

QuantumLock

Post-Quantum Cryptographic Proof for Critical AI and Sovereign
Decisions

Institutional White Paper

SoftQuantus innovative OÜ

Company Registry: 17048927
<https://softquantus.com>

Version 1.0 January 2026

Contents

Executive Summary

The growing integration of artificial intelligence systems in defense, security, and space domains creates a new legal and institutional requirement: the ability to demonstrate, over the long term, the integrity, traceability, and legitimacy of critical algorithmic decisions.

QuantumLock is a post-quantum cryptographic proof system designed to transform AI processing into verifiable proof objects, deployable in sovereign, classified, or air-gapped environments.

Key Points

- **Legal compliance:** Meets EU AI Act (Art. 12) requirements for traceability and record-keeping
- **Post-quantum security:** Protection against "harvest now, decrypt later" threats via NIST/ETSI/ANSSI standards
- **Digital sovereignty:** On-premise, air-gapped deployment without cloud dependencies
- **Robust legal proof:** Independently verifiable artifacts suitable for institutional audits

1 Context and Sovereign Stakes

1.1 AI in Critical State Domains

Artificial intelligence systems are now deployed in contexts where decisions engage national security, territorial defense, and the State's technological sovereignty:

- **Defense and military operations:** Intelligence analysis, tactical decision support, autonomous systems
- **Security and surveillance:** Threat detection, behavioral analysis, cyber defense
- **Space and critical infrastructure:** Satellite control, vital infrastructure management
- **Sovereign research:** Classified projects, strategic capability development

In these contexts, the **legal accountability** and **institutional legitimacy** of automated decisions become as critical as their technical performance.

1.2 Risk Transformation: From Security to Proof

Traditionally, security for critical AI systems focuses on:

- Perimeter protection (firewalls, network isolation)
- Access control (authentication, authorization)
- Operational supervision (SIEM logs, monitoring)

However, these mechanisms do not answer a fundamental question:

How can we prove, years after the fact, that an algorithmic decision was made with integrity, in a verified context, according to an established chain of responsibility?

This requirement for **enduring proof** becomes imperative in three situations:

1. **Judicial review:** Administrative appeals, litigation, parliamentary inquiries
2. **Inter-agency audit:** Compliance verification, certification, responsibility transfer
3. **Post-incident investigation:** Forensic analysis, decision chain reconstruction

2 Legal Framework and Regulatory Compliance

2.1 EU AI Act: The Traceability Obligation

The European Regulation on Artificial Intelligence (AI Act) imposes explicit obligations for **record-keeping** and traceability on high-risk AI systems (Article 12):

"High-risk AI systems shall be designed in such a way to enable the automatic recording of events (logs) over their lifetime. Logging capabilities shall ensure a level of traceability appropriate [...]"

This obligation specifically targets:

- Identification of input data used
- Traceability of automated decisions
- Capability for ex post audit and verification
- Evidential preservation of processing

Problem: Conventional application logs do not constitute robust legal evidence because they are:

- Modifiable after the fact (lack of cryptographic immutability)
- Dependent on execution context (vulnerable to system compromise)
- Not independently verifiable (tied to the original platform)

2.2 Defense Legal Framework and State Accountability

Beyond the AI Act, AI systems used in defense and security contexts are subject to specific requirements:

- **Administrative law:** Obligation to justify decisions, precautionary principle
- **Defense law:** Chain of command, operational responsibility
- **Democratic oversight:** Parliamentary inquiries, verification commissions
- **International law:** Conventions on the use of autonomous systems

These frameworks converge toward the same requirement: the ability to **demonstrate** compliance and integrity of processing, not merely to **assert** it.

2.3 Digital Sovereignty and Technological Control

Dependence on cloud providers or proprietary platforms for integrity proof constitutes an unacceptable sovereign risk in critical state contexts.

Organizations such as **defense AI agencies** require solutions that are:

- Deployable in air-gapped environments
- Verifiable without external infrastructure
- Controllable at algorithmic and cryptographic levels
- Sustainable against technological evolution (notably quantum)

3 Post-Quantum Threat and Cryptographic Horizon

3.1 The "Harvest Now, Decrypt Later" Risk

Large-scale quantum computers, while not yet operational, represent an imminent threat to current cryptography. Security agencies identify the following risk:

HNDL Threat: Adversaries can capture encrypted data today with the intention of decrypting it tomorrow, when quantum computers become available.

For defense AI systems, this threat is particularly critical:

- Algorithmic decisions must remain confidential for decades
- AI models constitute long-term strategic assets
- Retrospective compromise of proofs would invalidate all audit capability

3.2 NIST, ETSI, and ANSSI Post-Quantum Standards

In response to this threat, standardization bodies have finalized the first post-quantum cryptography standards:

NIST FIPS 203/204/205 (August 2024):

- **ML-KEM** (Module-Lattice-Based Key Encapsulation Mechanism): Key establishment
- **ML-DSA** (Module-Lattice-Based Digital Signature Algorithm): Digital signatures
- **SLH-DSA** (Stateless Hash-Based Digital Signature Algorithm): Stateless signatures

ANSSI Recommendations (2024):

- Gradual migration toward post-quantum cryptography
- Hybrid approach during transition phase
- Prioritization of long-lived data

ETSI TS 103 744:

- Migration profiles for critical systems
- Crypto-agility mechanisms
- Secure transition protocols

3.3 Migration Imperative for Critical AI Systems

Defense AI systems present characteristics that make post-quantum migration urgent:

Characteristic	PQC Implication
Long lifespan	High HNDL vulnerability
Strategic value	Priority target for capture
Classified data	Requirement for enduring confidentiality
Deferred audit	Need for quantum-resistant proof

Table 1: Critical AI system characteristics and post-quantum implications

4 Limitations of Current Approaches

4.1 SIEM Logs: Operational Security, Not Legal Proof

Security Information and Event Management (SIEM) systems are the standard for operational supervision. However, they present structural limitations for legal proof:

Criterion	Classic SIEM	QuantumLock Proof
Immutability	Not guaranteed	Cryptographically sealed
Independent verification	System-dependent	Offline verifiable
Long-term integrity	Vulnerable (classic crypto)	Quantum-resistant
Portability	Platform-bound	Autonomous artifact
Legal value	Limited (modifiable)	Evidential (inviolable)

Table 2: SIEM vs. cryptographic proof comparison

4.2 Blockchain: Unnecessary Decentralization, Limited Performance

Blockchain solutions are sometimes proposed to guarantee immutability. However, they present major drawbacks for sovereign environments:

- **Network dependency:** Incompatible with air-gapped environments
- **Unwanted decentralization:** The State must retain complete control
- **Performance:** High latency, significant computational cost
- **Confidentiality:** Model unsuitable for classified data
- **Classic crypto:** Vulnerable to quantum threat

4.3 Application Logging: Insufficient for Legal Accountability

Logging mechanisms integrated into AI frameworks (TensorFlow, PyTorch) do not constitute proof:

- Absence of cryptographic protection for logs
- No algorithmic chain of custody
- Vulnerability to post-facto modifications
- Not designed for independent verification

Finding: There exists a gap between operational security mechanisms and requirements for enduring legal proof.

5 QuantumLock Proof Model

5.1 Fundamental Principles

QuantumLock is based on four architectural principles:

1. **Cryptographic proof:** Each AI execution generates a mathematically provable artifact
2. **Algorithmic chain of custody:** Complete traceability of model, data, and context
3. **Independent verification:** Proofs are validatable without access to the original system
4. **Post-quantum resistance:** Protection against future cryptographic threats

5.2 Anatomy of an Evidence Bundle

A QuantumLock Evidence Bundle contains the following elements:

Evidence Bundle Structure

1. **Model fingerprint:** Cryptographic hash of architecture and weights
2. **Execution context:** Parameters, runtime environment, configuration
3. **Governance metadata:** Version, provenance, deployment authorization
4. **Input data:** Hash of inputs or secure representation
5. **Algorithmic result:** AI output or produced decision
6. **Qualified timestamp:** Cryptographic temporal proof
7. **Signature chain:** Multi-level post-quantum signatures
8. **Verification metadata:** Information for independent validation

5.3 Verifiable Properties

A QuantumLock Evidence Bundle enables cryptographic verification of:

- **Integrity:** The bundle has not been modified since creation
- **Authenticity:** The bundle was generated by the authorized system
- **Non-repudiation:** The issuer cannot deny having produced the bundle
- **Temporality:** The timestamp is reliable and inviolable
- **Compliance:** The model used corresponds to the certified version
- **Traceability:** The chain of responsibility is complete and verifiable

These properties are verifiable **years after the fact**, without dependency on the original system, and resist quantum threats.

5.4 Generation and Verification Workflow

Phase 1: Proof Generation (at execution time)

1. QuantumLock system captures execution context
2. Computation of cryptographic fingerprints (model, data, environment)
3. Qualified timestamping via temporal certification authority
4. Multi-level post-quantum signature
5. Assembly of autonomous Evidence Bundle
6. Secure storage (evidential file system, HSM)

Phase 2: Independent Verification (audit, investigation)

1. Recovery of Evidence Bundle
2. Verification of post-quantum signatures
3. Timestamp validation
4. Fingerprint integrity control
5. Chain of custody reconstruction
6. Production of verification report

Critical characteristic: Verification can be performed offline, in air-gapped environments, by third parties (inquiry commission, external auditor, judicial authority).

6 Post-Quantum Cryptography and Crypto-Agility

6.1 Recommended Cryptographic Schemes

QuantumLock implements NIST standards finalized in August 2024:

For key establishment:

- **ML-KEM-768** (FIPS 203): Security equivalent to AES-192, optimal performance
- Use for encryption of sensitive data in the bundle

For digital signatures:

- **ML-DSA-65** (FIPS 204): Primary signature, security/performance balance
- **SLH-DSA-128s** (FIPS 205): Backup signature, based solely on hash functions

6.2 Hybrid Approach for Transition

In accordance with ETSI TS 103 744 and ANSSI recommendations, QuantumLock supports a **hybrid** approach during the transition phase:

- **Dual signatures:** Combination of classic algorithms (RSA-4096, ECDSA) and post-quantum (ML-DSA)
- **Progressive protection:** Migration without disrupting existing systems
- **Compatibility:** Verification possible with or without PQC capabilities
- **Controlled evolution:** Planned transition to pure PQC

This ensures that:

1. Current proofs remain verifiable with existing infrastructure
2. New proofs are already protected against quantum threat
3. Migration is reversible and controllable

6.3 Crypto-Agility: Anticipating Future Evolutions

QuantumLock is designed according to **crypto-agility** principles:

Crypto-agility: Ability to change cryptographic algorithms without architectural re-design, in response to evolving threats or standards.

Implemented mechanisms:

- **Algorithmic abstraction:** Cryptographic primitives are interchangeable
- **Scheme versioning:** Each bundle indicates algorithms used
- **Multi-algorithm support:** Verification possible with different combinations
- **Planned evolution:** Migration roadmap integrated into the system

This guarantees proof sustainability even if algorithms are later deprecated or replaced.

7 Reference Architecture

7.1 Deployment Modes

QuantumLock supports three deployment modes adapted to institutional constraints:

Mode 1: Sovereign On-Premise

- Complete deployment in organization's datacenter
- Total control of cryptographic infrastructure
- Integration with existing PKI
- Storage on institutional evidential file systems

Mode 2: Air-Gapped (classified environments)

- Operation without external connectivity
- Timestamping via certified local temporal source
- Complete offline verification
- Bundle transfer via secure physical media

Mode 3: Hybrid (defense in depth)

- Local proof generation
- Replicated archiving (local + external digital vault)
- Double timestamping (internal TSA + external qualified TSA)
- Redundancy to guarantee sustainability

7.2 Architectural Components

QuantumLock architecture is built around five main components:

1. **Evidence Generator:** Capture and cryptographic sealing module
2. **Timestamping Authority:** Qualified timestamping service (internal or external)
3. **Signature Engine:** Post-quantum signature module (HSM-backed)
4. **Evidence Store:** Evidential and enduring storage system
5. **Verification Toolkit:** Independent validation tools

7.3 Integration with Existing AI Systems

QuantumLock integrates **non-intrusively** into MLOps pipelines:

For training:

- Capture of training metadata (hyperparameters, datasets, metrics)
- Sealing of final model
- Generation of cryptographic compliance certificate

For inference:

- Capture of execution context
- Recording of inputs/outputs (or their fingerprints)
- Generation of proof bundle
- Minimal performance impact (<5% overhead)

For governance:

- Integration with model management systems (MLflow, etc.)
- Export of bundles to institutional archiving systems
- Verification APIs for audit tools

8 Defense and Security Use Cases

8.1 Inter-Agency Audit

Context: A parliamentary commission or oversight authority wishes to audit the use of a defense AI system years after its deployment.

Requirements:

- Prove that the model used corresponds to the certified version
- Demonstrate integrity of processing performed
- Reconstruct chain of responsibility
- Verify regulatory compliance

QuantumLock contribution:

- Provision of Evidence Bundles for audited period
- Independent cryptographic verification by auditors
- Complete reconstruction of execution history
- Production of verifiable compliance report

8.2 Post-Incident Investigation

Context: An algorithmic decision led to an operational incident. An investigation must establish the causal chain.

Requirements:

- Identify the exact model and version used
- Reconstruct execution context (data, parameters)
- Verify absence of system compromise
- Establish precise chronology of events

QuantumLock contribution:

- Timestamped Evidence Bundle of incriminated execution
- Cryptographic proof of model integrity
- Complete traceability of inputs and context
- Incontestable factual basis for investigation

8.3 Inter-System Responsibility Transfer

Context: An AI system developed by agency A must be transferred and operated by agency B, with compliance guarantee.

Requirements:

- Certify integrity of transferred model
- Prove compliance with specifications

- Establish continuity of chain of trust
- Enable future audit by agency B

QuantumLock contribution:

- Model cryptographic certification bundle
- Independently verifiable compliance proof
- Secure transfer of chain of custody
- Preserved audit capability for receiving agency

8.4 AI Act Compliance for High-Risk Systems

Context: An AI system classified as "high-risk" under the AI Act must demonstrate compliance with traceability obligations (Article 12).

Requirements:

- Automatic event recording
- Traceability of decisions and their context
- Evidential preservation of logs
- Audit capability by supervisory authorities

QuantumLock contribution:

- Automatic generation of compliant Evidence Bundles
- Cryptographic proof of traceability
- Enduring archive resistant to tampering
- Verification mechanism for competent authorities

9 Evaluation and Acceptance Criteria

9.1 Verifiable Security Properties

A compliant QuantumLock system must demonstrate the following properties:

Property	Verification Criterion
Post-quantum resistance	Use of ML-DSA-65 or SLH-DSA compliant with NIST FIPS 204/205
Immutability	Impossibility to modify bundle without invalidating signatures
Independent verification	Validation possible without access to original system
Non-repudiation	Cryptographic proof of issuer
Reliable timestamping	Qualified TSA or certified temporal source
Long-term integrity	Bundle preservation on evidential media (10+ years)

Table 3: Security properties and verification criteria

9.2 Performance and Scalability

Critical AI systems require performance compatible with operational constraints:

- **Generation latency:** <100ms per Evidence Bundle
- **Computational overhead:** <5% on AI inference
- **Bundle size:** 10-50 KB (metadata + signatures)
- **Throughput:** Support for 1000+ inferences/second
- **Verification:** <50ms per bundle in offline mode

9.3 Regulatory Compliance Criteria

AI Act Article 12 Compliance Matrix:

AI Act Requirement	QuantumLock Mechanism
Automatic event recording	Automated context capture
Operation traceability	Evidence Bundle with complete meta-data
Input data identification	Cryptographic hash of inputs
Reliable timestamping	Qualified TSA or certified clock
Audit capability	Independent offline verification
Appropriate preservation	Long-term evidential storage

Table 4: Compliance with AI Act Article 12 requirements

10 Roadmap and Integration

10.1 Deployment Phases

Phase 1: Proof of Concept (3 months)

- Deployment on test environment with non-critical AI model
- Generation and verification of proof bundles
- Performance and integration validation
- Security audit and cryptographic review

Phase 2: Operational Pilot (6 months)

- Integration with critical AI system in production
- Deployment in hybrid mode (parallel generation, no interruption)
- Collection of user feedback and optimization
- Preparation of certification documentation

Phase 3: Generalization (12 months)

- Extension to all high-risk AI systems
- Integration with AI governance processes
- Training of technical and legal teams
- Establishment of enduring archiving

10.2 Technical Integration

With MLOps pipelines:

- REST API for integration into training workflows
- Python SDK for automatic capture during inference
- Plugins for popular frameworks (TensorFlow, PyTorch, ONNX)
- Connectors for governance tools (MLflow, Weights & Biases)

With security infrastructure:

- Integration with existing PKI
- Export to SIEM for correlation (metadata only)
- Storage on evidential file systems (WORM)
- HSM support for critical signature keys

With archiving systems:

- Automatic export to digital vaults
- Configurable retention according to classification
- Replication and backup mechanisms
- Compatibility with electronic archiving standards

10.3 Training and Support

QuantumLock deployment requires support for three populations:

Technical teams:

- Integration and configuration training
- Troubleshooting and maintenance guide
- API and SDK documentation

Security and governance managers:

- Awareness of cryptographic proof stakes
- Audit and verification processes
- Legal framework and regulatory compliance

Auditors and oversight authorities:

- Independent verification training
- Use of Verification Toolkit
- Interpretation of compliance reports