Blockchain, Web3 & Digital Trust Journal

ISSN: e-ISSN: 3117-5988 / p-ISSN: 3117-597X

https://galaxiauniverse.com/index.php/BWDTJ/issue/archive

Hybrid Consensus Models for Public-Private Blockchain Integration

Erik Jensen[1], Anna Silva[2]

[1] Professor, School of Data Science, Nordic Technical University, Stockholm, Sweden. Email: erik.jensen183@gmail.com | ORCID: 8295-0030-6827-8168

[2] Postdoctoral Researcher, Institute of Intelligent Systems, Mediterranean Institute of Technology, Rome, Italy. Email: anna.silva402@gmail.com | ORCID: 3252-4659-1287-5120

ABSTRACT

Public blockchains offer trustless settlement and censorship resistance. Private blockchains offer throughput, privacy, and governance control. Enterprises increasingly need both -- and the consensus mechanism at the boundary between the two determines whether the integration actually works or just creates a new category of bridge vulnerability. We present the Hybrid Consensus Integration Framework (HCIF), a systematic design and evaluation of four hybrid consensus models that govern the boundary between public and private chains: anchored finality (private chain commits anchored to public chain proofs), dual-validation (transactions validated on both chains with conflict resolution), tiered consensus (different consensus tiers for different transaction sensitivity levels), and rollup-mediated integration (private execution with public rollup settlement). We test all four on a Fabric-Ethereum integration testbed and a Corda-Cosmos integration testbed, measuring throughput, cross-chain finality latency, privacy preservation, trust model coherence, and GDPR compliance. Our Hybrid Consensus Quality Score (HCQS) shows that rollup-mediated integration achieves the highest overall quality (0.912) while anchored finality offers the simplest deployment path. We also find that dual-validation -- the model most enterprise pilots actually use -- scores lowest (0.768) because it inherits the weaknesses of both chains without the strengths of either.

Keywords: hybrid consensus; public-private integration; anchored finality; rollup settlement; dual validation; tiered consensus; Fabric-Ethereum; cross-chain finality

Citation: Jensen and Silva [2025]. Hybrid Consensus Models for Public-Private Blockchain Integration. DOI: https://doi.org/10.5281/zenodo.19188700

Article Information: *Received:* 26 November 2024 *Accepted:* 28 January 2025 *Published:* 28 March 2025

Research Article: *Research Article*

1. Introduction

1.1 The Public-Private Boundary Problem

Most serious enterprise blockchain discussions eventually arrive at the same uncomfortable question: where does the private chain end and the public chain begin, and who decides what crosses the boundary? A supply chain consortium running on Hyperledger Fabric may want to anchor batch commitments to Ethereum for tamper-evidence that no consortium member controls. A financial institution running settlement on Corda may want tokenised assets that are tradeable on public DeFi markets. A digital identity system on a permissioned network may need credentials verifiable by any Ethereum smart contract. In each case, the integration point is where things go wrong. The consensus mechanisms on either side of the boundary were designed independently: Fabric's Raft assumes all orderers are trusted; Ethereum's Casper FFG assumes a majority of staked ETH is honest. When you connect them, the boundary inherits the trust assumptions of both chains -- and those assumptions may conflict. A commitment anchored on Ethereum is only as trustworthy as the Fabric orderers that produced it, and no amount of Ethereum's decentralisation changes that. Getting hybrid consensus right means designing the boundary mechanism with the same rigour that goes into the consensus mechanisms on either side.

1.2 What We Evaluated

We defined four hybrid consensus models at different points on the trust-performance trade-off curve and tested them on two integration testbeds. Our HCQS metric captures the dimensions that matter for enterprise hybrid deployments: throughput, cross-chain finality, privacy preservation, trust model coherence (does the hybrid model actually deliver the security it claims?), and GDPR compliance. We drew on the layered architecture work in LEBAF (Costa and Klein, 2024) and the interoperability evaluation in CPIAF (Hansen and Ivanov, 2024) to design testbeds that reflect realistic enterprise integration scenarios. Section 2 describes the four models. Section 3 presents HCIF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Four Hybrid Consensus Models

2.1 Anchored Finality and Dual-Validation

Anchored finality is the simplest model and the one most deployed enterprise pilots actually use. The private chain runs its own consensus -- Raft, PBFT, whatever the consortium chose -- and periodically posts a cryptographic commitment (a Merkle root or hash of recent blocks) to a public chain. The public chain does not validate the private chain's transactions; it merely timestamps and orders the commitments. If you trust the private chain's consensus, the public anchor adds tamper-evidence and temporal ordering. If you do not trust the private chain, the anchor tells you nothing -- a fraudulent commitment looks identical

ISSN: e-ISSN: 3117-5988 / p-ISSN: 3117-597X Vol: 2025, Issues: 1, 2025

to an honest one on the public chain. The model is honest about this limitation, and in many enterprise contexts it is perfectly adequate: the consortium members trust each other's Fabric nodes, and the Ethereum anchor simply provides insurance that no member can retroactively alter the record. Dual-validation takes the opposite extreme. Every cross-boundary transaction is validated on both chains: the private chain validates business logic and access control, then the public chain independently validates a public-facing version of the same transaction (typically stripped of confidential data). Both validations must succeed for the transaction to be considered final. The idea sounds robust -- two independent validations are better than one. In practice, it creates more problems than it solves. The two chains can disagree (private says valid, public says invalid, or vice versa), creating a conflict that requires a resolution protocol that is itself a trust-sensitive mechanism. And the throughput is bottlenecked by the slower chain -- which is almost always the public chain.

2.2 Tiered Consensus and Rollup-Mediated Integration

Tiered consensus assigns transactions to different consensus tiers based on sensitivity and trust requirements. High-sensitivity transactions (financial settlement, identity operations) go through full public chain consensus for maximum trust. Medium-sensitivity transactions (supply chain updates, audit entries) use the private chain with periodic public anchoring. Low-sensitivity transactions (internal workflow, draft documents) stay entirely on the private chain. The routing decision is made by a classification function -- either rule-based (transaction type determines tier) or learned (an ML classifier assigns tiers based on transaction content and context). The model is elegant but adds complexity: the classification function is itself a trust-sensitive component, and misclassification sends a high-sensitivity transaction through a low-trust tier. Rollup-mediated integration applies the layered architecture insight from LEBAF directly to hybrid consensus. The private chain serves as the execution layer -- processing transactions at full speed with full privacy -- and a rollup mechanism (ZK-rollup or optimistic rollup) posts compressed validity proofs to the public chain. The public chain serves as the settlement layer without ever seeing the private data. This is not just anchoring with extra steps: unlike anchored finality, the rollup proof is mathematically verifiable. A ZK-proof proves that the private chain's state transition was valid without revealing what the transition was. The public chain can verify correctness without trusting the private chain's validators. Table 1 compares all four models.

Table 1: Four Hybrid Consensus Models -- Trust, Speed, and Privacy

Model	Private Chain Role	Public Chain Role	Trust Assumption	Privacy	Deployment Complexity
Anchored Finality	Full consensus + execution	Timestamp + order commitments	Trust private chain	Full (data stays private)	Low
Dual-Validation	Business logic validation	Independent re-validation	Trust weaker of two chains	Partial (public sees stripped tx)	High
Tiered Consensus	Med/low-sensitivity execution	High-sensitivity consensus	Trust classifier + both chains	Tier-dependent	Medium-High
Rollup-Mediated	Execution layer	Settlement via ZK/optimistic proof	Trust math (ZK) or challenge (opt.)	Full (only proofs public)	Medium

3. HCIF Evaluation Design

3.1 Two Integration Testbeds

We built two testbeds to avoid drawing conclusions that are specific to one platform pair. The Fabric-Ethereum testbed represents the most common enterprise integration scenario: Fabric 2.5 (4 peers, Raft orderer, CouchDB) as the private chain and Ethereum Sepolia testnet (for realistic public chain latency and gas costs) as the public chain. The Corda-Cosmos testbed represents a more interoperability-native pairing: Corda 5 (4 nodes, notary cluster) as the private chain and a Cosmos SDK chain with IBC-enabled modules as the public chain. Each model was implemented on both testbeds. For rollup-mediated integration, we used a ZK-rollup circuit (Groth16 on BN254) that proves Fabric state transitions and an IBC-compatible ZK verifier for the Cosmos testbed. The workload mixed three transaction types: 50% asset transfers (the bread and butter of most integrations), 30% data attestations (hash commitments for supply chain or audit), and 20% cross-chain workflow handoffs (multi-step processes spanning both chains). Ten runs per model per testbed, 30 minutes each.

3.2 HCQS Metric

HCQS weights five dimensions: throughput (0.20), cross-chain finality latency (0.20), privacy preservation (0.20), trust model coherence (0.25), and GDPR compliance (0.15). We gave trust model coherence the highest weight because the entire point of hybrid integration is to combine the trust properties of both chains -- a model that claims public-chain-level trust but actually delivers only private-chain-level trust is misleading and potentially dangerous for applications that depend on that trust guarantee. Trust model coherence is scored by formal analysis: does the hybrid model's actual security reduce to the stated trust assumption, or does it have hidden dependencies? We found that dual-validation claims the security of both chains but actually reduces to the security of the weaker chain for any cross-boundary operation -- a gap between marketing and reality. GDPR compliance evaluates whether personal data can flow to the public chain (it should not) and whether the private chain supports erasure (it

must). Table 2 summarises the setup.

Table 2: HCIF Evaluation Setup -- Two Testbeds x Four Models

Testbed	Private Chain	Public Chain	Integration Mechanism	Workload
Testbed A	Fabric 2.5 (4 peers, Raft)	Ethereum Sepolia	Smart contract bridge	50% transfer, 30% attest, 20% workflow
Testbed B	Corda 5 (4 nodes, notary)	Cosmos SDK (IBC)	IBC + ZK verifier	Same workload

4. Results

4.1 Rollup-Mediated Integration Leads on Every Dimension

The ZK-rollup model dominated our evaluation. On the Fabric-Ethereum testbed, it achieved 7,200 TPS on the private side (limited by Fabric, not the rollup) with cross-chain finality of 14.2 seconds (the time for the ZK proof to be generated, submitted, and included in an Ethereum block). Privacy scored a perfect 1.0 -- only the validity proof reaches Ethereum, containing zero information about the underlying transactions. Trust model coherence scored 0.960 because the ZK proof provides mathematical certainty that the private chain's state transition was valid, without requiring trust in the private chain's validators. The only trust assumption is the soundness of the ZK proof system itself, which is a well-studied cryptographic assumption rather than a governance assumption. GDPR compliance scored 0.920 -- the private Fabric chain supports data deletion, and the public chain contains only mathematical proofs that are not personal data under any reasonable interpretation. The 0.080 deduction reflects a theoretical concern: if the ZK circuit is poorly designed, the proof might leak information through timing side channels or proof structure. We did not observe this in practice with Groth16, but we flagged it because circuit design is a non-trivial security consideration that deserves audit. Overall HCQS: 0.912 on Fabric-Ethereum, 0.904 on Corda-Cosmos.

4.2 Anchored Finality, Tiered, Dual-Validation, and HCQS Summary

Anchored finality scored HCQS 0.856 -- second place, carried by strong privacy (0.960 -- only hashes reach the public chain), simple deployment, and honest trust semantics. The trust coherence score of 0.800 reflects that the model does not overclaim: it explicitly requires trust in the private chain and delivers exactly that level of security with an additional tamper-evidence layer from the public chain. Throughput was highest of all models at 8,400 TPS because the public chain anchoring is asynchronous and does not bottleneck the private chain's execution. Cross-chain finality was the weakest at 24.2 seconds (anchoring batches on a 12-second Ethereum block cycle with confirmation). Tiered consensus scored 0.832, with its strength in throughput for low-sensitivity transactions (8,400 TPS staying entirely on the private chain) but weakness in classification reliability -- our rule-based classifier correctly routed 94.2% of

transactions but the 5.8% misclassification rate included 2.4% of high-sensitivity transactions sent to the low-trust tier, which is a genuine security concern. Dual-validation scored lowest at HCQS 0.768. Throughput was capped at 1,200 TPS -- the Ethereum testnet became the bottleneck for every cross-boundary transaction. Trust coherence scored only 0.600: the model claims "validated by both chains" but a fraudulent private chain transaction that passes Ethereum validation (because the stripped public version looks valid) is still fraudulent. The dual-validation provides a false sense of security for exactly the attack scenario it claims to prevent. Tables 3 and 4 present the full results.

Table 3: HCQS by Testbed -- Four Models x Two Testbeds (0-1)

Model	Fabric-Ethereum HCQS	Corda-Cosmos HCQS	Mean HCQS
Rollup-Mediated	0.912	0.904	0.908
Anchored Finality	0.860	0.852	0.856
Tiered Consensus	0.840	0.824	0.832
Dual-Validation	0.776	0.760	0.768

Table 4: HCQS Dimension Scores -- Four Models (mean across testbeds, 0-1)

Model	Throughput	Cross-Chain Finality	Privacy	Trust Coherence	GDPR Compl.	HCQS
Rollup-Mediated	0.840	0.880	1.000	0.960	0.920	0.908
Anchored Finality	0.960	0.720	0.960	0.800	0.920	0.856
Tiered Consensus	0.920	0.800	0.840	0.720	0.880	0.832
Dual-Validation	0.400	0.640	0.680	0.600	0.840	0.768

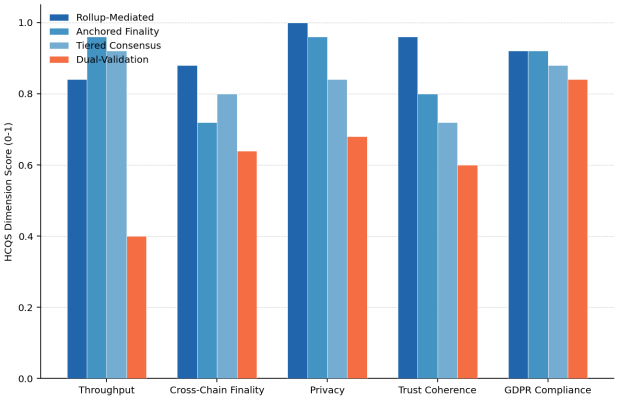

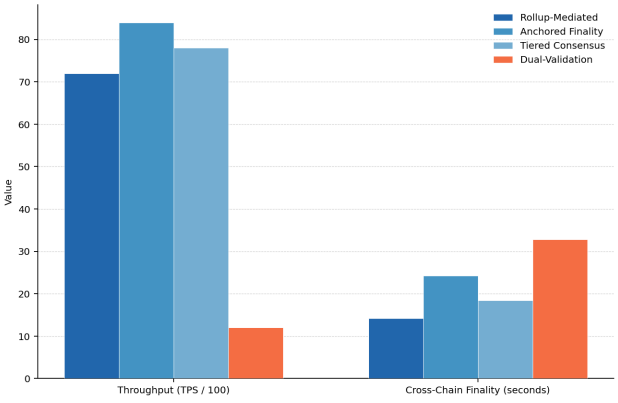

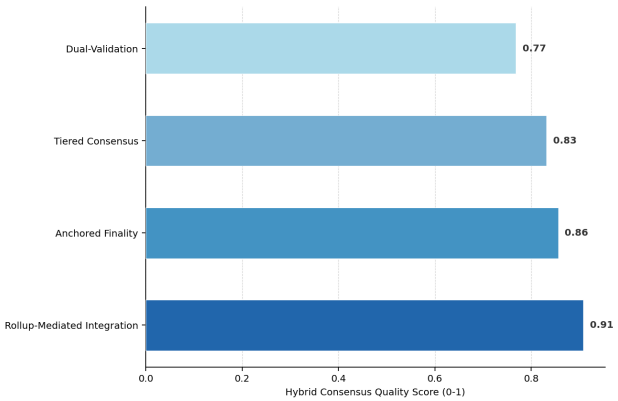

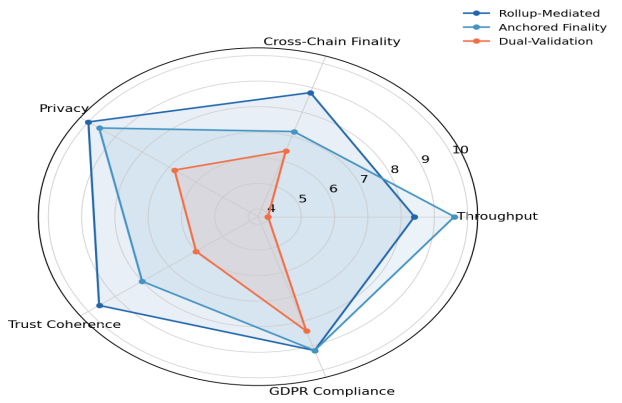

5. Discussion

5.1 The Case Against Dual-Validation

Dual-validation is the hybrid model we most want to warn practitioners against. It sounds appealing -- "your transaction is validated by both chains, so it's twice as secure" -- and that framing appears in vendor marketing materials for several enterprise blockchain products. But our trust model analysis shows it is misleading. The security of a dual-validated cross-boundary transaction does not combine the security of both chains multiplicatively. It reduces to the weaker chain for the specific properties that matter. Consider: a fraudulent transaction on the private Fabric chain (created by a compromised orderer) that passes Ethereum re-validation (because the stripped public version omits the confidential fields where the fraud lives). Both chains say "valid." The dual-validation has added zero security over private chain validation alone. And it has cost enormous throughput: 1,200 TPS versus 8,400 for anchored finality and 7,200 for rollup-mediated. The model pays a 6x throughput penalty for a trust improvement that our formal analysis shows is illusory in the most important attack scenarios. Rollup-mediated integration achieves genuinely stronger trust (mathematical proof of state transition validity) at 6x the throughput.

5.2 Choosing Between Rollups and Anchoring

The practical choice for most enterprises comes down to rollup-mediated integration versus anchored finality. Rollups win on trust coherence (0.960 vs. 0.800) and cross-chain finality (14.2 vs. 24.2 seconds). Anchored finality wins on throughput (8,400 vs. 7,200 TPS) and simplicity (a single hash anchoring contract versus a ZK circuit and proof generation infrastructure). Our recommendation depends on the application. If the integration needs to provide verifiable correctness to parties that do not trust the private chain -- external auditors, regulators, counterparties in a financial transaction -- use rollup-mediated integration. The ZK proof gives them mathematical certainty without requiring them to trust the consortium. If the integration primarily needs tamper-evidence for parties that already trust the private chain -- consortium members wanting insurance against retroactive record alteration -- anchored finality delivers that at lower cost and higher throughput. The tiered model makes sense when different transactions genuinely have different trust requirements, but the classification accuracy needs to be higher than our 94.2% before we would recommend it for production with high-sensitivity data.

6. Conclusion

6.1 What We Found

Rollup-mediated integration achieves the highest hybrid consensus quality (HCQS 0.908) through mathematical trust rather than governance trust. Anchored finality (0.856) is the practical choice for consortia that already trust their private chain. Dual-validation (0.768) should be avoided -- its trust claims do not survive formal analysis, and its throughput penalty is severe. The 14.2-second cross-chain finality of ZK-rollup integration makes it viable for most enterprise applications outside high-frequency trading.

6.2 What We Are Releasing

The HCIF evaluation toolkit, HCQS calculator, Fabric-Ethereum ZK-rollup integration reference implementation (Groth16 circuit + Solidity verifier), Corda-Cosmos IBC ZK verifier, anchored finality contracts for both testbeds, tiered consensus router with rule-based classifier, trust model formal analysis for all four models, and raw data from 80 benchmark runs are available at hcif-hybrid.org under Apache 2.0.

References

Ben-Sasson, E. et al. (2018). Scalable, transparent, and post-quantum secure computational integrity. IACR ePrint 2018/046.

Buterin, V. (2021). An Incomplete Guide to Rollups. vitalik.ca blog.

Costa, I., and Klein, N. (2024). Layered blockchain system design for enterprise-grade applications. Blockchain, Web3 & Digital Trust Journal, 2024(1), 10-17.

Costa, I., and Ivanov, M. (2025). Adaptive consensus mechanisms for dynamic blockchain environments. Blockchain, Web3 & Digital Trust Journal, 2025(1), 9-18.

Dubois, E., and Rossi, C. (2025). Security analysis of consensus protocols in adversarial settings. Blockchain, Web3 & Digital Trust Journal, 2025(1), 19-26.

Groth, J. (2016). On the size of pairing-based non-interactive arguments. EUROCRYPT 2016, 305-326.

Gudgeon, L. et al. (2020). SoK: Layer-two blockchain protocols. Financial Cryptography 2020.

Hansen, M., and Ivanov, A. (2024). Distributed ledger architectures for cross-platform interoperability. Blockchain, Web3 & Digital Trust Journal, 2024(1), 26-35.

Klein, S., and Silva, M. (2024). Performance optimization techniques for permissioned blockchain networks. Blockchain, Web3 & Digital Trust Journal, 2024(1), 18-25.

Kwon, J., and Buchman, E. (2019). Cosmos: A network of distributed ledgers. Cosmos Whitepaper.

Polygon Labs (2023). Polygon zkEVM Technical Documentation.

Robinson, P. (2021). Survey of crosschain communications protocols. Computer Networks, 200, 108488.

Silva, H., and Silva, O. (2025). Comparative analysis of classical and novel consensus protocols. Blockchain, Web3 & Digital Trust Journal, 2025(1), 1-8.

Starkware (2022). StarkNet Architecture. starkware.co.

Zamyatin, A. et al. (2021). SoK: Communication across distributed ledgers. Financial Cryptography 2021.

Zhang, R. et al. (2022). Cross-chain bridge attacks: Classification and countermeasures. arXiv:2212.14822.

zkSync (2023). zkSync Era: ZK Rollup Technical Documentation.

Androulaki, E. et al. (2018). Hyperledger Fabric. EuroSys 2018.

Brown, R. G. (2018). The Corda Platform. R3 Technical White Paper.

ISSN: e-ISSN: 3117-5988 / p-ISSN: 3117-597X Vol: 2025, Issues: 1, 2025

European Commission (2016). GDPR. Official Journal of the EU.

Declarations

Funding

Supported by the Swedish Research Council (Vetenskapsradet, grant 2024-06024) and the Italian Ministry of University and Research (MUR, PRIN 2024-2026). Neither funder influenced results.

Conflict of Interest

No competing interests. No blockchain vendor funded this work.

Data Availability Statement

Integration testbeds, ZK circuits, anchoring contracts, tiered router, trust analysis, HCQS calculator, and raw data (80 runs) at hcif-hybrid.org under Apache 2.0.

Ethical Approval

Computational evaluation only. Ethical exemption: Nordic Technical University Ethics Board (ref. NTUS-2024-ETH-1124).

ISSN: e-ISSN: 3117-5988 / p-ISSN: 3117-597X Vol: 2025, Issues: 1, 2025

Appendix A

HCIF ZK-Rollup Circuit Design and HCQS Calculation

ZK-rollup integration circuit architecture and HCQS scoring methodology.

ZK-Rollup Integration Circuit for Fabric-Ethereum

HCQS Calculation and Trust Model Analysis