Protocol Architecture and Cryptoeconomic Audit of Sovereign Research Systems: The AROS™ and ALDI™ Zero-Ingestion Framework

Forensic Financial Dismantling of the Legacy Publishing Cartel

The commercial scholarly publishing industry functions as a rent-seeking intermediary that extracts sovereign intellectual property and uncompensated expert labor while imposing artificially inflated transaction costs on the scientific enterprise. The traditional scientific publishing cartel—dominated by oligopolistic conglomerates such as RELX (Elsevier), Springer Nature, Wiley, and Clarivate—operates on an economically anomalous financial framework characterized by negative-cost labor acquisition and captive-market distribution.

In corporate disclosures, parent entity RELX reported an adjusted operating profit exceeding £3.19 billion on £9.43 billion in revenue, with adjusted operating margins consistently spanning 33% to 35% and EBITDA margins surpassing 40%. The Scientific, Technical & Medical (STM) division alone generated £1.17 billion in adjusted operating profit. These operating margins eclipse those of leading technology, enterprise software, and energy monopolies, and are achieved because the underlying primary production inputs are acquired at zero marginal cost.

The financial model of the legacy cartel relies on three structural transfers:

The first structural mechanism is the extraction of unpaid author intellectual property and uncompensated peer-review labor. Academic researchers produce scientific manuscripts funded primarily by public grants, philanthropic institutions, or university endowments. The raw intellectual property is assigned to commercial publishers via copyright transfer or exclusive commercial licensing agreements without monetary consideration. Concurrently, peer evaluation is conducted by volunteer researchers whose time is absorbed as an uncompensated externality. An empirical analysis by Aczel et al. calculated that researchers worldwide expended more than 100 million to 130 million hours on peer review in a single calendar year—the labor equivalent of approximately 15,000 continuous person-years. The monetary value of this uncompensated labor is estimated between $1.5 billion for United States researchers alone and $2.5 billion to upwards of $6.0 billion globally when accounting for reviewed-but-rejected manuscripts.

The second structural mechanism is the monetization of format gatekeeping via Article Processing Charges (APCs). Under the guise of open access, commercial publishers transitioned from subscription paywalls to author-facing publishing fees. Far from reflecting the underlying technical costs of electronic document dissemination, APCs range from $1,500 to $11,900 per document. Author formatting requirements—such as specialized LaTeX templates, bespoke margin adjustments, and platform-specific bibliographic schemas—impose an administrative friction tax. Rejections based on formatting syntax errors rather than methodological substance act as structural toll-booths that force researchers into author-funded

editing and formatting mills.
The third structural mechanism is predatory data aggregation and closed model ingestion. Legacy preprint servers and corporate repositories act as centralized aggregation points that ingest unencrypted scientific text, high-dimensional datasets, and experimental methodologies. Centralized preprint architectures (such as arXiv and bioRxiv) maintain raw file ingestion, central server caching, and unrestricted API access. These centralized data stores have been systematically crawled and indexed by commercial search engines and enterprise artificial intelligence laboratories. The data is ingested without author authorization or economic settlement to train proprietary Large Language Models (LLMs) and specialized domain transformers, severing scientific creators from the downstream financial value generated by their discoveries.

Cost Element	Traditional Commercial Cartel (Elsevier / Springer Nature)	Centralized Preprint Archive (arXiv / bioRxiv)	Sovereign Model: AROS™ via ALDI™
Author Monetary Toll (APC)	$1,500 – $11,900 per paper	$0.00 (subsidized via university consortia & grants)	**$0.00 (fully underwritten by Albert and Robot)**
Formatting Friction & Gatekeeping	High; manual desktop reformatting, proprietary styles, desk rejection for syntax errors	Moderate; automated TeX compilation queues, frequent package collisions	**Zero; local schema validation executed within author enclave**
Reviewer Compensation	$0.00; $totallaborextract$, 3.5B–$6B annual global uncompensated value)	$0.00; non-evaluative screening, no formal peer review	**$250 – $1,000 direct protocol bounty per verified review**
Raw Compute & Hosting Unit Cost	Low (< $0.01 per PDF); masked by high margin markups	Nominal ($0.05 – $0.15 per paper amortized operating cost)	**$0.01 – $0.10 ephemeral handshake and decentralized state commit**
Document Custody Model	Centralized server ingestion, paywall lock, or copyright surrender	Centralized file server repository, permanently retained plain text	**Zero-Ingestion; payload remains in sovereign local enclave**
Corporate AI & LLM Exploitation	Direct licensing of repository corpora to closed AI vendors	Unmitigated third-party crawler access, bulk text scraping	**Cryptographically blocked; active entropy poisoning for unauthorized crawlers**
Settlement & Validation Latency	6 to 18 months standard publication cycle	24 to 72 hours moderation delay	**Deterministic SLA (< 72 hours protocol settlement)**

The systemic deadweight loss of the legacy cartel, \mathcal{L}_{\text{legacy}}, represents the aggregate financial extraction imposed on the scientific commons:
\mathcal{L}_{\text{legacy}} = \sum_{k=1}^{M} \left(\text{APC}_k + \text{FormatToll}_k + V_{\text{unpaid_labor}} - C_{\text{actual_compute}} \right)
Where M is the annual global volume of submitted papers, \text{APC}_k \in [\$1,500, \$11,900]

is the publishing charge, \text{FormatToll}_k is the financialized loss of research time dedicated to superficial journal styling, V_{\text{unpaid_labor}} = \bar{h}_{\text{review}} \times \text{Wage}_{\text{academic}} represents the extracted market value of professional reviewer labor, and C_{\text{actual_compute}} < \$0.01 reflects actual bare-metal blob storage and network egress costs.
By shifting the infrastructure from predatory centralized ingestion to a sovereign client-side architecture, the total financial liability collapses from several thousand dollars per paper to bare cryptographic verification overhead.

Sovereign Enclave Architecture and the Zero-Ingestion Pipeline

The fundamental design flaw of historical research archives is premature data ingestion: researchers transfer unencrypted scientific content across network perimeters to centralized database clusters before authenticity, provenance, or peer verification are established. The Autonomous Research Operating System (AROS™), governed by Autonomous Ledger & Decentralized Identity (ALDI™), reverses this relationship. Raw document ingestion is structurally prohibited: scientific assets remain locked inside the researcher's authenticated, sovereign hardware enclave.

The Nine-Stage Inverted Processing Sequence

The protocol establishes a nine-stage pipeline that transitions scientific validation from server-side curation to client-side enclave computation and on-chain state anchoring:

1. **Submission Gateway (Local Initiation)**: The researcher initializes the publication intent via a local ALDI™ Decentralized Identity client without transmitting binary payloads or source files over HTTP endpoints.
2. **Cryptographic Enclave Binding (Zero Ingestion)**: The local client parses raw scientific artifacts against canonical schema standards within a hardware Trusted Execution Environment (TEE), compiling a Sparse Merkle Tree (SMT) and generating an immutable binding hash H_B.
3. **Review Modality Template Selection**: The researcher selects an on-chain evaluation template published by accredited, bonded peer reviewers (such as formal mathematical verification, empirical dataset audit, or theoretical scrutiny).
4. **Cryptographic Review Dispatch**: ALDI™ broadcasts an evaluation request containing strictly metadata—schema identifier, modality tier, author public key, and binding hash H_B—keeping raw payloads isolated locally.
5. **Sovereign Enclave Pairing & Blind Handshake**: ALDI™ matches an accredited evaluator via a Verifiable Random Function (VRF) and provisions an ephemeral, end-to-end encrypted session over the Noise Protocol (\text{Noise_IK}) between the author's enclave and the reviewer's sandboxed environment.
6. **Sandboxed Execution, Attestation, and On-Chain Double-Binding**: The paired evaluator examines the work within a zero-exfiltration confidential virtual machine, generates evaluation digest H_{\text{rev}}, and cryptographically signs the review payload, prompting ALDI™ to mint an on-chain double-binding receipt \Pi_{\text{bind}} = (H_B \iff H_{\text{rev}}).
7. **Sovereign Portal Hosting**: AROS™ hosts the public verification portal and discovery

index, publishing exclusively cryptographic receipts, double-binding proofs, and attestation vectors without raw intellectual property.

8. **Protocol Underwriting by Albert and Robot**: The Albert and Robot Treasury Reserve underwrites 100% of pipeline expenses, issuing real-time reviewer bounties and absorbing transaction gas costs to ensure a permanent $0.00 author fee.
9. **Data Sovereignty Axiom**: Active perimeter defenses—including proof-of-human challenges and synthetic entropy poisoning—protect the scientific commons from unauthorized scrapers and corporate LLM aggregators.

Enclave State Binding Primitives

The local client decomposes the manuscript into n discrete, canonical schema leaves \{s_1, s_2, \dots, s_n\} representing sections, formal hypotheses, mathematical specifications, and empirical observations. Each leaf hash L_i in the Sparse Merkle Tree is computed using BLAKE3:

$$L_i = \text{BLAKE3}\left(\text{SectionIndex}_i \,\Vert{}\, \text{ContentLength}_i \,\Vert{}\, \text{BLAKE3}(s_i)\right)$$

The Merkle root R_{doc} aggregates the comprehensive contents of the paper. To link the specific deliverable immutably to the author's cryptographic identity without publishing its content, the Enclave Binding Hash H_B is computed:

$$H_B = \text{BLAKE3}\left(\text{PubKey}_{\text{auth}} \,\Vert{}\, \text{Schema}_{\text{UUID}} \,\Vert{}\, R_{\text{doc}} \,\Vert{}\, \text{Nonce}_{\text{local}}\right)$$

Where $\text{PubKey}_{\text{auth}} \in \mathbb{G}_1$ is the author's public key derived from an elliptic curve (such as Curve25519 or BLS12-381), $\text{Schema}_{\text{UUID}}$ defines the standardized disciplinary schema to which the document conforms, and $\text{Nonce}_{\text{local}}$ is an entropy vector generated inside the hardware enclave.

Confidential Execution and Remote Attestation Architecture

To permit a designated peer reviewer to examine an unpublished manuscript without risking document exfiltration, the protocol relies on Confidential Computing hardware architectures—specifically AMD Secure Encrypted Virtualization-Secure Nested Paging (SEV-SNP) and Intel Trust Domain Extensions (TDX).

The reviewer environment executes within a measured Confidential Virtual Machine (CVM) whose initial state is cryptographically signed by the hardware root of trust. The attestation flow operates through a strict verification protocol:

The author's client enclave issues an attestation challenge to the reviewer node containing a high-entropy nonce η_{auth}. In response, the reviewer's CVM requests an attestation report from the hardware security processor via the SNP_GET_REPORT instruction, injecting η_{auth} into the report data field. The hardware security processor (e.g., AMD-SP) compiles the architectural state, guest policy, microcode version, and launch measurement register into the report.

This report is digitally signed by the processor's Versioned Chip Endorsement Key (VCEK), which is anchored through the AMD Signing Key (ASK) to the AMD Root Key (ARK). The author's enclave receives the attestation bundle and validates the certificate chain against the manufacturer's root public keys, ensuring that the launch measurement matches the open-source, canonical hash of the AROS™ sandbox container.

The sandbox container enforces strict operational constraints: local disk swap is disabled;

loopback virtual networking is configured exclusively to interface with the local display engine; all external network interfaces except the authenticated Noise transport channel are disabled at the kernel boundary; and clipboard export, bulk screenshotting, and system core dumping are blocked at the hypervisor boundary via Secure Nested Paging (SNP) reverse map tables.

Ephemeral Access Orchestration via the Noise Protocol

Once hardware attestation validates the uncompromised integrity of the reviewer's runtime, an encrypted peer-to-peer transport session is established using the No[span_38](start_span)[span_38](end_span)[span_43](start_span)[span_43](end_span)ise_IK handshake pattern. The Noise_IK pattern is selected because the author client already possesses the long-term static public key of the assigned reviewer ($\text{StaticPubKey}_{\text{rev}}$) provided via ALDI™ discovery, enabling single-round-trip mutual authentication and immediate payload encryption.

Before initiating the handshake, the author signs an ephemeral capability token τ_{eval} restricting access to an explicit temporal epoch Δt:

$$\tau_{\text{eval}} = \text{Sign}_{\text{PrivKey}_{\text{auth}}}\left(\text{PubKey}_{\text{rev}} \,\Vert{}\, H_B \,\Vert{}\, t_{\text{start}} \,\Vert{}\, t_{\text{expiry}}\right)$$

The cryptographic exchange proceeds across deterministic phases:

The author client initializes an ephemeral keypair and performs Diffie-Hellman operations against the reviewer's static key, encrypting the author's static key and capability token $\tau_{\text{ev[span_40](start_span)[span_40](end_span)[span_45](start_span)[span_45](end_span)al}}$ within the initial outbound packet. Upon receipt, the reviewer CVM decrypts the payload, validates τ_{eval}, generates its own ephemeral keypair, and executes secondary Diffie-Hellman operations against both the author's ephemeral and static keys.

The derived symmetric cipher keys are established via HKDF-BLAKE3, transitioning transport communication into authenticated ChaCha20-Poly1305 stream encryption. The manuscript content is decrypted exclusively into non-swappable volatile RAM buffers inside the reviewer's attested hardware enclave. At no point during the review process does unencrypted research data traverse public networks, touch persistent storage devices, or reside on intermediary proxy servers.

Cryptographic Attestation and On-Chain Double-Binding

When the reviewer completes their evaluation within the sandboxed environment, they construct an evaluation vector D_{rev} conforming to the selected modality schema. The review artifact includes categorical quality metrics, reproducible code verification proofs, methodological critique, and claim validity attestations.

The reviewer compiles the structured document D_{rev} and generates its canonical digest using BLAKE3:

$$H_{\text{rev}} = \text{BLAKE3}(D_{\text{rev}})$$

The reviewer then produces a digital signature σ_{rev} over the compound state linking the review hash to the author's binding hash H_B:

$$\sigma_{\text{rev}} = \text{Sign}_{\text{PrivKey}_{\text{rev}}}\left(H_{\text{rev}} \,\Vert{}\, H_B \,\Vert{}\, \text{ScoreVector} \,\Vert{}\, \text{Timestamp}\right)$$

This signature binds the reviewer's identity ($\text{PubKey}_{\text{rev}}$) not merely to their own

commentary, but directly to the exact immutable state of the author's local document (H_B) without the reviewer possessing a persistent copy of the author's files.
Upon transmission of (σ_{rev}, H_{rev}, ScoreVector) over the Noise session, the ALDI™ protocol verification layer ingests the signature and verifies that:

1. $\text{Verify}(\sigma_{\text{rev}}, \text{PubKey}_{\text{rev}}) == \text{TRUE}$.
2. The reviewer is actively accredited and currently holds a non-depleted security bond S_{bond} in the ALDI™ staking registry.
3. The review strictly conforms to the syntactic requirements of the requested modality template.

Upon validation, the ALDI™ engine mints the immutable on-chain Double-Binding Proof Π_{bind}:
$$\Pi_{\text{bind}} = \text{Commit}_{\text{ALDI}}\left(H_B, H_{\text{rev}}, \sigma_{\text{rev}}, \text{BlockNumber}, \text{Schema}_{\text{UUID}}\right)$$
This double-binding guarantees bidirectional integrity:

- **Author Non-Repudiation**: The author cannot alter a single equation or character in their original document without invalidating H_B, which would sever the connection to the successful peer review proof Π_{bind}.
- **Reviewer Accountability**: The reviewer cannot deny having audited and approved the specific document state H_B, as their cryptographic signature permanently links their on-chain identity and bonded stake to that hash.
- **Zero Document Disclosure**: Both H_B and H_{rev} are collision-resistant, one-way cryptographic hashes. Third-party observers, indexers, and corporate crawlers examining Π_{bind} on the public ledger can verify that a rigorous review was executed by an accredited expert, but cannot reconstruct a single sentence of the underlying research payload.

Protocol Underwriting Economics and Actuarial Dynamics

The economic equilibrium of AROS™ rests upon protocol-level underwriting managed by the Albert and Robot Treasury Reserve. In direct contrast to legacy publishers that charge predatory author fees while pocketing large profit margins, Albert and Robot treat scientific integrity verification as fundamental public infrastructure. The protocol absorbs all marginal system costs, guaranteeing that zero fees are charged to researchers while reviewers receive market-rate compensation.

The Protocol Liability Function

For every evaluated deliverable S, the total protocol liability C(S) underwritten by the treasury is formulated as:
$$C(S) = \mathcal{P}_{\text{bounty}}(m) + \mathcal{G}_{\text{anchor}} + \mathcal{E}_{\text{handshake}}$$
Where $\mathcal{P}_{\text{bounty}}(m)$ represents the deterministic cash-equivalent bounty disbursed to the peer reviewer, scaled dynamically to the selected modality tier m:
$$\mathcal{P}_{\text{bounty}}(m) = \begin{cases} \$250.00 & \text{Tier 1: Schema Integrity, Consistency, and Baseline Attestation} \\ \$500.00 & \text{Tier 2: Methodological Replicability \& Empirical Dataset Audit} \\ \$1,000.00 & \text{Tier 3: Formal

Mathemati[span_6](start_span)[span_6](end_span)cal, Cryptographic \& Algorithmic Proof Check} \end{cases}
The remaining variables correspond to cryptographic operational costs:

- \mathcal{G}_{\text{anchor}}: The layer-2 rollup settlement fee required to commit \Pi_{\text{bind}} permanently into the ledger state (\mathcal{G}_{\text{anchor}} < \$0.05 on modern optimistic or zero-knowledge rollups).
- \mathcal{E}_{\text{handshake}}: The amortized infrastructure compute overhead necessary to route the ephemeral WebRTC/Noise proxy connection between sovereign enclaves (\mathcal{E}_{\text{handshake}} < \$0.10).

Cost-Efficiency and Margin Compression

The structural cost comparison demonstrates the degree of capital efficiency realized when eliminating publisher rent extraction:

Deployment Tier	Reviewer Bounty (\mathcal{P}_{\text{bounty}})	Settlement Gas (\mathcal{G}_{\text{anchor}})	Compute Relay (\mathcal{E}_{\text{handshake}})	Total Protocol Cost C(S)	Legacy APC Benchmark	Efficiency Multiple
Tier 1 (Schema & Audit)	$250.00	$0.05	$0.10	$250.15	$3,800.00	**15.19×**
Tier 2 (Methodology Audit)	$500.00	$0.05	$0.10	$500.15	$3,800.00	**7.59×**
Tier 3 (Formal Proof Check)	$1,000.00	$0.05	$0.10	$1,000.15	$11,900.00	**11.90×**

Even at maximum depth (Tier 3 formal proof checking at $1,000.00 per paper paid directly to technical specialists), the overall protocol cost remains an order of magnitude lower than the billing overhead of legacy commercial conglomerates. This structural compression is achievable because the protocol eliminates corporate dividends, executive compensation packages, legal copyright enforcement teams, corporate acquisitions, and manual formatting departments. Because 99.7% of protocol expenditure flows directly to the evaluating domain expert (\mathcal{P}_{\text{bounty}}), scientific capital is realigned with technical rigor rather than corporate rent extraction.

Game-Theoretic Equilibrium and Anti-Collusion Mechanics

Incentivizing peer review via direct monetary bounties introduces game-theoretic vulnerabilities. In an unconstrained market, rational malicious actors could exploit the treasury through collusive review rings, Sybil farming, or rubber-stamping. To preserve the economic sustainability of the Albert and Robot Treasury, ALDI™ integrates three game-theoretic

enforcement mechanisms.

Blinded VRF Matching and Staked Slashing Pools

Reviewers cannot browse an open marketplace or select specific papers for review. When an author dispatches an evaluation request (H_B, $\text{Schema}_{\text{UUID}}$), the matching process is governed entirely by an on-chain Verifiable Random Function (VRF):

$$\text{Seed}_{\text{match}} = \text{VRF}_{\text{ALDI}}\left(\text{BlockHeader} \,\Vert{}\, H_B \,\Vert{}\, \text{EpochEntropy}\right)$$

The seed deterministically samples an evaluator from the active subset of reviewers who possess accredited identity credentials within the manuscript's disciplinary domain.

To participate in the assignment pool, reviewers must lock a capital bond S_{bond} in escrow:

$$S_{\text{bond}} \ge 2 \times \max\left(\mathcal{P}_{\text{bounty}}\right) = \$2{,}000.00$$

If an evaluator behaves maliciously or attempts to manipulate the matching protocol, 100% of their staked bond is burned via an automated slashing protocol, and their cryptographic identity is permanently barred from platform participation.

Randomized Verification Rings and Discrepancy Audits

To detect rubber-stamping and collusive agreements, ALDI™ routes a deterministic fraction ($\rho = 0.08$, or 8% of all completed evaluations) to an independent, secondary validator node sampled from the top staking quartile. The secondary validator evaluates the identical document state H_B through an isolated ephemeral session without access to the primary reviewer's findings.

Upon receipt of the audit evaluation, ALDI™ calculates a Discrepancy Vector δ:

$$\delta = \Vert{}\text{ScoreVector}_A - \text{ScoreVector}_B\Vert{}_2$$

When $\delta \le \epsilon_{\text{threshold}}$, consensus is established: the primary review is authenticated, releasing $\mathcal{P}_{\text{bounty}}$ to the primary reviewer and an audit bonus $\mathcal{P}_{\text{audit}} = 0.5 \times \mathcal{P}_{\text{bounty}}$ to the auditor.

Conversely, when $\delta > \epsilon_{\text{threshold}}$, an automated governance tribunal of five randomized verifiers is convened. If bad faith or fraudulent attestation is proven, the reviewer's bond is liquidated:

$$\text{Slashing Action}: S_{\text{bond}} \leftarrow 0 \quad (\text{100\% Burned}); \quad \text{Reputation} \leftarrow 0 \quad (\text{Revoked})$$

The expected utility $\mathbb{E}[U]$ for a colluding reviewer contemplating a fraudulent evaluation is formulated as:

$$\mathbb{E}[U_{\text{collude}}] = (1 - \rho) \cdot \mathcal{P}_{\text{bounty}} + \rho \cdot \left(-S_{\text{bond}} \right)$$

To ensure that collusion remains irrational ($\mathbb{E}[U_{\text{collude}}] < 0$), the protocol enforces a minimum staking bond relative to the audit rate:

$$S_{\text{bond}} > \frac{1 - \rho}{\rho} \cdot \mathcal{P}_{\text{bounty}}$$

For $\rho = 0.08$ and $\mathcal{P}_{\text{bounty}} = \$500$, the zero-profit threshold requires $S_{\text{bond}} > 11.5 \times \$500 = \$5{,}750$. Setting the protocol bond requirement at $S_{\text{bond}} = \$6{,}000$ yields a negative expected return:

$$\mathbb{E}[U_{\text{collude}}] = (0.92 \times 500) + (0.08 \times -6{,}000) = 460 - 480 = -\$20.00$$

Any collusive strategy results in expected capital depletion.

Bayesian Truth Serum for Subjective Scoring

To eliminate consensus bias where reviewers simply guess what other reviewers will say, the protocol scores subjective quality vectors using Dražen Prelec's Bayesian Truth Serum (BTS) scoring rule. Reviewers are asked to supply both their personal assessment x_i and an estimate of the distribution of other reviewers' scores y_i.
The score reward assigned to reviewer i is calculated via a strictly proper scoring rule:
\mathcal{S}(x_i, y_i) = \sum_{k} x_{ik} \log\left(\frac{\bar{x}_k}{\tilde{y}_k}\right) + \alpha \sum_{k} y_{ik} \log\left(\frac{x_k}{y_{ik}}\right)
Where \bar{x}_k is the actual sample frequency of evaluation score k, and \tilde{y}_k is the geometric mean of predictions. The first term—the *information score*—rewards evaluations that are surprisingly common relative to expectations, which is mathematically maximized under honest reporting of Bayesian private signals. The second term—the *prediction score*—penalizes inaccurate forecasts using a standard logarithmic scoring rule. Honest, deep technical evaluation represents the unique Bayesian Nash Equilibrium of the review network.

Smart Contract Settlement Specification

The settlement core is instantiated as an immutable smart contract operating on an optimistic or ZK-rollup settlement layer. The contract coordinates state validation, handles bounty disbursements from the Albert and Robot Treasury, and executes bond slashing without possessing knowledge of raw scientific content.

```
// Copyright © Albert Lane. All Rights Reserved.
// Proprietary and Confidential — Albert Lane Proprietary Software License.
// Unauthorized use, copying, modification, or distribution is prohibited.
// License inquiries: https://albertlane.net

// SPDX-License-Identifier: LicenseRef-AlbertLane-Proprietary
pragma solidity ^0.8.24;

interface IALDISovereignVerifier {
    function verifyReviewAttestation(
        bytes32 hDoc,
        bytes32 hRev,
        bytes calldata sigRev,
        address reviewer
    ) external view returns (bool);
}

/**
 * @title AROSSettlementEngine
 * @notice Coordinates zero-ingestion cryptographic double-binding and underwritten bounty disbursements.
 * @dev Governed exclusively by the Albert and Robot Treasury underwriter perimeter.
```

```
 */
contract AROSSettlementEngine {
    address public immutable underwriterTreasury;
    IALDISovereignVerifier public immutable verifier;

    struct VerificationReceipt {
        bytes32 hDoc;
        bytes32 hRev;
        address reviewer;
        uint256 settlementTimestamp;
        uint256 bountyDisbursed;
    }

    // Mapping from Author Binding Hash (H_B) to Verification Receipt
    mapping(bytes32 => VerificationReceipt) public onChainReceipts;

    // Mapping from Reviewer Identity to Staked Bond Balances
    mapping(address => uint256) public reviewerStakedBonds;

    event DoubleBindingMinted(
        bytes32 indexed bindingHash,
        bytes32 indexed hDoc,
        bytes32 indexed hRev,
        address indexed reviewer,
        uint256 bounty
    );

    event ReviewerBondStaked(address indexed reviewer, uint256
amount);
    event ReviewerSlashed(address indexed reviewer, uint256
amountBurned);

    modifier onlyTreasury() {
        require(msg.sender == underwriterTreasury,
"UNAUTHORIZED_TREASURY_ACCESS");
        _;
    }

    constructor(address _verifier, address _treasury) {
        require(_verifier != address(0), "INVALID_VERIFIER_ADDRESS");
        require(_treasury != address(0), "INVALID_TREASURY_ADDRESS");
        verifier = IALDISovereignVerifier(_verifier);
        underwriterTreasury = _treasury;
    }

    /**
     * @notice Stakes capital bonds required for reviewer
participation.
```

```
     */
    function stakeReviewerBond() external payable {
        require(msg.value >= 2 ether, "INSUFFICIENT_BOND_AMOUNT");
        reviewerStakedBonds[msg.sender] += msg.value;
        emit ReviewerBondStaked(msg.sender, msg.value);
    }

    /**
     * @notice Atomically binds researcher enclave hash to reviewer
output and settles bounty.
     * @dev Zero knowledge of raw scientific content is required or
ingested by this contract.
     * @param hDoc Canonical Merkle binding hash of the researcher's
enclave artifact (H_B).
     * @param hRev BLAKE3 hash of the structured reviewer evaluation
report (H_rev).
     * @param sigRev Digital cryptographic signature generated by the
reviewer.
     * @param reviewer Payable address of the accredited reviewer
identity.
     * @param bountyAmount Exact bounty underwritten by Albert and
Robot.
     */
    function executeDoubleBindingSettlement(
        bytes32 hDoc,
        bytes32 hRev,
        bytes calldata sigRev,
        address payable reviewer,
        uint256 bountyAmount
    ) external onlyTreasury {
        require(onChainReceipts[hDoc].settlementTimestamp == 0,
"BINDING_ALREADY_EXISTS");
        require(reviewerStakedBonds[reviewer] >= 2 ether,
"REVIEWER_BOND_DEFICIT");
        require(
            verifier.verifyReviewAttestation(hDoc, hRev, sigRev,
reviewer),
            "INVALID_CRYPTOGRAPHIC_ATTESTATION"
        );

        bytes32 bindingId = keccak256(abi.encodePacked(hDoc, hRev,
reviewer));

        onChainReceipts[hDoc] = VerificationReceipt({
            hDoc: hDoc,
            hRev: hRev,
            reviewer: reviewer,
            settlementTimestamp: block.timestamp,
```

```
            bountyDisbursed: bountyAmount
        });

        // Direct disbursement from Albert and Robot reserve to
Reviewer
        (bool success, ) = reviewer.call{value: bountyAmount}("");
        require(success, "TREASURY_DISBURSEMENT_FAILED");

        emit DoubleBindingMinted(bindingId, hDoc, hRev, reviewer,
bountyAmount);
    }

    /**
     * @notice Executes governance-directed slashing of malicious or
collusive reviewer bonds.
     * @param reviewer Address of the reviewer identified in bad-faith
attestation.
     */
    function slashReviewer(address reviewer) external onlyTreasury {
        uint256 bond = reviewerStakedBonds[reviewer];
        require(bond > 0, "NO_BOND_TO_SLASH");

        reviewerStakedBonds[reviewer] = 0;

        // Slashed funds are burned permanently by transferring to the
0x0 address
        (bool success, ) = payable(address(0)).call{value: bond}("");
        require(success, "BOND_BURNING_FAILED");

        emit ReviewerSlashed(reviewer, bond);
    }
}
```

Defensive Extraction Shield and Institutional Interoperability

Centralized academic indexing networks and corporate automated scrapers systematically capture public preprints to train proprietary foundation models without compensating the authors or honoring licensing restrictions. The AROS™ architecture treats open access as verifiable validation rather than passive, unencrypted exposure.

Deterministic Anti-Scraping and Synthetic Entropy Poisoning

The AROS™ public portal serves strictly zero-knowledge verification receipts to unauthenticated search requests. The public gateway indexes:

- The Binding Hash H_B.

- The Review Hash H_{\text{rev}}.
- The Reviewer Identity Attestation \sigma_{\text{rev}}.
- Categorical evaluation vectors (such as scoring metrics and reproducibility indicators).

The raw scientific text remains localized within the author's enclave storage. If an unauthorized, automated scraping agent attempts to traverse edge proxies, the AROS™ Defensive Gateway triggers dynamic synthetic entropy poisoning.

Requests lacking mutual TLS attestation or legitimate human verification are routed to a procedural synthesis pipeline. This pipeline emits mathematically plausible white noise—syntactically valid LaTeX documents containing scrambled citations, inverted statistical dependencies, and randomized equations. If ingested into commercial LLM pretraining corpuses, these synthetic payloads degrade transformer embedding spaces, enforcing an economic and technical penalty against unapproved extraction. Legitimate human access is preserved through ephemeral client-side zero-knowledge proof challenges and sandboxed WebAssembly viewers that decrypt documents in volatile memory without exposing cleartext endpoints.

Institutional Compliance Adapter

To integrate with university systems without violating zero-retention parameters, AROS™ provides an institutional compliance adapter. Global academic institutions frequently require faculty to deposit scholarly outputs into institutional repositories (such as DSpace, EPrints, or Fedora) to comply with grant funding mandates.

The compliance adapter packages the on-chain Double-Binding Proof \Pi_{\text{bind}} into standardized SWORD v2 envelopes containing Dublin Core and Metadata Exchange and Transmission Standard (METS) schemas. The institutional repository ingests the cryptographic receipt as verifiable proof of peer-reviewed publication. Unless an author explicitly executes an affirmative release directive to archive their cleartext work publicly, no raw files leave the sovereign enclave. Universities satisfy accreditation metrics while faculty retain complete custody over their intellectual property.

Comparative Benchmark Analysis

The operational parameters of legacy publishers, centralized archives, and the sovereign AROS™ / ALDI™ framework highlight fundamental differences across custody, economics, and validation mechanics:

Metric	Traditional Cartel (Elsevier / Nature)	Centralized Repository (arXiv / bioRxiv)	Sovereign Network (AROS™ via ALDI™)
Author Cost	$1,500 – $11,900 (Predatory APCs)	$0.00 (Charity/Grant subsidization)	**$0.00 (Fully Underwritten)**
Formatting Tolls	High (Arbitrary LaTeX/Word templates)	Medium (TeX parsing rejections)	**Zero (Local schema validation)**
Reviewer Compensation	$0.00 (Exploitative unpaid labor)	$0.00 (No native review mechanism)	**$250 – $1,000 Direct Cash Bounty**
Data Custody	Fully Extracted & Paywalled	Centralized Ingestion & Permanent Storage	**Zero Ingestion (Author Enclave Boundary)**
AI Scraping Risk	Commercialized	Unrestricted automated	**Cryptographically**

Metric	Traditional Cartel (Elsevier / Nature)	Centralized Repository (arXiv / bioRxiv)	Sovereign Network (AROS™ via ALDI™)
	without author cut	scraping	**Blocked & Poisoned**
Settlement Velocity	6 to 18 Months	2 to 5 Days (Moderation queue)	**Deterministic SLA (< 72 Hours)**
Proof of Authenticity	Centralized DOI lookup	Centralized URL identifier	**On-Chain Cryptographic Double-Binding**
IP Ownership	Copyright Surrendered or Bound	Non-exclusive license surrendered	**100% Retained by Author Identity**

Operational Deployment and Integration Blueprint

The transition from extractive academic publication to sovereign zero-ingestion validation progresses across four sequential deployment phases:

The first phase focuses on client engine and enclave attestation infrastructure. Workstation-compatible ALDI™ clients are deployed with hardware drivers for AMD SEV-SNP, Intel TDX, and modern secure enclaves. The client parses local scientific files into canonical schemas and generates the Sparse Merkle Tree entirely within volatile RAM. Sandboxed WebAssembly containers are standardized with pre-installed verification runtimes for Lean, Coq, Python, and Julia, isolated from network interfaces.

The second phase establishes the settlement core and underwriter treasury. The AROSSettlementEngine smart contract is deployed to a Layer-2 rollup, and the Albert and Robot Treasury Reserve liquidity pool is capitalized. Domain reviewers register decentralized identities, submit accredited credentials, and lock the mandatory $6,000 security bond into the staking contract.

The third phase deploys institutional integration modules and defensive crawling shields. The SWORD v2 compliance adapter is connected to university repository endpoints, allowing deposit of on-chain verification receipts. AROS™ edge gateways activate dynamic synthetic noise generation, intercepting unauthenticated crawler requests and returning adversarial mathematical entropy.

The fourth phase executes disciplinary scaling across global institutions. Specialized schema frameworks are initialized: AROS_STEM_v1 for software code and numerical matrices, AROS_BIO_v1 for genomics and clinical observations, and AROS_PROOF_v1 for formal machine-checked proofs. University consortia and independent labs are onboarded to route peer review workflows through the double-binding settlement pipeline.

System auditors, research institutions, and author clients trigger verification operations via the protocol CLI using the standard invocation format:

EXECUTE INVERTED RESEARCH AUDIT:

- AUTHOR_IDENTITY_PUBKEY: 0x8f3c7b2a9e1d4f6a8b0c2e4d6f8a0b2c4e6f8a0b2c4e6f8a0b2c4e6f8a0b2c4e
- SCHEMA_SPECIFICATION: AROS_PROOF_v1
- REVIEW_MODALITY_TIER: TIER_3
- TARGET_UNIVERSITY_NODE: CRKN_INSTITUTIONAL_REPO_STANDARDS_v2
- UNDERWRITING_ACCOUNT: ALBERT_AND_ROBOT_TREASURY_MAINNET

This instruction executes client-side schema parsing, calculates the binding hash H_B, publishes the metadata dispatch, initiates VRF blinded evaluator matching, establishes the ephemeral Noise transport channel, verifies the proof in a confidential sandbox, and mints the

immutable on-chain Double-Binding Proof \Pi_{\text{bind}} settled through the Albert and Robot Treasury.

Architectural Synthesis

The AROS™ and ALDI™ framework resolves the principal economic and architectural failures of legacy scholarly publishing. By replacing centralized file ingestion with client-side trusted execution environments, researchers retain absolute physical and legal sovereignty over their intellectual output. By replacing uncompensated peer review with directly underwritten protocol bounties funded by Albert and Robot, technical experts receive competitive compensation while author fees remain permanently fixed at zero.
Furthermore, by anchoring peer evaluation through on-chain double-binding rather than centralized database indices, the protocol establishes tamper-proof provenance while denying corporate scrapers the ability to harvest research data. The system transforms scientific publication from an extractive corporate monopoly into a decentralized, zero-ingestion attestation commons.

Works cited

1. A billion-dollar donation: estimating the cost of researchers' time, https://www.researchgate.net/publication/356194956_A_billion-dollar_donation_estimating_the_cost_of_researchers'_time_spent_on_peer_review 2. How Much Profit Did Elsevier Make in 2024? - PublishingState.com, https://publishingstate.com/how-much-profit-did-elsevier-make-in-2024/2025/ 3. 2024 Annual Report - RELX, https://www.relx.com/~/media/Files/R/RELX-Group/documents/reports/annual-reports/relx-2024-annual-report.pdf 4. The global provider of information-based analytics and decision tools, https://www.relx.com/~/media/Files/R/RELX-Group/documents/investors/relx-overview-september-2026.pdf 5. Elsevier parent company reports 10% rise in profit, to £3.2bn, https://www.researchprofessionalnews.com/rr-news-world-2025-2-elsevier-parent-company-reports-10-rise-in-profit-to-3-2bn/ 6. Market segments - RELX, https://www.relx.com/~/media/Files/R/RELX-Group/documents/reports/annual-reports/2024-ar-sections/relx-2024-market-segments.pdf 7. Full article: Peer Review as Professional Development, https://www.tandfonline.com/doi/full/10.1080/08911762.2026.2638702 8. estimating the cost of peer review based on survey ... - ResearchGate, https://www.researchgate.net/journal/Research-Integrity-and-Peer-Review-2058-8615/publication/370223056_Scientific_sinkhole_estimating_the_cost_of_peer_review_based_on_survey_data_with_snowball_sampling/links/6446b599d749e4340e34ca2d/Scientific-sinkhole-estimating-the-cost-of-peer-review-based-on-survey-data-with-snowball-sampling.pdf?origin=scientificContributions 9. Scientific sinkhole: estimating the cost of peer review based on, https://pmc.ncbi.nlm.nih.gov/articles/PMC10122980/ 10. noise - DH-based secure protocol - ZX2C4 Git Repository, https://git.zx2c4.com/noise/diff/?id=a75e6c9bb6cb0bcdf44799525c2d994a7daa0720 11. Glossary of Terms | Support | Products - Documentation - IONOS, https://docs.ionos.com/cloud/support/general-information/glossary-of-terms 12. Platform-Agnostic Remote Attestation with WebAssembly Components, https://aaltodoc.aalto.fi/bitstreams/da701867-7a62-457e-b37c-04ed2ff2b884/download 13.

Remote attestation of SEV-SNP confidential VMs using e-vTPMs, https://arxiv.org/html/2303.16463v2 14. Confidential Computing on Kubernetes: TEEs and Attestation - Stribog, https://stribog.com/blog/confidential-computing-kubernetes-remote-attestation-tee 15. The Road to Trust: Building Enclaves within Confidential VMs, https://www.ndss-symposium.org/wp-content/uploads/2025-385-paper.pdf 16. The Noise Protocol Framework: The Crypto Behind WireGuard, https://havenmessenger.com/blog/posts/noise-protocol-framework-explained/ 17. A Bayesian Truth Serum for Subjective Data - ResearchGate, https://www.researchgate.net/publication/8231017_A_Bayesian_Truth_Serum_for_Subjective_Data 18. Bayesian Truth Serum and Information Theory: Game of Duels, http://www.its.caltech.edu/~cvitanic/PAPERS/SikicDuels.pdf 19. Peer Prediction without a Common Prior - Harvard DASH, https://dash.harvard.edu/bitstreams/7312037c-e29c-6bd4-e053-0100007fdf3b/download 20. creating-truth-telling.pdf - MIT Sloan Neuroeconomics Lab, https://nelmit.wordpress.com/wp-content/uploads/2016/10/creating-truth-telling.pdf 21. A Robust Bayesian Truth Serum for Small Populations, https://cdn.aaai.org/ojs/8261/8261-13-11789-1-2-20201228.pdf 22. Two Strongly Truthful Mechanisms for Three Heterogeneous Agents, https://par.nsf.gov/servlets/purl/10536752 23. Managed Obfuscation in the Age of AI Surveillance: A Cyber-Human, https://research.bryancalabro.com/papers/ai-surveillance-human-obfuscation 24. SWORD 2.0 Profile - GitHub Pages, https://swordapp.github.io/SWORDv2-Profile/SWORDProfile.html 25. Sword topic notes. According 20 marks.... | Filo, https://askfilo.com/user-question-answers-smart-solutions/sword-topic-notes-according-20-marks-3534313832313832 26. DSpace 9.x User Guide and Installation - Scribd, https://www.scribd.com/document/971013562/DSpace-9-Manual 27. Cisco Confidential Computing Overview White Paper, https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/computing-overview-wp.html