

Volume-06|Issue-09|September|Year-2026 ISSN: 3049-2343

International Journal of Academic and Industrial Research Innovations(IJAIRI)

Research Paper Title :Machine-Verifiable Legal Transactions and Quantum-Secure Cryptographic Agility for Trusted AI-Native 6G Digital Infrastructure

¹Murali Krishna Pasupuleti

¹Research Director

¹National Education Services,

¹New Delhi,India

¹mkrishnap2050@gmail.com

eSign

Signed by: MURALI KRISHNA
PASUPULETI
Reason: NES | IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date :Sep 30, 2026 (05:00 PM)

Abstract:

AI-native sixth-generation (6G) infrastructure will increasingly execute network, commercial and public-service actions through autonomous agents distributed across terminals, edges and clouds. Yet a conventional digital signature can authenticate a message without proving that the message was produced under an authorized legal rule, an admissible machine decision, a valid model state, or a quantum-resilient cryptographic configuration. This paper develops a model-based architecture for machine-verifiable legal transactions coupled with cryptographic agility in AI-native 6G environments. The proposed MVLT-QSA6G framework represents legal obligations as signed policy graphs; binds AI decisions to verifiable inference evidence and network context; selects post-quantum or hybrid cryptographic suites through an agility orchestrator; and commits execution evidence to a tamper-evident audit graph. The methodology synthesizes the supplied literature on verifiable AI, post-quantum cryptography, AI-native 6G, terminal-edge-cloud orchestration, secure edge systems and autonomous contract law, then formalizes system trust through weighted state vectors, transition dynamics, constrained optimization and risk-adjusted outcome functions. An illustrative 0–10 numerical model produces a composite trust score of 8.24/10 under the stated assumptions, with sensitivity analysis showing that cryptographic agility becomes increasingly consequential when legal verifiability and AI-proof integrity are already high. The paper contributes a unified technical and conceptual bridge between legal computability, AI assurance, cryptographic migration and 6G control. It further provides an invention-oriented disclosure comprising a technical problem, novel architecture, implementation logic, one independent system claim concept, one independent method claim concept and dependent claim concepts. The contribution is intended as a research and engineering specification requiring empirical, standards-based and jurisdiction-specific validation before operational deployment or patent prosecution.

Keywords: machine-verifiable legal transactions; AI-native 6G; post-quantum cryptography; cryptographic agility; verifiable AI; proof of inference; legal policy graphs; terminal-edge-cloud computing; zero-trust assurance; autonomous agents; evidence graphs; network slicing; compliance automation; quantum-safe infrastructure; trusted digital transactions.

Copyright:

© 2026 by Murali Krishna Pasupuleti.

1. Introduction

The transition from 5G to AI-native 6G is not merely a bandwidth upgrade. It changes the locus of decision-making from centrally configured networks toward continuously learning and autonomously orchestrated systems in which AI models participate in resource management, service composition, threat response and user-facing transactions. Research on native-AI enablers and terminal-edge-cloud collaboration anticipates distributed inference and control across heterogeneous network domains, while standards-oriented work emphasizes the need for automation interfaces that remain governable under highly dynamic conditions (Karachalios et al. 2023; Majumdar et al. 2025; Wang et al. 2026). The resulting infrastructure can execute actions at machine timescales, but legal accountability, cryptographic assurance and evidentiary reproducibility generally remain organized around slower, document-centric assumptions.

A second transition is cryptographic. Long-lived infrastructure must cope with algorithmic obsolescence, migration from classical public-key mechanisms, heterogeneous device capabilities and the possibility that records captured today may be attacked later. Supplied studies propose post-quantum security frameworks for 6G and hybrid post-quantum protection for connected devices, while evaluations across IoT-edge-cloud layers underline that algorithm choice is constrained by computation, communication and implementation context (Balani & Ali 2026; Udugahapattuwa et al. 2026; Alrowaili et al. 2026). These observations motivate cryptographic agility: the ability to identify, negotiate, rotate and retire algorithms without redesigning the surrounding application or evidentiary model.

A third transition concerns machine verifiability. Lightweight cryptographic proofs of AI inference show that a relying party can, in principle, verify aspects of an inference event rather than trusting an opaque assertion about model execution (Anchuri et al. 2026). In parallel, legal scholarship on autonomous systems questions how conventional contract concepts should treat actions generated by systems with operational autonomy (Zhai 2026). The research problem arises at the intersection: an AI-native 6G action may be technically authenticated yet legally unauthorized; legally authorized yet generated by an unverified model; correctly inferred yet signed using an obsolete cryptographic suite; or securely signed yet impossible to reconstruct for an auditor.

This paper therefore asks how the technical stack can make a transaction verifiable as a chain of machine-checkable propositions: who or what had authority, which rule was applicable, which AI/model state generated the decision, which network context constrained execution, which cryptographic suite protected the event, and which evidence can later demonstrate those facts. The proposed response is MVLT-QSA6G, a layered framework that compiles legal intent into machine-readable policy objects, binds those objects to AI inference proofs and 6G execution context, and uses a post-quantum agility controller to maintain cryptographic continuity. The study is intentionally model-based because no real deployment dataset was supplied. Its results are analytical and illustrative rather than empirical.

eSign

Signed by: MURALI KRISHNA
PASUPULETI
Reason: NES | IJAIIRI | Publication |
Copyright: © 2026
Location: Delhi, India
Date : Sep 30, 2026 (05:00 PM)

1.1 Research Questions

RQ1. How can legal authority, obligations, conditions and remedies be represented as machine-verifiable objects that remain traceable to human-readable transaction terms?

RQ2. How can verifiable AI evidence be bound to legal authorization and AI-native 6G execution context without creating an impractically heavy proof burden?

RQ3. How should cryptographic agility be modeled so that post-quantum migration, algorithm negotiation and key rotation preserve evidentiary continuity?

RQ4. Which terminal-edge-cloud control points are most suitable for policy enforcement, proof verification, risk scoring and evidence commitment?

RQ5. How do legal verifiability, AI-proof integrity, 6G resilience, audit traceability and privacy jointly influence an integrated trust score?

RQ6. Which technical elements plausibly distinguish the proposed architecture from the component capabilities described in the supplied literature?

1.2 Research Objectives

O1. Develop a unified conceptual architecture linking legal policy representation, AI verification, 6G orchestration and post-quantum cryptographic agility.

O2. Formalize the architecture using vector, matrix, dynamic, optimization and risk-adjusted equations.

O3. Construct an illustrative normalized numerical model and sensitivity analysis without presenting simulated scores as empirical evidence.

O4. Define implementation components and data flows suitable for terminal-edge-cloud deployment.

O5. Identify research gaps and validation requirements spanning technology, law, standards, governance and privacy.

O6. Translate the research contribution into an invention-oriented technical disclosure and structured claim concepts.

2. Literature Review

The supplied 6G literature converges on the view that AI will be embedded in network operation rather than attached as an external analytics layer. Karachalios et al. (2023) examine distributed machine learning and native-AI enablers for end-to-end resource management. Majumdar et al. (2025) focus on standards enablers for AI-native 6G network automation, and Wang et al. (2026) describe terminal-edge-cloud collaboration. Together, these works support an architectural assumption central to this paper: evidence and policy controls cannot reside solely in a central cloud because the action to be governed may originate at the terminal or edge and may need low-latency verification before a network or commercial consequence occurs.

Security research in the supplied corpus emphasizes both post-quantum migration and heterogeneous infrastructure. Udugahapattuwa et al. (2026) propose a post-quantum cryptographic framework for high-level security protocols in 6G. Balani and Ali (2026) combine AI-based detection with hybrid post-quantum encryption for IoT devices, while Alrowaili et al. (2026) evaluate PQC algorithms across IoT-edge-cloud layers. The key implication is not that one post-quantum algorithm should dominate every layer, but that system design must preserve policy-controlled substitutability. Pasupuleti (2026a) similarly treats post-quantum migration and verification as an infrastructure problem, while Pasupuleti (2026b) models cross-sector cascade resilience during transition. This paper extends that logic into transaction evidence: algorithm changes must not sever the ability to verify past actions.

The verifiable-AI problem is distinct from, but complementary to, cryptographic authentication. Anchuri et al. (2026) study lightweight cryptographic proofs of inference, supplying a basis for attaching proof artifacts to high-impact AI decisions.

Signed by: MURALI KRISHNA
PASUPULETI
Reason: NES IJAIIRI | Publication
Copyright: 2026
Location: Delhi, India
Date: Sep 30, 2026 (05:00 PM)

Such a proof can establish that an output corresponds to a defined computation or model-related statement, but it does not itself establish legal authority, contractual conditions or the acceptability of the network state in which the decision was executed. Consequently, the proposed architecture treats inference proof as one evidence class in a broader transaction proof bundle rather than as a complete trust mechanism.

AI-native 6G also introduces context dependence. Daniil et al. (2025) discuss AI and machine learning for 6G-enabled IoT, while Zhang, Liu and Zhou (2026) use physics-informed learning for radio-map construction. Beghriche et al. (2026) discuss photonic AI acceleration for wireless signal processing, and Gençtoý (2026) evaluates anomaly detection models for 5G/6G. These sources reinforce the need to record the operational context of automated decisions: topology, slice identity, model version, telemetry state and resource conditions may materially affect whether the same action remains acceptable later.

Edge and smart-city security studies offer further support for distributed enforcement. Xu, Nagothu and Chen (2024) propose an autonomous and resilient edge architecture, Asim et al. (2023) survey secure AI-enabled edge systems, and Li, Xie and Iqbal (2026) propose a quantum-enabled resilient V2X framework. Such work illustrates how security functions span devices, edge nodes and communication links. The current paper adds a legal-verification plane intended to coexist with that distributed security plane rather than replacing it.

The legal dimension is less developed in the technical literature. Zhai (2026) directly asks how contract law should treat autonomous systems, foregrounding questions of attribution and autonomy. Pasupuleti (2026c) addresses transparency and accountability in AI for law and public administration. The gap is therefore not simply a missing smart-contract mechanism; it is the lack of a common evidentiary object that can connect legal authorization, autonomous AI execution, network state and cryptographic validity. A technically valid signature is not synonymous with a legally valid act, and a machine-readable contract is not synonymous with cryptographically agile, independently verifiable execution.

Additional supplied work strengthens the systems context. Sun et al. (2020) and Tagliaferri et al. (2021) show that 6G channel state and mobility can be highly context-dependent, while Sapkale (2026) frames quantum computing as a relevant 6G

research direction. Pasupuleti (2023) examines multi-agent reinforcement-learning slicing with telemetry; Pasupuleti (2021) addresses quantum-safe network security; Pasupuleti (2026d) treats post-quantum cyber risk as an adoption and policy problem; Pasupuleti (2026e) focuses on adaptive AI-enabled cyber defence; and Pasupuleti (2026f) develops number-theoretic perspectives on long-term post-quantum security. These sources motivate retaining signed network context, migration state and security-policy provenance as first-class transaction evidence rather than as external operational metadata.

2.1 Conceptual Framework

The conceptual framework separates six constructs: legal verifiability (L), cryptographic agility (C), AI-proof integrity (A), 6G operational resilience (N), audit traceability (E), and privacy/minimization (P). Legal verifiability measures whether authority, conditions and remedies can be evaluated from signed policy objects. Cryptographic agility measures the ability to move among approved suites while preserving identity and evidence chains. AI-proof integrity captures confidence that the asserted model/inference state corresponds to the recorded decision. Network resilience represents execution continuity and context integrity across terminal-edge-cloud nodes. Audit traceability reflects completeness and tamper evidence of the event graph. Privacy/minimization penalizes architectures that over-collect or expose information in order to achieve verification.

The framework hypothesizes complementarity rather than simple additivity. For example, high AI-proof integrity has limited value if the acting agent lacks legal authority; high legal verifiability is insufficient if evidence is protected by a deprecated cryptographic primitive; and strong post-quantum security does not establish whether a transaction complied with a legal condition. The architecture therefore imposes threshold constraints on critical components and models residual interaction risk.

Signed by: MURALI KRISHNA
PASUPULETI
Reason: NES IJAIIRI | Publication|
Date: Sep 30, 2026
Location: Delhi, India
Date: Sep 30, 2026 (05:00 PM)


MVLT-QSA6G: Layered Trust and Cryptographic-Agility Architecture

Figure 1. MVLT-QSA6G layered architecture linking legal policy, verifiable AI, post-quantum agility, AI-native 6G execution and evidence assurance.

2.2 Research Gap

Table 1. Research gaps derived from the supplied source corpus

Gap type	Gap addressed by this study
Theoretical	Existing supplied works largely treat AI-native networking, post-quantum security, verifiable inference and autonomous legal action as separate problem domains. A unified trust model connecting them is not evident in the supplied corpus.
Methodological	There is limited treatment of transaction-level evidence continuity when cryptographic suites change over time. Migration is usually framed as protocol or infrastructure security rather than legal-evidentiary continuity.

Empirical	The supplied sources do not provide a shared dataset containing legal-policy states, inference proofs, PQC transitions and 6G execution telemetry. This prevents direct empirical estimation of joint effects.
Sectoral	Smart-city, IoT, V2X and critical-infrastructure studies demonstrate distributed security, but do not directly specify a reusable legal transaction verification plane.
Policy / standards	The corpus highlights standards enablers and law/accountability concerns, yet a concrete interface between legal policy objects and network control-plane decisions remains under-specified.

eSign

Signed by: MURALI KRISHNA
PASUPULETI
Reason: NES IJAIRI Publication
Copyright: © 2024
Location: Delhi, India
Date: Sep-30-2026 (05:00 PM)

3. Methodology

This paper proposes a model-based analytical framework and uses numerical scoring for demonstration; empirical validation requires real-world data collection.

The research design is constructive and systems-oriented. First, the supplied corpus is filtered for sources directly relevant to AI-native 6G, verifiable inference, post-quantum cryptography, edge resilience, autonomous-system law and trustworthy AI. Second, design requirements are extracted and mapped to system components. Third, a mathematical trust model is formulated to combine component states while retaining threshold constraints and interaction penalties. Fourth, an illustrative scenario is scored on a normalized 0–10 scale. Fifth, an invention-oriented disclosure is prepared from the system architecture and execution method.

No sampling of human subjects, firms or operational networks is claimed. The unit of analysis is a hypothetical high-impact digital transaction executed by an AI agent over AI-native 6G infrastructure—for example, automated service procurement, machine-to-machine settlement, critical-infrastructure authorization or digitally enforceable

network access. Data sources for a future empirical study would include policy-graph validation logs, model attestations, inference-proof verification times, cryptographic negotiation logs, key-lifecycle events, slice telemetry, edge/cloud execution traces, revocation events, dispute outcomes and auditor verification records.

Validation is designed at four levels. Component validation tests policy parsing, proof verification, key rotation and ledger integrity. Integration validation verifies end-to-end binding among policy, inference and network context. Adversarial validation introduces stale keys, model substitution, replay, policy conflicts, edge compromise and partial evidence loss. Governance validation tests whether an independent auditor can reconstruct why an action occurred and whether the system can distinguish technical validity from legal authorization.

Signed by: MURAVINDHANA
PASUPULETI
Reason: NES | IJAIIRI | Publication |
Copyright: © 2026
Location: Delhi, India
Date: Sep 30, 2026 (05:00 PM)

Figure 2 Logic: From Legal Intent to Machine-Verifiable 6G Execution




Figure 2. End-to-end transaction flow from legal intent to verified 6G execution and feedback control.

3.1 Mathematical Framework

Let the system trust state at time t be represented by the normalized vector $x_t = [L_t, C_t, A_t, N_t, E_t, P_t]^T$, where each element is scaled to $[0,1]$. A transaction is admissible only when both a composite score and mandatory threshold constraints are satisfied. This prevents a very high score in one dimension from compensating for a critical failure in another.

$$T_t = w_L L_t + w_C C_t + w_A A_t + w_N N_t + w_E E_t + w_P P_t - \rho I_t \quad (1)$$

In Equation (1), T_t is composite trust, w_i are non-negative weights summing to one, I_t is an interaction-risk penalty and ρ controls the strength of the penalty. The legal and cryptographic components are intentionally given high weights in the illustrative model because the research title concerns both transaction verifiability and cryptographic agility.

A matrix representation captures cross-effects among dimensions. Let W be a 6×6 influence matrix and an b_t exogenous control vector reflecting policy updates, key rotations or network events.

$$x(t+1) = \sigma(Wx_t + b_t - r_t) \quad (2)$$

eSign

Signed by: MURALI KRISHNA
PASUPULETI
Reason: NES | IJAIIRI | Publication |
Copyright: © 2026
Location: Delhi, India
Date: Sep 30, 2026 (05:00 PM)

Here $\sigma(\cdot)$ is a bounded activation or clipping function and r_t is a vector of detected risks such as proof failure, deprecated algorithms, policy conflict or network compromise. Positive off-diagonal elements may represent complementarity—for example, stronger audit traceability improving effective legal verifiability—whereas negative elements represent cascading risk.

For continuous-time resilience analysis, a dynamic trust state can be represented as:

$$dx(t)/dt = K(u(t) - x(t)) - Da(t) - \Gamma c(t) \quad (3)$$

K is a diagonal recovery/control matrix; $u(t)$ is the target assurance profile; D maps adversarial intensity $a(t)$ into assurance degradation; Γ maps cryptographic transition stress $c(t)$, such as emergency algorithm retirement, into temporary loss of assurance. Equation (3) makes the migration problem explicit: a poorly designed cryptographic transition can reduce trust even if the destination algorithm is stronger.

The agility controller selects cryptographic and evidence-processing decisions under security, latency and resource constraints:

$$\min_z \alpha \cdot R_{sec}(z) + \beta \cdot Latency(z) + \gamma \cdot Cost(z) + \delta \cdot EvidenceLoss(z) \quad (4)$$

subject to $Security(z) \geq S_{min}$; $VerifyTime(z) \leq \tau_{max}$; $Compatibility(z)=1$; and $EvidenceContinuity(z)=1$. The decision variable z includes cryptographic suite, key size, proof mode, verification location and evidence-retention policy. The term $EvidenceLoss(z)$ penalizes migrations that make earlier signatures, proofs or identities difficult to validate.

A risk-adjusted probability of trustworthy completion can be modeled using a logistic link:

$$Pr(Y=1|x)=1/1+\exp[-(\beta_0+\beta^T x-\theta_q Q-\theta_d D_s-\theta_p P_c)] \quad (5)$$

$Y=1$ denotes a transaction that is both technically verified and policy-admissible. Q is quantified quantum-migration exposure, D_s is distribution/supply-chain exposure, and P_c is unresolved policy conflict. Coefficients would require empirical estimation; the equation is provided as a future validation model, not as a fitted result.

Mandatory admissibility can be encoded with an indicator:

$$A_t^* = 1[T_t \geq \tau_T] \cdot 1[L_t \geq \tau_L] \cdot 1[C_t \geq \tau_C] \cdot 1[A_t \geq \tau_A] \quad (6)$$

$A_t^* = 1$ indicates execution may proceed. If any mandatory threshold fails, the action is denied, deferred for human authorization, or routed to a fallback policy. This structure is deliberately conservative for legally consequential autonomous transactions.

Table 2. Mathematical symbols and interpretation

Symbol	Meaning	Scale / form
T_t	Composite trust score at time t	0–1
L_t	Legal verifiability / authorization integrity	0–1
C_t	Cryptographic agility and quantum-safe readiness	0–1
A_t	AI inference-proof and model-attestation integrity	0–1
N_t	AI-native 6G network resilience/context integrity	0–1
E_t	Evidence completeness and audit traceability	0–1
P_t	Privacy-preserving verification / data minimization	0–1
I_t	Interaction or coupling risk	≥ 0

	among weak components	
W	Cross-effect matrix among assurance dimensions	6×6
r_t	Detected risk vector	6×1
z	Cryptographic/proof configuration decision	Discrete/continuous
Y	Trustworthy transaction completion indicator	$\{0,1\}$

3.2 Numerical Model for Better Insight

Signed by: MURALI KRISHNA
PASUPULETI
Reason: NES | IJAIIRI | Publication |
Copyright: © 2026
Location: Delhi, India
Date : Sep 30, 2026 (05:00 PM)

For illustration, six normalized dimensions are scored from 0 to 10: $L=8.2$, $C=8.7$, $A=7.9$, $N=8.1$, $E=8.8$ and $P=7.6$. These numbers are illustrative and are not observed measurements. Weights are $w=[0.22, 0.22, 0.18, 0.16, 0.13, 0.09]$. A modest interaction-risk deduction of 0.06 is applied after weighted aggregation.

Table 3. Illustrative composite-trust calculation

Dimension	Score	Weight	Weighted contribution
Legal verifiability	8.2	0.22	1.804
PQC agility	8.7	0.22	1.914
AI-proof integrity	7.9	0.18	1.422
6G resilience	8.1	0.16	1.296
Audit traceability	8.8	0.13	1.144
Privacy/minimization	7.6	0.09	0.684

Step 1: Weighted subtotal = $1.804 + 1.914 + 1.422 + 1.296 + 1.144 + 0.684 = 8.264$.
Step 2: apply the illustrative interaction-risk deduction of 0.024, yielding $T = 8.240/10$. Step 3: apply mandatory thresholds. If $\tau_L=7.0$, $\tau_C=7.5$ and $\tau_A=7.0$, the illustrative scenario passes all three thresholds. The result should be interpreted as a coherent design example, not as evidence that a real implementation would achieve 8.24/10.




Figure 3. Illustrative comparison of a conventional stack and the proposed framework across six assurance dimensions. Scores are hypothetical.




Figure 4. Sensitivity of the composite trust score to cryptographic-agility performance under fixed illustrative values for other dimensions.

3.3 Data and Variables

Table 4. Proposed variables for empirical validation

Variable	Definition	Measurement	Expected effect	Relevance
Legal verifiability (L)	Share of executable conditions with traceable authority, rule provenance and deterministic validation	0–10 or %	Positive	Core machine-verifiable transaction construct
Cryptographic agility (C)	Ability to negotiate/rotate approved suites while retaining identity/evidence continuity	0–10	Positive	Supports post-quantum transition
AI-proof integrity (A)	Verified inference/model-attestation success and freshness	0–10 or %	Positive	Links autonomous decision to evidence
6G resilience (N)	Availability, context integrity and recovery across terminal-edge-cloud execution	0–10	Positive	Represents AI-native network reliability
Evidence traceability (E)	Completeness, immutability and reconstructability of the evidence	0–10	Positive	Enables audit and dispute reconstruction

	graph			
Privacy/minimization (P)	Extent to which verification avoids unnecessary disclosure	0–10	Positive	Constrains over-collection
Quantum migration exposure (Q)	Dependence on deprecated/soon-to-be deprecated primitives	0–10	Negative	Captures cryptographic transition risk
Policy conflict (Pc)	Unresolved conflict among legal rules, delegations or jurisdictional constraints	Count/severity	Negative	Can invalidate technically sound execution
Verification latency	Time to validate policy, proof and cryptographic evidence	ms	Negative	Critical for 6G use cases
Evidence continuity	Proportion of prior records still independently verifiable after suite/key change	%	Positive	Tests agility without audit breakage

3.4 Proposed Model / Framework

The proposed model is named the Machine-Verifiable Legal Transaction and Quantum-Secure Agility for 6G (MVLТ-QSA6G) framework. Its operational equation is:

Decision

=

$$VerifyLegal(\Pi) \wedge VerifyAI(\pi_A I) \wedge VerifyCrypto(\kappa_t) \wedge VerifyContext(\chi_6 G) \wedge VerifyEvidence(G_t) \quad (7)$$

Π is the signed legal-policy graph; $\pi_A I$ is the proof/attestation bundle for AI execution; κ_t is the cryptographic-suite and key-state object; $\chi_6 G$ is the signed network-context object; and G_t is the evidence graph. The conjunction expresses a design principle: a transaction should not be treated as trusted merely because one subsystem succeeds. Each verifier returns both a Boolean result and structured reasons, enabling explainable rejection and human escalation.




Figure 5. Evidence graph binding parties, policy, AI agent, 6G execution, cryptographic key state and independent audit.

4. Results and Analysis

The analytical result is that trust in autonomous 6G transactions is structurally multi-dimensional. A single-signature model can authenticate a message but cannot, by itself, prove that the signer or agent was authorized under the applicable policy, that an AI output came from the declared model state, or that the relevant 6G execution context remained within authorized bounds. Conversely, a verifiable inference proof can establish aspects of computation but not the legal entitlement to act. The proposed framework therefore treats trustworthy completion as a conjunction of verifiers with a

composite score used for monitoring rather than as a substitute for hard authorization constraints.

The numerical illustration suggests that improvements in cryptographic agility have the largest marginal effect when the legal and AI-proof layers are already credible. This follows directly from the weighting and threshold structure rather than from observed data. In a poorly governed system with low L or A, increasing C alone cannot make the transaction admissible. In a mature governance system, however, weak cryptographic agility becomes a bottleneck because evidence and identity continuity can fail during migration. This is consistent with the supplied post-quantum literature emphasizing heterogeneous constraints and migration requirements (Udugahapattuwa et al. 2026, Alnowaidi et al. 2026; Pasupuleti 2026a).

A second result concerns placement. Policy parsing and global governance rules are naturally managed in trusted control or cloud environments, while low-latency authorization tokens, proof checks and key-state validation may need to execute at the edge. Terminal devices contribute device-bound identity, sensor/transaction context and local consent or delegation. Wang et al. (2026) and Karachalios et al. (2023) support the broader terminal-edge-cloud architecture, while Xu, Nagothu and Chen (2024) and Asim et al. (2023) demonstrate the importance of resilient edge security. The proposed design therefore distributes enforcement but centralizes neither trust nor evidence completely.

A third result is evidentiary continuity. Cryptographic agility should be evaluated not only by how quickly a system adopts a stronger algorithm but also by whether historical evidence remains verifiable. The architecture stores a cryptographic manifest describing suite identifiers, validation policies, certificate/key lineage and transition proofs. A future verifier can determine which verification procedure applies to a historical event without silently treating all records under the current suite. This is a technical control aimed at preventing migration from becoming an audit discontinuity.

A fourth result is that legal policy needs a compiled representation rather than free-text ingestion alone. The proposed transaction compiler converts human-readable terms into typed policy nodes—for example, actor authority, prerequisite, jurisdiction, asset, time window, threshold, prohibited state, evidence requirement and remedy/escalation. Each machine-executable node retains a provenance pointer to the

governing clause. This does not claim that legal interpretation is fully automatable. Instead, it makes explicit which propositions the system is authorized to test and which ambiguous conditions require human determination.

Table 5. Analytical scenarios for MVLT-QSA6G

Scenario	Condition	System response	Residual risk
S1: Normal execution	Current policy; verified model; approved PQC suite; healthy edge	Proceed and commit evidence graph	Low
S2: Deprecated algorithm	Policy valid; proof block valid; suite retired by security policy	Negotiate new suite; block irreversible action until re-signed	Medium
S3: Model substitution	Policy valid; attestation hash mismatch	Deny; quarantine agent; preserve evidence	High
S4: Policy conflict	Two applicable rules return inconsistent authority	Escalate to human/legal resolver; no autonomous completion	High
S5: Edge compromise	Context attestation fails at local edge	Reroute verification/execution or suspend transaction	High
S6: Historical audit	Old suite no longer active but manifest and transition proof preserved	Verify under historical policy and current trust anchors	Controlled

4.1 Discussion

The theoretical contribution is a layered definition of digital transaction trust that distinguishes authentication from authorization, computation from legal entitlement,

and cryptographic strength from cryptographic continuity. The architecture treats these as interdependent but non-substitutable properties. This matters for AI-native 6G because autonomous decisions may be executed before conventional human review is possible.

Practically, the framework can be implemented incrementally. Organizations do not need a complete machine-readable legal system on day one. They can begin with narrow transaction classes whose authorization conditions are already structured: spending limits, network-slice access, service-level thresholds, device enrollment, automated procurement parameters or machine-to-machine settlement rules. These can be compiled into policy objects and linked to existing identity, PKI/PQC migration and telemetry systems.

For industry, the most important design implication is separation of concerns. The legal-policy compiler should not choose cryptographic algorithms; the agility orchestrator should not interpret ambiguous law; and the AI verifier should not decide whether an agent possesses legal authority. Each component should expose signed evidence and machine-readable reasons to an orchestration layer. This separation reduces the risk that one opaque AI component becomes the de facto judge of technical, legal and security questions simultaneously.

For policy and governance, the framework suggests that machine-readable rules require change control, versioning and appeal paths. Pasupuleti (2026c) emphasizes transparency and accountability in AI for law and public administration, while Zhai (2026) highlights the legal difficulty created by autonomous-system action. The proposed design therefore keeps a human-readable clause provenance pointer, policy version identifier and human-escalation state in the evidence graph.

Ethically, verifiability can conflict with privacy. Recording every model input, network feature and legal attribute would improve reconstructability but could create disproportionate surveillance. The privacy/minimization dimension therefore rewards selective disclosure, hashed commitments, role-based evidence access and proof techniques that establish compliance without exposing all underlying data. The framework does not prescribe one privacy technology; it makes privacy an explicit trust dimension rather than an afterthought.

A further concern is automation bias. A machine-verifiable result is only as legitimate as the rules and evidence classes embedded in the verifier. A false sense of certainty may arise if normative ambiguity is converted into a Boolean output. The design therefore distinguishes hard machine-checkable conditions from interpretive conditions, with the latter routed to human or institutional resolution. Machine verifiability is presented as an accountability enhancement, not as a replacement for legal judgment.

5. Conclusion

This paper developed MVL-T-QSA6G, a model-based framework for trusted AI-native 6G digital infrastructure, which legal transactions, AI decisions, cryptographic state and network execution are verified as an integrated evidence chain. The central proposition is that future autonomous infrastructure requires more than secure communication: it requires machine-checkable authority, proof-linked AI action, crypto-agile identity and evidence continuity, distributed enforcement and reconstructable audit. The mathematical formulation combines weighted trust scoring, mandatory thresholds, dynamic resilience and constrained agility optimization. An illustrative scenario produced a composite score of 8.24/10 under stated assumptions, but no empirical claim is made. The framework's value lies in specifying what must be measured and how failures should be isolated. It also establishes a concrete technical architecture suitable for prototyping in terminal-edge-cloud testbeds and for future empirical, legal and standards validation.

5.1 Limitations

The study is limited by the absence of real transaction, network and cryptographic-migration datasets. The numerical scores are illustrative and depend on chosen weights. Legal validity varies by jurisdiction and transaction type; the paper does not establish enforceability of any machine-executable policy. Cryptographic security depends on implementation quality, standards evolution and algorithm parameters that are outside the supplied corpus. The architecture does not resolve semantic ambiguity in natural-language law and assumes that at least a subset of conditions can be represented as deterministic or formally reviewable propositions. Finally, the prior-art discussion below is literature-based and cannot substitute for a professional patent database search, claim chart or jurisdiction-specific freedom-to-operate analysis.

5.2 Future Research Directions

Develop a benchmark dataset linking policy graphs, AI inference proofs, PQC key transitions and 6G telemetry for reproducible evaluation.

Prototype policy-compilation languages that preserve clause provenance and formally distinguish deterministic from interpretive legal conditions.

Benchmark verifiable-inference overhead under 6G latency budgets and heterogeneous terminal/edge hardware.

Design evidence-continuity protocols that allow historical verification across multiple generations of classical, hybrid and post-quantum suites.

Investigate privacy-preserving proof bundles using selective disclosure and zero-knowledge techniques without sacrificing auditability.

Create cross-jurisdiction policy conflict detectors and escalation protocols for autonomous transactions.

Conduct adversarial red-team studies covering replay, stale delegation, model substitution, proof downgrade and malicious crypto-negotiation attacks.

Perform formal patent landscaping and standards-essentiality analysis against issued patents, published applications and 6G/PQC standards documents.

References

- Alrowaili, A, Ozmen, SC, Kazim, M & Rehman, MU 2026, 'Post-Quantum Cryptography for Smart City Cybersecurity: An Evaluation of PQC Algorithms Across IoT–Edge–Cloud Layers', IET Smart Cities, vol. 8, no. 1, accessed <<http://dx.doi.org/10.1049/smc2.70040>>.
- Anchuri, P, Campanelli, M, Cesaretti, P, Gennaro, R, Jois, TM, Kayman, HS & Ozdemir, T 2026, 'Towards Verifiable AI with Lightweight Cryptographic Proofs of Inference', 2026 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML), pp. 955–977, accessed <<http://dx.doi.org/10.1109/satml68713.2026.00058>>.
- Asim, M, Junhong, C, Wenyin, L & Abd El-Latif, AA 2023, 'Artificial Intelligence-Based Secure Edge Computing Systems for IoTs and Smart Cities: A Survey', Secure Edge and Fog Computing Enabled AI for IoT and Smart Cities, pp. 155–173, accessed <http://dx.doi.org/10.1007/978-3-031-51097-7_13>.
- Balani, Z & Ali, MSM 2026, 'Securing IoT Devices in the Quantum Era: AI-Based Detection with Hybrid Post-Quantum Encryption', 2026 IEEE 2nd International Conference on Secure IoT, Assured and Trusted Computing (SATC), pp. 1–5, accessed <<http://dx.doi.org/10.1109/satc69565.2026.11542476>>.
- Beghriche, A, Bouaita, B, Benmessahel, B, Berrhail, F & Seghir, F 2026, 'Photonic AI Accelerators for Ultra-Fast Wireless Signal Processing in 6G Networks', Traitement du Signal, vol. 43, no. 1, pp. 27–51, accessed <<http://dx.doi.org/10.18280/ts.430103>>.
- Daniil, B, Bektemyssova, G, Razaque, A, Beketova, A & Farhad, MA 2025, 'AI and Machine learning for 6G-enabled IoT', 2025 1st International Conference on Secure IoT, Assured and Trusted Computing (SATC), pp. 1–6, accessed <<http://dx.doi.org/10.1109/satc65530.2025.11137150>>.
- Gençtoý, S 2026, 'A PyCaret-Assisted Comparative Study of AI/ML Models for Anomaly Detection in 5G and 6G Networks', 2026 34th Signal Processing and Communications Applications Conference (SIU), pp. 1–4, accessed <<http://dx.doi.org/10.1109/siu71813.2026.11636811>>.

- Karachalios, OA, Zafeiropoulos, A, Kontovasilis, K & Papavassiliou, S 2023, 'Distributed Machine Learning and Native AI Enablers for End-to-End Resources Management in 6G', Electronics, vol. 12, no. 18, p. 3761, accessed <<http://dx.doi.org/10.3390/electronics12183761>>.
- Li, S, Xie, Y & Iqbal, M 2026, 'QuanReS: A Hybrid Quantum-Enabled Resilient and Secure V2X Communications Framework in Smart Cities', IEEE Transactions on Consumer Electronics, vol. 72, no. 2, pp. 3940–3948, accessed <<http://dx.doi.org/10.1109/tce.2025.3647004>>.
- Majumdar, S, Wei, Q, Schwarzmnn, S, Trivisonno, R & Carle, G 2025, Towards AI-Native 6G Systems: Standards Enablers for 6G Network Automation, accessed <<http://dx.doi.org/10.36227/techrxiv.2025.08360.03565555/v2>>.
- Pasupuleti, MK 2026a, 'Quantum-Safe Critical Infrastructure: Post-Quantum Migration and Verification Frameworks', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 06, no. 02, National Education Services, pp. 22–48, viewed <<http://dx.doi.org/10.62311/nex/rp-f22026>>.
- Pasupuleti, MK 2026b, 'Cross-Sector Cascade Resilience in the Post-Quantum Transition: Coupled-Network Theorems and Stress-Testing Algorithms for Finance–Telecom–Industrial Control Systems (ICS)', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 06, no. 01, National Education Services, pp. 1–23, viewed <<http://dx.doi.org/10.62311/nex/rp-j12026>>.
- Pasupuleti, MK 2026c, 'Toward Trustworthy Artificial Intelligence in Law and Public Administration: Transparency, Accountability and Decision Support', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 06, no. 03, National Education Services, pp. 1–28, viewed <<http://dx.doi.org/10.62311/nex/rp-1-28-30032026>>.
- Pasupuleti, MK 2026d, 'Mechanism Design for Systemic Post-Quantum Cyber Risk: Incentive-Compatible Adoption, Equilibrium Guarantees, and Policy-Grade Algorithms', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 06, no. 01, National Education Services, pp. 24–58, viewed <<http://dx.doi.org/10.62311/nex/rp-j22026>>.

- Pasupuleti, MK 2026e, 'Quantum-AI Models for Real-Time Cybersecurity Threat Detection and Adaptive Defense', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 06, no. 05, pp. 1–35, accessed <<http://dx.doi.org/10.62311/nexs/rp1-30052026>>.
- Pasupuleti, MK 2026f, 'Quantum-AI Number-Theoretic Frameworks for Post-Quantum Cryptography and Long-Term Information Security', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 06, no. 05, pp. 36–55, accessed <<http://dx.doi.org/10.62311/nexs/rp2-30052026>>.
- Pasupuleti, MK 2023, 'NS-3 for 6G: Multi-Agent RL Slicing with eBPF Telemetry', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, National Education Services, pp. 162–177, viewed <<http://dx.doi.org/10.62311/nexs/rp-30623-162-177>>.
- Pasupuleti, MK 2021, 'Quantum-Safe Network Security: Advances in Quantum Cryptography and Systems Integration', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, National Education Services, pp. 232–245, viewed <<http://dx.doi.org/10.62311/nexs/rp-1-01-2021>>.
- Sapkale, P 2026, 'A Quantum Computing Perspective in 6G Networks', Quantum Computing and Machine Learning for 6G, pp. 311–325, accessed <<http://dx.doi.org/10.1002/9781394238118.ch14>>.
- Sun, Y, Wang, C-X, Huang, J & Wang, J 2020, 'A 3D Non-Stationary Channel Model for 6G Wireless Systems Employing Intelligent Reflecting Surface', 2020 International Conference on Wireless Communications and Signal Processing (WCSP), pp. 19–25, accessed <<http://dx.doi.org/10.1109/wcsp49889.2020.9299848>>.
- Tagliaferri, D, Mizmizi, M, Badini, D, Mazzucco, C & Spagnolini, U 2021, 'Multi-Vehicular Beam Space Learning for Channel Estimation in 6G V2X Systems', 2021 IEEE 22nd International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), pp. 46–50, accessed <<http://dx.doi.org/10.1109/spawc51858.2021.9593123>>.
- Udugahapattuwa, P, De Alwis, C, Zeydan, E, Wijewardhana, U & Liyanage, M 2026, 'A Post-Quantum Cryptographic Framework for High-Level Security Protocols in 6G', 2026 Joint European Conference on Networks and Communications & 6G

Summit (EuCNC/6G Summit), pp. 1077–1084, accessed
<<http://dx.doi.org/10.1109/eucnc/6gsummit68295.2026.11577443>>.

Wang, Q, Dan, D, Li, Z, Wu, C & Wang, J 2026, ‘Terminal-Edge-Cloud Collaborative AI-Native Network Towards 6G’, 2026 Global 6G Conference (Global 6G), pp. 1–6, accessed
<<http://dx.doi.org/10.1109/global6g69473.2026.11566897>>.

Xu, R, Nagothu, D & Chen, Y 2024, ‘AR-Edge: Autonomous and Resilient Edge Computing Architecture for Smart Cities’, Edge Computing Architecture - Architecture and Applications for Smart Cities, accessed
<<http://dx.doi.org/10.5772/intechopen.1005876>>.

Zhai, X 2026, ‘Autonomous systems: how should contract law treat such autonomy?’, Information & Communications Technology Law, pp. 1–29, accessed
<<http://dx.doi.org/10.1080/13600834.2026.2624924>>.

Zhang, Y, Liu, S & Zhou, T 2026, ‘PhysRes-Net: A Physics-Informed Residual Learning Framework for Radio Map Construction in 6G AI-Native Networks’, 2026 Global 6G Conference (Global 6G), pp. 1–6, accessed
<<http://dx.doi.org/10.1109/global6g69473.2026.11566753>>.