

Article

Not peer-reviewed version

ExecMesh: Cryptographically Verifiable AI Provenance for Regulatory Compliance

[Panagiotis Karmiris](#) *

Posted Date: 5 December 2025

doi: 10.20944/preprints202511.1085.v2

Keywords: zero-knowledge proofs; regulatory compliance; verifiable compute; hybrid verification; AI provenance; audit trails; FDA compliance; EU AI Act; GxP integration



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

ExecMesh: Cryptographically Verifiable AI Provenance for Regulatory Compliance

Panagiotis Karmiris

Independent Researcher; unbinder@msn.com

Abstract

ExecMesh introduces **cryptographically verifiable computation** as a foundational primitive for regulatory compliance and audit trail requirements in AI/ML systems [1–3]. By combining commitment-based verification with secure multi-party oracles and a two-tier regulatory architecture, ExecMesh enables enterprises to meet FDA, SEC, and EU AI Act requirements while maintaining the benefits of decentralized infrastructure. **Immediate Value Proposition:** ExecMesh provides immediate value as an audit trail and provenance layer for regulated AI systems, independent of advances in zero-knowledge proof technology. Even without full verification of large neural networks, the system delivers cryptographic guarantees for data integrity, execution timestamps, and pipeline reproducibility—meeting core regulatory requirements today.

Keywords: zero-knowledge proofs; regulatory compliance; verifiable compute; hybrid verification; AI provenance; audit trails; FDA compliance; EU AI Act; GxP integration

Key Contributions:

- A two-tier architecture separating permissionless verification from permissioned financialization
- Hybrid verification combining ZK proofs (simple tasks) with optimistic verification (complex tasks)
- Federated oracle system with economic security guarantees for compute pricing
- Enterprise-first go-to-market targeting FDA-regulated AI systems with drop-in integration for existing GxP/QMS workflows
- Realistic technical scope acknowledging current zkSNARK limitations [4,5]
- Optional financialization roadmap (see Appendix D) for future compute-backed DeFi primitives

Target Market: Pharmaceutical AI, medical devices, financial services, and autonomous vehicles requiring cryptographic audit trails for regulatory compliance [2,6]. Addressable market: \$2B+ in compliance spending (2026–2028). **Note:** Financial primitives (compute credits as collateral, futures markets) are described in Appendix D and represent Phase 4+ functionality, dependent on regulatory approvals and demonstrated enterprise adoption.

1. Introduction

1.1. The Regulatory Compliance Crisis in AI

AI systems in healthcare, finance, and autonomous vehicles face an unprecedented regulatory challenge: proving to regulators that computational work was performed correctly, data was not tampered with, and model outputs are reproducible [1–3,6]. Traditional centralized logging systems are insufficient because:

- **Logs can be tampered with:** Centralized databases do not provide cryptographic guarantees of integrity.
- **No global auditability:** Regulators must trust the organization's internal systems rather than verifying claims independently.
- **Expensive manual audit trails:** Pharmaceutical companies spend millions on manual documentation for FDA submissions.
- **Lack of provenance tracking:** Training pipelines involving multiple vendors create verification gaps.

1.2. ExecMesh: The Enterprise Audit Trail Layer

ExecMesh provides a **drop-in audit trail layer** that integrates with existing enterprise workflows while providing cryptographic guarantees. Unlike blockchain compute marketplaces, ExecMesh is purpose-built for regulatory compliance:

- **Works with your existing infrastructure:** No need to migrate computation—ExecMesh wraps existing processes with verification.
- **Integrates with GxP/QMS:** Designed to complement, not replace, existing quality management systems.
- **FDA-ready exports:** One-click PDF generation of cryptographically anchored audit trails for regulatory submissions.
- **No blockchain expertise required:** Enterprise API abstracts all blockchain complexity.

1.3. Core Innovation: Useful Even Without Full ZK Verification

A critical design decision: ExecMesh delivers regulatory value **independent of advances in zero-knowledge proof technology**. Even though full verification of GPT-scale models is currently infeasible [7,8], ExecMesh provides:

1. **Data Integrity Proofs:** Cryptographic guarantees that training datasets were not tampered with.
2. **Execution Timestamps:** Immutable records of when computation occurred (critical for patent priority, audit trails).
3. **Pipeline Provenance:** Verifiable chains of custody for multi-step workflows (data cleaning → training → validation).
4. **Reproducibility Guarantees:** Proof that identical inputs and configurations were used across runs.
5. **Dispute Resolution:** Cryptographic evidence for arbitration when vendors disagree.

Future improvements in zkML technology (Halo2, Plonky2, recursive SNARKs [7–9]) will enable migration of more workload into fully verified circuits, but the core value proposition remains intact even if those improvements are slower than expected.

1.4. Structure of This Document

Core Protocol (Sections 2–12): Production-ready regulatory compliance infrastructure available today.

Future Financialization (Appendix D): Optional DeFi integration (compute credits as collateral, futures markets) dependent on Phase 4+ adoption and regulatory approvals. These features are aspirational and not required for core regulatory value.

1.5. The Opportunity

Three convergent forces create a unique market opening:

1. **AI Regulation Wave** – FDA guidance on AI/ML medical devices, EU AI Act enforcement (2026), and SEC model validation requirements create mandatory demand for stronger, evidence-backed audit trails [2,3,6].
2. **Layer-2 Economics** – Sub-\$0.01 transaction costs on Optimism and Arbitrum make on-chain verification economically viable for enterprise workloads [10,11].
3. **Integration with Existing Systems** – ExecMesh complements rather than replaces existing GxP, QMS, and model risk management frameworks.

ExecMesh provides cryptographically verifiable computation infrastructure that enterprises can use for regulatory compliance while maintaining compatibility with traditional quality management processes.

2. Design Goals and Assumptions

2.1. Design Goals

G1: Verifiable Execution with Minimized Trust. Each task's integrity must be provable via succinct cryptographic evidence. We minimize, but do not eliminate, trust assumptions where technically infeasible [4,5].

G2: Regulatory Compliance by Design. System architecture enables FDA 21 CFR Part 11, SEC, and EU AI Act compliance through cryptographic audit trails and immutable record-keeping [1–3].

G3: Drop-In Integration for Enterprises. Works with existing infrastructure, GxP/QMS processes, and quality management workflows without requiring blockchain expertise.

G4: Hybrid Verification. Support both zkSNARK proofs (for simple tasks) and optimistic verification (for complex ML workloads) within a unified framework, aligning with the current performance envelope of modern proof systems [4,7].

G5: Enterprise-First Economics. Sustainable revenue from compliance-driven customers, not speculative marketplace dynamics.

G6: Optional Financialization. DeFi features (described in Appendix D) are strictly optional and do not affect core regulatory compliance value.

2.2. Assumptions

Assumption 1 (Honest Majority). *Most participants act rationally to preserve reputation; economic penalties and challenge periods deter malicious behavior.*

Assumption 2 (Federated Oracle Trust). *Price feeds from multiple independent sources with economic staking create acceptable security for collateral valuation. Full trustlessness is not achievable for off-chain data [12].*

Assumption 3 (L2 Viability). *Transaction costs on rollups remain below \$0.10 per verification, enabling economical enterprise deployment [10,11].*

Assumption 4 (zkSNARK Security). *Groth16 on BN254 remains secure under the discrete log assumption through 2030; quantum migration planned for 2030+ [4,5].*

2.3. Explicit Assumptions and Limitations

We acknowledge the following limitations in ExecMesh v4.0:

Technical Limitations

1. **ZK Circuit Complexity:** Full verification of large AI models (>100M parameters) is not feasible with current technology. We use hybrid verification (commitment + challenge period) for complex tasks [7,8,13].
2. **Oracle Trust:** Our price oracles rely on a federated trust model with economic security, not pure cryptographic trustlessness [12].
3. **L2 Dependency:** Protocol relies on Layer 2 networks (Optimism, Arbitrum) for cost-effective verification. L2 security assumptions apply [10,11].
4. **Groth16 Quantum Vulnerability:** Current cryptography is vulnerable to quantum computers (expected 2035+). See Section ?? for a migration plan.

Economic Limitations

1. **Initial Liquidity:** Enterprise-first strategy means limited marketplace liquidity in Year 1–2.
2. **Regulatory Compliance Costs:** \$1M+ annually for full financial layer compliance (Phase 4+).
3. **Collateral Volatility:** Compute credit values (if/when financialized) are subject to market fluctuations; liquidation risk for borrowers.

Regulatory Limitations

1. **US-Only Initially:** Financial features (Appendix D) restricted to US users (Phase 4+).
2. **Accredited Investor Requirement:** Fractional ownership limited to accredited investors (Reg D).
3. **CFTC Registration Timeline:** Futures market delayed until 2027 pending regulatory approval.

These limitations do not prevent ExecMesh from delivering immediate value in regulatory compliance and audit trail use cases.

3. Worked Example: FDA Compliance Workflow

This section demonstrates how a pharmaceutical company uses ExecMesh for FDA 21 CFR Part 11 and AI/ML SaMD compliance with a concrete, end-to-end workflow.

3.1. Scenario: AI-Powered Pathology Diagnostic

Company: PathAI Pharma (hypothetical)

Product: AI model for detecting cancer in histopathology images

Regulatory Requirement: FDA premarket approval (PMA) requiring complete audit trail of:

1. Training dataset provenance and integrity
2. Model training reproducibility
3. Validation/testing procedures
4. Post-market surveillance of inference outputs

3.2. Phase 1: Dataset Ingestion and Integrity

Step 1a: Data Collection

- PathAI receives 50,000 de-identified histopathology images from partner hospitals
- Each image is hashed locally: $h_i = \text{SHA-256}(\text{image}_i)$
- Dataset manifest created: $\{h_1, h_2, \dots, h_{50000}\}$

Step 1b: ExecMesh Registration

Listing 1: Dataset Registration API Call

POST /api/v1/datasets/register

```
Body: {
  "datasetId": "PMA-2026-001-TRAINING",
  "imageHashes": ["0xabc123...", "0xdef456...", ...],
  "timestamp": "2026-01-15T10:30:00Z",
  "hospitalSources": ["Hospital-A", "Hospital-B", ...]
}
```

Step 1c: On-Chain Commitment

- ExecMesh creates Merkle tree of all image hashes
- Root hash committed to Ethereum L2 (Optimism): $\text{Root} = \text{MerkleRoot}(\{h_1, \dots, h_{50000}\})$
- Transaction hash: 0x789abc... (permanent, immutable record)
- Gas cost: \$0.02 (single transaction for entire dataset)

Regulatory Value: FDA auditor can verify that training data was not tampered with by:

1. Recomputing image hashes from provided dataset
2. Reconstructing Merkle tree
3. Comparing root hash to on-chain commitment

3.3. Phase 2: Model Training with Provenance

Step 2a: Training Job Specification

Listing 2: Training Job Submission

POST /api/v1/training/submit

```
Body: {
  "modelId": "ResNet50-Cancer-v1",
  "datasetId": "PMA-2026-001-TRAINING",
  "hyperparameters": {
    "learning_rate": 0.001,
    "batch_size": 32,
    "epochs": 100
  },
  "dockerImage": "pathaimodels/resnet50:v2.1",
  "seed": 42
}
```

Step 2b: Training Execution

- PathAI runs training on their existing GPU cluster (no migration required)

- ExecMesh monitoring agent records:
 - Start time: 2026-01-20T09:00:00Z
 - End time: 2026-01-22T14:30:00Z
 - Final model weights hash: $h_{\text{model}} = \text{SHA-256}(\text{weights.pth})$
 - Training logs hash: $h_{\text{logs}} = \text{SHA-256}(\text{training.log})$

Step 2c: Commitment Proof Submission

```
// Simple commitment proof (no full zkSNARK needed)
struct TrainingRecord {
    bytes32 datasetRootHash;          // Links to Phase 1
    bytes32 modelWeightsHash;
    bytes32 hyperparametersHash;
    uint256 startTime;
    uint256 endTime;
    address executor;                // PathAI's verified address
    bytes executorSignature;         // Cryptographic signature
}
```

Listing 3: On-Chain Training Record

Gas cost: \$0.005

Regulatory Value: FDA can verify reproducibility by:

1. Obtaining identical dataset (verified via Phase 1 hashes)
2. Using identical hyperparameters and seed (committed on-chain)
3. Rerunning training
4. Comparing final model weights hash

3.4. Phase 3: Validation and Testing

Step 3a: Validation Dataset Registration

- Separate 10,000-image validation set registered (same process as Phase 1)
- Committed to chain with different dataset ID: PMA-2026-001-VALIDATION

Step 3b: Inference Pipeline Verification

Listing 4: Validation Run

```
POST /api/v1/inference/batch
Body: {
  "modelId": "ResNet50-Cancer-v1",
  "modelWeightsHash": "0x789def...",
  "datasetId": "PMA-2026-001-VALIDATION",
  "outputFormat": "predictions_with_confidence"
}
```

Step 3c: Results Commitment

- For each validation image:
 - Input hash: $h_{\text{input}} = \text{SHA-256}(\text{image})$
 - Output hash: $h_{\text{output}} = \text{SHA-256}(\text{prediction} \parallel \text{confidence})$
 - Pair $(h_{\text{input}}, h_{\text{output}})$ stored on-chain
- Accuracy metrics: {Sensitivity: 94.2%, Specificity: 96.8%}
- Metrics hash committed alongside results

Gas cost (10,000 images, batched): \$2.50 total

Regulatory Value: FDA can spot-check any subset of validation results:

1. Request specific images from PathAI
2. Verify image hash matches on-chain commitment
3. Rerun inference with committed model weights
4. Verify output hash matches commitment

3.5. Phase 4: FDA Submission Package

Step 4a: Audit Trail Export

Listing 5: Generate FDA Submission Package

GET /api/v1/audit/**export**

Query Parameters:

```
projectId=PMA-2026-001
format=pdf
includeProofs=true
```

Step 4b: Generated PDF Contents

1. **Executive Summary**
 - Dataset: 50,000 images from 15 hospitals
 - Model: ResNet50 architecture
 - Performance: 94.2% sensitivity, 96.8% specificity
2. **Section A: Data Provenance**
 - Merkle root: 0xabc123...
 - Ethereum L2 transaction: 0x789abc...
 - Block number: 12,345,678
 - Timestamp: 2026-01-15T10:30:00Z
 - Verification instructions for FDA auditor
3. **Section B: Training Provenance**
 - Hyperparameters (committed on-chain)
 - Model weights hash
 - Training duration: 53.5 hours
 - On-chain commitment transaction
4. **Section C: Validation Results**
 - Per-image results table (sample):

Table 1. Validation Results (Sample) with On-Chain Commitments.

Image ID	Input Hash	Output Hash	On-Chain Tx
IMG-0001	0xabc...	0xdef...	0x111...
IMG-0002	0x123...	0x456...	0x222...
...

4. **Section D: Cryptographic Verification Guide**
 - Step-by-step instructions for FDA auditor
 - Links to open-source verification tools
 - Example verification commands

3.6. Phase 5: Post-Market Surveillance

After FDA approval, PathAI deploys model to production hospitals. ExecMesh continues providing compliance value:

Ongoing Inference Logging

- Every inference in production: $(h_{\text{input}}, h_{\text{output}}, \text{timestamp})$ committed
- Batched daily to reduce costs: \$5/day for 1,000 inferences
- Enables retrospective audit of any flagged case

Adverse Event Investigation If a patient outcome raises questions:

1. Hospital provides case ID
2. PathAI looks up commitment: found on-chain at block 15,234,567
3. Verifies exact input image and model output used
4. Demonstrates to FDA/hospital that model performed as validated

3.7. Total Cost and Timeline

Costs:

- Setup and integration: 2 engineer-weeks (\$20k labor)
- On-chain commitments (one-time): \$15 total
- Ongoing (1 year post-market): \$1,825 ($\$5/\text{day} \times 365$)
- Total Year 1: \$21,840

Traditional Manual Audit Trail Cost: \$500k–2M annually for comparable documentation quality [6].

Timeline:

- ExecMesh integration: 1–2 weeks
- No disruption to existing ML workflows
- FDA submission package generation: 1 hour (automated)

3.8. Key Takeaways

1. **Drop-in integration:** PathAI didn't migrate compute—ExecMesh wrapped existing processes.
2. **No full ZK proofs needed:** Commitments (hashes) provide sufficient regulatory guarantees for most steps.
3. **Economically viable:** Sub-\$25k annual cost vs. \$500k+ for manual documentation.
4. **Cryptographically verifiable:** FDA can independently verify claims without trusting PathAI's internal systems.
5. **Backward compatible:** Complements existing GxP and 21 CFR Part 11 processes rather than replacing them.

4. System Architecture

5. Financial Primitives

6. Economic Model

7. Security Model & Risk Mitigation

8. Regulatory Requirements and Compliance Mapping

9. Discussion: Limitations, Criticisms, and Research Agenda

10. Regulatory Strategy: Two-Tier Architecture

11. Technical Implementation

12. Roadmap & Future Work

13. Conclusions

ExecMesh v4.0 provides a pragmatic, production-ready solution for cryptographic verification of computational work in regulated industries. The system delivers immediate compliance value through commitment-based verification and audit trail generation, independent of advances in full zero-knowledge proof technology for large AI models.

Our hybrid verification model, federated oracle design, explicit regulatory mapping, and two-tier regulatory architecture enable enterprise adoption today. The worked FDA example (Section 3) demonstrates how pharmaceutical companies can integrate ExecMesh with existing GxP workflows for sub-\$25k annual cost—replacing \$500k+ manual audit processes with cryptographically verifiable automation.

Key design decisions ensure practical deployment:

- **Drop-in integration:** Works with existing infrastructure; no compute migration required
- **Enterprise-first:** Sustainable revenue from compliance-driven customers with non-optional needs
- **Realistic scope:** Core value independent of zkML breakthroughs
- **Optional financialization:** DeFi features (Appendix D) available but not required

The enterprise-first go-to-market strategy targets 3–5 anchor customers in Year 1 (pharmaceutical AI, medical devices), expanding horizontally to financial services and autonomous vehicles in Year 2+. Financial primitives described in Appendix D represent Phase 4+ optionality, dependent on demonstrated enterprise adoption and regulatory approvals.

This is a foundational protocol for verifiable AI: work is auditable, provenance is cryptographic, and regulatory compliance is built-in by design.

14. Acknowledgments

This work builds upon research from Satoshi Nakamoto, Vitalik Buterin et al., Gavin Wood, Jens Groth, Eli Ben-Sasson and collaborators, and the Aave and Uniswap teams, among many others [4,5,7,14–18]. Special thanks to the critical reviewers whose feedback significantly strengthened this whitepaper.

Appendix A Mathematical Proofs & Formal Verification

Appendix A.1 Collateral Safety Theorem

Theorem A1 (Enhanced Collateral Safety). *For any locked compute credit with collateralization ratio $R > 150\%$ and graduated liquidation thresholds, the vault remains solvent under price drops up to $(R - 100)/R$ with partial liquidation restoring health before full insolvency.*

Proof. Let V_0 be the initial credit value and D the debt issued. By definition $V_0 = R \times D$ with $R > 1.5$. The graduated liquidation system triggers at $H = 1.1$ (110% collateralization):

$$V_{\text{partial}} = 1.1 \times D.$$

The safe decline to partial liquidation is:

$$\frac{V_0 - V_{\text{partial}}}{V_0} = \frac{RD - 1.1D}{RD} = \frac{R - 1.1}{R}.$$

For $R = 1.5$, this equals $\frac{1.5-1.1}{1.5} \approx 26.7\%$. Partial liquidation then restores health to $H = 1.3$, providing additional buffer. The system remains solvent through this graduated approach. \square \square

Appendix A.2 Proof Composition Soundness

Theorem A2 (Soundness of Composed Proofs). *If individual proofs π_1, \dots, π_n are valid and the composition proof π_c references their output hashes using a collision-resistant hash function H , then acceptance of π_c implies the prover used the genuine dependency outputs with negligible probability of forgery.*

Proof. Let $D = \{d_i\}$ be dependency outputs and $h_i = H(d_i)$ their hashes. Each π_i proves knowledge of d_i such that $H(d_i) = h_i$ with soundness error ϵ_{zk} . The composed proof π_c proves knowledge of inputs whose hashes equal $\{h_i\}$ and a transformation producing output O . By the collision-resistance of H , the probability of finding $d'_i \neq d_i$ such that $H(d'_i) = h_i$ is negligible ϵ_{cr} . Thus the probability that π_c verifies but uses incorrect dependencies is bounded by

$$\Pr[\text{forgery}] \leq \epsilon_{zk} + \epsilon_{cr} + \text{negl}(\lambda),$$

which is negligible in security parameter λ . \square \square

Appendix A.3 Oracle Manipulation Resistance

Theorem A3 (Oracle Manipulation Resistance). *Given $n \geq 3$ independent price sources with at least $2/3$ honest sources, the weighted median calculation resists manipulation requiring control of more than $n/2$ sources.*

Proof. Let S be the set of price sources with $|S| = n$, and $H \subset S$ be the honest sources with $|H| \geq \lceil 2n/3 \rceil$. The weighted median calculation sorts sources by price and selects the price where cumulative weight reaches 50% of total. To manipulate the median downward, an attacker must control enough sources to push the cumulative weight of lower prices past 50%. This requires controlling sources with total weight $> 50\%$. But honest sources have total weight $\geq 2/3$ of total weight, which already exceeds 50%, so the median must include at least one honest source price. Thus manipulation requires controlling $> 50\%$ of total weight, which contradicts the $2/3$ honest assumption. \square \square

Appendix B Circuit Specifications & Benchmarks

Appendix C API Reference & Integration Guide

Appendix D Future Financialization Roadmap (Phase 4+)

Important Context

This appendix describes **aspirational DeFi features** that represent Phase 4+ functionality (2027+), dependent on:

1. Demonstrated enterprise adoption and revenue from core compliance business (Phases 1–3)
2. Regulatory approvals (CFTC registration for futures, SEC guidance on compute credits as securities)
3. Proven demand from compute providers and buyers for financial instruments
4. Mature oracle infrastructure with demonstrated manipulation resistance

These features are optional and do not affect the core regulatory compliance value proposition. ExecMesh delivers immediate value as an audit trail layer (Sections 1–12) independent of whether any DeFi integration occurs.

Regulatory Warnings

- Compute credits used as loan collateral may be classified as securities in some jurisdictions
- Compute futures require CFTC registration as commodity derivatives (US)
- Fractional ownership tokens are likely securities under Howey Test
- KYC/AML requirements apply to all financial features
- Accredited investor restrictions may apply (Regulation D)

Timeline: Earliest possible deployment is Q4 2027, subject to regulatory approvals.

Revenue Model Clarification

The financial projections in this appendix (e.g., "\$50B DeFi TVL") assume:

- ExecMesh has already achieved \$5M+ annual revenue from 25+ enterprise compliance customers
- Successful deployment and 2+ years of operation without security incidents
- Multiple third-party DeFi protocols have integrated compute credits
- Liquid secondary markets for compute credits exist

Do not interpret these projections as Year 1–3 forecasts. They represent Phase 4+ potential if all prerequisites are met.

Appendix D.1 Introduction: Compute as an Asset Class

Cloud computing represents a \$200B+ annual market, yet unlike other commodities (oil, gold, wheat), computational resources cannot be:

- Traded as financial instruments
- Used as loan collateral
- Hedged against price fluctuations

- Owned fractionally by retail investors

ExecMesh's **optional financial layer** (Phase 4+) introduces compute-backed financial primitives, creating a new asset class with characteristics similar to bonds, REITs, and commodity futures.

Critical dependencies:

1. Regulatory approval for each financial product
2. Proven enterprise adoption establishing baseline compute credit value
3. Mature oracle infrastructure (\$100M+ in economic security)
4. Integration with at least 2 major DeFi protocols (Aave, Compound, etc.)

Appendix D.2 Compute Credits: The Base Instrument

Appendix D.2.1 Definition

A **Compute Credit** is an ERC-1155 NFT representing cryptographically verified computational work, issued after successful completion of a job on the core ExecMesh protocol.

Appendix D.2.2 Properties

Fungibility Credits of same type/quality are interchangeable

Verifiability ZK proof or commitment ensures work was performed correctly

Liquidity Tradable on secondary markets (OpenSea, Uniswap) if DeFi integration occurs

Intrinsic Value Represents actual computational capacity

Appendix D.2.3 Valuation Formula

$$V(\text{credit}) = C_{\text{market}} \times Q_{\text{factor}} \times R_{\text{score}} \times T_{\text{decay}}$$

where:

- C_{market} = Current market rate for computation type
- Q_{factor} = Quality factor based on provider tier (0.8–1.2)
- R_{score} = Provider reputation score (0.0–1.0)
- T_{decay} = Time decay factor for expiring credits

Appendix D.3 DeFi Collateral Integration

Phase 4 Feature: ExecMesh enables Compute Credits to function as collateral in major DeFi lending protocols (Aave, Compound, MakerDAO), subject to:

- DAO governance approval from each protocol
- Demonstrated price stability and liquidity
- Regulatory classification as non-securities (or compliance with securities laws)

Appendix D.3.1 Architecture

```
1 interface IComputeCollateralAdapter {
```

```

2 // Called by lending protocol to value collateral
3 function getCollateralValue(uint256 tokenId)
4     external view returns (uint256);
5
6 // Liquidate under-collateralized position
7 function liquidate(address borrower, uint256 amount)
8     external;
9
10 // Check if position is healthy
11 function isHealthy(address borrower)
12     external view returns (bool);
13 }

```

Listing 6: Compute Collateral Adapter Interface (Phase 4+)

Appendix D.3.2 Collateralization Process

1. **Deposit:** Provider deposits Compute Credit NFTs into vault
2. **Valuation:** Oracle determines current market value
3. **Borrowing:** Provider can borrow up to LTV ratio (typically 75–80%)
4. **Interest:** Borrower pays interest (e.g., 6% APY) to lenders
5. **Liquidation:** If collateral value drops, graduated liquidation occurs

Appendix D.3.3 Example Use Case

Provider Alice:

- Owns 10 Compute Credits valued at \$10,000 total
- Deposits in ExecMesh Vault
- Borrows \$7,500 USDC (75% LTV)
- Pays 6% APY interest (\$450/year)
- Uses borrowed funds for business operations
- Repays loan when cash flow improves

Appendix D.3.4 Risk Management: Graduated Liquidation

Multi-tier liquidation thresholds minimize losses:

Table A2. Graduated Liquidation Safeguards (Phase 4+).

Health Factor	State	Action	Penalty
$H > 1.5$	Healthy	None	0%
$1.3 < H \leq 1.5$	Warning	Email/alert	0%
$1.1 < H \leq 1.3$	At Risk	Collateral top-up required	0%
$1.0 < H \leq 1.1$	Critical	Partial liquidation (50%)	2%
$H \leq 1.0$	Underwater	Full liquidation	5%

Appendix D.4 Compute Futures Market

Phase 5 Feature (2027+): Trade future computational capacity at predetermined prices, subject to CFTC registration.

Appendix D.4.1 Concept

- Price risk hedging for customers
- Capacity monetization for providers
- Market-based price discovery
- Liquidity provision by speculators

Appendix D.4.2 Contract Specification

Table A3. Compute Futures Contract Specification (Phase 5+)

Parameter	Specification
Underlying	Specific compute type (e.g., "100 GPU-hours A100")
Contract Size	Standardized units
Expiry	1, 3, 6, 12 months
Settlement	Physical delivery or cash settlement
Margin	10–20% of contract value

Appendix D.4.3 Example: Customer Hedging

Pharmaceutical company scenario:

- Needs \$500k compute in 6 months
- Current spot price: \$5/GPU-hour
- 6-month futures: \$5.50/GPU-hour
- Action: Buy futures contract locking in \$5.50 price

Outcome if spot price rises to \$8:

- Without hedge: Pays \$800k
- With hedge: Pays \$550k (locked-in price)
- Savings: \$250k

Appendix D.5 Market Size and Economic Impact

Aspirational Phase 4+ Projections

Table A4. Compute-Backed Assets: Total Addressable Market (Aspirational Phase 4+).

Asset Class	Current Market	Compute Potential (Phase 4+)
Corporate Bonds	\$10T	\$2–5T
REITs	\$4T	\$1–2T
Commodity Futures	\$2T	\$500B–1T
Total	\$16T	\$3.5–8T

Critical caveat: These numbers assume ExecMesh has already:

1. Demonstrated 3+ years of security and regulatory compliance
2. Achieved \$50M+ annual revenue from enterprise customers

3. Obtained all necessary regulatory approvals
4. Integrated with major DeFi protocols

Do not interpret as near-term projections.

Appendix D.6 Integration with Existing DeFi Protocols

Phase 4+ Dependencies

Each DeFi integration requires:

- DAO governance vote approving compute credits as collateral
- Custom adapter contract (e.g., AaveComputeAdapter)
- Demonstrated price stability (\$100M+ TVL in compute credits)
- Oracle security (\$100M+ in economic security from staked operators)

Appendix D.6.1 Example: Aave V3 Integration

```

1 contract AaveComputeAdapter is IAaveV3Adapter {
2     IComputePriceOracle public oracle;
3
4     function getAssetPrice(address asset)
5         external view returns (uint256) {
6         return oracle.getPrice(asset);
7     }
8
9     function liquidationCall(
10        address collateral,
11        address debt,
12        address user,
13        uint256 debtToCover
14    ) external {
15        // Graduated liquidation for Compute Credits
16        uint256 liquidationAmount = calculateGraduatedAmount(
17            collateral,
18            debtToCover
19        );
20
21        _liquidate(user, liquidationAmount);
22    }
23 }

```

Listing 7: Aave Compute Adapter (Phase 4+)

Appendix D.7 Regulatory Considerations

Phase 4+ Compliance Requirements

Appendix D.7.1 Securities Classification

Compute Credits Likely utility tokens (represent actual compute capacity), but may be securities in some jurisdictions

Fractional Tokens Almost certainly securities under Howey Test (investment contract)

Futures Contracts Commodity derivatives (CFTC jurisdiction in U.S.)

Appendix D.7.2 Required Registrations

1. **KYC/AML:** Required for transactions exceeding \$10,000
2. **Accredited Investor:** May be required for fractional tokens
3. **Commodity Trading:** CFTC registration for futures platform
4. **Tax Reporting:** 1099 forms for U.S. users
5. **GDPR:** Data protection for European users

Appendix D.7.3 Risk Disclosures

All Phase 4+ users must acknowledge:

- Market volatility risk
- Liquidation risk in leveraged positions
- Smart contract vulnerabilities
- Regulatory uncertainty
- No deposit insurance

Appendix D.8 Phase 4+ Timeline and Prerequisites

Earliest Deployment: Q4 2027

Phase 4 Prerequisites (2027)

- \$5M+ annual revenue from 25+ enterprise customers (Phase 1–3)
- 2+ years operating history without major security incidents
- Oracle network with \$100M+ economic security
- At least one major DeFi protocol partnership (Aave or Compound)
- KYC/AML provider integration
- Legal budget: \$1M+ for regulatory compliance

Phase 5 Prerequisites (2028+)

- CFTC registration as Designated Contract Market (DCM)
- \$50M+ TVL in compute credits demonstrating price stability
- Secondary market liquidity (Uniswap V3 pools with \$10M+ depth)
- Futures market beta with \$1M+ daily volume

Appendix D.9 Summary: Optional Financialization

The DeFi features described in this appendix are:

1. **Strictly optional:** Core ExecMesh (Sections 1–12) delivers value without any financial integration
2. **Dependent on enterprise success:** Only viable after proven adoption from compliance customers
3. **Subject to regulatory approval:** Deployment timeline depends on CFTC, SEC, and international regulators
4. **High capital requirements:** Require \$100M+ in economic security and \$1M+ annual legal/compliance costs
5. **Long-term vision:** Represent Phase 4–5 functionality (2027–2028+), not Year 1–3

Do not conflate the core compliance value proposition with these aspirational financial features. ExecMesh succeeds as an audit trail layer even if no DeFi integration ever occurs.

References

1. U.S. Food and Drug Administration. Electronic Records; Electronic Signatures. *Code of Federal Regulations*, Title 21, Part 11, 2023. 21 CFR Part 11.
2. European Parliament.; Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council on Artificial Intelligence (AI Act). Official Journal of the European Union, 2024.
3. Board of Governors of the Federal Reserve System.; Office of the Comptroller of the Currency. SR 11-7: Guidance on Model Risk Management. Supervisory Letter SR 11-7, 2011.
4. Groth, J. On the Size of Pairing-Based Non-Interactive Arguments. In Proceedings of the Advances in Cryptology – EUROCRYPT 2016. Springer, 2016, Vol. 9666, *Lecture Notes in Computer Science*, pp. 305–326. https://doi.org/10.1007/978-3-662-49896-5_11.
5. Ben-Sasson, E.; Chiesa, A.; Genkin, D.; Tromer, E.; Virza, M. SNARKs for C: Verifying Program Executions Succinctly and in Zero Knowledge. In Proceedings of the Advances in Cryptology – CRYPTO 2013. Springer, 2013, Vol. 8043, *Lecture Notes in Computer Science*, pp. 90–108. https://doi.org/10.1007/978-3-642-40084-1_6.
6. U.S. Food and Drug Administration. Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning-Based Software as a Medical Device. Discussion paper and request for feedback, 2019.
7. Ben-Sasson, E.; Bentov, I.; Horesh, Y.; Riabzev, M. Scalable, Transparent, and Post-Quantum Secure Computational Integrity. *IACR Cryptology ePrint Archive* **2018**, 2018, 46.
8. Farmer, B.; Alefirenko, M.; et al. Plonky2: Fast Recursive Proofs with Plonk and FRI. Polygon Zero technical report, 2022.
9. Bowe, S.; Grigg, J.; Hopwood, D.; et al. Halo 2: Recursive Proof Composition without a Trusted Setup. Zcash protocol documentation and design notes, 2020. Zcash Foundation / Electric Coin Company design.
10. Optimism Collective. Optimism: Documentation and Protocol Overview. Online documentation, 2024.
11. Offchain Labs. Arbitrum: Developer Documentation. Online documentation, 2024.
12. Nazarov, S.; Ellis, S.; Juels, A.; et al. Chainlink: A Decentralized Oracle Network. Whitepaper, 2017.
13. Gabizon, A.; Williamson, Z.J.; Ciobotaru, O. PLONK: Permutations over Lagrange-bases for Oecumenical Noninteractive arguments of Knowledge. *IACR Cryptology ePrint Archive*, Report 2019/953, 2019.
14. Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System. *Bitcoin.org Whitepaper* **2008**.
15. Buterin, V. A Next-Generation Smart Contract and Decentralized Application Platform. *Ethereum Whitepaper*, 2014.
16. Wood, G. Ethereum: A Secure Decentralised Generalised Transaction Ledger. *Ethereum Yellow Paper*, 2014.
17. Aave Protocol. Aave Protocol v2: Documentation. Protocol documentation, 2020.
18. Adams, H.; Zinsmeister, N.; Salem, M.; Keefer, R.; Robinson, D. Uniswap v3 Core. In Proceedings of the Uniswap v3 Core Whitepaper, 2021.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.