

SCALABLE BLOCKCHAIN ARCHITECTURE FOR SECURE AND EFFICIENT SUPPLY CHAIN TRANSACTIONS

Yandamuri Shasikala¹, Dr. Alok Kumar Gupta²

Research Scholar, Department of Engineering & Technology, GLOCAL University,
Saharanpur, Uttar Pradesh, India¹

Supervisor, Department of Engineering & Technology, GLOCAL University, Saharanpur,
Uttar Pradesh, India²

ABSTRACT

Modern supply networks have widened the gap between the traceability that regulators and consumers demand and the visibility that conventional enterprise systems provide. Centralised enterprise resource planning ledgers are fast and inexpensive but constitute a single point of compromise; first-generation public blockchains deliver immutability at the cost of throughput, latency, and energy. This paper presents the design and empirical evaluation of ChainSCM, a secure and scalable permissioned blockchain framework that integrates echelon-based dynamic sharding, a Proof-of-Traceable-Authority consensus that couples proof-of-authority leader election with Byzantine-fault-tolerant committee finality, content-addressed off-chain payload storage with on-chain hash anchoring, and attribute-based access control. Three empirical instruments were used: a thirty-day instrumented pilot on a three-echelon pharmaceutical and agri-food testbed that processed 88.1 million events across four deployment arms; a survey of 180 supply-chain professionals drawn from 62 organisations, a 72.0 percent response rate; and a security stress test of 1200 injected events per attack vector per arm. The proposed framework sustained 3456 transactions per second at sixty-four validator nodes, three times the rate of a Fabric-style permissioned baseline, while committing in 0.465 s and growing the ledger by only 2.11 GB per day. Tampering was detected at 99.6 percent, above the 98.5 percent of the permissioned baseline. Perceived traceability, security, scalability, and interoperability jointly explained 53.2 percent of the variance in adoption intention, with traceability strongest. Taken together, the results demonstrate that the classical security-scalability trade-off can be substantially relaxed, although interoperability with legacy systems, cited by 71.0 percent of organisations, remains the binding barrier.

Keywords: *Blockchain¹, dynamic sharing², permissioned distributed ledger³, proof-of-traceable-authority⁴, smart contracts⁵, supply chain traceability⁶, system scalability⁷.*

1. INTRODUCTION

A. BACKGROUND AND MOTIVATION

Global supply chains are extensive, multi-tier, multi-jurisdictional networks in which a product may cross dozens of organisational boundaries before reaching a consumer. Each boundary introduces an information discontinuity: records are held in proprietary enterprise resource planning and warehouse management systems that were never designed to exchange verifiable evidence with trading partners [1], [2]. The consequences are well documented, from counterfeit pharmaceuticals and mislabelled food to unverifiable sustainability claims and recalls that take days rather than minutes [3], [4]. Blockchain technology offers a shared, append-only, tamper-evident ledger whose entries are agreed by consensus rather than asserted by a single authority [5], [6]. Systematic reviews report growing interest across logistics, transport, and manufacturing contexts [5], [7], and conceptual frameworks now link adoption to measurable operational outcomes [8], [9].

B. PROBLEM STATEMENT AND RESEARCH GAP

Two limitations explain why pilots rarely reach production. The first is technical. Public proof-of-work chains offer strong tamper evidence but settle in minutes and consume energy that high-volume logging cannot justify [14], [16]; permissioned platforms such as Hyperledger Fabric improve latency yet degrade as the validator set grows, because endorsements must reach every peer [18], [19]. Sharding, which partitions state and validation across committees, restores scalability in public settings [23], [24], but its use in permissioned networks, where the natural partition is the supply chain echelon rather than a random validator assignment, is underexplored. The second is empirical: the literature is dominated by architectural proposals and conceptual frameworks [6], [7], and existing benchmarks typically use generic workloads rather than supply chain event streams [22]. Few report instrumented measurement of throughput, latency, storage growth, and attack resistance for one system under a common workload.

C. OBJECTIVES AND CONTRIBUTIONS

This study therefore asks whether a framework that combines echelon sharding, hybrid on-chain and off-chain storage, and a lightweight authority-based consensus can deliver the security properties expected of a distributed ledger while sustaining the throughput that operational supply chains require. The paper makes four contributions. First, it specifies the ChainSCM architecture and its Proof-of-Traceable-Authority consensus protocol. Second, it reports a thirty-day instrumented pilot benchmarking the framework against three reference deployments under one workload. Third, it combines these measurements with a practitioner survey to test whether technical gains translate into adoption intention. Fourth, it quantifies the residual barriers that separate technical feasibility from deployment. Section II surveys related work, Section III describes the design and methodology, Section IV presents the data and analysis, Section V discusses the findings against prior work, and Section VI concludes.

2. LITERATURE SURVEY

Research on blockchain in supply chain management has evolved along three streams. The first examines application value. Early work established the conceptual case for distributed ledgers in logistics and

manufacturing, identifying transparency, traceability, and disintermediation as the principal sources of value [1], [5], [6]. The proposition is attractive because it reframes traceability as a property of the network rather than of any single firm's records. Design-science studies subsequently demonstrated traceability systems for agri-food [12], [13] and pharmaceutical [9], [11] contexts, and syntheses of this literature conclude that benefits are contingent on inter-organisational collaboration rather than on the ledger alone [4], [7]. A parallel body of work applies institutional, contingency, and sensemaking theory to explain why firms experiment with the technology and what they expect from it [8], [10]. The second stream addresses the performance problem. Foundational studies of consensus show that crash-tolerant and Byzantine-tolerant protocols behave very differently as committee size grows [20], [21], and systematic benchmarking confirmed that no design optimises throughput, latency, and fault tolerance simultaneously [22]. The implication drawn by most subsequent work is that consensus must be simplified before it can be scaled. Two responses dominate.

The first replaces probabilistic finality with authority-based or stake-based leader election [25], [26]; the second partitions the ledger through sharding so that validation effort grows sub-linearly with participation [23], [24]. Both were developed for public, adversarial settings and assume validators can be assigned to committees at random, an assumption that fails when validators are known trading partners with materially different roles. The third stream concerns security, privacy, and data management. Because commercial transactions cannot be published in the clear, designers combine permissioned membership services with selective disclosure and move bulky payloads off the ledger while retaining cryptographic anchors on-chain [15], [27], [28]. Access control, key management, and the tension between immutability and privacy are consistently identified as the hardest problems in this stream [17], [29]. Finally, a substantial body of work documents adoption barriers, including interoperability, cost, skills, governance, and regulatory uncertainty [1], [2], [9]. Most of these studies rely on expert opinion or case narratives rather than instrumented measurement, and the few that report performance do so for a single architecture. What remains largely absent is an empirical study that measures a sharded, permissioned design under a realistic supply chain workload and relates measured performance to the perceptions of the practitioners who would have to operate it. This paper addresses that gap.

3. METHODOLOGY

This study adopted a mixed-methods design in which technical and perceptual measurement were treated as complementary evidence about one framework [29]. An instrumented pilot generated objective performance data, a structured survey captured practitioner judgement, and a controlled stress test probed security; the three were analysed separately and triangulated. The pilot followed a repeated-measures design in which four arms executed an identical workload for thirty consecutive days, so each day acted as a blocking factor. The arms were a centralised enterprise resource planning ledger, a proof-of-work public chain, a Fabric-style permissioned baseline, and the proposed framework. The framework was implemented as five layers, as shown in Fig. 1. Layer one ingests EPCIS 2.0 events, sensor telemetry, and enterprise extracts through adapters that attach hardware-rooted device identities. Layer two groups' validators into echelon shards covering manufacturing, logistics, and retail, and runs Proof-of-Traceable-Authority, in which a proof-of-authority schedule elects a leader that proposes an ordered batch while a rotating Byzantine-fault-tolerant committee certifies finality; cross-shard transfers commit atomically. Layer three stores payloads as encrypted, content-

addressed objects and writes only hashes and Merkle roots to the ledger, with attribute-based encryption governing disclosure. Layer four hosts provenance and compliance contracts, and layer five exposes dashboards and regulator portals. Survey instruments used five-point Likert scales with 4 items per construct, adapted from validated instruments and pre-tested on twenty-two practitioners [29]. Composite scores averaged the items, and analysis used three stages: reliability using Cronbach's alpha; bivariate association using Pearson correlations; and multivariate estimation using ordinary least squares regression with adoption intention as the criterion [29]. Pilot metrics were compared using one-way analysis of variance, Welch's t-tests for pairwise contrasts, and Cohen's d for effect magnitude [30]; stress-test counts used chi-square tests. Normality and homogeneity of variance were examined, and where violated nonparametric equivalents were computed; conclusions were unchanged.




Fig. 1. Layered architecture of the proposed ChainSCM framework, showing the five functional layers (L1-L5) and the cross-cutting security and governance services that span them

Figure 1 shows the five layers of the framework and the services that cut across them. Data enter at L1, where EPCIS 2.0 events, sensor telemetry, and enterprise extracts are signed with hardware-rooted device identities. L2 carries the scalability innovation: validators are grouped into echelon shards, and Proof-of-Traceable-Authority elects a leader per shard while a rotating committee certifies finality, so validation effort grows with shard size, not network size. L3 keeps only hashes and Merkle roots on-chain; L4 encodes business logic; L5 serves users and regulators. The lower bar shows that identity, audit, and anomaly detection are properties enforced everywhere.

4. DATA COLLECTION AND ANALYSIS

A. DATA SOURCES AND SAMPLE COMPOSITION

Data were collected over a four-month window through three instruments. The pilot testbed executed an identical workload on all four arms for 30 consecutive days, ingesting a mean of 2.94 million signed events per

day and 88.1 million events in total. The survey invited 250 practitioners and retained 180 complete responses, a response rate of 72.0 percent, drawn from 62 organisations across six industry sectors. The stress test injected 1200 adversarial events per vector per arm using a scripted adversary with full network access but no valid credentials. Table 1 summarizes the composition of the survey sample, which was deliberately weighted towards respondents with operational rather than purely academic responsibility for supply chain systems.

Table 1: Composition of the Survey Sample (N = 180)

Characteristic	Category	n	%
Job role	Supply chain / procurement manager	44	24.4
	Logistics and warehouse operations	39	21.7
	IT / systems architect	34	18.9
	Quality, compliance, and audit	29	16.1
	Operations director / general manager	21	11.7
	Other	13	7.2
Industry sector	Pharmaceuticals and healthcare	48	26.7
	Food and agricultural commodities	44	24.4
	Automotive and engineering	33	18.3
	Retail and fast-moving consumer goods	30	16.7
	Electronics	15	8.3
	Other	10	5.6
Organisation size (employees)	Fewer than 250	41	22.8
	250 - 999	52	28.9
	1,000 - 4,999	47	26.1
	5,000 or more	40	22.2
Professional experience	Under 5 years	46	25.6
	5 - 10 years	67	37.2
	11 - 15 years	43	23.9
	More than 15 years	24	13.3
Prior blockchain exposure	None	38	21.1
	Pilot or proof of concept	79	43.9
	Limited production use	41	22.8
	Scaled production use	22	12.2
Region	India	74	41.1
	Asia-Pacific (excluding India)	38	21.1

	Europe	36	20.0
	North America	22	12.2
	Middle East and Africa	10	5.6

Note. Percentages are computed on 180 valid responses and sum to 100 within each characteristic. Respondents were screened for involvement in supply chain systems decisions.

The sample is well suited to the research question. Nearly four fifths of respondents had already been exposed to blockchain in some form, so their evaluations reflect experience rather than novelty, and the distribution across organisation sizes avoids the large-enterprise bias that characterises much of the adoption literature [2]. Sector concentration in pharmaceuticals and food is appropriate because both are subject to traceability mandates and were represented in the pilot testbed.

B. PILOT PERFORMANCE MEASUREMENTS

Table 2 reports the pilot results. All five metrics differed significantly across arms, with very large effects. The proposed framework sustained 3456 transactions per second against 1138 for the Fabric-style baseline, a threefold improvement, and against 23.2 for the proof-of-work chain. Commit latency fell to 0.465 s, roughly one third of the permissioned baseline and one thirty-first of the proof-of-work chain, while remaining above the 0.290 s of the centralised ledger, which pays no coordination cost. Ledger growth was the most striking result: hybrid anchoring reduced it to 2.11 GB per day, 66.5 percent below the permissioned baseline and an order of magnitude below the proof-of-work chain. Operating cost followed the same pattern at USD 0.61 per thousand transactions.

Table 2: Thirty-Day Pilot Performance (Mean $\pm$ Sd) By Deployment Arm

Metric	T1 Centralised	T2 PoW chain	T3 Fabric- style	T4 Proposed	F(3,116)	η^2	d (T4- T3)
Sustained throughput (tx/s)	5204 $\pm$ 183	23 $\pm$ 2	1138 $\pm$ 109	3456 $\pm$ 145	9,733.4	0.996	+18.06
Commit latency (s)	0.290 $\pm$ 0.035	14.434 $\pm$ 2.198	1.351 $\pm$ 0.164	0.465 $\pm$ 0.075	1,168.3	0.968	-6.97
CPU utilisation (%)	39.9 $\pm$ 4.2	88.0 $\pm$ 5.9	63.6 $\pm$ 5.3	51.6 $\pm$ 4.0	522.2	0.931	-2.54
Ledger growth (GB/day)	3.70 $\pm$ 0.33	18.16 $\pm$ 0.92	6.30 $\pm$ 0.49	2.11 $\pm$ 0.20	5,173.6	0.993	-11.18
Operating cost (USD/1,000 tx)	0.43 $\pm$ 0.06	9.89 $\pm$ 0.70	1.75 $\pm$ 0.14	0.61 $\pm$ 0.08	4,756.1	0.992	-10.03

Note. All ANOVA comparisons are significant at $p < 0.001$. d is Cohen's d for the contrast between the proposed framework and the Fabric-style baseline; positive values favour T4 for throughput, negative values favour T4 for latency, CPU utilisation, storage, and cost.

Pairwise contrasts confirm that the differences are substantial rather than merely significant. Against the Fabric-style baseline, the proposed framework recorded Cohen's d of +18.06 for throughput, -6.97 for latency, and -11.18 for ledger growth, all far beyond the conventional threshold for a large effect [30]. A paired comparison of daily latency means returned $t(29) = -32.4$, $p < 0.001$, confirming that the advantage held on every day of the pilot rather than in aggregate only.




Fig. 2. Scalability of the proposed framework: sustained throughput as the validator set grows from sixteen to one hundred and twenty-eight nodes, compared with two reference deployments.

Figure 2 isolates the scalability claim by holding the workload constant while enlarging the validator set. The proof-of-work chain declines from an already negligible base, and the Fabric-style baseline loses about a third of its capacity between sixteen and one hundred and twenty-eight nodes as broadcast traffic grows. The proposed framework retains most of its sixteen-node throughput at one hundred and twenty-eight nodes, losing about thirteen percent, because a new organisation extends a shard rather than joining every consensus round. The shaded bands show the ranking is stable across repetitions and never overlaps, so the ordering is not an artefact of noise. The scaling experiment reported in Fig. 2 explains the mechanism behind these gains. Because a new organisation joins a shard rather than a single global consensus group, the marginal cost of participation falls: throughput declined by only 13.8 percent between sixteen and one hundred and twenty-eight nodes, against 34.7 percent for the permissioned baseline, while latency grew by only 0.24 s.




Fig. 3. Distribution of end-to-end commit latency for each deployment arm across the thirty-day pilot; the vertical axis uses a logarithmic scale.

Figure 3 complements the throughput evidence by showing latency dispersion, which matters operationally because a system is judged by its worst days rather than its average. The proof-of-work chain is both slow and highly variable, reflecting probabilistic finality. The Fabric-style baseline is stable but centred above one second, with excursions beyond two seconds that warehouse staff would notice. The proposed framework shows the tightest distribution of the three blockchain arms and sits close to the centralised ledger, which remains marginally faster because it incurs no coordination cost; that gap is the price of shared custody, and it is small.

C. MEASUREMENT MODEL, RELIABILITY, AND ASSOCIATION

Table 3 presents the measurement model. All five constructs exceeded the conventional reliability threshold, with Cronbach's alpha between 0.866 and 0.900. Respondents rated security highest (3.99) and interoperability lowest (3.31), a pattern that anticipates the barrier ranking discussed below. All bivariate correlations were positive and significant at $p < 0.001$, and the strongest association with adoption intention was traceability ($r = 0.598$), followed by security ($r = 0.580$), scalability ($r = 0.502$), and interoperability ($r = 0.420$).

Table 3: Descriptive Statistics, Reliability, and Inter-Construct Correlations

Construct	Items	Mean	SD	α	1	2	3	4	5
Perceived traceability (TRC)	4	3.89	0.59	0.892	1.00	0.458	0.396	0.317	0.598
Perceived security and trust (SEC)	4	3.99	0.54	0.879	0.458	1.00	0.435	0.297	0.580
Perceived scalability (SCA)	4	3.69	0.64	0.892	0.396	0.435	1.00	0.394	0.502
Interoperability (INT)	4	3.31	0.65	0.900	0.317	0.297	0.394	1.00	0.420
Adoption intention (ADI)	4	3.82	0.55	0.866	0.598	0.580	0.502	0.420	1.00

Note. Columns 1-5 correspond to TRC, SEC, SCA, INT, and ADI respectively. All off-diagonal correlations are significant at $p < 0.001$ (two-tailed, $n = 180$). Scale range is 1 to 5.

Fig. 5 visualises the strongest of these relationships. The correlation matrix also indicates that the four predictors are positively intercorrelated but not collinear; the highest predictor intercorrelation is 0.458, well below the level at which standard errors become unstable, so the regression estimates reported next can be interpreted separately. Respondents also ranked the obstacles to adoption: interoperability with legacy ERP/WMS (71.0 percent); upfront integration cost (64.5 percent); shortage of in-house blockchain skills (58.1 percent); unclear data-governance / liability rules (53.2 percent); partner-network onboarding resistance (46.8 percent); regulatory uncertainty on cross-border data (37.1 percent).




Fig. 4. Relationship between perceived traceability and adoption intention, with the ordinary least squares fit; each marker represents one respondent.

Figure 4 plots perceived traceability against adoption intention for all respondents and overlays the fitted regression line. The upward slope shows that practitioners who judge a system to deliver better traceability are also more willing to deploy it, which is the behavioural bridge between the engineering results in Tables II and V and the adoption outcome the framework targets. The scatter is deliberately wide: the correlation is moderate, not determinative, and several respondents with high traceability scores remain cautious adopters. That residual variance is why interoperability, cost, and governance appear as separate predictors in the regression model rather than as noise.

D. HYPOTHESIS TESTING

Table 4 reports the hypothesis tests. The regression model was significant overall, explaining 53.2 percent of the variance in adoption intention (adjusted $R^2 = 0.521$, $F(4, 175) = 49.64$, $p < 0.001$). All four hypotheses were supported. Traceability exerted the largest standardised effect ($\beta = 0.342$, $p < 0.001$), followed by security ($\beta = 0.302$, $p < 0.001$), scalability ($\beta = 0.175$, $p = 0.005$), and interoperability ($\beta = 0.153$, $p = 0.009$). The fifth hypothesis, that the proposed framework resists tampering more effectively than the permissioned baseline, was also supported ($\chi^2(1) = 7.42$, $p = 0.006$).

Table 4: Hypothesis Tests and Multiple Regression on Adoption Intention

H	Relationship / test	B	SE	β	t	p	Decision
H1	TRC $\rightarrow$ ADI	0.321	0.057	0.342	5.64	< 0.001	Supported
H2	SEC $\rightarrow$ ADI	0.307	0.063	0.302	4.90	< 0.001	Supported
H3	SCA $\rightarrow$ ADI	0.151	0.053	0.175	2.84	0.005	Supported
H4	INT $\rightarrow$ ADI	0.130	0.049	0.153	2.65	0.009	Supported
H5	Tamper detection T4 > T3	$\chi^2 = 7.42$	—	—	—	0.006	Supported
	Model fit	$R^2 = 0.532$	adj. $R^2 = 0.521$	F = 49.64	p < 0.001	n = 180	

Note. OLS regression with adoption intention as the criterion; B is the unstandardised coefficient, SE the standard error, and β the standardised coefficient. H5 is tested by a chi-square test of independence on the stress-test counts.

The unstandardized coefficients allow a practical reading: a one-point increase in perceived traceability is associated with a 0.321-point increase in adoption intention on the five-point scale, holding the other predictors constant. The ordering of effects is theoretically sensible and matches the qualitative emphasis of the traceability literature [9], [13], in which the ability to reconstruct provenance quickly is treated as the primary operational benefit that justifies the cost of shared infrastructure.

E. SECURITY STRESS TEST AND COMPARATIVE BENCHMARK

Table 5 and Fig. 4 report the security evaluation. The proposed framework detected or prevented at least 99.2 percent of injected events on every vector, meeting the 99 percent service-level target throughout. The largest margins over the Fabric-style baseline appeared in unauthorised read attempts (99.3 percent against 96.2 percent) and availability under denial of service (99.3 percent against 95.0 percent), both of which benefit from attribute-based encryption and from sharding, since a shard-level flood does not exhaust global capacity. The centralised ledger failed most severely on availability and Sybil resistance.

Table 5: Comparative Benchmark against Representative Deployment Architectures

Configuration	Architecture	Consensus	Throughput (tx/s)	Latency (s)	Ledger growth (GB/day)	Tamper detection (%)	Max effective nodes	Ref.
A. Centralised ERP ledger (pre-blockchain)	Centralised DB	n/a (single authority)	5,210	0.28	3.80	61.2	1	[legacy]
B. PoW public chain (Bitcoin/Ethereum-class)	Public, unsharded	Proof-of-Work	23	14.62	18.20	98.8	unbounded*	[14],[16]
C. Fabric-style permissioned baseline	Permissioned, unsharded	PBFT / Raft	1,148	1.31	6.40	98.5	64	[18],[19]
D. Sharded permissioned design	Permissioned, static shards	BFT + PoS	2,074	0.79	4.35	98.9	96	[23],[24]
E. IPFS-	Hybrid	PoA	862	1.86	1.95	97.6	32	[27],[28]

anchored traceability stack	on/off chain							8]
F. Proposed ChainSCM framework	Sharded + off-chain anchoring	PoTA (PoA + BFT finality)	3,412	0.44	2.10	99.6	128	this work

Note. All configurations were deployed on the same testbed under the same event workload; configurations A-E are reference implementations of designs reported in the cited literature. Max effective nodes is the largest validator set at which a configuration retained at least 90 percent of its sixteen-node throughput. * Subject to probabilistic finality.




Fig. 5. Security stress test results: percentage of injected adversarial events detected or prevented, by attack vector and deployment arm.

Figure 5 compares the four arms across five attack vectors, and the shape matters more than the averages. The centralised ledger fails on exactly the properties that motivate distributed ledgers: it is weakest against availability attacks and Sybil injection, because a single compromised authority can rewrite history. The two blockchain arms already perform well, so the scope for improvement is narrow and must be measured carefully, which is why the trial used twelve hundred events per vector. The proposed framework is the only arm meeting the ninety-nine percent target on every vector, and its advantage is concentrated in confidentiality and availability.

5. DISCUSSION

A. INTERPRETATION OF THE PRINCIPAL FINDINGS

Taken together, the results support a single argument: the security and scalability objectives of supply chain blockchain design are less in tension than the literature assumes, provided that partitioning follows the structure of the supply chain. Echelon sharding is not a generic scaling trick borrowed from public chains; it works here because custody transfers are overwhelmingly local, so most validation occurs within a shard and cross-shard commits are rare. Hybrid storage contributes the second gain, and it is the more decisive one economically, because ledger growth rather than consensus dominates the long-run cost of an operational chain. The survey evidence closes the loop: performance improvements translate into adoption intention primarily through perceived traceability, the benefit practitioners can observe directly. Yet the model leaves nearly half the variance in intention unexplained, and the barrier ranking shows why: interoperability with legacy systems, integration cost, and skills shortages are organisational constraints that no protocol change resolves. Third, the two mechanisms are complementary rather than interchangeable: sharding protects throughput as the network grows, whereas anchoring protects cost as history accumulates. The practical implication is that architects should treat adoption as a co-design problem in which the ledger and the integration layer are specified together.

B. CRITICAL ANALYSIS AND COMPARISON WITH PAST WORK

These findings must be read against prior work with some care. The throughput advantage over a Fabric-style baseline is a factor of three, substantial but far smaller than the order-of-magnitude gains claimed for sharded public chains such as OmniLedger and RapidChain [23], [24]. That gap is explained less by implementation quality than by workload: those systems shard randomly across anonymous validators and measure simple payment transfers, whereas endorsements bound a permissioned supply chain workload long before consensus does. Our result is therefore consistent with rather than superior to the public-chain literature, and it suggests that echelon partitioning yields its benefit by reducing broadcast rather than by parallelising cryptography. The absolute latency gain, 0.465 s against 1.351 s, is modest but material where warehouse scanning is gated on confirmation. Compared with design-science traceability systems such as the RFID and blockchain architecture of Tian [13] and its successors [12], which anchor every scanned event on-chain, our hybrid anchoring cuts storage by 67 percent while preserving the same tamper-evidence guarantee, a trade that only pays off at the volumes we measured. The barrier ranking also replicates earlier work: the barrier categories identified by Saberi et al. [1] reappear here, with inter-organisational obstacles, interoperability above all, outranking purely technical ones, and the critical success factors of Hastig and Sodhi [9] locate the same problem in collaboration and supply chain practice rather than in protocol design. A further caveat: our testbed modelled three echelons of one product family, so networks with shared suppliers would generate more cross-shard traffic than we measured. Finally, and most importantly, the caution of Wüst and Gervais [17] applies: a centralised ledger remained marginally faster in our own measurements, and the case for blockchain rests on verifiable shared custody among parties who do not trust one another, not on raw speed.

6. CONCLUSION

This paper presented ChainSCM, a secure and scalable blockchain framework for supply chain management, and evaluated it empirically through a thirty-day instrumented pilot, a survey of 180 practitioners, and a controlled security stress test. The framework sustained 3456 transactions per second at sixty-four validator nodes, committed in 0.465 s, limited ledger growth to 2.11 GB per day, and detected 99.6 percent of tampering attempts, outperforming a Fabric-style permissioned baseline on every dimension while retaining the tamper evidence of a distributed ledger. Perceived traceability, security, scalability, and interoperability together explained 53.2 percent of the variance in adoption intention, with traceability the strongest predictor. Echelon sharding and hybrid anchoring relaxed this trade-off but did not remove it: the centralised ledger remained marginally faster, and the advantage narrows as shards become interdependent. Two limitations bound these claims: the evaluation used one testbed and a scripted adversary, so governance behaviour under commercial incentive is untested, and the survey is cross-sectional, so causality cannot be established. Future work will extend the pilot to live multi-enterprise deployments, investigate adaptive shard reconfiguration under seasonal demand, and examine governance mechanisms through which interoperability with legacy enterprise systems can be achieved without reintroducing a trusted intermediary. A reproducibility package of workload generators and measurement scripts accompanies the study. Taken together, the results position echelon-sharded, hash-anchored permissioned ledgers as a practical middle path between centralised speed and distributed assurance, and they indicate where further engineering effort should be directed. These findings generalise most readily to echelon-structured, high-volume supply networks.

7. REFERENCES

- [1] S. Saberi, M. Kouhizadeh, J. Sarkis, and L. Shen, "Blockchain technology and its relationships to sustainable supply chain management," *Int. J. Prod. Res.*, vol. 57, no. 7, pp. 2117–2135, 2019, doi: 10.1080/00207543.2018.1533261.
- [2] S. Kamble, A. Gunasekaran, and H. Arha, "Understanding the blockchain technology adoption in supply chains—Indian context," *Int. J. Prod. Res.*, vol. 57, no. 7, pp. 2009–2033, 2019, doi: 10.1080/00207543.2018.1518610.
- [3] N. Kshetri, "Blockchain's roles in meeting key supply chain management objectives," *Int. J. Inf. Manag.*, vol. 39, pp. 80–89, 2018.
- [4] M. M. Queiroz, R. Telles, and S. H. Bonilla, "Blockchain and supply chain management integration: a systematic review of the literature," *Supply Chain Manag.*, vol. 25, no. 2, pp. 241–254, 2020.
- [5] M. Pournader, Y. Shi, S. Seuring, and S. C. L. Koh, "Blockchain applications in supply chains, transport and logistics: a systematic review of the literature," *Int. J. Prod. Res.*, vol. 58, no. 7, pp. 2063–2081, 2020.
- [6] H. Treiblmaier, "The impact of the blockchain on the supply chain: a theory-based research framework and a call for action," *Supply Chain Manag.*, vol. 23, no. 6, pp. 545–559, 2018.

- [7] R. Cole, M. Stevenson, and J. Aitken, "Blockchain technology: implications for operations and supply chain management," *Supply Chain Manag.*, vol. 24, no. 4, pp. 469–483, 2019.
- [8] Y. Wang, M. Singgih, J. Wang, and M. Rit, "Making sense of blockchain technology: how will it transform supply chains?," *Int. J. Prod. Econ.*, vol. 211, pp. 221–236, 2019.
- [9] G. M. Hastig and M. S. Sodhi, "Blockchain for supply chain traceability: business requirements and critical success factors," *Prod. Oper. Manag.*, vol. 29, no. 4, pp. 935–954, 2020.
- [10] S. Nandi, J. Sarkis, A. A. Hervani, and M. M. Helms, "Redesigning supply chains using blockchain-enabled circular economy and COVID-19 experiences," *Sustain. Prod. Consum.*, vol. 27, pp. 1–22, 2021.
- [11] A. Ghadge, A. Bourlakis, S. Kamble, and S. Seuring, "Blockchain implementation in pharmaceutical supply chains: a review and conceptual framework," *Int. J. Prod. Res.*, vol. 61, no. 19, pp. 6633–6651, 2023.
- [12] Y. Kayikci, N. Subramanian, M. Dora, and M. K. Singh, "Food supply chain in the era of Industry 4.0: blockchain technology implementation opportunities and impediments from the perspective of people, process, performance, and technology," *Prod. Plan. Control*, vol. 33, no. 2–3, pp. 301–321, 2022.
- [13] F. Tian, "An agri-food supply chain traceability system for China based on RFID & blockchain technology," in *Proc. 13th Int. Conf. Serv. Syst. Serv. Manag. (ICSSSM)*, Kunming, China, 2016, pp. 1–6, doi: 10.1109/ICSSSM.2016.7538424.
- [14] S. Nakamoto, "Bitcoin: a peer-to-peer electronic cash system," 2008. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [15] K. Christidis and M. Devetsikiotis, "Blockchains and smart contracts for the Internet of Things," *IEEE Access*, vol. 4, pp. 2292–2303, 2016.
- [16] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, "An overview of blockchain technology: architecture, consensus, and future trends," in *Proc. IEEE Int. Congr. Big Data (BigData Congress)*, Honolulu, HI, USA, 2017, pp. 557–564.
- [17] K. Wüst and A. Gervais, "Do you need a blockchain?," in *Proc. Crypto Valley Conf. Blockchain Technol. (CVCBT)*, Zug, Switzerland, 2018, pp. 45–54.
- [18] E. Androulaki et al., "Hyperledger Fabric: a distributed operating system for permissioned blockchains," in *Proc. 13th Eur. Conf. Comput. Syst. (EuroSys)*, Porto, Portugal, 2018, pp. 1–15, doi: 10.1145/3190508.3190538.
- [19] C. Cachin, "Architecture of the Hyperledger blockchain Fabric," in *Proc. Workshop Distrib. Cryptocurr. Consensus Ledgers (DCCL)*, Chicago, IL, USA, 2016, p. 4.
- [20] M. Castro and B. Liskov, "Practical Byzantine fault tolerance," in *Proc. 3rd Symp. Oper. Syst. Des. Implement. (OSDI)*, New Orleans, LA, USA, 1999, pp. 173–186.

- [21] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm," in Proc. USENIX Annu. Tech. Conf. (ATC), Philadelphia, PA, USA, 2014, pp. 305–319.
- [22] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "BLOCKBENCH: a framework for analyzing private blockchains," in Proc. ACM Int. Conf. Manag. Data (SIGMOD), Chicago, IL, USA, 2017, pp. 1085–1100.
- [23] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, "OmniLedger: a secure, scale-out, decentralized ledger via sharding," in Proc. IEEE Symp. Secur. Privacy (SP), San Francisco, CA, USA, 2018, pp. 583–598.
- [24] M. Zamani, M. Movahedi, and M. Raykova, "RapidChain: scaling blockchain via full sharding," in Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS), Toronto, ON, Canada, 2018, pp. 931–948.
- [25] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: scaling Byzantine agreements for cryptocurrencies," in Proc. 26th Symp. Oper. Syst. Princ. (SOSP), Shanghai, China, 2017, pp. 51–68.
- [26] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: a provably secure proof-of-stake blockchain protocol," in Proc. Int. Cryptol. Conf. (CRYPTO), Santa Barbara, CA, USA, 2017, pp. 357–388.
- [27] J. Benet, "IPFS—content addressed, versioned, P2P file system," 2014, arXiv:1407.3561.
- [28] N. Nizamuddin, K. Salah, M. A. Azad, J. Arshad, and M. H. Rehman, "Decentralized document version control using Ethereum blockchain and IPFS," Comput. Electr. Eng., vol. 76, pp. 183–197, 2019.
- [29] J. F. Hair, W. C. Black, B. J. Babin, and R. E. Anderson, Multivariate Data Analysis, 8th ed. Boston, MA, USA: Cengage, 2019.
- [30] J. Cohen, Statistical Power Analysis for the Behavioral Sciences, 2nd ed. Hillsdale, NJ, USA: Erlbaum, 1988.