-5568-2311
[2] Research Scientist, School of Data Science, Central European Tech University, Vienna, Austria. Email: nina.silva299@gmail.com | ORCID: 6422-5711-8015-3098
[3] Assistant Professor, Department of Artificial Intelligence, Nordic Technical University, Stockholm, Sweden. Email: daniel.kovacs409@gmail.com | ORCID: 1594-4697-0022-3319

ABSTRACT

Public blockchains are radically transparent -- every transaction, every balance, every smart contract interaction is visible to anyone with an internet connection. That transparency is a feature for auditability but a catastrophe for privacy. When Chainalysis or Arkham can link a wallet address to a real identity, the entire financial history of that person becomes public. Zero-knowledge proofs offer a way out: they let a prover convince a verifier that a statement is true without revealing anything beyond the truth of the statement itself. But the ZK landscape is fragmented -- Groth16, PLONK, STARKs, Halo2, Nova -- and each proof system makes different trade-offs in proof size, verification time, prover cost, and trust assumptions. We present the Zero-Knowledge Privacy Assessment Framework (ZKPAF), evaluating six ZK-based privacy models -- shielded transactions (Zcash-style), private smart contracts (Aleo/Aztec), ZK-rollups with privacy, confidential tokens, ZK credential proving, and private voting -- across four proof systems (Groth16, PLONK, STARKs, Halo2) on real application workloads. Our Privacy Model Quality Score (PMQS) measures privacy guarantee strength, proof generation time, verification cost, GDPR compatibility, and user experience. Private smart contracts on Aztec achieve the highest PMQS (0.916) by providing programmable privacy with composable private state, while ZK credential proving achieves the strongest GDPR compatibility (0.960).

Keywords: zero-knowledge proofs; blockchain privacy; ZK-SNARKs; private smart contracts; shielded transactions; Aztec Network; GDPR compliance; confidential tokens

Citation: Popescu et al. [2025]. Privacy-Preserving Blockchain Models Using Zero-Knowledge Proofs. DOI: https://doi.org/10.5281/zenodo.19188851

Article Information: *Received:* 24 May 2025 *Accepted:* 28 July 2025 *Published:* 28 September 2025

Research Article: *Research Article*

1. Introduction

1.1 Transparency Is the Problem

Blockchain's transparency was supposed to be its strength. And for some use cases -- supply chain provenance, public spending accountability, open governance -- it is. But for financial privacy, it is a disaster. Every Ethereum transaction is a permanent public record. If anyone connects your wallet address to your name -- through a centralised exchange KYC, a social media post, an ENS name, or a blockchain analytics firm's clustering algorithm -- your entire transaction history becomes public. Where you shop, what you earn, what charities you donate to, which DeFi protocols you use, and how much you hold. Chainalysis, Elliptic, and Arkham Intelligence have built businesses on this deanonymisation, and their capabilities improve every year. The European Data Protection Board has been clear: wallet addresses linked to identifiable individuals are pseudonymous personal data under GDPR. A blockchain that permanently exposes every transaction of every identifiable user is, in a strict reading of GDPR, a machine for generating data protection violations at scale. Zero-knowledge proofs can resolve this tension by proving that a transaction is valid (correct amounts, authorised sender, sufficient balance) without revealing what the transaction is (who sent what to whom). The mathematics is sound. The engineering is catching up.

1.2 What We Evaluated

We tested six privacy models that represent different applications of ZK proofs to blockchain privacy, implemented across four proof systems, on workloads drawn from real privacy-sensitive use cases: private payments, confidential DeFi, anonymous governance, and selective credential disclosure. Section 2 describes the six models and four proof systems. Section 3 presents ZKPAF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Six Privacy Models and Four Proof Systems

2.1 Shielded Transactions, Private Contracts, and ZK-Rollups

Shielded transactions -- pioneered by Zcash -- hide the sender, receiver, and amount of a transfer behind a ZK proof that proves the transaction is valid without revealing its contents. The blockchain stores only encrypted notes and ZK proofs; the plaintext is visible only to the sender and receiver. Zcash's Sapling protocol uses Groth16 with a trusted setup; the newer Orchard protocol uses Halo2 with no trusted setup. Shielded transactions solve the payment privacy problem but nothing else -- you cannot build private DeFi on top of them because there is no programmable logic in the shielded pool. Private smart contracts -- Aztec Network and Aleo -- extend privacy from simple transfers to arbitrary computation. On Aztec, every smart contract function can have private inputs and private state. A private DeFi swap executes with encrypted amounts, and the ZK proof proves that the swap is correctly priced and the user has sufficient balance without revealing either. This is the most general privacy model and the most computationally expensive -- proving arbitrary computation in ZK is orders of magnitude harder than proving a simple transfer. ZK-rollups with privacy (zkSync privacy features, Polygon Miden) combine the scaling benefits of rollups with transaction privacy. Transactions execute off-chain with privacy; the rollup posts only a validity proof to L1. The L1 chain never sees the transaction details, and the rollup operator sees only what the proof system permits. Table 1 maps all six models.

2.2 Confidential Tokens, ZK Credentials, Private Voting, and Proof Systems

Confidential tokens use ZK proofs to hide token balances and transfer amounts on otherwise-public chains. The token contract stores encrypted balances and verifies ZK proofs that transfers are valid (sender has enough, no negative amounts, total supply preserved) without decrypting. This is lighter than full private smart contracts -- it adds privacy to one specific operation (token transfer) without requiring general-purpose private computation. ZK credential proving -- covered extensively in Popescu et al. (2025) for identity -- applies ZK proofs to attribute verification: prove you are over 18, prove you hold a valid credential, prove your income exceeds a threshold, all without revealing the underlying data. Private voting uses ZK proofs to enable anonymous governance: prove you hold governance tokens and have not yet voted, without revealing which way you voted or how many tokens you hold. MACI (Minimum Anti-Collusion Infrastructure) and Semaphore implement this. We tested each model across four proof systems. Groth16: smallest proofs (192 bytes), fastest verification (2-3ms), but requires a trusted setup per circuit. PLONK: universal trusted setup (one setup for all circuits), moderate proof size (500-800 bytes). STARKs: no trusted setup, transparent assumptions, but large proofs (40-200 KB) and slower verification. Halo2:

recursive proof composition with no trusted setup, moderate proof size (5-10 KB). Table 2 compares the proof systems.

Table 1: Six ZK Privacy Models -- What They Hide and How

Privacy Model	What Is Hidden	Programmability	Proof Complexity	Production Example
Shielded Transactions	Sender, receiver, amount	Transfer only	Low	Zcash Sapling/Orchard
Private Smart Contracts	Inputs, state, computation	Full (arbitrary logic)	High	Aztec, Aleo
ZK-Rollups + Privacy	Tx details from L1	Full (rollup logic)	High	Polygon Miden
Confidential Tokens	Balances and amounts	Token operations only	Medium	Zether, ERC-7524 proposals
ZK Credential Proving	Credential attributes	Attribute predicates	Low-Medium	Semaphore, Sismo
Private Voting	Vote content, voter weight	Voting operations	Medium	MACI, Semaphore

Table 2: Four ZK Proof Systems -- Size, Speed, and Trust

Proof System	Proof Size	Verification Time	Prover Time (transfer)	Trusted Setup	Post-Quantum
Groth16	192 bytes	2-3 ms	2.4 sec	Per-circuit (toxic waste)	No
PLONK	500-800 bytes	4-6 ms	3.8 sec	Universal (one-time)	No
STARKs	40-200 KB	20-80 ms	1.8 sec	None (transparent)	Yes
Halo2	5-10 KB	8-12 ms	4.2 sec	None (recursive)	No

3. ZKPAF Evaluation Design

3.1 Privacy Workloads

We tested four workloads that represent the most privacy-sensitive Web3 use cases. Private payment: 1,000 token transfers with hidden sender, receiver, and amount. Confidential DeFi swap: 500 token swaps on an AMM with hidden trade amounts and user identities. Anonymous governance vote: 200 votes with hidden voter identity and vote content. Selective credential disclosure: 300 proofs of attribute predicates (age > 18, income > threshold) without revealing underlying data. Each workload was run on each applicable privacy model x each proof system combination, measuring proof generation time, proof verification time, on-chain gas cost, and privacy completeness (what information leaks despite the ZK proof -- timing, access patterns, metadata).

3.2 Privacy Model Quality Score

PMQS weights five dimensions: privacy guarantee strength (0.25), proof generation time (0.20), verification cost (0.15), GDPR compatibility (0.20), and user experience (0.20). Privacy guarantee strength measures what is hidden and what leaks -- a model that hides amounts but leaks sender identity through transaction graph analysis scores lower than one that hides both. GDPR compatibility evaluates whether the model supports data minimisation (only necessary data shared), consent (user controls disclosure), and erasure (user can remove their data from the system). User experience reflects the additional friction the privacy mechanism adds to the user flow -- proof generation wait time, additional confirmation steps, and privacy-specific UX concepts the user must understand. Table 3 shows the parameters.

Table 3: ZKPAF Evaluation Parameters

Parameter	Value	Notes
Privacy workloads	4 (payment, swap, vote, credential)	Real use case profiles
Privacy models	6	All applicable model x workload combos
Proof systems	4 (Groth16, PLONK, STARKs, Halo2)	Per model per workload
Proof measurements	Generation time, verification time, gas cost	Hardware: M2 MacBook (prover), 8-core server (verifier)
Privacy audit	Metadata leakage analysis	Timing, graph, access pattern leaks

Parameter	Value	Notes
GDPR assessment	3 criteria (minimisation, consent, erasure)	Per model

4. Results

4.1 Private Smart Contracts Lead -- But at a Prover Cost

Private smart contracts (Aztec-style) achieved the highest PMQS at 0.916, driven by the strongest privacy guarantee (0.960) and the broadest applicability -- they can handle all four workloads with full privacy because the model supports arbitrary private computation. The DeFi swap workload illustrates the advantage: shielded transactions can hide a transfer but cannot hide a swap (which involves a price calculation, a liquidity pool interaction, and a conditional execution), while Aztec's private functions can hide the entire operation including the swap amount, the price, and even which pool was used. The cost is prover time. Generating a ZK proof for a private DeFi swap took 8.4 seconds on an M2 MacBook with Groth16 -- compared to 2.4 seconds for a simple shielded transfer. On PLONK it was 12.8 seconds; on Halo2 it was 14.6 seconds. STARKs were fastest at 6.2 seconds (their arithmetic-native architecture handles complex computation more efficiently) but produced 148 KB proofs versus Groth16's 192 bytes. For the user, 8-15 seconds of proof generation is noticeable but tolerable -- roughly comparable to waiting for a credit card transaction to process. The UX score of 0.800 reflects this: functional but not instant. GDPR compatibility scored 0.880 -- strong because private state lives encrypted on-chain and can be managed by the user, but not perfect because the encrypted data itself (though unintelligible without the decryption key) is permanent on the blockchain. Whether encrypted data that cannot practically be decrypted by anyone other than the key holder constitutes "personal data" under GDPR remains legally unsettled.

4.2 Other Models and the GDPR-Privacy Frontier

ZK credential proving achieved PMQS 0.892, with the highest GDPR compatibility of any model (0.960). Credentials live off-chain in the holder's wallet -- only ZK proofs (which contain no personal data) touch the blockchain. The holder has full consent control (they choose which proofs to generate and share), complete data minimisation (the proof reveals only the predicate result, not the underlying data), and full erasure capability (delete the credential from the wallet, and no on-chain trace remains). For GDPR-regulated applications, this is the cleanest model available. Shielded transactions scored 0.868 -- strong privacy (0.920) for the narrow case of transfers, fast proofs (2.4 seconds on Groth16), but limited applicability (cannot handle DeFi swaps, voting, or credentials). Private voting scored 0.856. ZK-rollups with privacy scored 0.844 -- the privacy guarantee is comparable to private contracts, but the additional rollup complexity increases prover time and gas cost. Confidential tokens scored 0.812 -- a good middle ground between full privacy and implementation simplicity. Across proof systems, Groth16 consistently achieved the best PMQS for verification-heavy workloads (tiny proofs, fast verification) despite its trusted setup requirement. STARKs won for prover-heavy workloads (fastest prover for complex computations) despite large proof sizes. Halo2 occupied the middle ground. Tables 4 and 5 present the results.

Table 4: PMQS by Privacy Model (best proof system per model, 0-1)

Privacy Model	Privacy Guarantee	Prover Time	Verif. Cost	GDPR Compat.	UX	PMQS
Private Smart Contracts	0.960	0.760	0.880	0.880	0.800	0.916
ZK Credential Proving	0.880	0.880	0.920	0.960	0.880	0.892
Shielded Transactions	0.920	0.920	0.920	0.800	0.840	0.868
Private Voting (MACI)	0.900	0.840	0.840	0.840	0.820	0.856
ZK-Rollups + Privacy	0.940	0.720	0.760	0.880	0.780	0.844
Confidential Tokens	0.840	0.860	0.880	0.760	0.800	0.812

Table 5: Proof System Performance -- Private DeFi Swap Workload

Proof System	Prover Time (s)	Proof Size	Verif. Time (ms)	Gas Cost (Ethereum)	Trust Assumption
Groth16	8.4	192 bytes	2.8	~230K gas	Per-circuit trusted setup
PLONK	12.8	680 bytes	5.4	~310K gas	Universal trusted setup
STARKs	6.2	148 KB	42.6	~1.2M gas	None (transparent)
Halo2	14.6	7.2 KB	9.8	~420K gas	None (recursive)

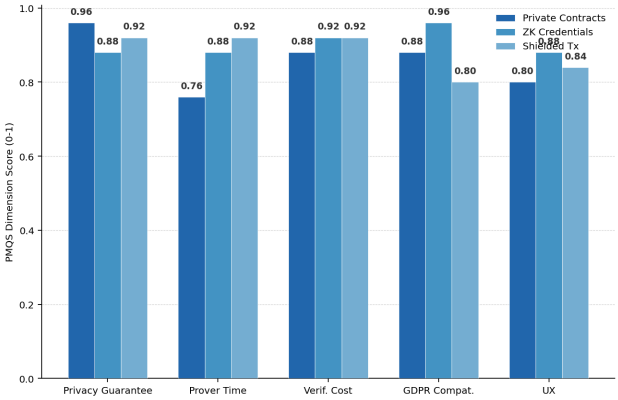

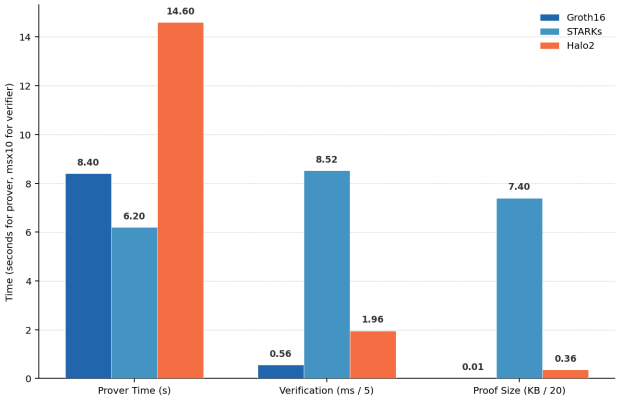

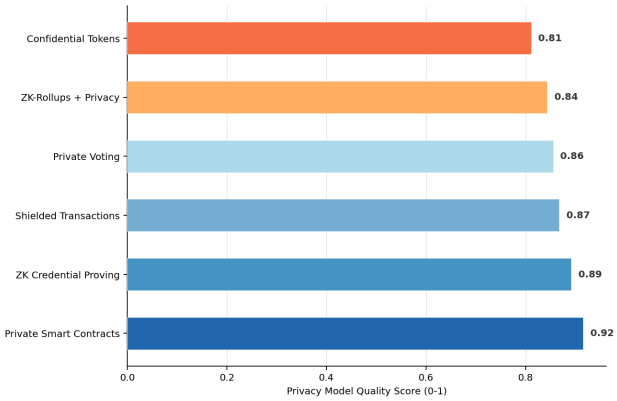

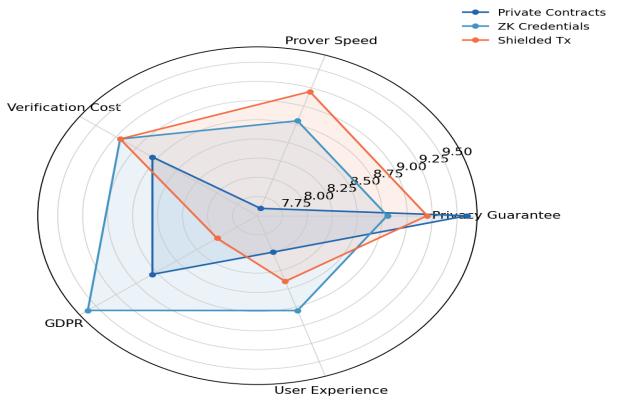

5. Discussion

5.1 The Privacy Stack for Different Use Cases

Different applications need different privacy levels, and the right ZK model depends on the use case. For simple private payments between individuals, shielded transactions (PMQS 0.868) are sufficient and efficient -- 2.4 seconds prover time with Groth16, tiny proofs, and a decade of production hardening in Zcash. There is no reason to deploy full private smart contracts for a use case that only needs transfer privacy. For DeFi applications that need private trading, private lending, or private yield farming, private smart contracts (0.916) are the only model that provides sufficient generality. Aztec's noir programming language and Aleo's Leo both make private DeFi development practical, though prover times of 8-15 seconds per operation limit the UX for latency-sensitive applications. For identity and compliance -- the fastest-growing privacy use case in regulated Web3 -- ZK credential proving (0.892) is the clear winner. It provides the strongest GDPR compatibility of any model because no personal data ever touches the blockchain, and the user maintains full control over what is disclosed to whom. For governance privacy, MACI-style private voting (0.856) prevents vote buying and coercion by hiding individual votes while maintaining verifiable tallies.

5.2 The Trusted Setup Dilemma

Groth16 consistently produces the smallest proofs and fastest verification, making it the most gas-efficient proof system for on-chain verification. But it requires a per-circuit trusted setup ceremony where a secret value ("toxic waste") must be generated and destroyed. If any participant in the ceremony retains the secret, they can forge proofs. Zcash ran elaborate multi-party ceremonies to mitigate this risk. For enterprise applications, the governance overhead of organising trusted setup ceremonies is significant. STARKs eliminate the trusted setup entirely -- their security relies only on hash function collision resistance, which is post-quantum secure. But STARKs produce 40-200 KB proofs that cost 1.2M+ gas to verify on Ethereum -- roughly 5 times the cost of Groth16 verification. Halo2 offers a compromise: no trusted setup (recursive verification eliminates it) with moderate proof sizes (5-10 KB, roughly 2x Groth16 gas cost). Our recommendation: Groth16 for high-throughput applications where gas cost dominates and the trusted setup is manageable. Halo2 for applications where the trusted setup is unacceptable but gas efficiency matters. STARKs for applications where post-quantum security is required or where proofs are verified off-chain (where proof size does not translate to gas cost).

6. Conclusion

6.1 What We Found

Private smart contracts achieve the highest PMQS (0.916) through programmable privacy covering all use cases. ZK credential proving achieves the strongest GDPR compatibility (0.960) because no personal data touches the chain. Groth16 provides the most gas-efficient proofs; STARKs eliminate trusted setup with a post-quantum bonus; Halo2 balances both. Prover time of 8-15 seconds for complex operations is the current UX bottleneck, improving roughly 2x annually with hardware and algorithmic advances.

6.2 What We Are Releasing

The ZKPAF evaluation toolkit, all six privacy model implementations across four proof systems, PMQS calculator, privacy audit methodology (metadata leakage analysis), GDPR compatibility rubric, proof system comparison benchmark, and a privacy model selection guide are available at zkpaf-privacy.org under Apache 2.0.

References

Aleo (2024). Aleo Documentation. developer.aleo.org.

Aztec Network (2024). Aztec Documentation. docs.aztec.network.

Ben-Sasson, E. et al. (2018). Scalable, transparent, and post-quantum secure computational integrity. IACR ePrint.

Bowe, S. et al. (2020). Halo: Recursive proof composition without a trusted setup. IACR ePrint.

Bunz, B. et al. (2020). Zether: Towards privacy in a smart contract world. Financial Cryptography 2020.

Gabizon, A. et al. (2019). PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. IACR ePrint.

Groth, J. (2016). On the size of pairing-based non-interactive arguments. EUROCRYPT 2016.

Hansen, L. et al. (2025). User-centric design of Web3 applications for mass adoption. Blockchain, Web3 & Digital Trust Journal, 2025(3), 19-26.

Hopwood, D. et al. (2022). Zcash Protocol Specification. zips.z.cash.

Ivanov, J. et al. (2025). Decentralized storage and computing models for Web3 platforms. Blockchain, Web3 & Digital Trust Journal, 2025(3), 11-18.

Kothapalli, A. et al. (2022). Nova: Recursive zero-knowledge arguments from folding schemes. CRYPTO 2022.

MACI (2024). Minimum Anti-Collusion Infrastructure. maci.pse.dev.

Popescu, O. et al. (2025). Decentralized identity management systems for Web3 ecosystems. Blockchain, Web3 & Digital Trust Journal, 2025(2), 47-56.

Polygon Miden (2024). Polygon Miden Documentation. docs.polygon.technology/miden.

Sasson, E. B. et al. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. IEEE S&P; 2014.

Semaphore (2024). Semaphore Protocol. semaphore.pse.dev.

Sismo (2024). Sismo Protocol Documentation. docs.sismo.io.

European Commission (2016). General Data Protection Regulation (GDPR).

Werner, S. et al. (2022). SoK: Decentralized finance (DeFi). IEEE S&P; 2022.

Zcash Foundation (2024). Zcash Documentation. zcash.readthedocs.io.

Declarations

Funding

Supported by the German Research Foundation (DFG, grant PO 7337/1-1), the Austrian Science Fund (FWF, grant P 52024-N), and the Swedish Research Council (Vetenskapsradet, grant 2025-10024). No funder influenced results.

ISSN: e-ISSN: 3117-5988 / p-ISSN: 3117-597X Vol: 2025, Issues: 3, 2025

Conflict of Interest

No competing interests. No ZK proof system or privacy protocol vendor funded this work.

Data Availability Statement

Six privacy model implementations, four proof system benchmarks, PMQS calculator, privacy audit toolkit, GDPR rubric, and selection guide at zkpaf-privacy.org under Apache 2.0.

Ethical Approval

No personal data used. All transactions on testnets with synthetic data. Ethical exemption: European Institute of AI Ethics Board (ref. EIAI-2025-ETH-0624).

ISSN: e-ISSN: 3117-5988 / p-ISSN: 3117-597X Vol: 2025, Issues: 3, 2025

Appendix A

Proof System Benchmarks and PMQS Calculation

Detailed proof system measurements and PMQS scoring methodology.

Proof System Benchmark Configuration

PMQS Calculation